Ultra-M UCS 240M4單硬碟故障 — 熱插拔過程 — CPAR

目錄

簡介

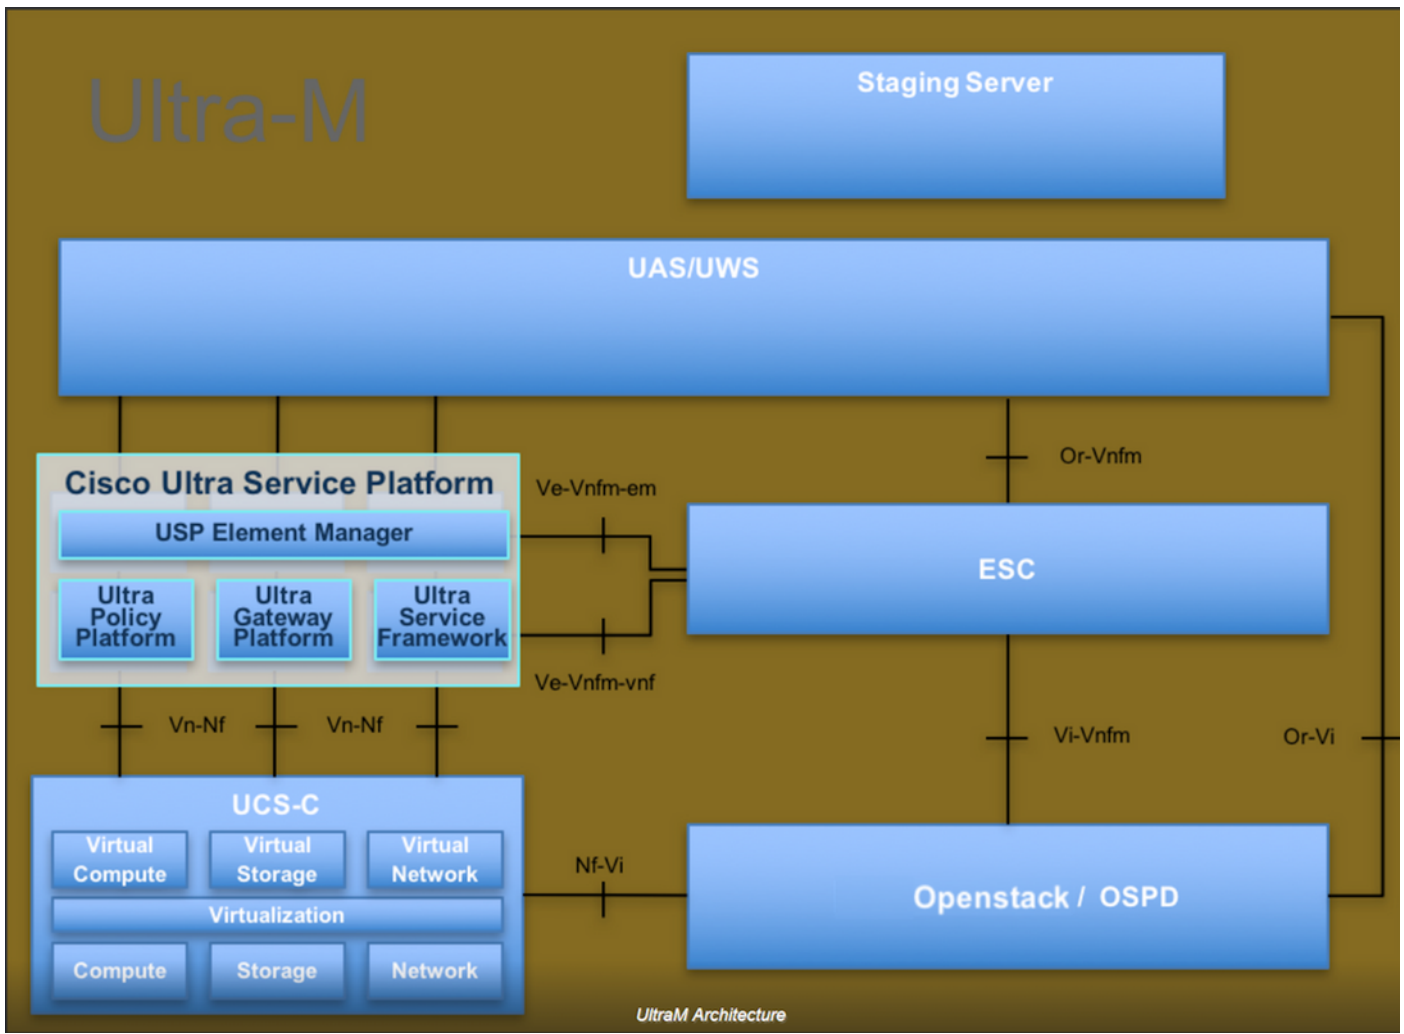本文檔介紹在Ultra-M設定中更換伺服器中有故障的硬碟驅動器(HDD)所需的步驟。

