

Network Working Group
Request for Comments: 4026
Category: Informational

L. Andersson
T. Madsen
Acreo AB
March 2005

Provider Provisioned Virtual Private Network (VPN) Terminology

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

The widespread interest in provider-provisioned Virtual Private Network (VPN) solutions lead to memos proposing different and overlapping solutions. The IETF working groups (first Provider Provisioned VPNs and later Layer 2 VPNs and Layer 3 VPNs) have discussed these proposals and documented specifications. This has lead to the development of a partially new set of concepts used to describe the set of VPN services.

To a certain extent, more than one term covers the same concept, and sometimes the same term covers more than one concept. This document seeks to make the terminology in the area clearer and more intuitive.

Table of Contents

1.	Introduction	3
2.	PPVPN Terminology	3
3.	Provider Provisioned Virtual Private Network Services	4
3.1.	Layer 3 VPN (L3VPN)	4
3.2.	Layer 2 VPN (L2VPN)	4
3.3.	Virtual Private LAN Service (VPLS)	4
3.4.	Virtual Private Wire Service (VPWS)	4
3.5.	IP-Only LAN-Like Service (IPLS)	5
3.6.	Pseudo Wire (PW)	5
3.7.	Transparent LAN Service (TLS)	5
3.8.	Virtual LAN (VLAN)	6
3.9.	Virtual Leased Line Service (VLLS)	6
3.10.	Virtual Private Network (VPN)	6
3.11.	Virtual Private Switched Network (VPSN)	6

4.	Classification of VPNs	7
5.	Building Blocks	8
5.1.	Customer Edge Device (CE)	8
5.1.1.	Device Based CE Naming	9
5.1.2.	Service Based CE Naming	9
5.2.	Provider Edge (PE)	10
5.2.1.	Device Based PE Naming	10
5.2.2.	Service Based PE Naming	10
5.2.3.	Distribution Based PE Naming	11
5.3.	Core	11
5.3.1.	Provider Router (P)	11
5.4.	Naming in Specific Internet Drafts	11
5.4.1.	Layer 2 PE (L2PE)	11
5.4.2.	Logical PE (LPE)	12
5.4.3.	PE-CLE	12
5.4.4.	PE-Core	12
5.4.5.	PE-Edge	12
5.4.6.	PE-POP	12
5.4.7.	VPLS Edge (VE)	12
6.	Functions	12
6.1.	Attachment Circuit (AC)	12
6.2.	Backdoor Links	13
6.3.	Endpoint Discovery	13
6.4.	Flooding	13
6.5.	MAC Address Learning	13
6.5.1.	Qualified Learning	13
6.5.2.	Unqualified Learning	13
6.6.	Signalling	13
7.	'Boxes'	14
7.1.	Aggregation Box	14
7.2.	Customer Premises Equipment (CPE)	14
7.3.	Multi-Tenant Unit (MTU)	14
8.	Packet Switched Network (PSN)	14
8.1.	Route Distinguisher (RD)	15
8.2.	Route Reflector	15
8.3.	Route Target (RT)	15
8.4.	Tunnel	15
8.5.	Tunnel Multiplexor	16
8.6.	Virtual Channel (VC)	16
8.7.	VC Label	16
8.8.	Inner Label	16
8.9.	VPN Routing and Forwarding (VRF)	16
8.10.	VPN Forwarding Instance (VFI)	16
8.11.	Virtual Switch Instance (VSI)	17
8.12.	Virtual Router (VR)	17
9.	Security Considerations	17
10.	Acknowledgements	17
11.	Informative References	17

Authors' Addresses	19
Full Copyright Statement	20

1. Introduction

A comparatively large number of memos have been submitted to the former PPVPN working group, and to the L2VPN, L3VPN, and PWE3 working groups, which all address the same problem space; provider provisioned virtual private networking for end customers. The memos address a wide range of services, but there is also a great deal of commonality among the proposed solutions.

This has led to the development of a partial set of new concepts used to describe this set of VPN services. To a certain extent, more than one term covers the same concept, and sometimes the same term covers more than one concept.

This document proposes a foundation for a unified terminology for the L2VPN and L3VPN working groups. In some cases, the parallel concepts within the PWE3 working group are used as references.

