VIRTUAL PRIVATE NETWORKS

Submitted to

Dr. Prabhaker Mateti

By,

Karthik Mohanasundaram

Table of Contents

Virtual Private Networks

Abstract

This report describes on a general basis the term 'Virtual Private Networks' and its practical applications and the theory behind them. All the material provided in this report is a compilation of various articles present in the Internet. I do not claim authorship for the material presented below.

Educational Objective

The main objective of this study is to understand the concept 'Virtual Private Networks' and experiment them to see them working in a real environment and understand the concept and it's practical importance in a more detailed manner.

Introduction

A virtual private network (VPN) is the extension of a private network that encompasses links across shared or public networks like the Internet. A VPN connection enables you to send/receive data between two computers across a shared or public internetwork in a manner that emulates the properties of a point-to-point private link. A virtual private network can also be defined as an on-demand connection between two computers in different locations. It consists of the two computers (one computer at each end of the connection) and a route, or *tunnel*, over a public or private network. To ensure privacy and secure communication, data transmitted between the two computers is encrypted by the Point-to Point Protocol (PPP) (a remote access protocol) and then routed over a dial-up or LAN connection by a PPTP device.

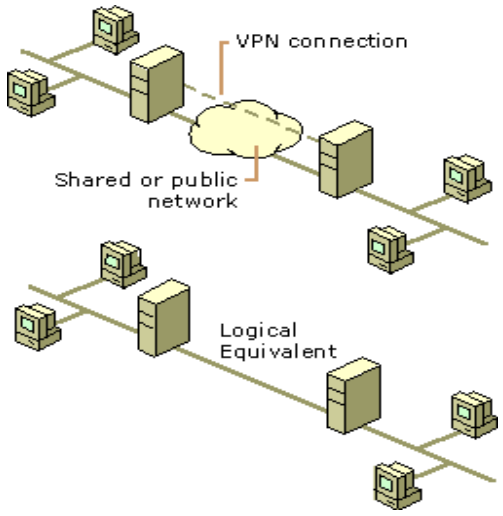

Figure 1: Virtual Private Network

VPN allows users to connect from anywhere in the world to any computer as long as they have legal permissions to access that computer. This technique is mainly useful for companies that want to interconnect all their offices, which are geographically apart. This reduces costs for the companies on a huge scale, as they need not maintain a dedicated link between their offices. A routed VPN connection across the public infrastructure is logically a dedicated WAN link. They can use the public infrastructure [Internet] to transfer all their data in a secure way to the other end. The VPN functions as a Point-to-Point connection between the client's computer and the corporate server that the client is trying to access. The infrastructure over which the data is transferred is irrelevant as the data sent over this public network is encrypted and it is sent on a private dedicated link that is not visible to the other users on the public network. The VPN connection actually gives the users access to the internal private network behind the VPN server. So users working from home actually get access to the company's internal network and they can work from home on the private network of the company itself. The VPN technology provides a basic platform over which all the private data is encapsulated in a packet and also the technique for the packet to traverse through the public network securely to the other end of the tunnel connection.

For the VPN connection to be established the client tries to connect to the VPN server and the VPN server authenticates the client that is attempting to connect by verifying whether the client has appropriate permissions to access the local network. The VPN client also authenticates the VPN server for protection purposes and after mutual authentication is successful the client is connected on to the internal network behind the VPN server. The data sent over the VPN tunnel is verified for integrity using a cryptographic checksum based on an encryption key known to the client and the server. The encrypted packet sent over the public network contains the private data encapsulated within it. This packet is unintelligible to anyone who doesn't have the knowledge of the encryption key used by the client and the server. The key length is an important factor in the encryption process as complex computation techniques can be used to calculate this encryption key.

Elements of a VPN connection

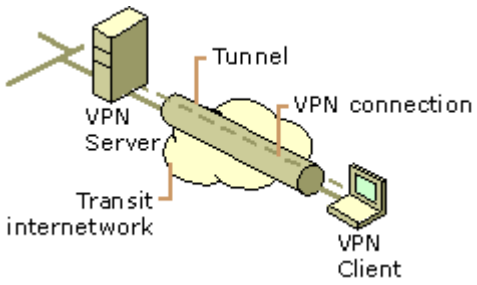

Figure 2: Components of a VPN Connection

A typical VPN connection encrypts the data that is communicated between the client and the server. The VPN connection involves the components explained below:

VPN Server:

A computer that accepts connection attempts from the client machines. The server mainly authenticates the clients into the network. It is the gateway between the client and the private network. The server has a virtual interface to which all the clients connect. Once the connection is established a virtual interface is created on the client side also. All the clients connected to the server connect to this virtual interface of the server. The VPN server maybe running any operating system [either Linux or Windows NT/2000]

VPN Client:

A computer that initiates a connection with the VPN server is called a VPN client. As soon as a connection is established a virtual interface is created on the client side. This virtual interface is connected with the virtual interface on the server side. All the encrypted data packets are sent over this virtual interface. This interface-interface connection forms the tunnel on top of the public network.

Tunnel:

The virtual interface formed on the client and server side at the time of connection establishment over which the encapsulated encrypted data travels is called the Tunnel.

