

Virtualized CUDB Virtual Machine Recovery

OPERATING INSTRUCTION

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	Description	1
1.2	Target Groups	2
1.3	Revision Information	3
1.4	Typographic Conventions	4
2	Reboot and Rebuild the VM	5
2.1	Reboot the VM	5
2.2	Rebuild the VM	6
3	Actions for Planned Infrastructure Maintenance Activities	9
3.1	Identify All Affected VMs	9
3.2	Recovery of Multiple VMs in Parallel	11
3.3	Cloud Administration Actions After Infrastructure Maintenance	14
3.4	Prepare the VMs for Operation	14
	Glossary	17
	Reference List	19





1 Introduction

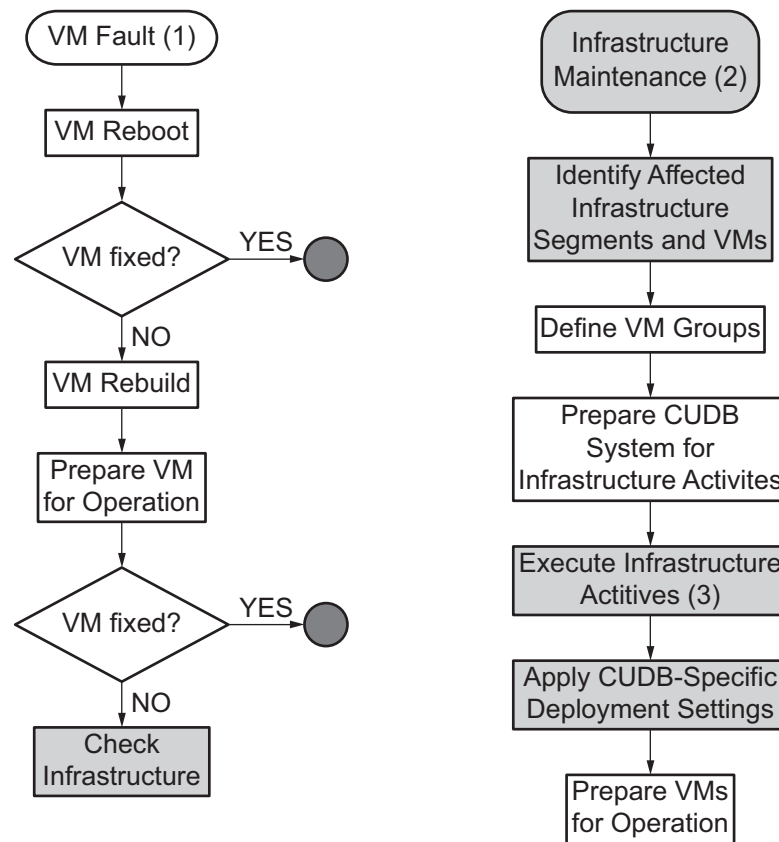
This document describes how to recover a Virtual Machine (VM) in an Ericsson Centralized User Data Base (CUDB) node deployed on a cloud infrastructure.

1.1 Description

This Operating Instruction (OPI) describes how to recover a VM in a virtualized CUDB node. Specifically, it describes the following procedures:

- Manually recovering the VM by means of rebooting or rebuilding it.
- Preparing the CUDB node and system to gracefully handle and recover from a planned infrastructure maintenance activity.

The major steps of the recovery procedures listed above are shown in Figure 1.



(1) Where reboot is recommended in CUDB procedures, or after other procedures have been exhausted.

(2) Triggered by VM faults, or other reasons.

(3) Aligned with VM grouping.

 Cloud Infrastructure Level

Figure 1 VM Recovery Procedures

1.2 Target Groups

This document is intended for system administrators operating CUDB systems. For some of the actions described in the document, cloud administration role is also required. The cloud administrator is the cloud service provider who delivers the cloud service and executes required actions on the cloud infrastructure.



1.3 Revision Information

Rev. A

Initial release.

Rev. B

Editorial changes only.

Rev. C

Other than editorial changes the document has been revised as follows:

- Section 2.2 on page 6: Added a note regarding restoring stored procedures after blade replacement.

Rev. D

Other than editorial changes the document has been revised as follows:

- Section 3.4.1 on page 14: Updated Step 3 in the SC VM preparation procedure.
- Section 3.4.2 on page 15: Removed obsolete step.

