

CEE Graceful Shutdown and Power-On

Cloud Execution Environment

OPERATING INSTRUCTIONS

Copyright

© Ericsson AB 2017, 2018. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

All trademarks mentioned herein are the property of their respective owners. These are shown in the document Trademark Information.



Contents

1	Introduction	1
1.1	Scope	1
1.2	Prerequisites	1
2	Procedure	4
2.1	Preparation	4
2.2	Identify Compute Hosts	5
2.3	Graceful Shutdown of CEE Region	6
2.4	Graceful Shutdown of External Storage	11
2.5	Graceful Shutdown of Switches, Routers, DC-GWs	11
2.6	Removal of Power Feed	12
3	Power-On of CEE Region	12
4	Power-On of CEE Compute Servers	15
5	Power-On of External Storage	16
6	Starting of Atlas VM	17
7	CEE Health Check Procedure	17
8	Starting of Tenant VMs	17
	Reference List	18





1 Introduction

This document provides instructions on how to gracefully shut down and power on the Cloud Execution Environment (CEE) region.

1.1 Scope

This document describes how to gracefully shut down and power on the CEE region. This procedure has been tested and verified on Dell configurations.

1.2 Prerequisites

This section provides information on the documents, tools, and conditions that apply to the procedure.

1.2.1 Documents

Before starting this procedure, ensure that the following documents have been read and understood:

- Personal Health and Safety Information
- System Safety Information

1.2.2 Tools

The following tools are needed:

- An Electrostatic Discharge (ESD) wrist strap (part number LYB 250 01/14)
- A computer capable of Secure Shell (SSH) logon to the virtual Cloud Infrastructure Controller (vCIC)
- RJ45 cable to connect the kickstart server to the control switch

1.2.3 Data

A site-specific IP and VLAN plan, based on the document IP and VLAN plan, Reference [1], is required. The address variables used throughout this document are summarized in Table 1.



Table 1 IP Variables Based on the IP and VLAN plan

VLAN	Variable Name	Factory Default IP Address Allocation
cee_ctrl_sp	<traffic_switch_a_static_ip>	192.168.2.2
cee_ctrl_sp	<traffic_switch_b_static_ip>	192.168.2.3
fuel_ctrl_sp	<fuel_static_ip>	192.168.0.11

Other site-specific data is listed in Table 2.

Table 2 Other Variables Used

Resource	Variable Name	Additional Information
Out-of-band (OOB) management IP address for the physical servers	<oob_management_ip_addresses_for_the_enclosure>	Use the primary IP address of the relevant enclosure.
External IP address of the vCIC	<vcic_address>	
Personal username to the vCIC	<personal_user>	
Password for the used personal user on the vCIC		
Password for the root user on vFuel		
Password to the OOB management device for user Administrator		Use the password of the relevant enclosure. This password can be found on the pull-out stickers attached to the OOB management device. The OOB management device is indicated by the green Active indicator.
A username with administrator privileges to the OOB management interface of the enclosure to be used at the GUI	<username>	
Password for user <username>		



Table 2 Other Variables Used

Resource	Variable Name	Additional Information
Port number to be used on the traffic switches for the new host	<port_number_on_traffic_switch>	The port number must be the same on the A and B switches. Use the hardware installation and installation instruction documents to determine the relevant port number for each new host: Extreme X460 Configuration, Extreme X670V Configuration, Extreme X770 Configuration.
Password for user network_admin to log on to the traffic switches	<password_for_network_admin>	
Name for each host	<host_name>	Host names are specified by the following scheme: compute-<enclosure_number>-<slot_number>

1.2.4 Conditions

Before starting this procedure, ensure that the following conditions are met:

- A work order for the graceful shutdown of CEE is received or the document is referred from another procedure.
- The IP addresses and credentials (root access) for SSH connections to the following devices are known. See also section Section 1.2.3 on page 1.
 - vCIC node
 - vFuel node
 - Atlas VM
 - Compute servers
 - OOB management interface (for example, iLO or iDRAC)
- All keys to the site are available and site access is granted.



2 Procedure

This procedure describes how to gracefully shut down and power on the CEE system.

