2

Basic Virtual Private Network Deployment

Before discussing the features of Windows 2000 tunneling technology, it is important to establish the terminology that one should be familiar with. The terminology is not specific to Windows 2000 and can be applied to almost any VPN-related product. After defining the terminology this book uses, this chapter discusses one all-important question: Why deploy a Virtual Private Network? It is important to first understand the needs of your environment and then decide whether tunneling will fulfill those needs. This chapter also covers many common attacks that occur over networks to help you understand why it is important to protect your servers. Finally, it covers basic tunnel network designs.

Terminology

The first step is to define some VPN terminology. You should be familiar with the following terms:

- **VPN server (also known as a tunnel server).** A computer that accepts VPN connections from VPN clients. A VPN server can provide remote access VPN connections or a router-to-router (site-to-site) VPN connection. It is the VPN server that is connected to the public network. This book primarily refers to Windows 2000 as the tunnel server, but there are many other types of tunnel servers.
- **VPN client (also known as a tunnel client.)** A computer that initiates a VPN connection to a VPN server. A VPN client can be an individual computer that obtains a remote access VPN connection or a router that obtains a router-to-router VPN connection. This book primarily covers Windows 2000, Windows 98, and Windows NT 4 as VPN clients.

- **Tunnel.** The logical link between the tunnel client and the tunnel server. This link is where the data is encrypted and encapsulated. It is possible to create a tunnel and send the data through the tunnel without encryption, but that is not a recommended VPN connection type because the data being sent can be intercepted and read.
- **Edge server.** This tunnel server is the outermost server on the company's private network. Typically, anything "behind" this server (on the corporate network) is "open frame" traffic and can be readily intercepted. If frames are captured on the private network, the security of the traffic is compromised, even though the network is using a tunnel to the edge server. This scenario does not, therefore, have end-to-end security. An edge server can be a firewall, or it can be a specific system that does nothing but handle tunnel traffic.
- **End-to-end security.** A path that is encrypted from the client all the way to the actual destination server has end-to-end security. Because the technology needed for a practical implementation of end-to-end security has just been released, most designs currently use a specific tunnel server on the edge of the corporate network. If you have complete security, it will not matter if frames are captured anywhere in the path because they maintain their encryption at all points in their journey. At this time, most designs use a specific tunnel server on the edge of the corporate network and have encryption only between the client and the tunnel server.
- **Voluntary tunnel.** A user or client computer can issue a VPN request to configure and create a voluntary tunnel. In this case, the user's computer is a tunnel endpoint and acts as the tunnel client. The client must have the appropriate tunnel protocol installed. Many network designs require this because the corporate networks do not generally control home LANs, and having the tunnel clients as the actual endpoints reduces the potential security risks.
- **Compulsory tunnel.** A tunnel configured and created by a VPN-capable dial-up access server. With a compulsory tunnel, the user's computer is not a tunnel endpoint. Another device, the remote access server, between the user's computer and the tunnel server is the tunnel endpoint, which acts as the tunnel client. This configuration allows multiple clients on the branch office or home LAN to use the tunnel concurrently. It is possible to share a single tunnel to multiple computers.

Design Considerations

Before you roll out a VPN solution, you must be very aware of an obvious question that many administrators fail to ask themselves: Why am I rolling out this VPN? There are many, many reasons to roll out a VPN, the most popular of which are detailed in the following sections. As you consider VPN deployment, you should recognize the benefits listed here and see if they answer the question of why you should roll out a VPN.

Network Access Is Too Expensive

Given the increasing costs associated with hardware, software, and upgrade costs, companies are looking for solutions that provide technological advances while impacting the least on the bottom line cost to their business. Compared to the alternatives, using VPNs saves a corporation a great deal of money. A remote user can dial a local connection to the Internet from any location and then use the connection to establish a link to corporate resources. This enables the company to outsource the investment of modem banks and the continued cost of maintaining and upgrading the modems. Additionally, this avoids the usage costs of long-distance phone links by remote users.

