Global Information Assurance Certification Paper

GIAC CERTIFICATION
SANS SECURITY ESSENTIALS
GSEC PRACTICAL ASSIGNMENT
Version 1.2f

VPN : Another Network Security Solution

BY :

MD. AZZAD MD GHAUTH

"THIS REPORT IS PREPARED AS A PREREQUISITE OF OBTAINING GIAC SECURITY ESSENTIALS CERTIFICATION"

Table of Contents

VPN - Another Network Security Solution

Introduction

For years, voice and data services were delivered using what the telephone companies called virtual private networks. In fact, the phone companies consider just about all software-defined networks VPNs.

But the current generation of VPNs is very different. The definition of VPN is a combination of tunneling, encryption, authentication, and access control technologies and services used to carry traffic over the Internet, a managed IP network or a provider's backbone.

The traffic reaches these backbones using any combination of access technologies, including T1, frame relay, ISDN, ATM or simple dial access. The term virtual private network refers to a private connection between two computers to transport private data over a shared or public network – often the Internet. The term VPN commonly refers to a point-to-point networking configuration that offers encryption, tunneling, authentication, and access control. Because transmitting sensitive data over the Internet is inherently unsafe, encryption hardware and firewall equipment are usually included in a VPN, plus a gateway and/or router. Many vendors are offering some combination of these devices and services in one box. Often VPNs replace the expensive leased lines that connect the remote sites of an enterprise, so VPN users may need broadband line speeds. Remote or mobile users are the major considerations in planning most VPNs. The performance they receive is unpredictable, since they're using modem connections from all sorts of providers, and they need reliable, secure access to the VPN at all hours.

The general idea behind using a VPN is that a company or an organization reduces the recurring telecommunications charges that are incurred when connecting remote users and branch offices to resources in corporate headquarters.

When it comes to remote-access VPNs, the basic concept is to give telecommuters and mobile workers a way to get back to a corporate network over the Internet or a service provider's backbone. In a remote-access VPN, a user dials into a service provider's point of presence (POP), establishes a tunnel back to headquarters over that provider's network or the Internet, and authenticates himself or herself to gain access to the corporate network. That is in contrast to the traditional dial-access approach whereby a user dials into a bank of modems, a remote-access server or concentrator located within the corporate headquarters.

Another use of VPN is for the site-to-site connectivity or in other word Intranet VPN. Branch offices are connected to corporate headquarters through tunnels that transport traffic over the Internet or via a provider's backbone.

Another use of VPN is extranets. The basic idea of VPN-based extranets is to use the access control and authentication services with a VPN implementation to deny or grant customers or business partner's access to specific information that is important to do business. With VPN-based extranet application, the outside party would get to the corporate Firewall by tunneling across the Internet or a service provider's services.

Firewalls and VPNs

Many firewalls have some kind of VPNs — encrypted firewall-to-firewall tunnels. All traffic between one firewall and another is encrypted, stuck inside of another IP packet, and sent over the Internet. At the remote site, the firewall pulls the encrypted payload out of the IP packet and decrypts it to get the original IP packet, which is forwarded to the final destination.

Firewall-to-Firewall With Controlled Access

As VPNs become more widespread on the Internet, and VPN establishment more automatic, most VPNs will be for privacy between sites, without a complete trust relationship between those sites. By "complete trust relationship," we mean the relationship that exists between business partners, suppliers and consumers, providers and customers. Privacy for the communications may be desirable or needed, but an Internet firewall can be used to control or prohibit access to the internal, private network.
As an example let us say ABC Company has 10 resellers. ABC could set up VPNs between ABC's Internet gateway and each of those 10 resellers. All e-mail, file transfers, everything, would go over an encrypted channel. Any TCP/IP traffic to other systems on the Internet would be sent, as usual, in the clear. Further, limited access will be allowed through the corporate firewall to selected servers inside the network security perimeter for the resellers.

