Network Working Group
Internet Draft
Proposed Status: Informational
Expires: December 2005

Tomonori Takeda (Editor)
NTT

June 2005

Framework and Requirements for Layer 1 Virtual Private Networks
draft-takeda-l1vpn-framework-04.txt

Status of this Memo

Copyright Notice

Abstract

This document provides a framework and service level requirements 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.

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 [RFC4026].

In addition, following new terms are used within this document.

- Virtual link: A provider network TE link advertised to customers in routing information for purposes which include path computation. A data link may or may not exist between the two end points of a virtual link.

- Virtual node: A provider network logical node advertised to customers in routing information. A virtual node may represent a single physical node, or multiple physical nodes and links.

- VPN end point: A CE's data plane interface, which is connected to a PE device, and which is part of the VPN membership. Note that a data plane interface is associated with a TE link end point. For example, if a CE router's interface is a channelized interface (defined in SONET/SDH), a channel in the channelized interface can be a data plane interface.

- VPN connection (or connection in the L1VPN context): A connection between a pair of VPN end points. Note that in some scenarios, a connection may be established between a pair of Cs (customer

devices), using this CE-CE VPN connection as a segment or forwarding adjacency.

Note that following terms are aligned with PPVPN terminology [RFC4026], and in this document, have a meaning in the context of L1VPNs, unless otherwise specified.

- 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 does not have to have the capability to switch at layer 1, but it must be capable of receiving a layer 1 signal and either switching it or terminating it with adaptation. A CE device may also be attached to one or more C devices on the customer site.

- 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 Time Division Multiplex (TDM) switch, an Optical Cross-Connect (OXC), a Fiber Switch (FXC), or a PE device may be an EPL (Ethernet Private Line) type of device, that maps Ethernet frames onto layer 1 connections.

- 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 switch, OXC, or FXC.

- 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.

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

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), Time Division Multiplex (TDM) capable switches, or Fiber Switches (FXCs) 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 [RFC4026] 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 (CE and PE devices). It is at the interface between CE and PE devices that the L1VPN service is provided.

Note that one of the fundamental differences between L1VPNs and L2/L3 VPNs is that in L1VPNs data plane connectivity does not guarantee control plane connectivity (and vice versa). CE-PE control plane connectivity is essential, and CE-CE data plane connectivity is maintained by signaling mechanisms based on this control plane connectivity. The provision of CE-CE control plane connectivity over the provider network is also a unique aspect of the L1VPN services, by which control packets can be exchanged between CEs over the control plane of the provider network.

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).

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

Current E-NNIs include features to exchange routing information, as well as to facilitate requests for end-to-end services.

The UNIs and E-NNIs, 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 (i.e., a control link shared by multiple VPNs) at a UNI/E-NNI, and these messages must be disambiguated to determine the L1VPN to which they apply.

Furthermore, there is no clear defined way to restrict connectivity among CEs (or over a UNI/E-NNI). In addition, E-NNIs allow routing information exchange, but there is no clear defined way to allow limited routing information exchange (i.e., a specific set of routing information is distributed to a specific set of CEs).

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.

This document presents the requirements and architectures developed within the ITU-T for better understanding within the IETF and to further cooperation between the two bodies.

Some work related to the L1VPN solution space has already been done within the IETF. This document sets a framework of requirements and architecture into which 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 offered 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: The quality of the service delivered to customers, e.g., the number of error-seconds per month.

The layer 1 services may be categorized based on the combination of connectivity features (data plane) and service control capability features (control plane) available to the customer. A CE is associated with the service interface between a customer site and the provider 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 management plane of the provider network is responsible for the management of both the network infrastructure and the end-user connections. For dynamic services, the management plane of the provider network is only responsible for the configuration of the infrastructure; end-user connections are established dynamically via the control plane of the provider network upon customer request.

Note that the ITU-T allows the second categorization of service type to embrace a variety of control 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). Note the first category (above) would include L1VPNs with only two CEs as a special case.

Compared to the first category of service, the L1VPN service has features such as connectivity restriction, a separate policy per VPN, and distribution of membership information.

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

In large carrier networks providing various kinds of service, it is often the case that multiple service networks are supported over a shared transport network. L1VPNs are expected to support this type of network architecture. Namely, by applying L1VPNs, multiple internal service networks (which may be managed and operated separately) can be supported over a shared layer 1 transport network controlled and managed by GMPLS. In addition, L1VPNs can support capabilities to offer innovative services to external clients.

Some more specific deployment scenarios are as follows.

4.3.1 Multi-Service Backbone

A multi-service backbone is characterized in terms such that each service department of a carrier that receives the carrier's L1VPN service provides a different kind 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's 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 layer 1 payload is

not known by the service provider.

The benefit is that the same layer 1 core network resources are shared by multiple services. A large capacity backbone network (data 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.

Note that it is also possible to control and manage these service networks and the layer 1 core network by using GMPLS as a unified control plane, instead of using L1VPNs. However, using L1VPNs is beneficial in the following points.

- Independent address space for each of the service networks.
- Network isolation (topology information isolation, fault isolation among service networks).
- Independent layer 1 resource view for each of the service networks.
- Independent policies that could be applied for each of the service networks.

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, less control of the L1VPN service is given at the service demarcation points. For example, customers of an L1VPN service receive:

- A more limited view of the L1VPN service provider network.
- More limited control over the 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 layer 1 service, e.g., providing a long haul bandwidth between cities. The L1VPN customer can construct its own network using layer 1 resources supplied by the L1VPN provider, usually with flexibility to modify topologies, and utilize dedicated control plane functionalities.

4.3.3 Layer 1 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 route redundancy and dynamic carrier selection based on the price.

The second tier service provider can support a function that enables a layer 1 resource trading service. Using resource information published by its core service providers, a second tier service provider can decide how to best use the core 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 (i.e., to improve availability).

4.3.4 Inter-SP L1VPN

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.

4.3.5 Other Scenarios

There could be more complex L1VPN scenarios such as the case where one or both CE-PE links of a L1VPN connection are not static, but are based on L1VPN connections in their own right provided by the same or different L1VPN service provider.

5. Reference Models

Figure 5.1 describes the L1VPN reference model.

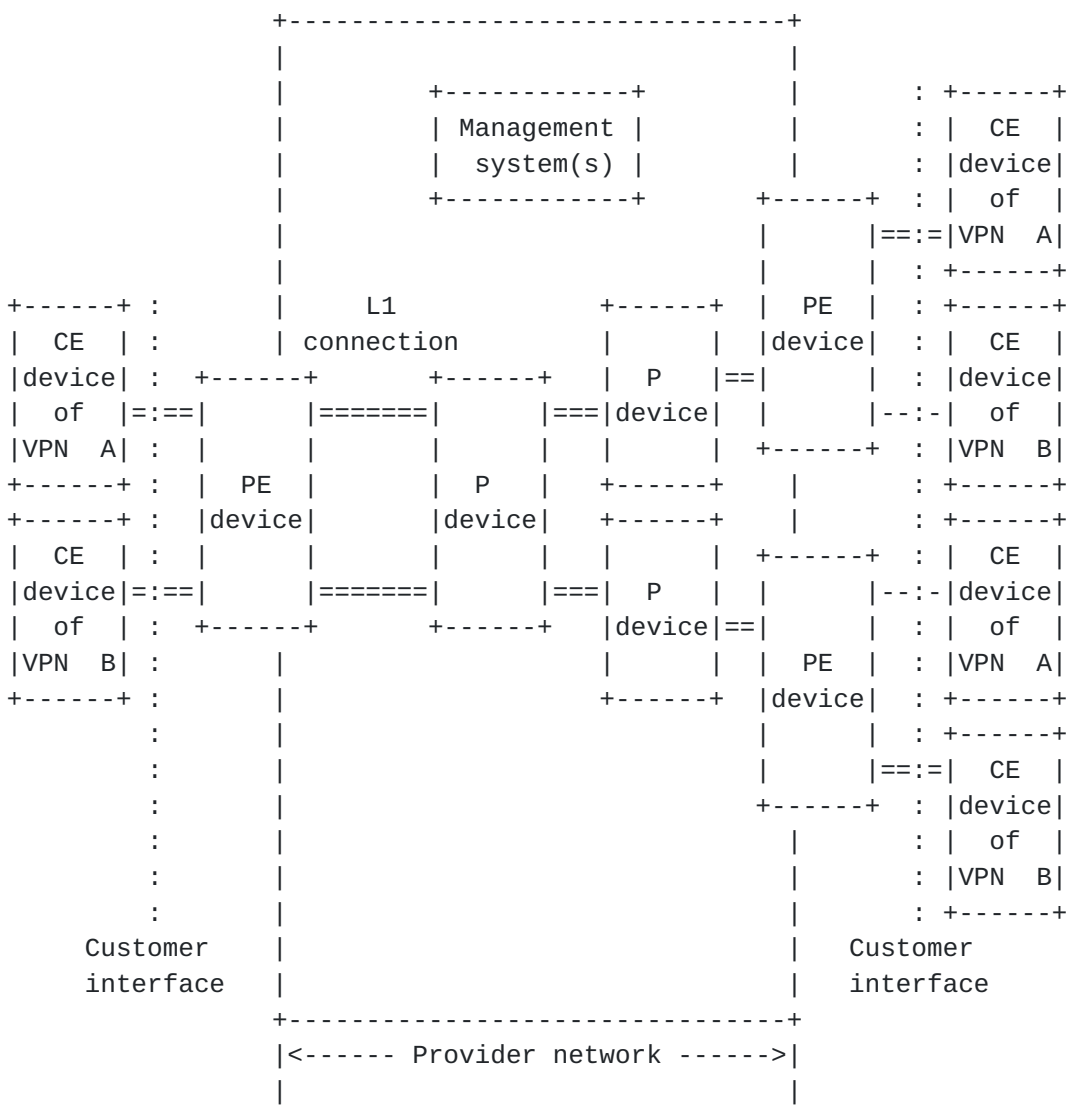

Figure 5.1: L1VPN reference model

In a L1VPN, layer 1 connections are provided between CEs' data plane interfaces within the same VPN. In Figure 5.1, a connection is provided between the left-hand CE of VPN A and the upper right-hand CE of VPN A, and another connection is provided between the left-hand CE of VPN B and lower right-hand CE of VPN B (shown as "=" mark). These layer 1 connections are called VPN connections.

5.1 Management Systems

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

6. Generic Service Description

This section describes generic L1VPN services. More detailed service descriptions are provided through specific service models in section 7.

6.1 CE Construct

- The CE device may support more than one customer VPN.
- CE-PE data plane links (between data plane interfaces) may be shared by multiple VPNs.

Note that it is necessary to disambiguate control plane messages exchanged between CE and PE if the CE-PE relationship is applicable to more than one VPN. This makes it possible to determine to which VPN such control plane messages apply. Such disambiguation might be achieved by allocating a separate control channel to each VPN (either using a separate physical channel, a separate logical channel (e.g., IP tunnel), or using separate addressing) or by extending the signaling and routing protocols to allow them to identify the correct VPN.

6.2 Generic Service Features

L1VPN has the following two generic service features.

- Connectivity restriction: Layer 1 connectivity is provided to a limited set of CEs' data plane interfaces, called VPN end points. (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 that can be supported by Generalized MPLS (GMPLS) protocols enabled networks. These models are derived from the generic service description presented above.

Such layer 1 networks are managed and controlled using GMPLS signaling as described in [RFC3471] and [RFC3473], and GMPLS routing as described in [GMPLS-ROUTING]. 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 (CE and PE devices). Such modifications are discussed in [L1VPN-APP]. A CE and a PE are connected by one or more

data links. The ends of each link are usually represented as GMPLS-capable interfaces.

Note that in this document, service models are classified by the semantics of information exchanged over the customer interface.

7.1 Management-based Service Model

Figure 7.1 describes the management-based service model.

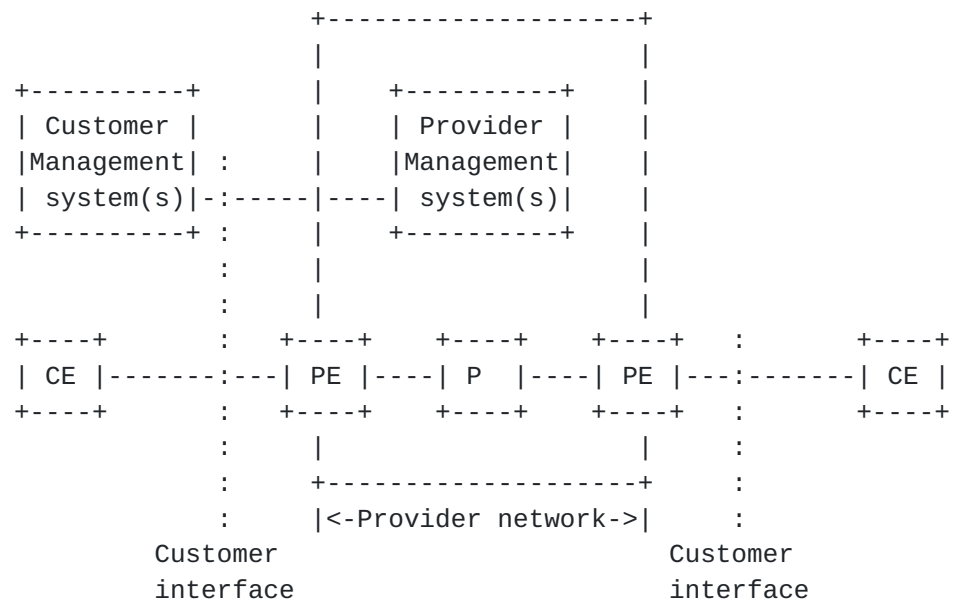

Figure 7.1: Management-based service model

In this service model, customer management systems and provider management systems communicate with each other. Customer management systems access provider management systems to request layer 1 connection setup/deletion between a pair of CEs. Customer management systems may obtain additional information, such as resource availability information and monitoring information, from provider management systems. 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 controlled by GMPLS may need to be studied further in [L1VPN-APP].

7.2 Signaling-based Service Model (Basic Mode)

In this service model, the CE-PE interface's functional repertoire is

limited to path setup signalling only. The provider's network is not involved in distribution of customer network's routing information.

7.2.1 Overlay Service Model

Figure 7.2 describes the Overlay service model.