Rev. E

Other than editorial changes the document has been revised as follows:

- Section 2.2 on page 6: Updated description.

Rev. F

Other than editorial changes the document has been revised as follows:

- **Preparations for PLDB or DSG VM Recovery:** Removed section.
- Section 3.2.2 on page 13: Removed former **Step 2**.
- Section 3.3 on page 13: Updated procedure.
- Section 3.4.2 on page 15: Removed case from the list of applicable steps to prepare the recovered PLDB or DSG VM for operation.

Rev. G

Other than editorial changes the document has been revised as follows:



- Section 3.1 on page 9: Added attention.

1.4 Typographic Conventions

Typographic conventions can be found in the following document:

- *Typographic Conventions*



2 Reboot and Rebuild the VM

If a VM is having issues and all applicable CUDDB recovery procedures outlined in *CUDDB Troubleshooting Guide*, Reference [1] have been performed, but the VM still did not recover, it can be rebooted from the cloud infrastructure (see Section 2.1 on page 5). If rebooting does not solve the issue, or if the VM must be reinstalled, the VM can also be rebuilt (see Section 2.2 on page 6).

2.1 Reboot the VM

Perform the following steps to reboot the VM from the cloud infrastructure.

Attention!

Any custom data not saved in the `/local` or `/local2` folders will be lost after rebooting.

In case of using Cloud Execution Environment (CEE), follow the below steps to reboot the VM:

1. Login to the Atlas Dashboard.
2. Select the appropriate project in the **Current Project** field, then select **Project** in the **View** field.
3. Choose the **Instances** category.
4. Identify the VM to reboot.
5. In the **Actions** column of the identified VM, select the action **Soft Reboot Instance** to use graceful shutdown, or the action **Hard Reboot Instance** to use non-graceful shutdown. Refer to the “Openstack End User Guide” or the “Atlas Dashboard End User Guide” documents of the CEE Customer Product Information (CPI) for more details.
6. While the reboot process is ongoing, the **Status** column of the instance will show `REBOOT`. Once it becomes `ACTIVE`, the processes in the VM will begin to start up.
7. If the issue still persists after the VM has fully started up, rebuild the VM by performing the steps of Section 2.2 on page 6.

Refer to the “Atlas Dashboard End User Guide” document in the CEE documentation for more information on how to reboot VMs if using CEE. In case



of using a different cloud solution, refer to the solution-specific documentation for more information.

2.2 Rebuild the VM

VMs are rebuilt during node installation to ensure the automatic recovery of System Controllers (SCs), or when any VM is reinstalled.

Attention!

All data on the VM will be lost when it is about to be reinstalled.

In case of using CEE, follow the steps below to rebuild the affected VM:

1. Login to the Atlas Dashboard.
2. Select the appropriate project in the **Current Project** field, then select **Project** in the **View** field.
3. Choose the **Instances** category.
4. Identify the VM to rebuild.
5. In the **Actions** column of the identified VM, select the action **Rebuild Instance**.
6. In the **Rebuild Instance** window, choose the **Payload** image option from the **Select Image** drop-down box to allow the VM to boot from network.
7. Once the image is chosen, select **Rebuild**.
8. While the rebuild process is ongoing, the **Status** column of the instance will show **REBUILD**. Once it becomes **ACTIVE**, the processes in the VM will begin to start up.
9. Depending on the type of the VM, perform the applicable steps below to prepare the recovered VM for operation:
 - If the rebuilt VM is an SC, wait until the synchronization between the SCs is completed. Use the following command to check the synchronization status:

```
cat /proc/drbd
```

Once the synchronization is finished, any custom `crontab` jobs and their definitions (or similar tasks) which are not deployed by default in CUDB will be lost. If necessary, redeploy them after the procedure is completed.