Firewall-to-Firewall With Open Access

A common configuration for VPNs between Internet Firewalls currently is in a trust relationship between offices of the same company. For example let's look at the ABC corporate network.
ABC has two offices in Kuala Lumpur, one in Singapore, two in Manila, one in London, and one in Jakarta. The Internet is used as the ABC corporate backbone. Without a VPN, all traffic between ABC offices would be vulnerable to disclosure as it flowed over the Internet. Further, using less secure TCP/IP services (such as file and directory sharing) would leave the ABC corporate network vulnerable to attack using those services.

With VPN connectivity using VPN-enabled Internet Firewalls between the sites, all traffic is encrypted, and so private. Additionally, when the trust is added where they are derived from all sites being administered by the same organization, all having the same security policy implemented, and all being under the same management organization, we can, over this VPN, allow all network services. In this way the flow is virtually going around the firewall, though actually the communications flow is still under the protection

of the firewalls. In this way it is possible to extend the network security perimeter to include those other offices. All sites are now virtually on the same LAN, with a virtual network perimeter.

Firewall-to-Remote System

The same sort of VPN technology can be used between a firewall and a single site. Often this is used to allow private access to a corporate from mobile users, working at a customer site, or connected in from home or a hotel. As with firewall-to-firewall VPNs, these can be set up with controlled or open access. Controlled access is useful for clients, customers, and partners needing access to particular systems on the inside of the security perimeter for particular services at particular hours of the day. Open access is useful for employees on the road who need to get access to shared files, printers, etc. on the inside of the network security perimeter. With a VPN they can do this securely.

The Need for Standardization

Many commercial firewalls today support VPN functionality. Two VPN-enabled firewalls can be used to establish a VPN. When this technology was first developed, there were no standards. Consequently, vendors established their own ways of implementing IP encryption. Some vendors went with a mechanism called swIPe (for Software IP Encryption). This was freely available in source code form with no licensing restrictions, so it was attractive because 1) it was already written (although it needed to be ported to platforms other than SunOS) and 2) it was the only mechanism approaching a standard, albeit "de facto," that existed. Still, one vendor's VPN using swIPe didn't work with others and it was not possible to establish an inter-vendor VPN.
In order for VPNs to become widely and routinely used, sites using differing encryption products need to be able to communicate.

There are currently three major tunneling protocols for VPNs. They include the Point-to-Point Tunneling Protocol (PPTP), Internet Protocol Security (IPSec) and Layer 2 Tunneling Protocol (L2TP).

The major tunnel protocol used for the enterprise VPN is IPSec, while L2TP is typically utilized by SPs to provide remote dialup VPN access for customers. The benefit of L2TP remote access is that it uses PPP for encapsulation and does not require installation of an extra package on the remote client.

IPSec (RFC 2401) with additional RFCs 2402/2412 for codifying the essential components has become the de facto industry standard for an IP-based VPN infrastructure. IPSec supports two different topologies in IPv4: client implementation and gateway implementation. The client implementation is referred to as the bump in the stack (BITS) because it is inserted between the IP stack and the local network drivers. Since source code access for the IP stack is not required, this implementation approach is appropriate for use with legacy systems. It can be installed in Windows 95, 98, and NT; Macintosh, and Unix operating systems, depending on the specific vendor

implementation. When Windows 2000 is released, it will have Microsoft's implementation of IPSec built in. It will be applicable for remote access or for supplying a secure network within a LAN, such as in a IS department with multiple personnel spread throughout a large enterprise network and a secure communication requirement for a client/server application. In this case, one would implement BITS on every client and server.

The gateway implementation is referred to as the bump in the wire (BITW) because there is a piece of equipment running IPSec on the edge of the network being protected. The gateway implementation may be in software or hardware. An example of a software implementation is IPSec running on a firewall or router. An example of a hardware implementation is an appliance running IPSec in silicon (black box), which is much faster because of the lack of OS overhead. A VPN requires at least one gateway implementation at the central site or the remote LAN, and a client implementation for each remote site (road warrior or single PC site).

