

Element Manager after Ultra 6.3 Release - Architectural Changes and How to Troubleshoot the EM Issues

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Abbreviations](#)

[Ultra Element Manager after Ultra 6.3 Release - Architectural Changes](#)

[Number of UEM VM Instances is Configurable as of 6.3 Release](#)

[ZooKeeper Distribution](#)

[Introduction of Keepalived for HA](#)

[Decouple SCM from NCS Process](#)

[EM Service Runs on Master Node Only](#)

[Steps to Troubleshoot Element Manager Related Issues](#)

[UEM Setup with 2 VM's - Process Startup Sequence and Log Location](#)

[UEM Setup with 3 VM's - Process Startup Sequence and Log Location](#)

[Summary of UEM Processes](#)

Introduction

This document describes the changes in the Element Manager (EM) architecture that are introduced as part of the 6.3 UltraM release.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- STAROs
- Ultra-M base architecture

Components Used

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

Prior to Ultra 6.3 release, for Ultra Element Manager to work there was a requirement to create 3 UEM VMs. The 3rd one was not in use and was there to help form ZooKeeper cluster. As of the 6.3 release, this design has changed.

Abbreviations

Abbreviations used in this article:

VNF	Virtual Network Function
CF	Control Function
SF	Service Function
ESC	Elastic Service Controller
VIM	Virtual Infrastructure Manager
VM	Virtual Machine
EM	Element Manager
UAS	Ultra Automation Services
UUID	Universally Unique Identifier
ZK	Zoo Keeper

Ultra Element Manager after Ultra 6.3 Release - Architectural Changes

This document describes these 5 changes that are introduced as part of the 6.3 UltraM release:

Number of UEM VM Instances is Configurable as of 6.3 Release

Prior to 6.3 release, 3 UEM VM were mandatory. You could see that with **nova list** after the core tenant file is sourced:

```
[root@POD]# openstack server list --all
+-----+-----+-----+-----+-----+
| ID | Name | Status | Networks | Image Name |
+-----+-----+-----+-----+-----+
.....
| fae2d54a-96c7-4199-a412-155e6c029082 | vpc-LAASmme-em-3 | ACTIVE | orch=192.168.12.53;
mgmt=192.168.11.53 | ultra-em |
| c89a3716-9028-4835-9237-759166b5b7fb | vpc-LAASmme-em-2 | ACTIVE | orch=192.168.12.52;
mgmt=192.168.11.52 | ultra-em |
| 5f8cda2c-657a-4ba1-850c-805518e4bc18 | vpc-LAASmme-em-1 | ACTIVE | orch=192.168.12.51;
mgmt=192.168.11.51 | ultra-em |
```

This configuration snapshot (from vnf.conf file) was used:

```
vnfc em
health-check enabled
health-check probe-frequency 10
health-check probe-max-miss 6
```

```
health-check retry-count 6
health-check recovery-type restart-then-redeploy
health-check boot-time 300
vdu vdu-id em
  number-of-instances 1 --> HERE, this value was previously ignored in pre 6.3 releases
connection-point eth0
...
```

Regardless of the number of instances specified in this command, the number of spun VM's was always 3. In other words, the **number-of-instances** value was ignored.

As of 6.3, this gets changed - **the configured value can be 2 or 3.**

When you configure 2, the 2 UEM VMs are created.

When you configure 3, the 3 UEM VMs are created.

```
vnfc em
health-check enabled
health-check probe-frequency 10
health-check probe-max-miss 6
health-check retry-count 3
health-check recovery-type restart
health-check boot-time 300
vdu vdu-id vdu-em
vdu image ultra-em
vdu flavor em-flavor
  number-of-instances 2 --> HERE
connection-point eth0
....
```

This configuration would result in 2 VM's as seen with **nova list**.

```
vnfc em
health-check enabled
health-check probe-frequency 10
health-check probe-max-miss 6
health-check retry-count 3
health-check recovery-type restart
health-check boot-time 300
vdu vdu-id vdu-em
vdu image ultra-em
vdu flavor em-flavor
  number-of-instances 2 --> HERE
connection-point eth0
....
```

