INTERNET DRAFT
Internet Engineering Task Force
Document:
draft-augustyn-ppvpn-l2vpn-requirements-02.txt
February 2003
Category: Informational
Expires: August 2003

W. Augustyn

Y. Serbest
SBC

(Co-Editors)

Service Requirements for Layer 2 Provider Provisioned Virtual Private Networks

Status of this memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC 2026 ([RFC-2026]).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This document is a product of the IETF's Provider Provisioned Virtual Private Network (ppvpn) working group. Comments SHOULD be addressed to WG's mailing list at ppvpn@ppvpn.francetelecom.com. The charter for ppvpn may be found at http://www.ietf.org/html.charters/ppvpn-charter.html

Abstract

This document provides requirements for Layer 2 Provider Provisioned Virtual Private Networks (PPVPNs). It first provides taxonomy and terminology and states generic and general service requirements. It covers point to point VPNs referred to as Virtual Private Wire Service (VPWS), as well as multipoint to multipoint VPNs also known as Virtual Private LAN Service (VPLS). Detailed requirements are

expressed from a customer as well as a service provider perspective.

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 ([RFC-2119]).

Table of Contents

1 Contributing Authors

This document was the combined effort of several individuals. The following are the authors that contributed to this document:

Waldemar Augustyn
Marco Carugi
Giles Heron
Vach Kompella
Marc Lasserre
Pascal Menezes
Hamid Ould-Brahim
Tissa Senevirathne
Yetik Serbest

2 Introduction

This section describes the scope and outline of the document.

2.1 Scope of this document

This document provides requirements for Layer 2 Provider Provisioned Virtual Private Networks (L2-PPVPN). It identifies requirements that MAY apply to one or more individual approaches that a Service Provider MAY use for the provisioning of a Layer 2 VPN service. The content of this document makes use of the terminology defined in [TERMINOLOGY] and common components for deploying Layer 2 PPVPNs (will be referred to as L2VPNs) described in [PPVPN-L2-FR]. The technical specifications to provide L2VPN services are outside the scope of this document. The framework document [PPVPN-L2-FR] and several documents, which explain technical approaches providing Layer 2 PPVPN services, are available to cover this aspect.

This document describes requirements for two types of L2VPNs: 1. Virtual Private Wire Service (VPWS), and 2. Virtual Private LAN Service (VPLS). The approach followed in this document distinguishes L2VPN types as to how the connectivity is provided (point-point or

multipoint-multipoint) as detailed in [PPVPN-L2-FR].

This document is intended as a "checklist" of requirements that will provide a consistent way to evaluate and document how well each individual approach satisfies specific requirements. The applicability statement documents for each individual approach SHOULD document the results of this evaluation.

In the context of provider provisioned VPNs, there are two entities involved in operation of such services, the Provider and the Customer. The Provider engages in a binding agreement with the Customer as to the behavior of the service in normal situation as well as exceptional situations. Such agreement is known as Service Level Specification (SLS) which is part of the Service Level Agreement (SLA) established between the Provider and the Customer.

A proper design of L2VPNs aids formulation of SLSs in that it provides means for proper separation between CE/PE, allows proper execution of the SLS offer, and supports flexible and rich set of capabilities.

This document provides requirements from both the ProviderÆs and the CustomerÆs point of view. It begins with common customerÆs and service providerÆs point of view, followed by a customerÆs perspective, and concludes with specific needs of a Service Provider (SP). These requirements provide high-level L2VPN features expected by a SP in provisioning L2VPNs, which include SP requirements for security, privacy, manageability, interoperability and scalability, in addition to service provider projections for number, complexity, and rate of change of customer VPNs over the next several years.

2.2 Outline

The outline of the rest of this document is as follows. Section 3 defines terminology. Section 4 provides common requirements that apply to both customer and service providers respectively. Section 5 states requirements from a customer perspective. Section 6 states network requirements from a service provider perspective. Section 7 states service provider management requirements. Section 8 describes the engineering requirements, particularly control and data plane requirements. Section 9 provides security considerations. Section 10 lists acknowledgements. Section 11 provides a list of references cited herein. Section 12 lists the editorsÆ addresses.

3 Definitions and Taxonomy

3.1 Definitions

The terminology used in this document is defined in [TERMINOLOGY].

The Layer 2 PPVPN framework document [PPVPN-L2-FR] further describes these concepts in the context of a reference model that defines layered service relationships between devices and one or more levels of tunnels.

3.2 Taxonomy of Layer 2 PPVPN Types

The requirements distinguish two major L2VPNs models, a Virtual Private Wire Service (VPWS), and a Virtual Private LAN Service (VPLS).

The following diagram shows a L2VPN reference model.

