

Network Working Group
Request for Comments: 4176
Category: Informational

Y. El Mghazli, Ed.
Alcatel
T. Nadeau
Cisco
M. Boucadair
France Telecom
K. Chan
Nortel
A. Gonguet
Alcatel
October 2005

Framework for Layer 3 Virtual Private Networks (L3VPN)
Operations and Management

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 a framework for the operation and management of Layer 3 Virtual Private Networks (L3VPNs). This framework intends to produce a coherent description of the significant technical issues that are important in the design of L3VPN management solutions. The selection of specific approaches, and making choices among information models and protocols are outside the scope of this document.

Table of Contents

1.	Introduction	2
1.1.	Terminology	2
1.2.	Management functions	4
1.3.	Reference Models	5
2.	Customer Service Operations and Management	7
2.1.	Customer Service Management Information Model	7
2.2.	Customer Management Functions	8
2.2.1.	Fault Management	8
2.2.2.	Configuration Management	9
2.2.3.	Accounting	9
2.2.4.	Performance Management	10
2.2.5.	Security Management	10
2.3.	Customer Management Functional Description	11
2.3.1.	L3VPN Service Offering Management	11
2.3.2.	L3VPN Service Order Management	12
2.3.3.	L3VPN Service Assurance	12
3.	Provider Network Manager	12
3.1.	Provider Network Management Definition	12
3.2.	Network Management Functions	13
3.2.1.	Fault Management	13
3.2.2.	Configuration Management	14
3.2.3.	Accounting	17
3.2.4.	Performance Management	17
3.2.5.	Security Management	17
4.	L3VPN Devices	18
4.1.	Information Model	18
4.2.	Communication	18
5.	Security Considerations	19
6.	Acknowledgements	19
7.	Normative References	19

1. Introduction

1.1. Terminology

In this document, the following terms are used and defined as follows:

VPN:

Virtual Private Network. A set of transmission and switching resources that will be used over a shared infrastructure to process the (IP) traffic that characterizes communication services between the sites or premises interconnected via this VPN. See [RFC4026].

L3VPN:

An L3VPN interconnects sets of hosts and routers based on Layer 3 addresses. See [RFC4026].

VPN Instance:

From a management standpoint, a VPN instance is the collection of configuration information associated with a specific VPN, residing on a PE router.

VPN Site:

A VPN customer's location that is connected to the Service Provider network via a CE-PE link, which can access at least one VPN.

VPN Service Provider (SP):

A Service Provider that offers VPN-related services.

VPN Customer:

Refers to a customer that bought VPNs from a VPN service provider.

Customer Agent:

Denotes the entity that is responsible for requesting VPN customer-specific information.

Service Level Agreement(SLA):

Contractual agreement between the Service Provider and Customer, which includes qualitative and quantitative metrics that define service quality guarantees and retribution procedures when service levels are not being met.

Service Level Specifications (SLS):

Internally-focused service performance specifications used by the Service Provider to manage customer service quality levels.

1.2. Management functions

For any type of Layer-3 VPN (PE or CE-based VPNs), having a management platform where the VPN-related information could be collected and managed is recommended. The Service and Network Management System may centralize information related to instances of a VPN and allow users to configure and provision each instance from a central location.

An SP must be able to manage the capabilities and characteristics of their VPN services. Customers should have means to ensure fulfillment of the VPN service to which they subscribed. To the extent possible, automated operations and interoperability with standard management protocols should be supported.

Two main management functions are identified:

A customer service management function:

This function provides the means for a customer to query, configure, and receive (events/alarms) customer-specific VPN service information. Customer-specific information includes data related to contact, billing, site, access network, IP address, routing protocol parameters, etc. It may also include confidential data, such as encryption keys. Several solutions could be used:

- * Proprietary network management system
- * SNMP manager
- * PDP function
- * Directory service, etc.

A provider network management function:

This function is responsible for planning, building, provisioning, and maintaining network resources in order to meet the VPN service-level agreements outlined in the SLA offered to the customer. This mainly consists of (1) setup and configuration of physical links, (2) provisioning of logical VPN service configurations, and (3) life-cycle management of VPN service, including the addition, modification, and deletion of VPN configurations.