2. PPVPN Terminology

The concepts and terms in this list are gathered from Internet Drafts sent to the L2VPN and L3VPN mailing lists (earlier the PPVPN mailing list) and RFCs relevant to the L2VPN and L3VPN working groups. The focus is on terminology and concepts that are specific to the PPVPN area, but this is not strictly enforced; e.g., some concepts and terms within the PWE3 and (Generalized) MPLS areas are closely related. We've tried to find the earliest uses of terms and concepts.

This document is intended to fully cover the concepts within the core documents from the L2VPN and L3VPN working groups; i.e., [L3VPN-REQ], [L2VPN-REQ], [L3VPN-FRAME], [L2VPN], and [RFC3809]. The intention is to create a comprehensive and unified set of concepts for these documents and, by extension, for the entire PPVPN area. To do so, it is also necessary to give some of the development the concepts of the area have been through.

The document is structured in four major sections. Section 4 lists the different services that have been or will be specified Section 5 lists the building blocks that are used to specify those services Section 6 lists the functions needed in those services. Section 7 lists some typical devices used in customer and provider networks.

3. Provider Provisioned Virtual Private Network Services

In this section, we define the terminology that relates the set of services to solutions specified by the L2VPN and L3VPN working groups. The "pseudo wire" concept, which belongs to the PWE3 working group, is included for reference purposes. For requirements in provider provisioned VPNs, see [L3VPN-REQ].

All terms and abbreviations are listed together with a brief description of the service. The list is structured to give the more general information first and the more specific later. The names of services for which the IETF is working on solutions have been moved to the top of the list. Older and more dated terminology has been pushed toward the end of the list.

3.1. Layer 3 VPN (L3VPN)

An L3VPN interconnects sets of hosts and routers based on Layer 3 addresses; see [L3VPN-FRAME].

3.2. Layer 2 VPN (L2VPN)

Three types of L2VPNs are described in this document: Virtual Private Wire Service (VPWS) (Section 3.4); Virtual Private LAN Service (VPLS) (Section 3.3); and IP-only LAN-like Service (IPLS) (Section 3.5).

3.3. Virtual Private LAN Service (VPLS)

A VPLS is a provider service that emulates the full functionality of a traditional Local Area Network (LAN). A VPLS makes it possible to interconnect several LAN segments over a packet switched network (PSN) and makes the remote LAN segments behave as one single LAN. For an early work on defining a solution and protocol for a VPLS, see [L2VPN-REQ], [VPLS-LDP], and [VPLS].

In a VPLS, the provider network emulates a learning bridge, and forwarding decisions are taken based on MAC addresses or MAC addresses and VLAN tag.

3.4. Virtual Private Wire Service (VPWS)

A Virtual Private Wire Service (VPWS) is a point-to-point circuit (link) connecting two Customer Edge devices. The link is established as a logical through a packet switched network. The CE in the customer network is connected to a PE in the provider network via an Attachment Circuit (see Section 6.1); the Attachment Circuit is either a physical or a logical circuit.

The PEs in the core network are connected via a PW.

The CE devices can be routers, bridges, switches, or hosts. In some implementations, a set of VPWSs is used to create a multi-site L2VPN network. An example of a VPWS solution is described in [PPVPN-L2VPN].

A VPWS differs from a VPLS (Section 3.3) in that the VPLS is point to multipoint, while the VPWS is point to point. See [L2VPN].

3.5. IP-Only LAN-Like Service (IPLS)

An IPLS is very like a VPLS (see Section 3.3), except that

- o it is assumed that the CE devices (see Section 5.1) are hosts or routers, not switches,
- o it is assumed that the service will only have to carry IP packets, and supporting packets such as ICMP and ARP (otherwise layer 2 packets that do not contain IP are not supported); and
- o the assumption that only IP packets are carried by the service applies equally to IPv4 and IPv6 packets.

While this service is a functional subset of the VPLS service, it is considered separately because it may be possible to provide it by using different mechanisms, which may allow it to run on certain hardware platforms that cannot support the full VPLS functionality [L2VPN].

