

Backup and Restore Guideline for Virtual and Cloud Deployment

Ericsson Dynamic Activation 1

USER GUIDE

Copyright

© Ericsson AB 2017. 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	Purpose and Scope	1
1.2	Target Groups	1
1.3	Typographic Conventions	1
1.4	Prerequisites	1
2	Backup and Restore	2
2.1	Backing Up and Restoring VM (KVM/VMware)	3
2.1.1	VM Backup	3
2.1.2	Restore VM Backup	3
2.2	Creating and Restoring VM Instance Snapshots (CEE)	4
2.2.1	Create Mandatory VM Instance Snapshot	4
2.2.2	Create Additional VM Instance Snapshot	5
2.2.3	Restore Mandatory VM Instance Snapshot	5
2.2.4	Restore Additional VM Instance Snapshot	6
2.3	Creating and Rebuilding VM Instance Snapshots (OpenStack)	6
2.3.1	Create VM Instance Snapshot	6
2.3.2	Rebuild VM Instance Snapshot	7
2.4	Backing Up and Restoring Dynamic Activation Software	8
2.4.1	Backing Up Dynamic Activation Software	8
2.4.2	Restoring Dynamic Activation Software	8
2.5	Backing Up and Restoring Zookeeper	9
2.5.1	Backing Up Zookeeper	9
2.5.2	Restoring Zookeeper	10
2.6	Backing Up and Restoring Cassandra (OAuth or Resource Configuration Related Data)	11
2.6.1	Backing Up Cassandra	11
2.6.2	Restoring Cassandra	12
	Reference List	13





1 Introduction

This document provides high-level guidelines on how to back up and restore Ericsson Dynamic Activation (EDA) on virtual and cloud deployments.

1.1 Purpose and Scope

The backup and restoration functionality in Dynamic Activation is not a complete solution with a separate backup media. Backup files created by the system must therefore be transferred to the secondary storage.

1.2 Target Groups

The target groups for this document are as follows:

- System Administrator
- Network Administrator
- Network Supervision Administrator

1.3 Typographic Conventions

Typographic conventions are described in the document *Library Overview*, Reference [1].

1.4 Prerequisites

To use this document fully, users must meet the following prerequisites:

- Knowledge about the Dynamic Activation product
- Knowledge about Linux operating system
- Knowledge about Kernel Based Virtualized Machine (KVM)
- Knowledge about VMware.
- Knowledge about Cloud Execution Environment (CEE)
- Knowledge about OpenStack.



2 Backup and Restore

The following backup and restore procedures are available:

- **VM Backup and Restore (KVM/VMware)**
 - Backing up VM – includes creating a backup of all virtual disks for each VM in the Dynamic Activation deployment, see Section 2.1.1 on page 3.
 - Restoring VM – includes restoring a previously backup of all virtual disks for each VM in the Dynamic Activation deployment, see Section 2.1.2 on page 3.
- **Creating and Restoring VM Instance Snapshots (CEE/OpenStack)**
 - Creating a VM instance snapshot – includes creating a snapshot of a VM instance in the cluster.
 - Restoring VM instance snapshot – includes restoring a previous snapshot of a VM instance in the cluster.

For detailed information on CEE deployment, see Section 2.2 on page 4.

For detailed information, on OpenStack deployment, see Section 2.3 on page 6.

- **Subset Backup and Restore**

A subset backup is used for backing up application data, and Dynamic Activation Software:

- Dynamic Activation Software, see Section 2.4.1 on page 8.
- Zookeeper, see Section 2.5.1 on page 9.
- Cassandra for OAuth or Recourse Configuration, see Section 2.6.1 on page 11.

A subset restore is used for restoring application data, and Dynamic Activation Software:

- Dynamic Activation Software, see Section 2.4.2 on page 8.
- Zookeeper, see Section 2.5.2 on page 10.
- Cassandra for OAuth or Recourse Configuration, see Section 2.6.2 on page 11.

- **Processing Log Data**

The Processing Log Data is handled separately.



For information on how to export Processing Log Data, see *System Administrators Guide for Virtual and Cloud Deployment*, Reference [2].

Note: All backup and export files need to be stored on a safe remote server.

2.1 Backing Up and Restoring VM (KVM/VMware)

The section describes the backup and restore procedures on all persistent storage, for each VM, in the Dynamic Activation deployment.

2.1.1 VM Backup

Caution!

To guarantee a clean backup, the VM needs to be shut down before taking a VM backup.

