                   Generic Requirements for Provider Provisioned
                       Virtual Private Networks (PPVPN)

Status of this Memo

Copyright Notice

Abstract

   This document describes 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.
   All PPVPN technologies are expected to meet the umbrella set of
   requirements described in this document.

Table of Contents

1. Introduction

   This document is an output of the design team formed to develop
   requirements for PPVPNs in the Provider Provisioned Virtual Private
   Networks (PPVPN) working group and provides requirements that are
   generic to both Layer 2 Virtual Private Networks (L2VPN) and Layer 3
   Virtual Private Networks (L3VPN). This document discusses generic
   PPVPN requirements categorized as service, provider and engineering
   requirements. These are independent of any particular type of PPVPN
   technology. In other words, all PPVPN technologies are expected to
   meet the umbrella set of requirements described in this document.
   PPVPNs may be constructed across single or multiple provider networks
   and/or Autonomous Systems (ASes). In most cases the generic
   requirements described in this document are independent of the
   deployment scenario. However, specific requirements that differ
   based on whether the PPVPN is deployed across single or multiple
   providers (and/or ASes) will be pointed out in the document.
   Specific requirements related to Layer 3 PPVPNs are described in
   [L3REQTS]. Similarly, requirements that are specific to layer 2
   PPVPNs are described in [L2REQTS].

   The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
   "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
   document are to be interpreted as described in [RFC2119].

1.1. Problem Statement

   Corporations and other organizations have become increasingly
   dependent on their networks for telecommunications and data
   communication. 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 Frame Relay (FR) or Asynchronous
   Transfer Mode (ATM) connections, and more recently 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. For example, for
service providers that need to support multiple types of VPN
services, it may be undesirable to require a completely different
solution approach for each type of VPN service.

1.2. Deployment Scenarios

   There are three different deployment scenarios that need to be
   considered for PPVPN services:

   1. Single-provider, single-AS: This is the least complex scenario,
      where the PPVPN service is offered across a single service
      provider network spanning a single Autonomous System.

   2. Single-provider, multi-AS: In this scenario, a single provider may
      have multiple Autonomous Systems (for e.g., a global Tier-1 ISP
      with different ASes depending on the global location, or an ISP

that has been created by mergers and acquisitions of multiple
networks). This scenario involves the constrained distribution of
routing information across multiple Autonomous Systems.

3. Multi-provider: This scenario is the most complex, wherein trust
   negotiations need to be made across multiple service provider
   backbones in order to meet the security and service level
   agreements for the PPVPN customer. This scenario can be
   generalized to cover the Internet, which comprises of multiple
   service provider networks. It should be noted that customers can
   construct their own VPNs across multiple providers. However such
   VPNs are not considered here as they would not be "Provider-
   provisioned".

A fourth scenario, "Carrier's carrier" VPN may also be considered.
In this scenario, a service provider (for example, a Tier 1 service
provider) provides VPN service to another service provider (for
example, a Tier 2 service provider), which in turn provides VPN
service on its VPN to its customers. In the example given above, the
Tier 2 provider's customers are contained within the Tier 2
provider's network, and the Tier 2 provider itself is a customer of
the Tier 1 provider's network. Thus, this scenario is not treated
separately in the document, because all of the single provider
requirements would apply equally to this case.

It is expected that many of the generic requirements described in
this document are independent of the three deployment scenarios
listed above. However, specific requirements that are indeed
dependent on the deployment scenario will be pointed out in this
document.

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 Service Provider focus group whose intentions were
   to present Service Provider view on the general requirements for
   PPVPN. A significant set of requirements were directly taken from
   previous work by the PPVPN WG to develop requirements for Layer 3
   PPVPN [L3REQTS]. The existing work in the L2 requirements area has
   also influenced the contents of this document [L2REQTS].

   Besides the editor, the following are the authors that contributed to
   this document:

      Loa Andersson (loa@pi.se)
      Ron Bonica (ronald.p.bonica@mci.com)
      Dave McDysan (dave.mcdysan@mci.com)
      Junichi Sumimoto (j.sumimoto@ntt.com)
      Muneyoshi Suzuki (suzuki.muneyoshi@lab.ntt.co.jp)
      David Meyer (dmm@1-4-5.net)
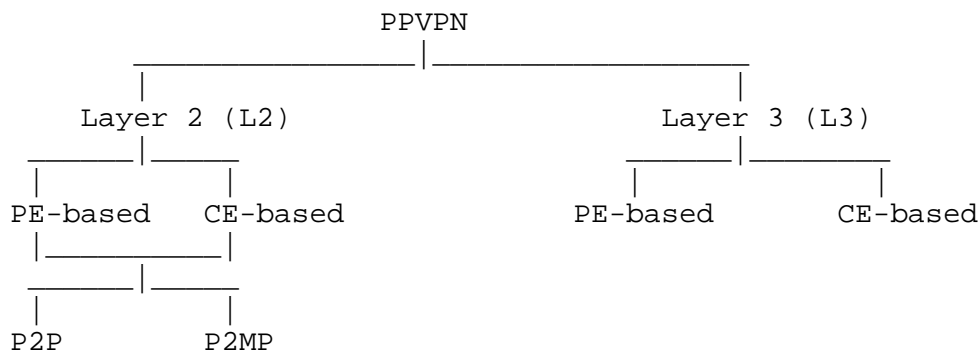   Marco Carugi (marco.carugi@nortelnetworks.com)

        Yetik Serbest (yetik_serbest@labs.sbc.com)
        Luyuan Fang (luyuanfang@att.com)
        Javier Achirica (achirica@telefonica.net)

3. Definitions and Taxonomy

   The terminology used in this document is defined in [TERMINOLOGY].
   In addition the following terminology is used:

   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.

```
                              PPVPN
                                |
               _____|_____
              |                                   |
        Layer 2 (L2)                        Layer 3 (L3)
        _____|_____                        _____|_____
       |             |                      |               |
    PE-based     CE-based               PE-based        CE-based
       |_____|
        _____|_____
       |             |
      P2P          P2MP
```

   The figure above presents a taxonomy of PPVPN technologies. PE-based
   and CE-based Layer 2 VPNs may also be further classified as point-to-
   point (P2P) or point-to-multipoint (P2MP). It is also the intention
   of the working group to have a limited number of solutions, and this
   goal must be kept in mind when proposing solutions that meet the
   requirements specified in this document. Definitions for CE-based
   and PE-based PPVPNs can be obtained from [L3FRAMEWORK]. Layer 2
   specific definitions can be obtained from [L2FRAMEWORK].

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. As mentioned before, each of these
   requirements apply equally across each of the three deployment
   scenarios unless stated otherwise.

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 PPVPNs) or
   PEs (for PE-based PPVPNs), or different POPs, or even different
   service providers.

2. Tunnel redundancy

   Redundant tunnels may be set up between the PEs (in a PE-based
   PPVPN) or the CEs (in a CE-based PPVPN) so that if one tunnel
   fails, VPN traffic can continue to flow across the other tunnel
   that has already been set-up in advance.

   Tunnel redundancy may be provided over and above physical
   diversity. For example, a single site may be connected to two CEs
   (for CE-based PPVPNs) or two PEs (for PE-based PPVPNs). Tunnels
   may be set up between each of the CEs (or PEs as the case may be)
   across different sites.

   Of course, redundancy means additional resources being used, and
   consequently, management of additional resources, which would
   impact the overall scaling of the service.

   It should be noted that it is difficult to guarantee high
   availability when the VPN service is across multiple providers,
   unless there is a negotiation between the different service
   providers to maintain the service level agreement for the VPN
   customer.

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.
For example, in the case of routing, route flapping or routing loops
MUST be avoided in order to ensure 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 or point-to-multipoint 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 applicable to L3 VPNs while broadcast
is more applicable to 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.

