

High Memory Utilization

Cloud Execution Environment

OPERATING INSTRUCTIONS

Copyright

© Ericsson AB 2016–2018. 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

1	Introduction	1
1.1	Alarm Description	1
1.2	Prerequisites	2
2	Procedure	2
2.1	Severity MINOR	3
2.2	Severity MAJOR and CRITICAL	3



High Memory Utilization



1 Introduction

This instruction concerns alarm handling.

1.1 Alarm Description

The High Memory Utilization alarm is issued by the Managed Object (MO) Host.

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 memory utilization is high.	This alarm is sent when the memory utilization is high and exceeds a set threshold level.	The memory utilization is higher than expected, more memory is needed.	This is a dimensioning and configuration fault.	The system capacity can be degraded when the memory utilization is exceeds the threshold.

Note: The High Memory Utilization alarm can appear as a result of the maintenance activity.

The alarm attributes are listed in Table 2.

Table 2 Alarm Attributes

Attribute Name	Attribute Value
Major Type	193
Minor Type	2031689
Managed Object Class	Host
Managed Object Instance	Region=<region_name>, Equipment=1, Host=<name>
Specific Problem	High memory utilization
Event Type	equipmentAlarm (5)
Probable Cause	systemResourcesOverload (207)



Attribute Name	Attribute Value
Additional Text	Available memory fell below the threshold, alarm is cleared when the available memory exceeds the threshold;uuid=<hw_uuid_of_corresponding_server>
Severity	<ul style="list-style-type: none">• MINOR (5): When the available memory is less than 1 GB and more than 0.5 GB.• MAJOR (4): When the available memory is less than 0.5 GB and more than 0.25 GB.• CRITICAL (3): When the available memory is less than 0.25 GB.

1.2 Prerequisites

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

1.2.1 Documents

The following document is used in this procedure:

— Data Collection Guideline

1.2.2 Tools

No tools are required.

1.2.3 Conditions

Before starting this procedure, ensure that the following conditions are met:

- 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.
3. Check the available memory by executing the following command:

```
/etc/zabbix/scripts/check_free_memory.sh
```

- If the available memory is more than 1 GB, the alarm has ceased, exit this procedure.
- If the available memory is less than 1 GB and more than 0.5 GB, the alarm severity remains MINOR. Return to Step 2.
- If the available memory is less than 0.5 GB and more than 0.25 GB, the alarm severity is MAJOR. Continue with Section 2.2 on page 3.
- If the available memory is less than 0.25 GB, the alarm severity is CRITICAL. Continue with Section 2.2 on page 3.

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 host:

- If the alarm is related to a compute node, continue with Section 2.2.1 on page 3.
- If the alarm is related to a vCIC node, continue with Section 2.2.2 on page 5.

2.2.1 Procedure for Compute Node

Do the following at the maintenance center:

1. Log in to the affected compute node using SSH:

```
ssh <compute_name>
```
2. Find out which VM is using the largest amount of memory on the node:



```
ps -eo pid,rss,cmd --sort rss | grep [q]emu.*instance⇒
- | awk '{print "PID:$1 " MEMORY:$2 " NAME:$6}'
```

Example output:

```
PID:43397 MEMORY:59380 NAME:instance-000002f9
PID:25380 MEMORY:59560 NAME:instance-000002f6
PID:47210 MEMORY:60136 NAME:instance-000002fc
```

3. Log out from the compute node.

4. Perform either of the following steps:

- Log in to a vCIC by using SSH:

```
ssh <admin_user>@<vcic_address>
```

- Or log in to Atlas by using SSH:

```
ssh <admin_user>@<atlas_address>
```

5. Execute `source openrc`.

6. Use the following command to obtain the instance ID:

```
nova list --fields name,status,host,instance_name|⇒
grep <affected_host>
```

7. If the VMs High Availability (HA) policy allows it to be migrated, migrate the VM to another node which has sufficient memory space. Use the command `nova migrate <server>` with the appropriate parameters.

8. Wait for 10 minutes and check the active alarm list.

- If the alarm persists, continue removing VMs from the node (continuing always with the VM that takes up the most memory on the node) by repeating Step 1 to Step 7.
- If there are no more VMs left on the node (or the only VMS on the node are VMs that cannot be migrated due to their HA policy), and the alarm still persists, continue with Step 9.
- If the alarm ceases, continue with Step 11.

9. Restart the node:

```
reboot
```

Note: Never reboot a compute node running vFuel.

Wait 15 minutes for the restart process to complete.

10. Log in to the affected compute node using SSH:



```
ssh <compute_name>
```

- If logging in is not possible, contact the next level of maintenance support and exit this procedure.
- If logging in is successful, continue with Step 11.

11. Collect troubleshooting data as described in the [Data Collection Guideline](#).
12. Consult the next level of maintenance support. Further actions are outside the scope of this instruction.
13. The job is completed.

2.2.2

Procedure for vCIC Node

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, contact the next level of maintenance support and exit this procedure.
 - b. If logging in was successful, collect troubleshooting data as described in the [Data Collection Guideline](#).
 - c. Restart the node with the following command:
reboot
2. 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 the VMs are running on the node, proceed with Step 3.
 3. Consult the next level of maintenance support. Further actions are outside the scope of this instruction.
 4. The job is completed.