CHAPTER 11
NETWORK SECURITY

Chapter Summary

This chapter describes why networks need security and how to provide it. The first step in any security plan is risk assessment, understanding the key assets that need protection, and assessing the risks to each. A variety of steps can be taken to prevent, detect, and correct security problems due to disruptions, destruction, disaster, and unauthorized access.

Learning Objectives

After reading this chapter, students should:

- Be familiar with the major threats to network security
- Be familiar with how to conduct a risk assessment
- Understand how to ensure business continuity
- Understand how to prevent intrusion

Key Terms

access card
access control list (ACL)
account
Advanced Encryption Standard (AES)
adware
algorithm
anomaly detection
antivirus software
application-level firewall
asymmetric encryption
authentication
authentication server
availability
backup controls
biometric system
brute-force attack
business continuity
central authentication
certificate
certificate authority (CA)
ciphertext
computer forensics
confidentiality, integrity, and availability (CIA)
continuous data protection (CDP)
controls
corrective control
cracker
Data Encryption Standard (DES)
DDoS agent
DDoS handler
decryption
denial-of-service (DoS) attack
desktop management
detective control
disaster recovery drill
disaster recovery firm
disaster recovery plan
disk mirroring
distributed denial-of-service (DDoS) attack
eavesdropping
encryption
entrapment
fault-tolerant server
financial impact
firewall
hacker
honey pot
host-based
information warfare
integrity
Internet Key Exchange (IKE)
intrusion prevention systems (IPS)
IP Security Protocol (IPSec)
IP spoofing
IPS management console
IPS sensor
IPSec transport mode
IPSec tunnel mode
Kerberos
key
key management
misuse detection
NAT firewall
network address translation (NAT)
network-based IPS
one-time password
online backup
packet-level firewall
passphrase
password
patch
phishing
physical security
plaintext
Pretty Good Privacy (PGP)
preventive control
private key
public key
public key encryption
public key infrastructure (PKI)
RC4
recovery controls
redundancy
redundant array of independent disks (RAID)
risk assessment
risk assessment frameworks
risk mitigation controls
risk score
rootkit
RSA
Secure Sockets Layer (SSL)
secure switch
security hole
security policy
smart card
sniffer program
social engineering
something you are
something you have
something you know
spyware
symmetric encryption
threat
threat scenario
time-based token
token
traffic analysis
traffic anomaly analyzer
traffic anomaly detector
traffic filtering
traffic limiting
triple DES (3DES)
Trojan horse
uninterruptible power supply (UPS)
user profile
user authentication
virus
worm
zero-day attack

Chapter Outline

1. Introduction
    a. Why Networks Need Security
    b. Types of Security Threats
    c. Network Controls
2. Risk Assessment
    a. Develop risk measurement criteria
    b. Inventory IT assets
    c. Identify Threats
    d. Document Existing Controls
    e. Identify Improvements
3. Ensuring Business Continuity
    a. Virus Protection
    b. Denial-of-Service Protection
    c. Theft Protection
    d. Device Failure Protection
    e. Disaster Protection
4. Intrusion Prevention
    a. Security Policy
    b. Perimeter Security and Firewalls
    c. Server and Client Protection
    d. Encryption
    e. User Authentication
    f. Preventing Social Engineering
    g. Intrusion Prevention Systems
    h. Intrusion Recovery
5. Best Practice Recommendations
6. Implications for Management
7. Summary

Answers to Textbook Exercises

Answers to End-of-Chapter Questions

1. What factors have brought increased emphasis on network security?

   Both business and government were concerned with security long before the need for computer-related security was recognized. They always have been interested in the physical protection of assets through means such as locks, barriers, and guards. The introduction of computer processing, large databases, and communication networks has increased the need for security. For many people, security means preventing unauthorized access, such as preventing a hacker from breaking into your computer. Security is more than that, however. It also includes being able to recover from temporary service problems (e.g., a circuit breaks) or from natural disasters (e.g., fire, earthquake).

   The factors that have brought increased emphasis on network security are:

   - Numerous legal actions involving officers and directors of organizations
   - Pronouncements by government regulatory agencies requiring controls
   - Losses associated with computer frauds are greater on a per incident basis than those not associated with computers
   - Recent highly publicized cases of viruses and criminally instigated acts of penetration
   - Data is a strategic asset
   - The rise of the Internet with opportunities to connect computers anywhere in the world (increased potential vulnerability of the organization's assets)
   - Highly publicized denial-of-service incidents

2. Briefly outline the steps required to complete a risk assessment.

   A. Develop risk measurement criteria
   B. Inventory IT assets
   C. Identify threats
   D. Document existing controls
   E. Identify improvements

3. Name and describe the main impact areas. Who should be responsible for assessing what is meant by low/medium/high impact for each of the impact areas? Explain your answer.

   1. Financial – revenue and expenses
   2. Productivity – business operations
   3. Reputation – customer perceptions
   4. Safety – health of customers and employees
   5. Legal – potential for fines and litigation

   Business leaders should make the decisions on the impact of each impact area because these are business decisions.

4. What are some of the criteria that can be used to rank security risks?

   Importance can be based on number of criteria such as which would have the greatest dollar loss, be the most embarrassing, be the most prone to liability judgments, and have the highest probability of occurrence. The relative importance of a threat to your organization depends upon your business. A bank for example, is more likely to be a target of fraud than a restaurant with an electronic marketing site on the Web. The criteria will also depend on the industry in which the organization works.

   Some other criteria that can be used to rank risk in a data communication network are:

   - Most damaging, most dangerous, most risky.
   - Most sensitive, most critical to organization, most likely to cause political problems
   - Most costly to recover, most difficult to recover, most time consuming to recover
   - Greatest delay, most likely to occur

5. What are the most common security threats? What are the most critical? Why?

   Some of the more common security threats include viruses, theft of equipment, theft of information, device failure, natural disaster, sabotage, and denial of services. The most critical will be the ones for a particular organization that have the highest impact score. This will vary based on the industry, geographic locations, etc.

6. Explain the purpose of threat scenarios. What are the steps in preparing threat scenarios?

   Threat scenarios describe how an asset can be compromised by one specific threat. An asset can be compromised by more than one threat, so it is common to have more than one threat scenario for each asset. The purpose is to begin preparation for mitigation of that threat.