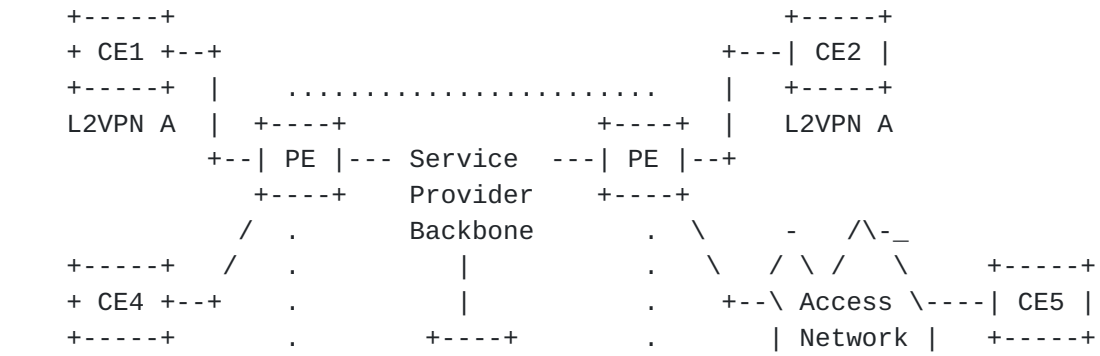

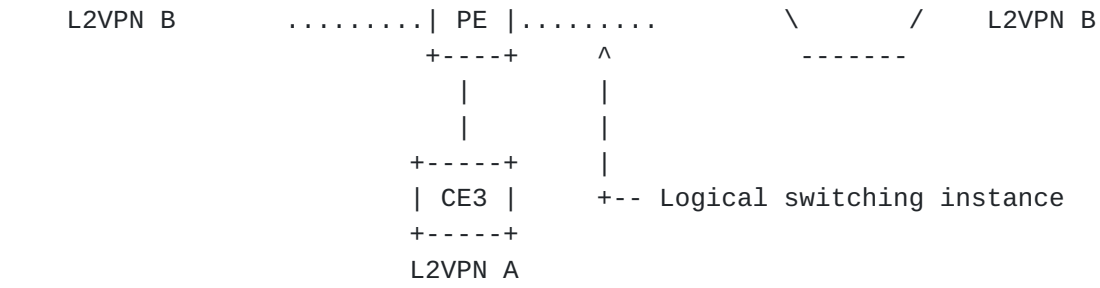

Figure 1 L2VPN Reference Model

3.3 VPWS

The PE devices provide a logical interconnect such that a pair of CE devices appear to be connected by a single logical Layer 2 circuit. PE devices act as Layer 2 circuit switches. Layer 2 circuits are then mapped onto tunnels in the service provider network. These tunnels can either be specific to a particular VPWS, or shared among several VPWSs. VPWS applies for all services including Ethernet, ATM, Frame Relay etc. In Figure 1, L2VPN B represents a VPWS case.

Each PE device is responsible for allocating customer Layer 2 frames to the appropriate VPWS and for proper forwarding to the intended destinations.

3.4 VPLS

In case of VPLS, the PE devices provide a logical interconnect such

that CE devices belonging to a specific VPLS appear to be connected by a single LAN. End-to-end VPLS consists of a bridge module and a LAN emulation module ([PPVPN-L2-FR]). A VPLS can contain a single VLAN or multiple VLANs. A variation of this service is IPLS, which is limited to supporting only customer IP traffic.

In a VPLS, a customer site receives layer 2 service from the SP. The PE is attached via an access connection to one or more CEs. The PE performs forwarding of user data packets based on information in the Layer 2 header, such as a MAC destination address. The CE sees a bridge.

In VPLS, the PE can be viewed as containing a Virtual Switching Instance (VSI) for each Layer 2 VPN that it serves. A CE device attaches, possibly through an access network, to a "Bridge" module of a PE. Within the PE, the Bridge module attaches, through an "Emulated LAN Interface" to an Emulated LAN. For each VPLS, there is an Emulated LAN instance. In Figure 1, the top PE routers maintain separate Emulated LAN instances for VPLS A and VPLS B. The Emulated LAN consists of "VPLS Forwarder" module (one per PE per VPLS instance) connected by pseudowires, where the pseudowires may be traveling through Packet Switched Network (PSN) tunnels over a routed backbone. VSI is a logical entity that contains a VPLS forwarder module and part of the bridge module relevant to the VPLS instance [PPVPN-L2-FR]. Hence, the VSI terminates pseudo-wires for interconnection with other VSIs and also terminates attachment circuits for accommodating CEs. A VSI includes the forwarding

information base for a L2VPN [PPVPN-L2-FR] which is the set of information regarding how to forward Layer 2 frames received over the attachment circuit from the CE to VSIs in other PEs supporting the same L2VPN (and/or to other attachment circuits), and contains information regarding how to forward Layer 2 frames received from pseudowires to attachment circuits. Forwarding information bases can be populated dynamically (such as by source MAC address learning) or statically (e.g., by configuration). Each PE device is responsible for proper forwarding of the customer traffic to the appropriate destination(s) based on the forwarding information base of the corresponding VSI.

4 Service Requirements Common to Customers and Service Providers

This section contains requirements that apply to both the customer and the provider, or are of an otherwise general nature.

4.1 Scope of emulation

L2VPN protocols SHOULD NOT interfere with existing Layer 2 protocols and standards of the Layer 2 network the customer is managing. If

they may impact customer Layer 2 protocols that are sent over the VPLS, then these impacts MUST be documented.

Some possibly salient differences between VPLS and a real LAN are:

- The reliability MAY likely be less, i.e., the probability that a message broadcast over the VPLS is not seen by one of the bridge modules in PEs is higher than in a true Ethernet.
- If sequencing is not turned on, BPDUs on a pseudowire may get out of order with respect to data packets and with respect to each other.
- VPLS frames can get duplicated if the sequencing option isn't turned on. The data frames on the pseudowires are sent in IP datagrams, and under certain failure scenarios, IP networks can duplicate packets. If the pseudowire data transmission protocol does not ensure sequence of data packets, frames can be duplicated or received out of sequence. If the customerÆs BPDU frames are sent as data packets, then BPDU frames can be duplicated or mis-sequenced.
- Delayed delivery of packets (e.g., more than half a second) rather than dropping them could have adverse effect on the performance of the service.
- 802.3x Pause frames will not be transported over a VPLS, as the bridge module ([PPVPN-L2-FR]) in the PE terminates them.
- Since the IPLS solution aims at transporting encapsulated traffic (rather than Layer 2 frames themselves), the IPLS solution is NOT REQUIRED to preserve the Layer 2 Header transparently from CE to CE. For example, Source MAC address may not be preserved by the IPLS solution.

The interaction between L2VPN and the customer equipment SHOULD comply with existing native protocols and specifications. In case, a L2VPN solution supports only a subset of these specifications, the exceptions MUST be documented.

4.2 Traffic Types

A VPLS MUST support unicast, multicast, and broadcast traffic. It is desirable to support efficient replication of broadcast and multicast traffic.

4.3 Topology

A SP network MAY be realized using one or more network tunnel topologies to interconnect PEs, ranging from simple point-to-point to distributed hierarchical arrangements. The typical topologies include:

- Point-to-point

- o Point-to-multipoint, a.k.a. hub and spoke
- o Any-to-any, a.k.a. full mesh
- o Mixed, a.k.a. partial mesh
- o Hierarchical

Regardless of the SP topology employed, the service to the customers MUST retain the connectivity type implied by the type of L2VPN. For example, a VPLS SHOULD allow multipoint to multipoint connectivity even if implemented with point to point circuits. This requirement does not imply that all traffic characteristics (such as bandwidth, QoS, delay, etc.) be necessarily the same between any two end points of a L2VPN. It is important to note that SLS requirements of a service have a bearing on the type of topology that can be used.

To the extent possible, a L2VPN service SHOULD be capable of crossing multiple administrative boundaries.

To the extent possible, the L2VPN services SHOULD be independent of access network technology.

