

VM Unavailable

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	3
2.1	Actions	3
3	Additional Information	4
3.1	Workaround for CPU Pinning Overlap	4
3.2	Cease the Alarm Manually	5



VM Unavailable



1 Introduction

This instruction concerns alarm handling.

1.1 Alarm Description

The VM Unavailable alarm is issued by the Managed Object (MO) VM.

The alarm is issued for a Virtual Machine (VM) when the compute node, where the VM is running, becomes unavailable. The alarm is issued for all VMs that are hosted on the compute node failing the periodical availability test three consecutive times. Alarms are also issued for VMs which are in ERROR state. Only VMs that have the High Availability (HA) policy set to ha-offline are evacuated.

The severity of the alarm is MINOR.

The possible alarm causes and fault locations are explained in Table 1.

Table 1 Alarm Causes

Alarm Cause	Description	Fault Reason	Fault Location	Impact
The evacuation of a VM has started. Or the VM became unavailable and the HA policy was not set to ha-offline.	The VM became unavailable due to the failure of the compute node it was running on. This can be a temporary fault if the VM has a HA policy defined. In this case the Cloud Execution Environment (CEE) starts the automated recovery procedure. If no HA policy was defined for the VM, it can be a permanent fault, as the automated recovery process will not start.	The compute node, on which the VM is running, has failed the periodical availability test three consecutive times.	Compute node	The VM is temporarily or permanently not available for the tenant.

The alarm attributes are listed in Table 2.

Table 2 Alarm Attributes

Attribute Name	Attribute Value
Major Type	193
Minor Type	2031702
Managed Object Class	VM



Attribute Name	Attribute Value
Managed Object Instance	Region=<name_of_the_region>, CeeFunction=1, Tenant=<tenant_uuid>, VM=<vm_uuid>
Specific Problem	VM Unavailable
Event Type	other (1)
Probable Cause	underlyingResourceUnavailable (165)
Additional Text	{"host": <hostname_of_the_node>}
Severity	MINOR (5)

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

No conditions.



2 Procedure

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

2.1 Actions

In case of managed VMs (when `ha-policy` is defined as `ha-offline` or `managed-on-host`), no actions are required when the VM Unavailable alarm is issued, since the automatic recovery is performed by CEE through VM evacuation. The alarm ceases automatically when the VM becomes available.

In case of unmanaged VMs (when `ha-policy` is set to `unmanaged` or is not defined), Continuous Monitoring High Availability (CM-HA) does not detect that the VM is available after a manual recovery, performed with the commands `nova redefine` and `nova start`. In this case, the alarm will remain. To cease the alarm manually, see Section 3.2 on page 5.

Note: In case the `nova-compute` service is down, Nova does not have information on VM availability and therefore cannot update VM status. For example, a VM that was `ACTIVE` before a `nova-compute` service failure will keep this status in the printouts, even though the VM Unavailable alarm shows it is offline. The status is updated when the `nova-compute` service is up and running again.

If the evacuation of the VM fails and the VM cannot be recovered, a VM Evacuation Failed alarm is issued. For more information about the VM Evacuation Failed alarm, refer to [VM Evacuation Failed](#).

If policy settings prevent evacuation, the VM will not be evacuated. However, if the `ha-policy` is set to `managed-on-host`, the VM will be recovered when the host is online again.



3 Additional Information

The alarm is ceased for a VM in the following situations:

- The evacuation is finished, and the VM became available again.
- Or the evacuation was unsuccessful, but the VM is recovered manually and became available again.
- Or the VM is recovered automatically after the host is up again.

3.1 Workaround for CPU Pinning Overlap

In case of VM evacuation, the physical CPU allocation on the original host for the VM is replicated on the destination host. If another VM is using these CPUs in the destination host, overlapping CPU allocation occurs. In this case, intersecting CPU allocation must be detected and the affected VMs must be migrated. After migration, new CPUs are allocated for the VM.

Do the following:

1. Log on to vFuel. For more information, refer to the [CEE Connectivity User Guide](#).
2. Change to the directory `/usr/share/ericsson-orchestration/scripts` using the following command:
`cd /usr/share/ericsson-orchestration/scripts`
3. Execute the script using the following command:
`./check-cpupin.py`

An example of the output is the following:

```
+-----+-----+-----+
| Instance A | Instance B | Intersected CPUs |
+-----+-----+-----+
| instance-0000028d | instance-0000028a | [10] |
+-----+-----+-----+
```

Identify the instances with overlapping CPU allocation from the printout and save the names of the affected instances for a later step.

4. Log on to one of the vCICs. For more information, refer to the [CEE Connectivity User Guide](#).
5. Find the affected VM by executing the following command:
`nova list --fields=name,instance_name --all-tenants`

An example of the output is the following:



ID	Name	Instance
329808c2-ecfd-40e9-825b-e0041d057811	vAPP1	instance-
6b772556-c446-4291-aa26-f8a663227ed0	vAPP2	instance-

Identify the affected VMs based on the instance names recorded in Step 3. Save the corresponding VM name or VM ID for a later step.

- Identify the VM that was evacuated in relation to the alarm by executing the following command:
nova instance-action-list <vm_name_or_id>
 where <vm_name> refers to the name of the affected VM recorded from the printout of Step 5.

An example of the output is the following:

Action	Request_ID	Message	Start_Time
create	req-7f8cd389-7f50-4efb-8dc6-bad6f0582046	-	2017-05-15T13:42:40.000000
evacuate	req-b73b2002-c62d-48fd-978c-c782da97571f	-	2017-05-15T13:50:13.000000

Identify the evacuated VM based on the printout. If several affected VMs have the evacuate action in the command printout, identify the VM evacuated in relation to the alarm based on the timestamp.

- Migrate the affected VM by executing the following command:
nova migrate <vm_name_or_id>

Wait until migration finishes, and check that VM status is VERIFY_RESIZE:

```
nova show <vm_name_or_id> | grep status
```

- Confirm the migration by executing the following command:
nova resize-confirm <vm_name>

Note: It can take five minutes to run the resize-confirm command.

If there is an ERROR (Conflict) error, rerun the resize-confirm command.

3.2 Cease the Alarm Manually

To cease the VM Unavailable alarm manually, do the following:

- Log on to one of the vCICs. For more information, refer to the [CEE Connectivity User Guide](#).
- Remove the event from the CM-HA database by using the following command:

```
mysql -e 'delete from cmha.event where vm="<vm_uuid>"'
```



For <vm_uuid>, use the value indicated in the Managed Object Instance attribute of the alarm, see Table 2.

An example of the command is the following:

```
mysql -e 'delete from cmha.event where vm="b3592207-3153-41ff-8d81-c1aaf8564af4"'
```

3. Remove the alarm from the watchmen history by using the following command:

```
watchmen-client create-event --stateful --source "<managed_object_instance>" =>  
--major-type 193 --minor-type 2031702 --severity CLEARED --event-type other =>  
--probable-cause underlayingResourceUnavailable --specific-problem =>  
"VM Unavailable"
```

For <managed_object_instance>, use the values indicated in the Managed Object Instance attribute of the alarm, see Table 2.

In the command, the values are used in the following format:

```
--source "Region=<name_of_the_region>,CeeFunction=1=>  
,Tenant=<tenant_uuid>,VM=<vm_uuid>"
```

An example value is the following:

```
--source "Region=dc033,CeeFunction=1,Tenant=c190c72ae3=>  
ac4ff2a1597789d1f75339,VM=b3592207-3153-41ff-8d81-c1aaf=>  
8564af4"
```