

Motherboard Replacement in Ultra-M UCS 240M4 Server - CPS

Contents

[Introduction](#)

[Background Information](#)

[Abbreviations](#)

[Workflow of the MoP](#)

[Motherboard Replacement in Ultra-M Setup](#)

[Motherboard Replacement in Compute Node](#)

[Identify the VMs Hosted in the Compute Node](#)

[Graceful Power Off](#)

[Compute Node Hosts CPS/ESC VMs](#)

[ESC Backup](#)

[Backup ESC Database](#)

[Replace Motherboard](#)

[Restore the VMs](#)

[Compute Node Hosts CPS, ESC](#)

[Restore the CPS VMs](#)

[Motherboard Replacement in OSD Compute Node](#)

[Put CEPH in Maintenance mode](#)

[Identify the VMs Hosted in the Osd-Compute Node](#)

[Graceful Power Off](#)

[Case 1. OSD-Compute node Hosts ESC](#)

[Replace Motherboard](#)

[Move CEPH out of Maintenance Mode](#)

[Restore the VMs](#)

[Case 1. OSD-Compute node hosting ESC or CPS VMs](#)

[Motherboard Replacement in Controller Node](#)

[Verify Controller Status and put cluster in Maintenance mode](#)

[Replace Motherboard](#)

[Restore Cluster Status](#)

Introduction

This document describes the steps required to replace a faulty motherboard of a server in an Ultra-M setup that hosts CPS Virtual Network Functions (VNFs).

