INTERNATIONAL TELECOMMUNICATION UNION

TELECOMMUNICATION STANDARDIZATION SECTOR

STUDY PERIOD 2001-2004

STUDY GROUP 13

TD 25 (WP3/13)

English only

Original: English

Question: 11/13 Geneva, 3-12 February 2004

TEMPORARY DOCUMENT

Source: Q.11 Rapporteur

Title: Document summarizing Layer 1 VPN work in SG13 up to now (attached to this SG13 meeting liaison to IETF on L1 VPN)

1. Introduction

This is a document summarizing the Layer 1 VPN services work done inside ITU-T SG13 up to now, and presented in the form of IETF Internet draft. This document exists to demonstrate the service level requirements for L1VPNs and to create a framework for L1VPNs within the existing Internet architectures. As such, this document is for better understanding of L1VPNs within the IETF and for further cooperation.

It has been attached to this SG13 meeting liaison to IETF on L1 VPN).

Contact: Marco Carugi
Nortel Networks Europe
UK

Tel: +33 1 6955 7027
Fax: +33 1 6955 3058
Email: marco.carugi@nortelnetworks.com

Network Working Group Tomonori Takeda (Editor)
Internet Draft NTT
Expires: August 2004 February 2004

Framework for Layer 1 Virtual Private Networks
draft-takeda-l1vpn-framework-00.txt

Status of this Memo

Abstract

This document provides a framework for Layer 1 Virtual Private Networks (L1VPNs). This framework is intended to aid in developing and standardizing protocols and mechanisms to support interoperable L1VPNs.

The document examines motivations for L1VPNs, high level (service level) requirements, and outlines some of the architectural models that might be used to build L1VPNs.

0. Summary

(This section to be removed before publication as an RFC.)

0.1. Summary

This document describes a framework for Layer 1 VPNs (L1VPNs).

L1VPNs provide services over layer 1 networks, such as WDM and TDM networks. This document provides a framework for L1VPNs and the realization of the framework by those networks being controlled by GMPLS protocols.

0.2. Where does it fit in the picture of the IETF Work

Services may be provisioned across layer 1 networks using GMPLS protocols. L1VPNs may be managed and operated using these protocols as described in this document. GMPLS protocols were developed within the IETF using IP addressing and based on IP and other Internet protocols. The IETF continues to work with GMPLS protocols, enhancing them and applying them to new requirements.

VPN related work areas might also have points of interaction with the content of this document.

0.3. Justification

This document exists to demonstrate the service level requirements for L1VPNs and to create a framework for L1VPNs within the existing Internet architectures. As such, this document is the justification for better understanding of L1VPNs within the IETF.

Study Group 13 of the ITU-T has been investigating the service level requirements for L1VPNs with input from major network service providers and equipment vendors. There is a strong feeling within SG13 that the desirability of L1VPN services is growing and that there is a need for a minimum set of common approaches that will lead to interoperable solutions.

0.4. Related Internet Documents

Much of the background work for this document has been directly developed within the ITU-T and is presented as [Y.1312] and [Y.L1VPNARCH]. However, some Internet drafts are related to this topic and the following three ones are relevant ones.

- o draft-ouldbrahim-ppvpn-gvpn-bgpgmpls-04.txt (October 2003)
  "GVPN Services: Generalized VPN Services using BGP and GMPLS Toolkit"
  This draft describes a suite of port-based Provider-provisioned VPN services called Generalized VPNs (GVPNs) that uses BGP as a VPN auto-discovery and GMPLS as a signaling mechanism.

- o draft-ietf-ccamp-gmpls-overlay-02.txt (October 2003)
  "GMPLS UNI: RSVP Support for the Overlay Model"
  This memo addresses the application of GMPLS to the overlay model. In one section, the memo provides a description of how the overlay

model may be used to support VPN connections across a core GMPLS network.

- draft-andersson-ppvpn-terminology-04.txt (September 2003) "PPVPN terminology"
  This draft sets out terminology common to all Provider Provisioned VPNs. Although this draft specifically targets L2VPNs and L3VPNs, the terminology may be used to L1VPNs as well.

