

# SW Installation in Multi-Server Deployment

Cloud Execution Environment

## INSTALLATION INSTRUCTIONS

**Copyright**

© Ericsson AB 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

|                       |  |           |
|-----------------------|--|-----------|
| <b>1</b>              | <b>Introduction</b>  | <b>1</b>  |
| 1.1                   | Prerequisites  | 1         |
| 1.2                   | Time Required  | 3         |
| <b>2</b>              | <b>Pre-Installation Steps</b>  | <b>3</b>  |
| 2.1                   | PCI Passthrough BIOS Configuration   | 3         |
| 2.2                   | CEE Reinstallation Fails with BSP  | 4         |
| <b>3</b>              | <b>Temporary pre-Installation Steps</b>  | <b>4</b>  |
| 3.1                   | ScaleIO Backend Networks Are Not Configured on CEE-Managed Extreme Traffic Switch                      | 5         |
| 3.2                   | Audit Log Contains Several Unfiltered CM-HA Related Events   | 5         |
| <b>4</b>              | <b>Install CEE Software in Server System</b>   | <b>6</b>  |
| <b>5</b>              | <b>Temporary Installation Steps</b>  | <b>8</b>  |
| <b>6</b>              | <b>Post-Installation Steps</b>   | <b>8</b>  |
| 6.1                   | vcpu_pin_set Parameter in /etc/nova/nova.conf Is Set to Empty when Assigning CPU Values in config.yaml | 8         |
| 6.2                   | SR-IOV Traffic Switch Ports Not Configured   | 9         |
| <b>7</b>              | <b>Temporary post-Installation Steps</b>   | <b>10</b> |
| <b>8</b>              | <b>Migrate vFuel to CEE Region</b>   | <b>11</b> |
| <b>9</b>              | <b>Post-Installation Activities</b>  | <b>12</b> |
| 9.1                   | Isolating Compute Hosts According to VirtIO Queue Size   | 13        |
| 9.2                   | Identify UID and GID for cinder User on vCICs  | 17        |
| <b>10</b>             | <b>Error Handling</b>  | <b>18</b> |
| <b>Appendix</b>       |  |           |
| <b>11</b>             | <b>CA and NBI Certificates for Secure HTTPS Access</b>   | <b>20</b> |
| <b>Reference List</b> |  | <b>22</b> |





# 1 Introduction

This document is part of the installation flow for the Cloud Execution Environment (CEE) multi-server deployment and describes how to install CEE software in a CEE region. Complete this procedure when directed here from [CEE Installation](#):

1. Start the procedure in [CEE Installation](#).
2. Continue with this document when directed here from [CEE Installation](#).
3. Return to [CEE Installation](#) and carry out the remaining steps.

For the complete installation flow, refer to section [Installation Flow in CEE Installation](#).

This instruction assumes that a kickstart server is used. For the installation and testing of the kickstart server, refer to [Preparation of Kickstart Server](#).

## 1.1 Prerequisites

This section describes the prerequisites that must be fulfilled before CEE software can be installed.

### 1.1.1 Documents

Activities in the following documents must be performed before the steps in this instruction are performed:

- Relevant sections of [CEE Installation](#)
- [Preparation of Kickstart Server](#)
- [Configuration File Guide](#)
- Limitations in [Limitations and Workarounds for Cloud Execution Environment \(CEE\)](#), Reference [3], must be known and considered

### 1.1.2 Hardware and Software Required

Before starting this procedure, make sure that the following software and hardware is available:

- CEE software release tarball is available, for more information refer to the [Configuration File Guide](#).

The required software can be downloaded from the SW Gateway. If you have problems with the download procedure, contact the next level of support.



— In case of using an unmanaged ScaleIO storage solution:

- ScaleIO server must be already installed and connected to the Data Center (DC).
- Check that the ScaleIO Data Client (SDC) kernel module version is compatible with the software versions of the Meta Data Manager (MDM) cluster and the ScaleIO Data Servers (SDSs). For issues with software compatibility, contact the next level of support.

For more information about the configuration of unmanaged ScaleIO in CEE, refer to the [Configuration File Guide](#).

The recommended installation method described in this document is using a kickstart server with Linux OS. For more information, refer to [Preparation of Kickstart Server](#).

