

# Ultra-M UCS 240M4单硬盘故障 — 热插拔程序 — CPS

## 目录

[简介](#)

[背景信息](#)

[缩写](#)

[MoP的工作流](#)

[单硬盘故障](#)

[计算服务器上的单硬盘故障](#)

[控制器服务器上的单硬盘故障](#)

[OSD-Compute服务器上的单硬盘故障](#)

[OSPD服务器上的单硬盘故障](#)

## 简介


本文档介绍在托管思科策略套件(CPS)虚拟网络功能(VNF)的Ultra-M设置中更换服务器中故障硬盘所需的步骤。

## 背景信息

Ultra-M是预打包和验证的虚拟化移动数据包核心解决方案，旨在简化VNF的部署。OpenStack是Ultra-M的虚拟化基础设施管理器(VIM)，由以下节点类型组成：

- 计算
- 对象存储磁盘 — 计算 ( OSD — 计算 )
- 控制器
- OpenStack平台 — 导向器(OSPD)

Ultra-M的高级体系结构和涉及的组件如下图所示：




UltraM架构

**注意：**考虑Ultra M 5.1.x版本以定义本文档中的步骤。本文档面向熟悉Cisco Ultra-M平台的思科人员，并详细介绍在更换OSPD服务器时在OpenStack级别执行所需的步骤。

## 缩写

VNF	虚拟网络功能
ESC	弹性服务控制器
MOP	程序方法
OSD	对象存储磁盘
硬盘	硬盘驱动器
SSD	固态硬盘
VIM	虚拟基础设施管理器
虚拟机	虚拟机
EM	元素管理器
UAS	超自动化服务
UUID	通用唯一Identifier

## MoP的工作流



## 单硬盘故障

1. 每台裸机服务器将配置两个HDD驱动器，以在Raid 1配置中充当引导磁盘。如果单个HDD发生故障，由于存在Raid 1级冗余，故障HDD驱动器可以热插拔。
2. 请参阅以下步骤以更换UCS C240 M4服务器上的故障组件：[更换服务器组件](#)

3.如果单个HDD出现故障，则仅热交换故障HDD，因此更换新磁盘后无需BIOS升级过程。

4.更换磁盘后，请等待磁盘之间的数据同步。可能需要几个小时才能完成。

5.在基于OpenStack(Ultra-M)的解决方案中，UCS 240M4裸机服务器可承担以下角色之一：计算、OSD — 计算、控制器和OSPD。

6.在这些服务器角色中处理单个HDD故障所需的步骤相同，本节介绍在磁盘热插拔之前要执行的运行状况检查。

## 计算服务器上的单硬盘故障

1.如果UCS 240M4（充当计算节点）中发现HDD驱动器故障，请在启动故障磁盘的热插拔过程之前执行这些运行状况检查。

2.确定此服务器上运行的VM，并验证功能状态是否正常。

### 确定托管在计算节点中的虚拟机

确定托管在计算服务器上的VM，并验证它们是否处于活动状态并正在运行。

计算服务器包含CPS VM/弹性服务控制器(ESC)虚拟机组合：

```
[stack@director ~]$ nova list --field name,host | grep compute-8
| 507d67c2-1d00-4321-b9d1-da879af524f8 | VNF2-DEPLOYM_XXXX_0_c8d98f0f-d874-45d0-af75-
88a2d6fa82ea | pod1-compute-8.localdomain | ACTIVE |
| f9c0763a-4a4f-4bbd-af51-bc7545774be2 | VNF2-DEPLOYM_c2_0_df4be88d-b4bf-4456-945a-
3812653ee229 | pod1-compute-8.localdomain | ACTIVE |
| 75528898-ef4b-4d68-b05d-882014708694 | VNF2-ESC-ESC-
0 | pod1-compute-8.localdomain | ACTIVE |
```

**注意：**在此处显示的输出中，第一列对应于通用唯一IDentifier(UUID)，第二列是VM名称，第三列是VM所在的主机名。

## 运行状况检查

1.登录计算节点中托管的ESC并检查状态。

```
[admin@VNF2-esc-esc-0 esc-cli]$ escadm status
0 ESC status=0 ESC Master Healthy
```

2.登录计算节点中托管的UAS并检查状态。

```
ubuntu@autovnf2-uas-1:~$ sudo su
root@autovnf2-uas-1:/home/ubuntu# confd_cli -u admin -C
Welcome to the ConfD CLI
admin connected from 127.0.0.1 using console on autovnf2-uas-1
autovnf2-uas-1#show uas ha
uas ha-vip 172.18.181.101
autovnf2-uas-1#
autovnf2-uas-1#
```

```

autovnf2-uas-1#show uas
uas version 1.0.1-1
uas state ha-active
uas ha-vip 172.18.181.101
INSTANCE IP    STATE  ROLE
-----
172.18.180.4  alive  CONFD-SLAVE
172.18.180.5  alive  CONFD-MASTER
172.18.180.8  alive  NA

```

```

autovnf2-uas-1#show errors
% No entries found.

