

Network Working Group
Request for Comments: 5251
Category: Standards Track

D. Fedyk, Ed.
Nortel
Y. Rekhter, Ed.
Juniper Networks
D. Papadimitriou
Alcatel-Lucent
R. Rabbat
Google
L. Berger
LabN
July 2008

Layer 1 VPN Basic Mode

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document describes the Basic Mode of Layer 1 VPNs (L1VPNs). L1VPN Basic Mode (L1VPN BM) is a port-based VPN. In L1VPN Basic Mode, the basic unit of service is a Label Switched Path (LSP) between a pair of customer ports within a given VPN port topology. This document defines the operational model using either provisioning or a VPN auto-discovery mechanism, and the signaling extensions for the L1VPN BM.

Table of Contents

1. Introduction	3
1.1. Conventions Used in This Document	4
2. Layer 1 VPN Service	4
3. Addressing, Ports, Links, and Control Channels	7
3.1. Service Provider Realm	7
3.2. Layer 1 Ports and Index	7
3.3. Port and Index Mapping	8
4. Port-Based L1VPN Basic Mode	10
4.1. L1VPN Port Information Tables	11
4.1.1. Local Auto-Discovery Information	12
4.1.2. PE Remote Auto-Discovery Information	12
4.2. CE-to-CE LSP Establishment	14
4.3. Signaling	15
4.3.1. Signaling Procedures	15
4.3.1.1. Shuffling Sessions	16
4.3.1.2. Stitched or Nested Sessions	17
4.3.1.3. Other Signaling	18
4.4. Recovery Procedures	19
5. Security Considerations	20
6. References	21
6.1. Normative References	21
6.2. Informative References	22
7. Acknowledgments	23

1. Introduction

This document describes the Basic Mode of Layer 1 VPNs (L1VPN BM) that is outlined in [RFC4847]. The applicability of Layer 1 VPNs is covered in [RFC5253]. In this document, we consider a layer 1 service provider network that consists of devices that support GMPLS (e.g., Lambda Switch Capable (LSC) devices, optical cross-connects, Synchronous Optical Network / Synchronous Digital Hierarchy (SONET/SDH) cross-connects, etc.). We partition these devices into P (provider) and PE (provider edge) devices. In the context of this document we will refer to the former devices as just "P", and to the latter devices as just "PE". The Ps are connected only to the devices within the provider's network. The PEs are connected to the other devices within the network (either Ps or PEs), as well as to the devices outside of the service provider network. We'll refer to such other devices as Customer Edge (CE) devices. An example of a CE would be a GMPLS-enabled device that is either a router, an SDH cross-connect, or an Ethernet switch.

[RFC4208] defines signaling from the CE to the PE. In [RFC4208], the term "Core Node (CN)" corresponds to P and PE nodes, and the term "Edge Node (EN)" corresponds to CE nodes. We additionally define an "edge Core Node" corresponding to a PE.

Figure 1 illustrates the components in an L1VPN network.




Figure 1: Generalized Layer 1 VPN Reference Model

This document specifies how the L1VPN Basic Mode service can be realized using off-line provisioning or VPN auto-discovery, Generalized Multi-Protocol Label Switching (GMPLS) Signaling [RFC3471], [RFC3473], Routing [RFC4202], and LMP [RFC4204] mechanisms.

L1VPN auto-discovery has similar requirements [RFC4847] to L3VPN auto-discovery. As with L3VPNs, there are protocol choices to be made with auto-discovery. Section 4.1.1 deals with the information that needs to be discovered.

GMPLS routing and signaling are used without extensions within the service provider network to establish and maintain LSC or SONET/SDH (Time Division Multiplexing (TDM)) connections between service provider nodes. This follows the model in [RFC4208].

In L1VPN Basic Mode, the use of LMP facilitates the population of the Port Information Tables of the service provider. Indeed, LMP MAY be used as an option to automate local CE-to-PE link discovery. LMP also MAY augment routing (in enhanced mode) as well as failure handling capabilities.

Consideration of inter-AS and inter-provider L1VPNs requires further analysis beyond the scope of this document.

1.1. 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 [RFC2119].

This document expects that the reader is familiar with the terminology defined and used in [RFC3945], [RFC3471], [RFC3473], [RFC3477], [RFC4201], [RFC4202], [RFC4204], [RFC4208], and the documents referenced therein.

2. Layer 1 VPN Service

Layer 1 VPN services on the interfaces of customer and service provider ports MAY be any of the Layer 1 interfaces supported by GMPLS. Since the mechanisms specified in this document use GMPLS as the signaling mechanism, and since GMPLS applies to both SONET/SDH (TDM) and LSC interfaces, it follows that L1VPN services include (but are not restricted) to LSC- or TDM-based equipment. Note that this document describes Basic Mode L1VPNs and as such requires that:

- (1) GMPLS RSVP-TE is used for signaling both within the service provider (between PEs), as well as between the customer and the service provider (between CE and PE);
- (2) GMPLS Routing on the CE-to-PE link is outside the scope of the Basic Mode of operation of L1VPN; see [RFC4847].