1. Shut down the VM that uses the persistent storage to be backed up, this to avoid data corruption.

Note: Shutdown of VM before a backup is mandatory, this to avoid data corruption.

2. Copy the virtual disk to an external backup storage.
3. When the copy process is done, startup the VM.

For detailed information regarding the procedure on VMWare, refer to Reference [3].

For detailed information regarding the procedure on KVM, refer to each host OS system guide.

2.1.2 Restore VM Backup

Caution!

To revert to a backup, the VM **MUST** be shut down. This is because of the risk of data corruption.

1. Shutdown VM that is to be restored.



2. Copy the virtual disk from the external backup storage to the storage used by the hypervisor.
3. Deploy the VM with the disks that were copied in previous step.
4. Start the restored VM.

2.2 Creating and Restoring VM Instance Snapshots (CEE)

In Cloud Execution Environment (CEE) deployment, VM instance backup is achieved by creating a snapshot of the instance. The snapshot can be used to restore the VM instance.

Note: Creating and restoring snapshots on the mandatory VM instances are different from that on additional VMs.

- Minimum 3 mandatory VMs are node-1, node-2, and node-3.
- Additional VMs (if any) are node-4, node-5, and forward.

2.2.1 Create Mandatory VM Instance Snapshot

This section describes how to create VM instance snapshot on node-1, node-2, or node-3.

1. Remove the VM instance from load balancer to be able to do a controlled shutdown of the instance.

For instructions, refer to section **Remove VM from Load Balancer** in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [2].

2. Log on to the VM instance to take snapshot of, and run the following command:

```
# /opt/dve/bin/CEESnapshotManagement.py --preSnapshot
```

3. Take a cold snapshot of the VM instance by following the instruction in **CEE 16A R5A** library in <http://cpistore.internal.ericsson.com/>.
4. When the snapshot is taken and the VM instance is running up again, enable all service by running the command from the VM instance:

```
# /opt/dve/bin/CEESnapshotManagement.py --postSnapshot
```

5. Add the VM instance to load balancer.

For instructions, refer to section **Add VM to Load Balancer** in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [2].

6. Check in the HAProxy GUI, to ensure that the VM instance is receiving traffic.



7. Log on to Atlas, choose **Computer > Image**, make sure that the created snapshot is listed in the image list.

2.2.2 Create Additional VM Instance Snapshot

This section describes how to create VM Instance snapshot on node-4 or forward.

1. Remove the VM instance from load balancer to be able to do a controlled shutdown of the instance.

For instructions, refer to section **Remove VM from Load Balancer** in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [2].

2. Take a cold snapshot of the VM instance by following the instruction in **CEE 16A R5A** library in <http://cpistore.internal.ericsson.com/>.

Note: In CEE, the VM instances have similar names as:
 node-4: `<hostname>-resource-0`
 node-5: `<hostname>-resource-1`
 and so forth.

3. Add the VM instance to load balancer.

For instructions, refer to section **Add VM to Load Balancer** in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [2].

4. Check in the HAProxy GUI, to ensure that the VM instance is receiving traffic.
5. Log on to Atlas, choose **Computer > Image**, make sure that the created snapshot is listed there.

2.2.3 Restore Mandatory VM Instance Snapshot

This section describes how to restore VM instance snapshot on node-1, node-2, or node-3.

1. In Atlas, choose **Compute > Instances**.
2. Locate the faulty VM instance and choose **Rebuild Instance** in the drop-down list.
3. In the pop-up dialog, select a correct snapshot image to restore from.
4. Enter the password of the logged in Atlas user, and then click **Rebuild**.
5. After the rebuilt image is running up, log on to the restored VM instance and run the following command to restore the Cassandra database:

```
# /opt/dve/bin/CEESnapshotManagement.py --restoreNode
```



Note: This command can take a long execution time, depending on how much data that were stored in the database.

6. To check the restoring progress, run:

```
# watch -n 1 nodedtool netstats
```

Look for lines similar as:

```
Receiving 91 files, 38132864075 bytes total.  
Already received 39 files, 27013761893 bytes total
```

Where “Receiving” refers to the total bytes to be transferred, and “Already” refers to the bytes has been transferred.

2.2.4 Restore Additional VM Instance Snapshot

This section describes how to restore VM Instance snapshot on node-4 or forward.

1. In Atlas, choose **Compute > Instances**.
2. Locate the faulty VM instance and choose **Rebuild Instance** in the drop-down list.

