1

Introduction

To learn something new, and review it
from time to time, is quite joyful.
—CONFUCIUS

As long as human beings have had the ability to communicate, we have had the need to keep certain conversations private. No matter the medium, one technology or another has been invented to hide content from unwanted listeners: from whispering to enciphering to scrambling pay TV channels. Private conversations discriminate between the intended audience and all others.

There are, in general, two ways to make a conversation private: *physical separation*, where only the intended audience can access the signal, and *obfuscation*, where—even though many might detect the signal—only the intended audience can understand the message. When the communication happens in a public medium, obfuscation is the only solution.

In a historical sense, the Internet is a recent phenomenon, yet it has had such a profound impact on the way people communicate that it ranks among the greatest hallmarks in the evolution of communication. The Internet has fundamentally changed both social and commercial interactions. For businesses in particular, the Internet is rapidly becoming *the* communications medium of choice. Yet conducting business requires private communications, and the Internet is a public medium. In a *virtual private network* (VPN), various networking technologies are applied toward the goal of providing private communications within the public Internet infrastructure.

A VPN is a concept composed of two parts: a *virtual network* overlaid on top of the ubiquitous interconnection of the Internet[1] and a *private network* for confidential communications and exclusive usage.

In VPNs, "virtual" implies that there is no physical network infrastructure dedicated to the private network. Instead, a single physical network infrastructure is shared among various logical networks. For example, you can use the same network access circuit to access the Internet, to connect different corporate sites, and to connect to another business's network. This virtual network allows the construction of additional logical networks by changing device configuration only. This approach is faster to deploy and is less costly than employing dedicated physical infrastructures.

Perhaps even more important is the "private" aspect of the VPN. The very purpose of a private network is to keep the data—and sometimes even the act of communicating the data—confidential so that it can be received only by the intended audience. This privacy ensures that advantages you gain by using a public infrastructure do not come at the expense of data security.

Therefore, a VPN is defined as a logical network that is created within a shared infrastructure while retaining the properties of a private network; the communication across this logical network is kept private, and the quality of the communication channel is maintained. The aim of VPNs is to use the public Internet to enable private communication to be conducted securely and reliably across the globe.

VPNs are applicable to a wide variety of users—anyone requiring private communication over a public medium. Although there is certainly much motivation outside the corporate world, business communication offers a particularly compelling case for the application of VPNs.

1.1 Business Communication

There are many types of business communication. Broadly speaking, business communication can be classified into three categories:

1. VPN has several other meanings, such as software-defined telephone network and frame relay networks. Unless otherwise noted, we use VPN to mean an Internet-based VPN.

- **Internal communication** The message is limited to selected internal audiences. For example, a corporation may periodically distribute an updated company employee directory to all its employees. Confidentiality is essential.
- **Selected external communication** The message is intended for selected external audiences. For example, a retail store may want to order a product from its supplier. Although not all communications of this type are considered proprietary, one company's business with another is generally confidential.
- **Communication with public and other external audiences** The message is intended for general public consumption. Sometimes, the wider audience the message reaches, the better. For example, a company may place a 30-second commercial during a sporting event to reach a large audience. At other times, a targeted message is designed to cater to a specific audience to maximize its impact. This type of communication is generally not confidential.

Businesses have traditionally used specialized technologies for these different types of communication and have managed them separately.

The Convergence of Business Communication

Although businesses have a variety of communication types—and hence the need for different modes of communication—the digitization of information, and the creation of computer networks to deliver it, has been a unifying factor. Internal memos are now emails, and employee directories are kept in databases. Orders can be placed online. The World Wide Web provides a means for publishing sophisticated product brochures. Although there will always be the need for traditional forms of information dissemination, much business communication is converging on a digital network.

The computer networking technologies are also converging. There used to be many types and formats of computer networks, each developed by a different vendor. IBM offered Systems Networking Architecture (SNA) for its mainframe and minicomputers. Digital had DECNET, used in the once-popular VAX computing environment. In the PC environment, Novell's Netware was dominant and still is fairly widely used for PC interconnections. Nonetheless, with the development of the Internet, most computer networks have migrated to an IP-based infrastructure.

IP—the Internet Protocol—serves as the common format for all connected network devices on the Internet.

