

# Network Description and Configuration for Virtual and Cloud Deployment

## Ericsson Dynamic Activation 1

---

### DESCRIPTION

**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

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Purpose and Scope	1
1.1.1	In Scope	1
1.1.2	Out of Scope	1
1.2	Target Groups	1
1.3	Typographic Conventions	2
1.4	Reader's Guideline	2
<b>2</b>	<b>Networks</b>	<b>3</b>
2.1	Dual Stack	3
2.2	External Physical Connectivity	3
2.3	Physical and Logical Network Setup	4
2.4	IP Allocation	6
<b>3</b>	<b>Load Balancing</b>	<b>9</b>
3.1	Keepalived and HAProxy	9
3.2	Keepalived/HAProxy Router	10
3.3	Real Server	10
<b>4</b>	<b>High Availability</b>	<b>11</b>
<b>5</b>	<b>Firewall Configuration</b>	<b>13</b>
<b>6</b>	<b>Node Configuration</b>	<b>15</b>
	<b>Reference List</b>	<b>17</b>





# 1 Introduction

This document gives detailed information about the network configuration of Ericsson Dynamic Activation (EDA) on Virtual and Cloud platforms.

It refers to the configuration of the Network Infrastructure Controller (NICs) and network setup of the hosting hypervisor.

## 1.1 Purpose and Scope

This section contains information about what is in scope and what has been left out.

### 1.1.1 In Scope

- Network Infrastructure Overview.
- Logical subnetworks: application-related traffic network and Operation Administration and Maintenance (OAM) network.
- Connectivity (logical and IP design).
- Support of primary - backup configuration of Keepalived/HAProxy load balancer.
- Configuration of the host servers.

### 1.1.2 Out of Scope

- Dynamic Activation application software configurations.
- Parts of the configuration of the customer network outside the system.

## 1.2 Target Groups

The target groups for this document are as follows:

- System Administrator
- Network Administrator
- Network Supervision Administrator

For more detailed information about the target groups presented in the list above, see *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, see *Glossary of Terms and Acronyms*, Reference [2].

## 1.4      **Reader's Guideline**

All examples throughout this document refer to Dynamic Activation Virtual Machines (VMs).



## 2 Networks

The network infrastructure for Dynamic Activation in a virtual and cloud environment is divided into three networks: Internal, Operating & Maintenance (O&M) and Provisioning traffic.

For cloud deployment, the O&M and Provisioning traffic must be combined on one interface. For virtual deployment, two different interfaces can be used.

To establish communication between the Virtual Machine (VM) and the IP backbone, either two or three virtual networks must be defined, (depending on if separate traffic network is used), and connected to the host.

### 2.1 Dual Stack

Dynamic Activation deployment supports IPv4/IPv6 Dual stack. It enables the possibility to use both IPv4, and IPv6, Northbound and Southbound simultaneous.

**Note:** Dynamic Activation does not support pure IPv6 deployment.

Dual stack is not supported for Ericsson Cloud Execution Environment (ECEE) or OpenStack deployments.

### 2.2 External Physical Connectivity

Provisioning traffic destined to and sourced by Dynamic Activation can be separated from OAM traffic in links dedicated to application traffic. This is achieved by either physical separated links or logical separation by use of VLANs.

The KVM/VMware connectivity is depicted in figure Figure 1.

The CEE/OpenStack connectivity is depicted in figure Figure 2.




Figure 1 Dynamic Activation Connectivity Overview KVM/VMware




Figure 2 Dynamic Activation Connectivity Overview CEE/OpenStack

## 2.3 Physical and Logical Network Setup

Provisioning traffic destined to, and sourced by VMs can be separated by the use of VLAN tagging.

The connectivity is depicted in the following figure:





**Note:** The figure below show an example of how to setup Dynamic Activation in a virtual or cloud environment.

CEE and OpenStack deployments use common uplinks for O&M and Traffic.




