Internet Engineering Task Force
INTERNET DRAFT
Informational

CY Lee
M. Higashiyama

March 2003

CE-based Virtual Private LAN

<draft-lee-ce-based-vpl-02.txt>

Abstract

This draft describes how a Virtual Private LAN (VPL) can be realized by setting up point-to-point tunnels from CE (Customer Edge) or CLE (Customer Located Equipment) to CE/CLE over a PSN, and bridging Ethernet traffic at the CE/CLE. Regardless of the access technology used (e.g. DSL or Ethernet) or the geographic distribution of CEs/CLEs, an end customer may use an emulated LAN service via an Ethernet interface, as long as the CEs/CLEs of the emulated LAN are reachable over the PSN. When used in conjunction with IPSec, the proposal allows secure transmission of Ethernet traffic, from one site to another site of a VPL, over the PSN.

1. Introduction

CE-based VPL over different types of tunneling technologies has been used for a number of years now, and could be viewed as a proven technology. A network user provisions the required tunnels (or circuits) at a CE to remote CE(s) and the CEs bridge Ethernet traffic over the tunnels.

2. Topology

2.1 Reference Model

Figure 1 shows CE-based VPL topology. In a CE-based VPL, tunnels are setup between sites of a VPL. A CE-based VPL can contain a single IEEE 802.1q VLAN or multiple VLANs. Each site has either a CE or CLE connected to the PSN. In a PPVPL, the provider provisions the VPL, a tunnel MAY be setup from CLE to CLE, and the CLEs MAY be owned by the provider, or the tunnels MAY be setup from CE to CE, and the CEs MAY be owned by the customer. In a CPVPL, a customer provisions its own VPL, a point to point tunnel from CE to CE may be provisioned by the customer at CEs or alternatively, a point to point tunnel may be provisioned by the provider from PE to PE. A tunnel appears as a virtual port or interface to the bridge entity in a CE.

At a CE, Ethernet traffic from a VPL is encapsulated in for e.g. a L2TPv3 or GRE or IPSec tunnel or FR VC or ATM VCC and transported over the IP/FR/ATM network to another CE of the VPL. The receiving CE decapsulates the Ethernet frame, and bridges the frame from virtual port to the destination node in the VPL.

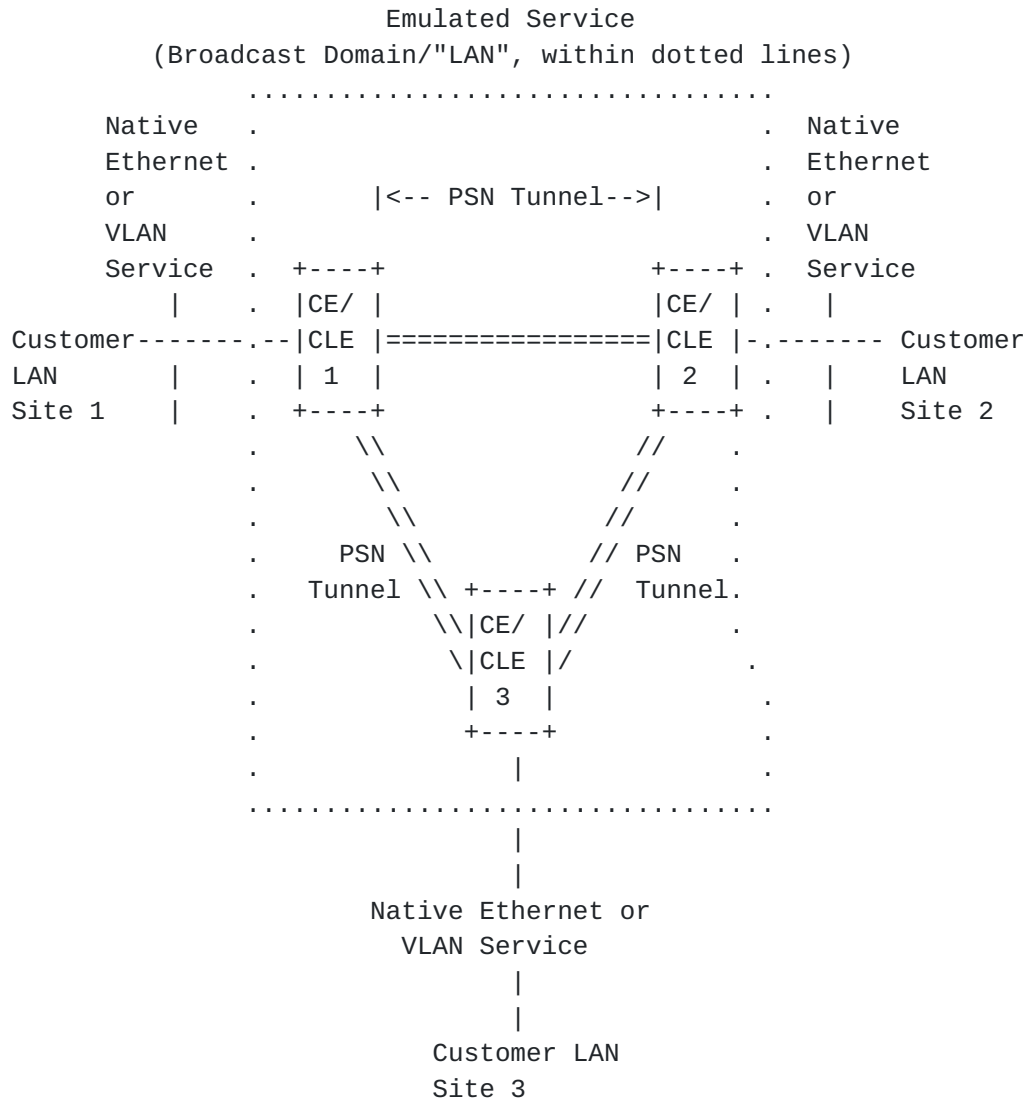
Emulated Service
(Broadcast Domain/"LAN", within dotted lines)
Native Ethernet or VLAN Service
Native Ethernet or VLAN Service
|<-- PSN Tunnel-->|
Customer LAN Site 1
Customer LAN Site 2
PSN Tunnel
PSN Tunnel
CE 1
CLE 1
CE 2
CLE 2
CE 3
CLE 3
Native Ethernet or VLAN Service
Customer LAN Site 3