此過程適用於NEWTON版本的Openstack環境，其中ESC不管理CPAR，並且CPAR直接安裝在部署在Openstack上的虛擬機器(VM)上。

背景資訊

Ultra-M是經過預先打包和驗證的虛擬化移動資料包核心解決方案，旨在簡化虛擬網絡功能(VNF)的部署。 OpenStack是適用於Ultra-M的Virtual Infrastructure Manager(VIM)，包含以下節點型別：

- 計算
- 對象儲存磁碟 — 計算（OSD — 計算）
- 控制器
- OpenStack平台 — 導向器(OSPD)

Ultra-M的高級體系結構及涉及的元件如下圖所示：

UltraM Architecture

本文檔面向熟悉Cisco Ultra-M平台的思科人員，詳細說明了更換OSPD伺服器時在OpenStack級別上需要執行的步驟。

　　附註：Ultra M 5.1.x版本用於定義本文檔中的過程。

縮寫

VNF	虛擬網路功能
MoP	程式方法
OSD	對象儲存磁碟
OSPD	OpenStack平台導向器
硬碟	硬碟驅動器
固態硬碟	固態驅動器
VIM	虛擬基礎架構管理員
虛擬機器	虛擬機器
EM	元素管理器
UAS	Ultra自動化服務
UUID	通用唯一識別符號