### 1.1.3

#### Tools

The following hardware tools are required:

- Kickstart server, refer to [Preparation of Kickstart Server](#)
- In case of BSP hardware: two Ethernet cables (DensiShield-RJ45, TSR 491 603/x) to connect the kickstart server to the system to be installed. In case of Dell hardware: one RJ45-RJ45 cat6 Ethernet cable to connect the kickstart server to the system to be installed.

### 1.1.4

#### Installation Data

The following data is needed:

Table 1 Installation Data

| Data Type    | Description  |
|--------------|--|
| Passwords    | Initial vFuel server root user password is <code>r00tme</code> (used for installation only)  |
| Certificates | <p>Certificates for the vCIC and Atlas Northbound Interfaces (NBIs), see Section 11 on page 20</p> <p>Obtain certificates for the NBIs from an authorized Certification Authority (CA) before starting the installation process, because after installation it is <b>not</b> possible to replace certificates with ones issued on a different domain name.</p> |



| Data Type    | Description   |
|--------------|---|
| yaml files   | Site-specific config.yaml in /mnt/cee_config, refer to Preparation of Kickstart Server and Configuration File Guide |
|              | Site-specific switching configuration file, refer to Configuration File Guide                                       |
|              | Cabling scheme file (hardware-specific from CEE_RELEASE/cabling_scheme/)  |
|              | Neutron configuration file (hardware and configuration specific, see CEE_RELEASE/neutron/)                          |
|              | Host networking configuration file (hardware-specific file from CEE_RELEASE/host_net_templates/)                    |
| IP addresses | The local version of IP and VLAN Plan updated with customer and site-specific values                                |
|              | IP address for the kickstart server   |
|              | IP addresses for vFuel in networks fuel_ctrl_sp and subrack_ctrl_sp, refer to the site-specific IP and VLAN Plan    |

## 1.2 Time Required

The expected execution time for the installation procedure is around four hours, in case all prerequisites are available and depending on the hardware used.

# 2 Pre-Installation Steps

This section describes steps that must be executed in this release prior to the actual installation.

## 2.1 PCI Passthrough BIOS Configuration

**Note:** This prerequisite is applicable for HDS hardware platforms.

In case of HDS hardware, the BIOS must be configured manually to use PCI passthrough. If CEE installation includes PCI passthrough for the compute nodes, SR-IOV must be Disabled in BIOS on the affected interfaces on HDS platform before CEE installation. The same virtualization mode must be used on each port or NIC of a compute server. Therefore on compute hosts configured for PCI Passthrough, SR-IOV support must be disabled for the entire compute host. The BIOS configuration of the HDS servers must be done by the data center owner.



For more information refer to HDS documentation Hyperscale Datacenter System 8000 Customer Documentation, Reference [2].

**The PCI passthrough feature can only be used with compatible NICs, that is, on NICs where the SR-IOV virtualization mode can be disabled.** Intel Fortville NICs are **not** compatible (XXV710-AM1, XXV710-AM2, XL710-AM1, XL710-AM2, XL710-BM1, XL710-BM2, X710-AM2, X710-BM2, XL710-QDA1, XL710-QDA2, X710-DA2, and X710-DA4).

**Note:** It is not possible to configure both SR-IOV and PCI passthrough on the same compute host.

## 2.2 CEE Reinstallation Fails with BSP

**Note:** This workaround is applicable for BSP hardware platforms.

CEE installation checks if a BSP tenant exists with the same name as stated in the config.yaml. If it already exists, the CEE installation removes the existing configuration of that BSP tenant.

When reinstalling CEE with BSP R10 or later, the removal of the CEE tenant fails if any IPV4 VLAN interface on the router belongs to the CEE tenant.

The installcee script fails with the following error message:

```
installcee.sh.error: Failed to configure BSP HW: "Failed to delete tenant "CEE". is it really configured in BSP?  
Traceback (most recent call last):  
File "/usr/bin/setup_cmx", line 9, in <module>  
load_entry_point('ericsson-cloud==1.9.3', 'console_scripts', 'setup_cmx')()  
File "build/bdist.linux-x86_64/egg/ericsson_cloud/cmx/setup_cmx.py", line 64, in main  
File "build/bdist.linux-x86_64/egg/ericsson_cloud/cmx/bsp_module.py", line 1901, in remove_cee_config  
File "build/bdist.linux-x86_64/egg/ericsson_cloud/cmx/bsp_module.py", line 1969, in remove_tenant  
ValueError: Failed to delete tenant "CEE". is it really configured in BSP?"  
installcee.sh.info: A background task exited with error. Waiting for any other jobs to complete before exiting  
installcee.sh.error: Failed to configure switches and/or external storage
```

Workaround: Remove all IPv4 VLAN interfaces on the router that belong to the CEE tenant before reinstalling CEE.

