CISSP

Certified Information Systems Security Professional

Student Guide

Version 1.0

Introduction

Overview

Networking has grown exponentially from its first inception to today's Internet which is nothing more than a vast network spanning all nations in every part of the globe. The Internet connects individuals, groups, corporations, universities, and governments in a way that is both very simplistic and very open. The Internet is an information exchange using a common infrastructure, protocols, and applications. The same is true of the individual entities that comprise the Internet. Yet, where the Internet is an open network, the individual enterprises that comprise it choose not to be open.

Objectives

Upon completing this module, you will be able to:

- Explain the purpose of information security
- Define the CISSP CBK
- Explain security concerns
- Explain the CIA Triad

Outline

The module contains these lessons:

- Information Security
- The CISSP CBK
- Security Concerns
- The CIA Triad

Information Security

Overview

Information security was created to provide networked entities with a solution to the problem of security. Yet, security is far more complicated than allowing or not allowing access to a system or network. Security is a blanket that covers all aspects of protection to an entity. Encompassed security includes items such as fences, walls, and security guards to access control, PBX fraud, and virtual private networks.

This course was created based on a common body of knowledge (CBK) that many security related individuals have amassed over years of in-the-field best practices.

Importance

The Certified Information Systems Security Professional (CISSP) certification is considered the gold standard in the information security realm. Individuals who posses this certification give prospective employers an invaluable tool in validating a candidate's expertise in the all-encompassing realm of securing an enterprise.

Objectives

Upon completing this lesson, you will be able to:

- Understand the CISSP credential
- Understand what the CISSP is
- Understand why you need the CISSP
- Understand how to obtain the CISSP

Outline

The lesson contains these topics:

- The CISSP
- What Is It?
- Why Do I Need It?
- How Do I Get It?

The CISSP

This section discusses the reason the CISSP certification was created and which entities created it.

Computer information security is now more important to companies who have an Internet presence. These companies have seen many successful attacks against large corporations and web sites such as Microsoft, Apple, Google, and Yahoo. They have also seen the devastating monetary damages that such attacks inevitably leave in their wake. No company, no matter how big or how small, is immune to information system attacks. For this reason, most companies do not take security lightly and have asked for help from the general Internet community. The International Information Systems Security Certification Consortium (ISC)² was created to establish a credential that perspective employers can use to determine the eligibility of security-related candidates.

What Is It?

This section will discuss what the CISSP certification is and what benefit it provides the computer community.

The Certified Information Systems Security Professional (CISSP) certification is a like a security badge that prospective employers look for in candidates who wish to obtain a security related position. The certification gives the employer assurance that the candidate is highly qualified in the security profession, has at least three years of experience in the security field, and has enough knowledge to pass a comprehensive security exam covering 10 different domains in the security field.

Why Do I Need It?

This section will discuss why an individual would need or what they would hope to gain from obtaining the CISSP credential.

Skills of a CISSP certified professional:

- Can implement solid security practices
- Can perform in depth risk analysis
- Can configure proper access rights and permissions
- Can implement access control
- Can secure data as it crosses the network
- Can implement proper change control
- Understands methods used to attack resources
- Understands the system development life cycle
- Can perform security audits
- Can develop a business continuity plan
- Understands laws on and about computer crime

.

In today's world where security incidents happen daily and devastating incidents happen much too regularly, a majority of corporations desperately search for talented and experienced security professionals to help them protect their networks and resources. A CISSP certification identifies you as a highly sought out security professional who has successfully met a predefined standard of experience and knowledge. In keeping this certification current, you show your interest in keeping up-to-date in the latest security knowledge and related events that help you achieve high standards for securing the enterprise and its information.

In obtaining this certification, you show current or perspective employers that you can…

- Implement solid security practices
- Perform in depth risk analysis
- Configure proper access rights and permissions
- Implement access control
- Secure data as it crosses the network
- Implement proper change control
- Recognize methods used to attack resources
- Recognize the system's development life cycle
- Perform security audits

- Develop a business continuity plan
- Comprehend the laws on and about computer crime

How Do I Get It?

This section will discuss how one would obtain the CISSP credential.

CISSP Certified

- Have general knowledge on all 10 domains that cover the CBK
- Pass a 250 multiple question exam with a score of at least 70 percent
- Have 3 years of direct security related experience to qualify to take the exam

To pass the CISSP exam, you must have general knowledge on the ten domains that cover the Common Body of Knowledge (CBK). You do not have to be an expert in all ten domains in every subject, but in order to pass the test, you must have general knowledge of many different subjects within the CBK.

The exam is made up of 250 multiple-choice questions. You are given six hours to complete it. Remember that each question will have four choices with only one correct answer. Of the 250 questions, only 225 are scored, the remaining 25 questions will be used for research purposes. You will not be able to discern which of the 250 questions are the 25 used for research. In order to pass the exam, you must correctly answer at least 70 percent of the 225 questions. At this time, the cost to take the exam is $499 for early registration (payment received 16 calendar days in advance of the exam date), or $599 if you make full payment less than 16 calendar days prior to the test.

To quality to take the CISSP exam, you must supply information that proves that you have a minimum of four years, three if you have a degree in computer science, of direct work related experience in one or more of the ten domains that make up the CBK. This prerequisite ensures that anyone who receives the CISSP certification has real-world experience to offer perspective employers.

Summary

The key points discussed in this lesson are:

- The CISSP credential
- Describing CISSP
- The need for CISSP
- Obtaining the CISSP

The CISSP CBK

Overview

Multiple professional associations formed the International Information Systems Security Certification Consortium (ISC)² in 1989 to develop an accepted industry standard for the practice of information security. To achieve their goal, any individual certified in information security would need to master a certain amount of knowledge to accomplish his or her job. To that end, the Common Body of Knowledge (CBK) was created. The CBK is continually updated to stay current in the rapidly changing atmosphere of information security.

Importance

The CISSP CBK is a compilation of all security information that has relevance to the security information professional collected over time and from around the globe. Today, 10 domains are defined in the CBK, which correlates to how this course is designed:

- Access Control Systems and Methodology
- Telecommunications, Network and Internet Security
- Security Management Practices
- Application Development Security
- Cryptography
- Security Architecture and Models
- Operations Security
- Business Continuity Planning
- Law, Investigation, and Ethics
- Physical Security

Objectives

Upon completing this lesson, you will be able to:

- Understand what the CBK is
- Understand why the CBK was created
- Identity the 10 domains of the CBK

Learner Skills and Knowledge

To fully benefit from this lesson, you must have these prerequisite skills and knowledge:

- Basicinformation security knowledge

Outline

The lesson contains these topics:

- CBK – What Is It?
- CBK – Why Was It Created?
- CBK – What Are the 10 Domains?

CBK – What Is It?

This section will discuss what the Common Body of Knowledge (CBK) is and what security domains make it up.

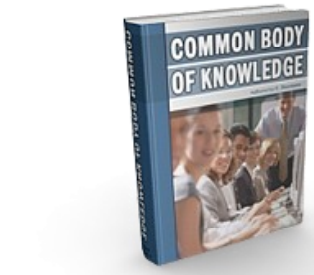

The Compilation and Distillation of All Pieces of Information Security

Information Security:

- Access Control Systems and Methodology
- Application and Systems Development Security
- Business Continuity Planning
- Disaster Recovery Planning
- Cryptography
- Law, Investigation, and Ethics
- Operations Security
- Physical Security
- Security Architecture and Models
- Security Management Practices
- Telecommunications and Network Security

The Common Body of Knowledge (CBK) is the compilation and distillation of all pieces of information security. The information that represents the collective body of knowledge (CBK) important to an Information Security Specialist was collected from states and nations. The CISSP CBK was created in 1989 in anticipation of a security specialization that was to become the CISSP.

CBK – Why Was It Created?

This section will discuss the reason for the creation of the Common Body of Knowledge.

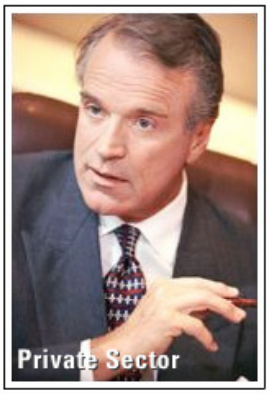

In order to create a security specialization, many security professionals in both the public and private sector banded together to identify key topics they believed should be required of a security specialist. The creators of the CBK identified those topics they felt a professional of information security should master and discuss intelligently with peers.

Throughout its years of existence, the CBK has been updated and reorganized, but the main concepts of security have remained unchanged.

CBK – What Are the 10 Domains?

This section will discuss what domains comprise the Common Body of Knowledge and what topics make up each domain.

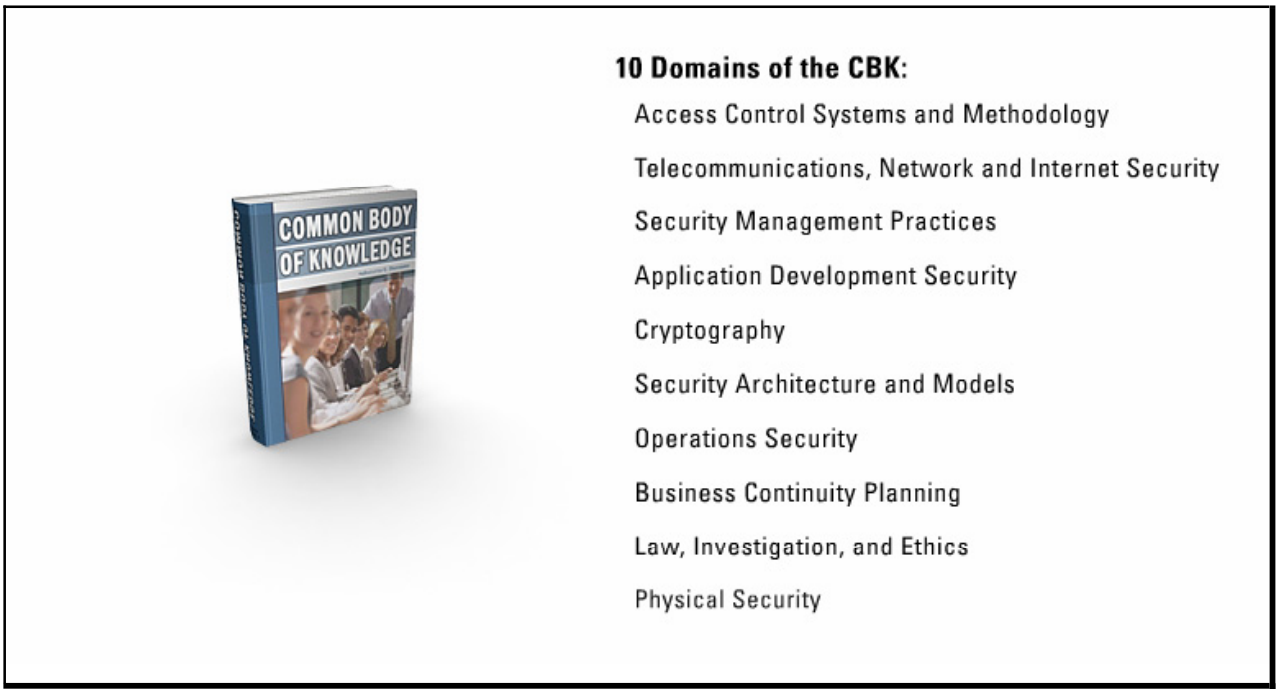

The 10 domains of the CBK include the following:

Access Control Systems and Methodology - This domain examines mechanisms and methods used to enable administrators and managers to control what users can access, their authorization capabilities, and the auditing and monitoring of activities.

- Access control security models
- Identification and authentication technologies and techniques
- Access control administration
- Data ownership
- Attack methods

Telecommunications, Network, and Internet Security - This domain examines external, internal, public, and private communications systems, including network structures, devices, protocols, and remote access and administration.

- OSI model and layers
- Local area network (LAN), metropolitan area network (MAN), and wide area network (WAN) technologies
- Internet, intranet, and extranet issues
- Virtual private networks (VPNs), firewalls, routers, bridges, and repeaters
- Network topologies and cabling

- Attack methods

Security Management Practices - This domain examines the identification of company assets, the proper way to determine the necessary level of protection required, and what type of budget to develop for security implementation with the goal of reducing threats and monetary loss.

- Data classification
- Policies, procedures, standards, and guidelines
- Risk assessment and management
- Personnel security, training, and awareness

Application Development Security - This domain examines the security components with operating systems and applications and how to best develop and measure their effectiveness. This domain also looks at software life cycles, change control, and application security.

- Data warehousing and data mining
- Different development practices and their risks
- System storage and processing components
- System development life cycle
- Malicious code

Cryptography - This domain examines methods and techniques for disguising data for protection purposes. The encryption domain involves cryptography techniques, approaches, and technologies.

Symmetric versus asymmetric algorithms and uses

- Public key infrastructure (PKI), Kerberos, and hashing functions
- Encryption protocols and implementation
- Attack methods

Security Architecture and Models - This domain examines concepts, principles, and standards for designing and implementing secure applications, operating systems, and systems. This domain covers international security measurement standards and their meaning for different types of platforms.

- Operating states, kernel functions, and memory mapping
- Security models, architectures, and evaluations
- Evaluation criteria including Trusted Computer Security Evaluation Criteria (TCSEC), Information Technology Security Evaluation Criteria (ITSEC), and Common Criteria
- Common flaws in applications and systems

Operations Security - This domain examines controls over personnel, hardware, systems, and auditing and monitoring techniques. This domain also addresses possible abuse channels and how to recognize and address them.

- Administrative responsibilities pertaining to personnel and job functions
- Maintenance concepts of antivirus, training, auditing, and resource protection activities
- Directive, prevention, detective, corrective, and recovery controls
- Standards compliance and due care concepts

Business Continuity Planning - This domain examines the preservation of business activities when faced with disruptions or disasters. This knowledge includes the identification of real risks, proper risk assessment, and countermeasure implementation.

- Business resource identification and value assignment
- Business impact analysis and prediction of possible losses
- Unit priorities and crisis management
- Plan development, implementation, and maintenance

Law, Investigation, and Ethics - This domain examines computer crimes, laws, and regulations. This knowledge includes techniques in investigating a crime, evidence gathering, and handling procedures. It also covers how to develop and implement an incident-handling program.

- Laws, regulations, and crimes
- Licensing and software piracy
- Export and import laws and issues
- Evidence types and admissibility into court
- Incident handling

Physical Security - This domain examines threats, risks, and countermeasures to protect facilities, hardware, data, media, and personnel. This information includes facility selection, authorized entry methods, and environmental and safety procedures.

- Restricted areas, authorization methods, and controls
- Motion detectors, sensors, and alarms
- Intrusion detection
- Fire detection, prevention, and suppression
- Fencing, security guards, and security badge types

Summary

The key points discussed in this lesson are:

- What the CBK is
- Why the CBK was created
- The CBK's 10 domains

Security Concerns

Overview

This lesson is designed to give the reader information on the various aspects of threats to the enterprise. Information security professionals must be able to discern what portion of the population poses a threat, what areas on the network are more susceptible to attack, and how the attacks are perpetrated.

Importance

A professional in information security must assess the threats posed to a corporation, identify from where the threats emanate, the type of threat an attack is, and how to mitigate the threat.

Objectives

Upon completing this lesson, you will be able to:

- Define security terminology
- Describe external attacks
- Describe internal attacks
- Describe structured attacks
- Describe unstructured attacks
- Describe modern trends in security

Outline

The lesson contains these topics:

- Security Definitions
- External Attacks
- Internal Attacks
- Structured Attacks
- Unstructured Attacks
- Modern Trends in Security

Security Definitions

This section will discuss the definitions of security related terms a practitioner of security would need to understand.

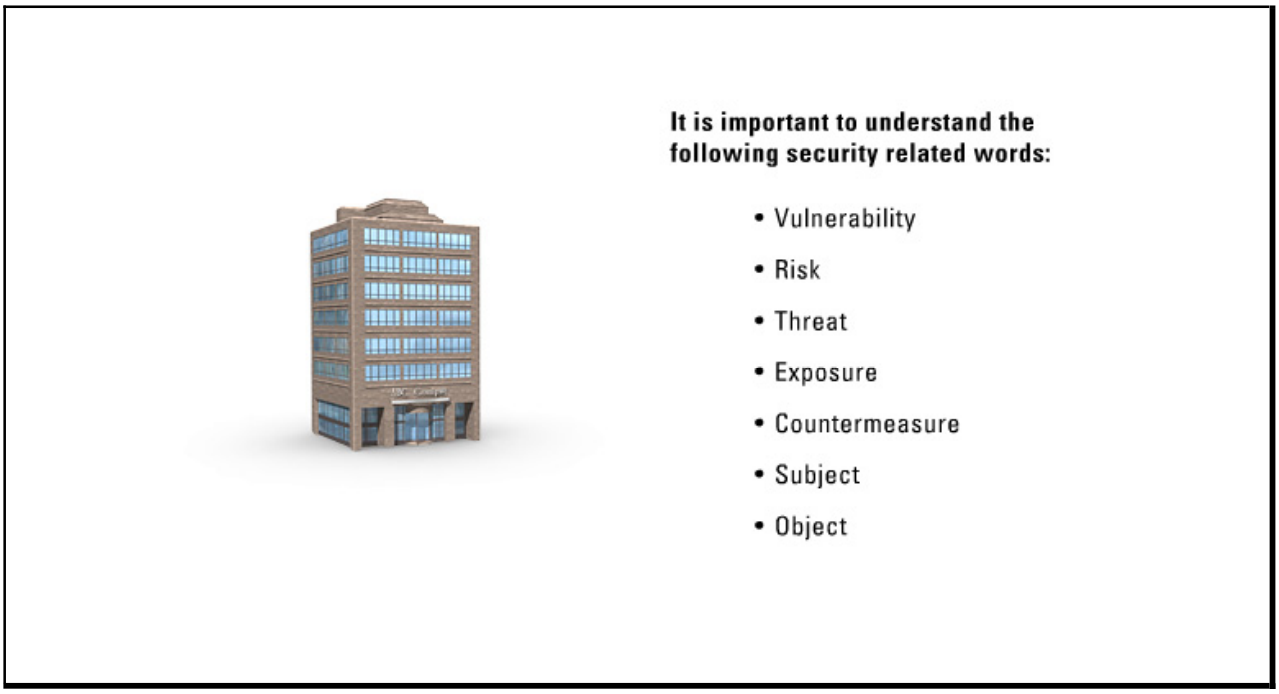

As a security professional, you need to understand the meaning of security related words.

Vulnerability - A point of weakness within a system where threats can occur

Risk - The quantifiable likelihood of a threat taking advantage of vulnerability in a system, or the probability that a threat will exploit a vulnerability

Threat - Something that is a source of danger; capabilities, intentions, and attack methods of adversaries that can exploit or cause harm to a system

Exposure - The potential compromise associated with an attack exploiting a corresponding vulnerability

Countermeasure - Reducing the impact of an attack, detecting the occurrence of an attack, and/or assisting in the recovery from an attack

Subject - Generally a person, process, or device that causes information to flow among objects.

Object - A passive entity containing or receiving information; Access to an object usually implies access to the information that it contains

External Attacks

This section will discuss what external attacks are, who perpetrates them, and how they affect security in the enterprise.

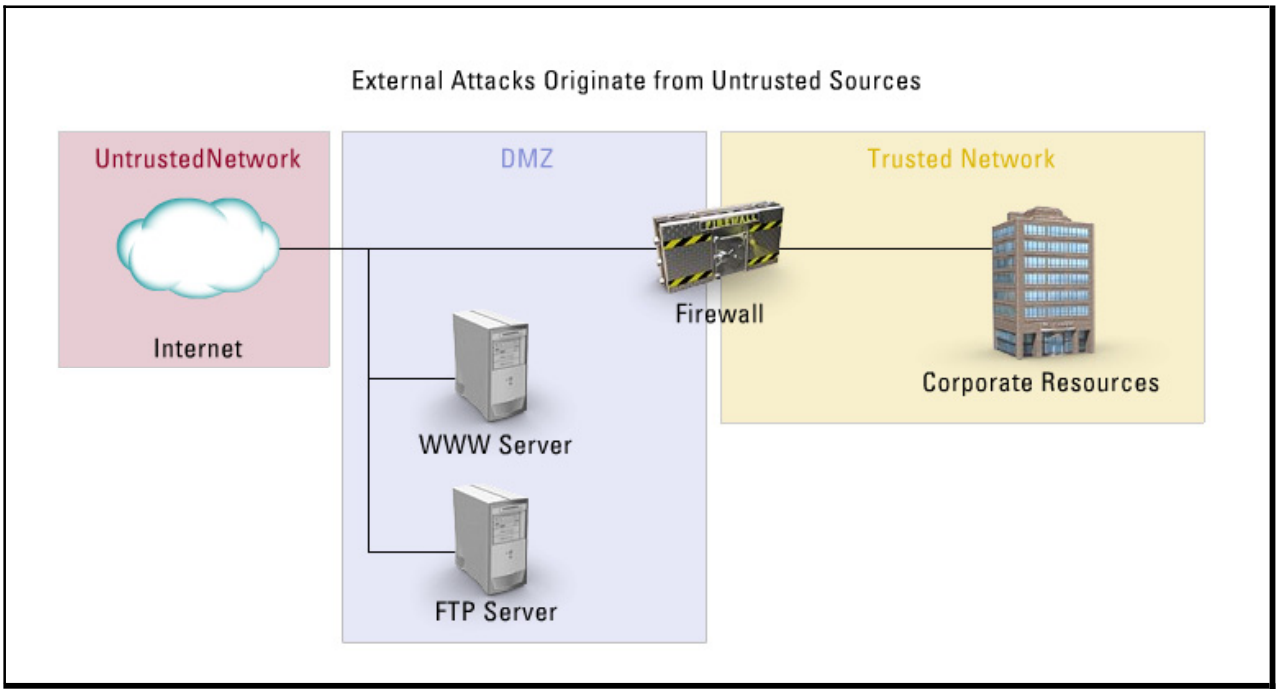

External attacks come from outside the trusted network. Edge devices connecting to untrusted devices create a demilitarized zone (DMZ) that establishes the boundaries of the corporate or trusted network. If an attack occurs from any untrusted point, it can be categorized as an external based attack. You need to realize that many attacks occur that are sourced from the untrusted network. Bots, scripts, and exploited hosts running malicious code are constantly sending malicious packets on the network destined to specific or random IP addresses. But, most if not all of these types of attacks, are easily mitigated with filtering or employing firewalls on the edge devices.

Internal Attacks

This section will discuss what internal attacks are, who perpetrates them, and how they affect security in the enterprise.

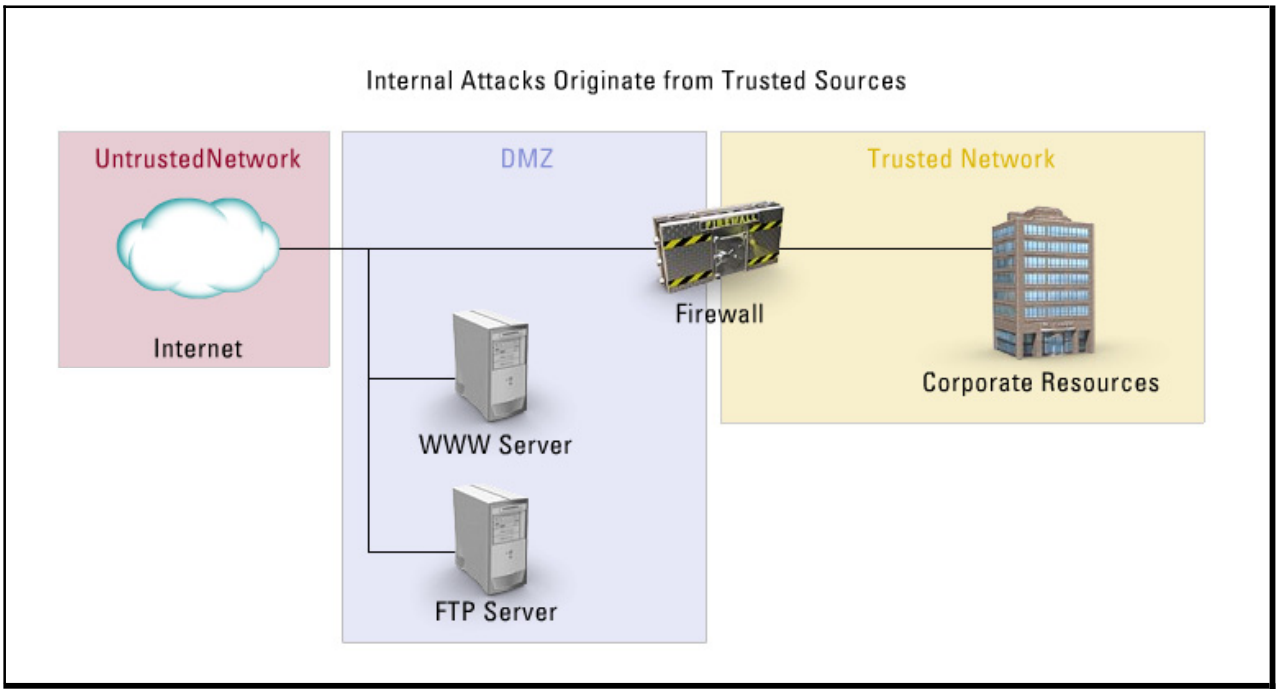

Internal attacks come from within the trusted network. Edge devices connecting to untrusted devices create a demilitarized zone (DMZ) that establishes the boundaries of the corporate or trusted network. If an attack occurs from within any trusted point, security categorizes it as an internal attack.

Note	The security professional needs to realize that more than 70 percent of all successful attacks occur from inside the trusted network.

Because the source is within the network does not mean that employees are always perpetrating these attacks. The internal attack means that the actual base of the attack occurred from a point that is trusted in the network. For example, an attacker may have compromised a trusted system, and from that vantage point, they could then attack other internal systems. Because of the inherent trust of the internal network, the attacker has a better chance at successfully carrying out an attack.

Structured Attacks

This section will discuss what structured attacks are, who perpetrates them, and how they affect security in the enterprise.

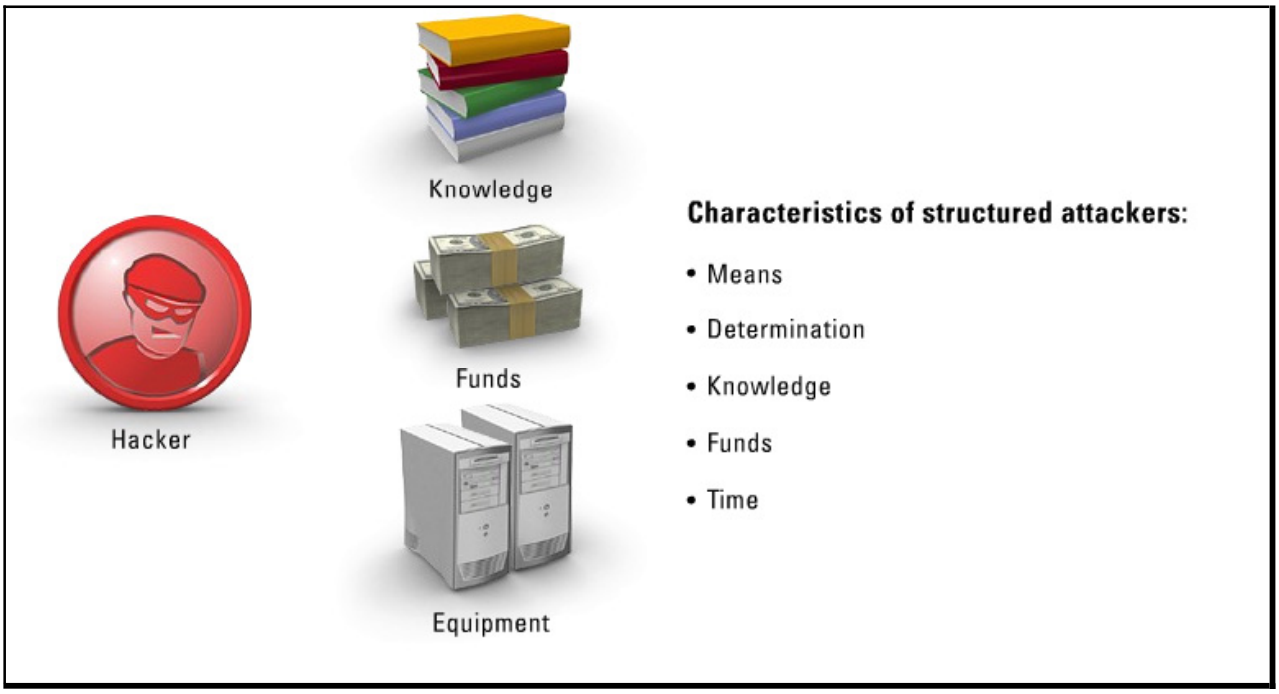

A structured attack is an attack that is carried out by an individual who has discernable characteristics. Those characteristics are means, determination, knowledge, funds, and time. Even if a device is in a secure location, unless it is completely powered off, it can never be 100 percent secure. If someone or some organization has the means usually equipment, the determination, the technical knowledge as well as knowledge of the subject, the funds, and time, they will eventually succeed in the attack.

The job of the information security specialist is to make the attacker's job as difficult as possible and provide countermeasures to identity attacks and attackers before they complete their quest. To accomplish this task, the information security specialist must use a layered approach to security, which is what this course will discuss.

Unstructured Attacks

This section will discuss what unstructured attacks are, who perpetrates them, and how they affect security in the enterprise.

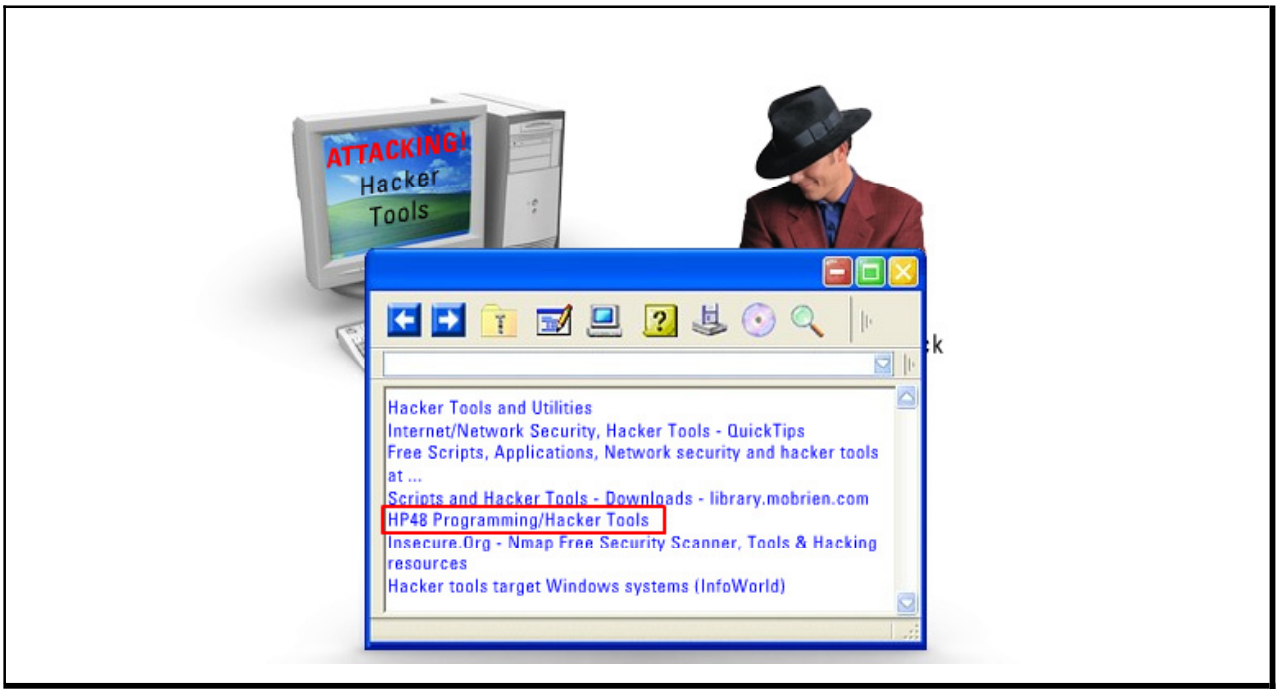

An unstructured attack is one carried out by an individual who does not have sufficient knowledge and experience to carry out an attack unaided. Because the Internet contains such a vast repository of hacker tools and utilities, a simple Google search will provide anyone with pre-built tools that can be executed with a simple click of the mouse.

Modern Trends in Security

Security in the enterprise is a constantly evolving entity that takes its queue from evolving technology. This section will discuss how modern trends in technology change trends in security.

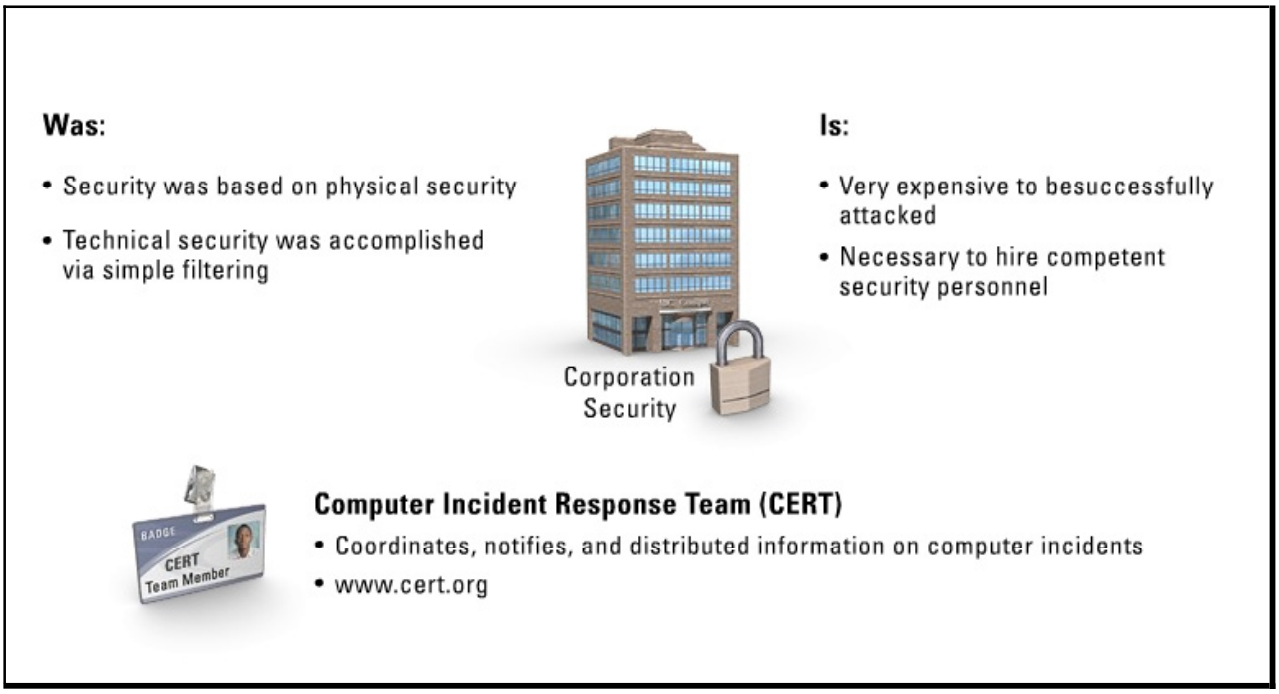

A decade ago, security in the corporate world was heavily dedicated around physical security in and around the headquarters. Security of the technical infrastructure was usually accomplished by means of simple filtering on edge devices. Now, corporations are very aware of the monetary and public relations damage such events incur. Hiring on staff security personnel or outsourcing security is now a matter of fact, not choice, if the company wishes to stay in business.

Because attacks happen so often and new exploits occur daily, many security conscious individuals and companies have formed their own security organizations. These organizations coordinate the notification and distribution of information about incidents of computer crime. One such company is CERT (Computer Incident Response Team) at www.cert.org. This group's prime directives are listed:

- The coordination, notification, and distribution of information of computer incidents to appropriate parties through a predefined escalation path.
- Mitigating risk to the enterprise by minimizing the disruptions to normal business activities and the costs associated with mitigating the incident.
- Assembling teams of technical personnel to investigate the potential vulnerabilities and to resolve specific intrusions.

Additional activities of CERT involve the management of network logs, the management of the resolution of an incident, management of the remediation of vulnerability, and the posting of event reports to appropriate parties.

Summary

The key points discussed in this lesson are:

- Security terminology
- External attacks
- Internal attacks
- Structured attacks
- Unstructured attacks
- Modern Trends in security

The CIA Triad

Overview

An enterprise's lifeblood is the often-proprietary information coursing through its network and stored on its servers. If disreputable or criminal individuals obtain this information, the very life of the enterprise could be at stake. Because so much is riding on the well being of this information, enterprises protect the information, so trusted users can access it when needed and un-trusted individuals cannot. Security goes beyond access. We need to make sure that the data is not altered in transit, when viewed, or when stored on the server.

Importance

In protecting an enterprise's information resources, an information security professional must establish and maintain a strict security defense that ensures three requirements: the information is available when needed, integrity of the information is high, and the information is kept confidential.

Objectives

Upon completing this lesson, you will be able to:

- Identify the fundamentals of security
- Discuss confidentiality as it pertains to security
- Discuss integrity as it pertains to security
- Discuss availability as it pertains to security

Outline

The lesson contains these topics:

- Fundamentals of Security
- Confidentiality
- Integrity
- Availability

Fundamentals of Security

In order to understand security, you must understand why security is necessary and what can be done to secure an object, whether it is a program on a server, a protocol, or the company facility itself.

CIA Triad

Security Component of the CIA Triad
Differs from Company to Company

The main three objectives in any security plan are confidentiality, integrity, and availability. These three security mechanisms are collectively referred to as the CIA triad. The security component of the CIA triad differs from company to company, as each company has different security goals, and reaching those goals must fall within the boundaries of the companies overall business plan. All security components, controls, and safeguards are implemented to mitigate any threat to the principles of the CIA triad.

Confidentiality

Certain data or information in any corporation needs to be secure from prying eyes. This section will discuss what confidentiality is, how confidentiality can protect data, and what the threats to confidentiality are.

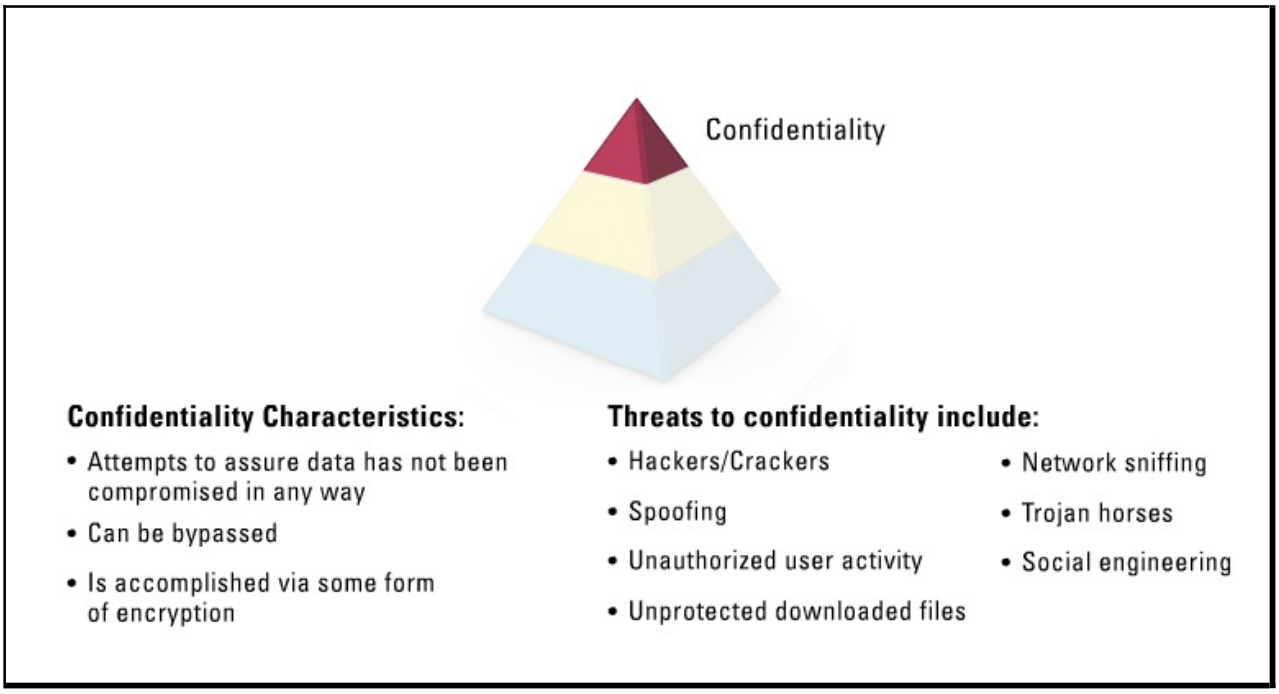

Confidentiality provides a degree of assurance that data has not been compromised, made available or disclosed to unauthorized individuals, processes, or other entities. In essence, it assures that data can only be read or understood between trusted parties. Confidentiality can be breached bypassed by someone shoulder surfing, sniffing or network monitoring, stealing passwords, or social engineering.

Social engineering is easily accomplished because it preys on one of the virtues we hope to instill in our children from a very young age, to help people. If an attacker can pose as a trusted individual asking for assistance or offering to help, the typical response of others will be to trust the person. This tactic can and usually leads to a sharing of confidential information. Mitigating this threat can only be accomplished by means of stringent rules of disclosure and security rule training.

In the electronic world as data crosses the wire, confidentiality is accomplished through some form of encryption.

Threats to confidentiality include the following:

- Hackers/crackers
- Masqueraders/spoofing
- Unauthorized user activity
- Unprotected downloaded files
- Network sniffing

- Trojan horses
- Social engineering

Integrity

Data that has been modified or has the possibility of being modified in the server or in transit is data that cannot be trusted. This section will discuss what integrity is and how it can be used to provide a level of protection for data in the enterprise.

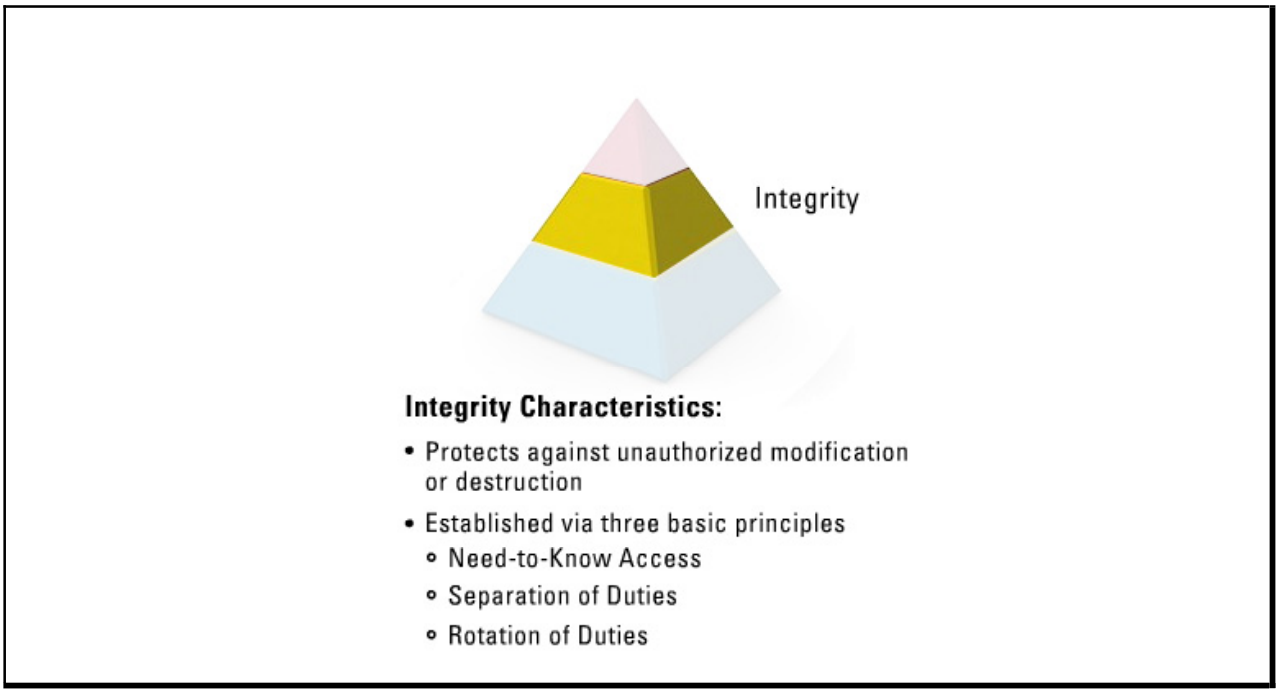

Integrity in the formal security model includes the issue of protecting against unauthorized modification or destruction of information. Good integrity includes the assurance that data leaving point A and arriving at point B arrives without modification. Good integrity assures that point A and point B are indeed who they claim to be.

There are three basic principles that are used to establish integrity in the enterprise:

- **Need-to-Know Access** - Users should be granted access only to those files and programs they absolutely need to fulfill their duties. This status is the least privileged.
- **Separation of Duties** - Use this principle to ensure that no single person has control of a transaction from beginning to end. Make sure that two or more people should be responsible for an entire transaction.
- **Rotation of Duties** - Job responsibilities should be periodically changed so that users will find collaboration more difficult to exercise complete control of a transaction or subvert one for fraudulent purposes.

Availability

Availability to resources is vital to any operational entity. This section will discuss availability as it relates to resources in the enterprise, what is required to achieve availability, and how availability can be defeated.

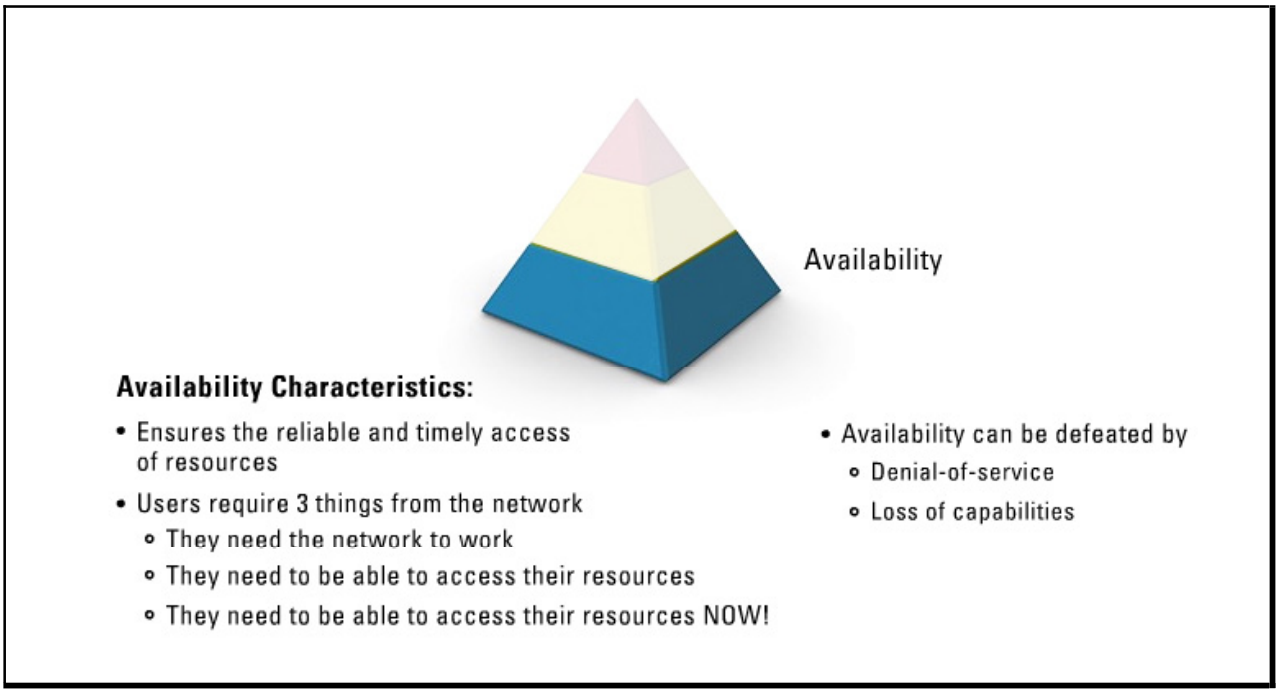

Availability is the attribute that ensures the reliable and timely access of resources to authorized individuals. Network engineers have three principles regarding users' needs ingrained into them from management:

- They need the network to work
- They need to be able to access their resources
- They need to be able to access their resources, now!

Network engineers can provide sufficient bandwidth from users to their resources; they can provide Quality of Service (QoS) guarantees to ensure high priority traffic reaches the destination first, they can provide high-availability and redundancy in the network, but they cannot ensure that the system from which data is being requested is operational and able to reply to requests. Availability addresses all those questions, and as you can see, covers a large portion of the network in general.

There are two facets of availability that are normally discussed:

- **Denial-of-Service (DoS)** - Actions by users or attackers that tie up computing resources in such a way that renders the system unusable or unable to reply to authorized users
- **Loss of Capabilities** - When natural disasters (fire, flood, earthquake) or human action (bombs, strikes, malicious code) create loss of data processing capabilities

Summary

The key points discussed in this lesson are:

- Fundamentals of security
- Confidentiality
- Integrity
- Availability

Access Control Systems and Methodology

Overview

Controlling access to an enterprise's information and data processing resources is critical in any enterprise. An information security professional must be able to describe access control concepts and methodologies used as well as how they are implemented in a centralized or decentralized environment.

Objectives

Upon completing this module, you will be able to:

- Describe the fundamental characteristics of access control
- Explain identification and authentication
- Explain the function of passwords
- Define access control techniques
- Define access control administration
- Explain monitoring and intrusion detection

Outline

The module contains these lessons:

- Access Control Overview
- Identification and Authentication
- Passwords
- Access Control Techniques
- Access Control Administration
- Monitoring and Intrusion Detection

Access Control Overview

Overview

The enterprise implements access controls to ensure the confidentiality, integrity, and availability of corporate resources. This lesson will discuss access rights and permissions and how to maintain access control. You will also learn how to revoke access control for an entity when needed.

Importance

It is important that an information security professional be able to identify the access control security tools and technologies used in an enterprise in order to avoid risks, exposures, and vulnerabilities. It is also important to be able to describe any auditing mechanism that is used to analyze behavior, use, and content of the enterprise IT system.

Objectives

Upon completing this lesson, you will be able to:

- Define access rights and permissions
- Define establishment parameters
- Define maintenance access control
- Define revocation steps
- Define accountability

Outline

The lesson contains these topics:

- Access Rights and Permissions
- Establishment
- Maintenance
- Revocation
- Accountability

Access Rights and Permissions

To provide a level of security in the network, objects such as printers, files and other resources need a protection mechanism, so they cannot be viewed or modified in an unauthorized manner. Providing access rights and permissions to subjects and objects in the network can provide a required amount of access control in the network.

Access control is the heart of access rights and permissions. Access controls are the mechanisms that specify what users on a system can or cannot do. Access control includes what actions a user can perform and what resources a user may have access to. Access controls are the countermeasures for ensuring that only those users with the proper need and proper authority can access the resource.

Establishment

For access control to be utilized, the subject must first be properly identified or established.

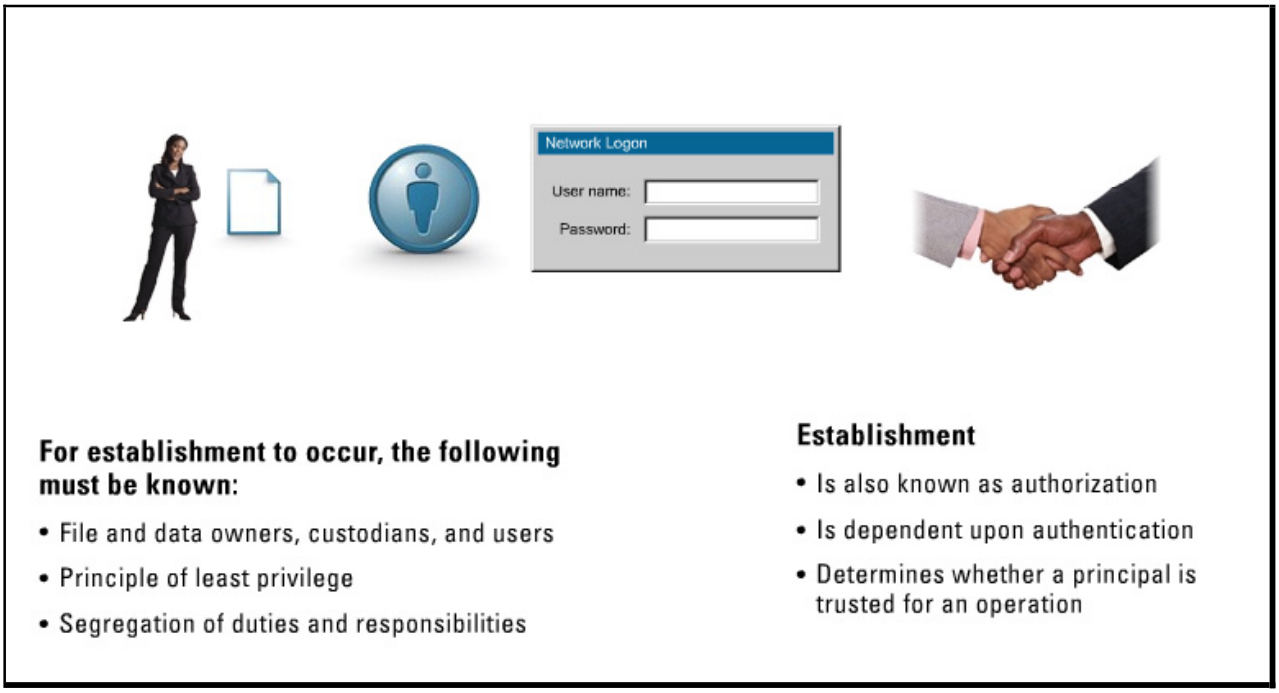

Establishment, also known as authorization, determines whether a particular principal, who has been authenticated as the source of a request to do something, is trusted for that operation. Authorization may also include controls on the time at which this action can take place or which IP address may request it.

In order for establishment to occur the following must be used and known:

- **File and Data Owners, Custodians, andUsers** - All information generated or used must have a designated owner. The owner determines the appropriate classification and access controls. The owner is also responsible for ensuring appropriate controls for the storage, handling, and distribution of the data. Custodians receive permission from the owners and manage the everyday care of the data such as backups. Users are the subject that requires data (the object) to perform their jobs.
- **The Principle of Least Privilege** - Requires that a user be given no more privilege than necessary to perform a job. Ensuring the least privilege requires identifying what the user's job is, determining the minimum set of privileges required to perform that job, and restricting the user to a domain with only necessary privileges.
- **Segregation of Duties and Responsibilities** - Requires that for particular sets of transactions, no single individual be allowed to execute transactions within the set. The transactions can either be static or dynamic.

Maintenance

This section will discuss the mechanisms used in order to maintain access control in the enterprise.

Maintaining access control in the enterprise requires several components for each category of access control. There are three main categories of access control:

Administrative

- **Policies and procedures** - A high-level plan that lays out management's plan on how security should be practiced in the company. It defines what actions are not acceptable and what level of risk the company is willing to accept.
- **Personnel controls** - Indicate how employees are expected to interact with corporate security, and how non-compliance will be enforced.
- **Supervisor structure** - Defines the overall company hierarchy. Each employee has a supervisor they report to and that supervisor has a superior they report to. This chain of command dictates who is responsible for each employee's actions.
- **Security awareness training** - Users are usually the weakest chain in the security chain. Proper training on security issues can instill access control usage on the network.
- **Testing** - Test access controls on the network to determine their effectiveness (or ineffectiveness).

Physical

- **Network segregation** - Defining segregation points can help enforce access controls on ingress or egress to the segment.
- **Perimeter security** - Defines how the perimeter of the company will be enforced such as guards, security badges, fences, gates.

- **Computer controls** - Defines the physical controls on computer systems such as locks on systems to deter theft of internal parts, removal of floppy to deter copying.
- **Work area separation** - Separation of work areas based on type of use such as server room, wiring closets, experimental room.
- **Data backups** - This physical control is used to ensure access to information in case of system failure or natural disaster.
- **Cabling** - Protecting the cabling from electrical interference, crimping, and sniffing.

Technical

- **System access** - Controls that determine how resources on a system are accessed such as MAC architecture, DAC architecture, username/password, RADIUS, TACACS+, Kerberos.
- **Network architecture** - Defines logical network segmentation to control how different network segments communicate.
- **Network access** - Defines access controls on routers, switches, and network interface cards, and bridges. Access control lists, filters, AAA, and firewalls would be used here.
- **Encryption and protocols** - A technical control that encrypts traffic as it courses through untrusted network segments. Protocols could include IPSec, L2TP, PPTP, SSH, SSL/TLS.
- **Control zone** - A specific area in the enterprise that surrounds and protects network devices that emit electrical signals. Electrical signals emanate from all computer systems and travel a certain distance before being drowned out by interference from other electrical fields. Control zones are both a technical and physical control.
- **Auditing** - Tracks activity as resources are being used in the enterprise.

Revocation

Access control for an entity sometimes needs to be revoked. This section will discuss when revocation is necessary and how it should be performed.

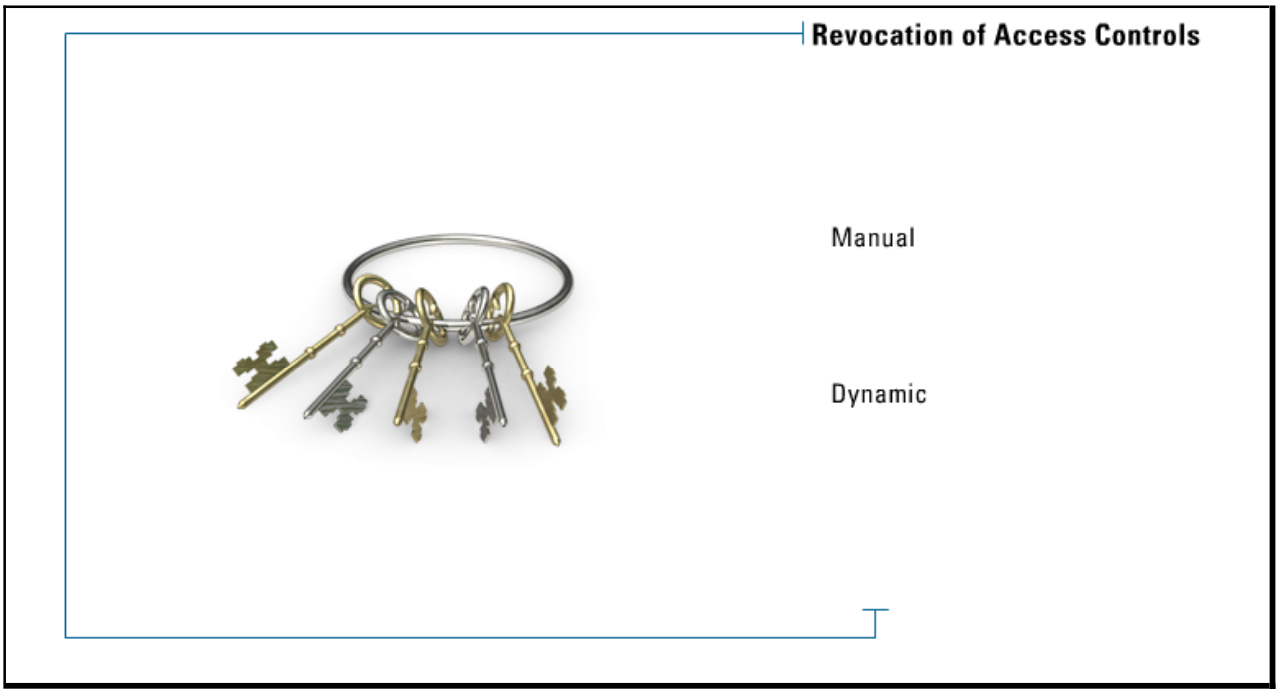

Revocation of access controls can be pre-determined as in the case of employee termination or dynamic for intrusion detection. When an employee is terminated, many things must occur such as disabling the employee account, changing passwords on equipment, and removing access to system accounts. Revocation can also be configured on a dynamic basis as in the case of intrusion detection, where traffic has been deemed hostile, and the source of the traffic is filtered from entering the network.

Accountability

Having the ability to account for actions taken by an entity provides a level of trust in the network. This lesson will discuss what accountability is and how it can be performed in the network.

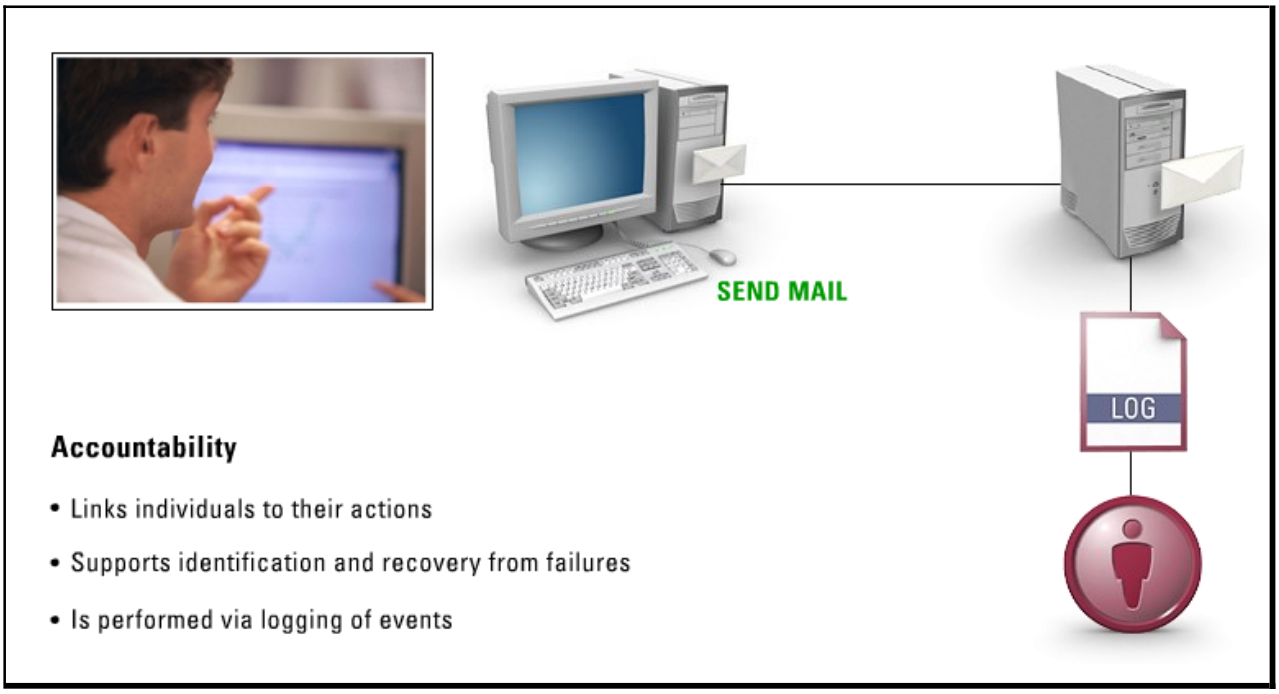

Accountability insures that of the system links individuals to their interactions within an IT product, thus supporting identification of and recovery from unexpected or unavoidable failures of the control objectives. Accountability is performed by logging of events on a system basis such as UNIX's syslog service, or on a network basis, such as traffic logging or SNMP traps.

Summary

The key points discussed in this lesson are:

- Access rights and permissions
- Establishment
- Maintenance access control
- Revocation steps
- Accountability

Identification and Authentication

Overview

System security requires measures to ensure that only users with proper need and authority can access the enterprise system resources. The most common way of achieving these measures in today's world is validation of credentials. Identification and authentication are the tools that implement this security event.

Importance

Information security professionals need to identify points of interest in the network. These points of interest are where confidential company resources reside. To gain access to these resources, the users need to use proper identification and authentication. The security professional should identify the mechanisms and characteristics of various identification methods.

Objectives

Upon completing this lesson, you will be able to:

- Define identification and authentication techniques
- Describe performance measures
- Define knowledge based systems
- Describe passwords
- Describe PINs
- Describe pass phrases
- Identify characteristics-based access control
- Describe biometrics
- Explain behavior systems
- Describe a single sign-on (SSO)
- Describe tickets
- Define kerberos

- Explain SESAME sign-on devices
- Describe thin clients
- Define scripts
- Define one-time passwords
- Describe token-based systems
- Explain synchronous token devices
- Explain the asynchronous OTP method
- Describe smart card devices
- Describe key cards

Outline

The lesson contains these topics:

- Identification and Authentication Techniques
- Performance Measures
- Knowledge-Based Systems
- Passwords
- PINs
- Pass Phrases
- Characteristics-Based Access Control
- Biometrics
- Behavior
- Single Sign-On (SSO)
- Tickets
- Kerberos
- SESAME
- Thin Clients
- Scripts
- One-Time Passwords
- Token-Based Method
- Synchronous
- Asynchronous
- Smart Card
- Key Card

Identification and Authentication Techniques

This section will discuss techniques used to provide for identification and authentication of subjects as they attempt to access objects in the enterprise.

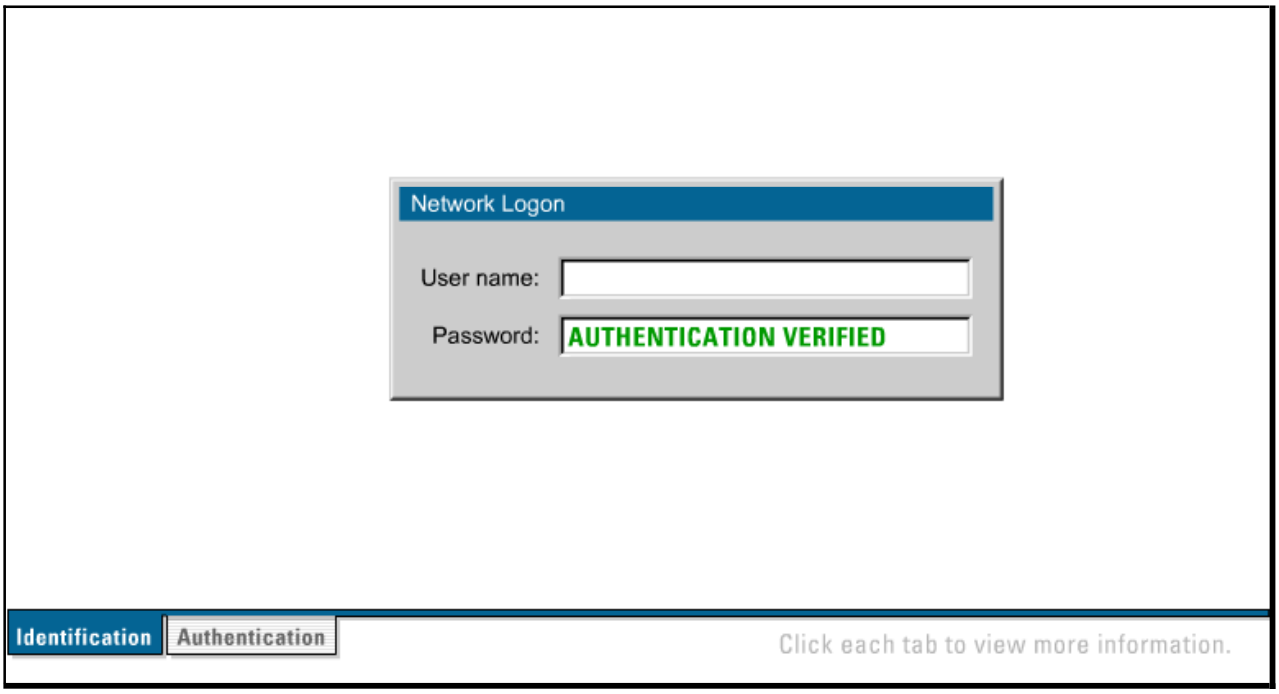

Identification - The act of a user professing an identity to a system usually in the form of a logon.

Authentication - The verification that the user's claimed-identity is valid. The authentication is usually implemented through a password at logon time.

Authentication is based on the following three factor types:

- **Something you know** - Password, PIN, mother's maiden name, passcode, etc.
- **Something you have** - ATM card, smart card, token, key, ID Badge, driver license, or passport.
- **Something you are** - Also known as biometrics: Fingerprint, voice scan, iris scan, retina scan, body odor, DNA.

Multi-factor authentication uses two or more authentication types to provide increased security.

Two-factor authentication systems require a user to provide two of the three types of authentication.

- ATM card + PIN
- Credit card + signature
- PIN + fingerprint
- Username + Password (NetWare, Unix, NT default)

Three-factor authentication offers the highest security.

- password + token + Fingerprint
- PIN + driver license + voice scan

Performance Measures

Fast performing authentication mechanisms usually are weak in comparison to slower performing authentication mechanisms. This section will discuss the performance measures used in various forms of authentication.

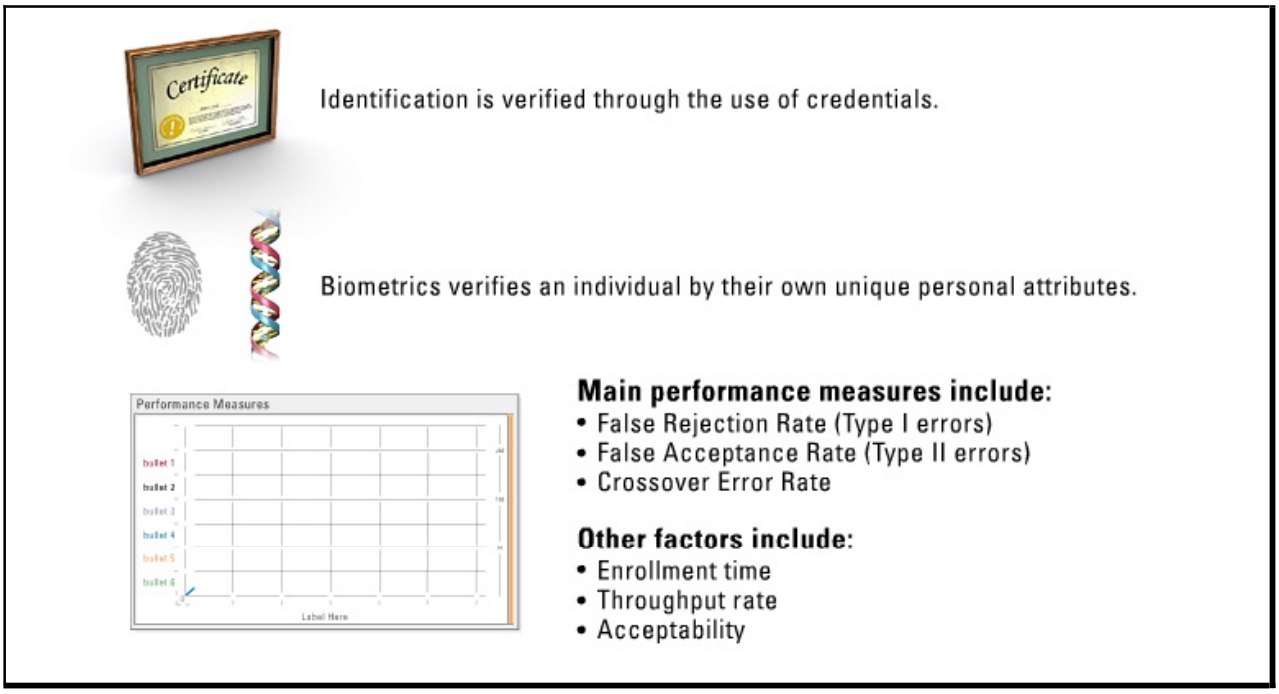

Identification describes a method of ensuring that a subject, user, program, or process is the entity it claims to be. Identification can be verified through the use of a credential. Biometrics verifies an individual's identity by a unique personal attribute, which is one of the most effective and accurate methods of verifying identification.

Three main performance measures:

- **FRR/False Rejection Rate or Type I Error** - The percentage of valid subjects that are falsely rejected.
- **FAR/False Acceptance Rate or Type II Error** - The percentage of invalid subjects that are falsely accepted.
- **CER/Crossover Error Rate** - The percent in which the False Rejection Rate equals the False Acceptance Rate.

Other factors that must be considered:

- **Enrollment time** - The time it takes to initially register with a system by providing samples of the biometric characteristic to be evaluated.
- **Throughput rate** - The rate at which individuals can be processed and identified or authenticated by a system.
- **Acceptability** - Considerations of privacy, invasiveness, and psychological and physical comfort when using the system.

Knowledge-Based Systems

This section will discuss what knowledge based authentication systems are and how they can be used in the enterprise.

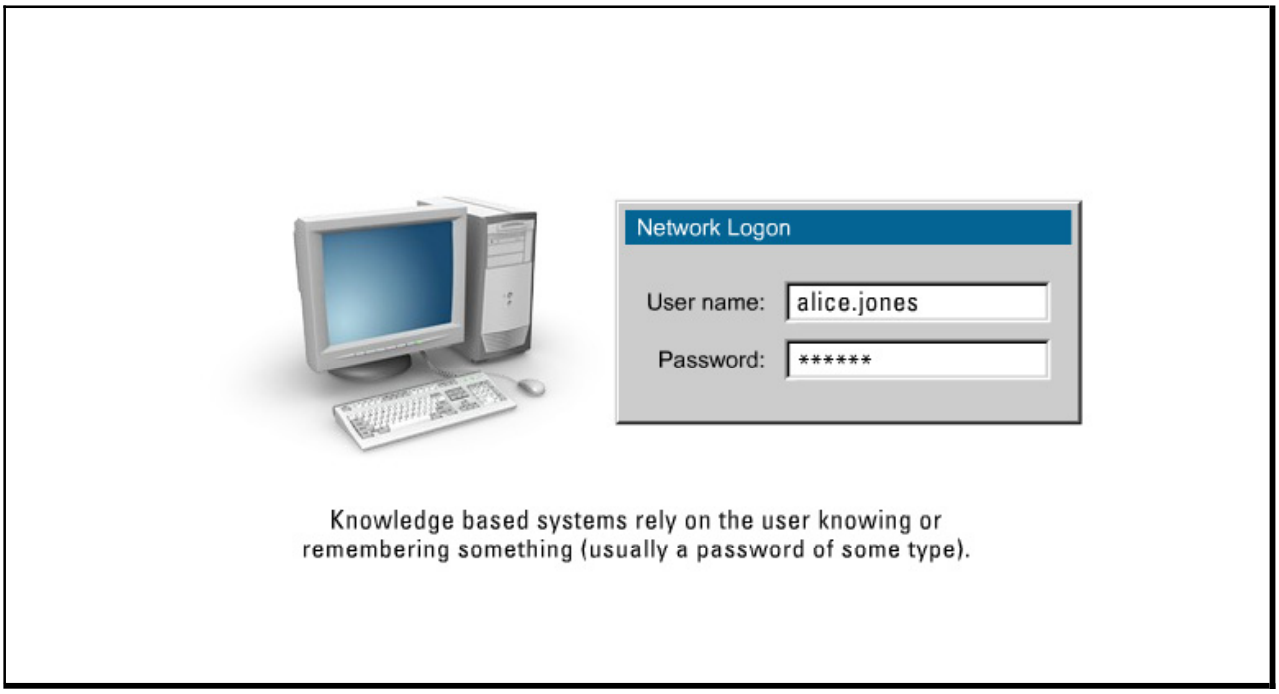

Knowledge based systems rely on the user knowing or remembering something. Password knowledge based systems use several schemes:

- **User Selected** - The user has free reign on the type and length of the password
- **Generated** - The system explicitly generates a password for the user
- **Token generated** - The system works in conjunction with a token card to provide a password for the user
- **Default** - The system provides a default password
- **Composition** - Combination of two, totally unrelated words
- **Pass phrases** - Good way of having very strong passwords

Password Management includes the following issues:

- Lifetime Considerations
- Cost of replacement
- Risk of compromise
- Guessing attacks
- Number of times used
- Password Changing Considerations
- Sixty days regular user
- Thirty days privilege users

- Fifteen days security officer

Use Security Policies to control password management issues.

Passwords

This section will discuss what passwords are, how they are used in the enterprise, and what can be done to strengthen a password scheme.

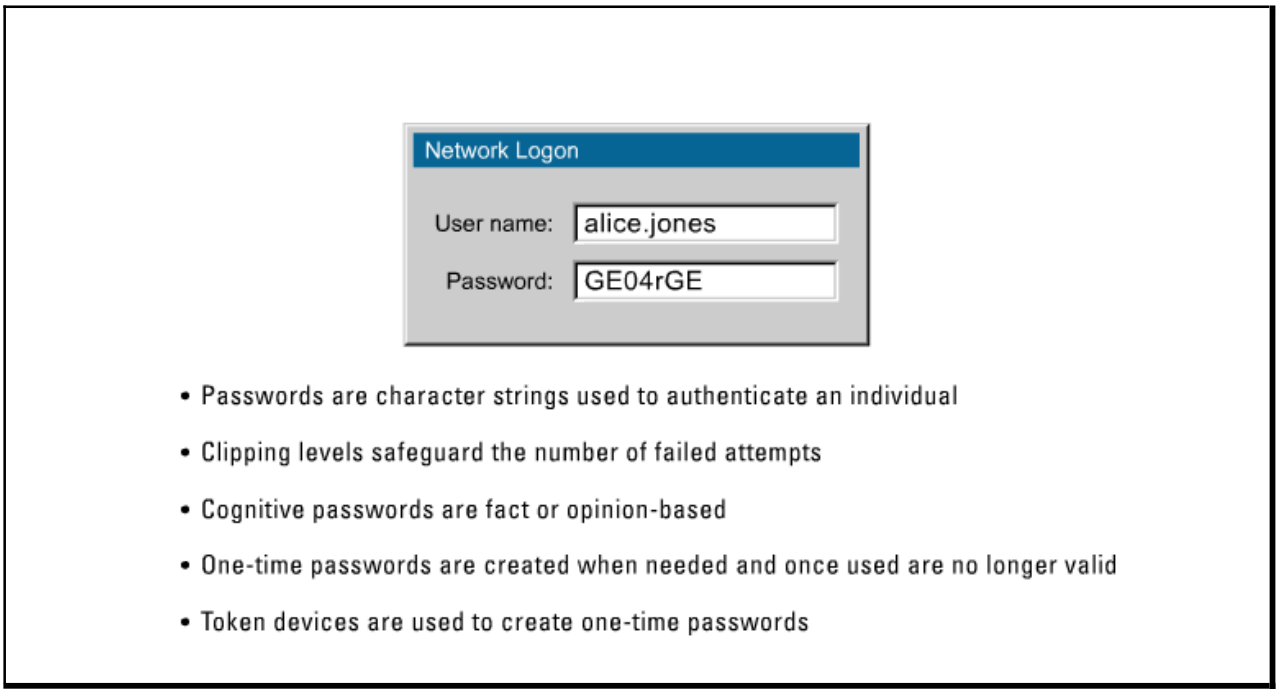

A password is a protected string of characters used to authenticate an individual. To limit the number of failed login attempts, system administrators create a clipping level, which when reached, locks the account. When using a password system for authentication, you may want to use the following for security-based reasons:

- **Password checkers** - Tests the security of user-chosen passwords.
- **Password Generators** - Generators that produce users' passwords.
- **Password Aging** - Expiration dates for passwords.
- **Limit Login Attempts** - Threshold set to allow only a certain number of unsuccessful login attempts.

A cognitive password is a fact or opinion based password that is used to verify an individual's identity. Many companies ask for your mother's maiden name or city you were born as a means to verify your identity.

A one-time password (also called a dynamic password) is one that is created when needed and once used is no longer valid.

Token devices can also be used to generate one-time passwords. Token devices use a challenge/response scheme in which the user enters the challenge key in the device with another mechanism that creates a password valid for one time or for a certain amount of time.

Administrators use a few methods in a token device scheme:

- **Synchronous token device** - Synchronizes with the authentication service by using time or an event as the core piece of the authentication process.

- **Time based synchronous token device** - The device and the authentication service must hold the exact same time within their internal clocks.
- **Event-synchronization** - The user may need to initiate the logon sequence on the computer and push a button on the token device.
- **Asynchronous token device** - Uses challenge-response scheme to communicate with the authentication service.

PINs

This section will discuss what Personal Identification Numbers (PINs) are, and why they may be important in an enterprise.

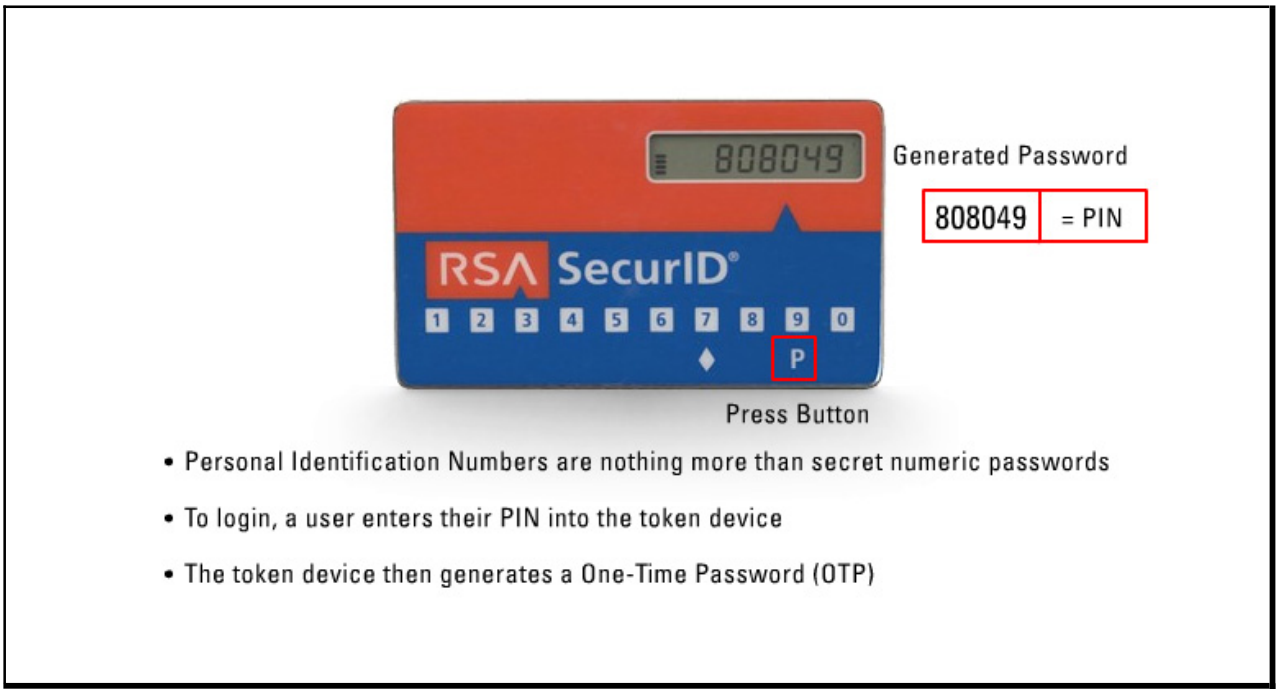

When using multi-factor authentication, the user provides two separate parameters for authentication: usually something he or she knows such as a password, User ID, or PIN, and something he or she has such as a token. Personal Identification Numbers (PINs) are in this sense nothing more than a secret numeric password.

The token might be software embedded in the user's PC or a separate handheld processor, a small device the size of a credit card. The token has the user's unique and secret key within it.

In order to successfully login, the user enters their PIN into the token. The token then generates a one-time password for the system.

Pass Phrases

This section will discuss the pros and cons of implementing a pass phrase system for authentication in an enterprise.

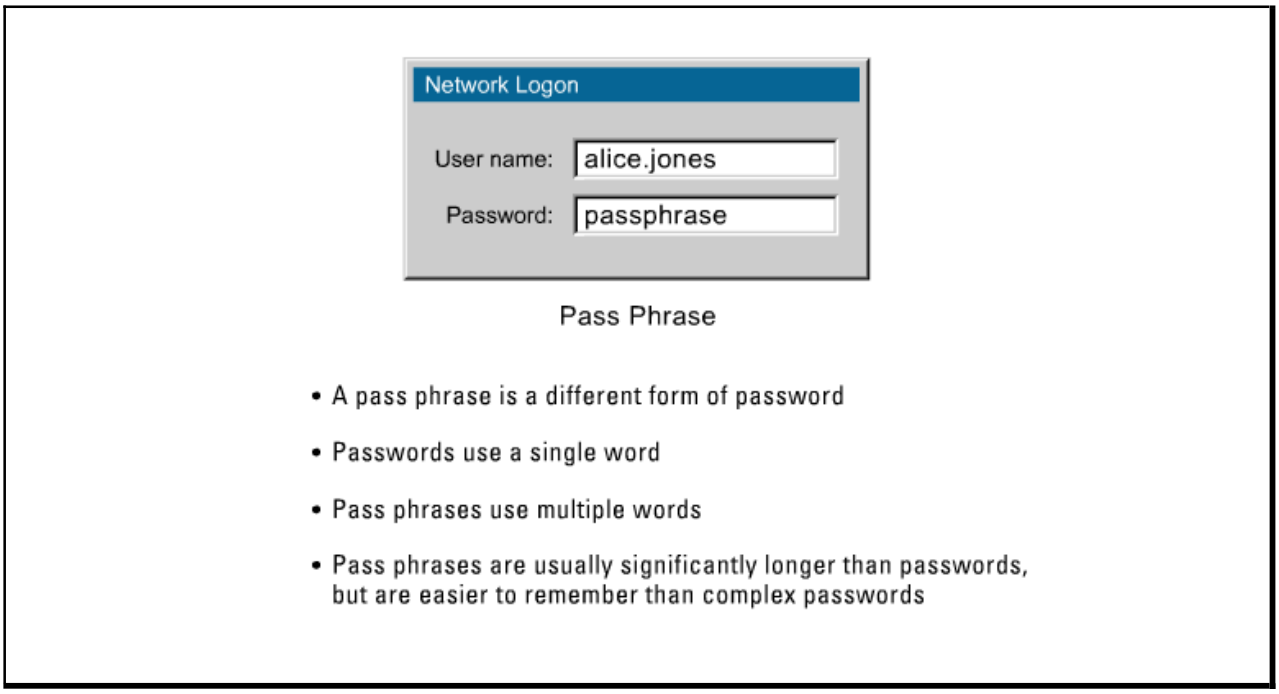

You might have heard the term pass phrase tossed about in your conversations or in a security seminar. Distinguishing between a pass phrase and a password is simply a measure of added complexity. Traditional passwords are a user-chosen word. Increasingly complex passwords include the use of different characters and cases such as mIP@$$W0rd, thus making passwords more difficult for an attacker to figure out and also more difficult for the user to remember.

A pass phrase on the other hand consists of more than one word, which may or may not form a complete sentence. This difference means a pass phrase will contain spaces and be significantly longer than a password. For example a pass phase might be “Let me EY into the system” or “Some say love is blind”. Both would be very difficult for an attacker to guess or perform brute force attacks against.

There are advantages and disadvantages to using either passwords or pass phrases, and in the end, it usually comes down to either personal choice or company mandate.

Characteristics-Based Access Control

Authentication performed today is usually performed with a password that can be forgotten. Characteristics-based access control performs authentication in a completely different way. This section will discuss those different ways and how they provide greater security to access control in the enterprise.

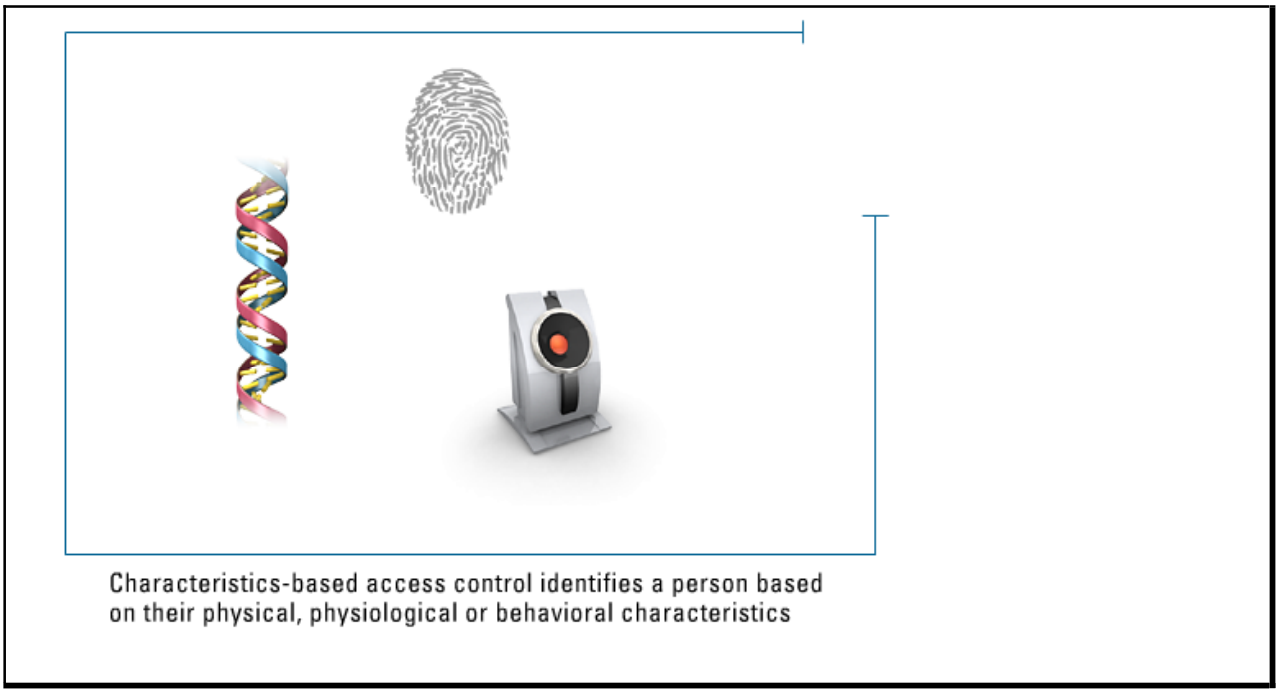

Characteristics-based access control identifies a person based on their physical, physiological or behavioral characteristics

Characteristics-based access control is an automated means of identifying or authenticating the identity of a living person based on their physical, physiological, or behavioral characteristics. Another term we use for characteristics-based access control is called biometrics.

Advantages of characteristic-based biometrics:

- They cannot be lent like a physical key or token, and they cannot be forgotten like a password.
- Good compromise between ease of use, template size, cost, and accuracy.
- They contain enough inherent variability to enable unique identification even in very large databases.
- They last forever or until the owner has an amputation or dismemberment.
- They make network login and authentication effortless.

Disadvantages:

- They are still relatively expensive per user.
- Companies and products that supply the biometrics are often new and immature.
- They have no common API or other standard.
- Users have some hesitancy accepting the concept.

Privacy issues:

- **Tracking and surveillance** - Ultimately, the system has the ability to track a person's movement from hour to hour.
- **Anonymity** - Biometrics linked to databases could dissolve much of the user's anonymity when they travel and access services.
- **Profiling** - Compilation of transaction data about a particular person creates a picture of that person's travels, preferences, affiliations, or beliefs.

Biometrics

This section will discuss the various forms of biometric authentication mechanisms in common use today.

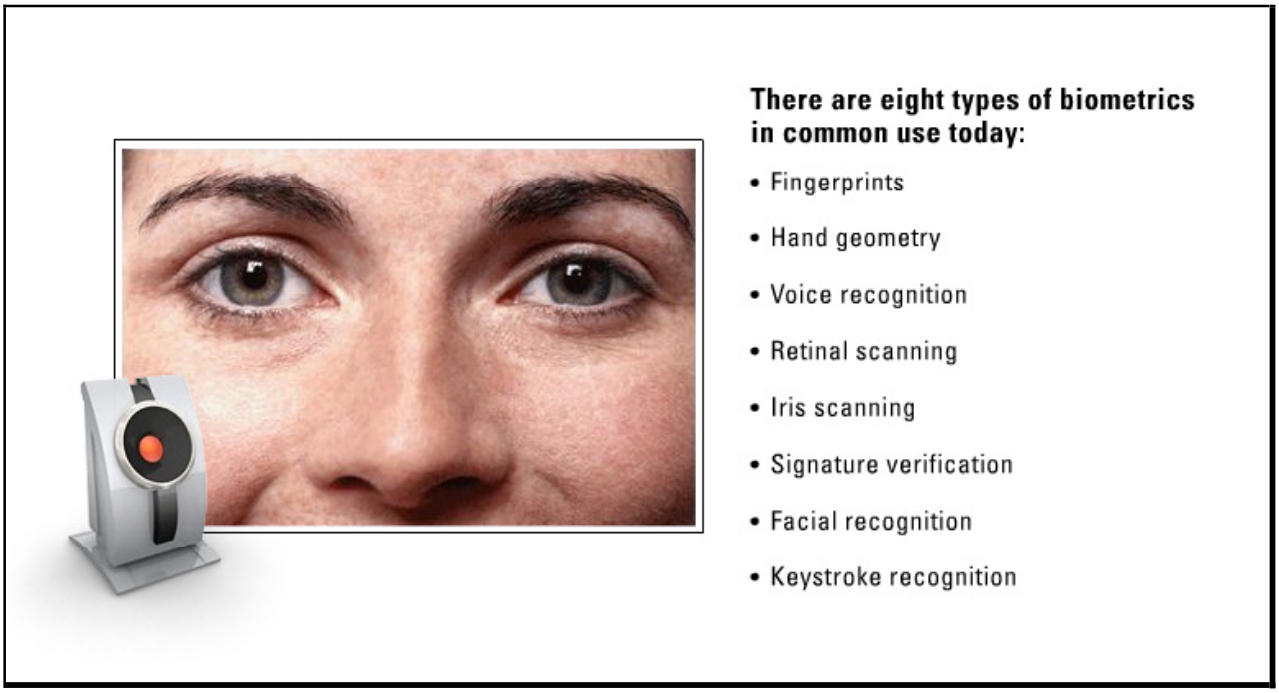

There are eight types of biometric measurements in common use today:

Fingerprint Verification is perhaps the best-known type of biometric measurement. Fingerprint scanning products are the most common type on the market today. Properly implemented, fingerprints offer potential for high accuracy. In addition, the readers tend to be small and easy to incorporated into a keyboard; they have a relatively low cost, and integration is usually easy. Some potential problems can arise, however. Cuts or dirt on the finger can cause some systems not to recognize a valid fingerprint. Some scanners require precise placement of the finger. Finally, some systems do not detect a real finger from some sort of copy. To overcome this shortcoming, some fingerprint scanners will scan for pulse as well as the fingerprint.

Hand Geometry measures the physical characteristics of the user's hand and fingers. Hand geometry is one of the most established methods and typically offers a good balance of performance and ease of use. Hand geometry is most widely used in physical access control and time/attendance systems. It is not currently in wide deployment for computer security applications primarily because it requires a large scanner.

Voice Recognition is perhaps the method most desirable to users since everyone seems to want to talk to computers. In practice, implementation is extremely difficult. While recent advances in voice recognition have greatly improved the technology, it remains subject to problems. Local acoustics, background noise, microphone quality, the common cold, anxiety, being in a hurry, and anger can alter the human voice enough to make voice recognition difficult or impossible. Further, voice recognition systems tend to have the most difficult and time-consuming enrollment process and require the most space for template storage.

Retinal Scanning is well established and can provide high accuracy. User acceptance may be a problem, however. Responses to using the scanner have been, "You're not shooting a laser into my eye!" In reality, retinal scanners do not employ a laser, but scan using low intensity light and are considered quite safe. One drawback is that the user must look directly into the retinal reader. Eyeglass wearers may find the device inconvenient. In public applications, there may also be concerns with the spread of germs because of the need for physical contact with the retinal scanner. Another issue can be the need to focus on a given point for the scan. Failure to focus correctly causes a significant impact on accuracy.

Iris Scanning overcomes most of the problems of retinal scanners. Because the iris, the colored part of the eye, is visible from a distance, direct contact with the scanner is not required, nor does the user have to remove his or her eyeglasses. The technology works by scanning the unique random patterns of the iris. Interestingly, the method does not rely on the iris color since the camera used is black-and-white. This feature is important because of the popularity of colored contact lenses; some vendors claim their systems will work with colored contacts and even through non-reflective sunglasses.

Signature Verification enjoys a synergy the other technologies because people are used to signing for things. The technology provides a greater feeling of normalcy. While signature verification has proved to be relatively accurate, very few products available implement the technology.

Facial Recognition is one of the newest biometric methods. The technology has attracted much attention. Unfortunately, extravagant claims proved difficult to substantiate cooling much of the enthusiasm. While matching two static images is not difficult, picking an individual out of a group, as some systems claim to be able to do, is more difficult. Progress continues to be made with this young technology, but to date, facial recognition systems have had limited success in practical applications.

Keystroke Recognition uses the fact that the way and the manner in which we type on our computer keyboard varies from individual to individual, and is in fact, unique enough to be considered to be a behavioral biometric. Keystroke recognition is completely a software-based solution. No new hardware needs to be installed; all that is needed is the existing computer and keyboard that the individual is currently using.

Order of effectiveness from most to least secure:

- Iris Scanning
- Retinal Scanning
- Hand Geometry
- Fingerprint Verification
- Voice Recognition
- Facial Recognition
- Signature Verification
- Keystroke Recognition

Behavior

This section will discuss behavioral based biometrics and how they can fail over time.

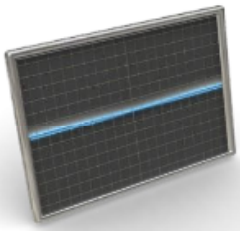

Behavior-based biometrics are less expensive:

- Voice recognition
- Signature verification
- Keystroke recognition

Physiological biometrics offer greater accuracy:

- Fingerprints
- Hand geometry
- Retinal scanning
- Iris scanning
- Facial recognition

Biometrics can change over time

When discussing one's behavioral characteristic, such as one's signature, voice, or keystroke dynamics, we need to remember that these characteristics are influenced by both controllable actions and less controllable psychological factors, meaning they can change over time.

Although behavior-based biometrics can be less expensive and less threatening to users, physiological traits tend to offer greater accuracy and security. In any case, both techniques provide a significantly higher level of identification than passwords or smart cards alone.

Single Sign-On (SSO)

In today's departmentalized enterprise scheme, its possible that a user needs to have multiple authentication credentials to access the various servers and resources in the enterprise. This section will discuss how single sign-on can be used to alleviate the problem with the departmentalized approach.

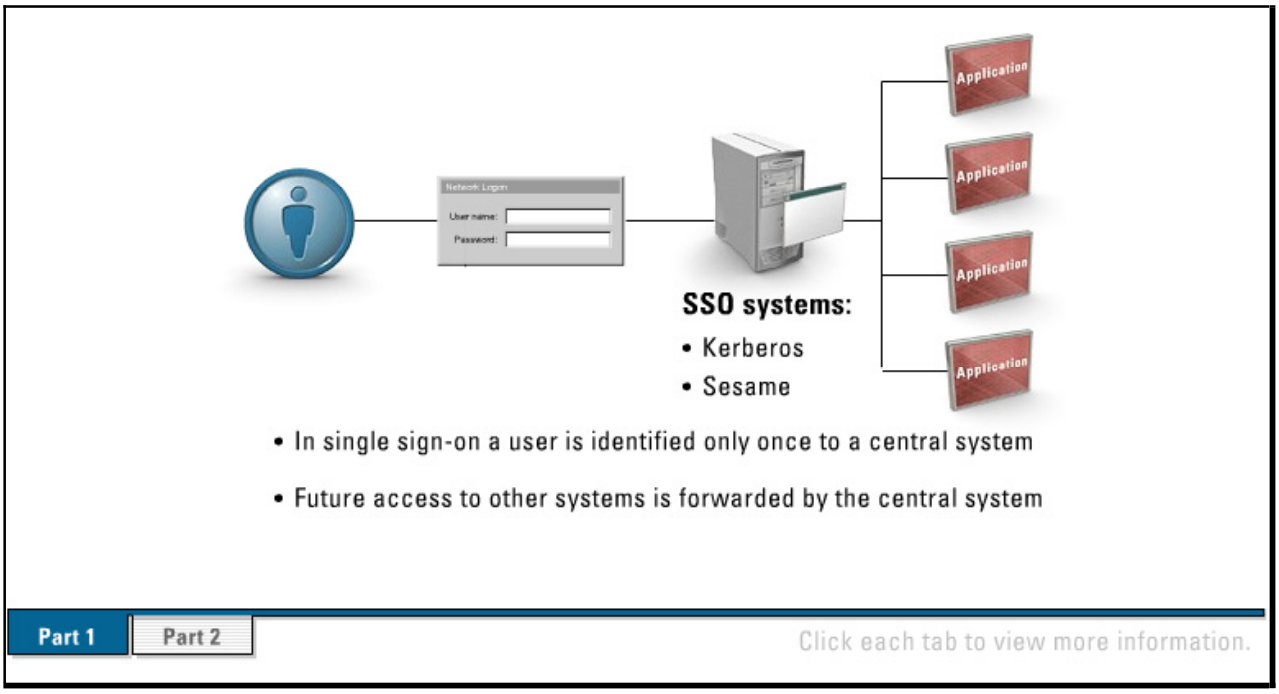

Single sign-on addresses the cumbersome situation of logging on multiple times to access different resources located on different systems. In the single sign-on scheme, users identify only once to a central system, then information needed for future system access to other resources is forwarded by the central system.

Pros

- More efficient user log-on process
- The ability to use stronger passwords
- User has a single password for entire enterprise resources
- A single strong password can be remembered and used
- A single user account is created and can be quickly created or removed

Cons

- Once user has logged on, they can freely roam the network
- Hard to implement and get working

Examples of single sign-on systems:

- **Kerberos (MIT project Athena)** - A trusted, third party authentication protocol that was developed at MIT. Using symmetric key cryptography, it authenticates clients to other entities on a network of which a client requires services.

- **SESAME (Secure European System for Applications in a Multivendor Environment)** - Addresses the weaknesses in Kerberos. Uses public key cryptography for the distribution of the secret keys and provides additional access control support. It uses the Needham-Schroeder protocol.

Tickets

Single sign-on authentication schemes require a mechanism to be used between the user, the central authentication server, and the various resources the user needs access to. This section will discuss the mechanism used to communicate authentication parameters between these entities.

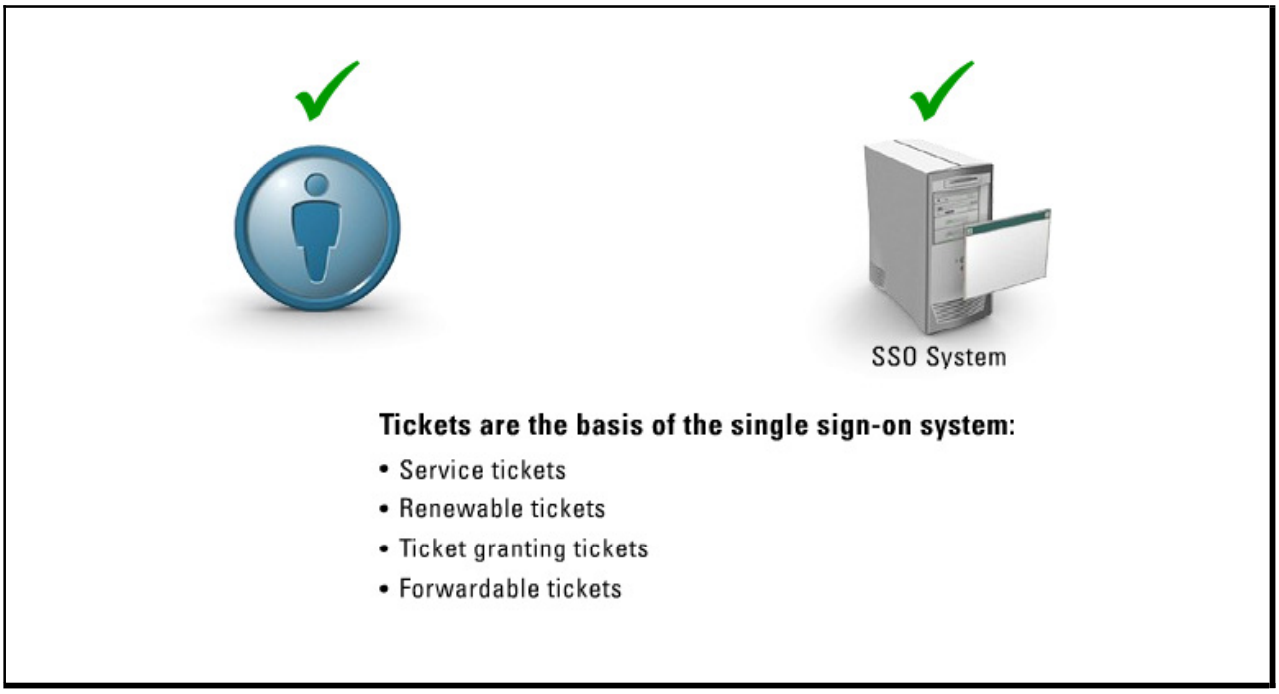

Tickets are the basis of creating a single sign-on scheme. Tickets are, in essence, passwords used to authenticate to a system or service.

Different tickets can be described as one of the following:

- **Service ticket** - A ticket that authenticates a principal to a service.
- **Renewable tickets** - In some cases, an application or service may want to have tickets that are valid for an extended period of time. However, the extended time could allow someone to steal these credentials, which would be valid until the ticket expired. Renewable tickets allow for applications to obtain tickets that are valid for extended periods while lessening the chances for theft.
- **Ticket granting tickets (TGT)** - A ticket that allows access to the ticket granting service. Ticket granting tickets are passed to the principal after the principal has completed a successful request. In a Windows 2000 environment, a user logs on to the network and the Key Distribution Center (KDC) verifies the principal's name and encrypted password, and then send a ticket granting ticket to the user.
- **Forwardable tickets** - Forwardable tickets allow a server to pass on the credentials of the requester to another service. For this pass-on to happen, the initial TGT must have been requested with the forwardable option, and the server must be allowed to delegate credentials.

Kerberos

This section will discuss the Kerberos single sign-on authentication mechanism, and how it can be implemented in the enterprise.

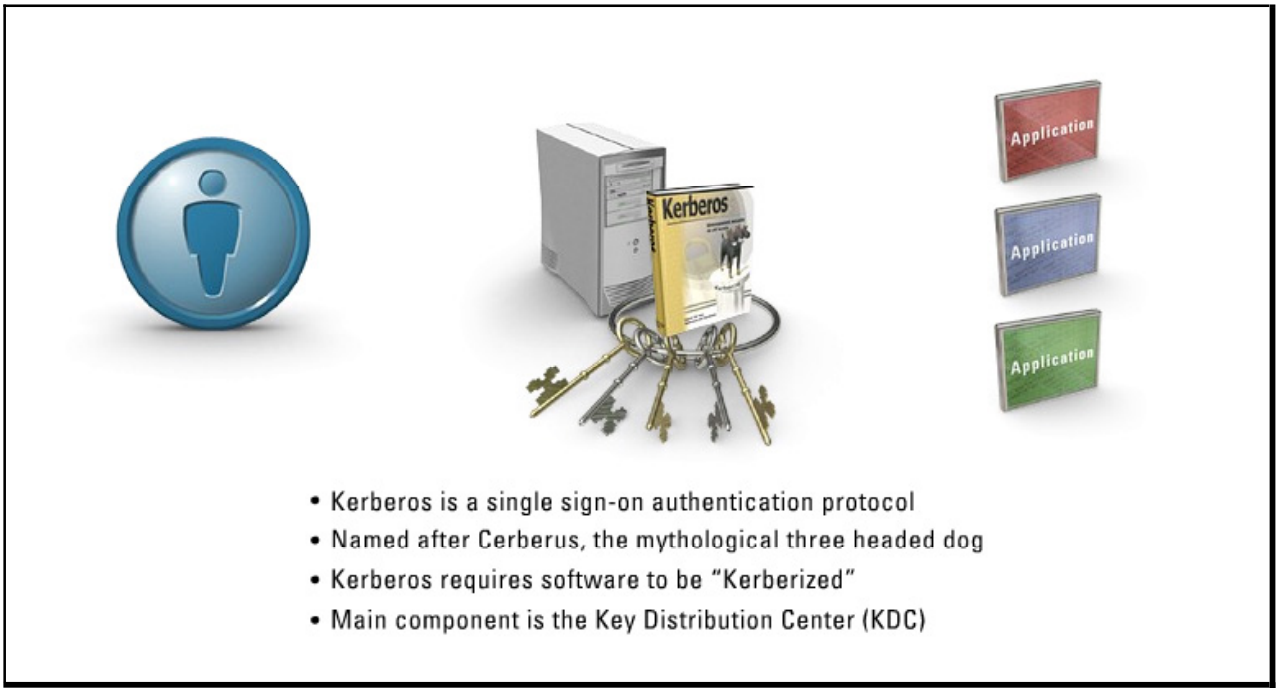

Kerberos is part of MIT's project Athena and is currently in version 5 of its existence. Kerberos version 4 was very simple and efficient; Kerberos version 5 is more complex and adds greater functionality. Kerberos is a single sign-on authentication protocol that uses a secret key based service, which can be used for network wide authentication. Named after Cerberus the mythological three-headed dog that guards the entrance to Hades, Kerberos requires authentication software to be modified to work (Kerberized).

Kerberos uses symmetric key cryptography to provide end-to-end security for users. Its main component is the Key Distribution Center (KDC), which is used to hold all users' and services' cryptographic keys. It provides authentication services, as well as key distribution functionality. The KDC provides security services to entities referred to as principals that can be users, applications, or services.

A ticket is generated by the KDC and given to a principal when that principal needs to authenticate to another principal.

A KDC provides security services for a set of components and principals. This is called a realm in Kerberos. The KDC is usually divided into two components, the Authentication Server (AS) and the Ticket Granting Server (TGS). The AS is the part of the KDC that authenticates a principle, while the TGS is the part of the KDC that makes the tickets and hands them out to the principles. Parties then use these tickets (session keys) for message encryption.

For Kerberos to work the following must be carried out:

- Each user must have an account on the KDC.
- The KDC must be a trusted server in a secured location.

- The KDC must share a DES key with each user.
- When a user wants to access a host or application, they request a ticket from the KDC via login and generate an authenticator that validates the tickets.
- To gain access, the user provides their ticket and authenticator to the application, which processes them for validity and will then grant access.

There are drawbacks to running Kerberos, which include the following:

- Each piece of software must be Kerberized.
- The KDC is a single point of failure.
- Secret keys are temporarily stored on users' workstations.
- Session keys are decrypted and reside on the users' workstations.
- Is vulnerable to password guessing.
- Network traffic is not protected.
- The AS must be able to handle a huge amount of requests.
- Requires all systems have synchronized time clocks.
- Relies on UDP which is often blocked by many firewalls.
- Kerberos v4 binds tickets to a single network address for a host. Host with multiple NIC's will have problems using tickets.
- When a user changes his password, it changes the secret key and the KDS needs to be updated.

Weaknesses in the Kerberos v4 RNG allowed secret keys to be easily guessed in some circumstances.

Some applications that use Kerberos include the following:

- Telnet
- FTP
- RSH
- NFS

SESAME

This section will discuss the SESAME single-sign authentication mechanism used in the European Union (EU).

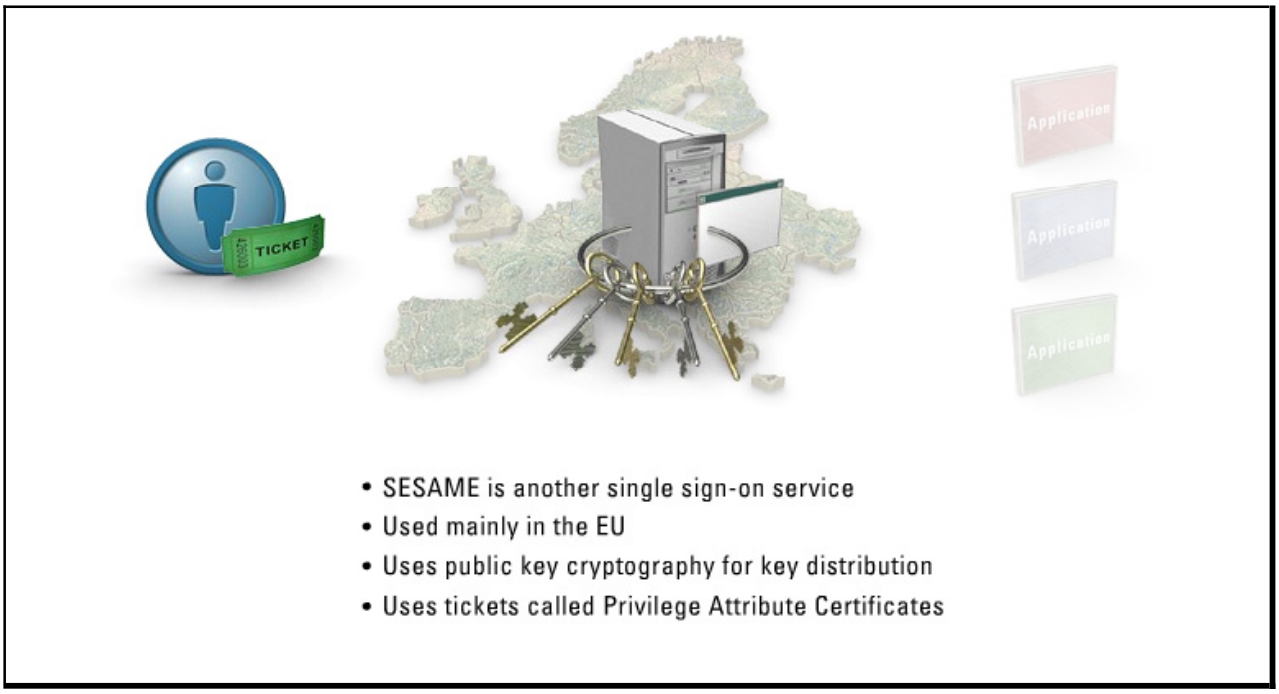

The European Commission funded a research and development project called The Secure European System for Applications in a Multi-vendor Environment (SESAME). SESAME is another single sign-on service that uses public key cryptography for the distribution of secret keys. Like Kerberos, it also uses a ticket for authorization, which in SESAME is called a Privilege Attribute Certificate.

To access the network system a user must perform the following:

- Authenticate to an Authentication Server (AS) to get a cryptographic protected token. The token is then used to prove the user's identity.
- Then user then presents the token to a Privileged Attribute Server (PAS) to obtain a guaranteed set of access rights contained in a Privileged Attribute Certificate (PAC). The PAC is digitally signed to prevent tampering and cryptographically protected from the point it leaves the PAS to the final requested resource.
- The user presents the PAC to a requested application/resource whenever access is needed.
- The requested application makes an access control decision according to the user's security attribute contained in the PAC.

Thin Clients

This section will discuss what thin clients are, why they would be used in the enterprise and the pros and cons of having them.

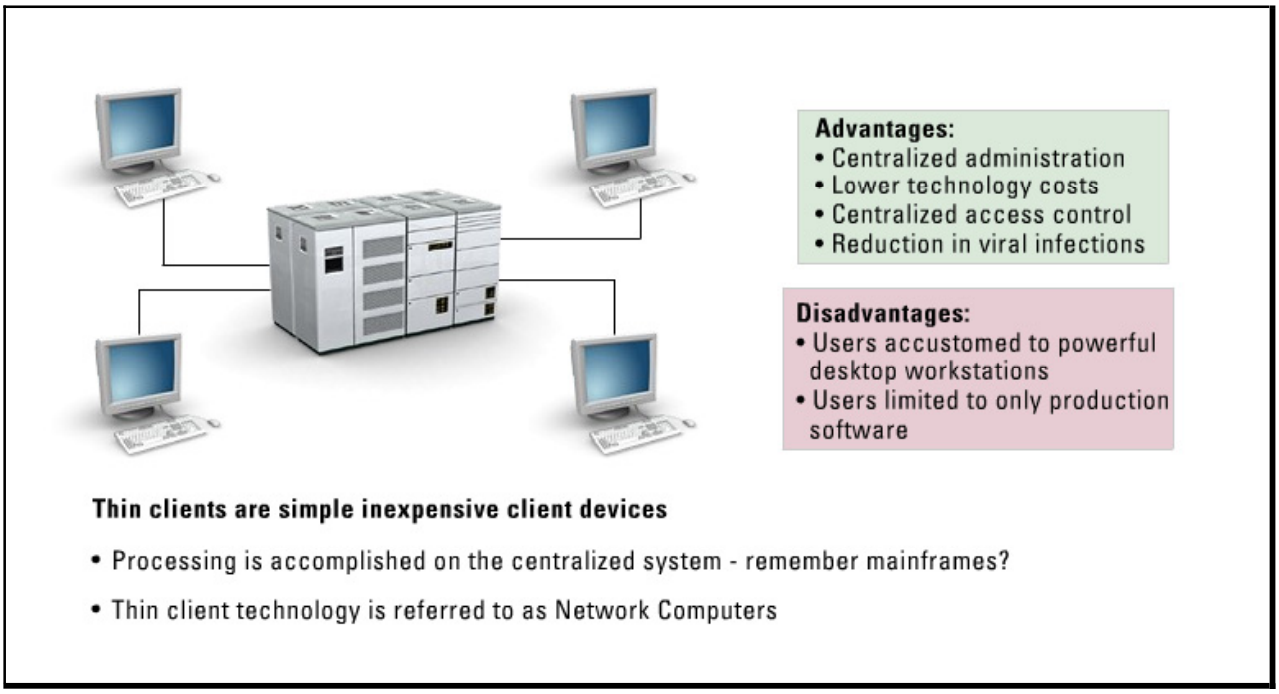

In a distributed networking environment, costs are always a concern, and thin clients fit perfectly in this solution. A **thin client** is a simple inexpensive client device that takes the place of the typical desktop computer. The thin client connects to the server over the network to process applications, access files, print, and perform other network services. All of the processing work is completed on the centralized system similar to a mainframe and is not local to the thin client. In fact, many thin clients can only send and receive keystroke information rather than actual data.

Advantages of thin clients include the following:

- Centralized administration
- Lower technology costs
- Centralized access control
- Reduction in viral infections

Disadvantages of thin clients include the following:

- Users accustomed to having powerful workstations on the desktop
- Users limited to only production software (no web browsing, listening to music, or playing games)

Thin client technology is also referred to as Network Computers (NCs). An NC is any type of computer system with minimal processing power, memory, and storage, which requires connection to a network resource in order to function.

Scripts

This section will discuss a method used to automate the often-laborious task of authenticating to multiple servers with different authentication accounts.

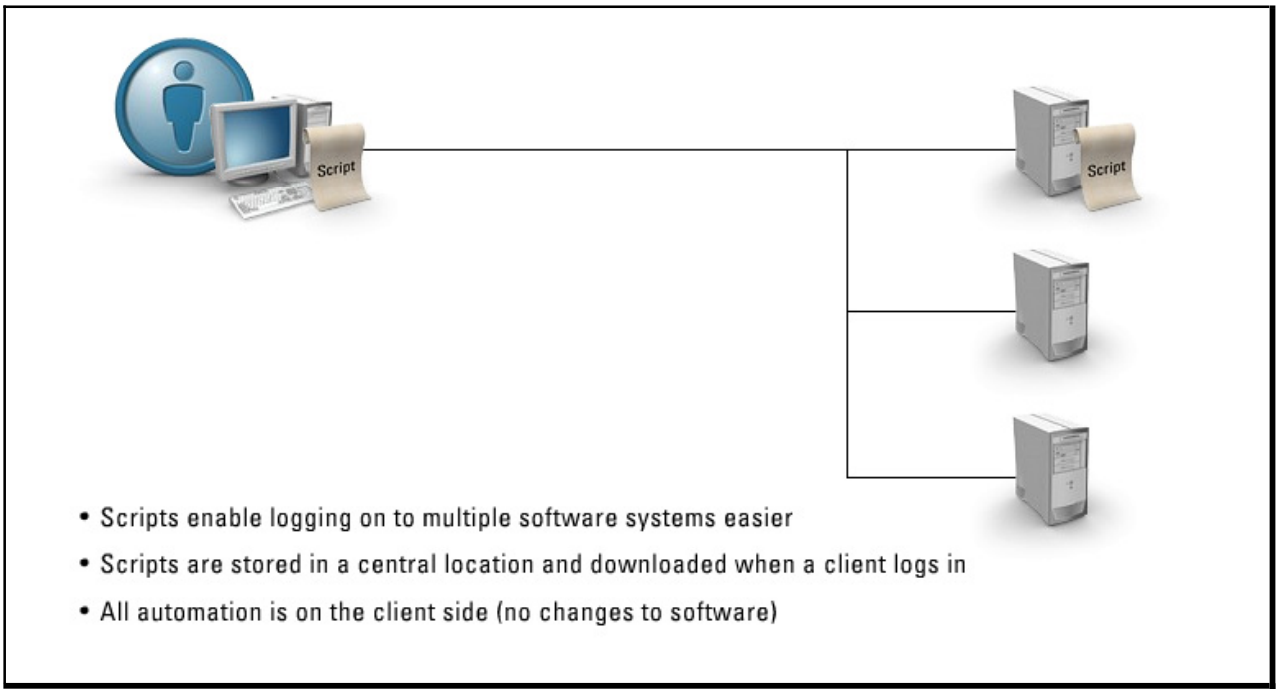

Single sign-on can also be accomplished through the use of scripts. In this solution, a user profile and automation scripts are stored encrypted in a central location and downloaded to the client when a user logs in. All automation is on the client side. No changes are required in the applications, application setup, or application administration.

Most client script SSO packages work on all versions of Windows and support sign-on automation for any application that runs on a Windows' machine, including terminal emulators, DOS, Windows, and web applications.

The client is one small executable that is easy to install and deploy and usually does not require administrator rights or rebooting. The basic premise is very simple:

- A user provides his or her normal login access credentials.
- Upon successful authentication, the user's profile and automation scripts are downloaded from the network. The automation scripts apply security policies and provide the user's authentication credentials to the application when called upon.
- The user executes the application, which first calls the script to provide authentication parameters to the application.
- After successful authentication, the application opens.

One-Time Passwords

This section will discuss what one-time passwords (OTPs) are and why they would be used for authentication in a network.

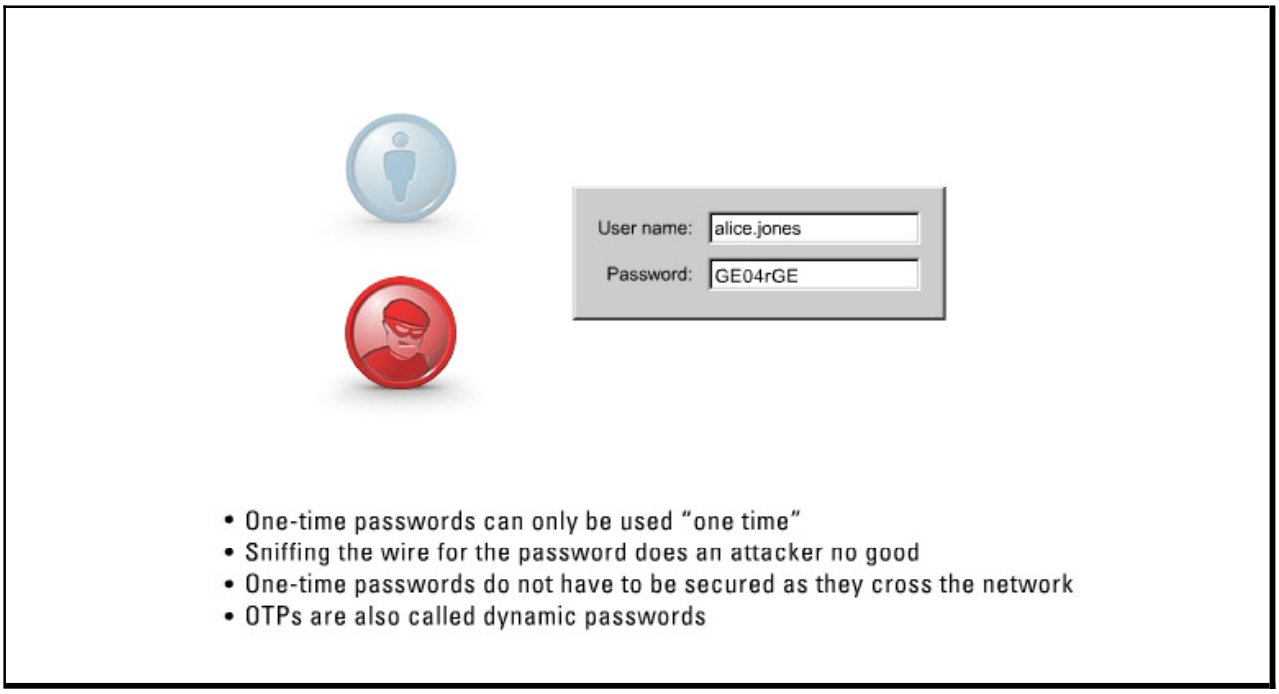

A **one-time password (OTP)** or a dynamic password is very different from the conventional password that can be used many times; an OTP can be used only once. A plus for the OTP is the fact that passwords to do not have to be secured as they cross the network. If an attacker sniffs the wire and manages to extract the password, it will do them no good as once the password is used, it cannot be used again.

An OTP scheme can print out a list of valid passwords than can only be used once, but this scheme is neither easy to use nor very secure, as it is a list on paper. Most OTP schemes use the services of a handheld token device.

Token-Based Method

This section will discuss what token-based authentication methods are, how they work, and what advantages they provide to an enterprise.

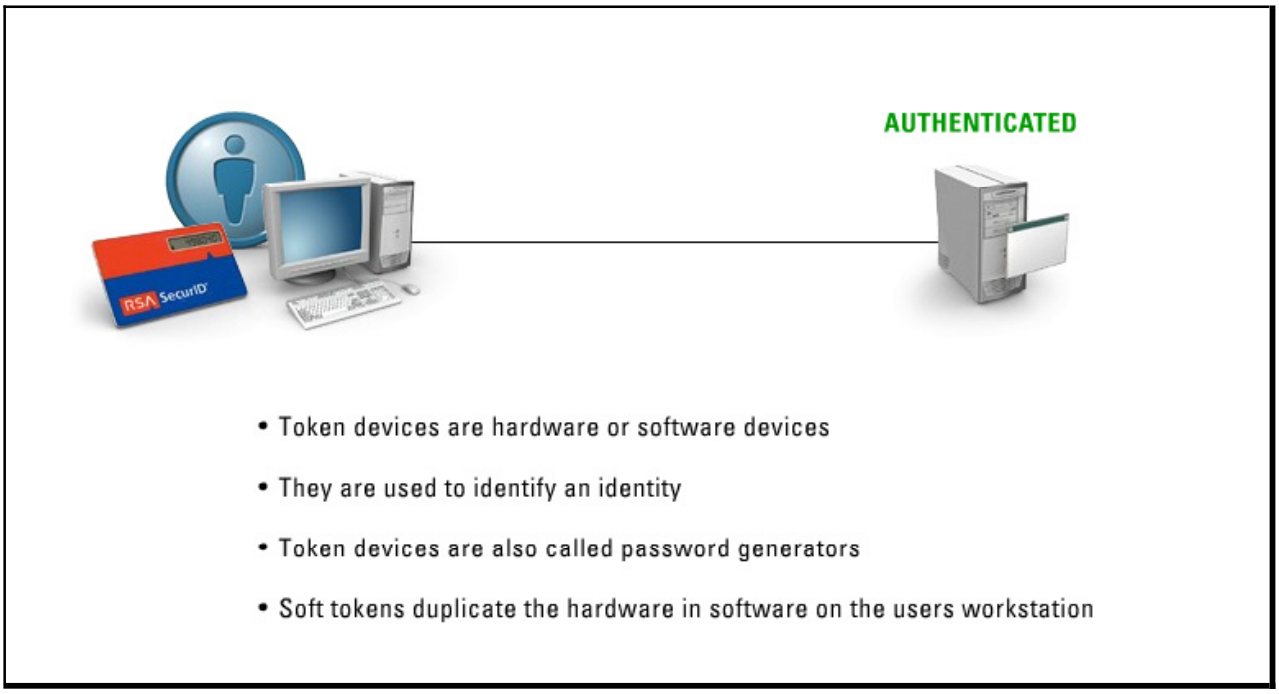

A **token-based** OTP scheme uses software or a hardware object to identify an identity in an authentication process. The token device, also called a password generator, usually has at least an LCD display and a keypad. Some token mechanisms have neither, which implies that possession alone along with the user's base secret is enough to gain access.

Instead of physical tokens, many corporations use what are called soft tokens, in which no hardware token is used at all. All algorithms that reside in a physical token are duplicated in a software application that is called up when the user needs to authenticate.

Synchronous

This section will discuss the various forms of synchronous token devices.

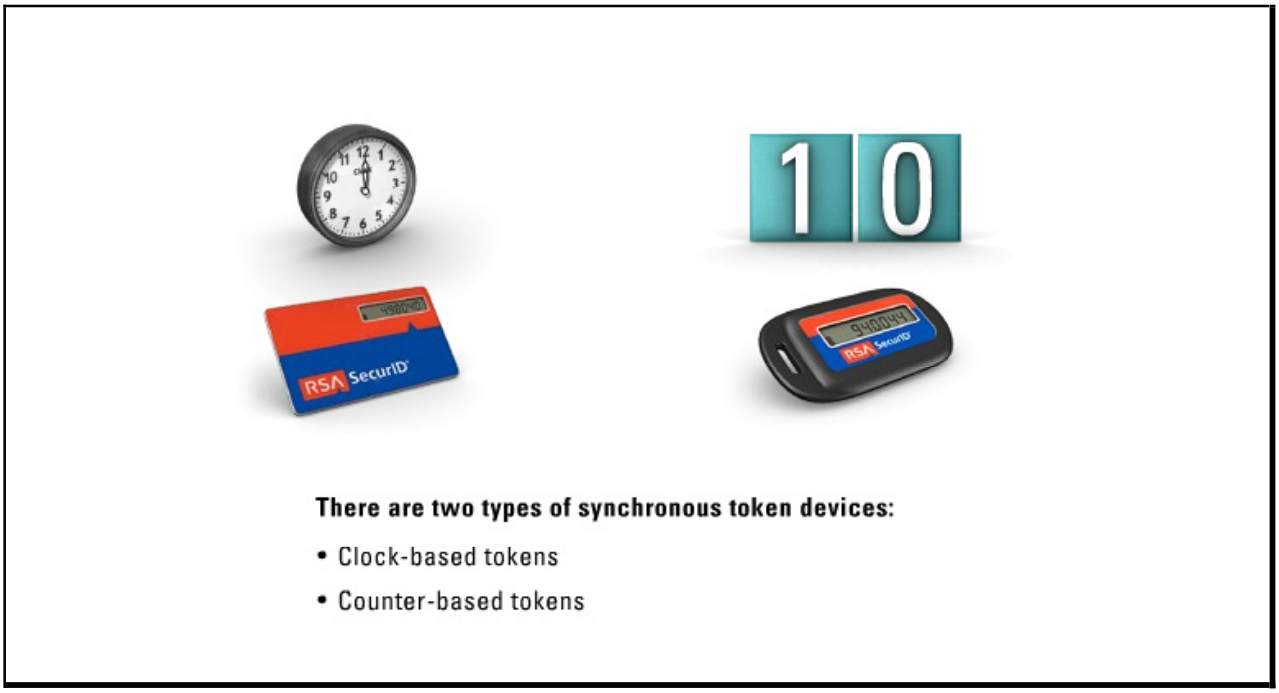

Two basic types of **synchronous token** devices carry out OTP authentication. Synchronous token methods are synchronized with an external source, usually a clock or a counter.

- **Clock-based tokens** - In this method, the token relies on an internal clock. This clock combined with a base secret key is used to generate a time-based password each time the user needs one. To validate the received password, the server performs the same function and determines the value of the password based on its time. If they match, the user is authenticated. The user has a certain window of time in which to use the password, usually about 60 seconds.
- **Counter-based tokens** - In this method, the administrator first inserts a specific base secret and internal counter into the user's token device and copies this information to the server device. To active the OTP service, the user simply pushes a button on the token. Once the user pushes the button, the token increments the internal counter combine with the owner's unique base secret, and the device computes a one-way hash result. The token formats the result and displays it for the user as the OTP.

Asynchronous

This section will discuss the asynchronous one-time password authentication method.

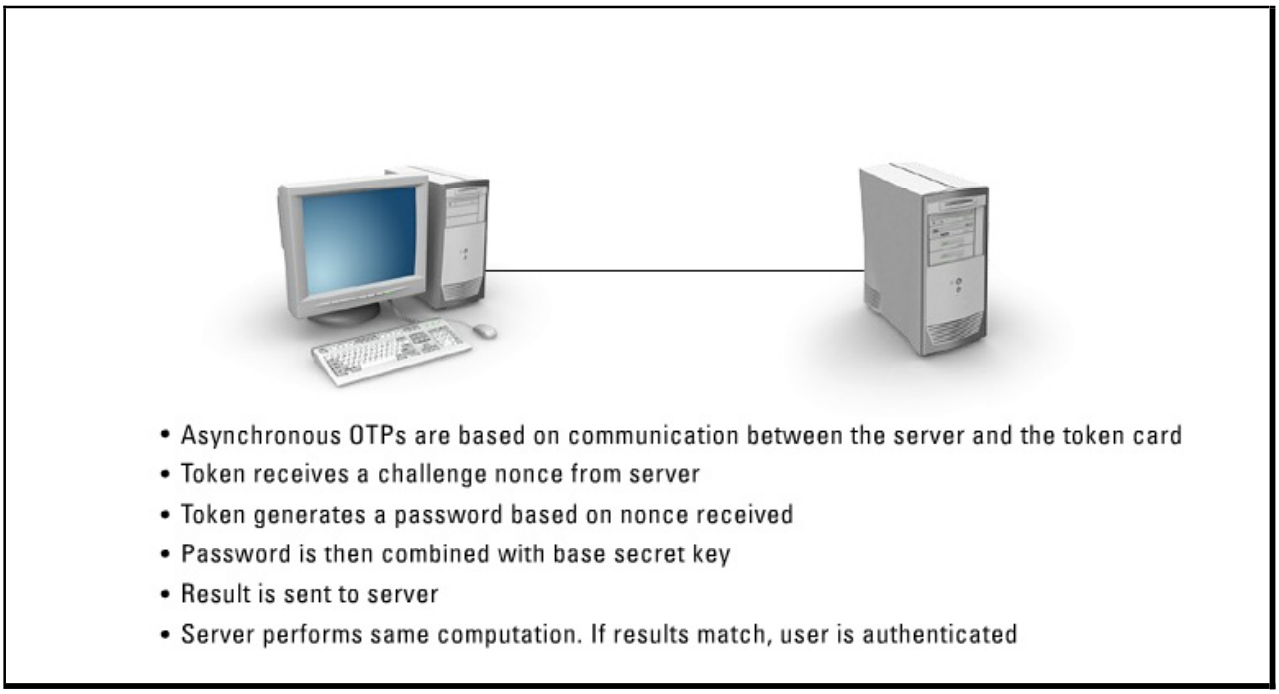

There is another type of OTP token mechanism, and it is the asynchronous OTP method. Here, the token generates a password based on a challenge or nonce, a random number, from the server. The number and password are then combined with a base secret key within the token. The user responds to the server's challenge, the nonce, using the result of this combination as its reply.

The steps for an asynchronous OTP request are as follows:

- The user requests a login, which the server replies with a challenge number.
- The user enters the challenge number into the token.
- The challenge number and base secret are combined inside the token and displayed for the user.
- The user enters the displayed number as the OTP for authentication.
- The server performs the same function and if the numbers match the user is authenticated.

Smart Card

In order to provide for greater security, higher mathematical computations must be performed. For example, instead of supplying a simple password, you could supply a very secure one-time password. Remembering how to compute these one-time passwords can be difficult to say the least. Smart cards provide the advantage of high mathematical computation when authentication occurs.

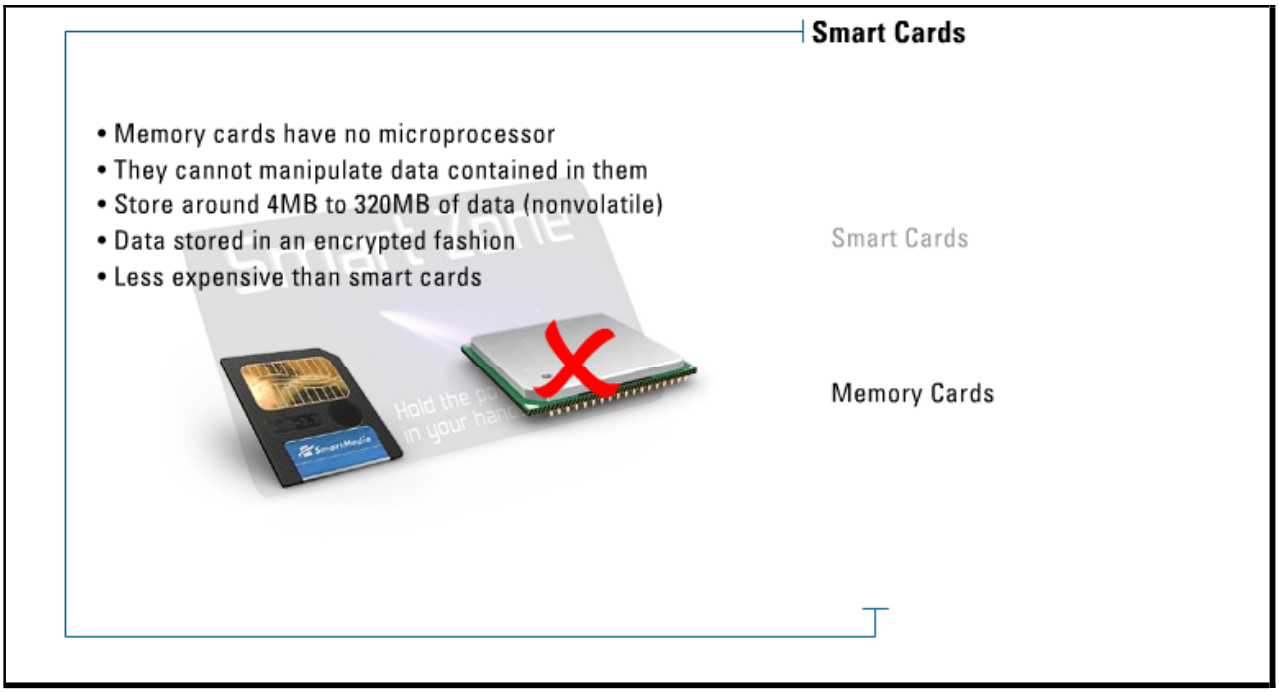

Smart cards are credit card-sized plastic cards that are similar in shape and size to a credit card. They are referred to as smart cards because they have an integrated microprocessor that allows manipulation and calculations to be performed on data stored in them. This capability means cryptographic functions for encrypting or decrypting passwords can be easily computed. In general, smart cards fall into two categories: contact or contactless.

- **Contact cards** need to be physically inserted in a card reader and make a physical connection with metallic contacts on the smart card.
- **Contactless cards** use an electromagnetic signal and an antenna on the card to create the connection between the card and the card reader.

Memory cards are different from smart cards in that memory cards have no microprocessor, which means they cannot manipulate data contained in them. Memory cards store from four to 320MB of data in a non-volatile fashion. All data stored on the memory card is in an encrypted format. The encryption means if the card is lost or stolen, the data stored on it cannot be retrieved. Because memory cards do not contain microprocessor chips, they are less expensive than smart cards and have a correspondingly lower level of security.

Pros of the smart card:

- Capability to store information with a high degree of security and portability
- Hacker resistant storage

- Offer an enterprise wide authentication system because the user can use the card for all authentication mechanisms
- Provides multi-factor authentication

Cons of the smart card:

- Susceptible to invasive or non-invasive attacks
- Invasive - attacks on the card render it inoperative
- Non-invasive - attacks information on the card without damaging it
- Lacks global standards for how data is embedded on the card
- Susceptible to “Optical Fault Induction Attacks”- A camera’s electronic flash with a microscope can be used to reveal private data directly from a smart card’s microprocessor.

Key Card

Some forms of authentication are cheap and easy to install and use. The key card provides one such method. This section will discuss what key cards are and how they can provide a secure and cost effective means of authentication.

Key cards are very inexpensive plastic cards that look much like a credit card. They have a magnetic strip that can be used to encode a minimal amount of information. Normally, data stored on the card is encrypted as a security precaution against the card being lost or stolen. You may have used a key card if you have visited a hotel in the last couple years. Instead of a key to open your room, the hotel provided you with a card that is slid into a reader that authenticates and unlocks the door.

Summary

The key points discussed in this lesson are:

- Identification and Authentication Techniques
- Performance Measures for authentication
- Knowledge Based authentication
- Passwords
- PINs
- Pass Phrases
- Characteristics-based Access Control
- Biometric methods
- Behavior methods
- Single Sign-On (SSO)
- Tickets
- Kerberos access
- SESAME access
- Thin client devices
- Scripts
- One-time passwords
- Token-based access
- Synchronous access
- Asynchronous access
- Smart card access

Passwords

Overview

Everyone uses some form of password many times every day in a business environment. Passwords are the keys that open the door to our computer systems, email, and resources. As such, passwords are high profile targets that attackers actively seek. Understanding the various password mechanisms, strength characteristics, and exploits is a fundamental necessity of any information security specialist.

Importance

Passwords are the most common element of protecting resources in use today. Therefore, identifying the strength characteristics of the password themselves, as well as the mechanisms used to validate them, are essential to an information security professional.

Objectives

Upon completing this lesson, you will be able to:

- Describe password selection
- Explain password management
- Explain password control
- Define file and data ownership and custodianship
- Describe methods of attack
- Explain a brute force attack
- Identify a password dictionary attack
- Explain denial-of-service
- Define spoofing
- Explain sniffers

Outline

The lesson contains these topics:

- Selection
- Management

- Control
- File and Data Ownership and Custodianship
- Methods of Attack
- Brute Force
- Password Dictionary Attack
- Denial-of-Service
- Spoofing
- Sniffers

Selection

This section will discuss the various ways a password system can be made to be secure in the enterprise.

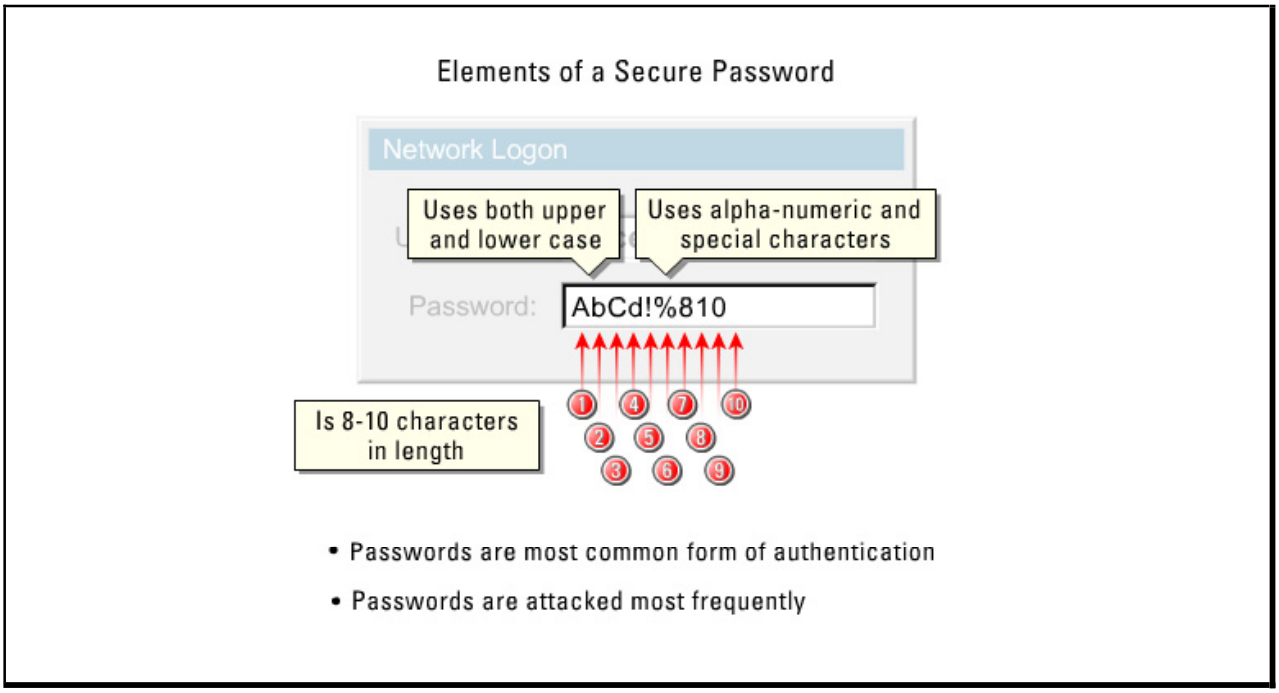

The password is the most commonly used form of authentication in the corporate industry today. The password is also the most commonly attacked authentication mechanism because of the inherent problem associated with them. The password's problem, which is also its greatest strength, is that a password is just a word. It can be any word, from a name to a title, to any common phrase, verb or noun.

All that is required to authenticate is to match the password with the password stored in the authentication server. If a match occurs, access is given. Because access is only a word, information security specialists must enforce a secure password scheme on users. This scheme can be done through software when the password is first entered. If the password is less than a certain length, is easily found in a dictionary, or the password only contains alphabetic characters the system can be configured to reject the password and ask the user to try another to meet the criteria for a secure password.

A secure password is a oxymoron; all passwords can be eventually cracked given enough time. Certain precautions can be taken to produce passwords that have a reasonable chance of thwarting attack. These precautions include the following:

- The length of the password should be a minimum eight to 10 characters
- Use both upper and lower case
- Use alphanumeric and special characters

Note	These are examples of special characters: !@#$%&

Management

This section will discuss the various management mechanisms that can be used to provide a secure password scheme in the enterprise.

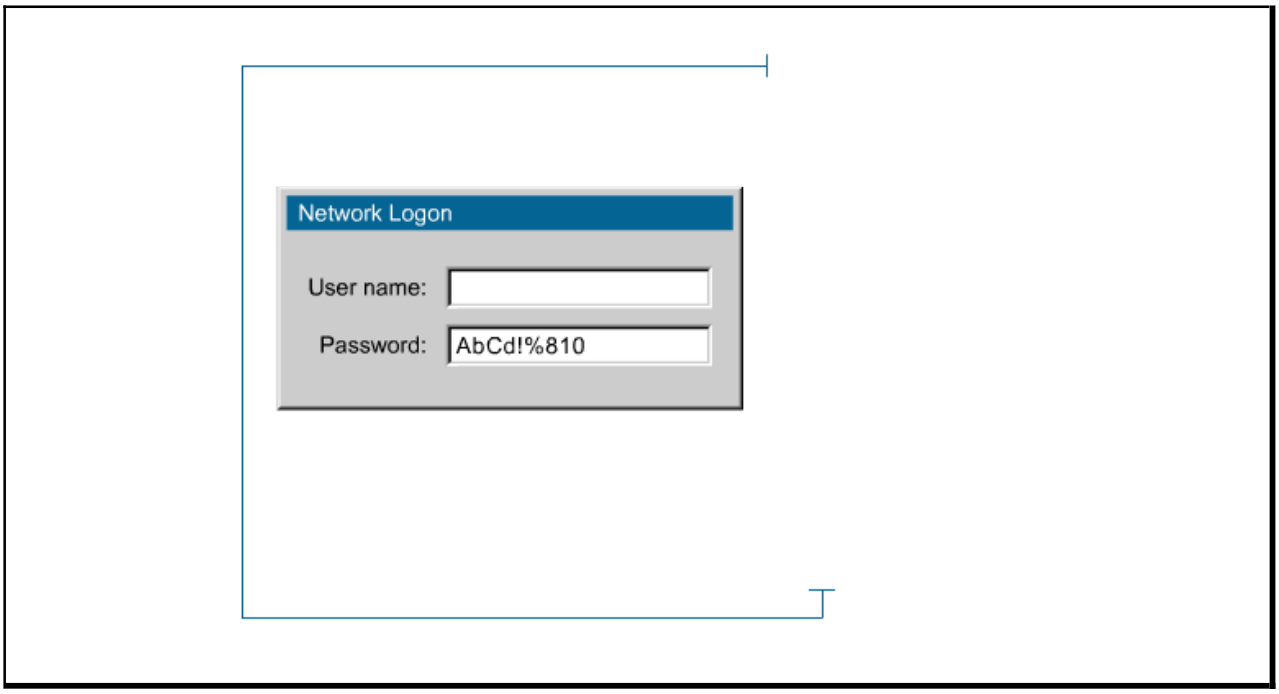

Any password can be compromised because they are usually:

- **Insecure** - Given the choice, people will choose easily remembered words and hence easily guessed passwords such as names of relatives, pets, phone numbers, birthdays, hobbies, etc.
- **Easily broken** - Programs such as crack, SmartPass, PWDUMP, NTCrack and L0phtcrack can easily decrypt Unix, NetWare, and NT passwords. Dictionary attacks succeed because users choose easily guessed passwords.
- **Inconvenient** - In an attempt to improve security, organizations will issue a computer-generated passwords that are difficult, if not impossible, to remember. Often, the user usually ends up writing them down and usually sticking them to their monitor thus compromising security.
- **No repudiation** - Unlike a written signature, when a transaction is signed with only a password, real proof of the identity of the individual that made the transaction is unavailable.

The best passwords are those that are both easy to remember and hard to crack using a dictionary attack. The best way to create passwords that fulfill both criteria is to use two or more small-unrelated words or phonemes, ideally with special characters and numbers. Good examples would be W@y0utThere (the 0 is a zero) or 1Glass&BOOK.

Do not allow users to use:

- Common names, DOB, spouse, phone number, etc.
- Words found in dictionaries
- The word “password” as a password
- System defaults - force a change on first authentication

Control

This section will discuss the forms of password control that can be placed on a system.

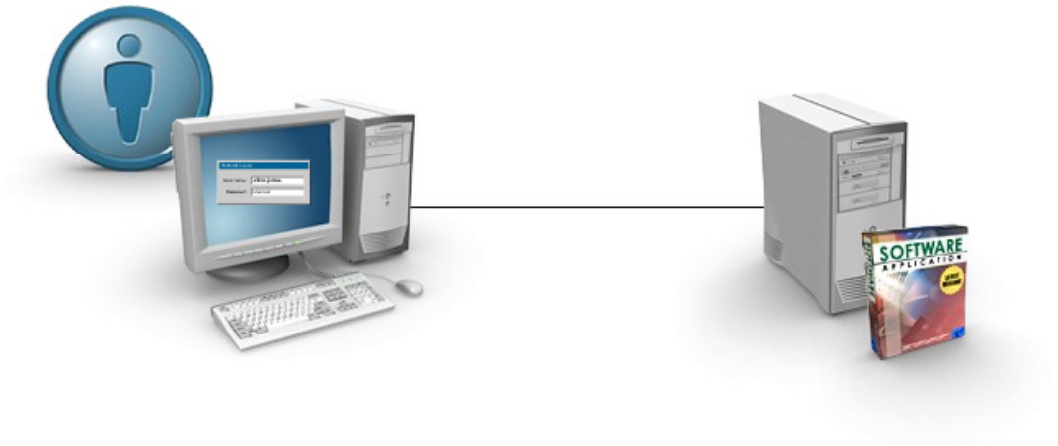

Password Control is Accomplished Through Software

Password control is usually accomplished through software. Windows NT, 2000, and almost all variants of UNIX allow the administrator to create stringent restrictions on the password used by users. Examples would include the following:

- Expiration date for the password
- Forced alphanumeric configuration
- Forced Upper and lower case letters
- Minimum length for a password
- No reuse of old passwords
- Unsuccessful login limit
- Concurrent usage limit
- Auditing
- Last login date used in the banner

File and Data Ownership and Custodianship

This section will discuss the differences among a file owner, data ownership, and a data custodian.

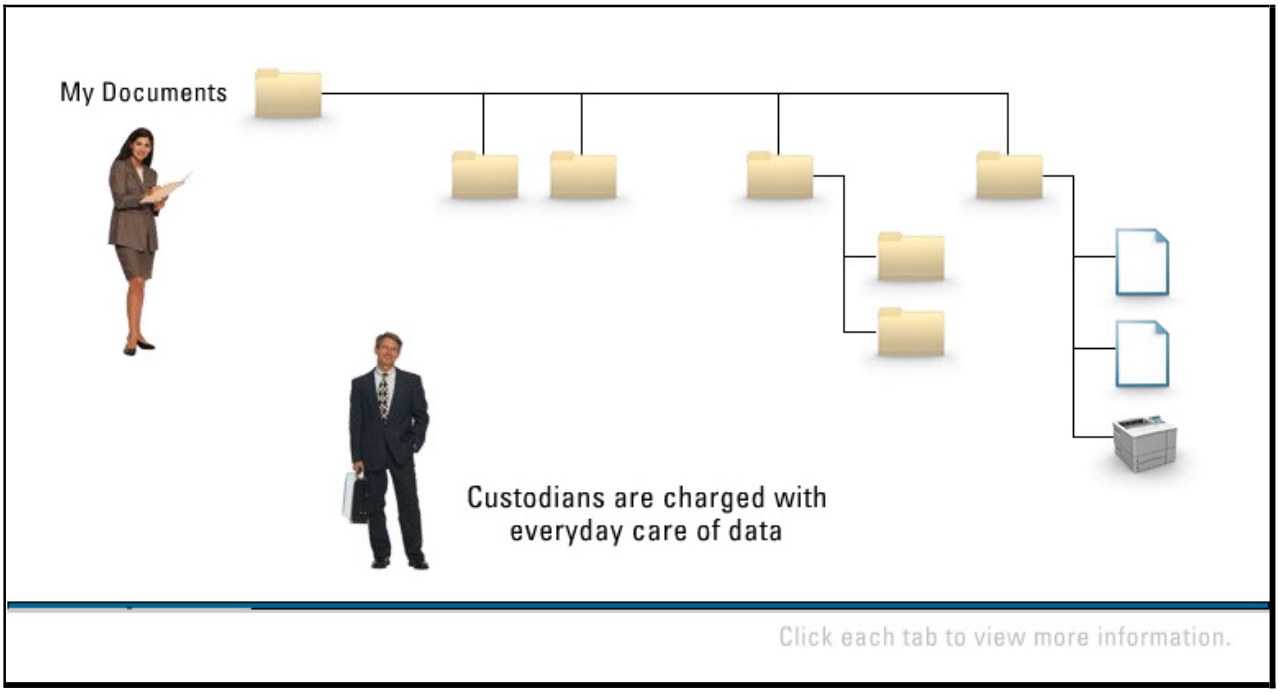

All information generated or used must have a designated owner. The owner must determine the appropriate classification and access controls. The owner is also responsible for ensuring appropriate controls for the storage, handling, and distribution of the data.

Custodians are charged by the owners for the everyday care of their data, which includes backups and general care.

Implementation of data classification requires support from higher management. Policy enforcement at the highest level is critical to the security success.

Policies that should be considered should include the following:

- Define information as an asset of the business unit.
- Declare local business managers as owners of the information.
- Establish Information Systems staff as custodians of the information.
- Clearly define roles and responsibilities.
- Determine data classification criteria.
- Determine controls for each classification.

Methods of Attack

This section will discuss the various forms of attack used against passwords and password files.

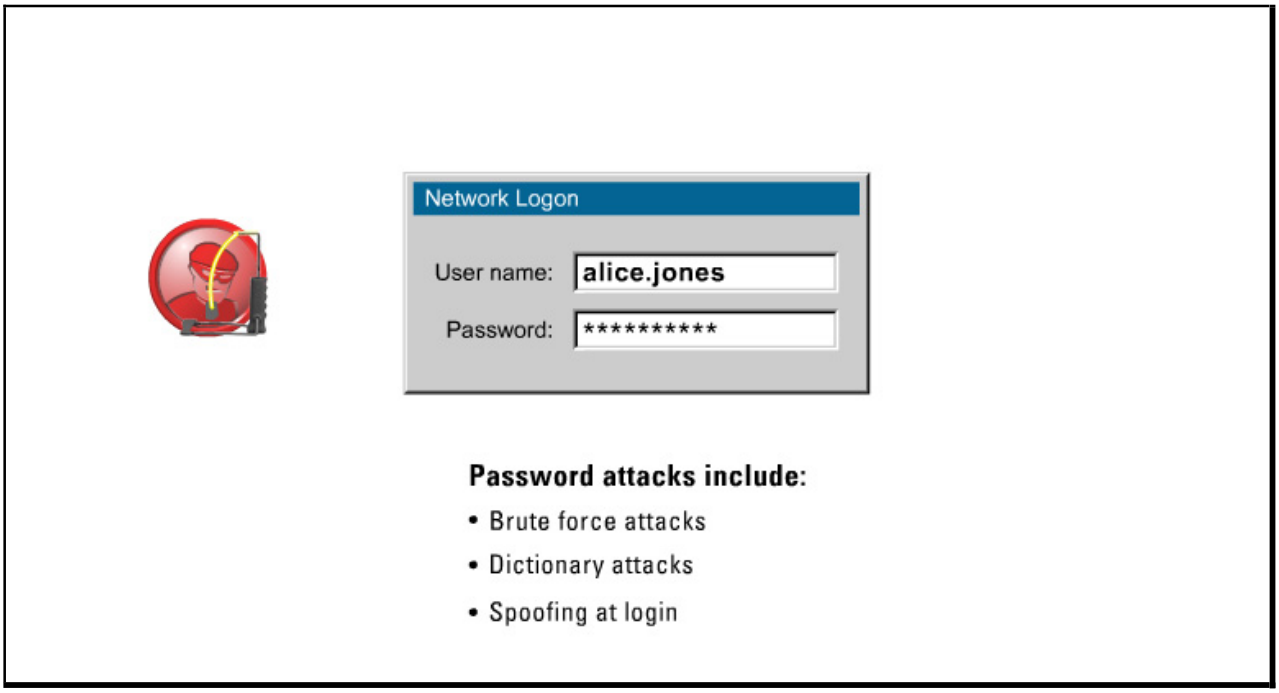

Methods of attack against passwords would include the following:

- **Brute Force Attack** - A brute force attack is an attack that continually tries different inputs to achieve the predefined goal of finding the correct password. Brute force attacks are also used in war dialing efforts.
- **Dictionary Attack** - Dictionary attacks are programs that enable an attacker to identify user credentials. The program is fed lists of commonly used words or combinations of characters, and the program applies these values to a logon prompt.
- **Spoofing at Login** - Spoofing at login is accomplished via a program that presents a fake login screen, to obtain user credentials.

Brute Force

This section will discuss what brute force attacks against passwords are and how they are accomplished.

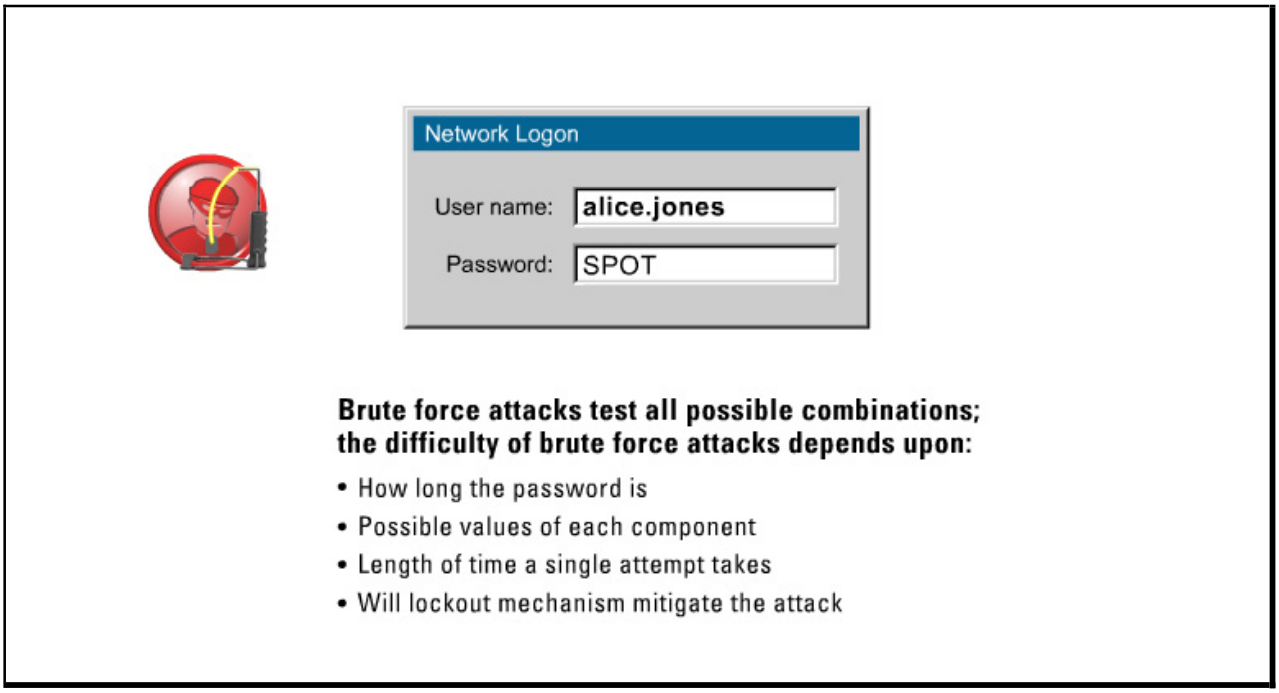

A brute force attack uses a software program that tests all combinations of potential passwords. Typically, the software tries a as the password, then b, then c, etc. until all possible single character possibilities have been tried. The software then attempts all two-character possibilities, and so on until the correct password is found.

The difficulty of conducting a brute force attack depends on several factors:

- The length of the password
- The number of possible values a password may have
- The length of time needed for a single attempt
- A lockout mechanism that disallows attempts after a certain number of invalid attempts

Brute force programs can take hours, days, weeks, or more to find the correct password. Having a password of sufficient length is important in mitigating this type of attack. In addition, imposing a delay of say 20 seconds between failed attempts or locking the account after 10 failed attempts deters this type of attack.

Brute force attacks can be conducted in-band, such as at a login prompt on a web page, or out-of-band on the attacker's workstation when the password file has been stolen.

Brute force programs include:

- L0phtCrack
- Brutus
- WebCracker

Password Dictionary Attack

This section will discuss what password dictionary attacks against passwords are and how they are accomplished.

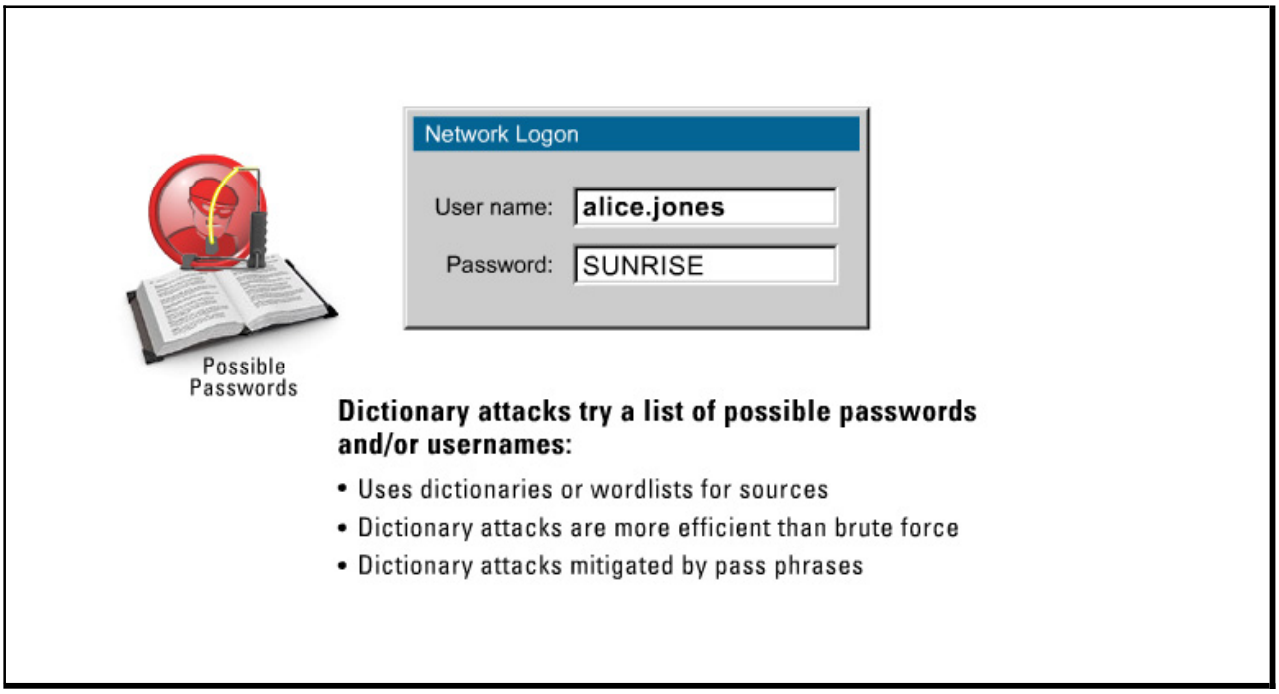

The dictionary password attack involves trying a list of possible passwords. These passwords and possibly usernames are located in a file called a dictionary or wordlist. Dictionary attacks are usually more efficient than the brute force attack because users generally use passwords that are easy to remember. Dictionary attacks are usually mitigated by systems that use pass phrases instead of passwords.

In most cases, dictionary attacks find the password more quickly than brute force attack;, however, a brute force attack will eventually obtain the correct password.

Improving the chances of a dictionary attack involves improving the dictionary. Creating a larger dictionary or using more dictionaries improves the attacker's odds of success. For example, using technical and foreign language dictionaries increases the overall chance of an attack succeeding. Another method of improving the chances of success in a dictionary attack is using string manipulation. Here, each password in the dictionary will be slightly manipulated and tried as a separate authentication attempt. For example, a single dictionary word might be 'crazy', but with string manipulation the following attempts would also be tried: yzarc (crazy backwords), Crazy, CRAZY, crazy1, crazy12, crazy123, and crazy1234.

Dictionary attack programs include:

- Crack
- John the Ripper
- UnSecure
- Brutus

Denial-of-Service

This section will discuss how Denial of Service (DoS) attacks against a password system can be carried out.

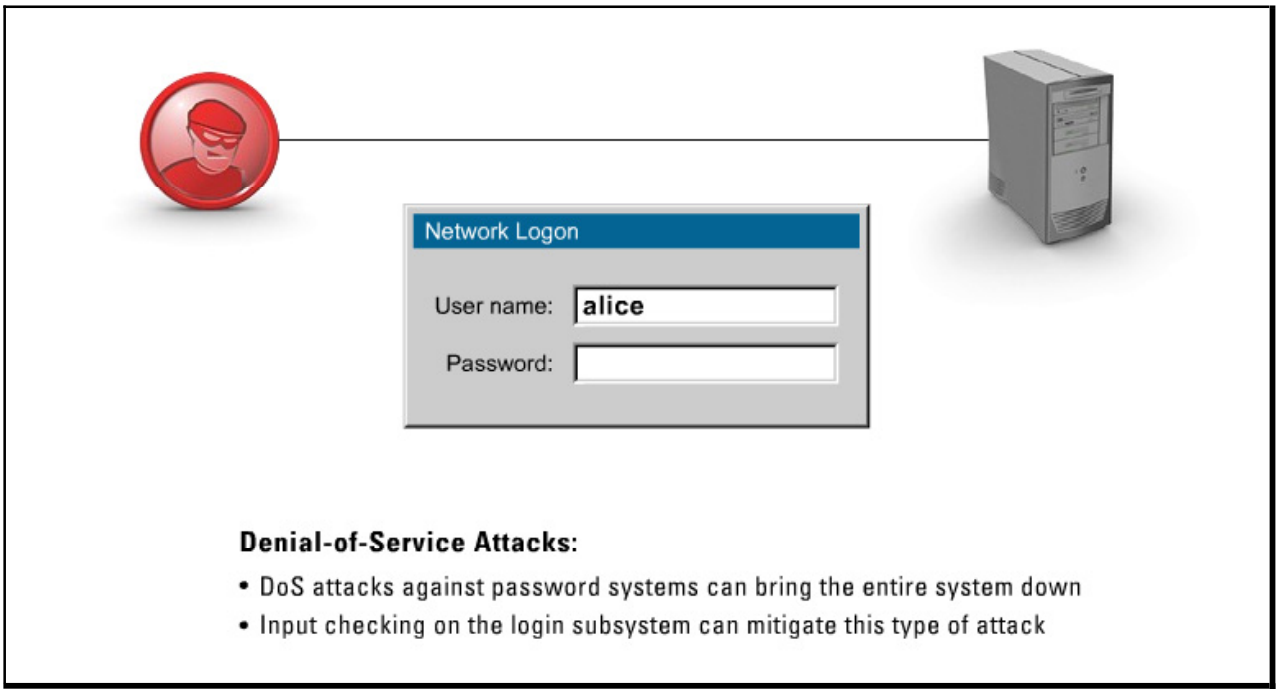

A **Denial-of-Service (DoS)** can be characterized as an attack on the operating system that renders the target unable to reply reliably, if at all. An example of a DoS attack against a password system is an exploit in a certain version of RADIUS running on Windows NT, Linux, and other UNIX based systems. In this attack, when an attacker appends a certain amount of spaces after the username, the RADIUS system crashes, keeping users from logging in.

To mitigate this type of problem, input-checking included in the login subsystem can easily stop this type of attack.

Spoofing

This section will discuss how spoofing attacks against a password system is perpetrated.

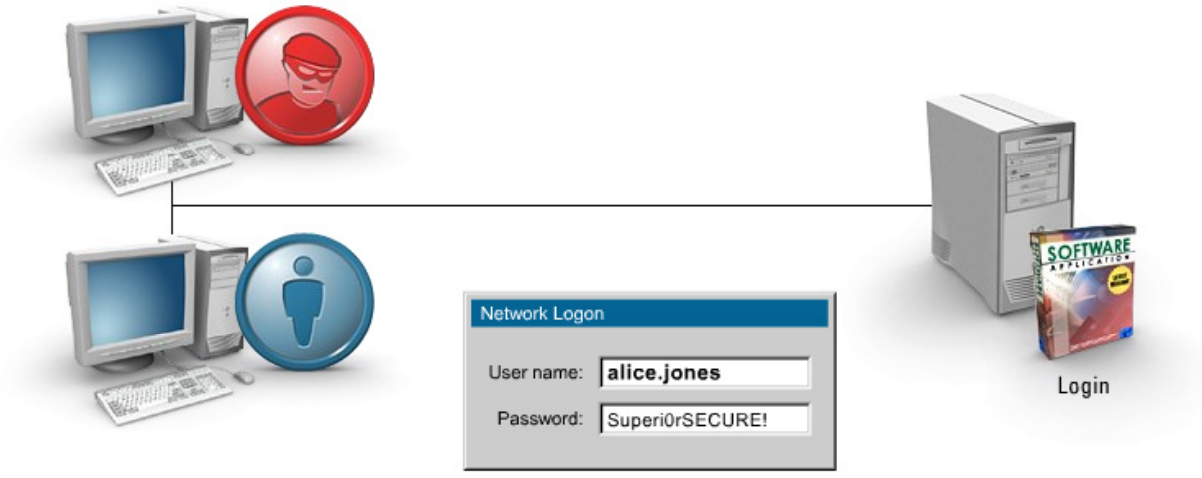

Password spoofing attacks pretend to be a trusted process

A **spoofing attack** on a password system is one in which one person or process pretends to be another person or process that has more privileges. An example would be a fake login screen also called a Trojan horse login. In this attack, the attacker obtains low-level access to the system and installs a malicious code that mimics the user login screen. On the next attempt to login, the user enters his username and password into the fake login screen. The malicious code then stores the username and password in a certain location or may even email the information to an email account. The Trojan horse then calls the correct login process to execute. To the user, the entry appears to be an incorrect or mistyped username or password and he or she will try again. When they do, of course, they are let into the system.

Sniffers

This section will discuss how sniffing can be used to obtain passwords and what can be done to mitigate this type of threat.

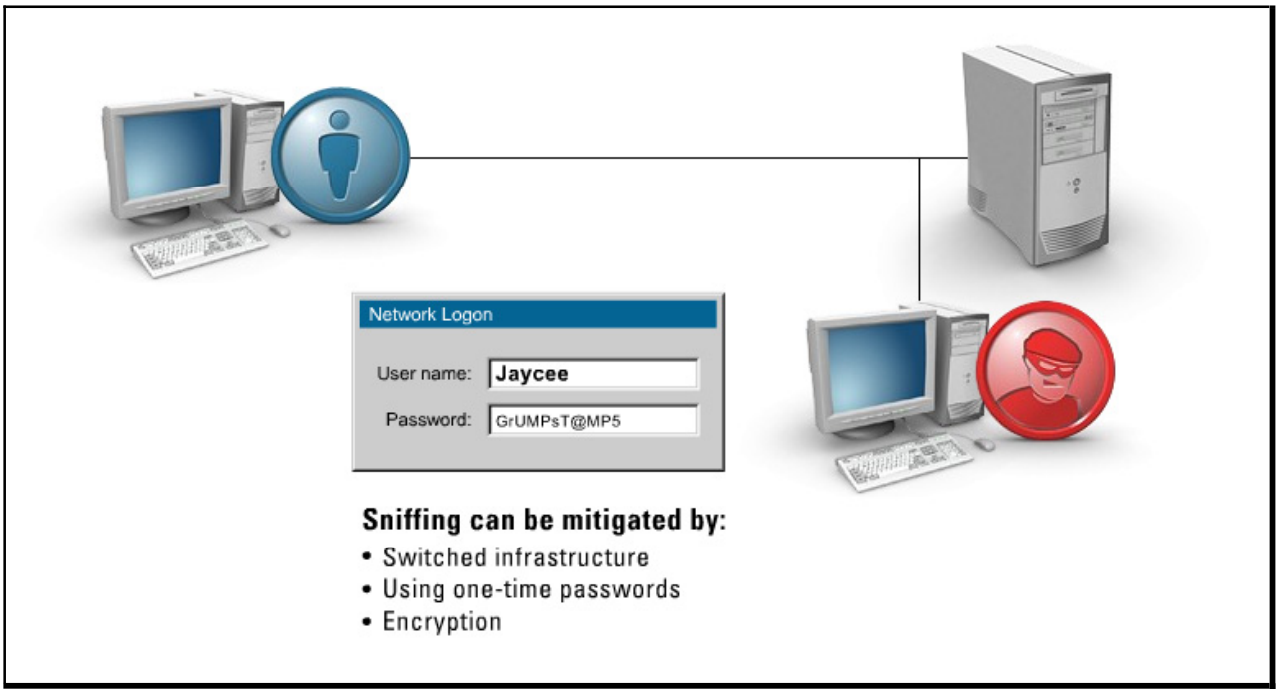

The act of sniffing is best described as a program or device that monitors data traveling over a network. Sniffing is hard to detect because as a passive attack, it only receives information and never sends out information. The goal of sniffing is to capture sensitive information such as a password in order to perform a replay attack at a later time.

Mitigation against sniffing attacks can include using a switched infrastructure, using one-time passwords, or enabling encryption.

Summary

The key points discussed in this lesson are:

- Password selection
- Password management
- Password control
- File and data ownership and custodianship
- Methods of attack
- Brute force attack
- Password dictionary attack
- Denial-of-service
- Spoofing
- Sniffers

Access Control Techniques

Overview

Access control mechanisms vary in part on how the enterprise is adopted, as well as how the individual members in the enterprise are ranked. For example, a manufacturing facility will have a different access control model than a military installation. Different models provide differing levels of security based on the tenets of the implementer. This lesson will identify the most successfully used access control models and will explain how and when the models would be used.

Importance

Understanding the access control model used in an installation is critical. The information allows the information security professional insight into the security mentality of the enterprise as well as a clear picture of how information dissemination occurs between entities.

Objectives

Upon completing this lesson, you will be able to:

- Explain the Discretionary Access Control (DAC) model
- Explain the Mandatory Access Control (MAC) model
- Define the Lattice-based access control model
- Define the Rule-based access control model
- Explain Role-based access control model
- Explain how restricted interfaces provide security
- Explain how non-discretionary access control provides security
- Explain how access control lists provide security
- List security models
- Explain the principles of the Bell-LaPadula model
- Explain the principles of the Biba model
- Explain the principles of the Clark-Wilson model
- Explain the principles of the non-interference model
- Explain the principles of the state machine model

- Explain the principles of the access matrix model
- Explain the principles of the information flow model
- Define the rule of least privilege approach
- Define the separation of duties approach
- Define the rotation of duties approach
- Define network segregation
- Define a control zone

Outline

The lesson contains these topics:

- DAC
- MAC
- Lattice-Based Access Control
- Rule-Based Access Control
- Role-Based Access Control
- Restricted Interfaces
- Non-Discretionary Access Control
- Access Control Lists
- Security Models
- Bell-LaPadula
- Biba
- Clark-Wilson
- Non-Interference
- State Machine
- Access Matrix Model
- Information Flow Model
- Rule of Least Privilege
- Separation of Duties
- Rotation of Duties
- Network Segregation
- Control Zone

DAC

Access control is the collection of mechanisms that permits managers of a system to exercise discretion or a restraining influence over the behavior, use, and content of a system. It permits management to specify what users can do, which resources they can access, and what operations they can perform on a system.

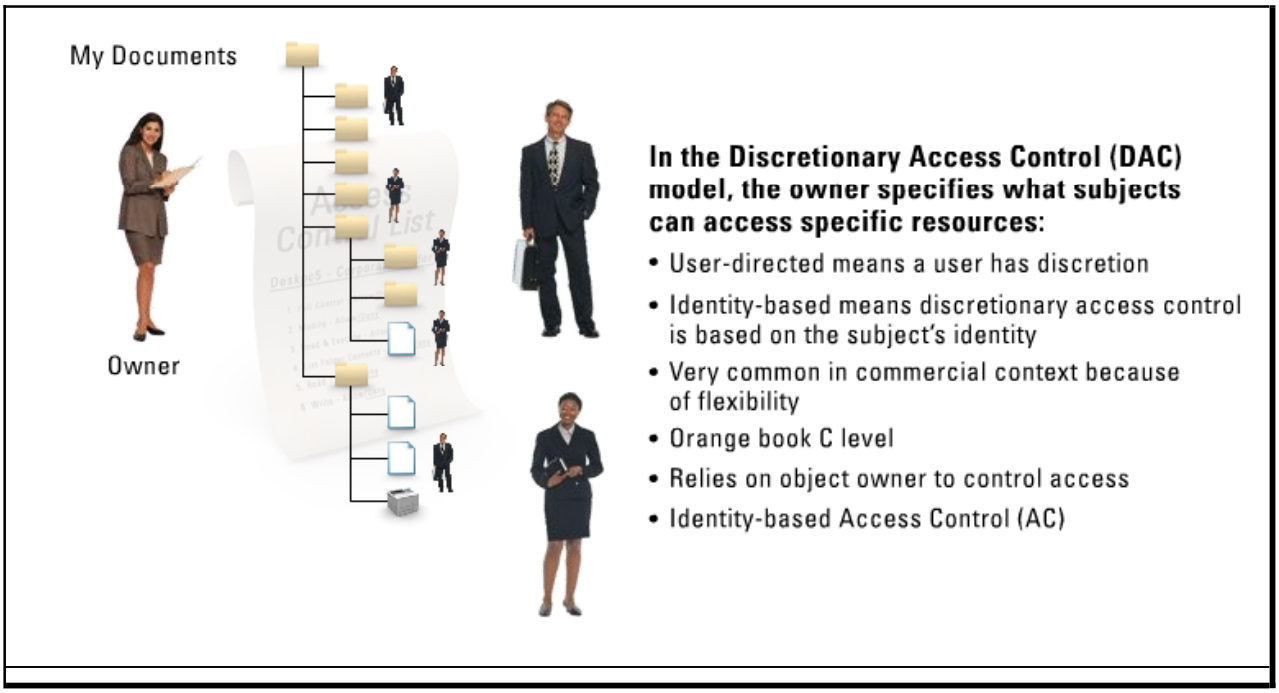

Access control models vary, and they dictate how subjects access objects. These models use various access control technologies and security mechanisms to enforce the rules and objectives of the model. One such model is the Discretionary Access Control (DAC) model. In the DAC model, the owner of the resource specifies what subjects can access specific resources. In other words, if you are using a DAC model, you decide how you want to protect and share your data. The most common implementation of DAC is through access control lists (ACLs). Other relevant information about the DAC model includes the following:

- Restricted access based on the authorization granted to the user
- Definitions based the Orange book C-level
- Separation and protection of prime users from unauthorized data
- Use by Unix, NT, NetWare, Linux, and Vines
- Reliance on the object owner to control access

MAC

In a Mandatory Access Control (MAC) model, users do not have much freedom to decide who can access their data files. MAC defines an imposed access control level.

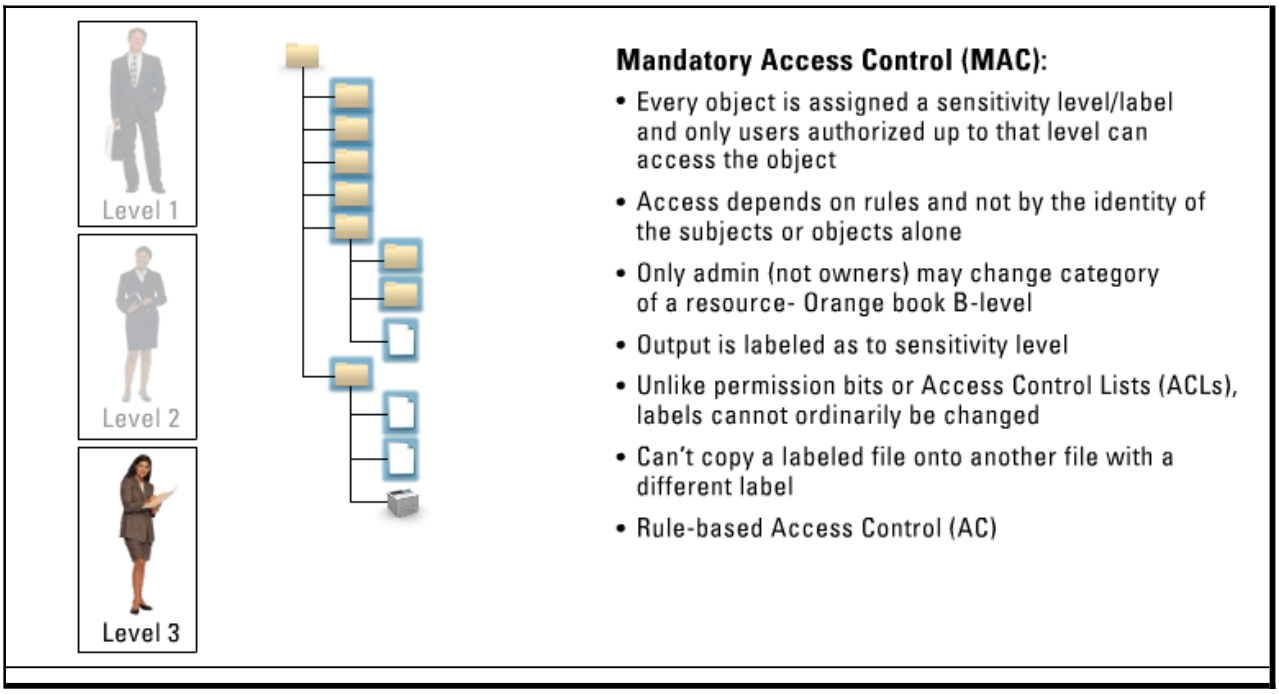

In this type of control system, decisions are based on privilege (clearance) of subject (user) and sensitivity (classification) of an object (file) through the use of labeling. In other words, the system decides how data will be shared. For example, the military classifies a document at secret. A user can be granted the secret privilege and have access to objects with this classification or lower as long as they have a "need to know". Other relevant information about the MAC model includes:

- The system assigns sensitivity levels, AKA labels.
- Every object is given a sensitivity label and is accessible only to users who are cleared up to that particular level.
- Only the administrators, not object owners, may change the object level.
- MAC is generally more secure than DAC.
- Parameters are defined by the Orange book B-level.
- Criticality of security levels, such as military installations, define MAC use.
- MAC systems are hard to program, configure, and implement.

Remember that MAC relies on the system to control access. For example, if a file is classified as confidential, MAC will prevent anyone from writing secret or top secret information into that file. Also, all output, such as print jobs, floppies, and other magnetic media must be labeled with its sensitivity level.

Lattice-Based Access Control

A lattice is a partially ordered set in which all finite subsets have the least upper boundary and greatest lower boundary. We can apply this same upper and lower boundary function to access control.

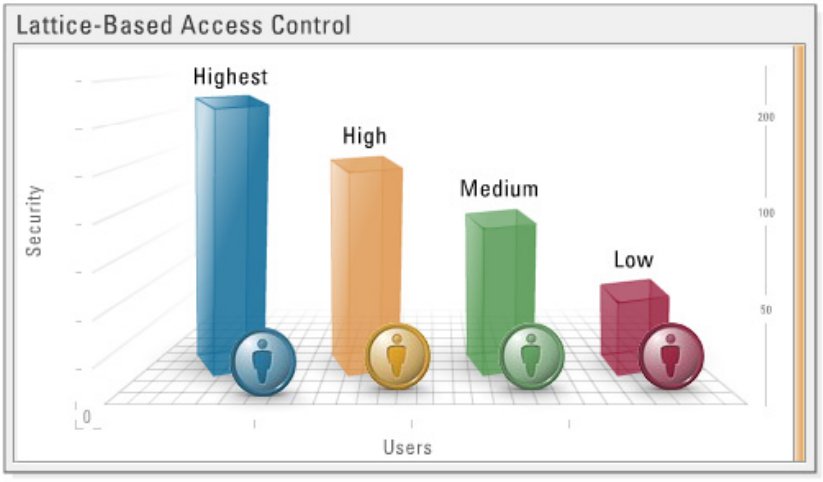

Lattice controls use an upper and lower boundary for access control

To apply this concept to access control, a pair of elements is the subject and the object. The subject has the greatest lower boundary and the least upper boundary of access rights to an object. This approach allows the administrator to combine objects from different security classes and determine the appropriate classification for the result by showing that any combination of security objects must maintain the lattice rule between objects.

Example: A <= A, If A <= B and B <= C, then A <= C.

The example shows that lattice-based access control is a way of setting upper and lower boundaries on what a subject can and cannot access.

Rule-Based Access Control

Rule-based access control is a type of MAC because the access to data is determined by rules or the use of classification labels, and not by the identity of the subjects and objects alone.

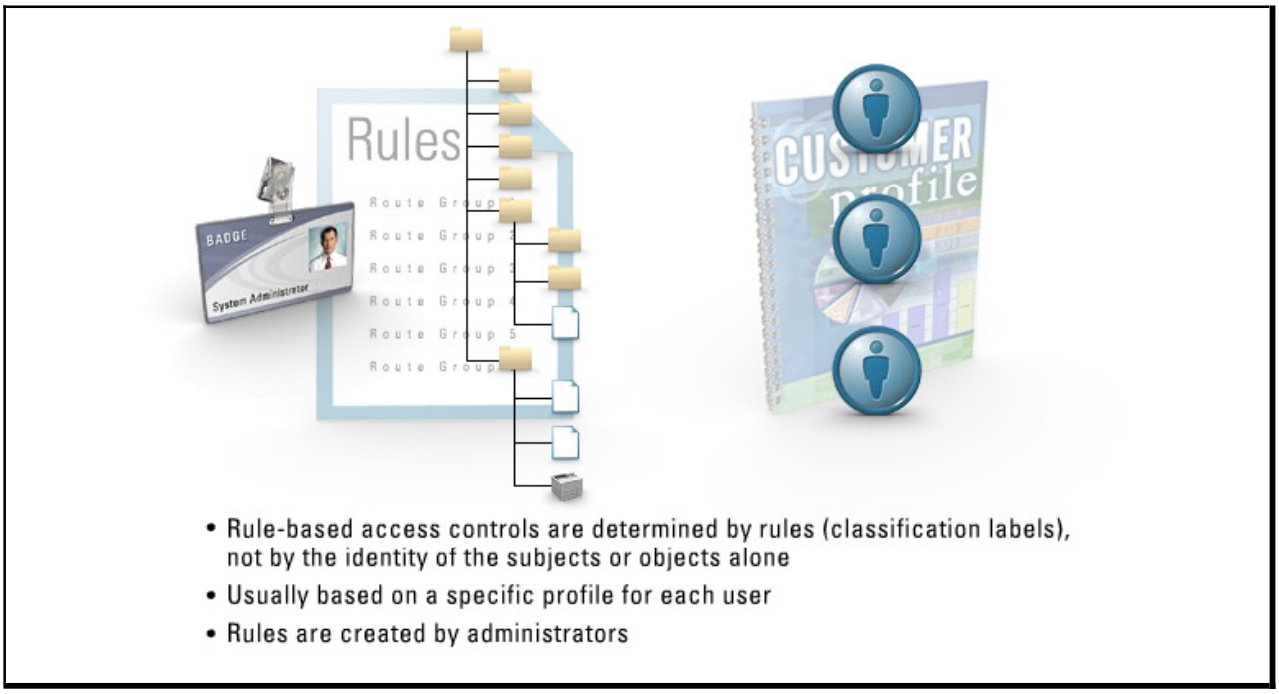

Rule-based access control is usually based on a specific profile for each user, allowing information to be easily changed for only one user. In essence, specific rules created by administrators indicate what can and cannot happen to an object.

Role-Based Access Control

In a **role-based access control** (RBAC) model, access decisions are based on the roles that individual users have as part of an organization.

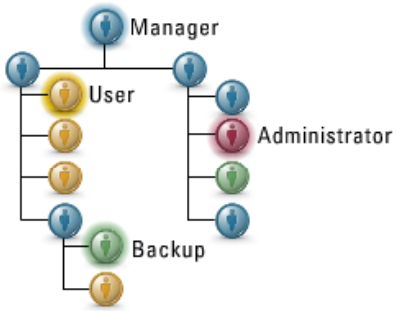

Organizational Structure

- Role-based access (RBAC) controls are based on roles individuals have in the organization
- Rights are grouped by role name (administrator, power user, backup, user, etc.)
- RBAC corresponds to an organization's structure
- RBAC is a type of MAC
- Also called nondiscretionary access control

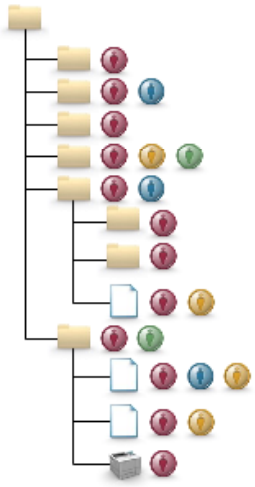

File Structure

Access rights are grouped by role name, and the use of resources is restricted to individuals authorized to assume the associated role. This model allows security to be managed at a level that corresponds closely to the organization's structure. Users with similar jobs are pooled into logical groups for the purposes of controlling access, and access is provided according to business requirements. RBAC is a type of MAC, and it is also commonly called **nondiscretionary access control**.

Security administration with RBAC consists of determining the operations that must be executed by persons in particular jobs and assigning employees to the proper roles.

Restricted Interfaces

The restricted interface model restricts users' access abilities by not allowing them to request certain functions, information or by not allowing access to specific system resources.

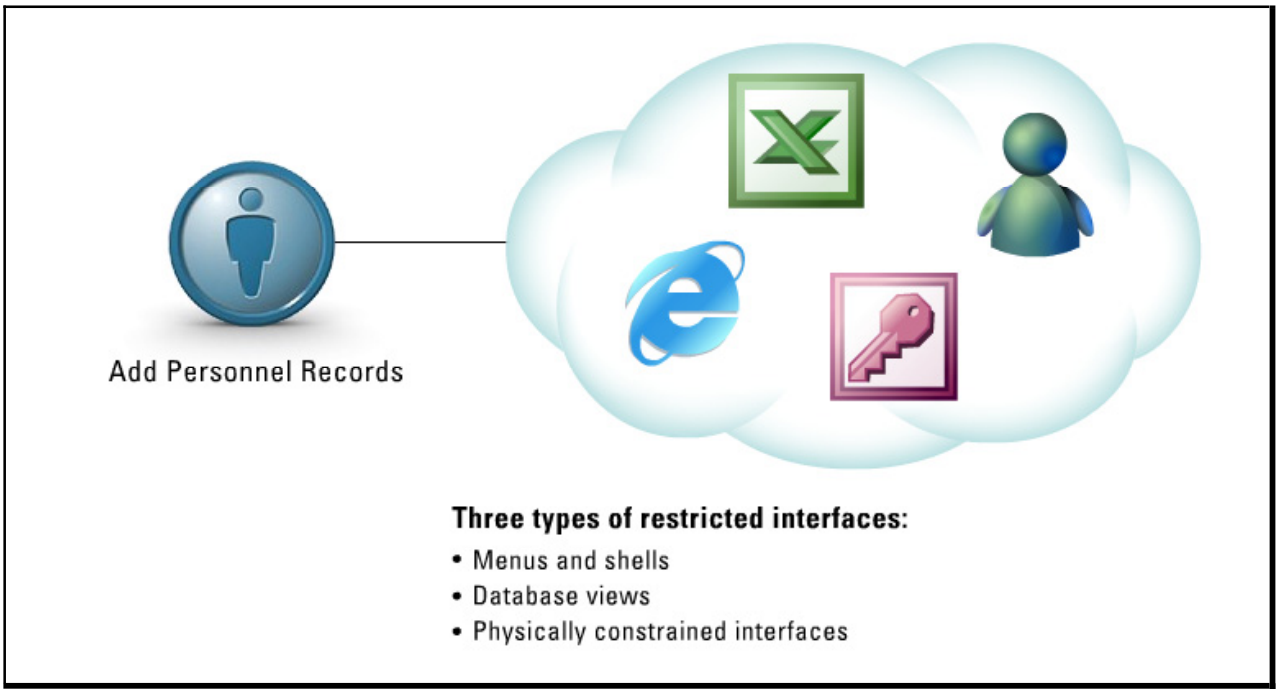

Three types of restricted interfaces are common:

- **Menus and shells** - Users are only given the options of the commands they can execute.
- **Database views** - User access to data is restricted by mechanisms.
- **Physically constrained interfaces** - User access is limited by providing certain keys on a keypad or touch buttons on a screen.

Non-Discretionary Access Control

This topic differentiates the discretionary access control (DAC) model and the non-discretionary access control model.

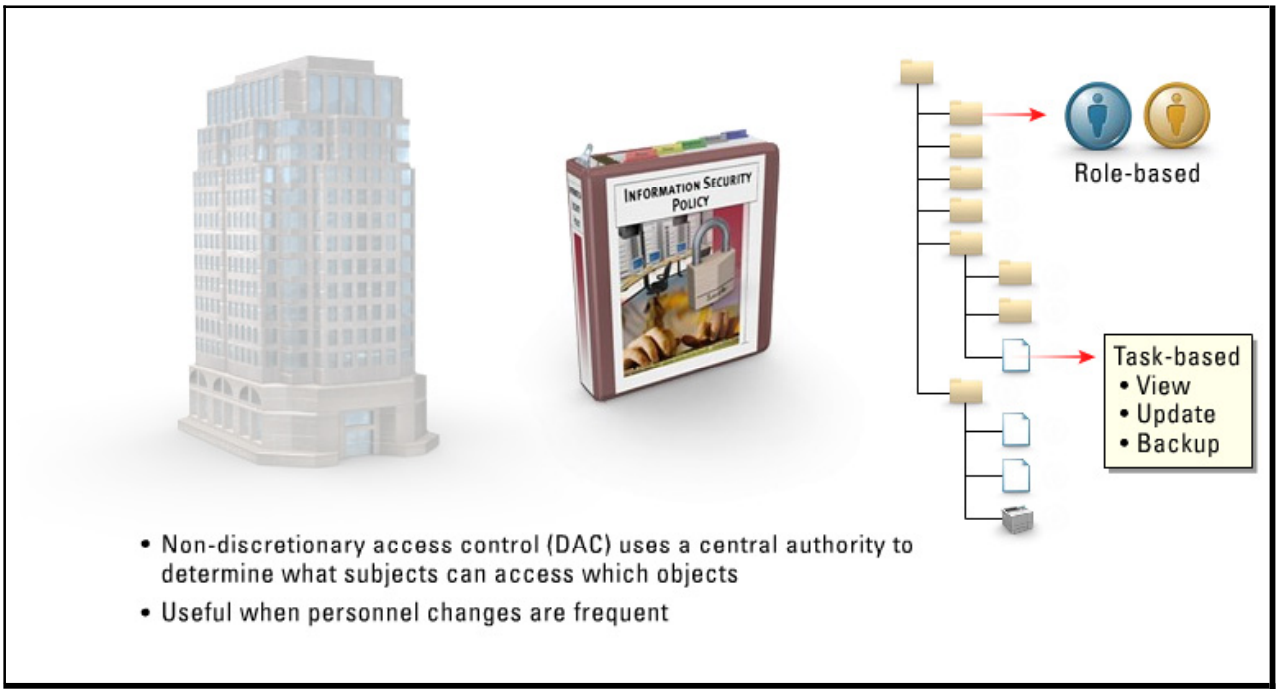

In the **discretionary access control (DAC)** model, the users' unique type of resource access to directories and files limits access. For example, Monica can be allowed to both read and change a file, while Ed can be restricted to only reading the file. In contrast, nondiscretionary access control implies that all users accessing a resource receive the same rights despite the share level set for the resource.

In the **non-discretionary access control** model, a central authority determines what subjects may access certain objects, based on the organizational security policy. These controls may be based on the individual's role (role-based) or the subject's responsibilities (task-based). This model is useful in organizations where frequent personnel changes are common.

Access Control Lists

Access control lists (ACLs) are a method of coordinating access to resources based on the listing of permitted or denied users, network addresses, or groups for each resource.

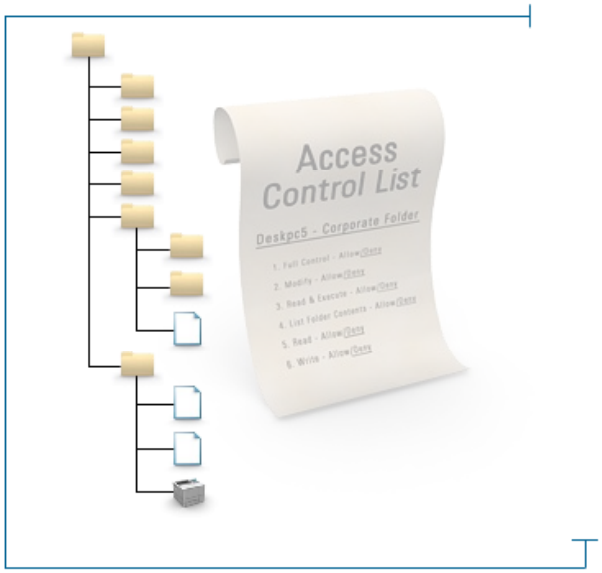

Access control lists permit or deny access based on permissions on a list

ACLs are basically a table of permissions dictating which subjects can access which objects, such as a file or directory.

Each object has a security attribute that identifies its access control list. The list has an entry for each system user with access privileges. The most common privileges include the ability to read a file or all the files in a directory, to write to the file or files, and to execute the file if it is an executable file, or program. Microsoft Windows NT/2000, Novell's NetWare, Digital's OpenVMS, and Unix-based systems are among the operating systems that use access control lists. The list is implemented differently by each operating system.

Basic types of access include read, write, create, execute, modify, delete, and rename.

Security Models

Security is best enforced with a multi-level security system. The best systems prevent users from reading information classified at a level for which they are not cleared.

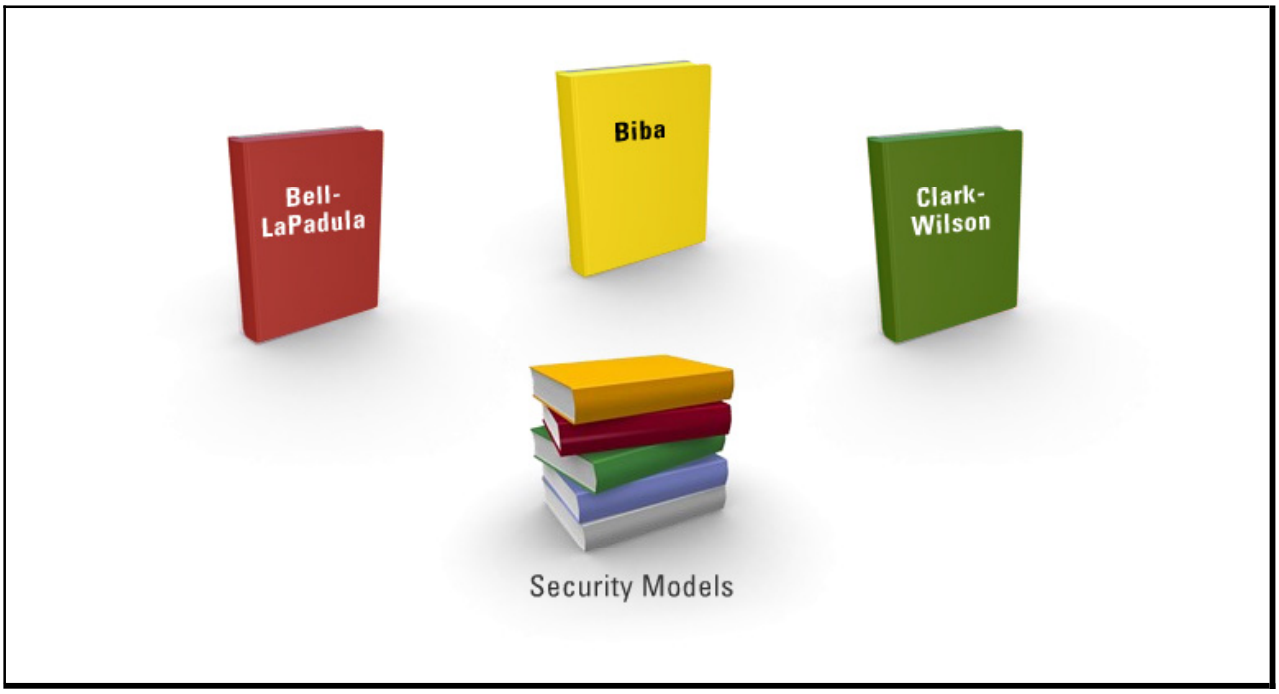

To assist in the creation of this multi-level security system, many security models are created to aid in the design and analysis of a secure computer system.

Security models include the following:

- Bell-LaPadula Model
- Biba Model
- Clark-Wilson Model

Bell-LaPadula

The Bell-LaPadula (BLP) model is built on the state machine concepts. Its main focus is on confidentiality. This concept defines a set of allowable states in a system.