The procedure contains the following activities:

1. CEE health check procedure, see Section 2.1.1 on page 4
2. CEE components backup, see Section 2.1.2 on page 4
3. Identifying vFuel master and vCIC hosts, see Section 2.2 on page 5
4. Graceful shutdown of CEE region, see Section 2.3 on page 6
5. Graceful shutdown of external storage, see Section 2.4 on page 11
6. Graceful shutdown of switches, routers, DC-GWs, see Section 2.5 on page 11
7. Power-on of the CEE region and other infrastructure hardware, see Section 3 on page 12
8. Restore of CEE software components from the backup as needed
9. Health check of CEE hardware and software components, see Section 7 on page 17

2.1 Preparation

2.1.1 CEE Health Check Procedure

Before starting the graceful shutdown of CEE, it is strongly recommended to perform a health check of the CEE region following the procedure detailed in the [Health Check Procedure](#).

2.1.2 CEE Components Backup

Do the following:

1. Back up the following entities and export all those backups to an external location according to the guides:
 - [CIC Domain Data Backup](#)
 - [Atlas Backup](#)
 - [Fuel Synchronization](#)



Note: For a full list of backup and restore documents in CEE and their use cases, refer to the Overview section of the [Disaster Recovery](#) document.

2. Save the configuration of site switches and DC-GWs in case these switches and routers have to be shut down.

2.2 Identify Compute Hosts

Before starting the graceful shutdown procedure, the following compute hosts must be identified:

- Compute server hosting the primary vCIC
- Compute servers hosting the secondary vCICs
- Compute server hosting the vFuel master

Before starting the graceful shutdown procedure, the primary vCIC and the compute servers hosting vCICs and/or the vFuel master must be identified.

Note: Make a note of the vFuel master and the vCIC hosts, as the information is needed in both the graceful shutdown and power on procedures.

2.2.1 Identifying Primary vCIC

To identify the primary vCIC:

1. Log on to any one of the vCICs using root user through SSH.
2. Issue the following command to identify the primary vCIC:
`crm_mon -1 -rf | grep -i current`

In the example below, cic-2 is the primary vCIC, and cic-1 and cic-3 are secondary vCICs:

```
root@cic-1:~# crm_mon -1 -rf |grep -i current
Current DC: cic-2.domain.tld (version 1.1.14-70404
b0) - partition with quorum
```

3. Identify the compute server hosting the primary vCIC, see Section 2.2.2 on page 5.

Note: Make a note of the primary vCIC, as the information is needed in both the graceful shutdown and the power on procedures.

2.2.2 Identify Compute Servers Hosting vCIC and/or vFuel Master

To identify which compute servers host the vFuel master and each vCIC, respectively:

1. Log on to each compute host using root user through SSH.
2. Issue the following `virsh` command:



```
virsh list --all
```

```
root@compute-0-2:~# virsh list --all
```

Id	Name	State
1	fuel_master	running
2	cic-1_vm	running

Example 1 Compute Hosting vCIC and vFuel Master

Note: Make a note of the vFuel master and the vCIC hosts, as the information is needed in both the graceful shutdown and the power on procedures.

2.3 Graceful Shutdown of CEE Region

Do the following:

1. Log on to any one of the vCICs using root user through SSH.

```
ssh ceeadm@<vcic_public_ip_address>
ceedm@cic-1:~$ sudo -i
```

2. Stop the CM-HA service.

Stop the CM-HA service from any one of the vCICs.

```
root@cic-1:~# crm resource status p_cmha
root@cic-1:~# crm resource stop p_cmha
```

3. Stop all the Nova-managed VMs.

- a. Stop the tenant VMs.

Stop all the tenant VMs using the nova command for all tenants.

```
root@cic-1:~# nova list --all-tenants
root@cic-1:~# nova stop <uuid>
```

Note: Stop the VMs using nova and **not** virsh destroy or virsh shutdown.

After the command has been executed, all the VMs are in SHUTOFF status.

- b. Stop the Atlas VM.

```
root@cic-1:~# nova list --all-tenants
root@cic-1:~# nova stop <uuid>
```

Stop the Atlas VM using the nova command.



Note: Stop the VMs using nova and **not** virsh destroy or virsh shutdown.

After the command has been executed, the Atlas VM is in SHUTOFF status.

4. Power off the compute servers only hosting tenant VMs and the Atlas VM.

Power off (shut down) the compute servers hosting tenant VMs and the Atlas VM.

Note: Do not shut down the compute servers that host vCIC or vFuel. See Section 2.2 on page 5.