3.6. Pseudo Wire (PW)

The PWE3 working group within the IETF specifies the pseudo wire technology. A pseudo wire is an emulated point-to-point connection over a packet switched network that allows the interconnection of two nodes with any L2 technology. The PW shares some of the building blocks and architecture constructs with the point-to-multipoint solutions; e.g., PE (see Section 5.2) and CE (see Section 5.1). An early solution for PWs is described in [TRANS-MPLS]. Encapsulation formats readily used in VPWS, VPLS, and PWs are described in [ENCAP-MPLS]. Requirements for PWs are found in [RFC3916], and [PWE3-ARCH] presents an architectural framework for PWs.

3.7. Transparent LAN Service (TLS)

TLS was an early name used to describe the VPLS service. TLS has been replaced by VPLS, which is the current term.

3.8. Virtual LAN (VLAN)

The term VLAN was specified by IEEE 802.1Q; it defines a method of differentiating traffic on a LAN by tagging the Ethernet frames. By extension, VLAN is used to mean the traffic separated by Ethernet frame tagging or similar mechanisms.

3.9. Virtual Leased Line Service (VLLS)

The term VLLS has been replaced by term VPWS. VLLS was used in a now dated document intended to create metrics by which it should have been possible to compare different L2VPN solutions. This document has now expired, and the work has been terminated.

3.10. Virtual Private Network (VPN)

VPN is a generic term that covers the use of public or private networks to create groups of users that are separated from other network users and that may communicate among them as if they were on a private network. It is possible to enhance the level of separation (e.g., by end-to-end encryption), but this is outside the scope of IETF VPN working group charters. This VPN definition is from [RFC2764].

In the [L3VPN-FRAME], the term VPN is used to refer to a specific set of sites as either an intranet or an extranet that have been configured to allow communication. Note that a site is a member of at least one VPN and may be a member of many.

In this document, "VPN" is also used as a generic name for all services listed in Section 3.

3.11. Virtual Private Switched Network (VPSN)

The term VPSN has been replaced by the term VPLS. The requirements have been merged into the L3VPN [L3VPN-REQ] and L2VPN [L2VPN-REQ] requirements.

4. Classification of VPNs

The terminology used in [RFC3809] is defined based on the figure below.




Figure 1

The figure above presents a taxonomy of PPVPN technologies. Some of the definitions are given below:

CE-based VPN: A VPN approach in which the shared service provider network does not have any knowledge of the customer VPN. This information is limited to CE equipment. All the VPN-specific procedures are performed in the CE devices, and the PE devices are not aware in any way that some of the traffic they are processing is VPN traffic (see also [L3VPN-FRAME]).

PE-Based VPNs: A Layer 3 VPN approach in which a service provider network is used to interconnect customer sites using shared resources. Specifically, the PE device maintains VPN state, isolating users of one VPN from users of another. Because the PE device maintains all required VPN states, the CE device may behave as if it were connected to a private network. Specifically, the CE in a PE-based VPN must not require any changes or additional functionality to be connected to a PPVPN instead of a private network.

The PE devices know that certain traffic is VPN traffic. They forward the traffic (through tunnels) based on the destination IP address of the packet, and optionally based on other information in the IP header of the packet. The PE devices are themselves the tunnel endpoints. The tunnels may make use of various encapsulations to send traffic over the SP network (such as, but not restricted to, GRE, IP-in-IP, IPsec, or MPLS tunnels) [L3VPN-FRAME].

Virtual Router (VR) style: A PE-based VPN approach in which the PE router maintains a complete logical router for each VPN that it supports. Each logical router maintains a unique forwarding table and executes a unique instance of the routing protocols. These VPNs are described in [L3VPN-VR].

BGP/MPLS IP VPNs: A PE-based VPN approach in which the PE router maintains a separate forwarding environment and a separate forwarding table for each VPN. In order to maintain multiple forwarding table instances while running only a single BGP instance, BGP/MPLS IP VPNs mark route advertisements with attributes that identify their VPN context. These VPNs are based on the approach described in [RFC2547bis].

RFC 2547 Style: The term has been used by the L3VPN to describe the extensions of the VPNs defined in the informational RFC 2547 [RFC2547]. This term has now been replaced by the term BGP/MPLS IP VPNs.

