Network Working Group
Request for Comments: 4111
Category: Informational

L. Fang, Ed.
AT&T Labs.

July 2005

Security Framework for Provider-Provisioned Virtual Private Networks (PPVPNs)

Status of This Memo

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

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document addresses security aspects pertaining to Provider-Provisioned Virtual Private Networks (PPVPNs). First, it describes the security threats in the context of PPVPNs and defensive techniques to combat those threats. It considers security issues deriving both from malicious behavior of anyone and from negligent or incorrect behavior of the providers. It also describes how these security attacks should be detected and reported. It then discusses possible user requirements for security of a PPVPN service. These user requirements translate into corresponding provider requirements. In addition, the provider may have additional requirements to make its network infrastructure secure to a level that can meet the PPVPN customer's expectations. Finally, this document defines a template that may be used to describe and analyze the security characteristics of a specific PPVPN technology.

Table of Contents

1. Introduction 2
2. Terminology 4
3. Security Reference Model 4
4. Security Threats 6
   4.1. Attacks on the Data Plane 7
   4.2. Attacks on the Control Plane 9
5. Defensive Techniques for PPVPN Service Providers 11
   5.1. Cryptographic Techniques 12
   5.2. Authentication 20
   5.3. Access Control Techniques 22
   5.4. Use of Isolated Infrastructure 27

5.5. Use of Aggregated Infrastructure 27
5.6. Service Provider Quality Control Processes 28
5.7. Deployment of Testable PPVPN Service 28
6. Monitoring, Detection, and Reporting of Security Attacks 28
7. User Security Requirements 29
7.1. Isolation 30
7.2. Protection 30
7.3. Confidentiality 31
7.4. CE Authentication 31
7.5. Integrity 31
7.6. Anti-replay 32
8. Provider Security Requirements 32
8.1. Protection within the Core Network 32
8.2. Protection on the User Access Link 34
8.3. General Requirements for PPVPN Providers 36
9. Security Evaluation of PPVPN Technologies 37
9.1. Evaluating the Template 37
9.2. Template 37
10. Security Considerations 40
11. Contributors 41
12. Acknowledgement 42
13. Normative References 42
14. Informative References 43

1. Introduction

Security is an integral aspect of Provider-Provisioned Virtual Private Network (PPVPN) services. The motivation and rationale for both Provider-Provisioned Layer-2 VPN and Provider-Provisioned Layer-3 VPN services are provided by [RFC4110] and [RFC4031]. These documents acknowledge that security is an important and integral aspect of PPVPN services, for both VPN customers and VPN service providers. Both will benefit from a PPVPN Security Framework document that lists the customer and provider security requirements related to PPVPN services, and that can be used to assess how much a particular technology protects against security threats and fulfills the security requirements.

First, we describe the security threats that are relevant in the context of PPVPNs, and the defensive techniques that can be used to combat those threats. We consider security issues deriving both from malicious or incorrect behavior of users and other parties and from negligent or incorrect behavior of the providers. An important part of security defense is the detection and report of a security attack,

which is also addressed in this document. Special considerations engendered by IP mobility within PPVPNs are not in the scope of this document.

Then, we discuss the possible user and provider security requirements for a PPVPN service. Users expectations must be met for the security characteristics of a VPN service. These user requirements translate into corresponding requirements for the providers offering the service. Furthermore, providers have security requirements to protect their network infrastructure, securing it to the level required to provide the PPVPN services in addition to other services.

Finally, we define a template that may be used to describe the security characteristics of a specific PPVPN technology in a manner consistent with the security framework described in this document. It is not within the scope of this document to analyze the security properties of specific technologies. Instead, our intention is to provide a common tool, in the form of a checklist, that may be used in other documents dedicated to an in-depth security analysis of individual PPVPN technologies to describe their security characteristics in a comprehensive and coherent way, thereby providing a common ground for comparison between different technologies.

It is important to clarify that this document is limited to describing users' and providers' security requirements that pertain to PPVPN services. It is not the intention to formulate precise "requirements" on each specific technology by defining the mechanisms and techniques that must be implemented to satisfy such users' and providers' requirements.

This document is organized as follows. Section 2 defines the terminology used in the document. Section 3 defines the security reference model for security in PPVPN networks. Section 4 describes the security threats that are specific of PPVPNs. Section 5 reviews defense techniques that may be used against those threats. Section 6 describes how attacks may be detected and reported. Section 7 discusses the user security requirements that apply to PPVPN services. Section 8 describes additional security requirements on the provider to guarantee the security of the network infrastructure providing PPVPN services. In Section 9, we provide a template that may be used to describe the security characteristics of specific PPVPN technologies. Finally, Section 10 discusses security considerations.

2. Terminology

This document uses PPVPN-specific terminology. Definitions and details specific to PPVPN terminology can be found in [RFC4026] and [RFC4110]. The most important definitions are repeated in this section; for other definitions, the reader is referred to [RFC4026] and [RFC4110].

CE: Customer Edge device, a router or a switch in the customer network interfacing with the service provider's network.

P: Provider Router. The Provider Router is a router in the service provider's core network that does not have interfaces directly toward the customer. A P router is used to interconnect the PE routers. A P router does not have to maintain VPN state and is thus VPN unaware.

PE: Provider Edge device, the equipment in the service provider's network that interfaces with the equipment in the customer's network.

PPVPN: Provider-Provisioned Virtual Private Network, a VPN that is configured and managed by the service provider (and thus not by the customer itself).

SP: Service Provider.

VPN: Virtual Private Network, which restricts communication between a set of sites using an IP backbone shared by traffic that is not going to or coming from those sites.

3. Security Reference Model

This section defines a reference model for security in PPVPN networks.

A PPVPN core network is the central network infrastructure (P and PE routers) over which PPVPN services are delivered. A PPVPN core network consists of one or more SP networks. All network elements in the core are under the operational control of one or more PPVPN service providers. Even if the PPVPN core is provided by several service providers, it appears to the PPVPN users as a single zone of trust. However, several service providers providing a common PPVPN core still have to secure themselves against the other providers. PPVPN services can also be delivered over the Internet, in which case the Internet forms a logical part of the PPVPN core.

A PPVPN user is a company, institution or residential client of the PPVPN service provider.

A PPVPN service is a private network service made available by a service provider to a PPVPN user. The service is implemented using virtual constructs built on a shared PPVPN core network. A PPVPN service interconnects sites of a PPVPN user.

Extranets are VPNs in which multiple sites are controlled by different (legal) entities. Extranets are another example of PPVPN deployment scenarios wherein restricted and controlled communication is allowed between trusted zones, often via well-defined transit points.

This document defines each PPVPN as a trusted zone and the PPVPN core as another trusted zone. A primary concern is security aspects that relate to breaches of security from the "outside" of a trusted zone to the "inside" of this zone. Figure 1 depicts the concept of trusted zones within the PPVPN framework.

```
+------------+                            +------------+
| PPVPN      +----------------------------+      PPVPN |
| user           PPVPN                               user |
| site       +---------------------XXX-----+       site |
+------------+  +------------------XXX--+  +------------+
                |   PPVPN core     | |  |
                +------------------| |--+
                                   | |
                                   | +------\
                                   +--------/  Internet
```

Figure 1: The PPVPN trusted zone model

In principle, the trusted zones should be separate. However, PPVPN core networks often offer Internet access, in which case a transit point (marked "XXX" in the figure) is defined.

The key requirement of a "virtual private" network (VPN) is that the security of the trusted zone of the VPN is not compromised by sharing the core infrastructure with other VPNs.

Security against threats that originate within the same trusted zone as their targets (for example, attacks from a user in a PPVPN to other users within the same PPVPN, or attacks entirely within the core network) is outside the scope of this document.

Also outside the scope are all aspects of network security that are independent of whether a network is a PPVPN network or a private

network. For example, attacks from the Internet to a web server inside a given PPVPN will not be considered here, unless the provisioning of the PPVPN network could make a difference to the security of this server.

4. Security Threats

This section discusses the various network security threats that may endanger PPVPNs. The discussion is limited to threats that are unique to PPVPNs, or that affect PPVPNs in unique ways. A successful attack on a particular PPVPN or on a service provider's PPVPN infrastructure may cause one or more of the following ill effects:

- observation, modification, or deletion of PPVPN user data,

- replay of PPVPN user data,

- injection of non-authentic data into a PPVPN,

- traffic pattern analysis on PPVPN traffic,

- disruption of PPVPN connectivity, or

- degradation of PPVPN service quality.

