

Emergency Recovery Procedure

Cloud Execution Environment (CEE) 6

EMERGENCY RECOVERY

Abstract

The objective of this document is to act as collection of emergency recovery procedures for CEE 6 in the Ericsson Cloud System solution.

Copyright

© Ericsson 2017. 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.

Trademarks

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



Contents

1	Introduction	3
1.1	Scope	3
1.2	Target Group	3
2	Prerequisites	3
2.1	Hardware and Software Tools	3
2.2	System Backups	4
2.3	Data Collection	4
2.4	Conditions	4
2.5	Site Emergency Kit	4
3	Definitions.....	5
3.1	Abbreviations	5
4	Procedures for Specific Problem Types	5
4.1	Loss of Connectivity to the System	5
4.2	Loss of Connectivity due to TLS Certificate	6
4.3	vCIC Is Down	7
4.4	vCIC Is Unreachable via SSH.....	8
4.5	vFuel Is Down	9
4.6	MySQL Database Is Down	10
4.7	MongoDB Database Is Down	10
4.8	Glance Image Upload Failed	10
4.9	VM Instantiation Failed	11
4.10	Problem Type: SR-IOV Related Problems	11
4.11	Power Outage in CEE System	15
4.12	Neutron Agents Are Down	15
4.13	Nova Services Are Down.....	16
4.14	Cinder Services Are Down.....	16
4.15	Glance Image Not Shown	17
5	Contact Support.....	17
6	Change Information	18
7	References.....	19
Appendix 19		
8	Additional Information	19
8.1	How to Connect.....	20
8.2	List vCIC and Compute Nodes	21



1 Introduction

This document provides a systematic approach for resolving a system emergency experienced with Cloud Execution Environment (CEE) 6.

Typically, an emergency procedure is required for conditions that make communication or normal management with alarm handling impossible. In a worst case scenario, a procedure is required to restore the product.

The system is assumed to have been in a fully working state before the problem started. Therefore no troubleshooting procedures that relate to faulty configuration or incorrect software version or hardware version, or both, are explained.

Some steps that have been identified as risky from an In Service Performance (ISP) point of view are avoided in this document. When such steps are necessary, it is recommended to contact the next level of support, as described in section Contact Support.

1.1 Scope

This procedure is performed in order to recover the system from a major failure, or to restore the Operation and Maintenance (O&M) connection with the node if necessary, or both. This procedure is valid for the following product and release:

CEE R6

1.2 Target Group

The recovery actions described in this procedure are expected to be executed by Ericsson local or global support organizations, or both.

2 Prerequisites

This section states the prerequisites for performing the emergency recovery procedure.

2.1 Hardware and Software Tools

The following hardware and software tools are required:

1. Tools specified in SW Installation in Single Server Deployment, Reference [1] and SW Installation in Multi-Server Deployment, Reference [2]



2. Computer with Secure Shell (SSH) connection to the virtual Cloud Infrastructure Controller (vCIC)

2.2 System Backups

Some of the recovery actions include recovery from the previous known stable state. In order to facilitate such a recovery, make sure the following is available:

- 1 Fuel VM cold standby (refer to Fuel Synchronization, Reference [3])
- 2 Atlas backup (refer to Atlas Backup, Reference [4] and Atlas Restore, Reference [5])

2.3 Data Collection

Ensure that recent activities for O&M, software updates or feature introduction within CEE and the tenant VMs are known, and the data needed for performing the emergency recovery actions has been collected as described in the Data Collection Guideline, Reference [6].

The purpose of data collection is to secure that recovery actions can be executed, but it does not necessarily include the full set of troubleshooting data needed for identifying the root cause of the problem.

2.4 Conditions

Before starting the procedure collect the following information to identify possible scenarios and recovery procedures:

1. Information about CEE and Ericsson Cloud System (ECS) solution used, for example, product name, software version, platform, operating system and hardware type.
2. How to connect to CEE. For more information, refer to CEE Connectivity User Guide, Reference [7].
3. Access to the respective user guides of all the hardware components involved in the ECS solution.