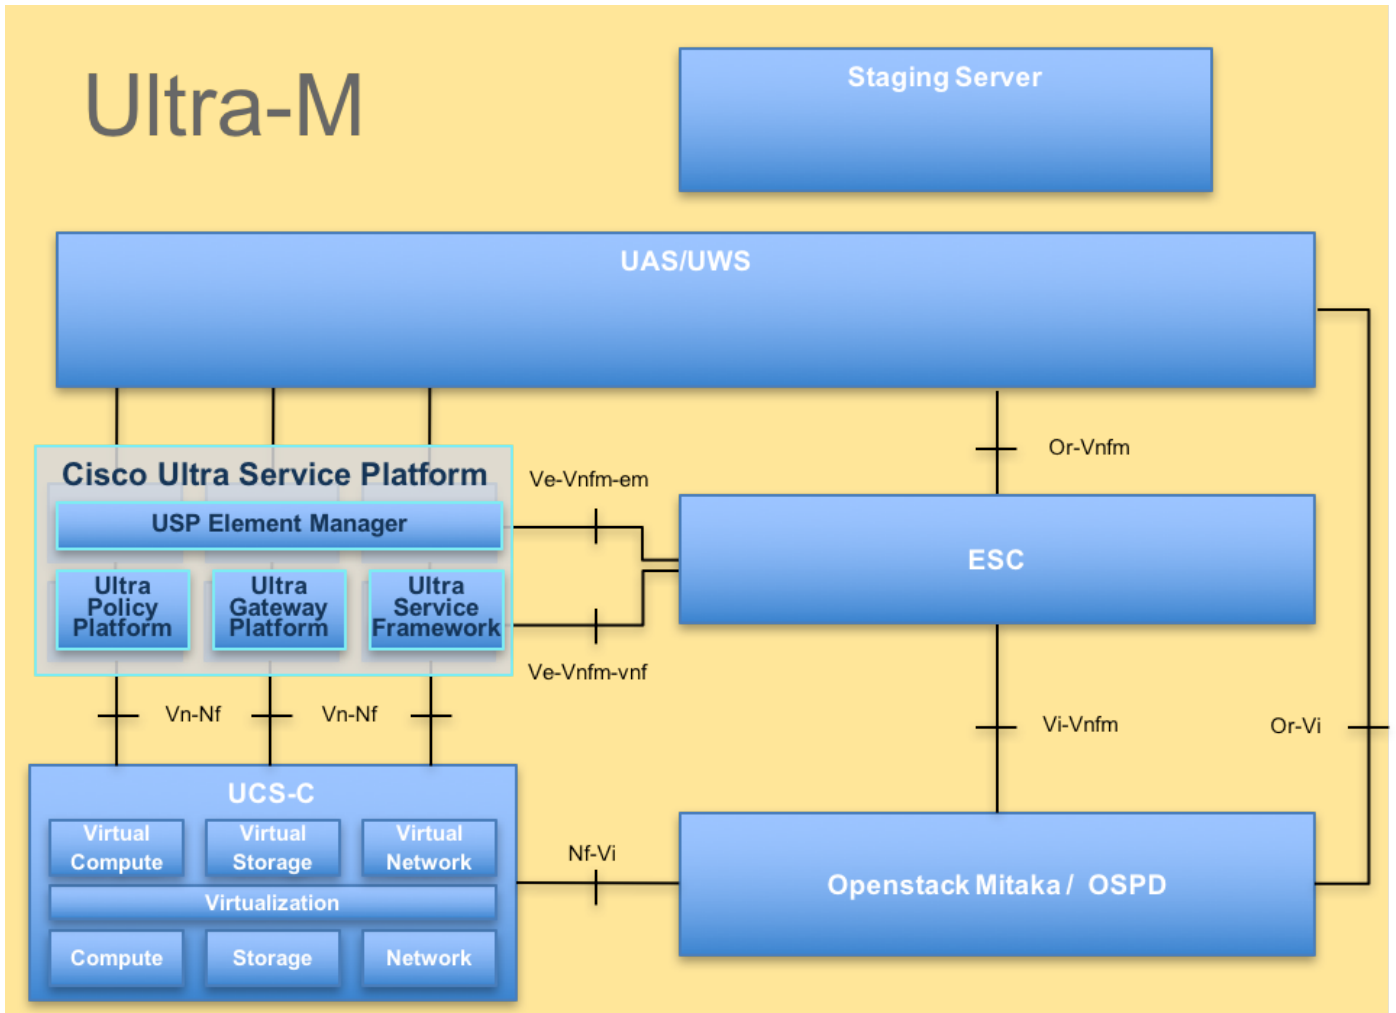
Background Information

Ultra-M is a pre-packaged and validated virtualized mobile packet core solution that is designed in order 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:



This document is intended for Cisco personnel who are familiar with Cisco Ultra-M platform and it details the steps that are required to be carried out at OpenStack and StarOS VNF level at the time of the motherboard replacement in a server.


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

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



Motherboard Replacement in Ultra-M Setup

In an Ultra-M setup, there can be scenarios where a motherboard replacement is required in the following server types: Compute, OSD-Compute and Controller.

Note: The boot disks with the openstack installation are replaced after the replacement of the motherboard. Hence there is no requirement to add the node back to overcloud. Once the server is powered ON after the replacement activity, it would enroll itself back to the overcloud stack.

Motherboard Replacement in Compute Node

Before the activity, the VMs hosted in the Compute node are gracefully shutdown. Once the Motherboard has been replaced, the VMs are restored back.

Identify the VMs Hosted in the Compute Node

Identify the VMs that are hosted on the compute server.

The compute server contains CPS or Elastic Services Controller (ESC) 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 |
| f9c0763a-4a4f-4bbd-af51-bc7545774be2 | VNF2-DEPLOYM_c1_0_df4be88d-b4bf-4456-945a-
3812653ee229 | pod1-compute-8.localdomain |
| 75528898-ef4b-4d68-b05d-882014708694 | VNF2-ESC-ESC-
0 | pod1-compute-8.localdomain |
```

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. The parameters from this output will be used in subsequent sections.

Graceful Power Off

Compute Node Hosts CPS/ESC VMs

Step 1. Log in to the ESC node that corresponds to the VNF and check the status of the VMs.

```
[admin@VNF2-esc-esc-0 ~]$ cd /opt/cisco/esc/esc-confd/esc-cli
[admin@VNF2-esc-esc-0 esc-cli]$ ./esc_nc_cli get esc_datamodel | egrep --color
"<state>|<vm_name>|<vm_id>|<deployment_name>"
<snip>
<state>SERVICE_ACTIVE_STATE</state>
    <vm_name>VNF2-DEPLOYM_c1_0_df4be88d-b4bf-4456-945a-3812653ee229</vm_name>
    <state>VM_ALIVE_STATE</state>
    <vm_name>VNF2-DEPLOYM_c3_0_3e0db133-c13b-4e3d-ac14-
    <state>VM_ALIVE_STATE</state>
```

```

<deployment_name>VNF2-DEPLOYMENT-em</deployment_name>
  <vm_id>507d67c2-1d00-4321-b9d1-da879af524f8</vm_id>
  <vm_id>dc168a6a-4aeb-4e81-abd9-91d7568b5f7c</vm_id>
  <vm_id>9ffec58b-4b9d-4072-b944-5413bf7fcf07</vm_id>
  <state>SERVICE_ACTIVE_STATE</state>
  <vm_name>VNF2-DEPLOYM_XXXX_0_c8d98f0f-d874-45d0-af75-88a2d6fa82ea</vm_name>
  <state>VM_ALIVE_STATE</state>

```

<snip>

Step 2. Stop the CPS VMs one-by-one with the use of its VM Name. (VM Name noted from section Identify the VMs hosted in the Compute Node).