Private Networks

To meet their information infrastructure needs, corporations have invested heavily in internal networks called *intranets*. Intranets serve the employees at the corporate site, but not employees on the road or telecommuting from home. To accommodate the remote access needs of "road warriors" and telecommuters, companies have set up remote access servers to extend intranets into the field. Usually, a bank of modems allows these users to dial in through public switched telephone networks (PSTNs). Furthermore, employees at branch offices require access to the same information and the same resources, so private lines are used to interconnect the various sites to make one corporatewide intranet.

Special arrangements are sometimes made to allow business partners to have limited access to some part of the corporate intranet.[2] These networks, usually called *extranets*, provide the means to improve the efficiency of business information flow.

Each form of access to the intranet, as shown in Figure 1-1, is a separate private networking solution. This is true even when some aspects of each solution, such as the underlying networking protocols used, are the same. Each form of access also has its own requirements for privacy—requirements that are met by keeping data transmission on separate dedicated channels.

Public Networks

It is also imperative for a corporation to exchange information outside the established private networks. This requires access to a public networking infrastructure such as the Internet.

In addition, the public network opens a new avenue of commerce. It is now unthinkable for a corporation not to have a presence in the World Wide Web. For

2. Here and elsewhere, we use the term *business partner* to mean external corporate organizations—such as vendors of parts or supplies—that work closely with your business and to which you give limited access to certain records.

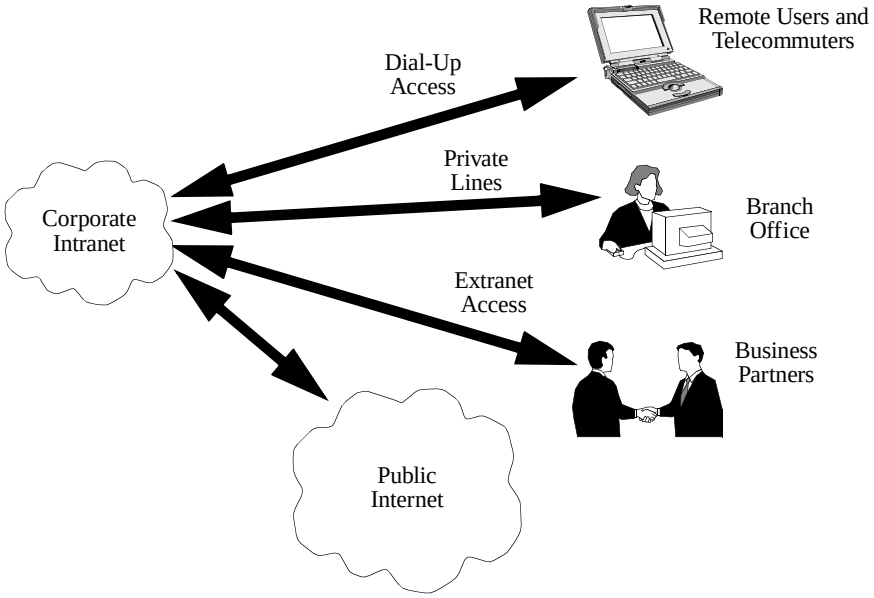

Figure 1-1 Separate private networks

many companies, such as Amazon.com, there is no “brick and mortar” storefront. The only place where they face customers is in cyberspace.

Virtual Private Networks

Protection of private corporate information is of utmost importance when designing an information infrastructure. However, the separate private networking solutions are expensive and cannot be updated quickly to adapt to changes in business requirements. The Internet, on the other hand, is inexpensive but does not by itself ensure privacy. Virtual private networking, as shown in Figure 1-2, is the collection of technologies applied to a public network—the Internet—to provide solutions for private networking needs. VPNs use obfuscation through secure tunnels, rather than physical separation, to keep communications private.

