

Network Working Group
Request for Comments: 4665
Category: Informational

W. Augustyn, Ed.
Y. Serbest, Ed.
AT&T
September 2006

Service Requirements for Layer 2
Provider-Provisioned Virtual Private Networks

Status of This Memo

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

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document provides requirements for Layer 2 Provider-Provisioned Virtual Private Networks (L2VPNs). 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 both a customer as well as a service provider perspectives.

Table of Contents

1. Introduction	4
1.1. Scope of This Document	4
1.2. Outline	5
2. Conventions used in this document	5
3. Contributing Authors	5
4. Definitions and Taxonomy	5
4.1. Definitions	5
4.2. Taxonomy of L2VPN Types	6
4.3. VPWS	6
4.4. VPLS	7
5. Service Requirements Common to Customers and Service Providers ..	7
5.1. Scope of emulation	8
5.2. Traffic Types	8
5.3. Topology	8
5.4. Isolated Exchange of Data and Forwarding Information	9
5.5. Security	9
5.5.1. User Data Security	10
5.5.2. Access Control	10
5.6. Addressing	11
5.7. Quality of Service	11
5.7.1. QoS Standards	11
5.7.2. Service Models	11
5.8. Service Level Specifications	12
5.9. Protection and Restoration	12
5.10. CE-to-PE and PE-to-PE Link Requirements	12
5.11. Management	12
5.12. Interoperability	12
5.13. Inter-working	13
6. Customer Requirements	13
6.1. Service Provider Independence	13
6.2. Layer 3 Support	13
6.3. Quality of Service and Traffic Parameters	14
6.4. Service Level Specification	14
6.5. Security	14
6.5.1. Isolation	14
6.5.2. Access Control	14
6.5.3. Value-Added Security Services	15
6.6. Network Access	15
6.6.1. Physical/Link Layer Technology	15
6.6.2. Access Connectivity	15
6.7. Customer Traffic	17
6.7.1. Unicast, Unknown Unicast, Multicast, and Broadcast forwarding	17
6.7.2. Packet Re-ordering	17
6.7.3. Minimum MTU	17
6.7.4. End-point VLAN Tag Translation	18

6.7.5. Transparency	18
6.8. Support for Layer 2 Control Protocols	18
6.9. CE Provisioning	19
7. Service Provider Network Requirements	19
7.1. Scalability	19
7.1.1. Service Provider Capacity Sizing Projections	19
7.1.2. Solution-Specific Metrics	19
7.2. Identifiers	19
7.3. Discovering L2VPN Related Information	19
7.4. Quality of Service (QoS)	20
7.5. Isolation of Traffic and Forwarding Information	20
7.6. Security	21
7.7. Inter-AS/SP L2VPNs	22
7.7.1. Management	22
7.7.2. Bandwidth and QoS Brokering	22
7.8. L2VPN Wholesale	23
7.9. Tunneling Requirements	23
7.10. Support for Access Technologies	23
7.11. Backbone Networks	24
7.12. Network Resource Partitioning and Sharing Between L2VPNs	24
7.13. Interoperability	24
7.14. Testing	25
7.15. Support on Existing PEs	25
8. Service Provider Management Requirements	26
9. Engineering Requirements	26
9.1. Control Plane Requirements	26
9.2. Data Plane Requirements	27
9.2.1. Encapsulation	27
9.2.2. Responsiveness to Congestion	27
9.2.3. Broadcast Domain	27
9.2.4. Virtual Switching Instance	27
9.2.5. MAC Address Learning	27
10. Security Considerations	28
11. Acknowledgements	28
12. References	29
12.1. Normative References	29
12.2. Informative References	29

1. Introduction

This section describes the scope and outline of the document.

1.1. Scope of This Document

This document provides requirements for provider-provisioned Layer 2 Virtual Private Networks (L2VPN). It identifies requirements that MAY apply to one or more individual approaches that a Service Provider (SP) may use for the provisioning of a Layer 2 VPN service. The content of this document makes use of the terminology defined in [RFC4026] and common components for deploying L2VPNs described in [RFC4664].

