Network Working Group   L. Andersson
Internet-Draft   T. Madsen
Expires: March 11, 2005   Acreo AB
                                                     September 10, 2004

Provider Provisioned VPN terminology
draft-ietf-l3vpn-ppvpn-terminology-04.txt

Status of this Memo

Copyright Notice

Abstract

The wide spread interest in provider-provisioned 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 been discussed these proposals and
documented specifications. This has lead to the development of a
partly new set of concepts used to describe the set of VPN services.
To a certain extent there is more than one term covering the same

concept and sometimes the same term covers more than one concept.
This document seeks to fill the need to make the terminology in the
area clearer and more intuitive.

Table of Contents

[1](#). **Introduction**

A comparatively large number of memos has been submitted to the
former PPVPN working group, and to the L2VPN, L3VPN and PWE3 working
groups that 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 lead to the development of a partly new set of concepts used
to describe this set of VPN services. To a certain extent there is
more than one term covering the same concept and sometimes the same
term covers more than one concept.

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 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 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 the core
   documents from the L2VPN and L3VPN working groups, i.e.
   [I-D.ietf-l3vpn-requirements] , [I-D.ietf-l2vpn-requirements],
   [I-D.ietf-l3vpn-framework], [I-D.ietf-l2vpn-l2-framework], and
   [I-D.ietf-l3vpn-generic-reqts]. 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 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 [I-D.ietf-l3vpn-requirements].

   In this section all abbreviations are listed together with short
   information on tyhe service. The list is structured to give the more
   general information first and more specific later. The names of
   services for which IETF are working on solution for have as far as
   feasible been moved to the top of the list. Older and dated
   terminology has been pushed towrd the en of the list.

3.1 Layer 3 VPN (L3VPN)

   An L3VPN interconnects sets of hosts and routers based on Layer 3
   addresses, see [I-D.ietf-l3vpn-framework].

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. 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
   [I-D.ietf-l2vpn-requirements], [I-D.ietf-l2vpn-vpls-ldp], and
   [I-D.ietf-l2vpn-vpls-bgp].

   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 PE's 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
[I-D.kompella-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
[I-D.ietf-l2vpn-l2-framework].

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 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.
o   the assumption that only IP packets is 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 using
different mechanisms, which may allow it to run on certain hardware
platforms that cannot support the full VPLS functionality
[I-D.ietf-l2vpn-l2-framework].

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 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 Section 5.2) and CE (see
Section 5.1). An early solution for PWs is described in

[I-D.martini-l2circuit-trans-mpls]. Encapsulation formats readily
used in VPWS, VPLS and PWs are described in
[I-D.martini-l2circuit-encap-mpls]. Requirements for PWs are found
in [I-D.ietf-pwe3-requirements] and [I-D.ietf-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, it was used e.g., in the now dated draft-lasserre-tls-mpls-00.txt. This draft has been merged into the working group VPLS solution using LDP as a signalling protocol. The term 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 to differentiate 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 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 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 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, this is however outside the scope of IETF VPN working group charters. This VPN definition is from [RFC2764].

In the [I-D.ietf-l3vpn-framework] 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.
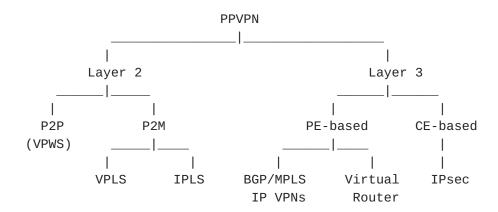
3.11 Virtual Private Switched Network (VPSN)

The term VPSN has been replaced by the term VPLS. The VPSN
abbreviation was used e.g. in the now dated draft-vkompella-ppvpn-
vpsn-reqmts-00.txt. The requirements has been merged into the L3VPN
[I-D.ietf-l3vpn-requirements] and L2VPN [I-D.ietf-l2vpn-requirements]
requirements.

4. Classification of VPNs

   The terminology used in [I-D.ietf-l3vpn-generic-reqts] is defined
   based on the figure below.


```
                            PPVPN
              _____|_____
              |                                   |
           Layer 2                             Layer 3
         _____|_____                        _____|_____
         |          |                        |            |
        P2P        P2M                     PE-based    CE-based
       (VPWS)    _____|____               _____|____      |
                 |        |               |         |      |
                VPLS     IPLS          BGP/MPLS   Virtual  IPsec
                                        IP VPNs    Router
```


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

                             Figure 1

   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 [I-D.ietf-l3vpn-framework]).

   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 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) [I-D.ietf-l3vpn-framework].

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 [I-D.ietf-l3vpn-vpn-vr].

BGP/MPLS IP VPNs: 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 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 [I-D.ietf-l3vpn-rfc2547bis].

RFC 2547 Style: The term has been used by the L3VPN to describe the
extensions of the VPNs defined in the informational 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 [I-D.ietf-l3vpn-rfc2547bis] and Virtual Routers
[I-D.ietf-l3vpn-vpn-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 it were physical
boxes, common for all services.

However, for different reasons this is an over-simplification. 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
[I-D.ietf-l3vpn-framework]. 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 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
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 CE's. There are many expamples of this type of
naming and we have not had the ambition to cover every possible
example, but only cover the most frequently used functional names.
As this is functional names it is quite possible that on single piece
of equipment are the platform for more than one type of function,
e.g. a router might at the same time be both a L2VPN-CE and a
L3VPN-CE. It might also be that the functions that is need to for a
L2VPN-CE or L3VPN-CE is 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
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 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 e.g. 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 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)

The P is defined as a router in the core network that does not 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 the now 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 now 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 has 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 the SP; may be used to interconnect CE
   devices in multiple-homing arrangements [I-D.ietf-l3vpn-framework].

