PCRF替代OSD-Compute UCS 240M4

目錄

簡介

本文檔介紹在託管思科策略套件(CPS)虛擬網路功能(VNF)的Ultra-M設定中更換有故障的osd-compute伺服器所需的步驟。

背景資訊

本文檔面向熟悉Cisco Ultra-M平台的思科人員，詳細說明了在OSD-Compute Server更換時在OpenStack和CPS VNF級別需要執行的步驟。

> **附註**：Ultra M 5.1.x版本用於定義本文檔中的過程。

運行狀況檢查

在替換Osd-Compute節點之前，請務必檢查Red Hat OpenStack平台環境的當前狀態。建議您檢查當前狀態，以避免在計算替換過程開啟時出現複雜情況。

來自OSPD

```
[root@director ~]$ su - stack
[stack@director ~]$ cd ansible
[stack@director ansible]$ ansible-playbook -i inventory-new openstack_verify.yml  -e
```

```
platform=pcrf
```

步驟1.從每15分鐘生成的超健康報告中驗證系統運行狀況。

```
[stack@director ~]# cd /var/log/cisco/ultram-health
```

檢查檔案ultrum_health_os.report。

唯一的服務應顯示為XXX狀態是neutron-sriov-nic-agent.service。

步驟2.檢查rabbitmq是否對所有控制器運行，這些控制器又從OSPD運行。

```
[stack@director ~]# for i in $(nova list| grep controller | awk '{print $12}'| sed
's/ctlplane=//g') ; do (ssh -o StrictHostKeyChecking=no heat-admin@$i "hostname;sudo rabbitmqctl
eval 'rabbit_diagnostics:maybe_stuck().'" ) & done
```

步驟3.驗證是否已啟用石碑。

```
[stack@director ~]# sudo pcs property show stonith-enabled
```

對於所有控制器，驗證PC狀態

- 所有控制器節點均在haproxy-clone下啟動
- 所有控制器節點在galera下都是Master
- 所有控制器節點均在Rabbitmq下啟動
- 在redis下，1個控制器節點為Master,2個從節點為Slave

來自OSPD

```
[stack@director ~]$ for i in $(nova list| grep controller | awk '{print $12}'| sed
's/ctlplane=//g') ; do (ssh -o StrictHostKeyChecking=no heat-admin@$i "hostname;sudo pcs status"
) ;done
```

步驟4.驗證所有openstack服務是否處於活動狀態，從OSPD運行以下命令：

```
[stack@director ~]# sudo systemctl list-units "openstack*" "neutron*" "openvswitch*"
```

步驟5.驗證控制器的CEPH狀態為HEALTH_OK。

```
[stack@director ~]# for i in $(nova list| grep controller | awk '{print $12}'| sed
's/ctlplane=//g') ; do (ssh -o StrictHostKeyChecking=no heat-admin@$i "hostname;sudo ceph -s" )
;done
```

步驟6.檢驗OpenStack元件日誌。尋找任何錯誤：

```
Neutron:
[stack@director ~]# sudo tail -n 20 /var/log/neutron/{dhcp-agent,l3-agent,metadata-
agent,openvswitch-agent,server}.log

Cinder:
[stack@director ~]# sudo tail -n 20 /var/log/cinder/{api,scheduler,volume}.log

Glance:
[stack@director ~]# sudo tail -n 20 /var/log/glance/{api,registry}.log
```

步驟7.從OSPD為API執行這些驗證。

```
[stack@director ~]$ source
```

```
[stack@director ~]$ nova list

[stack@director ~]$ glance image-list

[stack@director ~]$ cinder list

[stack@director ~]$ neutron net-list
```

步驟8.檢驗服務的運行狀況。

```
Every service status should be “up”:
[stack@director ~]$ nova service-list

Every service status should be “ :-)”:
[stack@director ~]$ neutron agent-list

Every service status should be “up”:
[stack@director ~]$ cinder service-list
```

備份

在進行恢復時，思科建議使用以下步驟備份OSPD資料庫。

步驟1.執行Mysql轉儲。

```
[root@director ~]# mysqldump --opt --all-databases > /root/undercloud-all-databases.sql
[root@director ~]# tar --xattrs -czf undercloud-backup-`date +%F`.tar.gz /root/undercloud-all-databases.sql
/etc/my.cnf.d/server.cnf /var/lib/glance/images /srv/node /home/stack
tar: Removing leading `/' from member names
```

此過程可確保在不影響任何例項可用性的情況下替換節點。

步驟2.從群集管理器虛擬機器備份CPS虛擬機器：

```
[root@CM ~]# config_br.py -a export --all /mnt/backup/CPS_backup_$(date +\%Y-\%m-\%d).tar.gz

or

[root@CM ~]# config_br.py -a export --mongo-all --svn --etc --grafanadb --auth-htpasswd --haproxy /mnt/backup/$(hostname)_backup_all_$(date +\%Y-\%m-\%d).tar.gz
```

確定OSD計算節點中託管的VM

確定託管於計算伺服器上的VM:

步驟1。計算伺服器包含彈性服務控制器(ESC)。

```
[stack@director ~]$ nova list --field name,host,networks | grep osd-compute-1
| 50fd1094-9c0a-4269-b27b-cab74708e40c | esc | pod1-osd-compute-0.localdomain
| tb1-orch=172.16.180.6; tb1-mgmt=172.16.181.3
```

附註：此處顯示的輸出中，第一列對應於通用唯一識別符號(UUID)，第二列是VM名稱，第三列是存在VM的主機名。此輸出的引數將在後續章節中使用。

附註：如果要替換的OSD-Compute節點已完全關閉且不可訪問，請繼續至標題為「從新星聚合清單中刪除Osd-Compute節點」的部分。否則，請從下一部分繼續。

步驟2.驗證CEPH是否有允許刪除單個OSD伺服器的可用容量。

```
[root@pod1-osd-compute-0 ~]# sudo ceph df
GLOBAL:
   SIZE       AVAIL      RAW USED     %RAW USED
    13393G     11804G        1589G         11.87
POOLS:
   NAME        ID    USED      %USED     MAX AVAIL     OBJECTS
   rbd         0          0         0         3876G           0
   metrics     1     4157M      0.10         3876G      215385
   images      2     6731M      0.17         3876G         897
   backups     3         0         0         3876G           0
   volumes     4      399G      9.34         3876G      102373
   vms         5      122G      3.06         3876G       31863
```