The technical specifications to provide L2VPN services are outside the scope of this document. The framework document [RFC4664] and several other documents, which explain technical approaches providing L2VPN services, such as [VPLS_LDP], [VPLS_BGP], and [IPLS], 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 [RFC4664].

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 document 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 a normal situation as well as in 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 SLSEs in that it provides means for proper separation between Customer Edge (CE) and Provider Edge (PE), allows proper execution of the SLS offer, and supports a flexible and rich set of capabilities.

This document provides requirements from both the Provider's and the Customer's point of view. It begins with common customer's and service provider's point of view, followed by a customer's

perspective, and concludes with specific needs of an SP. These requirements provide high-level L2VPN features expected by an SP in provisioning L2VPNs, which include SP requirements for security, privacy, manageability, interoperability, and scalability.

1.2. Outline

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

2. Conventions used in this document

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

3. 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

4. Definitions and Taxonomy

4.1. Definitions

The terminology used in this document is defined in [RFC4026]. The L2VPN framework document [RFC4664] 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.

4.2. Taxonomy of L2VPN Types

The requirements distinguish two major L2VPN models, a Virtual Private Wire Service (VPWS), and a Virtual Private LAN Service (VPLS).

The following diagram shows an L2VPN reference model.




Figure 1. L2VPN Reference Model

4.3. VPWS

The PE devices provide a logical interconnect such that a pair of CE devices appears 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 SP network. These tunnels can either be specific to a particular VPWS, or be shared among several services. 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.

4.4. 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 ([RFC4664]). A VPLS can contain a single VLAN or multiple VLANs ([IEEE_802.1Q]). A variation of this service is IPLS ([RFC4664]), 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. In Figure 1, L2VPN A represents a VPLS case.

The details of VPLS reference model, which we summarize here, can be found in [RFC4664]. In VPLS, the PE can be viewed as containing a Virtual Switching Instance (VSI) for each L2VPN 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. The Emulated LAN consists of VPLS Forwarder module (one per PE per VPLS service instance) connected by pseudo wires (PW), where the PWs 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 service instance [RFC4664]. Hence, the VSI terminates PWs for interconnection with other VSIs and also terminates Attachment Circuits (ACs) (see [RFC3985] for definition) for accommodating CEs. A VSI includes the forwarding information base for an L2VPN [RFC4664] which is the set of information regarding how to forward Layer 2 frames received over the AC from the CE to VSIs in other PEs supporting the same L2VPN service (and/or to other ACs), and it contains information regarding how to forward Layer 2 frames received from PWs to ACs. 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.

5. Service Requirements Common to Customers and Service Providers

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

5.1. 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 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.
- VPLS frames can get duplicated if the PW sequencing option isn't turned on. The data frames on the PWs are sent in IP datagrams, and under certain failure scenarios, IP networks can duplicate packets. If the PW data transmission protocol does not ensure sequence of data packets, frames can be duplicated or received out of sequence. If the customer's Bridge Protocol Data Unit (BPDU) frames are sent as data packets, then BPDU frames can be duplicated or mis-sequenced, although this may not create any problems for Real-Time Streaming Protocol (RSTP).
- 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 ([RFC4664]) 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 will probably not be preserved by the IPLS solution.

5.2. Traffic Types

A VPLS MUST support unicast, multicast, and broadcast traffic. Support for efficient replication of broadcast and multicast traffic is highly desirable.

5.3. 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
- Point-to-multipoint, a.k.a. hub and spoke
- Any-to-any, a.k.a. full mesh
- Mixed, a.k.a. partial mesh
- 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 MUST allow multipoint-to-multipoint connectivity even if it is implemented with point-to-point circuits. This requirement does not imply that all traffic characteristics (such as bandwidth, QoS, delay, etc.) necessarily be the same between any two end points of an 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, an L2VPN service SHOULD be capable of crossing multiple administrative boundaries.

