

# Health Check Procedure

Cloud Execution Environment

OPERATING 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	Scope	1
1.2	Target Groups	1
1.3	Prerequisites	1
<b>2</b>	<b>Overview</b>	<b>3</b>
2.1	When To Perform a Health Check	3
2.2	File Location	3
2.3	Script Execution	4
2.4	Modules	4
2.5	Modules with [SKIPPED] Verdict	18
<b>3</b>	<b>Report Problems</b>	<b>20</b>





# 1 Introduction

This document is to help support engineers check that the Cloud Execution Environment (CEE) operates in a fault-free state, and to detect issues that can affect normal operation.

## 1.1 Scope

The process is applicable to all CEE configurations.

## 1.2 Target Groups

This document is intended for both internal and external customers monitoring system health:

- Support organization personnel
- Customer operation and maintenance (O&M) personnel

## 1.3 Prerequisites

This section describes the prerequisites of performing the health check procedure.

### 1.3.1 Documents

Before starting the procedure, ensure that the following documents are available:

- Data Collection Guideline
- CEE Connectivity User Guide

### 1.3.2 Conditions

Before performing a health check, ensure that the following conditions are met:

- Users of this document must be familiar with commands and tools within CEE and OpenStack.
- Access to deployment-specific credentials must be available.
- There are no ongoing CEE maintenance activities. For more information about when the health check procedure is recommended, see Section 2.1 on page 3.



### 1.3.3

#### User Access

root access to the vFuel node is required. The procedure below can only be executed in root. For more information, refer to the [CEE Connectivity User Guide](#).



## 2 Overview

This document covers the procedures for checking the health of CEE and detecting issues before they become threats to the system.

“Health” in the context of this document means that CEE is running, provides the required functionality, and is available for the users. Health condition is evaluated by executing several checks. These checks are based on the information collected from printouts. The verdict for each check is displayed on the console and a detailed result is stored in the log file.

### 2.1 When To Perform a Health Check

The time needed to execute a health check depends on factors such as the complexity of the check or the system performance. If a check takes a long time, it is possible that CEE is not functioning correctly.

Checks need to be executed before and after an update of CEE or other maintenance activities are performed. Perform a health check in the following cases:

- As part of CEE maintenance activities
- As part of the CEE deployment
- As part of SW update and rollback procedure
- As part of CEE region expansion
- Before starting a troubleshooting procedure
- Before and after emergency recovery procedure
- Every time it is recommended in another procedure

**Note:** CEE collects a large amount of In-Service Performance (ISP) and Fault Management (FM) data. Alarms must be available in the management system (Atlas).

### 2.2 File Location

Before starting the health check procedure, make sure that the following script is available on vFuel:

```
/usr/bin/healthcheck.py
```

The log file can be found in the following directory:

```
/var/log/healthcheck/<log_filename>
```



## 2.3 Script Execution

Execute the script on vFuel as root using one of the following commands:

- The following command runs the script on the entire CEE region, and displays the verdicts for the respective checks:

```
[root@fuel ~]# healthcheck.py
```

- The `--help` argument displays the usage of `healthcheck.py`:

```
[root@fuel ~]# healthcheck.py --help
```

- The following command displays the verdicts for the services running on a specific node:

```
[root@fuel ~]# healthcheck.py --node <node_name>
```

For example:

```
healthcheck.py --node cic-1
```

This command displays the verdicts for the services running on `cic-1`.

- The following command displays the verdicts for the respective service status in the cluster:

```
[root@fuel ~]# healthcheck.py --service <service_name>
```

For example:

```
healthcheck.py --service nova
```

This command displays the verdicts for the nova service.

### 2.3.1 Execution Time

The execution time of the health check procedure depends on the size of the CEE region. The procedure takes approximately two minutes for smaller deployments and approximately nine minutes for large deployments.

## 2.4 Modules

This section describes the health check modules that are included in the `healthcheck.py` script.