4.4 Isolated Exchange of Data and Forwarding Information

L2VPN solutions SHALL define means that prevent CEs in a L2VPN from interaction with unauthorized entities.

L2VPN solutions SHALL avoid introducing undesired forwarding information that could corrupt the L2VPN forwarding information base.

A means to constrain, or isolate, the distribution of addressed data to only those VPLS sites determined either by MAC learning and/or configuration MUST be provided.

The internal structure of a L2VPN SHOULD not be advertised nor discoverable from outside that L2VPN.

4.5 Security

A range of security features SHOULD be supported by the suite of L2VPN solutions. Each L2VPN solution SHOULD state which security

features it supports and how such features can be configured on a per customer basis.

A number of security concerns arise in the setup and operation of a L2VPN, ranging from mis-configurations to attacks that may be launched on a L2VPN. This section lists some potential security hazards. There MUST be methods available to protect against the following situations.

- Protocol attacks
  - o Excessive protocol adjacency setup/teardown
  - o Excessive protocol signaling/withdrawal
- Resource Utilization
  - o Forwarding plane replication (VPLS)
  - o Looping (VPLS primarily)
  - o MAC learning table size limit (VPLS)
- Unauthorized access
  - o Unauthorized member of VPN
  - o Incorrect customer interface
  - o Incorrect service delimiting VLAN tag
  - o Unauthorized access to PE
- Tampering with signaling
  - o Incorrect FEC signaling
  - o Incorrect label assignment
  - o Incorrect signaled VPN parameters (e.g., QoS, MTU, etc.)
- Tampering with data forwarding
  - o Incorrect MAC learning entry
  - o Incorrect label
  - o Incorrect customer facing encapsulation
  - o Incorrect pseudo-wire encapsulation
  - o Hijacking pseudowires using the wrong tunnel
  - o Incorrect tunnel encapsulation

4.5.1 User data security

L2VPN solution MUST provide traffic separation between different L2VPNs.

In case of VPLS, VLAN Ids MAY be used as service delimiters. When used in this manner, they MUST be honored and traffic separation MUST be provided.

4.5.2 Access control

A L2VPN solution MAY also have the ability to activate the appropriate filtering capabilities upon request of a customer.

4.6 Addressing

A L2VPN solution MUST support overlapping addresses of different L2VPNs. For instance, customers SHOULD not be prevented from using

the same MAC addresses and/or the same VLAN Ids when used with different L2VPNs. Actually, for VLANs, there are two cases. First, a L2VPN is oblivious to customer VLANs. In this case, customers can have overlapping VLAN Ids. Second, VLAN Ids MAY be used as service delimiters, in which case it depends on whether the SP assigns the VLANs or not. If it does, then there is no overlapping. If it

doesn't, then overlapping VLAN Ids can occur and the SP has to put safeguards in place to avoid this situation.

4.7 Quality of Service

To the extent possible, L2VPN QoS SHOULD be independent of the access network technology.

4.7.1 QoS Standards

As provided in [PPVPN-REQTS] a L2PVN SHALL be able to support QoS in one or more of the following already standardized modes:

- Best Effort (support mandatory for all PPVPN types)
- Aggregate CE Interface Level QoS (i.e., ôhoseö level)
- Site-to-site, or ôpipeö level QoS

Note that all cases involving QoS MAY require that the CE and/or PE perform shaping and/or policing.

Mappings or translations of Layer 2 QoS parameters into Packet Switched Network QoS (e.g., DSCPs or MPLS EXP field) as well as QoS mapping based on VC (e.g., FR/ATM or VLAN) MAY be performed in order to provide QoS transparency. The actual mechanisms for these mappings or translations are outside the scope of this document. In addition, the Diffserv support of underlying tunneling technologies (e.g., [RFC3270] or [RFC3308]) and the Intserv model ([RFC2205]) MAY be used. As such, the L2VPN SLS requirements should be supported by appropriate core mechanisms.

4.7.2 Service Models

A service provider MUST be able to offer QoS service to a customer for at least the following generic service types: managed access VPN service or an edge-to-edge QoS service. The details of the service models can be found in [PPVPN-REQTS] and in [L3REQTS]. In L2VPN service, both DSCP and 802.1p fields MAY be used for this purpose.

4.8 Service Level Specifications

For a L2VPN service, the capabilities for Service Level Specification (SLS) monitoring and reporting stated in [PPVPN-REQTS] SHOULD be provided as appropriate.

4.9 Protection and Restoration

The L2VPN service infrastructure SHOULD provide redundant paths to assure high availability. The reaction to failures SHOULD result in an attempt to restore the service using alternative paths.

The intention is to keep the restoration time small. The restoration time MUST be less than the time it takes the CE devices, or customer

Layer 2 control protocols as well as Layer 3 routing protocols, to

detect a failure in the L2VPN.

4.10 CE-to-CE and CE-to-PE link requirements

The CE-to-PE links MAY either be direct physical links, e.g. 100BaseTX, T1/E1 TDM or logical links, e.g. ATM PVC, or RFC2427-encapsulated link, or transport networks carrying Ethernet, or a Layer 2 tunnel that go through a layer 3 network (e.g., L2TP sessions), over which Layer 2 traffic is carried.

Layer 2 frames MAY be tunneled through a layer 3 backbone from PE to PE, using one of a variety of tunneling technologies (e.g., IP-in-IP, GRE, MPLS, L2TP, etc.).

4.11 Management

Standard interfaces to manage L2VPN services MUST be provided (e.g., standard SNMP MIBs). These interfaces SHOULD provide access to configuration, verification and runtime monitoring protocols.

Service management MAY include the TMN 'FCAPS' functionalities, as follows: Fault, Configuration, Accounting, Provisioning, and Security, as detailed in [L3REQTS].

The ITU-T Telecommunications Management Network (TMN) model has the following generic requirements structure:

- o Engineer, deploy and manage the switching, routing and transmission resources supporting the service, from a network perspective (network element management);
- o Manage the L2VPNs deployed over these resources (network management);
- o Manage the L2VPN service (service management);
- o Manage the L2VPN business, mainly provisioning, administrative and accounting information related to the L2VPN service customers (business management).

4.12 Interoperability

Multi-vendor interoperability at network element, network and service levels among different implementations of the same technical solution SHOULD be guaranteed (that will likely rely on the completeness of the corresponding standard). This is a central requirement for SPs and customers.