2.5 Site Emergency Kit

For more information, contact the next level of support, as described in section Contact Support.



3 Definitions

Each of the procedures specified in this document contains an impact matrix that includes the following information:

Deployment Deployment scenario in which the problem is experienced. CEE 6 supports two deployment scenarios: single server deployment and multi-server deployment with High Availability (HA).

Hardware type Hardware type in which the problem is experienced.

Components Area of CEE functionality or component in which the problem is experienced.

Recovery actions intending to recover the system from the observed problem are clarified for the specified problem types.

3.1 Abbreviations

Refer to Glossary of Terms and Acronyms, Reference [8].

4 Procedures for Specific Problem Types

4.1 Loss of Connectivity to the System

Connectivity to CEE through its NBI (HA proxy public VIP or public vCIC IP addresses) is lost.

Impact matrix:

<i>Deployment</i>	Single server, multi-server
<i>Hardware type</i>	BSP, Dell,HDS
<i>Components</i>	All components in CEE, tenant traffic is probably not impacted

Recovery Action:

1. Check the traffic switches and make sure that the ports connected to the public interfaces are up and the ports are part of the `cee_om_sp` VLAN. If the problem is observed in the traffic switch, then fix it and exit this procedure. If there is no problem observed with the traffic switch, continue with the next step.



2. Connect to the compute host that hosts vCIC either through SSH connection or through serial connection.
3. Connect to the console of the vCIC:

```
virsh console cic-<x>_VM
```
4. Check the OVS port configuration in the vCIC to ensure that the host networking setup of vCIC is as designed:

```
sudo ovs-vsctl show
```
5. After login, get the IP of the bridge by executing the following command:

```
ifconfig br-ex
```

Ping the failed node from a vCIC which is still alive. While sending ping requests, run the following command multiple times on the compute node where the failed node is running. Check if the RX/TX counters are increasing:

```
ovs-appctl dpctl/show -s netdev@ovs-netdev | grep vhostuser-vcic-<x> -A4
```

If the counters are increasing, proceed to step 6.

If the counters are not increasing, restart the virtual switch on the compute node by issuing the following command:

```
service openvswitch-switch restart
```

6. In case there are missing ports in OVS configuration, or if the problem still persists, collect troubleshooting data as described in the Data Collection Guideline, Reference [6]. Contact the next level of maintenance support, as described in section Contact Support.

4.2 Loss of Connectivity due to TLS Certificate

Connectivity to CEE through its NBI (to HA proxy public VIP or public vCIC IP addresses) is lost due to corrupt or expired Transport Layer Security (TLS) certificate. For more information refer to Expiring Certificate, Reference [9].

Impact matrix:

Deployment	Single server, multi-server
Hardware type	BSP, Dell,HDS
Components	All components in CEE, tenant traffic is probably not impacted

Recovery Action:

1. Confirm that the problem observed is due to the TLS certificate used.



2. Obtain correct TLS certificates and store it in the directory `/mnt/cee_config/` on vFuel.

Note: In single server deployments vFuel must be available for this recovery action.
3. If required, update `/mnt/cee_config/config.yaml` with the appropriate certificate file names.
4. Install the certificate in the vFuel node by executing the below command in the Fuel master node:

```
fuel node --node <node_ids> --tasks eri_cert_copy_master
eri_copy_controller_certificates --force
```

Where `<node_ids>` is the list of vCIC IDs from the `fuel2 node list` printout.
5. If the problem persists, contact the next level of maintenance support, as described in section Contact Support.

4.3 vCIC Is Down

One or more vCICs are observed to be down.

