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

Tomonori Takeda (Editor)
NTT

November 2006

Applicability analysis of Generalized Multiprotocol Label Switching (GMPLS) protocols for the Layer 1 Virtual Private Network (L1VPN) Enhanced Mode

draft-ietf-l1vpn-applicability-enhanced-mode-00.txt

Status of this Memo

Abstract

This document provides an applicability analysis on the use of Generalized Multiprotocol Label Switching (GMPLS) protocols and mechanisms to satisfy the requirements of the Layer 1 Virtual Private Network (L1VPN) Enhanced Mode.

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 Enhanced Mode of operation features exchange of routing information between the layer 1 network and the customer domain.

In addition, this document identifies areas where additional protocol extensions or procedures are needed to satisfy the requirements of the L1VPN Enhanced Mode, and provides guidelines for potential extensions.

Table of Contents

1. Contributors

The details of this document are 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)

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 shows the applicability of existing Generalized Multiprotocol Label Switching (GMPLS) protocols and mechanisms to the Layer 1 Virtual Private Network (L1VPN) Enhanced Mode. In addition, this document identifies several areas where additional protocol extensions or modifications are needed to meet the L1VPN Enhanced Mode service requirements set out in [L1VPN-FW].

In particular, this document shows section by section (from Section 5 to 8) the applicability of GMPLS protocols and mechanisms to each sub-model of the Enhanced Mode mentioned in [L1VPN-FW], along with additional work areas needed to fully support the requirements for each sub-model.

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

As will be described in this document, support of the Overlay Extension service model and the Virtual Node service model are well covered by existing protocol mechanisms already described in other documents, with only minor protocol extensions required. The Virtual Link service model and the Per-VPN Peer service model are not explicitly covered by existing documents, but can be realized by extending current GMPLS protocols and mechanisms as described in this document.

The following section lists areas where additional work may be required to fully support the requirements for each sub-model. Some of the requirements are optional, therefore the additional work is also optional.

Commonalities of mechanisms over various sub-models, as well as over the L1VPN Basic Mode need to be considered. Also, various mechanisms should be coordinated in such a way that services are provided in a fully functional manner.

3.1 Work Items

This list of additional work areas is a summary derived from the main body of this document. The list will be updated in later versions of this document along with the development of the additional or enhanced requirements and increased understanding of the issues. As work progresses on protocol extensions, this list will be updated to remove completed items, and the body of this document will be updated to describe the analysis of protocol extensions. The intention is that this whole section is removed when the work has been completed and this document progresses to become an RFC.

- o Routing representation (how a VPN should be represented in routing, e.g., single area, multi-area, multi-AS). See Sections 6.3 and 7.3.
- o Signaling and routing for support of the Per-VPN Peer service model. See Section 8.3.

3.2 Existing Solutions Drafts

This section lists existing solution documents that describe how the L1VPN Enhanced Mode may be constructed using the mechanisms of GMPLS. This document draws on those solutions and explains their applicability and suggests further extensions to make the solutions more closely match the framework described in [L1VPN-FW]. Further solution documents may be listed in a future version of this document.

- o [GVPN] describes a suite of port-based Provider-provisioned VPN services called Generalized VPNs (GVPNs) that use BGP for VPN auto-discovery and GMPLS as a signaling mechanism.

- o [L1VPN-BM] addresses the L1VPN Basic Mode signaling.

Note that although [L1VPN-BM] specifies signaling mechanisms for L1VPN Basic Mode, it is applicable to the L1VPN Enhanced Mode, unless otherwise specified. Therefore, main focus of this document is how to realize routing related information exchange between a CE and a PE.

4. General Guidelines

This section provides general guidelines for L1VPN solutions. Note that applicability to specific sub-models will be separately described in following sections.

One important general guideline is that protocol mechanisms should be re-used where possible. This means that solutions should be incremental, building on existing protocol mechanisms rather than developing wholly new protocols. Further, as service models are extended or developed resulting in the requirement for additional functionalities, deltas should be added to the protocol mechanisms rather than developing new techniques. [L1VPN-FW] describes how the

service models can be seen to provide "cascaded" functionality, and this should be leveraged to achieve re-use of protocol extensions so that, for example, it is highly desirable that the same signaling protocols and extensions are used in both the Basic Mode and the Enhanced Mode.

In addition, the following are general guidelines.

- The support of L1VPNs should not necessitate any change to core (P) devices. Therefore, any protocol extensions made to facilitate L1VPNs need to be made in a backward compatible way allowing GMPLS aware P devices to continue to function.
- Customer (C) devices not directly involved in providing L1VPN services should also be protected from protocol extensions made to support L1VPNs. Again, such protocol extensions need to be backward compatible. Note however, that some L1VPN service models allow for VPN connectivity between C devices rather than between CE devices: in this case, the C devices may need to be aware of protocol extensions.
- Solutions should aim to minimize the protocol extensions on CE devices.
- Solutions should be scalable and manageable. Solutions should not require L1VPN state to be maintained on the P devices as much as possible.
- Solutions should be secure. Providers should be able to screen and protect information based on their operational policies.
- Solutions should provide an operational view of the L1VPN for the customer and provider. There should be a common operational and management perspective in regard to other (L2 and L3) VPN services.