All PPVPN approaches SHALL support both IPv4 and IPv6 traffic.
Specific L2 traffic types (e.g., ATM, Frame Relay and Ethernet) SHALL
be supported via encapsulation in IP or MPLS tunnels in the case of
L2VPNs.

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. It should be noted that the disclosure of any
signaling/routing information across an extranet MUST be filtered per
the extranet agreement between the organizations participating in the
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. These security features SHOULD be implemented based on the
framework outlined in [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 attacks to the PE or CE CPUs, access
connection congestion, TCP SYN attacks and ping attacks.

Some security mechanisms (such as use of IPsec on a CE-to-CE basis)
may be equally useful regardless of the scope of the VPN. Other
mechanisms may be more applicable in some scopes than in others. For
example, in some cases of single-provider single-AS VPNs, the VPN
service may be isolated from some forms of attack by isolating the

infrastructure used for supporting VPNs from the infrastructure used
for other services. However, the requirements for security are
common regardless of the scope of the VPN service.

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. User data security using encryption is
especially desirable in the multi-provider scenario.

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, cryptographic mechanisms
or similar mechanisms to apply policy-based access control. Access
control MUST also be applicable between CE-CE, PE-PE and CE-PE. Such
access control mechanisms are desirable in the multi-provider
scenario.

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. This is particularly important for multi-
service provider deployment scenarios. However, this will also be
important in single-provider multi-AS scenarios.

4.6. Topology

   A VPN SHOULD support arbitrary, customer-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.

   Multiple VPNs per customer site SHOULD be supported without requiring
   additional hardware resources per VPN. This SHOULD also include a
   free mix of L2 and L3 VPNs.

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

4.7. Addressing

   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.

   Support for private addresses as described in [RFC1918], as well as
   overlapping customer addresses SHALL be supported. One or more VPNs
   for each customer can be built over the same infrastructure without
   requiring any of them to renumber. The solution MUST NOT use NAT on
   the customer traffic to achieve that goal. Interconnection of two
   networks with overlapping IP addresses is outside the scope of this
   document.

   A VPN service SHALL be capable of supporting non-IP customer
   addresses via encapsulation techniques, 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 technical approach for supporting VPNs 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. The
   extent to which any specific VPN service will support QoS is up to
   the service provider. In many cases single-provider single-AS VPNs
   will offer QoS guarantees. Support of QoS guarantees in the multi-
   service-provider case will require cooperation between the various
   service providers involved in offering the service.

It should be noted that QoS mechanisms in the multi-provider scenario
REQUIRES each of the participating providers to support the
mechanisms being used, and as such, this is difficult to achieve.

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. Hence in most cases PE-PE QoS is not a major issue. As far
as access QoS is concerned, there are two directions of QoS
management that may 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, as applicable to the
specific type of PPVPN (L2 or L3). However, 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. As mentioned above, this is difficult to achieve in a
  multi-provider VPN offering.

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.

   In many cases the guaranteed levels for Service Level Objective (SLO)
   parameters may depend upon the scope of the VPN. For example, one
   level of guarantee might be provided for service within a single AS.
   A different (generally less stringent) guarantee might be provided
   within multiple ASs within a single service provider. At the current
   time, in most cases specific guarantees are not offered for multi-
   provider VPNs, and if guarantees were offered they might be expected
   to be less stringent still.

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 and, where applicable, with
   non-VPN Internet traffic. Mechanisms SHOULD be provided to prevent
   any specific VPN from taking up available network resources and
   causing others to fail. SLAs to this effect SHOULD be provided to
   the customer.

   Similarly, resources used for control plane mechanisms are also
   shared. When the service provider's control plane is used to
   distribute VPN specific information and provide other control
   mechanisms for VPNs, there SHALL be mechanisms to ensure that control
   plane performance is not degraded below acceptable limits when
   scaling the VPN service, or during network events such as failure,
   routing instabilities etc. Since a service provider's network would
   also be used to provide Internet service, in addition to VPNs,
   mechanisms to ensure the stable operation of Internet services and
   other VPNs SHALL be made in order to avoid adverse effects of
   resource hogging by large VPN customers.