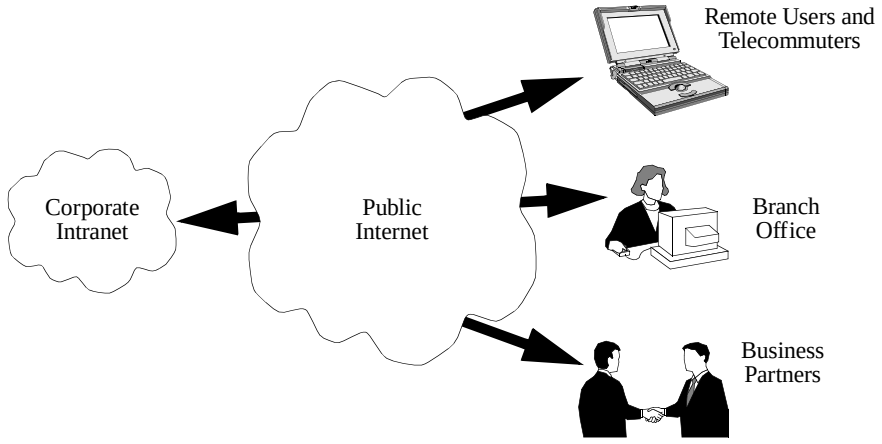

Figure 1-2 Virtual private networking

1.2 VPN Motivation

Why is it useful to employ virtual private networks for business communication? After all, separate private networks have been set up to serve the specific communication needs of many businesses. What advantages do you gain by converting the existing separate private networks to an Internet-based VPN?

Ubiquitous Coverage

The Internet offers far wider coverage compared with the private data network infrastructures offered by telecommunication providers. Adding new destinations to a private network means adding new circuits. Unlike the Internet, which has public and private peering points all over the world, few interconnection agreements exist between the service providers. Thus, the coverage of a private network is limited.

The Internet, on the other hand, is a vast interconnection of heterogeneous networks. Any host connected to a network that is connected to the Internet is in turn connected to any other host connected to a network connected to the Internet.

Cost Reduction

Another advantage gained by using an Internet-based VPN is cost reduction based on the system's economy of scale. Simply put, it eliminates the need to purchase and maintain several special-purpose infrastructures to serve the different types of communication needs within a corporation.

Security

VPNs use cryptographic technology to provide data confidentiality and integrity for the data in transit. Authentication and access control restrict access to corporate network resources and services.

In traditional private networks, the security of the data during transit relies on the telecommunication service provider's physical security practices for data confidentiality. For example, frame relay networks have no built-in provision for encrypting data frames. Consequently, data frames, if intercepted, can be easily decoded. In VPNs, you need not trust the perceived physical security of the telecommunication service provider. Instead, data is protected by cryptography.

E-Commerce

More and more business is being conducted using the Internet. Electronic commerce is not only a major new method of retailing merchandise (called "B2C" for business-to-consumer e-commerce), but it is also a way for businesses to trade goods and services among themselves (called "B2B" for business-to-business e-commerce). Interconnectivity of businesses is essential, and the Internet is the logical choice for the interconnection technology.

E-commerce must be secure. Private networks use physical separation for security, but it is impractical to have a separate infrastructure for each customer or B2B partner. Therefore, a closed, inflexible private network is not well suited for supporting e-commerce. A public infrastructure is more flexible but lacks security. VPNs provide both interconnectivity and security.

1.3 The VPN Market

VPNs, in one form or another, are becoming a crucial component of corporate networking solutions. Corporate networks use the Internet for various forms of business communication, and, for many organizations, VPN technologies are used to conduct private and commercial activities. Indeed, the trend is to migrate existing private corporate networks to Internet-based VPNs, and newly created corporate networks are increasingly using the Internet as their shared infrastructure.

To meet these needs, there has been tremendous growth in VPN offerings, which we separate into two categories: VPN products and VPN services. We will also discuss barriers to the development and deployment of VPN products and services.

VPN Products

VPN products are the hardware and software that make VPNs possible. One way to classify VPN products is based on how the product protects corporate resources. A *VPN gateway* is a stand-alone device that enables authorized access to the protected network resources. The resources are not located on the same physical device with the VPN gateway. A *VPN client*, on the other hand, is installed on the same network device it is supposed to protect. Usually, the client is a software package installed on the host computer.

VPNs require at least two cooperating devices. The communication path between these devices can be viewed as a secure tunnel across an insecure Internet infrastructure. Wrapped around this tunnel is a series of functions, including authentication, access control, and data confidentiality and encryption.

Depending on how these functions are implemented, VPN products can also be separated into two categories:

- **Software-based** Special software is added on top of a general computing platform, such as a UNIX or Windows operating system, to enable the use of VPN functions.
- **Hardware-based** Special hardware augmented with software is used to provide VPN functions. Sometimes, VPN functions are added to a hardware-based network device such as a router or a firewall. In other cases, VPN

functions are built from the ground up, and routing and firewall capabilities are added.

Many vendors network equipment are adding VPN products to their product lines. Some vendors add VPN functions to their existing products, and others build specialized VPN devices from the ground up.