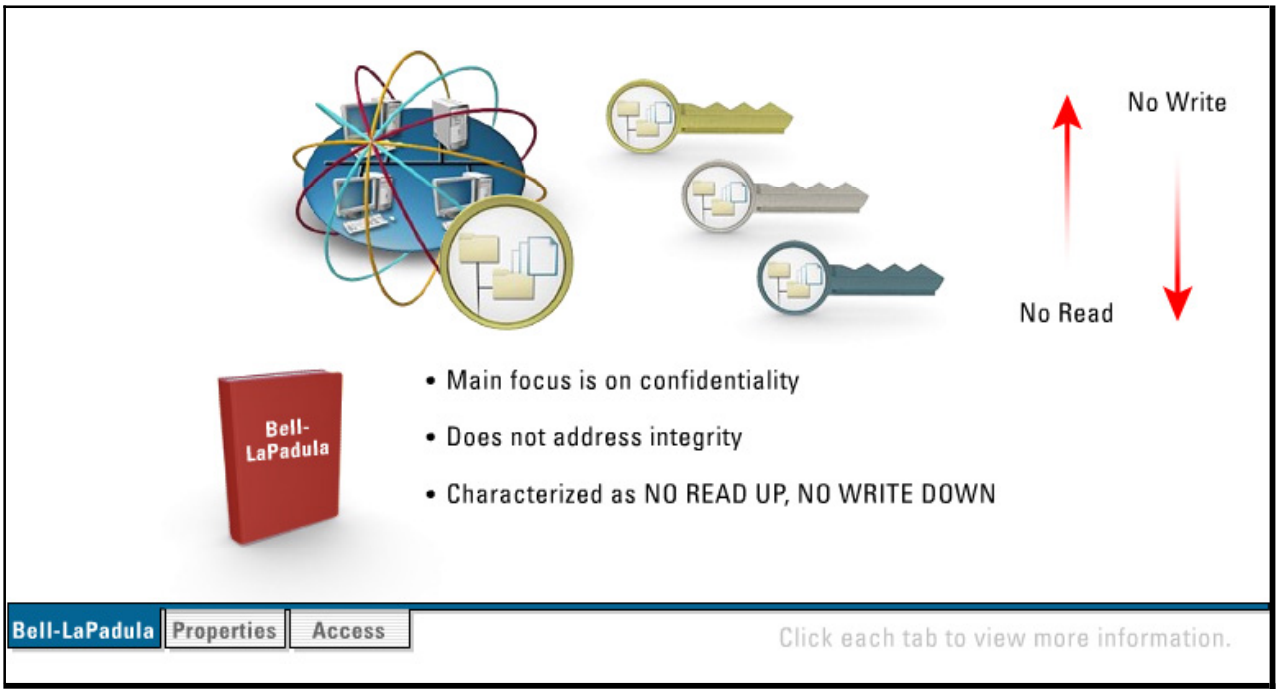

The transition from one state to another after receiving input is defined by transition functions. The objective of this model is to ensure that the initial state is secure, and that the transitions always result in a secure state. BLP defines a secure state through three multilevel properties:

- **Simple Security Property (SS)** - States that reading of information by a subject at a lower level from an object at a higher level is not permitted (no read up).
- *** property ("star property")** - States that writing of information by a subject at a higher level to an object at a lower level is not permitted (no write down).
- **Discretionary Security property (DS)** - Uses an access matrix to specify discretionary access controls.

This model prevents users and processes from reading above their security level. In addition, it prevents processes within any given classification from writing data associated with a lower classification. The "no write down" prevents placing data that is not sensitive, but that is contained in a sensitive document into a less sensitive file.

The BLP model addresses concerns about system security and leakage of classified information.

The model provides confidentiality and does not address integrity of the data that the system maintains.

Biba

The Biba model is latticed-based and uses the less than or equal to relationship.

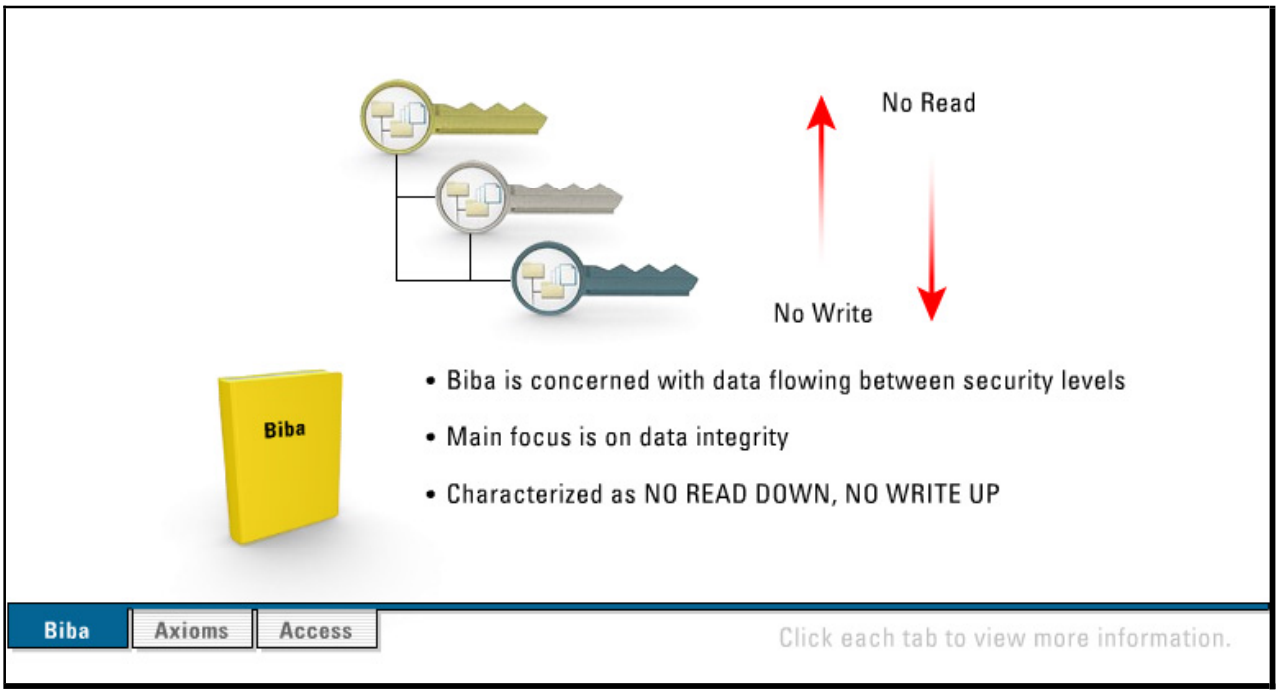

The Biba model addresses data flowing from one security level to another, and its main focus is on data integrity. Biba specifies the two following integrity axioms:

- **Simple Integrity Axiom** - States that a subject at one level of integrity is not permitted to observe (read) an object of a lower integrity (no read down).
- *** ("star")Integrity Axiom** - States that an object at one level of integrity is not permitted to modify (write to) an object of a higher level of integrity (no write up). For example, if a process can write above its security level, trustworthy data could be contaminated by the addition of less trustworthy data.

A subject at one level of integrity cannot invoke a subject at a higher level of integrity.

Because the Biba model uses a lattice of integrity levels, the various levels form the basis of expressing integrity policies that refer to the corruption of "clean" high-level entities by "dirty" low-level entities. Basically, Biba states that information can only flow downward. The simple integrity property and the * (star) integrity property ensure that clean subjects and objects cannot be contaminated by dirty information.

Clark-Wilson

The Clark-Wilson model emphasizes integrity, both internal and external consistency.

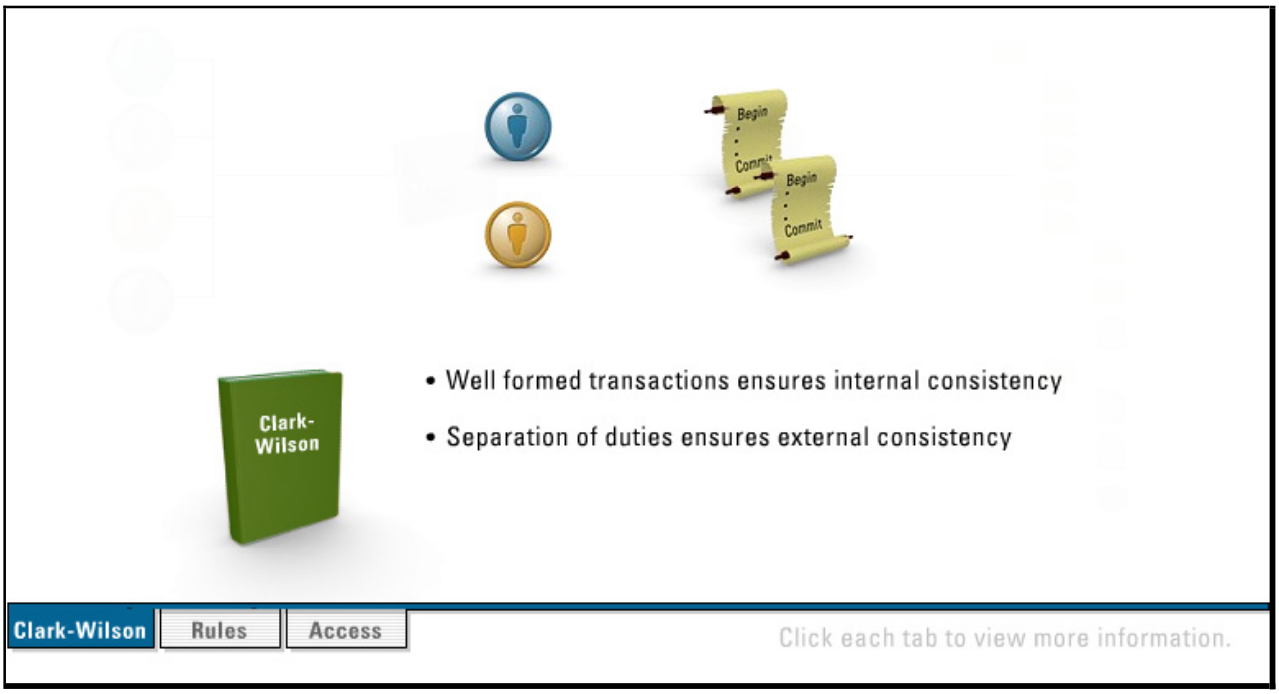

Clark-Wilson uses well-formed transactions, separation of duties, and the labeling of subjects and objects with programs to maintain integrity. Clark-Wilson identifies three rules of integrity:

- Unauthorized users should make no changes.
- The system should maintain internal and external consistency.
- Authorized users should make no unauthorized changes.

Internal consistency in the C-W model refers to the properties of the internal state of a system that can be enforced by the operating system, while external consistency is defined as the relationship of the internal state of a system to the real world to be enforced by means outside the operating system. These consistency checks ensure that…

- The process does what it is expected to do every time it is run.
- The data in the system is consistent with the value of similar data in the real world.

There are two mechanisms used to enforce integrity in the Clark-Wilson model:

- **Well-formed transactions** - Data and data processes can only be changed by a specific set of trusted programs. Users then have access to the programs and not the data directly.
- **Separation of duties** - Without separation of duties, users would need to collaborate to manipulate data or penetrate the system

Having well-formed transactions preserves and ensures internal consistency because users can only change the process and data in ways that ensure internal consistency. Having separation of duties ensures external consistency.

Non-Interference

The non-interference model relates to the information flow model by restricting information flow. Basically, this concept ensures that any action that takes place at a higher security level does not affect, or interfere, with actions that take place at a lower level.

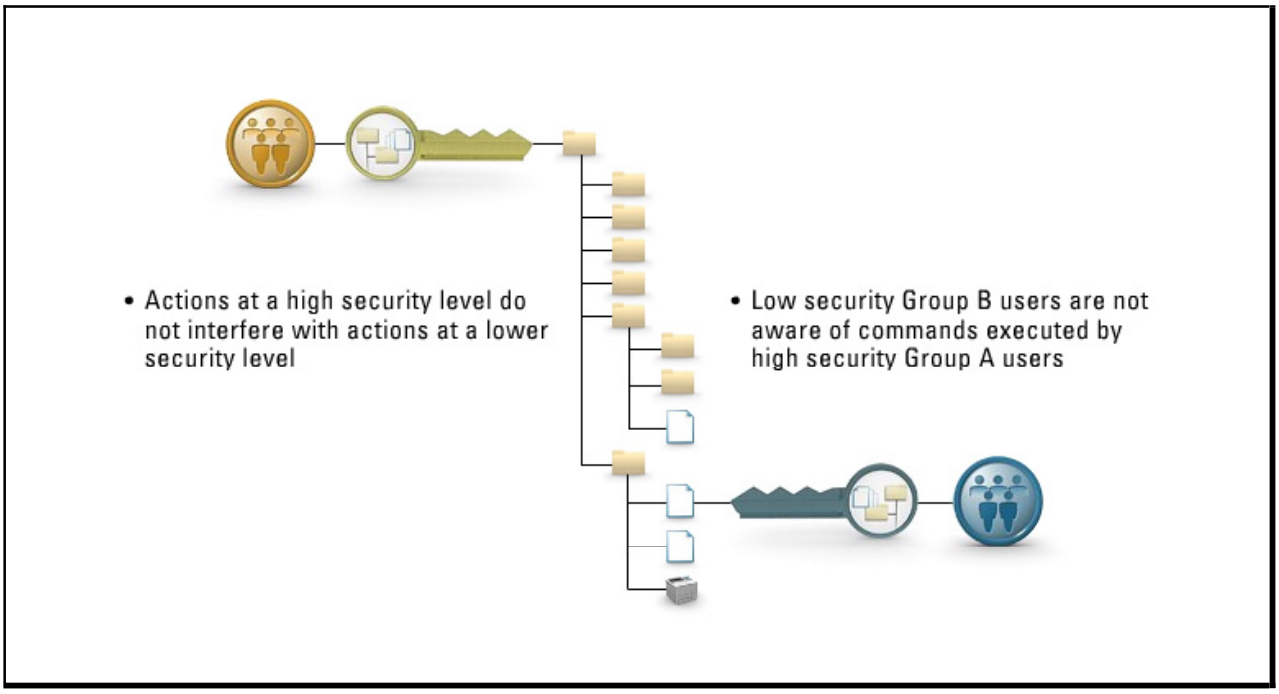

This model does not concern itself with the flow of data in the system, but rather concerns itself with what a subject knows about the state of the system.

The basic principle of this model ensures that a group of users (A), who are using the commands (C), do not interfere with the users in group (B), who are using the commands (D).

In this model, the lower level users in group (B) are not aware of the commands executed by users at a higher security level (A) and should not be affected by those commands in any way. If group (B) was aware of the commands issued by group (A), they might be able to deduce too much information about the activities of the higher security group.

State Machine

A state machine model captures the state of a system. A state can change only at discrete points in time by a clock or input event.

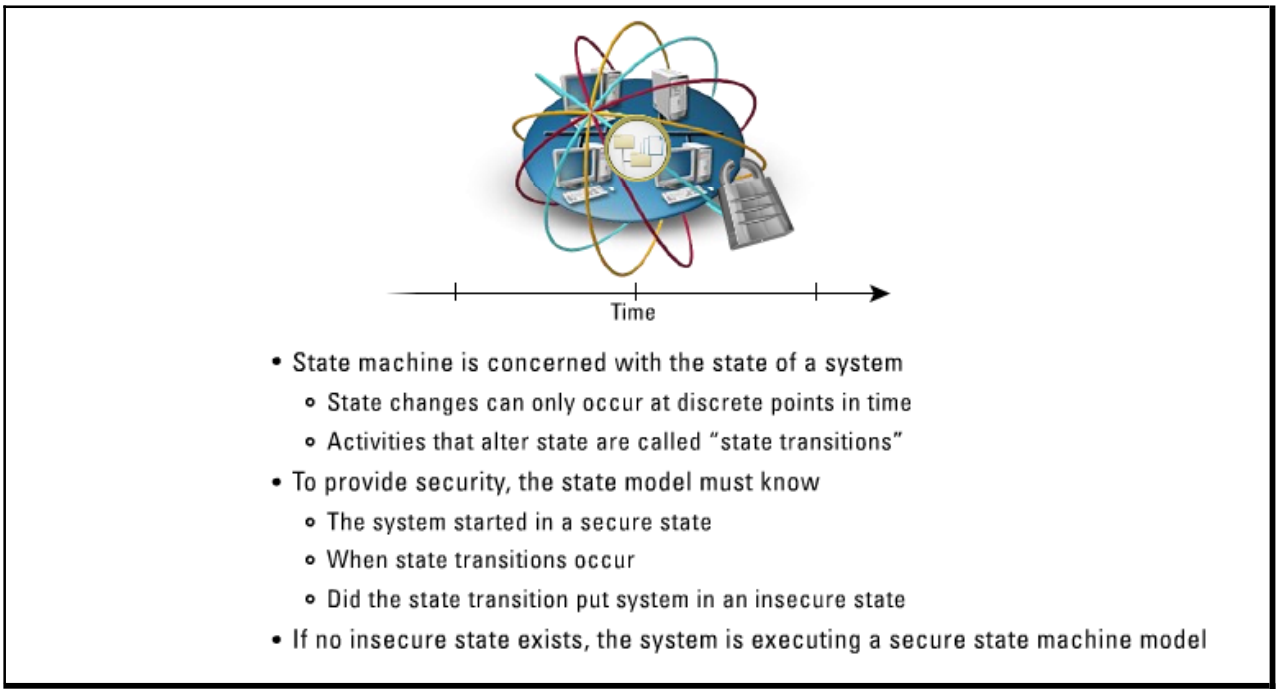

In order to verify the security state of a system, all current permissions and all current instances of subjects accessing objects must be captured. Basically, the state of the system is a snapshot of a system in one moment of time. Activities that alter this state are referred to as state transitions.

In order for the state machine model to provide security it must answer the following questions:

- Did the system start in a secure state?
- Did a state transition occur?
- Did the state transition put the system in an insecure state?
- If no insecure state exists, is the system executing a secure state machine model?

A system that is employing a state machine model will be in a secure state in each and every instance of its existence. It will boot up securely, execute commands and transactions securely, and allow subjects to access resources in secure states.

Access Matrix Model

The access matrix model is based on the concept of subjects and objects.

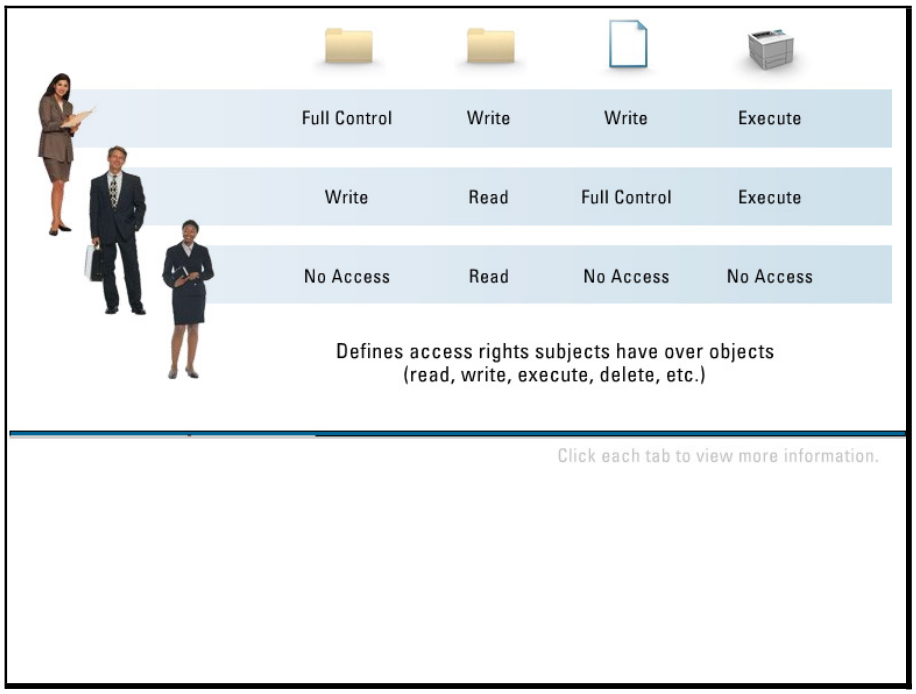

The subject is any entity, either a user or application, capable of accessing an object. An object is anything that is controlled, such as files, databases, and programs. The access matrix is used to define access rights and capabilities that subjects have over objects, such as read, write, execute, or delete.

The matrix consists of four major parts:

- A list of objects
- A list of subjects
- A function T that returns an objects type
- The matrix itself, with objects making the columns and the subjects making the rows

A sample access control matrix would look like the following:

	UserMonica Directory	UserEd Directory	UserSteph Directory	PrinterKoa
Monica	Full ControlWrite		Write	Execute
Ed	Write	Full ControlRead		Execute
Stephanie	No access	No Access	Full ControlExecute	
Jaycee	No access	Read	No access	No access

Information Flow Model

The **information flow model** is a variant of the access control model. In the access control model, the security state of the system is defined in terms of the security attributes of subjects and objects.

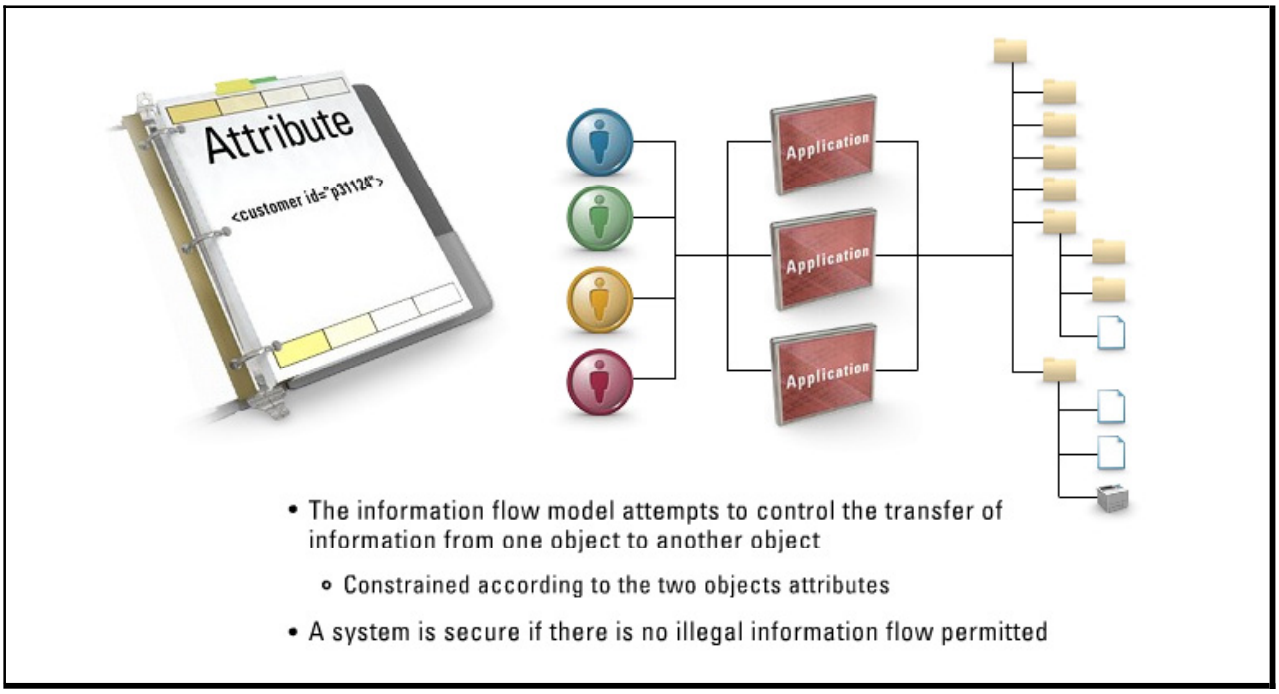

The access level a subject has on an object is determined by comparing their security attributes, rather than looking them up in a matrix.

In the information flow model, the control of the transfer of information from one object into another object is constrained according to the two objects' security attributes.

The information flow model can deal with any kind of information flow, not only the direction of the flow. It looks at insecure informational flow that can happen at the same level and between objects along with the flow between different levels. A system is secure if there is no illegal information flow permitted.

Rule of Least Privilege

The **rule of least privilege** is one of the most fundamental principles of INFOSEC and can be defined as a policy that limits both the system's users and processes to access only those resources necessary to perform assigned functions.

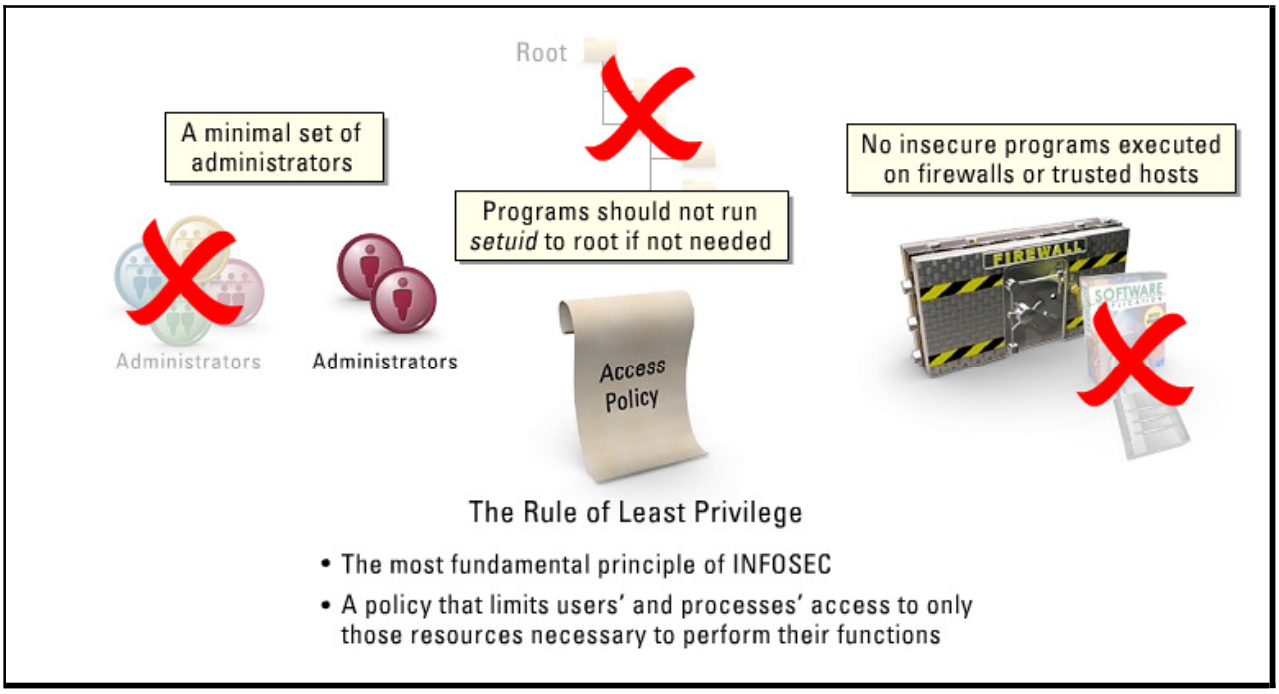

When an administrator applies the rule of least privilege, he or she identifies each user's job and outlines the minimum set of privileges required to perform that job. The administrator applies restrictions on the user to access only those resources required to accomplish the job – no more and no less.

Examples of implementing least privilege:

- Ensure that only a minimal set of users have root or administrator access.
- Make file group-writable to some group and make the program run setgid to that group, rather than setuid to root.
- Not running insecure programs on the firewall or other trusted hosts.

Separation of Duties

The objective of separation of duties is to ensure that a single person acting alone cannot compromise the company's security in any way.

Separation of duties ensures a single person cannot compromise a company's risk

Activities should be identified as high risk, medium risk, and low risk according to risk analysis. Next, high-risk activities should be broken up into different parts and distributed to different individuals. In this way, the company does not need to put a dangerously high level of trust on certain individuals, and if fraud were to take place, collusion would need to occur. This preventative measure means that more than one person would have to be involved in the fraudulent activity.

Rotation of Duties

Rotation of duties or job rotation is the act of having more than one person fulfilling the tasks of one position within the company.

Rotation of Duties

Enabling job rotation allows the company to have more than one person who understands the tasks and responsibilities of a specific job title, which provides personnel redundancy if a person leaves the company or is absent. Job rotation also helps when attempting to identify internal fraudulent activity.

Network Segregation

Network segregation can be accomplished through physical or technological means. The idea is to separate certain users, servers, or other resources from each other.

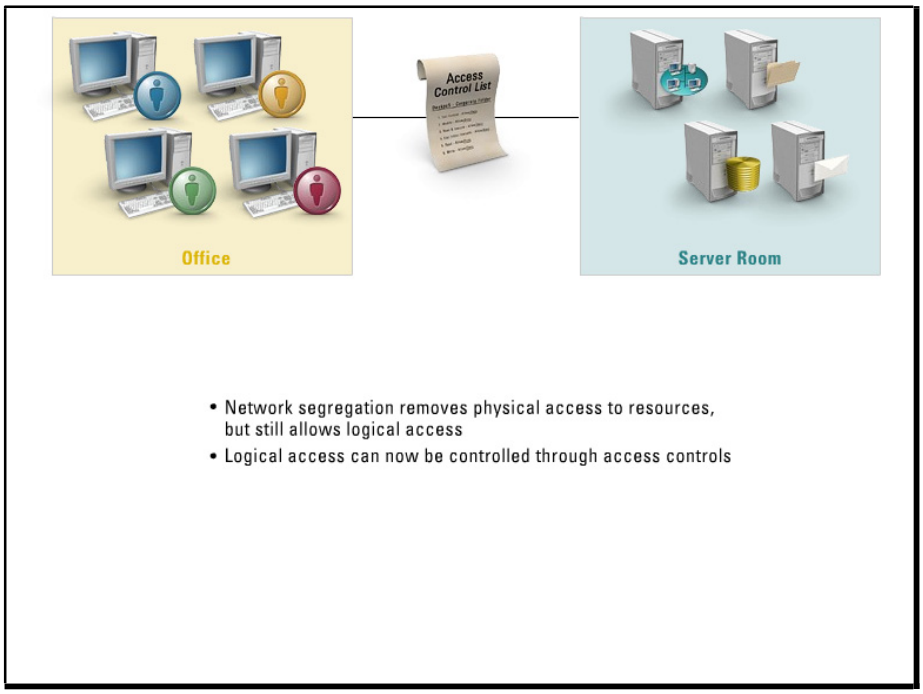

If access is required between any of the separated entities, a dematerialized zone exists to filter data. For example, by separating database, file, and HR servers into a single server room, physical access is removed from users. The system allows logical access to the servers through the network, but can now do so through a firewall that exists on the segment that separates users from servers. Now, with proper access controls the owners can identity and authorize only those users who require access to their servers.

Control Zone

A **control zone** is an emanation security feature used to defeat the capture of emanating electrical signals.

A control zone is used to defeat capture of emanating electrical signals

As an electrical device processes its instructions and stores data into memory or on the hard drive, electrical signals are emitted from the system. With the proper equipment and software, an attacker can see exactly what data is traversing the wires internally on the system. A control zone defeats this type of attack as the control zone creates a security perimeter that is constructed to protect against unauthorized access to data or the compromise of sensitive information.

Some companies use special materials in their facility walls or server cabinets to contain any stray electrical signals.

Summary

The key points discussed in this lesson are:

- Discretionary access control model
- Mandatory access control model
- Lattice-based access control model
- Rule-based access Control model definition
- Role-based access Control model
- Restricted interfaces security
- Non-Discretionary Access Control security
- Access control lists security
- Security models
- Principles of the Bell-LaPadula model
- Principles of the Biba model
- Principles of the Clark-Wilson model
- Principles of the Non-Interference model
- Principles of the State Machine model
- Principles of the Access Matrix model
- Principles of the Information Flow model
- Least privilege approach
- Separation of duties approach
- Rotation of duties approach
- Network segregation
- Control zones

Access Control Administration

Overview

Access control administration is critical, as it is the central point of authentication and authorization in an enterprise. Whether the access control mechanism is centralized or decentralized is dependent upon what the organization is trying to accomplish in its security goals. This lesson will identify the differing access control models as well as the most commonly used protocols.

Importance

Understanding the differences in access control administration is important to the information security professional, as its mechanisms and algorithms must be properly understood to effectively protect the enterprise.

Objectives

Upon completing this lesson, you will be able to:

- Explain centralized access control
- Define RADIUS
- Define TACACS+
- Explain decentralized access control

Outline

The lesson contains these topics:

- Centralized
- RADIUS
- TACACS+
- Decentralized

Centralized

In a centralized access control methodology, one entity either a department or individual is responsible for granting all users access to resources.

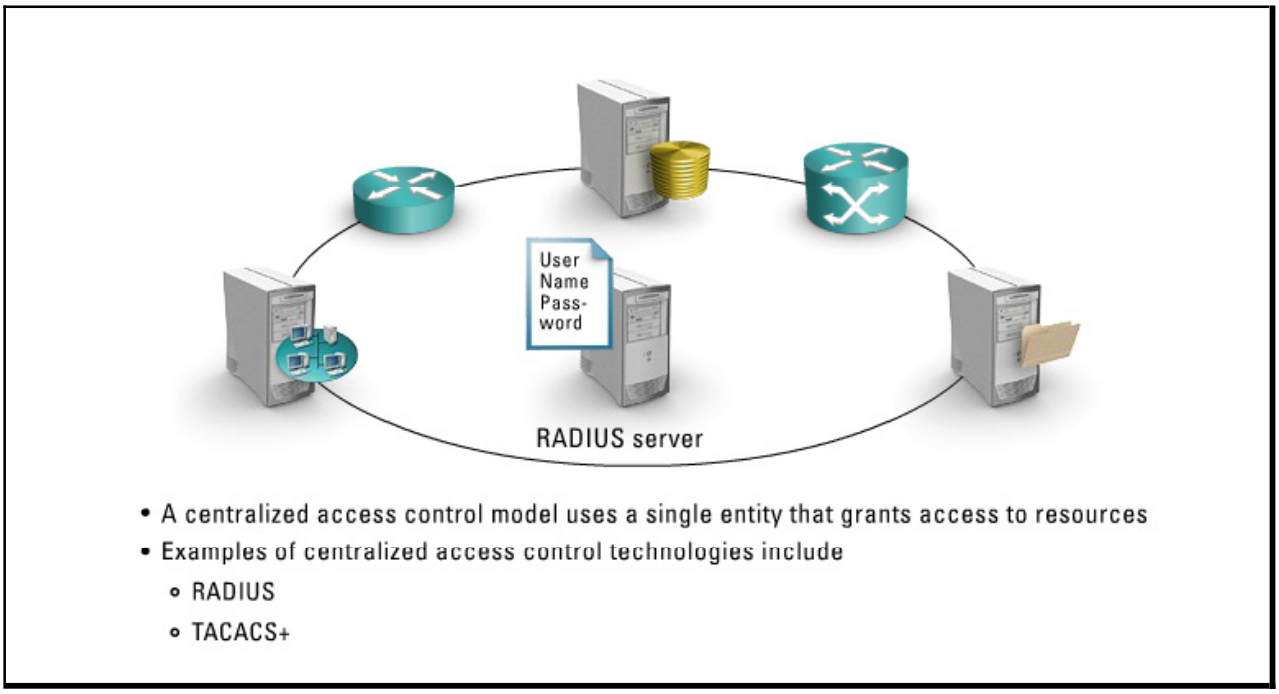

This control methodology provides two large benefits to the company:

- A consistent and uniform method of controlling users' access rights.
- A scalable solution where access control is centralized.

Examples on centralized access control technologies include:

RADIUS (Remote Authentication Dial-In User Service) - A client/server protocol and software that enables RAS to communicate with a central server to authenticate dial-in users and authorize their access to requested systems

TACACS+ (Terminal Access Controller Access Control System Plus) - An authentication protocol that allows a RAS to forward a user's logon credentials to an authentication server. TACACS is an unencrypted protocol, and therefore, less secure than the later TACACS+ and RADIUS protocols.

There are three generations of TACACS:

- **TACACS** - Combines authentication and authorization. Considered end-of-life.
- **XTACACS** - Separates authentication, authorization and accounting processes. Considered end-of-life.
- **TACACS+** - Separates authentication, authorization and accounting processes, with extended two-factor user authentication.

RADIUS

RADIUS (Remote Authentication Dial In User Server) is a protocol for carrying authentication, authorization, and accounting information between a Network Access Server (NAS), which desires to authenticate its links, and a shared Authentication Server. RADIUS was adopted as a standard protocol by the Internet Engineering Task Force.

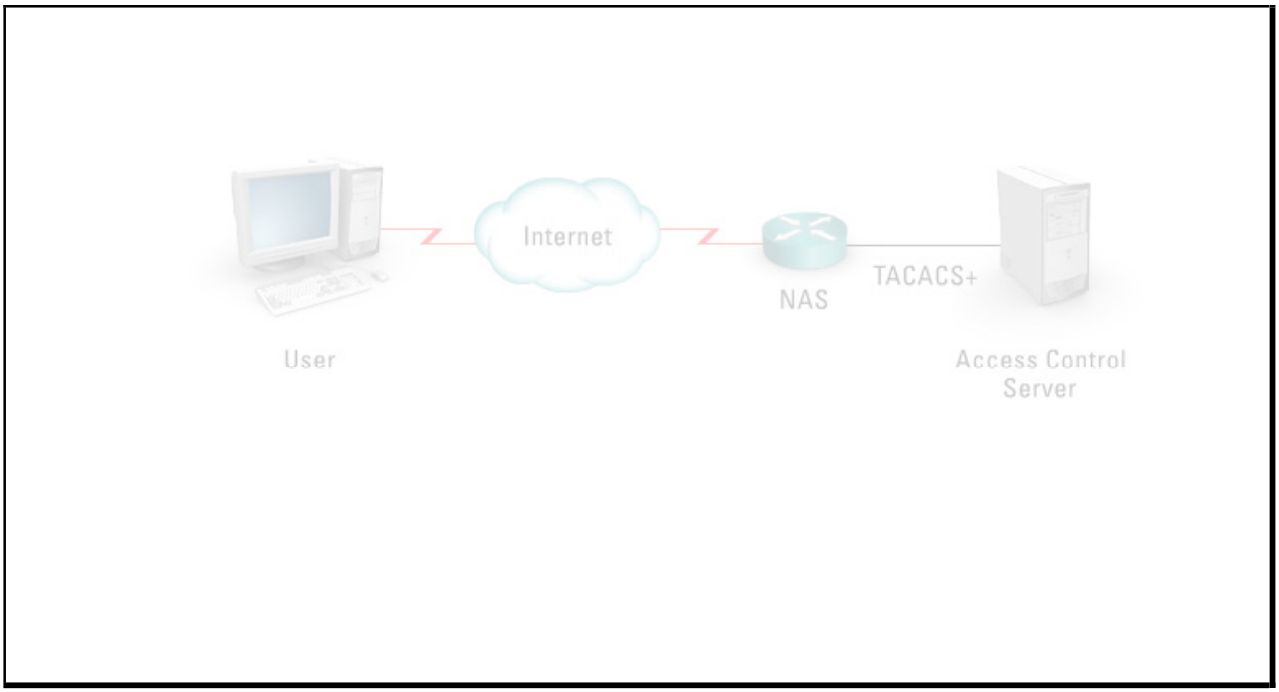

In this model, the client sends its authentication requests to a central RADIUS server that contains all of the users' authentication and network service access information, their network ACLs. A NAS operates as a client of RADIUS. RADIUS is a fully open protocol and is distributed in source code format. It can be modified to work with any security system that is currently available. It can also be used with TACACS+ and Kerberos and provides PAP or CHAP remote node authentication.

Features of RADIUS:

- Uses the Client/Server model.
- Transactions between the client and the RADIUS server are authenticated through the use of a shared secret, which is never sent over the network.
- Officially uses UDP ports 1812 (authentication) and 1813 (accounting).
- Earlier implementations used UDP ports 1645 and 1646.
- Encrypts only the password.
- Has very strong in accounting features.

RADIUS packets are identified by certain codes in the header:

- 1 - Access-Request
- 2 - Access-Accept
- 3 - Access-Reject

- 4 - Accounting-Request
- 5 - Accounting-Response
- 11 - Access-Challenge

TACACS+

TACACS+ is a client and server protocol for handling authentication, authorization, and accounting messages.

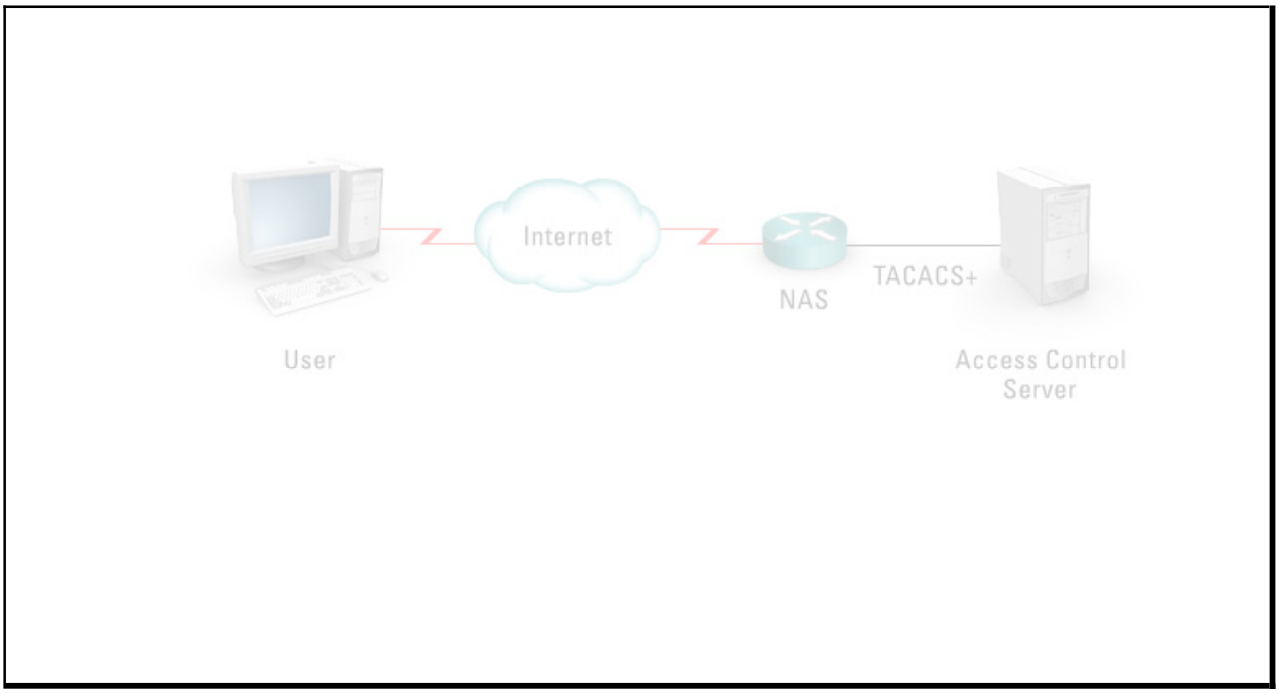

TACACS+ is the latest Cisco implementation. It provides attribute control or authorization and accounting. Authorization can be done on a per-user or per-group basis and is dynamic. User passwords are administered in a central database, which provides an easily scalable network security solution.

TACACS+ provides for separate and modular authentication, authorization, and accounting facilities. TACACS+ allows for a single access control server with the TACACS+ daemon to provide each service with authentication, authorization, and accounting independently. This single access control server means that each service can be tied into its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

The goal of TACACS+ is to provide a methodology for managing multiple network access points from a single management service.

TACACS+ has the following attributes:

- It uses a two-factor password authentication mechanism.
- The user has ability to change password.
- It uses TCP port 49.
- It encrypts entire payload.
- TACACS+ services are in the public domain and can be bundled in the OS of network devices.

After sending a request to the TACACS+ server for authentication, the NAS will receive one of the following responses:

- **Accept** - The user is authenticated and service may begin

- **Reject** - The user has failed authenticate
- **Error** - An error has occurred during authentication
- **Continue** - The user is prompted for additional authentication information

Decentralized

In a decentralized access control methodology, people closer to the resources control access. This methodology does not provide uniformity and fairness across the organization.

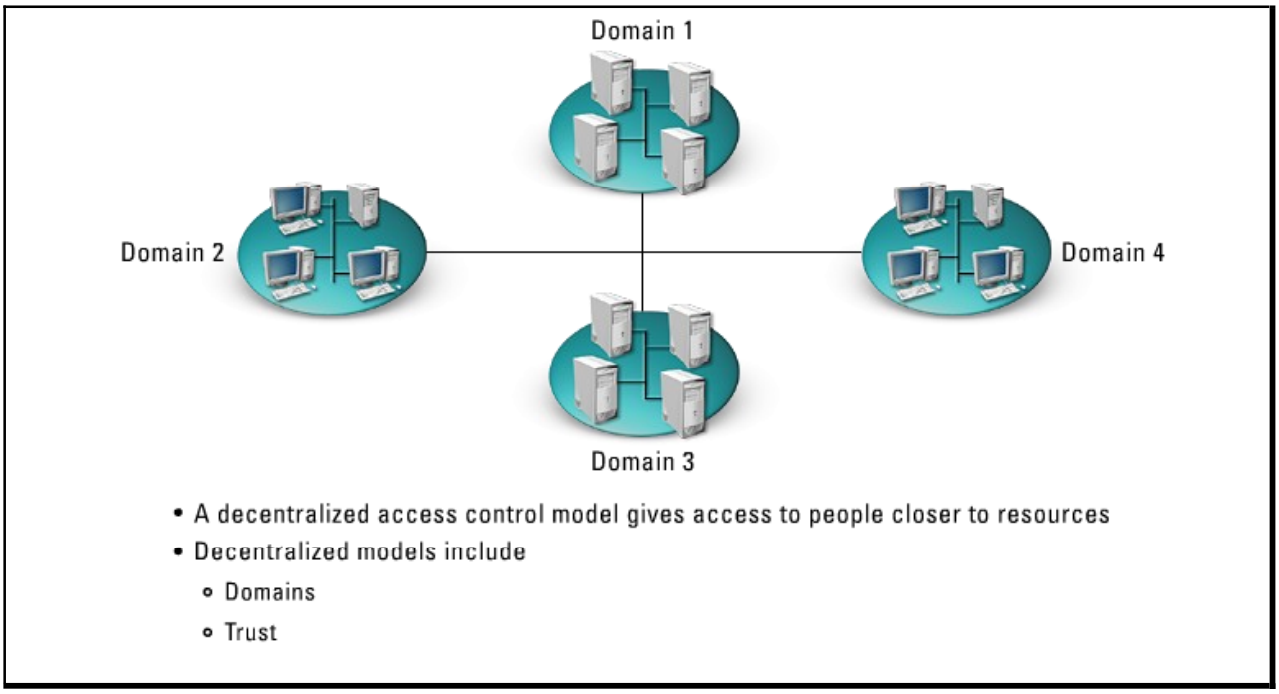

In this methodology of access control, the owners or creators of files maintain access control. This approach provides a greater flexibility for an individual administrator but comes as a tradeoff with less consistent implementation of access control policy. Examples include:

- **Domains** - A set of objects and subjects that have access rights for defined operations
- **Trust** - A trusted-computer system: all objects/subjects/operations OK

A security domain is a single domain of trust that shares a single security policy and a single management source. Usually, domains are defined as a sphere of influence. Domains are based on trust, but trust relationships sometimes can be compromised if proper care is not taken. Each security domain is different because different policies and management govern it. This approach makes scaling the organization very difficult.

Security domains are used within operating systems and applications to ensure that rogue activities do not accidentally damage important system files or processes. Segmenting memory spaces and addresses protect a security level. A security domain can also be described as the resources available to a user.

Summary

The key points discussed in this lesson are:

- Centralized access control
- RADIUS
- TACACS+
- Decentralized access control

Monitoring and Intrusion Detection

Overview

Attackers, if given enough time, system knowledge, and money will gain access to an enterprise's most valuable resources. Being able to identify the perpetrator before they can gain access is crucial. Monitoring your network for intrusions becomes vital for good security. Being able to identify the crumbs left by an attacker after they perform their reconnaissance or during the initial attack is vital in assuring that valuable resources are not stolen. If resources are stolen, the information left behind can identify what was stolen, when it was stolen, how it was stolen, and hopefully who stole it.

Importance

It is important that all information security professionals understand how data can be unlawfully obtained and how to counteract the measures the attacker uses to obtain the resource.

Objectives

Upon completing this lesson, you will be able to:

- Name the types of intrusions
- Define TEMPEST
- Name methods of intrusion prevention
- Name ways to identify intruders
- Name methods of intrusion detection
- Define NIDS and HIDS
- Name the types of data extraction
- Explain intrusion recognition
- Explain the use of traffic in an attack
- Explain the use of HoneyPots and HoneyNests
- Explain how a signature is used in attack identification

- Explain the use of intrusion reactive responses, alarms, and signals
- Explain the use of audit trails
- Explain the use of violation reports
- Explain the use of penetration testing

Outline

The lesson contains these topics:

- Types of Intrusions
- TEMPEST
- Intrusion Prevention
- Identification
- Intrusion Detection
- NIDS and HIDS
- Data Extraction
- Recognition
- Traffic
- HoneyPots and HoneyNests
- Attack Signature Identification
- Intrusion Reactive Response, Alarms, and Signals
- Audit Trails
- Violation Reports
- Penetration Testing

Types of Intrusions

Intrusion detection is the attempt of a system to identity and isolate computer and network attacks by observing system logs, audit data, or traffic on the wire.

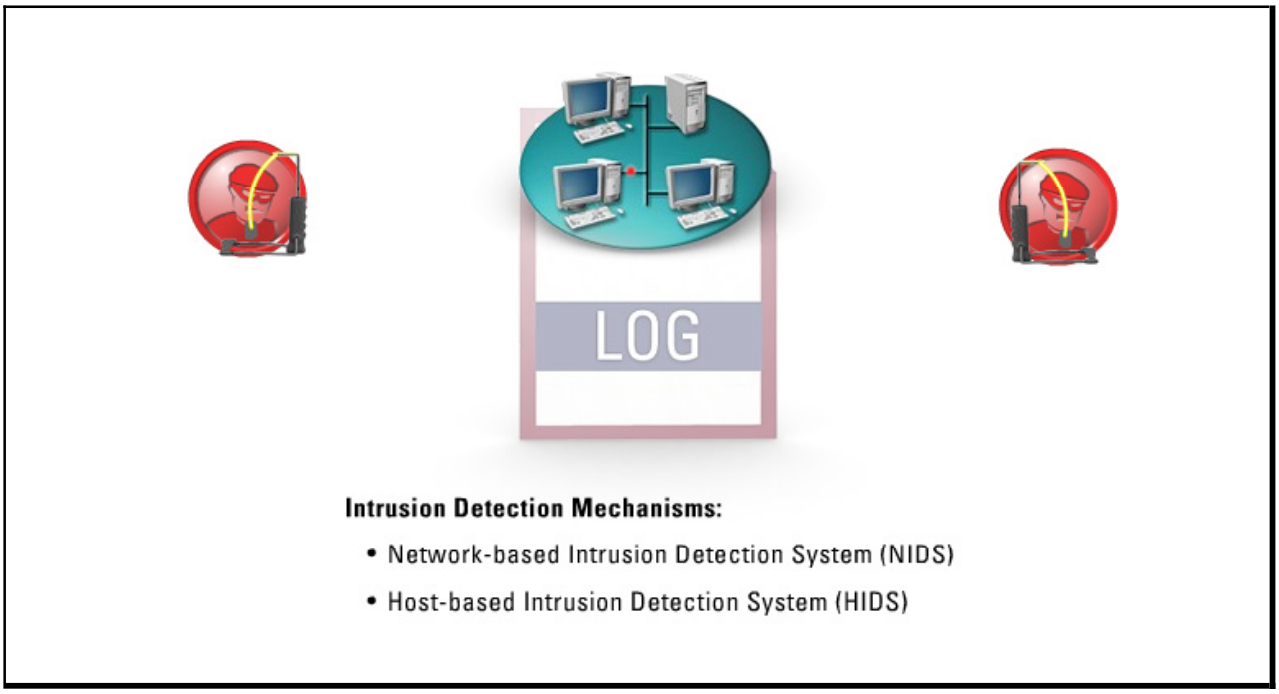

Audit data refers to audit records of all activities on a system and is the most common data for an Intrusion Detection System (IDS) to analyze. An IDS has three basic components: a sensor (agent), an analyzer, and a security interface (also called the director). The sensor collects information and forwards it to the analyzer. The analyzer receives this data and attempts to ascertain if the data constitutes and attack or intrusion. The security interface, which is usually a separate device, displays the output to the security administrator configures the sensors in the network.

There are two basic types of intrusion detection mechanisms:

- Network-based Intrusion Detection Systems (NIDS)
- Host-based Intrusion Detection Systems (HIDS)

Intrusion detection devices attempt to identify any of the following types of intrusions:

- Input Validation Errors
- Buffer Overflow
- Boundary Conditions
- Access Validation Errors
- Exceptional Condition Handling Errors
- Environmental Errors
- Configuration Errors
- Race Conditions

In addition, attacks against IP, passwords, DOS or DDOS, man-in-the-middle, port redirection, viruses, and Trojan horses can also be detected.

TEMPEST

TEMPEST stands for Telecommunications Electronics Material Protected from Emanating Spurious Transmissions.

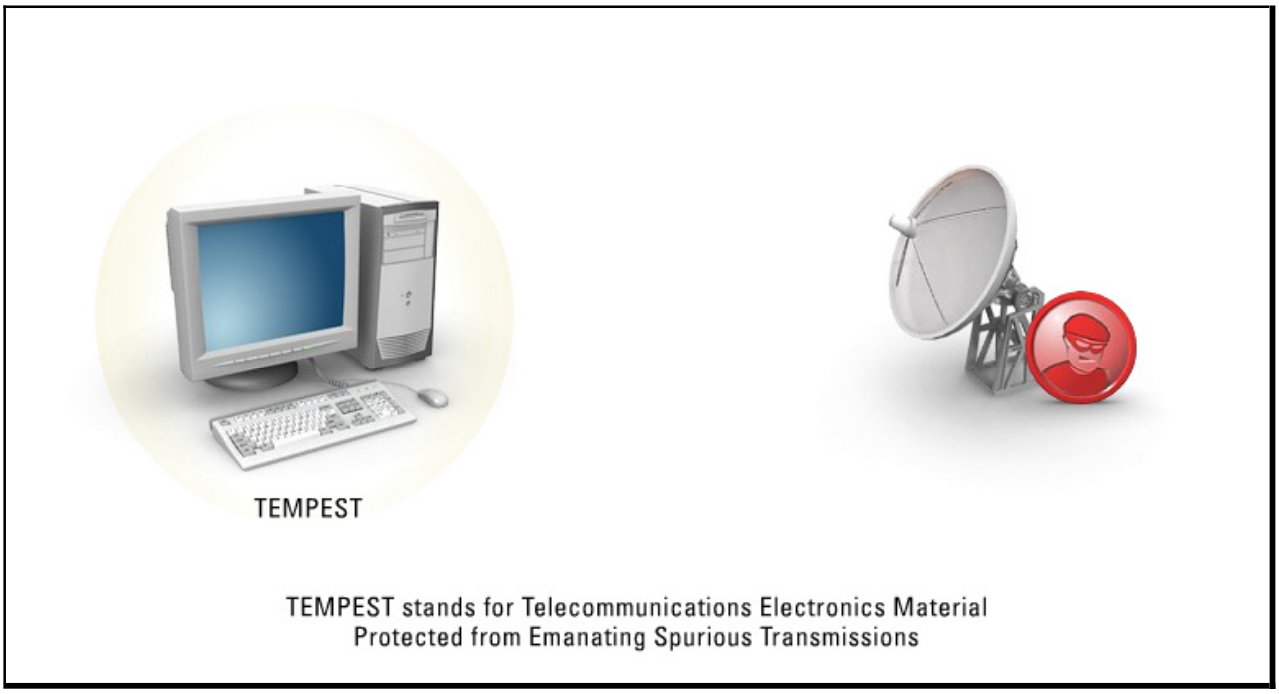

Emanations are any electronic signals radiating from electronic devices. For example, monitors, CPUs, transmission media, keyboards, and power cables are sources of electronic signals. Attackers now have methods of intercepting and analyzing these emanations to obtain sensitive corporate information. With sophisticated equipment, data can be readable at a few hundred yards. Also, any type of RF (radio frequency) device is susceptible to emanation interception. To counteract and study these types of intrusions, the U.S. government created a secret project named TEMPEST (Telecommunications Electronics Material Protected from Emanating Spurious Transmissions).

TEMPEST certified equipment, which encases the hardware into a tight, metal construct, shields the computer system from electromagnetic emanations. TEMPEST hardware is extremely expensive and can only be serviced by certified technicians. Even though, many highly secure government and military installations provide whole rooms and buildings that are TEMPEST-certified.

Intrusion Prevention

Intrusion prevention is the step beyond intrusion detection. Intrusion prevention turns the intrusion detection systems developed for spotting attacks into more useful products that can stop intruders cold.

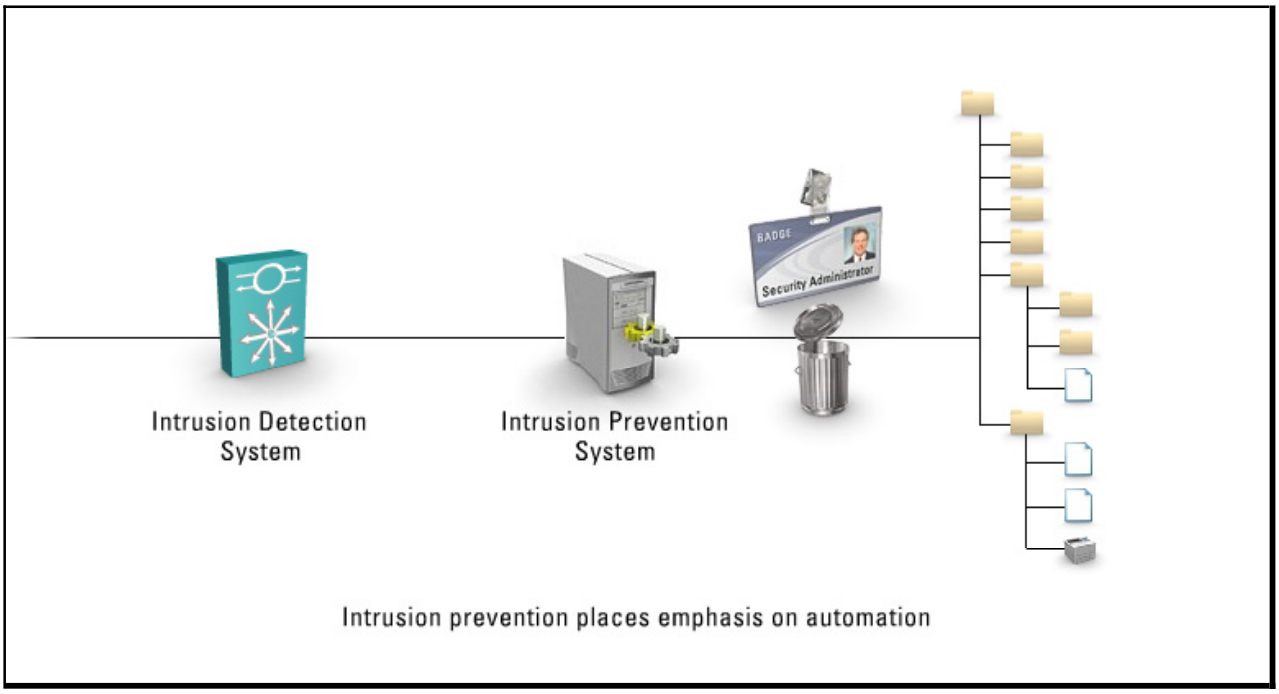

Intrusion prevention places emphasis on automation

The emphasis of intrusion prevention is on automation. Automatically determining with a vulnerability-assessment function that a network or particular system is vulnerable, the intrusion prevention system automatically blocks or shuns the packet, thwarting the attack. Basically, you are giving intelligence to intrusion detection. Detection has several steps:

- The intrusion detection system determines a packet (or packets) is malicious and sends the information to the intrusion prevention system.
- The intrusion prevention system analyzes the packet and determines many things, such as what type of attack it is, what process is it attacking, and who is the intended target.
- Next, the intrusion prevention system connects to the target and attempts to ascertain if it is vulnerable to the attack.
- If the system is not vulnerable to the attack, there is no need to contact the security administrator.
- If the system is vulnerable to the attack, the packet can be dropped and all further communication between the attacker and the target can be blocked. The system administrator is then contacted.
- Furthermore, the intrusion prevention system can identify and automatically install the updates or patches required to protect the target from further attack.

Identification

The ability to accurately identify any attack against a system is critical to any intrusion detection system.

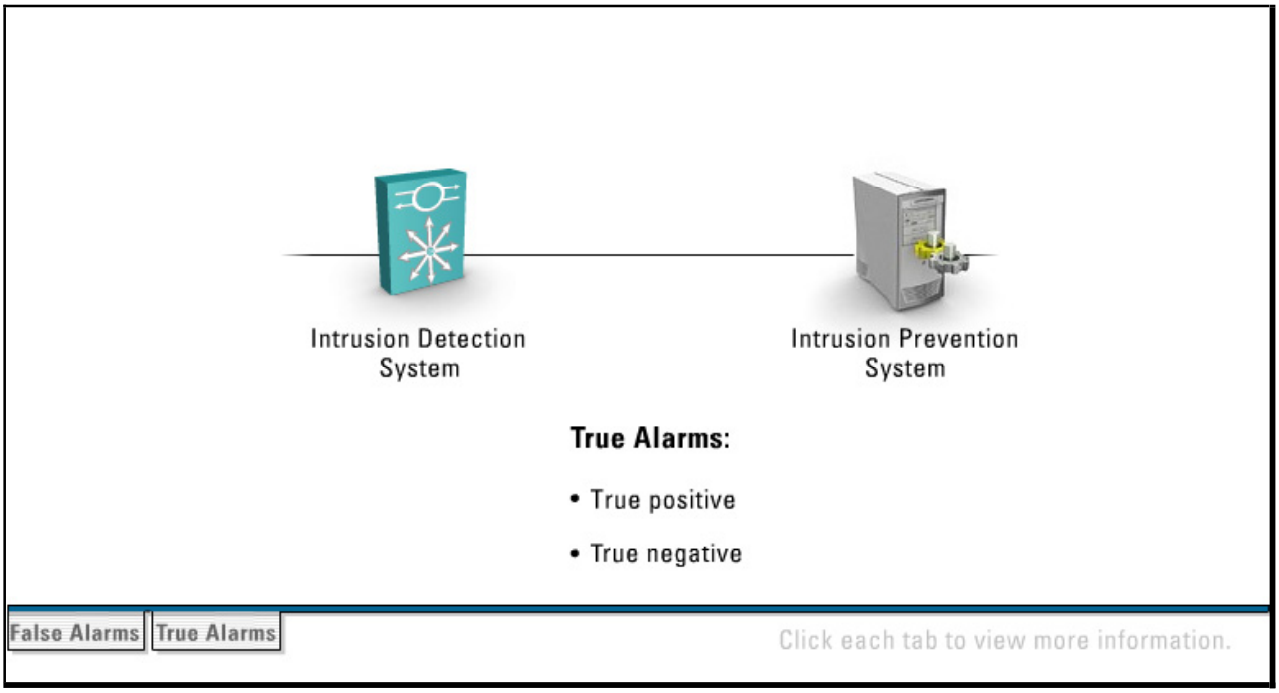

False alarms can cause the security administrator many headaches as well as many hours of unproductive and wasted time. Therefore, determining true alarms from false alarms is vital. Administrators find two forms of false alarms as well as two forms of true alarms.

False Alarms

- **False positive** - Normal or benign traffic that is inaccurately identified as an attack.
- **False negative** - Malicious traffic that should be identified as an attack and is not.

True Alarms

- **True positive** - Malicious traffic correctly identified as an attack.
- **True negative** - Normal or benign traffic is accurately identified as "OK"

Any intrusion detection system should have very high marks for detecting true alarms and very low marks on the false side.

Intrusion Detection

The actual IDS system is responsible for monitoring the network or system for attacks.

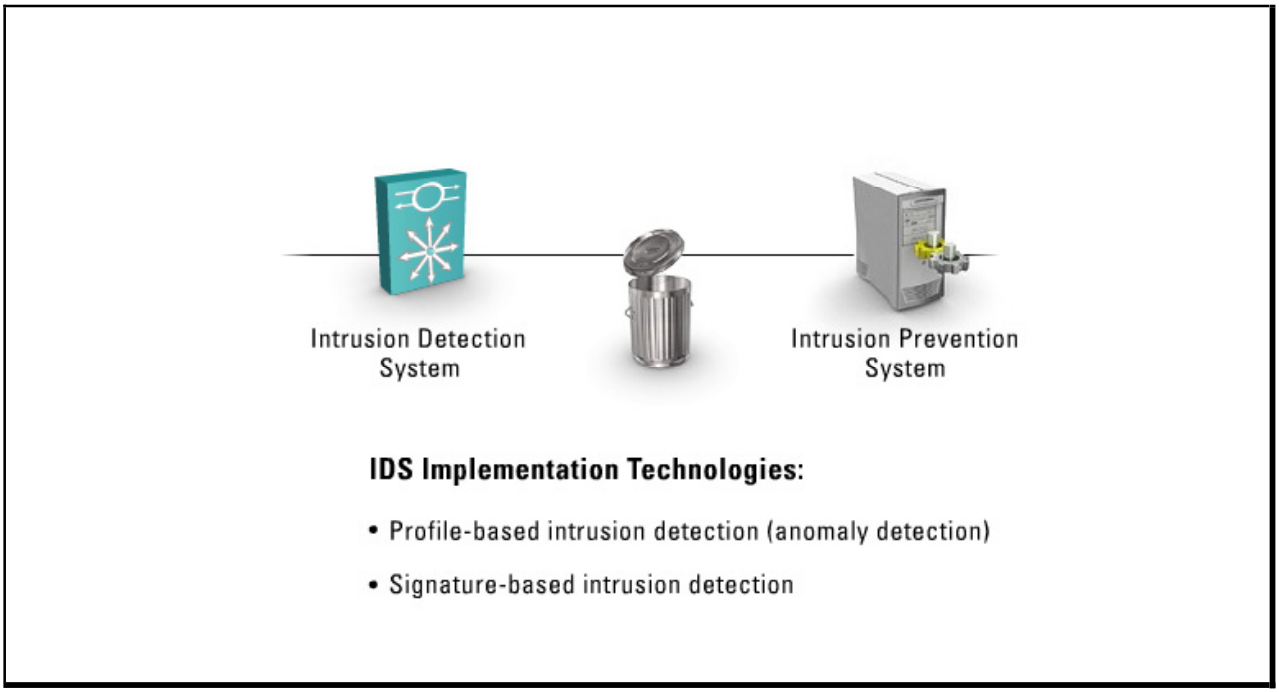

IDS has two types of implemented technologies:

- **Profile-based Intrusion Detection** - Also known as anomaly detection. In profile-based detection, an alarm is generated when activity on the network goes outside of the profile. A profile is a baseline of what should be considered normal traffic for each system running on the network. A problem exists because most systems do not follow a consistent profile. What is normal today, might not be normal tomorrow.
- **Signature-based Intrusion Detection** - In signature-based detection, a signature or set of rules is used to determine intrusion activity. An alarm is generated when a specific pattern of traffic is matched or a signature is triggered.

Typical responses to an attack include the following:

- Terminating the session (TCP resets)
- Block offending traffic (usually implemented with ACLs)
- Creating session log files
- Dropping the packet

NIDS and HIDS

Network-based intrusion detection (NIDS) and Host-based intrusion detection (HIDS) are complementary technologies.

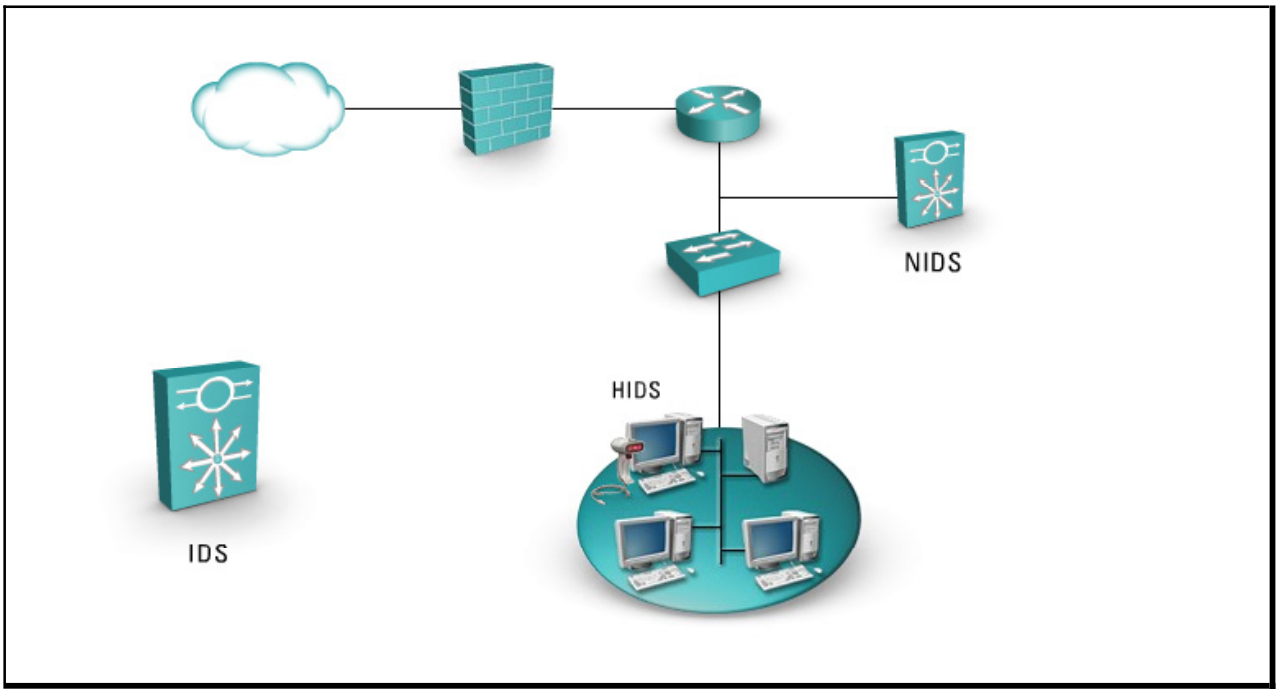

Both should be used to provide in-depth protection against attacks. A few features and issues with each technology include the following items:

NIDS

- Protects an entire network segment
- Is usually a passive device on the network and users are unaware of its existence
- Cannot detect malicious code in encrypted packets
- Is cost effective for mass protection
- Requires its own sensor for each network segment

HIDS

- Protects a single system.
- Uses system resources such as the CPU and memory from system.
- Provides application level security.
- Provides day-one security as a shunt between high and low level processes
- Intrusion detection is performed after decryption.
- Used on servers and sensitive workstations, but is costly for mass protection

Data Extraction

Extracting data from the network can be accomplished in two ways with network-based intrusion detection: passive extraction and active or inline extraction.

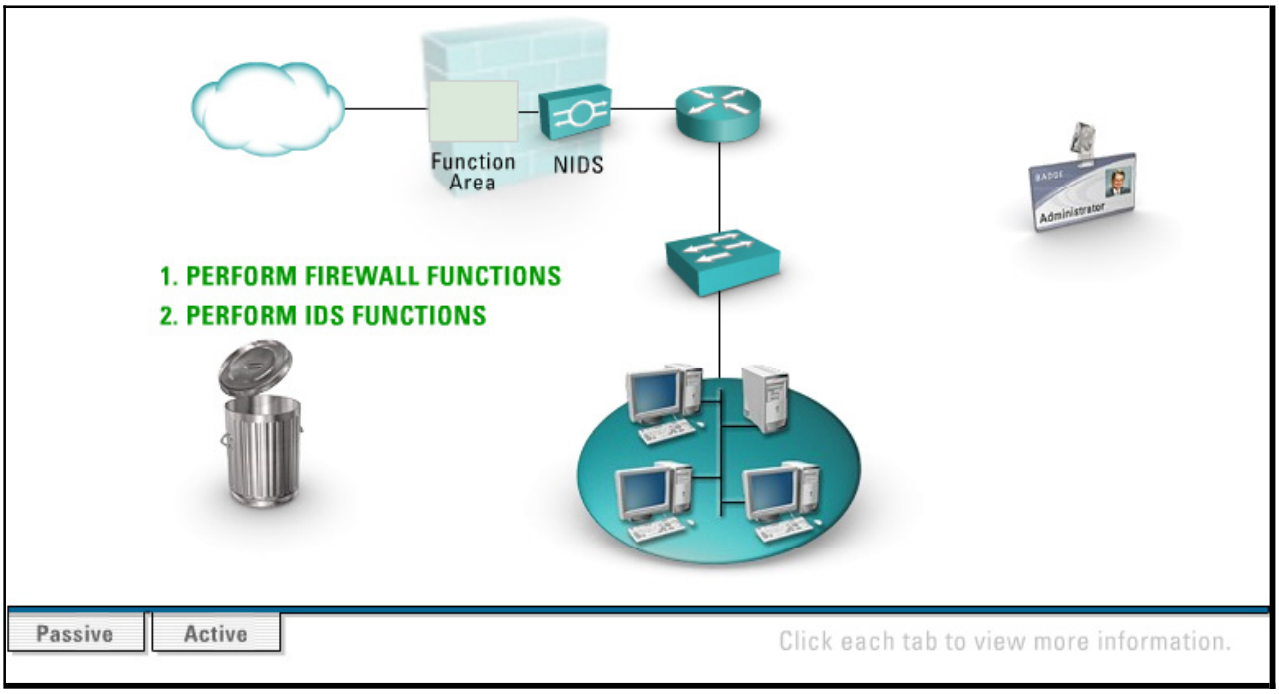

In the **passive mode**, a sensor or agent makes a copy and passively sniffs traffic on the wire. The sensor can afford to expend all its CPU and memory energies on determining if the packet is malicious or not. It does not perform any other function on the packet. Once the packet is deemed OK, it is flushed from memory. If the packet is deemed an attack, the sensor may act to protect the host and contact the director to let the security manager know of the attack. A sensor using the passive method of extraction cannot drop malicious packets, because the original packet would be already routed to the destination, while the copy is studied for intrusion.

Passive mode sensors are usually connected to switches for connectivity to the network. This arrangement can be a problem if not addressed because switches are intelligent and will only forward frames to the correct destination port, unlike hubs which forward a frame out all ports. To allow the sensor to receive frames not destined to it, the switch can be configured to copy frames from a port, ports, or VLAN to the port the sensor to which it is connected.

In the **active** or **inline extraction mode**, a device, which is usually a router or firewall, receives the packet, performs its security, QoS, routing, and other functions, and then turns the packet over to the IDS, which resides in the device. If the packet is not deemed a threat, it can be forwarded as normal. If the packet is deemed malicious, the device, since it is in the line of communication, has the ability to drop the packet. In this way, the packet will never reach its destination. The drawback with the active mode of intrusion detection is the fact that the router or firewall has a primary purpose, namely a router routes and a firewall firewalls. The IDS functionality detracts valuable CPU and memory resources from its primary purpose. For this reason, most active IDS devices only have a limited number of signatures used for

intrusion detection, usually around 100 or less, whereas the passive sensor can check over one thousand signatures.

Recognition

Signatures are the primary method used to recognize intrusions on the network.

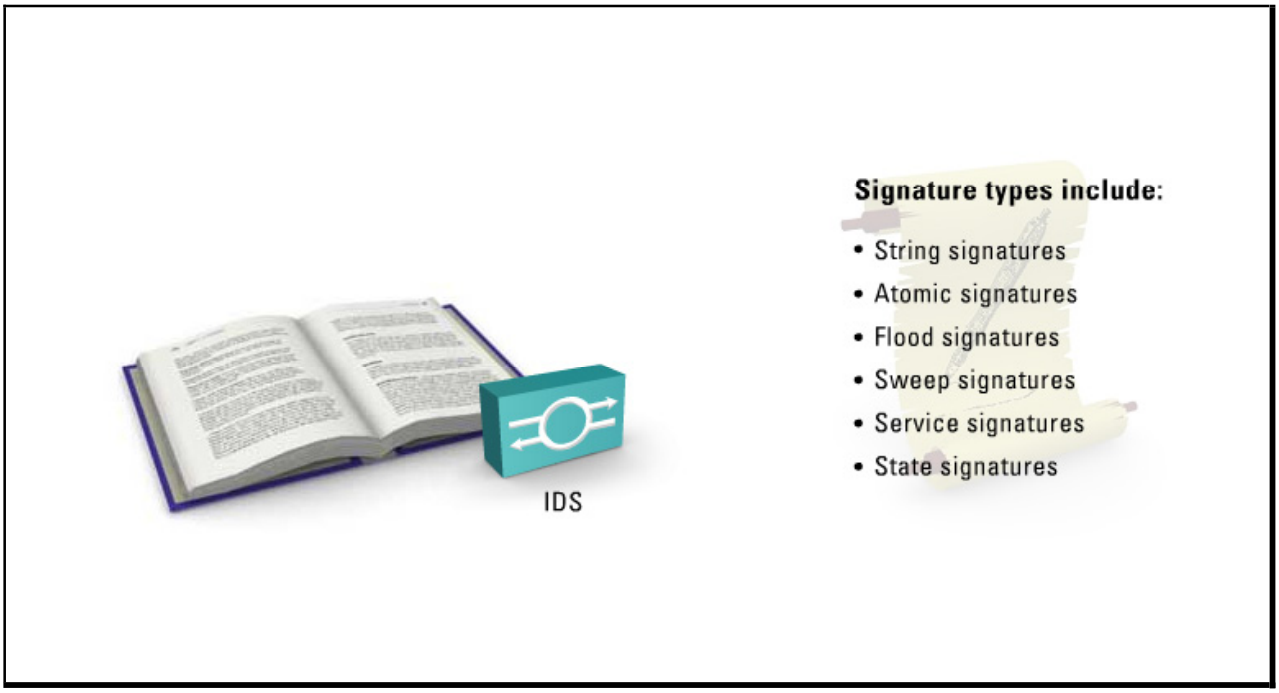

A signature is a set of rules used to detect typical intrusive activity. As sensor scans packets, the analyzer uses signatures to detect known attacks and respond with predefined actions. Administrators use many general signature recognition methods to identify different types of attacks; each method can identify the level of an attack, usually low, medium, or high: low is usually an informational attack (scanning for active IP addresses) and high is usually a malicious attack that can cause devastating results if successful (DoS). Typical recognition methods include the following:

- **String signatures** - Triggers on a particular string in a packet
- **Atomic signatures** - Triggers on a single packet condition
- **Flood signatures** - Triggers on detected DoS traffic
- **Sweep signatures** - Triggers on network reconnaissance attacks
- **Service signatures** - Triggers on layers 5, 6, and 7 attacks
- **State signatures** - Triggers on state-based attacks

Traffic

Sensors ratings are typically based on the amount of traffic they can analyze without dropping packets.

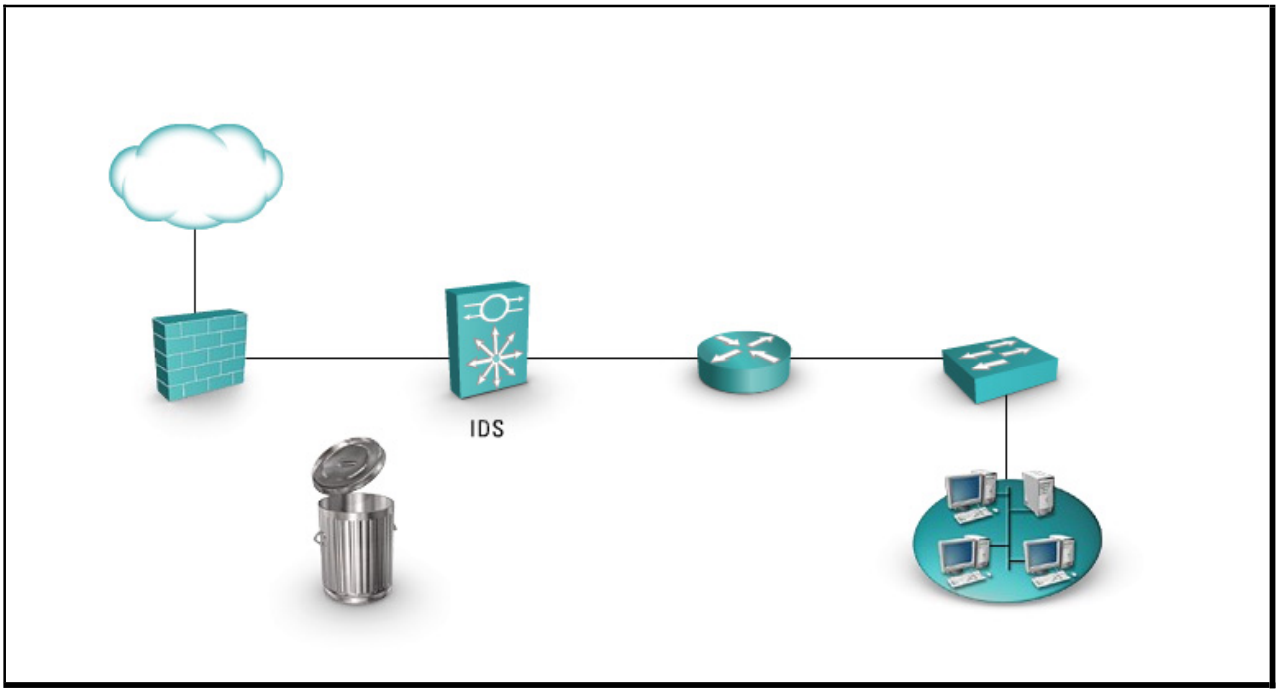

If the amount of traffic exceeds this rating, packets will be dropped by the sensor and will never be analyzed. Attackers are well aware of this limitation and will try to flood the network with large amounts of benign traffic while attempting to perform their malicious activity. All the while, they are hoping their attacks are dropped by the sensor.

Rate limiting counteracts this type of attack. When rate limiting is enabled, only a certain amount of traffic can enter the network, thereby assuring the sensor that it will receive no more traffic than it can safely analyze. This method, of course, has its own drawbacks, namely a waste of valuable bandwidth when the sensor can analyze less than the amount of traffic the network can accommodate.

In general, the best way is to obtain a sensor or sensors that can analyze traffic at the same speed the network injects traffic.

Honeypots and Honeynets

This topic will differentiate honeypots and honeynets.

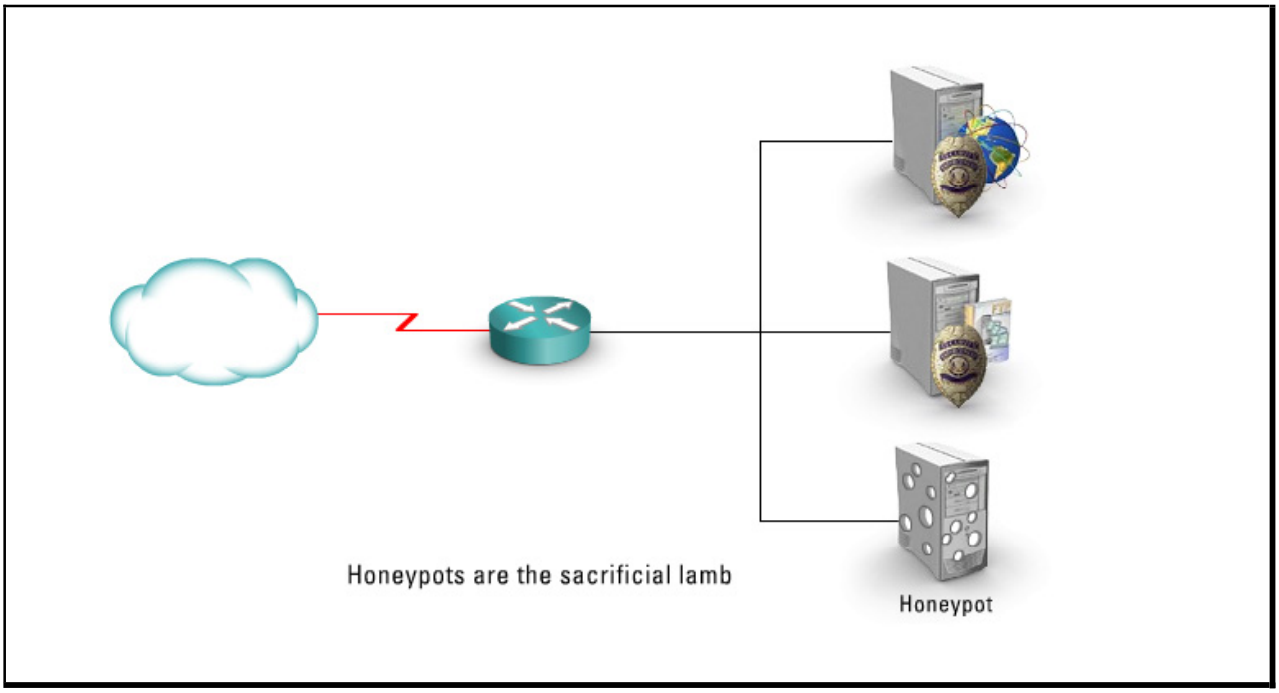

A **honeypot** is an information system resource that is given up as the sacrificial lamb to attackers. It is a system that is intentionally configured to be less secure than other systems. Data placed on the system is never sensitive, but might be mocked up to look important. To provide the best results, honeypots should be configured with the same services and applications as production servers, but with holes that allow attackers easier access.

The goals of the honeypot are many:

- First is the hope that an attacker will choose the 'easy' target, not production servers.
- Second, the security team can analyze the attack methods used by the attacker and better secure the production systems.
- Third, new tools and tactics can be analyzed before they hit other systems.
- Fourth, since no one is really using this system, any attempt to gain access must be coming from an attacker.
- Last, the security team can attempt to identify the attacker.

Remember, that honeypots cannot be used in a court of law to bring charges against the attacker. Since the data the intruders are taking from the honeypot does not cause financial loss to the business, courts have determined that attackers caught infiltrating a honeypot cannot be liable. No matter, the real benefit from the honeypot comes from learning a better understanding of the new and latest attacks and using them to better secure the network.

Honeypots also have legal issues, which include issues of entrapment, privacy, and liability. Since you will not be using the honeypot to capture a criminal, and you are not a law-enforcement official entrapment does not apply. Privacy laws in the US may limit your right to capture data about an attacker,

even when the attacker is breaking into your honeypot. Using banners on honeypots stating that an attacker consents to logging will waive the attackers privacy rights. A banner such as the following are legally acceptable:

```
############################################################
# !READ BEFORE CONTINUING!
#
# This system is for the use of authorized users only.
# By using this computer you are consenting to having
# all of your activity on this system monitored and
# disclosed to others.
############################################################
```

One of the challenges becomes which port to banner. For proper due diligence, all ports and services that are normally bannered should have a banner. Another issue is liability. What happens if the attacker compromises the honeypot and uses it to attack and cause harm to another entity? The argument is that if the security administrator had taken proper precautions to keep the system secure, the attacker would not have been able to harm another system, so the security administrator shares the fault for any damage occurred during the attack. An easy countermeasure would be to disallow any session initiated from the honeypot at the router or firewall. In this way, attackers can reach the system and attempt to gain access, but they will not be able to initiate attacks from the system.

A **honeynet** is used to perform the same function as the honeypot on a much larger scale. A honeynet is a network segment placed behind a firewall that captures all inbound and outbound data. The firewall limits the amount of malicious traffic than can leave the honeynet. The data is now contained, captured, and controlled. Standard production systems are used on the honeynet, in order to give the attacker the look and feel of a real system. Remember that since the honeynet has no production value, it should never be generating or receiving traffic. Thus, all ingress and egress traffic must be considered suspicious.

Attack Signature Identification

Signatures are the tools used to identify attacks on the network. To explain how a signature works, take a look at an example of a simple directory traversal exploit.

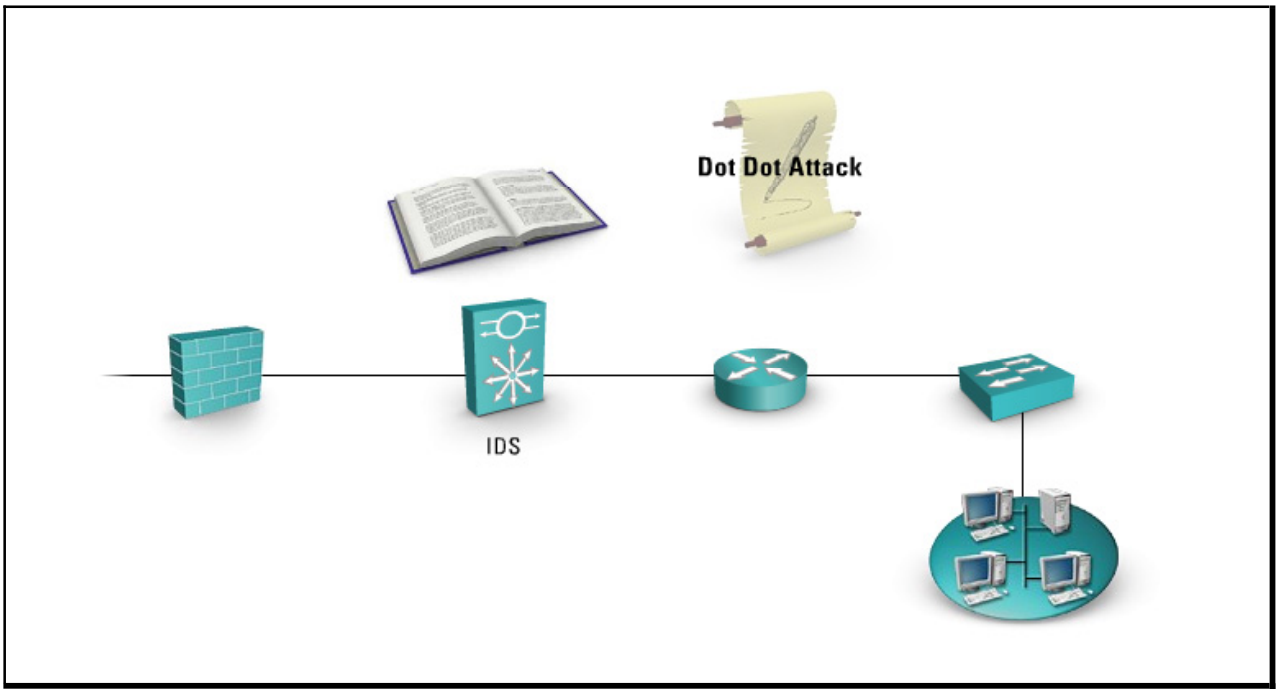

In this exploit, the attacker attempted to view directory contents on a server while connected through an http session. The actual exploit is the “dot dot” attack ../, so a simple signature can be created that only looks at ingress http packets (TCP port 80) and looks in the payload for the string../. If the string produces a match, the signature is triggered and some action can take place.

Many of the thousands of exploits have signatures to detect them. Another simple exploit is the impossible IP packet, also called the LAND attack after the name of the program that generates this attack. A LAND attack consists of a stream of TCP SYN packets that have the source IP address and TCP port number set to the same value as the destination address and port number of the attacked host. Some implementations of TCP/IP cannot handle this theoretically impossible condition, causing the operating system to go into a loop, as it tries to resolve repeated connections to itself. A signature to detect this type of attack would first perform a test to verify that the attack is a TCP-based packet, and then would check the source and destination IP fields. If the signatures match, the system would determine that the IP packet is impossible, and the signature would be triggered. This signature would not have to validate any TCP flags or TCP port numbers because if the source and destination match, the packet is automatically impossible.

Intrusion Reactive Response, Alarms, Signals

Intrusion detection has at its heart, the analyzer that uses signatures to confirm whether an attack is occurring based on the packet that is currently being checked. Equally important, however, is what is done when a signature has fired.

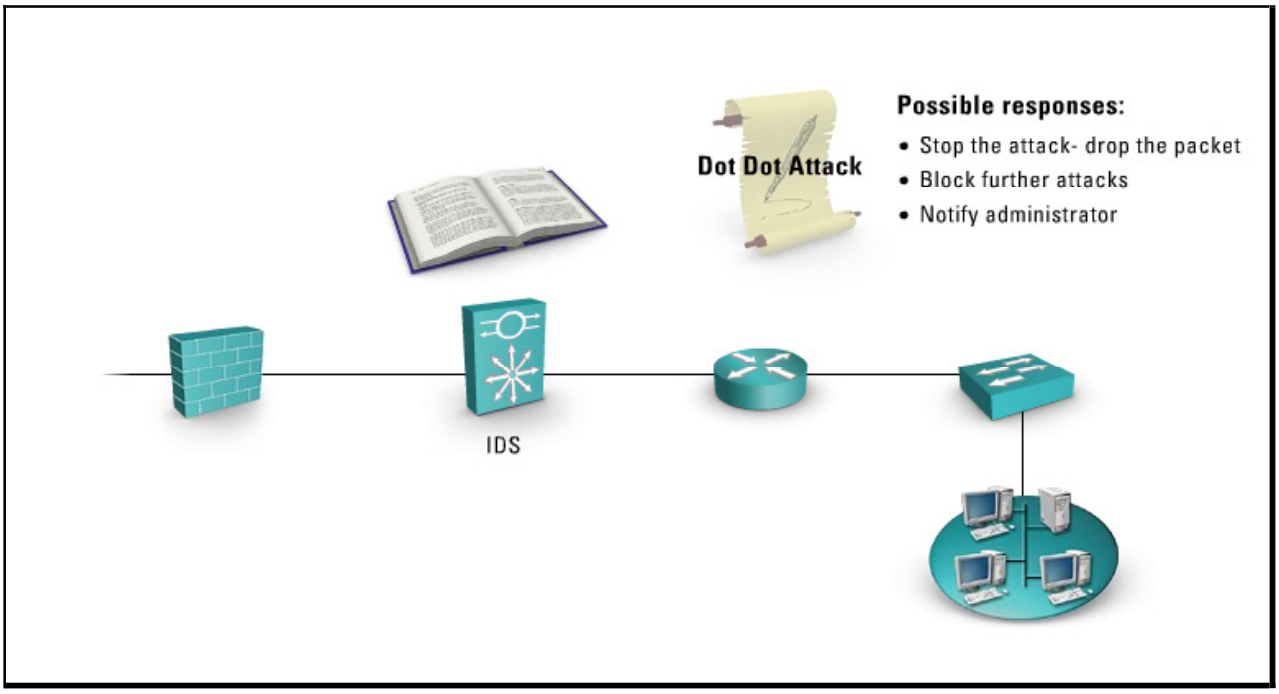

What are the possible responses that the sensor can take to protect the network? In general, you want to do three things when an attack occurs:

- Stop the attack
- Block any further attacks
- Notify the administrator that an attack occurred

These responses seem simple to carry out, but, in reality, they can be difficult to achieve. In the case of stopping the attack, if we use an active or inline device, we can simply drop the packet, but what if we use a passive sensor? Remember that passive sensing devices only look at copies, while the original packet proceeds toward the destination. Using the passive sensor means that if the sensor detected an attack, the packet cannot be stopped. The sensor can act in response to a fired signature to stop further attacks, but in the interim a few or many packets have been received by the target.

Another response from the sensor would be send TCP resets to the attacking source and the target destination. This method is usually faster in stopping the attack. Sending TCP packets is much faster than notifying additional devices to block the attacker using access lists, which would have to be created and applied to an interface.

Notifying the administrator can also become a challenge, in terms of providing information that an attack(s) has occurred. For instance, when an attacker sends 10 malicious packets to a target, the administrator would presumably receive 10 log files, email, and pages, one for each attack. When the attacker sends thousands of malicious packets into the network, the administrator would be bombarded

with thousands of IDS events. Extrapolating further, the administrator would be bombarded with an unmanageable number of IDS events if several attackers each sent thousands of malicious packets. Since the administrator is responsible for checking each one of the attacks and determining what service on what system was attacked, the administrator needs to determine if data were stolen, deleted, or modified. As you can see, responding to an attack implies much work.

Audit Trails

Audit trails are records of events that occurred on a computer system from the operating systems, applications, or user activities.

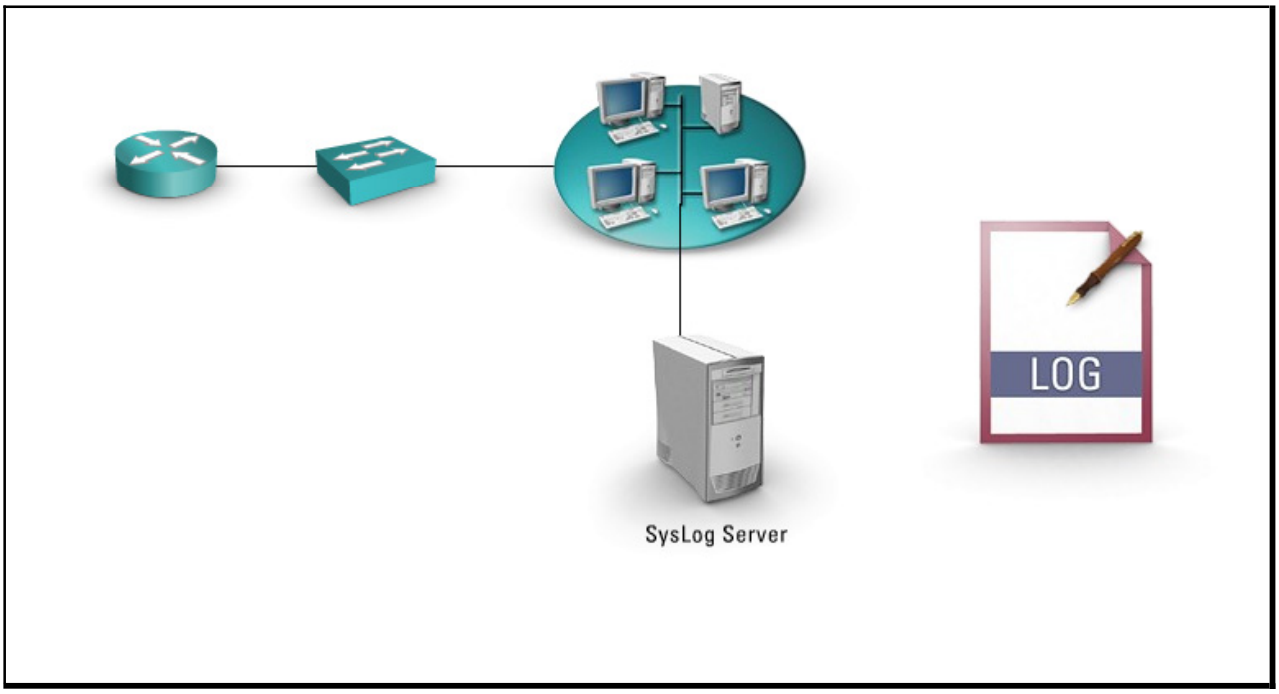

An auditing system generates the audit trail by monitoring all system activity. Audit trails are very important in the security world for items such as the following:

- **User accountability** - A user's actions can be tracked in the audit trail, which mandates that users become accountable for their actions while on the system. This approach is a great deterrent to users who would attempt to circumvent existing security policies.
- **Event reconstruction** - Audit trails can reconstruct events after a security events occurred. Reviewing the audit trails can provide the security administrator with information on the event, such as what, when, how, and why the event occurred, and of course, who was using the system when the event occurred. This information is vital in attempting to ascertain the extent of damage that occurred on the system in question.
- **Activity monitoring** - Audit trails are a great help when attempting to determine what caused a problem to occur. This real-time monitoring can help in the detection of problems such as disk failures, memory and CPU over utilization, or network outages.
- **Intrusion detection** - Audit trails can assist in the operation of intrusion detection if they record appropriate events.

Violation Reports

Audit trails are essentially giant log files that contain a myriad of information, some informational, some vital, with many between. Filtering and researching the audit trails is a time consuming and often tedious task. For this reason, most security professionals rely on reports that perform the filtering of the audit logs automatically for them. One of the most important reports produced is the violations report.

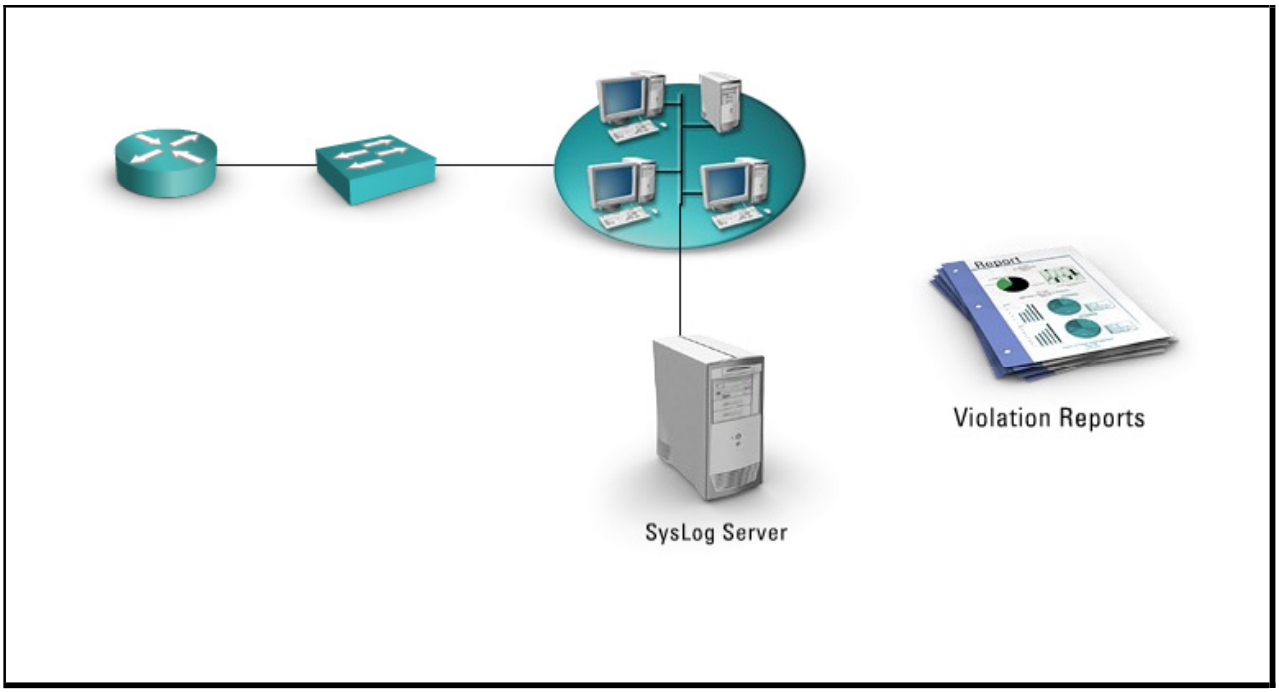

Violation reports are used to identify activities that portend breaches or attempted breaches in security access controls. For example, the violations report will show when someone makes numerous login attempts and failed using different passwords on a secure system.

All security administrators and IT staff should constantly review the violations reports to identify unauthorized access attempts.

Penetration Testing

After implementation of a new security device, the administrators needs to test whether the device is working and that it provides the security for which it was purchased. Penetration tests are used to test the security. Penetration tests simulate an attack on a network or information system at the request or permission of the owner.

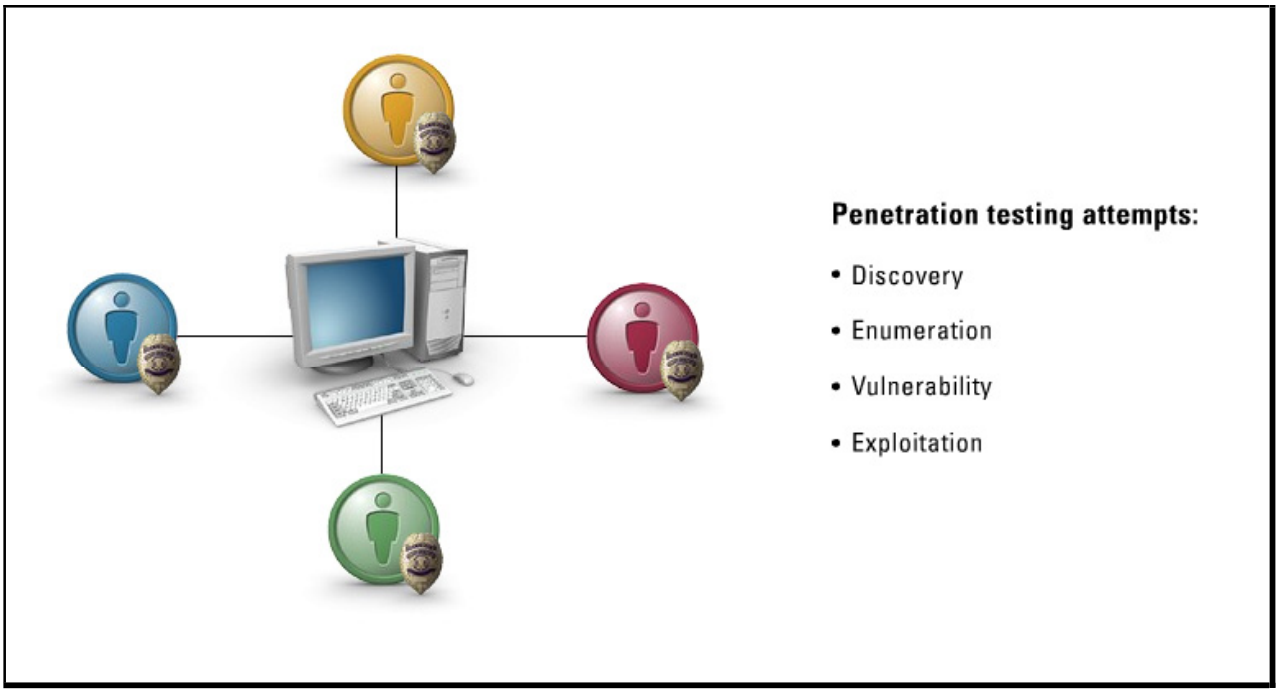

The goal of the penetration test is to identify weaknesses in the information system and apply the proper countermeasures, updates, or patches to protect the system. Areas to be tested should include all prevalent environments in the corporation including the Internet, Intranet, Extranet, or remote dial-in connections.

Penetration testing must be performed by those individuals who have the same level of knowledge and tools as those of the attacker community. Usually, outside teams conduct penetration testing for the larger enterprises. These teams specifically train in conducting penetration tests. For the penetration test to be useful, the team must be clearly defined methodologies and goals. The following at a minimum should be tested:

- Discovery and footprint analysis
- Exploitation
- Physical security assessment
- Social engineering

To provide the most realistic external penetration test, many companies choose to use a zero-knowledge test, which is a penetration test where the penetration team has no prior knowledge about the target network. This test usually begins with reconnaissance in which the penetration team attempts to gather a significant amount of information on the target network. A full-knowledge attack is one where the penetration team has full knowledge of the information system. Generic steps in penetration tests include the following:

- **Discovery** - gathers and documents information about the target system
- **Enumeration** - uses more intrusive methods to gain even more information about the target
- **Vulnerability** - maps the profile of the environment to known vulnerabilities
- **Exploitation** - attempts to gain privileges using vulnerabilities identified

Summary

The key points discussed in this lesson are:

- The types of intrusions
- TEMPEST
- Methods of intrusion prevention
- Ways to identify intruders
- Methods of intrusion detection
- NIDS and HIDS
- Types of data extraction
- Intrusion recognition
- Use of traffic in an attack
- Use of HoneyPots and HoneyNests
- How a signature is used in attack identification
- Use of intrusion reactive responses, alarms, and signals
- Use of audit trails
- Use of violation reports
- Use of penetration testing

Telecommunications, Network, and Internet Security

Overview

Gaining unauthorized access to company resources is the goal of many attackers. How they achieve that goal does not matter. To that end, many attackers who are stifled at their attempts to gain access at one end of the network will just try the other end and a different attack. Understanding how to counteract attacks on the PBX, through the email system, and over the network are critical in securing the enterprise.

The Telecommunications and Network Security domain encompasses the structures, transmission methods, transport formats, and security measures used to provide integrity, availability, authentication, and confidentiality for transmission over private and public communications networks and media.

Objectives

Upon completing this module, you will be able to:

- Describe email and facsimile security, PBX fraud and abuse, secure voice communications, and spamming
- Explain Internet, intranet, and extranet security
- Show TCP/IP affects security
- Show how LANs, WANs, and VPNs affect security
- Explain network layer security protocols
- Show how the Transport Layer security can be compromised
- Explain Application Layer security protocols

Outline

The module contains these lessons:

- Security Overview

- Internet, Intranet, and Extranet Security
- TCP/IP
- LANs, WANs, and VPNs
- Network Layer Security Protocols
- Transport Layer Security
- Application Layer Security Protocols

Security Overview

Overview

This lesson will discuss fundamental security topics such as email and facsimile security, PBX fraud and abuse, secure voice communications, and spamming.

Importance

Information security professionals need to understand the mechanics of attacks on the PBX, through the email system, and over the network in order to properly secure these points of interest in the enterprise.

Objectives

Upon completing this lesson, you will be able to:

- Explain email and facsimile security
- List the types of PBX fraud and abuse
- Explain secure voice communications
- Define spamming

Outline

The lesson contains these topics:

- Email and Facsimile Security
- PBX Fraud and Abuse
- Secure Voice Communications
- Spamming

Email and Facsimile Security

Email protocols such as the Simple Mail Transfer Protocol, (SMTP), Post Office Protocol version 3 (POP3), and Internet Message Access Protocol (IMAP), and applications such as UNIX's sendmail and Microsoft Exchange are usually the first targets hit by attackers who wish to gain entry into a network.

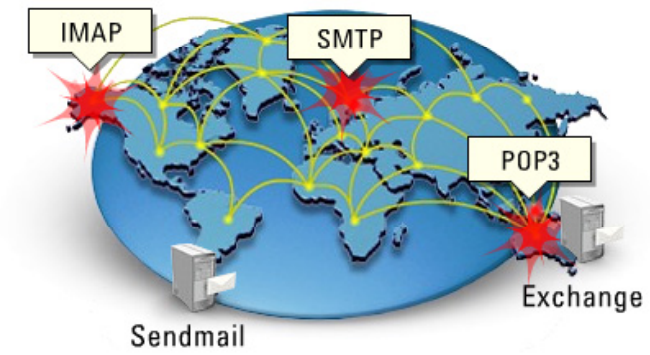

Email protocols are usually the first attacked

These protocols were created to transport email across the network to a repository where the receiver could pull the data at their leisure. SMTP is the protocol used by the sender to send email to the SMTP server, while the receiver used POP3 to retrieve the email from the server. We require two separate protocols because to get email to the correct server, the mail exchanger using SMTP might need to forward the email through many SMTP servers to arrive at the correct destination. When receiving data from an email server, the receiver will know exactly from what server to pull the information. This approach means that quite a bit of intelligence for mail relaying is needed to use SMTP, but the receiver using POP3 needs very little intelligence to pull the data from a known server.

SMTP and POP3 were created without thought to security. Consequently, these protocols with the applications that rely on them were extremely easy to exploit to penetrate a system. Today, these protocols and applications when secured correctly are very difficult to attack with success, but if not secured correctly, can lead to many security-related issues.

Facsimile or fax services are being eclipsed by email services, but they are still preferred when a legal record of transmission and delivery is required. Three technologies incorporate a typical fax machine with a scanner that converts a page into a graphical image, a modem that transmits the data across the Public Switched Telephone Network (PSTN), and a printer to print the resultant image.

The administrator should use exercise great care when accepting a fax as genuine because its integrity may be in question because no data validation or authentication exists between sender and receiver. A fax machine should not be used for confidential information when integrity of the information is paramount. In an effort to reduce this risk, many parties, callers and receivers, often physically watch over the fax

machine in order to capture the expected fax. However, because facsimile is wide open from a security perspective, a wire tap could be used by attackers to intercept faxes.

PBX Fraud and Abuse

Because court decisions make your company, not the carrier, responsible for PBX fraud, the security administrator should take proper steps to secure the PBX system.

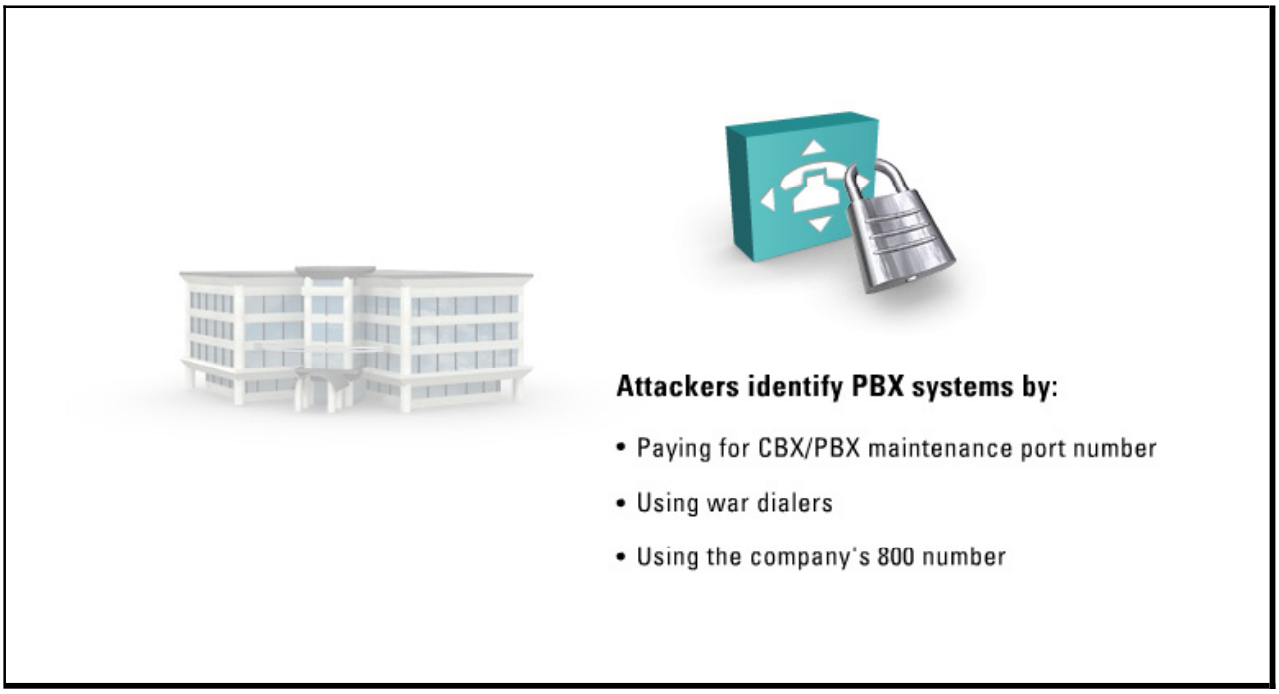

PBX fraud or communication theft can be perpetrated from remote distances by highly skilled and technically sophisticated attackers who have little fear of being detected, much less caught and prosecuted. These criminals conduct a growing business selling access to communications systems all over the world. Your company is responsible because only the company can differentiate legitimate calls from fraudulent ones. There are three ways attackers can find a particular PBX system:

- Attackers pay for a CBX/PBX maintenance port number and password.
- Attackers scan using war dialers to find systems equipped with modems.
- The company's phone number or 800 number make the company known to the attacker.

A war dialer is an application that repeatedly dials different numbers sequentially in an attempt to identify which numbers have a modem attached to them. Some war dialer applications can also identify the particular operating system running on the system and may also conduct automated penetration testing. In essence, a war dialer is used to identify potential targets, which the attacker can then begin attacking in a more meticulous manner.

Protecting your maintenance/administrator port on your PBX is critical, as it is the door that leads into the rest of the system. If an attacker gains access through the maintenance port to system software and can control your voice mail, Direct Inward Service Access (DISA) or other PBX features are unsecured because the attacker owns them.

Secure Voice Communications

Prior to the mid-1960s, nearly all voice communications systems were based on analog implementations. Analog systems are generally unable to implement secure encryption and decryption functions, but in the 1970s, with the advent of digital communications, secure voice communications were developed for military applications.

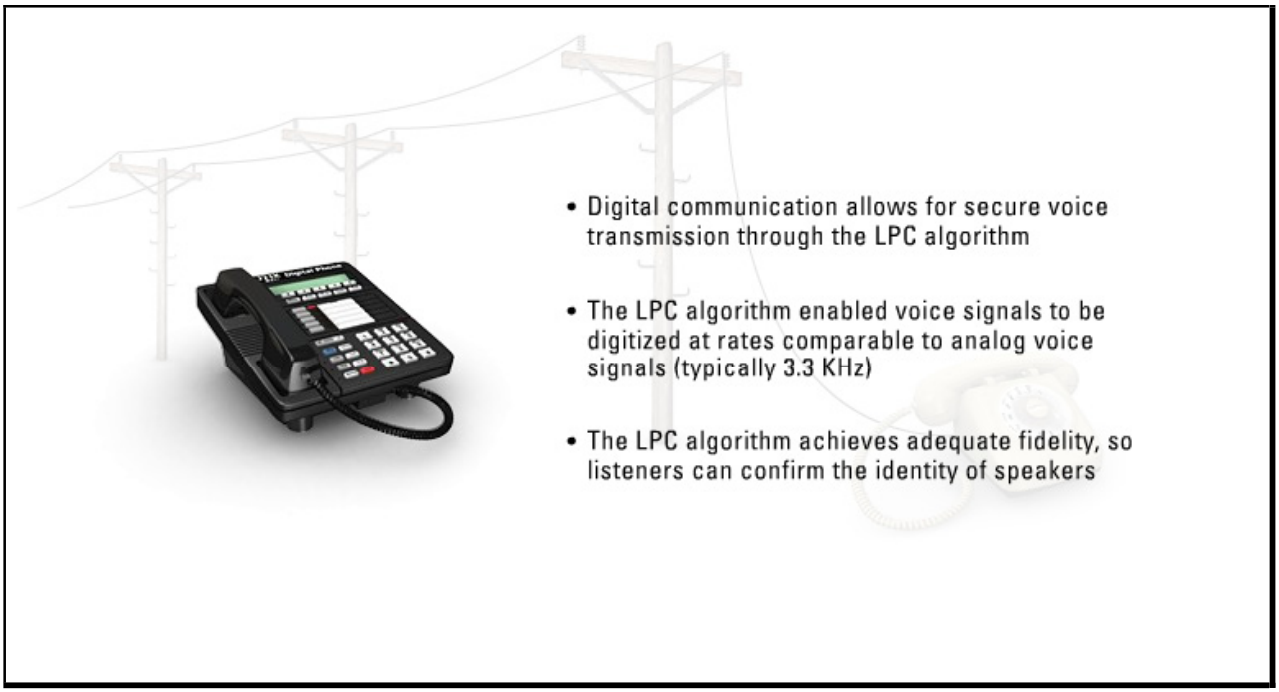

The Linear Predictive Coding (LPC) algorithm enabled voice signals to be digitized at rates comparable to analog voice signals (typically 3.3 KHz). This algorithm also achieved adequate fidelity, so that listeners could confirm identity of speakers, which is very important for integrity and authentication checks.

The LPC algorithm worked by sampling analog signals at a rate of at least 8,000 times a second and quantitized or digitized the signal with a precision of at least eight bits per sample. This algorithm produced a digital bit stream with a rate of 64,000 bits per second (64 Kbps). The advantage of this digital system is that the digital bit stream can be encrypted yielding an apparently random bit stream of 64 Kbps. Its disadvantage it is that the signal requires a bandwidth of at least 64 Kbps, which is a factor of 15 to 20 times the bandwidth of the original analog speech signal. This approach meant that to provide security for a 3.3 KHz analog signal, a 64 Kbps digital signal was required, which at the time, prohibited its use for all but the most important applications. Nowadays, with the advent of high compression algorithms, the digital signal can drop to as low as eight Kbps.

Spamming

Spamming or just spam can be defined as an inappropriate attempt to use a mailing list, USENET, or other networked communications facility as if it was a broadcast medium by sending the same message, unsolicited to a large number of people. The term probably comes from a famous Monty Python skit, which featured the word spam repeated over and over.

Spam costs in productivity and bandwidth usage for consumers and businesses were in the billions of dollars in the year 2002. Because of its rampant use, the U.S. government enacted the CAN-SPAM Act of 2003. This regulation prohibited spamming with imprisonment and/or fines when certain criteria were met, such as sending 20 or more falsified electronic mail messages or 10 or more falsified domain name registrations, with the volume of spam exceeding certain limits, or damages over $5000 were incurred. The Federal Trade Commission (FTC) has the task of enforcing its provisions.

This bill permits e-mail marketers to send spam as long is the following criteria were met:

- An opt-out mechanism
- A functioning return e-mail address
- A valid subject line indicating it is an advertisement
- The legitimate physical address of the mailer

Anti-spam activists greeted the new law with dismay and disappointment. Internet activists who work to stop spam stated that the Act would not prevent any spam. In fact, the regulation appeared to give federal approval to the practice fueling fears that spam would increase as a result of the law.

Summary

The key points discussed in this lesson are:

- Email and facsimile security
- Types of PBX fraud and abuse
- Secure voice communications
- Spamming

Internet, Intranet, and Extranet Security

Overview

The network edge is the point where the enterprise touches the untrusted network. At these points, attacks should be identified and diverted. This lesson will identify the various devices used to protect the enterprise at the network edge.

Importance

Security professionals need to understand the various methods used to protect the enterprise at diverse points in the network.

Objectives

Upon completing this lesson, you will be able to:

- List firewall and proxy technologies
- Explain firewall architecture
- Explain network services configuration
- Explain network address translation
- Compare port address translation to NAT
- Explain the use of routers and gateways

Outline

The lesson contains these topics:

- Firewalls and Proxies
- Firewall Architecture
- Network Services
- Network Address Translation
- Port Address Translation

- Routers and Gateways
- Intranets and Extranets

Firewalls and Proxies

Firewalls restrict access from one network to another, internally or externally.

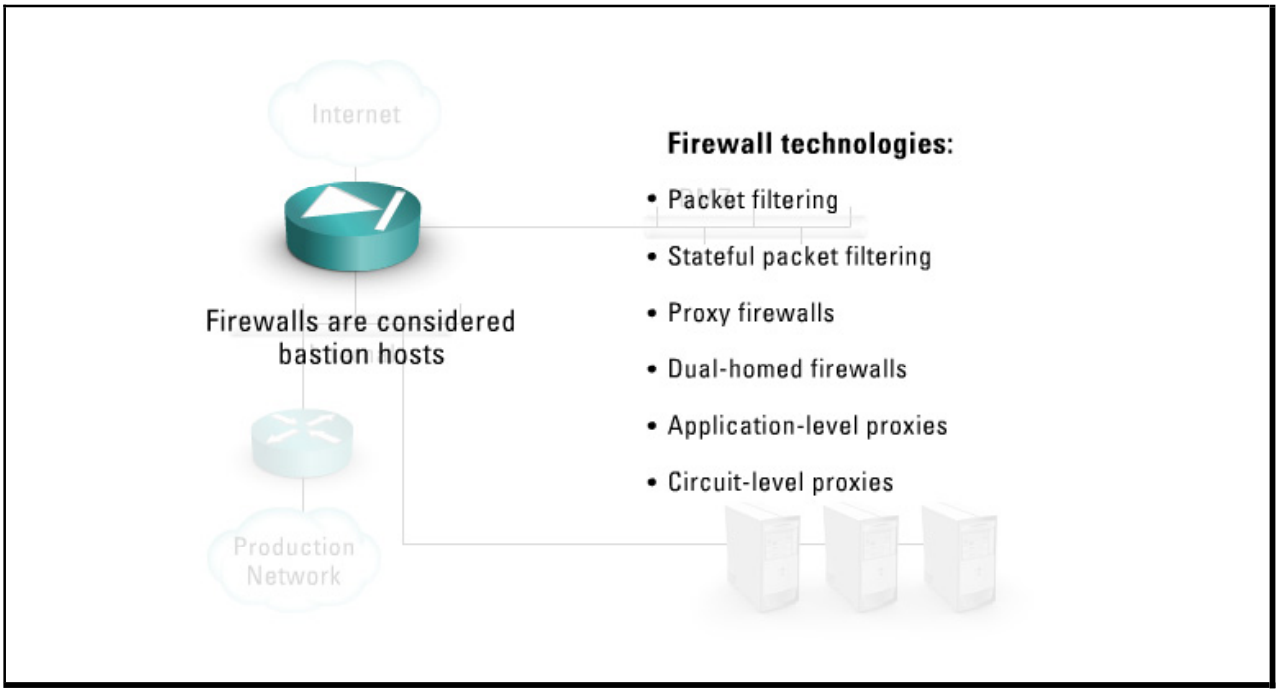

Firewalls are usually placed in the demilitarized zone (DMZ), which is a network segment that is located between the protected and the unprotected networks. A few types of technologies can be used singly to create a firewall; although, most modern firewalls take advantage of more than one technology:

- **Packet filtering** - A method of controlling what data can flow into and out of a network. Packet filtering firewalls use access control lists (ACLs) to filter traffic based on Layer 3 and/or Layer 4 information. Packet filtering provides high performance but does not keep track of session states.
- **Stateful packet filtering** - Stateful packet filters use a state table to identify session state. If a packet was initiated from a trusted interface, its state was identified by source IP, destination IP, source port, destination port, and TCP flags. If a packet returning from an untrusted interface matches an existing session in the state table, the filter allowed the packet through. If a packet cannot be matched to an existing entry in the state table, it will be dropped as in the case of an unsolicited initial packet from the Internet. Stateful packet filters work at the network layer.
- **Proxy firewalls** - Proxy firewalls stand between a trusted and untrusted network and actually make the connection, each way, on behalf of the source. In this way, connections on the Internet all look as if they are initiated from the proxy firewall. Proxy firewalls work at the application layer, meaning they have to decapsulate all the way to Layer 7 in the OSI model. This approach gives them the ability to provide very high security, but at the cost of CPU. Proxy firewalls were often very slow, but with new technology and very high CPU rates proxy, firewalls now work very well.
- **Dual-homed firewall** - A dual-homed host has two interfaces, one facing the external network and the other facing the internal network. To use a dual-homed host as a firewall, the administrator must disable the routing functionality thus disallowing internal and external systems to communicate directly with each other. IP traffic between them becomes completely blocked. Proxy firewalls usually use the dual-home firewall functionality.

- **Application-level proxies** - Application-level proxies work at the application layer to inspect the entire packet and make access decisions based on the actual content of the packet. Application-level proxies understand different services and protocols and the commands that are used within them since they are application dependent. Like proxy firewalls, application-level proxies have to decapsulate all the way to Layer 7 in the OSI model except there must be one application-level proxy per service. Application-level proxies support multi-channeled applications, such as FTP and H.323.
- **Circuit-level proxies** - Create a circuit between the client computer and the server. Unlike the application-level proxy, a circuit-level proxy is one that creates a circuit between the client and the server without interpreting the application protocol. This approach means the circuit-level proxy can handle a wide variety of protocols and services. The proxy knows the source and destination addresses and makes access decisions based on this information. Circuit-level proxies work at the network layer. SOCKS is an example of a circuit-level proxy gateway that provides a secure channel between two TCP/IP computers. SOCKS uses sockets t o represent and keep track of individual connections. A socket is nothing more than a source IP address with the corresponding source TCP or UDP port used in the session. SOCKS does not provide detailed protocol-specific control. For example, FTP communicates on the control channel, but it sends data on the data channel. Since the circuit-level proxy has no application level knowledge of this transaction, it breaks. Circuit-level proxies usually support only single channel applications.

Firewalls are considered bastion hosts. A **bastion host** is a gateway between an inside network and an outside network that is fully exposed to attack. Used as a security measure, the bastion host is designed to defend against attacks aimed at the more trusted networks. Depending on a network's complexity and configuration, a single bastion host may stand guard by itself, or may be part of a larger security system with different layers of protection. Some network administrators will also use sacrificial lambs as bastion hosts, these systems are deliberately exposed to potential hackers to both delay and facilitate tracking of attempted break-ins.

Firewall Architecture

Typical firewall architecture consists of three interfaces: a high security interface that connects to the trusted internal network, a low security interface that connects to the untrusted network, Internet, and a medium security interface that connects to the DMZ.

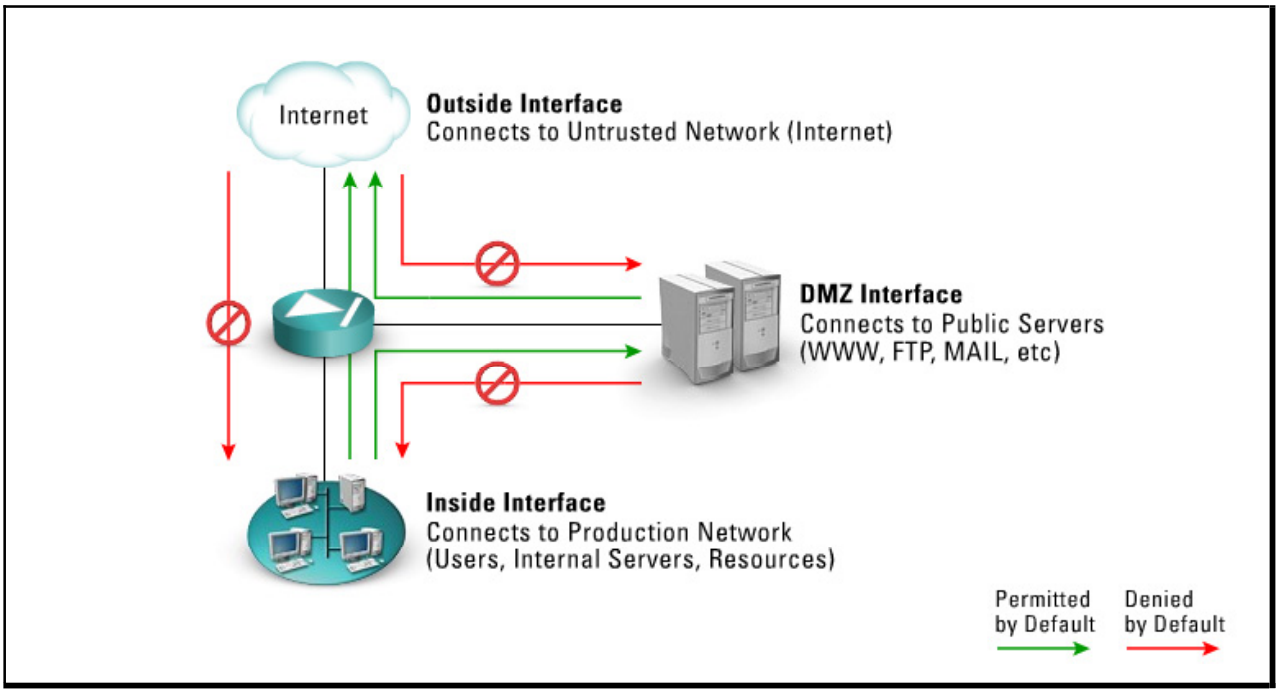

Many architectures support only two interfaces and many support more than three. From a security perspective, you can configure three or more interfaces in the same fashion. In front of the firewall, the perimeter or screening router is the first line of defense. This router provides basic security functions such as RFC 1918 and RFC 2827 filtering as well as performing other important features such as rate limiting, and QoS.

The default action of any firewall should be to implicitly deny any packets not explicitly allowed to reach the DMZ or Inside Network. Packets destined to the Inside Network should only be allowed if they are responding to a session initiated from the inside. The DMZ, though, has publicly available servers that never initiate connections, such as web servers, mail servers, and FTP servers. In this case, traffic can be initiated from the Internet and reach them, but for security reasons, only allow the minimum services to reach them. For example, the web server should only allow incoming connections to TCP port 80, access to all other ports should be blocked.

Network Services

For networking to function as expected, two items absolutely must be configured to operate correctly. The two items are the NOS and the DNS.

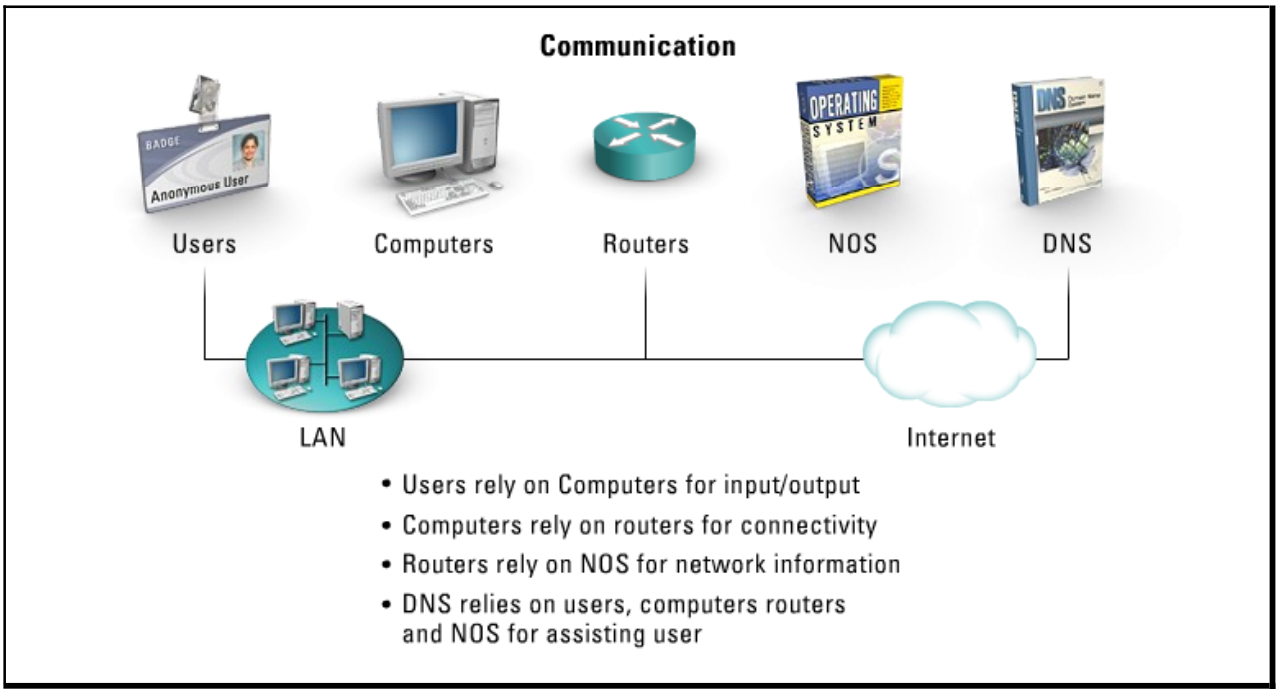

The Network Operation System (NOS) - The NOS is software specifically created to enable a computer to interact with the surrounding network. The NOS was first designed to read layer three information, determine the destination address, and forward the packet to the next logical hop to reach it. This hop required sophisticated new protocols and services to handle the job. Routing protocols allow for dynamic learning of network resources. The router became the central focal point for the NOS and expanded its support to provide much more than simple routing.

The Domain Name Service (DNS) - DNS is a method of resolving hostnames. In today's network, IPv4 can address over four billion IP addresses. The exact number is 2^32. IPv6 is making its way over the horizon and can support 2^128 IP addresses. Humans are not very good at remembering IP addresses, such as 199.107.69.75, but they do have a much easier time remembering names, such as www.usa.gov. DNS is the system by which a name is resolved into an IP address in order to facilitate communications. In DNS, networks are split up into zones, and the DNS server that holds the files for one of these zones is said to be the authoritative name server for that particular zone. Ideally, a primary and secondary DNS server should be available for each zone.

Network Address Translation

In the early 1990s, when the Internet community realized the finite number of IP addresses, they put their heads together to come up with ideas to correct the problem. Some people in the community began working on updating the IP protocol. In this way, they could create an IP addressing scheme that would never run out of IP addresses by recommending a 128-bit address. This approach would eventually take the form of IP version 6. However, IPv6 would be very difficult to implement as the world was already well entrenched in IPv4. Others identified the problem and created **Network Address Translation (NAT)** that would take care of the problem now.

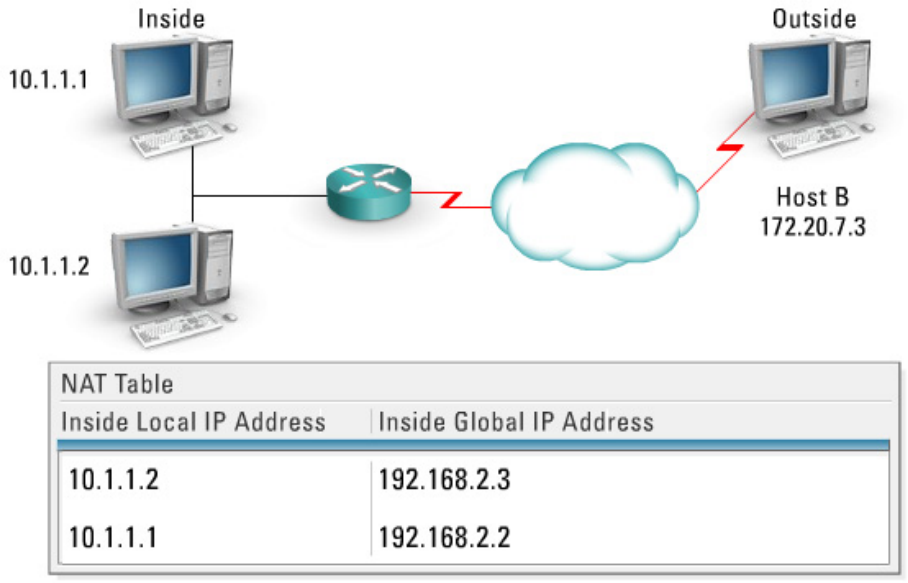

NAT Table

Inside Local IP Address	Inside Global IP Address
10.1.1.2	192.168.2.3
10.1.1.1	192.168.2.2

For example, say a company has 500 users. NAT allows a company to purchase a small amount of public IP addresses and share those IP addresses among all of their users when they wish to access information on the Internet. These 500 users all use the Internet, but only for short times. Internally, all 500 users have IP addresses in the private, RFC 1918 address space, by which the user communicate in the Intranet. When these users attempt to access resources on the Internet, their private address is translated to a public address for the duration of the session. When the session ends, the public IP address is placed back in the pool for others to use.

A great benefit of NAT happened to be one concerned with security. Since the Intranet uses private IP addresses, which are not routable on the Internet, attackers could not access them. Internet users know that a company has limited, for example 64, private IP addresses, and the only way to reach resources in the company is through those addresses. Unfortunately for any attacker, these IP addresses are constantly being reused by different inside hosts, as the NAT device reassigns each IP address to different client when sessions end. The constantly changing address frustrates an attack on an internal client.

Port Address Translation

Port Address Translation (PAT), also called NAT overload and Network Address Port Translation (NAPT), works like NAT with one difference. While NAT uses a pool of IP addresses, PAT uses a single IP address.

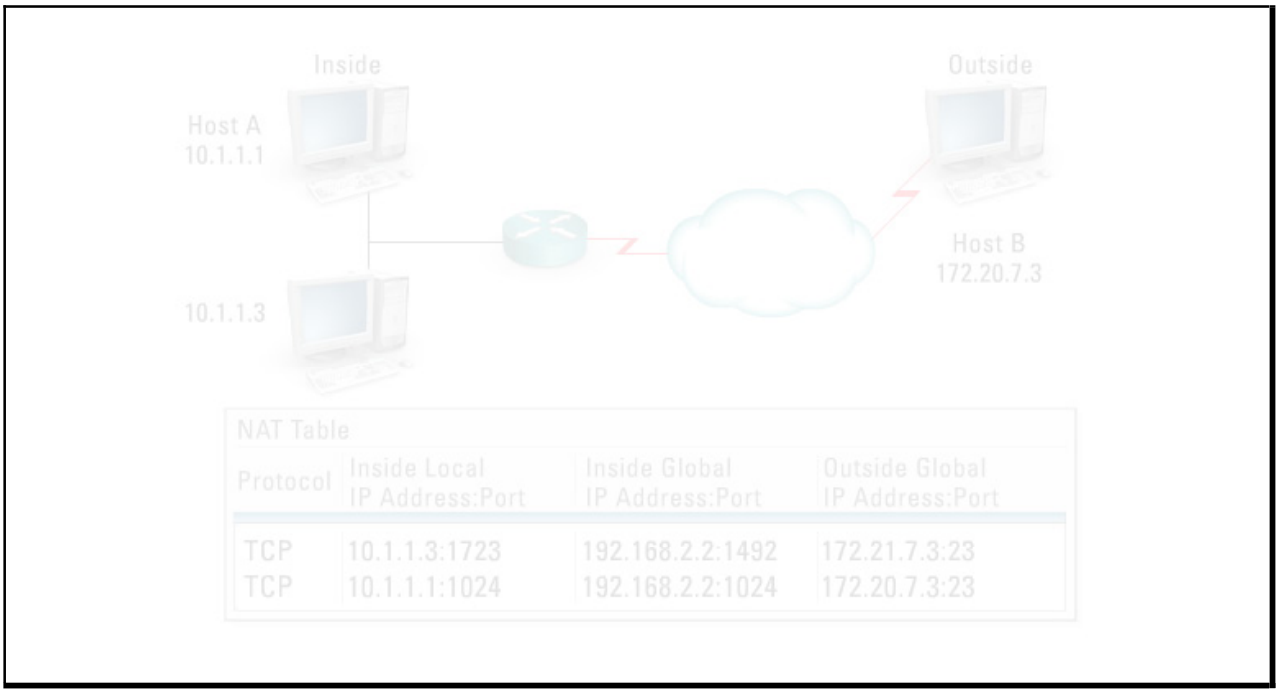

All internal users share the same IP address, as they attempt to access Internet resources, but all users can access the Internet at the same time. To make each session unique, even though they use the same IP address, the server assigns each user a different source port. An IP address plus a port is called a socket. As long as the system has unique sockets, the users have unique sessions.

With NAT, you only translate source IP addresses, but with PAT you translate source port as well as source IP addresses. As the port field is a 16-bit parameter, a single PAT address can theoretically handle up to 65,536 different internal user connections.

Multimedia applications can create problems with PAT since many multimedia applications dynamically assign ports when connections are created. PAT and multimedia applications can collide with one another in assigning ports. Good practices suggest that the security administrator not use PAT if multimedia applications are used on the Internet.

Routers and Gateways

A router is nothing more than a computer with two or more network interface cards built into it.

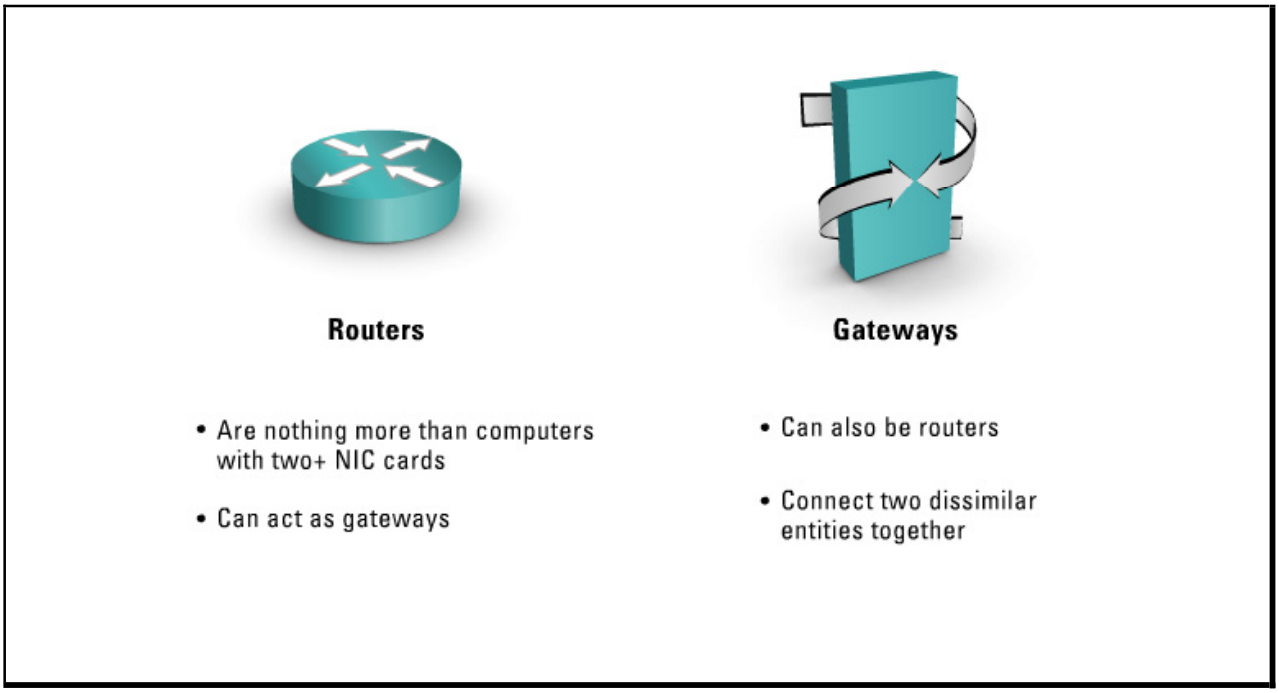

The Network Operating System, NOS, differentiates a router from a workstation. The NOS is the brain of the router that learns about routes to destinations from other routers; talks to other routers using ICMP; receives packets for forwarding; performs QoS, filtering, and rate limiting; builds VPN tunnels; encapsulates; encrypts; forwards (routes); and performs hundreds of other tasks necessary in today's high paced computing environment. Routers are the devices that connect LANs together, which connects corporations, which connect to everything else, making the Internet available to everyone. Routers are in essence the skeleton of the Internet.

Routers can also be gateways. Gateways connect dissimilar entities together. For example, a local hospital may want to connect their Ethernet network to a distance office, but the office exceeds Ethernet's maximum distance. The network engineer can use a microwave connection from the main office to the branch office, but how does he connect the microwave interface to the Ethernet router? And how can Ethernet translate microwave signals and vice versa? Since the devices do not easily connect, the engineer uses gateway. In this case, the gateway is a device that has an Ethernet port used to connect to the main office network and a microwave transceiver used to connect to the remote office. Since data formats between microwave and Ethernet are completely different, the gateway works at layer 7 in the OSI model to completely reconstruct the microwave data into an Ethernet packet the LAN can understand. The gateway also functions in the reverse direction to translate Ethernet packets to microwaves. Another application of a gateway could connect an IP network to Voice over IP (VoIP), yet offnet calls would be required. In this case, a gateway is needed to connect the IP network to the dissimilar analog PSTN network.

Intranets and Extranets

This topic differentiates intranets and extranets.

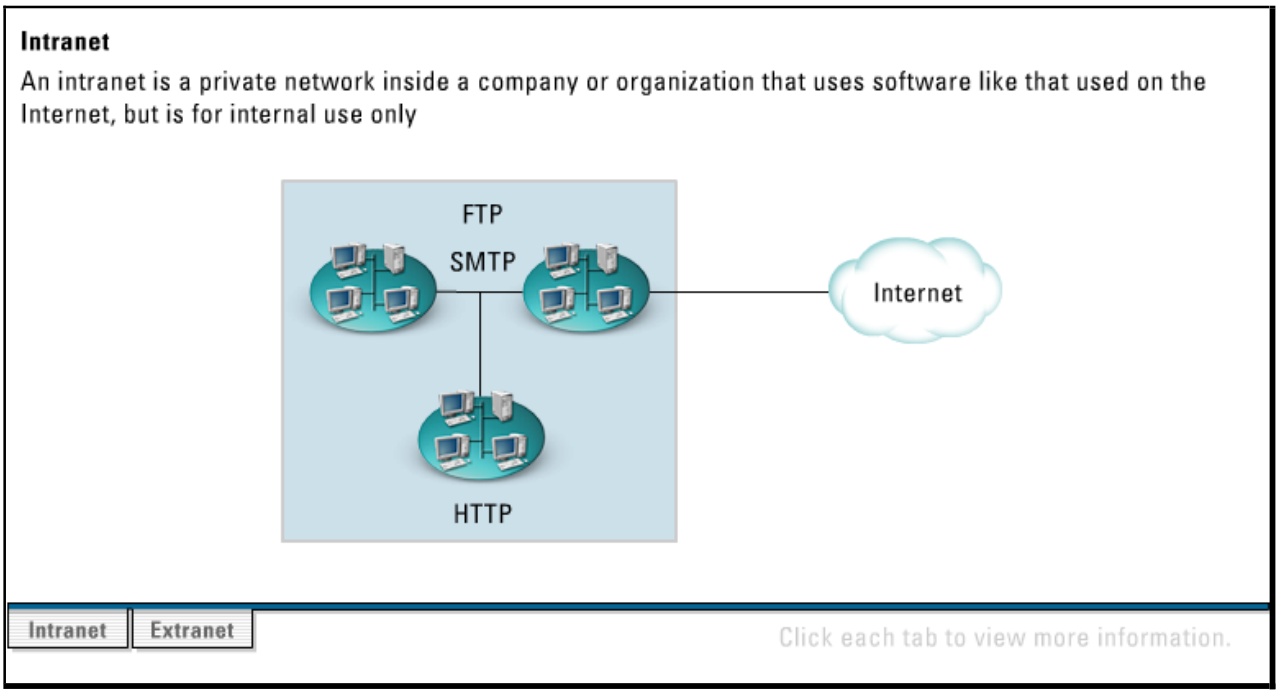

An **intranet** is a private network inside a company or organization that uses software like that used on the Internet, but is for internal use only; an intranet is not accessible to the public. Companies use intranets to manage projects, provide employee information, and distribute data and information.

An **extranet** is a private network that uses the Internet protocols and the public telecommunication system to share a business's information, data, or operations with external suppliers, vendors or customers. An extranet can be thought of as the external part of a company's intranet.

Summary

The key points discussed in this lesson are:

- Firewall and proxy tools
- Firewall architecture
- Network services configurations
- Network address translation
- Port address translation
- Routers and gateways

TCP/IP

Overview

The TCP/IP stack was created to connect various facilities in an open manner. FTP was used for file transfer, SMTP for mail transfer, and Telnet for virtual terminal access. At the time, security was not an issue as the more important goal of just getting connectivity was paramount. As the Internet grew, so did the community who used the Internet for mischievous purposes. In this open environment, Internet hackers and attackers were born. TCP/IP became the field in which good, the general Internet community, fought bad, the attacker community.

Importance

Information security professionals need a thorough understanding of the TCP/IP protocol stack to identify how a hacker attack against the protocol was carried out.

Objectives

Upon completing this lesson, you will be able to:

- Explain TCP/IP characteristics and vulnerabilities
- List TCP commands
- Explain the purpose of port sweeps
- Explain the purpose of evasive sweeps
- Methods of OS identification
- Define UDP
- Name Internetwork agencies
- List OSI layers and characteristics
- Explain the functions of Layer 1 – Physical Layer
- Explain the functions of Layer 2 – Datalink Layer
- Explain the functions of Layer 3 – Network Layer
- Explain the functions of Layer 4 – Transport Layer
- Explain the functions of Layer 5 – Session Layer
- Explain the functions of Layer 6 – Presentation Layer

- Explain the functions of Layer 7 – Application Layer

Outline

The lesson contains these topics:

- TCP/IP Characteristics and Vulnerabilities
- TCP
- Port Sweeps
- Evasive Sweeps
- OS Identification
- UDP
- Internetwork Agencies
- OSI Layers and Characteristics
- Layer 1 – Physical Layer
- Layer 2 – Datalink Layer
- Layer 3 – Network Layer
- Layer 4 – Transport Layer
- Layer 5 – Session Layer
- Layer 6 – Presentation Layer
- Layer 7 – Application Layer

TCP/IP Characteristics and Vulnerabilities

Protocol exploitation in the IP world is vast, prevalent, and on-going. The IP protocol stack was birthed under the name Transport Control Protocol/Internet Protocol (TCP/IP) in 1973.

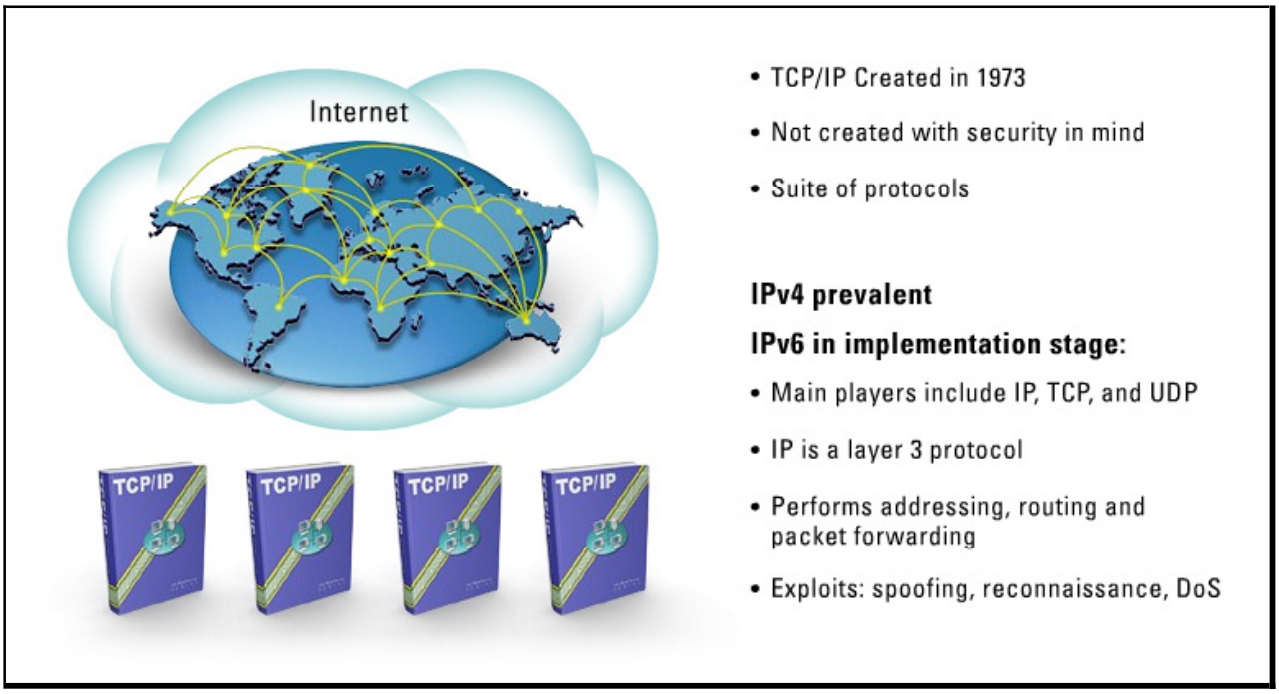

TCP/IP was not created with security in mind; in fact, obtaining open access was the primary concern of the fathers of this new protocol. As the TCP/IP protocol evolved into a solid information transfer technology, the Department of Defense (DoD) took note, and in 1976, required TCP/IP as the protocol of choice for its ARPANET project. The ARPANET would eventually move out of the hands of the DoD and into the hands of the National Science Foundation (NSF), where it would be renamed the Internet.

TCP/IP is a suite of protocols used to transmit data across local area networks (LANs), wide area networks (WANs), and across the Internet. The IP protocol is in its fourth revision, which is identified in the identity version field in all IP packets. IP version 6 is in its implementation stage, but has not of yet been implemented in any large degree in the Internet.

There are three main protocols that support most of the services provided by TCP/IP, those are IP, TCP, and UDP.

IP: A Layer 3 protocol. The main task is to support internetwork addressing, routing, and packet forwarding. IP is a connectionless protocol that envelops data passed to it from the Transport Layer (TCP or UDP). IP is exploited in attacks such as spoofing, information gathering, reconnaissance, and some forms of DoS. Attackers can indirectly attack an IP using ping sweeps. Ping sweeps are used to determine which hosts on a particular segment are alive. If an attacker can determine which IP addresses contain valid hosts, they can now whittle down the address space to those machines that are running. Using port sweeps, the attacker determines what services are running on the devices.

TCP

Transport Control Protocol (TCP) is a host-to-host, Layer 4 transport protocol that provides a reliable, connection-oriented service for IP traffic.

OSI 7 Layer Model

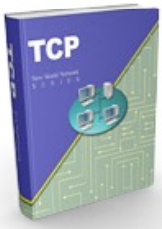

TCP – Layer 4 transport protocol:

- Provides reliable, connection-oriented service
- Well defined session establishment (3-way handshake)
- Exploits use all portions of the TCP header:
  - TCP flags
  - Acknowledgement
  - Sequence numbers

A session using TCP has a very well defined session establishment phase with a three-way handshake. In this handshake, each host attempts to initialize send/receive parameters with its peer, so both parties know exactly how the data transfer will proceed. To invoke this three-way handshake the initiator will send a synchronization (SYN) request to its peer, basically saying "I would like to transfer data with you," and asking "Would you like to transfer data with me?" If the opposite end agrees to transfer data, the initiator will reply with an acknowledgement (ACK). At this point, traffic would be flowing in one way, from initiator to destination, but TCP is a full duplex protocol meaning that the destination must perform the same synchronization (SYN) request and receive an acknowledgement (ACK) in return.

The process would look something like:

Initiator ------> SYN Receiver Start one way session establishment
Initiator ACK <------ Receiver One way session established
Initiator SYN <------ Receiver Start one way session establishment
Initiator ------> ACK Receiver One way session established

These messages may look like a four-way handshake; by combining the receivers two separate messages into a single message we get the following:

Initiator ------> SYN Receiver (Start one way session establishment)
Initiator SYN/ACK <------ Receiver (One way session established)
Initiator ------> ACK Receiver (One way session established)

This dialog is the complete session-establishment, three-way handshake (SYN – SYN/ACK – ACK).

After establishing the session, data transfer can take place as all specific parameters on how data is to be transferred has occurred during the three-way handshake.

When the initiator or the receiver has completed the data transfer, the sender will signal an end to the session by sending a finish (FIN) to the opposite end. The opposite end, should reply with its own FIN to complete the session close phase.

At a high level, this process is how all TCP based sessions occur. However, other possibilities may occur during data transfer, such as a segment getting lost during transit, or making sure a segment is not buffered on the receiving end and introduces latency. Because of these necessities and others, the TCP protocol has additional flags, other than the three already discussed. The following are TCP Protocol flags:

- Synchronization (SYN) – asks for a one way session
- Acknowledgement (ACK) – acknowledges the other party
- Finish (FIN) – end of transmission
- Push (PSH) – tells the receiver to deliver the segment directly to the application without buffering
- Reset (RST) – tells the receiver to drop the connection
- Urgent (URG) – marks a segment as high priority

Attackers have made use of the different TCP flags for reasons other than what the creators had in mind. For example, an attacker can perform stealthy port sweeps by not following the three-way handshake rule. Other different types of attacks that target the TCP protocol are SYN floods, ACK DoS attacks, and session hijacking using TCP sequence number manipulation to name a few.

Attacks generated against TCP are not relegated to only the TCP flags, but to any portion of the TCP protocol such as TCP sequence numbers, acknowledgement numbers, header length, and windows size.

TCP can be identified as protocol number six in the Layer 3 field protocol, used to identify the next level, Layer 4 protocols.

Port Sweeps

Port sweeps determine which services are active on a particular host. In the evolution of an attack, the attacker first uses a ping sweep to determine which servers are alive and reachable on the network.

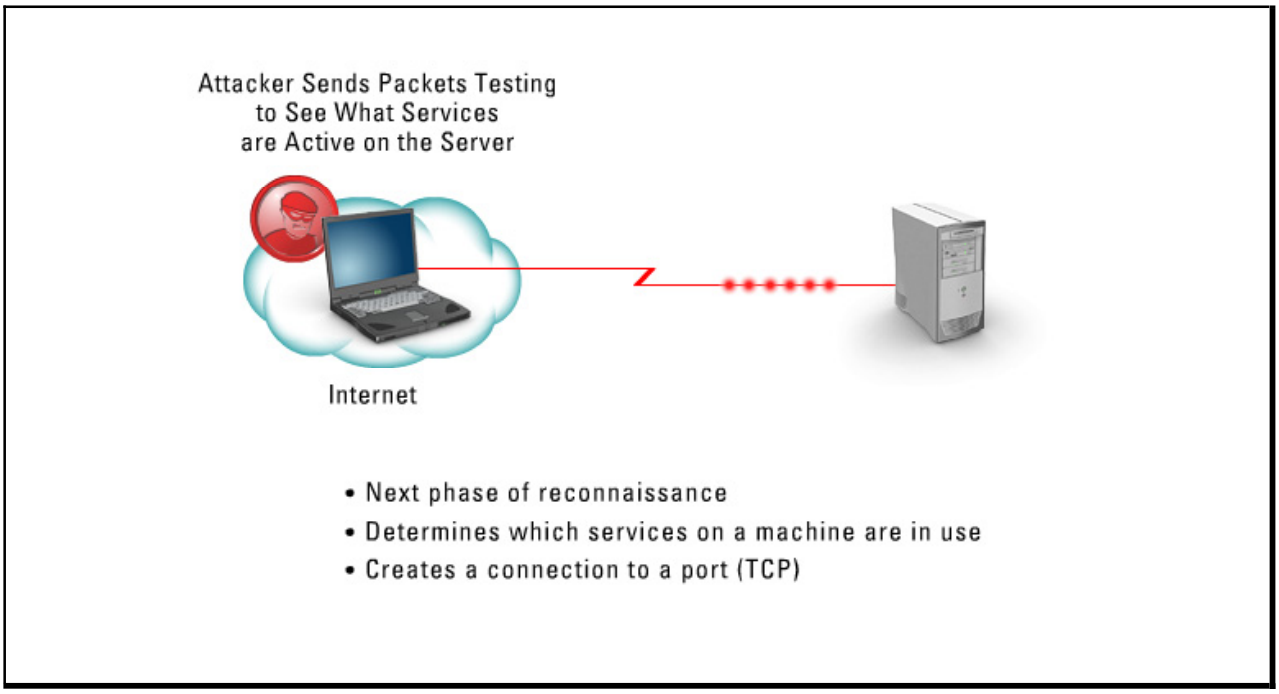

Now that the attacker knows which IP addresses he or she can attack, he or she will perform a port sweep on the systems that are alive. Using this method, the attacker can methodically map which services are running on particular hosts. After gaining this information, the attacker will then attempt to attack vulnerabilities in the active service.

Services on hosts are tied to port numbers. There are 65,536 possible ports than a single host can be listening on, and that range is divided into three ranges.

Well-known Ports - Well-known ports have been assigned by the IANA for specific usages. They are in the range 0 – 1023. Some examples of assigned ports include the following:

- FTP (control) TCP port 21
- SSH TCP port 22
- Telnet TCP port 23
- Domain UDP port 53
- www-http TCP port 80

Registered Ports - Registered ports fall in the range 1024 – 49151. An application developer can attempt to obtain a registered port by my submitting an application to the IANA.

Dynamic/Private Ports - Private or Dynamic Ports are those that fall in the range 49152 – 65535. They can be used by anyone for private use with their applications.

Only when a service is running and listening on a port can an attack occur on that port. If the service is not running and not listening, it cannot be attacked. This fact brings us to one of the first steps to securing

a device: turn off all unnecessary services. If you have a web server running, it needs to listen on TCP port 80. It should not be listening on any other port as the open listening port can introduce the possibility of an attacker gaining access on that port. Stopping attackers on the single TCP port 80 should prove to adequate work.

In a port sweep, the administrator runs a port scanning utility against the web server to verify that only TCP port 80 is listening. If you find other services listening, you will need to research how to disable them. Examples of port scanning utilities include Nmap, Nessus, IPEye, and SuperScan.

Evasive Sweeps

A problem attackers find when scanning networks is that their activity can be easily logged and tracked with a connection to a particular host. In an attempt to evade detection, attackers have delved into the mysteries of the IP protocol suite and exploited some weaknesses that can help them avoid detection. These evasive scan techniques are called stealth scans, and they work by never making a connection; thus not leaving a fingerprint. A connection is created when a full three-way handshake is completed (SYN->SYN/ACK->ACK).

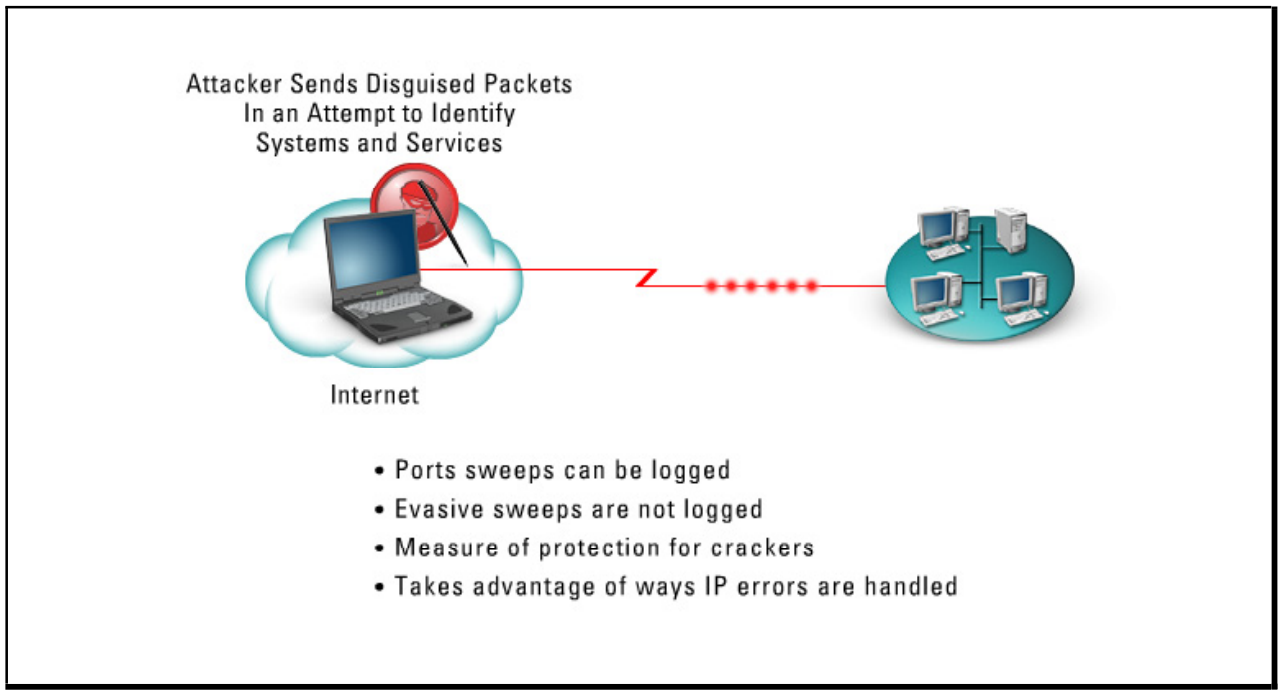

Attackers use different types of stealth scans; some hide the actual attack in a deluge of garbage packets; some perform attacks over time to hide their trail. The most effective type of stealth scans are those that take advantage of known protocol weaknesses, such as SYN stealth scans or FIN stealth scans. In using protocol weaknesses, the attacker is taking advantage of the way an IP based host handles errors. The attacker will send a packet to the receiving host modifying the packet. For example, the attacker may set the FIN flag in the TCP header. When the receiving host receives the packet, it notices the FIN flag set, telling the receiving host to close the connection. The receiving host that never made a connection to the attacker's workstation believes this to be an error in communication, so it sends an error message to the attacker. The message from the receiving host states whether the TCP service port is unavailable, or if it is available, simply ignores the packet. Either way, a connection is never created, meaning a log message is never generated, but the attacker now knows whether or not, a particular service is running on the target host.

Note Examples of evasive port scanning utilities include Nmap, IPEye, SuperScan, and AWSPS.

OS Identification

In order for an attacker to effectively generate attacks on a target system, he or she must know on which operating system the target is running. Being able to target an operating system decreases the attacker's chance of being detected. Discovering the operating system running on a target system is often referred to as a process called **enumeration**.

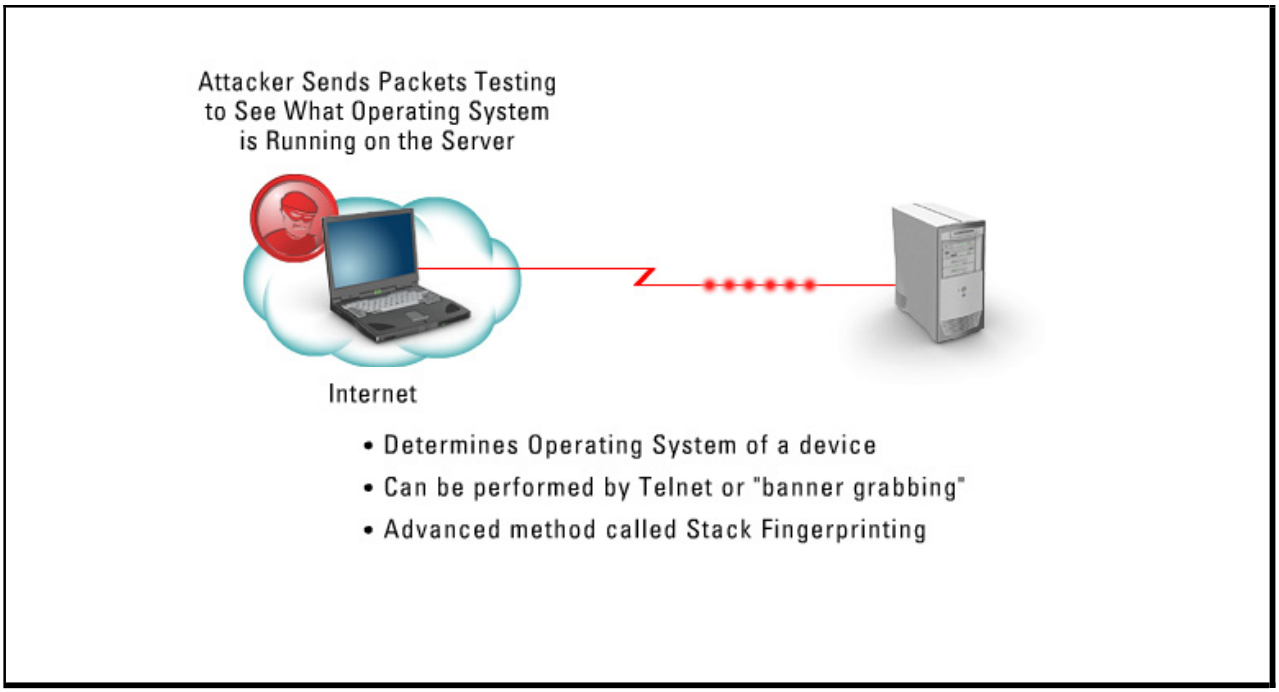

Enumeration can enable an attacker to compromise a system in a relatively short amount of time because the more an attacker knows about a target system, the greater his or her chances are of launching a successful attack. The attacker only needs to attempt to match the operating system against a list of known vulnerabilities.

Enumerating an OS in the old days was relatively easy; the attacker had to merely Telnet to the target, and the target would display its OS on the display. If that approach did not work, the attacker could try banner grabbing. With banner grabbing, the attacker examines the response from certain services like Telnet, FTP, or HTTP. Different operating systems would give different responses making it fairly easy to identify a system.

Today, attackers perform something called active stack fingerprinting. With this approach, the attacker attempts to enumerate an OS by probing its stack. The premise of the attack is the same as banner grabbing, except it is performed on the IP stack. This attack works because different programmers implemented the IP standards in different fashions. For example, if an attacker sends a TCP packet to a target with the TCP FIN flag set, the standard says the OS should not reply. Some implementations such as Microsoft Windows NT, however, return a FIN/ACK, while others might send a RST. By actively probing the stack, an attacker can very accurately determine which OS the target is running.

Note Examples of OS identification utilities include Nmap and Queso.

UDP

The **User Datagram Protocol(UDP)** acts as a transport mechanism like TCP. UDP was created to serve different types of data flows.

- Layer 4 protocol
- Connectionless
- No handshake
- Protocol number 17 in the layer 3 field "protocol"
- Cracker use for UDP bombs and attacks
- Mainly used for voice and video traffic

Since TCP has mechanisms for session creation, tear down, retransmission, and timeout. While UDP is not the fastest protocol in the world, in some situations, the system does not need the entire overhead of TCP. For example, a video stream from a single source to a single destination does not need a feedback mechanism and retransmitting a lost packet would not be feasible. The extra overhead of TCP would actually be a burden.

UDP was created as a poor man's transport protocol. It is connectionless, meaning it has no session establishment. If a segment is lost, UDP doesn't care, because it has no way to find out whether segments have been lost in transit. UDP's claim to fame is that it transmits data very fast in a very efficient manner. Since UDP is much less robust than TCP, it can be implemented in code with a much smaller footprint, making it much less susceptible to attack. Although many UDP based attacks are available, UDP is much easier to guard against exploitation than its sibling TCP is.

Note The most common type of UDP attacks are UDP floods, bombs, storms, and malformed UDP attacks.

UDP can be identified as protocol number 17 in the Layer 3 protocol field, which is used to identify the next level Layer 4 protocol.

Data or PDUs can be called different names: segments, datagrams or frames. Different terms describe data at different layers. Data flowing down the OSI model are referred to as Protocol Data Units (PDUs). PDUs encapsulated at different layers have their own names that describe their functions. When discussing PDUs at the Transport Layer, Layer 4, we call the PDUs segments. When discussing PDUs at the Network Layer, Layer 3, we call them packets or datagrams. When discussing PDUs at the Datalink Layer, Layer 2, we call them frames.

Internetwork Agencies

To promote the health, standardization, and security of the Internet, many agencies have been formed; each one provides needed services to the Internet community.

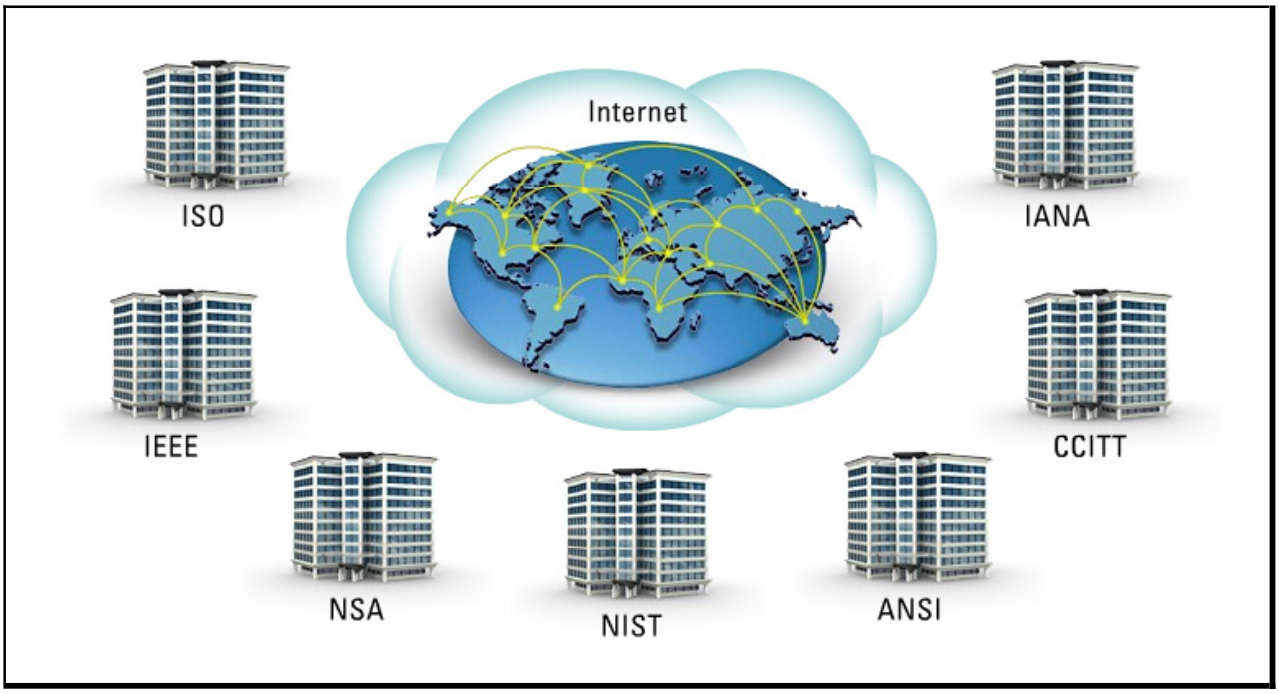

Some of the more familiar Internetwork agencies are listed below:

- ISO - International Organization for Standardization
- IEEE - Institute of Electrical and Electronics Engineers
- NSA - National Security Agency
- NIST - National Institute for Standards and Technology
- ANSI - American National Standards Institute
- CCITT - International Telegraph and Telephone Consultative Committee
- IANA – Internet Assigned Numbers Authority

ISO, International Organization for Standardization, is the world's largest developer of standards. Although ISO's principal activity is the development of technical standards, ISO standards also have important economic and social repercussions. ISO standards make a positive difference, not just to engineers and manufacturers for whom they solve basic problems in production and distribution, but to society as a whole. The ISO is comprised of 146 counties that are part of the national standards institute and is a non-government organization. The OSI model is one of the greatest achievements of the ISO.

The **IEEE** and its predecessors, the AIEE, American Institute of Electrical Engineers, and the IRE, Institute of Radio Engineers, date back to 1884. This organization has advanced the theory and application of electro-technology and allied sciences, served as a catalyst for technological innovation, and supported the needs of its members through a wide variety of programs and services. The IEEE 802 standards are some of the greatest achievements of the IEEE.

The **National Security Agency** (NSA)/Central Security Service is the United State's cryptologic organization. It coordinates, directs, and performs highly specialized activities to protect U.S. information systems and produce foreign intelligence information. A high technology organization, NSA works on the frontiers of communications and data processing. It is also one of the most important centers of foreign language analysis and research within the government. It is said to be the largest employer of mathematicians in the United States and perhaps the world.

NIST, National Institute of Standards and Technology, an agency of the Commerce Department's Technology Administration, was founded in 1901 as the nation's first federal physical science research laboratory. Over the years, the scientists and technical staff at NIST have made solid contributions to image processing, DNA diagnostic chips, smoke detectors, and automated error-correcting software for machine tools.

ANSI, American National Standards Institute, facilitates the development of American National Standards (ANS) by accrediting the procedures of standards developing organizations (SDOs). These groups work cooperatively to develop voluntary national consensus standards. Accreditation by ANSI signifies that the procedures used by the standards body in connection with the development of American National Standards meet the Institute's essential requirements for openness, balance, consensus, and due process.

The **CCITT**, Consultative Committee for International Telephony and Telegraphy, now known as the ITU-T, International Telecommunications Union, the Telecommunication Standardization Sector, is the primary international body for fostering cooperative standards for telecommunications equipment and systems. It is located in Geneva, Switzerland.

IANA, Internet Assigned Numbers Authority, is the organization under the Internet Architecture Board (IAB) of the Internet Society. Under a contract from the U.S. government, it oversees the allocation of Internet Protocol (IP) addresses to Internet Service Providers (ISPs). IANA also has had responsibility for the registry for any "unique parameters and protocol values" for Internet operation. These parameters and values include port numbers, character sets, and MIME media access types.

Because the Internet is now a global network, the U.S. government has withdrawn its oversight of the Internet, previously contracted out to IANA, and lent its support to a newly formed organization with global, non-government representation, the Internet Corporation for Assigned Names and Numbers (ICANN). ICANN has now assumed responsibility for the tasks formerly performed by the IANA.

OSI Layers and Characteristics

The ISO created the Open System Interconnection (OSI) model for worldwide communications that defines a networking framework for implementing protocols in seven layers.

OSI 7 Layer Model

Layer 1: Physical Layer - Transmits bit stream on physical medium. Specifies the electrical, mechanical, procedural, and functional requirements for the physical link.

Examples: 802.3, 802.5, v.35, v.24, EIA/TIA-232, EIA/TIA-449, NRZ, NRZI

Control is passed from one layer to the next, starting at Layer 7, the Application Layer, in one station, and proceeding to the bottom layer, the Physical Layer. The seven layers of the OSI model are listed below:

Layer 7: **Application Layer** - Provides specific services for applications such as file transfer.

Examples: FTP, TFTP, HTTP, SNMP, SMTP, DNS, NFS

Layer 6: **Presentation Layer** - Provides data representation between systems.

Examples: JPEG, GIF, MPEG, MIDI, ASCII, EBCDIC, HTML

Layer 5: **Session Layer** - Establishes, maintains, and manages sessions as well as synchronization of the data flow.

Examples: NetBIOS, RPC, X Windows, SSH, SSL/TLS

Layer 4: **Transport Layer** - Provides end-to-end transmission integrity.

Examples: TCP, UDP, SPX, NetBEUI

Layer 3: **Network Layer** - Switches and routes information units. Determines the best way to transfer data. Routers operate at this layer.

Examples: IP, IPX

Layer 2: **Data link Layer** - Provides transfer of units of information to the other end of physical link. Handles physical addressing, network topology, error notification, delivery of frames, and flow control. Bridges and switches operate at this layer.

Examples: Ethernet, Token-ring, Frame relay, PPP, HDLC, ATM

Layer 1: **Physical Layer** - Transmits bit stream on physical medium. Specifies the electrical, mechanical, procedural, and functional requirements for the physical link.

Examples: 802.3, 802.5, v.35, v.24, EIA/TIA-232, EIA/TIA-449, NRZ, NRZI

Tip	Helpful mnemonic: All People Seem To Need Data Processing

The OSI model does not correspond exactly to the TCP/IP model:

OSI Model	TCP/IP Model
Application Presentation Session	Application
Transport	Host-to-host
Network	Internet
DataLink Physical	Network Interface

Layer 1 - Physical Layer

Layer 1 or the **Physical Layer** in the OSI model specifies the electrical, mechanical, procedural, and functional requirements for activating, maintaining, and deactivating the physical link between end systems.

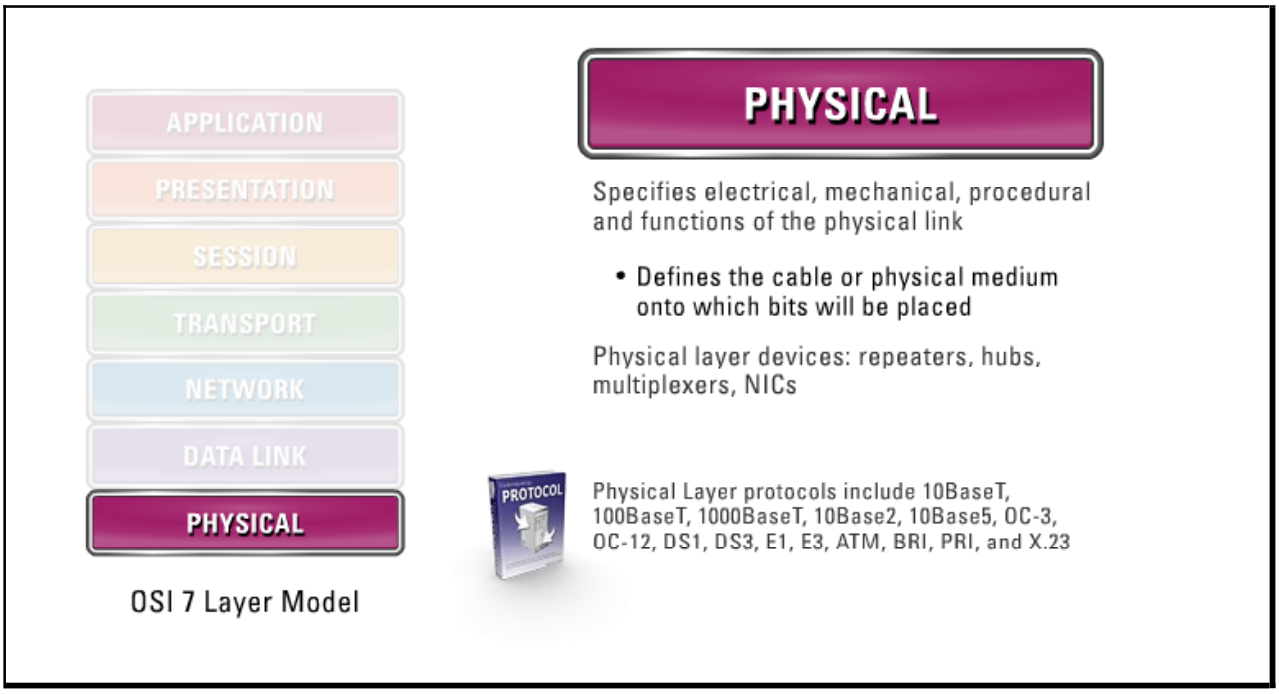

Examples of physical link characteristics include voltage levels, data rates, maximum transmission distances, and physical connectors. It essentially converts bits into voltage for transmission over the physical medium. Basically, the Physical Layer defines the cable or physical medium, such as thinnet, thicknet, or unshielded twisted pair (UTP).

Devices operating at the Physical Layer include repeaters, hubs, multiplexers, and the physical Network Interface card (NIC).

Physical Layer protocols include 10BaseT, 100BaseT, 1000BaseT, 10Base2, 10Base5, OC-3, OC-12, DS1, DS3, E1, E3, ATM, BRI, PRI, and X.23.

Layer 2 - DataLink Layer

Layer 2 or the **Data Link Layer** handles physical addressing, network topology, line discipline, error notification, orderly delivery of frames, and optional flow control.

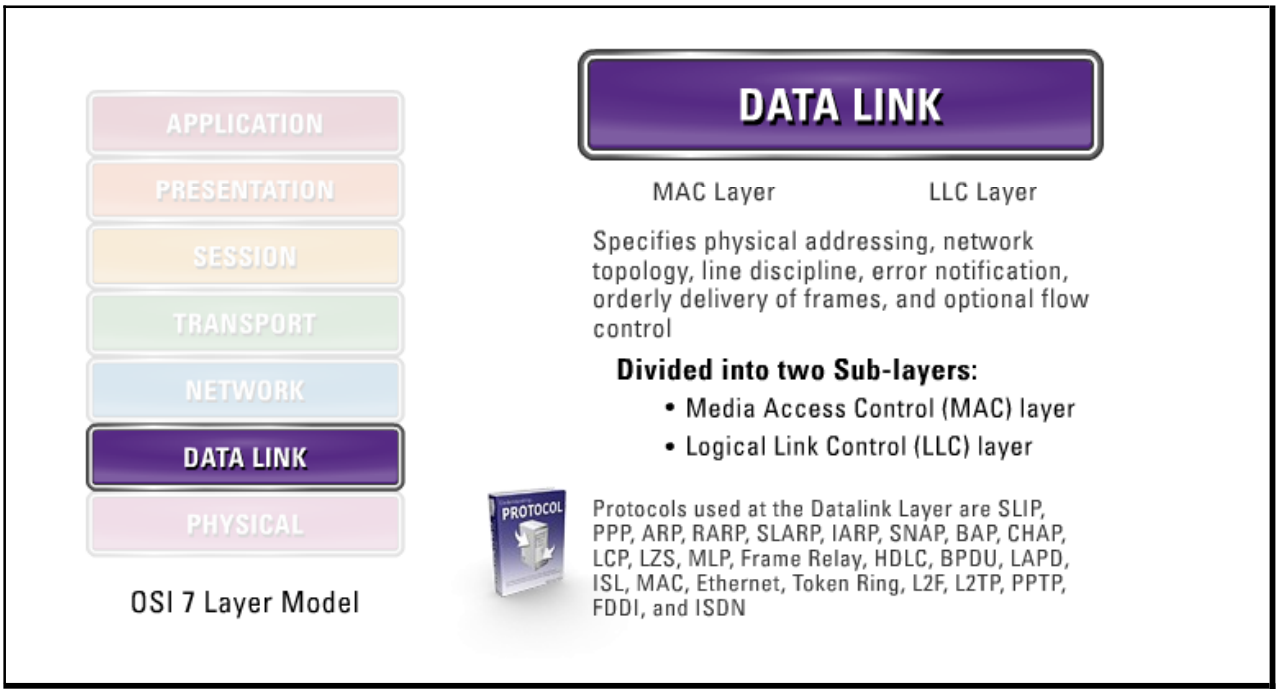

Layer 2 is divided into two sub-layers:

- Media Access Control (MAC) - refers downward to lower layer hardware functions
  - MAC address is physical address, unique for LAN interface card.
  - The MAC address is burned into the Read Only Memory (ROM).
  - MAC address is 48-bit address viewed as 12 hexadecimal digits.
  - The first six digits identify vendor, provided by IEEE.
  - The second six digits are unique and assigned by the vendor.
- Logical Link Control (LLC) - refers upward to higher layer software functions
  - Presents a uniform interface to upper layers.
  - Enables upper layers to gain independence over LAN media access.
  - Upper layers use network addresses rather than MAC addresses.
  - Provide optional connection, flow control, and sequencing services.

A major protocol working at the Datalink Layer is the Address Resolution Protocol (ARP); this protocol is used for mapping IP addresses to physical machine addresses (MAC). RARP or Reverse ARP allows them to request an IP address from the MAC address.

Each network technology has defined electronic signaling and bit patterns.

Bridges and switches work at the Datalink Layer.

Protocols used at the Datalink Layer are SLIP, PPP, ARP, RARP, SLARP, IARP, SNAP, BAP, CHAP, LCP, LZS, MLP, Frame Relay, HDLC, BPDU, LAPD, ISL, MAC, Ethernet, Token Ring, L2F, L2TP, PPTP, FDDI, and ISDN.

Layer 3 - Network Layer

The **Network Layer** or Layer 3 in the OSI model inserts information into the packet header, so it can be properly routed across the network.

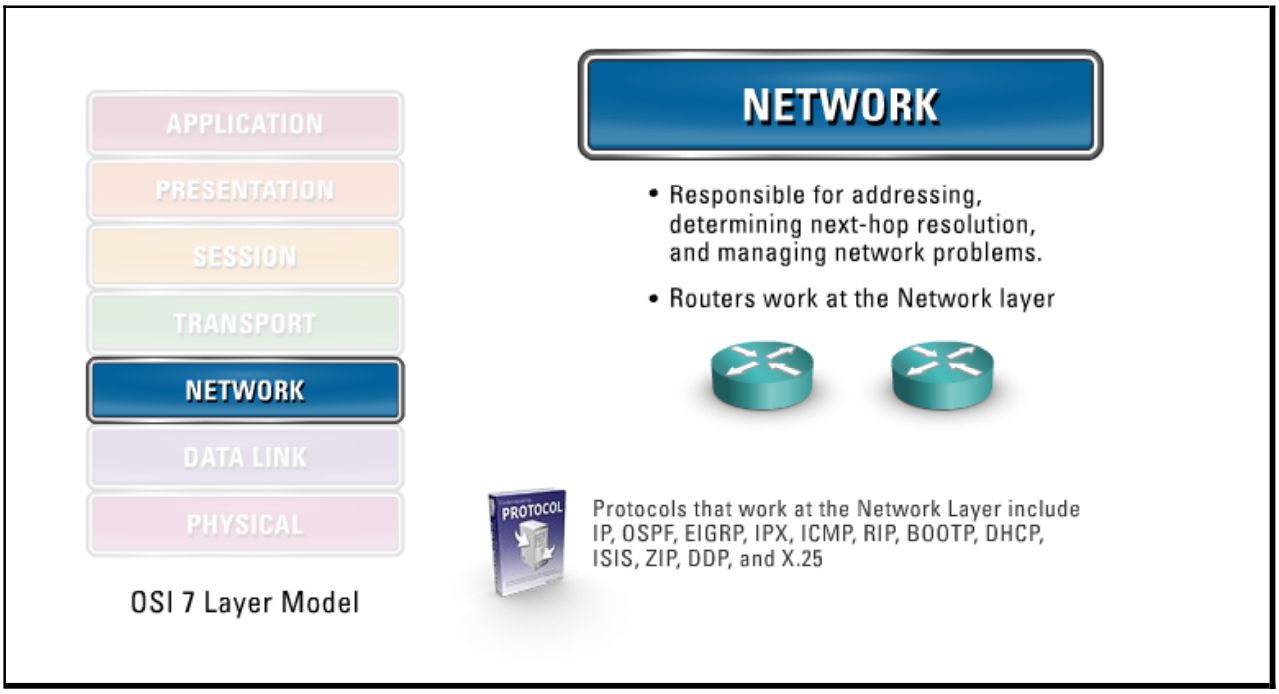

The Network Layer is responsible for addressing, determining next-hop resolution for a destination IP address, and managing network problems such as congestion. The Network Layer breaks down the segment it received from the Transport Layer into smaller maximum transmission units (MTU), which are the largest packets that can traverse a network. At the receiving end, the Network Layer is responsible for reassembling the data.

The Network Layer controls the operation of the subnet. Determining how packets are routed from source to destination is key design issue. Routes could be based on static tables that are wired into the network and rarely changed. The routes could also be determined at the start of each conversation, for example, a terminal session. Finally, routes could be highly dynamic, determined anew for each packet, to reflect the current network load.

If too many packets are present in the subnet at the same time, they will get in each other's way, forming bottlenecks. The control of such congestion also belongs to the Network Layer.

Protocols that work at the Network Layer include IP, OSPF, EIGRP, IPX, ICMP, RIP, BOOTP, DHCP, ISIS, ZIP, DDP, and X.25.

Layer 4 - Transport Layer

Layer 4 or the **Transport Layer** provides end-to-end data transport services and establishes the logical connection between two communicating computers.

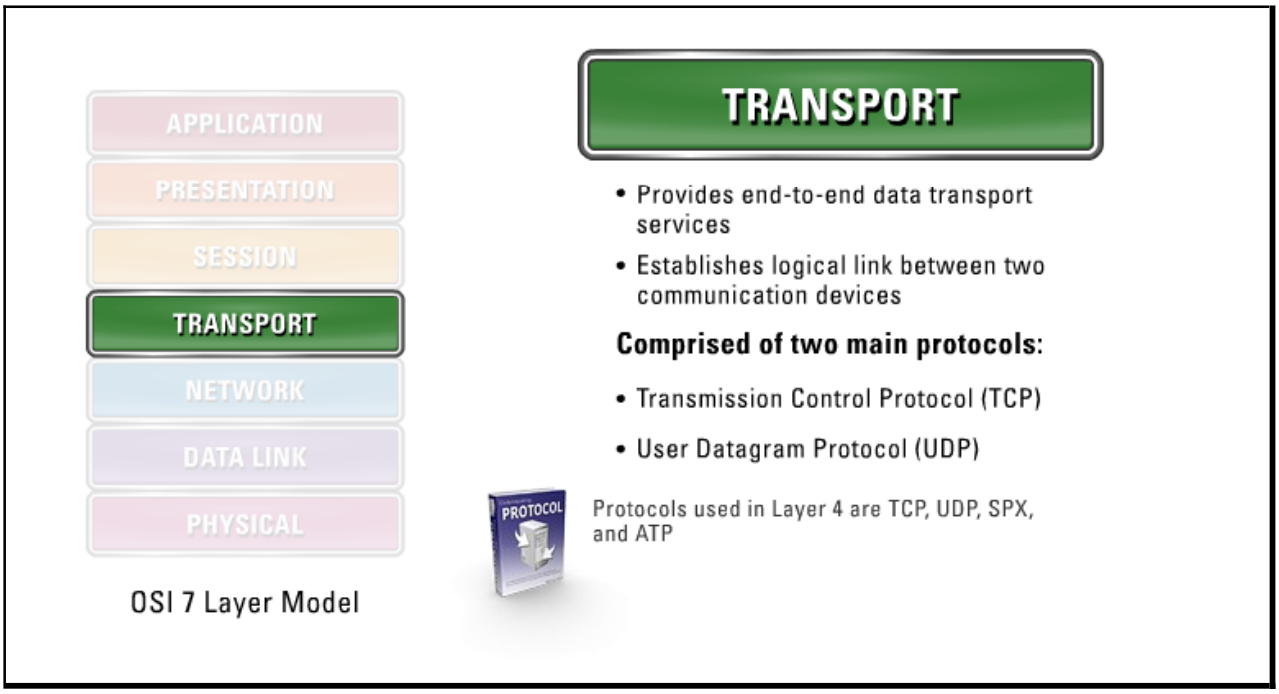

Basically, the Transport Layer manages the flow of data between parties across the network. The basic function of the Transport Layer is to accept data from the Session Layer, split it up into smaller units, pass these units to the Network Layer, and ensure that the pieces all arrive correctly at the other end. Furthermore, all the processing must be done efficiently in a way that isolates the Session Layer from the inevitable changes in the hardware technology.

The Transport Layer subdivides user-buffer into network-buffer sized datagrams and enforces desired transmission control. Two transport protocols, Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) sit at the Transport Layer. Reliability and speed are the primary difference between these two protocols. TCP establishes connections between two hosts on the network through sockets, determined by the IP address and port number. TCP keeps track of the packet delivery order and the packets that must be resent. Maintaining this information for each connection makes TCP a stateful protocol. UDP on the other hand provides a low overhead transmission service, but with less error checking, making it stateless.

Protocols used in Layer 4 are TCP, UDP, SPX, and ATP.

Layer 5 - Session Layer

Layer 5 or the **Session Layer** is used to establish a connection between two computers, to maintain the connection during the transferring of data, and to control the release of this connection.

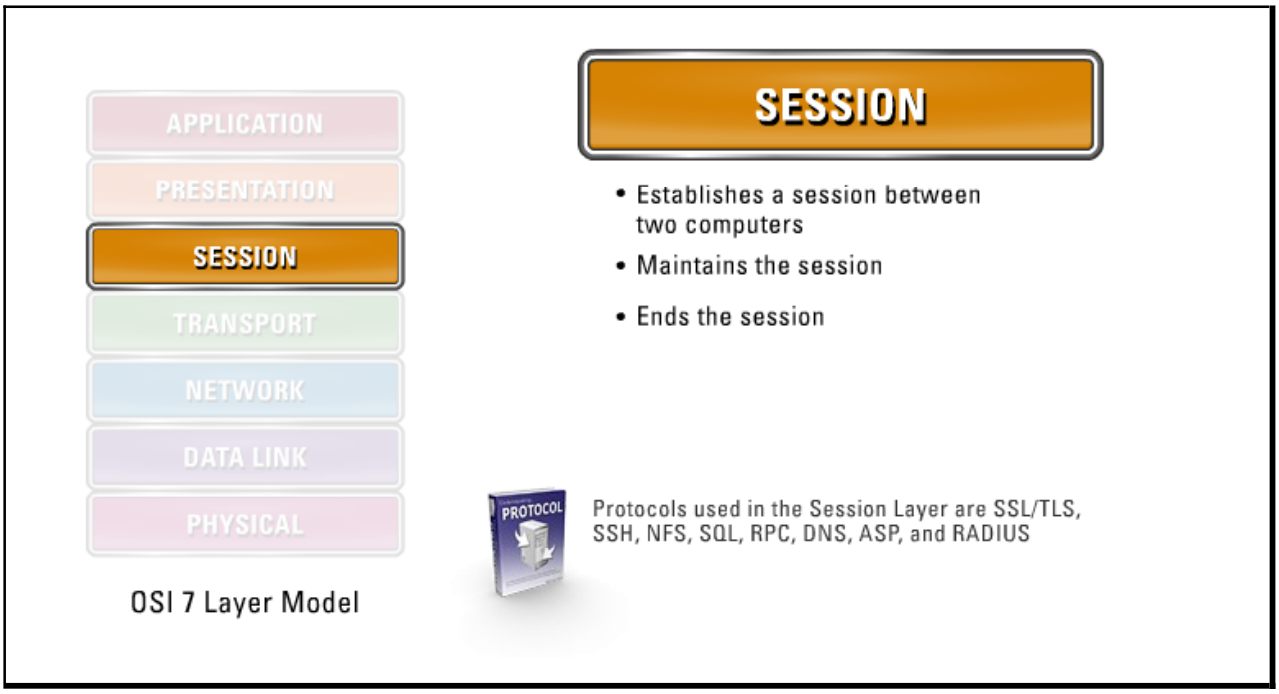

Basically, the Session Layer establishes, maintains, and ends sessions across the network. It is responsible for name recognition, identification, so only designated parties can participate in the session. It also provides synchronization services by planting checkpoints into the data stream. The checkpoints help if a session fails because only data after the most recent checkpoint need be transmitted.

The Session Layer allows users on different machines to establish sessions between them. A session allows ordinary data transport, as does the Transport Layer, but it also provides some enhanced services useful in some applications. A session might be used to allow a user to log into a remote time-sharing system or to transfer a file between two machines.

One of the services of the Session Layer is to manage dialogue control. Sessions can allow traffic to go in both directions at the same time, or in only one direction at a time. If traffic can only go one way at a time, the Session Layer can help keep track of whose turn it is.

Protocols used in the Session Layer are SSL/TLS, SSH, NFS, SQL, RPC, DNS, ASP, and RADIUS.

Layer 6 - Presentation Layer

Layer 6 or the **Presentation Layer** provides a common means of representing data in a structure that can be properly processed by the end system.

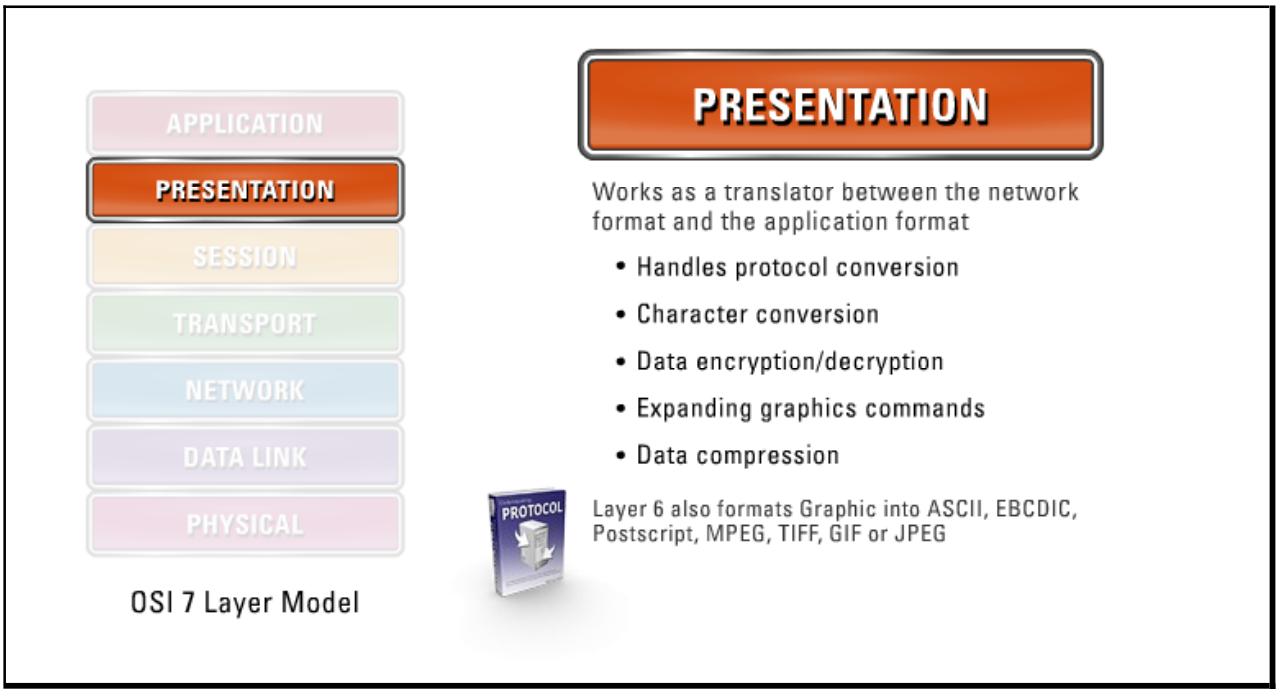

The Presentation Layer works as a translator between the network format and the application format. This approach configures different formats from all sources into a common uniform format that the rest of the OSI model can understand. The Presentation Layer handles protocol conversion, character conversion, data encryption/decryption, expanding graphics commands, and data compression.

Layer 6 also formats Graphic into ASCII, EBCDIC, Postscript, MPEG, TIFF, GIF or JPEG.

Layer 7 - Application Layer

Layer 7 or the **Application Layer** can be considered the applications interface, not the user interface, this application directly supports the services that, in turn, directly support user applications.

