

System Expansion for Virtual and Cloud Deployment

Ericsson Dynamic Activation 1

OPERATING INSTRUCTIONS

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 Group | 1 |
| 1.3 | Typographic Conventions | 1 |
| 2 | Expansion | 1 |
| 2.1 | Prerequisites for Expanding a Dynamic Activation Cluster | 2 |
| 2.2 | Guidelines for the Readers | 2 |
| 2.3 | Expansion of a Dynamic Activation Cluster | 3 |
| 2.3.1 | Expansion in Virtualized Deployment | 4 |
| 2.3.2 | Expansion in Cloud Deployment - ECEE | 10 |
| 2.3.3 | Expansion in Cloud Deployment - OpenStack | 14 |
| 2.4 | Installing HSS Validator Plug-in (Optional) | 17 |
| 2.5 | SSL Configuration (Optional) | 18 |
| 2.6 | Modify Notification Rules | 18 |
| 2.7 | Creating Administrative Users | 18 |
| 3 | Backup | 18 |
| | Reference List | 19 |





1 Introduction

This document describes the Kernel Based Virtual Machine (KVM), VMware, Ericsson Cloud Execution Environment (ECEE) and OpenStack expansion process for Ericsson Dynamic Activation (EDA).

1.1 Purpose and Scope

The purpose of this document is to give detailed description about how to expand a cluster with another virtual Dynamic Activation instance and how to verify the expansion.

1.2 Target Group

The target groups for this document are as follows:

- Ericsson installation engineers
- Other Dynamic Activation related engineers

The target groups are described in more detail in the *Library Overview*, Reference [1].

1.3 Typographic Conventions

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

For information about abbreviations used throughout this document refer to *Glossary of Terms and Acronyms*, Reference [2].

2 Expansion

This section contains instructions on how to expand a Dynamic Activation cluster.



2.1 Prerequisites for Expanding a Dynamic Activation Cluster

Attention!

Make sure to have a valid Dynamic Activation 1 license reflecting the number of VMs that the cluster will be expanded to.

- Make sure to have the Open Virtual Appliance (OVA) for Dynamic Activation accessible.
- If a Validator Plug-in is to be used (optional), make sure to have the HSS Validator Plug-in software accessible.

Contact [Ericsson support](#) for details about the applicable HSS Validator Plug-in version.

- Before scaling out the Dynamic Activation system, please consider if there is a need to scale out any other parts of the solution first.

2.2 Guidelines for the Readers

- All VMs in a Dynamic Activation installation have a type name. The type name is always node-n, for example. node-1, node-2, node-3.
- The type name gives an indication of what services that are configured, and that are executing on the VM.

Node-1, node-2, and node-3 possess individual, unique configuration data, and must be on different physical hosts, to achieve high-availability characteristics.

- The load-balancers are not running on all VMs. They are only on node-1 and node-2.
- Node-4 and onwards, are all identical when it comes to which services that are configured and activated.



2.3 Expansion of a Dynamic Activation Cluster

Attention!

For commercial deployment by using KVM and VMware ESXi, it is mandatory to set up persistent block storage for node 1-3, see section **Add Block Storage Device to VM** in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [6].

In ECEE deployment, ephemeral disk is used.

Note: The following needs to be performed for every added Virtual Machine (VM).

If using KVM, make sure:

- That the host system meets the prerequisites described in *Requirements on Virtualization and Cloud Infrastructure*, Reference [4].
- That the network bridge interfaces are configured as described in *Network Description and Configuration for Virtual and Cloud Deployment*, Reference [5].
- That the Dynamic Activation OVA is available on the host file system.
- That the KVM host is time synchronized through NTP.
- That the `libvirt` management tools are available on the KVM host.

For more information, see *Requirements on Virtualization and Cloud Infrastructure*, Reference [4].

- That the `genisoimage` rpm is installed on each host.

For more information, see *Requirements on Virtualization and Cloud Infrastructure*, Reference [4].

If using VMware, make sure:

- That the host system meets the prerequisites described in *Requirements on Virtualization and Cloud Infrastructure*, Reference [4].
- That the virtual network interfaces are configured as described in *Network Description and Configuration for Virtual and Cloud Deployment*, Reference [5].
- That the VMware ESXi host is time synchronized through NTP

If using ECEE, make sure:



- There is a working ECEE being used for deployment
- That all required variables needed from ECEE is gathered.
- User and password to ECEE (Atlas GUI) are available.