Note: In CEE, the VM instances have similar names as:

```
node-4: <hostname>-resource-0  
node-5: <hostname>-resource-1  
and so forth.
```

3. In the pop-up dialog, select a correct snapshot image to restore from.
4. Enter the password of the logged in Atlas user, and then click **Rebuild**.

2.3 Creating and Rebuilding VM Instance Snapshots (OpenStack)

In OpenStack deployment, VM instance backup is achieved by creating a snapshot of the instance. The snapshot can be used to rebuild the VM instance.

2.3.1 Create VM Instance Snapshot

1. Remove the VM instance from load balancer.

For instructions, refer to section **Remove VM from Load Balancer** in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [2].

2. Log on to the VM that the snapshot is to be taken on.
3. Shut down the VM:



```
# shutdown -h now
```

4. Log on to the OpenStack GUI as a user with correct rights and choose **Compute > Instances**.
5. Wait for the VM to shut down (it takes several minutes).
6. Locate the instance to back up and in the drop-down list (to the right) choose **Create Snapshot**.
7. Enter a name in the **Snapshot Name** box.
8. Click **Create Snapshot** and wait until the image has been created.
9. Start up the VM on the right panel.
10. Add the VM instance to load balancer.

For instructions, refer to section **Add VM to Load Balancer** in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [2].

11. Check in the HAProxy GUI, to ensure that the VM instance is receiving traffic.

2.3.2 Rebuild VM Instance Snapshot

1. Log on to the VM that is to be rebuilt and run the following command:

```
# shutdown -h now
```

2. Log on to the OpenStack GUI as a user with correct rights and choose **Compute > Instances**.
3. Make sure that the VM is shutoff before proceeding with the next step (it may take several minutes).
4. Locate the instance to rebuild and in the drop-down list (to the right) choose **Rebuild Instance**.
5. In the **Select Image** drop-down list, choose the image to rebuild from. **Disk Partition** must be set to *Automatic*.

Note: It is important to choose the correct image in the **Select Image** drop-down list.

6. Click **Rebuild Instance** and wait until the rebuild is done.
7. Start the VM after the rebuild is finished.



2.4 Backing Up and Restoring Dynamic Activation Software

This section describes the backup and restore procedures on the Dynamic Activation software.

Everything in the following directories is included:

- /home/bootloader
- /home/dveinstaller
- /home/actadm
- /home/dvecli
- /home/casadm
- /opt/ericsson/activation/bootloader

2.4.1 Backing Up Dynamic Activation Software

This section describes how to perform a backup of the Dynamic Activation software.

The backup script is installed by use of the `/opt/dve/bin/pgSoftwareBackup.sh` script.

To perform a full backup of the Dynamic Activation software:

1. Log in as user `root` on node-1.
2. Run the following command to back up the Dynamic Activation software:

```
# /opt/dve/bin/pgSoftwareBackup.sh backup </path/>  
<identifier_for_filename>
```

For example:

```
# /opt/dve/bin/pgSoftwareBackup.sh backup /home/ config
```

3. Backup command is finished when the following printout is displayed:

```
Backup stored as /home/dveinstaller/config-BootloaderBackup-20150212_124650.tar.gz
```

2.4.2 Restoring Dynamic Activation Software

This section describes how to restore the Dynamic Activation software from a backup.

1. Log in as user `root` on node-1.



2. Make sure that the backup `tar.gz` file is stored on the node to be restored.
3. Run the following command to restore the Dynamic Activation software:

```
# /opt/dve/bin/pgSoftwareBackup.sh restore <full_path_and_backup_tar.gz_file_name>.
```

When prompted, enter **yes** and press **Enter**

4. Log on to `node-1`, if currently logged in to another node.
5. Run the following command to activate the node that is to be restored:

```
# bootloader.py node activate --host <node-n>
```

Note: Do not use *all* parameter, which cannot be used in any cases other than new installations.

<node-n> is the name of the node that is to be activated, for example `node-1`.

6. Repeat Step 1 through Step 5 on other nodes that are to be restored.

Attention!

Wait for each node to be activated before starting with the next one, otherwise traffic disturbances occur.

2.5 Backing Up and Restoring Zookeeper

This chapter describes how to back up and restore the Zookeeper in Dynamic Activation. This involves, for example, Network Element and user configurations.

Note: Backing up and restoring Zookeeper-data is a software version-dependent procedure. The Zookeeper-data backed up from a specific system must only be restored on the same specific software release. This to ensure full functionality of the restored data.