1.3. Reference Models

- o Engineer, deploy and manage the switching, routing, and transmission resources supporting the service from a network perspective (network element management);
- o Manage the VPNs deployed over these resources (network management);
- o Manage the VPN service (service management);

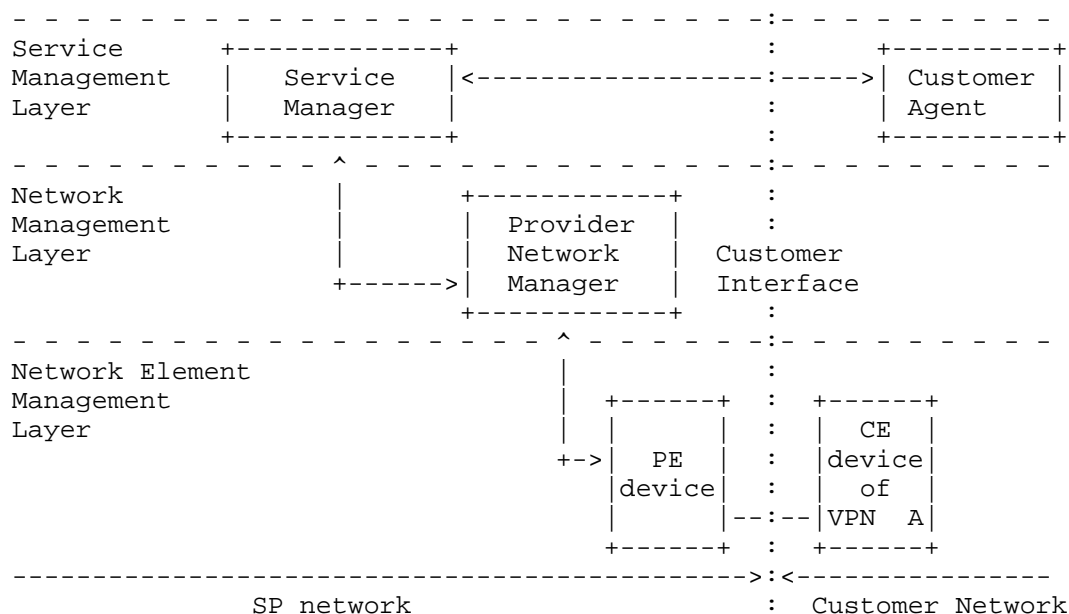


Figure 1: Reference Model for PE-based L3VPN Management

Figure 2: Reference Model for CE-based L3VPN Management

2. Customer Service Operations and Management

Services offered by providers can be viewed from the customer's or the provider's perspective. This section describes service management from the customer's perspective, focusing on the Customer Management function.

The Customer Management function's goal is to manage the service-based operations like service ordering, service subscription, activation, etc.

The Customer Management function resides in the L3VPN service manager at the Service Management Layer (SML). It mainly consists of defining the L3VPN services offered by the SP, collecting and consolidating the customer L3VPN services requirements, as well as performing some reporting for the customer. This function is correlated with the Network Management function at the Network Management Layer (NML) for initiating the L3VPN services provisioning, and getting some service reporting.

2.1. Customer Service Management Information Model

This section presents a framework that is used for L3VPN customer service management at the SML. The information framework represents the data that need to be managed, and the way they are represented. At the SML, the information framework that is foreseen is composed of Service Level Agreements (SLA) and Service Level Specifications (SLS).

Services are described through Service Level Agreements (SLA), which are contractual documents between customers and service providers. The technical part of the service description is called the Service Level Specification (SLS). The SLS groups different kinds of parameters. Some are more related to the description of the transport of the packets, and some to the specification of the service itself.

A Service Level Specification (SLS) may be defined per access network connection, per VPN, per VPN site, and/or per VPN route. The service provider may define objectives and the measurement intervals, for at least the SLS, using the following Service Level Objective (SLO) parameters:

- o QoS and traffic parameters
- o Availability for the site, VPN, or access connection
- o Duration of outage intervals per site, route, or VPN

- o Service activation interval (e.g., time to turn up a new site)
- o Trouble report response time interval
- o Time to repair interval
- o Total incoming/outgoing traffic from a site or a (VPN) route, or that has transited through the whole VPN
- o Measurement of non-conforming incoming/outgoing traffic (compliance of traffic should deserve some elaboration because of many perspectives - security, QoS, routing, etc.) from a site or a (VPN) route, or that has transited through the whole VPN

The service provider and the customer may negotiate contractual penalties in the case(s) where the provider does not meet a (set of) SLS performance objective(s).

Traffic parameters and actions should be defined for incoming and outgoing packets that go through the demarcation between the service provider premises and the customer's premises. For example, traffic policing functions may be activated at the ingress of the service provider's network, while traffic shaping capabilities could be activated at the egress of the service provider's network.

2.2. Customer Management Functions

This section presents detailed customer management functions in the traditional fault, configuration, accounting, performance, and security (FCAPS) management categories.

2.2.1. Fault Management

The fault management function of the Customer Service Manager relies upon the manipulation of network layer failure information, and it reports incidents to the impacted customers. Such reports should be based upon and related to the VPN service offering to which the customer is subscribed. The Customer Management function support for fault management includes:

- o Indication of customer's services impacted by failure
- o Incident recording or logs
- o Frequency of tests
- o Ability to invoke probes from the customer and provider

- o Ability to uncover faults before the customer notices them

2.2.2. Configuration Management

The configuration management function of the Customer Manager must be able to configure L3VPN service parameters with the level of detail that the customer is able to specify, according to service templates defined by the provider.

A service template contains fields which, when instantiated, yield a definite service requirement or policy. For example, a template for an IPsec tunnel [RFC2401] would contain fields such as tunnel end points, authentication modes, encryption and authentication algorithms, shared keys (if any), and traffic filters.

Other examples: a BGP/MPLS-based VPN service template would contain fields such as the customer premises that need to be interconnected via the VPN, and a QoS agreement template would contain fields such as one-way transit delay, inter-packet delay variation, throughput, and packet loss thresholds.

2.2.3. Accounting

The accounting management function of the Customer Manager is provided with network layer measurements information and manages this information. The Customer Manager is responsible for the following accounting functions:

- o Retrieval of accounting information from the Provider Network Manager
- o Analysis, 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 "Dynamic Bandwidth Management" service, then the schedule and the amount of the bandwidth 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 [RFC2975].

2.2.4. Performance Management

From the Customer Manager's perspective, performance management includes functions involved in the determination of the conformance level with the Service Level Specifications, such as QoS and availability measurements. The objective is to correlate accounting information with performance and fault management information to produce billing that takes into account SLA provisions for periods of time where the service level objectives are not met.

The performance information should reflect the quality of the subscribed VPN service as perceived by the customer. This information could be measured by the provider or controlled by a third party. The parameters that will be used to reflect the performance level could be negotiated and agreed upon between the service provider and the customer during the VPN service negotiation phase.

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.

2.2.5. Security Management

From the Customer Manager's perspective, the security management function includes management features to guarantee the security of the VPN. This includes security of devices, configuration data, and access connections. Authentication and authorization (access control) also fall into this category.

2.2.5.1. Access Control

Management access control determines the privileges that a user has for particular applications and parts of the network. Without such control, only the security of the data and control traffic is protected (leaving the devices providing the L3VPN network unprotected) among other equipment or resources. Access control capabilities protect these devices to ensure that users have access to only those resources and applications they are granted to use.

2.2.5.2. Authentication

Authentication is the process of verifying the identity of a VPN user.

2.3. Customer Management Functional Description

This section provides a high-level example of an architecture for the L3VPN management framework, with regard to the SML layer. The goal is to map the customer management functions described in Section 2.2 to architectural yet functional blocks, and to describe the communication with the other L3VPN management functions.

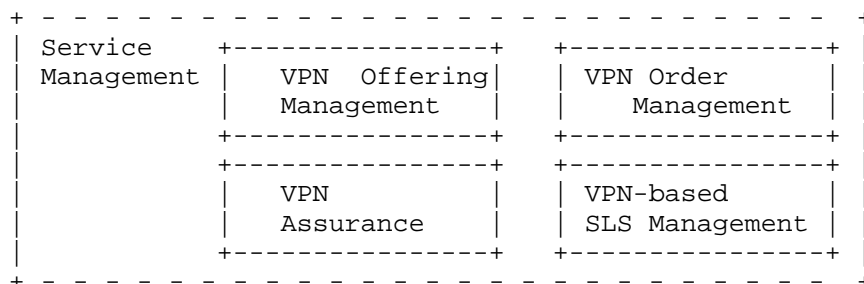


Figure 3: Overview of the Service Management

A customer must have a means to view the topology, operational state, order status, and other parameters associated with the VPN service offering that has been subscribed.

All aspects of management information about CE devices and customer attributes of an L3VPN, manageable by a SP, should be capable of being configured and maintained by an authenticated, authorized Service manager.

A customer agent should be able to make dynamic requests for changing the parameters that describe a service. A customer should be able to receive responses from the SP network in response to these requests (modulo the existence of necessary agreements). Communication between customer Agents and (VPN) service providers will rely upon a query/response mechanism.

A customer who may not be able to afford the resources to manage its CPEs should be able to outsource the management of the VPN to the service provider(s) supporting the network.

2.3.1. L3VPN Service Offering Management

Hopefully, the deployment of a VPN addresses customers' requirements. Thus, the provider must have the means to advertise the VPN-based services it offers. Then, the potential customers could select the service to which they want to subscribe. Additional features could be associated to this subscription phase, such as the selection of a

level of quality associated to the delivery of the VPN service, the level of management of the VPN service performed by the SP, security options, etc.

2.3.2. L3VPN Service Order Management

This operation aims at managing the requests initiated by the customers and tracks the status of the achievement of the related operations. The activation of the orders is conditioned by the availability of the resources that meet the customer's requirements with the agreed guarantees (note that it could be a result of a negotiation phase between the customer and the provider).

2.3.3. L3VPN Service Assurance

The customer may require the means to evaluate the fulfillment of the contracted SLA with the provider. Thus, the provider should monitor, measure, and provide statistical information to the customer, assuming an agreement between both parties on the measurement methodology, as well as the specification of the corresponding (set of) quality of service indicators.

3. Provider Network Manager

3.1. Provider Network Management Definition

When implementing a VPN architecture within a domain (or a set of domains managed by a single SP), the SP must have a means to view the physical and logical topology of the VPN premises, the VPN operational status, the VPN service ordering status, the VPN service handling, the VPN service activation status, and other aspects associated with each customer's VPN.

From a provider's perspective, the management of a VPN service consists mainly of:

- o Managing the customers (the term "customer" denotes a role rather than the end user, thus an SP could be a customer) and end-users in terms of SLA
- o Managing the VPN premises (especially creating, modifying, and deleting operations, editing the related information to a specific link, or supervising the AAA [RFC2903] [RFC2906] operations)
- o Managing the CE-PE links (particularly creating, modifying, and deleting links, editing the related information to a specific VPN)

- o Managing the service ordering, such as Quality of Service, in terms of supported classes of service, traffic isolation, etc.

Currently, proprietary methods are often used to manage VPNs. The additional expense associated with operators having to use multiple, proprietary, configuration-related management methods (e.g., Command Line Interface (CLI) languages) to access such systems is not recommended, because it affects the overall cost of the service (including the exploitation costs), especially when multiple vendor technologies (hence multiple expertise) are used to support the VPN service offering. Therefore, devices should provide standards-based interfaces. From this perspective, additional requirements on possible interoperability issues and availability of such standardized management interfaces need to be investigated.

3.2. Network Management Functions

In addition, there can be internal service provided by the SP for satisfying the customer service requirements. Some of these may include the notion of dynamic deployment of resources for supporting the customer-visible services, high availability service for the customer that may be supported by automatic failure detection, and automatic switchover to back-up VPNs. These are accomplished by inter-working with the FCAPS capabilities of the Provider Network Manager.

3.2.1. Fault Management

The Provider Network Manager support for fault management includes:

- o Fault detection (incidents reports, alarms, failure visualization)
- o Fault localization (analysis of alarms reports, diagnostics)
- o Corrective actions (data path, routing, resource allocation)

Since L3VPNs rely upon a common network infrastructure, the Provider Network Manager provides a means to inform the Service Manager about the VPN customers impacted by a failure in the infrastructure. The Provider Network Manager should provide pointers to the related customer configuration information to contribute to the procedures of fault isolation and the determination of corrective actions.

It is desirable to detect faults caused by configuration errors, because these may cause VPN service to fail, or not meet other requirements (e.g., traffic and routing isolation). One approach

could be a protocol that systematically checks that all constraints have been taken into account, and that consistency checks have been enforced during the tunnel configuration process.

A capability that aims at checking IP reachability within a VPN must be provided for diagnostic purposes.

A capability that aims at checking the configuration of a VPN device must be provided for diagnostic purposes.

