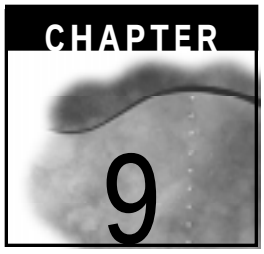


# Firewalls and Virtual Private Networks

## Introduction

In Chapter 8, we discussed the issue of security in remote access networks. In this chapter we will consider how security is applied in remote access networks. In particular, we will discuss how *firewalls* are used to protect corporate resources from outside intruders and how *virtual private networks* enable branch offices and remote users to access the corporate network in a secure manner via non-secure public networks.

## Firewall Protection

As stated in Chapter 8, remote access networking is a necessity in most corporations. For most of these corporations, the Internet is the virtual backbone of their enterprise network, interconnecting an organization's corporate network and those of its business partners and customers. The Internet also provides an inexpensive way to link branch offices, telecommuters, and mobile workers to the corporate network. Unfortunately, linking a corporate network to the Internet exposes it to the outside world. While the security mechanisms discussed in Chapter 8 can deter unauthorized access to the corporate network via the Internet, they are not foolproof. Extra

steps are usually taken, in conjunction with those discussed in Chapter 8, to protect the confidential information located in the corporate network from external unauthorized users.

This protection is provided by implementing appropriate network access control policies and using firewalls to enforce them. A firewall is a security system that controls access to a protected network, such as a private corporate network. The network is being protected from an untrusted public network, such as the Internet. As a result, a firewall is located so that every access request from a public network to the protected network must pass through the firewall, eliminating the need for individual protection of every server and host in the protected network.

A firewall is typically located the point the network connects to the Internet. This location permits the firewall to provide authentication and other security services to remote users in order to prevent unauthorized users from logging in to the network. Figure 9.1 illustrates a firewall-controlled access to the corporate network from the Internet.

For a firewall to be effective, companies first need to define their network security policy. A network security policy identifies the resources that need protection and the threats against them. It then defines how they can be used and who can use them, and stipulates the actions to be taken when the policies are violated. A policy is a set of rules against which arriving packets are tested. Examples of such rules include what IP traffic the organization wants to allow into its network, what source addresses should be excluded from the network, and what destination addresses within the network can be accessed from outside the network. Specific actions to be taken include *accept packet* and *reject packet*. The firewall is responsible for filtering traffic according to the security policy.

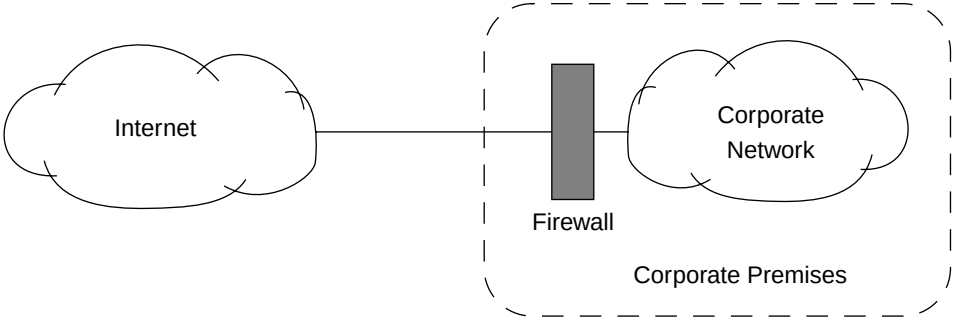


Figure 9.1 Firewall-controlled access from the Internet.

## Types of Firewalls

Firewalls can be classified into three basic categories: *packet filters*, *proxy servers* (which include *application gateways* and *circuit-level gateways*), and *stateful packet filters*. There is a fourth category that is essentially a hybrid of the three main categories. For example, a firewall may be a combination of the application gateway and packet filter, or a proxy server and a stateful packet filter. Figure 9.2 illustrates the different types of firewalls.

### *Packet Filters*

A packet filter is a firewall that inspects each packet for user-defined filtering rules to determine whether to pass or block it. For example, the filtering rule might require all Telnet requests to be dropped. Using this information, the firewall will block all packets that have a port number 23 (the default port number for Telnet) in their header. Filtering rules can be based on source IP address, destination IP address, Layer 4 (that is, TCP/UDP) source port, and Layer 4 destination port. Thus, a packet filter makes decisions based on the network layer and the transport layer.

Packet filters are fast and can be easily implemented in existing routers. Unfortunately, they are the least secure of all firewalls. One disadvantage of packet filters is that they have no logging facility that can be used to detect when a break-in has occurred. Also, a packet filtering firewall grants or denies access to the network according to the source and destination addresses and the source and destination ports. Unfortunately, these ports can be spoofed. As a result, anyone can access network resources once access has been granted to an authorized user.

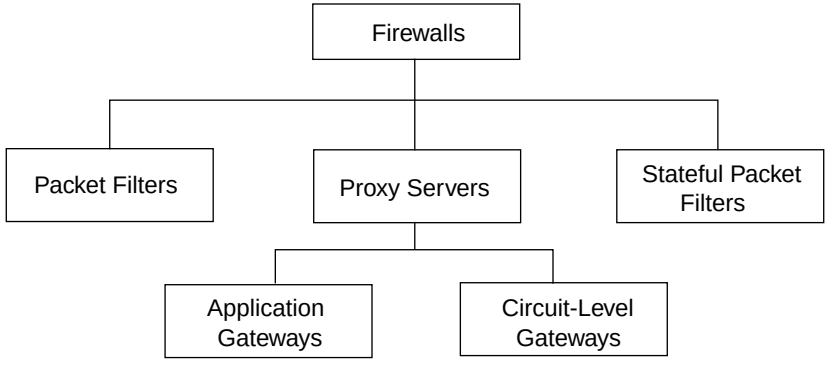


Figure 9.2 Firewall classification.

### Proxy Servers

A proxy service is an application that redirects users' requests to the actual services based on an organization's security policy. All communication between a user and the actual server occurs through the proxy server. Thus, a proxy server acts as a communications broker between clients and the actual application servers. Because it acts as a checkpoint where requests are validated against specific applications, a proxy server is usually processing intensive and can become a bottleneck under heavy traffic conditions.

Proxy servers can operate at either the application layer or the transport layer. Thus, there are two classes of proxy servers: application gateways, which operate at the application layer; and circuit-level gateways, which operate at the transport layer.

### Application Gateways

An application gateway is a proxy server that provides access control at the application layer. It acts as an application-layer gateway between the protected network and the untrusted network. Because it operates at the application layer, it is able to examine traffic in detail and, therefore, is considered the most secure type of firewall. It can prevent certain applications, such as FTP, from entering the protected network. It can also log all network activities according to applications for both accounting and security audit purposes.

