

Internet Draft
Expiration Date: April 2003
Category: Informational

Ananth Nagarajan
Sprint
(Editor)
October 2002

Generic Requirements for Provider Provisioned VPN
<[draft-nagarajan-ppvnp-generic-reqts-01.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of \[RFC-2026\]](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts. Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Abstract

This document described generic requirements for Provider Provisioned Virtual Private Networks (PPVPN). The requirements are categorized into service requirements, provider requirements and engineering requirements. These requirements are not specific to any particular type of PPVPN technology, but rather apply to all PPVPN technologies.

Conventions used in this 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 [\[RFC-2119\]](#).

1. Introduction

1.1. Summary for Sub-IP area

This document is an output of the design team formed to develop requirements for PPVPNs in the PPVPN working group. As such this work fits within the scope of the PPVPN working group. This document discusses generic PPVPN requirements categorized as service, provider and engineering requirements. These are independent of any particular type of PPVPN technology.

1.2. Problem Statement

Corporations and other organizations have become increasingly dependent on their networks for tele- and datacommunication. The data communication networks were originally built as Local Area Networks (LAN). Over time the possibility to interconnect the networks on different sites has become more and more important. The connectivity for corporate networks has been supplied by service providers, mainly as FR or ATM connections, but over recent years also as Ethernet and IP-based tunnels. This type of network, interconnecting a number of sites over a shared network infrastructure is called Virtual Private Network (VPN). If the sites belong to the same organization, the VPN is called an Intranet. If the sites belong to different organizations that share a common interest, the VPN is called an Extranet.

Customers are looking for service providers to deliver data and telecom connectivity over one or more shared networks, with service level assurances in the form of security, QoS and other parameters.

In order to provide isolation between the traffic belonging to different customers, mechanisms such as Layer 2 connections or Layer 2/3 tunnels are necessary. When the shared infrastructure is an IP network, the tunneling technologies that are typically used are IPsec, MPLS, L2TP, GRE, IP-in-IP etc.

Traditional Internet VPNs have been based on IPsec to provide security over the Internet. Service providers are now beginning to deploy enhanced VPN services that provide features such as service differentiation, traffic management, Layer 2 and Layer 3 connectivity, etc. in addition to security. Newer tunneling mechanisms have certain features that allow the service providers to provide these enhanced VPN services.

The VPN solutions we define now must be able to accommodate the

traditional types of VPNs as well as the enhanced services now being deployed. They need to be able to run in a single service provider's network, as well as between a set of service providers and across the Internet. In doing so the VPNs should not be allowed to violate basic Internet design principles or overload the Internet core routers or accelerate the growths of the Internet routing tables. Specifically, Internet core routers shall not be required to maintain VPN-related information, regardless of whether the Internet routing protocols are used to distribute this information or not. In order to achieve this, the mechanisms used to develop various PPVPN solutions shall be as common as possible with generic Internet infrastructure mechanisms like discovery, signaling, routing and management. At the same time, existing Internet infrastructure mechanisms shall not be overloaded.

Another generic requirement from a standardization perspective is to limit the number of different solution approaches to a specific type of VPN to as small a number as possible.

1.3. Outline of this document

This document describes generic requirements for Provider Provisioned Virtual Private Networks (PPVPN). The document contains several sections, with each set representing a significant aspect of PPVPN requirements. [Section 2](#) lists authors who contributed to this document. [Section 3](#) defines terminology and presents a taxonomy of PPVPN technologies. The taxonomy contains two broad classes, representing Layer 2 and Layer 3 VPNs. Each top level VPN class contains subordinate classes. For

example, the Layer 3 VPN class contains a subordinate class of PE-based Layer 3 VPNs. Sections [4](#), [5](#), [6](#) describe generic PPVPN requirements.

The requirements are broadly classified under the following categories:

- 1) Service requirements - Service attributes that the customer can observe or measure. For example, does the service forward frames or route datagrams? What security guarantees does the service provide? Availability and stability are key requirements in this category.
- 2) Provider requirements - Characteristics that Service Providers use to determine the cost-effectiveness of a PPVPN service. Scaling and management are examples of Provider requirements.
- 3) Engineering requirements - Implementation characteristics that make service and provider requirements achievable. These can be

further classified as:

3a) Forwarding plane requirements - e.g., requirements related to router forwarding behavior.

3b) Control plane requirements - e.g., requirements related to reachability and distribution of reachability information.

3c) Requirements related to the commonality of PPVPN mechanisms with each other and with generic Internet mechanisms.