## 3 Temporary pre-Installation Steps

This section describes temporary pre-installation workarounds that are needed for this release. Carry out these workarounds before starting the installation.





### 3.1 ScaleIO Backend Networks Are Not Configured on CEE-Managed Extreme Traffic Switch

**Note:** This prerequisite is applicable for HPE and Dell hardware platforms.

The ScaleIO backend networks `sio_be_san_pda` and `sio_be_san_pdb` are not configured on CEE-managed Extreme traffic switches, unless they have mixed traffic and storage switch functions. As a result, CEE deployment fails.

Associated trouble report: HW66117

Workaround: Configure the ScaleIO frontend and backend networks in the `config.yaml` to use iSCSI, as seen below:

```
ericsson:
  storage:
    scaleio:
      frontend_networks: ['iscsi-left', 'iscsi-right']
      backend_networks: ['iscsi-left', 'iscsi-right']
```

### 3.2 Audit Log Contains Several Unfiltered CM-HA Related Events

**Note:** This prerequisite is applicable for HPE, Dell, BSP, single server, and HDS hardware platforms.

Excessive audit logging is triggered when CM-HA logs to the infrastructure nodes, because all program executions during shell initialization are logged, not only the session start / end events. The information in these logs is not useful, and therefore not intended for the audit trail.

Associated trouble report: HW74686.

Workaround: Do the following:

Before CEE deployment, adjust the audit configuration template `/var/www/nailgun/plugins/ericsson_logging-1.0/deployment_scripts/puppet/modules/ericsson_audit_logging/templates/auditd/audit.rules.erb` on vFuel:

1. Insert the below lines before the line that begins with `# Monitoring for all` :

```
-a exit,never -F auid=1100 -F arch=b64 -S execve
```

```
-a exit,never -F auid=1100 -F arch=b32 -S execve
```

This will exclude auditing program executions for the CM-HA user having UID of 1100 on all CEE systems.



## 4 Install CEE Software in Server System

This section describes how to install CEE in the server system.

1. Log on to the kickstart server.
2. In case of using BSP hardware, connect to BSP from the kickstart server and check the BSP software release and BSP backups.

COM CLI interface is available via E-DBG Ethernet interface using the IP address defined as `lct_ip` in `config.yaml`.

Example:

```
ssh advanced@10.0.10.2 -p 2024
(password for advanced user)
show ManagedElement=1,SystemFunctions=1,SwInventory=1,active
show ManagedElement=1,SystemFunctions=1,BrM=1,⇒
BrmBackupManager=1,BrmBackup=1
exit
```

The value of `SWVERSION` cannot be lower than the software version indicated in CEE on BSP, Reference [1]. The name of backup (backupname) must contain the suffix specified in `cmx_switch.yaml` so backup name is suggested to be `<cee_region_name>_<yyyymmdd>_preCEETenant`. Mismatch between the suffix in the backup name and in `cmx_switch.yaml` causes the CEE installation script to exit with error.

3. Check that vFuel is running in the kickstart server:

```
virsh list --all
```

Examples:

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

| Id    | Name        | State   |
|-------|-------------|---------|
| ----- |             |         |
| 2     | fuel_master | running |

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

| Id | Name | State |
|----|------|-------|
|----|------|-------|



```
-----
-      fuel_master          shut off
```

In case vFuel is in shut off state, start vFuel and wait until booting is complete:

```
virsh start fuel_master
```

4. Log on to vFuel using SSH. For changing the password, refer to the [System Hardening Guideline](#).

```
ssh root@<vfuel-ip-address-in-network-fuel_ctrl_sp>
```

5. Verify that the correct time zone, time and date have been set by using the below command:

```
date
```

6. Change the working directory to /opt/ecs-fuel-utils with following command:

```
cd /opt/ecs-fuel-utils
```

7. Set up a screen session to ensure that the installation process is not interrupted:

```
# screen -S installcee -L
```

If the connection to vFuel is lost, log on to vFuel again and reattach the screen session with the below command:

```
# screen -r installcee
```

**Note:** The nohup option can cause installation failure and must not be used.

8. Execute the following:

```
./installcee.sh
```

The time required for command execution is approximately four hours.

