

Service Requirements for Layer 3 Provider Provisioned Virtual Private Networks (PPVPNs)

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 Internet Society (2005).

Abstract

This document provides requirements for Layer 3 Virtual Private Networks (L3VPNs). It identifies requirements applicable to a number of individual approaches that a Service Provider may use to provision a Virtual Private Network (VPN) service. This document expresses a service provider perspective, based upon past experience with IP-based service offerings and the ever-evolving needs of the customers of such services. Toward this end, it first defines terminology and states general requirements. Detailed requirements are expressed from a customer perspective as well as that of a service provider.

Table of Contents

1.	Introduction.	3
1.1.	Scope of This Document.	4
1.2.	Outline	5
2.	Contributing Authors.	5
3.	Definitions	5
3.1.	Virtual Private Network	6
3.2.	Users, Sites, Customers, and Agents	6
3.3.	Intranets, Extranets, and VPNs.	6
3.4.	Networks of Customer and Provider Devices	7
3.5.	Access Networks, Tunnels, and Hierarchical Tunnels.	7
3.6.	Use of Tunnels and Roles of CE and PE in L3VPNs	8
3.6.1.	PE-Based L3VPNs and Virtual Forwarding Instances	8
3.6.2.	CE-Based L3VPN Tunnel Endpoints and Functions	10

3.7.	Customer and Provider Network Management.	10
4.	Service Requirements Common to Customers and Service Providers	11
4.1.	Isolated Exchange of Data and Routing Information	11
4.2.	Addressing.	12
4.3.	Quality of Service.	12
4.3.1.	QoS Standards	12
4.3.2.	Service Models.	13
4.4.	Service Level Specification and Agreements.	14
4.5.	Management.	14
4.6.	Interworking.	15
5.	Customer Requirements	15
5.1.	VPN Membership (Intranet/Extranet).	15
5.2.	Service Provider Independence	16
5.3.	Addressing.	16
5.4.	Routing Protocol Support.	16
5.5.	Quality of Service and Traffic Parameters	16
5.5.1.	Application Level QoS Objectives.	17
5.5.2.	DSCP Transparency	17
5.6.	Service Level Specification/Agreement	18
5.7.	Customer Management of a VPN.	18
5.8.	Isolation	18
5.9.	Security.	19
5.10.	Migration Impact.	19
5.11.	Network Access.	19
5.11.1.	Physical/Link Layer Technology.	20
5.11.2.	Temporary Access.	20
5.11.3.	Sharing of the Access Network	20
5.11.4.	Access Connectivity	20
5.12.	Service Access.	23
5.12.1.	Internet Access	23
5.12.2.	Hosting, Application Service Provider	24
5.12.3.	Other Services.	24
5.13.	Hybrid VPN Service Scenarios.	24
6.	Service Provider Network Requirements	24
6.1.	Scalability	24
6.2.	Addressing.	25
6.3.	Identifiers	25
6.4.	Discovering VPN Related Information	26
6.5.	SLA and SLS Support	26
6.6.	Quality of Service (QoS) and Traffic Engineering.	27
6.7.	Routing	27
6.8.	Isolation of Traffic and Routing.	28
6.9.	Security.	28
6.9.1.	Support for Securing Customer Flows	28
6.9.2.	Authentication Services	29
6.9.3.	Resource Protection	30
6.10.	Inter-AS (SP)VPNs	30

6.10.1.	Routing Protocols	31
6.10.2.	Management.	31
6.10.3.	Bandwidth and QoS Brokering	31
6.10.4.	Security Considerations	32
6.11.	L3VPN Wholesale	32
6.12.	Tunneling Requirements.	33
6.13.	Support for Access and Backbone Technologies.	33
6.13.1.	Dedicated Access Networks	34
6.13.2.	On-Demand Access Networks	34
6.13.3.	Backbone Networks	35
6.14.	Protection, Restoration	35
6.15.	Interoperability.	35
6.16.	Migration Support	36
7.	Service Provider Management Requirements.	36
7.1.	Fault Management.	37
7.2.	Configuration Management.	37
7.2.1.	Configuration Management for PE-Based VPNs.	38
7.2.2.	Configuration Management for CE-Based VPNs.	39
7.2.3.	Provisioning Routing.	39
7.2.4.	Provisioning Network Access	39
7.2.5.	Provisioning Security Services.	40
7.2.6.	Provisioning VPN Resource Parameters.	40
7.2.7.	Provisioning Value-Added Service Access	40
7.2.8.	Provisioning Hybrid VPN Services.	41
7.3.	Accounting.	41
7.4.	Performance Management.	42
7.4.1.	Performance Monitoring.	42
7.4.2.	SLA and QoS Management Features	42
7.5.	Security Management	43
7.5.1.	Resource Access Control	43
7.5.2.	Authentication.	43
7.6.	Network Management Techniques	44
8.	Security Considerations	44
9.	Acknowledgments	45
10.	References.	45
10.1.	Normative References.	45
10.2.	Informative References.	46
	Authors' Addresses	49
	Full Copyright Statement	50

1. Introduction

This section describes the scope and outline of the document.

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 RFC 2119 ([RFC2119]).

1.1. Scope of This Document

This document provides requirements specific to Layer 3 Virtual Private Networks (L3VPN). (Requirements that are generic to L2 and L3 VPNs are contained in [RFC3809].)

This document identifies requirements that may apply to one or more individual approaches that a Service Provider may use to provision a Layer 3 (e.g., IP) VPN service. It makes use of the terminology and common components for Layer 3 VPNs as defined in [L3VPN-FR] and of the generic VPN terminology defined in [PPVPN-TERM].

The specification of technical means to provide L3VPN services is outside the scope of this document. Other documents are intended to cover this aspect, such as the L3 VPN framework document [L3VPN-FR] and several sets of documents, one for each technical approach for providing L3VPN services.

Technical approaches targeted by this document include the network-based (PE-based) L3VPN category (aggregated routing VPNs [2547bis] and virtual routers [PPVPN-VR]) and the CE-based L3VPNs category [CE-PPVPN][IPSEC-PPVPN]. The document distinguishes L3VPN categories as to where the endpoints of tunnels exist, as detailed in the L3VPN framework document [L3VPN-FR]. Terminology describing whether equipment faces a customer or the service provider network is used to define the various types of L3VPN solutions.

This document is intended as a "checklist" of requirements, providing a consistent way to evaluate and document how well each approach satisfies specific requirements. The applicability statement documents for each approach should present the results of this evaluation. This document is not intended to compare one approach to another.

This document provides requirements from several points of view. It begins with some considerations from a point of view common to customers and service providers not covered in the generic provider provisioned VPN requirement document [RFC3809], continues with a customer perspective, and concludes with specific needs of a Service Provider (SP).

The following L3VPN deployment scenarios are considered within this document:

1. Internet-wide: VPN sites attached to arbitrary points in the Internet.

2. Single SP/single AS: VPN sites attached to the network of a single provider within the scope of a single AS.
3. Single SP/multiple ASes: VPN sites attached to the network of a single provider consisting of multiple ASes.
4. Cooperating SPs: VPN sites attached to networks of different providers that cooperate with each other to provide the VPN service.

The above deployment scenarios have many requirements in common. These include SP requirements for security, privacy, manageability, interoperability, and scalability, including service provider projections for number, complexity, and rate of change of customer VPNs over the next several years. When requirements apply to a specific deployment scenario, the above terminology is used to state the context of those particular requirements.

1.2. Outline

The outline of the rest of the document is as follows: Section 2 lists the contributing authors. Section 3 provides definitions of terms and concepts. Section 4 provides requirements common to both customers and service providers that are not covered in the generic provider provisioned VPN requirement document [RFC3809]. Section 5 states requirements from a customer perspective. Section 6 states network requirements from a service provider perspective. Section 7 states service provider management requirements. Section 8 describes security considerations. Section 9 lists acknowledgments. Section 10 provides a list of references cited herein. Section 11 lists the authors' addresses.

2. Contributing Authors

This document is the combined effort of the two co-editors and the following contributing authors:

Luyuan Fang
Ananth Nagarajan
Junichi Sumimoto
Rick Wilder

3. Definitions

This section provides the definition of terms and concepts used throughout the document. Terminology used herein is taken from [PPVPN-TERM] and [L3VPN-FR].

3.1. Virtual Private Network

"L3 Virtual Private Network" (L3VPN) refers to the L3 communication between a set of sites making use of a shared network infrastructure.

"Provider Provisioned VPN" (PPVPN) refers to VPNs for which the service provider participates in management and provisioning of the VPN.

3.2. Users, Sites, Customers, and Agents

User: A user is an entity (e.g., a human being using a host, a server, or a system) authorized to use a VPN service.

Site: A site is a set of users that have mutual L3 (i.e., IP) reachability without use of a specific service provider network. A site may consist of a set of users that are in geographic proximity. Note that a topological definition of a site (e.g., all users at a specific geographic location) may not always conform to this definition. For example, two geographic locations connected via another provider's network would also constitute a single site as communication between the two locations does not involve the use of the service provider offering the L3 VPN service.

Customer: A single organization, corporation, or enterprise that administratively controls a set of sites.

Agent: A set of users designated by a customer who has the authorization to manage a customer's VPN service offering.

3.3. Intranets, Extranets, and VPNs

Intranet: An intranet restricts communication to a set of sites that belong to one customer. An example is branch offices at different sites that require communication with a headquarters site.