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, but also on the low end - i.e., a
   VPN with three sites and three users should be as viable as a VPN
   with hundreds of sites and thousands of users.

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 O(10^4) VPNs within four
   years.

   A PPVPN solution SHOULD be scalable to support 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
   (O(10^5)), with typical values being in the O(10^3) range. 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 is at least twice the number
   of site interfaces.

   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 single-provider multi-AS VPNs,
   multi-provider 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.
   It should be noted that the scaling numbers used in this document
   must be treated as typical examples as seen by the authors of this
   document. These numbers are only representative and different
   service providers may have different requirements for scaling.
   Further discussion on service provider sizing projections is in
   Section 5.1.1. Please note that the terms "user" and "site" are as
   defined in Section 3. It should also be noted that the numbers given

below would be different depending on whether the scope of the VPN is
single-provider single-AS, single-provider multi-AS, or multi-
provider. Clearly, the larger the scope, the larger the numbers that
may need to be supported. However, this also means more management
issues. The numbers below may be treated as representative of the
single-provider case.

5.1.2.1. 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 SHOULD scale from 1 user per site to $O(10^4)$ per site. L2
   VPNs SHOULD scale from 1 user to $O(10^3)$ per site for point-to-point
   VPNs and to $O(10^4)$ for point-to-multipoint VPNs.

5.1.2.2. Number of sites per VPN

   The number of sites per VPN clearly depends on the number of users
   per site. VPNs SHOULD scale from 2 to $O(10^3)$ sites per VPN. These
   numbers are usually limited by device memory.

5.1.2.3. Number of PEs and CEs

   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 in a PE-based
   VPN. Similarly, in a CE-based VPN, the number of CEs is 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. Typically, this number SHOULD be as
   low as possible in order to make the VPN cost effective and
   manageable.

5.1.2.4. 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 SHOULD be
   able to support a few or even a single site on the low end and
   $O(10^4)$ sites on the high end. This number is also limited by device
   memory. Implementations of PPVPN solutions may be evaluated based on
   this requirement, because it directly impacts cost and manageability
   of a VPN.

5.1.2.5. Number of VPNs in the network

   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). This
   requirement also effectively places a requirement on the number of
   tunnels that SHOULD be supported in the network. For a PE-based VPN,
   the number of tunnels is of the same order as the number of VPNs.
   For a CE-based VPN, the number of tunnels in the core network may be
   fewer, because of the possibility of tunnel aggregation or
   multiplexing across the core.

5.1.2.6. Number of VPNs per customer

   In some cases a service provider may support multiple VPNs for the
   same customer of that service provider. For example, this may occur
   due to differences in services offered per VPN (e.g., different QoS,
   security levels, or reachability) as well as due to the presence of
   multiple workgroups per customer. It is possible that one customer
   will run up to O(100) VPNs.

5.1.2.7. Number of addresses and address prefixes per VPN

   Since any VPN solution SHALL support private customer addresses, the
   number of addresses and address prefixes are important in evaluating
   the scaling requirements. The number of address prefixes used in
   routing protocols and in forwarding tables specific to the VPN needs
   to scale from very few (for smaller customers) to very large numbers
   seen in typical Service Provider 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.

5.1.3. Solution-Specific Metrics

   Each PPVPN solution SHALL document its scalability characteristics in
   quantitative terms. A VPN solution SHOULD quantify the amount of
   state that a PE and P device has to support. This SHOULD be stated
   in terms of the order of magnitude of the 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.

Another metric is that of complexity. In a PE-based solution the PE
is more complex in that it has to maintain tunnel-specific
information for each VPN, but the CE is simpler since it does not
need to support tunnels. On the other hand, in a CE-based solution,
the CE is more complex since it has to implement routing across a
number of tunnels to other CEs in the VPN, but the PE is simpler
since it has only one routing and forwarding instance. Thus, the
complexity of the PE or CE SHOULD be noted in terms of their
processing and management functions.

