L3VPN Working Group
Internet-Draft

Loa Andersson
Tove Madsen
TLA¡group

Expiration Date: March 2004

25 September, 2003

PPVPN terminology
<draft-andersson-ppvpn-terminology-04.txt>

Status of this Memo

Abstract

The provider provisioned VPN solutions has attracted a great deal of interest. Memos proposing different and overlapping solution have been discussed on the PPVPN mailing list and in the Working Group meetings. This has lead to a development of a partly new set of concepts used to describe the set of VPN services. To a certain extent there are more than one term covering the same concept and sometimes the same term covers more than on concept. The terminology needs to be made clearer and more intuitive. This document seeks to fill at least part of that need.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 [RFC2119].

Table of contents

1. Introduction

There are a comparatively large number of memos being submitted to the former PPVPN, and L2VPN, L3VPN and PWE3 working groups that all addresses 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 lead to a development of a partly new set of concepts used to describe this set of VPN services. To a certain extent there are more than one term covering the same concept and sometimes the same term covers more than one concept. The terminology needs to be made clearer and more intuitive.

This document seeks to fill at least part of the need and proposes a foundation for a unified terminology for the L2VPN, L3VPN working groups; in some cases the parallel concepts within the PWE3 working group is 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 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. there are concepts and terms within the PWE3 and (Generalized) MPLS areas that are closely related. We've tried to find the earliest use of terms and concepts.

This document intends to fully cover the concepts within five core documents from the L2VPN and L3VPN working groups the "Generic Requirements for Provider Provisioned VPN" [GENERIC], the "A Framework for Layer 3 Provider Provisioned Virtual Private Networks" [L3VPN-frmwrk], the "Service requirements for Layer 3 Provider Provisioned Virtual Private Networks" [PPVPN-req], the "L2VPN Framework [L2VPN-frmwrk] and "Service Requirements for Layer 2 Provider Provisioned Virtual Private Networks" [L2VPN-req]. The intention is to create a comprehensive and unified set of concepts for these documents, and by extension for the entire PPVPNarea.Todosoitisalsonecessarytogivesomeofthe development the concepts of the area have been through.

The document is structured in four major sections. Section 4 lists the different services that has been/will be specified, Section 5 lists the building blocks that is used to specify those services, section 6 lists the functions needed in those services and section 7 list 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 concept "pseudo wire" that belongs to the PWE3 working group is included for reference purposes. For requirements in provider provisioned VPNs see [PPVPN-req].

In this section all abbreviations are listed in alphabetic order.

3.1 IP-only LAN-like Service (IPLS)

An IPLS is very like a VPLS (see 3.8), except that:

- it is assumed that the CE devices (see 5.1) are hosts or routers, not switches

- it is assumed that the service will only need to carry IP packets, and supporting packets such as ICMP and ARP; otherwise layer 2 packets which do not contain IP are not supported.

While this service is a functional subset of the VPLS service, it is considered separately because it may be possible to provide it using different mechanisms, which may allow it to run on certain hardware platforms that cannot support the full VPLS functionality [PPVPN-L2-frmwrk].

3.2 Layer 2 VPN (L2VPN)

Three types of L2VPNs are described in this document, Virtual Private Wire Service (VPWS) (section 3.11), VPLS Virtual Private LAN Service (VPLS)(section 3.8), and IP-only LAN-like Service (IPLS).

3.3 Layer 3 VPN (L3VPN)

An L3VPN is a solution that interconnects several sets of hosts and routers and allows them to communicate based on L3 addresses, see [L3VPN-frmwrk].

3.4 Pseudo Wire (PW)

The PWE3 working group within IETF specifies the pseudo wire technology. A pseudo wire is an emulated point-to-point connectivity over a packet switched network that gives the possibility to interconnect 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 5.2) and CE (see 5.1). An early solution for PWs is described in [martini-tran]. Encapsulation formats readily used in VPWS, VPLS and PWs are described in [martini-encap]. Requirements for PWs are found in [PWE3-req] and [PWE3-frmwrk] present an architectural framework for PWs.

3.5 Transparent LAN Service (TLS)

TLS was an early name used to describe the VPLS service, it was used e.g. in the now dated draft-lasserre-tls-mpls-00.txt. It has been replaced by VPLS, which is the current term.

3.6 Virtual LAN (VLAN)

A VLAN is a way of separating traffic on a LAN, e.g. between different departments within a company. This acronym is not defined by former PPVPN working group, but is defined by IEEE 802.1Q. The VLANID is used to mark an Ethernet frame with a tag to create user groups on a LAN.