Contents

1. Contributors

This document is based heavily on the work of ITU-T Study Group 13 Question 11. SG13/Q11 has been investigating the service requirements and architecture for Layer 1 VPNs for some time, and this document is a summary and development of the conclusions they have reached. As such, ITU-T SG13 should be seen as a major contributor to this document.

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.

Raymond Aubin (Nortel)
Marco Carugi (Nortel)
Ichiro Inoue (NTT)
Hamid Ould-Brahim (Nortel)
Tomonori Takeda (NTT)

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The reader is assumed to be familiar with the terminology in [RFC3031], [RFC3209], [RFC3471], [RFC3473], [GMPLS-ROUTING] and [PPVPN-TERM].

3. Introduction

The document examines motivations for Layer 1 Virtual Private Networks (L1VPNs), provides high level (service level) requirements, and outlines some of the architectural models that might be used to build L1VPNs.

The objective of the document is mainly to present the requirements and architecture work in this field that has been undertaken within the ITU-T.

L1VPNs provide services over layer 1 networks. This document provides a framework for L1VPNs and the realization of the framework by those networks being controlled by GMPLS protocols.

3.1 Overview

3.1.1 Network Topology

The layer 1 network, made of Optical Cross-Connects (OXCs) or Time Division Multiplex (TDM) capable switches, may be seen as consisting of provider edge (PE) devices that give access from outside of the network, and provider (P) devices that operate only within the core of the network. Similarly, outside the layer 1 network is the customer network consisting of customer (C) devices with access to the layer 1 network made through customer edge (CE) devices.

A CE and 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.

3.1.2 Introducing Layer 1 VPNs

The concept of a provider provisioned VPN (PPVPN) has been established through many previous documents such as [L2VPN-FRAME] and [L3VPN-FRAME]. Terminology for PPVPNs is set out in [PPVPN-TERM] with special reference to layer 2 and layer 3 VPNs.

The realization of Layer 1 VPNs (L1VPNs) can be based on extensions of the concepts of the PPVPN to the layer 1 network. It must be understood that meeting the requirements set out in this document may necessitate modifications to the existing mechanisms both for the control plane within the layer 1 network and for service provisioning at the edge of the network between the CE and PE devices. It is at this interface (between CE and PE devices) that the L1VPN service is provided.

3.1.3 Current Technologies for Dynamic Layer 1 Provisioning

Pre-existing efforts at standardization have focused on the provision of dynamic connections within the layer 1 network (signaling and routing), and the interfaces for requesting services between the CE and PE or between PEs at network boundaries (UNI and E-NNI respectively).

No change in principle would be required to the operation within the network, and the E-NNI is not in scope for current L1VPN considerations. But the UNI is very relevant since it is a means by which the CE can make service requests to the PE to establish services (that is, connections) across the layer 1 network to remote CEs.

Current UNIs include features to facilitate requests for end-to-end (that is, CE to CE) service requests that include the specification of constraints such as explicit paths, bandwidth requirements, protection needs, and (of course) destinations.

The UNIs, however, do not provide a sufficiently high level of service to support VPNs without some additions. For example, there is no way to distinguish between control messages received over a shared control link at a UNI, and these messages must be disambiguated with respect to the L1VPN to which they apply.

Further, there is currently no leakage of routing information across the PE to CE boundary. While this restriction may be considered desirable from the perspective of network separation, VPN operation may benefit from the dynamic exchange of routing information between CEs that provide access to the VPNs.

In order that L1VPNs can be supported in a fully functional manner, these deficiencies and other requirements set out later in this document must be addressed.

3.2 Relationship with ITU-T

This document is based on the work of the ITU-T Study Group 13 Question 11. This group has been researching and specifying both the requirements and the architecture of L1VPNs for some time. In this context, this document is a representation of the findings of the ITU-T, and a presentation of those findings in terms and format that are familiar to the IETF.

In particular, this document is limited to the areas of concern of the IETF. That is, it is limited to layer 1 networks that utilize IP as the underlying support for their control plane.

