

Framework and Requirements for Layer 1 Virtual Private Networks

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

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.

Table of Contents

1. Introduction	3
2. Terminology	3
3. Overview	5
3.1. Network Topology	5
3.2. Introducing Layer 1 VPNs	5
3.3. Current Technologies for Dynamic Layer 1 Provisioning	6
3.4. Relationship with ITU-T	7
4. Motivations	8
4.1. Basic Layer 1 Services	8
4.1.1. L1VPN for Dynamic Layer 1 Provisioning	9
4.2. Merits of L1VPN	9
4.2.1. Customer Merits	9
4.2.2. Provider Merits	10
4.3. L1VPN Deployment Scenarios	10
4.3.1. Multi-Service Backbone	11
4.3.2. Carrier's Carrier	11
4.3.3. Layer 1 Resource Trading	12
4.3.4. Inter-AS and Inter-SP L1VPNs	12

4.3.5. Scheduling Service	13
5. Reference Model	14
5.1. Management Systems	15
6. Generic Service Description	15
6.1. CE Construct	15
6.2. Generic Service Features	16
7. Service Models	16
7.1. Management-Based Service Model	17
7.2. Signaling-Based Service Model (Basic Mode)	17
7.2.1. Overlay Service Model	18
7.3. Signaling and Routing Service Model (Enhanced Mode)	19
7.3.1. Overlay Extension Service Model	20
7.3.2. Virtual Node Service Model	20
7.3.3. Virtual Link Service Model	21
7.3.4. Per-VPN Peer Service Model	22
8. Service Models and Service Requirements	22
8.1. Detailed Service Level Requirements	24
9. Recovery Aspects	25
9.1. Recovery Scope	25
9.2. Recovery Resource Sharing Schemes	26
10. Control Plane Connectivity	27
10.1. Control Plane Connectivity between a CE and a PE	27
10.2. Control Plane Connectivity between CEs	28
11. Manageability Considerations	29
12. Security Considerations	31
12.1. Types of Information	32
12.2. Security Features	32
12.3. Scenarios	33
13. Acknowledgements	34
14. Contributors	34
15. Normative References	35
16. Informative References	35

1. Introduction

This 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 based on the work undertaken within Question 11 of Study Group 13 of 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 Generalized Multi-Protocol Label Switching (GMPLS) protocols.

Use of GMPLS protocols for providing L1VPN services has several advantages, such as:

- Flexible network operation.
- Use of standardized protocols.
- Use of common control and measurement plane protocols applicable to various layer 1 networks, including Time Division Multiplexing (TDM) networks and optical networks.

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], [RFC4202], [RFC3945], [RFC4208], and [RFC4026].

In this context, a Layer 1 Network is any transport network that has connectivity and/or switching using spatial switching (e.g., incoming port or fiber to outgoing port or fiber), lambda-switching, or time-division-multiplex-switching.

A Layer 1 VPN (L1VPN) is a service offered by a core layer 1 network to provide layer 1 connectivity between two or more customer sites, and where the customer has some control over the establishment and type of the connectivity. An alternative definition is simply to say that an L1VPN is a VPN whose data plane operates at layer 1. Further details of the essence of an L1VPN are provided in Section 3.

In addition, the following new terms are used within this document:

- Virtual link: A provider network Traffic Engineering (TE) link advertised to customers in routing information for purposes that include path computation. A direct 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 the links between them.
- VPN end point: A Customer Edge (CE) device's data plane interface, which is connected to a Provider Edge (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 C (Customer) devices using this CE-CE VPN connection as a segment or forwarding adjacency defined in [RFC4206].

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

- CE 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, Time Division Multiplexing (TDM) switch, router, and layer 2 switch. A CE device does not have to have the capability to switch at layer 1, but it is capable of receiving a layer 1 signal and either switching it or terminating it with adaptation. A CE device may be attached to one or more C devices on the customer site, and it may be a host using a layer 1 connection directly.
- PE 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 switch, an Optical Cross-Connect (OXC) (see [RFC3945]), or a Photonic Cross-Connect (PXC) (see [RFC3945]). Alternatively, a PE device may be an Ethernet Private Line (EPL) type of device that maps Ethernet frames onto layer 1 connections (by means of Ethernet over TDM etc.).

- P (Provider) device: A P device is a provider device that is connected only to other provider devices (P or PE devices). A layer 1 P is a TDM switch, OXC, or PXC.
- 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.
- Membership information: A list of CE-PE TE link addresses belonging to the same VPN. Membership information contains the association of a CE, a PE, and a VPN.

3. Overview

3.1. Network Topology

The layer 1 network, made of OXCs, TDM switches, or PXCs may be seen as consisting of PE devices that give access from outside of the network, and P devices that operate only within the core of the network. Similarly, outside the layer 1 network is the customer network consisting of C devices with access to the layer 1 network made through 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.

A layer 1 connection is provided between a pair of CEs. Such a connection follows the hierarchy defined in [RFC4206]. That is, a CE-CE connection may be nested in a lower layer connection (e.g., VC3 connection over STM1 connection). Likewise, the switching capabilities of the interfaces of the CEs, PEs, and Ps on which a connection is routed, follow the hierarchy defined in [RFC4206].

3.2. Introducing Layer 1 VPNs

The concept of a PPVPN has been established through many previous documents such as [RFC4664] and [RFC4110]. Terminology for PPVPNs is set out in [RFC4026] with special reference to layer 2 and layer 3 VPNs.

The realization of 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

extensions 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 the fundamental difference between L1VPNs and L2/L3 VPNs is that in L1VPNs, data plane connectivity does not guarantee control plane connectivity (and vice versa). But CE-PE control plane connectivity is required for L1VPN services provisioned through the control plane, and CE-CE data plane connectivity is maintained by signaling mechanisms based on this control plane connectivity. Furthermore, the provision of CE-CE control plane connectivity over the provider network is also required for certain levels of L1VPN service, and this can be achieved by the exchange of control packets between CEs over the control plane of the provider network. This aspect is discussed further in Section 10.2.

3.3. 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 definition of interfaces for requesting services between the user and the layer 1 network over the User-Network Interface (UNI), and between networks across the External Network-Network Interface (E-NNI) (see [RFC3945], [RFC4208], [RFC4139], and [RFC4258]).

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 may be applied in the context of L1VPNs. For example, the UNI may be applied between the CE and the PE, and the E-NNI may be applied between PEs (inter-AS/SP L1VPNs), or between the CE and the PE.

However, the existing UNI and E-NNI specifications do not provide sufficient parameters 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. A control link is an IP link used for establishing a control channel between nodes.

Another example is that there is no clearly defined way of distributing membership information to be used in combination with UNI/E-NNI. This function is necessary in order to discover the existence and location of the CEs to be connected by L1 connections. Distribution of membership information is typically done by the provider, and may be realized by mechanisms such as static provisioning, or by piggybacking on routing protocols (e.g., see Section 4.2.1 of [RFC4110]). Note that the method chosen for distribution of membership information depends on the solution used for supporting L1VPNs, which is outside of the scope of this document.

Furthermore, customer addressing realms may overlap with each other, and may also overlap with the service provider addressing realm. This requires address mapping mechanisms, but such mechanisms are not well defined in existing UNI/E-NNI specifications.

Lastly, there is no clearly 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 clearly 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 for L1VPNs to be supported in a fully functional manner, these additional capabilities and other requirements set out later in this document must be addressed.

Note that inter-AS/SP L1VPNs require additional analysis beyond the focus of this document.

3.4. Relationship with ITU-T

The foundation of this document is based on the work of the ITU-T Study Group 13, Question 11, such as [Y.1312] and [Y.1313]. This group has been researching and specifying both the requirements and the architecture of L1VPNs for some time. In this context, the foundation of 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 foundation of 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.

4. Motivations

The general benefits and desirability of VPNs have been described many times and in many places ([RFC4110] and [RFC4664]). 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 [RFC3945].

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 criteria that the provider defines, possibly agreed with each customer. To achieve the required level of availability for the customer connections the service provider's network may use restoration or protected resources [RFC4427].
- 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 (i.e., the customer is in control of when the signaled part is established).

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 service types 1 and 2, 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.

This document does not preclude other advanced services and topology support, such as point-to-multipoint (P2MP) services, as part of the layer 1 services, but these are for further study.

4.1.1. L1VPN for Dynamic Layer 1 Provisioning

Private network services in the second category in Section 4.1 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 that the first category in Section 4.1 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, and distribution of membership information applied to a specific group.

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 a layer 1 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 a layer 1 network. So, for example, by sharing the layer 1 network infrastructure with many other users, the customer sites can be connected together across the layer 1 network without bearing the full cost of deploying and managing the layer 1 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, 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 layer 1 network allows the provider to make more comprehensive use of the spare (that is, previously unused) resources within the network. This could be achieved by applying a network model where the provider is responsible for deciding how resources are used and for provisioning of the connection through the layer 1 network.

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. 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 using 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 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). The layer 1 transport network and each service network belong to the same organization, but may be managed separately. 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 transport 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 for each service department. 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 transport network by using GMPLS in the integrated model [RFC3945] 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.

These points may apply to the management plane functionalities as well as to the control plane functionalities.

4.3.2. Carrier's Carrier

A carrier's carrier is characterized such that one carrier that receives another carrier's L1VPN service provides its own services. In this scenario, two carriers are in different organizations. 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, while separating the control functions for each customer carrier.

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 in Section 4.3.2, 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 of service 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-AS and Inter-SP L1VPNs

In addition to the scenarios where a single connection between two CEs is routed over a single service provider as mentioned in Section 4.3.2, it is possible that a connection is routed over multiple ASes within a service provider (called inter-AS L1VPN) or over multiple service providers (called inter-SP L1VPN).

The inter-AS L1VPN scenario can be used to construct a single L1VPN from network resources administered by different domains of a single

service provider. These administrative domains might not usually have a collaborative relationship at layer 1, and so the inter-AS L1VPN offers a new business model for joint delivery of services to a customer. Consideration of inter-AS L1VPNs requires further analysis beyond the scope of this document.

The inter-SP 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, and this scenario contains many more manageability, security, privacy, policy, and commercial issues than the more simple inter-AS L1VPN case. Consideration of inter-SP L1VPN requires further analysis beyond the scope of this document.

4.3.5. Scheduling Service

In some deployment scenarios, customers of L1VPN services may wish to set up layer 1 connections not on-demand, but at a planned time in the future. Or, even though customers of L1VPN services may wish to use layer 1 connections on-demand, they can tolerate some delay, for example, due to lack of resources at that moment.

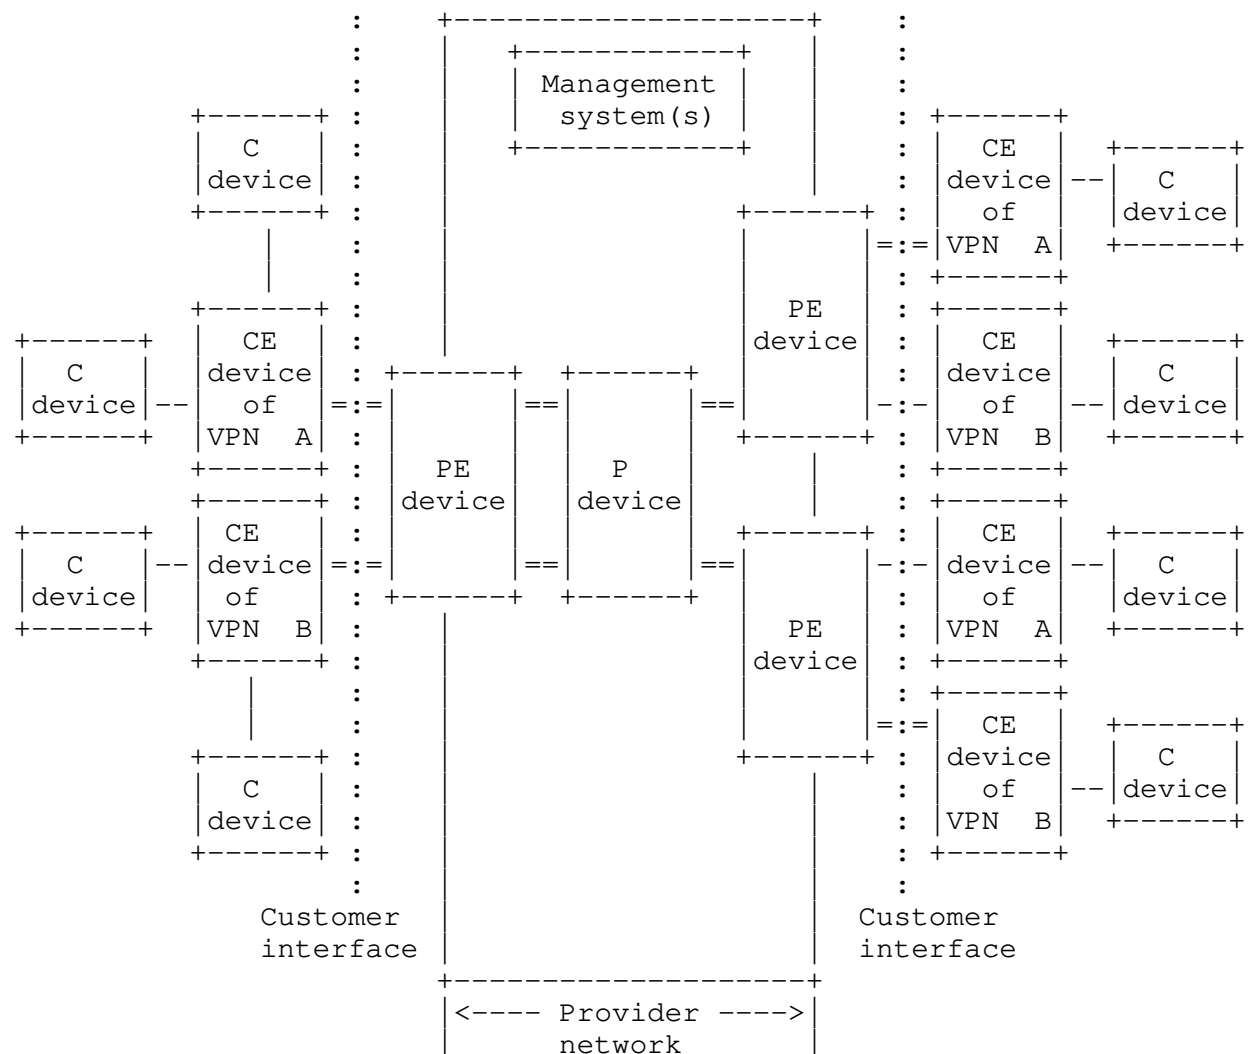
In those scenarios, the provider can reserve bandwidth at a specified time in the future, and can establish the VPN connections according to a schedule. This makes it possible to use bandwidth more efficiently over time (i.e., support more demand). This service, the scheduling service, may be used to support customers who use layer 1 connections for data backup applications, content delivery applications, and some other applications.

Furthermore, customers may be able to specify when to release layer 1 connections in advance. By considering this information, the provider may be able to further engineer scheduling, which leads to still more efficient bandwidth usage.

Note that scheduling of L1VPN services requires time-scoped resource management, which is not well considered in current GMPLS protocols and requires the support of the management plane. In addition, offering scheduling service and on-demand service on the same infrastructure needs careful consideration.

5. Reference Model

Figure 5.1 describes the L1VPN reference model.



Key: ==== Layer 1 Connection -- link

Figure 5.1: L1VPN Reference Model

In an 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.

Note that as mentioned in Section 3.1, these VPN connections follow the hierarchy defined in [RFC4206].

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. Sections 7 and 11 provide more detail.

6. Generic Service Description

This section describes generic L1VPN services. Detailed 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 such as IP tunnel, or using separate addressing).

A customer addressing realm consists of CE-PE TE link addresses and CE-PE control channel addresses as well as customer site addresses (C and CE addresses). Customer addressing realms may overlap, and may also overlap with the service provider addressing realm.

NATs or firewalls might reasonably be placed at customer interfaces, or between administrative domains within the core network. Addressing in the L1VPN model must handle such eventualities. Traversal of NATs and firewalls within the customer network might have implications for L1VPN services that connect C devices, and is for further study.

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 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 [RFC4202]. It must be understood that meeting the requirements set out in this document may necessitate extensions 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). 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. The customer interface may be instantiated by the CE-PE control plane communication and/or the management plane communication between the customer management systems(s) and the provider management system(s). Note that how to realize a CE-PE control channel is discussed in Section 10.1. Customer management system(s) and provider management systems(s) may communicate by utilizing the CE-PE control channel(s).

7.1. Management-Based Service Model

Figure 7.1 describes the Management-based service model.

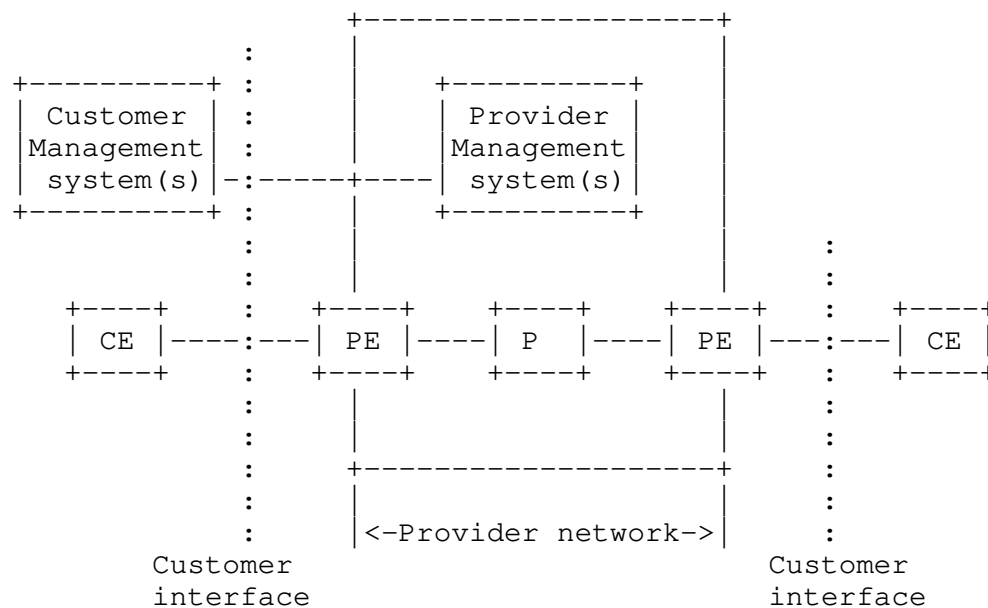


Figure 7.1: Management-Based Service Model

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, mechanisms to support soft permanent connections can be applied. However, interfaces between management systems are not within the scope of this document.

