實用的HyperFlex一般知識

目錄

[簡介](#)
[常用縮寫](#)
[HyperFlex VMware VMNIC訂購](#)
[3.5之前的訂購](#)
[3.5後訂購](#)
[融合節點上的SCVM與計算節點上的SCVM](#)
[融合節點](#)
[CPU資源保留](#)
[記憶體資源保留](#)
[計算節點](#)
[不正常群集方案](#)
[案例 1:節點關閉](#)
[案例 2:磁碟關閉](#)
[案例 3:節點和磁碟均未關閉](#)
[如何使用命令列介面(CLI)檢查SED群集](#)
[HX維護模式與ESXi維護模式](#)
[常見問題](#)
[思科HyperFlex M4和M5伺服器上的SCVM安裝在何處？](#)
[群集可以容忍多少個故障節點？](#)
[具有5個或更多節點的群集](#)
[具有3和4個節點的群集](#)
[如果其中一個SCVM關閉會發生什麼情況？虛擬機器是否繼續運行？](#)
[SCVM上的VMware硬體版本已更新。現在怎麼辦？](#)

簡介

本檔案介紹管理員應掌握的有關思科HyperFlex(HX)的一般知識。

常用縮寫

SCVM =存儲控制器虛擬機

VMNIC =虛擬機網路介面卡

VNIC =虛擬網路介面卡

SED =自加密驅動器

VM =虛擬機

HX = **HyperFlex**

HyperFlex VMware VMNIC訂購

VMNIC位置在HX 3.5版及更高版本中已修訂。

3.5之前的訂購

在版本3.5之前，VNIC是根據VNIC編號分配的。

VNIC	虛擬交換器(vSwitch)
VNIC 0和VNIC 1	vSwitch-hx-inband-mgmt
VNIC 2和VNIC 3	vSwitch-hx-storage-data
VNIC 4和VNIC 5	vSwitch-hx-vm-network
VNIC 6和VNIC 7	vMotion