5. Building Blocks

Starting with specifications of L3VPNs (e.g., the 2547 specification [RFC2547] and [RFC2547bis] and Virtual Routers [L3VPN-VR]), a way of describing the building blocks and allocation of functions in VPN solutions was developed. The building blocks are often used in day-to-day talk as if they were physical boxes, common for all services.

However, for different reasons, this is an oversimplification. Any of the building blocks could be implemented across more than one physical box. How common the use of such implementations will be is beyond the scope of this document.

5.1. Customer Edge Device (CE)

A CE is the name of the device with the functionality needed on the customer premises to access the services specified by the former PPVPN working group in relation to the work done on L3VPNs [L3VPN-FRAME]. The concept has been modified; e.g., when L2VPNs and CE-based VPNs were defined. This is addressed further in the sub-sections of this section.

There are two different aspects that have to be considered in naming CE devices. One could start with the type of device that is used to implement the CE (see Section 5.1.1). It is also possible to use the service the CE provides whereby the result will be a set of "prefixed CEs", (see Section 5.1.2).

It is common practice to use "CE" to indicate any of these boxes, as it is very often unambiguous in the specific context.

5.1.1. Device Based CE Naming

5.1.1.1. Customer Edge Router (CE-R)

A CE-R is a router in the customer network interfacing the provider network. There are many reasons to use a router in the customer network; e.g., in an L3VPN using private IP addressing, this is the router that is able to do forwarding based on the private addresses. Another reason to require the use of a CE-R on the customer side is that one wants to limit the number of MAC-addresses that need to be learned in the provider network.

A CE-R could be used to access both L2 and L3 services.

5.1.1.2. Customer Edge Switch (CE-S)

A CE-S is a service aware L2 switch in the customer network interfacing the provider network. In a VPWS or a VPLS, it is not strictly necessary to use a router in the customer network; a layer 2 switch might very well do the job.

5.1.2. Service Based CE Naming

The list below contains examples of how different functionality has been used to name CEs. There are many examples of this type of naming, and we only cover the most frequently used functional names. As these are functional names, it is quite possible that on a single piece of equipment there are platforms for more than one type of function. For example, a router might at the same time be both a L2VPN-CE and a L3VPN-CE. It might also be that the functions needed for a L2VPN-CE or L3VPN-CE are distributed over more than one platform.

5.1.2.1. L3VPN-CE

An L3VPN-CE is the device or set of devices on the customer premises that attaches to a provider provisioned L3VPN; e.g., a 2547bis implementation.

5.1.2.2. VPLS-CE

A VPLS-CE is the device or set of devices on the customer premises that attaches to a provider provisioned VPLS.

5.1.2.3. VPWS-CE

A VPWS-CE is the device or set of devices on the customer premises that attaches to a provider provisioned VPWS.

5.2. Provider Edge (PE)

A PE is the name of the device or set of devices at the edge of the provider network with the functionality that is needed to interface with the customer. Without further qualifications, PE is very often used for naming the devices since it is made unambiguous by the context.

In naming PEs there are three aspects that we need to consider, the service they support, whether the functionality needed for service is distributed across more than one device and the type of device they are build on.

5.2.1. Device Based PE Naming

Both routers and switches may be used to implement PEs; however, the scaling properties will be radically different depending on which type of equipment is chosen.

5.2.1.1. Provider Edge Router (PE-R)

A PE-R is a L3 device that participates in the PSN (see Section 8) routing and forwards packets based on the routing information.

5.2.1.2. Provider Edge Switch (PE-S)

A PE-S is a L2 device that participates in for example a switched Ethernet taking forwarding decision packets based on L2 address information.

5.2.2. Service Based PE Naming

5.2.2.1. L3VPN-PE

An L3VPN-PE is a device or set of devices at the edge of the provider network interfacing the customer network, with the functionality needed for an L3VPN.

5.2.2.2. VPWS-PE

A VPWS-PE is a device or set of devices at the edge of the provider network interfacing the customer network, with the functionality needed for a VPWS.

5.2.2.3. VPLS-PE

A VPLS-PE is a device or set of devices at the edge of the provider network interfacing the customer network, with the functionality needed for a VPLS.