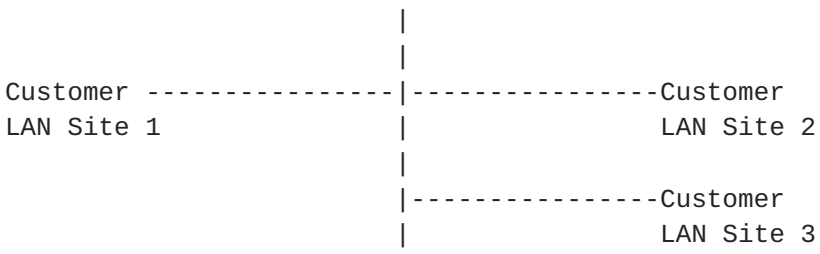

Fig 1 Emulated Ethernet Segment

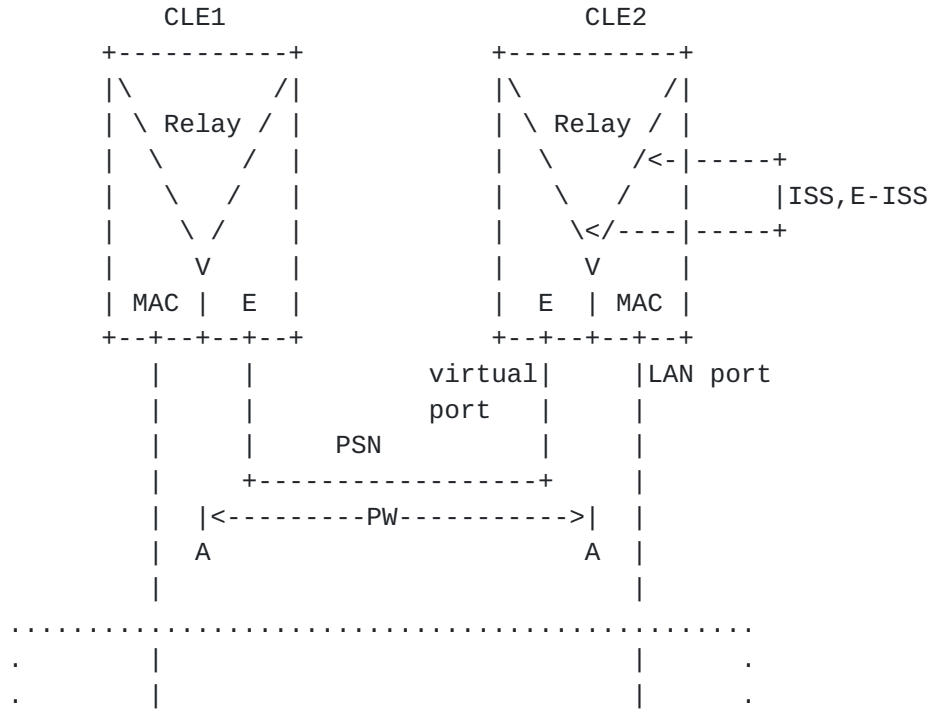

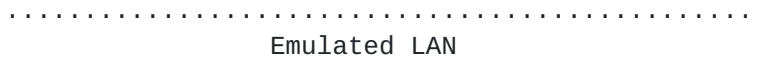

Figure 2

The Encapsulation (E) Entity, is responsible for encapsulating frames to be transported over the PSN. The Internal Sublayer Service (ISS) and Enhanced Internal Sublayer Service (E-ISS), i.e. the indication and request primitives provided by a MAC entity to the MAC Relay Entity within a CE is as defined in 802.1d and 802.1q. The E entity provides similar service primitives as the MAC entity, except that the handling of the Frame Check Sequence (FCS) parameter may be changed.

The Relay is a bridge as defined in 802.1d and optionally 802.1q. CLEs process BPDUs, as defined in 802.1d and optionally 802.1q, and MAC control frames as defined in IEEE specifications.

If a modified bridge (performing reverse split horizon forwarding on a full-meshed of PWs, I.e. traffic received on a PW is not forwarded to other PWs) is used, the modified bridge would reside between the Relay and the VP. The use of a modified bridge in a CLE is being reviewed, as it may introduce routing or briding issues in the customer's network.

2.2 Terminology

CE	Customer Edge [PPVPN-REQ]
CLE	Customer Located Equipment
CPVPL	Customer Provisioned VPL
LCCE	L2TP Control Connection Endpoint [L2TPv3]
MAN	Metropolitan Area Network
P	Provider's Network Equipment (excluding PE)
PE	Provider Edge [PPVPN-REQ]
PPVPL	Provider Provisioned VPL
PSN	Packet Switched Network
PW	Pseudo-Wire
VPL	Virtual Private LAN

3. Control Plane

For the CE-based VPL, control plane use [ETH-PW-L2TPv3]. Pseudo-wire Ethernet over L2TPv3 is specified in [ETH-PW-L2TPv3] and describes how network devices emulate ethernet over an IP network using L2TPv3 as the tunneling protocol. CEs supporting [ETH-PW-L2TPv3]in a PPVPL or CPVPL are applications of [ETH-PW-L2TPv3].

Below additional functions of these applications are required.
- Tunnel Endpoints Information Configuration

- VPL Monitoring

3.1 Tunnel Endpoints Information Configuration

The required tunnel endpoint information at an LCCE (CE) are the IP addresses and End Identifiers of peer LCCEs, and authentication keys.

The tunnel endpoints information may be pre-configured or remotely provisioned or, a mechanism to discover and distribute the tunnel endpoints information may be used or a mechanism where tunnel endponts information are retrieved from a server may be used.

3.2 Discovery

Multiple tunnels to other CE sites can be automatically configured on CEs if a tunnel endpoint information discovery mechanism is used.

To avoid having to provision deployed CEs, a mechanism to auto discover and distribute VPL site information is useful. [CE-AUTOCONFIG] or a directory query approach similar to [VPLS-DNS] are examples of mechanisms that may be used for this purpose. In particular for [CE-AUTOCONFIG], the Authentication Server/RADIUS approach is applicable to both a PPVPL and CPVPL, while the DHCP approach is only applicable to a PPVPL where the CEs are within one domain (e.g for VPLs offered by a network provider).