2. Contributing Authors

This document was the combined effort of several individuals that were part of the PPVPN requirements design team. A significant set of requirements were directly taken from previous work by the PPVPN WG to develop requirements for Layer 3 PPVPN [L3REQTS]. The following are the authors that contributed to this document:

Loa Andersson
Ron Bonica
Dave McDysan
Junichi Sumimoto
Muneyoshi Suzuki
David Meyer
Marco Carugi
Yetik Serbest
Luyuan Fang
Javier Achirica

3. Definitions and Taxonomy

The terminology used in this document is defined in [[TERMINOLOGY](#)].





The figure above presents a taxonomy of PPVPN technologies. Some of the definitions that are not covered in [[TERMINOLOGY](#)] are given below:

CE-based VPN: A VPN approach in which the shared service provider network does not have any knowledge of the customer VPN. This information is limited to CE equipment.

PE-Based VPNs: A layer 3 VPN approach in which a service provider network is used to interconnect customer sites using shared resources. Specifically the PE device maintains VPN state, isolating users of one VPN from users of another VPN. Because the PE device maintains all required VPN state, the CE device may behave as if it were connected to a private network. Specifically, the CE in a PE-based VPN must not require any changes or additional functionality to be connected to a PPVPN instead of a private network.

Virtual Router (VR) style: A PE-based VPN approach in which the PE router maintains a complete logical router for each VPN that it supports. Each logical router maintains a unique forwarding table and executes a unique instance of the routing protocols. These VPNs are described in [PPVPN-VR].

[RFC 2547](#) Style: A PE-based VPN approach in which the PE router maintains separate forwarding environment for each VPN and a separate forwarding table for each VPN. In order to maintain multiple forwarding table instances while running only a single routing protocol instance, [RFC 2547](#) style VPNs mark route advertisements with attributes that identify their VPN context. These VPNs are based on the approach described in [RFC2547bis].

4. Service requirements

These are the requirements that a customer can observe or measure, in order to verify if the PPVPN service that the Service Provider (SP) provides is satisfactory.

4.1. Availability

VPN services must have high availability. VPNs that are distributed over several sites require connectivity to be maintained even in the event of network failures or degraded service.

This can be achieved via various redundancy techniques such as:

1. Physical Diversity

A single site connected to multiple CEs (for CE-based PPVPN) or PEs (for PE-based PPVPNs), or different POPs, or even different service providers

or

2. via tunnel redundancy.

4.2. Stability

In addition to availability, VPN services must also be stable. Stability is a function of several components such as VPN routing, signaling and discovery mechanisms, in addition to tunnel stability. Stability of the VPN service is directly related to the stability of the mechanisms and protocols used to establish the service. It should also be possible to allow network upgrades and maintenance procedures without impacting the VPN service.

4.3. Traffic types

VPN services must support unicast (or point to point) traffic and should support any-to-any traffic including multicast and broadcast traffic. In the broadcast model, the network delivers a stream to all members of a subnetwork, regardless of their interest in that stream. In the multicast model, the network delivers a stream to a set of destinations that have registered interest in the stream. All destinations need not belong to the same subnetwork. Multicast is more appropriate for L3 VPNs while broadcast is more appropriate for

L2VPNs. It is desirable to support multicast limited in scope to an intranet or extranet. The solution should be able to support a large number of such intranet or extranet specific multicast groups in a scalable manner.

A PPVPN shall support both IPv4 and IPv6 traffic.

4.4. Data isolation

The PPVPN must support forwarding plane isolation. The network must never deliver user data across VPN boundaries unless the two VPNs participate in an intranet or extranet.

Furthermore, if the provider network receives signaling or routing information from one VPN, it must not reveal that information to another VPN unless the two VPNs participate in an intranet or extranet.

4.5. Security

A range of security features should be supported by the suite of PPVPN solutions in the form of securing customer flows, providing authentication services for temporary, remote or mobile users, and the need to protect service provider resources involved in supporting a PPVPN[VPN SEC]. Each PPVPN solution should state which security features it supports and how such features can be configured on a per customer basis. Protection against Denial of Service (DoS) attacks is a key component of security mechanisms. Examples of DoS attacks include mail spamming, access connection congestion, TCP SYN attacks, ping attacks and intrusion attempts such as Trojan horse attack.

4.5.1. User data security

PPVPN solutions that support user data security should use standard methods (e.g., IPsec) to achieve confidentiality, integrity, authentication and replay attack prevention. 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.