Log on to the compute host using SSH from vFuel or access the server using the serial console through the OOB management interface.

shutdown -h now

Note: If the shutdown of the compute hangs at some point and the server never shuts down, the server needs to be shut down through the OOB management interface. It is recommended to have console connection towards the server to monitor the shutdown process. If it hangs (that is, does not shut down after 10-15 minutes), shut down the server through the CLI or GUI of the OOB management.

5. If vFuel is migrated into the CEE region, proceed with Step 6.

If vFuel is placed outside the CEE region, proceed with Step 7.

6. If vFuel is migrated into the CEE region, shut down the vCIC VMs and the compute server hosting them.

- a. Shut down each secondary vCIC VM.

Connect to the compute server hosting the secondary vCIC and shut down the VM with virsh. Perform this step on both compute servers hosting secondary vCIC VMs.

```
root@compute-0-2:~# virsh list --all
```

Id	Name	State
1	fuel_master	running
2	cic-1_vm	running

```
root@compute-0-2:~# virsh shutdown <instance-id/name>
```

- b. Shut down the compute servers hosting the secondary vCIC VMs, as identified in Section 2.2 on page 5.

Wait a few seconds to make sure that the vCIC VM has been shut down.



Perform the shutdown of the compute server with `shutdown`.

Perform this step on both compute servers hosting secondary vCIC VMs.

Note: Do not shut down the server if the vFuel master is running on the server. See Section 2.2 on page 5.

```
root@compute-0-3:~# virsh list --all
```

Id	Name	State
-	cic-3_vm	shut off

```
root@compute-0-3:~# shutdown -h now
```

- c. Shut down the primary vCIC VM, as identified in Section 2.2 on page 5.

Connect to the compute server hosting the primary vCIC and shutdown the VM with `virsh`.

```
root@compute-0-1:~# virsh list --all
```

Id	Name	State
2	cic-1_vm	running

```
root@compute-0-1:~# virsh shutdown <instance-id/name>
```

- d. Shut down the compute server hosting the primary vCIC VM.

Note: Make a note of which compute server is hosting the primary vCIC VM. This vCIC has to be started first during power-on.

Wait a few seconds to make sure that the vCIC VM has been shut down.

Perform the shutdown of the compute server with `shutdown`.

Note: Do not shut down the server if the vFuel master is running on the server. See Section 2.2 on page 5.

```
root@compute-0-1:~# virsh list --all
```

Id	Name	State
-	cic-1_vm	shut off

```
root@compute-0-1:~# shutdown -h now
```

- e. Shut down the vFuel master VM, as identified in Section 2.2 on page 5.



Note: Make a note of which compute server is hosting the vFuel master VM. This compute host has to be started first during power-on.

Wait a few seconds to make sure that all the vCIC VMs have been shut down.

Check the status of all the nodes and vCIC VMs from vFuel with `fuel node`.

Connect to the compute server hosting the vFuel master, and shut down the vFuel master VM with `virsh`.

At this stage you are logged out from the shell prompt.

```
root@compute-0-2:~# virsh list --all
```

Id	Name	State
1	fuel_master	running
-	cic-1_vm	shut off

```
root@compute-0-2:~# virsh shutdown <instance-id/name>
```

- f. Shut down the compute server hosting the vCIC VM and vFuel master.

Since the vFuel was shut down in Step e, it is no longer possible to reach the compute server from vFuel.

Log on to the computer server using the serial console through the OOB management interface. Power off the server by using `shutdown`.

```
root@compute-0-2:~# virsh list --all
```

Id	Name	State
-	fuel_master	shut off
-	cic-1_vm	shut off

```
root@compute-0-2:~# shutdown -h now
```

- g. Continue with Section 2.4 on page 11.

7. If vFuel is placed outside the CEE region, shut down the vCIC VMs and the compute server hosting them. Repeat this step for both secondary vCICs.

- a. Shut down the secondary vCIC VM.

Connect to the compute server hosting the secondary vCIC and shut down the VM with `virsh`.

```
root@compute-0-2:~# virsh list --all
```



```
Id      Name                                     State
-----
2      cic-1_vm                                running
```

```
root@compute-0-2:~# virsh shutdown <instance-id/name>
```

- b. Shut down the compute server hosting the secondary vCIC VM.

Wait a few seconds to make sure that the vCIC VM has been shut down.