3.3 VPL Monitoring

The session keep-alive mechanism of L2TPv3 can serve as a link status monitoring mechanism for the point to point tunnels (session) that make up the VPL. Testing of reachability of nodes in the VPL from different sites may be performed at irregular intervals.

4. Data Plane

4.1 Encapsulation

Please see [ETH-L2TPv3].

4.2 Forwarding

A CE learns MAC addresses from the customer facing ports and the virtual interfaces (or the tunnels to remote LCCE sites of a VPL). When a new MAC address is learned, the MAC address is associated with the virtual interface or ports where the frame arrives. When a frame with the cached MAC address is received, the CE knows which virtual interface or port to forward the frame to. When a frame with a new MAC address is received, a CE floods the frame to all other ports or virtual interfaces, except the interface where the frame is received

from.

The learning, bridging, filtering and forwarding procedures are as defined in [802.1d] and [802.1q], except that the ports on a switch in this case can be a virtual interface as well as a physical port.

CEs belonging to the same VPL learn, store, manage VPL forwarding information and bridges traffic within the VPL, PEs do not have to learn MAC addresses from different VPLs, hence this approach scales for large number of VPLs and total MAC addresses in a network.

Bridging within a VPL does not affect other VPLs (or customers). A CE bridge which is not functioning correctly will only affect a VPL. In contrast, if bridging is also performed at PEs, a malfunctioning CE may cause network instability and affect other VPLs as well. Hence a CE-based VPL would be operationally stabler.

In a network based VPL, as the number of customers/VPLS and total MAC addresses grow in a provider's network, existing devices in the network will need to be upgraded or replaced by new devices. A CE-based VPL approach scales as the number of VPLs and total number of MAC addresses in VPLs grows and allows CEs in different MAN to be interconnected seamlessly.

New VPLs can be added transparently in an IP/MPLS network, without having to upgrade PEs with bridging functions.

CEs of a VPL may be located witin one domain or in different domains as long as the CEs have IP reachability to each other. CEs of a VPL in different MAN can send VPL traffic to each other without resorting to VLAN Stacking as long as there is IP reachability in the network.

Tunnels may be setup using L2TP with IPSec [L2TP-IPSec] to provide secure transmission of traffic from CE to CE in an IP PSN.

5. References

5.1 Normative References

[802.1D] IEEE, "ISO/IEC 15802-3:1998,(802.1D, 1998 Edition), Information technology --Telecommunications and information exchange between systems --IEEE standard for local and metropolitan area networks --Common specifications-Media access control (MAC) Bridges", June, 1998.

[802.1Q] ANSI/IEEE Standard 802.1Q, "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks", 1998 .

[802.3] IEEE, "ISO/IEC 8802-3: 2000 (E), Information technology--Telecommunications and information exchange between systems --Local and metropolitan area networks --Specific requirements --Part 3:

Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications", 2000.

[L2TPv3] Lau, J., Townsley, M., Valencia, A., Zorn, G., Goyret, I., Pall, G., Rubens, A., Palter, B., "Layer Two Tunneling Protocol "L2TP"", (draft-ietf-l2tpext- l2tp-base-01.txt), work in progress, July 2001.

[L2TP-IPSEC] RFC 3193, B. Patel,B. Aboba,W. Dixon, G. Zorn, S. Booth "Securing L2TP using IPSec"

[ETH-L2TPv3] Aggarwal, et al., Transport of Ethernet Frames over L2TPv3, draft-ietf-l2tpext-pwe3-ethernet-00.txt, October 2002.

[ETH-PW-L2TPv3] CY Lee, M Higashiyama, Ethernet Pseudo-Wire over L2TPv3, November 2002

5.2 Informative References

[L2VPN-REQ] W. Augustyn, Y. Serbest, Service Requirements for Layer 2 Provider Provisioned Virtual Private Networks, February 2003

[BCP] Mitsuru H. and Baker, "PPP Bridging Control Protocol (BCP)", RFC 2878, July 2000.

[L2TP] Townsley, W., Valencia, A., Rubens, A., Singh Pall, G., Zorn, G., Palter, B., "Layer Two Tunneling Protocol (L2TP)", RFC 2661 August 1999

Internet Architectural Guidelines and Philosophy. http://www.ietf.org/internet-drafts/draft-ymbk-arch-guidelines-05.txt

[EOL2TP] M. Higashiyama, "Ethernet Over L2TP", (draft-higashiyama-eol2tp-01.txt),

[PPVPN-REQ] M. Carugi,D. McDysan, L. Fang, F. Johansson, Ananth Nagarajan, J. Sumimoto, R. Wilder, "Service requirements for Layer 3 Provider Provisioned Virtual Private Networks" (draft-ietf-ppvpn-requirements-04.txt)

[CEVPN] De Clercq J., et al., "Provider Provisioned CE-based Virtual Private Networks using IPsec", draft-ietf-ppvpn-ce-based-01.txt, work in progress.

[Kompella] Kompella, K., Leelanivas, M., Vohra, Q., Bonica, R., Metz, E., Ould-Brahim, H., Achirica, J., Z., "MPLS-based Layer 2 VPNs", (draft-kompella- ppvpn-l2vpn-00.txt), work in progress, July 2001.

[Martini-encap] Martini, L., El-Aawar, N., Tappan, D., Rosen, E., Jayakumar, J., Vlachos, D., Liljenstolpe, C., Heron, G., Kompella, K., Vogelsang, S., Shirron, J., Smith, T., Radoaca, V., Malis, A., Sirkay, V., Cooper, D., "Encapsulation Methods for Transport of

Layer 2 Frames Over IP and MPLS Networks", (draft-martini-l2circuit-encap-mpls-03.txt), work in progress, July 2001.

[PWE3-frame] Pate, P., Xiao, X., So, T., Malis, A., Nadeau, T., White, C., Kompella, K., Johnson, T., "Framework for Pseudo Wire Emulation Edge-to-Edge (PWE3)" (draft- pate-pwe3-framework-02.txt), work in progress, July 2001.