Application gateways can also hide information. Since all requests for services in the protected network pass through the application gateway, it can provide network address translation (or IP address hiding) functionality and conceal IP addresses in the protected network from the Internet by replacing the IP address of every outbound packet (that is, packets going from the protected network to the Internet) with its own IP address. Network address translation also permits unregistered IP addresses to be freely used in the protected network because the gateway will map them to its own IP address when the users attempt to communicate with the outside world.

### Circuit-Level Gateways

A circuit-level gateway is a proxy server that validates TCP and UDP sessions before allowing a connection or circuit through the firewall. It is actively involved in the connection establishment and does not allow packets to be forwarded until the necessary access control rules have been satisfied.

A circuit-level gateway is not as secure as an application gateway because it validates TCP and UDP sessions without full knowledge of the applications that use these transport services. Moreover, once a session has been established, any application can run across that connection. This behavior exposes the protected network to attacks from intruders. Unlike a circuit-level gateway, an application gateway can differentiate the applications that need to be blocked from those that can be allowed to pass through the gateway.

### *Stateful Packet Filters*

Although the application gateway provides the best security among the preceding firewalls, its intensive processing requirement slows down network performance. A stateful packet filtering gateway attempts to provide tight security without compromising performance. Unlike the application gateway, it checks the data that passes through at the network layer but does not process it. The firewall maintains state information for each session, where session states include a combination of communication phase and the endpoint application state. When a stateful packet filtering gateway receives a data packet, it checks the packet against the known state of the session. If the packet deviates from the expected session state, the gateway blocks the rest of the session.

## Firewall Architectures

Firewall architecture refers to the manner in which firewall components are arranged to provide effective protection against unauthorized users. It is usually defined after the network security policy has been defined because it is supposed to be a model that enforces the security policy.

The network security policy is enforced at defensible boundaries within the network called *perimeter networks*. A corporate network usually contains multiple perimeter networks that can be classified into three groups: the *outermost perimeter* network, one or more *internal perimeter* networks, and the *innermost perimeter* network. The outermost perimeter network provides a boundary between corporate resources (that need to be protected) and external resources (resources the corporation cannot control). Internal perimeter networks represent boundaries within the corporate network that need additional security. Figure 9.3 shows the relationships among the three types of perimeter networks.

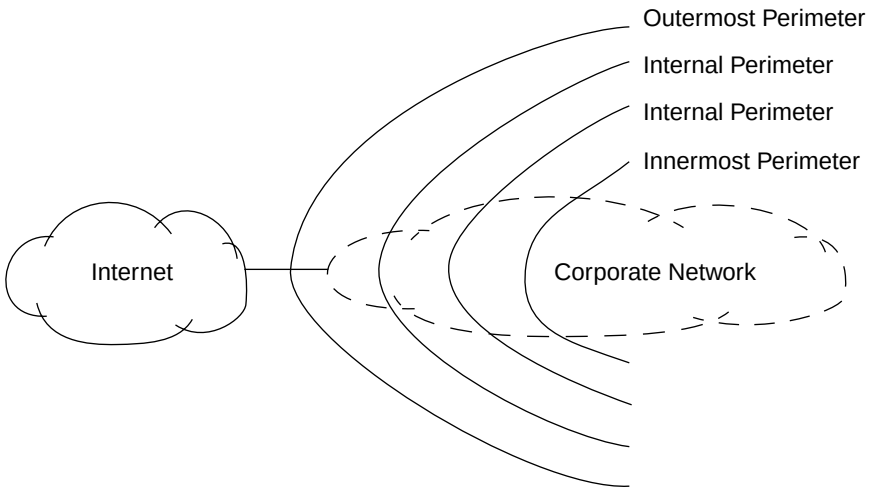


Figure 9.3 Perimeter networks.

Figure 9.3 provides only one potential firewall configuration; not all networks have three levels of perimeter networks. Organizations need to configure their firewalls to meet their network security policy. The following section discusses the three most commonly used firewall architectures.

### *Dual-Homed Host Firewall*

A dual-homed host firewall is a type of the multi-homed host firewall. In the dual-homed host firewall, the host (which provides the firewall functionality) has two interfaces: One interface is connected to the private network and the other interface is connected to the Internet (or some other untrusted network). Thus, all IP traffic from the Internet must pass through the firewall before arriving at a host in the private network. Similarly, an internal host can communicate with external hosts (that is, hosts in the Internet) via the dual-homed host.

Direct communication that bypasses the dual-homed host is blocked. This means that the IP forwarding capability of the dual-homed host is disabled to ensure that IP packets from one network are not be directly routed to the other network. The dual-homed host cannot operate as a router. However, disabling IP packet forwarding ensures that the Internet and the private network are logically disconnected so that even when system problems occur the firewall cannot *fail open*. Data can only pass

through the firewall via its application proxies, not through the operating system layer. Figure 9.4 illustrates a dual-homed host firewall.

The dual-homed host can be used to completely block access to the private network when it provides proxy services (such as Telnet, FTP and HTTP), as shown in Figure 9.4, the server providing the particular service is located between a packet filtering router, which may be present, and the dual-homed host. This arrangement prevents intruders from accessing systems protected by the dual-homed host.

### *Screened Host Firewall*

Unlike the dual-homed host firewall architecture, which allows the host firewall to be connected to two networks, the screened host firewall architecture allows the host providing the firewall (called a *bastion* host) to connect only to the private network. A separate screening router is placed between the host and the Internet. Thus, the screened host firewall is a combination of a packet filtering router and an application gateway.

The screening router performs a packet filtering function and is configured so that the bastion host is the only host in the private network that can be accessed from the Internet. Extra security can be provided in the screening router by configuring it to block traffic to specific ports on the bastion host. Figure 9.5 illustrates the screened host firewall.