- If the rebuilt VM belongs to a Data Store Unit Group (DSG), and replication issues are detected after all the processes in the VM started up, perform a combined unit data backup and restore as described in the “Performing Combined Unit Data Backup and Restore” section of *CUDB Backup and Restore Procedures*, Reference [2].
- If the rebuilt VM belongs to a PLDB, execute the following command after the Network Databases (NDBs) are started, and the database cluster server connections are OK:

```
cudbPrepareStore --pl
```

Note: After finishing the rebuild procedure, the stored procedures are not restored so it is recommended to recreate them with the following command:

```
cudbManageStore -p -o restorestoredprocedures
```

Refer to the “Atlas Dashboard End User Guide” document in the CEE documentation for more information on how to rebuild VMs if using CEE. In case of using a different cloud solution, refer to the solution-specific documentation for more information.





3 Actions for Planned Infrastructure Maintenance Activities

This section describes how to prepare the CUDB node and system to gracefully handle and recover from a planned maintenance activity on the infrastructure level. The major steps of this procedure are as follows:

1. Identifying all affected VMs (see Section 3.1 on page 9).
2. Preparing the affected VMs for recovery (see Section 3.2 on page 11).
3. Performing post-maintenance cloud administration actions (see Section 3.3 on page 13).
4. Preparing the recovered VMs for operation (see Section 3.4 on page 14).

3.1 Identify All Affected VMs

An infrastructure segment can host one or several VMs. If a maintenance activity on the infrastructure level is planned, find out which infrastructure segment is affected to identify which VMs must be prepared for it. In most cases, the affected infrastructure segment is known ahead, and based on that information, the cloud administrator can easily identify all affected VMs. However, it can happen that an unexpected failure in the infrastructure occurs, and the specific infrastructure segment has yet to be identified, based on the VMs with faults in their behavior.

Attention!

If the identified VM to be recovered belongs to a DSG replica, that DSG replica will result in a degraded state. If a DSG master replica is degraded due to hardware or software failure, the system selects a new master for the DSG provided that at least one other slave replica of the DSG is available, not degraded and with a replication delay below than 3 seconds; otherwise smooth mastership change will not happen until the previous conditions are met. Therefore, if the master remains in the degraded replica, the infrastructure maintenance or VM recovery must be performed in low traffic periods if possible. If not, repair the other replica so the smooth mastership change occurs and the mastership is removed from the degraded one.

Note also that if Automatic Mastership Change (AMC) is disabled, mastership will not be returned to preferred location until AMC is enabled or mastership change is done manually.



3.1.1 Identify Affected Infrastructure Segment

If the infrastructure segment has not yet been identified, it can be identified by the infrastructure position of the faulty VMs. To do so, provide the VM instance name(s), the VM instance Universally Unique Identifier(s) (UUIDs), or both to the cloud administrator. This information can be gathered as follows:

- The **instance name** can be obtained from the cloud infrastructure.
- The **instance UUID** can be obtained either from the CUDB system, if an alarm was raised for that VM (from the alarm description), or from the cloud infrastructure.

In case of using CEE, identify the instance name, the instance UUID, or both as described below:

- ☐ The **instance name** of the VM can be obtained from the **Instances** page of the Atlas Dashboard. Its format is `<tenant>_<nameOfFailingVirtualMachine>`. For example, the SC_2_1 VM instance name of the CUDB_VNF01 tenant would be CUDB_VNF01_SC_2_1.
- ☐ The **instance UUID** of the VM can be obtained from the Atlas Dashboard by choosing the instance name identified above, and checking the ID value under the **Information** section of the **Overview** tab.

Once one or both of the above data is available, contact the cloud administrator, and provide them the VM instance name(s), instance UUID(s), or both.

3.1.2 Cloud Administration Actions to Identify Infrastructure Position

After obtaining the VM(s) instance name, instance UUID, or both (as described in Section 3.1.1 on page 9), cloud administrators must identify the infrastructure position. Depending on the user interface used, perform the applicable procedure described below.

In case of using the Atlas Dashboard GUI, identify the VM infrastructure position with the following steps:

1. Login to the Atlas Dashboard.
2. Choose the **Instances** category, and search for the instance using the provided instance name.
3. Look for the compute host name, which is located left to the name of the provided instance name in the **Host** column. Ignore the `.domain.tld` suffix.

In case of using OpenStack command line tools, perform the following steps in a Cloud Infrastructure Controller (CIC):

1. Execute the below command to show the details of the specific VM:



```
nova show <instance_UUID>
```

2. Check the command output for the infrastructure position. The information is stated under the `OS-EXT-SRV-ATTR:host` field. Ignore the `.domain.tld` suffix.