[VPLS] Lasserre, M, Kompella, V, et al, "Virtual Private LAN Services over MPLS" draft-lasserre-vkompella-ppvpn-vpls-01.txt, March 2002

[VPLS-DNS] Heinanen, "DNS/LDP Based VPLS". draft-heinanen-dns-ldp-vpls-00.txt, January 2002.

[CE-AUTOCONFIG] CY Lee, "CE Auto-Configuration", (draft-lee-ppvpn-ce-auto-config-01.txt), work in progress, July 2002

[RADIUS] Rigney, C., Willens, S., Rubens, A. and W. Simpson, "Remote Authentication Dial in User Service (RADIUS)", RFC 2865, June 2000.

6. Security Considerations

It is recommended to use IPsec in conjunction with L2TPv3 to secure communication. Please see the Appendix for further information on Security Considerations.

7. IANA Consideration

A new PW Type for CE-based VPL is required.

8. Acknowledgment

The authors would like to thank Jeremy deClercq and Jeanne DeJaegher for their helpful comments on the initial version of this draft. The draft benefited from discussions with Sasha Cirkovic, Jeff Smith, Raymond Chang, Roy Nighswander, Neil Harrison, Alexis Berthillier, Dean Welsh and Arnold Jansen.

9. Author's Address

Cheng-Yin Lee Alcatel 600 March Rd, Ottawa Ontario, Canada K2K 2E6
e-mail: Cheng-Yin.Lee@alcatel.com

Mitsuru Higashiyama Anritsu Corporation 1800 Onna, Atsugi-shi, Kanagawa-prf., 243-8555 Japan e-mail: Mitsuru.Higashiyama@yy.anritsu.co.jp

10. Full Copyright Statement

Appendix

This Appendix describes how CE-based VPL (CEVPL) satisfies [L2VPN-REQ]. The sections are numbered as in [L2VPN-REQ].

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.

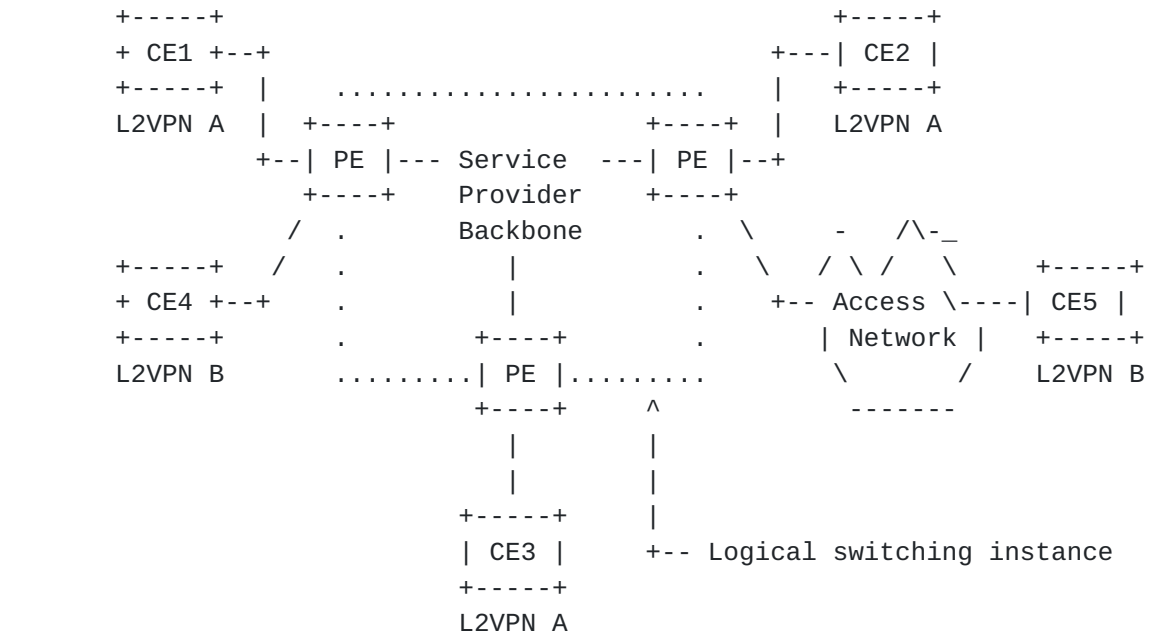

Figure 1 L2VPN Reference Model

4. Service Requirements Common to CUstomers and Service Providers

4.1 Scope of emulation

CEVPL interoperates with the existing protocols and standards (802.1d/q, 802.3 control frames handling) of the Ethernet network the

customer is managing.

Some possibly salient differences between CEVPL and a real LAN are:
- the reliability MAY likely be less -- a message sent to a remote node may not be received by the remote bridge because the remote bridge may be temporarily not reachable -- a message broadcast over a CEVPL using a spanning tree may not be seen by one of the remote CE if the PSN is partitioned -- the probability that a message broadcast over a CEVPL using a spanning tree is not seen by one of the remote CE is lower than if reverse split horizon forwarding is used (which can occur whenever a PSN tunnel is broken)

- if sequencing is not turned on, BPDUs on a PW may get out of oerder
- Ethernet frames can get duplicated or received out of sequence if the sequencing option is not turned on. - 802.3 Pause frames will be handled by a CE/CLE as defined in 802.3

CEVPL interoperates with customer equipment as specified in IEEE specifications.

4.2 Traffic Types

CE VPL support unicast, multicast, and broadcast traffic. It is desirable that the CE/CLE support efficient replication of broadcast and multicast traffic.

4.3 Topology

A CEVPL allows multipoint to multipoint connectivity over point to point tunnels setup for different types topologies e.g. :

- o 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

Not all traffic characteristics (such as bandwidth, QoS, delay, etc.) would necessarily be the same between any two end points of a CEVPL. The SLS requirements of a service would have a bearing on the type of topology that can be used.

A CEVPL is capable of crossing multiple administrative boundaries.

CEVPL services are independent of access network technology.

4.4 Isolated Exchange of Data and Forwarding Information

CEVPL has means to allow CE to authenticate each other during tunnel setup. If a CE needs to obtain configuration information from e.g. a server, a CE configuration solution which allows authentication and authorization SHOULD be used.

CEVPL implemented according to specification SHOULD not introduce undesired forwarding information that could corrupt the L2VPN forwarding information base.