Tunneling Protocols:

Communication standards used to manage tunnel connections and also the encapsulation of private data. The standard protocols used in VPN connections are:

- → Point-to-Point Tunneling Protocol [PPTP]
- → Layer Two Tunneling Protocol [L2TP] and
- → Internet Protocol Security [IPSec]

Transit Internetwork:

The shared or public network over which the data packet with the encrypted private data traverses is called the Transit Internetwork. The transit network can be the Public Internet or a private Intranet.

VPN Connections:

Basically VPN connections can be split into three categories based on the Transit Internetwork over which the network traffic is routed. They are:

- → Intranet-Based VPN Connections
- → Internet-Based VPN Connections and
- → Combined Internet and Intranet VPN Connections

The two kinds of VPN connections widely used are:

- ♦ Router-to-Router VPN Connection and
- ♦ Remote access VPN connection

Router-to-Router VPN Connection:

In the Router-to-Router connection two routers connect to each other. A Router is the server and another router is the client. The routers authenticate themselves mutually and after the authentication is successful the private networks behind these routers are connected to each other and the machines that are part of these two networks communicate with each other. The data packets that travel through the tunnel are not originating from both the client and server but the machines behind these routers. This connection is very useful for connecting offices as the authentication done is only between the two routers and this reduces the number of individual connections established between the machines of both the offices. The machines behind these routers can communicate between any machines on the other side.

Remote Access VPN Connection:

A typical Remote Access VPN connection is shown in Figure 2. In this kind of connection a remote client tries to connect to the server. The server and the client authenticate each other mutually and the client is authorized to access the server as well as the private network

behind the server. The data packets originate from the remote client machine and end in the network behind the VPN server.

Kind of VPN Systems

There are various ways of implementing a VPN solution. The different ways are:

- → SSH & PPP
- → PPTP
- → IPSec
- → CIPE

SSH & PPP:

This system uses SSH to create the tunnel connection through which the encrypted data packets are transferred and then use pppd to run TCP/IP traffic through the tunnel. This system is best described in the VPN HOW-TO given in the references section.

PPTP:

PPTP stands for Point-to-Point Tunneling Protocol proposed by Microsoft for implementing VPN solutions. PPTP encapsulates PPP frames into IP datagrams for transmission. The PPTP consists of two communication channels between the client and the VPN server. One channel is dedicated for creating, maintaining and terminating the Tunnel [link management] and another channel that takes care of the encrypted private network traffic is carried. The PPTP protocol does not take care of encrypting the data packets and the data channel that carries the data uses a modified version of the IP protocol 47 (GRE) Generic Encapsulation Protocol that takes care of the encryption. The PPTP protocol has several security issues to be considered and is not very secure. PPTP is documented in RFC 2637.

IPSec:

IPSec stands for Internet Protocol Security. This is a standard set of protocols for implementing secure VPN connections. The IPSec protocol gives an endpoint-endpoint communication between the client and the server. The IPSec VPN implementation consists of two communication channels. One channel handles the authentication and encryption key details are passed. There are one or more channels that handle the encrypted private data traffic. The IPSec protocol is documented in RFC 2402. The channel that takes care of the key

exchange is a standard UDP connection and the data channel operates on the IP protocol number 50 [ESP]. Windows 2000 provides IPSec implementation for Windows based systems and FreeS/WAN provides the Linux version of IPSec implementation.

CIPE:

CIPE is very well suitable for enterprise VPN setups as it provides kernel level encryption for the network traffic. This connection tunnels IP packets inside encrypted UDP packets. CIPE uses the well known Blowfish and IDEA crypto algorithms, which have a key length of 128 bits.

Experimental Setup:

The Experiment was carried out on Intel machines running Linux Mandrake 7.0 with 2.2.19 version kernel. The general setup for experimenting with a VPN connection is given below.

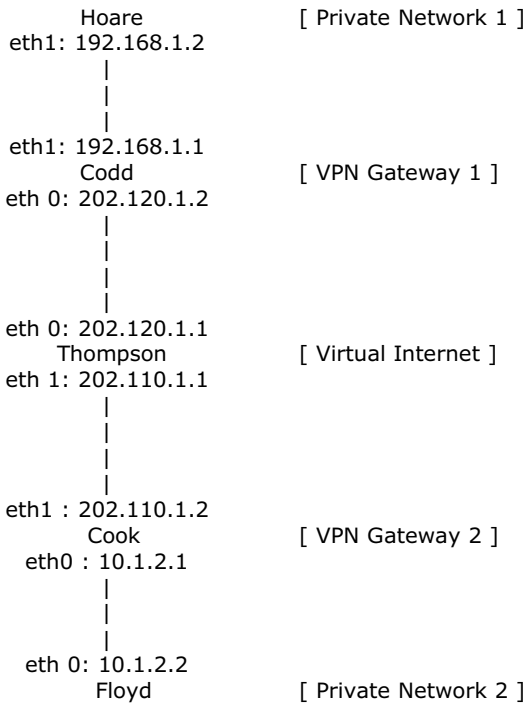

Figure 3: Experimental Setup

Experiment Setup Description:

The Experimental setup used to test the VPN connections is shown above in Fig 3. The setup contains 5 machines in total. The machines used are:

Hoare - Private Network Machine
Codd - VPN gateway 1
Thompson - Virtual Public Internet
Cook - VPN gateway 2
Floyd - Private Network Machine

The two VPN gateways are used for setting up the Router-to-Router VPN connection. The same machines can be used for the Remote Access VPN connection also. The machine Thompson acts as a virtual public internet. The machine Floyd has eth0 of Cook as it's default gateway and in the same fashion the machine Hoare has eth1 of Codd as it's default gateway.

