

Virtual IP Address Management

DESCRIPTION

Copyright

© Ericsson AB 2015. 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
2	Functions and Concepts	3
2.1	Types of Operation	8
3	Managed Object Model	11
4	Configuration Management	13
5	Fault Management	15
6	Security Management	17





1 Introduction

This document provides an overview of the management model and concepts associated with the Virtual IP Address Management managed area.

A managed area is represented by a group of Managed Object Classes (MOCs) within the Managed Object Model (MOM).





2 Functions and Concepts

The Virtual IP (VIP) Address Management concerns VIP addresses based on an embedded VIP framework.

VIP addressing is an established concept in the areas of resilient and scalable networking. However, different types of VIP addressing techniques are used in the industry. Some techniques fail to support common protocols and networking scenarios because of inherent limitations with the adopted technical methods. IPsec and SCTP are, for example, known to be difficult to accommodate with simple methods where the end-to-end IP protocol transparency is not upheld.

The term “VIP address” refers to a type of VIP technique preserving end-to-end-protocol transparency. The forwarding of IP packets is verbatim. That is, IP packets with a VIP address, used as a source or destination IP address in the IP packet header, are forwarded without any modification of information in the IP packet header. For example, no form of Network Address Translation (NAT) occurs with the type of VIP method that concerns this document. The inherent protocol complexities often imposed by NAT are thus avoided.

Address mobility in the VIP method covered in this document is not based on the Proxy Address Resolution Protocol (Proxy ARP) or similar types of ARP-based protocols. This avoids the limitations that come with solutions based on these protocols.

Purpose of VIP Addressing

The purpose of using VIP addressing is to achieve (some or all) of the following objectives:

- Hiding of internal addressing schemes to External Networks.
- Fault tolerance support to protect against network topology and intercommunication failure.
- Load sharing of traffic across Equal-Cost Multipaths (ECMPs).
- Dynamically controlled, on-the-fly, announcement and retraction of VIP addresses.
- VIP address migration to support local reconfiguration or resiliency cases.
- Geographical redundancy by global VIP address migration.
- Geographical redundancy by “any cast”.



The first three objectives in the list are typically relevant to most application product bundles. The other objectives depend on specific application product offerings and network deployment scenarios.

VIP Address

A VIP address is an IP address that does not represent a specific physical network interface (port). Instead a VIP address, which can be used as a source IP address or as a destination IP address, can be presented in the two following ways:

- Destination IP address. A VIP address represents the address of a functional entity inside a Network Element (NE) supporting one or more services. The services can be reached by many remote communicating clients using a VIP address as a destination IP address.
- Source IP address. A VIP address represents the address of a functional entity that originates one or more parallel service requests to one or more remote servers in the External Network.

For example, the VIP addresses can be used with typical client-server communication where the provided services would be associated with particular TCP or UDP port numbers.

The VIP addresses can also be used with dynamic routing in some deployments. In particular product bundlings, VIP addressing is offered with embedded support for the routing protocol Open Shortest Path First (OSPF) to attract IP traffic with VIP addresses. Thus, traffic can be attracted across a collection of physical paths and achieve “ECMP traffic load sharing” together with a resilient method for handling failover situations.

Abstract Load Balancer

The VIP addresses are configured to an Abstract Load Balancer (ALB) in the management model. The ALB is a logical container holding configured VIP addresses. Inside an NE, the logical function of an ALB can be distributed across a collection of physically separated processing units.

Figure 1 shows a set of ALBs, each connected to a separate VPN, and configured with a single IPv4 address as VIP address.




Figure 1 ALBs with VIP Addresses

In the External Network VIP, addresses are IP addresses that can be routed in the Data Communication Network (DCN). For example, IP packets, from remote clients in the DCN, with a VIP destination address are routed to a corresponding ALB.

Through an ALB, the VIP addresses can be used to address servers in a processing cluster, typically located inside a telecom node. For example, a collection of servers all listening to the same TCP port number, whereupon incoming TCP connections are load balanced over the servers.