If using OpenStack, make sure:

- There is a working OpenStack being used for deployment
- That all required variables needed from OpenStack is gathered.
- User and password to OpenStack GUI are available.

2.3.1 Expansion in Virtualized Deployment

This section describes how to expand a Dynamic Activation cluster for virtualized deployment.

2.3.1.1 Preparing Deployment Artifacts

When the prerequisites specified in Section 2.3 on page 2 are met, continue with continue with the following step-list:

1. Download and start the EDA Deployment Manager:

- 1 Save the zip file, [EDA_Deployment_Manager.zip](#) to, a local area on a local machine
- 2 Unpack the zip file.
- 3 Double-click the .jar file.

Note: Requires Oracle's JAVA version 1.8.0_40 or higher (about 200 KB in size)

If not having Oracle's JAVA version 1.8.0_40 or later, there is a possibility to create an Oracle JAVA independent .bat file and run the tool. To create such file:

Note: This alternative works only on windows.

- 1 Download (do not install) the appropriate Oracle JRE for the machine that the EDA Deployment Manager tool will be run from, and store it in the same folder as where the ActivationDeploymentArtifactManager.jar file was unpacked.
- 2 In the same directory, create the EDA_DeploymentManager.bat file, with the following content:

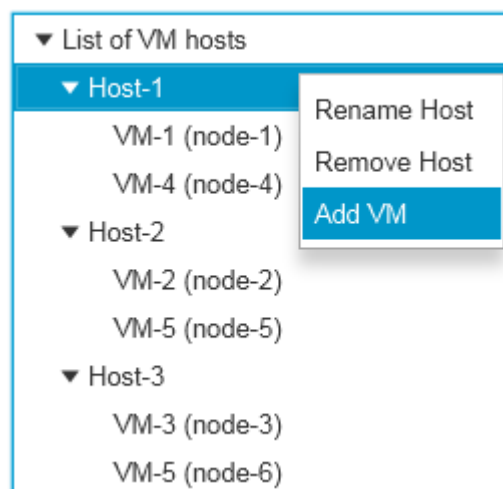
```
<Path to java.exe file in downloaded JRE> -jar
```



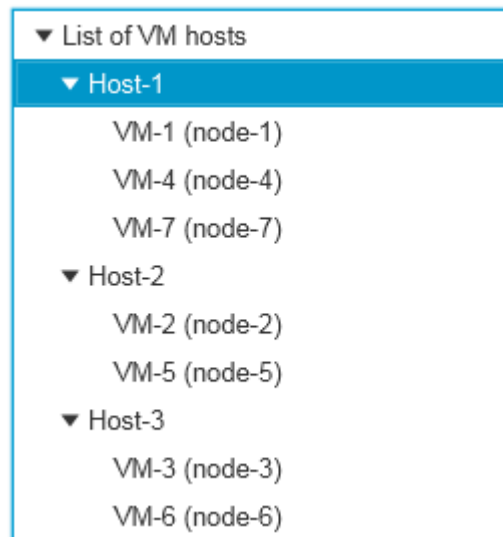

2. Open the deployment schema file, `.ds`, that was prepared for the initial installation. If the file does not exist, re-enter all values in the EDA Deployment Manager tool. For detailed information, see **Using EDA Deployment Manager** in *Software Installation for Virtual and Cloud Deployment*, Reference [3].
3. Add a VM to the desired host.

To add a VM, right-click on a host, for example **Host-1** and select **Add VMs**:

Add VM - Before configuration:



Add VM - After configuration:



4. Fill in the VM-specific data:



VM-specific data

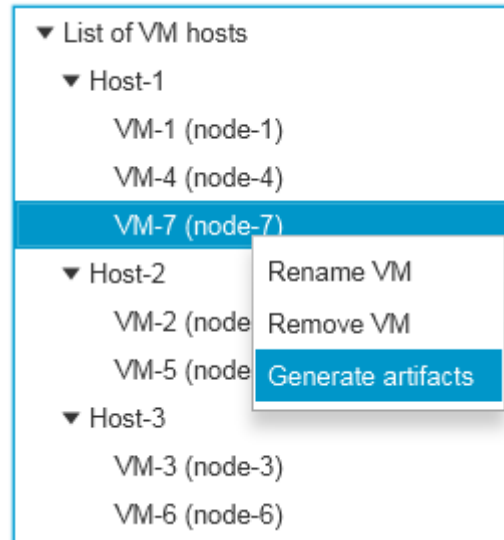
| | | |
|-------------------------------------|-----------------------------------|---|
| VM ID | <input type="text" value="VM-7"/> | ? |
| VM Host Name | <input type="text"/> | ? |
| # of vCPUs (0 = default) | <input type="text" value="0"/> | ? |
| Amount of RAM (GB) (0 = default) | <input type="text" value="0"/> | ? |
| O&M IPv4 (combined = incl. traffic) | <input type="text"/> | ? |
| -----> IPv6 (optional) | <input type="text"/> | ? |
| Traffic IPv4 | <input type="text"/> | ? |
| -----> IPv6 (optional) | <input type="text"/> | ? |
| Private IPv4 | <input type="text"/> | ? |
| -----> IPv6 (optional) | <input type="text"/> | ? |

5. Generate VM artifacts:

Right-click on the new VM and select **Generate artifacts**



Generate VM artifacts



6. Store the generated artifacts (`bootstrap.iso`) on an empty directory on a local machine.
7. From the EDA Deployment Manager tool, save the updated deployment schema file, `.ds`, on an appropriate storage area. This file can, for example be used for further cluster expansion.
8. If using KVM continue with Section 2.3.1.2 on page 7.
If using VMware continue with Section 2.3.1.3 on page 8.

2.3.1.2 Deploying Artifacts - KVM

The following instruction is used to deploy Dynamic Activation:

1. Create one folder on the KVM host that will hold deployment artifacts, as well as the image for the VM. Skip this step if the KVM host already contains at least one VM.

Note: Make sure that there is enough disk space. The space required is at least 130 GB times (*) the number of VMs that will be running on that specific host.
2. Copy or move the node-n folder to the KVM host.

Note: The folder itself needs to be copied, not just its contents.
3. If not already present, on the KVM host, extract the `<Software_Package>.tar.gz` file (EDA System KVM&Cloud SW) in the same folder as where the node-n folder was previously stored, see section **Preparing Deployment Artifacts** Step 6:



```
# tar xvf <Software_Package>.tar.gz
```

4. On the KVM host where the node-n folder resides, run the `deploy.sh` script, contained in the extracted (EDA System KVM&Cloud SW) file, to define and start the newly added VM:

```
# KVM/deploy.sh
```

Note: It is not supported to expand the cluster with more than one VM node at a time.

Depending on the host performance and number of VMs in the cluster, the deployment process can vary in time to complete. Verify that all processes are operational with the command in Step 2.

5. Continue with Section 2.3.1.4 on page 8.

2.3.1.3 Deploying Artifacts - VMware

1. Use the template created during the installation and create the new VM.
2. In VMware vSphere, upload the `bootstrap.iso` files that were created with EDA Deployment Manager section **Preparing Deployment Artifacts** Step 6, to a datastore that is connected to the hypervisor.
3. Connect the `bootstrap.iso` to the virtual CD/DVD on the added VM.
4. Power on the added VM, wait a few minutes to let Puppet configure the new VM.
5. Continue with Section 2.3.1.4 on page 8.

2.3.1.4 Activating Dynamic Activation on the Expanded Node

Note: Make sure all VMs are deployed on each respective physical host before proceeding with the following instruction.

1. Log in as root (**username:** root, **password:** rootroot) on the newly added node, and check that all configurations are completed:

```
# systemctl status puppet.service
```

Note: The host name of the node is not necessarily `node-1`. The name is depending on the values that were previously filled in, by use of the EDA Deployment Manager tool.

Example of a Successful Output:

```
Finished catalog run in 5.99 seconds
```

2. Check that all processes are up on all VMs:

```
# 3ppmon status --host all
```



Every entry must have status UP or, if the process is not supposed to run on the node, have a dash (-).

3. Run the following commands to activate the configuration and start the traffic test:
 - a. Change directory:


```
# cd /var/log/installfiles/<Prod_Number>-<Version>
```
 - b. Deploy the EDA package. Run the following command from node-1:


```
# ./ema deploy -p EMA
```
4. By default, test mode is enabled on the expanded node. Use the test port 8888 or 8989 to verify if it is possible to send and receive traffic on the expanded node.
5. Go back to normal mode.

From node-1, run the following commands:

```
# bootloader.py config remove --parameter @REGISTER_S
SERVICES@

# bootloader.py config remove --parameter @REGISTER_TES
T_SERVICES@
```

Note: All manual changes, for example in application dependent configuration files, added or upgraded RPMs, must be performed on the added node.

6. Enable traffic on the new node.

From node-1, run the following command:

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

<hostname> is the hostname of the node that is to be activated.
```

7. From node-1, run the following command to check that no errors exist and that all bindings are OK:


```
# bootloader.py node status --host all
```
8. Change password for the root account on the expanded VM:


```
# ssh <hostname of new node> passwd
```
9. To add more nodes in the cluster, repeat the procedure starting from Section 2.3.1.1 on page 4.

2.3.2 Expansion in Cloud Deployment - ECEE

This section describes how to expand a Dynamic Activation system that is deployed in ECEE.

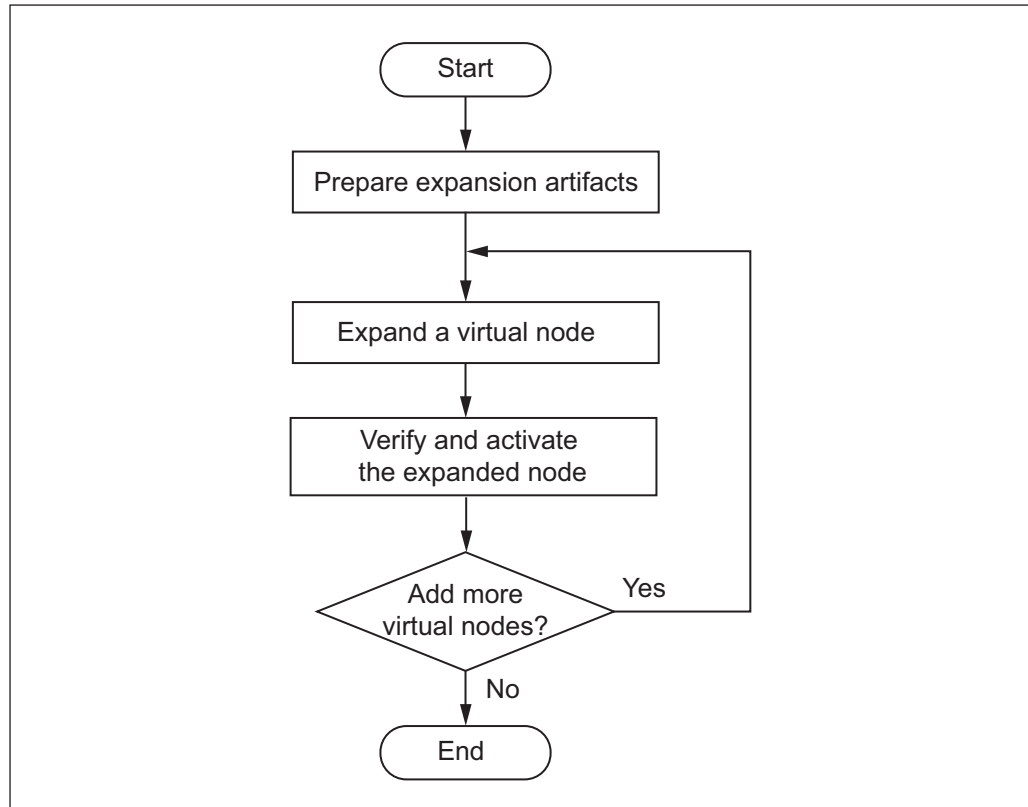


Figure 1 Expansion Workflow in ECEE

2.3.2.1 Prepare Expansion Artifacts

Before expanding the existing Dynamic Activation in ECEE, do as follows:

1. Prepare the `vEDA_expansion.yaml` file to a location that can be reached from ECEE.

This file can be found in the Dynamic Activation cloud deployment package `EDA_KVM_CLOUD_SW-<version>.tar.gz`. For more information, refer to *Software Installation for Virtual and Cloud Deployment*, Section **Unpack Deployment Package - CEE**, Reference [3].

2. Check the status of the existing Dynamic Activation.
 - a. From any node, check 3PP processes on all VMs.

As user `root`:

```
# 3ppmon status --host all
```



All 3PP processes must have status of either `UP` or `-` (a dash), depending whether the process is supposed to run on the node.

Active alarms must be `0`.

- b. From node-1, check Dynamic Activation application processes.

As user `root`:

```
# bootloader.py node status --host all
```

There must be no errors, and all bindings must be correct.

2.3.2.2 Expand Virtual Dynamic Activation Node

Use ECEE GUI Atlas to add one virtual node (VM instance) to the existing Dynamic Activation.

1. Log on to Atlas as a user with correct rights.
2. Choose **Orchestration > Catalog > Upload**.
3. Enter an application name.
4. In **Type**, select **HOT**.
5. In **Application Source** drop-down list, select **From File**, and then use **Choose File** to select the `vEDA_expansion.yaml` file that was prepared in Section 2.3.2.1 on page 10.
6. Click **Upload**, and wait until the file is uploaded.
7. In **Orchestration > Catalog**, check the application:
 - **Status** must be `Active`
8. Click **Launch**, and then **Next**.
9. Set the following configurations:
 - **A Stack Name**
 - The password of the logged in user
 - Name of instance/hostname
10. In **Name of the vEDA image** drop-down list, select the image of the Dynamic Activation that is to be expanded.
11. In **vEDA flavor** drop-down list, select a flavor name.

Note: The flavor name starts with the stack name that was created for the existing Dynamic Activation infrastructure.



12. Set the **Number of vEDA Existing Instances**, which is used for forming a correct hostname for the expanded VM instance.

For example, if entering 4, the expanded VM instance will be named as node-5 in the Dynamic Activation cluster.

Attention!

Count all VM instances that is currently part of the Dynamic Activation cluster to be expanded.

Wrong number causes expansion failure.

13. Enter the following IP addresses of the existing nodes:

- Internal IP addresses of node-1, node-2, and node-3.
- Internal and external VIP addresses

The IP information can be found as follows:

- a. Log on as user `root` to the master VM (with type name node-1).
- b. Run the following command.

```
# cat /etc/hosts
```

14. Set the following configurations:

- Time zone of the VM
- NTP server IP
- DNS server IP

15. Click **Launch**, and wait until the added VM instance is launched.

Note: It takes approximate 7 minutes to finish deploying Dynamic Activation in an expanded node.

16. Choose **Orchestration > Stacks**, select the created stack name, and check the **Events** tab:

- **Status** must be `Create Complete`
- **Status Reason** must be `Stack CREATE completed successfully`



2.3.2.3

Verify and Activate the Expanded Node

To verify and activate the expanded node, do as follows:

1. Log on as user `root` to the master VM (with type name `node-1`).
2. SSH from `node-1` to the expanded node.
3. Check that the log file `/var/log/cloud-init.log` contains the following text:

```
*** Activation finished"
```

```
# ssh <hostname of expanded node> cat /var/log/cloud-init.log
```

4. By default, test mode is enabled on the expanded node. Use the test port `8888` or `8989` to verify if it is possible to send and receive traffic on the expanded node.
5. Set the expanded node to normal mode.

From `node-1`, run the following commands:

```
# bootloader.py config remove --parameter @REGISTER_SERVICES@
```

```
# bootloader.py config remove --parameter @REGISTER_TEST_SERVICES@
```

Note: All manual changes, for example in application dependent configuration files, added or upgraded RPMs, must be performed on the expanded node.

6. (Optional) If an HSS Validator Plug-in is needed, install it on the expanded node. See Section 2.4 on page 17.
7. Enable traffic on the new node.

From `node-1`, run the following commands:

```
# bootloader.py node activate --host <hostname of the new node>
```

8. From `node-1`, run the following command to check that no errors exist and all bindings are OK:

```
# bootloader.py node status --host all
```

9. From `node-1`, change password for the root account on the expanded node:

```
# ssh <hostname of new node> passwd
```

10. To add more nodes in the cluster, repeat the procedure starting from Section 2.3.2.2 on page 11.

2.3.3 Expansion in Cloud Deployment - OpenStack

This section describes how to expand a Dynamic Activation system that is deployed in OpenStack.

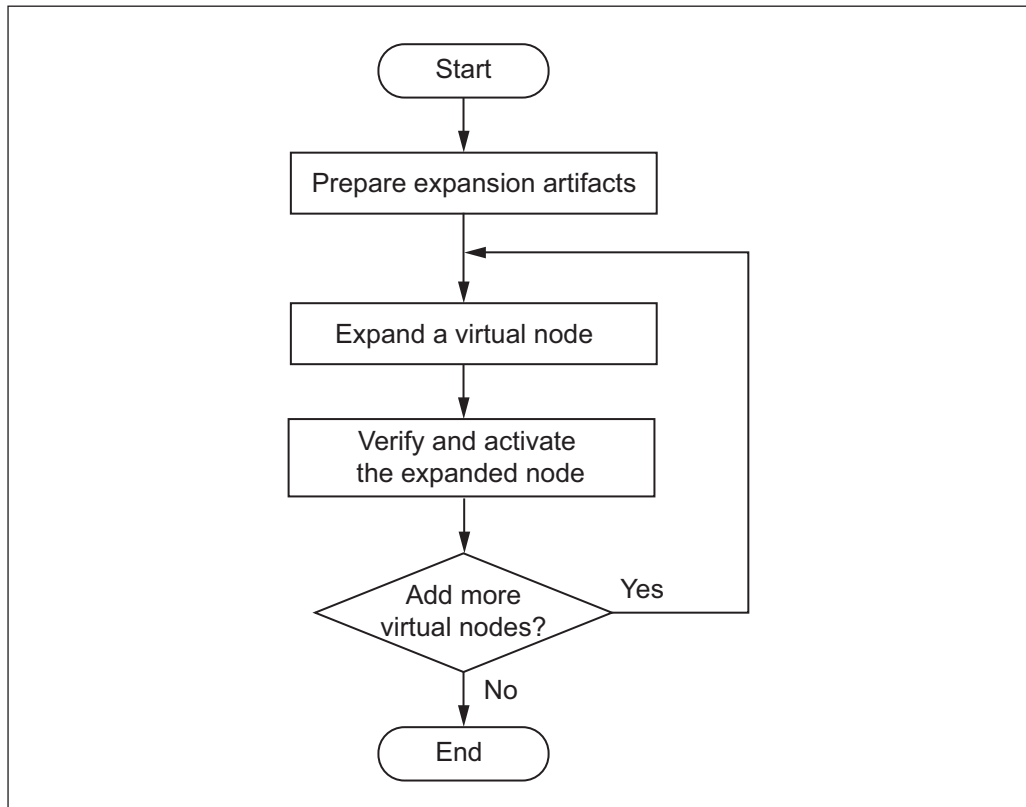


Figure 2 Expansion Workflow in OpenStack

2.3.3.1 Prepare Expansion Artifacts

Before expanding the existing Dynamic Activation in OpenStack, do as follows:

1. Prepare the `vEDA_expansion.yaml` file to a location that can be reached from OpenStack.

This file can be found in the Dynamic Activation cloud deployment package `EDA_KVM_CLOUD_SW-<version>.tar.gz`. For more information, refer to *Software Installation for Virtual and Cloud Deployment*, Section **Unpack Deployment Package** - OpenStack, Reference [3].

2. Check the status of the existing Dynamic Activation.
 - a. From any node, check 3PP processes on all VMs.



As user `root`:

```
# 3ppmon status --host all
```

All 3PP processes must have status of either `UP` or `-` (a dash), depending whether the process is supposed to run on the node.

Active alarms must be `0`.

- b. From node-1, check Dynamic Activation application processes.

As user `root`:

```
# bootloader.py node status --host all
```

There must be no errors, and all bindings must be correct.

2.3.3.2 Expand Virtual Dynamic Activation Node

Use OpenStack GUI to add one virtual node (VM instance) to the existing Dynamic Activation.

1. Log on to the OpenStack GUI as a user with correct rights.
2. Choose **Project > Orchestration > Stacks > Launch Stack**.
3. In **Template Source** drop-down list, select **File**, and then use **Template File Browse** to select the `vEDA_expansion.yaml` file that was prepared in Section 2.3.3.1 on page 14.
4. Click **Next**.
5. Set the following configurations:
 - **A Stack Name**
 - Creation Timeout (minutes). Default value is 60 minutes.
 - The password of the logged in user
 - Name of instance/hostname
6. In **Name of the vEDA image** enter the name of the Dynamic Activation image that is to be expanded.
7. In **vEDA flavor** drop-down list, select a flavor name.
8. Set the **Number of vEDA Existing Instances**, which is used for forming a correct hostname for the expanded VM instance.

For example, if entering `4`, the expanded VM instance will be named as `node-5` in the Dynamic Activation cluster.



Attention!

Count all VM instances that is currently part of the Dynamic Activation cluster to be expanded.

Wrong number causes expansion failure.

9. Enter the following IP addresses of the existing nodes:

- Internal IP addresses of node-1, node-2, and node-3.
- Internal and external VIP addresses

The IP information can be found as follows:

- a. Log on as user `root` to the master VM (with type name node-1).
- b. Run the following command.