The screening router may be configured to permit or block connections between internal hosts and the Internet. Its basic function is to filter traffic

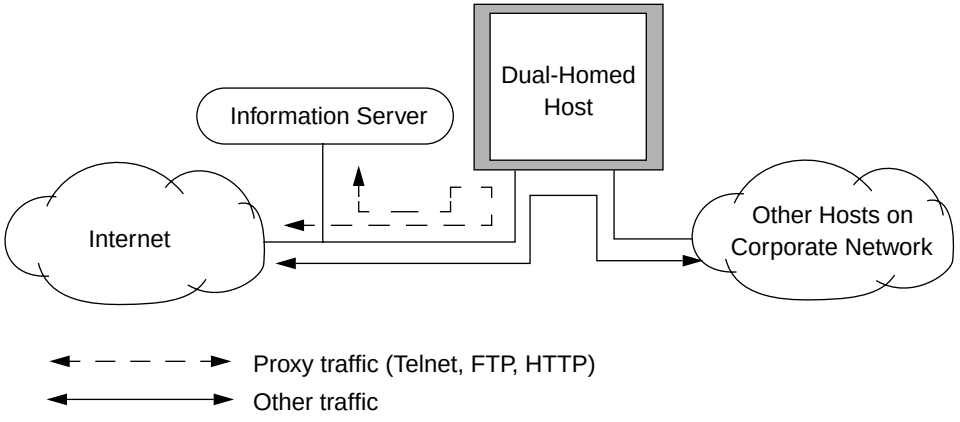


Figure 9.4 Dual-homed host firewall architecture.

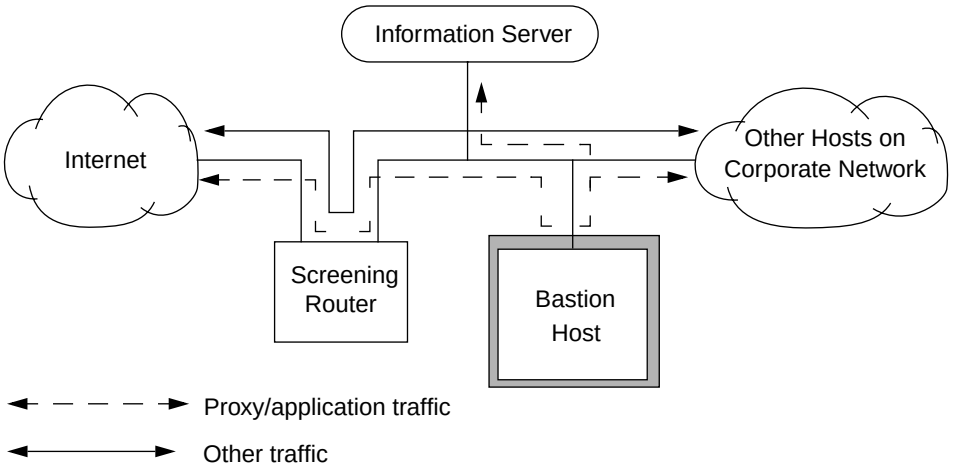


Figure 9.5 Screened host firewall.

classes that have been defined as security risks in the security policy before they arrive at the bastion and other internal hosts. Since the bastion host is the most exposed host in the private network, it is usually the most protected host. Generally, there are two or more bastion hosts in a network.

One of the advantages of this architecture is that a public information server providing FTP, Telnet, and HTTP services can be placed on the network segment between the screening router and the bastion host. If stronger security is desired, the bastion host can be configured to run proxy services that require both internal and external users to access the information server through the bastion host. In fact, one of the main functions of the bastion host is to act as a proxy server for various services including FTP, HTTP, Telnet, DNS, and SMTP.

One of the problems with the screened host firewall architecture is that once an attacker breaks through the bastion host, all the hosts in the private network are exposed to the attacker. In a dual-homed host firewall, however, it is impossible to pass through the dual-homed host without a corresponding proxy server. Unlike the dual-homed host firewall, the screened host firewall requires the screening router and the bastion host to be configured

### *Screened Subnet Firewall*

The screened subnet firewall can be considered an extension of the screened host firewall. Like the screened host firewall, it uses a screening

router (called the outer or external router) and a bastion host. However, this firewall, which is also called the *Demilitarized Zone* (DMZ), creates an extra layer of security by adding a perimeter network that further isolates the private network from the Internet. The firewall defines a DMZ demarcated by the outer router and an internal router, where the latter is placed closer to the private network than the outer router. The DMZ is an inner screened subnet bounded by the internal router and the outer router. The bastion host and information server are then located within the DMZ, as shown in Figure 9.6.

The DMZ can be considered an isolated network between the private network and the Internet. The outer router protects the network from external attacks by restricting access to systems in the DMZ. It also blocks traffic to the Internet from unauthorized sources in the private network. The internal router manages DMZ access to the private network by passing only traffic from the bastion host to the hosts in the private network that are not in the DMZ, and vice versa.

For an attack to reach any internal host located outside the DMZ, it must break through both routers. In addition, the architecture reveals only the DMZ network to the outside world and keeps the private network hidden. However, there is a drawback to this architecture: two routers and the bastion host need to be properly configured.

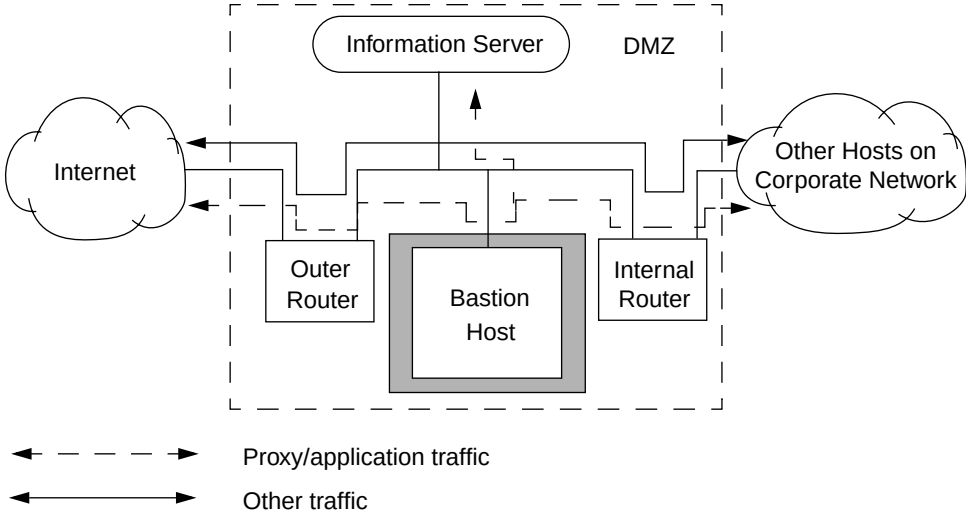


Figure 9.6 Screened subnet firewall.

## Virtual Private Networks

A virtual private network (VPN) provides a secure connection between a sender and a receiver over a public non-secure network such as the Internet. A secure connection is generally associated with private networks. (A private network is a network that is owned, or at least controlled via leased lines, by an organization.) Using the techniques discussed later in this chapter, a VPN can transform the characteristics of a public non-secure network into those of a private secure network. VPNs reduce remote access costs by using public network resources. Compared to other solutions, including private networks, a VPN is inexpensive.