7.2. Signaling-Based Service Model (Basic Mode)

In this service model, the CE-PE interface's functional repertoire is limited to path setup signaling only. The provider's network is not involved in distribution of customer network's routing information.

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

7.2.1. Overlay Service Model

Figure 7.2 describes the Overlay service model.

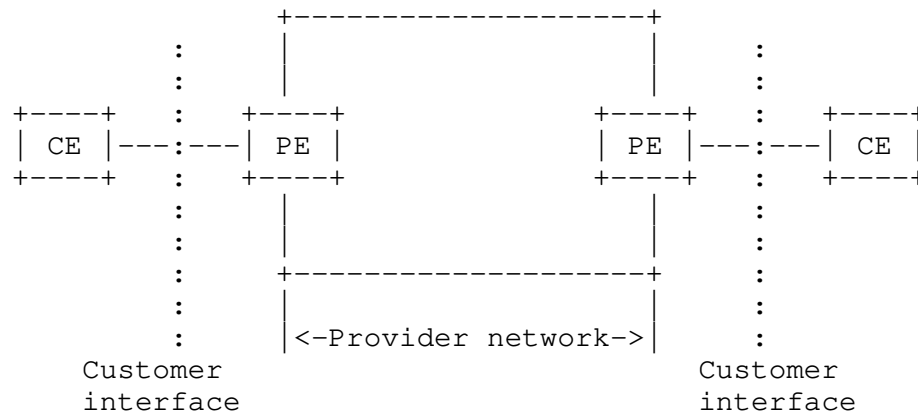


Figure 7.2: Overlay Service Model

In this service model, the customer interface is based on the GMPLS UNI Overlay [RFC4208]. The CE requests layer 1 connection setup/deletion to a remote CE. There is no routing protocol running (i.e., no routing neighbor/peering relationship) between a CE and a 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.

In this model, membership information needs to be configured on PEs, so that the PE that receives a Path message from the ingress CE can identify the remote PE connected to the egress CE. Distribution of membership information between PEs is typically done by the provider, and may be realized by mechanisms such as static provisioning, or by piggybacking on routing protocols (auto-discovery).

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.

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 customer network routing information (i.e., reachability or TE information in remote customer sites), or parameters of the part of the provider's network dedicated to the customer.

By allowing CEs to obtain customer network routing information, a so-called N-square routing problem could be solved.

In addition, by using the received traffic engineering-based routing information, a customer can use traffic engineering capabilities. 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.

As with the Signaling-based service model, 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 (i.e., the level of information details that a customer is allowed to receive in routing and signaling).

7.3.1. Overlay Extension Service Model

This service model complements the Overlay service model. In this service model, a CE receives a list of CE-PE TE link addresses to which it can request a VPN connection (i.e., membership information). This may include additional information concerning these TE links (e.g., switching type). 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.

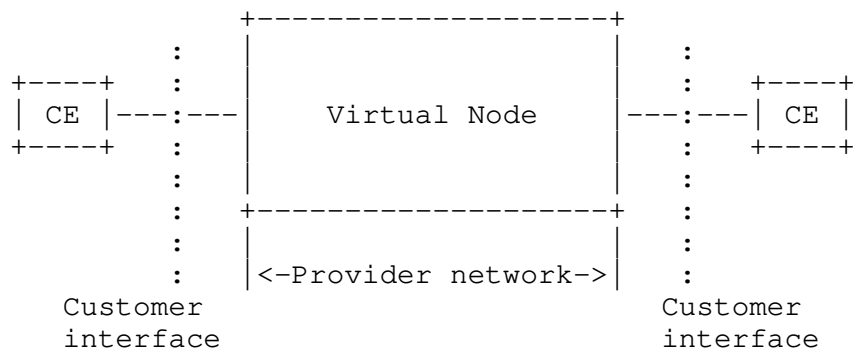


Figure 7.3: Virtual Node Service Model

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. The CE receives routing information about CE-PE links and the customer network (i.e., 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. A virtual link is assigned to each VPN and disclosed to the corresponding CEs. As such, the CE receives routing information about CE-PE links, customer network (i.e., remote customer sites), as well as virtual links assigned to each VPN. 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.