```

[admin@VNF2-esc-esc-0 esc-cli]$ ./esc_nc_cli vm-action STOP VNF2-DEPLOYM_c1_0_df4be88d-b4bf-4456-945a-3812653ee229

```

```

[admin@VNF2-esc-esc-0 esc-cli]$ ./esc_nc_cli vm-action STOP VNF2-DEPLOYM_XXXX_0_c8d98f0f-d874-45d0-af75-88a2d6fa82ea

```

Step 3. After it stops, the VMs must enter the SHUTOFF state.

```

[admin@VNF2-esc-esc-0 ~]$ cd /opt/cisco/esc/esc-confd/esc-cli
[admin@VNF2-esc-esc-0 esc-cli]$ ./esc_nc_cli get esc_datamodel | egrep --color
"<state>|<vm_name>|<vm_id>|<deployment_name>"
<snip>
<state>SERVICE_ACTIVE_STATE</state>
  <vm_name>VNF2-DEPLOYM_c1_0_df4be88d-b4bf-4456-945a-3812653ee229</vm_name>
  <state>VM_SHUTOFF_STATE</state>
  <vm_name>VNF2-DEPLOYM_c3_0_3e0db133-c13b-4e3d-ac14-
  <state>VM_ALIVE_STATE</state>
<deployment_name>VNF2-DEPLOYMENT-em</deployment_name>
  <vm_id>507d67c2-1d00-4321-b9d1-da879af524f8</vm_id>
  <vm_id>dc168a6a-4aeb-4e81-abd9-91d7568b5f7c</vm_id>
  <vm_id>9ffec58b-4b9d-4072-b944-5413bf7fcf07</vm_id>
  <state>SERVICE_ACTIVE_STATE</state>
  <vm_name>VNF2-DEPLOYM_XXXX_0_c8d98f0f-d874-45d0-af75-88a2d6fa82ea</vm_name>
  <state>VM_SHUTOFF_STATE</state>
<snip>

```

Step 4. Log in to the ESC hosted in the compute node and check if it is in the master state. If yes, switch the ESC to standby mode:

```

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

```

```

[admin@VNF2-esc-esc-0 ~]$ sudo service keepalived stop
Stopping keepalived: [ OK ]

```

```

[admin@VNF2-esc-esc-0 ~]$ escadm status
1 ESC status=0 In SWITCHING_TO_STOP state. Please check status after a while.

```

```

[admin@VNF2-esc-esc-0 ~]$ 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!

```

ESC Backup

Step 1. ESC has 1:1 redundancy in UltraM Solution. 2 ESC VMs are deployed and supports single failure in UltraM. i.e. the system is recovered if there is a single failure in the system.

Note: If there is more than a single failure, it is not supported and may require redeployment of the system.

ESC backup details:

- Running configuration
- ConfD CDB DB
- ESC Logs
- Syslog configuration

Step 2. The frequency of ESC DB backup is tricky and needs to be handled carefully as ESC monitors and maintains the various state machines for various VNF VMs deployed. It is advised that these backups are performed after following activities in given VNF/POD/Site.

Step 3. Verify the health of ESC is good for using **health.sh** script.

```
[root@auto-test-vnfm1-esc-0 admin]# escadm status
0 ESC status=0 ESC Master Healthy

[root@auto-test-vnfm1-esc-0 admin]# health.sh
esc ui is disabled -- skipping status check
esc_monitor start/running, process 836
esc_mona is up and running ...
vimmanager start/running, process 2741
vimmanager start/running, process 2741
esc_confd is started
tomcat6 (pid 2907) is running... [ OK ]
postgresql-9.4 (pid 2660) is running...
ESC service is running...
Active VIM = OPENSTACK
ESC Operation Mode=OPERATION

/opt/cisco/esc/esc_database is a mountpoint

===== ESC HA (MASTER) with DRBD =====

DRBD_ROLE_CHECK=0
MNT_ESC_DATABASE_CHECK=0
VIMMANAGER_RET=0
ESC_CHECK=0
STORAGE_CHECK=0
ESC_SERVICE_RET=0
MONA_RET=0
ESC_MONITOR_RET=0

=====

ESC HEALTH PASSED
```

Step 4. Take a backup of the Running configuration and transfer the file to the backup server.

