

8 Networks Security, Internet Security and Firewalls

8.1	INTRODUCTION.....	8-1
8.2	PHYSICAL NETWORK SECURITY.....	8-3
8.2.1	Overview.....	8-3
8.2.2	LAN Access Security.....	8-4
8.2.3	A Note about Hub versus Switch Security.....	8-5
8.2.4	Remote Access Security – General Considerations	8-5
8.2.5	Dial-Up Network Security.....	8-6
8.2.6	AAA Servers and Remote Access Dial-in User Service (RADIUS)	8-9
8.3	PUBLIC INTERNET SECURITY	8-11
8.3.1	The Public Internet at a Glance	8-11
8.3.2	Public Internet Security Issues	8-13
8.3.3	ISPs and NSPs as Players in the Public Internet Security.....	8-14
8.3.4	Internet Security Through VPNs and IPSec	8-17
8.4	CORPORATE INTRANET AND EXTRANET SECURITY.....	8-22
8.4.1	Overview of Corporate Network Security Issues.....	8-22
8.4.2	What are Intranets and Extranets?	8-22
8.4.3	Corporate Intranet and Extranet Architectures for Security	8-23
8.5	FIREWALLS – PROTECTING THE ENTERPRISE PERIMETER	8-26
8.5.1	Overview.....	8-26
8.5.2	Firewall Configurations	8-27
8.5.3	Firewall Administration, Design Issues and Policies	8-30
8.6	SHORT CASE STUDIES AND EXAMPLES	8-32
8.6.1	Combining Network and Physical Security.....	8-32
8.6.2	beBetter Networks Chooses a FireWall.....	8-33
8.6.3	Secure Automated Change Management for a University Network	8-33
8.6.4	From Site Finder to Internet Ownership	8-34
8.6.5	Using SATAN to Reduce Network Security Exposures.....	8-34
8.6.6	Digital Networks Tie Together National and Local Security	8-35
8.7	SUGGESTED REVIEW QUESTIONS BEFORE PROCEEDING.....	8-36

8.1 Introduction

Network security has assumed an increasingly important role for enterprises in this age of universal connectivity where customers, employees, business partners, and suppliers are interconnected in cyberspace. The increased reliance on the Internet, in particular, has

introduced several security risks due to the almost unlimited supply of hackers, intruders, and eavesdroppers. Internet vulnerabilities were highlighted by the 1998 annual report from the Computer Emergency Response Team (CERT) that listed over 1300 reported security incidents on the Internet affecting nearly 20,000 sites [CERT99]. Due to the increased worms, viruses and distributed denial-of-access attacks (estimated to be rising 120 percent every year), some pressure is building to let the private sector run the Internet [Cooper 2003].

Although network security and Internet security are often intermixed, we will attempt to differentiate between the two by noting that the Internet is a network of networks; i.e., it is a network that interconnects several physical networks. Thus each physical network needs to be protected, as well as the interconnection between these networks. Technically, the Internet consists of layer 3 and 4 protocols (TCP/IP) that support data, voice and video applications over almost all types of physical networks at layers 1 and 2 (see Figure 8-1). Thus network security needs to be considered at the physical network layers plus the Internet (TCP/IP) layers. The topics in this chapter include physical network security with an emphasis on remote access networks (Section 8.2), and the Internet security that covers public Internet (Section 8.3) and corporate intranets and extranets (Section 8.4). An important issue in network security is to protect enterprise assets from the intruders and hackers roaming the Net¹. For this, “Firewalls” are erected to regulate and control the traffic around the enterprise perimeter (Section 8.5).

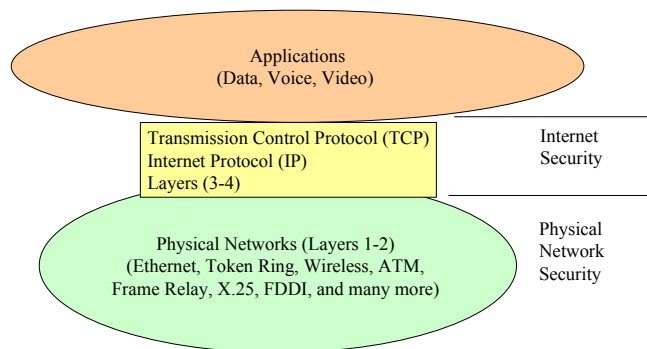


Figure 8-1: Network and Internet Security

Chapter Highlights

- Network security can be discussed in terms of physical network security and Internet security
- Physical network security includes remote access networks. The most common technologies used for security of remote access networks are RADIUS servers and VPN.
- Internet security includes public Internet and corporate intranet and extranet security.
- Increased attacks on the Internet have fueled the debate about making the public Internet privately owned.
- VPN and IPSec are the most widely used technologies for Internet security.
- Extranet security is managed by the extranet owners.

¹ In the popular press, the term “the Net” is used to indicate the Internet.

- Firewalls are erected to regulate and control the traffic around the enterprise perimeter.
- Firewalls provide many options and configurations to implement security.

8.2 Physical Network Security

8.2.1 Overview

Security at the physical network level involves protecting the information that is transported over wired or wireless communication links. We will discuss wireless security in the next chapter, so let us concentrate on the wired physical links such as coaxial cables, dial-up lines, T1 lines, fiber links, etc., which transport the local and wide area network traffic. Specifically, this involves:

- Local area network (LAN) access – how you will access services over a LAN in your office
- Remote access – how you will access your remote office services from your home over a dial-up or DSL/cable modem connection

These topics are discussed in this section. The higher-level issues of accessing the Internet over the physical networks are discussed in the next section because, as indicated above, Internet security needs to be viewed at higher layers than the physical network layers. It is useful here to revisit the tradeoffs in higher versus lower layers of security. To illustrate this, let us assume that you are in a branch office and you need to access your email from a remote corporate site and you want this access to be private. The main tradeoffs of higher versus lower layers of security are:

- Lower-layer security secures a physical network segment (e.g., a LAN segment). But once you leave that network segment, your traffic is traversing on other physical network segments that may not be secure. On the other hand, securing a given network segment protects *all* traffic (phone, IP, FTP) on that segment. Thus if a LAN uses encryption, then all traffic on that LAN is encrypted. For example, if your branch office has a small LAN that uses encryption in the office, then you are protected in that office but not necessarily after the messages leave the branch office. To carry the encrypted traffic between the branch office and the corporate office, the corporate office must be able to decrypt the messages, or the messages have to be decrypted *before* leaving the branch office.
- Internet-layer security assures that the IP traffic is encrypted and authenticated. Even if lower layer network segments are not encrypted, the IP packets over these segments are encrypted. So the intruders cannot understand the data even if they can eavesdrop on it. Thus if a VPN is available between your branch office and the corporate office, then the traffic over the Internet is encrypted. However, *all* traffic – not just the email – will be encrypted over the VPN.
- Application security allows you to protect the traffic between clients and servers of specific applications. For example, SSL protects communications between Web browsers and Web servers, PGP secures email activities, and SET is used to secure

electronic transactions. Thus, if PGP is used between the corporate email server and branch office clients, then the email is protected even if the branch office has no protection or no VPN exists between the branch offices and corporate office.

In most cases, securing at all layers is not advisable because it adds significant overhead. The trick is to decide at what layer to use the security. We will revisit this issue several times again.

8.2.2 LAN Access Security

LANs were introduced in the marketplace in the early 1980s. In the early stages, there was no agreement on LAN technologies, and many LAN vendors were pushing proprietary solutions. For example, in the early 1980s, more than 50 vendors were marketing LANs on different devices, using different communication media, protocols, and topologies. The IEEE (Institute of Electrical and Electronics Engineers) 802 Committee on LANs was formed in 1980 to develop standards for LANs². Thanks to the work of this committee and market pressures for interoperability, the LAN landscape has simplified considerably, with only a few surviving vendors and technologies. The following figure represents a typical LAN configuration with possible paths that need to be protected. Typically an IEEE802.3 (Ethernet) LAN provides services to various users and may be connected to mainframes or other LANs. Access to a LAN is either through a direct connection to the LAN or through a LAN Server.




Figure 8-2: A Typical Local Area Network

Technically, LAN access is achieved through protocols that belong to the IEEE802 family such as the Ethernet (802.3), Token Bus (802.4), Token Ring (802.5), and others. Ethernet and its variants (Fast Ethernet, Gigabit Ethernet) are by far the most widely deployed. The IEEE 802.10 working group was formed in May 1988 to address LAN and MAN security. IEEE802.10 specified *Standards for Interoperable LAN/MAN*

² 802 indicates February of 1980, the date when the IEEE 802 Committee was formed.

Security (SILS) that are compatible with the IEEE 802 and the OSI standard specifications. Unfortunately, SILS has not been implemented widely, thus we do not discuss SILS. Vendor-specific security packages for LAN access are available from vendors such as 3Com and Cisco.

At the time of this writing, most wired Ethernet LAN security vulnerabilities are quite well covered by a wide range of suppliers. For this reason, we will not discuss them further here. As we will see in the next chapter, the situation is quite different in wireless LANs.

8.2.3 A Note about Hub versus Switch Security

Given that Ethernet LANs have become common in corporations, the decision of shared versus switched LANs is the primary design choice to be made in enterprise LANs. Shared LANs use a hub that provides a common cable to be shared by multiple workstations. For small networks, shared LANs with hubs work fine, but for larger networks, switched LANs are becoming popular. In a switched LAN, a switch is used to directly connect workstations. Figure 8-3 shows a typical Ethernet LAN design in which hubs as well as switches are used.

A security tradeoff between hubs versus switches is worth noting here. A sniffer in a hub can look at all the traffic on the hub because the hub is shared by all users. On the other hand, a sniffer on a switch can only intercept the traffic between the two participants. Thus switches tend to be a bit more secure. However, sniffers have been designed that attack the entire switch.

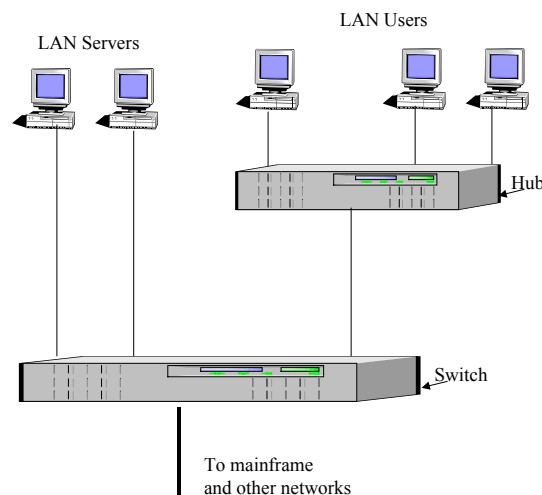


Figure 8-3: A Typical LAN Design With Hub and a Switch

8.2.4 Remote Access Security – General Considerations

Remote access is very popular at present because it enables enterprises to increase productivity of employees and provide 24x7 support to customers. Convenience of telecommuting and network-enabled computers in every office and home is driving the growth of remote access computing. In addition, an increase in e-commerce and B2B

trade is driving the need for remote access to new levels. Employees, customers, suppliers and business partners seldom work from the same place over an extended period of time. According to Gartner, more than 100 million people use remote access. The type of remote users are:

- Telecommuting employees who need fast access to corporate resources and enterprise applications for an extended period of time to do their work
- Travelling employees who need access, typically, to email
- Authorized partners and suppliers who need access to corporate inventories, catalogs, and other resources needed for B2B trade
- Customers who may need access to online order processing systems and order status
- Network and system administrators who may need to diagnose and correct network and computer system problems without having to drive to work

Despite its obvious benefits, remote access introduces many security and integrity problems. For example, an intruder could pose as a network administrator and shut down the network, or pose as a business partner to access proprietary information. Thus the network administrators need to be authenticated very carefully before being allowed to vary network devices online and offline, and business partners need to be rigorously screened before being allowed any access. These and other security features need to be embedded in the remote access technologies such as dial-up networks and authentication servers discussed in the next sections.