CEVPL constrains, or isolates, the distribution of addressed data to only those VPLS sites determined either by MAC learning.

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

4.5 Security

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. Methods to protect against the following situations are listed in >>.

- Protocol attacks
  - o Excessive protocol adjacency setup/teardown
    >> This may not be applicable to CEVPL.
  - o Excessive protocol signaling/withdrawal
    >> Setting a limit on the tunnel setup frequency needs to be specified in CEVPL. A provider may also ingress rate-limit tunnel setup messages.

- Resource Utilization
  - o Forwarding plane replication (VPLS)
    >> emulated LAN topology/network engineering to limit the number of CEs that a CE shall be connected to
  - o Looping (VPLS primarily)
    >> STP
  - o MAC learning table size limit (VPLS)
    >> Same requirement as all bridges in the LAN, the MAC table is not shared by other customers
- Unauthorized access
  - o Unauthorized member of VPN
    >> A remote CE is authenticated during tunnel setup
  - o Incorrect customer interface
    >> A remote CE is authenticated during tunnel setup
  - o Incorrect service delimiting VLAN tag
    >> N/A
  - o Unauthorized access to PE
    >> A PE may be protected from authorized in the same as an existing NE in a provider's network
- Tampering with signaling
  - o Incorrect FEC signaling
    >> N/A
  - o Incorrect PW ID assignment

>> See [L2TPv3]
  o Incorrect signaled VPN parameters (e.g., QoS, MTU, etc.)
>> See [L2TPv3]
- Tampering with data forwarding
  o Incorrect MAC learning entry
>> A CE can only be updated by authorized entites
  o Incorrect Session ID/PW ID
>> See [L2TPv3]
  o Incorrect customer facing encapsulation
>> See [L2TPv3]
  o Incorrect pseudo-wire encapsulation
>> See [L2TP3]
  o Hijacking pseudowires using the wrong tunnel
>> See [L2TPv3]
  o Incorrect tunnel encapsulation
>> See [L2TPv3]

4.5.1 User data security

CEVPL provides traffic separation between different VPL.

in 0 4.5.2 Access control

A CEVPL MAY have the ability to activate the appropriate filtering capabilities upon request of a customer, if CE configuration or management such as LMI is used.

4.6 Addressing

CEVPL supports overlapping addresses of different L2VPNs. For instance, customers are not prevented from using the same MAC addresses and/or the same VLAN Ids when used with different L2VPNs. A CEVPL is oblivious to customer VLANs, I.e. customers can have overlapping VLAN Ids.

4.7 Quality of Service

A CEVPL QoS SHOULD be independent of the access network technology.

4.7.1 QoS Standards

A CEVPL 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., hose level)
- Site-to-site, or pipe level QoS

This MAY require that the CE and/or PE perform shaping and/or policing.

Mappings or translations of Layer 2 QoS parameters into PacketSwitched work 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 CEVPL service, 802.1p fields shall be mapped to DSCP.

4.8 Service Level Specifications

For a CEVPL, the monitoring and reporting of the VPL is described in "VPL Monitoring". Further work is required to fulfil the capabilities for Service Level Specification (SLS) monitoring and reporting stated in [PPVPN-REQTS].

4.9 Protection and Restoration

To assure high availability, the underlying PSN SHOULD be able to provide fast failover to alternative paths or rerouting of traffic.

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 Ethernet traffic is carried.

Ethernet frames MAY be tunneled through a layer 3 backbone from CE/CLE to CE/CLE, using the following tunneling technologies (e.g., IP-in- IP, GRE, MPLS, L2TP, etc.).

4.11 Management

Standard interfaces to manage Ethernet services is provided (e.g., standard SNMP MIBs). These interfaces shall provide access to configuration, verification and runtime monitoring protocols.

The Service management MAY include the TMN 'FCAPS' functionalities,

as follows: Fault, Configuration, Accounting, Provisioning, and Security, as detailed in [L3REQTS]. Accounting is out of scope of this draft.

4.12 Interoperability

If CEVPL is standardized, it should promote interoperability among CEs/CLEs from different vendors. CEVPL is defined to interoperate with existing 802.1d/q devices at customers' LAN sites.

A CEVPL is agnostic to different access technologies. For instance, customer access connections to a CEVPL service MAY be different at different CE/CLE devices (e.g. Ethernet, DSL, ATM)

4.13 Inter-working

CEVPL may inter-work with PE-based VPLS or Provider Bridges if the customer facing port of a CLE is connected to the customer facing side of a PE. The scalability of this inter-working is dependent on the scalability of PE-based VPLS or Provider Bridges.

CEVPL may inter-work with Hybrid VPLS [HYBRID-VPLS] if the tunnel from a CE/CLE is terminated at a PE. The PE switch the decapsulated traffic onto an Attachment Circuit. The CE/CLE at the end of the Attachment Circuit bridges the traffic accordingly. Inter-working in this case scales as well as Hybrid VPLS. This allows CEs/CLEs with IP access to peer with CEs/CLEs with Ethernet, FR or ATM access. The SLA for CEs/CLEs with different types of access technologies may vary and dependent on the customer's network requirements.

In both inter-working scenarios, customer traffic isolation is maintained. CEVPL may encrypt transmission of traffic over an IP network from CE/CLE to CE/CLE but when interworking with other VPLS solutions which terminates tunnels at PEs, the traffic is encrypted from a CE/CLE to a PE.

CEVPL is able to work independently of other VPLS solutions in a network (Ships in the Night). Sites belonging to different VPLS networks shall continue to have emulated LAN service while CEVPL inter-working is being introduced.

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. Sites of a CEVPL is able to span multiple AS and SP networks, but still be able to act and to appear as a single, homogenous emulated LAN from a customer point of view.

A customer might also start with a VPL 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, a 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, if all AS and SP gives similar and consistent treatment of 802.1q p bits, DSCP and COS, but this is not generally practiced.

5.2 Layer 3 Support

A CEVPL is agnostic to customer layer 3 traffic (e.g., IPv4/v6, IPX, Appletalk) encapsulated within Layer 2 frames.

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 a CEVPL service provide the QoS applicable to his or her application, which can range from voice and interactive video to multimedia applications. Hence, best-effort as well as delay and loss sensitive traffic MUST be supported over a L2VPN service. The support of this requirement is dependent on the PSN QoS support.