A CE is connected to a PE via one or more links. In the context of this document, a link is a GMPLS Traffic Engineering (TE) link construct, as defined in [RFC4202]. In the context of this document, a TE link is a logical construct that is a member of a VPN, hence introducing the notion of membership to a set of CEs forming the VPN. Interfaces at the end of each link are limited to either TDM or LSC as supported by GMPLS. More specifically, a <CE, PE> link MUST be of the type <X, LSC> or <Y, TDM> where X = PSC (Packet Switch Capable), L2SC (Layer 2 Switch Capable), or TDM and Y = PSC or L2SC. In case the LSP is not terminated by the CE, X MAY also = LSC and Y = TDM. One of the applications of a L1VPN connection is to provide a "virtual private lambda" or similar. In this case, the CE is truly the endpoint in GMPLS terms, and its switching capability on the TE link is not relevant (although its Generalized Protocol Identifier (GPID) MUST be signaled and identical at both CEs, i.e., head-end and tail-end CE).

Likewise, PEs could be any Layer 1 devices that are supported by GMPLS (e.g., optical cross-connects, SDH cross-connects), while CEs MAY be devices at layers 1, 2, and 3, such as an SDH cross-connect, an Ethernet switch, and a router, respectively).

Each TE link MAY consist of one or more channels or sub-channels (e.g., wavelength or wavelength and timeslot, respectively). For the purpose of this discussion, all the channels within a given link MUST have similar shared characteristics (e.g., switching capability, encoding, type, etc.), and MAY be selected independently from the CE's point of view. Channels on different links of a CE need not have the same characteristics.

There MAY be more than one TE link between a given CE-PE pair. A CE MAY be connected to more than one PE (with at least one port per PE). And, conversely, a PE MAY have more than one CE from different VPNs connected to it.

If a CE is connected to a PE via multiple TE links and all the links belong to the same VPN, these links (referred to as component links) MAY be treated as a single TE link using the link bundling constructs [RFC4201].

In order to satisfy the requirements of the LlVPN Basic Mode, it is REQUIRED that for a given CE-PE pair at least one of the links between them has at least one data bearing channel, and at least one control bearing channel, or that there is IP reachability between the CE and the PE that could be used to exchange control information.

A point-to-point link has two end-points: one on the CE and one on the PE. This document refers to the former as "CE port", and to the latter as "PE port". From the above, it follows that a CE is connected to a PE via one or more ports, where each port MAY consist of one or more channels or sub-channels (e.g., wavelength or wavelength and timeslot, respectively), and all the channels within a given port have shared similar characteristics and can be interchanged from the CE's point of view. Similar to the definition of a TE link, in the context of this document, ports are logical constructs that are used to represent a grouping of physical resources that are used to connect a CE to a PE on a per-LlVPN basis.

At any point in time, a given port on a PE is associated with at most one LlVPN, or, to be more precise, with at most one Port Information Table maintained by the PE (although different ports on a given PE could be associated with different LlVPNs, or, to be more precise, with different Port Information Tables). The association of a port with a VPN MAY be defined by provisioning the relationship on the service provider devices. In other words, the context of a VPN membership in Basic Mode is enforced through service provider control.

It is REQUIRED that the interface (that is between the CE and PE and that is used for the purpose of signaling) be capable of initiating/processing GMPLS protocol messages [RFC3473] and of following the procedures described in [RFC4208].

An important goal of LlVPN service is the ability to support what is known as "single-ended provisioning", where the addition of a new port to a given LlVPN involves configuration changes only on the PE that has this port. The extension of this model to the CE is outside the scope of the LlVPN BM.

Another important goal in the LlVPN service is the ability to establish/terminate an LSP between a pair of (existing) ports within an LlVPN from the CE devices without involving configuration changes in any of the service provider's devices. In other words, the VPN topology is under the CE device control (assuming that the underlying PE-to-PE connectivity is provided and allowed by the network).

The mechanisms outlined in this document aim to achieve the above goals. Specifically, as part of the L1VPN service offering, these mechanisms (1) enable the service provider to restrict the set of ports to which a given port could be connected and (2) enable a CE to establish the actual LSP to a subset of ports. Finally, the mechanisms allow arbitrary L1VPN topologies to be supported; including topologies ranging from hub-and-spoke to full mesh point-to-point connections. Only point-to-point links are supported.

The exchange of CE routing or topology information to the service provider is out of scope for L1VPN BM mode.

3. Addressing, Ports, Links, and Control Channels

GMPLS-established conventions for addressing and link numbering are discussed in [RFC3945]. This section builds on those definitions for the L1VPN case where we now have customer and service provider addresses in a Layer 1 context.