The intention of this document is to present the requirements and architectures developed within the ITU-T to the IETF for better understanding and further cooperation between the two bodies.

Some work related to L1VPN solution space has already been done within the IETF. This document intends to set a framework of requirements and architectures into which all possible solutions can fit.

4. Motivations

In this discussion many merits and motivations may be taken for granted.

The general benefits and desirability of VPNs has been described many times and in many places. This document does not dwell on the merits of VPNs as such, but focuses entirely on the applicability of the VPN concept to layer 1 networks.

Similarly, the utility and value of a control plane for the

configuration, management and operation of a layer 1 network is well-rehearsed.

4.1 Basic Layer 1 Services

Basic layer 1 services may be characterized in terms that include:

- Connectivity: Between a pair of CEs.
- Capacity: For example, the bit rate for a TDM service or the capacity of a lambda.
- Transparency: For example, for an SDH network, overhead transparency.
- Availability: The percentage of time that the quality of the service meets the agreed criteria. To achieve the required level of availability for the customer connections the service provider's network may use restoration or protected resources.
- Performance: For example, the number of error-seconds per month.

The layer 1 services may be categorized based on the combination of connectivity features (U-plane) and service control capability features (C-plane) available to the customer. A CE is associated with the service interface between a customer site and the network, and the categorization can be seen in the context of this service interface as follows.

1. A single connection between a pair of CEs.

   - Static Service
     The classic private line service achieved through a permanent connection.

   - Dynamic Service
     Either a switched connection service, or a customer-controlled soft permanent connection service

2. Multiple connections among a set of CEs.

   - Static Service
     A private network service consisting of a mesh of permanent connections.

   - Dynamic Service
     A dynamic private network service consisting of any combination of switched connection services and customer-controlled soft permanent connection services.

For both service types, connections are point-to-point, and can be permanent, soft-permanent, or switched. For a static service, the network is responsible for the management of both the network

infrastructure and the end user connections. For dynamic services, the network is only responsible for the configuration of the infrastructure; end user connections are established dynamically by the network.

Note that the ITU-T allows the second categorization of service type to embrace a variety of C-plane types.

4.1.1 L1VPN for Dynamic Layer 1 Provisioning

Private network services in the second category (above) can be enhanced so that multiple private networks are supported across the layer 1 network as virtual private networks. These are Layer 1 Virtual Private Networks (L1VPNs).

Compared to the first type of service, the L1VPN service has features such as a separate policy per VPN, and distribution of information about which CEs can participate in which VPNs.

4.2 Merits of L1VPN

4.2.1 Customer Merits

From the customer's perspective, there are two main benefits to a L1VPN. These benefits apply over and above the advantages of access to a dynamically provisioned network.

- The customer can outsource the direct management of an optical network by placing the VPN management in the control of a third party. This frees the customer from the need to configure and manage the connectivity information for the CEs that participate in the VPN.

- The customer can make small-scale use of an optical network. So, for example, by sharing access to the optical network with many other users, the customer sites can be connected together across the optical network without bearing the full cost of deploying and managing the optical network.

To some extent, the customer may also gain from the provider's benefits (see below). That is, if the provider is able to extract more value from the layer 1 network, and provide better differentiated services, the customer will benefit from lower priced services that are better tailored to the customer's needs.

4.2.2 Provider Merits

The provider benefits from the customer's perception of benefits.

In particular, the provider can build on dynamic, on-demand services by offering new VPN services and off-loading the CE-to-CE configuration requirements from the customers

Additionally, a more flexible VPN structure applied to the optical network allows the provider to make more comprehensive use of the spare (that is, previously unused) resources within the network. In particular, since the PE could be responsible for routing the connection through the optical network, the optical network can reclaim control of how resources are used and adjust the paths so that optimal use is made of all available resources.

4.3 L1VPN Deployment Scenarios

4.3.1 Multi-Service Backbone

A multi-service backbone is characterized in terms such that one service department of a carrier receiving the carrier's L1VPN service provides different kinds of higher-layer service. The customer receiving the L1VPN service (i.e. each service department) can offer its own services whose payloads can be any layer (e.g. ATM, IP, TDM). From the L1VPN service provider point of view, these services are not visible and are not part of the L1VPN service. That is, the type of service being carried within the L1 payload is not known by the service provider.