Impact matrix:

Deployment	Single server, multi-server
Hardware type	BSP, Dell,HDS
Components	All components in CEE, tenant traffic is probably not impacted

Recovery Action:

1. Connect to the compute host that hosts vCIC either through SSH connection or through serial connection.
2. Make sure that the vCIC is defined in the Libvirt of the compute host:

```
virsh list --all
```
3. Attempt to start the vCIC manually:

```
virsh start cic-<x>_VM
```
4. If the vCIC does not start, contact the next level of maintenance support, as described in section Contact Support.
5. If the CIC VM can be started, connect to the vCIC through SSH.

Note: In case vCIC starts up, but SSH connection is not possible, proceed with section Loss of Connectivity to the System.
6. Check the status of the vCIC in the cluster:

```
crm_mon -l -rf
```



Note: It will take a few minutes for all the resources to come up in the cluster.

7. If vCIC is in `online` status in the cluster, but the resources did not join the cluster, check the status of the Corosync and Pacemaker services:

```
service corosync status
service pacemaker status
service corosync-notifyd status
```
8. If any of the Corosync or Pacemaker services are not running, start those services:

```
service <service_name> start
```
9. Check the cluster status and monitor whether the resources are starting up in the cluster:

```
crm_mon -l -rf
```
10. If the problem persists, contact the next level of maintenance support, as described in section Contact Support.

4.4 vCIC Is Unreachable via SSH

One or more vCICs are observed to be unreachable with SSH. vCIC is pingable and accessible via console but console session is unresponsive.

Impact matrix:

Deployment	Single server, multi-server
Hardware type	BSP, Dell,HDS
Components	All components in CEE, tenant traffic is probably not impacted

Recovery Action:

1. Connect to the compute host that hosts vCIC either through SSH connection or through serial connection.
2. Make sure that the vCIC is defined in the Libvirt of the compute host:

```
virsh list --all
```
3. Destroy the CIC VM manually:

```
virsh destroy cic-<x>_VM
```
4. Start the CIC VM manually:

```
virsh start cic-<x>_VM
```
5. If the CIC VM doesn't start, contact the next level of maintenance support, as described in section Contact Support.
6. If the CIC VM can be started, connect to the vCIC through SSH.



Note: In case vCIC starts up but SSH connection is not possible, proceed with section Loss of Connectivity to the System.

7. Check the status of the vCIC in the cluster:

```
crm_mon -l -rf
```

Note: It will take few minutes for all the resources to come up in the cluster.

8. If vCIC is in `online` status in the cluster, but the resources did not join the cluster, check the status of the Corosync and Pacemaker services:

```
service corosync status
service pacemaker status
service corosync-notifyd status
```

9. If any of the Corosync or Pacemaker services are not running, start those services:

```
service <service_name> start
```

10. Check the cluster status and monitor whether the resources are starting up in the cluster:

```
crm_mon -l -rf
```

11. If the problem persists, contact the next level of maintenance support, as described in section Contact Support.

4.5 vFuel Is Down

vFuel instance reports running status but is inaccessible via SSH after a power event on the compute host hosting vFuel.

Impact matrix:

Deployment	Single server, multi-server
Hardware type	BSP, Dell,HDS
Components	All components in CEE, tenant traffic is probably not impacted

Recovery Action:

1. Connect to the compute host that hosts vFuel either through SSH connection or through serial connection.
2. Check that vFuel is present and in running state on the compute host:

```
virsh list --all
```
3. Log into vFuel console:

```
virsh console fuel_master
```
4. If the vFuel console message is “if you would like to access maintenance mode enter root password, or hit Ctrl + D to exit”; choose to exit, and the normal boot process starts.



5. If the vFuel boots normally, connect to vFuel through SSH.

Note: In case vFuel boots up but SSH connection is not possible, proceed with section Loss of Connectivity to the System.

6. If the problem persists, contact the next level of maintenance support, as described in section Contact Support.

4.6 MySQL Database Is Down

MySQL database is observed to be down in one or more vCICs.