Branch offices benefit tremendously from using VPNs because the tunnel link provides a link to the corporate servers and resources in both directions without expensive leased lines and equipment. Another benefit is the ability to set up branch offices quickly and effectively because links are not dependent on any particular line provider. The setup time is often cut from 6–12 weeks for a circuit to as little as 24 hours for dialup connections or 2–3 weeks for DSL or cable modems. Additionally, it is possible to get faster links to the Internet, resulting in faster branch office connectivity through the VPN. Another benefit is that if the needs of a branch office change, it is much cheaper to change the VPN connection than to go through all the changes needed with leased lines.

Most corporations are concerned about the security of their data at one level or another, and this issue influences network design because captured data can cost companies a tremendous amount of money. For example, if you ran a bank, you would certainly want to protect your data very closely. Likewise, a manufacturing company would want to protect product designs from its competitors.

There is no accurate gauge of how much money is at risk given the varying levels of protection you can employ for the data over your network links. The network designer must coordinate with the various teams within the corporation to segment and identify the critical data and/or servers. Then, with Windows 2000, he or she can define policies and security zones based on this design.

If a user in California wants to use his laptop computer to access his corporate network in New York, the only option for doing so without a VPN is to use a 1-800 number or, even worse, to make a direct, long-distance call. More and more people are doing this not only because they are traveling for business, but also because home offices are increasingly common. There are many reasons for this trend; for example, studies show that most home-office employees are more productive, and that this allows for more flexible hiring. Compaq, for example, has thousands of employees who work full-time from their homes. It is becoming more common to link such branch offices with VPNs because they have a number of unique requirements that the laptop users do not have.

To avoid long-distance calls completely, remote users can obtain local access from a global Internet service provider (ISP). Through the link to the Internet, they can establish a VPN to a corporate tunnel server. Most ISPs allow for VPN traffic, and although some ISPs use proprietary software, it is generally possible to configure clients for effective communication. Even consumer-oriented ISPs such as AOL can be configured to support VPN connections to corporate networks. As always, however, it is critical that you read the ISP contract to make sure they allow for this.

If a company has branch offices at various locations around the world, leased lines connecting those sites are very expensive. It is now possible to create a link to the Internet at each site, and through those links establish a VPN that provides the same services the leased lines did. This option is scalable and reconfigurable, and it functions independently from the ISP that is being used.

The previously described approach can also be used if a company wants to connect two disjointed networks to enable vendor access. For example, Compaq and Microsoft have a tunnel server that allows for connection to both networks in both directions. This allows a client that might be working for two corporations to effectively communicate with both companies through tunneling technologies. This can be accomplished with separate sessions or concurrently.

Many companies use VPN technologies instead of leased lines to link partner sites and networks. This has all the benefits of a leased line, yet it can dramatically reduce the cost of the connection. It is completely possible to allow all clients on both networks to connect to resources on either network.

Data Security Concerns

Security is an issue now more than ever, and Windows 2000 offers features that help alleviate the security issues within an infrastructure. A network environment involves many levels of security. When considering the complete security situation, you must look at passive network monitoring, authentication for network access, data modification, application-based security, and more. Because this book focuses primarily on network design, I will discuss network-oriented security and how it relates to the VPN design.

With Windows 2000, developers can create a secure route from source to destination and back. This enables security-conscious organizations to use Windows 2000 in ways that previously required custom hardware solutions or that were not even possible. This level of encryption is available from the integrated features of IPSec and the related policies configuring the connection parameters.

The ultimate goal of a secure network is end-to-end security so that if network traffic is captured at any point in the transmission, the corporation can be sure that traffic is secure. This type of encryption is not specific to a VPN. It can also be applied to LANs, WANs, and tunnels.

