Emergency Recovery Procedure

ERP R6

EMERGENCY RECOVERY

Abstract
The objective of this document is to act as collection of emergency recovery procedures for CEE R6 in Ericsson Cloud System solution.

Contents

1 Introduction

This Emergency Recovery Procedure provides a systematic approach for resolving a system emergency experienced with CEE R6.

Typically, an emergency procedure is required for conditions that make communication or normal management and 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, see Contact Support.

1.1 Scope

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

1. CEE R6

1.2 Target Group

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

2 Pre-requisites

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

2.1 Hardware and Software Tools

The following hardware and software are required:

1. Tools specified in SW Installation, User guide

2. Computer with SSH connection to the Cloud Infrastructure Controller (CIC)

2.2 System Backups

Some of the recovery actions will include recovery from the previous known stable state. In order to facilitate such recovery, the below should be available:

1. Fuel VM Cold Stand-by (Backup and Restore)

2. Atlas backup

2.3 Data Collection

Recent activities (Operation and Maintenance (O&M), software (SW) updates or feature introduction) within CEE and the tenant VM's are known, and the data needed for performing the emergency recovery actions has been collected with help of the Data Collection Guideline, User Guide.

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

2.4 Conditions

Before starting the procedure, get familiar with the following information types that can be issued and help with identifying scenarios and recovery procedures:

1. Information about CEE and ECS solution used, for example, product name, software version, platform, operating system and hardware type.

2. How to connect to CEE, refer 'How to Connect' section of Data Collection Guideline, User Guide.

3. Access to the respective user guides of all the hardware components involved in the ECS solution

2.5 Site Emergency Kit

Please contact the next level of support, see Contact Support.

3 Definitions

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

Deployment	Deployment scenario to which the problem is applicable. CEE R6 supports two deployment scenarios: Single server deployment, multi-server HA deployment.
Hardware type	Hardware type to which the problem is applicable
Components	Area of CEE functionality/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

BGW Border Gateway

CEE Cloud Execution Environment

CIC Cloud Infrastructure Controller

ECS Ericsson Cloud System

4 Procedures for Problem Types

4.1 Loss of connectivity to the system

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

Impact matrix:

Deployment	Single-server, Multi-server HA deployments
Hardware type	BSP, Dell, HP
Components	All components in CEE, tenant traffic shouldn't be 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 member of cee_om_sp VLAN. If the problem is observed to be with 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. If it is possible, connect to the compute host that hosts vCIC either through ssh connection or through serial connection.

3. Connect to the console of vCIC using:
   ```
   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
   ```

Start pinging the failed node from a CIC which is still alive. While sending ping requests, run the following command on the Compute multiple times, where the failed node is running and 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 by issuing the following command:

```
service openvswitch-switch restart
```

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

4.2 Loss of Connectivity due to SSL Certificate

Connectivity to CEE through its NBI (to HA proxy public VIP or public vCIC IP addresses) is lost due to corrupt or expired SSL certificate.

Impact matrix:

Deployment	Single-server, Multi-server HA deployments
Hardware type	BSP, Dell, HP
Components	All components in CEE, tenant traffic shouldn't be impacted

Recovery Action:

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

2. Obtain correct SSL certificates from the customer and store it in the directory `/mnt/cee_config/` of vFUEL.

   Note: For Single-server deployments, vFUEL should be available for this recovery action.

3. If required, update `/mnt/cee_config/config.yaml` with the appropriate certificate file names.

4. Reinstall the SSL certificates by running the ansible playbooks from vFuel using the following command:
   ```
   cd /usr/share/ericsson-orchestration/playbooks/
   openstack-ansible setup-all.yml
   ```

5. If the problem still persists, then contact next level of maintenance support, see Contact Support.

4.3 CIC VM (vCIC) is down

One or more vCIC is observed to be down.

Impact matrix:

Deployment	Single-server, Multi-server HA deployments
Hardware type	BSP, Dell, HP
Components	All components in CEE, tenant traffic shouldn't be impacted

Recovery Action:

1. If it is possible, connect to the compute host that hosts vCIC either through ssh connection or through serial connection.