Note however that 3 IP address requirement remained the same. That is, in the **EM** part of the config (vnf.conf file) the 3 IP address are still **mandatory**:

```
vnfc em
health-check enabled
health-check probe-frequency 10
health-check probe-max-miss 6
health-check retry-count 3
health-check recovery-type restart
health-check boot-time 300
```

```

vdu vdu-id vdu-em
vdu image ultra-em
vdu flavor em-flavor
number-of-instances 2 ---> NOTE NUMBER OF INSTANCES is 2
connection-point eth0
virtual-link service-vl orch
virtual-link fixed-ip 172.x.y.51 --> IP #1
!
virtual-link fixed-ip 172.x.y.52 --> IP #2
!
virtual-link fixed-ip 172.x.y.53 --> IP #3
!

```

This is needed for ZK to work 3 instances of ZK are required. Every instance requires an IP address. Even though the 3rd instance is not effectively used, the 3rd IP is allocated to the 3rd, so-called **Arbiter** ZK instance (see *Diff.2* for more explanation).

What effect this has on the orchestration network?

There is always going to be 3 ports created in the orchestration network (to bind the 3 mentioned IP addresses).

```

[root@POD# neutron port-list | grep -em_

| 02d6f499-b060-469a-b691-ef51ed047d8c | vpc-LAASmme-em_vpc-LA_0_70de6820-9a86-4569-b069-
46f89b9e2856 | fa:16:3e:a4:9a:49 | {"subnet_id": "bf5dea3d-cd2f-4503-a32d-5345486d66dc",
"ip_address": "192.168.12.52"} |
| 0edcb464-cd7a-44bb-b6d6-07688a6c130d | vpc-LAASmme-em_vpc-LA_0_2694b73a-412b-4103-aac2-
4be2c284932c | fa:16:3e:80:eb:2f | {"subnet_id": "bf5dea3d-cd2f-4503-a32d-5345486d66dc",
"ip_address": "192.168.12.51"} |
| 9123f1a8-b3ea-4198-9ea3-1f89f45dfe74 | vpc-LAASmme-em_vpc-LA_0_49ada683-a5ce-4166-aeb5-
3316fe1427ea | fa:16:3e:5c:17:d6 | {"subnet_id": "bf5dea3d-cd2f-4503-a32d-5345486d66dc",
"ip_address": "192.168.12.53"} |

```

ZooKeeper Distribution

Prior to 6.3 ZK was used to form the cluster, hence this requirement is for 3rd VM. That requirement has not changed. However, for the setups where 2 UEM VM's are used, a 3rd ZK instance is hosted on the same set of VM's:

Prior to 6.3 and after 6.3 in a setup with 3 UEM VMs:

UEM VM1: hosting Zk instance 1

UEM VM2: hosting Zk instance 2

UEM VM3: hosting Zk instance 3

In 6.3 and later where 2 VMs only:

UEM VM1: hosting Zk instance 1 & Zk instance 3

UEM VM2: hosting Zk instance 2

UEM VM3: not exist

See picture 1. at the bottom of this article for detailed graphical representation.

Useful Zk commands:

To see Zk mode (leader/follower):

```
/opt/cisco/usp/packages/zookeeper/current/bin/zkServer.sh status
ZooKeeper JMX enabled by default
Using config: /opt/cisco/usp/packages/zookeeper/current/bin/../conf/zoo.cfg
Mode: leader
```

To check if Zk is running:

```
echo stat | nc IP_ADDRESS 2181
```

How to find the Ip address of Zk instance:

```
Run 'ip addr' from EM
In the /opt/cisco/em/config/ip.txt there are all the 3IP's
From vnf.conf file
From 'nova list' look for orchestration IP
For 2 EM's the arbiter IP can be found also in /opt/cisco/em/config/proxy-params.txt
```

How to check status of the Zk instance:

```
echo stat | nc 192.168.12.51 2181 | grep Mode
Mode: follower
```

You can run this command from one Zk for all other Zk instances (even they are on different VM)!