In the past, many corporations used parallel LANs to separate traffic, and in doing so, prevented people from easily eavesdropping on the separated LAN. This was accomplished by leaving no network ports open and available on the "secure LAN," which is typical with standard networks. This approach was very expensive because the corporations not only had to run two completely separate networks, but also had to have separate hardware and staff to maintain the two networks. Additionally, it is practically impossible to always guarantee that the secure LAN is completely secure because of all the advanced ways to eavesdrop on networks.

Corporations can now encrypt individual groups of traffic or resources through one policy while encrypting other resources through another policy. This creates *security zones*. A financial database, for example, would be in a high security zone, whereas a printer would be in a low security zone. Once the developer defines which resources are to be in which zones, he can create Active Directory configuration group policies to reflect this design.

With Windows 2000, an organization can implement an aggressive protection mechanism against attacks. It's probably unlikely that someone will spend a great deal of time tracking VPN traffic, but it's a good idea to consider potential security flaws in a network design and try to eliminate them.

The first step in implementing protection is to find out the company's network topology; in a large organization, this can be a big challenge. The developer needs to track down each segment, each firewall, each link to the Internet, and even unauthorized departmental links or servers.

If a developer is going to roll out an organizationally secure network environment, everybody in the corporation has to play along; otherwise, the developer can never guarantee a secure network. Even one user who has a phone line connection to his Windows 98 desktop could be running the Plus Pack dial-in server on his computer, which represents a serious security risk. The network administrator has many tools to help prevent this, and one of the most important is computer policies. Your security team usually defines computer policies for your organization. In some industries, regulatory agencies provide industry guidelines.

Tip

In my experience, the most important issue with ensuring that the network is secure is to fully understand the network design. This includes everything from the actual wires down to the policies affecting your clients. It is critical to document the configuration of the network so that traffic can be monitored and understood. Any unusual or unpredicted traffic can then be investigated. Documentation is the number one issue because this type of monitoring and maintenance will typically include a number of teams. ◆

Methods of Attack on Network Traffic

It seems that every few days we hear about yet another type of virus or attack compromising data and bringing down networks. In recent years, many products address the need to be proactive with attacks. These products can look for trends or signs of new attacks that can affect your network or servers. This approach is very different from a typical virus scanner that waits for the virus to attack a system before trying to clean that system.

Most available intrusion detection software looks for all types of known attacks. If it detects one, it notifies the administrator, or starts a routine that tries to protect the network resources, or both. Many of these intrusion packages are available on the market, but I believe it is generally best to go with a name-brand solution.

As a network administrator, you might face any of the following types of attacks:

- Denial-of-service attacks
- Address spoofing
- Session hijacking
- Sniffers
- Compromise key attacks
- Data modifications

- Man in the middle
- Replay attacks
- Brute force
- Password guessers and dictionary attacks
- Social attacks

Each of these is discussed in the following sections.

Denial-of-Service Attacks

The denial-of-service attack is the most common form of security breach. In a denial-of-service attack, the attacker floods a network interface with traffic to make the server so busy that it cannot answer requests. In another form of attack, the attacker sends specific invalid packets to a computer that can cause a computer's operating system to crash. This is not productive, and the attacker does not get information from the attack. He has only one aim—to prevent you from using your own equipment.

In the past, not much could be done to eliminate denial-of-service attacks. Most people tried to isolate the source of the attack and then configure their firewalls to block all traffic from that source. The obvious problem with this is that the attack would be successful until an administrator physically intervened. Windows 2000 allows the server to be configured in ways that dramatically reduce the chances of denial-of-service attacks. The server can be configured with filters and, better yet, IPSec policies that discard traffic defined as unnecessary.

No server that is available on a network is completely immune to denial-of-service attacks, but with a properly configured server, the risk of this type of attack can be reduced to a more palatable level.

Address Spoofing

Given the design of TCP/IP, attackers can "spoof" their target systems into thinking that packets originated from places they did not. A computer on the outside of your firewall can spoof a computer on the inside of your firewall, making your firewall believe that any related traffic is originating from inside the corporate network. This could then enable the rogue computer to access internal resources without being detected.