**Note:** The script automatically omits modules if their preconditions are not met. In this case, no output is generated for the module.

The following verdicts are possible for each module:





**Note:** During the health check procedure, only one verdict is returned for each module.

<b>[PASSED]</b>	The requirements of the specific module are met.
<b>[FAILED]</b>	The requirements of the specific module are not met.
<b>[WARNING]</b>	This verdict signals a possible system issue and requires the attention of the user.
<b>[SKIPPED]</b>	There are commands included in the <code>healthcheck.py</code> script for troubleshooting purposes that do not return a [PASSED], [FAILED], or [WARNING] verdict. The script executes these commands and prints the results in the log file without displaying a verdict in the console output. For a list of modules that return a [SKIPPED] verdict, see Section 2.5 on page 18.

The console output displays the name of the module and the verdict. For example:

```
2018-03-20 07:21:33 : 192.168.0.28(cic-1) > Checking GALERA CLUSTER STATUS [PASSED]
```

Example 1 Console Module Output

The log file contains the detailed description of each check. For example:

```
2018-03-20 07:23:22 : 192.168.0.28(cic-1) > Checking GALERA CLUSTER STATUS
Command: mysql -sN -hlocalhost -D information_schema -e "SELECT
variable_value FROM global_status WHERE variable_name = 'WSREP_CLUSTER_STATUS'"
Result:
Primary
Verdict:
[PASSED]
```

Example 2 Log File Module Printout

## 2.4.1 Check For Alarms

Performance and Fault Management alarms are reported by Watchmen to Atlas or the Ericsson Cloud Manager (ECM). If Atlas or ECM are not available, use the `watchmen-client` command to fetch the active alarm list. For more information, refer to the [CEE CLI Guide](#).

Table 1 Check For Alarms

Verdict	Meaning
[PASSED]	The output does not contain Major or Minor alarms.
[FAILED]	The output contains Major or Minor alarms.
[WARNING]	Not applicable.



### Implications

Check the active alarms and act according to the relevant Operating Instructions (OPIs) in each case.

Check and analyze the alarm history for intermittent or persistent issues.

## 2.4.2 Check Alarm History

Performance and Fault Management alarms are reported by Watchmen to Atlas or the Ericsson Cloud Manager (ECM). If Atlas or ECM are not available, use the `watchmen-client` command to fetch the active alarm list. For more information, refer to the [CEE CLI Guide](#).

Table 2 Check Alarm History

Verdict	Meaning
[PASSED]	The output contains less than six Major alarms.
[FAILED]	Not applicable.
[WARNING]	The output contains more than five Major alarms.

### Implications

Check the active alarms and act according to the relevant OPIs in each case.

Check and analyze the alarm history for intermittent or persistent issues.

## 2.4.3 Check Uptime For Alarms

Table 3 Check Alarm Uptime

Verdict	Meaning
[PASSED]	The periodic uptime measurements show that the uptime for an alarm is more than one day.
[FAILED]	Not applicable.
[WARNING]	The periodic uptime measurements show that the uptime for an alarm is less than one day.



## 2.4.4 Verify vCIC Connectivity

Table 4 Verify vCIC Connectivity

Verdict	Meaning
[PASSED]	Connectivity between one vCIC and the other two vCICs is established and present.
[FAILED]	Connection from one vCIC to either of the other two vCICs is lost.
[WARNING]	Not applicable.

## 2.4.5 Check the Presence of Crash and Core Dumps

Table 5 Check Crash and Core Dumps

Verdict	Meaning
[PASSED]	No crash or core dumps are present.
[FAILED]	Crash or core dumps, or both are present.
[WARNING]	Not applicable.

## 2.4.6 Check Pacemaker

Pacemaker is a cluster resource manager.

Table 6 Check Pacemaker

Verdict	Meaning
[PASSED]	There are no vCICs or cluster resources in FAILED or STOPPED state.
[FAILED]	There are vCICs or cluster resources in FAILED or STOPPED state.
[WARNING]	Not applicable.