```

3.如果运行状况检查正常，请继续执行故障磁盘热插拔过程，并等待数据同步，因为可能需要几个小时才能完成。请参阅：[更换服务器组件](#)

4.重复这些运行状况检查步骤，以确认托管在计算节点上的虚拟机的运行状况已恢复。

## 控制器服务器上的单硬盘故障

1.如果UCS 240M4（充当控制器节点）中发现HDD驱动器故障，请在启动故障磁盘的热插拔过程之前执行这些运行状况检查。

2.检查控制器上的Pace maker状态。

3.登录其中一个活动控制器并检查Pacemaker状态。所有服务必须在可用控制器上运行，并在故障控制器上停止。

```

[heat-admin@pod1-controller-0 ~]$ sudo pcs status
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod1-controller-0 (version 1.1.15-11.e17_3.4-e174ec8) - partition with quorum
Last updated: Thu Jun 28 07:53:06 2018          Last change: Wed Jan 17 11:38:00 2018 by root
via cibadmin on pod1-controller-0

3 nodes and 22 resources conimaged

Online: [ pod1-controller-0 pod1-controller-1 pod1-controller-2 ]

Full list of resources:

ip-10.2.2.2      (ocf::heartbeat:IPAddr2):      Started pod1-controller-0
ip-11.120.0.42 (ocf::heartbeat:IPAddr2):      Started pod1-controller-1
ip-11.119.0.42 (ocf::heartbeat:IPAddr2):      Started pod1-controller-2
ip-11.120.0.50 (ocf::heartbeat:IPAddr2):      Started pod1-controller-0
ip-11.118.0.48 (ocf::heartbeat:IPAddr2):      Started pod1-controller-1
ip-192.200.0.102 (ocf::heartbeat:IPAddr2):      Started pod1-controller-2
Clone Set: haproxy-clone [haproxy]
  Started: [ pod1-controller-0 pod1-controller-1 pod1-controller-2 ]
Master/Slave Set: galera-master [galera]
  Masters: [ pod1-controller-0 pod1-controller-1 pod1-controller-2 ]
Clone Set: rabbitmq-clone [rabbitmq]
  Started: [ pod1-controller-0 pod1-controller-1 pod1-controller-2 ]
Master/Slave Set: redis-master [redis]
  Masters: [ pod1-controller-0 ]
  Slaves: [ pod1-controller-1 pod1-controller-2 ]
openstack-cinder-volume (systemd:openstack-cinder-volume):      Started pod1-controller-0
my-ipmilan-for-controller-0 (stonith:fence_ipmilan):      Started pod1-controller-1
my-ipmilan-for-controller-1 (stonith:fence_ipmilan):      Started pod1-controller-2

```

```
my-ipmilan-for-controller-2 (stonith:fence_ipmilan): Started pod1-controller-0
```

Daemon Status:

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

#### 4.检查活动控制器中的MariaDB状态。

```
[stack@director] nova list | grep control
| 4361358a-922f-49b5-89d4-247a50722f6d | pod1-controller-0 | ACTIVE | - | Running |
ctlplane=192.200.0.102 |
| d0f57f27-93a8-414f-b4d8-957de0d785fc | pod1-controller-1 | ACTIVE | - | Running |
ctlplane=192.200.0.110 |
```

```
[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
*** 192.200.0.152 ***
```

```
Variable_name      Value
wsrep_local_state_comment Synced
Variable_name      Value
wsrep_cluster_size      2
```

```
*** 192.200.0.154 ***
Variable_name      Value
wsrep_local_state_comment Synced
Variable_name      Value
wsrep_cluster_size      2
```

验证每个活动控制器都存在以下线路：

```
wsrep_local_state_comment: Synced
```

```
wsrep_cluster_size: 2
```