The benefit is that the same L1 core network resources are shared by multiple services. A large capacity backbone network (U-Plane) can be built economically by having the resources shared by multiple services usually with flexibility to modify topologies, while separating the control functions. Thus, each customer can select a specific set of features that are needed to provide their own service.

4.3.2 Carrier's Carrier

A carrier's carrier is characterized in terms such that one carrier that receives another carrier's L1VPN service provides its own services. In this scenario, two carriers may be in different organizations (or may be separately managed within the same organization). It is, therefore, expected that the information provided at the service demarcation points is more limited than in the multi-service backbone case. Similarly, more less control of the L1VPN service is given at the service demarcation points. For example, customers of L1VPN service receive:

- more limited view of L1VPN service provider network

- more limited control over L1VPN service provider network.

One of the merits is that each carrier can concentrate on a specific service. For example, the customer of the L1VPN service may focus on L3 services, e.g. providing secure access to the Internet, leaving the L1VPN provider to focus on the L1 service, i.e. providing a long haul bandwidth between cities. The L1VPN customer can construct its own network using L1 resources supplied by the L1VPN provider, usually with flexibility to modify topologies, and utilize dedicated C-Plane functionalities.

4.3.3 L1 Resource Trading

In addition to the scenarios where the second tier service provider is using a single core service provider as mentioned above, it is possible for the second tier provider to receive services from more than one core service provider. In this scenario, there are some benefits for the second tier service provider such as dynamic carrier selection based on the price and route redundancy.

The second tier service provider can support a function that enables a L1 resource trading service. Using resource information published by its core service providers, a second tier service providers can decide how to best use those providers. For example, if one core service provider is no longer able to satisfy requests for service, an alternate service provider can be used. Or the second tier service provider could choose to respond to price changes over time.

Another example of second tier service provider use is to reduce exposure to failures in each provider (improve availability).

4.3.4 Inter-SP L1 VPN

In addition to the scenarios where a single connection between two CEs is routed over a single service provider, it is possible that a connection is routed over multiple service providers. This service scenario is called Inter-SP L1VPN.

This scenario can be used to construct a single L1VPN from services provided by multiple regional providers. There could be a variety of business relationships among providers and customers.

5. Reference models

Figure 5.1 describes the L1VPN reference model.