3.5後訂購

在3.5及更新版本中，VNIC是根據媒體訪問控制(MAC)地址分配的。因此，沒有特定的分配順序。

如果執行從3.5版本升級到3.5或更高版本的操作，則維持VMNIC訂購。

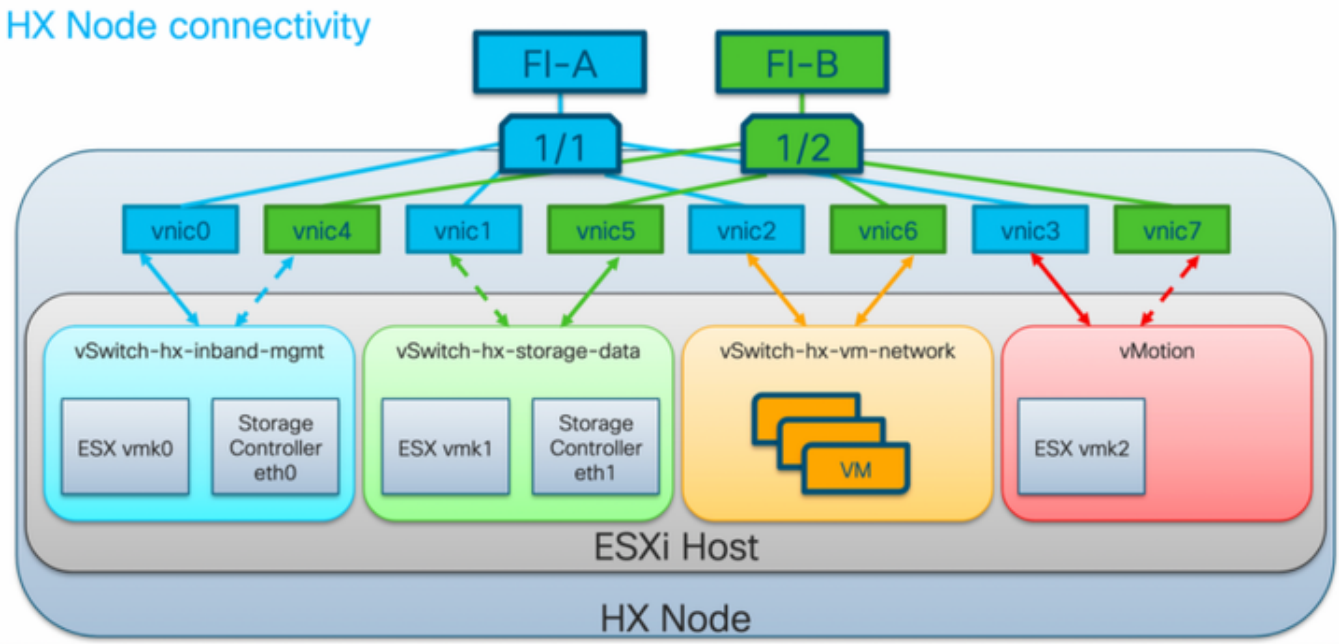

附註：對於HX Hyper-V，這將不適用，因為Hyper-V使用一致的裝置命名(CDN)。

融合節點上的SCVM與計算節點上的SCVM

SCVM同時駐留在融合節點和計算節點上，它們之間存在差異。

融合節點

CPU

SCVMCisco HXHyperFlexCPUVM(CPU)ESXiCPUVMSCVMCPUVMVMCPU

vCPU數量	股份	預留	限制
8	低	10800 MHZ	無限制

SCVMCisco HXHyperFlexVMVMESXiVM

伺服器型號	訪客記憶體量	保留所有訪客記憶體
HX 220c-M5SX HXAF 220c-M5SX HX 220c-M4S HXAF220c-M4S	48 GB	是
HX 240c-M5SX HXAF 240c-M5SX HX240c-M4SX HXAF240c-M4SX	72 GB	是
HX240c-M5L	78 GB	是

計算節點

僅計算節點具有輕量SCVM。它只配置了一個1024MHz的vCPU和512 MB的記憶體預留。

擁有計算節點的目的主要是維護vCluster Distributed Resource Scheduler™(DRS)設定，以確保DRS不會將使用者VM移回融合節點。

不正常群集方案

在以下情況中，HX群集可能會變為不正常狀態。

案例 1:節點關閉

當節點關閉時，群集將進入不正常狀態。在群集升級期間或伺服器進入維護模式時，節點應處於關閉狀態。

```
root@SpringpathController:~# stcli cluster storage-summary --detail
<snip>
current ensemble size:3
# of caching failures before cluster shuts down:2
minimum cache copies remaining:2
minimum data copies available for some user data:2
current healing status:rebuilding/healing is needed, but not in progress yet. warning:
insufficient node or space resources may prevent healing. storage node 10.197.252.99is either
down or initializing disks.
minimum metadata copies available for cluster metadata:2
# of unavailable nodes:1
# of nodes failure tolerable for cluster to be available:0
health state reason:storage cluster is unhealthy.storage node 10.197.252.99 is unavailable.
# of node failures before cluster shuts down:2
# of node failures before cluster goes into readonly:2
# of persistent devices failures tolerable for cluster to be available:1
# of node failures before cluster goes to enospace warn trying to move the existing data:na
# of persistent devices failures before cluster shuts down:2
```

```
# of persistent devices failures before cluster goes into readonly:2
# of caching failures before cluster goes into readonly:na
# of caching devices failures tolerable for cluster to be available:1
resiliencyInfo:
messages:
----------------------------------------
Storage cluster is unhealthy.
----------------------------------------
Storage node 10.197.252.99 is unavailable.
----------------------------------------
state: 2
nodeFailuresTolerable: 0
cachingDeviceFailuresTolerable: 1
persistentDeviceFailuresTolerable: 1
zoneResInfoList: None
spaceStatus: normal
totalCapacity: 3.0T
totalSavings: 5.17%
usedCapacity: 45.9G
zkHealth: online
clusterAccessPolicy: lenient
dataReplicationCompliance: non_compliant
dataReplicationFactor: 3
```