```
[root@auto-test-vnfm1-esc-0 admin]# /opt/cisco/esc/confd/bin/confd_cli -u admin -C

admin connected from 127.0.0.1 using console on auto-test-vnfm1-esc-0.novalocal
auto-test-vnfm1-esc-0# show running-config | save /tmp/running-esc-12202017.cfg
auto-test-vnfm1-esc-0#exit

[root@auto-test-vnfm1-esc-0 admin]# ll /tmp/running-esc-12202017.cfg
```

```
-rw-----. 1 tomcat tomcat 25569 Dec 20 21:37 /tmp/running-esc-12202017.cfg
```

Backup ESC Database

Step 1. Login to ESC VM and run this command before you take the backup.

```
[admin@esc ~]# sudo bash
[root@esc ~]# cp /opt/cisco/esc/esc-scripts/esc_dbtool.py /opt/cisco/esc/esc-
scripts/esc_dbtool.py.bkup
[root@esc esc-scripts]# sudo sed -i "s,'pg_dump','usr/pgsql-9.4/bin/pg_dump,"
/opt/cisco/esc/esc-scripts/esc_dbtool.py
```

```
#Set ESC to mainenance mode
```

```
[root@esc esc-scripts]# escadm op_mode set --mode=maintenance
```

Step 2. Check ESC mode and ensure it is in maintenance mode.

```
[root@esc esc-scripts]# escadm op_mode show
```

Step 3. Backup database using database backup restore tool available in ESC.

```
[root@esc scripts]# sudo /opt/cisco/esc/esc-scripts/esc_dbtool.py backup --file
scp://<username>:<password>@<backup_vm_ip>:<filename>
```

Step 4. Set ESC Back to Operation Mode and confirm the mode.

```
[root@esc scripts]# escadm op_mode set --mode=operation
```

```
[root@esc scripts]# escadm op_mode show
```

Step 5. Navigate to the scripts directory and collect the logs.

```
[root@esc scripts]# /opt/cisco/esc/esc-scripts
```

```
sudo ./collect_esc_log.sh
```

Step 6. To create a snapshot of the ESC first shutdown the ESC.

```
shutdown -r now
```

Step 7. From OSPD create an image snapshot.

- ```
nova image-create --poll esc1 esc_snapshot_27aug2018
```

Step 8. Verify that the snapshot is created

```
openstack image list | grep esc_snapshot_27aug2018
```

Step 9. Start the ESC from OSPD

```
nova start esc1
```

Step 10. Repeat the same procedure on standby ESC VM & transfer the logs to the backup server.



Step 11. Collect syslog configuration backup on both the ESC VMS and transfer them to the backup server.

```
[admin@auto-test-vnfm2-esc-1 ~]$ cd /etc/rsyslog.d
[admin@auto-test-vnfm2-esc-1 rsyslog.d]$ls /etc/rsyslog.d/00-escmanager.conf
00-escmanager.conf

[admin@auto-test-vnfm2-esc-1 rsyslog.d]$ls /etc/rsyslog.d/01-messages.conf
01-messages.conf

[admin@auto-test-vnfm2-esc-1 rsyslog.d]$ls /etc/rsyslog.d/02-mona.conf
02-mona.conf

[admin@auto-test-vnfm2-esc-1 rsyslog.d]$ls /etc/rsyslog.conf
rsyslog.conf
```

## Replace Motherboard

Step 1. The steps in order to replace the motherboard in a UCS C240 M4 server can be referred from:

[Cisco UCS C240 M4 Server Installation and Service Guide](#)

Step 2. Log in to the server with the use of the CIMC IP.

Step 3. Perform BIOS upgrade if the firmware is not as per the recommended version used previously. Steps for BIOS upgrade are given here:

[Cisco UCS C-Series Rack-Mount Server BIOS Upgrade Guide](#)

## Restore the VMs

### Compute Node Hosts CPS, ESC

#### Recovery of ESC VM

Step 1. ESC VM is recoverable if the VM is in error or shutdown state do a hard reboot to bring up of the impacted VM. Run these steps to recover ESC.

Step 2. Identify the VM which is in ERROR or Shutdown state, once identified hard-reboot the ESC VM. In this example, reboot auto-test-vnfm1-ESC-0.