VPNs are not new. In fact, they have been used in telephone networks for years and have become more prevalent since the development of the intelligent network. Frame relay networks, which have been around for some time, are VPNs. Virtual private networks are only new to IP networks such as the Internet. Therefore, some authors use the terms *Internet VPN* and *virtual private data network* to distinguish the VPN described in this chapter from other VPNs. In this book, the term VPN refers to Internet VPN.

The goal of a VPN is to provide a secure passage for users' data over the non-secure Internet. It enables companies to use the Internet as the virtual backbone for their corporate networks by allowing them to create secure *virtual* links between their corporate office and branch or remote offices via the Internet. The cost benefits of VPN service have prompted corporations to move more of their data from private WANs to Internet-based VPNs.

A VPN uses data encryption and other security mechanisms to prevent unauthorized users from accessing data, and to ensure that data cannot be modified without detection as it flows through the Internet. It then uses the *tunneling* process to transport the encrypted data across the Internet. Tunneling is a mechanism for encapsulating one protocol in another protocol. In the context of the Internet, tunneling allows such protocols as IPX, AppleTalk, and IP to be encrypted and then encapsulated in IP. Similarly, in the context of VPNs, tunneling disguises the original network layer protocol by encrypting the packet and enclosing the encrypted packet in an IP *envelope*. This IP envelope, which is an IP packet, can then be transported securely across the Internet. At the receiving side, the envelope is removed and the data it contains is decrypted and delivered to the appropriate access device, such as a router.

Simply put, a corporation is creating a private *tunnel* through the Internet for the secure delivery of its data. The tunnel enables the

corporation to create a virtual WAN, within the Internet, that is cheaper than private WANs and safe from intruders.

VPNs also provide quality of service (QoS) guarantees, which usually specify an upper bound on the average packet delay in the network. The guarantees may also include specification of a lower bound on the bandwidth available to the user. These guarantees are developed through service level agreements (SLAs) with the service provider. Most service providers have their own private backbone IP networks. Therefore, they are in a better position to meet QoS guarantees. However, such networks do not provide the global coverage as the Internet. Sometimes a service provider makes private peering arrangements with other service providers, whose private IP networks cover more area. This arrangement allows the service provider to hand off its high-priority traffic (that is, traffic with stringent QoS requirements) to networks with a guaranteed QoS instead of putting such high-priority traffic on the Internet where delay is unpredictable.

From the preceding discussion, we can define a VPN by the following relationship:

VPN = Tunneling + Security + QoS Guarantees

The remainder of the chapter deals with different tunneling techniques used to create VPNs.

## Advantages of VPNs

As stated earlier, VPNs are inexpensive; they provide remote and mobile user with access to the corporate network at the price of a local call.

VPNs also provide the framework for corporate *intranets* and *extranets*. Corporations can exploit the global nature of the Internet and use VPNs to link all branch offices into private networks called intranets. A corporation can make certain sections of its intranet accessible to its vendors and strategic partners by means of the extranet. Both intranets and extranets are discussed in greater detail in Chapter 10.

VPN tunnels permit non-routable protocols to be delivered to specific LAN segments in the corporate intranet. In this way, VPNs enable legacy applications to use the intranet.

VPNs have also contributed significantly to the increased use of private IP addresses. Since applications are tunneled rather than routed across the WAN, companies can assign their own addresses, as long as these addresses are not advertised to the outside world, as discussed in Chapter 5.

## Types of VPNs

Currently there are three types of VPNs. While their goal is to leverage the Internet as a private enterprise backbone network, each of them addresses the needs of a different interest group in the enterprise. The three types of VPNs are as follows:

1. *Access VPNs* provide remote users such as road warriors (or mobile users), telecommuters, and branch offices with reliable access to corporate networks.
2. *Intranet VPNs* allow branch offices to be linked to corporate headquarters in a secure manner.
3. *Extranet VPNs* allow customers, suppliers, and partners to access corporate intranet in a secure manner.

Because of their growing importance in corporate networking, both intranet VPNs and extranet VPNs are discussed in greater detail in Chapter 10. Most of the discussion in this chapter deals with access VPNs, which are the building blocks for intranet VPNs and extranet VPNs.

## VPN Architectures

A VPN consists of our components: a VPN client, a network access server (NAS), a tunnel terminating device (or *VPN server*), and a VPN protocol. In a typical access VPN connection, a remote user (or VPN client) initiates a PPP connection with the ISP's NAS via the public switched telephone network (PSTN). An NAS is a device that terminates dial-up calls over analog (basic telephone service) or digital (ISDN) circuits. The NAS is owned by the ISP, and is usually implemented in the ISP's POP. After the user has been authenticated by the appropriate authentication method, the NAS directs the packet to the tunnel that connects both the NAS and the VPN server. The VPN server may reside in the ISP's POP or at the corporate site, depending on the VPN model that is implemented. (VPN models are discussed later in this chapter.) The VPN server recovers the packet from the tunnel, unwraps it, and delivers it to the corporate network. Figure 9.7 illustrates VPN architecture.

There are four tunneling protocols used to establish VPNs, and three are extensions of the Point-to-Point Protocol (PPP):

- Point-to-Point Tunneling Protocol (PPTP)
- Layer 2 Forwarding (L2F)

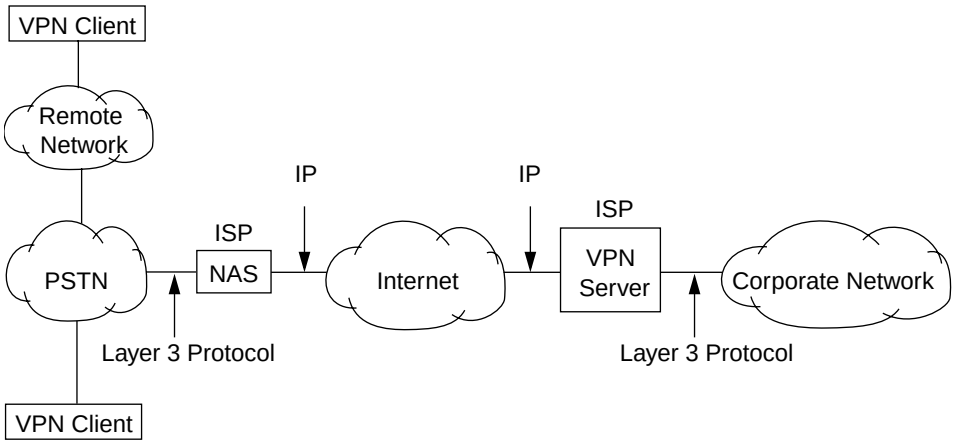


Figure 9.7 Virtual private network architecture.