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

5.4. Isolated Exchange of Data and Forwarding Information

L2VPN solutions SHALL define means that prevent CEs in an 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 an L2VPN SHOULD not be advertised or discoverable from outside that L2VPN.

5.5. Security

A range of security features MUST be supported by the suite of L2VPN solutions. Each L2VPN solution MUST 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 an L2VPN, ranging from misconfigurations to attacks that can be launched on an L2VPN and can strain network resources such as memory space, forwarding information base table, bandwidth, and CPU processing.

This section lists some potential security hazards that can result due to mis-configurations and/or malicious attacks. 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 PW label assignment
 - o Incorrect signaled VPN parameters (e.g., QoS, MTU, etc.)
- Tampering with data forwarding
 - o Incorrect MAC learning entry
 - o Incorrect PW label
 - o Incorrect AC identifier
 - o Incorrect customer facing encapsulation
 - o Incorrect PW encapsulation
 - o Hijacking PWs using the wrong tunnel
 - o Incorrect tunnel encapsulation

5.5.1. User Data Security

An 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.

5.5.2. Access Control

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

5.6. Addressing

An L2VPN solution MUST support overlapping addresses of different L2VPNs. For instance, customers MUST NOT be prevented from using the same MAC addresses with different L2VPNs. If a service provider uses VLANs as service delimiters, the L2VPN solution MUST ensure that VLAN Ids cannot overlap. If VLANs are not used as service delimiters, L2VPN solutions MAY allow VLAN Ids to overlap.

5.7. Quality of Service

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

5.7.1. QoS Standards

As provided in [RFC3809], an L2VPN SHALL be able to support QoS in one or more of the following already standardized modes:

- Best Effort (support mandatory for all provider-provisioned VPN types)
- Aggregate CE Interface Level QoS (i.e., 'hose' level)
- Site-to-site, or 'pipe' 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 PSN 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.

5.7.2. Service Models

A service provider may desire 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 [RFC3809] and in [RFC4031].

In L2VPN service, both DSCP ([RFC2474]) and 802.1p ([IEEE_802.1D]) fields may be used for this purpose.

5.8. Service Level Specifications

For an L2VPN service, the capabilities for Service Level Specification (SLS) monitoring and reporting stated in [RFC3809] SHOULD be provided.

5.9. Protection and Restoration

The L2VPN service infrastructure SHOULD provide redundant paths to ensure 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.

5.10. CE-to-PE and PE-to-PE Link Requirements

The CE-to-PE links MAY be

- direct physical links (e.g., 100BaseTX, and T1/E1 TDM),
- logical links (e.g., ATM PVC, and RFC2427-encapsulated link),
- transport networks carrying Ethernet,
- a Layer 2 tunnel that goes through a Layer 3 network (e.g., L2TP sessions).

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.).

5.11. Management

Standard interfaces to manage L2VPN services MUST be provided (e.g., standard SNMP MIB Modules). 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, Performance, and Security, as detailed in [ITU_Y.1311.1].

5.12. Interoperability

Multi-vendor interoperability, which corresponds to similar network and service levels among different implementations, at the network element SHOULD be guaranteed. This will likely rely on the completeness of the corresponding standard.

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 use of the L2VPN service.

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

5.13. Inter-working

Inter-working scenarios among different solutions providing L2VPN services are highly desirable. It is possible to have cases that require inter-working or interconnection between customer sites, which span network domains with different L2VPN solutions or different implementations of the same approach. 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.

6. Customer Requirements

This section captures requirements from a customer perspective.

6.1. Service Provider Independence

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

A customer might also start with an L2VPN provided in a single AS with a certain SLS but then ask for an expansion of the service spanning multiple ASes and/or multiple-SPs. In this case, as well as for all kinds of multi-AS and multiple-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.

6.2. Layer 3 Support

With the exception of IPLS, an L2VPN service SHOULD be agnostic to customer's Layer 3 traffic (e.g., IP, IPX, Appletalk) encapsulated within Layer 2 frames.

IPLS MUST allow transport of customer's IPv4 and IPv6 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.