Extranet: An extranet allows the specification of communication between a set of sites that belong to different customers. In other words, two or more organizations have access to a specified set of each other's sites. Examples of extranets include multiple companies cooperating in joint software development, a service provider having access to information from the vendors' corporate sites, different companies, or universities participating in a consortium. An extranet often has further restrictions on reachability, for example, at a host and individual transport level.

Note that an intranet or extranet can exist across a single service provider network with one or more ASes, or across multiple service provider networks.

L3 Virtual Private Network (L3VPN): An alternative definition of VPN refers to a specific set of sites that have been configured to allow communication as either an intranet or an extranet. Note that a site is a member of at least one VPN and may be a member of many VPNs.

3.4. Networks of Customer and Provider Devices

L3VPNs are composed of the following types of devices.

Customer Edge (CE) device: A CE device faces the users at a customer site. The CE has an access connection to a PE device. It may be a router or a switch that allows users at a customer site to communicate over the access network with other sites in the VPN. In a CE-based L3VPN, as intended in this document (provider-provisioned CE-based VPN), the service provider manages (at least partially) the CE device.

Provider Edge (PE) device: A PE device faces the provider network on one side and attaches via an access connection over one or more access networks to one or more CE devices. It participates in the Packet Switched Network (PSN) in performing routing and forwarding functions.

Note that the definitions of Customer Edge and Provider Edge do not necessarily describe the physical deployment of equipment on customer premises or a provider point of presence.

Provider (P) device: A device within a provider network that interconnects PE (or other P) devices but does not have any direct attachment to CE devices. The P router does not keep VPN state and is VPN unaware [PPVPN-TERM].

Packet Switched Network (PSN): A (IP or MPLS [RFC3031]) network through which the tunnels supporting the VPN services are set up [PPVPN-TERM].

Service Provider (SP) network: An SP network is a set of interconnected PE and P devices administered by a single service provider in one or more ASes.

3.5. Access Networks, Tunnels, and Hierarchical Tunnels

VPNs are built between CEs by using access networks, tunnels, and hierarchical tunnels across a PSN.

Access connection: An access connection provides connectivity between a CE and a PE. This includes dedicated physical circuits, virtual circuits (such as Frame Relay), ATM, Ethernet (V)LAN, or IP tunnels (e.g., IPsec, L2TP [RFC2661]).

Access network: An access network provides access connections between CE and PE devices. It may be a TDM network, an L2 network (e.g., FR, ATM, and Ethernet), or an IP network over which access is tunneled (e.g., by using L2TP).

Tunnel: A tunnel between two entities is formed by encapsulating packets within another encapsulating header for the purposes of transmission between those two entities in support of a VPN application. Examples of protocols commonly used for tunneling are GRE, IPsec, IP-in-IP tunnels, and MPLS.

Hierarchical Tunnel: Encapsulating one tunnel within another forms a hierarchical tunnel. The innermost tunnel protocol header defines a logical association between two entities (e.g., between CEs or PEs) [VPNTUNNEL]. Note that the tunneling protocols need not be the same at different levels in a hierarchical tunnel.

3.6. Use of Tunnels and Roles of CE and PE in L3 VPNs

This section summarizes the points where tunnels terminate and the functions implemented in the CE and PE devices that differentiate the two major categories of L3VPNs for which requirements are stated, namely PE-based and CE-based L3VPNs. See the L3VPN framework document for more detail [L3VPN-FR].

3.6.1. PE-Based L3VPNs and Virtual Forwarding Instances

In a PE-based L3VPN service, a customer site receives IP layer (i.e., layer 3) service from the SP. The PE is attached via an access connection to one or more CEs. The PE forwards user data packets based on information in the IP layer header, such as an IPv4 or IPv6 destination address. The CE sees the PE as a layer 3 device such as an IPv4 or IPv6 router.

Virtual Forwarding Instance (VFI): In a PE-based L3VPN service, the PE contains a VFI for each L3 VPN that it serves. The VFI terminates tunnels for interconnection with other VFIs and also terminates access connections for accommodating CEs. VFI contains information regarding how to forward data received over the CE-PE access connection to VFIs in other PEs supporting the same L3VPN. The VFI includes the router information base and the forwarding information base for an L3VPN [L3VPN-FR]. A VFI enables router functions dedicated to serving a particular VPN, such as separation of

forwarding and routing and support for overlapping address spaces. Routing protocols in the PEs and the CEs interact to populate the VFI.

The following narrative and figures provide further explanation of the way PE devices use tunnels and hierarchical tunnels. Figure 1.1 illustrates the case where a PE uses a separate tunnel for each VPN. As shown in the figure, the tunnels provide communication between the VFIs in each of the PE devices.




Figure 1.1. PE Usage of Separate Tunnels to Support VPNs

Figure 1.2 illustrates the case where a single hierarchical tunnel is used between PE devices to support communication for VPNs. The innermost encapsulating protocol header provides the means for the PE to determine the VPN for which the packet is directed.




Figure 1.2. PE Usage of Shared Hierarchical Tunnels to Support VPNs

3.6.2. CE-Based L3VPN Tunnel Endpoints and Functions

Figure 1.3 illustrates the CE-based L3VPN reference model. In this configuration, typically a single level of tunnel (e.g., IPsec) terminates at pairs of CEs. Usually, a CE serves a single customer site, and therefore the forwarding and routing is physically separate from all other customers. Furthermore, the PE is not aware of the membership of specific CE devices to a particular VPN. Hence, the VPN functions are implemented with provisioned configurations on the CE devices, and the shared PE and P network is used to only provide the routing and forwarding that supports the tunnel endpoints on between CE devices. The tunnel topology connecting the CE devices may be a full or partial mesh, depending on VPN customer requirements and traffic patterns.




Figure 1.3. CE-Based L3VPN

3.7. Customer and Provider Network Management

Customer Network Management Function: A customer network management function provides the means for a customer agent to query or configure customer-specific information, or to receive alarms regarding his or her VPN. Customer-specific information includes data related to contact, billing, site, access network, IP address,

and routing protocol parameters. It may use a combination of proprietary network management system, SNMP manager, or directory service (e.g., LDAP [RFC3377] [RFC2251]).

Provider Network Management Function: A provider network management function provides many of the same capabilities as a customer network management system across all customers. This would not include customer confidential information, such as keying material. The intent of giving the provider a view comparable to that of the customer is to aid in troubleshooting and problem resolution. Such a system also provides the means to query, configure, or receive alarms regarding any infrastructure supporting the L3VPN service. It may use a combination of proprietary network management system, SNMP manager, or directory service (e.g., LDAP [RFC3377] [RFC2251]).

4. Service Requirements Common to Customers and Service Providers

Many of the requirements that apply to both the customer and the provider and are of an otherwise general nature, or that apply to both L2 and L3VPNs, are described in [RFC3809]. This section contains requirements that are not covered in [RFC3809] and that are specific to L3VPNs.

4.1. Isolated Exchange of Data and Routing Information

A mechanism must be provided for isolating the distribution of reachability information to only those sites associated with a VPN.

L3VPN solutions shall define means that prevent routers in a VPN from interacting with unauthorized entities and that avoid introducing undesired routing information that could corrupt the VPN routing information base [VPN-CRIT].

A means must be provided to constrain or isolate the distribution of addressed data to only those VPN sites determined by either routing data and/or configuration.

A single site shall be capable of being in multiple VPNs. The VPN solution must ensure that traffic is exchanged only with sites in the same VPN.

The internal structure of a VPN should not be advertised or discoverable from outside that VPN.

Note that isolation of forwarded data or exchange of reachability information to only those sites that are part of a VPN may be viewed as a form of security - for example, [Y.1311.1], [MPLSSEC].

4.2. Addressing

IP addresses must be unique within the set of sites reachable from the VPNs of which a particular site is a member.

A VPN solution must support IPv4 and IPv6 as both the encapsulating and encapsulated protocol.

If a customer has private or non-unique IP addresses, then a VPN service SHOULD be capable of translating such customer private or non-unique IP addresses for communicating with IP systems having public addresses.

4.3. Quality of Service

To the extent possible, L3VPN QoS should be independent of the access network technology.

4.3.1. QoS Standards

A non-goal of the L3VPN WG effort (as chartered) is the development of new protocols or extension of existing ones. An L3VPN shall be able to support QoS in one or more of the following already defined modes:

- Best Effort (mandatory support for all L3VPN types)
- Aggregate CE Interface Level QoS ("hose" level QoS)
- Site-to-site ("pipe" level QoS)
- Intserv (i.e., RSVP) signaled
- Diffserv marked
- Across packet-switched access networks

Note that all cases involving QoS may require that the CE and/or PE perform shaping and/or policing.

L3VPN CEs should be capable of supporting integrated services (Intserv) for certain customers in support of session applications, such as switched voice or video. Intserv-capable CE devices shall support the following Internet standards:

- Resource reSerVation Protocol (RSVP) [RFC2205]
- Guaranteed Quality of Service providing a strict delay bound [RFC2212]
- Controlled Load Service providing performance equivalent to that of an unloaded network [RFC2211]

L3VPN CE and PE should be capable of supporting differentiated service (Diffserv). Diffserv-capable L3VPN CE and PE shall support the following per hop behavior (PHB) [RFC2475] types:

- Expedited Forwarding (EF) - The departure rate of an aggregate class of traffic from a device that must equal or exceed a configured rate [RFC3246].
- Assured Forwarding (AF) - A means for a provider Diffserv (DS) domain to offer different levels of forwarding assurances for IP packets received from a customer DS domain. Four AF classes are defined, where each AF class implies allocation in each DS node of a certain amount of forwarding resources (e.g., buffer space and bandwidth) [RFC2597].