To connect to the Zk cli - now must use the IP (rather than localhost earlier):

```
/opt/cisco/usp/packages/zookeeper/current/bin/zkCli.sh -server <ip>:2181
```

You can use same command to connect to other Zk instances (even they are on different VM)!

Some useful command you can run once you connect to ZkCli:

```
ls /config/vdus/control-function
ls /config/element-manager
ls /
ls /log
ls /stat
get /config/vdus/session-function/BOOTxx
```

Introduction of Keepalived for HA

With the previous releases, ZK leader election framework used to determine master EM. That is not the case anymore as Cisco has moved to the keepalived framework.

What is keepalived and how it works?

Keepalived is [Linux based software](#) used for load balancing and high-availability to Linux system and Linux based infrastructures.

It is already used in ESC for HA.

In EM, Keepalived is used to decouple the NCS from Zk cluster state.

Keepalived process runs only on the first two instances of the EM and would determine the master

state for NCS process.

To check if the keepalived process is running:

```
ps -aef | grep keepalived  
(must return the process ID)
```

Why change?

In an earlier implementation, the (NCS/SCM) master node selection was closely integrated with Zk cluster status (the first instance to lock on the /em in Zk database was elected master). This creates problems when Zk loses connectivity with the cluster.

Keepalived is used to maintain Active/Standby UEM Cluster on VM basis.

NCS maintains the configuration data.
Zookeeper maintains the operational data.

Decouple SCM from NCS Process

In releases prior to 6.3, SCM component was bundled with NCS. That means when the NCS started, the SCM started as well (as consequence). In this release, this is now decoupled and SCM is a separate process for itself.

Commands to check the NCS and SCM services & processes.

To be executed from the ubuntu command line

```
ps -aef | grep ncs  
ps -aef | grep scm  
  
sudo service show ncs  
sudo service scm status
```