"Spoofing" is possible because the TCP/IP protocol suite is independent of the lower layers of the OSI network model. Due to the way packet routing works and how headers are constructed, it is virtually impossible to guarantee the genuine source of a packet. If traffic is sent from Point A to Point B, Point B simply assumes the traffic did in fact come from Point A.

Denial-of-service attacks can be performed by flooding servers with packets whose source addresses cannot be replied to. Spoofing the source IP address of the packets that are unreachable does this. The most common way is to spoof a private address, because when the server responds to the traffic, the server receives an ICMP Destination Unreachable message. An easy way to reduce the risk of this is to add to the server a filter that accepts all packets except when the source address is a standard private address such as 10.0.0.0, 172.16.0.0, or 192.168.0.0. For more information on setting up filters, see Chapter 10, "Routing and Filtering."

Session Hijacking

Session hijacking occurs when a session between Source A and Server B is intercepted and copied by an attacker. The attacker usually intercepts a TCP session between the two machines. Because the authentication typically occurs at the beginning of the session with non-encrypted traffic, it allows the attacker's system to participate in the conversation between the two systems.

An attacker can then configure a computer that identifies itself with an address of an actual computer or server that should be part of the conversation. It is possible for the attacker's computer to impersonate an email server in your corporation or to simply fool the valid systems into passing their IP packets to it.

During the time an impersonation is occurring, all clients attempting to send or receive email talk to the attacker rather than to the actual server. During that time, the attacker's system collects all information that is being sent to it. By the time the network administrator is alerted that the mail server is not functioning properly, the attacker's imposter system has probably been shut down, and the attacker has collected all data. This involves a simple capture of data.

Session hijacking can be prevented by encrypting the data traveling on the network. If captured data was analyzed, the attacker would have to decrypt the data to make it readable. Additionally, the original computers would discard traffic that was intercepted by an unknown system.

Sniffers, Lack of Privacy

A network sniffer is a device that captures all traffic going over the network. In the early days of networks, sniffers were very expensive and fairly rare, but now that Windows NT and Windows 2000 (even the Workstation) ship with a rudimentary version of a network sniffer, sniffers are very common. With Ethernet switching, the problem is reduced because each segment of the Ethernet network is isolated by port.

A network sniffer can see all traffic that travels over the local network. If the traffic is encrypted, the sniffer simply shows unreadable data. But if the data is open and non-encrypted, everything is readable. A network sniffer is one of the most powerful tools for diagnosing network problems, but it can be a powerful tool for an attacker as well.

Compromise Key Attack

Compromise key attacks occur when an attacker obtains the key used for the encryption and decryption of the data and then captures that data. Because the key has been compromised, the attacker can then decrypt and read the data. This is common when a set string is used for the key because the key does not change for long periods of time. The obvious way to avoid this attack is to use a dynamic key instead of a static one.

This type of attack was more common in the past, when keys were passed over the network through various unprotected methods. This type of attack was also common when the system was compromised and the attacker took the key from it. This is still possible, of course, if shared secrets are being used. It is up to the administrator of the system to guarantee that this key does not become available.

In a dynamically changing key environment, such as the Internet Key Exchange (IKE), it is much less likely that the key can be compromised because it is never manually handled. Additionally, the key changes periodically, usually based on either a time schedule or a specified amount of data, so even if the key is compromised, only part of the data is readable.

IKE uses an approach called Perfect Forward Secrecy (PFS), wherein the key exchange uses one long-term key and generates short-term keys as required. Even if an attacker acquires the long-term key, he or she cannot compromise the data without the short-term key as well. You will learn more about IKE in Chapter 6, "Internet Protocol Security (IPSec)."

Data Modification

This occurs when traffic sent from Point A to Point B is intercepted in transit, modified, and then forwarded to Point B. Point B never knows that the traffic was changed since it left Point A, and Point A never knows the traffic was changed before it reached Point B.