Check that the printout is the following:

```
Ericsson CEE installed successfully
```



## 5 Temporary Installation Steps

There are no temporary installation procedures in the current release.

## 6 Post-Installation Steps

This section describes steps that must be executed in this release after a successful installation. Carry out these workarounds after the installation script has run successfully.

### 6.1 `vcpu_pin_set` Parameter in `/etc/nova/nova.conf` Is Set to Empty when Assigning CPU Values in `config.yaml`

**Note:** This workaround is applicable for Dell, BSP, and HDS hardware platforms.

When all CPUs have been manually assigned in the `config.yaml`, Nova configuration receives the parameter `vcpu_pin_set` as empty, meaning that there are no free vCPUs available at OpenStack user level. In such a scenario Nova interprets that all CPUs on the compute host, including OVS and host OS, are free to use. As a result, there is no CPU pinning on these compute hosts and VMs allocated on them share the vCPUs with the infrastructure. This scenario is mostly applicable to GEP5 boards where both vCIC and vFuel are co-hosted. It is recommended to disable nova-compute service in these hosts so no tenant VMs are allocated on these hosts.

Associated trouble report: HV63386

Workaround: Perform the below procedure.

1. Collect the nodes where there are no CPUs left for Nova, using the below command:  
  

```
nova hypervisor-show <compute_host_name> | grep vcpus_used
```
2. If the above command returns 0 for `vcpus_used`, disable the Nova service by executing the following command:  
  

```
nova service-disable <compute_host> <binary> --reason <reason>
```
3. Issue the command `nova service-list` and check if the status of these compute hosts is set to disabled.



## 6.2 SR-IOV Traffic Switch Ports Not Configured

**Note:** This workaround is applicable for Dell hardware platforms.

Dell R620 and R630 traffic switch ports are not configured when a VM is booted in a Neutron network used for SR-IOV. Even if CEE Neutron is installed with managed Extreme switches, the Modular Layer 2 (ML2) mechanism driver does not understand to cover SR-IOV based Neutron ports.

Associated trouble report: HV18706

Workaround:

To create SR-IOV network or networks, the following command can be used. The parameter values can be different from the ones used in this example:

```
neutron net-create --provider:physical_network=pool_0000_41_00_0 =>
--provider:network_type=vlan --provider:segmentation_id=3666 sriov_net1
```

Manual configuration of the Extreme switch can be done by using the following command. The value of `segmentation_id` must be used for the tag parameter.

```
create vlan <sriov_vlan_name>
configure vlan <sriov_vlan_name> tag 3666
configure vlan <sriov_vlan_name> add ports <sriov_port_1> =>
<sriov_port_2> tagged
```

The configuration of the traffic switch can be done when the CEE region has been installed, if the VLAN specifications are available at that time. It is suggested to use a VLAN tag for the Neutron network that is not included in the range used for the default physical network. If external connectivity must be configured as well, add the port or ports that connect to the DC-GW (BGW) as shown below:

```
configure vlan <sriov_vlan_name> add ports <sriov_port_1> =>
<sriov_port_2> <port_to_bgw> tagged
```

The configuration changed manually on the traffic switches can be saved on each traffic switch. However, if the consistency checker function is triggered for other reasons, it will restore the previous configuration that does not contain the SR-IOV settings. Therefore, the changed configuration also has to be uploaded to the TFTP server.

To save the configuration change and upload it to the TFTP server, perform the following for Switch A, then repeat for Switch B:

1. Save the configuration change on the traffic switch.

Example:



```
* DC203_SWA_X670V.6 # save configuration "ericsson_backup"
The configuration file ericsson_backup.cfg already exists.
Do you want to save configuration to ericsson_backup.cfg and overwrite it? (y/N) Yes
Saving configuration on master ..... done!
Configuration saved to ericsson_backup.cfg successfully.

The current selected default configuration database to boot up the system
(ericsson_active.cfg) is different than the one just saved (ericsson_backup.cfg).
Do you want to make ericsson_backup.cfg the default database? (y/N) No
Default configuration database selection cancelled.
DC203_SWA_X670V.7 #
```

## 2. Upload the configuration change to the TFTP server:

- a. Change permissions 0644 to 0666 for the ericsson\_backup.cfg in /var/lib/tftpboot/extreme\_conf.