PPTP implementation of VPN:

The PPTP implementation was experimented with the server running on a Linux machine and the client being a Windows 98 machine. The Linux version of the PPTP server was downloaded, installed and configured properly.

The machine Codd acts as the VPN server and the machine Cook is the VPN client that is trying to connect to the VPN server. The machine is running PPPD daemon and Cook is running Windows 98 uses dial-up networking to connect to the VPN server. A new connection is created with the VPN adapter and Cook uses the ip-number of Codd as the telephone number to dial out.

The PPPD server configuration files are given below:

- /etc/ppp/options

--snip--

```
debug
name testvpn
auth
require-chap
proxyarp
```

--snip--

- /etc/pptpd.conf

--snip--

```
localip         10.1.2.1-240
remoteip        192.168.1.1-240
```

--snip--

- /etc/ppp/chap-secrets

--snip--

```
#username      servername      password        ip

bob            testvpn         billy           10.1.2.1
```

--snip--

After the entire above configuration files were setup then the pptpd daemon was started and the client was successful in establishing a connection with the username 'bob' and password 'billy'. The Cook machine successfully logged onto the VPN server Codd. To browse through the files present in the machines behind the VPN server Samba needs to be installed and configured.

In the file /etc/ppp/options we need to mention the name of the server that the user is trying to connect. For our testing purpose the name of the VPN server used here is called testvpn. The same server name is used in the file /etc/ppp/chap-secrets also as the user has all his accounts applicable to different servers configured in that file.

In the file /etc/pptpd.conf file the localip of the machine and the remote ipaddress range of machines to which it is trying to connect to is mentioned so that a route is formed from this machine to the other machines. So the VPN connection was successfully established.

IPSec Implementation of VPN:

The IPSec implementation was carried out using the ipsec setup got from the FreeS/WAN website. The site offers complete documentation and also example configuration files that can be used to setup VPN networks effectively. The IPSec implementation gives users an end-end encryption rathen than the link encryption provided by PPTP. IPSec implementation is considered to be more secure than the PPTP implementation of VPN's. The detailed

explanation of the setup is given in the IPSec documentation given in their website. The configuration files that were used for experimentation purpose are given below.

- /etc/ipsec.conf

--snip--

```
# /etc/ipsec.conf - FreeS/WAN IPsec configuration file

# More elaborate and more varied sample configurations can be found
# in FreeS/WAN's doc/examples file, and in the HTML documentation.
# basic configuration
config setup
       # THIS SETTING MUST BE CORRECT or almost nothing will work;
       # %defaultroute is okay for most simple cases.
       interfaces=%defaultroute
       # Debug-logging controls:  "none" for (almost) none, "all" for lots.
       klipsdebug=none
       plutodebug=none
       # Use auto= parameters in conn descriptions to control startup actions.
       plutoload=%search
       plutostart=%search
       # Close down old connection when new one using same ID shows up.
       uniqueids=yes


# connection details
conn codd-cook
       left=202.110.1.2
       leftsubnet=10.1.0.0/24
       leftnexthop=202.110.1.1
       right=202.120.1.2
       rightnexthop=202.120.1.1
       rightsubnet=192.168.1.0/24
       keyringtries=0
       auth=esp
       authby=rsasig
       leftrsasigkey = xx.. [key]
       rightrsasigkey = xx .. [key]
```

--snip--

- /etc/ipsec.secrets

--snip--

```
# This file holds shared secrets or RSA private keys for inter-Pluto
# authentication.  See ipsec_pluto(8) manpage, and HTML documentation.

# RSA private key for this host, authenticating it to any other host
# which knows the public part.  Suitable public keys, for ipsec.conf, DNS,
# or configuration of other implementations, can be extracted conveniently
# with "ipsec showhostkey".
: RSA  {
       # RSA 2048 bits   localhost.localdomain   Wed Nov 14 10:06:27 2001
       # for signatures only, UNSAFE FOR ENCRYPTION
       #pubkey= xx ..
       #IN KEY 0x4200 4 1 xx ..
```

```
        # (0x4200 = auth-only host-level, 4 = IPSec, 1 = RSA)
        Modulus: xx ..
        PublicExponent: 0x03
        # everything after this point is secret
        PrivateExponent: xx ..
        Prime1: xx ..
        Prime2: xx ..
        Exponent1: xx ..
        Exponent2: xx ..
        Coefficient: xx ..
        }
# do not change the indenting of that "}"

Note : [xx ..] represents actual RSA keys that can be generated as per commands
that are given in the actual documentation given in HTML format on the
FreeS/WAN webpage.
```

--snip--

After all these configuration files are configured on both the gateway machines the VPN connection is established and the configuration tested. To establish the connection the following steps are followed:

On one gateway, start IPSec with: *ipsec auto –up name*

where name is the connection name that is specified in the ipsec.conf

To shutdown the connection on both the gateways type:

ipsec auto –down name

Note: the down command is issued on both the gateways between which the connection is established.