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 [I-D.ietf-l3vpn-requirements]. 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 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 is typically an L2 switch placed by a service provider in a
   building where several customers of that service provider are
   located. The term wwere introduced in an Internet Draft specifying a
   VPLS solution with function distributed between the MTU and the PE in
   the contect of a [I-D.ietf-l2vpn-vpls-bgp].

   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 [I-D.ietf-l3vpn-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 [I-D.ietf-l3vpn-framework].

8.3 Route Target (RT)

A Route Target attribute [I-D.ietf-l3vpn-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 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 [I-D.ietf-l3vpn-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 8.5). 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 it possible to decide to which instance
   of a service a packet belongs and from which sender it was received.
   In [I-D.kompella-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
   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 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 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 [I-D.ietf-l3vpn-framework].

8.11 Virtual Switch Instance (VSI)

In a layer 2 context a VSI is a virtual switching instance that
serves one single VPLS [I-D.ietf-l2vpn-l2-framework]. 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
[I-D.ietf-l3vpn-vpn-vr].

9. **Security Considerations**

   This is a terminology document and as such doesn't have direct
   security implications. Security considerations will be specific to
   the solutions, framework 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 out of Word became to
   painful. Avri Doria has been very helpful in guiding is in the use
   of XML.

11 **Informative References**

   [I-D.ietf-l2vpn-l2-framework]
             Andersson, L. and E. Rosen, "Framework for Layer 2 Virtual
             Private Networks (L2VPNs)",
             draft-ietf-l2vpn-l2-framework-05 (work in progress), June
             2004.

   [I-D.ietf-l2vpn-requirements]
             Augustyn, W. and Y. Serbest, "Service Requirements for
             Layer 2 Provider Provisioned Virtual Private Networks",
             draft-ietf-l2vpn-requirements-01 (work in progress),
             February 2004.

   [I-D.ietf-l2vpn-vpls-bgp]
             Kompella, K., "Virtual Private LAN Service",
             draft-ietf-l2vpn-vpls-bgp-02 (work in progress), May 2004.

   [I-D.ietf-l2vpn-vpls-ldp]
             Lasserre, M. and V. Kompella, "Virtual Private LAN
             Services over MPLS", draft-ietf-l2vpn-vpls-ldp-04 (work in
             progress), August 2004.

   [I-D.ietf-l3vpn-bgpvpn-auto]
             Ould-Brahim, H., Rosen, E. and Y. Rekhter, "Using BGP as
             an Auto-Discovery Mechanism for Layer-3 and Layer-2 VPNs",
             draft-ietf-l3vpn-bgpvpn-auto-04 (work in progress), May

2004.

     [I-D.ietf-l3vpn-framework]
               Callon, R. and M. Suzuki, "A Framework for Layer 3
               Provider Provisioned Virtual Private Networks",
               draft-ietf-l3vpn-framework-00 (work in progress), July
               2003.

[I-D.ietf-l3vpn-generic-reqts]
        Nagarajan, A., "Generic Requirements for Provider
        Provisioned Virtual Private Networks",
        draft-ietf-l3vpn-generic-reqts-03 (work in progress),
        February 2004.

[I-D.ietf-l3vpn-requirements]
        Carugi, M. and D. McDysan, "Service requirements for Layer
        3 Virtual Private Networks",
        draft-ietf-l3vpn-requirements-02 (work in progress), July
        2004.

[I-D.ietf-l3vpn-rfc2547bis]
        Rosen, E., "BGP/MPLS IP VPNs",
        draft-ietf-l3vpn-rfc2547bis-01 (work in progress),
        September 2003.

[I-D.ietf-l3vpn-vpn-vr]
        Knight, P., Ould-Brahim, H. and B. Gleeson, "Network based
        IP VPN Architecture using Virtual Routers",
        draft-ietf-l3vpn-vpn-vr-02 (work in progress), April 2004.

[I-D.ietf-pwe3-arch]
        Bryant, S. and P. Pate, "PWE3 Architecture",
        draft-ietf-pwe3-arch-07 (work in progress), March 2004.

[I-D.ietf-pwe3-requirements]
        Xiao, X., "Requirements for Pseudo-Wire Emulation
        Edge-to-Edge (PWE3)", draft-ietf-pwe3-requirements-08
        (work in progress), January 2004.

[I-D.kompella-ppvpn-l2vpn]
        Kompella, K., "Layer 2 VPNs Over Tunnels",
        draft-kompella-ppvpn-l2vpn-02 (work in progress), June
        2002.

[I-D.martini-l2circuit-encap-mpls]
        Martini, L., "Encapsulation Methods for Transport of Layer
        2 Frames Over IP and MPLS Networks",
        draft-martini-l2circuit-encap-mpls-07 (work in progress),
        June 2004.

   [I-D.martini-l2circuit-trans-mpls]
            Martini, L. and N. El-Aawar, "Transport of Layer 2 Frames
            Over MPLS", draft-martini-l2circuit-trans-mpls-14 (work in
            progress), June 2004.

   [RFC2547] Rosen, E. and Y. Rekhter, "BGP/MPLS VPNs", RFC 2547, March

                1999.

   [RFC2764] Gleeson, B., Heinanen, J., Lin, A., 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

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment