

Virtual IP Addressing

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

1	Understanding Virtual IP Addressing	1
1.1	Key Virtual IP Addressing Concepts	1
1.2	Virtual IP Address	2
1.3	Abstract Load Balancer	3
1.4	Flow Policies	4
1.5	Named VIP Address	5
1.6	Equivalent Virtual IP Address	5
2	Basic Virtual IP Addressing Procedures	8
3	Virtual IP Addressing-related Alarms and Events	9
4	Security Management	10





1 Understanding Virtual IP Addressing

1.1 Key Virtual IP Addressing Concepts

The Virtual IP (VIP) Address Management here described concerns a VIP addressing system framework for embedded Virtual IP deployment in Network Elements (NE). This Virtual IP system framework is used by several applications across different Network Elements (NE) provided by Ericsson.

The concept of Virtual IP (VIP) addressing implies the use of an IP address that does not correspond to a specific hardware networking interface. In general this is an established networking practice typically used in conjunction with resilient and scalable NEs. However, a variety of methods are used in the industry and the main distinguishing property of the used VIP technique can be described in the following way:

With the VIP addressing method here concerned, internally inside the NE, packet processing and transport is verbatim and therefore this method is particularly suitable for networking scenarios with protocols which by design assumes that end-to-end IP protocol transparency is preserved. For instance, this is typically the case with application protocols that are based on IPsec or SCTP.

The supported verbatim packet processing inside the NE implies that an IP packet which contains a VIP address in the IP packet header, used either as a source IP address or destination IP address, is forwarded without any modification of the information in the IP packet header. For example, no form of Network Address Translation (NAT) occurs in the NE with the used VIP addressing method.

The term “VIP address” henceforth refers to VIP addressing with the here stated method. The VIP addresses are within the said system framework configured to Abstract Load Balancers (ALBs) which are logical software entities embedded in the NE. The external interfaces of an ALB are called the front-end interfaces and are terminated in the Front End Elements (FEEs) of the ALB. Furthermore, the front-end interfaces are connected to external gateways which are connected with the DCN.

The virtual IP addressing framework uses the following listed key concepts:

- Virtual IP Address; this is an IP address that is not tied to a specific hardware networking interface.
- Abstract Load Balancers; are logical containers holding configuration structures for VIP addresses and Front End Elements (FEEs).
 - Front End Elements; belong to an ALB and terminates the front-end interfaces of the FEEs which provides connectivity to external gateways that are connected to a DCN.



- Flow Policies; are IP packet filters that are configured within the scope of an ALB for separating the incoming traffic flows to different internal functions in the NE.
- Named Virtual IP Address; are configuration-wise a type of VIP Addresses that can be referred to by user defined names instead of actual IP address values.
- Equivalent Virtual IP Source Address; is a VIP address which is used as an alias address for a source VIP address.

The purpose of using VIP addressing in a NE as here described, is to achieve some of (or all of), the following objectives:

- Hiding of internal addressing schemes to External Networks.
- Achieve a decoupling of an IP address from a specific hardware interface.
- Fault tolerance support to protect against network topology rupture and intercommunication failure.
- Load sharing of traffic across Equal-Cost Multipaths (ECMPs).
- Dynamical routing (OSPF), on-the-fly, announcement, and retraction of VIP addresses.
- VIP address migration to support local reconfiguration or resiliency failover cases.
- Geographical redundancy by global VIP address migration.
- Geographical redundancy by “VIP any cast”.

Which objectives are desired depends on NE configuration and the type of application that is to be deployed in the NE.

The Virtual IP Address Management managed area, *Evip*, can be found in the Managed Object Model (MOM). For general information about the MOM, Managed Object Classes (MOCs), cardinality, and related concepts, refer to *Managed Object Model User Guide*.

1.2 Virtual IP Address

A Virtual IP address (VIP address) is an IP address that can be used 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 destination IP address which matches a VIP address which is configured in the NE.



- **Source IP address.** A VIP address represents the address of a functional entity in the NE 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.

A VIP address is an IP address, but does not represent a specific physical network interface (port).

The IP address value of a VIP address is configured using industry typical textual notation formats. For IPv4, the decimal dot notation format is used, and for IPv6, the canonical textual form is used.

1.3 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 the structures for configured VIP addresses. Inside a Network Element (NE), the logical function of an ALB can be distributed across a collection of physically separated processing units.