Impact matrix:

Deployment	Single server, multi-server
Hardware type	BSP, Dell,HDS
Components	OpenStack components in CEE, Watchmen, CM-HA

Recovery Action:

Contact the next level of maintenance support, as described in section Contact Support.

4.7 MongoDB Database Is Down

MongoDB database is observed to be down in one or more vCICs.

Impact matrix:

Deployment	Single server, multi-server
Hardware type	BSP, Dell,HDS
Components	Ceilometer

Recovery Action:

Contact the next level of maintenance support, as described in section Contact Support.

4.8 Glance Image Upload Failed

One or more vCICs are observed to be down.

Impact matrix:

Deployment	Single server, multi-server
-------------------	-----------------------------



Hardware type	BSP, Dell, HDS
Components	Glance, Swift

Recovery Action:

Contact the next level of maintenance support, as described in section Contact Support.

4.9 VM Instantiation Failed

One or more vCIC is observed to be down.

Impact matrix:

Deployment	Single server, multi-server
Hardware type	BSP, Dell, HDS
Components	All components in CEE, tenant traffic is probably not impacted

Recovery Action:

Contact the next level of maintenance support, as described in section Contact Support.

4.10 Problem Type: SR-IOV Related Problems

Basic Data

Symptoms:	<ul style="list-style-type: none"> - SR-IOV traffic is not working - SR-IOV Ethernet interfaces in the compute node are operational - Corresponding tenant VMs are in <i>active and running state</i> - All compute hosts are operational
-----------	---

Prerequisites and Additional Data Needed:

1. Access to compute hosts
2. Access to vCIC hosts
3. Access to traffic switches



This method can be used if SR-IOV tenant traffic is not working.

Recovery Action: Restore SR-IOV Traffic

Severity Level:	3 (severe)
Network Impact:	Ethernet packet tracing might be required. Possible changes in traffic switches, if required.
Risks:	No risk known
Execution time:	30 minutes
Expected Outcome:	SR-IOV traffic for tenants is restored

1. Make sure that the SR-IOV interfaces have been passed to the VMs successfully.

- a. Check VM details in the vCIC host:

```
nova show <vm_name_or_vm_uuid>
```

Note down the fields `OS-EXT-SRV-ATTR:host`, `OS-EXT-SRV-ATTR:instance_name` and `flavor` from the answer printout, as it will be used in the subsequent steps.

- b. Determine the number of SR-IOV interfaces requested for the VM by checking the flavor key used for the VM:

```
nova flavor-show <flavor_name>
```

`flavor_name` is derived from the answer printout of step 12. In the answer printout of `nova flavor-show` check the `extra_specs` field and count the number of Virtual Functions (VFs) requested for the VMs. This can be done by summing up the `pci_passthrough:alias` devices requested.

For example, the below flavor indicates that there are two VFs requested from alias `pool_83_00_0` and two VFs from alias `pool_83_00_1`. So in total 4 VFs are requested for the VM.

```
nova flavor-show m1.sriov
```

Property	value
OS-FLV-DISABLED:disabled	False
OS-FLV-EXT-DATA:ephemeral	200
disk	300
extra_specs	{"pci_passthrough:alias":
pool_83_00_0:2,pool_83_00_1:2"	}
id	6
name	m1.sriov



os-flavor-access:is_public	True
ram	4096
rxtx_factor	1.0
swap	
vcpus	4
-----+-----	
-----+	

- c. Connect to the vFuel node and check the `config.yaml` to determine whether the compute host identified in the field `os-EXT-SRV-ATTR:host` from the answer printout in step 12 has SR-IOV enabled:

```
cat /mnt/cee_config/config.yaml
```

The compute host element in the `config.yaml` must have the following SR-IOV configurations:

```
sriov:
  devices:
    - pci_address: "83:00.0"
      bandwidth: 10000000
    - pci_address: "83:00.1"
      bandwidth: 10000000
  vf : 8
```

Note down the PCI address of the physical interface used for SR-IOV in the compute host.