MoP工作流

```
┌─────────────────────────┐
│  Identify the  Faulty HDD │
│          disk            │
└─────────────────────────┘
            │
            ▼
┌─────────────────────────┐
│  Perform the openstack role│
│   specific health checks  │
└─────────────────────────┘
            │
            ▼
┌─────────────────────────┐
│   Hot swap the faulty HDD  │
│          disk            │
└─────────────────────────┘
            │
            ▼
┌─────────────────────────┐
│   Wait for the  data sync to │
│         complete          │
└─────────────────────────┘
            │
            ▼
┌─────────────────────────┐
│  Verify the newly added disk│
│  joins the RAID1 redundancy │
│   of BOOTOS & health is    │
│           good            │
└─────────────────────────┘
```
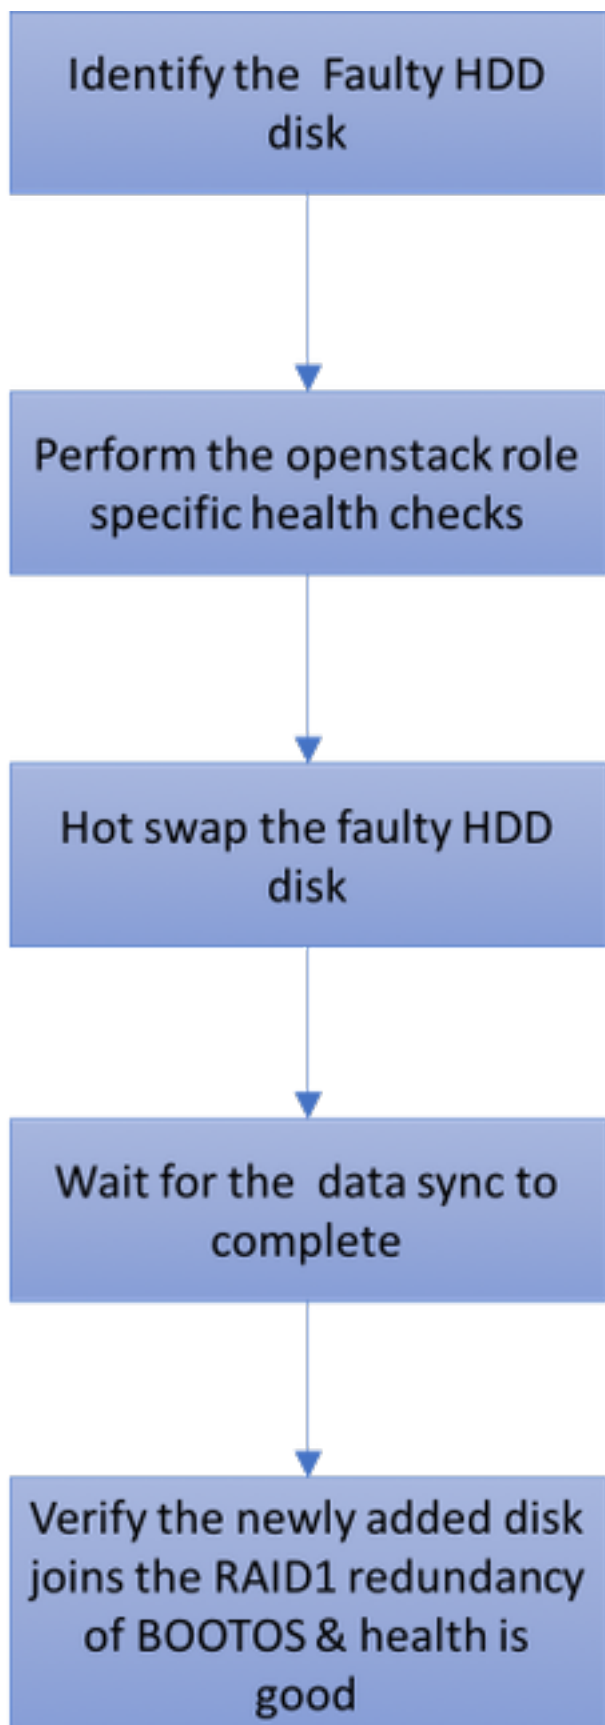
單HDD故障

1. 每台裸機伺服器都配置有兩個HDD驅動器，以便在Raid 1配置中充當BOOT DISK。在單HDD故障的情況下，由於存在RAID 1級冗餘，故障的HDD驅動器可以熱交換。
2. 有關更換UCS C240 M4伺服器上故障元件的步驟，請參閱更換伺服器元件。

3. 在單HDD故障的情況下，僅熱更換有故障的硬碟，因此更換新磁碟後不需要執行BIOS升級過程。
4. 更換磁碟後，必須等待磁碟之間的資料同步。這可能需要幾個小時才能完成。
5. 在基於Openstack(Ultra-M)的解決方案中，UCS 240M4裸機伺服器可以承擔以下角色之一：計算、OSD — 計算、控制器和OSPD。在每一個伺服器角色中處理單硬碟故障所需的步驟都相同，本節介紹了在熱插拔磁碟之前要執行的運行狀況檢查。

Compute Server上的單HDD故障

1. 如果在充當計算節點的UCS 240M4中觀察到硬碟驅動器故障，請在執行故障磁碟的熱交換之前執行此運行狀況檢查。
2. 確定在此伺服器上運行的VM，並驗證這些功能的狀態是否良好。

確定託管於計算節點中的虛擬機器

確定託管在計算伺服器上的虛擬機器，並驗證它們是否處於活動狀態且正在運行。

```
[stack@director ~]$ nova list
| 46b4b9eb-a1a6-425d-b886-a0ba760e6114 | AAA-CPAR-testing-instance | pod2-stack-compute-
4.localdomain |
```