### Implications

If any resources are in FAILED state, then they must be acted upon.

If there are any resources that are in STOPPED state, that can be because of dependencies on resources in a FAILED state.



## 2.4.7 Check Galera Cluster Status

The Galera cluster provides certification based replication among the cluster nodes.

Table 7 Check Galera Cluster Status

Verdict	Meaning
[PASSED]	All vCIC hosts belong to the Primary component.
[FAILED]	Not all vCIC hosts belong to the Primary component.
[WARNING]	Not applicable.

## 2.4.8 Check Galera Cluster Size

Table 8 Check Galera Cluster Size

Verdict	Meaning
[PASSED]	The number in the output equals the number of vCIC nodes.
[FAILED]	The number in the output does not equal the number of vCIC nodes.
[WARNING]	Not applicable.

## 2.4.9 Check Galera Desired State

Table 9 Check Galera Desired State

Verdict	Meaning
[PASSED]	The number in the output is 4 (Galera desired state).
[FAILED]	The number in the output is not 4 (Galera desired state).
[WARNING]	Not applicable.

## 2.4.10 Check That vCICs Are Not in Maintenance Mode

Table 10 Check That vCICs Are Not in Maintenance Mode

Verdict	Meaning
[PASSED]	The output contains <code>runlevel N 2</code> (indicating that the system is in Multiuser mode).



Verdict	Meaning
[ FAILED ]	The output does not contain runlevel N 2.
[ WARNING ]	Not applicable.

### 2.4.11 Check Neutron Agents

Neutron agents mentioned are monitored by CEE ISP.

For systems using Software Defined Networking (SDN) tightly integrated into CEE, the Neutron DHCP agent and the Neutron Open vSwitch agent are not visible in the agent list. In this case the agent list can be empty.

For systems not using SDN, all agents are expected to be alive.

**Note:** The Neutron DHCP agent has to be active on one vCIC host for systems not using SDN, but it can be present on other vCIC hosts with down status (displayed as xxx).

Table 11 Check Neutron Agents

Verdict	Meaning
[ PASSED ]	The output does not contain False/xxx.
[ FAILED ]	The output contains False/xxx.
[ WARNING ]	Not applicable.

#### Implications

If any of the Neutron agents are down, collect data related to the problem, as described in the [Data Collection Guideline](#). For the recovery procedure, refer to the relevant section of the [Emergency Recovery Procedure](#).

### 2.4.12 Verify Disk Space Utilization

Table 12 Verify Disk Space Utilization

Verdict	Meaning
[ PASSED ]	The disk space utilization is less than or equals to 80% per partition.
[ FAILED ]	The disk space utilization is more than 80% per partition.
[ WARNING ]	Not applicable.



### 2.4.13 Check iSCSI Multipath Connection to VNX

**Note:** Centralized block storage based on EMC VNX hardware was phased out as CEE-certified configuration.

Table 13 Check iSCSI Multipath Connection to VNX

Verdict	Meaning
[PASSED]	The output does not contain FAILED/FAULT string.
[FAILED]	The output contains FAILED/FAULT string.
[WARNING]	Not applicable.

### 2.4.14 Check the State of Extreme Switches from Neutron Perspective

**Note:** Only applicable for systems using Extreme switches dynamically configured by CEE. Extreme switches are monitored by Fault Management. Check for alarms.

Table 14 Check the State of Extreme Switches from Neutron Perspective

Verdict	Meaning
[PASSED]	The output does not contain False, Error or Inactive states.
[FAILED]	The output contains False, Error or Inactive states.
[WARNING]	Not applicable.

### 2.4.15 Check the State of Ethernet Interfaces

**Note:** Not applicable for single server platform.

OVS bonding must show that Ethernet ports are enabled. The list of all checked interfaces can be found in the log file.