3.1. Service Provider Realm

It is REQUIRED that a service provider, or a group of service providers that collectively offer L1VPN service, have a single addressing realm that spans all PE devices involved in providing the L1VPN service. This is necessary to enable GMPLS mechanisms for path establishment and maintenance. We will refer to this realm as the service provider addressing realm. It is further REQUIRED that each L1VPN customer have its own addressing realm with complete freedom to use private or public addresses. We will refer to such realms as the customer addressing realms. Customer addressing realms MAY overlap addresses (i.e., non-unique address) with each other, and MAY also overlap addresses with the service provider realm.

3.2. Layer 1 Ports and Index

Within a given L1VPN, each port on a CE that connects the CE to a PE has an identifier that is unique within that L1VPN (but need not be unique across several L1VPNs). One way to construct such an identifier is to assign each port an address that is unique within a given L1VPN, and use this address as a port identifier. Another way to construct such an identifier is to assign each port on a CE an index that is unique within that CE, assign each CE an address that is unique within a given L1VPN, and then use a tuple <port index, CE address> as a port identifier. Note that both the port and the CE address MAY be an address in several formats. This includes, but is not limited to, IPv4 and IPv6. This identifier is part of the

Customer addressing Realm and is used by the CE device to identify the CE port and the CE remote port for signaling. CEs do not know or understand the service provider realm addresses.

Within a service provider network, each port on a PE that connects that PE to a CE has an identifier that is unique within that network. One way to construct such an identifier is to assign each port on a PE an index that is unique within that PE, assign each PE an IP address that is unique within the service provider addressing realm, and then use a tuple <port index, PE IPv4 address> or <port index, PE IPv6 address> as a port identifier within the service provider network. Another way to construct such an identifier is to assign an IPv4 or IPv6 address that is unique within the service provider addressing realm to each such port. Either way, this IPv4 or IPv6 address is internal to the service provider network and is used for GMPLS signaling within the service provider network.

As a result, each link connecting the CE to the PE is associated with a CE port that has a unique identifier within a given L1VPN, and with a PE port that has a unique identifier within the service provider network. We'll refer to the former as the Customer Port Identifier (CPI), and to the latter as the Provider Port Identifier (PPI).

3.3. Port and Index Mapping

This document requires that each PE port that has a PPI also has an identifier that is unique within the L1VPN customer addressing realm of the L1VPN associated with that port. One way to construct such an identifier is to assign each port an address that is unique within a given L1VPN customer addressing realm, and use this address as a port identifier. Another way to construct such an identifier is to assign each port an index that is unique within a given PE, assign each PE an IP address that is unique within a given L1VPN customer addressing realm (but need not be unique within the service provider network), and then use a tuple <port index, PE IP address> that acts as a port identifier. We'll refer to such port identifier as the VPN-PPI. See Figure 2.

For L1VPNs, it is a requirement that service provider operations are independent of the VPN customer's addressing realm and the service provider addressing realm is hidden from the customer. To achieve this, we define two identifiers at the PE, one customer facing and the other service provider facing. The PE IP address used for the VPN-PPI is independent of the PE IP address used for the PPI (as the two are taken from different address realms, the former from the customer's addressing realm and the latter from a VPN service provider's addressing realm). If for a given port on a PE, the PPI

and the VPN-PPI port identifiers are unnumbered, then they both could use exactly the same port index. This is a mere convenience since the PPI and VPN_PPI can be in any combination of valid formats.




Figure 2: Customer/Provider Port/Index Mapping

Note, as stated earlier, that IP addresses used for the CPIs, PPIs, and VPN-PPIs could be either IPv4 or IPv6 format addresses.

For a given link connecting a CE to a PE:

- If the CPI is an IPv4 address, then the VPN-PPI MUST be an IPv4 address as well since VPN-PPIs are created from the customer address space. If the CPI is a <port index, CPI IPv4 address> tuple, then the VPN-PPI MUST be a <port index, PE IPv4 address> tuple for the same reason.
- If the CPI is an IPv6 address, then the VPN-PPI MUST be an IPv6 address as well since VPN-PPIs are created from the customer address space. If the CPI is a <port index, CPI IPv6 address> tuple, then the VPN-PPI MUST be a <port index, PE IPv6 address> tuple for the same reason.

Note: for a given port on the PE, whether the VPN-PPI of that port is an IP address or a <port index, PE IP address> is independent of the format of the PPI of that port.

This document assumes that assignment of the PPIs is controlled solely by the service provider (without any coordination with the L1VPN customers), while assignment of addresses used by the CPIs and VPN-PPIs is controlled solely by the administrators of L1VPN. This provides maximum flexibility. The L1VPN administrator is the entity that controls the L1VPN service specifics for the L1VPN customers. This function may be owned by the service provider but may also be performed by a third party who has agreements with the service provider. And, of course, each L1VPN customer could assign such addresses on its own, without any coordination with other L1VPNs.