If the ipsec auto –up command is successful then you can check the status of the VPN connection by:

ipsec look

This command displays the tunnel details and also the interface details.

CIPE Implementation of VPN:

CIPE is very well suitable for enterprise VPN setups as it provides kernel level encryption for the network traffic. This connection tunnels IP packets inside encrypted UDP packets. CIPE uses the well known Blowfish and IDEA crypto algorithms, which have a key length of 128 bits.

CIPE implementation was carried out on the same test network with the appropriate configurations. The configuration files are given below:

- /etc/cipe/options

--snip--

#options for Machine Codd

```
device              cipcb0
ptpaddr             192.168.110.1
ipaddr              192.168.120.1
me                  202.120.1.2:9000
peer                202.110.1.2:9000
key                 xx ..
```

--snip--

The only important configuration file is the options file. The first step in establishing the connection is to invoke the CIPE daemon, which reads this configuration file. The daemon brings up the CIPE interface whose name is specified in the options file by the parameter 'device'. The ptpaddr and ipaddr are the addresses of the VPN interfaces. The address of both these parameters should not be assigned to any other physical interface in the network.

Bringing up the VPN connection involves invoking the CIPE daemon and running the /etc/rc.d/rc.cipe script. This script is given as a standard script in the Linux+CIPE Masquerading mini How-To. This document as a whole explains the overall setup and configuration of the CIPE VPN implementation.

Troubles Faced during the Experiment:

The troubles that were faced during this experimentation are documented below:

- The VPN masquerading scripts that were explained in the documents were all for old version of Linux and the development of these scripts and patches were stopped after the 2.2.x kernels. This was a major drawback to the experiment, as I had to switch back to the old kernel for this experiment purpose.

- The patches that needed to be applied as specified in the documentation gave problems when patching them with the kernel. Troubleshooting them took some time and they proved unfruitful at the end.

- The lack of proper documentation for the various configuration setups was a hindrance in the quick understanding of concepts involved in the setups. I was learning from my mistakes and that took some precious time of the experiment.

- All the implementations involved lots of understanding and all the documents were outdated and not updated recently. The documentation demanded a more thorough understanding of the topic from the audience.

- The basic elementary understanding of the topic can be gained from numerous websites but the implementation details demanded more details, which were not easily available and had to be learned on a trial and error basis.

Conclusion

Thus the concept 'Virtual Private Networks' was successfully studied, understood and experimented in the private laboratory.

Future Work

The work that I intend to do in the near future are:

→ Prepare a very thorough and proper documentation of all the implementations that I had experimented with so that it could be useful for someone else who is also trying the experiment for the first time.

→ Implement all the VPN connection types on a more general basis so that the configuration can be implemented on a large scale and useful non-commercial purposes.

Frequently Asked Questions [FAQ's]

This is a collection of basic questions that are frequently asked on the topic VPN. The collection given below is compiled from various websites.

What is a virtual private network?

A virtual private network is essentially a system that allows two or more private networks to be connected over a publically accessible network, such as the Internet. It usually consists of an encrypted tunnel of some kind, although a VPN can take several forms, using different combinations of hardware and software technologies. They can exist between an individual machine and a private network, or a remote LAN and a private network. Security features differ from product to product, but most security experts agree that VPNs include encryption, strong authentication of remote users or hosts, and mechanisms for hiding or masking information about the private network topology from potential attackers on the public network.

What are the basic features of VPNs?

Aside from supporting basic LAN interfaces, a good VPN should have high-availability features such as redundant power supplies. Also, all VPNs require some kind of authorization protocol and encryption, although some companies may choose to opt out of the latter. Other advanced functions can be useful, such as data compression, routing ability, network address translation, bandwidth management capabilities and fail-over redundancy. When purchasing a ready-made VPN package from a solutions provider, it is often possible to get other bundled services to compliment the network, such as voice over IP and other hosted applications.

What types of VPNs are there, and what are their relative advantages and disadvantages?
Despite the large (and rapidly expanding) number of VPN products, all fall into three broad categories:

- Hardware-based Systems
- Firewall-based VPN's and
- Standalone VPN Applications

Most hardware-based VPN systems are encrypting routers. They are secure and easy to use, since they provide the nearest thing to "plug and play" encryption equipment available. They provide the highest network throughput of all VPN systems, since they don't waste processor overhead in running an operating system or other applications. However, they may not be as flexible as software-based systems. The best hardware VPN packages offer software-only clients for remote installation, and incorporate some of the access control features more traditionally managed by firewalls or other perimeter security devices.

Firewall-based VPN's take advantage of the firewall's security mechanisms, including restricting access to the internal network. They also perform address translation; satisfy requirements for strong authentication; and serve up real-time alarms and extensive logging. Most commercial firewalls also "harden" the host operating system kernel by stripping out dangerous or unnecessary services, providing additional security for the VPN server. OS protection is a major plus, since very few VPN application vendors supply guidance on OS security. Performance may be a concern, especially if the firewall is already loaded -- however, some firewall vendors offer hardware-based encryption processors to minimize the impact of VPN management on the system.

Software-based VPNs are ideal in situations where both endpoints of the VPN are not controlled by the same organization (typical for client support requirements or business partnerships), or when different firewalls and routers are implemented within the same organization. At the moment, standalone VPNs offer the most flexibility in how network traffic is managed. Many software-based products allow traffic to be tunneled based on address or protocol, unlike hardware-based products, which generally tunnel all traffic they handle, regardless of protocol. Tunneling specific traffic types is advantageous in situations where remote sites may see a mix of traffic some that needs transport over a VPN (such as entries to a database at headquarters) and some that doesn't (such as Web surfing). In situations where performance requirements are modest (such as users connecting over dial-up links), software-based VPNs may be the best choice.

Software-based systems are generally harder to manage than encrypting routers. They require familiarity with the host operating system, the application itself, and appropriate security mechanisms. And some software VPN packages require changes to routing tables and network addressing schemes.

Why would a company use a VPN?

A VPN service is an economical alternative to setting up a private network with expensive leased lines, as it can use existing IP infrastructure and equipment to connect remote users and offices. For offices with great distances between them, VPNs are ideal because they can provide connectivity for almost any location in the world, and without incurring long-distance charges. Also, the flexibility and relative simplicity of VPNs allows small- to medium-sized businesses the option to switch to a different provider, increase bandwidth, or add more offices to the network more freely than with other schemes.

How do companies use VPN's?

Once a company connects to a VPN server, it can either use the same applications that it normally uses to connect to the Internet, or it can purchase or rent the appropriate devices, depending on the scope of the network. It can then be used to connect LANs in different sites, or give customers, clients and consultants access to corporate resources, provided they have compatible software and can be authenticated. Often VPNs are useful for mobile workers such as salespeople, for home workers or day extenders.

Are extranets and VPN's the same thing?

Not really. An extranet is basically a glorified Web site, which allows clients or partners access to the corporate intranet for highly specific, often administrative functions. For example, an online newspaper's extranet might allow advertisers to change banner ads on its site. A VPN uses a protocol that allows a remote PC full access to a company's network neighbourhood, as if it were actually in the home office. Although extranets take a variety of forms, some of which can resemble a VPN, they do not have the same function. However, using a more sophisticated authentication and segmentation method, a company can build a separate extranet application on its VPN, possibly saving money in the process.

How do VPN's save money?

By using a relatively cheap local dial-up or broadband connection, companies using VPNs save on telecommunications costs, and also reduce long-distance phone charges. They also cut down on operational costs by outsourcing the management of equipment used for remote access, as well as reducing the number of access line running into a corporate site. In some cases, the company can "borrow" the necessary hardware from a VPN service provider, at no extra charge. Finally, a VPN can theoretically alleviate the support burden, as the public service provider is generally responsible for supporting its dial-up customers.

What about VPN performance?

There are a number of factors that can contribute to the VPN's performance. While some of the issues may be related to the hardware or software applications being used, much of it depends on the Internet itself. The availability and speed of IP services may differ from one area to the next, as well as the actual provider. Because of this, most VPN providers will not offer a guarantee on the latency of packets moving across the network. Performance also depends somewhat on the encryption scheme being used, as well as the client's ability to process it. Highly encrypted data takes considerably longer to transmit, especially on larger packets being sent through a dial-up line.

What are some common tunneling protocols?

The most popular tunneling protocols for VPNs are the Point-to-Point Tunneling Protocol (PPTP), Internet Protocol Security (IPSec), and Layer 2 Tunneling Protocol (L2TP), which combines PPTP and Cisco Systems' Layer-2 Forwarding (L2F). SOCKS 5 is yet another approach, which follows a proxy server model and is considered among the most secure. Companies with very low security requirements may consider other alternatives, such as Secure Shell (SSH).

What type of encryption can be used?

Modern VPNs can use just about any common encryption technology available, and equipment vendors usually give their customers the choice. Triple DES and 3DES seem to be the standards in North America, although in some countries encryption strength is regulated by legislation, and must use a less robust technology. Whether hardware- or software-based, all VPN providers offer some sort of encryption scheme, which can often be customized to suit the buyer.

How are VPN users authenticated?

VPNs usually take some sort of firewall, often a surprisingly simple "plug-and-play" solution provided by a vendor. The system is installed on as many LANs as needed, and keys are exchanged between the users in order to provide authentication. All VPNs require that an access device be configured to recognize and authenticate remote users. A wide number of techniques and products, both hardware- and software-based, are available from vendors. Stronger and more advanced authentication techniques, such as tokens or regulated access levels, can also be implemented.

What is PPTP? Is it a good VPN solution?

