Virtual private networks (VPNs) are becoming more and more popular as a means for mobile users to securely connect to their corporate networks. Essentially, VPNs establish a *tunnel* of encrypted traffic between two computers. All traffic passing through the tunnel is protected by the encryption, and the traffic is decrypted at the destination, where it can be used as normal network traffic. VPNs are most frequently used over public networks, such as the Internet.

There are two basic types of VPNs. The first type is a *client-to-server* VPN and is when a single remote user establishes a VPN with some type of VPN gateway device on the corporate network. Client-to-server VPNs are similar to dial-up connections, only the user "dials up" with a VPN instead of with a phone line. VPN gateway devices can accept VPNs from several users at once, and the VPN makes it seem to the user as if his remote computer is directly attached to the corporate network. Many companies use client-to-server VPNs for their mobile users: Users simply dial into a local ISP and then establish a VPN to the corporate network.

The other type of VPN is a *server-to-server* VPN, also called *router-to-router*. In this type of VPN, two server computers or routers establish a VPN connection to one another for the purposes of connecting two remote networks to one another. Traditionally, remote networks were connected with truly private connections, such as point-to-point T1 lines, frame-relay lines, and so forth. VPNs enable corporations to connect branch offices over Internet connections instead.

Both types of VPNs can be built from one of two common VPN protocols. The oldest VPN protocol is the *Point to Point Tunneling Protocol (PPTP)*. Originally developed by Microsoft, PPTP provides strong encryption and tunneling capabilities, includes integrated support for user authentication, and is very easy to configure and maintain. Unfortunately, PPTP is pretty much unique to the Windows platform, making it less desirable in mixed-OS networks. Another downside to PPTP is security: PPTP encrypts only the data sent through the tunnel; it does not encrypt the tunnel itself. Although that might seem to be a pretty fine point, keep in mind that hackers can easily determine when PPTP traffic is in use because the tunnel itself is unencrypted. That knowledge enables hackers to immediately begin focusing on breaking the encryption on the tunnel's contents, meaning you're effectively surrendering half the battle to protect your data. Of course, PPTP's unencrypted tunnels enable PPTP to easily pass through firewalls performing Network Address Translation (NAT) because the tunnel's header packets can be translated more easily.

A somewhat newer VPN protocol is the *Layer 2 Tunneling Protocol (L2TP)*. L2TP provides only the tunneling portion of the VPN and works in conjunction with IPSec to provide the encryption portion of the VPN. L2TP/IPSec tunnels are considered somewhat more secure because the entire tunnel is encrypted, making it difficult for hackers to determine that L2TP is even in use. Unfortunately, passing L2TP/IPSec tunnels through a NAT device such as a firewall is difficult, and Windows Server 2003 includes one of the first implementations of L2TP/IPSec to support NAT. A major downside to L2TP/IPSec is the difficulty of configuring VPNs. IPSec is primarily designed to use digital encryption certificates for encryption, rather than the dynamically generated keys PPTP uses. IPSec's encryption requirement means you have to issue certificates to VPN users, manage certificate trust and revocation lists, and much more. IPSec does support the use of *shared secrets*—essentially preconfigured passwords—instead of certificates, but the encryption afforded by shared secrets is much less secure than that provided by certificates. L2TP/IPSec tunnels are supported by Windows 2000 and later Microsoft operating systems, as well as by a huge array of non-Microsoft software and operating systems, making L2TP more universally accepted.

Typically, VPN software on a server requires you to identify an internal and an external interface. The *external* interface accepts incoming VPN connections, whereas the *internal* interface provides a connection to the corporate network. The same network adapter can act as both the internal and external interface, accepting VPN connections and internal network connections, but this configuration is considered less secure. The most secure way to use a VPN server is to place it between two firewalls. The outer firewall connects directly to the Internet and ensures that only VPN traffic is permitted to pass from the Internet to the VPN server's external interface. The inner firewall connects to the corporate network and allows only traffic originating from the VPN server's internal interface to pass to the corporate network. The VPN server thus sits on a *demilitarized zone*, which is a special network designed to act as a secure buffer between the internal network and the Internet.

The Windows Server families implement VPN functionality through the Routing and Remote Access Service (RRAS). This software enables administrators to define ports to receive VPN connections, define router-to-router connections, and so forth. Windows also includes client VPN functionality in the base Dial Up Networking software, which enables users to create and manage client-to-server VPN connections.