8.2.5 Dial-Up Network Security

Dial-up network access protocols include the *serial line IP (SLIP)* and the *point-to-point protocol (PPP)*. For all practical purposes, PPP is most commonly used for dial-up access. Thus, we discuss this level of security for dial-up networks using PPP.

Let us take the example of corporate network access security from remote users. Suppose that Pat is in Miami attending a conference (she says) and she wants to access the corporate mainframe and Intranet in Detroit to access her email. The Detroit corporate office also needs to be accessed from branch offices in Miami and Atlanta. It is assumed that the Detroit office has an *AAA (Authentication, Authorization, Accounting) server* that verifies the users before they can access the mainframe. The AAA servers, described in section 8.2.6, can be accessed from remote users through dial-up or other choices (see Figure 8-4). From an end-user point of view, here are the basic choices:

- Pat can use a regular public-switched-telephone-network (PSTN) or ISDN connection (if available) to connect to a modem pool or any other RAS (remote access service) in Detroit. The AAA server will authenticate and set up a PPP connection. This approach is simple but it can be expensive (long distance call) and is not secure because the data is not encrypted.
- The branch offices can use a network service provider (NSP) or Internet service provider (ISP) network to connect to the office in Detroit. Let us assume that the Atlanta office uses this choice as shown in Figure 8-4. The NSP/ISP will have a *point-of-presence (POP)* in Detroit and also in Atlanta so that the Atlanta folks can make a local call. This is cheaper than the first option but does not use encryption unless done at a higher level (i.e., your application encrypts the data before sending it over the line). The AAA server can do any authentication before completing the connection. Note that the POP is sometimes also referred to as a *NAS (Network*

Access Server), Network Access Point (NAP), Front-end Processor (FEP), or a dial-in server. To avoid confusion, we will use the term POP whenever needed.

- The branch offices can use a VPN (Virtual Private Network) for accessing the Detroit office (shown for Miami office in Figure 8-4). The main advantage of VPN is that the data can be encrypted by the VPN. To use VPN, your office has to subscribe to a VPN provider – i.e., the Detroit office is connected to a VPN point-of-presence (POP) for encryption. As a user, you will connect to a VPN POP in Miami over a local line and then use a secure ID card that will be used to authenticate you and then also encrypt your data. Since Pat is in Miami, she could also use the Miami POP for secure access to the Detroit office. Many network service providers such as UUNET support VPN POPs.

Naturally, the VPN choice is more secure. VPNs basically encrypt and encapsulate the network data into a “tunnel” that carries the encrypted data. There are different approaches to tunneling (see the sidebar “Tunneling Protocols For Security – A Technical Discussion” for more details). We will take a closer look at VPNs in the next section.

Let us take a quick look at the underlying technologies (you can skip this paragraph if it makes you drowsy). The network access authentication in all of these options is typically done at the link layer (layer 2). This means that the authentication takes place at the beginning of the session, with periodic re-authentications if desired. In addition, the POP is one hop from the client³. The client and POP use the link layer protocol over a point-to-point link (i.e., Ethernet or PPP). Once the client is connected to a POP, the POP and AAA server use the AAA protocol, which runs over IP. Thus the AAA server can be many hops from the client and the POP. As part of the POP and AAA exchange, the POP tells the AAA server what kind of access is desired (VPN, xDSL, 802.11, etc.) and the AAA server responds with appropriate attributes, including the type of tunneling to be used. Microsoft's PPTP (Point-to-Point Tunneling Protocol) is a common commercially available implementation that supports a variety of authentication and encryption options.

³ The technical reason for this one-hop is that the link layer only supports connections over one network segment (i.e., two machines on the same LAN or a telephone call. It is not routable because it is at a lower layer [layer 2] and the routing is done at layer 3 and above).

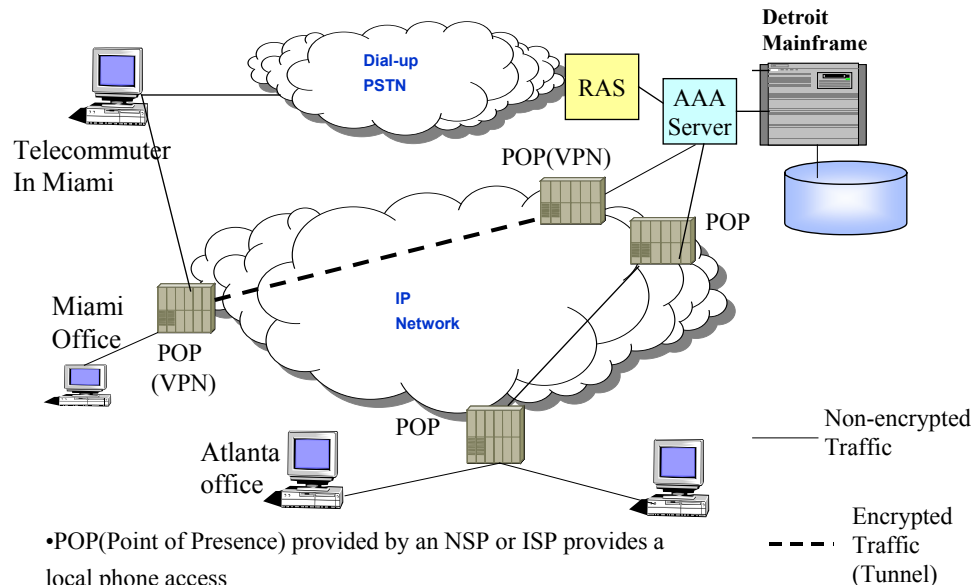


Figure 8-4: Remote Access Options

Tunneling Protocols For Security – A Technical Discussion

Simply stated, a tunnel provides a path for different types of payloads. For example, an IP packet can be used as a tunnel to carry different types of network data (802.4, 802.5, etc.). VPN networks consist of authenticated and encrypted tunnels over a shared data network (typically, an IP network). The tunnels are set up between a POP and a tunnel terminating device on the destination network.

VPN tunnels can be developed to carry lower or higher layers of network data (see Figure 8-5). A VPN tunnel can encapsulate network layer data (IP, IPX) inside PPP and then encapsulate the entire package inside a tunnel protocol (typically, IP – but it could be ATM). This is commonly known as “layer 2 tunneling (L2T)” because the passenger is layer 2 (PPP). Alternatively, the network layer data can be directly encapsulated in a layer 3 tunneling protocol such as 3Com's Virtual Tunneling Protocol (VTP). This is called layer 3 tunneling because the passenger is at the network layer (IP or IPX). This is shown in the figure below.

Which approach is better? There seems to be a consensus that layer 2 tunneling protocol as specified by the IETF Working Group L2TP is the best way to proceed [Oppliger 2000]. Microsoft's PPTP (Point-to-Point Tunneling Protocol) is a common commercially available implementation of L2TP (L2T protocol). MS PPTP supports a variety of authentication and encryption options.

The link (www.drizzle.com/~aboba/IEEE/BURP-BOF.zip) has a presentation at IETF on how authentication is done for network access and why this is most often handled at layer 2 (PPP, IEEE 802.1X) rather than at layer 1 (802.11) or at layer 3 (Mobile IP) or

higher. A detailed discussion of tradeoffs between different tunneling protocols can be found in Oppliger [2000].

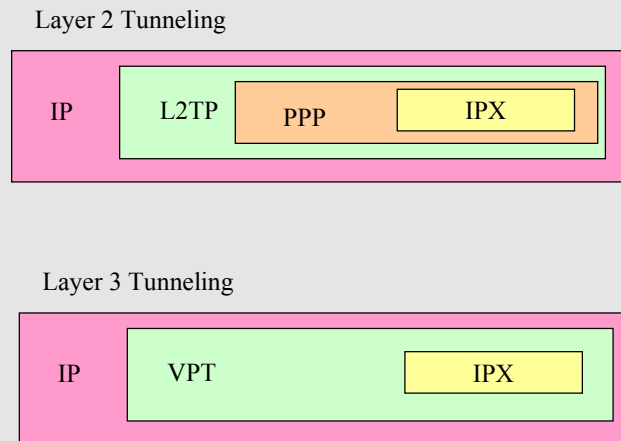


Figure 8-5: Layer 2 Versus 3 Tunneling

8.2.6 AAA Servers and Remote Access Dial-in User Service (RADIUS)

Authentication, authorization, and accounting (AAA) is used for controlling access to computer resources, enforcing policies, auditing usage, and providing the information necessary to bill for services. Authentication, authorization, and accounting services are often provided by a dedicated AAA server, a program that performs these functions. Examples of these services (see the previous chapter for more details) follow:

- Authentication identifies a user, typically by having the user enter a valid user name and valid password before access is granted. The process of authentication is based on each user having a unique set of criteria for gaining access. The AAA server compares a user's authentication credentials with other user credentials stored in a database. If the credentials match, the user is granted access to the network. Otherwise, network access is denied. As discussed in the previous chapter, asymmetric cryptography can be used for authentication by using a digital signature.
- Authorization ensures that the user does only what he/she is allowed to do (e.g., read but not update a database). After logging into a system, for instance, the user may try to issue commands. The authorization process determines whether the user has the authority to issue such commands. Authorization enforces the policies that determine the types of activities, resources, or services a user is permitted. Usually, authorization occurs within the context of authentication. Once you have authenticated a user, they may be authorized for different types of access or activity.
- Accounting measures the resources a user consumes during access. This can include the amount of system time, disk space used, network connection time, or the amount of data a user has sent and/or received during a session. Accounting is based on

logging of session statistics and usage information and is used for authorization control, billing, trend analysis, resource utilization, and capacity planning.

Commercially available AAA servers are server programs that are installed on a computer (e.g., a Unix box) to handle user requests for access to computer resources. For each user request, they invoke the needed authentication, authorization, and accounting (AAA) services. The AAA server typically interacts with network access and gateway servers and with databases and directories containing user information.

Over the years, several AAA servers with proprietary services have been deployed. The current standard by which devices or applications communicate with an AAA server is the **Remote Authentication Dial-In User Service (RADIUS)**. The RADIUS specifications define the authorization, authentication, and accounting features in great detail (see the sidebar “Sources of Information for RADIUS”). RADIUS is a widely deployed protocol to enable network access authentication, authorization and accounting (AAA). Numerous free as well as industrial strength RADIUS-based AAA servers are available from a large number of vendors such as Cisco, Cistron, AdvancedRadius, and others. However, some problems with RADIUS have been uncovered. Although there are numerous problems with RADIUS, including issues with security and transport, it is likely that RADIUS will continue to be widely used for a while because it is simple, efficient and easy to implement (and there are no other choices). The main advantage of RADIUS is that it can fit into the most inexpensive embedded devices due to its simplicity. The main issues with RADIUS are:

- Accounting: RADIUS runs on UDP and does not include retransmission or accounting record retention policy, and does not support application-layer acknowledgments or error messages. This makes RADIUS accounting unreliable for usage-based billing services, particularly in inter-domain usage (such as roaming) where there can be substantial packet losses. As a result, usage-based billing is often done with SNMP (Simplified Network Management Protocol) today.
- Security. RADIUS includes its own application-layer integrity protection and authentication, as well as confidentiality services. Initial RADIUS standards (RFC 2865) specified that RADIUS Access Requests need not be authenticated or integrity-protected. This raised security concerns. The situation was improved somewhat in RFC 2869, which requires that all messages involved in a conversation include authentication and integrity protection. RADIUS security has also been poor when dealing with cleartext password authentication. RFC 2865 has specified improvements for this. Basically, RADIUS security is improving, slowly but surely. A possible approach is to run RADIUS over IPsec. Unfortunately, many embedded systems do not run IPsec, so RADIUS/IPsec is not widely used today.