A consortium that included Ascend Communications, 3Com/Primary Access, ECI Telematics, U.S. Robotics and Microsoft originally developed the PPTP specification. The protocol was originally designed as an encapsulation mechanism, to allow the transport of non-TCP/IP protocols (such as IPX) over the Internet using Generic Routing Encapsulation (GRE). The specification itself is fairly generic, and allows for a variety of authentication mechanisms and encryption algorithms. Note that these security features were added later, not built in from the beginning.

Several vendors have created PPTP systems. However, the vast majority of PPTP users implement the Microsoft version. The following discussions of PPTP security issues are specific to the Microsoft implementation, which features:

- PPTP server -- NT 4.0 or later
- PPTP clients -- Win 95/98/NT; WFW, Macintosh with 3rd party hardware
- Authentication/authorization mechanisms limited to NT domain security; manage access to non-NT domain resources via network segregation, RADIUS (maybe)

PPTP can be used to control access to the private network via NT domain security controls (user- and group-level access to domain resources), and by segregating resources on the corporate network. With the release of the Internet Authentication Services update for NT 4.0, RADIUS may be used to perform PPTP authentication -- but it is unknown whether or not the authorization and access control features of RADIUS are also supported.
Use of PPTP requires that IP forwarding be enabled on the NT server.
Setting up a PPTP system requires configuring the Remote Access Server capability on the NT server, adding routing functionality to the RAS system, applying several newly released security patches, configuring the PPTP-specific registry keys and hardening the server itself.

<u>Security Concerns:</u>

- Flawed encryption mechanism -- non-random keys, session keys weak hash of user password, key lengths too short (non-configurable)
- Bad password management in mixed Win95/NT environment; static passwords easily compromised
- Vulnerable to server spoofing attacks because packet authentication not implemented, easy denial-of-service attacks even inside firewalls
- MS claims cryptographic weaknesses not yet exploited

The initial release of PPTP used the MSCHAP mechanism for end-user authentication. After numerous criticisms that MSCHAP was easily compromised, especially in situations when Windows 95 was the client operating system, Microsoft released a patch to the original authentication protocol. To quote the Microsoft WebSite: "This new protocol provides mutual authentication, stronger initial data encryption keys, and different encryption keys for the transmit and receive paths. To minimize the risk of password compromise during MSCHAP exchanges, MSCHAP V2 drops support for the MSCHAP password change V1, and will not transmit the LMHash encoding of the password. ...For VPN connection requests, a Windows NT server will offer MSCHAP V2 before offering the legacy MSCHAP. Updated Windows clients (all platforms) will accept MSCHAP V2 when it is offered." (August 18, 1998) Microsoft also added a new registry key, SecureVPN, which forces incoming VPN connection requests to use the new authentication mechanism. These changes should prevent a PPTP client from indicating using the older, LMHash mechanism. However, any independent reviewer has not yet verified the effectiveness of these patches.

[Note that the dependence of PPTP authentication on MSCHAP makes it vulnerable to attacks using l0phtcrack [http://www.l0oht.com/l0phtcrack/] -- so it's the only VPN tool with its own l0pht hyperlink!]

Also note that although Microsoft describes PPTP as using either 40-bit or 128-bit encryption, their use of the user's password to create a session key, rather than a randomly generated key, greatly reduces the strength of the encryption process. *None of the recent security releases addresses this issue.*

Microsoft claims to have improved the mechanism that generates session keys (which is based on a hash of the user's password). If this is true, it helps protect against hijacking attacks, as well as making brute force crypto attacks harder. NB: even this enhancement does not improve the cryptographic weakness, which is based on the flawed decision to use passwords to generate keys. Remember, no matter how strong an encryption algorithm is, it can be compromised via a brute-force attack. The only protection against brute force is a long key length, with purely random keys - not what Microsoft has implemented. This enhancement has not been verified (as of November 1998) by any third-party evaluator.
[http://kubarb.phsx.ukans.edu/~tbird/vpn/PPTPrefs.html]

What is IPSec? How does it relate to VPNs and firewalls?

IPSec is an evolving standard for secure private communications over the Internet. Normal IPv4 packets consist of headers and payload, both of which contain information of value to an attacker. The header contains source and destination IP addresses, which are required for

routing but may be spoofed or altered in what are known as "man-in-the-middle" attacks; the payload consists of information which may be confidential to a particular organization. IPSec provides mechanisms to protect both header and payload data. The IPSec Authentication Header (AH) digitally signs the outbound packet, both data payload and headers, with a hash value appended to the packet, verifying the identity of the source and destination machines and the integrity of the payload. The IPSec Encapsulating Security Payload (ESP) guarantees the integrity and confidentiality of the data in the original message by combining a secure hash and encryption of either the original payload by itself, or the headers and payload of the original packet.

How does IPSec work with network address translation (NAT)?

NAT is incompatible with Authentication Header protocol, whether used in transport or tunnel mode. An IPSec VPN using AH protocol digitally signs the outbound packet, both data payload and headers, with a hash value appended to the packet. When using AH protocol, packet contents (the data payload) are not encrypted.