```
[root@tb1-baremetal scripts]# nova list | grep auto-test-vnfm1-ESC-
| f03e3cac-a78a-439f-952b-045aea5b0d2c | auto-test-vnfm1-ESC-
0 | ACTIVE | - | running | auto-testautovnf1-
uas-orchestration=172.57.12.11; auto-testautovnf1-uas-
management=172.57.11.3
|
| 79498e0d-0569-4854-a902-012276740bce | auto-test-vnfm1-ESC-
1 | ACTIVE | - | running | auto-testautovnf1-
uas-orchestration=172.57.12.15; auto-testautovnf1-uas-
management=172.57.11.5
|
```

```
[root@tb1-baremetal scripts]# [root@tb1-baremetal scripts]# nova reboot --hard f03e3cac-a78a-439f-952b-045aea5b0d2c\
```

Request to reboot server <Server: auto-test-vnfm1-ESC-0> has been accepted.

```
[root@tb1-baremetal scripts]#
```

Step 3. If ESC VM is deleted and needs to be brought up again.

```
[stack@pod1-ospd scripts]$ nova list |grep ESC-1
| c566efbf-1274-4588-a2d8-0682e17b0d41 | vnfl-ESC-ESC-
1 | ACTIVE | - | running | vnfl-
UAS-uas-orchestration=172.168.11.14; vnfl-UAS-uas-
management=172.168.10.4
```

```
[stack@pod1-ospd scripts]$ nova delete vnfl-ESC-ESC-1
Request to delete server vnfl-ESC-ESC-1 has been accepted.
```

Step 4. From OSPD, check new ESC VM is ACTIVE/running:

```
[stack@pod1-ospd ~]$ nova list|grep -i esc
| 934519a4-d634-40c0-a51e-fc8d55ec7144 | vnfl-ESC-ESC-
0 | ACTIVE | - | running | vnfl-
UAS-uas-orchestration=172.168.11.13; vnfl-UAS-uas-
management=172.168.10.3
| 2601b8ec-8ff8-4285-810a-e859f6642ab6 | vnfl-ESC-ESC-
1 | ACTIVE | - | running | vnfl-
UAS-uas-orchestration=172.168.11.14; vnfl-UAS-uas-
management=172.168.10.6
```

#Log in to new ESC and verify Backup state. You may execute health.sh on ESC Master too.

```
...
#####
ESC on vnfl-esc-esc-1.novalocal is in BACKUP state.
#####
```

```
[admin@esc-1 ~]$ escadm status
0 ESC status=0 ESC Backup Healthy
```

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

ESC HEALTH PASSED

```
[admin@esc-1 ~]$ cat /proc/drbd
version: 8.4.7-1 (api:1/proto:86-101)
GIT-hash: 3a6a769340ef93b1ba2792c6461250790795db49 build by mockbuild@Build64R6, 2016-01-12
13:27:11
1: cs:Connected ro:Secondary/Primary ds:UpToDate/UpToDate C r-----
ns:0 nr:504720 dw:3650316 dr:0 al:8 bm:0 lo:0 pe:0 ua:0 ap:0 ep:1 wo:f oos:0
```

Step 5. If ESC VM is unrecoverable and requires the restore of the database, please restore the database from the previously taken backup.

Step 6. For ESC database restore, we have to ensure the esc service is stopped before restoring the database; For ESC HA, execute in secondary VM first and then the primary VM.

```
service keepalived stop
```

Step 7. Check ESC service status and ensure everything is stopped in both Primary and Secondary VMs for HA

```
escadm status
```

Step 8. Execute the script to restore the database. As part of the restoration of the DB to the newly created ESC instance, the tool will also promote one of the instances to be a primary ESC, mount its DB folder to the drbd device and will start the PostgreSQL database.

```
/opt/cisco/esc/esc-scripts/esc_dbtool.py restore --file
scp://<username>:<password>@<backup_vm_ip>:<filename>
```

Step 9. Restart ESC service to complete the database restore.

For HA execute in both VMs, restart the keepalived service

```
service keepalived start
```

Step 10. Once the VM is successfully restored and running; ensure all the syslog specific configuration is restored from the previous successful known backup. Ensure that it is restored in all the ESC VMs

```
[admin@auto-test-vnfm2-esc-1 ~]$
[admin@auto-test-vnfm2-esc-1 ~]$ cd /etc/rsyslog.d
[admin@auto-test-vnfm2-esc-1 rsyslog.d]$ls /etc/rsyslog.d/00-escmanager.conf
00-escmanager.conf
```

```
[admin@auto-test-vnfm2-esc-1 rsyslog.d]$ls /etc/rsyslog.d/01-messages.conf
01-messages.conf
```

```
[admin@auto-test-vnfm2-esc-1 rsyslog.d]$ls /etc/rsyslog.d/02-mona.conf
02-mona.conf
```

```
[admin@auto-test-vnfm2-esc-1 rsyslog.d]$ls /etc/rsyslog.conf
rsyslog.conf
```

Step 11. If the ESC needs to be rebuilt from OSPD snapshot use below command using snapshot taken during backup.