- d. Connect to the compute host (which can be identified in the field `os-EXT-SRV-ATTR:host` from the answer printout in step 12 where the VM is hosted and check the SR-IOV VF and physical function (PF) details:

```
dpdk_nic_bind.py -status
nic_bind.sh -l
```

Ensure that the following conditions are met:

- The PFs SR-IOV PCI address identified in step 13 is bound to the IXGBE driver
- There are corresponding VFs from the PFs
- The VFs are bound to the VFIO driver

Note down the interface device name of the PFs.

- e. Make sure that the VM is hosted by the hypervisor in the compute host using the following Libvirt command:

```
virsh list -all
```

VM `os-EXT-SRV-ATTR:instance_name` identified in step 12 must be running in the compute host.



- f. Check the XML file for the VM generated by Libvirt in the compute host:

```
virsh dumpxml <instance_name>
```

In the XML file of the VM there must be an “n” number of `hostdev` elements, where “n” corresponds to the number of SR-IOV interfaces requested for the VM that is identified in step 12 .

```
<hostdev mode='subsystem' type='pci' managed='yes'>
  <driver name='vfio' />
  <source>
    <address domain='0x0000' bus='0x83' slot='0x10' function='0x6' />
  </source>
  <alias name='hostdev2' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x06'
function='0x0' />
</hostdev>
```

Address within the `<source>` element represents the VF PCI address and the address in the `<hostdev>` element represents the PCI address passed to the guest OS. Note down this address passed to the guest as it will be used in the subsequent step.

- g. Connect to the VM (guest OS) and make sure that the SR-IOV VFs are passed as PCI devices and detected by the guest OS. The commands to check this depend on the guest OS, however in most types of guest OS below command will work:

```
lspci
```

There must be an “n” number of VFs visible in the PCI list with the corresponding PCI address noted down in step 14 .

Example:

```
00:06.0 Ethernet controller: Intel Corporation 82599 Ethernet
Controller Virtual Function (rev 01)
00:07.0 Ethernet controller: Intel Corporation 82599 Ethernet
Controller Virtual Function (rev 01)
00:08.0 Ethernet controller: Intel Corporation 82599 Ethernet
Controller Virtual Function (rev 01)
00:09.0 Ethernet controller: Intel Corporation 82599 Ethernet
Controller Virtual Function (rev 01)
```

2. Connect to the traffic switches and make sure that the ports connected to the SR-IOV physical interfaces in the compute hosts are enabled and active. Commands depend on what type of switches are used as traffic switches. In case of using Extreme switches, run the below command:

```
show ports no-refresh
```

If the ports are not enabled and active, enable them. In case of using Extreme switches, run the below command:

```
enable ports <port_number>
```



3. Perform Ethernet packet tracing in the traffic switches for the SR-IOV ports and check whether traffic from the VM is reaching the traffic switches.

If the packets from the guest can reach the traffic switches, then the problem is probably outside CEE. Contact the cloud administrator to have the network checked.

4. If the problem persists, contact the next level of maintenance support, as described in section Contact Support.

Actions to Prevent Further Incidents

Collect necessary data for troubleshooting according to the Data Collection Guideline, Reference [6]. In case an official or permanent fix is applicable, contact the next level of maintenance support, as described in section Contact Support.

4.11 Power Outage in CEE System

Impact matrix:

Deployment	Multi-server
Hardware type	BSP, Dell, HDS
Components	OpenStack components

Recovery Action:

1. Recover the power supply of the CEE system.
2. Wait for 30 minutes for the automatic service restoration of CEE.
3. Refer to the Health Check Procedure, Reference [10].