4.5.2. Access control

A PPVPN solution may also have the ability to activate the appropriate filtering capabilities upon request of a customer. A filter provides a mechanism so that access control can be invoked at the point(s) of communication between different organizations involved in an extranet. Access control can be implemented by a firewall, access control lists on routers or similar mechanisms to apply policy-based access control to transit traffic. Access control must also be applicable between CE-CE, PE-PE and CE-PE.

4.5.3. Site authentication and authorization

A PPVPN solution requires authentication and authorization of the following:

- temporary and permanent access for users connecting to sites (authentication and authorization BY the site)
- the site itself (authentication and authorization FOR the site)

4.5.4. Inter domain security

The VPN solution must have appropriate security mechanisms to prevent the different kinds of Distributed Denial of Service (DDoS) attacks mentioned earlier, misconfiguration or unauthorized accesses in inter domain PPVPN connections.

4.6. Topology

A VPN should support arbitrary, customer agent defined inter-site connectivity, ranging, for example, from hub-and-spoke, partial mesh to full mesh topology. These can actually be different from the topology used by the service provider. To the extent possible, a PPVPN service should be independent of the geographic extent of the deployment.

A VPN solution should support multiple VPNs per customer site without requiring additional hardware resources per VPN. It should also support a free mix of L2 and L3 VPNs per customer site.

To the extent possible, the PPVPN services should be independent of access network technology.

4.7. Addressing (support for private, overlapping addresses)

Each customer resource must be identified by an address that is unique within its VPN. It need not be identified by a globally unique address.

A VPN service shall be capable of supporting non-IP customer addresses if it is a Layer 2 VPN (e.g., Frame Relay, ATM, Ethernet). Support for non-IP Layer 3 addresses may be desirable in some cases, but is beyond the scope of VPN solutions developed in the IETF, and therefore, this document.

4.8. Quality of Service

A PPVPN shall be able to support QoS via IETF standardized mechanisms such as Diffserv. Support for best-effort traffic shall be mandatory for all PPVPN types.

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

The need to provide QoS will occur primarily in the access network, since that will often be the bottleneck. This is likely to occur since the backbone effectively statistically multiplexes many users, and is traffic engineered or includes capacity for restoration and growth. There are two directions of QoS management that must be considered in any PPVPN service regarding QoS:

- From the CE across the access network to the PE
- From the PE across the access network to CE

PPVPN CE and PE devices should be capable of supporting QoS across at least the following subset of access networks, although, to the extent possible, the QoS capability of a PPVPN should be independent of the access network technology :

- ATM Virtual Connections (VCs)
- Frame Relay Data Link Connection Identifiers (DLCIs)
- 802.1d Prioritized Ethernet
- MPLS-based access
- Multilink Multiclass PPP
- QoS-enabled wireless (e.g., LMDS, MMDS)
- Cable modem
- QoS-enabled Digital Subscriber Line (DSL)

Different service models for QoS may be supported. Examples of PPVPN

QoS service models are:

- Managed access service : Provides QoS on the access connection between CE and the customer facing ports of the PE. No QoS support is required in the provider core network in this case.
- Edge-to-edge QoS : Provides QoS across the provider core, either between CE pairs or PE pairs, depending on the tunnel demarcation points. This scenario requires QoS support in the provider core network.

4.9. Service Level Agreement and Service Level Specification Monitoring and Reporting

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 interval for at least the SLS using the following Service Level Objective (SLO) parameters:

- QoS and traffic parameters for the Intserv flow or Diffserv class [Y.1541]
- Availability for the site, VPN, or access connection
- Duration of outage intervals per site, route or VPN
- Service activation interval (e.g., time to turn up a new site)
- Trouble report response time interval
- Time to repair interval
- Total traffic offered to the site, route or VPN
- Measure of non-conforming traffic for the site, route or VPN
- Delay and delay variation (jitter) bounds
- Packet ordering, at least when transporting L2 services sensitive to reordering (e.g., ATM).

The above list contains items from [Y.1241], as well as other items typically part of SLAs for currently deployed VPN services [FRF.13]. See [[RFC3198](#)] for generic definitions of SLS, SLA, and SLO.

The provider network management system shall measure, and report as

necessary, whether measured performance meets or fails to meet the above SLS objectives.

The service provider and the customer may negotiate a contractual arrangement that includes a Service Level Agreement (SLA) regarding compensation if the provider does not meet an SLS performance objective. Details of such compensation are outside the scope of this document.