Example:

```
[root@fuel extreme_conf]# cd /var/lib/tftpboot/extreme_conf/
[root@fuel extreme_conf]# ls -l
total 680
-rw-r--r-- 1 root root 345459 Dec 13 18:56 DC203_SWA_X670V_ericsson_backup.cfg
-rw-r--r-- 1 root root 345519 Dec 13 18:56 DC203_SWB_X670V_ericsson_backup.cfg
[root@fuel extreme_conf]#
[root@fuel extreme_conf]# chmod 0666 DC203_SWA_X670V_ericsson_backup.cfg
```

- b. Upload the ericsson\_backup.cfg from the traffic switch to the TFTP server:

```
tftp put <fuel_tftp_server> vr <vr_name> <file_path_on_
_switch> <tftp_file_path>
```

- c. Revert the permissions 0666 to 0644 for the ericsson\_backup.cfg on the TFTP server.

Example:

```
[root@fuel extreme_conf]# cd /var/lib/tftpboot/extreme_conf/
[root@fuel extreme_conf]# chmod 0644 DC203_SWA_X670V_ericsson_backup.cfg
```

## 7 Temporary post-Installation Steps

There are no temporary post-installation procedures in the current release.

Depending on the CEE release and configuration, workarounds can apply to newly deployed CEE regions. Verify that all relevant workarounds in Limitations



and Workarounds for Cloud Execution Environment (CEE), Reference [3] are considered and performed.

## 8 Migrate vFuel to CEE Region

This section describes how to migrate vFuel to a CEE region.

**Note:** This section is applicable for migrating vFuel in a Linux OS.

Perform the below steps:

1. Log on to the kickstart server.
2. Execute the below script:

```
CEE_RELEASE/scripts/migrate_fuel.sh
```

An example of the output is:



```
./migrate_fuel.sh
migrate_fuel.sh.info: Checking current Fuel state
migrate_fuel.sh.info: Preparing to migrate Fuel
migrate_fuel.sh.info: Fuel will be migrated to compute-0-4 (192.168.0.23)
migrate_fuel.sh.info: The vFuel image and the Domin XML will also be prepared =>
on compute-0-6 (192.168.0.25)
migrate_fuel.sh.info: Shutting down current Fuel
migrate_fuel.sh.info: Waiting for Fuel to complete shutdown
migrate_fuel.sh.info: Copying Fuel disk image to compute-0-4 (172.30.160.1)
sending incremental file list
fuel_br3160.qcow2
 68,730,224,640 100% 107.69MB/s   0:10:08 (xfr#1, to-chk=0/1)
migrate_fuel.sh.info: Copying Fuel disk image to compute-0-6 (172.30.160.2)
sending incremental file list
fuel_br3160.qcow2
 68,730,224,640 100% 103.17MB/s   0:10:35 (xfr#1, to-chk=0/1)
migrate_fuel.sh.info: Starting new vFuel inside CEE region on compute-0-4 (192.168.0.23)
migrate_fuel.sh.info: Waiting for new vFuel to start up
migrate_fuel.sh.info: Waiting for new vFuel to be ready
migrate_fuel.sh.info: New vFuel ready
migrate_fuel.sh.info: Performing post migrate actions
migrate_fuel.sh.info: Post migrate actions done
migrate_fuel.sh.info: Fuel is successfully migrated to compute-0-4 (192.168.0.23)
```

## 9 Post-Installation Activities

Execute the following steps after the installation:

1. Disconnect the kickstart server.
2. Verify the version of CEE by executing the command `cat /etc/cee_version.txt` on the Fuel master node.

The output has the following format:

```
RELEASE=CEE CXC1737883_4-<build_number>
NAME=Mitaka on Ubuntu 14.04
VERSION=R6-<r-state>-<specific_build_number>-9.0
```





An example of the output is the following:

```
[root@fuel ~]# cat /etc/cee_version.txt
```

```
RELEASE=CEE CXC1737883_4-4280
NAME=Mitaka on Ubuntu 14.04
VERSION=R6-R4A02-35547a3-9.0
```

```
[root@fuel ~]#
```

Verify the CEE version by comparing `build_number` and `r-state` to the Product Revision Information for Cloud Execution Environment (CEE), Reference [6].