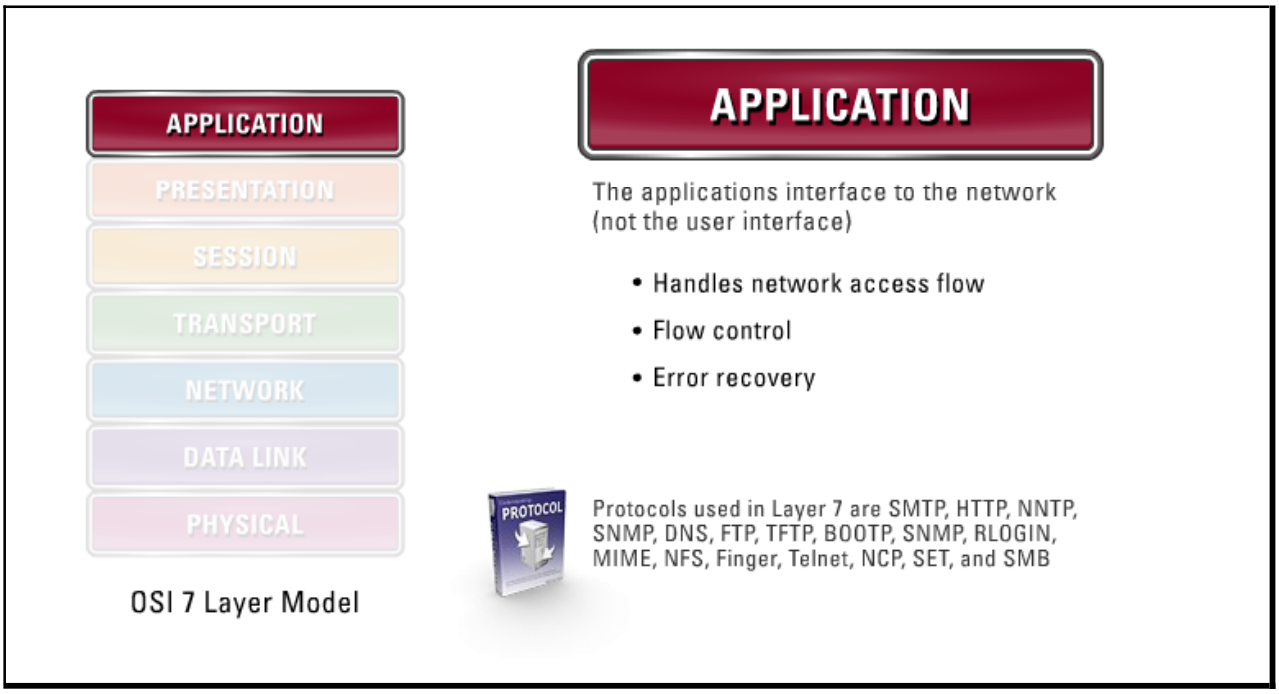

Its purpose is to handle network access flow, flow control, and error recovery.

Protocols used in Layer 7 are SMTP, HTTP, NNTP, SNMP, DNS, FTP, TFTP, BOOTP, SNMP, RLOGIN, MIME, NFS, Finger, Telnet, NCP, SET, and SMB.

Summary

The key points discussed in this lesson are:

- TCP/IP characteristics and vulnerabilities
- TCP
- Purpose of port sweeps
- Purpose of evasive sweeps
- Methods of OS identification
- UDP
- Internetwork agencies
- OSI layers and characteristics
- Layer 1 – Physical Layer
- Layer 2 – Datalink Layer
- Layer 3 – Network Layer
- Layer 4 – Transport Layer
- Layer 5 – Session Layer
- Layer 6 – Presentation Layer
- Layer 7 – Application Layer

LANs, WANs, and VPNs

Overview

Data has to move over a medium to get from point A to point B. This medium will most likely be copper, glass, or air as in the case of wireless. LANs and WANs are composed of discreet network segments using some or many of these network technologies. LANs are usually composed of high-speed Ethernet links, while WANs use lower speed links over provider networks. These different approaches to connectivity use different algorithms based on the separate technologies. This chapter will discuss the different connectivity options as well as the protocols they support.

Importance

Understanding how data travels over the medium using a specific protocol is essential in counteracting malicious attacks on the network level. Different topologies use different technologies, each with their own security concerns.

Objectives

Upon completing this lesson, you will be able to:

- Explain data topologies
- List physical media characteristics
- Describe coaxial cable
- Compare twisted pair to coaxial
- Explain the features of fiber optics
- List the LAN signaling types
- Define the LAN transmission protocols
- List LAN transmission methods
- Explain network topologies
- Define routing in Bus
- Define routing in STAR
- Define routing in Ring
- Define routing in MESH

- List LAN media access methods
- List transmission types
- List LAN devices
- List WAN devices
- List WAN technologies
- Compare circuit switched and packet switched networks
- Explain packet switched technologies
- Define remote access
- Explain when to use SLIP
- Explain when to use PPP
- Explain when to use PAP
- Explain when to use CHAP
- Explain when to use EAP
- Define the Ethernet
- Explain how a Token Ring works
- Explain how X.25 works
- Explain how Frame Relay works
- Explain how SDLC works
- Explain how HDLC works
- Explain how LAPB works
- Explain how ISDN works
- Explain how xDSL works
- Explain how Cable Modems works
- Explain how VPN works

Outline

The lesson contains these topics:

- Data Topologies
- Physical Media Characteristics
- Coaxial
- Twisted Pair
- Fiber Optics
- LAN Signaling Types
- LAN Transmission Protocols
- LAN Transmission Methods
- Network Topologies

- Bus
- STAR
- Ring
- MESH
- LAN Media Access Methods
- Transmission Types
- LAN Devices
- WAN Devices
- WAN Technologies
- Circuit Switched vs. Packet Switched Networks
- Packet Switched Technologies
- Remote Access
- SLIP
- PPP
- PAP
- CHAP
- EAP
- Ethernet
- Token Ring
- X.25
- Frame Relay
- SDLC
- HDLC
- LAPB
- ISDN
- xDSL
- Cable Modems
- VPN

Data Topologies

This topic discusses LAN and WAN technologies.

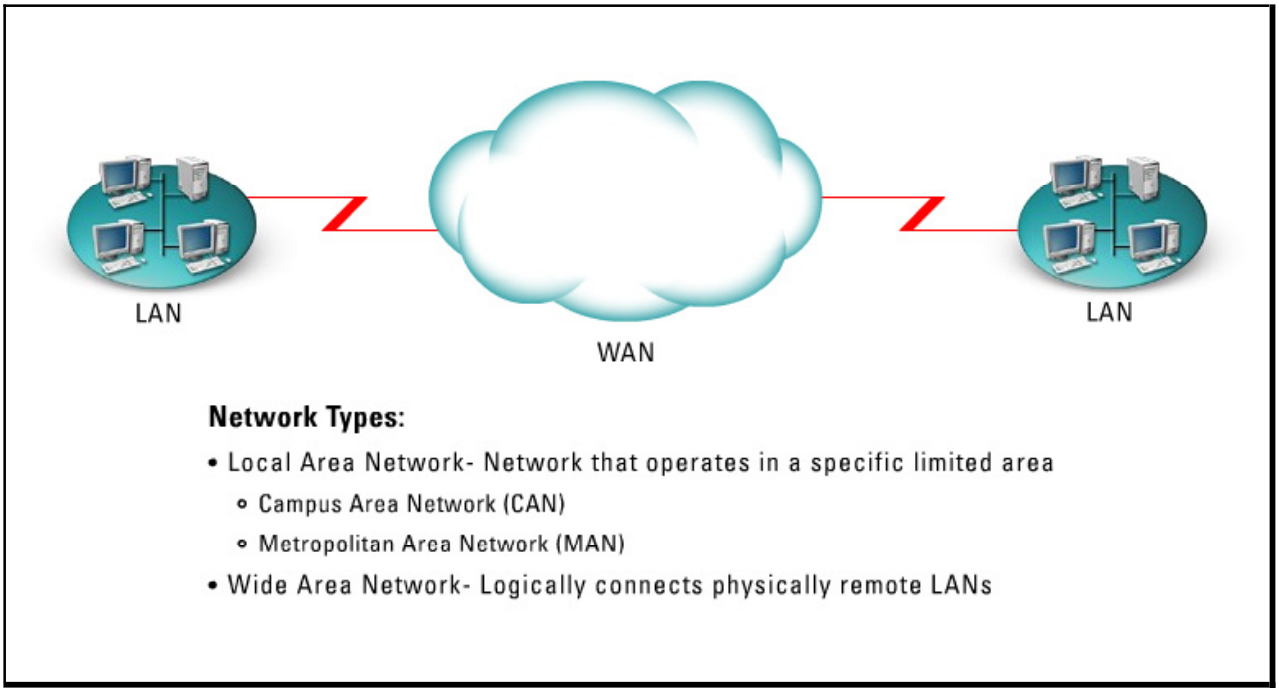

LAN, Local Area Network - A discrete network designed to operate in a specific limited area. LAN has two common types:

- **CAN**, Campus Area Network, is a typically large campus network that connects multiple buildings with each other across a high performance, switched backbone on the main campus.
- **MAN**, Metropolitan Area Network, while not often used as a description, is essentially a LAN that extends over a city-wide metropolitan area.

WAN, Wide Area Network - A network of subnetworks that physically or logically interconnects LANs over a large geographic area. A WAN is basically everything outside a LAN.

Physical Media Characteristics

Physically, few types of media can contain the data being transmitted.

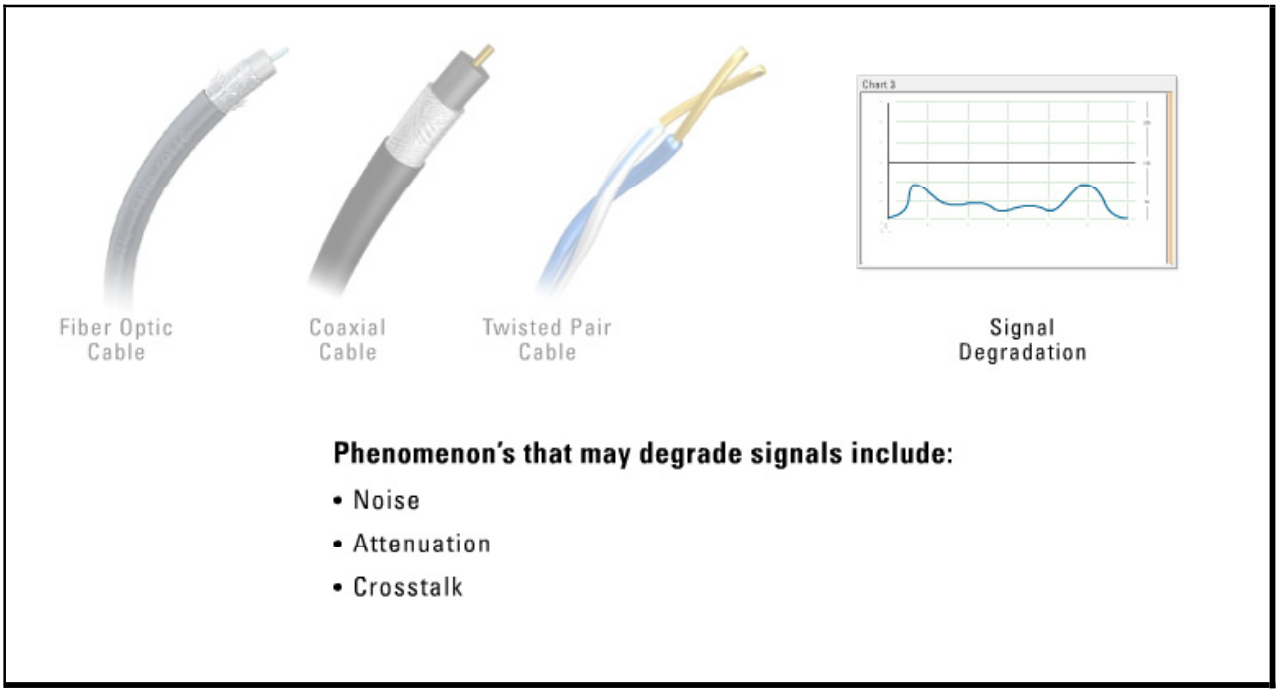

The media include the following:

- **Fiber Optic** - Refers to the medium and the technology associated with transmission of information as light impulses along a glass or plastic medium. Pros: very high transmission speeds, very long distances, hard to wire tap. Cons: Expensive, difficult to install.
- **Coaxial Cable** - Uses a copper wire, shielded for electromagnetic interference and the elements. Typically used in older thinnet networks. Coaxial cable in the data world uses a baseband technology, which effectively limits traffic rates to half duplex. Pros: high resistance to EMI, higher bandwidth, longer cable runs. Cons: Bulky, more expensive, limited to half duplex
- **Twisted Pair** - Ordinary copper wire that connects home and business computers together. Two types are Shielded twisted pair (STP) and Unshielded twisted pair (UTP). Pros: less expensive, easier to work with, can be configured for full duplex (switches). Cons: less resistant to EMI, shorter cable runs.

Each physical medium is susceptible in one degree or another to natural phenomenon that degrades the signal being sent:

- **Noise** - The receiving end will not receive the data in the form that was originally transmitted. Can be caused by motors, computers, copy machines, florescent lightning, and microwave ovens.
- **Attenuation** - The loss of signal strength as it travels or the loss caused by cable breaks and cable malfunctions.
- **Crosstalk** - When electrical signals of one wire spill over to another wire. UTP is much more vulnerable to this than STP or coaxial.
- **Plenum** - Network cabling placed in an area to meet a specific fire rating. Plenum ensures that no harmful chemicals are released in the case of a fire.

- **Pressurized conduits -** Encapsulation of wires, so an attempt to access a wire will change the pressure of the conduit and sound an alarm or send a message to the administrator.

Coaxial

Coaxial cable is a copper conductor surrounded by a shielding layer and grounding wire encased in a protective outer jacket.

Coaxial Cable

Types of coaxial cable:

- 50 ohm cable – Ethernet
- 75 ohm cable – Video

Cable specifications:

- 10Base2 (thinnet)- RG8 or RG58
  - 10Mbps
  - baseband
  - 200 meters (185 actually)
- 10Base5 (thicknet)- RG-6, RG-11, or RG-59
  - 10Mbps
  - baseband
  - 500 meters

Because of the extra shielding, coaxial cable is more resistant to electromagnetic interference (EMI), which allows for high bandwidth usage and longer cable runs. Two types of coaxial cable used in data networks:

- 50 ohm cable - Very widely used in Ethernet networks
- 75 ohm cable - An international standard, used for all video applications

Coaxial cable types include the following:

- 10Base2 or thinnet - This cable is 10Mbs, baseband, with a maximum distance of 185 meters. It uses the RG8 or RG58 cable.
- 10Base5 or thicknet - This cable is 10Mbs, baseband, with a maximum distance of 500 meters. It uses the RG6, RG11, or RG-59 cable.

Twisted Pair

Twisted pair cabling is cabling that has multiple pairs of copper wires twisted around each other, surrounded by an outer protective jacket.

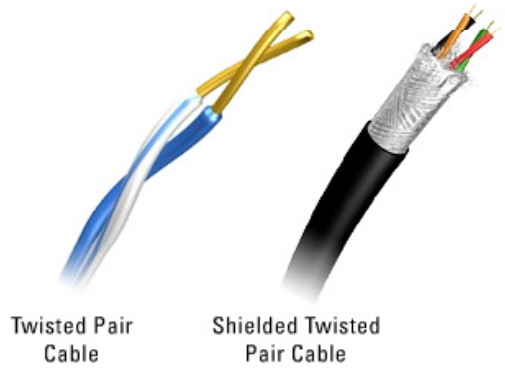

Twisted Pair Cable

Shielded Twisted Pair Cable

Categories of twisted pair includes:

- Category 1 - Standard telephone wire
- Category 2 - Data to 4Mbps
- Category 3 - Data to 10Mbps
- Category 4 - Data to 16Mbps (token ring)
- Category 5 - Data to 100Mbps
- Category 6 - Data to 155Mbps
- Category 7 - Data to 1Gbps

When the outer protective jacket is a foil shielding, we call this type of cable, Shielded Twisted Pair (STP). Without the foil shielding is called Unshielded Twisted Pair (UTP). The twists in each pair of cables is very important because the greater the twist, the better the EMI protection to. UTP has different category ratings based on the number of twists, type of insulation, quality of the copper, and the shielding of the wire.

The different categories of Twisted Pair wire are listed below:

Category 1 - Standard telephone communications wire

Category 2 - Data usage to 4Mbs

Category 3 - Data usage to 10Mbs (10Base-T)

10Mbps, baseband, Twisted pair, 100 meters

Category 4 - 16Mbps (Token Ring)

Category 5 - Data usage to 100Mbps (100Base-TX and FDDI)

100Mbps, baseband, Twisted pair, 100 meters

Category 6 - Data usage to 155Mbps

Category 7 - Data usage to 1Gbps

Fiber Optics

Fiber Optic cables can carry a heavy load more easily than copper cables. They are long, thin strands of very pure glass about the diameter of a human hair.

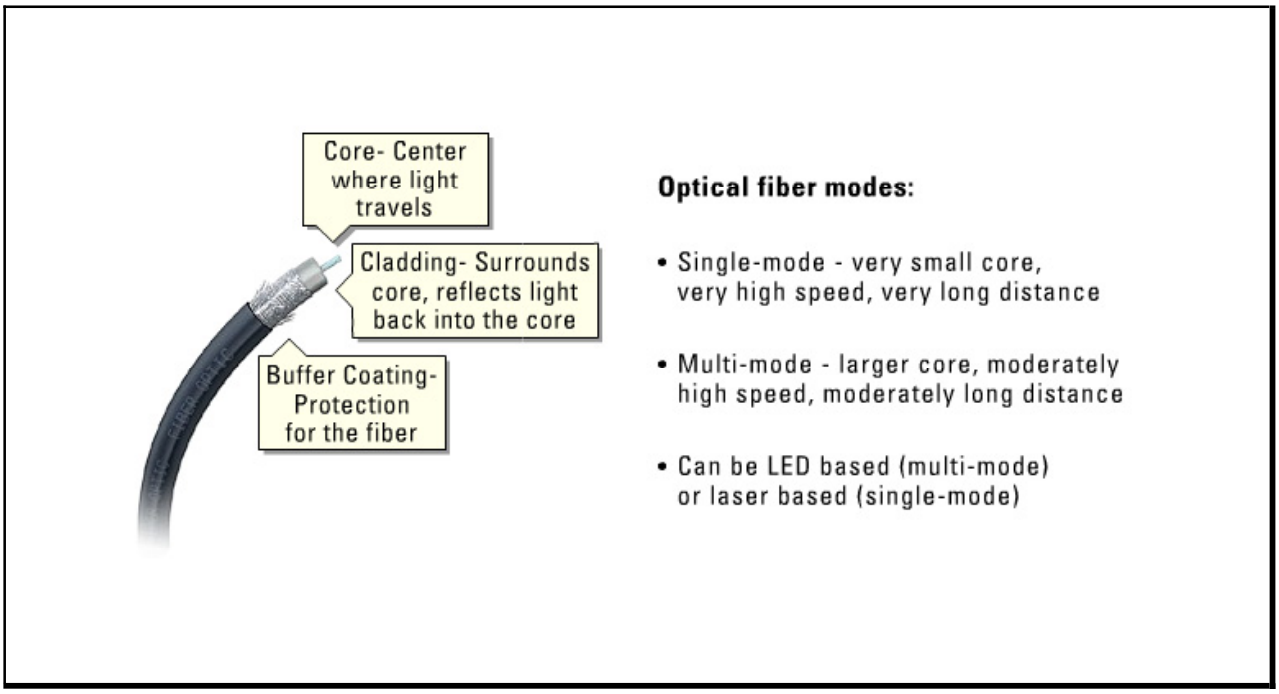

Fiber optics is most commonly used for infrastructure backbones, server farms, or connections that need large amounts of bandwidth. The main drawbacks are cost and the high level of expertise needed to install.

If you look closely at a single optical fiber, you will see the following parts:

- **Core** - Thin glass center of the fiber where the light travels
- **Cladding** - Outer optical material surrounding the core that reflects the light back into the core
- **Buffer coating** - Plastic coating that protects the fiber from damage and moisture

Hundreds or thousands of these optical fibers are arranged in bundles in optical cables. The bundles are protected by the cable's outer covering, called a jacket.

Optical fibers come in two types:

- **Single-mode fibers** - Have small cores (about 9 microns in diameter)
- **Multi-mode fibers** - Have larger cores (about 62.5 microns in diameter)

Some optical fibers are made from plastic and have very large cores (one mm in diameter).

Fiber optics work on the principle called total internal reflection. Light bounces inside the core along the length of the cable. In multi-mode fibers, some reflections bounce up and down the cable in very tight angles, while others might bounce at much larger angles. Tight angle reflections take more time to reach the destination than larger angle reflections. Because of this phenomenon, we can say that the same signal can have multiple modes of reflection, each appearing at the destination at slightly different times. This

multi-mode problem limits the maximum length that data can travel without being completely distorted by its own reflections.

Single mode fibers because of their small diameter only allow a single mode of reflection in the core. Because only one reflection travels down the core, the distances used with single mode can be much greater than those of multi-mode fiber. Optical transmitters are either LED based or laser-based. Lasers provide much more bandwidth and can travel longer distances. LEDs do not provide as much bandwidth, cannot travel as long, but are much cheaper.

LAN Signaling Types

Data can be transmitted in a myriad of ways over many different mediums. Certain medium can handle only a single stream of data at one time. Only one station can transmit data onto the cable at any one time; all other stations receive the data transmitted and accept it if it is directed towards that node. This type of transmission system is called a baseband system.

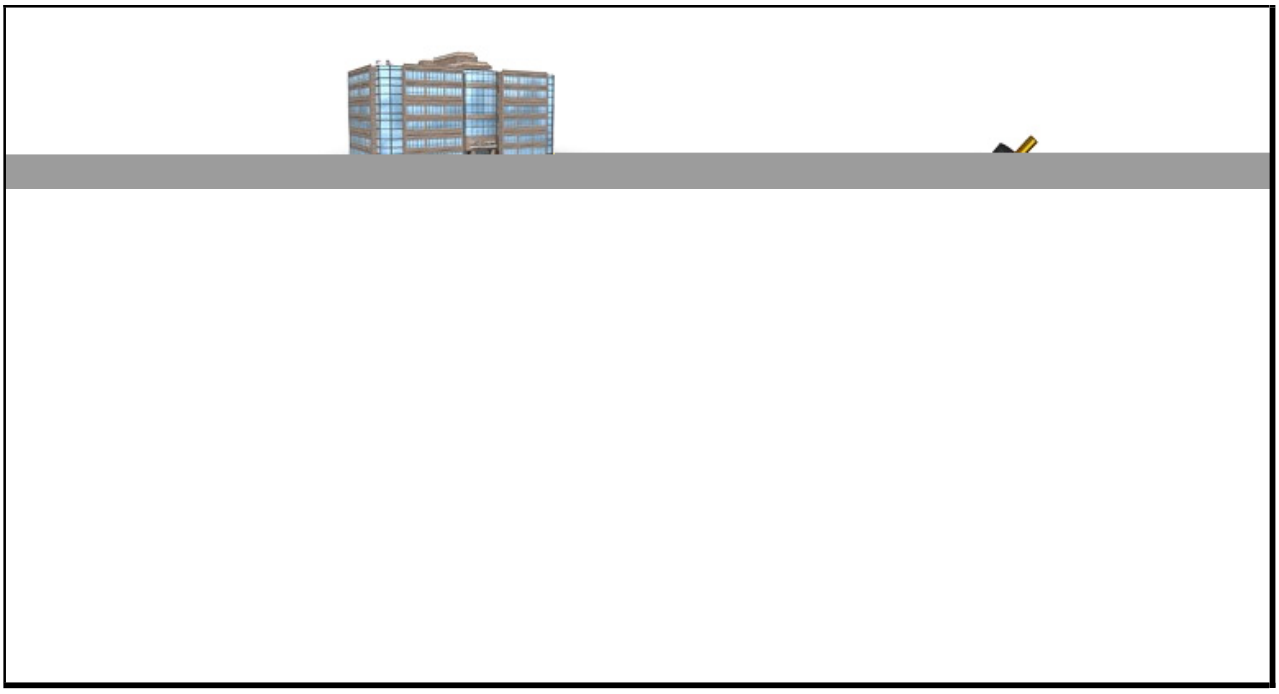

Ethernet and Token Ring are examples of baseband transmission systems. Broadband systems on the other hand can carry more than one transmitted signal on the medium at one time. They do this by creating different channels that are logically separated in the medium. Cable TV is a good example of a broadband system. One cable provides many different channels. Examples of broadband systems in the data world include Digital Subscriber Line (DSL) and cable modems.

- **Baseband** is a digital signal, serial bit stream
- **Broadband** can be an analog signal or cable TV technology

LAN Transmission Protocols

LAN transmission protocols are the rules for communication between computers on a LAN.

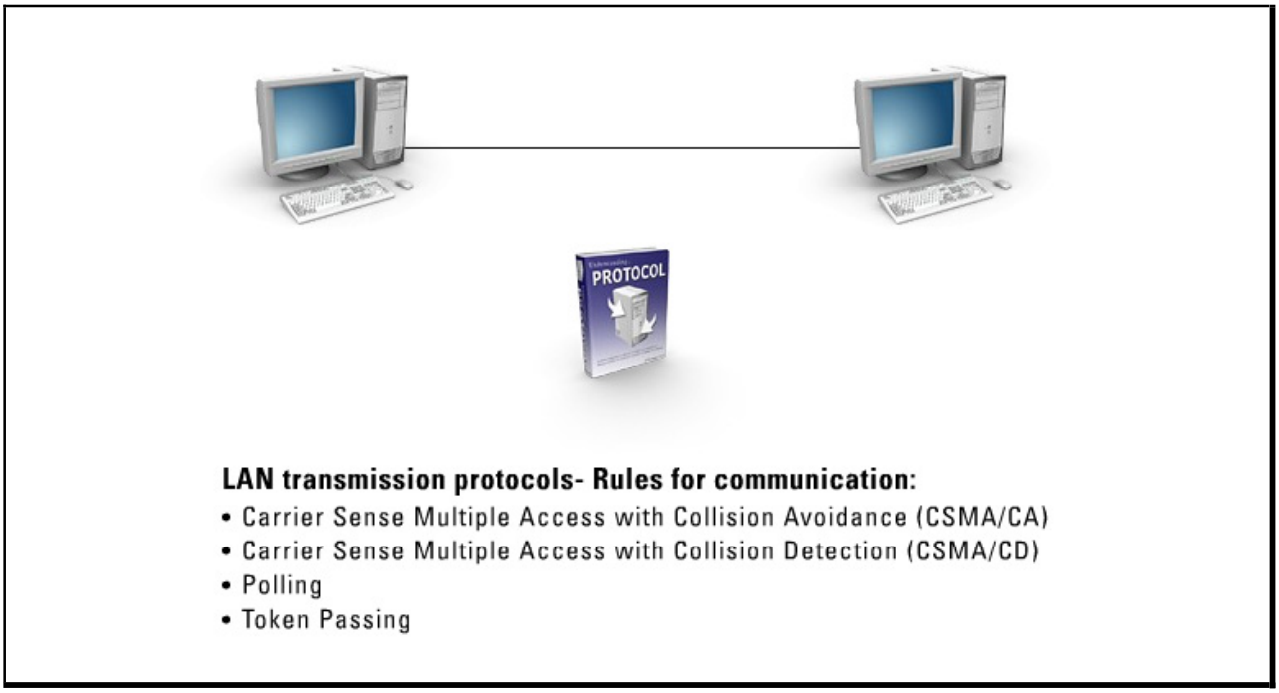

These rules oversee the various steps of communication, such as formatting the data frame, timing and sequencing a packet delivery, and resolution of error states.

- **Carrier Sense Multiple Access** - Is the foundation of the Ethernet communications protocol. It has two functional variations: CSMA/CA (Collision Avoidance) and the Ethernet standard, CMSA/CD (Collision Detection). . In CSMA, a workstation continually monitors the medium while waiting to send a packet then transmits the packet when it thinks the line is free. If the workstation does not receive an acknowledgment from the destination to which it sent the packet, it assumes a collision has occurred and it resends the packet. This principle is defined as persistent carrier sense. Another version of CSMA is called no persistent carrier sense where a workstation waits a random amount of time before resending a packet, thus resulting in fewer errors.
- **Carrier-Sense Multiple Access with Collision Avoidance** - In this variation of CSMA, workstations are attached by two coaxial cables. Each coaxial cable carries data signals in only one direction. A workstation monitors its receive cable to see if the carrier is busy. It then communicates on its transmit cable if no carrier was detected. Thus, the workstation transmits its intention to send when it feels the line is clear because of a precedence based upon pre-established tables. Pure CSMA does not have a feature to avoid the problem of one workstation dominating a conversation.
- **Carrier-Sense Multiple Access with Collision Detection** - Under the Ethernet CSMA/CD media-access process, any computer on a CSMA/CD LAN can access the network at any time. Before sending data, CSMA/CD hosts listen for traffic on the network. A host wanting to send data waits until it does not detect any traffic before it transmits. Ethernet enables any host on a network to transmit whenever the network is quiet. In addition, the transmitting host also constantly monitors the wire to make sure that no other hosts begin transmitting. If the host detects another signal on the wire, it then sends out an extended jam signal that causes all nodes

on the network segment to stop sending data. These nodes respond to that jam signal by waiting a small random amount of time before attempting to transmit again. CSMA/CD was created to overcome the problem of collisions that occur when packets are simultaneously transmitted from different nodes. Collisions occur when two hosts listen for traffic, and upon hearing none, both transmit simultaneously. In this situation, both transmissions are damaged and both hosts must retransmit at a later time.

- **Polling** - In the polling transmission method, a primary workstation checks a secondary workstation regularly at predetermined times to see if it has data to transmit. Secondary workstations are not permitted to transmit until they are given permission by the primary host. Polling is commonly used in large mainframe environments where hosts are polled to see if they need to transmit. Because polling is very inexpensive, networks that are low-level and peer-to-peer typically use it.
- **Token Passing** - Used in Token Ring, FDDI, and Attached Resource Computer Network networks, stations in token passing networks cannot transmit until they receive a special frame called a token. This arrangement prevents the collision problems that are present in CSMA. Token passing networks will work if large bandwidth consuming applications are commonly used on the network. Token Ring and IEEE 802.5 are two principal examples of token-passing networks. Token passing networks move a small frame called a token around the network. Possession of this token grants the right to transmit. If a node receiving the token has no info to transmit, it passes the token to the next node. Unlike CSMA/CD networks, such as Ethernet, token passing networks are deterministic and calculate the maximum time that will pass before any end station will be able to transmit.

LAN Transmission Methods

There are three basic LAN transmission methods used in today's networks: unicast, broadcast, and multicast.

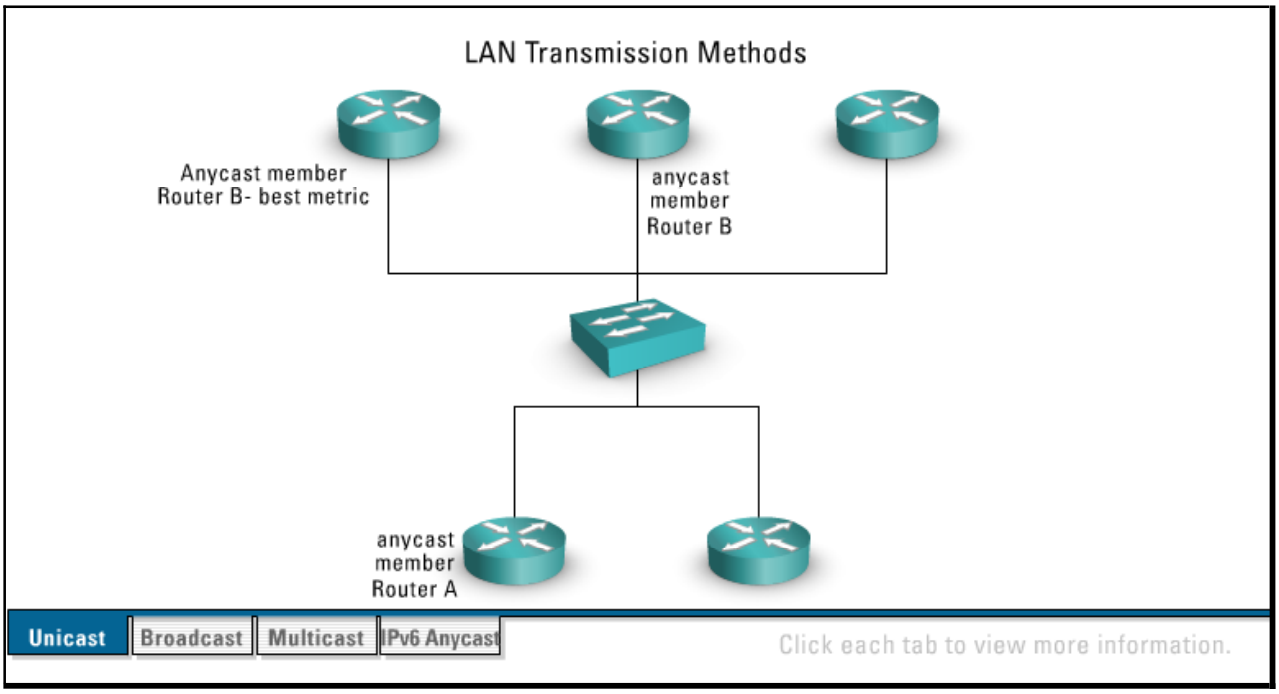

Unicast - The packet is sent from a single source to a single destination address. All LANs such as Ethernet, Token Ring, frame relay, and ATM, and IP networks support the unicast transfer mode.

Broadcast - In this communication method, communication takes place between one point and all other points. The packet sent is received and accepted by all other nodes in the network. An issue with this transmission mode is that each node must spend valuable CPU time reading the contents of each transmission.

Multicast - The source packet is copied and sent to specific multiple destinations on the network. In this communication mechanism, multiple sources send packets to multiple receivers. Only those stations that have signaled to the network their wish to receive the packet will receive it.

IPv6, the next generation of IP, has no support for the broadcast transmission method. Instead a new flavor of transmission call the Anycast method is deployed:

- An IPv6 anycast address is an address assigned to more than one interface, typically belonging to different nodes.
- A packet sent to an anycast address is routed to the nearest interface having that address, according to the routing protocols' measure of distance.
- Anycast addresses are allocated from the unicast address space, using any of the defined unicast address formats.
- An assigned anycast address has a longest address prefix P that identifies the topological region in which all interfaces belonging to that anycast address reside.
- Anycast address allocation is handled by IANA: http://www.iana.org/assignments/ipv6-anycastaddresses

Network Topologies

This topic discusses typical LAN network topologies.

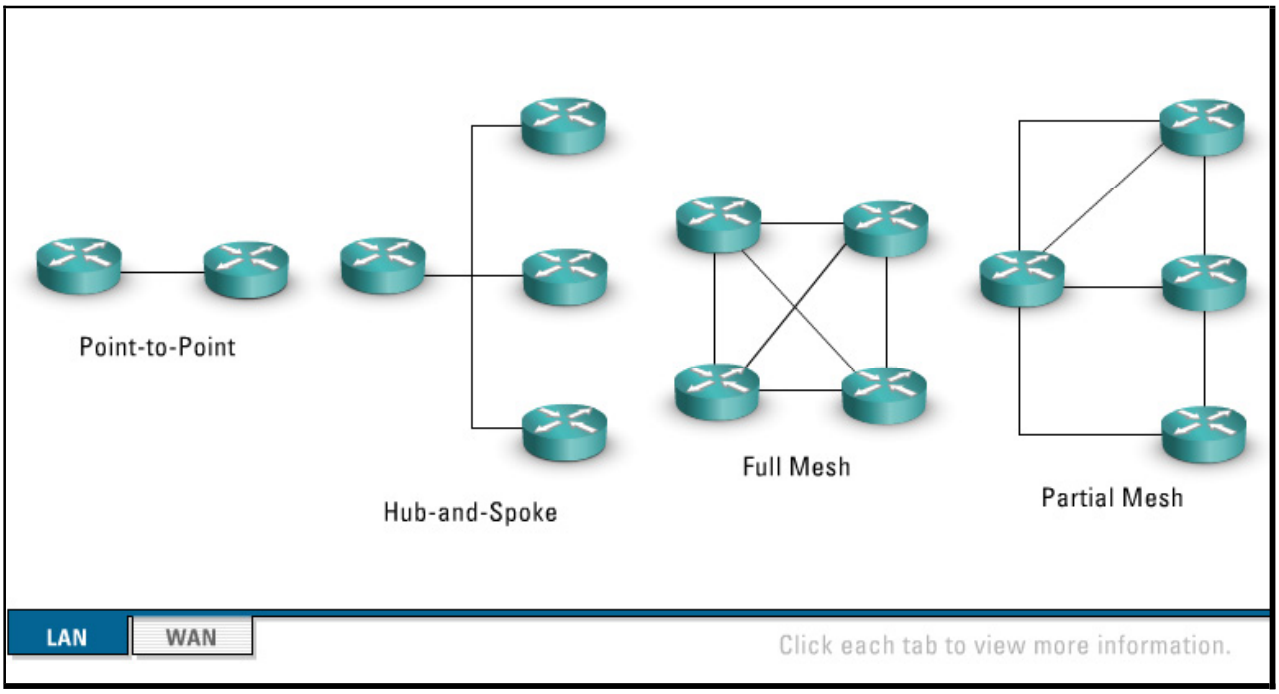

Bus - In this type of topology, a single cable runs the entire distance of the network. Nodes connect to the network by attaching drop points on the cable. In this topology, data transmitted across the cable is received by all nodes attached to the cable. Bus has two types of topologies:

- **Linear bus** - A single cable running linearly across the network
- **Tree** - The single cable divides into branches containing nodes

Ring - In this type of topology, a series of nodes connect in a closed loop.

Star - In this type of topology, all nodes connect to a central switch.

Mesh - In this type of topology, the network's purpose is greater redundancy. To create more redundancy, more connections are required within the network, essentially creating loops in the network. The drawbacks of the mesh network are added complexity to handle the loops To accommodate the loops, more intelligence is required in the network for a loop prevention mechanism The loop prevention is implemented at Layer 2 with spanning-tree (802.1D), or at Layer 3 routing protocols.

Typical WAN network topologies include the following topologies:

- **Point-to-Point** - In this type of topology, a single node connects directly to another node.
- **Hub and Spoke** - In this type of topology, a single node connects to multiple nodes.
- **Full Mesh** - In this type of topology, every node is connected to every other node.
- **Partial Mesh** - In this type of topology, some nodes connect to more than one other node.

Bus

In a bus topology, all transmissions of the network nodes travel the full length of the cable and all stations receive each transmission. This means the physical media is shared between devices.

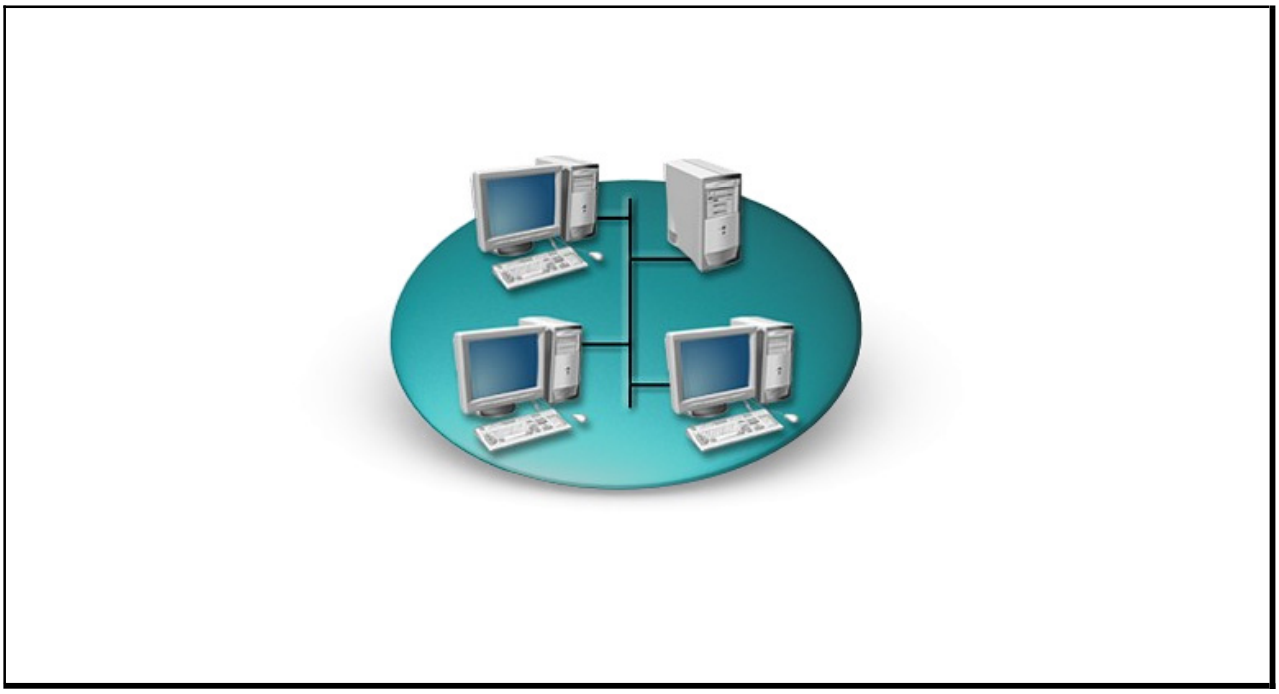

When multiple devices attempt to access the network at the same time, some method must be used to prevent a collision (CSMA/CD). Ethernet primarily uses this topology. This topology does have some faults. For example, if any station on the bus experiences errors an entire bus can cease to function.

Star

In a star topology, the nodes of a network are connected to a central LAN device, a hub or switch, directly.

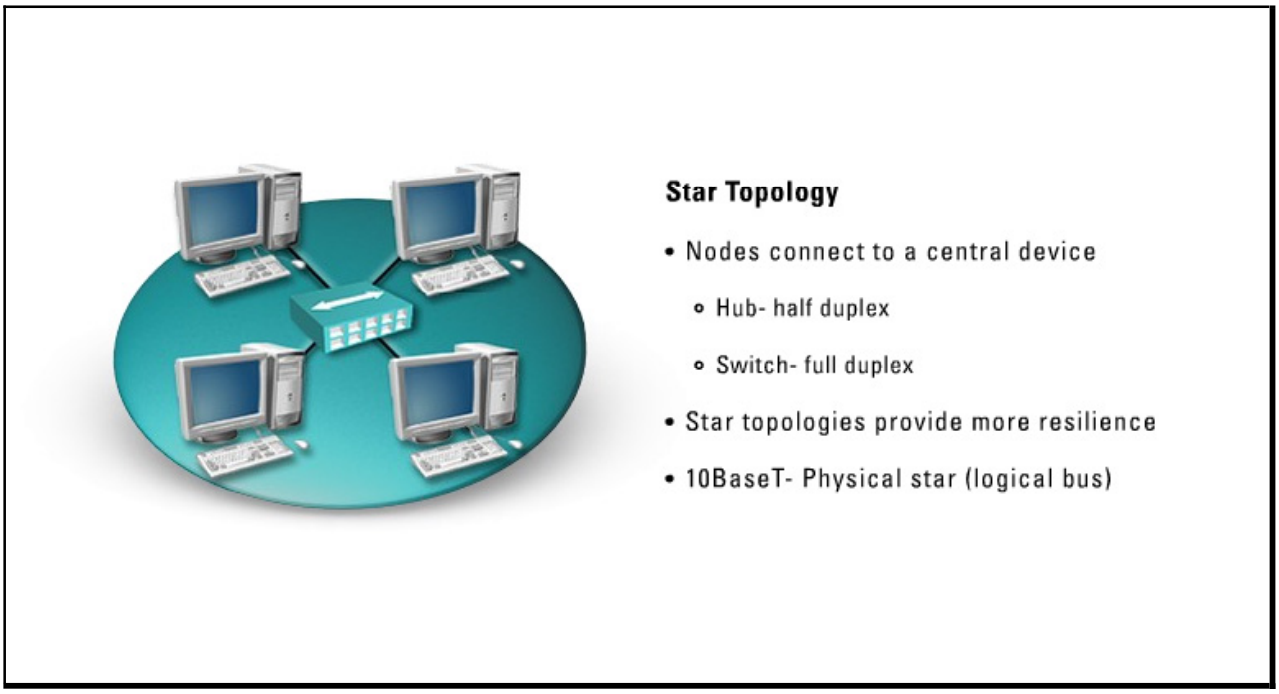

Several logical bus and ring topologies can be configured in a STAR topology. Although Ethernet is logically thought of as a bus topology since its first implementations were thinnet and thicknet on a bus, 10BaseT is actually wired as a Star topology providing more resilience for the entire topology, especially when a station experiences errors. Physically, 10BaseT is wired as a Star, but logically it remains an Ethernet bus.

Ring

In a ring topology, the network nodes are connected by unidirectional transmission links to form a closed loop.

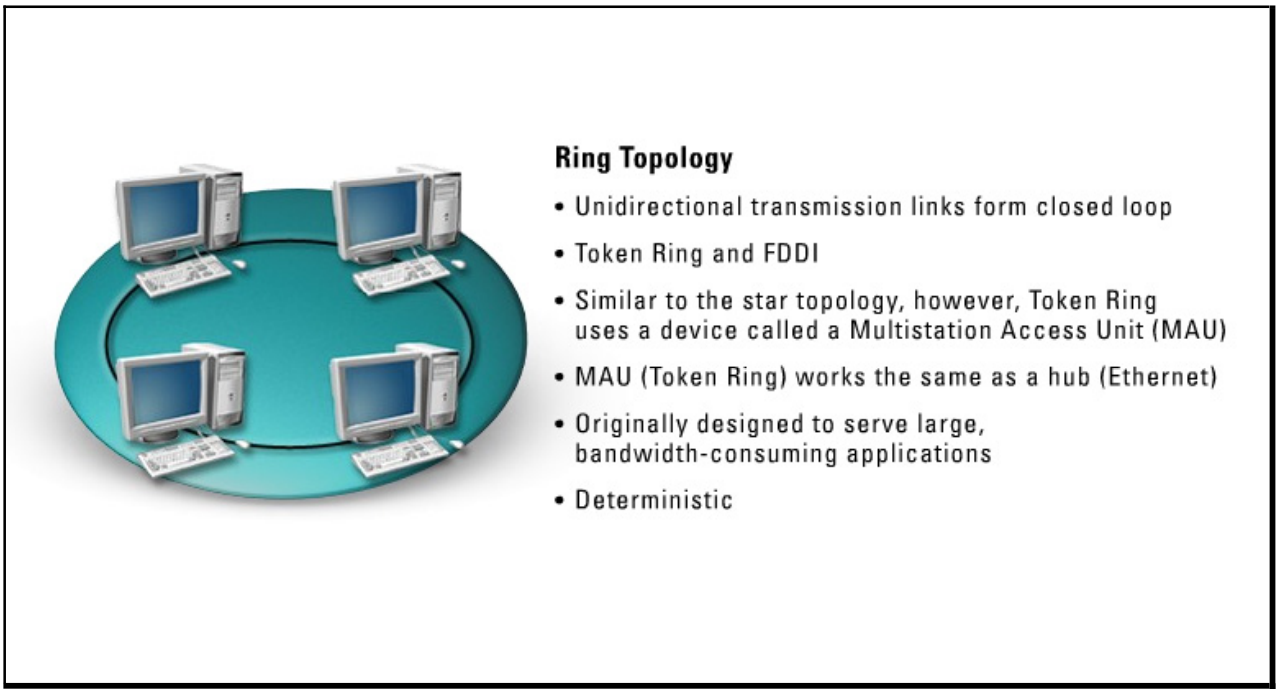

Token Ring and FDDI both use this topology. Ring topologies are similar to bus topologies, except they transmit in one direction only from station to station. Typically, ring architectures use separate physical ports and wires for transmit and receive.

Mesh

In a mesh topology, all the nodes connect to all other nodes in a network. This topology may be used to create backbone redundant paths.

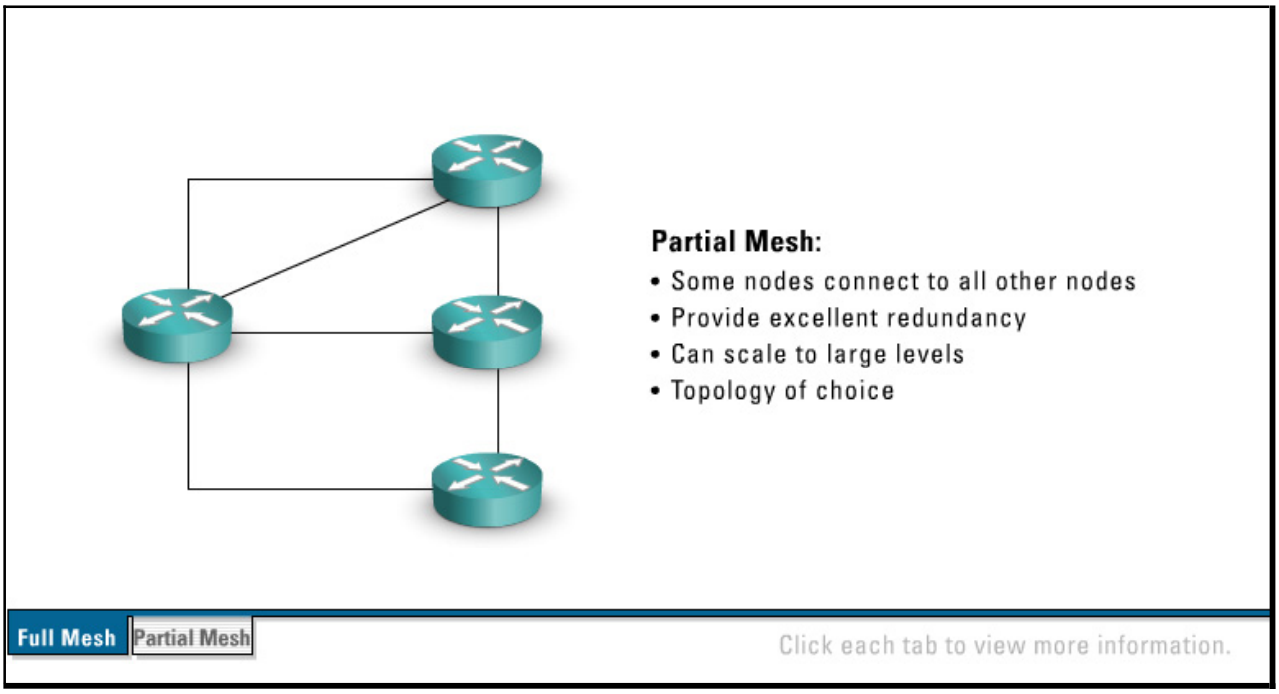

A full mesh topology has every node connected to every other node. A partial mesh topology may connect multiple full mesh networks together. Full mesh networks provide the greatest amount of redundancy in the network, but cannot scale to a very large level. Partial mesh networks can scale to large levels, provide excellent redundancy, and are the topology of choice in WAN designs.

LAN Media Access Methods

LAN media access methods control the use of a network in its Physical and Data Link Layers.

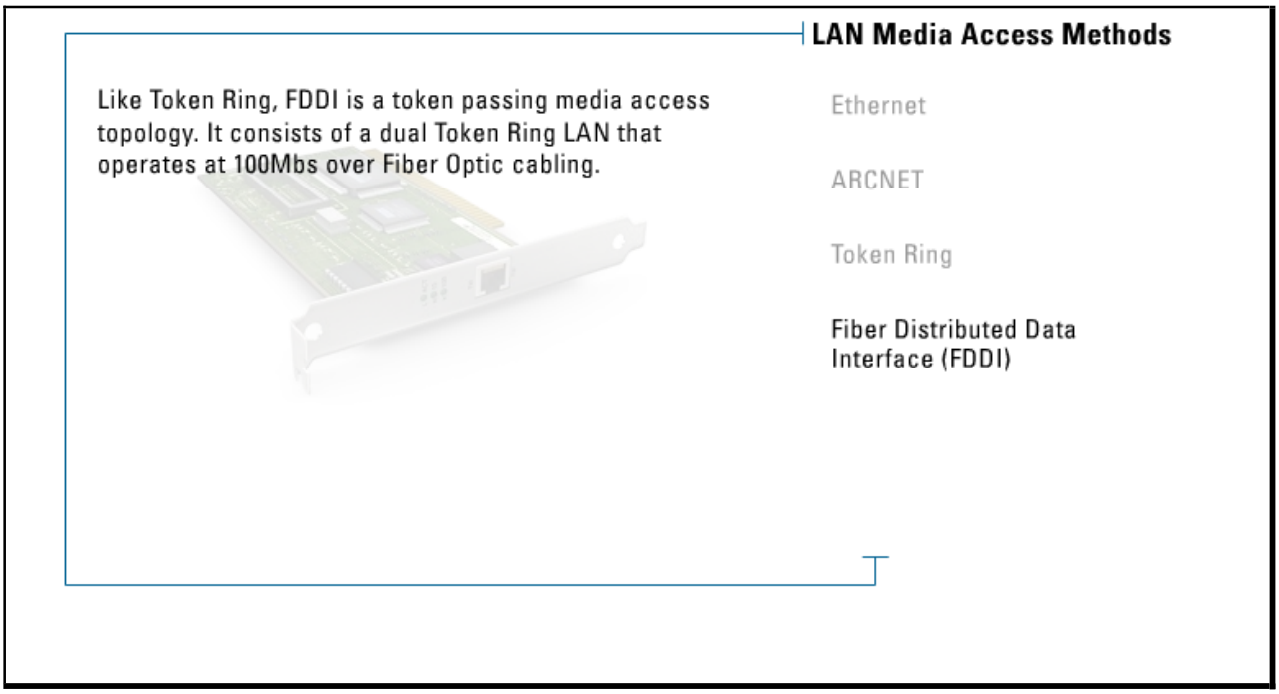

Ethernet - The Ethernet media access method transports data to the LAN using CSMA/CD. Currently, Ethernet often refers to all CSMA/CD LANs. Ethernet was designed to serve in networks with bursty, occasionally heavy traffic requirements.

Ethernet defines a bus topology LAN with three cable standards:

- Thinnet, also known as 10base2, is a coaxial cable with segments of up to 185 meters.
- Thicknet, also known as 10base5, is a coaxial cable with segments up to 500 meters.
- Unshielded Twisted Pair, usually in UTP. All hosts are connected using an unshielded twisted pair cable connected to a central device such as a hub or switch. UTP has three common variations: 10baseT, 100baseT, and 1000baseT.

ARCNET - ARCnet is one of the earliest LAN technologies. It uses a token passing access method in a Star technology on coaxial cable. ARCnet provides predictable and slow network performance. An issue with ARCnet is the need to manually enter the node address of each station during installation, thus creating the possibility of duplicate and conflicting nodes.

Token Ring - IBM originally developed the Token Ring network in the 1970s. Token Ring is second only to Ethernet in general LAN popularity. The term, Token Ring, refers to both IBM's Token Ring network and IEEE 802.5 networks. In a Token Ring network, a device called Multistation Access Unit (MAU) connects all end stations. One station on a Token Ring network is designated the Active Monitor. The Active Monitor makes sure no more than one token is on the ring at any given time. If a transmitting station fails, it probably is not able to remove a token as it makes its way back onto the ring. In this case, the active monitor steps in and removes the token and generates a new one.

Fiber Distributed Data Interface (FDDI) - Like Token Ring, FDDI is a token passing media access topology. It consists of a dual Token Ring LAN that operates at 100Mbs over Fiber Optic cabling. FDDI employs a token-passing media access method with dual counter rotating rings. Only one ring, however, is active at any given time. If a break or outage occurs, the ring will then wrap back the opposite direction, keeping the ring intact.

Major advantages:

- Can operate long distances at high speeds with no interference.
- Provides predictable, deterministic delays, and permits several tokens to be present. Copper distributed data interface can be used with a UTP cable to connect servers or other stations into the ring instead of fiber. Hpwever, copper is inferior in length and EMI.

Transmission Types

AM, FM, and Digital are the three methods to encode a transmission signal.

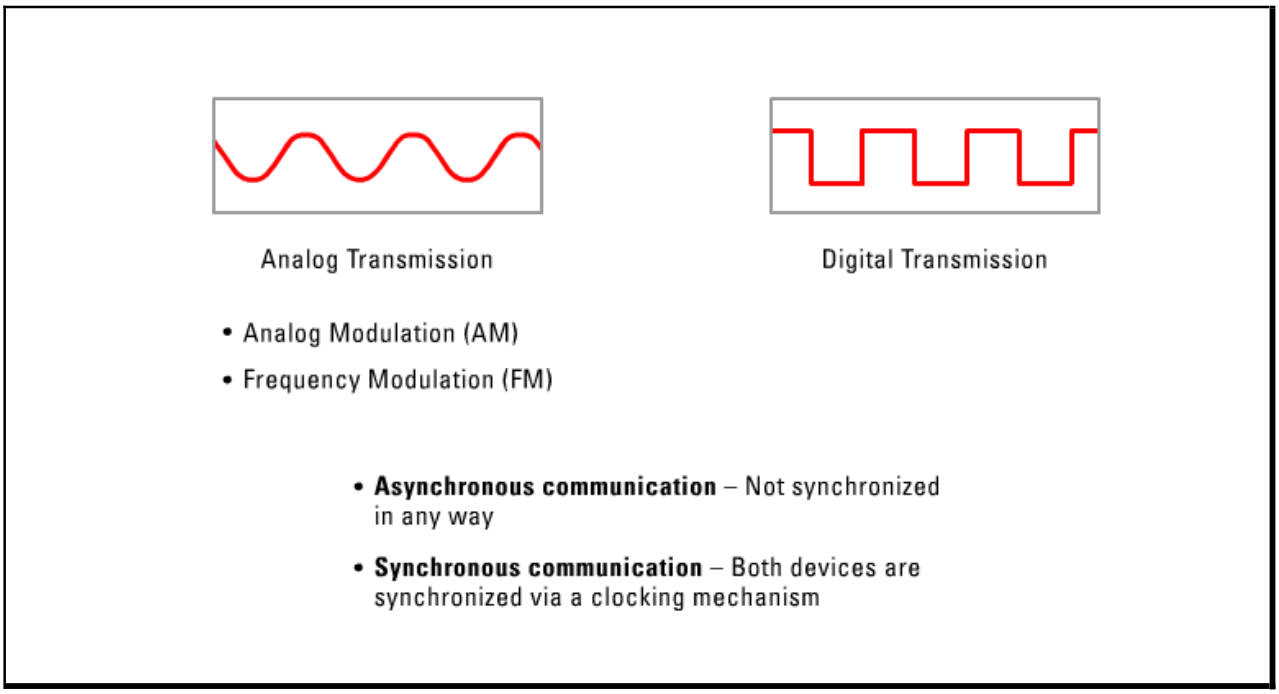

Amplitude modulation (AM) and frequency modulation (FM) are both analog modulation schemes. The third method is digital modulation.

- **Analog transmission signals** - Representing changing values as continuously variable quantities. In analog, information is transmitted by modulating a continuous transmission signal, such as varying its frequency to add or take away data.
  - **Analog modulation** - A transmission technique in which the amplitude of the carrier varies in accordance with the signal.
  - **Frequency modulation** - A method of transmission in which the carrier frequency varies in accordance with the signal.
- **Digital transmission signals** - While analog signals typically vary smoothly and continuously over time, digital signals are present at discrete points in time, that is they are either on or they are off.

Asynchronous communication - Takes place when two devices are not synchronized in any way,; the transmitter and receiver clock are independent of each other and are not synchronized. The sender can send data at anytime and the receiving end must always be ready. An asynchronous link communicates data as a series of characters of fixed size and format. Each character is preceded by a start bit and followed by 1 to 2 stop bits. Parity is often added to provide some limited protection against errors occurring on the link. The use of independent transmit and receive clocks constrains transmission to relatively short characters usually less than eight bits, and moderate data rates of around or less than sixty-four kilobytes per second (64 kbps). The asynchronous transmitter delimits each character with a start sequence and a stop sequence. The start bit (0), data is usually eight bits plus parity, and stop bit(s) are transmitted using a shift register clocked at the nominal data rate.

Synchronous communication - Takes place between two devices that are synchronized, usually with a clocking mechanism. Data transfer is performed as a stream of bits. While asynchronous communication sends smallish blocks of data with many control bits for error correction, synchronous techniques use big blocks of data with control bits only at the start and end of the entire transmission. Because synchronous transmission has minimal error checking, synchronous communicating devices must be timed to perfection, and they need a clean line. A modem used over an analog line will not handle synchronous communication well because a sputter on the line would throw the message out of sync.

LAN Devices

System designers and administrators use many devices on the LAN to enhance communication between nodes.

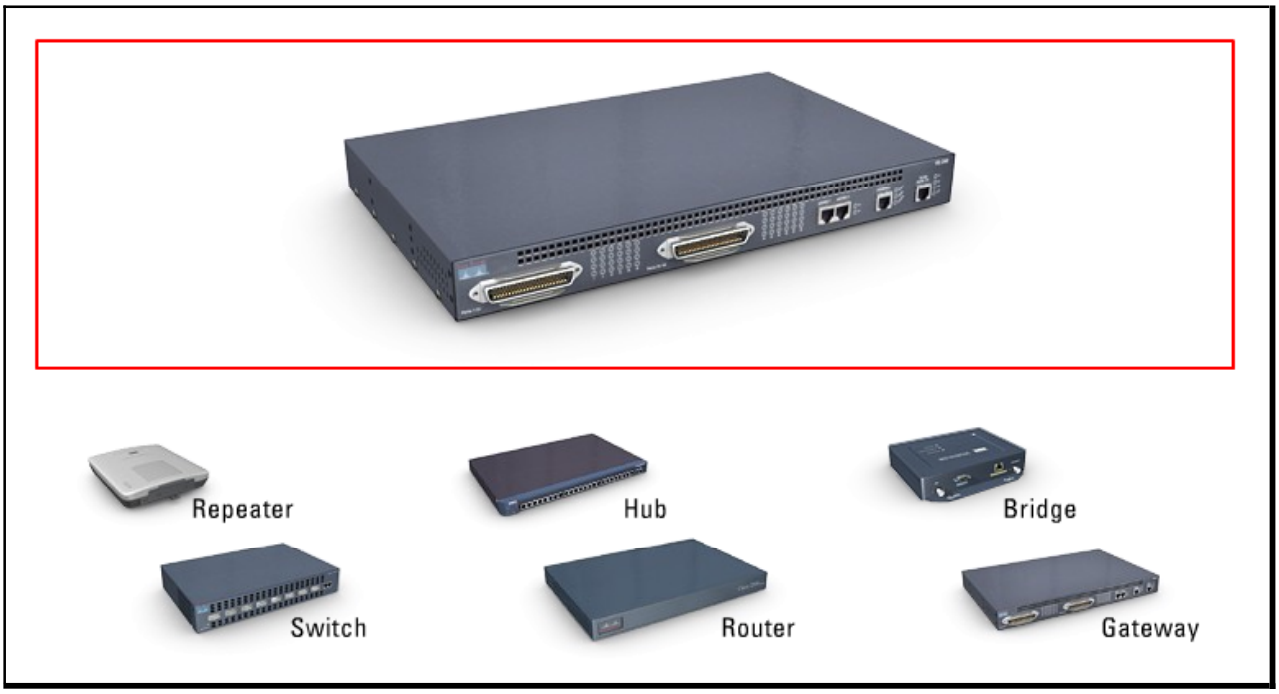

Some of the more basic devices on the LAN include the following:

Repeaters - Simple layer 1 device that extends a signal by amplification. Repeaters have no intelligence.

Hubs - Connect multiple nodes together. Hubs are basically multi-port repeaters.

Bridges - Amplify data signals and add some intelligence to the signal. A bridge forwards the data it receives on one port to all other ports if the MAC address of the destination computer is not on the local network segment. If the destination computer is on the local network segment, it does not forward the data. Bridges operate at Layer 2, which means they are MAC oriented.

Switches - Switches are similar to bridges in a logical view. Bridges typically use software to bridge data, whereas switches use hardware in the form of Application Specific Integrated Circuits (ASICs) to switch data. Because of the additional hardware required, switches are more expensive than bridges. Switches also operate at layer 2, which means they are MAC oriented.

Routers - Routers work at Layer 3 in the OSI model, which means they are IP oriented. They forward packets based on the destination IP address. Routers require knowledge of all LANs to which destination packets will be forwarded. If the destination does not exist in the router's routing table, the packet will be dropped.

Gateways - Gateways work at Layer 7 in the OSI model; they connect unlike media or topologies.

WAN Devices

This topic discusses devices on the WAN.

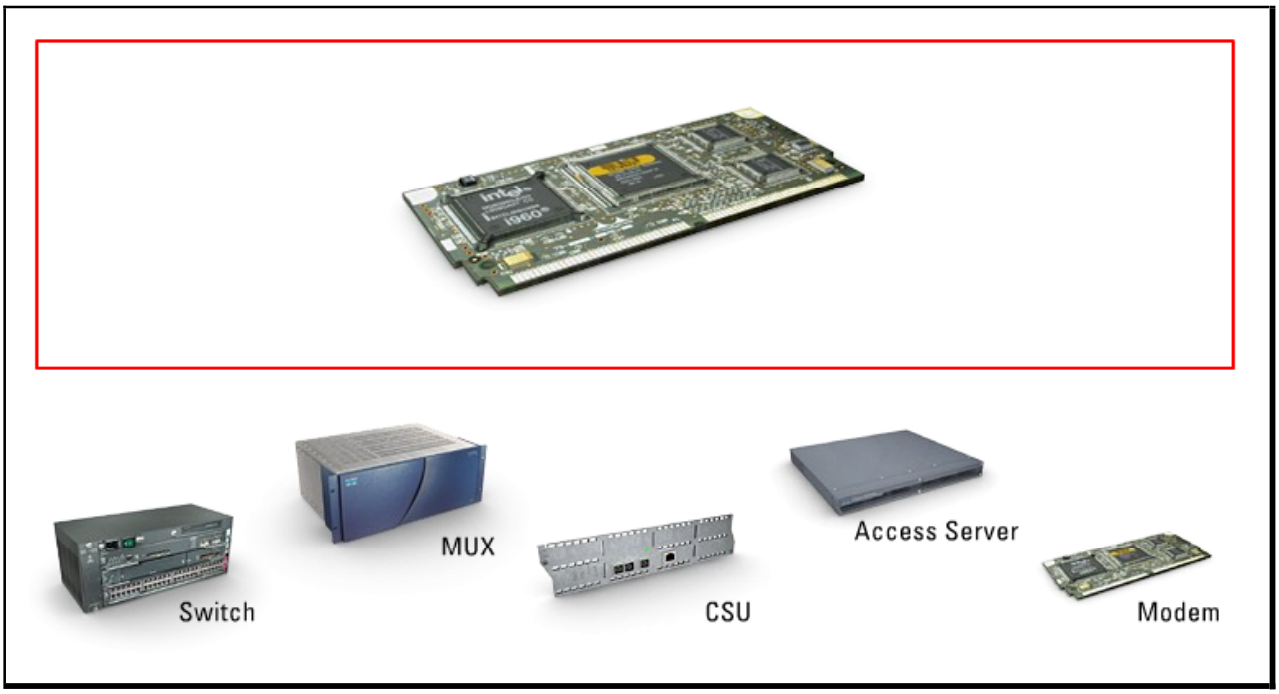

Devices on the WAN can include the following:

Multiplexers - Commonly referred to as a mux; a multiplexer is a device that enables more than one signal to be sent out simultaneously over one physical circuit.

WAN switches - Multiport networking devices that are used in carrier networks. They connect private data over public data circuits by using digital signals. Typically, frame relay switches and ATM switches prevail.

Access Servers - An access server is a server that provides dial-in and dial-out connections to a network.

Modems - A modulator/demodulator (modem) is a device that interprets digital and analog signals thus enabling data to be transmitted over voice grade phone lines. The modem translates digital signals to an analog form that can be transmitted over an analog communications medium. The modem converts these analog signals back to their digital form at the destination.

Channel Service Unit - A digital interface device used to terminate the physical interface on a DTE device to the interface of a DCE device in a switched carrier network. These devices connect to the closest telephone company switch in a central office. The channel service unit supports all co-directional timing patterns, loopback modes, and data and control signals.

WAN Technologies

Like LAN protocols, WAN protocols are the rules for communicating between computers on a WAN. Because the WAN more often than LAN connects networks together, these protocols address the issues involved with communications between many large and disparate networks.

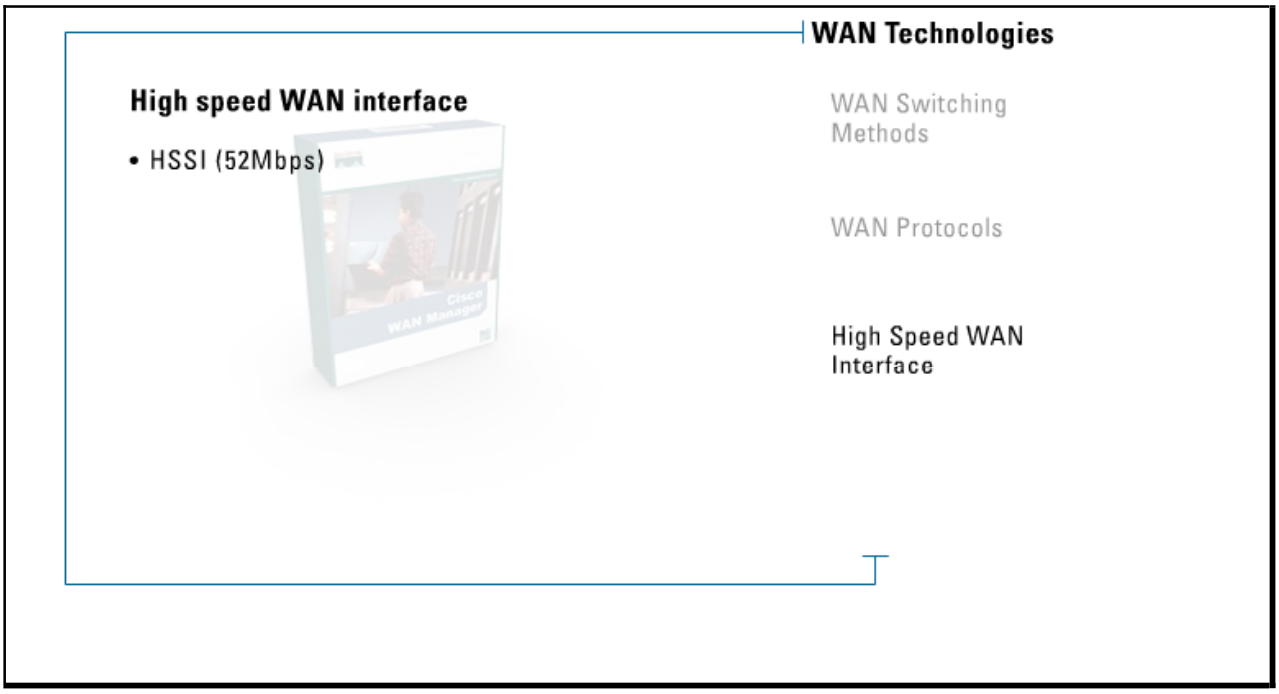

Almost every WAN protocol is designed to run on a specific WAN technology.

- **Circuit switching technology** - Circuit switching sets up a virtual connection between nodes that acts like a dedicated circuit. Once the circuit is up, it is dedicated for use only between the two nodes. After the session ends, the circuit is torn down.
- **Packet switching technology** - In packet switching, the conversation between nodes is broken up in fragments, with each fragment having all information required to get from sender to receiver. The network then acts upon each independent packet.

A **CSU/DSU** - When digital equipment connects a LAN network to a WAN network, the connection needs a Channel Service Unit / Data Service Unit. The DSU converts digital signals to be transmitted over the telephone company's digital lines. The CSU is the unit that connects the network directly to the telephone company's line.

The Customer Premise Equipment (CPE) provides a digital interface for DTE - Data Terminal Equipment.

The Service Provider provides an interface to the DCE - Data Circuit-Terminating Equipment device.

X.25 is an older WAN switching protocol that defines how devices and networks establish and maintain connections. Data in X.25 is divided into 128 bytes and encapsulated in High-level Data Link Control (HDLC) frames. The frames are then addressed and forwarded across the carrier switches. X.25 was created to combat noisy lines and long distances; hence, it had high overhead in error detection and recovery. As lines become more and more stable, the additional overhead attributed to X.25 became a severe bottleneck. Frame relay was created to address the bottleneck issues.

Frame relay is a streamlines version of X.25, offering fewer of the robust capabilities. Frame relay is a WAN protocol that operates at the physical and data link layers. Designed for use across Integrated Services Digital Network (ISDN) interfaces, frame relay is used over a variety of other interfaces as well. It uses packet-switching technology and implements a concept called the CIR (committed information rate). The CIR is the level of service the service provider guarantees to always be available to the customer. Two main types of equipment facilitate frame relay:

- DTE/Data Terminal Equipment which is Customer owned.
- DCE/Data Circuit-Terminating Equipment which can be the service provider's or phone company's

Frame relay provides two types of services:

- PVC or Permanent virtual circuit which works like a private line for a customer with agreed-upon bandwidth availability.
- SVC or switched virtual circuits which require steps similar to a dial-up and connection procedure.

ATM - Asynchronous Transfer Mode is a cell switching technology. In ATM, data is segmented into fixed size cells, 53 bytes, instead of variable-size packets. Five (5) bytes are used for the ATM header, which leaves 48 bytes for the payload. ATM is a high-speed networking technology used for LAN, WAN and service provider connections. ATM uses virtual circuits that act like dedicated paths between the source and destination. These virtual circuits can guarantee bandwidth and QoS.

SMDS - Switched Multi-megabit Data Service is a high-speed, packet-switched, datagram-based WAN technology used to enable customers to extend their LANs across MANs and WANs. It is a connectionless protocol and can provide bandwidth on demand. SMDS can use fiber or copper-based media. It supports speeds of 1.544 Mbps over Digital Signal level 1 (DS-1) transmission facilities or 44.736Mbps over Digital Signal level 3 (DS-3) transmission facilities. SMDS data units are large enough to encapsulate entire IEEE 802.3, IEEE 802.5, and Fiber Distributed Data Interface (FDDI) frames.

HDLC - High-level Data Link Control is a bit-oriented data link layer protocol used for transmission over synchronous lines. Developed by the International Organization for Standardization (ISO), it falls under the ISO standards 3309 and ISO 4335. It supports half duplex and full duplex communication lines, point-to-point, and multi-point networks. It works with primary stations that contact secondary stations to establish data transmission. HDLC is an open standard protocol, but is implemented is slightly different fashions from vendor to vendor. Because of the differences, HDLC implementations are usually incompatible with each other.

SDLC - Networks with dedicated, leased lines with permanent physical connections use Synchronous Data Link Control. .SDLC provides the polling media access technology, which is a mechanism that enables secondary stations to communicate on the network. SDLC is not a peer-to-peer protocol like HDLC, frame relay, or X.25. An SDLC network is made up of a primary station that controls all communications and one or more secondary stations. SDLC is capable of full duplex operation, but nearly all practical applications are strictly half duplex, allowing either the primary or one of the secondaries to transmit at any one time, but never both.

HSSI - High-Speed Serial Interface is a point-to-point DTE/DCE interface developed by Cisco Systems and T3plus Networking. It is used to connect multiplexers and routers to high-speed communication WAN services like ATM and frame relay. HSSI defines both electrical and physical interfaces on DTE and DCE devices. It operates at the physical layer of the OSI reference model. Its maximum signaling rate is 52 Mbps and maximum cable length is 50 feet.

Circuit Switched vs. Packet Switched Networks

This topic differentiates circuit-switched networks and packet-switched networks.

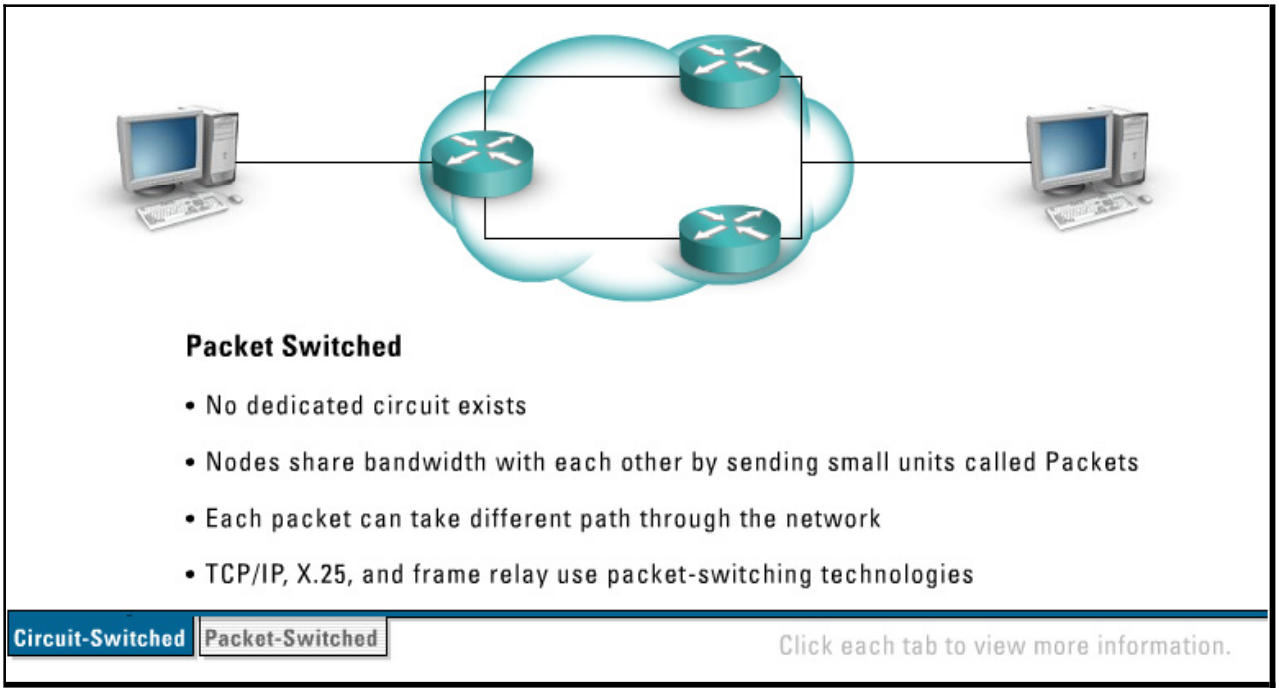

Circuit Switched Networks - Defined as a switching system in which a dedicated physical circuit path must exist between the sender and receiver for the duration of the transmission. A circuit switched network describes a type of WAN that consists of a physical, permanent connection from one point to another. This older technology is the main choice for communications that need to be on constantly and have a limited scope of distribution. Used heavily in telephone company networks, the Public Switched Telephone Network (PSTN) is the best example of a circuit switched network.

Packet Switching Networks - A networking method where nodes share bandwidth with each other by sending small units called packets. Encapsulation and fragmentation occurs to break large messages into smaller packets that are individually switched through the network toward their destination. Individual packets can follow different routes to the destination; however, once all of the packets forming a message arrive at the destination, they are recompiled into the original message. TCP/IP, X.25, and frame relay use packet-switching technologies.

Tip	Circuit-switching systems are ideal for communications that require data to be transmitted in real-time. Packet-switching networks are more efficient if some amount of delay is acceptable.

Circuit-switching networks are sometimes called connection-oriented networks. Although packet switching is essentially connectionless, a packet switching network can be made connection-oriented by using a higher-level protocol. TCP, for example, makes IP networks connection-oriented.

ATM attempts to combine the best of both worlds by guarantying delivery of circuit-switched networks and providing robust and efficient packet-switched networks.

Packet Switched Technologies

Packet switched networks can be far more inexpensive than dedicated circuits because they create virtual circuits which are used as needed.

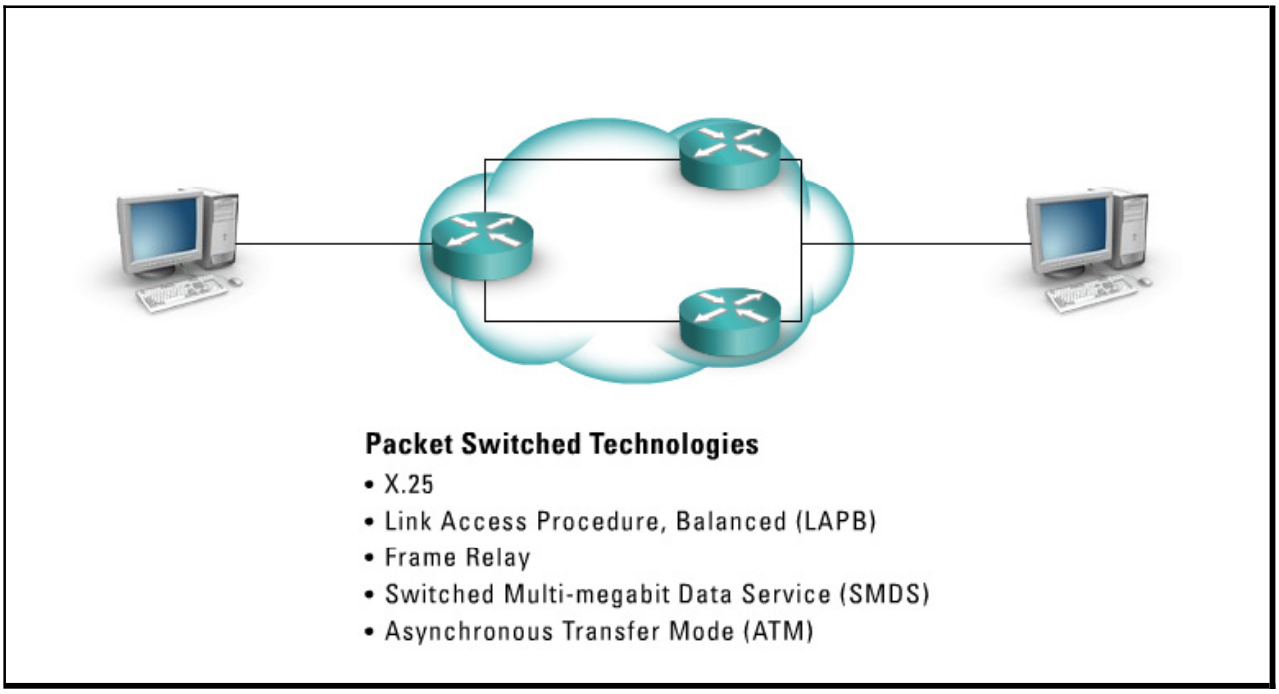

The following are examples of packet switching Networks.

X.25 - The first packet switching network, X.25, defines the point-to-point communication between data terminal equipment and data circuit terminating equipment or a Data Service Unit/Channel Service Unit (DSU/CSU). The DSU/CSU, in turn, supports both switched virtual circuits and permanent virtual circuits. X.25 defines how WAN devices are established and maintained. X.25 was designed to operate effectively regardless of the type of systems that are connected to the network. X.25 is used much more overseas.

Link Access Procedure Balanced - Created for use with X.25, LAPB defines frame types and is capable of retransmitting, exchanging, and acknowledging frames as well as detecting out of sequence or missed frames.

Frame Relay - A high performance WAN protocol that operates at the physical and data link layers of the OSI. Originally designed for use across ISDN interfaces, it is currently used with a variety of other interfaces and is a major standard for high-speed WAN communications. Frame relay is an upgrade from X.25 and LAPB. It is the fastest and requires access to a high quality digital network infrastructure.

Switched Multi-megabit Data Service - SMDS is a high-speed technology that is used over public switched networks. It is ideal for companies that need to exchange large amounts of data with other enterprises over WANs on a bursty or incontiguous basis, by providing connectionless bandwidth upon demand.

Asynchronous Transfer Mode - ATM is a very high bandwidth, low delay technology that uses both switching and multiplexing. It uses a 53 byte, fixed size cells instead of frames like Ethernet. It can allocate bandwidth upon demand, making it a solution for bursty applications.

Remote Access

Remote access is the ability to use dial-up and other Remote Access Service (RAS) technologies to authenticate over a public medium and connect to a corporate LAN.

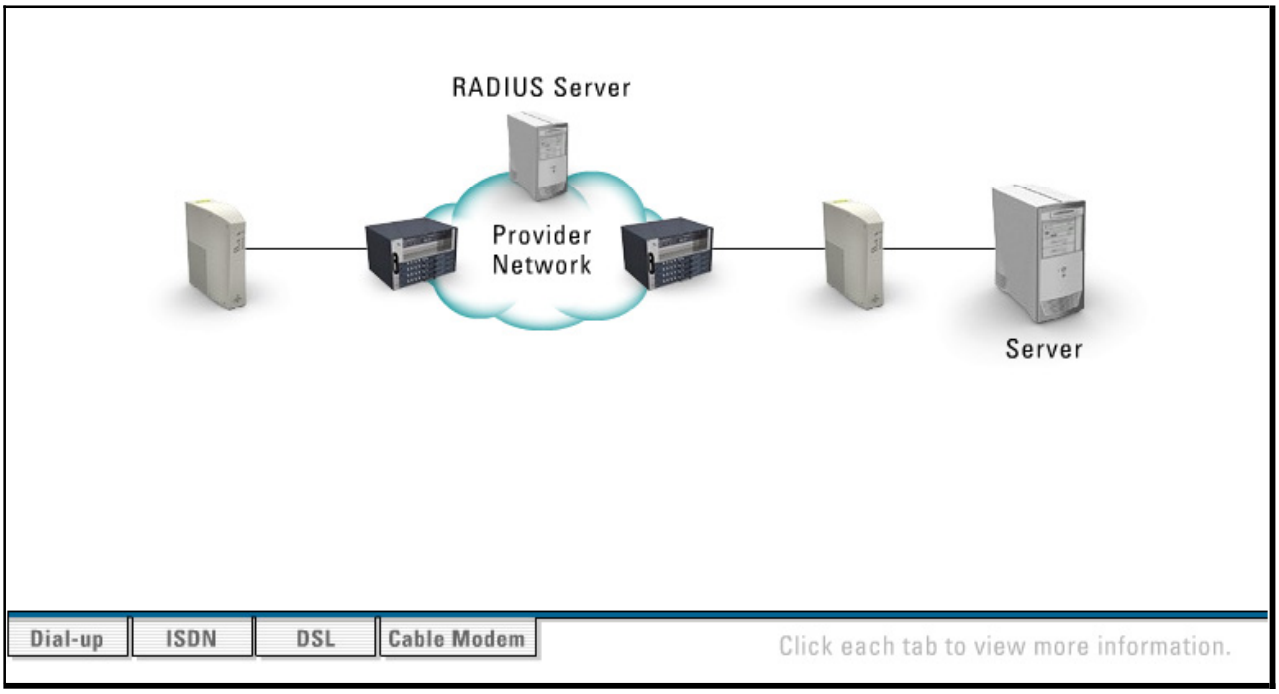

Remote access technologies include the following:

- Dial-up
- Integrated Services Digital Network (ISDN)
- Digital Subscriber Line (DSL)
- Cable modem

Because the user wants access over an untrusted network, authentication, integrity, and encryption are important aspects for security specialists. The most common protocol for these parameters is the Point-to-Point Protocol (PPP). It allows for the authentication of the user and can provide additional services, such as multilink, callback, link fragmentation, etc.

Ethernet does not provide for authentication upon link establishment, which means that Ethernet-based technologies such as DSL and cable modem cannot authenticate who is connecting to their network. To mitigate this problem, DSL and some cable modem providers are using a new protocol, based upon PPP. That protocol is called Point-to-Point Protocol over Ethernet (PPPoE). This marriage allows an Ethernet-based network to support authentication and other PPP features when used over Ethernet.

SLIP

Serial Line Internet Protocol (SLIP) provides point-to-point serial connections running only TCP/IP.

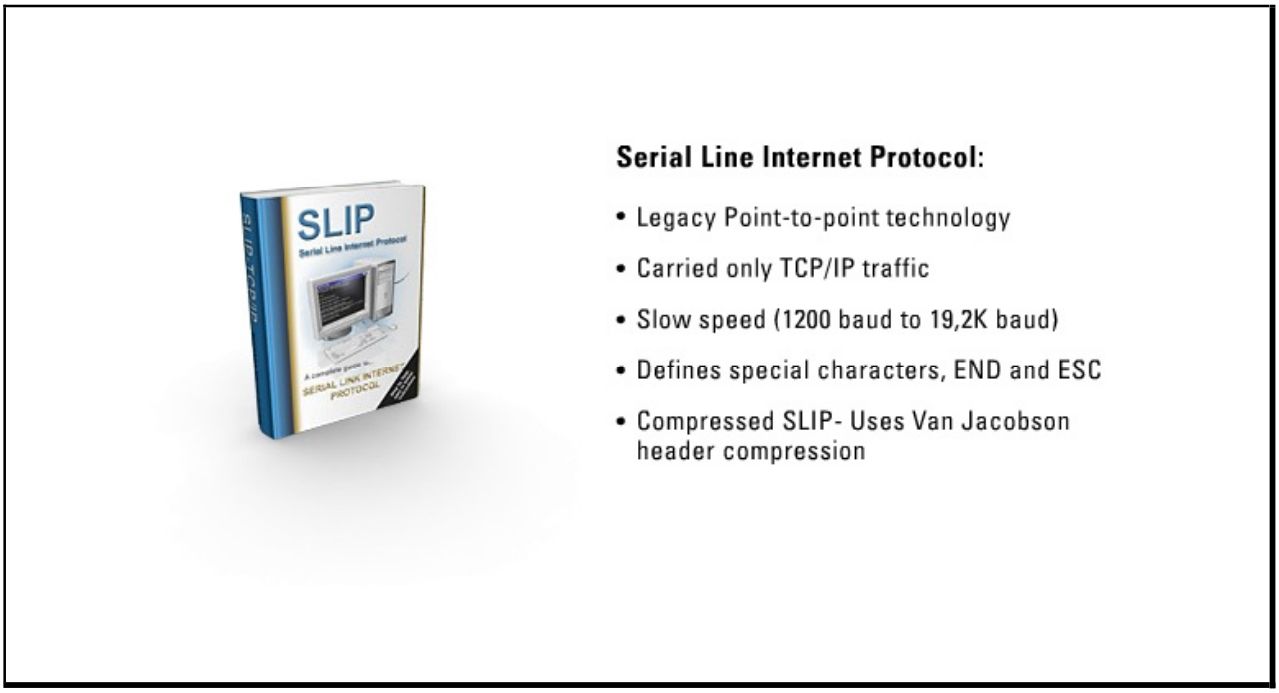

SLIP is an older technology, rarely used and configured to run on dedicated serial links and dial-up over relatively slow line speeds of 1200 baud to 19.2 Kbps. SLIP is a packet framing protocol: SLIP defines a sequence of characters that frame IP packets on a serial line. It does not provide addressing, packet type identification, error detection/correction or compression mechanisms.

The SLIP protocol defines two special characters, END and ESC. END is octal 300 (decimal 192) and ESC is octal 333 (decimal 219). To send a packet, a SLIP host simply starts sending the data in the packet. If a data byte is the same code as END character, a two-byte sequence of ESC and octal 334 (decimal 220) is sent instead. If the data byte is the same as an ESC character, a two-byte sequence of ESC and octal 335 (decimal 221) is sent instead. When the last byte in the packet has been sent, an END character is then transmitted.

Compressed Serial Line IP (CSLIP) performs the Van Jacobson header compression on outgoing IP packets. This compression improves throughput for interactive sessions noticeably.

Today, the **Point-to-Point Protocol (PPP)** largely replaces SLIP, which is more feature-rich and flexible.

PPP

The **Point-to-Point Protocol (PPP)** is the Internet Standard for transmission of IP packets over serial lines that provide far more features than its predecessor, the SLIP protocol.

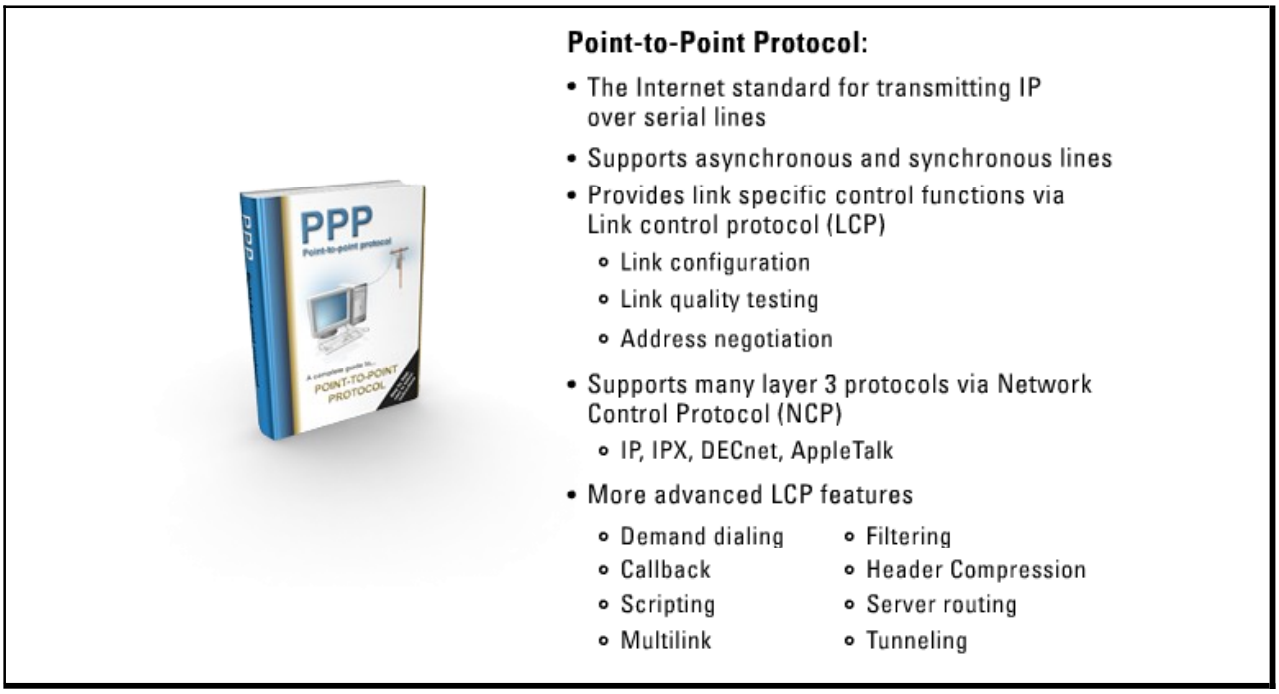

PPP supports asynchronous and synchronous lines. PPP also established a standard for the assignment and management of IP addresses, asynchronous (start/stop) and bit-oriented synchronous encapsulation, network protocol multiplexing, link configuration, link quality testing, error detection, and option negotiation for such capabilities as network layer address negotiation and data-compression negotiation. PPP supports these functions by providing an extensible Link Control Protocol (LCP) and a family of Network Control Protocols (NCPs) to negotiate optional configuration parameters and facilities. In addition to IP, PPP supports other protocols, including Novell's Internetwork Packet Exchange (IPX) and DECnet.

Some of the more advanced features in PPP include the following features:

- **Demand dialing** - Brings up a PPP interface and dials the phone when packets are queued for delivery; brings the interface down after some period of inactivity.
- **Callback** - Is sometimes called camping. Callback brings up a PPP interface whenever it goes down, to keep a line up.
- **Scripting** - Negotiates through a series of prompts or intermediate connections to bring up a PPP link, much like the sequence of events used to bring up a UUCP link.
- **Multilink** - Configures several PPP lines to the same destination and performs load sharing between them. Standardized as Multilink PPP (RFC 1990).
- **Filtering** - Selects which packets to send down a link or decides whether to bring up a demand-dial link based on IP or TCP packet type or TOS (e.g., do not dial the phone for ICMP ping packets).

- **Header compression** - Is version of SLIP that compresses the data for transmission. TCP header compression, according to RFC1144, is marginally useful on high-speed lines, and is essential for low-speed lines.
- **Server routing** - Accepts incoming PPP connections, which might well also include doing the right things with routing.
- **Tunneling** - Builds a virtual network over a PPP link across a TCP stream through an existing IP network.
- **Extra escaping** – Are byte-stuffing characters outside the negotiated asyncmap, configurable in advance but not negotiable.

PPP is capable of operating across any DTE/DCE interface, such as EIA/TIA-232-C, EIA/TIA-422, and V.35.

PAP

The **Password Authentication Protocol** (PAP) is a protocol that provides an optional authentication phase before proceeding to the Network-Layer Protocol phase.

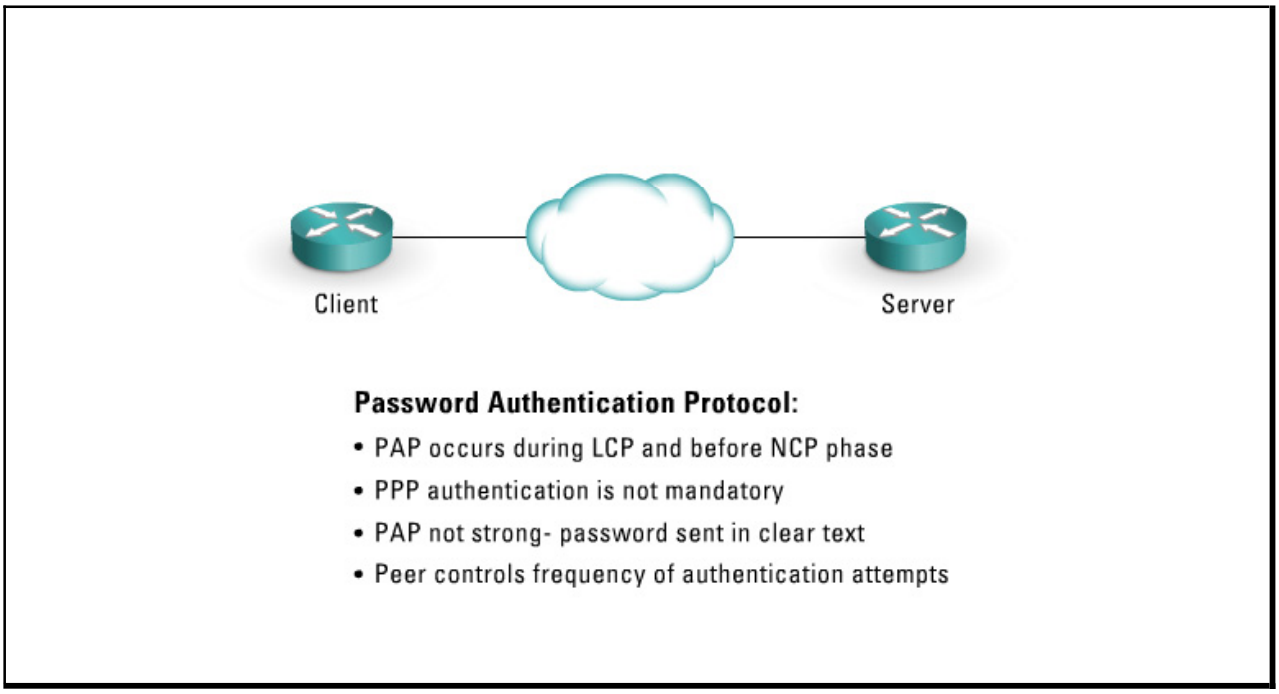

In order to establish communications over a point-to-point link, each end of the PPP link must first send LCP packets to configure the data link during Link Establishment phase. PAP authenticates the link.

By default, authentication is not mandatory. If authentication of the link is desired, an implementation must specify the Authentication-Protocol Configuration Option during Link Establishment phase.

PAP provides a simple method for the peer to establish its identity using a 2-way handshake at initial link establishment. After the Link Establishment phase is complete, an ID/Password pair is repeatedly sent by the peer to the authenticator until the authenticator acknowledges or the connection is terminated.

PAP is not a strong authentication method. Passwords are sent over the circuit in the clear, with no protection from playback or repeated trial and error attacks. The peer is in control of the frequency and timing of the attempts.

Any implementations that include a stronger authentication method such as CHAP, must offer to negotiate that method prior to PAP. To appropriately use this authentication method, a plaintext password must be available to simulate a login at a remote host. In such use, this method provides a similar level of security to the usual user login at the remote host.

CHAP

Challenge Handshake Authentication Protocol (CHAP) is a PPP authentication protocol used for remote logon, usually between a client and server or Web browser and Web server.

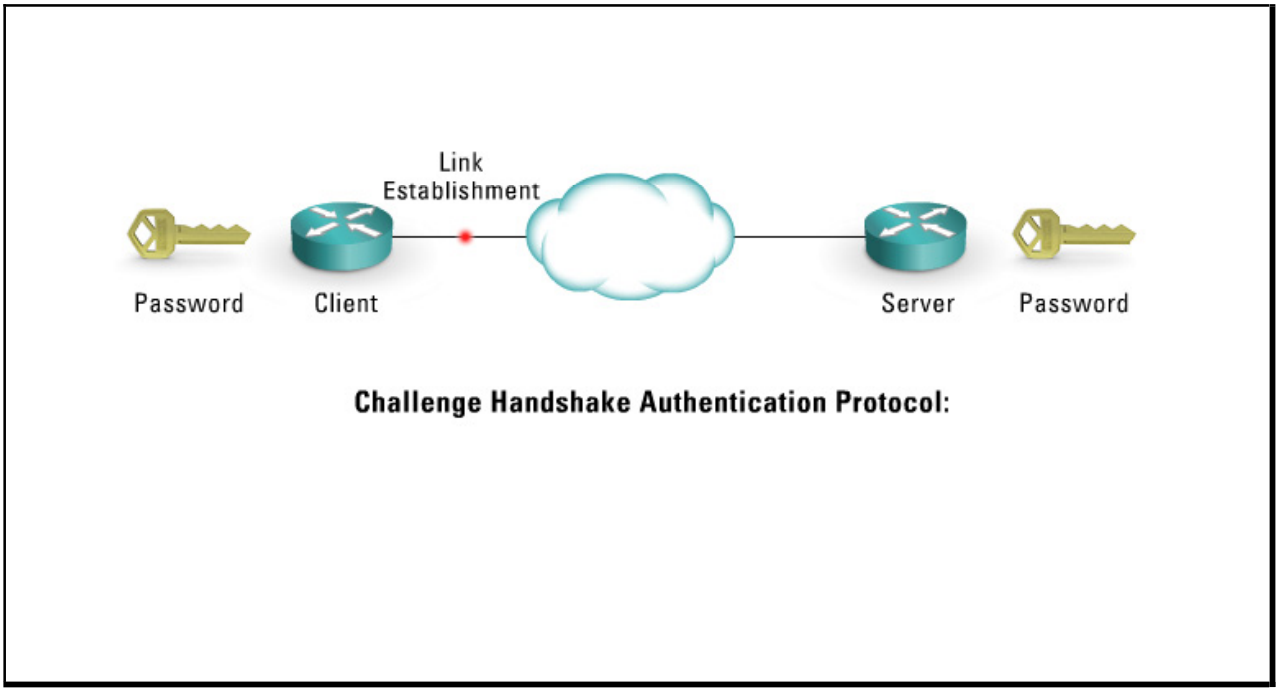

A challenge/response is a security mechanism for verifying the identity of a person or process without revealing a secret password that is shared by the two entities. It is also referred to as a three-way handshake. An important concept related to CHAP is that the client must prove to the server that it knows a shared secret without actually revealing the secret by sending the secret across the wire and revealing it to an eavesdropper. CHAP provides a mechanism for doing this.

1. After the Link Establishment phase is complete, the authenticator sends a challenge message to the peer.
2. The peer responds with a value calculated using a one-way hash function.
3. The authenticator checks the response against its own calculation of the expected hash value. If the values match, the authentication is acknowledged; otherwise, the connection should be terminated.

At random intervals, the authenticator sends a new challenge to the peer, and repeats steps 1 to 3.

These schemes are often called proof of possession protocols. The challenge requires that an entity prove possession of a shared key or one of the key pairs in a public key scheme.

This procedure is repeated throughout the session to verify that the correct client is still connected. Repeating these steps prevents someone from stealing the client's session by replaying information that was intercepted on the line.

CHAP provides protection against playback attack by the peer through the use of an incrementally changing identifier and a variable challenge value. The use of repeated challenges limits the time of exposure to any single attack. The authenticator controls of the frequency and timing of the challenges.

EAP

The **Extensible Authentication Protocol (EAP)** is an authentication framework that supports multiple authentication methods. EAP typically runs directly over data link layers such as Point-to-Point Protocol (PPP) or IEEE 802, without requiring IP.

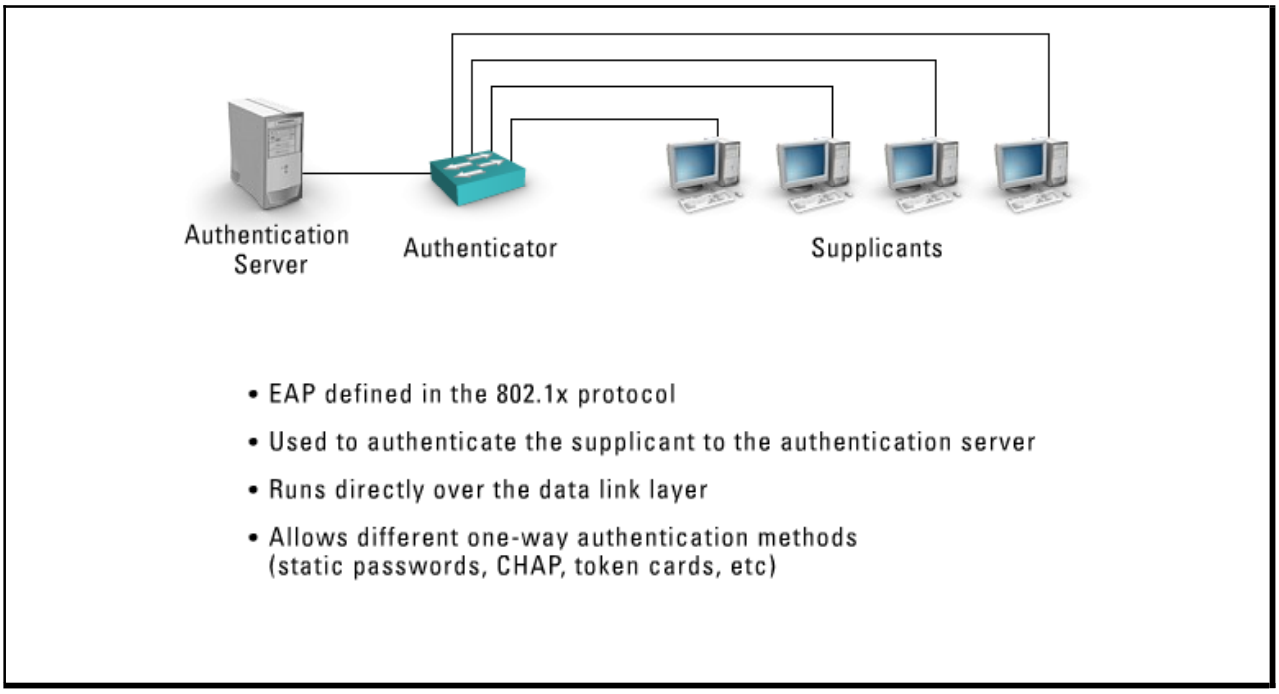

EAP is a general protocol for PPP authentication and supports multiple authentication mechanisms, such as token cards, Kerberos, one-time passwords, certificates, public key authentication and smart cards. IEEE 802.1x specifies how EAP should be encapsulated in LAN frames. EAP does not select a specific authentication mechanism at Link Control Phase, but rather postpones the selection until the Authentication Phase. This approach allows the authenticator to request more information before determining the specific authentication mechanism. This approach also permits the use of a back-end server that actually implements the various mechanisms while the PPP authenticator merely passes through the authentication exchange.

In wireless communications using EAP, a user requests connection to a WLAN through an Access Point, which then requests the identity of the user and transmits that identity to an authentication server such as RADIUS. The server asks the AP for proof of identity; the AP gets the proof from the user and then sends it back to the server to complete the authentication.

There are three entities in an EAP solution:

- **Authenticator** - Is the host requiring the authentication. The authenticator specifies the authentication protocol to be used in the Configure-Request during the Link Establishment phase.
- **Backend authentication server** - Is an entity that provides an authentication service to an authenticator. When used, this server typically executes EAP methods for the authenticator.
- **EAP server** - Is the entity that terminates the EAP authentication method with the peer. In the case where no backend authentication server is used, the EAP server is part of the authenticator.

In the case where the authenticator operates in pass-through mode, the EAP server is located on the backend authentication server.

Ethernet

The term, **Ethernet**, refers to the family of local-area network (LAN) products covered by the IEEE 802.3 standard that defines what is commonly known as the CSMA/CD protocol.

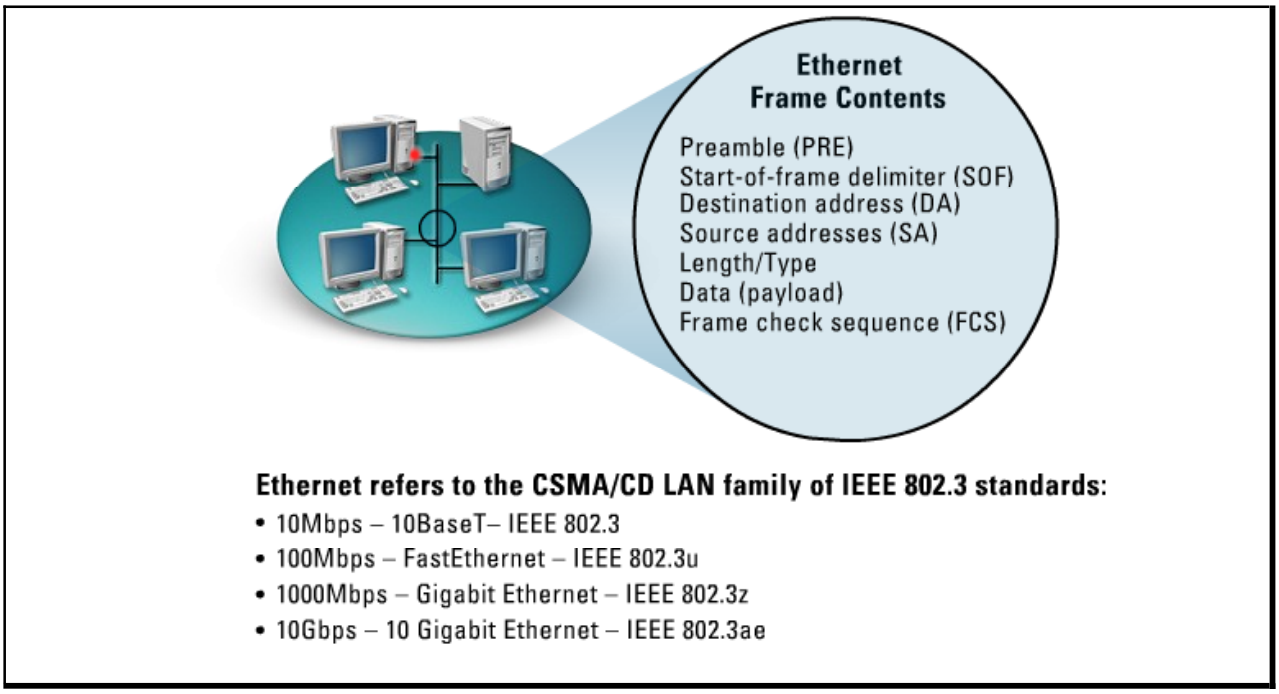

Four data rates are currently defined for operation over optical fiber and twisted-pair cables:

- 10 Mbps - 10Base-T Ethernet
- 100 Mbps - Fast Ethernet
- 1000 Mbps - Gigabit Ethernet
- 10 Gbps - 10Gigabit Ethernet

As with all IEEE 802 protocols, the ISO data link layer is divided into two IEEE 802 sublayers, the Media Access Control (MAC) sublayer and the MAC-client sublayer. The IEEE 802.3 physical layer corresponds to the ISO physical layer.

The MAC-client sublayer may be one of the following:

- Logical Link Control (LLC), if the unit is a DTE. This sublayer provides the interface between the Ethernet MAC and the upper layers in the protocol stack of the end station. The LLC sublayer is defined by IEEE 802.2 standards.
- Bridge entity, if the unit is a DCE. Bridge entities provide LAN-to-LAN interfaces between LANs that use the same protocol, such as Ethernet to Ethernet, and also between different protocols such as Ethernet to Token Ring. Bridge entities are defined by IEEE 802.1 standards.

Because specifications for LLC and bridge entities are common for all IEEE 802 LAN protocols, network compatibility becomes the primary responsibility of the particular network protocol.

The MAC layer controls the node's access to the network media and is specific to the individual protocol. All IEEE 802.3 MACs must meet the same basic set of logical requirements, regardless of whether they include one or more of the defined optional protocol extensions. The only requirement for communication

that does not require optional protocol extensions between two network nodes is that both MACs must support the same transmission rate.

The 802.3 Physical Layer is specific to the transmission data rate, the signal encoding, and the type of media interconnecting the two nodes. Gigabit Ethernet, for example, is defined to operate over either twisted-pair or optical fiber cable, but each specific type of cable or signal-encoding procedure requires a different physical layer implementation.

The MAC sublayer has two primary responsibilities:

- Data encapsulation, including frame assembly before transmission, and frame parsing/error detection during and after reception
- Media access control, including initiation of frame transmission and recovery from transmission failure

The IEEE 802.3 standard defines a basic data frame format that is required for all MAC implementations, plus several additional optional formats that are used to extend the protocol's basic capability. The basic data frame format contains the seven fields:

- **Preamble (PRE)** - Consists of seven bytes. The PRE is an alternating pattern of ones and zeros that tells receiving stations that a frame is coming, and that provides a means to synchronize the frame-reception portions of receiving physical layers with the incoming bit stream.
- **Start-of-frame delimiter (SOF)** - Consists of one byte. The SOF is an alternating pattern of ones and zeros, ending with two consecutive 1-bits indicating that the next bit is the left-most bit in the left-most byte of the destination address.
- **Destination address (DA)** - Consists of six bytes. The DA field identifies which station(s) should receive the frame. The left-most bit in the DA field indicates whether the address is an individual address, indicated by a zero (0), or a group address as indicated by a one (1). The second bit from the left indicates whether the DA is globally administered, indicated by a zero (0) or locally administered, indicated by a one (1). The remaining 46 bits are a uniquely assigned value that identifies a single station, a defined group of stations, or all stations on the network.
- **Source addresses (SA)** - Consists of six bytes. The SA field identifies the sending station. The SA is always an individual address and the left-most bit in the SA field is always zero (0).
- **Length/Type** - Consists of four bytes. This field indicates either the number of MAC-client data bytes that are contained in the data field of the frame, or the frame type ID if the frame is assembled using an optional format. If the Length/Type field value is less than or equal to 1500, the number of LLC bytes in the Data field is equal to the Length/Type field value. If the Length/Type field value is greater than 1536, the frame is an optional type frame, and the Length/Type field value identifies the particular type of frame being sent or received.
- **Data** - Is a sequence of n bytes of any value, where n is less than or equal to 1500. If the length of the Data field is less than 46, the Data field must be extended by adding a filler, a pad, sufficient to bring the Data field length to 46 bytes.
- **Frame check sequence (FCS)** - Consists of four bytes. This sequence contains a 32-bit cyclic redundancy check (CRC) value, which is created by the sending MAC and is recalculated by the receiving MAC to check for damaged frames. The FCS is generated over the DA, SA, Length/Type, and Data fields.

Whenever an end station MAC receives a transmit-frame request with the accompanying address and data information from the LLC sublayer, the MAC begins the transmission sequence by transferring the LLC information into the MAC frame buffer.

- The preamble and start-of-frame delimiter are inserted in the PRE and SOF fields.

- The destination and source addresses are inserted into the address fields.
- The LLC data bytes are counted, and the number of bytes is inserted into the Length/Type field.
- The LLC data bytes are inserted into the Data field. If the number of LLC data bytes is less than 46, a pad is added to bring the Data field length up to 46.

An FCS value is generated over the DA, SA, Length/Type, and Data fields and is appended to the end of the Data field.

After the frame is assembled, actual frame transmission will depend on whether the MAC is operating in half-duplex or full-duplex mode.

The IEEE 802.3 standard currently requires that all Ethernet MACs support half-duplex operation in which the MAC can be either transmitting or receiving a frame, but it cannot be doing both simultaneously. Full-duplex operation is an optional MAC capability that allows the MAC to transmit and receive frames simultaneously.

Token Ring

The **Token Ring** network was originally developed by IBM in the 1970s. It is still IBM's primary local-area network (LAN) technology.

	Data Rates	Stations/ segment	Topology	Media	Signaling	Access Method	Encoding
IBM Token Ring	4 - 16Mbps	260 STP 72 UTP	Star	Twisted Pair	Baseband	Token Passing	Differential Manchester
IEEE 802.5	4 - 16Mbps	250	Not Specified	Twisted Pair	Baseband	Token Passing	Differential Manchester

Token Ring:

- IBM (token ring) & IEEE 802.5
- In a token ring network, all stations connect directly to a Multistation Access Unit.
- Can only transmit when possessing token

The related IEEE 802.5 specification is almost identical to and completely compatible with IBM's Token Ring network. In fact, the IEEE 802.5 specification was modeled after IBM Token Ring, and it continues to shadow IBM's Token Ring development. The term Token Ring generally is used to refer to both IBM's Token Ring network and IEEE 802.5 networks. This chapter addresses both Token Ring and IEEE 802.5.

Token Ring and IEEE 802.5 networks are basically compatible; although, the specifications differ in minor ways. IBM's Token Ring network specifies a star, with all end stations attached to a device called a Multistation Access Unit (MSAU). In contrast, IEEE 802.5 does not specify a topology, although virtually all IEEE 802.5 implementations are based on a star. Other differences exist, including media type and routing information field size. For instance, IEEE 802.5 does not specify a media type; although, IBM Token Ring networks use twisted-pair wire.

IBM Token Ring network stations are directly connected to MSAUs, which can be wired together to form one large ring. Patch cables connect MSAUs to adjacent MSAUs while lobe cables connect MSAUs to stations. MSAUs include bypass relays for removing stations from the ring.

Token-passing networks move a small frame, called a token, around the network. Possession of the token grants the right to transmit. If a node receiving the token has no information to send, it passes the token to the next end station. Each station can hold the token for a maximum period of time.

If a station possessing the token does have information to transmit, it seizes the token, alters one (1) bit of the token and turns the token into a start-of-frame sequence. The station appends the information that it wants to transmit and sends this information to the next station on the ring. While the information frame is circling the ring, no token is on the network, and other stations wanting to transmit must wait.

Therefore, collisions cannot occur in Token Ring networks. If early token release is supported, a new token can be released when frame transmission is complete.

Unlike CSMA/CD networks like Ethernet, token-passing networks are deterministic, meaning that an administrator or attacker can calculate the maximum time that will pass before any end station will be capable of transmitting.

Token Ring networks employ several mechanisms for detecting and compensating for network faults. For example, one station in the Token Ring network is selected to be the active monitor. This station, which potentially can be any station on the network, acts as a centralized source of timing information for other ring stations and performs a variety of ring-maintenance functions. One of these functions is the removal of continuously circulating frames from the ring. When a sending device fails, its frame may continue to circle the ring. The circling frame can prevent other stations from transmitting their own frames and essentially can lock up the network. The active monitor can detect such frames, remove them from the ring, and generate a new token.

The IBM Token Ring network's star topology also contributes to overall network reliability. Because all information in a Token Ring network is seen by active MSAUs, these devices can be programmed to check for problems and selectively remove stations from the ring, if necessary.

A Token Ring algorithm called beaconing detects and tries to repair certain network faults. Whenever a station detects a serious problem with the network, such as a cable break, it sends a beacon frame that defines a failure domain. This domain includes the station reporting the failure, its nearest active upstream neighbor (NAUN), and everything in between. Beaconing initiates a process called autoreconfiguration, in which nodes within the failure domain automatically perform diagnostics in an attempt to reconfigure the network around the failed areas. Physically, the MSAU can accomplish this through electrical reconfiguration.

X.25

X.25 is an International Telecommunication Union-Telecommunication Standardization Sector (ITU-T) protocol standard for WAN communications. It defines how connections between user devices and network devices are established and maintained.

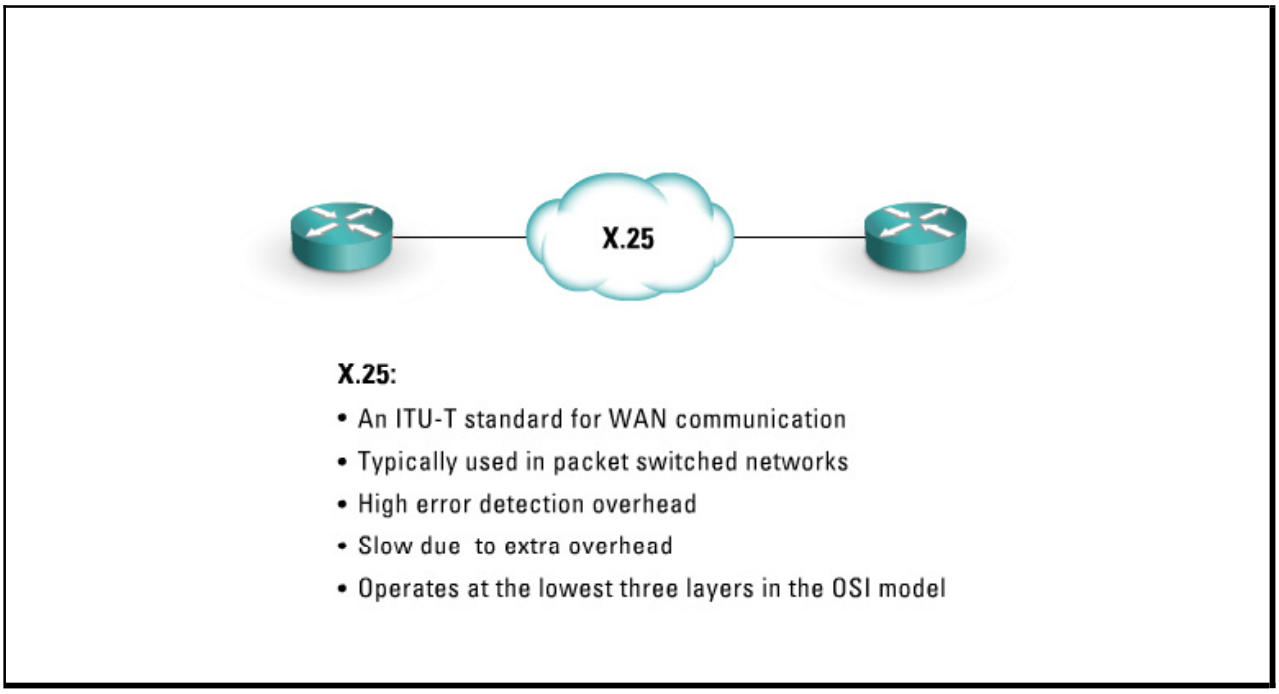

X.25 is designed to operate effectively regardless of the type of systems connected to the network. It is typically used in the packet-switched networks (PSNs) of common carriers, such as the telephone companies. Subscribers are charged based on their use of the network. The development of the X.25 standard was initiated by the common carriers in the 1970s. At that time, they needed WAN protocols capable of providing connectivity across public data networks (PDNs). X.25 is now administered as an international standard by the ITU-T.

X.25 network devices fall into three general categories: data terminal equipment (DTE), data circuit-terminating equipment (DCE), and packet-switching exchange (PSE). Data terminal equipment devices are end systems that communicate across the X.25 network. They are usually terminals, personal computers, or network hosts, and they are located on the premises of individual subscribers. DCE devices are communications devices such as modems and packet switches that provide the interface between DTE devices and a PSE. They are generally located in the carrier's facilities. PSEs are switches that compose the bulk of the carrier's network. They transfer data from one DTE device to another through the X.25 PSN.

The Packet Assembler/Disassembler (PAD) is a device commonly found in X.25 networks. PADs are used when a DTE device, such as a character-mode terminal, is too simple to implement the full X.25 functionality. The PAD is located between a DTE device and a DCE device, and it performs three primary functions: buffering by storing data until a device is ready to process it, packet assembly, and packet disassembly. The PAD buffers data sent to or from the DTE device. It also assembles outgoing

data into packets, add an X.25 header, and forwards them to the DCE device. Finally, the PAD disassembles incoming packets and removes the x.25 header before forwarding the data to the DTE.

The X.25 protocol suite maps to the lowest three layers of the OSI reference model. The following protocols are typically used in X.25 implementations: Packet-Layer Protocol (PLP), Link Access Procedure, Balanced (LAPB), and those among other physical-layer serial interfaces such as EIA/TIA-232, EIA/TIA-449, EIA-530, and G.703.

Frame Relay

Frame Relay is a high-performance WAN protocol that operates at the physical and data link layers of the OSI reference model. Originally was designed for use across Integrated Services Digital Network (ISDN) interfaces, today Frame Relay is used over a variety of other network interfaces as well.

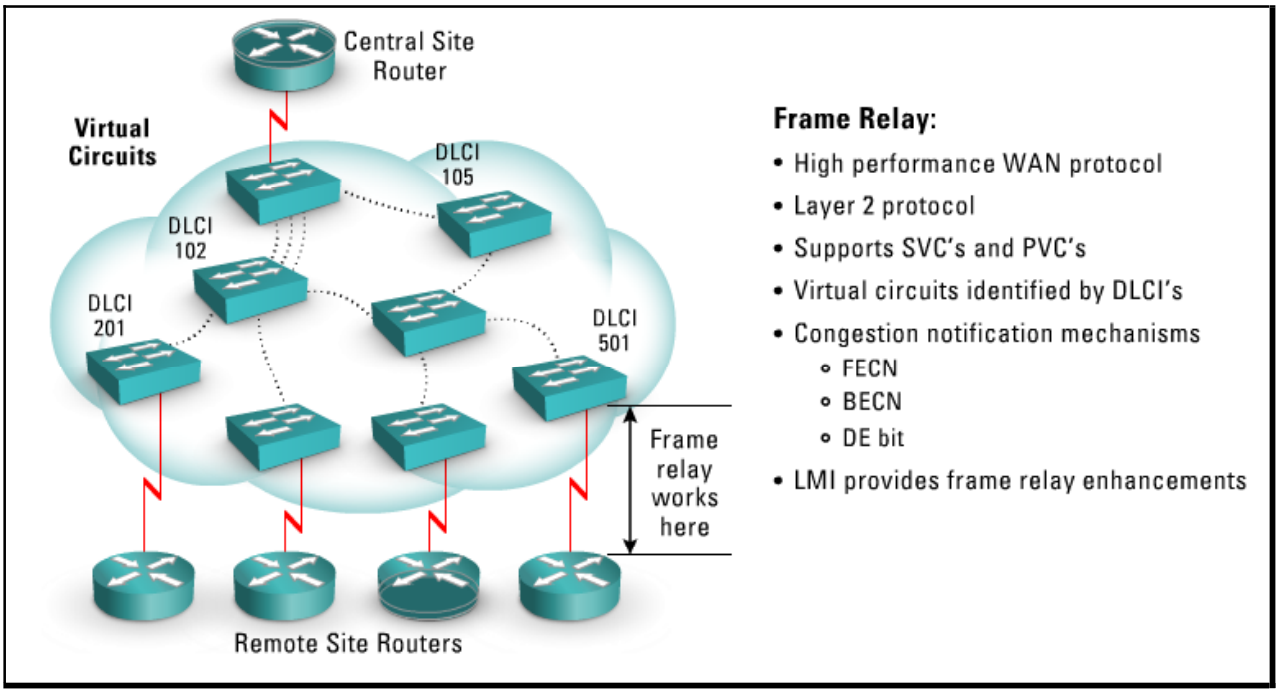

Frame Relay often is described as a streamlined version of X.25, offering fewer of the robust capabilities, such as windowing and retransmission of last data that are offered in X.25. Frame Relay typically operates over WAN facilities that offer more reliable connection services and a higher degree of reliability than the facilities available during the late 1970s and early 1980s that served as the common platforms for X.25 WANs. As mentioned earlier, Frame Relay is strictly a Layer 2 protocol suite, whereas X.25 provides services at Layer 3, the Network Layer, as well. This enables Frame Relay to offer higher performance and greater transmission efficiency than X.25, and makes Frame Relay suitable for current WAN applications, such as LAN interconnection.

Frame Relay provides connection-oriented data link layer communication, so a defined communication exists between each pair of devices and these connections are associated with a connection identifier. This service is implemented by using a Frame Relay virtual circuit, a logical connection created between two data terminal equipment (DTE) devices across a Frame Relay packet-switched network (PSN). A virtual circuit can pass through any number of intermediate DCE devices (switches) located within the Frame Relay PSN.

Frame Relay virtual circuits fall into two categories: switched virtual circuits (SVCs) and permanent virtual circuits (PVCs).

Switched Virtual Circuits (SVCs) are temporary connections used in situations requiring only sporadic data transfer between DTE devices across the Frame Relay network. A communication session across an SVC consists of the following four operational states:

- **Call setup** - The virtual circuit between two Frame Relay DTE devices is established.

- **Data transfer** - Data is transmitted between the DTE devices over the virtual circuit.
- **Idle** - The connection between DTE devices is still active, but no data are transferred. If an SVC remains in an idle state for a defined period of time, the call can be terminated.
- **Call termination** - The virtual circuit between DTE devices is terminated.

After the virtual circuit is terminated, the DTE devices must establish a new SVC if there is additional data to be exchanged.

Permanent Virtual Circuits (PVCs) are permanently established connections that are used for frequent and consistent data transfers between DTE devices across the Frame Relay network. Communication across a PVC does not require the call setup and termination states that are used with SVCs. PVCs always operate in one of the following two operational states:

- **Data transfer** - Data is transmitted between the DTE devices over the virtual circuit.
- **Idle** - The connection between DTE devices is active, but no data is transferred. Unlike SVCs, PVCs will not be terminated under any circumstances when in an idle state.

DTE devices can begin transferring data whenever they are ready because the circuit is permanently established.

Frame Relay virtual circuits are identified by Data-Link Connection Identifiers (DLCIs). Frame relay DLCI values are typically assigned by the Frame Relay service provider and have local significance, which means that their values are unique in the LAN, but not necessarily in the Frame Relay WAN.

Frame Relay reduces network overhead by implementing simple congestion-notification mechanisms rather than explicit, per-virtual-circuit flow control. Frame Relay typically is implemented on reliable network media, so data integrity is not sacrificed because flow control can be left to higher-layer protocols. Frame Relay implements two congestion-notification mechanisms:

- Forward-explicit congestion notification (FECN)
- Backward-explicit congestion notification (BECN)

FECN and BECN each is controlled by a single bit contained in the Frame Relay frame header. The Frame Relay frame header also contains a Discard Eligibility (DE) bit that is used to identify less important traffic that can be dropped during periods of congestion.

The FECN bit is part of the Address field in the Frame Relay frame header. When a DTE device sends Frame Relay frames into the network, it initiates the FECN mechanism. If the network is congested, DCE switches set the value of the frames' FECN bit to one (1). When the frames reach the destination DTE device, the Address field with the FECN bit set indicates that the frame experienced congestion in the path from source to destination. The DTE device can relay this information to a higher-layer protocol for processing. Depending on the implementation, flow control may be initiated, or the indication may be ignored.

The BECN bit is part of the Address field in the Frame Relay frame header. DCE devices set the value of the BECN bit to one (1) in frames traveling in the opposite direction of frames with their FECN bit set. This informs the receiving DTE device that a particular path through the network is congested. The DTE device then can relay this information to a higher-layer protocol for processing. Depending on the implementation, flow-control may be initiated, or the indication may be ignored.

The Local Management Interface (LMI) is a set of enhancements to the basic Frame Relay specification. Cisco Systems, StrataCom, Northern Telecom, and Digital Equipment Corporation developed the LMI in 1990. It offers a number of features (called extensions) for managing complex

internetworks. Key Frame Relay LMI extensions include global addressing, virtual circuit status messages, and multicasting.

The LMI global addressing extension gives Frame Relay data-link connection identifier (DLCI) values global rather than local significance. DLCI values become DTE addresses that are unique in the Frame Relay WAN. The global addressing extension adds functionality and manageability to Frame Relay internetworks. Individual network interfaces and the end nodes attached to them, for example, can be identified by using standard address-resolution and discovery techniques. In addition, the entire Frame Relay network appears to be a typical LAN to routers on its periphery.

SDLC

IBM developed the **Synchronous Data Link Control (SDLC)** protocol in the mid-1970s for use in Systems Network Architecture (SNA) environments. SDLC was the first link layer protocol based on synchronous, bit-oriented operation.

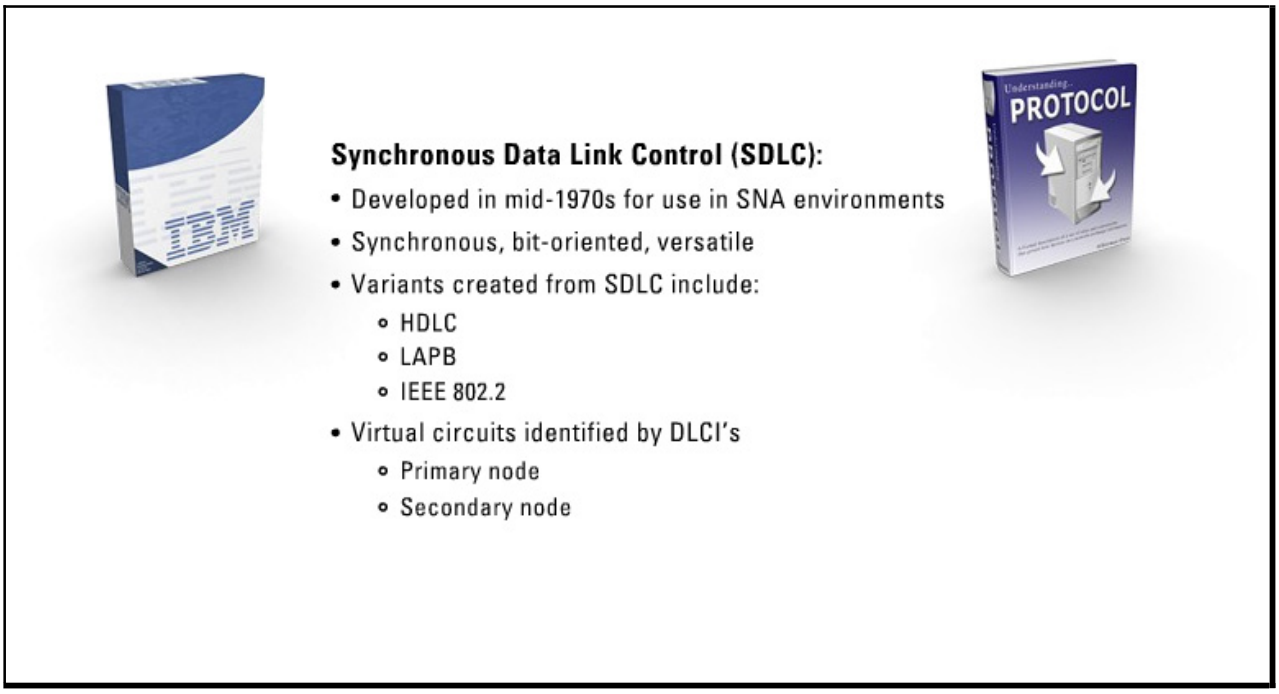

After developing SDLC, IBM submitted it to various standards committees. The International Organization for Standardization (ISO) modified SDLC to create the High-Level Data Link Control (HDLC) protocol. The International Telecommunication Union-Telecommunication Standardization Sector (ITU-T; formerly CCITT) subsequently modified HDLC to create Link Access Procedure (LAP) and then Link Access Procedure, Balanced (LAPB). The Institute of Electrical and Electronic Engineers (IEEE) modified HDLC to create IEEE 802.2. Each of these protocols has become important in its domain, but SDLC remains the primary SNA link layer protocol for WAN links.

SDLC supports a variety of link types and topologies. It can be used with point-to-point and multipoint links, bounded and unbounded media, half-duplex and full-duplex transmission facilities, and circuit-switched and packet-switched networks.

SDLC identifies two types of network nodes: primary and secondary. Primary nodes control the operation of other stations, called secondaries. The primary polls the secondaries in a predetermined order, and secondaries can then transmit if they have outgoing data. The primary also sets up and tears down links and manages the link while it is operational. Secondary nodes are controlled by a primary, which means that secondaries can send information to the primary only if the primary grants permission.

SDLC primaries and secondaries can be connected in four basic configurations:

- **Point-to-point** - Involves only two nodes, one primary and one secondary.
- **Multipoint** - Involves one primary and multiple secondaries.

- **Loop** - Involves a loop topology, with the primary connected to the first and last secondaries. Intermediate secondaries pass messages through one another as they respond to the requests of the primary.
- **Hub go-ahead** - Involves an inbound and an outbound channel. The primary uses the outbound channel to communicate with the secondaries. The secondaries use the inbound channel to communicate with the primary. The inbound channel is daisy-chained back to the primary through each secondary.

HDLC

ISO's **High-level Data Link Control (HDLC)** is a standard that corresponds to Layer 2, the Data Link Layer, of the OSI 7-layer architecture. It is responsible for the error-free movement of data between network nodes.

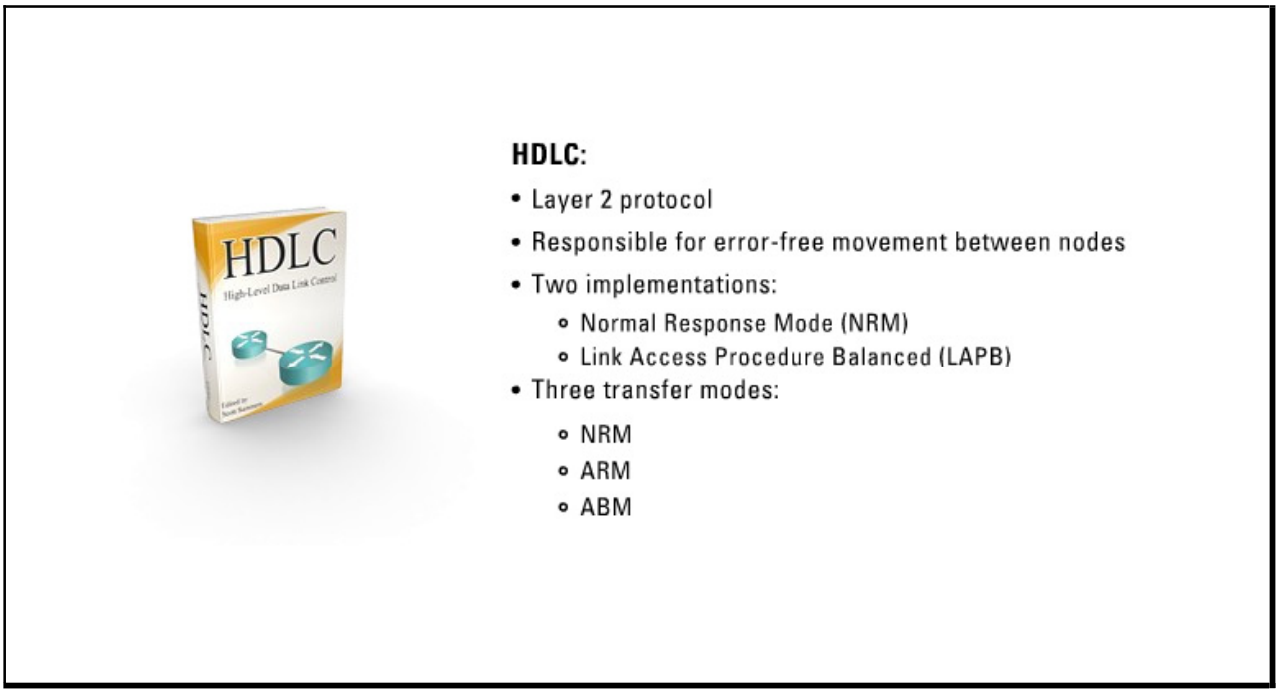

The job of the HDLC layer is to ensure that data passed up to the next layer has been received exactly as transmitted, error free, without loss, and in the correct order. Another important job is flow control, which ensures stations.

HDLC LAPB is a very efficient protocol. A minimum of overhead is required to ensure flow control, error detection, and recovery. If data is flowing in both directions (full duplex), the data frames themselves carry all the information required to ensure data integrity.

The data is transmitted only as fast as the receiver can receive it.

There are two distinct HDLC implementations: HDLC, NRM, and HDLC Link Access Procedure Balanced (LAPB). Usually, when referring to HDLC, people mean LAPB or some variation. LAPB is a bit-oriented synchronous protocol that provides complete data transparency in a full-duplex point-to-point operation. It supports a peer-to-peer link in that neither end of the link plays the role of the permanent master station. HDLC NRM, on the other hand, has a permanent primary station with one or more secondary

The concept of a frame window is used to send multiple frames before receiving confirmation that the first frame has been received correctly. This means that data can continue to flow in situations where there may be long turnaround time lags without stopping to wait for an acknowledgement. This kind of situation occurs, for instance, in satellite communication. The window sizes vary, but are typically seven frames for most terrestrial lines and up to 128 frames for satellite links.

There are three categories of frames:

- Information frames transport data across the link and may encapsulate the higher layers of the OSI architecture.
- Supervisory frames perform the flow control and error recovery functions.
- Unnumbered frames provide the link initialization and termination.

HDLC shares the frame format of SDLC, and HDLC fields provide the same functionality as those in SDLC. Also, as in SDLC, HDLC supports synchronous, full-duplex operation.

HDLC differs from SDLC in several minor ways, however. First, HDLC has an option for a 32-bit checksum. Also, unlike SDLC, HDLC does not support the loop or hub go-ahead configurations.

The major difference between HDLC and SDLC is that SDLC supports only one transfer mode, whereas HDLC supports three:

- **Normal response mode (NRM)** - This transfer mode is also used by SDLC. In this mode, secondaries cannot communicate with a primary until the primary has given permission.
- **Asynchronous response mode (ARM)** - This transfer mode enables secondaries to initiate communication with a primary without receiving permission.
- **Asynchronous balanced mode (ABM)** - ABM introduces the combined node, which can act as a primary or a secondary, depending on the situation. All ABM communication occurs between multiple combined nodes. In ABM environments, any combined station can initiate data transmission without permission from any other station.

LAPB

Link Access Procedure, Balanced (LAPB) is a data link layer protocol used to manage communication and packet framing between data terminal equipment (DTE) and the data circuit-terminating equipment (DCE) devices.

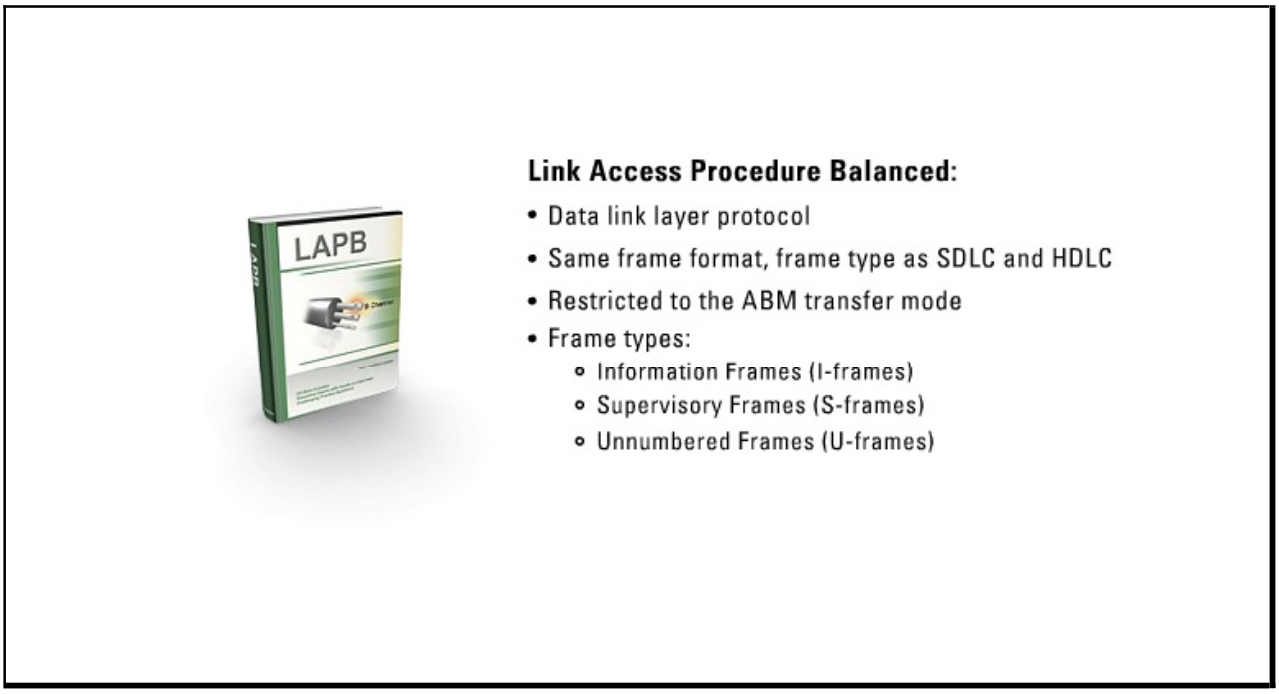

LAPB is best known for its presence in the X.25 protocol stack. LAPB shares the same frame format, frame types, and field functions as SDLC and HDLC. Unlike either of these, however, LAPB is restricted to the ABM transfer mode and is appropriate only for combined stations. Also, LAPB circuits can be established by either the data terminal equipment (DTE) or the data circuit-terminating equipment (DCE). The station initiating the call is determined to be the primary, and the responding station is the secondary.

LAPB's Frame Types:

- **I-Frames** - Information frames - Carry upper-layer information and some control information. I-frame functions include sequencing, flow control, and error detection and recovery. I-frames carry send and receive sequence numbers.
- **S-Frames** - Supervisory Frames - Carry control information. S-frame functions include requesting and suspending transmissions, reporting on status, and acknowledging the receipt of I-frames. S-frames carry only receive sequence numbers.
- **U-Frames** - Unnumbered Frames - Carry control information. U-frame functions include link setup and disconnection, as well as error reporting. U-frames carry no sequence numbers.

ISDN

ISDN, which stands for **Integrated Services Digital Network**, is a system of digital phone connections, which has been available for over a decade. This system allows voice and data to be transmitted simultaneously across the world using end-to-end digital connectivity.

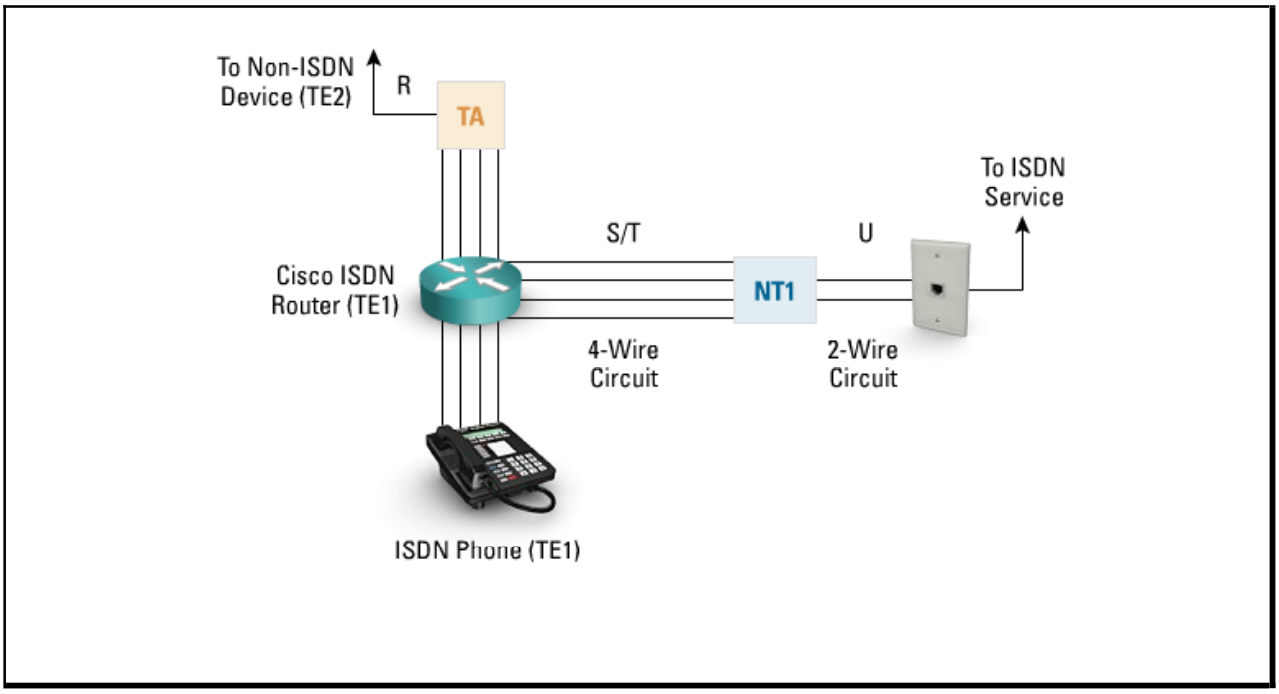

With ISDN, voice and data are carried by bearer channels (B channels) occupying a bandwidth of 64 Kbps. Some switches limit B channels to a capacity of 56 Kbps. A data channel (D channel) handles signaling at 16 Kbps or 64 Kbps, depending on the service type.

There are two basic types of ISDN service: Basic Rate Interface (BRI) and Primary Rate Interface (PRI). BRI consists of two 64 Kbps B channels and one 16 Kbps D channel for a total of 144 Kbps. This basic service is intended to meet the needs of most individual users.

PRI is intended for users with greater capacity requirements. Typically the channel structure is 23 B channels plus one 64 Kbps D channel for a total of 1536 Kbps. In Europe, PRI consists of 30 B channels plus one 64 Kbps D channel for a total of 1984 Kbps.

In the U.S., the telephone company provides its BRI customers with a U interface. The U interface is a two-wire, single pair, interface from the phone switch, the same physical interface provided for POTS lines. It supports full-duplex data transfer over a single pair of wires; therefore, only a single device can be connected to a U interface. This device is called a Network Termination 1 (NT-1). The situation is different elsewhere in the world, where the phone company is allowed to supply the NT-1, and thereby the customer is given an S/T interface.

The NT-1 is a relatively simple device that converts the two-wire U interface into the 4-wire S/T interface. The S/T interface supports multiple devices; up to seven devices can be placed on the S/T bus. While it is still a full-duplex interface, a pair of wires receives data and another pair transmits data. Today, many devices have NT-1s built into their design. This configuration has the advantage of making

the devices less expensive and easier to install, but often reduces flexibility by preventing additional devices from being connected.

Technically, ISDN devices must go through a Network Termination 2 (NT-2) device, which converts the T interface into the S interface. This configuration is possible because the S and T interfaces are electrically equivalent. Virtually all ISDN devices include an NT-2 in their design. The NT-2 communicates with terminal equipment, and handles the Layer 2 and Layer 3 ISDN protocols. Devices most commonly expect either a U interface connection with a built-in NT-1, or an S/T interface connection.

Devices that connect to the S/T or S interface include ISDN capable telephones and FAX machines, video teleconferencing equipment, bridge/routers, and terminal adapters. All devices that are designed for ISDN are designated Terminal Equipment 1 (TE1). All other communication devices that are not ISDN capable, but have a POTS telephone interface, also called the R interface, including ordinary analog telephones, FAX machines, and modems. These devices are designated Terminal Equipment 2 (TE2). A Terminal Adapters (TA) connects a TE2 to an ISDN S/T bus.

2B1Q (2 Binary 1 Quaternary) is the most common signaling method on U interfaces. This protocol is defined in detail in 1988 ANSI spec T1.601. In summary, 2B1Q provides the following:

- Two bits per baud
- 80 kilobaud (baud = 1 modulation per second)
- Transfer rate of 160Kbps

xDSL

xDSL ranges from 6.1 Mbps to 155 Mbps inbound, and from 600 Kbps to 15 Mbps outbound. The *x* is a wildcard that can be ADSL, asynchronous, or SDSL, synchronous.

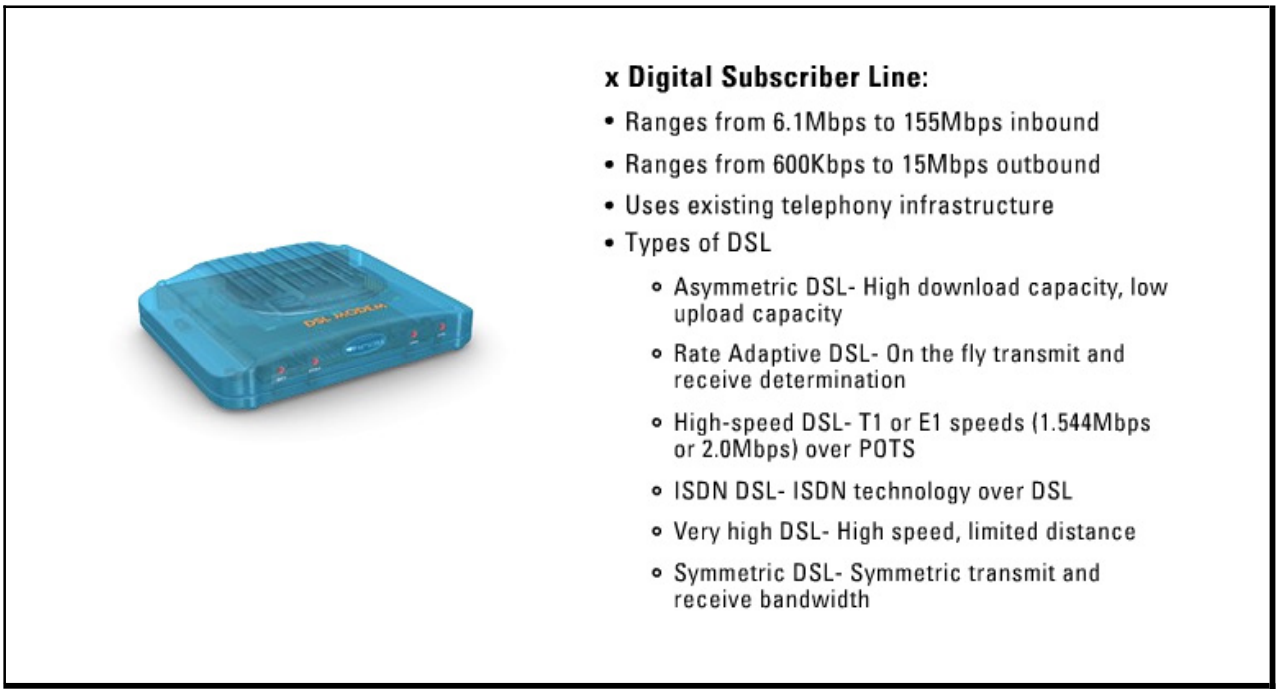

xDSL uses digital encoding to provide more bandwidth over existing twisted-pair telephone lines (POTS). Many iterations of xDSL allow the phone to be used for data communication at the same time it is being used to transmit data. This technique works because phone conversations use frequencies below 4 KHz, above which xDSL tends to operate. Several types of xDSL modems come with splitters for using voice and data concurrently.

xDSL connections use frequencies of more than 4000 KHz to achieve their great bandwidth. This bandwidth comes at the expense of attenuation. The two most popular types of line coding, CAP and DMT, use lower frequencies and, therefore, cannot support longer loops between the user and the phone company.

- **Asymmetric Digital Subscriber Line (ADSL)** is asymmetric because of its relatively high capacity to download data when compared to its lower upload capacity. ADSL allows you an 18,000-foot loop from the phone company and is capable of transmitting at speeds of up to 8 Mbps over ordinary twisted copper pairs. ADSL allows for a splitter box that lets users talk on the telephone at the same time data is being transmitted. The asymmetric speed of ADSL is appropriate for home users who typically draw more from the Internet than they send out to it. ADSL uses carrierless amplitude phase modulation (CAP) or discrete multitone (DMT).
- The speed of a **Rate Adaptive Digital Subscriber Line (R-ADSL)** is dependent on the capacity of the twisted pair it's running over. R-ADSL allows for on-the-fly determinations of proper transmit and receive speeds based upon the quality of the connection, length of the loop, and the type of wire being used in the loop. It should be used in situations in which the quality of the line connection is variable or affected by weather. R-ADSL also allows for a splitter. It transmits data using CAP.

- **High-Speed Digital Subscriber Line (HDSL)** is the result of early 1990s research into approaching T1 and E1 speeds (1.5Mbps and 2.0Mbps, respectively) over POTS. HDSL uses the same encoding methods employed by ISDN and employs two sets of phone lines. It also employs a higher number of bits per baud. The incoming and outgoing speeds of HDSL are identical.
- **ISDN Digital Subscriber Line (IDSL)** technology ports ISDN functionality to DSL. It permits data speeds of 128Kbps over ordinary twisted-pair phone lines in loops of 18,000 feet. IDSL is capable of using the same hardware as ISDN. IDSL has the advantage of being able to use any transport protocol that ISDN can, such as PPP or Frame Relay. IDSL uses the same 2B1Q line coding that ISDN does. IDSL does not support voice transmission.
- **Very High Digital Subscriber Line (VDSL)** suffers from extremely high attenuation, resulting in loop lengths of only about 3,000 feet.
- **Symmetric Digital Subscriber Line (SDSL)** is symmetrical in that the incoming and outgoing bandwidth is the same. SDSL can duplicate T1 or E1 speeds over normal twisted-pair (copper) phone cable for distances up to 11,000 feet (looped, so you must be within about 1 wire mile of your telephone company). SDSL uses carrierless amplitude phase modulation (CAP).

Cable Modems

The term cable modem refers to a modem that operates over the ordinary cable TV network cables.

Cable Modem:

- Data over television cable
- Additional equipment required for two-way communication in provider plant
- Typical speeds of 3Mbps to 50Mbps
- Very long distances (100km or more)
- Dominant standard is DOCSIS
- Downstream rates of 27Mbps to 36Mbps
- Upstream rates of 320Kbps to 10Mbps

Basically you just connect the Cable Modem to the TV outlet for your cable TV, and the cable TV operator connects a Cable Modem Termination System (CMTS) in his end (the Head-End). As cable TV is normally a simplex circuit, one way, the cable operators must install additional equipment to allow the two-way communications that are required to take place.

Actually the term Cable Modem is a bit misleading, as a Cable Modem works more like a Local Area Network (LAN) interface than as a modem.

A Cable Modem connection is something between a modem, low speed, unlimited distance, and Ethernet, a high-speed, limited distance. The speed is typically 3-50 Mbps and the distance can be 100 km or even more. The Cable Modem Termination System (CMTS) can talk to all the Cable Modems (CM's), but the Cable Modems can only talk to the CMTS. If two Cable Modems need to talk to each other, the CMTS will have to relay the messages.

First generation Cable Modems uses various proprietary protocols making it impossible for the CATV network operators to use multiple vendors Cable Modems on the same system. Around 1997, three standards emerged. DAVIC/DVB were first with a European standard, closely followed by MCSN with a US standard (DOCSIS). IEEE came last with 802.14, and clearly lost the 1st round. IEEE is now trying to leap-frog the two other standards by focusing on the next generation standards.

The dominant US standard is Data Over Cable Service Interface Specification (DOCSIS) even though it has not gone through any formal/independent standards body yet. This standard is very much driven by the wish of the large cable operators to have cable modems sold through the retail channel. Initially, the chip manufacturer Broadcom played an important role, by pushing the standard and the level of chip integration at a very fast pace. As a result, the complexity of the standard is generally agreed to be much

higher than what is strictly required, and is even growing. DOCSIS specifies downstream traffic transfer rates between 27 and 36 Mbps and upstream traffic transfer rates between 320 Kbps and 10 Mbps.

VPNs

Virtual Private Networks (VPNs) are secure private connections created using a public network. They are virtual in the sense that the public network is seen as a single hop between networks allowing the two networks to be virtually connected. They are private in the sense that data sent over the public network cannot be viewed but untrusted personnel. Encryption techniques create the privacy.

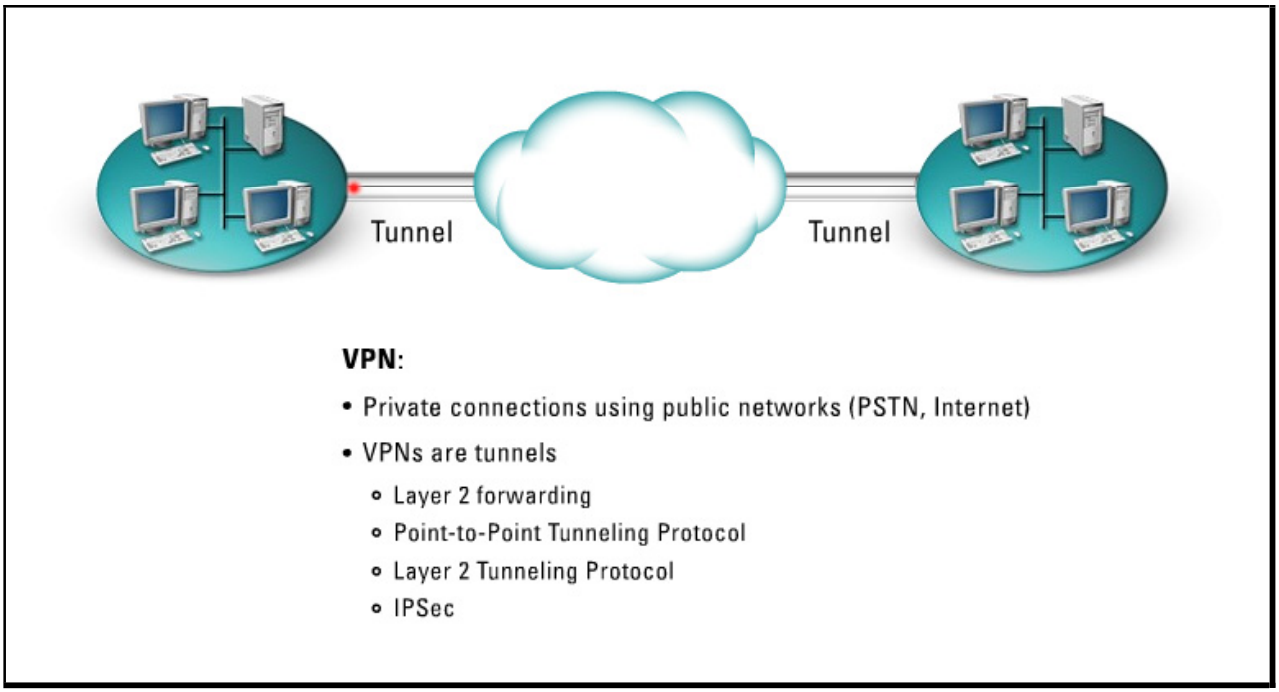

The four main VPN protocols are in use today:

- **Layer two Forwarding (L2F)** is a protocol developed by Cisco that supports the creation of secure virtual private dial-up networks (VPDNs) over the Internet.
- **Point to Point Tunneling Protocol (PPTP)** is a network protocol developed by Microsoft that enables the secure transfer of data from a remote client to a private enterprise server by creating a VPN across TCP/IP-based data networks. PPTP supports on-demand, multiprotocol, virtual private networking over public networks, such as the Internet.
- **Layer 2 Tunnel Protocol (L2TP)** is an Internet Engineering Task Force (IETF) standard that combines the best features of two existing tunneling protocols: Cisco's Layer 2 Forwarding (L2F) and Microsoft's Point-to-Point Tunneling Protocol (PPTP).
- **IPSec** - The Security Architecture for the Internet Protocol is designed to provide interoperable, high quality, cryptographically based security for IPv4 and IPv6. The set of security services offered includes access control, connectionless integrity, data origin authentication, detection and rejection of replays, a form of partial sequence integrity, confidentiality through encryption, and limited traffic flow confidentiality. The IP layer provides these services, offering protection in a standard fashion for all protocols that may be carried over IP, including IP itself.

Summary

The key points discussed in this lesson are:

- Data topologies
- Physical media characteristics
- Coaxial
- Twisted pair
- Features of fiber optics
- LAN signaling types
- LAN transmission protocols
- LAN transmission methods
- Network topologies
- Routing in Bus
- Routing in STAR
- Routing in Ring
- Routing in MESH
- LAN media access methods
- Transmission types
- LAN devices
- WAN devices
- WAN technologies
- Circuit switched and packet switched networks
- Packet switched technologies
- Remote access
- SLIP
- PPP
- PAP
- CHAP
- EAP
- Ethernet
- Token Ring
- X.25
- Frame Relay
- SDLC
- HDLC

- LAPB
- ISDN
- xDSL
- Cable Modems
- VPN

Network Layer Security Protocols

Overview

Confidentiality and integrity of data as it crosses a medium is important for many enterprise institutions. To that end, many approaches have been created to protect data in transit. These approaches usually occur at Layer 2 or Layer 3 in the OSI stack and will be discussed in this lesson.

Importance

The information security professionals need to understand the different methods of securing data en-route. They need to understand the protocols and algorithms that make up these methods in order to understand how their systems can be attacked and what countermeasures can be implemented to suppress the attacks.

Objectives

Upon completing this lesson, you will be able to:

- Define IPSec Authentication and Confidentiality
- Explain L2TP
- Explain PPTP
- Define SKIP
- Define swIPe
- List Network Attacks and Countermeasures

Outline

The lesson contains these topics:

- IPSec Authentication and Confidentiality
- L2TP
- PPTP
- SKIP

- swIPe
- Network Attacks and Countermeasures

IPSec Authentication and Confidentiality

IPSec operates at the network layer, and it enables multiple and simultaneous tunnels, unlike the point-to-point connections of the previous standards.

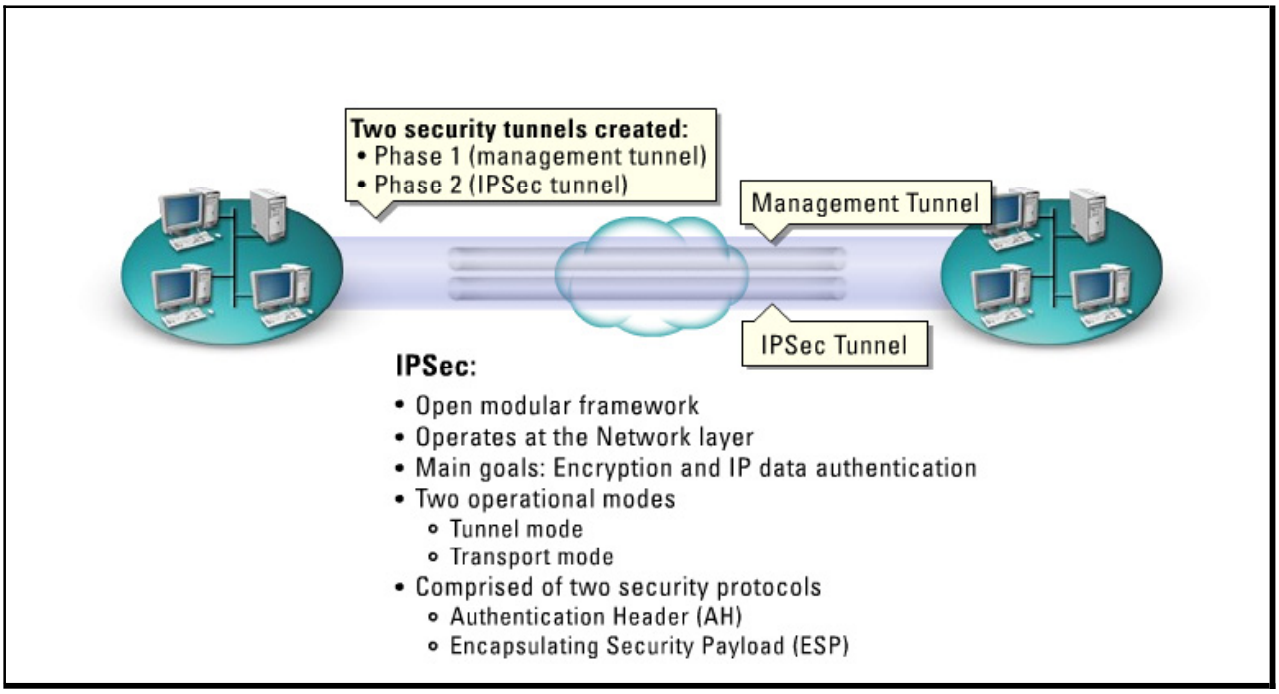

IPSec has the functionality to encrypt and authenticate IP data. It is built into the new IPv6 standard and is used as an add-on to the current IPv4. While PPTP and L2TP were aimed more at dial up VPNs, IPSec focuses more on network-to-network connectivity.

IPSec-compatible VPN devices, installed on a network's perimeter (VPN gateway) encrypt the traffic between networks or nodes by creating a secure tunnel through the unsecured network, the Internet. Because they employ IPSec encryption, they only work with IP, thus they are not multi-protocol. IPSec devices have 2 operational modes:

- **Tunnel mode** - The entire layer 3 packet is encrypted and encapsulated in an IPSec packet
- **Transport mode** - Only the layer 3 payload is encrypted leaving the IP header visible.

IPSec is an open, modular framework that provides much flexibility. It has two basic security protocols:

- **Authentication Header (AH)**: Is the authenticating protocol that provides no confidentiality services (encryption). Services of AH include the following:
  - Connectionless integrity
  - Data origin authentication
  - Anti-replay protection
- **Encapsulating Security Payload (ESP)**: Is an authenticating and encrypting protocol. Services include the following:
  - Connectionless integrity
  - Data origin authentication

- Anti-replay protection
- Confidentiality

The ESP header is inserted after the IP header and before the upper layer protocol header in the transport mode or before an encapsulated IP header, the tunnel mode.

A security association (SA) is a relationship between two or more entities that describes how the entities will use security services to communicate securely. The Internet Key Exchange (IKE) protocol negotiates the IPSec security associations (SAs). This process requires that the IPSec systems first authenticate themselves to each other and establish shared keys. IKE identifies two phases to complete the IPSec tunnel:

- In phase 1 of this process, IKE creates an authenticated, secure channel between the two IKE peers, called the IKE security association. The Diffie-Hellman key agreement is always performed in this phase.
- In phase 2, IKE negotiates the IPSec security associations and generates the required key material for IPSec. The sender offers one or more transform sets that are used to specify an allowed combination of transforms with their respective settings. The sender also indicates the data flow to which the transform set is to be applied. The sender must offer at least one transform set. The receiver then sends back a single transform set that indicates the mutually agreed-upon transforms and algorithms for this particular IPSec session. A new Diffie-Hellman agreement may be done in phase 2, or the keys may be derived from the phase 1 shared secret.

IKE Phase 1 has three methods to authenticate IPSec peers:

- **Pre-shared keys** - A key value entered into each peer manually (out of band) and used to authenticate the peer.
- **RSA signatures** - Uses a digital certificate authenticated by an RSA signature.
- **RSA encrypted nonces** - Uses RSA encryption to encrypt a nonce value, a random number generated by the peer, and other values.

IPSec involves many component technologies and encryption methods. Yet IPSec's operation can be broken down into five main steps:

1. Interesting traffic initiates the IPSec process. Traffic is deemed interesting when the IPSec security policy configured in the IPSec peers starts the IKE process.
2. IKE phase 1 - IKE authenticates IPSec peers and negotiates IKE SAs during this phase, setting up a secure channel for negotiating IPSec SAs in phase 2.
3. IKE phase 2 - IKE negotiates IPSec SA parameters and sets up matching IPSec SAs in the peers.
4. Data transfer - Data is transferred between IPSec peers based on the IPSec parameters and keys stored in the SA database.
5. IPSec tunnel termination - IPSec SAs terminate through deletion or by timing out.

L2TP

The **Layer 2 Tunnel Protocol (L2TP)** is an emerging Internet Engineering Task Force (IETF) standard that combines the best features of two existing tunneling protocols: Cisco's Layer 2 Forwarding (L2F) and Microsoft's Point-to-Point Tunneling Protocol (PPTP).

Layer 2 Tunneling Protocol:

- Combines the best of L2F and PPTP
- Supports multiple protocols
- ISP hosts the LAC
- Customer is the LNS

L2TP is an extension to the Point-to-Point Protocol (PPP), an important component for VPNs. VPNs allow users and telecommuters to connect to their corporate intranets or extranets. VPNs are cost-effective because users can connect to the Internet locally and tunnel back to connect to corporate resources. This approach not only reduces overhead costs associated with traditional remote access methods, but also improves flexibility and scalability.

Traditional dial-up networking services only support registered IP addresses limiting the types of applications that are implemented over VPNs. L2TP supports multiple protocols and unregistered and privately administered IP addresses over the Internet. This approach allows the existing access infrastructure, such as the Internet, modems, access servers, and ISDN terminal adapters (TAs), to be used.

Using L2TP tunneling, an Internet Service Provider (ISP) or other access service can create a virtual tunnel to link customer's remote sites or remote users with corporate home networks. The L2TP access concentrator (LAC) located at the ISP's point of presence (POP) exchanges PPP messages with remote users and communicates by way of L2TP requests and responses with the customer's L2TP network server (LNS) to set up tunnels. L2TP passes protocol-level packets through the virtual tunnel between end points of a point-to-point connection. The ISP's POP accepts frames from remote users stripped of any linked framing or transparency bytes, encapsulated in L2TP and forwarded over the appropriate tunnel. The customer's home gateway accepts these L2TP frames, strips the L2TP encapsulation, and processes the incoming frames for the appropriate interface.

PPTP

The **Point-to-Point Tunneling Protocol(PPTP)** works at the data link layer. Designed for individual client to server connections, it enables only a single point-to-point connection per session. This standard is very common with asynchronous connections that use Windows 9x, NT, XP, or 2000 clients.

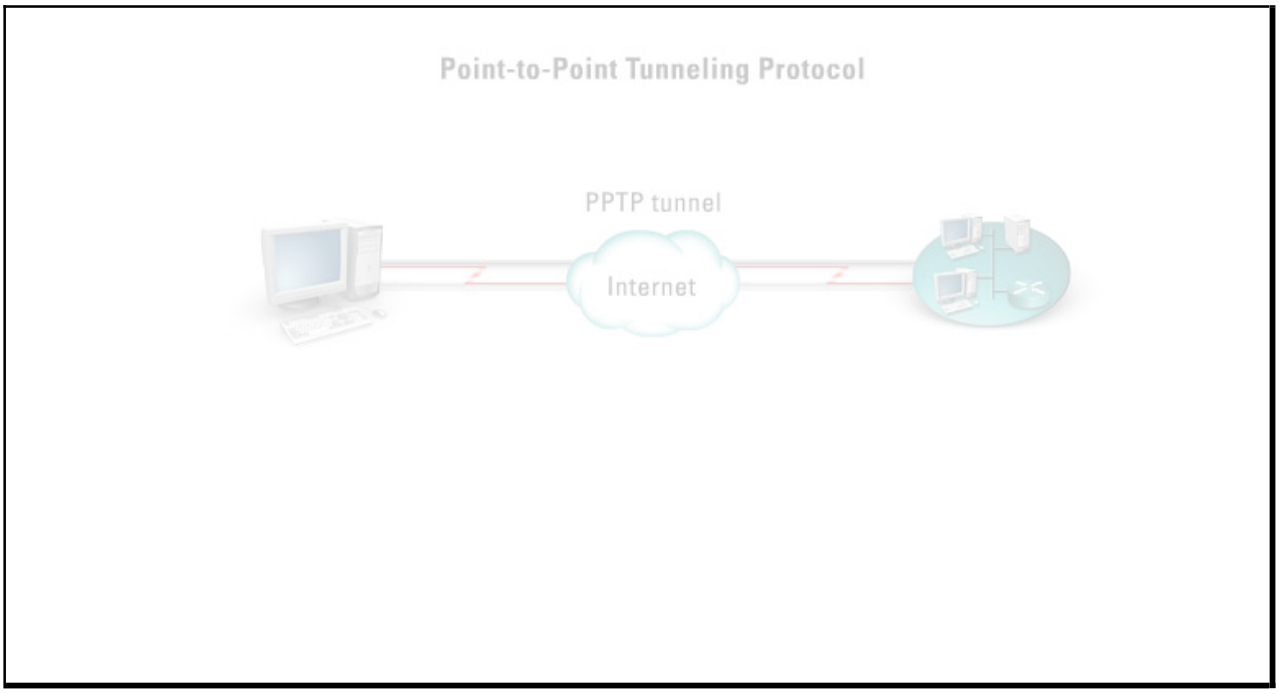

PPTP uses native point-to-point protocol authentication and encryption services. PPTP uses a Generic Routing Encapsulation (GRE) -like mechanism to provide a flow and congestion-controlled encapsulation datagram service for carrying PPP packets.

PPTP is most frequently implemented in Microsoft Windows servers and clients. It is multi-protocol, uses PAP, CHAP, or MS-CHAP user authentication, compresses data for efficient transmissions, and employs end-to-end encryption. Dial up VPNs are LAN remote access servers that have multi-protocol VPN services implemented using PPTP. They are commonly used by ISPs.

PPTP enables a low-cost, private connection to a corporate network through the public Internet. This approach is particularly useful for people who work from home or people who travel and must access their corporate networks remotely to check e-mail or perform other activities. Rather than dial a long distance number to remotely access a corporate network, with PPTP, a user could dial a local phone number using V.34 modem or ISDN for an Internet service provider point of presence. That PPTP session could provide a secure connection through the Internet back to the corporate network.

The PPTP protocol is designed to perform the following tasks:

- Query the status of Communication Servers
- Provide In-Band management
- Allocate channels and place outgoing calls
- Notify NT Server on incoming calls
- Transmit and Receive User Data with flow control in both directions

- Notify NT Server on disconnected calls.

A client-server architecture is defined in order to decouple functions that exist in current Network Access Servers (NAS) and support Virtual Private Networks (VPNs).

The PPTP Network Server (PNS) is designed to run on a general-purpose operating system while the client, referred to as a PPTP Access Concentrator (PAC), operates on a dial access platform.

SKIP

Usage of networks has expanded beyond LAN or WAN boundaries to encompass virtual networks of arbitrary size and composition. Three obstacles currently impede the growth of network services. The obstacles include the lack of security, authentication, and registration of the service data. A protocol has been developed and presented to the IETF Security Working Group - the **Simple Key Management Protocol for IP (SKIP)** protocol.

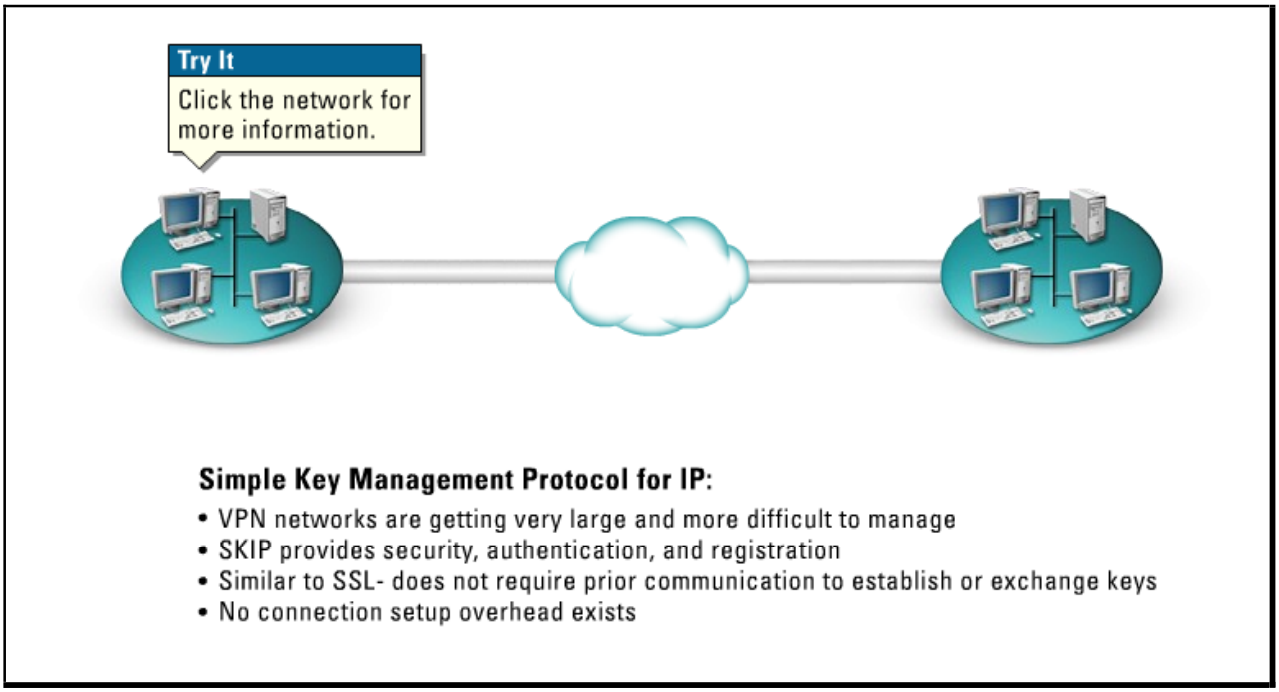

SKIP uses Diffie-Hellman 1024 bit public keys for authentication and long term key set up, session and traffic key generation, and it uses DES, RC2, and RC4 -based traffic encryption. Unique to this protocol is the lack of exchange of unencrypted keying material over the network, pipelining of traffic key generation and on-the-fly traffic key changing.

Basically, SKIP is a security technology that provides high availability in encrypted sessions as in the case of crashed gateways. SKIP is similar to SSL, except that it requires no prior communication to establish or exchange keys on a session-by-session basis. Therefore, no connection setup overhead exists, and new key values are not continually generated.

swIPe

The **swIPe** IP Security Protocol is a network-layer protocol for the IP protocol suite. swIPe provides confidentiality, integrity, and authentication of network traffic.

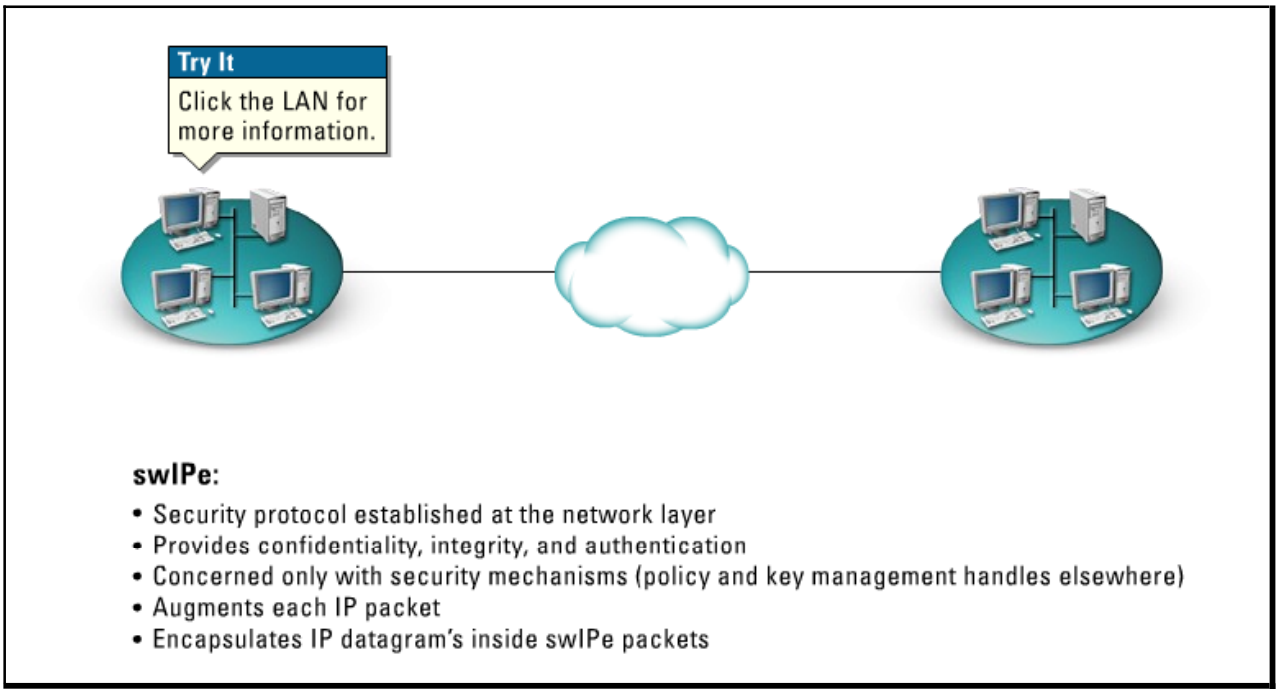

It can also provide both end-to-end and intermediate-hop security. swIPe is concerned only with security mechanisms; policy and key management are handled outside the protocol.

For the Internet continue to grow in size and features, its users need to feel confident of its safety without hiding behind impenetrable draconian barriers. The existing internetworking protocols, including IP, are deficient in three areas:

- Lack of source authentication
- Lack of data integrity
- Lack of data confidentiality

swIPe is a network-layer security protocol that provides the mechanisms to solve the three problems listed above. It works by augmenting each packet with a cryptographically strong authenticator and/or encrypting the data to be sent.

swIPe works by encapsulating each IP datagram to be secured inside a swIPe packet. A swIPe packet starts with a header that contains identifying data and authentication information; the header is followed by the original IP datagram, which in turn is followed by any padding required by the security processing. Depending on the negotiated policy, the sensitive part of the swIPe packet, the authentication information and the original IP datagram, may be encrypted.

A swIPe system consists of three conceptual entities: the protocol engine, the key management engine, and the policy engine.

Network Attacks and Countermeasures

This topic discusses the general classes of network abuse.

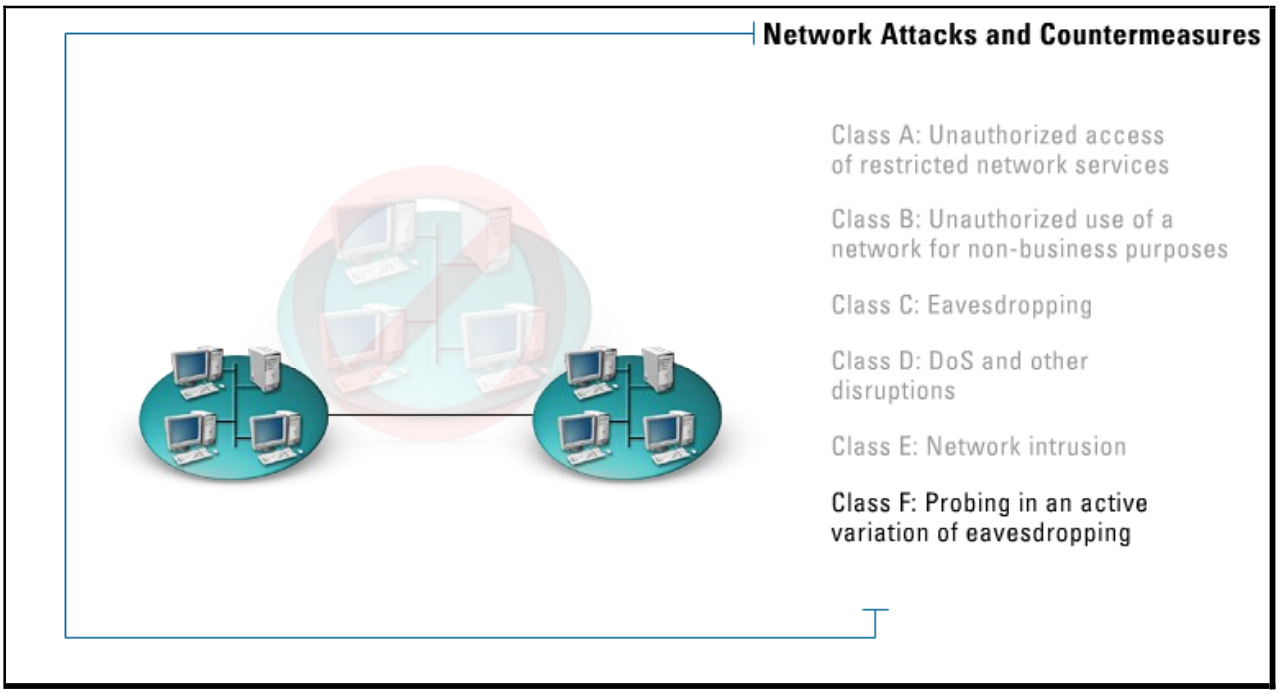

Class A: Unauthorized access of restricted network services - This type of usage is called logon abuse. It refers to legitimate users accessing networked services that would normally be restricted to them. Unlike network intrusion, this type of abuse focuses primarily on those users who may be internal to the network, legitimate users of a different system, or users who have a lower security classification. Authentication and authorization are the main countermeasures to this type of network attack.

Class B: Unauthorized use of a network for non-business purposes - This style of network abuse refers to non-business or personal use of a network by otherwise authorized users, such as internet surfing to inappropriate content sites such as porn sites. According to ISC2 code of ethics, the use of networked services for other than business purposes is abuse. Internet filtering applications such as Websense, N2H2, and SmartFilter are used to mitigate this type of network abuse.

Class C: Eavesdropping - This type of network attack consists of unauthorized interception of network traffic. Eavesdropping attacks occur through the interception of network traffic, such as sniffing. Tapping refers to the physical interception of a transmission medium such as the splicing of cable. Passive eavesdropping is covertly monitoring or listening to transmissions that are unauthorized by either sender or receiver. Active eavesdropping is tampering with a transmission to create a covert signaling channel, or actively probing the network for infrastructure information. Encryption and a switched infrastructure are used to mitigate this type of threat.

Class D: DOS and other disruptions - These attacks create service outages because of the saturation of networked resources. This saturation can be aimed at network devices, servers, or bandwidth. This attack can also be used as a diversion to enable an intentional hack to gain info from a different part of the system by diverting the company's information technology resources elsewhere. Traffic filtering, rate filtering, and firewalls help protect against this type of attack.

Class E: Network Intrusion - This type of attack refers to the use of unauthorized access to break into a network from an external source. Unlike a login abuse attack, the intruders are not considered to be known to the company. Also known as a penetration attack, this approach exploits known security vulnerabilities. Spoofing or piggybacking refers to an attacker gaining unauthorized access by using a legit account. Attacks can also include backdoors, and intrusions from dial up or external network connections. Using an Intrusion Detection System (IDS) can mitigate these types of threats.

Class F: Probing is an active variation of eavesdropping - Probing is usually used to give an attacker a road map of the network in preparation for an intrusion or a DOS attack. It can list available services. If the attacker is inside the network, he can use a sniffer to see what services are being utilized, thus can acquire knowledge of what services to exploit. Probing can be manual or automatic. Manual checks are performed using telnet for banner grabbing. Automated probing tools such as Nmap abound on the Internet. Using an Intrusion Detection System (IDS), firewalling as well as disabling unused ports and services on workstations and servers can mitigate these types of threats.

Summary

The key points discussed in this lesson are:

- IPSec Authentication and Confidentiality
- L2TP
- PPTP
- SKIP
- swIPe
- Network Attacks and Countermeasures

Transport Layer Security

Overview

Network Layer security, Layer 3 in the OSI model, requires two endpoints to agree on the way data is to be secured before the data is sent. These endpoints would then have security software manually configured to support the agreed to method. What happens when two endpoints wish to have a secure connection, but have not or cannot be configured to support security at the network layer? In this instance, security at the Transport Layer, layer 4 in the OSI model, can come into play. This lesson will discuss the two most often used protocols that secure data at the Transport Layer.

Importance

The information security professional needs to understand exactly how data is secured using Transport Layer algorithms. Equally important is to understand how this method of security can be compromised and what can be done to prevent it.

Objectives

Upon completing this lesson, you will be able to:

- Explain SSL
- Explain TLS
- Explain SSH

Outline

The lesson contains these topics:

- SSL
- TLS
- SSH

SSL

The primary goal of the **Secure Sockets Layer(SSL)** Protocol is to provide privacy and reliability between two communicating applications.

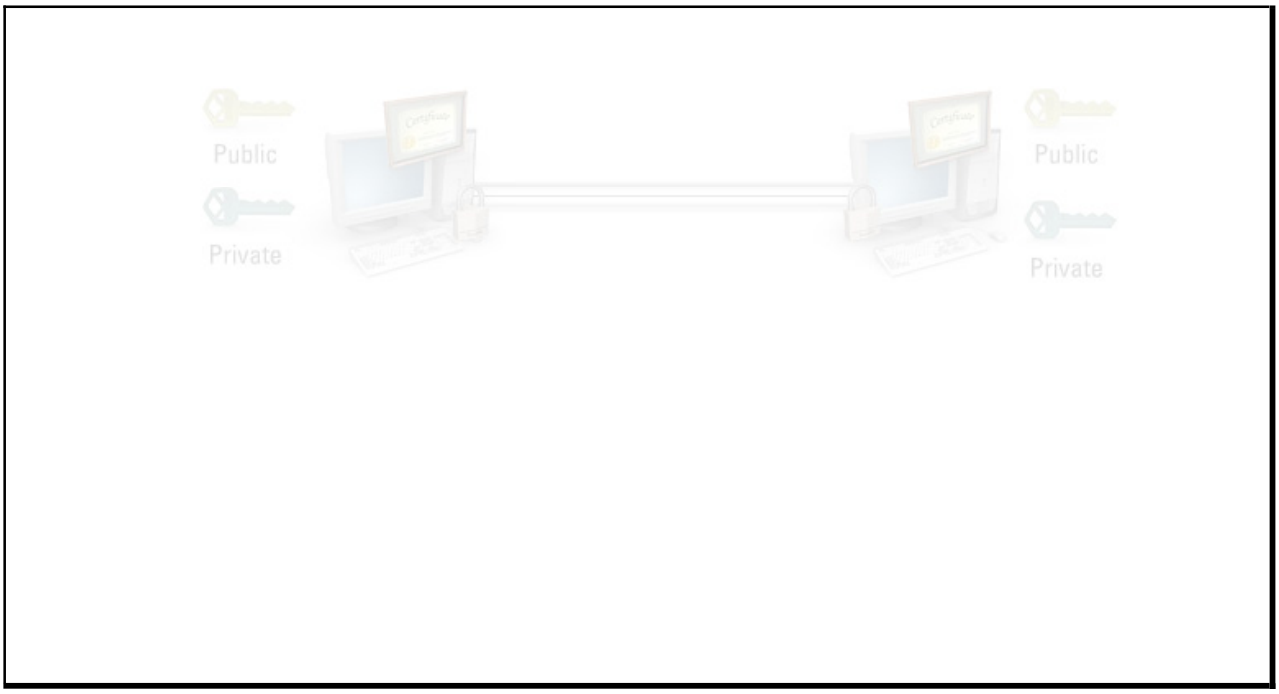

The protocol is composed of two layers. Working at the lowest level, layered on a reliable transport protocol such as TCP is the SSL Record Protocol. The SSL Record Protocol encapsulates various higher-level protocols. One encapsulated protocol, the SSL Handshake Protocol, allows the server and client to authenticate each other and to negotiate an encryption algorithm and cryptographic keys before the application protocol transmits or receives its first byte of data. One advantage of SSL is that it is application protocol independent. A higher-level protocol can layer on top of the SSL Protocol transparently. SSL is considered a session layer protocol, and each session is stateful.

The SSL protocol provides connection security that has three basic properties:

- The connection is private. Encryption is used after an initial handshake to define a secret key. Symmetric cryptography is used for data encryption in the form of DES or RC4.
- The peer's identity can be authenticated using asymmetric or public key cryptography in the form of RSA or DSS.
- The connection is reliable. Message transport includes a message integrity check using a keyed MAC. Use secure hash functions such as SHA and MD5 for MAC computations.

The goals of SSL are cryptographic security between parties, interoperability between applications and extensibility.

In digital signing, one-way hash functions are used as input for a signing algorithm. In RSA signing, a 36-byte structure of two hashes, one SHA and one MD5, is signed, encrypted with the private key.

In block cipher encryption, every block of plaintext encrypts to a block of ciphertext. Because it is unlikely that the plaintext will break neatly into the necessary block size of 64 bits, the encryption program pads out the end of short blocks with some regular pattern, usually all zeroes.

In public key encryption, one-way functions with secret trapdoors encrypt the outgoing data. Data encrypted with the public key of a given key pair can only be decrypted with the private key, and vice-versa.

SSL uses a 40-bit key size for the RC4 stream encryption algorithm, which is considered an adequate degree of encryption for commercial exchange.

TLS

SSL was developed by Netscape Communications Corporation in 1994 to secure transactions over the World Wide Web. Soon after, the Internet Engineering Task Force (IETF) began work to develop a standard protocol that provided the same functionality. They used SSL 3.0 as the basis for that work that became the **Transport Layer Security (TLS)** protocol.

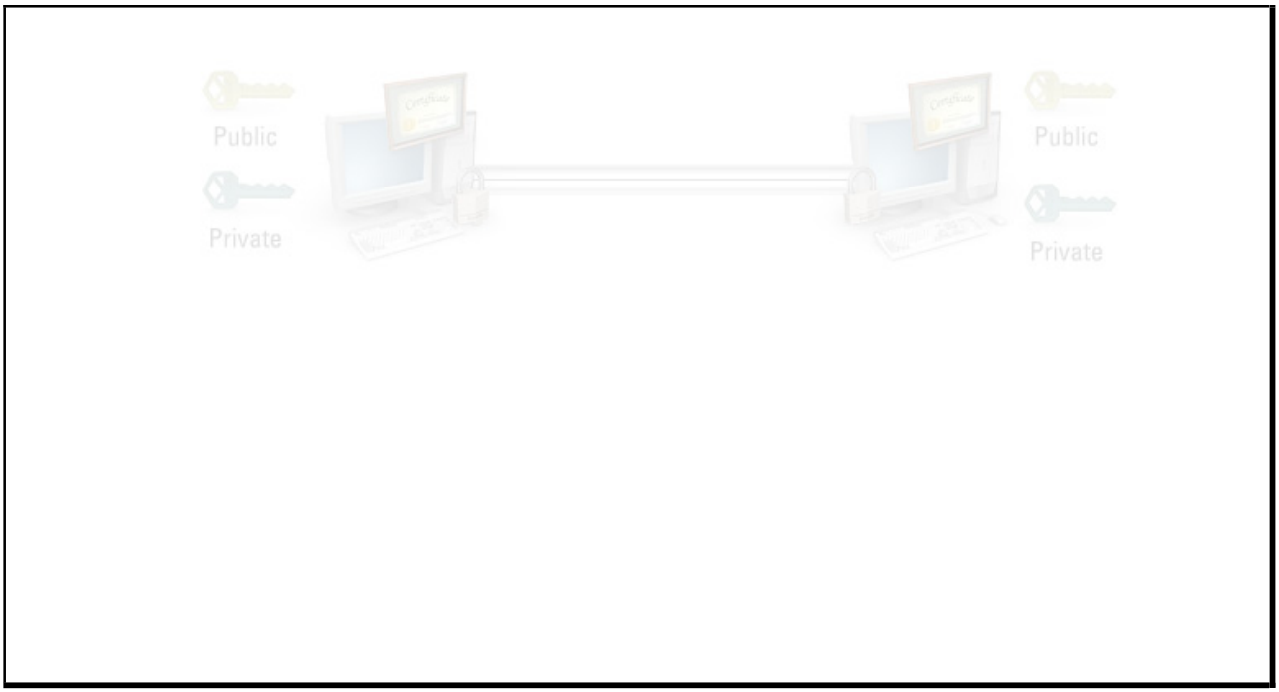

TLS and SSL are the most widely recognized as the protocols that provide secure HTTP (HTTPS) for Internet transactions between Web browsers and Web servers. Other application level protocols, such as File Transfer Protocol (FTP), Lightweight Directory Access Protocol (LDAP), and Simple Mail Transfer Protocol (SMTP) can use TLS/SSL. It enables server authentication, client authentication, data encryption, and data integrity over networks such as the World Wide Web.

The primary goal of the TLS Protocol is to provide privacy and data integrity between two communicating applications. The protocol includes two layers: the TLS Record Protocol and the TLS Handshake Protocol. At the lowest level, layered on top of some reliable transport protocol such as TCP is the TLS Record Protocol. The TLS Record Protocol provides connection security that has two basic properties:

1. The connection is private. Symmetric cryptography is used for data encryption with DES, RC4. The keys for this symmetric encryption are generated uniquely for each connection and are based on a secret negotiated by another protocol such as the TLS Handshake Protocol. The Record Protocol can also be used without encryption.
2. The connection is reliable. Message transport includes a message integrity check using a keyed MAC. MAC computations use secure hash functions such as SHA, MD5. The Record Protocol can operate without a MAC, but is generally only used in this mode while another protocol is using the Record Protocol as a transport for negotiating security parameters.

The TLS Record Protocol is used for encapsulation of various higher-level protocols. One such encapsulated protocol, the TLS Handshake Protocol, allows the server and client to authenticate each other and to negotiate an encryption algorithm and cryptographic keys before the application protocol transmits or receives its first byte of data. The TLS Handshake Protocol provides connection security that has three basic properties:

- The peer's identity can be authenticated using asymmetric, or public key, cryptography such as RSA, DSS. This authentication can be made optional, but is generally required for at least one of the peers.
- The negotiation of a shared secret is secure: the negotiated secret is unavailable to eavesdroppers, and for any authenticated connection, the secret cannot be obtained even by an attacker who can place himself in the middle of the connection.
- The negotiation is reliable: no attacker can modify the negotiation communication without being detected by the parties to the communication.

One advantage of TLS is that it is application protocol independent. Higher-level protocols can layer on top of the TLS Protocol transparently. The TLS standard, however, does not specify how protocols add security with TLS; the decisions on how to initiate TLS handshaking and how to interpret the authentication certificates exchanged are left up to the judgment of the designers and implementers of protocols who run on top of TLS.

The differences between this protocol and SSL 3.0 are not dramatic, but they are significant enough that TLS 1.1, TLS 1.0, and SSL 3.0 do not interoperate.

SSH

Secure Shell (SSH) is a program to log into another computer over a network, to execute commands in a remote machine, and to move files from one machine to another.

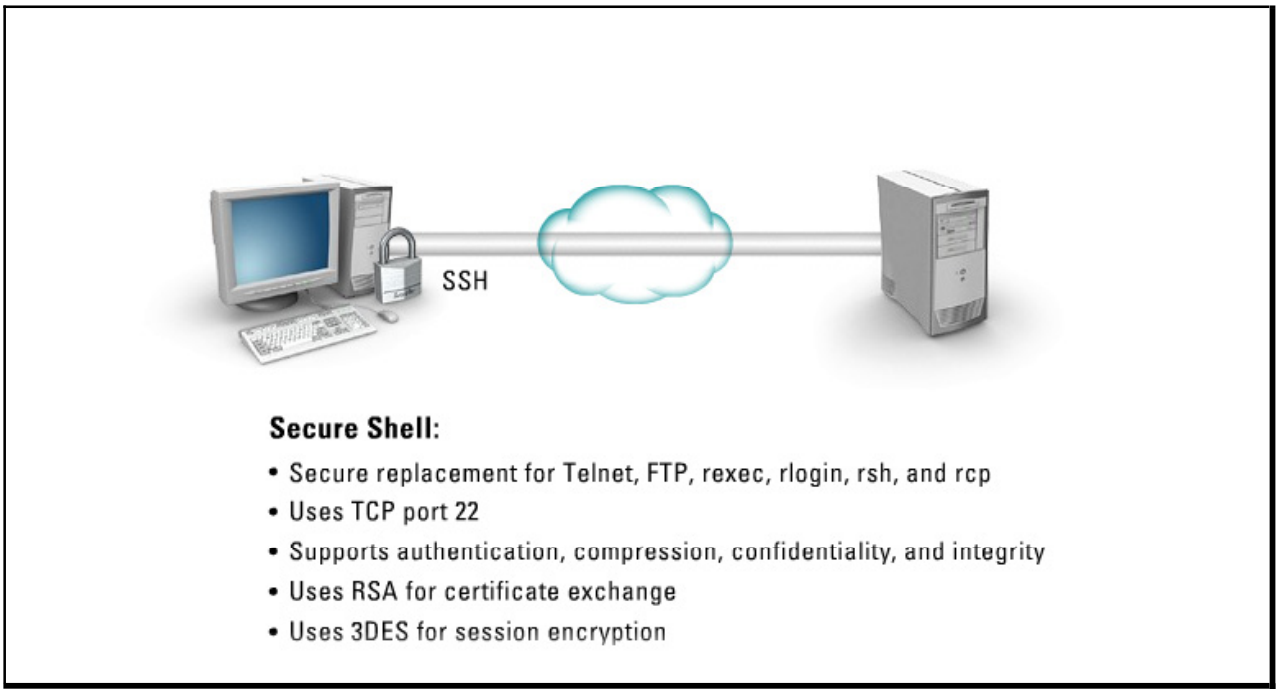

It provides strong authentication and secure communications over unsecured channels. It is intended as a replacement for telnet, ftp, rexec, rlogin, rsh, and rcp. SSH has been designated the use of TCP port 22.