This document also requires IP connectivity between the CE and the PE as specified earlier, which is used for the control channel between CE and PE. This connectivity could be either a single IP hop, which could be realized by either a dedicated link or by an L2 VPN, or an IP private network, such as an L3VPN. The only requirement on this connectivity is an unambiguous way to correlate a particular CE-to-PE control channel with a particular L1VPN. When such a channel is realized by a dedicated link, such a link should be associated with a particular L1VPN. When such channel is realized by an L2VPN, a distinct L2VPN should be associated with an L1VPN. When such channel is realized by an L3VPN, a distinct L3VPN should be associated with an L1VPN.

We'll refer to the CE's address of this channel as the CE Control Channel Address (CE-CC-Addr), and to the PE's address of this channel as the PE Control Channel Address (PE-CC-Addr). Both CE-CC-Addr and PE-CC-Addr are REQUIRED to be unique within the L1VPN they belong to, but are not REQUIRED to be unique across multiple L1VPNs. Control channel addresses are not shared amongst multiple VPNs. Assignment of CE-CC-Addr and PE-CC-Addr is controlled by the administrators of the L1VPN.

Multiple ports on a CE could share the same control channel only as long as all these ports belong to the same L1VPN. Likewise, multiple ports on a PE could share the same control channel only as long as all these ports belong to the same L1VPN.

4. Port-Based L1VPN Basic Mode

An L1VPN is a port-based VPN service where a pair of CEs could be connected through the service provider network via a GMPLS-based LSP within a given VPN port topology. It is precisely this LSP that forms the basic unit of the L1VPN service that the service provider network offers. If a port by which a CE is connected to a PE consists of multiple channels (e.g., multiple wavelengths), the CE could establish LSPs to multiple other CEs in the same VPN over this single port.

In the L1VPN, the service provider does not initiate the creation of an LSP between a pair of CE ports. The LSP establishment is initiated by the CE. However, the SP, by using the mechanisms/toolkit outlined in this document, restricts the set of other CE ports, which may be the remote endpoints of LSPs that have the given port as the local endpoint. Subject to these restrictions, the CE-to-CE connectivity is under the control of the CEs themselves. In other words, the SP allows a L1VPN to have a certain set of

topologies (expressed as a port-to-port connectivity matrix). CE-initiated signaling is used to choose a particular topology from that set.

For each L1VPN that has at least one port on a given PE, the PE maintains a Port Information Table (PIT) associated with that L1VPN. This table contains a list of <CPI, PPI> tuples for all the ports within its L1VPN. In addition, for local PE ports of a given L1VPN, the tuples also include the VPN-PPIs of these ports.

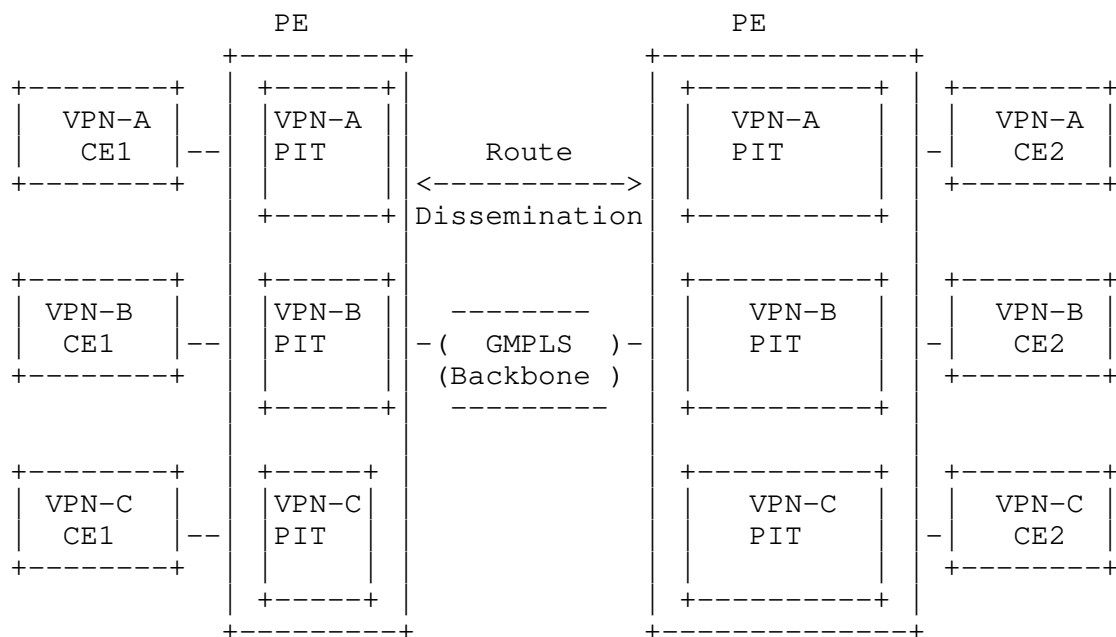


Figure 3: Basic Mode L1VPN Service

4.1. L1VPN Port Information Tables

Figure 3 illustrates three VPNs, VPN-A, VPN-B, and VPN-C, with their associated PITs. A PIT consists of local information as well as remote information. It follows that a PIT on a given PE is populated from two information sources:

1. The information related to the CEs' ports that are attached to the ports local to that PE.
2. The information about the CEs connected to the remote PEs.