It is possible for an attacker to set up a system that intercepts traffic and looks for data patterns or key words. The process would then replace this value with bogus data. For example, a process could search for the standard format of a Visa card number. It could save the valid Visa account number to a database and then replace the original value with a bogus account number. This would prevent the transaction from being successful, and the attacker would then have the valid Visa account number.

As with many of these attacks, the obvious solution is to ensure that data traveling over the network is encrypted. Then pattern searches would never come up with anything except encrypted and unreadable data.

Man in the Middle

This is data modification, except that the data does not have to be modified. It can be analyzed, stored, or witnessed. The "man in the middle" has complete power over the data as it goes from Point A to Point B. The key to a man in the middle is that neither Point A nor Point B ever knows that the data has been intercepted. But the man in the middle has gained complete power over the data. This enables any type of data to be compromised and stolen without detection.

It is possible for the attacker to remove the original public key sample and save it to a database. After doing this, he or she can create a bogus private/public key pair, attach it to the original message, and send the message on to the destination. A poorly configured environment could use the bogus public key that appears to be from the original system to encrypt messages. This allows the attacker to decrypt any message sent with the bogus private key he or she generated.

The key difference between this and passive network monitoring is that the attacker has complete control of the data between the source and destination. Man in the middle attacks are typically a combination of several types of attacks.

Replay Attack

A replay attack occurs when an attacker intercepts and records traffic from Point A to Point B and then retransmits to Point B using the information gained from the earlier interception. The replayed traffic is out of order and unexpected, which often causes a number of unpredictable results at the destination computer.

Even if the attacker is unable to read the intercepted data, he can replay the message at a later time. He captures the traffic and then retransmits it over the wire to the destination computer. A great example of this occurred recently when I was listening to streaming music over the Internet. My office coworker, who didn't like my taste in music, captured the stream for a few seconds and retransmitted the stream over the same wire. This completely confused my audio player, and it promptly errored, stopping the "offensive" music.

Brute Force Attack

A brute force attack occurs when data is encrypted and, typically, the algorithm is a known algorithm. However, because the attacker does not know the key, after he or she captures the encrypted data, he or she must use a powerful computer to crack the encryption by guessing at the key.

In the past, this was not a serious threat. But with today's processors and the speed at which the key can be guessed, a brute force attacker should be considered a serious threat. A hacker can now buy a kit (such as L0pthcrack to crack Windows NT passwords) to build his own hardware solution specifically designed for key cracking.

One of the most obvious ways you can make it more difficult for someone to commit a brute force attack is to increase the frequency with which keys are changed. You should consider this when defining your computer policies. Many brute force attack utilities are available on the Internet.

Password Guessers and Dictionary Attacks

As the name implies, a password guesser is a routine in which a computer tries to guess passwords, usually by selecting words from the dictionary and then appending common characters. It is important to educate your network users about the importance of using secure passwords, passwords that include different cases, numbers, and even punctuation marks. Even with advanced user education, it is generally a good idea to enforce policies so that the passwords users choose must meet predefined criteria.

Social Attack Strategies

Probably the most common way to bypass security is to take advantage of people. You can have the highest level of encryption, the most complicated password, and the best laid security plans of any organization, but if an attacker poses as a help desk employee and convinces a user to give his or her password over the phone, all your well-laid network plans are for naught.

Social attacks are not limited to telephone calls; users who write passwords on notepaper or Post-it Notes run the risk of having their passwords collected by coworkers, office visitors, or the janitor.

A common practice in cities is for hackers to dig through dumpsters and collect Post-it Notes to search for written passwords. The best defense is to make sure you implement a training plan for all network users. There are other strategies, such as smart cards and certificate-based strategies, that help prevent this problem, but it all boils down to the important fact that your users need to understand the risks and to treat security as a high priority.