4.10. Network Resource Partitioning and Sharing between VPNs

Network resources such as memory space, FIB table, bandwidth and CPU processing shall be shared between VPNs. Mechanisms should be provided to prevent any specific VPN from taking up available network resources and causing others to fail.

5. Provider requirements

This section describes operational requirements for a cost-effective, profitable VPN service offering.

5.1. Scalability

The scalability for VPN solutions has many aspects. The list below is intended to comprise of the aspects that PPVPN solutions should address. Clearly these aspects in absolute figures are very different for different types of VPNs - i.e., a point to point service has only two sites, while a VPLS or L3VPN may have a larger number of sites. It is also important to verify that PPVPN solutions not only scales on the high end, a VPN with three sites and three users should be as viable as a VPN with hundreds of sites and thousands of users.

Terminology:

Site: a geographical location with one or more users or one or more servers or a combination of servers and users.

User: the end user equipment (hosts), e.g., a workstation.

Note: Further discussion on [Section 5.1.1](#) and 5.1.2 is needed.

5.1.1. Service Provider Capacity Sizing Projections

A PPVPN solution should be scalable to support a very large number of VPNs per Service Provider network. The estimate is that a large service provider will require support for on the order of 10,000 VPNs within four years.

A PPVPN solution should be scalable to support of a wide range of number of site interfaces per VPN, depending on the size and/or structure of the customer organization. The number of site interfaces should range from a few site interfaces to over 50,000 site interfaces per VPN.

A PPVPN solution should be scalable to support of a wide range of number of routes per VPN. The number of routes per VPN may range from just a few to the number of routes exchanged between ISPs (on the order of 100,000). The high end number is especially true considering the fact that many large ISPs may provide VPN services to smaller ISPs or large corporations. Typically, the number of routes per VPN are twice the number of site interfaces or larger.

A PPVPN solution should support high values of the frequency of configuration setup and change, e.g., for real-time provisioning of an on-demand videoconferencing VPN or addition/deletion of sites.

Approaches should articulate scaling and performance limits for more complex deployment scenarios, such as inter-AS(S) VPNs and carriers' carrier. Approaches should also describe other dimensions of interest, such as capacity requirements or limits, number of interworking instances supported as well as any scalability implications on management systems.

A PPVPN solution should support a large number of customer interfaces on a single PE (for PE-based PPVPN) or CE (for CE-based PPVPN) with current Internet protocols.

5.1.2. VPN Scalability aspects

This section describes the metrics for scaling PPVPN solutions, points out some of the scaling differences between L2 and L3 VPNs. Further discussion on service provider sizing projections are in [Section 5.1.2](#).

5.1.2.1. Number of users per VPN

The number of users per VPN is the combination of servers and hosts connected to the VPN. It needs to scale from a handful to extremely high numbers.

Clearly there is a possible trade off between the number of users and the VPN technology chosen.

L3 VPNs must scale from 2 users per VPN to $O(10^5)$ users per VPN. L2 VPNs must scale from 2 users to a few hundred.

5.1.2.2. Number of users per site

The number of users per site follows the same logic as for users per VPN. Further, it must be possible to have single user sites connected to the same VPN as very large sites are connected to.

L3 VPNs must scale from 1 user per site to thousands of users per site. L2 VPNs must scale from 1 user to a few hundred per site.

5.1.2.3. Number of sites per VPN

The number of sites per VPN is clearly limited by the number of users for a L2 VPN. The largest number of sites in a L2VPN would be equal to the largest number of users, distributed one per site. For L3 VPNs number of sites per VPN should scale from 2 to very large numbers that are usually limited by router memory.

5.1.2.4. Number of PEs

The number of PEs that supports the same set of VPNs, i.e., the number of PEs that needs to directly exchange information on VPN de-multiplexing information is clearly a scaling factor. This number is driven by the type of VPN service, and also by whether the service is within a single AS/domain or involves a multi-SP or multi-AS network.

5.1.2.5. Number of sites per PE

The number of sites per PE needs to be discussed based on several different scenarios. On the one hand there is a limitation to the number of customer facing interfaces that the PE can support. On the other hand the access network may aggregate several sites connected on comparatively low bandwidth on to one single high bandwidth interface on the PE. The scaling point here is that the PE must be able to support a few or even a single site on the low end and $O(10^4)$ sites on the high end. Various PPVPN solutions may be evaluated based on this requirement.