步驟3.驗證osd-compute伺服器上的ceph osd樹狀態是否為up。

```
[heat-admin@pod1-osd-compute-0 ~]$ sudo ceph osd tree
ID WEIGHT   TYPE NAME                          UP/DOWN REWEIGHT PRIMARY-AFFINITY
-1 13.07996 root default
-2  4.35999     host pod1-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 pod1-osd-compute-2
```

```
 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 pod1-osd-compute-1
 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
```

步驟4. CEPH進程在osd-compute伺服器上處於活動狀態。

```
[root@pod1-osd-compute-0 ~]# systemctl list-units *ceph*
UNIT                                LOAD   ACTIVE SUB     DESCRIPTION
var-lib-ceph-osd-ceph\x2d11.mount loaded active mounted /var/lib/ceph/osd/ceph-11
var-lib-ceph-osd-ceph\x2d2.mount  loaded active mounted /var/lib/ceph/osd/ceph-2
var-lib-ceph-osd-ceph\x2d5.mount  loaded active mounted /var/lib/ceph/osd/ceph-5
var-lib-ceph-osd-ceph\x2d8.mount  loaded active mounted /var/lib/ceph/osd/ceph-8
ceph-osd@11.service               loaded active running Ceph object storage daemon
ceph-osd@2.service                loaded active running Ceph object storage daemon
ceph-osd@5.service                loaded active running Ceph object storage daemon
ceph-osd@8.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
```

步驟5.禁用並停止每個ceph例項，從osd中刪除每個例項並解除安裝目錄。對每個ceph例項重複上述操作。

```
[root@pod1-osd-compute-0 ~]# systemctl disable ceph-osd@11

[root@pod1-osd-compute-0 ~]# systemctl stop ceph-osd@11

[root@pod1-osd-compute-0 ~]# ceph osd out 11
marked out osd.11.

[root@pod1-osd-compute-0 ~]# ceph osd crush remove osd.11
removed item id 11 name 'osd.11' from crush map

[root@pod1-osd-compute-0 ~]# ceph auth del osd.11
updated

[root@pod1-osd-compute-0 ~]# ceph osd rm 11
removed osd.11

[root@pod1-osd-compute-0 ~]# umount /var/lib/ceph/osd/ceph-11

[root@pod1-osd-compute-0 ~]# rm -rf /var/lib/ceph/osd/ceph-11
```

(或)

步驟6. Clean.sh腳本可用於同時執行上述任務。

```
[heat-admin@pod1-osd-compute-0 ~]$ sudo ls /var/lib/ceph/osd
ceph-11 ceph-3 ceph-6 ceph-8
[heat-admin@pod1-osd-compute-0 ~]$ /bin/sh clean.sh

[heat-admin@pod1-osd-compute-0 ~]$ cat clean.sh
#!/bin/sh
set -x
CEPH=`sudo ls /var/lib/ceph/osd`
for c in $CEPH
do
  i=`echo $c |cut -d'-' -f2`
```

```
  sudo systemctl disable ceph-osd@$i || (echo "error rc:$?"; exit 1)
  sleep 2
  sudo systemctl stop ceph-osd@$i || (echo "error rc:$?"; exit 1)
  sleep 2
  sudo ceph osd out $i || (echo "error rc:$?"; exit 1)
  sleep 2
  sudo ceph osd crush remove osd.$i || (echo "error rc:$?"; exit 1)
  sleep 2
  sudo ceph auth del osd.$i || (echo "error rc:$?"; exit 1)
  sleep 2
  sudo ceph osd rm $i || (echo "error rc:$?"; exit 1)
  sleep 2
  sudo umount /var/lib/ceph/osd/$c || (echo "error rc:$?"; exit 1)
  sleep 2
  sudo rm -rf /var/lib/ceph/osd/$c || (echo "error rc:$?"; exit 1)
  sleep 2
done

sudo ceph osd tree
```

在所有OSD進程都進行了遷移/刪除之後，節點可以從超雲中刪除。

附註：刪除CEPH後，VNF HD RAID進入「降級」狀態，但hd-disk仍然必須可供訪問。

正常斷電

將ESC遷移到備用模式

步驟1.登入到計算節點中託管的ESC並檢查它是否處於主狀態。如果是，將ESC切換到備用模式。

```
[admin@esc esc-cli]$ escadm status
0 ESC status=0 ESC Master Healthy


[admin@esc ~]$ sudo service keepalived stop
Stopping keepalived:                                       [  OK  ]

[admin@esc ~]$ escadm status
1 ESC status=0 In SWITCHING_TO_STOP state. Please check status after a while.

[admin@esc ~]$ sudo reboot
```

```
Broadcast message from admin@vnf1-esc-esc-0.novalocal
      (/dev/pts/0) at 13:32 ...
The system is going down for reboot NOW!
```

步驟2.從Nova聚合清單中刪除Osd-Compute節點。

- 列出nova聚合，並根據計算伺服器所承載的VNF確定對應於計算伺服器的聚合。通常，其格式為<VNFNAME>-EM-MGMT<X>和<VNFNAME>-CF-MGMT<X>

```
[stack@director ~]$ nova aggregate-list
+----+------+-------------------+
| Id | Name | Availability Zone |
+----+------+-------------------+
| 3 | esc1 | AZ-esc1 |
| 6 | esc2 | AZ-esc2 |
| 9 | aaa | AZ-aaa |
+----+------+-------------------+
```

在本例中，osd-compute伺服器屬於esc1。因此，相應的聚合將是**esc1**

步驟3.從標識的聚合中刪除osd-compute節點。

```
nova aggregate-remove-host
```

```
[stack@director ~]$ nova aggregate-remove-host esc1 pod1-osd-compute-0.localdomain
```

步驟4.驗證是否已從聚合中刪除osd-compute節點。現在，請確保該主機未列在聚合下。

```
nova aggregate-show
```

```
[stack@director ~]$ nova aggregate-show esc1
[stack@director ~]$
```