Conversely, in the opposite direction, VIP addresses are used as source IP addresses by clients, typically when inside a telecom node and originating requests to remote servers in the DCN. For example, a client establishing a TCP connection with a remote server located in an external DCN.

An ALB is configured with one or more VIP addresses, which can be IPv4 or IPv6 addresses, or both. VIP addresses must be known to the External Network realm, that is, the VIP addresses of an ALB must addressing-wise be possible to route within a DCN. Each ALB can have several IPv4 and IPv6 addresses configured as VIP addresses, for example, a collection of non-contiguous addresses.

In typical deployments, 1–5 VIP addresses are configured to an ALB. However, up to 40 VIP addresses can be configured to an ALB, or 4 ALBs with 10 VIP addresses in each. The extra VIP addresses can be used as VIP equivalent



source addresses, which can be understood as alias addresses. More VIP addresses can be added in runtime, if needed.

A typical NE is configured with several networks used for separate purposes, such as separate Virtual Private Networks (VPN). Here overlapping IP addresses can occur, as each VPN is regarded as a separate network realm. Therefore, a separate ALB must be configured for each such VPN.

The VIP addresses configured to different ALBs in an NE must not overlap. That is, the VIP addresses configured to one ALB must not be reused for another ALB. However, other IP addresses than the specific VIP addresses configured to the ALBs can overlap between the separated networks (VPNs). Up to eight ALBs can be configured in the NE.

VIP Address and VIP Equivalent Source Address

In the most basic form, an ALB contains one VIP address, which is often sufficient. In the incoming traffic direction to the ALB, this VIP address can be regarded as a collective IP destination address, shared by services offered by the NE through this VIP address. In the outgoing direction, this VIP address can be regarded as a collective IP source address shared by a collection of clients located inside the NE.

More VIP addresses can be added to an ALB, if needed. An added VIP address must be routable in the connected DCN. The reasons for adding a VIP address to an ALB are typically one of the following:

- A high outgoing traffic volume causing a demand for more TCP or UDP port numbers. A resource that is allocated per IP address. In this case, the extra VIP address (or addresses) must be configured as a “VIP equivalent source address”.
- Grouping a set of services by an extra IP address so that the said services can, addressing-wise, be separated.

The purpose of a VIP equivalent source address is to overcome a high traffic bottleneck situation. This occurs when all ephemeral port numbers used for outgoing connections are consumed for the VIP address of an ALB. For example, for each new outgoing TCP connection a new ephemeral TCP source port number is consumed on the client side. Therefore, in a situation with a high rate of new connections, the available ephemeral port numbers can all have been consumed.

If an extra VIP address is configured as a VIP equivalent address in the same ALB, the bottleneck situation is avoided automatically. This is because the clients in the NE can continue to set up new TCP connections. The new connections are then given the VIP equivalent source address as source IP address. Hence, for outgoing connections in a high traffic situation, a VIP equivalent source address can replace any VIP address in the ALB as source IP address.



Figure 2 shows a set of ALBs, one ALB with both a VIP address and a VIP equivalent source address, and another ALB with two autonomous VIP addresses.