3.7 Virtual Leased Line Service (VLLS)

The VLLS has been replaced by VPWS. It was used in now dated draft-ppvpn-metrics.00.txt.

3.8 Virtual Private LAN Service (VPLS)

A VPLS is a provider service that emulates the full functionality of a traditional Local Area Network. 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], [Lasserre-vkompella], and [Kompella-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.

3.9 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 may communicate among them as if they were on a private network. The level of separation is possible to enhance e.g. by end-to-end encryption, this is however outside the scope

of IETF VPN working group charters. This VPN definition is from [RFC2764].

In the [L3VPN-frmwrk] 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 VPNs.

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

3.10 Virtual Private Switched Network (VPSN)

A VPSN is replaced by VPLS. The VPSN abbreviation was used e.g. in the now dated draft-vkompella-ppvpn-vpsn.reqmts-00.txt.

3.11 Virtual Private Wire Service (VPWS)

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

The PE's in the core network is 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 [L2VPN].

A VPWS differs from a VPLS (section 4.8) in that the VPLS is point to multipoint, while the VPWS is point to point. See [L2VPN-frmwrk].

4. Classification of VPNs

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

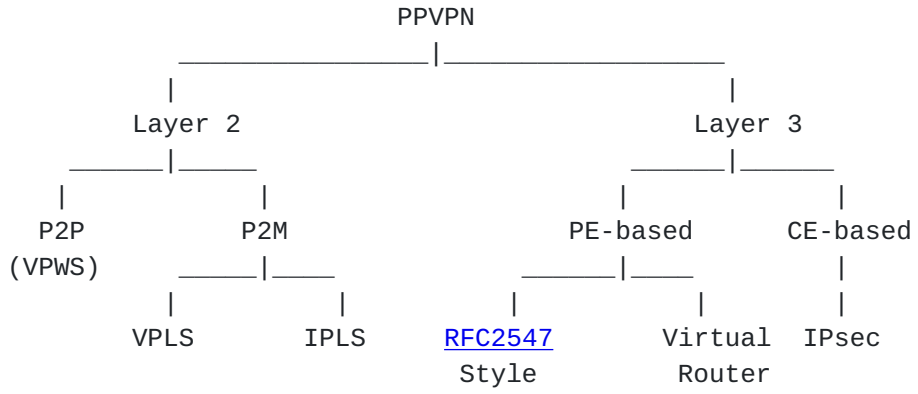

The figure above presents a taxonomy of PPVPN technologies. Some of the definitions are given below:

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-frmwrk]).

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 VPN. Because the PE device maintains all required VPN state, 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 on 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-frmwrk].

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 [PPVPN-VR].

RFC 2547 Style: A PE-based VPN approach in which the PE router maintains separate forwarding environment for each VPN and a separate forwarding table for each VPN. In order to maintain multiple forwarding table instances while running only a single routing protocol instance, RFC 2547 style VPNs mark route advertisements with attributes that identify their VPN context. These VPNs are based on the approach described in [RFC2547bis].

5. Building blocks

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

However, for different reasons this is to over-simplify. 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.

There are two different aspects that need 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 and with the result it 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, since 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 want to limit the number on MAC-addresses that needs to be learnt 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 is just examples and it will be extended as the number of services increases.

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 the customer. PE, without further qualifications, 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 which type of equipment that 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)

APE-SisaL2devicethatparticipatesine.g.aswitched
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 it is in the VPLS/VPWS cases sometimes desired to distribute the functions in the VPLS/VPWS-PE across more than one device, e.g. is it feasible to allocate MAC address learning on a comparatively small and in-expensive device close to the customer site, while participation in the PSN signalling and set up 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 decision at the ingress of the provider network.

5.3 Core

5.3.1 Provider router (P)

ThePisdefinedasarouterinthecorenetworkthatdoesnot have interfaces directly towards a customer. Hence a P router does not need to keep VPN state and is VPN un-aware.

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(s) or uplinks. The LPE will appear to the core network as a single PE. In this document the devices that constitutes the LPE is called N-PE and U-PE.

5.4.3 PE-CLE

An alternative name for the U-PE suggested in now dated 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 now dated 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 deviceusedbyaprovidernetworktohandoffaVPLStoa customer. In this document the VE is called a VPLS-PE.

This name has dated.

6. Functions

In this section we have grouped a number of concepts and terms that has 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 the SP; may be used to interconnect CE devices in multiple-homing arrangements [L3VPN-frmwrk].

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 [PPVPN-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 and L3 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 learnt 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 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 is part of the provider network and 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 [DTLS] is typically an L2 switch placed by a service provider in a building where customers of that service provider are located.