Figure 3 Physical and Logical Network Setup Example

**Note:** In CEE deployments, bridges are managed automatically by the CEE environment.

- Bridge `internal->bond0` handles internal traffic between the VMs, for example zookeeper data replication.
- Bridge `br-oam->bond1` handles O&M traffic, for example ESA alarms and GUI access. It can also be used for provisioning traffic, if combining O&M and traffic on the same interface.
- Bridge `br-trf->bond1` handles provisioning traffic and connection towards network elements.

The two bridges, in the example above, `br-oam` and `br-trf`, share physical Network Interface Cards (NICs). There are two ways to perform this network setup:

The two bridges, in the example above, `br-oam` and `br-trf`, share physical Network Interface Cards (NICs). The NIC bridge in the hypervisor is set up with VLAN tagging. This means that the VM sends untagged packages to the hypervisor, which in turn handles the VLAN tagging.



The Keepalived and HAProxy solutions are used for load balancing. For more informations, see Section 3 on page 9.

Only VMs where the Keepalived/HAProxy routers reside are receiving provisioning traffic. These are the only VMs that need to have connection towards the traffic network. The VMs acting as Real servers, require only the internal network, `internal > bond0` and O&M network `br-oam > bond1`.

## 2.4 IP Allocation

Dynamic Activation on virtualized and cloud deployment uses an IP addressing schema as described in Table 2. The internal network is defined by private IP addresses not accessible from external networks.

**Note:** It is important that the Keepalived/HAProxy routers reside on different hosts. This for redundancy purposes.

For virtualized deployment, if the NE addresses defined in the GUI are located in the same subnet as the VM:s O&M network, static routing needs to be defined from node-3 to node-x towards these elements. These routes should point towards VIP internal gateway.

*Table 1 Load Balancing Allocation Example*

Name	Node	Host	Description
Keepalived/HAProxy Primary	node-1	1	The first node acts as the primary router and as a Real server.
Keepalived/HAProxy Backup	node-2	2	The second node acts as the backup router and as a Real server.
Real Server	node-3	3	The third node acts as a Real server.

**Note:** For CEE the internal addresses are dynamically allocated using DHCP.

*Table 2 IP Allocation Example*

IP Address	VM	Interface	Type	Comments
<b>Internal Network</b>				
192.168.0.0/24	-	-	-	The Internal Network
.1	node-1	eth0	Internal	IP of the first VM
.2	node-2	eth0	Internal	IP of the second VM
.3	node-3	eth0	Internal	IP of the third VM
.254	node-1, node-2	eth0	Internal	The virtual IP of the Keepalived/HAProxy for the internal network which also works as the default gateway for the Real servers
<b>External Network</b>				



IP Address	VM	Interface	Type	Comments
VIP_TRF	node-1, node-2	eth2	External	The virtual IP of the Keepalived/HAProxy for the external provisioning network
ETH2_IP	node-1, node-2	eth2	External	IP of the node in the provisioning network
ETH2_TRF_GW	node-1, node-2	eth2	External	Default Gateway used for outgoing provisioning traffic in communication with network elements <sup>(1)</sup>
<b>External O&amp;M Network</b>				
VIP_OAM	node-1, node-2	eth1	External	The virtual IP of the Keepalived/HAProxy for the external O&M traffic and, if not using traffic separation, also provisioning traffic
ETH1_IP	node-(1, 2..., n)	eth1	External	IP of the node in the O&M network
ETH1_OAM_GW	node-(1, 2..., n)	eth1	External	Default Gateway used for outgoing O&M traffic and, if not using traffic separation, also provisioning traffic

(1) Only applicable if traffic separation is used. Traffic separation is not supported for CEE and OpenStack deployments.