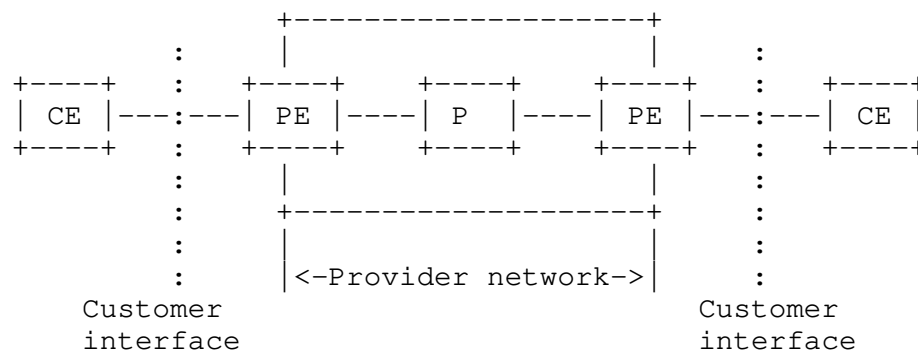


Figure 7.5: Per-VPN Peer Service Model

This service model is a generalization and combination of the Virtual Link service model and the Virtual Node service model mentioned in Sections 7.3.2 and 7.3.3 respectively.

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, customer network (i.e., 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 between them.

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 what 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 (optionally includes TE information of the associated CE-PE TE links)

- Customer network routing information (reachability only, or may include TE information)

- Provider network routing information (TE information)

Note that link management information (e.g., LMP [RFC4204]) may be exchanged between a CE and a PE, but this is orthogonal to the definition of the service models.

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 customers receive provider network routing information in the form of virtual link, customers must be able to specify such links for a VPN connection over the provider network in signaling.

8.1. Detailed Service Level Requirements

In addition to the requirements set out in table 1, more detailed service requirements are provided below. They are generally common to the various service models, except where indicated.

- Selection of layer 1 service class: Customers MAY be allowed to specify a layer 1 service class (e.g., availability level) for a VPN connection. Further details are described in Section 9.

- 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., failure notification by RSVP-TE, data plane alarm notification through the management plane, notification of connection setup rejection causes). Note that this does not prevent customers from using Operations and Management (OAM) mechanisms for, or on, their VPN connections. 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 (through the management plane).
- 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 12.
- Filtering: Unnecessary information (e.g., information concerning other VPNs) MUST NOT be provided to each customer. This applies particularly to the Signaling and Routing service model, but is also relevant to the Signaling-based service model and to the Management-based service model. Further details are described in Section 12.

9. Recovery Aspects

9.1. Recovery Scope

GMPLS provides various recovery techniques for use in different recovery scenarios [RFC4427]. The provider network may apply these recovery techniques to protect VPN connections as part of the L1VPN service, for example as follows:

- o PE-PE recovery

The provider network constitutes a recovery domain, and the recovery scope is the PE-PE part of the CE-CE VPN connection.

It should be possible for the provider network to hide the provider network recovery operation from the customer. Namely, it should be possible to configure the provider network to not notify the customer when a failure occurs and a PE-PE recovery operation successfully repairs the failure. Further, when PE-PE recovery fails and the failure should be notified to the customer, it should be possible for the provider network to hide its internal topology.

- o CE-PE recovery

The recovery scope is either or both of the ingress and egress CE-PE links of the CE-CE VPN connection.

- o CE-CE recovery

The recovery scope is the entire CE-CE VPN connection.

When a failure needs to be notified to a customer so that the customer can initiate recovery operation, it should be possible for the provider network to hide its internal topology.

These recovery schemes may be applied in combination.

Customers may be allowed to specify the desired recovery level in a connection setup request. Furthermore, the customer may be allowed to specify the desired recovery level in a way that is agnostic of the recovery technique (e.g., when the recovery operation does not require cooperation between the provider network and the customer network). In such cases, the provider network must translate the specified recovery level into specific recovery techniques, based on operational policies. This allows enhanced recovery techniques above and beyond the GMPLS specifications to be used in the provider network.

9.2. Recovery Resource Sharing Schemes

The provider network may support various recovery resource sharing schemes, such as the following:

- o Shared recovery

When the provider network supports shared recovery (e.g., shared mesh restoration [RFC4427]), the provider network may provide

sharing recovery resources between VPN connections that serve with only the same VPN, a specific set of VPNs, or any VPN. The default mode is sharing recovery resources with any VPN.

o Extra traffic

GMPLS recovery mechanisms support extra traffic. Extra traffic allows the transfer of preemptable traffic on the recovery resources when these resources are not being used for the recovery of protected normal traffic [RFC4427].

