

Network Working Group
Internet Draft
Proposed Status: Informational
Expires: April 2007

Tomonori Takeda (Editor)
NTT
October 2006

Applicability Statement for Layer 1 Virtual Private Networks (L1VPNs) Basic Mode

[draft-takeda-l1vpn-applicability-basic-mode-00.txt](http://www.ietf.org/drafts/takeda-l1vpn-applicability-basic-mode-00.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 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. Contributors.....	2
2. Terminology.....	3
3. Introduction.....	3
4. Basic Mode Overview.....	3
5. Supported Network Types.....	4
5.1 Data Plane.....	4
5.2 Control Plane.....	4
6. Addressing.....	5
7. Provider Control of its Infrastructure.....	5
7.1 Provisioning Model.....	5
7.2 PE-PE Segment Control.....	6
7.2.1 Path Computation and Establishment.....	6
7.2.2 Resource Management.....	7
7.2.3 Consideration of CE-PE TE information.....	7
7.3 Connectivity Restriction.....	8
8. Customer Control of its VPN.....	8
8.1 Topology Control.....	8
8.2 Note on Routing.....	8
9. Scalability, Resiliency.....	9
9.1 Scalability.....	9
9.2 Data Plane Resiliency.....	10
9.3 Control Plane Resiliency.....	11
10. Security.....	11
10.1 Topology Confidentiality.....	11
10.2 External Control of the Provider Network.....	12
10.3 Data Plane Security.....	12
10.4 Control Plane Security.....	13
11. Management.....	14
12. References.....	14
12.1 Normative References.....	14
12.2 Informative References.....	14
13. Acknowledgments.....	15
14. Author's Addresses.....	15
15. Intellectual Property Consideration.....	16
16. Full Copyright Statement.....	17

1. Contributors

This document is the result of contributions from several authors who are listed here in alphabetic order. Contact details for these authors can be found in a separate section near the end of this document.

Deborah Brungard (AT&T)

T.Takeda, et al.

Expires April 2007

[Page 2]

Adrian Farrel (Old Dog Consulting)
Hamid Ould-Brahim (Nortel Networks)
Dimitri Papadimitriou (Alcatel)
Tomonori Takeda (NTT)

2. Terminology

The reader is assumed to be familiar with the terminology in [\[RFC3031\]](#), [\[RFC3209\]](#), [\[RFC3471\]](#), [\[RFC3473\]](#), [\[RFC4202\]](#), [\[RFC4026\]](#) and [\[L1VPN-FW\]](#).

3. 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 [\[L1VPN-FW\]](#).

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 [\[L1VPN-BM\]](#), [\[L1VPN-BGP-DISC\]](#), and [\[L1VPN-OSPF-DISC\]](#), 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.

4. Basic Mode Overview

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

Note that in L1VPNs, routing operates within the provider network and may be used by PEs to exchange information specific to the VPNs 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-PE segment of the CE-CE VPN connection that is controlled and computed by the provider (PE-PE segment control). On the other hand, the VPN itself, constructed from a set of CEs and the VPN connections provided by the provider, is

under the control of each customer. This 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 VPN to a third party, including to the provider itself. There is a confidentiality requirement between the provider and each customer.

[[L1VPN-BM](#)], which expands [[RFC4208](#)], specifies GMPLS signaling to establish CE-CE VPN connections.

[[L1VPN-BGP-DISC](#)] and [[L1VPN-OSPF-DISC](#)] specify alternative mechanisms to exchange membership information between PEs, based on BGP and OSPF respectively.

5. Supported Network Types

5.1 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 (OXC), or Photonic Cross-Connects (PXC). 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 OXC). The provider network may provide a single type of VPN 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 [[L1VPN-FW](#)] and [[L1VPN-BM](#)], 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 VPN, or to multiple VPNs, and a PE may support one or more VPNs through a single CE or through multiple CEs.

5.2 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 [[L1VPN-FW](#)] and [[L1VPN-BM](#)], 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-PE relationship is applicable to more than one VPN. This makes it possible to determine to which VPN such control plane messages apply. Such disambiguation can be achieved by allocating a separate control channel to each VPN (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 10](#).

[6. Addressing](#)

As described in [[L1VPN-BM](#)], 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 re-use addresses used by the provider network, and may re-use 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 a CE-PE control channel addressing realm overlap. Furthermore, once a VPN connection has been established, the L1VPN Basic Mode does not enforce any restriction on address assignment for this VPN connection (treated as a link) for customer network operation (e.g., IP network, MPLS network).

[7. Provider Control of its Infrastructure](#)

[7.1 Provisioning Model](#)

As described in [[L1VPN-BM](#)], for each VPN that has at least one customer-facing port on a given PE, the PE maintains a Port Information Table (PIT) associated with that VPN. 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 VPN. In addition, for local PE ports of a given VPN 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 VPNs is added to a PE, PIT entries associated to those VPNs need to be configured on the PE. Section 4 of [[L1VPN-BM](#)] specifies such procedures:

- If no PIT exists for the VPN 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 remote PE associated with the VPN. For such purpose, manual configuration, or some sort of auto-discovery mechanisms can be used. [[L1VPN-BGP-DISC](#)] and [[L1VPN-OSPF-DISC](#)] specifies alternative auto-discovery mechanisms.
- In addition, remote PIT information associated with the VPN 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, [[L1VPN-BGP-DISC](#)] and [[L1VPN-OSPF-DISC](#)].

When VPN 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.

[7.2](#) PE-PE Segment Control

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

[7.2.1](#) Path Computation and Establishment

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

Examples of PE-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-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 management systems).

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-PE segments for a specific pair of PEs. When a PE receives a Path message from a CE for a VPN connection, a PE needs to determine which PE-PE segment to use. In such cases, the provider may want to control:

- Which VPN uses which PE-PE VPN segment.
- Which CE-CE VPN connection uses which PE-PE VPN segment.

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

- Mapping between individual PIT entries and PE-PE segments.
- Use of a Path Key ID [[Conf-Segment](#)] supplied by the provider to the CE, and signaled by the CE as part of the VPN 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-CE VPN connection to a pre-established PE-PE segment.

[7.2.2 Resource Management](#)

The provider network may operate resource management based on various policies. These policies can be applied per VPN or per VPN 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 VPNs or customers. Such a function might be achieved within the scope of the Basic Mode using resource affinities [[RFC3209](#)], but the details of per-VPN resource models (especially in terms of CE-PE routing) are considered as part of the Enhanced Mode.

[7.2.3 Consideration of CE-PE TE information](#)

[L1VPN-OSPF-DISC] and [[BGP-TE](#)] allow CE-PE TE link information to be injected into the provider network. 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-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 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 9.2](#), consideration of the full feature set enabled by dual-homing (such as resiliency) is out of scope of the L1VPN Basic Mode.

[7.3](#) Connectivity Restriction

The L1VPN Basic Mode allows restricting connection establishment between CEs belonging to the same VPN for policy reasons (including VPN security). Since the PIT at each PE is associated with a VPN, 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 VPN connections are pointlessly attempted.

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

[8](#). Customer Control of its VPN

[8.1](#) Topology Control

In the L1VPN Basic Mode, VPN connection topology is controlled by the customer. That is, a customer can request setup/deletion/modification of VPN connections using signaling mechanisms specified in [\[L1VPN-BM\]](#).

Also note that if there are multiple CE-PE TE links (single-homed or multi-homed), a customer can specify which CE-PE TE link to use to support any VPN connection. Alternatively, a customer may let the provider choose the CE-PE TE link at the egress side, as described in [Section 6.2.3](#).

[8.2](#) 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 VPN-specific way, perhaps through configuration or through a directory server.

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

Furthermore, if a VPN is used by a customer to operate a private Multiprotocol Label Switching (MPLS) or GMPLS network, the customer may wish to treat a VPN connection as a Traffic Engineering (TE) link, and this requires a CE-CE control channel. Note that a Forwarding Adjacency [[RFC4206](#)] cannot be formed from the CE-CE VPN 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 VPN connection. Otherwise, CE-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 [[L1VPN-FW](#)].

9. Scalability, Resiliency

9.1 Scalability

There are several factors that impact scalability.

- o Number of VPNs (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 VPNs supported in a provider network is governed by how the ID associated with a VPN 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-PE TE links for each VPN

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 [[L1VPN-OSPF-DISC](#)] floods membership information not only among PEs, but also to all P nodes. This may lead to scalability concerns, compared to [[L1VPN-BGP-DISC](#)], 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-PE TE link information, and not customer routing information, which is quite different from the mode of operation of a L3VPN. Therefore, the scalability concern is considered to be less problematic.

o Number of VPN connections

With the increase of this number, information to be maintained on each PE/P increases. When stitching or nesting is used, states to be maintained at PE increase 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, 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.

[9.2](#) Data Plane Resiliency

The L1VPN Basic Mode supports following data plane recovery techniques [[L1VPN-BM](#)].

o PE-PE segment recovery

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

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

- The CE-PEs link may be protected at a lower layer and GMPLS signaling may request the use of link-level protection
- If the PE-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-PE LSPs.

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

Note that it is not possible to protect only CE-PE portion or PE-PE portion by link protection because the CE-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 CE-PE portion by link recovery and PE-PE portion by segment recovery at the same time.

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

[9.3 Control Plane Resiliency](#)

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

[10. Security](#)

Security considerations are described in [[L1VPN-FW](#)], and this section describes how these considerations are addressed in the L1VPN Basic Mode.