It is useful to consider that threats to a PPVPN, whether malicious or accidental, may come from different categories of sources. For example they may come from:

- users of other PPVPNs provided by the same PPVPN service provider,

- the PPVPN service provider or persons working for it,

- other persons who obtain physical access to a service provider site,

- other persons who use social engineering methods to influence behavior of service provider personnel,

- users of the PPVPN itself, i.e., intra-VPN threats (such threats are beyond the scope of this document), or

- others, i.e., attackers from the Internet at large.

In the case of PPVPNs, some parties may be in more advantageous positions that enable them to launch types of attacks not available to others. For example, users of different PPVPNs provided by the

same service provider may be able to launch attacks that those who are completely outside the network cannot.

Given that security is generally a compromise between expense and risk, it is also useful to consider the likelihood of different attacks. There is at least a perceived difference in the likelihood of most types of attacks being successfully mounted in different environments, such as

- in a PPVPN contained within one service provider’s network, or

- in a PPVPN transiting the public Internet.

Most types of attacks become easier to mount, and hence more likely, as the shared infrastructure that provides VPN service expands from a single service provider to multiple cooperating providers, and then to the global Internet. Attacks that may not be sufficiently likely to warrant concern in a closely controlled environment often merit defensive measures in broader, more open environments.

The following sections discuss specific types of exploits that threaten PPVPNs.

4.1. Attacks on the Data Plane

This category encompasses attacks on the PPVPN user’s data, as viewed by the service provider. Note that from the PPVPN user’s point of view, some of this might be control plane traffic, e.g., routing protocols running from PPVPN user site to PPVPN user site via an L2 PPVPN.

4.1.1. Unauthorized Observation of Data Traffic

This refers to "sniffing" VPN packets and examining their contents. This can result in exposure of confidential information. It can also be a first step in other attacks (described below) in which the recorded data is modified and re-inserted, or re-inserted unchanged.

4.1.2. Modification of Data Traffic

This refers to modifying the contents of packets as they traverse the VPN.

4.1.3. Insertion of Non-authentic Data Traffic: Spoofing and Replay

This refers to the insertion into the VPN (or "spoofing") of packets that do not belong there, with the objective of having them accepted as legitimate by the recipient. Also included in this category is

the insertion of copies of once-legitimate packets that have been recorded and replayed.

4.1.4. Unauthorized Deletion of Data Traffic

This refers to causing packets to be discarded as they traverse the VPN. This is a specific type of Denial-of-Service attack.

4.1.5. Unauthorized Traffic Pattern Analysis

This refers to "sniffing" VPN packets and examining aspects or meta-aspects of them that may be visible even when the packets themselves are encrypted. An attacker might gain useful information based on the amount and timing of traffic, packet sizes, source and destination addresses, etc. For most PPVPN users, this type of attack is generally considered significantly less of a concern than are the other types discussed in this section.

4.1.6. Denial-of-Service Attacks on the VPN

Denial-of-Service (DoS) attacks are those in which an attacker attempts to disrupt or prevent the use of a service by its legitimate users. Taking network devices out of service, modifying their configuration, or overwhelming them with requests for service are several of the possible avenues for DoS attack.

Overwhelming the network with requests for service, otherwise known as a "resource exhaustion" DoS attack, may target any resource in the network, e.g., link bandwidth, packet forwarding capacity, session capacity for various protocols, and CPU power.

DoS attacks of the resource exhaustion type can be mounted against the data plane of a particular PPVPN by attempting to insert (spoof) an overwhelming quantity of non-authentic data into the VPN from outside of that VPN. Potential results might be to exhaust the bandwidth available to that VPN or to overwhelm the cryptographic authentication mechanisms of the VPN.

Data plane resource exhaustion attacks can also be mounted by overwhelming the service provider’s general (VPN-independent) infrastructure with traffic. These attacks on the general infrastructure are not usually a PPVPN-specific issue, unless the attack is mounted by another PPVPN user from a privileged position. For example, a PPVPN user might be able to monopolize network data plane resources and thus to disrupt other PPVPNs.)

4.2. Attacks on the Control Plane

This category encompasses attacks on the control structures operated by the PPVPN service provider.

4.2.1. Denial-of-Service Attacks on Network Infrastructure

Control plane DoS attacks can be mounted specifically against the mechanisms that the service provider uses to provide PPVPNs (e.g., IPsec, MPLS) or against the general infrastructure of the service provider (e.g., P routers or shared aspects of PE routers.) Attacks against the general infrastructure are within the scope of this document only if the attack happens in relation to the VPN service; otherwise, they are not a PPVPN-specific issue.

Of special concern for PPVPNs is denial of service to one PPVPN user caused by the activities of another. This can occur, for example, if one PPVPN user’s activities are allowed to consume excessive network resources of any sort that are also needed to serve other PPVPN users.

The attacks described in the following sections may each have denial of service as one of their effects. Other DoS attacks are also possible.

4.2.2. Attacks on Service Provider Equipment via Management Interfaces

This includes unauthorized access to service provider infrastructure equipment, in order, for example, to reconfigure the equipment or to extract information (statistics, topology, etc.) about one or more PPVPNs.

This can be accomplished through malicious entrance of the systems, or as an inadvertent consequence of inadequate inter-VPN isolation in a PPVPN user self-management interface. (The former is not necessarily a PPVPN-specific issue.)

4.2.3. Social Engineering Attacks on Service Provider Infrastructure

Attacks in which the service provider network is reconfigured or damaged, or in which confidential information is improperly disclosed, may be mounted through manipulation of service provider personnel. These types of attacks are PPVPN-specific if they affect PPVPN-serving mechanisms. It may be observed that the organizational split (customer, service provider) that is inherent in PPVPNs may make it easier to mount such attacks against provider-provisioned

VPNs than against VPNs that are self-provisioned by the customer at the IP layer.

4.2.4. Cross-Connection of Traffic between PPVPNs

This refers to events where expected isolation between separate PPVPNs is breached. This includes cases such as:

- a site being connected into the "wrong" VPN,

- two or more VPNs being improperly merged,

- a point-to-point VPN connecting the wrong two points, or

- any packet or frame being improperly delivered outside the VPN it is sent in.

Misconnection or cross-connection of VPNs may be caused by service provider or equipment vendor error, or by the malicious action of an attacker. The breach may be physical (e.g., PE-CE links misconnected) or logical (improper device configuration).

Anecdotal evidence suggests that the cross-connection threat is one of the largest security concerns of PPVPN users (or would-be users).

4.2.5. Attacks against PPVPN Routing Protocols

This encompasses attacks against routing protocols that are run by the service provider and that directly support the PPVPN service. In layer 3 VPNs this, typically relates to membership discovery or to the distribution of per-VPN routes. In layer 2 VPNs, this typically relates to membership and endpoint discovery. Attacks against the use of routing protocols for the distribution of backbone (non-VPN) routes are beyond the scope of this document. Specific attacks against popular routing protocols have been widely studied and are described in [RFC3889].

4.2.6. Attacks on Route Separation

"Route separation" refers here to keeping the per-VPN topology and reachability information for each PPVPN separate from, and unavailable to, any other PPVPN (except as specifically intended by the service provider). This concept is only a distinct security concern for layer-3 VPN types for which the service provider is involved with the routing within the VPN (i.e., VR, BGP-MPLS, routed version of IPsec). A breach in the route separation can reveal topology and addressing information about a PPVPN. It can also cause

black hole routing or unauthorized data plane cross-connection between PPVPNs.

4.2.7. Attacks on Address Space Separation

In layer-3 VPNs, the IP address spaces of different VPNs have to be kept separate. In layer-2 VPNs, the MAC address and VLAN spaces of different VPNs have to be kept separate. A control plane breach in this addressing separation may result in unauthorized data plane cross-connection between VPNs.

4.2.8. Other Attacks on PPVPN Control Traffic

Besides routing and management protocols (covered separately in the previous sections), a number of other control protocols may be directly involved in delivering the PPVPN service (e.g., for membership discovery and tunnel establishment in various PPVPN approaches). These include but may not be limited to:

- MPLS signaling (LDP, RSVP-TE),
- IPsec signaling (IKE) ,
- L2TP,
- BGP-based membership discovery, and
- Database-based membership discovery (e.g., RADIUS-based).

Attacks might subvert or disrupt the activities of these protocols, for example, via impersonation or DoS attacks.

5. Defensive Techniques for PPVPN Service Providers

The defensive techniques discussed in this document are intended to describe methods by which some security threats can be addressed. They are not intended as requirements for all PPVPN implementations. The PPVPN provider should determine the applicability of these techniques to the provider’s specific service offerings, and the PPVPN user may wish to assess the value of these techniques in regard to the user’s VPN requirements.

The techniques discussed here include encryption, authentication, filtering, firewalls, access control, isolation, aggregation, and other techniques.