Summary of Network Attack Methods

The most important thing to realize about network attacks is that with unsecured network traffic—which is what most people have at this point—you do not have anything protecting your data. If a person on your network, authorized or unauthorized, is listening to traffic on the wire, he can capture anything that is transmitted. Without any encryption in place, even a novice can capture all email from various computers and read through private email as it is being transmitted over the wire. This works in contrast to an attack on a server. It is much harder to attack a server because, at the very least, server-based permissions or access control lists (ACLs) will serve as protection.

The goal of a VPN is to protect the data that is transmitted over a network so vulnerable applications like Telnet, FTP, and POP3 will be protected. Although today's network administrators protect against simple captures of usernames and passwords, with a complete network encryption solution, all data can be protected. This level of security will protect the corporate data over all types of network links, both LAN and WAN.

Virtual Private Network Deployment

Now that you know the ways in which VPNs can benefit you, let's look at some typical network deployments with VPNs.

To avoid using long distance or 1-800 numbers (and their associated costs), an offsite client can use the Internet to establish a VPN connection to the corporate network, as shown in Figure 2.1.

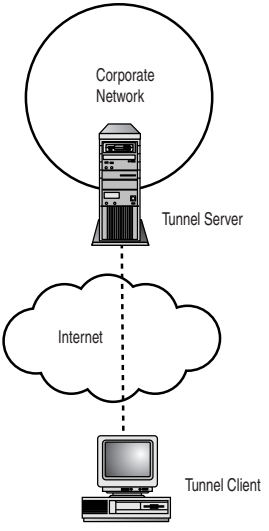

Figure 2.1 *Remote access over the Internet.*

Two sites with dedicated or dial-up links to an ISP can have VPN links. As an example of this, consider the branch office connections shown in Figure 2.2.

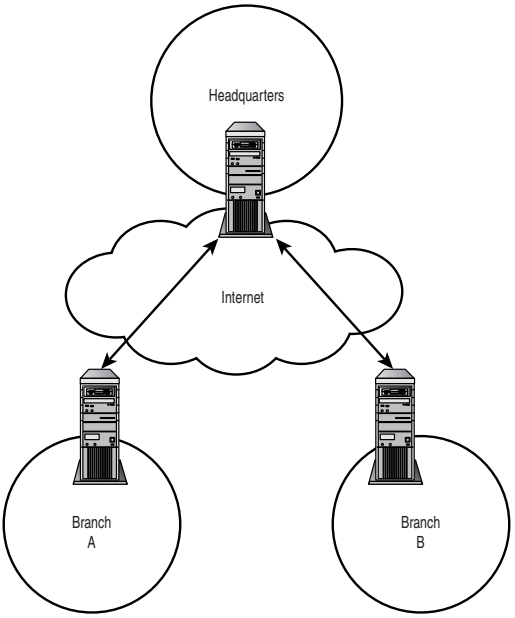

Figure 2.2 *Connecting sites over the Internet.*

Tunneling over a corporate network is basically the same as remote access over the Internet, except that both the client and the destination networks are on the corporate network. This enables users to access secure or hidden networks, as shown in Figure 2.3.

Tunneling can be used to connect two or more secure or hidden networks on the same corporate network based on account security. As an example of this, consider how two hidden networks are connected by a VPN connection over the intranet, as shown in Figure 2.4.

In addition, a tunnel linking two networks can have another tunnel within the tunnel, as shown in Figure 2.5. This setup is sometimes used to solve multiprotocol issues and is sometimes needed for unusual network designs. Another reason to use a tunnel within a tunnel is to encrypt IP traffic for end-to-end security.