6.3. Quality of Service and Traffic Parameters

QoS is expected to be an important aspect of an 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 PWS (e.g., SONET emulation) to voice, interactive video, and multimedia applications. Hence, best-effort as well as delay and loss sensitive traffic MUST be supported over an L2VPN service. A customer application SHOULD experience consistent QoS independent of the access network technology used at different sites connected to the same L2VPN.

6.4. 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 MIB Modules).

6.5. Security

6.5.1. Isolation

An 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.

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

6.5.2. Access Control

An 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.

6.5.3. Value-Added Security Services

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

L2VPN services MUST NOT interfere with the security mechanisms employed at Layer 3 and higher layers by customers. Layer 2 security mechanisms, such as 802.10b ([IEEE_802.10]) and 802.1AE ([IEEE_802.1AE]), MAY inhibit L2VPN services, when the service delimiting VLAN Ids are encrypted.

6.6. 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.

6.6.1. Physical/Link Layer Technology

L2VPN solutions 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.

6.6.2. Access Connectivity

Various types of physical connectivity scenarios MUST be supported, such as multi-homed sites, backdoor links between customer sites, and devices homed to two or more SP networks. In case of VPLS, IEEE 802.3ad-2000 link aggregation SHOULD be supported. L2VPN solutions SHOULD support at least the types of physical or link-layer connectivity arrangements shown in Figures 2 - 4 (in addition to the case shown in Figure 1). As in Figure 2, a CE can be dual-homed to an SP or to two different SPs via diverse access networks.




Figure 2. Dual-Homed Access of CE Devices

Resiliency of the L2VPN service can be further enhanced as shown in Figure 3, where CE's connected via a "back door" connection, connect to the same SP or to different SPs.




Figure 3. Backdoor Links Between CE Devices

Arbitrary combinations of the above methods, with a few examples shown in Figure 4, SHOULD be supported by any L2VPN solution.




Figure 4. Combination of Dual-Homing and Backdoor Links for CE Devices

6.7. Customer Traffic

6.7.1. 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.

6.7.2. Packet Re-ordering

During normal operation, the queuing and forwarding policies SHOULD preserve packet order for packets with the same QoS parameters.

6.7.3. 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 use IP fragmentation if it presents reassembled packets at VPLS customer edge devices.

6.7.4. End-point VLAN Tag Translation

The L2VPN service MAY support translation of customers' AC 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 AC 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 [IEEE_802.1s]) and can break the proper operation.

6.7.5. Transparency

The L2VPN service is intended to be transparent to Layer 2 customer networks. An 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 Layer 2 frames themselves), the IPLS solution MUST not alter the packets encapsulated inside Layer 2 frames that are transported by the IPLS. However, the IPLS solution is NOT REQUIRED to preserve the Layer 2 header transparently from CE to CE. For example, Source MAC address might 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.

6.8. 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 with a loop-free topology or appropriate forwarding rules. Control protocols such as Spanning Tree (STP) or similar protocols 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.

6.9. CE Provisioning

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

7. Service Provider Network Requirements

This section describes requirements from an SP perspective.

7.1. Scalability

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

7.1.1. Service Provider Capacity Sizing Projections

[RFC3809] lists projections regarding L2VPN sizing and scalability requirements and metrics. The examples are provided in [RFC3809].

7.1.2. Solution-Specific Metrics

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

7.2. Identifiers

An SP domain **MUST** be uniquely identified at least within the set of all interconnected SP networks when supporting an 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 an L2VPN spans multiple SPs (e.g., [RFC2685]). Globally unique identifiers facilitate the support of inter-AS/SP L2VPNs.

7.3. Discovering L2VPN Related Information

Configuration of PE devices (i.e., U-PE and N-PE [RFC4664]) is a significant task for an SP. Solutions **SHOULD** provide methods that dynamically allow L2VPN information to be discovered by the PEs to minimize the configuration steps.

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

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

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

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 an L2VPN solution.

7.4. Quality of Service (QoS)

A significant aspect of a provider-provisioned VPN is support for QoS. An 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 [RFC4031].

