Network Working Group
Internet Draft
Proposed Status: Informational
Expires: January 2006

Tomonori Takeda (Editor)
NTT

July 2005

Applicability analysis of GMPLS protocols to Layer 1 Virtual Private Networks

draft-takeda-l1vpn-applicability-03.txt

Status of this Memo

Copyright Notice

Abstract

This document provides an applicability analysis on the use of Generalized Multiprotocol Label Switching (GMPLS) protocols and mechanisms to satisfy the requirements of Layer 1 Virtual Private Networks (L1VPNs).

In addition, this document identifies areas where additional protocol extensions or procedures are needed to satisfy the requirements of L1VPNs, and provides guidelines for potential extensions.

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], [GMPLS-RTG], [RFC4026] and [L1VPN-FW].

3. Introduction

This document shows the applicability of existing Generalized Multiprotocol Label Switching (GMPLS) protocols and mechanisms to Layer 1 Virtual Private Networks (L1VPNs). In addition, this document identifies several areas where additional protocol extensions or modifications are needed to meet the L1VPN service requirements set out in [L1VPN-FW].

In particular, this document shows section by section (from section 5 to 7) the applicability of GMPLS protocols and mechanisms to each L1VPN service model mentioned in [L1VPN-FW], along with additional work areas needed to fully support the requirements for each service model. Note that management aspects, some of which are common over various service models, are described separately in section 8.

Additional, non-normative information regarding network usage of L1VPN service models is provided in the appendix.

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 Management-based service model, the Signaling-based service model, the Overlay Extension service model and the Virtual Node service model are well covered by existing 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.

Also, as will be described, the following are possible work areas where additional work may be required to fully support the requirements for each L1VPN service model. Some of the requirements are optional therefore the additional work is also optional. Also, some items below may have more than one existing mechanism (with possible extensions). For those items, the required work is to choose the minimum set of mechanisms.

Commonalities of mechanisms over various service models 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, it is expected that 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.

- MIB module for SPC
- Resource management per VPN
- Signaling mechanisms
- VPN membership information exchange within the provider network
- CE-PE TE link information exchange within the provider network
- VPN membership information exchange between a CE and a PE
- CE-PE TE link information exchange between a CE and a PE
- Routing representation (how a VPN should be represented in routing, e.g., single area, multi area, multi AS)
- Control plane routing (routing information exchange related to control plane topology, per-VPN control packet routing)
- Signaling and routing for support of the Per-VPN Peer service model
- Management aspects (fault management, configuration management, security management)
- MIB modules for protocol extensions

3.2. Existing Solutions Drafts

This section lists existing solutions documents that describe how L1VPNs 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 solutions 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 [GMPLS-UNI] addresses the application of GMPLS to the overlay model. The document provides a description of how the overlay model may be used to support VPN connections across a core GMPLS network.

4. General Guidelines

This section provides general guidelines for L1VPN solutions. Note that applicability to specific service 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 Signaling-based service model and the Signaling and Routing service model.

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.
- It should be considered 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.
- 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.

Note that some deployments may wish to support multiple L1 connection types (such as VC3, VC4, etc.) at the same time. Specific functionalities may need to be considered for these scenarios. This is for further study.

5. Applicability to Management-based Service Model

5.1. Overview of the Service Model

The customer and the provider communicate via a management interface. The provider management system(s) communicate with the PE/P to set up a connection.

Note that in this service model the PE-PE connections may be signaled using GMPLS under management control at the ingress PE, or may be statically provisioned through management control of the PEs and Ps. Thus, it remains appropriate to describe signaling and routing mechanisms within this service model.

5.2. Applicability of Existing Solutions

SNMP MIB modules are one way to realize connection setup/deletion/modification from the management system(s). In particular, GMPLS-LSR-STD-MIB [LSR MIB] can control static connections, while GMPLS-TE-STD-MIB [TE MIB] can control signaled connections.

As indicated in [L1VPN-FW], the specification of interface(s) between management system(s) (i.e. customer and provider) is out of the scope of this document.

5.3. Additional Work Area(s)

The following additional work areas are identified to support the Management-based Service Model.

o MIB module for SPC (Soft Permanent Connection)

  The notion of an SPC only applies if the PE-PE connection is signaled.

  There are no required extensions to the MIB modules to support the static parts of the connections (CE-PE links) since they can be managed as normal static links using [LSR-MIB].