In the context of L1VPNs, extra traffic is applied for CE-CE VPN connections, or PE-PE part of CE-CE VPN connections. The latter case may be applied only when there is hierarchy (i.e., CE-CE VPN connection is nested on top of PE-PE connection). In this section, the latter aspect is analyzed.

When the provider network allows a CE-CE VPN connection to be set up as "extra traffic", it means that the VPN connection may use a PE-PE connection that protects some other CE-CE VPN connection. In such a case the provider network may restrict extra traffic CE-CE VPN connection to use resources (i.e., the PE-PE connections) that:

- protect VPN connections from the same VPN as the extra traffic connection.
- are used for a specific set of VPNs.
- are available for any VPN.

The default mode is to support preemptable traffic on recovery resources reserved for any VPN.

10. Control Plane Connectivity

10.1. Control Plane Connectivity between a CE and a PE

In the Signaling-based service model and the Signaling and Routing service model, there must be a control channel (IP-level connectivity) between a CE and its PE. The instantiation of the control channel may differ depending on addressing and security.

As stated in Section 6.1, it is necessary to disambiguate control plane messages exchanged between the CE and PE if the CE-PE relationship is applicable to more than one VPN. Furthermore, private addresses may be assigned to CE-PE control channels.

Security aspects of the CE-PE control channel are discussed in Section 12.

10.2. Control Plane Connectivity between CEs

A customer network connected by VPN connections may be controlled by MPLS or GMPLS, and the VPN connections may be treated as TE links within the customer network. In such cases, there must be control plane (IP-level) connectivity between the CEs, so that control messages, such as signaling and routing messages, can be exchanged between the CEs. Furthermore, in some recovery techniques, Notify message exchange is needed between the ingress and egress of the VPN connection, which requires control plane connectivity between the CEs. There are several potential ways to achieve this.

o Use of VPN connections as in-band control channels

If the CEs have the ability to inject control messages into the VPN connections and to extract the messages at the far end of the VPN connections, then control messages can be exchanged in-band. For example, when a VPN connection is a Packet Switch Capable (PSC) TE link in the customer network, this operation is transparent to the L1VPN service provider.

o Use of overhead associated with the VPN connections

If the VPN connection provides connectivity in the customer network at a different switching capability (implying network technology layer) from that used by the provider network to support the CE-PE and PE-PE connectivity, then the customer network can utilize any overhead available within the VPN connection as a control channel to connect the CEs. For example, if a VPN connection provides a TDM TE link in the customer network but is supported by a technology such as lambda or fiber, then the CEs may utilize the overhead (DCC) as a control channel, if the network supports transparent transfer of such overhead. This operation is transparent to the L1VPN service provider.

o Use of control-channel-specific VPN connections

A customer establishes VPN connections dedicated as control channels. This operation is transparent to the L1VPN service provider, but since control plane traffic is likely to be relatively low compared with the capacity of VPN connections, this may be an expensive solution for the customer.

- o Use of separate network

A customer may utilize another network and network service, such as private line service, L3VPN service, L2VPN service, or Internet access service, to establish CE-CE control channel connectivity. This operation is transparent to the L1VPN service provider.

- o Use of CE-PE control channels

In the Signaling-based service model, and the Signaling and Routing service model, there must be control plane (IP-level) connectivity between the CE and PE, as described in Section 10.1.

By utilizing this, CE-CE control message exchange could be realized as part of the service provided by the L1VPN service provider. Namely, the provider network transfers control messages received over the CE-PE control channel to the other side of the provider network and delivers them through the PE-CE control channel. The realization of this within the provider network is up to the operator, but where the provider network uses a GMPLS control plane, the customer control plane messages could be forwarded through the provider control plane, perhaps using IP tunnels.

Care must be taken to protect the provider network and other customers from Denial of Service (DoS) attack. Traffic saturation over the control plane network needs to be carefully managed as well. Note that if private addresses are assigned to the CE-PE control channels, the provider network must support VPN-scoped routing and forwarding for control messages.

11. Manageability Considerations

Manageability considerations for GMPLS are described in existing documents, such as [RFC3945]. Also, manageability considerations for L3VPN are described in existing documents, such as [RFC4176]. These manageability considerations should also be applied in L1VPNs, and these aspects are described in this section. In addition, there are some specific manageability considerations for L1VPNs, such as configuration and accounting.

- o Fault management

The provider network MUST support fault management. It MUST support liveness detection, and monitoring and verification of correct operation.

When a failure occurs, the provider network SHOULD correlate the failure. Also, it SHOULD be able to detect which customer is affected by the failure.

If the provider network can resolve failures without intervention from the customer network, it MUST be possible to configure the provider network to not report failures to the customers. However, it MAY be part of an agreement between a customer and provider that failures are reported to the customer, regardless.

o Configuration management

The provider network MUST support configuration management, such as the following.