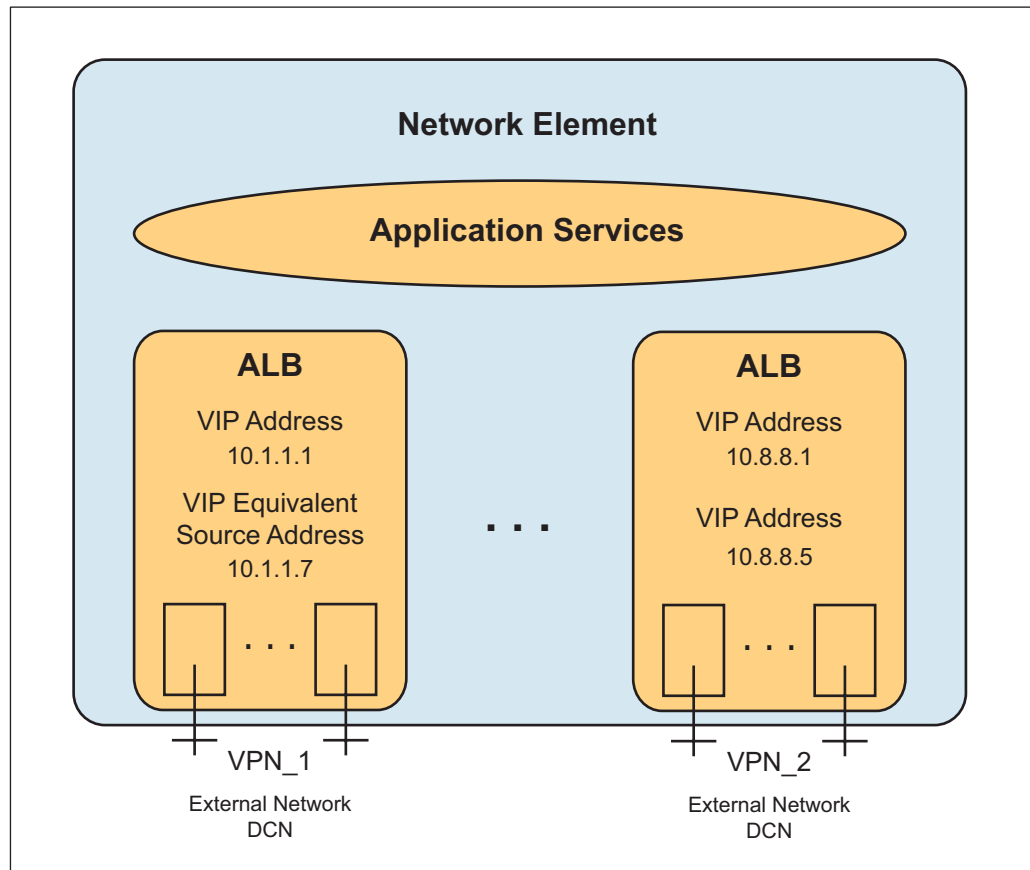


Figure 2 ALBs with VIP Addresses and VIP Equivalent Source Address

The extra VIP address (non-equivalent source address) can be used as a method for grouping a set of services so that they can be separated based on an IP address. For example, separated when routed in the DCN, so that the packets of two VIP addresses of the same ALB in the DCN network travel through different preferred paths. This is sometimes referred to as shared fate path diversity and is typically used with SCTP. For this purpose, policies to separate the traffic and service processing would typically be configured in the Customer-Premises Equipment (CPE) of the External Network and internally in the NE. If this type of diversity arrangement is used with TCP or UDP traffic, any extra VIP address in the ALB must not be configured as a VIP equivalent source address.

Grouping a set of services by an extra VIP address can also be used to separate the processing of application services in the NE based on configured “flow policies”, which considers the destination VIP address.



Flow Policies

The incoming traffic to an ALB has a destination IP address, matching one of the VIP addresses that has been configured to the ALB.

Flow policies are configured filters. The filters can segregate incoming traffic into different IP flows and direct these flows to internal functions, which represent the application services. Typically, such functions are allocated to a collection of different internal processing units. Target pools and socket groups are two internal system mechanisms used to abstract distributed internal functions associated with the application services.

Incoming traffic selected by a flow policy must match all defined matching attributes (logical AND) of the flow policy.

2.1 Types of Operation

Virtual IP Address Management supports the following operations:

- Add a VIP address to an ALB

A typical example is to add an extra VIP address to an ALB with an existing VIP address. The procedure in *Add Virtual IP Address* provides further details on how to perform this operation.

- Add a VIP equivalent source address to an ALB with an existing single VIP address

A VIP equivalent address is a VIP address with a specific attribute configured in this way. This is needed when there are insufficient ephemeral port numbers of the first configured VIP address, because of a high rate of outgoing connections. Typically, for reasons of symmetry regarding incoming and outgoing traffic cases, all afterwards added addresses to this ALB would be VIP equivalent source addresses and must then have this attribute set.