5.2.3. Distribution Based PE Naming

For scaling reasons, in the VPLS/VPWS cases sometimes it is desired to distribute the functions in the VPLS/VPWS-PE across more than one device. For example, is it feasible to allocate MAC address learning on a comparatively small and inexpensive device close to the customer site, while participation in the PSN signalling and setup of PE to PE tunnels are done by routers closer to the network core.

When distributing functionality across devices, a protocol is needed to exchange information between the Network facing PE (N-PE) (see Section 5.2.3.1) and the User facing PE (U-PE) (see Section 5.2.3.2).

5.2.3.1. Network Facing PE (N-PE)

The N-PE is the device to which the signalling and control functions are allocated when a VPLS-PE is distributed across more than one box.

5.2.3.2. User Facing PE (U-PE)

The U-PE is the device to which the functions needed to take forwarding or switching decisions at the ingress of the provider network.

5.3. Core

5.3.1. Provider Router (P)

The P is defined as a router in the core network that does not have interfaces directly toward a customer. Therefore, a P router does not need to keep VPN state and is VPN unaware.

5.4. Naming in Specific Internet Drafts

5.4.1. Layer 2 PE (L2PE)

L2PE is the joint name of the devices in the provider network that implement L2 functions needed for a VPLS or a VPWS.

5.4.2. Logical PE (LPE)

The term Logical PE (LPE) originates from a dated Internet Draft, "VPLS/LPE L2VPNs: Virtual Private LAN Services using Logical PE Architecture", and was used to describe a set of devices used in a provider network to implement a VPLS. In a LPE, VPLS functions are distributed across small devices (PE-Edges/U-PE) and devices attached to a network core (PE-Core/N-PE). In an LPE solution, the PE-edge and PE-Core can be interconnected by a switched Ethernet transport network or uplinks. The LPE will appear to the core network as a single PE. In this document, the devices that constitutes, the LPE are called N-PE and U-PE.

5.4.3. PE-CLE

An alternative name for the U-PE suggested in the expired Internet Draft, "VPLS architectures".

5.4.4. PE-Core

See the origins and use of this concept in Section 5.4.2.

5.4.5. PE-Edge

See the origins and use of this concept in Section 5.4.2.

5.4.6. PE-POP

An alternative name for the U-PE suggested in the expired Internet Draft, "VPLS architectures".

5.4.7. VPLS Edge (VE)

The term VE originates from a dated Internet Draft on a distributed transparent LAN service and was used to describe the device used by a provider network to hand off a VPLS to a customer. In this document, the VE is called a VPLS-PE. This name is dated.

6. Functions

In this section, we have grouped a number of concepts and terms that have to be performed to make the VPN services work.

6.1. Attachment Circuit (AC)

In a Layer 2 VPN the CE is attached to PE via an Attachment Circuit (AC). The AC may be a physical or logical link.

6.2. Backdoor Links

Backdoor Links are links between CE devices that are provided by the end customer rather than by the SP; they may be used to interconnect CE devices in multiple-homing arrangements [L3VPN-FRAME].

6.3. Endpoint Discovery

Endpoint discovery is the process by which the devices that are aware of a specific VPN service will find all customer facing ports that belong to the same service.

The requirements on endpoint discovery and signalling are discussed in [L3VPN-REQ]. It was also the topic in a now dated Internet Draft reporting from a design team activity on VPN discovery.

6.4. Flooding

Flooding is a function related to L2 services; when a PE receives a frame with an unknown destination MAC address, that frame is send out over (flooded) every other interface.

6.5. MAC Address Learning

MAC address learning is a function related to L2 services; when PE receives a frame with an unknown source MAC address, the relationship between that MAC-address and interface is learned for future forwarding purposes. In a layer 2 VPN solution from the L2VPN WG, this function is allocated to the VPLS-PE.

6.5.1. Qualified Learning

In qualified learning, the learning decisions at the U-PE are based on the customer Ethernet frame's MAC address and VLAN tag, if a VLAN tag exists. If no VLAN tag exists, the default VLAN is assumed.

6.5.2. Unqualified Learning

In unqualified learning, learning is based on a customer Ethernet frame's MAC address only.

6.6. Signalling

Signalling is the process by which the PEs that have VPNs behind them exchange information to set up PWs, PSN tunnels, and tunnel multiplexers. This process might be automated through a protocol or done by manual configuration. Different protocols may be used to establish the PSN tunnels and exchange the tunnel multiplexers.