5.2. Management

A service provider MUST have a means to view the topology,
operational state, service 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.

In the multi-provider scenario, it is unlikely that participating
providers would provide each other a view to the network topology and
other parameters mentioned above. However, each provider MUST ensure
via management of their own networks that the overall VPN service
offered to the customers are properly managed. In general the
support of a single VPN spanning multiple service providers requires
close cooperation between the service providers. One aspect of this
cooperation involves agreement on what information about the VPN will
be visible across providers, and what network management protocols
will be used between providers.

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

5.2.1. Customer Management of a VPN

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

All aspects of management information about CE devices and customer
attributes of a PPVPN manageable by an SP SHOULD be capable of being
configured and maintained by the customer after being authenticated
and authorized.

A customer 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 as service is a "Dynamic Bandwidth management"
capability, that enables real-time response to customer requests for
changes of allocated bandwidth allocated to their VPN(s). A possible
outcome of giving customers such capabilities is Denial of Service
attacks on other VPN customers or Internet users. This possibility
is documented in the Security Considerations section.

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 and scaling
   capabilities of the VPN solution.

   It should also be noted that specific tunneling techniques may not be
   feasible depending on the deployment scenario. In particular, there
   is currently very little use of MPLS in the inter-provider scenario.
   Thus, native MPLS support may be needed between the service
   providers, or it would be necessary to run MPLS over IP or GRE. It
   should be noted that if MPLS is run over IP or GRE, some of the other
   capabilities of MPLS, such as Traffic Engineering, would be impacted.
   Also note that a service provider MAY optionally choose to use a
   different encapsulation for multi-AS VPNs than is used for single AS
   VPNs. Similarly, a group of service providers may choose to use a
   different encapsulation for multi-service provider VPNs than for VPNs
   within a single service provider.

   For Layer 2 VPNs, solutions SHOULD utilize the encapsulation
   techniques defined by the Pseudo-Wire Emulation Edge-to-Edge (PWE3)
   Working Group, 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 influencing the SLA/SLS guarantees
   associated with the service. The SP SHOULD, however, be REQUIRED to
   provide per-VPN management, tunnel maintenance and other maintenance
   required in order to meet the SLA/SLS.

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.

A 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 or impact general Internet traffic handling
in any way.

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 as common as possible to all VPN solutions, SHOULD be
defined. However, these mechanisms SHOULD NOT adversely affect the
cost, scalability or stability of a service by being overly complex,
or by increasing layers in the protocol stack.

Mechanisms to protect the SP network from effects of misconfiguration
of VPNs SHOULD be provided. This is especially of importance in the
multi-provider case, where misconfiguration could possibly impact
more than one network.

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. This is
important in the case of multi-provider VPNs as well as single-
provider multi-AS VPNs.

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, 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 core routers. Similarly,
   well-known L2 technologies SHOULD be used in VPNs offering L2
   services, without imposing risks to the Internet routers. A solution
   MUST be implementable without requiring additional functionality to
   the P devices in a network, and minimal functionality to the PE in a
   PE-based VPN and CE in a CE-based VPN.

   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 over an L3VPN as if it were a customer
   site SHALL be supported. However, interworking of Layer 2
   technologies is not required, and is outside the scope of the working
   group, and therefore, of this document.

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

7. Security Considerations

Security requirements for Provider Provisioned VPNs have been
described in Section 4.5. In addition, the following considerations
need to be kept in mind when a provider provisioned VPN service is
provided across a public network infrastructure that is also used to
provide Internet connectivity. In general, the security framework
described in [VPN-SEC] SHOULD be used as far as it is applicable to
the given type of PPVPN service.

