ERICSSON

High CPU Load

Cloud Execution Environment

OPERATING INSTRUCTIONS

1/1543-CRA 119 1861/5 Uen D

Contents

1 Introduction

This instruction concerns alarm handling.

1.1 Alarm Description

The alarm is issued by the Managed Object (MO) `Host`.

The alarm is sent if the CPU workload, CPU utilization, or both exceed the threshold configured in the monitoring tool for triggering the alarm. The alarm ceases if the triggering measures go under the threshold configured for ceasing.

Note: Generally, the configured threshold for ceasing is lower than the threshold for triggering the alarm.

The possible alarm causes and the corresponding fault reasons, fault locations, and impacts are described in Table 1.

Table 1 Alarm Causes

Alarm Cause	Description	Fault Reason	Fault Location	Impact
The CPU workload exceeds the configured threshold.	The alarm is sent if the CPU workload or CPU utilization or both exceed the configured threshold.	• Software issue • The volume of tasks exceeds the maximum CPU capacity	Compute node or vCIC node.	The system capacity can be degraded causing loss of payload.
The CPU utilization exceeds the configured threshold.				
Both the CPU workload and utilization exceed the configured thresholds.				

Note: The *High CPU Load* alarm can appear as a result of network disturbances, or a maintenance activity on infrastructure or application level. If a maintenance activity is ongoing, wait until it is completed and five additional minutes.

The alarm attributes are listed in Table 2.

Table 2 Alarm Attributes

Attribute Name	Attribute Value
Major Type	`193`
Minor Type	`2031688`
Managed Object Class	`Host`

Managed Object Instance	`Region=<region_name>,`
	`Equipment=1,`
	`Host=<name>`

Specific Problem	`High CPU load`
Event Type	`equipmentAlarm (5)`
Probable Cause	`systemResourcesOverload (207)`
Additional Text	`The average load per CPU or the CPU utilization or both exceeded the configured thresholds during the measuring period;uuid=<HW_UUID_of_corresponding_server>`(1)
Severity	• `MINOR (5)`: The average CPU load is higher than 1.1 and the CPU utilization is under 80% for the measuring time period. • `MAJOR (4)`: The average CPU load is higher than 1.2 and the CPU utilization is over 80%, or the alarm is MINOR for at least 10 minutes. • `CRITICAL (3)`: The average CPU load is higher than 1.3 and CPU utilization is over 95%, or the alarm is `MAJOR` for at least 10 minutes.

(1) The format of this field is expected to change in CEE R6.

1.2 Prerequisites

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

1.2.1 Documents

The following documents are used in the procedure:

- *Data Collection Guideline*
- *Performance Management Northbound API*
- *Region Expansion*

1.2.2 Tools

No tools are required.

1.2.3 Conditions

- No ongoing maintenance activities on application level are assumed.
- SSH credentials for vCIC node and compute node are available.

2 Procedure

This section describes the procedure to follow when this alarm is received.

Based on the severity indicated in the alarm text, continue with the relevant section:

- If the severity is `MINOR`, continue with Section 2.1 on page 3.
- If the severity is `MAJOR` or `CRITICAL`, continue with Section 2.2 on page 3.

2.1 Severity MINOR

If the alarm severity is `MINOR`, do the following at the maintenance center:

1. Check if any related alarms are active. Act on any related alarms.
2. Wait 10 minutes for the alarm to cease.
   - If this alarm ceases, exit this procedure.
   - If the alarm severity increases to `MAJOR` or `CRITICAL`, continue with Section 2.2 on page 3.

Note: The Graphical User Interface (GUI) of the Zabbix monitoring tool or the performance management northbound API shows the actual CPU load and utilization, see Section 3 on page 7.

2.2 Severity MAJOR and CRITICAL

If the alarm severity is `MAJOR` or `CRITICAL`, continue with the relevant section depending on the type of the reported node:

- If the alarm is related to a **compute node**, continue with Section 2.2.1 on page 4.
- If the alarm is related to a **vCIC node**, continue with Section 2.2.2 on page 6.

2.2.1 Procedure for Compute Nodes

Do the following at the maintenance center:

1. Perform either of the following steps:
   - Log in to a vCIC using SSH:

     ssh *<admin-user>@<vcic_address>*
   - Or log in to Atlas using SSH:

     ssh *<admin-user>@<atlas_address>*
2. Investigate the total resource use on the available nodes. Use the below commands:
   nova hypervisor-stats
   nova host-describe *<hostid>*
   - If there are not enough resources in the region or if they are too fragmented to move VMs, refer to *Region Expansion* to install additional compute servers and increase the number of compute nodes. Exit this procedure.
   - If there are enough resources to migrate VMs, start migrating to decrease CPU load or CPU utilization or both.
     - In case of MAJOR severity, start with the VM that is using the least amount of CPU resource on the node issuing the alarm.