案例 2:磁碟關閉

當磁碟不可用時，群集將進入不正常狀態。當資料分發到其他磁碟時，應清除該情況。

```
root@SpringpathController:~# stcli cluster storage-summary --detail
<snip>
current ensemble size:3
# of caching failures before cluster shuts down:2
minimum cache copies remaining:2
minimum data copies available for some user data:2
current healing status:rebuilding/healing is needed, but not in progress yet. warning:
insufficient node or space resources may prevent healing. storage node is either down or
initializing disks.
minimum metadata copies available for cluster metadata:2
# of unavailable nodes:1
# of nodes failure tolerable for cluster to be available:0
health state reason:storage cluster is unhealthy. persistent device disk
[5000c5007e113d8b:0000000000000000] on node 10.197.252.99 is unavailable.
# of node failures before cluster shuts down:2
# of node failures before cluster goes into readonly:2
# of persistent devices failures tolerable for cluster to be available:1
# of node failures before cluster goes to enospace warn trying to move the existing data:na
# of persistent devices failures before cluster shuts down:2
# of persistent devices failures before cluster goes into readonly:2
# of caching failures before cluster goes into readonly:na
# of caching devices failures tolerable for cluster to be available:1
resiliencyInfo:
messages:
----------------------------------------
Storage cluster is unhealthy.
----------------------------------------
Persistent Device Disk [5000c5007e113d8b:0000000000000000] on node 10.197.252.99 is unavailable.
----------------------------------------
state: 2
nodeFailuresTolerable: 0
cachingDeviceFailuresTolerable: 1
persistentDeviceFailuresTolerable: 1
zoneResInfoList: None
```

```
spaceStatus: normal
totalCapacity: 3.0T
totalSavings: 8.82%
usedCapacity: 45.9G
zkHealth: online
clusterAccessPolicy: lenient
dataReplicationCompliance: non_compliant
dataReplicationFactor: 3
```

案例 3:節點和磁碟均未關閉

當節點和磁碟均未關閉時，群集可能會進入不正常狀態。如果正在進行重建，則會出現此情況。

```
root@SpringpathController:~# stcli cluster storage-summary --detail
<snip>
resiliencyDetails:
        current ensemble size:5
        # of caching failures before cluster shuts down:3
        minimum cache copies remaining:3
        minimum data copies available for some user data:2
current healing status:rebuilding is in progress, 98% completed.          minimum metadata copies
available for cluster metadata:2
        time remaining before current healing operation finishes:7 hr(s), 15 min(s), and 34
sec(s)
        # of unavailable nodes:0
        # of nodes failure tolerable for cluster to be available:1
        health state reason:storage cluster is unhealthy.
        # of node failures before cluster shuts down:2
        # of node failures before cluster goes into readonly:2
        # of persistent devices failures tolerable for cluster to be available:1
        # of node failures before cluster goes to enospace warn trying to move the existing
data:na
        # of persistent devices failures before cluster shuts down:2
        # of persistent devices failures before cluster goes into readonly:2
       # of caching failures before cluster goes into readonly:na
        # of caching devices failures tolerable for cluster to be available:2
resiliencyInfo:
    messages:
        Storage cluster is unhealthy.
    state: 2
    nodeFailuresTolerable: 1
    cachingDeviceFailuresTolerable: 2
    persistentDeviceFailuresTolerable: 1
    zoneResInfoList: None
spaceStatus: normal
totalCapacity: 225.0T
totalSavings: 42.93%
usedCapacity: 67.7T
clusterAccessPolicy: lenient
dataReplicationCompliance: non_compliant
dataReplicationFactor: 3
```

如何使用命令列介面(CLI)檢查SED群集

如果無法訪問HX Connect，則可以使用CLI檢查群集是否為SED。

```
# Check if the cluster is SED capable
root@SpringpathController:~# cat /etc/springpath/sed_capability.conf
sed_capable_cluster=False

# Check if the cluster is SED enabled root@SpringpathController:~# cat /etc/springpath/sed.conf
sed_encryption_state=unknown

root@SpringpathController:~# /usr/share/springpath/storfs-appliance/sed-client.sh -l
WWN,Slot,Supported,Enabled,Locked,Vendor,Model,Serial,Size
5002538c40a42d38,1,0,0,0,Samsung,SAMSUNG_MZ7LM240HMHQ-00003,S3LKNX0K406548,228936
5000c50030278d83,25,1,1,0,MICRON,S650DC-800FIPS,ZAZ15QDM000822150Z3,763097
500a07511d38cd36,2,1,1,0,MICRON,Micron_5100_MTFDDAK960TCB_SED,17261D38CD36,915715
500a07511d38efbe,4,1,1,0,MICRON,Micron_5100_MTFDDAK960TCB_SED,17261D38EFBE,915715
500a07511d38f350,7,1,1,0,MICRON,Micron_5100_MTFDDAK960TCB_SED,17261D38F350,915715
500a07511d38eaa6,3,1,1,0,MICRON,Micron_5100_MTFDDAK960TCB_SED,17261D38EAA6,915715
500a07511d38ce80,6,1,1,0,MICRON,Micron_5100_MTFDDAK960TCB_SED,17261D38CE80,915715
500a07511d38e4fc,5,1,1,0,MICRON,Micron_5100_MTFDDAK960TCB_SED,17261D38E4FC,915715
```