2. Make sure that the CIC VM is defined in the Libvirt of the compute host:
   ```
   virsh list --all
   ```

3. Attempt to start the CIC VM manually:
   ```
   virsh start cic-x_VM
   ```

4. If the CIC VM doesn't start, then contact next level of maintenance support, see Contact Support.

5. If the CIC VM can be started, then connect to the vCIC through SSH.

   Note: In case if vCIC starts up, but SSH connection is not possible then refer Loss of connectivity to the system.

6. Check the status of the vCIC in the cluster:
   ```
   crm_mon -1 -rf
   ```

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

7. If vCIC is ONLINE in the cluster, but didn't the resources didn't join the cluster, then check the status of the Corosync/Pacemaker services:
   ```
   service corosync status
   service pacemaker status
   service corosync-notifyd status
   ```

8. If any of the Corosync/Pacemaker services are not running, attempt to start those services using the command:
   ```
   service <service_name> start
   ```

9. Check the cluster status and monitor whether the resources are starting up in the cluster:
   ```
   crm_mon -1 -rf
   ```

10. If the problem still persists, then contact next level of maintenance support, see Contact Support.

4.4 CIC VM (vCIC) is unreachable via SSH

One or more vCIC is observed to be unreachable via ssh. vCIC is pingable and accessible via console but console session is unresponsive.

Impact matrix:

Deployment	Single-server, Multi-server HA deployments
Hardware type	BSP, Dell, HP
Components	All components in CEE, tenant traffic shouldn't be impacted

Recovery Action:

1. If it is possible, connect to the compute host that hosts vCIC either through ssh connection or through serial connection.

2. Make sure that the CIC VM 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, then contact next level of maintenance support, see Contact Support.

6. If the CIC VM can be started, then connect to the vCIC through SSH.

   Note: In case if vCIC starts up, but SSH connection is not possible then refer Loss of connectivity to the system.

7. Check the status of the vCIC in the cluster:
   ```
   crm_mon -1 -rf
   ```

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

8. If vCIC is ONLINE in the cluster, but didn't the resources didn't join the cluster, then check the status of the Corosync/Pacemaker services:
   ```
   service corosync status
   service pacemaker status
   service corosync-notifyd status
   ```

9. If any of the Corosync/Pacemaker services are not running, attempt to start those services using the command:
   ```
   service <service_name> start
   ```

10. Check the cluster status and monitor whether the resources are starting up in the cluster:
    ```
    crm_mon -1 -rf
    ```

11. If the problem still persists, then contact next level of maintenance support, see Contact Support.

4.5 Fuel VM (vFuel) is down

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

Impact matrix:

Deployment	Single-server, Multi-server HA deployments
Hardware type	BSP, Dell, HP
Components	All components in CEE, tenant traffic shouldn't be impacted

Recovery Action:

1. Connect to the compute host that hosts vFuel either through ssh connection or through serial connection.

2. Check that the vFuel VM is present and in state running on the compute host:
   ```
   virsh list --all
   ```

3. Log into vFuel console:
   ```
   virsh console fuel_master
   ```

4. If the vFuel VM 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 VM boots normally, then connect to vFuel through SSH.

   Note: In case if vFuel boots up, but SSH connection is not possible then refer Loss of connectivity to the system.

6. If the problem still persists, then contact next level of maintenance support, see Contact Support.

4.6 MySQL database is down

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

Impact matrix:

Deployment	Single-server, Multi-server HA deployments
Hardware type	BSP, Dell, HP
Components	Openstack components in CEE, watchmen, CM-HA

Recovery Action:

1. Connect to one of the vCIC through SSH

2. Check the status of the MySQL resource in the cluster:
   ```
   crm_mon -1 -rf
   ```

3. If the MySQL resource is observed to be down in one or two vCIC, connect to the vCIC through SSH and check the MySQL process status:
   ```
   ps -ef | grep sql
   ```
   If there is a process with '`mysql-wss start`', it indicates that Corosync/pacemaker is attempting to start the MySQL resource on that vCIC.

4. If the attempt to start the MySQL by Corosync/pacemaker doesn't succeed, then cleanup the MySQL database directory in the vCIC. Identify a location that has enough free space to contain the /var/lib/mysql/ directory contents before proceeding. If you are unable to back up the /var/lib/mysql/ directory contents, then contact the next level of maintenance support, see Contact Support.