7. 'Boxes'

We list a set of boxes that will typically be used in an environment that supports different kinds of VPN services. We have chosen to include some names of boxes that originate outside the protocol specifying organisations.

7.1. Aggregation Box

The aggregation box is typically an L2 switch that is service unaware and is used only to aggregate traffic to more function rich points in the network.

7.2. Customer Premises Equipment (CPE)

The CPE equipment is the box that a provider places with the customer. It serves two purposes: giving the customer ports to plug in to and making it possible for a provider to monitor the connectivity to the customer site. The CPE is typically a low cost box with limited functionality and, in most cases, is not aware of the VPN services offered by the provider network. The CPE equipment is not necessarily the equipment to which the CE functions are allocated, but it is part of the provider network and is used for monitoring purposes.

The CPE name is used primarily in network operation and deployment contexts and should not be used in protocol specifications.

7.3. Multi-Tenant Unit (MTU)

An MTU is typically an L2 switch placed by a service provider in a building where several customers of that service provider are located. The term was introduced in an Internet Draft specifying a VPLS solution with function distributed between the MTU and the PE in the context of a [VPLS].

The MTU device name is used primarily in network operation and deployment contexts and should not be used in protocol specifications, as it is also an abbreviation used for Maximum Transmit Units.

8. Packet Switched Network (PSN)

A PSN is the network through which the tunnels supporting the VPN services are set up.

8.1. Route Distinguisher (RD)

A Route Distinguisher [RFC2547bis] is an 8-byte value that, together with a 4 byte IPv4 address, identifies a VPN-IPv4 address family. If two VPNs use the same IPv4 address prefix, the PEs translate these into unique VPN-IPv4 address prefixes. This ensures that if the same address is used in two different VPNs, it is possible to install two completely different routes to that address, one for each VPN.

8.2. Route Reflector

A route reflector is a network element owned by a Service Provider (SP) that is used to distribute BGP routes to the SP's BGP-enabled routers [L3VPN-FRAME].

8.3. Route Target (RT)

A Route Target attribute [RFC2547bis] can be thought of as identifying a set of sites or, more precisely, a set of VRFs (see Section 8.9).

Associating a particular Route Target with a route allows that route to be placed in all VRFs used for routing traffic received from the corresponding sites.

A Route Target attribute is also a BGP extended community used in [RFC2547] and [BGP-VPN]. A Route Target community is used to constrain VPN information distribution to the set of VRFs. A route target can be perceived as identifying a set of sites or, more precisely, a set of VRFs.

8.4. Tunnel

A tunnel is connectivity through a PSN that is used to send traffic across the network from one PE to another. The tunnel provides a means to transport packets from one PE to another. Separation of one customer's traffic from another customer's traffic is done based on tunnel multiplexers (see Section 8.5). How the tunnel is established depends on the tunnelling mechanisms provided by the PSN; e.g., the tunnel could be based on the IP-header, an MPLS label, the L2TP Session ID, or the GRE Key field.

8.5. Tunnel Multiplexor

A tunnel multiplexor is an entity that is sent with the packets traversing the tunnel to make it possible to decide which instance of a service a packet belongs to and from which sender it was received. In [PPVPN-L2VPN], the tunnel multiplexor is formatted as an MPLS label.

8.6. Virtual Channel (VC)

A VC is transported within a tunnel and identified by its tunnel multiplexor. A virtual channel is identified by a VCI (Virtual Channel Identifier). In the PPVPN context, a VCI is a VC label or tunnel multiplexor, and in the Martini case, it is equal to the VCID.

8.7. VC Label

In an MPLS-enabled IP network, a VC label is an MPLS label used to identify traffic within a tunnel that belongs to a particular VPN; i.e., the VC label is the tunnel multiplexor in networks that use MPLS labels.

8.8. Inner Label

"Inner label" is another name for VC label (see Section 8.6).

8.9. VPN Routing and Forwarding (VRF)

In networks running 2547 VPN's [RFC2547], PE routers maintain VRFs. A VRF is a per-site forwarding table. Every site to which the PE router is attached is associated with one of these tables. A particular packet's IP destination address is looked up in a particular VRF only if that packet has arrived directly from a site that is associated with that table.