```
nova rebuild --poll --name esc_snapshot_27aug2018 esc1
```

Step 12. Check the status of the ESC after rebuild is complete.

```
nova list --fileds name,host,status,networks | grep esc
```

Step 13. Check ESC health with below command.

```
health.sh
```

```
Copy Datamodel to a backup file
```

```
/opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli get esc_datamodel/opdata > /tmp/esc_opdata_`date
+%Y%m%d%H%M%S`.txt
```

## Restore the CPS VMs

The CPS VM would be in error state in the nova list:

```
[stack@director ~]$ nova list |grep VNF2-DEPLOYM_s9_0_8bc6cc60-15d6-4ead-8b6a-10e75d0e134d
| 49ac5f22-469e-4b84-badc-031083db0533 | VNF2-DEPLOYM_s9_0_8bc6cc60-15d6-4ead-8b6a-10e75d0e134d
| ERROR | - | NOSTATE |
```

Recover the CPS VM from the ESC:

```
[admin@VNF2-esc-esc-0 ~]$ sudo /opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli recovery-vm-action DO
VNF2-DEPLOYM_s9_0_8bc6cc60-15d6-4ead-8b6a-10e75d0e134d
[sudo] password for admin:
```

Recovery VM Action

```
/opt/cisco/esc/confd/bin/netconf-console --port=830 --host=127.0.0.1 --user=admin --
privKeyFile=/root/.ssh/confd_id_dsa --privKeyType=dsa --rpc=/tmp/esc_nc_cli.ZpRCGiieuW
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="1">
 <ok/>
</rpc-reply>
```

Monitor the yangesc.log:

```
admin@VNF2-esc-esc-0 ~]$ tail -f /var/log/esc/yangesc.log
```

```
...
14:59:50,112 07-Nov-2017 WARN Type: VM_RECOVERY_COMPLETE
14:59:50,112 07-Nov-2017 WARN Status: SUCCESS
14:59:50,112 07-Nov-2017 WARN Status Code: 200
14:59:50,112 07-Nov-2017 WARN Status Msg: Recovery: Successfully recovered VM [VNF2-
DEPLOYM_s9_0_8bc6cc60-15d6-4ead-8b6a-10e75d0e134d].
```

## When ESC Fails to Start VM

Step 1. In some cases, ESC will fail to start the VM due to an unexpected state. A workaround is to perform an ESC switchover by rebooting the Master ESC. The ESC switchover will take about a minute. Execute health.sh on the new Master ESC to verify it is up. When the ESC becomes Master, ESC may fix the VM state and start the VM. Since this operation is scheduled, you must wait 5-7 minutes for it to complete.

Step 2. You can monitor /var/log/esc/yangesc.log and /var/log/esc/escmanager.log. If you do NOT see VM getting recovered after 5-7 minutes, the user would need to go and do the manual recovery of the impacted VM(s).

Step 3. Once the VM is successfully restored and running; ensure all the syslog specific configuration is restored from the previous successful known backup. Ensure it is restored in all the ESC VMs.

```
root@autotestvnmflesc2:/etc/rsyslog.d# pwd
/etc/rsyslog.d
```

```
root@autotestvnmflesc2:/etc/rsyslog.d# ll
```

```
total 28
drwxr-xr-x 2 root root 4096 Jun 7 18:38 ./
drwxr-xr-x 86 root root 4096 Jun 6 20:33 ../]
-rw-r--r-- 1 root root 319 Jun 7 18:36 00-vnmf-proxy.conf
-rw-r--r-- 1 root root 317 Jun 7 18:38 01-ncs-java.conf
-rw-r--r-- 1 root root 311 Mar 17 2012 20-ufw.conf
-rw-r--r-- 1 root root 252 Nov 23 2015 21-cloudinit.conf
-rw-r--r-- 1 root root 1655 Apr 18 2013 50-default.conf
```

```
root@abautotestvnmlem-0:/etc/rsyslog.d# ls /etc/rsyslog.conf
rsyslog.conf
```

## Motherboard Replacement in OSD Compute Node

Before the activity, the VMs hosted in the Compute node are gracefully shutdown and the CEPH is put into maintenance mode. Once the Motherboard has been replaced, the VMs are restored back and CEPH is moved out of maintenance mode.

### Put CEPH in Maintenance mode

Step 1. Verify ceph osd tree status are up in the server