   In order to prepare for threat scenarios, the following steps must be followed:
   1. name the asset
   2. describe the threat
   3. explain the consequence (violation of confidentiality, integrity or availability)
   4. estimate the likelihood of this threat happening (high, medium, low)

7. What is the purpose of the risk score and how is it calculated?

   Risk scores are used to compare the risk scores among all the different threat scenarios to help us identify the most important risks we face. It is calculated by multiplying the impact score by the likelihood (using 1 for low likelihood, 2 for medium likelihood, and 3 for high likelihood).

8. In which step of the risk assessment should existing controls be documented?

   Documenting existing controls is the fourth step in the process, between identifying threats and identifying improvements.

9. What are the four possible risk control strategies? How do we pick which one to use?

   The risk control strategies are to accept the risk, mitigate it, share it, or defer it. Selection of a strategy depends on things such as the impact (positive or negative) of the risk, the likelihood of the event occurring, and the cost.

10. Why is it important to identify improvements that are needed to mitigate risks?

    It is important to identify improvements that are needed to mitigate risks because risks are always changing and responses (including technologies) are changing as well.

11. What is the purpose of a disaster recovery plan? What are the major elements of a typical disaster recovery plan?

    A disaster recovery plan should address various levels of response to a number of possible disasters and should provide for partial or complete recovery of all data, application software, network components, and physical facilities. The most important element of the disaster recovery plan are backup and recovery controls that enable the organization to recover its data and restart its application software should some portion of the network fail.

    Major elements of a typical disaster recovery plan are:

    - The name of the decision-making manager who is in charge of the disaster recovery operation. A second manager should be indicated in case the first manager is unavailable.
    - Staff assignments and responsibilities during the disaster.
    - A pre-established list of priorities that states what is to be fixed first.
    - Location of alternative facilities operated by the company or a professional disaster recovery firm and procedures for switching operations to those facilities using backups of data and software.
    - Recovery procedures for the data communication facilities (WAN, MAN, BN, and LAN), servers and application systems. This includes information on the location of circuits and devices, whom to contact for information, and the support that can be expected from vendors, along with the name and telephone number of the person to contact.
    - Action to be taken in case of partial damage, threats such as a bomb threat, fire, water or electrical damage, sabotage, civil disorders, or vendor failures.
    - Manual processes to be used until the network is functional.
    - Procedures to ensure adequate updating and testing of the disaster recovery plan
    - Storage of the data, software, and the disaster recovery plan itself in a safe area where they cannot be destroyed by a catastrophe. This area must be accessible, however, to those who need to use the plan.

12. What is a computer virus?   What is a worm?

    A computer virus is an executable computer program that propagates itself (multiplies), uses a carrier (another computer program), may modify itself during replication, is intended to create some unwanted event. Viruses cause unwanted events -- some are harmless (such as nuisance messages), others are serious (such as the destruction of data). Most viruses attach themselves to other programs or to special parts on disks. As those files execute or are accessed, the virus spreads. Some viruses change their appearances as they spread, making detection more difficult.

    Macro viruses, viruses that are contained in documents or spreadsheet files, can spread when an infected file simply is opened. Macro viruses are the fastest growing type of virus, accounting for more than 75 percent of all virus problems.

    A worm is a special type of virus that spreads itself without human intervention. Worms spread when they install themselves on a computer and then send copies of themselves to other computers, sometimes by e-mail, sometimes via security holes in software.

13. Explain how a denial-of-service attack works.

    A DOS attacks works by an attacker attempting to disrupt the network by flooding it with messages so that the network cannot process messages from normal users. The simplest approach is to flood a Web server, mail server, and so on, with incoming messages. The server attempts to respond to these, but there are so many messages that it cannot.

14. How does a denial-of-service attack differ from a distributed denial-of-service attack?

    While the source of a denial-of-service (DoS) attack could be a single computer, a distributed denial-of-service (DDoS) attack could involve hundreds of computers on the Internet simultaneously sending messages to a target site. A DDoS hacker plants DDoS agent software on these computers and then controls the agents with DDoS handler software, which can send instructions to the agent software on the computers controlled by the hacker for purposes of launching a coordinated attack.

15. What is a disaster recovery firm? When and why would you establish a contract with them?

    Disaster recovery firms provide second level support for major disasters. Building a network that has sufficient capacity to quickly recover from a major disaster such as the loss of an entire data center is beyond the resources of most firms. Therefore, contracts with disaster recovery firms are established.

    Disaster recovery firms provide a full range of services. At the simplest, they provide secure storage for backups. Full services include a complete networked data center that clients can use when they experience a disaster. Once a company declares a disaster, the disaster recovery firm immediately begins recovery operations using the backups stored on site and can have the organization's entire data network back in operations on the disaster recovery firm's computer systems within hours. Full services are not cheap, but compared to the potentially millions of dollars that can be lost per day from the inability to access critical data and application systems, these systems quickly pay for themselves in time of disaster.

16. What is online backup?

Online backup allows you to back up data to a server across the Internet. Generally, software is installed on the client which allows the user to select which files/folders to backup.

17. People who attempt intrusion can be classified into four different categories. Describe them.

There are four types of intruders who attempt to gain unauthorized access to computer networks. The first are casual computer users who have only a limited knowledge of computer security. They simply cruise along the Internet trying to access any computer they come across. Their unsophisticated techniques are the equivalent of trying doorknobs, and only those networks that leave their front doors unlocked are at risk.

The second type of intruders are experts in security, but whose motivation is the thrill of the hunt. They break into computer networks because they enjoy the challenge. Sometimes they also enjoy showing off for friends or embarrassing the network owners. Fortunately, they usually cause little damage and make little attempt to profit from their exploits.

The third type of intruder is the most dangerous. They are professional hackers who break into corporate or government computer for specific purposes, such as espionage or fraud. Less than 5 percent of intrusions by these professionals are detected, unless of course, they have been hired to destroy data or disrupt the network.

The fourth type of intruder is also very dangerous. These are organization employees who have legitimate access to the network, but who gain access to information they are not authorized to use. This information could be used for their own personnel gain, sold to competitors, or fraudulently changed to give the employee extra income. Most security break-ins are caused by this type of intruder.

With a denial-of-service attack, a hacker attempts to disrupt the network by sending messages to the network that prevent other's messages from being processed. The simplest approach is to flood a server with incoming messages. When an external computer sends a message to a computer connected with a proxy server, it addresses the message to the proxy server. The proxy server receives the incoming message, and determines if the packet should be permitted inside. A proxy server receiving the flood of messages will log the attack and discard the messages (does not permit the messages inside.

18. There are many components in a typical security policy. Describe three important components.

Major elements of a security policy are:

- The name of the decision-making manager who is in charge of security.
- An incident reporting system and a rapid response team that to respond to security breaches in progress.
- A risk assessment with priorities as to which components are most important.
- Effective controls placed at all major access points into the network to prevent or deter access by external agents.

- Effective controls placed within the network to ensure internal users cannot exceed their authorized access.
- An acceptable use policy that explains to users what they can and cannot do.
- A plan to routinely train users on security policies and build awareness of security risks.
- A plan to routinely test and update all security controls that includes monitoring of popular press and vendor reports of security holes.

19. What are the three major aspects of intrusion prevention (not counting the security policy)?

    The three main aspects of preventing unauthorized access: securing the network perimeter, securing the interior of the network, and authenticating users.

20. How do you secure the network perimeter?

    There are three basic access points into most organizational networks: from LANs, the Internet, and WLANs. One important element of preventing unauthorized users from accessing an internal LAN is through physical security. A firewall is commonly used to secure an organization's Internet connection. NAT is a common security measure that can be used as well.

21. What is physical security and why is it important?

    Physical security refers to policies and procedures that are designed to prevent outsiders from gaining access to the organization's offices, server room, or network equipment facilities. Good security requires implementing the proper access controls so that only authorized personnel can enter closed areas where servers and network equipment are located or access the network. Network components themselves also have a level of physical security. Computers can have locks on their power switches or passwords that disable the screen and keyboard.

22. What is eavesdropping in a computer security sense?

    Eavesdropping refers to the process of unauthorized tapping into a computer network through local cables that are not secured behind walls or in some other manner.

23. What is a sniffer?

    A sniffer program records all messages received for later (unauthorized) analysis. A computer with a sniffer program could then be plugged into an unattended hub or bridge to eavesdrop on all message traffic.

24. How do you secure dial-in access?

    Some dial-up modem controls include changing the modem telephone numbers periodically, keeping the telephone numbers confidential, and requiring the use of computers that have an electronic identification chip for all dial-up ports. Another strategy is to use a call-back modem.

25. What is a firewall?

A firewall is a router, gateway, or special purpose computer that examines packets flowing into and out of a network and restricts access to the organization's network. The network is designed so that a firewall is placed on every network connection between the organization and the Internet. No access is permitted except through the firewall. Some firewalls have the ability to detect and prevent denial-of-service attacks, as well as unauthorized access attempts. Two commonly used types of firewalls are packet level, and application level.

26. How do the different types of firewalls work?

A packet-level firewall examines the source and destination address of every network packet that passes through it. It only allows packets into or out of the organization's networks that have acceptable source and destination addresses.

An application level firewall acts as an intermediate host computer between the Internet and the rest of the organization's networks. Anyone wishing to access the organization's networks from the Internet most login to this firewall, and can only access the information they are authorized for based on the firewall account profile they access.

The NAT firewall uses an address table to translate the private IP addresses used insidethe organization into proxy IP addresses used on the Internet

27. What is IP spoofing?

IP spoofing means to fool the target computer (and any intervening firewall) into believing that messages from the intruder's computer are actually coming from an authorized user inside the organization's network. Spoofing is done by changing the source address on incoming packets from their real address to an address inside the organization's network. Seeing a valid internal address, the firewall lets the packets through to their destination. The destination computer believes the packets are from a valid internal user and processes them.

The goal of an intruder using IP spoofing is to send packets to a target computer requesting certain privileges be granted to some user (e.g., setting up a new account for the intruder or changing access permission or password for an existing account). Such a message would not be accepted by the target computer unless it can be fooled into believing that the request is genuine.

28. What is a NAT firewall and how does it work?

The *NAT firewall* (sometimes referred to as a proxy server) uses an address table to translate the private IP addresses used inside the organization into proxy IP addresses used on the Internet.
When a computer inside the organization accesses a computer on the Internet, the NAT firewall changes the source IP address in the outgoing IP packet to its own address.
When the external computer responds to the request, it addresses the message to the NAT firewall's IP address. The NAT firewall receives the incoming message, and after ensuring the packet should be permitted inside, changes the destination IP address to the private IP address of the internal computer and changes the TCP port id to the correct port id before

transmitting it on the internal network. This way, systems outside the organization never see the actual internal IP addresses, and thus they think there is only one computer on the internal network.

29. What is a security hole and how do you fix it?

Many commonly used operating system have major security problems (called security holes) well known to potential intruders; UNIX systems are among the worst. Many security holes have been documented and "patches" are available from vendors to fix them, but network managers may be unaware of all the holes or simply forget to regularly update their systems with new patches.

Many security holes are highly technical; for example, sending a message designed to overflow a network buffer, thereby placing a short command into a very specific memory area that unlocks a user profile. Others are rather simple, but not obvious.

Other security holes are not really holes, but simply policies adopted by computer vendors that open the door for security problems, such as computer systems that come with a variety of pre-installed user accounts.

30. Explain how a Trojan horse works.

Trojans are remote access management consoles that enable users to access a computer and manage it from afar. Trojans are often concealed in other software that unsuspecting users download over the Internet. Music and video files shared on the Internet are common carriers of Trojans. When the user downloads and plays the music file, it plays normally and the attached Trojan software silently installs a small program that enables the attacker to take complete control of the user's computer, so the user is unaware that anything bad has happened.

31. Compare and contrast symmetric and asymmetric encryption.

A *symmetric (or single key) encryption algorithm* is one in which the key used to encrypt a message is the same as the one used to decrypt it. Both parties to the transmission must possess the same key for encryption and decryption. The key must be kept secret, leading to a need for key management.

An *asymmetric algorithm* is one in which the key used to decrypt a message is *different* from the one used to encrypt it. *Public key encryption* is the most common for asymmetric encryption. , there are two keys. One key (called the public key) is used to encrypt the message and a second, very different private key is used to decrypt the message. The net result is that if two parties wish to communicate with one another, there is no need to exchange keys beforehand. All public keys are published in a directory. Each knows the other's public key from the listing in the public directory and can communicate encrypted information immediately. The key management problem is reduced to the on-site protection of the private key.

32. Describe how symmetric encryption and decryption works.

A symmetric algorithm is an algorithm in which the key used to decrypt a particular bit stream is the same as the one used to encrypt it. Using any other key produces plaintext that appears as random as the ciphertext. No keys are exchanged between the sender and the receiver.

Encryption is the process of disguising information into ciphertext, whereas decryption is the process of restoring it to readable form (plaintext). An encryption system has two parts: the algorithm itself and the key, which personalizes the algorithm by making the transformation of data unique. Two pieces of identical information encrypted with the same algorithm but with different keys produce completely different ciphertexts. When using most encryption systems, communicating parties must share this key. If the algorithm is adequate and the key is kept secret, acquisition of the ciphertext by unauthorized personnel is of no consequence to the communicating parties.

33. Describe how asymmetric encryption and decryption works.

In asymmetric encryption and decryption there are two keys. One key (called the*public key*) is used to encrypt the message and a second, very different *private key* is used to decrypt the message. Public key systems are based on one-way functions. Even though you originally know both the contents of your message and the public encryption key, once it is encrypted by the one-way function, the message cannot be decrypted without the private key. One-way functions, which are relatively easy to calculate in one direction, are impossible to "uncalculate" in the reverse direction.

All public keys are published in a directory. When Organization A wants to send an encrypted message to Organization B, it looks through the directory to find its public key. It then encrypts the message using B's public key. This encrypted message is then send through the network to Organization B, which decrypts the message using its private key.

34. What is key management?

Key management is concerned with dispersing and storing keys carefully. Because the DES algorithm is known publicly, the disclosure of a secret key can mean total compromise of encrypted messages. Managing this system of keys can be challenging, especially with symmetric algorithms.

35. How does DES differ from 3DES? From RC4? From AES?

DES uses a 56-bit key while 3DES uses a 168-bit key (3 x 56).

RC4 uses keys from 40 to 256 bits in length.

AES uses the Rijndael algorithm and has key sizes of 128, 192, and 256 bits.

36. Compare and contrast DES and public key encryption.

DES is a symmetric algorithm, which means that the key used to decrypt a particular bit stream is the same as the one used to encrypt it. Using any other key produces plaintext that appears as random as the ciphertext. Because the DES algorithm is known publicly, the disclosure of a secret key can mean total compromise of encrypted messages. Managing this system of keys can be challenging.

Public key encryption is inherently different from secret key systems like DES. because it is asymmetric; there are two keys. One key (called the public key) is used to encrypt the message and a second, very different private key is used to decrypt the message. Public key systems are based on one-way functions. Even though you originally know both the contents of your message and the public encryption key, once it is encrypted by the one-way function, the message cannot be decrypted without the private key. One-way functions, which are relatively easy to calculate in one direction, are impossible to "uncalculate" in the reverse direction. Public key encryption is one of the most secure encryption techniques available, excluding special encryption techniques developed by national security agencies.

Note: DES key length is 56 bits (168 bits for 3DES) while private key length for public key encryption is 512 or 1,024 bits.

The American government has tried to develop a policy to require key escrow. With key escrow, any organization using encryption must register its keys with the government. This enables the government, after receiving a legally authorized search warrant, to decrypt and read any messages sent by that organization. Without key escrow, the government is worried that criminal organizations will use encryption to prevent police agencies from performing legally authorized wiretaps. Free speech advocates are concerned that key encryption will be abused by police agencies who will illegally monitor transactions by otherwise innocent citizens. Key escrow is a good idea if it is used to combat the criminal organizations, but key escrow is not a good idea if it increases "big brother" and abuses our right to free speech.

37. Explain how authentication works.

Public key encryption permits authentication (or digital signatures). When one user sends a message to another, it is difficult to legally prove who actually sent the message. Legal proof is important in many communications, such as bank transfers and buy/sell orders in currency and stock trading, which normally require legal signatures. Thus a digital signature or authentication sequence is used as a legal signature on many financial transactions. This signature is usually the name of the signing party plus other key-contents such as unique information from the message (e.g., date, time, or dollar amount). This signature and the other key-contents are encrypted by the sender using the private key. The receiver uses the sender's public key to decrypt the signature block and compares the result to the name and other key contents in the rest of the message to ensure a match.

38. What is PKI and why is it important?

PKI stands for Public Key Infrastructure. PKI refers to the encryption infrastructure that has developed around the most popular form of asymmetric encryption (also called public key encryption) called RSA. RSA was invented at MIT in 1977. The patent expired on the technology in 2000 and many new companies have now entered the market and public key software has dropped in price.

Public key encryption is different from symmetric single key systems. Because pubic key encryption is asymmetric, there are two keys. One key (called the public key) is used to encrypt the message and a second, very different private key is used to decrypt the message. Public key encryption is one of the most secure encryption techniques available.

39. What is a certificate authority?

A certificate authority (CA) is a trusted organization that can vouch for the authenticity of the person or organization using authentication (e.g., VeriSign). A person wanting to use a CA registers with the CA and must provide some proof of identify. There are several levels of certification, ranging from a simple confirmation from valid email address to a complete police-style background check with an in-person interview. The CA issues a digital certificate that is the requestor's public key encrypted using the CA's private key as proof of identify. This certificate is then attached to the user's email or Web transactions in addition to the authentication information. The receiver then verifies the certificate by decrypting it with the CA's public key -- and must also contact the CA to ensure that the user's certificate has not been revoked by the CA.

40. How does PGP differ from SSL?

<u>Pretty Good Privacy</u> (PGP) is freeware public key encryption package developed by Philip Zimmermann that is often used to encrypt e-mail. Users post their public key on Web pages, for example, and anyone wishing to send them an encrypted message simply cuts and pastes the key off the Web page into the PGP software which encrypts and sends the message.

<u>Secure Sockets Layer</u> (SSL) operates between the application layer software and the transport layer. SSL encrypts outbound packets coming out of the application layer before they reach the transport layer and decrypts inbound packets coming out of the transport layer before they reach the application layer. With SSL, the client and the server start with a handshake for PKI authentication and for the server to provide its public key and preferred encryption technique to the client (usually RC4, DES or 3DES). The client then generates a key for this encryption technique, which is sent to the server encrypted with the server's public key. The rest of the communication then uses this encryption technique and key.