- Layer 2 Tunneling Protocol (L2TP)
- IP Security (IPSec) Protocol Suite

When both PPTP and L2F were submitted to the IETF, the organization decided to combine the features of both protocols into a common tunneling protocol, which is now L2TP. The Layer 2 Tunneling Protocol was in the draft stage at the time of writing.

The four protocols can be classified broadly into two groups: Layer 2 tunneling protocols and Layer 3 tunneling protocols. PPTP, L2F, and L2TP are Layer 2 tunneling protocols, while IPSec is a Layer 3 tunneling protocol.

## Layer 2 Tunneling Protocols

Layer 2 tunneling protocols operate at the data link layer (or Layer 2). They encapsulate Layer 3 packets in Layer 2 PPP before encapsulating them in IP. They use the security provided by PPP. Thus, they perform user authentication using existing PPP authentication protocols, namely PAP and CHAP. No specific provision is made for data encryption, which may be performed by the user prior to requesting VPN service.

The general structure of the tunnel creation process for Layer 2 tunneling protocols is shown in Figure 9.8. A client initiates a PPP connection by dialing up the NAS, which is typically implemented in the ISP's POP. The client then uses the logical control protocol (LCP) negotiation to establish the PPP connection. Once the PPP connection is

established, the NAS uses either PAP or CHAP to authenticate the client. If the authentication is successful, the NAS attempts to open a PPP connection to the VPN server using LCP negotiation. The VPN server will authenticate the NAS using PAP or CHAP. The client and the VPN server will then use the network control protocol (NCP) to negotiate the network layer protocol. This completes the tunneling process.

Note that different tunneling protocols handle this process in different ways. However, in principle, they are closely related to the preceding process.

### *PPTP*

PPTP is a protocol developed by Microsoft and a group of network equipment vendors including Ascend Communications and 3Com. It permits IPX,

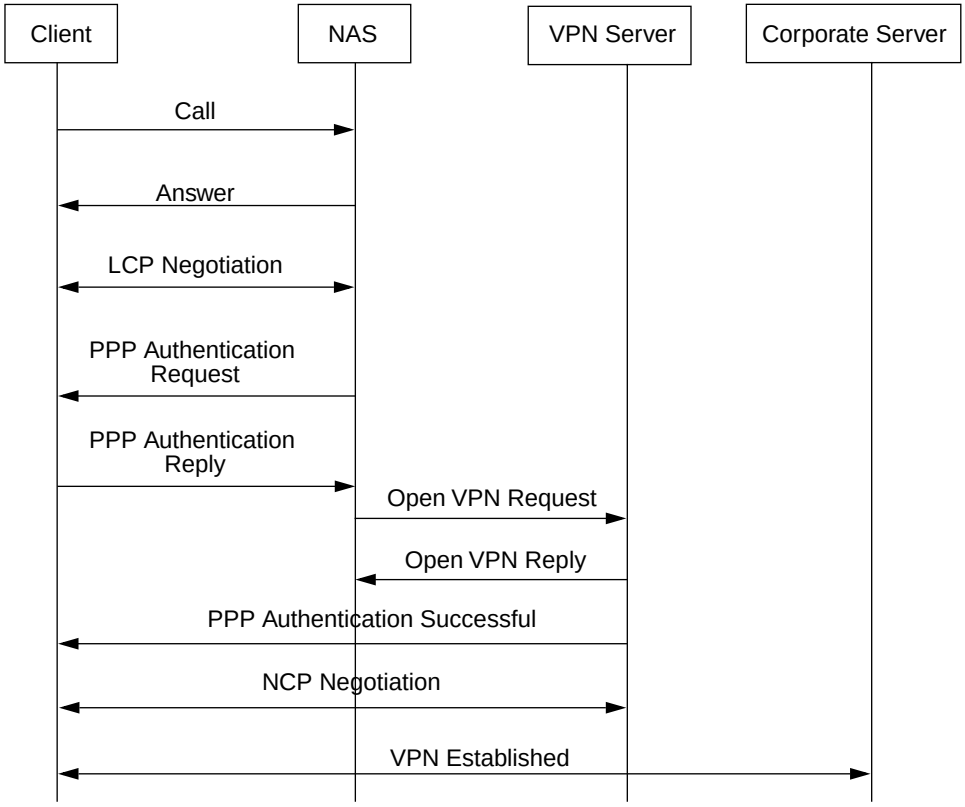


Figure 9.8 Tunnel creation process for Layer 2 VPNs.

NetBEUI, and IP packets to be encapsulated inside IP packets, enabling non-IP applications to run over the Internet. As an extension of PPP, it handles only point-to-point connections; it does not support point-to-multipoint connections. More importantly, PPTP is an IP-centric protocol that is designed only for IP networks.

In PPTP, the NAS that allows the remote user to initiate a VPN call is called the *PPTP access concentrator* (PAC), and the VPN server is called the *PPTP network server* (PNS). Figure 9.9 illustrates the components of a PPTP-based VPN.

PPTP does not provide packet-by-packet encryption. Instead, it relies on PPP's native encryption capability in PAP and CHAP. A PPTP packet is encapsulated in Generic Routing Encapsulation (GRE), which is then carried over IP. PPTP separates the control and data channels into a control stream that runs over TCP and a data stream that runs over GRE. Figure 9.10 illustrates the PPTP packet format. The PPP payload consists of the data and its TCP and IP headers.

PPTP is a proprietary protocol, but most VPNs to date are based on this protocol. PPTP uses TCP, which allows it to support flow control. It also supports a rate control mechanism that limits the amount of data in transit, minimizing the need for retransmission due to dropped packets. This results in better use of the bandwidth.

### *L2F*

L2F is a proprietary protocol that was developed by Cisco Systems. It is protocol-independent and can run over X.25, frame relay, and ATM networks.

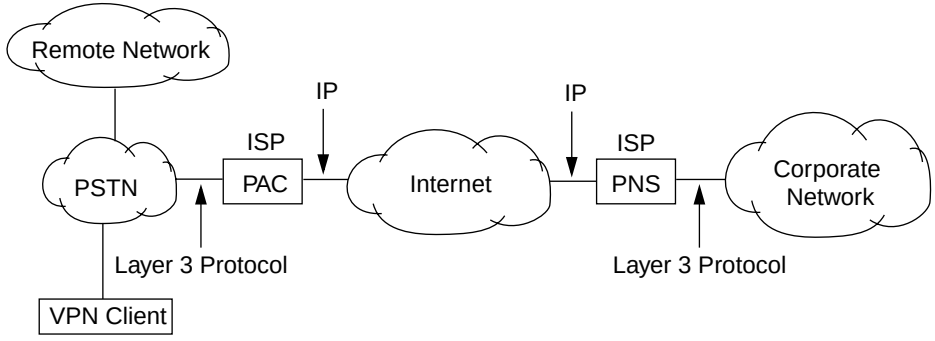


Figure 9.9 PPTP VPN components.

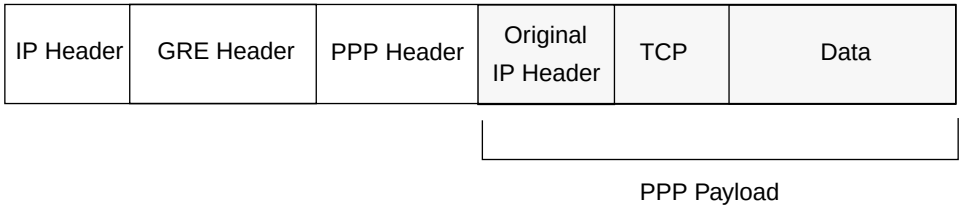


Figure 9.10 PPTP packet format.

It supports private IP, IPX, and AppleTalk, and uses UDP for Internet tunneling. L2F defines many connections within a tunnel, allowing a tunnel to support many connections.

In L2F, the VPN server is called the *Home Gateway*. L2F uses PPP for dial-up user authentication. However, it also supports other authentication schemes including RADIUS and TACACS+. Figure 9.11 illustrates the components of an L2F-based VPN.

Unlike PPTP, L2F defines its own encapsulation header, which is not dependent on IP and GRE. This capability permits L2F to work in different types of networks. Figure 9.12 illustrates the format of an L2F packet. The SLIP/PPP payload is encapsulated in an L2F packet with an L2F header and an optional L2F checksum as the trailer.

### *L2TP*

As stated earlier, L2TP combines the features of PPTP and L2F. Unlike PPTP, which runs over TCP, L2TP runs over UDP and does not use GRE. Because

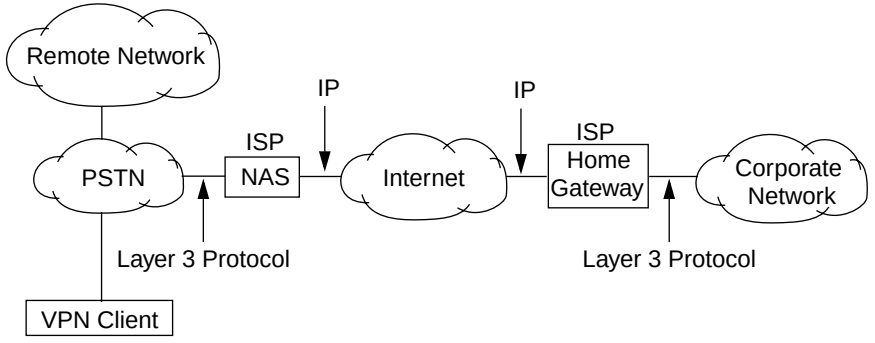


Figure 9.11 L2F VPN components.

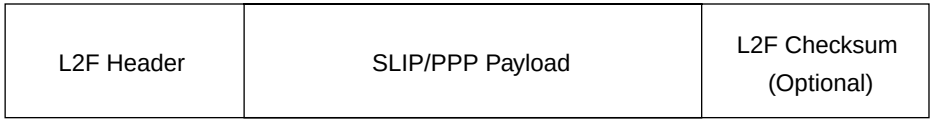


Figure 9.12 L2F packet format.

many firewalls do not support GRE, L2TP is more firewall-friendly than PPTP. In L2TP, the NAS is called the *L2TP access concentrator* (LAC) and the VPN server is called the *L2TP network server* (LNS). Figure 9.13 illustrates the components of an L2TP-based VPN.

L2TP uses PPP dial-up links and, as a result, also uses PAP and CHAP for authentication. However, it allows the use of RADIUS for user authentication. L2TP permits multiple tunnels to be created between the same pair of end-points, allowing the user to create different tunnels with different qualities of service between the same end-points. Figure 9.14 illustrates the format of an L2TP packet.

L2TP is supported by many vendors, and is expected to be the predominant protocol once it becomes a standard. L2TP relies on IPSec to perform data encryption. If L2TP discovers that IPSec is not supported at the remote end, it uses the less secure PPP encryption. The encryption may be performed by the user's workstation or by the LAC, depending on the VPN model that is used. (VPN models are discussed later in this chapter.)

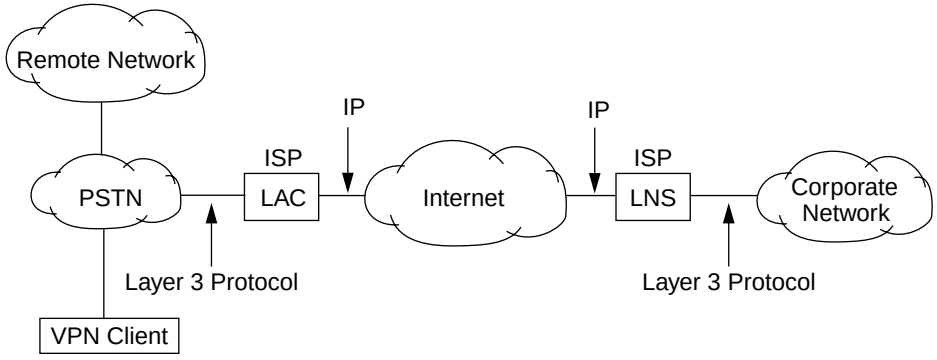


Figure 9.13 L2TP VPN components.

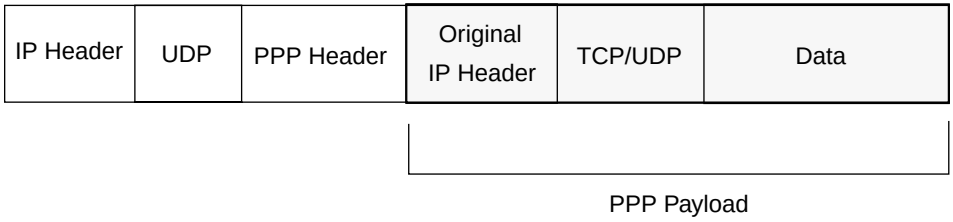


Figure 9.14 L2TP packet format.

## Layer 3 Tunneling Protocols

As discussed in Chapter 8, IPSec was originally designed to add security to the TCP/IP. IPSec provides packet-level authentication, integrity, and confidentiality by adding two security headers: the Authentication Header (AH), which provides header integrity and authentication without confidentiality; and the Encapsulating Security Payload (ESP) which provides integrity, authentication, and confidentiality to payload (or IP datagrams). ESP permits packet-by-packet encryption and uses a standards-based encryption key management protocol. Thus, IPSec VPN can be created with either AH or ESP or both. AH does not provide data encryption and is useful in those environments where only authentication is required. More importantly, since authentication is not regulated, AH is the preferred method for VPNs that cross U.S. borders. AH also has a lower processing overhead than ESP. However, when data encryption is desired, ESP is used.

One drawback to using IPSec is that it supports only IP. However, PPTP, L2F, and L2TP can support non-IP traffic, such as IPX and AppleTalk, because they are Layer 2 protocols. Unlike IPSec, Layer 2 tunneling protocols support individual dial-up access because they use PPP user authentication, which permits seamless dial-up connections through ISPs. IPSec is designed for security protection between routers and firewalls; it does not provide user authentication.

# VPN Models

There are two ways to implement Layer 2 VPNs, which are referred to as the *NAS-initiated VPN*, and the *client-initiated VPN*. In both models, the VPN client initiates a remote dial-up to the ISP's POP. However, the major difference lies in the extent of the tunnel.

## NAS-Initiated VPN

In NAS-initiated VPN, a VPN client initiates a dial-up session with the ISP's NAS. The NAS assigns the user an IP address independent of the user's IP address for the local network. The NAS is responsible for tunneling the packet through the Internet to the VPN server. Thus, the VPN connection extends only between the NAS and the VPN server.

A NAS-initiated VPN is also called a *compulsory VPN* because the client does not participate in its creation and is compelled to use it. All encryption occurs between the NAS and the VPN server and the two functional entities form the end-points of the tunnel. A compulsory VPN is created without the user's consent, meaning that the VPN is transparent to the user. One of the advantages of the NAS-initiated VPN is that it can support multiple connections, which reduces the overhead associated with establishing one VPN for each connection. However, the connection between the client and the NAS occurs outside the tunnel, making the VPN vulnerable to attacks.

Figure 9.15 illustrates a NAS-initiated L2TP VPN. This model can also be considered an *out-sourced VPN* in which the ISP manages the remote access needs of a corporation. It is particularly useful in those environments where the information technology staff is not equipped to handle VPN management.

## Client-Initiated VPN

In a client-initiated VPN, the VPN client is VPN-enabled (that, the VPN software is already installed). The VPN client dials up the ISP's local POP to establish a PPP session. Then, using the Internet connection, the client

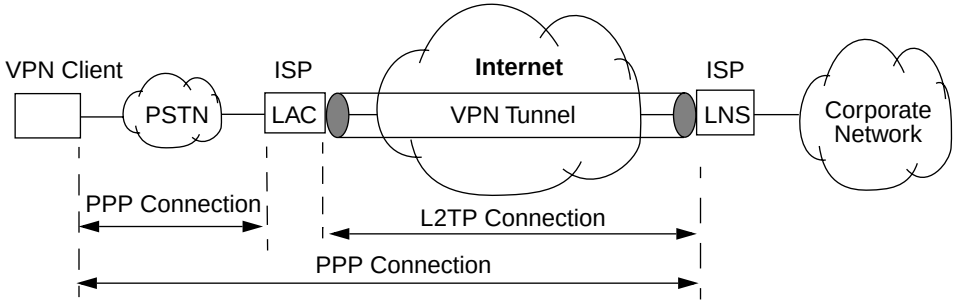


Figure 9.15 NAS-initiated L2TP VPN model.

establishes a VPN connection with the VPN server. In this model, the tunnel extends from the VPN client to the VPN server. The NAS is not involved in the tunnel establishment. A client-initiated VPN is also called a *voluntary VPN* because the user determines when and where to establish the VPN. The user is responsible for the necessary encryption between the client and the VPN server.

The VPN server may reside in either the ISP's network or in the corporate network. If the VPN server resides in the ISP's network, the contents of the tunnel are delivered to the corporate network via the WAN that the corporate router uses to access the ISP's POP. This WAN may be a frame relay network or an ISDN network. If the VPN server resides in the corporate network, the tunnel extends from the client directly into the corporate network. Figure 9.16 illustrates a client-initiated L2TP VPN with the VPN server in an ISP network and in a corporate network. In the case where the VPN server resides in the corporate network, the VPN is essentially managed in-house.

## Comparison of the Models

The client-initiated model permits remote users with VPN-enabled clients to dial into any ISP's POP and establish a tunnel to the corporate network. The ISP's POP is not required to be VPN-enabled; that is, the NAS functionality does not need to be available at the POP. In the NAS-initiated model, the ISP's POP must have NAS functionality. Unlike the NAS-initiated model, the client-initiated model does not bind a corporation to any one ISP; it permits a corporation to change ISPs without changing its addressing scheme. The client-initiated model does not lend itself easily to hacker intrusion. The model does not require a company to hand its authentication database over to the ISP. Since a company controls both ends of the tunnel, it can impose any desired requirement for user authentication to prevent unauthorized users from gaining access to its resources. In the NAS-initiated model, a company is required to give its authentication database to the ISP. Thus, a hacker can penetrate the ISP under a false identity.

NAS-initiated VPN permits VPN clients to support tunneling without software or hardware upgrades. However, the model requires the ISP to manage both the addressing and authentication processes. A corporation may have to restructure its addressing scheme to match the ISP's scheme. Also, the ISP must be constantly informed of changes in the network.

A Layer 3 VPN, such as IPSec, is more suited to the NAS-initiated model because it is designed primarily to provide security between a router and a firewall. In general, the NAS functionality is implemented in a router and the