```
[heat-admin@pod1-osd-compute-1 ~]$ 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

Step 2. Log in to the OSD Compute node and put CEPH in the maintenance mode.

```
[root@pod1-osd-compute-1 ~]# sudo ceph osd set norebalance
```

```
[root@pod1-osd-compute-1 ~]# sudo ceph osd set noout
```

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

```
cluster eb2bb192-b1c9-11e6-9205-525400330666
health HEALTH_WARN
noout,norebalance,sortbitwise,require_jewel_osds flag(s) set
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 e194: 12 osds: 12 up, 12 in
flags noout,norebalance,sortbitwise,require_jewel_osds
pgmap v584865: 704 pgs, 6 pools, 531 GB data, 344 kobjects
1585 GB used, 11808 GB / 13393 GB avail
704 active+clean
client io 463 kB/s rd, 14903 kB/s wr, 263 op/s rd, 542 op/s wr
```

**Note:** When CEPH is removed, VNF HD RAID goes into the Degraded state but hd-disk must still be accessible

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

Identify the VMs that are hosted on the OSD compute server.

The compute server contains Elastic Services Controller (ESC) or CPS 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 |
| f9c0763a-4a4f-4bbd-af51-bc7545774be2 | VNF2-DEPLOYM_c1_0_df4be88d-b4bf-4456-945a-
3812653ee229 | pod1-compute-8.localdomain |
| 75528898-ef4b-4d68-b05d-882014708694 | VNF2-ESC-ESC-
0 | pod1-compute-8.localdomain |
| f5bd7b9c-476a-4679-83e5-303f0aae9309 | VNF2-UAS-uas-
0 | pod1-compute-8.localdomain |
```

**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. The parameters from this output will be used in subsequent sections.

## Graceful Power Off

### Case 1. OSD-Compute node Hosts ESC

Procedure to gracefully power of ESC or CPS VMs is same irrespective of whether the VMs are hosted in Compute or OSD-Compute node.

Follow steps from "Motherboard Replacement in Compute Node" to gracefully power off the VMs.

## Replace Motherboard

Step 1. The steps in order to replace motherboard in a UCS C240 M4 server can be referred from:

[Cisco UCS C240 M4 Server Installation and Service Guide](#)

Step 2. Log in to the server with the use of the CIMC IP

3. Perform BIOS upgrade if the firmware is not as per the recommended version used previously. Steps for BIOS upgrade are given here:

[Cisco UCS C-Series Rack-Mount Server BIOS Upgrade Guide](#)

## Move CEPH out of Maintenance Mode

Log into the OSD Compute node and move CEPH out of the maintenance mode.

```
[root@pod1-osd-compute-1 ~]# sudo ceph osd unset norebalance
[root@pod1-osd-compute-1 ~]# sudo ceph osd unset noout
```

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

```
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 e196: 12 osds: 12 up, 12 in
flags sortbitwise,require_jewel_osds
pgmap v584954: 704 pgs, 6 pools, 531 GB data, 344 kobjects
1585 GB used, 11808 GB / 13393 GB avail
704 active+clean
client io 12888 kB/s wr, 0 op/s rd, 81 op/s wr
```

## Restore the VMs

### Case 1. OSD-Compute node hosting ESC or CPS VMs

Procedure to restore of CF/ESC/EM/UAS VMs is same irrespective of whether the VMs are hosted in Compute or OSD-Compute node.

Follow steps from "Case 2. Compute Node Hosts CF/ESC/EM/UAS" to restore the VMs.

## Motherboard Replacement in Controller Node

### Verify Controller Status and put cluster in Maintenance mode

From OSPD, login to controller and verify pcs is in good state – all three controllers Online and galera showing all three controllers as Master.

```
[heat-admin@pod1-controller-0 ~]$ sudo pcs status
```

```
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod1-controller-2 (version 1.1.15-11.e17_3.4-e174ec8) - partition with quorum
Last updated: Mon Dec 4 00:46:10 2017 Last change: Wed Nov 29 01:20:52 2017 by hacluster via
crmd on pod1-controller-0
```

```
3 nodes and 22 resources configured
```

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

```
Full list of resources:
```

```
ip-11.118.0.42 (ocf::heartbeat:IPaddr2): Started pod1-controller-1
ip-11.119.0.47 (ocf::heartbeat:IPaddr2): Started pod1-controller-2
ip-11.120.0.49 (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]
ip-11.120.0.47 (ocf::heartbeat:IPaddr2): Started 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-2]
Slaves: [pod1-controller-0 pod1-controller-1]
```