41. How does SSL differ from IPSec?

SSL differs from IPSec in that SSL is focused on Web applications, while IPSec can be used with a much wider variety of application layer protocols.

42. Compare and contrast IPSec tunnel mode and IPSec transfer mode.

- IPSec *transport mode* provides only encryption of the message payload, while *tunnel mode* additionally encrypts the final destination by encrypting the entire IP packet which is then included in a new added packet that is address to an IPSec agent rather than to the true final destination.
- In *transport mode* leaves the IP packet header unchanged so it can be easily routed through the Internet. It adds an additional packet (either an Authentication Header (AH)

or an Encapsulating Security Payload (ESP)) at the start of the IP packet that provides encryption information for the receiver.

- In *tunnel mode*, the newly added IP packet conceals the final destination (which is encrypted since it just identifies the IPSec encryption agent as the destination, not the final destination. As with the transport mode, encryption information is added in the form of an AH or ESP. When the IPSec packet arrives at the encryption agent, the encrypted packet is decrypted and sent on its way. In tunnel mode, attackers can only learn the endpoints of the tunnel, not the ultimate source and destination of the packets

43. What are the three major ways of authenticating users? What are the pros and cons of each approach?

The three major ways to authenticate users is to base account access on something you know, something you have, or something you are.

The most common approach is something you know—usually, a password. Requiring passwords provides mid-level security, at best; it won't stop the professional intruder, but it will slow amateurs.

More and more systems are requiring users to enter a password in conjunction with something they have, such as a smart card. Intruders must have access to both before they can break in.

In high-security applications, a user may be required to present something they are—such as a finger, hand, or retina of their eye for scanning by the system. These biometric systems scan the user to ensure that the user is the sole individual authorized to access the network account.

44. What are the different types of one-time passwords and how do they work?

Using a one-time password users connect to the network as usual, and after the user's password is accepted, the system generates a one-time password. The user must enter this password to gain access, otherwise the connection is terminated.

Other systems provide users with a unique number that must be entered into a separate handheld device (called a token system), which in turn displays the password for the user to enter. Other systems used time based tokens in which the one-time password is changed every 60 seconds. The user has a small device (often attached to a key chain) that is synchronized with the server and displays the one-time password.

With any of these systems, an attacker must know the user's account name, password, and have access to the user's password device before he or she can login.

45. Explain how a biometric system can improve security. What are the problems with it?

In high security applications, a user may be required to present something they are, such as a finger, hand, or the retina of their eye for scanning by the system. These biometric systems scan the user to ensure that user is the sole individual authorized to access the network account. While someone can obtain someone else's password and access card, it is most difficult to acquire another person's handprint or eye retina print. While most biometric systems are developed for high security users, several low cost biometric systems to recognize fingerprints are now on the market.

46. Why is the management of user profiles an important aspect of a security policy?

Each user's profile specifies what data and network resources he or she can access, and the type of access allowed (read only, write, create, delete).

47. How does network authentication work and why is it useful?

Instead of logging into a file server or application server, network authentication requires that users login to an authentication server. This server checks the user id and password against its database and if the user is an authorized user, issues a certificate. Whenever the user attempts to access a restricted service or resource that requires a user id and password, the user is challenged and his or her software presents the certificate to the authentication server. If the authentication server validates the certificate then the service or resource lets the user in. In this way, the user no longer needs to enter his or her password for each new service or resource he or she uses. This also ensures that the user does not accidentally give out his or her password to an unauthorized service—it provides mutual authentication of both the user and the service or resource.

48. What is social engineering? Why does it work so well?

Social engineering refers to breaking security simply by asking. For example, hackers routinely phone unsuspecting users and imitate someone else (e.g., a technician, a boss, a network expert) and ask for a password. Most security experts no longer test for social engineering attacks; they know from past experience that social engineering will eventually succeed in any organization and therefore assume that hackers can gain access at will to "normal" user accounts. Training end users not to divulge passwords may not eliminate social engineering attacks, but it may reduce its effectiveness so that hackers give up and move on to easier targets.

A skilled social engineer is like a good con artist, he can manipulate people.

49. What techniques can be used to reduce the chance that social engineering will be successful?

Most security experts no longer test for social engineering attacks; they know from experience that social engineering will eventually succeed in any organization and therefore assume that attackers can gain access at will to normal user accounts. Training end users not to divulge passwords may not eliminate social engineering attacks, but it may reduce its effectiveness so that hackers give up and move on to easier targets. Acting out social engineering skits in front of users often works very well; when a group of employees sees how they can be manipulated into giving out private information, it becomes more memorable and they tend to become much more careful.

50. What is an intrusion prevention system?

Assuming that prevention efforts will not be sufficient to avoid all intrusions, intrusion prevention systems (IPSs) can be used to monitor networks, circuits, and particular applications and report detected intrusions.

51. Compare and contrast a network-based IPS, a host-based IPS, and an application-based IPS.

In each case the IPS reports intrusions to an IPS management console:

- The *network-based IPS* monitors key network circuits through IPS sensors that are placed on the key circuits to monitors all network packets on that circuit.
- The *host-based IPS* monitors a server and incoming circuits. It is installed on the server that it is monitoring.
- An *application-based IPS* is a specialized host-based IPS that monitors one application on its server such as a Web server.

52. How does IPS anomaly detection differ from misuse detection?

Anomaly detection compares monitored activities with a known "normal" set of activities for a stable network environment while misuse detection compares monitored activities with signatures of prior known attacks. Anomaly detection looks for extreme changes in certain kinds of behavior while misuse detection guards against a repeat of prior intrusions.

53. What is computer forensics?

Computer forensics is the use of computer analysis techniques to gather evidence for criminal and/or civil trials and includes the following steps:

- Identify potential evidence.
- Preserve evidence by making backup copies and use those copies for all analysis.
- Analyze the evidence.
- Prepare a detailed legal report for use in prosecutions.

54. What is a honey pot?

A honey pot is a server that contains highly interesting fake information available only through illegal intrusion to "bait" or "entrap" the intruder and also possibly divert the hacker's attention from the real network assets. The honey pot server has sophisticated tracking software to monitor access to this information that allows the organization and law enforcement officials to trace and document the intruder's actions. If the hacker is subsequently found to be in possession of information from the honey pot, that fact can be used in prosecution.

55. What is desktop management?

Desktop management refers to security measures at the individual client level. Strong desktop management may include the use of thin clients (perhaps even network PCs that lack hard disks). Centralized desktop management, in which individual users are not permitted to change the settings on their computers with regular reimaging of computers to prevent Trojans and viruses and to install the most recent security patches. All external software downloads will likely be prohibited.

56. A few security consultants have said that broadband and wireless technologies are their best friends. Explain.

If the role of the security consultants is to penetrate security, then broadband and wireless technologies offer important advantages. When guided media are used, physical access or proximity is necessary for interception.

As is mentioned in Chapter 3 of the text (Physical Layer), "wireless media (radio, infrared, microwave, and satellite) are the least secure because their signals are easily intercepted." Chapter 9 (The Internet) also warns regarding wireless LANs: "Because anyone within range of a WLAN can receive transmissions, eavesdropping a serious threat. IEEE 802.11 encrypts all transmissions using a 40-bit encryption scheme so that only those computers that have the key can decode and read the messages." This chapter documents the limitations of a 40-bit key.

Broadband technologies use broadcast signals. The signal in the medium is a composite (in terms of modulation technique and whatever multiplexing is present) of all the separate communications supported by that circuit at the instant of sampling. Anyone with physical access to the circuit can monitor the composite signal and, as in the case of wireless transmission, eavesdrop. In any broadcast protocol, all computers connected to the medium "hear" all the messages but only act to receive and process the messages with their own addresses (respectively) as the destination. You might encourage students to think of the sound in a room (let us say, at a party or social gathering) as everything transmitted through the medium of air. You can hear many conversations. Typically you only answer when someone says your name (= destination address for the utterance!) However, you might overhear some conversation of interest and decide to eavesdrop, even though the conversation is not intended for your ears. Clearly the broadcast medium has higher security risks.

57. Most hackers start their careers breaking into computer systems as teenagers. What can we as a community of computer professionals do to reduce the temptation to become a hacker?

Computer professionals can reduce the temptation to become a hacker by increasing the costs involved in the activity. This can be done by rigorously following and enforcing security measures throughout the organization. In addition, professionals can work with high schools and colleges to identify the most promising computer students and get them involved in more productive organizational tasks and projects.

58. Some experts argue that CERT's posting of security holes on its Web site causes more security break-ins than it prevents and should be stopped. What are the pros and cons on both sides of this argument? Do you think CERT should continue to post security holes?

While CERT's posting obviously disseminate technical information that could inform someone who wished to exploit vulnerabilities (and thus increase risk to those who do not take advantage of the service), it is necessary to help conscientious security operations reduce risk. I assume the information about the security holes would be disseminated by various means (to interested hackers) even if CERT's postings were not available.

The CERT postings reduce the time it takes for potential victims to identify new security hole threats and respond to them, thus helping to reduce the adverse impacts of viruses and worms which can be quickly propagated around the world through the Internet.

59. What is one of the major risks of downloading unauthorized copies of music files from the Internet (aside from the risk of jail, fines, and lawsuits)?

    One of the major risks of downloading files of this nature is the potential to introduce a Trojan Horse into the network.

60. Although it is important to protect all servers, some servers are more important than others. What server(s) are the most important to protect and why?

    To answer this question, return to the risk assessment process and ask which server loss will cause the greatest damage to the organization. Is the highest priority assigned to...?

    - Loss of customer/client information
    - Loss of online access
    - Interruption of or loss of support for activity (manufacturing, health care, etc.)
    - Invasion of privacy/loss of confidentiality
    - Integrity of accounts
    - Proprietary secrets

Mini-Cases

I. Belmont State Bank

Perform a risk assessment.

The steps of the risk assessment include:
1. Develop risk measurement criteria
2. Inventory IT assets
3. Identify threats
4. Document existing controls
5. Identify improvements

Step 1:
Student answers can vary considerably in such a large case. Following is a sample of what they may consider.

Impact areas	*Priority*	*Low Impact*	*Medium Impact*	*High Impact*
Financial	*High*			*An outsider gains control of a server.*
Productivity	*Medium*		*The website is affected by a DOS attack and slows customer web responses*	
Reputation	*Medium*			*A security event that affects our bank is headlined in the news*
Safety	*Low*			
Legal	*Medium*			*A security breach occurs and all of our customer records are potentially accessed.*

Step 2:
These can be classified as hardware, circuits, network software, client software, org data, and mission-critical applications.

Step 3:
Some of the top threats include viruses, theft of equipment and information, device failure, and natural disaster.

Step 4:
Existing controls may include:

1. *Disaster recovery plan*
2. *Halon fire system in server room; sprinklers in the rest of building*
3. *Not on or below ground level*
4. *UPS*
5. *Contract guarantees from interexchange carriers*
6. *Extra fiber backbone laid in different conduits*
7. *Virus checking software present and updated*
8. *Extensive user training about viruses and reminders in monthly newsletter*
9. *Strong password software*
10. *Extensive user training about password security and reminders in monthly newsletter*
11. *Application-layer firewall*

Step 5:
Improvements can include any number of updates to controls. An example would include moving from physical keys to card access or biometrics systems.

II. Western Bank

Design the key security hardware and software the bank should use.

Student answers will vary considerably for this answer. Questions to consider when developing this plan: What assets are being protected? Consider the hardware, software, and encryption options documented in the chapter. Specify any physical security needed. Compare firewalls and NAT proxy servers. Is there a use for dial-in security in this environment? What level of encryption is needed and for what transmissions? Should the bank have a misuse detection or anomaly detection IDS? Does the bank need a honey pot?

The students' answers should include numerous levels of security. An answer similar to what is seen in Figure 11.16 would be good, with the web server in the DMZ.

III. Classic Catalog Company, Part 1

Perform a risk assessment.

Student answers will vary, but be similar to those in minicase 1.

IV. Classic Catalog Company, Part 2

Outline a brief business continuity plan, including controls to reduce the risks in advance and develop a disaster recovery plan.

Have the students develop a business continuity plan that accounts for natural disasters, theft, viruses, and denial-of-service attacks. Use the elements of a good disaster recovery plan given in the chapter to develop a plan for Classic Catalog. Discuss back up procedures, firewall design, employee password procedures and intrusion preventions for this firm. How will you design the redundancy for the company? Do you recommend any outsourcing in this situation?

V. Classic Catalog Company, Part 3

Outline a brief security policy and controls that would prohibit unauthorized access.