Perform the shutdown of the compute server with shutdown.

```
root@compute-0-2:~# virsh list --all
```

```
Id      Name                                     State
-----
-      cic-1_vm                                shut off
```

```
root@compute-0-2:~# shutdown -h now
```

- c. Shut down the primary vCIC VM.

Connect to the compute server hosting the primary vCIC and shutdown the VM with virsh.

```
root@compute-0-1:~# virsh list --all
```

```
Id      Name                                     State
-----
2      cic-1_vm                                running
```

```
root@compute-0-1:~# virsh shutdown <instance-id/name>
```

- d. Shut down the compute server hosting the primary vCIC VM.

Wait a few seconds to make sure that the vCIC VM has been shut down.

Perform the shutdown of the compute server with shutdown.

```
root@compute-0-1:~# virsh list --all
```

```
Id      Name                                     State
-----
-      cic-1_vm                                shut off
```

```
root@compute-0-1:~# shutdown -h now
```

- e. Shut down the vFuel VM.



Wait a few seconds to make sure that all the vCIC VMs have been shut down.

Check the status of all the nodes and vCIC VMs from vFuel with `fuel node`.

Connect to the kickstart server hosting the vFuel and shut down the vFuel VM with `virsh`.

```
root@fuelhost:~# virsh list --all
```

Id	Name	State
1	fuel_vm	running

```
root@fuelhost:~# virsh shutdown <instance-id/name>
```

- f. Shut down the kickstart server hosting the vFuel.

Note: This step is optional and to be performed if required.

Log on to the kickstart server hosting the vFuel and shut down the server with `shutdown`.

```
root@fuelhost:~# virsh list --all
```

Id	Name	State
-	fuel_master	shut off

```
root@fuelhost:~# shutdown -h now
```

2.4 Graceful Shutdown of External Storage

The software and hardware management of external storage is outside the scope of the CEE documentation. Refer to the relevant product documentation regarding the power off and shutdown of the external storage.

2.5 Graceful Shutdown of Switches, Routers, DC-GWs

The software and hardware management of external switches, routers, DC-GWs is outside the scope of the CEE documentation. Refer to the relevant product documentation regarding the power off and shutdown of the external switches, routers, DC-GWs.

It is strongly recommended to make a backup of the complete switch and router configuration.



2.6 Removal of Power Feed

The removal of the power feed to servers, enclosures, and switches is optional and to be done if required.

3 Power-On of CEE Region

Do the following:

1. Restore the power feed to the switches, routers, DC-GWs.

Note: It is recommended to restore the power feed individually for the servers that will be powered on. Many servers have the option to start as soon as the power feed is connected, however, it can lead to uncontrolled power up of all the servers at once.

2. Power on the kickstart or compute server hosting the vFuel.

The vFuel VM has to be powered on first.

If vFuel is migrated into the CEE region, proceed with Step a.

If vFuel is placed outside the CEE region, proceed with Step d.

- a. Power on the compute server hosting the vFuel master VM.

Log on to the compute server hosting the vFuel master identified in Section 2.2.2 on page 5. Access the compute server through the OOB management interface.

Power on the server and wait until the server has been booted successfully. Once the server has been booted successfully, the vFuel VM is launched and booted automatically. If there are any vCIC hosted on the same server, the vCIC VM is also launched and booted automatically.

Wait enough time for the VMs to be operational.

```
root@compute-0-2:~# virsh list --all
```

Id	Name	State
1	fuel_master	running
2	cic-1_vm	running

- Log on to the kickstart server hosting vFuel.