The technical solution MUST be multi-vendor interoperable not only within the SP network infrastructure, but also with the customer's network equipment and services making usage of the L2VPN service.

A L2VPN solution SHOULD NOT preclude different access technologies. For instance, customer access connections to a L2VPN service MAY be different at different CE devices (e.g., Frame Relay, ATM, 802.1d, MPLS).

4.13 Inter-working

Inter-working scenarios among different solutions, providing L2VPN services, are highly desirable. Inter-working SHOULD be supported in a scalable manner.

Inter-working scenarios MUST consider at least traffic isolation, security, QoS, access, and management aspects. This requirement is essential in the case of network migration, to ensure service continuity among sites belonging to different portions of the network.

5 Customer Requirements

This section captures requirements from a customer perspective.

5.1 Service Provider Independence

Customers MAY require L2VPN service that spans multiple administrative domains or service provider networks. Therefore, a L2VPN service MUST be able to span multiple AS and SP networks, but still to act and to appear as a single, homogenous L2VPN from a customer point of view.

A customer might also start with a L2VPN provided in a single AS with a certain SLS but then ask for an expansion of the service spanning multiple ASs/SPs. In this case, as well as for all kinds of multi-AS/SP L2VPNs, L2VPN service SHOULD be able to deliver the same SLS to all sites in a VPN regardless of the AS/SP to which it homes.

5.2 Layer 3 Support

With the exception of IPLS, a L2VPN service SHOULD be agnostic to customerÆs layer 3 traffic (e.g., IP, IPX, Appletalk) encapsulated within Layer 2 frames.

IPLS MUST allow transport of IPv4 and IPv6 customerÆs traffic encapsulated within Layer 2 frames. IPLS SHOULD also allow CEs to run ISIS and MPLS protocols transparently among them when those are used in conjunction with IP.

5.3 Quality of Service and Traffic Parameters

QoS is expected to be an important aspect of a L2VPN service for some customers.

A customer requires that the L2VPN service provide the QoS applicable to his or her application, which can range from pseudo-wires (e.g., SONET emulation) to voice and interactive video, and multimedia applications. Hence, best-effort as well as delay and loss sensitive traffic MUST be supported over a L2VPN service.
A customer application SHOULD experience consistent QoS independent of the access network technology used at different sites connected

to the same L2VPN.

5.4 Service Level Specification

Most customers simply want their applications to perform well. A SLS is a vehicle for a customer to measure the quality of the service that SP(s) provide. Therefore, when purchasing a service, a customer requires access to the measures from the SP(s) that support the SLS.

Standard interfaces to monitor usage of L2VPN services SHOULD be provided (e.g., standard SNMP MIBs).

5.5 Security

5.5.1 Isolation

A L2VPN solution MUST provide traffic as well as forwarding information base isolation for customers similar to that obtained in private lines, FR, or ATM services.

A L2VPN service MAY use customer VLAN identifications as service delimiters. In that case, they MUST be honored and traffic separation MUST be provided.

5.5.2 Access control

A L2VPN solution MAY have the mechanisms to activate the appropriate filtering capabilities upon request of a customer. For instance, MAC and/or VLAN filtering MAY be considered between CE and PE for a VPLS.

5.5.3 Value added security services

A L2VPN solution MAY provide value added security services such as encryption and/or authentication of customer packets, certificate management, and similar.

Security measures employed by a L2VPN service SHOULD NOT restrict implementation of customer based security add-ons.

5.6 Network Access

Every packet exchanged between the customer and the SP over the access connection MUST appear as it would on a private network providing an equivalent service to that offered by the L2VPN.

5.6.1 Physical/Link Layer Technology

L2VPNs SHOULD support a broad range of physical and link layer access technologies, such as PSTN, ISDN, xDSL, cable modem, leased line, Ethernet, Ethernet VLAN, ATM, Frame Relay, Wireless local loop, mobile radio access, etc. The capacity and QoS achievable MAY

be dependent on the specific access technology in use.

5.6.2 Access Connectivity

Various types of physical connectivity scenarios MUST be supported, such as multi-homed sites, backdoor links between customer sites, devices homed to two or more SP networks. In case of VPLS, multi-link access for CE devices SHOULD be supported. L2VPN solutions SHOULD support at least the types of physical or link-layer connectivity arrangements shown in Figure 2 (in addition to the case

shown in Figure 1). For example, in case (b) a CE MAY connect to two different SPs via diverse access networks. Resiliency MAY be further enhanced as shown in case (d), where CE's, connected via a "back door" connection, connect to different SPs. Furthermore, arbitrary combinations of the above methods, with a few examples shown in cases (e) and (f) SHOULD be supportable by any L2VPN solution.

