

Internet Working Group
Internet Draft

Don Fedyk, Nortel
Yakov Rekhter, Juniper Networks
(Editors)

Expiration Date: September 2006
Standards Track

March 2006

Layer 1 VPN Basic Mode

[draft-fedyk-l1vpn-basic-mode-01.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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

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

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

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

Abstract

This draft describes the basic mode of Layer 1 VPNs (L1VPN BM) that is port based VPNs. In L1VPN BM, the basic unit of service is a Label Switched Path (LSP) between a pair of customer ports within a given VPN port-topology. This draft 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
2. Layer 1 VPN Services.....	4
3. Addressing, Ports, Links and Control Channels.....	6
3.1 Service Provider Realm.....	6
3.2 Layer 1 Ports and Index.....	6
3.3 Port and Index Mapping.....	7
4. Port Based L1VPN Basic Mode.....	9
4.1 L1VPN Port Information Tables.....	10
4.1.1. Local Auto-Discovery Information.....	10
4.1.2. PE Remote Auto-Discovery Information.....	11
4.2 CE to CE LSP Establishment.....	12
4.3 Signaling.....	13
4.3.1 Signaling Procedures.....	13
4.3.1.1 Shuffling Sessions.....	14
4.3.1.2 Stitched or Nested Sessions.....	15
4.3.1.3 Other Signaling.....	15
4.4 Recovery Procedures.....	16
5. Security Considerations.....	17
6. IANA Considerations.....	17
7. Intellectual Property Considerations.....	17
8. References.....	18
8.1 Normative References.....	18
8.2 Informative References.....	18
9. Author's Addresses.....	19
10. Disclaimer of Validity.....	20
11. Full Copyright Statement.....	20

Change from .00 version.

- Many Editorial Changes
- Removed NSAPs no support
- Made Auto-Discovery Generic added discovery information
- Improved LMP Descriptions
- Clarification around Stitching and Shuffling
- Added the L1vpn Globally unique identifier for Signaling
- Clarified recovery Procedures

1. Introduction

This draft describes the basic mode of Layer 1 VPNs (L1VPN BM) that outlined in [[L1VPN-Framework](#)]. In this document, we consider a service provider network that consists of devices that support GMPLS (e.g., Lambda Switch Capable devices, Optical Cross Connects, 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 services provider network. We'll refer to such other devices as Client 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.

The [[GMPLS-OVERLAY](#)] draft is the basis for signaling from the CE to the PE. In the [[GMPLS-OVERLAY](#)] draft the terms Core Node (CN) and Edge Node (EN) correspond to PE and CE respectively.




Figure 1: Generalized Layer 1 VPN Reference Model

This draft specifies how the L1VPN Basic Mode (BM) service can be realized using provisioning or VPN auto-discovery, Generalized Multi-Protocol Label Switching (GMPLS) Signaling [[GMPLS-RSVP-TE](#)], Routing [[GMPLS-Routing](#)] and LMP [[GMPLS-LMP](#)] mechanisms.

The L1VPN auto-discovery has similar requirements to the L3VPNS

auto-discovery [[L3VPN-REQ](#)]. As with L3VPNs there are protocol options to be made with auto-discovery. [Section 4.1.1](#) deals with the information need to be discovered. GMPLS routing and signaling are
Fedyk & Rekhter. [Page 3]

used without extensions within the services provider network to establish and maintain Lambda Switch Capable (LSC) or SONET/SDH (TDM) connections between service provider nodes. This follows the model in [GMPLS-Overlay].

In L1VPN BM, the use of LMP facilitates the population of the services provider port information tables. Indeed, LMP may be used as an option to automate local CE-PE link discovery. LMP also may augment routing in extended mode as well as failure handling capabilities.

2. Layer 1 VPN Services

Layer 1 services on the interfaces of customer and services provider ports could be any of the L1 interfaces supported by GMPLS. Since the mechanisms specified here use GMPLS as the signaling mechanism, and since GMPLS applies to both SONET/SDH (TDM) and Lambda Switch Capable (LSC) interfaces, it results that L1VPN services includes but is not restricted to Lambda Switch Capable or TDM based equipment. Note that this document describes Basic Mode L1 VPNs and as such assumes 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-PE link is outside the scope of the basic mode of operation of L1VPN see [[L1VPN-Framework](#)].

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 [[GMPLS-ROUTING](#)]. 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 type LSC or TDM interfaces that are supported by GMPLS. More specifically a [CE,PE] link must be of the type [X,LSC] or [Y,TDM] where X = PSC, L2SC, or TDM and Y = PSC or L2SC, in case the LSP is not terminated by the CE, X may also = LSC and Y = TDM (the latter case is outside the scope of this document). Likewise, PEs could be any L1 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 we assume that all the channels within a given link have similar shared characteristics (e.g., switching capability, encoding, type, etc_), and can be selected independently from the CE's point of view. Channels on different links of a CE

need not have the same characteristics.

Fedyk & Rekhter.

[Page 4]

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 each 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, for the purpose of this document, these links (referred to as component links) MAY be treated as a single TE link using the link bundling constructs [[LINK-BUNDLING](#)].

A link may have only data bearing channels, or only control bearing channels, or both. For the purpose of this discussion 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 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. In the context of this document we'll refer 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 L1VPN basis.

At any point in time, a given port on a PE is associated with at most one L1VPN, 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 L1VPNs, 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 services provider devices. In other words the context of a VPN membership in Basic mode is enforced through service provider control.

This document assumes that the interface between the CE and PE used for the purpose of signaling is capable of initiating/processing GMPLS protocol messages [[GMPLS-RSVP-TE](#)] and follows the procedures described in [[GMPLS-OVERLAY](#)].

An important goal of L1VPN services is the ability to support what is known as "single ended provisioning", where the addition of a new port to a given L1VPN 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 L1VPN BM.

Another important goal in the L1VPN service is the ability to establish/terminate an LSP between a pair of (existing) ports within Fedyk & Rekhter.

[Page 5]

a L1VPN from the CE devices without involving configuration changes in any of the services provider's devices. In other words, the VPN topology is under the CE device control (provided that the underlying PE-PE connectivity is provided and allowed by the network).

The mechanisms outlined in this document aim at achieving these 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, (2) enable a CE to establish the actual LSP to a subset of ports. Finally, the mechanisms allow arbitrary L1VPN topologies to be supported ranging from hub-and-spoke to full mesh point to point connections. Other more advanced service and topology support such as point to multi point (P2MP) services etc. is for further study.

The L1VPN BM mode does not specify the exchange of CE routing or topology information to the services provider.

3. Addressing, Ports, Links and Control Channels

GMPLS established conventions for Addressing and link numbering are discussed in the [GMPLS-Arch] documents. This section builds on those definitions for the L1VPN case where we now have Customer and services provider addresses in a Layer 1 Context.

3.1 Service Provider Realm

This document assumes 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. This document further assumes that each L1VPN customer has its own addressing realm. We will refer to such realms as the customer addressing realms. Customer addressing realms may overlap with each other, and may also overlap with the service provider addressing 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 Fedyk & Rekhter.

[Page 6]

the CE port and the CE remote port for signaling. CEs do not know or understand the services 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 services 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 assumes 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.

For L1VPNs it is a requirement that services provider operations are independent of the VPN customer's addressing realm and the services provider addressing realm is hidden from the customer. To achieve this we have created two identifiers at the PE, one customer facing and the other services 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 services provider's addressing realm and the latter from a VPN customer'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.

Fedyk & Rekhter.

[Page 7]




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/IPv6 address, then the VPN-PPI has to be an IPv4/IPv6 address as well. And if the CPI is a <port index, CPI IPv4/IPv6 address>, then the VPN-PPI must be a <port index, PE IPv4/IPv6 address>. However, for a given port on PE, whether the VPN-PPI of that port is an IP address or a <port index, PE IPv4/IPv6 address> is independent of whether the PPI of that port is an IP address or a <port index, PE IPv4/IPv6 address>.

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. 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 could assign such addresses on its own, without any coordination with other L1VPNs.

This document also assumes that there is an IP control channel between the CE and the PE. This channel could be either a single IP hop, or a tunnel (GRE or IP-in-IP) or an IP private network, or even an IP VPN for example. 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.

Fedyk & Rekhter.

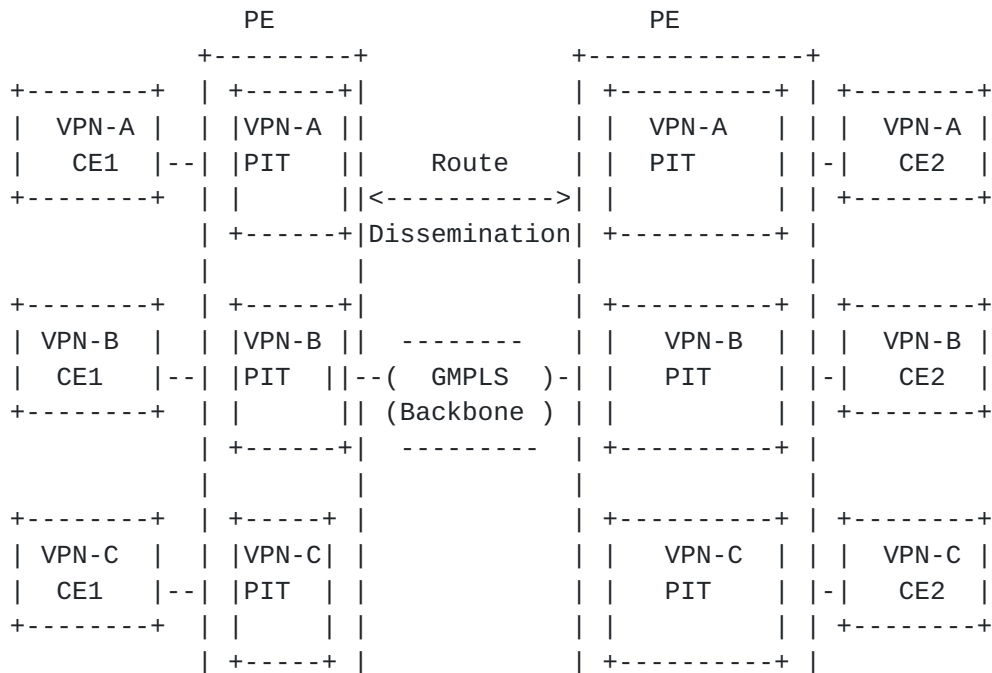
[Page 8]

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 PE ports. This is done rather by the CEs, which attach to the ports. However, the SP, by using the mechanisms/toolkit outlined in this document, restricts the set of other PE 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. A PIT 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

A PIT consists of local information as well as remote information. It follows that PIT on a given PE is populated from two information sources:

1. The information related to the CEs' ports 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 which 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 in that entries for a L1VPN are only required on a PE if that PE locally supports that L1VPN by having CEs belonging to that VPN attached to the PE. 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 ie the remote CEs of this VPN could 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 remote CPI and the VPN-PPI. In many cases if LMP is used between the CE and PE, LMP can exchange this information. Other mechanism may also be used but discussion of these mechanisms is outside the scope of this document.

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-PPI's 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.

The encoding of the auto-discovery information uses BGP address family identifiers (AFIs), and defines a new AFI for L1VPN (to be assigned by the IANA). This information should be consistent regardless of the mechanism use to distribute the information. [[L1VPN-BGP-AD](#)], [[L1VPN-OSPF-AD](#)].

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

```

+-----+
|      Length (1 octet)      |
+-----+
|      PPI Length (1 octet)  |
+-----+
|      PPI (variable)       |
+-----+
|      CPI AFI (2 octets)    |
+-----+
|      CPI (length)         |
+-----+
|      CPI (variable)       |
+-----+

```

Figure 4: Auto-Discovery Information

The use and meaning of these fields are as follows:

Length:

A one octet field whose value indicates the length of the <PPI, CPI> Information tuple in octets.

PPI Length:

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

PPI field:

A variable length field that contains the value of the PPI (either an address or <port index, address> tuple

CPI AFI field:

A two octets field whose value indicates address family of the CPI.

CPI Length:

A once octet field whose value indicates the length of the CPI field.

CPI (variable):

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. 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 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-PE link, in addition to their CPI and PPI may also have other information 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 services provider network (e.g., support RSOH 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.

Fedyk & Rekhter.

[Page 12]

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 request originated by the CE contains the CPI of the port on the CE that CE 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 need not establish an LSP to every target port that CE knows about - it is a local CE matter to select a subset of target ports to which the 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 [[GMPLS-OVERLAY](#)].

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

Fedyk & Rekhter. [Page 13]

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.

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 ERO.

4.3.1.1 Shuffling Sessions

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. The egress PE performs the reverse mapping - it maps PPIs carried in the Session and the Sender Template object into the appropriate CPIs, and then sends the Path message to the egress CE that has the target port. This translation of addresses and session ids is termed shuffling and driven by the L1VPN Port information tables. This must be performed for all RSVP-TE Messages at the PE edges. In this case there is one CE to CE session.

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 out 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 Fedyk & Rekhter.

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 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-PE LSP in all its detail. The PEs may filter the RSVP-TE signaling removing information about the provider topology and replacing it with a view of a virtual link.

4.3.1.2 Stitched or Nested Sessions

Another option which starts out similarly is 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. An additional field is required in the RSVP-TE CE to CE path message. The L1vpn Globally unique identifier is carried to the egress PE in order to identify the Egress PIT table to be used to find CPI and the matching VPN-PPI.

There are a couple of options depending 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 [[GMPLS-STICHING](#)] apply. If the CE to CE LSPs and the PE to PE LSPs are of not the same switching type or of different but compatible capacity the LSPs may be Nested and the procedures for [[GMPLS-HIERARCHY](#)] apply. The Stitched and Nested LSP signaling are analogous procedures. Both result in two sessions: a CE to CE session and a PE to PE session.

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-PE FA-LSP to the destination that meets the requested parameters.
- Establishing a compliant PE-PE LSP.

At this point the CE's view is a single LSP 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-PE LSP in all its detail. The PEs do not have filter the RSVP-TE signaling removing information about the provider topology because the PE-PE signaling is not visible to the CE nodes.

4.3.1.3 Other Signaling

Fedyk & Rekhter.

[Page 15]

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.

An ingress or an egress PE may include a RECORD_ROUTE object and remove/filter the RRO from the received Path message before forwarding it. Further an egress or an ingress PE may remove/filter the RRO from a Resv message before forwarding it. Filtering an RRO consists in editing its content and including only the subobjects based on local or global policy. This allows the ingress/egress CE to be aware of the selected link and labels on the egress/ingress CE side, respectively, for the switched connections constituting this L1VPN.

The exact format of the extensions is TBD in a future revision.

4.4 Recovery Procedures

Signaling:

A CE requests network protected (from PE-to-PE) LSP by using [[SEGMENT-REC](#)] technique. Dynamic identification of merge nodes is supported via the LSP Segment Recovery Flags carried in the Protection object (see Section 6.2 of [[SEGMENT-REC](#)]).

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:

- o) Upstream notification is indicated via the inclusion of a Notify Request Object in the corresponding Path message.

- o) 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 state block. The PE SHOULD also include a Notify Request object in the outgoing Path

Fedyk & Rekhter.

[Page 16]

or Resv message. The outgoing Notify Node Address MAY be updated based on local policy. This means that a PE upon reception of this object from the CE MAY update its value.

If the ingress CE includes a NOTIFY_REQUEST object into the Path message, the ingress PE MAY replace the received 'IPv4 Notify Node Address' by its own selected 'IPv4 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 [[GMPLS-RSVP-TE](#)]). The format of the NOTIFY_REQUEST object is defined in [[GMPLS-RSVP-TE](#)].

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

Since association of a particular port with a particular L1VPN (or to be more precise with a particular PIT) is done by the service provider as part of the service provisioning process (and thus can't be altered via signaling between CE and PE), and since signaling between CE and PE is assumed to be over a private network (and thus can't be spoofed by entities outside the private network), the solution described in this document doesn't require authentication in signaling.

6. IANA Considerations

This document makes no requests for IANA action.

7. Intellectual Property Considerations

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 Fedyk & Rekhter.

[Page 17]

this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

8. References

8.1 Normative References

- [L1VPN-REQ] Ould-Brahim, H., Rekhter, Y., et al., "Service Requirements for Optical Virtual Private Networks", work in progress.
- [GMPLS-OVERLAY] Swallow, G., et al., "Generalized Multiprotocol Label Switching(GMPLS)User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", [RFC 4208](#).
- [GMPLS-STICHING] A. Ayyangar, Ed., J.P. Vasseur, "Label Switched Path Stitching with Generalized MPLS Traffic Engineering", work in progress.
- [GMPLS-HIERARCHY] Kompella, K. and Y. Rekhter, "LSP Hierarchy with Generalized MPLS TE", work in progress.
- [GMPLS-ARCH] Mannie, E. (Editor), "Generalized Multi-protocol Label Switching Architecture", [RFC3945](#), October 2004.
- [L1VPN-Framework] Takeda, T (editor), "Framework and Requirements for Layer 1 Virtual Private Networks", work in progress.
- [SEGMENT-REC] Berger, L., Bryskin, I., Papadimitriou, D. Farrel, A., " GMPLS Based Segment Recovery", work in progress.

8.2 Informative References

- [GMPLS-SIGNALING] Berger, L. (editor), "Generalized MPLS -Signaling Functional Description", January 2003, [RFC3471](#).
- [GMPLS-RSVP-TE] Berger, L. (editor), "Generalized MPLS Signaling - RSVP-TE Extensions", [RFC3473](#), January 2003.
- [GMPLS-ROUTING] Kompella, K., Rekhter, Y., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4202](#), October 2005.
- [GMPLS-HIERARCHY] Kompella, K., Rekhter, Y., "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", [RFC 4206](#), October 2005.

[LINK-BUNDLING] Kompella, K., Rekhter, Y., Berger, L., "Link Bundling in MPLS Traffic Engineering (TE)", [RFC 4201](#), October 2005.

Fedyk & Rekhter.

[Page 18]

[BGP-VPN-AUTODISCOVERY] Ould-Brahim, H., Rosen, E., Rekhter, Y.,
"Using BGP as an Auto-Discovery Mechanism for Layer-3 and Layer-2
VPNs", [draft-ietf-l3vpn-bgpvpn-auto-05.txt](#), work in progress

[GMPLS-LMP] J. Lang (Editor), "Link Management Protocol (LMP)," [RFC 4204](#), October 2005.

[L3VPN-REQ] A. Nagarajan, "Generic Requirements for Provider
Provisioned Virtual Private Networks (PPVPN)" [RFC 3809](#), June
2004.

[L1VPN-BGP-AD] Ould-Brahim, H., Fedyk, D., Rekhter, Y., "BGP-based
Auto-Discovery for L1VPNs", work in progress.

[L1VPN-OSPF-AD] Bryskin, I., Berger, Lou "OSPF Based L1VPN Auto-
Discovery", work in progress.

9. Acknowledgments

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

9. Author's Addresses

Don Fedyk
Nortel Networks
600 Technology Park
Billerica, Massachusetts
01821 U.S.A
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)
Fr. Wellesplein 1,
B-2018 Antwerpen, Belgium
Phone: +32 3 240-8491
Email: dimitri.papadimitriou@alcatel.be

Richard Rabbat
Fujitsu
1240 East Arques Ave, MS 345

Sunnyvale, CA 94085

Email: richard@us.fujitsu.com

Fedyk & Rechter.

[Page 19]

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

10. Disclaimer of Validity

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.

11. 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.