Osd計算節點刪除

不論計算節點中託管的VM，本節中提到的步驟都是通用的。

從超雲中刪除

步驟1.建立名為delete_node.sh的指令碼檔案，其內容如圖所示。請確保提到的模板與用於堆疊部署的**deploy.sh**指令碼中使用的模板相同。

```
 delete_node.sh

openstack overcloud node delete --templates -e /usr/share/openstack-tripleo-heat-
templates/environments/puppet-pacemaker.yaml -e /usr/share/openstack-tripleo-heat-
```

```
templates/environments/network-isolation.yaml -e /usr/share/openstack-tripleo-heat-
templates/environments/storage-environment.yaml -e /usr/share/openstack-tripleo-heat-
templates/environments/neutron-sriov.yaml -e /home/stack/custom-templates/network.yaml -e
/home/stack/custom-templates/ceph.yaml -e /home/stack/custom-templates/compute.yaml -e
/home/stack/custom-templates/layout.yaml -e /home/stack/custom-templates/layout.yaml --stack
```

```
[stack@director ~]$ source stackrc
[stack@director ~]$ /bin/sh delete_node.sh
+ openstack overcloud node delete --templates -e /usr/share/openstack-tripleo-heat-
templates/environments/puppet-pacemaker.yaml -e /usr/share/openstack-tripleo-heat-
templates/environments/network-isolation.yaml -e /usr/share/openstack-tripleo-heat-
templates/environments/storage-environment.yaml -e /usr/share/openstack-tripleo-heat-
templates/environments/neutron-sriov.yaml -e /home/stack/custom-templates/network.yaml -e
/home/stack/custom-templates/ceph.yaml -e /home/stack/custom-templates/compute.yaml -e
/home/stack/custom-templates/layout.yaml -e /home/stack/custom-templates/layout.yaml --stack
pod1 49ac5f22-469e-4b84-badc-031083db0533
Deleting the following nodes from stack pod1:
- 49ac5f22-469e-4b84-badc-031083db0533
Started Mistral Workflow. Execution ID: 4ab4508a-c1d5-4e48-9b95-ad9a5baa20ae

real   0m52.078s
user   0m0.383s
sys    0m0.086s
```

步驟2.等待OpenStack堆疊操作變為COMPLETE狀態。

```
[stack@director ~]$  openstack stack list
+--------------------------------------+------------+-----------------+----------------------+--
--------------------+
| ID                                   | Stack Name | Stack Status    | Creation Time        |
Updated Time         |
+--------------------------------------+------------+-----------------+----------------------+--
--------------------+
| 5df68458-095d-43bd-a8c4-033e68ba79a0 | pod1 | UPDATE_COMPLETE | 2018-05-08T21:30:06Z | 2018-
05-08T20:42:48Z |
+--------------------------------------+------------+-----------------+----------------------+--
--------------------
```

從服務清單中刪除Osd-Compute節點

從服務清單中刪除計算服務。

```
[stack@director ~]$ source corerc
[stack@director ~]$ openstack compute service list | grep osd-compute-0
| 404 | nova-compute     | pod1-osd-compute-0.localdomain     | nova     | enabled | up    |
2018-05-08T18:40:56.000000 |
```

```
openstack compute service delete
```

```
[stack@director ~]$ openstack compute service delete 404
```

刪除中子代理

刪除計算伺服器的舊關聯中子代理和open vswitch代理。

```
[stack@director ~]$ openstack network agent list | grep osd-compute-0
| c3ee92ba-aa23-480c-ac81-d3d8d01dcc03 | Open vSwitch agent | pod1-osd-compute-0.localdomain
| None              | False  | UP    | neutron-openvswitch-agent |
| ec19cb01-abbb-4773-8397-8739d9b0a349 | NIC Switch agent   | pod1-osd-compute-0.localdomain
| None              | False  | UP    | neutron-sriov-nic-agent   |
```

openstack network agent delete

```
[stack@director ~]$ openstack network agent delete c3ee92ba-aa23-480c-ac81-d3d8d01dcc03
[stack@director ~]$ openstack network agent delete ec19cb01-abbb-4773-8397-8739d9b0a349
```

從Nova和Ironic資料庫中刪除

從nova清單中刪除一個節點以及諷刺的資料庫，然後對其進行驗證。

```
[stack@director ~]$ source stackrc

[stack@al01-pod1-ospd ~]$ nova list | grep osd-compute-0
| c2cfa4d6-9c88-4ba0-9970-857d1a18d02c | pod1-osd-compute-0 | ACTIVE | -          | Running
| ctlplane=192.200.0.114 |

[stack@al01-pod1-ospd ~]$ nova delete c2cfa4d6-9c88-4ba0-9970-857d1a18d02c
```

nova show

```
[stack@director ~]$ nova show pod1-osd-compute-0 | grep hypervisor
| OS-EXT-SRV-ATTR:hypervisor_hostname  | 4ab21917-32fa-43a6-9260-02538b5c7a5a
```

ironic node-delete

```
[stack@director ~]$ ironic node-delete 4ab21917-32fa-43a6-9260-02538b5c7a5a
[stack@director ~]$ ironic node-list (node delete must not be listed now)
```