5.1.2.6. Number of VPNs in the network

The number of sites in a network should not be a scaling issue. The number of VPNs should scale linearly with the size of the access network and with the number of PEs. As mentioned in [Section 5.1.1](#), the number of VPNs in the network should be $O(10^4)$.

5.1.2.7. Number of VPNs per PE

This is a function of number of sites per PE and number of VPNs per site. It is estimated that the number of VPNs needs to be at least one order of magnitude higher than the number of sites per PE.

5.1.2.8. Number of VPNs per site

On a customer's main site it is fully conceivable that the number of VPNs could be fairly large, both service diversification and separation of different work groups contributes to this. It is possible that one customer will run up to $O(100)$ VPNs.

5.1.2.9. Number of addresses per VPN

Since any VPN solution shall support private customer addresses, the number of addresses supported for a L3 VPN needs to scale from very few (for smaller customers) to very large numbers seen in typical SP backbones. The high end is especially true considering that many Tier 1 SPs may provide VPN services to Tier 2 SPs or to large corporations. For a L2 VPN this number would be on the order of

addresses supported in typical native Layer 2 backbones.

5.1.2.10. Number of addresses per PE

This is a function of the number of VPNs per site and the number of addresses per site.

5.1.3. Solution-Specific Metrics

Each PPVPN solution shall document its scalability characteristics in quantitative terms. Two examples are provided below as an illustration.

The following example applies to the number of tunnels necessary in various devices in the network. In a PE-based VPN, edge-to-edge tunnels (PE-to-PE) need to be established, while in a CE-based VPN, end-to-end tunnels between pairs of CEs are necessary. Therefore, fewer tunnels need to be maintained in a PE-based solution and scalability could be improved over that of CE-based VPNs. The other tradeoff is that in a PE-based solution, the CE is simple at the expense of complex PE devices, while on the other hand, in a CE-based solution, the PE devices remain simple while the CE devices are more complex.

A scalable PE-based solution should quantify the amount of state that a PE and P device must support. This should be stated in terms of the total number of VPNs and site interfaces supported by the service provider. Ideally, all VPN-specific state should be contained in the PE device for a PE-based VPN. Similarly, all VPN-specific state should be contained in the CE device for a CE-based VPN. In all cases, the backbone routers (P devices) shall not maintain VPN-specific state as far as possible.

5.2. Management

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, the service provider 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.

VPN devices should provide standards-based management interfaces wherever feasible.

Service Provider Network Management System (NMS) requirements mainly fall in the traditional fault, configuration, accounting, performance, and security (FCAPS) management categories. These requirements are available in detail in [Y.1311.1].

6. Engineering requirements

These requirements are driven by implementation characteristics that make service and provider requirements achievable.

6.1. Forwarding plane requirements

VPN solutions should not pre-suppose or preclude the use of IETF developed tunneling techniques such as IP-in-IP, L2TP, GRE, MPLS or IPsec. The separation of VPN solution and tunnels will facilitate adaptability with extensions to current tunneling techniques or development of new tunneling techniques. It should be noted that the choice of the tunneling techniques may impact the service capabilities of the VPN solution.

For Layer 2 VPNs, solutions should utilize the encapsulation techniques defined by PWE3, and should not impose any new requirements on these techniques.

PPVPN solutions must not impose any restrictions on the backbone traffic engineering and management techniques. Conversely, backbone engineering and management techniques must not affect the basic operation of a PPVPN, apart from impacting the SLA/SLS guarantees associated with the service.

By definition, VPN traffic should be segregated from each other, and from non-VPN traffic in the network. After all, VPNs are a means of dividing a physical network into several logical (virtual) networks. VPN traffic separation should be done in a scalable fashion. However, safeguards should be made available against misbehaving VPNs to not affect the network and other VPNs.

VPN solution should not impose any hard limit on the number of VPNs provided in the network.

6.2. Control plane requirements

The plug and play feature of a VPN solution with minimum configuration requirements is an important consideration. The VPN solutions should have mechanisms for protection against customer interface and/or routing instabilities so that they do not impact other customers' services.

A VPN should be provisioned with minimum number of steps. For instance, a VPN need not be configured in every PE. For this to be accomplished, an auto-configuration and an auto-discovery protocol, which should be common to all VPN solutions, should be defined. However, these mechanisms should not affect the cost, scalability or stability of a service by being overly complex, or by increasing layers in the protocol stack.