2.5.1 Backing Up Zookeeper

To perform a backup of Zookeeper:

1. Log in as an administrator on either `node-1` or `node-2`.
2. Start the backup process by executing the following script:

```
$ sudo zookeeperBackup.sh backup </path> <identifier_for_filename>
```



For example:

```
$ sudo zookeeperBackup.sh backup /home/ config
```

3. When the backup process is started, answer **yes** on the following question:

```
Backup will use approximately 1065161108 kb of disk  
space. Do you want to continue? (yes/no):
```

4. When the backup process is done, the following printout is displayed:

```
Backup stored as /home/config-zookeeperBackup-2015021  
2_154150.tar.gz
```

2.5.2

Restoring Zookeeper

Note: This involves stopping and starting the Dynamic Activation application on all nodes, which results in traffic downtime.

When the configuration data has been restored, alarms that are no longer valid can still persist in the system. Use the `fmsendmessage` command to remove those alarms manually. For more information, refer to *System Administrators Guide for Virtual and Cloud Deployment*, Reference [2].

1. Log in as an administrator on node-1 and run the following command:

```
$ sudo zookeeperBackup.sh restore <full_path_and_backup_t  
ar.gz_file_name>
```

Caution!

This causes traffic downtime on the Dynamic Activation Application.

For example:

```
$ sudo zookeeperBackup.sh restore /home/config-zookeepe  
rBackup-20150209_154150.tar.gz
```

When prompted enter **yes** and press **Enter**

The restore command is finished when the following printout is displayed:

```
Restore finished: /home/config-zookeeperBackup-2015020  
9_154250.tar.gz
```

2. Go to the Dynamic Activation GUI and check the user and network element configurations.



2.6 Backing Up and Restoring Cassandra (OAuth or Resource Configuration Related Data)

2.6.1 Backing Up Cassandra

To perform a backup of Cassandra:

1. Log in as an administrator on node-1 in the cluster.
2. Start the backup process by executing the backup script depending on what to back up:

- **OAuth**

```
$ sudo -u actadm cassandra_backup.py backup oauth
```

- **Resource Configuration**

```
$ sudo -u actadm cassandra_backup.py backup scm
```

- **OAuth and Resource Configuration**

```
$ sudo -u actadm cassandra_backup.py backup all
```

3. Enter the full path where the backup is to be stored:

Enter full path to backup folder: `/<path>`

For example:

Enter full path to backup folder: `/home/actadm/backup/`

4. When the backup process is done, for each node on which Cassandra is installed, one backup file is stored in the `backup` folder.

For example, in a four-node cluster the following files are to be stored in the `backup` folder of node-1, node-2, and node-3:

```
$ ls -l /home/actadm/backup
total 1032
-rw-r----- 1 actadm activation 349702 Mar  2 07:28 node-1_201703
-rw-r----- 1 actadm activation 349270 Mar  2 07:28 node-2_201703
-rw-r----- 1 actadm activation 350804 Mar  2 07:28 node-3_201703
```

Note: For Virtual or Cloud deployment, Cassandra is installed on the first three nodes even if there are more than three nodes in the cluster.



2.6.2 Restoring Cassandra

Note: Restoring Cassandra requires stopping and starting the Dynamic Activation application on all nodes, which results in traffic downtime.

Restoring also means that all data in the Cassandra database tables are to be overwritten by backup data. The affected data are OAuth and Resource Configuration related data.

1. Log in as an administrator on node-1 and disable the Dynamic Activation application processes by executing the following command:

Caution!

The following command causes traffic downtime on the Dynamic Activation application.

```
$ bootloader.py node stop --host all
```

2. Run the following command to start the Cassandra restore process:

Caution!

The following command causes all old OAuth or Resource Configuration related data to be replaced in the Cassandra database.

```
$ sudo -u actadm cassandra_backup.py restore
```

3. Enter the full path where the backup is stored:

Enter full path to backup folder: `/<path>`

For example:

Enter full path to backup folder: `/home/actadm/backup/`

4. Log on as an administrator on node-1 and enable the Dynamic Activation application processes by executing the following command:

```
$ bootloader.py node start --host all
```




Reference List

Ericsson Documents

- [1] *Library Overview*, 18/1553-CSH 109 628 Uen
- [2] *System Administrators Guide for Virtual and Cloud Deployment*, 3/1543-CSH 109 628 Uen

Other References

- [3] VMware vSphere Data Protection Documentation, https://www.vmware.com/support/pubs/vdr_pubs.html