安裝新的計算節點

有關安裝新UCS C240 M4伺服器的步驟和初始設定步驟，請參閱：Cisco UCS C240 M4伺服器安裝和服務指南

步驟1.安裝伺服器後，將硬碟插入相應插槽中作為舊伺服器。

步驟2.使用CIMC IP登入到伺服器。

步驟3.如果韌體與先前使用的推薦版本不一致，請執行BIOS升級。此處提供了BIOS升級步驟：Cisco UCS C系列機架式伺服器BIOS升級指南

步驟4.檢驗物理驅動器的狀態。它一定是無**限的好東西**。

步驟5.使用RAID級別1從物理驅動器建立虛擬驅動器。

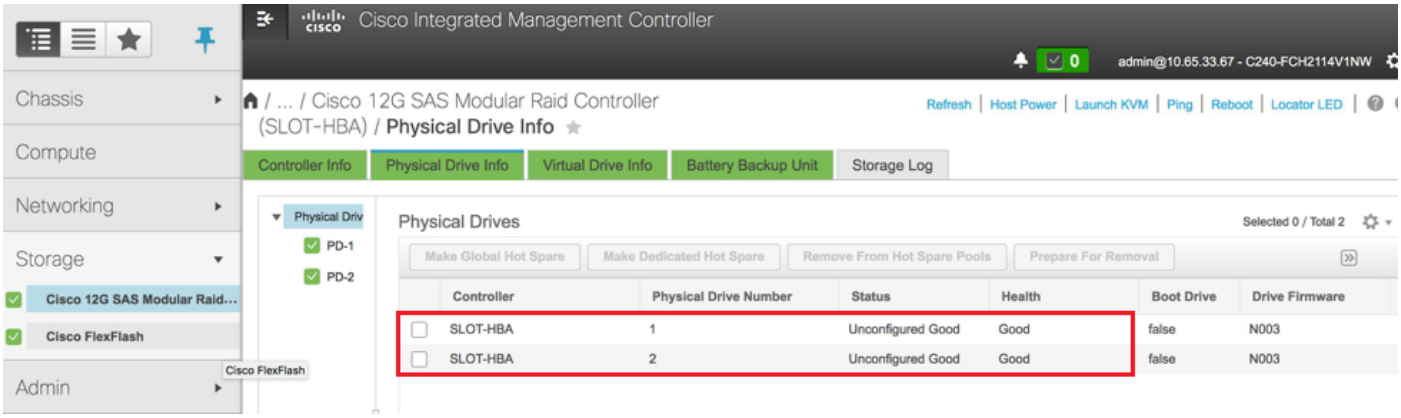

步驟6.導航到儲存部分並選擇Cisco 12G Sas模組化Raid控制器，然後驗證raid控制器的狀態和運行狀況，如下圖所示。

> **附註**：上述影象僅供說明之用，在實際OSD計算CIMC中，您會看到插槽[1,2,3,7,8,9,10]中的七個物理驅動器處於未配置的良好狀態，因為沒有從它們建立虛擬驅動器。