Nothing is ever 100% secure. Defense therefore protects against those attacks that are most likely to occur or that could have the most dire consequences. Absolute protection against these attacks is seldom achievable; more often it is sufficient to make the cost of a successful attack greater than what the adversary would be willing to expend.

Successful defense against an attack does not necessarily mean that the attack must be prevented from happening or from reaching its target. In many cases, the network can instead be designed to withstand the attack. For example, the introduction of non-authentic packets could be defended against by preventing their introduction in the first place, or by making it possible to identify and eliminate them before delivery to the PPVPN user’s system. The latter is frequently a much easier task.

5.1. Cryptographic Techniques

PPVPN defenses against a wide variety of attacks can be enhanced by the proper application of cryptographic techniques. These are the same cryptographic techniques that are applicable to general network communications. In general, these techniques can provide confidentiality (encryption) of communication between devices, authentication of the identities of the devices, and detection of a change of the protected data during transit.

Privacy is a key part (the middle name!) of any Virtual Private Network. In a PPVPN, privacy can be provided by two mechanisms: traffic separation and encryption. This section focuses on encryption; traffic separation is addressed separately.

Several aspects of authentication are addressed in some detail in a separate "Authentication" section.

Encryption adds complexity, and thus it may not be a standard offering within every PPVPN service. There are a few reasons for this. Encryption adds an additional computational burden to the devices performing encryption and decryption. This may reduce the number of user VPN connections that can be handled on a device or otherwise reduce the capacity of the device, potentially driving up the provider’s costs. Typically, configuring encryption services on devices adds to the complexity of the device configuration and adds incremental labor cost. Encrypting packets typically increases packet lengths, thereby increasing the network traffic load and the likelihood of packet fragmentation, with its increased overhead. (Packet length increase can often be mitigated to some extent by data compression techniques, but with additional computational burden.) Finally, some PPVPN providers may employ enough other defensive techniques, such as physical isolation or filtering/firewall techniques, that they may not perceive additional benefit from encryption techniques.

The trust model among the PPVPN user, the PPVPN provider, and other parts of the network is a key element in determining the applicability of encryption for any specific PPVPN implementation.

In particular, it determines where encryption should be applied, as follows.

- If the data path between the user's site and the provider's PE is not trusted, then encryption may be used on the PE-CE link.

- If some part of the backbone network is not trusted, particularly in implementations where traffic may travel across the Internet or multiple provider networks, then the PE-PE traffic may be encrypted.

- If the PPVPN user does not trust any zone outside of its premises, it may require end-to-end or CE-CE encryption service. This service fits within the scope of this PPVPN security framework when the CE is provisioned by the PPVPN provider.

- If the PPVPN user requires remote access to a PPVPN from a system that is not at a PPVPN customer location (for example, access by a traveler), there may be a requirement for encrypting the traffic between that system and an access point on the PPVPN or at a customer site. If the PPVPN provider provides the access point, then the customer must cooperate with the provider to handle the access control services for the remote users. These access control services are usually implemented by using encryption, as well.

Although CE-CE encryption provides confidentiality against third-party interception, if the PPVPN provider has complete management control over the CE (encryption) devices, then it may be possible for the provider to gain access to the user's VPN traffic or internal network. Encryption devices can potentially be configured to use null encryption, to bypass encryption processing altogether, or to provide some means of sniffing or diverting unencrypted traffic. Thus, a PPVPN implementation using CE-CE encryption has to consider the trust relationship between the PPVPN user and provider. PPVPN users and providers may wish to negotiate a service level agreement (SLA) for CE-CE encryption that will provide an acceptable demarcation of responsibilities for management of encryption on the CE devices.

The demarcation may also be affected by the capabilities of the CE devices. For example, the CE might support some partitioning of management or a configuration lock-down ability, or it might allow both parties to verify the configuration. In general, if the managed CE-CE model is used, the PPVPN user has to have a fairly high level of trust that the PPVPN provider will properly provision and manage the CE devices.

5.1.1. IPsec in PPVPNs

IPsec [RFC2401] [RFC2402] [RFC2406] [RFC2407] [RFC2411] is the security protocol of choice for encryption at the IP layer (Layer 3), as discussed in [RFC3631]. IPsec provides robust security for IP traffic between pairs of devices. Non-IP traffic must be converted to IP packets, or it cannot be transported over IPsec. Encapsulation is a common conversion method.

In the PPVPN model, IPsec can be employed to protect IP traffic between PEs, between a PE and a CE, or from CE to CE. CE-to-CE IPsec may be employed in either a provider-provisioned or a user-provisioned model. The user-provisioned CE-CE IPsec model is outside the scope of this document and outside the scope of the PPVPN Working Group. Likewise, data encryption that is performed within the user's site is outside the scope of this document, as it is simply handled as user data by the PPVPN. IPsec can also be used to protect IP traffic between a remote user and the PPVPN.

IPsec does not itself specify an encryption algorithm. It can use a variety of encryption algorithms with various key lengths, such as AES encryption. There are trade-offs between key length, computational burden, and the level of security of the encryption. A full discussion of these trade-offs is beyond the scope of this document. In order to assess the level of security offered by a particular IPsec-based PPVPN service, some PPVPN users may wish to know the specific encryption algorithm and effective key length used by the PPVPN provider. However, in practice, any currently recommended IPsec encryption offers enough security to substantially reduce the likelihood of being directly targeted by an attacker. Other, weaker, links in the chain of security are likely to be attacked first. PPVPN users may wish to use a Service Level Agreement (SLA) specifying the service provider's responsibility for ensuring data confidentiality rather than to analyze the specific encryption techniques used in the PPVPN service.

For many of the PPVPN provider's network control messages and some PPVPN user requirements, cryptographic authentication of messages without encryption of the contents of the message may provide acceptable security. With IPsec, authentication of messages is provided by the Authentication Header (AH) or by the Encapsulating Security Protocol (ESP) with authentication only. Where control messages require authentication but do not use IPsec, other cryptographic authentication methods are available. Message authentication methods currently considered to be secure are based on hashed message authentication codes (HMAC) [RFC2104] implemented with a secure hash algorithm such as Secure Hash Algorithm 1 (SHA-1) [RFC3174].

One recommended mechanism for providing a combination confidentiality, data origin authentication, and connectionless integrity is the use of AES in Cipher Block Chaining (CBC) Mode, with an explicit Initialization Vector (IV) [RFC3602], as the IPsec ESP.

PPVPNs that provide differentiated services based on traffic type may encounter some conflicts with IPsec encryption of traffic. As encryption hides the content of the packets, it may not be possible to differentiate the encrypted traffic in the same manner as unencrypted traffic. Although DiffServ markings are copied to the IPsec header and can provide some differentiation, not all traffic types can be accommodated by this mechanism.

5.1.2. Encryption for Device Configuration and Management

For configuration and management of PPVPN devices, encryption and authentication of the management connection at a level comparable to that provided by IPsec is desirable.

Several methods of transporting PPVPN device management traffic offer security and confidentiality.

- Secure Shell (SSH) offers protection for TELNET [STD8] or terminal-like connections to allow device configuration.

- SNMP v3 [STD62] provides encrypted and authenticated protection for SNMP-managed devices.

- Transport Layer Security (TLS) [RFC2246] and the closely-related Secure Sockets Layer (SSL) are widely used for securing HTTP-based communication, and thus can provide support for most XML- and SOAP-based device management approaches.

- As of 2004, extensive work is proceeding in several organizations (OASIS, W3C, WS-I, and others) on securing device management traffic within a "Web Services" framework. This work uses a wide variety of security models and supports multiple security token formats, multiple trust domains, multiple signature formats, and multiple encryption technologies.

- IPsec provides the services with security and confidentiality at the network layer. With regard to device management, its current use is primarily focused on in-band management of user-managed IPsec gateway devices.

5.1.3. Cryptographic Techniques in Layer-2 PPVPNs

Layer-2 PPVPNs will generally not be able to use IPsec to provide encryption throughout the entire network. They may be able to use IPsec for PE-PE traffic where it is encapsulated in IP packets, but IPsec will generally not be applicable for CE-PE traffic in Layer-2 PPVPNs.

Encryption techniques for Layer-2 links are widely available but are not within the scope of this document or IETF documents in general. Layer-2 encryption could be applied to the links from CE to PE, or it could be applied from CE to CE, as long as the encrypted Layer-2 packets can be handled properly by the intervening PE devices. In addition, the upper-layer traffic transported by the Layer-2 VPN can be encrypted by the user. In this case, confidentiality will be maintained; however, this is transparent to the PPVPN provider and is outside the scope of this document.

5.1.4. End-to-End vs. Hop-by-Hop Encryption Tradeoffs in PPVPNs

In PPVPNs, encryption could potentially be applied to the VPN traffic at several different places. This section discusses some of the tradeoffs in implementing encryption in several different connection topologies among different devices within a PPVPN.

Encryption typically involves a pair of devices that encrypt the traffic passing between them. The devices may be directly connected (over a single "hop"), or there may be intervening devices that transport the encrypted traffic between the pair of devices. The extreme cases involve hop-by-hop encryption between every adjacent pair of devices along a given path or "end-to-end" encryption only between the end devices along a given path. To keep this discussion within the scope of PPVPNs, we consider the "end to end" case to be CE to CE rather than fully end to end.

Figure 2 depicts a simplified PPVPN topology, showing the Customer Edge (CE) devices, the Provider Edge (PE) devices, and a variable number (three are shown) of Provider core (P) devices that might be present along the path between two sites in a single VPN, operated by a single service provider (SP).

```
Site_1---CE---PE---P---P---P---PE---CE---Site_2
```

Figure 2: Simplified PPVPN topology

Within this simplified topology and assuming that P devices are not to be involved with encryption, there are four basic feasible configurations for implementing encryption on connections among the devices:

1) Site-to-site (CE-to-CE): Encryption can be configured between the two CE devices, so that traffic will be encrypted throughout the SP's network.

2) Provider edge-to-edge (PE-to-PE): Encryption can be configured between the two PE devices. Unencrypted traffic is received at one PE from the customer's CE; then it is encrypted for transmission through the SP's network to the other PE, where it is decrypted and sent to the other CE.

3) Access link (CE-to-PE): Encryption can be configured between the CE and PE, on each side (or on only one side).

4) Configurations 2) and 3) can be combined, with encryption running from CE to PE, then from PE to PE, and then from PE to CE.

Among the four feasible configurations, key tradeoffs in considering encryption include the following:

- Vulnerability to link eavesdropping: Assuming that an attacker can observe the data in transit on the links, would it be protected by encryption?

- Vulnerability to device compromise: Assuming an attacker can get access to a device (or freely alter its configuration), would the data be protected?

- Complexity of device configuration and management: Given Nce, the number of sites per VPN customer, and Npe, the number of PEs participating in a given VPN, how many device configurations have to be created or maintained and how do those configurations scale?

- Processing load on devices: How many encryption or decryption operations must be done, given P packets? This influences considerations of device capacity and perhaps end-to-end delay.

- Ability of SP to provide enhanced services (QoS, firewall, intrusion detection, etc.): Can the SP inspect the data in order to provide these services?

These tradeoffs are discussed below for each configuration.

1) Site-to-site (CE-to-CE) Configurations

- o Link eavesdropping: Protected on all links.
- o Device compromise: Vulnerable to CE compromise.
- o Complexity: Single administration, responsible for one device per site (Nce devices), but overall configuration per VPN scales as Nce**2.
- o Processing load: on each of two CEs, each packet is either encrypted or decrypted (2P).
- o Enhanced services: Severely limited; typically only DiffServ markings are visible to SP, allowing some QoS services.

2) Provider edge-to-edge (PE-to-PE) Configurations

- o Link eavesdropping: Vulnerable on CE-PE links; protected on SP's network links.
- o Device compromise: Vulnerable to CE or PE compromise.
- o Complexity: Single administration; Npe devices to configure. (Multiple sites may share a PE device, so Npe is typically much less than Nce.) Scalability of the overall configuration depends on the PPVPN type: If the encryption is separate per VPN context, it scales as Npe**2 per customer VPN. If the encryption is per PE, it scales as Npe**2 for all customer VPNs combined.
- o Processing load: On each of two PEs, each packet is either encrypted or decrypted (2P).
- o Enhanced services: Full; SP can apply any enhancements based on detailed view of traffic.

3) Access link (CE-to-PE) Configuration

- o Link eavesdropping: Protected on CE-PE link; vulnerable on SP's network links.
- o Device compromise: Vulnerable to CE or PE compromise.
- o Complexity: Two administrations (customer and SP) with device configuration on each side (Nce + Npe devices to configure), but as there is no mesh, the overall configuration scales as Nce.

- o Processing load: On each of two CEs, each packet is either encrypted or decrypted. On each of two PEs, each packet is either encrypted or decrypted (4P).

- o Enhanced services: Full; SP can apply any enhancements based on detailed view of traffic.

4) Combined Access link and PE-to-PE (essentially hop-by-hop).

- o Link eavesdropping: Protected on all links.

- o Device compromise: Vulnerable to CE or PE compromise.

- o Complexity: Two administrations (customer and SP), with device configuration on each side (Nce + Npe devices to configure). Scalability of the overall configuration depends on the PPVPN type. If the encryption is separate per VPN context, it scales as Npe**2 per customer VPN. If the encryption is per-PE, it scales as Npe**2 for all customer VPNs combined.

- o Processing load: On each of two CEs, each packet is either encrypted or decrypted. On each of two PEs, each packet is both encrypted and decrypted (6P).

- o Enhanced services: Full; SP can apply any enhancements based on detailed view of traffic.

Given the tradeoffs discussed above, a few conclusions can be reached.

- Configurations 2 and 3, which are subsets of 4, may be appropriate alternatives to 4 under certain threat models. The remainder of these conclusions compare 1 (CE-to-CE) with 4 (combined access links and PE-to-PE).

- If protection from link eavesdropping is most important, then configurations 1 and 4 are equivalent.

- If protection from device compromise is most important and the threat is to the CE devices, both cases are equivalent; if the threat is to the PE devices, configuration 1 is best.

- If reducing complexity is most important and the size of the network is very small, configuration 1 is the best. Otherwise, the comparison between options 1 and 4 is relatively complex , based on a number of issues such as, how close the CE to CE communication is to a full mesh, and what tools are used for key management. Option 1 requires configuring keys for each CE-CE

pair that is communicating directly. Option 4 requires configuring keys on both CE and PE devices but may offer benefit from the fact that the number of PEs is generally much smaller than the number of CEs.

Also, under some PPVPN approaches, the scaling of 4 is further improved by sharing the same PE-PE mesh across all VPN contexts. The scaling characteristics of 4 may be increased or decreased in any given situation if the CE devices are simpler to configure than the PE devices, or vice versa. Furthermore, with option 4, the impact of operational error may be significantly increased.

- If the overall processing load is a key factor, then 1 is best.

- If the availability of enhanced services support from the SP is most important, then 4 is best.

As a quick overall conclusion, CE-to-CE encryption provides greater protection against device compromise, but it comes at the cost of enhanced services and with additional operational complexity due to the Order(n**2) scaling of the mesh.

This analysis of site-to-site vs. hop-by-hop encryption tradeoffs does not explicitly include cases where multiple providers cooperate to provide a PPVPN service, public Internet VPN connectivity, or remote access VPN service, but many of the tradeoffs will be similar.

5.2. Authentication

In order to prevent security issues from some denial-of-service attacks or from malicious misconfiguration, it is critical that devices in the PPVPN should only accept connections or control messages from valid sources. Authentication refers to methods for ensuring that message sources are properly identified by the PPVPN devices with which they communicate. This section focuses on identifying the scenarios in which sender authentication is required, and it recommends authentication mechanisms for these scenarios.

Cryptographic techniques (authentication and encryption) do not protect against some types of denial-of-service attacks, specifically, resource exhaustion attacks based on CPU or bandwidth exhaustion. In fact, the processing required to decrypt or check authentication may in some cases increase the effect of these resource exhaustion attacks. Cryptographic techniques may, however, be useful against resource exhaustion attacks based on exhaustion of state information (e.g., TCP SYN attacks).

5.2.1. VPN Member Authentication

This category includes techniques for the CEs to verify that they are connected to the expected VPN. It includes techniques for CE-PE authentication, to verify that each specific CE and PE is actually communicating with its expected peer.

5.2.2. Management System Authentication

Management system authentication includes the authentication of a PE to a centrally-managed directory server when directory-based "auto-discovery" is used. It also includes authentication of a CE to its PPVPN configuration server when a configuration server system is used.

5.2.3. Peer-to-Peer Authentication

Peer-to-peer authentication includes peer authentication for network control protocols (e.g., LDP, BGP), and other peer authentication (i.e., authentication of one IPsec security gateway by another).

5.2.4. Authenticating Remote Access VPN Members

This section describes methods for authentication of remote access users connecting to a VPN.

Effective authentication of individual connections is a key requirement for enabling remote access to a PPVPN from an arbitrary Internet address (for instance, by a traveler).

There are several widely used standards-based protocols to support remote access authentication. These include RADIUS [RFC2865] and DIAMETER [RFC3588]. Digital certificate systems also provide authentication. In addition, there has been extensive development and deployment of mechanisms for securely transporting individual remote access connections within tunneling protocols, including L2TP [RFC2661] and IPsec.

Remote access involves connection to a gateway device, which provides access to the PPVPN. The gateway device may be managed by the user at a user site, or by the PPVPN provider at any of several possible locations in the network. The user-managed case is of limited interest within the PPVPN security framework, and it is not considered at this time.

When a PPVPN provider manages authentication at the remote access gateway, this implies that authentication databases, which are usually extremely confidential user-managed systems, will have to be

referenced in a secure manner by the PPVPN provider. This can be accomplished through proxy authentication services, which accept an encrypted authentication credential from the remote access user, pass it to the PPVPN user’s authentication system, and receive a yes/no response as to whether the user has been authenticated. Thus, the PPVPN provider does not have access to the actual authentication database, but it can use it on behalf of the PPVPN user to provide remote access authentication.

Specific cryptographic techniques for handling authentication are described in the following sections.

5.2.5. Cryptographic Techniques for Authenticating Identity

Cryptographic techniques offer several mechanisms for authenticating the identity of devices or individuals. These include the use of shared secret keys, one-time keys generated by accessory devices or software, user-ID and password pairs, and a range of public-private key systems. Another approach is to use a hierarchical Certificate Authority system to provide digital certificates.

This section describes or provides references to the specific cryptographic approaches for authenticating identity. These approaches provide secure mechanisms for most of the authentication scenarios required in operating a PPVPN.

5.3. Access Control Techniques

Access control techniques include packet-by-packet or packet flow - by - packet flow access control by means of filters and firewalls, as well as by means of admitting a "session" for a control/signaling/management protocol that is being used to implement PPVPNs. Enforcement of access control by isolated infrastructure addresses is discussed elsewhere in this document.

We distinguish between filtering and firewalls primarily by the direction of traffic flow. We define filtering as being applicable to unidirectional traffic, whereas a firewall can analyze and control both sides of a conversation.

There are two significant corollaries of this definition:

- Routing or traffic flow symmetry: A firewall typically requires routing symmetry, which is usually enforced by locating a firewall where the network topology assures that both sides of a conversation will pass through the firewall. A filter can then operate upon traffic flowing in one direction without considering traffic in the reverse direction.

- Statefulness: Because it receives both sides of a conversation, a firewall may be able to obtain a significant amount of information concerning that conversation and to use this information to control access. A filter can maintain some limited state information on a unidirectional flow of packets, but it cannot determine the state of the bi-directional conversation as precisely as a firewall can.

5.3.1. Filtering

It is relatively common for routers to filter data packets. That is, routers can look for particular values in certain fields of the IP or higher level (e.g., TCP or UDP) headers. Packets that match the criteria associated with a particular filter may be either discarded or given special treatment.

In discussing filters, it is useful to separate the filter characteristics that may be used to determine whether a packet matches a filter from the packet actions that are applied to packets that match a particular filter.

o Filter Characteristics

  Filter characteristics are used to determine whether a particular packet or set of packets matches a particular filter.

  In many cases, filter characteristics may be stateless. A stateless filter determines whether a particular packet matches a filter based solely on the filter definition, on normal forwarding information (such as the next hop for a packet), and on the characteristics of that individual packet. Typically, stateless filters may consider the incoming and outgoing logical or physical interface, information in the IP header, and information in higher layer headers such as the TCP or UDP header. Information in the IP header to be considered may, for example, include source and destination IP address, Protocol field, Fragment Offset, and TOS field. Filters may also consider fields in the TCP or UDP header such as the Port fields and the SYN field in the TCP header.

  Stateful filtering maintains packet-specific state information to aid in determining whether a filter has been met. For example, a device might apply stateless filters to the first fragment of a fragmented IP packet. If the filter matches, then the data unit ID may be remembered, and other fragments of the same packet may then be considered to match the same filter. Stateful filtering is more commonly done in firewalls, although firewall technology may be added to routers.

- o Actions Based on Filter Results

  If a packet, or a series of packets, match a specific filter, then there are a variety of actions that may be taken based on that filter match. Examples of such actions include:

  - Discard

    In many cases, filters may be set to catch certain undesirable packets. Examples may include packets with forged or invalid source addresses, packets that are part of a DoS or DDoS attack, or packets that are trying to access forbidden resources (such as network management packets from an unauthorized source). Where such filters are activated, it is common to silently discard the packet or set of packets matching the filter. The discarded packets may also be counted and/or logged, of course.

  - Set CoS

    A filter may be used to set the Class of Service associated with the packet.

  - Count Packets and/or Bytes

  - Rate Limit

    In some cases, the set of packets that match a particular filter may be limited to a specified bandwidth. Packets and/or bytes would be counted and forwarded normally up to the specified limit. Excess packets may be discarded or marked (for example, by setting a "discard eligible" bit in the IP ToS field or the MPLS EXP field).

  - Forward and Copy

    It is useful in some cases not only to forward some set of packets normally, but also to send a copy to a specified other address or interface. For example, this may be used to implement a lawful intercept capability, or to feed selected packets to an Intrusion Detection System.

- o Other Issues Related to Packet Filters

  There may be a very wide variation in the performance impact of filtering. This may occur both due to differences between implementations, and due to differences between types or numbers

of filters deployed. For filtering to be useful, the performance of the equipment has to be acceptable in the presence of filters.

The precise definition of "acceptable" may vary from service provider to service provider and may depend on the intended use of the filters. For example, for some uses a filter may be turned on all the time in order to set CoS, to prevent an attack, or to mitigate the effect of a possible future attack. In this case it is likely that the service provider will want the filter to have minimal or no impact on performance. In other cases, a filter may be turned on only in response to a major attack (such as a major DDoS attack). In this case a greater performance impact may be acceptable to some service providers.

A key consideration with the use of packet filters is that they can provide few options for filtering packets carrying encrypted data. Because the data itself is not accessible, only packet header information or other unencrypted fields can be used for filtering.

5.3.2. Firewalls

Firewalls provide a mechanism for control over traffic passing between different trusted zones in the PPVPN model, or between a trusted zone and an untrusted zone. Firewalls typically provide much more functionality than filters, as they may be able to apply detailed analysis and logical functions to flows and not just to individual packets. They may offer a variety of complex services, such as threshold-driven denial-of-service attack protection, virus scanning, or acting as a TCP connection proxy. As with other access control techniques, the value of firewalls depends on a clear understanding of the topologies of the PPVPN core network, the user networks, and the threat model. Their effectiveness depends on a topology with a clearly defined inside (secure) and outside (not secure).

Within the PPVPN framework, traffic typically is not allowed to pass between the various user VPNs. This inter-VPN isolation is usually not performed by a firewall, but it is a part of the basic VPN mechanism. An exception to the total isolation of VPNs is the case of "extranets", which allow specific external access to a user’s VPN, potentially from another VPN. Firewalls can be used to provide the services required for secure extranet implementation.

In a PPVPN, firewalls can be applied between the public Internet and user VPNs, in cases where Internet access services are offered by the provider to the VPN user sites. In addition, firewalls may be applied between VPN user sites and any shared network-based services offered by the PPVPN provider.

Firewalls may be applied to help protect PPVPN core network functions from attacks originating from the Internet or from PPVPN user sites, but typically other defensive techniques will be used for this purpose.

Where firewalls are employed as a service to protect user VPN sites from the Internet, different VPN users, and even different sites of a single VPN user, may have varying firewall requirements. The overall PPVPN logical and physical topology, along with the capabilities of the devices implementing the firewall services, will have a significant effect on the feasibility and manageability of such varied firewall service offerings.

Another consideration with the use of firewalls is that they can provide few options for handling packets carrying encrypted data. As the data itself is not accessible, only packet header information, other unencrypted fields, or analysis of the flow of encrypted packets can be used for making decisions on accepting or rejecting encrypted traffic.

5.3.3. Access Control to Management Interfaces

Most of the security issues related to management interfaces can be addressed through the use of authentication techniques described in the section on authentication. However, additional security may be provided by controlling access to management interfaces in other ways.

Management interfaces, especially console ports on PPVPN devices, may be configured so that they are only accessible out of band, through a system that is physically or logically separated from the rest of the PPVPN infrastructure.

Where management interfaces are accessible in-band within the PPVPN domain, filtering or firewalling techniques can be used to restrict unauthorized in-band traffic from having access to management interfaces. Depending on device capabilities, these filtering or firewalling techniques can be configured either on other devices through which the traffic might pass, or on the individual PPVPN devices themselves.

