

Network Working Group
Request for Comments: 5253
Category: Informational

T. Takeda, Ed.
NTT
July 2008

Applicability Statement for
Layer 1 Virtual Private Network (L1VPN) Basic Mode

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.

Abstract

This document provides an applicability statement on the use of Generalized Multiprotocol Label Switching (GMPLS) protocols and mechanisms to support Basic Mode Layer 1 Virtual Private Networks (L1VPNs).

L1VPNs provide customer services and connectivity at Layer 1 over Layer 1 networks. The operation of L1VPNs is divided into the Basic Mode and the Enhanced Mode, where the Basic Mode of operation does not feature any exchange of routing information between the Layer 1 network and the customer domain. This document examines how GMPLS protocols can be used to satisfy the requirements of a Basic Mode L1VPN.

Table of Contents

1. Introduction	3
1.1. Terminology	3
2. Basic Mode Overview	3
3. Supported Network Types	4
3.1. Data Plane	4
3.2. Control Plane	4
4. Addressing	5
5. Provider Control of Its Infrastructure	5
5.1. Provisioning Model	5
5.2. PE-to-PE Segment Control	6
5.2.1. Path Computation and Establishment	6
5.2.2. Resource Management	7
5.2.3. Consideration of CE-to-PE Traffic Engineering Information	8
5.3. Connectivity Restriction	8
6. Customer Control of Its L1VPN	8
6.1. Topology Control	8
6.2. Note on Routing	9
7. Scalability and Resiliency	10
7.1. Scalability	10
7.2. Data Plane Resiliency	11
7.3. Control Plane Resiliency	12
8. Security Considerations	12
8.1. Topology Confidentiality	12
8.2. External Control of the Provider Network	13
8.3. Data Plane Security	13
8.4. Control Plane Security	14
9. Manageability Considerations	15
10. References	15
10.1. Normative References	15
10.2. Informative References	16
11. Acknowledgments	17

1. Introduction

This document provides an applicability statement on the use of Generalized Multiprotocol Label Switching (GMPLS) protocols and mechanisms to Basic Mode Layer 1 Virtual Private Networks (L1VPNs) as specified in [RFC4847].

The operation of L1VPNs is divided into the Basic Mode and the Enhanced Mode. The Basic Mode of operation does not feature any exchange of routing information between the Layer 1 network and the customer domain, while the Enhanced Mode of operation features exchange of routing information between the Layer 1 network and the customer domain.

The main GMPLS protocols and mechanisms applicable to the L1VPN Basic Mode are described in [RFC5251], [RFC5195], and [RFC5252], along with several other documents referenced within this document.

Note that discussion in this document is focused on areas where GMPLS protocols and mechanisms are relevant.

1.1. Terminology

The reader is assumed to be familiar with the terminology in [RFC3031], [RFC3209], [RFC3471], [RFC3473], [RFC4202], [RFC4026], and [RFC4847].

2. Basic Mode Overview

As described in [RFC4847], in the Basic Mode service model, there is no routing exchange between the Customer Edge (CE) and the Provider Edge (PE). CE-to-CE L1VPN connections (i.e., the CE-to-CE VPN connection in RFC 4847) are set up by GMPLS signaling between the CE and the PE, and then across the provider network. A L1VPN connection is limited to the connection between CEs belonging to the same L1VPN.

Note that in L1VPNs, routing operates within the provider network. Also note that routing may be used by PEs to exchange information specific to the L1VPNs supported by the provider network (e.g., membership information).

In the L1VPN Basic Mode, the provider network is completely under the control of the provider. This includes the PE-to-PE segment of the CE-to-CE L1VPN connection that is controlled and computed by the provider (PE-to-PE segment control). On the other hand, the L1VPN itself, constructed from a set of CEs and the L1VPN connections provided by the provider, is under the control of each customer. This control includes that a customer can request between which CEs a

connection is to be established (topology control). Note that a customer may outsource the management of its L1VPN to a third party, including to the provider itself. There is a confidentiality requirement between the provider and each customer.

[RFC5251], which extends [RFC4208], specifies GMPLS signaling to establish CE-to-CE L1VPN connections.

[RFC5195] and [RFC5252] specify alternative mechanisms to exchange L1VPN membership information between PEs, based on BGP and OSPF, respectively.

3. Supported Network Types

3.1. Data Plane

The provider network can be constructed from any type of Layer 1 switches, such as Time Division Multiplexing (TDM) switches, Optical Cross-Connects (OXCs), or Photonic Cross-Connects (PXCs). Furthermore, a PE may be an Ethernet Private Line (EPL) type of device, that maps Ethernet frames onto Layer 1 connections (by means of Ethernet over TDM, etc.). The provider network may be constructed from switches providing a single switching granularity (e.g., only VC3 switches), or from switches providing multiple switching granularities (e.g., from VC3/VC4 switches, or from VC3 switches and OXCs). The provider network may provide a single type of L1VPN connection (e.g., VC3 connections only), or multiple types of connection (e.g., VC3/VC4 connections, or VC3 connections and wavelength connections).

A CE does not have to have the capability to switch at Layer 1, but it must be capable of receiving a Layer 1 signal and either switching it or terminating it with adaptation.

As described in [RFC4847] and [RFC5251], a CE and a PE are connected by one or more links. A CE may also be connected to more than one PE, and a PE may have more than one CE connected to it.

A CE may belong to a single L1VPN, or to multiple L1VPNs, and a PE may support one or more L1VPNs through a single CE or through multiple CEs.

3.2. Control Plane

The provider network is controlled by GMPLS. L1VPN Basic Mode provider networks are limited to a single AS within the scope of this document. Multi-AS Basic Mode L1VPNs are for future study.

As described in [RFC4847] and [RFC5251], a CE and a PE need to be connected by at least one control channel. It is necessary to disambiguate control plane messages exchanged between a CE and a PE if the CE-to-PE relationship is applicable to more than one L1VPN. This makes it possible to determine to which L1VPN such control plane messages apply. Such disambiguation can be achieved by allocating a separate control channel to each L1VPN (either using a separate physical channel, a separate logical channel such as an IP tunnel, or using separate addressing).

GMPLS allows any type of control channel to be used, as long as there is IP level reachability. In the L1VPN context, instantiation of a control channel between a CE and a PE may differ depending on security requirements, etc. This is discussed in Section 8.

4. Addressing

As described in [RFC5251], the L1VPN Basic Mode allows that customer addressing realms overlap with each other, and also overlap with the service provider addressing realm. That is, a customer network may reuse addresses used by the provider network, and may reuse addresses used in another customer network supported by the same provider network. This is the same as in any other VPN model.

In addition, the L1VPN Basic Mode allows CE-to-PE control channel addressing realms to overlap. That is, a CE-to-PE control channel address (CE's address of this control channel and PE's address of this control channel) is unique within the L1VPN that the CE-to-PE control channel belongs to, but not necessarily unique across multiple L1VPNs.

Furthermore, once a L1VPN connection has been established, the L1VPN Basic Mode does not enforce any restriction on address assignment for this L1VPN connection (treated as a link) for customer network operation (e.g., IP network, MPLS network).

5. Provider Control of Its Infrastructure

5.1. Provisioning Model

As described in [RFC5251], for each L1VPN that has at least one customer-facing port on a given PE, the PE maintains a Port Information Table (PIT) associated with that L1VPN. A PIT provides a cross-reference between Customer Port Indices (CPIs) and Provider Port Indices (PPIs) and contains a list of <CPI, PPI> tuples for all the ports within the L1VPN. In addition, for local PE ports of a given L1VPN, the PE retains an identifier known as the VPN-PPI, and this is stored in the PIT with the <CPI, PPI> tuples.

When a new CE belonging to one or more L1VPNs is added to a PE, PIT entries associated to those L1VPNs need to be configured on the PE. Section 4 of [RFC5251] specifies such procedures:

- If no PIT exists for the L1VPN on the PE, a new PIT is created by the provider and associated with the VPN identifier.
- The PIT (new or pre-existing) is updated to include information related to the newly added CE. The VPN-PPI, PPI, and CPI are installed in the PIT. Note that the PPI is well-known by the PE, but the CPI must be discovered either through manual configuration or automatically by mechanisms such as the Link Management Protocol (LMP) [RFC4204]. In addition, a CE-to-PE control channel needs to be configured.
- The updated PIT information needs to be configured in the PITs on the remote PE associated with the L1VPN. For such purposes, manual configuration or some sort of auto-discovery mechanisms can be used. [RFC5195] and [RFC5252] specify alternative auto-discovery mechanisms.
- In addition, remote PIT information associated with the L1VPN needs to be configured on this PE if the PIT has been newly created. Again, this can be achieved through manual configuration or through auto-discovery; see [RFC5195] and [RFC5252].

When L1VPN membership of an existing CE changes, or when a CE is removed from a PE, similar procedures need to be applied to update the local and remote PITs.

5.2. PE-to-PE Segment Control

In the L1VPN Basic Mode, a PE-to-PE segment of a CE-to-CE L1VPN connection is completely under the control of the provider network.

5.2.1. Path Computation and Establishment

A PE-to-PE segment of a CE-to-CE L1VPN connection may be established based on various policies. Those policies can be applied per L1VPN or per L1VPN connection. The policy is configured by the provider, possibly based on the contracts with each customer.

Examples of PE-to-PE segment connection establishment policies supported in the L1VPN Basic Mode are as follows.

- Policy 1: On-demand establishment, on-demand path computation
- Policy 2: On-demand establishment, pre-computed path
- Policy 3: Pre-establishment, pre-computed path

In each policy, the PE-to-PE path may be computed by the local PE or by a path computation entity outside of the local PE (e.g., a Path Computation Element (PCE) [RFC4655] or a management system).

In policies 2 and 3, pre-computation of paths (and pre-establishment if applicable) can be done at the network planning phase, or just before signaling (e.g., triggered by an off-line customer request). As the result of pre-computation (and pre-establishment), there could be multiple PE-to-PE segments for a specific pair of PEs. When a PE receives a Path message from a CE for a L1VPN connection, a PE needs to determine which PE-to-PE segment to use. In such cases, the provider may want to control:

- Which L1VPN uses which PE-to-PE L1VPN segment.
- Which CE-to-CE L1VPN connection uses which PE-to-PE L1VPN segment.

The former requires mapping between the PIT and the PE-to-PE segment. The latter requires some more sophisticated mapping method, for example:

- Mapping between individual PIT entries and PE-to-PE segments.
- Use of a Path Key ID [CONF-SEG] supplied by the provider to the CE, and signaled by the CE as part of the L1VPN connection request.

The L1VPN Basic Mode does not preclude usage of other methods, if applicable.

In policy 3, stitching or nesting is necessary in order to map the CE-to-CE L1VPN connection to a pre-established PE-to-PE segment.

5.2.2. Resource Management

The provider network may operate resource management based on various policies. These policies can be applied per L1VPN or per L1VPN connection. The policy is configured by the provider, possibly based on the contracts with each customer.

For example, a provider may choose to partition the resources of the provider network for limited use by different L1VPNs or customers. Such a function might be achieved within the scope of the Basic Mode using resource affinities [RFC3209], but the details of per-L1VPN resource models (especially in terms of CE-to-PE routing) are considered as part of the Enhanced Mode.

5.2.3. Consideration of CE-to-PE Traffic Engineering Information

[RFC5252] and [BGP-TE] allow CE-to-PE Traffic Engineering (TE) link information to be injected into the provider network, and in particular to be exchanged between PEs. This may be helpful for the ingress PE to prevent connection setup failure due to lack of resources or incompatible switching capabilities on remote CE-to-PE TE links.

Furthermore, the L1VPN Basic Mode allows a remote CE to be reached through more than one TE link connected to the same PE (single-homed) or to different PEs (dual-homed). In such cases, to facilitate route choice, the ingress CE needs to initiate signaling by specifying the egress CE's router ID, not the egress CPI, in the Session Object and the Explicit Route Object (ERO), if present, so as to not constrain the choice of route within the provider network. Therefore, the CE's router ID needs to be configured in the PITs.