The IPSec protocol provides tunnel authentication and encryption. Since IPSec was designed to be algorithm-independent, several acceptable options of encryption and authentication algorithms are supported, allowing an optimal choice to be made for a specific VPN's required security level. For authentication, the supported algorithms are HMAC MD5 (normal usage) and for a higher level of security authentication, it is HMAC SHA1 (FIPS-180-1). Encryption algorithms are DES, 3DES, RC5, Cast, Idea, and Blowfish, with others being added. DES and 3DES are mandatory, but not all vendors support additional encryption algorithms, so if a higher level of encryption is required the individual vendors' products must be researched.

IPSec was designed so that once an encryption algorithm is chosen, the actual keying material used by the algorithm can be regenerated inside the tunnel in a requested time frame (once every hour, each 10Mbs, etc.) to change the session keys, effectively changing the encryption. Therefore, even though DES has been broken in less than 23 hours, it is still usually acceptable for use if the session key is automatically changed every hour. Even if the tunnel data is intercepted and subjected to a brute force attack, only one hour of data would be compromised.

The IPSEC architecture includes authentication (how to know if the site communicating to the secured site really is who it claims to be) and encryption. These mechanisms can be used together or independently.

Authentication Header

As described in RFC 1826,
The Authentication Header is a mechanism for providing strong integrity and authentication for IP datagrams. It might also provide non-repudiation, depending on which cryptographic algorithm is used and how keying is performed. For example, use of an asymmetric digital signature algorithm, such as RSA, could provide non-repudiation. The IP Authentication Header seeks to provide security by adding authentication

information to an IP datagram. This authentication information is calculated using all of the fields in the IP datagram (including not only the IP Header but also other headers and the user data), which do not change in transit. Fields or options which need to change in transit (e.g., "hop count", "time to live", "ident", "fragment offset", or "routing pointer") are considered to be zero for the calculation of the authentication data. This provides significantly more security than is currently present in IPv4 and might be sufficient for the needs of many users.

Encapsulating Security Payload (ESP)

As described in RFC 1847,
ESP is a mechanism for providing integrity and confidentiality to IP datagrams. In some circumstances it can also provide authentication to IP datagrams. The mechanism works with both IPv4 and IPv6. ESP is a mechanism for providing integrity and confidentiality to IP datagrams. It may also provide authentication, depending on which algorithm and algorithm mode are used. The IP Authentication Header may be used in conjunction with ESP to provide authentication. Users desiring integrity and authentication without confidentiality should use the IP Authentication Header (AH) instead of ESP.

ESP in IPSEC can be used to encrypt the entire IP datagram or just information at the transport layer (TCP, UDP, etc.).

Again quoting from the RFC,

In Tunnel-mode ESP, the original IP datagram is placed in the encrypted portion of the Encapsulating Security Payload and that entire ESP frame is placed within a datagram having unencrypted IP headers. The information in the unencrypted IP headers is used to route the secure datagram from origin to destination. An unencrypted IP Routing Header might be included between the IP Header and the Encapsulating Security Payload. In Transport-mode ESP, the ESP header is inserted into the IP datagram immediately prior to the transport-layer protocol header (e.g., TCP, UDP, or ICMP). In this mode bandwidth is conserved because there are no encrypted IP headers or IP options.

Different implementations of IPSec offer a variety of authentication methods, including shared secret, token cards or digital certificates. Shared secret is fairly easy to utilize for a small number of endpoints (clients and/or gateways). Token cards work very well for large intranet implementations, but for a large extranet implementation the easiest method is to use a digital certificate (Public Key Infrastructure).

A Public Key Infrastructure (PKI)starts with a Certificate Authority (CA), a software package operated in a high security area, that issues digital certificates. It also includes a directory service to make the certificate widely available. When implementing a PKI, you can either contract for the service or implement the operation in house. The decision to contract or purchase must be based not only on cost, but also more importantly on security policy and requirements. Do you have full control of the PKI or do you let some one else operate it for you?

Besides a CA, a PKI also includes, at minimum, a X.509v3-compatible database. The CA operator issues the digital certificates to the end entity (EE)—IPSec endpoints in this case—and records the information in the database. When a certificate is either compromised or is no longer correct for some reason, it is listed by the CA operator on a Certificate Revocation List (CRL). Each time an IPSec endpoint checks the validity of a certificate that is presented for authentication, it checks the CRLs to see if that certificate is listed. If the certificate is listed, it is no longer valid and the IPSec endpoint will reject it.

Written Certificate Policy (CP) is needed regardless of whether CA is operated or outsourced. The CP delineates the requirements (for example, a certificate must be requested in person and requires two forms of ID, one a picture ID) to receive a certificate from the CA and/or a level of authority (for example, this certificate allows signature authority for one million dollars). For an IPSec endpoint, the CP defines what information must be submitted to the CA for certification. The CP should also specify the requirements that the CA must meet for security reasons.

To successfully implement a CA, the operator of the CA must write a Certificate Practice Statement (CPS), which responds to the CP spelling out how the operation of the CA matches the CP requirements. They are important for two reasons. First, it ensures optimal security by requiring written documents for both operational guidance and for audit purposes. Second, if there are need to cross-certify (that is, be treated as an equal and able to accept certificates) with a CA operated by someone else, both the CP and CPS are required to ensure that both certificates are considered equal in the required aspects.

Key Management: The Challenge and the Emerging Standards

Key exchange mechanisms — how to decide on and communicate what cryptographic keys to use to secure the cryptographic communication — and keeping that exchange secure (private) is critical to private communications and the establishment of VPNs. Most implementations of VPNs used a manual key exchange mechanism with out of band communication. This often meant, software run manually to come up with a random key for communication, communicated over encrypted e-mail, the telephone, or via floppy disk. It should be clear that this mechanism does not scale. There are frameworks and mechanisms being developed to do this work. The IPSEC effort is focusing on ISAKMP and Oakley.

ISAKMP

The Internet Security Association & Key Management Protocol (ISAKMP) provides a framework for Internet key management. It also defines the protocol for the negotiation of security attributes (algorithm to use, length of key).

Oakley

The Oakley Key Determination Protocol uses a hybrid Diffie-Hellman (a public key) technique to establish session keys.
ISAKMP and Oakley Combined
As stated above, both private key exchange and a mechanism for negotiating keys, algorithms, etc. are required for automatic establishment of VPNs. Combining ISAKMP and Oakley provides this.

Oakley defines a method to establish an authenticated key exchange. This includes how payloads are constructed, the information they carry, and the order in which they are processed and how they are used. While Oakley defines "modes", ISAKMP defines "phases". The relationship between the two is very straightforward. ISAKMP's phase 1 is where the two ISAKMP peers establish a secure, authenticated channel with which to communicate. This is called the ISAKMP Security Association (SA). Phase 2 is where Security Associations are negotiated on behalf of services such as IPSEC or any other service, which needs key material and/or parameter negotiation.

Challenges of VPN Deployment

There are two areas that make VPN deployment challenging : They are

1. Security
2. Performance

Good VPN performance results from the interaction of many complex network devices. How the devices behave in concert is usually unknown. We rarely know the answers to the following questions before starting deployment :

- will network devices from different vendors work well with each other ?
- Are performance levels likely to be sufficient for remote users and offices, especially when compared to what local employees are getting ?
- what's the maximum number of simultaneous tunnels that can connect to a VPN gateway ?
- if aVPN is being deployed, is special hardware needed for acceleration or QoS ?
- Can users run the same networked applications as before and still get acceptable performance ?
- If VPN service is being purchased, is the service provider giving the level of service that was in the contract ?