**Table 3 IPv6 Allocation Example (Only applicable for KVM and VMware deployments)**

IP Address	VM	Interface	Type	Comments
<b>Internal Network</b>				
2001:0DB8:10:d3::/64	-	-	-	The Internal Network
:1	node-1	eth0	Internal	IP of the first VM
:2	node-2	eth0	Internal	IP of the second VM
:3	node-3	eth0	Internal	IP of the third VM
:254	node-1, node-2	eth0	Internal	The virtual IP of the Keepalived/HAProxy for the internal network which also works as the default gateway for the Real servers
<b>External Network</b>				
VIP_TRF_IPv6	node-1, node-2	eth2	External	The virtual IP of the Keepalived/HAProxy for the external provisioning network
ETH2_IPv6	node-1, node-2	eth2	External	IP of the node in the provisioning network



IP Address	VM	Interface	Type	Comments
ETH2_TRF_GW_IPv6	node-1, node-2	eth2	External	Default Gateway used for outgoing provisioning traffic in communication with NEs <sup>(1)</sup>
<b>External O&amp;M Network</b>				
VIP_OAM_IPv6	node-1, node-2	eth1	External	The virtual IP of the Keepalived/HAProxy for the external O&M traffic and, if not using traffic separation, also provisioning traffic
ETH1_IPv6	node-(1, 2..., n)	eth1	External	IP of the node in the O&M network
ETH1_OAM_GW_IPv6	node-(1, 2..., n)	eth1	External	Default Gateway used for outgoing O&M traffic and, if not using traffic separation, also provisioning traffic

(1) Only applicable if traffic separation is used.



## 3 Load Balancing

This section contains information about Virtual IP and Load Balancing.

The load balancing functionality is divided between the Keepalived and HAProxy solutions.

Virtual IP is handled by the Keepalived solution, and load balancing is handled by the HAProxy solution.

### 3.1 Keepalived and HAProxy

Keepalived uses the Virtual Router Redundancy Protocol (VRRP) to achieve a non-single point of failure.

HAProxy is used to achieve distributed load over the VMs. It works in a primary/backup setup, and to achieve a non-single point of failure setup. The primary/backup Keepalived/HAProxy load balancers will be installed on two physically separated hosts, see Figure 4.

**Note:** CEE and OpenStack deployments use common uplinks for O&M and Traffic.

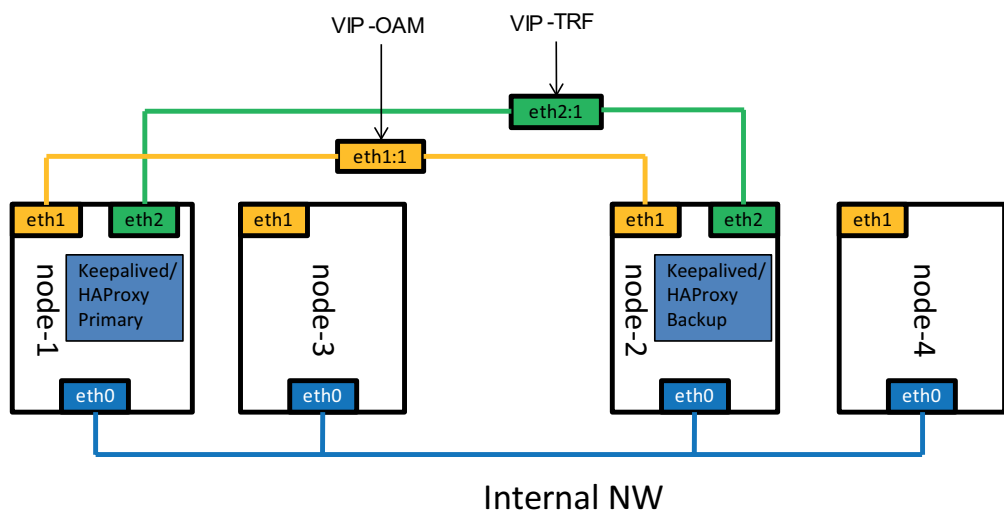


Figure 4 Load Balancing

For more information about Keepalived, see [Keepalived Loadbalancing & High-Availability](#), Reference [6].