Unlike telnet, SSH has a strong method of performing client authentication. It supports authentication, compression, confidentiality, and integrity. SSH uses RSA for certificate exchange and triple DES for session encryption.

SSH protects against:

- IP spoofing, where a remote host sends out packets that pretend to come from another, trusted host
- SSH even protects against a spoofer on the local network, who can pretend he is your router to the outside
- IP source routing during which a host can pretend that an IP packet comes from another, trusted host
- DNS spoofing during which an attacker forges name server records
- Interception of cleartext passwords and other data by intermediate hosts
- Manipulation of data by people in control of intermediate hosts as in man-in-the-middle attacks

To use SSH, the user has to authenticate herself to the remote host, and the remote host has to authenticate itself to the user. The user authenticates herself by encrypting a message with her private key, something only the user can do. The remote host decrypts the message with a public key. The remote host then knows that the user and only the user sent the message. The remote host daemon then encrypts a message with its private key, and the user client decrypts the message with the host's public key. The

client then knows that the host is really the host. User set up of SSH consists of sending the user's public key to the host and getting the host's public key.

Summary

The key points discussed in this lesson are:

- SSL
- TLS
- SSH

Application Layer Security Protocols

Overview

Corporations can impose security restrictions on data traveling to or from a particular system using a very specific protocol. For example, a corporation might be displaying a client's social security number or credit card number for the benefit of the user in their web browser. Obviously, this data if sent in clear text is susceptible to wire tapping attacks. For this reason, the corporation can mandate that all Web traffic to this server be secured using a particular security algorithm, which is usually based on SSL/TLS. This lesson will discuss the various protocols that can secure data at the Application Layer, Layer 7 of the OSI model.

Importance

The information security professional needs to understand the different ways that data can be secured at the Application Layer. This information will allow the security professional a better chance of identifying hostile activity against the application, server, and session.

Objectives

Upon completing this lesson, you will be able to:

- Explain Secure Electronic Transactions (SET)
- Define Privacy Enhanced Mail (PEM)
- Explain the difference between S-HTTP and HTTPS
- Explain the features of MIME
- Explain the features of S/MIME
- Explain the premise of Pretty Good Privacy (PGP)
- Explain what Cookies do
- Explain the use of RAID
- List the RAID levels

Outline

The lesson contains these topics:

- Secure Electronic Transactions (SET)
- Privacy Enhanced Mail (PEM)
- S-HTTP vs. HTTPS
- MIME
- S/MIME
- Pretty Good Privacy (PGP)
- Cookies
- RAID
- RAID levels

Secure Electronic Transactions (SET)

The **Secure Electronics Transaction (SET)** originated by Visa and MasterCard supports the authentication of both the sender and the receiver, and it ensures content privacy of credit card numbers and other personal information using digital certificates and signatures over the public Internet. SET uses TCP port 257.

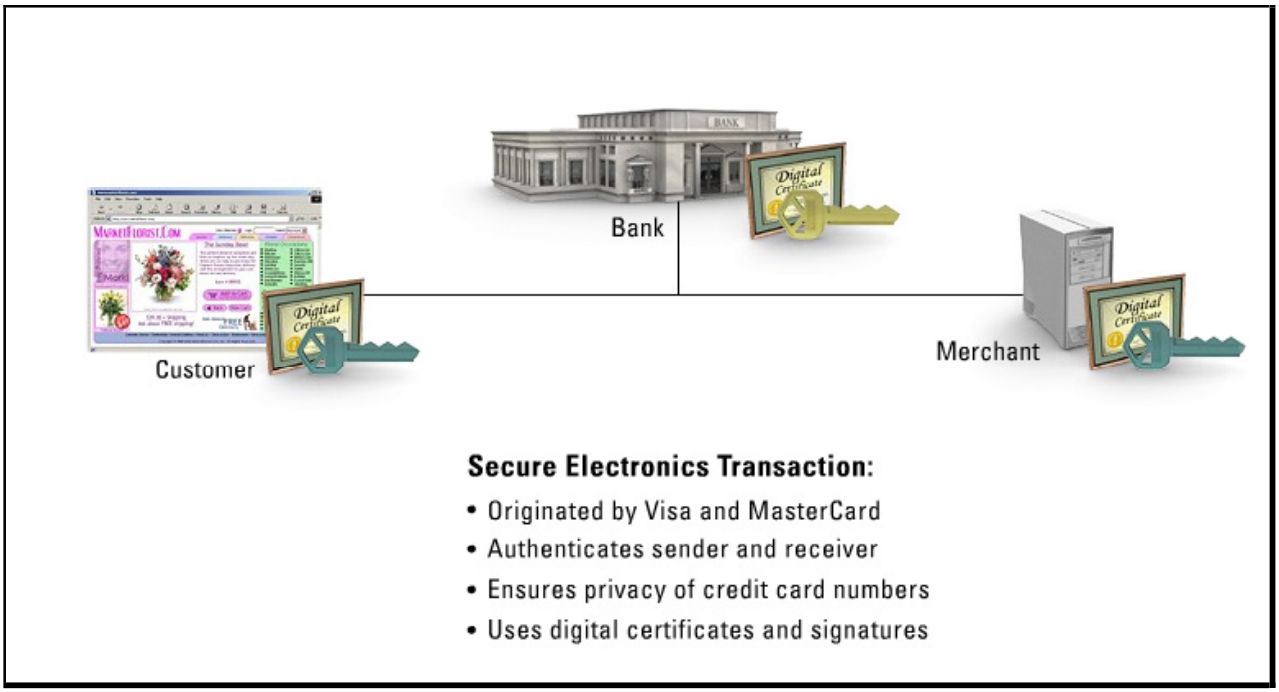

SET is comprised of three main parts: the electronic wallet, a digital certificate, the software running on the merchant's server at its web site, and the payment server that is located at the merchant's bank.

For SET to work the customer must have a SET enabled browser like Netscape or Internet Explorer, and the transaction provider either the bank or store must have a SET enabled server.

1. The customer opens a MasterCard or Visa bank account. In this case, any issuer of a credit card is some kind of bank.
2. The customer receives a digital certificate. This electronic file functions as a credit card for online purchases or other transactions. It includes a public key with an expiration date. It has been through a digital switch to the bank to ensure its validity.
3. Third-party merchants also receive certificates from the bank. These certificates include the merchant's public key and the bank's public key.
4. The customer places an order over a Web page, by phone, or some other means.
5. The customer's browser receives and confirms from the merchant's certificate that the merchant is valid.
6. The browser sends the order information. This message is encrypted with the merchant's public key, the payment information, which is encrypted with the bank's public key, which can not be read by the merchant, and information that ensures the payment can only be used with this particular order.

7. The merchant verifies the customer by checking the digital signature on the customer's certificate. This verification may be done by referring the certificate to the bank or to a third-party verifier.
8. The merchant sends the order message along to the bank. This message includes the bank's public key, the customer's payment information, which the merchant cannot decode, and the merchant's certificate.
9. The bank verifies the merchant and the message. The bank uses the digital signature on the certificate with the message and verifies the payment part of the message.

The bank digitally signs and sends authorization to the merchant, who can then fill the order.

Because of SET's overhead, additional steps and software required, SET has been largely overtaken by SSL/TLS.

Privacy Enhanced Mail (PEM)

To provide privacy for electronic mail, mechanisms are needed to assure both sender and receiver that messages are confidential, that messages are from an authentic source, that messages have not been altered or corrupted, and that the sender cannot repudiate or disown the message. The Privacy Enhanced Mail (PEM) protocol was one of the first standards for securing the text of e-mail messages.

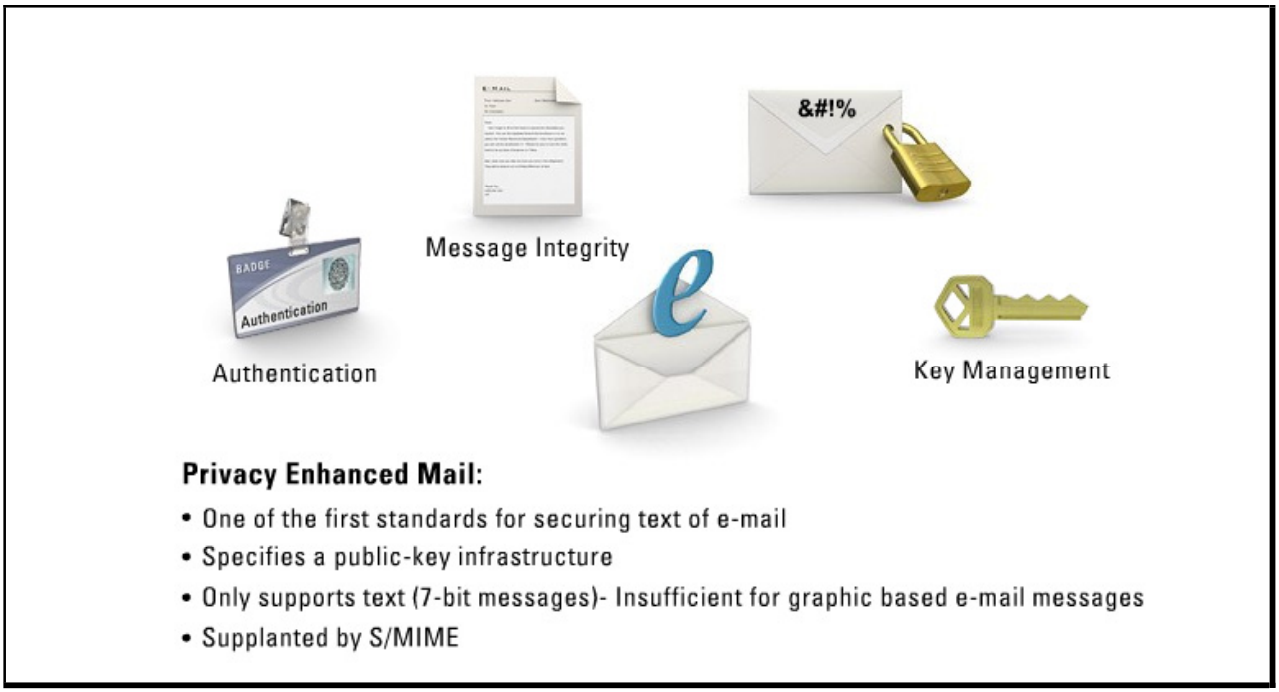

PEM was defined by the IETF as a way to encrypt 7-bit text messages. It also defined a hierarchical structure for distributing and verifying digital signatures. PEM specifies a public-key infrastructure for key exchange over large networks like the Internet. However, the specification was deficient and newer standards have been developed. Basically, PEM was created by the IETF to act for email in a similar fashion as IPSEC does to IP.

PEM provides for:

- Authentication provided by MD2 or MD5
- Message integrity using the X.509 standard for certificate structure and format
- Encryption provided by DES in Cipher Block Chaining (CBC) mode
- Key management provided by RSA signatures

When Multipurpose Internet Mail Extension (MIME) was introduced as a way to add binary attachments to e-mail, PEM became less important because of its support for only 7-bit text messages. PEM was then extended with MIME Object Security Standard (MOSS), a protocol with PEM compatibility and support for MIME attachments. However, MOSS was difficult to implement and use.

S/MIME has now become the de facto standard for securing mail messages that has generally replaced PGP (Pretty Good Privacy) and PEM (Privacy-Enhanced Mail).

S-HTTP vs. HTTPS

The **Secure Hypertext Transport Protocol (S-HTTP)** was developed to provide secure communication over HTTP. S-HTTP encrypts messages with session keys to provide integrity and sender authentication capabilities. It can also support multiple encryption modes and types. S-HTTP uses public key technology and symmetric encryption when an individual message needs to be encrypted.

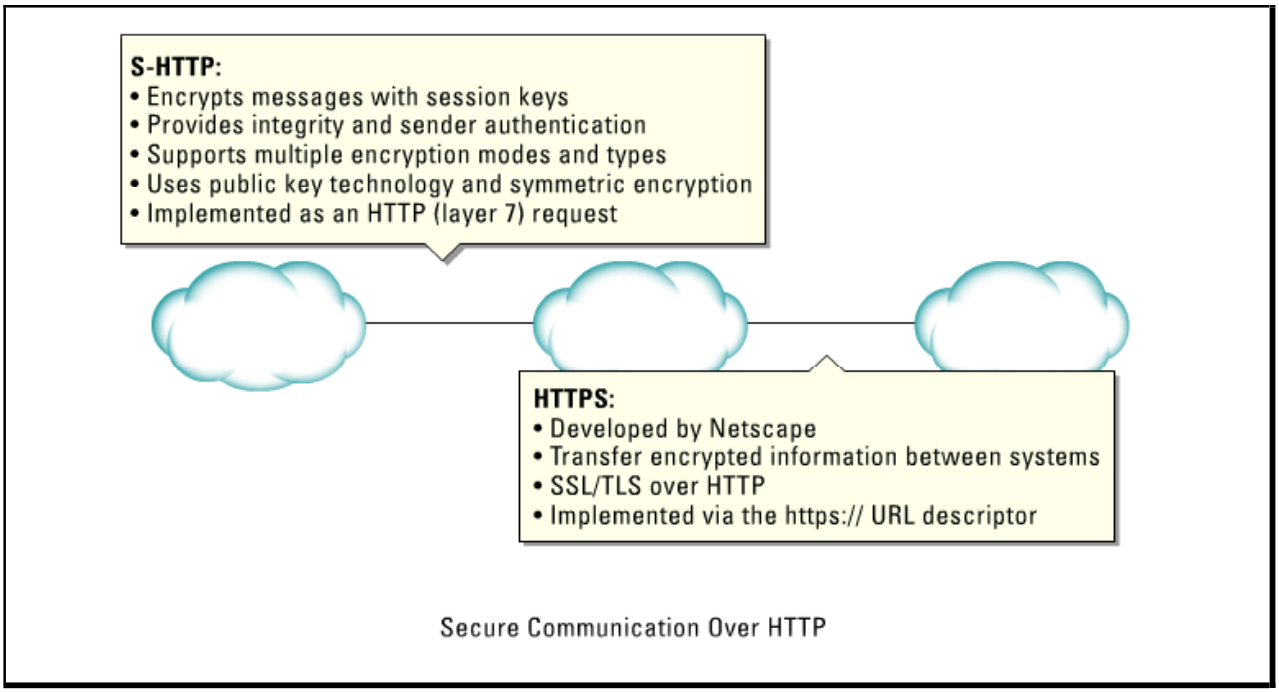

Secure Communication Over HTTP

A Secure HTTP message is a request or status line, followed by other headers that must be RFC-822 compliant, and some content. The content can be raw data, a Secure HTTP message, or an HTTP message. The request line is defined as:

```
Secure * Secure-HTTP/1.1 to which the response must be:
      Secure-HTTP/1.1 200 OK
```

The lines' definition precludes an attacker from seeing the success or failure of a given request. Secure HTTP takes a generally paranoid attitude to all information, leaking as little as possible.

S-HTTP has been overtaken by the SSL protocol (HTTPS). However, HTTPS and S-HTTP have very different designs and goals, so it is possible to use the two protocols together. Whereas HTTPS is designed to establish a secure connection between two computers, S-HTTP is designed to send individual messages securely.

The Secure Hypertext Transfer Protocol (HTTPS) developed by Netscape is a communications protocol designed to transfer encrypted information between computers over the World Wide Web. HTTPS is HTTP using a Secure Socket Layer (SSL). A secure socket layer is an encryption protocol invoked on a Web server that uses HTTPS. HTTPS has been designated the use of TCP port 443.

Most implementations of the HTTPS protocol involve online purchasing or the exchange of private information. Accessing a secure server often requires some sort of registration, login, or purchase.

For example, a customer may use a browser to visit a Web site to view an online catalog. When he is ready to order, he will be given a Web page order form with a Uniform Resource Locator (URL) that starts with https://. When he clicks Send, to send the page back to the catalog retailer, the browser's HTTPS layer will encrypt it. The acknowledgement the customer receives from the server will also travel in encrypted form, arrive with an https:// URL, and the browser's HTTPS sublayer will decrypted the form.

HTTPS and SSL support the use of X.509 digital certificates from the server so, if necessary, a user can authenticate the sender.

MIME

The **Multipurpose Internet Mail Extensions (MIME)** protocol is defined in RFC 1521 and RFC 1522.

When e-mail first came into existence, email messages were meant to be pure text only messages. As the Internet started to grow, people wanted to share more that just text. They also wanted to share graphic files, audio files, and HTTP, but they didn't want them to be seen as attachments; they wanted them to be seen dynamically when the email document was opened. The problem manufactures had was in regards to the multitude of graphic and audio formats, as well as the different platforms and operating systems available. A standard was needed in order for all platforms to display email messages in the same manner across systems. That standard turned out to be the Multipurpose Internet Mail Extensions (MIME) protocol.

MIME allowed a one-time modification to email reading programs that enabled the program to display a wide variety of types of messages. This email extension allows the user to view dynamic multi-type email messages full of color, sound, animations, and moving graphics.

S/MIME

MIME allowed emails to display features that they never could have before without regard to security. Confidential email, although quite impressive, was still subject to the same old hacks, such as sniffing and replay. Businesses required a secure way of sending MIME data, one that guaranteed that confidentiality, and integrity of the email message was kept. **Secure Multipurpose Internet Mail Extensions (S/MIME)** was created for this purpose.

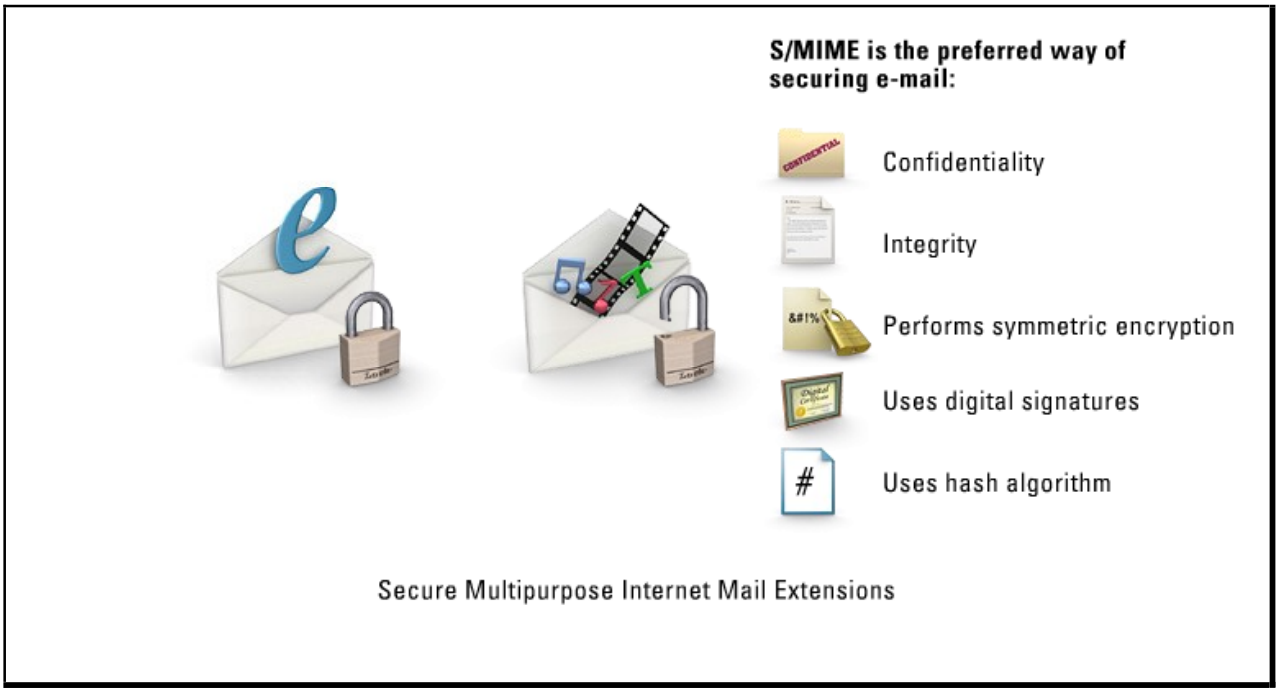

S/MIME is a standard for encrypting and digitally signing electronic mail that contains attachments and providing secure data transmissions. S/MIME provides confidentiality through the user's encryption algorithm with RC2, DES, 3DES; integrity through the user's hashing algorithm with MD5, SHA1; authentication through the use of X.509 public key certificates; and non-repudiation through cryptographically signed messages with RSA.

Using S/MIME is the preferred way of securing e-mail, as it traverses the unfriendly world of the Internet. S/MIME version 2 is described in RFC 2311, and S/MIME version 3 is described in RFC 2633.

Pretty Good Privacy (PGP)

Pretty Good Privacy (PGP) is a technology created by Phil R. Zimmermann in response to the 1991 Senate Bill 266. This ominous anti-crime bill had a measure in it that stated all encryption software must have a back door to allow the U.S. Government to decrypt messages sent between parties.

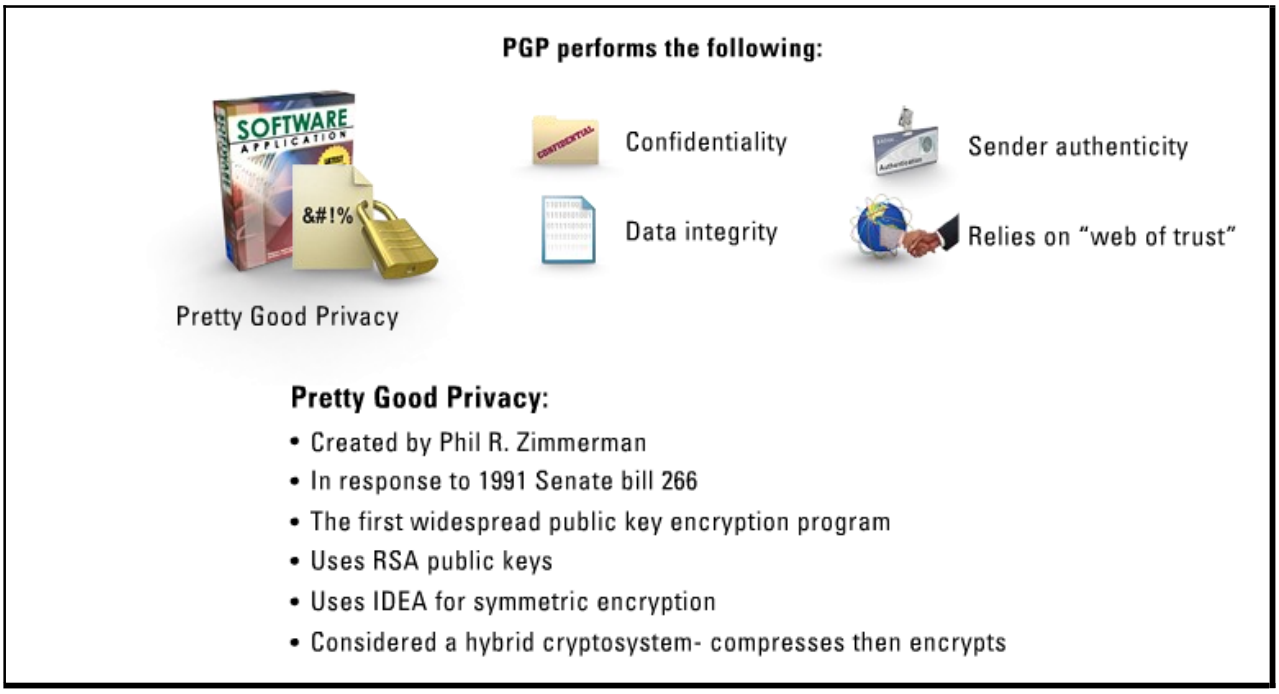

Being a staunch supporter of civil rights, Phil created a crypto system in which no one except the two parties could read their email messages. PGP became the first widespread public key encryption program. It uses RSA public key encryption for key management, and IDEA symmetric cipher for bulk encryption of data.

PGP works using a public key cryptosystem. With this method, each party creates an RSA public/private key pair. One of these keys is kept private, the private key, and one is given out to anyone in the public Internet, the public key. What one key encrypts, only its partner private key can decrypt. This approach means if user X obtains user Y's public key and encrypts a message destined to user Y using its public key, the only person in the universe who can decrypt the message would be user Y, as he has the corresponding private key.

PGP is a hybrid cryptosystem because before encryption is performed, the email data is first compressed. Compression not only makes an email message smaller, it also removes any patterns found in plain text that mitigates many cryptanalysis techniques that look for these patterns. PGP performs the following security measures: Confidentiality, data integrity, and sender authenticity.

PGP relies on a web of trust in its key management approach. Everyone can generate a PGP key by himself. If you want to know if a given key really belongs to the person stated in the key, you have to verify that information. Finding the person is very easy if a person knows who the person is who created the key. If that person is unknown, finding the key can be difficult. PGP provides a method to establish trust with a certain public key. Each public key you obtain will be assigned a trust level when you place it in your key ring. This trust level tells PGP how much you trust key certificates done by this key. The

lower the trust value the more key certificates are necessary to validate a key. Trust levels include the following:

- **Untrusted** - Key certificates using this value are ignored
- **Marginal** - At least 2 keys with marginal trust have to sign another, a third key, to make this third key a valid key
- **Complete** - At least one key with complete trust has to sign another key to make the key valid
- **Ultimate** - If you have a secret key for a public key, this key is ultimately trusted. Every key you sign with an ultimate trusted key becomes valid

As keys get signed and submitted to PGP key servers, others can benefit from such signatures. All these signatures build a kind of web, which is why we call it a web of trust.

Cookies

Netscape Corporation developed a mechanism that helps alleviate the stateless nature of the HTTP protocol. As users surf a particular web site, each request is treated as an entirely new interaction. The web server has no idea if a particular request is from a new client or from an existing client walking their site. This stateless behavior makes it difficult for web sites to create things like shopping carts that must remember what items were ordered, how many items, and the part number of the item, as the user moves from page to page identifying new material to purchase. Cookies were created to solve this problem.

Cookies:

- Created by Netscape
- A way to alleviate the stateless nature of HTTP
- Tiny text files placed on a computer system
- Identifies individual, information entered in a form, etc.
- Cannot gather information
- Can be used to store confidential information or track movements

Cookies are tiny text files that many web sites place on the user's computer to help identity the individual or store some type of information, such as data you entered in a form. The Cookies themselves are recipients of information. They cannot gather information. Only programs can gather information. Cookies are written and read on a computer by either java scripts included in the HTML code, or by commands sent by server side programs.

Cookies can be used to store much more than harmless time saving data; however, they can be used to store confidential information or track the movements a user makes through a web site or sites. For this reason, many people do not like or allow cookies to be created on their system.

If cookies store authentication parameters, session IDs and information, they can greatly assist the web server because the server never has to perform a database lookup. The browser sends all the relevant information to the server. The problem is that cookies are sent in clear text, and if a cracker sniffs the wire and obtains this information, they have all the proper credentials they need to authenticate to the server.

RAID

The Redundant Array of Inexpensive Disk (RAID) system is used for fault tolerance against hard drive crashes. Its secondary benefit includes improved system performance because data can be simultaneously gathered from multiple hard drives concurrently.

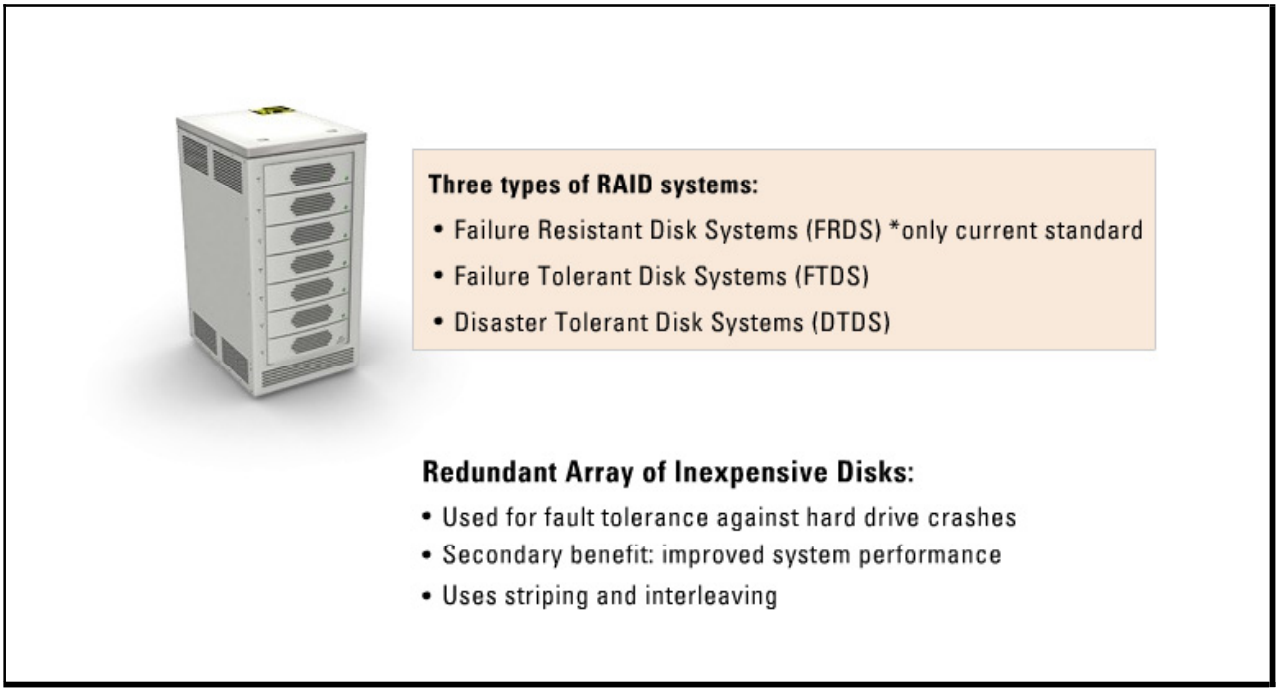

RAID employs the technique of striping that partitions each drive's storage space into units ranging from a sector of 512 bytes up to several megabytes. The stripes of all disks are then interleaved and addressed in order. RAID can be performed in both hardware and software implementations.

The RAID advisory board identifies three types of RAID systems

- **Failure Resistant Disk Systems** (FRDS) - the only current standard. Provides the ability to reconstruct the contents of a failed disk onto a replacement disk.
- **Failure Tolerant Disk Systems** (FTDS) – a system that protects against loss of data due to failure of any single component.
- **Disaster Tolerant Disk Systems** (DTDS) – a system that consists of two or more independent zones, either of which could provide access to stored information

The characteristics of the three classes of RAID include:

FRDS (meets criteria 1 to 6 at a minimum)

1.Provides protection against data loss and loss of access due to a drive failure

2.Provides for reconstruction of failed drive content to a replacement drive

3.Provides protection against data loss due to a "write hole"

4.Provides protection against data loss due to a host or a hosts I/O bus failing

5.Provides protection against data loss due to a replaceable unit that has failed

6.Provides for replaceable unit monitoring and failure indication

FTDS (meets criteria 1 to 15 at a minimum)

7.Provides for automatic disk swap and hot swap capabilities

8.Provides protection against data loss due to a cache failure

9.Provides protection against data loss due to an external power failure

10.Provides protection against data loss due to temperatures above maximum operating range

11.Provides for replaceable unit and environmental failure warning

12.Provides protection against loss of access to data due to device channel failure

13.Provides protection against loss of access to data due to controller module failure

14.Provides protection against loss of access to data due to cache failure

15.Provides protection against loss of access to data due to power supply failure

DTDS (meets criteria 1-21 at a minimum)

16.Provides protection against loss of access to data due to a host or hosts I/O bus failing

17.Provides protection against loss of access to data due to an external power failure

18.Provides protection against loss of access to data due to component replacement

19.Provides protection against loss data and loss of access to data due to multiple disk failures

20.Provides protection against loss of access to data due to a zone failure

21.Provides for long-distance protection against loss of data due to zone failure

RAID Levels

This topic discusses RAID levels.

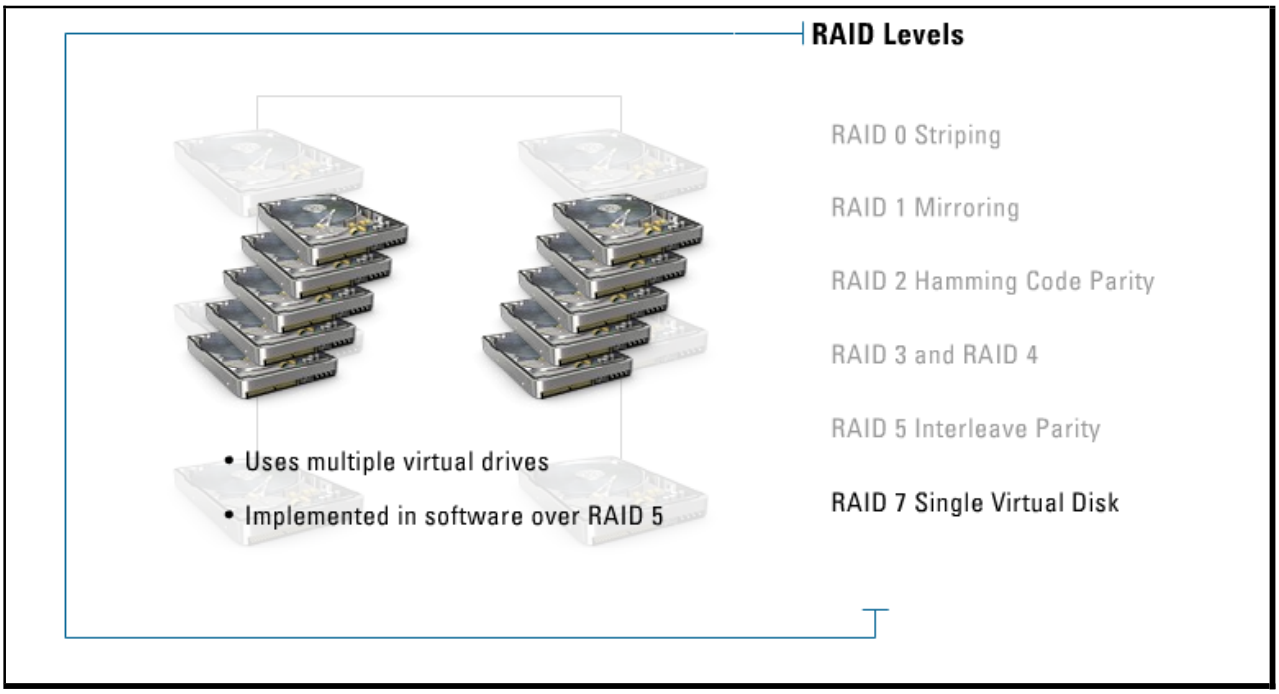

RAID 0Striping creates one large logical disk by striping or reading multiple disks at the same time. This level improves system performance dramatically as it provides parallel reading and writing on serial hard drive systems. It does have some draw backs as there is no redundancy and no fault tolerance.

RAID 1Mirroring duplicates data on other disks, usually in a one-to-one ratio. This level provides both redundancy of data and fault tolerance, but comes at a heavy price as you double the cost of storage media.

RAID 2 Hamming Code Parity provides parity information created by using a special hamming code with multiple disks for fault tolerance. RAID level 2 was usually implemented in a rigid 39 disk array system, where 32 disks were used for data, and 7 disks were used for recovery purposes. RAID 2 is not used anymore, as it has been replaced by more flexible levels.

RAID 3 and RAID 4 are used to stripe data across multiple drives (RAID 0), while providing parity (RAID 2) across all drives. This level provides increased system performance while providing redundancy, but in a single parity drive implementation performance can be drop significantly.

RAID 3 – Byte level

RAID 4 – Block level

RAID 5Interleave Parity allows reads and writes across all drives to be implemented concurrently and usually implemented in a 3-5 drive system. It is the most popular RAID implementation, as it uses RAID 0 striping and RAID 2 parity interleaving. If one drive in a RAID 5 system fails, data from the failed drive can be reconstructed using information from the remaining drives.

RAID 7 Single Virtual Disk allows multiple drives to function as a single virtual disk. This level of RAID is usually implemented in software over RAID level 5 hardware and enables the drive array to continue to operate if any disk or any path to any disk fails.

Summary

The key points discussed in this lesson are:

- Secure Electronic Transactions (SET)
- Privacy Enhanced Mail (PEM)
- The difference between S-HTTP and HTTPS
- Features of MIME
- Features of S/MIME
- Premise of Pretty Good Privacy (PGP)
- Cookies
- Use of RAID
- RAID levels

Security Management Practices

Overview

Security Management defines many security related management issues that an enterprise will undertake. Security issues such as how to classify data, how to implement change control, what to list as risks to the enterprise, and how to implement countermeasures are but a few of the numerous issues security must manage.

Objectives

Upon completing this module, you will be able to:

- Explain change control management
- Explain data classification
- Recognize good employment policies and practices
- Explain risk management
- Explain each role and its responsibilities

Outline

The module contains these lessons:

- Security Overview
- Data Classification
- Employment Policies and Practices
- Risk Management
- Roles and Responsibilities

Security Overview

Overview

There are three core components that make up the foundation of a corporation's security program: risk management, security policies, and security education. These three components make up the core of security management, and therefore, are central to a company's computer and information security. This lesson will discuss the goal of security management, control measures, and change control management techniques.

Importance

Without a thorough understanding of security management, the information security professional has little chance of securing the enterprise as a whole. This module will discuss the practices every security professional must master to provide an effective security foundation to the enterprise.

Objectives

Upon completing this lesson, you will be able to:

- List security objectives and definitions
- Explain change control and management
- Define controls
- Define administrative controls
- Define technical controls
- Define the physical controls
- Explain a planning horizon

Outline

The lesson contains these topics:

- Security Objectives and Definitions
- Change Control and Management
- Controls
- Administrative

- Technical
- Physical
- Planning Horizon

Security Objectives and Definitions

The goal of security management is to protect the company's assets.

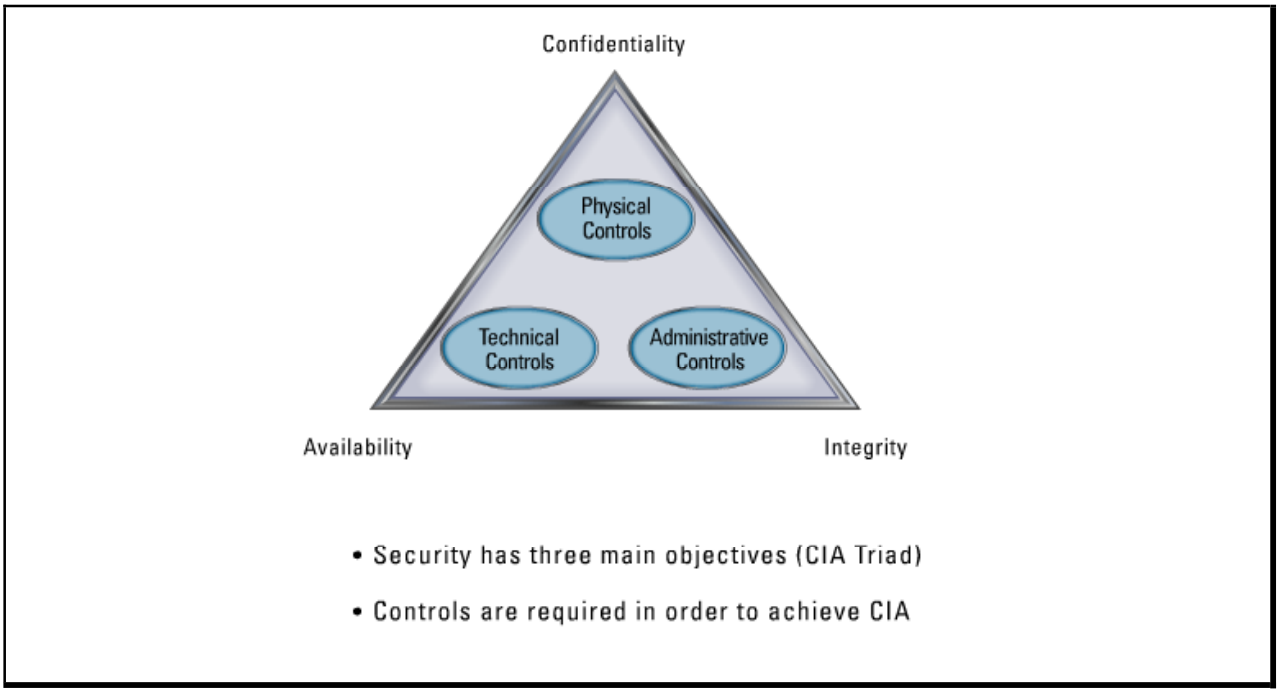

The goal of security management is to protect the company's assets. Risk management identifies these assets, discovers the risks that threaten them, and provide an estimate if possible damage or loss occurs. The result of performing risk analysis is the development of applicable security policies that should be in place to protect the company's assets. Security education then disperses the relevant security policies to each and every person in the company.

Security management utilized three control types to protect the company:

1. **Administrative controls** - are controls that include the development and publication of security policies, standards, procedures, personnel screening, system activity monitoring, change control procedures, and security awareness training.
2. **Technical controls** - are controls that consist of logical access control mechanisms, password and resource management, identification and authentication methods, configuration of the network, and relevant security devices.
3. **Physical controls** - are controls that allow individual access into a facility, locking systems, protecting the perimeter of the facility, intrusion detection, and environmental controls.

Fundamentally, security management is concerned with three objectives identified as the CIA triad:

- **Confidentiality** - The protection of information systems to prevent unauthorized persons, resources, and processes to obtain access.
- **Integrity** - The protection of information systems from intentional or accidental unauthorized change.
- **Availability** - The assurance that any information system can be accessible by authorized personnel whenever needed.

Definitions that you must know include the following items:

- **Vulnerability** - Is a software, hardware or procedural weakness that may provide the attacker the open door he is seeking to enter a computer or network to have unauthorized access to resources within the environment.
- **Threat** - Is any potential danger to information or systems.
- **Risk** - Is the likelihood of a threat agent taking advantage of a vulnerability.
- **Exposure** - Is an instance of being exposed to possible loss from a threat agent
- **Countermeasure or safeguard** - Mitigates the potential risk.
- **Risk Management** - Is the process of identifying, assessing, reducing risks to an acceptable level, and implementing the right mechanisms to maintain that level of risk.
- **Top-down approach** - An approach when the initiation, support, and direction of security come from top management and work their way through middle management to staff members.
- **Bottom-up approach** - An approach when security programs are developed by IT without getting proper management support and direction.
- **Operational goals** - Daily goals
- **Tactical goals** - Mid-term goals
- **Strategic goals** - Long-term goals

Change Control and Management

Undocumented or unapproved changes to production systems can severely impact the security integrity of any information system. Change control management helps to alleviate these problems by authorizing changes to production systems.

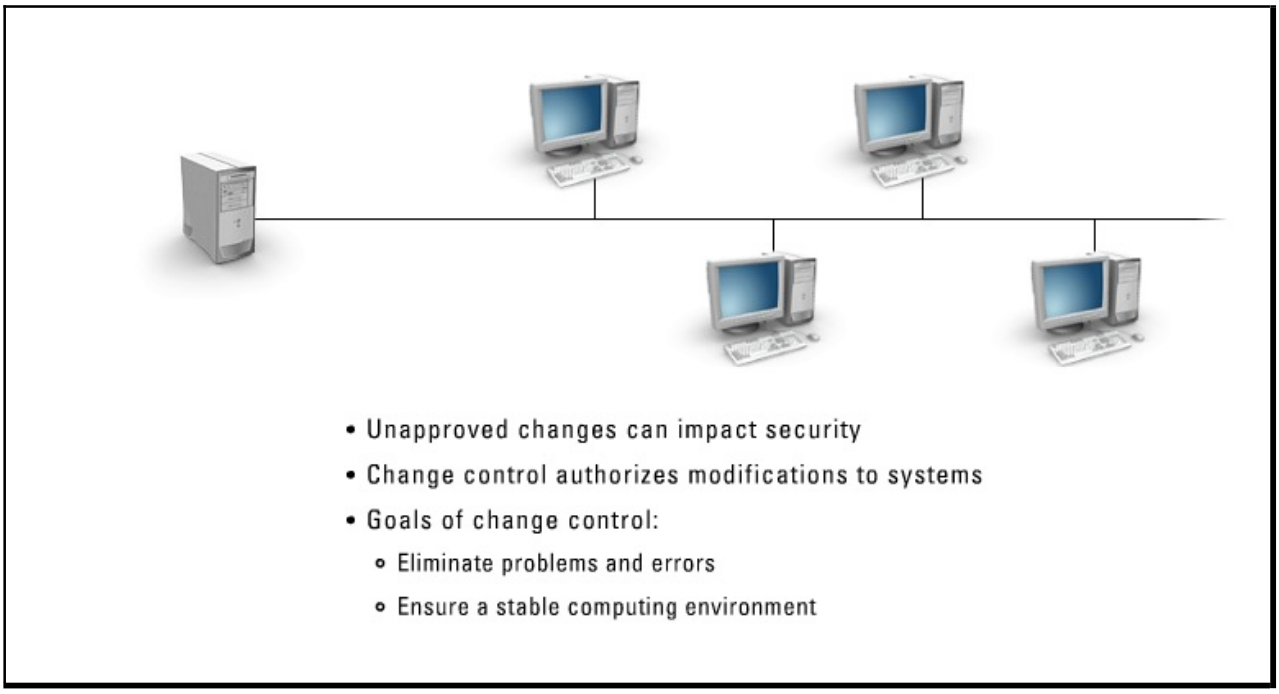

These changes can include items such as OS and application software, modifications of existing applications, upgrading or patching software, installing new images for routers, and adding switches and firewalls. The purpose of change control is to manage changes within the computing environment.

The following is an example of a generic change control outline:

1. Change request is submitted to management
2. The change request is analyzed and validated
3. Implementation of the change is analyzed
4. Implementation costs are analyzed
5. Analysis and change recommendations are documented
6. Change control board approves or disapproves final change request
7. Changes are made and relevant documentation reflects changes
8. Quality control validates and approves change

The goal of change management is to eliminate problems and errors and ensure the entire computing environment is stable. To meet these goals, it is important to implement the following:

- Ensure orderly change
- Inform relevant computing members of changes taking place
- Analyze changes
- Reduce possible loss of service (changes cause system outages)

- Application process required for any change
- Cataloging the change
- Scheduling the change
- Implementing the change
- Reporting change to proper personnel

Controls

The operation environment in an information system requires the use of controls to manage the organizations information security environment.

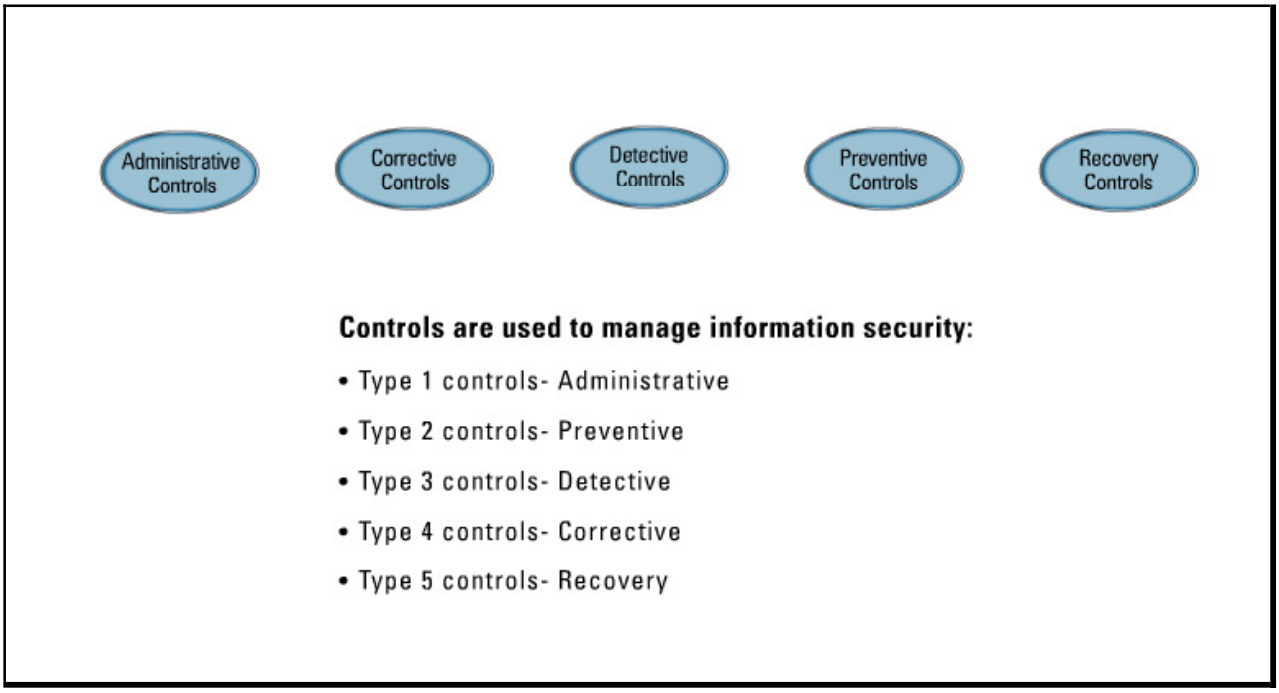

The following is a list of control measures.

- **Administrative controls** - Type 1 controls advise employees of the expected behavior when interacting with the company's information systems; they are usually implemented as company policies and guidelines.
- **Preventive controls** - Type 2 controls preclude actions violating company policy. Included in preventive controls are the following:
  - Administrative controls
  - Technical controls
  - Physical controls
- **Detective controls** - Type 3 controls identify and potentially react to security violations. Detective controls involve the use of practices, processes, and tools.
- **Corrective controls** - Type 4 controls react to detection of an incident in order to reduce or eliminate the opportunity for the unwanted event to occur. Examples include procedures for keeping anti-virus software up to date, additional training for members of the security force, additional security awareness training for employees, use of two-factor authentication, and implementation of more sophisticated firewalls.
- **Recovery controls** - Type 5 controls restore a system or operation of the system when an incident occurs that results in the compromise of integrity or availability of a computing system. Recovery controls are often associated with business continuity and disaster recovery planning. Other examples include fault tolerant systems, RAID, and resending lost or corrupted messages.

Administrative

Administrative controls include the creation of policies, standards, procedures, and guidelines. They also include the screening of employees and non-employees, security awareness training, system activity monitoring, and change control procedures.

A security policy is an administrative control used to provide a high-level plan stating management's intentions about how security should be practiced within an organization. This policy identifies what actions are acceptable, and what level of risk the company is willing to accept.

Personnel controls are a form of administrative control that explains to employees the expected behavior when interacting with security mechanisms. Personal controls include noncompliance issues pertaining to these expectations.

- Separating of duties is a personnel control enforced, so no one person can carry out critical tasks that could prove to be detrimental to the company.
- Rotation of duties is a personnel control that ensures that more than one person can fulfill the obligations of more than one position.

Security awareness training is an administrative control that trains personnel on any new or old technology. The training informs personnel on what may detrimental to the company and how to mitigate those threats.

Supervisor structure is an administrative control that ensures that management staff members are responsible for employees and take a vested interest in their actions. This control is used to aid in the fight against fraud as well as enforcing proper access control.

Testing is an administrative control used to test all security controls and mechanisms on a periodic basis to ensure they properly support the security policy, goals, and objectives set for them.

Technical

Technical controls include logical access control mechanisms, password and resource management, identification and authentication controls, security devices, and the configuration of the network.

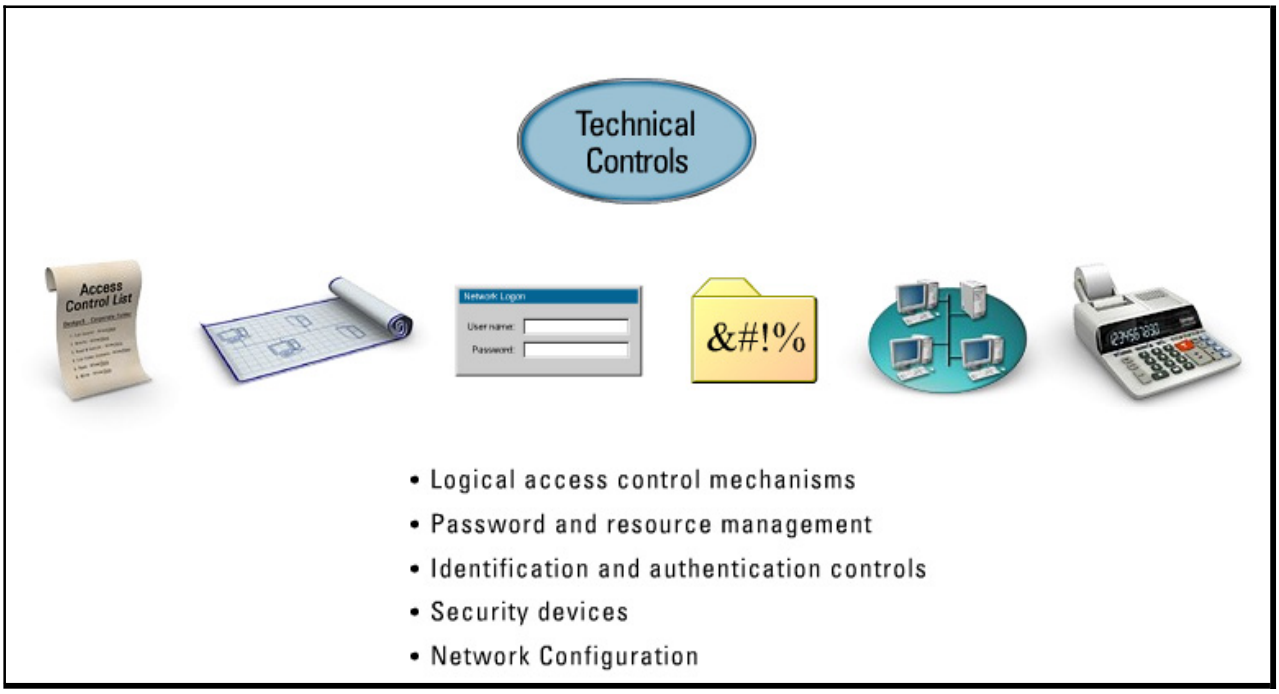

System access controls are a form of technical control used to limit system accessed. If the system is using a MAC-based architecture, the clearance level of a user must be identified and compared to the resource classification level. If the system is using a DAC- based architecture, the operating system must check whether a user has been granted access to the resource. Other types of system access controls include Kerberos, TACACS+, and RADIUS.

Network architecture is a technical control used to provide segregation and protection of a computing environment. Networks can be physically segregated by walls and location. They can be logically segregated by implementing VLANs and using filtering firewalls.

Network access is a technical control that limits who can and who cannot access a system and what an individual can do when they have access.

Encryption is a technical control used to protect information as it passes over an untrusted medium or while stored on a system.

A Control zone is a technical control that specifies a specific area that surrounds and protects network devices that emit electrical signals.

Auditing is a technical control used to track activity within a network or on a computing device.

Physical

Physical controls include controlling individual access into a facility, locking systems, and removing unnecessary mass storage devices such as floppy or CD ROM drives. Physical controls protect the perimeters of the facility, monitor for intrusion, and control the environment.

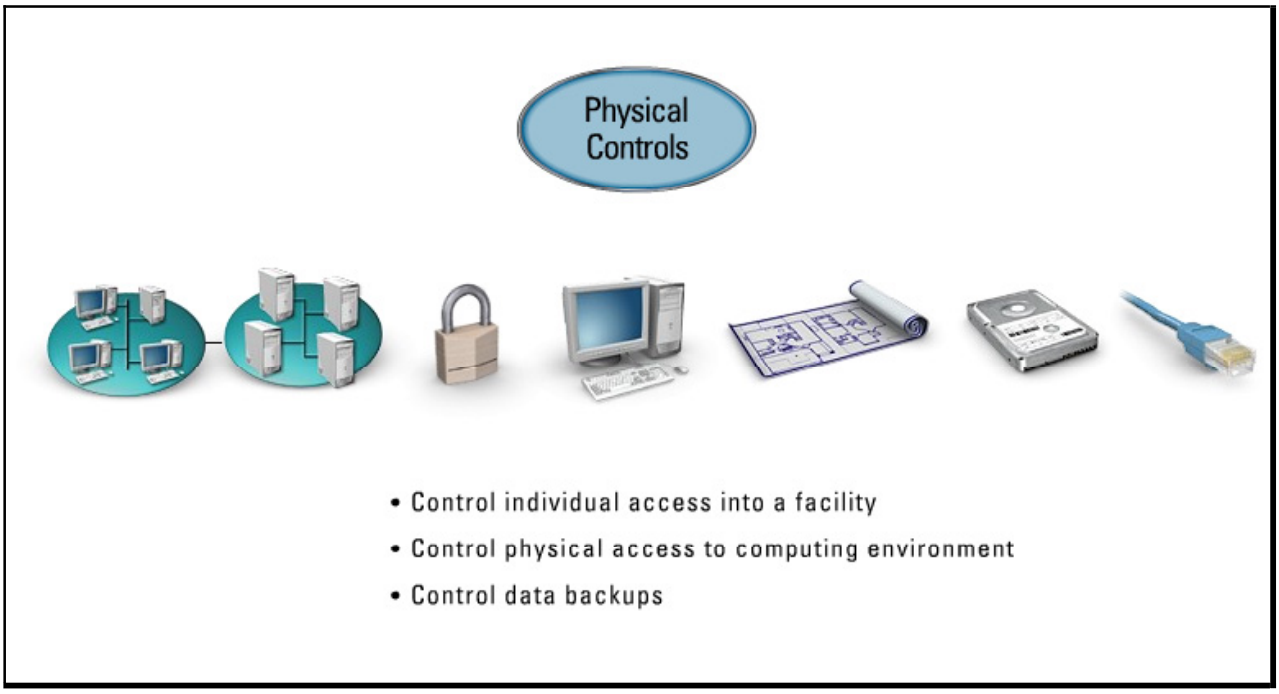

Network Segregation is a physical control that segregates computing devices such as server farms, HR department information, IT department records, routers and switches. Each area would have the necessary physical controls to ensure that only the right individuals can physically enter those areas.

Perimeter security controls the perimeter of a facility. Security can take the form of security guards viewing badge information, closed circuit TVs that scan employee parking lots, fences that surround a building, motion detectors, sensors, and alarms.

Computer controls are physical controls that eliminate improper access to or from a computer. Examples would be locks on covers, removing floppy or CD ROM drives, or reducing the electrical emissions from the computer.

Work area separation is a technical control that limits the number of individuals can access certain areas of a facility. Security badges and keys are examples of work area separation.

Data backups are physical controls providing access to information in cases of an emergency or a disruption of the computing system.

Cabling is a physical control to maintain cable integrity in the network. The security professional could eliminate any instances of crosstalk, exposed cable, or possibilities of cables being cut, burnt, crimped or sniffed.

Planning Horizon

A planning horizon is defined as the length of time a security model projects into the future.

Security can be broken down in three strategic goals that stretch forward in time:

- Daily goals
- Mid-term goals
- Long-term goals

Summary

The key points discussed in this lesson are:

- Security objectives and definitions
- Change control and management
- Controls
- Administrative controls
- Technical controls
- Physical controls
- Planning horizon

Data Classification

Overview

To secure data by any method requires that the security professional classify the data in some fashion. The data itself can have a security marking, the location of the data in a certain folder or directory can classify the security level of the data or some other method can be used. This lesson will identify the objectives and criteria of data classification as well as some of the more common classification methods.

Importance

Understanding why and how data is classified is important in protecting the data from threats.

Objectives

Upon completing this lesson, you will be able to:

- Name the objectives of a classification scheme
- List the criteria by which data is classified
- List the major classification criteria
- Name the information classification roles

Outline

The lesson contains these topics:

- Classification Scheme Objectives
- Data Classification Criteria
- Major Classification Criteria
- Information Classification Roles

Clas

most cost-effective manner.

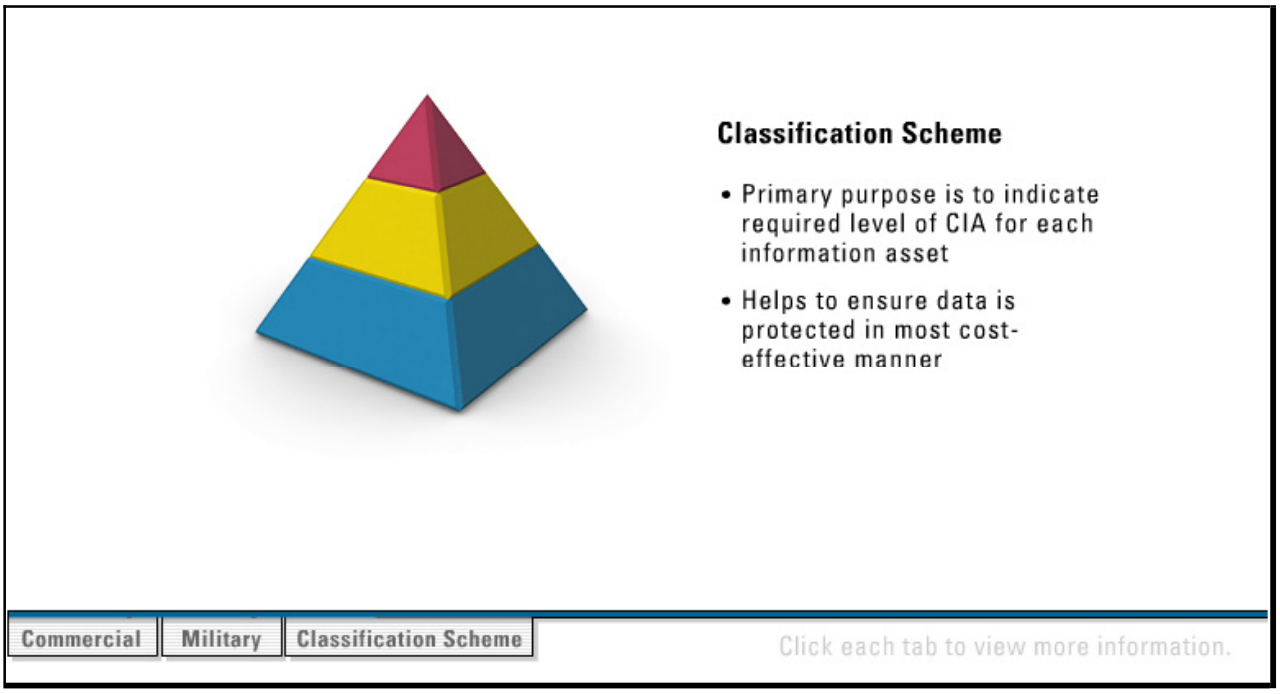

Data owners are responsible for defining the security level of the data.

Common classification levels (from highest to the lowest level):

- Commercial business
  - Confidential
  - Private
  - Sensitive
  - Public
- Military
  - Top secret
  - Secret
  - Confidential
  - Sensitive but unclassified
  - Unclassified

Data Classification Criteria

In order to properly classify an information asset, all parties must agree on a criteria scheme to identify what information is classified into which security level.

Organizations must come up with their own criteria, but some of the more common criteria parameters are listed here:

- Data usefulness
- Age of the data
- Monetary value of the data
- Damage that could be caused if data was disclosed
- Damage that could be caused if data was altered or corrupted
- Regulatory laws required on specific data
- Effects of data on national security
- Who are the data owners or manipulators
- Where data is stored
- Who has permission to backup data

After the data is classified based on its criteria, the security professional needs to determine how each classification should be protected. Items such as access control, identification, and labeling need to be specified with an agreement on how the data in a specific classification will be stored, maintained, backup up, transmitted, and destroyed.

Major Classification Criteria

This topic discusses major classification criteria.

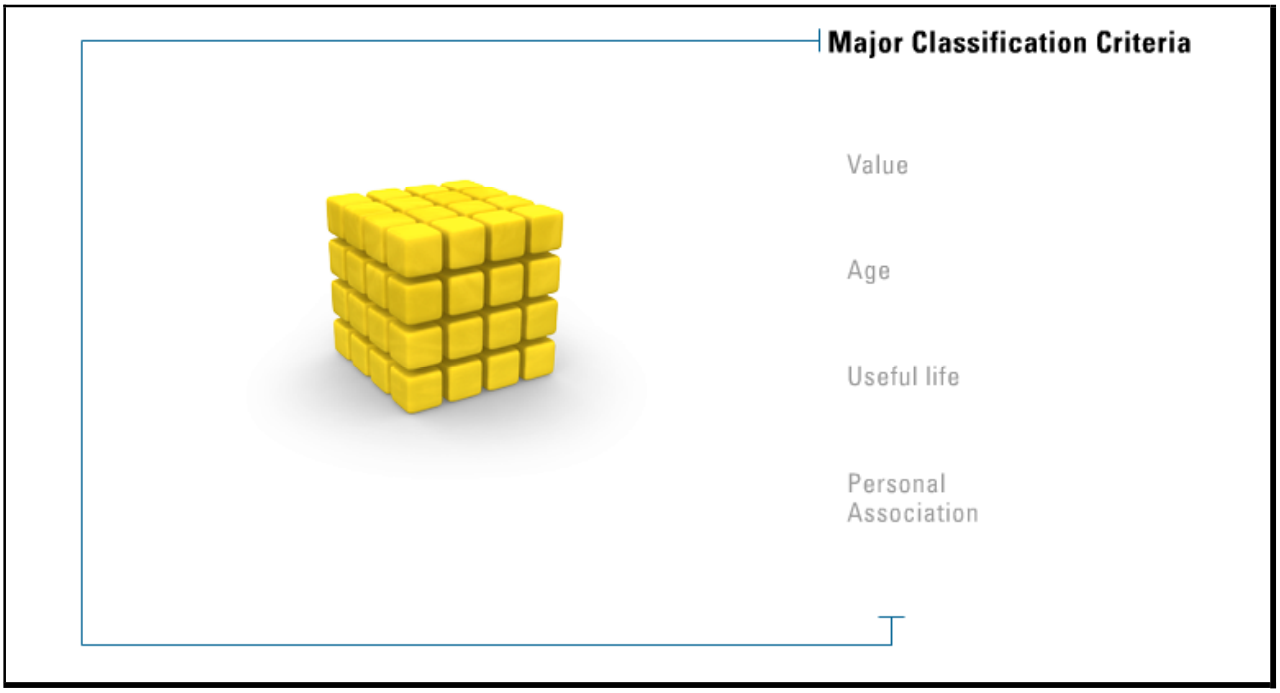

The following include the major classification criteria:

- **Value** - number one criteria, if it is valuable it should be protected
- **Age** - value of data lowers over time, automatic de-classification
- **Useful life** - if the information is made obsolete it can often be de-classified
- **Personal Association** - if the data contains personal information it should remain classified

Sharing classified information may be required in the event of the following:

- **Court order** - May be required by a court order
- **Government contracts** - Government contractors may need to disclose classified information
- **Senior level approval** - Senior executives may approve release

Information Classification Roles

This topic discusses information classification roles.

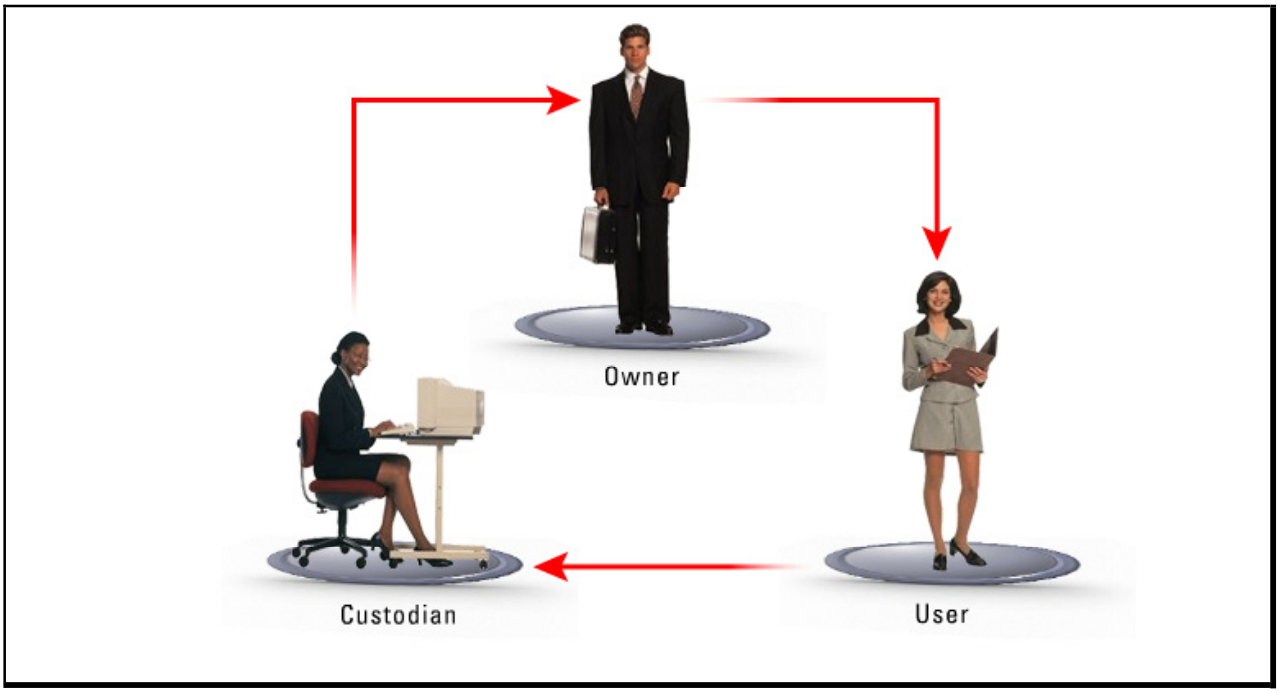

Information classification roles include the following:

- Owner
  - May be executive or manager
  - Has final corporate responsibility of the data protection
  - Makes determination of classified level
  - Reviews classification level regularly for appropriateness
  - Delegates responsibility of data protection to the Custodian
- Custodian
  - Generally IT systems personnel
  - Runs regular backups and testing recovery
  - Performs restoration when required
  - Maintains records in accordance with the classification policy
- User
  - Is anyone that routinely uses the data
  - Must follow operating procedures
  - Must take due care to protect the information
  - Must use computing resources of the company for company purposes only

Summary

The key points discussed in this lesson are:

- Classification scheme objectives
- Data classification criteria
- Major classification criteria
- Information classification roles

Employment Policies and Practices

Overview

Hiring an individual seems like a very simple thing to do, but as anyone in the Human Resources team can testify, hiring can be very involved. Items like a security clearance, badge creation, domain authentication credentials, and a home directory must be identified and completed, all without an error that could lead to a breach in security. All this work is just for hiring an employee. Terminating an employee can be just as involved. This lesson will discuss the various policies and procedures for dealing with any aspect of employee relations.

Importance

Understanding the security policies, procedures, and guidelines when hiring or terminating an employee is essential to decrease the possibility of a malicious or accidental internal attack. User awareness during time of employment and role responsibilities are also essential to bound resource use in the enterprise.

Objectives

Upon completing this lesson, you will be able to:

- Explain the purpose of policies, standards, guidelines, and procedures
- Name types of objectives
- Name policy types
- Define regulatory policies
- Define advisory policies
- Define informative policies
- Define baselines and guidelines for policies
- Explain background checks and security clearances
- Guidelines for employment agreements
- Purpose of security awareness programs

- Explain security's role in hiring practices
- Explain security's role in the terminating practices
- Explain security's role in job descriptions
- Explain role responsibilities for security
- Explain the separation of duties and responsibilities
- Explain the purpose of job rotations

Outline

The lesson contains these topics:

- Policies, Standards, Guidelines, and Procedures
- Objectives
- Policy Types
- Regulatory
- Advisory
- Informative
- Baselines and Guidelines
- Background Checks and Security Clearances
- Employment Agreements
- Awareness Programs
- Hiring Practices
- Terminating Practices
- Job Descriptions
- Role Responsibilities
- Separation of Duties and Responsibilities
- Job Rotations

Policies, Standards, Guidelines, and Procedures

The CIA triad defines the objectives of any security program. The threats to these objectives include unauthorized access, modification, destruction, and disclosure of company assets. These threats can occur accidentally, or they can be perpetrated intentionally.

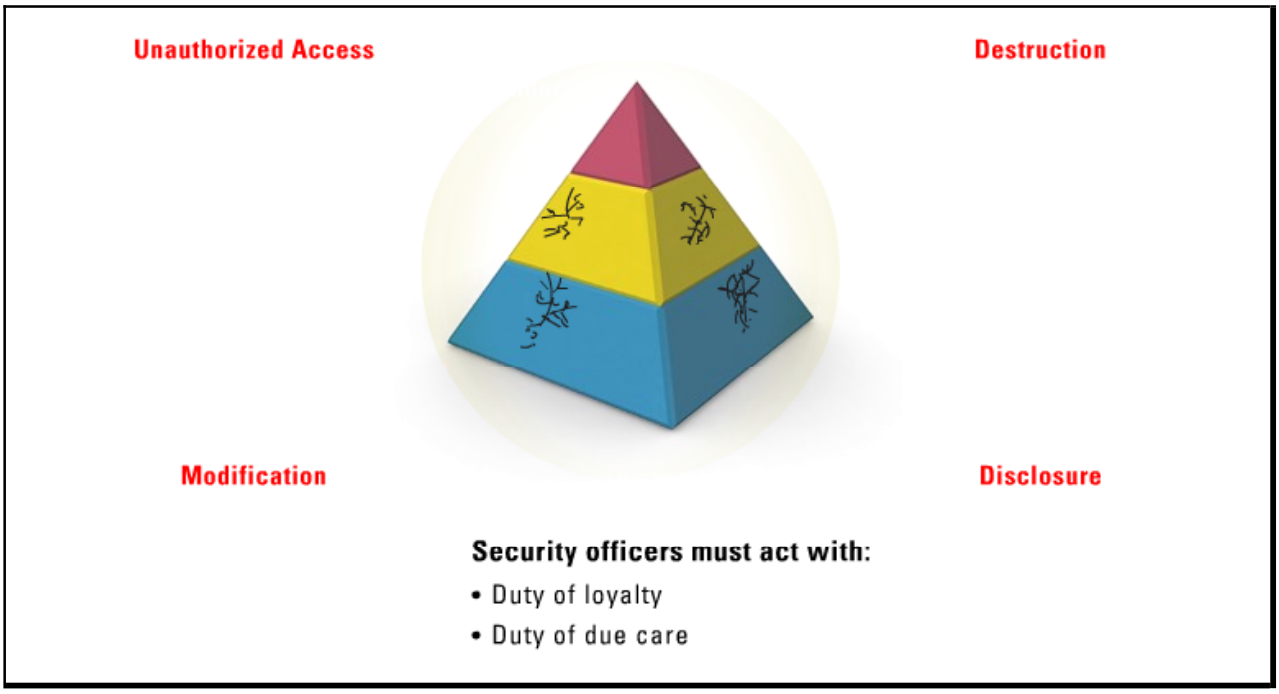

In order for the goals of the CIA triad to be realized, all employees must understand the corporate security strategies. Providing company policies, standards, guidelines and procedures for all employees in a company will help to ensure greater security.

The company charges the security group, led by the Information Security Officer, with the responsibility of creating these documents and disseminating them throughout the corporation. Remember, many corporations have legal requirements in the form of laws and regulations to protect private records from becoming public. Generally, security officers and their directors are required to perform two specific duties:

- Duty of loyalty
- Duty of due care

Objectives

The objectives to carry out the specific duties of loyalty and due care are in place to protect the company and its consumers.

Senior Executive

Potentially Profitable Company Information

Duty of Loyalty and Due Care Objectives:

- Used to protect the company and its consumers
- Based on legal concepts:
  - Conflict of interest
  - Confidentiality
  - Duty of fairness
  - Corporate opportunity
- Security officers and directors need to act:
  - In good faith
  - In the best interest of the company
  - With proper due care and due diligence

The basic principle of the duty of loyalty is to ensure that a senior manager could not use his position to make personal profit or gain by means of the protected information. Duty of loyalty and due care policies use certain legal concepts:

- **Conflict of interest** - An individual must divulge any outside interests or relationships that might conflict with the company's interests.
- **Confidentiality** - Any company matter must be kept confidential until it is made public by the appropriate personnel.
- **Duty of fairness** - When an individual is presented with a conflict of interest, the individual has the obligation to act in the best interest of all parties.
- **Corporate opportunity** - When an individual is presented with advanced notice on mergers, acquisitions, or patents, the individual may not use this information for personal gain.

Security officers and directors also assume a duty to act carefully in fulfilling important company tasks. In general these personnel need to act in the following ways:

- In good faith
- In a way he or she reasonably believes is in the best interest of the company
- With the proper care that an ordinarily prudent person in a like position would exercise in similar circumstances

These requirements form the basis of due diligence and help establish an effective security program for the company. The seven elements that capture these basic functions include the following:

- Appointing a high level manager to oversee compliance using policies, standards, and guidelines.

- Establish the policies, standards, guidelines for the workforce.
- Communicate the policies, standards, guidelines, to all employees.
- Enforce the policies, standards, guidelines with appropriate disciplinary measures.
- Implement procedures for corrections when violations occur.
- Exercise care when granting authority to employees.
- Verify that compliance policies are being carried out.

Policy Types

A security policy is a document that contains senior management's directives to create an information security program to protect the corporation's assets, establish security related goals and security measures, as well as target and assign responsibilities.

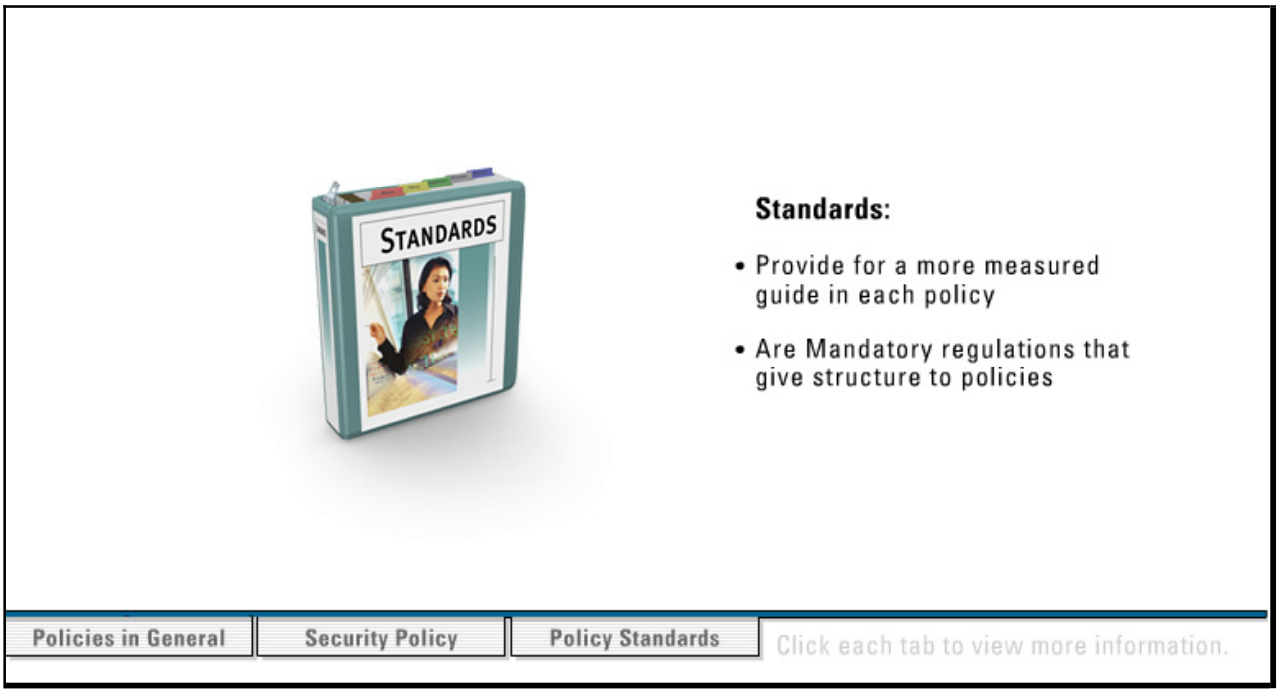

The policy defines the corporation's high-level information security beliefs. In these terms, a policy can then be described as a high-level statement of a company's security beliefs, goals, and objectives, as well as the general means for their realization for a specified subject area.

In general, policies are brief, technical, and solution independent documents. They provide the necessary authority for the establishment and implementation of technology and solution specific standards. Policies remain relevant for a substantial amount of time and are usually updated or revised only when a fundamental change to the organization's operations take place.

Standards, on the other hand, provide for a more measured guide in each of the policy areas. Standards are mandatory regulations designed to provide policies with the support structure and specific actions that are required to be meaningful and effective.

Regulatory

Regulatory policies are those enforced to meet legal compliance.

The data covered under California's Senate Bill 1386 is an individual's name in combination with any one or more of the following items:

- Social security number
- Drivers license number or California ID number
- Account numbers

Regulatory Policies are:

- Enforced to meet legal compliance
- Laws, bills, regulations, statutes, etc.

An example of a regulatory policy is California's Senate Bill 1386. This Bill was enacted as a reaction to a security breach in which hackers broke into a government computer system that maintained information on 265,000 California state employees. This bill added new sections to the California Civil Code of the Information Practices Act. The new regulations require notification to California residents of any breach to the security of a system including UC systems where there is reasonable suspicion that an unauthorized person has acquired unencrypted sensitive information. The data covered under this law is an individual's first name or first initial and last name in combination with any one or more of the following items:

- Social security number
- Drivers license number or California ID number
- Account number, credit card number, in combination with an security code, access code, or password

Advisory

Advisory policies are those policies that define a required behavior with sanctions.

Advisory Policy

- Policies that define a required behavior
- Usually very strong termed suggestions
- Not rigorously enforced
- Tend to apply to more experienced users

For example, to curtail any unauthorized or harmful attacks to its facility, a university might issue an advisory policy stating that it reserves the right to limit, restrict, or extend computing privileges and access to its information resources.

Informative

Informative policies are those which are not enforceable, but can be regulated.

Informative Policy

- Policies that are not enforceable
- Have no implied actions that are expected
- Have no penalty if not followed
- "For information only"

For example, the IT staff at the University might encourage their students to create web pages on their campus servers, but the maximum size of the student's web pages will be determined by the sites webmaster.

Baselines and Guidelines

A baseline provides for the minimal level of security necessary throughout an organization. This baseline needs to be established before the company security architecture can be developed. Standards are usually developed from baselines meaning that baselines can be considered the abstraction of standards.

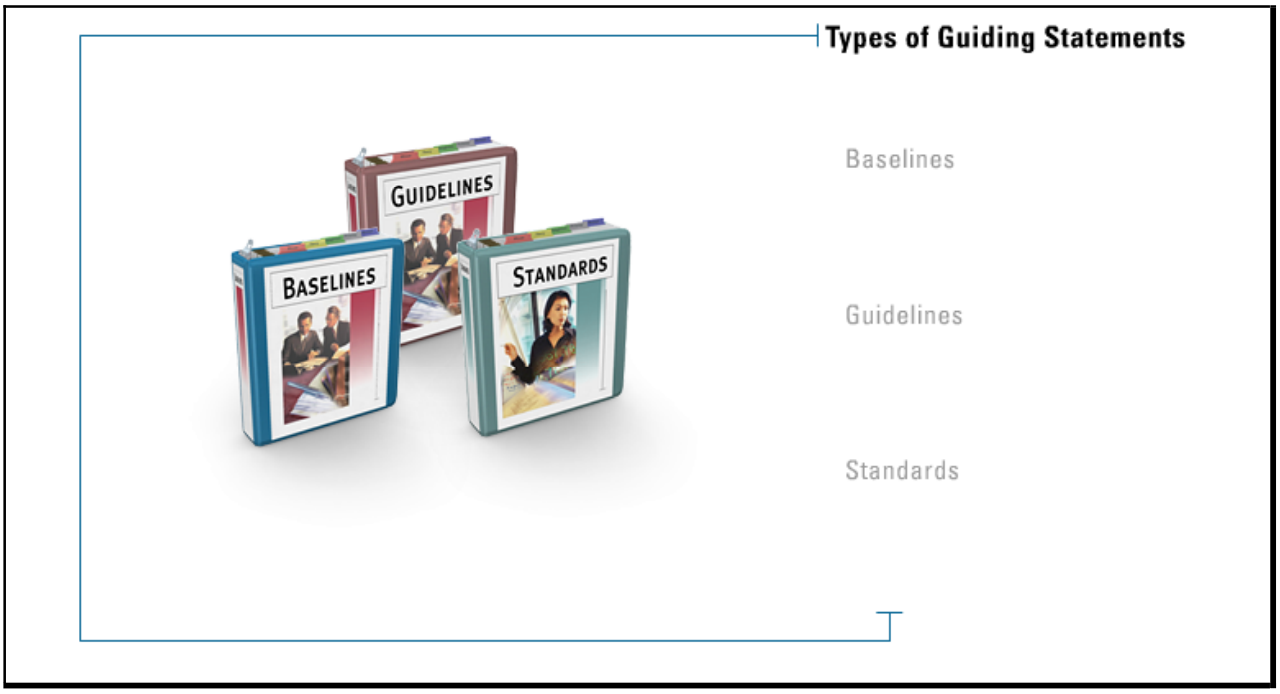

Usually baselines are platform specific implementations that are necessary to provide the minimum level of security to an information system.

Guidelines are the recommended action and operations guides provided to users, the IT staff, operations staff, and others when a specific standard does not apply. In other words, guidelines deal with the methodologies of securing computers and their software. Guidelines address the gray areas that a user confronts.

While guidelines are general approaches that provide the necessary flexibility for any unforeseen circumstances, standards are specific mandatory activities.

Background Checks and Security Clearances

These days no one disputes the value of employee screening. Background checks are the first line of defense to securing any workplace.

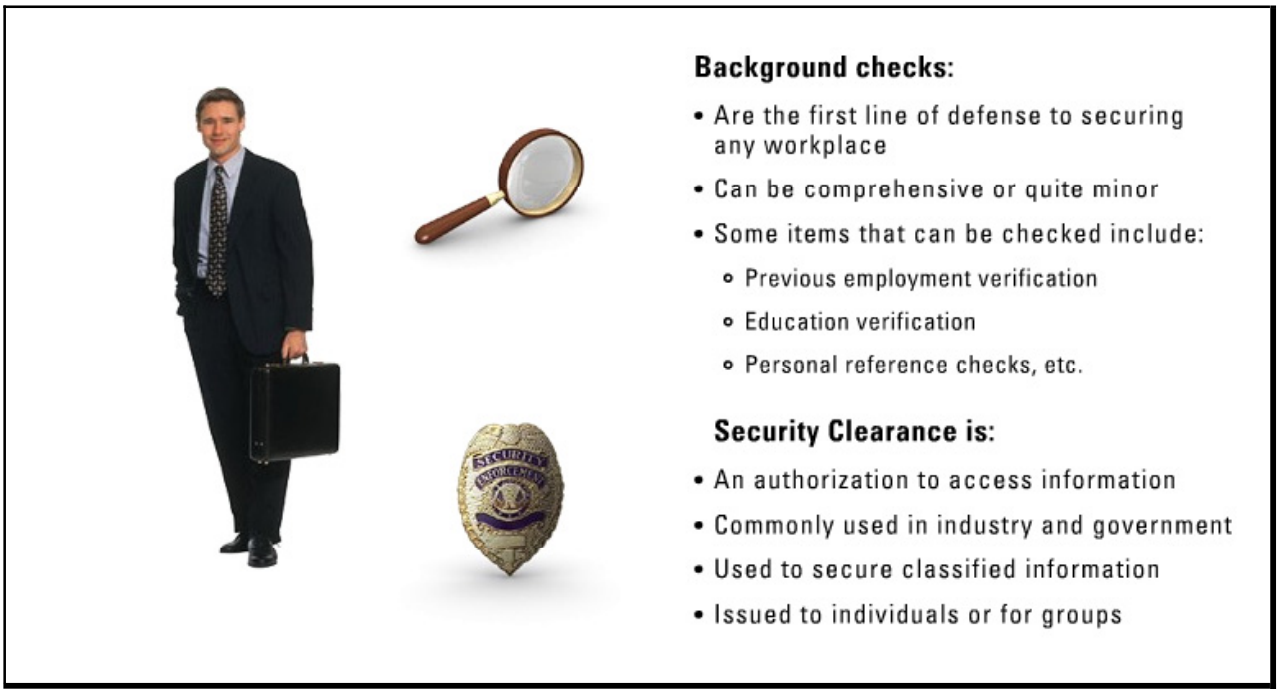

Background checks can be very comprehensive or quite minor. Some items that can be check for background validation include the following items:

- Previous employment verification
- Education verification
- Personal reference checks
- Credit history checks
- Criminal court history
- Civil court history
- Federal court records
- Motor vehicle report
- Worker's compensation report
- Trace reports
- Professional license verification

A security clearance is an authorization that allows access to information that would otherwise be forbidden. Security clearances are commonly used in industry and government, and many jobs in information technology require security clearances. When a security clearance is required to access specific information, the information is said to be classified. Security clearances can be issued for individuals or for groups.

In some organizations, corporations, and agencies, the security policy requires multiple levels of security clearances. The terms used for the classifications, and the requirements imposed on applicants vary among countries and entities. In the United States, the two most common government security clearances are known as Secret and Top Secret. In order to obtain a security clearance, a person must meet certain requirements concerning character, criminal record, and credit history. A higher the level of the security clearance has more stringent requirements.

Employment Agreements

In today's dot.com-to-global-giant world, the employment agreement has become a key negotiating document. Once reserved for only the higher-level executives, the employment agreement is now being used for nearly all new hires.

Employment Agreements:

- Have become key negotiating documents
- Are now being used for all new hires
- Address the following issues:
  - The definition of the business of the employer
  - The scope of employment
  - Limitations on what the employee can or cannot do
  - Identifying what is ***considered confidential*** information and what are ***trade secrets***
  - The right to assign
  - Termination
  - Salary, bonuses, and stock options

Companies using employment agreements are attempting to address the following issues:

- The definition of the business of the employer
- The scope of employment
- Limitations on what the employee can or cannot do
- Identifying what is considered confidential information and what are trade secrets
- The right to assign
- Termination
- Salary, bonuses, and stock options

Awareness Program

Although the vulnerability that people present to a company can never be totally eliminated, a well-planned security awareness program can help to reduce the risk to an acceptable level. To have effective security, employees need to understand their role in protecting information and information assets.

Security Awareness Program:

- Can help reduce risk to an acceptable level
- Helps users understand their role in protecting assets
- Primary objective- educate users on their responsibilities to CIA or company resources
- Helps reinforce the overall security policy

The primary objective of a security awareness program is to educate users on their responsibility to help protect the confidentiality, availability, and integrity of the organization's information and information assets. Information security is everyone's responsibility, not just the IT security department's responsibility. Users need to understand not only on how to protect the organization's information, but why it is important to protect that information. People are often the weakest link in a security chain because they are not trained nor are they generally aware of security's methods and purpose.

A security awareness program should reinforce security policy and other information security practices that are supported by the organization.

Hiring Practices

Having documentation for new employees created and ready is important even in the hiring process. A clear end-user document, acceptable-use policy, and need to know policy for permissions will make the task of adding new users and educating them a smoother process.

In the process of hiring new staff, the human resources department prepares a set of guidelines including the conduct of the staff. HR often presents guidelines on issues of sexual harassment and discrimination that must be understood by the new employee. At this time, HR or IT security should also convey the acceptable-use policies regarding the network and discuss the procedures for acquiring a computer account with an emphasis on the responsibilities of the end-user.

During the hiring process, HR notifies the IT department of new users and their job roles. Following the security procedures, administrators can create these new accounts and assign the appropriate permissions to them.

The employee may require direct information about how to create strong passwords and the type of privileges he or she will be granted. Most of this information should already exist in a policy document for that particular role in the company. Additionally, the IT staff may also help new employees access the system with end-user education.

Terminating Practices

As in the hiring process, the termination process should also be managed. Generally, employees that have planned termination dates should be restricted from all sensitive network areas. The administration team begins to plan for the employee's departure and will execute the procedures already developed.

Terminated Employees

Documentation for new hires:

- Planned termination users should be restricted from sensitive areas
- Administration team should have plan for employees departure (procedures should be created)
- Terminated employees:
  - Should no longer have access to network
  - Should have their accounts disabled (not deleted)
  - After a certain amount of time account can be safely deleted

As part of the termination process make sure the employee no longer has access to the network. In many cases, the administration team will simply delete the employee's computer account; however, this practice has limitations.

Many times the employee's computer account may own files and folders that could be lost when the account is deleted; in other cases, the employee's email will no longer function perhaps resulting in a customer communication problem.

The best practice is to disable the account and not delete it until a certain amount of time has passed that allows the administration team to deal with the changes. A disabled account cannot be logged onto by the old user, but is available to administrators if they need it. No files or folders are lost, and administrators can check or redirect the email. After a planned amount of time, perhaps 30 to 45 days, the account can safely be removed from the system.

Job Descriptions

Human Resources use job descriptions for requisitions and advertisements to help fill jobs within the organization. In the information security context, job descriptions define the roles and responsibilities for each employee.

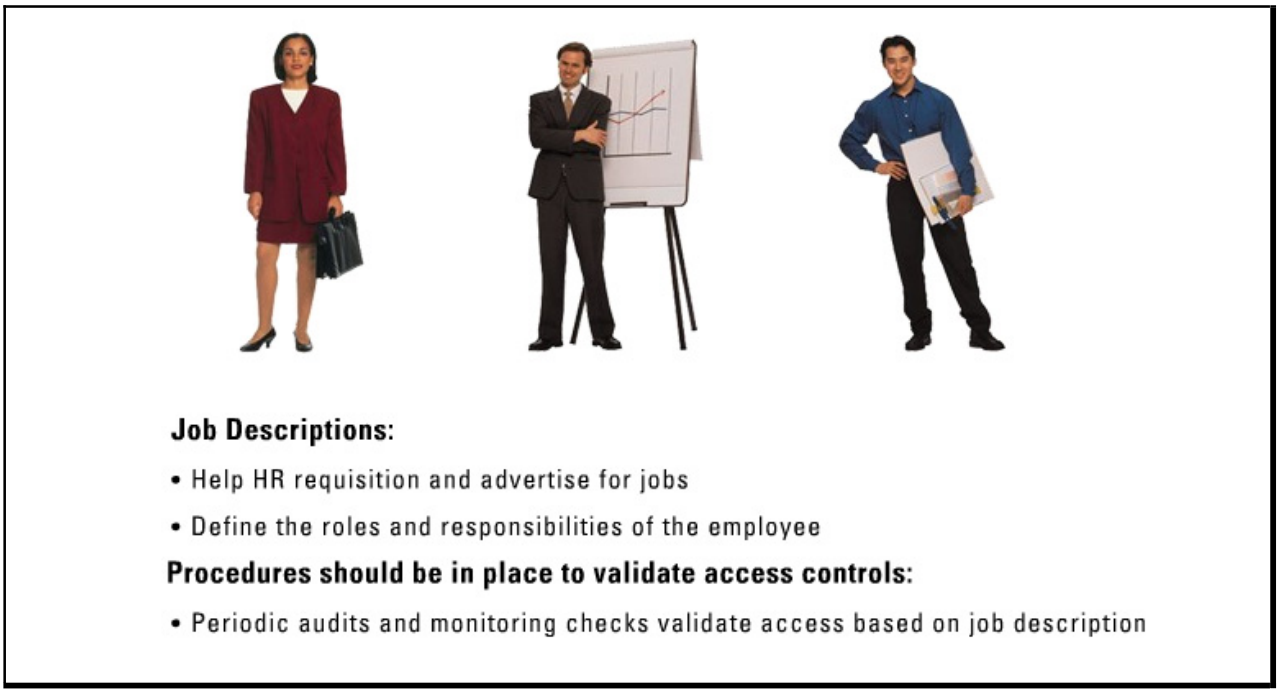

As part of the roles and responsibilities, procedures set the various access controls to ensure that the user can obtain access to only those company resources required to fulfill their job.

During periodic audits and monitoring, finding a user who accesses information beyond his job description may be an indication of a minor or a major security problem. For example, a contractor working on the development of the new database system should not be able to access accounting data. A danger lies in giving access to sensitive information to unauthorized personnel. When the job descriptions are not properly maintained or when a job description is informally changed without changing the official job description, security will find enforcing security policies difficult. A policy to change job descriptions before changing access control lists is imperative.

Role Responsibilities

From senior management to the everyday users, everyone who has access to your organization's systems and networks is responsible for their role in maintaining security as set by the policies. Understanding these roles and responsibilities is key to creating and implementing security policies and procedures.

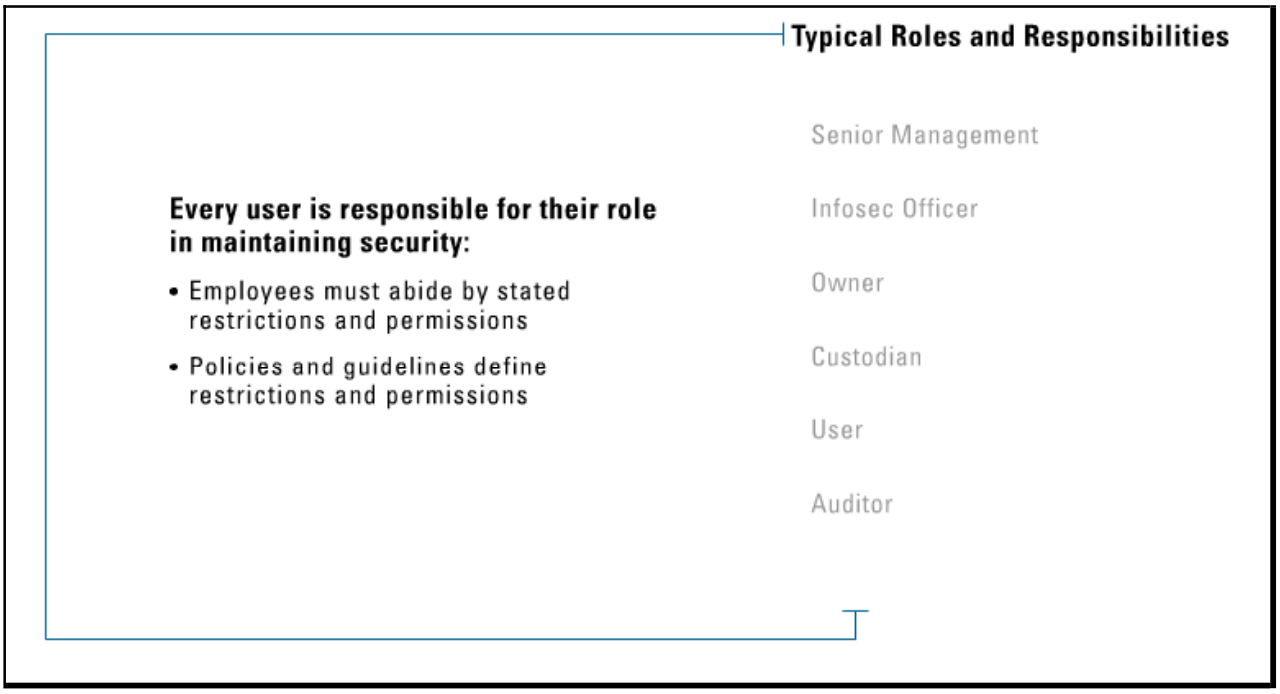

Every employee must abide by certain restrictions and permissions that ensure the security context of the company. These security permissions are set out in policies and guidelines set forth in the company policies and are detailed for each role that must be filled in the company.

Typical Roles and Responsibilities include the following:

- **Senior Management** - Has ultimate responsibility for security
- **Infosec Officer** - Has the functional responsibility for security
- **Owner** - Determines the data classification
- **Custodian** - Preserves C.I.A.
- **User** - Performs in accordance with stated policy
- **Auditor** - Examines security

Separation of Duties and Responsibilities

One of the most common security procedures in larger corporate environments is the concept if separation of duties and need to know policies. The separation of duties restricts security access by job role; it also clearly restricts the possibilities of computer crime.

A single or small group of staff is responsible for specific security duties. This administrative team may consist of enterprise administrators who have security access to the entire network, and sub-group of administrators who only have access to portions of the network under restricted security access.

An example of a sub-group would be administrators who have the ability to create and manage new user accounts but do not have the ability to delete or manage the users email, or set specific security permissions. Only an Enterprise administrator has the role and permissions to delete accounts or to monitor email. This separation of duties ensures that opportunities for mistakes are limited. The limits clearly lay the keys to the company at the foot of the most trusted administrative personnel.

Many management tools are provided for administrators to meet this role-based administrative control. As an example, Windows2000 and Server 2003 from Microsoft use Active Directory to permit the separation of duties for administrators.

Separation of duties also guards against collusion, when more than one person would need to work together to cause some type of destruction or fraud. Limiting duties and permissions reduces the probability of collusion.

Job Rotations

Job rotation limits the amount of time one employee works in a position. Allowing someone to have total control over certain assets can result in the misuse of information, the potential modification of data, and fraud.

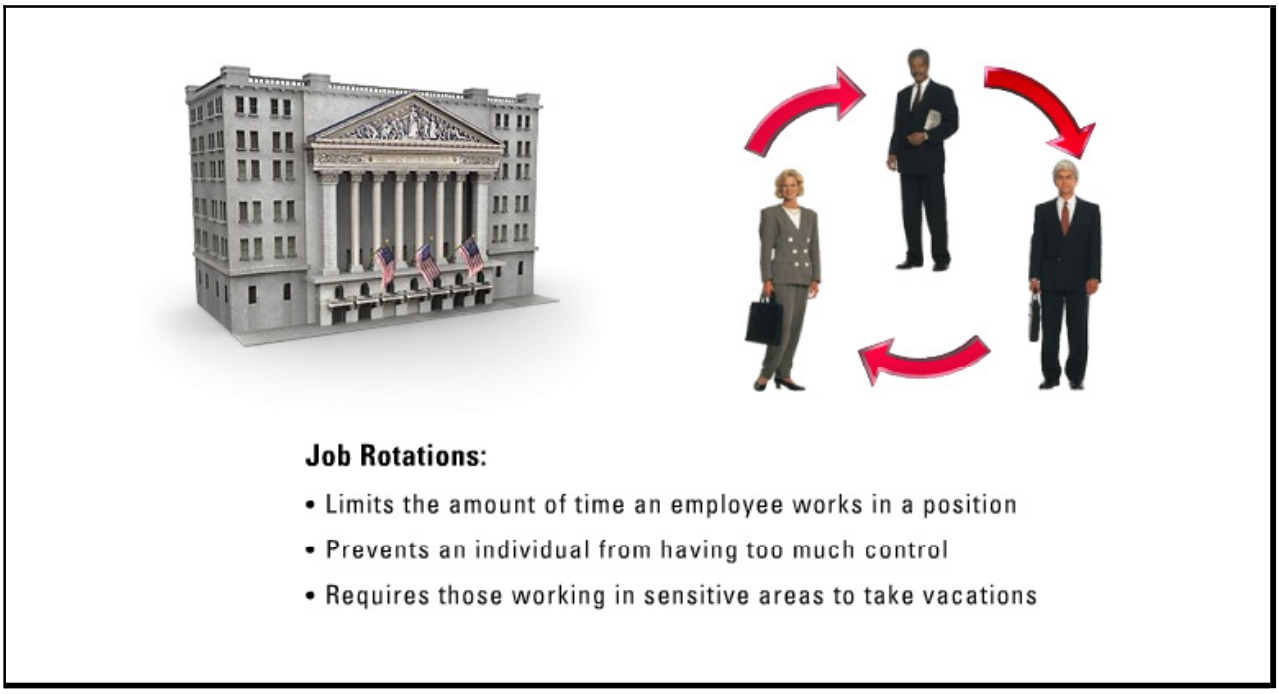

Enforced job rotation prevents a single individual from having too much control because the job rotation limits the time that person may have to build the control that could place information assets at risk.

Another aspect of job rotation is to require those working in sensitive areas to take vacations. By having some of the employees leave the work place, others can step in and provide another measure of oversight. Some companies, such as financial organizations, require their employees to take their vacations during the calendar or fiscal year.

Summary

The key points discussed in this lesson are:

- Purpose of policies, standards, guidelines, and procedures
- Types of objectives
- Policy types
- Regulatory policies
- Advisory policies
- Informative policies
- Baselines and guidelines for policies
- Background checks and security clearances
- Employment agreements
- Security awareness programs
- Security's role in hiring practices
- Security's role in the terminating practices
- Security's role in job descriptions
- Role responsibilities for security
- Separation of duties and responsibilities
- Purpose of job rotations

Risk Management

Overview

To calculate the chance of experiencing an unwelcome outcome, as well as the degree of its severity, the security administrator must calculate the elements of risk and their relationship to one another. Risk management provides weight mitigation factors and provides risk assessment and solutions. This lesson will categorize the processes of identifying, analyzing and assessing, mitigating, or transferring risk in an enterprise.

Importance

The information security profession needs to identify risks in the enterprise as well as calculate the potential losses if loss or theft should occur.

Objectives

Upon completing this lesson, you will be able to:

- Explain the principles of risk management
- List the elements of planning
- Define strategic planning
- Explain the principles of tactical security management
- Explain the principles of operational plans
- Define threat analysis
- Explain the tenets of social engineering
- Explain probability determination
- Explain vulnerability analysis
- Define asset valuation
- List risk assessment tools and techniques
- Explain risk analysis
- Compare qualitative and quantitative risk assessment
- Explain Delphi technique methods

- Define a single occurrence loss
- List formulas for risk calculations
- Name the elements in a benefit analysis
- List the factors in counter measure assessment and evaluation
- Explain the difference between total risk and residual risk
- Explain risk reduction, assignment, and acceptance
- Define information risk management (IRM)

Outline

The lesson contains these topics:

- Principles of Risk Management
- Planning
- Strategic Planning
- Tactical Security Management
- Operational Plans
- Threat Analysis
- Social Engineering
- Probability Determination
- Vulnerability Analysis
- Asset Valuation
- Risk Assessment Tools and Techniques
- Risk Analysis
- Qualitative vs. Quantitative Risk Assessment
- Delphi Technique
- Single Occurrence Loss
- Calculations
- Benefit Analysis
- Counter Measure Assessment and Evaluation
- Total Risk and Residual Risk
- Risk Reduction, Assignment, and Acceptance
- Information Risk Management (IRM)

Principles of Risk Management

Risk management is the process of assessing risk and applying mechanisms to reduce, mitigate, or manage risks to the information assets.

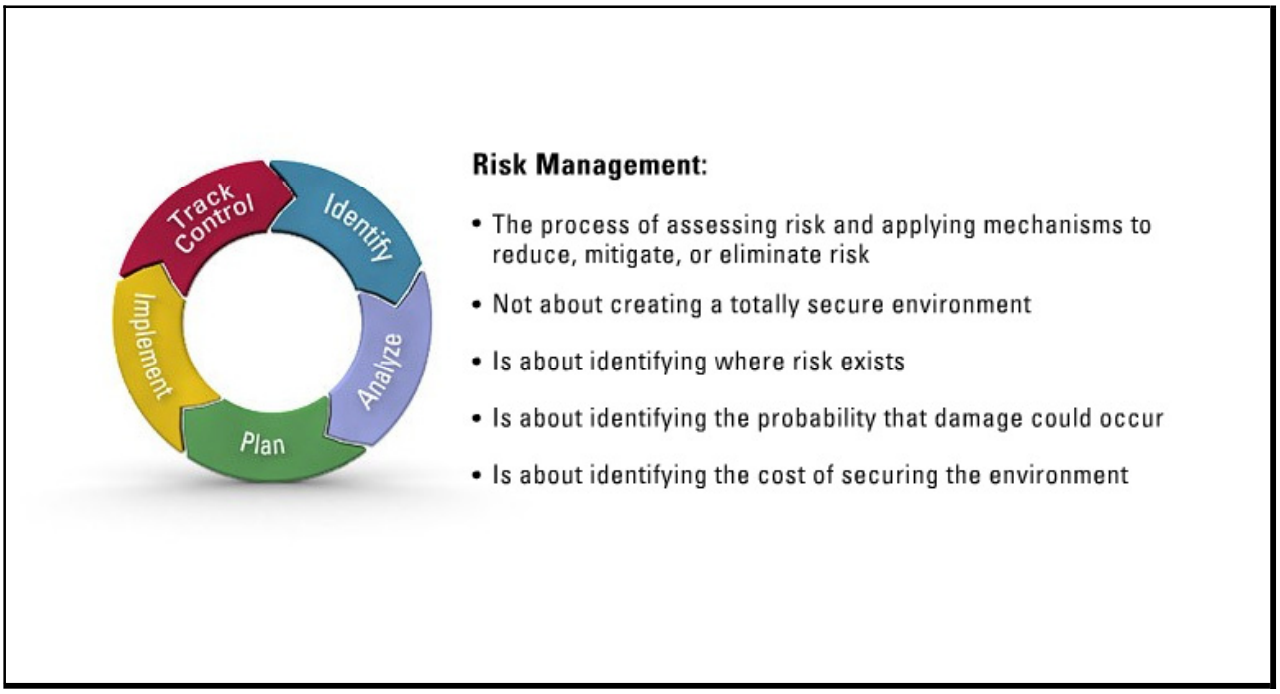

Risk management is not about creating a totally secure environment. Its purpose is to identify where risks exist, the probability that the risks could occur, the damage that could be caused, and the costs of securing the environment. Even if there is a risk to information assets, risk management can determine that it would cost more to secure the asset than if it was damaged or disclosed.

Risk management is not as straightforward as finding the risk and quantifying the cost of loss. Because risks can come from varying sources, an information asset can have several risks. For example, sales data stored on a network disk has the risk of unauthorized access from internal or external users, loss from a software or hardware failure, or inaccessibility because of a network failure.

Planning

Risk management looks at the various possibilities of loss, determines what would cause the greatest loss, and applies controls appropriately. A risk manager might want to reduce all the risk to zero. This reaction is a natural emotional response to trying to solve risk. However, preventing unauthorized access from internal users while trying to ensure accessibility of the data makes risk inevitable. Here, the risk manager must look at the likelihood of the risk and either look for other mitigations or accept it as a potential loss to the organization.

Assessing risk for information security involves considering the types of loss in a risk category, and how that loss might occur as a risk factor. The categories of risk are listed here:

- **Damage** - Results in physical loss of an asset or the inability to access the asset as in the case of a cut in a network cable.
- **Disclosure** - Disclosing critical information regardless of where or how it was disclosed.
- **Losses** - Can be permanent or temporary, including the altering of data or the inability to access data.
- **Risk Factor**
- **Physical damage** - Can result from natural disasters or other factors as in the case of a power loss or vandalism.
- **Malfunctions** - The failure of systems, networks, or peripherals
- **Attacks** - Purposeful acts whether from the inside or outside. Misuse of data, as in unauthorized disclosure, is an attack on that information asset.
- **Human errors** - Usually considered accidental incidents as compared to attacks that are purposeful incidents.

- **Application errors** - Failures of the application, including the operating system. Application errors are usually accidental errors while exploits of buffer overflows or viruses are considered attacks.

Every analyzed information asset has at least one risk category associated with one risk factor. Not every asset has more than one risk category or more than one risk factor. The real work of the risk analysis is to properly identify these issues.

Strategic Planning

Strategic planning produces fundamental security decisions and actions that shape and guide what the company believes risk management is, why it is needed, and how it can be achieved.

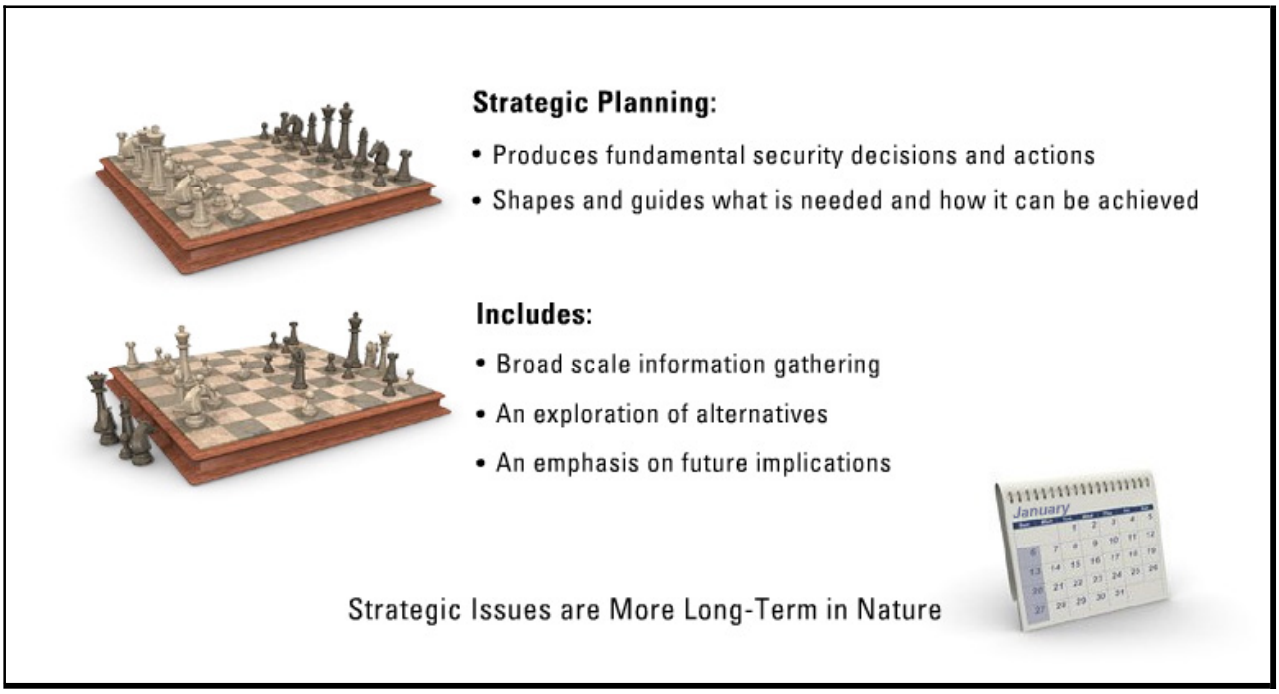

It includes broad-scale information gathering, an exploration of alternatives, and an emphasis on the future implications of present decisions. Top level security managers engage chiefly in strategic planning or long range planning. These managers answer such questions as "What is the threat to this organization?" "What does this organization have to do to secure its assets?" Top level managers clarify the mission of the organization and set its goals.

Strategic planning is the process of developing and analyzing the organization's security mission, overall goals, general security strategies, and implementing security throughout the enterprise. A strategy is a course of action created to achieve a long-term goal. The time length for strategies is arbitrary, but can be two, three, or perhaps as many as five years long, determined by how far in the future the organization is committing its resources. Goals focus on desired changes. They are the ends that the organization strives to attain. Traditionally, strategic planning is done annually.

Tactical Security Management

Tactical security management addresses the daily operations that keep the company a viable and effective enterprise.

Areas in the tactical management arena include moving equipment, physical location of equipment, configuration of firewalls and other perimeter devices, revising SOPs and policy documents, and procedures for ensuring the SOP staff is effectively performing its job.

Top level managers set very general goals that require more than one year to achieve. Examples of long-term goals include long-term growth, improved customer service, and increased profitability. Middle managers interpret these goals and develop departmental tactical plans that can be accomplished within one year or less. In order to develop tactical plans, middle management needs detailed reports on the financial, operational, market, and external environment. Tactical plans have shorter time frames and narrower scopes than strategic plans. Tactical planning provides the specific ideas for implementing the strategic plan. It is the process of making detailed decisions about what to do, who will do it, and how to do it.

Operational Plans

Supervisors implement operational plans that are short-term and deal with the day-to-day work of the team.

Operational Plans:

- Operations plans that are short-term
- Deal with day-to-day work
- Aligned with long-term goals
- Can be achieved in less than one year
- Support tactical plans

Operational Plans include:

- Policies
- Procedures
- Methods
- Rules

Short-term goals are aligned with the long-term goals and can be achieved within one year. Supervisors set standards, form schedules, secure resources, and report progress. They need very detailed reports about operations, personnel, materials, and equipment. The supervisor interprets higher management's plans as they apply to his or her unit. Thus, operational plans support tactical plans. They are the supervisor's tools for executing daily, weekly, and monthly activities.

Operational plans include policies, procedures, methods, and rules. These terms imply different degrees of scope. A policy is a general statement designed to guide employees' actions in recurring situations. It establishes broad limits, provides direction, but permits some initiative and discretion on the part of the supervisor. Thus, policies are guidelines.

A procedure is a sequence of steps or operations describing how to carry out an activity and usually involves a group. More specific than a policy, a procedure establishes a customary way of handling a recurring activity. It gives less discretion to the supervisor in its application. An example of a procedure is the sequence of steps in routing of parts. A method sets up the manner and sequence of accomplishing a recurring, individual task. Almost no discretion is allowed. An example of a method is the steps in cashing a check.

A rule is an established guide for conduct. Rules include definite things to do and not to do. There are no exceptions to the rules. An example of a rule is "No Smoking."

Threat Analysis

Threat analysis identifies threats and defines a cost-effective risk mitigation policy for a specific architecture, functionality, and configuration.

What it Identifies | What it Involves | What it's Required For

Click each tab to view more information.

It involves the mapping of assets, modeling of threats, and building of a mitigation plan that lowers system risk to an acceptable level. The mitigation plan is comprised of countermeasures considered to be effective against the identified threats.

Threat analysis is required for the following instances:

- Complex software systems that integrate multiple infrastructures and technologies.
- Customized application solutions built on standard products.
- All other cases where it is unacceptable to implement pre-compiled to-do lists provided by a software vendor or standards committee.

Threat analysis should be used in the earliest possible stages of system design, and thereafter, as an ongoing process throughout the system's lifecycle of development, integration, change requests, and problem management.

Social Engineering

Social engineering attacks are based on the fact that people trust each other. Using this truism as a basis of attack is one of the most overlooked vulnerabilities in the security world.

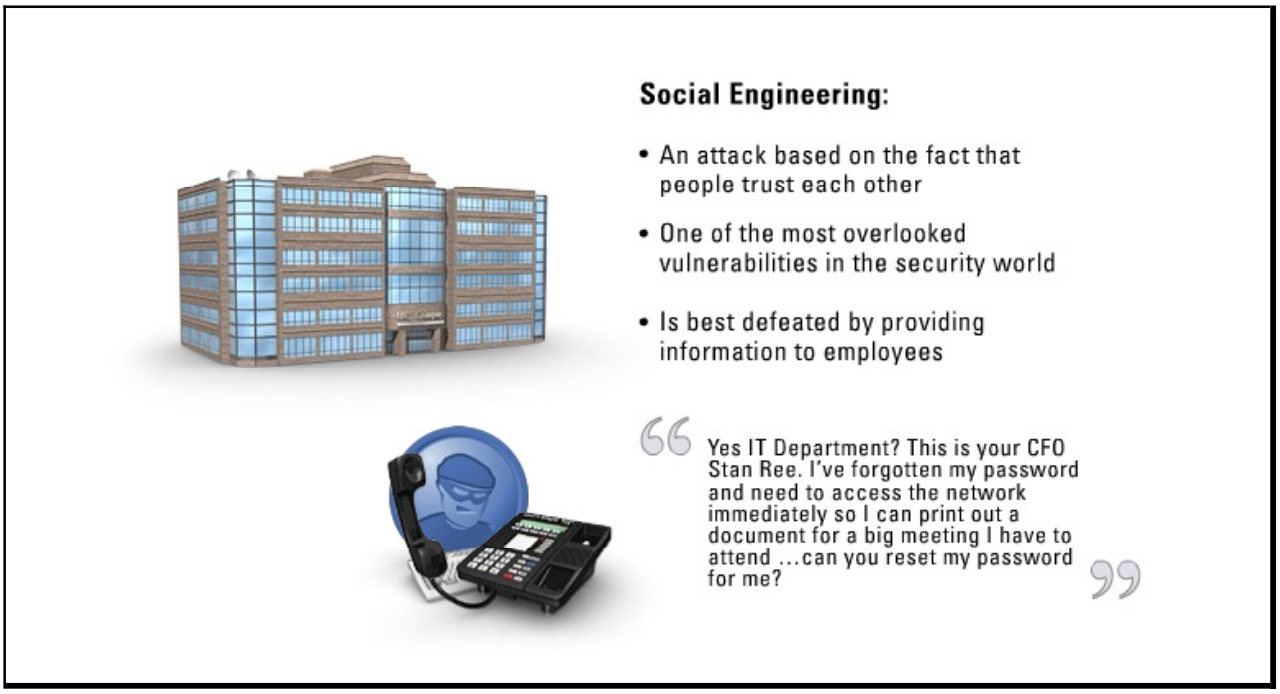

Providing information to employees is the best way to deter this type of attack. All employees in your company should have basic security training, so as not to be vulnerable to this type of attack.

The following story is an example of a social engineering attack:

Cracker Joe does some reconnaissance of the ABC Company and learns that the CFO is named Stan Ree. He learned this information simply by asking HR. He then walks into ABC Company, pretending to have lost his badge, and accesses the facility with the help of a friendly employee. He finds an empty room with a computer already hooked to the company network, and phones the IT Department pretending to be Stan Ree. He tells the IT Department employee that he forgot his password and needs immediate access to the network to obtain a document for his big meeting. Being the helpful citizen, the IT technician resets Stan Ree's password to one of his choice, with a reminder to "…not forget it this time". From there Cracker Joe uses his regular hacking tools to obtain administrator status, creates a few administrator accounts, logs off, and exits the building. Cracker Joe now can obtain remote access into the company using the forged administrator accounts.

Probability Determination

Probability determination is the likelihood a threat will be realized.

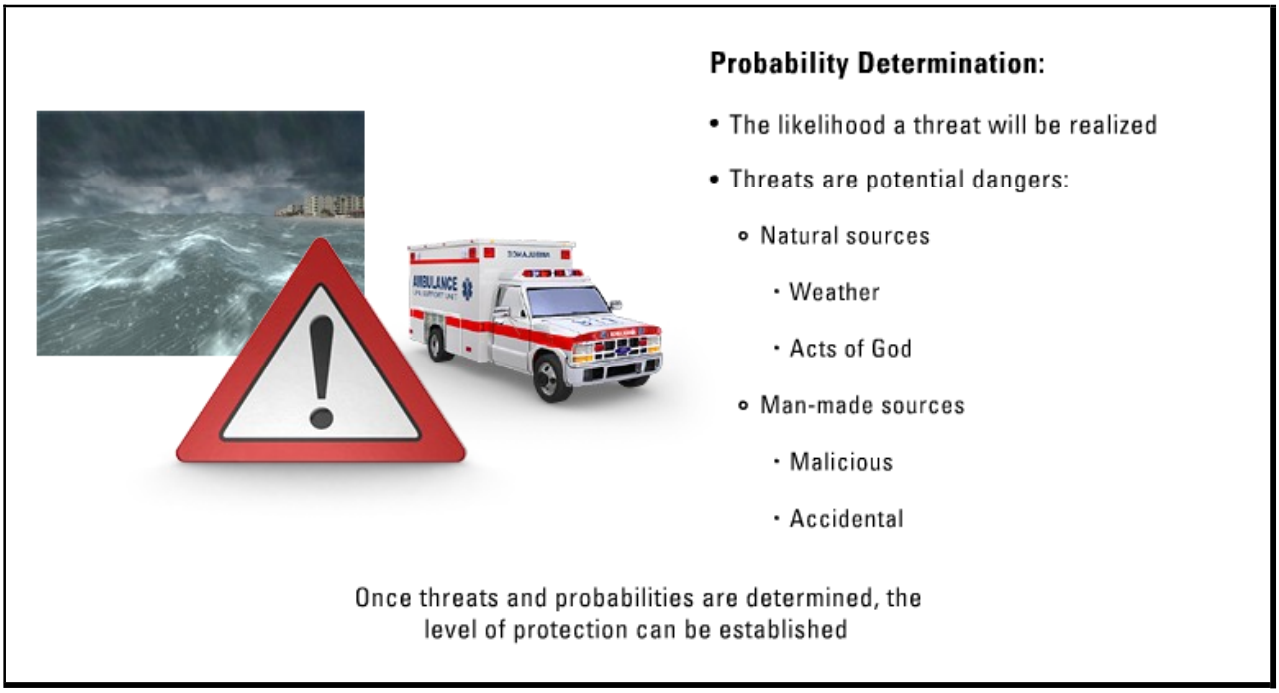

A threat can be identified as any potential danger:

- Natural Source
  - Weather
  - Acts of God
- Man-made source
  - Malicious
  - Accidental

Vulnerability Analysis

Computer security statistics show that over 80 percent of all computer-related fraud is committed by insiders. Insiders often have a motive and a means to strike against a company. Corporations that have Internet connectivity open their systems to a greater number of potential attackers. Most corporations have multiple access points to corporate resources, some of these access points are known, some of them are unknown. Security penetration and vulnerability analysis will uncover any potential exploits.

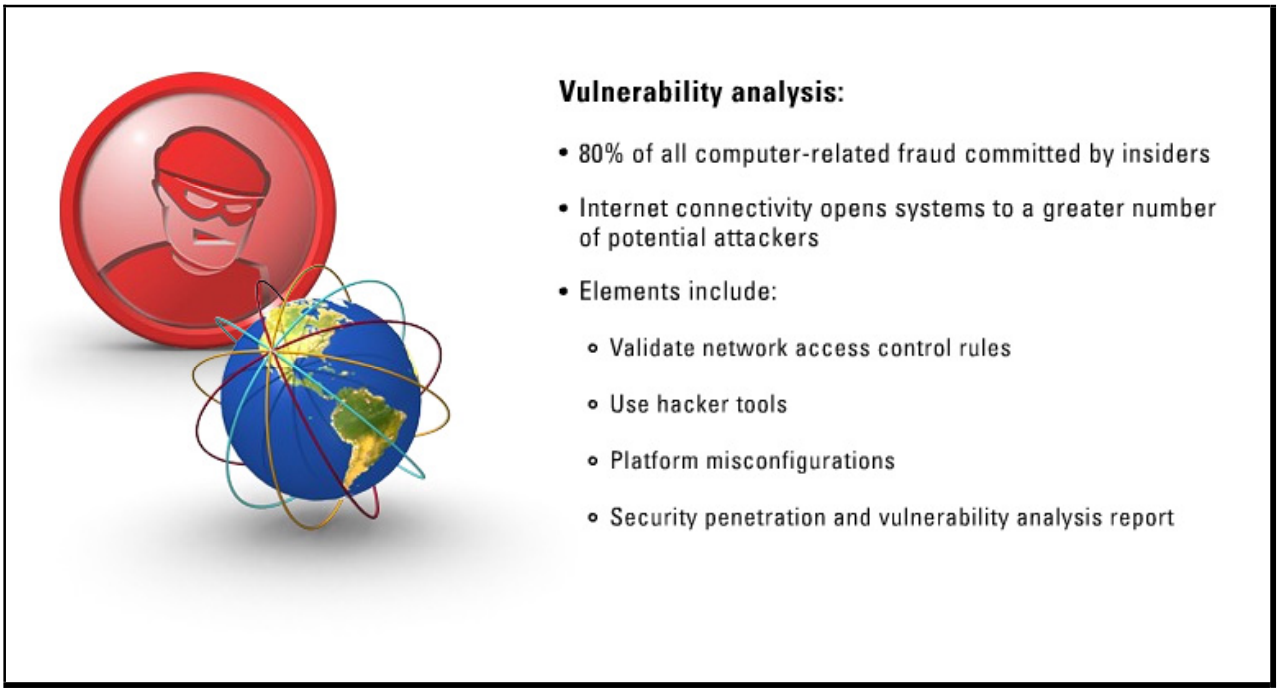

The following is a list of elements in a vulnerability analysis:

- **Validate the Network Access ControlRules** - Since most network access control products are user configurable and prone to human errors, security engineers should validate the rule set against the corporate security policy and the known Internet attacks.
- **Use Hacker Tools** - The most effective way to break into a system is to use the tools that hackers use. Security engineers should use a wide variety of custom written tools to uncover weaknesses in the firewall and router configurations. The most susceptible machines are the publicly accessible hosts such as the World Wide Web, mail, news, and anonymous ftps.
- **Platform misconfigurations** - Optionally, this analysis can test for platform misconfigurations such as NFS, NIS, and .rhosts. It will also verify the strengths of user passwords.
- **Security Penetration and Vulnerability Analysis Report** - The report will document any exploits found in penetration testing and suggest possible security solutions to alleviate the vulnerability.

Asset Valuation

Once the security manager lists the assets, the threats, and vulnerabilities he or she needs to go through an asset valuation process.

Asset valuation:

- Determines the value of assets
- Usually done by creating a matrix
- Value can be determined:
  - In terms of production
  - Research and development
  - Criticality to the tangibles and intangibles of the business model

The security manager needs to determine the value of the assets. The evaluation can be completed by creating a matrix with the value of the assets in terms of high, medium, and low value to the organization, based on definitions. For each asset, the security manager needs to consider the organization's total cost, initial and ongoing, for the full life cycle of the asset. He or she needs to determine the value of the asset is in terms of production, research and development, and criticality to the tangibles and intangibles of the business model. By determining the value of the asset is in the marketplace including intellectual property rights, the security manager can determine the replacement cost the and how to safeguard the asset.

Risk Assessment Tools and Techniques

Certain tools and techniques can be used to perform risk assessment.

They include the following items:

- Risk Analysis
- Identifying Threats and Vulnerabilities
- Asset Valuation
- Qualitative Risk Analysis
- Countermeasure Selection and Evaluation

Risk Analysis

Risks have a loss potential. If a threat agent actually exploits the vulnerability, a company would lose something.

Risk Analysis has three main goals:

- Identify risks
- Quantify the impact of potential threats
- Provide an economic balance between the impact of the risk and the cost of the countermeasure

Risk analysis identifies risk and quantifies the possible damage that can occur to the information assets to determine the most cost-effective way to mitigate the risks. A risk analysis also assesses the possibility that the risk will occur in order to weigh the cost of mitigation. Information security professionals would like to create a secure, risk-free environment. However, it might not be possible to do so without a significant cost. The security manager has to weigh the costs of protection against the potential costs of loss.

Qualitative vs. Quantitative Risk Assessment

Risk evaluation teams use qualitative and quantitative methods to assess risk. The team doing the risk analysis needs to determine which approach is best.

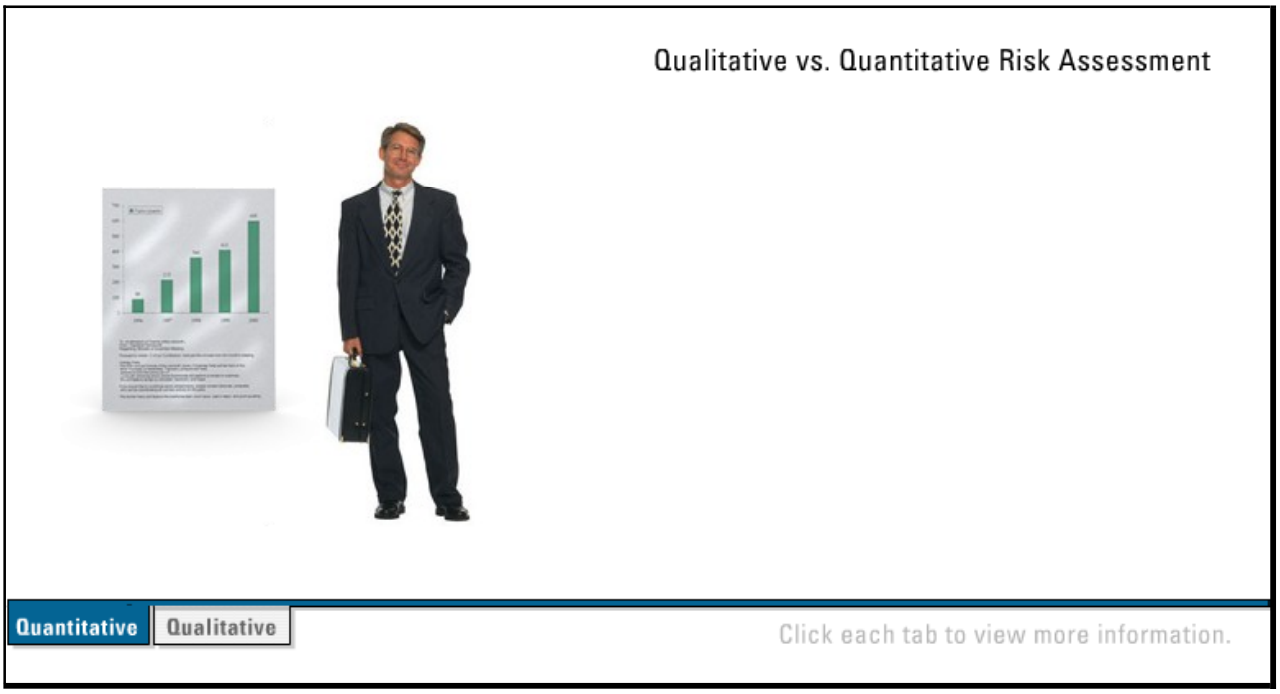

A quantitative approach to risk analysis uses monetary values to assess risk. The quantitative approach cannot quantify every asset and every threat. The values at the extremes, whether high or low, tend to not reflect the reality of the quantitative analysis.

A qualitative risk analysis is a more subjective analysis that ranks threats, countermeasures, and their effectiveness on a scoring system rather than by assigning dollar values. The qualitative approach uses a scoring system to determine risk relative to the environment. The qualitative analysis is good for understanding the severity of the risk analysis relative to the environment and is easier for some people to understand. Teams use various methods of qualitative analysis from group decisions such as the Delphi technique to using surveys and interviews for their ranking system.

Doing a qualitative risk analysis is a bit different from a quantitative analysis. In a quantitative analysis, the analyst does not have to be an expert in the business of the organization or have an extensive knowledge of the systems. Using his basic knowledge, he can analyze the basic business processes and use formulas to assess value to the asset and threats. Qualitative analysts are experts in the systems and the risks being investigated. They use their expertise and the users of the system, to give the threats appropriate ranks.

To do a qualitative risk analysis, the analyst identifies and analyzes the major threats and the scenarios for the possible sources of the threat. The scores generated in this analysis show the likelihood of the threat occurring, the potential for the severity, and the degree of loss. Additionally, the potential countermeasures are analyzed by ranking them for their effectiveness.

When the qualitative analysis is completed, the scores for the threat are compared to the countermeasures. If the scores for the countermeasure are greater than the threat, the countermeasure will be more effective

in protecting the asset. However, since the analysis is a subjective analysis, the meanings of the rankings are also open to interpretation.

Delphi Technique

The Delphi Technique was originally conceived as a way to obtain the opinion of experts without necessarily bringing them together face to face. The security world uses the Delphi technique as a group decision method. This technique ensures that each member of a group gives an honest opinion of what he or she thinks the result to a particular risk will be.

The Delphi Technique:

- Conceived as a way to obtain expert opinions without experts being face-to-face
- Security world uses as a group decision and problem-solving method
- Ensures that each member gives an honest opinion on any particular risk threat

The Delphi technique is another way of obtaining group input for ideas and problem-solving. Unlike the nominal group process, the Delphi does not require face-to-face participation. It uses a series of carefully designed questionnaires interspersed with information summaries and feedback from preceding responses.

In a planning situation, the Delphi can be used to get the following information:

- Develop a number of alternatives.
- Assess the social and economic impacts of rapid company growth.
- Explore underlying assumptions or background information leading to different judgments.
- Seek out information on which agreement may later be generated.
- Correlate informed judgments on a subject involving many disciplines.
- Educate respondents on the diverse and interrelated elements of a topic.

Single Occurrence Loss

This topic discusses a loss called a single occurrence loss.

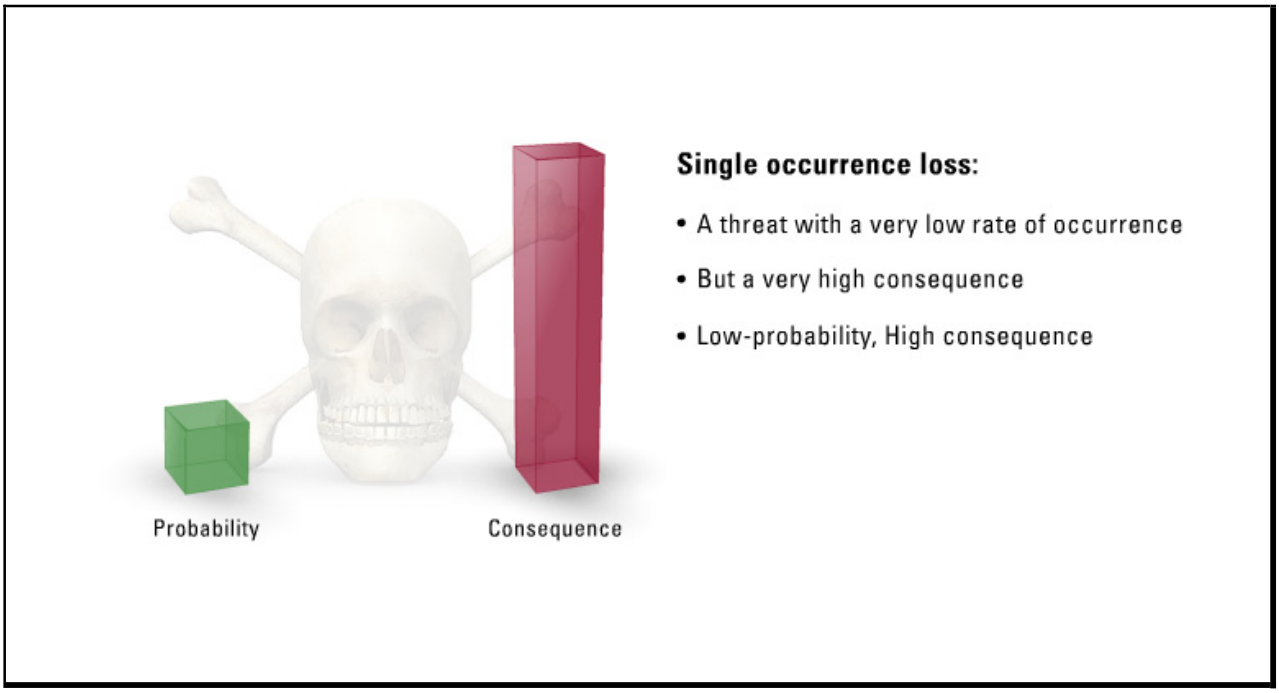

When a threat has a very low rate of occurrence that is difficult to estimate, but the threat would cause a very high loss were it to occur, the result would be referred to as a low-probability, high-consequence risk. This type of loss is often called a single occurrence loss.

Calculations

This topic discusses commonly used risk calculation formulas.

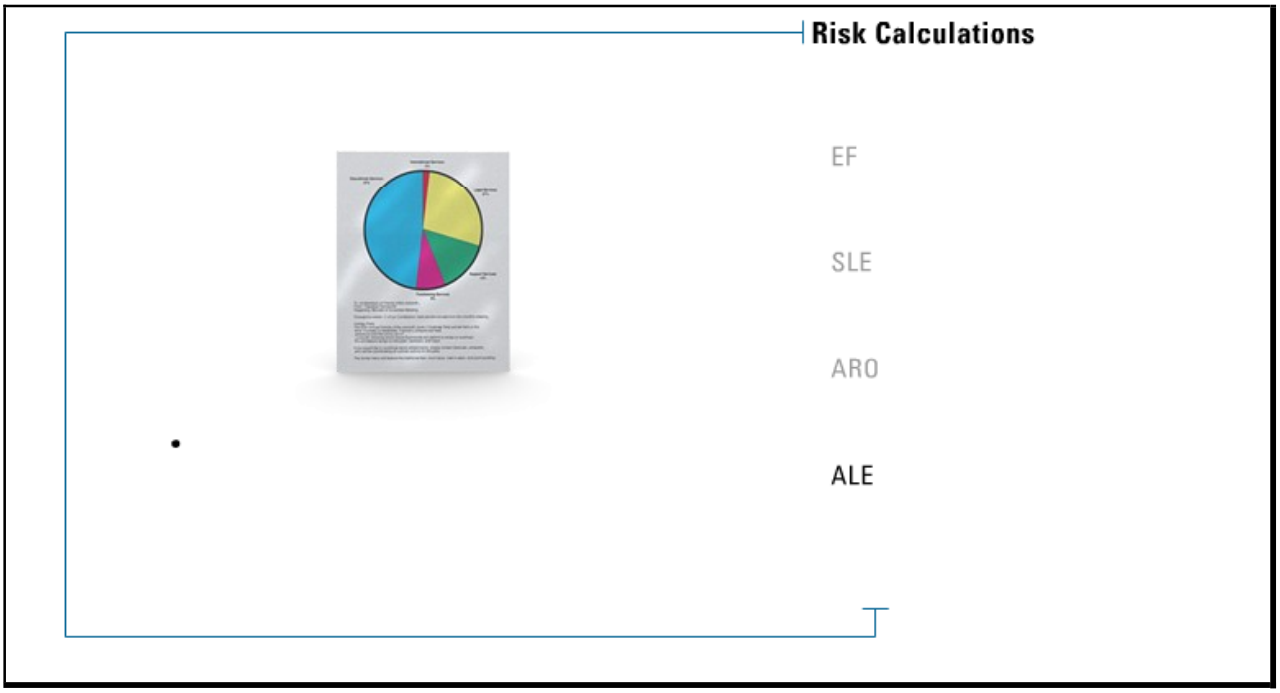

Commonly used risk calculation formulas include the following:

- **EF (Exposure Factor)** = Percentage of asset loss caused by identified threat.
- **SLE (Single Loss Expectancy)** = Asset value * Exposure Factor
- **ARO (Annualized Rate of Occurrence)** = Estimated frequency a threat will occur within a year.
- **ALE (Annualized Loss Expectancy)** = Single Loss Expectancy * Annualized Rate of Occurrence

Cost-Benefit Analysis

The information from the risk analysis allows the risk manager to perform a cost-benefit analysis (CBA), to compare the cost of the safeguards to the costs of not adding the safeguards.

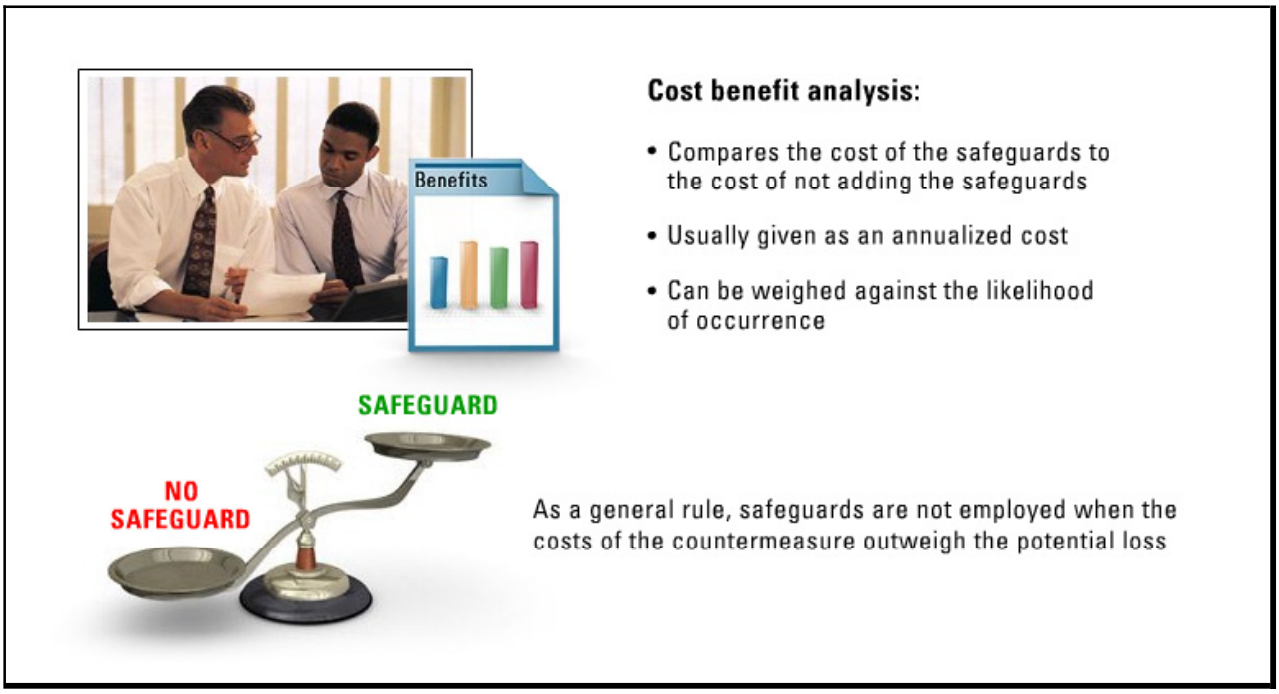

Costs are usually given as an annualized cost and can be weighed against the likelihood of occurrence. As a general rule, safeguards are not employed when the costs of the countermeasure outweighs the potential loss.

The following table has a possible analysis of an information asset worth $10,000 should it be lost.

Cost of Countermeasure	Gain/(Loss)	Analysis
$0	($10,000)	By doing nothing, if the asset is lost, there could be a complete loss that costs $10,000.
$5,000	$5,000	If the countermeasure costs $5,000, you will gain $5,000 in providing the protection by mitigating the loss.
$10,000	$0	The cost of the countermeasure equals the cost of the asset. Here, you might weigh the potential for the countermeasure to be needed before making a decision.
$15,000	($5,000)	With the countermeasure costing more than the asset, the benefit does not make sense in this case in terms of financial cost.

Countermeasure Selection and Evaluation

Organizations employ countermeasures, or safeguards, to protect information assets. In selecting the proper countermeasures, the security manager needs to find a countermeasure that is also the most cost-effective. Determining the most cost-effective countermeasure is called a cost/benefit analysis.

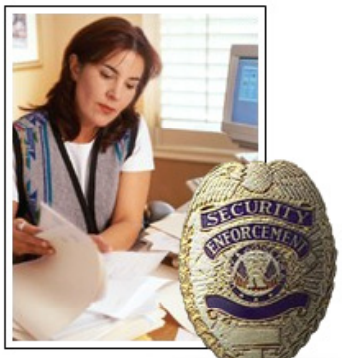

When selecting:

- Security manager needs to select the most cost-effective countermeasure
- Uses a cost/benefit analysis
  - Looks at the annualized loss expectancy
  - Compares ALE to the annual cost of safeguard
  - And ALE after the countermeasure is installed

Value of Countermeasure =
ALE (without countermeasure) – Cost (safeguard) – ALE (with countermeasure)

A cost/benefit analysis looks at the annualized loss expectancy ALE, compares that to the annual cost of the safeguard, and the ALE after the countermeasure is installed to determine whether the costs show a benefit for the organization. The calculation can be written as follows:

Value of Countermeasure = ALE (without countermeasure) – Cost (safeguard) – ALE (with countermeasure)

Choosing a countermeasure for the amount of cost is a purely business way of analyzing risk. However, the security professional understands that regardless of the cost, the countermeasure is not worth using unless it protects the asset. Information security professionals should work with business people to select the most effective countermeasure that will function to properly protect the asset.

Total Risk and Residual Risk

This topic differentiates total risk and residual risk.

Total risk is incurred when no safeguard is in place to protect the asset. Residual risk is the value of the risk after implementing the countermeasure.

- **Total Risk** = Threats * Vulnerability * Asset value
- **Residual Risk** = Total Risk * Controls Gap

Risk Reduction, Assignment, and Acceptance

Risk assessment tells the organization what the risks are; the organization determines how to manage the risks. Risk management is the trade-off that an organization makes regarding that risk. Not every risk could be mitigated. It is the job of management to decide how that risk is handled.

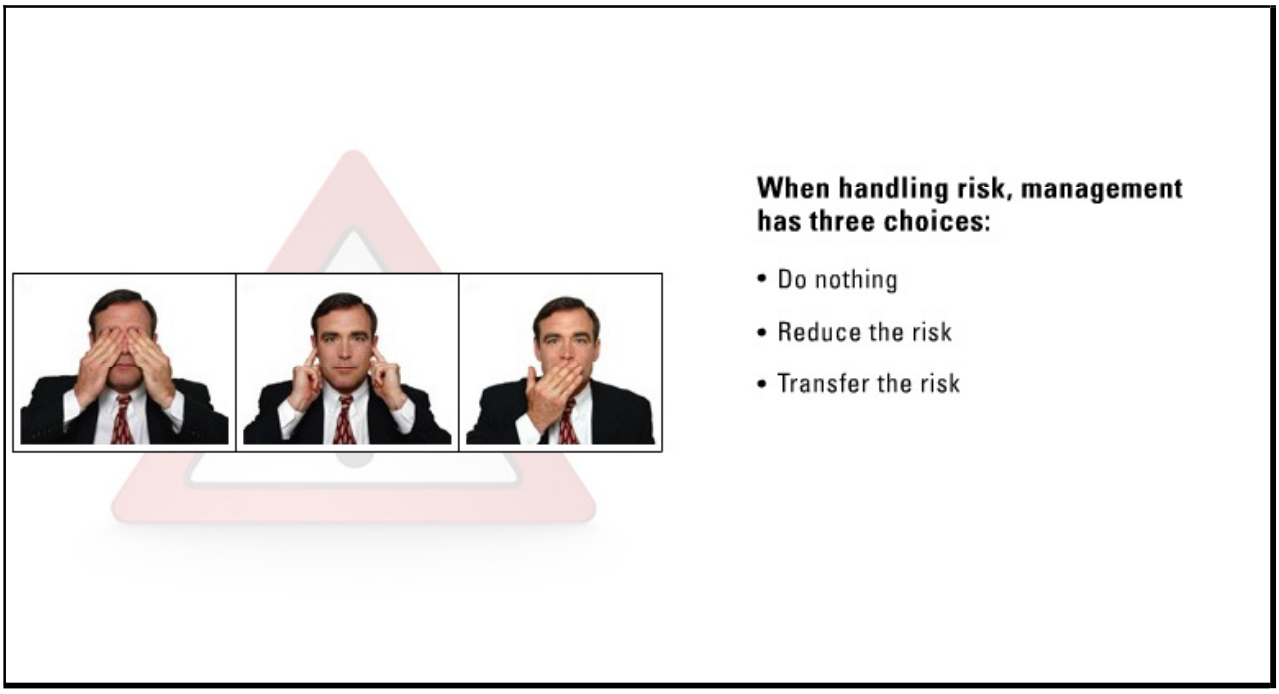

In basic terms, the choices are one of three options:

- **Do nothing** - To do nothing means that management must accept the risk and the potential loss if the threat occurs.
- **Reduce the risk** - To reduce the risk, management implements a countermeasure and accepts the residual risk.
- **Transfer the risk** - To transfer the risk, management purchases insurance against the damage.

These decisions can be made only after identifying the assets, analyzing the risk, and determining countermeasures. Management uses these steps to make the proper decisions based on the risks found during this process.

Information Risk Management (IRM)

Information Resource Management (IRM) encapsulates the concept that information is a major corporate resource and must be managed using the same basic principles used to manage other assets.

Information Risk Management:

- Information is a major resource and must be managed using the same basic principles used to manage other assets
- Effective management of data improves availability, accessibility, and utilization of data
- Key functions include:
  - Data administration
  - Records management

This approach means the effective management and control of data and information as a shared resource to improve the availability, accessibility, and utilization of the data and information. Data administration and records management are key functions of information resource management.

Summary

The key points discussed in this lesson are:

- Principles of risk management
- Elements of planning
- Strategic planning
- Principles of tactical security management
- Principles of operational plans
- Threat analysis
- Tenets of social engineering
- Probability determination
- Vulnerability analysis
- Asset valuation
- Risk assessment tools and techniques
- Risk analysis
- Qualitative and quantitative risk assessment
- Delphi technique methods
- Single occurrence loss
- Formulas for risk calculations
- Elements of benefit analysis
- Counter measure assessment and evaluation
- Difference between total risk and residual risk
- Risk reduction, assignment, and acceptance
- Information risk management (IRM)

Roles and Responsibilities

Overview

Corporate resources are owned by the end user or by the IT staff. A clear understanding of exact resource ownership, defined by the roles and responsibilities of all entities in the enterprise, is an important element of security.

Importance

The information security professional needs to understand the ownership responsibilities of corporate resources in order to provide an effective security solution to protect the company resources.

Objectives

Upon completing this lesson, you will be able to:

- Name individuals responsible for security
- Define the role of information stewards
- Define the role of the custodians
- Explain security awareness training
- Explain security management planning

Outline

The lesson contains these topics:

- Individuals
- Stewards
- Custodians
- Security Awareness Training
- Security Management Planning

Individuals

Each individual in a company is responsible for taking certain security related precautions to protect company assets.

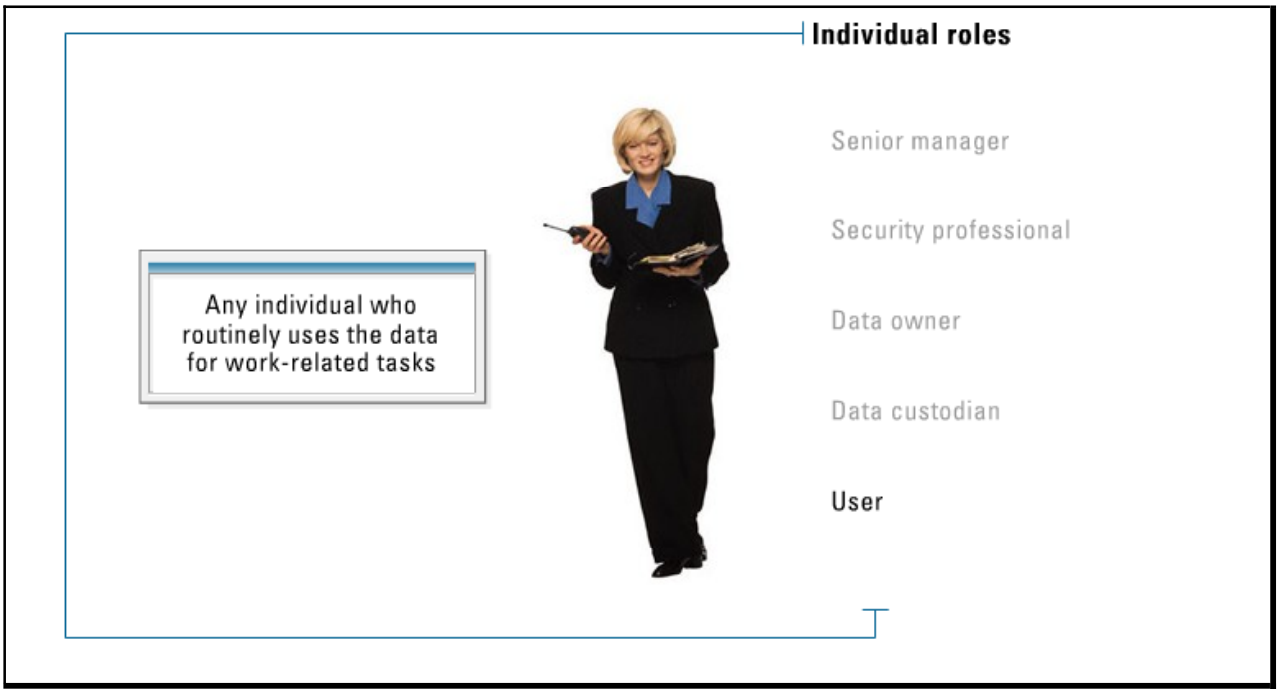

From senior management to end users, layers of responsibility include the following roles:

- **Senior Manager** - Ultimately responsible for security of the organization and the protection of its assets.
- **Security Professional** - Functionally responsible for security and carries out sensitive directives proclaimed by his direct manager.
- **Data Owner** - Usually a member of senior management and ultimately responsible for the protection and use of the data. Decides upon the classification of the data. Is responsible for and alters these classifications if the business needs arise. Will delegate the responsibility of the day-to-day maintenance of the data, which is the responsibility of the data custodian.
- **Data Custodian** - Is given the responsibility of the maintenance and protection of the data.
- **User** - Any individual who routinely uses the data for work-related tasks. Must have the necessary level of access to the data to perform the duties within his position and is responsible for following operational security procedures to ensure the data's C.I.A. to others.

Stewards

Information stewards are senior business unit managers with the authority for acquiring, creating, and maintaining information systems within their assigned area of control.

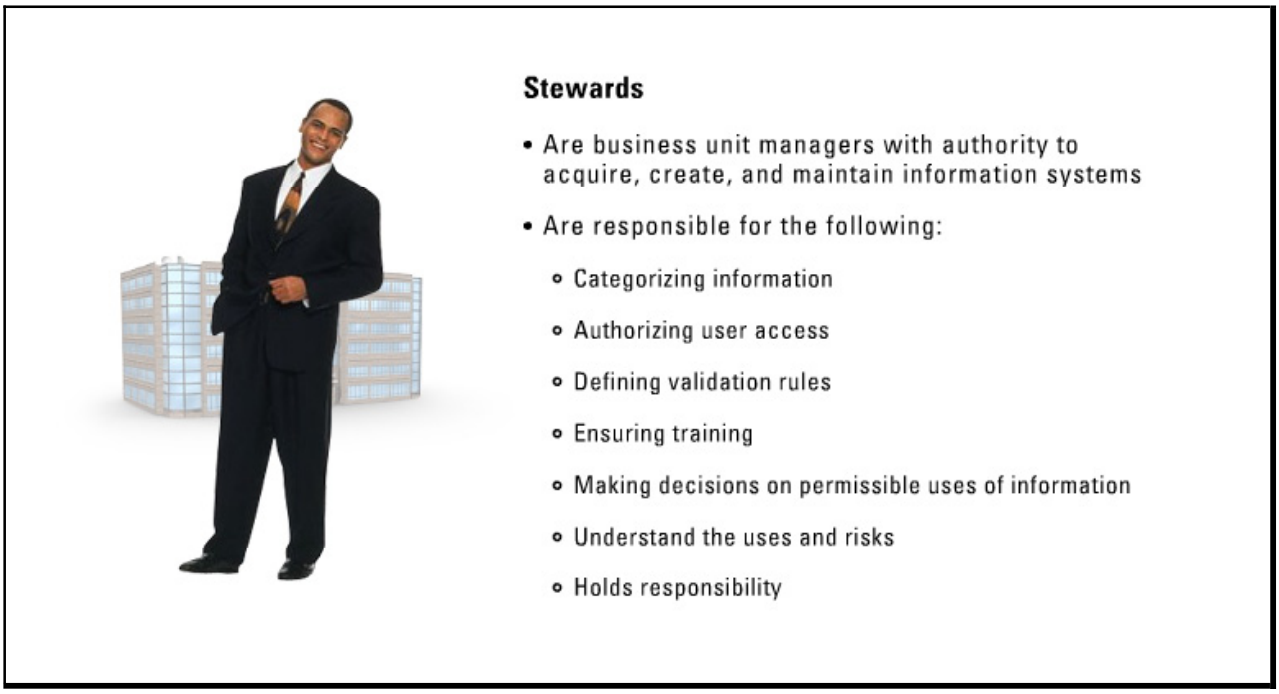

Stewards are responsible for the following duties:

- Categorizing the information for which they have been designated a Steward using categories defined in the Data Classification Policy.
- Authorizing user access to information based on the need-to-know.
- Defining the validation rules used to verify the correctness and acceptability of input data.
- Insuring that users who enter the system or modify data have a sufficient level of training.
- Assisting with contingency planning efforts and categorizing information or specific application systems according to a criticality scale.
- Making decisions about the permissible uses of information.
- Understanding the uses and risks associated with the information for which they are accountable.
- The information steward's role holds responsibility for the consequences associated with improper disclosure, insufficient maintenance, inaccurate classification labeling, and other security related control deficiencies pertaining to the information for which they are the designated a steward.

Custodians

Information Custodians are individuals, often staff within the Information Systems Department or departmental systems administrators, in physical or logical possession of information from Stewards.

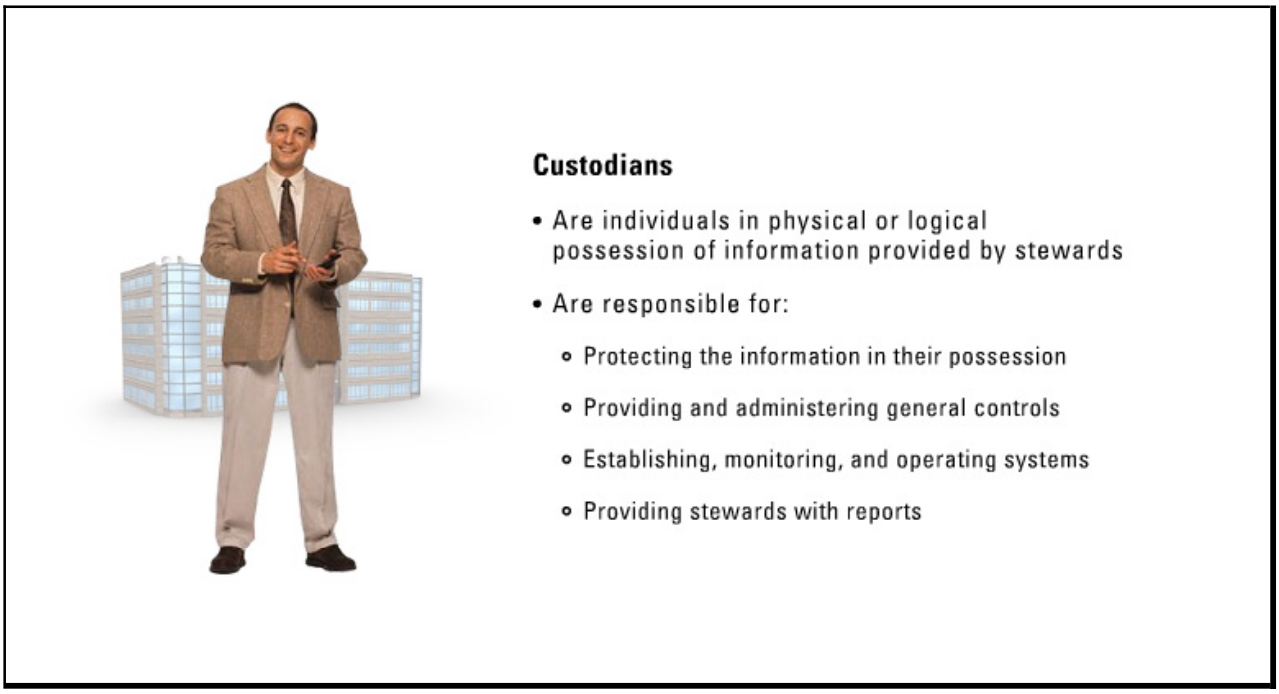

Custodians are responsible for the following tasks:

- Protecting the information in their possession from unauthorized access, alteration, destruction, or usage.
- Providing and administering general controls such as back-up and recovery systems consistent with company policies and standards.
- Establishing, monitoring, and operating information systems in a manner consistent with policies and standards
- Providing Stewards with reports about the resources consumed on their behalf through a charge-back system and reports indicating user activities.
- Changing the production information in their possession only after receiving explicit and temporary permission from either the steward or an authorized user.

Security Awareness Training

The importance of security awareness training and education cannot be overstated. Security professionals need to train all the stakeholders on the policy standards and procedures and reinforce role responsibilities for maintaining the security environment.

Security Awareness Training

- Cannot be overstated
- All stakeholders must embrace the security policy
- Not an easy task
- Requires clear communication
- Includes:
  - Security-related training
  - Awareness training for specific departments
  - Technical training for IT
  - Advanced Infosec training for security personnel
  - Security training for senior managers

All stakeholders need to embrace the policy as an integral part of their jobs. This task is not easy. Over the last decade, the commitment to security by industry-leading companies has been viewed as lacking. The results are products that have insufficient security measures installed into environments that further weaken the information security program. The dichotomy can be vexing.

Security awareness training requires clear communication. Having a technically competent communicator for the security department could be helpful. This person would do the training, educate the department to the concerns of its users, and act as a liaison between users and the department. Having someone who can communicate helps raise the confidence level users should have for the department. Types of security awareness training include the following items:

- Security-related job training for operators
- Awareness training for specific departments or personnel groups with security sensitive positions
- Technical security training for IT support personnel and system administrators
- Advanced InfoSec training for security practitioners and information system auditors
- Security training for senior managers, functional managers and business unit managers

Security Management Planning

Planning for information security includes preparation to create information security policies that will be the guidance for the entire information security program.

Security Management Planning

- Includes preparation to create security policies
- The guidance for the entire information security program
- Includes education

To create the policy, security managers should plan to perform a risk analysis on the information assets to be protected. The risk analysis will identify the assets, determine risks to them, and assign a value to their potential loss. Using this approach, management can make decisions on the policies that best protect those assets by minimizing or mitigating the risks.

Over the years, network technology has changed how information assets are protected, because information security programs are becoming more a very what why because why dynamic. Tactics to maintain security change with the release of every new operating system and with every new communications enhancement. In the past, data was stored and accessed through mainframes where all the controls were centralized. Networked systems change this paradigm by distributing data across the network.

Additionally, network protocols invented to share information and not secure information have compounded security issues. In the beginning, security was left up to each system's manager in a small society of network users. As technology grew, the information assets became less centralized and security managers confronted the problem of maintaining the integrity of the network while the information was used on the networked systems. Although centralized management of servers and information security is growing, information security managers needs to take into account everywhere the information assets touch.

The final aspect of information security management is education. Management is responsible for supporting the policy not only with its backing, but also by including policies and the backing for educating users on those policies. Through security awareness training, users should know and understand their roles under the policies.

Summary

The key points discussed in this lesson are:

- Individuals responsible for security
- Role of information stewards
- Role of the custodians
- Security awareness training
- Security management planning

Application Development Security

Overview

Applications are an asset to an enterprise, as they allow very precise solutions to be implemented exactly to the enterprise's specifications. If the application is poorly written, an attacker may find errors in the code or sub-routines that he or she can exploit in a way the programmer did not intend. This module will discuss how to increase the security of application development.

Objectives

Upon completing this module, you will be able to:

- Describe the various types of programs
- List the types of malicious code
- List security attack methods
- List database and data warehousing models
- Explain knowledge-based systems
- Explain the systems development life cycle
- Name threats to data security

Outline

The module contains these lessons:

- Application Development Introduction
- Malicious Code
- Methods of Attack
- Databases and Data Warehousing
- Knowledge-Based Systems
- Systems Development Life Cycle
- Security and Protection

Application Development Introduction

Overview

Application issues are an essential component of security. This lesson discusses open source and closed source code and the various types of programs.

Importance

Understanding the ways that applications are created and can be exploited is essential to the security information professional. Application exploits can cause serious breaches in security and loss of revenue.

Objectives

Upon completing this lesson, you will be able to:

- Identify application security issues
- Differentiate open source and closed source code
- Explain what an assembler does
- Explain what a compiler does
- Explain what an interpreter does

Outline

The lesson contains these topics:

- Application Issues
- Open Source vs. Closed Source Code
- Assemblers
- Compilers
- Interpreters

Application Issues

Application security issues can arise from Java, ActiveX controls, database security, data warehousing and data mining security, malicious code and other attack methods, and various application development methods.

You must study and establish countermeasures for any portion of an application that can come under attack by a threat agent.

Open Source vs. Closed Source Code

This topic differentiates open source code and closed source code.

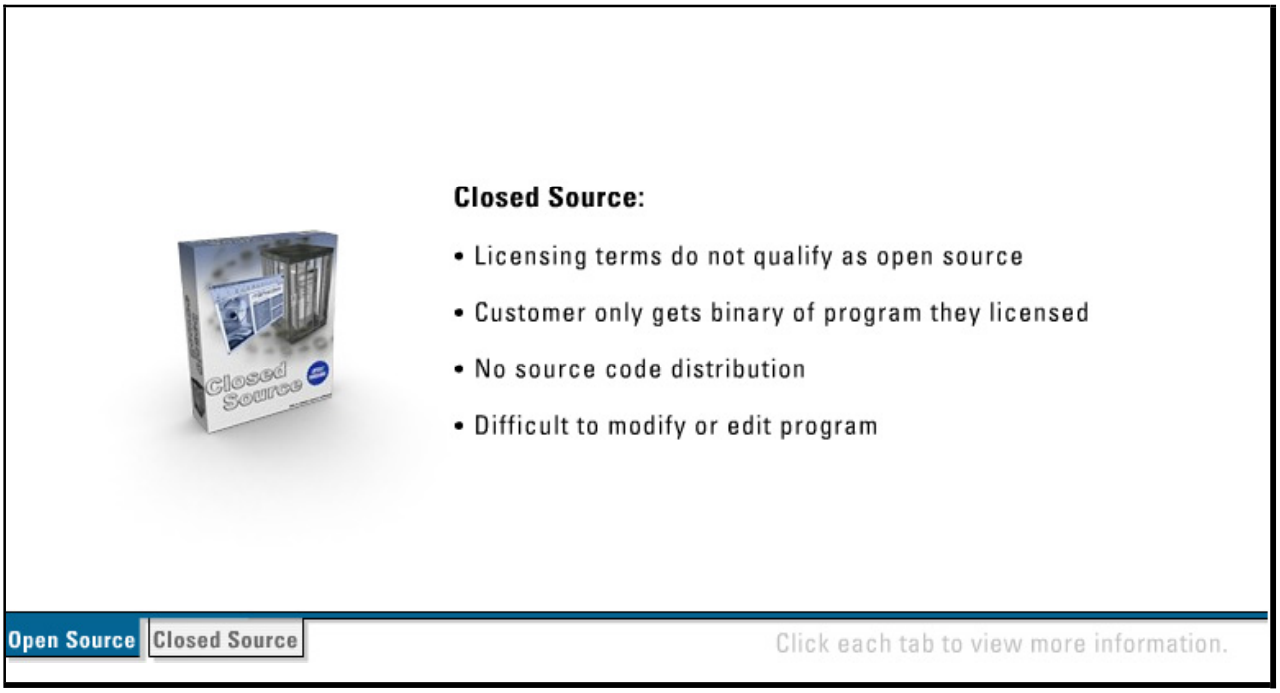

Open source code is a term for a program whose license gives users the freedom to run the program for any purpose, to study and modify the program, and to redistribute copies of either the original or modified program without having to pay royalties to previous developers. Two of the best examples of open source code software are the GNU/Linux operating system and the Apache Web Server.

Closed source code refers to any program whose licensing terms do not qualify as open source. Generally, closed source means that the customer will only receive a binary version of the computer program that he or she licensed with no copy of the program's source code. Software modifications are then difficult because the usual way to modify a program is to edit its source code and then compile the code. The source code in this development model is considered a trade secret of the company, so parties who may receive access to the source code, such as colleges, have to sign non-disclosure agreements.

Note	Much of academic and scientific programming is based on free exchange of source code, as scientists freely share materials and methods in order to replicate experiments.

Closed source code still dominates commercial software development, but in the last few years, through the success of open source projects like Linux, KDE, and Apache, more open source coding is available. Seeing the success of these open source projects, corporate thinking is undergoing a transformation.

Assemblers

An **assembler** is a program that takes basic computer instructions and converts them into a pattern of bits that the computer's processor can use to perform its basic operations. Some people call these instructions assembly language.

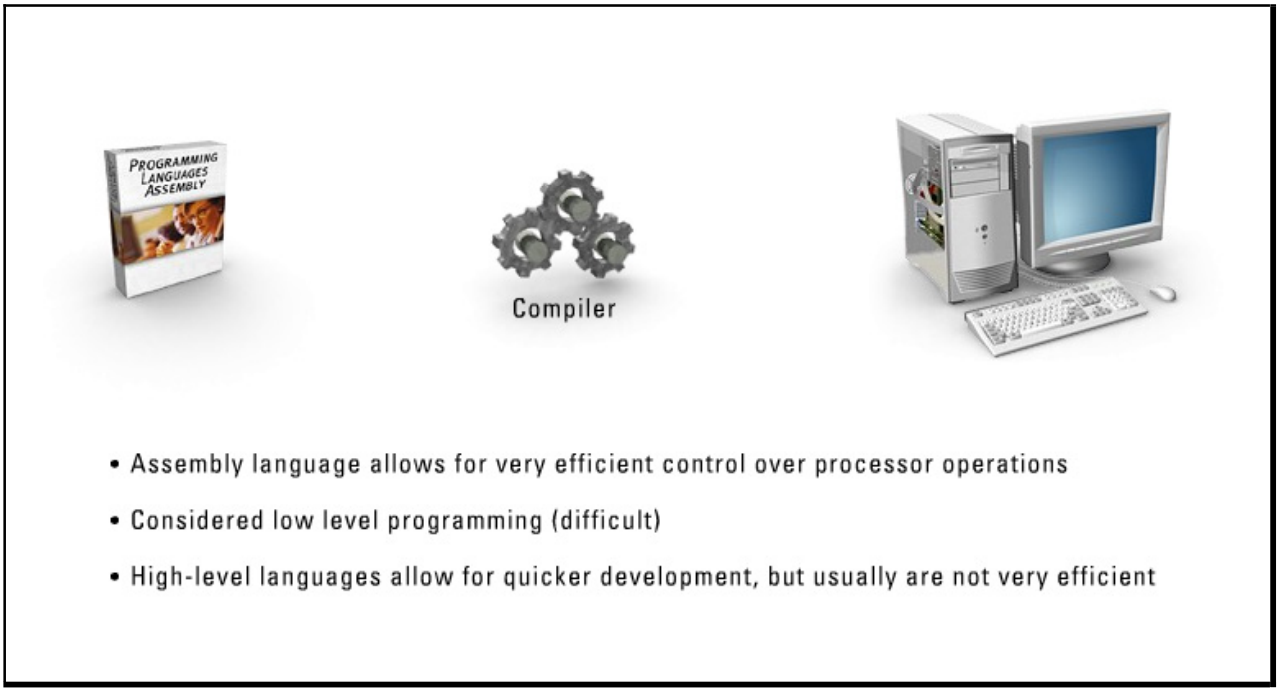

In the earliest computers, programmers actually wrote programs in machine code; assembly language or instruction sets were soon developed to speed up programming. Today, programmers only use assembler programming where they need very efficient control over processor operations. It requires knowledge of a particular computer's instruction set, however. Historically, programmers wrote code in higher-level languages such as COBOL, FORTRAN, PL/I, and C. These languages make coding faster and are easier to learn than the assembly languages.

The program that processes the source code written in assembly languages is called a compiler. Like the assembler, a compiler takes higher-level language statements and reduces them to machine code.

A compiler is distinguished from an assembler by the fact that each input statement does not, in general, correspond to a single machine instruction or fixed sequence of instructions. A compiler may support such features as automatic allocation of variables, arbitrary arithmetic expressions, control structures (such as FOR and WHILE loops), variable scope, input/output operations, higher-order functions, and portability of source code. In fact, some compilers output assembly language that is then converted to machine language by a separate assembler.

Compilers

A **compiler** is a special program that processes statements written in a particular programming language.

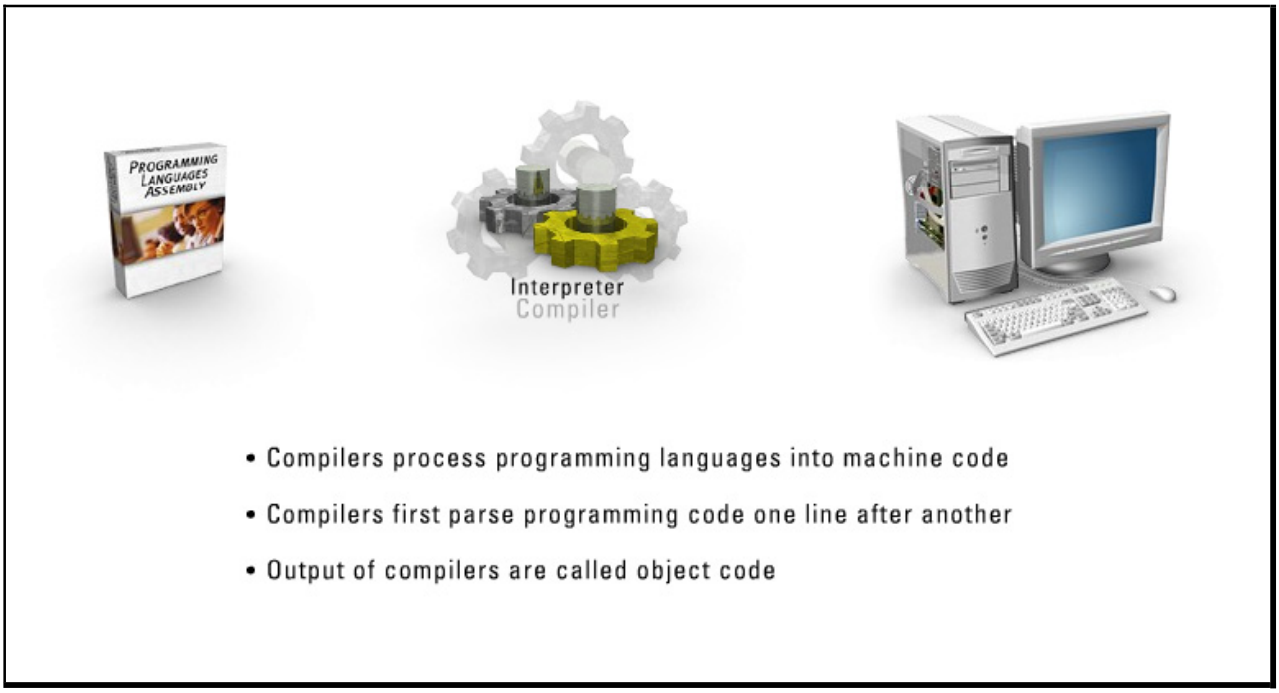

The compiler turns the code into machine language or a code that a computer's processor uses. Typically, a programmer writes language statements in a language such as Pascal or C one line at a time using an editor. The resulting file contains what are called the source statements. The programmer then runs the appropriate language compiler, specifying the name of the file that contains the source statements.

When executing or running, the compiler first parses or analyzes all of the language statements syntactically one after the other. In one or more successive stages or passes, it builds the output code, making sure that statements that refer to other statements are referred to correctly in the final code. Traditionally, the output of the compilation has been called object code or sometimes an object module. (Note that the term "object" here is not related to object-oriented programming.) The object code is machine code that the processor can process or execute one instruction at a time.

More recently, the Java programming language, a language used in object-oriented programming, has introduced the possibility of compiling output, called bytecode. This bytecode can run on any computer system platform that has a Java virtual machine or bytecode interpreter to convert the bytecode into instructions that can be executed by the actual hardware processor. Using this virtual machine, the bytecode can optionally be recompiled at the execution platform by a just-in-time compiler.

A compiler works with third generation languages (3GL) and higher-level languages. An assembler works on programs written using a processor's assembly language.

Interpreters

An **interpreter** is a program that reads textual commands from the user or from a file and then executes them.

EXAMPLE

```
type config.sys
```

- Interpreters take textual commands and interprets them into machine code the CPU can use
- Command interpreters include:
  - MSDOS command windows
  - Unix shells
  - Router exec prompts
- Program interpreters exist for all high level languages, from BASIC to C++

Some commands may be executed directly within the interpreter. An example of this execution is setting variables or control constructs. Other commands may cause the interpreter to load and execute other files.

Note	UNIX's command interpreters are known as shells.

When an IBM PC is booted, BIOS loads and runs the MS-DOS command interpreter into memory from the file COMMAND.COM found on a floppy disk or hard disk drive. The commands that COMMAND.COM recognizes such as COPY, DIR, and PRN, are called internal commands, in contrast to external commands, which are called executables.

It takes an interpreter less time to interpret code than the total time required to compile and run code. This feature is especially important when prototyping and testing code; an edit-interpret-debug cycle can often be much shorter than an edit-compile-run-debug cycle.

Interpreting code is slower than running the compiled code because the interpreter must analyze each statement in the program each time it is executed. The interpreter then performs the desired action whereas the compiled code just performs the action. This run-time analysis is known as interpretive overhead. Access to variables is also slower in an interpreter because the mapping of identifiers to storage locations must be done repeatedly at run time rather than at compile time.

Summary

The key points discussed in this lesson are:

- You must study and establish countermeasures for any portion of an application that can come under attack by a threat agent.
- Open source code is a term for a program whose license gives users the freedom to run the program for any purpose, to study and modify the program, and to redistribute copies of either the original or modified program without having to pay royalties to previous developers. Closed source means that the customer will only receive a binary version of the computer program that he or she licensed with no copy of the program's source code.
- An assembler is a program that takes basic computer instructions and converts them into a pattern of bits that the computer's processor can use to perform its basic operations.
- A compiler is a special program that processes statements written in a particular programming language.
- An interpreter is a program that reads textual commands from the user or from a file and then executes them.

Malicious Code

Overview

Everyday employees download, install, and use applications in corporate America. Attackers are well aware of this fact and often attempt to gain access to corporate resources. Attackers covertly apply malicious code into applications to gain access to resources or cause undesirable events to occur. This lesson will discuss the various ways attackers disseminate malicious code into relatively benign applications to gain entrance into enterprise resources.

Importance

The information security professional needs to understand the tactics and methods embedded in malicious code that attackers use to gain access to corporate resources.

Objectives

Upon completing this lesson, you will be able to:

- Define malicious code and its associated terms
- Define malware
- Define jargon
- Define covert channels
- Define hoaxes
- Explain the differences among hackers, crackers, and phone phreaks

Outline

The lesson contains these topics:

- Definitions
- Malware
- Jargon
- Covert Channels
- Hoaxes
- Hackers, Crackers, and Phone Phreaks

Definitions

Malicious code can be defined as any computer program or part of a computer program that is designed to take an action that the end user would not knowingly permit. This action might be intended as a joke such as a funny message appearing on the screen, or the code may render other computer programs useless. At its worst, malicious code may cause the deletion of an entire hard disk.

A computer **virus** is transmitted from computer to computer much like a biological virus passes from person to person. Like a biological virus, the computer virus is able to replicate itself, often using some of the host's resources to do so. Viruses therefore require a host program in order to function. According to Fred Cohen, an expert in information protection, a virus is "…a computer program that can infect other computer programs by modifying them in such a way as to include a (possibly evolved) copy of itself." Note that any program that transmits itself as defined above is considered a virus under this definition, whether or not its intent is malicious. Viruses may be further divided into several different types, which are discussed below.

File infectors piggyback on (usually executable) program files. The virus replicates whenever the infected file is executed, and attaches itself to more files. A "direct-action" file infector attaches itself to programs each time a program infected by it is launched, while a "resident" file infector lives in the computer's random access memory (RAM) and infects programs as they are opened. Either way, the virus gets around.

Boot sector infectors are not the huge problem that they used to be, mainly because manufacturers have begun adding special protection to the boot sectors of their storage media. The goal of a boot sector infector is to modify the section of the disk that is used to start up the computer. For this reason, it is important to boot a computer that may be infected with a virus from a "clean" or an uninfected disk.

To stand the best chance at longevity, viruses should be difficult to detect and remove. Virus authors have accomplished this goal in a number of interesting ways.

A **stealth virus** actually hides the damage it has done. In order to operate, the virus modifies critical files or the boot sector of a hard disk. When the operating system calls for the status of the modified files, the stealth virus sends back a phony version, hoping to fool the system (including any active virus detection software) into believing that everything is just fine. Since the stealth virus must be resident in memory, however, sometimes the virus detection software will detect that.

A **polymorphic virus** produces copies of itself that are different from the original, but still function to deliver whatever payload the virus has to offer. The hope here is that a virus scanner will not detect all instances of the virus, leaving a few behind.

An **armored virus** is difficult to trace, because it is encoded with special tricks or decoys that make it difficult to understand its code and its source.

Unlike a virus, a **worm** does not require a host program in order to survive and propagate itself, as it is a self-contained program. It sends itself, or parts of itself, to other computers, most often through a network connection. Once introduced to a network, the worm looks for other machines with a specific security breach or "hole," installs itself on those machines, and then begins replicating from them. A famous example of a worm is the Code Red worm, which replicated itself more than 250,000 times in nine hours on July 19, 2001.

Similar to a worm, a **Trojan horse** is a complete computer program unto itself. Its uniqueness stems from the fact that it is usually consciously installed and launched by the unaware user. The Trojan horse purports to do one thing, but in reality does another. It may be disguised as a game or free program, for example, but actually does some form of damage. Unlike worms and viruses, the Trojan horse cannot reproduce independently.

Malicious code can be detected by:

- File size increase
- Many unexpected disk accesses
- Change in update or modified timestamps

Malware

Malware is short for "malicious software" and is described as being any software designed specifically to damage or disrupt a system.

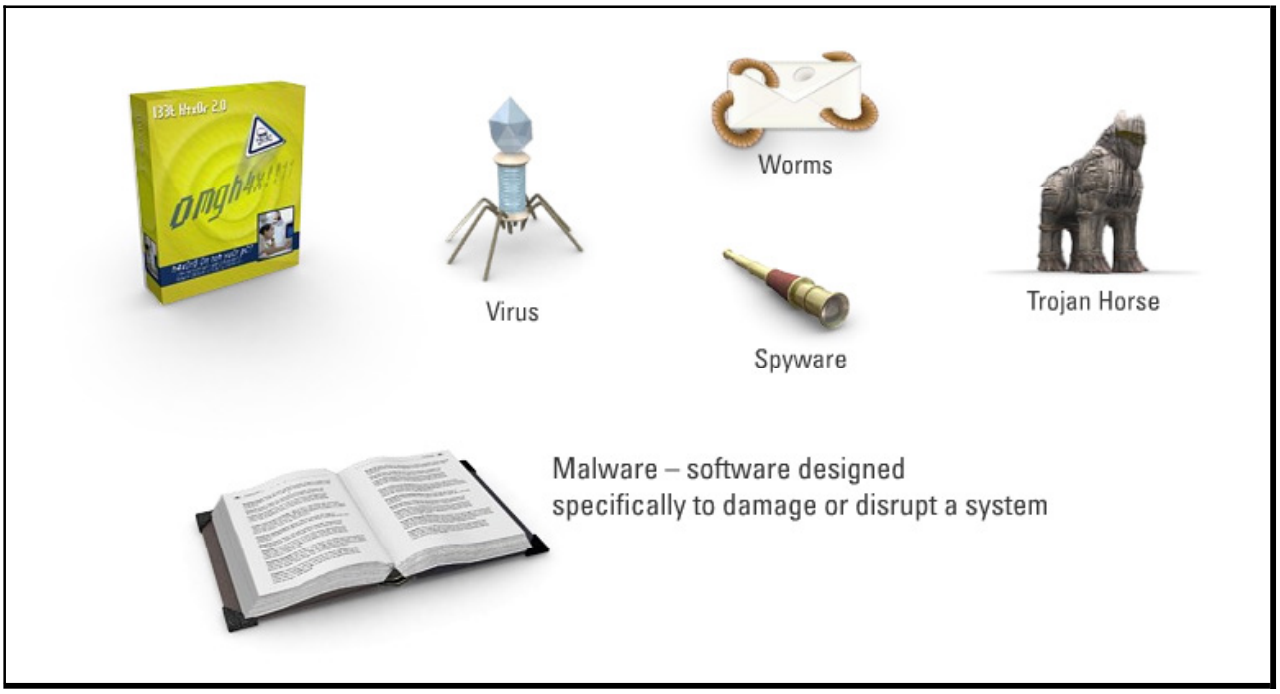

Examples of malware are viruses, worms, Trojan horses, and spyware. Spyware is software that gathers information about a computer user without the user's permission.

Jargon

Jargon can be defined as any words or definitions of computing terms that are "created" and accepted by the general public.

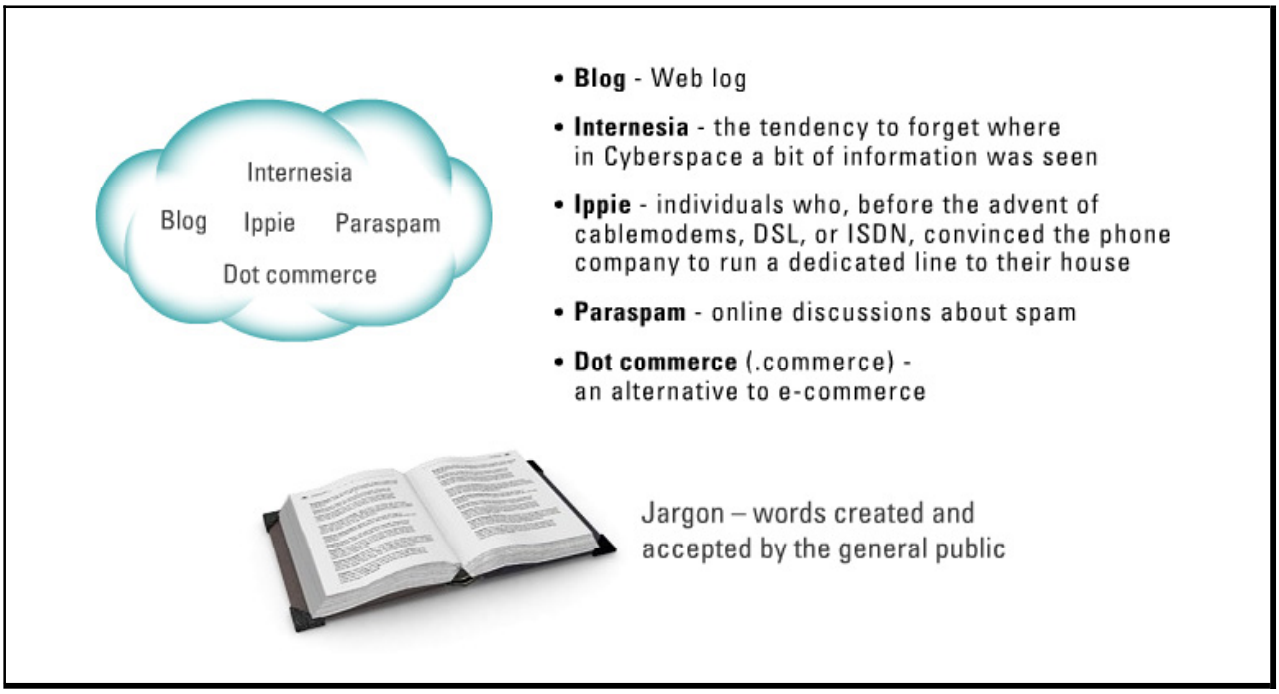

These words are usually specialized for some type of technical terminology and characterized for a particular subject. For example, the jargon word **blog** (short for "web log") is a type of web page that serves as a publicly accessible personal journal (or log) for an individual. Typically updated daily, blogs often reflect the personality of the author. Blog software usually has archives of old blogs, and is searchable. Frequently, web pages use blog software to provide information on many topics, although the content is usually personal and requires VERY careful evaluation.

Covert Channels

A **covert channel** can be defined as a communications channel that uses entities not normally viewed as data objects to transfer information from one subject to another.

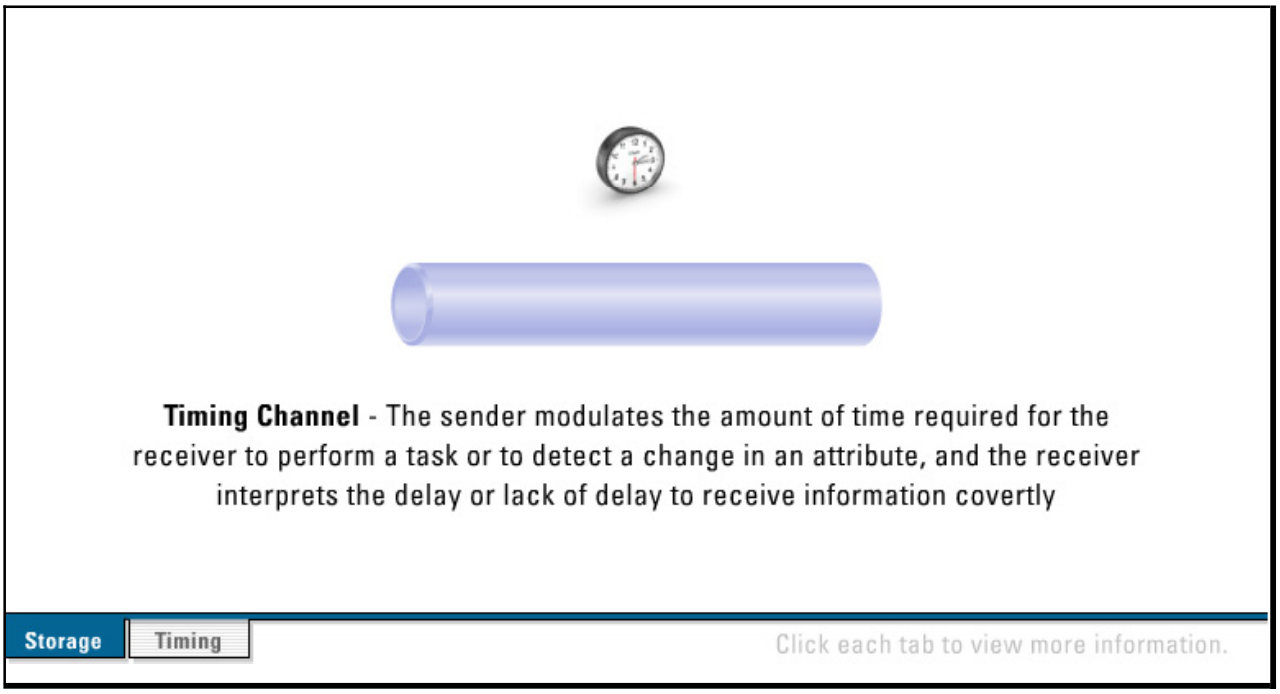

There are two forms of covert channels:

- **Storage Channels** - A potential covert channel is a storage channel if its scenario of use involves the direct or indirect writing of a storage location by one process and the direct or indirect reading of the storage location by another process.
- **Timing Channels** - A potential covert channel is a timing channel if its scenario of use involves a process that signals information to another by modulating its own use of system resources (e.g., CPU time) in such a way that this manipulation affects the real response time observed by the second process.

Hoaxes

A **hoax** is an act, a document, or an artifact intended to deceive the public. Examples range from relatively benign instances of trickery, such as April Fool's Day pranks, to scientific fraud on a grand scale, such as the Piltdown Man hoax of the early 20th century. Computer hoaxes are false statements meant to frighten users and usually end with a sentence or two urging readers to send an e-mail on to everyone they know. Hoaxes do nothing more than scare people into performing some rash action; it is, therefore, crucial to inform end users about hoaxes.

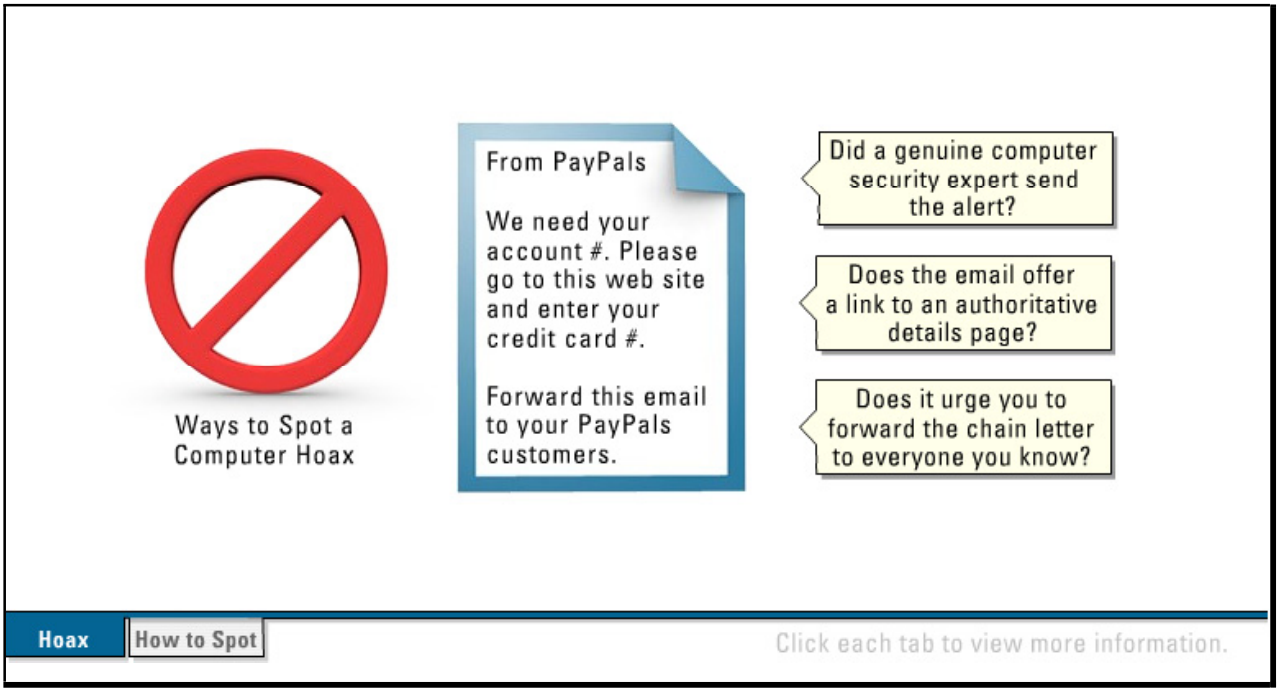

The advent of the Internet has provided an unparalleled platform for hoaxers by placing inexpensive, easy-to-use self-publishing tools at everyone's disposal. E-mail hoaxes spread false information quickly by encouraging recipients to forward fake documents, chain letter-style, to everyone they know. Web hoaxes consist of bogus web sites designed to fool users into believing they are visiting legitimate home pages, which in fact present false or misleading information.

Regardless of the specific form it may take, what distinguishes any hoax from mere error or folklore is that it is deliberately deceptive.

There are several ways to spot a computer virus alert that is a hoax:

- Did a genuine computer security expert send the alert?
- Does it urge you to forward the chain letter to everyone you know?
- Does the e-mail offer a link to an authoritative details page?

Hackers, Crackers, and Phone Phreaks

This topic will define hackers, crackers, and phone phreaks.

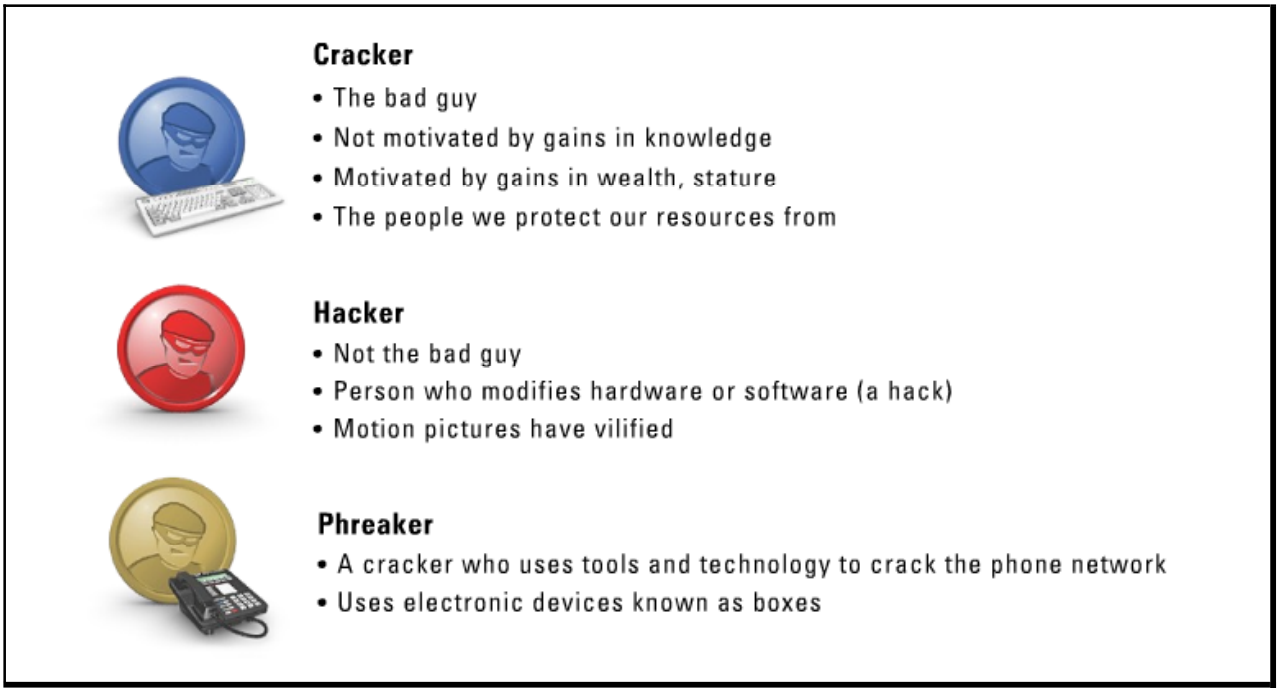

Film and television have taken the term **hacker** out of its original context and vilified it. When discussing information security with experienced security specialists, it is only polite to use the correct term when referring to a hacker or cracker. Hackers are not the bad guys (although motion pictures will have you think otherwise). Hackers are people who want to know how things work. To find out how things work, they dissect code, introduce new code, and generally just play with software (or hardware). Their intentions are not malicious but instead can sometimes be viewed as helpful. More times than not, they find a problem or potential problem that crackers may use for different, possibly malicious, purposes.

Crackers, or attackers, are not motivated by a gain in knowledge, but instead are motivated by a gain in wealth or stature. They want what others have and use technology to get what they want. If they cannot obtain a particular resource, they sometimes try to make it so nobody can get it, which is called a denial of service attack. Essentially, crackers are the people from which you are protecting your resources. They are smart, motivated, and knowledgeable in their art. You must know the methodologies, tools, and strategies of crackers if you are to stop their attacks on your resources.

Phonephreaks are those individuals who use the tools and technology to crack the phone network in order to make long distance phone calls for free. The tools of the phone phreak are electronic devices known as boxes:

- **Blue Boxes** - A device that simulates a tone in order to trick the telephone company's system into thinking the user is authorized for long distance service, which enables the user to make the call.
- **Red Boxes** - Simulates the sound of coins being dropped into a payphone.
- **Black Boxes** - Manipulates the line voltage to receive a toll-free call.

Summary

The key points discussed in this lesson are:

- Malicious code can be defined as any computer program or part of a computer program that is designed to take an action that the end user would not knowingly permit.
- Malware is short for “malicious software” and is described as being any software designed specifically to damage or disrupt a system.
- Jargon can be defined as any words or definitions of computing terms that are “created” and accepted by the general public.
- A covert channel can be defined as a communications channel that uses entities not normally viewed as data objects to transfer information from one subject to another.
- E-mail hoaxes spread false information quickly by encouraging recipients to forward fake documents, chain letter-style, to everyone they know.
- Hackers are people who want to know how things work. Crackers, on the other hand, are motivated by a gain in wealth or stature. You must know the methodologies, tools, and strategies of crackers if you are to stop their attacks on your resources.
- Phone phreaks are those individuals who use the tools and technology to crack the phone network in order to make long distance phone calls for free.

Methods of Attack

Overview

Attackers attempting to gain access to corporate resources for profit or fun use their knowledge and experience in ways others had not thought possible. But, once the security world obtains an understanding of the exploit used, the application, algorithm, or protocol is updated to mitigate the threat. Attackers then try different avenues of attack, which leads to an endless exploit/mitigation loop. This war between attackers and information security professionals has been going on for the past two plus decades. This lesson will discuss the various attack methods in use today and how security professionals mitigate their effectiveness in corporate networks.