A PIT MAY be populated via provisioning or by auto-discovery procedures. When provisioning is used, the entire table MAY be populated by provisioning commands either at a console or by a management system that may have some automation capability. As the network grows, some form of automation is desirable.

For local information between a CE and a PE, a PE MAY leverage LMP to populate the <CPI, VPN-PPI> link information. This local information also needs to be propagated to other PEs that share the same VPN. The mechanisms for this are out of scope for this document, but the information needed to be exchanged is described in Section 4.1.1.

The PIT is by nature VPN-specific. A PE is REQUIRED to maintain a PIT for each L1VPN for which it has member CEs locally attached. A PE does not need to maintain PITs for other L1VPNs. However, the full set of PITs with all L1VPN entries for multiple VPNs MAY also be available to all PEs.

The remote information in the context of a VPN identifier (i.e., the remote CEs of this VPN) MAY also be sent to the local CE belonging to the same VPN. Exchange of this information is outside the scope of this document.

4.1.1. Local Auto-Discovery Information

The information that needs to be discovered on a PE local port is the local CPI and the VPN-PPI.

This information MAY be configured; or, if LMP is used between the CE and PE, LMP MAY be used to exchange this information.

Once a CPI has been discovered, the corresponding VPN-PPI maps in a local context to a VPN identifier and a corresponding PPI. One way to enforce a provider-controlled VPN context is to pre-provision VPN-PPIs with a VPN identifier. Other policy mechanisms to achieve this are outside the scope of this document. In this manner, a relationship of a CPI to a VPN and PPI port can be established when the port is provisioned as belonging to the VPN.

4.1.2. PE Remote Auto-Discovery Information

This section provides the information that is carried by any auto-discovery mechanism, and is used to dynamically populate a PIT. The information provides a single <CPI, PPI> mapping. Each auto-discovery mechanism will define the method(s) by which multiple <CPI, PPI> mappings are communicated, as well as invalidated.

This information should be consistent regardless of the mechanism used to distribute the information [RFC5195], [RFC5252].

The format of encoding a single <PPI, CPI> tuple is:




Figure 4: Auto-Discovery Information

The use and meaning of these fields are as follows:

PPI Length:

A one-octet field whose value indicates the length of the PPI field.

PPI:

A variable-length field that contains the value of the PPI (either an address or <port index, address> tuple). Note, PPI is always encoded consistently within a provider domain so the format of the PPI field is implicit within a given provider network.

CPI AFI:

A two-octet field whose value indicates the address family of the CPI. This value is taken from [RFC1700].

CPI Length:

A one-octet field whose value indicates the length of the CPI field.

CPI:

A variable-length field that contains the CPI value (either an address or <port index, address> tuple).

<PPI, CPI> tuples MUST also be associated with one or more globally unique identifiers associated with a particular VPN. A globally unique identifier can encode a VPN-ID, a route target, or any other globally unique identifier. The globally unique identifiers are under control of network providers. Uniqueness within a service provider administrative domain is sufficient for Basic Mode operation. In the case of multiple provider networks (which is beyond the scope of this document), the globally unique identifier need only be unique and consistent between the those providers. In this document, we specify a generic encoding format for the globally unique identifier common to all the auto-discovery mechanisms. However, each auto-discovery mechanism will define the specific method(s) by which the encoding is distributed and the association with a <PPI, CPI> tuple is made. The encoding of the globally unique identifier associated with the VPN is:

```
+-----+
| L1VPN globally unique identifier (8 octets) |
+-----+
```

Figure 5: Auto-Discovery Globally Unique Identifier Format

4.2. CE-to-CE LSP Establishment

In order to establish an LSP, a CE needs to identify all other CEs in the CE's L1VPN that it wants to connect to. A CE may already have obtained this information through provisioning or through some other schemes (such schemes are outside the scope of this document).

Ports associated with a given CE-to-PE link MAY also have other information, in addition to their CPI and PPI, associated with them that describes characteristics and constraints of the channels within these ports, such as encoding supported by the channels, bandwidth of a channel, total unreserved bandwidth within the port, etc. This information could be further augmented with the information about certain capabilities of the service provider network (e.g., support regeneration section overhead (RSOH), Data Communications Channel (DCC) transparency, arbitrary concatenation, etc.). This information is used to ensure that ports at each end of an LSP have compatible characteristics, and that there are sufficient unallocated resources to establish an LSP between these ports.

It may happen that for a given pair of ports within an L1VPN, each of the CEs connected to these ports would concurrently try to establish an LSP to the other CE. If having a pair of LSPs between a pair of ports is viewed as undesirable, the way to resolve this is to require

the CE with the lower value of the CPI to terminate the LSP originated by the CE. This option could be controlled by configuration on the CE devices.