5.4. Use of Isolated Infrastructure

One way to protect the infrastructure used for support of VPNs is to separate the VPN support resources from the resources used for other purposes (such as support of Internet services). In some cases, this may require the use of physically separate equipment for VPN services, or even a physically separate network.

For example, PE-based L3 VPNs may be run on a separate backbone not connected to the Internet, or they may use separate edge routers from those used to support Internet service. Private IP addresses (local to the provider and non-routable over the Internet) are sometimes used to provide additional separation.

It is common for CE-based L3VPNs to make use of CE devices that are dedicated to one specific VPN. In many or most cases, CE-based VPNs may make use of normal Internet services to interconnect CE devices.

5.5. Use of Aggregated Infrastructure

In general it is not feasible to use a completely separate set of resources for support of each VPN. One of the main reasons for VPN services is to allow sharing of resources between multiple users, including multiple VPNs. Thus, even if VPN services make use of a separate network from Internet services, there will still be multiple VPN users sharing the same network resources. In some cases, VPN services will share the use of network resources with Internet services or other services.

It is therefore important for VPN services to provide protection between resource use by different VPNs. Thus, a well-behaved VPN user should be protected from possible misbehavior by other VPNs. This requires that limits be placed on the amount of resources that can be used by any one VPN. For example, both control traffic and user data traffic may be rate limited. In some cases or in some parts of the network where a sufficiently large number of queues are available, each VPN (and, optionally, each VPN and CoS within the VPN) may make use of a separate queue. Control-plane resources such as link bandwidth and CPU and memory resources may be reserved on a per-VPN basis.

The techniques that are used to provision resource protection between multiple VPNs served by the same infrastructure can also be used to protect VPN services from Internet services.

The use of aggregated infrastructure allows the service provider to benefit from stochastic multiplexing of multiple bursty flows and may

also, in some cases, thwart traffic pattern analysis by combining the data from multiple VPNs.

5.6. Service Provider Quality Control Processes

Deployment of provider-provisioned VPN services requires a relatively large amount of configuration by the service provider. For example, the service provider has to configure which VPN each site belongs to, as well as QoS and SLA guarantees. This large amount of required configuration leads to the possibility of misconfiguration.

It is important for the service provider to have operational processes in place to reduce the potential impact of misconfiguration. CE-to-CE authentication may also be used to detect misconfiguration when it occurs.

5.7. Deployment of Testable PPVPN Service

This refers to solutions that can readily be tested for correct configuration. For example, for a point-point VPN, checking that the intended connectivity is working largely ensures that there is not connectivity to some unintended site.

6. Monitoring, Detection, and Reporting of Security Attacks

A PPVPN service may be subject to attacks from a variety of security threats. Many threats are described in another part of this document. Many of the defensive techniques described in this document and elsewhere provide significant levels of protection from a variety of threats. However, in addition to silently employing defensive techniques to protect against attacks, PPVPN services can add value for both providers and customers by implementing security-monitoring systems that detect and report on any security attacks that occur, regardless of whether the attacks are effective.

Attackers often begin by probing and analyzing defenses, so systems that can detect and properly report these early stages of attacks can provide significant benefits.

Information concerning attack incidents, especially if available quickly, can be useful in defending against further attacks. It can be used to help identify attackers and their specific targets at an early stage. This knowledge about attackers and targets can be used to further strengthen defenses against specific attacks or attackers, or to improve the defensive services for specific targets on an as-needed basis. Information collected on attacks may also be useful in identifying and developing defenses against novel attack types.

Monitoring systems used to detect security attacks in PPVPNs will typically operate by collecting information from Provider Edge (PE), Customer Edge (CE), and/or Provider backbone (P) devices. Security monitoring systems should have the ability to actively retrieve information from devices (e.g., SNMP get) or to passively receive reports from devices (e.g., SNMP notifications). The specific information exchanged will depend on the capabilities of the devices and on the type of VPN technology. Particular care should be given to securing the communications channel between the monitoring systems and the PPVPN devices.

The CE, PE, and P devices should employ efficient methods to acquire and communicate the information needed by the security monitoring systems. It is important that the communication method between PPVPN devices and security monitoring systems be designed so that it will not disrupt network operations. As an example, multiple attack events may be reported through a single message, rather than allow each attack event to trigger a separate message, which might result in a flood of messages, essentially becoming a denial-of-service attack against the monitoring system or the network.

The mechanisms for reporting security attacks should be flexible enough to meet the needs of VPN service providers, VPN customers, and regulatory agencies. The specific reports will depend on the capabilities of the devices, the security monitoring system, the type of VPN, and the service level agreements between the provider and customer.

7. User Security Requirements

This section defines a list of security-related requirements that the users of PPVPN services may have for their PPVPN service. Typically, these translate into requirements for the provider in offering the service.

The following sections detail various requirements that ensure the security of a given trusted zone. Since in real life there are various levels of security, a PPVPN may fulfill any or all of these security requirements. This document does not state that a PPVPN must fulfill all of these requirements to be secure. As mentioned in the Introduction, it is not within the scope of this document to define the specific requirements that each VPN technology must fulfill in order to be secure.

7.1. Isolation

A virtual private network usually defines "private" as isolation from other PPVPNs and the Internet. More specifically, isolation has several components, which are discussed in the following sections.

7.1.1. Address Separation

A given PPVPN can use the full Internet address range, including private address ranges [RFC1918], without interfering with other PPVPNs that use PPVPN services from the same service provider(s). When Internet access is provided (e.g., by the same service provider that is offering PPVPN service), NAT functionality may be needed.

In layer-2 VPNs, the same requirement exists for the layer 2 addressing schemes, such as MAC addresses.

7.1.2. Routing Separation

A PPVPN core must maintain routing separation between the trusted zones. This means that routing information must not leak from any trusted zone to any other, unless the zones are specifically engineered this way (e.g., for Internet access.)

In layer-2 VPNs, the switching information must be kept separate between the trusted zones, so that switching information of one PPVPN does not influence other PPVPNs or the PPVPN core.

7.1.3. Traffic Separation

Traffic from a given trusted zone must never leave this zone, and traffic from another zone must never enter this zone. Exceptions are made where zones are is specifically engineered that way (e.g., for extranet purposes or Internet access.)

7.2. Protection

The common perception is that a completely separated "private" network has defined entry points and is only subject to attack or intrusion over those entry points. By sharing a common core, a PPVPN appears to lose some of these clear interfaces to networks outside the trusted zone. Thus, one of the key security requirements of PPVPN services is that they offer the same level of protection as private networks.

7.2.1. Protection against Intrusion

An intrusion is defined here as the penetration of a trusted zone from outside. This could be from the Internet, another PPVPN, or the core network itself.

The fact that a network is "virtual" must not expose it to additional threats over private networks. Specifically, it must not add new interfaces to other parts outside the trusted zone. Intrusions from known interfaces such as Internet gateways are outside the scope of this document.

7.2.2. Protection against Denial-of-Service Attacks

A denial-of-service (DoS) attack aims at making services or devices unavailable to legitimate users. In the framework of this document, only those DoS attacks are considered that are a consequence of providing network service through a VPN. DoS attacks over the standard interfaces into a trusted zone are not considered here.

The requirement is that a PPVPN is not more vulnerable against DoS attacks than it would be if the same network were private.

7.2.3. Protection against Spoofing

It must not be possible to violate the integrity of a PPVPN by changing the sender identification (source address, source label, etc) of traffic in transit. For example, if two CEs are connected to the same PE, it must not be possible for one CE to send crafted packets that make the PE believe those packets are coming from the other CE, thus inserting them into the wrong PPVPN.

7.3. Confidentiality

This requirement means that data must be cryptographically secured in transit over the PPVPN core network to avoid eavesdropping.

7.4. CE Authentication

Where CE authentication is provided, it is not possible for an outsider to install a CE and pretend to belong to a specific PPVPN to which this CE does not belong in reality.

7.5. Integrity

Data in transit must be secured in such a manner that it cannot be altered or that any alteration may be detected at the receiver.

7.6. Anti-replay

Anti-replay means that data in transit cannot be recorded and replayed later. To protect against anti-replay attacks, the data must be cryptographically secured.

Note: Even private networks do not necessarily meet the requirements of confidentiality, integrity, and anti-reply. Thus, when private and "virtually private" PPVPN services are compared, these requirements are only applicable if the comparable private service also included these services. However, the fact that VPNs operate over a shared infrastructure may make some of these requirements more important in a VPN environment than in a private network environment.

8. Provider Security Requirements

In this section, we discuss additional security requirements that the provider may have in order to secure its network infrastructure as it provides PPVPN services.