Importance

Understanding the "tools of the trade" in use by attackers today gives the security professional a glimpse into the mind of the attacker and a more thorough understanding of how to mitigate each threat.

Objectives

Upon completing this lesson, you will be able to:

- Identify attacks seen on the Internet
- Define trapdoors
- Define buffer overflows
- Define brute force attacks
- Define denial of service attacks
- Define dictionary attacks
- Define spoofing
- Define pseudo flaws
- Identify reasons for crackers to alter authorized code
- Define flooding

Outline

The module contains these lessons:

- Internet Attacks
- Trapdoors
- Buffer Overflows
- Brute Force Attacks
- Denial of Service Attacks
- Dictionary Attacks
- Spoofing
- Pseudo Flaws
- Alteration of Authorized Code
- Flooding

Internet Attacks

Many different attacks can occur on all points of an information system. Simple attacks such as IP spoofing and the ping-of-death are easily thwarted today by proper filtering and system updates. This section will discuss some of the more prevalent attacks seen on the Internet.

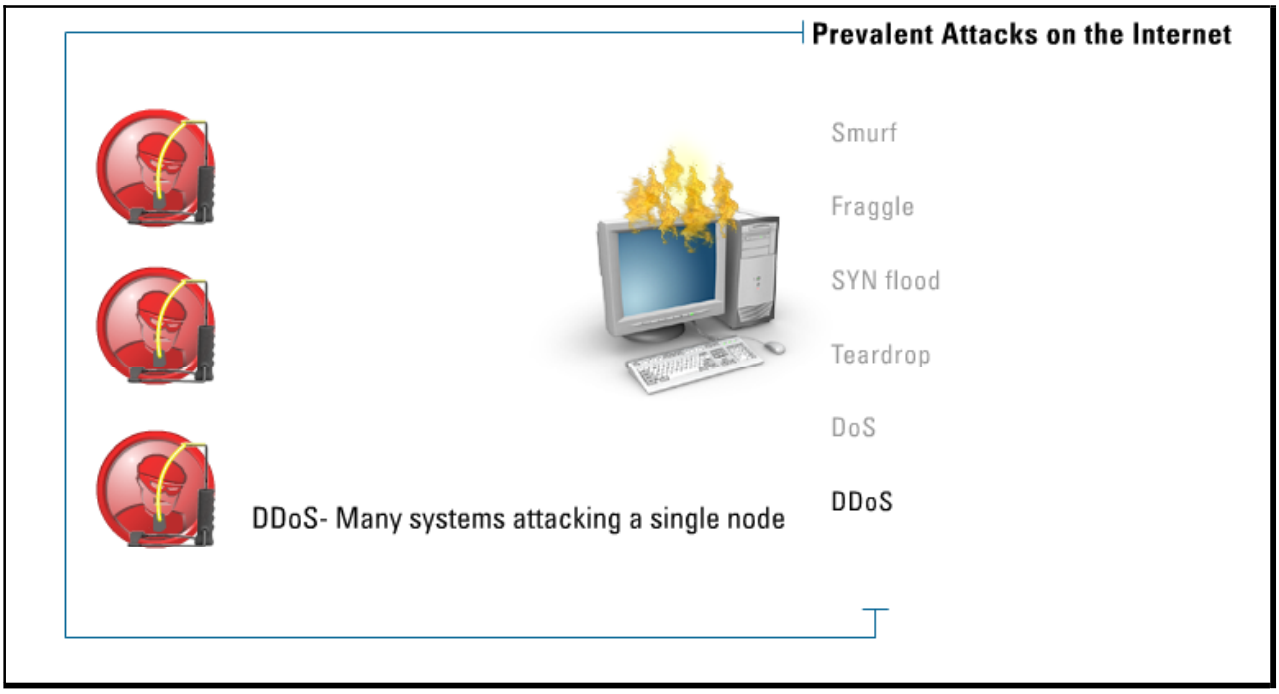

Smurf: An attack with three entities: the attacker, the victim, and the amplifying network. The attacker spoofs, or changes the source Internet Protocol (IP) address in a packet header, to make an Internet Control Message Protocol (ICMP) ECHO packet seem as though it originated at the victim's system. This ICMP ECHO message is broadcasted to the amplifying network, where all active nodes send replies to the source (the victim). The victim's system and network become overwhelmed by the large amounts of ECHO replies.

Fraggle: A Fraggle attack is the same type of attack as the Smurf attack, except Fraggle usesthe User Datagram Protocol (UDP) as its weapon of choice. The attacker broadcasts a spoofed UDP packet to the amplifying network, which in turn replies to the victim's system

SYN Flood: In a SYN flood attack, the victim is continually sent SYN messages with spoofed source addresses. The victim will commit the necessary resources to set up this communication socket, and it will send its Synchronize-Acknowledge (SYN-ACK) message, waiting for the ACK message in return. No ACK will ever return and if enough SYNs are sent, resources could become depleted.

Teardrop: In a teardrop attack, the attacker's IP puts a confusing offset value in the second or later fragment. If the receiving operating system does not have a plan for this situation, it can cause the system to crash.

Denial of Service (DoS): A denial of service (DoS) attack consumes the victim's bandwidth or resources, causing the system to crash or stop processing other packets.

Distributed Denial of Service (DDoS): The distributed denial of service (DDoS) attack is a logical extension of the DoS attack. The attacker creates master controllers that can in turn control slaves/zombie machines, all of which can be configured to attack a single node.

DNS DoS Attacks: In a DNS DoS attack, a record at a domain name server (DNS) server is replaced with a new record pointing at a fake/false IP address.

Cache Poisoning: In a cache poisoning attack, the attacker inserts data into the cache of the server instead of replacing the actual records.

Trapdoors

Trapdoors are also referred to as one-way functions. In the computer context, they are mathematical functions easy to compute in one direction, but very difficult or impossible in the other; they are often used in cryptography. Two examples of trapdoor functions are public key cryptography and message digests.

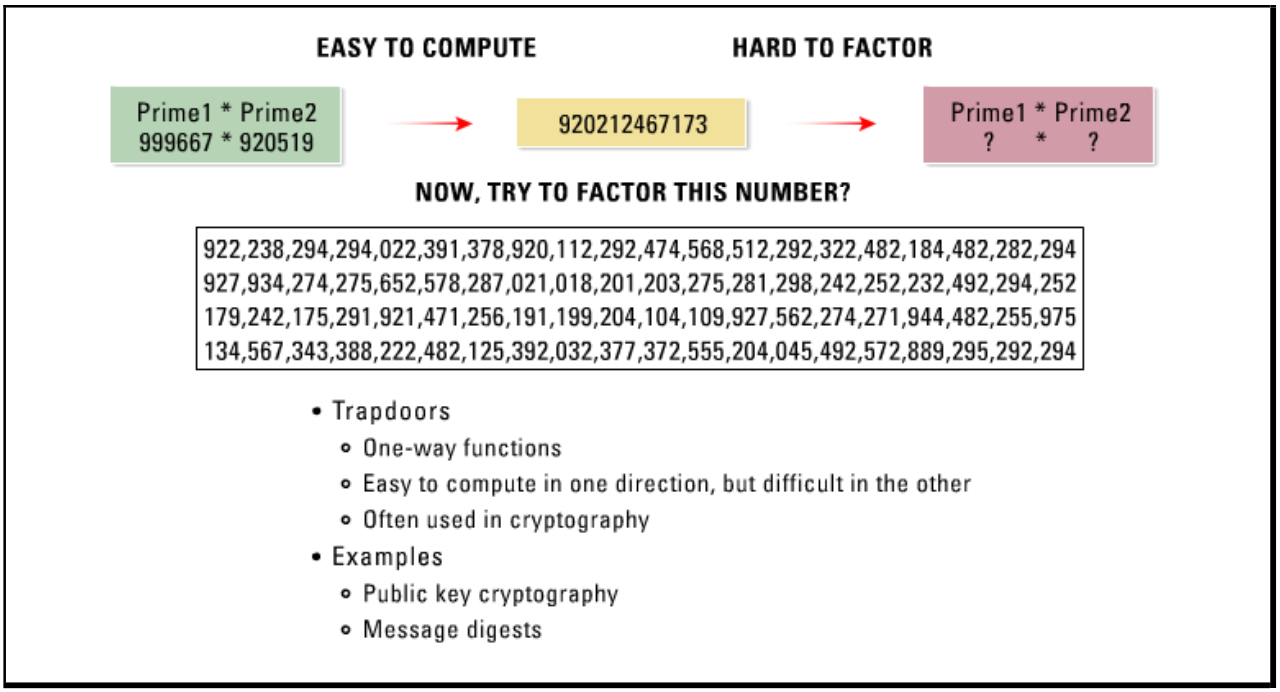

Public key cryptography relies on the multiplication of two large prime numbers. It is easy to multiply two large prime numbers to get a result, but not very easy to have a result and figure out which two prime numbers were used to obtain it. In order to obtain the values, the result would have to be factored, which is very time consuming. It can be done, but even the world's fastest supercomputer would take eons to determine the values.

Message digests are another great example of trapdoor functions. With message digests, you take an input and place it into an algorithm. This value can be anything: a phrase, a text document, a dictionary, an IP packet, etc. But, with message digests, no matter what the size of the input into the algorithm, the result will always be a value of a certain size, usually 128 bits (MD-5) or 160 bits (SHA1). Message digests can also be described as being lossy algorithms. If you take a 1500-byte packet, enter it into the message digest algorithm, and obtain a 128-bit result, there will be no way to "decrypt" the 128-bit value back into its 1500-byte form. Many people wrongly refer to message digests as encryption algorithms. Message digests are used as integrity checks, not for encryption. Once the input is placed into the algorithm, there is no way to get it back, hence the trapdoor one-way function.

Buffer Overflows

A **buffer overflow** is a software-based attack created when a program does not check the length of data that is inputted into it, which will then be processed by the CPU.

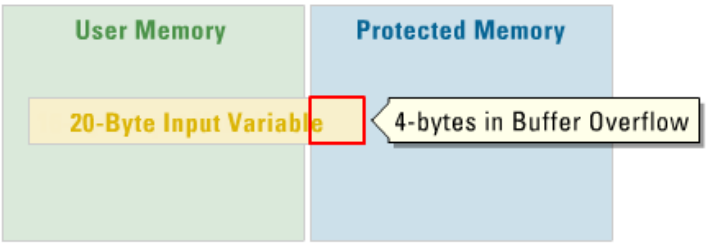

Buffer overflow is created when programmers do not perform input validation

A buffer overflow exists when a particular program attempts to store more information in a buffer "memory storage" than it was intended to hold. Since the buffer was only intended to hold a certain amount of data, the additional data overflows into a different area of memory. It is this different area of memory where overflows cause the problem.

For example, say you have two areas of adjacent memory; one area is assigned to normal user privilege mode and another area is assigned to superuser privilege mode. Under normal circumstances, the normal privilege service cannot enter into and execute commands in the superuser area of memory. If this is attempted, the operating system will deny access to the user service as it does not have the correct privilege level. Now, take a situation where a buffer overflow occurs in the user service. Data is written to memory and overflows into the area of memory where the service running at the superuser privilege is located. But, the attacker who wrote the code knew the exact boundary where the two memory locations reside, and the code that was written to the superuser area included a command to execute an EXEC session for any user connecting to a particular port, say 9999. When the overflow commands are executed via normal operations of the program running in superuser mode, the hack will take effect. The cracker needs only to connect to port 9999 to have superuser level access to the EXEC session. From here he will have complete control of the system.

Note	Buffer overflows are also referred to as "smashing the stack".

Brute Force Attacks

Brute force attacks occur when a cracker attempts to obtain the correct password for an account by trying every conceivable value hoping to stumble across the correct one.

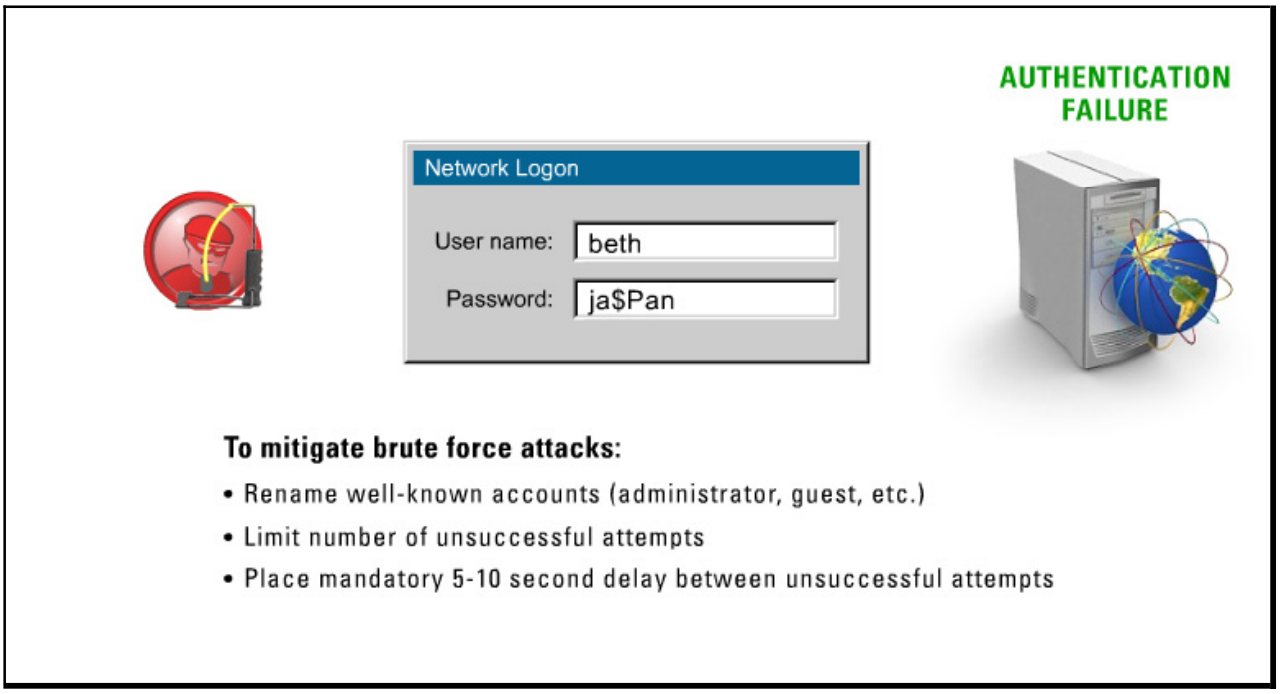

For example, the cracker may try to access the administrator account with the password of "a," and then try a password of "b," and then try a password of "c," and so on until the cracker has tried every conceivable value for a single digit. Next, the cracker tries "aa," "ab," and so on until the cracker eventually stumbles across the correct password.

Administrators have known about brute force attacks for many, many years and have come up with ways to mitigate these types of attacks. One of the easiest methods is to rename the administrator account to something else. In this way the cracker must know two things, the account name and the password. Administrators will also create passwords of at least eight characters in length. This technique helps because it takes time to brute force an attack on a password that is at least eight characters long. Hopefully, the administrator will notice the attack and take precautionary steps to block the cracker.

Attackers have in turn known about the precautionary steps administrators take to stop brute force attacks on live computers, so they have simply changed their tactics. Instead of attempting a brute force attack directly on the system, attackers attempt to first exploit some weakness in the OS and obtain the encrypted password database. This could be the shadow password file on UNIX or the SAM database on Microsoft Windows. Once a cracker obtains this file, he or she then attempts to perform an offline brute force attack on the password files.

Denial of Service Attacks

Denial of Service (DoS) attacks are carried out by attackers with an intent to stop legitimate users from accessing certain resources. Their intent is malicious and not designed to obtain information.

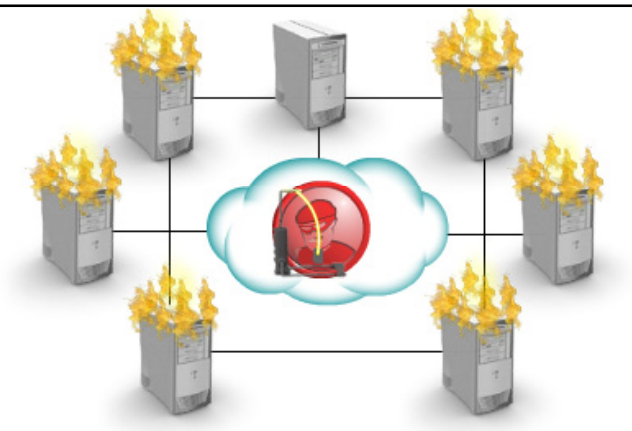

- DoS attacks- Most formidable of attacks to deal with
- Involve very large amounts of traffic
- DoS mitigation includes:
  - Rate limiting
  - NIDS
  - HIDS

DoS attacks are usually the most formidable of attacks to deal with as they usually involve very large amounts of traffic that may or may not 'look' on the wire as valid transmissions. Knowing how these attacks are sculpted and executed will allow network administrators to better deter them on their networks. Mitigation of DoS attacks can be performed at the ISP egress router into the company via rate limiting, via NIDS, HIDS, and by have up to date security patches and hot fixes installed on all critical servers and systems.

Dictionary Attacks

Dictionary attacks are another form of brute force attacks and take advantage of a well-known flaw in the password authentication scheme. That flaw is the fact that many people use common words as the password for an account. Attackers exploit this fact by using a source for common words (the dictionary) to try to obtain a password for an account. They simply try every possible word in the dictionary until a match is found.

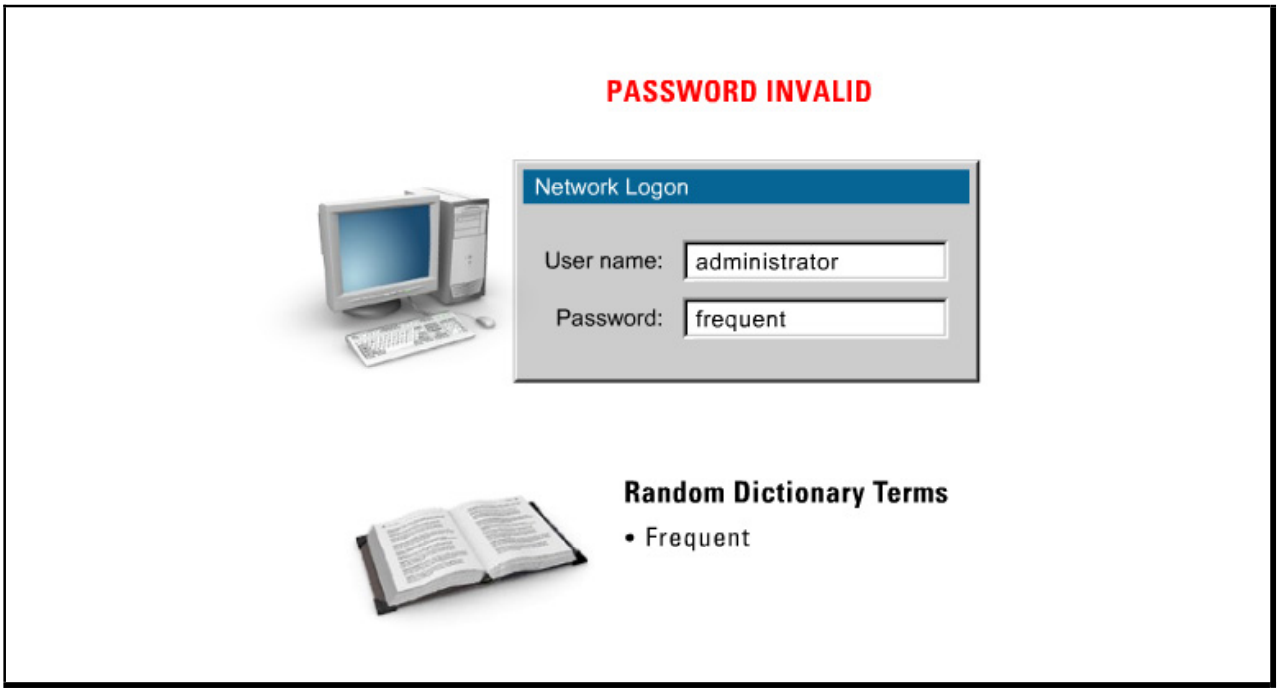

Dictionary attacks can be performed using ready-made username files and ready-made password files. In this modified dictionary attack, the attacker uses two database files, one with proper names and a second with common password phrases, such as letmein, admin, or super. The program will attempt to use the first proper name in the first database for the username field, and then will attempt each word in the common passwords file. If no match is made, the second proper name is used, and again all common passwords will be attempted, and so on until a match is made.

Note	Common brute force attacking utilities include L0phtCrack, Brutus, and John the Ripper.

Spoofing

Attackers can use many different types of spoofing attacks, but they all use spoofing for one reason, which is to impersonate another host. Sometimes the attacker does not care who he or she is impersonating; the attacker only cares that the packet he or she is transmitting does not identify him or her. Other times the attacker knows exactly what host he or she wants to impersonate and wants the return traffic to reach this host.

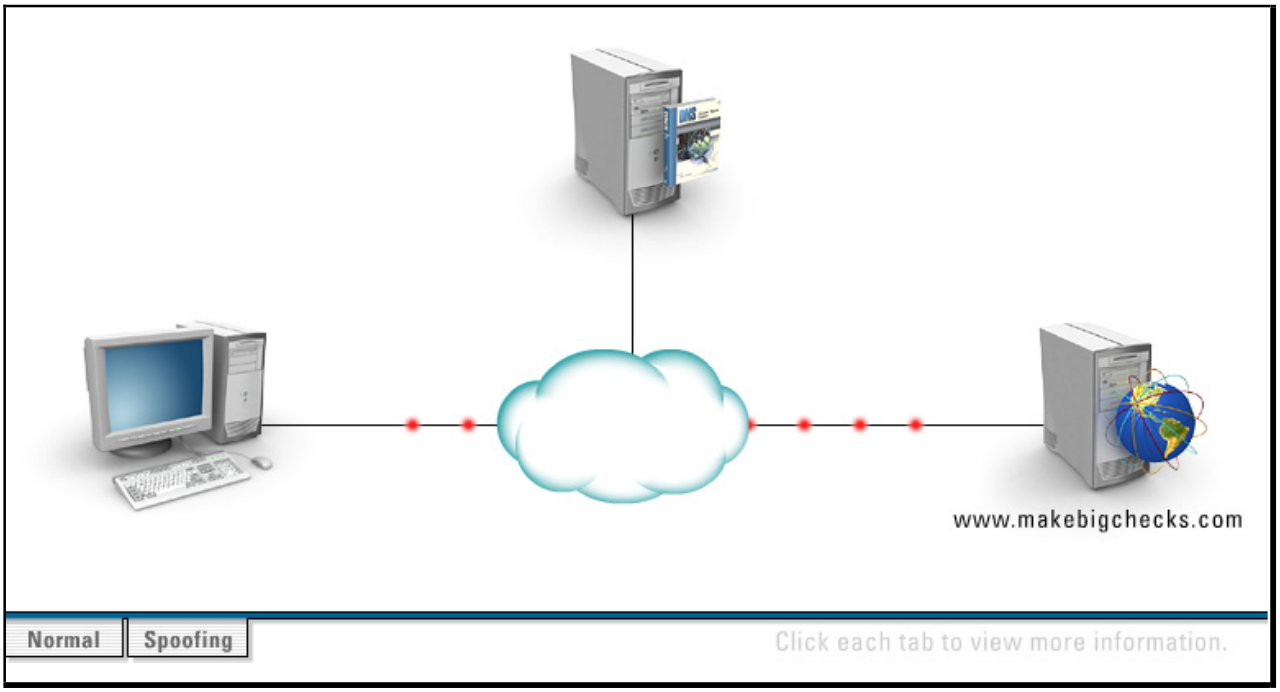

DNS spoofing attacks work by convincing the target machine that the machine that it wants to contact (for example, www.makebigchecks.com) is the machine of the attacker. When the target issues a DNS query, it could be intercepted and replied with the spoofed IP address, or the query could reach the DNS server, which has been tampered with in order to give the IP address of the cracker's host, rather than the real server's IP address. Either way the target receives a false IP address of the target and will attempt to contact it.

In a Transmission Control Protocol (TCP) takeover attack, the cracker will attempt to insert malicious data into an already existing TCP session between two hosts. In this type of attack, the attacker is either attempting to inject false data into the conversation, or take over the session completely. This type of attack is usually used in conjunction with a DoS attack to stop the host it is impersonating from sending any further packets.

The DoS attack against the impersonated host will itself be using spoofed packets. In this way, the attacker will hide his or her identity from the host he or she took over the TCP session from, while the opposite end still believes its ongoing session is with the original host.

Pseudo Flaws

A **pseudo flaw** is an apparent loophole deliberately implanted in an operating system or program as a trap for intruders.

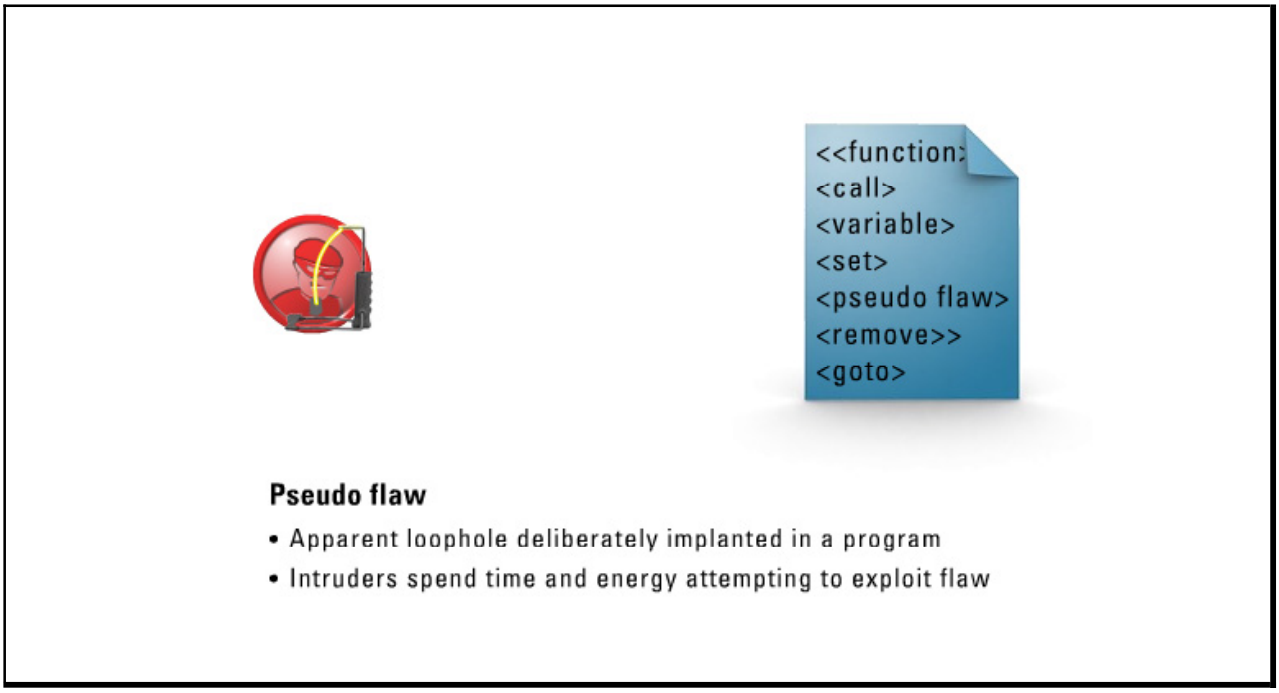

Pseudo flaws are inserted into programs to get attackers to spend time and energy attempting to uncover weaknesses in programs that they hope will allow them to gain access to other parts of the system. Because these are deliberate flaws, the attacker can spend weeks attempting to exploit the flaw, before he or she becomes discouraged and moves on to different parts of the program.

Alteration of Authorized Code

Attackers often write small programs that create a patch in authorized code.

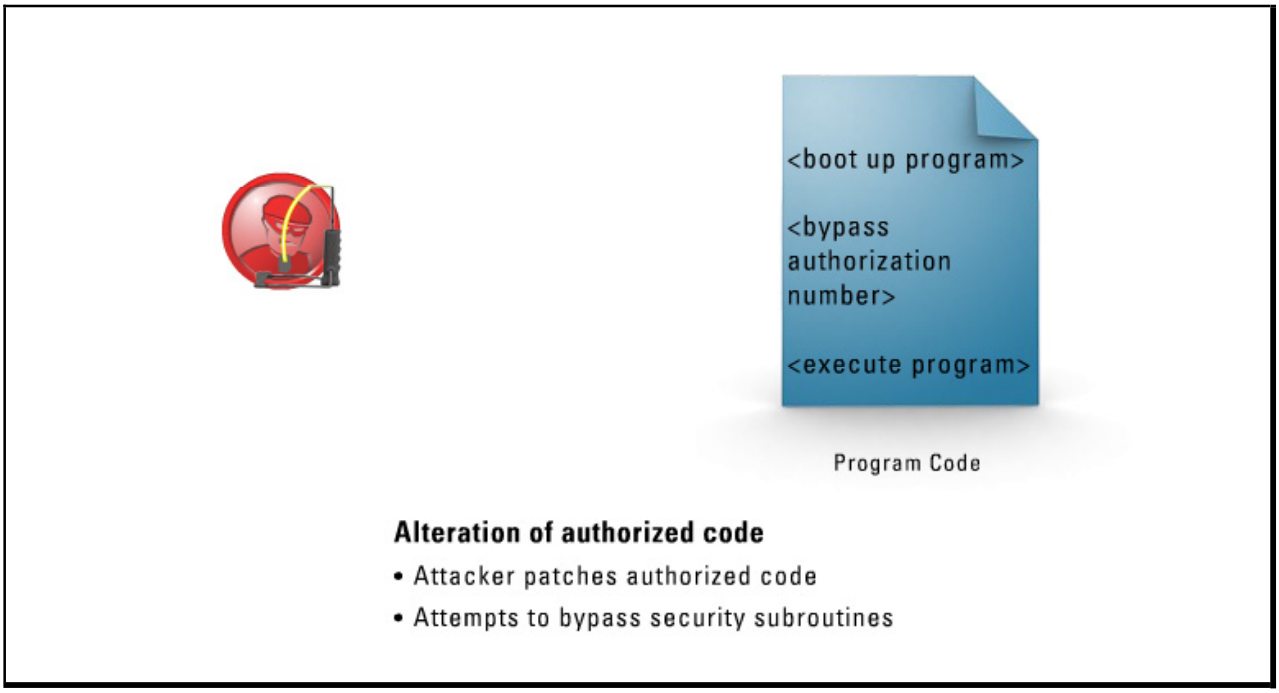

Take a program that will not execute until the user enters a valid serial number or authorization code. The attacker does not have this information, yet still wants to execute the program. Using his or her knowledge of programming and off-the-shelf software, the attacker can identify where in the program the subroutine that performs authorization is called from. The attacker then writes a program that modifies that very same area of the program, but instead of calling the authorization subroutine, the instructions are now a series of NOPs (no operations). This alteration of authorized code simply bypasses the authorization subroutine and begins executing the program.

Flooding

Flooding is the process of overwhelming some portion of the information system. This could be bandwidth on a serial link or memory in a router or server.

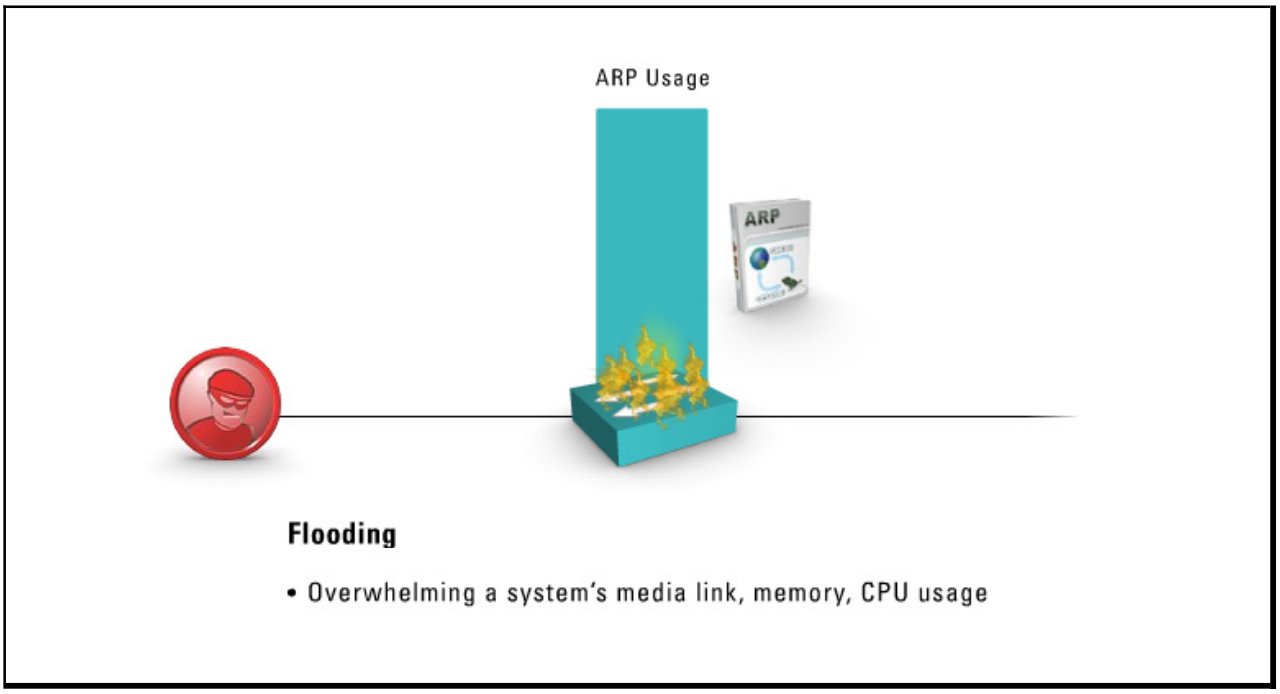

There are many uses of flooding for attackers. Attackers could hide their attacks in a flood of 'random' attack packets, they could attempt to overwhelm a switch's Address Resolution Protocol (ARP) table, or they could perform DoS attacks. SYN floods are an example of flooding used in a DoS attack.

SYN floods take advantage of TCP's three-way-handshake. In this DoS attack, the attacker sends many thousands of half-formed or embryonic TCP connection requests (SYN packets), usually with a spoofed source address, to the target server. The server that receives these connection requests sets aside a small amount of memory for each connection, and replies with an SYN-ACK to the spoofed address. The spoofed host (if it exists) receives the SYN-ACK packet and discards it. This leaves the server with an 'open' or a half-formed connection, which will remain so for three minutes as it waits for the connection to complete.

A few open connections will not cause harm to a server, but thousands upon thousands of open connections, each using a small amount of memory, will quickly consume all available resources on the server. When all resources are consumed, the server will no longer respond to the SYN requests of the attacker. Unfortunately, the server will also not respond to any SYN request from a valid user, which is the DoS the attacker is trying to perform.

Summary

The key points discussed in this lesson are:

- Many different attacks can occur on all points of an information system.
- Two examples of trapdoor functions are public key cryptography and message digests.
- A buffer overflow exists when a particular program attempts to store more information in a buffer "memory storage" than it was intended to hold.
- Brute force attacks occur when a cracker attempts to obtain the correct password for an account by trying every conceivable value hoping to stumble across the correct one.
- DoS attacks are usually the most formidable of attacks to deal with as they usually involve very large amounts of traffic that may or may not 'look' on the wire as valid transmissions.
- Dictionary attacks take advantage of a well-known flaw in the password authentication scheme. That flaw is the fact that many people use common words as the password for an account.
- Attackers can use many different types of spoofing attacks, but they all use spoofing for one reason, which is to impersonate another host.
- A pseudo flaw is an apparent loophole deliberately implanted in an operating system program as a trap for intruders.
- Attackers often write small programs that create a patch in authorized code.
- Flooding is the process of overwhelming some portion of the information system.

Databases and Data Warehousing

Overview

Databases are the repositories of an enterprise's daily life. Whatever happens in the enterprise, whether it is through accounts payable, accounts receivable, billing, trade requests, etc., will eventually be entered into a database. Without readily available database information, an enterprise would come to a screeching halt. This makes database care essential to the well being of the overall enterprise. This lesson will discuss the various models of data warehousing and what protocols are being used to access the information contained in databases.

Importance

It is important for the information security professional to understand the flow of data into and out of a database in order to effectively secure the database from attack or unwanted exposure.

Objectives

Upon completing this lesson, you will be able to:

- Identify the fundamental elements of the Database Management System (DBMS)
- Describe the hierarchical data model
- Describe the relational database model
- Describe the distributed data model
- Identify the main components of the SQL database
- Describe the network data model
- Describe the object-oriented database model
- Identify the goal of Open Database Connectivity (ODBC)
- Describe the key feature of Object Linking and Embedding (OLE)
- Identify the main characteristics of Extensible Markup Language (XML)
- Describe the functionality of Java Database Connectivity (JDBC)

- Describe the key feature of ActiveX Data Objects (ADO)
- Describe the functionality of online transaction processing (OLTP)
- Define data mining
- Define inference
- Define polyinstantiation

Outline

The lesson contains these topics:

- DBMS Architecture
- Hierarchical Model
- Relational Model
- Distributed Model
- Structured Query Language (SQL)
- Network Database Management Model
- Object-Oriented Database Model
- Open Database Connectivity (ODBC)
- Object Linking and Embedding (OLE)
- Extensible Markup Language (XML)
- Java Database Connectivity (JDBC)
- ActiveX Data Objects (ADO)
- Online Transaction Processing (OLTP)
- Data Mining
- Inference
- Polyinstantiation

DBMS Architecture

A **Database Management System (DBMS)** is a suite of application programs that typically manage large, structured sets of persistent corporate data.

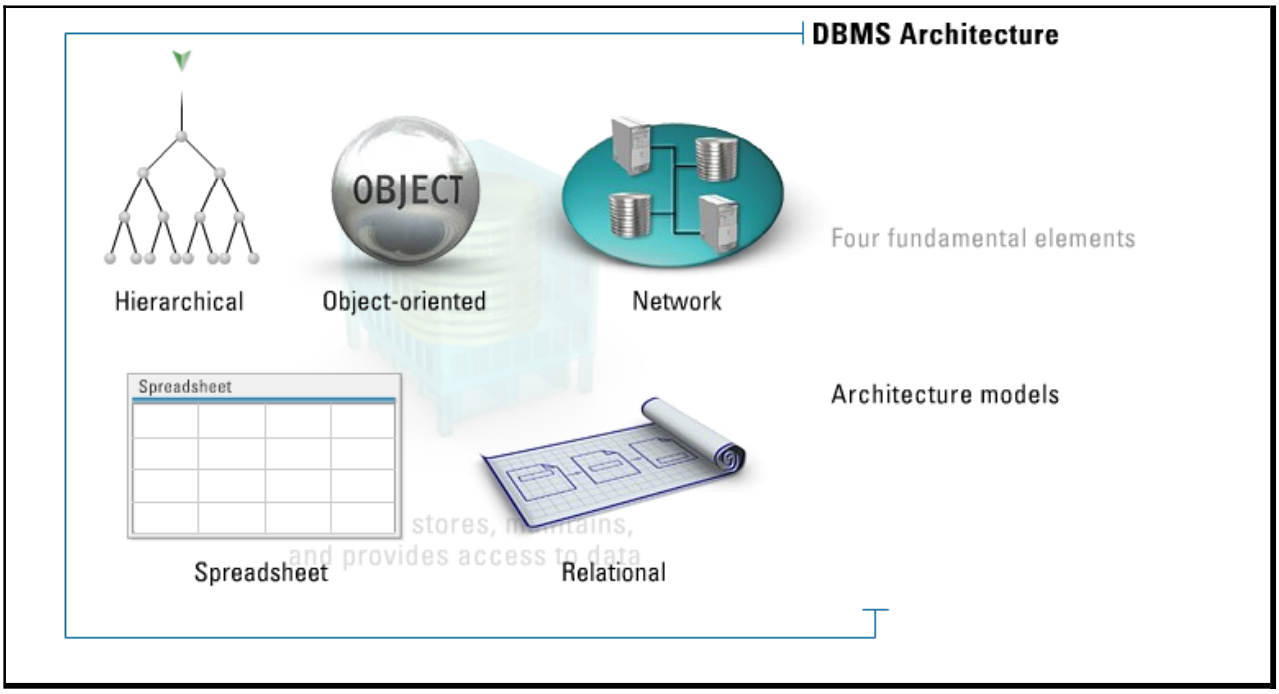

A DBMS is created to store, maintain, and provide access to this data using ad hoc query capabilities. The DBMS will provide the structure for the data and some form of language for accessing and manipulating the data. The main objective of the DBMS is to store data and allow users to view, edit, and manage the data.

A DBMS typically has four fundamental elements:

- Database
- Hardware
- Software
- End users

The data in the DBMS consist of individual entities and entities with relationships linking them together. Database models map or organize the data entities and should be created to provide for:

- **Transaction Persistence** - The state of the database must be the same after a transaction has occurred
- **Fault Tolerance and Recovery** - If a failure occurs, the data should remain in their original state
  - **Rollback** - A method of recovery that backs out an incomplete or invalid transaction
  - **Shadow** - A method of recovery that uses transaction logging to identify the last good transaction
- **Sharing by Multiple Users** - All data must be available to multiple users at the same time
- **Security Controls**- Access controls and integrity checking

Several architectural models exist for databases:

- **Hierarchical** - A method that uses the parent/child relationship through trees
- **Spreadsheet** - Simple databases created from spreadsheet programs (Excel, Lotus 123, Quattro Pro, etc.)
- **Object-Oriented** - Similar to OO programming languages, where data are stored as objects
- **Relational** - A method based on set theory and predicate logic, which provides a high level of abstraction
- **Network** - An extended form of the hierarchical model; represents its data in the form of a network of related records and sets, forming a network of links

Hierarchical Model

The **hierarchical data model** combines related records and fields in a logical tree structure. The records can have one child, many children, or no children. This model is useful for mapping one-to-many relationships.

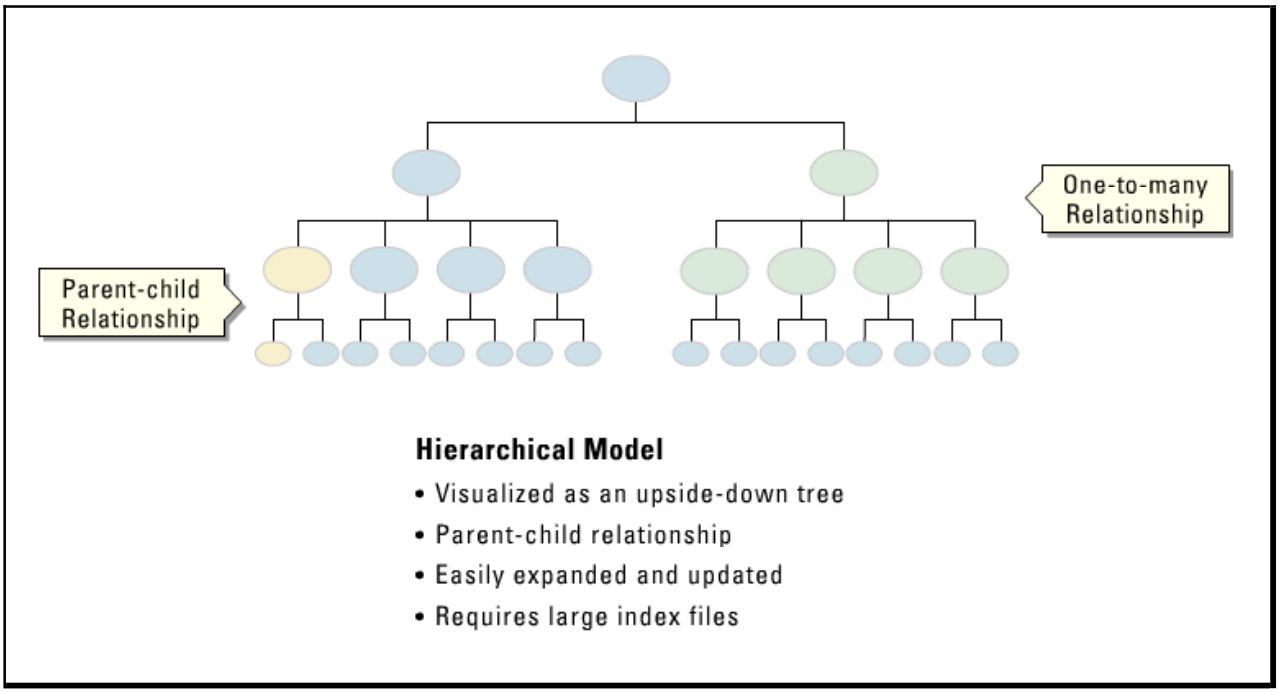

Perhaps the most intuitive way to visualize this type of relationship is by visualizing an upside down tree of data. In this tree, a single table acts as the "root" of the database from which other tables "branch" out.

You can think of relationships in such a system in terms of children and parents such that a child may only have one parent but a parent can have multiple children. Parents and children are tied together by links called "pointers" (perhaps physical addresses inside the file system). A parent will have a list of pointers to their children.

This child/parent rule assures that data is systematically accessible. To get to a low-level table, you start at the root and work your way down through the tree until you reach your target. Of course, as you might imagine, one problem with this system is that the user must know how the tree is structured in order to find anything.

Hierarchical database systems are easily expanded and updated. But due to their nature of growth, they require large index files and must be frequently maintained.

Relational Model

The **relational database model** uses attributes (columns) and tuples (rows) to contain and organize information. A primary key is a field that links all the data within a record to a corresponding value, while a foreign key is a field that represents a reference to an entry in some other table.

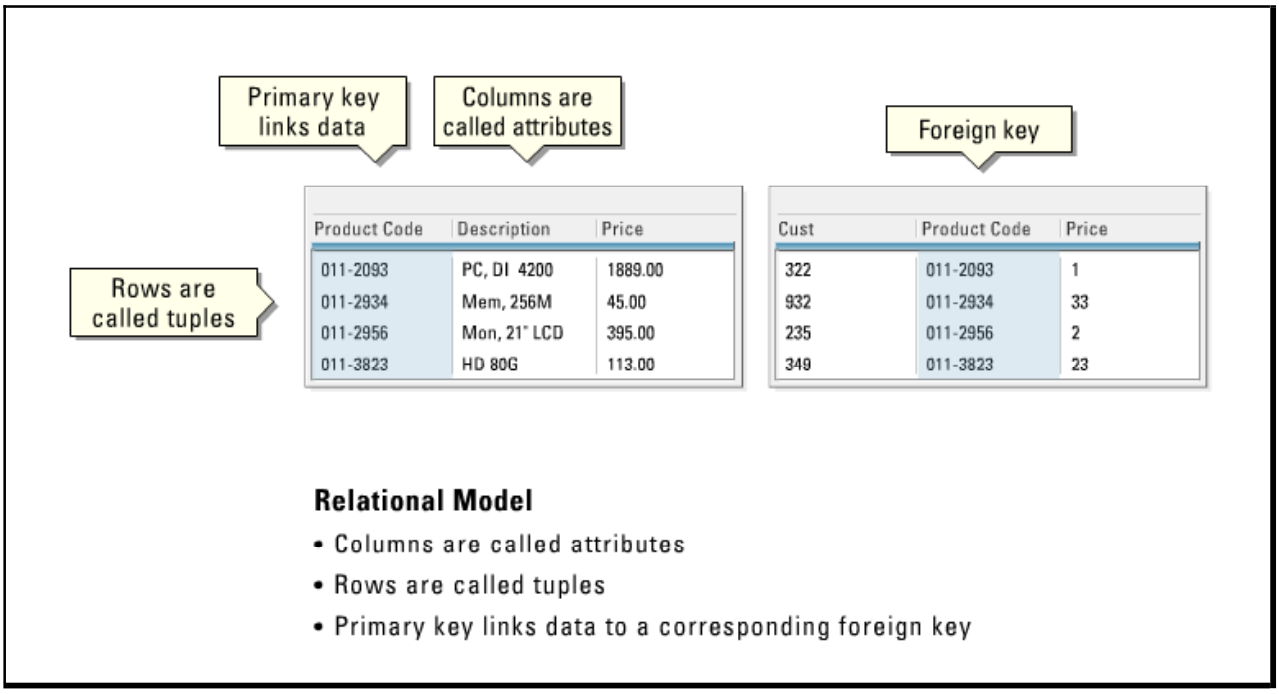

Relational database components:

- **Data Definition Language (DDL)** - Defines the structure and schema of the database
  - **Structure** - Table size, key placement, views, and data element relationships
  - **Schema** - The type of data that will be held and manipulated and their properties
- **Data Manipulation Language (DML)** - All the commands that enable a user to view, manipulate, and use the database
- **Query Language** - Enables users to make requests of the database
- **Report Generator** - Produces printouts of data in a user-defined manner

Properties of relational tables:

- Values are atomic
- Each row is unique
- Column values are of the same kind
- The sequence of columns is insignificant
- The sequence of rows is insignificant
- Each column has a unique name

Distributed Model

In the **distributed data model**, data are stored in more than one database, but are logically connected. This model enables different administrators to manage different databases, although one person or group must manage the entire logical database.

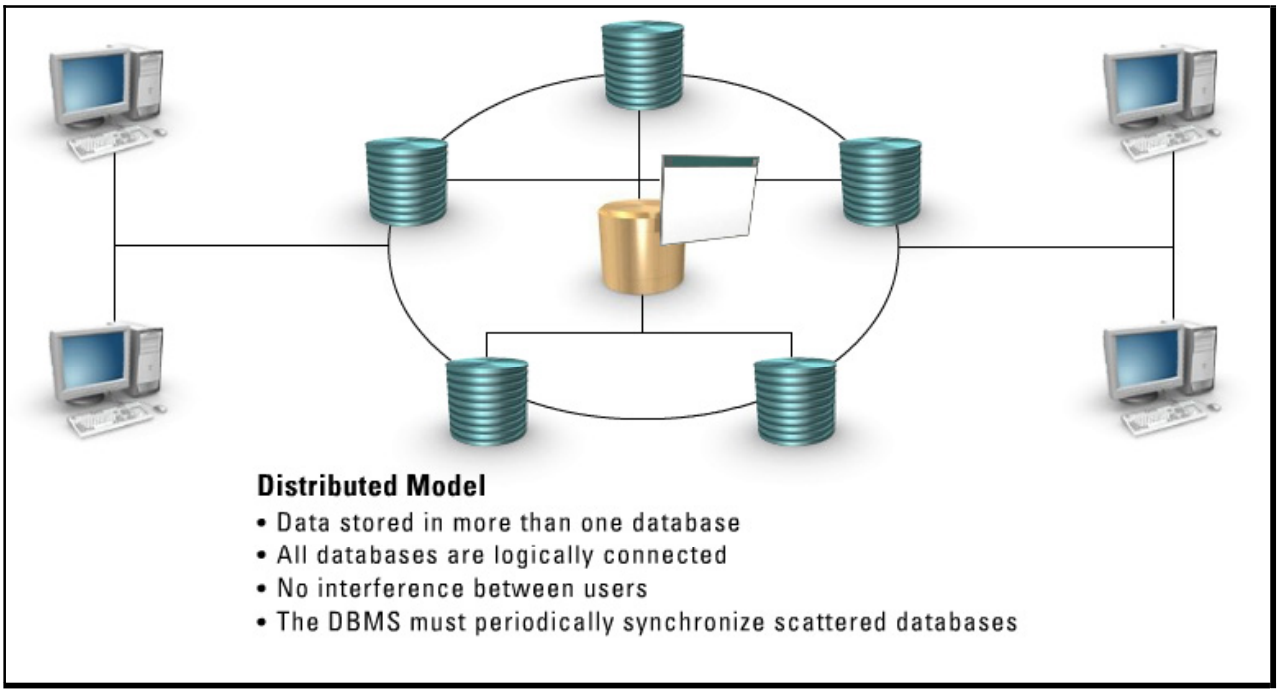

Because the database is distributed, different users can access it without interfering with one another. However, the DBMS must periodically synchronize the scattered databases to make sure that they all have consistent data.

Structured Query Language (SQL)

Structured Query Language (SQL) is one of the standardized languages used in the relational model. SQL is a standard interactive and programming language for getting information from and updating a database. Although SQL is both an American National Standards Institute (ANSI) and an International Organization for Standardization (ISO) standard, many database products support SQL with proprietary extensions to the standard language.

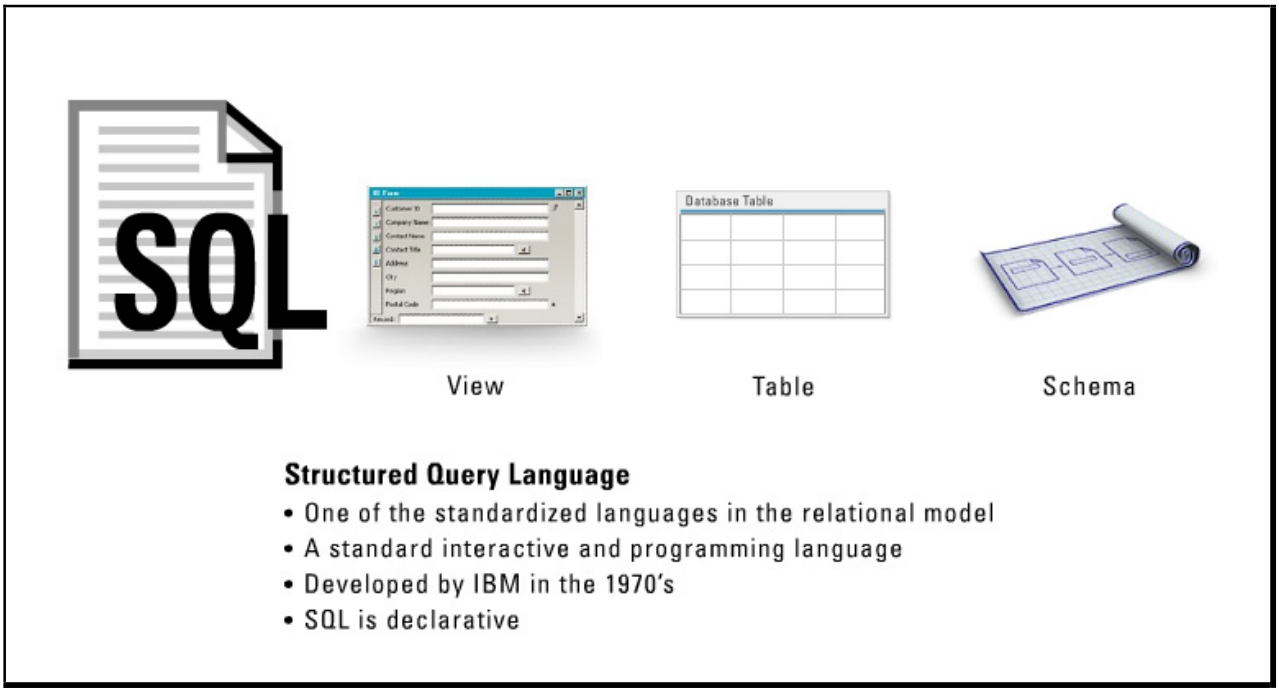

IBM developed SQL in the mid-1970s as a way to get information into and out of relational database management systems. A fundamental difference between SQL and standard programming languages is that SQL is declarative. You specify what kind of data you want from the database and the Relational Database Management System (RDBMS) is responsible for figuring out how to retrieve the data.

The main components of the SQL database include:

- **Schemas** - Schemas describe the overall structure of the database, including the access control scheme to be used
- **Tables** - The table is the entire set of data configured as columns and rows
- **Views** - Views are used to define what information a user can view in the tables

Network Database Management Model

The popularity of the **network data model** coincided with the popularity of the hierarchical data model.

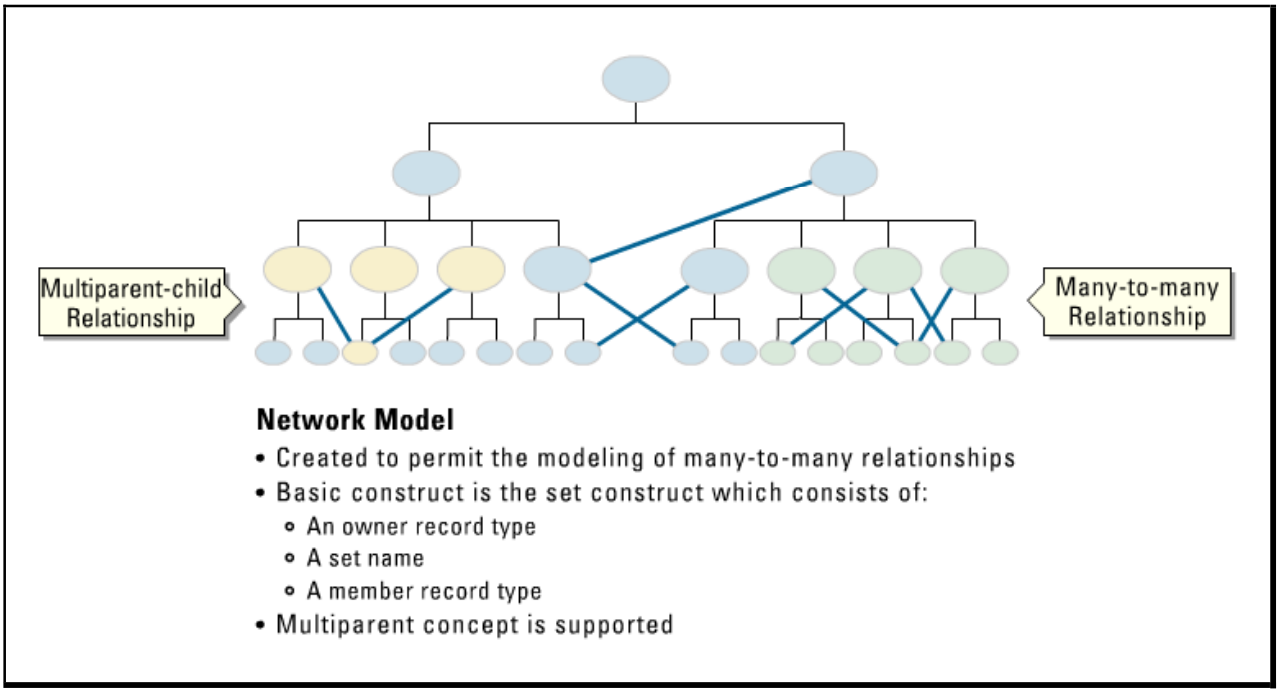

In the hierarchical model, some data sets were more naturally modeled with more than one parent per child. So, the network model was created to permit the modeling of many-to-many relationships in data. In 1971, the Conference on Data Systems Languages (CODASYL) formally defined the network model. The basic data modeling construct in the network model is the set construct.

A set consists of:

- An owner record type
- A set name
- A member record type

A member record type can have that role in more than one set; therefore the multiparent concept is supported. An owner record type can also be a member or an owner in another set. The data model is a simple network, and link and intersection record types (called junction records by IDMS) may exist, as well as sets between them. Thus, the complete network of relationships is represented by several pairwise sets; in each set, some (one) record type is the owner and one or more record types are members.

Object-Oriented Database Model

The **object-oriented database model** adds database functionality to object programming languages. This module brings much more than persistent storage of programming language objects.

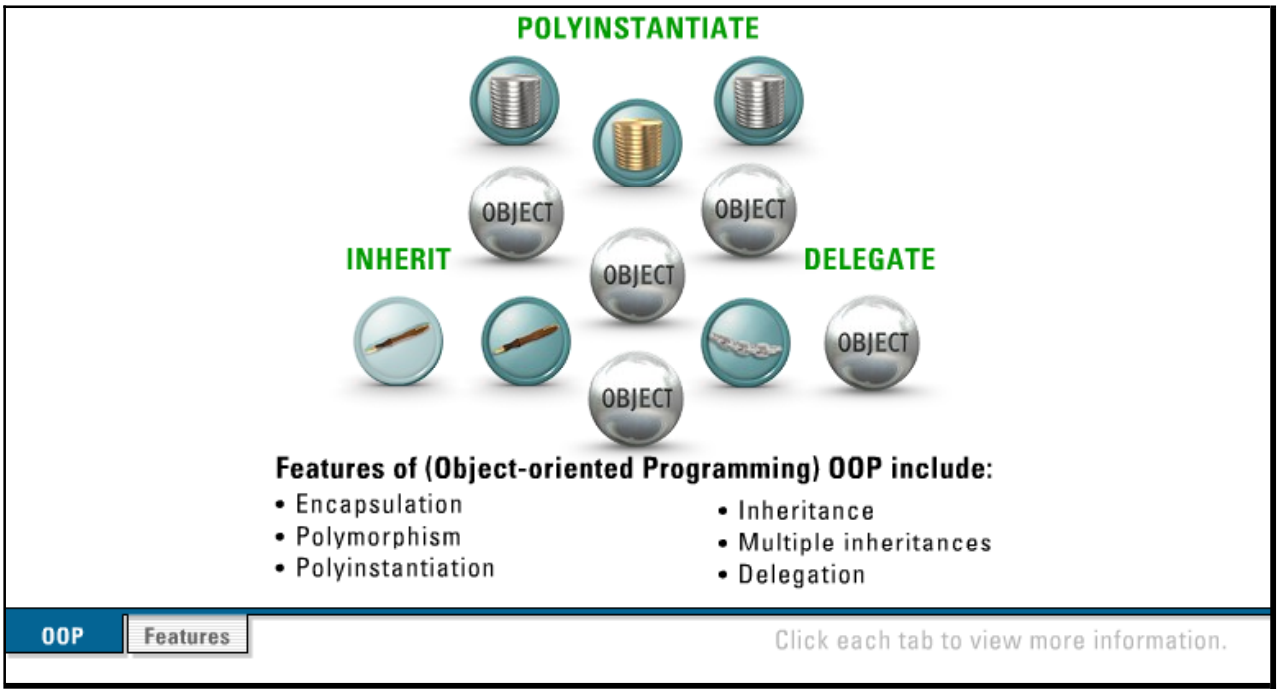

An object DBMS extends the semantics of the C++, Smalltalk, and Java object programming languages to provide full-featured database programming capability, while retaining native language compatibility. A major benefit of this approach is the unification of the application and database development into a seamless data model and language environment. As a result, applications require less code and use more natural data modeling, and code bases are easier to maintain. Object developers can write complete database applications with a modest amount of additional effort.

In contrast to a relational DBMS where a complex data structure must be flattened out to fit into tables or joined together from those tables to form the in-memory structure, an object DBMS has no performance overhead to store or retrieve a web or hierarchy of interrelated objects. This one-to-one mapping of object programming language objects to database objects has two benefits over other storage approaches: it provides higher performance management of objects, and it enables better management of the complex interrelationships between objects. This makes an object DBMS better suited to support applications such as financial portfolio risk analysis systems, telecommunications service applications, World Wide Web document structures, design and manufacturing systems, and hospital patient record systems, which all have complex data relationships.

Features of object oriented programming (OOP) include:

- **Encapsulation** - Hides internal data and operations
- **Polymorphism** - Makes copies of objects and makes changes to those copies
- **Polyinstantiation** - Creates multiple distinct differences between data within objects to discourage lower-level subjects from learning information at a higher level of security

- **Inheritance** - Shares properties and attributes
- **Multiple inheritances** - Is the situation where a class inherits the behavioral characteristics of more than one parent class
- **Delegation** - Is the forwarding of a request by an object to another object or delegate; this forwarding is necessitated by the fact that the object receiving the request does not have a method to service the request

Open Database Connectivity (ODBC)

Open Database Connectivity (ODBC) is a standard database access method developed by the SQL Access group in 1992.

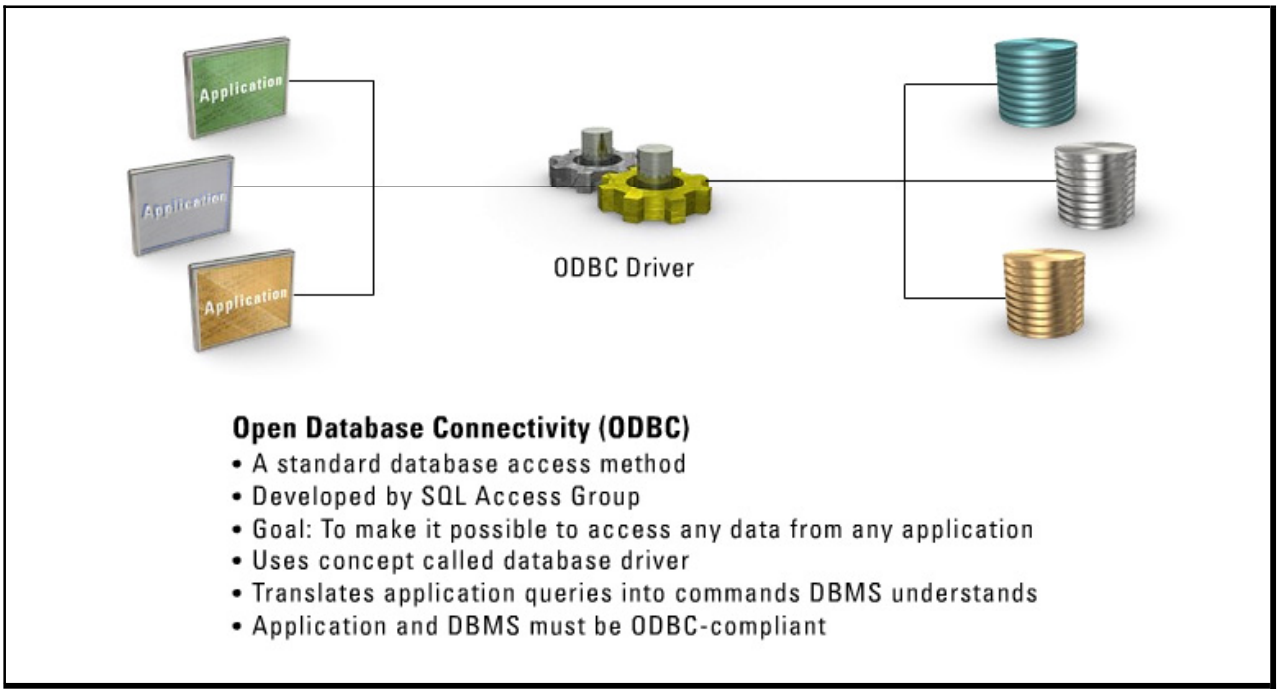

The goal of ODBC is to make it possible to access any data from any application, regardless of which DBMS is handling the data. ODBC manages this by inserting a middle layer, called a database driver, between an application and the DBMS. The purpose of this layer is to translate the application's data queries into commands that the DBMS understands. For this to work, both the application and the DBMS must be ODBC-compliant – that is, the application must be capable of issuing ODBC commands and the DBMS must be capable of responding to them. Since version 2.0, the standard supports SAG SQL.

Object Linking and Embedding (OLE)

Object Linking and Embedding (OLE) is a Microsoft technology that allows you to link elements from different applications within each other. For example, you can embed an Excel spreadsheet or chart inside a PowerPoint presentation. When you change the spreadsheet or chart, it changes inside the PowerPoint

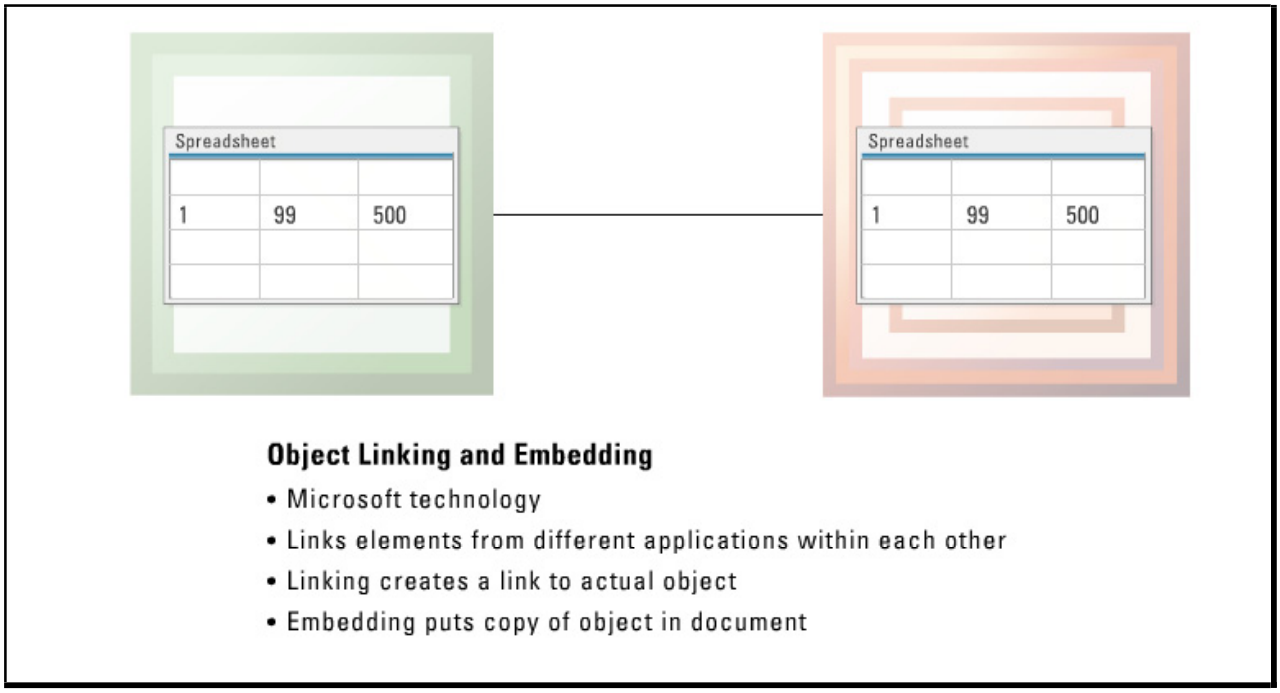

presentation.

OLE allows objects to be linked to and embedded in other documents. Linking creates a link to the actual object; embedding puts a copy of the object into the document. You can usually access the program an object was created with in order to edit the linked or embedded object just by clicking the object. This capability is much more advanced than just taking a screenshot of the data you want and pasting it into another program as a graphic that has no relation to the original data.

Extensible Markup Language (XML)

Extensible Markup Language (XML) is a flexible way to create standard information formats and share both the format and the data on the World Wide Web.

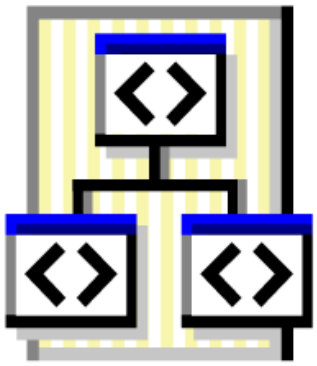

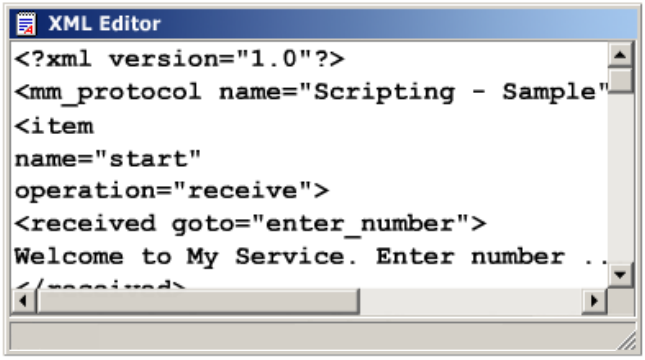

Use scroll bar to view all text.

Extensible Markup Language

- A flexible way to create standard information formats
- Developed by the W3C
- XML documents made up of storage units called entities
- Markup encodes a description of layout and structure
- Requires software module called XML processor

The W3C developed the XML specification. XML is a pared-down version of Standard Generalized Markup Language (SGML), which was designed especially for Web documents. It allows designers to create their own customized tags, enabling the definition, transmission, validation, and interpretation of data between applications and between organizations.

For example, the House of Representatives has recently issued a list of XML tags to be used for Web forms on Member Web sites and other Web sites that send e-mail to congressional offices. The purpose of these forms is to enable correspondence management systems (CMS) and other software to easily identify and process certain types of information – such as name, city, state, zip code, issue, etc. – which will help make the software more efficient and more effective.

XML documents are made up of storage units called entities, which contain either parsed or unparsed data. Parsed data is made up of characters, some of which form character data, and some of which form markup. Markup encodes a description of the document's storage layout and logical structure. XML provides a mechanism to impose constraints on the storage layout and logical structure.

A software module called an XML processor is used to read XML documents and provide access to their content and structure.

Java Database Connectivity (JDBC)

Java Database Connectivity (JDBC) is a standardized DB interface for Java. This technology allows you to write an application once, and use it with any SQL database that has a JDBC driver.

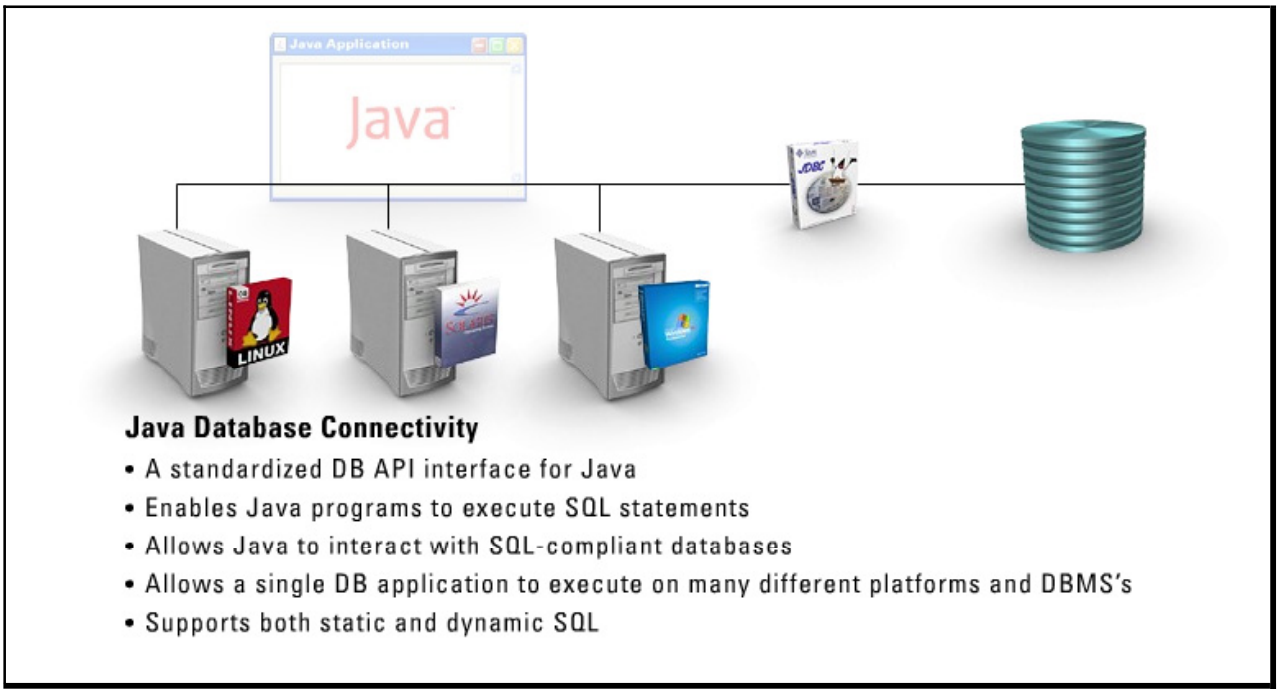

JDBC is a Java application programming interface (API) that enables Java programs to execute SQL statements. This capability allows Java programs to interact with any SQL-compliant database. Since nearly all relational DBMSs support SQL, and because Java itself runs on most platforms, JDBC makes it possible to write a single database application that can run on different platforms and interact with different DBMSs.

JDBC is a Java API (a standard set of Java classes) that provides for vendor-independent access to relational data. The JDBC classes provide standard features such as simultaneous connections to several databases, transaction management, simple queries, manipulation of pre-compiled statements with bind variables, calls to stored procedures, streaming access to long column data, access to the database dictionary, and descriptions of cursors. JDBC supports both static and dynamic SQL.

ActiveX Data Objects (ADO)

ActiveX Data Objects (ADO) is an API from Microsoft that lets a programmer writing Windows applications get access to a relational or non-relational database from both Microsoft and other database providers.

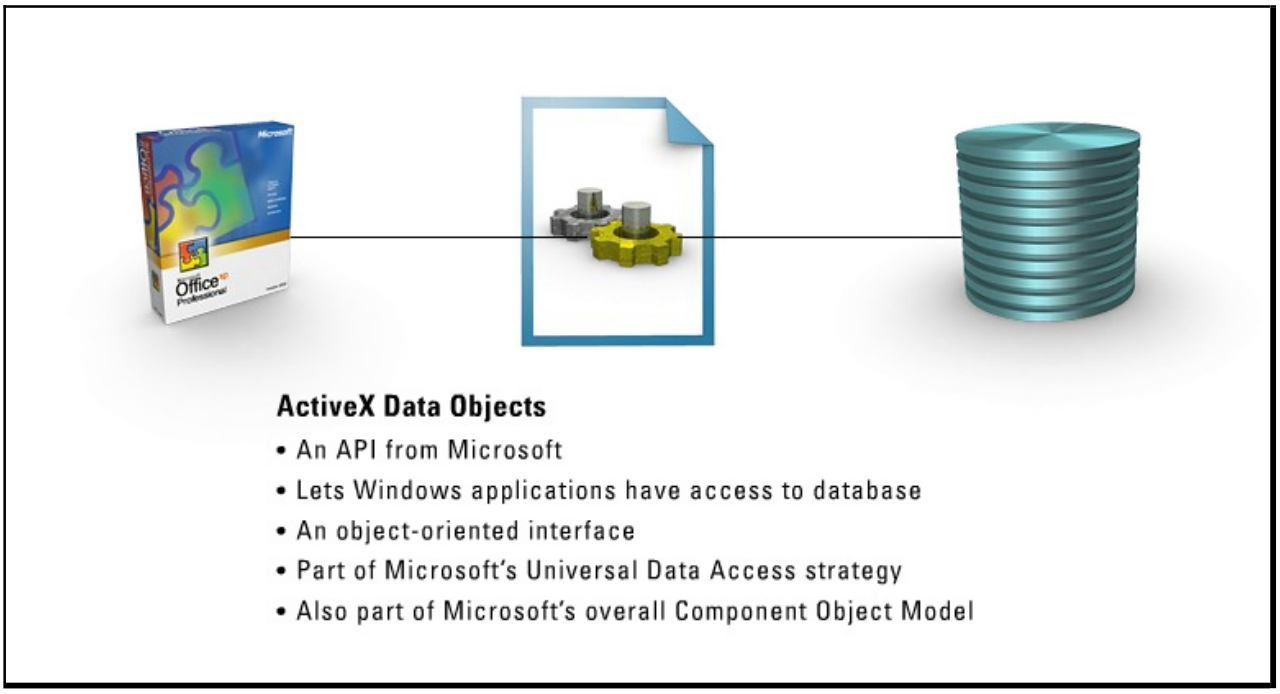

For example, if you want to write a program that would provide users of your Web site with data from an IBM DB2 database or an Oracle database, you could include ADO program statements in a Hypertext Markup Language (HTML) file that you then identify as an Active Server Page. Then, when a user requests the page from the Web site, the page sent back would include appropriate data from a database, which would be obtained using ADO code.

Like Microsoft's other system interfaces, ADO is an object-oriented programming interface. It is also part of an overall data access strategy from Microsoft called Universal Data Access. Microsoft believes that rather than trying to build a universal database as IBM and Oracle have suggested, finding a way to provide universal access to various kinds of existing and future databases is a more practical solution. In order for this to work, Microsoft and other database companies would need to provide a "bridge" program between the database and Microsoft's OLE DB, the low-level interface to databases.

OLE DB is the underlying system service that a programmer using ADO is actually using. A feature of ADO, Remote Data Service, supports "data-aware" ActiveX control in Web pages and efficient client-side cache. As part of ActiveX, ADO is also part of Microsoft's overall Component Object Model (COM), its component-oriented framework for putting programs together.

Online Transaction Processing (OLTP)

Databases must often allow the real-time processing of SQL transactions to support e-commerce and other time-critical applications. This type of processing is known as **online transaction processing (OLTP)**.

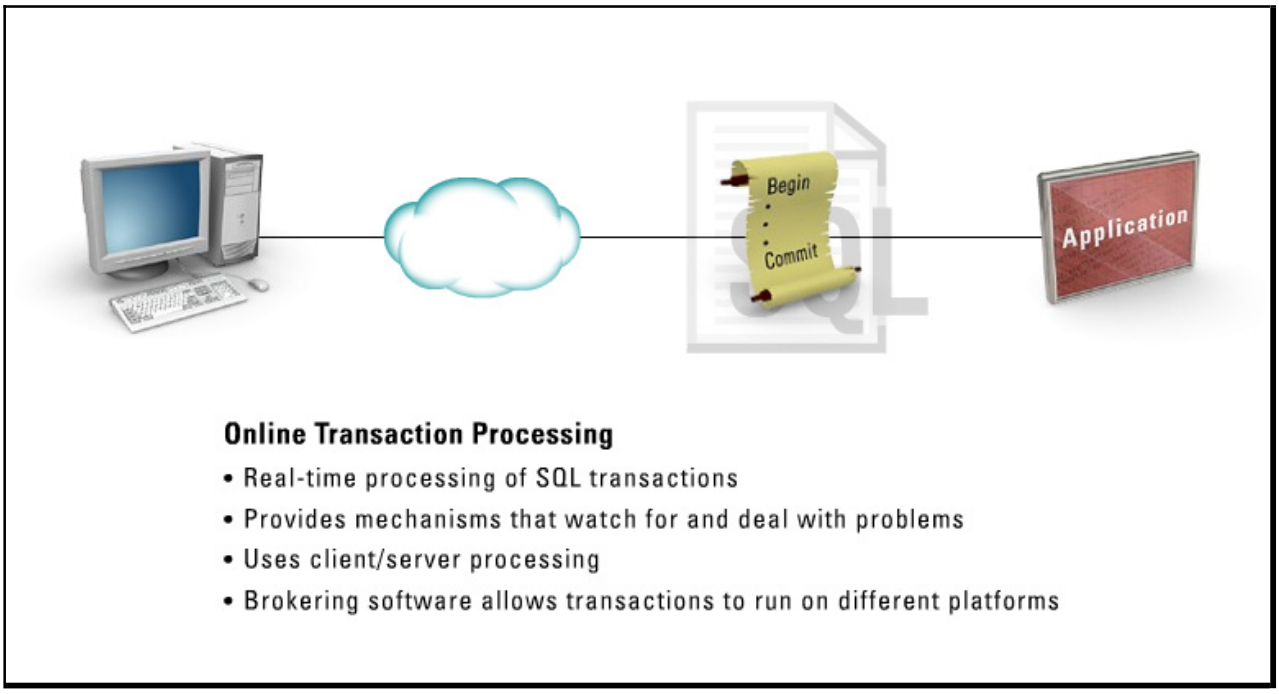

OLTP is a common term used to describe any form of transaction processing involving communication devices and data processing environments. OLTP provides mechanisms that watch for problems and deal with them appropriately when they do occur.

Today's online transaction processing increasingly requires support for transactions that span a network and may include more than one company. For this reason, new OLTP software uses client/server processing and brokering software that allows transactions to run on different computer platforms in a network.

Data Mining

Data mining is the process of discovering information in data warehouses by running queries on the data.

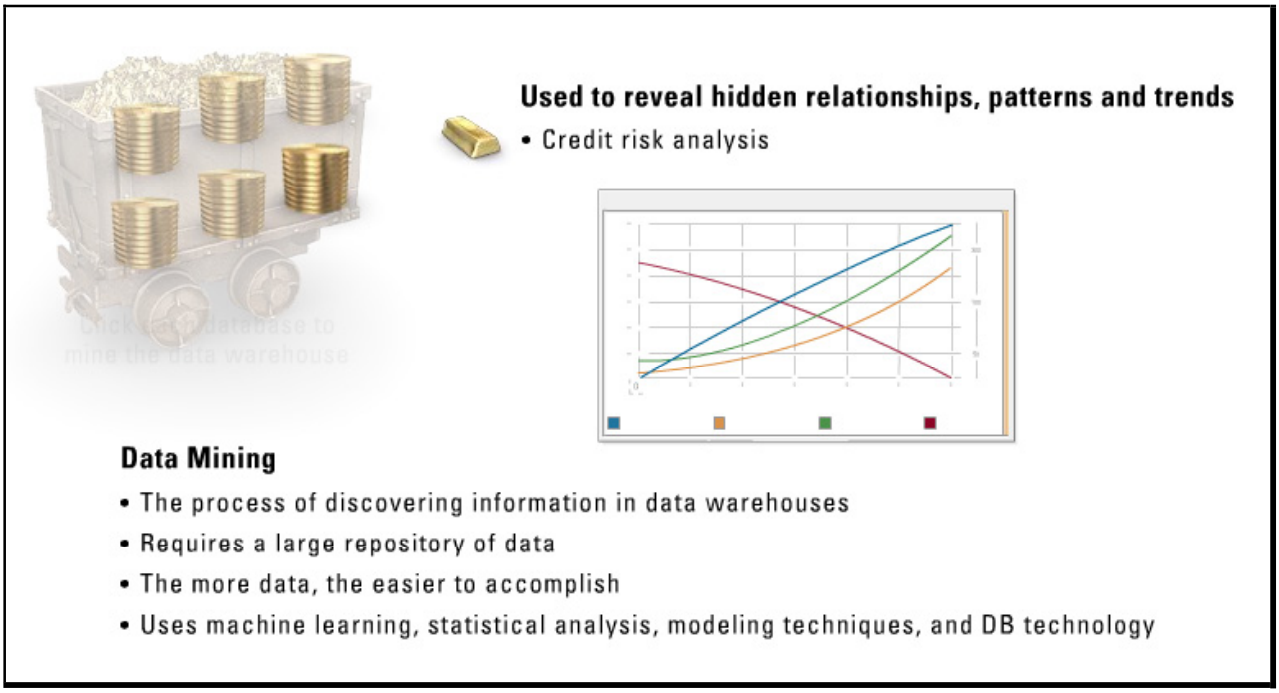

To perform data mining, a large repository of data is required because data mining is used to reveal hidden relationships, patterns, and trends in the data warehouse. The more data, the easier it will be to accomplish these tasks.

Note	Many people consider data mining a decision-making technique that is based on a series of analytical techniques taken from the fields of mathematics, statistics, cybernetics, and genetics.

Using a combination of machine learning, statistical analysis, modeling techniques, and database technology, data mining finds patterns and subtle relationships in data and infers rules that allow the prediction of future results. Typical applications include market segmentation, customer profiling, fraud detection, evaluation of retail promotions, and credit risk analysis.

Inference

Inference is defined as the reasoning involved in drawing a conclusion or making a logical judgment on the basis of circumstantial evidence and prior conclusions rather than on the basis of direct observation.

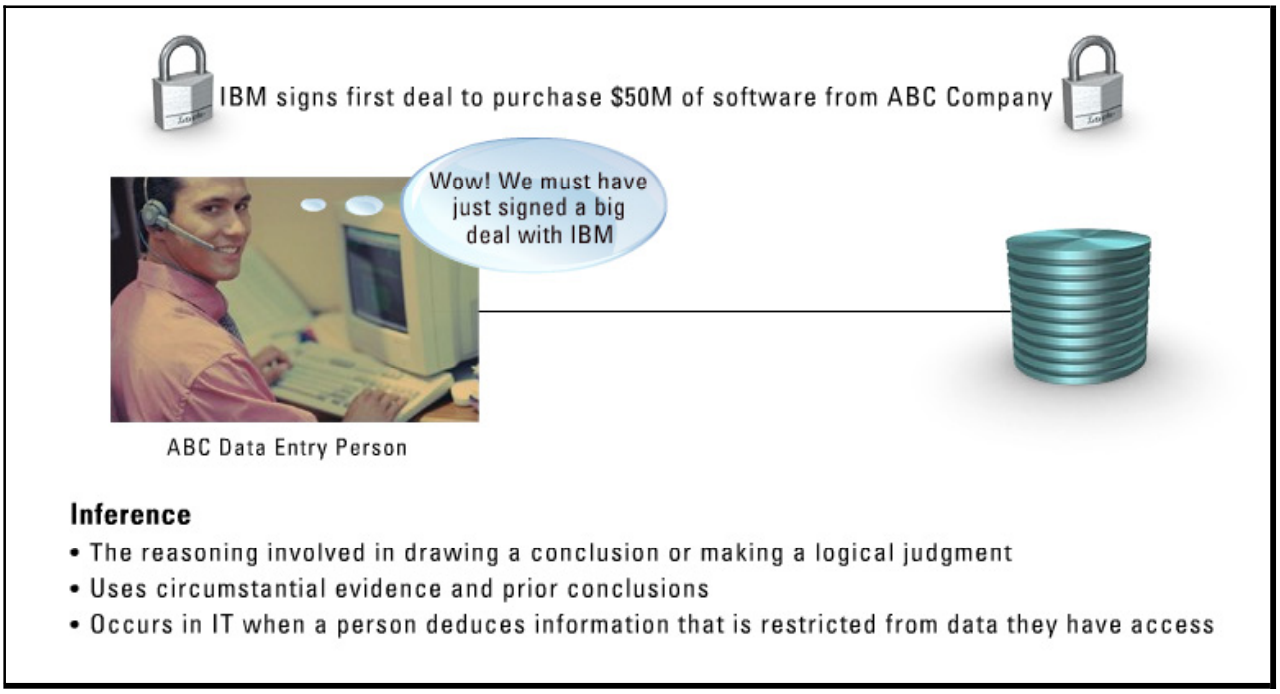

In the information technologies realm, inference occurs when a person deduces information that is restricted from data to which he or she has access. You can see this event occur when data at a lower security level indirectly portrays data at a higher level.

Polyinstantiation

Polyinstantiation in the database world is referred to as an environment characterized by information stored in more than one location in the database.

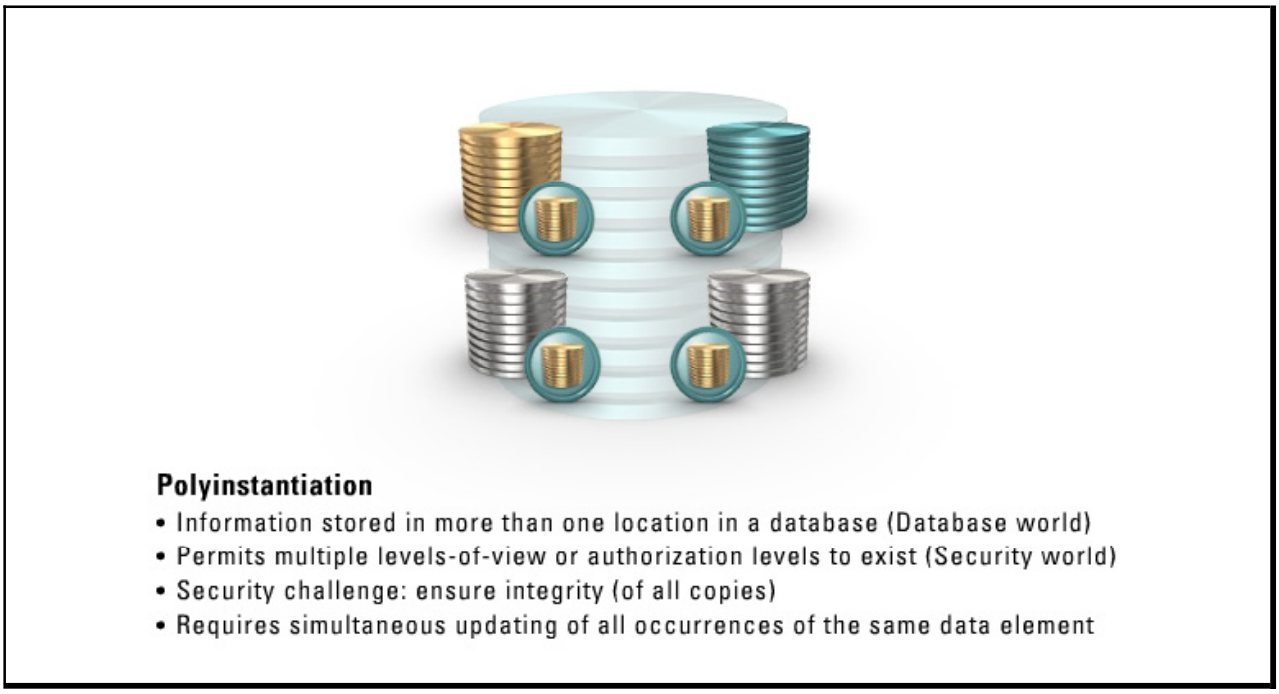

Within a security context, polyinstantiation would permit multiple levels-of-view and authorization levels to exist. The security challenge here is to ensure the integrity (of all copies) of the information in the database. This challenge would require a method of simultaneously updating all occurrences of the same data element, otherwise, integrity cannot be guaranteed.

Summary

The key points discussed in this lesson are:

- A DBMS is a suite of application programs that typically manage large, structured sets of persistent corporate data.
- The hierarchical data model combines related records and fields in a logical tree structure. The records can have one child, many children, or no children.
- The relational database model uses attributes (columns) and tuples (rows) to contain and organize information. A primary key is a field that links all the data within a record to a corresponding value, while a foreign key is a field that represents a reference to an entry in some other table.
- In the distributed data model, data are stored in more than one database, but are logically connected.
- SQL is one of the standardized languages used in the relational model.
- The basic data modeling construct in the network model is the set construct.
- The object-oriented database model adds database functionality to object programming languages.
- The goal of ODBC is to make it possible to access any data from any application, regardless of which DBMS is handling the data.
- OLE is a Microsoft technology that allows you to link elements from different applications within each other.
- XML is a flexible way to create standard information formats and share both the format and the data on the World Wide Web.
- JDBC is a standardized DB interface for Java. This technology allows you to write an application once, and use it with any SQL database that has a JDBC driver.
- ADO is an API from Microsoft that lets a programmer writing Windows applications get access to a relational or non-relational database from both Microsoft and other database providers.
- OLTP is a common term used to describe any form of transaction processing involving communication devices and data processing environments.
- Data mining is the process of discovering information in data warehouses by running queries on the data.
- In the information technologies realm, inference occurs when a person deduces information that is restricted from data to which he or she has access.
- Polyinstantiation in the database world is referred to as an environment characterized by information stored in more than one location in the database.

Knowledge-Based Systems

Overview

In today's world where information abounds everywhere, it is becoming increasingly difficult to sort or mine through the vast amounts of data to obtain the information needed. Knowledge-based systems use artificial intelligence (AI) to solve this dilemma and many other types of problems. This lesson will discuss what knowledge-based systems are and how they work.

Importance

Understanding how technology is evolving will allow information security professionals to gain a better understanding of data and CPU logic, which will allow them to better arm themselves and their enterprises against future attacks.

Objectives

Upon completing this lesson, you will be able to:

- Describe expert systems
- Describe neural networks

Outline

The lesson contains these topics:

- Expert Systems
- Neural Networks

Expert Systems

Expert systems are artificial intelligence systems used to emulate human knowledge in order to solve problems.

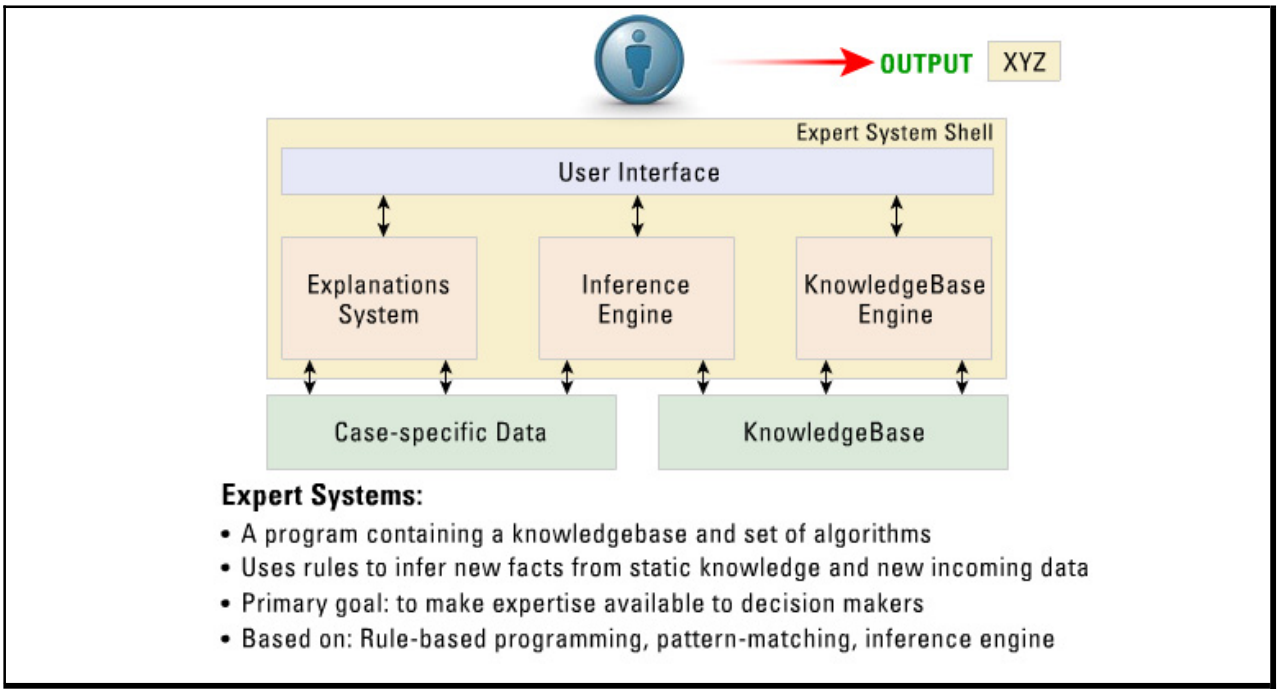

An expert system is a computer program containing a knowledge base and set of algorithms and rules used to infer new facts from static knowledge and new incoming data.

The primary goal of expert systems is to make expertise available to decision makers and technicians who need answers quickly. The right amount of expertise is not always available at the right place and the right time. Portable computers loaded with in-depth knowledge of specific subjects can bring a decade's worth of knowledge to a problem. The same systems can assist supervisors and managers with situation assessment and long-range planning.

These knowledge-based applications of artificial intelligence have enhanced productivity in business, science, engineering, and the military. Each new deployment of an expert system yields valuable data regarding what works in what context, thus fueling the AI research that provides even better applications.

Expert systems are based on:

- **Rule-Based Programming** - A common way of developing expert systems
- **Pattern Matching** - Based on if-then logic units
- **Inference Engine** - A mechanism that automatically matches facts against patterns and determines which rules are applicable

Neural Networks

A **neural network** is an electronic model based on the neural structure of the brain. It attempts to replicate the basic functions of neurons and their circuitry to solve problems in a new way.

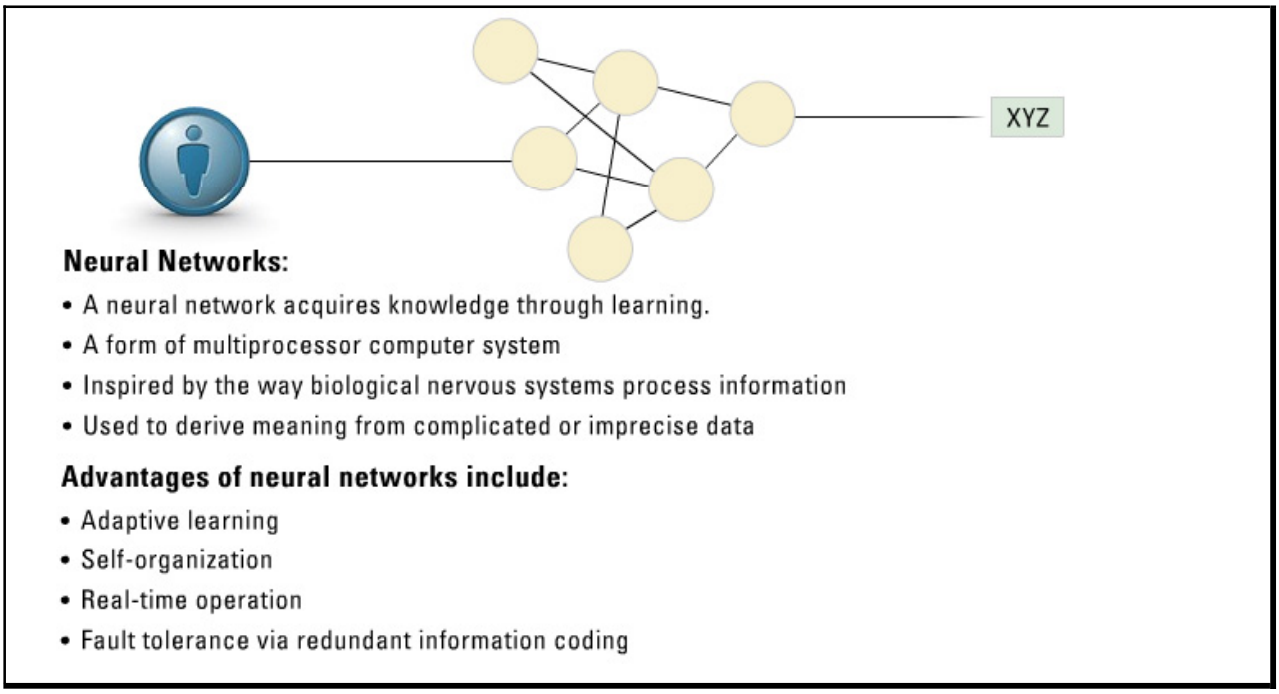

Neural networks are a form of multiprocessor computer systems, with:

- Simple processing elements
- A high degree of interconnection
- Simple scalar messages
- Adaptive interaction between elements

An **Artificial Neural Network** (ANN) is an information-processing paradigm that is inspired by the way biological nervous systems, such as the brain, process information. The key element of this paradigm is the novel structure of the information processing system. It is composed of a large number of highly interconnected processing elements (neurons) working in unison to solve specific problems. ANNs, like people, learn by example. An ANN is configured for a specific application, such as pattern recognition or data classification, through a learning process. Learning in biological systems involves adjustments to the synaptic connections that exist between the neurons. This is true of ANNs as well.

Neural networks, with their remarkable ability to derive meaning from complicated or imprecise data, can be used to extract patterns and detect trends that are too complex to be noticed by either humans or other computing techniques. A trained neural network can be thought of as an "expert" in the category of information it has been given to analyze. This expert can then be used to provide projections given new situations of interest and answer "what if" type questions. Other advantages include:

- **Adaptive Learning**: An ability to learn how to do tasks based on the data given for training or initial experience

- **Self-Organization**: An ANN can create its own organization or representation of the information it receives during learning time
- **Real Time Operation**: ANN computations may be carried out in parallel; special hardware devices that take advantage of this capability are being designed and manufactured
- **Fault Tolerance via Redundant Information Coding**: Partial destruction of a network leads to the corresponding degradation of performance; however, some network capabilities may be retained even with major network damage

Neural networks take a different approach to problem solving than that of conventional computers. Conventional computers use an algorithmic approach (i.e., the computer follows a set of instructions in order to solve a problem). Unless the specific steps that the computer needs to follow are known the computer cannot solve the problem. That restricts the problem solving capability of conventional computers to problems that we already understand and know how to solve. But computers would be so much more useful if they could do things that we don't exactly know how to do.

Neural networks process information in a similar way the human brain does. The network is composed of a large number of highly interconnected processing elements (neurons) working in parallel to solve a specific problem. Neural networks learn by example. They cannot be programmed to perform a specific task. The examples must be selected carefully otherwise useful time is wasted or even worse the network might be functioning incorrectly. The disadvantage is that because the network finds out how to solve the problem by itself, its operation can be unpredictable.

On the other hand, conventional computers use a cognitive approach to problem solving; the way the problem is to be solved must be known and stated in small unambiguous instructions. These instructions are then converted to a high-level language program and then into machine code that the computer can understand. These machines are totally predictable; problems are due to a software or hardware fault.

Neural networks and conventional algorithmic computers are not in competition but rather complement each other. There are tasks more suited to an algorithmic approach like arithmetic operations and tasks that are more suited to neural networks. Even more, a large number of tasks require systems that use a combination of the two approaches in order to perform at maximum efficiency. (Normally, a conventional computer is used to supervise the neural network.)

Summary

The key points discussed in this lesson are:

- Expert systems are artificial intelligence systems used to emulate human knowledge in order to solve problems.
- A neural network is an electronic model based on the neural structure of the brain. It attempts to replicate the basic functions of neurons and their circuitry to solve problems in a new way.

Systems Development Life Cycle

Overview

You can best protect a system in development if you plan and manage it with security in mind during its entire life cycle. Security risks and security-related events eventually begin appearing during development and having a sound model to follow will allow you to deal with these issues in a far more efficient and secure manner. This lesson will discuss the various states a system in development goes through in its lifetime.

Importance

Understanding the systems development life cycle will allow the information security professional to provide insight and better security to systems being developed with or without security in the forefront.

Objectives

Upon completing this lesson, you will be able to:

- Identify guidelines for designing controls
- Define the three primary approaches for software development
- List the phases of development
- Describe the project initiation phase
- Describe the functional requirements phase
- Describe the system design specifications phase
- Describe the build/development phase
- Describe the acceptance phase
- Describe the testing and evaluation phase
- Describe the certification and accreditation processes
- Describe the installation phase
- Describe the post-installation phase

- Describe the revisions and replacement phase

Outline

The lesson contains these topics:

- Systems Development Controls
- Requirements Determination
- Project Specifications Development
- Project Initiation
- Functional Requirements Definition
- System Design Specifications
- Build/Development
- Acceptance
- Testing and Evaluation Controls
- Certification and Accreditation
- Installation
- Post Installation
- Revisions and Replacement

Systems Development Controls

You can break **systems development controls** into two distinct controls, which are designing controls and controlling the development process.

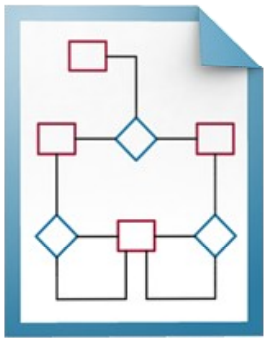

Systems Development

Two distinct controls:

- Designing controls- building suitable applications
- Controlling the development process- impose control and discipline onto the process

Systems analyst needs to consider:

- Testing
- File conversion
- Reconciliations
- Controls over amendments
- Training of users and documentation

Systems Development Controls

When designing controls, the systems analyst should build suitable operational (applications) controls into the system. The systems analyst should meet the following guidelines:

- All data due for processing should be processed
- Situations with potential for errors should be minimized
- Errors should be detected, located, and corrected as early as possible
- Controls should not interrupt the flow of data through the system
- Controls must meet cost-benefit criteria
- Controls should be part of the organization's overall security strategy

When controlling the development process, some approaches to development (methodologies) impose control and discipline onto the process (signing off each stage).

The systems analyst should consider these matters:

- Testing
  - Program testing
  - System testing
  - Acceptance testing
- File conversion
  - Planning
  - Follow-up of errors

    - Checking back to old files
- Reconciliations
- Controls over amendments
- Training of users and documentation

Requirements Determination

Several System Development Methods (SDMs) have evolved to satisfy the different requirements for software programming.

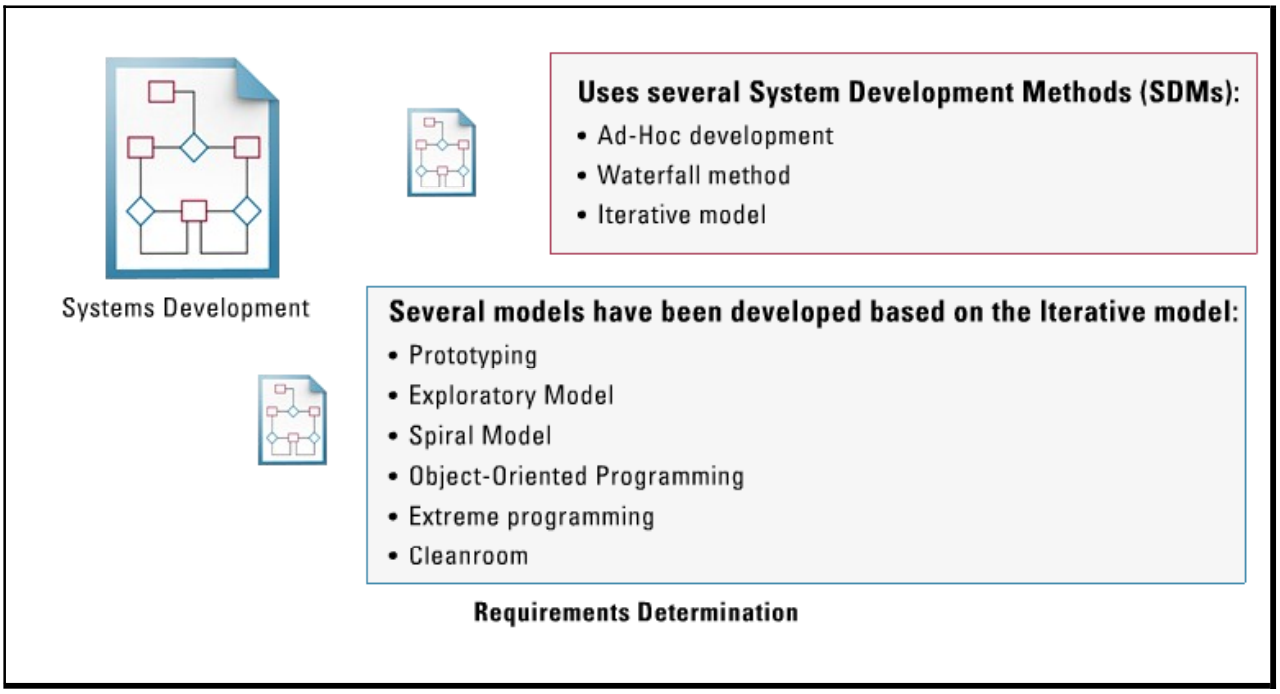

The term "SDM" is used to describe the various models software developers use to provide guidance during the analysis, design, development, and maintenance of information systems. There are three primary approaches for software development:

- Ad-Hoc development
- Waterfall method
- Iterative model

When development takes place in a chaotic, an unpredictable or a haphazard manner, it is called Ad-Hoc development. **Ad-Hoc development** relies entirely on the skills of the individual writing the software. Normally, schedules, budgets, and software functionality are inconsistent in that performance can be predicted only by the individual rather than organizational capability.

The **Waterfall method** was developed in the early 1970s and provided a sense of order to the systems development process. In this model, each phase contains a list of activities that must be performed and documented before the next phase begins. A distinct disadvantage of the Waterfall method is that it does not always scale well for large and complex projects. Also, because each phase is dependent upon the prior phase being completed, it can inhibit a development team from pursuing concurrent phases or activities. Because of this factor the Waterfall method is not good for projects that must be developed in quick turnaround time periods (less than six months).

The Waterfall method created a demand for new models of software development that provided for greater flexibility. The **iterative model** fragments a project into smaller components and each component follows the regular Waterfall model. The iterative model thus allows for successive refinements of requirements, design, and coding. One distinct advantage of the iterative model is that system users can

provide feedback earlier in the process, which brings up a challenge unto itself; how to ensure involvement from the user community. Several models have been developed from the iterative model:

- **Prototyping** - Build a simplified version (prototype) of the application, release it for review, and use the feedback from the users to build a second version, etc.
    - **Rapid Application Development (RAD)** - A form of prototyping that requires strict time limits on each phase and requires application tools that enable quick development.
    - **Joint Analysis Development (JAD)** - A management process that helps developers work directly with users to develop a working application; JAD facilitates the gathering of a team of users, expert system developers, and technical experts.
    - **Modified Prototype Model (MPM)** - MPM allows for the basic functionality of a desired system or component to be formally deployed in a quick time frame; it is ideal for Web application development.
- **Exploratory Model** - In this model, a set of requirements is built with what is currently available. The distinguishing characteristic of this model is the lack of precise specifications.
- **Spiral Model** - A combination of both the Waterfall model and the prototyping model with a new component – risk assessment.
- **Object-Oriented Programming (OOP)** - Programming is organized around objects rather than actions. OOP is a programming method that makes a self-sufficient object, which is a block of pre-assembled code in a self-contained module.
    - **Reuse Model** - In this model, an application is built from existing components.
    - **Component-Based Development** - This model uses standardized building blocks to assemble, rather than develop, an application.
- **Extreme Programming** - A discipline of software development based on values of simplicity, communication, feedback, and courage. Its goal is to bring the entire development team together and provide enough feedback to determine if the team can develop software for the situation.
- **Cleanroom** - Named after the process of cleaning the waste from the wafer after it has been made. The development of high-quality software where a method of controlling defects (bugs) is paramount. Its goal is to write software correctly the first time (defect prevention is its focus).

Protection Specifications Development

After you select the proper software development model, you need to follow several phases of development to ensure proper protection of the system.

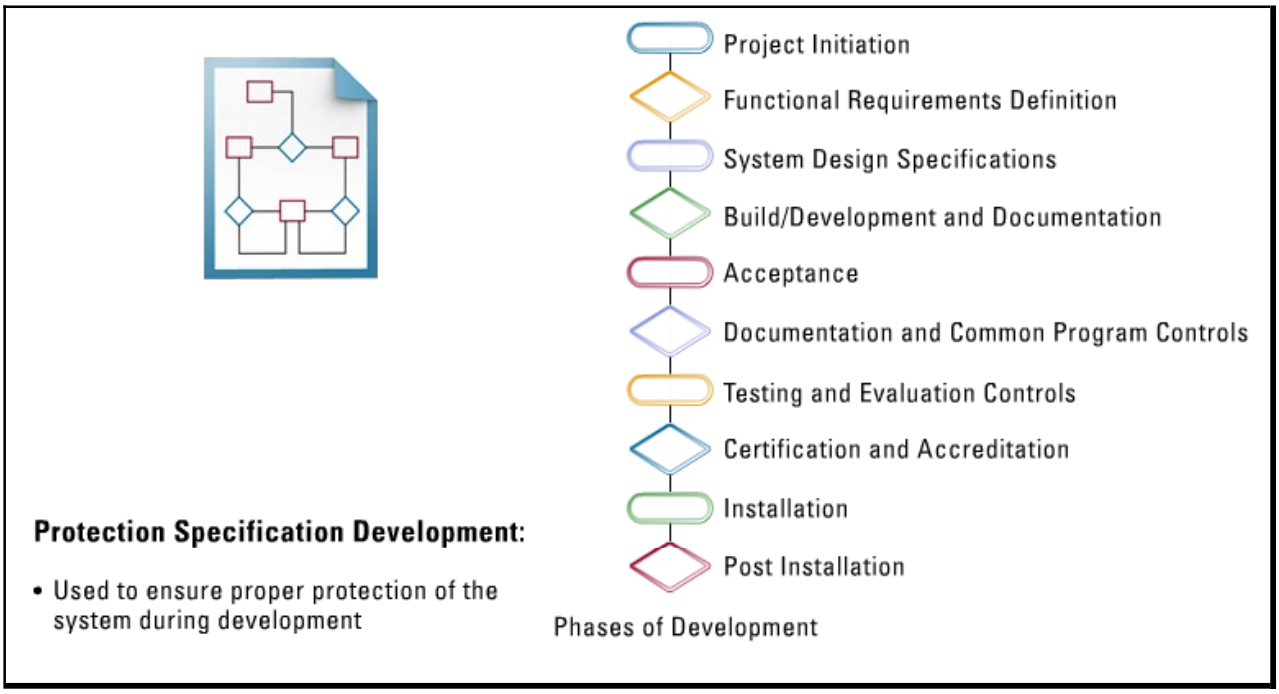

Phases of development:

- **Project Initiation** - Business needs are identified along with a proposed technical solution.
- **Functional Requirements Definition** - A comprehensive analysis of current and possible future functional requirements to ensure that the new system will meet end-user needs.
- **System Design Specifications** - Includes all activities related to designing the system and software. The system architecture, system outputs, and system interface are designed.
- **Build/Development and Documentation -** Source code is generated, test scenarios and test cases are developed, unit and integration testing is conducted, and the program and system are documented for maintenance and turned over to acceptance testing and production.
- **Documentation and Common Program Controls -** The controls used when editing data within a program. Includes the type of logging the program should be doing and how the program versions should be stored.
- **Acceptance** - An independent developer group develops test data and tests code to ensure that the code will function within the proper environment and that it will meet all functional and security requirements.
- **Testing and Evaluation Controls** - Testing all changes with known good data and evaluating all outputs.
- **Certification and Accreditation** - Certification is used to evaluate the security stance of the software or system against a set of security standards. Certification also involves how well the system performs its functional requirements. Accreditation is performed after certification and is the authorization of software to be implemented in a production status, in a specific environment, for a specific period of time.

- **Implementation** - The new system is transitioned from the acceptance phase into the live production environment.
- **Post Installation** - After the system is in general use, monitoring the performance of the system and ensuring continuity of operations are the goals of this phase.

Project Initiation

When a project begins it is important that you identify the business needs (functional requirements) along with any current proposed technical solution. You should identify and document this information in the project outline along with the project objectives, scope, strategies, cost, schedule, and other business-related factors. Management will use this project plan document as the basis for approval of the project.

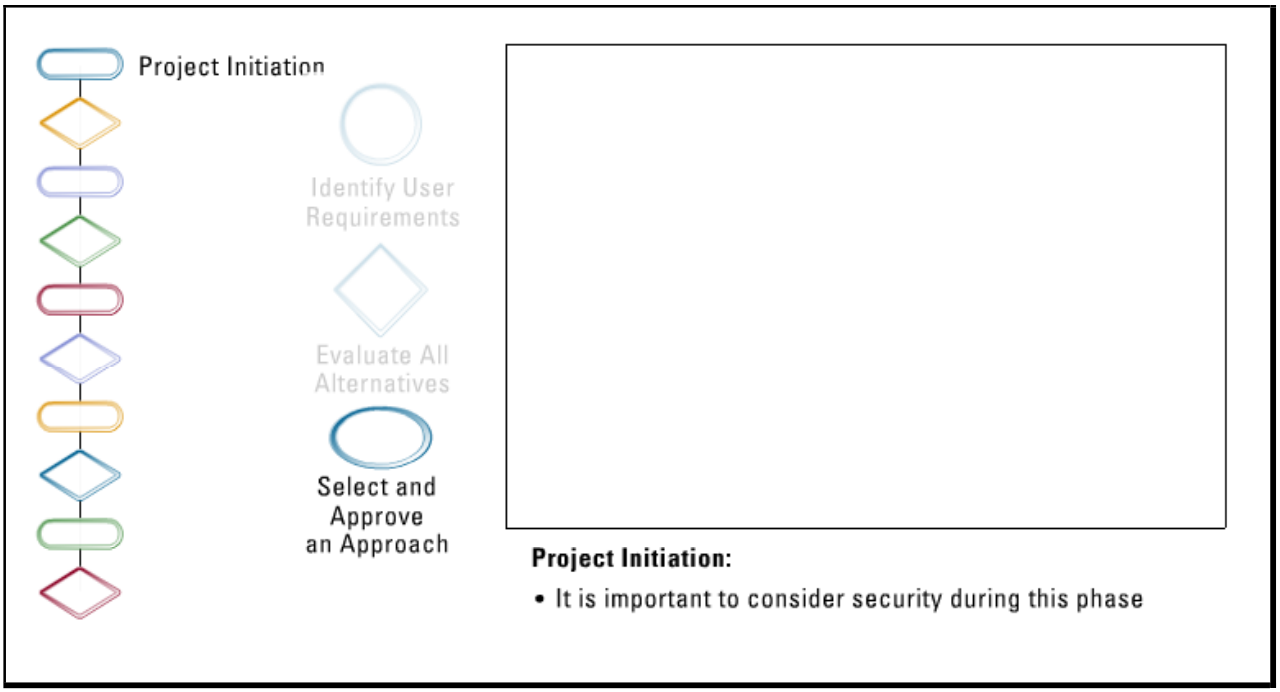

It is important that you consider security during this phase as well as any service level agreements (SLAs) that you need to meet. A SLA is a contract that defines the technical support or business parameters that an application service provider must provide to their clients. Normally, the SLA defines the measures and performance and any consequences for failing to abide to them.

The security checklist during the project initiation phase includes:

- Does the application (or data) have special value or require protection?
- Has the system owner determined the information's value? What are the assigned classifications?
- Will the application potentially risk exposure of sensitive information?
- Will any output (display/printed) require special measures?
- Will data be generated in a public facility? Will controlled areas be required for operation?

Functional Requirements Definition

During the functional requirements phase, the project team should conduct a comprehensive analysis of current and possible future functional requirements to make sure that the new system will meet all end-user requirements.

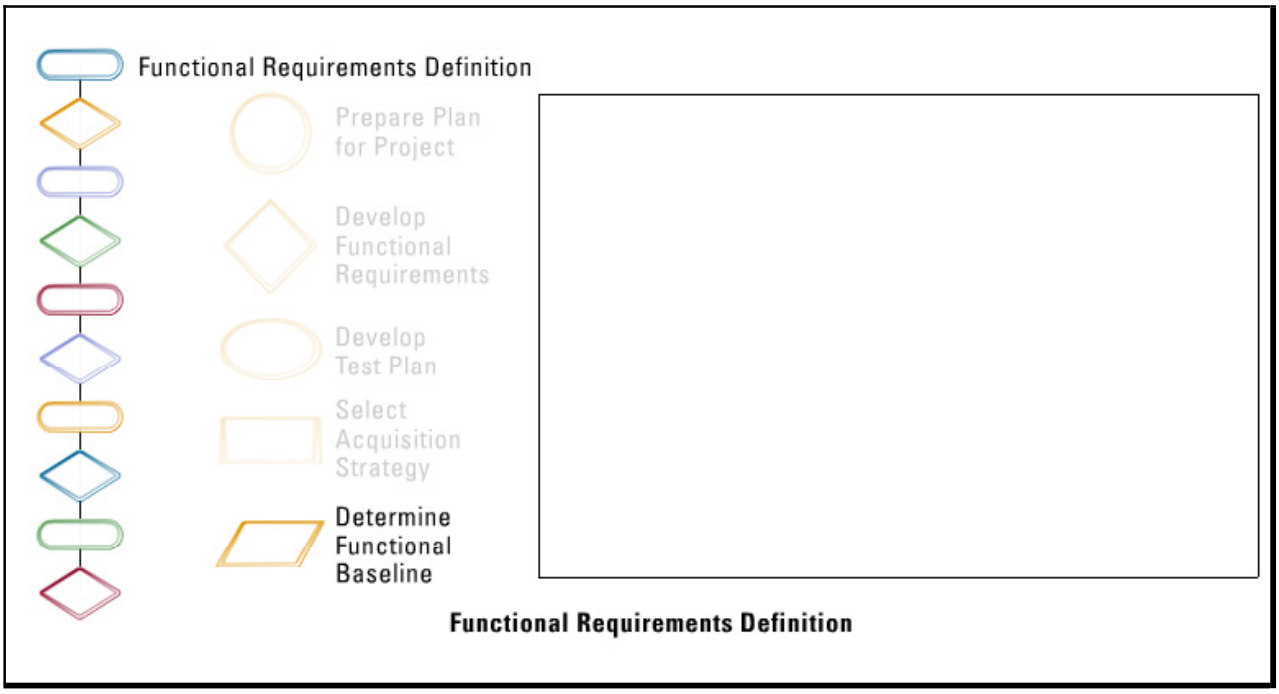

The team must also be sure to review the documents created from the project initiation phase and make any necessary changes to them. Normally, this phase is part of the project initiation phase.

System Design Specifications

The system design specifications phase should include all activities related to designing the entire system and necessary software.

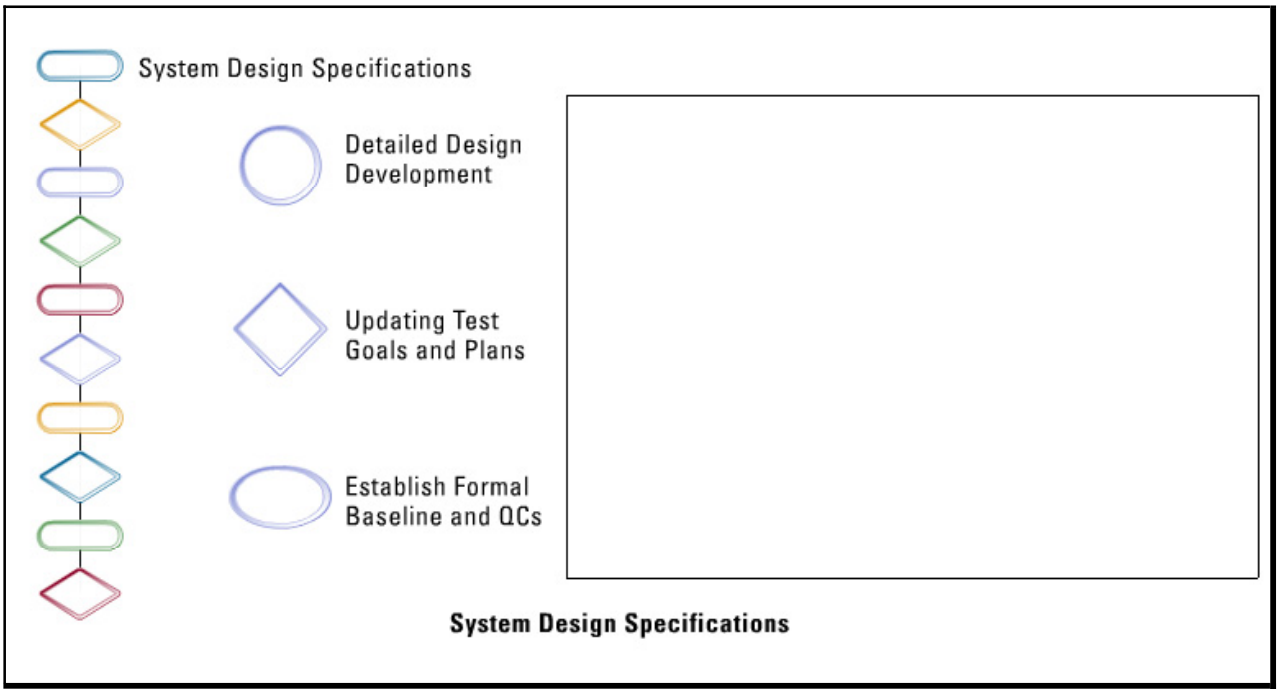

You must design the system architecture, system outputs, and system interfaces per system specifications. You must establish data input, data flow, and data output requirements with designed security features in mind.

Build/Development

During the build/development phase, you generate the program's source code as well as proper test scenarios and test cases.

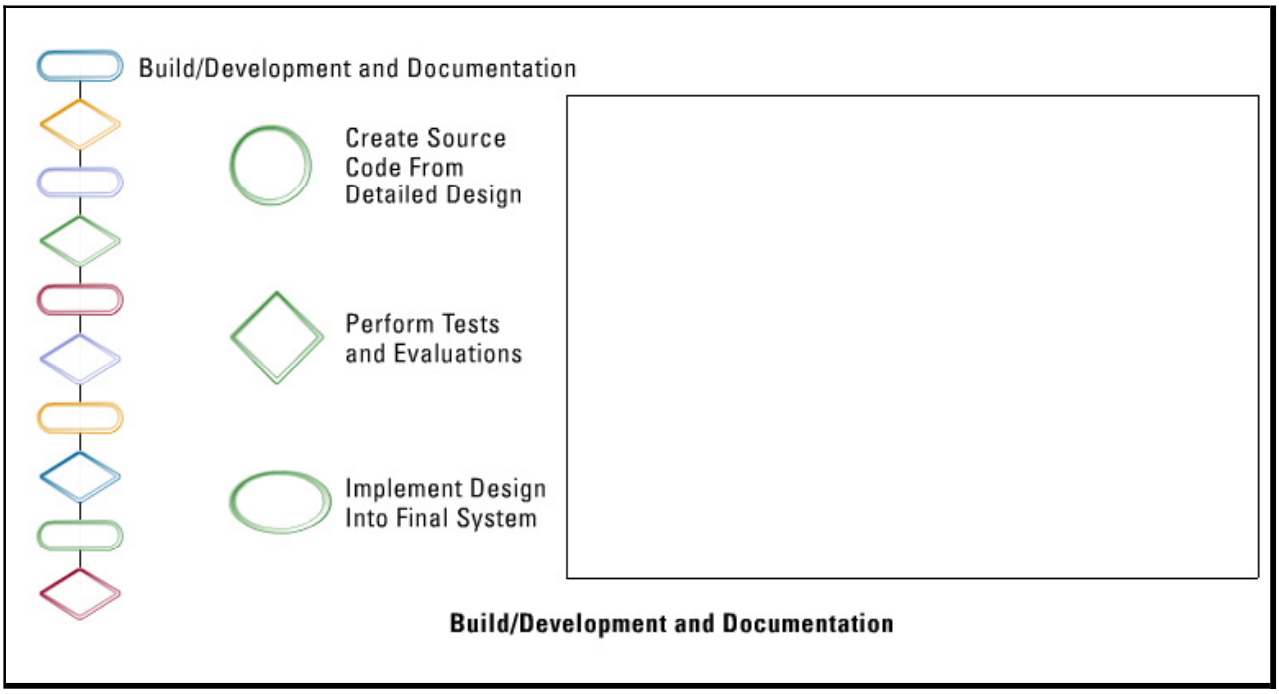

You should conduct unit testing during this phase. Finally, you should document the program and system for maintenance and for turnover to acceptance testing and production during this phase.

Acceptance

The acceptance phase allows for an independent group to develop test data and test the code to ensure that it will function within the organization's environment and within specification, and that it will meet all the functional and security requirements.

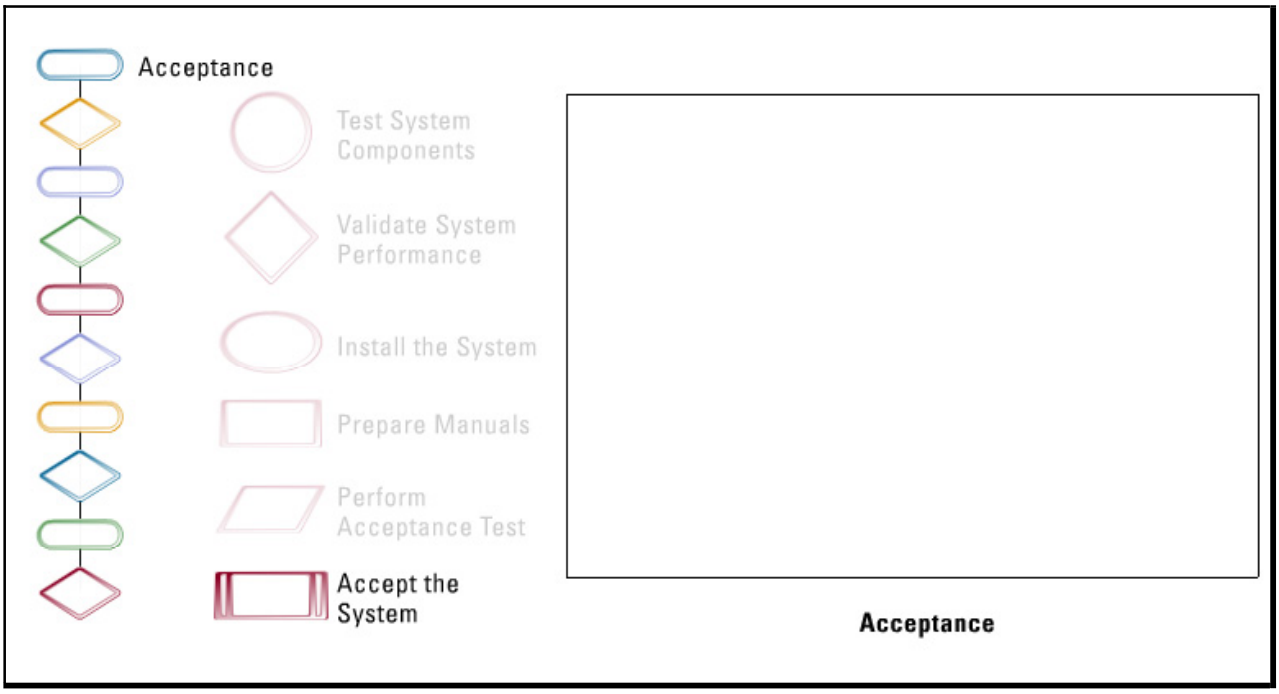

To prevent any separation of duties issues, an independent group must test the code during all applicable stages of development. Remember, the goal of security testing is to ensure that the application meets all security requirements and specifications.

Security testing should be conducted such that all design and implementation flaws that would allow a user to violate the software security policy and requirements are uncovered. Simulating the production environment is essential to ensure test validity. During this phase, you should also create a security certification package as well as any user documentation.

Testing and Evaluation Controls

During the testing and evaluation phase, you can use these guidelines as appropriate to your environment.

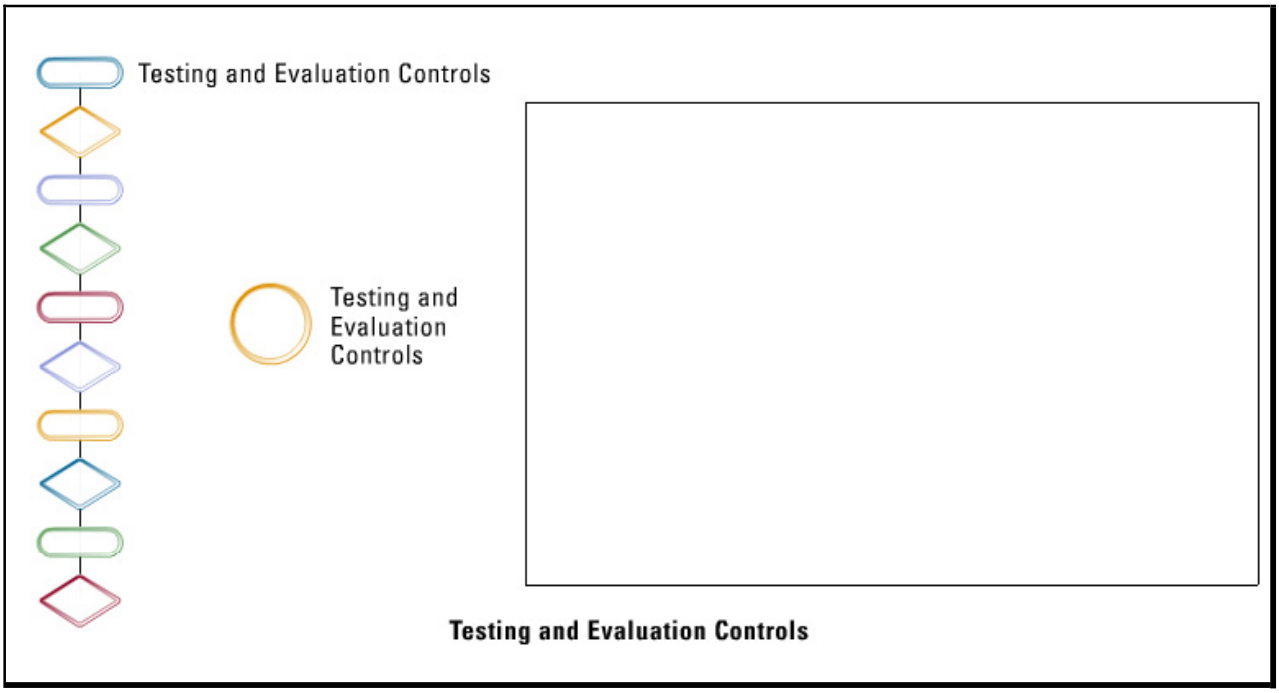

Guidelines:

- Test data should include data at both ends of the acceptable data ranges as well as data beyond the expected and allowable data points.
- Always test with known good data.
- Always validate data before and after each test, and review the test data to ensure they have not been inadvertently modified.
- Always check data bounds – field size, time, date, etc. – to help eliminate any buffer overflow.
- Make sure to sanitize any test data to ensure they will not expose any sensitive information.
- You should not use production data in tests until you are preparing for the final user acceptance test.

Testing controls should include the following:

- Make sure to test all changes.
- Make sure your management team confirms the results of all tests.
- Have the program librarian retain all implementation test data.
  - Use this data to test modifications.
- Use a separate copy of production data for any parallel runs.
  - Use copies of master files, never production versions.

Certification and Accreditation

Certification is the process of evaluating the security stance of the new software against a narrow set of predetermined security standards.

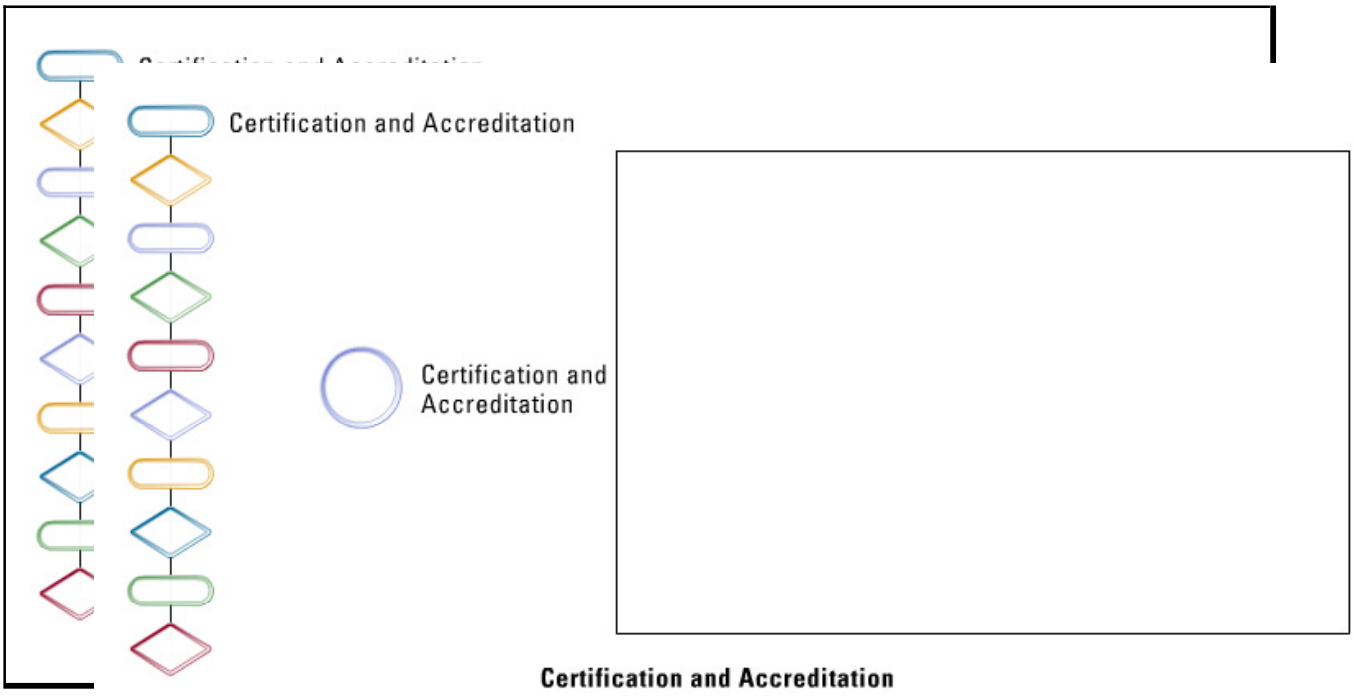

Certification and Accreditation

Certification is also the process of determining how well the system performs its intended functional requirements. Make sure to document the evaluation results in some type of security certification paper. This document should contain an analysis of the technical and non-technical security features and their countermeasures as well as the extent that the software meets the security requirements for its intended mission and operational environment. A certified security officer should then verify the software has been tested and that it meets all applicable policies, regulations, and standards for securing information systems. You must note all exceptions and supply them to the accreditation officer.

Once the accreditation officer reviews the certification document, he or she will authorize the software to be implemented in a production status but only in a specific environment and only for a specific period of time.

- **Provisional accreditation** is specified as being valid for only a specific time period and should outline any required change to the application, system, or accreditation document.
- **Full accreditation** means no changes are required for making the accreditation decision.

Installation

During the installation phase you will transition the new system from the acceptance phase into the live production environment.

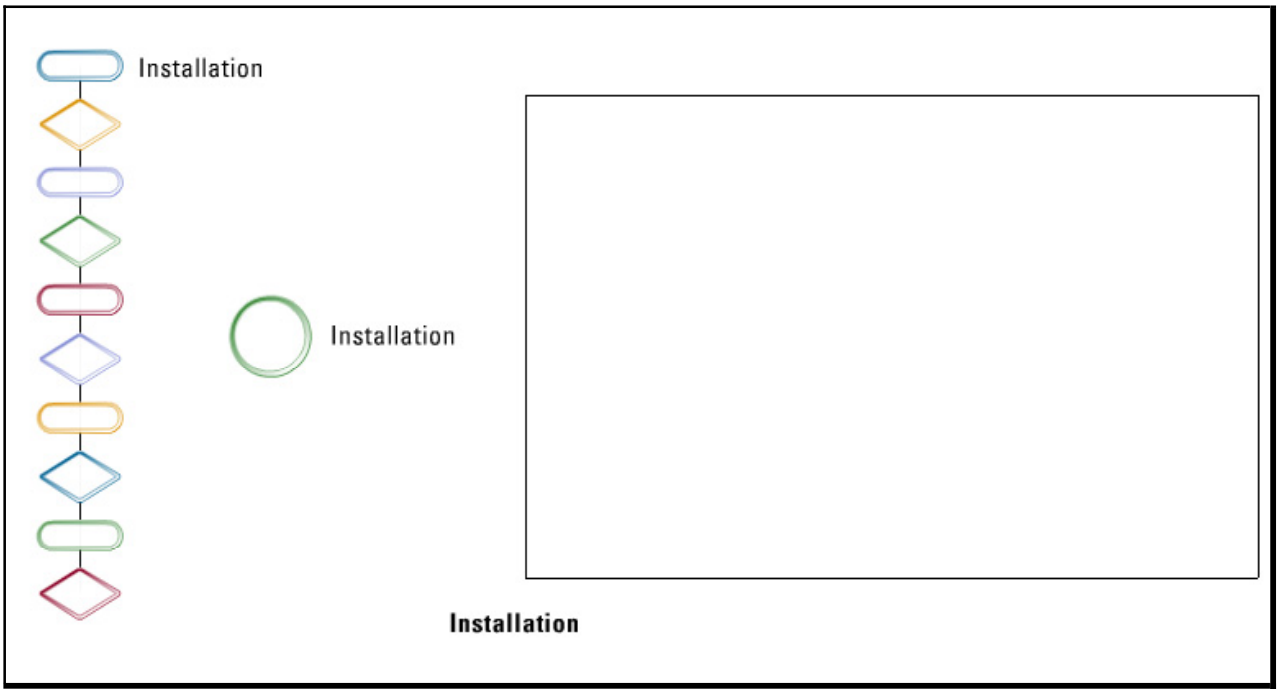

You should perform these activities during this phase:

- Obtain the required security accreditation (if not already included in the accreditation process)
- Train the end-user population according to the implementation and training schedule
- Implement the system (installation, data conversions, etc.)

You must control data conversion and data entry, and you should only allow the necessary people to access the system during this process. Also, you must have appropriate controls in place to reconcile and validate the accuracy of information you enter into the new system.

Post Installation

During the post-installation phase the new system should be in general use throughout the organization.

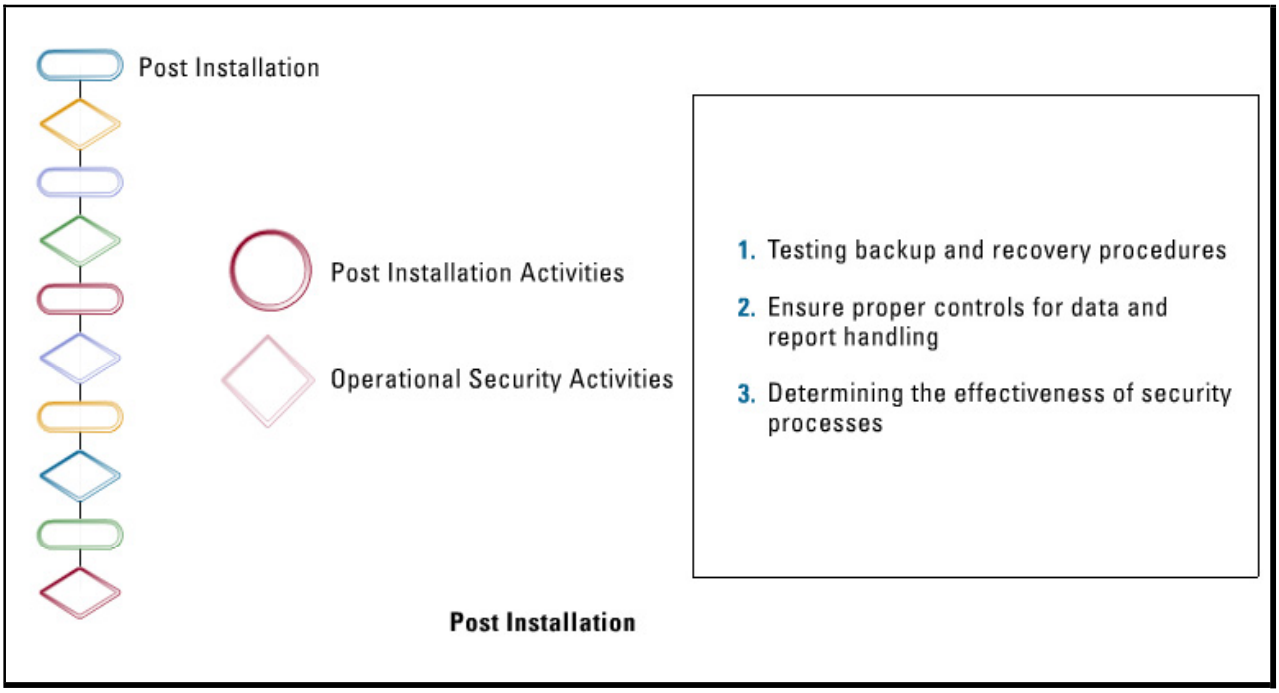

Activities required during this phase include:

- Monitoring the performance of the system
- Ensuring continuity of operations
  - Detecting defects or weaknesses
  - Managing and preventing system problems
  - Recovering from system problems
  - Implementing system changes

Operational security activities during this phase include:

- Testing backup and recovery procedures
- Ensuring proper controls for data and report handling
- Determining the effectiveness of security processes

Maintenance activities during this phase include:

- Periodic risk analysis and recertification of sensitive applications when significant changes occur. Significant changes include:
  - Changes in data sensitivity
  - Relocation
  - A major change to the physical environment (new equipment, external interfaces, operating system software, or application software)

It is vital that verification of any procedure or functionality does not disable or circumvent the security features.

Revisions and Replacement

In the revisions and replacement phase, the new system is already in production mode and hardware and software baselines should be subject to periodic evaluations and audits.

Revisions and Replacement:

- System is in production
- Periodically evaluate hardware and software baselines
- If a flaw or defect is identified:
  - Follow same SDM for repair
  - Record changes in change management system
- Revisions should include security planning and procedures
- Conduct periodic application audits
- Document any security incidents

If you determine any flaw or defect and changes to the application are required, the application must follow the same SDM and be recorded in a change management system.

To avoid any future problem, revision reviews should include security planning and procedures. Make sure you conduct periodic application audits and include the documentation of any security incidents when they occur. It is much easier to justify any future system enhancement when you have supporting documentation of system failures.

Summary

The key points discussed in this lesson are:

- You can breaksystems development controls into two distinct controls, which are designing controls and controlling the development process.
- Several SDMs have evolved to satisfy the different requirements for software programming.
- After you select the proper software development model, you need to follow several phases of development.
- When a project begins it is important that you identify the business needs (functional requirements) along with any current proposed technical solution.
- During the functional requirements phase, the project team should conduct a comprehensive analysis of current and possible future functional requirements to make sure that the new system will meet all end-user requirements.
- The system design specifications phase should include all activities related to designing the entire system and necessary software.
- During the build/development phase, you generate the program's source code as well as proper test scenarios and test cases.
- The acceptance phase allows for an independent group to develop test data and test the code to ensure that it will function within the organization's environment and within specification, and that it will meet all the functional and security requirements.
- During the testing and evaluation phase, you must follow several guidelines as appropriate to your environment.
- Certification is the process of evaluating the security stance of the new software against a narrow set of predetermined security standards.
- During the installation phase you will transition the new system from the acceptance phase into the live production environment.
- During the post-installation phase the new system should be in general use throughout the organization.
- In the revisions and replacement phase, the new system is already in production mode and hardware and software baselines should be subject to periodic evaluations and audits.

Security and Protection

Overview

A plethora of applets, agents, and viral code can attack computer systems. These attacks can be benign or catastrophic to the system. This lesson will discuss the more common computer attacks seen and how you can mitigate them.

Importance

Understanding how certain attacks are perpetrated, the way they infiltrate a system, and how they harm it is important to the information security professional. Through this understanding, the security professional will be able to identify different types of attacks and mitigate their effects.

Objectives

Upon completing this lesson, you will be able to:

- Identify common types of computer attacks
- Define threat agents
- Define mobile code
- Describe the features of the sandbox
- Identify the advantages and disadvantages of Java applets
- Describe the functionality of ActiveX
- Define viruses
- Define file-infector viruses
- Define boot sector viruses
- Definemultipartite viruses
- Define macro viruses
- Define script viruses
- Define encrypted and polymorphic viruses
- Define worms
- Define Trojan horses

- Define logic bombs
- Describe the functionality and benefits of antivirus software
- Describe the functionality of DCOM

Outline

The lesson contains these topics:

- Types of Computer Attacks
- Threat Agents
- Mobile Code
- Sandbox
- Java
- ActiveX
- Viruses
- File-Infector Viruses
- Boot Sector/System Infector Viruses
- Multipartite Viruses
- Macro Viruses
- Script Viruses
- Encrypted and Polymorphic Viruses
- Worms
- Trojan Horses
- Logic Bombs
- Antivirus software
- Distributed Component Object Model (DCOM)

Types of Computer Attacks

Many types of computer attacks can affect your system. From social engineering attacks to virus attacks to Time of Check/Time of Use (TOC/TOU) exploits, each attack has a prevention mechanism or countermeasure you can use to thwart or minimize its success.

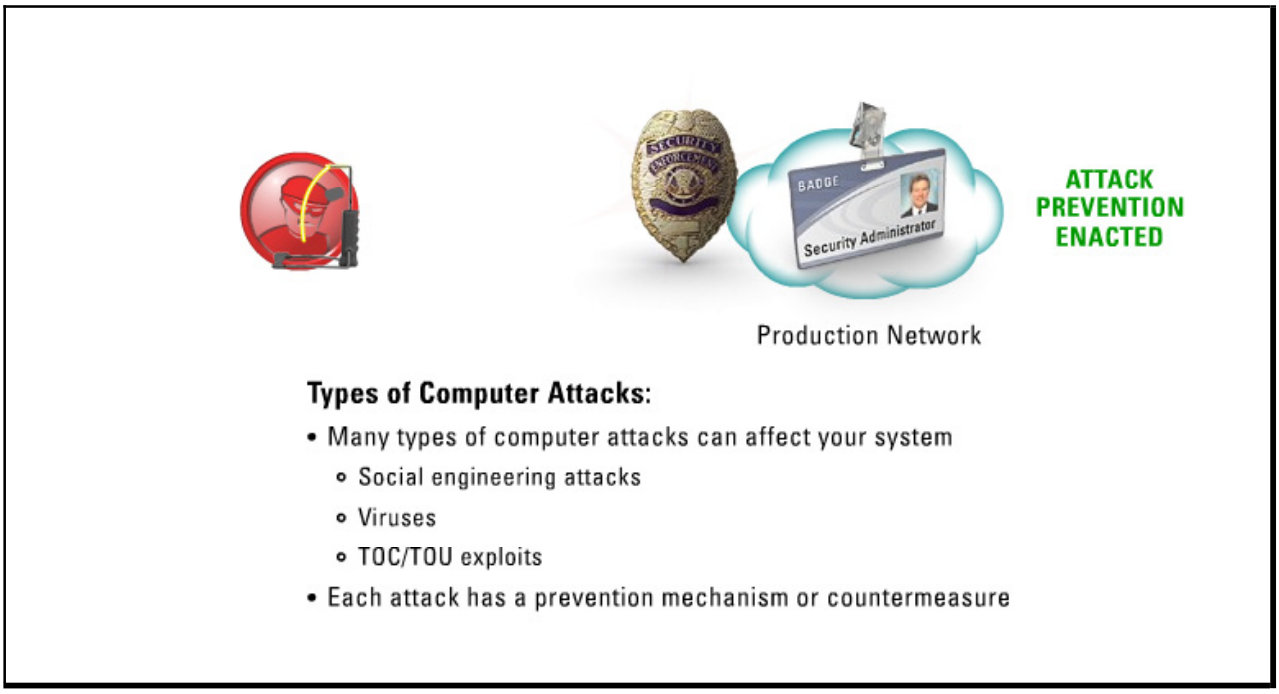

For example, the best method to prevent social engineering attacks is to make users aware of this type of threat and give them proper procedures for handling unusual requests for information or clearance.

To avoid TOC/TOU problems, especially file-based issues, the system programmer should avoid any file system call that takes a filename for an input, instead of a file handle or a file descriptor. When using file descriptors, it is easy to ensure that the file being used does not change after it has been called. Also, files that are to be used should be kept in their own directory, where the directory is only accessible by the Universal ID (UID) of the program performing the requested file operation. When you use these symbolic names, attackers are not able to exploit a race condition, unless they already have the proper UID.

Threat Agents

A **threat agent** is any entity that threatens an information system by exploiting vulnerabilities.

Threat agents are any entity that threatens information security:

- Can be human
- Can be programmatic (an error or malware)
- Can be a natural disaster

Audits:

- Look at potential threat agents to an asset
- Determine which factors may result in risk
- Identify countermeasures to mitigate threat

Threat agents can be human, programmatic (such as an error or Malware), or a natural disaster. Audits look at all the potential threat agents to an asset, determine which factors may result in the risk to the asset, and identify what countermeasures can be configured to mitigate the threat.

Mobile Code

Mobile code, also called executable content, is any software that can be transmitted across a network from a remote source to a local system, and then executed on the local system.

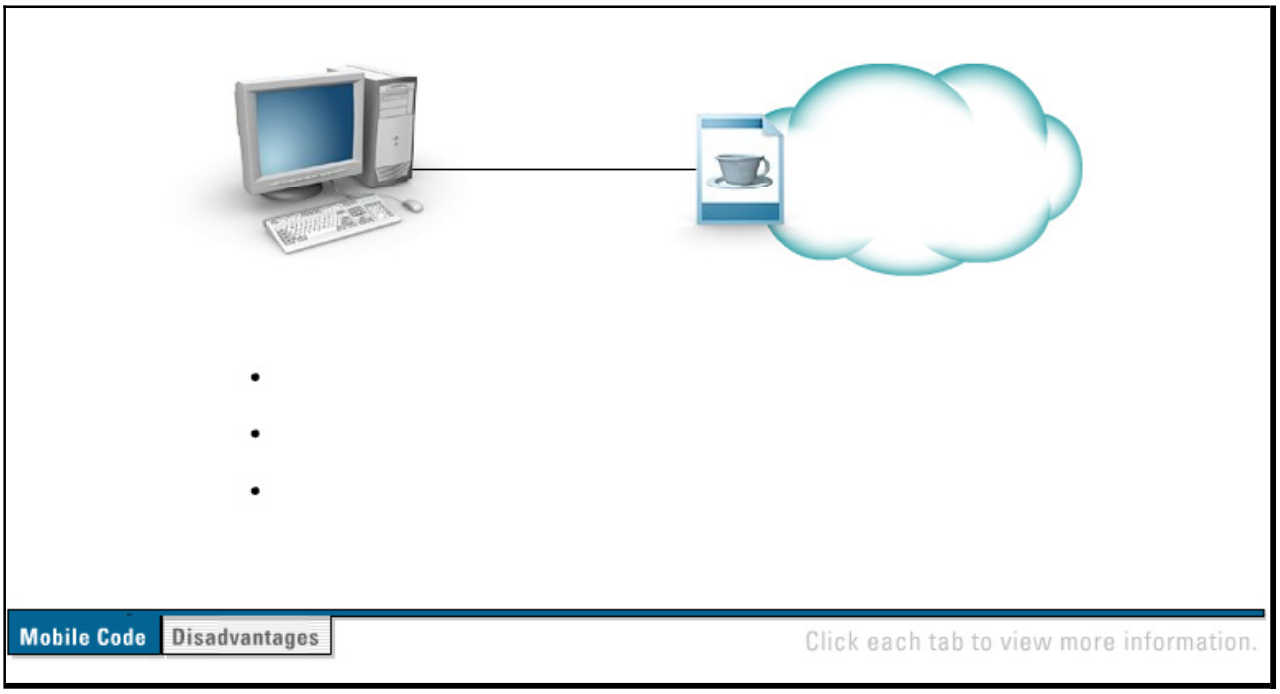

Mobile code is usually transferred via a user action, but can also be transmitted without the explicit action of the user. In general, code can be transmitted to the local system via these two ways:

- Attachments to e-mail messages
- Web pages

An example of mobile code is Java applets written in the Java programming language. Since Java is a mobile code, it does not matter what operating system the local system is using. As long as the system is Java-compliant, it can execute the code. It does not matter if the system is Microsoft Windows-based, Apple Macintosh-based, or any variant of UNIX, which is the power behind mobile code.

But, there is a drawback to executing code on a local system, which may have been downloaded from potentially untrusted sources. Executing Java code inside a Web browser could possibly reveal information that is on the local hard drive. Also, it is possible to use the Hypertext Transport Protocol (HTTP) header to report any information that Web browsers provide, such as the last addressed Internet Protocol (IP) address, the local system's IP address, username and password, and browser type.

Sandbox

The **sandbox** is one of the control mechanisms for mobile code. A sandbox will provide a protective area for program execution.

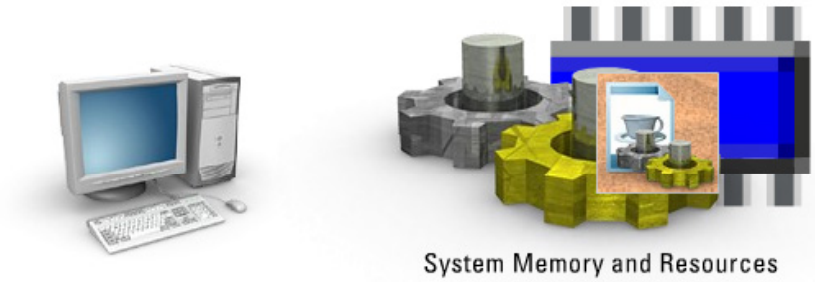

Sandbox:

- One of the control mechanisms for mobile code
- A protective area for program execution
- Limits the amount of memory and processor resources the program can consume
- If program exceeds limits, the Web browser will terminate the process
- Usually created on the client side to protect resource usage from Java applets
- The Java Security Manager is the entity that restricts untrusted code to the sandbox
- Trusted code always resides outside the sandbox

The sandbox provides this protective area by limiting the amount of memory and processor resources the program can consume. If a program does exceed these limits, the Web browser will terminate the process and log an error code. The sandbox is usually created on the client side to protect the resource usage from Java applets; this can ensure the safety of the browser's performance.

The Java Sandbox Security model, which is included in Java application programming interface (API) releases, provides an area for the Java code to do what it needs to do, such as restricting the bounds of this area. Remember, that a sandbox cannot confine code and its behavior without some type of enforcement mechanism. When using Java, the Java Security Manager is the entity that makes sure all restricted code stays in the sandbox. Trusted code will always reside outside the sandbox, while untrusted code is confined within it.

Note	By default, Java applications always reside outside the sandbox, while Java applets are confined within it. This fact ensures that an untrusted application cannot gain access to any system resource.

Java

With the release of more powerful markup languages that include support for forms and scripts, Web pages have become bright, vivid, and eye catching. However, you still need to download the pages statically. Java changed this fact by enabling programs called applets, which you can download and execute directly on your system. Java is a programming language created by Sun Microsystems that has some very interesting properties.

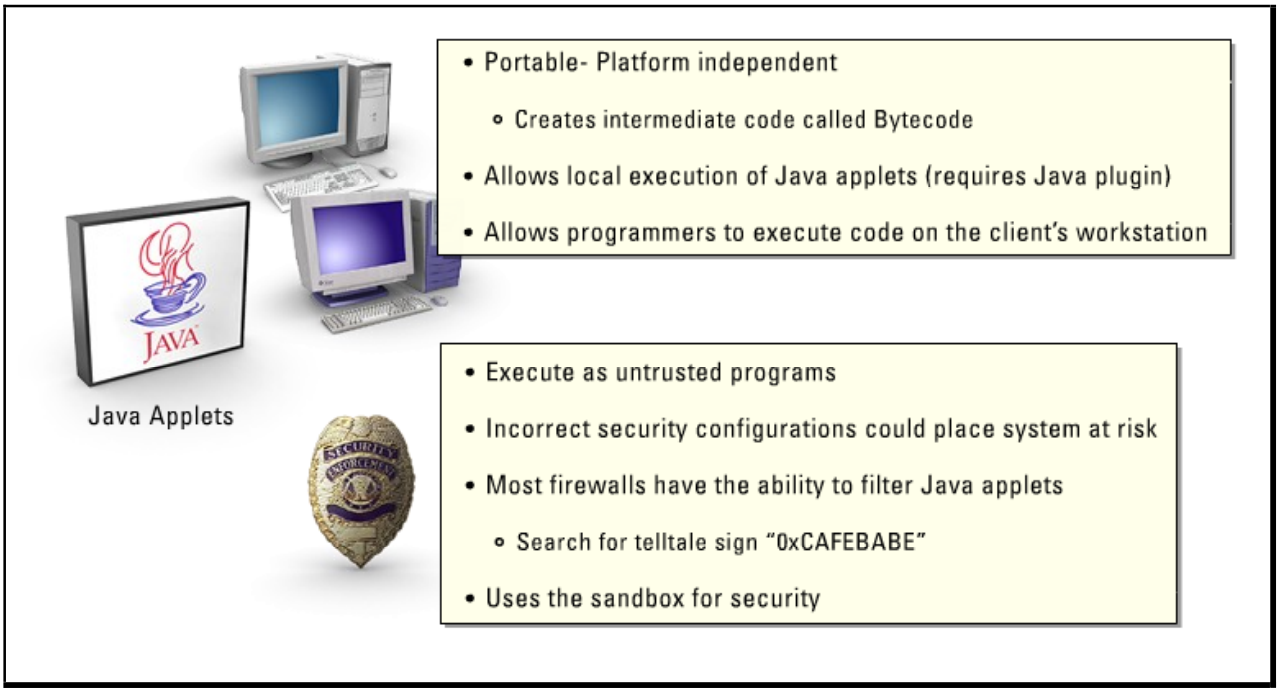

Sun created Java to be portable so that programs can be dynamically loaded over the network and run locally on any operating system with the Java plugin. This feature allows programmers to actually execute programming code on the client's workstation.

With the increased power of Web content using Java also came a potential problem. Applets will execute any code for which the author programmed them. Users surfing the Web now have to worry about potentially hostile Java applets writing malicious code to random access memory (RAM), sending confidential data to unknown destinations, or erasing all data on the hard drive.

Java applets execute as untrusted programs, meaning they have very limited access to client memory and CPU resources. However, if the client incorrectly configures security settings, the client's system as well as the entire network could be at risk of attack. For this reason, many security administrators do not let Java applets pass from unknown servers. Most firewalls have the ability to filter out Java applets as they attempt to pass the trusted interface. They do so by searching for the telltale sign "0xCAFEBABE," which is the Java Virtual Machine (JVM) file type identity number.

Java is platform independent because it creates intermediate code, bytecode, which is not processor specific (mobile code). The JVM then converts the bytecode to the correct machine code necessary for execution on the specific platform. Java applets use a security scheme that employs a sandbox to limit the applet's access to specific areas within the user's system and to protect these areas from malicious or poorly written applets.

ActiveX

Microsoft's **ActiveX** is a set of technologies built on the Component Object Model (COM) that enables software components, regardless of the programming language they were created in, to work together (such as Java and Visual Basic).

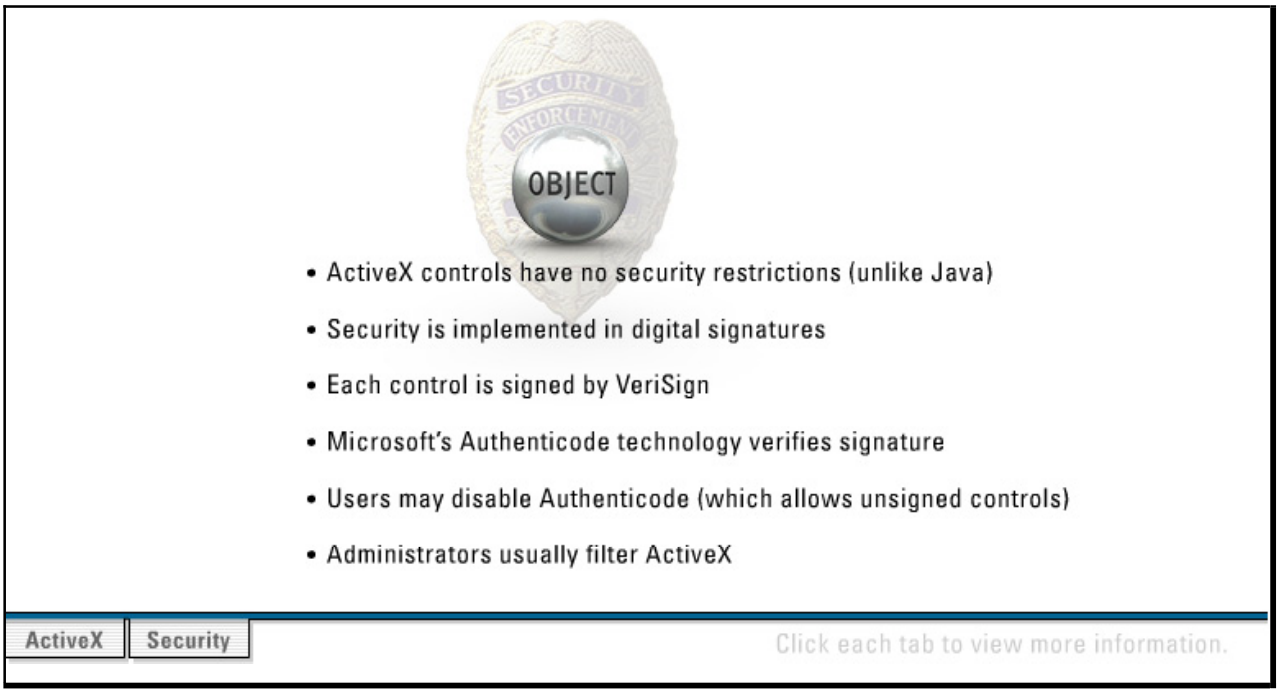

ActiveX controls, formerly called Object Linking and Embedding (OLE), are reusable, stand-alone components. ActiveX controls are the interactive objects in a Web page that provide user-controllable functions to enhance the experience of visiting a Web site.

Java is normally configured to execute in a protected memory area. Critical areas such as the file system or the boot sector are strictly off-limits. Theoretically, this makes it impossible for applets built using Java to damage a computer or its contents. ActiveX, on the other hand, has no such restrictions, allowing controls to reside on a system and use any of its resources, even writing to protected memory and the hard drive.

ActiveX security is implemented in digital signatures. Here, each control is packaged with a digital signature signed by VeriSign. Microsoft's Authenticode technology then verifies the signature with one of its built-in root certificates to make sure the control was not tampered with before downloading. However, users can disable Authenticode, which enables unsigned controls to be downloaded without warning, which is when problems begin to occur. For this reason, security administrators filter ActiveX in much the same way as they filter Java applets.

Viruses

A **virus** is a software structure of computer code that attaches itself to a program or file so it can spread from computer to computer, infecting as it travels.

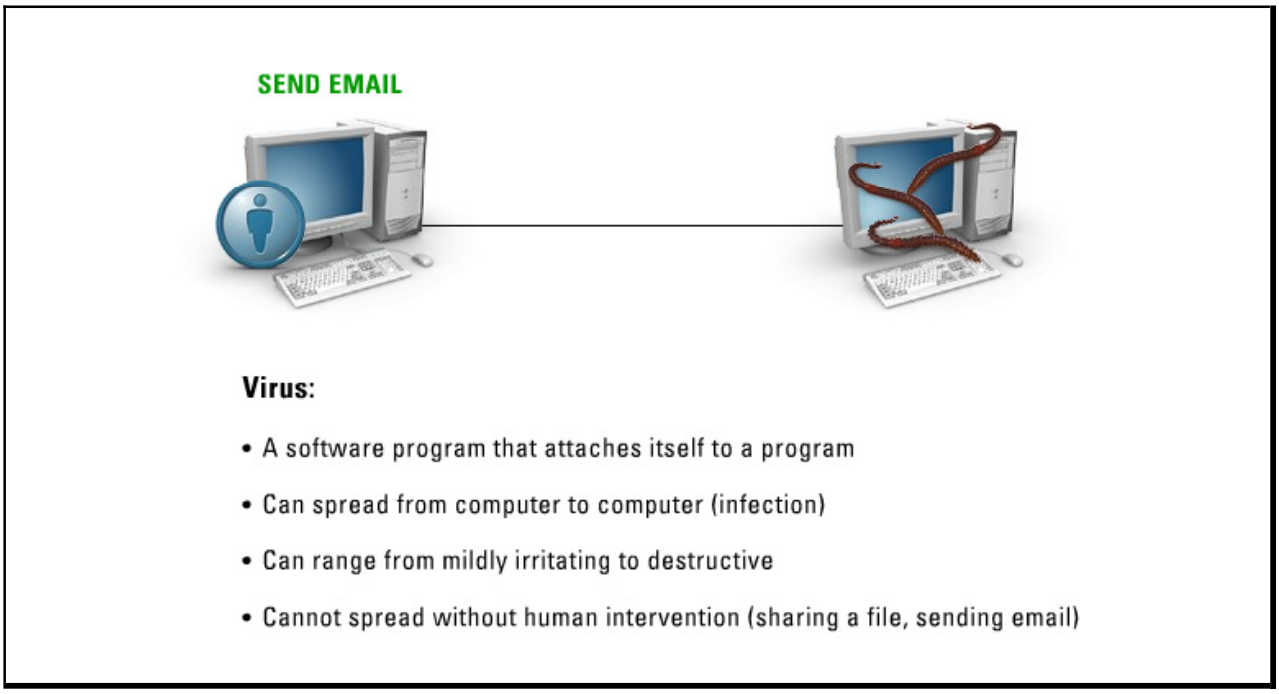

Viruses can damage your software, hardware, and files. A virus is written with the express intention of replicating itself. A virus attempts to spread from computer to computer by attaching itself to a host program.

Just as human viruses range in severity from the Ebola virus to the common cold, computer viruses range from the mildly irritating to the downright destructive. The good news is that a true virus does not spread without human action to move it along, such as sharing a file or sending an e-mail.

File-Infector Viruses

File-infector viruses attach themselves to executable program code, usually .com or .exe files. However, some of the more virulent examples can infect any program for which operating system execution is enabled, including .sys, .ovl, .prg, and .dll files.

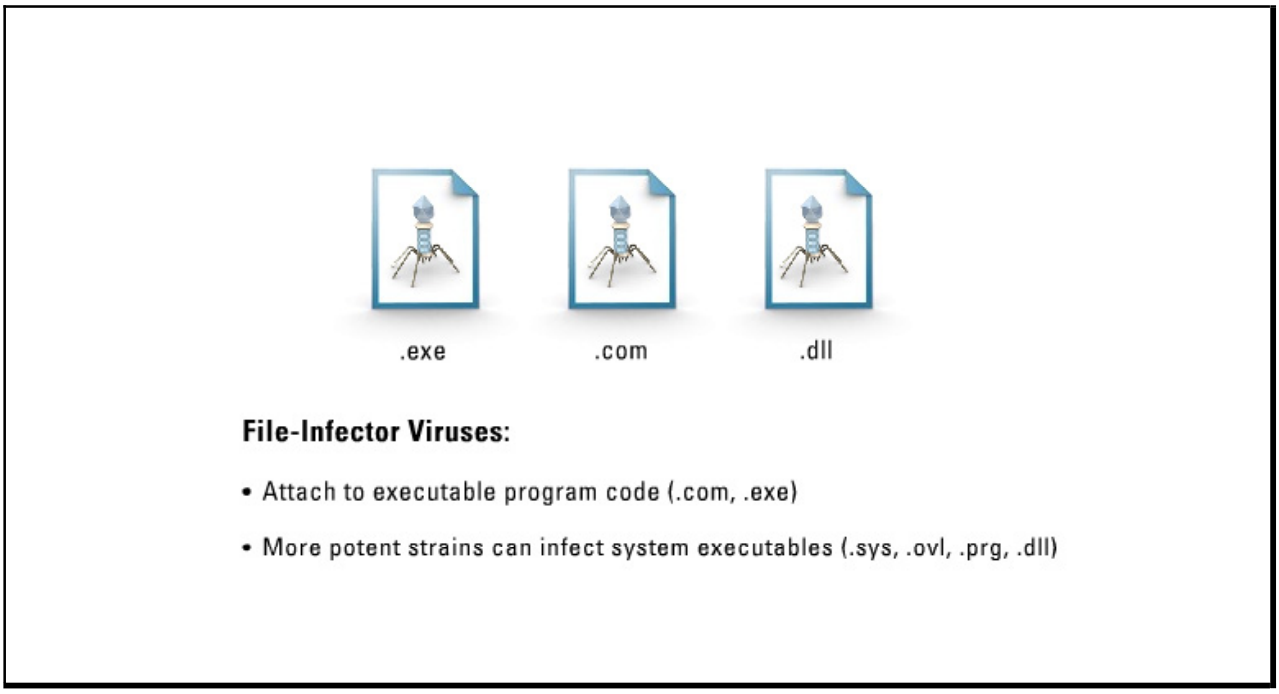

When the program is loaded, the virus piggybacks on the legitimate code to access the host system. Malicious Trojan horse code, which is contained within apparently harmless applications, can be propagated as part of a file infector virus.

Boot Sector/System Infector Viruses

System or **boot sector viruses** work by infecting executable code found in system areas. They target the disk operating system (DOS) boot sector on floppies or the master boot record on hard drives.

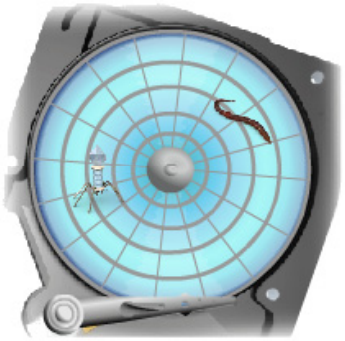

Boot Sector/System Infector Viruses:

- Work by infecting executable code found in system areas
- Target the disk operating system (DOS) boot sector on floppies or MBR on hard drives
- Infection occurs when booting from an infected floppy or compromised hard drive
- When infection occurs virus hides in memory, where it can infect any file
- Most devastating type of system virus is called a worm
- Worms do not change files (they place themselves in memory and propogate)

Once this type of infection has been loaded into a host system, through booting with an infected floppy disk or a compromised hard drive, the virus hides in memory, which makes it very difficult to detect.

When the virus has taken up residence in system memory it can infect any uninfected file that is executed. The only way to remove the virus is to turn off the power on the infected computer.

One of the most devastating types of system viruses is called a worm. Worms do not change files, but place themselves in active system memory where they propagate, usually without being seen by the user of the host computer. The presence of such viruses is often noticed only when their uncontrolled replication consumes system resources and brings the host to a standstill.

Multipartite Viruses

A **multipartite virus**, also called a multi-part virus, is a virus that attempts to attack both the boot sector and the executable, or program, files at the same time.

Multipartite Viruses:

- A virus that attempts to attack both the boot sector and executable at the same time
- When infection of boot sector occurs, virus will infect files (and vice versa)

When the virus attaches to the boot sector, it will in turn infect the system's files, and when the virus attaches to the files, it will in turn infect the boot sector.

Caution	This type of virus can re-infect a system over and over again if you do not eradicate all parts of the virus.

Macro Viruses

Macro viruses are among the most common type of virus, but do not typically have substantially destructive payloads. Generally created using Visual Basic scripts, macros 'infect' a Microsoft Office or similar application and cause their payloads to be triggered when the application is started.

Macro viruses are computer viruses that use an application's own macro programming language to distribute themselves. These macros have the potential to inflict damage to the document or to other computer software. These macro viruses can infect Word files, as well as any other application that uses a programming language.

Unlike previous viruses, macro viruses do not infect programs; they infect documents and templates. Opening a document or template that contains a macro virus will infect your system, and the virus will spread to other documents and templates you may have on your system. Some macro viruses are not harmful, but they can be annoying. However, there are some macro viruses that can be very destructive. Also, Word macro viruses can be spread across platforms; for example, the macro virus can infect files on the Windows platform, as well as files on the Macintosh platform.

It is not always easy to determine whether you have a macro virus. If you are familiar with the Word macros you have on your system, you can look through the various macros for ones that you do not recognize. It is possible that one, or more, of them are part of a macro virus that has infected your system. Some examples of these types of macro names are: AAAZAO, AAAZFS, AutoOpen, FileSaveAs, and PayLoad.

Script Viruses

Script viruses are a subset of file viruses, written in a variety of script languages (VBS, JavaScript, BAT, PHP, etc.). They either infect other scripts (e.g., Windows or Linux command and service files) or form a part of multi-component viruses. Script viruses are able to infect other file formats, such as HTML, if the file format allows the execution of scripts.

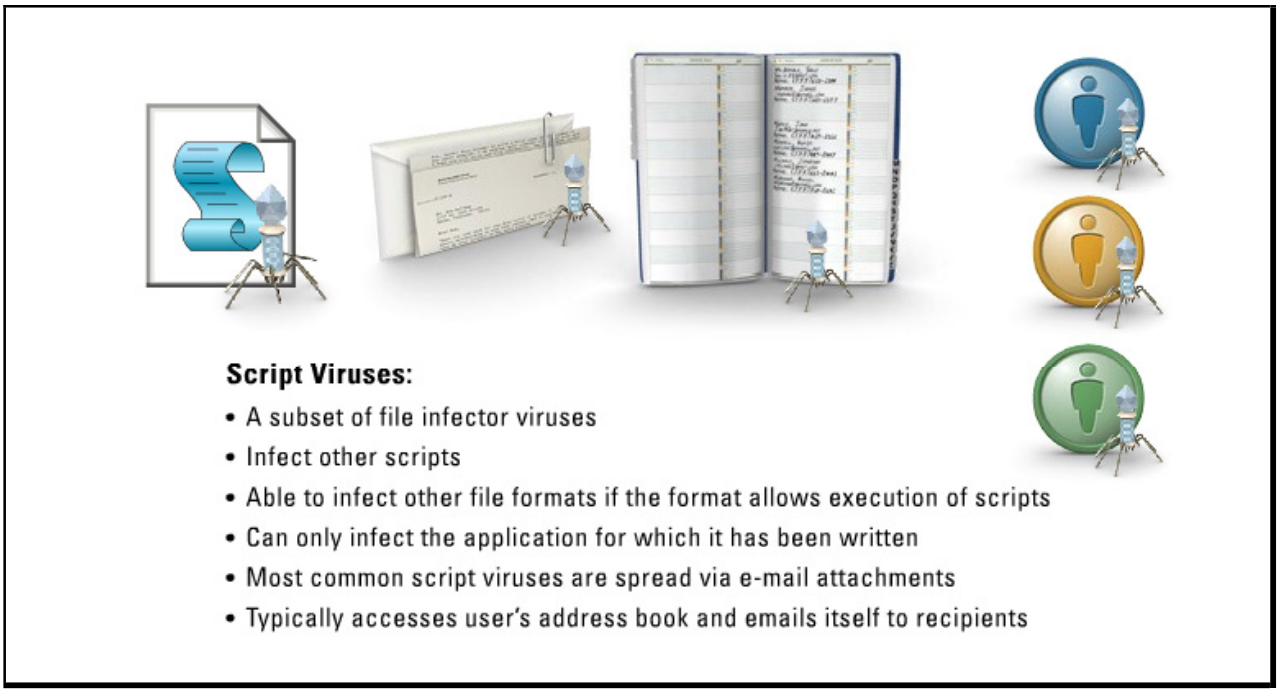

Scripts access a user's address book (and any other specific services) using a specific COM-based programming interface. This means that a script virus will only affect the application for which it has been written. For example, the Melissa and Love Bug viruses were written to specifically target the Microsoft Outlook Windows client. These script viruses would not spread if executed on a computer that was not using Outlook. This important fact explains why these viruses did not affect users of other email applications, such as Novell GroupWise.

The most common script viruses are spread via e-mail attachments. An e-mail message is composed and the script virus is attached to the message. In many cases, the script file extension is 'hidden' by simply taking advantage of a common Windows feature, which makes the script file appear to be a 'safe' attachment, such as an image file. (Windows uses a file association system that associates file extensions with the program that is used to run or open a file. The Windows Script Host, which actually executes Visual Basic Script code, is associated with script files using either the .VBS or .VBE file extension.) When the recipient of the email message double-clicks the attachment, the script is executed by the Windows script host. A script virus typically accesses the user's address book in order to email itself to a large number of recipients, thus continuing the cycle of infection.

Encrypted and Polymorphic Viruses

A simple virus that merely replicates itself is the easiest to detect. If a user launches an infected program, the virus gains control of the computer and attaches a copy of itself to another program file. After it spreads, the virus transfers control back to the host program, which functions normally. Yet, no matter how many times a simple virus infects a new file or floppy disk, for example, the infection always makes an exact copy of itself. Anti-virus software needs only search, or scan, for a telltale sequence of bytes (known as a signature) found in the virus.

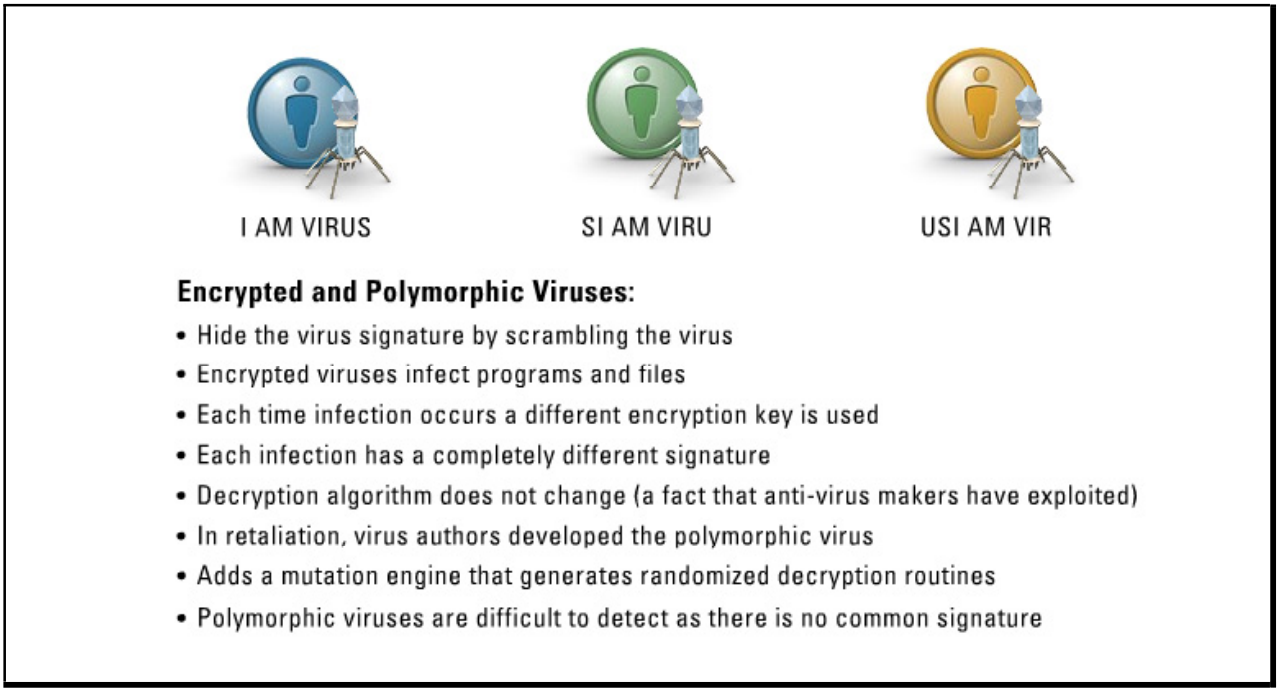

In response, virus authors began encrypting viruses. The idea was to hide the fixed signature by scrambling the virus, making it unrecognizable to a virus scanner.

An **encrypted virus** consists of a virus decryption routine and an encrypted virus body. If a user launches an infected program, the virus decryption routine first gains control of the computer, and then decrypts the virus body. Next, the decryption routine transfers control of the computer to the decrypted virus.

An encrypted virus infects programs and files as any simple virus does. Each time it infects a new program, the virus makes a copy of both the decrypted virus body and its related decryption routine, encrypts the copy, and attaches both to a target.

To encrypt the copy of the virus body, an encrypted virus uses an encryption key that the virus is programmed to change from infection to infection. As this key changes, the scrambling of the virus body changes, making the virus appear different from infection to infection. This makes it extremely difficult for anti-virus software to search for a virus signature extracted from a consistent virus body.

However, the decryption routines remain constant from generation to generation – a weakness that anti-virus software quickly evolved to exploit. Instead of scanning just for virus signatures, virus scanners were modified to also search for the telltale sequence of bytes that identify a specific decryption routine.

In retaliation, virus authors developed the **polymorphic virus**. Like an encrypted virus, a polymorphic virus includes a scrambled virus body and a decryption routine that first gains control of the computer,

and then decrypts the virus body. However, a polymorphic virus adds a third component to these two components – a mutation engine that generates randomized decryption routines that change each time a virus infects a new program.

In a polymorphic virus, the mutation engine and virus body are both encrypted. When a user runs a program infected with a polymorphic virus, the decryption routine first gains control of the computer, and then decrypts both the virus body and the mutation engine. Next, the decryption routine transfers control of the computer to the virus, which locates a new program to infect.

At this point, the virus makes a copy of both itself and the mutation engine in RAM. The virus then invokes the mutation engine, which randomly generates a new decryption routine that is capable of decrypting the virus, yet bears little or no resemblance to any prior decryption routine. Next, the virus encrypts this new copy of the virus body and mutation engine. Finally, the virus appends this new decryption routine, along with the newly encrypted virus and mutation engine, onto a new program.

Polymorphic viruses are more difficult to detect by scanning because each copy of the virus looks different than the other copies. A polymorphic virus changes its virus signature (i.e., its binary pattern) every time it replicates and infects a new file in order to keep from being detected by an antivirus program.

Polymorphic computer viruses are the most complex and difficult viruses to detect, often requiring anti-virus companies to spend days or months creating the detection routines needed to catch a single polymorphic.

Worms

A **worm** is a sub-class of the virus and is designed to copy itself from one computer to another, but worms do so automatically by taking control of features on the computer that can transport files or information, like e-mail. Once you have a worm in your system it can travel alone.

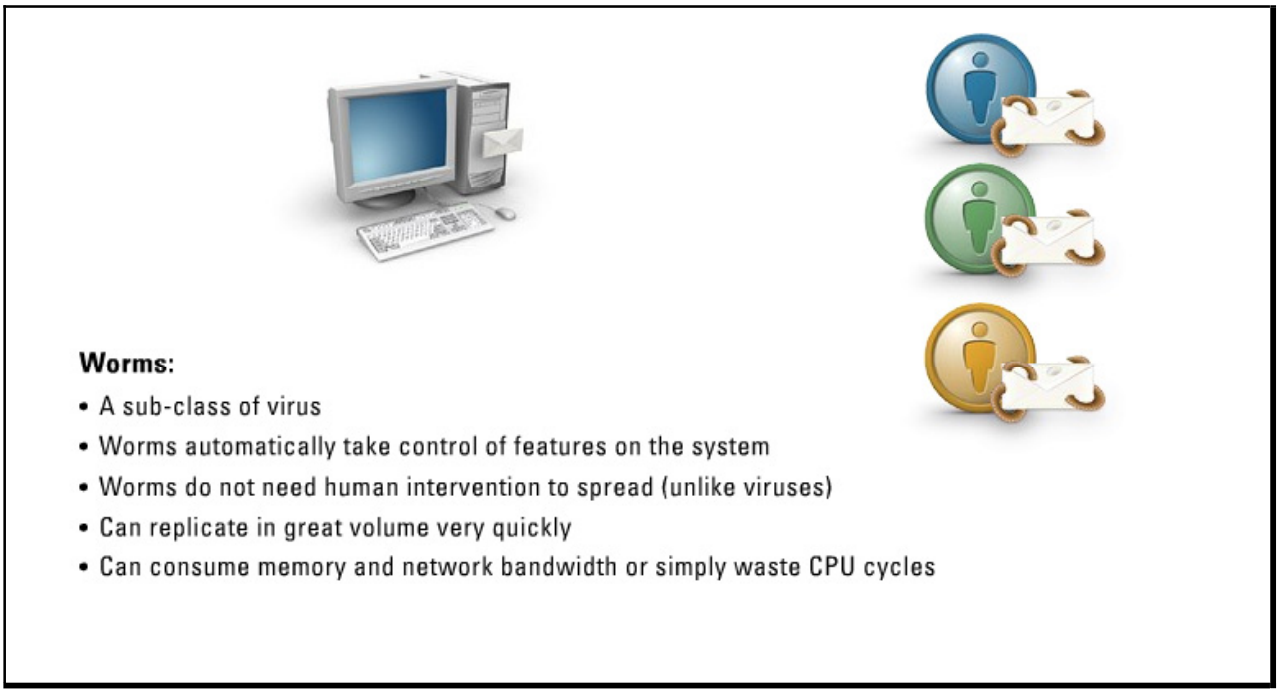

Worms are dangerous because they can replicate in great volume. For example, a worm could send out copies of itself to everyone listed in your e-mail address book, and their computers would then do the same, causing a domino effect of heavy network traffic that could slow down business networks and the Internet as a whole. When new worms are unleashed, they spread very quickly, clogging networks and possibly making you (and everyone else) wait twice as long as normal to view Web pages on the Internet. Because worms commonly use e-mail as their transport between hosts, worms are commonly spread very quickly among friends, relatives, and co-workers.

Note Common examples of worms are the Sasser worm and the Blaster worm.

A worm generally spreads without user action and distributes complete copies (possibly modified) of itself across networks. A worm can consume memory and network bandwidth, or simply cause the CPU on the system to waste CPU cycles, thus causing a computer to stop responding.

Because worms do not need to travel via a host program or file, they can also tunnel into your system and allow somebody else to take control of your computer remotely.

Trojan Horses

Unlike viruses or worms, **Trojan horses** have no replicating abilities; instead, Trojan horses are usually embedded into games or other software programs that look innocent.

Trojan Horses:

- Have no replicating abilities
- Instead are usually embedded into games or other innocent looking programs
- When program is opened, the Trojan executes its payload
- Most Trojans install software to allow backdoor access to a system

Once you start the application, the Trojan executes and performs its instructed duties. Most Trojans today install a piece of software that allows backdoor access to the system. These backdoor applications usually take the form of the software packages BackOrifice or SubSeven.

Just as the mythological Trojan horse appeared to be a gift, but turned out to contain Greek soldiers who overtook the city of Troy, today's Trojan horses are computer programs that appear to be useful software, but instead they compromise your security and can potentially cause a lot of damage. A recent Trojan horse came in the form of an e-mail that included attachments claiming to be a manufacturer's security update, but turned out to be a virus that attempted to disable antivirus and firewall software.

Trojan horses spread when people are lured into opening a program because they think it comes from a legitimate source. Trojan horses can also be included in software that you download for free. You should instruct end users in your organization to never download software from a source that they do not trust.

Logic Bombs

A **logic bomb**, like a real bomb, can ignite a destructive payload when the correct trigger is hit.

Logic Bombs:

- Ignite destructive payloads when a trigger is hit
- Once triggered, logic bombs can do anything

The trigger for a logic bomb is up to the author of the code, but can be something like a certain hour on the clock, a certain date on the calendar, the number of times a program was executed, certain mouse movements, etc. But, once the logic bomb triggers, it can do anything from erasing the hard drive to randomly changing data in the system or on the hard drive.

Logic bombs that randomly destroy data are tricky to identify, as it is difficult to know the bomb has been triggered. Even if you can identity the logic bomb after its detonation, it is usually too late to salvage any corrupted data. For this reason, it is a very good idea to maintain accurate backups of your valuable data and maintain adequate anti-virus protection across the enterprise that can detect logic bombs.

Antivirus Software

Just as you can use vaccines to stop or prevent infection of a biological virus, you can use antivirus software in much the same manner. **Antivirus software** is a software program that either comes installed on your computer or that you purchase and install yourself.

When you open and run an infected program, you might not know that you have contracted a virus. Your computer may slow down, stop responding, or crash and begin restarting every few minutes. Sometimes a virus will attack the files you need to boot up a computer. In this case, you might press the power button and find yourself staring at a blank screen. All of these symptoms are common signs that your computer has a virus – although they could also be caused by hardware or software problems that have nothing to do with having a virus.

Antivirus software helps protect your computer against most viruses, worms, Trojan horses, and other unwanted invaders that can make your computer "sick." Viruses, worms, and the like often perform malicious acts, such as deleting files, accessing personal data, or using your computer to attack other computers.

Antivirus software can help keep your computer healthy as long as you make sure to keep your antivirus software up to date. Companies are constantly creating antivirus updates as attackers introduce new viruses, worms, and Trojan horses into the computing world.

Tip	Nothing will guarantee the security of your computer systems 100 percent. However, you can continue to improve computer security in your organization by keeping all computer software up to date and maintaining a current antivirus software subscription.

Distributed Component Object Model (DCOM)

Today's programmers are creating applications with software systems that are based on the use of distributed objects, such as the Common Object Request Broker Architecture (CORBA), Java Remote Method Invocation (JRMI), Enterprise JavaBeans (EJB), and Distributed Component Object Model (DCOM).

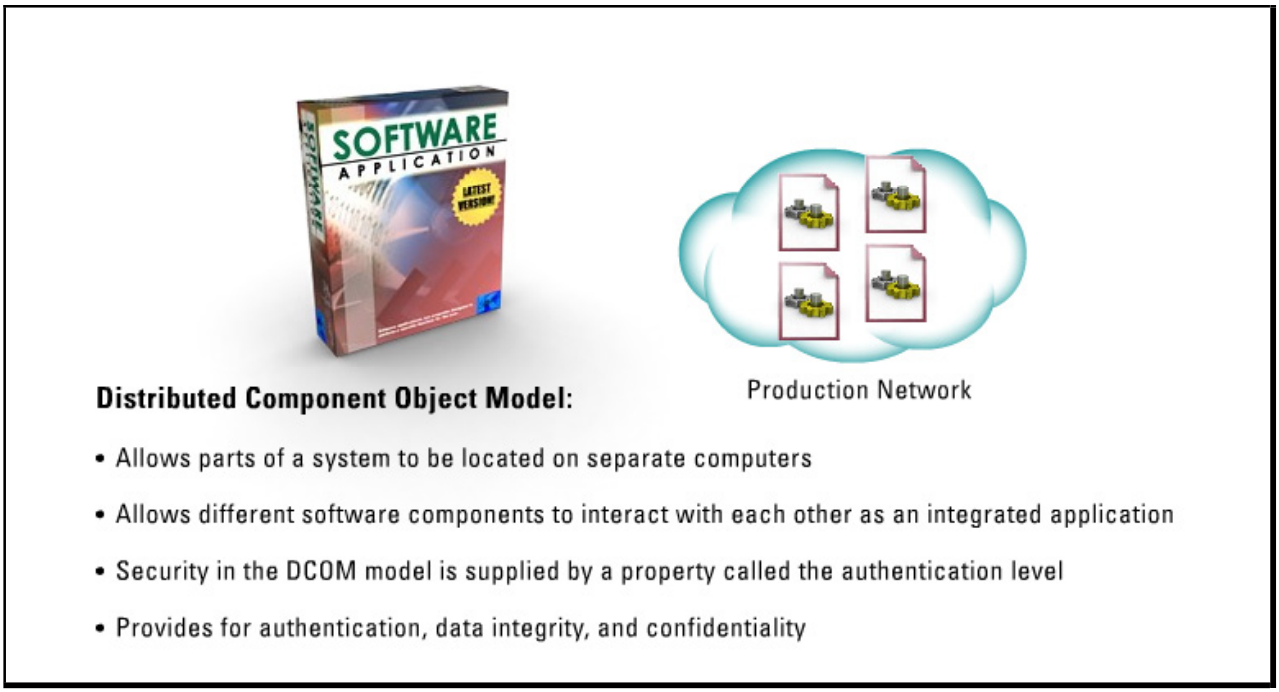

A distributed object-oriented system allows parts of the system to be located on separate computers within an enterprise network. The object system itself is a compilation of reusable self-contained objects of code designed to perform specific business functions.

Microsoft created **DCOM** for use on their Microsoft platforms. DCOM allows different software components to interact with each other as an integrated application. Security in the DCOM model is supplied by a single property called the authentication level, which provides for authentication, data integrity, and confidentiality. Authentication levels can only be applied to server objects, and each object can have its own level set. Higher levels of DCOM provide for additional security, but at a greater cost.

Summary

The key points discussed in this lesson are:

- Many types of computer attacks can affect your system. From social engineering attacks to virus attacks to Time of Check/Time of Use (TOC/TOU) exploits, each attack has a prevention mechanism or countermeasure you can use to thwart or minimize its success.
- A threat agent is any entity that threatens an information system by exploiting vulnerabilities.
- Mobile code, also called executable content, is any software that can be transmitted across a network from a remote source to a local system, and then executed on the local system.
- The sandbox is one of the control mechanisms for mobile code. A sandbox will provide a protective area for program execution.
- Java is a programming language created by Sun Microsystems that has some very interesting properties.
- Microsoft's ActiveX is a set of technologies built on the Component Object Model (COM) that enables software components, regardless of the programming language they were created in, to work together (such as Java and Visual Basic).
- A virus is a software structure of computer code that attaches itself to a program or file so it can spread from computer to computer, infecting as it travels.
- File-infector viruses attach themselves to executable program code, usually .com or .exe files. However, some of the more virulent examples can infect any program for which operating system execution is enabled, including .sys, .ovl, .prg, and .dll files.
- System or boot sector viruses work by infecting executable code found in system areas. They target the disk operating system (DOS) boot sector on floppies or the master boot record on hard drives.
- A multipartite virus, also called a multi-part virus, is a virus that attempts to attack both the boot sector and the executable, or program, files at the same time.
- Macro viruses are among the most common type of virus, but do not typically have substantially destructive payloads.
- Script viruses are a subset of file viruses, written in a variety of script languages (VBS, JavaScript, BAT, PHP, etc.). They either infect other scripts (e.g., Windows or Linux command and service files) or form a part of multi-component viruses.
- An encrypted virus consists of a virus decryption routine and an encrypted virus body. A polymorphic virus adds a third component to these two components – a mutation engine that generates randomized decryption routines that change each time a virus infects a new program.
- A worm is a sub-class of the virus and is designed to copy itself from one computer to another, but worms do so automatically by taking control of features on the computer that can transport files or information, like e-mail.
- Unlike viruses or worms, Trojan horses have no replicating abilities; instead, Trojan horses are usually embedded into games or other software programs that look innocent.
- A logic bomb, like a real bomb, can ignite a destructive payload when the correct trigger is hit.
- Antivirus software can help keep your computer healthy as long as you make sure to keep your antivirus software up to date.

- DCOM allows different software components to interact with each other as an integrated application.

Cryptography

Overview

Cryptography is the science of taking plaintext (or cleartext) and encoding it into ciphertext via a process called encryption. This process allows data in its ciphertext persona to travel across hostile networks (the Internet) without concern of loss of confidentiality. When the ciphertext reaches its destination it can then be decrypted back into its native plaintext form. Cryptography has four main concerns:

- Confidentiality
- Integrity
- Authentication
- Non-repudiation

Objectives

Upon completing this module, you will be able to:

- Define key cryptographic terms
- Describe cryptographic technologies
- Describe the mechanics of message authentication
- Describe the mechanics of certificate authorities

Outline

The module contains these lessons:

- Cryptographic Terms and Technologies
- Message Authentication
- Certificate Authority

Cryptographic Terms and Technologies

Overview

Cryptography is the art of garbling data so the data look nothing like their original form, and then being able to restore the data back to their original form at some future time. To garble the data, plain text is encrypted using a special value called a key. This key produces the garbled data, called ciphertext. Both ends of the secure link must know the encrypting algorithm as well as the key used to encrypt/decrypt the data. An attacker can sniff ciphertext, but the attacker cannot decipher the message without the correct algorithm and key. This lesson discusses key cryptographic terms and technologies.

Importance

It is important for the information security professional to understand the mechanics of cryptography, which will allow him or her to better secure sensitive information as it passes through untrusted or even trusted networks.

Objectives

Upon completing this lesson, you will be able to:

- Define key cryptographic terms
- Define encryption
- Define ciphers
- Define cryptanalysis
- Describe the logistics of symmetric algorithms
- Describe the logistics of the Data Encryption Standard
- Describe the logistics of the Data Encryption Standard 3
- Describe the logistics of the Advanced Encryption Standard
- Describe the logistics of additional symmetric encryption algorithms
- Describe the logistics of asymmetric algorithms

- Describe the logistics of the RSA asymmetric algorithm
- Describe the logistics of the Digital Signature Algorithm
- Describe the logistics of the Diffie-Hellman asymmetric algorithm
- Define elliptic curve cryptography
- Define man-in-the-middle attacks
- Define block and stream ciphers
- Define one-time cipher keys (pads)
- Describe the logistics ofElectronic Code Book mode
- Describe the logistics of Cipher Block Chaining mode
- Describe the logistics ofCipher Feedback mode
- Describe the logistics ofOutput Feedback mode
- Identify the six categories of cryptanalytic attacks

Outline

The lesson contains these topics:

- Key Terms
- Encryption
- Ciphers
- Cryptanalysis
- Symmetric Algorithms
- Data Encryption Standard
- Data Encryption Standard 3
- Advanced Encryption Standard
- Other Symmetric Encryption Algorithms
- Asymmetric Algorithms
- RSA Asymmetric Algorithm
- Digital Signature Standard
- Diffie-Hellman
- Elliptic Curve Cryptography
- Man-in-the-Middle Attack
- Block Ciphers
- Stream Ciphers
- One-Time Cipher Keys (Pads)
- Electronic Code Book
- Cipher Block Chaining
- Cipher Feedback

- Output Feedback
- Attacking Encryption

Key Terms

This topic defines key cryptographic terms.