Sources of Information for RADIUS

www.ietf.org/rfc/rfc2865.txt – RADIUS authentication (Proposed Standard, RFC 2865)
www.ietf.org/rfc/rfc2866.txt – RADIUS accounting (Informational, RFC 2866)
www.ietf.org/rfc/rfc3162.txt – RADIUS for IPv6 (Proposed Standard, RFC 3162)
www.untruth.org/~josh/security/radius/radius-auth.html – RADIUS security analysis
www.ietf.org/internet-drafts/draft-ietf-aaa-transport-05.txt – Diameter Transport Specification (IETF draft, work in progress, discusses how to make AAA transport

really reliable)

8.3 Public Internet Security

8.3.1 The Public Internet at a Glance

As mentioned in the previous chapter, the *Internet* is a network based on the TCP/IP protocol stack. The term “Public Internet,” or just Internet, is used to refer to a large collection of TCP/IP networks that are tied together through network interconnectivity devices such as routers and gateways that are spread around the globe and are not owned by a single entity. The Internet model, shown in Figure 8-6, shows the typical Internet stack. Although there are several players in the Internet, IP (Internet Protocol) is at the heart of the Internet. In fact, the Internet is defined as a network of networks that is supported by the Internet Protocol (IP). The computers on the public Internet have publicly-known Internet Protocol (IP) addresses that are used to exchange information over the public Internet. An IP address serves the same purpose in the Internet as a telephone number does in the telephone system – it provides a unique address for someone to contact you. Millions of computers are interconnected through thousands of networks that use different underlying physical network technologies in different parts of the world. All these computers and networks are tied together through a global IP network with unique IP addresses.

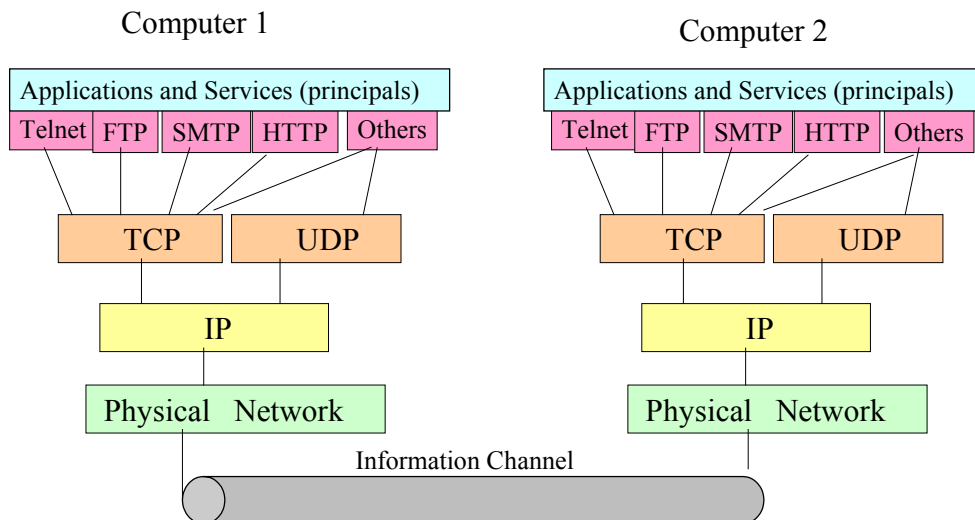


Figure 8-6: The Internet Stack

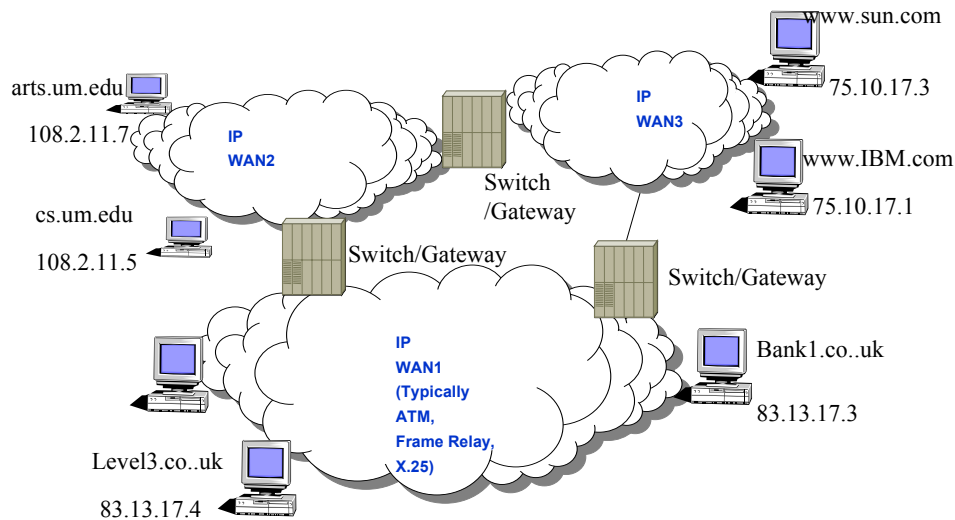
While IP addresses are important for the technical architecture of the Internet, i.e., Domain Naming Service (DNS) is of key importance for end users because DNS helps the users to locate different resources by using names such as www.ibm.com instead of IP addresses such as 22.111.292.13. DNS defines hierarchical naming structures which



are much easier to remember than the IP addresses. For example, the machine with an IP address of 135.25.7.82 may have a domain name of shoeshop.com. A user “mills” may have an email address mills@shoeshop.com. The DNS naming structures define the organization type, organization name, etc. The last word in the domain name identifies an organization type or a country. Consider, for example, the following domain names:

telcordia.com = commercial company Telcordia
 ibm.com = commercial company IBM
 um.edu = educational institution University of Michigan
 omg.org = organization OMG (Object Management Group)
 waterloo.ca = Waterloo University in Canada
 lancs.ac.uk = Lancaster University in UK
 ansa.co.uk = ANSA consortium in UK
 iona.ie = Iona Corporation in Ireland

The Internet uses a large number of domain name servers that translate domain names to IP addresses (the IP routers only understand IP addresses). Although many DNSs may exist, there are about a dozen root DNSs that contain all the DNS addresses.



- DNS (Domain Name Services) translates cs.um.edu to 108.2.11.5
- Telnet cs.um.edu = Telnet 108.2.11.5
- FTP cs.um.edu = FTP 108.2.11.5

Figure 8-7: Partial View of Internet

Figure 8-7 shows a conceptual and partial view of the Internet that was shown in the previous chapter. This Internet view shows three networks (a university network with two computers, a commercial company network, and a network in the UK). Each computer (“host”) on this network has an IP address and also has been assigned a domain name. The Internet is very heterogeneous (i.e., different computers, different physical networks.) However, to the users of this network, it provides a set of uniform TCP/IP services (TCP/IP hides many details). Let us use this simple Internet view to illustrate the key Internet capabilities from a security point of view.

Since the Internet is based on TCP/IP, the applications and services provided by TCP are also available on the Internet. From an end-user point of view, the following services have been, and still are, used very heavily on the Internet:

- Email
- Telnet
- FTP

Electronic mail on the Internet is based on the **Simple Mail Transfer Protocol (SMTP)**. This TCP-based protocol is the Internet electronic mail exchange mechanism. Email is still one of the most heavily used services on the Internet. Users have email addresses such as johnm@cs.um.edu, hevner@sun.com and howard@bank1.co.uk. If the Internet fails, your email cannot be delivered.

Terminal emulation is used to remotely log on to other machines. **Telnet** is used to provide terminal access to hosts; it runs on top of TCP. Let us assume that a user “joe” on cs.um.edu needs to remotely log on to the bank1.co.uk machine to run a program “directory.” The user would use the following steps (the steps are explained through comments marked /* */):

```
> telnet bank1.co.uk      /* invoke Telnet. Could have typed "telnet 85.13.17.3."*/
bank1> enter login: joe   /* prompt from bank1 for login ID. joe is ID */
bank1> password: xxxx    /* prompt from bank1 for password */
bank1> directory         /* run the program "directory" */
bank1> exit              /* quit telnet */
```

File transfer is used for bulk of data transfer over the Internet. The **File Transfer Protocol (FTP)** provides a way to transfer files between hosts on the Internet. Let us assume that a user “garner” on “sun.com” needs to transfer a file from the host arts.um.edu. The following steps would be used (the steps are explained through comments marked /* */):

```
> ftp arts.um.edu        /* invoke FTP. Could have typed "ftp 102.52..10.7"*/
arts> enter login: garner /* prompt from arts.um for login ID. garner is ID */
arts> password: xxxx     /* prompt from arts.um for password */
arts> get file1 file2     /* FTP file transfer command */
arts> exit (or quit)      /* quit FTP */
```

8.3.2 Public Internet Security Issues

Figure 8-8 shows the different security technologies used at different layers of the Internet. For the purpose of network security, we will concentrate on the first few layers (anything above TCP/UDP is at a higher layer than networks and will be discussed in later chapters). Let us start from the bottom layer, the physical network.

Since the public Internet consists of a very large number of physical networks that run under IP, the security of these physical networks is left to the owners. Since many alternate routes are followed by the public Internet, the failure of individual physical networks does not impact the overall public Internet. Similarly, IP is imbedded in a very large number of IP routers and these are also managed by the owners – the NSPs and the ISPs (see the next section). These service providers support VPN, IPsec and other techniques to keep their segment of users safe. If a large ISP such as AOL or NSP such as UUNET are assaulted, then a very large number of users would be affected.

It is naturally not possible to secure the millions of physical networks and the IP routers that comprise the public Internet; hence the public Internet is generally thought of being unsafe, and rightfully so.

There are some higher-level security issues also. As mentioned previously, the Internet uses a large number of domain name servers that translate domain names to IP addresses. Without these DNS servers, we could not access our email or websites. Although there are many DNS servers, there are about a dozen DNS servers at the core of the public Internet. The root DNS servers have been under constant denial of service attacks. For example, in October 2002, denial of service attacks brought 9 out of the 13 root DNS servers down, causing concerns about the public Internet and calls to make the Internet privately owned [Cooper 2003].

Denial of service attacks on individual applications on the public Internet can cause major problems because you cannot access your email or transfer files without these applications. Web technologies have made this even more serious. For many years, the Internet had been used mainly by researchers, teachers, scientists, students, and programmers to transfer files and send/receive electronic mail. These users relied on text-based commands to do their job. The Web is a set of services that run on top of the Internet to support Web browsing and surfing. The current advent of Web Services is making the corporate applications also reliant on the Internet. We will discuss Web security in a later chapter.