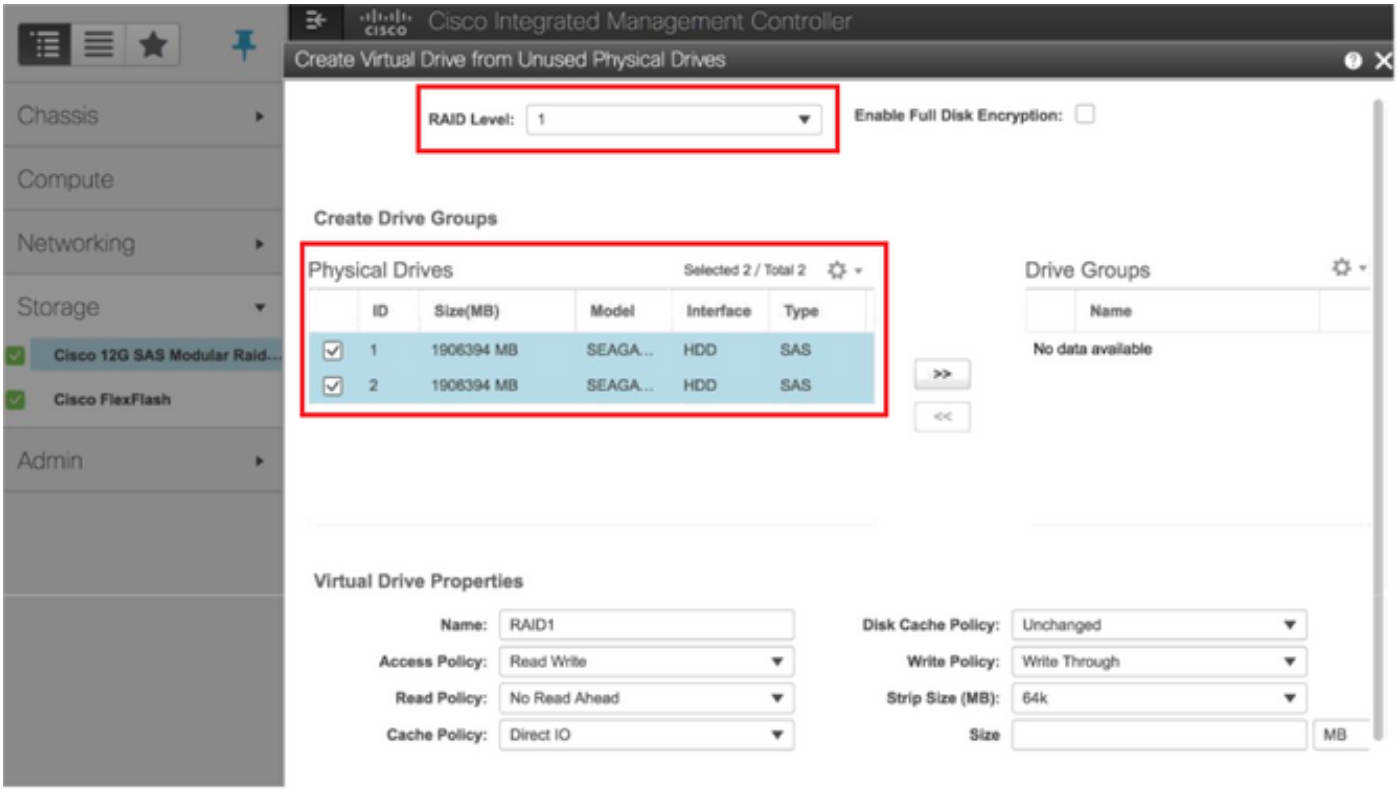

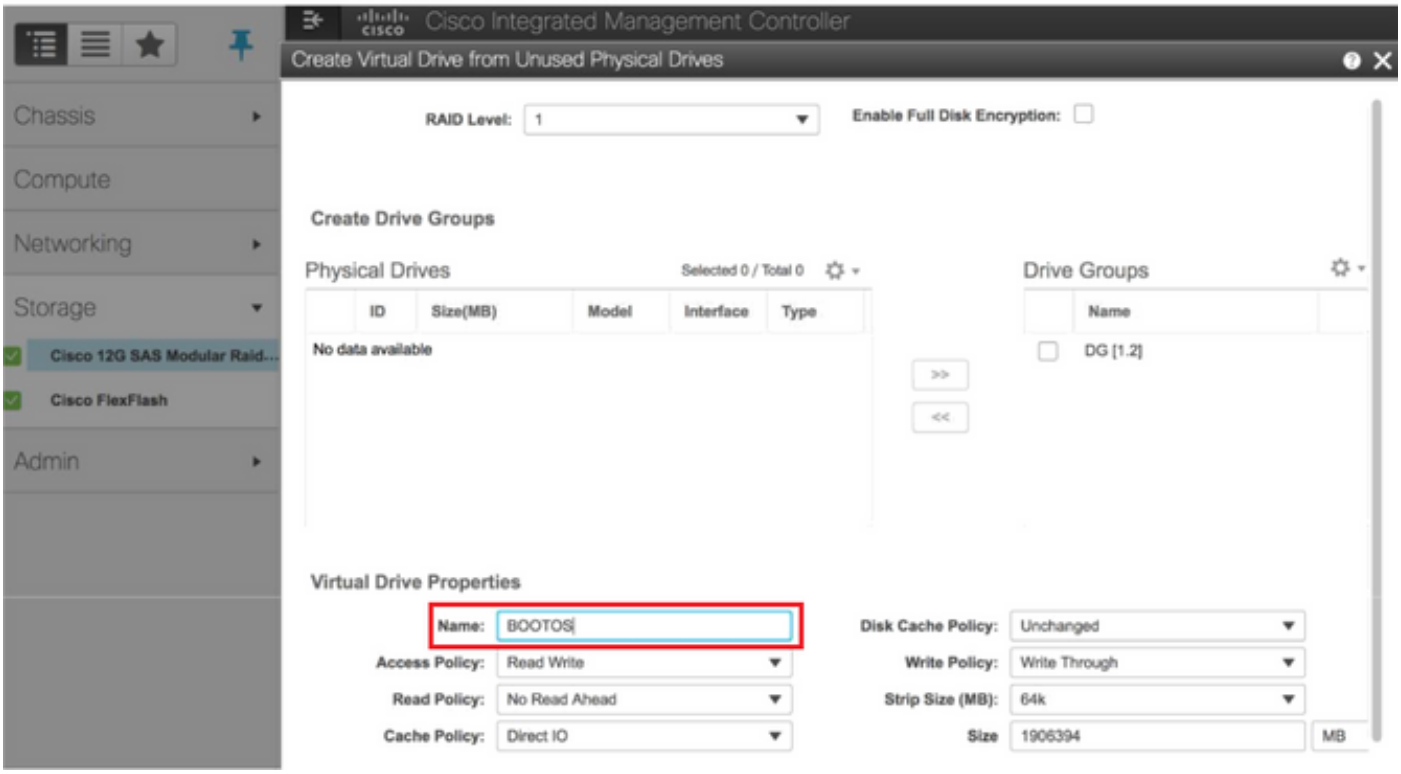

步驟7.現在從**Cisco 12G SAS模組化Raid控制器**下的控制器資訊中未使用的物理驅動器建立虛擬驅動器。

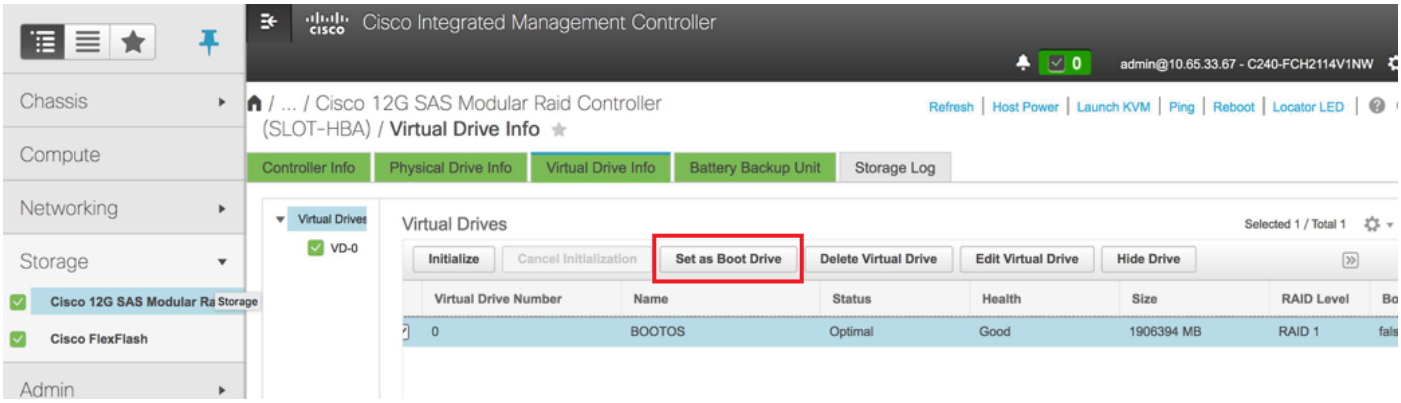

步驟8.選擇VD並配置設定為引導驅動器。

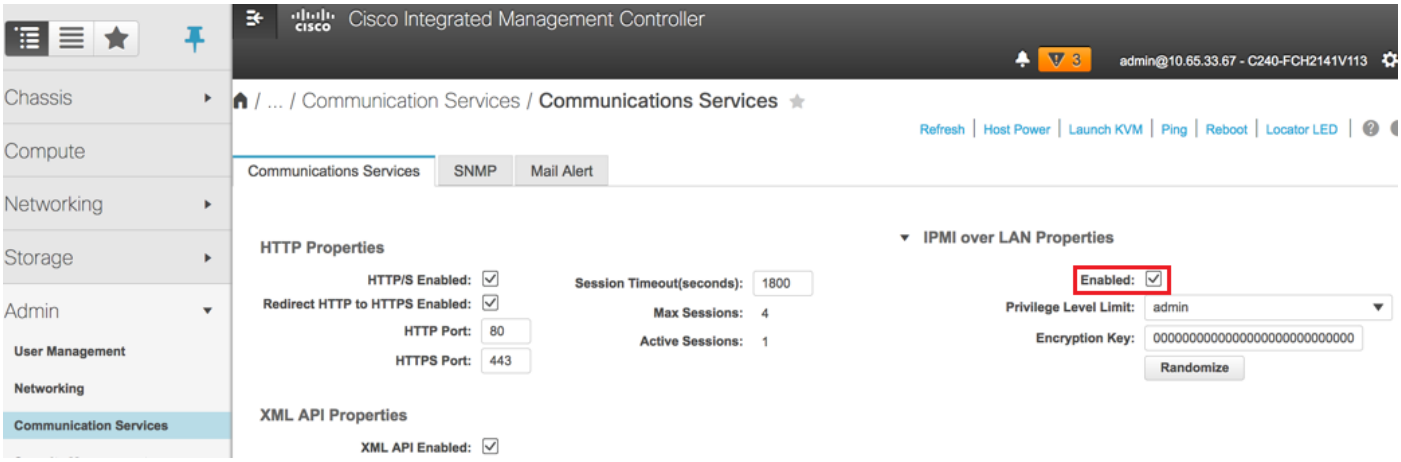

步驟9.從Admin頁籤下的Communication services啟用IPMI over LAN。

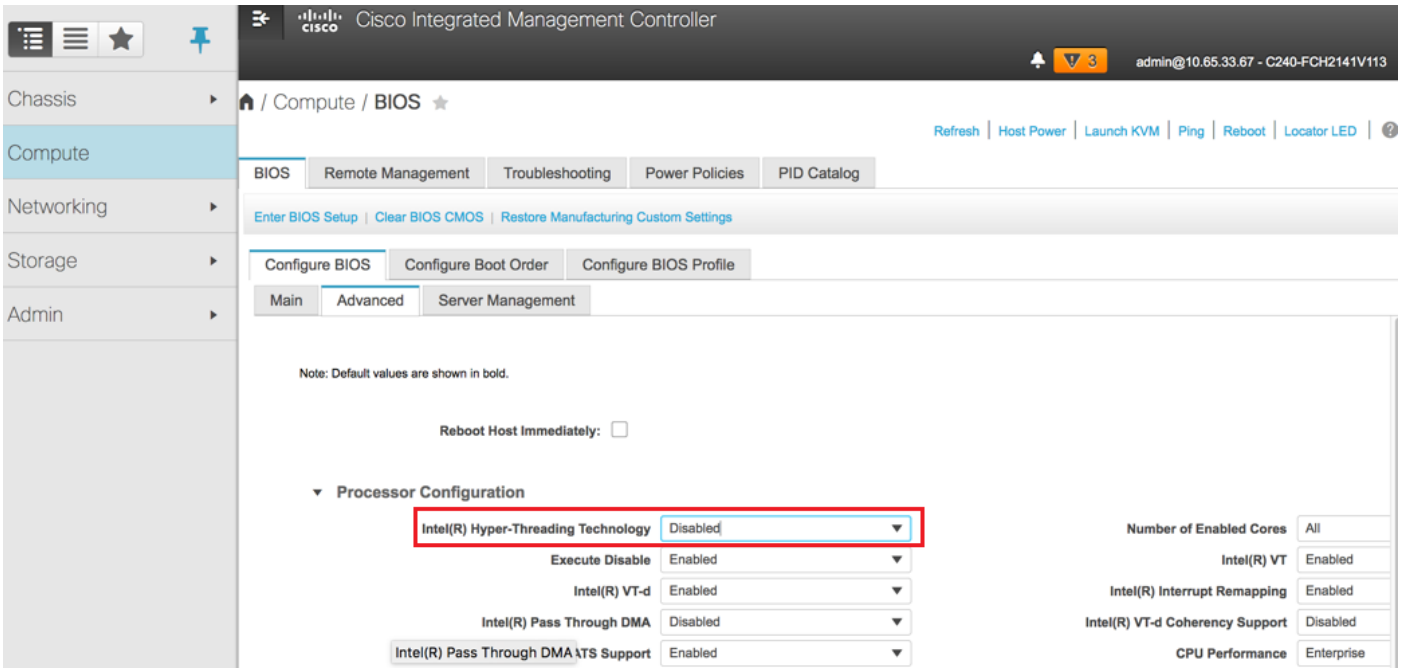

步驟10.在Compute節點下的Advanced BIOS配置中禁用超執行緒，如下圖所示。

步驟11.與使用物理驅動器1和2建立的BOOTOS VD類似，另外建立四個虛擬驅動器，作為

JOURNAL — 從物理驅動器號3

OSD1 — 從物理驅動器號7

OSD2 — 從物理驅動器號8

OSD3 — 從物理驅動器號9

OSD4 — 從物理驅動器號10

步驟7.最後，物理驅動器和虛擬驅動器必須相似。

> **附註**：此處顯示的影象和本節中提到的配置步驟是參考韌體版本3.0(3e)，如果您使用其他版本，可能會有細微的變化。

將新的OSD-Compute節點新增到Overcloud

不論計算節點託管的VM，本節中提到的步驟都是通用的。

步驟1.新增具有不同索引的計算伺服器。

建立一個**add_node.json**檔案，該檔案僅包含要新增的新計算伺服器的詳細資訊。 確保新的osd-compute伺服器的索引號之前未使用過。通常，遞增下一個最高計算值。

範例：最高驗前是osd-compute-0，因此在2-vnf系統的情況下建立了osd-compute-3。

> **附註**：請記住json格式。

```
[stack@director ~]$ cat add_node.json
{
   "nodes":[
       {
           "mac":[
               "<MAC_ADDRESS>"
           ],
           "capabilities": "node:osd-compute-3,boot_option:local",
           "cpu":"24",
           "memory":"256000",
           "disk":"3000",
           "arch":"x86_64",
           "pm_type":"pxe_ipmitool",
           "pm_user":"admin",
           "pm_password":"<PASSWORD>",
           "pm_addr":"192.100.0.5"