If the kickstart server is already powered on, start vFuel VM with `virsh`.

```
root@fuelhost:~# virsh list --all
```

```
root@fuelhost:~# virsh start <instance name>
```

- ```
[root@fuel ~]# fuel node
id | status | name | cluster | ip | mac | roles | =>
 pending_roles | online | group_id |
+-----+-----+-----+-----+-----+-----+-----+-----+
4 | ready | compute-0-2 | 1 | 192.168.0.23 | ec:f4:bb:cd:45:0c | compute, virt | =>
 | | | 1 | | | | |
1 | ready | compute-0-1 | 1 | 192.168.0.22 | ec:f4:bb:cd:42:98 | compute, virt | =>
 | | | 1 | | | | |
5 | ready | compute-0-5 | 1 | 192.168.0.26 | ec:f4:bb:cd:45:18 | compute | =>
 | | | 1 | | | | |
8 | ready | cic-3 | 1 | 192.168.0.27 | 56:bd:11:f2:cd:42 | controller, mongo | =>
 | | | 1 | | | | |
2 | ready | compute-0-4 | 1 | 192.168.0.25 | ec:f4:bb:cd:45:30 | compute | =>
 | | | 1 | | | | |
3 | ready | compute-0-3 | 1 | 192.168.0.24 | ec:f4:bb:cd:45:50 | compute, virt | =>
 | | | 1 | | | | |
6 | ready | cic-1 | 1 | 192.168.0.29 | fa:30:2d:96:16:40 | controller, mongo | =>
 | | | 1 | | | | |
7 | ready | cic-2 | 1 | 192.168.0.28 | 5e:1e:4b:ae:db:4d | controller, mongo | =>
 | | | 1 | | | | |

[root@fuel ~]# fuel env
id | status | name | release_id |
+-----+-----+-----+-----+
1 | operational | DC201 | 2 |
[root@fuel ~]# fuel-utils check_all
checking with command "systemctl is-active nailgun"
active
```



```

checking with command "! pgrep puppet"
nailgun is ready.
checking with command "egrep -q ^[2-4][0-9]? < <(curl --connect-timeout 1 -s -w '%{http_code}' =>
http://192.168.0.11:8777/ostf/not_found -o /dev/null)"
checking with command "! pgrep puppet"
ostf is ready.
checking with command "ps waux | grep -q 'cobblerd -F' && pgrep dnsmasq"
21708
checking with command "cobbler profile find --name=ubuntu* | grep -q ubuntu && cobbler profile =>
find --name=*bootstrap* | grep -q bootstrap"
checking with command "! pgrep puppet"
cobbler is ready.
checking with command "curl -f -L -i -u "naily:pn3RTJwd21ErE6ausLZZ06cJ" =>
http://127.0.0.1:15672/api/nodes 1>/dev/null 2>&1"
checking with command "curl -f -L -u "mcollective:bC1A0q9dHn07s9q0yuerbE8e" -s =>
http://127.0.0.1:15672/api/exchanges | grep -qw 'mcollective_broadcast'"
checking with command "curl -f -L -u "mcollective:bC1A0q9dHn07s9q0yuerbE8e" -s =>
http://127.0.0.1:15672/api/exchanges | grep -qw 'mcollective_directed'"
checking with command "! pgrep puppet"
rabbitmq is ready.
checking with command "PGPASSWORD=rRyo4SRiemptF6fmwjZ051is /usr/bin/psql -h =>
192.168.0.11 -U "nailgun" "nailgun" -c '\copyright' 2>&1 1>/dev/null"
checking with command "! pgrep puppet"
postgres is ready.
checking with command "ps waux | grep -q 'astuted'"
checking with command "curl -f -L -u "naily:pn3RTJwd21ErE6ausLZZ06cJ" -s =>
http://127.0.0.1:15672/api/exchanges | grep -qw 'nailgun'"
checking with command "curl -f -L -u "naily:pn3RTJwd21ErE6ausLZZ06cJ" -s =>
http://127.0.0.1:15672/api/exchanges | grep -qw 'naily_service'"
checking with command "! pgrep puppet"
astute is ready.
checking with command "ps waux | grep -q mcollectived"
checking with command "! pgrep puppet"
mcollective is ready.
checking with command "ps waux | grep -q nginx"
checking with command "! pgrep puppet"
nginx is ready.
checking with command "keystone --os-auth-url "http://192.168.0.11:35357/v2.0" =>
--os-username "nailgun" --os-password "ha8Nj4yzTTEXChLxbrMrCWKP" token-get &>/dev/null"
checking with command "! pgrep puppet"
keystone is ready.
checking with command "netstat -nl | grep -q 514"
checking with command "! pgrep puppet"
rsyslog is ready.
checking with command "netstat -ntl | grep -q 873"
checking with command "! pgrep puppet"
rsync is ready.
[root@fuel ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/mapper/os-root 9.5G 5.3G 3.8G 59% /
devtmpfs 2.9G 0 2.9G 0% /dev
tmpfs 2.9G 0 2.9G 0% /dev/shm
tmpfs 2.9G 305M 2.6G 11% /run
tmpfs 2.9G 0 2.9G 0% /sys/fs/cgroup
/dev/mapper/os-var 20G 7.0G 12G 38% /var
/dev/mapper/os-varlog 8.3G 1.2G 6.7G 16% /var/log
/dev/vda3 197M 113M 84M 58% /boot
/dev/vda2 200M 0 200M 0% /boot/efi
tmpfs 578M 0 578M 0% /run/user/0
tmpfs 578M 0 578M 0% /run/user/1100
tmpfs 578M 0 578M 0% /run/user/163
[root@fuel ~]#