4.3. Signaling

In L1VPN BM, a CE needs to be configured with the CPIs of other ports. Once a CE is configured with the CPIs of the other ports within the same L1VPN, which we'll refer to as "target ports", the CE uses a subset of GMPLS signaling to request the provider network to establish an LSP to a target port.

For inter-CE connectivity, the CE originates a request that contains the CPI of one of its ports that it wants to use for the LSP, and the CPI of the target port. When the PE attached to the CE that originated the request receives the request, the PE identifies the appropriate PIT, and then uses the information in that PIT to find out the PPI associated with the CPI of the target port carried in the request. The PPI should be sufficient for the PE to establish an LSP. Ultimately, the request reaches the CE associated with the target CPI (note that the request still carries the CPI of the CE that originated the request). If the CE associated with the target CPI accepts the request, the LSP is established.

Note that a CE needs not establish an LSP to every target port that the CE knows about, i.e., it is a local CE policy matter to select a subset of target ports to which that CE will try to establish LSPs.

The procedures for establishing an individual connection between two corresponding CEs is the same as the procedure specified for GMPLS overlay [RFC4208].

4.3.1. Signaling Procedures

When an ingress CE sends an RSVP Path message to an ingress PE, the source IP address in the IP packet that carries the message is set to the appropriate CE-CC-Addr, and the destination IP address in the packet is set to the appropriate PE-CC-Addr. When the ingress PE sends back to the ingress CE the corresponding Resv message, the source IP address in the IP packet that carries the message is set to the PE-CC-Addr, and the destination IP address is set to the CE-CC-Addr.

Likewise, when an egress PE sends an RSVP Path message to an egress CE, the source IP address in the IP packet that carries the message is set to the appropriate PE-CC-Addr, and the destination IP address in the packet is set to the appropriate CE-CC-Addr. When the egress CE sends back to the egress PE the corresponding Resv message, the

source IP address in the IP packet that carries the message is set to the CE-CC-Addr, and the destination IP address is set to the PE-CC-Addr.

In addition to being used for IP addresses in the IP packet that carries RSVP messages between CE and PE, CE-CC-Addr and PE-CC-Addr are also used in the Next/Previous Hop Address field of the IF_ID RSVP_Hop Object that is carried between CEs and PEs.

In the case where a link between CE and PE is a numbered non-bundled link, the CPI and VPN-PPI of that link are used for the Type 1 or 2 TLVs of the IF_ID RSVP_Hop Object that is carried between the CE and PE. In the case where a link between CE and PE is an unnumbered non-bundled link, the CPI and VPN-PPI of that link are used for the IP Address field of the Type 3 TLV. In the case where a link between CE and PE is a bundled link, the CPI and VPN-PPI of that link are used for the IP Address field of the Type 3 TLVs.

Additional processing related to unnumbered links is described in Sections 3 ("Processing the IF_ID RSVP_HOP object") and 4.1 ("Unnumbered Forwarding Adjacencies") of RFC 3477 [RFC3477].

When an ingress CE originates a Path message to establish an LSP from a particular port on that CE to a particular target port, the CE uses the CPI of its port in the Sender Template object. If the CPI of the target port is an IP address, then the CE uses it in the Session object. And if the CPI of the target port is a <port index, IP address> tuple, then the CE uses the IP address part of the tuple in the Session object, and the whole tuple as the Unnumbered Interface ID subobject in the Explicit Route Object (ERO).

There are two options for RSVP-TE sessions. One option is to have a single RSVP-TE session end to end where the addresses of the customer and the provider are swapped at the PE; this is termed shuffling. The other option is when stitching or hierarchy is used to create two LSP sessions, one between the provider PE(s) and another end-to-end session between the CEs.

4.3.1.1. Shuffling Sessions

Shuffling sessions are used when the desire is to have a single LSP originating at the CE and terminating at the far end CE. The customer addresses are shuffled to provider addresses at the ingress PE, and back to customer addresses at the egress PE by using the mapping provided by the PIT.

When the Path message arrives at the ingress PE, the PE selects the PIT associated with the L1VPN, and then uses this PIT to map CPIs carried in the Session and the Sender Template objects to the appropriate PPIs. Once the mapping is done, the ingress PE replaces CPIs with these PPIs. As a result, the Session and the Sender Template objects that are carried in the GMPLS signaling within the service provider network carry PPIs, and not CPIs.

At the egress PE, the reverse mapping operation is performed. The PE extracts the ingress/egress PPI values carried in the Sender Template and Session objects (respectively). The egress PE identifies the appropriate PIT to find the appropriate CPI associated with the PPI of the egress CE. Once the mapping is retrieved, the egress PE replaces the ingress/egress PPI values with the corresponding CPI values. As a result, the Session and the Sender Template objects (included in the GMPLS RSVP-TE Path message sent from the egress PE to the egress CE) carry CPIs, and not PPIs.