Note that, as described in Section 7.2, consideration of the full feature set enabled by dual-homing (such as resiliency) is out of scope of the L1VPN Basic Mode.

5.3. Connectivity Restriction

The L1VPN Basic Mode allows restricting connection establishment between CEs belonging to the same L1VPN for policy reasons (including L1VPN security). Since the PIT at each PE is associated with a L1VPN, this function can be easily supported. The restriction can be applied at the ingress PE or at the egress PE according to the applicable restriction policy, but note that applying the policy at the egress may waste signaling effort within the network as L1VPN connections are pointlessly attempted.

In addition, the L1VPN Basic Mode does not restrict use of any advanced admission control based on various policies.

6. Customer Control of Its L1VPN

6.1. Topology Control

In the L1VPN Basic Mode, L1VPN connection topology is controlled by the customer. That is, a customer can request setup/deletion/modification of L1VPN connections using signaling mechanisms specified in [RFC5251].

Also note that if there are multiple CE-to-PE TE links (single-homed or multi-homed), a customer can specify which CE-to-PE TE link to use to support any L1VPN connection. Alternatively, a customer may let the provider choose the CE-to-PE TE link at the egress side, as described in Section 5.2.3.

6.2. Note on Routing

A CE needs to obtain the remote CPI to which it wishes to request a connection. Since, in the L1VPN Basic Mode, there is no routing information exchange between a CE and a PE, there is no dynamic mechanism supported as part of the Basic Mode L1VPN service, and the knowledge of remote CPIs must be acquired in a L1VPN-specific way, perhaps through configuration or through a directory server.

If an L1VPN is used by a customer to operate a private IP network, the customer may wish to form routing adjacencies over the CE-to-CE L1VPN connections. The L1VPN Basic Mode does not enforce any restriction on such operation by a customer, and the use made of the L1VPN connections is transparent to the provider network.

Furthermore, if an L1VPN is used by a customer to operate a private Multiprotocol Label Switching (MPLS) or GMPLS network, the customer may wish to treat a L1VPN connection as a TE link, and this requires a CE-to-CE control channel. Note that a Forwarding Adjacency [RFC4206] cannot be formed from the CE-to-CE L1VPN connection in the Basic Mode because there is no routing exchange between CE and PE. That is, the customer network and the provider network do not share a routing instance, and the customer control channel cannot be carried within the provider control plane. But where the CE provides suitable adaptation (for example, where the customer network is a packet-switched MPLS or GMPLS network), the customer control channel may be in-band and a routing adjacency may be formed between the CEs using the L1VPN connection. Otherwise, CE-to-CE control plane connectivity may form part of the L1VPN service provided to the customer by the provider and may be achieved within the L1VPN connection (for example, through the use of overhead bytes) or through a dedicated control channel connection or tunnel. The options available are discussed further in Section 10.2 of [RFC4847].

7. Scalability and Resiliency

7.1. Scalability

There are several factors that impact scalability.

- o Number of L1VPNs (PITs) configured on each PE

With the increase of this number, information to be maintained on the PE increases. Theoretically, the upper limit of the number of L1VPNs supported in a provider network is governed by how the ID associated with a L1VPN is allocated, and the number of PITs configured on each PE is limited by this number. However, implementations may impose arbitrary limits on the number of PITs supported by any one PE.

- o Number of CE-to-PE TE links for each L1VPN

With the increase of this number, information to be maintained in each PIT increases. When auto-discovery mechanisms are used, the amount of information that an auto-discovery mechanism can support may restrict this number.

Note that [RFC5252] floods membership information not only among PEs, but also to all P nodes. This may lead to scalability concerns, compared to [RFC5195], which distributes membership information only among PEs. Alternatively, a separate instance of the OSPF protocol can be used just between PEs for distributing membership information. In such a case, Ps do not participate in flooding.

Note that in the L1VPN Basic Mode, a PE needs to obtain only CE-to-PE TE link information, and not customer routing information, which is quite different from the mode of operation of an L3VPN. Therefore, the scalability concern is considered to be less problematic.