       }
   ]
}
```

步驟2.匯入json檔案。

```
[stack@director ~]$ openstack baremetal import --json add_node.json
Started Mistral Workflow. Execution ID: 78f3b22c-5c11-4d08-a00f-8553b09f497d
Successfully registered node UUID 7eddfa87-6ae6-4308-b1d2-78c98689a56e
Started Mistral Workflow. Execution ID: 33a68c16-c6fd-4f2a-9df9-926545f2127e
```

```
Successfully set all nodes to available.
```

步驟3.使用上一步中提到的UUID運行節點內檢。

```
[stack@director ~]$ openstack baremetal node manage 7eddfa87-6ae6-4308-b1d2-78c98689a56e
[stack@director ~]$ ironic node-list |grep 7eddfa87
| 7eddfa87-6ae6-4308-b1d2-78c98689a56e | None | None                                    | power off
  | manageable         | False       |

[stack@director ~]$ openstack overcloud node introspect 7eddfa87-6ae6-4308-b1d2-78c98689a56e --
provide
Started Mistral Workflow. Execution ID: e320298a-6562-42e3-8ba6-5ce6d8524e5c
Waiting for introspection to finish...
Successfully introspected all nodes.
Introspection completed.
Started Mistral Workflow. Execution ID: c4a90d7b-ebf2-4fcb-96bf-e3168aa69dc9
Successfully set all nodes to available.

[stack@director ~]$ ironic node-list |grep available
| 7eddfa87-6ae6-4308-b1d2-78c98689a56e | None | None                                    | power off
  | available          | False       |
```