#### 5.检查活动控制器中的Rabbitmq状态。

```
[heat-admin@pod1-controller-0 ~]$ sudo rabbitmqctl cluster_status
Cluster status of node 'rabbit@pod1-controller-0' ...
[{nodes,[{disc,['rabbit@pod1-controller-0','rabbit@pod1-controller-1',
'rabbit@pod1-controller-2']}]},
{running_nodes,['rabbit@pod1-controller-2',
'rabbit@pod1-controller-1',
'rabbit@pod1-controller-0']},
{cluster_name,<<"rabbit@pod1-controller-0.localdomain">>},
{partitions,[]},
{alarms,[{'rabbit@pod1-controller-2',[]},
{'rabbit@pod1-controller-1',[]},
{'rabbit@pod1-controller-0',[]}]}
```

6.如果运行状况检查正常，请继续执行故障磁盘热插拔过程，并等待数据同步，因为可能需要几个小时才能完成。请参阅：[更换服务器组件](#)

7.重复这些运行状况检查步骤，以确认控制器上的运行状况已恢复。

## OSD-Compute服务器上的单硬盘故障

如果UCS 240M4中发现HDD驱动器故障，它充当OSD计算节点，请在启动故障磁盘的热插拔过程之前执行这些运行状况检查。

## 识别OSD计算节点中托管的虚拟机

### 1. 计算服务器包含ESC VM。

```
[stack@director ~]$ nova list --field name,host | grep osd-compute-1
| 507d67c2-1d00-4321-b9d1-da879af524f8 | VNF2-DEPLOYM_XXXX_0_c8d98f0f-d874-45d0-af75-
88a2d6fa82ea | pod1-compute-8.localdomain | ACTIVE |
| f9c0763a-4a4f-4bbd-af51-bc7545774be2 | VNF2-DEPLOYM_c1_0_df4be88d-b4bf-4456-945a-
3812653ee229 | pod1-compute-8.localdomain | ACTIVE |
| 75528898-ef4b-4d68-b05d-882014708694 | VNF2-ESC-ESC-
0 | pod1-compute-8.localdomain | ACTIVE |
| f5bd7b9c-476a-4679-83e5-303f0aae9309 | VNF2-UAS-uas-
0 | pod1-compute-8.localdomain | ACTIVE |
```

**注意：**在此处显示的输出中，第一列与(UUID)对应，第二列是VM名称，第三列是VM所在的主机名。

### 2. OSD-Compute服务器上的Ceph进程处于活动状态。

```
[root@pod1-osd-compute-1 ~]# 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
```

### 3.验证OSD ( 硬盘 ) 到日志(SSD)的映射是否正常。

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

### 4.验证Ceph运行状况和OSD树映射是否良好。

```
[heat-admin@pod1-osd-compute-3 ~]$ sudo ceph -s
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_OK
  1 mons down, quorum 0,1 pod1-controller-0,pod1-controller-1
monmap e1: 3 mons at {pod1-controller-0=11.118.0.10:6789/0,pod1-controller-1=11.118.0.11:6789/0,pod1-controller-2=11.118.0.12:6789/0}
  election epoch 28, quorum 0,1 pod1-controller-0,pod1-controller-1
osdmap e709: 12 osds: 12 up, 12 in
  flags sortbitwise,require_jewel_osds
pgmap v941813: 704 pgs, 6 pools, 490 GB data, 163 kobjects
  1470 GB used, 11922 GB / 13393 GB avail
    704 active+clean
client io 58580 B/s wr, 0 op/s rd, 7 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  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

-4  4.35999  host pod1-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

-5  4.35999  host pod1-osd-compute-3
  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
```

5.如果运行状况检查正常，请继续执行故障磁盘热插拔过程，并等待数据同步，因为可能需要几个小时才能完成。请参阅[更换服务器组件](#)

6.重复这些运行状况检查步骤，以确认托管在OSD-Compute节点上的虚拟机的运行状况已恢复。



## OSPD服务器上的单硬盘故障

1.如果UCS 240M4 ( 充当OSPD节点 ) 中观察到HDD驱动器故障 , 建议在启动故障磁盘的热插拔过程之前执行这些检查。

2.检查OpenStack堆栈和节点列表的状态。

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

3.检查OSPD节点中是否所有下云服务都处于加载、活动和运行状态。

```
[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-openvswitch-agent.service	loaded	active	running	OpenStack Neutron Open vSwitch Agent
neutron-ovs-cleanup.service	loaded	active	exited	OpenStack Neutron Open vSwitch Cleanup Utility
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.

37 loaded units listed. Pass --all to see loaded but inactive units, too.

To show all installed unit files use 'systemctl list-unit-files'.

4.如果运行状况检查正常，请继续执行故障磁盘热插拔过程，并等待数据同步，因为可能需要几个小时才能完成。请参阅[更换服务器组件](#)

5.重复这些运行状况检查步骤，以确认OSPD节点的运行状况已恢复。