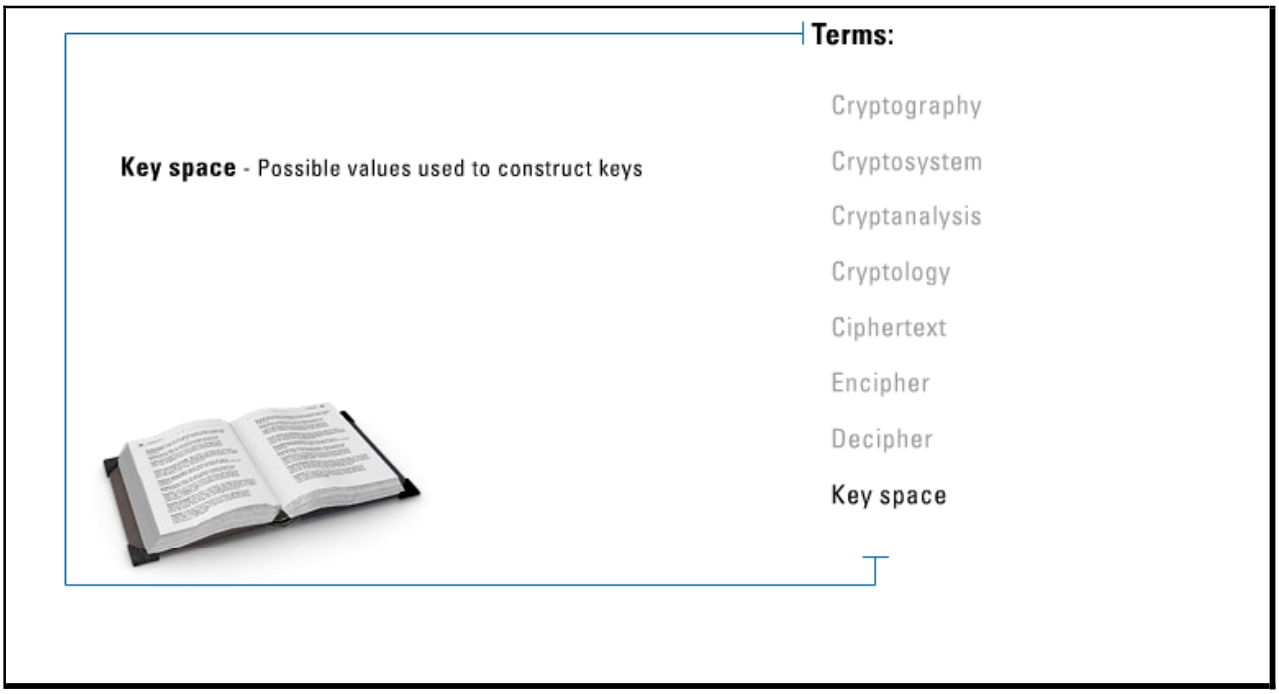

- **Algorithm** - Set of mathematical rules used in encryption and decryption
- **Cryptography** – Science of secret writing that enables you to store and transmit data in a form that is available only to the intended individuals
- **Cryptosystem** - Hardware or software implementation of cryptography that transforms a message to ciphertext and back to plaintext
- **Cryptanalysis** - Practice of obtaining plaintext from ciphertext without a key or breaking the encryption
- **Cryptology** - The study of both cryptography and cryptoanalysis
- **Ciphertext** - Data in an encrypted or unreadable format
- **Encipher** - Act of transforming data into an unreadable format
- **Decipher** - Act of transforming data into a readable format
- **Key** - Secret sequence of bits and instructions that governs the act of encryption and decryption
- **Key Clustering** - Instance when two different keys generate the same ciphertext from the same plaintext
- **Key space** - Possible values used to construct keys
- **Plaintext** - Data in readable format, also referred to as cleartext
- **Work Factor** -Estimated time, effort, and resources necessary to break a cryptosystem

Encryption

Encryption is the key algorithm used in cryptography.

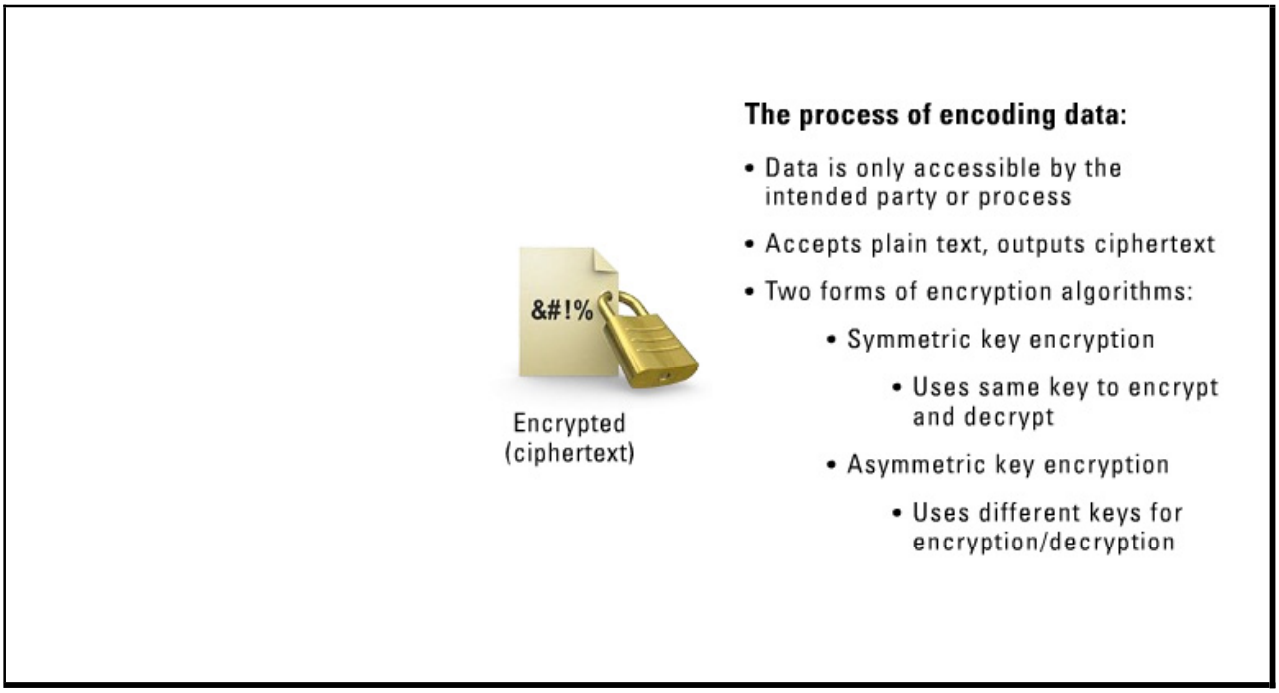

Encryption is the process of encoding data to ensure the data are accessible only by the intended recipient. Online credit card purchases are often encrypted. Decryption is the process of decoding the data. To encrypt data, break the plaintext data into pieces and insert the pieces into the encryption algorithm with an encryption key. The algorithm outputs the ciphertext that is sent to the peer. The peer performs the same algorithm in reverse using the same key.

The end result of encryption is that only the person who has the shared secret key can decrypt the ciphertext back into its plaintext form.

There are two types of encryption algorithms:

- **Symmetric Key Encryption** - This encryption method uses a shared secret key to both encrypt and decrypt data.
- **Asymmetric Key Encryption** - This encryption method uses two specially created mathematical keys. These keys have the interesting quality in that what one key encrypts, the other key can decrypt. The same key cannot both encrypt and decrypt the same data.

Ciphers

A **cipher** is any cryptographic system in which arbitrary symbols, or groups of symbols, represent units of plaintext of regular length, usually single letters, or in which units of plaintext are rearranged, or both, in accordance with certain predetermined rules.

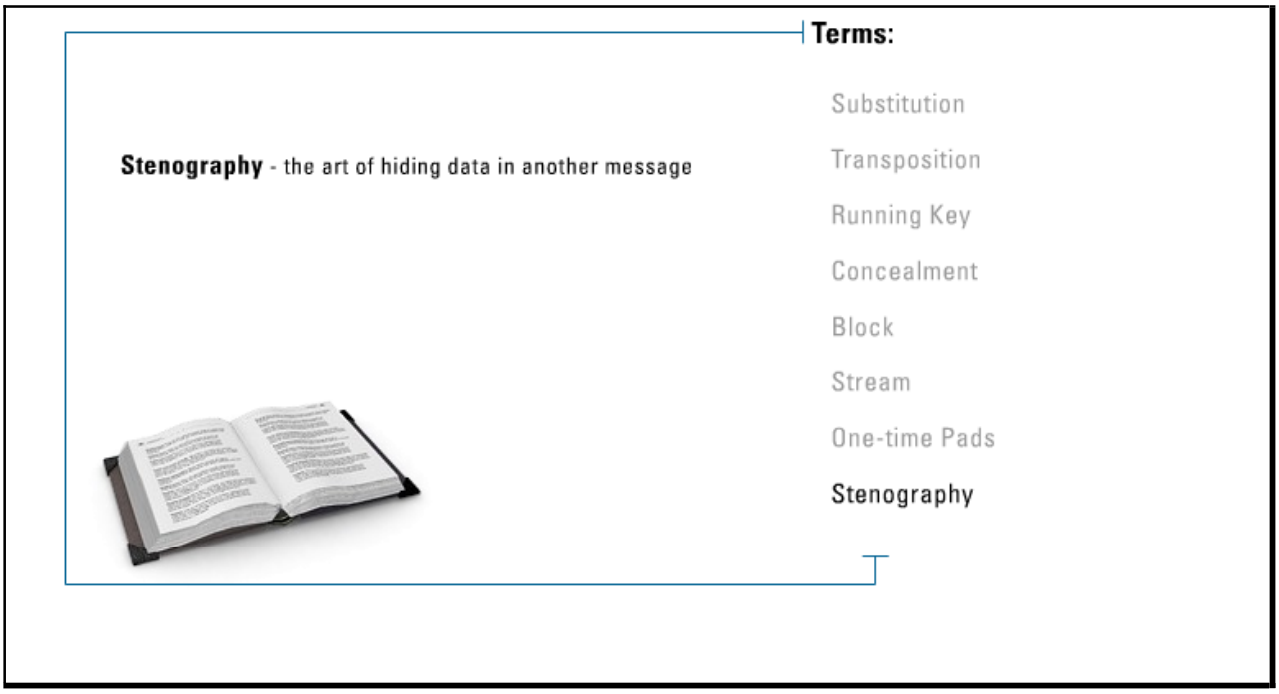

There are many types of ciphers in use today including:

- **Substitution Cipher** - Replaces bits, characters, or blocks of characters with different bits, characters, or blocks
- **Transposition Cipher** - Permutation is used, meaning that letters are scrambled; the key determines the positions to which the characters are moved
- **Running Key Cipher** - Uses steps in the physical world around us, like books (page, line number, and word count); each word is described by a sequence of numbers
- **Concealment Cipher** - Every X number of words within a text is part of the real message
- **Block Cipher** - A method of encrypting text (to produce ciphertext) in which a cryptographic key and algorithm are applied to a block of data (for example, 64 contiguous bits) at once as a group rather than to one bit at a time
- **Stream Cipher** - A method of encrypting text (to produce ciphertext) in which a cryptographic key and algorithm are applied to each binary digit in a data stream, one bit at a time

A **one-time pad** uses a keystream string of bits that is generated in a completely random fashion. The keystream will be the same length as the plaintext message. Because the entire keystream is completely random and used only once, a one-time pad is said to have perfect secrecy (one that is unable to be defeated by brute-force attacks).

Stenography is the art of hiding data in another message so that the very existence of the data is concealed. You can hide a message in a wave file, in a graphic, or in unused space on a hard drive or on sectors that are marked as unusable.

The National Security Agency (NSA) designed the **Clipper Chip**, which is a tamperproof chip for encrypting data. It uses the SkipJack encryption algorithm. Each Clipper Chip has a unique serial number, and a copy of the unit key is stored in the database under this serial number. The sending Clipper Chip generates a Law Enforcement Access Field (LEAF) value and includes it in the transmitted message. The Clipper Chip is based on an 80-bit key and a 16-bit checksum.

Cryptanalysis

Cryptanalysis refers to the study of ciphers, ciphertext, and cryptosystems in order to find weaknesses in them that will permit retrieval of the plaintext from the ciphertext, without necessarily knowing the key or the algorithm. The main method of cryptanalysis is **frequency analysis**, which is the analysis of the frequent patterns of letters used in messages and conversation.

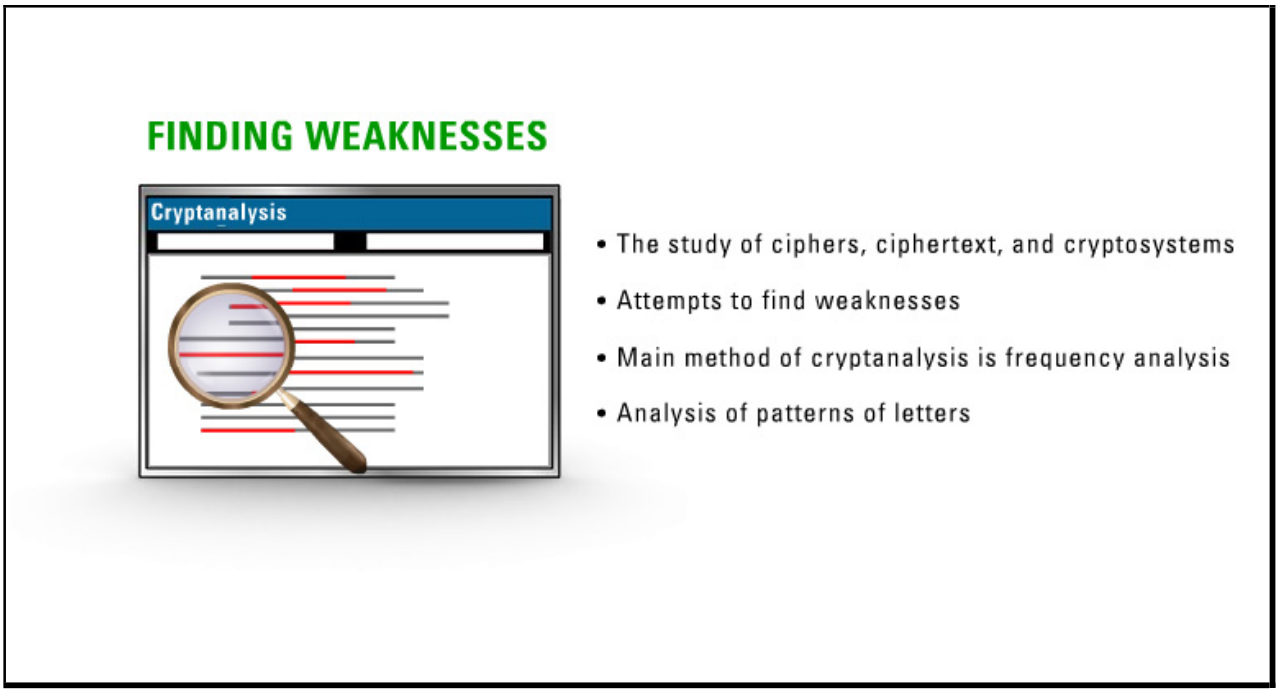

Frequency analysis is the basic tool for breaking classical ciphers in natural languages. Certain letters of the alphabet appear more frequently than others. For example, in English, "E" is likely to be the most common letter in any given sample of text. Similarly, the digraph "TH" is the most likely pair of letters, and so on. Frequency analysis relies on a cipher failing to hide these statistics. For example, in a simple substitution cipher (where each letter is simply replaced with another), the most frequent letter in the ciphertext would be a likely candidate for "E".

Note	Frequency analysis relies as much on linguistic knowledge as it does on statistics, but as ciphers became more complex, mathematics gradually became the predominant approach to cryptanalysis.

Symmetric Algorithms

Symmetric key encryption is an encryption system in which the sender and receiver of a message share a single, common key that they use to encrypt and decrypt the message, respectively.

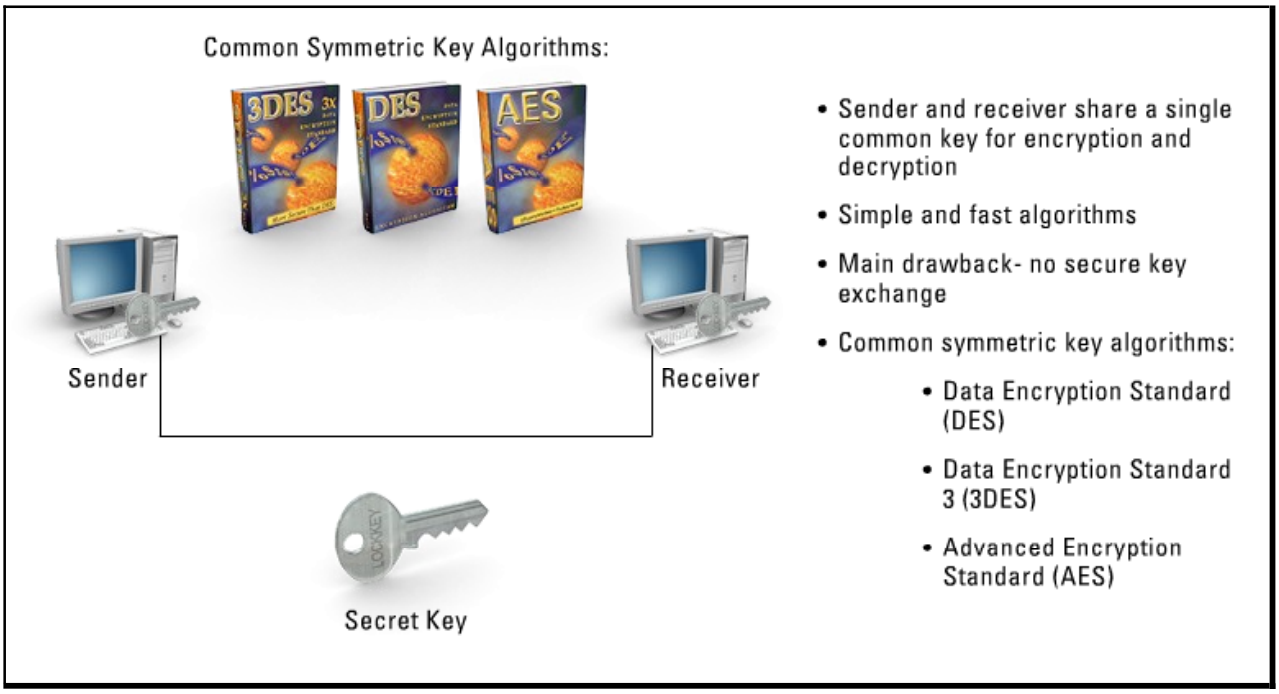

Symmetric-key systems are simple and fast, but their main drawback is that the two parties must somehow exchange the key in a secure way.

Note	Symmetric-key cryptography is sometimes called secret-key cryptography.

Common symmetric key algorithms include:

- Data Encryption Standard (DES)
- Data Encryption Standard 3 (3DES)
- Advanced Encryption Standard (AES)

These algorithms have withstood the test of time as cryptographers have attempted to crack the code and look for weaknesses to exploit. In order to break a symmetric algorithm, normally the attack is applied against the shared secret key and not the algorithm itself.

Other common symmetric algorithms include International Data Encryption Algorithm (IDEA), RC5, RC6, Blowfish, and CAST.

Data Encryption Standard

Data Encryption Standard(DES) is an algorithm certified by the National Institute of Standards and Technology (NIST), and was based on IBM's 128-bit Lucifer algorithm. DES is a block cipher—meaning it operates on plaintext blocks of a given size (64-bits) and returns ciphertext blocks of the same size. Thus, DES results in a permutation among the 2^{64} (read as "two to the 64th power") possible arrangements of 64 bits.

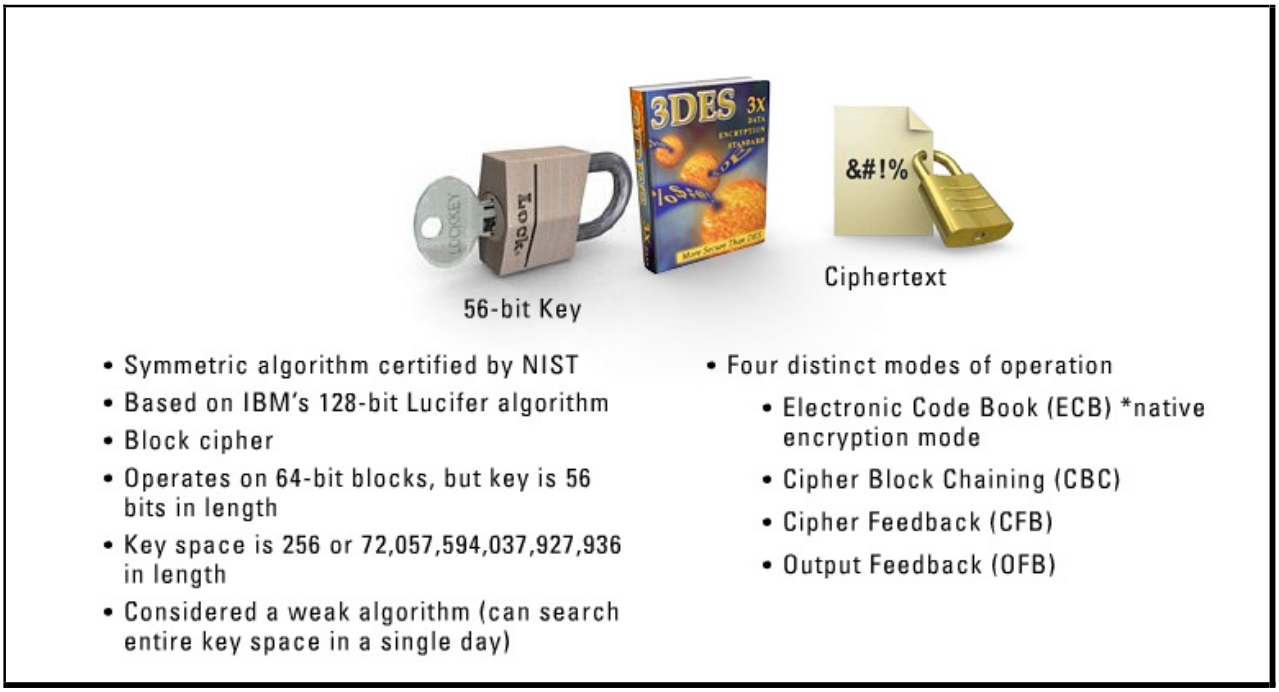

DES operates on the 64-bit blocks using key sizes of 56 bits. The keys are actually stored as being 64 bits long, but every eighth bit in the key is not used (i.e., bits numbered eight, 16, 24, 32, 40, 48, 56, and 64).

Because DES operates using a 56-bit key, the size of the key space is exactly 2^{56} or 72,057,594,037,927,936 possible values. Seventy-two quadrillion was a very large number in 1977 and would have taken computers back then hundreds of years to search the DES key space. In fact, it was considered so secure, the U.S. Department of Defense adopted it as a standard and restricted its exportation.

In today's computing environment, DES is considered a very weak algorithm. Searching the 72 quadrillion key space can be done in a relatively short time with modern computers. In 1999, the Electronics Frontier Foundation broke a DES key in less than one day using specially designed equipment. Even though, DES is still in wide use today, as it is a fast algorithm that provides reasonably secure transmission of everyday information.

DES has four distinct modes of operation:

- **ECB Mode/Electronic Code Book** - Native encryption mode. Provides the recipe of substitutions and permutations that will be performed on the block of plaintext. Data within a file does not have to be encrypted in a certain order. Used for small amounts of data, like challenge-response and key management tasks. Also used to encrypt PINs in ATM machines.

- **CBC Mode/Cipher Block Chaining -** Each block of text, the key, and the value based on the previous block are processed in the algorithm and applied to the next block of text.
- **CFB Mode/Cipher Feedback -** The previously generated ciphertext from the last encrypted block of data is inputted into the algorithm to generate random values. These random values are processed with the current block of plaintext to create ciphertext. This mode is used when encrypting individual characters is required.
- **OFB Mode/Output Feedback -** Functions like a stream cipher by generating a stream of random binary bits to be combined with the plaintext to create ciphertext. The ciphertext is fed back to the algorithm to form a portion of the next input to encrypt the next stream of bits.

Data Encryption Standard 3

As DES became more and more vulnerable, the Internet community decided a more secure algorithm was required. Because DES was normally based in hardware, a completely new algorithm was out of the question. As a result, **Data Encryption Standard 3 (3DES)** was created.

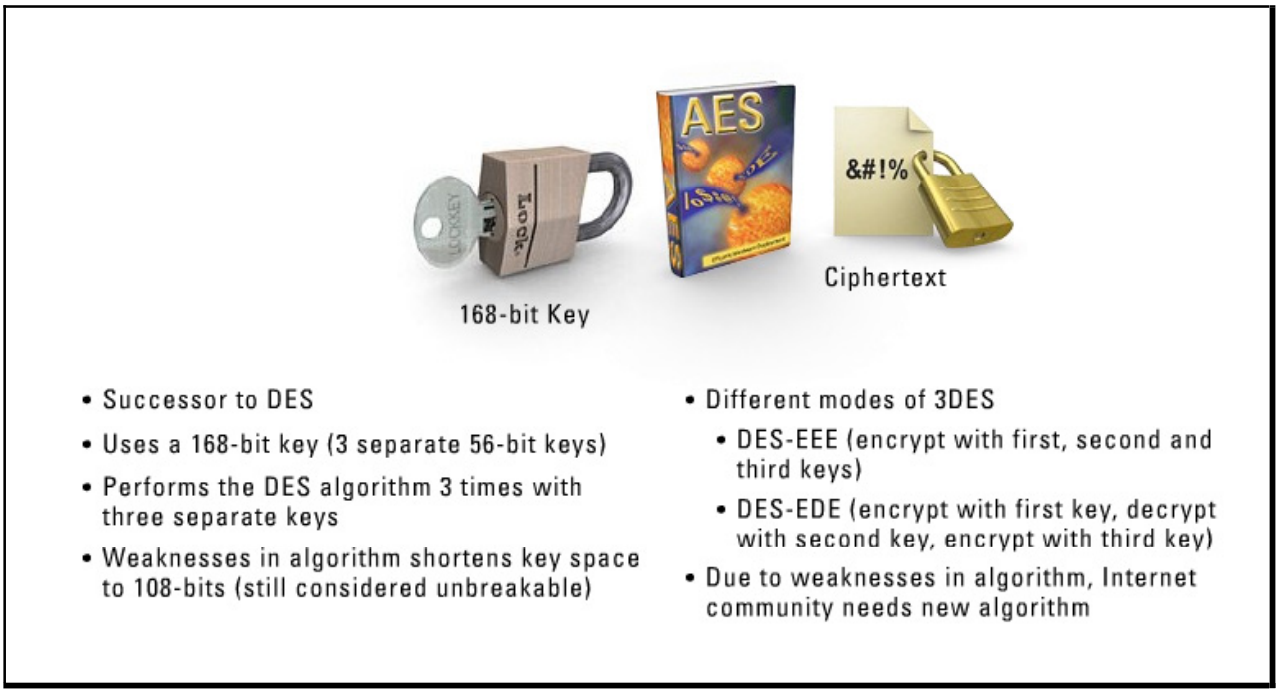

3DES uses a 168-bit key (actually it uses three 56-bit keys). In essence, the 3DES algorithm encrypts/decrypts data three times with three different keys, effectively creating a 168-bit key. But due to weaknesses in the algorithm, cryptographers discovered that they could apply shortcuts, which would bring the "useable" key space equal to approximately a 108-bit key space. Using 108 bits (2^{108}) produces an incredibly large key space and to this day, no one has successfully broken a 3DES key. 3DES uses 48 rounds in its computation and takes a heavy performance hit, as it can take up to three times longer than DES to perform encryption and decryption.

3DES uses the same basic machinery of DES three times over using three keys: k_1, k_2, and k_3. The plaintext (M) is encrypted using k_1. This result is encrypted with k_2 and the result is then encrypted with k_3 to get ciphertext (C).The equation is as follows:$C = E_{k3}(E_{k2}(E_{k1}(M)))$. This mode of using 3DES is called the DES-EEE mode since all three keys run in the encryption mode. The other mode is called DES-EDE, in which the second stage is run in decryption mode, i.e., $C = E_{k3}(D_{k2}(E_{k1}(M)))$.

3DES advantages:

- It is easy to implement in both hardware and software compared to other algorithms.
- It is based on DES, which is a very trusted cipher. DES has been studied thoroughly for over 25 years now and is proven to have sound basics. Although, the key length is too small now.
- It is much faster than public key cryptography methods like the RSA method. This is one of the main advantages of using a system like 3DES.

3DES disadvantages:

- Newer algorithms like RC6 and Blowfish are much faster than 3DES since they were built much later and with performance as an objective.
- The transmission of the secret key between users is unsafe. This is where public key cryptography excels.
- The new AES standard has been specified so most systems will likely shift to AES soon.

Due to the weaknesses in the 3DES algorithm and the computational overhead it required, the Internet community needed a new symmetric algorithm. In 1997, NIST sponsored a contest to see who the successor of DES would be. The winner of the contest would be judged based on speed and security of the algorithm. To qualify the winner was to give up all intellectual property of the algorithm.

Advanced Encryption Standard

The winner of the NIST contest was an algorithm named Rijndael, which was created by Joan Daemen and Vincent Jijmen. Rijndael, now renamed the **Advanced Encryption Standard (AES)**, is a variable block length and key length cipher. The number of rounds, or iterations of the main algorithm, can vary from 10 to 14 and depends on the block size and key length.

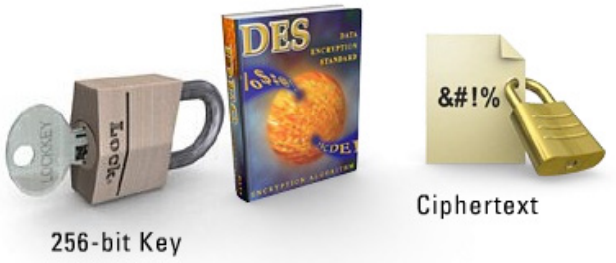

- Winner of NIST contest
- Originally named Rijndael
- Variable block length and key length cipher
- Current key lengths are 128, 192, or 256 bits
  - 128-bit key recommended
  - 256-bit key considered unbreakable

Current AES key lengths are those of 128, 192, or 256 bits used to encrypt blocks with lengths of 128, 192, or 256 bits, respectively. You can implement AES very efficiently on a wide range of processors and in hardware.

Today, key lengths of 128 bits are recommended, but key lengths of 192 or 256 bits provide for the utmost security now and in the future. To put a 256-bit key length in perspective, IBM's Blue Gene/C supercomputer is capable of achieving 1,000 teraflops, or 1,000 trillion calculations per second. For the sake of argument, assume that it takes 20 calculations to check a single key. Therefore, Blue Gene/C would be able to check 50 trillion (1,000 trillion divided by 20) keys per second. Now assume you have 50 trillion Blue Gene/C supercomputers, each checking 50 trillion keys per second. It would take all supercomputers working in concert almost 1.5×10^{28} trillion years to search just one percent of the entire key space!

Other Symmetric Encryption Algorithms

This topic discusses other symmetric encryption algorithms.

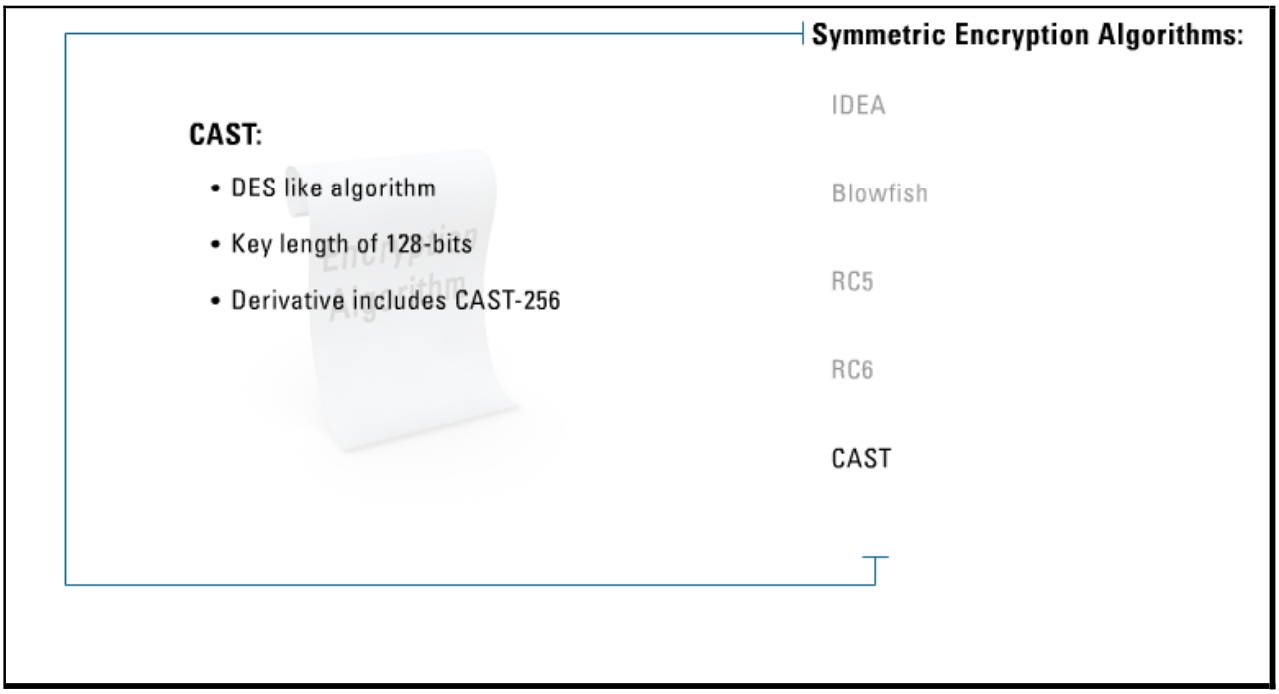

The **International Data Encryption Algorithm (IDEA)** is a block cipher that operates on 64-bit blocks of data. The key is 128 bits long. The 64-bit data block is divided into four 16-bit smaller blocks and each has eight rounds of mathematical functions performed on it. IDEA is used in the PGP encryption software. Many people consider IDEA to be superior to DES simply because of its larger 128-bit key size.

Blowfish is a block cipher that works on 64-bit blocks of data. The key length can be up to 448 bits and the data blocks go through 16 rounds of cryptographic functions.

RC5 is a very fast block cipher designed by Ron Rivest. It has a variety of parameters it can use for block size, key size, and the number of rounds used. Block sizes are 32 bits, 64 bits, and 128 bits, and key size is up to 2048 bits. Speed tests show a 50 percent increase over DES, using highly optimized versions of both algorithms.

RC6 is based on the RC5 block cipher algorithm and, just like RC5, it is a parameterized algorithm where the block size, the key size, and the number of rounds used are variable. The maximum size of the key is 2040 bits, which should certainly make it strong for quite a few years to come. Two new features added to RC6 include the inclusion of integer multiplication and the use of four four-bit working registers instead of two two-bit working registers.

CAST is a 128-bit encryption algorithm. It uses a DES-like Substitution-Permutation Network (SPN) cryptosystem, which appears to have good resistance to differential cryptanalysis. CAST-128 belongs to the class of encryption algorithms known as Feistel ciphers; thus, the overall operation of CAST is similar to the operation of DES. Cast uses a pair of subkeys per round: A 32-bit quantity key is used for masking and a five-bit quantity key is used as a “rotation” key. CAST has also been designed in a 256-bit mode, which has three different round functions as compared to its 128-bit cousin. CAST-256 has a block size of 128 bits and a variable key size (128, 160, 192, 224, or 256 bits).

Asymmetric Algorithms

Asymmetric algorithms, often called **public key algorithms**, perform encryption and decryption in a completely different way than symmetric algorithms.

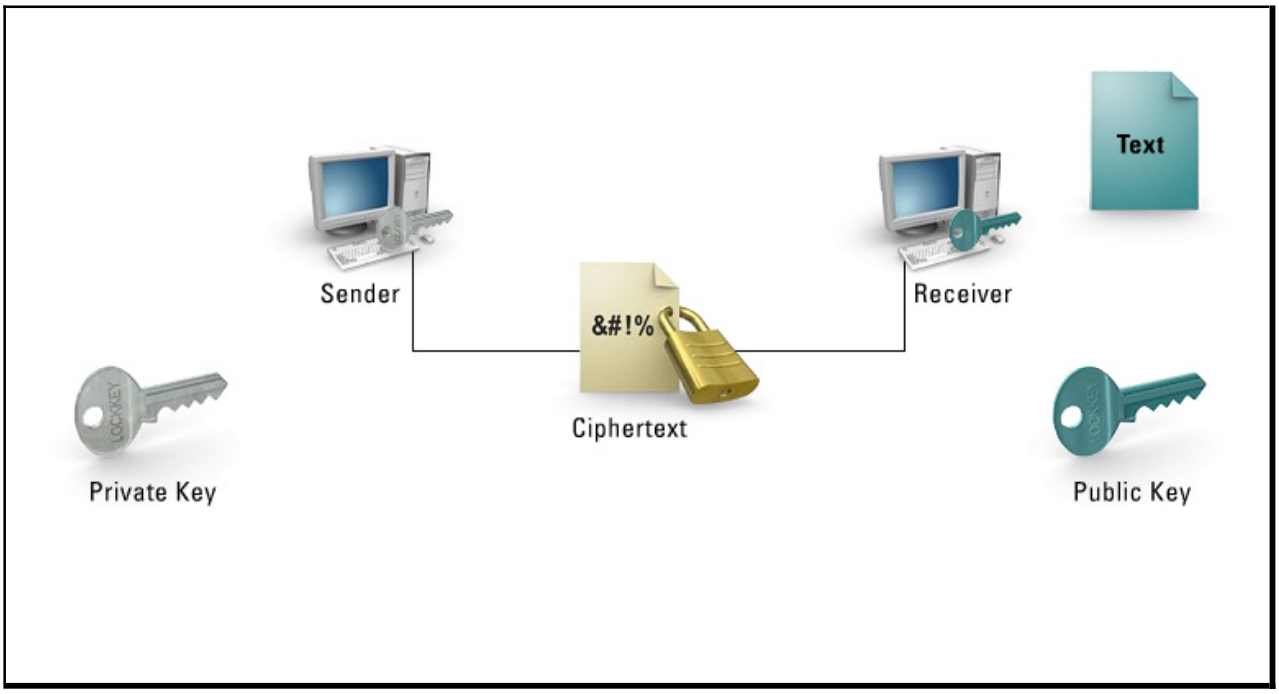

Asymmetric algorithms do not rely on a randomly generated shared encryption key that changes per session; instead, they create two static keys that never change. These static keys are completely different but mathematically bound to each other in the sense that one key decrypts what the other key encrypts. One key alone cannot encrypt and decrypt the same data.

This encryption method works by keeping one key private and giving the other key to anyone in the public Internet. Anyone can have the public key as it is useless without the private key.

For example, imagine peer X generates a public and private key pair, and then encrypts a message with its private key. Peer X then sends the ciphertext to peer Y. Peer Y obtains peer X's public key via some mechanism and then decrypts the message sent from peer X. This process may seem flawed because anyone who sniffs the wire and obtains peer X's public key can read the message. However, if peer Y obtains peer X's public key and encrypts a message with it, only peer X can decrypt the message.

The main problem with asymmetric algorithms is the fact that they are very slow. The reason they are so slow can be attributed to the fact that they all use very heavy mathematics to perform their functions. It is not practical to encrypt bulk data with asymmetric algorithms, but you can still use them to encrypt/decrypt small amounts of data, such as a hash value. Depending on which key you use to encrypt messages, asymmetric key algorithms can perform the following functions:

- **Secure Message Format** - Encrypted by the receiver's public key
- **Open Message Format** - Encrypted by the sender's private key
- **Secure and Signed Format** - Encrypted by the sender's private key and then encrypted with the receiver's public key

Asymmetric algorithms include the following:

- RSA
- DSA
- Diffie-Hellman

RSA Asymmetric Algorithm

Ronald Rivest, Adi Shamir, and Leonard Adleman developed the **RSA asymmetric algorithm** in 1977. RSA stands for the first letter of each of its inventors' last names.

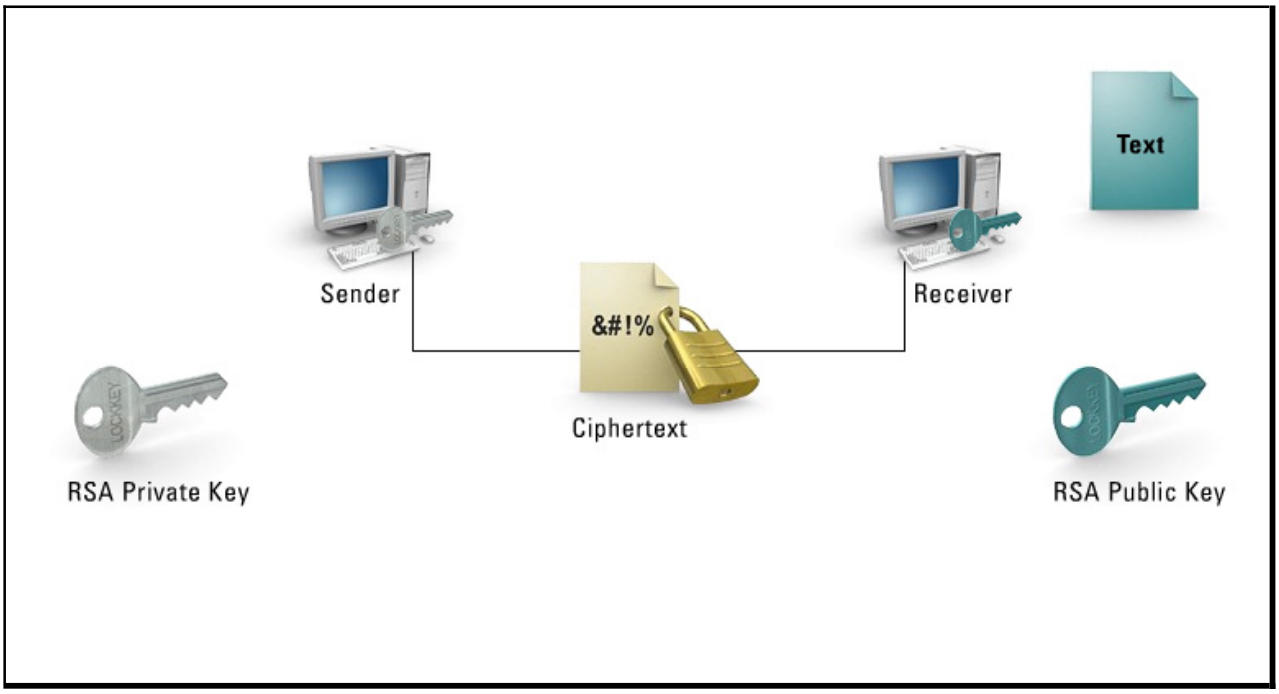

The math behind the RSA algorithm works as follows:

"Take two large prime numbers, p and q, and compute their product $n = pq$; n is called the modulus. Choose a number, e, less than n and relatively prime to $(p-1)(q-1)$, which means e and $(p-1)(q-1)$ have no common factors except 1. Find another number d such that $(ed-1)$ is divisible by $(p-1)(q-1)$. The values e and d are called the public and private exponents, respectively. The public key is the pair (n, e); the private key is (n, d). The factors p and q may be destroyed or kept with the private key."

Based on what is known today, it is very difficult for anyone to attempt to obtain the private key d from the public key (n, e). If someone could factor n into p and q, then one could obtain the private key d. This, however, is a very difficult thing to do, which is the basis for the security behind the RSA algorithm. A typical key size for RSA is 1024 bits.

The RSA algorithm is used in many Web browsers with Secure Sockets Layer (SSL), in Pretty Good Privacy (PGP) and government systems that use public key cryptosystems, and, of course, with IP Security (IPSec). You can use the RSA algorithm with IPSec for two discreet purposes:

- **Encryption** – Here, peer X uses peer Y's public key to encrypt data and then sends the data back to peer Y. Since only peer Y has the corresponding public key, peer Y can successfully decrypt the data.
- **Digital Signatures** – Here, peer X encrypts a hash value with a private key and then sends the data to peer Y. Peer Y obtains peer X's public key and decrypts the ciphertext to obtain the hash. Since peer Y used peer X's public key, only peer X could have encrypted the hash, hence, the encrypted hash must have come from peer X.

Digital Signature Standard

NIST created the **Digital Signature Standard (DSS)** in 1994. It specifies the **Digital Signature Algorithm** (DSA) as the algorithm for digital signatures.

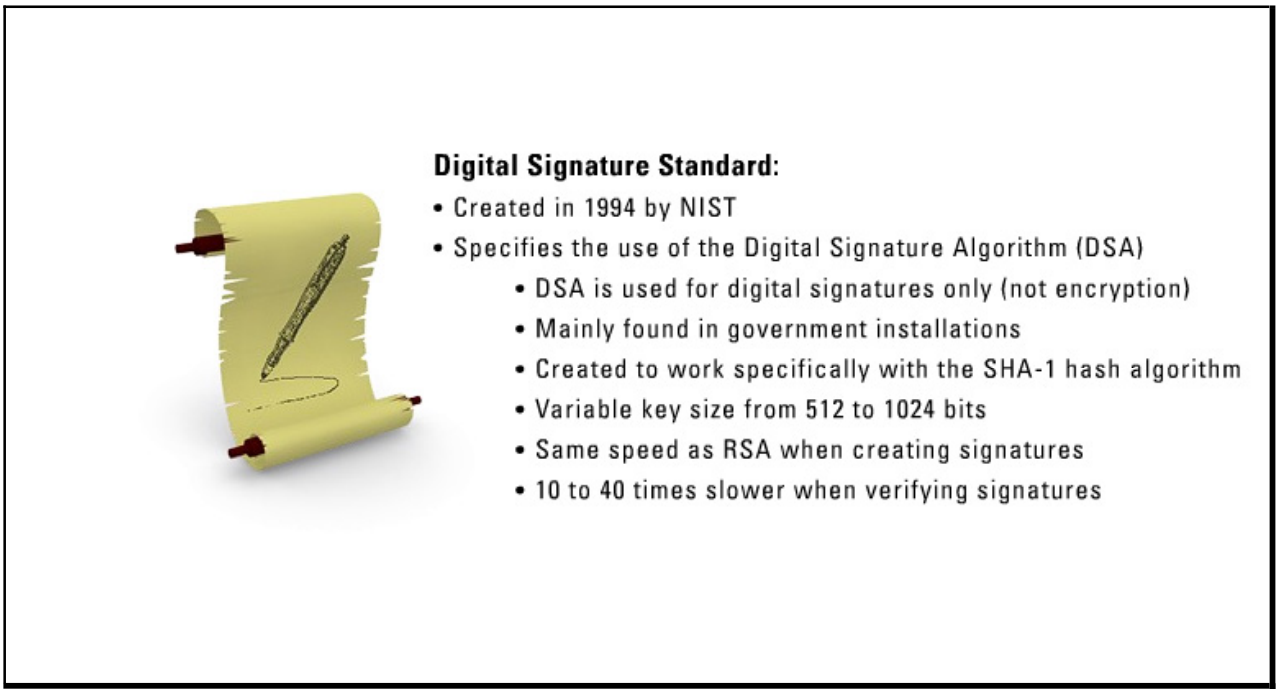

DSA is mainly found in government installations, and has been created to work specifically with the SHA-1 hash algorithm. DSA is for digital signatures only. It is not used for encryption, as is RSA.

DSA is a public key algorithm; the secret key operates on the message hash generated by SHA-1. To verify a signature, one re-computes the hash of the message, uses the public key to decrypt the signature, and then computes the results. The key size is variable and can be from 512 to 1024 bits. Using more than 768 bits is adequate for current computing capabilities.

DSA moves at roughly the same speed as RSA when creating signatures, but is 10 to 40 times slower when it comes to verifying signatures. Since verification takes place more frequently than creation, this is an issue worth noting when deploying DSA in any environment.

Diffie-Hellman

Whitfield Diffie and Martin Hellman created the **Diffie-Hellman (DH)** asymmetric algorithm in 1976. It was the first asymmetric algorithm ever created.

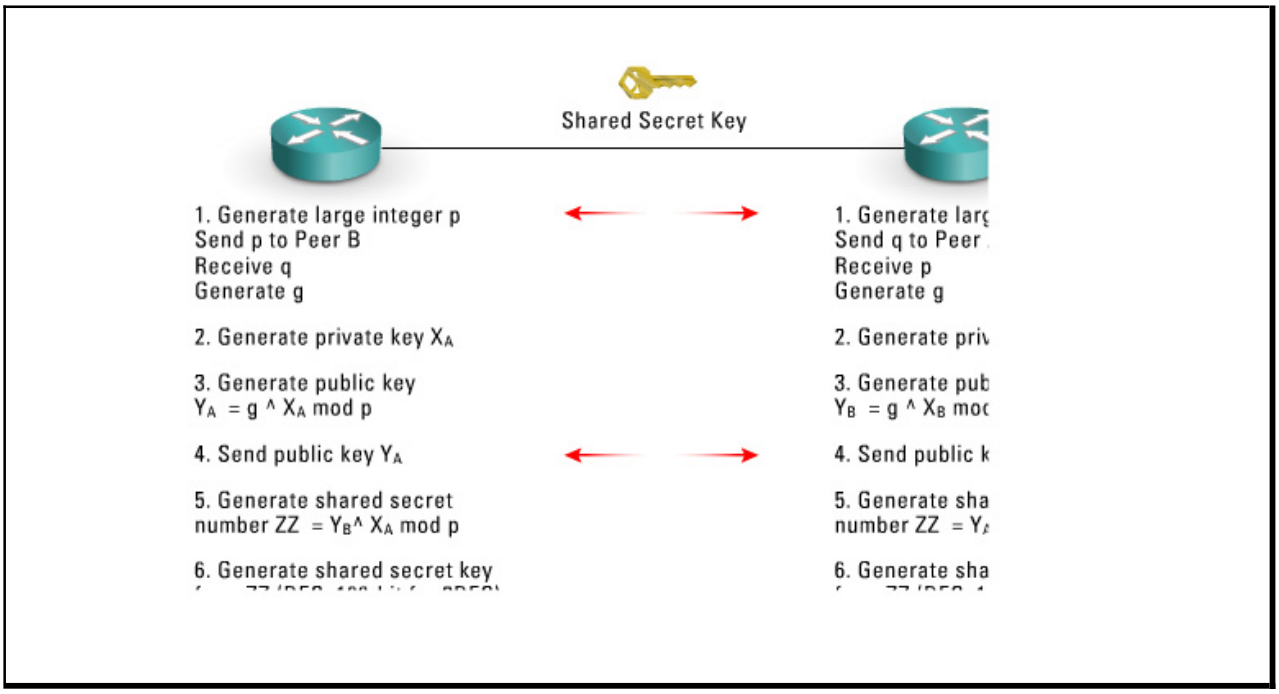

DH is not used for encryption or digital signatures, but is instead used to obtain a shared secret "key agreement" between two parties over an insecure medium such as the Internet. Each party can then use the shared secret key to encrypt bulk data using a symmetric key algorithm. DH works by sending large mathematical numbers over the Internet. No one can mathematically obtain the shared secret key even if they obtain the numbers being sent through the Internet. Only the two ends of the exchange using the DH algorithm can compute the shared secret key. The math for the algorithm is as follows:

"Suppose Alice and Bob want to agree on a shared secret using the DH key agreement protocol. They proceed as follows: First, Alice generates a random private value a, and Bob generates a random private value b. Both a and b are drawn from the set of integers $\{1, \ldots, p\text{-}2\}$. Then they derive their public values using parameters p and g and their private values. Alice's public value is g^a mode p and Bob's public value is g^b mode p, and Bob computes $g^{ba} = (g^a)^b$ mode p. Since $g^{ab} = g^{ba} = k$, Alice and Bob now have a shared secret key k."

The Diffie-Hellman key exchange is vulnerable to a man-in-the-middle attack. To avoid this potential problem, the two parties can authenticate themselves to each other by using a shared secret key, digital signatures, or public-key certificates.

When two systems need to create a shared secret between them, they can use the services of DH to obtain it. Many services need shared secret keys. The problem IPSec has with DH is that DH is computationally expensive. Too many CPU cycles are used to create all the shared secret keys needed. Therefore, to reduce the number of DH exchanges, IPSec will perform DH a single time and a number of shared keys will be derived from the original. These derived keys are identical on both sides and tagged such that all possible mechanisms that need a shared key will have one.

For example, the DH key (prime) will be *k*:

- *K1* (derived from *k*) used for process *a*
- *K2* (derived from *k*) used for process *b*
- *K3* (derived from *k*) used for process *c*
- *K4* (derived from *k*) used for process *d*
- *K5* (derived from *k*) used for process *e*

Elliptic Curve Cryptography

Elliptic curve cryptography (ECC) studies elliptical curves. Elliptical curves are simple functions that can be drawn as gently looping lines in the (X,Y) plane. The study of these elliptical curves has brought about some very interesting algorithms in the study of cryptography.

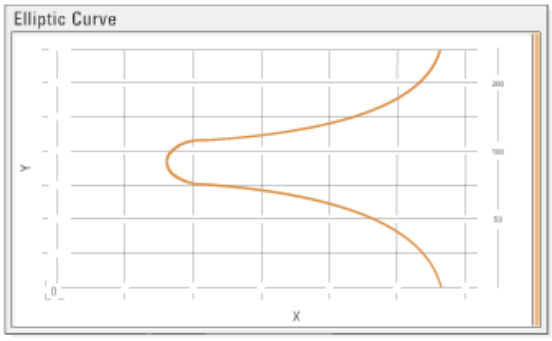

- The study of cryptography and elliptical curves
- Simple functions drawn as gently looping lines in the (X,Y) plane
- Most public key algorithms (except RSA) rely on discrete log problem
- Discrete logs are becoming easier
    - Larger key sizes are required to provide necessary security
- Elliptic curve discrete log techniques have not been improved upon

Normally, public key cryptography has based its strength on the computational difficultly of someone solving two problems: integer factoring, and discrete logarithms over finite fields. These two problems (factoring and discrete logarithms) appear to be very computationally difficult. But, there are subexponentialalgorithms for both problems. These algorithms use both the additive and multplicative structure of the fields (or rings) in which they operate.

Just about all public key cryptography algorithms, except RSA, rely on the difficulty of a discrete log problem. As discrete logs become easier, longer bit-lengths are required to provide the necessary security. Discrete logs are now much easier thanks to the discovery of what is called Number Field Sieve (NFS) techniques. But, elliptic curve discrete log techniques have not seen significant improvement in the past 20 years.

Man-in-the-Middle Attack

A **man-in-the-middle** (**MITM**) **attack** is an attack in which the attacker is able to read, and modify at will, messages between two parties without either party knowing that the link between them has been compromised. In order for this type of attack to succeed, the attacker must be in the direct line of communication between the two victims.

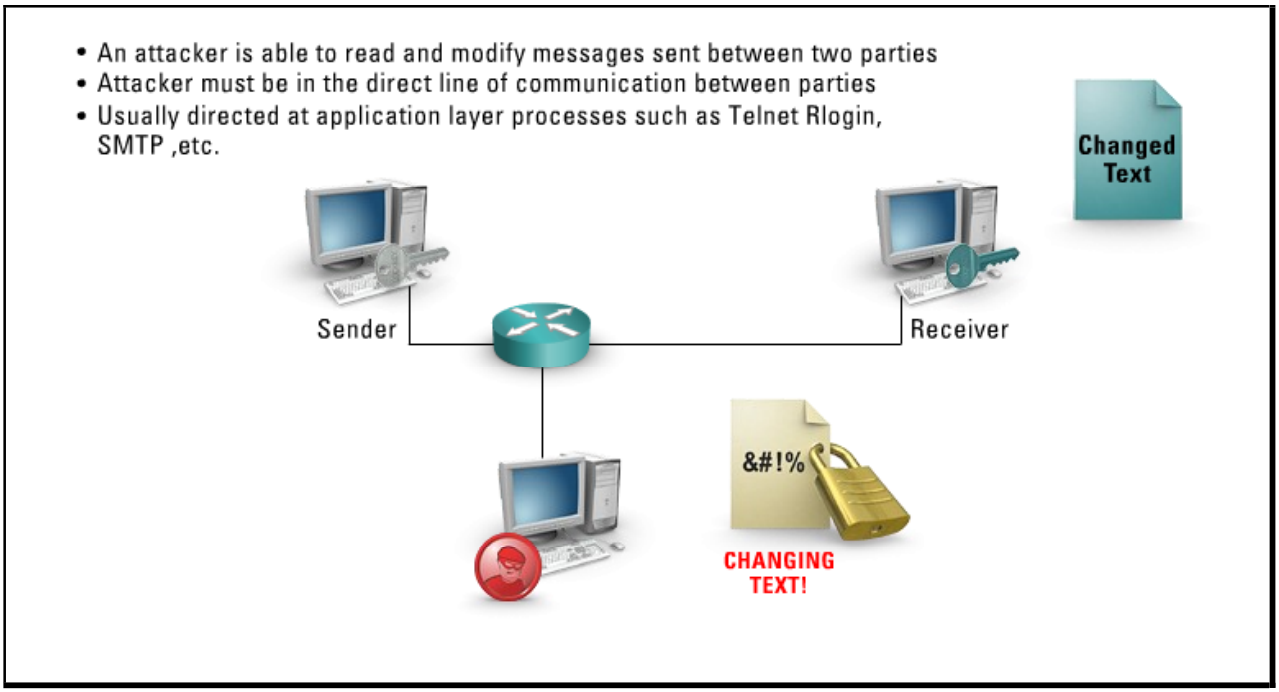

MITM attacks are usually directed at application layer processes such as Telnet, rlogin, Simple Mail Transfer Protocol (SMTP), File Transfer Protocol (FTP), Hypertext Transport Protocol (HTTP), and others. They allow an attacker to intercept confidential information, modify the data, and then send that data to the receiving party.

For example, if an attacker can obtain access to any router in the path between X and Y, the attacker can divert traffic between these victims to his or her local workstation. Once the attacker has the data, he or she can instruct his or her script to insert or remove data, and then send the data to the indended recipient. In this way, X and Y have no way of knowing that their data have been tampered with.

Block Ciphers

This topic discusses block ciphers.

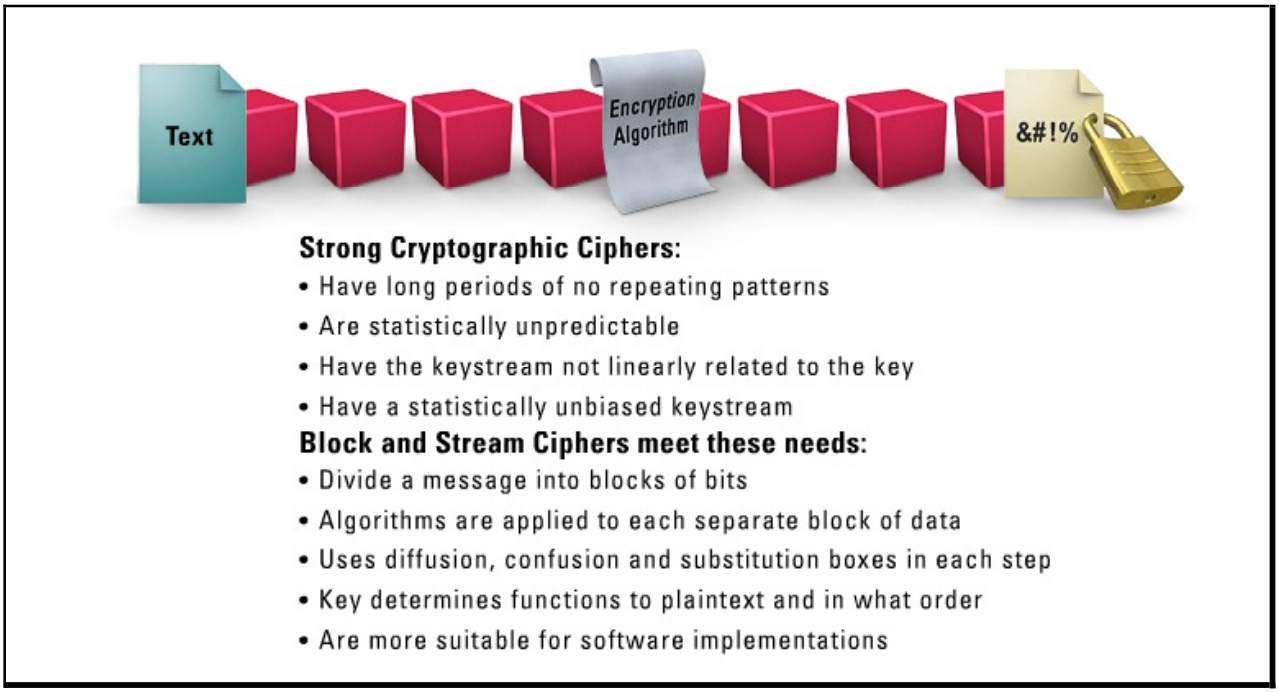

All cryptographic cipher algorithms need to have the following attributes to be considered strong:

- Long periods of no repeating patterns
- Statistically unpredictable
- The keystream is not linearly related to the key
- Statistically unbiased keystream (as many 0's as 1's)

To this end, two predominant cipher technologies have been created and are in general use: block ciphers and stream ciphers.

In **block ciphers**, the message is divided into blocks of bits. Block ciphers use diffusion and confusion in their methods as well as Substitution boxes (S-boxes) in each step. It is the key that determines what functions are applied to the plaintext and in what order. Block ciphers are more suitable for software implementations because they work with blocks of data, which are usually the width of a data bus (64 bits). Block ciphers sometimes work in a mode that emulates a stream cipher. The following terms are important when discussing block cipher algorithms:

- **Confusion** - Different unknown key values are used
- **Diffusion** - Putting the bits within the plaintext through many different functions so that they are dispersed throughout the algorithm
- **S-Box** - Contains a lookup table that instructs how the bits should be permuted or moved around; the key that is used in the decryption process dictates what S-boxes are used and in what order.

Stream Ciphers

A **stream cipher** is a type of symmetric encryption algorithm. You can design a stream cipher to be exceptionally faster than any block cipher.

- Exceptionally faster that block ciphers
- Treat the message as a stream of bits
- Key is input in the stream cipher (ensures randomness)
- More suitable for hardware implementations
- Generates keystream (a sequence of bits used as a key)
- Encryption is accomplished usually using a bitwise XOR
- Developed to approximate the action of the one-time pad

Stream ciphers treat the message as a stream of bits or bytes and perform mathematical functions on them individually. The key is a random value input into the stream cipher, which it uses to ensure the randomness of the keystream data. Stream ciphers are more suitable for hardware implementations because they encrypt and decrypt one bit at a time, but because they can be processor intensive (each bit must be manipulated), they usually work better at the silicon level.

A stream cipher generates what is called a keystream (a sequence of bits used as a key). Encryption is accomplished by combining the keystream with the plaintext, usually with the bitwise Exclusive OR (XOR) operation. The generation of the keystream can be independent of the plaintext and ciphertext, yielding what is termed a synchronous stream cipher, or it can depend on the data and their encryption, in which case the stream cipher is said to be self-synchronizing. Most stream cipher designs are for synchronous stream ciphers.

While block ciphers operate on large blocks of data, stream ciphers typically operate on smaller units of plaintext, usually bits. The encryption of any particular plaintext with a block cipher will result in the same ciphertext when the same key is used. With a stream cipher, the transformation of these smaller plaintext units will vary, depending on when they are encountered during the encryption process.

Stream ciphers were developed as an approximation to the action of the one-time pad. While contemporary stream ciphers are unable to provide the satisfying theoretical security of the one-time pad, they are at least practical.

One-Time Cipher Keys (Pads)

One-time pads are considered unbreakable, and each pad is used exactly once.

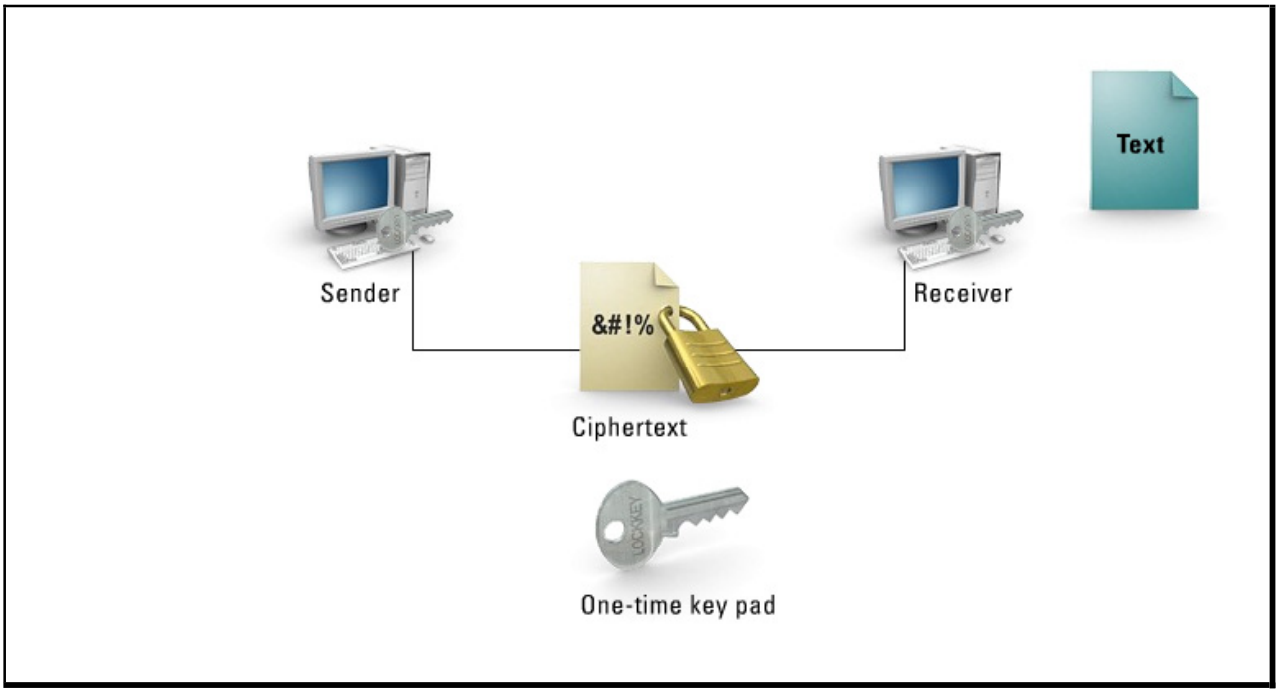

One-time pads use a truly non-repeating set of random bits that are combined in a bitwise XOR with the message to produce ciphertext. The random key is the same size as the message and is only used once. One-time pads are not completely without implementation difficulties, as it is very difficult to distribute the pads of random numbers to all the necessary parties.

Since the entire keystream is random, even an attacker with infinite computational resources can only guess the plaintext if he or she sees the ciphertext. Such a cipher is said to offer perfect secrecy, and the analysis of the one-time pad is seen as one of the cornerstones of modern cryptography.

Electronic Code Book

In the **Electronic Code Book(ECB)** mode, each plaintext block is encrypted independently with the block cipher.

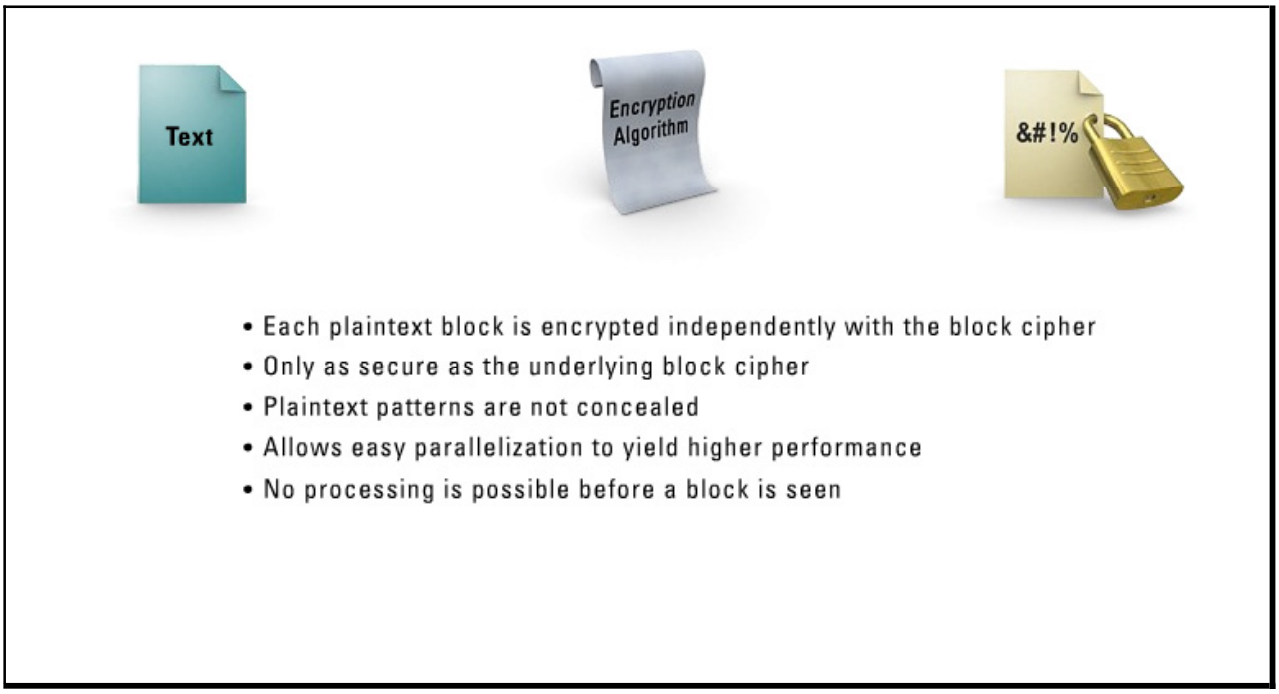

ECB mode is as secure as the underlying block cipher. However, plaintext patterns are not concealed. Each identical block of plaintext gives an identical block of ciphertext. Attackers can easily manipulate the plaintext by removing, repeating, or interchanging blocks. The speed of each encryption operation is identical to that of the block cipher. ECB allows easy parallelization to yield higher performance. Unfortunately, no processing is possible before a block is seen (except for key setup).

Cipher Block Chaining

In the **Cipher Block Chaining (CBC)** mode, each plaintext block is XOR'd with the previous ciphertext block and then encrypted. An initialization vector c_0 is used as a "seed" for the process.

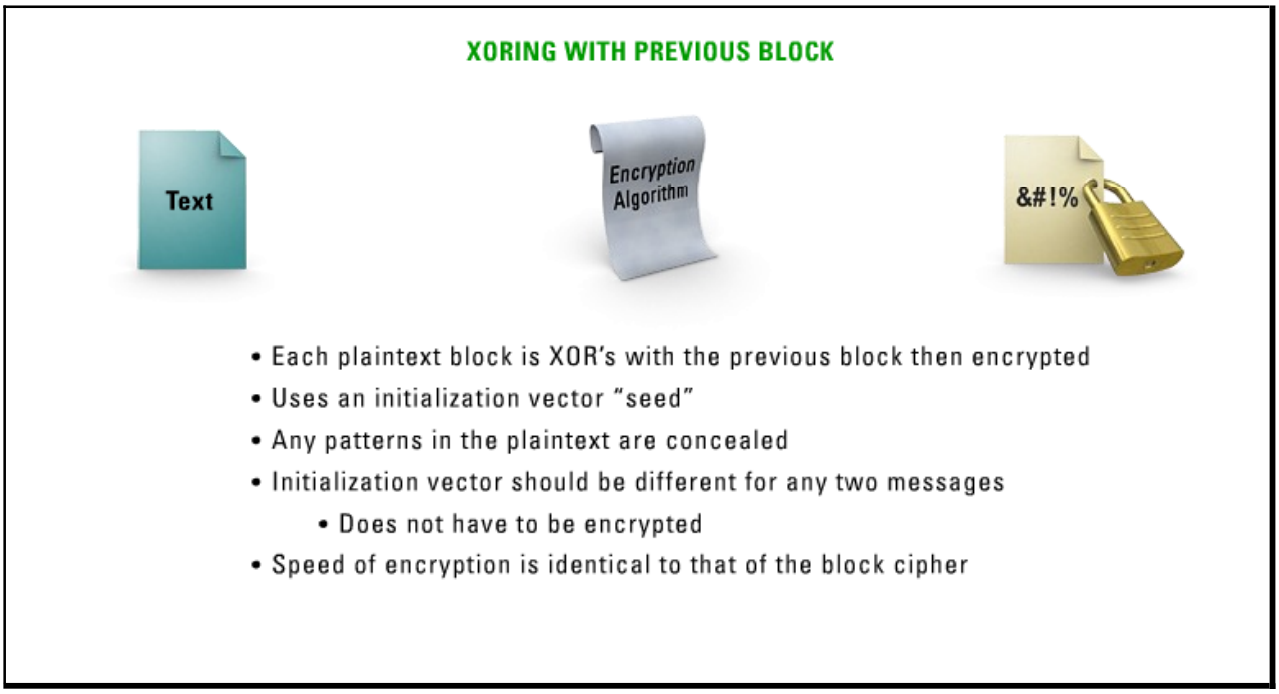

CBC mode is as secure as the underlying block cipher against standard attacks. In addition, any patterns in the plaintext are concealed by the XORing of the previous ciphertext block with the plaintext block. Note also that you cannot directly manipulate the plaintext except by removing blocks from the beginning or the end of the ciphertext. The initialization vector should be different for any two messages encrypted with the same key and is preferably randomly chosen. It does not have to be encrypted, and it can be transmitted with (or considered as the first part of) the ciphertext. The speed of encryption is identical to that of the block cipher, but the encryption process cannot be easily parallelized, although the decryption process can be.

Cipher Feedback (CFB)

In **Cipher Feedback(CFB)** mode, the previous ciphertext block is encrypted and the output produced is combined with the plaintext block using an XOR operation to produce the current ciphertext block. It is possible to define CFB mode so it uses feedback that is less than one full data block. An initialization vector c_0 is used as a "seed" for the process.

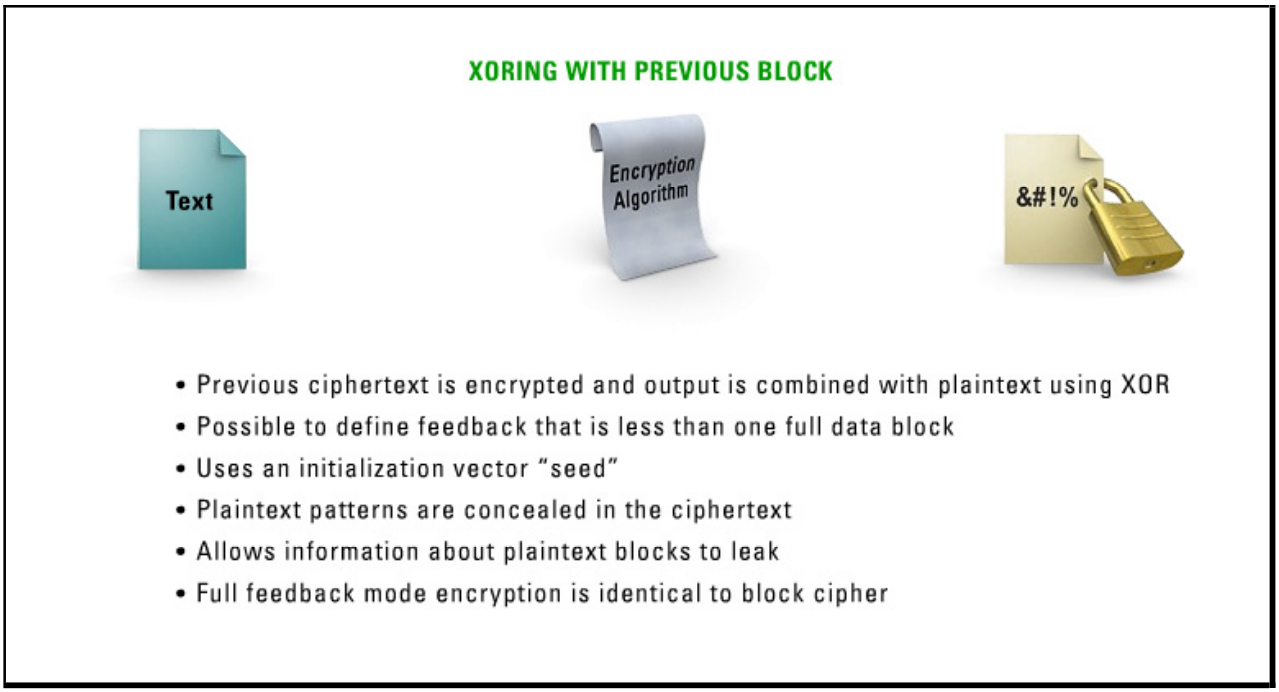

CFB mode is as secure as the underlying cipher, and plaintext patterns are concealed in the ciphertext by the use of the XOR operation. You cannot directly manipulate the plaintext except by removing blocks from the beginning or the end of the ciphertext. With CFB mode and full feedback, when two ciphertext blocks are identical, the outputs from the block cipher operation at the next step are also identical. This allows information about plaintext blocks to leak.

When using full feedback, the speed of encryption is identical to that of the block cipher, but the encryption process cannot be easily parallelized.

Output Feedback

Output Feedback (OFB) mode is similar to CFB mode except that the quantity XOR'd with each plaintext block is generated independently of both the plaintext and ciphertext. An initialization vector s_0 is used as a "seed" for a sequence of data blocks s_i, and each data block s_i is derived from the encryption of the previous data block s_{i-1}. The encryption of a plaintext block is derived by taking the XOR of the plaintext block with the relevant data block.

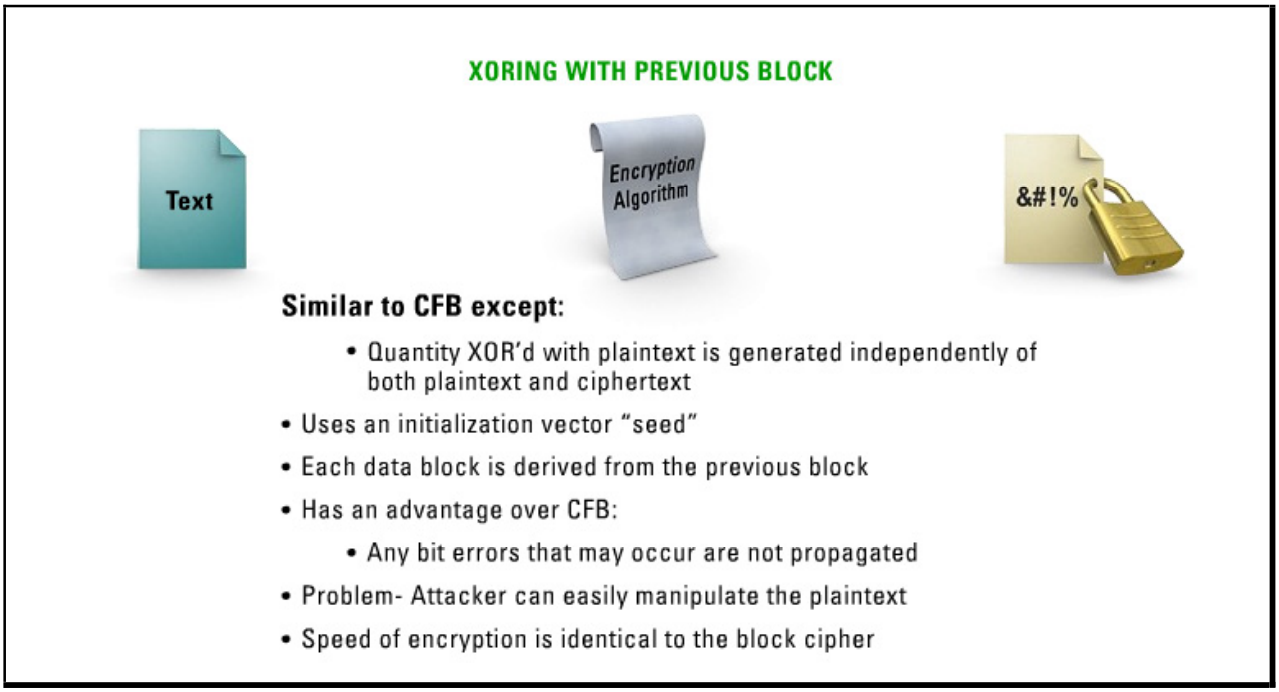

Feedback widths less than a full block are not recommended for security. OFB mode has an advantage over CFB mode in that any bit errors that might occur during transmission are not propagated to affect the decryption of subsequent blocks. The security considerations for the initialization vector are the same as in CFB mode.

A problem with OFB mode is that an attacker can easily manipulate the plaintext. Namely, an attacker who knows a plaintext block m_i may replace it with a false plaintext block x by XORing m_i Åx to the corresponding ciphertext block c_i. Similar attacks can take place in CBC and CFB modes, but in those attacks some plaintext block will be modified in a manner unpredictable by the attacker. Yet, the very first ciphertext block (that is, the initialization vector) in CBC mode and the very last ciphertext block in CFB mode are just as vulnerable to the attack as the blocks in OFB mode. You can prevent attacks of this kind by using, for example, a digital signature scheme.

The speed of encryption is identical to that of the block cipher. Even though the process cannot easily be parallelized, you can save time by generating the keystream before the data are available for encryption.

Attacking Encryption

Cryptanalytic attacks are generally classified into six categories that distinguish the kind of information the cryptanalyst has available to mount an attack. The categories of attack are listed here in increasing order of the quality of information available to the cryptanalyst, or, equivalently, in decreasing order of the level of difficulty to the cryptanalyst. The objective of the cryptanalyst in all cases is to be able to decrypt new pieces of ciphertext without additional information. The ideal for a cryptanalyst is to extract the secret key.

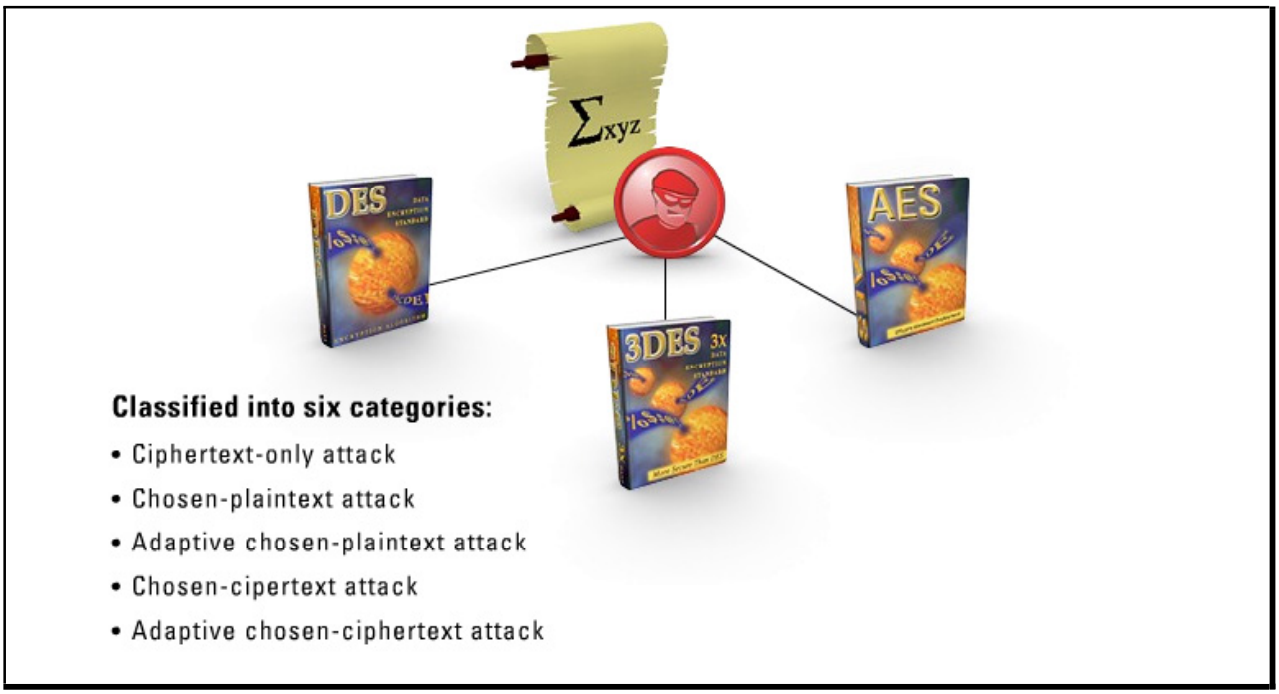

Ciphertext-Only Attack: The attacker obtains a sample of ciphertext, without the plaintext associated with it. These data are relatively easy to obtain in many scenarios, but a successful ciphertext-only attack is generally difficult, and requires a very large ciphertext sample.

Chosen-Plaintext Attack: The attacker has the plaintext and ciphertext and can choose the plaintext that gets encrypted.

Adaptive Chosen-Plaintext Attack: A special case of chosen-plaintext attack in which the cryptanalyst is able to choose plaintext samples dynamically, and alter his or her choices based on the results of previous encryptions.

Chosen-Ciphertext Attack: The attacker may choose a piece of ciphertext and attempt to obtain the corresponding decrypted plaintext. This type of attack is generally most applicable to public-key cryptosystems.

Adaptive Chosen-Ciphertext Attack: An attacker can mount an attack of this type in a scenario in which he or she has free use of a piece of decryption hardware, but is unable to extract the decryption key from it.

Known-Plaintext: The attacker has the plaintext and ciphertext of one or more messages.

Summary

The key points discussed in this lesson are:

- Cryptography is the art of garbling data so the data look nothing like their original form, and then being able to restore the data back to their original form at some future time.
- Encryption is the process of encoding data to ensure the data are accessible only by the intended recipient.
- A cipher is any cryptographic system in which arbitrary symbols, or groups of symbols, represent units of plaintext of regular length, usually single letters, or in which units of plaintext are rearranged, or both, in accordance with certain predetermined rules.
- Cryptanalysis refers to the study of ciphers, ciphertext, and cryptosystems in order to find weaknesses in them that will permit retrieval of the plaintext from the ciphertext, without necessarily knowing the key or the algorithm.
- Symmetric key encryption is an encryption system in which the sender and receiver of a message share a single, common key that they use to encrypt and decrypt the message, respectively.
- DES is a block cipher—meaning it operates on plaintext blocks of a given size (64-bits) and returns ciphertext blocks of the same size.
- The 3DES algorithm encrypts/decrypts data three times with three different keys, effectively creating a 168-bit key.
- The AES is a variable block length and key length cipher.
- Asymmetric algorithms do not rely on a randomly generated shared encryption key that changes per session; instead, they create two static keys that never change.
- The RSA algorithm is used in many Web browsers with SSL, in PGP and government systems that use public key cryptosystems, and, of course, with IPSec.
- NIST created the DSS in 1994. It specifies the DSA as the algorithm for digital signatures.
- DH is not used for encryption or digital signatures, but is instead used to obtain a shared secret "key agreement" between two parties over an insecure medium such as the Internet.
- ECC studies elliptical curves. Elliptical curves are simple functions that can be drawn as gently looping lines in the (X,Y) plane.
- A MITM attack is an attack in which the attacker is able to read, and modify at will, messages between two parties without either party knowing that the link between them has been compromised.
- In block ciphers, the message is divided into blocks of bits. Block ciphers use diffusion and confusion in their methods as well as S-boxes in each step.
- A stream cipher is a type of symmetric encryption algorithm. You can design a stream cipher to be exceptionally faster than any block cipher.
- One-time pads are considered unbreakable, and each pad is used exactly once.
- In the ECB mode, each plaintext block is encrypted independently with the block cipher.
- In the CBC mode, each plaintext block is XOR'd with the previous ciphertext block and then encrypted.
- In CFB mode, the previous ciphertext block is encrypted and the output produced is combined with the plaintext block using an XOR operation to produce the current ciphertext block.

- OFB mode is similar to CFB mode except that the quantity XOR'd with each plaintext block is generated independently of both the plaintext and ciphertext.
- Cryptanalytic attacks are generally classified into six categories that distinguish the kind of information the cryptanalyst has available to mount an attack.

Message Authentication

Overview

Attackers can modify data traveling over an untrusted network without the knowledge of the sender or receiver. You should use message authentication to validate that the data has not been modified in transit. Message authentication has three main goals:

- Validate message integrity
- Identify the originator of the message
- Identify the uniqueness of the message

This lesson will discuss the various algorithms you can use to perform message authentication.

Importance

Understanding the mechanics of message authentication is essential in the understanding of cryptography and the secure exchange of data over an untrusted network.

Objectives

Upon completing this lesson, you will be able to:

- Describe the logistics of hash algorithms
- Describe the logistics of the MD5 algorithm
- Describe the logistics of the SHA algorithm
- Identify the reason for the creation of the Hash Message Authentication Code
- Explain the two main sources of weakness in hash algorithms
- Describe the attack linked to the “birthday paradox”

Outline

The lesson contains these topics:

- Hash Functions
- Message Digest 5
- Secure Hash Algorithm

- Hash Message Authentication Code
- Attacks Against Hash Algorithms
- Birthday Paradox

Hash Functions

Many people confuse hashing with encryption, but this is an incorrect comparison. **Hash algorithms** are used to produce a "fingerprint" of some data by taking the variable length data and running them through an algorithm.

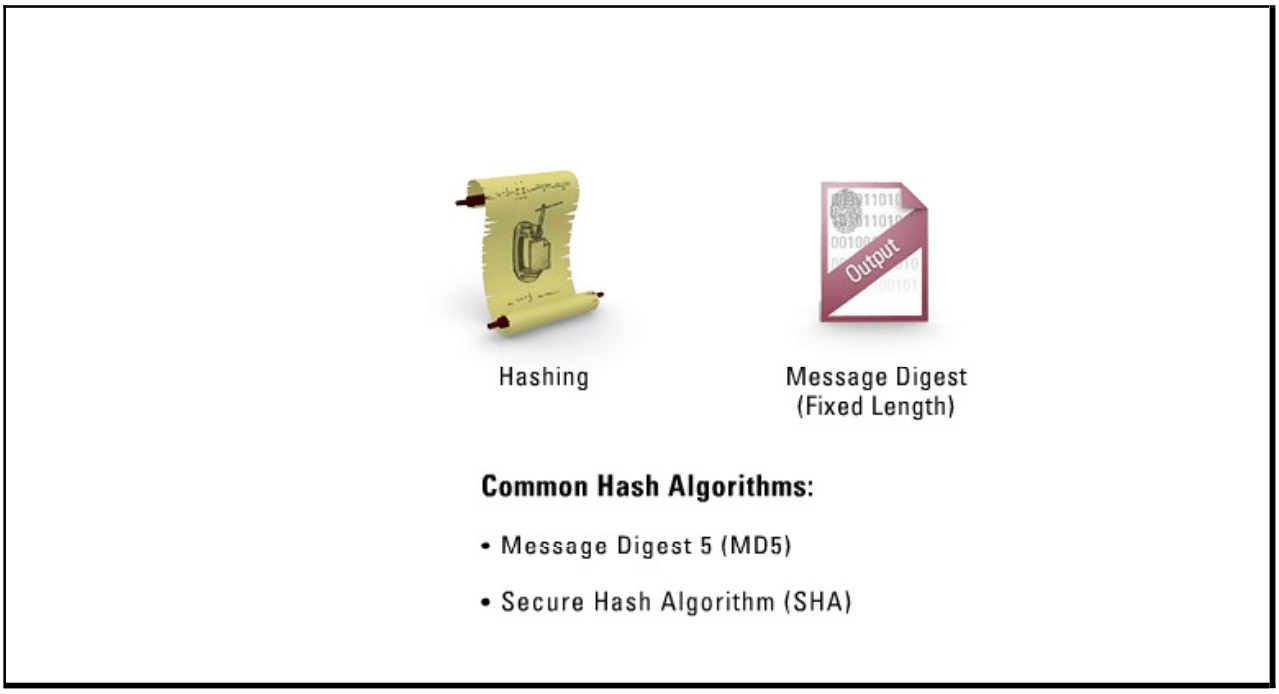

The output is a small fixed length value (called a message digest). It does not matter how many times the same input is run through the hash algorithm; the output value will always be the same. In this way, you can validate the integrity of data. If the hash value computed when the data are sent matches the hash value when the data are received, you can safely assume the data has not been modified in transit.

The basic requirements for a cryptographic hash function are as follows.

- The input can be of any length.
- The output has a fixed length.
- $H(x)$ is relatively easy to compute for any given x.
- $H(x)$ is one-way.
- $H(x)$ is collision-free.

A hash function H is said to be one-way if it is hard to invert, where "hard to invert" means that given a hash value h, it is computationally infeasible to find some input x such that $H(x) = h$.

The hash algorithms in general use today are:

- Message Digest 5 (MD5)
- Secure Hash Algorithm (SHA)

Message Digest 5

Ron Rivest of RSA Security invented the **Message Digest 5(MD5)** algorithm. It is described in RFC 1321.

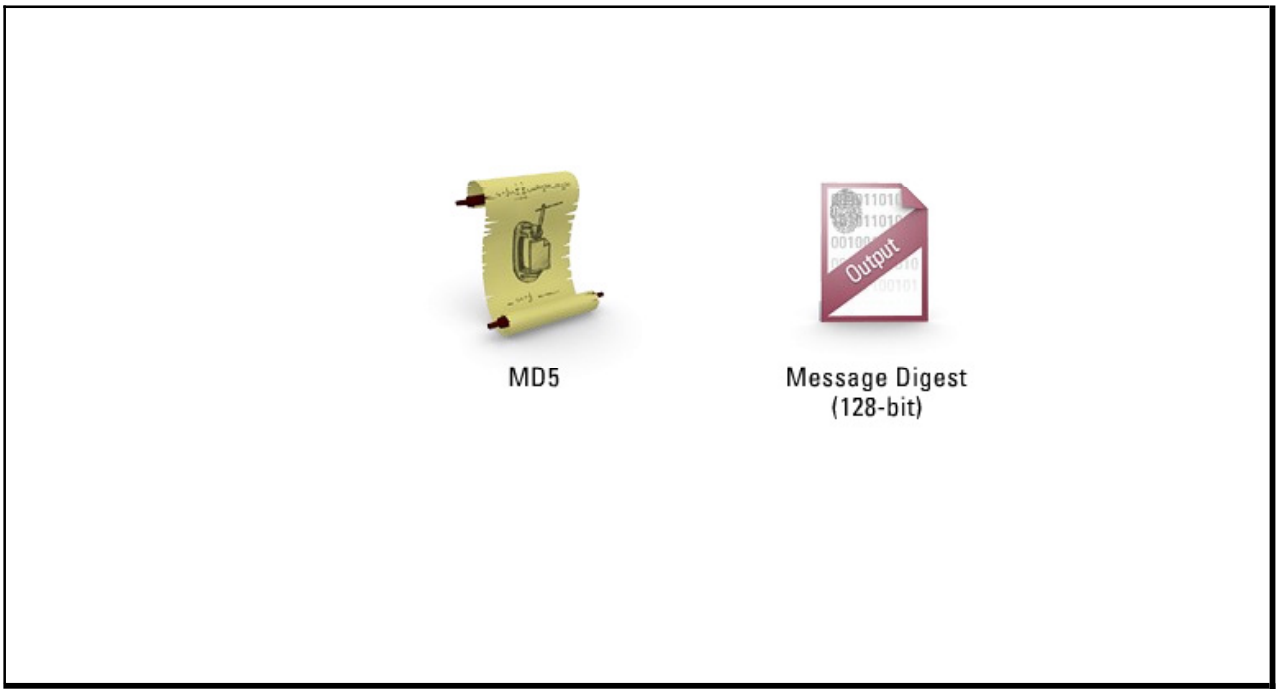

This algorithm takes a message of arbitrary length as input and produces a 128-bit "fingerprint" or message digest as output. A 128-bit algorithm means there are approximately 2^{128} possible values for any single message. Although it is technically possible to create a message to match a particular hash, the probability is so small that it is not worth considering.

For example, if you run a 64-byte Ethernet frame through the MD5 algorithm, you will receive a 128-bit value as output. If you run the same frame through the algorithm again, you will receive the exact same 128-bit value. If someone modifies even one single bit, however, the hash algorithm will compute a completely different 128-bit value.

The MD5 algorithm always outputs a 128-bit value regardless of the size of the input. A 1000-page Microsoft Word document and a 64-byte Ethernet frame each produces an output of 128 bits.

The predecessors to the MD5 algorithm are MD2 and MD4. While the structures of these algorithms are somewhat similar, the design of MD2 is quite different from that of MD4 and MD5. MD2 was optimized for eight-bit machines, whereas MD4 and MD5 were aimed at 32-bit machines.

Rivest developed MD2 in 1989. With MD2, the message is first padded so its length in bytes is divisible by 16. A 16-byte checksum is then appended to the message, and the hash value is computed on the resulting message. Cryptanalysis has discovered that collisions for MD2 can be constructed if the calculation of the checksum is omitted. This is the only cryptanalytic result known for MD2.

Rivest developed MD4 in 1990. Cryptographers quickly developed attacks on versions of MD4 with either the first or the last rounds missing. They have shown how collisions for the full version of MD4 can be found in under a minute on a typical PC. In recent work, it has been shown that a reduced version

of MD4 in which the third round of the compression function is not executed but everything else remains the same, is not in fact one-way. Clearly, MD4 should now be considered broken.

Secure Hash Algorithm

The MD5 algorithm proved to have some weaknesses in certain situations. Collisions making a well-known value match a particular hash output value were confirmed. Knowing there were possible weaknesses in the algorithm, the **Secure Hash Algorithm (SHA)** was created.

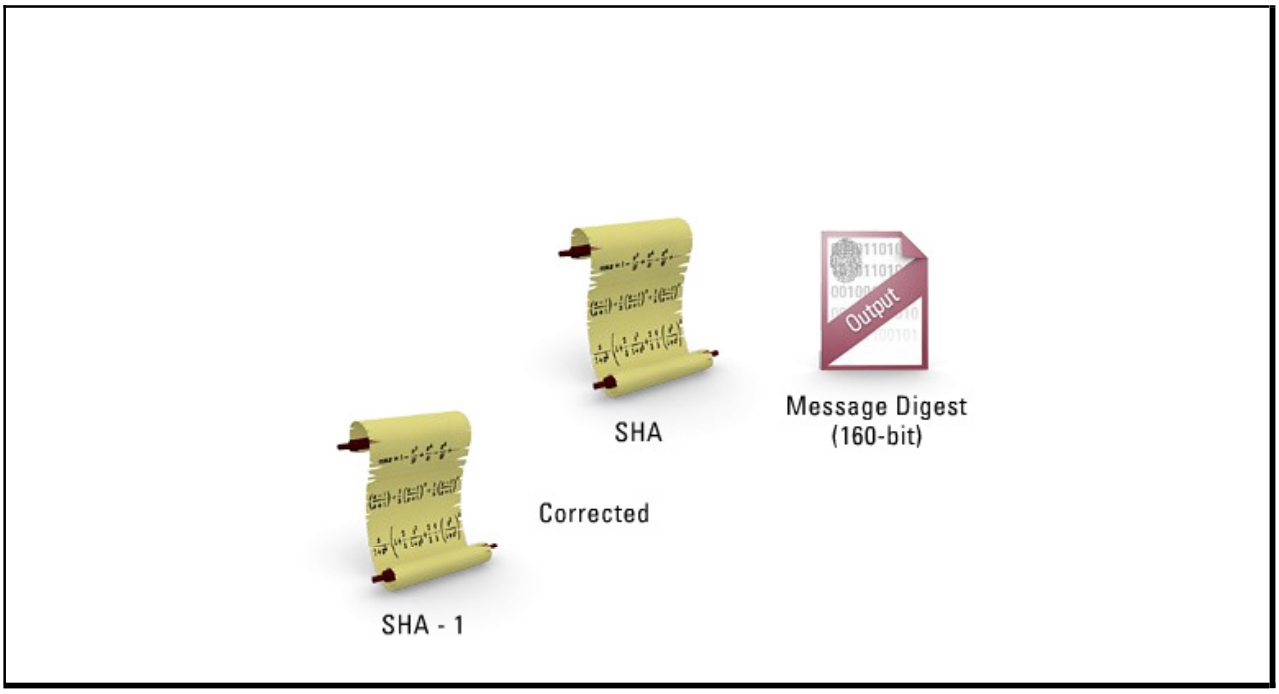

The National Institute of Standards and Technology (NIST) developed the SHA algorithm. Its design is very similar to the MD4 family of hash functions that Rivest developed.

The SHA algorithm takes a message of fewer than 2^{64} bits in length and produces a 160-bit message digest. The algorithm is slightly slower than MD5, but its larger message digest makes it more secure against brute-force collision and inversion attacks.

Note	SHA-1 is a revision to SHA that was published in 1994; the revision corrected an unpublished flaw in SHA.

Hash Message Authentication Code

Message digest algorithms have a drawback. If an attacker (man-in-the-middle) intercepts the message containing the data and the hash value, he or she can create a new message, calculate the correct hash value, append the new hash value to the new data, and send the message to the destination. The destination will separate the data from the hash, run the data through the hash algorithm, and compare the result with the received hash. Since they match, the receiver thinks the data is valid and accepts the message as being sent from its peer.

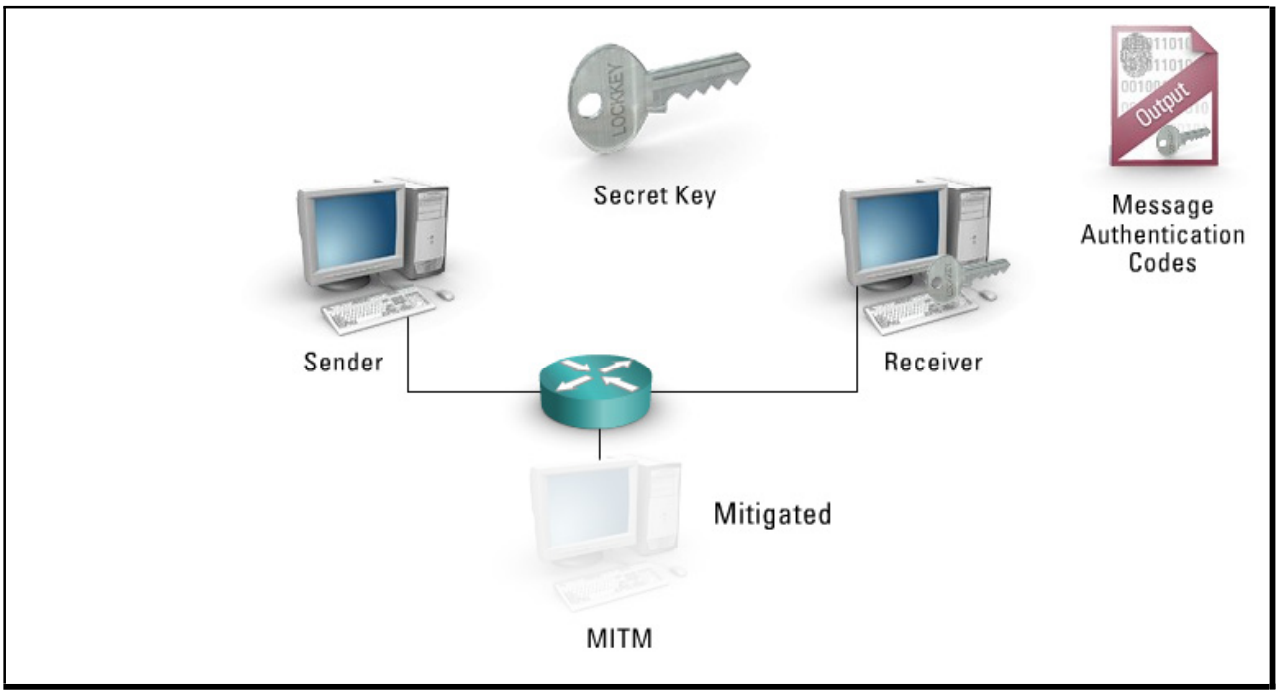

To mitigate this type of attack, a shared secret key known only between the two peers is also inserted into the hash algorithm. In this way, a random value (the key) unknown to anyone else is used to make sure that the man-in-the-middle attack cannot be successful. In effect, this key creates a built-in message authentication. Mechanisms that provide such integrity checks based on a secret key are called Message Authentication Codes (MACs).

To create the hash, the data and the shared secret key are inserted into the hash algorithm to obtain the output message digest. This is appended to the data and sent to the peer. Even if the data and hash algorithm are modified in transit, the receiver will calculate a different hash with the secret value, and discard the message.

Under certain circumstances the MD5 algorithm was shown to be susceptible to certain types of attack. An additional hash function was added to the algorithm to mitigate this problem. The additional hash function is called a **Hash Message Authentication Code (HMAC)**.

Note	When using the MAC function, MD5 is called HMAC-MD5 and SHA-1 is called HMAC-SHA-1.

Attacks Against Hash Algorithms

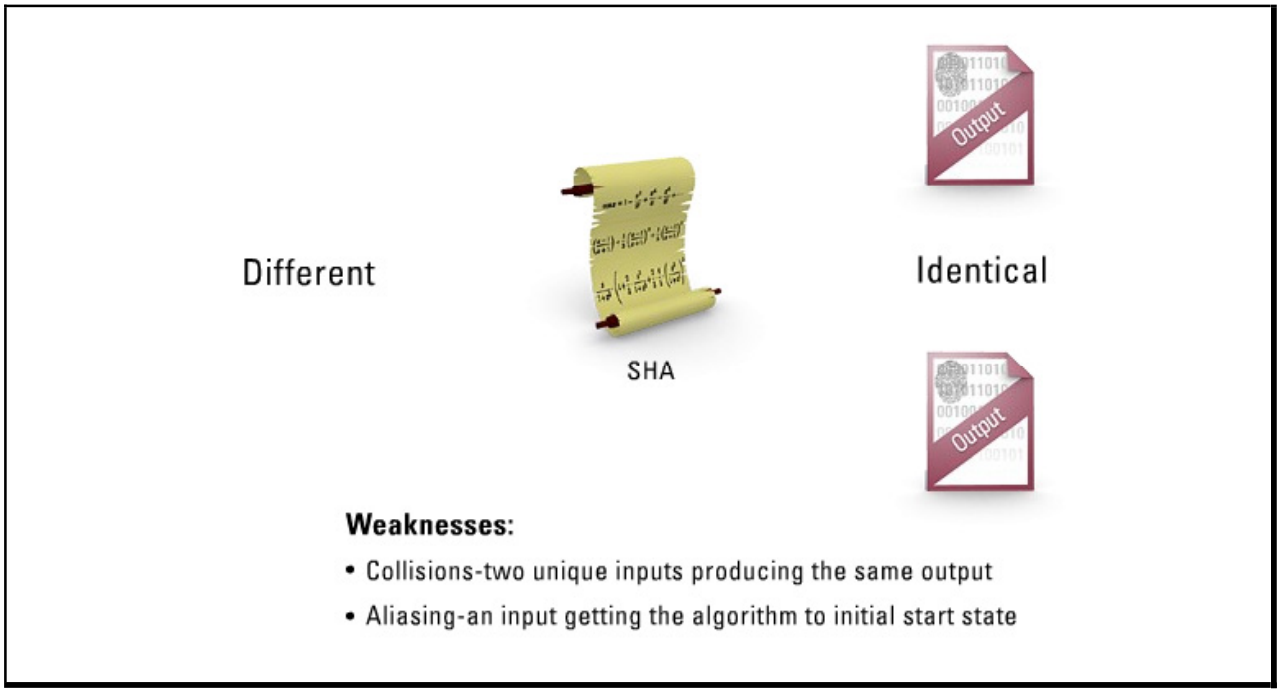

This lesson discusses attacks against hash algorithms.

The two main sources of weakness in a hash algorithm are:

- **Collisions** - The chance that two unique inputs will produce the same output
- **Aliasing** - The chance that an input can get the hash algorithm to its start state

Collisions are generally less of a problem as the bit size of the output hash increases (i.e., you are less likely to have a collision with a 160-bit hash than you are with a 128-bit hash). Incidentally, it is also possible to have a collision during the calculation of a hash value. This is known as a pseudo-collision. For instance, MD5 is said to have pseudo-collisions during some of its stages of calculation. MD4 was discovered to have collisions if either the first or last stage of calculation were left out, and eventually collisions were found for the entire algorithm.

Cryptographic hash algorithms aim to make it as difficult as possible to create a collision, meaning that their goal is to prevent someone who knows a hash value from being able to produce some data that will create the same hash value. Some hashes, mostly checksums, are actually designed with the reverse goal in mind. For instance, error correcting codes have the goal of allowing a system to see the code and the data, and if they do not match, modify the data so that they do.

Aliasing is a bit more of an interesting attack on hashing algorithms. Basically, the goal of someone mounting an aliasing attack is to find some sequence of inputs that when fed to the hashing algorithm will reset the internal state back to the start state. With this information, the attacker can make any message he or she wants have the same hash as the original. For instance, if it is found that a one-kilobyte sequence of 0xFF bytes causes the internal state of a specific algorithm to be the same as the start state, the attacker could intercept a message, put whatever data he or she likes at the beginning of it, follow it up with the kilobyte of 0xFF, and then put the original message at the end. Since the 0xFF bytes caused the hashing

algorithm to in effect start over, the hash of the modified message will be the same as the hash of the original one. A successful aliasing attack creates a collision, since it will lead to many possible inputs with the same hash value.

Birthday Paradox

There are different types of match attacks that do not attack the algorithms themselves, but instead attack probabilities. One such attack is the **birthday attack**.

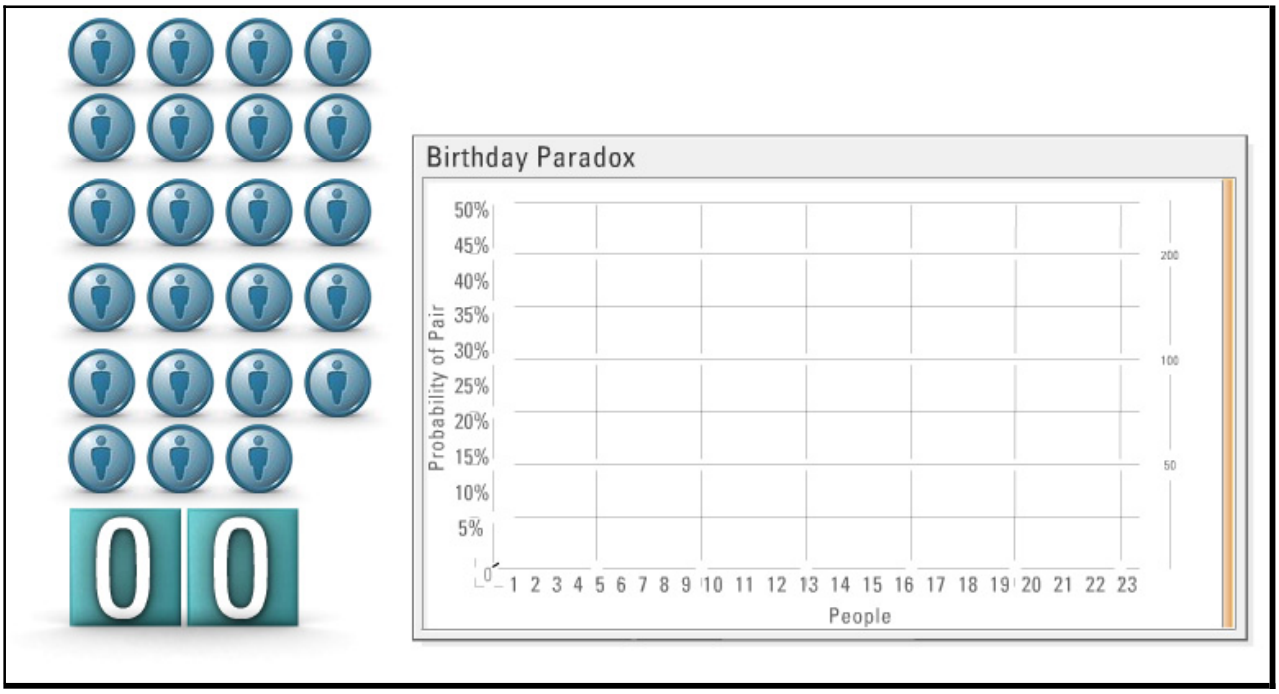

This attack is based on the so-called "birthday paradox," which is a well-known probability theory. The birthday paradox says that if 23 or more people are gathered in a room, there are better than even odds that some pair of them will share a common birthday. The odds of a single person having the same birthday as one other person is one in 365, but those odds improve for each person as the number of people increases.

Note	Birthday attacks are often used to find collisions in hash algorithms, such as MD5 and SHA-1.

Summary

The key points discussed in this lesson are:

- Hash algorithms are used to produce a “fingerprint” of some data by taking the variable length data and running them through an algorithm.
- The MD5 algorithm takes a message of arbitrary length as input and produces a 128-bit “fingerprint” or message digest as output. The MD5 algorithm always outputs a 128-bit value regardless of the size of the input.
- The SHA algorithm takes a message of fewer than 2^{64} bits in length and produces a 160-bit message digest. The algorithm is slightly slower than MD5, but its larger message digest makes it more secure against brute-force collision and inversion attacks.
- Under certain circumstances the MD5 algorithm was shown to be susceptible to certain types of attack. An additional hash function was added to the algorithm to mitigate this problem. The additional hash function is called HMAC.
- The two main sources of weakness in a hash algorithm are:
  - Collisions - The chance that two unique inputs will produce the same output
  - Aliasing - The chance that an input can get the hash algorithm to its start state
- There are different types of match attacks that do not attack the algorithms themselves, but instead attack probabilities. One such attack is the birthday attack.

Certificate Authority

Overview

Certificate authorities (CAs) allow users, applications, or other entities to unquestionably validate their identities. This validation is achieved via many of the cryptographic functions previously discussed. When a certificate is signed by a trusted third party (the CA), the sender can sign messages using information supplied on the certificate, thus providing at minimum data and source integrity. The receiver can validate the data using information it also received from the CA, thereby providing authentication and validation to the data and the source. Once all parties obtain a signed certificate with which to sign messages, along with the CA's root certificate, which is used to validate received messages, the CA can effectively be removed from the equation. This allows for the scalability of cryptography in the network.

Importance

It is important to understand the mechanics and logistics of using certificate authorities. This information allows the information security professional the ability to implement large secure private networks that are scalable.

Objectives

Upon completing this lesson, you will be able to:

- Describe the logistics of digital signatures
- Define non-repudiation
- Identify the key benefit of key escrow
- Describe the logistics of the Public Key Infrastructure
- Define guidelines for key generation, distribution, management, storage, and recovery
- Identify the two ways cryptanalysts attempt to break an encryption method

Outline

The lesson contains these topics:

- Digital Signatures
- Non-Repudiation
- Key Escrow

- Public Key Infrastructure
- Key Management and Distribution
- Key Generation
- Key Recovery
- Key Storage
- Key Strength

Digital Signatures

This topic discusses digital signatures.

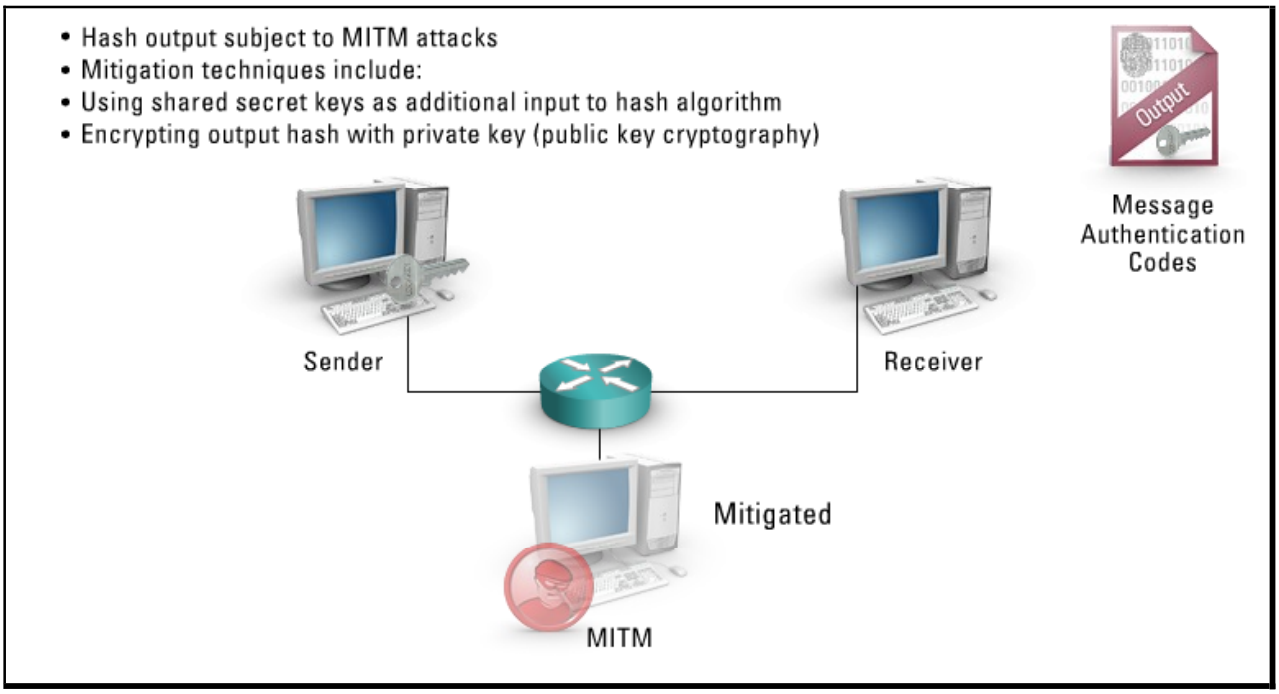

When a hash function is performed, and the hash output is appended to the original data that are sent to the peer, the data are subject to man-in-the-middle attacks. In this type of attack, the attacker intercepts the hash and data, creates his or her own data, hashes the data using the same algorithm, and sends the data to the original receiver. The receiver validates the hash and accepts it. To stop this type of attack, you can use one of the derived Diffie-Hellman (DH) shared secret keys to authenticate the data. Here, the data and the shared secret key are hashed to create the message digest. Only the peer on the opposite end, who also has the shared secret value, can create the same hash value. Attackers attempting to forge packets will not have the shared key, which means they will never compute a hash that can match the data.

In addition to stopping man-in-the-middle attacks, you can use public key cryptography to authenticate the hash. In this case, the hash is encrypted with the private key of the sender and then sent to the opposite peer. The opposite peer obtains the public key of the sender and decrypts the hash, and then performs the normal hash checking function to validate the data. Since only one person in the entire world has the private key that was used to encrypt the hash, it can be safely assumed that this person is indeed the entity that sent the message.

When you hash data and encrypt the resulting message digest using your private key, you create a **digital signature**. This means digital signatures will change for every packet that is sent, but can only be decrypted using the corresponding public key. In essence, this process validates that the hash the peer decrypted with the public key could have come from no one else in the world.

Non-Repudiation

Suppose an attacker sent an e-mail with a spoofed source e-mail address. The attacker sent the document posing as the sales manager, which confirmed the purchase of a large shipment of widgets to a company in the Netherlands. Before shipment was sent, the sales manager found out about it and stopped the order. This is possible because he could refute the fact that he sent the e-mail. This is called **repudiation**, denying that communication via e-mail (in this case) took place.

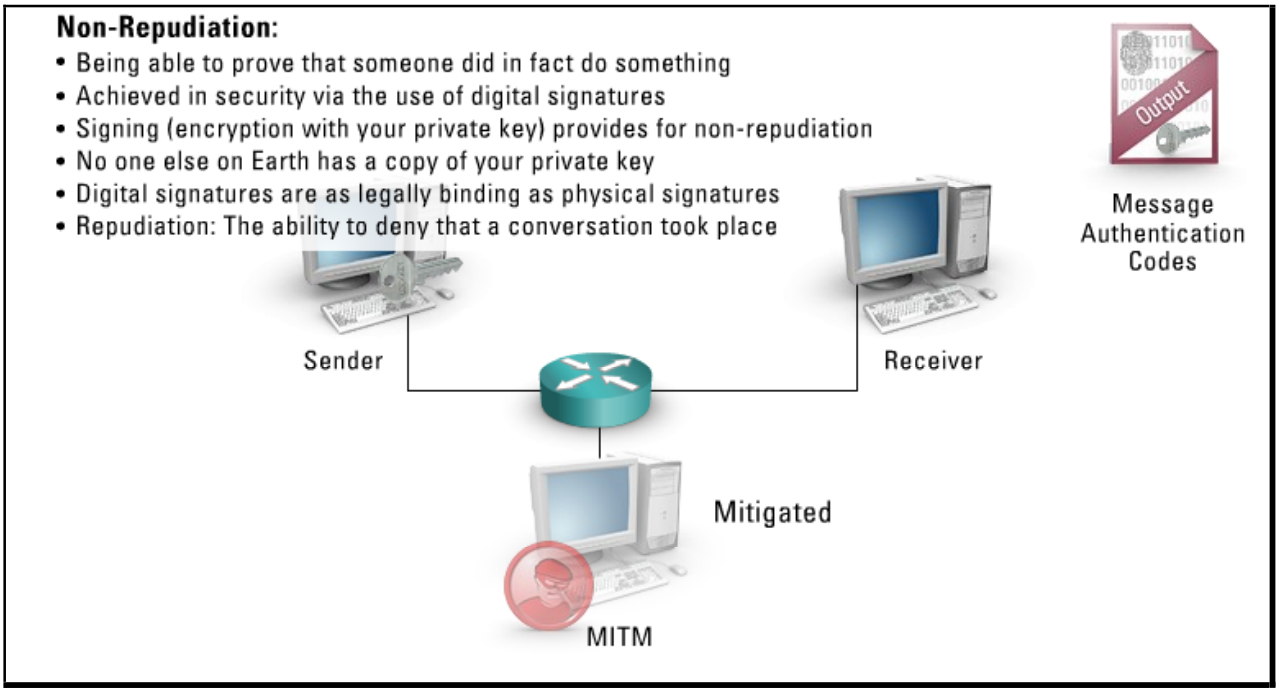

In the security world, it is as important to be able to prove that someone else did something, as it is to prove that you did not do something. If you had to take someone's denial at his or her word, the person could easily purchase products or services online, and then deny the purchase to avoid payment. Being able to prove that someone did in fact do something is called **non-repudiation**.

Non-repudiation is achieved in the security world through the use of digital signatures. Remember, every person in the entire world can have a unique private/public key pair. If you send an e-mail and digitally sign it with your private key, you and only you could have sent it. Today, many state governments have passed their own digital signature legislation acts, which give documents marked with digital signatures the same law-binding contractual obligations as a written signature.

Key Escrow

This topic discusses key escrow.

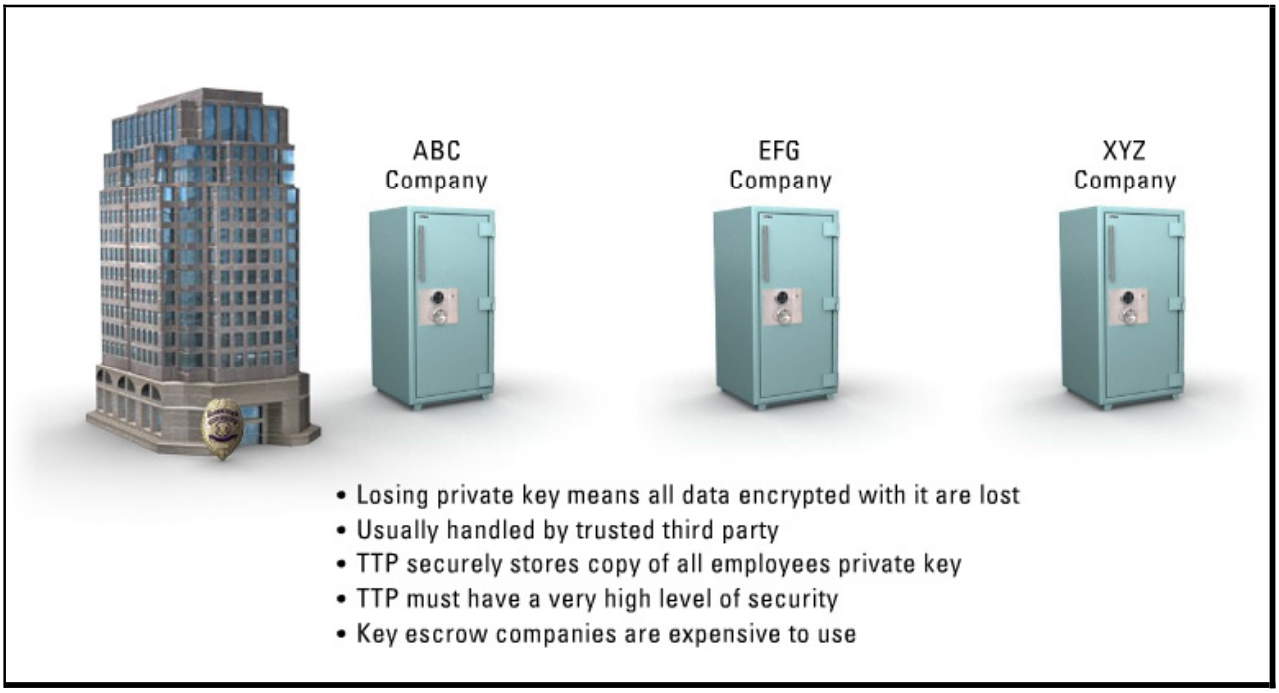

Assume you have decided to use keys and certificates to help secure your network and important data. In the process of generating your keys, you create a 4096-bit length key with an algorithm you wrote yourself. Your key protection is virtually unbreakable. But, what happens if an Internet worm attacks your system and destroys all the data on your system. Now, you have lost your public/private key pair. Obtaining your public key is not a problem as you most likely have sent it to many people, but how can you obtain your private key again? The answer is you cannot. There is no way to recreate the private key from the public key. All your data encrypted with your public key is now completely useless.

This scenario can happen in the corporate world, where employees are given their own private/public key pairs for security reasons. But, since all data rightly belongs to the company, the company would like to make sure that even if an employee loses his or her private key, the contents that require decryption with the private key are not lost. This is where key escrow takes place.

Key escrow is usually handled by a third party, who your company completely trusts not to take advantage of the fact that they have copies of all private keys in the corporation. There must be a very high level of security at the trusted third party location, which is why key escrow companies are expensive to use.

Public Key Infrastructure

A **Public Key Infrastructure (PKI)** enables users of an insecure public network such as the Internet to securely and privately exchange data and money through the use of a public and private cryptographic key pair that is obtained and shared through a trusted authority.

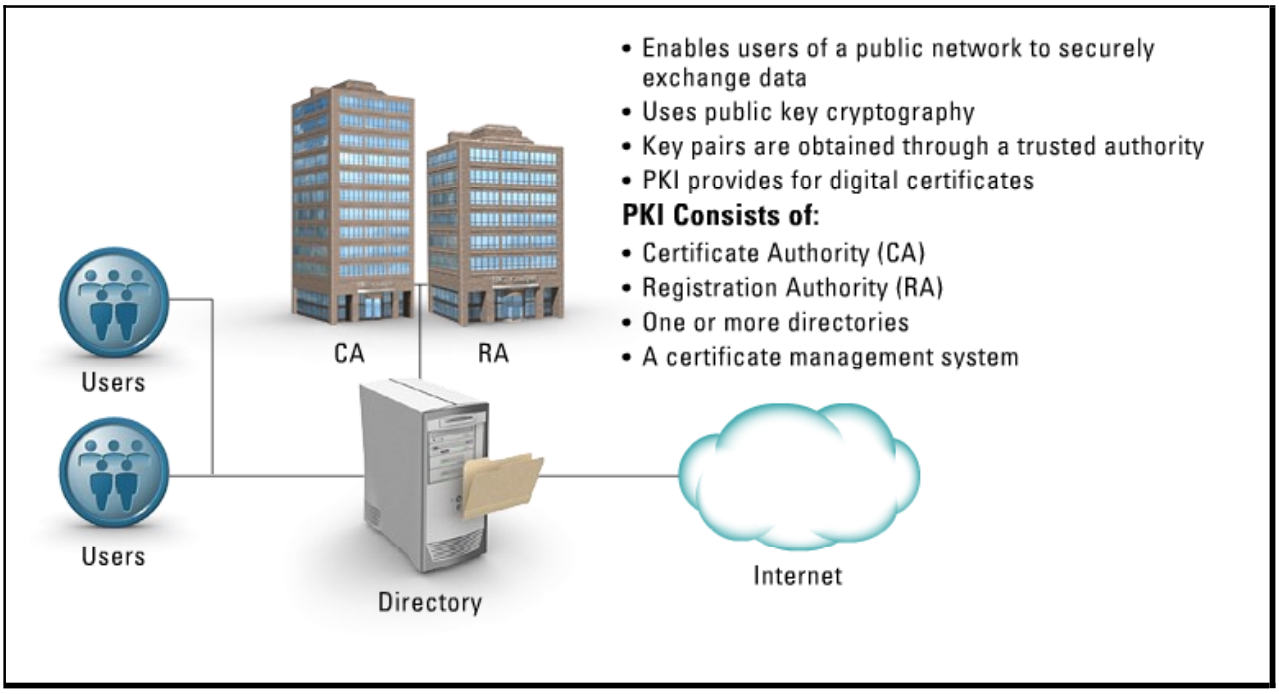

The public key infrastructure provides for a digital certificate that can identify an individual or an organization and directory services that can store and, when necessary, revoke the certificates. Although the components of a PKI are generally understood, a number of different vendor approaches and services are emerging.

The public key infrastructure assumes the use of public key cryptography, which is the most common method on the Internet for authenticating a message sender or encrypting a message. Traditional cryptography has usually involved the creation and sharing of a secret key for the encryption and decryption of messages. This secret or private key system has the significant flaw that if the key is discovered or intercepted by someone else, this person can easily decrypt messages. For this reason, public key cryptography and the public key infrastructure create the preferred approach on the Internet.

A public key infrastructure consists of:

- A **certificate authority(CA)** that issues and verifies digital certificates; a certificate includes the public key or information about the public key
- A **registration authority (RA)** that acts as the verifier for the certificate authority before a digital certificate is issued to a requestor
- One or more directories where the certificates (with their public keys) are held
- A certificate management system

Key Management and Distribution

This topic discusses rules for key management and distribution.

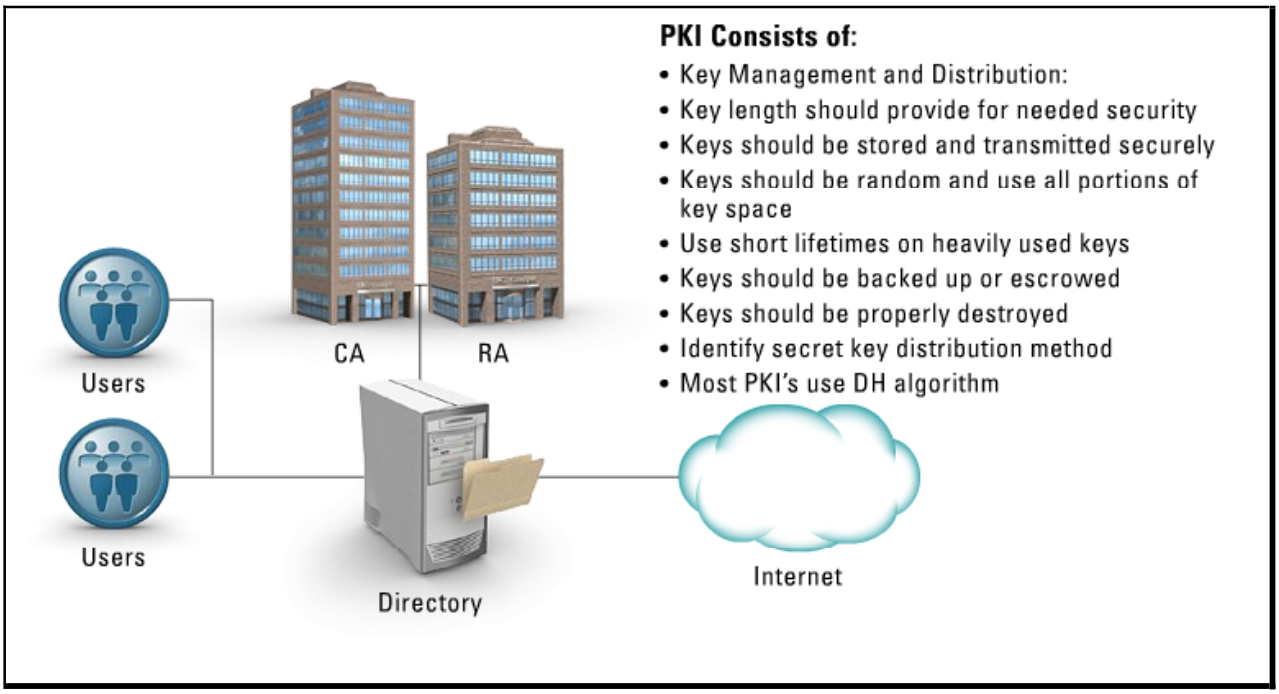

Key management rules include:

- The key length should be long enough to provide the necessary level of protection. Keys should be stored and transmitted by secure means.
- Keys should be extremely random and use the full spectrum of the keyspace.
- The key's lifetime should correspond with the sensitivity of the data it is protecting.
- The more the key is used, the shorter its lifetime should be.
- Keys should be backed up or escrowed in case of emergencies.
- Keys should be properly destroyed when their lifetimes come to an end.

Key distribution is the process of getting a key from the point from which it is first generated, to the point where it is intended for use. The key distribution process is more difficult in symmetric algorithms, where it is necessary to protect the key from disclosure in the process. Normally, this step is performed using a channel separate from the one in which traffic moves.

When dealing with cryptographic components of a PKI, key distribution for symmetric key encryption can become burdensome. Remember, for secure communication to take place based on a symmetric key system, you have to distribute the secret keys before any decryption can be done. The question now is, "How should you distribute the secret keys?"

- Via mail? This method is insecure and unreliable, and there is a time delay.
- Via courier? This method is unreliable and expensive, and there is a time delay.
- Use a secure point-to-point connection? This method is too costly and hard to change.
- Use public channels? With this method, virtually everyone can get access.

The method that most public key infrastructures adopt is to use the DH algorithm to agree upon the symmetric key between parties. Remember, though, that DH is susceptible to man-in-the-middle attacks and lacks authentication capabilities. So, you can use DH with the service of hashing algorithms and public/private keys to perform both integrity and authentication.

Now, you may be thinking about key distribution for pubic key cryptography. Well, you could have peer X send his or her public key to peer Y, but this method will create these questions:

- To how many people does peer X have to send his or her public key?
- How does peer Y trust that the key he or she received is truly the public key of peer X?

A fundamental issue in public key distribution is to make sure the key received does indeed belong to the entity that claims to have sent it. The answer is to use digital certificates. A digital certificate is a certificate issued by a mutual (by the owner of the key and whoever used the certificate) trusted third party. This digital certificate contains the identification information (not sensitive information) of the key owner, the public key, issuing authority, expiration date, etc.

Key Generation

In most cases, the PKI will perform the key generation as well as key distribution that are required to obtain a public/private key pair for all entities. But, it is entirely possible for the entity to create his or her own public/private key pair.

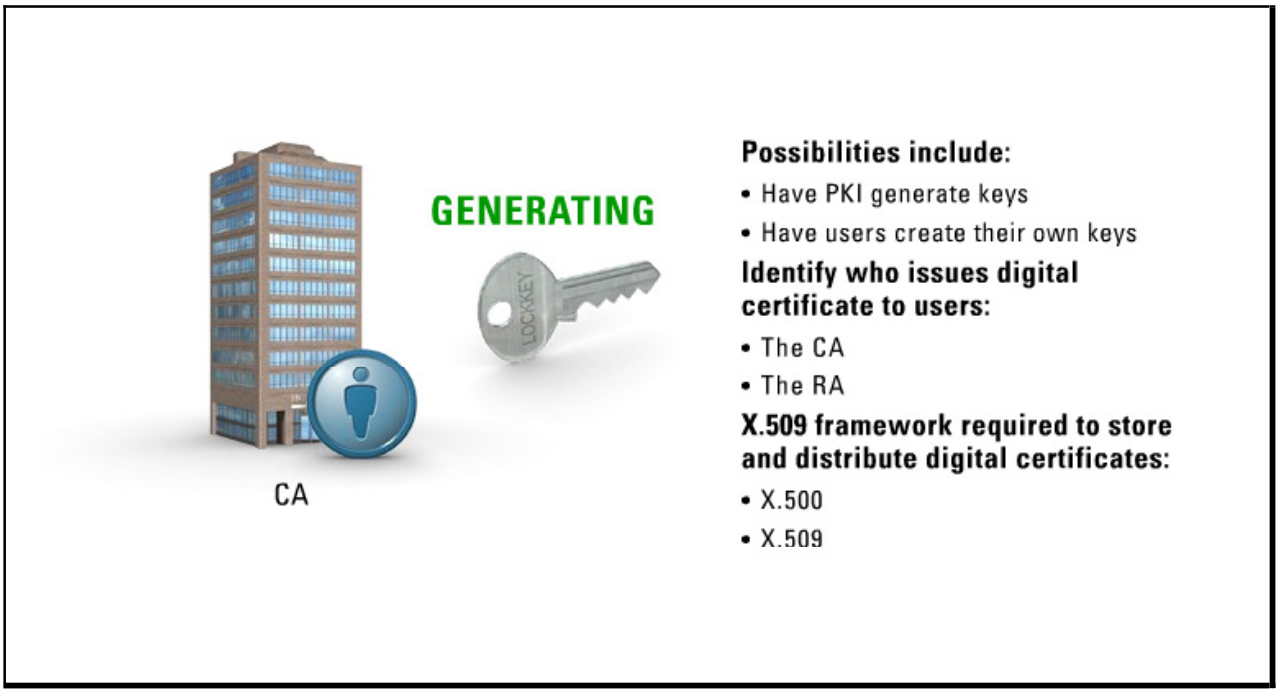

As an encrypted code or password, keys must be generated, distributed, and eventually replaced like any other security mechanism. You must handle the digital key with as much care as you would handle the key to your house. No unknown person should have any access to the key, or the lifecycle process, at any time.

The next step is to determine who issues the public key certificate for the end users and how it is used:

- A public key certificate issuing authority (certificate authority)
- The CA's public key is publicly available (and trusted)
- Validity of an individual's public key can be verified using the CA's public key
- Individual's public key can be retrieved from the public key certificate

Note You have to trust the CA, if you want to use the public key certificate services.

CAs are the entities that issue public key certificates. But, if the community base is too large, the CA can become overwhelmed with requests for public keys and other tasks that are part of its responsibilities (such as key deletion, regeneration, etc.) To offload some of the more mundane services from the CA, registration authorities (RAs) were created. RAs are the entities that verify the candidates who wish to receive a public key certificate.

To allow public key systems to work properly, a single trusted third party that manages the public key certificates is not enough. You need a repository to store public key certificates and distribute them as

needed. The International Organization for Standardization (ISO) has developed the X.509 framework for just this process.

- **X.500** is one or more directories, as specified in the X.500 standard, to hold the certificates.
- **X.509** is the protocol used to manage the certificates (e.g., how to obtain the certificates, how to update them, etc.).

Key Recovery

To provide for key recovery, many certificate authorities invoke a key recovery agent. You can use key recovery agents to recover keys, key components, or plaintext messages upon the receipt of an authorized request.

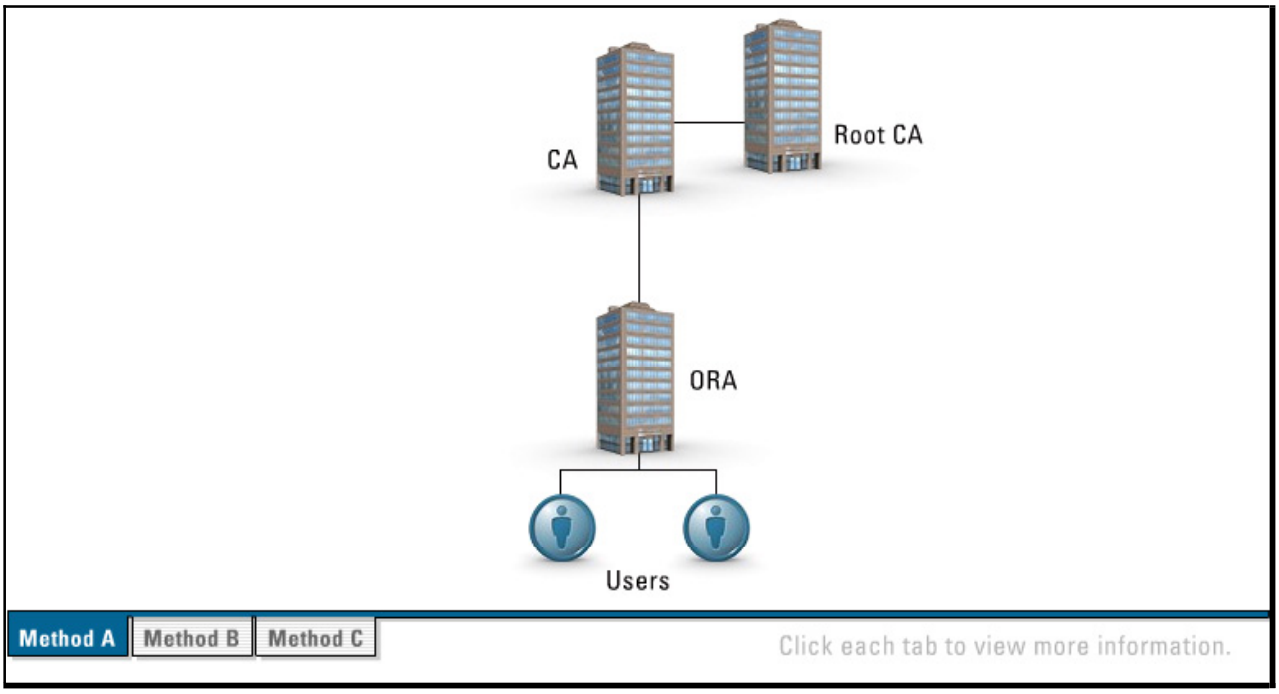

There are usually three methods for key recovery:

- In method A of the example infrastructure, a key has been stored and is directly accessible by the key recovery agent.
- In method B, key components have been stored at separate storage locations from which they may be retrieved.
- In method C, the key recovery agent does not need to store the user's key. For example, a message header could contain a session key that has been encrypted with a key known by the key recovery agent. Although the graphic shows the key recovery services as being provided by a remote entity, these services can be collocated with any of the elements shown here (e.g., RA, user).

Key Storage

There are many ways to store encryption keys, but all methods strive for one main goal and that is to protect the confidentiality of the key.

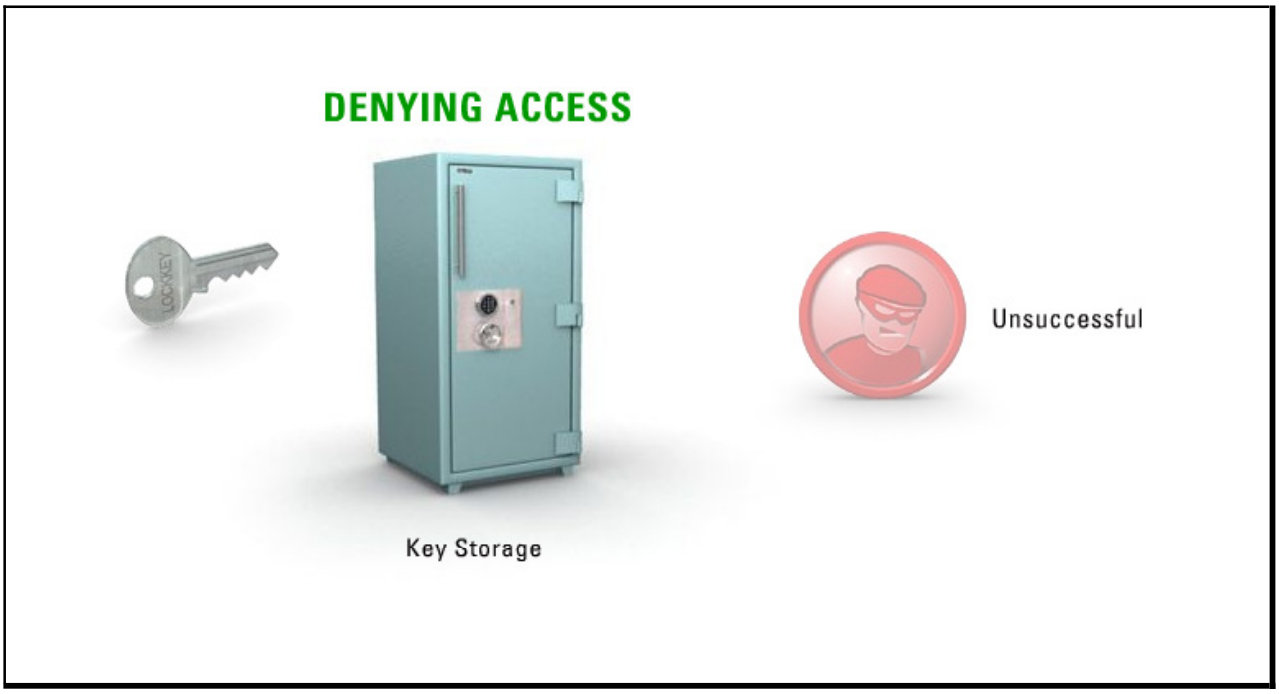

Some keys may be protected by the integrity of the storage mechanism itself. In other words, the mechanism is designed so that once the key is installed, it cannot be observed from outside the encryption machine itself. As a matter of fact, some key storage devices have been designed to self-destruct when subjected to forces that might in some way disclose the key. Other key storage mechanisms involve storing the key in encrypted form. In this way, even if the key itself is disclosed, the data protected by the key are secure, assuming that the key itself is sufficiently strong enough to withstand an attack.

Key Strength

There are two ways cryptanalysts attempt to break an encryption method: break the algorithm, or break the key.

Some people believe you best protect encryption algorithms when you keep the algorithm secret. But, this has been proven time after time to be false thinking. In the beginning, the algorithm may be secure, but eventually the cryptographic community will discover the algorithm. When they do find the algorithm, cryptanalysts like nothing more than to break the algorithm or find its weaknesses.

Because hiding the algorithm does not work, many cryptographers take the opposite approach. They advertise their algorithms to the public and invite the public to break them or expose their weaknesses. When the community does expose an algorithm's weaknesses, the inventor will then repair the algorithm and advertise it again. In the end, a secure algorithm will enter the public. Algorithms that have been tested by the cryptographic community are ultimately found to be free from algorithmic weaknesses. Because all of today's algorithms have been in the public domain, you can be reasonably sure that they are safe from attack. This means that attackers will have a very slim chance of breaking an algorithm, so instead their best bet would be to attack the key.

Key strength is directly proportional to the length of the key. You have seen that attackers can crack the Data Encryption Standard (DES), which uses a 56-bit keyspace, in less than one day and is therefore considered weak. Newer algorithms have increased key sizes and therefore provide greater key strength.

Below is a comparison of the time and money that would be required to break keys of different lengths.

Comparison of Time and Money Needed to Break Different Length Keys					
	Length of Cryptographic Key				
Cost	40-bit	56-bit	64-bit	80-bit	128-bit
$0.1 M	2 sec.	35 hrs.	1 yr.	70,000 yrs.	10E19 yrs.
$1.0 M	.2 sec.	3.5 hrs.	37 days	7,000 yrs.	10E18 yrs.
$100 M	2 millisec.	2 min.	9 hrs.	70 yrs.	10E16 yrs.
$1.0 B	.2 millisec.	13 sec.	1 hr.	7 yrs.	10E15 yrs.
$100 B	2 microsec.	1 sec.	32 sec.	24 days	10E13 yrs.

Source: B. Schneier, Applied Cryptography, 2nd ed., 1996.

Secret-Key and Public-Key Lengths for Equivalent Levels of Security	
Secret-Key Length	Public-Key Length
56 bits	384 bits
64 bits	512 bits
80 bits	768 bits
112 bits	1792 bits
128 bits	2304 bits

Summary

The key points discussed in this lesson are:

- When you hash data and encrypt the resulting message digest using your public key, you create a digital signature. This means digital signatures will change for every packet that is sent, but can only be decrypted using the corresponding public key.
- Being able to prove that someone did in fact do something is called non-repudiation. Non-repudiation is achieved in the security world through the use of digital signatures.
- Key escrow is usually handled by a third party, who your company completely trusts not to take advantage of the fact that they have copies of all private keys in the corporation. There must be a very high level of security at the trusted third party location, which is why key escrow companies are expensive to use.
- A PKI enables users of an insecure public network such as the Internet to securely and privately exchange data and money through the use of a public and private cryptographic key pair that is obtained and shared through a trusted authority.
- As an encrypted code or password, keys must be generated, distributed, and eventually replaced like any other security mechanism. You must handle the digital key with as much care as you would handle the key to your house. No unknown person should have any access to the key, or the lifecycle process, at any time.
- To provide for key recovery, many certificate authorities invoke a key recovery agent. You can use key recovery agents to recover keys, key components, or plaintext messages upon the receipt of an authorized request.
- There are many ways to store encryption keys, but all methods strive for one main goal and that is to protect the confidentiality of the key.
- Because all of today's algorithms have been in the public domain, you can be reasonably sure that they are safe from attack. This means that attackers will have a very slim chance of breaking an algorithm, so instead their best bet would be to attack the key. Key strength is directly proportional to the length of the key.

Security Architecture and Models

Overview

Organizations create networks to achieve three main goals:

- Give users access to corporate resources
- Give users bandwidth for data consumption and increased productivity
- Give users access to non-corporate resources

These goals along with the added complexity of securing the network make network architecture and design an extremely important facet of an enterprise's overall business objective.

Objectives

Upon completing this module, you will be able to:

- Identify the various hardware and software systems and how security is implemented in the network architecture
- Identify the different types of memory storage and their associated security concerns
- Identify the various security models and the controls used in securing today's networks
- Identify common flaws and security issues that have been exploited by attackers
- Identify the various types of timing attacks and how they can be mitigated

Outline

The module contains these lessons:

- Common Computer Architectures and Designs
- Storage Types
- Principles of Common Security Models
- Common Flaws and Security Issues with System Architectures and Designs
- Timing Attacks

Common Computer Architectures and Designs

Overview

This lesson will discuss the various hardware and software systems available and how you can implement security in your network architecture to achieve your organization's goals.

Importance

Understanding the principles of common computer architectures and designs is important to the creation of a secure network with many layers of defense.

Objectives

Upon completing this lesson, you will be able to:

- Describe the basic architecture for an information system
- Identify common addressing schemes
- Identify enhancements to the CPU
- Describe the logistics of protection rings
- Identify the four operating states of a system
- Define hardware, software, and firmware
- Define real machines
- Define virtual machines
- Define multistate machines
- Identify the different types of multitasking
- Define multiuser machines
- Describe the functions of network protocols
- Describe the functions of the resource manager

Outline

The lesson contains these topics:

- Principles of Common Computer Networks
- Addressing Schemes
- CPU Enhancements
- Protection Rings
- Operating States
- Hardware, Software, and Firmware Differences
- Machine Types
- Real Machines
- Virtual Machines
- Multistate Machines
- Multitasking
- Multiuser Machines
- Network Protocol Functions
- Resource Manager Functions

Principles of Common Computer Networks

When building an information system, you are required to balance various elements, such as capability, flexibility, performance, ease of use, and cost. Of course, the two most important elements you must balance are business requirements and security. Remember, you should always consider security from the beginning.

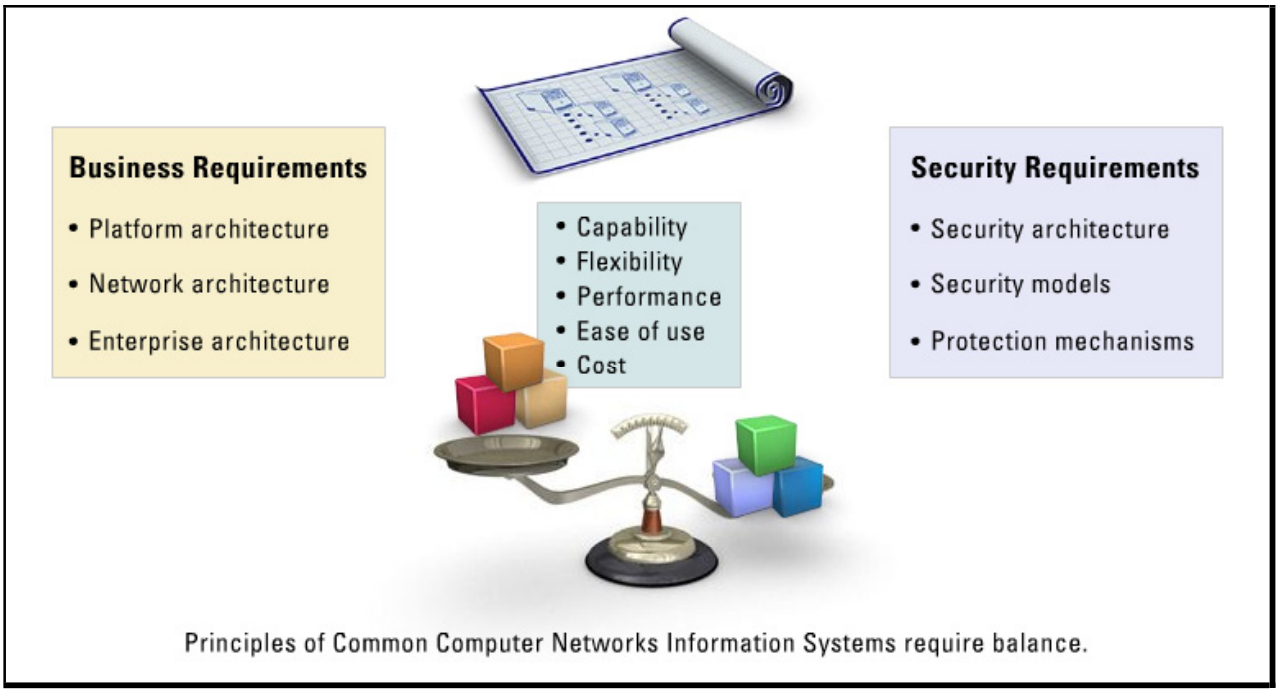

Principles of Common Computer Networks Information Systems require balance.

A **security architecture** will describe how a system is put together to satisfy security requirements. It is more of a design overview and not a description of the functions of the system. It will describe at an abstract level the relationships between key elements of the hardware, operating systems, applications, network, and other required components to protect the organization's interests. The security architecture should also describe how the functions in the system development process follow the defined security requirements.

A basic architecture for an information system can be:

- **Platform Architecture** - Encompasses the computer and how it manages system resources
  - Operating system and various utilities
  - CPU states
  - Memory management
  - Input/output (I/O) devices
  - Storage devices
- **Network Architecture** - Network design that identifies how entities will communicate with each other, whether it is an open system or a closed system
- **Enterprise Architecture** - A systematically derived and captured structural description of the mode of operation for any given enterprise

- **Security Models** - How the concept of security will be implemented into the design and analysis of secure computer systems; models include:
  - Bell-LaPadula security model
  - Biba security model
  - Clark-Wilson security model
- **Protection Mechanisms** - Define the mechanisms you will use to protect enterprise assets; you will use these assets to ensure the separation among objects in a system

Addressing Schemes

There are many different types of **addressing schemes** that CPUs can use. These addressing schemes are dependent upon many factors such as hardware used, CPU type, memory type, and addressing type.

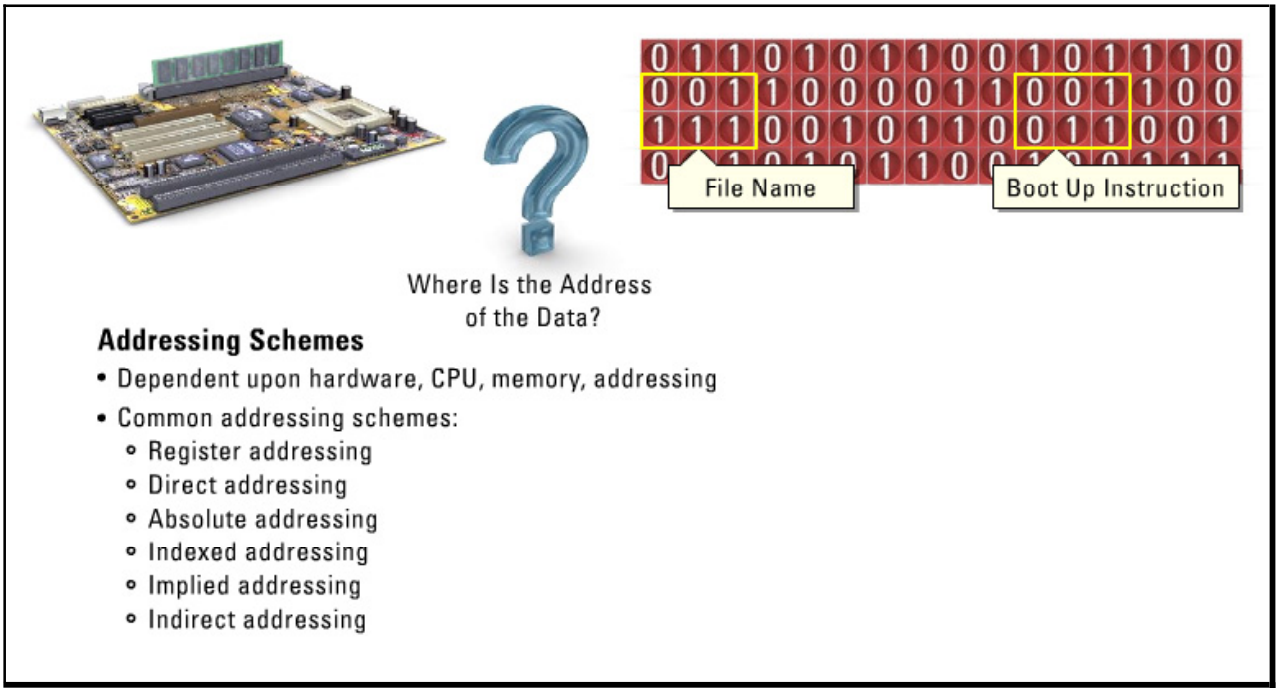

Common addressing schemes include:

- **Register Addressing** - Addressing registers within the CPU or registers in primary memory
- **Direct Addressing** - Addressing a portion of primary memory with an actual address of real memory
- **Absolute Addressing** - Addressing all the primary memory space
- **Indexed Addressing** - Adding the memory address to an index register in order to address a certain memory location
- **Implied Addressing** - Internal register where there is no need to supply an address
- **Indirect Addressing** - Address specified in the instruction contains final desired location

Note	Often, operating systems implement a memory protection scheme, which is used to prevent one program from modifying the memory contents of another.

CPU Enhancements

As CPU programming and creation have evolved, so have their enhancements.

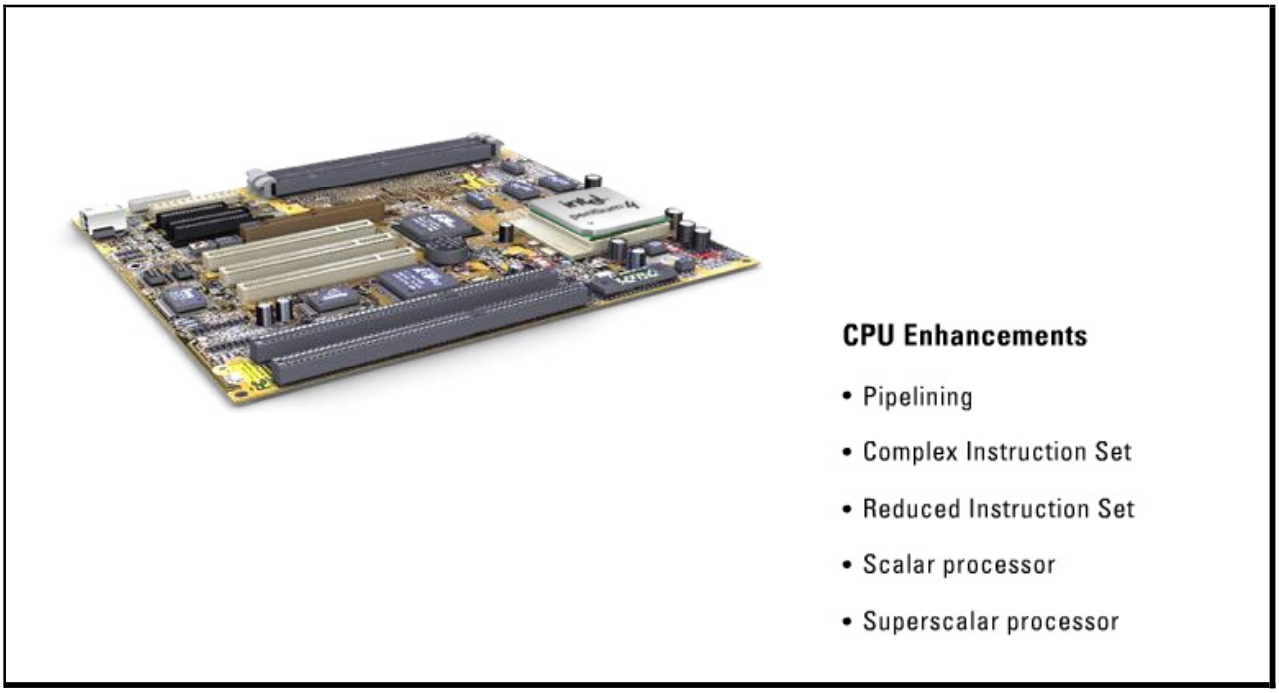

CPU enhancements:

- **Pipelining** - Increases performance by overlapping the steps of instructions
- **Complex Instruction Set** - Instructions perform many operations per instruction and are based on taking advantage of longer fetch times
- **Reduced Instruction Set** - Simpler instruction set that requires less clock cycles to complete, which results in faster processors that enable fetch times to equal decode and execute times
- **Scalar Processor** - A processor that executes one instruction at a time
- **Superscalar Processor** - A processor that enables concurrent execution of multiple instructions in the same pipeline

Protection Rings

Protection rings provide strict boundaries and definitions on what the processes that work within each ring can access and what commands they can successfully execute. The processes that operate within the inner rings have more privileges (privileged/supervisor mode) than the processes operating in the outer rings (user mode).

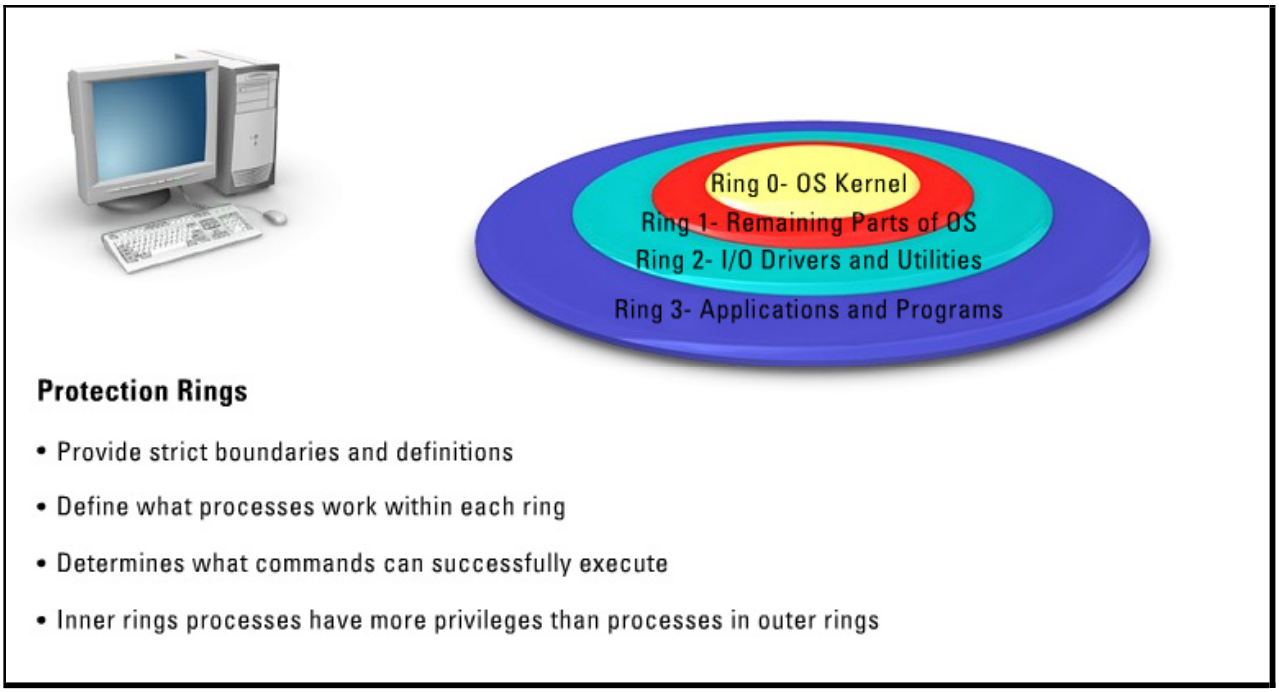

Protection rings support the integrity, availability, and confidentiality requirements of multitasking operating systems. Many systems use four protection rings:

- **Ring 0** - Operating System kernel
- **Ring 1** - Remaining parts of the OS
- **Ring 2** - I/O drivers and utilities
- **Ring 3** - Applications and programs

Protection rings provide an intermediate layer between subjects and objects, and are used for access control when a subject tries to access an object. It is the ring that determines the access level to sensitive system resources.

Operating States

There are different operating states in which a system can operate.

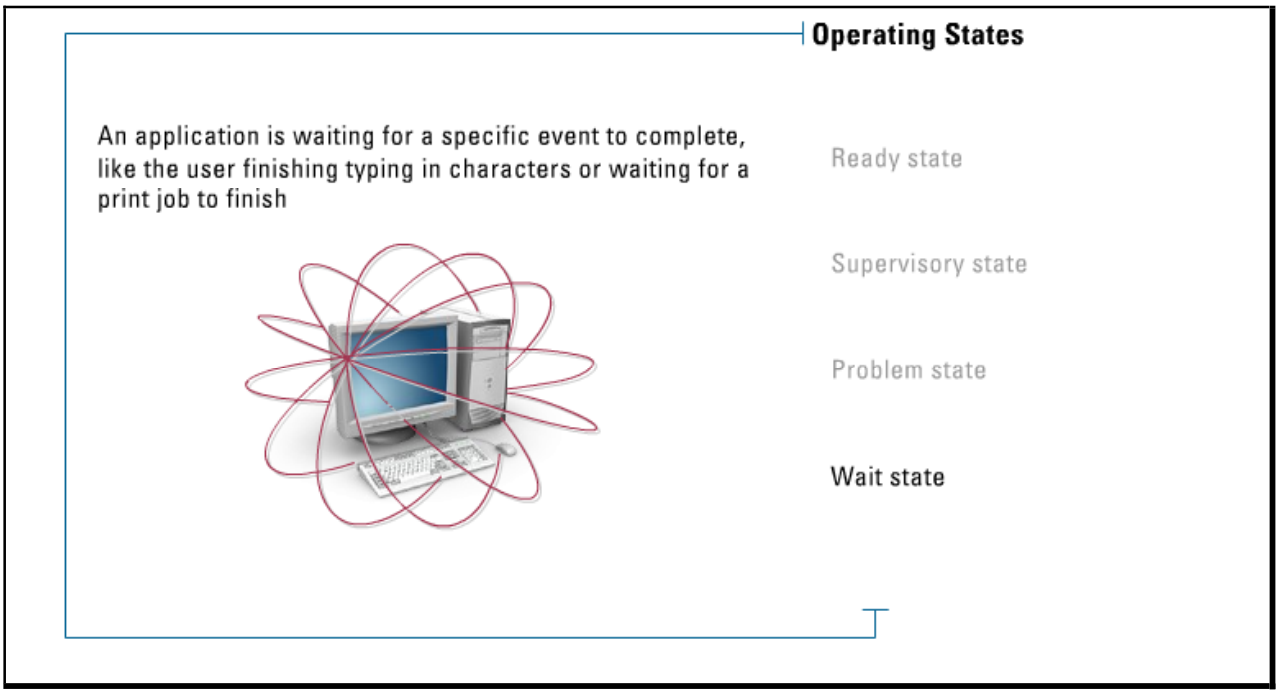

The four operating states a system can operate in include:

- **Ready State** - An application is ready to resume processing
- **Supervisory State** - The system is executing a system, or highly privileged, routine
- **Problem State** - The system is executing an application
- **Wait State** - An application is waiting for a specific event to complete, like the user finishing typing in characters or waiting for a print job to finish

Hardware, Software, and Firmware Differences

This topic discusses the differences among hardware, software, and firmware.

Hardware is any physical electronic device. Computers, adapter cards, and Ethernet cables are examples of hardware.

Software is a broad term for the programs running on hardware. Familiar kinds of software are operating systems, which provide overall control for computer hardware and applications; applications are optional programs used for a particular job. Software resides on disks and is brought into memory when it is needed.

Often a distinction is drawn between software and firmware. **Firmware** is software that is semi-permanently placed in hardware. It does not disappear when hardware is powered off, and is often changed by special installation processes or with administration tools. The memory firmware uses is very fast—making it ideal for controlling hardware where performance is important.

A software or firmware upgrade makes a permanent change—usually to improve features or performance, or to correct errors.

Machine Types

Computers can run many different applications, programs, and utilities at the same time. They must follow a model to make sure each process has its share of CPU cycles and memory to perform its function.

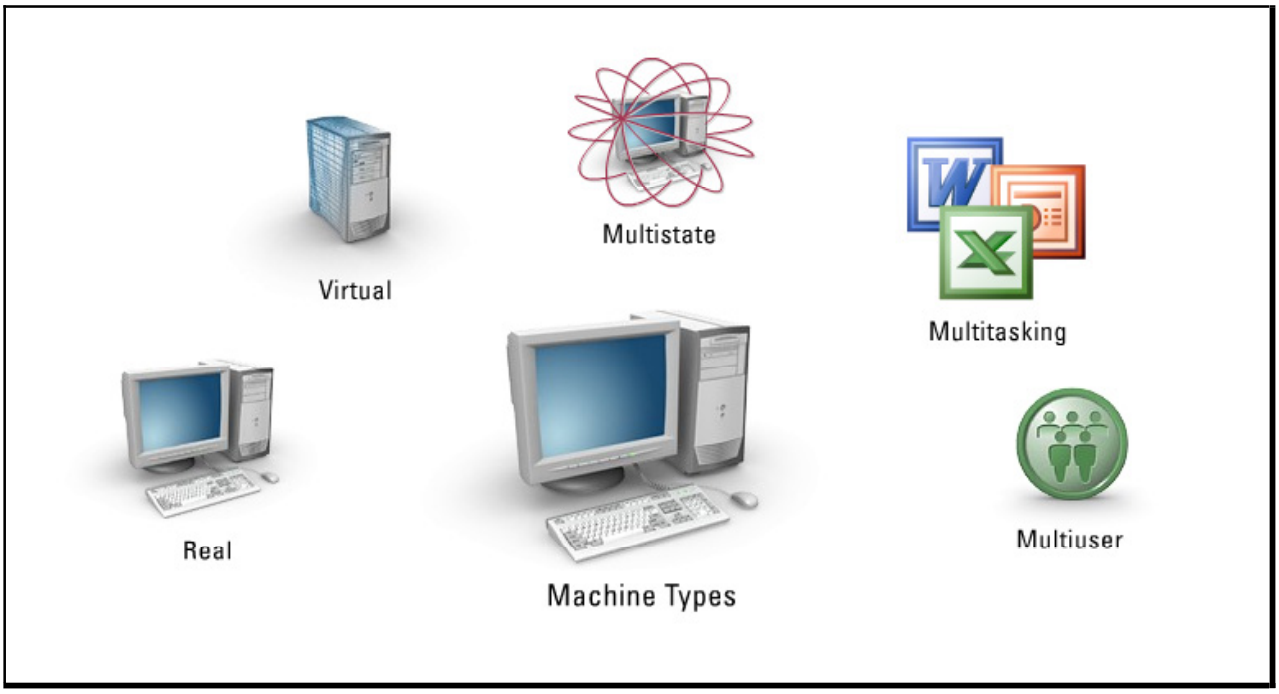

Machine types define how systems operate and can be any of the following:

- Real
- Virtual
- Multistate
- Multitasking
- Multiuser

Real Machines

A **real machine** is a computer running a single application and using all memory and CPU it has at its disposal to accomplish the task.

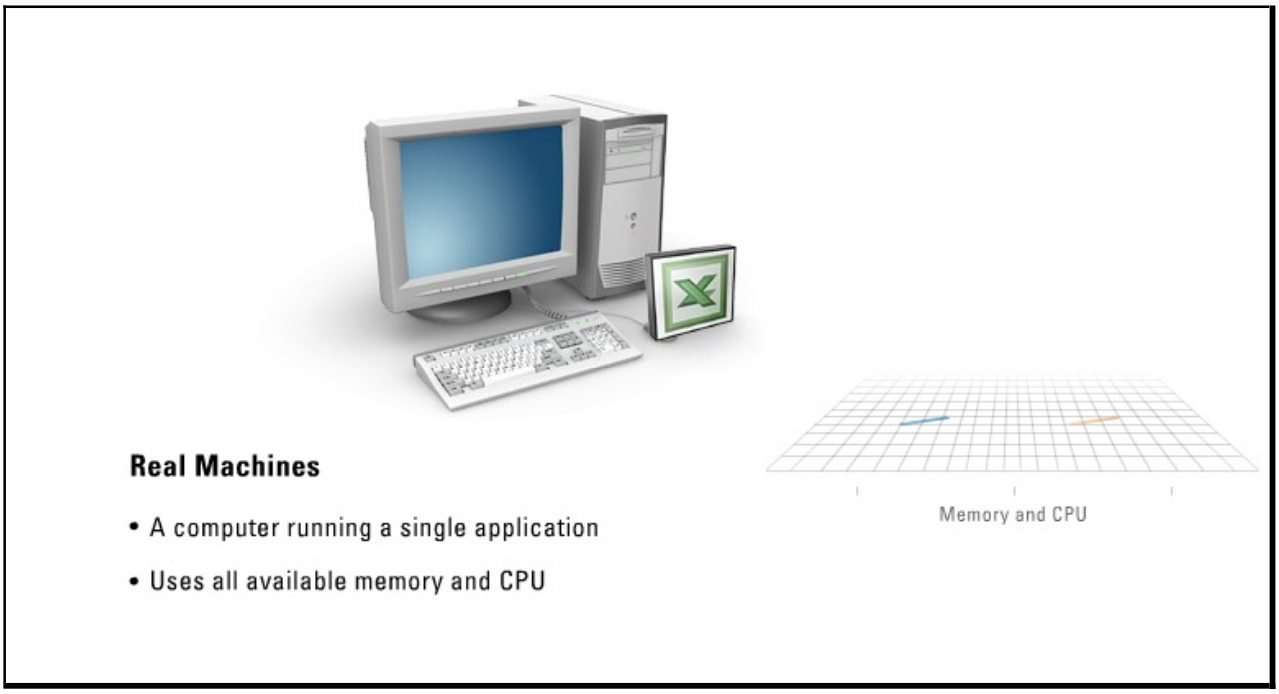

Unfortunately, most computers run more than one application or program and need to emulate the real machine.

Virtual Machines

Virtual machines create an environment that emulates the real machine.

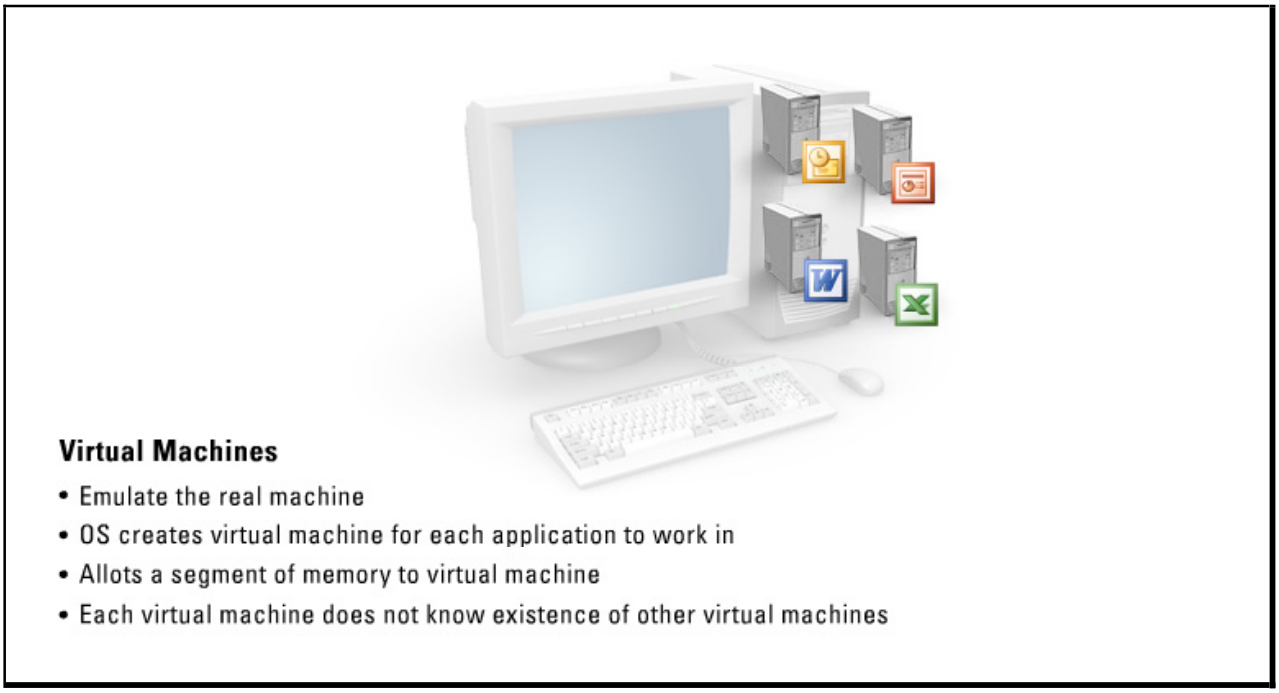

The operating system creates a virtual machine for each application to work in and allots each application a segment of virtual memory. Essentially, each application has its own virtual machine and virtual address segment and does not know that the other applications even exist.

Multistate Machines

Multistate machines are a byproduct of virtual machines.

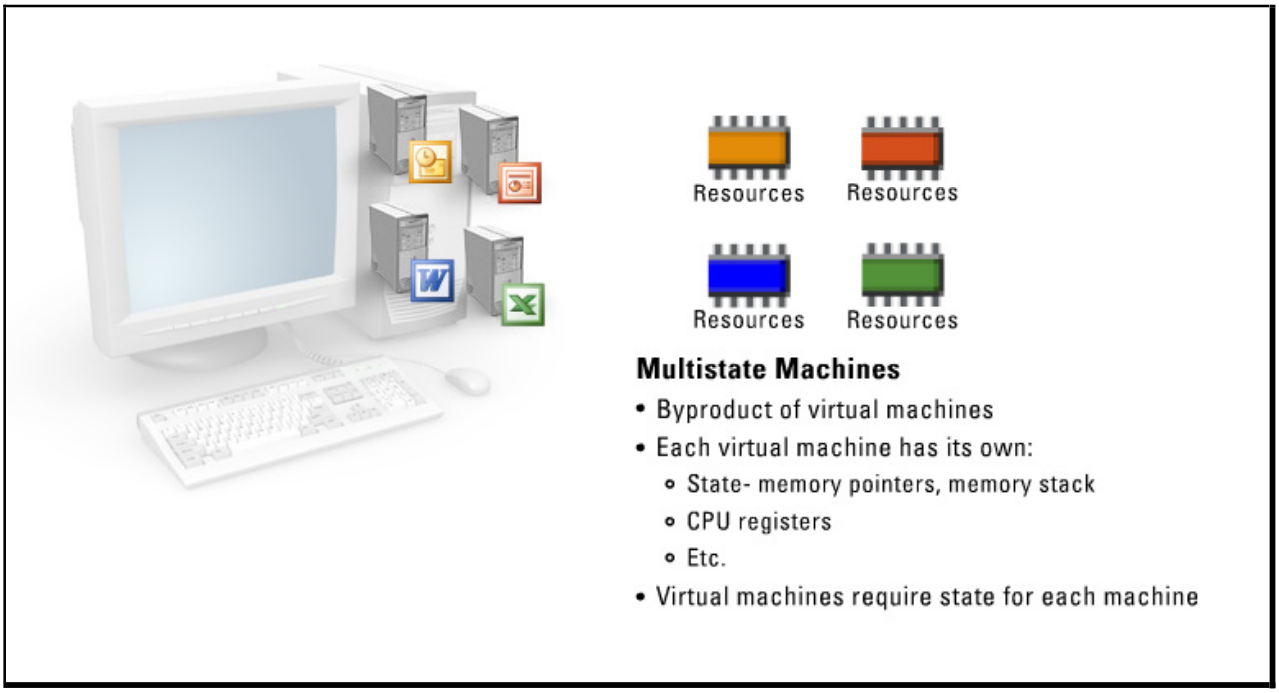

Each virtual machine has its own state, that is, memory pointers, memory stacks, CPU registers, etc., that are required to execute the application. Each virtual machine requires these state machines to be stored and retrieved when computing in its virtual machine.

Multitasking

Multitasking is the culmination of all virtual machines.

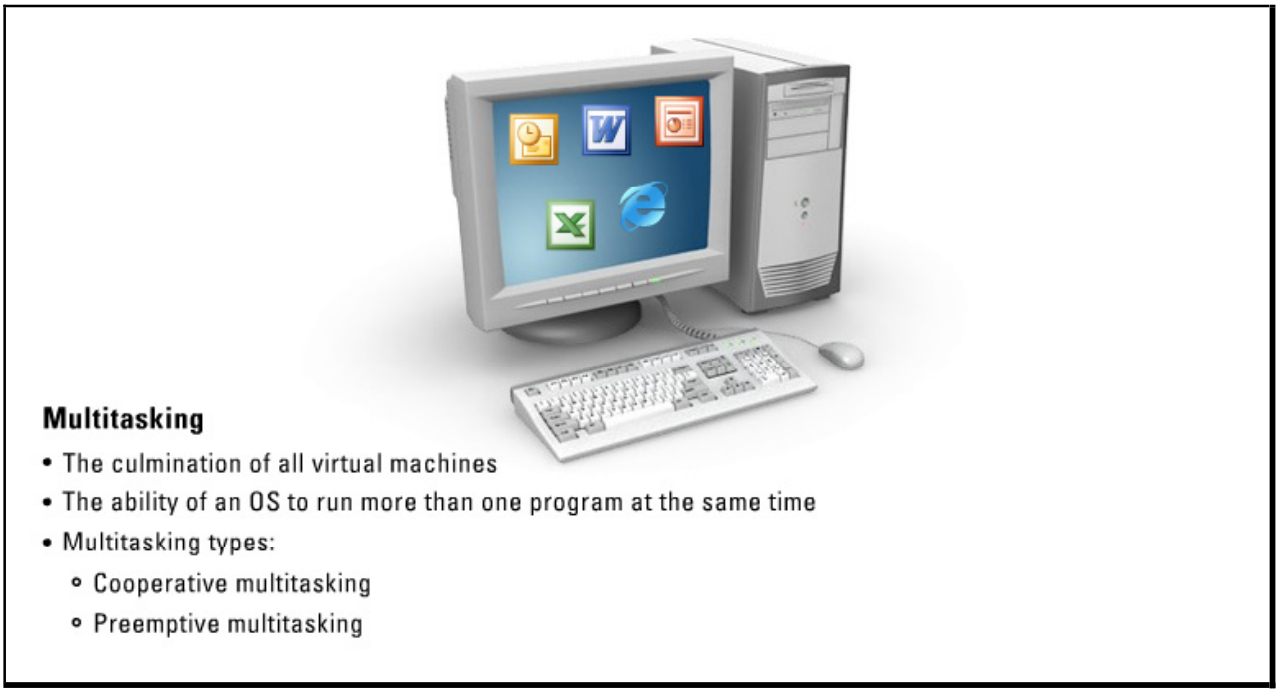

Multitasking is the ability of an operating system to run more than one program at the same time (hence the need for virtual machines).

There are different types of multitasking:

- **Cooperative multitasking** requires a program to be written to allow other programs to access the system.
- In **preemptive multitasking**, the system can suspend any program to allow other programs access. Preemptive multitasking provides better performance, as programs can switch with less overhead.

Multiuser Machines

Multiuser machines refer to computer systems that support two or more simultaneous users.

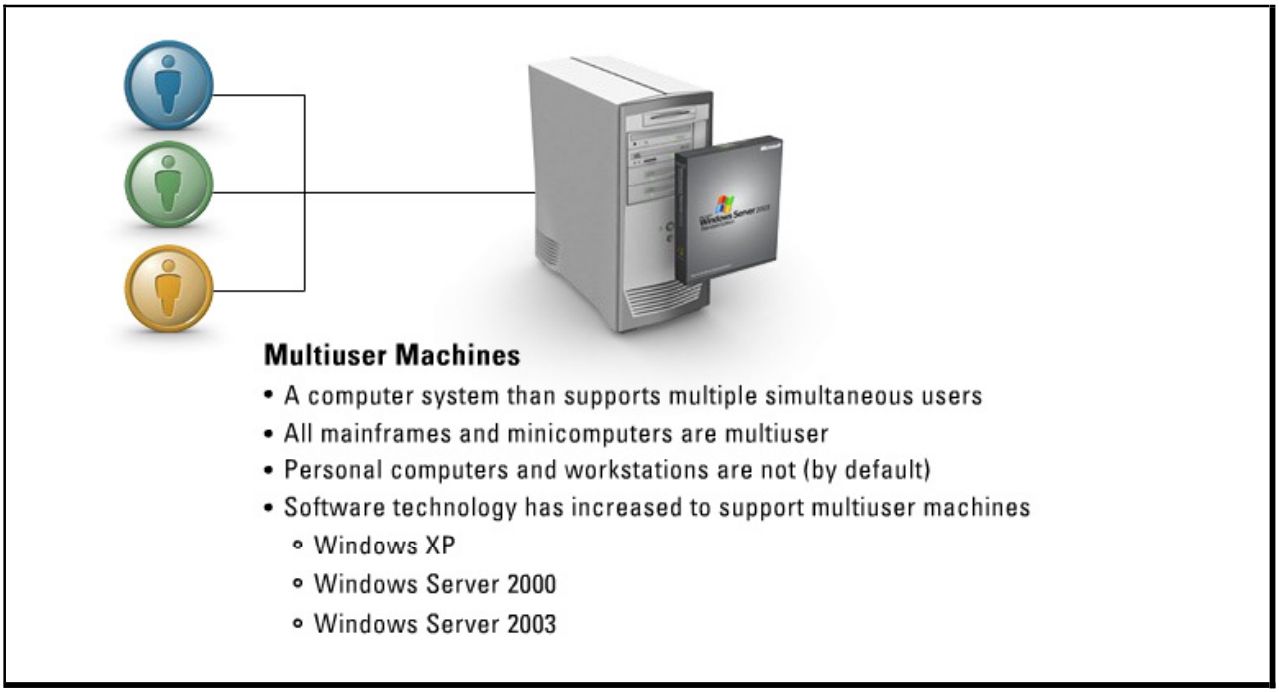

All mainframes and minicomputers are multiuser systems, but most personal computers and workstations are not. Software technology has increased to where multiuser systems are becoming more common. Windows XP has the ability for multiuser state machines to exist and execute simultaneous applications although only a single user has input/output capability at any one time. Windows Server 2000 and 2003 have the ability for terminal services, where many multiuser state machines can exist and operate independently by many remote users.

Network Protocol Functions

Network protocols function as a common language two systems use to communicate.

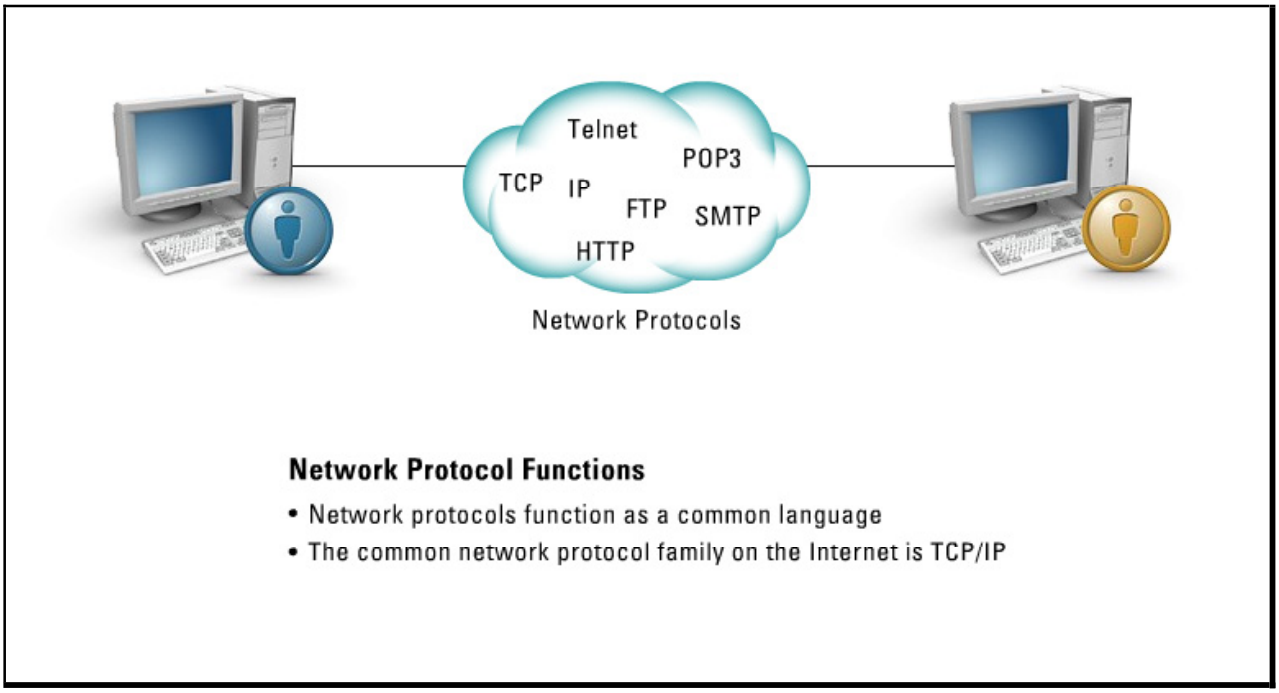

In a simple analogy, one person is asking, “Where is the airport?” in Spanish, but the other person only speaks Chinese. These two people will not be able to communicate, as they do not speak a common language. Network protocols are created to allow different computer systems to “speak” the same language so they can communicate. The common network protocols in use on the Internet are the Transmission Control Protocol/Internet Protocol (TCP/IP) family of protocols.

Resource Manager Functions

A **resource manager** makes sure that all subjects in a system have the required access to resources.

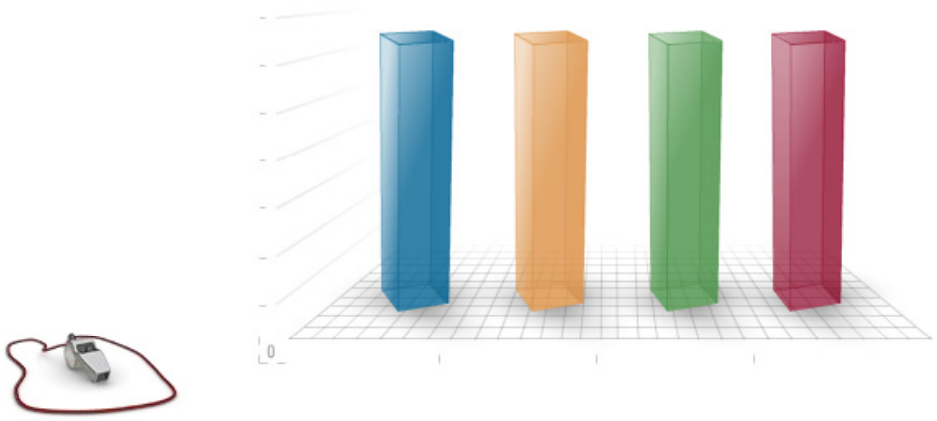

Resource Manager Functions

- Makes sure all subjects in a system have required access to resources
- Stops entities from consuming more than their share of resources
- Monitors if entity is no longer using its assigned resources

If one entity begins to consume a limited resource, the resource manager will police the situation to limit the amount of resources that entity can consume in order for all other entities to have access to the resource.

The resource manager will also monitor if an entity is no longer using its assigned resources. If a computer program is assigned resources, but locks up due to program malfunction or does not tear down and release the resources after use, the resource manager can release those resources so other programs or processes can consume them.

Summary

The key points discussed in this lesson are:

- A security architecture will describe how a system is put together to satisfy security requirements. It is more of a design overview and not a description of the functions of the system.
- There are many different types of addressing schemes that CPUs can use. These addressing schemes are dependent upon many factors such as hardware used, CPU type, memory type, and addressing type.
- As CPU programming and creation have evolved, so have their enhancements. These enhancements include: pipelining, complex instruction sets, reduced instruction sets, scalar processors, and superscalar processors.
- Protection rings provide strict boundaries and definitions on what the processes that work within each ring can access and what commands they can successfully execute.
- The different operating states in which a system can operate are ready, supervisory, problem, and wait.
- Hardware is any physical electronic device. Software is a broad term for the programs running on hardware. Firmware is software that is semi-permanently placed in hardware.
- Computers can run many different applications, programs, and utilities at the same time. They must follow a model to make sure each process has its share of CPU cycles and memory to perform its function.
- A real machine is a computer running a single application and using all memory and CPU it has at its disposal to accomplish the task.
- The operating system creates a virtual machine for each application to work in and allots each application a segment of virtual memory. Essentially, each application has its own virtual machine and virtual address segment and does not know that the other applications even exist.
- Multistate machines are a byproduct of virtual machines. Each virtual machine has its own state, that is, memory pointers, memory stacks, CPU registers, etc., that are required to execute the application.
- Multitasking is the ability of an operating system to run more than one program at the same time (hence the need for virtual machines).
- Multiuser machines refer to computer systems that support two or more simultaneous users.
- Network protocols function as a common language two systems use to communicate.
- A resource manager makes sure that all subjects in a system have the required access to resources.

Storage Types

Overview

You access data in a computer system from specialized memory chips. These chips have been specially designed to provide low latency when accessed as well as other specially designed features. This lesson will discuss the different types of memory storage as well as the security concerns inherent to them.

Importance

It is important that the information security professional understand how the system or end user accesses memory, and more importantly how an attacker can access the data stored in memory.

Objectives

Upon completing this lesson, you will be able to:

- Define common memory terms
- Define primary storage
- Define secondary storage
- Define real memory
- Define virtual memory

Outline

The lesson contains these topics:

- Memory
- Primary Storage
- Secondary Storage
- Real Memory
- Virtual Memory

Memory

This topic discusses random access memory and read-only memory.

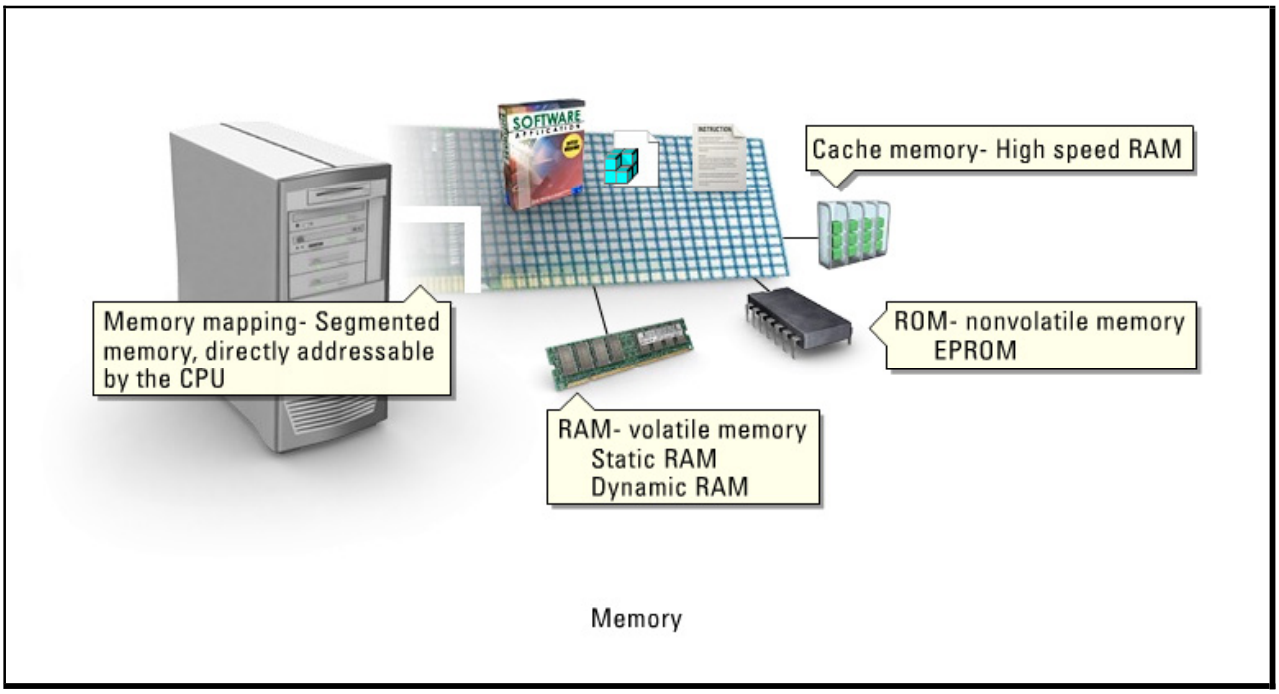

Memory

Random access memory(RAM) is a volatile memory, because when power is lost all information is lost. Types of RAM include:

- **Static RAM** - When it stores data, the data stays there without the need to be continually refreshed
- **Dynamic RAM** - Requires data held within it to be periodically refreshed because the data dissipates and decays

Read-only memory(ROM) is a nonvolatile memory. Software that is stored within ROM is called firmware.

Erasable and programmable read-only memory (EPROM) holds data that can be electrically erased or written to.

Cache memory is a part of RAM that is used for high-speed writing and reading activities.

Memory mapping can occur with real or primary memory. With memory mapping, memory is directly addressable by the CPU and used for the storage of instructions and data associated with the program that is being executed.

Primary Storage

Computer systems require many different types of devices to store information for short or long periods of time. Various storage solutions offer different access times, information volatility, and, of course, cost constraints, which affect the feasibility of use within the system.

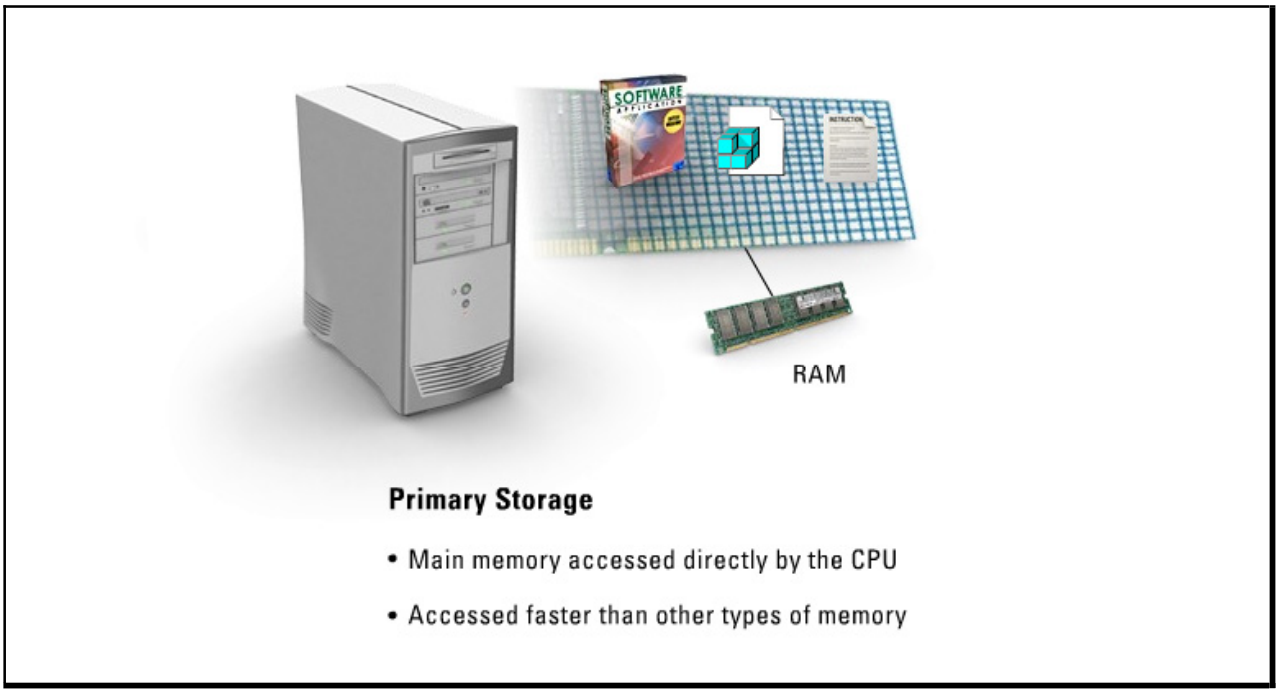

The types of storage devices include:

- Primary
- Secondary
- Real
- Virtual
- Volatile
- Nonvolatile
- Write-once read many

Primary storage is the main memory accessed directly by the computer's CPU. Systems access primary storage much faster than other types of memory. This type of memory is typically RAM. Primary storage is used to store data that is likely to be in active use, so it is usually faster than other types of storage.

Secondary Storage

Secondary storage usually comes in the form of nonvolatile storage mediums, such as disk drives that store the data until required by the system.

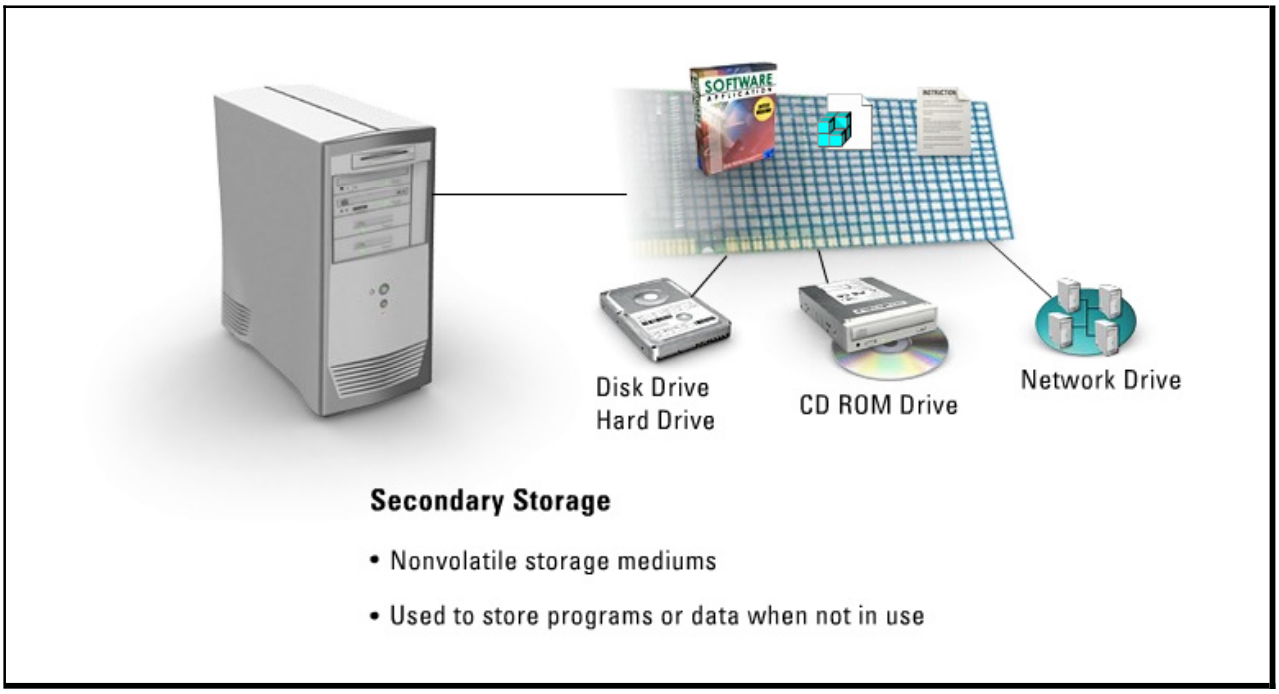

When data from secondary storage systems are required for use, either as an application program or data such as Word documents, the system will copy the required information from secondary storage to primary storage, where the CPU now has direct access and can manipulate or execute the information as needed.