Reports from network testing labs strongly indicate that a VPN deployment is likely to challenge the reliability and performance of the data networks we are using. Deploying a VPN usually causes some or all of the following :

1. Increased latency
2. Increased congestion

3. Decreased throughput
4. More failed connections
5. More lost data in streaming, multimedia traffic.

Networks evolve over time. New users and new networked applications change the balance and disrupt performance. VPNs radically change the way people access and use networks, so you want to make sure that the networks perform at least as well after VPNs are deployed as they did before the changes were made. I would only discuss performance issues as I think performance should be balanced accordingly as the security portion.

General Performance

VPNs have some inherent attributes that could obviously impact performance. Network traffic that would ordinarily travel over a corporate backbone is modified when sent over a VPN connection. To start, VPNs use encryption and tunneling, a combination that could affect performance in several ways.

First of all, VPNs require processing power to handle encryption and tunneling. Some managers are concerned that, for example, a telecommuter's workstation might not have the processing power to handle these tasks.

Early VPN implementors have not found this to be the case. Telecommuters with a Pentium-class PC who dial into a service provider's network using ISDN or analog modem lines seem to have more than enough processing power to handle even IP Security (IPSec) tunneling and encryption using long keys.

Still, there is the potential for a performance problem. If the user is running a processor-intensive application such as SAP, the application might slow down under the load of the VPN. However, it should be noted that most early users report no noticeable performance differences when running most common business applications.

In another area, there might be performance problems when telecommuters use high-speed services. It's one thing to encrypt and tunnel at data rates of 56 or 128 kilobits per second; it's quite another if a telecommuter is using a cable modem or DSL service that operates at 1 or 2 megabits per second or higher.

Whether a PC can handle the encryption and tunneling tasks at these data rates while also running other applications will depend on the processing demands of the specific application. To avoid problems, some early VPN adopters have chosen to use a hardware-based VPN device in a remote office when using these high-speed services.

Another performance issue unique to VPNs: the extra burden on devices as a result of excessive packet fragmentation, which can stress routers and impact VPN performance.

Packet fragmentation results from the overhead added to packets by encryption and tunneling protocols. A single packet on a LAN can sometimes fracture into as many as three packets.

In some cases, VPNs use triple-layer enveloping, where packets are inserted within packets inside another packet.

For example, an IPX packet traveling on a LAN could be inserted into an IP packet for travel over the Internet.

Often, a Point-to-Point Protocol header is added to the IP packet. A tunneling protocol, such as PPTP or IPSec, requires that another header be added to that packet. And then, if you layer on encryption, you can easily more than double each packet's nonuseful information.

Once a packet exceeds the preset packet-size limitations, it splits into two separate packets. To send them over a VPN link would essentially require starting the process over again. The end result is that you have many more packets traveling between sites using a VPN than if the traffic were just running over an IP backbone.

It is difficult to say exactly what this additional burden will be. Some early tests have found that as many as 30 percent more packets could be created as a result of this fragmentation process. In that case, existing internetworking devices might suffer performance problems.

VPNs share the same general reliability issues as any other type of network connectivity scenario. First there is the reliability of the device handling the VPN services. Just as is the case when connecting users or branch offices using the public switched telephone network, the device within the company walls that terminates the service must have a level of reliability that is appropriate for the level and type of traffic that it will be supporting.

If the VPN is primarily used to give remote users dial access to their e-mail at night, then Uninterrupted Power Supply is needed to keep a VPN device running.

However, if the majority of a company's business runs off of the VPN, redundant power supplies and cooling fans is important. If the device is modular, the ability to hot-swap a failed component without having to take the entire unit offline to replace the defective part is an advantage.

Such features have been common on traditional internetworking gear for a long time. Most high-end routers, servers and hubs already have these features. If VPNs are going to carry a large portion of your traffic, redundancy is needed in the device that establishes and terminates VPN tunnels, and performs the other tasks associated with a VPN.