5. Overlay Extension Service Model

5.1 Overview of the Service Model

This service model complements the Basic Mode and may assume all of the requirements, solutions and work items for that model.

In this service model, a CE receives from its attached PEs a list of TE link addresses to which it can request a VPN connection (i.e., membership information).

The CE may also receive some TE information concerning these CE-PE links within the VPN (e.g., switching type).

The CE does not receive any of the following from the PE.

- Routing information about the core provider network.
- Information about P device addresses.
- Information about P-P, PE-P or PE-PE TE links.

- Routing information about other customer sites. The CE may have access to routing information about the remainder of the VPN (C-C and CE-C links), but this is exchanged by control plane tunneling on the CE-CE connections and is not passed to the CE in the control plane exchange between PE and CE.

5.2 Applicability of Existing Solutions

The following are required in this service model (in addition to requirements in the L1VPN Basic Mode).

- VPN membership information exchange between a CE and PE.
- CE-PE TE link information exchange between a CE and a PE.

[GVPN] covers the requirement to exchange membership information between the CE and the PE based on BGP. Furthermore, [BGP-TE] allows the exchange of CE-PE TE link information between a CE and a PE.

5.3 Additional Work Area(s)

None.

6. Virtual Node Service Model

6.1 Overview of the Service Model

In this service model, there is a private routing exchange between the CE and the PE, or to be more precise between the CE routing protocol instance and the VPN routing protocol instance running on the PE. The provider network is considered as one private node from the customer's perspective. The routing information exchanged between the CE and the PE includes CE-PE TE link information, customer network (i.e., remote CE sites), and may include TE links (Forwarding Adjacencies) connecting CEs (or Cs) across the provider network as well as control plane topology information from the customer network (i.e., CE sites).

6.2 Applicability of Existing Solutions

The following are required in this service model.

- VPN routing
- Signaling: CE-CE Label Switching Path (LSP) setup, deletion, and modification

[GVPN] covers most of the requirements.

Specifically, [GVPN] handles VPN routing by a per VPN database called the GVSI (Generalized Virtual Switching Instance) held in each PE.

GVSIs are inter-connected by tunnel-based control channels, and routing adjacencies are established between them. BGP is used for auto-discovery of remote GVSIs (VPN auto-discovery) in the same VPN. GVSIs advertise VPN routing information by using a single ROUTER_ID to represent the provider network as one node.

It is also possible to use IGP-based auto-discovery (based on [L1VPN-OSPF-DISC]), instead of BGP-based auto-discovery.

Signaling mechanisms are covered by [L1VPN-BM].

6.3 Additional Work Area(s)

o Routing Representation

[GVPN] allows flexible routing configuration for each VPN (e.g., single IGP area, multiple IGP areas, or multiple ASes).

However, it may be valuable to consider how to represent a VPN in routing. This may simplify the solution (e.g., in terms of scalability). This requires further discussion.

7. Virtual Link Service Model

7.1 Overview of the Service Model

In this service model, virtual links are established between PEs. A virtual link is assigned to each VPN and disclosed to the corresponding CEs. The routing information exchanged between the CE and the PE includes CE-PE TE links, customer network (i.e., remote CE sites), virtual links (i.e., PE-PE links) assigned to each VPN, and may include CE-CE (or C-C) Forwarding Adjacencies as well as control plane topology from the customer network (i.e., CE sites).

7.2 Applicability of Existing Solutions

Currently, there is no solution document for this type of service model.

7.3 Additional Work Area(s)

Simple modifications of [GVPN], in addition to enhancements mentioned in Section 6.3, may realize this type of service model. Modifications could be:

- Do NOT modify the ROUTER_ID of the TE link information when advertising a CE-PE TE link to the CE (in the OSPF packet header as well as in the LSA header).

- Set up Forwarding Adjacency LSPs (FA-LSPs, GVSI-LSPs in [GVPN] terms) between PEs to construct virtual links, and advertise these FAs in VPN routing. Note these FAs (virtual links) may be assigned private addresses, which means customer assigned addresses (or that customers are allowed to configure addresses). This may require extensions to current IGP behavior.

There could be other ways to construct virtual links (e.g., virtual links without actually setting up an FA-LSP [MRN-REQ]).

Resource management for a dedicated data plane is a mandatory requirement for the Virtual Link service model. This could be realized by assigning pre-configured FA-LSPs to each VPN routing protocol instance (no protocol extensions needed) in order to instantiate the necessary FAs.