8.10. VPN Forwarding Instance (VFI)

VPN Forwarding Instance (VFI) is a logical entity that resides in a PE that includes the router information base and forwarding information base for a VPN instance [L3VPN-FRAME].

8.11. Virtual Switch Instance (VSI)

In a layer 2 context, a VSI is a virtual switching instance that serves one single VPLS [L2VPN]. A VSI performs standard LAN (i.e., Ethernet) bridging functions. Forwarding done by a VSI is based on MAC addresses and VLAN tags, and possibly on other relevant information on a per VPLS basis. The VSI is allocated to VPLS-PE or, in the distributed case, to the U-PE.

8.12. Virtual Router (VR)

A Virtual Router (VR) is software and hardware based emulation of a physical router. Virtual routers have independent IP routing and forwarding tables, and they are isolated from each other; see [L3VPN-VR].

9. Security Considerations

This is a terminology document and as such doesn't have direct security implications. Security considerations will be specific to solutions, frameworks, and specification documents whose terminology is collected and discussed in this document.

10. Acknowledgements

Much of the content in this document is based on discussion in the PPVPN design teams for "auto discovery" and "l2vpn".

Dave McDysan, Adrian Farrel, and Thomas Narten have carefully reviewed the document and given many useful suggestions.

Thomas Narten converted an almost final version of this document into XML, after extracting an acceptable version from Word became too painful. Avri Doria has been very helpful in guiding us in the use of XML.

11. Informative References

- [L2VPN] Andersson, L. and E. Rosen, "Framework for Layer 2 Virtual Private Networks (L2VPNs)", Work in Progress, June 2004.
- [L2VPN-REQ] Augustyn, W. and Y. Serbest, "Service Requirements for Layer 2 Provider Provisioned Virtual Private Networks", Work in Progress, October 2004.
- [VPLS] Kompella, K., "Virtual Private LAN Service", Work in Progress, January 2005.

- [VPLS-LDP] Lasserre, M. and V. Kompella, "Virtual Private LAN Services over MPLS", Work in Progress, September 2004.
- [BGP-VPN] Ould-Brahim, H., Rosen, E., and Y. Rekhter, "Using BGP as an Auto-Discovery Mechanism for Layer-3 and Layer-2 VPNs", Work in Progress, May 2004.
- [L3VPN-FRAME] Callon, R. and M. Suzuki, "A Framework for Layer 3 Provider Provisioned Virtual Private Networks", Work in Progress, July 2003.
- [RFC3809] Nagarajan, A., "Generic Requirements for Provider Provisioned Virtual Private Networks (PPVPN)", RFC 3809, June 2004.
- [L3VPN-REQ] Carugi, M. and D. McDysan, "Service requirements for Layer 3 Virtual Private Networks", Work in Progress, July 2004.
- [RFC2547bis] Rosen, E., "BGP/MPLS IP VPNs", Work in Progress, October 2004.
- [L3VPN-VR] Knight, P., Ould-Brahim, H. and B. Gleeson, "Network based IP VPN Architecture using Virtual Routers", Work in Progress, April 2004.
- [PWE3-ARCH] Bryant, S. and P. Pate, "PWE3 Architecture", Work in Progress, March 2004.
- [RFC3916] Xiao, X., McPherson, D., and P. Pate, "Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)", RFC 3916, September 2004.
- [PPVPN-L2VPN] Kompella, K., "Layer 2 VPNs Over Tunnels", Work in Progress, June 2002.
- [ENCAP-MPLS] Martini, L., "Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks", Work in Progress, September 2004.
- [TRANS-MPLS] Martini, L. and N. El-Aawar, "Transport of Layer 2 Frames Over MPLS", Work in Progress, June 2004.
- [RFC2547] Rosen, E. and Y. Rekhter, "BGP/MPLS VPNs", RFC 2547, March 1999.

[RFC2764] Gleeson, B., Lin, A., Heinanen, J., Armitage, G., and
A. Malis, "A Framework for IP Based Virtual Private
Networks", RFC 2764, February 2000.

Authors' Addresses

Loa Anderson
Acreo AB

EMail: loa@pi.se

Tove Madsen
Acreo AB

EMail: tove.madsen@acreo.se

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