**Note:** In general, the state of all Ethernet interfaces must be UP. See the module ETHERNET INTERFACES THAT ARE IN UP STATE in Section 2.5 on page 18 and check the log file for more information.

Table 15 Check the State of Ethernet Interfaces

Verdict	Meaning
[PASSED]	The output lists the Ethernet interfaces in enabled state.



Verdict	Meaning
[FAILED]	The output does not list the Ethernet interfaces in enabled state.
[WARNING]	Not applicable.

## 2.4.16 Check Zombie Processes

Table 16 Check Zombie Processes

Verdict	Meaning
[PASSED]	The output does not list zombie processes.
[FAILED]	The output lists zombie processes.
[WARNING]	Not applicable.

## 2.4.17 Check vFuel Status

Table 17 Check vFuel Status

Verdict	Meaning
[PASSED]	The procedure returns a printout with a list of vCIC and compute hosts. Each node must display the status column as ready.
[FAILED]	The status column of the output displays one of the following states: discover, provisioning, provisioned, deploying, error.
[WARNING]	Not applicable.

## 2.4.18 Check MongoDB Status

**Note:** Not applicable for single server platform.

Table 18 Check MongoDB Status

Verdict	Meaning
[PASSED]	The output contains start/running status.
[FAILED]	The output does not contain start/running status.
[WARNING]	Not applicable.



### 2.4.19 Check MongoDB Replication

**Note:** Not applicable for single server platform.

MongoDB has built-in functions for replication.

Table 19 Check MongoDB Replication

Verdict	Meaning
[PASSED]	The output displays one vCIC in PRIMARY state and at least one vCIC in SECONDARY state.
[FAILED]	The output displays no vCICs in SECONDARY replica state.
[WARNING]	Not applicable.

### 2.4.20 Check Nova Services

It is expected that `nova-scheduler`, `nova-conductor` and `nova-consoleauth` services are present and enabled on each vCIC. In addition, the `nova-compute` service is enabled on each compute host.

Table 20 Check Nova Services

Verdict	Meaning
[PASSED]	The output lists all services in enabled state.
[FAILED]	The output lists a service in down state.
[WARNING]	Not applicable.

#### Implications

If any of the Nova services are down, collect data related to the problem, as described in the [Data Collection Guideline](#). For the recovery procedure, refer to the relevant section of the [Emergency Recovery Procedure](#).

### 2.4.21 Verify RAM Utilization

CPU, RAM, and local disk usage is monitored by Fault Management. Check for alarms.

There must be at least 20% of RAM free on the vCIC.

**Note:** In compute hosts, the use of `ReservedHugePages` for VMs can result in close to 100% RAM usage.





Table 21 Verify RAM Utilization

Verdict	Meaning
[PASSED]	At least 20% of RAM is free on the vCIC.
[FAILED]	Less than 20% of RAM is free on the vCIC.
[WARNING]	Not applicable.

#### 2.4.22 Check the State of the SDN Controller

Only applicable for systems using SDN tightly integrated into CEE.

Table 22 Check SDN Controller State

Verdict	Meaning
[PASSED]	The output lists all services in OPERATIONAL mode.
[FAILED]	The output does not list all services in OPERATIONAL mode.
[WARNING]	Not applicable.

#### 2.4.23 Check ScaleIO Cluster Status

ScaleIO is used for Storage.

Table 23 Check ScaleIO Cluster Status

Verdict	Meaning
[PASSED]	The output displays the ScaleIO cluster in Normal state.
[FAILED]	The output does not display the ScaleIO cluster in Normal state.
[WARNING]	Not applicable.

#### 2.4.24 Check Swift Store on ScaleIO

Swift store on ScaleIO needs to be activated. To configure Swift to use ScaleIO as storage back end, the type value of `swift_on_backend_storage` must be set to `scaleio` in the `config.yaml`. For more information, refer to the [Configuration File Guide](#).