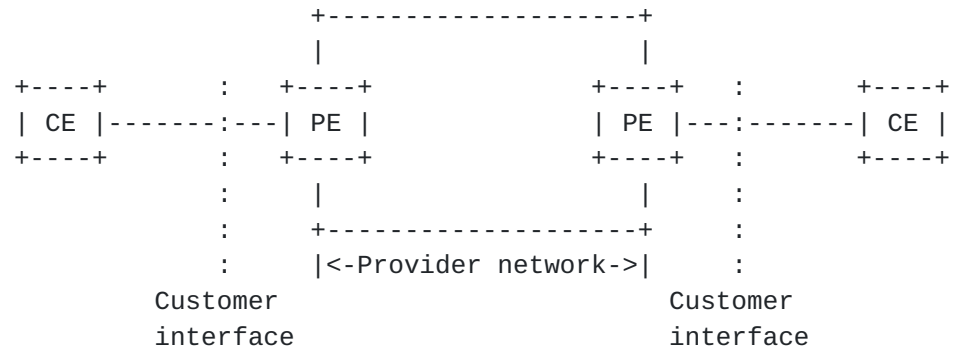

Figure 7.2: Overlay service model

In this service model, the customer interface is based on the GMPLS UNI Overlay [GMPLS-UNI]. The CE requests layer 1 connection setup/deletion to a remote CE. There is no routing between a CE and PE. The CE does not receive routing information from remote customer sites, nor routing information about the provider network. The CE's interface may be assigned a public or private address, that designates VPN end points.

There are various ways that customers perceive the provider network. In one example, the whole provider network may be considered as one node - the path specified and recorded in signaling messages reflects this. Note that this is distinct from the Virtual Node service model described in section 7.3.2 because such a model requires that the network is represented to the VPN sites as a virtual node - that is, some form of routing advertisement is implied, and this is not in scope for the Signaling-based service model.

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

7.3 Signaling and Routing Service Model (Enhanced Mode)

In this service model, the CE-PE interface provides the signaling capabilities as in the Basic Mode, plus permits limited exchange of information between the control planes of the provider and the customer to help such functions as discovery of reachability information in remote sites, or parameters of the part of the provider's network dedicated to the customer.

By allowing CEs to obtain reachability information, a so-called N-square routing problem could be solved [GVPN].

In addition, by using the received traffic engineering-based routing information, a customer can use traffic engineering capabilities within his portion of the provider network. For example, a customer can set up two disjoint connections between a pair of CEs. Another example is that a customer can request a connection between a pair of devices within customer sites, and not necessarily between CEs, with more effective traffic engineering.

As such, the customer interface is based on GMPLS signaling and mechanisms to exchange reachability/TE information. Typically, a routing protocol is used between a CE and PE, or more precisely between a CE and the VPN routing context instantiated on the PE. Link state routing information would be needed to implement the above two example scenarios. Some scenarios may be satisfied with reachability routing information only.

Note that this service model does not preclude the use of mechanisms other than routing protocols to exchange reachability/TE information. Details need to be studied in [L1VPN-APP].

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

Four specific types of the Signaling and Routing service model are the Overlay Extension service model, the Virtual Node service model, the Virtual Link service model and the Per-VPN Peer service model, depending on how customers perceive the provider network in routing and signaling.

7.3.1 Overlay Extension Service Model

This service model is a slight extension from the Overlay service model. In this service model, a CE receives a list of TE link addresses to which it can request a VPN connection (a list of addresses within the same VPN). This may include additional information concerning these TE links (e.g., switching type). Note, in the Overlay Extension service model, information a CE can receive is limited to information about the CE-PE TE link. Mechanisms other than routing could be used to exchange reachability/TE information between the CE and the PE.

7.3.2 Virtual Node Service Model

Figure 7.3 describes the Virtual Node service model.

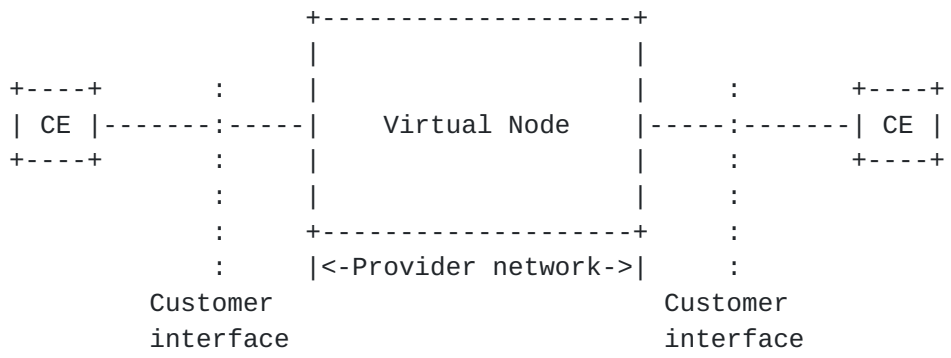

Figure 7.3: Virtual Node service model

In this type of service model, the whole provider network is represented as a virtual node (defined in section 2). The customer perceives the provider network as one single node, i.e., a Generalized Virtual Private Cross-Connect (GVPXC) [GVPN]. The CE receives routing information about CE-PE links and remote customer sites.

Note that in this service model, there must be one single virtual node, and this virtual node must be connected with every CE in the VPN.

7.3.3 Virtual Link Service Model

Figure 7.4 describes the Virtual Link service model.

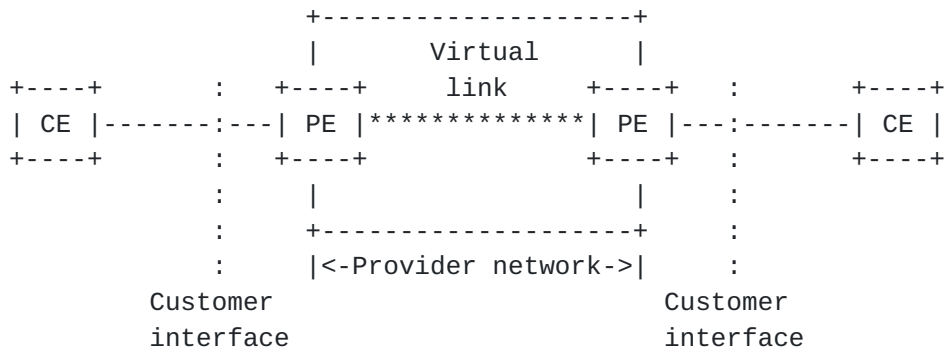

Figure 7.4: Virtual Link service model

In this service model, a virtual link is constructed between PEs. For the definition of a virtual link, please refer to terminology in section 2. The CE receives routing information about CE-PE links, remote customer sites, as well as virtual links. A special property of the virtual links used in this service model is that the provider network allocates data plane link resources for the exclusive use of each virtual link. The TE attributes of a virtual link are determined

according to data plane link resources allocated to this virtual link. Virtual links are an abstraction of the provider network to customers for administrative purposes as well as to exclude "unnecessary information".

Note that in this service model, both end points of each virtual link must be a PE device.

7.3.4 Per-VPN Peer Service Model

Figure 7.5 describes the Per-VPN Peer service model.

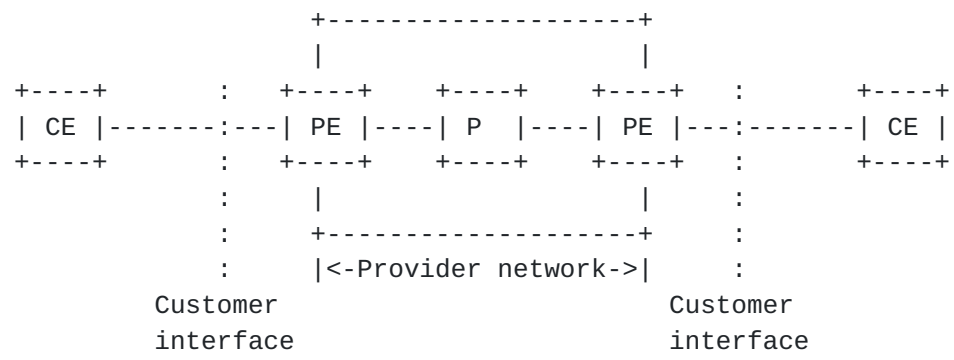

Figure 7.5: Per-VPN Peer service model

In this service model, the provider partitions the 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 about CE-PE links, remote customer sites, as well as partitioned portions of the provider network.

Note that PEs may advertise abstracted routing information about the provider network to CEs for administrative purpose as well as to exclude "unnecessary information". In other words, virtual links may be constructed between two nodes where direct data links do not exist, or virtual nodes may be constructed to represent multiple physical nodes and links.