Why this bothers NAT is the last part: a NAT device in between the IPSec endpoints will rewrite either the source or destination address with one of its own choosing. The VPN device at the receiving end will verify the integrity of the incoming packet by computing its own hash value, and will complain that the hash value appended to the received packet doesn't match. The VPN device at the receiving end doesn't know about the NAT in the middle, so it assumes that the data has been altered for nefarious purposes.

IPSec using Encapsulating Security Payload in *tunnel* mode encapsulates the entire original packet (including headers) in a new IP packet. The new IP packet's source address is the outbound address of the sending VPN gateway, and its destination address is the inbound address of the VPN device at the receiving end. When using ESP protocol with authentication, the packet contents (in this case, the entire original packet) are encrypted. The encrypted contents, but not the new headers, are signed with a hash value appended to the packet.

This mode (tunnel mode ESP with authentication) is compatible with NAT, because integrity checks are performed over the combination of the "original header plus original payload," which is unchanged by a NAT device. Transport mode ESP with authentication is also compatible with NAT, but is not often used by it. Since the hash is computed only over the original payload, original headers may be rewritten.

In addition, NAT may interfere with IPSec (both ESP and AH) if it prevents the two VPN gateways from successfully negotiating SAs using ISAKMP/IKE with certificates. X.509

certificates are signed by a trusted third party (called a Certificate Authority) in order to bind a user's or device's public key to some other identifying public characteristic. Once common identifying characteristic used for VPN gateway devices is external IP address.

If the two VPN gateways exchange signed certificates that bind each gateway's identity to its IP address, NAT address rewriting will cause IKE negotiation to fail.

Contributed by Craig Biggerstaff. [craig@blkbox.com]

Altiga (Cisco) has a revised version of their IPSec client that allows connections through NAT. They term it "IPSec over UDP". This may be somewhat less secure than true IPSec. Native IPsec requires that there be no change to the headers and NAT obviously breaks that rule. But Altiga's IPSec over UDP addresses a very practical concern - that many users of broadband Internet access are starting to be required by their ISPs to use NAT. PPTP seems to depend on who's NAT you are passing through. Seems to work fine going through say a SonicWall firewall wall. But PPTP going out through a Cisco NAT'ed router or firewall does not appear to work (if you are NAT'ing to a single routable IP address - hence overloading/PAT'ing.)

Contributed by Michael Medwid . [Michael.Medwid@ariba.com]

How do I decide whether a VPN is a good remote access solution for my organization?

To determine whether or not a VPN is a good answer to your company's needs for remote connectivity, consider your specific technical requirements, along with the pros and cons of VPN use.

Some advantages to using VPNs include:

- The ability to securely connect high-speed remote users over broadband technology, including Cable Modems and DSL lines, which before VPN's had not been possible. VPN's will work with any last mile technology as long as IP is run over the connection.
- No administrative headache for managing direct access telephone lines, T1 or PRI lines used for data, or for the RAS equipment (modems or other network access servers) terminating the phone calls
- Potential cost savings, especially if many of your remote users are located outside your local calling area.

Some disadvantages include:

- Potentially lower bandwidth available to remote users over a VPN connection as compared to a direct dial-in line.
- Inconsistent remote access performance due to changes in Internet connectivity. (To counteract this, you can have your users choose Service Providers that have higher levels of service, perhaps the same ISP from which you purchase your corporate Internet connection to keep traffic inside the same backbone.)
- No entrance in to the network if the Internet connection is broken (Some administrators choose to leave a limited amount of dial-in access for emergency access)

Contributed by Pete Davis. [pdavis@altiga.com]

How do I control who uses a VPN to access my site?

By definition, a VPN generally requires configuration of some sort of access device, either software or hardware-based, to setup a secure channel using private encryption and security parameters. A casual user can't just "use" my VPN, since some knowledge is required to allow the remote user or site access to my network (or even to begin a VPN handshake!).
Allowing VPN access only in conjunction with strong authentication also prevents an intruder from successfully authenticating to your network, even if they somehow configured/captured a VPN session.

Contributed by Matt McClung. [mmcclung@ndwcorp.com]

How do I control which sorts of network traffic are transmitted over my VPN?

Depending on the VPN solution being implemented, there are a few ways to control the type of traffic sent over a VPN session. Many VPN devices allow you to define a user or group based filter, which can control IP address and protocol/port services allowed through a tunnel. In addition, IPSec-based VPNs allow you to define a list of networks to which traffic can be passed (Security Associations). The first mechanism allows the administrator to limit access to specific networks/machines and applications on her network. The second usually provides fully connectivity to the private network.

Contributed by Pete Davis. [pdavis@altiga.com]

Should my VPN be terminated on my firewall, or directly on my private network?

There is no correct answer to this question, since it depends on your security requirements and network architecture. Two of the most common configurations for a VPN device providing Corporate Remote Access are to run a VPN device in parallel to an existing firewall, or behind an existing firewall. Terminating VPN sessions in front of a firewall or on a firewall itself is not as popular.

There are pros and cons for all implementations.