運行狀況檢查

步驟1.在作業系統(OS)級別運行命令**/opt/CSCOar/bin/arstatus。**

```
[root@aaa04 ~]# /opt/CSCOar/bin/arstatus
Cisco Prime AR RADIUS server running        (pid: 24834)
Cisco Prime AR Server Agent running         (pid: 24821)
Cisco Prime AR MCD lock manager running     (pid: 24824)
Cisco Prime AR MCD server running           (pid: 24833)
Cisco Prime AR GUI running                  (pid: 24836)
SNMP Master Agent running                  (pid: 24835)
[root@wscaaa04 ~]#
```
步驟2.在作業系統級別運行命令**/opt/CSCOar/bin/aregcmd**，然後輸入管理員憑據。驗證CPAR Health（CPAR運行狀況）是否為10/10，並退出CPAR CLI。

```
[root@aaa02 logs]# /opt/CSCOar/bin/aregcmd
Cisco Prime Access Registrar 7.3.0.1 Configuration Utility
Copyright (C) 1995-2017 by Cisco Systems, Inc.  All rights reserved.
Cluster:
User: admin
Passphrase:
Logging in to localhost
[ //localhost ]

    LicenseInfo = PAR-NG-TPS 7.2(100TPS:)

                 PAR-ADD-TPS 7.2(2000TPS:)
```

```
                    PAR-RDDR-TRX 7.2()

                    PAR-HSS 7.2()

    Radius/

    Administrators/
Server 'Radius' is Running, its health is 10 out of 10
--> exit
```

步驟3.運行命令netstat | grep diameter並驗證所有Diameter Routing Agent(DRA)連線已建立。

此處提到的輸出適用於預期存在Diameter連結的環境。如果顯示的連結較少，則表示與需要分析的DRA斷開連線。

```
[root@aa02 logs]# netstat | grep diameter
tcp        0            0 aaa02.aaa.epc.:77 mp1.dra01.d:diameter ESTABLISHED
tcp        0            0 aaa02.aaa.epc.:36 tsa6.dra01:diameter ESTABLISHED
tcp        0            0 aaa02.aaa.epc.:47 mp2.dra01.d:diameter ESTABLISHED
tcp        0            0 aaa02.aaa.epc.:07 tsa5.dra01:diameter ESTABLISHED
tcp        0            0 aaa02.aaa.epc.:08 np2.dra01.d:diameter ESTABLISHED
```

步驟4.檢查TPS日誌是否顯示CPAR正在處理的請求。突出顯示的值代表需要注意的TPS。

TPS的值不能超過1500。

```
[root@wscaaa04 ~]# tail -f /opt/CSCOar/logs/tps-11-21-2017.csv
11-21-2017,23:57:35,263,0
11-21-2017,23:57:50,237,0
11-21-2017,23:58:05,237,0
11-21-2017,23:58:20,257,0
11-21-2017,23:58:35,254,0
11-21-2017,23:58:50,248,0
11-21-2017,23:59:05,272,0
11-21-2017,23:59:20,243,0
11-21-2017,23:59:35,244,0
11-21-2017,23:59:50,233,0
```

步驟5.在name_radius_1_log中查詢任何「錯誤」或「警報」消息

```
[root@aaa02 logs]# grep -E "error|alarm" name_radius_1_log
```

步驟6.為了驗證CPAR進程使用的記憶體量，請運行命令：

```
top | grep radius

[root@sfraaa02 ~]# top | grep radius
27008 root      20    0 20.228g 2.413g  11408 S 128.3  7.7   1165:41 radius
```

此突出顯示的值必須小於7 Gb，這是應用程式級別允許的最大值。

步驟7.要驗證磁碟利用率，請運行命令df -h。

```
[root@aaa02 ~]# df -h
Filesystem                    Size  Used Avail Use% Mounted on
/dev/mapper/vg_arucsvm51-lv_root 26G   21G  4.1G  84% /
tmpfs                         1.9G  268K  1.9G   1% /dev/shm
/dev/sda1                     485M   37M  424M   8% /boot
```

```
/dev/mapper/vg_arucsvm51-lv_home 23G  4.3G  17G  21% /home
```
此總值必須低於80%，如果超過80%，則識別不必要檔案並進行清理。