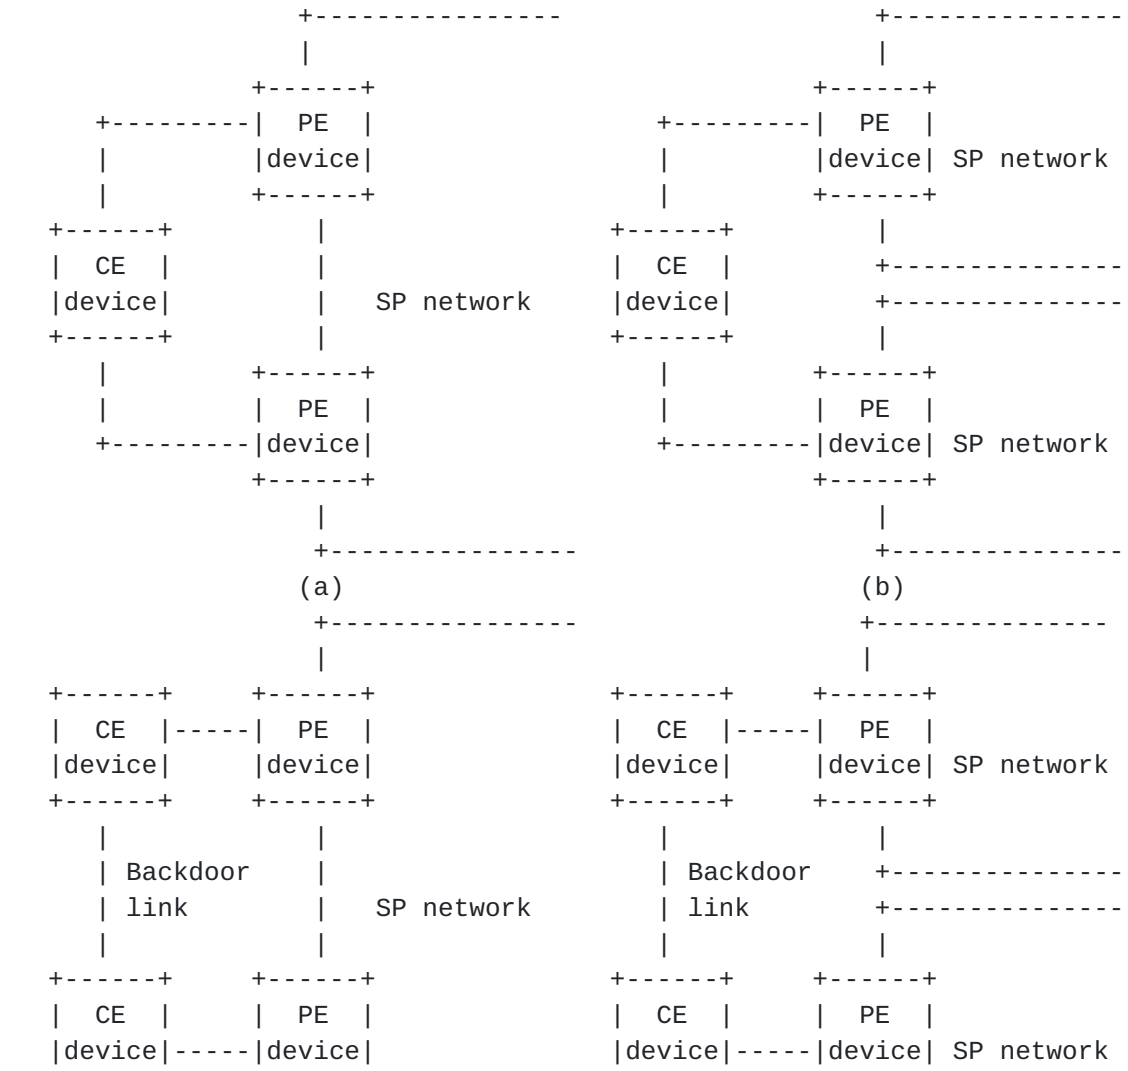

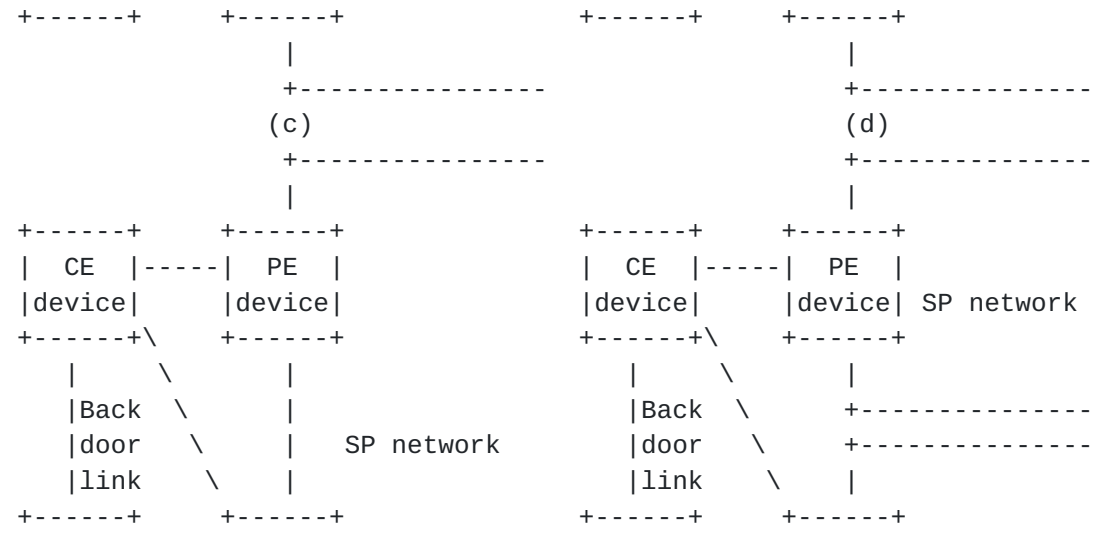

Figure 2 Representative types of access arrangements.

5.7 Customer traffic

5.7.1 Unicast, Unknown Unicast, Multicast, and Broadcast forwarding

A VPLS MUST deliver every packet at least to its intended destination(s) within the scope of the VPLS subject to the ingress policing and security policies.

5.7.2 Packet Re-ordering

The queuing and forwarding policies SHOULD preserve packet order for packets with the same QoS parameters.

5.7.3 Minimum MTU

A VPLS MUST support the theoretical MTU of the offered service.

The committed minimum MTU size MUST be the same for a given VPLS instance. Different L2VPN services MAY have different committed MTU sizes. If the customer VLANs are used as service delimiters, all VLANs within a given VPLS MUST inherit the same MTU size.

A VPLS MAY fragment packets as long as it is transparent to the customer.

5.7.4 End-point VLAN tag translation

The L2VPN service MAY support translation of customers' attachment circuit identifiers (e.g., VLAN tags, if the customer VLANs are used

as service delimiters). Such service simplifies connectivity of sites that want to keep their attachment circuit assignments or sites that belong to different administrative domains. In the latter case, the connectivity is sometimes referred to as Layer 2 extranet. On the other hand, it SHOULD be noted that VLAN tag translation affects the support for multiple spanning trees (i.e., 802.1s) and can break the proper operation.

5.7.5 Transparency

The L2VPN service is intended to be transparent to Layer 2 customer networks. A L2VPN solution SHOULD NOT require any special packet processing by the end users before sending packets to the provider's network.

If VLAN-ids are assigned by the SP, then VLANs are not transparent. Transparency does not apply in this case, as it is the same as FR/ATM service model.

Since the IPLS solution aims at transporting encapsulated traffic (rather than L2 frames themselves) the IPLS solution MUST not alter

the packets encapsulated inside Layer 2 frames which are transported by the IPLS. However, the IPLS solution is NOT REQUIRED to preserve the L2 Header transparently from CE to CE. For example, Source MAC address may not be preserved by the IPLS solution. The IPLS solution MAY remove Layer 2 headers for transport over the backbone when those can be reconstructed on egress without compromising transport of encapsulated traffic.

