

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

Contents

[Introduction](#)

[Background Information](#)

[Abbreviations](#)

[Workflow of the MoP](#)

[Single HDD Failure](#)

[Single HDD Failure on Compute Server](#)

[Single HDD Failure on Controller Server](#)

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

[Single HDD Failure on OSPD Server](#)

Introduction


This document describes the steps required to replace the faulty Hard Disk Drive (HDD) drive in the server in an Ultra-M setup that hosts StarOS Virtual Network Functions (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 depicted in this image:




UltraM Architecture This document is intended for the Cisco personnel who are familiar with 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.

Note: Ultra M 5.1.x release is considered in order to define the procedures in this document.

Abbreviations

VNF	Virtual Network Function
CF	Control Function
SF	Service 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 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. Procedure to replace a faulty component on UCS C240 M4 server can be referred from: [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 replacing new disks.
4. After replacing the disks, wait for the data sync between the disks. It might take 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. The steps required in order to handle the single HDD failure in each of these server roles is the same and this section describes the health checks to be performed before the hot swap of the disk.

Single HDD Failure on 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 finally perform the hot swap 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. There can be two possibilities:

1. The Compute server contains only SF VM.

```
[stack@director ~]$ nova list --field name,host | grep compute-10
| 49ac5f22-469e-4b84-badc-031083db0533 | VNF2-DEPLOYM_s8_0_8bc6cc60-15d6-4ead-8b6a-10e75d0e134d |
pod1-compute-10.localdomain | ACTIVE|
```

2. The Compute server contains CF/ESC/EM/UAS 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 |
| 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.

Health Checks:

1. Log in to the StarOS VNF and identify the card that corresponds to the SF or CF VM. Use the

UUID of the SF or CF VM identified from the section "Identify the VMs hosted in the Compute Node", and identify the card that corresponds to the UUID.

```
[local]VNF2# show card hardware
Tuesday might 08 16:49:42 UTC 2018
<snip>
```

Card 8:

```
Card Type          : 4-Port Service Function Virtual Card
CPU Packages       : 26 [#0, #1, #2, #3, #4, #5, #6, #7, #8, #9, #10, #11, #12, #13, #14,
#15, #16, #17, #18, #19, #20, #21, #22, #23, #24, #25]
CPU Nodes          : 2
CPU Cores/Threads  : 26
Memory             : 98304M (qvpc-di-large)
UUID/Serial Number : 49AC5F22-469E-4B84-BADC-031083DB0533
<snip>
```

```
[local]VNF2# show card hardware
Tuesday might 08 16:49:42 UTC 2018
<snip>
```

Card 2:

```
Card Type          : Control Function Virtual Card
CPU Packages       : 8 [#0, #1, #2, #3, #4, #5, #6, #7]
CPU Nodes          : 1
CPU Cores/Threads  : 8
Memory             : 16384M (qvpc-di-large)
UUID/Serial Number : F9C0763A-4A4F-4BBD-AF51-BC7545774BE2
<snip>
```

2. Verify the status of the card.

```
[local]VNF2# show card table
Tuesday might 08 16:52:53 UTC 2018
```

Slot	Card Type	Oper State	SPOF	Attach
-----	-----	-----	----	-----
1: CFC	Control Function Virtual Card	Active	No	
2: CFC	Control Function Virtual Card	Standby	-	
3: FC	4-Port Service Function Virtual Card	Active	No	
4: FC	4-Port Service Function Virtual Card	Active	No	
5: FC	4-Port Service Function Virtual Card	Active	No	
6: FC	4-Port Service Function Virtual Card	Active	No	
7: FC	4-Port Service Function Virtual Card	Active	No	
8: FC	4-Port Service Function Virtual Card	Active	No	
9: FC	4-Port Service Function Virtual Card	Active	No	
10: FC	4-Port Service Function Virtual Card	Standby	-	

3. Log in to the ESC hosted in the Compute node and check the status.

```
[local]VNF2# show card table
Tuesday might 08 16:52:53 UTC 2018
```

Slot	Card Type	Oper State	SPOF	Attach
-----	-----	-----	----	-----
1: CFC	Control Function Virtual Card	Active	No	
2: CFC	Control Function Virtual Card	Standby	-	
3: FC	4-Port Service Function Virtual Card	Active	No	
4: FC	4-Port Service Function Virtual Card	Active	No	
5: FC	4-Port Service Function Virtual Card	Active	No	
6: FC	4-Port Service Function Virtual Card	Active	No	
7: FC	4-Port Service Function Virtual Card	Active	No	
8: FC	4-Port Service Function Virtual Card	Active	No	