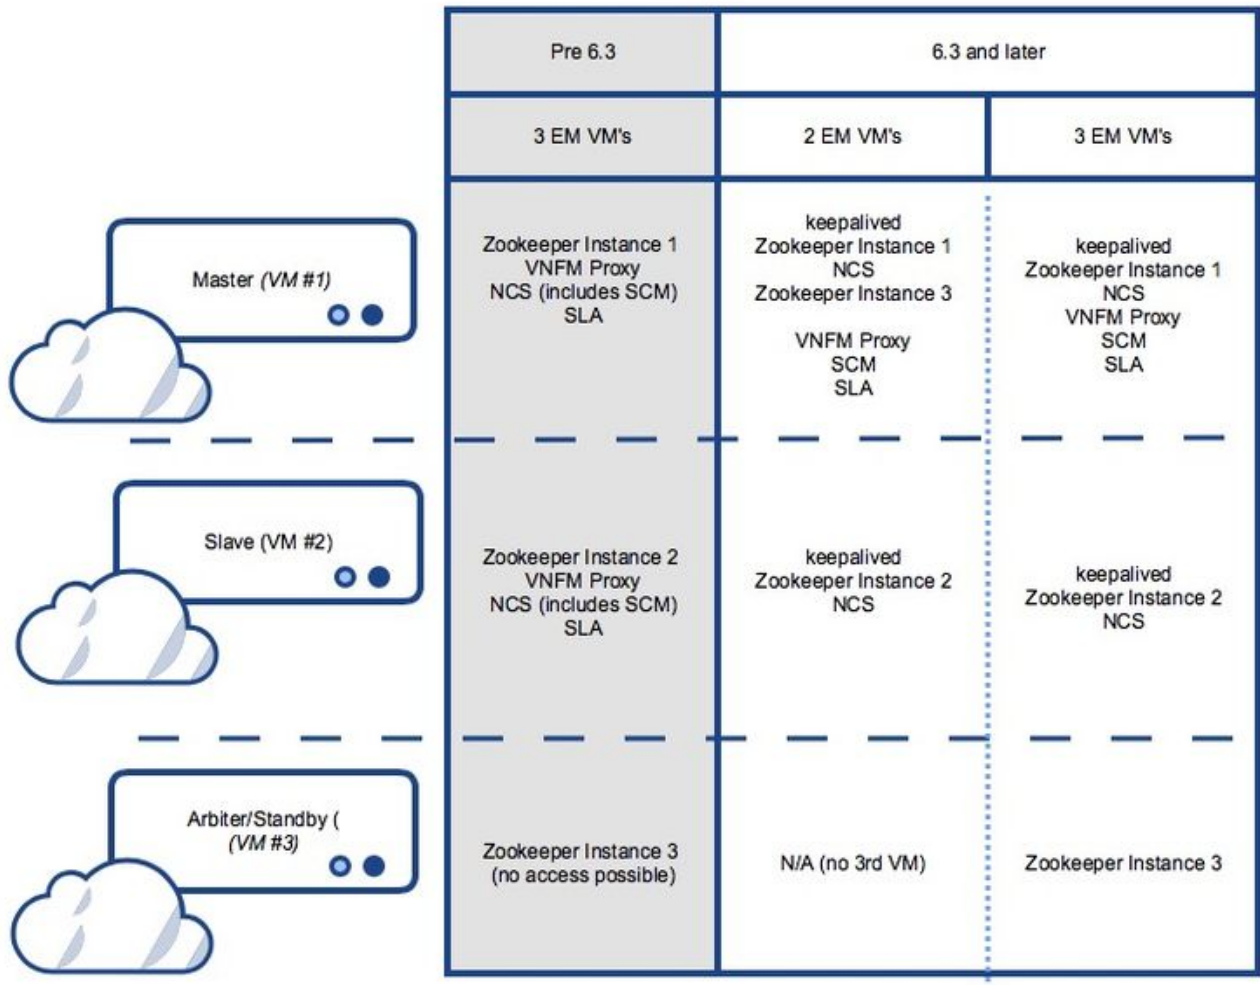
EM Service Runs on Master Node Only

Prior to 6.3 release, UEM services runs on both Master/Slave. As of 6.3 services runs on the master node only. This would impact the output displayed in **show ems**. As of 6.3, it is expected to see only one (master) node with this command, once logged in to the UEM CLI:

```
root@vpc-em-2:/var/log# sudo -i  
root@vpc-em-2:~# ncs_cli -u admin -C  
  
admin connected from 127.0.0.1 using console on vpc-LAASmme-em-2  
admin@scm# show ems  
EM VNFM  
ID SLA SCM PROXY VERSION  
-----  
52 UP UP UP 6.3.0 ==> HERE Only one EM instance is seen. In previous releases you were able to  
see 2 instances.  
Effectively all services would run on the master node, with exception of the NCS, and that is due to  
NCS requirements.
```

This image shows the summary of the possible services and VM distribution for Ultra Element

Manager



Steps to Troubleshoot Element Manager Related Issues

During bootup, this is the startup sequence:

UEM Setup with 2 VM's - Process Startup Sequence and Log Location

Master UEM:

- keepalived
- Zookeeper
- NCS
- Arbiter (3rd) instance of Zookeeper
- VNFM-Proxy
- SCM
- SLA

Slave UEM:

- keepalived
- Zookeeper
- NCS

UEM Setup with 3 VM's - Process Startup Sequence and Log Location

Master UEM:

- keepalived
- Zookeeper
- NCS
- VNFM-Proxy
- SCM
- SLA

Slave UEM:

- keepalived
- Zookeeper
- NCS

3rd UEM:

- Zookeeper

Summary of UEM Processes

This is the summary of UEM processes that you have to run.

You check status with **ps -aef | grep xx**

keepalived
arbiter
scm
sla
zoo.cfg
ncs

You can check status with **service xx status**, where xx:

zookeeper-arbiter
proxy
scm
sla
zk
ncs