Have students develop a security policy that defines the assets to be protected and the controls needed to do so. Use the figure in the chapter to design the security policy. Consider the physical layout of the company, relative to its networking choices. Would you recommend wireless? How would you ensure the security of the wireless system? What kinds of controls would you recommend to protect the office computers and warehouse computers?

VI. Classic Catalog Company, Part 4

What patching policy would you recommend for Classic Catalog?

Patching should be done on a regular, automated basis. Most systems today have the capability to automatically update. These updates should be periodically, manually checked to insure that they complete. In addition, network managers may be required to manually update some systems.

Do you anticipate that there will be any security issues associated with the patching of known network flaws? What is a good managerial approach to reducing the corporation's risk exposure from patching?

Many applications today include automatic patching or updating. These controls should be turned on to allow this. For those applications and systems that don't offer automatic updates, a schedule should be developed. For operating systems, patches should be searched for weekly. Depending on the level of importance of other applications, the patching schedule may be more or less often.

VII. Personal Password Storage and Protection

To help us not forget our many passwords, there are several companies that provide password managers. Find the top 5 password manager programs, compare their features and costs, and make a presentation of your findings to your classmates.

Following is a list of 5 top password managers based on the following article:

http://www.informationweek.com/applications/top-5-password-managers/d/d-id/1106016?page_number=1

Lastpass

KeePass

Clipperz

RoboForm

Windows 8 and Live ID

Next Day Air Service Case Study

1. Prepare a report outlining the major security threats faced by NDAS. Be sure to identify those that you think are major threats and those that are minor threats.

Obviously the president considers web page defacement as a major threat!
Other major threats could be anything that interrupts service traffic (as described in earlier chapters, such as communications failures or power supply interruptions. Potential disasters could afflict NDAS - some cities where NDAS has an office are prone to flood, tornado, or other natural disaster threats. There would be a need to safeguard information assets (transactions, customer information, inventory, etc.) With regard to Memphis, consider the potential for a disastrous earthquake (rare in this region).
Less urgent would be more general vigilance regarding potential internal threats.

2. Prepare a partial risk assessment for NDAS that includes their major assets, threats and controls. You will need to make some reasonable assumptions.

Follow the instructions in Chapter 11 to create a risk assessment. Consider the dispersed nature of NDAS operations (offices spread across the country in 16 different cities.) Focus on major threats as identified in the report generated for exercise 1.

3. Develop a set of security controls, for use in the NDAS main office and for its Web site, designed to control risks due to disruption, destruction, and disaster.

Cover policies, disaster recovery plan, fire and/or flood protection, decentralization of appropriate assets, use of contract services (outsourcing) for disaster recovery and decentralization, firewalls and NAT proxy servers. Unauthorized access carries with it the potential for disruption, destruction and disaster. IN addition, if an unauthorized user gains access and goes undetected it could pose a serious corporate sabotage threat to NDAS. For this reason security controls related to unauthorized access must include the coverage of policies, disaster recovery plan, scrutinization of contract services (outsourcing) at all levels and firewall process and procedures. In the event that an unauthorized user attempts to gain access to NDAS's network, the firewall system should have appropriate rules set to detect it and set off the alarm. Staff should respond immediately to identify the intruder and to take appropriate action to control the threat.

4. What are the main concerns regarding the accounts payable employees that are being laid off? How can the risks related to this layoff be mitigated?

As with any employee that is leaving the company, the employer should be concerned. For those employees being laid off, I suggest disabling their network account as soon as they are notified. This prevents a disgruntled employee from copying, modifying, or deleting files on the network. It is also recommended to go ahead and ask them to leave the building permanently so that they cannot do the same to paper documents.

Additional Content

Teaching Notes

I usually spend 4 hours of class time on this chapter.

My goals in this chapter are to discuss specific technologies and approaches to security. Risk assessment sounds more exciting than it really is. I try to express the importance of it on the students, even though it is a bit mundane. Likewise you could describe some of the organizations that can be hired to test security.

I try to emphasize the need for controls to address risks due to disruption, destruction, and disaster. Most students overlook these because preventing unauthorized access is much sexier and gets more attention from them and from the popular press. Unauthorized access is important, but it is less likely to occur than disruption, destruction, and disaster. Most organizations would cease to function without their networks, so networks are their lifeblood. One disaster that takes out the network can bankrupt the company if there is no recovery plan. I try to ensure that students can differentiate among controls to prevent problems, detect problems, and correct problems. Each is important, but the emphasis is different.

Security to prevent unauthorized access remains a rapidly changing area. Threats from outside have increased as the Internet has become a regular part of most organizations, and are approximately equal to the very real threat from the inside from current or former employees. Firewalls and proxy servers are interesting discussion topics. I use the section on proxy servers to reinforce the TCP/IP and Internet concepts from earlier chapters. This is a hot area to watch, given the many new security holes discovered in commercial products and news stories about break-ins and worms. Network security in general, and firewalls in particular are likely to change over the next few years and bear watching closely. The students can share information about their experiences with different virus check software packages and personal firewalls. A sniffer demonstration (if this can be arranged with help from your campus computer center) dramatically illustrates the vulnerability of data transmissions in networks. Likewise, having two student role play a social engineering episode can be very powerful (you could use the Focus box in the text as a pseudo script).

One of the most striking things to me is how the number of security incidents keeps growing.

I recommend students visit appropriate web sites and discuss what they find there:

- www.microsoft.com/security/bulletins/
- www.cert.org - the CERT® Coordinator Center at the Software Engineering Center, Carnegie Mellon University
- csrc.nist.gov/ - The National Institutes of Standards and Technology web site on Computer Resource Security
- www.happyhacker.org - Carolyn Meinel's site

War Stories

Backups Aren't Fool-Proof

(Objective: illustrate the importance of disaster recovery drills)
In August 1995, a fire destroyed our building at The University of Georgia. The fifth floor was completely destroyed by fire, and the remaining floors were badly damaged by 1.5 million gallons of water (it took 100 firefighters 8 hours to put out the fire). Our two Web servers were destroyed, but we had off-site backups, exactly as planned.