**Note:** This check is not executed when the `activation_mode` value of `swift_on_backend_storage` is set to `manual`.



Table 24 Check Swift Store in ScaleIO

Verdict	Meaning
[PASSED]	<ul style="list-style-type: none"><li>• The output of <code>cinder list</code> has the following pattern: CEE+cic-&lt;index&gt;.domain.tld+/dev/image/glance&lt;+integer&gt;</li><li>• There are as many physical volumes connected to the volume group image as the number of Cinder volumes displayed for a specific vCIC.</li><li>• The logical volume size of the Swift store in each vCIC equals the summarized size of the Cinder volumes for the specific vCIC.</li></ul>
[FAILED]	One or more conditions of [PASSED] is not met.
[WARNING]	Not applicable.

## 2.4.25 Check vFuel Services

The following services are checked: `astute`, `cobbler`, `keystone`, `mcollective`, `nailgun`, `nginx`, `ostf`, `postgres`, `rabbitmq`, `rsync`, `rsyslog`.

Table 25 Check vFuel Services

Verdict	Meaning
[PASSED]	The output lists all vFuel services.
[FAILED]	One or more vFuel services are not listed in the output. Any service missing from the list is not functional.
[WARNING]	Not applicable.

## 2.4.26 Check SDN VTEP Configuration: Check Number of TEPs

This module is only applicable for systems using SDN tightly integrated into CEE.



Table 26 SDN VTEP Configuration: Check Number of TEPs

Verdict	Meaning
[PASSED]	The number of TEP endpoints (lines in the output table) matches the number of compute nodes on vFuel.
[FAILED]	The number of TEP endpoints (lines in the output table) does not match the number of compute nodes on vFuel.
[WARNING]	Not applicable.

## 2.4.27 Check SDN VTEP Configuration: Check Number of VXLAN Tunnels

This module is only applicable for systems using SDN tightly integrated into CEE.

From each TEP endpoint there should be a VXLAN tunnel defined to all other TEP endpoints. If there are  $n$  TEP endpoints configured,  $n*(n-1)$  tunnels must be visible in the table. For example, 4 TEP endpoints require  $4*3 = 12$  tunnels. All tunnels must have trunk-state UP.

Table 27 SDN VTEP Configuration: Check Number of VXLAN Tunnels

Verdict	Meaning
[PASSED]	The number of tunnels visible in the table is $n*(n-1)$ .
[FAILED]	The number of tunnels is not $n*(n-1)$ .
[WARNING]	Not applicable.

## 2.4.28 Check SDN VTEP Configuration: Check Trunk State

This module is only applicable for systems using SDN tightly integrated into CEE.

Table 28 SDN VTEP Configuration: Check Trunk State

Verdict	Meaning
[PASSED]	The output shows that the trunk state is UP.
[FAILED]	The output shows that the trunk state is DOWN.
[WARNING]	Not applicable.



## 2.4.29 Check OpenStack Storage Components

Table 29 Check OpenStack Components

Verdict	Meaning
[PASSED]	There is a successful output: <ul style="list-style-type: none"><li>• At least one Glance image is returned.</li><li>• Cinder schedulers are enabled and up on all vCICs.</li><li>• In case of multi-server deployment, Ceilometer meters are listed.</li><li>• <code>cinder-volume</code> is present as one of the services when EMC ScaleIO is used.</li></ul>
[FAILED]	One or more components are missing or disabled.
[WARNING]	Not applicable.

### Implications

If any of the Cinder services are down, collect data related to the problem, as described in the [Data Collection Guideline](#). For the recovery procedure, refer to the relevant section of the [Emergency Recovery Procedure](#).

## 2.4.30 Check Cinder Scheduler Status

**Note:** Not applicable for single server platform.

The conditions of this check depend on whether OpenStack Cinder services are integrated in CEE. Cinder is integrated if the `configure_cinder_ericsson` attribute of the `ericsson_openstack_config` Fuel plugin is set to true or not defined in the `config.yaml`.