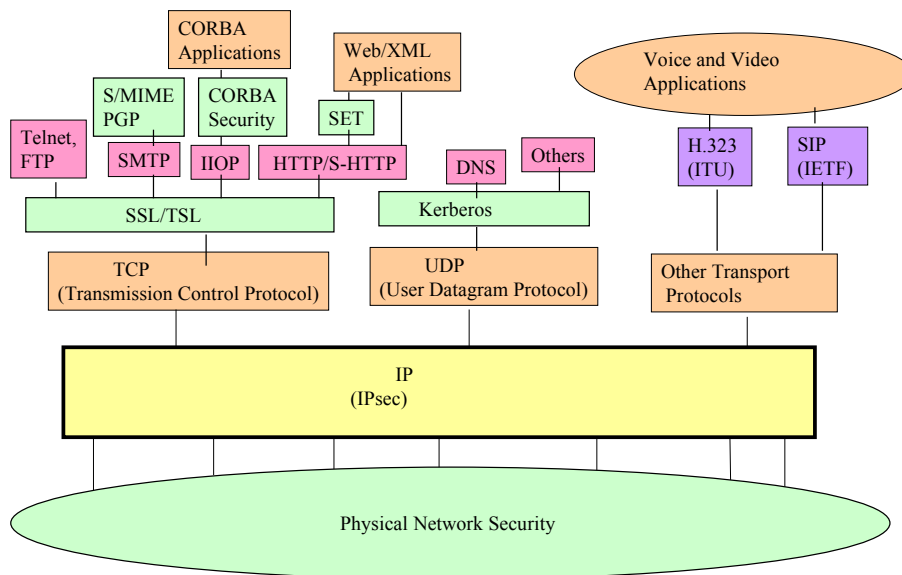


Figure 8-8: Security at Different Levels of the Internet

8.3.3 ISPs and NSPs as Players in the Public Internet Security

ISPs (Internet Service Providers) and NSPs (Network Service Providers) can play a major role in public Internet security by supporting different types of security features such as VPNs and secure dial-up services. These service providers can keep their segment of users safe and thus help to make the public Internet safe.

Let us look at different service providers before proceeding.

Simply stated, a service provider offers its customers a set of services based on an agreed-upon contract. The services can be business services such as physical site security, or technical such as Web hosting. Different service provider models, centered around the Internet, are quite popular to facilitate outsourcing. For example, businesses and consumers can rent services from service providers such as the following (see Figure 8-9):

- Network service providers (NSPs) that provide the network “pipe” (end-to-end network communication and routing services) for Ebusiness. Examples of NSPs are the telecommunications companies that include a variety of local exchange and long distance carriers.
- Internet Service Providers (ISPs) that support Web services and provide access to the public Internet. America Online is a well known example of an ISP.
- Platform Service Providers (PSPs) that provide the platform services (computing hardware, operating systems, basic middleware) needed to support e-commerce or other applications. PSPs in essence are similar to the old “computing centers” that provided the computing hardware/software for business applications. Due to the emphasis on ecommerce, PSPs are also referred to as CSP (commerce service providers). Examples of PSP/CSPs are Rightworks.com, CommerceOne, and Ariba.net.
- Application Service Providers (ASPs) host application components (mostly business-aware) that clients use over a wide area network. A very wide range of ASPs have emerged in recent years, with services that range from payroll to inventory control. For example, major software vendors such as SAP, Oracle, and Peoplesoft are becoming ASPs. We will discuss ASPs in more detail later.
- Business service providers (BSPs) that provide business services such as mail delivery, customer support, and building security.

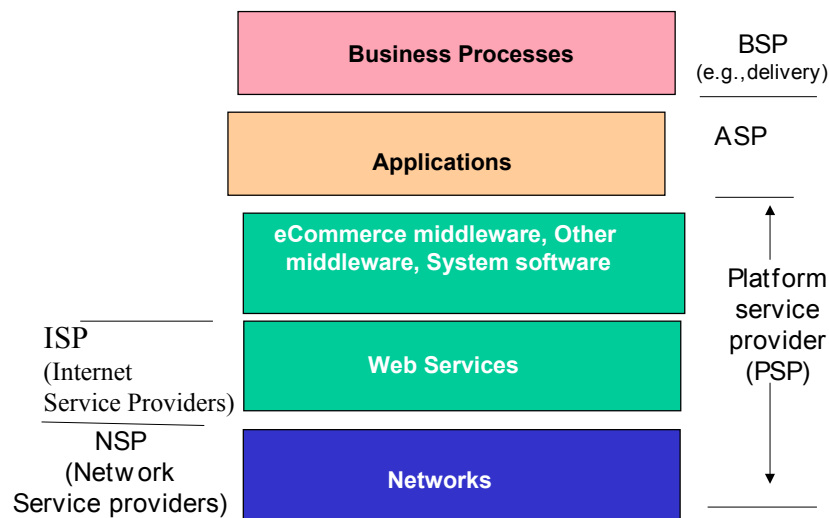


Figure 8-9: Service Providers

Let us consider Network Service Providers (NSPs) and the Internet Service Providers (ISPs) in some detail.

Network Service Providers (NSPs), also known as Internet Access Providers (IAPs), are the organizations that provide the physical network, i.e., give you a communication line and an access port on the Internet. You can think of IAPs as the local authorities that provide you with roads and signs to get you to the shopping malls. For dial-up users, an NSP provides several POPs (points of presence) that the users can dial into as a local call. An example of NSP is UUNET, which has POPs around the globe. I am a regular user of UUNET. A GUI shows me the phone number of the nearest POP (I just type in the name and country of the city). When I travel (I have traveled from South Africa to Budapest), I quickly locate a UUNET number and make a local call to reach my office computers.

Internet service providers (ISPs) go beyond the network pipe and offer Internet services such as email, Web hosting and Web surfing. Many ISPs also provide help in building Web sites. Many small Web content providers seek the help of ISPs to set up Web servers with appropriate security and backup/recovery. An ISP may rent a NSP (i.e., use an existing POP) or own a network, and thus may own POPs. Examples of ISPs are America Online and Asia online. Figure 8-10 shows a conceptual view and Figure 8-11 shows a physical view of ISPs and NSPs. From a security point of view, the ISP and the NSP can support various security services to provide security in the public Internet.

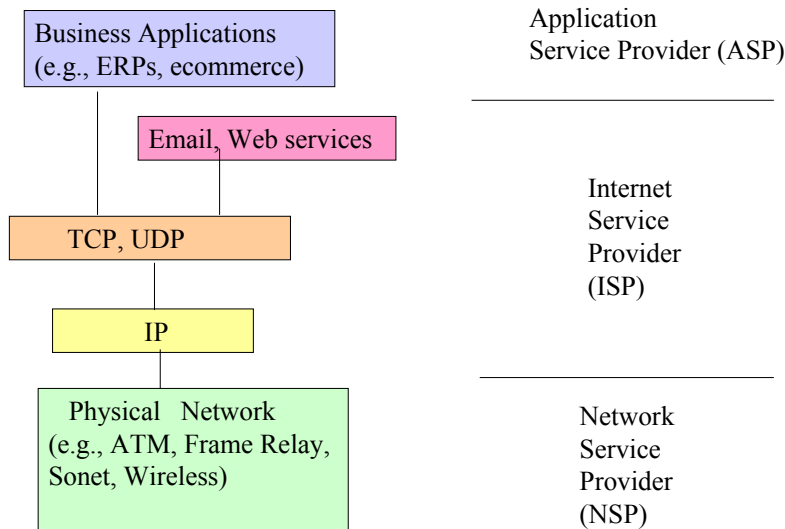


Figure 8-10: A Conceptual View of ISPs and NSPs

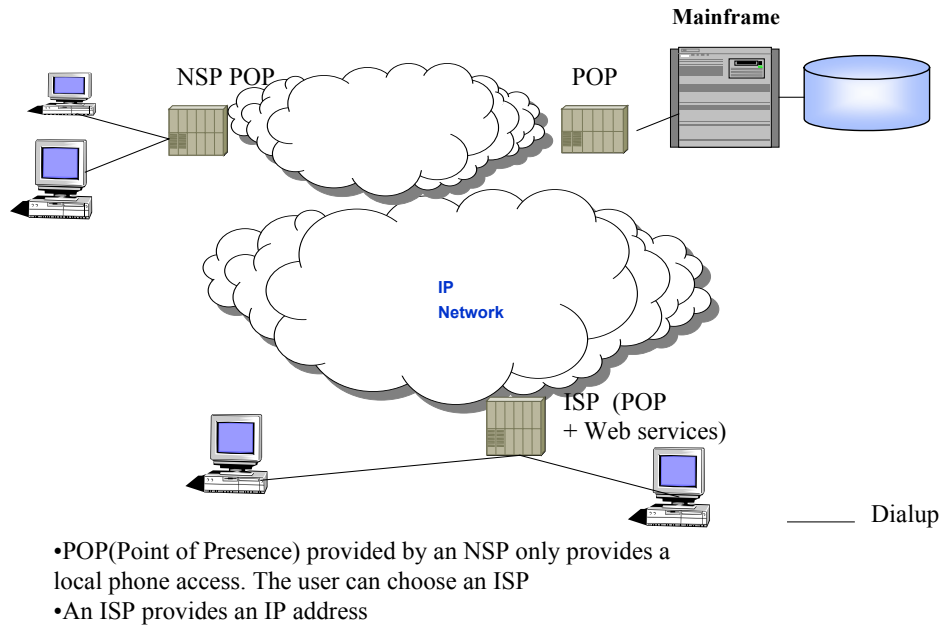


Figure 8-11: Physical View of ISPs and NSPs

8.3.4 Internet Security Through VPNs and IPSec

Virtual Private Networks (VPN), as described in a previous chapter, allow the users to setup a *private* network over a *public* network by using encryption. This makes the public network private – once your messages are encrypted, then the intruder cannot understand them even if he/she looks at them. IPSec provides standard protocols to assure that transportation of encrypted messages over a multitude of physical networks is smooth. We already reviewed VPNs and IPSec in a previous chapter. The following discussion recaps how VPNs and IPSec are used for Internet security and provides additional details of IPSec.

8.3.4.1 Virtual Private Networks (VPNs) and IPSec

A VPN, shown in Figure 8-12, provides dedicated, secure paths, or tunnels, over a network that is shared by other users. VPN networks consist of authenticated and encrypted tunnels over a shared data network (typically, an IP network). The tunnels are set up between a point of presence (POP), also called a network access point (NAP), and a tunnel terminating device on the destination network. Basically the POP server allows you to make a local call. A user initiates a dial-up session to a local POP, where a server authenticates the user and then establishes a tunnel through its Internet “cloud,” which terminates at the edge of the user's corporate network. Currently, a large number of companies offer VPN services. Examples are Shiva, telecommunication companies (e.g., Southwestern Bell and Nortel), and network service providers such as UUNET. Additional information about VPN can be found at the VPN Consortium website (www.vpnc.org).




Figure 8-12: A VPN

Most of the currently available VPNs are based on the IETF IPsec (RFC 2401) and related standards to transport encrypted messages over shared networks. IPsec is not, however, restricted to VPNs – a corporate LAN within a building can use IPsec by installing IPsec-compliant software on various routers. IPsec-compliant software encrypts and signs Headers and/or Data parts of an IP Packet and specifies security at the packet level, instead of the application level. It thus provides security without requiring changes to applications and is especially suitable for VPNs (see Figure 8-13). The principal feature of IPsec is that it can encrypt and/or authenticate all traffic at the IP level. Thus all applications that use IP (email, Web access, file transfer, etc.) can be secured.

VPNs are commonly used to overcome wireless link security weaknesses. However, there are some drawbacks of VPNs:

- Roaming between VPNs is not completely transparent.
- VPNs have to overcome firewall barriers.
- VPNs do not support multicasting.
- VPNs do introduce excessive overhead. Everything is encrypted, even if you are surfing the Web.




Figure 8-13: IPSec Conceptual View

8.3.4.2 IPSec Architecture – A Closer Look

The IPSec security architecture is centered around two IP constructs that are used to provide security services in IPv4 and IPv6. These headers are the “IP Authentication Header (AH)” and the “IP Encapsulating Security Payload (ESP)” header. To understand how these two constructs work, it is necessary first to look at security associations (SAs). The following discussion takes a quick look at SAs, AHs, and ESPs. Detailed discussion of these topics can be quite deep, dense and bitter (see Stallings [2000], for example).

Security Associations (SAs). A security association is a table or database record consisting of a set of security parameters that dictate security operations on one or more network connections. A security association is normally one-way. The SA tables are established on the receiving host and referenced by the sending host using an index parameter known as Security Parameter Index(SPI). Before a secure transmission can be established the SA tables are created on the sending and receiving hosts. The sending host looks up the appropriate SA before transmission and passes the resulting SPI to the receiving host. The receiving host uses the SPI and the destination address to look up the corresponding SA on its system. The receiving host uses the SA to authenticate the transmission. In addition, an AH implementation will always be able to use the SPI in combination with the Destination Address to determine the security association and related security configuration data for all valid incoming packets.

Since a security association is normally one-way, an authenticated communications session between two hosts will normally have two Security Parameter Indexes in use (one in each direction). The combination of a particular Security Parameter Index and a particular Destination Address uniquely identifies the Security Association. The Destination Address may be a unicast address (i.e., single destination) or a multicast group address. An implementation of the Authentication Header or the Encapsulating Security Payload MUST support the concept of a Security Association. A Security Association normally includes the following key parameters (all parameters are not shown here):

- Authentication algorithm and algorithm mode being used with the IP Authentication Header (required)
- Key(s) used with the authentication algorithm in use with the Authentication Header (required)
- Encryption algorithm being used with the ESP (required)

- Key(s) used with the encryption algorithm for ESP (required)
- Lifetime of the key or time when key change should occur (recommended)
- Lifetime of this Security Association (recommended)
- Sensitivity level (for example, Secret or Unclassified) of the protected data (required for all systems claiming to provide multi-level security, and recommended for all other systems)

The IP Authentication Header (AH) is designed to provide integrity and authentication of IP packets. The data integrity feature of AH ensures that the data is not modified in transit, and the authentication feature ensures that the received packet was, in fact, transmitted by the party identified as the source in the packet header. AH does not include confidentiality. Why? Confidentiality is not included so that the AH can be widely available on the Internet, even in locations where the use of encryption to provide confidentiality is regulated. The Authentication Header supports security between two or more hosts implementing AH directly, or between a host and an IPsec “gateway” (e.g., a firewall with IPsec software). Other combinations are also possible (e.g., gateway-to-gateway). The IPsec gateway can provide services on behalf of one or more hosts on a trusted subnet. For example, the firewall with IPsec software implements AH, while all of the systems behind the firewall do not need individual IPsec, as they take advantage of AH services between the firewall and external systems. Thus the firewall acts on behalf of all the systems behind it, and the AH is used only between the firewall and the external hosts. Security associations are used between the various players (hosts, firewalls) that use AHs.

The IP Encapsulating Security Payload (ESP) is designed to provide integrity, authentication, and confidentiality to IP packets. ESP is basically responsible for encrypting the content of the packets. In a manner similar to AH, the ESP supports security between two or more hosts implementing ESP either directly or between a host and an IPsec “gateway” (e.g., a firewall with IPsec software). Other combinations are also possible (e.g., gateway-to-gateway). ESP is used to encrypt data between various players. If there are no security gateways present in the connection, then two end systems that implement ESP may also use it to encrypt the user data (e.g., TCP or UDP) being carried between the two systems. ESP is designed to provide maximum flexibility so that users may select and use only the security that they need. Security associations are used between the various players (hosts, firewalls) that use ESPs.

8.3.4.3 IPsec Documentation

IPsec is a complex protocol suite that is described in several documents shown in Figure 8-14. The main document is the Architecture document which covers the general concepts, security requirements, definitions, and mechanisms defining IPsec technology. The ESP Protocol and AH Protocol documents cover the packet format and general issues regarding the respective protocols. These protocol documents also specify default and mandatory values to implement algorithms. The “Encryption Algorithm” document set describes how various encryption algorithms are used for ESP. The “Authentication Algorithm” document set describes how various authentication algorithms are used for both ESP and AH. The encryption algorithm and authentication algorithm document sets describe, respectively, the sample encryption and authentication algorithms being used (e.g., MD5, SHA). The “Key Management Documents” describe the IETF key management schemes. The Domain Of Interpretation (DOI) document plays a central role in the document hierarchy because it contains values needed for the other documents

to relate to each other. This includes, for example, encryption algorithms, authentication algorithms, and operational parameters such as key lifetimes. DOI serves as a cross-reference index. Details about these documents can be found at the IETF site (www.ietf.org).

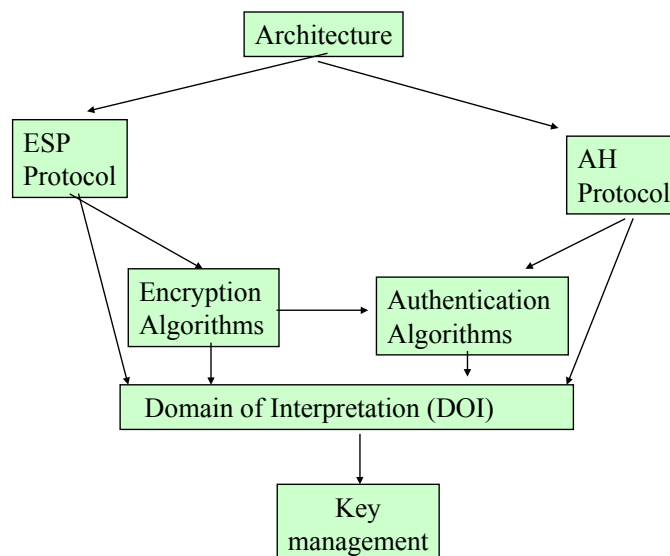


Figure 8-14: IPsec Documents

VPN and IPsec Information Sources

- www.vpnlabs.org – A good site for VPN information
- www.ietf.org/html.charters/ipsec-charter.html – Official website for IPsec
- www.bitpipe.com – white papers on different aspects of IPsec
- www.ietf.org/internet-drafts/draft-ietf-ipsec-dhcx-13.txt – Configuration of IPsec tunnel mode with DHCPv4 (IETF Proposed Standard)
- www.ietf.org/rfc/rfc3193.txt – Securing L2TP with IPsec (IETF Proposed Standard)
- www.ietf.org/internet-drafts/draft-ietf-ipsra-pic-04.txt – Legacy authentication within IPsec tunnel mode (PIC) (IETF draft, work in progress)
- www.counterpane.com/pptpv2-paper.html – Security analysis of PPTPv2
- www.counterpane.com/pptp-paper.html – Security analysis of PPTP
- www.microsoft.com/NTServer/Support/faqs/VPNSec_FAQ.asp – Microsoft point of view on PPTP
- www.ima.umn.edu/~pliam/xauth/ – Security analysis of XAUTH (shipping in most IPsec tunnel mode implementations)
- online.securityfocus.com – A good site for general information and articles on security

8.4 Corporate Intranet and Extranet Security

8.4.1 Overview of Corporate Network Security Issues

Securing a corporate network involves many issues. First, the external network has to be secure for customers, business partners, and employees. As stated previously, the public Internet is not secure and requires the use of VPN for security. The main problem is that there may be gaps in VPN support for very-long-haul communications. For example, communications between an office in New York and a partner in Hong Kong may go through several ISPs and NSPs with varying levels of VPN support. While IPSEC-based VPNs are addressing these issues, there are enough security gaps that need to be addressed in the public Internet. For this reason, extranets are adopted by some companies for B2B and C2B trade. In addition to the outside traffic, the internal network – the intranet – has to be secure so that malicious individuals cannot intrude. Thus intranets and extranets are important players in corporate network security, in addition to the public Internet security, of course.

8.4.2 What are Intranets and Extranets?

While the public Internet is not controlled by anyone and can raise security concerns, intranets and extranets are privately and jointly owned networks that can be controlled and secured by the owning organizations.

Intranets, also known as private Internets, are the IP networks that are used by corporations for their internal business, especially for exploiting Web technologies. Technically, an intranet uses the same technology as the public Internet – only it is smaller and privately owned (the physical network below the IP is privately owned) and thus hopefully better controlled and more secure. Thus any applications and services that are available on the public Internet are also available on the intranets. This is an important point for the Web because many companies are using Web technologies on their intranets for internal applications (e.g., employee information systems).

Extranets, also known as community of interest networks, are IP networks that are jointly owned by corporations for conducting secure business processing. These networks use the same Internet technologies, but the physical network is collectively owned by corporations to meet the security and reliability requirements imposed by the owners. As mentioned in the previous chapter, extranets are semi-private IP networks that are used to communicate within a group of interdependent communities of enterprises or trading partners. An extranet consists of a collection of Internet segments, each protected by firewalls, which are interconnected using secure leased lines across the remote locations. This solution provides security and guaranteed bandwidth, at the cost of leasing lines from telecom providers.

Corporations form extranets to conduct trade among members only in a secure and reliable environment as compared to the public Internet. An example of an extranet is the Automotive Network eXchange (ANX) network formed by manufacturing corporations (GM, Ford, Chrysler, and others). This extranet was formed by the Automotive Industry Action Group (AIAG) to provide a common communication infrastructure among automotive trading partners. ANX is intended for North America initially, with plans to

expand worldwide with over 3000 trading partners. The drivers for ANX are control communications cost with predictable service quality, support of a common set of applications (e.g., EDI, database lookup, Web, email, and Computer-Aided-Design [CAD] file transfer), and facilitated introduction of new applications (e.g., videoconferencing, interactive CAD) by using Internet technology.

Internet Role Players

Different individuals, groups and organizations play different roles in the Internet. To illustrate these roles, let us envision the Internet as an electronic shopping mall. Then we can discuss the following roles:

- Internet users are the people who visit the shopping mall (i.e., log on to the Internet). The Internet users are essentially the consumers of the services provided by the Internet.
- Content providers are the merchants (individuals, groups or organizations) that provide the products in the shopping mall (i.e., resources available on the Internet). You can think of these content providers as the merchants in the shopping mall.
- Internet access providers (IAPs) are the organizations that facilitate your access to the shopping mall (i.e., give you a communication line and an access port on the Internet). You can think of IAPs as the local authorities that provide you with roads and signs to get you to the shopping malls.
- **Internet service providers (ISPs)** are the individuals and organizations that help the content providers set up their shops in the shopping mall (i.e., help in building websites). Many small content providers seek the help of ISPs to set up Web servers with appropriate security and backup/recovery.

8.4.3 Corporate Intranet and Extranet Architectures for Security

Corporate Intranet architectures have changed over the years to provide secure access. Figure 8-15 shows a typical corporate Intranet in the 1990s. In this case, separate networks were maintained for the internal users, the intranet, the branch offices, and remote users. In particular, the remote users and branch offices accessed the corporate resources through dial-up or leased lines.