The first Web server backup was stored on another server in the Computer Science building. The morning after the fire, we got permission from the library to use their Web server. We simply downloaded our directory structure onto their Web sever and changed the University DNS server so that any request for our Web server was answered with the IP address for the library's Web server (we also changed their home page to provide a link to our new home page on their server, along with an explanatory note as to what had happened). The server was back up in less than 24 hours. When our network was restored 5 weeks later, we simply moved everyone back to our new server and changed the University DNS server entry back to our new IP address. No problems.

Break-in on Ginger

(Objective: describe a specific security break-in via a security hole)
Ginger is a UNIX Web server and mail server in our department that is used by a few faculty and students. One of those students, who ironically was taking my telecom class and doing a group project on network security, decided to break-in to Ginger. We were running an old version of UNIX with a security hole that we didn't know about. It turns out that anyone with an account on the server could access the password file for the entire computer. Even though the file was encrypted, this was still a big security hole.

The student copied the file onto a diskette and took it home. He used Crack, a software tool available on the Internet, to decrypt the password file. Crack is not very efficient, but if you let it run for a few days, it is often successful at decrypting passwords. In this case, it worked. The student was able to get the root password (i.e., the system administrator's password). Fortunately, all he did was poke around, but he could have destroyed the entire system.

Ginger was configured so that no one could log in as root (i.e., the system administrator) without first logging into a regular user account. Ginger was also configured to record all log-ins to the root account and the network administrator routinely checked the log files for unusual activity (which is not often done in most organization). He discovered that someone had logged into the root account after midnight one weekend. He knew he didn't do it and checked with the only other person with root access and discovered that he had not either. A quick check of the log files revealed that the student had logged into the root account from his own account. Had he used someone else's account or erased the file, it still would have been discovered, but would have been much harder to identify the hacker.

When confronted with the evidence, he quickly confessed, and implicated an accomplice (another student). For some unknown reason, the University student disciplinary board was extremely reluctant to prosecute the pair under the University's computer hacking policy. The only alternative was to notify the FBI and have them prosecuted under the Federal hacking law (a felony) which seemed too extreme. No action was taken.

Break-in on Blaze
(Objective: describe a specific security break-in via a bad user password)
Blaze is a UNIX Web server and mail server in our department that is used by a most of our faculty and a few students (it got its name after our fire). As part of our routine security detection process we log all telnet accesses and periodically check to see if the accesses seem reasonable.

One check indicated a lot of logins from Romania and Japan to accounts held by two of our faculty. We immediately checked with the account owners and quickly found that they were not in Japan or Romania and had not authorized users from there to access their accounts. The users whose account had been broken into had chosen very poor passwords (e.g., dawgs, the slang name of our football team and an easy word to guess).

We were not completely sure that password guessing had been the means of attack and did not know the extent of the penetration. Therefore, we took the server down, installed the newest version of the operating system (which had patched several security holes in the old version), and recreated all user accounts. We put in a new password security system that forced users to select passwords that were not words in a dictionary. Quite naturally, there was some user resistance but we felt that security concerns outweighed ease of use considerations.

AT&T Network Outage
(Objective; illustrate the need for a disaster recovery plan)
In early April 1998, AT&T's frame relay network in the United States crashed due to a broadcast storm. The problem started when a technician updated software on one of the 145 Cisco frame relay switches in the AT&T network. The switch software contained a bug that did not prevent a "loopback," a situation in which the switch receives its own broadcast message (you will recall that until recently, broadcast messages were unusual in switched networks, so much of this type of switch software had not been really tested in the field). The switch retransmitted the same broadcast message to other switches, and to itself. It turned out that all of the Cisco switches had the same bug so the entire network suddenly began transmitting and retransmitting the same broadcast message until the network collapsed.

It took more than a day to restore service to customers. For that day, businesses that depended on the network and had no disaster recovery plan could not do business. Many travel agents were unable to make bookings or check their customers' reservations. MasterCard, one of AT&T's largest frame relay customers, had a disaster recovery plan in place and a series of redundant circuits leased from another carrier. When AT&T's network failed, it switched over to the redundant circuits and continued to operate. MasterCard customers noticed only slight delays.

As well as being troublesome and embarrassing, to both Cisco and AT&T, the outage cost AT&T millions of dollars. AT&T had service level agreements (SLA) with most customers guaranteeing 99.99% network availability and promising to restore lost service within four hours. Because the outage exceeded the SLA, AT&T was required to pay penalties to customers.

Redundancy is Important, Even in Cabling

(Objective: illustrate why redundancy is needed)

Our college of business network is composed of networks in three adjacent buildings: Caldwell Hall (a classroom building), Brooks Hall (faculty offices), and Sanford Hall (classrooms). All but one server (file servers, Web servers, and mail servers) were located in Caldwell Hall. The original network design called for all three buildings to be connected to the campus backbone and to be connected to each other via our own college-wide internal backbone (see the figure below). However, due to cost considerations not all segments were actually built.

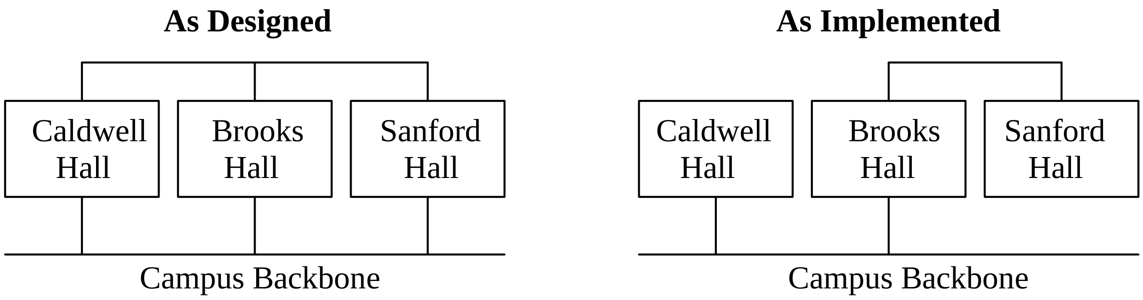

In early 1998, a backhoe cut the cable linking Caldwell Hall to the campus backbone. Suddenly, almost all of our severs were cut off from the Internet and from the rest of the college. The cable was restored in two days, but for those two days, none of the new electronic classrooms in Sanford Hall could access course material on our servers and none of the electronic classrooms in Caldwell Hall could access the Internet. Faculty and staff could only access their email from the classrooms and labs in Caldwell Hall, and even then could only send and receive mail to and from people on the same server.