3.2.2. Configuration Management

The Provider Network Manager must support configuration management capabilities in order to deploy VPNs. To do so, a Provider Network Manager must provide configuration management that provisions at least the following L3VPN components: PE, CE, hierarchical tunnels, access connections, routing, and QoS, as detailed in this section. If access to the Internet is provided, then this option must also be configurable.

Provisioning for adding or removing VPN customer premises should be as automated as possible.

Finally, the Provider Network Manager must ensure that these devices and protocols are provisioned consistently and correctly. The solution should provide a means for checking 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.

Requirements for L3VPN configuration management are:

- o The Provider Network Manager must support configuration of VPN membership.
- o The Provider Network Manager should use identifiers for SPs, L3VPNs, PEs, CEs, hierarchical tunnels, and access connections.
- o Tunnels must be configured between PE/CE devices. This requires coordination of tunnel identifiers, paths, VPNs, and any associated service information, for example, a QoS service.
- o Routing protocols running between PE routers and CE devices must be configured. For multicast services, multicast routing protocols must also be configurable.

- o Routing protocols running between PE routers, and between PE and P routers, must also be configured.

PE-based only:

- o Routing protocols running between PE routers and CE devices, if any, must be configured on a per-VPN basis. The Provider Network Manager must support configuration of a CE routing protocol for each access connection.
- o The configuration of a PE-based L3VPN should be coordinated with the configuration of the underlying infrastructure, including Layer 1 and 2 networks that interconnect components of an L3VPN.

3.2.2.1. Provisioning Routing-based Configuration Information

If there is an IGP running within the L3VPN, the Provider Network Manager must provision the related parameters. This includes metrics, capacity, QoS capability, and restoration parameters.

3.2.2.2. Provisioning Access-based Configuration Information

The Provider Network Manager must provision network access between SP-managed PE and CE equipment.

3.2.2.3. Provisioning Security Services-based Configuration Information

When a security service is requested, the Provider Network Manager must provision the entities and associated parameters involved in the provisioning of the service. For example, IPsec services, tunnels, options, keys, and other parameters should be provisioned at either the CE and/or the PE routers. In the case of an intrusion detection service, the filtering and detection rules should be provisioned on a VPN basis.

3.2.2.4. Provisioning VPN Resource Parameters

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

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

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 Service Level Specifications. Examples of QoS parameters are the response time and the probability of being able to service such a request.

Dynamic VPN resource allocation is crucial to cope with the frequent requests for changes that are expressed by customers (e.g., sites joining or leaving a VPN), as well as to achieve scalability. The PE routers should be able to dynamically assign the VPN resources. This capability is especially important for dial-up and wireless VPN services.

3.2.2.5. 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 detection, IP telephony and IP Centrex, application hosting, backup, etc. It is outside the scope of this document to define if and how these different services interact with the VPN service offering. However, the VPN service should 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 well-known IP services (e.g., DNS, NTP, RADIUS, etc.) needed for ordinary network operation and management. The provider should be able to provide IP services to multiple customers from one or many servers.

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

Managed firewalls may be supported on a per-VPN basis, although multiple VPNs will be supported by the same physical device. In such cases, managed firewalls should be provided at the access point(s) of the L3VPN. Such services may be embedded in the CE or PE devices, or implemented in stand-alone devices.

The Provider Network Manager should allow a customer to outsource the management of an IP service to the SP providing the VPN or to a third party.

The management system should support the collection of information necessary for optimal allocation of IP services in response to customers' orders, in correlation with provider-provisioned resources supporting the service.

If Internet access is provided, reachability to and from the Internet from/to sites within a VPN should be configurable by an SP. Configuring routing policy to control distribution of VPN routes advertised to the Internet may realize this.

3.2.2.6. Provisioning Hybrid VPN Services

Configuration of interworking L3VPN solutions should also be supported, taking security and end-to-end QoS issues into account.

3.2.3. Accounting

The Provider Network Manager is responsible for the measurements of resource utilization.

3.2.4. Performance Management

From the Provider Network Manager's perspective, performance management includes functions involved in monitoring and collecting performance data regarding devices, facilities, and services.

The Provider Network Manager must monitor the devices' behavior to evaluate performance metrics associated with an SLS. 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 Provider Network Manager must also monitor aspects of the VPN that are not directly associated with an SLS, such as resource utilization, status 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.