Here also, for the GMPLS RSVP-TE Path messages sent from the egress PE to CE, the source IP address (of the IP packet carrying this message) is set to the appropriate PE-CC-Addr, and the destination IP address (of the IP packet carrying this message) is set to the appropriate CE-CC-Addr.

At this point, the CE's view is a single LSP that is point-to-point between the two CEs with a virtual link between the PE nodes: CE-PE(-)PE-CE. The L1VPN PE nodes have a view of the PE-to-PE LSP segment in all its detail. The PEs MAY filter the RSVP-TE signaling, i.e., remove information about the provider topology and replace it with a view of a virtual link.

This translation of addresses and session IDs is termed shuffling and is driven by the L1VPN Port Information Tables (see Section 4). This MUST be performed for all RSVP-TE messages at the PE edges. In this case, there is one CE-to-CE session.

4.3.1.2. Stitched or Nested Sessions

Stitching or Nesting options are dependent on the LSP switching types. If the CE-to-CE and PE-to-PE LSPs are identical in switching type and capacity, the LSP MAY be stitched together and the procedures in [RFC5150] apply. If the CE-to-CE LSPs and the PE-to-PE LSPs are of not the same switching type, or are of different but compatible capacity, the LSPs MAY be Nested and the procedures for [RFC4206] apply. As both Stitched and Nested LSP signaling procedures involve a PE-to-PE session establishment compatible with CE session parameters, they are described together.

When the Path Message arrives at the ingress PE, the PE selects the PIT associated with the L1VPN, and then uses this PIT to map CPIs carried in the Session and the Sender Template objects to the appropriate PPIs. Once the mapping is done, a new PE-to-PE session is established with the parameters compatible with the CE session. Upon successful establishment of the PE-to-PE session, the CE signaling request is sent to the egress PE.

At the ingress PE, when stitching and nesting are used, a PE-to-PE session is established. This could be achieved by several means:

- Associating an already established PE-to-PE LSP or Forwarding Adjacency LSP (FA-LSP) to the destination that meets the requested parameters.
- Establishing a compliant PE-to-PE LSP segment.

At this point, the CE's view is a single LSP that is point-to-point between the two CEs with a virtual node between the PE nodes: CE-PE(-)PE-CE. The L1VPN PE nodes have a view of the PE-to-PE LSP segment in all its detail. The PEs do not have to filter the RSVP-TE signaling by removing information about the provider topology because the PE-to-PE signaling is not visible to the CE nodes.

4.3.1.3 Other Signaling

An ingress PE may receive and potentially reject a Path message that contains an Explicit Route Object and so cause the switched connection setup to fail. However, the ingress PE may accept EROs, which include a sequence of {<ingress PE (strict), egress CE CPI (loose)>}.

- Path message without ERO: when an ingress PE receives a Path message from an ingress CE that contains no ERO, it MUST calculate a route to the destination for the PE-to-PE LSP and include that route in an ERO, before forwarding the Path message. One exception would be if the egress core node were also adjacent to this core node.
- Path message with ERO: when an ingress PE receives a Path message from an ingress CE that contains an ERO (of the form detailed above), the former computes a path to reach the egress PE. It then inserts this path as part of the ERO before forwarding the Path message.

In the case of shuffling, the overlay rules for notification and RRO processing are identical to the User-Network Intercase (UNI) or Overlay Model [RFC4208], which state that Edge PE MAY remove/edit

Provider Notification and RRO objects when passing the messages to the CEs.

4.4. Recovery Procedures

Signaling:

A CE requests a network-protected LSP (i.e., an LSP that is protected from PE-to-PE) by using the technique described in [RFC4873]. Dynamic identification of merge nodes is supported via the LSP Segment Recovery Flags carried in the Protection object (see Section 6.2 of [RFC4873]).

Notification:

A Notify Request object MAY be inserted in Path or Resv messages to indicate the address of a CE that should be notified of an LSP failure. Notifications MAY be requested in both the upstream and downstream directions:

- Upstream notification is indicated via the inclusion of a Notify Request object in the corresponding Path message.
- Downstream notification is indicated via the inclusion of a Notify Request object in the corresponding Resv message.

A PE receiving a message containing a Notify Request object SHOULD store the Notify Node Address in the corresponding RSVP state block. The PE SHOULD also include a Notify Request object in the outgoing Path or Resv message. The outgoing Notify Node Address MAY be updated based on local policy. This means that a PE, upon receipt of this object from the CE, MAY update the value of the Notify Node Address.

If the ingress CE includes a Notify Request object into the Path message, the ingress PE MAY replace the received 'Notify Node Address' by its own selected 'Notify Node Address', and in particular the local TE Router_ID. The Notify Request object MAY be carried in Path or Resv messages (Section 7 of [RFC3473]). The format of the Notify Request object is defined in [RFC3473]. Per Section 4.2.1 of [RFC3473], Notify Node Addresses SHALL be set to either IPv4 or IPv6.