步驟8.驗證是否未生成「core」檔案。

- 當CPAR無法處理異常時，核心檔案會在應用程式崩潰時生成，它將在以下兩個位置生成：

```
[root@aaa02 ~]# cd /cisco-ar/
[root@aaa02 ~]# cd /cisco-ar/bin
```
這兩個位置中不能有任何核心檔案。如果找到，請提出Cisco TAC案例，以確定此類異常的根本原因，並附加核心檔案以進行調試。

- 如果運行狀況檢查正常，請繼續執行有故障的磁碟熱交換過程，並等待資料同步，因為需要幾個小時才能完成。

[更換伺服器元件](#)

- 重複運行狀況檢查過程，以確認託管在計算節點上的虛擬機器的運行狀況已恢復。

控制器伺服器上的單HDD故障

- 如果在充當控制器節點的UCS 240M4中觀察到HDD驅動器故障，請在執行故障磁碟的熱交換之前執行這些運行狀況檢查。

- 檢查控制器上的起搏器狀態。

- 登入其中一個作用中控制器並檢查心臟起搏器狀態。所有服務必須在可用控制器上運行並在出現故障的控制器上停止。

```
[heat-admin@pod2-stack-controller-0 ~]$ sudo pcs status
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod2-stack-controller-2 (version 1.1.15-11.el7_3.4-e174ec8) - partition with quorum
Last updated: Tue Jul 10 10:04:15 2018Last change: Fri Jul 6 09:03:35 2018 by root via
crm_attribute on pod2-stack-controller-0

3 nodes and 19 resources configured

Online: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]

Full list of resources:

ip-11.120.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
Clone Set: haproxy-clone [haproxy]
Started: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
Master/Slave Set: galera-master [galera]
Masters: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
ip-192.200.0.110(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
ip-11.120.0.44(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
ip-11.118.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
Clone Set: rabbitmq-clone [rabbitmq]
Started: [ pod2-stack-controller-0 pod2-stack-controller-1 pod2-stack-controller-2 ]
ip-10.225.247.214(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-1
Master/Slave Set: redis-master [redis]
Masters: [ pod2-stack-controller-2 ]
```

```
Slaves: [ pod2-stack-controller-0 pod2-stack-controller-1 ]
ip-11.119.0.49(ocf::heartbeat:IPaddr2):Started pod2-stack-controller-2
openstack-cinder-volume(systemd:openstack-cinder-volume):Started pod2-stack-controller-1

Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

- 檢查活動控制器中的MariaDB狀態。

```
[stack@director ~]$ nova list | grep control
| b896c73f-d2c8-439c-bc02-7b0a2526dd70 | pod2-stack-controller-0 | ACTIVE | - | Running |
ctlplane=192.200.0.113 |
| 2519ce67-d836-4e5f-a672-1a915df75c7c | pod2-stack-controller-1 | ACTIVE | - | Running |
ctlplane=192.200.0.105 |
| e19b9625-5635-4a52-a369-44310f3e6a21 | pod2-stack-controller-2 | ACTIVE | - | Running |
ctlplane=192.200.0.120 |

[stack@director ~]$ for i in 192.200.0.102 192.200.0.110 ; do echo "*** $i ***" ; ssh heat-
admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local_ state_comment'\" ; sudo mysql --
exec=\"SHOW STATUS LIKE 'wsrep_cluster_size'\""; done 192.200.0.110 ; do echo "*** $i ***" ; ssh
heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local_st5 192.200.0.110 ; do echo
"*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local_st ; do echo
"*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local_st3 ; do
echo "*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local_st ; do
echo "*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local_s1 ; do
echo "*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local_9 ; do
echo "*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_local2 ; do
echo "*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_loca. ; do
echo "*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_loc2 ; do
echo "*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_lo0 ; do echo
"*** $i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_l0 ; do echo "***
$i ***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep_. ; do echo "*** $i
***" ; ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsrep0 ; do echo "*** $i ***" ;
ssh heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsre. ; do echo "*** $i ***" ; ssh
heat-admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'wsr1 ; do echo "*** $i ***" ; ssh heat-
admin@$i "sudo mysql --exec=\"SHOW STATUS LIKE 'ws2 ; do echo "*** $i ***" ; ssh heat-admin@$i
"sudo mysql --exec=\"SHOW STATUS LIKE 'w0 ; do echo "*** $i ***" ; ssh heat-admin@$i "sudo mysql
--exec=\"SHOW STATUS LIKE '