Real Memory

Real memory is the memory that is given a definite storage location in memory and direct access to peripheral devices. Real memory is the real physical memory chips installed in the computer system.

Real Memory

- Given a definite storage location in memory
- Has direct access to peripheral devices
- Portion always reserved for kernel and its data
- The rest is termed available memory
- Programs can only execute in real memory

A portion of real memory is always reserved for the kernel and its data structures, which are dynamically allocated. The amount of real memory not reserved for the kernel and its data structures is termed available memory. Available memory is consumed by user processes and also nonkernel system processes such as network daemons. Because the size of the kernel varies depending on the number of interface cards, users, and values of the tunable parameters, available memory varies from system to system.

Virtual Memory

Real memory refers to the actual memory chips that are installed in the computer. All programs actually run in this physical memory. However, it is often useful to allow the computer to *think* that it has memory that is not actually there, in order to permit the use of programs that are larger than will physically fit in memory, or to allow multitasking (multiple programs running at once). This concept is called **virtual memory**.

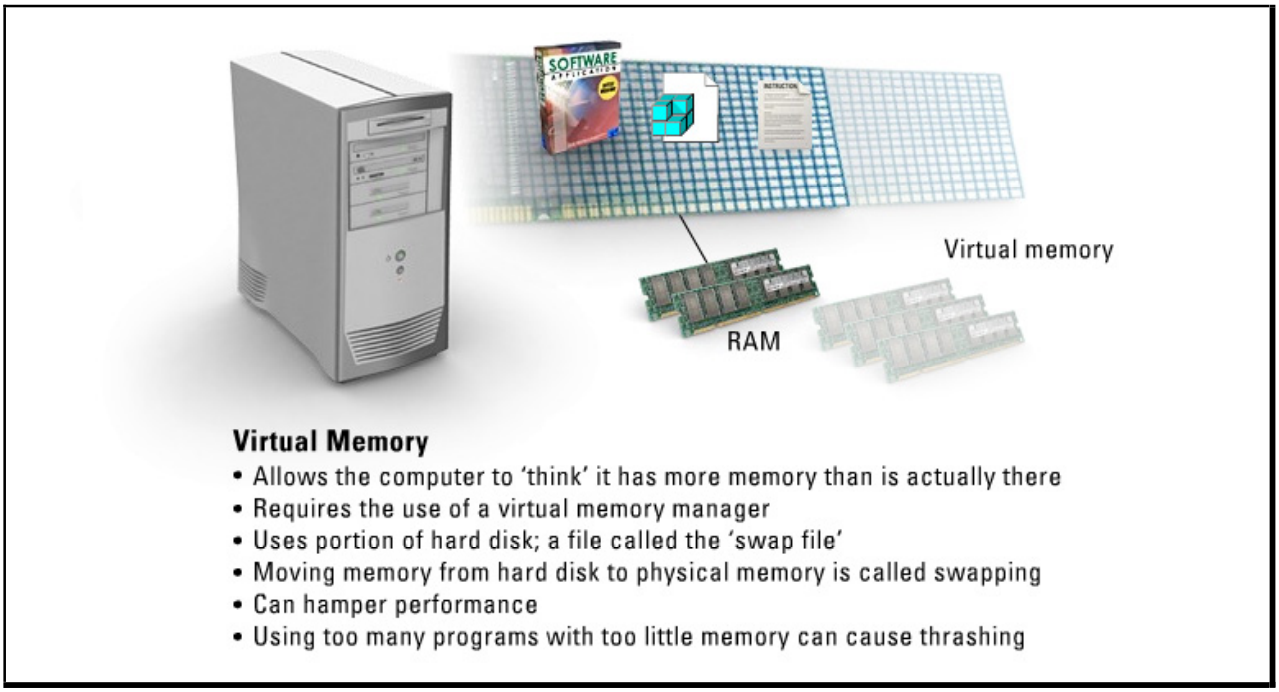

The way virtual memory works is relatively simple. Let us suppose the operating system needs 400 MB of memory to hold the operating system and all the programs that are running, but there are only 256 MB of RAM chips installed in the computer. The operating system sets up 400 MB of virtual memory and employs a virtual memory manager, a program designed to control virtual memory. The virtual memory manager sets up a file on the hard disk that is 144 MB in size (400 minus 256). The operating system then proceeds to use 400 MB worth of memory addresses. To the operating system, it appears as if 400 MB of memory exists. It lets the virtual memory manager worry about how to handle the fact that we only have 256 MB of real memory.

Of course, not all of the 400 MB will fit in the physical 256 MB that do exist. The other 144 MB reside on a disk, in the file controlled by the virtual memory manager. This file is called a swap file. Whenever the operating system needs a part of memory that is currently not in physical memory, the virtual memory manager picks a part of physical memory that has not been used recently, writes it to the swap file, and then reads the part of memory that is needed from the swap file and stores it into real memory in place of the old block. This process is called **swapping**. The blocks of memory that are swapped around are called **pages**.

Virtual memory can also hamper performance. The larger the virtual memory is compared to the real memory, the more swapping has to occur to the hard disk. The hard disk is much, much slower than the system memory. Trying to use too many programs at once in a system with too little memory will result in constant disk swapping, called **thrashing**.

Summary

The key points discussed in this lesson are:

- RAM is a volatile memory, because when power is lost all information is lost. ROM is a nonvolatile memory. Software that is stored within ROM is called firmware.
- Primary storage is the main memory accessed directly by the computer's CPU. Systems access primary storage much faster than other types of memory. This type of memory is typically RAM.
- Secondary storage usually comes in the form of nonvolatile storage mediums, such as disk drives that store the data until required by the system.
- Real memory is the memory that is given a definite storage location in memory and direct access to peripheral devices.
- It is often useful to allow the computer to *think* that it has memory that is not actually there, in order to permit the use of programs that are larger than will physically fit in memory, or to allow multitasking (multiple programs running at once). This concept is called virtual memory.

Principles of Common Security Models

Overview

Security in the enterprise is of wide concern to all security professionals as well as corporate leaders and stockholders. To this end, many security models have been devised, modified, and improved upon to help enterprises achieve the goal of a secure network. This lesson will discuss the various security models as well as controls used in securing today's networks.

Importance

It is important to understand what others have discovered or implemented in securing their networks. This allows the information security professional the ability to learn from others and not make the same mistakes. Understanding security models is vital to the information security professional in protecting enterprise resources.

Objectives

Upon completing this lesson, you will be able to:

- Identify the reason for conducting a security evaluation
- Define certification and accreditation
- Differentiate open and closed systems
- Differentiate active and passive protection
- Identify three security technologies operating systems use to protect software security features
- Define Trusted Computing Base
- Define Reference Monitors and Security Kernels
- Describe the logistics of Mandatory Access Controls
- Describe the logistics of Discretionary Access Controls
- Identify an ideal environment for IPSec architecture
- Identify the security topics included in the Trusted Computer System Evaluation Criteria

- Identify the main reason for the creation of the Rainbow Series
- Identify the security topics included in the Red Book
- Identify the security topics included in the Information Technology Security Evaluation Criteria
- Identify the main reason for the creation of the Common Criteria standard
- Define security model
- Describe the logistics of the state machine security model
- Describe the logistics of the Bell-LaPadula security model
- Describe the logistics of the Biba security model
- Describe the logistics of the Clark-Wilson security model

Outline

The lesson contains these topics:

- Security Evaluation
- Certification and Accreditation
- Open and Closed Systems
- Active and Passive Protection
- Controls
- Trusted Computing Base
- Reference Monitors and Security Kernels
- Mandatory Access Controls
- Discretionary Access Controls
- Internet Protocol Security Architecture
- Trusted Computer System Evaluation Criteria
- Rainbow Series
- Trusted Network Interpretation
- Information Technology Security Evaluation Criteria
- Common Criteria
- Security Models
- State Machine Models
- Bell-LaPadula Model
- Biba Model
- Clark-Wilson model

Security Evaluation

You can use a **security evaluation** to examine the security-relevant parts of your system, meaning the trusted computing base (TCB), access control mechanisms, reference monitor, kernel, and other protection mechanisms.

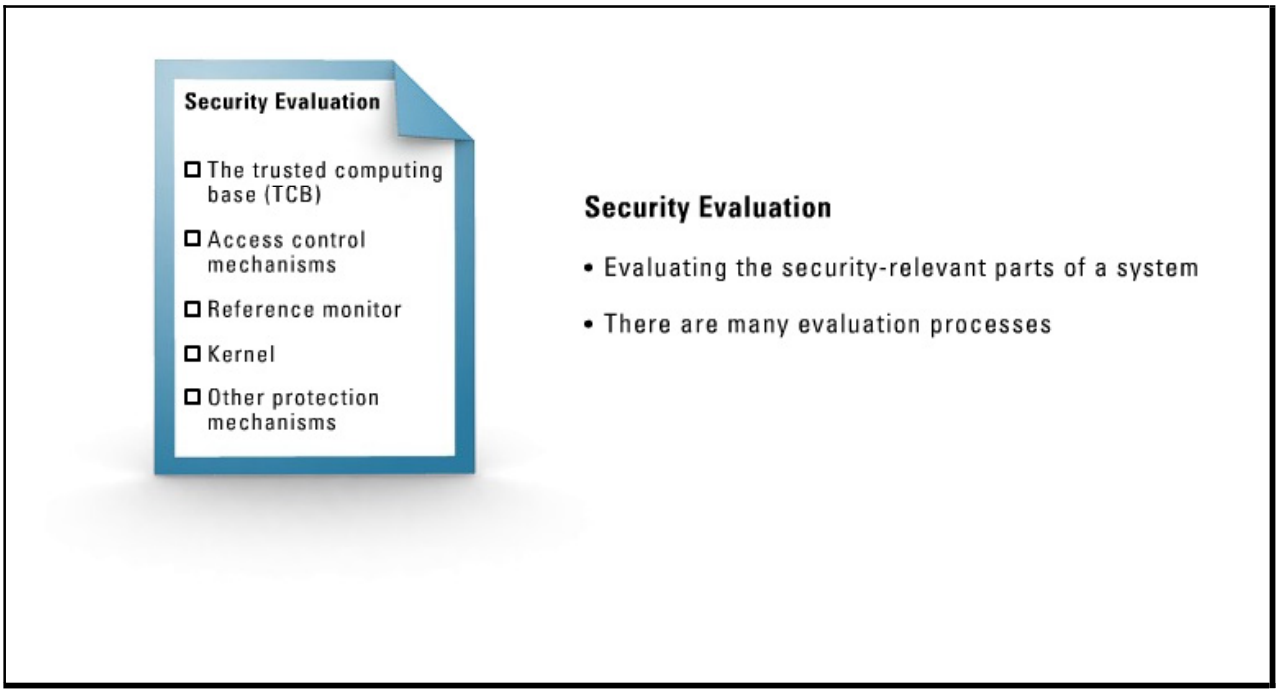

There are different methods of evaluating and assigning security and trust levels to systems. There is more than one type of security evaluation process because the methodologies and ideologies have evolved over time and because different parts of the world look at computer security differently and rate some aspects of security higher than others.

Note	The term “trusted computing base” originated from the Orange Book and does not address the level of security a system provides, but instead addresses the level of trust a system provides.

Certification and Accreditation

The technical evaluation of the security components and their compliance for the purpose of accreditation is called **certification**. The certification process can use evaluation mechanisms, risk analysis mechanisms, verification, testing, and auditing techniques to assess the appropriateness of a specific system, which processes a certain classification level of information within a particular environment.

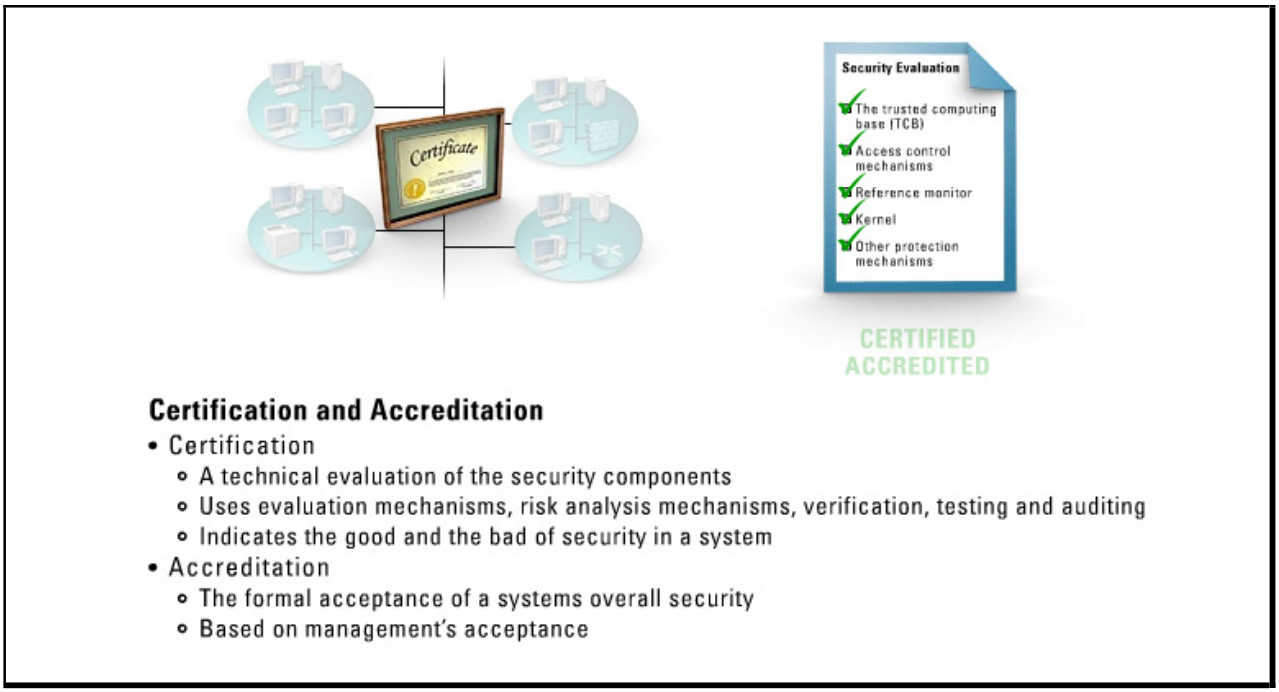

Certification is performed when a company performs tests on software configurations, hardware, firmware, design, implementation procedures, system procedures, and physical and communication controls.

The outcome of the certification process will indicate the good and the bad about the security protection level and all mechanisms that support it within the system. If the outcome of certification looks good, management would then start the accreditation process.

The formal acceptance of the adequacy of a system's overall security by management is called **accreditation**. When the certification information is presented to management, it is up to management to ask the proper security related questions, review reports and findings, and decide upon the acceptance of the safeguards and if any corrective actions are required.

Once management is satisfied with the system's overall security as it is presented, they will make a formal accreditation statement. When they do this, management is stating that they understand the level of protection the system will provide in its current environment and understand the security risks associated with installing a new system.

Note	Certification is the process of assessing the security mechanisms, and accreditation is management 's official acceptance of the information in the certification process findings.

Open and Closed Systems

This topic discusses the features of open and closed systems.

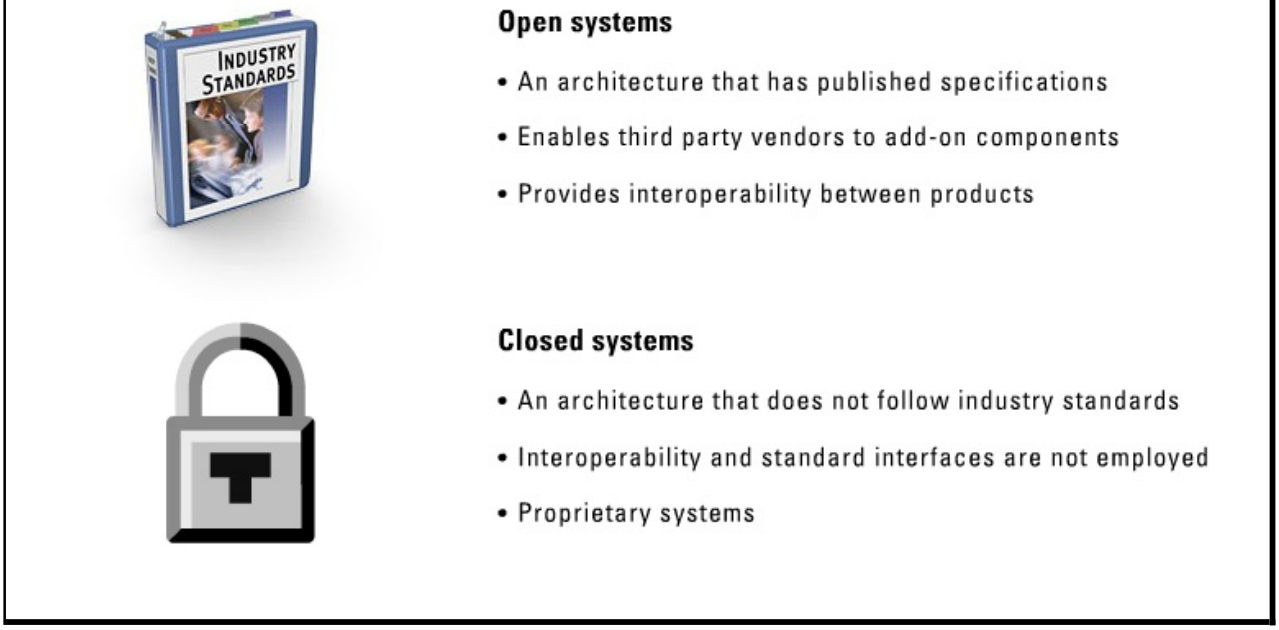

An **open system** is an architecture that has published specifications, which enables third-party vendors to develop add-on components and devices. Open systems provide interoperability between products by different vendors of different operating systems, applications, and hardware devices.

A **closed system** is an architecture that does not follow industry standards. In a closed system, interoperability and standard interfaces are not employed to enable easy communication between different types of systems and add-on features. They are proprietary, meaning that the system can only communicate with like systems.

Active and Passive Protection

This topic discusses the features of active and passive protection.

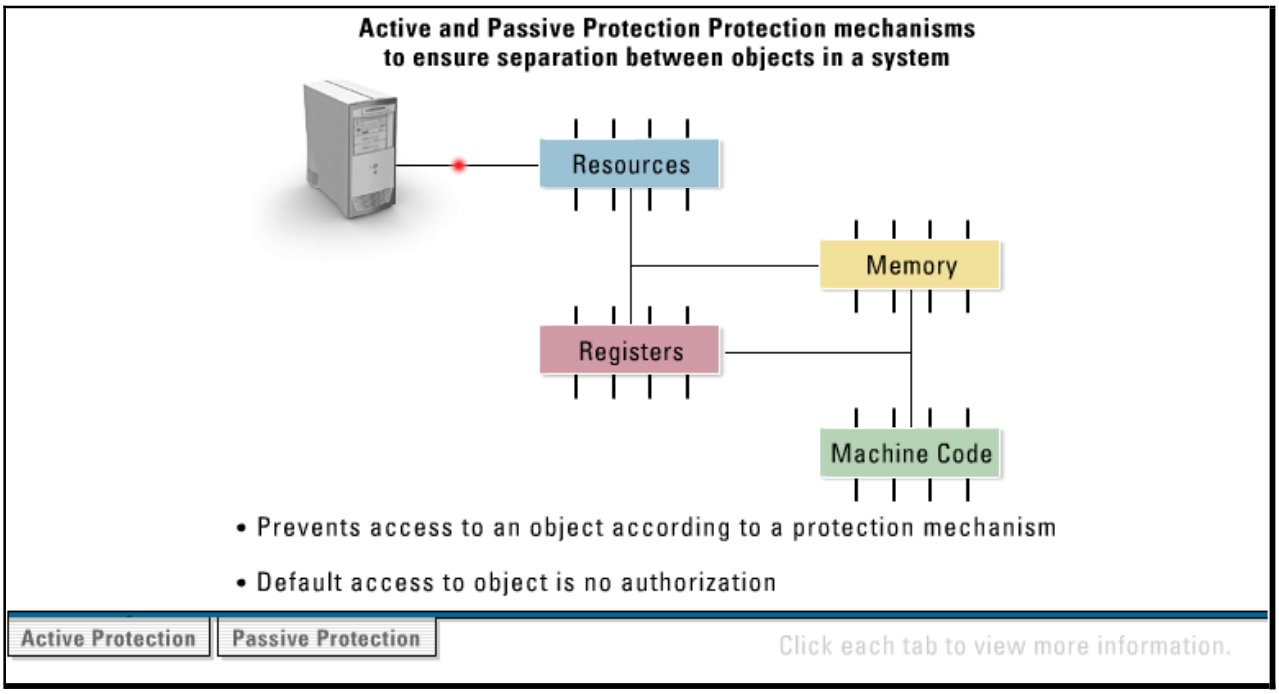

You use a protection mechanism to ensure the separation between objects in a system. There are two categories of protection mechanisms: active and passive. You use **active protection** mechanisms to prevent access to an object if, according to the protection mechanism, the access is not authorized. As an example, an active protection mechanism is when memory protection is provided by the hardware architecture, where access to objects at certain addresses may be controlled depending on such criteria as the current processor state and other attributes of the process attempting to reference the object in question.

Passive protection mechanisms, on the other hand, are those that prevent or detect unauthorized use of the information associated with an object, even if access to the object itself is not prevented. In most cases, passive protection mechanisms use cryptographic techniques to prevent unauthorized disclosure of information or checksum techniques to detect unauthorized alteration of an object.

You often implement protection mechanisms in operating systems, hardware, or firmware. You can use them to control access to primitive objects such as memory, machine code instructions, and registers, which are the functions used to control input/output (I/O) operations.

Controls

The operating system is the only line of defense between users, programs, applications, utilities, and the low level functions that are the core security features of the system.

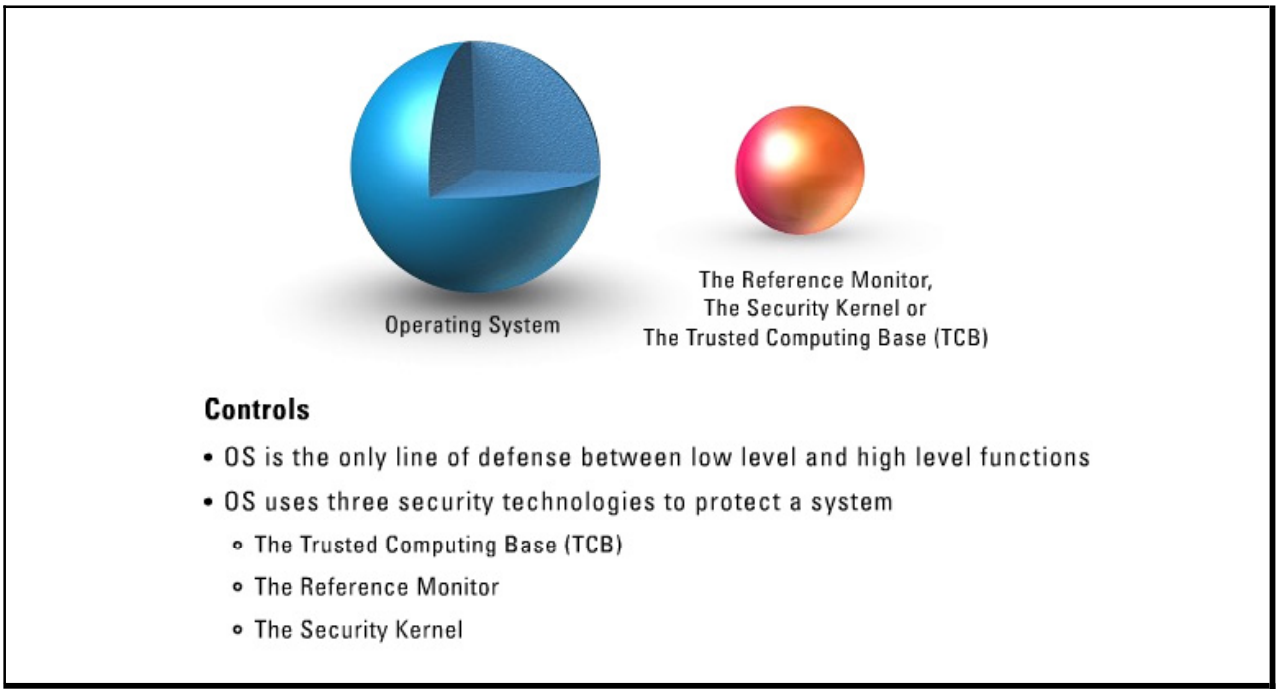

Operating systems use three security technologies to protect security features of the software:

- The Trusted Computing Base (TCB)
- The Reference Monitor
- The Security Kernel

Trusted Computing Base

The **Trusted Computing Base (TCB)** is defined as the total combination of protection mechanisms within a computer system. These mechanisms include all hardware, software, firmware, and processes, etc., that, combined, are responsible for enforcing a security policy.

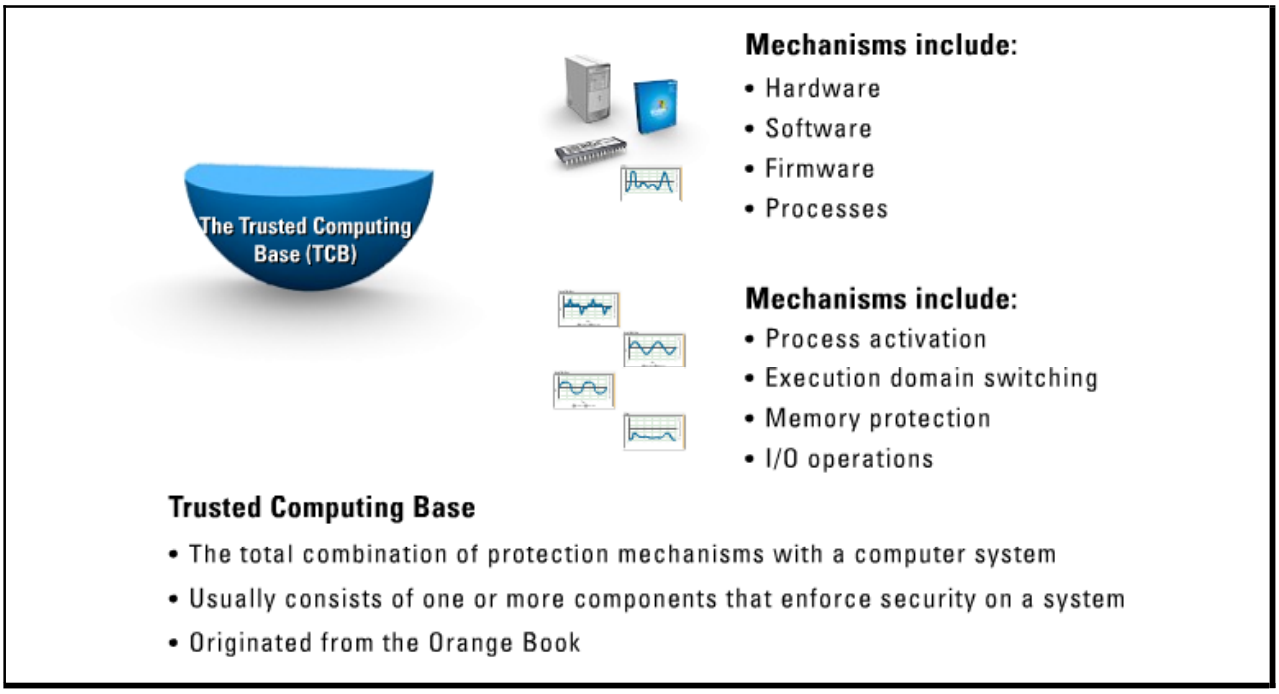

The TCB usually consists of one or more components that together enforce a unified security policy over a system. To correctly enforce a security policy, the TCB must depend solely on the mechanisms within it and on the correct input of parameters related to the security policy by system administrators.

The TCB originated from the Orange Book and monitors four basic functions:

- Process activation
- Execution domain switching
- Memory protection
- I/O operations

Note	If an operating system implements the TCB, it usually follows the reference monitor concept.

Reference Monitors and Security Kernels

A **Reference Monitor** is an abstract machine, which mediates all the access subjects have to objects in order to ensure that the subjects have the necessary access rights and to protect the objects from unauthorized access and destructive modification.

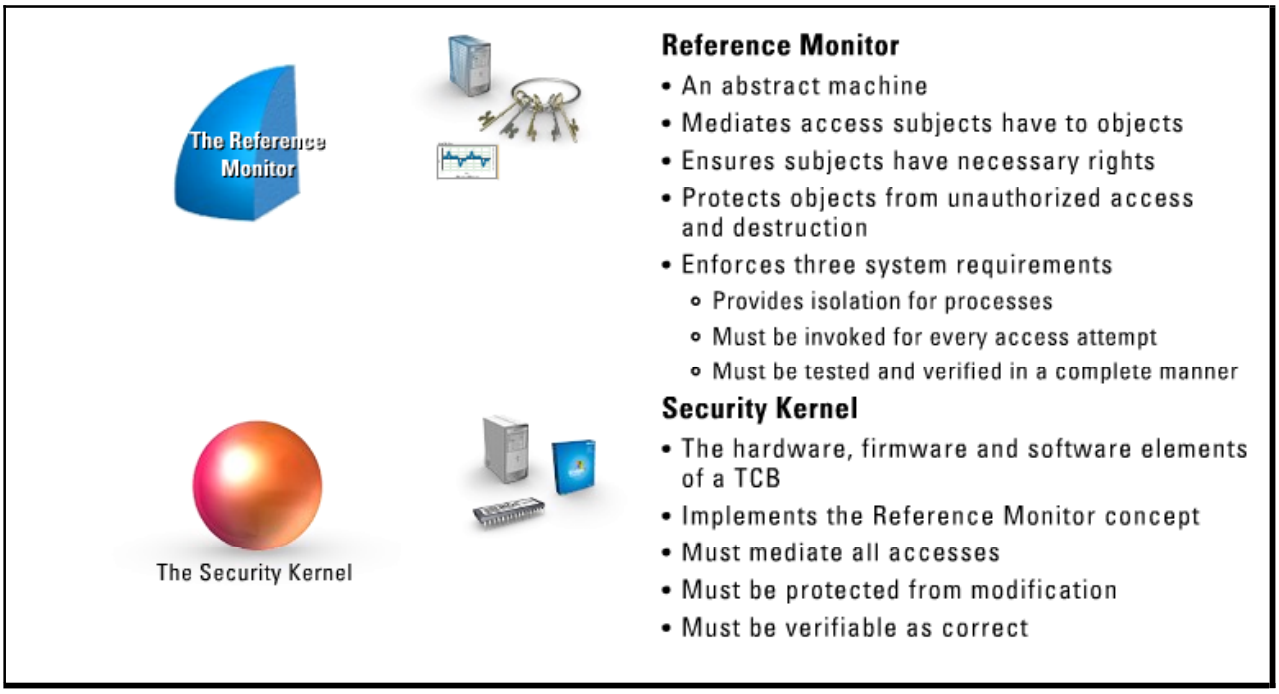

The Reference Monitor is an access control concept, not an actual physical component. It is the core of the TCB and is the most commonly used approach to building trusted computing systems. The Reference Monitor enforces three system requirements:

- The system must provide isolation for the processes carrying out the Reference Monitor concept, and they must be tamperproof.
- The reference monitor must be invoked for every access attempt and must be impossible to circumvent. Thus, the reference monitor must be implemented in a complete and foolproof way.
- The system must be small enough to be able to be tested and verified in a complete and comprehensive manner.

The **Security Kernel** is the hardware, firmware, and software elements of a TCB implementing the Reference Monitor concept. The Security Monitor must satisfy the following principles:

- It must mediate all accesses (completeness).
- It must be protected from modification (isolation).
- It must be verifiable as correct (verifiable).

In essence, you must implement the Security Monitor in such a fashion that you prevent users from modifying the operating system. If a user was able to disable or circumvent the protection mechanisms by modifying the OS, he or she could compromise the integrity of the system.

Mandatory Access Controls

Information systems that process highly sensitive data use **Mandatory Access Controls(MACs)** as part of their access control policies. MACs are considered a policy-based control in an information system and are used to assign sensitivity labels to all system subjects and all system objects.

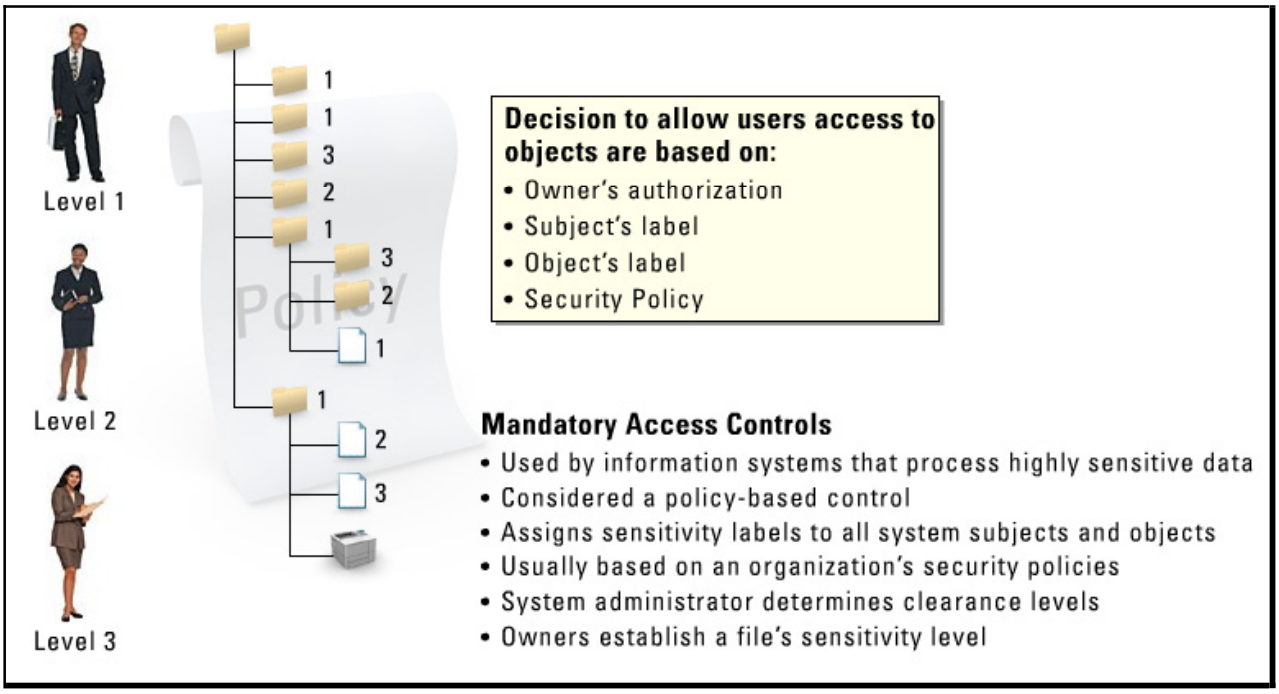

MACs use sensitivity labels to determine who can access what information. MACs are usually based on an organization's security policies, which means the system administrator and the owner of the information manage the policies. The system administrator should install the clearance levels and the owner should be responsible for establishing the files' sensitivity levels. MACs require that object owners be allowed to access objects only in accordance with labeling. The decisions to allow users to access objects are based on an owner's authorization, the subject's label, the object's label, and the security policy.

Tip	Essentially, MACs limit user access so that a user at a lower level cannot access information labeled at a higher level.

Discretionary Access Controls

Discretionary Access Controls(DACs) are policies that restrict access to files and other system resources based on the identity and assignment of the user and/or the groups to which the user belongs. DACs are considered a policy-based control because they permit resource owners to specify who has access and what privileges other users will have. A simple DAC is one that is implemented using an Access Control List (ACL).

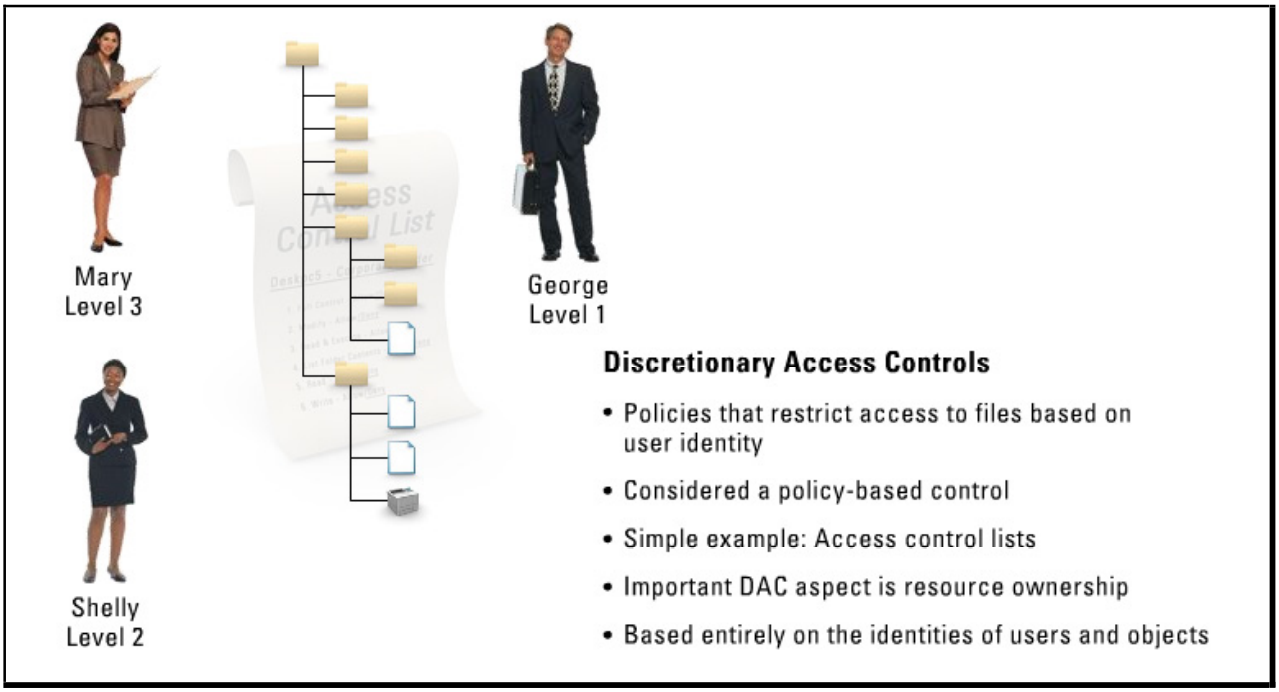

The file structure on a storage system can also use a DAC. For example, in a tree-structured file system, there are directories and files. You can apply DACs to both of these elements. This type of control will allow users to not only specify who has access, but also what type of access others will have. In most operating systems, there are generic types of access to files—none, read, write, execute, delete, change, and full control.

An important DAC aspect is resource ownership. In some systems, the user who creates the file is considered the owner of the file. This user is then granted full control over the file, including the ability to set permissions on the file. If you have a user who is not the owner of the file, he or she does not have permissions to set access controls.

DAC mechanisms are based entirely on the identities of users and objects in the system.

Internet Protocol Security Architecture

This topic discusses the **Internet Protocol Security (IPSec) architecture**.

Internet Protocol Security Architecture

- IPSec
- Controls used over and above any controls provided by the network
- Typically implemented on end systems
- Secure integrity of transmissions
- Usually only necessary on highly critical systems

Controls that could be used over and above any controls provided in the network (i.e., telecommunications) system are another important control mechanism. These are the controls that would typically be implemented on the end system (host) as part of an overall application security strategy to secure the integrity of the transmissions. These types of controls are usually only necessary on highly critical systems where the protection of data is paramount. IPSec architecture is one such control.

Trusted Computer System Evaluation Criteria

The U.S. Department of Defense developed the **Trusted Computer System Evaluation Criteria (TCSEC)**. These criteria are used to evaluate products to assess if they contain the security properties they claim to have and evaluate if the products are appropriate for specific applications or functions. The evaluation criteria were published in a book with an orange cover, which is why it is usually called the **Orange Book**.

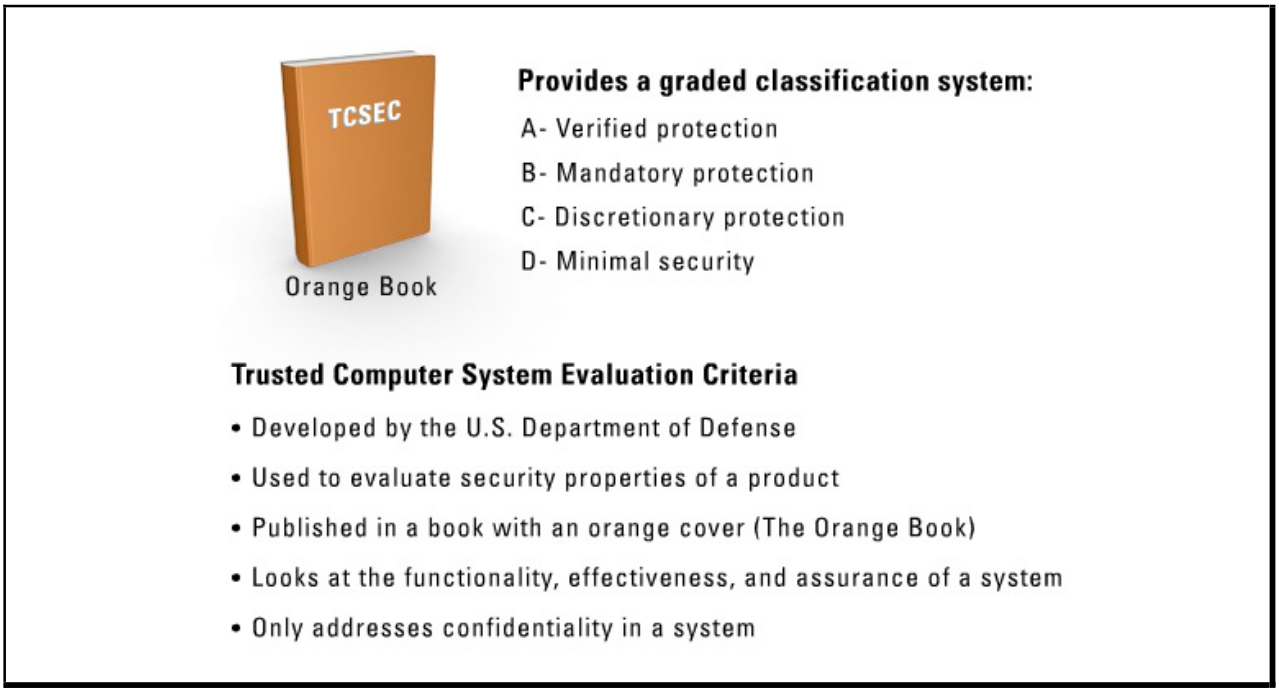

The Orange Book looks at the functionality, effectiveness, and assurance of a system during its evaluation, and it uses classes that were devised to address typical patterns of security requirements. TCSEC provides a graded classification of systems that is divided into hierarchical divisions of security levels:

- A - Verified protection
- B - Mandatory protection
- C - Discretionary protection
- D - Minimal security

The highest classification level is level A, and it represents the highest level of security; level D then represents the lowest level of security. Each of the above divisions can also have one or more numbered classes, each of which have a corresponding requirement that must be met for a system to achieve that particular rating. In this system, the classes with higher numbers indicate a greater degree of trust and security. These criteria include the following topics:

- **Security Policy** - Must be explicit, well defined, and enforced by the mechanisms within the system
- **Identification** - Individual subjects must be uniquely identified
- **Labels** - Access control labels must be properly associated with objects

- **Documentation** - Includes test, design, and specification documents, and user guides and manuals
- **Accountability** - Audit data must be captured and protected to enforce accountability
- **Life Cycle Assurance** - Software, hardware, and firmware must be able to be tested individually to ensure that each component enforces the security policy in an effective manner throughout its lifetime
- **Continuous Protection** - The security mechanisms and the system as a whole must continuously perform predictably and acceptably in different situations

Remember that these categories are evaluated independently, but their assigned ratings do not specify each of these objectives individually. The rating is a sum total of these items. When a consumer is interested in certain products and systems with a particular security rating, he or she can check the Evaluated Product List (EPL) to find out the security level assigned to a specific product.

Evaluation levels:

- D - Minimal protection
- C1 - Discretionary security protection
- C2 - Controlled access protection
- B1 - Labeled security
- B2 - Structured protection
- B3 - Security domains
- A1 - Verified design

Remember that TCSEC only addresses confidentiality in a system, not integrity, availability, or authenticity. Functionality of the security mechanisms and the assurance of those mechanisms are not evaluated separately, but rather combined and rated as a whole.

Problems with the Orange Book include:

- It is based on the old Bell-LaPadula model.
- It stands alone, offering no way to network systems.
- Systems take a long time (one to two years) to certify.
    - Any changes (i.e., hot fixes, service packs, and patches) break the certification.
- It has not adapted to changes in client-server and corporate computing.
- Certification is expensive.
- For the most part, it is not used outside of the government sector.

Rainbow Series

This topic introduces a collection of books called the Rainbow Series.

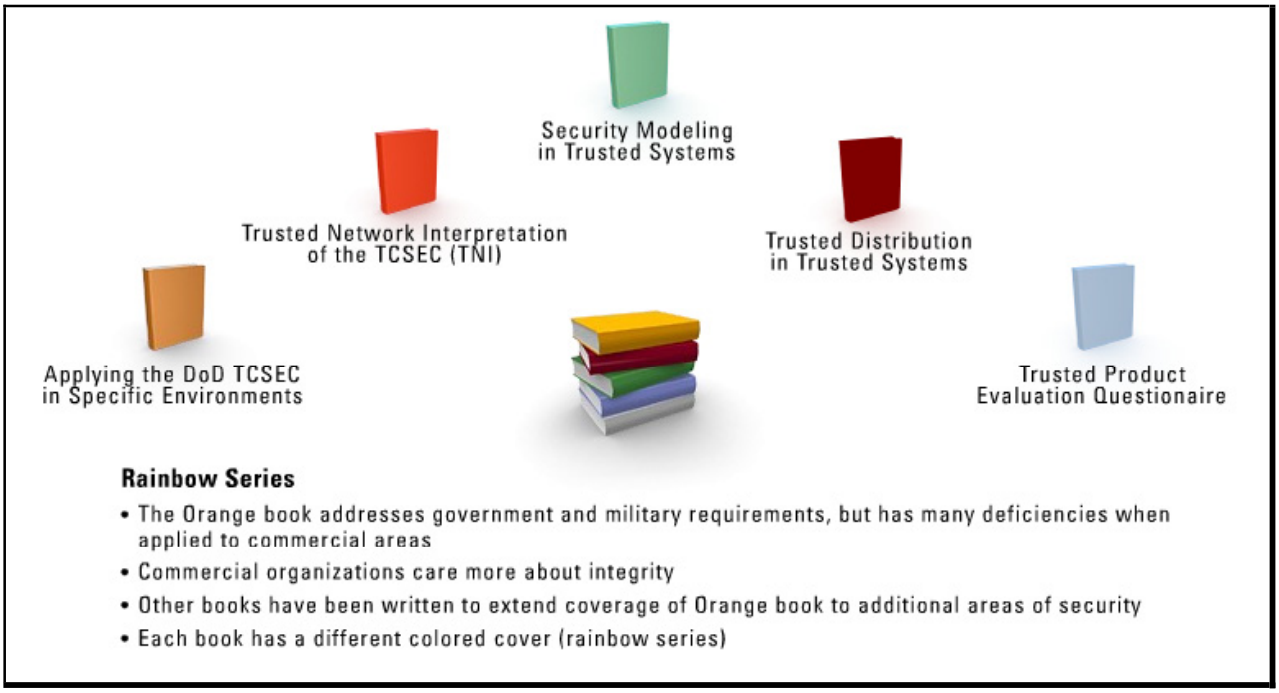

The Orange Book addresses government and military requirements and expectations regarding their computer systems, but the Orange Book has many deficiencies when applied to commercial areas. The Orange Book also places great emphasis on controlling which users can access a system, but virtually ignores controlling what those users do with the information once they have authorization to access it. On the other hand, commercial organizations have expressed more concern about the integrity of their data. Because of these differing goals, the Orange Book is a better evaluation tool for government and military systems.

Other books have been written to extend the coverage of the Orange Book into additional areas of security. These books are collectively called the **Rainbow Series** because each book has a different colored cover.

Trusted Network Interpretation

The **Trusted Network interpretation (TNI)**, also called the **Red Book**, is used to extend the Orange Book to networks. Remember that the Orange Book addresses single-system security, but networks are many systems, each with connectivity to the other one. The Red Book addresses isolated local area network (LAN) and wide area network (WAN) systems.

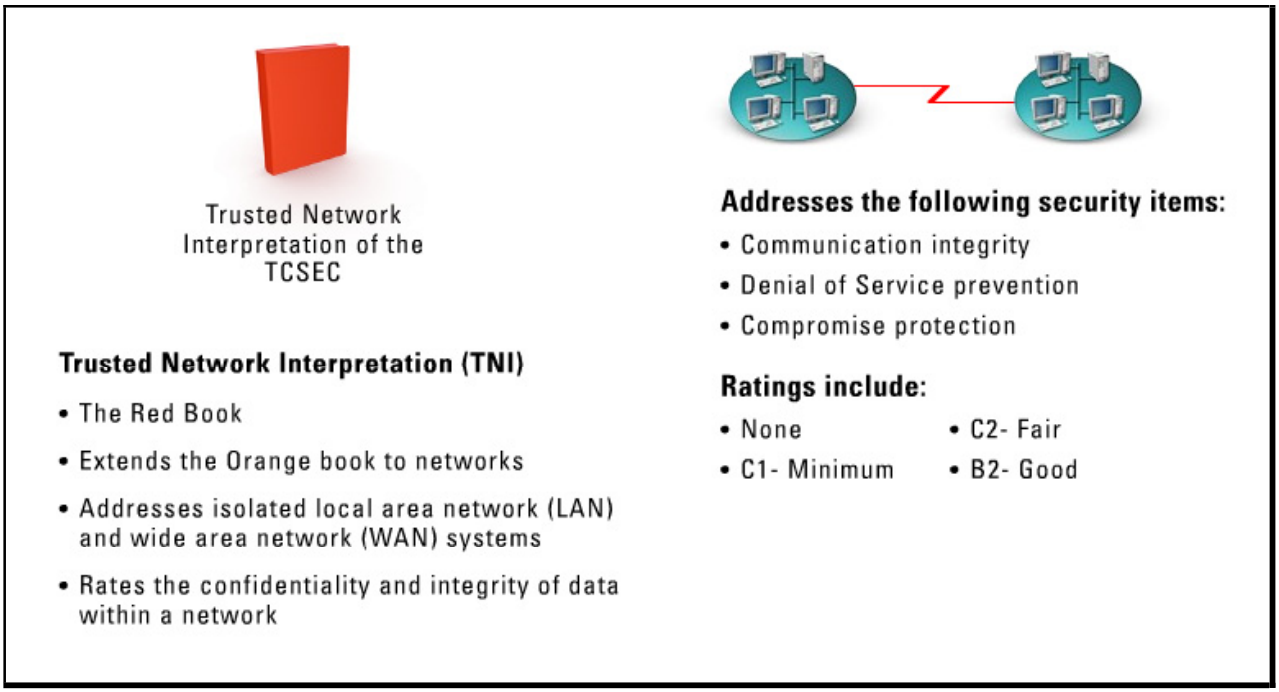

The Red Book, like the Orange Book, does not supply details on how to implement security mechanisms, but instead provides a framework for securing different types of networks. In a network, the subject could be a workstation and an object could be a network service on a server. The Red Book rates the confidentiality and integrity of data and operations within a network, as well as the network products themselves.

The Red Book addresses the following security items:

- Communication integrity
  - Authentication
  - Message integrity
  - Non-repudiation
- Denial of service prevention
  - Continuity of operation
  - Network management
- Compromise protection
  - Data confidentiality
  - Traffic flow confidentiality
  - Selective routing

The Red Book ratings are:

- None
- C1- Minimum
- C2- Fair
- B2- Good

Information Technology Security Evaluation Criteria

European countries created the **Information Technology Security Evaluation Criteria (ITSEC)** in an attempt to establish a single standard for evaluating security attributes of computer systems. ITSEC is only used in Europe. The ITSEC is Europe's version of the Orange Book and Rainbow Series from the United States.

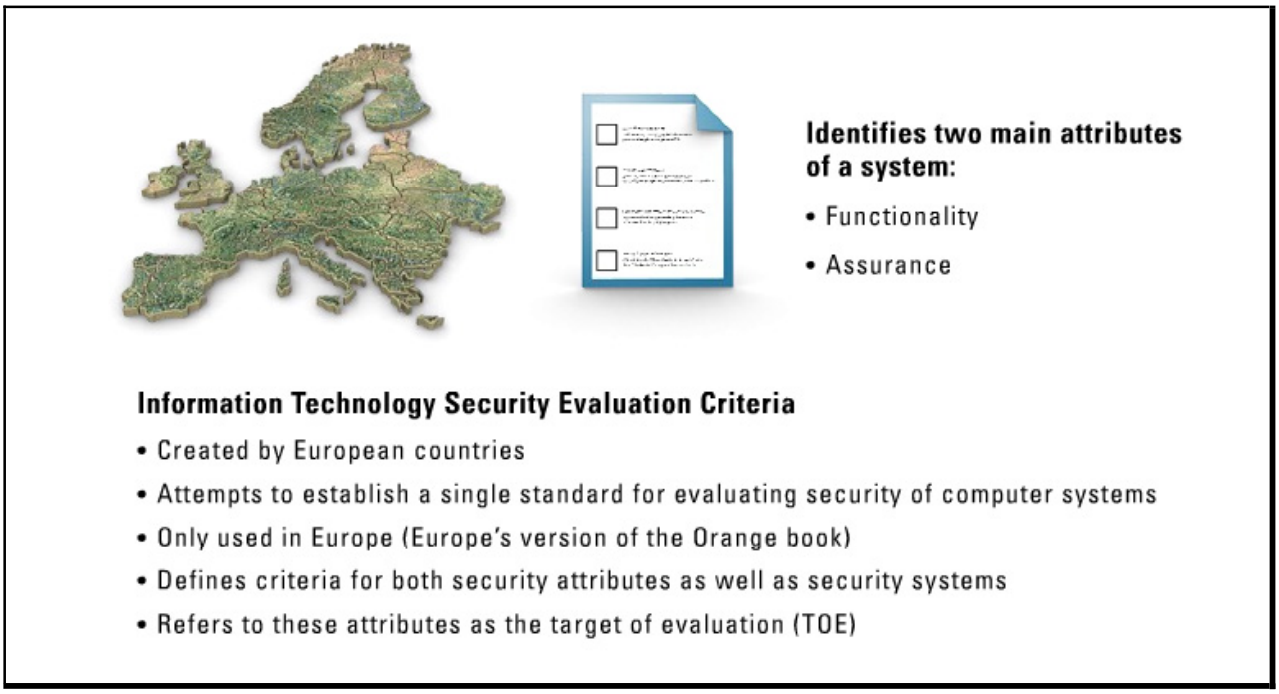

The ITSEC identifies two main attributes of a system for security evaluation:

- Functionality
- Assurance

Because of differing underlying mechanisms providing for functionality, it is possible for two systems to have the same functionality, but still have different assurance levels. Because of this possibility, the ITSEC rates these two attributes separately, whereas TCSEC lumps them together and assigns them a single rating (D through A1).

ITSEC defines criteria for both security attributes as well as security systems, and refers to these attributes as the target of evaluation (TOE).

Common Criteria

Because the Orange Book was used in the United States and the ITSEC was used in Europe, any attempt to provide consistent security across borders was a management nightmare. A new common set of security evaluation criteria was required that could be used worldwide. The International Organization for Standardization (ISO) created a standard called the **Common Criteria**.

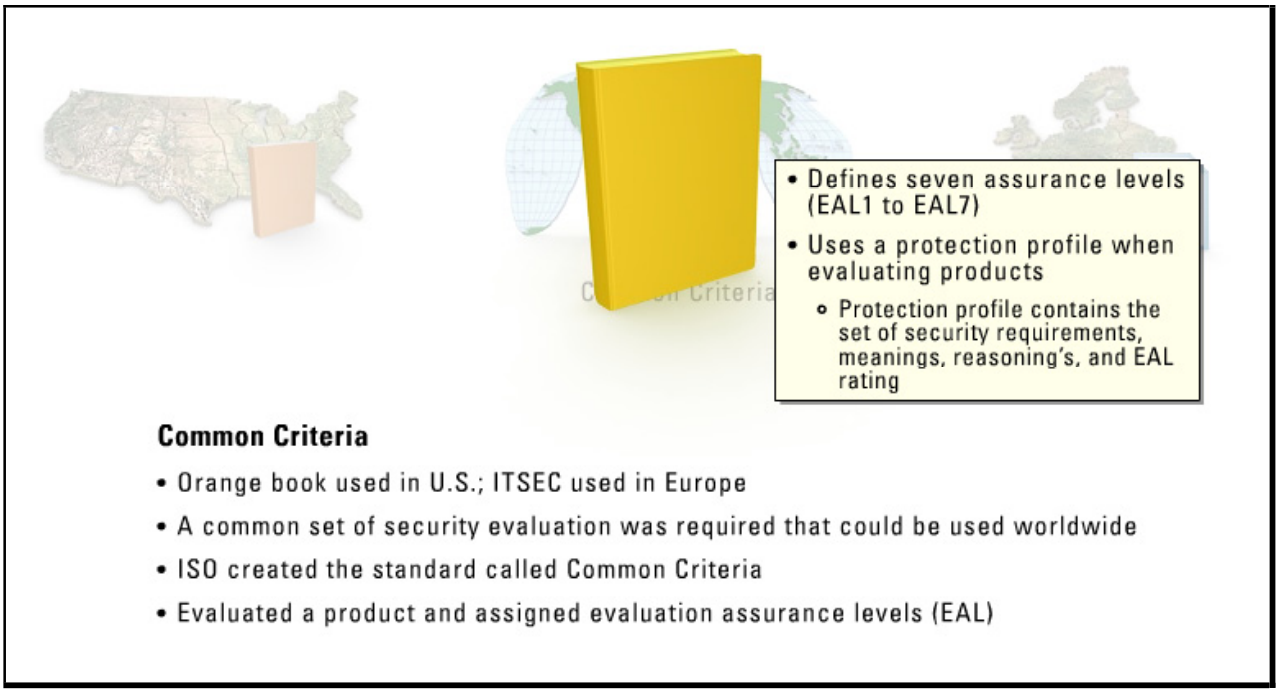

The Common Criteria standard is carried out on a product, which is then assigned an evaluation assurance level (EAL). The Common Criteria standard defines seven assurance levels, from EAL1, where functionality testing takes place, to EAL7, where thorough testing is performed and the system design is verified.

When evaluating products, the Common Criteria standard uses a mechanism called a **protection profile**. The protection profile contains the set of security requirements, their meaning and reasoning, and the corresponding EAL rating. The profile also contains the environmental assumptions, the objectives, and the functional and assurance level expectations.

The protection profile contains five sections:

- Descriptive elements
- Rationale
- Functional requirements
- Development assurance requirements
- Evaluation assurance requirements

Security Models

You can best design and analyze a secure system using a well-defined security model. A security model is a symbolic representation of a security policy; policymakers use it to map their desires to a set of rules that a computer system should follow. In essence, the security policy provides the abstract goals, and the security model provides the dos and don'ts necessary to fulfill those goals.

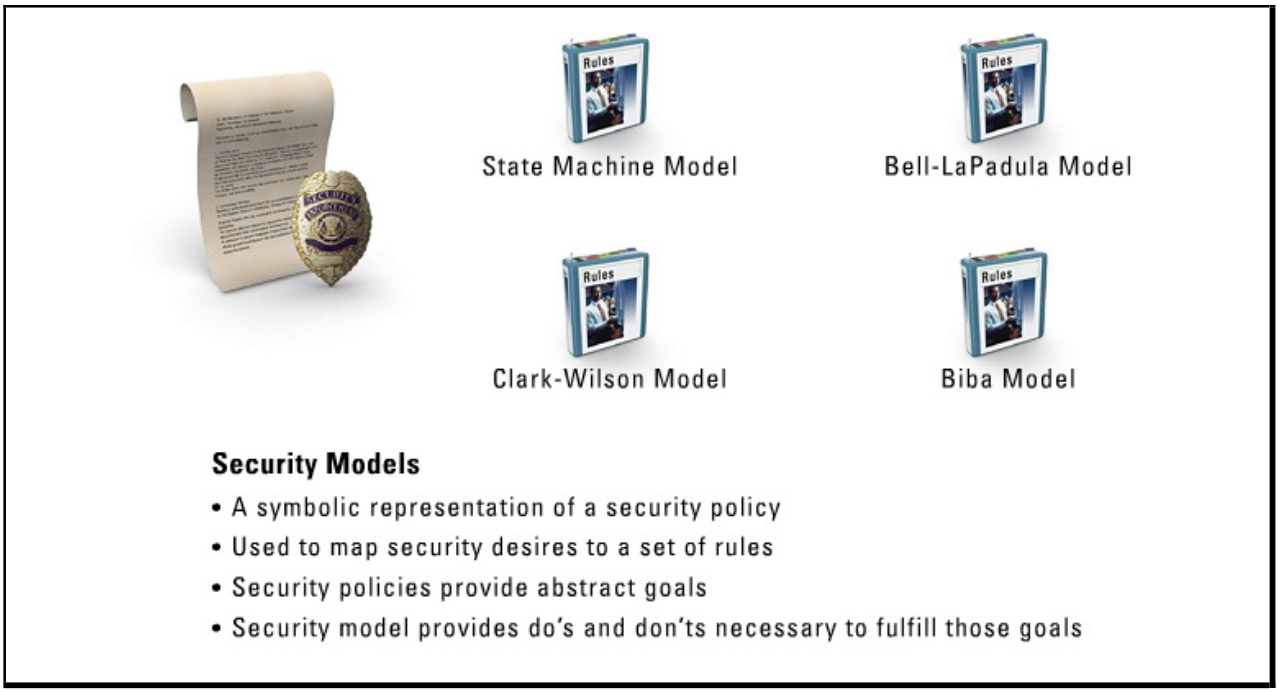

Some security models are created to enforce rules to protect the confidentiality of objects, while other models enforce rules to protect the integrity of objects. The following is a list of security models that have been developed to enforce security policies:

- State Machine Models
- Bell-LaPadula Model
- Biba Model
- Clark-Wilson Model

State Machine Models

A **state machine model** is one that verifies the security of a system by establishing and protecting the system state. The state is captured when a system gathers all current permissions and all current instances of subjects accessing objects. Maintaining the security of this state is performed when the system ensures that subjects can only access objects by means that are concurrent with the security policies.

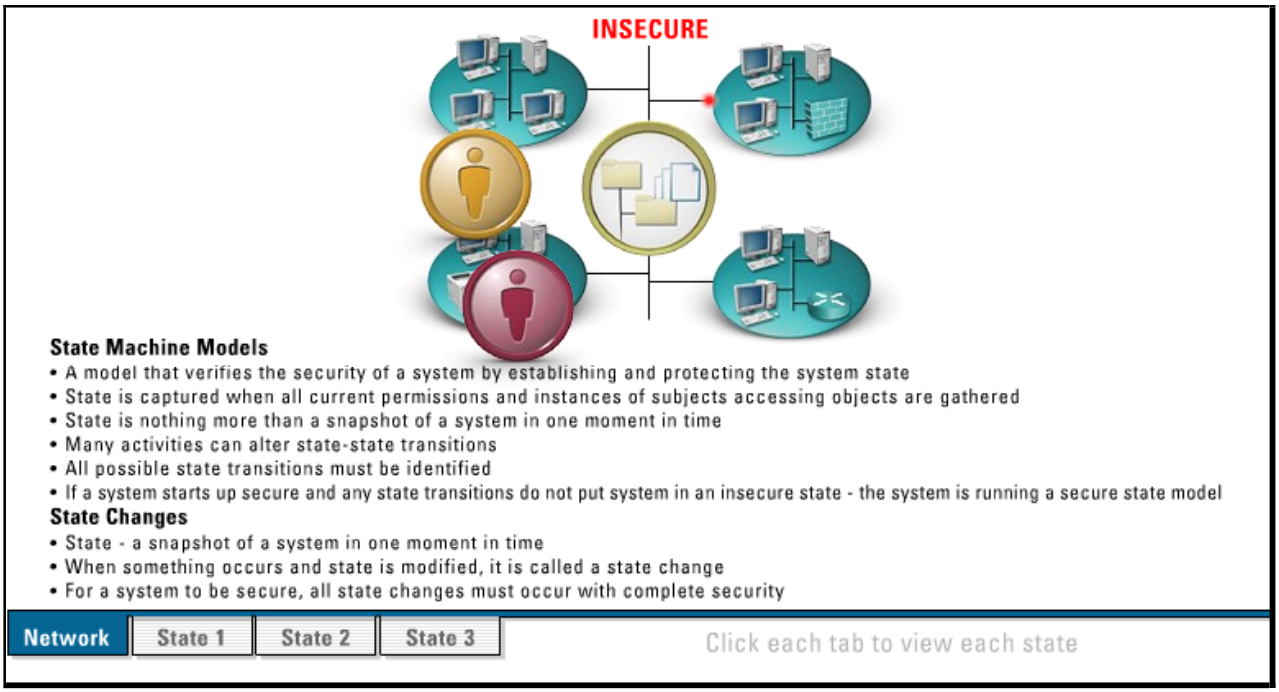

The state of a system is nothing more than a snapshot of the system in one moment of time. There are many activities occurring in the system that can alter this state, and when they occur we get state transitions. You must identify all possible state transitions that may occur in the system. If the developers of the state machine can validate the system starts up in a secure state, and any state transitions that occur do not put the system into an insecure state, then the system is said to be operating in a secure state model.

Note	It is said that a system implements a secure state model if each and every instance of its existence provides for a secure state. This means that the system must boot up in a secure state, execute commands in a secure state, execute transactions in a secure state, and allow subjects to access resources in a secure state.

Bell-LaPadula Model

The U.S. Military created the **Bell-LaPadula model** in the 1970s when they were concerned about the security of their time-sharing mainframe systems. Security of these systems and leakage of classified information were of utmost importance; hence the model's main goal was to prevent secret information from being accessed in an unauthorized manner.

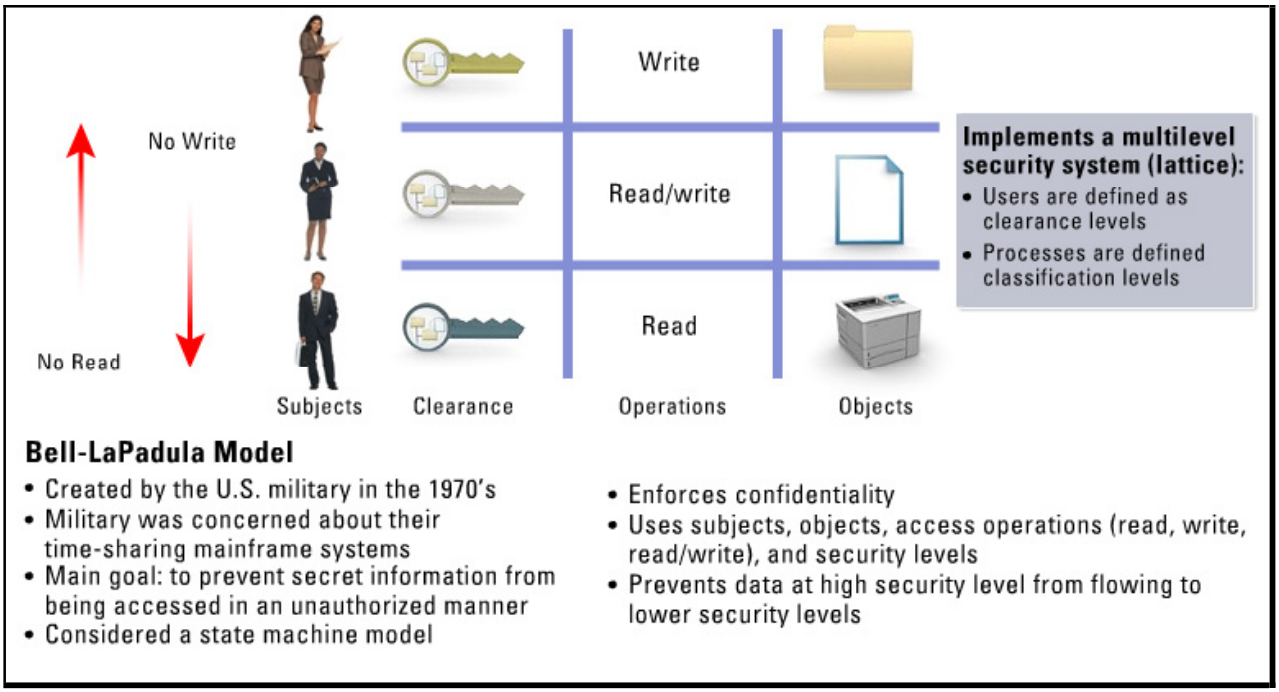

The Bell-LaPadula model implements a multilevel security system because users with different clearances use the U.S. Military's systems, and the systems process data with different classification levels. It should be noted that government installations use different data classification schemes than commercial organizations. The government uses top secret, secret, and unclassified schemes while commercial organizations classify their data as confidential, proprietary, and public. The level at which information is classified determines the handling procedure that should be used. These classifications form a lattice, which contains an upper and a lower bound of authorized access. For example, a user who has top secret clearance can access top secret, secret, and unclassified data. Top secret is the upper bound and unclassified is the lower bound.

Note MAC models are based on a lattice of security labels.

Bell-LaPadula is a state machine model that is used to enforce all confidentiality aspects of access control. In this model, the user's clearance is compared to the object's classification; if the clearance is higher or equal to the object's classification, then the user can access the object without violating the security policy. In order to enforce policy, the Bell-LaPadula model uses subjects, objects, access operations (read, write, and read/write), and security levels. Subjects and objects can reside at different security levels. If you properly implement and enforce this model, it has been mathematically proven to prevent data at a higher security level from flowing to a lower security level.

The Bell-LaPadula model is an information flow security model, which means that information cannot flow from an object of lesser or non-comparable classification.

There are two main rules used to enforce the Bell-LaPadula model:

- The simple security rule
- The *-property rule

The simple security rule states that a given security level cannot read data that resides at a higher security level. The *-property rule states that a user in a given security level cannot write information to a lower security level. The simple security rule is enforced as a “no read up” rule and the *-property rule is enforced as a “no write down” rule. Together, these rules indicate into which states the system can go.

It is important to understand that the Bell-LaPadula model was developed to make sure secrets stay secret; thus it provides confidentiality. The Bell-LaPadula model does not address integrity of the data that the system maintains.

The following is a list of cons to the Bell-LaPadula model:

- It addresses confidentiality only, not integrity.
- There is no management of access control (there are no mechanisms to modify access rights).
- It does not address the problem of covert channels.
- It does not address file sharing (used in more modern systems).

Biba Model

The **Biba model** was developed after the Bell-LaPadula model. This model was created to address the integrity of data. Integrity is threatened when users at lower security levels are able to write to objects at higher security levels and when users can read data at lower levels.

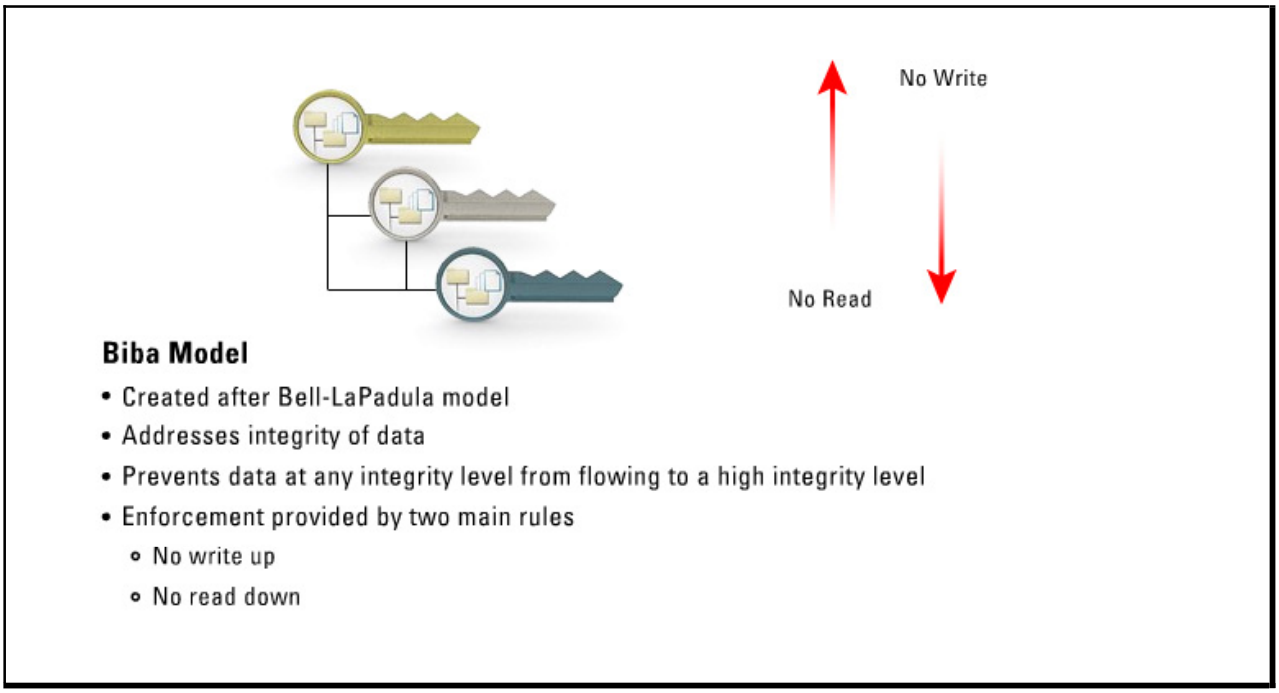

The Biba model prevents data at any integrity level from flowing to a higher integrity level. To enforce this model, the Biba model provides for two main rules:

- **No write up** - A user cannot write data to an object at a higher integrity level
- **No read down** - A user cannot read data from a lower integrity level

Clark-Wilson Model

The **Clark-Wilson model** is also used to provide for integrity of systems, but does so using an entirely different approach. It focuses on preventing authorized users from making unauthorized data modifications, committing fraud, and producing errors within commercial applications.

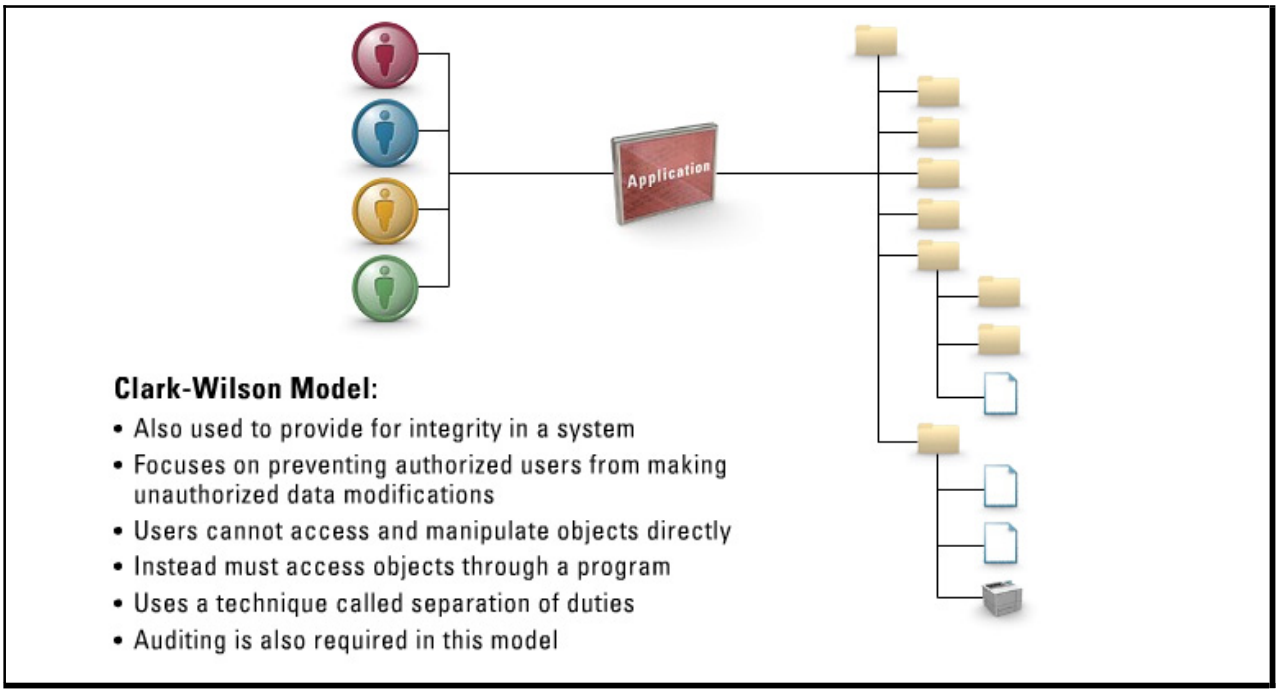

In the Clark-Wilson model, users cannot access and manipulate objects directly, but instead must access the object through a program. This process provides another layer of protection between the user and the object and further restricts the type of actions that can take place on that object, which further protects the object's integrity.

The Clark-Wilson model uses a technique called separation of duties; this technique divides an operation into different parts and requires different users to perform each part. You can use this process to prevent authorized users from making unauthorized data modifications, which again protects the integrity of the data. Auditing is also required in this model to track the information coming in from the outside of the system.

Summary

The key points discussed in this lesson are:

- You can use a security evaluation to examine the security-relevant parts of your system.
- The technical evaluation of the security components and their compliance for the purpose of accreditation is called certification. The formal acceptance of the adequacy of a system's overall security by management is called accreditation.
- An open system is an architecture that has published specifications, which enables third-party vendors to develop add-on components and devices. A closed system is an architecture that does not follow industry standards.
- You use active protection mechanisms to prevent access to an object if, according to the protection mechanism, the access is not authorized. Passive protection mechanisms, on the other hand, are those that prevent or detect unauthorized use of the information associated with an object, even if access to the object itself is not prevented.
- The operating system is the only line of defense between users, programs, applications, utilities, and the low level functions that are the core security features of the system.
- The TCB is defined as the total combination of protection mechanisms within a computer system.
- A Reference Monitor is an abstract machine, which mediates all the access subjects have to objects in order to ensure that the subjects have the necessary access rights and to protect the objects from unauthorized access and destructive modification. The Security Kernel is the hardware, firmware, and software elements of a TCB implementing the Reference Monitor concept.
- MACs are considered a policy-based control as an information system that provides MACs must assign sensitivity labels to all system subjects and all system objects.
- DACs are policies that restrict access to files and other system resources based on the identity and assignment of the user and/or the groups to which the user belongs.
- IPSec architecture is usually only necessary on highly critical systems where the protection of data is paramount.
- TCSEC are used to evaluate products to assess if they contain the security properties they claim to have and evaluate if the products are appropriate for specific applications or functions. The evaluation criteria were published in a book with an orange cover, which is why it is usually called the Orange Book.
- Other books have been written to extend the coverage of the Orange Book into additional areas of security. These books are collectively called the Rainbow Series because each book has a different colored cover.
- The TNI, also called the Red Book, is used to extend the Orange Book to networks.
- European countries created the ITSEC in an attempt to establish a single standard for evaluating security attributes of computer systems.
- Because the Orange Book was used in the United States and the ITSEC was used in Europe, any attempt to provide consistent security across borders was a management nightmare. A new common set of security evaluation criteria was required that could be used worldwide. The IS created a standard called the Common Criteria.
- A security model is a symbolic representation of a security policy; policymakers use it to map their desires to a set of rules that a computer system should follow.

- A state machine model is one that verifies the security of a system by establishing and protecting the system state.
- The Bell-LaPadula model implements a multilevel security system because users with different clearances use the U.S. Military's systems, and the systems process data with different classification levels.
- The Biba model was developed after the Bell-LaPadula model. This model was created to address the integrity of data begin threatened when users at lower security levels are able to write to objects at higher security levels and when users can read data at lower levels.
- The Clark-Wilson model is also used to provide for integrity of systems, but does so using an entirely different approach. It focuses on preventing authorized users from making unauthorized data modifications, committing fraud, and producing errors within commercial applications.

Common Flaws and Security Issues with System Architectures and Designs

Overview

Attackers pour through thousands of lines of code in applications in an attempt to subvert some code for their own use. This lesson will discuss common flaws and security issues that attackers have exploited.

Importance

Understanding how attackers attempt to infiltrate a network by subverting code is vital to the understanding of attacks as well as how to mitigate them.

Objectives

Upon completing this lesson, you will be able to:

- Define the two types of covert channels
- Describe the security flaw regarding memory
- Describe the logistics of input and parameter checking
- Define backdoors and maintenance hooks

Outline

The lesson contains these topics:

- Covert Channels
- Memory
- Input and Parameter Checking
- Backdoors and Maintenance Hooks

Covert Channels

A **covert channel** is another term for a confinement problem, which is an information flow issue. A covert channel is a way for an entity to receive information in an unauthorized manner. It is an information flow that is not controlled by a security mechanism.

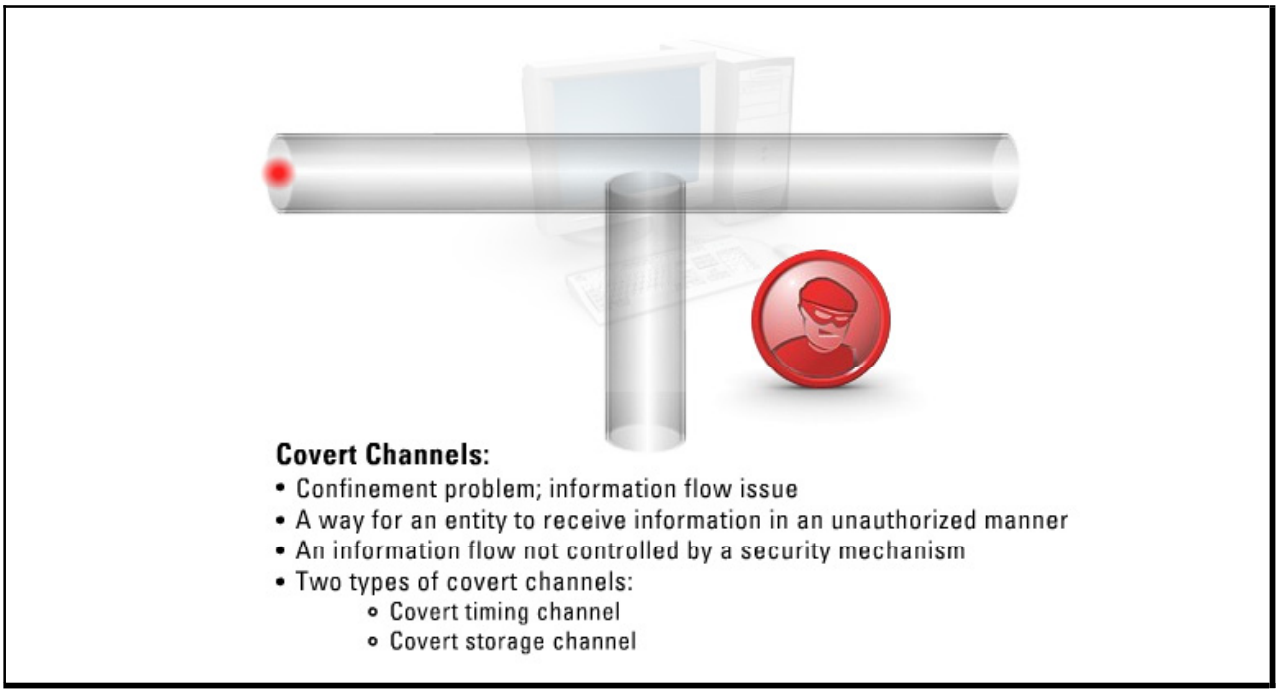

There are two types of covert channels:

- **Covert Timing Channel** - One process relays information to another by modulating its use of system resources
  - Accessing the hard drive
  - Using excessive CPU cycles
  - Head placement on hard drive track
- **Covert Storage Channel** - One process writes data to a storage location and another process directly, or indirectly, reads it

Memory

When memory is allocated to a process it is eventually deallocated when the service completes its task, and may be reallocated to yet another process and so on. It is quite possible and very probable that residual information may remain when a section of memory is reassigned to a new process after a previous process is finished with it. This situation can lead to a security violation.

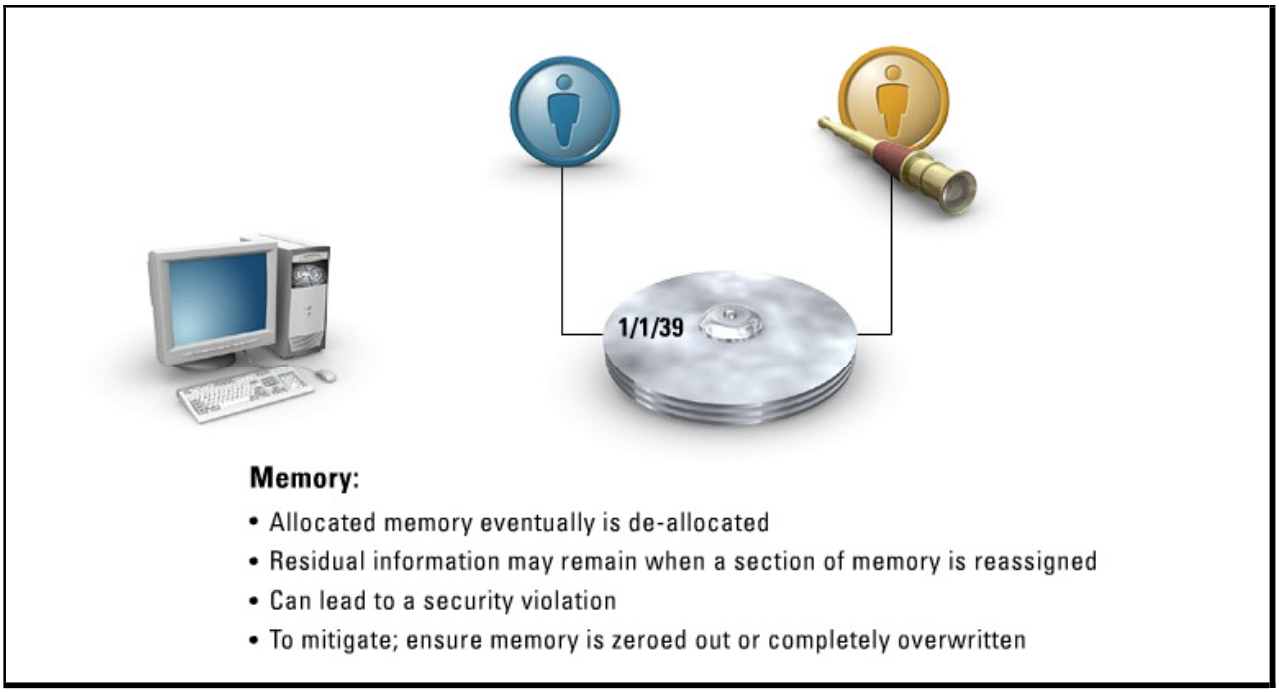

To mitigate this type of problem, the operating system must ensure that memory is zeroed out or completely overwritten before it can be accessed by any new process. Thus, there is no residual information in memory that is carried over from one process to another.

Input and Parameter Checking

When data is inputted into a system, there is a possibility that the data is outside the value that is "proper" for the parameter in question.

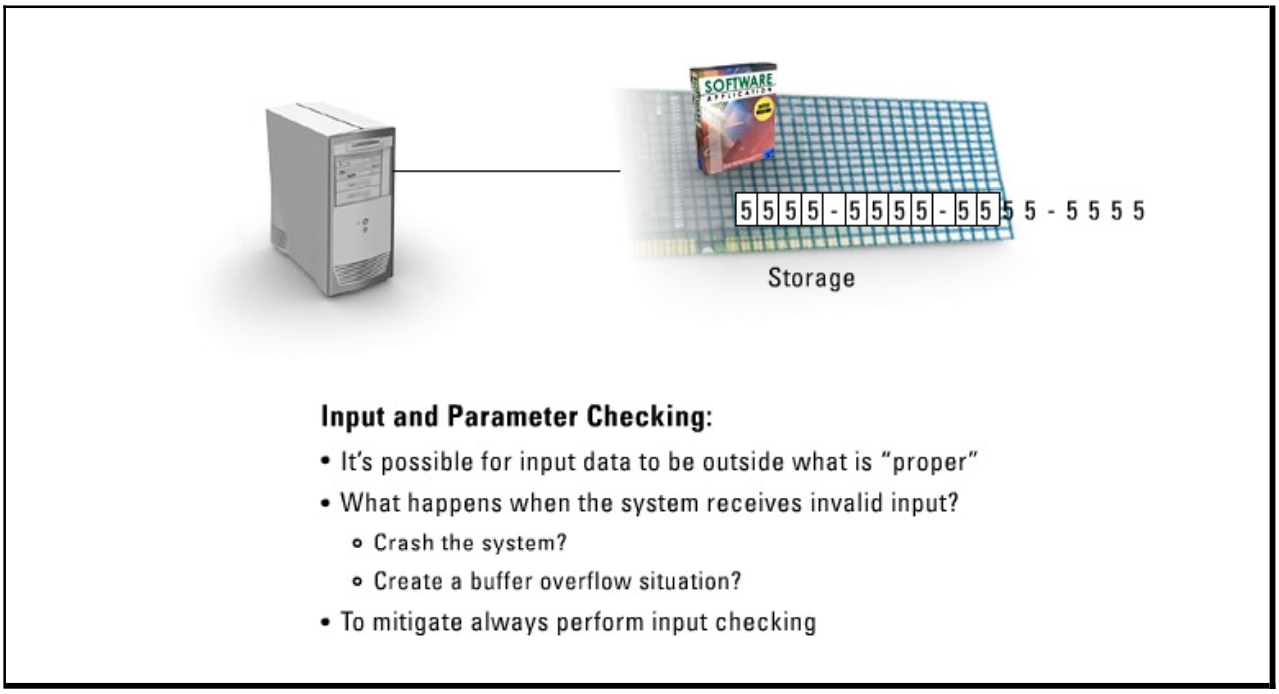

For example, if a field to be populated is one requiring the entry of a phone number, the system would expect to receive something like (602)555-1234. But, what happens when a value of 555-5555-5555-5555-5555-5555-5555 or a value of 10,000 bytes of code is inputted into the system. How will the system respond to this obviously invalid input? If the developer has implemented proper input and parameter checking, the system will reject the value or might not even allow more than 10 numeric-only characters. If the developer has not implemented proper input validation techniques, then it is possible for anything to happen. The system might crash (DoS), or worse a buffer-overflow condition might exist that may allow a savvy attacker to gain complete system control.

Backdoors and Maintenance Hooks

This topic discusses backdoors and maintenance hooks.

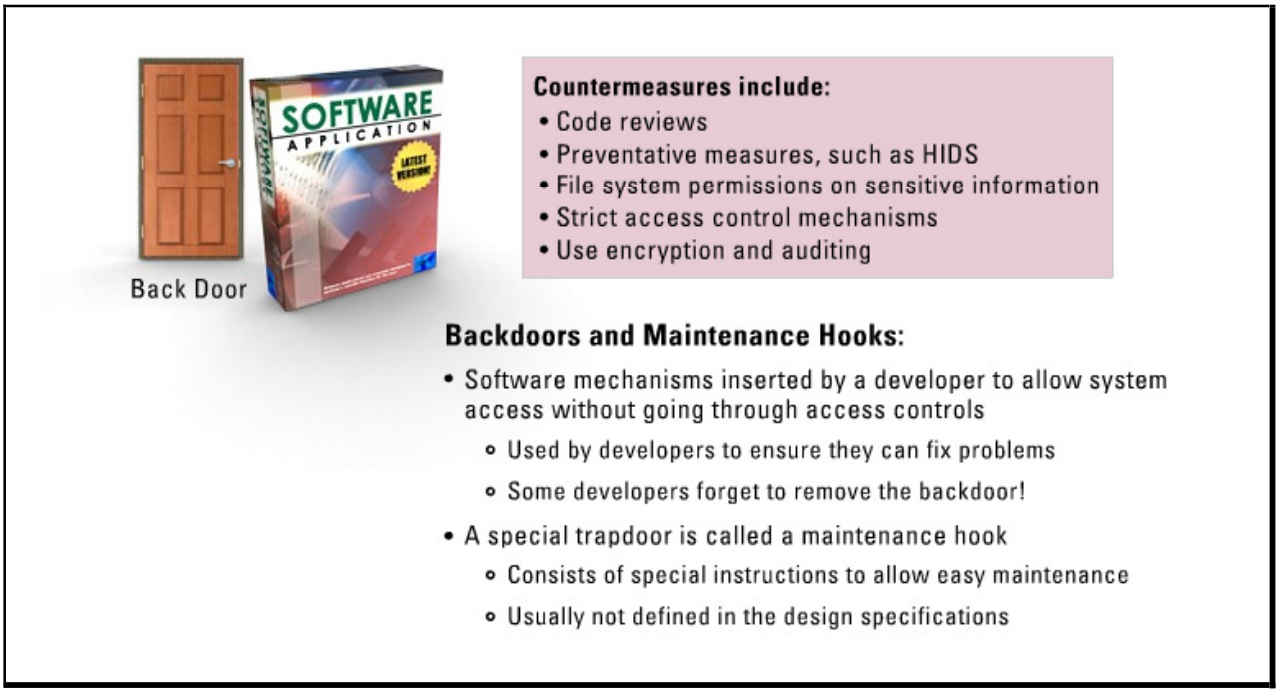

Backdoors or **trapdoors** are developer inserted software mechanisms that enable system access without going through the access control mechanism of the software. Many developers install this code to be sure they can fix a problem when the access control section of the program is damaged or fails.

Note	Security conscious developers require a password to activate the trapdoor.

Unfortunately, some developers forget to remove these backdoors when the program enters production mode. Now, anyone able to discover the trapdoor can easily obtain unauthorized access.

A special type of trapdoor is called a **maintenance hook**, and it consists of special instructions in software to allow easy maintenance and additional future development. Maintenance hooks are usually not defined in the design specifications and frequently allow system entry at unusual points, often without the usual security checks. For this reason, they can become a security risk if not removed before the system enters production.

Countermeasures to backdoors and maintenance hooks include:

- Code reviews and unit and integration testing should always be looking out for backdoors.
- Put preventative measures against backdoors in place, such as a host intrusion detection system.
- Use file system permissions to protect configuration files and sensitive information from being modified.
- Use strict access control mechanisms.
- Conduct file system encryption and auditing.

Summary

The key points discussed in this lesson are:

- A covert channel is a way for an entity to receive information in an unauthorized manner. It is an information flow that is not controlled by a security mechanism.
- The operating system must ensure that memory is zeroed out or completely overwritten before it can be accessed by any new process. Thus, there is no residual information in memory that is carried over from one process to another.
- When data is inputted into a system, there is a possibility that the data are outside the value that is "proper" for the parameter in question. Developers must implement proper input and parameter checking.
- Backdoors or trapdoors are developer inserted software mechanisms that enable system access without going through the access control mechanism of the software. Many developers install this code to be sure they can fix a problem when the access control section of the program is damaged or fails. A special type of trapdoor is called a maintenance hook, and it consists of special instructions in software to allow easy maintenance and additional future development.

Timing Attacks

Overview

A timing attack takes advantage of the way a system processes requests or performs various tasks. Timing attacks, also called asynchronous attacks, deal with how the system uses a timed sequence of events to complete a task. This lesson will discuss the various types of timing attacks and how you can mitigate them.

Importance

It is important that the information security professional understand the nature of timing attacks, how they exploit a system call, and what he or she can do to prevent them.

Objectives

Upon completing this lesson, you will be able to:

- Define timing attacks
- Identify ways to mitigate time-of-check versus time-of-use attacks
- Define state changes

Outline

The lesson contains these topics:

- Timing Issues
- Time-of-Check Versus Time-of-Use Attacks
- State Changes

Timing Issues

It is quite possible for an attacker to take advantage of the way a system processes requests and certain platform tasks. These types of attacks are called **timing** or **asynchronous attacks** and they deal with the timing difference of the sequence of steps a system uses to complete a task.

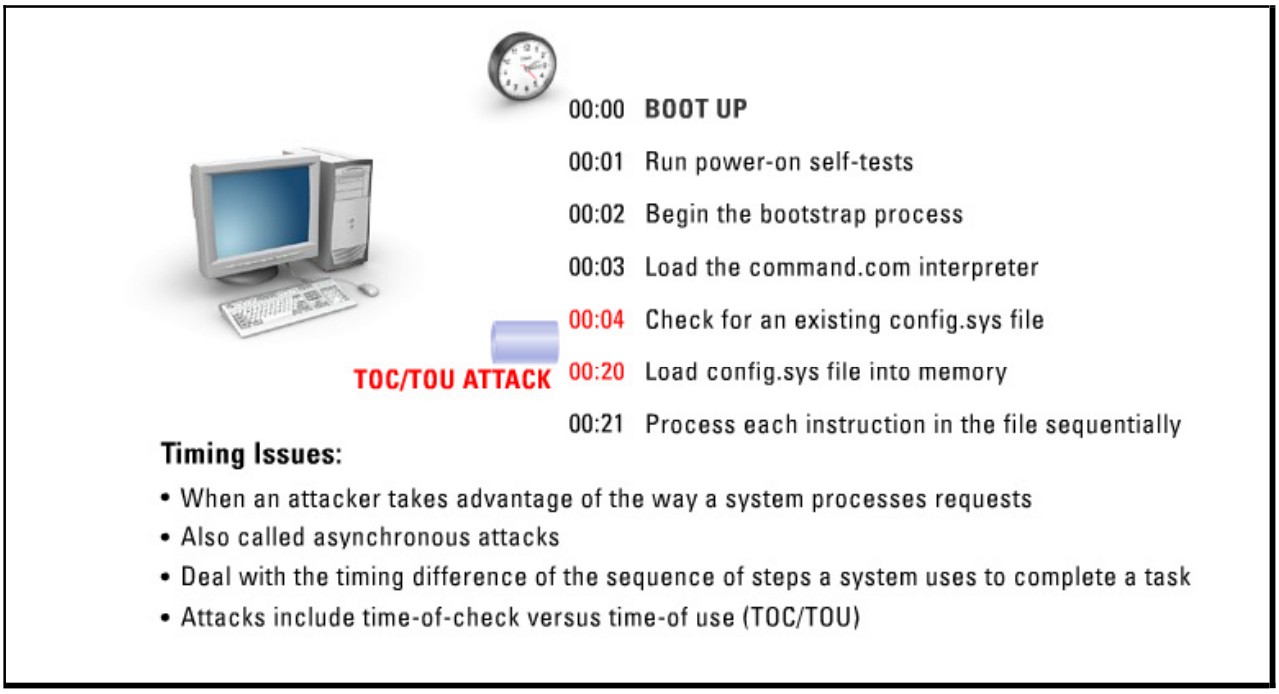

Take an example of a Windows platform that uses a config.sys file. When a system first boots up, it goes through its power-on self-tests, begins the bootstrap process, loads the command.com interpreter, and then checks to see if there is an existing config.sys file. If there is not a config.sys file, the system will continue to boot. If there is a config.sys file, the system will load it into memory and begin to process each instruction in the file sequentially. There is a timing difference from when the system checks to see if a config.sys file exists and actually accesses and opens the file.

A time-of-check versus time-of use (TOC/TOU) attack could conceivably replace the config.sys file with a different config.sys file that compromises the system before the system even loads its operating system.

Time-of-Check Versus Time-of-Use Attacks

Time-of-check versus time-of-use(TOC/TOU) attacks are called race conditions and occur when an attacker attempts to gain privilege to a system by "racing" the system to the resource it is attempting to access. The types of programming flaws that allow for race conditions occur when the system (or application) splits up the operations of verifying credentials and providing access to a resource.

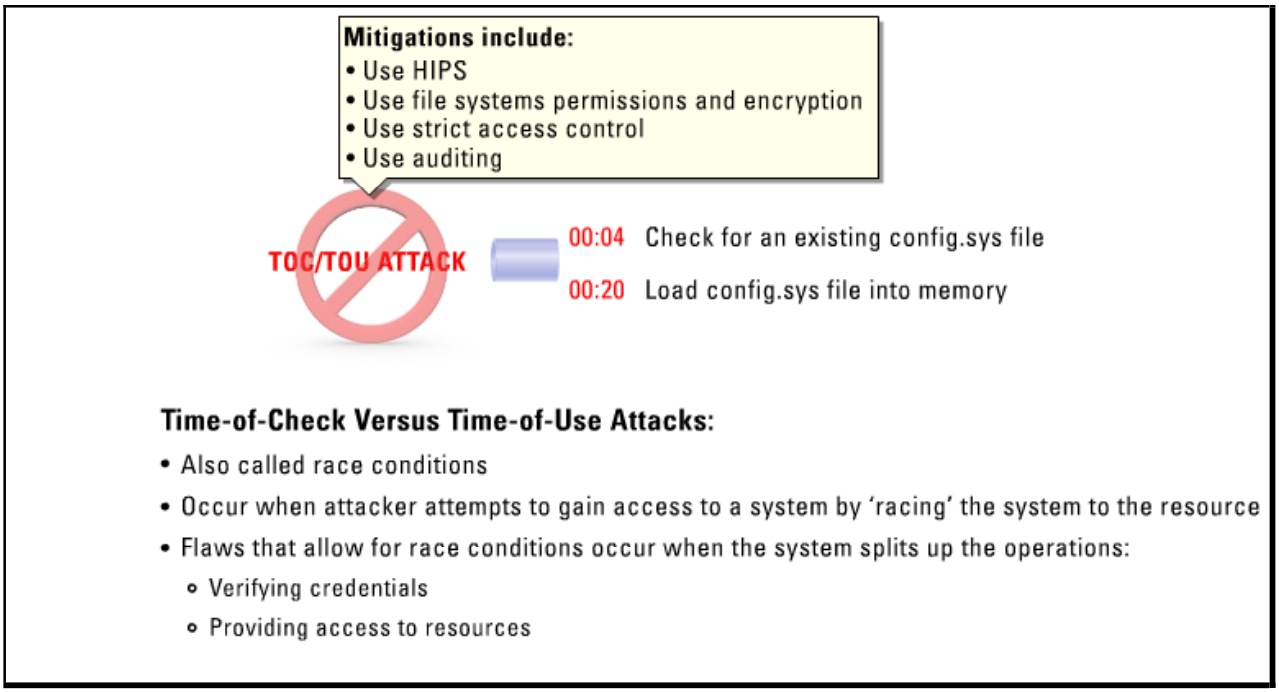

You can use the following list of items to protect a system against race conditions:

- Use host intrusion protection systems to watch for and stop these types of attacks.
- Set up file system permissions and encryption to protect sensitive files.
- Use strict access control to stop the attacker from accessing the system in the first place.
- Use auditing to identify patterns or indications of TOC/TOU attacks.

State Changes

A state when discussed in a computing environment is a snapshot of a system in one moment of time. You can describe all applications, programming variables, system variables, computing stacks, and memory segments that are in use at a single point in time as the state of that system. If something occurs and the state of the system is modified, it is called a **state change**.

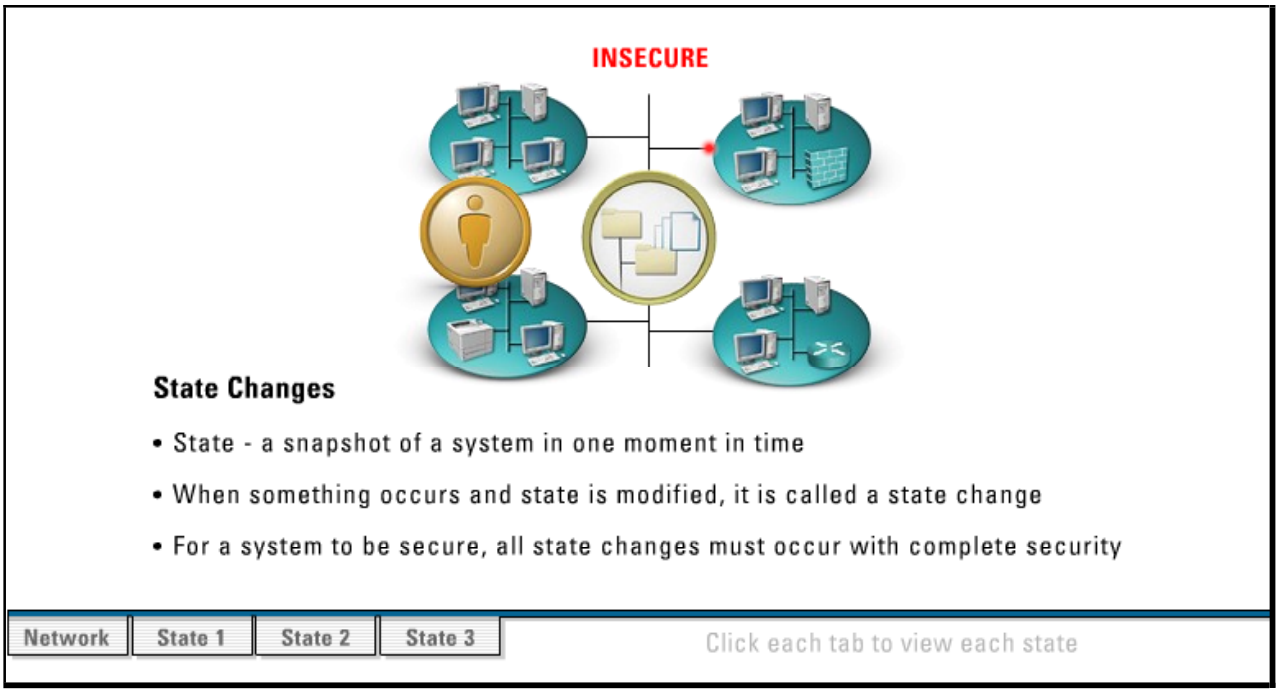

For a system to be secure, any and all state changes must occur with complete security. For example, in a multiprocessing system, the CPU is shared among many processes, each occurring individually in its own virtual machine. The flags and counters each virtual machine uses to keep the system variable locations known must constantly change when the CPU goes from virtual machine to virtual machine. If an attacker can subvert these flags and counters to point to another portion of the system, another process can be executed.

Summary

The key points discussed in this lesson are:

- It is quite possible for an attacker to take advantage of the way a system processes requests and certain platform tasks. These types of attacks are called timing or asynchronous attacks and they deal with the timing difference of the sequence of steps a system uses to complete a task.
- TOC/TOU attacks are called race conditions and occur when an attacker attempts to gain privilege to a system by "racing" the system to the resource it is attempting to access.
- If something occurs and the state of the system is modified, it is called a state change. For a system to be secure, any and all state changes must occur with complete security.

Operations Security

Overview

The operation of a network, system, or enterprise is anything that is required or takes place to keep the entity functioning. Operations security is the security overlay applied to keep the system secure.

Objectives

Upon completing this module, you will be able to:

- Identify various activities required to secure the operations of an entity
- Identify audits that help determine how susceptible a network or system is to attack
- Describe how violations occur and how to analyze them when they do occur
- Describe methods used to identify how an attack occurred and the attacker involved
- Identify tools used to monitor systems
- Identify what resources must be protected and how to protect them
- Identify the protocols used in e-mail systems, their insecurities, and how to mitigate them
- Identify the underlying protocols of the World Wide Web and how security plays a role in the Web
- Describe file transfer technologies and their security safeguards
- Identify the methods attackers use to obtain information about a system
- Describe the role of separation of duties and responsibilities in a company

Outline

The module contains these lessons:

- Operations Security Overview
- Security Audits
- Violation Analysis
- Auditing
- Monitoring

- Resource Protection
- E-Mail Security
- The Web
- File Transfer
- Anatomy of an Attack
- Separation of Duties and Responsibilities

Operations Security Overview

Overview

This lesson will discuss the various activities necessary to keep the operations of an entity up and running in a secure manner.

Importance

It is vital that the information security specialist understand security controls, auditing, violation analysis, and other methods of maintaining security in the operations of an enterprise.

Objectives

Upon completing this lesson, you will be able to:

- Identify the three critical requirements for information security controls
- Identify common attacks that can take place against an entity's operations
- Identify mechanisms that you can use to provide operations security
- Identify operations security tools
- Identify security guidelines for facilities

Outline

The lesson contains these topics:

- Information Security Controls
- Common Attacks
- Mechanisms
- Tools
- Facilities

Information Security Controls

This topic discusses information security controls.

Information security controls are in place to protect hardware, software, and media resources from:

- Threats in an operating environment
- Internal or external intruders
- Operators who are inappropriately accessing resources

There are three critical requirements for information security controls:

- Resource protection
- Privileged-entry controls
- Hardware controls

Common Attacks

This topic discusses common attacks that can take place against an entity's operations.

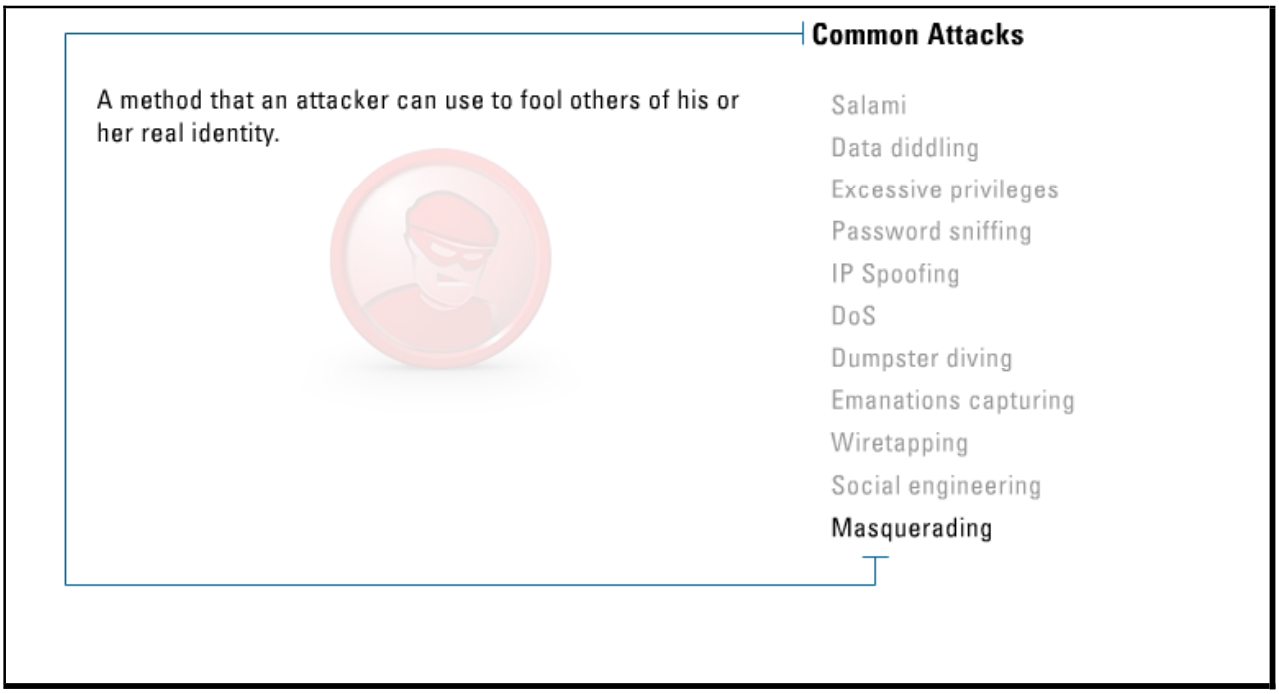

Common attacks include:

- **Salami** - Involves subtracting a small amount of funds from an account with the hope that such an insignificant amount would be unnoticed
- **Data Diddling** - Refers to the alteration of existing data; many times this modification happens before the data are entered into an application or as soon as the data complete processing and are outputted from an application
- **Excessive Privileges** - Occurs when a user has more computer rights, permissions, or privileges than what are required for the tasks she needs to fulfill
- **Password Sniffing** - Sniffing network traffic in the hopes of capturing passwords being sent between computers
- **IP Spoofing** - Manually change the Internet Protocol (IP) address within a packet to point to another address
- **Denial of Service (DoS)** - Denying others the service that the victim system usually provides
- **Dumpster Diving** - Refers to someone rummaging through another person's garbage for discarded documents, information, and other precious items that could then be used against that person or his or her company
- **Emanations Capturing** - Eavesdropping of the electrical waves emitted by an electrical device
- **Wiretapping** - Eavesdropping of communication signals
- **Social Engineering** - The art of tricking people and using the information they unknowingly supply in a malicious way
- **Masquerading** - A method that an attacker can use to fool others of his or her real identity

Mechanisms

This topic discusses mechanisms that you can use to provide operations security.

Mechanisms to provide operations security include:

- **Preventative Controls** - Are designed to lower the amount and impact of unintentional errors that are entering the system and to prevent unauthorized intruders from internally or externally accessing the system
- **Detective Controls** - Are used to detect an error once it has occurred
- **Corrective Controls/Recovery Controls** - Are implemented to mitigate the impact of a loss event through data recovery procedures
- **Deterrent Controls/Directive Controls** - Are used to encourage compliance with external controls
- **Application Controls** - Software application controls used to minimize and detect the software's operational irregularities
- **Transaction Controls** - Are used to provide control over the various stages of a transaction; types of transaction controls are input, processing, output, change, and test

You can also use Orange Book controls to provide operations security/operational assurance:

- System architecture
- System integrity
- Covert channel analysis
- Trusted facility management
- Trusted recovery

Tools

This topic discusses tools you can use in operations security.

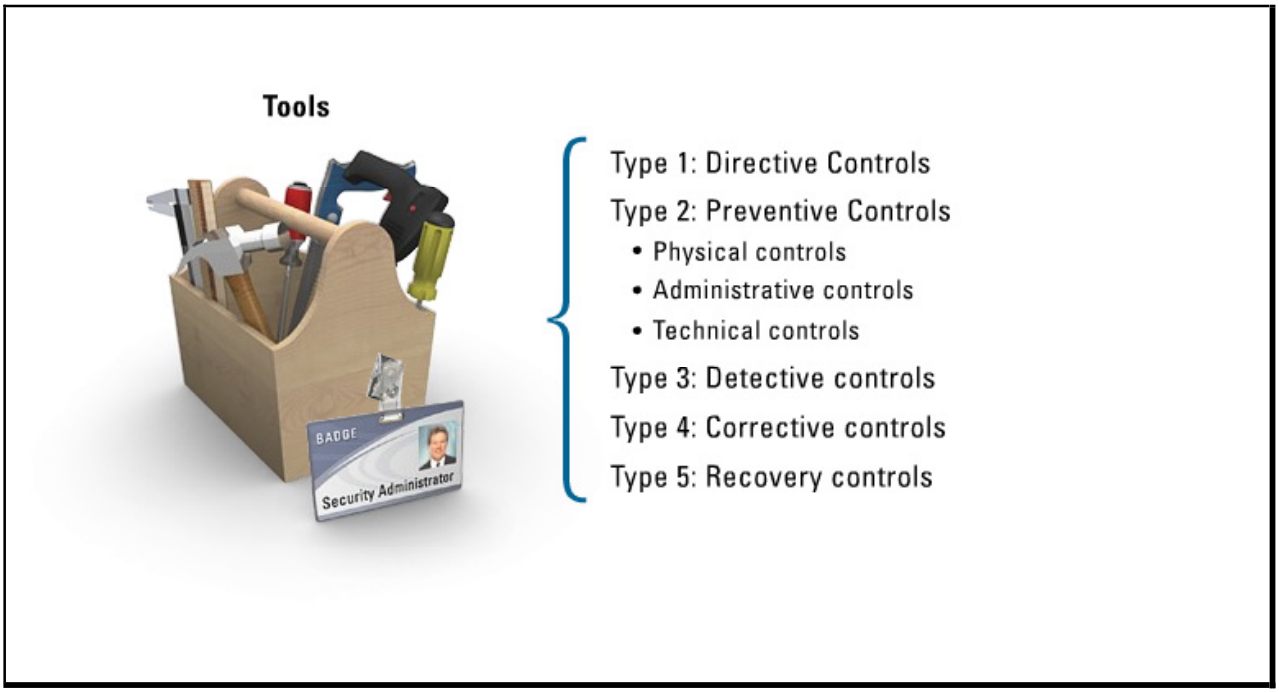

Tools:

- **Type 1: Directive Controls** - Intended to advise employees of the behavior expected of them; often called administrative controls
- **Type 2: Preventive Controls** - Intended to prevent the occurrence of undesirable activity
  - Physical controls
  - Administrative controls
  - Technical controls
- **Type 3: Detective Controls** - Involve the use of practices, processes, and tools that identify and react to security violations; examples include audit trails, intrusion detection system (IDS) software, logs, integrity checks, violation reports, etc.
- **Type 4: Corrective Controls** - Involve physical, administrative, and technical controls designed to react to the detection of an incident in order to reduce or eliminate the opportunity for the unwanted event to occur
- **Type 5: Recovery Controls** - Necessary to restore the system or operation to a normal operating state

Facilities

All data processing facilities and related equipment should be located within a structure so as to minimize any exposure to threats, such as water, fire, corrosive agents, and smoke. Other potential hazards can occur from neighboring areas, explosion or shock, and unobserved unauthorized access.

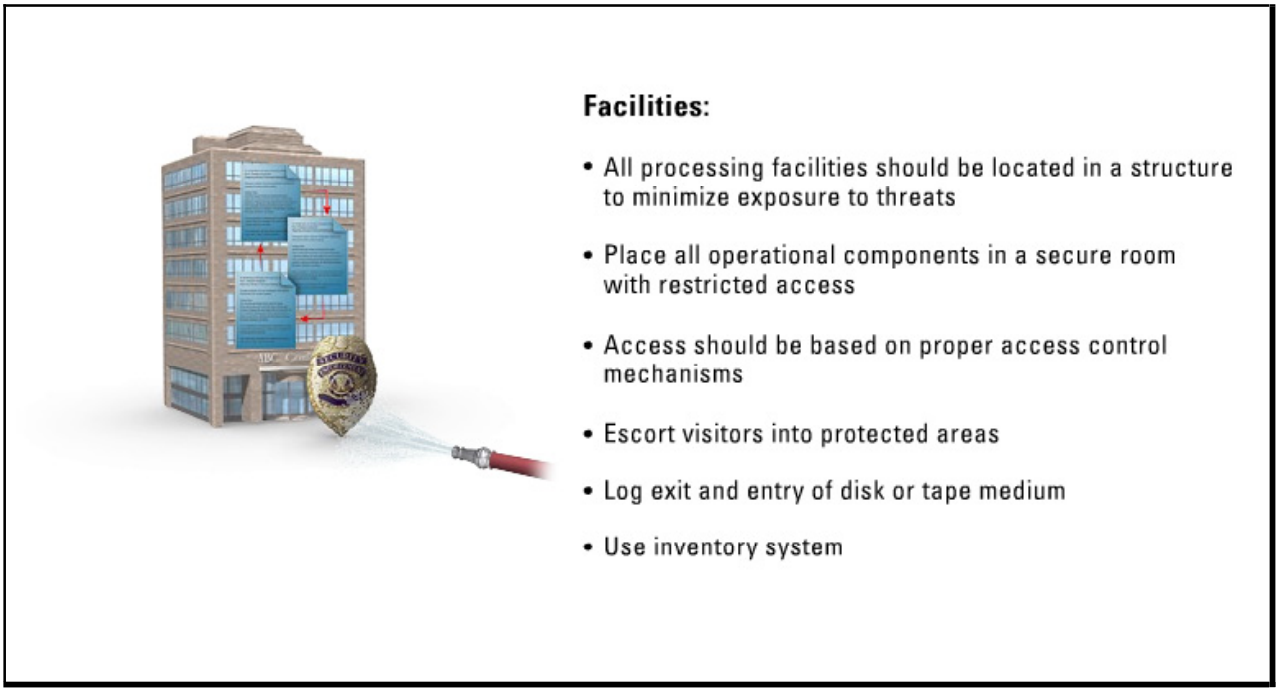

You should place all operational components in a secure room with restricted access. You must base access to the room on proper access control mechanisms (badge, ID card, security guard, etc.). To make sure visitors do not obtain access to sensitive equipment, knowledgeable operations personnel should escort them into these areas.

When any disk or tape medium is entering or leaving the restricted room, the item should be logged to ensure proper accountability. You should also use an inventory system that identifies what equipment is assigned to which person.

Summary

The key points discussed in this lesson are:

- There are three critical requirements for information security controls: resource protection, privileged-entry controls, and hardware controls.
- Common attacks include salami, data diddling, excessive privileges, password sniffing, IP spoofing, DoS, dumpster diving, emanations capturing, wiretapping, social engineering, and masquerading.
- Mechanisms to provide operations security include preventative controls, detective controls, corrective controls/recovery controls, deterrent controls/directive controls, application controls, and transaction controls.
- Tools you can use in operations security include directive controls, preventive controls, detective controls, corrective controls, and recovery controls.
- You should place all operational components in a secure room with restricted access. You must base access to the room on proper access control mechanisms.

Security Audits

Overview

You can use audits to gather relevant facts about systems, security procedures, events, or item counts. This lesson will discuss the audits you can use to determine how susceptible your network or system is to attack.

Importance

It is important for information security professionals to constantly monitor and audit their sensitive systems to determine their susceptibility to attack.

Objectives

Upon completing this lesson, you will be able to:

- List items you can evaluate with internal audits
- List items you can evaluate with external audits
- Describe the standard of due care and its related terms

Outline

The lesson contains these topics:

- Internal Audits
- External Audits
- Standard of Due Care

Internal Audits

A strong internal auditing program is the best way to identify and correct deficiencies before they result in regulatory actions against your company.

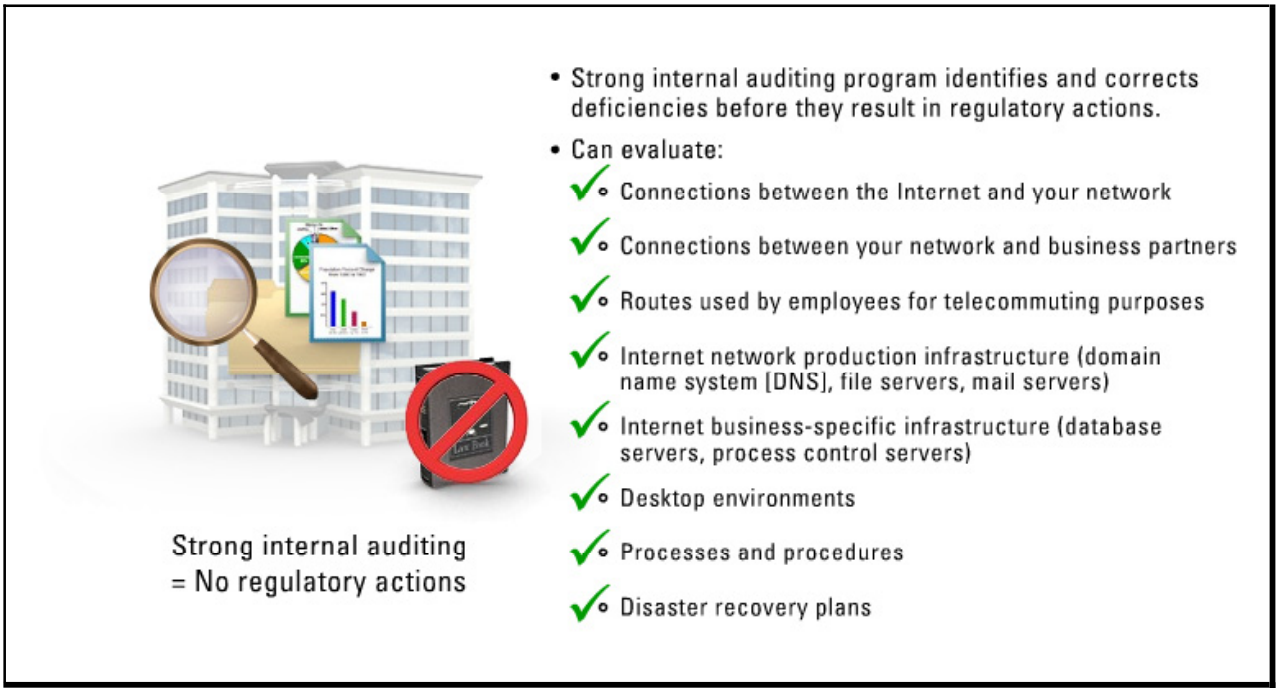

The configuration of your network is subject to change on a regular basis. Your network likely stores valuable organizational assets, which are often restricted to certain personnel. Databases secure human resources records, financial data, intellectual property, and the like. However, with inadequate security, it is trivial to extract valuable data from local or trusting networks. Different variables can affect your network security including new services, new applications, staffing changes, and new connections. These seemingly innocent alterations can accidentally open windows and backdoors. You must close these open invitations to attackers and opportunists.

Internal audits can evaluate:

- Connections between the Internet and your network
- Connections between your network and business partners
- Routes used by employees for telecommuting purposes
- Internet network production infrastructure (domain name system [DNS], file servers, mail servers)
- Internet business-specific infrastructure (database servers, process control servers)
- Desktop environments
- Processes and procedures
- Disaster recovery plans

External Audits

Internal auditing is one critical aspect of a security plan that can reduce the risk associated with new attack tools. However, many internal auditing projects, if they are being done at all, focus on high-level policy issues like weak passwords, directory and file permissions, and disaster-recovery procedures. Often, it is only the external audits that commonly test for the actual operating system and network service vulnerabilities being exploited by today's new hacking tools.

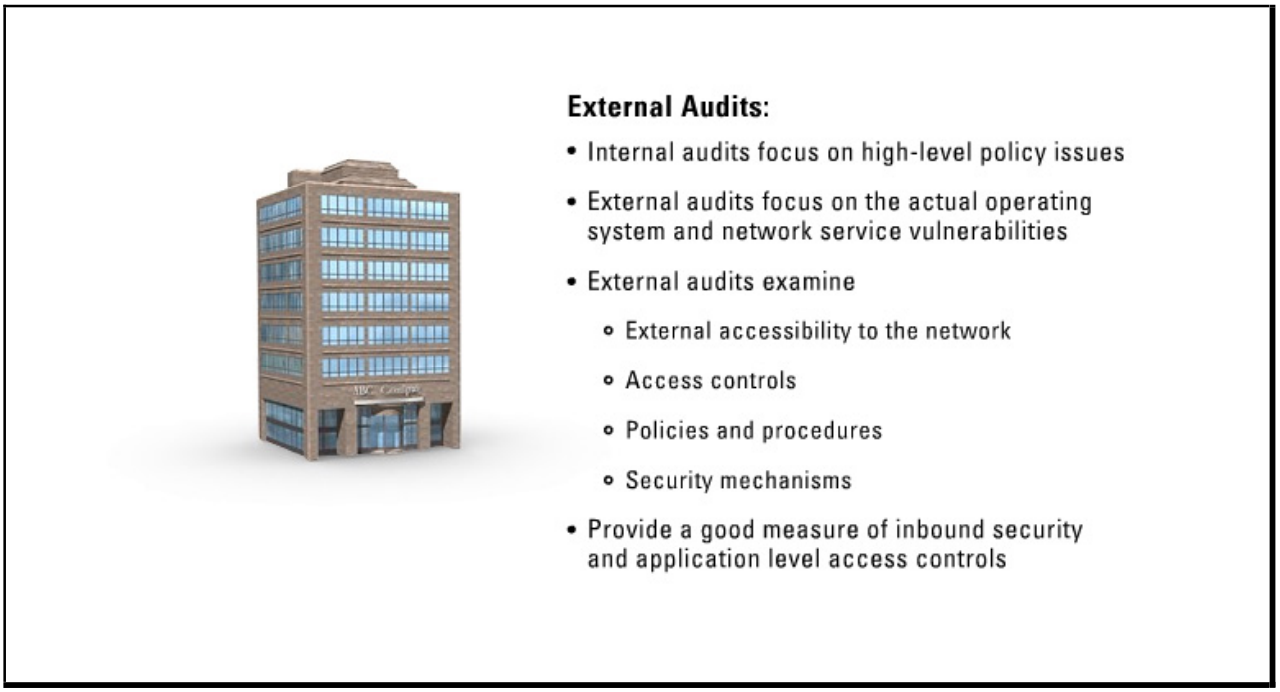

External audits examine external accessibility to the network, access controls, policies and procedures pertaining to those areas, and security mechanisms such as IDS and virus control. External audits provide a good measure of inbound security and application level access controls. External audits completely map the entire architecture of your computer network to determine what is available to potential hackers and then launch various intrusion tests pertinent to this mapped architecture.

However, the external audit does not provide the thorough, in-depth analysis that is required to understand all aspects of the system, data, and data flows. Therefore, you should regularly conduct internal audits. Internal audits assess the state of servers and workstations relative to known vulnerabilities.

Standard of Due Care

The steps that you take to show that your company has taken responsibility for its activities and has taken the necessary steps to help protect the company, its resources, and its employees are referred to as the **standard of due care**.

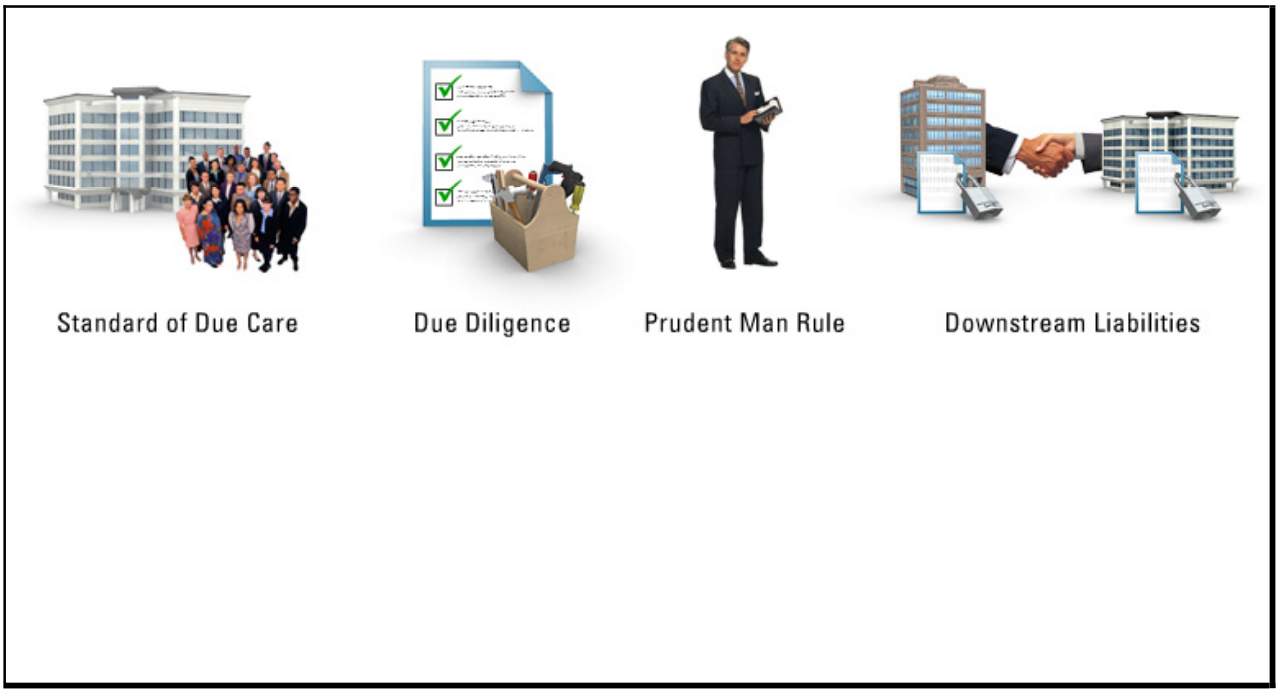

Due diligence includes the continual activities that you do to make sure the protection mechanisms are continually maintained and operational.

The **prudent man rule** states that a person is not liable for damages if the person performs duties that prudent people would exercise in similar circumstances.

The **downstream liabilities** standard states that when companies come together to work in an integrated manner, special care must be taken to ensure that each party promises to provide the necessary level of protection, liability, and responsibility needed, which should be clearly defined in the contract that each party signs.

The **legally recognized obligation** standard states that there is a standard of conduct expected of the company to protect others from unreasonable risks. If the company fails to conform to this standard, injury or damage could result. **Proximate causation** means that someone can prove that the damage that was caused was the company's fault.

Summary

The key points discussed in this lesson are:

- Internal audits assess the state of servers and workstations relative to known vulnerabilities.
- External audits examine external accessibility to the network, access controls, policies and procedures pertaining to those areas, and security mechanisms such as IDS and virus control.
- The steps that you take to show that your company has taken responsibility for its activities and has taken the necessary steps to help protect the company, its resources, and its employees are referred to as the standard of due care.

Violation Analysis

Overview

When a system or network has been violated it is important to gain as much information as possible about the attack. At a minimum, you should determine how the system was violated, what tactics were used in the violation, how information about the system was gathered, what information was stolen or destroyed, and, if possible, who attacked the system.

Importance

It is important that the information security professional understand how violations occur and how to analyze them when they do occur.

Objectives

Upon completing this lesson, you will be able to:

- Define violation analysis
- Identify potential hardware and software exposures
- Identify the effects of device address modification
- Identify the effects of an attacker rerouting output from your system to his or her system
- Identify the effects of an attacker obtaining supervisory access to equipment
- Identify the effects of an attacker bypassing system logs
- Identify the effects of an attacker shutting down a system
- Identify the effects of an attacker gaining physical access to a device
- Identify the effects of an attacker gaining access to a system's I/O functions
- Identify the effects of a network that is not properly secured
- Identify how an attacker can compromise the server boot-up sequence
- Identify the effects of an attacker highjacking a server's network address
- Identify the effects of an attacker stealing the password file
- Identify the effects of an attacker gaining access to the operating system or application user accounts
- Describe how you can use clipping levels as security measures

Outline

The lesson contains these topics:

- Violation Analysis Introduction
- Potential Hardware and Software Exposures
- Device Address Modification
- Rerouting Output
- Obtaining Supervisory Terminal Functions
- Bypassing System Logs
- System Shutdown/Downtime
- IPL from Tape, CD, or Floppy Disk
- I/O System Generation
- Network Vulnerability
- Server Boot-up Sequence
- Highjacking the Server's Network Address
- Stealing Password File/Table from the Server
- Gaining Access to OS or Application User Accounts
- Determining Clipping Levels

Violation Analysis Introduction

Violation analysis permits an organization to locate and understand specific trouble spots, both in security and usability.

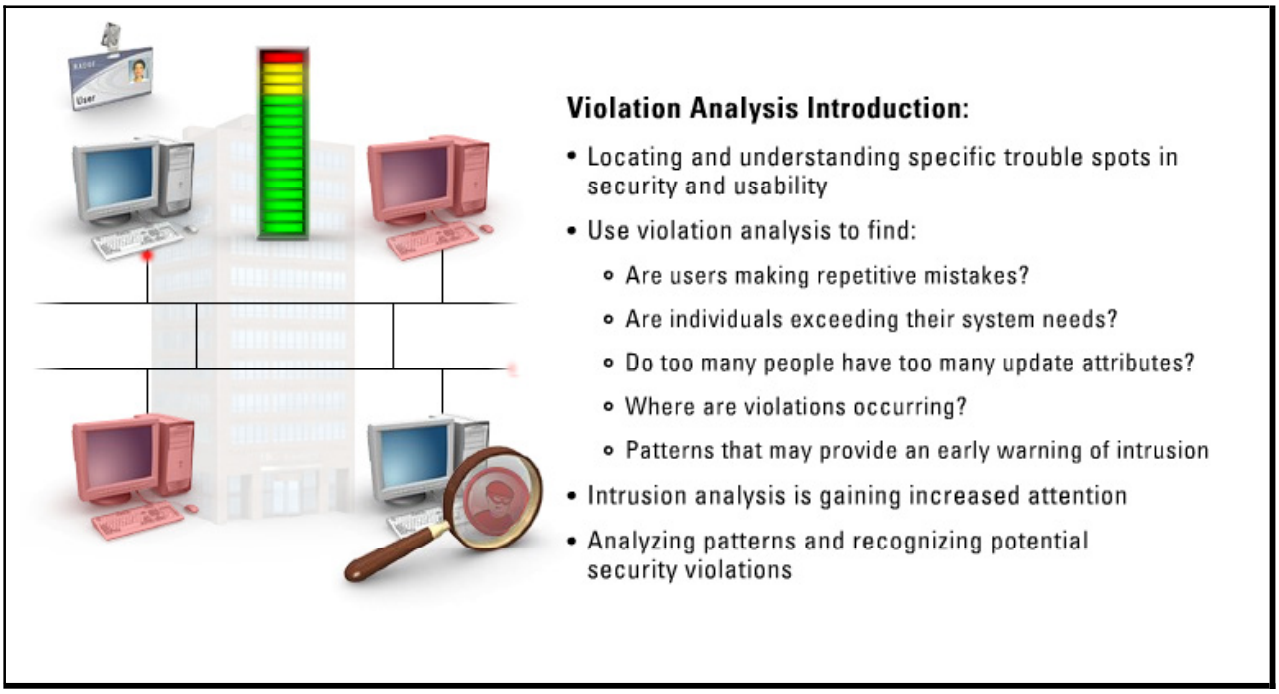

You can use violation analysis to find:

- The types of violations occurring
  - Are users making repetitive mistakes? This might be a sign of poor implementation or user training.
  - Are individuals exceeding their system needs? This might be an indication of weak control implementation.
  - Do too many people have too many update abilities? This might be a result of inadequate information security design.
- Where the violations are occurring, which might help identify program or design problems
- Patterns that can provide an early warning of serious intrusions (e.g., hackers or disgruntled employees)

A specialized form of violation examination, **intrusion analysis** (i.e., attempting to provide analysis of intrusion patterns), is gaining increased attention. As expert systems gain in popularity and ability, their use in analyzing patterns and recognizing potential security violations will grow. The need for such automated methods is based on the fact that intrusions continue to increase rapidly in quantity and intensity and are related directly to the increasing number of personal computers connected to various networks.

Note	The need for automated methods is not likely to diminish in the near future, at least not until laws surrounding computer intrusion are much more clearly defined and enforced.

Potential Hardware and Software Exposures

This topic discusses potential hardware and software exposures.

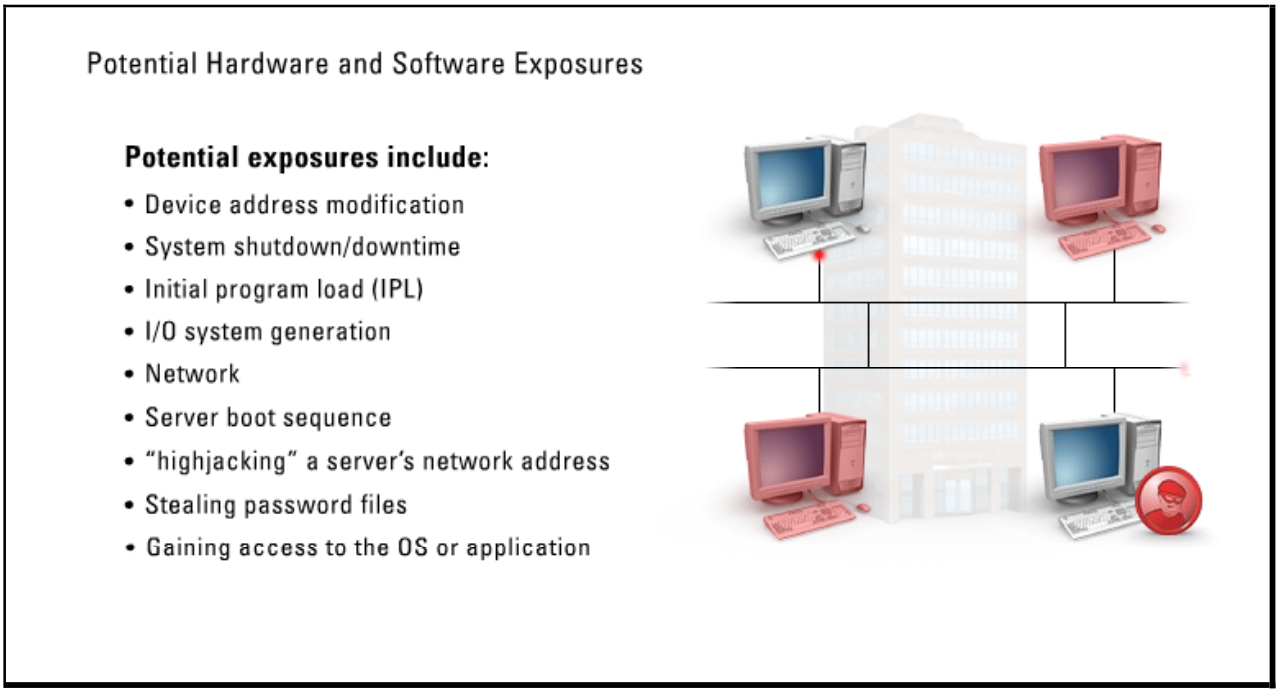

Potential exposures include:

- Device address modification, such as rerouting output, obtaining supervisory terminal access, or bypassing system logs
- System shutdown/downtime, such as shutdown handled from console/operations area
- IPL from tape, such as initial program load without security
- I/O system generation, including terminals/printers/disks and tape drives
- Network
- Server boot sequence from tape, CD, or floppy disk that can bypass O/S security
- "Highjacking" a server's network address in order to capture traffic to and from the server
- Stealing the password file from the server
- Gaining access to the OS or application user accounts

Device Address Modification

This topic discusses the effects of device address modification.

Device Address Modification:

- If an attacker can modify a system device address, they can cause many problems such as:
  - DoS attacks
  - Rerouting attacks
  - Man-in-the-middle attacks

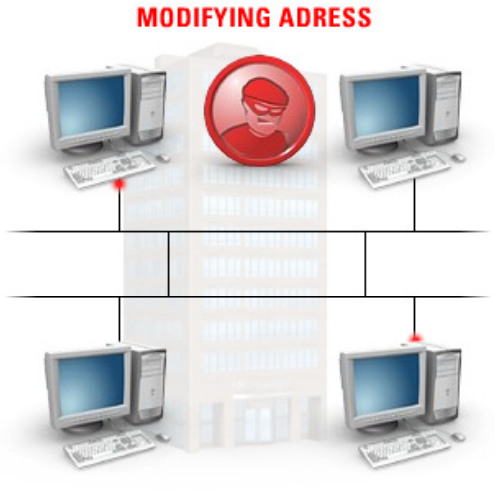

If an attacker can modify a system device address, the attacker can cause many types of problems, including denial of service (DoS) attacks as the end users can no longer reach their intended device, rerouting attacks where the attacker uses the system device address on his or her system, and man-in-the-middle attacks.

Rerouting Output

This topic discusses the effects of an attacker rerouting output from your system to his or her system.

Rerouting Output:

- If an attacker can reroute output from a system to his system, they can:
  - Steal passwords
  - Obtain sensitive information
  - Generally wreak havoc on the system

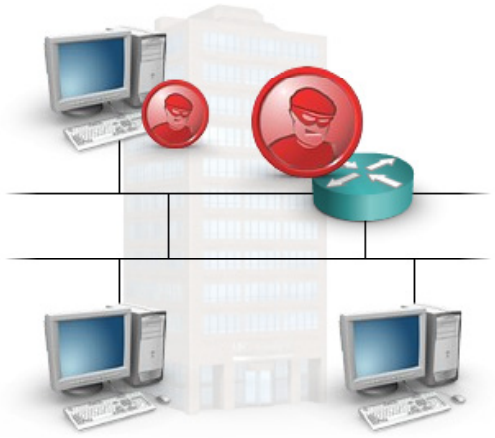

If an attacker can reroute output from your system to his or her system, the attacker can steal passwords, obtain sensitive information, and generally wreak havoc on your system.

Obtaining Supervisory Terminal Functions

The goal of most attackers is to obtain supervisory access to equipment.

Obtaining Supervisory Terminal Functions:

- The goal of most attackers
- Even gaining supervisory access for very short periods of time can cause major harm
- Attackers can:
  - Create additional supervisor accounts
  - Destroy valuable data
  - Do too many people have too many update attributes?

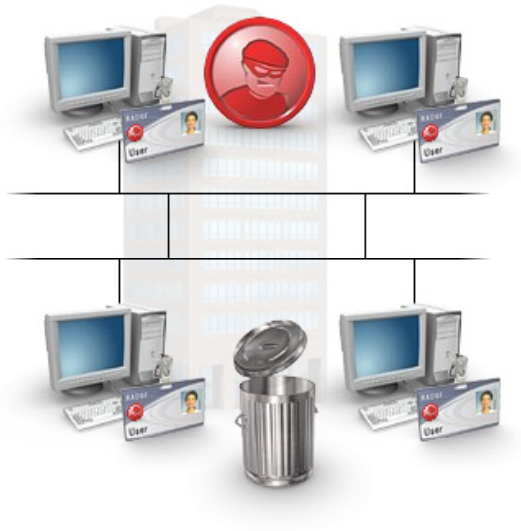

If the attacker does gain supervisory access, even for a very short period of time, he or she can create additional accounts on the system with supervisory access, open "doors" in the operating system, or destroy valuable data.

Bypassing System Logs

System logs are vital to administrators as they provide important information on what has happened to the system.

Bypassing System Logs:

- System logs provide important information on what has happened to a system
- Administrators need to be aware of these items
- If an attacker can bypass system logs, any mischief they do will go unnoticed
  - Errors might occur
  - Applications may halt
  - Services may fail
  - Services may be started

Errors might occur, applications may halt, services may fail, or services might be started. The administrator must be aware of these items to determine if a problem exists on the system or if an event occurred that is indicative of a problem about to occur. If an attacker can bypass system logs, then whatever mischief he or she has performed on the system is effectively hidden from the administrator.

System Shutdown/Downtime

If an attacker can gain minimal access to a system, he or she can immediately shut down the system, or invoke a script that shuts down the system at a certain time of day.

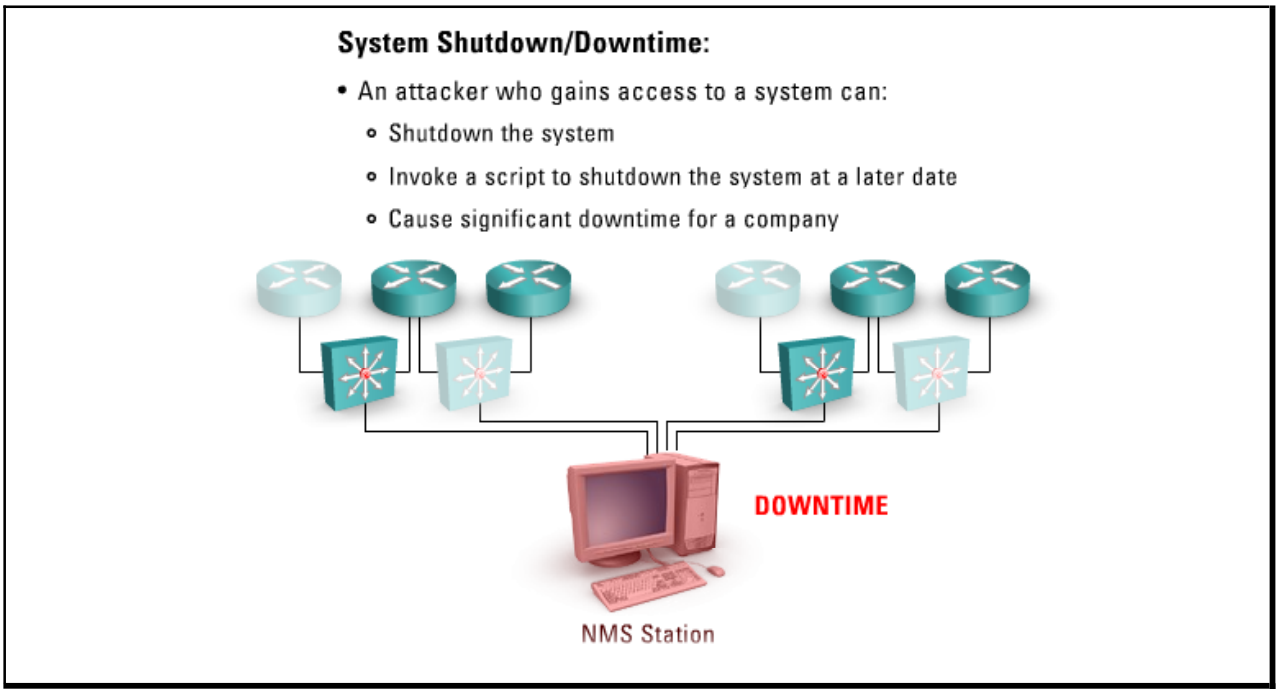

For example, one company used a network management application to manage its large base of routers and switches. A disgruntled employee obtained access to the main console and requested that all routers and switches be reloaded in sequential order during peak business hours. This request produced a long downtime for the company in which the company lost hundreds of thousands of dollars in revenue.

IPL from Tape, CD, or Floppy Disk

This topic discusses the effects of an attacker gaining physical access to a device.

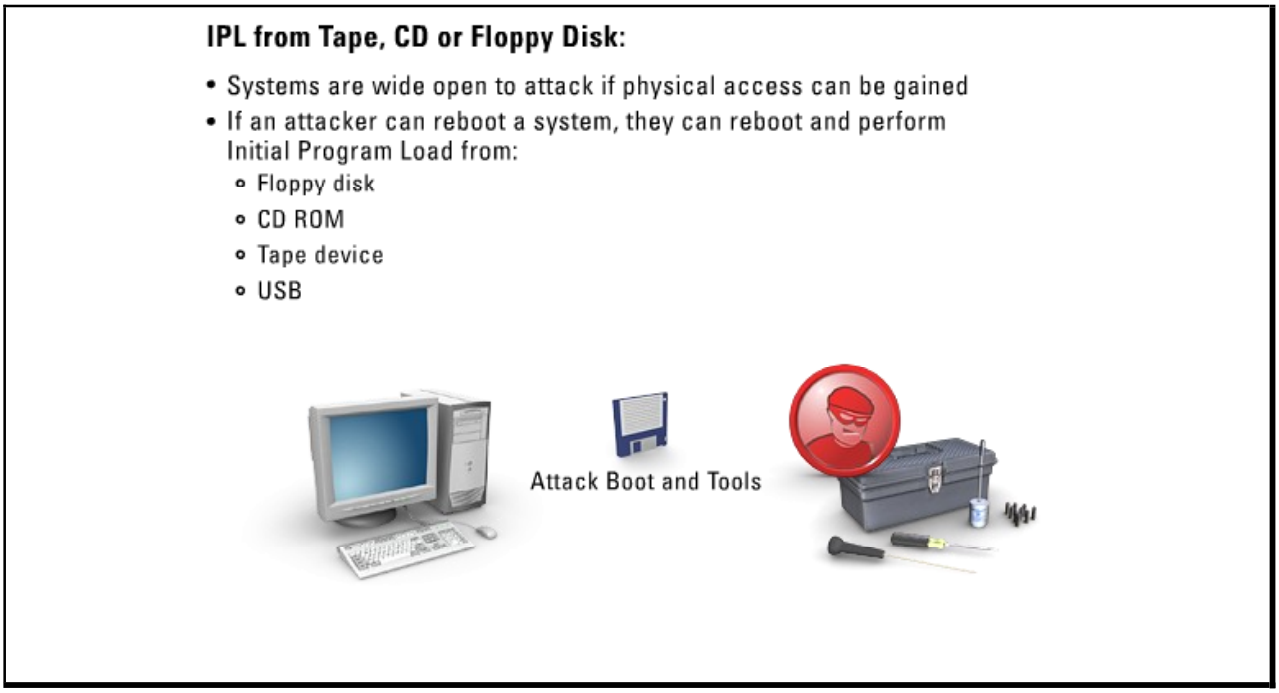

When operating systems are up to date on the latest patches, security updates, and hot fixes they are said to be fairly secure (until the next worm, virus, or buffer overflow is created). But no matter how secure a system is, it can be wide open if an attacker can gain physical access to it. If the attacker can gain physical access to the device, he or she can reboot the system (perform the initial program load [IPL]) and have it boot from a floppy disk, CD, or tape device, effectively bypassing all security measures of the operating system.

I/O System Generation

This topic discusses the effects of an attacker gaining access to a system's input/output (I/O) functions.

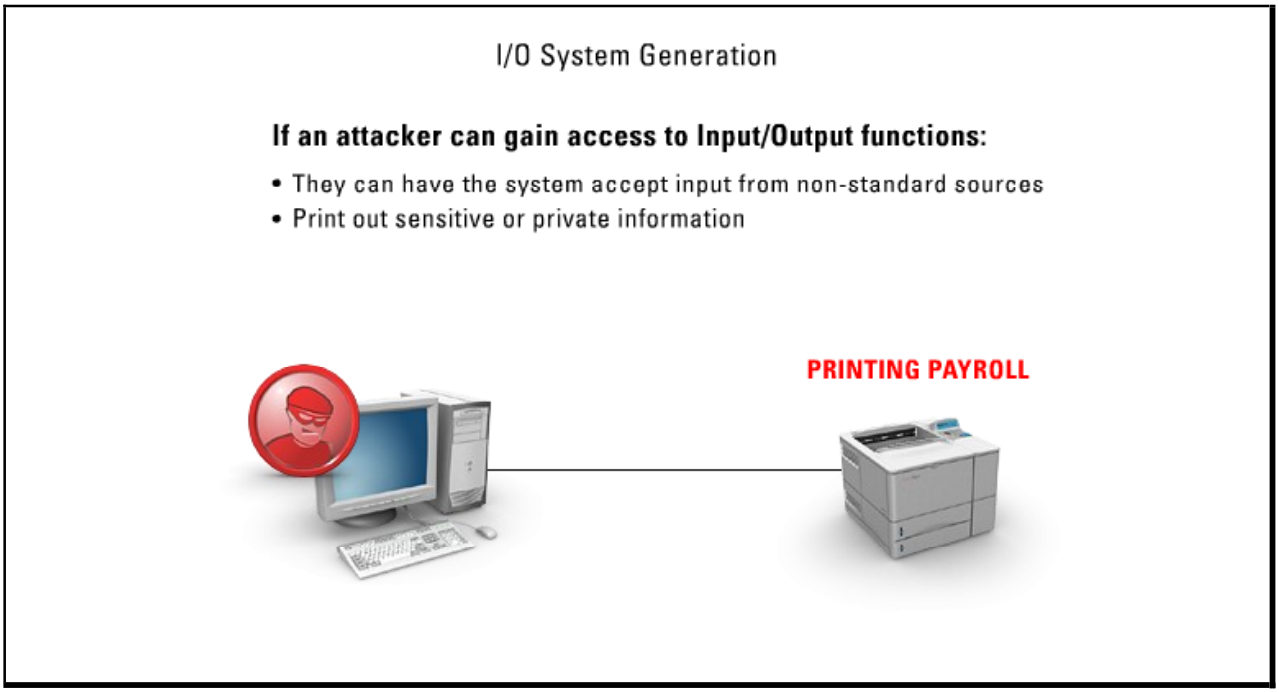

If an attacker can gain access to a system's I/O functions, he or she can have the system accept input from non-standard sources, or print out sensitive or private information.

Network Vulnerability

This topic discusses the effects of a network that is not properly secured.

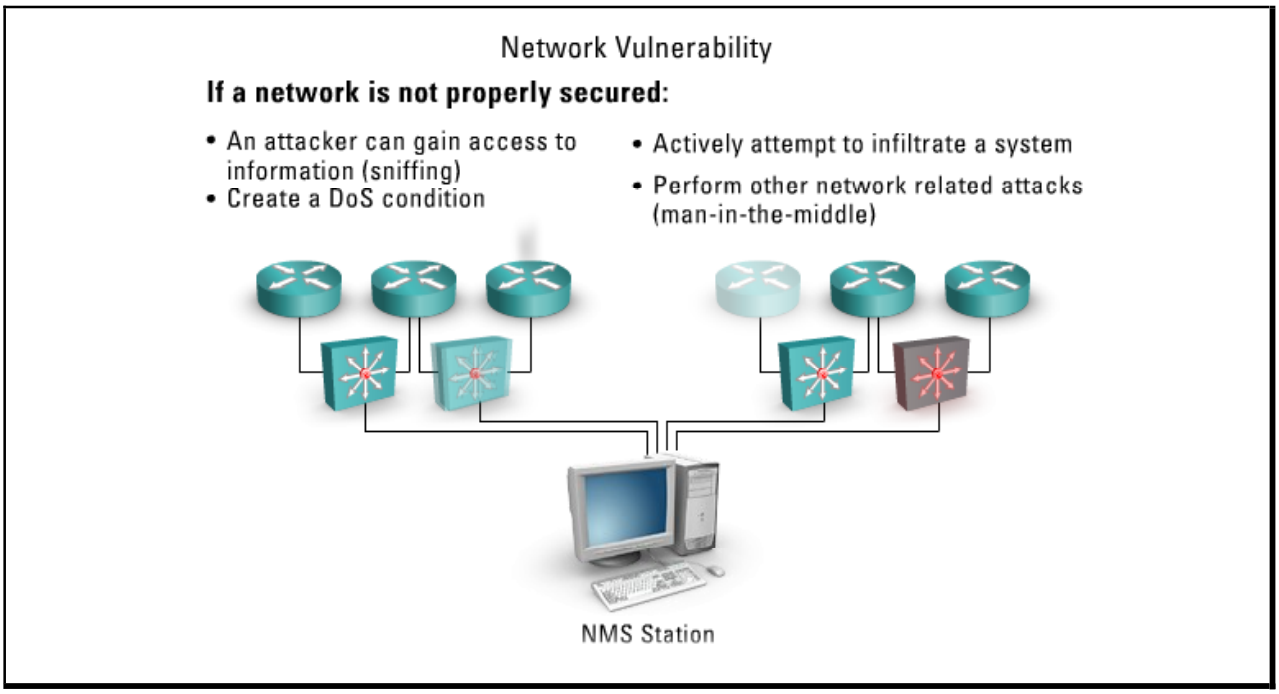

If the network is not properly secured, it is easy for an attacker to gain access to information (via sniffing), create a DoS condition, actively attempt to infiltrate a system, or perform any other attack that may involve access over the network medium.

Server Boot-up Sequence

This topic discusses the server boot-up sequence.

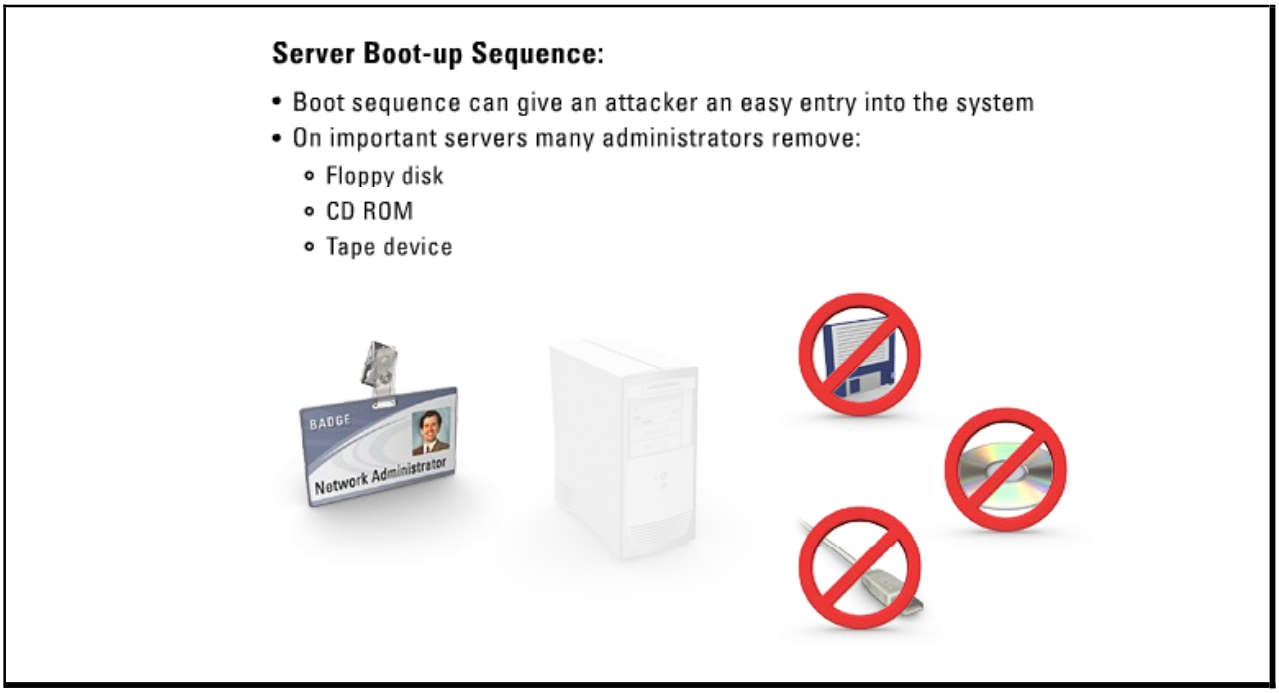

If the administrator of a system kept the boot order on a server to first boot from a floppy, then the CD ROM drive, and then the hard drive, it would be very easy for an attacker to place a floppy in the system and have it introduce a bug or Trojan into the system upon next reboot. For this reason, many administrators remove floppy drives and CD ROM drives from their important servers.

Highjacking Server ’s Network Address

If an attacker can “highjack” a server’s network address, he or she can begin to receive all information intended to be sent to the system.

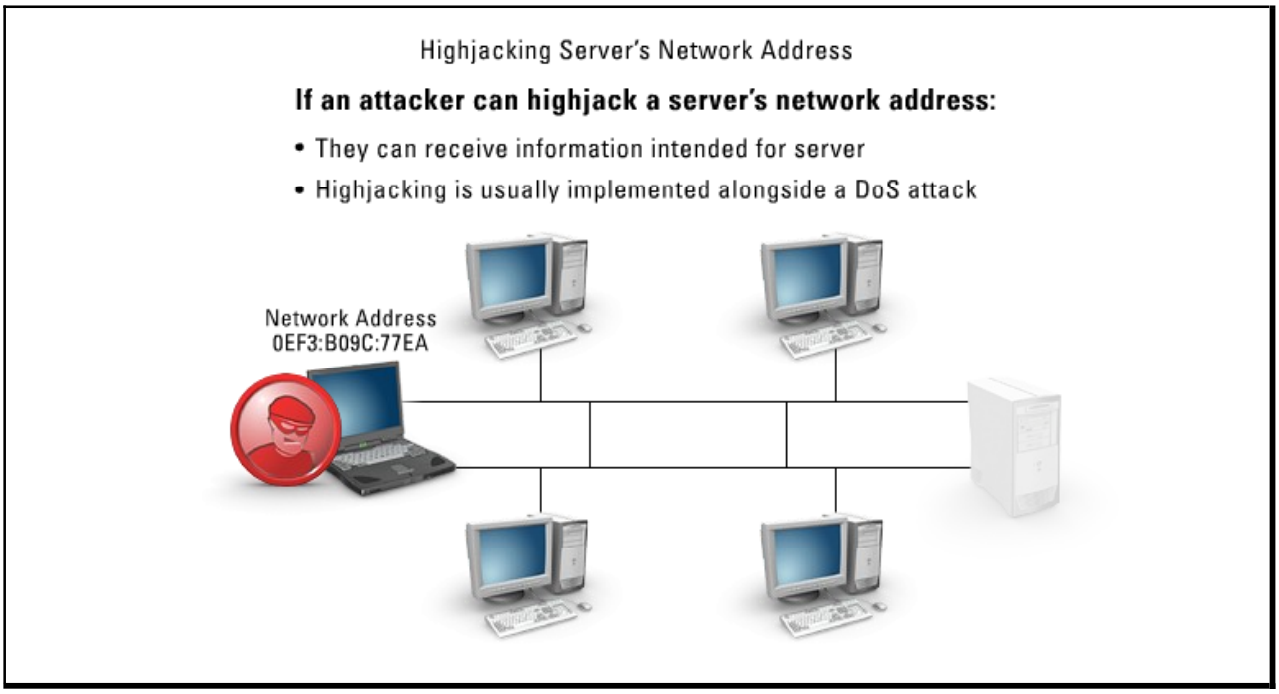

The attacker can obtain items such as usernames and passwords if the system is configured to present a login screen to users attempting to access it. In order for this type of attack to be successful, it is usually implemented alongside a DoS attack. When an attacker initiates a DoS attack against the server, his or her system is then brought up online and will begin to receive all packets that were destined to be sent to the server.

Stealing Password File/Table from the Server

This topic discusses the effects of an attacker stealing the password file.

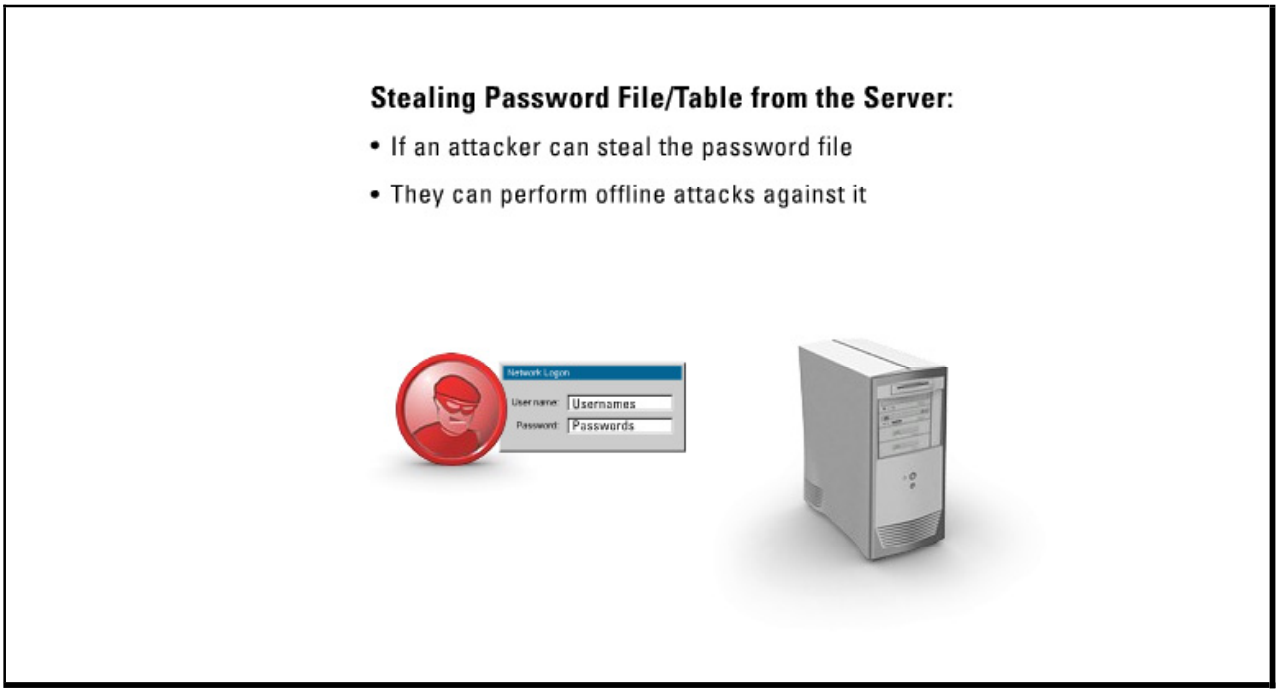

If an attacker is able to steal the password file off of a server, he or she could then perform offline attacks against the password file using an off-the-shelf password cracking program such as L0phtcrack.

Gaining Access to OS or Application User Accounts

This topic discusses how an attacker may try to gain access to the operating system or application user accounts.

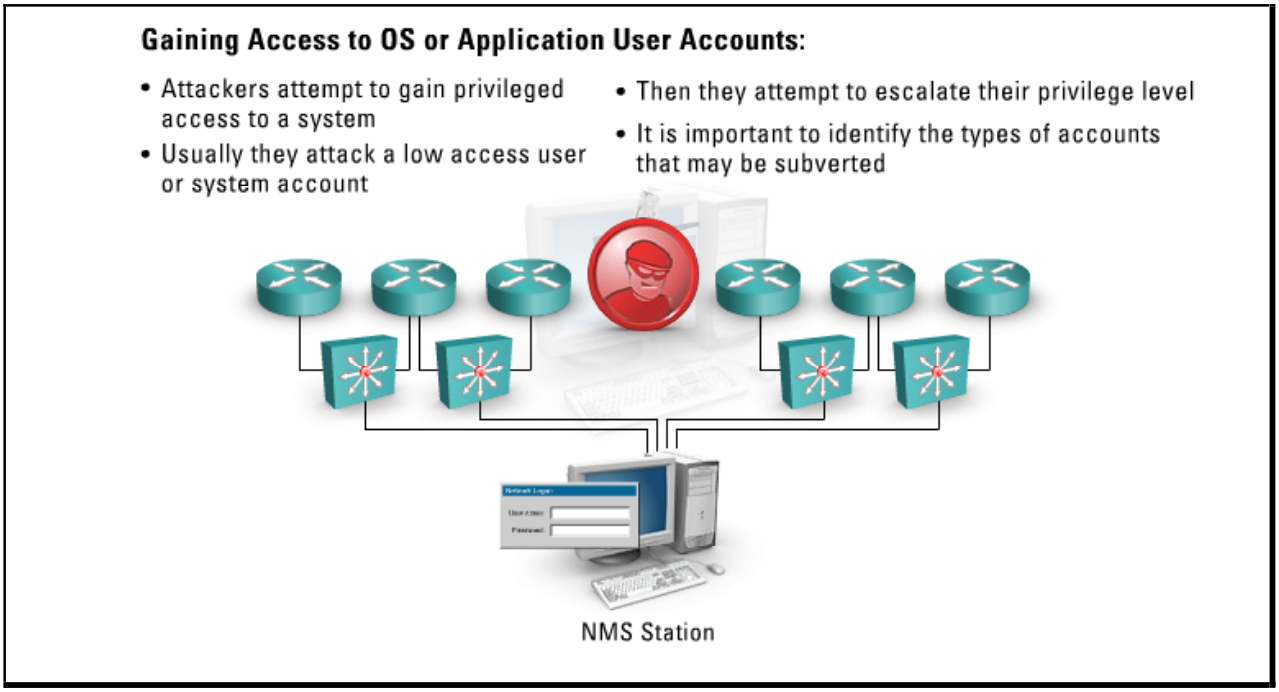

In order to infiltrate a system, an attacker will usually first crack an OS or user account, and then use that account to gain higher privileges, until eventually the attacker obtains root or administrator access. It is essential for you to identify the types of accounts that attackers may try to subvert or compromise.

Determining Clipping Levels

You can use **clipping levels** as security measures. It's important that security controls and mechanisms that are in place in your network, have a degree of transparency. That is, make sure it is easy for users to access and it is hard for attackers to get to the resource

Determining Clipping Levels

Set clipping levels on all sensitive systems:

- Number of failed login attempts
- Number of possible routes received from a neighbor
- Number of ICMP/UDP packets (DoS based attacks)
- Attempts to use disabled accounts

Once the clipping level is reached an action can be taken:

- Disable the service or account
- Block all further attempts
- Send a Syslog message
- Create and send a violation report

For example, you can set a clipping level by setting an operating parameter that will only allow a certain number of failed login attempts before an account is locked out. In addition, you can set it so the account is locked out for five minutes, a few hours, or even a full day after the clipping threshold has been exceeded.

Summary

The key points discussed in this lesson are:

- Violation analysis permits an organization to locate and understand specific trouble spots, both in security and usability.
- Potential hardware and software exposures include:
  - Device address modification, such as rerouting output, obtaining supervisory terminal access, or bypassing system logs
  - System shutdown/downtime, such as shutdown handled from console/operations area
  - IPL from tape, such as initial program load without security
  - I/O system generation, including terminals/printers/disks and tape drives
  - Network vulnerabilities
  - Server boot-up sequence from tape, CD, or floppy disk that can bypass O/S security
  - "Highjacking" a server's network address in order to capture traffic to and from the server
  - Stealing the password file from the server
  - Gaining access to the OS or application user accounts
- You can use clipping levels as security measures.

Auditing

Overview

Very few crimes are “perfect” crimes; there are almost always bits of evidence left behind that can point to or help identify the attacker. This lesson will discuss the methods you can use to identify how an attack occurred and possibly identify the attacker involved.

Importance

Understanding how to walk the audit trail is vital in determining what system was attacked, when the attack occurred, and what the attacker was after.

Objectives

Upon completing this lesson, you will be able to:

- Define audit trails and audit events
- Define individual accountability
- Describe the concept of reconstructing events
- Identify problem identification techniques

Outline

The lesson contains these topics:

- Audit Trails
- Individual Accountability
- Reconstruction of Events
- Problem Identification

Audit Trails

Audit activities result in log reviews and attempts to identify if critical systems have changed "state." The motivation for maintaining audit trails is to determine if sensitive business resources are being used for authorized purposes only.

Audit Trails:

- Result in log reviews and attempt to identify changed "state"
- Used to determine if sensitive business resources are being used for authorized purposes only
- Typically include information to establish:
  - What event occurred
  - When the event occurred
  - Who caused the event
  - How the event was detected
  - When the event was detected
- There are two types of audit events
  - Failure events
  - Successful events

Audit Trail

An **audit trail** typically includes sufficient information to establish:

- What event occurred
- When the event occurred
- Who caused the event
- How the event was detected
- When the event was detected

The audit trail consists of **audit events**. There are two types of audit events: successful events and failure events. Successful events indicate that a user successfully gained access to a resource. Failure events indicate that the individual was not successful at accessing the resource, but did attempt to try to gain access.

Individual Accountability

An audit provides valuable information that can help you determine if security violations did in fact take place and the scope of the damage experienced.

Individual Accountability:

- The assurance that an audit trail supports accountability by providing a trace of that user's actions
- Analyzed information can also provide:
  - Users accessing information that does not relate to their job function
  - Attempts being made to access specific areas of a system
  - Are there accounts that have consistent authentication failures?

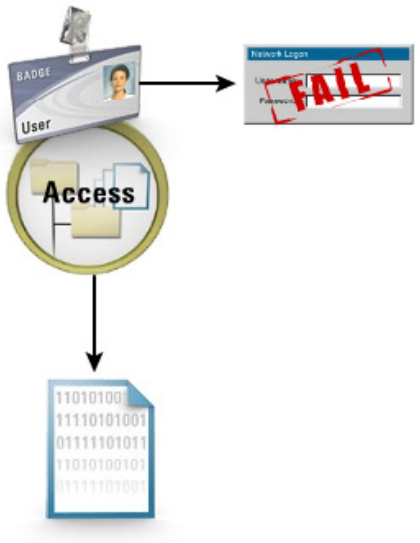

The analyzed information can also provide insight into areas such as:

- Are users accessing information that does not relate to their job functions?
- Are attempts being made to access specific areas of the system?
- Are there accounts that consistently have authentication failures?

The analysis of all such information will increase awareness of areas that you need to look at closely to prevent security violations.

Individual accountability is assured when an audit trail supports accountability of an individual by providing a trace of that user's actions.

Reconstruction of Events

This topic introduces the concept of reconstructing events.

Reconstruction of Events:

- Occur when you use an audit to support after-the-fact investigations
- Provide security incident/violation information such as:
  - How?
  - When?
  - What?

Reconstruction of events occurs when you use an audit to support after-the-fact investigations of how, when, and what occurred with respect to a security incident or violation.

Note	You can also useintrusion detection alongside auditing to provide supporting evidence as to when, how, and what occurred during a security incident.

Problem Identification

Problem identification and **problem resolution** are primary goals associated with auditing and monitoring. Monitoring contains the mechanisms, tools, and techniques that permit the identification of security events that could impact the operation of a computer facility.

Problem Identification:

Primary goals associated with auditing and monitoring include:

- Problem identification
- Problem resolution

Techniques for problem identification include:

- Intrusion detection: Used to analyze traffic patterns as well as functions as an intrusion detection system (IDS)
- Penetration testing: Violation processing using clipping levels

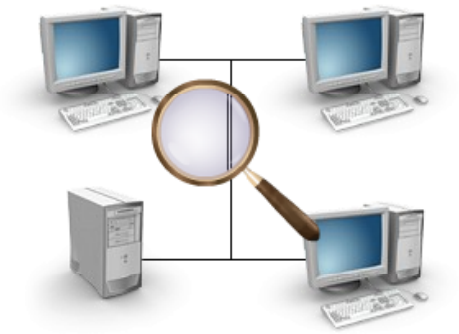

Techniques for problem identification include:

- **Intrusion detection**: Used to analyze traffic patterns as well as functions as an intrusion detection system (IDS)
- **Penetration testing:** Violation processing using clipping levels

Summary

The key points discussed in this lesson are:

- The motivation for maintaining audit trails is to determine if sensitive business resources are being used for authorized purposes only. There are two types of audit events: successful events and failure events.
- Individual accountability is assured when an audit trail supports accountability of an individual by providing a trace of that user's actions.
- Reconstruction of events occurs when you use an audit to support after-the-fact investigations of how, when, and what occurred with respect to a security incident or violation.
- Techniques for problem identification include intrusion detection and penetration testing.

Monitoring

Overview

Monitoring is an effective tool that attackers use to gain access to systems. This lesson will discuss what monitoring is and the tools used to monitor systems.

Importance

Understanding how attackers use monitoring to gather information from a system is vital for the information security professional.

Objectives

Upon completing this lesson, you will be able to:

- List monitoring techniques
- Define warning banners
- Define keystroke monitoring
- Define pattern recognition
- Identify the benefits of trend analysis
- List monitoring tools

Outline

The lesson contains these topics:

- Monitoring Introduction
- Warning Banners
- Keystroke Monitoring
- Pattern Recognition
- Trend Analysis
- Monitoring Tools

Monitoring Introduction

Monitoring contains the mechanisms, tools, and techniques that allow you to identify security events that could impact the operations of your computer facility.

Monitoring Introduction:

- Monitoring allows you to identify security events that could impact the operations of the company
- Monitoring techniques include:
  - Intrusion detection
  - Penetration testing
  - Scanning and probing
  - Demon dialing
  - Sniffing
  - Dumpster diving
  - Social engineering
  - Violation processing

Monitoring techniques include:

- Intrusion detection
- Penetration testing
- Scanning and probing
- Demon dialing
- Sniffing
- Dumpster diving
- Social engineering
- Violation processing using clipping levels

Warning Banners

Warning banners are banners that display at login or initial connection. Warning banners state that the system is for the exclusive use of authorized users and that their activity may be monitored.

Warning banners:

- Display at login or initial connection
- Used to state:
  - The system is for exclusive use of authorized users
  - All activity may be monitored
- Good from a legal perspective
- If you do employ banners, make sure:
  - They do not reveal system information
  - They do not provide any type of invitational wording

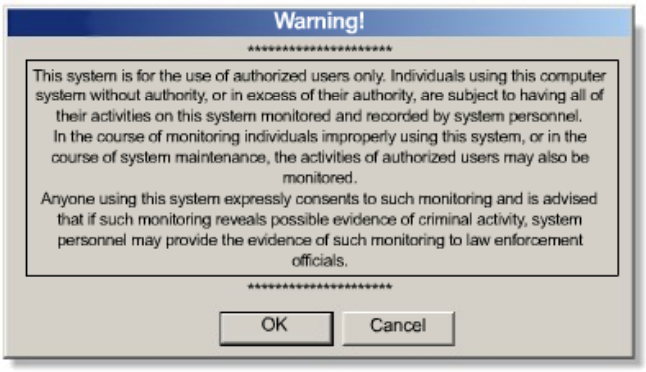

Warning banners are not foolproof, but are a good start, especially from a legal perspective in protecting the enterprise. If you do employ banners, always make sure that the banner does not reveal system information, i.e., OS, version, hardware, etc., and it does not provide any type of invitational wording, such as “welcome”.

An example of a warning banner would be:

```
                          **********
This system is for the use of authorized users only.
Individuals using this computer system without authority, or
in excess of their authority, are subject to having all of
their activities on this system monitored and recorded by
system personnel.

In the course of monitoring individuals improperly using this
system, or in the course of system maintenance, the activities
of authorized users may also be monitored.

Anyone using this system expressly consents to such monitoring
and is advised that if such monitoring reveals possible
evidence of criminal activity, system personnel may provide
the evidence of such monitoring to law enforcement officials.
                          **********
```

Keystroke Monitoring

Another type of monitoring is called **keystroke monitoring** where all keystroke activity is recorded. Keystroke monitoring can be performed on a specific sequence of keystrokes, such as when a user inputs his or her password, or it can be conducted on all keystroke activity.

Keystroke Monitoring:

- All keystroke activity is recorded
- Can be performed on:
  - Specific sequence of keystrokes
  - All keystroke activity
- If you use keystroke monitoring, be sure to:
  - Use it based on your organizations security policy
  - Advise all employees that it is occurring
  - Apply it to all employees in the organization

If you wish to perform keystroke monitoring, you should follow these guidelines to legitimize its use:

- Use keystroke monitoring based on your organization's information security policy
- Make sure that employees are well aware that keystroke monitoring is occurring
- Make sure that keystroke monitoring is applied to all employees in the organization

Pattern Recognition

The act of recognition can be divided into two broad categories: recognizing concrete items and recognizing abstract items. The recognition of concrete items involves the recognition of spatial and temporal items.

Pattern Recognition:

Recognition can be divided into two broad categories:

- Recognizing concrete items
  - Spatial items
    - Fingerprints
    - Weather maps
    - Pictures
    - Physical objects
  - Temporal items
    - Waveforms
    - Signatures
- Recognizing abstract items
  - Items that do not exists physically

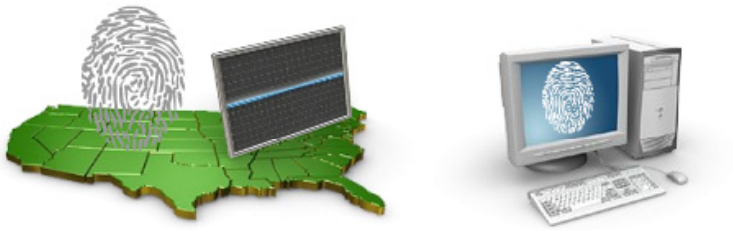

Examples of spatial items are fingerprints, weather maps, pictures, and physical objects. Examples of temporal items are waveforms and signatures. Recognition of abstract items involves the recognition of a solution to a problem, an old conversation or argument, etc.; in other words, recognizing items that do not exist physically.

Pattern recognition (also known as pattern classification) is a field within the area of computer science and can be defined as "the act of taking in raw data and taking an action based on the category of the data". It uses methods from statistics, machine learning, and other areas.

For the computer to recognize patterns, the patterns must be converted into digital signals and compared with patterns already stored in memory. Some uses of this technology occur in the areas of character recognition, voice recognition, handwriting recognition, and robotics.

Trend Analysis

This topic introduces the concept of trend analysis.

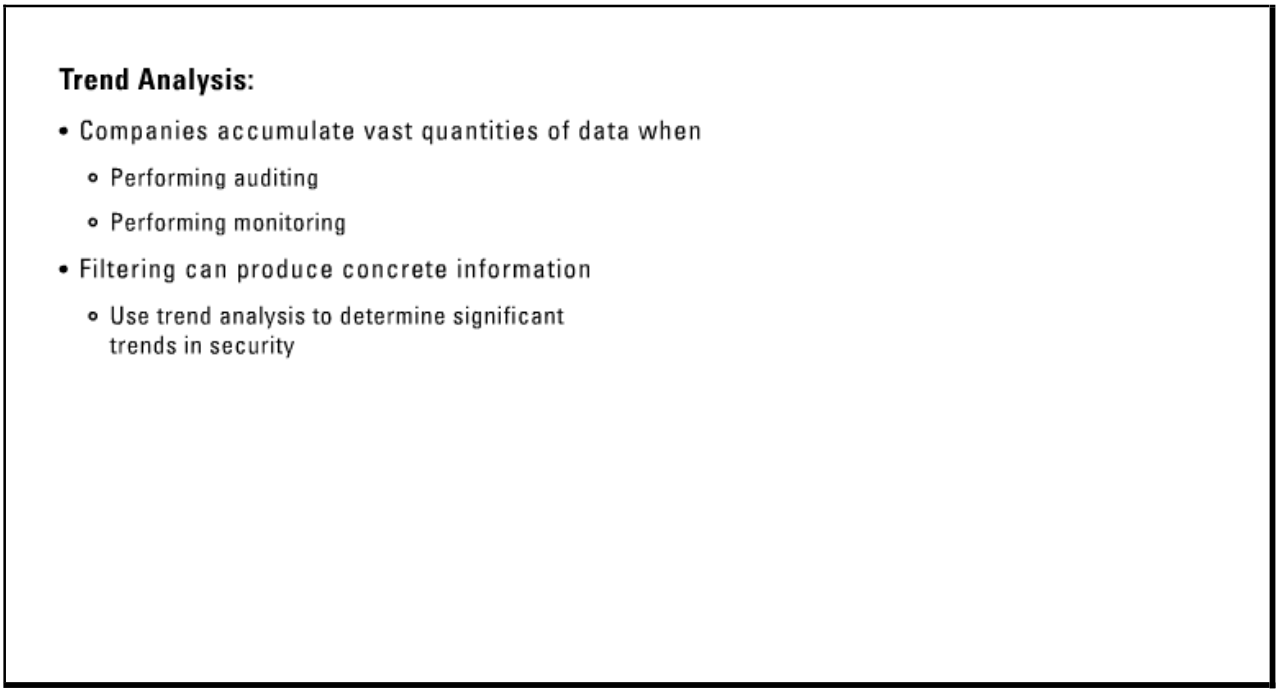

Many companies accumulate vast quantities of data when performing auditing and monitoring. Filtering the data can produce concrete information from a certain time or about a certain value. One such useful measure is performing trend analysis on the information. You can use trend analysis to determine or locate significant trends in security.

Monitoring Tools

This topic introduces tools you can use for monitoring.

Monitoring Tools:

Include tools such as:

- Network analyzers
- Protocol analyzers
- Trend analysis software
- Sniffer programs
- Auditing software
- Analysis software
- IDS
- Monitoring tools
- Keystroke monitors

Tools available for monitoring include:

- Network analyzers (LANSleuth, GNA)
- Protocol analyzers (Ethereal, Network Packet Analyzer)
- Trend analysis software (GPOWER, PASS, PATROL, WebTrends)
- Sniffer programs (Ethereal, SnoopAnalyzer, Network Probe, EtherSnoop)
- Auditing software (iInventory, eSMART, KnowledgeStorm)
- Analysis software (NetForm, SPSS)
- Intrusion detection systems (NetPatrol, Check Point InterSpect, Snort)
- Monitoring tools (Multi Router Traffic Grapher - MRTG, NetSaint, nPULSE)
- Keystroke monitors (Keylogger, Keyboard Monitor, PAL KeylogPro, KeyGhost)

Summary

The key points discussed in this lesson are:

- Monitoring contains the mechanisms, tools, and techniques that allow you to identify security events that could impact the operations of your computer facility.
- Warning banners are banners that display at login or initial connection. Warning banners state that the system is for the exclusive use of authorized users and that their activity may be monitored.
- Keystroke monitoring can be performed on a specific sequence of keystrokes, such as when a user inputs his or her password, or it can be conducted on all keystroke activity.
- For the computer to recognize patterns, the patterns must be converted into digital signals and compared with patterns already stored in memory.
- You can use trend analysis to determine or locate significant trends in security.
- There are a significant number of tools available to help you monitor your system.

Resource Protection

Overview

You must make the resources of your enterprise readily available to trusted users on the network, yet you must protect them from unauthorized access. This lesson will discuss what resources you must protect and how to protect them.

Importance

It is vital for the information security professional to understand how to protect enterprise resources in a secure and easy to use manner.

Objectives

Upon completing this lesson, you will be able to:

- Explain why enterprises create password files
- Explain why you must secure application program libraries
- Explain why you must secure source code
- Explain why you must secure vendor software

Outline

The lesson contains these topics:

- Password Files
- Application Program Libraries
- Source Code
- Vendor Software

Password Files

In order for users to authenticate with servers, there must be a shared secret password agreed upon between both entities. Users should remember their passwords, but servers also need to store them somewhere that is accessible to users, which is usually on the local hard drive. Instead of storing each individual user's password in a separate file, file servers store all user passwords in a single file called the password file. Since this single file stores all the passwords for a system, workgroup, or domain, it is essential that this password file be secure from prying eyes.

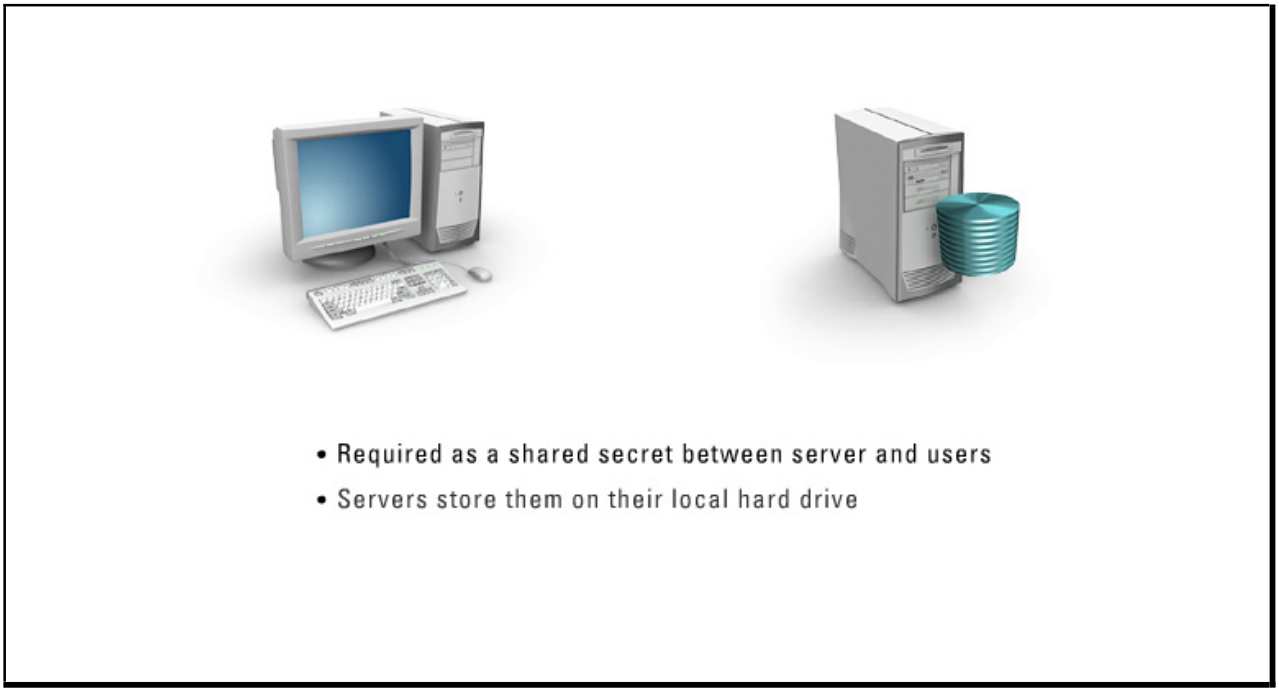

UNIX uses the /etc/passwd file to keep track of every user on the system. The /etc/passwd file contains the username, the real name, identification information, and basic account information for each user. Passwords are normally represented by a special encrypted format. The password itself is not stored in /etc/passwd. Instead, UNIX stores a value generated by using the password to encrypt a block of zero bits with a one-way function called crypt(). The result of the calculation is stored in the /etc/passwd file.

When you attempt to log into a UNIX system, UNIX does not actually decrypt your password in the /etc/passwd file and compare it to the one you typed. Rather, UNIX takes the password you typed and uses the crypt() function on another block of zero bits. UNIX then compares the result of the calculation with the value stored in /etc/passwd. If the values are equal, then you typed the correct password and the system lets you in. Almost all versions of UNIX support shadow password files. A shadow password file is a separate file, usually /etc/shadow, which contains the encrypted password. A shadow password file is protected so that it cannot be read by regular users, but can be read and written to by "root". Enabling shadow password files on your system is a good idea.

On each Windows NT/2000/XP/2003 system, the Security Account Manager (SAM) maintains a security account database. This database contains information about all user and group accounts on the machine. On networked systems, both individual workstations and collections of workstations can be grouped together to share a common user account database (SAM database).

Caution	If an attacker is able to obtain a password file, he or she can perform an offline attack against the system.

Application Program Libraries

Programmers create libraries of utilities, subroutines, programs, and applications.

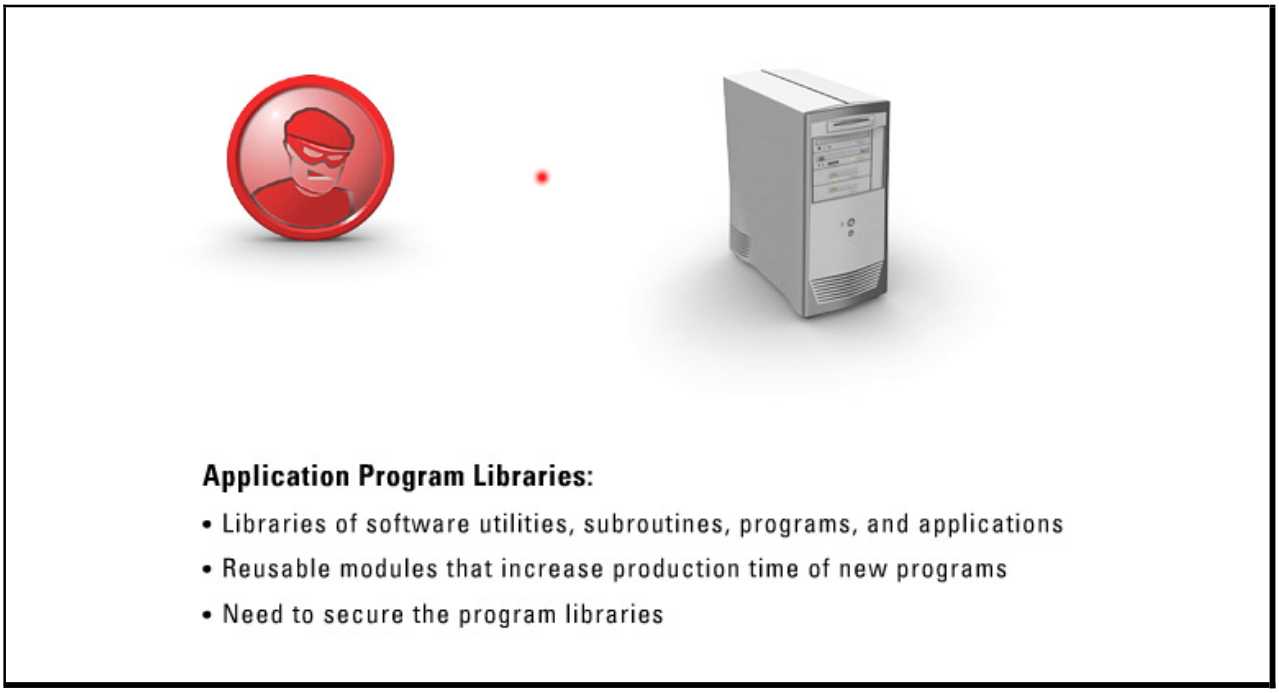

Once a programmer creates a tool, the programmer can then simply pull the tool (program) from the library and plug it into a different application. For this reason, you need to make sure that the program libraries are secure. If an attacker can obtain the code used to make an application, he or she can subvert the code to create a hole in the application.

Source Code

Initially, a programmer writes a program in a particular programming language. This form of the program is called the source program, or more generically, source code.

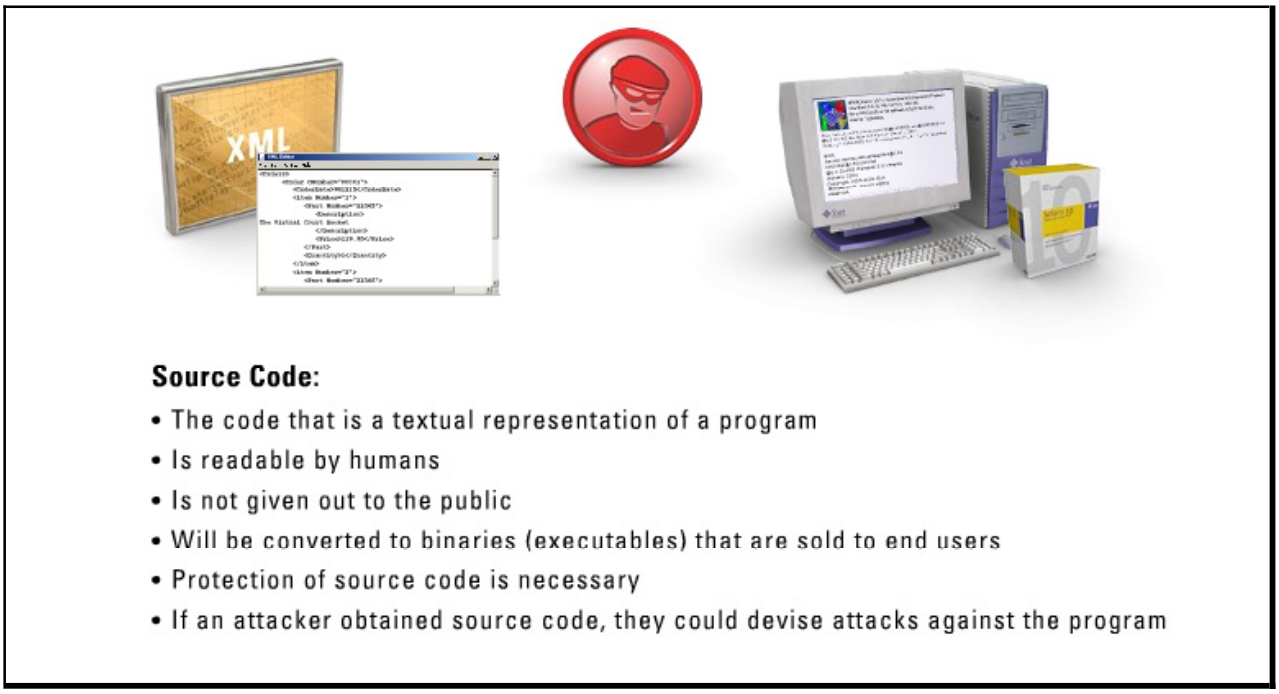

Source code is the only format that is readable by humans. When you purchase programs, you usually receive them in their machine-language format. This means that you can execute them directly, but you cannot read or modify them. Some software manufacturers provide source code, but this is useful only if you are an experienced programmer. If an attacker obtained the source code for an application, it would be much easier for the attacker to devise attacks against it.

Vendor Software

You must secure vendor software to protect the copyrights of the developers. You also need to protect vendor software to prevent the distribution of illegal copies, which could create legal issues for your enterprise.

Vendor Software:

- Must be secured to:
  - Protect the copyrights of developers
  - Prevent illegal distribution of copies (could create legal issues for your company)
- Warez- Compromising a company's public Internet server and placing illegal copies of software on it.
- For FTP servers:
  - Never allow read and write permissions on the same directory
  - Create one directory strictly for uploads (write capability)
  - Create one directory strictly for downloads (read capability)

If an attacker compromises a company's public Internet server, such as a file transfer protocol (FTP) server, the attacker can place illegally copied software on the server, and give access credentials to whomever he or she wishes. Then anybody can obtain the illegal software using the company's bandwidth and resources. This type of event is known as **Warez**.

Summary

The key points discussed in this lesson are:

- Instead of storing each individual user's password in a separate file, file servers store all user passwords in a single file called the password file. Since this single file stores all the passwords for a system, workgroup, or domain, it is essential that this password file be secure from prying eyes.
- You need to make sure that all program libraries are secure. If an attacker can obtain the code used to make an application, he or she can subvert the code to create a hole in the application.
- If an attacker obtained the source code for an application, it would be much easier for the attacker to devise attacks against it.
- You must secure vendor software to protect the copyrights of the developers. You also need to protect vendor software to prevent the distribution of illegal copies, which could create legal issues for your enterprise.

E-Mail Security

Overview

Electronic mail was created to allow disparate parties to communicate over long distances in a relatively short amount of time. Since the inception of e-mail, security was not a major concern, but that has drastically changed in today's marketplace. This lesson will discuss the protocols that make up the e-mail world, their insecurities, and what you can do to mitigate them.

Importance

Understanding e-mail systems and protocols is extremely important to the security specialist, as e-mail has become the de facto standard of communication in the workplace.

Objectives

Upon completing this lesson, you will be able to:

- Describe a typical e-mail infrastructure
- Describe the key features of the Multipurpose Internet Mail Extensionsprotocol
- Describe the key features of the Secure MIME protocol
- Describe the key features of the Pretty Good Privacy technology
- Identify the most common type of exploit that attackers use against e-mail servers and clients
- Describe the two types of spam messages seen on the Internet
- Define hoaxes
- Define SMTP relaying

Outline

The lesson contains these topics:

- Electronic Mail Security
- Multipurpose Internet Mail Extensions
- Secure MIME
- Pretty Good Privacy Technologies
- Electronic Mail Vulnerabilities

- Spam
- Hoaxes
- SMTP Relaying

Electronic Mail Security

E-mail access was one of the first protocols defined under the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite.

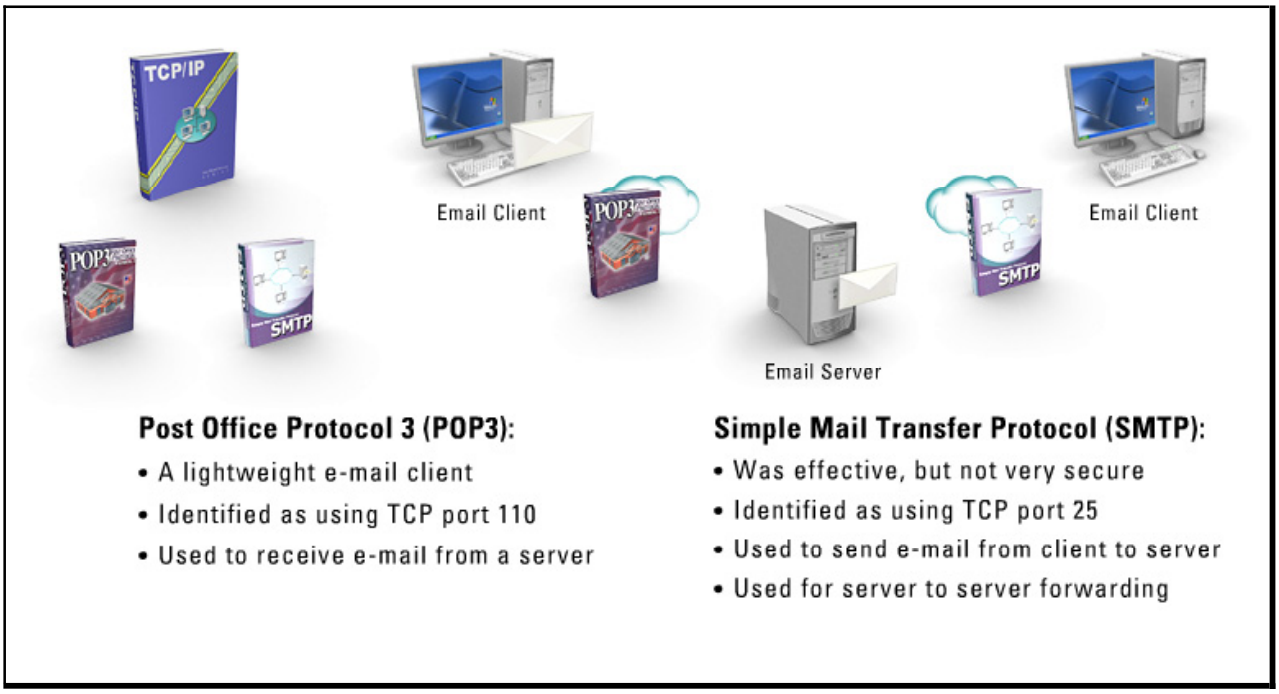

To obtain e-mail from a server, the **Simple Mail Transfer Protocol (SMTP)** was created. This protocol defines the mechanism a sender (e-mail client) uses to connect to, request, and send e-mail to the server (receiver). SMTP was created to be efficient and get the job done. SMTP was an effective protocol, but was soon to be found to be riddled with holes that attackers would soon begin to exploit.