       Note: The CPU usage of each VM can be checked in the output of the **nova** command.
     - In case of CRITICAL severity, start migrating VMs immediately. Migrate at least half of the VMs to decrease the CPU load or utilization.

     Note: Never migrate a vCIC.
   - Migrate the selected VMs to a node with available CPU resources, if they can be migrated. Use the below command:
     nova migrate *<server>*

     Verify the migration with the command:
     nova resize-confirm *<server>*

Note: VM migration will cause a VM restart.

3. Check the actual CPU load and utilization either by using the Performance Management Northbound API or the GUI of the Zabbix monitoring tool as described in Section 3 on page 7.

   - If the CPU load or utilization or both are high, continue with migrating the VMs.
   - If the CPU load or utilization reached the normal level, continue with Step 4.

4. Wait 10 minutes, then check the active alarm list and perform the relevant action:

   - If the alarm has ceased, exit this procedure.
   - If the alarm remains, migrate all remaining VMs from the node issuing the command:
     nova migrate *<server>*

     Verify the migration with the below command:

     nova resize-confirm *<server>*

5. If all VMs have been migrated from the node:

   a. Log in to the node by using SSH:

      ssh *<admin-user>@<node_address>*

      If logging in was not possible, continue with Step 7.

   b. If logging in was successful, collect troubleshooting data as described in the *Data Collection Guideline*. For alarm-specific logs, refer to the Table *Data Collection for Alarms and Alerts* in the *Data Collection Guideline*.

   c. Restart the node by using the command:
      reboot

6. Wait 15 minutes for the restart to complete.

   - If the alarm ceases, exit this procedure.
   - If the alarm reappears when the node has been restarted and VMs are running on the node, run the **check config** command on the vCIC to collect log files and perform data collection, as described in the *Data Collection Guideline*.

7. Consult the next level of maintenance support. Attach the previously collected `sosreport` or the screenshot of the running processes to the customer service request. Further actions are outside the scope of this instruction.

8. The job is completed.

2.2.2 Procedure for vCIC Nodes

Do the following at the maintenance center:

1. Log in to the node using SSH:

   ssh *<admin-user>@<vcic_address>*

   a. If logging in was not possible, continue with Step 3.

   b. If logging in was successful, collect troubleshooting data as described in the *Data Collection Guideline*. For alarm-specific logs, refer to the Table *Data Collection for Alarms and Alerts* in the *Data Collection Guideline*.

   c. Check if the other two vCICs are running.

      - If both of the two other vCICs are running, restart the node with the following command:
        reboot

      - Else, continue with Step 3.

2. Wait 15 minutes for the restart to complete.

   - If the alarm ceases, do the following:

     – Log in to another vCIC using SSH:

       ssh *<admin-user>@<vcic_address>*

     – Check that all three vCIC nodes are up in normal operation by issuing the command:
       crm status

       Verify that the response in the line starting with `Online:` contains all three vCIC nodes:
       `Online: [cic<id> cic<id> cic<id>]`

       - If any of the three vCICs is not running, continue with Step 3.

       - If both of the other two vCICs are running, exit the procedure.

   - If the alarm reappears when the node has been restarted and VMs are running on the node, run the following command:
     check config
     on the vCIC to collect log files and perform data collection, as described in the *Data Collection Guideline*.

     Continue with Step 3.

3. Consult the next level of maintenance support. Attach the previously collected `sosreport` or the screenshot of the running processes to the customer service request. Further actions are outside the scope of this instruction.

4. The job is completed.

3 Checking CPU Load and Utilization

To check the CPU load and CPU utilization, use either of the following tools:

- The GUI of the Zabbix monitoring tool, see Section 3.1 on page 7.
- The performance management northbound API, see Section 3.2 on page 7.

3.1 Zabbix Monitoring Tool

To access the Zabbix monitoring tool, use the address:

`https://192.168.2.22/zabbix`

The user group, user name, and password can be configured before deployment by setting the correct parameters in the `config.yaml` file. Refer to the *Configuration File Guide*.

The default user group is `CEEUserGroup`, the default user is `ceeuser`. The default password is generated during deployment, and can be found in `/etc/openstack_deploy/user_secrets.yml` under `zabbix_cee_user_password` on vFuel.

3.2 Performance Management Northbound API

To check CPU load in the performance management northbound API, refer to the section *Monitoring API* in the *Performance Management Northbound API*.