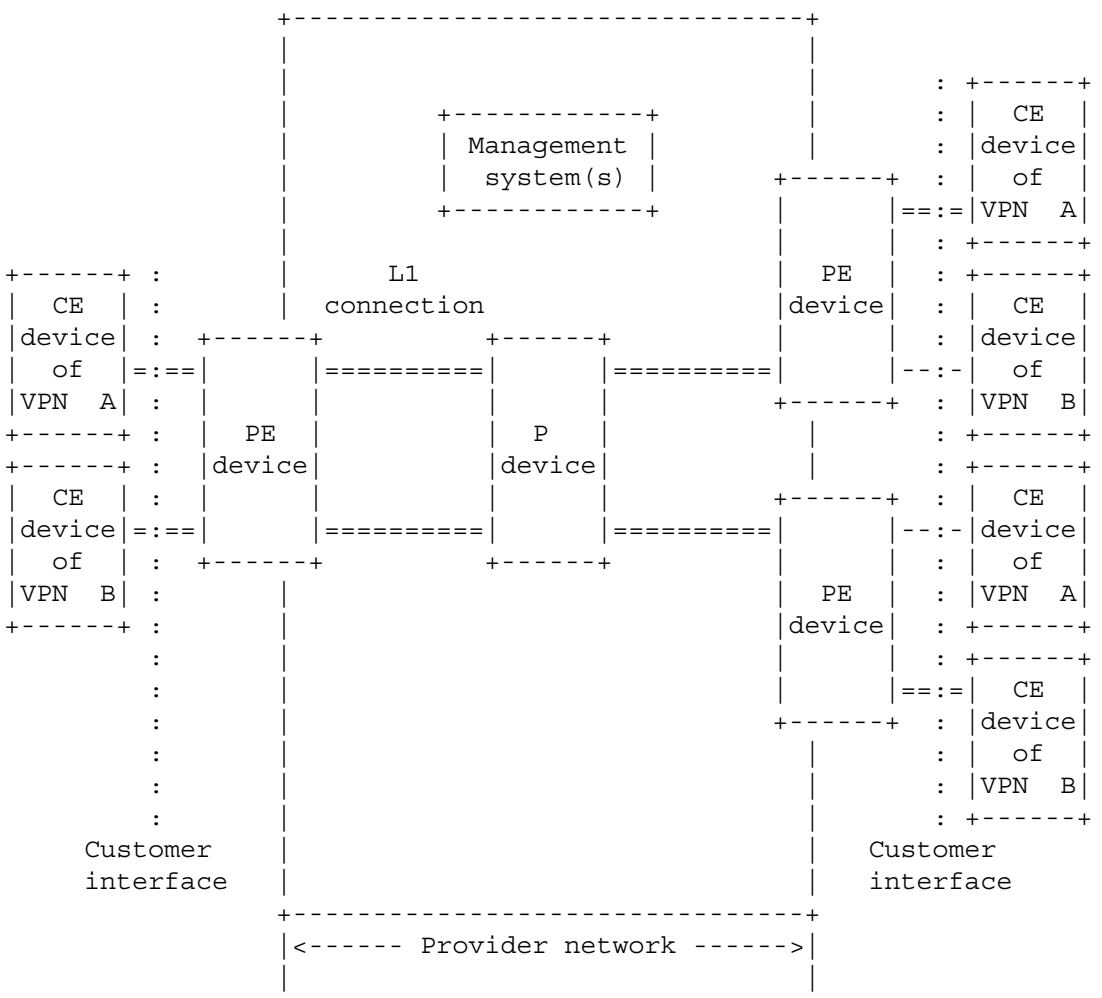

Figure 5.1: L1VPN reference model

In L1VPN, L1 connections are provided between CE's physical interfaces within the same VPN. In Figure 5.1, a connection is provided between the lefthand CE of VPN A and the upper righthand CE of VPN A, and another connection is provided between the lefthand CE of VPN B and lower righthand CE of VPN B (shown as "=" mark).

5.1 CE/PE/P Terminology

In the reference model, the following three types of network devices are described. Note that these terminologies are from PPVPN works [PPVPN-TERM].

o CE (Customer Edge) device

A CE device is a customer device that receives L1VPN service from the

provider. A CE device is connected to at least one PE device. A CE device can be a variety of devices, for example, TDM cross connect, router and L2 switch. A CE device may also be attached to one or more C devices on the customer site.

o PE (Provider Edge) device

A PE device is a provider device that provides L1VPN service to the customer. A PE device is connected to at least one CE device. A layer 1 PE device is a TDM or optical cross connect. Or a PE device may be an EPL (Ethernet Private Line) type of device, that maps Ethernet frames on L1 connections.

o P (Provider) device

A P device is a provider device, which is connected only to other provider devices (P or PE devices). A layer 1 P is a TDM or optical cross connect.

5.2 Customer/Provider Terminology

In this document, the following two types of administrative entities are described.

o Customer

A Customer has authority over a set of CE devices within the same VPN (e.g. the owner of CE devices). Note that a customer may outsource the management of CE devices to other organizations, including to the provider itself.

o Provider

A Provider has authority over the management of the provider network.

5.3 Management Systems

As shown in the reference model, a provider network may contain a management system(s). A management system(s) may support functions including provisioning, monitoring, billing and recording. Provider's management system(s) may also communicate with customer's management system(s) in order to provide services.

6. Generic Service Description

This section describes generic service descriptions. More detailed service description is described as specific service models in section 7.