步驟4.將IP地址新增到custom-templates/layout.yml的OsdComputeIP下。在這種情況下，當您替換osd-compute-0時，會將該地址新增到每個型別的清單末尾。

```
OsdComputeIPs:

   internal_api:

   - 11.120.0.43

   - 11.120.0.44

   - 11.120.0.45

   - 11.120.0.43   <<< take osd-compute-0 .43 and add here

   tenant:

   - 11.117.0.43

   - 11.117.0.44

   - 11.117.0.45

   - 11.117.0.43   << and here

   storage:

   - 11.118.0.43

   - 11.118.0.44

   - 11.118.0.45

   - 11.118.0.43   << and here

   storage_mgmt:
```

```
   - 11.119.0.43

   - 11.119.0.44

   - 11.119.0.45

   - 11.119.0.43   << and here
```

步驟5.運行以前用於部署堆疊的deploy.sh指令碼，以便將新的計算節點新增到超雲堆疊。

```
[stack@director ~]$ ./deploy.sh
++ openstack overcloud deploy --templates -r /home/stack/custom-templates/custom-roles.yaml -e
/usr/share/openstack-tripleo-heat-templates/environments/puppet-pacemaker.yaml -e
/usr/share/openstack-tripleo-heat-templates/environments/network-isolation.yaml -e
/usr/share/openstack-tripleo-heat-templates/environments/storage-environment.yaml -e
/usr/share/openstack-tripleo-heat-templates/environments/neutron-sriov.yaml -e
/home/stack/custom-templates/network.yaml -e /home/stack/custom-templates/ceph.yaml -e
/home/stack/custom-templates/compute.yaml -e /home/stack/custom-templates/layout.yaml --stack
ADN-ultram --debug --log-file overcloudDeploy_11_06_17__16_39_26.log --ntp-server 172.24.167.109
--neutron-flat-networks phys_pcie1_0,phys_pcie1_1,phys_pcie4_0,phys_pcie4_1 --neutron-network-
vlan-ranges datacentre:1001:1050 --neutron-disable-tunneling --verbose --timeout 180
…
Starting new HTTP connection (1): 192.200.0.1
"POST /v2/action_executions HTTP/1.1" 201 1695
HTTP POST http://192.200.0.1:8989/v2/action_executions 201
Overcloud Endpoint: http://10.1.2.5:5000/v2.0
Overcloud Deployed
clean_up DeployOvercloud:
END return value: 0

real    38m38.971s
user    0m3.605s
sys     0m0.466s
```

步驟6.等待openstack狀態變為COMPLETE。

```
[stack@director ~]$  openstack stack list
+--------------------------------------+------------+----------------+----------------------+--
--------------------+
| ID                                   | Stack Name | Stack Status    | Creation Time        |
Updated Time         |
+--------------------------------------+------------+----------------+----------------------+--
--------------------+
| 5df68458-095d-43bd-a8c4-033e68ba79a0 | pod1 | UPDATE_COMPLETE | 2017-11-02T21:30:06Z | 2017-
11-06T21:40:58Z |
+--------------------------------------+------------+----------------+----------------------+--
--------------------+
```

步驟7.檢查新的osd-compute節點是否處於活動狀態。

```
[stack@director ~]$ source stackrc
[stack@director ~]$ nova list |grep osd-compute-3
| 0f2d88cd-d2b9-4f28-b2ca-13e305ad49ea | pod1-osd-compute-3    | ACTIVE | -          | Running
  | ctlplane=192.200.0.117 |

[stack@director ~]$ source corerc
[stack@director ~]$ openstack hypervisor list |grep osd-compute-3
| 63 | pod1-osd-compute-3.localdomain    |
```

步驟8.登入新的osd-compute伺服器並檢查ceph進程。最初，狀態在HEALTH_WARN中作為ceph恢復。

```
[heat-admin@pod1-osd-compute-3 ~]$ sudo ceph -s

   cluster eb2bb192-b1c9-11e6-9205-525400330666

    health HEALTH_WARN

           223 pgs backfill_wait

           4 pgs backfilling

           41 pgs degraded

           227 pgs stuck unclean

           41 pgs undersized

           recovery 45229/1300136 objects degraded (3.479%)

           recovery 525016/1300136 objects misplaced (40.382%)

    monmap e1: 3 mons at {Pod1-controller-0=11.118.0.40:6789/0,Pod1-controller-
1=11.118.0.41:6789/0,Pod1-controller-2=11.118.0.42:6789/0}

           election epoch 58, quorum 0,1,2 Pod1-controller-0,Pod1-controller-1,Pod1-controller-2

    osdmap e986: 12 osds: 12 up, 12 in; 225 remapped pgs

           flags sortbitwise,require_jewel_osds

     pgmap v781746: 704 pgs, 6 pools, 533 GB data, 344 kobjects

           1553 GB used, 11840 GB / 13393 GB avail

           45229/1300136 objects degraded (3.479%)

           525016/1300136 objects misplaced (40.382%)

                477 active+clean

                186 active+remapped+wait_backfill

                 37 active+undersized+degraded+remapped+wait_backfill

                  4 active+undersized+degraded+remapped+backfilling
```

步驟9.但是，在很短時間（20分鐘）後，CEPH會返回到HEALTH_OK狀態。

```
[heat-admin@pod1-osd-compute-3 ~]$ sudo ceph -s

   cluster eb2bb192-b1c9-11e6-9205-525400330666

    health HEALTH_OK

    monmap e1: 3 mons at {Pod1-controller-0=11.118.0.40:6789/0,Pod1-controller-
1=11.118.0.41:6789/0,Pod1-controller-2=11.118.0.42:6789/0}

           election epoch 58, quorum 0,1,2 Pod1-controller-0,Pod1-controller-1,Pod1-controller-2
```

```
    osdmap e1398: 12 osds: 12 up, 12 in
           flags sortbitwise,require_jewel_osds
     pgmap v784311: 704 pgs, 6 pools, 533 GB data, 344 kobjects
           1599 GB used, 11793 GB / 13393 GB avail
                704 active+clean
 client io 8168 kB/s wr, 0 op/s rd, 32 op/s wr


[heat-admin@pod1-osd-compute-3 ~]$ sudo ceph osd tree
ID WEIGHT   TYPE NAME                          UP/DOWN REWEIGHT PRIMARY-AFFINITY
-1 13.07996 root default
-2        0     host pod1-osd-compute-0
-3  4.35999     host pod1-osd-compute-2
 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 pod1-osd-compute-1
 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
-5  4.35999     host pod1-osd-compute-3
 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
```