4.12 Neutron Agents Are Down

Impact matrix:

Deployment	Multi-server
Hardware type	BSP, Dell, HDS
Components	OpenStack components in CEE

Recovery Action:

1. Connect to one of the vCICs through SSH.



2. Check the status of Neutron agents:

```
neutron agent-list
```

3. If any neutron agents are down, log in to each affected node (vCIC or compute) and restart the services which are down:

```
service openvswitch-switch restart
```

```
crm resource status clone_neutron-dhcp-agent
```

4. If you cannot restore any of the services, contact the next level of support.

4.13 Nova Services Are Down

Impact matrix:

Deployment	Multi-server
Hardware type	BSP, Dell, HDS
Components	OpenStack components in CEE

Recovery Action:

1. Connect to one of the vCICs through SSH.

2. Check the status of the Nova services:

```
nova service-list
```

3. If any of the services are down, log in to the affected node (vCIC or compute) and restart the services which are down:

```
service nova-consoleauth restart
```

```
service nova-cert restart
```

```
service nova-conductor restart
```

```
service nova-scheduler restart
```

```
service nova-compute restart
```

4. If you cannot restore any of the services, contact the next level of support.

4.14 Cinder Services Are Down

Impact matrix:

Deployment	Multi-server
Hardware type	BSP, Dell, HDS



Components	OpenStack components in CEE
-------------------	-----------------------------

Recovery Action:

1. Connect to one of the vCIC through SSH:

```
cinder service-list
```

3. If any of the cinder services are down login to each affected node (vCIC or Cinder) and restart the services which are down:

```
service cinder-volume restart
```

```
service cinder-backup restart
```

```
service cinder-scheduler restart
```

4. If you cannot restore any of the services, contact the next level of support.

4.15 Glance Image Not Shown

Impact matrix:

Deployment	Multi-server
Hardware type	BSP, Dell,HDS
Components	OpenStack components in CEE

Recovery Action:

1. Connect to one of the vCICs through SSH and check the status of the Glance image:

```
glance image-list
```

3. None or at least one image is returned in the list.
4. If the above command resulted in error state, contact the next level of support.

5 Contact Support

Before contacting the next level of maintenance support, prepare answers to the following questions to reduce the time needed for recovery. Further actions are outside the scope of this instruction.

1. What hardware type and software versions are used?
2. What type of CEE deployment is used?



3. Are there any changes or adaptations done to the CEE SW?
4. Were any maintenance or configuration activities ongoing at the time of the problem?
5. Were any alarms active before the emergency situation occurred? If yes, what alarms were active?
6. What actions have been performed so far to recover from the emergency situation?

It is also essential to collect the following data:

1. A site specific IP and VLAN plan must be available, especially for the interface connections indicating the components involved
2. All relevant user names and passwords
3. Collect all the relevant logs as described in the Data Collection Guideline, Reference [6].

6 Change Information

Revision	Date	Author	Comment
A	2016-04-06	ESUCHEL	First version
B	2017-04-01	ERAKBOH	Updated version for CEE R6
C	2017-07-21	ERAKBOH	Sections Loss of Connectivity due to SSL Certificate, MySQL database is down, MongoDB database is down updated. Sections Power outage of CEE system, Neutron agents are down, Nova services are down, Cinder services are down, Glance image not shown added.



7 References

- [1]. [SW Installation in Single Server Deployment](#)
- [2]. [SW Installation in Multi-Server Deployment](#)
- [3]. [Fuel Synchronization](#)
- [4]. [Atlas Backup](#)
- [5]. [Atlas Restore](#)
- [6]. [Data Collection Guideline](#)
- [7]. [CEE Connectivity User Guide](#)
- [8]. [Glossary of Terms and Acronyms](#)
- [9]. [Expiring Certificate](#)
- [10]. [Health Check Procedure](#)

Appendix

8 Additional Information

This section describes how to log in to the vCIC from a remote location, how to log in to vFuel and list the hostnames and addresses of the vCIC and the compute nodes. For more information, refer to the CEE Connectivity User Guide, Reference [7].