A customer application SHOULD experience consistent QoS independent of the access network technology used at different sites connected to the same L2VPN. Again, this is dependent on the SP's network being able to support QoS consistently independent of access network technology.

5.4 Service Level Specification

Most customers simply want their applications to perform well. A SLS is a vehicle for a customer to measuere 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 CEVPL services SHALL be provided (e.g., standard SNMP MIBs).

5.5 Security

5.5.1 Isolation

A CEVPL service provides traffic as well as forwarding information base isolation for customers similar to that obtained in private lines, FR, or ATM services.

A CEVPL service MAY use customer VLAN identifications as service delimiters. In that case, traffic separation is provided by 802.1q.

5.5.2 Access control

A CEVPL 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/CLE and PE for a CEVPL.

5.5.3 Value added security services

CEVPL allows value added security services such as encryption and/or authentication of customer packets, but does 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

CEVPL is agnostic to 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, as long as the CEs/CLEs have reachability in the PSN. 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. CEVPL supports multi- link access for CE devices, 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/CLE MAY connect to two different SPs via diverse access networks. Resiliency MAY be further enhanced as shown in case (d), where CEs/CLEs, 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) is supported.

Note: In CEVPL, CEs/CLEs process BPDUs from other CEs/CLEs. A customer BPDU is tunneled from CE/CLE to CE/CLE and is transparent to the attached PE.

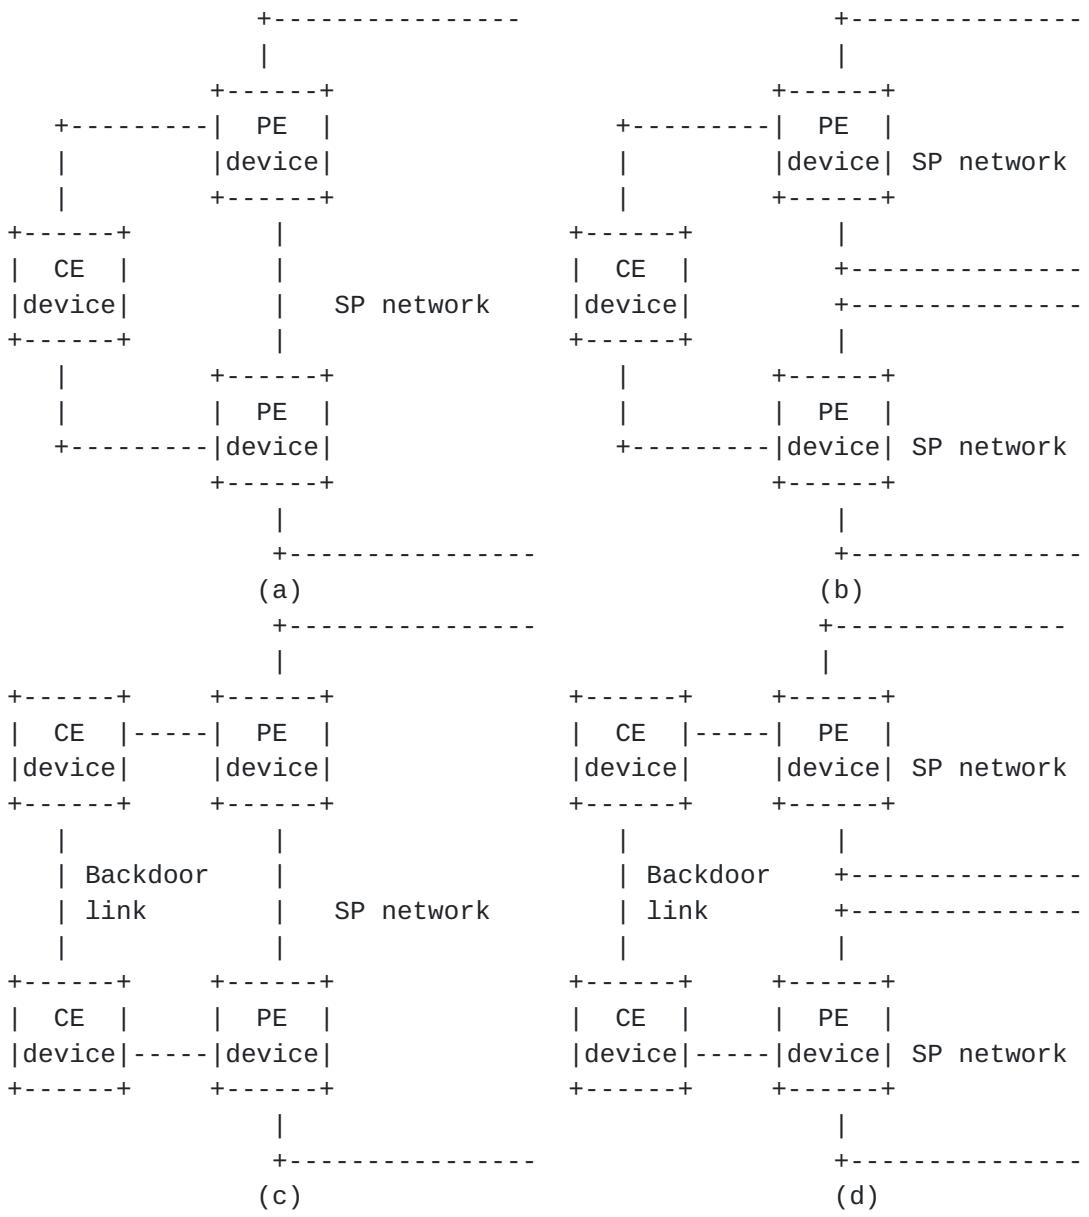
+------+
| CE |
|device|
+------+
+------+
| PE |
|device|
+------+
SP network
(a)
+------+
| CE |
|device|
+------+
+------+
| PE |
|device|
+------+
SP network
(b)
+------+
| CE |
|device|
+------+
Backdoor
link
+------+
| CE |
|device|
+------+
+------+
| PE |
|device|
+------+
SP network
(c)
+------+
| CE |
|device|
+------+
Backdoor
link
+------+
| CE |
|device|
+------+
+------+
| PE |
|device|
+------+
SP network
(d)