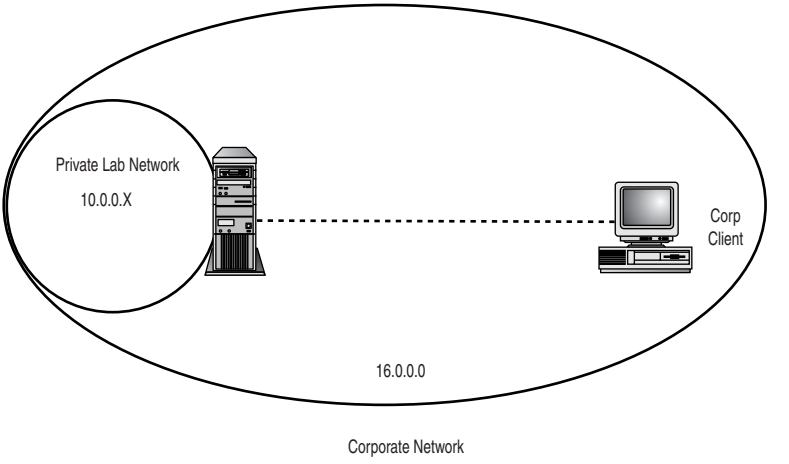

Figure 2.3 *Remote access over a corporate network.*

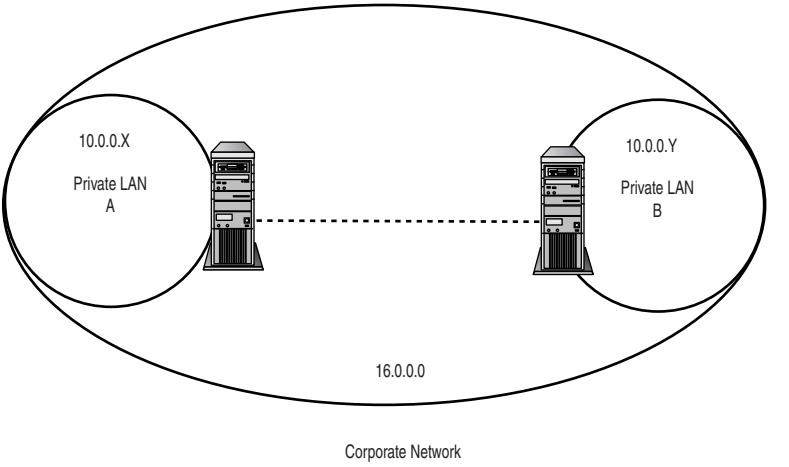

Figure 2.4 *Connecting corporate networks.*

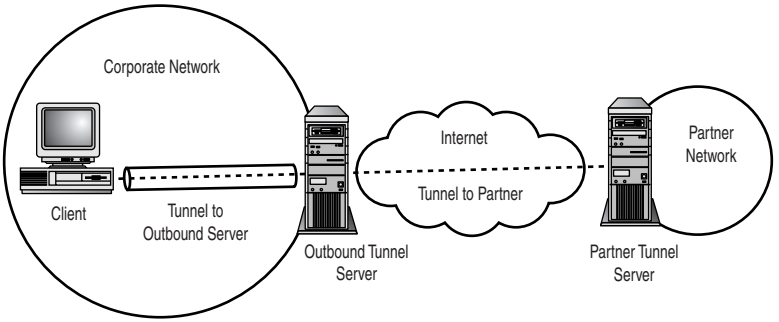

Figure 2.5 *Tunnel within a tunnel.*

If the ISP has a VPN service, ISP-provided VPN services can be used to link separate sites by initiating a tunnel from two network links through the ISP's VPN network. Such a configuration is shown in Figure 2.6.

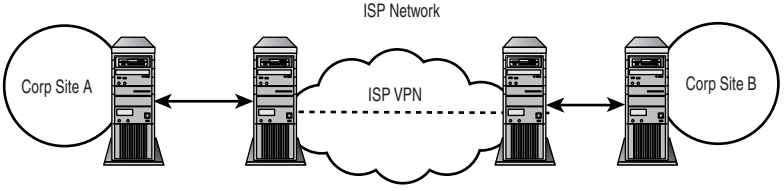

Figure 2.6 *Tunneling to an ISP-provided VPN service.*

Network Design Concepts with Tunneling

As you start to see how tunneling works, it is important that you understand how it fits into your environment. Although it is possible to implement most VPN designs because of the flexibility of the technology, some designs work better than others. It is up to the network architect to decide what is best for a particular environment. Several areas need to be addressed, such as network infrastructure, network topology, and firewalls.