A CE or PE device supporting an L3VPN service may classify a packet for a particular Intserv or Diffserv service based on one or more of the following IP header fields: protocol ID, source port number, destination port number, destination address, or source address.

For a specifiable set of Internet traffic, L3VPN devices should support Random Early Detection (RED) to provide graceful degradation in the event of network congestion.

4.3.2. Service Models

A service provider must be able to offer QoS service to a customer for at least the following generic service types: managed-access VPN service or edge-to-edge QoS VPN service [RFC3809]. More detail specific to L3VPNs is provided below.

A managed-access L3VPN service provides QoS on the access connection between the CE and the PE. For example, diffserv would be enabled only on the CE router and the customer-facing ports of the PE router. Note that this service would not require Diffserv implementation in the SP backbone. The SP may use policing for inbound traffic at the PE. The CE may perform shaping for outbound traffic. Another example of a managed-access L3VPN service is when the SP performs the packet classification and diffserv marking. An SP may provide several packet classification profiles that customers may select or may offer custom profiles based on customer specific requirements. In general, more complex QoS policies should be left to the customer for implementation.

An edge-to-edge QoS VPN service provides QoS from edge device to edge device. The edge device may be either PE or CE, depending on the service demarcation point between the provider and the customer. Such a service may be provided across one or more provider backbones.

The CE requirements for this service model are the same as the managed access VPN service. However, in this service QoS is provided from one edge of the SP network(s) to the other.

4.4. Service-Level Specification and Agreements

A generic discussion of SLAs is provided in [RFC3809]. Additionally, SLS measurements for quality based on the DiffServ scheme SHOULD be based on the following classification:

- A Point-to-Point SLS [Y.1311.1], sometimes also referred to as the "Pipe" model, defines traffic parameters in conjunction with the QoS objectives for traffic exchanged between a pair of VPN sites (i.e., points). A Point-to-Point SLS is analogous to the SLS typically supported over point-to-point Frame Relay or ATM PVCs or an edge-to-edge MPLS tunnel. The set of SLS specifications to all other reachable VPN sites would define the overall Point-to-Point SLS for a specific site.
- A Point-to-Cloud SLS [Y.1311.1], sometimes also referred to as the "Hose" model, defines traffic parameters in conjunction with the QoS objectives for traffic exchanged between a CE and a PE for traffic destined to a set (either all or a subset) of other sites in the VPN (i.e., the cloud), as applicable. In other words, a point-to-cloud SLS defines compliance in terms of all packets transmitted from a given VPN site toward the SP network on an aggregate basis (i.e., regardless of the destination VPN site of each packet).
- A Cloud-to-Point SLS (a case not covered by this SLS is where flows originating from multiple sources may congest the interface toward a specific site).

Traffic parameters and actions SHOULD be defined for packets to and from the demarcation between the service provider and the site. For example, policing may be defined on ingress, and shaping on egress.

4.5. Management

An SP and its customers MUST be able to manage the capabilities and characteristics of their VPN services. To the extent possible, automated operations and interoperability with standard management platforms SHOULD be supported.

The ITU-T Telecommunications Management Network (TMN) model has the following generic requirements structure:

- 0 Engineer, deploy, and manage the switching, routing, and transmission resources supporting the service, from a network perspective (network element management).
- 0 Manage the VPN networks deployed over these resources (network management).
 - o Manage the VPN service (service management).
 - o Manage the VPN business, mainly provisioning administrative and accounting information related to the VPN service customers (business management).

Service management should include the TMN 'FCAPS' functionalities, as follows: Fault, Configuration, Accounting, Provisioning, and Security, as detailed in section 7.

4.6. Interworking

Interworking scenarios among different solutions providing L3VPN services is highly desirable. See the L3VPN framework document for more details on interworking scenarios [L3VPN-FR]. Interworking SHOULD be supported in a scalable manner.

Interworking scenarios MUST at least consider traffic and routing isolation, security, QoS, access, and management aspects. This requirement is essential of network migration, to ensure service continuity among sites belonging to different portions of the network.

5. Customer Requirements

This section captures additional requirements from a customer perspective.

5.1. VPN Membership (Intranet/Extranet)

When an extranet is formed, a customer agent from each of the organizations first approves addition of a site to an extranet VPN as a business decision between the parties involved. The solution SHOULD provide a means for these organizations to control extranet communication involving the L3VPN exchange of traffic and routing information.

5.2. Service Provider Independence

Customers MAY require VPN service that spans multiple administrative domains or service provider networks. Therefore, a VPN service MUST be able to span multiple AS and SP networks, but still act and appear as a single, homogeneous VPN from a customer point of view.

A customer might also start with a VPN provided in a single AS with a certain SLA but then ask for an expansion of the service, spanning multiple ASes/SPs. In this case, as well as for all kinds of multi-AS/SP VPNs, VPN service SHOULD be able to deliver the same SLA to all sites in a VPN regardless of the AS/SP to which it homes.

5.3. Addressing

A customer requires support from an L3VPN for the following addressing IP assignment schemes:

- o Customer-assigned, non-unique, or [RFC1918] private addresses
- o Globally unique addresses obtained by the customer
- o Globally unique addresses statically assigned by the L3VPN service provider
- o On-demand, dynamically assigned IP addresses (e.g., DHCP), irrespective of whether the access is temporary (e.g., remote) or permanent (e.g., dedicated)

In the case of combined L3VPN service with non-unique or private addresses and Internet access, mechanisms that permit the exchange of traffic between the customer's address space and the global unique Internet address space MAY be supported. For example, NAT is employed by many customers and by some service providers today to meet this need. A preferred solution would be to assign unique addresses, either IPv4 or IPv6; however, some customers do not want to renumber their networks.

5.4. Routing Protocol Support

There SHOULD be no restriction on the routing protocols used between CE and PE routers, or between CE routers. At least the following protocols MUST be supported: static routing, IGP protocols such as RIP, OSPF, IS-IS, and BGP [L3VPN-FR].

5.5. Quality of Service and Traffic Parameters

QoS is expected to be an important aspect of an L3VPN service for some customers. QoS requirements cover scenarios involving an intranet, an extranet, and shared access between a VPN site and the Internet.

5.5.1. Application-Level QoS Objectives

A customer is concerned primarily that the L3VPN service provides his or her applications with the QoS and level of traffic so that the applications perform acceptably. Voice, interactive video, and multimedia applications are expected to require the most stringent QoS. These real-time applications are sensitive to delay, delay variation, loss, availability, and/or reliability. Another set of applications, including some multimedia and interactive video applications, high-performance web browsing, and file transfer intensive applications, requires near real time performance. Finally, best effort applications are not sensitive to degradation, that is they are elastic and can adapt to conditions of degraded performance.

The selection of appropriate QoS and service type to meet specific application requirements is particularly important to deal with periods of congestion in an SP network. Sensitive applications will likely select per-flow Integrated service (Intserv) with precise SLA guarantees measured on a per-flow basis. On the other hand, non-sensitive applications will likely rely on a Diffserv class-based QoS.

The fundamental customer application requirement is that an L3VPN solution MUST support both the Intserv QoS model for selected individual flows and Diffserv for aggregated flows.

A customer application SHOULD experience consistent QoS independent of the access network technology used at different sites connected to the same VPN.

5.5.2. DSCP Transparency

The Diffserv Code Point (DSCP) set by a user as received by the ingress CE SHOULD be capable of being relayed transparently to the egress CE (see section 2.6.2 of [RFC3270] and [Y.1311.1]). Although RFC 2475 states that interior or boundary nodes within a DS domain can change the DSCP, customer VPNs MAY have other requirements, such as

- o applications that use the DSCP in a manner differently from the DSCP solution supported by the SP network(s),
- o customers using more DSCPs within their sites than the SP network(s) supports,
- o support for a carrier's carrier service in which one SP is the customer of another L3VPN SP. Such an SP should be able to resell VPN service to his or her VPN customers independently of the DSCP mapping solution supported by the carrier's carrier SP.

Note that support for DSCP transparency has no implication on the QoS or SLA requirements. If an SP supports DSCP transparency, then that SP needs to carry only the DSCP values across its domain but MAY map the received DSCP to some other value for QoS support across its domain.

5.6. Service-Level Specification/Agreement

Most customers simply want their applications to perform well. An SLA is a vehicle for customer recourse in the event that SP(s) do not perform or manage a VPN service well in a measurable sense. Therefore, when purchasing service under an SLA, a customer agent

MUST have access to the measures from the SP(s) that support the SLA.

5.7. Customer Management of a VPN

A customer MUST have a means to view the topology, operational state, order status, and other parameters associated with his or her VPN.

Most aspects of management information about CE devices and customer attributes of an L3VPN manageable by an SP SHOULD be capable of being configured and maintained by an authenticated, authorized customer agent. However, some aspects, such as encryption keys, SHALL NOT be readable nor writable by management systems.

A customer agent SHOULD be able to make dynamic requests for changes to traffic parameters. A customer SHOULD be able to receive real-time response from the SP network in response to these requests. One example of such service is a "Dynamic Bandwidth management" capability that enables real-time response to customer requests for changes of allocated bandwidth allocated to his or her VPN [Y.1311.1].

A customer who may not be able to afford the resources to manage his own sites SHOULD be able to outsource the management of the entire VPN to the SP(s) supporting the VPN network.

5.8. Isolation

These features include traffic and routing information exchange isolation, similar to that obtained in VPNs based on Layer 1 and Layer 2 (e.g., private lines, FR, or ATM) [MPLSSEC].

5.9. Security

The suite of L3VPN solutions SHOULD support a range of security related features. Higher levels of security services, such as edge-to-edge encryption, authentication, or replay attack, should be supported. More details on customer requirements for security are described in [VPNSEC].

Security in an L3VPN service SHOULD be as transparent as possible to the customer, with the obvious exception of support for remote or temporary user access, as detailed in section 5.11.2.

L3VPN customers MUST be able to deploy their own internal security mechanisms in addition to those deployed by the SP, in order to secure specific applications or traffic at a granularity finer than that on a site-to-site basis.

If a customer requires QoS support in an L3VPN, then this request MUST be communicated to the SP either by using unencrypted fields or via an agreed security association. For example, applications could send RSVP messages in support of Intserv either in the clear or encrypted with a key negotiated with the SP. Another case is that where applications using an IPsec tunnel could copy the DSCP from the encrypted IP header to the header of the tunnel's IP header.

5.10. Migration Impact

Often, customers are migrating from an already deployed private network toward one or more L3VPN solutions. A typical private network scenario is CE routers connected via real or virtual circuits. Ideally, minimal incremental cost SHOULD result during the migration period. Furthermore, if necessary, any disruption of service SHOULD also be minimized.

A range of scenarios of customer migration MUST be supported. Full migration of all sites MUST be supported. Support for cases of partial migration is highly desirable [Y.1311.1] - that is, legacy private network sites that belong to the L3VPN service SHOULD still have L3 reachability to the sites that migrate to the L3VPN service.

5.11. Network Access

Every L3 packet exchanged between the customer and the SP over the access connection MUST appear as it would on a private network providing an equivalent service to that offered by the L3VPN.

5.11.1. Physical/Link Layer Technology

L3VPNs SHOULD support a broad range of physical and link-layer access technologies, such as PSTN, ISDN, xDSL, cable modem, leased line, Ethernet, Ethernet VLAN, ATM, Frame Relay, Wireless local loop, and mobile radio access. The capacity and QoS achievable may be dependent on the specific access technology in use.

5.11.2. Temporary Access

The VPN service offering SHOULD allow both permanent and temporary access to one or more L3VPNs for authenticated users across a broad range of access technologies. Support for remote or temporary VPN access SHOULD include ISDN, PSTN dial-in, xDSL, or access via another SP network. The customer SHOULD be able to choose from alternatives for authentication of temporary access users. Choices for access authentication are SP-provided, third-party, or customer-provided authentication.

A significant number of VPN users may not be permanently attached to one VPN site: in order to limit access to a VPN to authorized users, it is first necessary to authenticate them. Authentication SHALL apply as configured by the customer agent and/or SP where a specific user may be part of one or more VPNs. The authentication function SHOULD be used to invoke all actions necessary to join a user to the VPN automatically.

A user SHOULD be able to access an L3VPN via a network having generic Internet access.

Mobile users may move within an L3VPN site. Mobile users may also have temporary connections to different L3VPN sites within the same VPN. Authentication SHOULD be provided in both of these cases.

5.11.3. Sharing of the Access Network

In a PE-based L3VPN, if the site shares the access network with other traffic (e.g., access to the Internet), then data security in the access network is the responsibility of the L3VPN customer.

5.11.4. Access Connectivity


Various types of physical connectivity scenarios MUST be supported, such as multi-homed sites, backdoor links between customer sites, and devices homed to two or more SP networks. L3VPN solutions SHOULD support at least the types of physical or link-layer connectivity arrangements shown in Figure 2.1. Support for other physical connectivity scenarios with arbitrary topology is desirable.

Access arrangements with multiple physical or logical paths from a CE to other CEs and PEs MUST support redundancy and SHOULD support load balancing. Resiliency uses redundancy to provide connectivity between a CE site and other CE sites and, optionally, other services. Load balancing provides a means to perform traffic engineering so that capacity on redundant links is used to achieve improved performance during periods when the redundant component(s) are available.

For multi-homing to a single SP, load balancing capability SHOULD be supported by the PE across the CE to PE links. For example, in case (a), load balancing SHOULD be provided by the two PEs over the two links connecting to the single CE. In case (c), load balancing SHOULD be provided by the two PEs over the two links connecting to the two CEs.

In addition, the load-balancing parameters (e.g., the ratio of traffic on the multiple load-balanced links, or the preferred link) SHOULD be provisionable based on customer's requirements. The load-balancing capability may also be used to achieve resiliency in the event of access connectivity failures. For example, in case (b) a CE may connect to two different SPs via diverse access networks. Resiliency MAY be further enhanced as shown in case (d), where CEs connected via a "back door" connection connect to different SPs. Furthermore, arbitrary combinations of the above methods, with a few examples shown in cases (e) and (f), should be supportable by any L3VPN approach.

For multi-homing to multiple SPs, load balancing capability MAY also be supported by the PEs in the different SPs (clearly, this is a more complex type of load balancing to realize, requiring policy and service agreements between the SPs to interoperate).






Figure 2.1. Representative types of access arrangements

5.12. Service Access

Customers MAY also require access to other services, as described in this section.

5.12.1. Internet Access

Customers SHOULD be able to have L3VPN and Internet access across the same access network for one or more of the customer's sites.

Customers SHOULD be able to direct Internet traffic from the set of sites in the L3VPN to one or more customer sites that have firewalls, other security-oriented devices, and/or NATs that process all traffic between the Internet and the customer's VPN.

L3 VPN Customers SHOULD be able to receive traffic from the Internet addressed to a publicly accessible resource that is not part of the VPN, such as an enterprise's public web server.

As stated in section 5.3, if a customer L3VPN employs private or non-unique IP addresses, then network address translation (NAT) or a similar mechanism MUST be provided either by the customer or the SP in order to allow traffic exchange with devices outside the customer's L3VPN.

5.12.2. Hosting, Application Service Provider

A customer SHOULD be able to access hosting, other application services, or other Application Service Providers (ASP) over an L3 L3VPN service. This MAY require that an ASP participate in one or more VPNs with the customers that use such a service.

5.12.3. Other Services

In conjunction with a VPN service, a customer MAY also wish to have access to other services, such as DNS, FTP, HTTP, NNTP, SMTP, LDAP, VoIP, NAT, LDAP, Videoconferencing, Application sharing, E-business, Streaming, E-commerce, Directory, Firewall, etc. The resources that implement these services could be physically dedicated to each VPN. If the resources are logically shared, then they MUST have access separated and isolated between VPNs in a manner consistent with the L3VPN solution to meet this requirement.

5.13. Hybrid VPN Service Scenarios

Intranet or extranet customers have a number of reasons for wanting hybrid networks that involve more than one VPN solution type. These include migration, mergers, extranet customers with different VPN types, the need for different capabilities between different sets of sites, temporary access, and different availability of VPN solutions as provided by different service providers.

The framework and solution approaches SHOULD include provisions for interworking, interconnection, and/or reachability between different

L3VPN solutions in a way that does not overly complicate provisioning, management, scalability, or performance.

6. Service Provider Network Requirements

This section describes requirements from a service provider perspective.

6.1. Scalability

[RFC3809] lists projections of L3VPN sizing and scalability requirements and metrics related to specific solutions.

6.2. Addressing

As described in section 4.2, SPs MUST have support for public and private IP addresses, IPv4 and IPv6, for both unicast and multicast. In order to support this range of addressing schemes, SPs require the following support from L3VPN solutions.

An L3VPN solution MUST be able to assign blocks of addresses from its own public IP address space to L3VPN customer sites so that advertisement of routes to other SPs and other sites aggregates efficiently.

An L3VPN solution MUST be able to use address assignments made by a customer. These customer-assigned addresses may be public or private.

If private IP addresses are used, an L3VPN solution MUST provide a means for an SP to translate such addresses to public IP addresses for communication with other VPNs by using overlapping addresses or the Internet.

6.3. Identifiers

A number of identifiers MAY be necessary for SP use in management, control, and routing protocols. Requirements for at least the following identifiers are known.

An SP domain MUST be uniquely identified at least within the set of all interconnected SP networks when supporting a VPN that spans multiple SPs. Ideally, this identifier should be globally unique (e.g., an AS number).

An identifier for each VPN SHOULD be unique, at least within each SP's network. Ideally, the VPN identifier SHOULD be globally unique to support the case where a VPN spans multiple SPs (e.g., [RFC2685]).

A CE device SHOULD have a unique identifier, at least within each SP's network.

A PE device SHOULD have a unique identifier, at least within each SP's network.

The identifier of a device interconnecting SP networks MUST be unique within the set of aforementioned networks.

Each site interface SHOULD have a unique identifier, at least within each PE router supporting such an interface.

Each tunnel SHOULD have a unique identifier, at least within each router supporting the tunnel.

6.4. Discovering VPN Related Information

Configuration of CE and PE devices is a significant task for a service provider. Solutions SHOULD strive to contain methods that dynamically allow VPN information to be discovered (or learned) by the PE and/or CE to reduce configuration complexity. The following specific requirements apply to intra- and inter-provider VPNs [VPNDISC].

Every device involved in a VPN SHALL be able to identify and authenticate itself to other devices in the VPN. After learning the VPN membership, the devices SHOULD be able to exchange configuration information securely. The VPN information MUST include at least the IP address of the PE and may be extensible to provide additional information.

Each device in a VPN SHOULD be able to determine which other devices belong to the same VPN. Such a membership discovery scheme MUST prevent unauthorized access and allow authentication of the source.

Distribution of VPN information SHOULD be limited to those devices involved in that VPN.

In the case of a PE-based VPN, a solution SHOULD support the means for attached CEs to authenticate each other and verify that the SP's VPN network is correctly configured.

The mechanism SHOULD respond to VPN membership changes in a timely manner. This is no longer than the provisioning timeframe, typically on the order of minutes, and could be as short as the timeframe required for "rerouting", typically on the order of seconds.

Dynamically creating, changing, and managing multiple VPN assignments to sites and/or customers is another aspect of membership that MUST be addressed in an L3VPN solution.

6.5. SLA and SLS Support

Typically, a Service Provider offering an L3VPN service commits to specific Service Level Specifications (SLS) as part of a contract with the customer, as described in section 4.4 and [RFC3809]. Such a Service Level Agreement (SLA) implies SP requirements for measuring Specific Service Level Specifications (SLS) for quality, availability, response time, and configuration intervals.

6.6. Quality of Service (QoS) and Traffic Engineering

A significant aspect of an L3VPN is support for QoS. Since an SP has control over the provisioning of resources and configuration of parameters in at least the PE and P devices and, in some cases, in the CE device as well, the onus is on the SP to provide either managed QoS access service, or edge-to-edge QoS service, as defined in section 4.3.2.

Each L3VPN approach **MUST** describe the traffic engineering techniques available for an SP to meet the QoS objectives. These descriptions of traffic engineering techniques **SHOULD** quantify scalability and achievable efficiency. Traffic engineering support **MAY** be on an aggregate or per-VPN basis.

QoS policies **MUST** not be impacted by security mechanisms. For example, Diffserv policies **MUST** not be impacted by the use of IPSec tunnels using the mechanisms explained in RFC 2983 [RFC2983].

As stated in RFC 2475, a mapping function from customer provided Diffserv marking to marking used in an SP network should be provided for L3 VPN services.

If a customer requires DSCP transparency, as described in section 5.5.2, an L3VPN service **MUST** deliver the same value of DSCP field in the IP header received from the customer to the egress demarcation of the destination.

6.7. Routing

The distribution of reachability and routing policy **SHOULD** be constrained to the sites that are members of the VPN.

Optionally, the exchange of such information **MAY** use some form of authentication (e.g., MD5).

Functions to isolate the SP network and customer VPNs from anomalous routing behavior from a specific set of customer sites **SHOULD** be provided. Examples of such functions are controls for route flap dampening, filters that accept only prefixes configured for a specific CE, a maximum number of routes accepted for each CE, or a maximum rate at which route updates can be received from a CE.

When VPN customers use overlapping non-unique IP addresses, the solution **MUST** define a means to distinguish between such overlapping addresses on a per-VPN basis.

Furthermore, the solution SHOULD provide an option that either allows or prevents advertisement of VPN routes to the Internet.

Ideally, the choice of an SP's IGP SHOULD not depend on the routing protocol(s) used between PE and CE routers in a PE-based VPN.

Furthermore, it is desirable that an SP SHOULD have a choice regarding the IGP routing protocol.

The additional routing burden that an SP must carry should be articulated in each specific L3VPN solution.

6.8. Isolation of Traffic and Routing

The internal structure of an L3VPN network SHOULD not be visible to outside networks (e.g., the Internet or any connected VPN).

From a high-level SP perspective, a PE-based L3VPN MUST isolate the exchange of traffic and routing information to only those sites that are authenticated and authorized members of a VPN.

In a CE-based VPN, the tunnels that connect the sites effectively meet this isolation requirement if both traffic and routing information flow over the tunnels.

An L3VPN solution SHOULD provide a means to meet L3VPN QoS SLA requirements that isolates VPN traffic from the effects of traffic offered by non-VPN customers. Also, L3VPN solutions SHOULD provide a means to isolate the effects that traffic congestion produced by sites as part of one VPN can have on another VPN.

6.9. Security

This section contains requirements related to securing customer flows; providing authentication services for temporary, remote, or mobile users; and protecting service provider resources involved in supporting an L3VPN. More detailed security requirements are provided in [VPNSEC].

6.9.1. Support for Securing Customer Flows

In order to meet the general requirement for providing a range of security options to a customer, each L3VPN solution MUST clearly spell out the configuration options that can work together and how they can do so.

When a VPN solution operates over a part of the Internet, it should support a configurable option to support one or more of the following standard IPsec methods for securing a flow for a specified subset of a customer's VPN traffic:

- o Confidentiality, so that only authorized devices can decrypt it
- o Integrity, to ensure that the data has not been altered
- o Authentication, to ensure that the sender is indeed who he or she claims to be
- o Replay attack prevention.

The above functions SHOULD be applicable to "data traffic" of the customer, which includes the traffic exchanged between sites between temporary users and sites, and even between temporary users. It SHOULD also be possible to apply these functions to "control traffic", such as routing protocol exchanges, that are not necessarily perceived by the customer but are nevertheless essential to maintain his or her VPN.

Furthermore, such security methods MUST be configurable between different end points, such as CE-CE, PE-PE, and CE-PE. It is also desirable to configure security on a per-route or per-VPN basis [VPNSEC].

A VPN solution MAY support one or more encryption schemes, including AES, and 3DES. Encryption, decryption, and key management SHOULD be included in profiles as part of the security management system.

6.9.2. Authentication Services

A service provider MUST provide authentication services in support of temporary user access requirements, as described in section 5.11.2.

Furthermore, traffic exchanged within the scope of VPN MAY involve several categories of equipment that must cooperate to provide the service [Y.1311.1]. These network elements can be CE, PE, firewalls, backbone routers, servers, management stations, etc. These network elements learn about each other's identity, either via manual configuration or via discovery protocols, as described in section 6.4. When network elements must cooperate, these network elements SHALL authenticate peers before providing the requested service. This authentication function MAY also be used to control access to network resources.

The peer identification and authentication function described above applies only to network elements participating in the VPN. Examples include:

- traffic between a CE and a PE,
- traffic between CEs belonging to the same VPN,
- CE or PE routers dealing with route announcements for a VPN,
- policy decision point [RFC3198] and a network element, and
- management station and an SNMP agent.

For a peer authentication function, each L3VPN solution SHOULD describe where necessary, how it shall be implemented, how secure it must be, and the way to deploy and maintain identification and authentication information necessary to operate the service.

6.9.3. Resource Protection

Recall from the definitions in section 3.3 that a site can be part of an intranet with sites from the only same organization, can be part of an extranet involving sites from other organizations, can have access to the Internet, or can have any combination of these scopes of communication. Within these contexts, a site might be subject to various attacks coming from different sources. Potential sources of attack include:

- users connected to the supporting public IP backbone,
- users from the Internet, and
- users from temporary sites belonging to the intranet and/or extranet VPN the site is part of.

Security threats and risks that a site may encounter include the following:

- Denial of service, for example mail spamming, access connection congestion, TCP SYN attacks, and ping attacks
- Intrusion attempts, which may eventually lead to denial of service (e.g., a Trojan horse attack).

Additional threat scenarios are defined in [VPNSEC]. An L3VPN solution MUST state how it addresses each potential threat scenario.

The devices in the L3VPN network must provide some means of reporting intrusion attempts to the service provider resources.

6.10. Inter-AS (SP)VPNs

The scenario for VPNs spanning multiple Autonomous Systems (AS) or Service Providers (SP) requires standard solutions. The scenario where multiple ASes are involved is the most general case and is therefore the one described here. The scenarios of concern are the CE-based and PE-based L3VPNs defined in section 3.

In each scenario, all applicable SP requirements, such as traffic and routing isolation, SLAs, management, security, and provisioning. MUST be preserved across adjacent ASes. The solutions MUST describe the inter-SP network interface, encapsulation method(s), routing protocol(s), and all applicable parameters [VPNIW].

An essential pre-condition for an inter-AS VPN is an agreement between the ASes involved that spells out at least trust, economic, and management responsibilities.

The overall scalability of the VPN service MUST allow the L3VPN service to be offered across potentially hundreds of SPs, with the overall scaling parameters per SP given in [RFC3809].

6.10.1. Routing Protocols

If the link between ASes is not trusted, routing protocols running between those ASes MUST support some form of authentication. For example, the TCP option for carrying an MD5 digest may be used to enhance security for BGP [RFC2385].

BGP MUST be supported as the standard inter-AS routing protocol to control the path taken by L3VPN traffic.

6.10.2. Management

The general requirements for managing a single AS apply to a concatenation of ASes. A minimum subset of such capabilities as follows:

- Diagnostic tools (e.g., ping, traceroute)
- Secured access to one AS management system by another
- Configuration request and status query tools
- Fault notification and trouble-tracking tools

6.10.3. Bandwidth and QoS Brokering

When a VPN spans multiple ASes, a brokering mechanism is desired that requests certain SLA parameters, such as bandwidth and QoS, from the other domains and/or networks involved in transferring traffic to various sites. Although bandwidth and QoS brokering across multiple ASes is not common in today's networks, these may be desirable for maintaining SLAs in inter-AS VPNs. This section describes requirements for features that would facilitate these mechanisms. The objective is that a solution SHOULD be able to determine whether a set of ASes can establish and guarantee uniform QoS in support of an L3VPN.

The brokering mechanism can be a manual one, for example, in which one provider requests from another a specific set of bandwidth and QoS parameters for traffic going to and from a specific set of sites. The mechanism could also be an automated one where a device dynamically requests and receives certain bandwidth and SLA/QoS parameters. For instance, in the case of an L3VPN over MPLS, a PE may negotiate the label for different traffic classes to reach a PE residing in a neighboring AS. Or, it might be a combination of both. For additional detailed requirements on the automated approach, see [TE-INTERAS].

Brokering on a per VPN basis is not desirable as this approach would not scale. A solution MUST provide some means to aggregate QoS and bandwidth brokering requests between ASes. One method could be for SPs to make an agreement specifying the maximum amount of bandwidth for specific QoS parameters for all VPN customers using the SP network. Alternatively, such aggregation might be on a per hierarchical tunnel basis between PE routers in different ASes supporting an L3VPN service [TE-INTERAS].

6.10.4. Security Considerations

If a tunnel traverses multiple SP networks and passes through an unsecured SP, POP, NAP, or IX, then security mechanisms MUST be employed. These security mechanisms include encryption, authentication, and resource protection, as described in section 6.9, and security management, as covered in section 7.5. For example, a provider should consider using both authentication and encryption for a tunnel used as part of an L3VPN that traverses another service provider's network.

6.11. L3VPN Wholesale

The architecture MUST support the possibility of one service provider offering VPN service to another service provider. Another example is when one service provider sells L3VPN service at wholesale to another service provider, who then resells that VPN service to his or her customers.

The wholesaler's VPN MUST be transparent to the addressing and routing used by the reseller.

Support for additional levels of hierarchy (for example, three levels at which a reseller can again resell the VPN service to yet another VPN provider) SHOULD be provided.

The Carrier's Carrier scenario is the term used in this document for this category of L3VPN wholesale (although some scenarios of Inter-

AS/Inter-Provider VPN could possibly fall in this L3VPN wholesale category, too). Various carrier's carrier scenarios should be supported, such as when

- the customer carriers do not operate L3VPN services for their clients;
- the customer carriers operate L3VPN services for their clients, but these services are not linked with the L3VPN service offered by the Carrier's Carrier and
- the customer carriers operate L3VPN services for their clients, and these services are linked with the L3VPN service offered by the Carrier's Carrier ("Hierarchical VPNs" scenario).

6.12. Tunneling Requirements

Connectivity between CE sites or PE devices in the backbone SHOULD use a range of tunneling technologies, such as L2TP, IPSEC, GRE, IP-in-IP, and MPLS.

To set up tunnels between routers, every router MUST support static configuration for tunneling and MAY support a tunnel setup protocol. If employed, a tunnel establishment protocol SHOULD be capable of conveying information such as the following:

- Relevant identifiers
- QoS/SLA parameters
- Restoration parameters
- Multiplexing identifiers
- Security parameters

There MUST be a means to monitor the following aspects of tunnels:

- Statistics, such as amount of time spent in the up and down state.
- Count of transitions between the up and down state.
- Events, such as transitions between the up and down states.

The tunneling technology used by the VPN Service Provider and its associated mechanisms for tunnel establishment, multiplexing, and maintenance MUST meet the requirements on scaling, isolation, security, QoS, manageability, etc.

6.13. Support for Access and Backbone Technologies

This section describes requirements for aspects of access and backbone network technologies from an SP point of view.

Some SPs MAY desire that a single network infrastructure suffices for all services, public IP, VPNs, traffic engineering, and differentiated services [L2VPN].

6.13.1. Dedicated Access Networks

Ideally, the L3VPN service SHOULD be independent of physical, link layer, or even network technology of the access network. However, the characteristics of access networks MUST be accounted for when the QoS aspects of SLAs for VPN service offerings are specified.

6.13.2. On-Demand Access Networks

Service providers SHOULD be able to support temporary user access, as described in section 5.11.2, by using dedicated or dial-in access network technology.

L3VPN solutions MUST support the case where a VPN user directly accesses the VPN service through an access network connected to the service provider. They MUST also describe how they can support the case where one or more other service provider networks are used for access to the service provider supporting the L3VPN service.

Ideally, all information necessary to identify and authenticate users for an intranet SHOULD be stored and maintained by the customer. In an extranet, one customer SHOULD be able to maintain the authentication server, or the customers involved in the extranet MAY choose to outsource the function to a service provider.

Identification and authentication information could be made available to the service provider for controlling access, or the service provider may query a customer maintained server. Furthermore, one SP may act as access for the SP providing the VPN service. If the access SP performs identification and authentication on behalf of the VPN SP, an agreement MUST be reached on a common specification.

Support for at least the following authentication protocols SHALL be supported: PAP, CHAP, and EAP, as they are currently used in a wide range of equipment and services.

6.13.3. Backbone Networks

Ideally, the backbone interconnecting SP, PE, and P devices SHOULD be independent of physical and link layer technology. Nevertheless, the characteristics of backbone technology MUST be taken into account when specifying the QoS aspects of SLAs for VPN service offerings.

6.14. Protection, Restoration

When primary and secondary access connections are available, an L3VPN solution MUST provide restoration of access connectivity whenever the primary access link from a CE site to a PE fails. This capability SHOULD be as automatic as possible, that is, the traffic should be directed over the secondary link soon after failure of the primary access link is detected. Furthermore, reversion to the primary link SHOULD be dynamic, if configured to do so [VPN-NEEDS].

As mentioned in section 5.11.4, in the case of multi-homing, the load balancing capability MAY be used to achieve a degree of redundancy in the network. In the case of failure of one or more (but not all) of the multi-homed links, the load balancing parameters MAY be dynamically adjusted to redirect the traffic rapidly from the failed link(s) to the surviving links. Once the failed link(s) is (are) restored, the original provisioned load balancing ratio SHOULD be restored to its value prior to the failure.

An SP SHOULD be able to deploy protection and restoration mechanisms within his or her backbone infrastructure to increase reliability and fault tolerance of the VPN service offering. These techniques SHOULD be scalable, and therefore should strive not to perform such function in the backbone on a per-VPN basis.

Appropriate measurements and alarms that indicate how well network protection and restoration mechanisms are performing MUST be supported.

6.15. Interoperability

Service providers are interested in interoperability in at least the following scenarios:

- Facilitating use of PE and managed CE devices within a single SP network.
- Implementing L3VPN services across two or more interconnected SP networks.

- Achieving interworking or interconnection between customer sites using different L3VPN approaches or different implementations of the same approach.

Each approach MUST describe whether any of the above objectives can be met. If an objective can be met, the approach MUST describe how such interoperability could be achieved. In particular, the approach MUST describe the inter-solution network interface, encapsulation method(s), routing protocol(s), security, isolation, management, and all other applicable aspects of the overall VPN solution provided [VPN IW].

6.16. Migration Support

Service providers MUST have a graceful means to migrate a customer with minimal service disruption on a site-by-site basis to an L3VPN approach.

If L3VPN approaches can interwork or interconnect, then service providers MUST have a graceful means to migrate a customer with minimal service disruption on a site-by-site basis whenever interworking or interconnection is changed.

7. Service Provider Management Requirements

A service provider MUST have a means to view the topology, operational state, order status, and other parameters associated with each customer's VPN. Furthermore, an SP MUST have a means to view the underlying logical and physical topology, operational state, provisioning status, and other parameters associated with the equipment providing the VPN service(s) to its customers.

Currently, proprietary methods are often used to manage VPNs. The additional expense associated with operators using multiple proprietary management methods (e.g., command line interface (CLI) languages) to access such systems is undesirable. Therefore, devices SHOULD provide standards-based interfaces wherever feasible.

The remainder of this section presents detailed SP management requirements for a Network Management System (NMS) in the traditional fault, configuration, accounting, performance, and security (FCAPS) management categories. Much of this text was adapted from ITU-T Y.1311.1.

7.1. Fault Management

Support for fault management includes:

- indication of customers impacted by failure,
- fault detection (incidents reports, alarms and failure visualization),
- fault localization (analysis of alarms reports and diagnostics),
- incident recording or logs (creation and follow-through of trouble tickets), and
- corrective actions (traffic, routing, and resource allocation).

As PE-based VPNs rely on a common network infrastructure, the NMS MUST provide a means to inform the provider of the VPN customers impacted by a failure in the infrastructure. The NMS SHOULD provide pointers to the related customer configuration information to aid in fault isolation and determining corrective action.

Detecting faults caused by configuration errors is desirable, because these may cause VPN service failure or may disrupt other requirements (e.g., traffic and routing isolation). This is a likely case of compromised security [VPNSEC]. Detection of such errors is inherently difficult because the problem involves more than one node and may reach across a global perspective. One approach could be a protocol that systematically checks whether all constraints and consistency checks hold among tunnel configuration parameters at the various end points.

A capability to verify L3 reachability within a VPN MUST be provided for diagnostic purposes.

A capability to verify the parameter configuration of a device supporting an L3VPN MUST be provided for diagnostic purposes.

7.2. Configuration Management

Overall, the NMS must support a configuration necessary to realize the desired L3-reachability of an L3VPN. Toward this end, an NMS MUST provide configuration management to provision at least the following L3VPN components: PE, CE, hierarchical tunnels, access connections, routing, and QoS, as detailed in this section. If shared access to the Internet is provided, then this option MUST also be configurable.

As VPN configuration and topology are highly dependent on a customer's organization, provisioning systems MUST address a broad range of customer-specific requirements. The NMS MUST ensure that

these devices and protocols are provisioned consistently and correctly.

Provisioning for adding or removing sites SHOULD be as localized and automated as possible.

Configuration management for VPNs, according to service templates defined by the provider MUST be supported. A service template contains fields that, when used, yield a definite service requirement or policy. For example, a template for an IPsec tunnel would contain fields such as tunnel end points, authentication modes, encryption and authentication algorithms, pre-shared keys (if any), and traffic filters. An SLA template would contain fields such as delay, jitter, and throughput and packet loss thresholds, as well as end points over which the SLA has to be satisfied. In general, a customer's service order can be regarded as a set of instantiated service templates. This set can, in turn, be regarded as the logical service architecture of the customer's VPN.

Service templates can also be used by the provider to define the service architecture of the provider's own network. For example, OSPF templates could contain fields such as the subnets that form a particular area, the area identifier, and the area type. BGP service template could contain fields that, when used, would yield a BGP policy, such as for expressing a preference about an exit router for a particular destination.

The set of service templates SHOULD be comprehensive in that it can capture all service orders in some meaningful sense.

The provider SHOULD provide means to translate service templates into device configurations so that associated services can be provisioned.

Finally, the approach SHOULD provide means to check whether a service order is correctly provisioned. This would represent one method of diagnosing configuration errors. Configuration errors can arise due to a variety of reasons: manual configuration, intruder attacks, and conflicting service requirements.

7.2.1. Configuration Management for PE-Based VPNs

Requirements for configuration management unique to a PE-based VPN are as follows:

- o The NMS MUST support configuration of at least the following aspects of L3 PE routers: intranet and extranet membership, CE routing protocol for each access connection, routing metrics, and tunnels.

- o The NMS SHOULD use identifiers for SPs, L3VPNs, PEs, CEs, hierarchical tunnels, and access connections, as described in section 6.3.
- o Tunnels MUST be configured between PE and P devices. This requires coordination of identifiers of tunnels, hierarchical tunnels, VPNs, and any associated service information, for example, a QoS/SLA service.
- o Routing protocols running between PE routers and CE devices MUST be configured per VPN.
- o For multicast service, multicast routing protocols MUST also be configurable.
- o Routing protocols running between PE routers and between PE and P routers MUST also be configured.
- o The configuration of a PE-based L3VPN MUST be coordinated with the configuration of the underlying infrastructure, including Layer 1 and 2 networks interconnecting components of an L3VPN.

7.2.2. Configuration Management for CE-Based VPN

Requirements for configuration management unique to a CE-based VPN are as follows:

- o Tunnels MUST be configured between CE devices. This requires coordination of identifiers of tunnels, VPNs, and any associated service information, for example, a QoS/SLA service.
- o Routing protocols running between PE routers and CE devices MUST be configured. For multicast service, multicast routing protocols MUST also be configurable.

7.2.3. Provisioning Routing

A means for a service provider to provision parameters for the IGP for an L3VPN MUST be provided. This includes link level metrics, capacity, QoS capability, and restoration parameters.

7.2.4. Provisioning Network Access

A service provider MUST have the means to provision network access between SP-managed PE and CE, as well as the case where the customer manages the CE.

7.2.5. Provisioning Security Services

When a security service is requested, an SP MUST have the means to provision the entities and associated parameters involved with the service. For example, for IPsec service, tunnels, options, keys, and other parameters must be provisioned at either the CE or the PE. In the case of an intrusion detection service, the filtering and detection rules must be provisioned on a VPN basis.

7.2.6. Provisioning VPN Resource Parameters

A service provider MUST have a means to provision resources associated with VPN services dynamically. For example, in a PE-based service, the number and size of virtual switching and forwarding table instances must be provisionable.

Dynamic VPN resource assignment is crucial for coping with the frequent change requests from customers (e.g., sites joining or leaving a VPN), as well as for achieving scalability. The PEs SHOULD be able to dynamically assign the VPN resources dynamically. This capability is especially important for dial and wireless VPN services.

If an SP supports a "Dynamic Bandwidth management" service, then the provisioning system MUST be able to make requested changes within the ranges and bounds specified in the SLA. Examples of SLA parameters are response time and probability of being able to service such a request.

7.2.7. Provisioning Value-Added Service Access

An L3VPN service provides controlled access between a set of sites over a common backbone. However, many service providers also offer a range of value-added services. (for example, Internet access, firewall services, intrusion protection, IP telephony and IP Centrex, application hosting, and backup). It is outside of the scope of this document to define whether and how these different services interact with the VPN to solve issues such as addressing, integrity, and security. However, the VPN service MUST be able to provide access to these various types of value-added services.

A VPN service SHOULD allow the SP to supply the customer with different kinds of standard IP services, such as DNS, NTP, and RADIUS, that are needed for ordinary network operation and management. The provider SHOULD be able to provide IP services to multiple VPN customers.

A firewall function MAY be required to restrict access to the L3VPN from the Internet [Y.1311].

A managed firewall service MUST be carrier grade. For redundancy and failure recovery, a means for firewall fail-over should be provided. Managed firewall services that may be provided include dropping specified protocol types, intrusion protection, and traffic-rate limiting against malicious attacks.

Managed firewalls MUST be supported on a per-VPN basis, although multiple VPNs may be supported by the same physical device (e.g., in a PE-based solution). Managed firewalls SHOULD be provided at the major access point(s) for the L3VPN. Managed firewall services may be embedded in CE or PE device or implemented in standalone devices.

The NMS SHOULD allow a customer to outsource the management of an IP networking service to the SP providing the VPN or to a third party.

The NMS SHOULD support collection of information necessary for optimal allocation of IP services in response to customer orders.

Reachability to and from the Internet to sites within a VPN MUST be configurable by an SP. This could be controlled by configuring routing policy to control distribution of VPN routes advertised to the Internet.

7.2.8. Provisioning Hybrid VPN Services

Configuration of interworking or interconnection between L3VPN solutions SHOULD be also supported. Ensuring that security and end-to-end QoS issues are provided consistently SHOULD be addressed.

7.3. Accounting

Many service providers require collection of measurements regarding resource usage for accounting purposes. The NMS MAY need to correlate accounting information with performance and fault management information to produce billing that takes into account SLA provisions for periods of time when the SLS is not met.

An L3VPN solution MUST describe how the following accounting functions can be provided:

- Measurements of resource utilization.
- collection of accounting information.
- storage and administration of measurements.

Some providers may require near - real time reporting of measurement information and may offer this as part of a customer network management service.

If an SP supports a "Dynamic Bandwidth management" service, then the dates, times, amounts, and interval required to perform requested bandwidth allocation change(s) MUST be traceable for monitoring and accounting purposes.

Solutions should state compliance with accounting requirements, as described in section 1.7 of RFC 2975 [RFC2975].

7.4. Performance Management

Performance management MUST support functions involved with monitoring and collecting performance data for devices, facilities, and services, as well as determining conformance to SLS, such as QoS and availability measurements.

Performance management SHOULD also support analysis of important aspects of an L3VPN, such as bandwidth utilization, response time, availability, QoS statistics, and trends based on collected data.

7.4.1. Performance Monitoring

The NMS MUST monitor device behavior to evaluate performance metrics associated with an SLA. Different measurement techniques may be necessary depending on the service for which an SLA is provided. Example services are QoS, security, multicast, and temporary access. These techniques MAY be either intrusive or non-intrusive depending on the parameters being monitored.

The NMS MUST also monitor aspects of the VPN not directly associated with an SLA, such as resource utilization, state of devices, and transmission facilities, as well as control of monitoring resources such as probes and remote agents at network access points used by customers and mobile users.

7.4.2. SLA and QoS Management Features

The NMS SHOULD support SLAs between an SP and the various VPN customers according to the corresponding SLSes by measurement of the indicators defined within the context of the SLA, on a regular basis.

The NMS SHOULD use the QoS parameter measurement definitions, techniques, and methods as defined by the IETF IP Performance Metrics (IPPM) working group for delay, loss, and delay variation.

The NMS SHOULD support allocation and measurement of end-to-end QoS requirements to QoS parameters for one or more VPN network(s).

Devices supporting L3VPN SLAs SHOULD have real-time performance measurements that have indicators and threshold crossing alerts. Such thresholds should be configurable.

7.5. Security Management

The security management function of the NMS MUST include management features to guarantee the security of devices, access connections, and protocols within the L3VPN network(s), as well as the security of customer data and control as described in section 6.9.

7.5.1. Resource Access Control

Resource access control determines the privileges that a user has to access particular applications and VPN network resources. Without such control, only the security of the data and control traffic is protected, leaving the devices providing the L3VPN network unprotected. Access control capabilities protect these devices to ensure that users have access only to the resources and applications they are authorized to use.

In particular, access to the routing and switching resources managed by the SP MUST be tightly controlled to prevent and/or effectively mitigate a malicious attack. More detailed requirements in this area are described in [VPNSEC].

7.5.2. Authentication

Authentication is the process of verifying that the sender is actually who he or she claims to be. The NMS MUST support standard methods for authenticating users attempting to access management services.