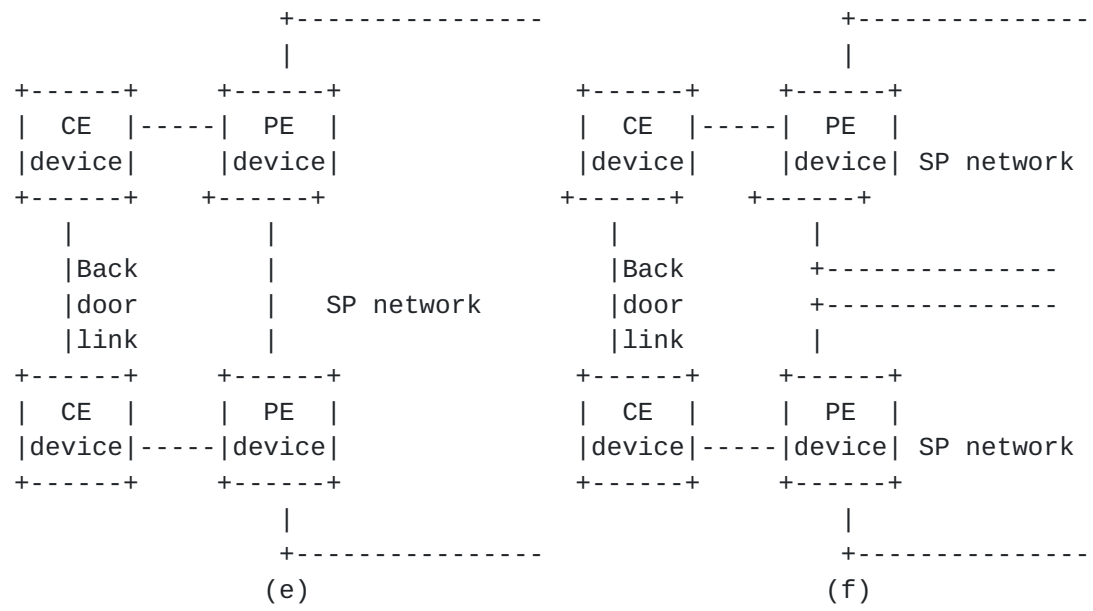

Figure 2 Representative types of access arrangements.

5.7 Customer traffic

5.7.1 Unicast, Unknown Unicast, Multicast, and Broadcast forwarding

CEVPL delivers every packet at least to its intended destination(s) within the scope of the VPL 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. Sequencing in [ETH-LT2Pv3] should be turned on to preserve frame order.

5.7.3 Minimum MTU

CEVPL can support the theoretical MTU of the offered service.

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

CEVPL MAY fragment packets and it is transparent to the customer.

5.7.4 End-point VLAN tag translation

A CEVPL service does not translate customers' VLAN tags, when the customer VLANs are used as service delimiter. It SHOULD be noted that VLAN tag translation (may be supported in other PE-based VPLS) affects the support for multiple spanning trees (i.e., 802.1s) but this is not an issue with CEVPL.

5.7.5 Transparency

A CEVPL service emulates a LAN and appear to customer devices as a bridged LAN. CEVPL does not require any special packet processing by the end users before sending packets to the provider's network.

CEVPL does not require VLAN-ids to be assigned by the SP, hence the issue with VLANs being not transparent if VLAN-ids are assigned by the SP, is not applicable here.

5.8 Support for Layer 2 Control Protocols

CLEs process L2 control protocols as defined in IEEE specifications.

CEVPL ensure that loops be prevented. This can be accomplished through Control protocols such as Spanning Tree (STP). A CEVPL does not require indications from customer Layer 2 control protocols, e.g. STP BPDU snooping, to improve the operation of a VPL, since CLEs participate in STP.

5.9 CE Provisioning

A CEVPL requires configuration of the tunnel endpoint information on CEs/CLES. Provisioning on CEs/CLEs can be minimized or automated using the methods described in "Discovery" of this draft or an LMI method.

6 Service Provider Network Requirements

This section describes the requirements from a service provider perspective. How CEVPL behaves wrt these requirements are described in text preceeded by >>.

6.1 Scalability

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

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.
>> The forwarding states for L2VPN forwarding in CEVPL are only created at CLEs. Hence the number of VPL that can be supported in a provider's network is independent of the number VPWSs or the number of VPLSs. The limitation to the number of VPLSs shall be the amount of traffic that can be handled by the provider's network.

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 CEVPL is capable of supporting number of site interfaces ranging from a few site interfaces to lower end O(10^2) site interfaces per VPLS. It is recommended that a LAN (including emulated LAN) be subnetted when then number of nodes (including CLEs) becomes large. It is not clear how well bridging works in the case of higher end O(10^2) site interfaces per VPLS. Note: In CEVPL, it is not necessary to have a full-meshed of tunnels from CLE to CLE - traffic is forwarded on a Spanning Tree. For e.g. a 5 site interfaces VPLS may require 4 tunnels in a hub and spoke configuration with a few additional tunnels for failover.

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.
>> Since MAC addresses are only stored in CLEs, CEVPLS can support a wide range of number of customer addresses per VPLS from just a few to a very large number such as 1000 or the number of MAC addresses in a typical LAN can be easily supported.

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.
>> With an auto-configuration mechanism, CEVPL can support high values of frequency of configuration setup and change.

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.
>> Since bridging and tunneling are performed on CLEs, CEVPL is transparent to the provider's PEs and Ps and hence can be deployed

and scale well across different domains. No inter-working is required across different domains.

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.
>> CEVPL scales 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.
>> CEVPL scales 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.
>> CEVPL is transparent to PEs and hence is independent of the size of the access network or the number of PEs, I.e, CEVPL scales wrt to these parameters.

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.

>> CEVPL work transparently across different domains, hence a globally unique domain identifier for CEVPL that spans different domains is not required.

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.
>> CE Auto-configuration referred to in CEVPL can be used to minimize 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.
>> CE Auto-configuration referred to in CEVPL has means to prevent unauthorized access and allows authentication of devices.