3. After installation, there is an active NeLS Server Communication Problem alarm, because the NeLS server is not configured and not available.
  - a. To configure the connection to the NeLS server, follow the instructions in the [Runtime Configuration Guide](#).
  - b. If the alarm is not cleared, follow the instructions in the [NeLS Server Communication Problem alarm OPI](#).
4. For disaster recovery purposes, after deployment, the installation media must be backed up outside the CEE region. For more information, refer to the document [Disaster Recovery](#).
5. If the customized QEMU with increased VirtIO queue size was configured in `config.yaml`, isolate the compute hosts according to VirtIO queue size, see Section 9.1 on page 13.
6. If PCI Passthrough technology is used, see [OpenStack Compute API in CEE](#) for usage within OpenStack.
7. Continue with the relevant sections of [CEE Installation](#).

## 9.1 Isolating Compute Hosts According to VirtIO Queue Size

**Note:** This section applies if increased VirtIO queue size was configured on any compute host(s) in `config.yaml` before deployment. For information on customized QEMU with increased VirtIO queue size, refer to the [Multi-Server System Dimensioning Guide, CEE 6](#).

Compute hosts with customized QEMU must be separated from compute hosts with standard QEMU, to allow scheduling for these two sets of hosts in a controlled manner. Host aggregates can be used to isolate the compute hosts:

- Isolate the compute hosts with increased VirtIO queue size, see Section 9.1.1 on page 13.
- Isolate the compute hosts with standard VirtIO queue size, see Section 9.1.2 on page 15.



### 9.1.1 Isolate Compute Hosts with Increased VirtIO Queue Size

To create host-aggregates for compute hosts with the customized QEMU (virtio\_queue\_size=1024), follow the below procedure.

1. List the current host aggregates:

```
nova aggregate-list
```

An example of the printout:

```
root@cic-1:~# nova aggregate-list
+-----+-----+-----+
| Id | Name | Availability Zone |
+-----+-----+-----+
+-----+-----+-----+
```

2. Create a host aggregate for the hosts with customized QEMU:

```
nova aggregate-create increased_qemu
```

An example of the printout:

```
root@cic-1:~# nova aggregate-create increased_qemu
+-----+-----+-----+-----+-----+
| Id | Name | Availability Zone | Hosts | Metadata |
+-----+-----+-----+-----+-----+
| 71 | increased_qemu | - | | |
+-----+-----+-----+-----+-----+
```

3. Add the compute hosts with the customized QEMU to the host aggregate:

```
nova aggregate-add-host increased_qemu compute-0-1.domain.tld
```

An example of the printout:

```
root@cic-1:~# nova aggregate-add-host increased_qemu compute-0-1.domain.tld
Host compute-0-1.domain.tld has been successfully added for aggregate 71
+-----+-----+-----+-----+-----+
| Id | Name | Availability Zone | Hosts | Metadata |
+-----+-----+-----+-----+-----+
| 71 | increased_qemu | - | 'compute-0-1.domain.tld' | |
+-----+-----+-----+-----+-----+
```

4. Update the metadata of the host aggregate:

```
nova aggregate-set-metadata increased_qemu virtio_queue_size=1k
```

An example of the printout:



```
root@cic-1:~# nova aggregate-set-metadata increased_qemu virtio_queue_size=1k
Metadata has been successfully updated for aggregate 71.
```

| Id | Name           | Availability Zone | Hosts                    | Metadata               |
|----|----------------|-------------------|--------------------------|------------------------|
| 71 | increased_qemu | -                 | 'compute-0-1.domain.tld' | 'virtio_queue_size=1k' |

5. Create the extra specs in a flavor, to be used for the VM instantiation. The extra specs must have the namespace `aggregate_instance_extra_specs`.

An example is `aggregate_instance_extra_spec:virtio_queue_size=1k`.

```
root@cic-1:~# nova flavor-key qemu-flavor set aggregate_instance_extra_specs:virtio_queue_size=1k
```

**Note:** In case of using the GUI, create a server resource template using the extra specs, then create a dummy VM using this server resource template to create the corresponding flavor.

VMs must use flavors with extra spec `virtio_queue_size=1k` in order to be scheduled in compute hosts with customized QEMU.

### 9.1.2

### Isolate Compute Hosts with Standard QEMU

To create host aggregates for compute hosts with the standard QEMU (`virtio_queue_size=256`), follow the below procedure.