*** 192.200.0.102 ***
Variable_nameValue
wsrep_local_state_commentSynced
Variable_nameValue
wsrep_cluster_size2
*** 192.200.0.110 ***
Variable_nameValue
wsrep_local_state_commentSynced
Variable_nameValue
wsrep_cluster_size2
```

- 驗證每個作用中控制器是否存在以下線路：

```
wsrep_local_state_comment: Synced

wsrep_cluster_size: 2
```

- 檢查作用中控制器中的Rabbitmq狀態。

```
[heat-admin@pod2-stack-controller-0 ~]$ sudo rabbitmqctl cluster_status
Cluster status of node 'rabbit@pod2-stack-controller-0' ...
[{nodes,[{disc,['rabbit@pod2-stack-controller-0',
'rabbit@pod2-stack-controller-1',
'rabbit@pod2-stack-controller-2']}]},
{running_nodes,['rabbit@pod2-stack-controller-1',
'rabbit@pod2-stack-controller-2',
'rabbit@pod2-stack-controller-0']},
{cluster_name,<<"rabbit@pod2-stack-controller-1.localdomain">>},
{partitions,[]},
{alarms,[{'rabbit@pod2-stack-controller-1',[]},
{'rabbit@pod2-stack-controller-2',[]},
{'rabbit@pod2-stack-controller-0',[]}]}]}
```

- 如果運行狀況檢查正常，請繼續執行有故障的磁碟熱交換過程，並等待資料同步，因為需要幾個小時才能完成。

[更換伺服器元件](#)

- 重複運行狀況檢查過程，以確認控制器上的運行狀況已恢復。

OSD-Compute Server上的單硬碟故障

- 如果在充當OSD-Compute節點的UCS 240M4中觀察到HDD驅動器故障，請在執行故障磁碟的熱交換之前執行運行狀況檢查。

1. 確定OSD計算節點中託管的VM
2. 識別託管在計算伺服器上的VM

```
[stack@director ~]$ nova list
| 46b4b9eb-a1a6-425d-b886-a0ba760e6114 | AAA-CPAR-testing-instance | pod2-stack-compute-
4.localdomain |
```

- CEPH進程在osd-compute伺服器上處於活動狀態。

```
[heat-admin@pod2-stack-osd-compute-1 ~]$ systemctl list-units *ceph*

UNIT LOAD ACTIVE SUB DESCRIPTION
var-lib-ceph-osd-ceph\x2d1.mount loaded active mounted /var/lib/ceph/osd/ceph-1
var-lib-ceph-osd-ceph\x2d10.mount loaded active mounted /var/lib/ceph/osd/ceph-10
var-lib-ceph-osd-ceph\x2d4.mount loaded active mounted /var/lib/ceph/osd/ceph-4
var-lib-ceph-osd-ceph\x2d7.mount loaded active mounted /var/lib/ceph/osd/ceph-7
ceph-osd@1.service loaded active running Ceph object storage daemon
ceph-osd@10.service loaded active running Ceph object storage daemon
ceph-osd@4.service loaded active running Ceph object storage daemon
ceph-osd@7.service loaded active running Ceph object storage daemon
system-ceph\x2ddisk.slice loaded active active system-ceph\x2ddisk.slice
system-ceph\x2dosd.slice loaded active active system-ceph\x2dosd.slice
ceph-mon.target loaded active active ceph target allowing to start/stop all ceph-mon@.service
instances at once
ceph-osd.target loaded active active ceph target allowing to start/stop all ceph-osd@.service
instances at once
ceph-radosgw.target loaded active active ceph target allowing to start/stop all ceph-
radosgw@.service instances at once
ceph.target loaded active active ceph target allowing to start/stop all ceph*@.service instances
at once
```

```
LOAD = Reflects whether the unit definition was properly loaded.
ACTIVE = The high-level unit activation state, i.e. generalization of SUB.
SUB = The low-level unit activation state, values depend on unit type.