6.1 CE Construct

- The CE device may contain multiple VPN instances.
- CE-PE physical links (between physical interfaces) may be shared by multiple VPNs. (assuming that each CE-PE logical link maps one-to-one to a VPN, and maps one-to-one or many-to-one to the physical link)

6.2 Generic Service Features

L1VPN has following two generic service features.

- Connectivity restriction: Layer 1 connectivity is provided to a limited set of CE's physical interfaces. (This set forms the L1VPN membership.)
- Per VPN control and management: Some level of control and management capability is provided to the customer. Details differ depending on service models described in section 7.

7. Service Models

This section describes Layer 1 VPN service models, derived from the generic service description presented above, that can be supported by Generalized MPLS (GMPLS) protocols enabled networks.

Such layer 1 networks are managed and controlled using GMPLS as described in [RFC3471] and [RFC3473]. It must be understood that meeting the requirements set out in this document may necessitate modifications to the existing GMPLS protocols both for the control plane within the layer 1 network and for service provisioning at the edge of the network between the CE and PE devices. A CE and a PE are connected by one or more GMPLS Traffic Engineering (TE) links as defined in [GMPLS-ROUTING]. The ends of each link are usually represented as GMPLS-capable interfaces.

7.1 Management-based Service Models

Figure 7.1 describes the management-based service models.

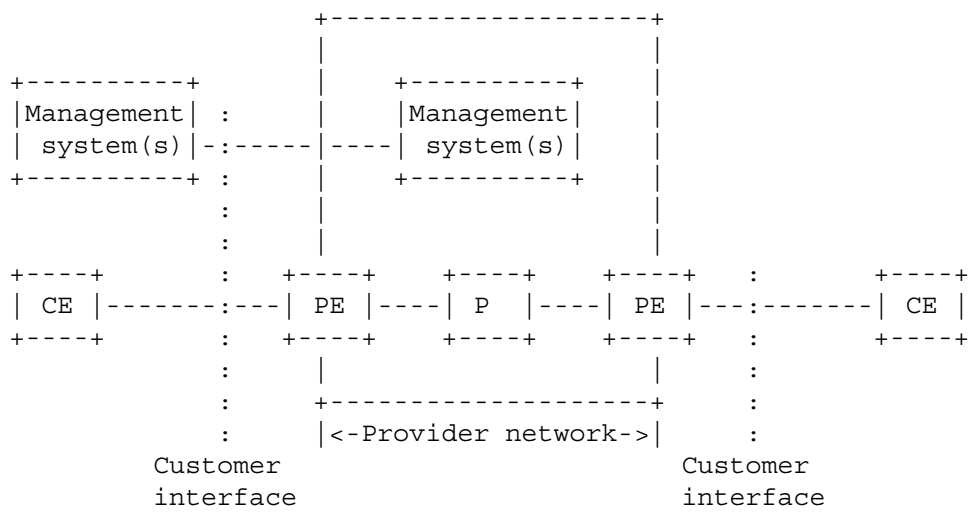

Figure 7.1: Management-based service models

In this service model, customer's management system(s) and provider's management system(s) communicate with each other. Customer's management system(s) access provider's management system(s) to request L1 connection setup/deletion between a pair of CEs. Customer's management system(s) may obtain additional information, such as resource availability information and monitoring information, from provider's management system(s). There is no control message exchange between a CE and PE.

The provider network may be based on GMPLS. In this case, existing protocols to meet this service model may need to be extended (e.g. to support soft permanent connections). However, interfaces between management systems are not within the scope of this document. Interfaces between management systems and network devices may need to be studied further.

7.2 Signaling-based Service Models (Overlay Service Models)

Figure 7.2 describes the signaling-based service models.