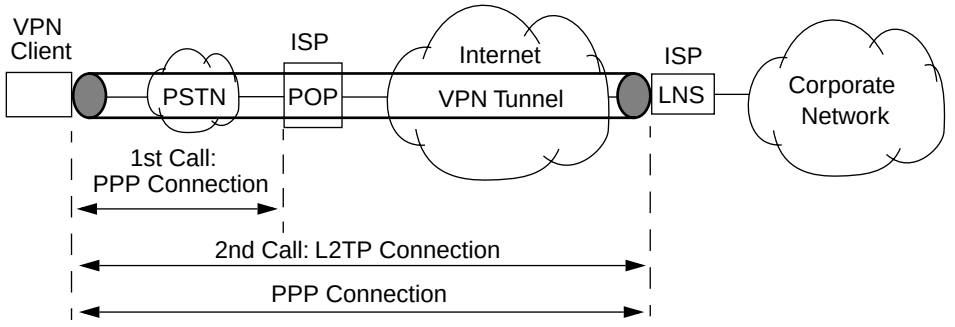


**(a) VPN server in ISP network**

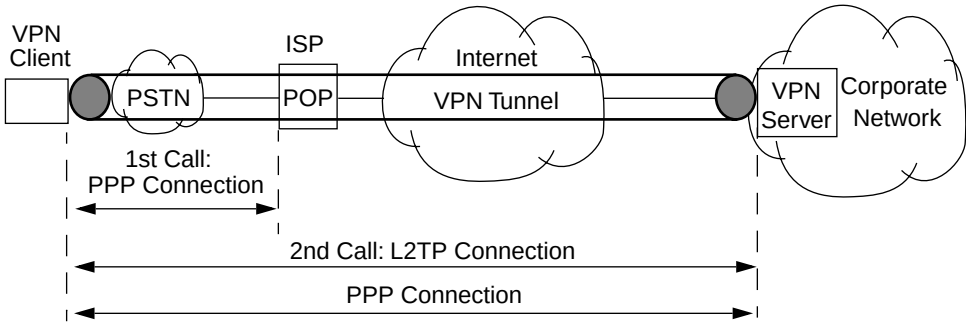


**(b) VPN server in corporate network**

Figure 9.16 Client-initiated L2TP VPN model.

VPN server functionality is implemented in a firewall. IPSec is an integral part of IPv6. Thus, unless a VPN client is running IPv6, it may not be economical for a legacy end-station implementing VPN client functionality to be equipped to provide the IPSec protocol stack. It may be more economical for the ISP's NAS to provide IPSec.

## Firewalls and VPNs

Firewalls and VPNs go hand in hand. Many firewall products provide encrypted firewall-to-firewall tunnels. In particular, it was stated that application gateways

provide IP address hiding by encapsulating one IP packet in another. This, by our definition, is the tunneling associated with VPNs.

Firewalls control access to corporate network resources and establish trust between the user and the network. Consider the network configuration in Figure 9.17. The firewall at each network controls access to resources in the network. However, the data transmitted between the two sites is still vulnerable to attack as it traverses the Internet.

On the other hand, VPNs are created to provide privacy between two sites; there is usually no trust between the two sites. A combination of firewalls and a VPN establishes trust and provides privacy between the two sites. This approach provides more security than using either firewalls at both sites or a VPN between the two sites. Figure 9.18 shows a VPN tunnel between two firewalls.

In the past, firewall products provided only firewall security service. However, many new firewall products now support VPN functionality. As stated earlier, both firewall functionality and VPN functionality are needed to establish effective security control.

## Summary

This chapter has presented the basic principles of firewalls and virtual private networks. A firewall provides access control to a protected network, protecting a company's private network from an untrusted public network. Every access request from a public network to the protected (that

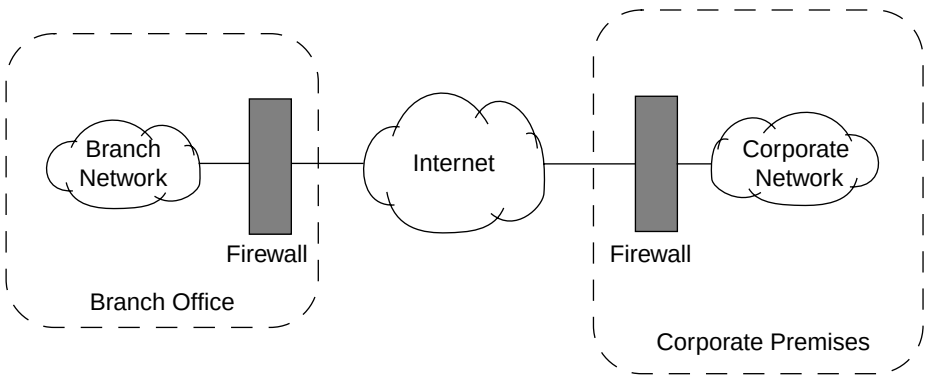


Figure 9.17 Firewalls providing authorized access to two networks.

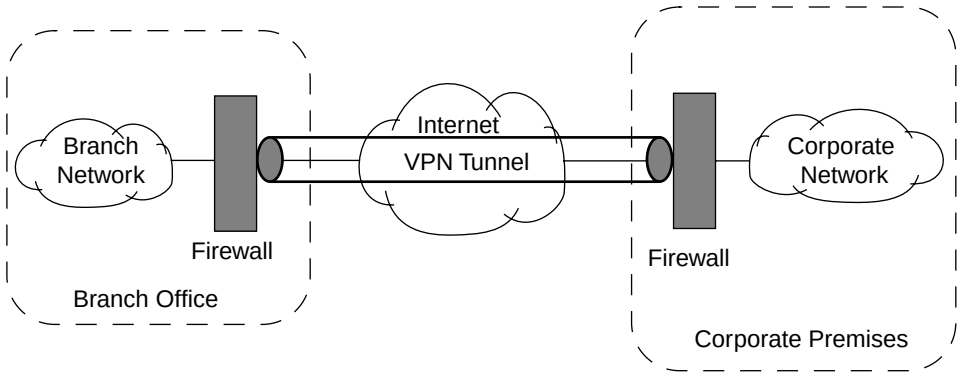


Figure 9.18 VPN tunnel between two firewalls.

is, corporate) network must pass through the firewall. Different types of firewalls and firewall architectures have also been discussed.

A virtual private network (VPN) provides secure connections using a public non-secure network, such as the Internet. VPNs reduce remote access costs by using public network resources that can be shared by many users. VPN technology has enabled companies to build intranets to link branch offices to the corporate network. Furthermore, the technology enables companies to deploy extranets that securely link corporate networks to those of their strategic partners, suppliers, and special customers.

VPNs are used in conjunction with firewalls to provide more comprehensive security protection for an organization. Firewalls control access to corporate network resources, establishing trust between the user and the network. However, the data transmitted between the user and the corporate network is still vulnerable to attack as it traverses the Internet. VPNs are created to provide privacy between two sites. Thus, combining the two technologies provides more effective access control and increases privacy.

## Further Reading

More detailed information on the topics discussed in this chapter can be found in the following references: [ATGVPN], [CHAP95], [CHECK1], [CHECK2], [CISCO5], [CISCO6], [COMPT], [IBMVPN], [IETL2T], [IETPPT], [KOS98A], [MICRO1], [MICRO2], [MICRO3], [RF1701], [RF2341], [SHIVA1], and [TIMEST].