7.5. Isolation of Traffic and Forwarding Information

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

An L2VPN solution SHOULD provide a means for meeting provider-provisioned VPN 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.

7.6. Security

The security requirements are stated in Section 6.5. The security requirements provided in [RFC3809] SHOULD be met. The security requirements, except Layer 3 and higher-layer dependent ones, specified in [RFC4031], SHOULD be met.

In addition, an SP network MUST be protected against malformed or maliciously constructed customer traffic. This includes but is 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 SP network SHOULD provide some means of reporting intrusion attempts to the SP, if the intrusion is detected.

When an L2VPN solution operates over a part of the Internet, it should support a configurable option to support one or more of the following standard IPsec methods for securing a customer's VPN traffic:

- Confidentiality, so that only authorized devices can decrypt it
- Integrity, to ensure that the data has not been altered
- Authentication, to ensure that the sender is indeed who he or she claims to be
- Replay attack prevention.

The above functions SHOULD be applicable to "data traffic" of the customer, which includes the traffic exchanged between sites. It SHOULD also be possible to apply these functions to "control traffic", such as routing or signaling protocol exchanges, that is not necessarily perceived by the customer but is nevertheless essential to maintain his or her VPN.

Furthermore, such security methods MUST be configurable between different end-points, such as PE-PE and PE-MTU, only in the case where L2VPN data traffic is carried over IP [RFC4023]. Methods to secure data flows at the native service layer (Layer-2), from CE-CE, CE-MTU and CE-PE, are outside the scope of this document. It is also desirable to configure security on a per-VPN basis.

A VPN solution MAY support one or more encryption schemes, including AES, and 3DES. Encryption, decryption, and key management SHOULD be included in profiles as part of the security management system.

7.7. Inter-AS/SP L2VPNs

All applicable SP requirements, such as traffic and forwarding information isolation, SLSEs, 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.

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

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

An L2VPN solution SHOULD support policies for proper selection of operational parameters coming from different ASes. Similarly, an L2VPN solution SHOULD support policies for selecting information to be disseminated to different ASes.

7.7.1. Management

The general requirements for managing a single AS apply to a concatenation of ASes. 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

7.7.2. Bandwidth and QoS Brokering

When an L2VPN spans multiple ASes, 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 ASes can establish and guarantee uniform QoS in support of a provider-provisioned VPN.

7.8. L2VPN Wholesale

The architecture MUST support the possibility of one SP's 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.

7.9. 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 SP 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.

7.10. Support for Access Technologies

The connectivity between PE and CE devices is referred to as an 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.

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.

7.11. Backbone Networks

Ideally, the backbone interconnecting the SP's 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 SLses for VPN service offerings.

7.12. Network Resource Partitioning and Sharing Between L2VPNs

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

7.13. 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.

7.14. 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 and locating faults.

Examples of testing mechanisms are as follows:

- Checking connectivity between "service-aware" network nodes
- Verifying data plane and control plane integrity
- 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 of customer's data packets, and the test packets MUST not propagate beyond the boundary of the SP network.

7.15. 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 SP and MAY have been designed primarily for Layer 3 services.

8. Service Provider Management Requirements

An SP desires to have a means to view the topology, operational state, and other parameters associated with each customer's L2VPN. Furthermore, the SP 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 MIB Modules), 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 [ITU_Y.1311.1].

9. Engineering Requirements

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

9.1. Control Plane Requirements

An 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 an 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 an L2VPN service grows.

An 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 an L2VPN service.

9.2. Data Plane Requirements

9.2.1. Encapsulation

An L2VPN solution SHOULD utilize the encapsulation techniques defined by PWE3 ([RFC3985]), and SHOULD not impose any new requirements on these techniques.

9.2.2. Responsiveness to Congestion

An L2VPN solution SHOULD utilize the congestion avoidance techniques defined by PWE3 ([RFC3985]).

9.2.3. Broadcast Domain

A separate Broadcast Domain MUST be maintained for each VPLS.

In addition to VPLS Broadcast Domains, an 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.