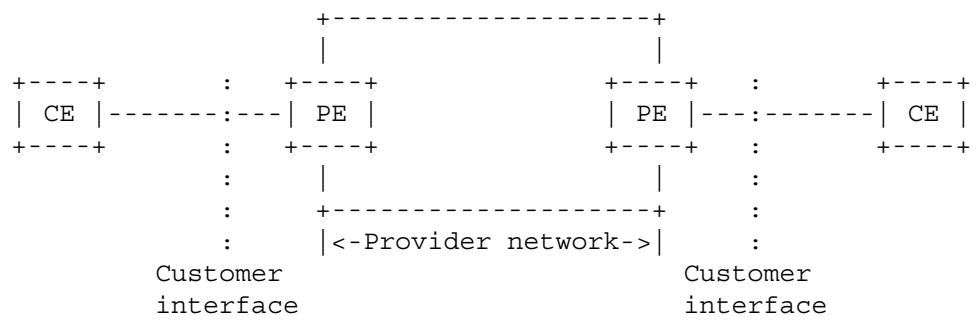

Figure 7.2: Signaling-based service models

In this service model, the customer interface is based on GMPLS UNI overlay. The CE requests L1 connection setup/deletion to a remote CE. There is no routing between a CE and PE. The CE does not receive routing information of remote CE sites, nor routing information of the provider network. The CE's interface may be assigned a public or private address, that designates connection end points.

A CE may optionally receive a list of TE link addresses to which it can request a connection (a list of addresses within the same VPN) (overlay ext.).

Note that in addition, there may be communication between customer's management system(s) and provider's management system(s) in order to provide detailed monitoring and fault information etc. to customers.

7.3 Signaling and Routing Service Models

In this service model, the customer interface is based on GMPLS signaling and routing. The CE requests L1 connection setup/deletion to a remote CE. There is routing between a CE and PE, or more precisely between a CE and the VPN routing context instantiated on the PE. By using traffic engineering-based routing information obtained, customers can use traffic engineering capabilities within his portion of the provider network.

For example, a customer can setup two disjoint connections between a pair of CEs. Another example is that a customer can request a connection between a pair of devices within CE sites, and not necessarily between CEs.

Note that in addition, there may be communication between customer's management system(s) and provider's management system(s) in order to provide detailed monitoring and fault information etc. to customers.

There are two more detailed signaling and routing service models, virtual link models and per VPN peer models.

7.3.1 Virtual Link Models

Figure 7.3 describes the virtual link models.

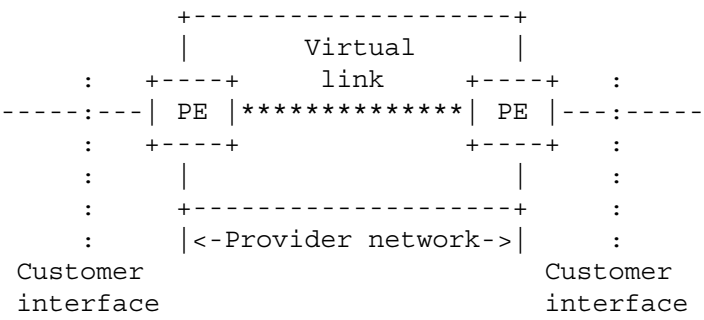

Figure 7.3: Virtual link models

In this service model, a virtual link is constructed between PEs. A virtual link is a TE link connecting two devices where a direct physical link does not exist. The CE receives routing information of PE-CE links, remote CE sites, as well as virtual links. A virtual link's TE attributes may be derived from physical links within the provider network.

As a special case, the provider may choose not to advertise virtual links to customers. The CE receives routing information of CE-PE links and remote CE sites only. This corresponds to advertising a whole provider network as one node, i.e. Generalized Virtual Private Cross-Connect (GVPXC) [GVPN].

7.3.2 Per VPN Peer Models

Figure 7.4 describes per VPN peer models.

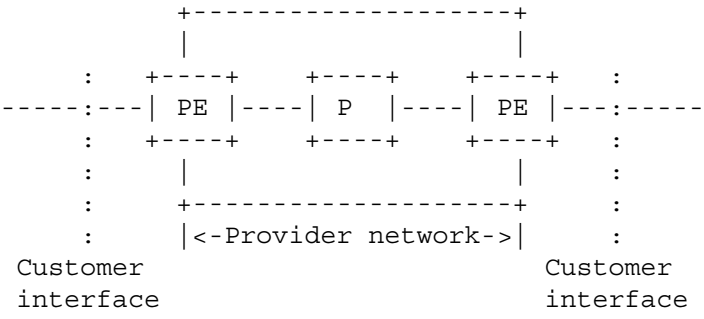

Figure 7.4: Per VPN peer models

In this service model, the provider partitions TE links within the

provider network per VPN, and discloses per VPN TE link information to corresponding CEs. As such, a CE receives routing information of PE-CE links, remote CE sites, as well as partitioned portions of the provider network.