For the signaled part of the connection (PE-PE), the ingress and egress PEs need to know which (static) CE-PE TE links to use. This information can be carried to the egress PE using egress control [EGRESS-CONTROL], but needs to be configurable at the ingress PE. There are two alternatives.

Option1: MIB module extension

Define two new MIB objects as part of the specification of the TE LSP [TE-MIB] to specify the ingress and egress CE-PE TE links to be used.

Option2: MIB object usage extension

Use the current MIB objects, but define new, extended meanings. There are two possible ways to do this.

(1) Set the mplsTunnelIngressLSRId in the mplsTunnelTable (that corresponds to the Tunnel Sender Address in the Sender Template object of RSVP-TE) to the ingress CE-PE TE link address. Set the mplsTunnelHopIpAddr of the final MplsTunnelHopEntry in the mplsTunnelHopTable to the egress CE-PE TE link address.

(2) Set the mplsTunnelHopIpAddr of the first MplsTunnelHopEntry in the mplsTunnelHopTable to the ingress CE-PE TE link address. This may require a new mplsTunnelHopAddrType value to be defined in order to give precise meaning. Set the mplsTunnelHopIpAddr of the final MplsTunnelHopEntry in the mplsTunnelHopTable to the egress CE-PE TE link address.

Detailed analysis of options 1 and 2 is for further study.

o Resource management per VPN

In the Management-based service model, the data plane may be managed to create two optional functional requirements.

- Resource management to create a dedicated per-VPN data plane. The provider network partitions link resources per-VPN for exclusive use by a particular VPN.
- Resource management to share part of the data plane among a specific sub-set of VPNs. The provider network assigns link resources to a specific sub-set of VPNs.

The default behavior, without this option, is that all resources are available for use by any VPN.

If either of these options are applied with a statically managed

PE-PE connection then the required function is a matter for policy within the network management tool for the core network. No extensions are required.

There are two alternatives to achieve this function for signaled PE-PE connections.

Option 1: Policy

A simple way to meet this requirement is to implement resource management functionalities as a policy solely in the entity that computes a path. No protocol extensions are needed because links and resources can be explicitly configured using [TE-MIB] and signaled using [RFC3473].

This scheme is especially effective when path computation is done in a centralized manner (e.g., in the management system(s)) and is similar to the policy applied to achieve these functional options using statically configures PE-PE connections.

Option 2: Routing extension

The other alternative is advertise the amount of resources available to each VPN using extensions to the TE information flooding performed by the routing protocol within the core provider network.

In this scheme, the PE/P can compute a path in a distributed way, thus this scheme is especially beneficial in the case of dynamic restoration (restoration that does not reserve backup resources in advance).

Note that link coloring might be used for this purpose, but this would eliminate the opportunity to use link coloring for other purposes (e.g., link coloring within VPNs).

Detailed analysis of options 1 and 2 is for further study.

o Other considerations

When path computation is done in a centralized entity (e.g., management system(s)), it is important that resource information is synchronized between the core provider network and such an entity.

6. Applicability to Signaling-based Service Model (Basic Mode)

6.1. Overlay Service Model

6.1.1. Overview of the Service Model

In this service model, there is no routing exchange between the CE and the PE. Connections are setup by GMPLS signaling between the CE and the PE, and then across the provider network.

Note that routing operates within the provider network and may be used by PEs to exchange information specific to the VPNs supported by the provider network.

6.1.2. Applicability of Existing Solutions

The following are required in this service model.

- VPN membership information exchange: CE-PE TE link address exchange between PEs, along with information associated with a VPN. The TE link addresses may be customer assigned private addresses.
- Signaling: CE-CE LSP setup, deletion and modification
- Others: Resource management per VPN etc.

[GVPN] and [GMPLS-UNI] cover most of the requirements.

Specifically, [GVPN] covers VPN membership information exchange by BGP running on the PEs. Customer assigned private addresses for customer site CEs are configured on the PE that provides VPN access to the customer site, and are exchanged by BGP along with a provider network address (which is reachable in the provider network's routing) and an ID associated with the VPN (i.e., Route Target). This allows PEs to perform address translation/mapping and connectivity restriction.

The other possibility is to use IGP based VPN membership information exchange (e.g., similar to as an AS external route, or based on [OSPF-NODE-ADDR], with extensions for VPN applications).