The PE device has a lot of functionality required for the successful
operation of the VPN service. The PE device is frequently also part
of the backbone providing Internet services, and is therefore
susceptible to security and denial of service attacks. The PE
control plane CPU is vulnerable from this point of view, and it may
impact not only VPN services but also general Internet services if
not adequately protected. In addition to VPN configuration, if
mechanisms such as QoS are provisioned on the PE, it is possible for
attackers to recognize the highest priority traffic or customers and
launch directed attacks. Care SHOULD be taken to prevent such
attacks whenever any value added services such as QoS are offered.

When a service such as "Dynamic Bandwidth Management" as described in
Section 5.2.1 is provided, it allows customers to dynamically request
for changes to their bandwidth allocation. The provider MUST take
care to authenticate such requests and detect and prevent possible
Denial-of-Service attacks. These DoS attacks are possible when a
customer maliciously or accidentally may cause a change in bandwidth
allocation that may impact the bandwidth allocated to other VPN
customers or Internet users.

Different choices of VPN technology have different assurance levels
of the privacy of a customer's network. For example, CE-based
solutions may enjoy more privacy than PE-based VPNs by virtue of
tunnels extending from CE to CE, even if the tunnels are not
encrypted. In a PE-based VPN, a PE has many more sites than those
attached to a CE in a CE-based VPN. A large number of these sites
may use [RFC1918] addresses. Provisioning mistakes and PE software
bugs may make traffic more prone to being misdirected as opposed to a
CE-based VPN. Care MUST be taken to prevent misconfiguration in all
kinds of PPVPNs, but more care MUST be taken in the case of PE-based
VPNs, as this could impact other customers and Internet services.
Similarly, there SHOULD be mechanisms to prevent the flooding of

Internet routing tables whenever there is a misconfiguration or
failure of PPVPN control mechanisms that use Internet routing
protocols for relay of VPN-specific information.

Different deployment scenarios also dictate the level of security
that may be needed for a VPN. For example, it is easier to control
security in a single provider, single AS VPN and therefore, expensive
encryption techniques may not be used in this case, as long as VPN
traffic is isolated from the Internet. There is a reasonable amount
of control possible in the single provider, multi AS case, although
care SHOULD be taken to ensure the constrained distribution of VPN
route information across the ASes. Security is more of a challenge
in the multi-provider case, where it may be necessary to adopt
encryption techniques in order to provide the highest level of
security.

8. References

8.1. Normative References

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

8.2. Informative References

   [TERMINOLOGY] Andersson, L., Madsen, T., "Terminology for Provider
                  Provisioned Virtual Private Networks", Work in
                  Progress.

   [L3FRAMEWORK] Callon, R., Suzuki, M., et al. "A Framework for Layer 3
                  Provider Provisioned Virtual Private Networks", Work in
                  Progress, March 2003.

   [L2FRAMEWORK] Andersson, L., et al. "Framework for Layer 2 Virtual
                  Private Networks (L2VPNs)", Work in Progress, March
                  2004.

   [L3REQTS]   Carugi, M., McDysan, D. et al., "Service Requirements
                  for Layer 3 Provider Provisioned Virtual Private
                  Networks", Work in Progress, April 2003.

   [L2REQTS]   Augustyn, W., Serbest, Y., et al., "Service
                  Requirements for Layer 2 Provider Provisioned Virtual
                  Private Networks", Work in Progress, April 2003.

   [Y.1241]   "IP Transfer Capability for the support of IP based
                   Services", Y.1241 ITU-T Draft Recommendation, March
                   2000.

   [RFC1918]   Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot,
                   G. and E. Lear, "Address Allocation for Private
                   Internets", BCP 5, RFC 1918, February 1996.

   [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.

   [VPN-SEC]   Fang, L., et al., "Security Framework for Provider
                   Provisioned Virtual Private Networks", Work in
                   Progress, February 2004.

   [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

10. Editor's Address

   Ananth Nagarajan
   Juniper Networks

   EMail: ananth@juniper.net

11. Full Copyright Statement

Intellectual Property

Acknowledgement