VPN Services

A corporation can either create and manage the VPN itself or purchase VPN services from a service provider. When a corporation creates its own VPN, it obtains only IP connectivity from the service provider. All other functions pertaining to the virtual private network service are managed by the corporation. These functions include the purchase and installation of equipment, network monitoring, and configuration management.

In the case of a contracted VPN service, the service provider attempts to mask the complexity of the VPN service. The idea is that the service provider, by virtue of being in the network service business, has the expertise to manage the Internet-based VPN. Because the service provider may operate networks for many different corporations, it has the advantage of economy of scale and can run a network operations center with 24×7 availability. This may not be economically feasible for a small company with limited resources. Additionally, Internet service providers (ISPs) control the network infrastructure, so they are better equipped to deal with problems that arise within the network infrastructure.

When you purchase a VPN service, one issue is who retains control of the network. The data being sent through the VPN is critical. Putting such critical data in the control of a service provider can be sensitive for the corporation. A trust relationship must exist between the service provider and the corporation.

Another issue is the quality of the service. Specific performance guarantees, called *service level agreements* (SLAs), are negotiated between the service provider and the customer. Various measures can be taken when the SLAs are not met.

VPN Barriers

There are several barriers to widespread deployment of VPNs. First is the lack of interoperability of IPsec implementations. IPsec is the Internet Engineering Task Force's (IETF) security standard for IP. Although IPsec was standardized in November 1998, many vendors' implementations of these complex protocols have not yet achieved full interoperability with each other, even if they claim to be IPsec-compliant. Also, the Public Key Infrastructure for the Internet (PKIX) standard—X.509 authentication adapted for use in the Internet—is still moving slowly in the IETF working group. (For more on X.509, see Chapter 6.) This important standard provides a strong certificate-based authentication mechanism, but it is not expected to be widely available in the Internet in the near term.

Second, the lack of widely used quality of service (QoS) standards, as well as the sparse deployment of QoS-capable infrastructures, has made it very difficult to guarantee the quality of Internet connectivity, especially when traffic traverses the infrastructures of multiple ISPs. Many time-sensitive applications require certain guarantees to function correctly. This is not a new problem, and several proposals are on the table, but none has established itself as a clear winner.

Third, the Internet infrastructure is still largely focused on providing connectivity and does not yet offer services beyond connectivity. Security services in support of VPNs must be constructed from additional hardware and software components. Furthermore, computer operating systems in general, and Microsoft Windows in particular, do not yet contain mature built-in security functionality.

1.4 VPN Technologies

Several key technologies are employed by VPNs, as illustrated in Figure 1-3. A virtual or overlay network relies on *tunneling*, a concept discussed in great detail in this book. With tunneling, a network looks as if there is a set of simple links between the several sites when, in reality, the links may be a set of complex routes through the Internet. *Authentication* is the process whereby the identities of the VPN users and devices are verified. *Access control* provides ways to ensure authorized use of private corporate resources. The data transmitted over the VPN must

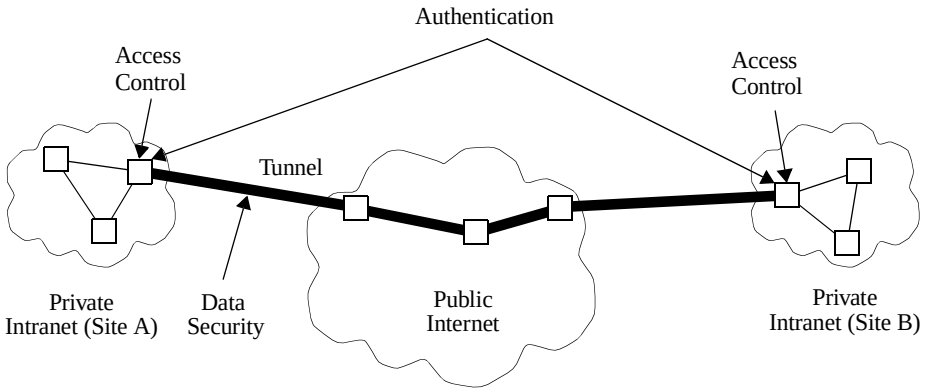

Figure 1-3 Key technologies of VPNs

be obscured from all but the intended recipient: cryptographic technologies provide *data security* for VPNs.

The technologies that form the building blocks of VPNs are the topics for chapters in Part II of this book.