- o Number of L1VPN connections

With the increase of this number, information to be maintained on each PE/P increases. When stitching or nesting is used, the state to be maintained at each PE increases compared to when connectivity is achieved without stitching or nesting.

However, in a Layer 1 core, this number is always bounded by the available physical resource because each LSP uses a separate label which is directly bound to a physical, switchable resource

(timeslot, lambda, or fiber). Thus, it can be safely assumed that the PEs/Ps can comfortably handle the number of LSPs that they may be called on to switch for a L1VPN.

7.2. Data Plane Resiliency

The L1VPN Basic Mode supports the following data plane recovery techniques [RFC5251].

o PE-to-PE segment recovery

The CE indicates to protect the PE-to-PE segment by including Protection Object specified in [RFC4873] in the Path message and setting Segment Recovery Flags. The CE may also indicate the branch and merge nodes by including a Secondary Explicit Route Object.

Depending on the signaling mechanisms used within the provider network, details on how to protect the PE-to-PE segment may differ as follows.

- If LSP stitching or LSP hierarchy are used to provision the PE-to-PE segment, then the PE-to-PE LSP may be protected using end-to-end recovery within the provider network.
- If the CE-to-CE L1VPN connection is a single end-to-end LSP (including if session shuffling is used), then the PE-to-PE LSP segment may be protected using segment protection [RFC4873].

o CE-to-PE recovery and PE-to-PE recovery via link protection

The CE indicates to protect ingress and egress CE-to-PE links as well as links within the provider network by including the Protection Object specified in [RFC3473] and setting Link Flags in the Path message.

- The ingress and egress CE-to-PE link may be protected at a lower layer.

Depending on the signaling mechanisms used within the provider network, details on how to protect links within the provider network may differ as follows.

- If the PE-to-PE segment is provided as a single TE link (stitching or hierarchy) so that the provider network can perform simple PE-to-PE routing, then the TE link may offer link-level protection through the instantiation of multiple PE-to-PE LSPs.

- The PE-to-PE segment may be provisioned using only link-protected links within the core network.

Note that it is not possible to protect only the CE-to-PE portion or the PE-to-PE portion by link protection because the CE-to-CE signaling request asks for a certain level of link protection on all links used by the LSP. Also, it is not possible to protect the CE-to-PE portion by link recovery and the PE-to-PE portion by segment recovery at the same time.

CE-to-CE recovery through the use of connections from one CE to diverse PEs (i.e., dual-homing) is not supported in the L1VPN Basic Mode.

7.3. Control Plane Resiliency

The L1VPN Basic Mode allows use of GMPLS control plane resiliency mechanisms. This includes, but is not limited to, control channel management in LMP [RFC4204] and fault handling in RSVP-TE ([RFC3473] and [RFC5063]) between a CE and a PE, as well as within the provider network.

8. Security Considerations

Security considerations are described in [RFC4847], and this section describes how these considerations are addressed in the L1VPN Basic Mode.

Additional discussion of GMPLS security can be found in [GMPLS-SEC].

8.1. Topology Confidentiality

As specified in [RFC5251], a provider's topology confidentiality is preserved by the Basic Mode. Since there is no routing exchange between PE and CE, the customer network can gather no information about the provider network. Further, as described in Section 4 of [RFC4208], a PE may filter the information present in a Record Route Object (RRO) that is signaled from the provider network to the customer network. In addition, as described in Section 5 of [RFC4208] and Section 4.4 of [RFC5251], when a Notify message is sent to a CE, it is possible to hide the provider internal address. This is accomplished by a PE updating the Notify Node Address with its own address when the PE receives a NOTIFY_REQUEST object from the CE.

Even in the case of pre-computed and/or pre-signaled PE-to-PE segments, provider topology confidentiality may be preserved through the use of path key IDs [CONF-SEG].

The customer's topology confidentiality cannot be completely hidden from the provider network. At the least, the provider network will know about the addresses and locations of CEs. Other customer topology information will remain hidden from the provider in the Basic Mode, although care may be needed to protect the customer control channel as described in Section 8.4.