Note that as in the case of the Virtual Node service model, solution details may differ depending on the routing representation. This requires further discussion.

8. Per-VPN Peer Service Model

8.1 Overview of the Service Model

In this service model, the provider partitions TE links within the provider network per VPN. The routing information exchanged between the CE and the PE includes CE-PE TE links, customer network (i.e., remote CE sites), as well as partitioned portions of the provider network, and may include CE-CE (or C-C) Forwarding Adjacencies and control plane topology from customer network (i.e., CE sites). Note that PEs may abstract routing information about the provider network and advertise it to CEs.

Note scalability must be carefully considered for advertising provider network routing information to the CE [INTER-DOMAIN-FW].

8.2 Applicability of Existing Solutions

Currently, there is no solution document for this type of service model.

8.3 Additional Work Area(s)

There are two approaches for this service model.

o Signaling and routing for support of the per-VPN Peer service model

  Option1: Virtual Link based approach

The Per-VPN Peer service model may be realized by extending the virtual link technique so that not only PEs but also Ps that contain end points of virtual links in the abstracted topology contain VPN routing instances. There may be no additional protocol extensions needed from the Virtual Link service model.

Option2: Virtual Node based approach

The Per-VPN Peer service model may be realized by extending the virtual node technique so that PEs selectively advertise provider internal TE links to CEs. There are several extensions needed for this.

Details are for further study.

9. Discussion

This section summarizes items for which existing solutions may need to be extended in order to fulfill the full set of L1VPN Enhanced Mode functionalities.

For the Overlay Extension service model and the Virtual Node service model, the existing solutions can be applied with few extensions.

As described in Sections 7.2 and 8.2, there are no existing solutions to support the Virtual Link service model and the Per-VP Peer service model. For the Virtual Link service model, however, minor extensions from existing solutions are expected to meet the requirements.

Note that the list of additional work areas will be updated in later versions of this document with the development of additional or enhanced requirements and further understanding of the issues.

o Routing representation
  - One building block for the Enhanced Mode
  - Further discussion required (single area, multi areas, multi ASes, etc.)
  - Impact: Details to be studied (routing etc.)

o Signaling and routing for support of the Per-VPN Peer service model
  - Details for further study

10. Manageability Considerations

Section 11 of [L1VPN-FW] describes manageability considerations for L1VPNs.

This document defines a following new manageability requirement specific for the L1VPN Enhanced Mode.

MIB modules MUST be available for any protocol extensions for the L1VPN Enhanced Mode.

A future revision of this document may cover more aspects.

11. Security Considerations

Section 12 of [L1VPN-FW] describes security considerations for L1VPNs. This document defines a following new security requirements specific for the L1VPN Enhanced Mode.

In the L1VPN Enhanced Mode, since there is a routing adjacency between a CE and a PE, care must be taken whether the provider network's control plane topology information is leaked to the CE. Due to security concerns, this is not recommended in general, and there must be a mechanism to prevent such operation.

A future revision of this document may cover more aspects.

12. References

12.1 Normative References

[RFC3668] Bradner, S., "Intellectual Property Rights in IETF Technology", BCP 79, RFC 3668, February 2004.

[L1VPN-FW] Takeda, T., Editor "Framework and Requirements for Layer 1 Virtual Private Networks", draft-ietf-l1vpn-framework, work in progress.

12.2 Informative References

For information on the availability of this document, please see http://www.itu.int.

[Y.1312] Y.1312 - Layer 1 Virtual Private Network Generic requirements and architecture elements, ITU-T Recommendation, September 2003.

For information on the availability of this document, please see http://www.itu.int.

[Y.1313] Y.1313 - Layer 1 Virtual Private Network service and network architectures, ITU-T Recommendation, July 2004.

[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., 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.

[RFC4026] Anderssion, L., and Madsen, T., "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

[GVPN] Ould-Brahim, H., and Rekhter, Y. Editors, "GVPN Services: Generalized VPN Services using BGP and GMPLS Toolkit", draft-ouldbrahim-ppvpn-gvpn-bgpgmpls, 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-fedyk-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.

[INTER-DOMAIN-FW] Farrel, A., et al., "A Framework for Inter-Domain MPLS Traffic Engineering", draft-ietf-ccamp-inter-domain-framework, work in progress.

[BGP-TE] Ould-Brahim, H., Fedyk, D., and Rekhter, Y., "Traffic Engineering Attribute", draft-fedyk-bgp-te-attribute, work in progress.

[MRN-REQ] Shiomoto, K., et al., "Requirements for GMPLS-based multi-region and multi-layer networks (MRN/MLN)", draft-ietf-ccamp-gmpls-mln-reqs, work in progress.

13. Acknowledgments

Authors would like to thank Marco Carugi, Ichiro Inoue, 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