Today, with many types of networking equipment, clustering is used to improve the chance that users will have access to critical data or services. For instance, a large number of companies use clustered database or file servers to provide an extra layer of protection when it comes to giving users access to those particular servers.

To that end, some VPN vendors have recently started offering a way to essentially cluster VPN devices. The clustering takes place at two levels. First, there is a fail-over capability, where one VPN device takes over if the primary device fails. This feature is supported in a few VPN products today and will increasingly be offered in the future.

VPN device clustering can also provide a second type of reliability-load sharing between two or more VPN devices. This feature helps in a couple of ways. First, the load-sharing services will frequently deal with a VPN device failure by moving whatever sessions the downed device was handling over to a working unit.

Second, load sharing better ensures that a user will be able to establish a tunnel. Often the greatest stress on a VPN device occurs when there are many requests to establish simultaneous tunnels. Typical examples include a large number of telecommuters all signing on at the same time in the morning, or numerous sites all being brought up at the same time on a Monday morning.

Without load sharing, a user could have trouble establishing a tunnel when a VPN device approaches its maximum capacity. Load sharing alleviates this problem. For example, there might be dual T1 lines between headquarters and a large regional office. In that way, if one line was cut or down for some reason, the traffic could be sent over the other. In some cases, the two lines were supplied by different service providers to add an extra level of assurance that traffic could flow between the sites.

With the growing popularity of public switched services, other approaches took hold. For example, it is quite common for a site to have a primary connect be a frame relay link with an ISDN dial backup line in reserve. In that way, the ISDN line is only brought up when there is a problem with the frame relay connection. Or, in some cases, the ISDN line is used to augment the bandwidth of the frame relay connection when there is a traffic surge.

If a service provider link goes down or the provider's network is not available, the VPN device will automatically recognize this and perform some action. The device could, for example, use a second provider's access service, or else the VPN device could dial up an ISDN connection to link to a VPN device in another office.

For VPNs to truly be used to carry a large volume of corporate data, they will eventually need to support these types of functions.

One way is to look for products that support clustering of VPN devices. As discussed above, some vendors let you increase support for more sessions by adding another unit to the network.

Another method that is becoming popular is the use of modular VPN devices. Some vendors are offering chassis-based VPN systems that allow a manager to start with a small VPN implementation and then scale up.

The typical approach to using these devices is to have modular cards that can be added to a chassis' empty slots. Each additional card has more processing power to handle applications such as encryption and tunneling.

With either the clustering or chassis-based approach, the problem that is being avoided is obsolescence. A fixed-capacity VPN device might be perfect when only a few users are connected, but would serve no purpose when more users are added. If there is no way to either expand the capacity of the device or cluster it with other VPN devices, the unit may need to be retired.

Conclusion

All of these need to be coordinated and communicated. The technologies that exist today, and are being implemented in firewall-based VPNs, provide these mechanisms and will allow a growth in deployment and use of VPNs extensively in the near future.

Bibliography

Virtual Private Network: How They Work, Roger W. Younglove, CISSP, Senior Network Systems Consultant, NetworkCare Security Services, Lucent Technologies.
URL : www.itpapers.com

VPN: The Basics, Does Everybody Know What A VPN Is ?, Salvatore Salamone
URL : www.itpapers.com

IT Managers Seek Answers To VPN Performance Queries, Salvatore Salamone
URL : www.itpapers.com

Testing and Monitoring Virtual Private Networks (VPNs) White Paper, Susan M. Pearsall and John Q. Walker, NetIQ.
URL: www.netiq.com

VPN Frequently Ask Questions
URL : www.internetweek.com/VPN/faq.html

Firewalls and Virtual Private Networks, Frederick Avolio, Trusted Information Systems, Inc.
URL : www.itpapers.com