Tunneling

Tunneling is defined as the *encapsulation* of a certain data packet (the original, or *inner* packet) into another data packet (the encapsulating, or *outer* packet) so that the inner packet is opaque to the network over which the outer packet is routed. The encapsulation is done in such a way that one or more protocol layers are repeated. For example, when an IP packet is put into a User Datagram Protocol (UDP) packet (which is itself put into an IP packet), we say that UDP is tunneling IP.

The need for tunneling arises when it is not appropriate for the inner packet to travel directly across the network for various reasons. For example, tunneling can be used to transport multiple protocols over a network based on some other protocol, or it can be used to hide source and destination addresses of the original packet. When tunneling is used for security services, an unsecure packet is put into a secure, usually encrypted, packet.

Two components can uniquely determine a network tunnel: the endpoints of the tunnel and the encapsulation protocol used to transport the original data packet within the tunnel. In most cases, a tunnel has two endpoints: one where the tunnel starts (encapsulation) and one where the tunnel ends (decapsulation). In the multicast case, a tunnel can have one starting point and multiple ending points, but we do not consider those tunnels in this book. In a VPN, authentication and access control decisions are made, and security services are negotiated and rendered, at these endpoints. Depending on the required services, different encapsulation protocols can be chosen.

Chapter 4 discusses details regarding tunneling technologies and the various standard tunneling protocols. One of the more important protocol suites, IPsec, is addressed in Chapter 5.

Authentication

Before any communication can be called private, each party must know the identity of the other. The same holds true for secure network communication: One network system must make sure that the other network system is the intended correspondent. The process of such identity verification is called *authentication*.

Authentication ensures that the data is indeed coming from the source it claims to be. In the traditional circuit-switched network, authentication is carried out only during the circuit provision phase. Verification of identities is implicit because the circuit is provisioned by the administrators, who presumably know the communicants. After the establishment of the circuit, the circuit, and thus the assumption of identity, remains in place indefinitely until it is explicitly torn down through another provisioning action.

When a secure VPN tunnel is established, the two endpoints (where the security service is negotiated and rendered) must authenticate each other.

Authentication methods can be broadly categorized into two kinds: two-party authentication and trusted third-party authentication. These methods are discussed in Chapter 6; Chapter 7 focuses on PKI, or public key infrastructure, the most important trusted third-party method.

Access Control

When the authentication process is completed, the communication entities can decide whether to allow the communication session to continue or to reject the session. In the case of VPNs, one purpose of secure communication is to allow authorized access to resources. When an access request is presented, the resource (or the security proxy responsible for safeguarding that resource) makes a decision as to whether to allow the access request to proceed. Usually, this *access control* procedure is performed at the endpoints of the tunnel.

An access control process has two aspects. The first aspect is the information on which the access control decision is made. Usually, this information includes the identity of the entity that is requesting access, the resources to be accessed, and rules governing the access. The complete identity information of the requester and resources can be presented at the time of the request, or it can be derived from stored information using data or credentials provided at the time of the request, such as username or IP address or both.

The second aspect is how the access control decision is made based on the information available. For example, the decision-making process can be carried out entirely at the location where the security service is negotiated and rendered, or the system can query a separate policy server for such decision. Having all the policies administered at a centralized server can make it easier to manage.

Chapter 8 addresses access control and explains how it can be performed in a VPN scenario.

Data Security

Data security touches on all VPN technology building blocks. Overall security is only as secure as its weakest link. If one link in the sequence that leads to the delivery of data is not secure, the entire process is not secure.

Because VPNs use a shared infrastructure to transport private traffic, it is possible for the data to be intercepted and even altered by others during its transit over the public infrastructure. Thus, strong encryption and data integrity should be applied to every data packet to make it opaque to interceptors. In addition, the packet delivery system should also be able to guard against replay attacks. In a *replay* attack, the attacker simply retransmits a previously captured packet. Because the original packet was authentic, it can pass all the cryptographic

checks. If the receiver inserts a replayed packet into the received data stream, the entire message becomes incorrect, thus denying normal service.

Having strong encryption algorithms and well-conceived keys helps to ensure that the integrity of the data cannot be compromised by the interceptor without knowledge of the cryptographic keys. Changing the encryption in the middle of a communication session—a technique sometimes called *over the air rekeying* (OTAR)—helps to further the guard against attacks on the keys.

The basic concepts of cryptography are discussed in Chapter 2. Chapter 5 discusses the IPsec protocol suite, in which data security techniques are employed for IP-based networks.