In the Per-VPN Peer service model, at least one virtual node corresponding to P devices (one single P or a set of Ps) must be visible to customers.

8. Service Models and Service Requirements

The service models mentioned in section 7 are related to which information is exchanged between CE and PE. In addition, service models differ in how data plane resources are allocated for each VPN.

Note that in the ITU-T documents, the term "U-Plane" is used instead

of "data plane".

o Data plane resource allocation

- Shared or dedicated:

  Shared means that provider network data plane links are shared by multiple (i.e., any or a specific set of) VPNs. (Data plane links are dynamically allocated to a VPN when a VPN connection is requested, and data plane links allocated to one VPN at one time can be allocated to another VPN at another time.)

  Dedicated means that provider network data plane links are partitioned per VPN. (Data plane links are statically allocated to one VPN and can not be used by other VPNs.)

o Information exchanged between CE and PE

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

Table 1 shows combination of service requirements and service models.

	Data plane shared	Data plane dedicated
Signaling	Overlay	Overlay
Signaling + Membership information	Overlay Extension	Overlay Extension
Signaling + Membership information + Customer network routing information	Virtual Node	Virtual Node
Signaling + Membership information + Customer network routing information + Provider network routing information	Not applicable	Virtual Link Per-VPN Peer

Table 1: Combination of service requirements and service models

As described in previous sections, the difference between the Virtual Link service model and the Per-VPN Peer service model is whether customers have visibility of P devices. In the Virtual Link service model, the end points of virtual links must be PE devices, thus P devices are not visible to customers. In the Per-VPN Peer service model, at least one virtual node corresponding to P devices (one single P, or a set of Ps) is visible to customers.

Note that when provider network routing information is provided to customers, customers must be able to specify explicit links for a VPN connection over the provider network.

8.1 Detailed Service Level Requirements

More detailed service requirements are provided below. They are generally common to the various service models, except where indicated.

- Selection of layer 1 class of service: Customers MAY be allowed to specify a layer 1 class of service (e.g., availability level) for a VPN connection.

- Reception of performance information: Customers MAY be allowed to receive performance information for their VPN connections (e.g., performance monitoring data). When data plane links are dedicated, customers MAY be allowed to receive performance information for links dedicated to them.

- Reception of fault information: Customers MAY be allowed to receive fault information for their VPN connections (e.g., failures, data plane alarms, rejections). When data plane links are dedicated, customers MAY be allowed to receive fault information for links dedicated to them.

- Reception of connection information: Customers MAY be allowed to receive information for current VPN connections.

- Reception of accounting information: Customers MUST be able to receive accounting information for each VPN.

- Specification of policy: Customers MAY be allowed to specify policies (e.g., path computation policies, recovery policies including parameters) for each VPN.

- Security: The communication between the customer and the provider MUST be secure. Further details are described in section 9.

- Filtering: Unnecessary information (e.g., information concerning

other VPNs) MUST NOT be provided to each customer. This applies particularly to Signaling and Routing service models, but is also relevant to Signaling-based service models and to Management-based service models. Further details are described in section 9.

- Requests/indications for arbitrary CE-CE control plane information delivery. All models that support routing exchanges MAY support the exchange of arbitrary CE-CE control plane information passed from CE to PE within routing protocol messages and delivered from PE to CE at the other side of the core network. In addition, some signaling models MAY allow directed signaling message exchange between CEs for hierarchical or stitched LSPs over CE-CE LSP.

9. Security Considerations

Security is clearly one of the essential requirements in L1VPNs. In this section, key security requirements are highlighted. Security considerations for L3VPNs and L2VPNs are described in existing documents, such as [L3VPN-FRAME] and [L2VPN-FRAME]. These security considerations should also be applied in L1VPNs, and these aspects are described in this section. In addition, there are some specific security considerations for L1VPNs, such as connectivity restriction and shared control links.