If Cinder services are not integrated, the following verdicts can be returned:

Table 30 Check Cinder Scheduler Status - Cinder Not Integrated

Verdict	Meaning
[PASSED]	The result of the check is 0.
[FAILED]	The result of the check is not 0.
[WARNING]	Not applicable.

If Cinder services are integrated, the following verdicts can be returned:



Table 31 Check Cinder Scheduler Status - Cinder Integrated

Verdict	Meaning
[PASSED]	The result of the check equals to the number of vCICs.
[FAILED]	The result of the check does not equal to the number of vCICs.
[WARNING]	Not applicable.

### 2.4.31 Check Cinder Volume Status

**Note:** Not applicable for single server platform.

The conditions of this check depend on whether OpenStack Cinder services are integrated in CEE. Cinder is integrated if the `configure_cinder_ericsson` attribute of the `ericsson_openstack_config` Fuel plugin is set to `true` or not defined in the `config.yaml`.

If Cinder services are not integrated, the following verdicts can be returned:

Table 32 Check Cinder Volume Status - Cinder Not Integrated

Verdict	Meaning
[PASSED]	The result of the check is 0.
[FAILED]	The result of the check is not 0.
[WARNING]	Not applicable.

If Cinder services are integrated, the following verdicts can be returned:

Table 33 Check Cinder Volume Status - Cinder Integrated

Verdict	Meaning
[PASSED]	The result of the check is 1.
[FAILED]	The result of the check is not 1.
[WARNING]	Not applicable.

### 2.4.32 Check Ethernet Statistics on All Compute Hosts

Table 34 Check Ethernet Statistics on All Compute Hosts

Verdict	Meaning
[PASSED]	The threshold for acceptable “lost” and “error” frames is 0.002%.



Verdict	Meaning
[FAILED]	The output shows more than 0.002% “lost” and “error” frames.
[WARNING]	Not applicable.

### 2.4.33 Check Ethernet Statistics on All vCIC Hosts and Compute Hosts

Table 35 Check Ethernet Statistics on All vCIC Hosts and Compute Hosts

Verdict	Meaning
[PASSED]	The threshold for acceptable “dropped” and “error” frames (RX-DRP,TX-DRP,RX-ERR,TX-ERR) is 0.01%.
[FAILED]	Not applicable.
[WARNING]	The output shows more than 0.01% “dropped” and “error” frames.

### 2.4.34 Check RabbitMQ Cluster Status

Table 36 Check RabbitMQ Cluster Status

Verdict	Meaning
[PASSED]	The output shows the number of vCIC nodes and does not contain ERROR or WARNING.
[FAILED]	The output shows a different number than the number of vCIC nodes or contains ERROR or WARNING, or both.
[WARNING]	Not applicable.

#### Implications

Perform further checks if you suspect that the queue disappears when, for example, OpenStack components are not running correctly. For example, one Ceilometer agent queue must be present for each compute host, and one Nova compute queue for each host.

## 2.5 Modules with [SKIPPED] Verdict

This section covers the modules that are included in the script for data collection and troubleshooting purposes.



The script executes these commands and prints the results in the log file without displaying a verdict in the console output.

The following modules return a [SKIPPED] verdict:

- ETHERNET INTERFACES THAT ARE IN UP STATE
- SERVICE STATUS
- RABBITMQ STATUS
- MONGODB SERVICE STATUS
- OVS SERVER STATUS
- OVS SWITCH STATUS
- NEUTRON NET-LIST
- NEUTRON AGENT LIST IN SDN
- RUN CEE\_SDNC\_SANITY\_CHECK\_FILE
- MEMORY USAGE



## 3 Report Problems

Collect data related to the occurring problems, refer to the [Data Collection Guideline](#).

For persistent problems, contact the next level of support.