   Find the amount of space needed to back up /var/lib/mysql:
   ```
   du -sh /var/lib/mysql/
   ```

   Move the contents of /var/lib/mysql/ to the backup location. Here the /var/lib/glance/ directory is used as an example
   ```
   mkdir /var/lib/glance/mysql
   mv /var/lib/mysql/* /var/lib/glance/mysql
   ```

   Observe the next attempt of the Corosync/pacemaker to start MySQL

   Note: Corosync/Pacemaker will periodically reattempt to start the MySQL resource. Each attempt will have a predefined timeout value after which a new attempt will be made.

5. If the problem still persists, then contact next level of maintenance support, see Contact Support.

4.7 MongoDB database is down

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

Impact matrix:

Deployment	Single-server, Multi-server HA deployments
Hardware type	BSP, Dell, HP
Components	Ceilometer

Recovery Action:

1. Check the status of MongoDB service in each of the vCIC
```
     for node in $(fuel node | awk '/cic-/ {print $5}'); do echo
"${node}: $(ssh -q ${node} "service mongodb status")"; done
```

2. In case if there are atleast one vCIC with MongoDB service running, goto step 12 in this recovery action. If none of the vCIC has MongoDB running, then continue with next step.

3. Connect to one of the vCIC through SSH

4. Create a new MongoDB directory under glance partition and modify the MongoDB configuration to point to the new MongoDB directory location.
```
     service mongodb stop
     mkdir /var/lib/glance/mongodb
     rsync -Pa /var/lib/mongo/mongodb/ /var/lib/glance/mongodb/
     sed -i.bak 's|/var/lib/mongo/mongodb|/var/lib/glance/mongodb|'
/etc/mongodb.conf
     service mongodb start
```

5. Connect to the started MongoDB instance and check whether that is the primary MongoDB instance in the replica set.
```
     mongo -u admin -p <db_password> admin !<for db_password check
the ceilometer part
     in astute.yaml of cic or check
     primary-mongo_y.yaml file in
/var/lib/ericsson/pre_deploy/deployment_x/)
     db.isMaster();
     rs.conf()
     rs.status()
     exit;
```

6. If that is not the primary MongoDB instance, then stop the MongodB and repeat the above steps in other CICs until the primary MongoDB instance is identified:
```
     service mongodb stop
     mkdir /var/lib/glance/mongodb
     rsync -Pa /var/lib/mongo/mongodb/ /var/lib/glance/mongodb/
     sed -i.bak 's|/var/lib/mongo/mongodb|/var/lib/glance/mongodb|'
     /etc/mongodb.conf
     service mongodb start
     mongo -u admin -p <db_password> admin !<for db_password check
the
     ceilometer part in astute.yaml of cic or check
     primary-mongo_y.yaml file in
/var/lib/ericsson/pre_deploy/deployment_x/)
     db.isMaster();
     rs.conf()
     rs.status()
     exit;
```

7. Once connected to the primary MongoDB instance and remove the contents of the Ceilometer:
```
     mongo -u admin -p <db_password> admin
     total_size = 0; db.adminCommand( { listDatabases: 1 }
     ).databases.forEach(function (database) {
     total_size += database.sizeOnDisk
     print(database.name +': '+ (database.sizeOnDisk / 1073741824)
     +' Gb')
     if (database.name == 'admin') {
     return true
     }
     db = db.getSiblingDB(database.name)
     db.getCollectionNames().forEach(function(collectionName) {
```

```
print('\t' + collectionName +': '+
(db[collectionName].stats().storageSize / 1073741824) + ' Gb')
})
}); print('total: '+ (total_size / 1073741824) + ' Gb')
use ceilometer
db["meter"].find()
db["event"].find()
db.stats()
db["meter"].remove({"timestamp":{"$lt":ISODate("2016-02-
03T19:30:07.805Z")}}); ! here modify the date as per the need!
db["event"].remove({"timestamp":{"$lt":ISODate("2016-02-
03T19:30:07.805Z")}}); ! here modify the date as per the need!
db.stats()
rs.status()
total_size = 0; db.adminCommand( { listDatabases: 1 }
).databases.forEach(function (database) {
total_size += database.sizeOnDisk
print(database.name +': '+ (database.sizeOnDisk / 1073741824)
+' Gb')
if (database.name == 'admin') {
return true
}
db = db.getSiblingDB(database.name)
db.getCollectionNames().forEach(function(collectionName) {
print('\t' + collectionName +': '+
(db[collectionName].stats().storageSize / 1073741824) + ' Gb')
})
}); print('total: '+ (total_size / 1073741824) + ' Gb')
exit;
```