This section first describes types of information to be secured. Then, security features or aspects are described. Finally, some considerations concerning scenarios where security mechanisms are applied is described.

9.1 Types of Information

It MUST be possible to secure the information exchanged between the customer and the provider. This includes data plane information, control plane information and management plane information. At layer 1, data plane information is normally assumed to be secured once connections are established, since those connections are dedicated to each VPN. In L1VPNs, VPN connections MUST be restricted to be used only within the same VPN, as described in section 6.2. Note that a customer may wish to assure data plane information security against not only other customers, but also the provider. In such case, the customer may wish to apply their own security mechanisms for data plane information (CE-CE security), as later described.

In addition, information contained in the provider network MUST be secured. This includes VPN service contract information, current VPN connection information, VPN membership information, and system information. Note these types of information MAY be accessible to authorized entities.

9.2 Security Features

Security features include the following:

o Data integrity

The information exchanged between the customer and the provider MUST be delivered unchanged.

o Confidentiality

The information exchanged between the customer and the provider MUST NOT be retrieved by the third party.

o Authentication

The entity requesting the service to the provider MUST be identified.

o Access control

Access to the information contained in the provider network MUST be restricted to the authorized entity.

9.3 Scenarios

There are two scenarios (or occasions) in which security mechanisms are applied. One is the service contract phase, where security mechanisms are applied once. The other is the service access phase, where security mechanisms are applied every time the service is requested.

o Service contract scenario (static)

This scenario includes the addition of new physical devices, such as CE devices, data links and control links. It MUST be guaranteed that these physical devices are connected to the right entity. In addition, authority to access specific information MAY be given to each customer as a part of service contract.

o Service access scenario (dynamic)

This scenario includes the reception of connection requests, routing information exchange requests, and management information retrieval requests. If a communication channel between the customer and the provider (control channel, management interface) is physically separate per customer, and the entity connected over this communication channel is identified in the service contract phase, the provider can ensure who is requesting the service. Also,

the communication channel could be considered as secure. However, when communication channel is physically shared among customers, security mechanisms MUST be available and SHOULD be enforced. Note that even in the case of physically separate communication channels, customers may wish to apply security mechanisms, such as IPsec, to assure higher security, and such mechanisms MUST be available.

When the entity requesting the service is identified, the provider MUST ensure that the request is authorized for that entity. This includes assuring that connection request is between VPN end points belonging to the same VPN.

Also note that customers may wish to apply their own security mechanisms for data plane information (CE-CE security). This includes IPsec for IP traffic.

10. Acknowledgements

The material in this document is based on the work of the ITU-T Study Group 13.

We would like to thank Dimitri Papadimitriou, Deborah Brungard, Yakov Rekhter, Alex Zinin, Igor Bryskin and Adrian Farrel for their useful comments and suggestions.

11. Normative References

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

12. 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.

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

[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., and Rekhter, Y. (editors), "Routing Extensions in Support of Generalized MPLS", draft-ietf-ccamp-gmpls-routing, work in progress.

[L2VPN-FRAME] Andersson, L., and Rosen, E. (editors), "Framework for Layer 2 Virtual Private Networks (L2VPNs)", draft-ietf-l2vpn-l2-framework, work in progress.

[L3VPN-FRAME] Callon, R., and Suzuki, M. (editors), "A Framework for Layer 3 Provider Provisioned Virtual Private Networks", draft-ietf-l3vpn-framework, 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.

[L1VPN-APP] T. Takeda (Ed.), "Applicability analysis of GMPLS protocols to Layer 1 Virtual Private Networks", draft-takeda-l1vpn-applicability, work in progress.

13. 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-Chateaufort
Les Jeunes Bois - MS CTF 32B5 - Chateaufort
78928 YVELINES Cedex 9 - FRANCE
Phone: +33 1 6955 7027
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

14. 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.

15. Full Copyright Statement