HX維護模式與ESXi維護模式

當需要在屬於HX群集的伺服器上執行維護活動時，應使用HX維護模式而不是ESXi維護模式。在使用HX維護模式時，SCVM會正常關閉；而使用ESXi維護模式時，SCVM會突然關閉。

當某個節點處於維護模式時，該節點將被視為關閉，即1個節點出現故障。

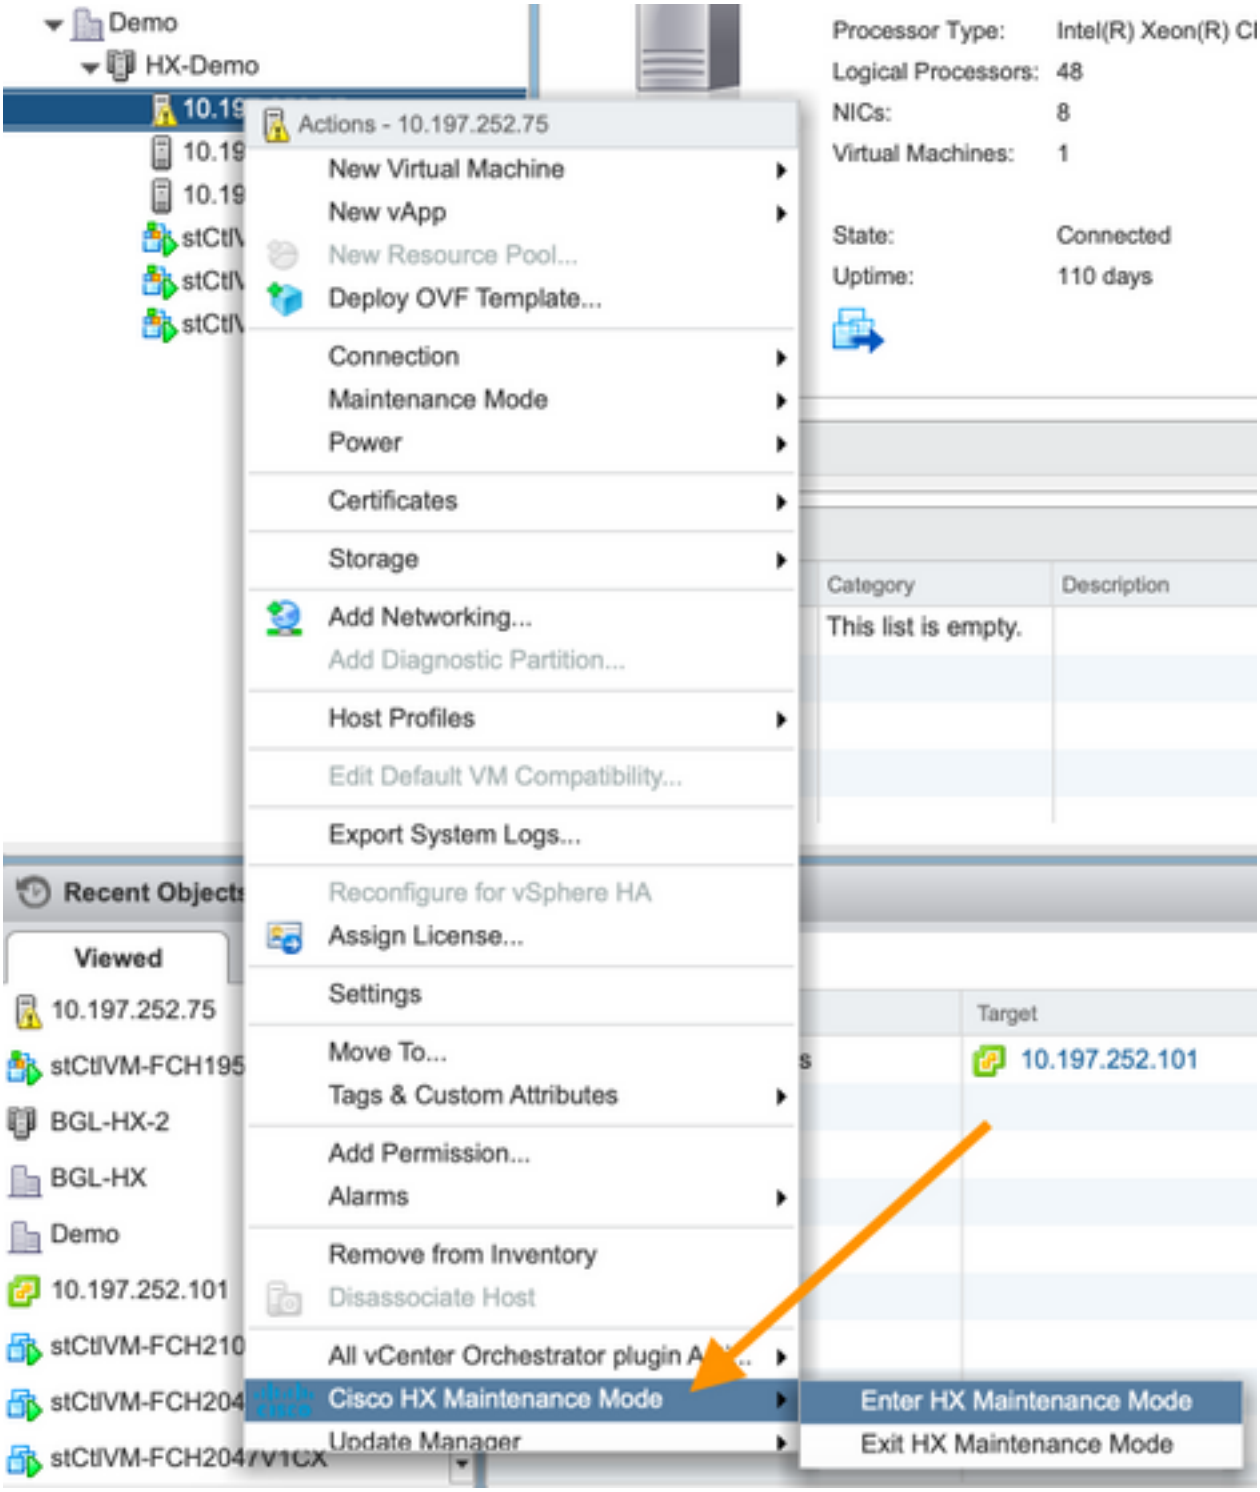

在將另一個節點移至維護模式之前，請確保群集顯示正常。

```
root@SpringpathController:~# stcli cluster storage-summary --detail
<snip>
current ensemble size:3
# of caching failures before cluster shuts down:3
minimum cache copies remaining:3
minimum data copies available for some user data:3
minimum metadata copies available for cluster metadata:3
# of unavailable nodes:0
# of nodes failure tolerable for cluster to be available:1
health state reason:storage cluster is healthy.
# of node failures before cluster shuts down:3
# of node failures before cluster goes into readonly:3
# of persistent devices failures tolerable for cluster to be available:2
```

```
# of node failures before cluster goes to enospace warn trying to move the existing data:na
# of persistent devices failures before cluster shuts down:3
# of persistent devices failures before cluster goes into readonly:3
# of caching failures before cluster goes into readonly:na
# of caching devices failures tolerable for cluster to be available:2
resiliencyInfo:
messages:
Storage cluster is healthy.
state: 1
nodeFailuresTolerable: 1
cachingDeviceFailuresTolerable: 2
<snip>
```

