

# Ultra-M UCS 240M4 Single HDD Failure - Hot Swap Procedure - CPS

## Contents

[Introduction](#)

[Background Information](#)

[Abbreviations](#)

[Workflow of the MoP](#)

[Single HDD Failure](#)

[Single HDD Failure on a Compute Server](#)

[Single HDD Failure on a Controller Server](#)

[Single HDD Failure on an OSD-Compute Server](#)

[Single HDD Failure on an OSPD Server](#)

## Introduction

This document describes the steps required in order to replace the faulty HDD drive in a server in an Ultra-M setup that hosts Cisco Policy Suite (CPS) Virtual Network Function (VNFs).

## Background Information

Ultra-M is a pre-packaged and validated virtualized mobile packet core solution designed to simplify the deployment of VNFs. OpenStack is the Virtualized Infrastructure Manager (VIM) for Ultra-M and consists of these node types:

- Compute
- Object Storage Disk - Compute (OSD - Compute)
- Controller
- OpenStack Platform - Director (OSPD)

The high-level architecture of Ultra-M and the components involved are as shown in this image:




UltraM Architecture

**Note:** Ultra M 5.1.x release is considered in order to define the procedures in this document. This document is intended for the Cisco personnel who are familiar with the Cisco Ultra-M platform and it details the steps required to be carried out at OpenStack level at the time of the OSPD Server replacement.

## Abbreviations

VNF	Virtual Network Function
ESC	Elastic Service Controller
MOP	Method of Procedure
OSD	Object Storage Disks
HDD	Hard Disk Drive
SSD	Solid State Drive
VIM	Virtual Infrastructure Manager
VM	Virtual Machine
EM	Element Manager
UAS	Ultra Automation Services
UUID	Universally Unique IDentifier

## Workflow of the MoP



## Single HDD Failure

1. Each Baremetal server will be provisioned with two HDD drives in order to act as Boot Disk in Raid 1 configuration. In case of single HDD failure, since there is Raid 1 level redundancy, the faulty HDD drive can be Hot Swapped.

2. Refer to the procedure in order to replace a faulty component on UCS C240 M4 server here: [Replacing the Server Components](#)
3. In case of single HDD failure, only the faulty HDD will be Hot Swapped and hence no BIOS upgrade procedure is required after you replace new disks.
4. After you replace the disks, wait for the data sync between the disks. It might take a couple of hours to complete.
5. In an OpenStack based (Ultra-M) solution, UCS 240M4 baremetal server can take up one of these roles: Compute, OSD-Compute, Controller and OSPD.
6. The steps required in order to handle the single HDD failure in each of these server roles are same and this section describes the health checks to be performed before the Hot Swap of the disk.

## Single HDD Failure on a Compute Server

1. If the failure of HDD drives is observed in UCS 240M4 which acts as a Compute node, perform these health checks before you initiate the Hot Swap procedure of the faulty disk.
2. Identify the VMs running on this server and verify the status of the functions are good.

### Identify the VMs Hosted in the Compute Node

Identify the VMs that are hosted on the Compute server and verify that they are active and running.

The Compute server contains CPS VMs/Elastic Services Controller (ESC) combination of VMs:

```
[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 |
```

**Note:** In the output shown here, the first column corresponds to the Universally Unique Identifier (UUID), the second column is the VM name and the third column is the hostname where the VM is present.

### Health Checks

1. Log in to the ESC hosted in the compute node and check the status.

```
[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 |
```

2. Log in to the UAS hosted in the compute node and check the status.

```
[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 |
```

3. If health checks are fine, proceed with the faulty disk Hot Swap procedure and wait for the data sync as it might take a couple of hours to complete. Refer to: [Replacing the Server Components](#)

4. Repeat these health check procedures in order to confirm that the health status of the VMs hosted on compute node are restored.

## Single HDD Failure on a Controller Server

1. If the failure of the HDD drives is observed in UCS 240M4 which acts as the Controller node, perform these health checks before you initiate the Hot Swap procedure of the faulty disk.

2. Check the Pacemaker status on the controllers.

3. Log in to one of the active controllers and check the Pacemaker status. All services must be running on the available controllers and stopped on the failed controller.

```
[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 |
```

4. Check the MariaDB status in the active controllers.

```
[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
```

Verify that these lines are present for each active controller:

```
[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
```

## 5. Check **Rabbitmq** status in the active controllers.

```
[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
```

6. If health checks are fine, proceed with faulty disk Hot Swap procedure and wait for the data sync as it might take a couple of hours to complete. Refer to: [Replacing the Server Components](#)

7. Repeat these health check procedures in order to confirm that the health status on the controller is restored.

## Single HDD Failure on an OSD-Compute Server

If the failure of HDD drives is observed in UCS 240M4 which acts as an OSD-Compute node, perform these health checks before you initiate the Hot Swap procedure of the faulty disk.

### Identify the VMs Hosted in the OSD-Compute Node

## 1. The Compute server contains 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 |
```

**Note:** In the output shown here, the first column corresponds to the (UUID), the second column is the VM name and the third column is the hostname where the VM is present.

## 2. Ceph processes are active on the OSD-Compute server.

```
[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 |
```

## 3. Verify that the mapping of OSD (HDD disk) to Journal (SSD) is good.

```
[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 |
```

## 4. Verify that the Ceph health and the OSD tree mapping is good.

```
[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 |
```

```
[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 |
```

5. If health checks are fine, proceed with the faulty disk Hot Swap procedure and wait for the data sync as it might take a couple of hours to complete. Refer to [Replacing the Server Components](#)

6. Repeat these health check procedures in order to confirm that the health status of the VMs hosted on OSD-Compute node are restored.

## Single HDD Failure on an OSPD Server

1. If the failure of the HDD drives is observed in UCS 240M4, which acts as an OSPD node, it is recommended you perform these checks before you initiate the Hot Swap procedure of the faulty disk.

2. Check the status of the OpenStack stack and the node list.

```
[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 |
```

3. Check if all the undercloud services are in loaded, active and running status from the OSPD node.

```
[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 |
```

4. If health checks are fine, proceed with faulty disk Hot Swap procedure and wait for the data sync as it might take a couple of hours to complete. Refer to [Replacing the Server Components](#)

5. Repeat these health check procedures in order to confirm that the health status of the OSPD node is restored.