In addition, [GVPN] and [GMPLS-UNI] suggest two signaling mechanisms for VPN connections.

o Shuffling [GVPN]

  Information carried in RSVP-TE messages identifying a LSP (i.e., SESSION and SENDER_TEMPLATE objects) is translated by the ingress and egress PE. There is one end-to-end session (i.e., CE-CE), but the identifiers of that session change along the path of the LSP.

o Nesting [GVPN][GMPLS-UNI][LSP HIER]

  When Path message arrives at the ingress PE, the ingress PE checks whether there is appropriate PE-PE connectivity. If there is not, it initiates a PE-PE FA-LSP. The CE-CE LSP is carried nested

hierarchically within the FA-LSP. There are two sessions (i.e., CE-CE and PE-PE).

LSP stitching [STITCHING] operates in a similar manner to LSP nesting. The properties of the PE-PE LSP segment are such that exactly one end-to-end LSP can be stitched to the LSP segment i.e., the PE-PE LSP and the CE-CE LSP correspond exactly one to one. There are two sessions (i.e., CE-CE and PE-PE).

LMP [LMP] may be running between a CE and a PE. In that case, the PE is able to obtain customer assigned private addresses on directly attached CEs automatically. This eliminates configuring manually the customer assigned private addresses on PEs, which are distributed by membership information exchange mechanisms.

6.1.3. Additional Work Area(s)

The following additional work areas are identified to support the Overlay service model.

o Signaling mechanisms

As described in section 6.1.2, [GMPLS-UNI] and [GVPN] suggest two signaling mechanisms for VPN connections.

Option 1: Shuffling

In this mechanism, objects need to be translated at the ingress and egress PEs. It is necessary to specify rules for this translation and mechanisms to ensure that the information is available in order to perform the translation.

Option 2: Nesting

In this mechanism, there is a need to set up a PE-PE FA-LSP.

In the case of nesting, PE-PE direct signaling message exchange takes place, and this message exchange may use the provider network addressing space, or the VPN addressing space. It may be necessary to specify an addressing space to be used.

When the provider network addressing space is used, there must be a mechanism to identify which VPN each message is associated with at PEs. Otherwise, the PE received the message is not able to proceed the message furthermore (i.e., session identification, and next hope resolution). This mechanism needs to be specified.

When the VPN addressing space is used by forming per-VPN control

channels between PEs, the identification of VPN is straightforward. However, the mechanisms to realize per-VPN control channels need to be specified (e.g., IP-based tunnel, physically separate control channels).

Detailed analysis, including under which condition signaling mechanisms (shuffling or nesting) should be used, is for further study.

o VPN membership information exchange within the provider network

As described in section 6.1.2, there are two existing mechanisms for the functional option of exchanging VPN membership information within the provider network. This model does not support VPN membership exchange between CE and PE (see section 7.1) and so such information is assumed to be configured within the provider network, usually on the PEs.

Option 1: BGP-based

[GVPN] specifies a BGP-based mechanism to realize VPN membership information exchange between PEs without informing core Ps. Configuration of this information is performed at the PEs that provide access to a VPN. There is no additional work required, except to update [GVPN] for detailed specification of format and encoding.

Option 2: IGP-based

OSPF allows AS external routes to be advertised. In addition, [OSPF-NODE-ADDR] extends OSPF-TE to advertise a router's local addresses. These mechanisms can be used to advertise CE-PE TE link addresses within the core provider network. In order to support customer assigned private addresses and connectivity restrictions, this mechanism needs to be extended to exchange information similar to an RT (Route Target) and possibly an RD (Route Distinguisher), along with CE-PE TE link addresses.

Detailed analysis of options 1 and 2 is for further study.

o Resource management per VPN

Section 5.2 describes how provider network resources can be partitioned for use by a single VPN or a sub-set of VPNs. Note that in option 1 of section 5.2, when path computation is done in a separate entity, the interface to the PCE (Path Computation Element) [PCE ARCH] may need to be extended for VPN identification.

Note also that it is also possible to apply resource partitioning

in the CEs and on the CE-PE links in this model. It will be necessary, however, to ensure consistent configuration through the network management tools of both the customer and provider equipment to provide this function.

- CE-PE TE link information exchange within the provider network

In the Signaling-based service model, it may be useful to consider not only TE link information within the provider network (PE-P, PE-PE TE links), but also remote CE-PE TE link information in path computation. This prevents connection setup failure due to lack of resources on remote CE-PE TE links. Therefore, CE-PE TE link information should be optionally propagated within the provider network to be used for path computation.

There are two alternatives for this.

Option1: BGP-based

[GVPN] describes potential use of BGP for exchanging CE-PE TE link information. Detailed protocol specifications are needed as additional work. This option is consistent with the BGP-based membership exchange described above.

Option 2: IGP-based

An alternative is to use IGP to advertise CE-PE TE links. Since a CE does not participate in routing protocol exchange with the provider network, TE link information must be properly constructed by the PE advertising full CE-PE TE link information. This option is consistent with the IGP-based membership exchange described above.

Detailed analysis of options 1 and 2 is for further study.

- Other considerations

Note, there could be a L1VPN solution where connectivity restriction, address translation/mapping etc. are performed not in PEs, but in other entities, such as a centralized policy server. In this case, the interface between the PE and the other entity may need to be specified. This could utilize existing mechanisms such as COPS or LDAP.

Also note that [GVPN] assumes that a PE and a CE communicate using separate control channels for each VPN (i.e., a CE-PE control channel is not shared by multiple VPNs). As described above, this facilitates easy separation of VPN signaling messages, but is achieved at the cost of extra configuration at the CE and PE. If

a shared control channel is desired in the [GVPN] solution, additional mechanisms such as VPN identification within signaling messages, may be required.

7. Applicability to Signaling and Routing Service Model (Enhanced Mode)

7.1. Overlay Extension Service Model

7.1.1. Overview of the Service Model

This service model is a slight extension from the Overlay service model (section 6.1) 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 (a list of CE addresses within the same VPN).

The CE may also receive some of 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.

7.1.2. Applicability of Existing Solutions

The following are required in this service model.

- 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 by BGP for overlay extension.

The other possibility is to use IGP based VPN membership information exchange (e.g., similar to as an AS external route, or based on [OSPF-NODE-ADDR], with extensions for VPN applications).

7.1.3. Additional Work Area(s)

o VPN membership information exchange between a CE and a PE

As described in section 7.1.2, there are two existing mechanisms based on which VPN membership information exchange is realized.

Option 1: BGP-based

[GVPN] suggests a BGP-based mechanism to realize VPN membership information exchange between a CE and a PE.

Option 2: IGP-based

OSPF allows AS external routes to be advertised. In addition, [OSPF-NODE-ADDR] extends OSPF-TE to advertise a router's local addresses. These mechanisms can be used to advertise CE-PE TE link addresses between a CE and a PE.

Detailed analysis of options 1 and 2 is for further study.

o CE-PE TE link information exchange between a CE and a PE

As just mentioned [GVPN] suggests a BGP-based mechanism to realize VPN membership information exchange. Such a mechanism does not extend well to carrying additional TE information about the CE-PE link either between PEs or between PE and CE because it is generally agreed that BGP should not be used to transport TE information.

However, there is no reason in principle why specific, tightly specified extensions should not be used to transport this additional information within the limited context of the L1VPN.

An alternative is to use an IGP mechanism to distribute this information. [GVPN] does not constrain the CE-PE routing protocol to be BGP, so this option could be used in either of the options listed for membership exchange.

Note that the additional membership and TE information might be considered as superfluous within the core provider network were it to be flooded by an IGP to all P devices. An option, in this case might be to run a separate instance of the IGP including only the CEs and PEs.

Mechanisms other than routing protocols could be used to exchange reachability/TE information between the CE and the PE.

7.2. Virtual Node Service Model

7.2.1. 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 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, 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 CE sites.

7.2.2. Applicability of Existing Solutions

The followings are required in this service model.

- VPN routing
- Signaling: CE-CE LSP setup, deletion, and modification
- Others: Resource management per VPN etc.

[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.

In addition, [GVPN] supports nested signaling (as in the case of the Signaling-based service model).

There are other ways to realize VPN auto-discovery. One such way is to use an IGP-based mechanism (e.g., based on [OSPF-CAP] or [OSPF-NODE-ADDR] with extensions). Other possibilities are to use a server based approach (e.g., DNS, based on [DNS DISCOVERY], RADIUS, based on [RADIUS DISCOVERY]) and multicast (e.g., based on [RFC2917]).