5.8 Support for Layer 2 Control Protocols

The L2VPN solution SHOULD allow transparent operation of Layer 2 control protocols employed by customers.

In case of VPLS, the L2VPN service MUST ensure that loops be prevented. This can be accomplished through a loop free topologies or appropriate forwarding rules. Control protocols such as Spanning Tree (STP) or similar could be employed. The L2VPN solution MAY use indications from customer Layer 2 control protocols, e.g. STP BPDU snooping, to improve the operation of a VPLS.

5.9 CE Provisioning

The L2VPN solution MUST require only minimal or no configuration on the CE devices, depending on the CE device that connects into the infrastructure.

6 Service Provider Network Requirements

This section describes requirements from a service provider perspective.

6.1 Scalability

This section contains projections regarding L2VPN sizing projections and scalability requirements and metrics specific to particular solutions.

6.1.1 Service Provider Capacity Sizing Projections

This section captures projections for scaling requirements over the next several years in terms of number of L2VPNs, number of interfaces per L2VPN, the size of forwarding information base per L2VPN, and the rate of L2VPN configuration changes. The examples are provided in [PPVPN-REQTS].

The numbers provided in this section are examples and MUST be treated as such. A L2VPN solution MAY scale much more than the examples provided here. Each requirement in this section MUST be considered independently.

A L2VPN solution SHOULD be scalable to support a very large number of L2VPNs per Service Provider network. The estimate is that a large service provider will require support O(10^5) VPWSs and O(10^4) VPLSs within the next four years.

A L2VPN solution SHOULD be scalable to support of a wide range of number of site interfaces per VPLS, depending on the size and/or

structure of the customer organization. The number of site interfaces SHOULD range from a few site interfaces to O(10^2) site interfaces per VPLS.

A L2VPN solution SHOULD be scalable to support a wide range of number of customer addresses (e.g., MAC) per VPLS. The number of customer addresses per VPLS MAY range from just a few (i.e., the number of sites when the CE devices are routers or when the service is IPLS) to a very large number such as 1,000 (i.e., when CE devices are switches). The number of customer addresses would be on the order of addresses supported in a typical native Layer 2 backbone.

A L2VPN solution SHOULD support high values of the frequency of configuration setup and change, e.g., for real-time provisioning of an on-demand videoconferencing or addition/deletion of sites.

Approaches SHOULD articulate scaling and performance limits for more complex deployment scenarios, such as inter-AS(S) L2VPNs and carriers' carrier. Approaches SHOULD also describe other dimensions of interest, such as capacity requirements or limits, number of inter-working instances supported as well as any scalability implications on management systems.

The number of users per VPLS is the combination of servers and hosts connected to the VPLS. It needs to scale from a handful to high numbers. A VPLS MUST scale from 2 users to a few hundred.

The number of users per VPLS interface follows the same logic as for users per VPLS. Further, it MUST be possible to have single user sites connected to the same VPLS as very large sites are connected to. VPLSs MUST scale from 1 user to a few hundred per site.

The number of sites per VPLS is clearly limited by the number of users for a VPLS. The largest number of sites in a VPLS would be equal to the largest number of users, distributed one per site.

The number of L2VPNs SHOULD scale linearly with the size of the access network and with the number of PEs.

6.1.2 Solution-Specific Metrics

Each L2VPN solution SHALL document its scalability characteristics in quantitative terms.

6.2 Identifiers

A SP domain MUST be uniquely identified at least within the set of all interconnected SP networks when supporting a L2VPN that spans multiple SPs. Ideally, this identifier SHOULD be globally unique (e.g., an AS number).

An identifier for each L2VPN SHOULD be unique, at least within each SP's network, as it MAY be used in auto-discovery, management (e.g, alarm and service correlation, troubleshooting, performance

statistics collection), and signaling. Ideally, the L2VPN identifier SHOULD be globally unique to support the case, where a L2VPN spans multiple SPs (e.g., [RFC2685]). Globally unique identifiers facilitate the support of inter-AS/SP L2VPNs.

6.3 Discovering L2VPN Related Information

Configuration of PE devices (i.e., U-PE and N-PE) is a significant task for a service provider. Solutions SHOULD provide methods that dynamically allow L2VPN information to be discovered by the PEs to minimize the configuration steps.

Each device in a L2VPN SHOULD be able to determine which other devices belong to the same L2VPN. Such a membership discovery scheme MUST prevent unauthorized access and allows authentication of the source.

Distribution of L2VPN information SHOULD be limited to those devices

involved in that L2VPN. A L2VPN solution SHOULD employ discovery mechanisms to minimize the amount of operational information maintained by the SPs. For example, if a SP adds or removes a customer port on a given PE, the remaining PEs SHOULD determine the necessary actions to take without the SP having to explicitly reconfigure those PEs.

A L2VPN solution SHOULD support the means for attached CEs to authenticate each other and verify that the service provider L2VPN is correctly configured.

The mechanism SHOULD respond to L2VPN membership changes in a timely manner. A "timely manner" is no longer than the provisioning timeframe, typically on the order of minutes, and MAY be as short as the timeframe required for "rerouting," typically on the order of seconds.

Dynamically creating, changing, and managing multiple L2VPN assignments to sites and/or customers is another aspect of membership that MUST be addressed in a L2VPN solution.

6.4 SLS Support

Typically, a SP offering a L2VPN service commits to specific Service Level Specifications (SLS) as part of a contract with the customer. Such a Service Level Agreement (SLA) drives the specific SP requirements for measuring Specific Service Level Specifications (SLS) for quality, availability, response time, and configuration intervals.

6.5 Quality of Service (QoS)

A significant aspect of a PPVPN is support for QoS. A SP has control over the provisioning of resources and configuration of parameters in at least the PE and P devices, and in some cases, the CE devices as well. Therefore, the SP is to provide either managed QoS access service, or edge-to-edge QoS service, as defined in [L3REQTS].

6.6 Isolation of Traffic and Forwarding Information

From a high level SP perspective, a L2VPN MUST isolate the exchange of traffic and forwarding information to only those sites that are authenticated and authorized members of a L2VPN.

A L2VPN solution SHOULD provide a means for meeting PPVPN QoS SLS requirements that isolates L2VPN traffic from the affects of traffic offered by non-VPN customers. Also, L2VPN solutions SHOULD provide a means so that traffic congestion produced by sites as part of one L2VPN does not affect another L2VPN.