The PPVPN service provider requirements defined here are the requirements for the PPVPN core in the reference model. The core network can be implemented with different types of network technologies, and each core network may use different technologies to provide the PPVPN services to users with different levels of offered security. Therefore, a PPVPN service provider may fulfill any number of the security requirements listed in this section. This document does not state that a PPVPN must fulfill all of these requirements to be secure.

These requirements are focused on 1) how to protect the PPVPN core from various attacks outside the core, including PPVPN users and non-PPVPN alike, both accidentally and maliciously, and 2) how to protect the PPVPN user VPNs and sites themselves. Note that a PPVPN core is not more vulnerable against attacks than a core that does not provide PPVPNs. However, providing PPVPN services over such a core may lead to additional security requirements, if only because most users are expecting higher security standards in a core delivering PPVPN services.

8.1. Protection within the Core Network

8.1.1. Control Plane Protection

- Protocol Authentication within the Core:

  PPVPN technologies and infrastructure must support mechanisms for authentication of the control plane. For an IP core, IGP and BGP

sessions may be authenticated by using TCP MD5 or IPsec. If an MPLS core is used, LDP sessions may be authenticated by using TCP MD5. In addition, IGP and BGP authentication should also be considered. For a core providing layer-2 services, PE to PE authentication may also be used via IPsec.

With the cost of authentication coming down rapidly, the application of control plane authentication may not increase the cost of implementation for providers significantly, and it will improve the security of the core. If the core is dedicated to VPN services and there are no interconnects to third parties, then it may reduce the requirement for authentication of the core control plane.

- Elements protection

  Here we discuss means to hide the provider’s infrastructure nodes.

  A PPVPN provider may make the infrastructure routers (P and PE routers) unreachable by outside users and unauthorized internal users. For example, separate address space may be used for the infrastructure loopbacks.

  Normal TTL propagation may be altered to make the backbone look like one hop from the outside, but caution should be taken for loop prevention. This prevents the backbone addresses from being exposed through trace route; however, it must also be assessed against operational requirements for end-to-end fault tracing.

  An Internet backbone core may be re-engineered to make Internet routing an edge function, for example, by using MPLS label switching for all traffic within the core and possibly by making the Internet a VPN within the PPVPN core itself. This helps detach Internet access from PPVPN services.

  PE devices may implement separate control plane, data plane, and management plane functionality in terms of hardware and software, to improve security. This may help limit the problems when one particular area is attacked, and it may allow each plane to implement additional security measurement separately.

  PEs are often more vulnerable to attack than P routers, since, by their very nature, PEs cannot be made unreachable to outside users. Access to core trunk resources can be controlled on a per-user basis by the application of inbound rate-limiting/shaping. This can be further enhanced on a per-Class of Service basis (see section 8.2.3).

In the PE, using separate routing processes for Internet and PPVPN service may help improve the PPVPN security and better protect VPN customers. Furthermore, if the resources, such as CPU and memory, may be further separated based on applications, or even on individual VPNs, it may help provide improved security and reliability to individual VPN customers.

Many of these were not particular issues when an IP core was designed to support Internet services only. Providing PPVPN services introduces new security requirements for VPN services. Similar consideration apply to L2 VPN services.

8.1.2. Data Plane Protection

PPVPN using IPsec technologies provides VPN users with encryption of secure user data.

In today's MPLS, ATM, and Frame Relay networks, encryption is not provided as a basic feature. Mechanisms can be used to secure the MPLS data plane and to secure the data carried over the MPLS core. Additionally, if the core is dedicated to VPN services and there are no external interconnects to third party networks, then there is no obvious need for encryption of the user data plane.

Inter-working IPsec/L3 PPVPN technologies or IPsec/L2 PPVPN technologies may be used to provide PPVPN users with end-to-end PPVPN services.

8.2. Protection on the User Access Link

Peer/Neighbor protocol authentication may be used to enhance security. For example, BGP MD5 authentication may be used to enhance security on PE-CE links using eBGP. In the case of an inter-provider connection, authentication/encryption mechanisms between ASes, such as IPsec, may be used.

WAN link address space separation for VPN and non-VPN users may be implemented to improve security in order to protect VPN customers if multiple services are provided on the same PE platform.

Firewall/Filtering: Access control mechanisms can be used to filter out any packets destined for the service provider's infrastructure prefix or to eliminate routes identified as illegitimate.

Rate limiting may be applied to the user interface/logical interfaces against DDoS bandwidth attack. This is very helpful when the PE device is supporting both VPN services and Internet services, especially when it supports VPN and Internet services on the same physical interfaces through different logical interfaces.

8.2.1. Link Authentication

Authentication mechanisms can be employed to validate site access to the PPVPN network via fixed or logical (e.g., L2TP, IPsec) connections. When the user wishes to hold the 'secret' associated to acceptance of the access and site into the VPN, then PPVPN based solutions require the flexibility for either direct authentication by the PE itself or interaction with a customer PPVPN authentication server. Mechanisms are required in the latter case to ensure that the interaction between the PE and the customer authentication server is controlled, for example, by limiting it simply to an exchange in relation to the authentication phase and with other attributes (e.g., optional filtering of RADIUS).

8.2.2. Access Routing

Mechanisms may be used to provide control at a routing protocol level (e.g., RIP, OSPF, BGP) between the CE and PE. Per-neighbor and per-VPN routing policies may be established to enhance security and reduce the impact of a malicious or non-malicious attack on the PE, in particular, the following mechanisms should be considered:

- Limiting the number of prefixes that may be advertised into the PE on a per-access basis . Appropriate action may be taken should a limit be exceeded; for example, the PE might shut down the peer session to the CE.

- Applying route dampening at the PE on received routing updates.

- Definition of a per-VPN prefix limit, after which additional prefixes will not be added to the VPN routing table.

In the case of inter-provider connection, access protection, link authentication, and routing policies as described above may be applied. Both inbound and outbound firewall/filtering mechanism may be applied between ASes. Proper security procedures must be implemented in inter-provider VPN interconnection to protect the providers' network infrastructure and their customer VPNs. This may be custom designed for each inter-Provider VPN peering connection, and both providers must agree on it.

8.2.3. Access QoS

PPVPN providers offering QoS-enabled services require mechanisms to ensure that individual accesses are validated against their subscribed QOS profile and are granted access to core resources that match their service profile. Mechanisms such as per-Class of Service rate limiting/traffic shaping on ingress to the PPVPN core are one option in providing this level of control. Such mechanisms may require the per-Class of Service profile to be enforced by marking, remarking, or discarding traffic that is outside of the profile.

8.2.4. Customer VPN Monitoring Tools

End users requiring visibility of VPN-specific statistics on the core (e.g., routing table, interface status, QoS statistics) impose requirements for mechanisms at the PE both to validate the incoming user and to limit the views available to that particular user’s VPN. Mechanisms should also be considered to ensure that such access cannot be used to create a DoS attack (either malicious or accidental) on the PE itself. This could be accomplished either through separation of these resources within the PE itself or via the capability to rate-limit such traffic on a per-VPN basis.

8.3. General Requirements for PPVPN Providers

The PPVPN providers must support the users’ security requirements as listed in Section 7. Depending on the technologies used, these requirements may include the following.

- User control plane separation: Routing isolation.

- User address space separation: Supporting overlapping addresses from different VPNs.

- User data plane separation: One VPN traffic cannot be intercepted by other VPNs or any other users.

- Protection against intrusion, DoS attacks and spoofing.

- Access Authentication.

- Techniques highlighted through this document identify methodologies for the protection of PPVPN resources and infrastructure.

Hardware or software bugs in equipment that lead to security breaches are outside the scope of this document.

9. Security Evaluation of PPVPN Technologies

This section presents a brief template that may be used to evaluate and summarize how a given PPVPN approach (solution) measures up against the PPVPN Security Framework. An evaluation using this template should appear in the applicability statement for each PPVPN approach.

9.1. Evaluating the Template

The first part of the template is in the form of a list of security assertions. For each assertion the approach is assessed and one or more of the following ratings is assigned:

- The requirement is not applicable to the VPN approach because ... (fill in reason).

- The base VPN approach completely addresses the requirement by ... (fill in technique).

- The base VPN approach partially addresses the requirement by ... (fill in technique and extent to which it addresses the requirement).

- An optional extension to the VPN approach completely addresses the requirement by ... (fill in technique).

- An optional extension to the VPN approach partially addresses the requirement by ... (fill in technique and extent to which it addresses the requirement).

- The requirement is addressed in a way that is beyond the scope of the VPN approach. (Explain.) (One example of this would be a VPN approach in which some aspect, such as membership discovery, is done via configuration. The protection afforded to the configuration would be beyond the scope of the VPN approach.).