The provider network is responsible for maintaining confidentiality of topology information between customers and across L1VPNs. Since there is no distribution of routing information from PE to CE in the Basic Mode, there is no mechanism by which the provider could accidentally, or deliberately but automatically, distribute this information.

8.2. External Control of the Provider Network

The provider network is protected from direct control from within customer networks through policy and through filtering of signaling messages.

There is a service-based policy installed at each PE that directs how a PE should react to a L1VPN connection request received from any CE. Each CE is configured at the PE (or through a policy server) for its membership of a L1VPN, and so CEs cannot dynamically bind to a PE or join a L1VPN. With this configuration comes the policy that tells the PE how to react to a L1VPN connection request (for example, whether to allow dynamic establishment of PE-to-PE connections). Thus, the provider network is protected against spurious L1VPN connection requests and can charge for all L1VPN connections according to the service agreement with the customers. Hence, the provider network is substantially protected against denial-of-service (DoS) attacks.

At the same time, if a Path message from a CE contains an Explicit Route Object (ERO) specifying the route within provider network, it is rejected by the PE. Thus, the customer network has no control over the resources in the provider network.

8.3. Data Plane Security

As described in [RFC4847], at Layer 1, data plane information is normally assumed to be secure once connections are established since the optical signals themselves are normally considered to be hard to intercept or modify, and it is considered difficult to insert data into an optical stream. The very use of an optical signal may be considered to provide confidentiality and integrity to the payload data. Furthermore, as indicated in [RFC4847], L1VPN connections are

each dedicated to a specific L1VPN by which an additional element of security for the payload data is provided.

Misconnection remains a security vulnerability for user data. If a L1VPN connection were to be misconnected to the wrong destination, user data would be delivered to the wrong consumers. In order to protect against mis-delivery, each L1VPN connection is restricted to use only within a single L1VPN. That is, a L1VPN connection does not connect CEs that are in different L1VPNs. In order to realize this, the identity of CEs is assured as part of the service contract. And upon receipt of a request for connection setup, the provider network assures that the connection is requested between CEs belonging to the same L1VPN. This is achieved as described in Section 5.3.

Furthermore, users with greater sensitivity to the security of their payload data should apply appropriate security measures within their own network layer. For example, a customer exchanging IP traffic over a L1VPN connection may choose to use IPsec to secure that traffic (i.e., to operate IPsec on the CE-to-CE exchange of IP traffic).

8.4 Control Plane Security

There are two aspects for control plane security.

First, the entity connected over a CE-to-PE control channel must be identified. This is done when a new CE is added as part of the service contract and the necessary control channel is established. This identification can use authentication procedures available in RSVP-TE [RFC3209]. That is, control plane entities are identified within the core protocols used for signaling, but are not authenticated unless the authentication procedures of [RFC3209] are used.

Second, it must be possible to secure communication over a CE-to-PE control channel. If a communication channel between the customer and the provider (control channel, management interface) is physically separate per customer, the communication channel could be considered as secure. However, when the communication channel is physically shared among customers, security mechanisms need to be available and should be enforced. RSVP-TE [RFC3209] provides for tamper-protection of signaling message exchanges through the optional Integrity object. IPsec tunnels can be used to carry the control plane messages to further ensure the integrity of the signaling messages.

Note that even in the case of physically separate communication channels, customers may wish to apply security mechanisms, such as

IPsec, to assure higher security, and such mechanisms must be available.

Furthermore, the provider network needs mechanisms to detect DoS attacks and to protect against them reactively and proactively. In the Basic Mode, this relies on management systems. For example, management systems collect and analyze statistics on signaling requests from CEs, and protect against malicious behaviors where necessary.

Lastly, it should be noted that customer control plane traffic carried over the provider network between CEs needs to be protected. Such protection is normally the responsibility of the customer network and can use the security mechanisms of the customer signaling and routing protocols (for example, RSVP-TE [RFC3209]) or may use IPsec tunnels between CEs. CE-to-CE control plane security may form part of the data plane protection where the control plane traffic is carried in-band in the L1VPN connection. Where the CE-to-CE control plane connectivity is provided as an explicit part of the L1VPN service by the provider, control plane security should form part of the service agreement between the provider and customer.