6.7 Security

The security requirements are stated in Section 4.5. The requirements provided in [PPVPN-REQTS] and in [L3REQTS] SHOULD be met as appropriate.

In addition, a SP network MUST be immune to malformed or maliciously constructed customer traffic. This includes but not limited to duplicate or invalid Layer 2 addresses, customer side loops, short/long packets, spoofed management packets, spoofed VLAN tags, high volume traffic.

The SP network devices MUST NOT be accessible from any L2VPN, unless specifically authorized. The devices in the PPVPN network SHOULD provide some means of reporting intrusion attempts to the SP, if the intrusion is detected.

6.8 Inter-AS (SP) L2VPNs

All applicable SP requirements, such as traffic and forwarding information isolation, SLS's, management, security, provisioning, etc. MUST be preserved across adjacent ASÆs. The solution MUST describe the inter-SP network interface, encapsulation method(s), routing protocol(s), and all applicable parameters [VPN-IW].

A L2VPN solution MUST provide the specifics of offering L2VPN services spanning multiple ASs and/or SPs.

A L2VPN solution MUST support proper dissemination of operational parameters to all elements of a L2VPN service in the presence of multiple ASs and/or SPs. A L2VPN solution MUST employ mechanisms for sharing operational parameters between different ASs

A L2VPN solution SHOULD support policies for proper selection of operational parameters coming from different ASs. Similarly, a L2VPN solution SHOULD support policies for selecting information to be disseminated to different ASs.

6.8.1 Management

The general requirements for managing a single AS apply to a concatenation of AS's. A minimum subset of such capabilities is the following:

- Diagnostic tools
- Secured access to one AS management system by another
- Configuration request and status query tools
- Fault notification and trouble tracking tools

6.8.2 Bandwidth and QoS Brokering

When a L2VPN spans multiple AS's, there is a need for a brokering

mechanism that requests certain SLS parameters, such as bandwidth and QoS, from the other domains and/or networks involved in transferring traffic to various sites. The essential requirement is that a solution MUST be able to determine whether a set of AS's can establish and guarantee uniform QoS in support of a PPVPN.

6.9 L2VPN Wholesale

The architecture MUST support the possibility of one SP offering L2VPN service to another SP. One example is when one SP sells L2VPN service at wholesale to another SP, who then resells that L2VPN service to his or her customers.

6.10 Tunneling Requirements

Connectivity between CE sites or PE devices in the backbone SHOULD be able to use a range of tunneling technologies, such as L2TP, GRE, IP-in-IP, MPLS, etc.

Every PE MUST support a tunnel setup protocol, if tunneling is used. A PE MAY support static configuration. If employed, a tunnel establishment protocol SHOULD be capable of conveying information, such as the following:

- Relevant identifiers
- QoS/SLS parameters
- Restoration parameters
- Multiplexing identifiers
- Security parameters

There MUST be a means to monitor the following aspects of tunnels:

- Statistics, such as amount of time spent in the up and down state
- Count of transitions between the up and down state
- Events, such as transitions between the up and down states

The tunneling technology used by the VPN Service Provider and its associated mechanisms for tunnel establishment, multiplexing, and maintenance MUST meet the requirements on scaling, isolation, security, QoS, manageability, etc.

Regardless of the tunneling choice, the existence of the tunnels and their operations MUST be transparent to the customers.

6.11 Support for Access Technologies

The connectivity between PE and CE devices is referred to as an Attachment Circuit (AC). ACs MAY span networks of other providers or public networks.

There are several choices for implementing ACs. Some popular choices

include Ethernet, ATM (DSL), Frame Relay, MPLS-based virtual circuits etc.

In case of VPLS, the AC MUST use Ethernet frames as the Service Protocol Data Unit (SPDU).

A CE access connection over an AC MUST be bi-directional in nature.

PE devices MAY support multiple ACs on a single physical interface. In such cases, PE devices MUST NOT rely on customer controlled parameters for distinguishing between different access connections. For example, if VLAN tags were used for that purpose, the provider would be controlling the assignment of the VLAN tag values and would strictly enforce compliance by the CEs.

An AC, whether direct or virtual, MUST maintain all committed characteristics of the customer traffic, such as QoS, priorities etc. The characteristics of an AC are only applicable to that connection.

6.12 Backbone Networks

Ideally, the backbone interconnecting SP PE and P devices SHOULD be independent of physical and link layer technology. Nevertheless, the characteristics of backbone technology MUST be taken into account when specifying the QoS aspects of SLSs for VPN service offerings.

6.13 Network Resource Partitioning and Sharing Between L2VPNs

In case network resources such as memory space, FIB table, bandwidth and CPU processing are shared between L2VPNs, the solution SHOULD guarantee availability of resources necessary to prevent any specific L2VPN instance from taking up available network resources and causing others to fail. The solution SHOULD be able to limit the resources consumed by a L2VPN instance. The solution SHOULD guarantee availability of resources necessary to fulfill the obligation of committed SLSs.

6.14 Interoperability

Service providers are interested in interoperability in at least the following scenarios:

- To facilitate use of PE and managed CE devices within a single SP network
- To implement L2VPN services across two or more interconnected SP networks
- To achieve inter-working or interconnection between customer sites using different L2VPN solutions or different implementations of the same approach

Each approach MUST describe whether any of the above objectives can be met. If an objective can be met, the approach MUST describe how such interoperability could be achieved.

6.15 Testing

The L2VPN solution SHOULD provide the ability to test and verify operational and maintenance activities on a per L2VPN service basis, and in case of VPLS, on a per VLAN basis if customer VLANs are used as service delimiters.

The L2VPN solution SHOULD provide mechanisms for connectivity verification, and for detecting/locating faults.

Examples of testing mechanisms are as follows:

- o Checking connectivity between "service-aware" network nodes
- o Verifying data plane and control plane integrity
- o Verifying service membership

The provided mechanisms MUST satisfy the following: the connectivity checking for a given customer MUST enable the end-to-end testing of the data path used by that customer's data packets and the test packets MUST not propagate beyond the boundary of the SP network.

6.16 Support on Existing PEs

To the extent possible, the IPLS solution SHOULD facilitate support of IPLS on existing PE devices that may be already deployed by the Service Provider and may have been designed primarily for ôLayer 3ö services.