1.5 VPN Solutions

In building a house, simply having the right materials does not mean that the house is built well. Similarly merely having the relevant technological building blocks is not enough to construct a virtual private network. A correct solution requires the right materials, and they must be put together in the right way.

Determining which kind of functions to provide at the tunnel endpoints, and how to implement these functions, is central to creating a VPN solution. When the VPN functions are implemented and integrated into an Internet device (e.g., a gateway), that device becomes a VPN device.

A VPN solution consists of multiple, appropriately configured VPN devices that are placed in the appropriate locations within the network. As with any network, after all the VPN devices are installed and configured, the network should be continually monitored and managed.

These aspects of an overall VPN solution are the topics of the chapters in Part III of this book.

VPN Gateways

The most common VPN device is the VPN gateway, which acts as the gatekeeper for network traffic to and from protected resources. Tunnels are established from the VPN gateway to other appropriate VPN devices serving as tunnel endpoints.

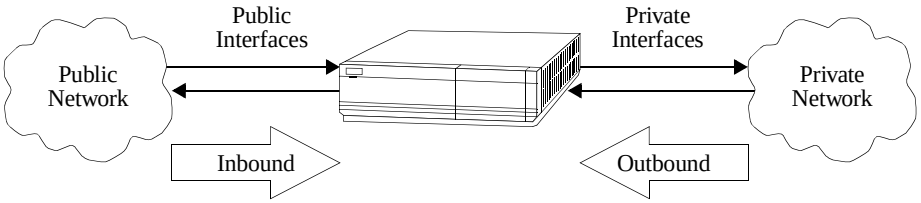

Figure 1-4 VPN gateway interfaces

A VPN gateway is usually located at the corporate network perimeter, and it acts on behalf of the protected network resources within the corporate intranet to negotiate and render security services. The gateway assembles the tunneling, authentication, access control, and data security functions into a single device. The details of how these functions are integrated within a VPN gateway are specific to a vendor's implementation. Sometimes these functions can be integrated into existing router or firewall products. Sometimes a VPN gateway can be a stand-alone device that performs pure VPN functions, without firewall or dynamic routing exchange capabilities.

In general, a VPN gateway has two or more network interfaces (see Figure 1-4). At least one network interface connects to the unsecure network; interfaces of this type are usually referred to as the public, or external, interfaces of the VPN gateway. The VPN gateway also has one or more network interfaces connecting to the secure corporate intranet, usually referred to as private, or trusted, interfaces.

The data traffic coming from the public interfaces is termed *inbound* traffic. Because the inbound traffic is from the unsecure network, it is thoroughly examined according to the security policies. Usually, only the traffic from the established secure tunnels should be processed by the VPN gateway. If no secure tunnel is found, the traffic should be dropped immediately. Depending on the implementation, an alert or alarm can be generated to notify the network management station.

At least one exception exists. The gateway must process traffic whose purpose is to negotiate and establish these tunnels according to the negotiation protocols. Other exceptions may apply, as in responding to certain diagnostic requests. In

general, however, the fewer exceptions there are, the more secure a gateway can be.

The data traffic coming from the private interfaces and exiting to the public interface is termed *outbound* traffic. Because the outbound traffic is from the private network, it is deemed secure a priori, even though many network attacks are generated within a corporate network. The outbound traffic is examined according to a set of policies on the gateway. If secure tunneling is required for the traffic, the VPN gateway first determines whether such a tunnel is already in place. If it is not, the gateway attempts to establish a new tunnel with the intended device—either another VPN gateway or simply some other device with secure tunneling capabilities. After the tunnel is established, the traffic is processed according to the tunnel rules and is sent into the tunnel. The traffic from the private interface can also be dropped if a policy cannot be found. Depending on how quickly a secure tunnel can be established, the VPN gateway may buffer the outbound packet before the secure tunnel is in place.

A VPN gateway implements some or all of the VPN technologies mentioned in the preceding section. Which functions are selected and how they are implemented is largely the choice of the device implementor. Chapter 9 is devoted to the details of issues surrounding the design, implementation, and evaluation of VPN gateways.

VPN Clients

The VPN client is software used for remote VPN access for a single computer or user. Unlike the VPN gateway—which is a specialized device and can protect multiple network resources at the same time—the VPN client software is usually installed on an individual computer and serves that computer only.