6.3. Control Plane Containment

The PPVPN control plane must include a mechanism through which the service provider can filter PPVPN related control plane information as it passes between Autonomous Systems. For example, if a service provider supports a PPVPN offering, but the service provider's neighbors do not participate in that offering, the service provider should not leak PPVPN control information into neighboring networks. Neighboring networks must be equipped with mechanisms that filter this information should the service provider leak it.

6.4. Requirements related to commonality of PPVPN mechanisms with each other and with generic Internet mechanisms

As far as possible, the mechanisms used to establish a VPN service should re-use well-known IETF protocols as far as possible, limiting the need to define new protocols from scratch. It should, however, be noted that the use of Internet mechanisms for the establishment and running of an Internet-based VPN service, shall not affect the stability, robustness, and scalability of the Internet or Internet services. In other words, these mechanisms should not conflict with the architectural principles of the Internet, nor should it put at risk the existing Internet systems. For example, IETF-developed routing protocols should be used for routing of L3 PPVPN traffic, without adding VPN-specific state to the Internet routers. Similarly, familiar L2 technologies should be used in VPNs offering L2 services, without imposing risks to the Internet routers. A solution must be implementable without requiring to add additional functionality to the devices in a network.

In addition to commonality with generic Internet mechanisms, infrastructure mechanisms used in different PPVPN solutions (both L2 and L3), e.g., discovery, signaling, routing and management, should be as common as possible.

6.5. Interoperability

Each technical solution is expected to be based on interoperable Internet standards.

Multi-vendor interoperability at network element, network and service levels among different implementations of the same technical solution should be ensured (that will likely rely on the completeness of the corresponding standard). This is a central requirement for SPs and customers.

The technical solution must be multi-vendor interoperable not only within the SP network infrastructure, but also with the customer's network equipment and services making usage of the PPVPN service.

Customer access connections to a PPVPN solution may be different at different sites (e.g., Frame Relay on one site and Ethernet on another)

Interconnection of a L2VPN to a L3VPN as if it were a customer site shall be supported.

Inter-domain interoperability - It should be possible to deploy a PPVPN solution across domains, Autonomous Systems, or the Internet.

7. Security Considerations

This document does not have any security considerations other than the security requirements described in [Section 4.5](#).

8. References

8.1. Normative References

- [RFC2026] Bradner, S., "The Internet Standards Process - Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [[RFC2119](#)] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997

8.2. Non-normative References

- [TERMINOLOGY] Andersson, L., Madsen, T., "Terminology for Provider Provisioned Virtual Private Networks", work in progress
- [PPVPN-FR] Callon, R., Suzuki, M., et al. "A Framework for Provider Provisioned Virtual Private Networks ", work in progress
- [PPVPN-VR] Ould-Brahim, H., Gleeson, B., et al. "Network based IP VPN Architecture using Virtual Routers", work in progress
- [L3REQTS] Carugi, M., McDysan, D. et al., "Service Requirements for Layer 3 Provider Provisioned Virtual Private Networks", work in progress
- [RFC2547bis] Rosen, E., Rekhter, Y. et al., "BGP/MPLS VPNs", work in progress.
- [Y.1241] "IP Transfer Capability for the support of IP based Services", Y.1241 ITU-T Draft Recommendation, March 2000
- [Y.1311.1] Carugi, M. (editor), "Network Based IP VPN over MPLS architecture", Y.1311.1 ITU-T Recommendation, May 2001 (<http://ppvnpn.francetelecom.com/ituRelated.html>)
- [Y.1311] Knightson, K. (editor), " Network based IP VPN Service - Generic Framework and Service Requirements ", Y.1311 ITU-T Draft Recommendation, May 2001 (<http://ppvnpn.francetelecom.com/ituRelated.html>)
- [[RFC 3198](#)] A. Westerinen et al, "Terminology for Policy-Based Management," November, 2001.
- [VPN SEC] J. De Clercq et al, "Considerations about possible security extensions to BGP/MPLS VPN," work in progress.
- [FRF.13] Frame Relay Forum, "Service Level Definitions Implementation Agreement," August, 1998.
- [Y.1541] "Network Performance Objectives for IP-based Services," Y.1541, ITU-T Recommendation.

9. Acknowledgements

This work was done in consultation with the entire design team for PPVPN requirements. A lot of the text was adapted from the Layer 3 requirements document produced by Layer 3 requirements design team. The authors would also like to acknowledge the constructive feedback from Alex Zinin.

10. Editor's Address

Ananth Nagarajan
Sprint
6220 Sprint Parkway
Overland Park, KS 66251
USA
E-mail: ananth.nagarajan@mail.sprint.com

11. Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