14 loaded units listed. Pass --all to see loaded but inactive units, too.
To show all installed unit files use 'systemctl list-unit-files'.
```

- 驗證OSD（HDD磁碟）到日誌(SSD)的對映是否正常。

```
[heat-admin@pod2-stack-osd-compute-1 ~]$ sudo ceph-disk list
/dev/sda :
/dev/sda1 other, iso9660
/dev/sda2 other, xfs, mounted on /
/dev/sdb :
/dev/sdb1 ceph journal, for /dev/sdc1
/dev/sdb3 ceph journal, for /dev/sdd1
/dev/sdb2 ceph journal, for /dev/sde1
/dev/sdb4 ceph journal, for /dev/sdf1
/dev/sdc :
/dev/sdc1 ceph data, active, cluster ceph, osd.1, journal /dev/sdb1
/dev/sdd :
/dev/sdd1 ceph data, active, cluster ceph, osd.7, journal /dev/sdb3
/dev/sde :
/dev/sde1 ceph data, active, cluster ceph, osd.4, journal /dev/sdb2
/dev/sdf :
/dev/sdf1 ceph data, active, cluster ceph, osd.10, journal /dev/sdb4
```

- 驗證ceph運行狀況和osd樹對映是否良好。

```
[heat-admin@pod2-stack-osd-compute-1 ~]$ sudo ceph -s
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_OK
monmap e1: 3 mons at {pod2-stack-controller-0=11.118.0.10:6789/0,pod2-stack-controller-
1=11.118.0.11:6789/0,pod2-stack-controller-2=11.118.0.12:6789/0}
election epoch 10, quorum 0,1,2 pod2-stack-controller-0,pod2-stack-controller-1,pod2-stack-
controller-2
osdmap e81: 12 osds: 12 up, 12 in
flags sortbitwise,require_jewel_osds
pgmap v23095222: 704 pgs, 6 pools, 809 GB data, 424 kobjects
2418 GB used, 10974 GB / 13393 GB avail
704 active+clean
client io 1329 kB/s wr, 0 op/s rd, 122 op/s wr


[heat-admin@pod2-stack-osd-compute-1 ~]$ sudo ceph osd tree
ID WEIGHT TYPE NAME UP/DOWN REWEIGHT PRIMARY-AFFINITY
-1 13.07996 root default
-2 4.35999 host pod2-stack-osd-compute-0
0 1.09000 osd.0 up 1.00000 1.00000
3 1.09000 osd.3 up 1.00000 1.00000
6 1.09000 osd.6 up 1.00000 1.00000
9 1.09000 osd.9 up 1.00000 1.00000
-3 4.35999 host pod2-stack-osd-compute-1
1 1.09000 osd.1 up 1.00000 1.00000
4 1.09000 osd.4 up 1.00000 1.00000
7 1.09000 osd.7 up 1.00000 1.00000
10 1.09000 osd.10 up 1.00000 1.00000
-4 4.35999 host pod2-stack-osd-compute-2
2 1.09000 osd.2 up 1.00000 1.00000
5 1.09000 osd.5 up 1.00000 1.00000
8 1.09000 osd.8 up 1.00000 1.00000
```

```
11 1.09000 osd.11 up 1.00000 1.00000
```

- 如果運行狀況檢查正常,請繼續執行有故障的磁碟熱插拔過程,並等待資料同步,因為它需要幾個小時才能完成。

[更換伺服器元件](#)

- 重複運行狀況檢查過程,以確認在OSD-Compute節點上託管的VM的運行狀況已恢復。

OSPD伺服器上的單HDD故障

- 如果在充當OSPD節點的UCS 240M4中觀察到HDD驅動器故障,請在執行故障磁碟的熱交換之前執行運行狀況檢查。

- 檢查openstack堆疊的狀態和節點清單。

```
[stack@director ~]$ source stackrc
[stack@director ~]$ openstack stack list --nested
[stack@director ~]$ ironic node-list
[stack@director ~]$ nova list
```

- 從OSP-D節點檢查是否所有底層雲服務都處於已載入、活動和運行狀態。