7.2.3. Additional Work Area(s)

The following additional work areas are identified to support the Virtual Node service model.

o Routing representation

  In the Virtual Node service model, one item that should be considered is how to represent a VPN in routing (e.g., single IGP area, multiple IGP areas, multiple ASes). Depending on the routing representation, solution details may differ (e.g., use of

auto-discovery). This requires further discussion.

o Resource management per VPN

See section 6.1.3

o Control plane routing

An explicit decision must be taken about whether the provider network's control plane topology information should be leaked to the CE. If it is, it may be necessary to separate the address spaces. Further, if control messages (e.g., BGP messages) can be transferred between CE sites using the provider network control plane, care must be taken over how to route per VPN control packets received from the CE.

7.3. Virtual Link Service Model

7.3.1. Overview of the Service Model

In this service model, virtual links are established between PEs. The routing information exchanged between the CE and the PE includes CE-PE TE links, CE sites, virtual links (i.e., PE-PE links), and may include CE-CE (or C-C) Forwarding Adjacencies as well as control plane topology from the CE sites.

7.3.2. Applicability of Existing Solutions

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

7.3.3. Additional Work Area(s)

Simple modifications of [GVPN], in addition to enhancements mentioned in section 7.2.3, 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 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.

Note there could be other ways to construct virtual links (e.g.,

virtual links without actually setting up a FA-LSP [MRN REQ]).

There is no additional work area beyond the work already identified for the Virtual Node service model mentioned in section 7.2.3, and that described above.

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

7.4. Per-VPN Peer Service Model

7.4.1. 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, 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 the 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].

7.4.2. Applicability of Existing Solutions

Currently, there is no solution document for this type of service model. However, [GVPN] provides several functionalities to meet this type of service model, as described in section 7.2.2. One way is to extend mechanisms for the Virtual Node service model. The other way is to extend mechanisms for the Virtual Link service model.

7.4.3. Additional Work Area(s)

As described in section 7.4.2, there are two approaches for this service model.

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

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

model

Option 1: Virtual node-based

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.

- Topology filtering

  The PE must choose TE links that are assigned to a specific VPN, and then advertise these TE links to a specific set of CEs corresponding to that VPN.

- Topology abstraction

  The PE may abstract routing information of the provider network, and then advertise abstracted topology information to the CE. It means that the PE may construct a TE link where a direct physical link does not exit, or the PE may construct a single node to represent multiple nodes and TE links.

  Note scalability must be carefully considered [INTER-DOMAIN FW].

- ERO/RRO expansion/modification

  The CE may specify an ERO with abstracted topology. The provider network must expand this ERO to match the provider network topology. Note this must be done even if a strict route is specified in the ERO passed from the CE.

  At the same time, when an RRO is requested, the RRO passed to the CE must be either edited to match the abstracted topology, or removed.

- Private address

  The provider network may support private addresses for routing information provided to the customer. This means that the customer is able to assign private addresses to a partitioned portion of the TE links within the provider network.

Option 2: Virtual link-based

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.

Detailed analysis of options 1 and 2 is for further study.

8. Management Aspects

8.1. Fault Management

The provider network may support various recovery techniques mentioned in [P&R TERM]. The customer may be allowed to specify the desired level of recovery in connection setup requests. The provider network may constitute a recovery domain (PE-PE recovery).

The following aspects need to be considered relative to L1VPNs.

o Shared recovery

When the provider network supports shared recovery (e.g., shared mesh restoration), the provider network may be able to support shared recovery only within the same VPN and/or shared recovery among multiple VPNs. The default mode is to be specified.

If the provider network supports both, the provider network must provide configuration tools for operators.

o Extra traffic

GMPLS recovery mechanisms support extra traffic. Extra traffic allows supporting preemptable traffic on recovery resources when these resources are not being used for the recovery of normal traffic [P&R TERM].

When the provider network supports extra traffic, the provider network may be able to support extra traffic only within the same VPN and/or extra traffic among multiple VPNs. The default mode is to be specified.

If the provider network supports both, the provider network must provide configuration tools for operators.

8.2. Configuration Management

Some VPN specific configuration aspects must be considered, such as:

o Configuration of resource management per VPN

Physical link resources may be dedicated, shared by a specific sub-set of VPNs, or shared by any VPNs. The provider network must

provide configuration tools for resource management per VPN.