```

## Example 2 Performing Health Check on vFuel VM



## 4 Power-On of CEE Compute Servers

Do the following:

1. Power on the compute servers running vCICs.

**Note:** The compute servers have to be powered on in sequential order and enough time interval has to be kept between the power-on of the servers.

- a. Power on the compute server running the primary vCIC.

Log on to the compute server hosting the primary vCIC identified in Section 2.2.1 on page 5. Access the compute server through the OOB management interface.

Power on the server and wait until the server has been booted successfully. Once the server has been booted successfully, the vCIC VM is launched and booted automatically.

Wait approximately 10 minutes for the vCIC to be operational. Use `virsh list` to check if the vCIC VM is operational.

**Note:** Skip this step if the compute server has already been powered on as part of powering on the vFuel VM, see Step a in Section 3 on page 12.

```
root@compute-0-2:~# virsh list --all
```

| Id | Name        | State   |
|----|-------------|---------|
| 1  | fuel_master | running |
| 2  | cic-1_vm    | running |

- b. Power on the compute server running the secondary vCIC.

Access the compute server through the OOB management interface.

Power on the server and wait until the server has been booted successfully. Once the server has been booted successfully, the vCIC VM is launched and booted automatically.

Wait approximately 10 minutes for the vCIC to be operational.

```
root@compute-0-3:~# virsh list --all
```

| Id | Name     | State   |
|----|----------|---------|
| 1  | cic-2_vm | running |



- c. Power on the compute server running the other secondary vCIC.

Access the compute server through the OOB management interface.

Power on the server and wait until the server has been booted successfully. Once the server has been booted successfully, the vCIC VM is launched and booted automatically.

Wait approximately 10 minutes for the vCIC to be operational.

```
root@compute-0-4:~# virsh list --all
```

| Id | Name     | State   |
|----|----------|---------|
| 1  | cic-3_vm | running |

- d. Start the p\_cmha service.

Once all the vCIC VMs have been booted and are in operation, start the p\_cmha resource on any one of the vCICs.

```
root@cic-1:~# crm resource status p_cmha
root@cic-1:~# crm resource start p_cmha
root@cic-1:~# crm resource status p_cmha
resource p_cmha is running on: cic-3.dc196.ericsson.se
```

- e. Perform health check of the vCICs.

Check and verify that all the crm services are operational on the vCICs.

Execute the following commands from any of the vCIC.

```
root@cic-1:~# crm_mon -lrf
```

2. Power on the remaining compute servers.

Access the compute server through the OOB management interface.

Power on the server and wait until the server has been booted successfully. It is recommended to power on the servers in sequential order.

## 5 Power-On of External Storage

If the CEE region is configured to have external storage (EMC ScaleIO), power on the external storage by following the relevant product documentation.



## 6 Starting of Atlas VM

Log on to any one of the vCICs and launch the Atlas VM with nova:

```
root@cic-1:~# nova list --all-tenants
root@cic-1:~# nova start <atlas_vm_name_or_uuid>
```

Once the VM is booted successfully, log on to the Atlas VM and check the status of services with systemctl.

**Note:** It can take some time for all the services to be operational.

```
atlasadm@atlas:~$ sudo -i
[sudo] password for atlasadm:
root@atlas:~# systemctl
```

## 7 CEE Health Check Procedure

Once all the compute servers are powered on, it is strongly recommended to perform health check of the CEE region following the procedure detailed in the [Health Check Procedure](#).

## 8 Starting of Tenant VMs

After completing the health check procedure, you can launch all the tenant VMs. Consider the booting order of the tenant VMs. Log on to any one of the vCICs and launch the tenant VMs with nova.

**Note:** Do not use any other command.

```
root@cic-1:~# nova list --all-tenants
root@cic-1:~# nova start <vm_name_or_uuid>
```



## Reference List

- [1] IP and VLAN plan, 2/102 62-CRA 119 1862/5