常見問題

思科HyperFlex M4和M5伺服器上的SCVM安裝在何處？

Cisco Hyperflex M4和M5伺服器之間的SCVM位置不同。下表列出了SCVM的位置，並提供了其他有用的資訊。

Cisco HX伺服器	ESXi	SCVM sda	快取固態驅動器(SSD)	內務管理SSD sdb1和sdb2
HX 220 M4	安全數位（SD卡）	SD卡上的3.5G	插槽2	插槽1
HX 240 M4	SD卡	在PCH控制的SSD上（esxi對此進行控制）	插槽1	在PCH控制的SSD上
HX 220 M5	M.2驅動器	M.2驅動器	插槽2	插槽1
HX 240 M5	M.2驅動器	M.2驅動器	後插槽SSD	插槽1

群集可以容忍多少個故障節點？

群集可以容忍的故障數取決於複製因子和訪問策略。

具有5個或更多節點的群集

當Replication Factor(RF)為3且Access Policy設定為Lenient時，如果2個節點出現故障，群集仍將處於讀/寫狀態。如果3個節點出現故障，則群集將關閉。

複製因子	訪問策略	失敗節點數		
		讀取/寫入	唯讀	關機
3	寬大	2	—	3
3	嚴格	1	2	3
2	寬大	1	—	2
2	嚴格	—	1	2

具有3和4個節點的群集

當RF為3且訪問策略設定為Lenient或Strict時，如果單個節點發生故障，群集仍然處於讀/寫狀態。如果2個節點出現故障，群集將關閉。

複製因子	訪問策略	失敗節點數		
		讀取/寫入	唯讀	關機
3	寬大還是嚴格	1	—	2
2	寬大	1	—	2
2	嚴格	—	1	2

3節點集群(RF:3，訪問策略：寬大)

圖形使用者介面(GUI)示例

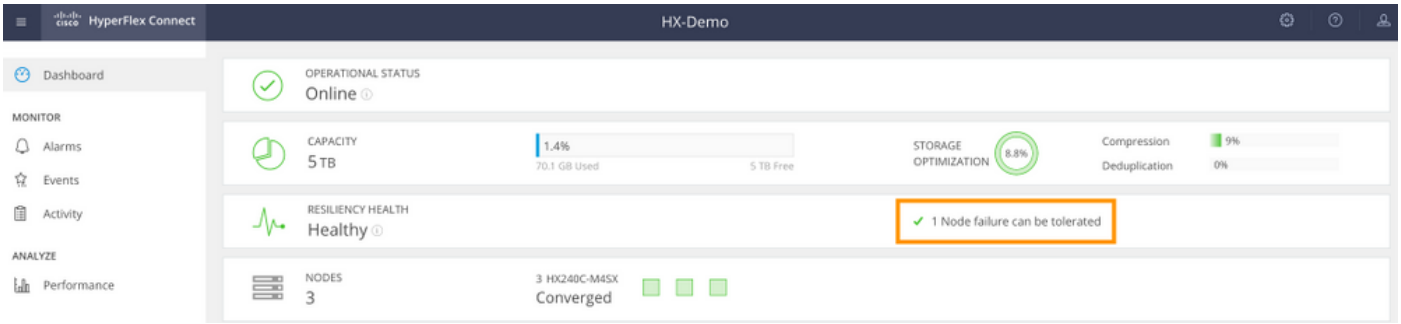

CLI範例

```
root@SpringpathController:~# stcli cluster storage-summary --detail
<snip>
current ensemble size:3
# of caching failures before cluster shuts down:3
minimum cache copies remaining:3
minimum data copies available for some user data:3
minimum metadata copies available for cluster metadata:3
# of unavailable nodes:0
# of nodes failure tolerable for cluster to be available:1
health state reason:storage cluster is healthy.
# of node failures before cluster shuts down:3
# of node failures before cluster goes into readonly:3
# of persistent devices failures tolerable for cluster to be available:2
# of node failures before cluster goes to enospace warn trying to move the existing data:na
# of persistent devices failures before cluster shuts down:3
# of persistent devices failures before cluster goes into readonly:3
# of caching failures before cluster goes into readonly:na
# of caching devices failures tolerable for cluster to be available:2
resiliencyInfo:
messages:
Storage cluster is healthy.
state: 1
<snip>
clusterAccessPolicy: lenient
```