Note that PEs may advertise abstracted routing information of the provider network to CEs, for administrative purpose, as well as for excluding "unnecessary information".

Note that when inter-area/AS solutions are available, it may be valuable to consider inter-area/AS interfaces to be the basis for customer interface.

8. Service Models and Service Requirements

Service models mentioned in section 7 is related to what information is exchanged between the CE and the PE. In addition, service models vary depending on how U-Plane resources are allocated for each VPN. Specifically, service models are described by combining following service requirements.

NOTE: Later version of this document may include more detailed service requirements from Y.1312.

o U-Plane resource allocation

- Shared or dedicated : Shared means that provider network physical links are shared by multiple VPNs.(Physical links are allocated to each VPN when connection is requested, and physical links allocated to one VPN at one time can be allocated to another VPN at another time.) Dedicated means that provider network physical links are partitioned per VPN. (Physical links allocated to one VPN can not be used by other VPNs.)

o Information exchanged between the CE and the PE

- Signaling
- Membership information : A list of TE link addresses within the same VPN (connection end points)
- Customer network routing information
- Provider network routing information

Table 1 shows combination of service requirements and service models.

	U-Plane shared	U-Plane dedicated
Signaling	Overlay	Overlay
Signaling + Membership information	Overlay (ext.)	Overlay (ext.)
Signaling + Membership information + Customer network routing information (Note1)	Virtual link	Virtual link
Signaling + Membership information + Customer network routing information + Provider network routing information	Not applicable	Per VPN peer

Table 1: Combination of service requirements and service models

Note1: In virtual link models, to be precise, PE-PE virtual link information, which is part of provider network routing information, is advertised from a PE to CE.

9. Security Considerations

TBD

10. Acknowledgements

The material in this document is based on the work of the ITU-T Study Group 13, and it has been submitted to meet the 00 ID deadline. Further related communication will be provided after the closure of Study Group 13 February 2004 meeting in the form of liaison statement.

11. Intellectual Property Consideration

result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

12. Normative References

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", RFC 2026, October 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

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

13. Informative References

[Y.L1VPNARCH] Y.l1vpnarch - Layer 1 Virtual Private Network service and network architectures, ITU-T draft Recommendation.

[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-UNI] Swallow, G., et al., "GMPLS UNI: RSVP Support for the Overlay Model", draft-ietf-ccamp-gmpls-overlay, work in progress.

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

[L2VPN-FRAME] Andersson, L., and Rosen, E. (editors), "L2VPN Framework", draft-ietf-l2vpn-l2-framework, work in progress.

[L3VPN-FRAME] Callon, R., et al., "A Framework for Layer 3 Provider Provisioned Virtual Private Networks, draft-ietf-l3vpn-framework, work in progress.

[PPVPN-TERM] Andersson, L., and Madsen, T., "PPVPN terminology", draft-andersson-ppvpn-terminology, 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.

14. Authors' Addresses

Raymond Aubin
Nortel Networks
P O Box 3511 Station C
Ottawa, ON K1Y 4H7 Canada
Phone: +1 (613) 763 2208
Email: aubin@nortelnetworks.com

Marco Carugi
Nortel Networks S.A.
Parc d'activites de Magny-Les Jeunes Bois CHATEAUFORT
78928 YVELINES Cedex 9 - FRANCE
Email: marco.carugi@nortelnetworks.com

Ichiro Inoue
NTT Network Service Systems Laboratories, NTT Corporation
3-9-11, Midori-Cho
Musashino-Shi, Tokyo 180-8585 Japan
Phone: +81 422 59 6076
Email: inoue.ichiro@lab.ntt.co.jp

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

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. Full Copyright Statement