Devices supporting L3VPN whose level of quality is defined by SLSes should have real-time performance measurements that have indicators and threshold crossing alerts. Such thresholds should be configurable.

3.2.5. Security Management

From the Provider Network Manager's perspective, the security management function of the Provider Network Manager must include management features to guarantee the preservation of the confidentiality of customers' traffic and control data, as described in [RFC3809].

3.2.5.1. Authentication Management

The Provider Network Manager must support standard methods for authenticating users attempting to access VPN 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.

Support for strong authentication schemes needs to be supported to ensure the security of both VPN access point-to-VPN access point (PE to PE) and client-to-VPN Access point (CE-to-PE) communications. This is particularly important to prevent VPN access point (VPN AP) spoofing. VPN Access Point Spoofing is the situation where an attacker tries to convince a PE or a CE that the attacker is the VPN Access Point. If an attacker succeeds, then the device will send VPN traffic to the attacker (who could forward it on to the actual (and granted) access point after compromising confidentiality and/or integrity).

In other words, a non-authenticated VPN AP can be spoofed with a man-in-the-middle attack, because the endpoints rarely verify each other. A weakly authenticated VPN AP may be subject to such an attack. However, strongly authenticated VPN APs are not subject to such attacks, because the man-in-the-middle cannot authenticate as the real AP, due to the strong authentication algorithms.

4. L3VPN Devices

4.1. Information Model

Each L3VPN solution must specify the management information (MIBs, PIBs, XML schemas, etc.) for network elements involved in L3VPN services. This is an essential requirement in network provisioning. The approach should identify any L3VPN-specific information not contained in a standards track MIB module.

4.2. Communication

The deployment of a VPN may span a wide range of network equipment, potentially including equipment from multiple vendors. Therefore, the provisioning of a unified network management view of the VPN shall be simplified by means of standard management interfaces and models. This will also facilitate customer self-managed (monitored) network devices or systems.

In cases where significant configuration is required whenever a new service is to be provisioned, it is important, for scalability reasons, that the NMS provides a largely automated mechanism for the relevant configuration operations. Manual configuration of VPN services (i.e., new sites, or re-provisioning existing ones) could lead to scalability issues, and should be avoided. It is thus important for network operators to maintain visibility of the complete picture of the VPN through the NMS system. This should be achieved by using standards track protocols such as SNMP. Use of proprietary command-line interfaces is not recommended.

5. Security Considerations

This document describes a framework for L3VPN Operations and Management. Although this document discusses and addresses some security concerns in Section 2.2.5 and Section 3.2.5 above, it does not introduce any new security concerns.

6. Acknowledgements

Special Thanks to Nathalie Charton, Alban Couturier, Christian Jacquenet, and Harmen Van Der Linde for their review of the document and their valuable suggestions.

7. Normative References

- [RFC2975] Aboba, B., Arkko, J., and D. Harrington, "Introduction to Accounting Management", RFC 2975, October 2000.
- [RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [RFC2903] de Laat, C., Gross, G., Gommans, L., Vollbrecht, J., and D. Spence, "Generic AAA Architecture", RFC 2903, August 2000.
- [RFC2906] Farrell, S., Vollbrecht, J., Calhoun, P., Gommans, L., Gross, G., de Bruijn, B., de Laat, C., Holdrege, M., and D. Spence, "AAA Authorization Requirements", RFC 2906, August 2000.
- [RFC3809] Nagarajan, A., "Generic Requirements for Provider Provisioned Virtual Private Networks (PPVPN)", RFC 3809, June 2004.
- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

[Y.1311] ITU, "Network-based IP VPN over MPLS architecture",
ITU-T Y.1311.1, 2001.

Authors' Addresses

Yacine El Mghazli (Editor)
Alcatel
Route de Nozay
Marcoussis 91460
France

EMail: yacine.el_mghazli@alcatel.fr

Thomas D. Nadeau
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough, MA 01719

Phone: +1-978-936-1470
EMail: tnadeau@cisco.com

Mohamed Boucadair
France Telecom
42, rue des Coutures
Caen 14066
France

EMail: mohamed.boucadair@francetelecom.com

Kwok Ho Chan
Nortel Networks
600 Technology Park Drive
Billerica, MA 01821
USA

EMail: khchan@nortel.com

Arnaud Gonguet
Alcatel
Route de Nozay
Marcoussis 91460
France

EMail: arnaud.gonguet@alcatel.fr

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.