Scalability is critical, as the number of nomadic/mobile clients is increasing rapidly. The authentication scheme implemented for such deployments MUST be manageable for large numbers of users and VPN access points.

Strong authentication schemes SHALL be supported to ensure the security of both VPN access point-to-VPN access point (e.g., PE to PE in a PE-based case) and client-to-VPN access point (e.g., CE-to-PE in a PE-based case) communications. This is particularly important for preventing VPN access point spoofing, a situation where an attacker tries to convince a PE or CE that the attacker is the VPN access point. If an attacker can convince a PE or CE device of this,

then that device will send VPN traffic to the attacker (who could forward it to the true access point after compromising confidentiality or integrity). In other words, a non-authenticated VPN AP can be spoofed with a man-in-the-middle attack, because the endpoints never verify each other. A weakly authenticated VPN AP may be subject to such an attack. Strongly authenticated VPN APs are not subject to such attacks, because the man-in-the-middle cannot be authenticated as the real AP due to the strong authentication algorithms.

7.6. Basis and Presentation Techniques of Management Information

Each L3VPN solution approach MUST specify the management information bases (MIB) modules for the network elements involved in L3VPN services. This is an essential requirement in network provisioning. The approach SHOULD identify any information not contained in a standard MIB related to FCAPS that is necessary to meet a generic requirement.

An IP VPN (Policy) Information model, when available, SHOULD reuse the policy information models being developed in parallel for specific IP network capabilities [IM-REQ]. This includes the QoS Policy Information Model [QPIM] and the IPSEC Configuration Policy Model [IPSECIM]. The IP VPN Information model SHOULD provide the OSS with adequate "hooks" to correlate service level specifications with traffic data collected from network elements. The use of policies includes rules that control corrective actions taken by OSS components responsible for monitoring the network and ensuring that it meets service requirements.

Additional requirements on VPN information models are given in reference [IM-PPVPN]. In particular, an information model MUST allow an SP to change VPN network dimensions with minimal influence on provisioning issues. The adopted model SHOULD be applicable to both small/medium size and large-scale L3VPN scenarios.

Some service providers MAY require systems that visually, audibly, or logically present FCAPS information to internal operators and/or customers.

8. Security Considerations

Security considerations occur at several levels and dimensions within L3VPNs, as detailed within this document. This section provides a summary with references to detailed supporting information [L3VPN-SEC] [VPNSEC].

The requirements in this document separate traditional notions of security requirements, such as integrity, confidentiality, and authentication, from issues such as isolating (or separating) the exchange of VPN data and control traffic between specific sets of sites (as defined in sections 3.3 and 4.1). Further detail on security requirements is given from the customer and service provider perspectives in sections 5.9 and 6.9, respectively. Further detail on data and control traffic isolation requirements are given from the customer and service provider perspectives in sections 5.1 and 6.8, respectively.

Furthermore, requirements regarding management of security from a service provider perspective are described in section 7.5.

9. Acknowledgements

The authors of this document would like to acknowledge the contributions from the people who launched the work on VPN requirements inside ITU-T SG13 and the authors of the original IP VPN requirements and framework document [RFC2764], as well as Tom Worster, Ron Bonica, Sanjai Narain, Muneyoshi Suzuki, Tom Nadeau, Nail Akar, Derek Atkins, Bryan Gleeson, Greg Burns, and Frederic Le Garrec. The authors are also grateful to the helpful suggestions and direction provided by the technical advisors, Alex Zinin, Scott Bradner, Bert Wijnen, and Rob Coltun. Finally, the authors wish to acknowledge the insights and requirements gleaned from the many documents listed in the references section. Citations to these documents were made only where the authors believed that additional insight could be obtained from reading the source document.

10. References

10.1. Normative References

- [RFC3377] Hodges, J. and R. Morgan, "Lightweight Directory Access Protocol (v3): Technical Specification", RFC 3377, September 2002.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, R., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

- [RFC2211] Wroclawski, J., "Specification of the Controlled-Load Network Element Service", RFC 2211, September 1997.
- [RFC2212] Shenker, S., Partridge, C., and R. Guerin, "Specification of Guaranteed Quality of Service", RFC 2212, September 1997.
- [RFC2251] Wahl, M., Howes, T., and S. Kille, "Lightweight Directory Access Protocol (v3)", RFC 2251, December 1997.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Service", RFC 2475, December 1998.
- [RFC2597] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597, June 1999.
- [RFC2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.
- [RFC2685] Fox, B. and B. Gleeson, "Virtual Private Networks Identifier", RFC 2685, September 1999.
- [RFC3246] Davie, B., Charny, A., Bennet, J.C., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, March 2002.
- [RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.
- [RFC3809] Nagarajan, A., "Generic Requirements for Provider Provisioned Virtual Private Networks (PPVPN)", RFC 3809, June 2004.

10.2. Informative References

- [2547bis] Rosen, E., Rekhter, Y. et al., "BGP/MPLS VPNs", Work in Progress.
- [IM-PPVPN] Lago, P., et al., "An Information Model for Provider Provisioned Virtual Private Networks", Work in Progress.

- [IM-REQ] Iyer, M., et al., "Requirements for an IP VPN Policy Information Model", Work in Progress.
- [IPSECIM] Jason, J., "IPsec Configuration Policy Model", Work in Progress.
- [CE-PPVPN] De Clercq, J., Paridaens, O., Krywaniuk, A., Wang, C., "An Architecture for Provider Provisioned CE-based Virtual Private Networks using IPsec", Work in Progress.
- [IPSEC-PPVPN] Gleeson, B., "Uses of IPsec with Provider Provisioned VPNs", Work in Progress.
- [L2VPN] Rosen, E., et al., "An Architecture for L2VPNs", Work in Progress.
- [MPLSSEC] Behringer, M., "Analysis of the Security of the MPLS Architecture", Work in Progress.
- [PPVPN-TERM] Andersson, L., Madsen, T., "PPVPN Terminology", Work in Progress.
- [L3VPN-SEC] Fang, L., et al., "Security Framework for Provider Provisioned Virtual Private Networks", Work in Progress.
- [L3VPN-FR] Callon, R., Suzuki, M., et al. "A Framework for Layer 3 Provider Provisioned Virtual Private Networks", Work in Progress.
- [PPVPN-VR] Knight, P., Ould-Brahim, H., Gleeson, B., "Network based IP VPN Architecture using Virtual Routers", Work in Progress.
- [QPIM] Snir, Ramberg, Strassner, Cohen and Moore, "Policy QoS Information Model", Work in Progress.
- [RFC2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.
- [RFC2764] Gleeson, B., Lin, A., Heinanen, J., Armitage, G., and A. Malis, "A Framework for IP Based Virtual Private Networks", RFC 2764, February 2000.
- [RFC2975] Aboba, B., Arkko, J., and D. Harrington, "Introduction to Accounting Management", RFC 2975, October 2000.

- [RFC2983] Black, D., "Differentiated Services and Tunnels", RFC 2983, October 2000.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC3198] Westerinen, A., Schnizlein, J., Strassner, J., Scherling, M., Quinn, B., Herzog, S., Huynh, A., Carlson, M., Perry, J., and S. Waldbusser, "Terminology for Policy-Based Management", RFC 3198, November 2001.
- [TE-INTERAS] Zhang, R., Vasseur, J.P., "MPLS Inter-AS Traffic Engineering requirements", Work in Progress.
- [VPNDISC] Squire, M. et al., "VPN Discovery Discussions and Options", Work in Progress.
- [VPNIW] Kurakami, H., et al., "Provider-Provisioned VPNs Interworking", Work in Progress.
- [VPNSEC] De Clercq, J., et al., "Considerations about possible security extensions to BGP/MPLS VPN", Work in Progress.
- [VPNTUNNEL] Worster, T., et al., "A PPVPN Layer Separation: VPN Tunnels and Core Connectivity", Work in Progress.
- [VPN-CRIT] Yu, J., Jou, L., Matthews, A., Srinivasan, V., "Criteria for Evaluating VPN Implementation Mechanisms", Work in Progress.
- [VPN-NEEDS] Jacquenet, C., "Functional needs for the deployment of an IP VPN service offering : a service provider perspective", Work in Progress.
- [Y.1311.1] Carugi, M. (editor), "Network Based IP VPN over MPLS architecture", Y.1311.1 ITU-T Recommendation, July 2001.
- [Y.1311] Knightson, K. (editor), "Network based VPNs - Generic Architecture and Service Requirements", Y.1311 ITU-T Recommendation, March 2002.

Authors' Addresses

Marco Carugi (co-editor)
Nortel Networks
Parc d'activites de Magny-Chateaufort
Les Jeunes Bois - MS CTF 32B5 - Chateaufort
78928 YVELINES Cedex 9 - FRANCE

Email: marco.carugi@nortel.com

Dave McDysan (co-editor)
MCI
22001 Loudoun County Parkway
Ashburn, VA 20147, USA

Email: dave.mcdysan@mci.com

Luyuan Fang
AT&T
200 Laurel Ave - Room C2-3B35
Middletown, NJ 07748 USA

Email: Luyuanfang@att.com

Ananth Nagarajan
Juniper Networks

Email: ananth@juniper.net

Junichi Sumimoto
NTT Communications Corporation
3-20-2 Nishi-Shinjuku, Shinjuku-ku, Tokyo 163-1421, Japan

Email: j.sumimoto@ntt.com

Rick Wilder
Alcatel

Email: rick.wilder@alcatel.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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 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.