8.1 How to Connect

8.1.1 Single Server Deployment

From a remote location only the vCIC can be reached through the NBI IP address on VLAN `cee_om_sp`. Check the site specific IP plan to identify the NBI IP address of the vCIC. vFuel is shut off in the single server deployment and hence is not available in general. From the vCIC it is possible to connect to the compute host that hosts the vCIC.

8.1.2 Multi-Server Deployment

From a remote location only the vCIC hosts can be reached through the vCIC public IP address (or addresses) on VLAN `cee_om_sp`. The system has three vCIC hosts with public IP addresses as defined by the IP addresses of the `cee_om_sp` network. The first IP of the `cee_om_sp` that is denoted by “start” of the `cee_om_sp` network in the site specific IP and VLAN Plan is normally assigned as the vrouter public IP and the second IP is assigned as the HA proxy public IP which is the CIC NBI IP. SSH connection to this HA proxy public IP is not possible. The next sequential IPs in the `cee_om_sp` network are assigned to the three vCIC hosts.

It is possible to connect to the vCIC host by using IdAM credentials as follows:

Operational Mode

When the vCICs are in operational mode, the user must log in with

```
<personal_user>:
ssh <personal_user>@<IP_of_any_one_of_the_vCICs>
```

Change to user `ceeadm` after you logged in to the vCIC with the following command:

```
su - ceeadm
```

Maintenance Mode

When the vCICs are in maintenance mode, the user must log in with maintenance user `ceeadm`.

To log in with `ceeadm`, issue the following command:

```
ssh ceeadm@<IP_of_any_one_of_the_vCICs>
```

OpenStack commands need OpenStack tenant, URL and password information which need to be updated to execute OpenStack commands.

From the vCIC host, it is possible to connect to the vFuel node. To log in to the vFuel node, issue the following command:

```
ssh ceeadm@<vfuel_address>
```

For the `vfuel_address` refer to the site specific IP and VLAN Plan.



Once in vFuel, it is possible to connect to any of the compute and vCIC hosts by using the hostname of the respective hosts in vFuel node printout as specified in section List CIC and Compute Nodes.

8.2 List vCIC and Compute Nodes

8.2.1 Single Server Deployment

In case of single server deployment, there is only one compute host that hosts the vCIC. vFuel is intended to be turned off and hence it is not possible to list the nodes in the CEE environment. When vFuel is turned on for any reason, for example for maintenance activities, it is possible to list the nodes in the CEE environment using the command:

```
sudo fuel node
```

8.2.2 Multi-Server Deployment

To display the hostnames and IP addresses of the vCIC nodes, issue the following command, while logged in to the vFuel node:

```
sudo fuel node
```

Note: From a remote location only the vCIC hosts can be reached. For more information, see section How to Connect.

Below is an example printout:

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

id	status	name	cluster	ip	mac	roles	pending_roles	online	group_id
20	ready	compute-0-3	3	192.168.0.25	00:17:a4:77:00:04	compute, virt		True	3
19	ready	compute-0-5	3	192.168.0.24	00:17:a4:77:00:08	compute		True	3
22	ready	cic-2	3	192.168.0.26	32:1c:70:b7:1b:49	controller, mongo		True	3
16	ready	compute-0-2	3	192.168.0.20	00:17:a4:77:00:02	compute, virt		True	3
24	ready	cic-1	3	192.168.0.27	56:56:b1:3d:0d:44	controller, mongo		True	3
18	ready	compute-0-1	3	192.168.0.22	00:17:a4:77:00:00	compute, virt		True	3
23	ready	cic-3	3	192.168.0.28	72:d1:87:71:f2:47	controller, mongo		True	3
25	ready	compute-0-4	3	192.168.0.23	00:17:a4:77:00:06	compute		True	3
17	ready	compute-0-6	3	192.168.0.21	00:17:a4:77:00:0a	compute		True	3

For identifying the compute hosts corresponding to the vCIC, check `/mnt/cee_config/config.yaml` in vFuel.