In case of using a cloud solution other than CEE, refer to the solution-specific documentation for more information on how to identify the infrastructure position.

3.1.3 Cloud Administration Actions to Identify Affected VMs

The cloud administrator can identify the affected VMs on the cloud infrastructure level from the cloud infrastructure segment. To do so, identify which VMs the infrastructure segment hosts.

In case of using CEE, follow the steps below to identify the VMs hosted by the specific infrastructure segment:

1. Login to the Atlas Dashboard.
2. Choose the **Compute Environment** category, and search for the identified compute host.
3. Choose the identified compute host to see the related details.
4. Check the instances running on the selected compute host. This information can be found under the **Instances** view of the opened compute host details.
5. Provide the instance names to the requesting tenants.

In case of using a cloud solution other than CEE, refer to the solution-specific documentation for more information on how to identify the affected VMs.

3.2 Recovery of Multiple VMs in Parallel

This section describes how to recover multiple VMs in parallel on a virtualized CUDB node.

3.2.1 VM Groups

Because of the virtualized CUDB node infrastructure deployment on host aggregates, multiple VMs can be affected by an infrastructure maintenance. An affected infrastructure segment (that is, “compute host”) can host either one SC, or one or several payload VMs for one CUDB node.



Do!

Because of the deployment rules described above, VMs of another virtualized CUDB node can be hosted on the same infrastructure segment, and can therefore also be affected. Take special care while identifying the affected VMs, and execute the recovery steps of Section 3.2.2 on page 13 on all affected virtualized CUDB nodes.

In the CUDB system, VMs are categorized into three distinct groups: SC, PLDB, and DSG. These groups can be further divided into groups of even-numbered and odd-numbered VMs. Considering the deployment of a typical virtualized CUDB node and the applied anti-affinity policies, the affected infrastructure segment is likely to host one of the following groups of VMs:

- SC_2_1
- SC_2_2
- Odd-numbered PLDB and odd-numbered DSG VMs.
- Even-numbered PLDB and even-numbered DSG VMs.

Note: The type of VMs in the last two groups can vary depending on the infrastructure availability during the deployment, but the redundancy over the infrastructure must be satisfied because of the applied anti-affinity policies. This means the following:

- The same infrastructure segment cannot host both VMs of the same DSG.
- The PLDB VMs must be evenly distributed. In other words, the majority of the PLDB VMs cannot be hosted by one infrastructure segment.

In case of a failure in multiple infrastructure segments (that is, more than one group of VMs of the above four groups must be recovered), consider the following rules for recovery:

- Do not execute recovery in parallel for different VM groups. Instead, recover one VM group at a time, with the priority as listed above.
- If the affected VMs belong to the same group, they can be recovered in parallel. To do so, perform the recovery in the following order:
 - First, recover fully one SC.
 - Continue with the parallel recovery of the odd-numbered PLDB and DSG VMs.



- Then, continue with the parallel recovery of the even-numbered PLDB and DSG VMs.

3.2.2 Execute Parallel Recovery

Do!

During VM recovery, always follow the order of groups exactly as listed in Section 3.2.1 on page 11. Deviations from the defined order can result in a major node outage.

Perform the following steps to recover multiple VMs belonging either to the same or a different VM group:

1. Identify all affected VMs inside the node to group them. Follow the steps of Section 3.1 on page 9 to do so.

Note: To ensure that the traffic handling capacity is enough during the recovery procedure, the number of VMs to replace in parallel should not exceed the configured value of the `redundancyLevel` attribute of the `CudbLdapAccess` class (refer to the “Class `CudbLdapAccess`” section of *CUDB Node Configuration Data Model Description*, Reference [3]. If this condition cannot be fulfilled (that is, the amount of VMs to recover is larger than the value of the `redundancyLevel` attribute), then it is recommended to perform the recovery only for the same number of VMs in parallel at a single time as the value of the `redundancyLevel` attribute. However, if recovery is performed during a low traffic period or in a maintenance window (when the degraded traffic handling capacity could still be enough), recovery may be executed in parallel for more VMs, than the value of the `redundancyLevel` attribute.