如果其中一個SCVM關閉會發生什麼情況？虛擬機器是否繼續運行？

警告：SCVM上不支援此操作。這僅用於演示目的。

附註：確保一次僅有一個SCVM關閉。此外，請確保群集在SCVM關閉之前處於正常狀態。此方案僅用於證明，即使SCVM關閉或不可用，VM和資料儲存仍可正常運行。

虛擬機器將繼續正常工作。下面是一個輸出示例，其中SCVM已關閉，但datastore仍然被裝載並且可用。

```
[root@node1:~] vim-cmd vmsvc/getallvms
Vmid Name File Guest OS Version Annotation
1 stCtlVM-F 9H [SpringpathDS-F 9H] stCtlVM-F 9H/stCtlVM-F 9H.vmx ubuntu64Guest vmx-13

[root@node1:~] vim-cmd vmsvc/power.off 1
Powering off VM:

[root@node1:~] vim-cmd vmsvc/power.getstate 1
Retrieved runtime info
Powered off

[root@node1:~] esxcfg-nas -l
Test is 10.197.252.106:Test from 3203172317343203629-5043383143428344954 mounted available
ReplSec is 10.197.252.106:ReplSec from 3203172317343203629-5043383143428344954 mounted available
New_DS is 10.197.252.106:New_DS from 3203172317343203629-5043383143428344954 mounted available
```

SCVM上的VMware硬體版本已更新。現在怎麼辦？

警告：SCVM上不支援此操作。這僅用於演示目的。

在**相容性 > 升級VM相容性**中通過編輯VM設定來升級VMware硬體版本是vSphere Web客戶端在SCVM上不受支援的操作。SCVM將在HX Connect中報告為離線。

```
root@SpringpathController0      UE:~# lsblk
NAME   MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda      8:0    0  2.5G  0 disk
`-sda1   8:1    0  2.5G  0 part /
sdb      8:16   0  100G  0 disk
|-sdb1   8:17   0   64G  0 part /var/stv
`-sdb2   8:18   0   24G  0 part /var/zookeeper


root@SpringpathController0      UE:~# lsscsi
[2:0:0:0]    disk    VMware   Virtual disk     2.0   /dev/sda
[2:0:1:0]    disk    VMware   Virtual disk     2.0   /dev/sdb


root@SpringpathController0      UE:~# cat /var/log/springpath/diskslotmap-v2.txt
1.11.1:5002538a17221ab0:SAMSUNG:MZIES800HMHP/003:S1N2NY0J201389:EM19:SAS:SSD:763097:Inactive:/de
v/sdc
1.11.2:5002538c405537e0:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
98:GXT51F3Q:SATA:SSD:3662830:Inactive:/dev/sdd
1.11.3:5002538c4055383a:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
88:GXT51F3Q:SATA:SSD:3662830:Inactive:/dev/sde
1.11.4:5002538c40553813:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
49:GXT51F3Q:SATA:SSD:3662830:Inactive:/dev/sdf
1.11.5:5002538c4055380e:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
44:GXT51F3Q:SATA:SSD:3662830:Inactive:/dev/sdg
1.11.6:5002538c40553818:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
54:GXT51F3Q:SATA:SSD:3662830:Inactive:/dev/sdh
1.11.7:5002538c405537d1:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
83:GXT51F3Q:SATA:SSD:3662830:Inactive:/dev/sdi
1.11.8:5002538c405537d8:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
90:GXT51F3Q:SATA:SSD:3662830:Inactive:/dev/sdj
1.11.9:5002538c4055383b:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
```

89:GXT51F3Q:SATA:SSD:3662830:**Inactive**:/dev/sdk
1.11.10:5002538c4055381f:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
61:GXT51F3Q:SATA:SSD:3662830:**Inactive**:/dev/sdl
1.11.11:5002538c40553823:Samsung:SAMSUNG_MZ7LM3T8HMLP-00003:S
65:GXT51F3Q:SATA:SSD:3662830:**Inactive**:/dev/sdm

注意：如果意外執行此操作，請致電思科支援尋求進一步幫助。需要重新部署SCVM。