8. Remove the content of the mongodb partition in the secondary nodes after backing it up to a temporary location (perform this on both the CICs where MongoDB instance is secondary):

```
mkdir /var/lib/glance/mongo
mv /var/lib/mongo/mongodb/* /var/lib/glance/mongo ! Now the
Mongo disk
partition should be empty !
service mongodb status
service mongodb start
service mongodb status
```

   Note: Now the secondary MongoDB instance should sync from the primary and hence should have the correct disk occupancy.

9. Now connect to the primary MongoDB instance and step down from primary role:

```
mongo -u admin -p <db_password> admin
rs.conf()
rs.stepDown(120) ! This will make this primary MongoDB
instance to step down from being Primary !
rs.conf() ! It will take some time for changing the Primary !
exit;
```

10. Once the primary is changed, now delete the contents and re-sync the mongo for the ‘previous-primary’ Mongodb instance:

```
mkdir /var/lib/glance/mongo
mv /var/lib/mongo/mongodb/* /var/lib/glance/mongo ! Now the
Mongo disk partition should be empty !
service mongodb status
service mongodb start
service mongodb status
```

11. If the MongoDB instance is recovered, then exit this recovery action. If the problem still persists, then contact next level of maintenance support, see Contact Support.

12. If atleast one vCIC has MongoDB service running, identify the primary instance of MongoDB by executing the following command in vFUEL:

```
for node in $(fuel node | awk '/cic-/ {print $5}'); do echo
"${node}: $(ssh -q ${node} "echo 'db.isMaster()' | mongo | grep
ismaster")"; done
```

13. Connect to the vCIC with primary MongoDB instance, through SSH and goto step 7 of this recovery action.

4.8 Not possible to upload Glance images

One or more vCIC is observed to be down.

Impact matrix:

Deployment	Single-server, Multi-server HA deployments
Hardware type	BSP, Dell, HP
Components	Glance, Swift

Recovery Action:

1. Contact next level of maintenance support, see Contact Support.

4.9 Not possible to instantiate VMs

One or more vCIC is observed to be down.

Impact matrix:

Deployment	Single-server, Multi-server HA deployments
Hardware type	BSP, Dell, HP
Components	All components in CEE, tenant traffic shouldn’t be impacted

Recovery Action:

1. Contact next level of maintenance support, see Contact Support.

4.10 Problem Type: SR-IOV related problems

Basic Data

Symptoms:	- SR-IOV traffic is not working - SR-IOV Ethernet interfaces in the compute node is working fine

	- Corresponding tenant VMs are in ACTIVE and RUNNING state - All compute hosts are ok

Prerequisites and Additional Data Needed

1. Access to the compute hosts
2. Access to the CIC hosts
3. Access to the traffic switches

This method can be used if,

1. 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 and 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 the VM details in CIC 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 a. In the answer printout of 'nova flavor-show' check the extra_specs field and count the number of VFs requested for the VMs. This can be done by summing up the pci_passthrough:alias devices requested.

      For example, the below flavor indicates, 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 Fuel 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 previous step a. has SR-IOV enabled

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

The compute host element in the config.yaml should have the SR-IOV configurations like below:

```
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 (can be identified in the field `os-EXT-SRV-ATTR:host` from the answer printout in step a) where the VM is being hosted and check the SR-IOV virtual function (VF) / physical function (PF) details

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

Make sure that the PFs SR-IOV PCI address identified in the step c is bind to IXGBE driver and there are corresponding VFs from the PFs. And the VFs are bound to 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 the step 1 should be running in the compute host.

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

```
virsh dumpxml <instance_name>
```

In the XML file of the VM, there should be ‘n’ number of hostdev elements, where ‘n’ corresponds to the number of SR-IOV interfaces requested for the VM that is identified in step b.