For more information about HAProxy, see [HAProxy The Reliable, High Performance TCP/HTTP Load Balancer](#), Reference [7].



## 3.2 Keepalived/HAProxy Router

On VMs (primary/backup) that are acting as Keepalived/HAProxy routers, three network interfaces are configured. If using traffic separation (only applicable for KVM/VMware deployments), two interfaces are used for external communications, one for O&M and one for provisioning traffic. The third interface is used for internal traffic distribution between the VMs.

The primary (active) router distributes the provisioning traffic among the Real servers. In a failover scenario the backup (inactive) router becomes active and distributes provisioning traffic to the Real servers.

Traffic sourced from the Keepalived/HAProxy router uses, if traffic separation is used (only applicable for KVM/VMware deployments), the provisioning traffic interface `eth2` as source address.

All nodes in the cluster are processing traffic, even those acting as loadbalancers.

## 3.3 Real Server

On the Real servers, two interfaces are created, one for O&M and one for internal traffic. The Real servers receive and transmit all provisioning traffic on the internal network.



## 4 High Availability

For High Availability (HA) in a virtual or cloud environment, there are certain criteria that needs to be fulfilled. For redundancy purposes, the integrator needs to understand, if deviating from it.

For Dynamic Activation on virtualized or cloud deployment to be redundant, there are two important processes that need to reside on physical separated hosts:

- Keepalived/HAProxy – The load balancer (on node-1 and node-2)
- Zookeeper – Configuration database (on node-1, node-2, and node-3)
- Cassandra - Storage and SCM configuration database (on node-1, node-2, and node-3)

The figure below illustrates an example of where the crucial processes can reside to obtain full high availability.

**Note:** CEE and OpenStack deployments use common uplinks for O&M and Traffic.




Figure 5 Processes Overview for Full HA

When dedicating the external interface, for example `eth0` and `eth3` for internal traffic, it is important that they are not using the same NIC in case of failure, otherwise both interfaces will go down.

Furthermore, the Keepalived/HAProxy routers are located on two different hosts, in this case `Host-1` and `Host-2`, for redundancy in case a host goes down.

Internally of Dynamic Activation, there exists a storage engine called Zookeeper, which runs in a separate process. Its servers run on three different VMs. In order for the zookeeper to work, at least two instances need to be operational at the same time. Hence, to obtain high-availability for Zookeeper it needs to run on three different hosts, otherwise two instances will go down at the same time in case of host failure. In the above figure, zookeeper resides in `node-1`, `node-2`, and `node-3`.

The external default gateway needs to be a VIP (Virtual IP) using a protocol, for example, VRRP (Virtual Router Redundancy Protocol) for redundancy in the customer network.





## 5 Firewall Configuration

A number of ports need to be open to allow Dynamic Activation to function properly.

This means that the firewall on the hypervisor must grant access to the ports needed by Dynamic Activation.

For a list of ports that need to be accessible on the VMs, refer to *System Administrators Guide for Virtual and Cloud Deployment*, Reference [4].





## 6 Node Configuration

Each independent node is automatically configured as part of the deployment process that is described in *Software Installation for Virtual and Cloud Deployment*, Reference [3].

The configuration files used for each node can be edited based on the input from the *Customer Questionnaire for Virtual and Cloud Deployment*, Reference [5].





## 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] *System Administrators Guide for Virtual and Cloud Deployment*, 3/1543-CSH 109 628 Uen
- [5] *Customer Questionnaire for Virtual and Cloud Deployment*, 2/1057-CSH 109 628 Uen

### Online References

- [6] *Keepalived Loadbalancing and High-Availabiliy* <http://keepalived.org/>
- [7] *HAProxy The Reliable, High Performance TCP/HTTP Load Balancer* <http://www.haproxy.org/>