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 a used abbreviation 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 togetherwitha4byteIPv4addressidentifiesaVPN-IPv4address family. If two VPNs use the same IPv4 address prefix, the PEs translates 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-frmwrk].

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 8.8).

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

A Route Target attribute is also a BGP extended community used in [RFC2547], and [BGPVPN-auto]. 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 mechanism 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 8.4). How the tunnel is established depends on the tunnelling mechanisms provided by the PSN, i.e. the tunnel could be based on e.g. the IP-header, an MPLS label, the L2TP Session ID, or on 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 possible to decide to which instance of a service a packet belongs and from which sender it was received. In [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 multiplexer. A virtual channel is identified by a VCI (Virtual Channel Identifier). In the PPVPN context a VCI is a VC label or tunnel multiplexer 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 multiplexer in networks that uses 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 VRF's. 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, which 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-frmwrk].

8.11 Virtual Switch Instance (VSI)

In a layer 2 context a VSI is a virtual switching instance that serves one single VPLS [L2VPN-frmwrk]. 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 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 [PPVPN-VR].

9. Acknowledgements

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

10. Authors' Contact

Loa Andersson
TLA-group
loa@pi.se

Tove Madsen
TLA-group
tove@niebelungen.net

11. Normative References

[GENERIC] Nagarajan, A (ed), "Generic Requirements for Provider Provisioned VPN", draft-ietf-l3vpn-generic-reqts-01.txt, Work in Progress, Internet Draft, Aug 2003

[L2VPN-frmwrk] Andersson, L. and Rosen, E., "L2VPN Framework", draft-ietf-l2vpn-l2-framework-01.txt, Work in Progress, Internet Draft, Sept 2003

[L3VPN-frmwrk] Callon, R. and Suzuki, M.," A Framework for Layer 3 Provider Provisioned Virtual Private Networks", draft-ietf-l3vpn-framework-00.txt, Work in Progress, Internet Draft, March 2003

[PPVPN-req] Carugi, M. and McDysan, D., "Service requirements for Layer 3 Provider Provisioned Virtual Private Networks", draft-ietf-l3vpn-requirements-00.txt, Work in Progress, Internet Draft, Oct 2003

[L2VPN-req] Augustyn, W., and Serbest, Y., "Service Requirements for Layer 2 Provider Provisioned Virtual Private Networks", draft-ietf-l2vpn-requirements-00.txt, Work in Progress, Internet Draft, May 2003

12. Non-Normative References

[BGPVPN-auto] Ould-Brahim, H., Rosen, E. and Rekhter, Y. "Using BGP as an auto-discovery mechanism for network-based VPNs", draft-ietf-l3vpn-bgpvpn-auto-00.txt, Work in progress, Internet Draft, July 2003

[kompella-VPLS] Kompella, K. "Virtual Private LAN Service", draft-ietf-l2vpn-vpls-bgp-00.txt, Work in Progress, Internet Draft, May 2003

[L2VPN] Kompella, K., et.al. "Layer 2 VPNs Over Tunnels", draft-kompella-ppvpn-l2vpn-03.txt, Work in Progress, June 2002

[lasserre-vkompella] Kompella, V. and Lasserre, M., "Virtual Private LAN Services over MPLS" draft-ietf-l2vpn-vpls-ldp-00.txt, Work in progress, Mar 2002

[martini-encap] Martini, L., et.al. "Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLSNetworks", draft-martini-l2circuit-encap-mpls-05.txt, Work in Progress, Internet Draft, April 2003

[martini-tran] Martini, L., et.al. "Transport of Layer 2 Frames Over MPLS", draft-martini-l2circuit-trans-mpls-11.txt, Work in progress, Internet Draft, April 2003

[PPVPN-VR] Ould-Brahim, H., et.al. "Network-based IP VPN using Virtual Routers", draft-ietf-l3vpn-vpn-vr-00.txt, Work in Progress, Internet Draft, July 2002

[PWE3-arch] Prayson, P. and Bryant, S., "PWE3 Architecture", draft-ietf-pwe3-arch-05.txt, Work in Progress, Internet Draft, August 2003

[PWE3-req] Xiao, X., "Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)", draft-ietf-pwe3-requirements-06.txt, Work in progress, Internet Draft, July 2003

[RFC2547] Rosen, E., et.al. "BGP/MPLS VPNs", rfc2547, March 1999

[RFC2547bis] Rosen, E., "BGP/MPLS IP VPNs", draft-ietf-l3vpn-rfc2547bis-01.txt, Work in Progress, Internet Draft, September 2003

[RFC2764] Gleeson, B., et.al. "A Framework for IP Based Virtual Private Networks", rfc2764, February 2000