恢復虛擬機器

新星聚合清單的新增內容

將osd-compute節點新增到聚合主機並驗證是否新增了主機。

nova aggregate-add-host

```
[stack@director ~]$ nova aggregate-add-host esc1 pod1-osd-compute-3.localdomain
```

nova aggregate-show

```
[stack@director ~]$ nova aggregate-show esc1
+----+------+-------------------+---------------------------------------+--------------------
--------------------+
| Id | Name | Availability Zone | Hosts | Metadata |
+----+------+-------------------+---------------------------------------+--------------------
--------------------+
| 3 | esc1 | AZ-esc1 | 'pod1-osd-compute-3.localdomain' | 'availability_zone=AZ-esc1',
'esc1=true' |
+----+------+-------------------+---------------------------------------+--------------------
--------------------+
```

恢復ESC虛擬機器

步驟1.從新星清單中檢查ESC VM的狀態並將其刪除。

```
stack@director scripts]$ nova list |grep esc
| c566efbf-1274-4588-a2d8-0682e17b0d41 | esc                                                    |
ACTIVE | -          | Running     | VNF2-UAS-uas-orchestration=172.168.11.14; VNF2-UAS-uas-
management=172.168.10.4
                                |
[stack@director scripts]$ nova delete esc
Request to delete server esc has been accepted.

If can not delete esc then use command: nova force-delete esc
```

步驟2.在OSPD中，導航到ECS-Image目錄，並確儲存在ESC版本的bootvm.py和qcow2（如果未將其移動到目錄）。

```
[stack@atospd ESC-Image-157]$ ll

total 30720136

-rw-r--r--. 1 root  root       127724 Jan 23 12:51 bootvm-2_3_2_157a.py

-rw-r--r--. 1 root  root           55 Jan 23 13:00 bootvm-2_3_2_157a.py.md5sum

-rw-rw-r--. 1 stack stack 31457280000 Jan 24 11:35 esc-2.3.2.157.qcow2
```

步驟3.建立映像。

```
[stack@director ESC-image-157]$  glance image-create --name ESC-2_3_2_157 --disk-format "qcow2"
--container "bare" --file /home/stack/ECS-Image-157/ESC-2_3_2_157.qcow2
```

步驟4.驗證ESC影象是否存在。

```
stack@director ~]$ glance image-list
+--------------------------------------+--------------------------------------+
| ID                                   | Name                                 |
+--------------------------------------+--------------------------------------+
| 8f50acbe-b391-4433-aa21-98ac36011533 | ESC-2_3_2_157|
| 2f67f8e0-5473-467c-832b-e07760e8d1fa | tmobile-pcrf-13.1.1.iso              |
| c5485c30-45db-43df-831d-61046c5cfd01 | tmobile-pcrf-13.1.1.qcow2            |
| 2f84b9ec-61fa-46a3-a4e6-45f14c93d9a9 | tmobile-pcrf-13.1.1_cco_20170825.iso |
| 25113ecf-8e63-4b81-a73f-63606781ef94 | wscaaa01-sept072017                  |
| 595673e8-c99c-40c2-82b1-7338325024a9 | wscaaa02-sept072017                  |
| 8bce3a60-b3b0-4386-9e9d-d99590dc9033 | wscaaa03-sept072017                  |
| e5c835ad-654b-45b0-8d36-557e6c5fd6e9 | wscaaa04-sept072017                  |
| 879dfcde-d25c-4314-8da0-32e4e73ffc9f | WSP1_cluman_12_07_2017               |
| 7747dd59-c479-4c8a-9136-c90ec894569a | WSP2_cluman_12_07_2017               |
+--------------------------------------+--------------------------------------+

[stack@ ~]$ openstack flavor list
+--------------------------------------+------------+--------+------+-----------+-------+-------
----+
| ID                                   | Name       |    RAM | Disk | Ephemeral | VCPUs | Is
Public |
+--------------------------------------+------------+--------+------+-----------+-------+-------
----+
| 1e4596d5-46f0-46ba-9534-cfdea788f734 | pcrf-smb   | 100352 |  100 |         0 |     8 | True
|
| 251225f3-64c9-4b19-a2fc-032a72bfe969 | pcrf-oam   |  65536 |  100 |         0 |    10 | True
|
| 4215d4c3-5b2a-419e-b69e-7941e2abe3bc | pcrf-pd    |  16384 |  100 |         0 |    12 | True
|
| 4c64a80a-4d19-4d52-b818-e904a13156ca | pcrf-qns   |  14336 |  100 |         0 |    10 | True
|
| 8b4cbba7-40fd-49b9-ab21-93818c80a2e6 | esc-flavor |   4096 |    0 |         0 |     4 | True
|
| 9c290b80-f80a-4850-b72f-d2d70d3d38ea | pcrf-sm    | 100352 |  100 |         0 |    10 | True
|
| e993fc2c-f3b2-4f4f-9cd9-3afc058b7ed1 | pcrf-arb   |  16384 |  100 |         0 |     4 | True
|
| f2b3b925-1bf8-4022-9f17-433d6d2c47b5 | pcrf-cm    |  14336 |  100 |         0 |     6 | True
|
+--------------------------------------+------------+--------+------+-----------+-------+-------
----+
```

步驟5.在影象目錄下建立此檔案並啟動ESC例項。

```
[root@director ESC-IMAGE]# cat esc_params.conf
openstack.endpoint = publicURL

[root@director ESC-IMAGE]./bootvm-2_3_2_157a.py esc --flavor esc-flavor --image ESC-2_3_2_157 --
net tb1-mgmt --gateway_ip 172.16.181.1 --net tb1-orch --enable-http-rest --avail_zone AZ-esc1 --
user_pass "admin:Cisco123" --user_confd_pass "admin:Cisco123" --bs_os_auth_url
http://10.250.246.137:5000/v2.0 --kad_vif eth0 --kad_vip 172.16.181.5 --ipaddr 172.16.181.4 dhcp
--ha_node_list 172.16.181.3 172.16.181.4 --esc_params_file esc_params.conf
```

bootvm.py ESC VMESC HAESC

步驟6.登入新的ESC並驗證備份狀態。

```
[admin@esc ~]$ escadm status
0 ESC status=0 ESC Backup Healthy

[admin@VNF2-esc-esc-1 ~]$ health.sh
============== ESC HA (BACKUP) ================================================
ESC HEALTH PASSED
```