- Configuration of virtual links

  For the Virtual Link service model and the Per-VPN Peer service model, the provider network must provide configuration tools for operators.

- Configuration of service model for each VPN

  When the provider network supports multiple service models, the provider network must provide configuration tools for operators.

8.3. Security Management

- CE-PE security

  When a CE-PE control channel is physically shared by multiple VPNs, security mechanisms need to be applied for data integrity and confidentiality of control messages exchanged. Furthermore, when a CE-PE control channel is dynamically setup, authentication need to be performed. The mechanisms to achieve these include IPsec.

  Denial of service attack is one significant security threat. The provider network must have mechanisms to detect denial of service attach, and to protect against it reactively and proactively.

  Details for additional work areas are for further study.

- CE-CE security

  The provider network must restrict connections between CEs in the same VPN. As such, the provider network must avoid mis-connection under any scenario, including failures, recovery and preemption.

  Furthermore, when customers want to assure security against the provider network, the customers may apply their own security mechanisms (CE-CE security). IPsec can be used for this purpose.

9. Discussion

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

Note that several of these items are in support of optional features. For the Management-based service model, the Signaling-based service model, the Overlay Extension service model and the Virtual Node service model, the existing solutions can be applied with few

extensions.

As described in sections 7.3.2 and 7.4.2, there are no existing solutions to support the Virtual Link service model and the Per-VPN 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 MIB module for SPC
  - Optional, but highly required for the Management-based service model
  - Two alternatives (MIB module extension or MIB object usage extension)
  - Impact: MIB module or none

o Resource management per VPN
  - Optional requirement for the Management-based, the Signaling-based, the Overlay extension and the Virtual Node service models
  - Mandatory requirement for the Virtual Link and the Per-VPN Peer service models (support of resource management for dedicated data plane)
  - Two alternatives (policy or routing extension)
  - For the Virtual Link service model, can be realized by no protocol extensions (assign pre-configured FA-LSPs to each VPN routing instance).
  - Impact: None or IGP

o Signaling mechanisms
  - Mandatory requirement for the Signaling-based service model and the Signaling and Routing service model
  - Two alternatives (shuffling or nesting)
  - Impact: Signaling

o VPN membership information exchange within the provider network
  - Mandatory requirement for the Signaling-based service model and the Overlay Extension service model
  - Two alternatives (BGP or IGP)
  - Impact: BGP or IGP

o CE-PE TE link information exchange within the provider network
  - Optional requirement for the Signaling-based service model and the Overlay Extension service model
  - Two alternatives (BGP or IGP)

- Impact: BGP or IGP

- VPN membership information exchange between a CE and a PE
  - Mandatory requirement for the Overlay Extension service model
  - Two alternatives (BGP or IGP)
  - Impact: BGP or IGP

- CE-PE TE link information exchange between a CE and a PE
  - Optional requirement for the Overlay Extension service model
  - Two alternatives (BGP or IGP)
  - Impact: BGP or IGP

- Routing representation
  - One building block for the Signaling and Routing service model
  - Further discussion required (single area, multi areas, multi ASes, etc.)
  - Impact: Details to be studied (routing, use of auto-discovery, etc.)

- Control plane routing
  - Optional requirement for the Signaling and Routing service model
  - Impact: Routing

- Signaling and routing for support of the Per-VPN Peer service model
  - Two options (virtual node-based, virtual link-based)
  - Impact: Routing, signaling (details to be studied)

- Management aspects
  - Default mode to be specified for shared recovery and extra traffic
  - Support of configuration tools mandatory for fault management and configuration management
  - Details for security management to be studied
  - Impact: mostly on operational tools (impacts on protocols to be studied)

10. Security Considerations

Section 8.3 describes security considerations.

11. IANA Considerations

This document defines no new protocols or extensions and makes no requests to IANA for registry management.

12. Acknowledgement

We would like to thank Marco Carugi, Ichiro Inoue and Takumi Ohba for their useful comments and suggestions.

13. Normative References

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

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

14. 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.

[GMPLS-UNI] Swallow, G., et al., "Generalize Multiprotocol Label Switching(GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", draft-ietf-ccamp-gmpls-overlay, work in progress.

[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.