2. In case of recovering SC group(s) or PLDB group(s), force the applications to move their primary connections to another CUDB node. This applies in case primary connections are established, or if the SC or PLDB VMs are affected.
3. Execute the recovery of the VMs exactly in the order as specified in Section 3.2.1 on page 11, skipping any group which has no affected VMs.
4. If all recoveries have been finished, and more VMs were recovered than the value of the `redundancyLevel` attribute, then perform a rolling restart of the LDAP Front Ends (FEs) with the `cudbLdapFeRestart` command.



3.3 Cloud Administration Actions After Infrastructure Maintenance

In case the infrastructure maintenance activities include the re-deployment of any infrastructure segment (for example in case when server replacement is performed), the cloud administrator must check if any applicable CUDB-specific deployment preparations must be executed on the affected infrastructure segment(s) (such as verifying that the ARP spoofing protection is turned off).

In case of using CEE, the cloud administrator must execute the following steps to verify that the ARP spoofing protection is turned off:

1. Execute the following command on each infrastructure segment (that is, “compute host”):

```
grep prevent_arp_spoofing /etc/neutron/plugins/ml2/ml2_conf.ini
```

2. In case the parameter is set to `True`, the ARP spoofing protection is on. Turn it off by setting the `prevent_arp_spoofing` parameter to `False` in `/etc/neutron/plugins/ml2/ml2_conf.ini` on all compute nodes. Then, execute the following command:

```
service neutron-plugin-openvswitch-agent restart
```

3.4 Prepare the VMs for Operation

This section describes how to prepare the VMs for operation after the infrastructure maintenance activities have been performed.

Do!

Before continuing with the steps of Section 3.4.1 on page 14 or Section 3.4.2 on page 15, verify with the cloud administrator that the applicable CUDB-specific deployment preparations of Section 3.3 on page 13 have been executed.

Note: If the VM did not join the cluster automatically after the infrastructure maintenance has been performed, refer to the Release Notes to check if any manual action is needed for the recovery of the network connectivity of the VMs. Some actions can involve cloud administration.



3.4.1 Prepare SC VM

Note: When recovering the SC, the *SAF, LOTC Disk Replication Consistency Failed*, Reference [7] alarm might appear. At the same time, if the infrastructure maintenance activities are taking more than 20 minutes to complete, then the *SAF, LOTC Disk Replication Communication Failed*, Reference [8] alarm might also appear. These alarms are expected during the VM recovery procedure on the SC, and should be automatically cleared when all recovery steps are executed. Refer to the corresponding alarm OPI for more information on these alarms.

If the VM to recover is an SC, then perform the following steps once the infrastructure maintenance has been finished:

1. Login to the other SC with the following command:

```
ssh root@<CUDB_Node_OAM_VIP_Address>
```

Refer to *CUDB Users and Passwords*, Reference [4] for more information on the default `root` password.

2. During the first boot, the new SC also synchronizes its replicated disk storage system with another SC. This process can take up to one hour, depending on the disk storage system size and network bandwidth. Use the following command to check the synchronization status:

```
cat /proc/drbd
```

3. On the recovered SC, restore the `crontab` jobs and their definitions, or similar tasks that are not deployed by default in CUDB. Edit cron configuration according to the cron configuration on the other SC.

3.4.2 Prepare PLDB or DSG VM

Depending on the type of VM, perform the applicable steps below to prepare the recovered PLDB or DSG VM for operation:

- ☐ In case the recovered VM belongs to a DSG, and replication is not recovered, then refer to the “Performing Combined Unit Data Backup and Restore” section of *CUDB Backup and Restore Procedures*, Reference [2].





Glossary

For the terms, definitions, acronyms and abbreviations used in this document, refer to *CUDB Glossary of Terms and Acronyms*, Reference [9].





Reference List

CUDB Documents

- [1] *CUDB Troubleshooting Guide*
- [2] *CUDB Backup and Restore Procedures*
- [3] *CUDB Node Configuration Data Model Description*
- [4] *CUDB Users and Passwords*, 3/00651-HDA 104 03/10
- [5] *CUDB Node Commands and Parameters*
- [6] *CUDB System Administrator Guide*
- [7] *SAF, LOTC Disk Replication Consistency Failed*
- [8] *SAF, LOTC Disk Replication Communication Failed*
- [9] *CUDB Glossary of Terms and Acronyms*