

High Local Disk 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
3	Check Disk Utilization	4
3.1	Performance Management Northbound API	4



High Local Disk Utilization



1 Introduction

This instruction concerns alarm handling.

1.1 Alarm Description

The High Local Disk 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 local disk utilization is high.	The alarm is sent when local disk utilization exceeds the hard-coded threshold level. ⁽¹⁾	The local disk utilization is higher than expected, more disk space is needed.	This is a dimensioning and configuration fault.	The system capacity can be degraded causing loss of payload.

(1) The alarm is raised when disk utilization exceeds 90% and ceases when utilization drops below 80%.

Note: The High Local Disk Utilization alarm can appear as a result of network disturbances or a maintenance activity. 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	2031690
Managed Object Class	Host
Managed Object Instance	Region=<region_name>, Equipment=1, Host=<name>
Specific Problem	High local disk utilization
Event Type	equipmentAlarm (5)
Probable Cause	resourceAtOrNearingCapacity (100541)



Attribute Name	Attribute Value
Additional Text	Measured value exceeded 90% on <file_system>, alarm is cleared when it goes below 80%;uuid=<hw_uuid_of_corresponding_server>
Severity	CRITICAL (3)

1.2 Prerequisites

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

1.2.1 Documents

Not applicable.

1.2.2 Tools

No tools are required.

1.2.3 Conditions

Before starting this procedure, ensure that SSH credentials for vCIC node and compute node are available.

2 Procedure

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

1. Check if any related alarms are active. Act on any related alarms.
2. Wait five minutes and check if the alarm has ceased. If this alarm ceased, exit this procedure.
3. Determine which partition is full, by running the following command:
df
Write down which partition is full.

Printout example:



CIC:

```
root@cic-1:/var/log# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
udev	13G	12K	13G	1%	/dev
tmpfs	2.5G	704K	2.5G	1%	/run
/dev/dm-4	50G	5.7G	41G	13%	/
none	4.0K	0	4.0K	0%	/sys/fs/cgroup
none	5.0M	0	5.0M	0%	/run/lock
none	13G	39M	13G	1%	/run/shm
none	100M	0	100M	0%	/run/user
/dev/vda3	196M	43M	144M	23%	/boot
/dev/mapper/logs-log	48G	47G	1.7G	97%	/var/log
/dev/mapper/image-glance	40G	1.8G	38G	5%	/var/lib/glance
/dev/mapper/mysql-root	40G	7.9G	30G	22%	/var/lib/mysql
/dev/mapper/mongo-mongodb	69G	13G	53G	20%	/var/lib/mongo

Compute:

```
root@compute-0-5:/var/log# df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
udev	5.4G	12K	5.4G	1%	/dev
tmpfs	6.3G	5.2M	6.3G	1%	/run
/dev/dm-2	50G	2.2G	45G	5%	/
none	4.0K	0	4.0K	0%	/sys/fs/cgroup
none	5.0M	0	5.0M	0%	/run/lock
none	32G	4.0K	32G	1%	/run/shm
none	100M	0	100M	0%	/run/user
/dev/sdb3	196M	53M	134M	29%	/boot
/dev/mapper/logs-log	40G	39G	1.9G	96%	/var/log
/dev/mapper/vm-nova	1.1T	30G	996G	3%	/var/lib/nova

4. Log in to the node using SSH:

```
ssh <admin_user>@<node_address>
```

5. Collect troubleshooting data as described in the [Data Collection Guideline](#).
6. Contact next level of maintenance support immediately.
7. The job is completed.



3 Check Disk Utilization

To check the disk utilization, use the performance management northbound API, see Section 3.1 on page 4.

3.1 Performance Management Northbound API

To check disk utilization in performance management northbound API, refer to the section Monitoring API in the [Performance Management Northbound API](#).