

Ultra-M UCS 240M4 Single HDD Failure - Hot Swappable Procedure - CPAR

Contents

[Introduction](#)

[Background Information](#)

[Abbreviations](#)

[MoP Workflow](#)

[Single HDD Failure](#)

[Single HDD Failure on Compute Server](#)

[Identify VMs Hosted in Compute Node](#)

[Health Checks](#)

[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) in a server in an Ultra-M setup.


This procedure applies for an Openstack environment with NEWTON version where ESC does not manage CPAR and CPAR is installed directly on the Virtual Machine (VM) deployed on the Openstack.

Background Information

Ultra-M is a pre-packaged and validated virtualized mobile packet core solution designed in order to simplify the deployment of Virtual Network Functions (VNFs). OpenStack is the Virtual 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:




This document is intended for the Cisco personnel 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
MoP	Method of Procedure
OSD	Object Storage Disks
OSPD	OpenStack Platform Director
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

MoP Workflow



Single HDD Failure

1. Each Baremetal server is 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. 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 is hot swapped and hence no BIOS upgrade procedure is required after the replacement of new disks.
4. After the replacement of the disks, you must wait for the data sync between the disks. It might take hours to complete.
5. In 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 for handling single HDD failure in each of these server roles are the same and the section here 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 are observed in UCS 240M4 which acts as compute node, do this health checks before you perform hot swap of the faulty disk.
2. Identify the VMs that runs on this server and verify that the status of the functions are good.

Identify VMs Hosted in Compute Node

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

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

Health Checks

Step 1. Run the command **/opt/CSCOar/bin/arstatus** at Operating System (OS) level.

```
[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 ~]#
```

Step 2. Run the command **/opt/CSCOar/bin/aregcmd** at OS level and enter the admin credentials. Verify that the CPAR Health is 10 out of 10 and the exit 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
```

Step 3. Run the command **netstat | grep diameter** and verify that all Diameter Routing Agent (DRA) connections are established.

The output mentioned here is for an environment where Diameter links are expected. If fewer links are displayed, this represents a disconnection from the DRA that needs to be analyzed.

```
[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
```

Step 4. Check that the TPS log shows requests being processed by CPAR. The values highlighted represent the TPS and those are the ones you need to pay attention to.

The value of TPS must not exceed 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
```

Step 5. Look for any “error” or “alarm” messages in name_radius_1_log

```
[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
```

Step 6. In order to verify the amount of memory that the CPAR process uses, run the command:

```
[root@wscaaa04 ~]# tail -f /opt/CSC0ar/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
```

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

This highlighted value must be lower than 7Gb, which is the maximum allowed at application level.

Step 7. In order to verify the disk utilization, run the command **df -h**.

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

This overall value must be lower than 80%, if it's more than 80% then identify the unnecessary files and clean it up.

Step 8. Verify that there is no 'core' file generated.

- Core file is generated in the event of application crash when CPAR is unable to handle an exception and it's generated in these two location:

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

There must not be any core files located in these two location. If found, raise a Cisco TAC case in order to identify the root cause of such exception and attach the core files for debugging.

- If health checks are fine, proceed with faulty disk hot swap procedure and wait for the data sync as it takes hours to complete .

[Replacing the Server Components](#)

- Repeat the 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

- If the failure of HDD drives are observed in UCS 240M4 which acts as Controller node, do these health checks before you perform hot swap of the faulty disk.
- Check the Pacemaker status on controllers.
- 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.

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

- Check MariaDB status in the active controllers.

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

- Verify that these lines are present for each active controller:

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

- Check **Rabbitmq** status in the active controllers.

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

- If health checks are fine, proceed with faulty disk hot swap procedure and wait for the data sync as it takes hours to complete .

[Replacing the Server Components](#)

- Repeat the health check procedures in order to confirm that the health status on controller are restored.

Single HDD Failure on OSD-Compute Server

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

1. Identify the VMs Hosted in the OSD Compute Node
2. Identify the VMs that are hosted on the **compute** server

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

- CEPH processes are active on the osd-compute server.

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

- Verify that the mapping of OSD (HDD disk) to Journal (SSD) are good.

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

- Verify that the ceph health and osd tree mapping are good.

```
[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 -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
```

- If health checks are fine, proceed with faulty disk hot swap procedure and wait for the data sync as it takes hours to complete.

[Replacing the Server Components](#)

- Repeat the 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 OSPD Server

- If the failure of HDD drives are observed in UCS 240M4 which acts as OSPD node, do the health checks before you perform hot swap of the faulty disk.
- Check the status of openstack stack and the node list.

```
[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
```

- Check if all the undercloud services are in loaded, active and running status from the OSP-D node.

```
[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

- If health checks are fine, proceed with faulty disk hot swap procedure and wait for the data sync as it takes hours to complete.

[Replacing the Server Components](#)

- Repeat the health check procedures in order to confirm that the health status of the OSPD node are restored.