```
[stack@director ~]$ systemctl list-units "openstack*" "neutron*" "openvswitch*"
UNIT LOAD ACTIVE SUB DESCRIPTION
neutron-dhcp-agent.service loaded active running OpenStack Neutron DHCP Agent
neutron-metadata-agent.service loaded active running OpenStack Neutron Metadata Agent
neutron-openvswitch-agent.service loaded active running OpenStack Neutron Open vSwitch Agent
neutron-server.service loaded active running OpenStack Neutron Server
openstack-aodh-evaluator.service loaded active running OpenStack Alarm evaluator service
openstack-aodh-listener.service loaded active running OpenStack Alarm listener service
openstack-aodh-notifier.service loaded active running OpenStack Alarm notifier service
openstack-ceilometer-central.service loaded active running OpenStack ceilometer central agent
openstack-ceilometer-collector.service loaded active running OpenStack ceilometer collection
service
openstack-ceilometer-notification.service loaded active running OpenStack ceilometer
notification agent
openstack-glance-api.service loaded active running OpenStack Image Service (code-named Glance)
API server
openstack-glance-registry.service loaded active running OpenStack Image Service (code-named
Glance) Registry server
openstack-heat-api-cfn.service loaded active running Openstack Heat CFN-compatible API Service
openstack-heat-api.service loaded active running OpenStack Heat API Service
openstack-heat-engine.service loaded active running Openstack Heat Engine Service
openstack-ironic-api.service loaded active running OpenStack Ironic API service
openstack-ironic-conductor.service loaded active running OpenStack Ironic Conductor service
openstack-ironic-inspector-dnsmasq.service loaded active running PXE boot dnsmasq service for
Ironic Inspector
openstack-ironic-inspector.service loaded active running Hardware introspection service for
OpenStack Ironic
openstack-mistral-api.service loaded active running Mistral API Server
openstack-mistral-engine.service loaded active running Mistral Engine Server
openstack-mistral-executor.service loaded active running Mistral Executor Server
openstack-nova-api.service loaded active running OpenStack Nova API Server
openstack-nova-cert.service loaded active running OpenStack Nova Cert Server
openstack-nova-compute.service loaded active running OpenStack Nova Compute Server
openstack-nova-conductor.service loaded active running OpenStack Nova Conductor Server
openstack-nova-scheduler.service loaded active running OpenStack Nova Scheduler Server
openstack-swift-account-reaper.service loaded active running OpenStack Object Storage (swift) -
Account Reaper
```

```
openstack-swift-account.service loaded active running OpenStack Object Storage (swift) - Account
Server
openstack-swift-container-updater.service loaded active running OpenStack Object Storage (swift)
- Container Updater
openstack-swift-container.service loaded active running OpenStack Object Storage (swift) -
Container Server
openstack-swift-object-updater.service loaded active running OpenStack Object Storage (swift) -
Object Updater
openstack-swift-object.service loaded active running OpenStack Object Storage (swift) - Object
Server
openstack-swift-proxy.service loaded active running OpenStack Object Storage (swift) - Proxy
Server
openstack-zaqar.service loaded active running OpenStack Message Queuing Service (code-named
Zaqar) Server
openstack-zaqar@1.service loaded active running OpenStack Message Queuing Service (code-named
Zaqar) Server Instance 1
openvswitch.service loaded active exited Open vSwitch

LOAD = Reflects whether the unit definition was properly loaded.
ACTIVE = The high-level unit activation state, i.e. generalization of SUB.
SUB = The low-level unit activation state, values depend on unit type.
lines 1-43
lines 2-44 37 loaded units listed. Pass --all to see loaded but inactive units, too.
To show all installed unit files use 'systemctl list-unit-files'.
lines 4-46/46 (END) lines 4-46/46 (END) lines 4-46/46 (END) lines 4-46/46 (END) lines 4-46/46
(END)
```

- 如果運行狀況檢查正常，請繼續執行有故障的磁碟熱插拔過程，並等待資料同步，因為它需要幾個小時才能完成。

更換伺服器元件

- 重複運行狀況檢查過程，以確認OSPD節點的運行狀況已恢復。