- Service mode/model configuration.
- Network representation configuration: Configuration of virtual node and virtual link.
- Resource allocation configuration: Dedicated, shared. See Section 8 for more detail.
- Recovery policy configuration: For example, recovery resource sharing schemes, such as shared recovery, extra traffic. See Section 9 for more detail.
- Membership configuration.
- Network/Element level configuration: For example, TE link configuration.

It SHOULD be possible for the provider network to verify that configuration is correctly made.

o Accounting management

The provider network MUST support accounting management. It MUST be able to record usage of VPN connections for each customer.

o Performance management

The provider network MUST support performance management.

In particular, it MUST support performance monitoring of parameters associated with the Service Level Agreement (SLA), such as bit error rate per VPN connection, and SLA verification.

In addition, it MUST support performance monitoring and analysis of parameters related to the network and equipment not directly associated with the SLA, such as network resource utilization.

- o Security management

The provider network MUST support security management. See Section 12 for details.

- o Management systems

In order to support various management functionalities, the provider network relies on management systems and related tools. GMPLS protocols and potential extensions of GMPLS MUST be able to work with management systems and related tools to provide such functionalities.

In particular, MIB modules for GMPLS protocols and potential extensions MUST be supported.

- o Management of customer networks

Customers MAY outsource management of their network (especially CEs and CE-CE links) to the provider network. In such case, the provider MUST be able to manage the customer network, as well as the provider network.

12. 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 [RFC4110], [RFC4111], and [RFC4664]. 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.

12.1. 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. That is, it is not possible to communicate unless there is a connection. Therefore, in L1VPNs, the main concern of data plane security is restricting VPN connections 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.

12.2. 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 disclosed to a third party.

- o Authentication

The entity requesting the service to the provider MUST be identified and have its identity authenticated, and the provider providing the service MUST also be identified and have its identity authenticated.

- o Access control

Access to the information contained in the provider network, which may be information about the customer networks or the existence of customers, as well as about the provider network, MUST be restricted to the authorized entity.

- o DoS attack detection and protection

The provider network MUST have mechanisms to detect DoS attack and to protect against it reactively and proactively.

12.3. 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 (e.g., attempts to establish a neighbor relationship in routing protocols, or command request via the management plane interface), 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. Examples of such security mechanisms include IPsec [RFC4302] and [RFC4303]. Note that even in the case of physically separate communication channels, customers may wish to apply security mechanisms 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 [RFC4302] and [RFC4303] for IP traffic.

13. 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, Adrian Farrel, and Ross Callon for their useful comments and suggestions.

Thanks to Mark Townsley, Dan Romascanu, and Cullen Jennings for helpful input during IESG review.

14. Contributors

The foundation of 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 the foundation of this document is a summary and development of the conclusions they have reached. Based on such material, the IETF and the L1VPN WG in particular have developed this framework and requirements for the support of L1VPNs by use of GMPLS protocols.

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)

15. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.
- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.
- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.
- [Y.1312] Y.1312 - Layer 1 Virtual Private Network Generic requirements and architecture elements, ITU-T Recommendation, September 2003, available from <<http://www.itu.int>>.

16. Informative References

- [Y.1313] Y.1313 - Layer 1 Virtual Private Network service and network architectures, ITU-T Recommendation, July 2004, available from <<http://www.itu.int>>.

- [RFC4110] Callon, R. and M. Suzuki, "A Framework for Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4110, July 2005.
- [RFC4111] Fang, L., Ed., "Security Framework for Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4111, July 2005.
- [RFC4139] Papadimitriou, D., Drake, J., Ash, J., Farrel, A., and L. Ong, "Requirements for Generalized MPLS (GMPLS) Signaling Usage and Extensions for Automatically Switched Optical Network (ASON)", RFC 4139, July 2005.
- [RFC4176] El Mghazli, Y., Ed., Nadeau, T., Boucadair, M., Chan, K., and A. Gonguet, "Framework for Layer 3 Virtual Private Networks (L3VPN) Operations and Management", RFC 4176, October 2005.
- [RFC4204] Lang, J., Ed., "Link Management Protocol (LMP)", RFC 4204, October 2005.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.
- [RFC4258] Brungard, D., Ed., "Requirements for Generalized Multi-Protocol Label Switching (GMPLS) Routing for the Automatically Switched Optical Network (ASON)", RFC 4258, November 2005.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [RFC4427] Mannie, E., Ed., and D. Papadimitriou, Ed., "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4427, March 2006.
- [RFC4664] Andersson, L., Ed., and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, September 2006.

Authors' Addresses

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

Tomonori Takeda, Editor
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

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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

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

Intellectual Property

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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