- The VPN approach does not meet the requirement.

9.2. Template

The following assertions solicit responses of the types listed in the previous section.

1. The approach provides complete IP address space separation for each L3 VPN.

2. The approach provides complete L2 address space separation for each L2 VPN.

3. The approach provides complete VLAN ID space separation for each L2 VPN.

4. The approach provides complete IP route separation for each L3 VPN.

5. The approach provides complete L2 forwarding separation for each L2 VPN.

6. The approach provides a means to prevent improper cross-connection of sites in separate VPNs.

7. The approach provides a means to detect improper cross-connection of sites in separate VPNs.

8. The approach protects against the introduction of unauthorized packets into each VPN
   a. in the CE-PE link,
   b. in a single- or multi-provider PPVPN backbone, or
   c. in the Internet used as PPVPN backbone.

9. The approach provides confidentiality (secrecy) protection for PPVPN user data
   a. in the CE-PE link,
   b. in a single- or multi-provider PPVPN backbone, or
   c. in the Internet used as PPVPN backbone.

10. The approach provides sender authentication for PPVPN user data.
    a. in the CE-PE link,
    b. in a single- or multi-provider PPVPN backbone, or
    c. in the Internet used as PPVPN backbone.

11. The approach provides integrity protection for PPVPN user data
    a. in the CE-PE link,
    b. in a single- or multi- provider PPVPN backbone, or
    c. in the Internet used as PPVPN backbone.

12. The approach provides protection against replay attacks for PPVPN user data
    a. in the CE-PE link,
    b. in a single- or multi-provider PPVPN backbone, or
    c. in the Internet used as PPVPN backbone.

13. The approach provides protection against unauthorized traffic pattern analysis for PPVPN user data
    a. in the CE-PE link,
    b. in a single- or multi-provider PPVPN backbone, or
    c. in the Internet used as PPVPN backbone.

14. The control protocol(s) used for each of the following functions provides message integrity and peer authentication

    a. VPN membership discovery.
    b. Tunnel establishment.
    c. VPN topology and reachability advertisement:
       i. PE-PE.
       ii. PE-CE.
    d. VPN provisioning and management.
    e. VPN monitoring, attack detection, and reporting.
    f. Other VPN-specific control protocols, if any (list).

The following questions solicit free-form answers.

15. Describe the protection, if any, the approach provides against PPVPN-specific DoS attacks (i.e., inter-trusted-zone DoS attacks):

    a. Protection of the service provider infrastructure against Data Plane or Control Plane DoS attacks originated in a private (PPVPN user) network and aimed at PPVPN mechanisms.

    b. Protection of the service provider infrastructure against Data Plane or Control Plane DoS attacks originated in the Internet and aimed at PPVPN mechanisms.

    c. Protection of PPVPN users against Data Plane or Control Plane DoS attacks originated from the Internet or from other PPVPN users and aimed at PPVPN mechanisms.

16. Describe the protection, if any, the approach provides against unstable or malicious operation of a PPVPN user network

    a. Protection against high levels of, or malicious design of, routing traffic from PPVPN user networks to the service provider network.

    b. Protection against high levels of, or malicious design of, network management traffic from PPVPN user networks to the service provider network.

c. Protection against worms and probes originated in the PPVPN user networks, sent toward the service provider network.

17. Is the approach subject to any approach-specific vulnerabilities not specifically addressed by this template? If so, describe the defense or mitigation, if any, that the approach provides for each.

10. Security Considerations

Security considerations constitute the sole subject of this memo and hence are discussed throughout. Here we recap what has been presented and explain at a very high level the role of each type of consideration in an overall secure PPVPN system. The document describes a number of potential security threats. Some of these threats have already been observed occurring in running networks; others are largely theoretical at this time.

DoS attacks and intrusion attacks from the Internet against service provider infrastructure have been seen. DoS "attacks" (typically not malicious) have also been seen in which CE equipment overwhelms PE equipment with high quantities or rates of packet traffic or routing information. Operational/provisioning errors are cited by service providers as one of their prime concerns.

The document describes a variety of defensive techniques that may be used to counter the suspected threats. All of the techniques presented involve mature and widely implemented technologies that are practical to implement.

The document describes the importance of detecting, monitoring, and reporting both successful and unsuccessful attacks. These activities are essential for "understanding one’s enemy", mobilizing new defenses, and obtaining metrics about how secure the PPVPN service is. As such, they are vital components of any complete PPVPN security system.

The document evaluates PPVPN security requirements from a customer perspective and from a service provider perspective. These sections re-evaluate the identified threats from the perspectives of the various stakeholders and are meant to assist equipment vendors and service providers, who must ultimately decide what threats to protect against in any given equipment or service offering.

Finally, the document includes a template for use by authors of PPVPN technical solutions for evaluating how those solutions measure up against the security considerations presented in this memo.

11. Contributors

The following people made major contributions to writing this document: Michael Behringer, Ross Callon, Fabio Chiussi, Jeremy De Clerque, Paul Hitchen, and Paul Knignt.

Michael Behringer
Cisco
Village d'Entreprises Green Side, Phone: +33.49723-2652
400, Avenue Roumanille, Bat. T 3 EMail: mbehring@cisco.com
06410 Biot, Sophia Antipolis
France

Ross Callon
Juniper Networks
10 Technology Park Drive Phone: 978-692-6724
Westford, MA 01886 EMail: rcallon@juniper.net

Fabio Chiussi Phone: 1 978 367-8965
Airvana EMail: fabio@airvananet.com
19 Alpha Road
Chelmsford, Massachusetts 01824

Jeremy De Clercq
Alcatel
Fr. Wellesplein 1, 2018 Antwerpen EMail: jeremy.de_clercq@alcatel.be
Belgium

Mark Duffy
Sonus Networks
250 Apollo Drive Phone: 1 978-614-8748
Chelmsford, MA 01824 EMail: mduffy@sonusnet.com

Paul Hitchen
BT
BT Adastral Park
Martlesham Heath Phone: 44-1473-606-344
Ipswich IP53RE EMail: paul.hitchen@bt.com
UK

Paul Knight
Nortel
600 Technology Park Drive Phone: 978-288-6414
Billerica, MA 01821 EMail: paul.knight@nortel.com

12. Acknowledgement

The author and contributors would also like to acknowledge the helpful comments and suggestions from Paul Hoffman, Eric Gray, Ron Bonica, Chris Chase, Jerry Ash, and Stewart Bryant.

13. Normative References

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

[RFC2246] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", RFC 2246, January 1999.

[RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[RFC2402] Kent, S. and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.

[RFC2406] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.

[RFC2407] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, November 1998.

[RFC2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.

[RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[RFC3588] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.

[RFC3602] Frankel, S., Glenn, R., and S. Kelly, "The AES-CBC Cipher Algorithm and Its Use with IPsec", RFC 3602, September 2003.

[STD62] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, December 2002.

Case, J., Harrington, D., Presuhn, R., and B. Wijnen, "Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3412, December 2002.

Levi, D., Meyer, P., and B. Stewart, "Simple Network Management Protocol (SNMP) Applications", STD 62, RFC 3413, December 2002.

Blumenthal, U. and B. Wijnen, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)", STD 62, RFC 3414, December 2002.

Wijnen, B., Presuhn, R., and K. McCloghrie, "View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3415, December 2002.

Presuhn, R., "Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3416, December 2002.

Presuhn, R., "Transport Mappings for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3417, December 2002.

Presuhn, R., "Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3418, December 2002.

[STD8] Postel, J. and J. Reynolds, "Telnet Protocol Specification", STD 8, RFC 854, May 1983.

14. Informative References

[RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.

[RFC2411] Thayer, R., Doraswamy, N., and R. Glenn, "IP Security Document Roadmap", RFC 2411, November 1998.

[RFC3174] Eastlake 3rd, D. and P. Jones, "US Secure Hash Algorithm 1 (SHA1)", RFC 3174, September 2001.

[RFC3631] Bellovin, S., Schiller, J., and C. Kaufman, "Security Mechanisms for the Internet", RFC 3631, December 2003.

[RFC3889] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", RFC 3889, October 2004.

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

[RFC4031] Carugi, M. and D. McDysan, Eds., "Service Requirements for Layer 3 Provider Provisioned Virtual Private Networks (PPVPNs)", RFC 4031, April 2005.

[RFC4110] Callon, R. and M. Suzuki, "A Framework for Layer 3 Provider Provisioned Virtual Private Networks", RFC 4110, July 2005.

Author’s Address

Luyuan Fang
AT&T Labs.
200 Laurel Avenue, Room C2-3B35
Middletown, NJ 07748

Phone: 732-420-1921
EMail: luyuanfang@att.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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

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

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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