Inclusion of a Notify Request object is used to request the generation of notifications upon failure occurrence but does not guarantee that a Notify message will be generated.

5. Security Considerations

Security for L1VPNs is covered in [RFC4847] and [RFC5253]. In this document, we discuss the security aspects with respect to the control plane.

The association of a particular port with a particular L1VPN (or to be more precise, with a particular PIT) is a configuration operation, generally done manually by the service provider as part of the service provisioning process. Thus, it cannot be altered via signaling between CE and PE. This means that the signaling cannot be used to deliver L1VPN traffic to the wrong customer. The operator should apply appropriate security mechanisms to the management and configuration process, and should consider data plane verification techniques to protect against accidental misconfiguration. The customer may also apply end-to-end (i.e., CE-to-CE) data plane connectivity tests over the L1VPN connection to detect misconnection. Data plane connectivity testing can be performed using the Link Management Protocol (LMP) [RFC4204].

Note that it is also possible to populate the local part of a PIT using auto-discovery through LMP. LMP may be secured as described in [RFC4204]. Signaling between CE and PE is assumed to be over a private link (for example, in-band or in-fiber) or a private network. Use of a private link makes the CE-to-PE connection secure at the same level as the data link described in the previous paragraphs. The use of a private network assumes that entities outside the network cannot spoof or modify control plane communications between CE and PE. Furthermore, all entities in the private network are assumed to be trusted. Thus, no security mechanisms are required by the protocol exchanges described in this document.

However, an operator that is concerned about the security of their private control plane network may use the authentication and integrity functions available in RSVP-TE [RFC3473] or utilize IPsec ([RFC4301], [RFC4302], [RFC4835], [RFC4306], and [RFC2411]) for the point-to-point signaling between PE and CE. See [MPLS-SEC] for a full discussion of the security options available for the GMPLS control plane.

Note further that a private network (e.g., Layer 2 VPN or Layer 3 VPN) might be used to provide control plane connectivity between a PE and more than one CE. In this scenario, it is RECOMMENDED that each L1 VPN customer have its own such private network. Then, the security mechanisms provided by the private network SHOULD be used to ensure security of the control plane communication between a customer and a service provider.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.
- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.
- [RFC4204] Lang, J., Ed., "Link Management Protocol (LMP)", RFC 4204, October 2005.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.
- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, May 2007.
- [RFC5150] Ayyangar, A., Kompella, K., Vasseur, JP., and A. Farrel, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", RFC 5150, February 2008.

6.2. Informative References

- [RFC1700] Reynolds, J. and J. Postel, "Assigned Numbers", RFC 1700, October 1994.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4201] Kompella, K., Rekhter, Y., and L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", RFC 4201, October 2005.
- [RFC4847] Takeda, T., Ed., "Framework and Requirements for Layer 1 Virtual Private Networks", RFC 4847, April 2007.
- [RFC2411] Thayer, R., Doraswamy, N., and R. Glenn, "IP Security Document Roadmap", RFC 2411, November 1998.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.
- [RFC4306] Kaufman, C., Ed., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.
- [RFC4835] Manral, V., "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 4835, April 2007.
- [RFC5195] Ould-Brahim, H., Fedyk, D., and Y. Rekhter, "BGP-Based Auto-Discovery for Layer-1 VPNs", RFC 5195, June 2008.
- [RFC5252] Bryskin, I. and L. Berger, "OSPF-Based Layer 1 VPN Auto-Discovery", RFC 5252, July 2008.
- [RFC5253] Takeda, T., Ed., "Applicability Statement for Layer 1 Virtual Private Network (L1VPN) Basic Mode", RFC 5253, July 2008.
- [MPLS-SEC] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", Work in Progress, February 2008.

7. Acknowledgments

The authors would like to thank Adrian Farrel, Hamid Ould-Brahim, and Tomonori Takeda for their valuable comments.

Sandy Murphy, Charlie Kaufman, Pasi Eronen, Russ Housley, Tim Polk, and Ron Bonica provided input during the IESG review process.

Authors' Addresses

Don Fedyk
Nortel Networks
600 Technology Park
Billerica, MA 01821
Phone: +1 (978) 288 3041
EMail: dwfedyk@nortel.com

Yakov Rekhter
Juniper Networks
1194 N. Mathilda Avenue
Sunnyvale, CA 94089
EMail: yakov@juniper.net

Dimitri Papadimitriou
Alcatel-Lucent
Fr. Wellesplein 1,
B-2018 Antwerpen, Belgium
Phone: +32 3 240-8491
EMail: Dimitri.Papadimitriou@alcatel-lucent.be

Richard Rabbat
Google Inc.
1600 Amphitheatre Pky
Mountain View, CA 95054
EMail: rabbat@alum.mit.edu

Lou Berger
LabN Consulting, LLC
Phone: +1 301-468-9228
EMail: lberger@labn.net

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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, THE IETF TRUST 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.