9. Manageability Considerations

Manageability considerations are described in [RFC4847]. In the L1VPN Basic Mode, we rely on management systems for various aspects of the different service functions, such as fault management, configuration and policy management, accounting management, performance management, and security management (as described in Section 8).

In order to support various management functionalities, MIB modules need to be supported. In particular, the GMPLS TE MIB (GMPLS-TE-STD-MIB) [RFC4802] can be used for GMPLS-based traffic engineering configuration and management, while the TE Link MIB (TE-LINK-STD-MIB) [RFC4220] can be used for configuration and management of TE links.

10. References

10.1. Normative References

- [RFC3031] Rosen, E., Viswanathan, A. and R. Callon, "Multiprotocol label switching Architecture", RFC 3031, January 2001.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

- [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.
- [RFC4026] Anderssson, L. and Madsen, T., "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.
- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, 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.
- [RFC4847] Takeda, T., Ed., "Framework and Requirements for Layer 1 Virtual Private Networks", RFC 4847, April 2007.
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, May 2007.
- [RFC5195] Ould-Brahim, H., Fedyk, D., and Y. Rekhter, "BGP-Based Auto-Discovery for Layer-1 VPNs", RFC 5195, June 2008.
- [RFC5251] Fedyk, D., Ed., Rekhter, Y., Ed., Papadimitriou, D., Rabbat, R., and L. Berger, "Layer 1 VPN Basic Mode", RFC 5251, July 2008.
- [RFC5252] Bryskin, I. and Berger, L., "OSPF-Based Layer 1 VPN Auto-Discovery", RFC 5252, July 2008.

10.2. Informative References

- [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.

- [RFC4220] Dubuc, M., Nadeau, T., and J. Lang, "Traffic Engineering Link Management Information Base", RFC 4220, November 2005.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4802] Nadeau, T., Ed., and A. Farrel, Ed., "Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base", RFC 4802, February 2007.
- [RFC5063] Satyanarayana, A., Ed., and R. Rahman, Ed., "Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart", RFC 5063, October 2007.
- [BGP-TE] Ould-Brahim, H., Fedyk, D., and Y. Rekhter, "Traffic Engineering Attribute", Work in Progress, January 2008.
- [CONF-SEG] Bradford, R., Ed., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Key-Based Mechanism", Work in Progress, May 2008.
- [GMPLS-SEC] Fang, L., Ed., " Security Framework for MPLS and GMPLS Networks", Work in Progress, February 2008.

11. Acknowledgments

The authors would like to thank Ichiro Inoue for valuable comments. In addition, the authors would like to thank Marco Carugi and Takumi Ohba for valuable comments in the early development of this document.

Thanks to Tim Polk and Mark Townsley for comments during IESG review.

Authors' Addresses

Tomonori Takeda (editor)
NTT Network Service Systems Laboratories, NTT Corporation
3-9-11, Midori-Cho
Musashino-Shi, Tokyo 180-8585 Japan
Phone: +81 422 59 7434
EMail: takeda.tomonori@lab.ntt.co.jp

Deborah Brungard
AT&T
Rm. D1-3C22 - 200 S. Laurel Ave.
Middletown, NJ 07748, USA
Phone: +1 732 4201573
EMail: dbrungard@att.com

Adrian Farrel
Old Dog Consulting
Phone: +44 (0) 1978 860944
EMail: adrian@olddog.co.uk

Hamid Ould-Brahim
Nortel Networks
P O Box 3511 Station C
Ottawa, ON K1Y 4H7 Canada
Phone: +1 (613) 765 3418
EMail: hbrahim@nortel.com

Dimitri Papadimitriou
Alcatel-Lucent
Francis Wellensplein 1,
B-2018 Antwerpen, Belgium
Phone: +32 3 2408491
EMail: dimitri.papadimitriou@alcatel-lucent.be

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.