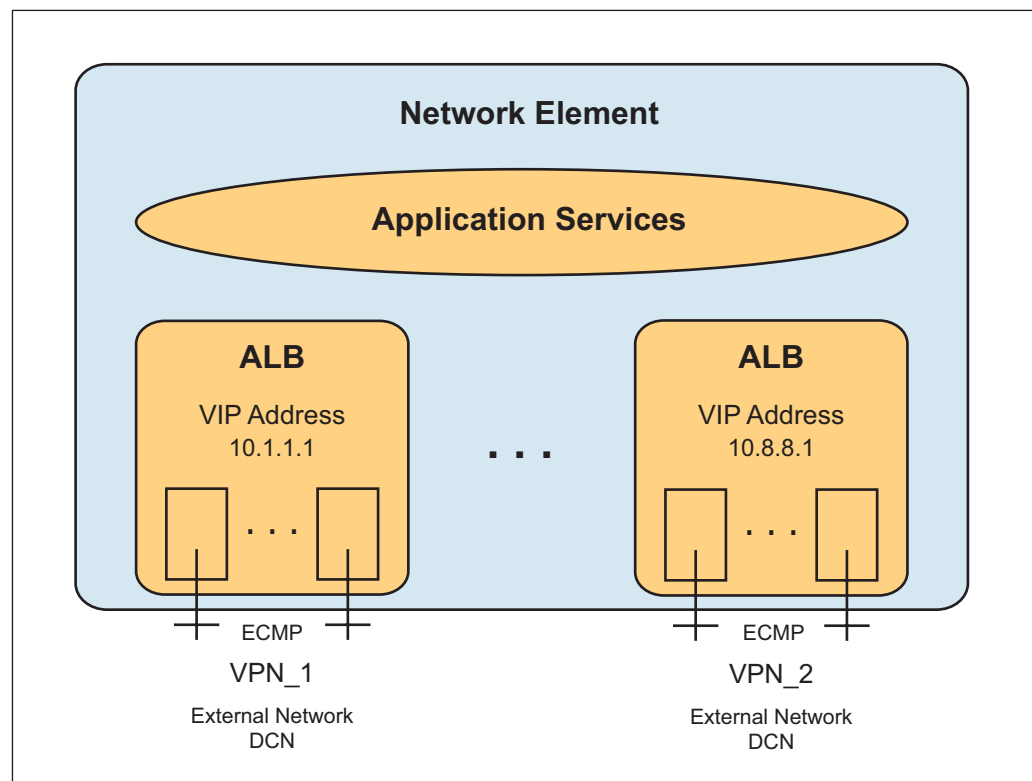


Figure 1 ALBs with VIP Addresses



In the External Network connected to a ALB, VIP addresses, are IP addresses of packets which addressing-wise can be routed in the Data Communication Network (DCN). For example, IP packets, from remote clients in the DCN that have a VIP destination address can thus be routed in DCN to the external interfaces of a corresponding ALB in the NE.

In the incoming traffic direction an ALB delivers traffic arriving from DCN to an application in the NE. In the outgoing direction an ALB delivers traffic, from an application inside the NE to gateway routers connected to the DCN.

The external interfaces of an ALB are called the front-end interfaces and are configured in the Front End Elements (FEEs) of the ALB. The FEEs are logical entities in the ALB and acts as IP routers which through the corresponding front-end interfaces are connected with one or two gateway routers which are present in the DCN.

Through an ALB, the VIP addresses can be used to address servers in a processing cluster located inside a NE, which typically corresponds to a telecom node.

For example, a collection of servers in the NE that are all listening to the same TCP port number, and upon arrival of incoming TCP traffic to the NE, the TCP connections are then load balanced over the said collection of servers.

Conversely, in the opposite direction towards the DCN, VIP addresses are used as source IP addresses by clients inside a NE, as they are originating requests to remote servers in the DCN. For example, a client inside the NE originating 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.

Moreover, each ALB can have several IPv4 and IPv6 addresses configured to it as VIP addresses. The VIP addresses can be chosen from a collection of non-contiguous IP addresses.

In case the DCN is comprised of more than one Virtual Private Network (VPN) connected to the NE using the same embedded framework for VIP addressing, then a separate ALB is used for each VPN connected to the Network Element.

1.4 Flow Policies

Flow Policies are packet filters that are configured within the scope of an ALB.

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

The Flow Policies can segregate the incoming traffic to an ALB 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. Flow Policies can be used to segregate and direct the incoming traffic from DCN to different application services. For example, a first Flow Policy can be configured to direct traffic with a first VIP destination address to a first target pool and a second Flow Policy can be configured to direct traffic with a second VIP destination address to a second target pool.