[10.1 Topology Confidentiality](#)

As specified in [[L1VPN-BM](#)], 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 [[L1VPN-BM](#)], 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 NOTIFY_REQUEST object from the CE.

Even in the case of pre-computed and/or pre-signaled PE-PE segments, provider topology confidentiality may be preserved through the use of path key IDs [[Conf-Segment](#)].

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 10.4](#).

The provider network is responsible for maintaining confidentiality of topology information between customers and across VPNs. 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.

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

[10.3](#) Data Plane Security

As described in [[L1VPN-FW](#)], at layer 1, data plane information is normally assumed to be secure once connections are established since those connections are dedicated VPNs. That is, it is not possible to communicate unless there is a connection.

In order to protect against mis-delivery, each VPN connection is restricted to use only within a single VPN. That is, a VPN connection does not connect CEs that are in different VPNs. 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 VPN. This is achieved as described in [Section 7.3](#).

Furthermore, customers can apply their own security mechanisms to protect data plane information (CE-CE security). This includes IPsec for IP traffic.

[10.4](#) Control Plane Security

There are two aspects for control plane security.

First, the entity connected over a CE-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](#)].

Second, it must be possible to secure communication over a CE-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. IPsec tunnels can further be used for this purpose.

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 need mechanisms to detect Denial of Service (DoS) attacks and to protect against them reactively and proactively. In the Basic Mode, this relies on management system(s).

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-CE control plane security may form part of the data plane protection where the control plane traffic is carried in-band in the VPN connection. Where the CE-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.

11. Management

Manageability considerations are described in [[L1VPN-FW](#)]. In the L1VPN Basic Mode, we rely on management system(s) 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 10](#)).

Details are for further study.

12. References

12.1 Normative References

- [L1VPN-FW] Takeda, T., Editor "Framework and Requirements for Layer 1 Virtual Private Networks", [draft-ietf-l1vpn-framework](#), work in progress.
- [RFC4208] Swallow, G., et al., "Generalize Multiprotocol Label Switching(GMPLS) User-Network Interface: Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model," [RFC4208](#), October 2005.
- [L1VPN-BM] Fedyk, D., and Rekhter, Y., Editors, "Layer 1 VPN Basic Mode", [draft-ietf-l1vpn-basic-mode](#), work in progress.
- [L1VPN-BGP-DISC] Ould-Brahim, H., Fedyk, D., and Rekhter, Y., "BGP-based Auto-Discovery for L1VPNs", [draft-ietf-l1vpn-bgp-auto-discovery](#), work in progress.
- [L1VPN-OSPF-DISC] Bryskin, I., and Berger, L., "OSPF Based L1VPN Auto-Discovery", [draft-ietf-l1vpn-ospf-auto-discovery](#), work in progress.
- [SEG-PROT] Berger, L., et al., "GMPLS Based Segment Recovery", [draft-ietf-ccamp-gmpls-segment-recovery](#), work in progress.
- [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.

12.2 Informative References

- [RFC3031] Rosen, E., Viswanathan, A. and R. Callon,

- "Multiprotocol label switching Architecture", [RFC 3031](#), January 2001.
- [RFC3471] Berger, L., Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", [RFC 3471](#), January 2003.
- [RFC3473] Berger, L., Editor "Generalized Multi-Protocol Label Switching (GMPLS) Signaling - Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.
- [RFC4202] Kompella, K., et al., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4202](#), October 2005.
- [RFC4204] Lang, J., "Link Management Protocol (LMP)", [RFC 4204](#), October 2005.
- [RFC4026] Andersson, L., and Madsen, T., "Provider Provisioned Virtual Private Network (VPN) Terminology", [RFC 4026](#), March 2005.
- [RFC4206] Kompella, K., Rekhter, Y., "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", [RFC 4206](#), October 2005.
- [BGP-TE] Ould-Brahim, H., Fedyk, D., and Rekhter, Y., "Traffic Engineering Attribute", [draft-fedyk-bgp-te-attribute](#), work in progress.
- [RFC4655] Farrel, A., Vasseur, JP, Ash, J., "Path Computation Element (PCE) Architecture", [RFC 4655](#), August 2006.
- [Conf-Segment] Bradford, R., Editor, "Preserving Topology Confidentiality in Inter-Domain Path Computation and Signaling", [draft-bradford-pce-path-key](#), work in progress.

[13. Acknowledgments](#)

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

[14. Author's Addresses](#)

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)
Francis Wellensplein 1,
B-2018 Antwerpen, Belgium
Phone: +32 3 2408491
Email: dimitri.papadimitriou@alcatel.be

Tomonori Takeda
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

15. Intellectual Property Consideration

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.

16. Full Copyright Statement

Copyright (C) The Internet Society (2006).

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

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