[LSP HIER] Kompella, K., Rekhter, Y., "LSP Hierarchy with Generalized MPLS TE", draft-ietf-mpls-lsp-hierarchy, work in progress.

[STITCHING] Ayyangar, A. (editor), "Label Switched Path Stitching with Generalized MPLS Traffic Engineering", draft-ietf-ccamp-lsp-stitching, 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.

[P&R TERM] Mannie, E., and Papadimitriou, D. (editors), "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", draft-ietf-ccamp-gmpls-recovery-terminology, work in progress.

[LSR MIB] Nadeau, T., et al., "Generalized Multiprotocol Label Switching (GMPLS) Label Switching Router (LSR) Management Information Base", draft-ietf-ccamp-gmpls-lsr-mib, work in progress.

[TE MIB] Nadeau, T., et al., "Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base", draft-ietf-ccamp-gmpls-te-mib, work in progress.

[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.

[GMPLS-RTG] Kompella, K., et al., "Routing Extensions in Support of Generalized MPLS", draft-ietf-ccamp-gmpls-routing, work in progress.

[EGRESS CONTROL] Berger, L., "GMPLS Signaling Procedure For Egress Control", RFC 4003, February 2005.

[OSPF-CAP] Lindem, A. (editor), "Extensions to OSPF for Advertising Optional Router Capabilities", draft-ietf-ospf-cap, work in progress.

[OSPF-NODE-ADDR] Aggarwal, R., Kompella, K., "Advertising a Router's Local Addresses in OSPF TE Extensions", draft-ietf-ospf-te-node-addr, work in progress.

[LMP] Lang, J., "Link Management Protocol (LMP)", draft-ietf-ccamp-lmp, work in progress.

[DNS DISCOVERY] Squire, M., et al., "Using DNS for VPN Discovery", draft-luciani-ppvpn-vpn-discovery (Expired).

[RADIUS DISCOVERY] Weber, G., Editor "Using RADIUS for PE-Based VPN Discovery", draft-ietf-l2vpn-radius-pe-discovery (Expired).

[RFC2917] Muthukrishnan, K., Malis, A., " A Core MPLS IP VPN Architecture", RFC2917, September 2000.

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

[PCE ARCH] Ash, J., et al., "Path Computation Element (PCE) Architecture", draft-ietf-pce-architecture, work in progress.

[MRN REQ] Shiomoto, K., et al., "Requirements for GMPLS-based multi-region and multi-layer networks", draft-shiomoto-ccamp-gmpls-mrn-reqs, work in progress.

15. Authors' 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

Appendix I: Network Usage of L1VPN Service Models

This appendix provides additional information concerning network usage of the L1VPN service models.

- Management-based service model

  In this model, the provider network can support non-GMPLS capable CEs. Therefore, this model is best suited when customer networks are non-GMPLS, e.g., legacy SONET/SDH and IP/MPLS networks.

  It is expected that the provider network requires no or minimal GMPLS extensions for L1VPN specific functions.

- Signaling-based service model

  In this model, by implementing GMPLS signaling functions in CEs, the customer can request an LSP setup/deletion/modification to the provider by signaling. Other customer site nodes (C devices) do not need to be GMPLS-capable. Customers will receive rapid failure notifications of an LSP by using notification mechanisms available in GMPLS RSVP-TE.

  There are some L1VPN specific extensions required within the provider network. Concerning customer network routing information, since only CE-PE TE link addresses are contained within the provider network, it is expected that there is less concern on scalability. Trust relationships between the customer and the provider may need to be carefully considered.

- Signaling and Routing service model

  In this model, a customer can seamlessly operate its VPN using end-to-end GMPLS. Therefore, this model is best suited when customer networks are operated by GMPLS.

  For the service model where the provider network's routing information is not provided to customers (i.e., Virtual Node

service model), a customer can outsource routing complexity within the provider network to the provider. On the other hand, in the service model where the provider network routing information is provided to customers (i.e., Virtual Link service model and Per-VPN Peer service model), customers play more of a role. For example, by allowing customers to assign SRLG IDs for virtual links, customers can compute and set up end to end disjoint LSPs in their VPN.

There are some L1VPN specific extensions required within the provider network. Concerning customer network routing information, since the customer network routing information is contained within the provider network, scalability must be carefully considered. Trust relationships between the customer and the provider may need to be carefully considered as well.

Intellectual Property Considerations

Full Copyright Statement