When static routing is used between the NE and the CPE, the CPE must be reconfigured with new static routes for the new VIP address. When OSPF is used to the CPE, a new added VIP address is automatically announced and service can start without any manual intervention on the CPE and network side.

Adding a VIP address to an ALB that is not a VIP equivalent address is typically only done as part of a major network reconfiguration activity.

The procedure in *Add Virtual IP Equivalent Source Address* provides further details on how to perform this operation.



Note: Any runtime modification of VIP addresses must be done in a coordinated way, which considers the impact on the External Network and application specifics. That is, impact regarding the behavior of the specific applications of the NE that use VIP addresses must be thoroughly understood.

Also, modifying VIP addresses can for some applications not take effect unless the application is restarted.

- Add a flow policy to an ALB

Target pools and socket groups are mutually exclusive attribute choices configured in a flow policy. A target pool or a socket group is a destination target for the segregated packet flows. For application services that are to be reached by TCP or UDP traffic, the attribute target pool is the relevant configuration choice, whereas socket groups are primarily used for SCTP traffic. The procedure in *Add Flow Policy* provides further details on how to perform this operation.

Note: Misconfiguration of flow policies can lead to black-holing of traffic and can cause a complete disruption of service.

Note: Removing VIP addresses can be dangerous and requires deep coordination with specific applications. For example, stale information in the form of port number resources must be purged.



3 Managed Object Model

The Virtual IP Address Management managed area is represented in the *Managed Object Model (MOM)* as follows:

```
ManagedElement
+-Transport
+-Evip
+-EvipAlbs
+-EvipAlb
+-EvipFlowPolicies
+-EvipFlowPolicy
+-EvipVips
+-EvipVip
```

For general information about the MOM, MOCs, Managed Objects (MOs), cardinality, and related concepts, refer to *Managed Object Model User Guide*.

The Virtual IP Address Management MOCs are described in Table 1.

Table 1 Virtual IP Address Management Managed Object Class Descriptions

Managed Object Class	Description
<i>Evip</i>	The root of the Virtual IP Address Management model.
<i>EvipAlbs</i>	Container for ALBs.
<i>EvipAlb</i>	Defines and handles an ALB.
<i>EvipFlowPolicies</i>	Container for flow policies.
<i>EvipFlowPolicy</i>	Defines a flow policy.
<i>EvipVips</i>	Container for VIP addresses.
<i>EvipVip</i>	Defines a VIP address or a VIP equivalent source address.





4 Configuration Management

Virtual IP Address Management is accessed using NETCONF or the Ericsson Command-Line Interface (ECLI) to manipulate the Management Information Base (MIB).

The following operations can be performed by the user and are described in an Operating Instruction using the ECCLI:

- *Add Virtual IP Address*
- *Add Virtual IP Equivalent Source Address*
- *Add Flow Policy*





5 Fault Management

The Virtual IP Address Management alarms are described in Table 2.

Table 2 Virtual IP Address Management Alarms

Alarm	Description
<i>eVIP, Gateway Unavailable</i>	Raised when contact is lost with an external gateway.
<i>eVIP, IPSEC Tunnel Fault</i>	Raised when an IPsec tunnel goes down ungracefully between a VIP-enabled cluster and a peer.
<i>eVIP, IKE Distribution Not Possible</i>	Raised when the distribution of the Internet Key Exchange (IKE) processes cannot be resolved and there are no available blades for every IKE instance.

The Virtual IP Address Management-related event is described in Table 3.

Table 3 Virtual IP Address Management Event

Event	Description
<i>eVIP, Configuration Fault</i>	Reported when Virtual IP Address Management detects a faulty configuration.





6 Security Management

One Virtual IP Address Management role is defined, named System Administrator.

Once authenticated as a System Administrator, full access is granted to the *Transport* MO, its attributes, and actions.

For more information, refer to *Security Management for ECLI, NETCONF, and SFTP Users*.