```
# cat /etc/hosts
```

10. Set the following configurations:

- Time zone of the VM
- NTP server IP
- DNS server IP

11. Click **Launch**, and wait until the added VM instance is launched.

Note: It takes approximate 5-10 minutes to finish deploying Dynamic Activation in an expanded node.

12. Choose **Orchestration** > **Stacks**, select the created stack name, and check the **Events** tab:

- **Status** must be `Create Complete`
- **Status Reason** must be `Stack CREATE completed successfully`

2.3.3.3 Verify and Activate the Expanded Node

To verify and activate the expanded node, do as follows:

1. Log on as user `root` to the master VM (with type name node-1).
2. SSH from node-1 to the expanded node.



3. Check that the log file `/var/log/cloud-init.log` contains the following text:

```
*** Activation finished
```

```
# ssh <hostname of expanded node> cat /var/log/cloud-init.log
```

4. By default, test mode is enabled on the expanded node. Use the test port 8888 or 8989 to verify if it is possible to send and receive traffic on the expanded node.
5. Set the expanded node to normal mode.

From node-1, run the following commands:

```
# bootloader.py config remove --parameter @REGISTER_SERVICES@
```

```
# bootloader.py config remove --parameter @REGISTER_TEST_SERVICES@
```

Note: All manual changes, for example in application dependent configuration files, added or upgraded RPMs, must be performed on the expanded node.

6. (Optional) If an HSS Validator Plug-in is needed, install it on the expanded node. See Section 2.4 on page 17.
7. Enable traffic on the new node.

From node-1, run the following commands:

```
# bootloader.py node activate --host <hostname of the new node>
```

8. From node-1, run the following command to check that no errors exist and all bindings are OK:

```
# bootloader.py node status --host all
```

9. From node-1, change password for the root account on the expanded node:

```
# ssh <hostname of new node> passwd
```

10. To add more nodes in the cluster, repeat the procedure starting from Section 2.3.2.2 on page 11.

2.4 Installing HSS Validator Plug-in (Optional)

For information on how to install an HSS Validator Plug-in, see **Installing HSS Validator Plug-in (Optional)** in *Software Installation for Virtual and Cloud Deployment*, Reference [3].



2.5 SSL Configuration (Optional)

For information on how to configure SSL, follow the instructions in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [6].

2.6 Modify Notification Rules

For more information, see section **Notification Rules File Administration** in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [6].

2.7 Creating Administrative Users

Administrative users need to be created for the newly added nodes. For details, refer to section **Create Administrative User** in *System Administrators Guide for Virtual and Cloud Deployment*, Reference [6].

3 Backup

When the system is expanded and properly configured, make a full backup to be able to revert to the original state when needed. Create a full backup as described in *Backup and Restore Guideline for Virtual and Cloud Deployment*, Reference [7].



Reference List

Ericsson Documents

- [1] *Library Overview*, 18/1553-CSH 109 628 Uen
- [2] *Glossary of Terms and Acronyms*, 0033-CSH 109 628 Uen
- [3] *Software Installation for Virtual and Cloud Deployment*, 4/1531-CSH 109 628 Uen
- [4] *Requirements on Virtualization and Cloud Infrastructure*, 2/2135-CSH 109 628 Uen
- [5] *Network Description and Configuration for Virtual and Cloud Deployment*, 1/1551-CSH 109 628 Uen
- [6] *System Administrators Guide for Virtual and Cloud Deployment*, 3/1543-CSH 109 628 Uen
- [7] *Backup and Restore Guideline for Virtual and Cloud Deployment*, 6/1553-CSH 109 628 Uen