Generally, VPN client software creates a secure path from the client computer to a designated VPN gateway. The secure tunnel enables the client computer to obtain IP connectivity to access the network resources protected by that particular VPN gateway.

VPN client software also must implement the same functions as VPN gateways—tunneling, authentication, access control, and data security—although these implementations may be simpler or have fewer options. For example, the VPN software usually implements only one of the tunneling protocols. Because the remote com-

puter does not act on behalf of any other users or resources, the access control can also be less complex.

Unlike the VPN gateway, in which all of the gateway's hardware and software is geared toward the VPN functionality, VPN client software is usually just an application running on a general-purpose operating system on the remote computer. Consequently, the client software should carefully consider its interactions with the operating system.

Today, most remote access is achieved through telephone dial-up. As broadband access (e.g., cable modems and digital subscriber lines, or DSL) deployment becomes more common, dedicated high-speed remote access will be popular. In many dial-up cases, the modem speed is relatively slow, so the VPN client software must perform IP data compression before encryption to increase the bandwidth performance.

One of the important concerns regarding VPN client software is the simplicity of installation and operation. Because client software is expected to be deployed widely on end users' machines, it must be easily installed and easily operated by regular computer users who may not know much about the operating system, software compatibility, remote access, or VPNs. A VPN gateway, on the other hand, is usually deployed on the company's corporate network site and is managed by information technology professionals.

Authentication on the client software can take several different approaches. The VPN software may have its own authentication mechanism, or it may inherit the authentication scheme from the operating system. It is often desirable to have a single sign-on for all the services on the desktop. Either digital certificates or some type of shared secret authentication scheme can be used.

In some cases, VPN client software can be designed to allow only specific application access to a server—for example, secure shell applications or a program working within a browser session. These types of clients are sometimes referred to as *thin* VPN clients.

Chapter 10 discusses the many details regarding the implementation and deployment of VPN client software.

VPN Network and Service Management

All networks require management to remain in good working condition; occasionally, the network topology and configuration must be changed according to business and application requirements. This management is especially important in the VPN scenario. Owing to the dynamic nature of VPNs, the networks require continuous network and service management.

We specifically use the term *service management* because, in many cases, a VPN is a service offered from telecommunication carriers or ISPs to their customers. The customers of the VPN service do not directly manage the VPN but rather view it as just another network service.

It is important for the service provider to guarantee a certain level of service quality according to its SLA. SLAs provide a factor that differentiates service providers. In the VPN space, SLAs are usually related to infrastructure availability and performance metrics. For example, in dedicated site-to-site intranet VPN, a certain level of availability (e.g., 99.9%) can be applied to the VPN gateways, and a certain amount of average latency between the VPN gateways can also be provided. In a remote access VPN, the modem availability and connection speed are the subjects of performance monitoring and guarantee.

Because a VPN is a secure network service, it is paramount for the network to be managed in a secure fashion. A network cannot be secure if its configuration can be altered without security checks and strong authentication mechanisms.

Traditionally, network management relies on the management information base (MIB) and the Simple Network Management Protocol (SNMP). Standard MIBs are geared toward device-specific information management. A VPN, on the other hand, by its nature, relates to more than one device. Therefore, new MIBs need to be developed (some are already under development) to address the VPN management problem.

A good VPN service management solution should take into consideration (but not be limited to) all the aspects we have mentioned. Chapter 11 addresses the many issues regarding VPN network and service management.

VPN Directions

VPN technology is still emerging, and, as with any new technology, it will experience continued developments. The performance of VPN devices will certainly improve dramatically with advances in component technology. Innovative ways to design and implement VPN functions will also be invented. Although it is impossible to predict the future, we can certainly make some observations on business and technology trends.

One such trend is the integration of VPN and firewall functions into a single device. In this way, you can manage security in a unified way rather than having separate policies and interfaces for the two devices. Additionally, incorporating routing and quality of service features into the VPN device will make it even more versatile.

Another important trend is the move toward adding intelligence to the network. In the telephone network, intelligence resides mostly within the switched telephone network; the telephone itself is a simple device. Currently in the Internet, the connected devices are computers that have substantial processing power. The routers and switches merely forward the packets without knowing what is inside them or how to process them accordingly. Having intelligence in the network enables service providers to offer value-added services to customers. We anticipate that VPNs will be among the first services to be supported by these more intelligent networks.

These and other trends in the VPN arena are discussed in Chapter 12.