Network Infrastructure

To start, the network administrator must make sure he or she is familiar with the current network environment. To roll out a secure network environment, you must have an accurate inventory of the current network topology. This can be a huge challenge, depending on the size of the network. A good way to track down external links is to analyze phone bills. The accounts from the phone company will show leased lines, as well as

links that might have been installed long ago. I have seen networks that have grown so much or have been involved in so many mergers that no one knows the details of all of the links. Until such data is discovered and documented, potential security problems and possible technical issues are inevitable.

Another problem the network administrator has to consider is the type of protocols that need to be supported. More and more networks are moving toward the exclusive use of TCP/IP. But there are certainly exceptions to this. Your VPN infrastructure might need to support IPX, AppleTalk, NetBEUI, and others. You must know which protocols are needed before you can implement the VPN rollout.

For example, you need to know whether it's necessary to continue to support IPX for legacy NetWare servers. If your tunnel server is a Windows 2000-based server and the clients still need to access NetWare servers, they must use either PPTP or L2TP to access the corporate network because those protocols permit IPX. This protocol requirement also affects the router configuration and the client configuration. This additional configuration consideration must be part of your deployment plan.

Network Topology

Network topology plays an important role in the design of your corporate VPN plan. First, you must have a thorough knowledge of your routing environment. Your VPN clients and servers have to fit into the routing infrastructure of your network. Otherwise, your routing infrastructure needs to be changed to reflect your VPN rollout. There are a tremendous number of options when it comes to routing with VPNs, and each network has different needs and, likely, different paths.

The network topology is critically important to the design of a VPN-based solution for branch office links. Your design must fit with the existing infrastructure, or the new routes will never work.

The network administrator must have a full understanding of TCP/IP and routing, as well as knowledge of the existing network. When tunnel servers are introduced, even for just client access, a number of network issues must be addressed. The plan needs to be defined for DDNS, DHCP scope options, routes, network load, external addresses, and more. None of these designs should be planned until the existing network infrastructure and topology is fully understood.

Firewalls

Your firewall configuration is very important to VPN design. In large organizations with a number of firewalls, a common problem is that not all the firewalls are configured in the same way. In fact, many companies have different types of firewalls, connected in various different ways. Because the firewalls affect the way you roll out your VPN, it is extremely important that the configuration be consistent throughout the organization. The firewalls need not be the same type, but you need to be able to get the same results from each.

Firewall configuration is also going to be a big issue because of demilitarized zones, protocol passthrough settings, and ACLs. It is extremely important that the network administrator coordinates and manages the firewall settings as they relate to VPNs. Because the firewall is the lifeline to the network, when a firewall is misconfigured, it can bring the network down. I cover this in detail in Chapter 8, "NAT and Proxy Servers."

Some network environments might not even have firewalls at remote sites. These networks rely on packet filtering (packet filtering is also a type of a "firewall" that blocks unauthorized traffic). Additionally, some remote access clients, either by network policy or manually, might want to initiate packet filtering at the time of the connection to the VPN. These strategies help ensure that no unauthorized traffic comes into the Internet-accessed corporate network resources through the VPN. Packet filtering issues need to be addressed based on the needs of the corporate network.

Many VPN network designs rely on both one-way and two-way initiated connections using the demand-dial routing features of Routing and Remote Access Service (RRAS). This dictates how your VPN routing environment should be defined. It also dictates how security zones are implemented from a routing point of view.

Summary

As you review the ways in which tunneling can be deployed, it is very important to consider how the technology can be deployed in your specific environment. The network infrastructures of no two companies are alike, and it would be impossible to create a guide that would cover all types of network configurations. By carefully considering the technical information about the features of tunneling in Windows 2000, the goals you are trying to achieve, and the existing configuration of your environment, you should be able to define a deployment plan that works very well.