9: FC	4-Port Service Function Virtual Card	Active	No
10: FC	4-Port Service Function Virtual Card	Standby	-

4. Log in to the EM hosted in the Compute node and check the status.

```
[local]VNF2# show card table
```

Tuesday might 08 16:52:53 UTC 2018

Slot	Card Type	Oper State	SPOF	Attach
1: CFC	Control Function Virtual Card	Active	No	
2: CFC	Control Function Virtual Card	Standby	-	
3: FC	4-Port Service Function Virtual Card	Active	No	
4: FC	4-Port Service Function Virtual Card	Active	No	
5: FC	4-Port Service Function Virtual Card	Active	No	
6: FC	4-Port Service Function Virtual Card	Active	No	
7: FC	4-Port Service Function Virtual Card	Active	No	
8: FC	4-Port Service Function Virtual Card	Active	No	
9: FC	4-Port Service Function Virtual Card	Active	No	
10: FC	4-Port Service Function Virtual Card	Standby	-	

5. Log in to the UAS hosted in the Compute node and check the status.

```
[local]VNF2# show card table
```

Tuesday might 08 16:52:53 UTC 2018

Slot	Card Type	Oper State	SPOF	Attach
1: CFC	Control Function Virtual Card	Active	No	
2: CFC	Control Function Virtual Card	Standby	-	
3: FC	4-Port Service Function Virtual Card	Active	No	
4: FC	4-Port Service Function Virtual Card	Active	No	
5: FC	4-Port Service Function Virtual Card	Active	No	
6: FC	4-Port Service Function Virtual Card	Active	No	
7: FC	4-Port Service Function Virtual Card	Active	No	
8: FC	4-Port Service Function Virtual Card	Active	No	
9: FC	4-Port Service Function Virtual Card	Active	No	
10: FC	4-Port Service Function Virtual Card	Standby	-	

6. If health checks are fine, proceed with faulty disk hot swap procedure and wait for the data sync as it takes hours to complete. Refer to [Replacing the Server Components](#).

7. 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 Controller Server

1. If the failure of HDD drives are observed in UCS 240M4 which acts as a Controller node, follow the health checks before you perform hot swap of the faulty disk.

2. Check the Pacemaker status on 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.

```
[local]VNF2# show card table
```

Tuesday might 08 16:52:53 UTC 2018

Slot	Card Type	Oper State	SPOF	Attach
1: CFC	Control Function Virtual Card	Active	No	

2: CFC	Control Function Virtual Card	Standby	-
3: FC	4-Port Service Function Virtual Card	Active	No
4: FC	4-Port Service Function Virtual Card	Active	No
5: FC	4-Port Service Function Virtual Card	Active	No
6: FC	4-Port Service Function Virtual Card	Active	No
7: FC	4-Port Service Function Virtual Card	Active	No
8: FC	4-Port Service Function Virtual Card	Active	No
9: FC	4-Port Service Function Virtual Card	Active	No
10: FC	4-Port Service Function Virtual Card	Standby	-

4. Check 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
```

5. 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
```

6. 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
```

7. If health checks are fine, proceed with faulty disk hot swap procedure and wait for the data sync as it takes hours to complete. Refer to [Replacing the Server Components](#).

8. Repeat these health check procedures in order to confirm the health status of the controller is restored.

Single HDD Failure on OSD-Compute Server

If the failure of HDD drives are observed in UCS 240M4 which acts as sn OSD-Compute node, perform these health checks before you perform hot swap of the faulty disk.

Identify the VMs Hosted in the OSD-Compute Node:

Identify the VMs that are hosted on the Compute server. There can be two possibilities:

1. The OSD-Compute server contains EM/UAS/Auto-Deploy/Auto-IT combination of VMs.

```
[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
```

2. The Compute server contains CF/ESC/EM/UAS combination of VMs.

```
[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.

3. 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 |

```

4. 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 |

```

5. 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 |

```

6. If health checks are fine, proceed with faulty disk hot swap procedure and wait for the data sync as it takes hours to complete. Refer to [Replacing the Server Components](#).

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

Single HDD Failure on OSPD Server

1. If the failure of HDD drives are observed in UCS 240M4 which acts as an OSPD node, perform these health checks before you initiate the hot swap of the faulty disk.

2. Check the status of 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 OSP-D 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 takes 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 nodes are restored.