```
<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 virtual functions (VFs) are passed as PCI devices and detected by guest OS. The commands to check this depends on the guest OS, however in most of the guest OS, below command should work:

```
lspci
```

There should be ‘n’ number of VFs visible as below, in the PCI list with the corresponding PCI address noted down in previous step f.

```
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 depends on what type of switches are used as traffic switches. For Extreme switches, below is the command to check

```
show ports no-refresh
```

If the ports are not enabled and active, then enable them by using the following command (in case of Extreme switches)

```
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 Guest can reach the traffic switches, then most likely the problem should be outside the CEE edge, contact the cloud administrator to get the network checked.

4. If the problem still persists, then contact next level of maintenance support, see Contact Support.

Actions to Prevent Further Incidents

Collect necessary data for troubleshooting according to the Data Collection Guideline and in case of an official or permanent fix is applicable, then contact next level of maintenance support, see Contact Support.

5 Contact Support

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

1. Which hardware type is being used and software versions are used?
2. What type of CEE deployment is used?
3. Is there any changes/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? If yes, what alarms were active?
6. What actions have been performed so far to recover from the emergency situation?

It is also essential that the following data is available:

1. A customer-specific IP and VLAN plan should be available, especially for the interface connections indicating the components involved.
2. All relevant user names and passwords.
3. If possible collect all the relevant logs as per Data Collection Guideline.

For detailed information, refer to Data Collection Guideline.

6 Change Information

Revision	Date	Author	Comment
A	2016-04-06	ESUCHEL	First version

7 References

[1]. Atlas Backup

[2]. Atlas Restore

[3]. CEE R6 Network Impact Report

[4]. CEE Architecture Description

[5]. CEE Infrastructure Backup

[6]. CEE Infrastructure Restore

[7]. CEE on BSP

[8]. CEE Technical Description

[9]. CIC Failed

[10]. Compute Host Failed

[11]. Data Collection Guideline

[12]. Fuel Synchronyzation

[13]. Health Check Procedure

[14]. Troubleshooting Guideline

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 Compute nodes.

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 customer IP plan to identify the NBI IP address of vCIC. vFuel will be shutoff in the single server deployment and hence wouldn't be available in general. From the vCIC it should be possible to connect to the compute host that hosts vCIC.

8.1.2 Multi-server deployment

From a remote location only the CIC hosts can be reached through the CIC public IP address(es) on VLAN `cee_om_sp`. The system has three CIC hosts with public IP addresses as defined by IP addresses of `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 customer IP and VLAN Plan is normally assigned as the vrouter public IP and the second IP will be assigned as the HA proxy public IP which is the CIC NBI IP. It should be noted that SSH connection to this HA proxy public IP will not be possible. The next sequential IPs in the cee_om_sp network will be assigned to the three CIC hosts.

It shall be possible to connect to the CIC host by using IdAM credentials as follows:

Operational Mode

When the CICs are in operational mode the user must log in with *<personal-user>*:

```
ssh <personal-user>@IP_of_any_one_of_the_CIC
```

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

```
su - ceeadm
```

Maintenance Mode

When the CICs 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_CIC
```

Openstack commands will need Openstack tenant, URL and password information which needs to be updated to execute the Openstack commands.

From the CIC host, it should be possible to connect to the Fuel node. To log in to the Fuel node, issue the following command:

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

For the *vfuel_address* refer to the customer IP and VLAN Plan.

Once in vFuel, it should be possible to connect to any of the compute and CIC hosts by using the hostname of the respective hosts in fuel node printout as specified in next chapter List CIC and Compute Nodes.

8.2 List CIC and Compute Nodes

8.2.1 Single-server deployment

In case of single-server deployment, there will be only one compute host that hosts the vCIC. vFuel is intended to be turned off and hence it won't be possible to list the nodes in the CEE environment. When vFuel is turned on for any reasons, like 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 CIC nodes issue the following command, while being logged in to the Fuel node:

```
sudo fuel node
```

Note: From a remote location only the CIC servers can be reached. For more information, refer 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.