- Placing a VPN device in parallel to an existing firewall requires no changes to an existing firewall infrastructure, but also means that you will have two entry points into your private network. On most VPN devices, verify that they block all non-VPN traffic to minimize the additional security risk. Depending on how your network is set up, this will probably also require the VPN device to do some sort of Address Translation, or to have the ability to redirect this traffic to an existing firewall.
- By placing a VPN device behind an existing firewall, you will be required to make changes to your firewall. You will also need a firewall smart enough to be able to configure a filter to pass the VPN traffic. Depending on how your network is set up, this may also allow you to make use of only one of the two or more Ethernet ports on your VPN device. This configuration is sometimes known as One-Arm-Routing.
- By placing a VPN device in front of your firewall, you will be terminating secure traffic in a public zone. You will need to assign addresses to users from a certain block of IP addresses and open a large hole in the firewall for access from these IP addresses. A potential advantage to doing this would be that you could then use your existing firewall to control the destination of traffic, but most VPN boxes will also allow you to do this. This type of application may make more sense for trading partner connectivity vs. remote access users.
- By doing VPN on an existing firewall, you add some intense processing to a device whose original purpose was simply speaking, to control network access. Some people like the simplicity of adding a service to an existing device on the network perimeter.

Contributed by Pete Davis. [pdavis@altiga.com]

How does the use of encryption affect the performance of a network connection?

The use of encryption adds some additional overhead to a session. Most VPN devices, whether hardware or software based, will be able to process encryption for connections up to 10baseT speeds. On a lower speed connection like a modem, VPN processing is much faster than delays introduced by the limited bandwidth availability. Often performance is potentially affected

more by packet loss and latency on bad Internet connections than by the encryption overhead.

Contributed by Pete Davis. [pdavis@altiga.com]

What is strong authentication?

In general, user authentication is based on the following principle: an entity has authenticating knowledge (what you know), possession of an authenticating device (what you have), or exhibits an required characteristic (what you are). Strong authentication requires that at least two of the three principles be demonstrated.

Contributed by Steve Acheson. [satch@cisco.com]

Why does key length matter? What is a brute force attack?

Encryption systems depend on two mechanisms to guarantee data confidentiality. The encryption algorithm provides the mathematical "rules" that convert the plain text message to a random ciphertext message. The algorithm provides steps for convolving the plain text message with an "encryption key," a block of (typically) alphanumeric data that introduces the random element into the ciphertext message. The longer the secret key is, the more time it takes for an attacker to test all possible values of the key - and determine the plain text content of the message. In other words, data that will be valuable to an attacker for a long time should be encrypted with longer keys than ephemeral data. This sort of attack is called a brute force attack.

References:

General Articles:

VPN How-to
[http://www.ibiblio.org/pub/Linux/docs/HOWTO/other-formats/html_single/VPN-HOWTO.html]

VPN-Masquerade-HOWTO.html
[http://www.ibiblio.org/pub/Linux/docs/HOWTO/other-formats/html_single/VPN-Masquerade-HOWTO.html]

Comparing IPSec and PPTP
[http://www.intranetjournal.com/foundation/tunneling.shtml]

VPN Mini How-To
[http://www.ibiblio.org/pub/Linux/docs/HOWTO/mini/other-formats/html_single/VPN.html]

Microsoft Support:

Virtual Private Network Home Page of Microsoft
[http://www.microsoft.com/windows2000/technologies/communications/vpn/default.asp]

Virtual Private Networks White Papers
[http://www.microsoft.com/serviceproviders/vpn_ras/default.asp]

Overview of Virtual Private Networks
[http://www.microsoft.com/technet/treeview/default.asp?url=/TechNet/prodtechnol/windows2000serv/reskit/intnetwk/part2/intch09.asp]

Installing and Configuring PPTP on a PPTP Server
[http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnras/html/instpptp.asp]

Step-by-Step guide to Internet Protocol Security
[http://www.microsoft.com/windows2000/techinfo/planning/security/ipsecsteps.asp]

Understanding PPTP [Windows NT 4.0]
[http://www.microsoft.com/technet/treeview/default.asp?url=/TechNet/prodtechnol/winntas/plan/pptpudst.asp]

Point-to-Point Tunneling Protocol [PPTP] FAQ
[http://www.microsoft.com/ntserver/ProductInfo/faqs/PPTPfaq.asp]

Linux Sites:

Linux Implementation of IPSec
[http://www.freeswan.org/]

Linux VPN Masquerade
[http://www.impsec.org/linux/masquerade/ip_masq_vpn.html]

PPTP Server for Linux
[http://poptop.lineo.com/]

PPTP Client for Linux
[http://cag.lcs.mit.edu/~cananian/Projects/PPTP/]

A Power Point Presentation on Virtual Private Networks
[http://bmrc.berkeley.edu/people/chaffee/advnet98/nat_vpn.ppt]

VPN bibliography of Docs and Published Articles
[http://www.infosyssec.net/infosyssec/secvpn1.htm]

Virtual Private Network Resources
[http://thewhir.com/vpn]

Tina Birds collection of Links on VPN
[http://kubarb.phsx.ukans.edu/~tbird/vpn/]

CIPE Home Page
[http://sites.inka.de/~bigred/devel/cipe.html]