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

1.5 Named VIP Address

There are two different ways of configuring a VIP address; the VIP address can be configured either as an Explicit VIP address or as a Named VIP address:

- **Explicit VIP Address:** Explicit VIP addresses are after creation, configuration-wise, always referred to by their explicit IP address value. An explicitly configured VIP address must always be referred to by its IP address value, for example 10.8.8.1, in a configured Flow Policy.
- **Named VIP Address:** Named VIP addresses are, configuration-wise, created with individual user-defined names, which, for configuration purposes, are used to indicate the VIP address by its given name instead of its IP address value. The names given to the VIP addresses must be unique within the scope of an ALB. VIP addresses configured as named VIP addresses must always, configuration-wise, be referred to by their given names in a configured Flow Policy.

The use of Named VIP Addresses allows for Flow Policies to be preconfigured ahead of obtained precise knowledge of the actual IP addresses to be filtered on.

1.6 Equivalent Virtual IP 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 an 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.

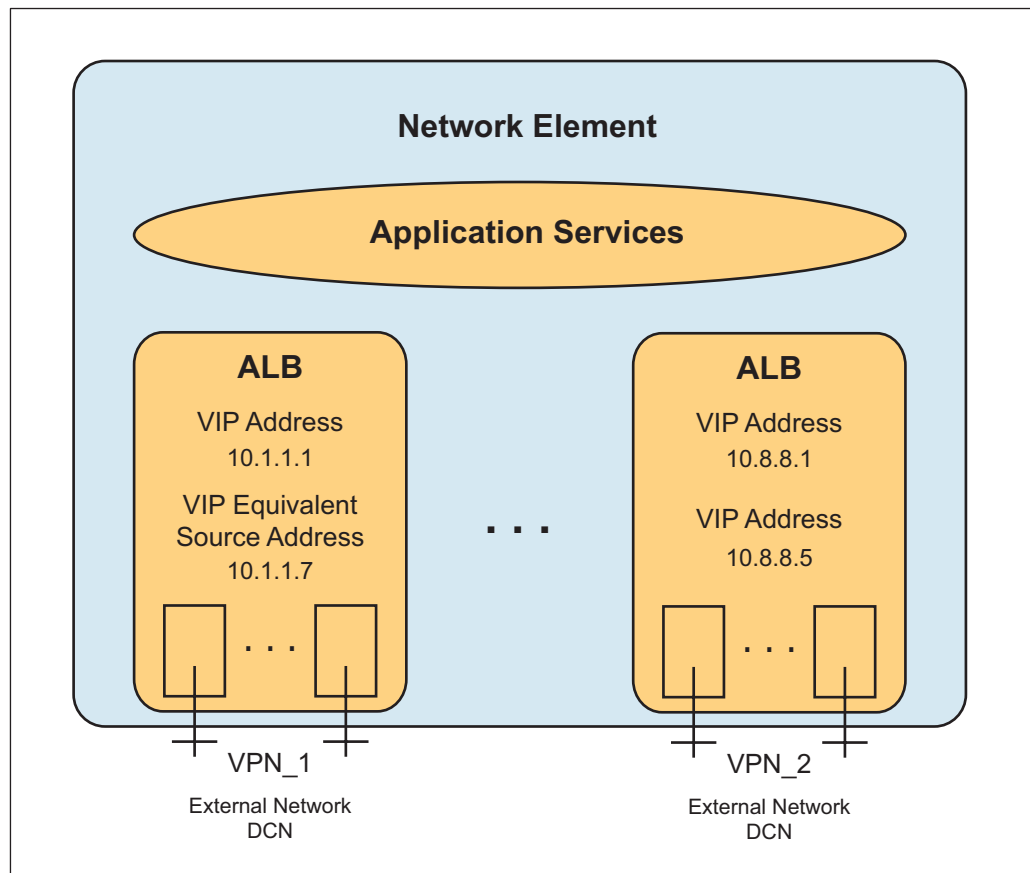


Figure 2 ALBs with VIP Addresses And VIP Equivalent Source Address

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.



2 Basic Virtual IP Addressing Procedures

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

Virtual IP Addressing supports the following operations described in OPIs:

- *Add Virtual IP Address*

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 when configuring an Explicit VIP address.

- *Add Named Virtual IP Address*

In case it has been deemed preferable to configure the VIP address as a Named VIP address, for example, in a situation where the precise IP address value of the VIP address is to be obtained later.

- *Add Virtual IP Equivalent Source 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.

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 Flow Policy*



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.

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

3 Virtual IP Addressing-related Alarms and Events

Table 1 Virtual IP Addressing Related 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.

Table 2 Virtual IP Addressing Related Events

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



4 Security Management

One Virtual IP Addressing 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 *User Management*.