1. List the current host aggregates:

```
nova aggregate-list
```

An example of the printout:

```
root@cic-1:~# nova aggregate-list
```

| Id | Name           | Availability Zone |
|----|----------------|-------------------|
| 3  | increased_qemu | -                 |

2. Create a host aggregate for the hosts with standard QEMU:

```
nova aggregate-create default_qemu
```

An example of the printout:

```
root@cic-1:~# nova aggregate-create default_qemu
```

| Id | Name         | Availability Zone | Hosts | Metadata |
|----|--------------|-------------------|-------|----------|
| 9  | default_qemu | -                 |       |          |



```
+-----+-----+-----+-----+-----+-----+
```

3. Add the rest of the compute hosts with standard QEMU to the default\_qemu (in this example the computes that are not in customized QEMU host aggregate increased\_qemu).

```
nova aggregate-add-host default_qemu compute-0-2.domain.tld
nova aggregate-add-host default_qemu compute-0-3.domain.tld
nova aggregate-add-host default_qemu compute-0-4.domain.tld
```

An example of the printout:

```
+-----+-----+-----+-----+-----+-----+
| Id | Name          | Availability Zone | =>
+-----+-----+-----+-----+-----+-----+
| 9  | default_qemu  | -                | =>
+-----+-----+-----+-----+-----+-----+
```

```
+-----+-----+-----+-----+-----+-----+
| Hosts                                     | Metadata |
+-----+-----+-----+-----+-----+-----+
| 'compute-0-2.domain.tld', 'compute-0-3.domain.tld', 'compute-0-4.domain.tld' |           |
+-----+-----+-----+-----+-----+-----+
```

4. Update the metadata of the host aggregate:

```
nova aggregate-set-metadata default_qemu virtio_queue_size=256
```

An example of the printout:

```
root@cic-1:~# nova aggregate-set-metadata default_qemu virtio_queue_size=256
Metadata has been successfully updated for aggregate 9.
```

```
+-----+-----+-----+-----+-----+-----+
| Id | Name          | Availability Zone | =>
+-----+-----+-----+-----+-----+-----+
| 9  | default_qemu  | -                | =>
+-----+-----+-----+-----+-----+-----+
```

```
+-----+-----+-----+-----+-----+-----+
| Hosts                                     | Metadata |
+-----+-----+-----+-----+-----+-----+
| 'compute-0-2.domain.tld', 'compute-0-3.domain.tld', 'compute-0-4.domain.tld' | 'virtio_queue_size=256' |
+-----+-----+-----+-----+-----+-----+
```

5. Update the extra specs of a flavor to be used for the VM instantiation. The extra specs must have the namespace aggregate\_instance\_extra\_specs:

An example is aggregate\_instance\_extra\_spec:virtio\_queue\_size=256.

```
root@cic-1:~# nova flavor-key qemu-flavor set aggregate_instance_extra_specs:virtio_queue_size=256
```

**Note:** In case of using the GUI, create a server resource template using the extra specs, then create a dummy VM using this server resource template to create the corresponding flavor.

An example of the flavor:



```

root@cic-2:~# nova flavor-show medium_default
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Property | Value |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| OS-FLV-DISABLED:disabled | False |
| OS-FLV-EXT-DATA:ephemeral | 0 |
| disk | 20 |
| extra_specs | {"hw:cpu_policy": "dedicated", "aggregate_instance_extra_specs": |
| id | b104c60a-158b-412e-af51-1b951db437bf |
| name | medium_default |
| os-flavor-access:is_public | True |
| ram | 2048 |
| rxtx_factor | 1.0 |
| swap | |
| vcpus | 2 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

virtio_queue_size": "256", "hw:mem_page_size": "1048576"}
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

VMs must use flavors with extra spec `virtio_queue_size=256` in order to be scheduled in compute hosts with standard QEMU.

## 9.2 Identify UID and GID for cinder User on vCICs

**Note:** This section is only applicable if `cinder-backup` service is enabled and configured with an external NFS storage backend.

When the Cinder backup feature is enabled with an external NFS storage backend, the `cinder` user needs write access permission on the NFS storage server.

To identify the corresponding UID and GID of the `cinder` user, do the following:

1. Log on to one of the vCICs:

```
ssh <personal_user>@<cic_address>
```

2. Run the following commands:

— UID:

```
id -u cinder
```

— GID:

```
id -g cinder
```



Example output:

```
root@cic-1:~# id -u cinder
125
root@cic-1:~# id -g cinder
131
```

## 10 Error Handling

In case of any errors during the installation procedure, follow the below steps:

1. Check the console for failure messages or reference to any logs that possibly contain failure messages. Refer to the [Configuration File Guide](#) for the location of logs.
2. Fix possible problems.
3. Copy the original network templates to the `/mnt/cee_config` directory.

**Note:** If this step is missed, VLANs and interfaces from the previous run will be used, which causes the newer configuration options to be skipped.

On the vFuel node issue the following command:

```
cp CEE_RELEASE/host_net_templates/host_nw_*.yaml
/mnt/cee_config/
```

4. Rerun `installcee.sh` and collect logs:

```
./installcee.sh 2>&1 | tee <file_name>.log
```

**Note:** The `installcee` script does not automatically delete an existing CEE Region (Fuel environment), so installation attempts with an existing Fuel environment will fail. In this case reinstall CEE with the below command:

```
./installcee.sh --force
```

5. The following scenarios are possible:
  - The cause of failure is identified, fixed, or the install succeeds.  
In this case, exit this procedure.
  - Or the cause of failure is not identified, fixed, or the install still fails for presumably the same reason.



In this case, proceed to Step 6.

6. Perform data collection according to the Data Collection Guideline.
7. Contact the next level of support.

# Appendix

## 11 CA and NBI Certificates for Secure HTTPS Access

Certification Authority (CA) and Northbound Interface (NBI) certificates are required for secure HTTPS access to CEE.

Make sure to perform the following tasks before starting the installation process:

1. Choose a unique hostname for the vCIC NBI.
2. Choose a unique hostname for the Atlas NBI.
3. Obtain certificates for the NBIs from an authorized Certification Authority (CA).

**Note:** It is **not** possible to replace the certificates with ones issued on a different domain name after installation.

The following certificate files are needed:

- CA certificate (or chain of certificates) of the organization issuing the Atlas NBI
- CA certificate (or chain of certificates) of the organization issuing the vCIC NBI
- Atlas NBI certificate
- vCIC NBI certificate

**Note:** Atlas and vCIC certificates can be issued by the same CA, or by two separate CAs.

The Common Name (CN) and at least one DNS entry in the Subject Alternate Name (SAN) attribute must contain the publicly known hostname chosen for the NBI, so that the certificate refers to this publicly known hostname. The private key belonging to the certificate cannot be encrypted.

4. Concatenate the vCIC NBI certificate and private key into a single PEM format under `/mnt/cee_config` on vFuel. Perform the same for the Atlas NBI.

ASCII format is preferred for the individual certificates.





**Note:** The `pkcs12` binary format is commonly used. This output format contains multiple entities in a single binary file and uses encryption. Issue the below command to convert it to PEM format:

```
openssl pkcs12 -in <input_file> -out <output_file>
-nodes
```

`-nodes` is needed to save the private key in unencrypted format.

In case other binary formats need to be converted, refer to Reference [4] or Reference [5].

5. Update the `config.yaml` file with the necessary information. Refer to the [Configuration File Guide](#) for updating the publicly known hostname and other relevant options in the `config.yaml` file.
6. Update the DNS resolver to contain the hostname and IP address pairs for the NBI.



## Reference List

- [1] CEE on BSP, 1/1551-CNA 403 3045/2
- [2] Hyperscale Datacenter System 8000 Customer Documentation, 2/1551-LZN 901 5032
- [3] Limitations and Workarounds for Cloud Execution Environment (CEE) 6.6, 5/109 21-AZE 102 01/5-12
- [4] [SSL Support](https://support.ssl.com/Knowledgebase/Article/View/19/0/der-vs-crt-vs-cer-vs-pem-certificates-and-how-to-convert-them), <https://support.ssl.com/Knowledgebase/Article/View/19/0/der-vs-crt-vs-cer-vs-pem-certificates-and-how-to-convert-them>
- [5] [Thawte Licensing](https://search.thawte.com/support/ssl-digital-certificates/index?page=content&id=SO26449), <https://search.thawte.com/support/ssl-digital-certificates/index?page=content&id=SO26449>
- [6] Product Revision Information for Cloud Execution Environment (CEE) 6.6, 109 21-AZE 102 01/5-12 Uen