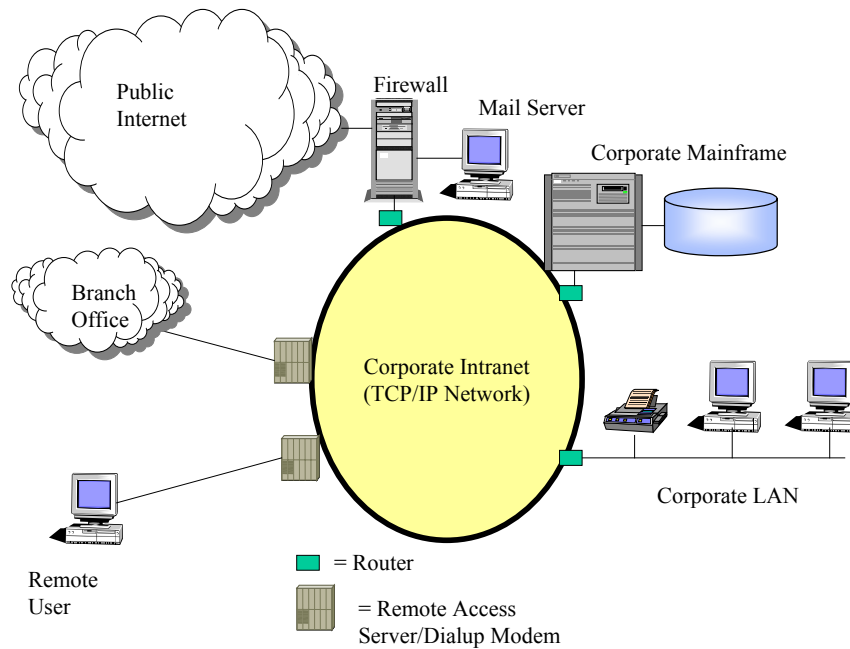


Figure 8-15: Past Corporate Network Architecture

Figure 8-16 shows a more recent view in which all external accesses go through the public Internet. The firewall has to be strong enough and the internal intranet has to have its own security to deter privacy, integrity, and availability attacks.

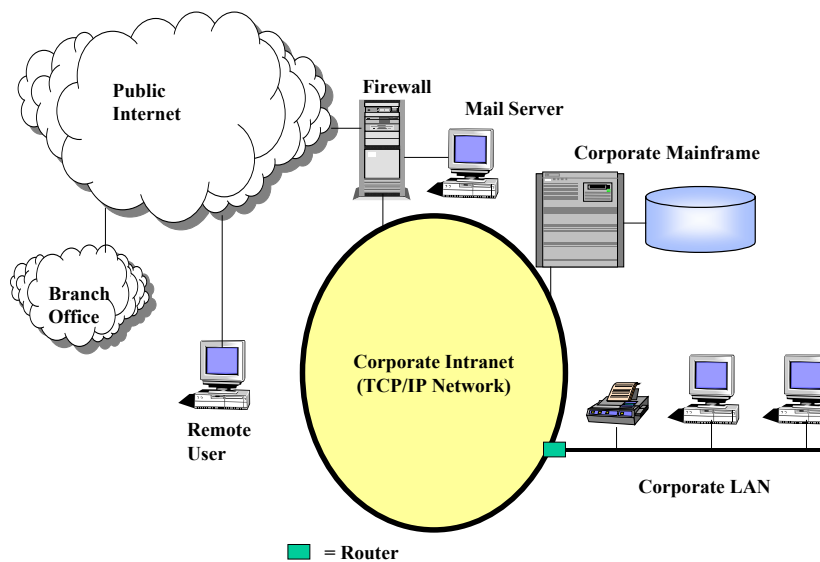


Figure 8-16: Current Corporate Network Architecture

Although Figure 8-16 protects the access to the corporate resources from outsiders by using a firewall, it does not secure the path between its remote users and the firewall over the public Internet. Figure 8-17 shows how a VPN can be used to secure the public Internet traffic *before* it reaches the firewall.

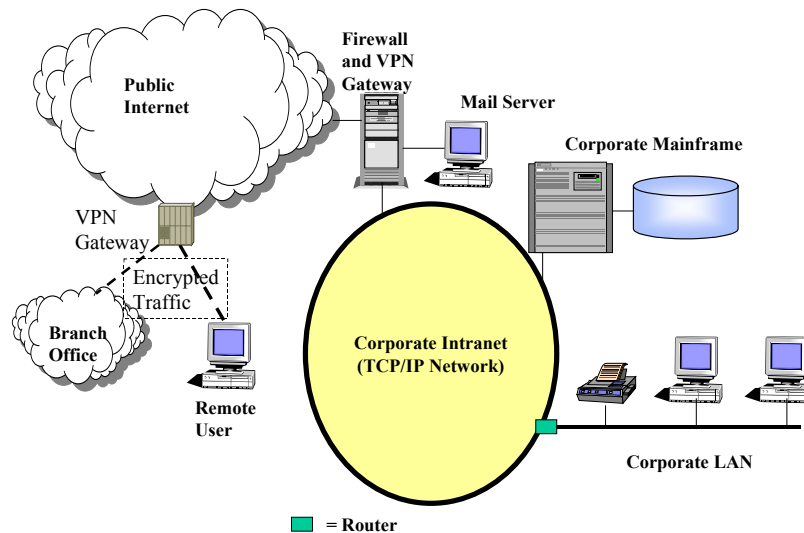


Figure 8-17: VPN Access for External Customers

Stronger external security can be achieved by forming extranets between business partners. Extranets maintain security by controlling access to the underlying network and thus bypassing the public Internet altogether. For example, the Automotive Network eXchange (ANX) network is controlled by manufacturing corporations (GM, Ford, Chrysler, and others). The architecture of ANX is shown in Figure 8-18. This network is shared by various certified service providers (CSPs) that support the ANX subscribed trading partners. Public Internet access is restricted through Internet Exchange Points that provide adequate security controls.

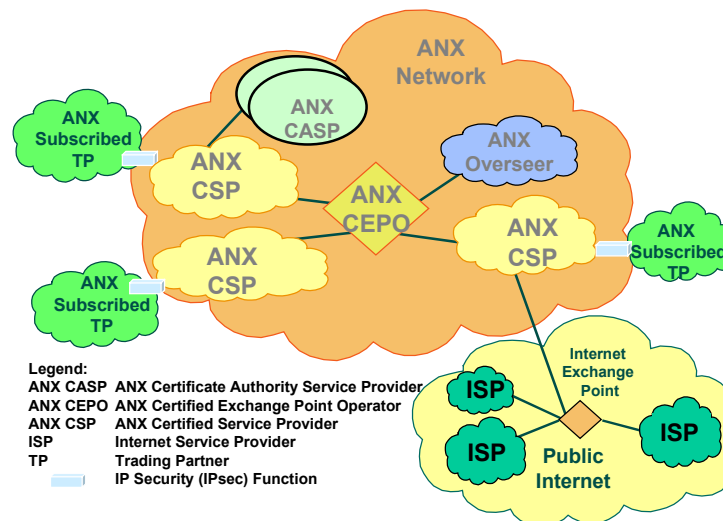


Figure 8-18: ANX Extranet Architecture

VPNs and extranets secure external networks but do little to secure the internal corporate resources. To strongly secure internal corporate resources, more than one firewall may be used. Figure 8-15 shows a common architecture that uses two firewalls: an outer firewall that exposes some services to the outside world and a second, inner firewall, that protects

the inner resources. The zone between the two firewalls is known as a demilitarized zone, or DMZ. Enterprises can put customer- and business partner-facing applications and databases in the DMZ and thus separate their back-end applications and databases by using the second firewall. In many cases, DMZs are used to compensate for the weaknesses of the external network. For example, if you are concerned that anyone from the public Internet can access your customer-facing applications, then a second firewall is highly desirable.

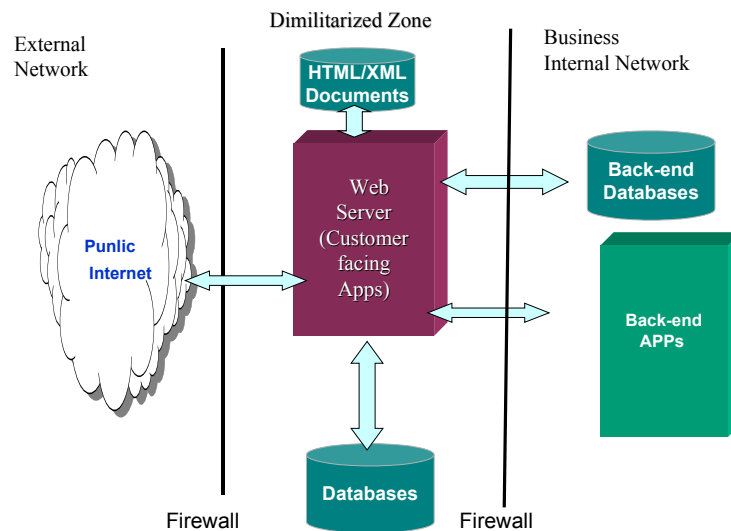


Figure 8-19: Dimilitarized Zones

8.5 Firewalls – Protecting the Enterprise Perimeter

8.5.1 Overview

Firewalls are network filters that police *who* enters and leaves an enterprise network and *what* gets in and out. A firewall is essentially a software package that is typically installed on network routers. This software checks each IP packet and determines if it should enter the system. The usual metaphor for a firewall is the medieval castle and its perimeter defense systems [Sheldon 1996]. The castle walls provide the perimeter defense, while the gatehouses and drawbridges provide “choke points” through which everyone must travel to enter or leave the castle. You can monitor and block access at these choke points. Interestingly, the castle proved quite capable of withstanding attacks until cannons (and later on, airplanes) came along.

Most firewalls are designed to protect the enterprise systems (e.g., internal networks and databases) from outsiders. They do not secure the path from your site to a remote office – that is accomplished through VPNs, for example. Thus it can be argued that firewalls do not provide network security, instead they provide “system security.” Discussions about

protecting enterprise systems usually focus on threats from the Internet, but internal users are also a threat. Indeed, many unauthorized activities are perpetrated by the internal users. In addition, organizations that connect with business partners over private networks create a potential avenue for attack. Users on the business partner's network may take advantage of the inter-company link to steal valuable information. Using the castle analogy again, the castle does not protect the monarch from mutinies from inside, or from sneaky attacks from “friends” who come to visit and want to stay forever – this is the well known Trojan horse example.

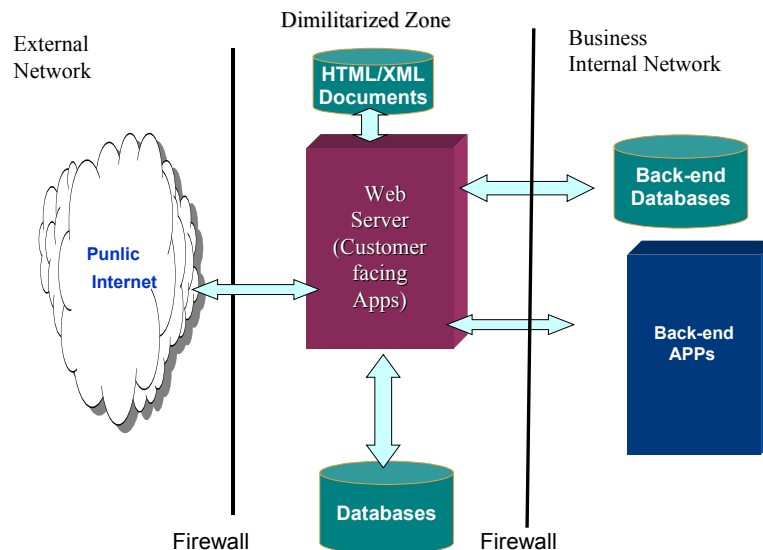


Figure 8-20: Conceptual View of Firewalls

Firewalls provide a logical and physical separation of the public Internet and internal IT systems. A good security design generally has two firewalls: an outer firewall that exposes some services to the outside world, and a second, inner firewall, that protects the inner resources (see Figure 8-20). The zone between the two firewalls is known as a demilitarized zone, or DMZ. You can put customer- and business partner-facing applications and databases in the DMZ and thus separate your back-end applications and databases by using the second firewall. Different types of firewalls are commercially available from vendors such as Cisco, Bay Networks, Ascend, Microsoft, Altavista and Borderware (secure computing). Some of the firewalls are small and inexpensive software packages that you can install on your laptop to keep intruders away from your machine, while others are expensive and sophisticated software systems that are intended to protect your organization networks. The following sections describe the various types of firewalls.