```
ip-10.84.123.35 (ocf::heartbeat:IPaddr2): Started pod1-controller-1
openstack-cinder-volume (systemd:openstack-cinder-volume): Started pod1-controller-2
my-ipmilan-for-controller-0 (stonith:fence_ipmilan): Started pod1-controller-0
my-ipmilan-for-controller-1 (stonith:fence_ipmilan): Started pod1-controller-0
my-ipmilan-for-controller-2 (stonith:fence_ipmilan): Started pod1-controller-0
```

Daemon Status:

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

Put the cluster in maintenance mode.

```
[heat-admin@pod1-controller-0 ~]$ sudo pcs cluster standby
```

```
[heat-admin@pod1-controller-0 ~]$ sudo pcs status
```

```
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod1-controller-2 (version 1.1.15-11.e17_3.4-e174ec8) - partition with quorum
Last updated: Mon Dec 4 00:48:24 2017 Last change: Mon Dec 4 00:48:18 2017 by root via
crm_attribute on pod1-controller-0
```

3 nodes and 22 resources configured

Node pod1-controller-0: standby

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

Full list of resources:

```
ip-11.118.0.42 (ocf::heartbeat:IPaddr2): Started pod1-controller-1
ip-11.119.0.47 (ocf::heartbeat:IPaddr2): Started pod1-controller-2
ip-11.120.0.49 (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-1 pod1-controller-2]
Stopped: [pod1-controller-0]
Master/Slave Set: galera-master [galera]
Masters: [pod1-controller-1 pod1-controller-2]
Slaves: [pod1-controller-0]
ip-11.120.0.47 (ocf::heartbeat:IPaddr2): Started 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-2]
Slaves: [pod1-controller-1]
Stopped: [pod1-controller-0]
ip-10.84.123.35 (ocf::heartbeat:IPaddr2): Started pod1-controller-1
```

```
openstack-cinder-volume (systemd:openstack-cinder-volume): Started pod1-controller-2
my-ipmilan-for-controller-0 (stonith:fence_ipmilan): Started pod1-controller-1
my-ipmilan-for-controller-1 (stonith:fence_ipmilan): Started pod1-controller-1
my-ipmilan-for-controller-2 (stonith:fence_ipmilan): Started pod1-controller-2
```

## Replace Motherboard

Step 1. The steps in order to replace the motherboard in a UCS C240 M4 server can be referred from:

[Cisco UCS C240 M4 Server Installation and Service Guide](#)

Step 2. Log in to the server with the use of the CIMC IP.



Step 3. Perform BIOS upgrade if the firmware is not as per the recommended version used previously. Steps for BIOS upgrade are given here:

[Cisco UCS C-Series Rack-Mount Server BIOS Upgrade Guide](#)

## Restore Cluster Status

Log in to the impacted controller, remove standby mode by setting **unstandby**. Verify controller comes Online with cluster and galera shows all three controllers as Master. This might take a few minutes.

```
[heat-admin@pod1-controller-0 ~]$ sudo pcs cluster unstandby

[heat-admin@pod1-controller-0 ~]$ sudo pcs status
Cluster name: tripleo_cluster
Stack: corosync
Current DC: pod1-controller-2 (version 1.1.15-11.e17_3.4-e174ec8) - partition with quorum
Last updated: Mon Dec 4 01:08:10 2017 Last change: Mon Dec 4 01:04:21 2017 by root via
crm_attribute on pod1-controller-0

3 nodes and 22 resources configured

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

Full list of resources:

ip-11.118.0.42 (ocf::heartbeat:IPaddr2): Started pod1-controller-1
ip-11.119.0.47 (ocf::heartbeat:IPaddr2): Started pod1-controller-2
ip-11.120.0.49 (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]
ip-11.120.0.47 (ocf::heartbeat:IPaddr2): Started 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-2]
Slaves: [pod1-controller-0 pod1-controller-1]
ip-10.84.123.35 (ocf::heartbeat:IPaddr2): Started pod1-controller-1
openstack-cinder-volume (systemd:openstack-cinder-volume): Started pod1-controller-2
my-ipmilan-for-controller-0 (stonith:fence_ipmilan): Started pod1-controller-1
my-ipmilan-for-controller-1 (stonith:fence_ipmilan): Started pod1-controller-1
my-ipmilan-for-controller-2 (stonith:fence_ipmilan): Started pod1-controller-2

Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enable
```