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.
>> Distribution of configuration information to CE is limited to CEs of a VPLS only. CE Auto-configuration describes ways to minimize the amount of operational information maintained and configured by SPs, for e.g. if a CE is removed, other CEs should be updated accordingly.

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

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.
>> Using a CE Auto-Configuration to dynamically update CEs of L2VPN membership change should allow CEVPL to respond in a "timely manner".

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.
>> CE Auto-Configuration can be used to dynamically update CEs of different L2VPN membership changes. In general, this could be addressed within a provisioning framework, independent of the CEVPL proposal.

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.
>> In CEVPL, the exchange of traffic and forwarding information only occurs among those sites that are authenticated and authorized members of the VPLS.

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.
>> A provider can use existing methods to treat L2VPN traffic with the appropriate priority or DSCP or CoS to differentiate them from non-VPN traffic. Rate-limiting at ingress to provider's network can also be used to prevent traffic from a L2VPN from exceeding its share of bandwidth use.

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.
>> In CEVPL, internal customer addresses or customer BPDUs are transparent to an SP network since trafffic is tunneled from CLE to CLE, and one customer/s CLE does not communicate with another customer's CLE. Hence malformed or maliciously consructed customer traffic shall not affect an SP network.

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.
>> There is no reason to grant access of SP network devices to L2VPN. Hence a provider may perform ingress filtering of traffic from L2VPN towards the SP network addresses (subnet). A CLE only need to be able to reach other CLEs of the same L2VPN.

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 ASes. 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.

>> CEVPL works the same within a domain as it is across different domains. It is not clear if there is any motivation for a CEVPL SP to share management of a CEVPL with another CEVPL SP, since the tunnels from a CE to another CE can be setup across different domains without another SP cooperation or intervention, as long as the different CEs are reachable.

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 CE Auto-Configuration mechanism which allows CEs to query a server accessible from different ASs can be used for this purpose.

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.
>> CEVPL allows one SP to sell a CEVPL at wholesale to another SP, who then resells the emulated LAN to the 2nd SP's customers. In this case, the "CE" would appear as a 802.1d/q bridge in the second SP's network, providing one emulated LAN to all the 2nd SP's customers. If the 2nd SP wishes to provide L2VPN to its customers', it is better for the 2nd SP to provide the L2VPN service over a PSN than over the wholesale L2VPN. The latter incurs additional encapsulation layers and may be susceptible to forwarding loops.

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.
>> Although CEVPL is specified for L2TPv3, it does not preclude the use of other tunneling technologies.

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

>> A CLE support a tunnel setup protocol. A CLE allows static configuration of the tunnel. The multiplexing identifier is conveyed, other information (TLV) may be added to L2TPv3 if necessary.

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

>> Using the functions in "VPL Monitoring" the above statistics to be collected.

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. >> The tunnels and their operations are transparent to customers.

6.11 Support for Access Technologies

>> A CE/CLE is connected to a PE via an access link. The access link MAY span networks of other providers or public networks. Some popular choices of access technologies include Ethernet, ATM (DSL), Frame Relay, MPLS-based virtual circuits etc. The access link, whether direct or virtual, MUST maintain all committed characteristics of the customer traffic, such as QoS, priorities etc. This provider may use existing methods to enforce this. The access technology used is transparent to CEVPL.

>> The CLE connection to the customer network (or CE) is typically Ethernet and is bi-directional in nature. The CLE uses Ethernet frames as the Service Protocol Data Unit (SPDU).

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.
>> CEVPL L2VPN operation is transparent to the provider's network. Hence L2VPN related resources and processing e.g. MAC table, Layer 2 protocol processing has no effect on the provider's network. A CEVPL may share bandwidth in the provider's network with another CEVPL. A provider may use rate-limiting at ingress to limit traffic from a L2VPN and existing methods to manage bandwidth sharing. The may be specified in a separate L2VPN QoS framework.

6.14 Interoperability

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

- To facilitate use of CLE and customer devices (CE)

>> The CLE and CE shall interoperate as specified in IEEE 802.1d

- To implement L2VPN services across two or more interconnected

SP
networks
>> As long as there is reachability from a CLE to another CLE, traffic in a CEVPL can be tunneled in one domain or across different domains. Hence there is no need to devise inter-working means when CEVPL span different SP networks.
- To achieve inter-working or interconnection between customer sites using different L2VPN solutions or different implementations
of the same approach
>> CEVPL may inter-work with other VPLS solutions if the customer facing side of a CLE in a CEVPL is connected to the customer facing side of a PE of other VPLS solutions.

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.

>> There are on-going work to address Ethernet "ping" and "traceroute" functions. This may be adapted for CEVPL in future.

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.
>> CEVPL shall provide standards-based interfaces for the above purpose.

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. >> Using a suitable CE Auto-Configuration, the provisioning of CEVPL can be minimized.

The infrastructure MAY employ manual configuration methods to provide this type of information. >> CEVPL does not preclude the use of manual configuration methods.

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. >> CE Auto-Configuration mechanisms only distribute reachability information to only those sites associated with a CEVPL.

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.

>> In CEVPL, L2VPN control plane processing occurs only in CLE/CE, hence L2VPN control plane resource consumption are restricted to CLEs and customer devices (including CE) only. The number of customer hosts that can be attached to a CEVPL is similar to the MAC table size supported by customer devices (and CLEs).

8.2 Data Plane Requirements

8.2.1 Encapsulation

>> CEVPL utilizes the encapsulation techniques defined by PWE3 ([PWE3-FR]), and does not impose any new requirements on these techniques.

8.2.2 Responsiveness to Congestion

>> CEVPL utilize thes congestion avoidance techniques defined by PWE3 ([PWE3-FR]).

8.2.3 Broadcast Domain

>> CEVPL maintains a separate Broadcast Domain 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. >> In the case of CEVPL, customer VLANs are switched as defined by 802.1q.

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.

>> In CEVPL, typically only one VSI may be required. Customer's VLAN traffic may be switched as defined in IEEE 802.1q.

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

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. >> CEVPL populates the FIB on a per VSI basis, typically only one VSI is required.

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. >> CEVPL security considerations are covered in the respective sections listed above.

IPSec can be also be applied after tunneling Layer-2 traffic to provide additional security. >> CEVPL allows the use of IPSec with L2TPv3.

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