8.5.2 Firewall Configurations

Firewalls appear in many configurations. In essence, any device that controls network traffic for security reasons can be called a firewall, and in fact the term “firewall” is used in a generic way. However, there are some basic types of firewalls that use different strategies for protecting network resources. The most basic firewall devices are built on routers and work in the lower layers of the network protocol stack. They provide packet

filtering and are often called screening routers. On the other extreme are high-end proxy server gateways that operate at the upper levels of the protocol stack (i.e., the application layer) and perform advanced monitoring and traffic control by looking at certain information inside packets. Most firewalls have been designed around these two approaches. A packet filtering firewall uses the “strip-search” method – the packets are first checked and then either dropped or allowed to enter based on various rules and specified criteria. A proxy service acts as an agent for a user who needs to access a system on the other side of the firewall. There are other types of firewalls also that act as circuit switches between the internal and external users. Figure 8-21 shows the three most common firewall configurations that are described below. In several real life situations, mixtures are used to protect enterprise assets. For example, routers are often used in conjunction with proxy gateways to build a multitiered defense system. Additional techniques are also used in modern firewalls. For example, some firewalls use a “stateful” inspection technique that is analogous to a gatekeeper remembering some defining characteristics of anyone leaving the premises and only allowing people back in with those characteristics. We will also look at these techniques.

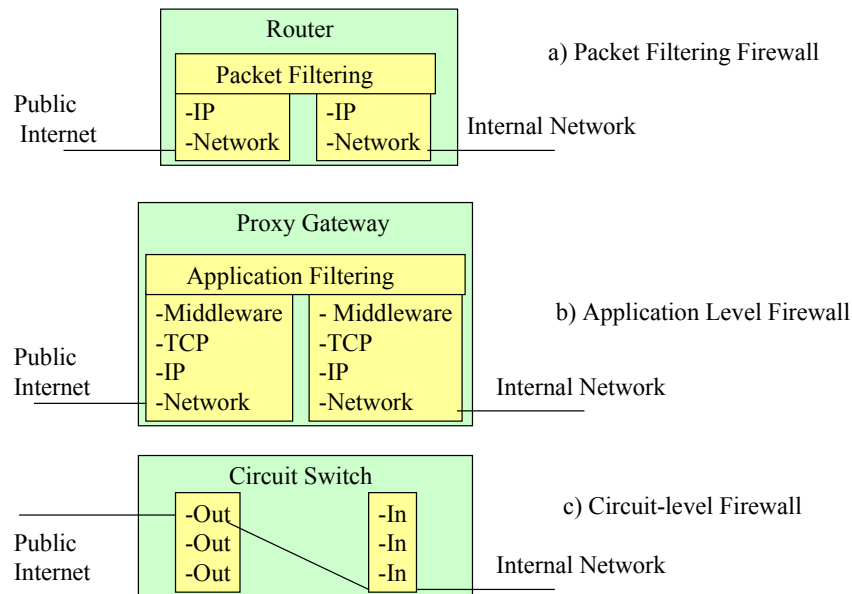


Figure 8-21: Typical Firewall Configurations

8.5.2.1 Packet-Filtering Firewalls

In this firewall, also known as a Router Firewall, the two networks are completely isolated but the packets transfer directly through routers (see Figure 8-21a). The screening routers use packet filter programs that look at information related to the hardware address of a computer, its IP address, and also the TCP port number to determine the application type, and then provide filtering based on that information. For example, email and HTTP traffic has certain TCP port numbers that are part of the packet. This is why and how your emails and HTTP traffic goes through most firewalls while other

traffic does not. A screening router may be a stand-alone routing device or a computer that contains two network interface cards (dual-homed system). Some packets may be sent to another program and repackaged before transmission. Network administrators program the device with a set of rules that define how packet filtering is done. Ports can also be blocked; for example, most firewalls block all applications except email and HTTP (Web) services. It is extremely important to define the rules that provide the highest level of protection to your enterprise resources, especially if the public Internet is connected to one side of the router. In these cases, two firewalls and a DMZ as described above are important (think of these as a castle with two walls).

8.5.2.2 Application Level Firewalls – The Proxy Gateways

A high level of protection can be provided by application-level proxy servers that act as firewalls (see Figure 8-21b). In general, proxies sit between clients and servers and act as “stand-ins.” For example, a proxy acts as a server to a client and vice versa. A proxy server is essentially an intermediate program that behaves as a server but in fact passes the requests to one or more real servers. In effect, it is a fake server. In most practical cases, the proxy server receives the client calls, does some processing (typically security checking) and then itself becomes a client to other servers. Web proxies are used commonly to serve as a front-end to many clients; i.e., instead of handling calls from all clients, the Web server handles calls from the proxy server that is front-ending many clients. In some cases, the proxy blocks all outside connections and only allows internal users to access the Internet. The only packets allowed back through the proxy are those that return responses to requests from inside the firewall. In other cases, both inbound and outbound traffic are allowed under strictly controlled conditions. The proxy service changes the IP address of the client packets to hide the ID of the client, then it acts as a proxy agent for the client on the Internet. Using proxies reduces the threat from hackers who monitor network traffic to glean information about computers on internal networks. The proxy hides the addresses of all internal computers.

Proxies can operate at several levels but the application-level proxy firewalls provide the most extensive packet analysis. When packets from the outside arrive at the firewall, it evaluates IP addresses and also looks at the data in the packets to stop hackers from hiding information in the packets. A typical application-level firewall can provide proxy services for applications and protocols like Telnet, FTP (file transfers), HTTP (Web services), and SMTP (email). As expected, application level firewalls introduce performance overhead but ensure higher level of security. In addition, a separate proxy must be installed for each application-level service (i.e., you need one proxy program for email, another for FTP, and so on). However, the proxy security policies can be much more powerful and flexible because all of the information in packets can be used by administrators to write the rules that determine how packets are handled by the gateway. It is easy to audit just about everything that happens on the gateway. Network administrators can choose to strip computer names to hide internal systems and evaluate the contents of packets for appropriateness and security.

8.5.2.3 Circuit-Level Firewalls

This type of firewall, shown in Figure 8-21c, provides a controlled network connection between internal and external systems. A virtual “circuit” exists between an internal client (e.g., marketing department LAN) and an external site (e.g., a valued customer). The firewall basically has the rules that switch the traffic from various input ports to the

output ports. Due to this one-to-one relationship, the firewall can do very focussed filtering and checking very efficiently. The connections are made at the TCP level. The circuit-level firewall does not allow end-to-end TCP connections. Instead, it sets up two TCP connections – one between itself and the internal user and the other between itself and the external user. After this, the firewall relays the messages between the in and out ports. The circuit-level firewall can be used as a stand-alone gateway or as a front-end to proxy servers. For example, the public Internet requests go through this circuit to the proxy server, and the proxy server delivers those requests to the Internet after changing the IP address. While traffic is allowed through this gateway, external systems never see the internal systems. This type of connection is often used to connect “trusted” internal users to the Internet.

8.5.2.4 Inspection and Filtering Techniques

The problem with proxies and application-level firewalls is that they must evaluate a great deal of information in thousands of packets. This affects performance and increases costs. A variety of approaches have been developed to address these issues. Some of the approaches involve combining the router, circuit-switch and proxies into the same firewall so that different levels of security can be used for different types of traffic. Some firewalls use more innovative techniques. For example, some firewalls use **stateful inspection**, where instead of examining the contents of each packet, the bit patterns of the packets are compared to packets that are already known to be trusted. For example, if you access some outside service, the server remembers things about your original request such as the port number, and source and destination address. This “remembering” is called saving the state. When the outside system responds to your request, the firewall server compares the received packets with the saved state to determine if they are allowed in. While stateful inspection provides speed and transparency, one of its biggest disadvantages is that inside packets make their way to the outside network, thus exposing internal IP addresses to potential hackers. Some firewall vendors are using stateful inspection and proxies together for added security.

8.5.3 Firewall Administration, Design Issues and Policies

The firewall controls access to the organization's internal networks (perimeter) by acting like a gatekeeper that examines each user's credentials before allowing access to the network. To create a good firewall, someone must write and maintain the internal rules identifying the people, applications, hosts, or addresses that are allowed or rejected in very fine detail. It is extremely important to use firewalls as control devices by clearly thinking through the following:

- What policies can be enforced by using the firewall
- Who will be responsible for specifying and implementing these policies
- Are adequate controls provided so that unauthorized users do not change firewall rules
- Are firewalls audited on a regular basis to assure continued compliance with policies

Once in place, a firewall requires constant monitoring and control. A firewall is not a stand-alone device – it needs to be managed and monitored on a regular basis in the context of the organization. Security policies and procedures must be put into place that reflect the overall management approach to security (we discussed the security management approaches in chapter 2). It is important to know what firewalls can and

cannot do. For example, while firewalls keep the external intruders out, they do not protect a corporation from internal saboteurs. You may need to separate departments, workgroups, divisions, or business partners using the same firewall technology, and you may need to implement encryption throughout your organization. Firewalls also do not protect against leaks from telecommuters who connect to the outside with a desktop modem. In addition, if some new threat comes along, your firewall might not be able to protect against it.

Weaknesses in firewalls are the main reason for enterprise security breaches. It is extremely important that intruders cannot find a hole in your firewall and do not hijack an administrative account. No firewall can protect against inadequate or mismanaged policies. If an internal user dials out through an unauthorized connection, then this provides a backdoor that must be closed somehow. Naturally, the firewall itself and the firewall policy are two distinct things that require their own planning and implementation. Firewalls need to implement complicated sets of policy rules that address internal and external access, remote user access, virus protection and avoidance, encryption requirements, program usage, and a number of other considerations. Most firewall products provide graphical interfaces for rule specification. These tools can help administrators and users know what type of activities are allowed on the network. If internal users find security policies too restrictive, they may bypass them by connecting to the Internet through a modem. Thus a tradeoff between ease of use and protection of enterprise assets is needed. Some of the guidelines for firewalls are [Sheldon 1996]:

- The traffic must be filtered to allow only authorized packets to pass. In particular, do not allow any passwords or internal addresses to cross the firewall.
- Do not run any services on the firewall except those specifically required to provide firewall services. In particular, firewalls must not be used for general-purpose file storage or to run programs, except for those required by the firewall.
- Use a DMZ approach with two firewalls if you need to provide services to the public. Put the public services in the DMZ zone protected by an outer firewall. The inner firewall should protect the back-end systems. A packet-filter router can be used in the outer firewall and an application-level gateway can be used in the inner firewall.
- For outbound connections, encryption should be used to hide transmitted information. If users are accessing the Web with Web browsers, then SSL (Secure Socket Layer) should be used (we will discuss SSL later).
- For inbound traffic, decide what you want to allow in, when, and why. Email is the usual requirement. HTTP traffic is also allowed but should be monitored closely.