9.2.4. Virtual Switching Instance

L2VPN PE devices MUST maintain a separate VSI per 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 PE 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.

9.2.5. 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 VSI basis; i.e., in the context of a VPLS and, if supported, in the context of VLANs therein.

10. Security Considerations

Security considerations occur at several levels and dimensions within L2VPNs, as detailed within this document.

The requirements based on security concerns and potential security hazards are detailed in Section 6.5. Further details on security requirements are given from the customer and service provider perspectives in Sections 6.5 and 7.6, respectively. In an analogous manner, further detail on traffic and routing isolation requirements are given from the customer and service provider perspectives in Sections 5.4 and 7.5, respectively. Safeguards to protect network resources such as CPU, memory, and bandwidth are required in Section 7.12.

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

In the case where an L2VPN service is carried over IP [RFC4023], traverses multiple SP networks and passes through an unsecured SP, POP, NAP, or IX, then security mechanisms MUST be employed. These security mechanisms include encryption, authentication, and resource protection, as described in section 5.5. For example, a provider should consider using both authentication and encryption for a tunnel used as part of an L2VPN that traverses another service provider's network.

11. Acknowledgements

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 appreciation 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 VPN requirements document produced by the Layer 3 VPN requirements design team.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

12.2. Informative References

- [VPLS_LDP] Lasserre, M., Kompella, V. "Virtual Private LAN Services over MPLS", Work in Progress.
- [VPLS_BGP] Kompella, K., Rekhter, Y. "Virtual Private LAN Service", Work in Progress.
- [IPLS] Shah, H., et al. "IP-Only LAN Service (IPLS)", Work in Progress.
- [IEEE_802.1Q] IEEE Std 802.1Q-1998, "Virtual Bridged Local Area Networks", 1998
- [RFC2205] Braden, R., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [RFC2685] Fox, B. and B. Gleeson, "Virtual Private Networks Identifier", RFC 2685, September 1999.
- [RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.
- [RFC3308] Calhoun, P., Luo, W., McPherson, D., and K. Peirce, "Layer Two Tunneling Protocol (L2TP) Differentiated Services Extension", RFC 3308, November 2002.

- [RFC3809] Nagarajan, A., "Generic Requirements for Provider Provisioned Virtual Private Networks (PPVPN)", RFC 3809, June 2004.
- [RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005.
- [RFC4031] Carugi, M. and D. McDysan, "Service Requirements for Layer 3 Provider Provisioned Virtual Private Networks (PPVPNs)", RFC 4031, April 2005.
- [RFC4664] Andersson, L. and E. Rosen, "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, September 2006.
- [IEEE_802.1D] ISO/IEC 15802-3: 1998 ANSI/IEEE Std 802.1D, 1998 Edition (Revision and redesignation of ISO/IEC 10038:98), "Part 3: Media Access Control (MAC) Bridges", 1998.
- [ITU_Y.1311.1] Carugi, M. (editor), "Network Based IP VPN over MPLS architecture", Y.1311.1 ITU-T Recommendation, May 2001.
- [IEEE_802.10] IEEE Std 802.10-1998 Edition (Revision IEEE Std 802.10-1992, incorporating IEEE Std 802.10b-1992, 802.10e-1993, 802.10f-1993, 802.10g-1995, and 802.10h-1997), "Standard for Interoperable LAN/MAN Security (SILS)", 1998.
- [IEEE_802.1AE] IEEE 802.1AE/D5.1, "Draft Standard for Local and Metropolitan Area Networks - Media Access Control (MAC) Security", P802.1AE/D5.1, January 19, 2006.
- [IEEE_802.1s] IEEE Std 802.1s-2002, "Virtual Bridged Local Area Networks-Amendment 3: Multiple Spanning Trees", 2002.

Editors' Addresses

Waldemar Augustyn

E-Mail: waldemar@wdmsys.com

Yetik Serbest

AT&T Labs

9505 Arboretum Blvd.

Austin, TX 78759

E-Mail: yetik_serbest@labs.att.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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

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

Intellectual Property

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

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

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

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