7 Service Provider Management Requirements

A service provider desires to have a means to view the topology, operational state, and other parameters associated with each customer's L2VPN. Furthermore, the service provider requires a means to view the underlying logical and physical topology, operational state, provisioning status, and other parameters associated with the equipment providing the L2VPN service(s) to its customers. Therefore, the devices SHOULD provide standards-based interfaces (e.g., L2VPN MIBs) wherever feasible.

The details of service provider management requirements for a Network Management System (NMS) in the traditional fault, configuration, accounting, performance, and security (FCAPS) management categories can be found in [Y.1311.1].

8 Engineering Requirements

These requirements are driven by implementation characteristics that make service and SP requirements achievable.

8.1 Control Plane Requirements

A L2VPN service SHOULD be provisioned with minimum number of steps. Therefore, the control protocols SHOULD provide methods for

signaling between PEs. The signaling SHOULD inform of membership, tunneling information, and other relevant parameters.

The infrastructure MAY employ manual configuration methods to provide this type of information.

The infrastructure SHOULD use policies to scope the membership and reachability advertisements for a particular L2VPN service. A mechanism for isolating the distribution of reachability information to only those sites associated with a L2VPN MUST be provided.

The control plane traffic increases with the growth of L2VPN membership. Similarly, the control plane traffic increases with the number of supported L2VPN services. The use of control plane resources MAY increase as the number of hosts connected to a L2VPN service grows.

A L2VPN solution SHOULD minimize control plane traffic and the consumption of control plane resources. The control plane MAY offer means for enforcing a limit on the number of customer hosts attached to a L2VPN service.

8.2 Data Plane Requirements

8.2.1 Encapsulation

A L2VPN solution SHOULD utilize the encapsulation techniques defined by PWE3 ([PWE3-FR]), and SHOULD not impose any new requirements on these techniques.

8.2.2 Responsiveness to Congestion

A L2VPN solution SHOULD utilize the congestion avoidance techniques defined by PWE3 ([PWE3-FR]).

8.2.3 Broadcast Domain

A separate Broadcast Domain MUST be maintained for each VPLS.

In addition to VPLS Broadcast Domains, a L2VPN service MAY honor customer VLAN Broadcast Domains, if customer VLANs are used as service delimiters. In that case, the L2VPN solution SHOULD maintain a separate VLAN Broadcast Domain for each customer VLAN.

8.2.4 Virtual Switching Instance

L2VPN Provider Edge devices MUST maintain a separate Virtual Switching Instance (VSI) per each VPLS. Each VSI MUST have capabilities to forward traffic based on customer's traffic parameters such as MAC addresses, VLAN tags (if supported), etc. as well as local policies.

L2VPN Provider Edge devices MUST have capabilities to classify incoming customer traffic into the appropriate VSI.

Each VSI MUST have flooding capabilities for its Broadcast Domain to facilitate proper forwarding of Broadcast, Multicast and Unknown Unicast customer traffic.

8.2.5 MAC address learning

A VPLS SHOULD derive all topology and forwarding information from packets originating at customer sites. Typically, MAC address learning mechanisms are used for this purpose. With IPLS, snooping of particular packets originating at customer sites and signaling might also be used.

Dynamic population of the Forwarding Information Base (e.g. via MAC address learning) MUST take place on a per Virtual Switching Instance (VSI) basis, i.e. in the context of a VPLS and, if supported, in the context of VLANs therein.

9 Security Considerations

Security considerations occur at several levels and dimensions within Layer 2 Provider Provisioned VPNs, as detailed within this document.

The requirements in this document separate the notion of traditional security requirements, such as integrity, confidentiality, and authentication as detailed in section 4.5 from that of isolating (or separating) the exchange of forwarded packets and exchange of forwarding information between specific sets of sites. Further details on security requirements are given from the customer and service provider perspectives in sections 5.5 and 6.7, respectively. In an analogous manner, further detail on traffic and routing isolation requirements are given from the customer and service provider perspectives in sections 4.4 and 6.6, respectively. Safeguards to protect network resources such as CPU, memory, and bandwidth are required in section 6.13.

IPSec can be also be applied after tunneling Layer-2 traffic to provide additional security.

10 Acknowledgments

The authors would like to acknowledge extensive comments and contributions provided by Loa Andersson, Joel Halpern, Eric Rosen, Ali Sajassi, Muneyoshi Suzuki, Ananth Nagarajan, Dinesh Mohan, Yakov Rekhter, Matt Squire, Norm Finn, Scott Bradner, and Francois Le Faucheur. The authors, also, wish to extend their appreciationÆs to

their respective employers and various other people who volunteered to review this work and provided feedback. This work was done in consultation with the entire Layer 2 PPVPN design team. A lot of the text was adapted from the Layer 3 requirements document produced by Layer 3 requirements design team.

11 References

11.1 Normative References

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

[TERMINOLOGY] Andersson, L, Madsen, T. "PPVPN Terminology", work in progress

11.2 Non-normative References

[GENERIC-REQTS] Nagarajan, A., et al. "Generic Requirements for Provider Provisioned VPN", work in progress

[PPVPN-L2-FR] Andersson, L, et al. "PPVPN L2 Framework", work in progress

[RFC3270] Le Faucheur, F., et al. "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.

[RFC3308] Calhoun, P., et al, "Layer 2 Tunneling Protocol (L2TP) Differentiated Services Extension", RFC 3308, November 2002.

[RFC2205] Braden, R., et al, "Resource ReSerVation Protocol (RSVP)", RFC 2205, September 1997.

[L3REQTS] Carugi, M., McDysan, D. et. al., "Service Requirements for Layer 3 Provider Provisioned Virtual Private Networks", work in progress

[Y.1311.1] Carugi, M. (editor), "Network Based IP VPN over MPLS architecture",Y.1311.1 ITU-T Recommendation, May 2001 (http://standards.nortelnetworks.com/ppvpn/relateditu.htm)

[RFC2685] Fox B., et al, "Virtual Private Networks Identifier", RFC 2685, September 1999.

[VPN-IW] H. Kurakami et al, "Provider-Provisioned VPNs Interworking," work in progress.

[PWE3-FR] Pate, P, et al. "Framework for Pseudo Wire Emulation Edge-to-Edge (PWE3)", work in progress

12 Editors' Addresses

Waldemar Augustyn

Email: waldemar@nxp.com

Yetik Serbest
SBC Technology Resources
9505 Arboretum Blvd.
Austin, TX 78759
Email: serbest@tri.sbc.com