Firewalls, even at their best, can only deter, but not completely prevent, network penetration from outsiders, and should be viewed as one element in an overall security plan. In order to deal effectively with Internet security, broader corporate policies and procedures, user responsibilities, and security awareness training may be required. We discussed these issues in chapter 2. In addition to management controls, firewalls may be augmented by using intrusion detection systems (IDSs). In fact, as discussed in a previous chapter, many new firewalls are combining firewall and IDS functions into a single package. See, for example, the offerings by ISS (www.iss.com).

Additional Information Sources About Firewalls

- Cheswick, W., and Bellovin, S. *Firewalls and Internet Security: Repelling the Wily Hacker*. Reading, MA: Addison-Wesley, 1994.
- Chapman, D., and Zwicky, E. *Building Internet Firewalls*. Sebastopol, CA: O'Reilly & Associates, 1995.
- Sheldon, T. "General Firewall White Paper," <http://secinf.net/info/nt/fw/firewall.html>, Nov. 1996.
- Stallings, W. *Network Security Essentials*. Prentice Hall, 2000.
- Stein, L. *Web security: A step-by-step Reference Guide*. Addison Wesley, 1998.
- *Microsoft Windows NT® Security Handbook*.
- Website: <http://Firewall.com>

8.6 Short Case Studies and Examples

8.6.1 Combining Network and Physical Security

Many universities are facing physical security threats (e.g., theft, vandalism) along with information security problems of eavesdropping and intrusions. Combining IT and physical security appears to be a good approach.

An example of such an effort can be seen at the Bishop's University in Canada. Bishop's is a university with 1,900 students who attend classes in 12 buildings spread out across the campus. A few specific trouble spots were identified as a result of a physical security analysis: the computer labs (computers were easily stolen), the parking lot, and the darker pathways between the buildings all invited mischief, or worse, assaults. The vandalism costs alone amounted to \$30,000 per year. Thus a solution costing around \$100K would take the university three years to recover its costs.

The core of the new security system involved deployment of closed circuit cameras around the campus and through the 12 buildings. The question was how best to do it. The approach adopted was using the university's IP network. Instead of building a completely new cabling infrastructure, the university decided to plug in cameras to the university's existing Internet network. This also provided them with a link to Bishop's central monitoring service and to other application service providers off-site. In addition, the university could convert pictures taken with existing analog cameras into digital images and send them through the network.

The solution was developed under the general umbrella of **AVVID (Architecture for Voice, Video and Integrated Data)** – a partnership between Cisco (which owns 80 per cent of the Internet market), ISS (a large-scale integrator of video surveillance equipment) and Sanyo (which supplies the IP surveillance solution). Solutions were developed jointly by the three partners.

Source: B. Toews, "AVVID Teammates," *Canadian Security Magazine*, December 2003.

8.6.2 beBetter Networks Chooses a FireWall

beBetter Networks specializes in employee productivity, performance, and retention. The company has developed a suite of video-enabled services which include management training and coaching, selection and career advancement testing, behavioral health and substance abuse counseling, nutrition and financial counseling, and crisis training and intervention. The company initially shared office space and network resources with another company. However, when it moved into its own offices, beBetter needed a secure firewall solution that would enable connectivity with the Internet while protecting the company's valuable and rapidly growing IT assets. The company wanted to publish content residing on servers inside the firewall for access via a Virtual Private Network (VPN). It also wanted support for videoconferencing protocols such as H.323.

The company chose the Microsoft Internet Security & Acceleration (ISA) Server for its Internet connectivity solution. For an estimated retail price of \$1500, the ISA Server provided support for virtual private networking and secure server publishing. This turned out to be a much cheaper solution than the Check Point FireWall-1 that costs close to \$10,000. The ISA server provides inspection and filtering at packet, circuit, and application-levels and securely exposes internal systems for access over the Internet. The built-in wizard makes it easy to expose internal content to the Internet. The company is using the VPN capabilities of the ISA Server to enable unrestricted yet secure access for employees while out of the office. The ISA Server provides management capabilities for the company's network administrators through a single user interface, and also provides capabilities for add-ons such as software for Web filtering.

Editor's comments: This case study was chosen to illustrate typical features found in firewalls and also for cost comparisons – a very rare piece of information.

Source: www.microsoft.com .

8.6.3 Secure Automated Change Management for a University Network

Effective configuration and change management is essential for managing changes across multi-vendor networks. In particular, making the changes, verifying them, and then pushing the new configurations to the devices should take very little time and must not introduce system vulnerabilities. Many products are commercially available for network change management and have been used by different organizations.

An example is St. Mary University, a private Catholic university located in San Antonio, Texas, that manages its network from a central site. The university, like many others, depends on its network to support educational, administrative, and recreational services, including Internet access, email, databases, and administrative services. The university network includes 120 network devices from a variety of hardware vendors and supports more than 5000 users daily, including students, faculty and other employees. The university used DeviceAuthority Configuration and Change Management from AlterPoint to introduce changes to all the switches in the campus residence halls. These switches are the most dynamic devices in the network and require daily management. After a successful initial deployment across the critical switches, St. Mary's rolled out

DeviceAuthority to the additional devices in their network. The product also delivers automatic daily updates on device changes and configurations throughout the network.

DeviceAuthority stores configuration data to enable immediate rollback to previous known good configurations in case of network faults. The product also adds a layer of control to the network. Administrators can designate various levels of authority for propagating changes across the network by proposing and reviewing changes before implementing them. It also encrypts the configurations via SSH, and alerts the administrators when configurations change.

To summarize, propagating configuration changes across the university network used to take them over a week. With automated systems such as DeviceAuthority, changes are made in 30 minutes.

Source: “AlterPoint™ Customer Case Study – St. Mary’s University” (www.alterpoint.com)

8.6.4 From Site Finder to Internet Ownership

Verisign is the registrar for the main database that keeps track of who owns which names in the .com and .net top-level domains – this database directs people to .com and .net addresses. The company temporarily suspended a new service that redirected misspelled or unassigned .com domain names to a search page it managed. This service was supposed to get around the common problem that requests for nonexistent or inactive domain names trigger error messages. But this service was opposed by many who contended that VeriSign's addition of a “wild card” feature interfered with spam filters and mail servers. VeriSign subsequently ordered a temporary suspension of the service.

The controversy attending the dispute over the Site Finder service revealed a cultural divide between the Internet technologists who helped guide the Internet in its infancy and the businesspersons who realized the Internet's commercial potential. The tension is about who will control the future of the Internet – current public ownership or the private sector. Verisign management believes that it is time to transfer the responsibility for operating the root servers from volunteers to the commercial sector because of the increased security attacks. For a discussion of tradeoffs about this topic, see the interview with Verisign's CEO (CNET News.com interview on October 17, 2003 published on Zdnet.com).

Source: C. Cooper, *CNET News.com*, interview aired October 17, 2003, 8:01 AM PT

8.6.5 Using SATAN to Reduce Network Security Exposures

A team of security experts were asked to define the security status of a large corporate TCP/IP network, consisting of over 14,000 systems, and reduce the level of security exposure quickly. Various tools were available but the team chose SATAN (Security Administrator Tool for Analyzing Networks) for its breadth of checks, ease of use, availability, speed, and modifiability. SATAN (<http://www.fish.com/satan/>) was a commonly used tool for network security analysis and audits in the mid 1990s.

The team developed a plan for how best to canvas all nodes. Due to the enormity of the project, the team decided to augment SATAN's reporting capabilities with a few Perl

scripts. The overall strategy was to check all registered IP addresses for network security problems by using SATAN. These problems were reported to the DNS administrators, who were responsible for getting the problems corrected or justified. The final security state of the systems was reported to upper management. As a result of this work, the security exposure was reduced from 40% to 4% in 6 months on all 14,000 systems. The SATAN system proved very useful in this effort. A detailed technical discussion of what was done and how/why it was done can be found in the original report.

Source: Nancy Cook and Marie Corbin, "Flirting With SATAN," March 1, 1996, www.porcupine.org/auditing

8.6.6 Digital Networks Tie Together National and Local Security

Digital technology is tying together national and local public safety and security systems. For example, DICE Corp. of Bay City, Mich., is deploying an alarm monitoring and communications system for the Department of Homeland Security. This nationwide system is intended to tie together federal and local emergency management and response organizations, including 911 centers, hospitals, weather services, military operations, and other governmental groups and businesses. National and local government agencies, and private organizations with emergency response concerns, subscribe to the system and request that certain kinds of data be routed to them. For example, a company specializing in distributing information about the availability of beds in hospitals has subscribed to the system.

During a disaster, the company will receive information from both first responders and hospitals, process the data, and inform emergency medical teams about which hospitals can and cannot accommodate patients. This avoids routing of ambulances to a hospital filled to capacity, with later re-routings which could cause deaths while looking for a suitable hospital. The system receives emergency information from a variety of sources, sorts the information, and forwards it to the appropriate national and local authorities. This could include hazardous material spills, natural disasters, automobile crashes, public health alerts, terrorist threats, railroad incidents, and virtually any other kind of local, regional, or national emergency.

The DICE network routes emergency information from the ground up to appropriate local and national decision-makers. The technology to send information up the chain and also to implement policy is developed by VistaScape Technology. For example, a U.S. Navy installation had concerns about potential terrorist activity against ships docked in its port. Navy policy required that guards in boats patrol the port at all times. In light of current threats, comprehensive coverage required too many guards and too many boats to be practical. In addition, guards in boats could initiate a response to an event only as fast as they could pass alerts along by phone, radio or PDA. VistaScape addressed this problem with a security data management system (SDMS) capable of automating the operation of a closed circuit television (CCTV) system along with the policy responses required by what the CCTV cameras see. Off-the-shelf infrared and visible spectrum cameras capture the information and SDMS monitors the video from the cameras. The cameras operate as sensors and transmit structured pieces of information representing events, objects, and locations.

Unlike a conventional CCTV system, SDMS does not simply transmit video to human monitors who must develop a response. Instead, the technology comprehends situations and makes measured responses based on policies programmed into the system. Suppose, for example, the system detects an unauthorized boat approaching the naval base. Programming would turn on the floodlights in that area of the port. If the boat continued to approach, an audible alarm would go off. Next, a public address system would broadcast a pre-recorded message requesting that the boat turn around immediately. At the same time, the system would page guards on port patrol boats over PDAs and initiate a human response. Such a system could continue to apply higher and higher-level criteria to security problems as the level of threat increases. At the same time, the system would inform the human chain of command.

Because the DICE system and the VistaScape system deal in digital information, both can communicate if necessary. An SDMS owner, for example, might subscribe to the DICE communications network. Once connected, SDMS could communicate with the Federal Emergency Management Agency (FEMA) at the national level, a hospital at the local level, or any other emergency responder on the system.

Source: M. Fickes, "Digital Networks Tie Together National and Local Security," *Access Control & Security Systems*, May 1, 2003. (Summarized with permission).

8.7 Suggested Review Questions Before Proceeding

- 1) What is network security and how is it related to Internet security?
- 2) What are AAA servers and how is RADIUS related to it?
- 3) What are the main issues in public Internet security and what role can the service providers play in making the public Internet safe?
- 4) What are VPNs? Explain through an example.
- 5) What is IPSec and how is it related to VPN? Explain through an example.
- 6) What are intranets and extranets and how do their security features differ from those of the public Internet?
- 7) What are firewalls and what are different types of firewalls? How are firewalls used in network security?