SMTP can be identified as using TCP port 25 on the network. Although SMTP is a workable solution, many smaller companies could not take the overhead hit required in order to keep their e-mail infrastructures in place. A new 'lightweight' protocol was needed in order for a single workstation to connect to a server and request its e-mail. The **Post Office Protocol version 3 (POP3)** was created for this exact fashion. POP3 is intended to permit a workstation to dynamically access a maildrop on a server host in a useful fashion.

In today's typical e-mail infrastructure, e-mail gets into an SMTP server via the SMTP protocol and users obtain their e-mail via the POP3 protocol. In short, SMTP is used to send e-mail from an e-mail client to an e-mail server and POP3 is used to receive e-mail from the e-mail server to the e-mail client. POP3 can be identified as using TCP port 110 on the network.

Multipurpose Internet Mail Extensions

This topic discusses the **Multipurpose Internet Mail Extensions (MIME)** protocol.

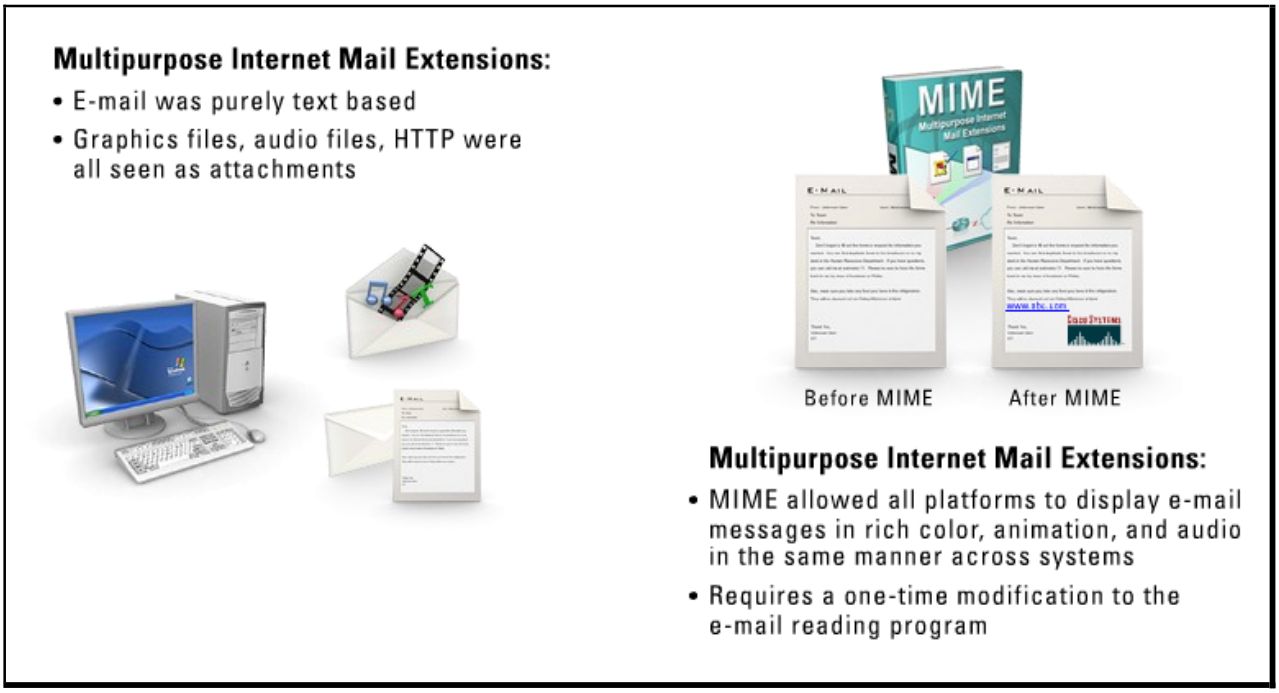

When e-mail first came into existence, e-mail messages were meant to be pure text only messages. As the Internet started to grow, people wanted to share more than just text; they also wanted to share graphic files, audio files, Hypertext Transport Protocol (HTTP), etc. However, they did not want these files to be seen as attachments. Rather, they wanted them to be seen dynamically when the e-mail document was opened. The problem manufactures had was there are a multitude of graphic and audio formats, as well as different platforms and operating systems. A standard was needed in order for all platforms to display e-mail messages in the same manner across systems. That standard turned out to be the MIME protocol, and is defined in RFC 1521 and RFC 1522.

MIME allows a one-time modification to e-mail reading programs that would enable the program to display a wide variety of messages types. This e-mail extension allows you to view dynamic multitype e-mail messages that include color, sound, animations, and moving graphics.

Secure MIME

MIME allowed users to display e-mail in a way they never could have before. Unfortunately, it did so without regard to security. E-mail was still subject to the same old hacks, such as sniffing and replay. A secure way of sending MIME data, one that guaranteed the confidentiality and integrity of the e-mail message, was necessary. **Secure MIME (S/MIME)** was created for this purpose.

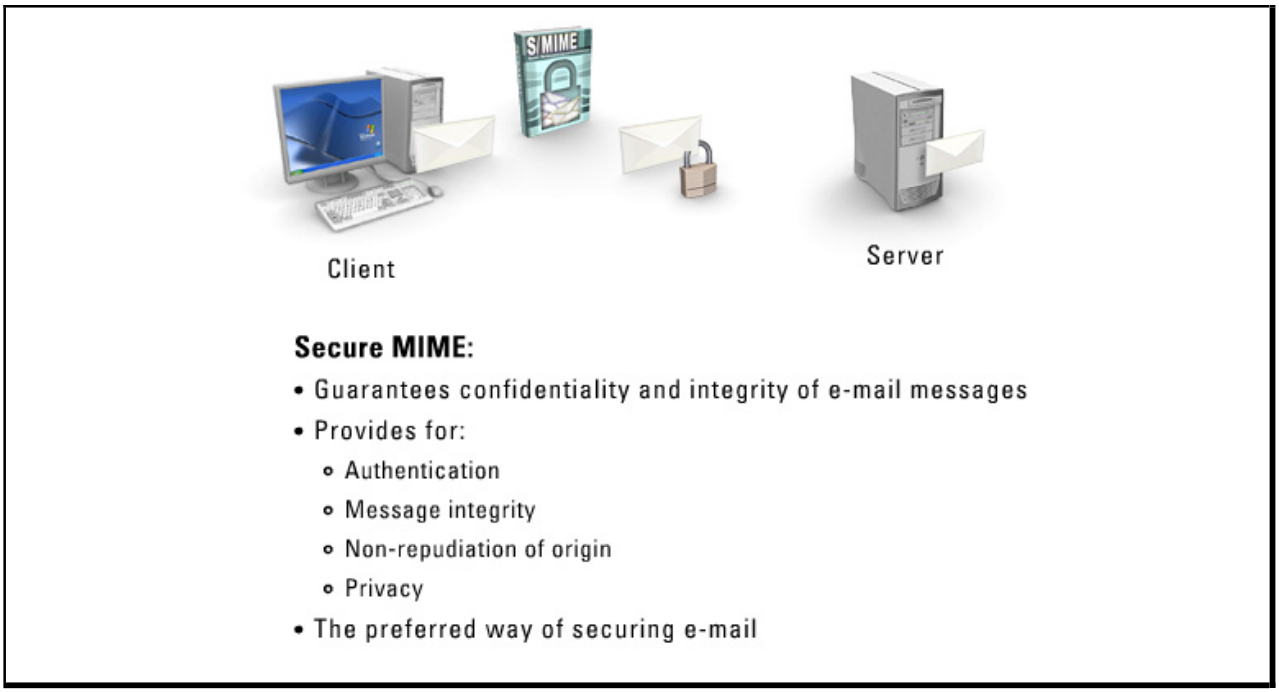

S/MIME provides cryptographic security services for electronic messaging applications by providing authentication, message integrity, non-repudiation of origin (using digital signatures), and privacy and data security (using encryption). Using S/MIME is the preferred way of securing e-mail as it traverses the unfriendly world of the Internet.

S/MIME version 2 is described in RFC 2311 and S/MIME version 3 is described in RFC 2633.

Pretty Good Privacy Technologies

Phil R. Zimmermann created the **Pretty Good Privacy (PGP)** technology in response to the 1991 Senate Bill 266. This ominous anti-crime bill had a measure in it that stated all encryption software must have a backdoor in it, which the U.S. Government could use to decrypt messages sent between parties. Being a staunch supporter of civil rights, Zimmermann created a cryptosystem in which no one except the two parties could read their e-mail messages.

PGP works using a public key cryptosystem. In this method, each party creates an RSA public/private key pair. One of these keys is kept private (the private key), and one is given out to anyone in the public Internet (the public key). What one key encrypts, only its partner private key can decrypt. This means if user X obtains user Y's public key and encrypts a message destined to user Y using its public key, the only person in the universe who can decrypt the message would be user Y, as he or she has the corresponding private key.

PGP is a hybrid cryptosystem in that before encryption is performed the e-mail data is first compressed. Compression not only makes an e-mail message smaller, it also removes any patterns found in plain text, which mitigate many cryptanalysis techniques that look for these patterns.

PGP performs the following security measures: confidentiality, data integrity, and sender authenticity.

Electronic Mail Vulnerabilities

E-mail is subject to many security exploits. Protocol exploits are the most common type of exploit that attackers use against e-mail servers and clients.

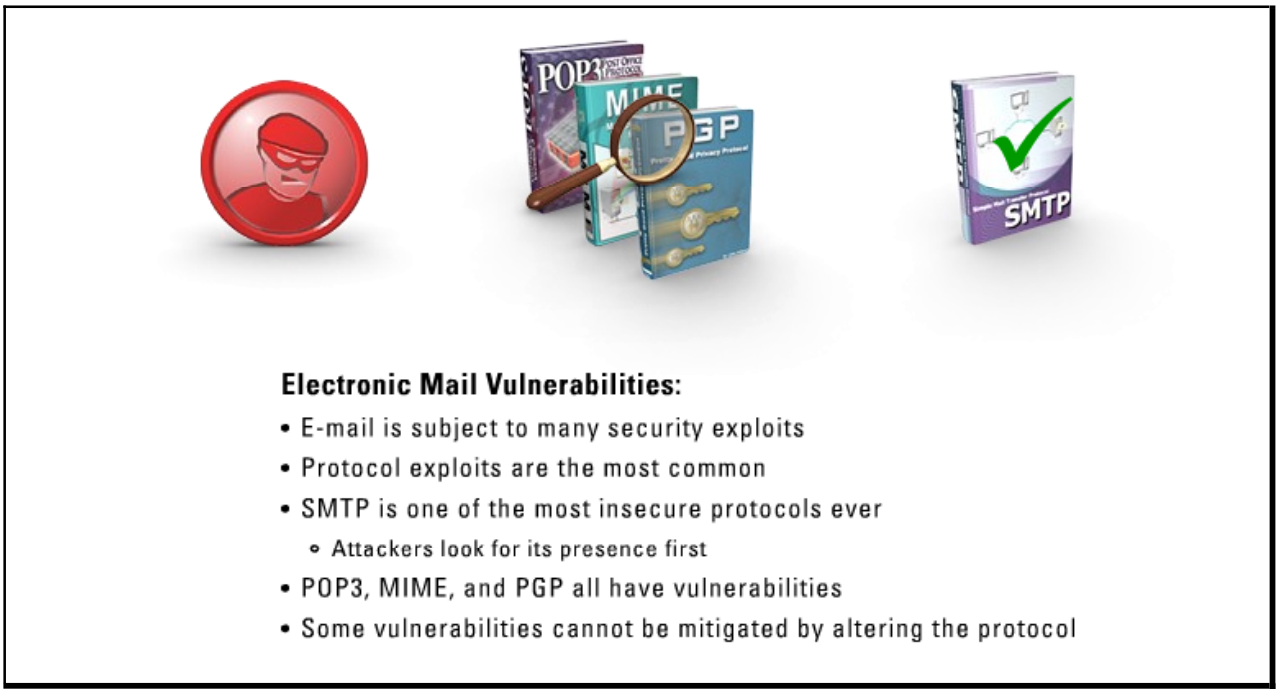

SMTP is one of the most insecure protocols ever created on the Internet as there are more cracks and exploits for it than almost any other protocol. The POP3 protocol, MIME types, and PGP also have vulnerabilities that attackers have exploited.

But there are some vulnerabilities to mail standards that cannot be mitigated by altering the protocol itself, which will be discussed over the next few pages.

Spam

Spam is defined as the unsolicited receipt of unwanted mail messages.

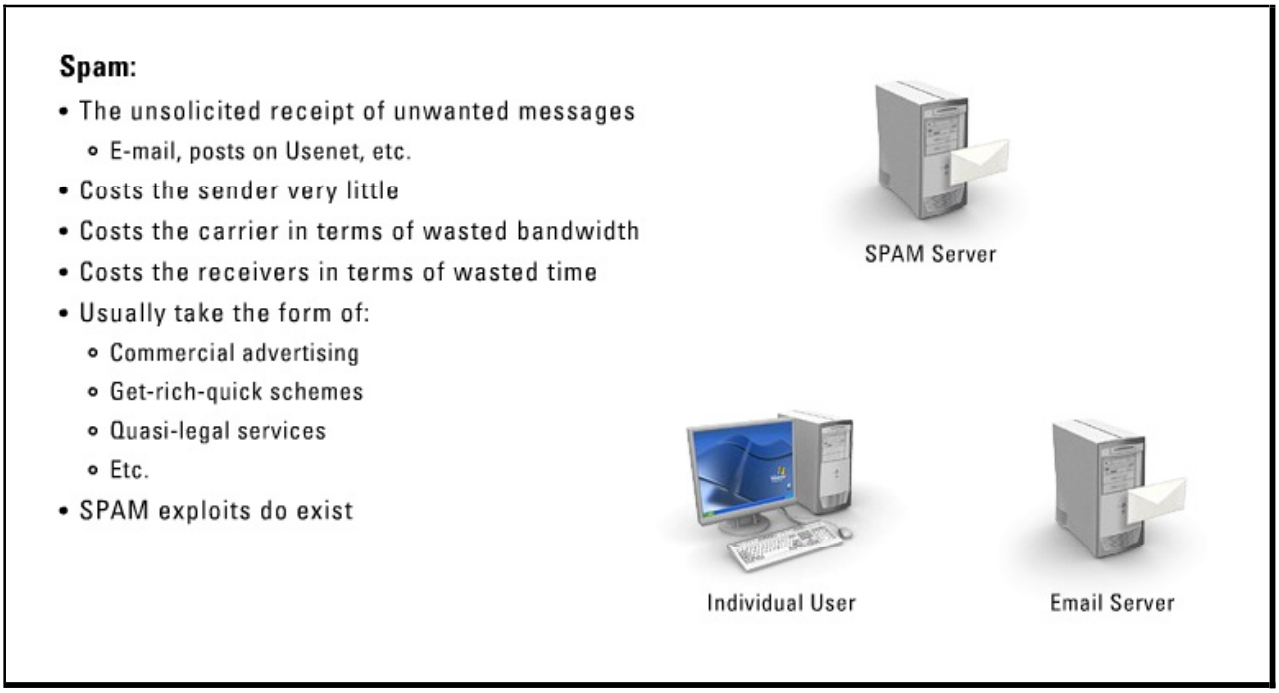

When a spammer sends large amounts of mail it costs the person very little, as most of the costs are paid for by the carriers in terms of wasted bandwidth. The recipient's e-mail server is also hit in terms of wasted resources, and the end user is hit in terms of wasted time looking at the junk mail.

Note	Most spam messages take the form of commercial advertising, often for dubious products, get-rich-quick schemes, quasi-legal services, and the like.

There are basically two types of spam messages seen on the Internet: usenet spam and e-mail spam. Usenet spam messages are those posts that are aimed at many newsgroups, while e-mail spam messages are those e-mail messages received by many parties. Individuals who scan usenet postings, steal Internet mailing lists, or search the Web for addresses often create e-mail spam lists. It has been commonplace for many companies to sell their customers' e-mail addresses to the highest bidder in order to turn a higher profit for their business.

Spam exploits do exist. These attacks are aimed at the e-mail server itself and are an attempt to deny service to normal users. Spam attacks are those that send very large amounts of forged and undeliverable mail messages to a server in order to use up all available resources, eventually crashing or halting the server.

Hoaxes

Hoaxes usually come in the form of an e-mail message.

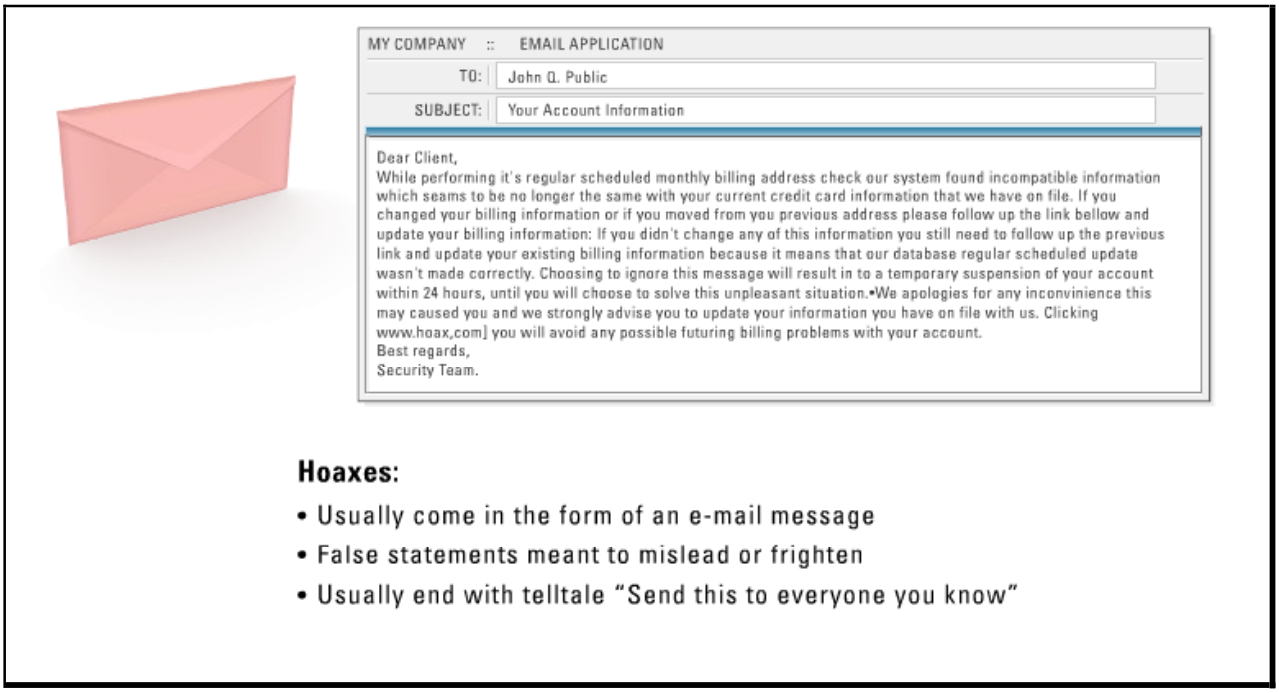

These messages are false statements meant to mislead or frighten users and usually end with the telltale "Send this to everyone you know". Users who receive warnings should always be wary when reading them. Also, always be sure of the legitimacy of the e-mail; you should always verify its accuracy before acting upon it.

Caution	Hoaxes do nothing more than scare people into performing some rash action so do not fall for them.

SMTP Relaying

SMTP relaying is the sending of e-mail messages via the SMTP protocol from one server to another.

Attackers use this function to send unsolicited e-mail (spam) and hide the real origin of their messages. Many Internet service providers (ISPs) also use SMTP relaying in their war against spam. Here the ISP allows SMTP traffic from their local inside customers, but blocks all SMTP traffic from the outside. Mail traffic on the inside is then forwarded toward their final destination. This process also has the added benefit of hiding the destination SMTP server behind the SMTP relay, making it less susceptible to attack.

Summary

The key points discussed in this lesson are:

- In today's typical e-mail infrastructure, e-mail gets into an SMTP server via the SMTP protocol and users obtain their e-mail via the POP3 protocol. In short, SMTP is used to send e-mail and POP3 is used to receive e-mail. POP3 can be identified as using TCP port 110 on the network.
- MIME allows a one-time modification to e-mail reading programs that would enable the program to display a wide variety of messages types. This e-mail extension allows you to view dynamic multitype e-mail messages that include color, sound, animations, and moving graphics.
- S/MIME provides cryptographic security services for electronic messaging applications by providing authentication, message integrity, non-repudiation of origin (using digital signatures), and privacy and data security (using encryption).
- PGP works using a public key cryptosystem. In this method, each party creates an RSA public/private key pair. One of these keys is kept private (the private key), and one is given out to anyone in the public Internet (the public key). What one key encrypts, only its partner private key can decrypt.
- Protocol exploits are the most common type of exploit that attackers use against e-mail servers and clients.
- Spam is defined as the unsolicited receipt of unwanted mail messages. There are basically two types of spam messages seen on the Internet: usenet spam and e-mail spam.
- Hoaxes usually come in the form of an e-mail message. These messages are false statements meant to mislead or frighten users and usually end with the telltale "Send this to everyone you know".
- SMTP relaying is the sending of e-mail messages via the SMTP protocol from one server to another.

The Web

Overview

The World Wide Web (WWW) is a wonderful source of news and information that anyone in the world can use. Unfortunately, attackers like to exploit the web for their own personal gain. This lesson will discuss the Web, its underlying protocols, and how security plays a role in it.

Importance

Understanding how attackers use the infrastructure and protocols of the World Wide Web is essential to all security specialists.

Objectives

Upon completing this lesson, you will be able to:

- Describe the origin of the World Wide Web
- Describe the features of the Hypertext Transport Protocol
- Describe the security concerns of instant messaging
- Describe the features and security concerns of the Common Gateway Interface
- Identify the characteristics of a strong password

Outline

The lesson contains these topics:

- Introduction
- Hypertext Transport Protocol
- Instant Messaging
- Common Gateway Interface
- Passwords

Introduction

In 1991, the National Science Foundation (NSF) lifted its ban on commercial access on the Internet, and the Wide Area Information Servers came online, which provided a mechanism for indexing and accessing items on the Internet. Tim Berners then posted a notice on the alt.hypertext newsgroup on where people could download his hypertext-based Web server and line mode browser.

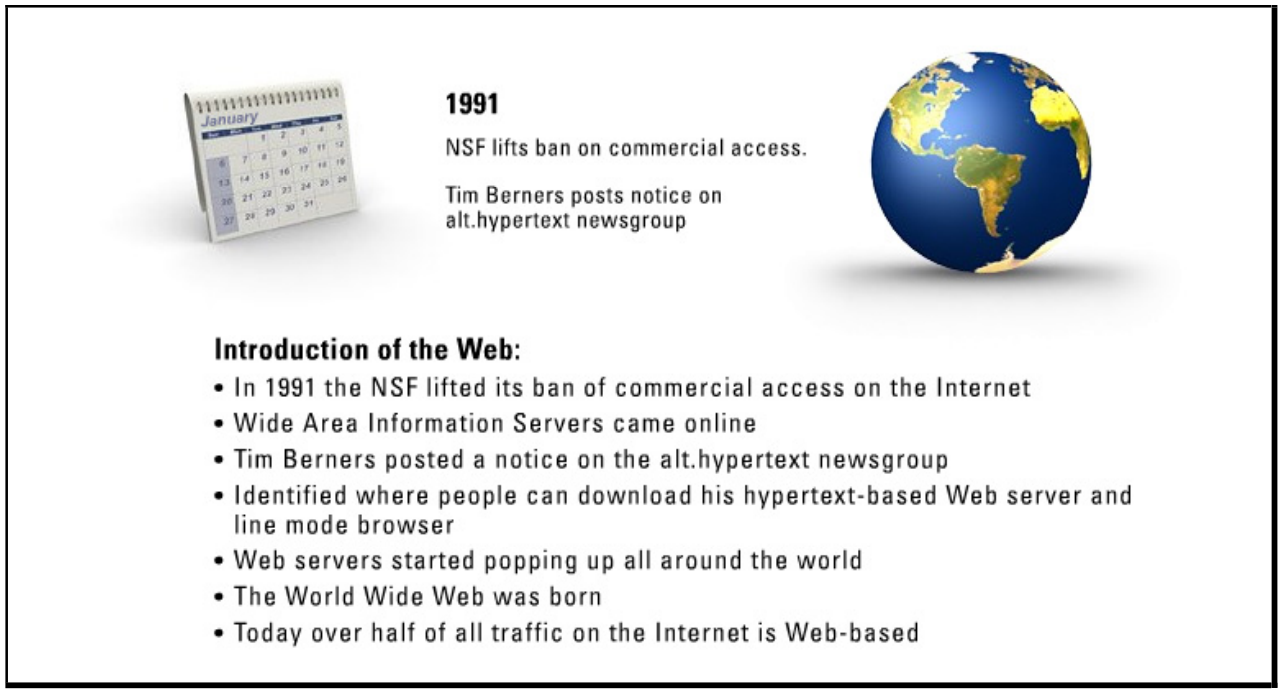

Web servers started popping up around the world almost immediately there after. Thus, the World Wide Web (WWW) was born. **HyperText Markup Language(HTML)** was used to display Web pages on the monitor.

Today, over half of all traffic on the Internet is Web-based, meaning Internet Protocol (IP) packets are carrying some type of **Hypertext Transport Protocol(HTTP)** traffic from server to client. This lesson will discuss the protocols you commonly see on the Internet and their security-related concerns.

Hypertext Transport Protocol

In 1992, the World Wide Web (WWW) consisted mainly of documents and links. Indexes were special documents that people could search. The search result was another document containing links to where you could find a particular document. You would use a simple protocol called "HTTP" to allow the browser program to perform your search request.

Hypertext Transport Protocol:

- In 1992, the Web consisted mainly of documents and links
- HTTP allowed the browser to perform a search request
- HTML was used to display the Web page on a monitor
- HTTP is used to carry HTML traffic across the Internet
- The first graphic based Web-browser was Mosaic
- HTTP is defined in RFC 2616
- Identified as TCP port 80 traffic

HTML was used to display a Web page on a monitor, while HTTP was used to carry HTML traffic across the Internet. Back in 1992, all Web traffic was text-based. In 1993, Mosaic, the first graphic-based Web browser, was released. HTML has evolved tremendously since 1991, but HTTP has remained, at its core, basically the same. HTTP still carries HTML payloads across the Internet. HTTP is defined in RFC 2616 and can be seen as using TCP port 80 on the Internet.

Instant Messaging

Instant messaging(IM) has become commonplace in the workforce as well as the home environment. IM protocols were created to allow anyone to communicate with anyone else under any possible configuration, which make them very difficult to control using normal security measures.

Instant Messaging:

- Created to allow anyone to communicate with anyone else in real time
- Difficult to control using normal security measures
- Do not support access control, confidentiality, and logging
- Allow users to transfer text messages as well as files
- Possibility that viruses, worms, and other malware can enter a network - bypassing any firewall or desktop security measure

IM applications are very difficult to secure as they inherently do not support access control, confidentiality, and logging. Blocking these applications is also very difficult as they are normally configured to hop from port to port, often using port 80 for communication, which must be opened for Internet access.

IM applications allow you to transfer text messages as well as files. Being able to transfer files means instant messengers can transfer viruses, worms, and other malware such as backdoor Trojan horses. Crackers can gain backdoor access to computers without opening a listening port, effectively bypassing any firewall or desktop security measures currently in place.

Note	Yahoo, Microsoft, and AOL all offer free IM software.

Common Gateway Interface

The **Common Gateway Interface (CGI)** is a standard for interfacing external applications with Web servers. In a normal HTML document, the client retrieves a static document, which is a constant unchangeable text file. A CGI program is different; it is executed in real time, so the output to the client is dynamic. This feature allows the client to affect how the Web page will look and feel.

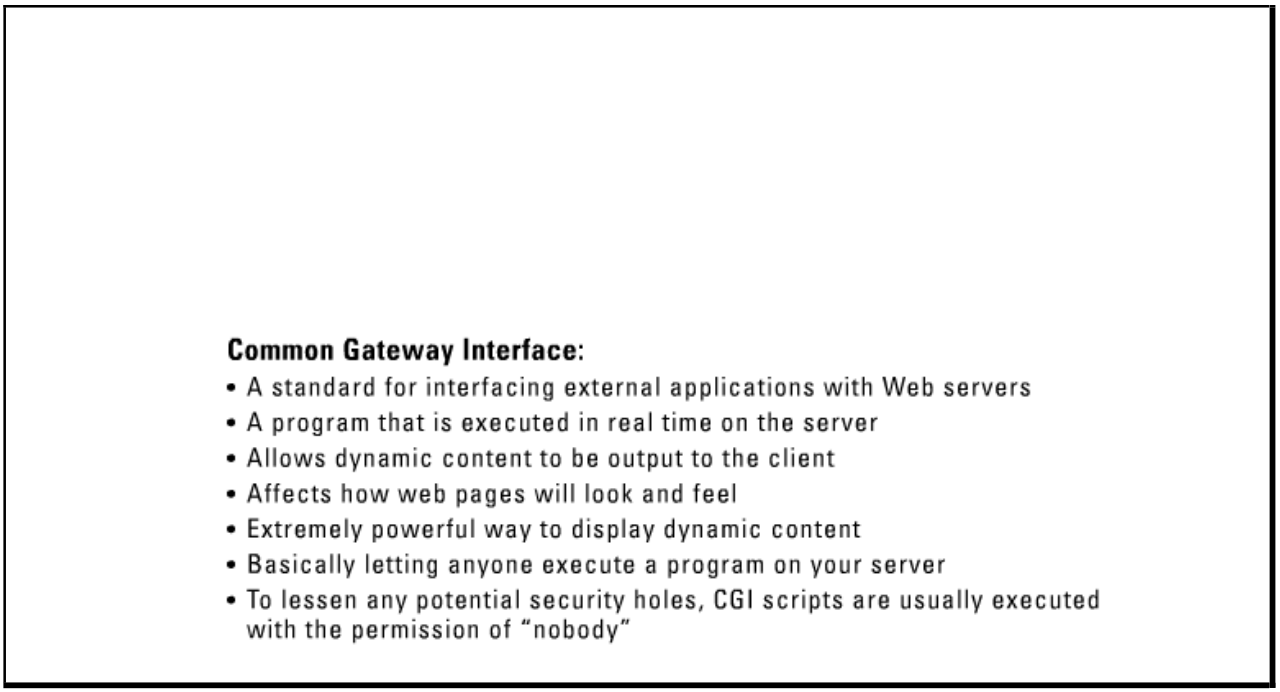

For example say you have an Oracle database that you want clients to query. The client will connect to a Web page and execute a CGI script. This script will ask for items for which to search. The CGI (gateway) will transmit information to the database engine, and then receive and display the results on the client's Web page.

CGI scripts are an extremely powerful way to display dynamic content for the client, but since CGI scripts are executable programs, they are basically equivalent to letting the world run a program on your server, which brings up many security concerns. Remember, if a cracker can compromise the CGI program, he or she can therefore gain access to the server and all its resources.

To lessen the potential security holes of CGI scripts, they are usually executed with the permission of "nobody", which gives the program very limited access to critical resources.

Passwords

One of the most important preventative security measures is the selection of good (strong) passwords for all accounts.

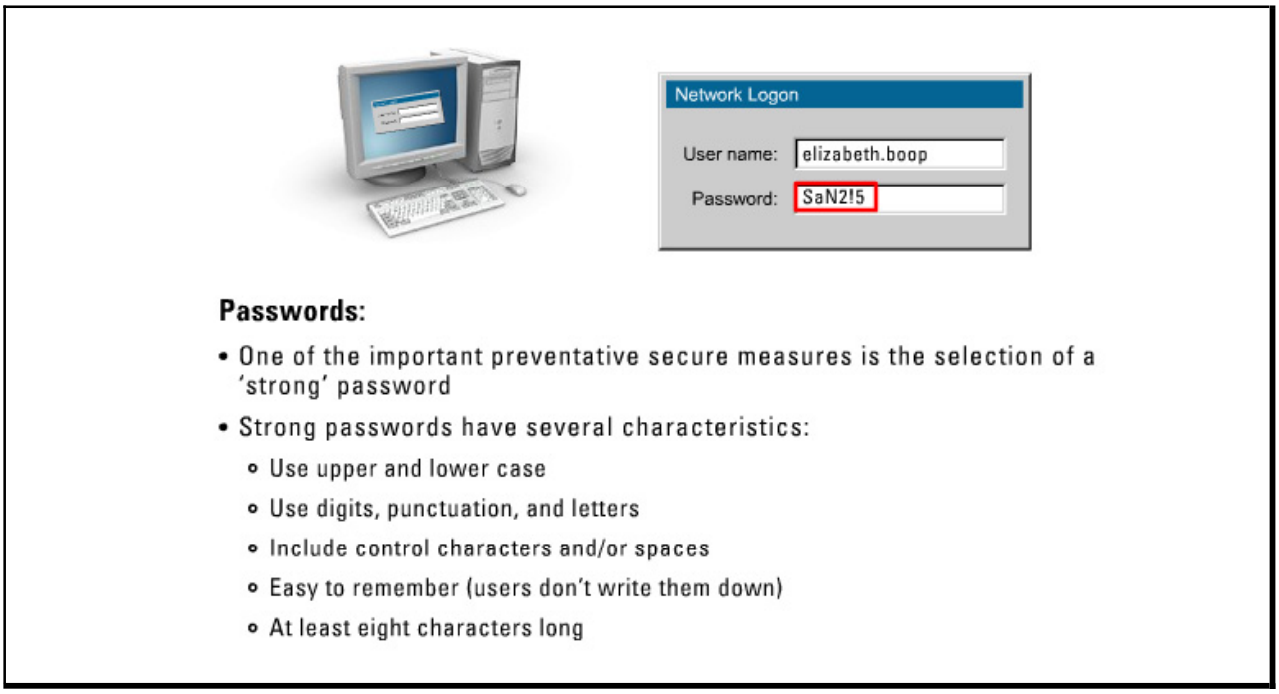

A good or strong password has several characteristics, which are as follows:

- Use of upper and lower case
- Use of digits, punctuation, and letters
- May include some control characters and/or spaces
- Easy to remember so the user does not have to write it down
- At least eight characters long (the longer the better)

Note	To anyone else, a strong password appears to be a seemingly random string of characters.

Why are eight-character passwords recommended? Using a very fast machine, passwords six characters or less can be matched in fewer than two days. Seven-character passwords can be matched in four months. By the time an eight-character password could be cracked, you should have changed the password to a new eight-letter string, thereby protecting your account.

Tip	An educational program that instructs users on the dangers of picking bad passwords versus good passwords is essential. Many system administrators will run a password guessing program to reveal easily guessed passwords. Well known password cracking software programs are LC4 and John the Ripper.

Use your operating system's configurable notification countdown schedule to inform users that their passwords will expire in XX days. Be sure to force the use of mixed case alphanumeric passwords.

Suggest strategies, such as using uncommon phrases and number-letter substitution strategies so they do not need to write their passwords down.

Summary

The key points discussed in this lesson are:

- In 1991, the National Science Foundation (NSF) lifted its ban on commercial access on the Internet, and the Wide Area Information Servers came online, which provided a mechanism for indexing and accessing items on the Internet.
- HTML has evolved tremendously since 1991, but HTTP has remained, at its core, basically the same. HTTP still carries HTML payloads across the Internet.
- IM protocols were created to allow anyone to communicate with anyone else under any possible configuration, which make them very difficult to control using normal security measures.
- CGI scripts are an extremely powerful way to display dynamic content for the client, but since CGI scripts are executable programs, they are basically equivalent to letting the world run a program on your server, which brings up many security concerns.
- One of the most important preventative security measures is the selection of good (strong) passwords for all accounts.

File Transfer

Overview

Being able to transfer files in the workplace, between friends, and over the Internet is important. Being able to transport these files in a secure manner is even more important. This lesson will discuss file transfer technologies and how security has been implemented to safeguard them.

Importance

Understanding the file transfer process and its protocols is important in understanding how to secure the process from attack or unauthorized access.

Objectives

Upon completing this lesson, you will be able to:

- Identify the reason for the creation of the File Transfer Protocol
- Describe the logistics of the File Transfer Protocol
- Identify the reason for the creation of Secure FTP
- Describe the logistics of anonymous FTP
- Describe the logistics of NetBIOS
- Describe the logistics of directory services
- Describe the logistics of the Lightweight Directory Access Protocol

Outline

The lesson contains these topics:

- Introduction
- File Transfer Protocol
- Secure FTP
- Anonymous FTP
- File Sharing
- Directory Services

- Lightweight Directory Access Protocol

Introduction

Computers were created to perform complex mathematical equations.

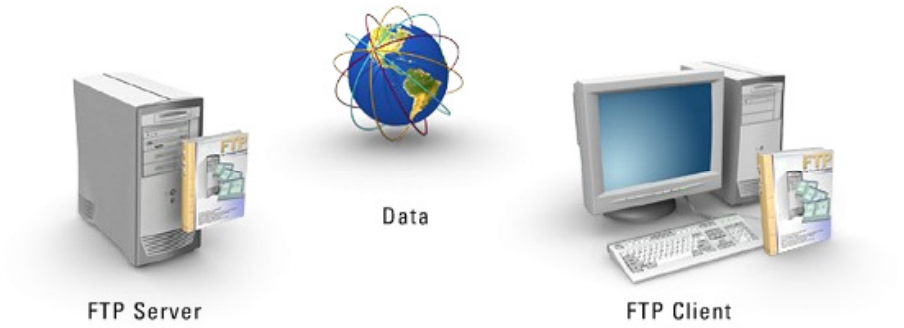

File Transfer Introduction:

- Programs and data in pre-networking days were transferred via the "sneaker net"
- Networks changed the way information was exchanged
- A standard method of file transfer was required
- The Internet community standardized on the File Transfer Protocol (FTP)

As time went on, computers became faster and were able to perform additional functions, which could be "programmed" into them. These programs were then made available to other computers, usually in the form of cards, tape, or floppy disk. Then came computer networking, which changed the way data information was exchanged. No longer was there a need for the "sneaker net" method of file transfer; an alternate electronic method now existed. Therefore, users needed a standard method of file transfer. In the local area network (LAN) community, different operating systems created their own methods for file sharing. In the Internet community, the File Transfer Protocol (FTP) was created.

File Transfer Protocol

The **File Transfer Protocol(FTP)** defined in RFC 959 was created to promote the sharing of files, programs, and data on the Internet in a manner that is efficient and reliable. It was also created to shield a user from the multitude of different file storage systems in use.

File Transfer Protocol:

- Defined in RFC 959
- Created to promote sharing of files and programs
- Shields users from the multitude of different file storage systems
- Has two ports defined for use:
  - TCP port 21 for control connections
  - TCP port 20 for data connections
- Runs in two different modes
- Standard mode:
  - Client initiates control connection
  - Server initiates data connection
- Passive mode:
  - Client initiates control connection
  - Client initiates data connection
- Uses a very insecure method of authentication
  - Usernames and passwords transmitted in clear text

FTP was created and added to the Transmission Control Protocol/Internet Protocol (TCP/IP) suite of protocols specifically for its file sharing capabilities. Since FTP runs on top of TCP, all errors and retransmissions were handled by TCP, which meant a smaller footprint (less code) protocol could be created.

FTP has two ports defined for its use. FTP uses TCP port 21 for the control connection, and TCP port 20 for the data connection. It can run in two different modes:

- **Standard Mode** - In standard mode, the client requests a file to be downloaded from the FTP server using TCP port 21. The server would then initiate an upload to the client on TCP port 20. In essence two sessions were created; one from the client to the server, then another from the server to the client, each using a different source and destination port.
- **Passive (PASV) Mode** - In passive mode, there are still two sessions being created, but the client initiates them both. The client requests a file to be downloaded from the FTP server using TCP port 21. The server responds with a data port the client should use to connect to. The client then initiates a different connection to the server's data port, and the file transfer begins.

FTP has a very insecure method of data retrieval. In order to download a file from a server, the client must first give credentials in the form of a username and password to the server. The problem with this method is that the username/password combination is sent across the wire in clear text. Anyone sniffing the wire can see the credentials being used as well as all traffic between the two systems as the data are also sent across in clear text.

Secure FTP

This topic discusses **Secure FTP (S/FTP)**.

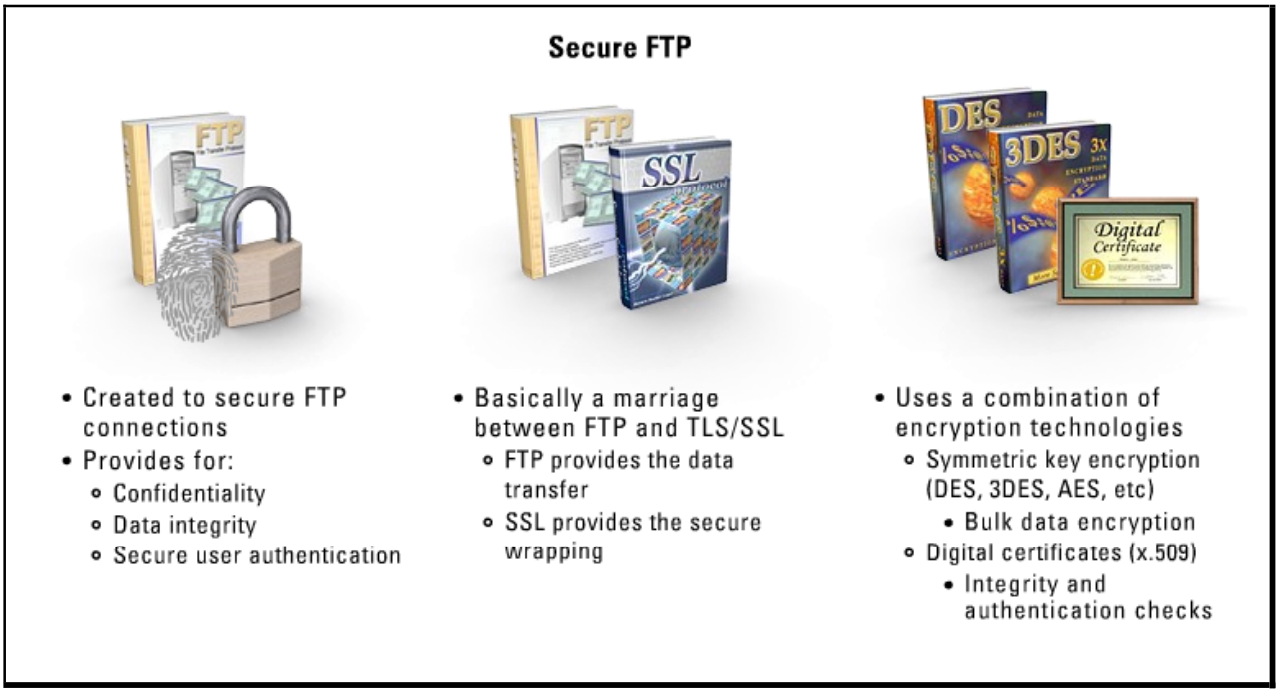

FTP control transmissions (used for setup and user authentication) are performed in the clear as are the FTP data transmissions. For simple everyday usage this is fine, but if you are going to transfer a database with very confidential data (such as credit card numbers), then you should definitely not use FTP. FTP needed to be updated to include some security measures, including confidentiality, data integrity, and secure user authentication. S/FTP was created to address these very issues. S/FTP is basically a marriage between FTP and Transport Layer Security/Secure Sockets Layer (TLS/SSL). FTP provides the data transfers, while SSL creates a secure wrapping for these transfers.

S/FTP uses a combination of encryption technologies (Data Encryption Standard [DES], Triple DES [3DES], Advanced Encryption Standard [AES], etc.) as well as digital (x.509) certificates for integrity and authentication checks. Public key cryptography using digital certificates is used to authenticate each end as well as perform a secure key exchange. The shared secret key will then be used to perform the bulk encryption/decryption of data.

Anonymous FTP

You can use **anonymous** or **blind FTP** when you have no need to identify or authenticate yourself to the FTP server.

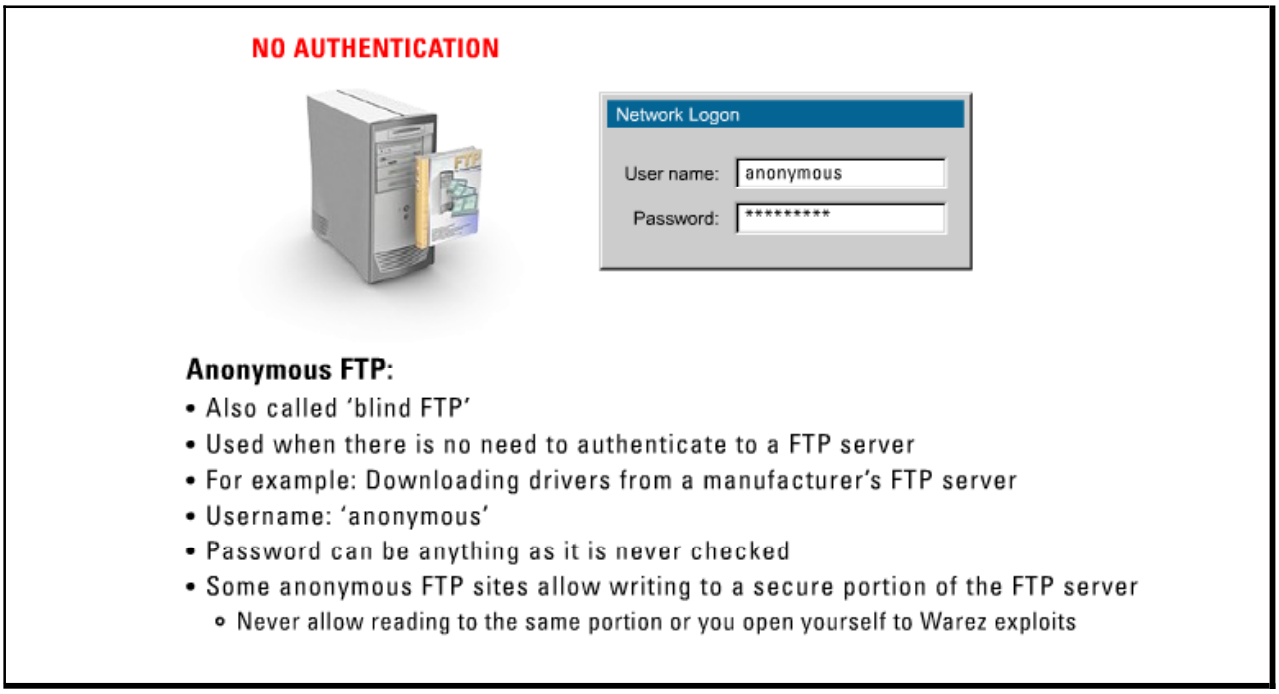

For example, many companies have drivers that often need to be updated on the client's system. Since they have thousands upon thousands of customers, giving each one a different username and password, or giving them all the same username and password can be an administrative nightmare. Instead it would be much easier for everyone to log on anonymously using a well-known username account.

That account is normally named "anonymous". When a user attempts to authenticate with an anonymous FTP server he or she simply supplies the username "anonymous". The password he or she supplies for this account can be anything, as the password itself is never checked. Once the user is authenticated, the user can then download the proper driver that he or she was after.

Some anonymous FTP sites also allow you to write data to a secure portion of the FTP server. This in itself is not a problem, but coupled with allowing users to also read from the secure portion of the site is when attackers can exploit anonymous FTP servers. Warez is the name of this exploit, and it is very simple in nature. Basically, if you allow someone to write data to your FTP server and other people to download the same files, you are opening yourself up to becoming a "pirated" FTP site. Users will upload popular high priced software to the site and make it available (usually through a pirate newsgroup) for download.

File Sharing

File sharing can occur in a multitude of fashions, including over a LAN, a wide area network (WAN), the Internet, or any other medium. File sharing became critical to a company's success when the first computer networks were created. Data, which were usually stored on a single machine, could now be passed from user to user during a project.

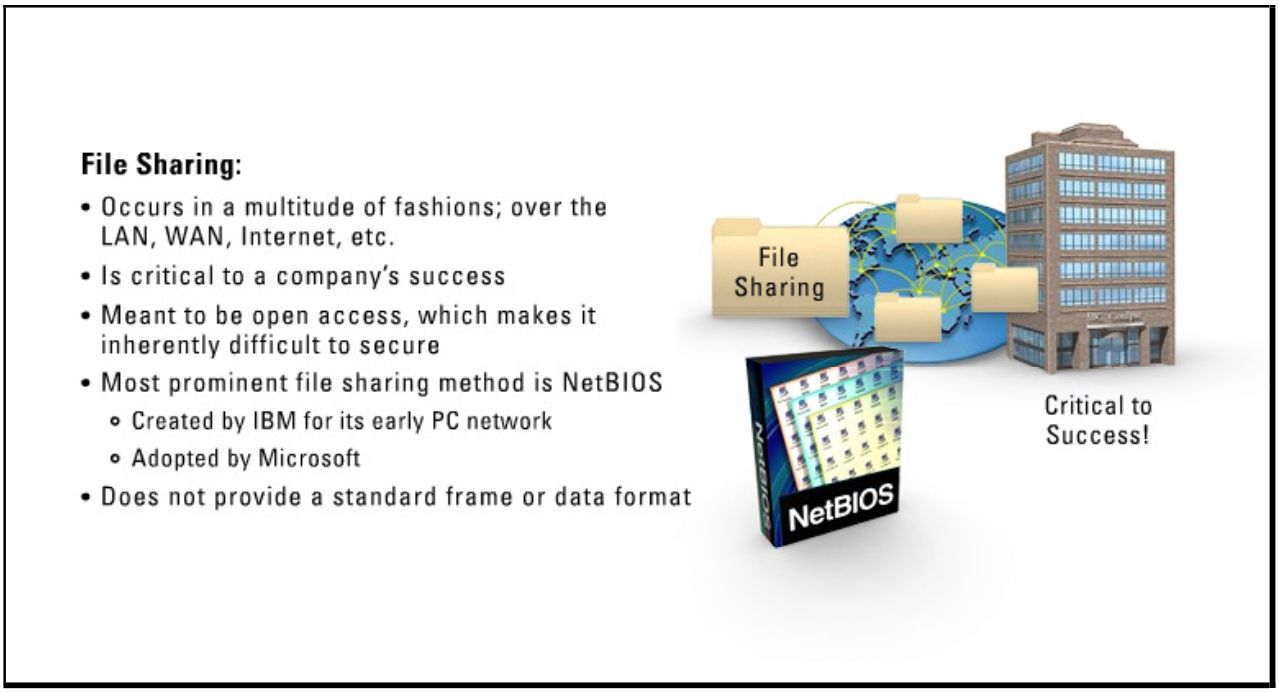

File sharing is meant to be open; you want others to be able to use your data. This fact makes it inherently very difficult to secure your data from those who you do not want to be able to access the data.

The most prominent file (and print) sharing method on the market is Microsoft's **NetBIOS** (Network Basic Input Output System) implementation. IBM actually created it for its early PC Network. Microsoft later adopted it, and it has since become a de facto industry standard. NetBIOS is an application programming interface (API) that augments the DOS BIOS by adding special functions for LANs.

NetBIOS provides the session t services described in the Open Systems Interconnection (OSI) model. NetBIOS sets up and maintains connections and is a non-routable protocol. NetBIOS uses broadcasts to spread information about servers. However, it does not provide a standard frame or data format for transmission. NetBIOS Extended User Interface (**NetBEUI**) provides a standard frame format. NetBEUI is a layer 3 protocol, which means in order to be passed across the WAN, NetBIOS frames must be encapsulated in another transport mechanism, such as TCP.

NetBIOS supports 3 services:

- Session service- This is a connection oriented protocol (peer-to-peer and client/server).
- Datagram service- This is used to deliver broadcasts and is thus connectionless.
- Name service- This means that no central name servers are required. There is a possibility to run NetBIOS over TCP/IP and then you need a name resolution-system. The two options for that are:

- The LMHOSTS file
- A WINS server (Windows Internet Name Service)

NetBIOS has a few well-defined ports:

- TCP port 137 is the NetBIOS Name Service port
- TCP port 138 is the NetBIOS Datagram Service port
- TCP port 139 is the NetBIOS Session Service port

Directory Services

When you want to find where something is located, you look in a directory. The same is true for computer networks.

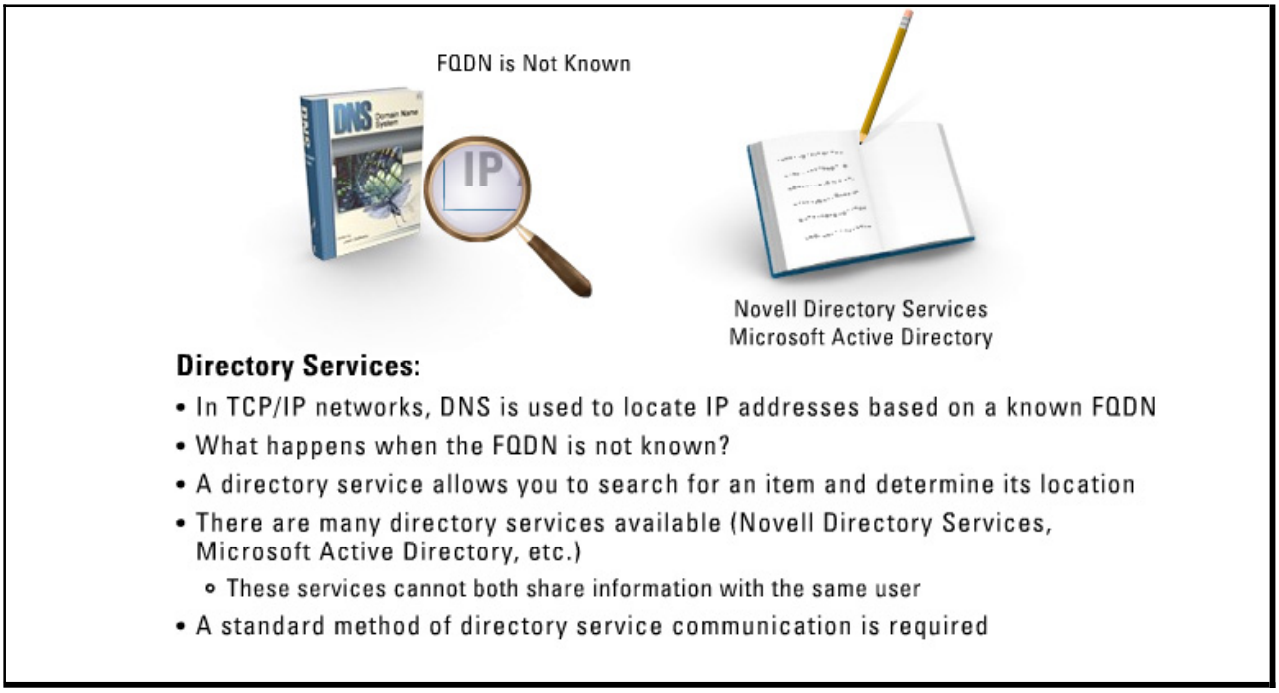

In TCP/IP networks, you use the domain name system (DNS) to locate an IP address based on its fully qualified domain name (FQDN). The problem with DNS is you must know the FQDN. What happens when you do not? A directory service will allow you to search for an item (host, individual, service, etc.) and determine its location. There are many directory services available for certain types of operating systems. For example, Novell uses the Novell Directory Services (NDS), and Microsoft uses Active Directory (AD). Unfortunately, these two directory services cannot share information with the same user. A user uses either NDS or AD; he or she cannot use both. In order for anyone using any type of computer system to use a directory, you must have a standard method of communication (a protocol).

Lightweight Directory Access Protocol

Lightweight Directory Access Protocol (LDAP) is a simple protocol that heterogeneous systems can use to locate information in a database and is defined in RFC 1777.

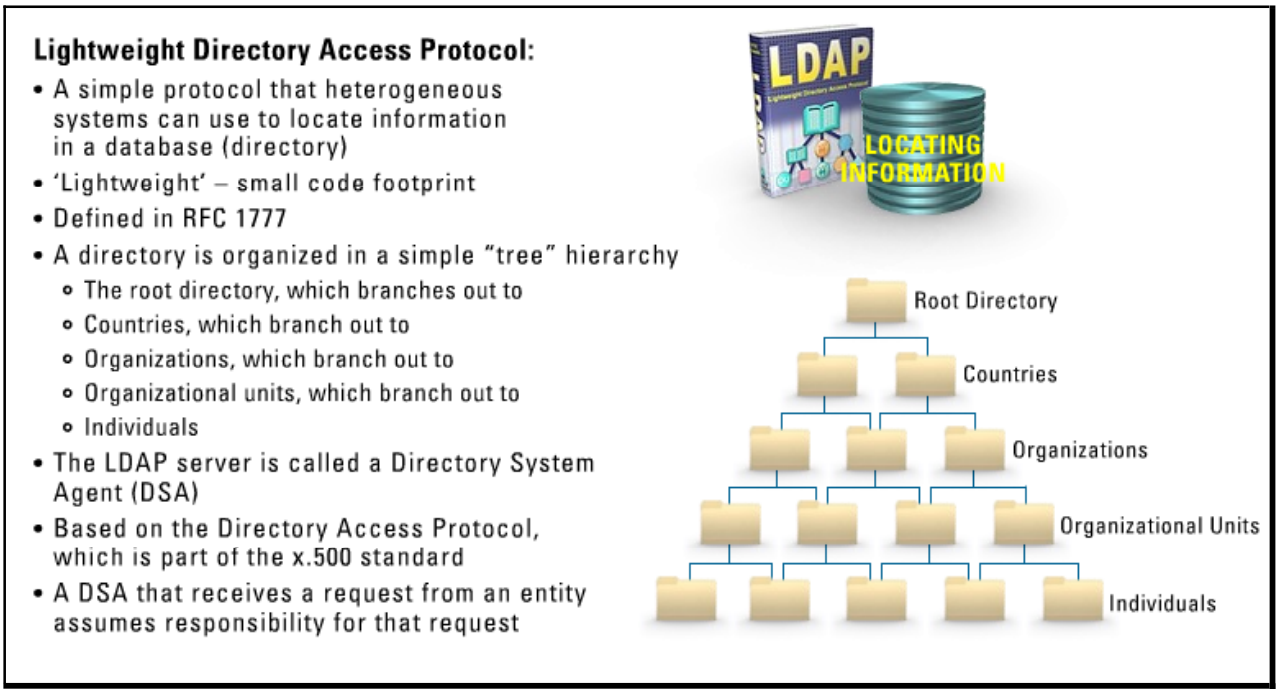

An LDAP directory is organized in a simple "tree" hierarchy consisting of the following levels:

- The root directory (the starting place or the source of the tree), which branches out to
- Countries, which branch out to
- Organizations, which branch out to
- Organizational units (divisions, departments, and so forth), which branch out to (and include an entry for)
- Individuals (which include people, files, and shared resources such as printers)

You can distribute LDAP directories among many servers, where each server can have a replicated version of the total directory that is synchronized periodically. The LDAP server itself is called a Directory System Agent (DSA) and is the entity that receives requests from users. The DSA that receives a user request takes responsibility for the request, and can pass the request to other DSAs as necessary, but always ensures that a single coordinated response is sent to the user.

LDAP is based upon the Directory Access Protocol (DAP), which is part of the x.500 standard. It is considered "lightweight" (smaller amount of code) because the initial version did not include security features. Secure LDAP (S/LDAP) uses the services of TLS/SSL.

Note	LDAP can be seen as using TCP port 389 on the Internet. S/LDAP can be seen as using TCP port 636 on the Internet.

Summary

The key points discussed in this lesson are:

- In the LAN community, different operating systems created their own methods for file sharing. In the Internet community, FTP was created.
- FTP was created to promote the sharing of files, programs, and data on the Internet in a manner that is efficient and reliable. It was also created to shield a user from the multitude of different file storage systems in use.
- FTP needed to be updated to include some security measures, including confidentiality, data integrity, and secure user authentication. S/FTP was created to address these very issues.
- You can use anonymous or blind FTP when you have no need to identify or authenticate yourself to the FTP server.
- File sharing became critical to a company's success when the first computer networks were created. Data, which were usually stored on a single machine, could now be passed from user to user during a project. The most prominent file (and print) sharing method on the market is Microsoft's NetBIOS implementation.
- A directory service will allow you to search for an item and determine its location. There are many directory services available for certain types of operating systems. For example, Novell uses the NDS, and Microsoft uses AD.
- LDAP is a simple protocol that heterogeneous systems can use to locate information in a database.

Anatomy of an Attack

Overview

Competent attackers will not simply start an attack against a system until they know exactly what they are facing. They will first try to gather as much information as possible at the network, system, and protocol levels. This lesson will discuss the primary reconnaissance attack methods attackers use to gather this information.

Importance

It is important that the information security professional understand how an attack is perpetrated in order to understand proper countermeasures and mitigation techniques.

Objectives

Upon completing this lesson, you will be able to:

- Define port sweeps and ways to mitigate them
- Describe the different types of stealth scans
- Describe ways crackers use to identify the OS running on a target system

Outline

The lesson contains these topics:

- Port Sweeps
- Evasive Sweeps
- OS Identification

Port Sweeps

Attackers use **port sweeps** to determine which services are active on a particular host.

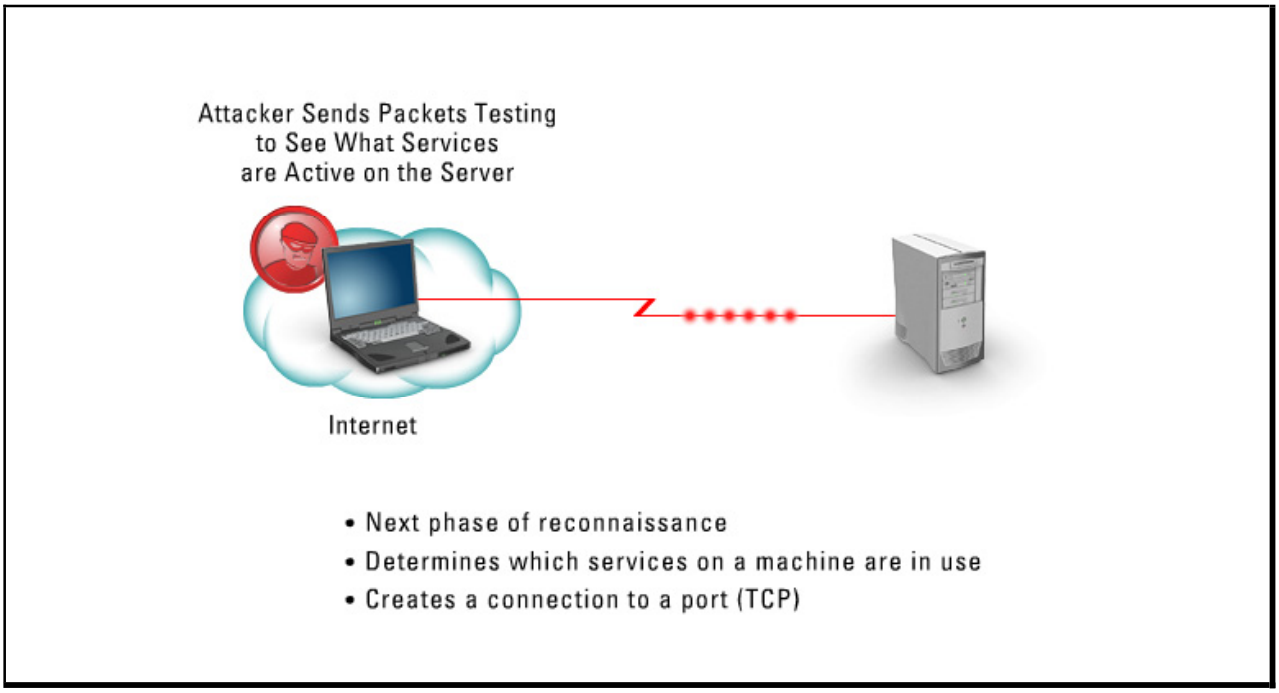

In the evolution of an attack, the cracker first determines which servers are alive and reachable on the network. Crackers perform this task using the ping sweep. Now that the cracker knows which IP addresses he or she can attack, the cracker will perform a port sweep on the systems that are alive. In this way, the cracker can methodically map which services are running on particular hosts. After gaining this information, the cracker will then attempt to attack vulnerabilities in the active service.

Services on hosts are tied to port numbers. There are 65536 possible ports that a single host can be listening on; these ports are divided into three ranges.

Well-Known Ports - Well-known ports are those that have been assigned by the Internet Assigned Numbers Authority (IANA) for specific usages; everyone should know and use these ports. They are in the range zero to 1023. Some examples include:

- FTP (control) TCP port 21
- SSH TCP port 22
- Telnet TCP port 23
- Domain UDP port 53
- www-http TCP port 80

Registered Ports - Registered ports fall in the range 1024 to 49151. An application developer can attempt to obtain a registered number by submitting an application to the IANA.

Dynamic/Private Ports - Private or dynamic ports fall in the range 49152 to 65535. Anyone can use them for private use with their applications.

It is important to remember that only when a service is running and listening on a port can an attack occur on that port. If the service is not running and not listening, a cracker cannot attack it. This brings up one

of the first steps to securing a device, which is to turn off all unnecessary services. If you have a Web server running, it needs to listen on TCP port 80. It should not be listening on any other port as you will introduce the possibility of an attacker gaining access on that port.

Say you have a Web server running and believe you have turned off all other unnecessary services. You can run a port scanning utility against the Web server to verify that only TCP port 80 is listening. If you find other services listening, you will need to research how to disable them.

Note Examples of port scanning utilities include Nmap, Nessus, IPEye, and SuperScan.

Evasive Sweeps

A problem crackers have seen when scanning networks is the fact that their activity can be easily logged and tracked when a connection is made to a particular host. In an attempt to evade detection, crackers have delved into the mysteries of the IP protocol suite and exploited some weaknesses than can help them avoid detection. These evasive scan techniques are called stealth scans and they work by never making a connection, and thus not leaving a 'fingerprint'. A connection is created when a full three-way handshake is completed (SYN->SYN/ACK->ACK).

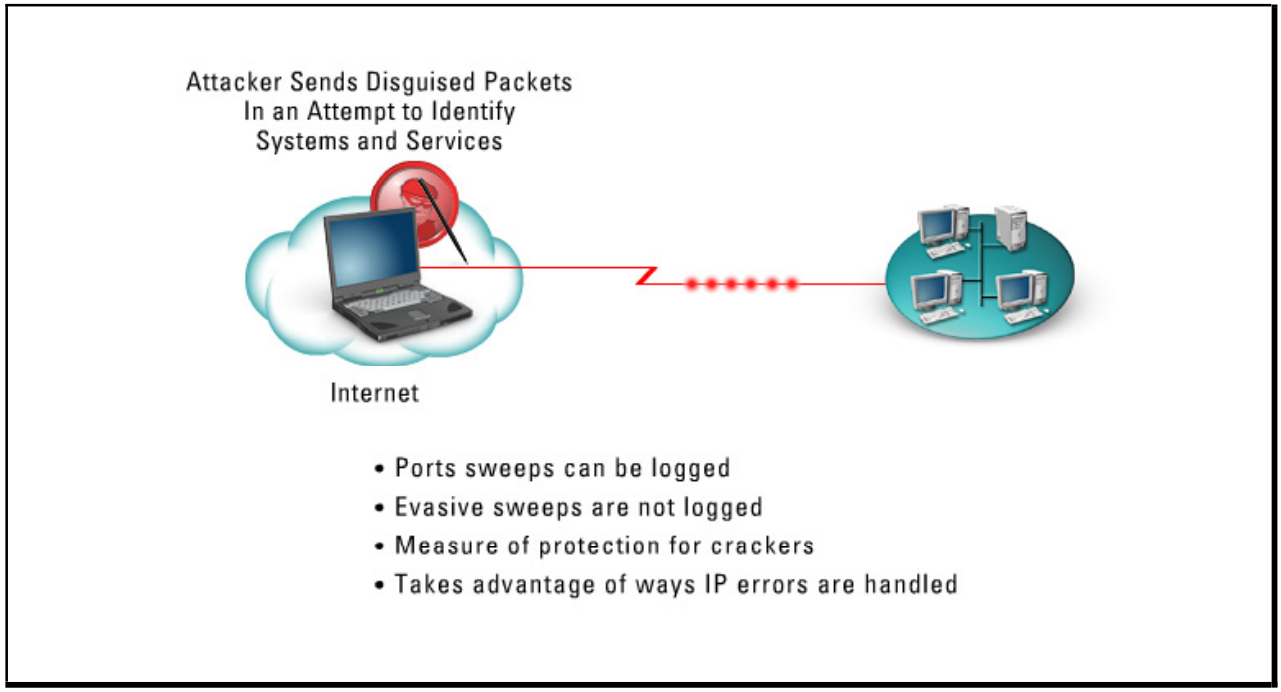

There are many different types of stealth scans. Some scans hide the actual attack in a deluge of 'garbage' packets, and some scans perform attacks over time to hide their trails. The most effective type of stealth scans are those that take advantage of known protocol weaknesses, such as SYN stealth scans or FIN stealth scans. Here the attacker is taking advantage of the way errors are handled in an IP-based host. The cracker will send a packet to the host he or she is trying to gather information about, but the cracker modifies the packet. Imagine the cracker setting the FIN flag in the TCP header. When the receiving host receives the packet, it notices the FIN flag set, which tells the receiving host to close the connection. The receiving host, which has never made a connection to the cracker's workstation, believes this to be an error in communication, so it sends an error message to the cracker that the TCP service port is unavailable; if it is available, it simply ignores the packet. Either way, a connection is never created, which means a log message is never generated, but the cracker now knows whether or not a particular service is running on the target host.

Note	Examples of evasive port scanning utilities include Nmap, IPEye, SuperScan, and AWSPS.

OS Identification

In order for a cracker to effectively generate attacks on a target system, he or she must know which operating system the target is running. Otherwise, the cracker will perform many more attacks on the system, which dramatically increases his or her chances of being detected.

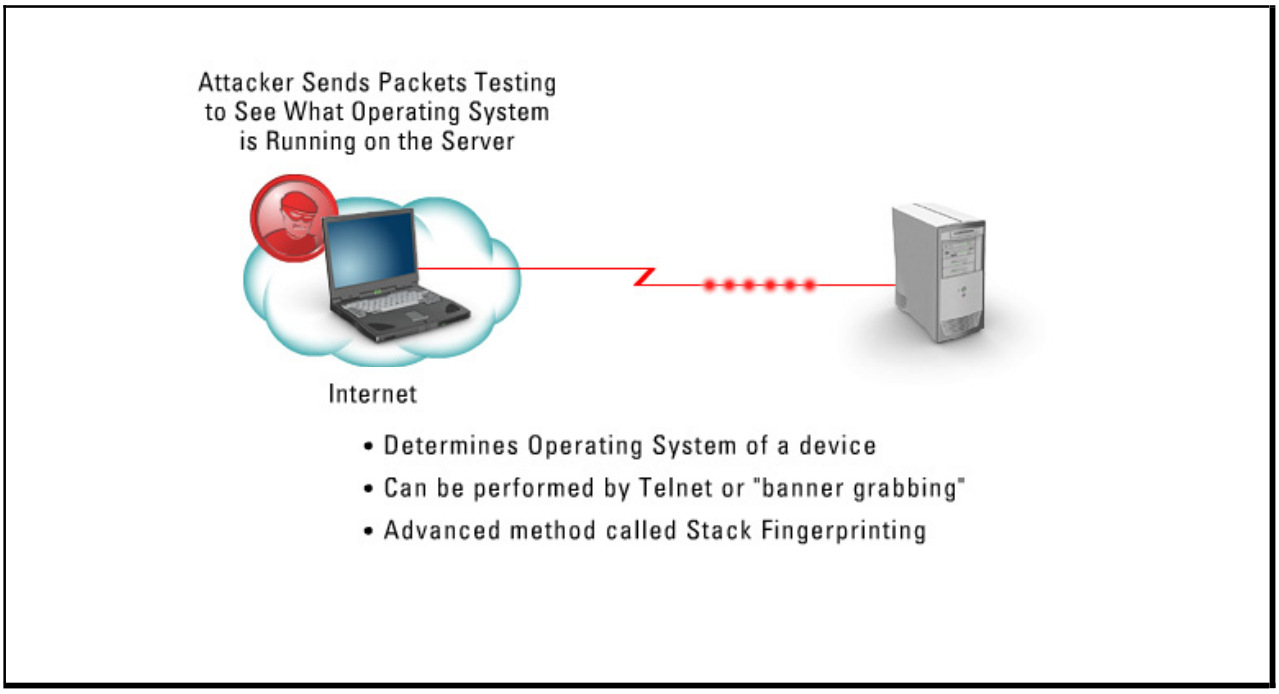

Discovering the operating system running on a target system is often referred to as **enumeration**; this process can enable a cracker to compromise a system in a relatively short amount of time. This is because the more a cracker knows about a target system, the greater his or her chances are of achieving a successful attack. All the cracker would have to do is attempt to match the operating system against a list of known vulnerabilities.

Enumerating an OS in the 'old' days was relatively easy; all you had to do was Telnet to the target and the target would display its OS. If that did not work, you could try banner grabbing. With banner grabbing, you examine the response from certain services like Telnet, FTP, or HTTP. Different operating systems would give different responses, which make it fairly easy to identity which was which.

Today, crackers perform something called active stack fingerprinting. Here the cracker attempts to enumerate an OS by probing its stack. The basis is the same as banner grabbing, except it is performed on the IP stack. This process works because different programmers implemented the IP standards in different fashions. For example, if a cracker sends a TCP packet to a target with the TCP FIN flag set, the standard says the OS should not reply. But, some implementations such as Microsoft Windows NT return a FIN/ACK, while others might send a RST. By actively probing the stack you can very accurately determine which OS the target is running.

Note	Examples of OS identification utilities include Nmap and Queso.

Summary

The key points discussed in this lesson are:

- Attackers use port sweeps to determine which services are active on a particular host. It is important to remember that only when a service is running and listening on a port can an attack occur on that port. If the service is not running and not listening, a cracker cannot attack it.
- There are many different types of stealth scans. Some scans hide the actual attack in a deluge of ‘garbage’ packets, and some scans perform attacks over time to hide their trails. The most effective type of stealth scans are those that take advantage of known protocol weaknesses, such as SYN stealth scans or FIN stealth scans.
- Discovering the operating system running on a target system is often referred to as enumeration; this process can enable a cracker to compromise a system in a relatively short amount of time. This is because the more a cracker knows about a target system, the greater his or her chances are of achieving a successful attack.

Separation of Duties and Responsibilities

Overview

For highly sensitive material or data it would be very easy for a single person who has access to all portions of the network to steal or copy the material. For example, if someone has access to both the accounts payable and accounts receivable, he or she can conceivably modify both ledgers to look like something was either not received or was ordered and received at a much lower cost. To defeat these situations, many companies apply a simple separation of duties. In this way, in order for the crime to be committed, collusion would have to take place. That is, two parties would have to agree to perform this crime against the company, which is much more unlikely to happen.

Importance

It is important that the information security professional understand the role of separation of duties and responsibilities in a company.

Objectives

Upon completing this lesson, you will be able to:

- Define separation of duties and rotation of duties
- Identify guidelines for selecting the appropriate physical, technical, and administrative controls
- Identify general guidelines for the rotation and separation of duties in the workplace
- Define the least privilege principle
- Describe security controls required for media and media storage devices
- Identify security policies for the media library

Outline

The lesson contains these topics:

- Separation and Rotation of Duties
- Selection of Controls

- Rotation and Separation of Duties Guidelines
- Least Privilege
- Media Labels
- Media Library Security

Separation and Rotation of Duties

This administrative control separates a process into component parts, with different users responsible for different parts of the process.

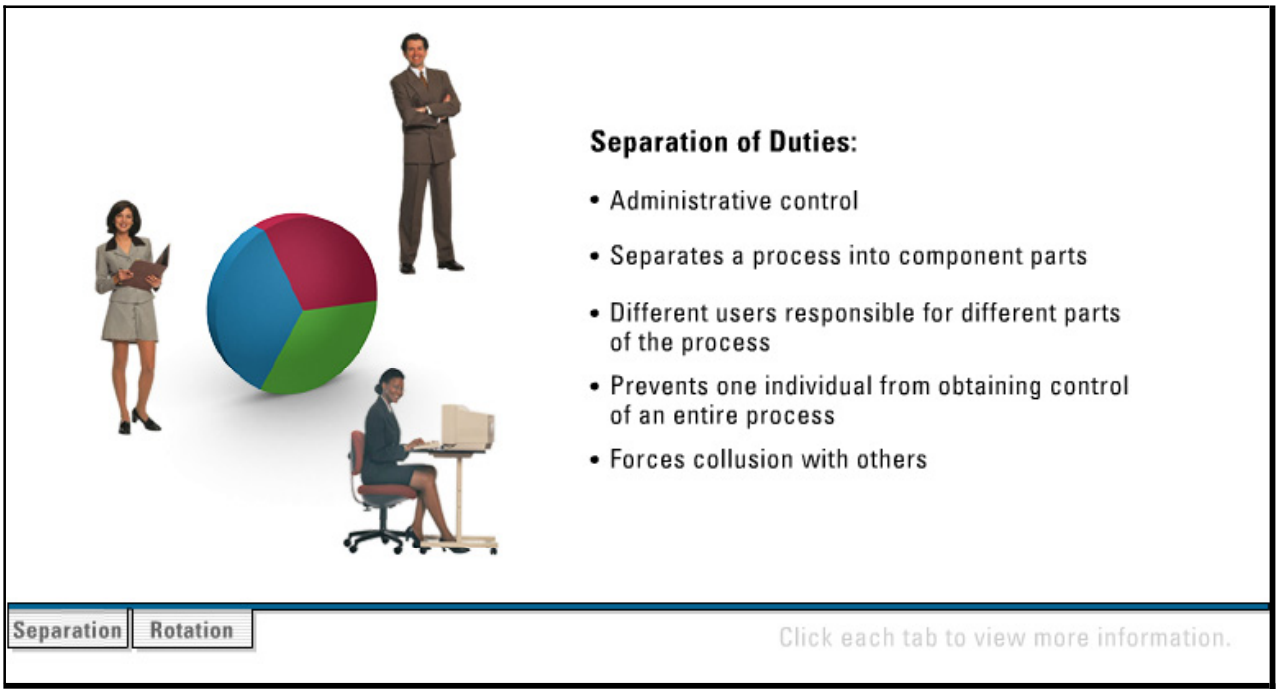

Judicious **separation of duties** prevents one individual from obtaining control of an entire process and forces collusion with others in order to manipulate the process for personal gain.

Rotation of duties, that is, moving employees from one job to another at random intervals, helps deter fraud. As a result of rotating duties, employees are also cross-trained to perform each other's functions in case of illness, vacation, or termination. It is also much easier for fraud to be identified if different people are performing checks and balances as an active part of their rotational duty.

Selection of Controls

You should review your organization's security policy to determine the confidentiality, integrity, and availability needs of the organization. You can then select the appropriate physical, technical, and administrative controls to provide the required level of information protection, as stated in the security policy.

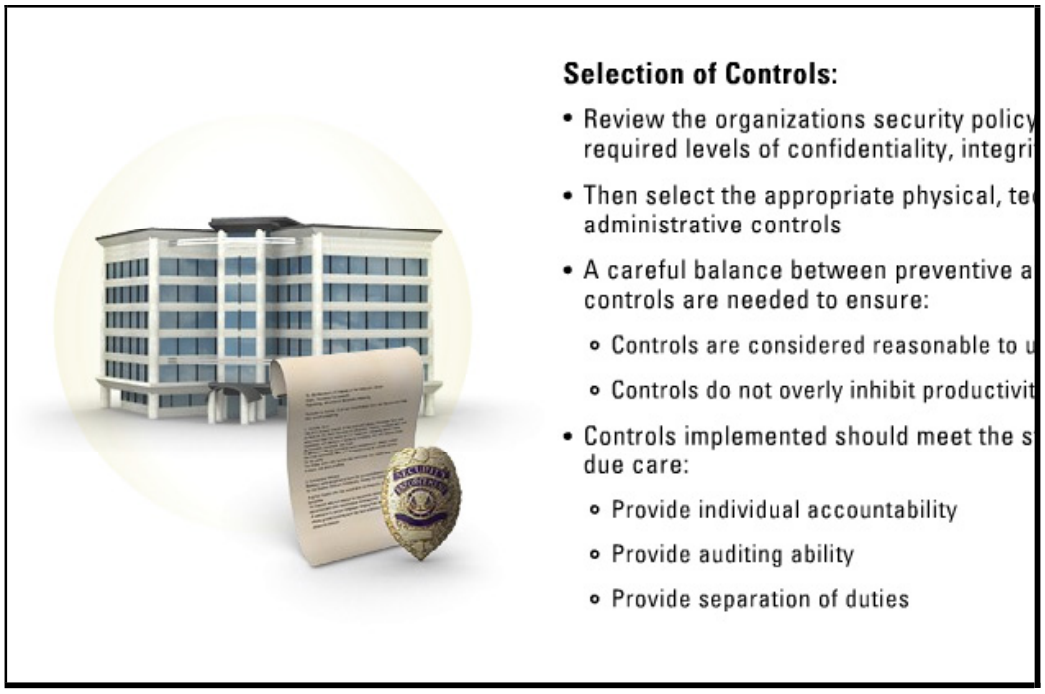

A careful balance between preventive and detective control measures is needed to ensure that users consider the security controls reasonable and to ensure that the controls do not overly inhibit productivity. You can identify the combination of physical, technical, and administrative controls best suited for your specific computing environment by completing a quantitative risk analysis. Because this is usually an expensive, tedious, and subjective process, however, an alternative approach—referred to as meeting the standard of due care—is often used.

Note	Controls that meet a standard of due care are those that would be considered prudent by most organizations in similar circumstances or environments.

Controls that meet the standard of due care generally are readily available for a reasonable cost and support the security policy of the organization; they include, at the least, controls that provide individual accountability, audit ability, and separation of duties.

Rotation and Separation of Duties Guidelines

This topic discusses the general guidelines that you should implement for rotation and separation of duties in the workplace.

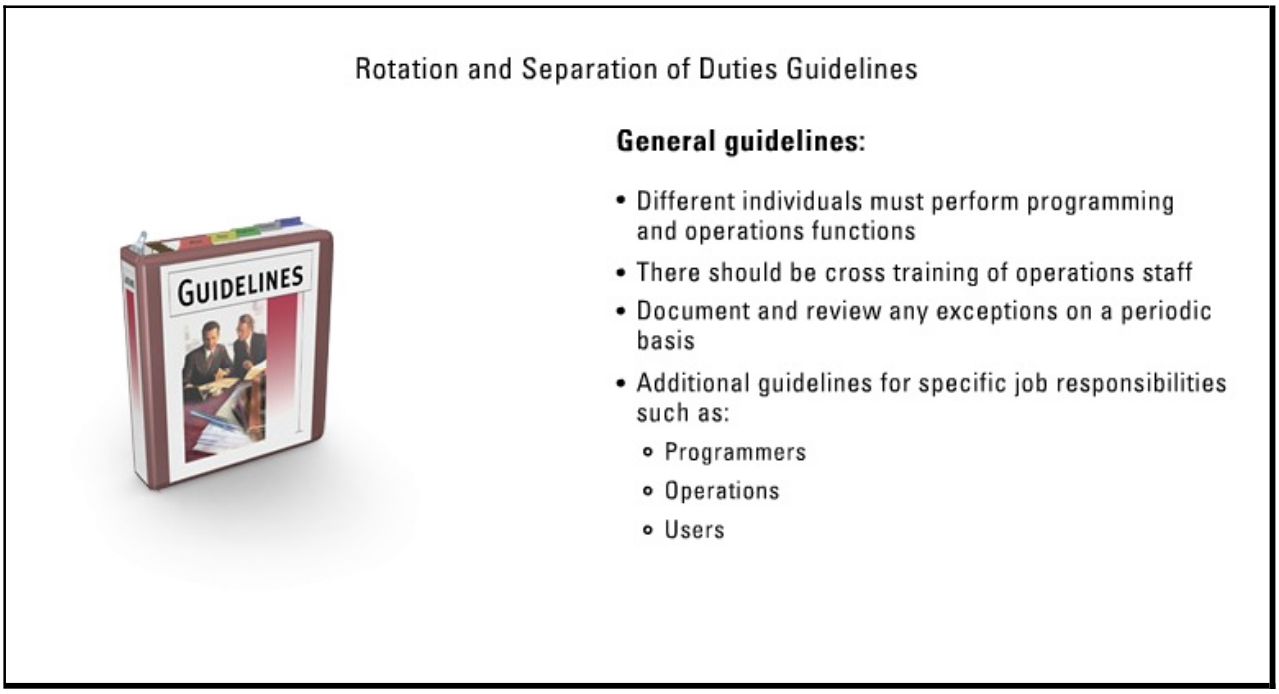

General guidelines:

- Different individuals must perform programming and operations functions.
- There should be cross training of operations staff to provide depth and backup, and to reduce individual dependence.
- The following guidelines regard separation of duties for various groups of employees. You should document and review any exceptions to these guidelines on a periodic basis for justification and risk analysis purposes.

Programmers:

- Programmers should not execute jobs in a production mode.
- Programmers should not control any transfers between programmer development libraries and production libraries.
- Programmers should not have update capabilities within any production application.

Operators:

- Operators should not have the ability to make changes to production application or system software libraries.
- Operators should not perform balancing activities, except those necessary for run-to-run controls.
- Operators should not have the ability to make changes to job control language (JCL) of scheduled jobs without proper notification and authorization.
- Operators should execute only those jobs/programs scheduled through the established procedures.

- Operators should not execute (outside of standard production processing) data or software-modifying system utilities without proper authorization and dual control.
- Operators should not override internal tape labels without supervisory approval.

Users:

- Data entry personnel should not prepare source documents for input.
- Someone, other than the input operator, should verify all data input, unless programmatically verified.
- The same person should not perform input and output duties.
- The same person should not post and balance general ledger and other sensitive entries.
- The person who prepared the original transaction should not review rejects or non-reads for reentry.
- Master file and other sensitive transaction changes should be under dual control.

Least Privilege

The idea behind the principle of **least privilege** is that you give a process only the authority it needs to accomplish its job. For example, consider a program needing to write a log to disk.

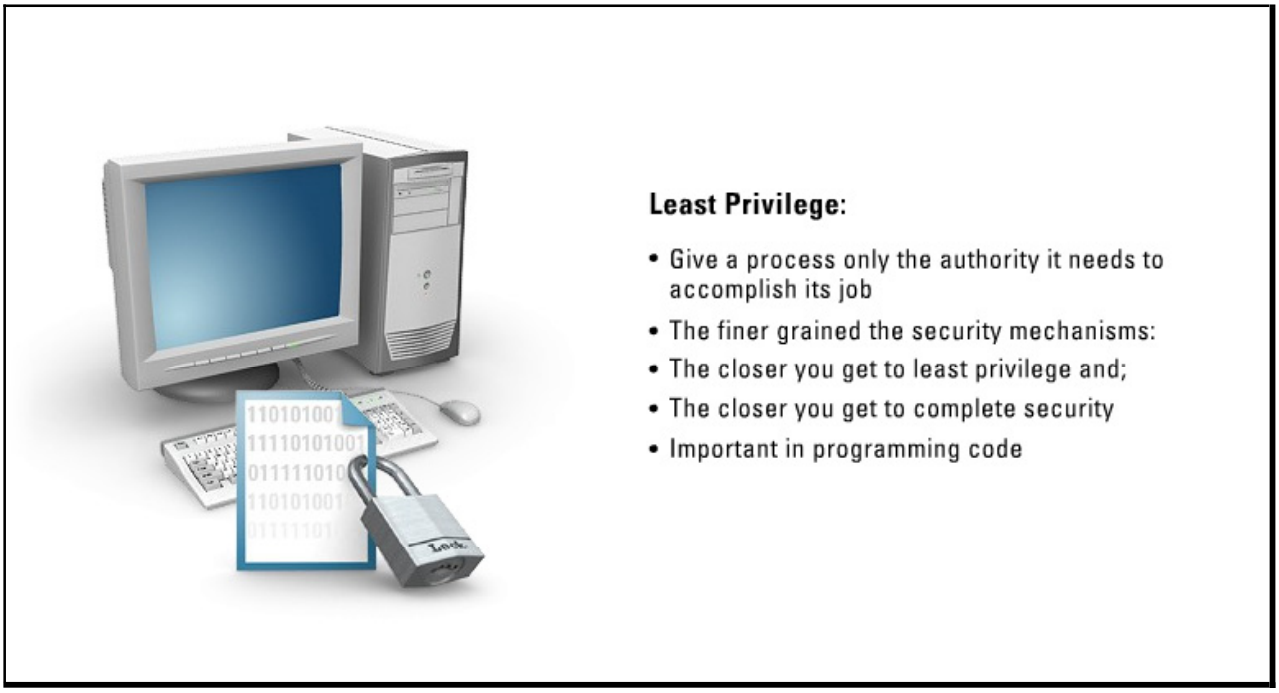

There are several possible privileges that could allow this task, including:

- Raw disk access
- Read/write rights to a single directory
- Read/write rights to a supplied file
- Rights to a supplied output stream

The finer grained the security mechanism, the closer you can get to the least privilege, and therefore the closer you get to complete security.

Note	Completely secure systems only exist theoretically, but practically speaking, you can get pretty close.

The principle of least privilege is important in programming code. When you execute a program, it will be loaded into memory and begin running at a privilege level assigned by the programmer. If the program is running at the highest system level and an attacker cracks the program, the attacker essentially has system level privileges. Now, if the same program was running at basic user privilege, it would be very difficult for a malicious user to do damage with it.

Media Labels

Media and media storage devices require their own set of security controls to ensure they are properly preserved and that the integrity, confidentiality, and availability of the data stored on them are not compromised.

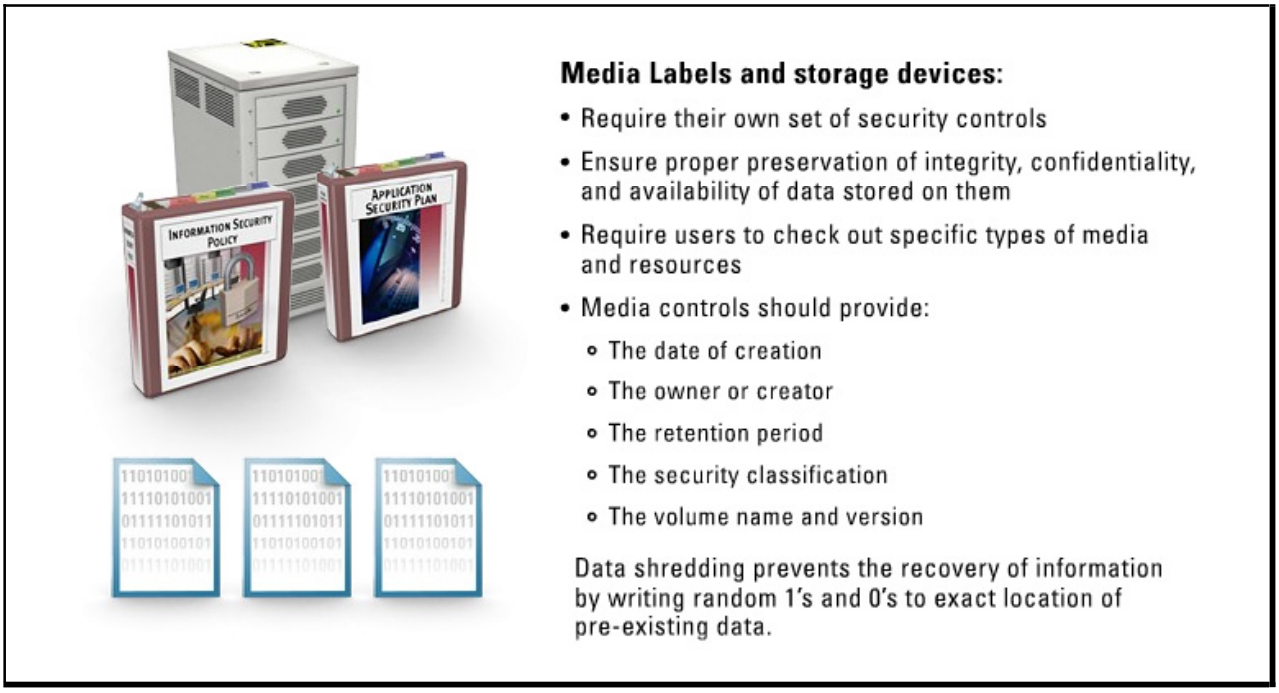

If the media library has grown to the extent that a media librarian is required, then you should require users to check out specific types of media and resources. The media controls in this case would require media labeling that should provide the following:

- The date of creation
- The owner or creator
- The retention period (how long the media is good for)
- The security classification
- The volume name and version

You should clearly mark and log all media, you should evaluate the integrity of the media, and when destruction is required, you should do it in a secure manner. When media is cleared or zeroed, it is said to be sanitized. There are different methods of sanitation including overwriting, degaussing, and destruction.

It is important to understand that even though data might be erased and zeroed (a series of zeros have been written over the data), it is quite possible for the original data to still be readable by sophisticated equipment and forensic software. This residual reading of data remanence can be of concern to high security installations.

Data shredding prevents the recovery of this information by writing binary 1's or 0's and/or random bit patterns to the exact location where pre-existing data were stored. If this process occurs a minimum of three times (three iterations), it would be very difficult for even the most sophisticated forensic equipment to read the original information. At seven iterations, it is considered impossible.

Media Library Security

You must secure storage areas that hold the media library in a manner consistent with the security policies of your organization.

Media Library Security:

- Secure any storage area that holds the media library according to security policy
- Paper based media libraries require additional precautions:
  - Increased threat of fire
  - Combustibles
- Access controls on the media library should:
  - Identify media sensitivity
  - Segregate classified media
  - User check-out privilege level
- Environmental controls to provide for adequate:
  - Temperature
  - Humidity

If the media library contains large amounts of paper material, you must allow for additional precautions including:

- Increased threat of fire
- Combustibles

Access controls to the media library should be in place including:

- Media sensitivity
- Segregation of classified media
- User check-out privilege level

Finally, environmental controls must be in place to control the temperature and humidity inside the storage area.

Summary

The key points discussed in this lesson are:

- Judicious separation of duties prevents one individual from obtaining control of an entire process and forces collusion with others in order to manipulate the process for personal gain. Rotation of duties, that is, moving employees from one job to another at random intervals, helps deter fraud.
- You should review your organization's security policy to determine the confidentiality, integrity, and availability needs of the organization. You can then select the appropriate physical, technical, and administrative controls to provide the required level of information protection, as stated in the security policy.
- You should follow several guidelines when you implement rotation and separation of duties in the workplace. You should document and review any exceptions to these guidelines on a periodic basis for justification and risk analysis purposes.
- The idea behind the principle of least privilege is that you give a process only the authority it needs to accomplish its job.
- Media and media storage devices require their own set of security controls to ensure they are properly preserved and that the integrity, confidentiality, and availability of the data stored on them are not compromised.
- You must secure storage areas that hold the media library in a manner consistent with the security policies of your organization.

Business Continuity Planning

Overview

Having the business continue in the event of a major catastrophe is essential to the livelihood of many enterprises today. Having business contingency plans (BCPs) prior to any unforeseen accident that devastates an enterprise's data and resources can alleviate concerns of the business surviving.

Objectives

Upon completing this module, you will be able to:

- Describe the process for creating a business continuity plan
- Identify strategies for recovering crucial business systems and data
- Identify the primary recovery strategies that you should follow to ensure business continuity
- Describe how assurance and trust of the business continuity plan is developed

Outline

The module contains these lessons:

- Business Continuity Plan Process
- Recovery
- Primary Strategies
- Assurance and Trust

Business Continuity Plan Process

Overview

This lesson will discuss the scoping and organization of a business continuity plan. This lesson will also discuss the business impact of establishing a business continuity plan in the enterprise.

Importance

It is important for the information security specialist to understand the stages of the business continuity plan process.

Objectives

Upon completing this lesson, you will be able to:

- List the elements of project scoping and planning
- Describe management's role in business continuity planning
- Identify the responsibilities of the business continuity plan team
- Identify the steps involved in conducting a business impact analysis

Outline

This lesson contains these topics:

- Project Scoping and Planning
- Role of Management
- BCP Team Responsibilities
- Business Impact Analysis

Project Scoping and Planning

Project scoping and planning mark the beginning of the business continuity plan (BCP) process.

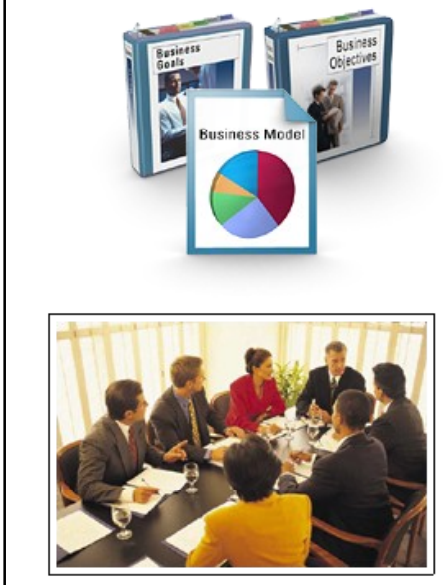

Project Scoping and Planning

Elements of project scoping and planning include:

- Scope and plan initiation
- Business impact analysis
- BCP development
- Plan approval and implementation

The BCP committee and senior management create the BCP.

The BCP committee is responsible for:

- Creating
- Implementing
- Testing

Senior management is responsible for all four phases.

Prime elements of project scoping and planning include:

- Scope and plan initiation
- Business impact analysis
- BCP development
- Plan approval and implementation

Together, the BCP committee and senior management create the BCP.

The BCP committee is responsible for creating, implementing, and testing the plan. This committee is made up of the security administrator and representatives from senior management, all functional business units, and the information systems department.

Senior management is ultimately responsible for all four phases of the plan.

Role of Management

Management must understand what the real risks are to the company, the consequences of those risks, and the potential loss values for each risk. Without fully understanding the ramifications of *not* having a contingency plan, management might only give 'lip-service' to disaster recovery and contingency planning.

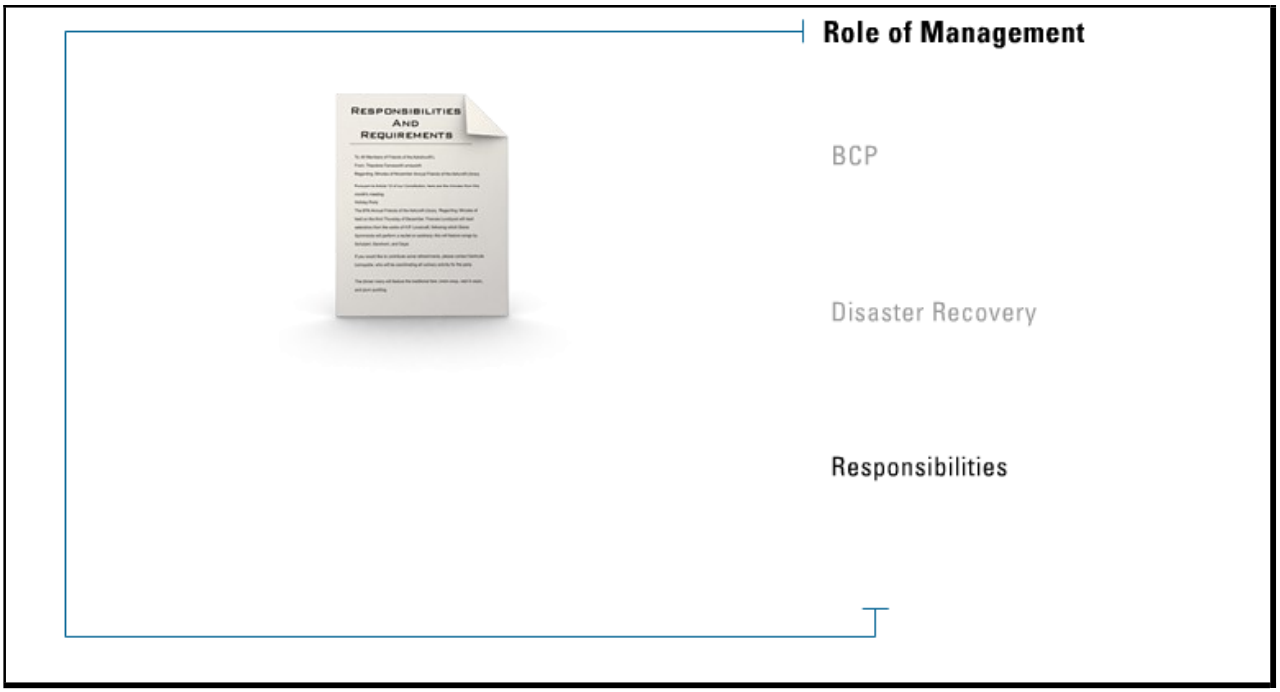

Without the support of management, the necessary monies and time will not be devoted to the BCP process, which is necessary to provide adequate protection in case of emergency. Executives may be held liable under various legal and regulatory statutes and could be sued by stockholders and customers if they do not practice due diligence and fulfill their responsibilities when it comes to disaster recovery.

Management has the following responsibilities:

- Full commitment
- Policy and goal setting
- Providing the necessary funds and company resources
- Taking responsibility for the outcome of the BCP process
- Appointing a team for the process

BCP Team Responsibilities

Once management has decided a contingency plan is necessary, they must appoint a specific team of people to perform the necessary task of developing the plan.

BCP Team Responsibilities

- BCP Team works with management to:
  - Develop the ultimate goals of the plan
  - Identify the critical business service
  - Identify the priorities of personnel and departments
- The BCP Team should include individuals who are familiar with every department
- The BCP Team has the following responsibilities:
  - Identify all regulatory and legal requirements
  - Identify all possible threats and risks
  - Provide an estimate on possible loss potential
  - Perform a Business Impact Analysis
  - Outline which processes must be operational before any other
  - Develop procedures to resume business after a disaster
  - Assign tasks to individuals that must be performed during a disaster
  - Document and train employees

This team will then work with management to develop the ultimate goals of the plan, identify the critical business services that must be dealt with first during a disaster, and identify the priorities of personnel and various departments. With management's support, the team will now have the goals, objectives, and priorities necessary to create and shape the contingency program.

The BCP team should include individuals who are familiar with every department or representatives from each department.

Note	It may be necessary to contact outside vendors when the creation of off-site solutions is required. If so, the BCP team must document and outline all agreements in the final solution.

The BPC team has the following responsibilities:

- Identify all regulatory and legal requirements that the team must meet
- Identify all possible threats and risks
- Provide an estimate on possible loss potential
- Perform a business impact analysis (BIA)
- Outline which systems, processes, and departments must be operational before any other
- Develop procedures necessary to resume business after a disaster
- Assign tasks to individuals that they must perform during a disaster
- Document and train employees

Business Impact Analysis

A business impact analysis (BIA) identifies the potential impact of a data processing outage.

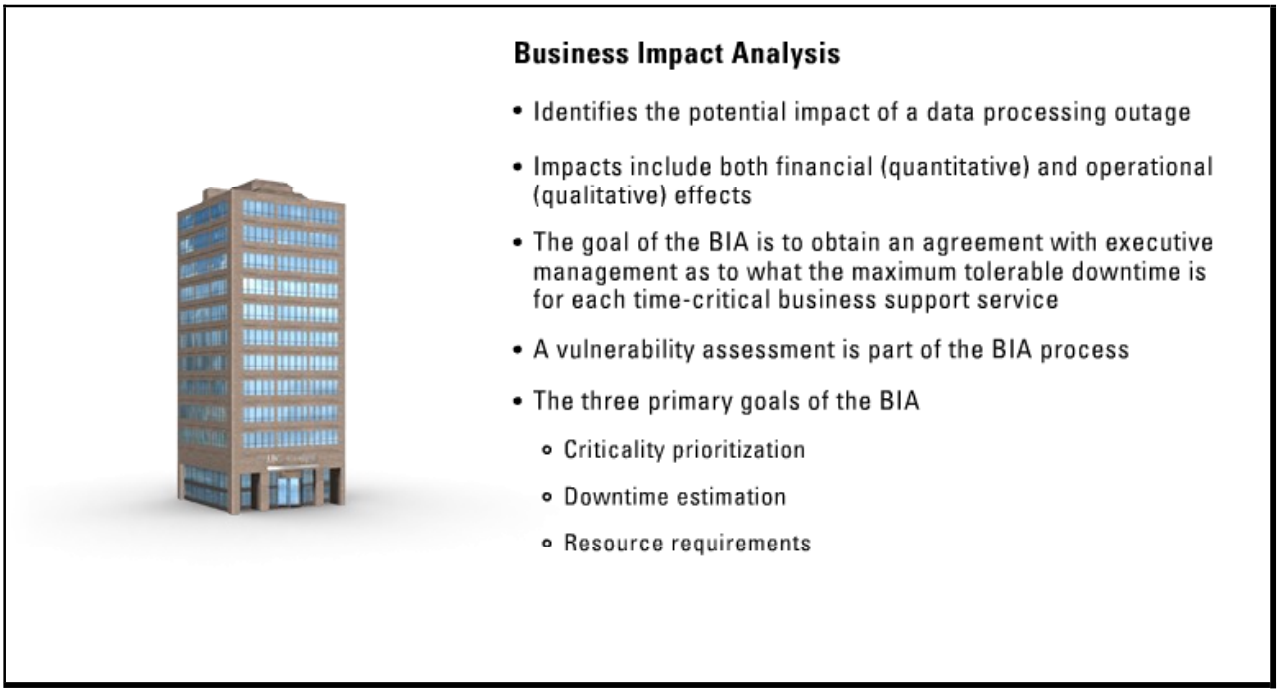

The impacts include both financial effects (quantitative) and operational effects (qualitative, such as the inability to respond to a customer). The goal of the BIA is to obtain an agreement with executive management as to what the maximum tolerable downtime (MTD) is for each time-critical business support service. MTD is also referred to as the maximum allowable outage (MAO).

A vulnerability assessment is often a part of the BIA process. It identifies the company's critical systems needed for survival and estimates the outage time that the company can tolerate as a result of a disaster or disruption.

The three primary goals of a BIA:

- Criticality Prioritization - The BIA must identify and prioritize every critical business unit process and evaluate the impact of a disruptive event.
- Downtime Estimation - The BIA must estimate the MTD that the business can tolerate and still remain a viable company.
- Resource Requirements - The BIA must identify the resource requirements for the critical processes, with the most time-sensitive processes receiving the most resource allocation.

The steps involved in conducting a BIA include the following:

1. Determine information-gathering techniques: Techniques include electronic or paper-based surveys and questionnaires, one-on-one interviews, group interviews, workgroups, videoconference meetings, etc.
2. Select interviewees: Identify which management and staff members within each business unit are necessary in determining all critical business processes.

3.Customize questionnaire: There is no set of standard BIA questions. Questions must be created based on the organization's traits and should identify all time-critical business processes and their support services.

4.Analyze information: Analyze information gathered in Step 3. Document results based on each business unit and create a summary that contains the information collected during the interview process.

5.Determine time-critical business functions: From the analysis performed in Step 4, identify all time-critical business functions.

6.Determine MTD: The MTD is the period of time a business function or process can remain interrupted before its ability to recover becomes questionable. Business processes requiring shorter time periods are considered more time-critical.

7.Prioritize critical business functions based on MTDs: When you identify all key business processes, you can then rank them to identify those of most critical importance.

8.Document and prepare a report for recovery recommendations: Create a report based on the result of the BIA for executive-level management. Once management approves the BIA report, you can continue to the next phase, which is developing recovery strategies.

Summary

The key points discussed in this lesson are:

- Project scoping and planning mark the beginning of the BCP process.
- Without the support of management, the necessary monies and time will not be devoted to the BCP process, which is necessary to provide adequate protection in case of emergency. Executives may be held liable under various legal and regulatory statutes and could be sued by stockholders and customers if they do not practice due diligence and fulfill their responsibilities when it comes to disaster recovery.
- Once management has decided a contingency plan is necessary, they must appoint a specific team of people to perform the necessary task of developing the plan. With management's support, the team will now have the goals, objectives, and priorities necessary to create and shape the contingency program.
- A BIA identifies the potential impact of a data processing outage. The goal of the BIA is to obtain an agreement with executive management as to what the MTD is for each time-critical business support service.

Recovery

Overview

Before disaster strikes and systems and resources are lost, you must plan and institute a recovery strategy. This lesson will discuss strategies that you can use to recover crucial business systems and data when a devastating loss occurs.

Importance

It is important for the information security specialist to understand the recovery strategies that he or she can implement when a critical loss of systems occurs.

Objectives

Upon completing this lesson, you will be able to:

- Explain the recovery strategy development process
- Identify the procedures of facility and supply recovery strategies
- List the elements of user-group recovery strategies
- List the elements of technical recovery strategies

Outline

The lesson contains these topics:

- Recovery Strategies
- Facility and Supply Recovery Strategies
- User Recovery Strategies
- Technical Recovery Strategies

Recovery Strategies

In the recovery strategies phase, you should identify all business recovery strategies. These recovery strategies consist of a set of predefined and management-approved actions that would be implemented in response to an unacceptable business interruption.

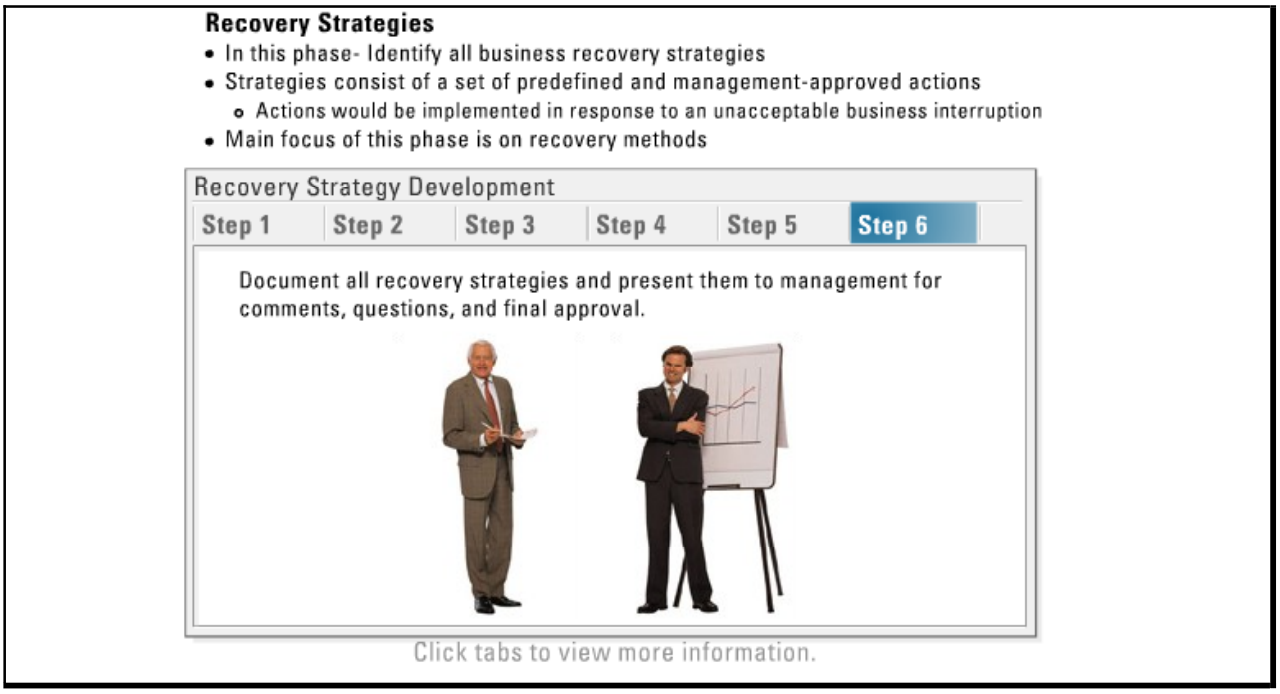

The main focus in the phase is on recovery methods that you will use to meet the timeframes established for the operation of the critical business functions. To develop recovery strategies, follow these steps:

1. Document each alternative recovery strategy along with its associated cost.
2. Obtain quotes for any necessary outside services, such as requests for proposals (RFPs) from outside vendors.
3. Create written agreements, which include definitions, terms, responsibilities, recovery requirements, costs, the technical approach to reach objectives, payment terms, and the deadline response data.
4. Evaluate all risk reduction and resumption strategies based on a full loss of the facility.
5. Identify all risk reduction measures and revise resumption priorities and timeframes.
6. Document all recovery strategies and present them to management for comments, questions, and final approval.

Facility and Supply Recovery Strategies

The purpose of the facility and supply recovery phase is to identify all recovery procedures for any alternate facility, including space required, security needed, fire protection requirements, infrastructure requirements, utility requirements, supply requirements, and environmental requirements.

Facility and Supply Recovery Strategies:

- In this phase- Identify all recovery procedures for any alternate facility including:

- Space required
- Security needed
- Fire protection requirements
- Infrastructure requirements
- Utility requirements
- Supply requirements
- Environmental requirements

The following procedures are part of the recovery plan:

- Determine minimum space required for work areas and conference rooms of critical business units.
- Determine minimum space required for all less critical resources.
- Determine security requirements at the recovery site.
- Determine fire protection requirements.
- Determine business furnishings and office equipment requirements.
- Determine infrastructure requirements such as construction needs, underground power and telephone lines, etc.
- Determine utility and environmental needs such as heat, ventilation, and air conditioning (HVAC), power, water, generators, uninterruptible power supplies (UPSs), etc.
- Determine office supplies required, such as forms, notepads, pens, staplers, etc.

User Recovery Strategies

User recovery strategies focus on manual procedures, vital records, and restoration procedures.

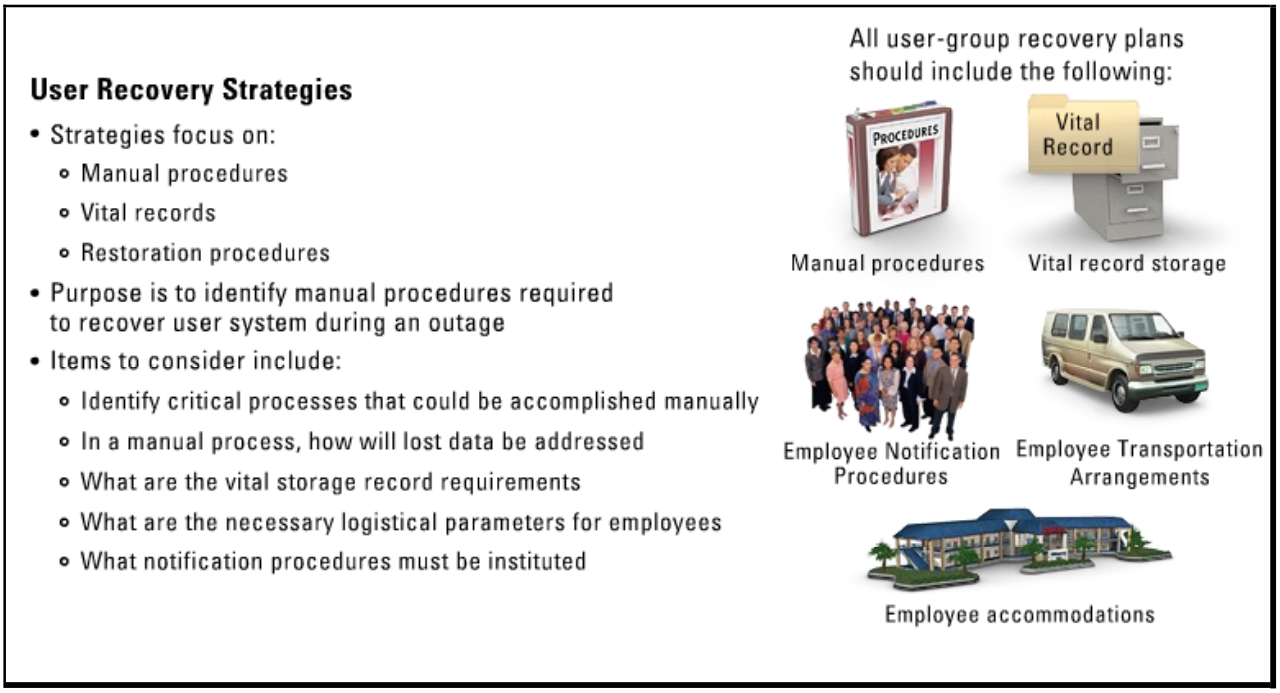

The purpose of a user recovery strategy is to identify manual procedures required to recover a user system during an outage. It is critical that recovery strategies establish methods to implement the process and maintain the records so that information can be easily and accurately updated to the electronic format when service is restored. Items that you should consider include the following:

- Identify critical processes that could be accomplished manually (e.g., writing checks instead of printing them).
- If a manual process can be implemented, how will lost data or transactions be addressed? What type of paper trail must be established?
- What are the vital storage record requirements?
- What are necessary logistical parameters for employees, such as transportation, housing, meals, and so on?
- What notification procedures must be instituted?

All user-group recovery plans should include the following:

- Manual procedures
- Vital record storage
- Employee notification procedures
- Employee transportation arrangements
- Employee accommodations

Technical Recovery Strategies

Technical recovery strategies are aimed at data center and network infrastructure recovery. This recovery facility is necessary when the primary business location is unavailable.

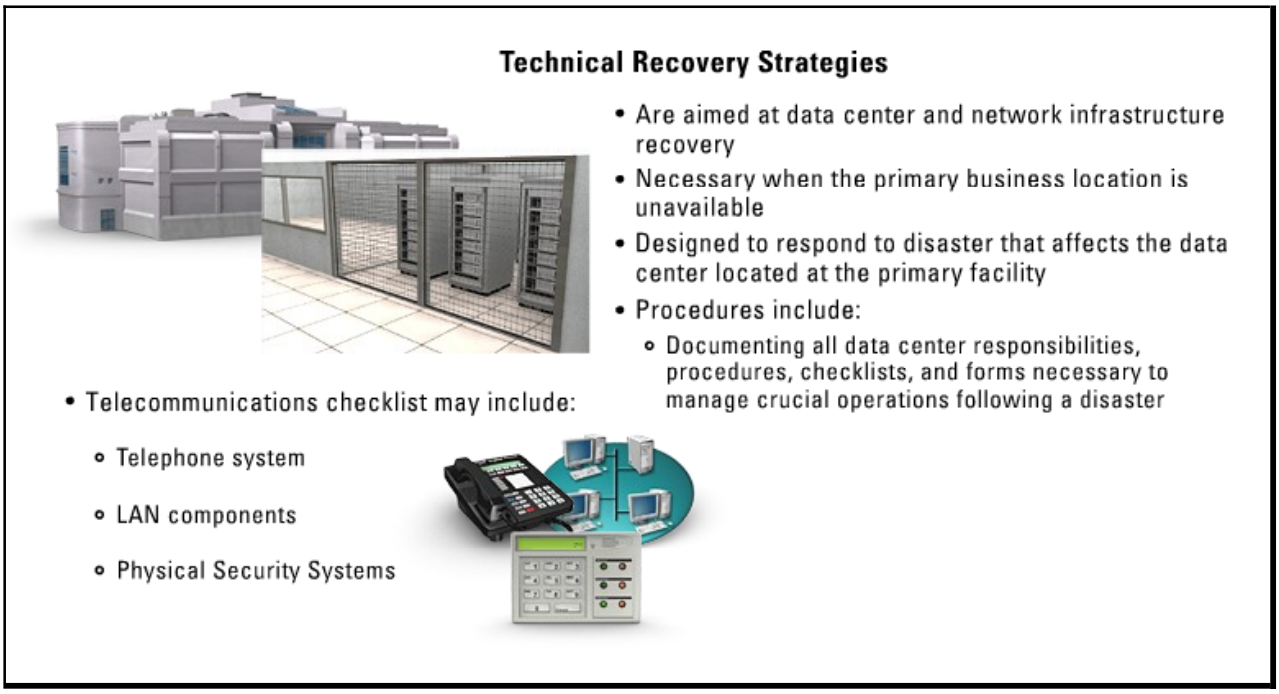

Data center recovery plans are designed to respond to a disaster that affects the data center located at the primary facility. You will need to document all data center responsibilities, procedures, checklists, and forms that are necessary to manage and control the recovery of crucial computer operations following a disaster.

Network and data communication recovery plans are designed to recover a network system that is separate from the main data center location, such as local area networks (LANs). You first need to determine the network support requirements, such as the hardware requirements. This includes computers and other LAN equipment, such as peripherals and cabling. You will also need to list any hardware requirements for connecting input and output devices to the recovery systems, such as Channel Service Units/Digital Service Units (CSUs/DSUs), switches, routers, firewalls, and so on.

Telecommunications recovery plans are essential, as communication to and from the recovery site is usually the core requirement of any business. A checklist for telecommunications may include the following:

- **Telephone Systems** - Key telephone lines, private branch exchanges (PBXs), voicemail systems, fax machines, paging systems, and cell phones
- **LAN Components** - Computer hardware, cable system, power supplies, modems, switches and routers, personal computers, teleconferencing equipment, and test and measurement equipment
- **Physical Security Systems** - Closed-circuit television (CCTV), motion detectors, and lighting controls

Summary

The key points discussed in this lesson are:

- In the recovery strategies phase, you should identify all business recovery strategies. These recovery strategies consist of a set of predefined and management-approved actions that would be implemented in response to an unacceptable business interruption.
- The purpose of the facility and supply recovery phase is to identify all recovery procedures for any alternate facility, including space required, security needed, fire protection requirements, infrastructure requirements, utility requirements, supply requirements, and environmental requirements.
- The purpose of a user recovery strategy is to identify manual procedures required to recover a user system during an outage.
- Technical recovery strategies are aimed at data center and network infrastructure recovery. This recovery facility is necessary when the primary business location is unavailable.

Primary Strategies

Overview

The business continuity plan must implement an emergency response procedure that you will enact when a disaster has happened to your enterprise. This lesson will discuss the logistics, procedures, and primary recovery strategies that you should follow to ensure business continuity.

Importance

It is important for the information security specialist to understand the primary strategies of disaster recovery.

Objectives

Upon completing this lesson, you will be able to:

- List the four primary technical recovery strategies
- Describe the logistics of service level agreements
- Describe the main categories of subscription services
- Identify the types of data backups
- Identify the three types of electronic vaulting

Outline

This lesson contains these topics:

- Primary Technical Recovery Strategies
- Service Level Agreements
- Subscription Services
- Data Recovery Strategies
- Electronic Vaulting

Primary Technical Recovery Strategies

This topic lists the primary technical recovery strategies.

Primary Technical Recovery Strategies

There are four primary technical recovery strategies:

- Subscription services
- Reciprocal or mutual aid agreements
- Multiple processing centers
- Service bureaus

There are four primary technical recovery strategies:

- Subscription services
- Reciprocal or mutual aid agreements
- Multiple processing centers
- Service bureaus

Service Level Agreements

Service level agreements (SLAs) are fundamental to business continuity. They define your minimum levels of availability from key suppliers, and often determine what actions will be taken in the event of serious disruption. Consequently, they require full consideration and attention and must be constructed extremely carefully.

Service Level Agreements

- Are fundamental to business continuity
- Define minimum levels of availability from key suppliers
- Often determine what actions will be taken in the event of a serious disruption
- The SLA agreement with supplier must be:
  - Wide in scope
  - Cover all key aspects of service
  - Fully embrace such issues as:
    - Problem management
    - Compensation
    - Warranties
    - Resolutions of disputes
    - Legal compliance

The SLA you reach with a provider must be wide in scope, covering all key aspects of the service. Typically, it will fully embrace such issues as problem management, compensation (often essential in terms of motivation), warranties and remedies, resolution of disputes, and legal compliance. It essentially frames the relationship, and determines the major responsibilities, both in times of normal operation and during an emergency situation.

Subscription Services

Subscription services are alternate facility sites you can use for recovery.

Subscription services are usually categorized as hot, warm, cold, mirror, or mobile sites.

- **Hot Sites** - Hot sites are fully configured sites with complete hardware and software provided by the service.
  - Advantages:
    - Can be operational in a very short period of time
    - This service assures exclusive use of the facility
    - The contract for this service usually includes test time for compatibility checks
    - Will usually have multiple sites available in case another site is not available
  - Disadvantages:
    - The most expensive to maintain
    - Possible contention for the hot site if a regional disaster occurs
    - Limited choices available if special or unusual hardware is required
- **Warm Sites** - Warm sites are similar to hot sites, except the expensive equipment is not available onsite. This type of site is usually available in a few hours after the equipment arrives.
  - Advantages:
    - Availability is usually assured for longer periods of time
    - Various convenient locations are usually available
    - Usually available for exclusive use of the organization
    - Less expensive than a hot site

        - Practical for organizations that use special or unusual hardware
    - Disadvantages:
        - Operational testing not available
        - Necessary resources might not be immediately available
        - More expensive than cold sites or in-house recovery sites
- **Cold Sites** - Cold sites do not contain any technical equipment or supplemental resources, except air conditioning, power, telecommunications, raised floors, and so on. All necessary computing equipment must be installed and tested before the facility can support the critical business functions.
    - Advantages:
        - Available for longer periods of time
        - Site can be in various locations
        - Site is for exclusive use of the organization
        - Less expensive than other options
        - Practical for organizations that use special or unusual hardware
    - Disadvantages:
        - Operational testing not available
        - Necessary resources might not be immediately available
        - Costs are more expensive than in-house facilities
- **Mirror Sites** - Mirror sites contain all necessary communication lines, and appropriate hardware that is fully operational and processes each transaction along with the primary site. In most situations, a high-speed link exists between the primary site and the mirror site to support the mirroring operations. This site can also be described as a full redundancy site.
    - Advantages:
        - No loss of data if a disaster occurs at the primary site
        - Site is instantly available for processing
    - Disadvantages:
        - Expensive to maintain
- **Mobile Sites** - A mobile site is usually a trailer complete with all necessary computing hardware than can be set up in a subscriber's parking lot and linked by a trailer sleeve for connectivity.
    - Advantages:
        - Travel is minimized for employees
        - Allows for a decentralized organization to engage a single vendor to service the entire organization
    - Disadvantages:
        - Operational testing might not be available
        - Necessary resources might not be immediately available
        - Costs are more expensive than in-house facilities

Data Recovery Strategies

The data recovery process is the most critical recovery process, as data processing capabilities are the life blood of any organization.

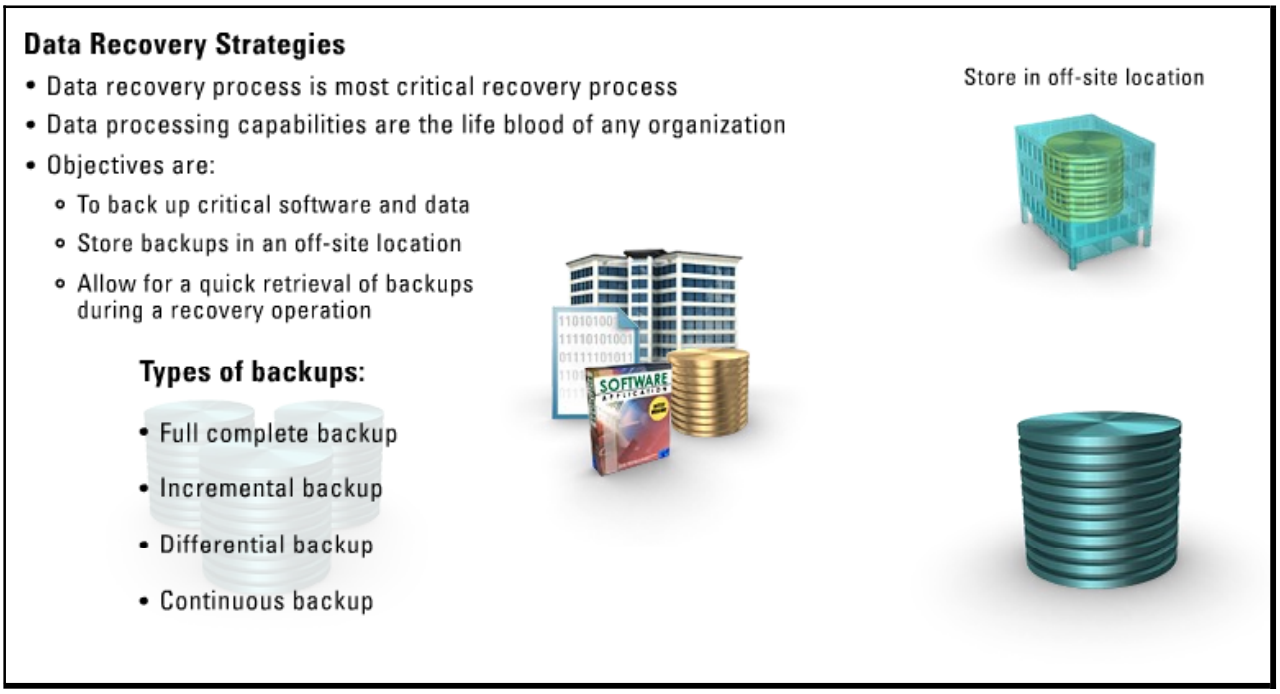

The objectives of data recovery are to back up critical software and data, store the backups in an off-site location, and allow for a quick retrieval of backups during the recovery operation. You should include operating system software, utilities, application software, production data, databases, and associated transaction logs in the backup. Backup and off-site storage is the process of backing up organizational data and required software, and securing the backup media. The types of backups include:

- **Full Complete Backup** - A complete drive level backup (system, application data, etc.)
- **Incremental Backup** - Includes all the files that have changed since the last backup was performed
- **Differential Backup** - Includes only those files that have changed since the previous full backup was performed
- **Continuous Backup** - Continuous or progressive backups that keep a database of all existing files and their locations on the backup media

Electronic Vaulting

The electronic bulk transfer of backup data to an alternate site is called **electronic vaulting**.

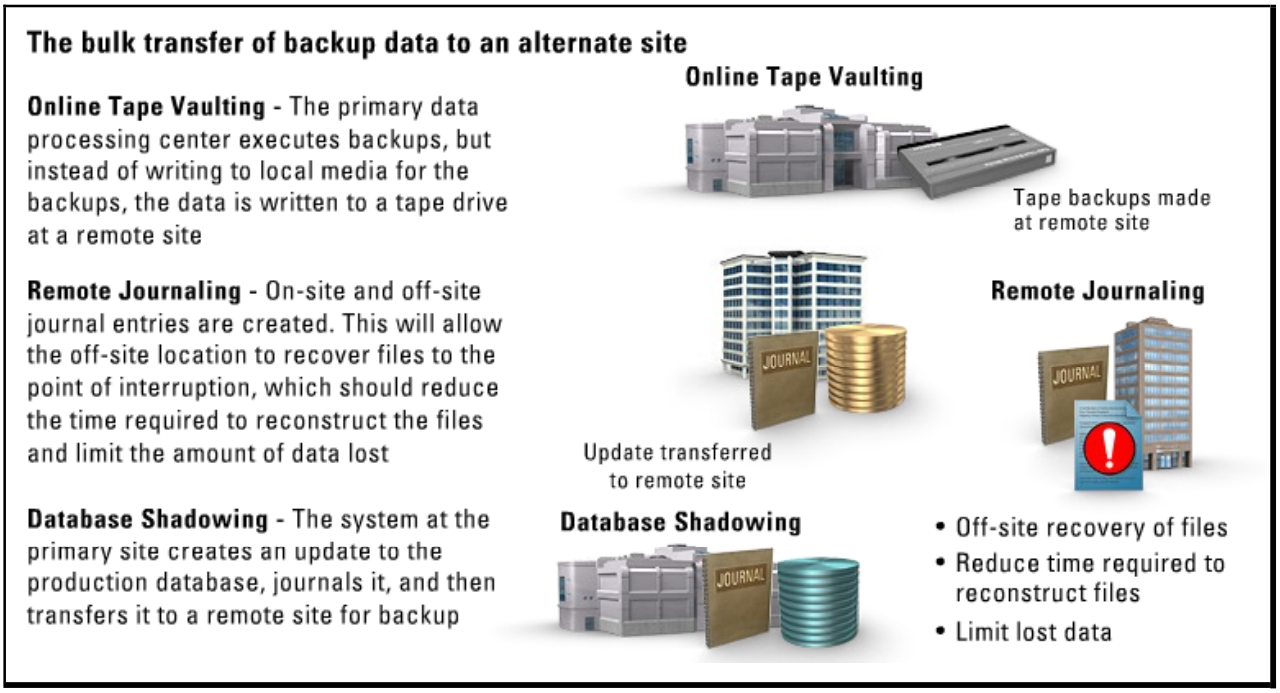

There are three types of electronic vaulting:

- **Online Tape Vaulting** - The primary data processing center executes backups, but instead of writing to local media for the backups, the data is written to a tape drive at a remote site.
- **Remote Journaling** - On-site and off-site journal entries are created. This will allow the off-site location to recover files to the point of interruption, which should reduce the time required to reconstruct the files and limit the amount of data lost.
- **Database Shadowing** - The system at the primary site creates an update to the production database, journals it, and then transfers it to a remote site for backup.

Summary

The key points discussed in this lesson are:

- There are four primary technical recovery strategies:
  - Subscription services
  - Reciprocal or mutual aid agreements
  - Multiple processing centers
  - Service bureaus
- The SLA you reach with a provider must be wide in scope, covering all key aspects of the service. It essentially frames the relationship, and determines the major responsibilities, both in times of normal operation and during an emergency situation.
- Subscription services are usually categorized as hot, warm, cold, mirror, or mobile sites.
- The data recovery process is the most critical recovery process, as data processing capabilities are the life blood of any organization.
- The electronic bulk transfer of backup data to an alternate site is called electronic vaulting.

Assurance and Trust

Overview

Providing assurance and trust in the business continuity plan is important. Knowing that the plan will work because you have tested it as well as kept it up to date with any changing business situations is critical when a disaster occurs.

Importance

It is important for the information security specialist to understand how to provide assurance and trust in the business continuity plan.

Objectives

Upon completing this lesson, you will be able to:

- List the goals of the testing strategy
- Describe the five main types of BCP testing strategies
- Identify plan maintenance techniques
- Identify the objectives of the BCP awareness training

Outline

The lesson contains these topics:

- Assurance, Trust, and Confidence Mechanisms
- BCP Testing Strategies
- Plan Maintenance Techniques
- BCP Awareness Training

Assurance, Trust, and Confidence Mechanisms

In the assurance and trust phase, you should implement plans for testing and maintaining the BCP, and perform awareness and training procedures.

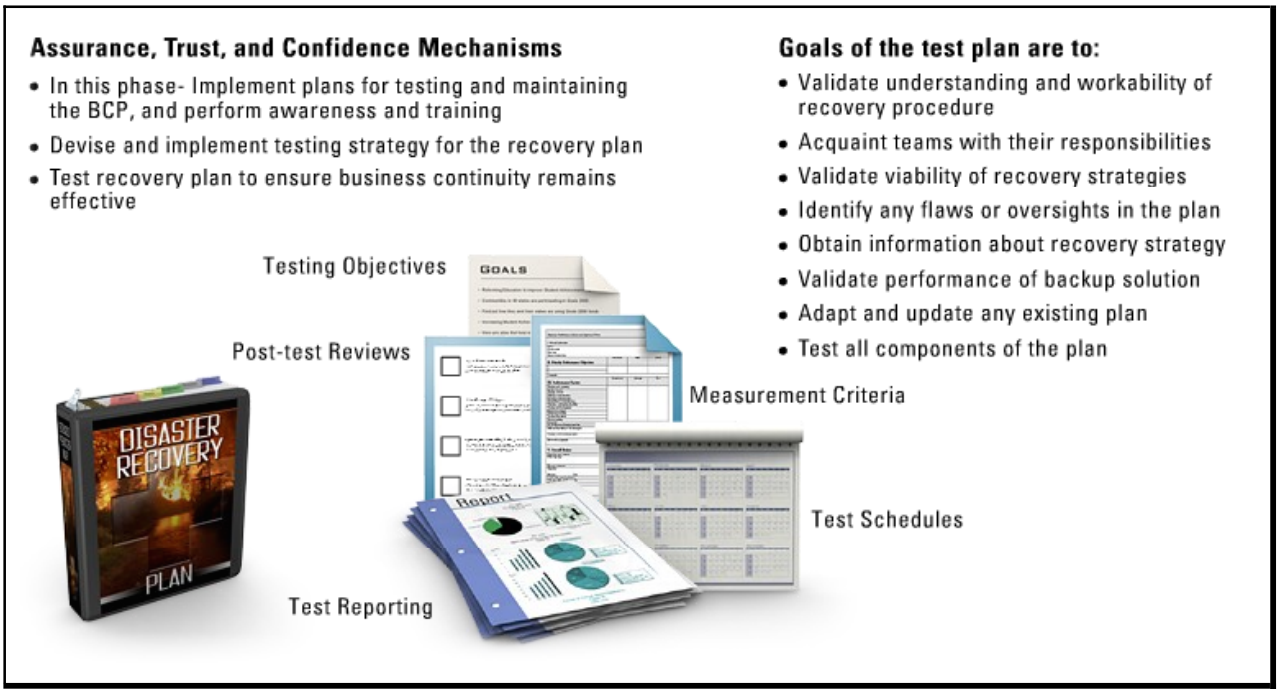

You must devise and implement a testing strategy for the recovery plan during this phase. You need to test the recovery plan to ensure that the business continuity capability remains effective. These test plans typically include the testing objectives, measurement criteria, test schedules, post-test reviews, and test reporting to management.

The goals of the test planning are to:

- Validate the understanding and workability of the documented recovery procedure
- Acquaint participants and teams with their responsibilities
- Validate viability of recovery strategies
- Identify any flaws or oversights in the plan
- Obtain information about the recovery strategy in action
- Validate performance of backup solution can meet production system demands
- Adapt and update any existing plan to meet new requirements
- Test all components of the plan

Before you perform the test, make sure to document the test goals, objectives, and scenario.

BCP Testing Strategies

This topic discusses BCP testing strategies.

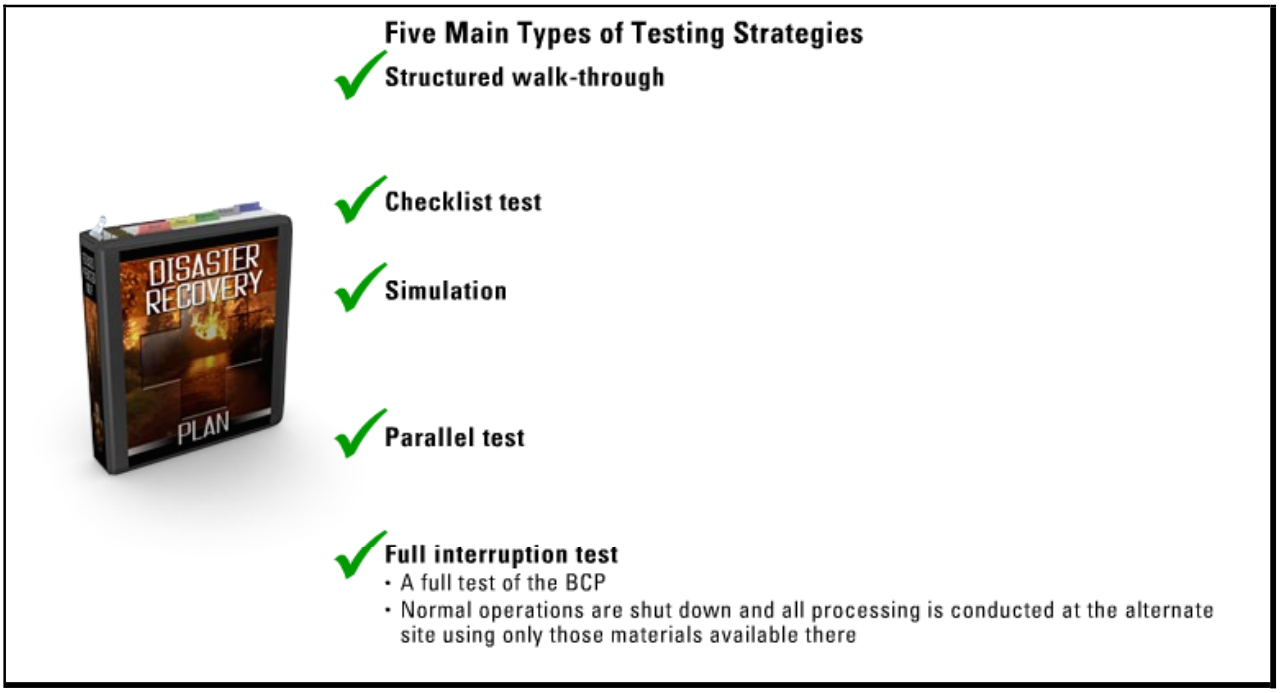

There are five main types of BCP testing strategies, which include:

- **Structured walk-through**
  - Process representatives meet to review the plan
  - Take a thorough look at each of the plan steps and procedures
  - Validate planned activities are accurately described
- **Checklist test**
  - Distribute copies of the test to each of the process areas
  - Each area reviews the plan
- **Simulation**
  - Operation and support members meet to practice execution of the BCP based on a test scenario
  - Only allow those materials and information that would be available in a real disaster
  - Simulation should continue up to the point of actual relocation to an alternate site
- **Parallel test**
  - An operational test of the BCP
  - Place critical systems in operation at alternate site to verify operation
  - Compare results with real operational output and note differences
- **Full interruption test**

- A full test of the BCP
- Normal operations are shut down and all processing is conducted at the alternate site using only those materials available there

Plan Maintenance Techniques

Any recovery plan must change when the goals of the business or organization change.

For a successful recovery to take place, you must perform strategic modifications to the plan that coincide with business changes.

Plan maintenance techniques include:

- As changes to business strategies occur, you must outline change management procedures to provide a method for regularly updating plan procedures
- Be sure to resolve all problems found during testing
- After evaluating the effectiveness of the recovery plan, auditors should report their findings
- Build maintenance procedures into the operation of the organization
- Provide a central repository for all updates
- Provide regular update reports to team members and management if necessary

BCP Awareness Training

Key components of a sound BCP include management's support of the development of the BCP as well as the provision of adequate financial and personnel resources required to develop a sound plan. Another key task that management must perform is providing a BCP awareness program.

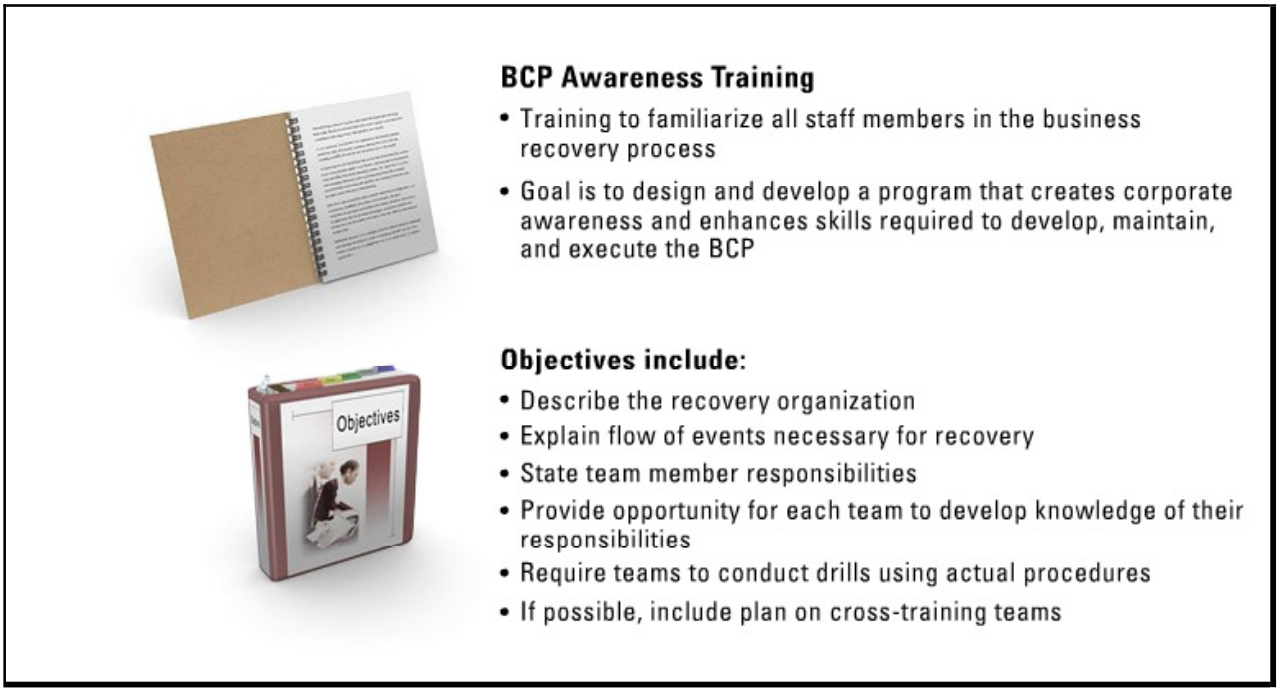

This training should familiarize all staff members in the business recovery process. The goal of this awareness training is to design and develop a program that creates corporate awareness and enhances the skills required to develop, implement, maintain, and execute the BCP.

The objectives of the awareness training should include:

- Describe the recovery organization
- Explain the flow of events necessary in the recovery process
- State team member responsibilities
- Provide an opportunity for each recovery team to develop an in-depth knowledge of their responsibilities
- Require teams to conduct drills using the actual checklists and procedures
- If possible, include a plan of cross-training teams

Summary

The key points discussed in this lesson are:

- You need to test the recovery plan to ensure that the business continuity capability remains effective. These test plans typically include the testing objectives, measurement criteria, test schedules, post-test reviews, and test reporting to management.
- There are five main types of BCP testing strategies, which include:
  - Structured walk-through
  - Checklist test
  - Simulation
  - Parallel test
  - Full interruption test
- For a successful recovery to take place, you must perform strategic modifications to the plan that coincide with business changes.
- The goal of the BCP awareness training is to design and develop a program that creates corporate awareness and enhances the skills required to develop, implement, maintain, and execute the BCP.

Computer Crime

Overview

Precise and reliable statistics on the amount of computer crime (or cybercrime) occurring today and the subsequent financial loss to victims are unknown. This is due to the fact that many of these crimes are not detected by the victims, and many of these crimes are never reported to authorities. But, there is a consensus among both law enforcement officials and security specialists that both the number of computer crime incidents and the sophistication of computer criminals are increasing rapidly.

Objectives

Upon completing this module, you will be able to:

- Identify types and key characteristics of computer crime
- Describe the major categories of laws
- Identify specific computer crime-related laws
- Identify the responsibilities of the information security specialist under the laws of due care
- Identify techniques for investigating computer crime and the associated ethical considerations

Outline

The module contains these lessons:

- Types of Computer Crime
- Major Categories of Laws
- Computer Crime-Related Laws
- Due Care
- Investigation and Ethics

Types of Computer Crime

Overview

Many types of attacks can take place against an enterprise. This lesson will discuss the types and key characteristics of computer crime.

Importance

It is important for the information security specialist to be able to identify each type of computer crime and its key characteristics.

Objectives

Upon completing this lesson, you will be able to:

- Describe the four general types of computer crime
- Differentiate breaches of communication, data security, and operations security
- Identify examples of computer crime
- List key characteristics of computer crime
- Identify key statistics of computer crime
- Identify patterns of computer crime
- Define patents, trademarks, and copyrights
- Define trade secrets

Outline

The lesson contains these topics:

- Four Types of Computer Crime
- Breaches of Communication, Data Security, and Operations Security
- Examples of Computer Crime
- Characteristics of Computer Crime
- Statistics of Computer Crime
- Patterns of Computer Crime

- Patents, Trademarks, and Copyrights
- Trade Secrets

Four Types of Computer Crime

There are four general types of computer crime.

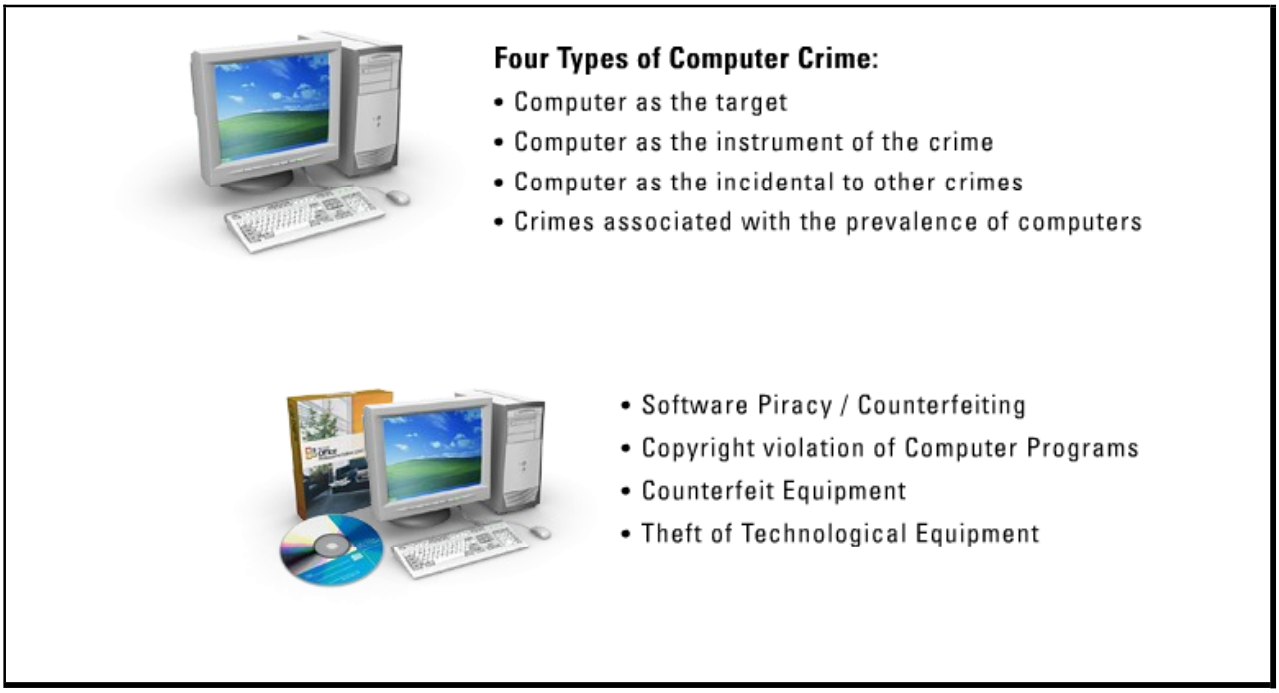

The four types include:

- Computer as the target
- Computer as the instrumentality of the crime
- Computer is incidental to other crimes
- Crimes associated with the prevalence of computers

However, in practice, multiple crimes, that is, concurrent criminality or lesser offenses, can occur during any given criminal transaction, resulting in an overlap among the classifications.

Crimes in which the **computer is the target** include such offenses as theft of intellectual property, theft of marketing information (e.g., customer lists, pricing data, or marketing plans), or blackmail based on information gained from computerized files (e.g., medical information, personal history, or sexual preference). These crimes also could entail sabotage of intellectual property, marketing, pricing, or personnel data or sabotage of operating systems and programs with the intent to impede a business or create chaos in an organization's operations.

Crimes in which the **computer is used as the instrument**: Essentially, the criminal introduces a new code (programming instructions) to manipulate the computer's analytical processes, thereby facilitating the crime. Another method involves converting legitimate computer processes for illegitimate purposes. Crimes in this category include fraudulent use of automated teller machine (ATM) cards and accounts; theft of money from accrual, conversion, or transfer accounts; credit card fraud; fraud from computer transactions (e.g., stock transfers, sales, or billings); and telecommunications fraud.

Crimes in which the **computer is incidental to other crimes**: In this category of computer crime, the computer is not essential for the crime to occur, but it is related to the criminal act. This means that the

crime could occur without the technology; however, computerization helps the crime to occur faster, permits processing of greater amounts of information, and makes the crime more difficult to identify and trace. Such crimes include money laundering and unlawful banking transactions, bulletin board systems (BBSs) supporting unlawful activity, organized crime records or books, and bookmaking.

In **crimes associated with the prevalence of computers**, the simple presence of computers, and notably the widespread growth of microcomputers, generates new versions of fairly traditional crimes. In these cases, technological growth essentially creates new crime targets. Software piracy/counterfeiting, copyright violation of computer programs, counterfeit equipment, black market computer equipment and programs, and theft of technological equipment fall into this category of computer crime.

Breaches of Communication, Data Security, and Operations Security

This topic differentiates breaches of communication, data security, and operations security.

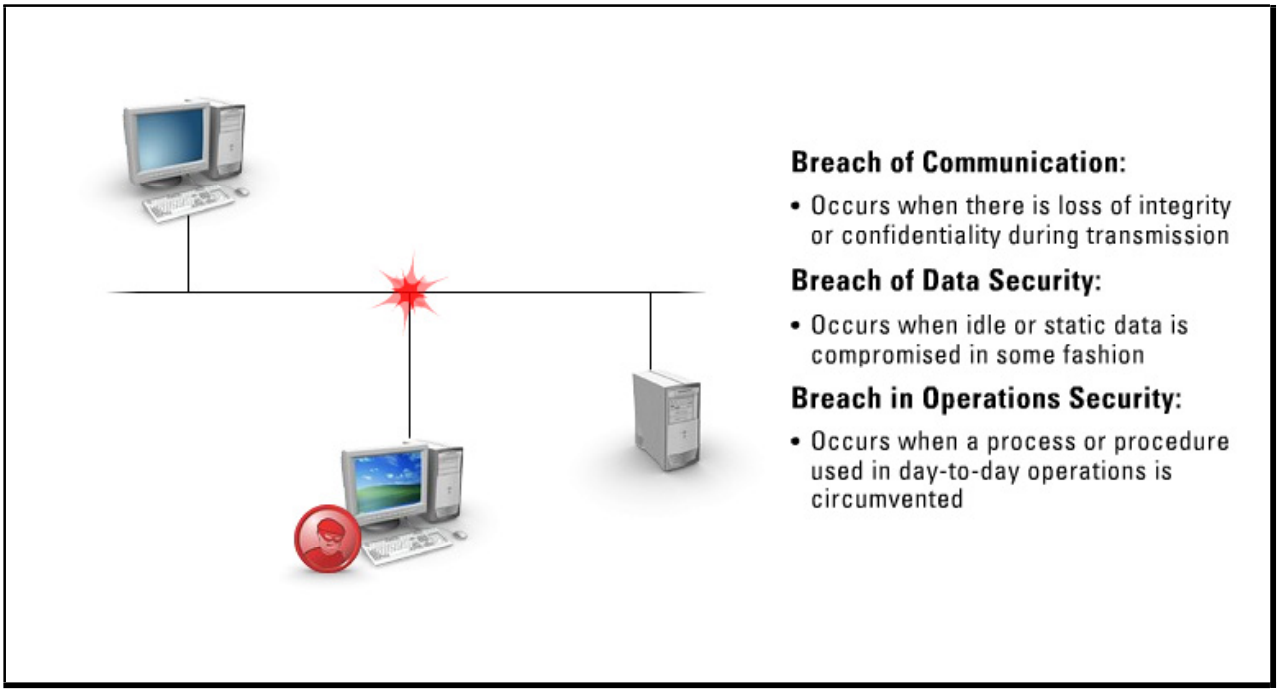

A **breach of communication** occurs when there is a loss of integrity or confidentiality during a transmission between parties. Sniffing can produce a breach of communication, but if the data are encrypted, then no breach can occur as the data are protected. A **breach of data security** occurs when idle or static data is compromised in some fashion. A **breach of operations security** occurs when a process or procedure that is used in the day-to-day operations of an enterprise is circumvented.

Examples of Computer Crime

Computer crimes can range from the catastrophic to the merely annoying.

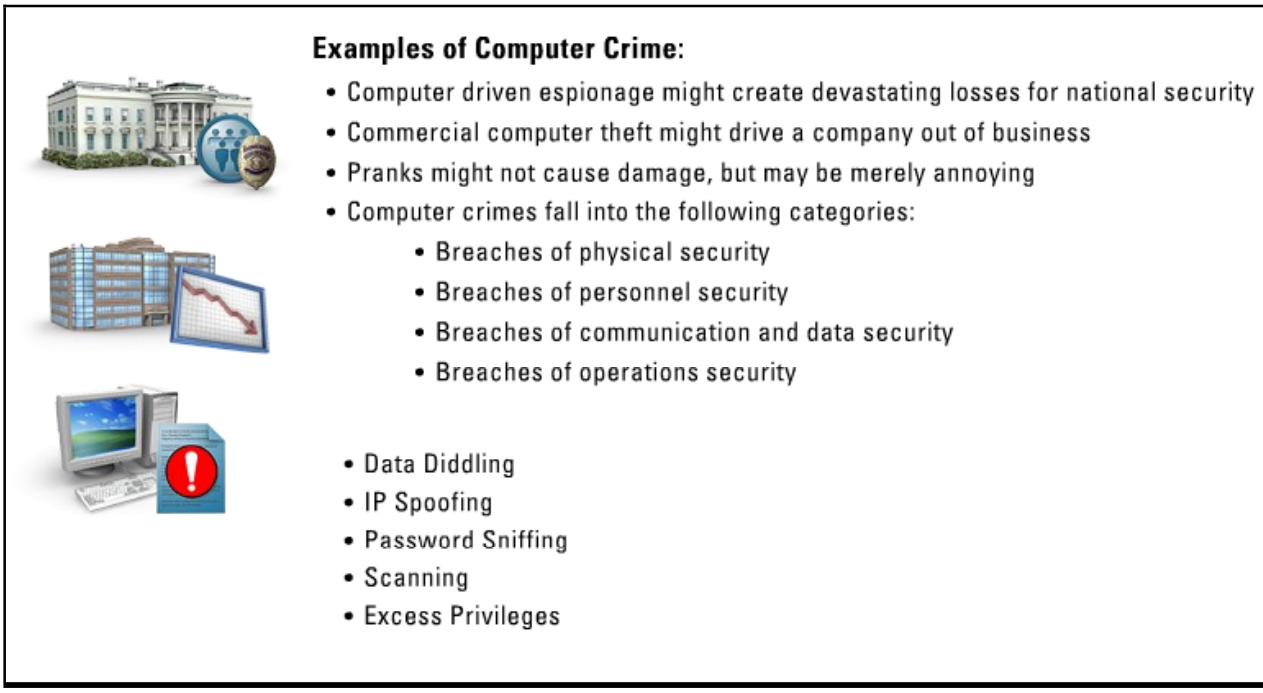

A case of computer-driven espionage might create devastating losses for national security. A case of commercial computer theft might drive a company out of business. A cracker's prank might not actually cause damage at all, but rather might cause a video game company or another computer user some annoyance. Some computer crimes are perpetrated for kicks, and some for social or political causes; others are the serious business of professional criminals.

Note	There is perhaps no other form of crime that cuts so broadly across the types of criminals and the severity of their offenses.

In general, computer crimes fall in the following categories:

- **Breaches of Physical Security**
  - Dumpster diving
  - Wiretapping
  - Eavesdropping on emanations
  - Denial of service (DoS)
- **Breaches of Personnel Security**
  - Masquerading
  - Social engineering
  - Harassment
  - Software piracy

- **Breaches of Communication and Data Security**
  - Unauthorized copying of data
  - Traffic analysis
  - Covert channels
  - Trapdoors
  - Session highjacking
  - Tunneling
  - Timing attacks
  - Trojan horses
  - Viruses and worms
  - Salamis and logic bombs
- **Breaches of Operations Security**
  - Data diddling
  - IP spoofing
  - Password sniffing
  - Scanning
  - Excess privileges

Characteristics of Computer Crime

This topic discusses key characteristics of computer crime.

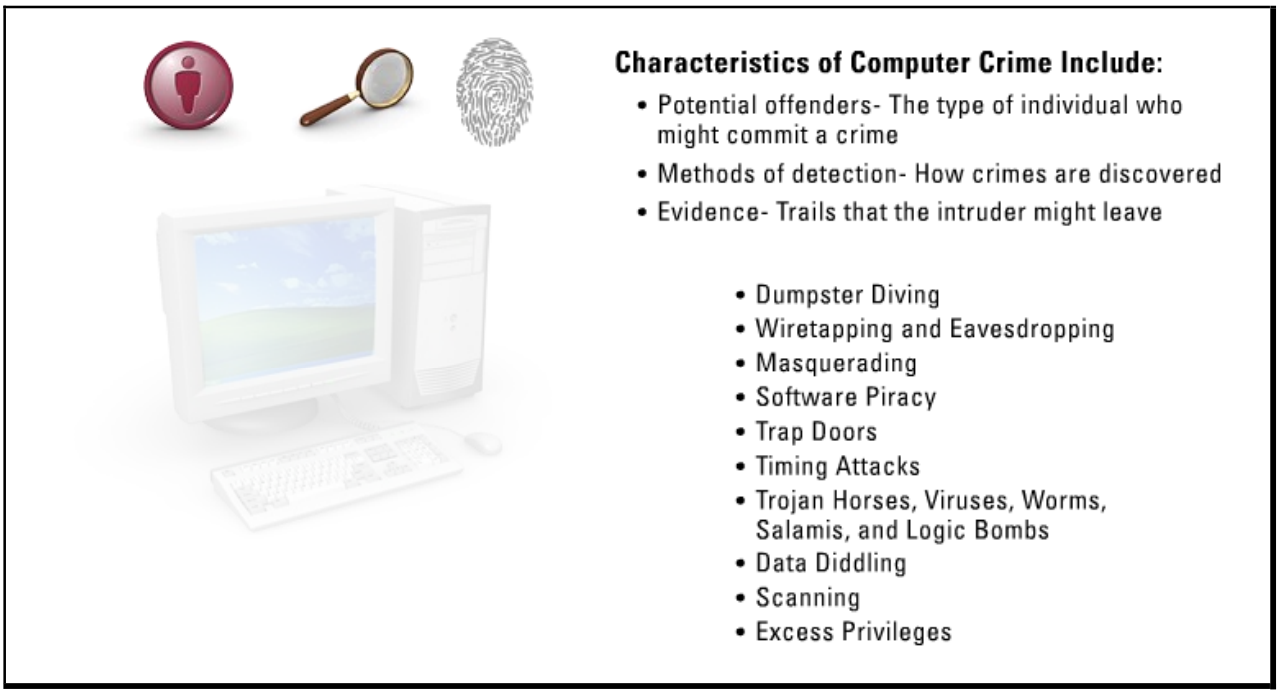

Before learning about the characteristics of computer crime, it is important to understand the following:

- **Potential Offenders** - The type of individual (e.g., programmers or spies) who might commit a crime of this type
- **Methods of Detection** - How such crimes are discovered (e.g., tracing equipment of various kinds or analyzing log files)
- **Evidence** - Trails that the intruder might leave, which might help in detection (e.g., system logs or telephone company records)

Dumpster Diving

Potential Offenders:

- System users
- Anyone able to access the trash area
- Anyone who has access to computer areas or areas used to store backups

Methods of Detection:

- Tracing proprietary information back to its source (e.g., memos with company names or logos)
- Observation (guards may actually see intruders in action)
- Testing an operating system to discover data left over after job execution

Evidence:

- Computer output media (e.g., may contain vendor name or identifying page numbers)
- Similar information produced in suspected ways in the same form

- Characteristics of printout or other media (e.g., type fonts or logos)

Wiretapping and Eavesdropping

Potential Offenders:

- Communications technicians and engineers
- Agents for competitors
- Communications employees, former employees, vendors, and contractors
- Agents for foreign intelligence services

Methods of Detection:

- Voice wiretapping methods
- Tracing the origin of the equipment used in the crime (e.g., monitoring equipment)
- Tracing computer output (e.g., disks and tapes) to the source
- Observation
- Discovery of stolen information

Evidence:

- Voice wiretapping
- Computer output forms
- Computer audit logs
- Computer storage media
- Characteristics of printout or other media (e.g., type fonts or logos)
- Manual after-hours sign-in/sign-out sheets

Masquerading

Potential Offenders:

- Potentially everyone

Methods of Detection:

- Analysis of audit logs and journals (e.g., a log shows that an authorized user apparently logged in, but it is known that the person was away at that time)
- Observation (e.g., an eyewitness saw an intruder at an authorized user's terminal)
- Password violations (e.g., a log shows repeated failed attempts to use an invalid password)
- Report by the person who has been impersonated (e.g., the authorized person logs in, and the system tells the person that he has had six unsuccessful logins since the last time he knows he actually logged in)

Evidence:

- Backups
- System audit logs
- Telephone company records (pen register and dialed number recorder [DNR] records)
- Violation reports from access control packages
- Notes and documents found in the possession of suspects

- Witnesses
- Excessively large phone bills (excessive message units may indicate that someone is using resources)

Software Piracy

Potential Offenders:

- Purchasers and users of commercial software
- Software pirates
- Employees who steal proprietary software

Methods of Detection:

- Observation
- Testimony of legitimate purchasers of software
- Search of users' facilities and computers

Evidence:

- Pictures of computer screens where pirated software is being executed
- The contents of memory in computers containing pirated software
- Copies of media on which pirated software is found
- Printouts produced by pirated software

Trap Doors

Potential Offenders:

- Systems programmers
- Applications programmers

Methods of Detection:

- Exhaustive testing
- Specific testing based on evidence
- Comparison of specifications to performance

Evidence:

- Programs that perform tasks not specified for them
- Output reports that indicate that programs are performing tasks not specified for them

Timing Attacks

Potential Offenders:

- Advanced system analysts
- Advanced computer programmers

Methods of Detection:

- System testing of suspected attack methods
- Complaints from system users that their jobs are not being performed efficiently
- Repeat execution of a job under normal and safe conditions

Evidence:

- Output that deviates from normally expected output of logs
- Computer operations logs

Trojan Horses, Viruses, Worms, Salamis, and Logic Bombs

Potential Offenders:

- Programmers who have detailed knowledge of a program
- Employees or former employees
- Vendor or contractor programmers
- Financial system programmers
- Computer users
- Computer operators
- Crackers

Methods of Detection:

- Comparison of program code with backup copies of the program
- Tracing of unexpected events of possible gain from the act to suspected perpetrators
- Detailed data analysis, including analysis of program code (e.g., you may detect a virus because a file increases in size when it is modified or because disk space decreases)
- Observation of financial activities of possible suspects (especially for salami attacks)
- Testing of suspect programs
- Examination of computer audit logs for suspicious programs or pertinent entries (e.g., log entries that show that many programs were updated at the same time) (especially for viruses)
- Transaction audits

Evidence:

- Output reports
- Unexpected results of running programs
- Computer usage and file request journals
- Undocumented transactions
- Analysis test program results
- Audit logs

Data Diddling

Potential Offenders:

- Participants in transactions being entered or updated
- Suppliers of source data
- Preparers of data
- Nonparticipants with access

Methods of Detection:

- Comparison of data
- Manual controls

- Analysis of computer validation reports
- Integrity tests
- Validation of documents
- Analysis of audit logs
- Analysis of computer output

Evidence:

- Data documents for source data, transactions, etc.
- Manual logs, audit logs, journals, etc.
- Backups and other computer media (e.g., tapes and disks)
- Incorrect computer output control violation alarms

Scanning

Potential Offenders:

- Malicious intruders
- Spies attempting to access systems for targeted data
- Criminals intent on committing fraud

Methods of Detection:

- Computer logs that show when telephone calls were received by the computer and when attempts were made
- Loss of data or transfer of funds or other assets
- Telephone company records

Evidence:

- Telephone company records (pen register and dialed number recorder (DNR) records)
- Possession of war dialing programs
- Computer logs
- Possession of information compromised as a result of scanning, including lists of telephone numbers

Excess Privileges

Potential Offenders:

- Programmers with access to Superzap-type programs
- Computer operations staff

Methods of Detection:

- Comparison of files with historical copies
- Examination of computer usage logs
- Discrepancies noted by those who receive reports

Evidence:

- Discrepancies in output reports
- Computer usage and file request journals

- Undocumented transactions

Statistics of Computer Crime

This topic discusses key computer crime statistics.

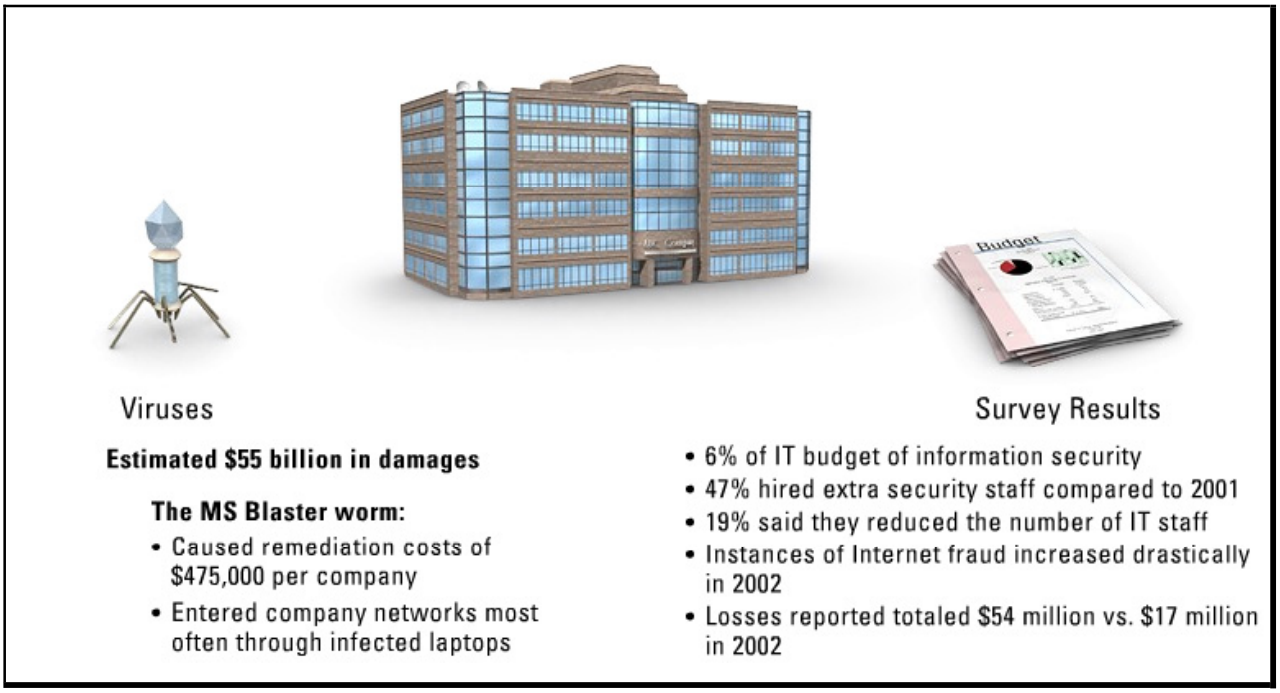

In January 2004:

- It was estimated that PC viruses cost businesses approximately $55 billion in damages in 2003.
- The same calculations were done in 2002 and 2001, at $20 to $30 billion and $13 billion, respectively.

In August 2003, a survey including 882 respondents determined that the MS Blaster worm:

- Caused remediation costs of $475,000 per company (median average that included hard, soft, and productivity costs) with larger node-count companies reporting losses up to $4,228,000.
- Entered company networks most often through infected laptops, then through virtual private networks (VPNs), and finally through misconfigured firewalls or routers.

In May 2003, a survey team determined:

- Financial services companies were spending approximately six percent of their IT budgets on information security.
- Forty-seven percent hired extra security staff compared with 2001.
- Only 19 percent of respondents said they had reduced the number of IT security staff, despite the slowdown in the economy.

In April 2003, a survey determined that:

- Instances of Internet fraud increased drastically in 2002 as compared to 2001.
- Losses reported by victims totaled $54 million, versus $17 million the year before, and complaints referred to law enforcement totaled 48,252, compared to 16,755 in 2001.

- Auction fraud and non-delivery of merchandise were the top two reported crimes, with credit and debit card fraud following them at 12 percent.

Patterns of Computer Crime

The Computer Crime and Security Survey is conducted by CSI with the participation of the San Francisco Federal Bureau of Investigation's (FBI) Computer Intrusion Squad. CSI, established in 1974, is a San Francisco-based association of information security professionals.

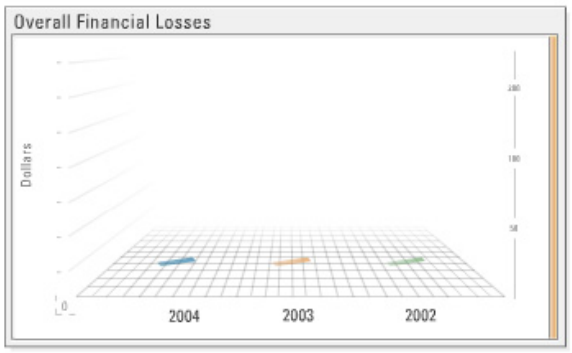

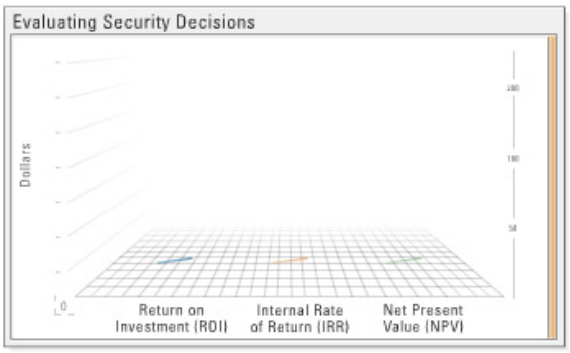

The aim of this effort is to raise the level of security awareness, as well as help determine the scope of computer crime in the United States. CSI has thousands of members worldwide and provides a wide variety of information and education programs to assist practitioners in protecting the information assets of corporations and governmental organizations.

Highlights of the 2002 Computer Crime and Security Survey include:

- Ninety percent of respondents (primarily large corporations and government agencies) detected computer security breaches within the last 12 months.
- Eighty percent of respondents acknowledged financial losses due to computer breaches.
- Forty-four percent (223 respondents) were willing and/or able to quantify their financial losses. These 223 respondents reported $455,848,000 in financial losses.
- As in previous years, the most serious financial losses occurred through theft of proprietary information (26 respondents reported $170,827,000 in losses) and financial fraud (25 respondents reported $115,753,000 in losses).
- For the fifth year in a row, more respondents (74 percent) cited their Internet connections as a frequent point of attack than cited their internal systems as a frequent point of attack (33 percent).
- Thirty-four percent reported the intrusions to law enforcement. (In 1996, only 16 percent acknowledged reporting intrusions to law enforcement.)

Highlights of the 2003 Computer Crime and Security Survey include:

- Overall financial losses from 530 survey respondents totaled $201,797,340. This is down significantly from 503 respondents reporting $455,848,000 last year. (Seventy-five percent of organizations acknowledged financial losses, though only 47 percent could quantify them.)
- The overall number of significant incidents remained roughly the same as last year, despite the drop in financial losses.
- Losses reported for financial fraud were drastically lower, at $9,171,400. This compares to nearly $116 million reported last year.
- As in prior years, theft of proprietary information caused the greatest financial loss ($70,195,900 was lost, with the average reported loss being approximately $2.7 million).
- In a shift from previous years, the second-most expensive computer crime among survey respondents was denial of service (DoS), with a cost of $65,643,300 – up 250 percent from last year's losses of $18,370,500.

Highlights of the 2004 Computer Crime and Security Survey include:

- Overall financial losses totaled from 494 survey respondents were $141,496,560. This is down significantly from 530 respondents reporting $201,797,340 last year.
- In a shift from previous years, the most expensive computer crime was denial of service. Theft of intellectual property, the prior leading category, was the second most expensive in 2004.
- Organizations are using metrics from economics to evaluate their security decisions. Fifty-five percent use Return on Investment (ROI), 28 percent use Internal Rate of Return (IRR), and 25 percent use Net Present Value (NPV).
- The vast majority of organizations in the survey do not outsource computer security activities. Among those organizations that do outsource some computer security activities, the percentage of security activities outsourced is quite low.

Patents, Trademarks, and Copyrights

This topic defines patents, trademarks, and copyrights.

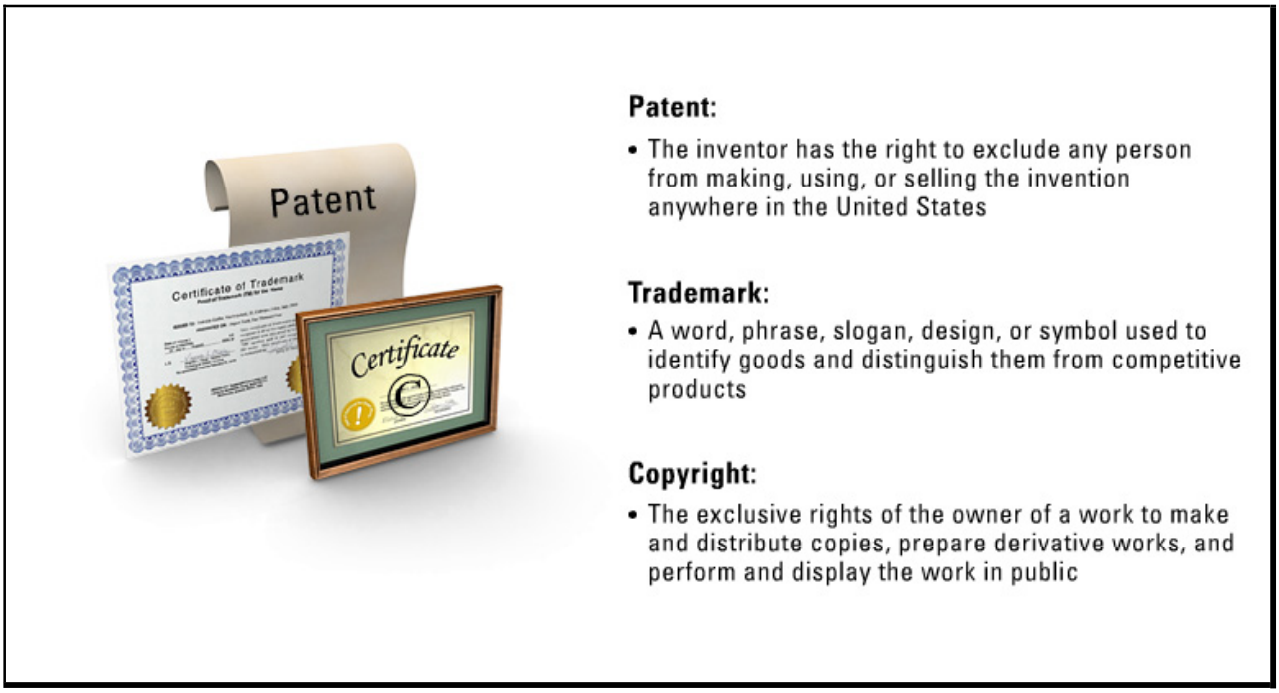

A **patent** is a legal document issued by the United States to an inventor. The inventor, as the owner of the patent, has the right to exclude any other person from making, using, or selling the invention covered by the patent anywhere in the United States for 17 years from the date the patent was issued. The U.S. Patent and Trademark Office issues a document to the inventor. This document contains a detailed description of the invention, how to make or use it, and what rights the inventor has.

A **trademark** is a word, phrase, slogan, design, or symbol used to identify goods and distinguish them from competitive products. You may register a trademark with the U.S. Patent and Trademark Office and similar offices worldwide. However, in the United States and in other countries with legal systems based on English common law, trademark rights also accrue through common law usage.

A **copyright** describes the exclusive rights of the owner of a work to make and distribute copies, prepare derivative works, and perform and display the work in public. (The last two rights mainly apply to plays, films, dances, and the like, but could also apply to software.)

A work, including a piece of software, is under copyright by default in most countries, whether or not it displays a copyright notice. However, a copyright notice may make it easier to assert ownership. The copyright owner is the person or company whose name appears in the copyright notice on the box, disk, screen, etc.

A copyright notice has three parts. The first part can be either a ‘c’ with a circle around it, the word ‘Copyright’, or the abbreviation ‘Copr’. (A ‘c’ in parentheses has no legal meaning.) This is followed by the name of the copyright holder and the year of first publication.

Trade Secrets

Trade secret protection stems from the common law and dates back to the 1800s.

A **trade secret** may consist of any formula, pattern, device, or compilation of information that a person uses in a business and which gives the person an opportunity to obtain an advantage over competitors who do not know or use it.

Note	A trade secret may be a formula for a chemical compound, a process of manufacturing, treating, or preserving materials, a pattern for a machine or other device, or a list of customers.

While patents and copyrights require you to disclose your information in the application process (information that eventually becomes public), trade secrets require you to actively keep the information secret.

Additional examples of trade secrets include customer identities and preferences, vendors, product pricing, marketing strategies, company finances, manufacturing processes, and other valuable information that makes a company competitive.

Summary

The key points discussed in this lesson are:

- There are four general types of computer crime:
  - Computer as the target
  - Computer as the instrumentality of the crime
  - Computer is incidental to other crimes
  - Crimes associated with the prevalence of computers
- A breach of communication occurs when there is a loss of integrity or confidentiality during a transmission between parties. A breach of data security occurs when idle or static data is compromised in some fashion. A breach of operations security occurs when a process or procedure that is used in the day-to-day operations of an enterprise is circumvented.
- Computer crimes can range from the catastrophic to the merely annoying. There is perhaps no other form of crime that cuts so broadly across the types of criminals and the severity of their offenses.
- In January 2004, it was estimated that PC viruses cost businesses approximately $55 billion in damages in 2003.
- The inventor, as the owner of the patent, has the right to exclude any other person from making, using, or selling the invention covered by the patent anywhere in the United States for 17 years from the date the patent was issued. A trademark is a word, phrase, slogan, design, or symbol used to identify goods and distinguish them from competitive products. A copyright describes the exclusive rights of the owner of a work to make and distribute copies, prepare derivative works, and perform and display the work in public.
- A trade secret may consist of any formula, pattern, device, or compilation of information that a person uses in a business and which gives the person an opportunity to obtain an advantage over competitors who do not know or use it.

Major Categories of Laws

Overview

Because of the tremendous amount of damages inflicted by computer attackers, the U.S. government has established laws and regulations against computer crime. Sentences that were once thought to be too light have become much harsher on criminals who openly and wantonly destroy or divulge sensitive corporate or personal information. This lesson will discuss the major categories of laws and specific laws related to information security and privacy.

Importance

It is important for the information security specialist to understand what laws are affected when a crime against the enterprise occurs.

Objectives

Upon completing this lesson, you will be able to:

- Describe common law systems
- Describe civil law systems
- Describe the major categories of laws
- Identify information security-related laws
- Identify privacy-related laws

Outline

The lesson contains these topics:

- Common Law Systems (United States)
- Civil Law Systems (Europe)
- Major Categories of Laws
- Information Security-Related Laws
- Privacy-Related Laws

Common Law Systems (United States)

The common law system originally developed under the auspices of the adversarial system in historical England from judicial decisions that were based in tradition, custom, and precedent. The form of reasoning used in common law is known as casuistry or case-based reasoning.

Common Law Systems (United States):

- Law based on tradition, custom, and precedent
- Form of reasoning is known as casuistry or case-based reasoning
- Written in statutes or codes
- Devised as:
- A means of compensating someone for wrongful acts known as torts
- A means of developing a body of law recognizing and regulating contracts
- Usually applies only to civil disputes
- Constitutes the basis of the legal system in England and Wales, The Republic of Ireland, the United States (except Louisiana and Puerto Rico), Canada (except Quebec) and other generally English-speaking countries

Common law may be unwritten or written in statutes or codes. The common law, as applied in civil cases (which are distinct from criminal cases), was devised as a means of compensating someone for wrongful acts known as torts, including both intentional torts and torts caused by negligence, and as a means of developing the body of law recognizing and regulating contracts.

Today, common law is generally thought of as applying only to civil disputes; originally it encompassed the criminal law before criminal codes were adopted in most common law jurisdictions in the late 19th century. The type of procedure practiced in common law courts is known as the adversarial system; this is also a development of the common law.

The common law constitutes the basis of the legal systems of: England and Wales, the Republic of Ireland, the United States (except Louisiana and Puerto Rico), Canada (except Quebec private law), Australia, New Zealand, South Africa, India, Malaysia, Singapore, Hong Kong, and many other generally English-speaking countries or Commonwealth countries.

Civil Law Systems (Europe)

The main alternative to the common law system is the civil law system, which is used in Continental Europe, the former Soviet bloc, and most of the rest of the world.

Civil Law Systems (Europe):

- Used in Continental Europe, the former Soviet bloc, and most of the rest of the world
- A legal tradition based on Roman Law

Difference between Civil Law and Common Law:

- Common Law was law developed by custom
- Developed out of the Roman law and specifically from broad legal principles and the interpretation of doctrinal writings rather than the application of facts to legal fictions
- Later, civil law become codified as customary laws, that were local compilations of legal principles recognized as normative

Civil or civilian law is a legal tradition that is the basis of the law in the majority of countries of the world, especially in continental Europe, but also in Quebec (Canada), Louisiana (United States), Japan, Latin America, and most former colonies of continental European countries.

Civil law is based on Roman law, especially the Corpus Juris Civilis of Emperor Justinian, as later developed through the Middle Ages by medieval legal scholars.

Originally, civil law was one common legal system in much of Europe, but with the development of nationalism in the Nordic countries in the 17th century and around the time of the French Revolution, it became fractured into separate national systems. This change was brought about by the development of separate national codes.

Civil law is primarily contrasted against common law, which is the legal system developed among Anglo-Saxon peoples, especially in England. The original difference is that, historically, common law was law developed by custom, beginning before there were any written laws and continuing to be applied by courts after there were written laws. Civil law, on the other hand, developed out of the Roman law of Justinian's Corpus Juris Civilis, and specifically from broad legal principles and the interpretation of doctrinal writings rather than the application of facts to legal fictions.

In later times, civil law became codified as droit coutumier, or customary laws, that were local compilations of legal principles recognized as normative. Sparked by the Age of Enlightenment, attempts to codify private law began during the second half of the 18th century, but civil codes with a lasting influence were promulgated only after the French Revolution, in jurisdictions such as France (with its Napoleonic Code), Austria, Quebec, Spain, the Netherlands, and Germany.

Thus, the difference between civil law and common law lies less in the mere fact of codification, but in the methodological approach to codes and statutes. In civil law countries, legislation is seen as the primary source of law. Thus, by default, courts base their judgments on the provisions of codes and statutes, from which solutions in particular cases are to be derived. Courts thus have to reason extensively on the basis of general principles of the code, or by drawing analogies from statutory provisions to fill lacunae. In contrast, in the common law system, cases are the primary source of law, while statutes are only seen as incursions into the common law and are thus interpreted narrowly.

Major Categories of Laws

This topic discusses the major categories of laws.

There are two major categories of laws:

- **Criminal Law**
- **Civil Law**
  - **Contract Law** - Laws pertaining to an agreement between two or more parties that creates in each party a duty to do or not do something and a right to performance of the other's duty or a remedy for the breach of the other's duty.
  - **Property Law** - Laws pertaining to property that derives from the work of the mind or intellect, specifically, an idea, an invention, a trade secret, a process, a program, data, a formula, a patent, a copyright, a trademark, an application, a right, or a registration.
  - **Tort Law** - Laws pertaining to tort. A tort occurs when someone deliberately or through carelessness causes harm or loss to another person or his or her property.

Criminal law (also known as penal law) is the body of law that regulates governmental sanctions (such as imprisonment and/or fines) as retaliation for crimes against the social order. The goal of this process is to achieve criminal justice.

According to criminal law, crimes are offences against the social order. In common law jurisdictions, there is a legal fiction that crimes disturb the peace of the sovereign. Government officials, as agents of the sovereign, are responsible for the prosecution of offenders. Hence, the criminal law "plaintiff" is the sovereign, which in practical terms translates into the monarch or the people.

The major objectives of criminal law are deterrence and punishment, while the major objective of civil law is individual compensation. Criminal offences consist of two distinct elements: the physical act (the actus reus, guilty act) and the requisite mental state with which the act is done (the mens rea, guilty mind).

Criminal law distinguishes crimes from civil wrongdoings such as tort or breach of contract. Criminal law has been seen as a system of regulating the behavior of individuals and groups in relation to societal norms at large whereas civil law is aimed primarily at the relationships between private individuals and their rights and obligations under the law.

Information Security-Related Laws

The primary federal statute regarding computer fraud is 18 U.S.C. Section 1030, which was amended in October 1996 to protect computer and data integrity, confidentiality, and availability.

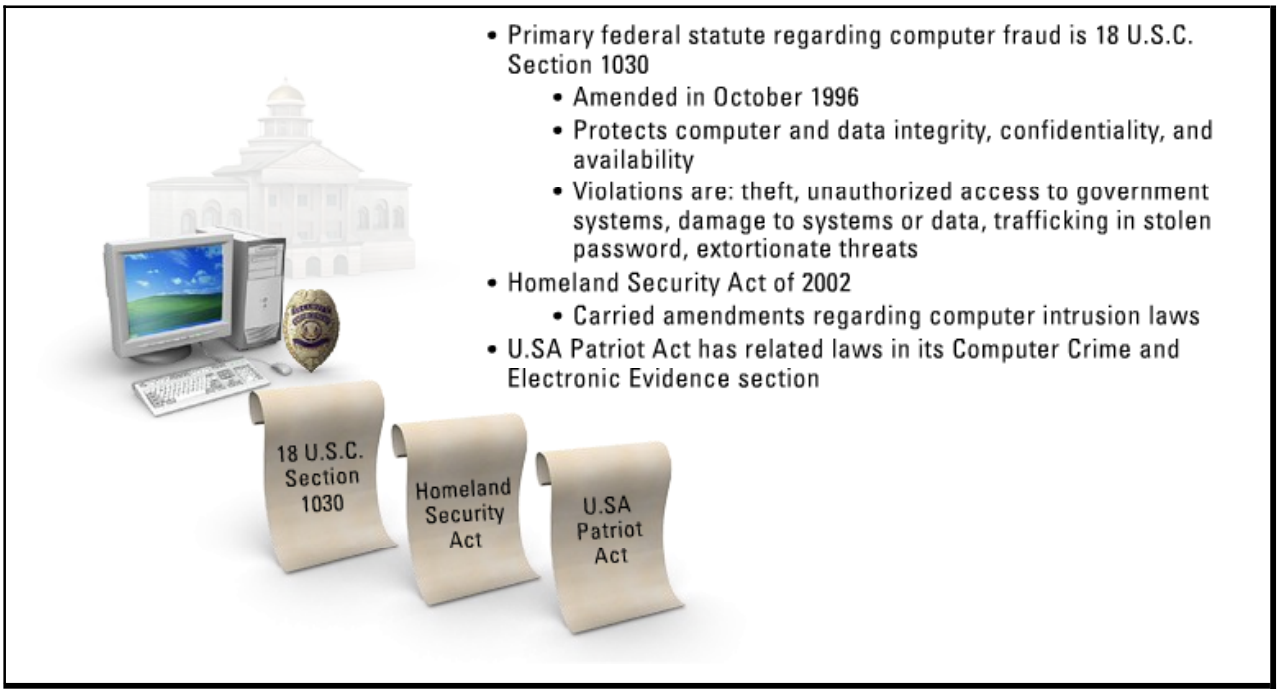

Examples of violations are:

- Theft of information from computers belonging to financial institutions or federal agencies, or computers used in interstate commerce
- Unauthorized access to government computers
- Damage to systems or data (intentionally or recklessly)
- Trafficking in stolen passwords
- Extortionate threats to damage computers

The Homeland Security Act of 2002 carried amendments regarding computer intrusion laws:

- Section 225, also known as the Cyber Security Enhancement Act of 2002
- Amendments to Sections 225 and 896 of the Homeland Security Act of 2002

The USA Patriot Act has related laws in its Computer Crime and Electronic Evidence section.

Other federal criminal codes related to computer intrusions include:

- 18 U.S.C. § 1029. Fraud and Related Activity in Connection with Access Devices
- 18 U.S.C. § 1362. Communication Lines, Stations, or Systems
- 18 U.S.C. § 2510 et seq. Wire and Electronic Communications Interception and Interception of Oral Communications
- 18 U.S.C. § 2701 et seq. Stored Wire and Electronic Communications and Transactional Records Access

- 18 U.S.C. § 3121 et seq. Recording of Dialing, Routing, Addressing, and Signaling Information

Privacy-Related Laws

This topic discusses privacy-related laws in the United States.

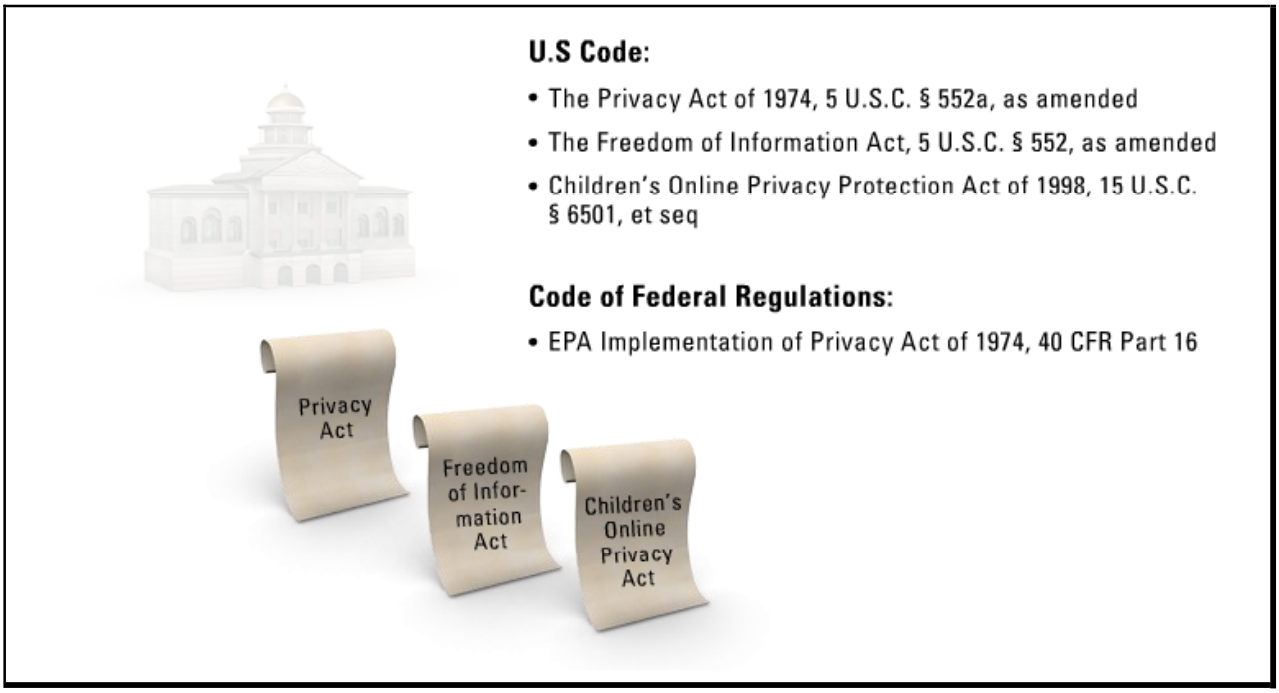

U.S. Code:

- The Privacy Act of 1974, 5 U.S.C. § 552a, as amended
- The Freedom of Information Act, 5 U.S.C. § 552, as amended
- Children's Online Privacy Protection Act of 1998, 15 U.S.C. § 6501, et seq.

Code of Federal Regulations:

- EPA Implementation of Privacy Act of 1974, 40 CFR Part 16

Office of Management and Budget Circulars, Memoranda, and Guidance:

- Management of Federal Information Resources, OMB Circular No. A-130
- Guidance on Interagency Sharing of Personal Data - Protecting Personal Privacy, OMB Memorandum M-01-05
- Privacy Policies and Data Collection on Federal Web Sites, OMB Memorandum M-00-13
- Privacy Policies on Federal Web Sites, OMB Memorandum M-99-18
- Privacy Guidance and Reference Materials
  - Privacy Act Implementation: Guidelines and Responsibilities, 40 FR 28948
  - Final Guidance Interpreting the Provisions of Public Law 100-503, Computer Matching and Privacy Protection Act of 1988, 54 FR 25818

Other Authorities:

- Protection of Individual Privacy, Federal Acquisition Regulation - Subpart 24.1
- Protection of Individual Privacy, EPA Acquisition Regulation - Subpart 1524.1

Summary

The key points discussed in this lesson are:

- The common law constitutes the basis of the legal systems of: England and Wales, the Republic of Ireland, the United States (except Louisiana and Puerto Rico), Canada (except Quebec private law), Australia, New Zealand, South Africa, India, Malaysia, Singapore, Hong Kong, and many other generally English-speaking countries or Commonwealth countries.
- Civil or civilian law is a legal tradition that is the basis of the law in the majority of countries of the world, especially in continental Europe, but also in Quebec (Canada), Louisiana (United States), Japan, Latin America, and most former colonies of continental European countries.
- There are two major categories of laws: criminal law and civil law.
- The primary federal statute regarding computer fraud is 18 U.S.C. Section 1030, which was amended in October 1996 to protect computer and data integrity, confidentiality, and availability.

Computer Crime-Related Laws

Overview

The U.S. government has devised protections against computer crimes in the form of laws and acts. This lesson will discuss the various legal laws created by the United States as well as international laws on computer crime that cross international borders.

Importance

It is important the information security specialist understand what laws and acts are enacted to protect an enterprise from computer crime.

Objectives

Upon completing this lesson, you will be able to:

- Identify key characteristics of U.S. privacy laws
- Identify the significant influences in the Privacy Act of 1974
- Identify the major revisions brought about by the Comprehensive Crime Control Act of 1984
- Define the reason for the enactment of the U.S. Medical Computer Crime Act of 1984
- Identify the actions prohibited by the Computer Fraud and Abuse Act
- Describe how the National Information Infrastructure Protection Act amended the Computer Fraud and Abuse Act
- Define federal interest computers
- Describe the protections provided under the Electronic Communications Privacy Act of 1986
- Identify the key points of the Computer Security Act of 1987
- Distinguish between NIST and NSA controls
- Identify the shortcoming of Section 342.1 of the Criminal Code of Canada
- Identify the key points of the U.S. Copyright Act and Canada's Bill C-17
- Describe the steps taken by the European Union to fight computer crime

- List the topics included in the appendix to Recommendation No. R(95) 13
- Describe the anti-bribery provisions of the Foreign Corrupt Practices Act of 1997

Outline

The lesson contains these topics:

- U.S. Privacy Laws
- U.S. Federal Privacy Act of 1974
- U.S. Federal Comprehensive Crime Control Act of 1984
- U.S. Medical Computer Crime Act of 1984
- U.S. Computer Fraud and Abuse Act of 1986
- National Information Infrastructure Protection Act of 1996
- Federal Interest Computers
- Electronic Communications Privacy Act of 1986
- Computer Security Act of 1987
- Distinguishing Between NIST and NSA Controls
- Canadian Criminal Code
- U.S. Copyright Act and Canada's Bill C-17
- European Union Laws
- Appendix to Recommendation No. R(95) 13
- Foreign Corrupt Practices Act of 1997

U.S. Privacy Laws

Privacy laws in the United States state that any private information collected on an individual must be done in a fair and lawful manner. The data collected can only be used for the purpose it was originally collected for and only for a reasonable amount of time. Also, if a company collects data on an individual, that individual has the right to receive a report outlining the data collected.

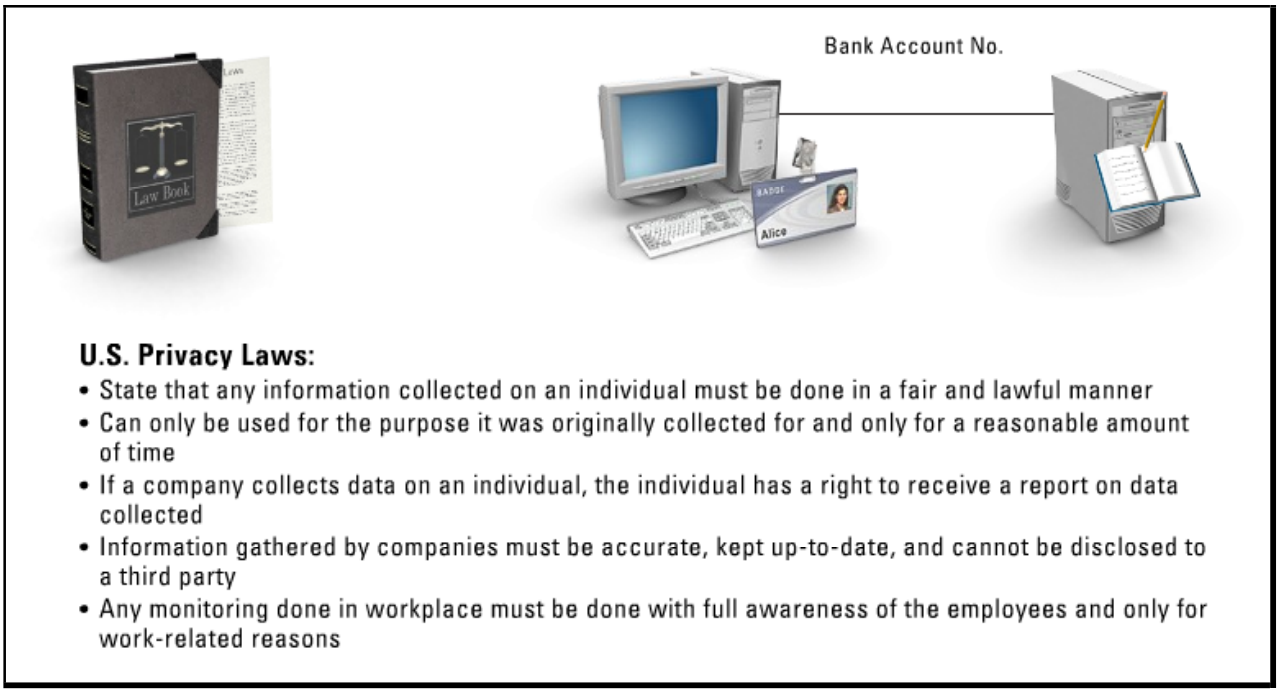

Other portions of current privacy laws state that information gathered by companies must be accurate, kept up-to-date, and cannot be disclosed to a third party unless authorized by statute or consent of the individual. If there is inaccurate information, people also have the right to make a correction to their personal information.

If any type of monitoring is put in place in the workplace, all employees must be fully aware that these surveillance practices are in use. The monitoring must only be put in place for work-related reasons, meaning that management does have the right to listen in on conversations between an employee and a customer, but not on personal conversations of the employee. Also, monitoring must happen in a consistent way, where all employees are subject to monitoring, and not just a few people.

U.S. Federal Privacy Act of 1974

The Privacy Act of 1974 states that federal agencies are required "to collect, maintain, use, or disseminate any record of identifiable personal information in a manner that assures…that adequate safeguards are provided to prevent misuse of such information."

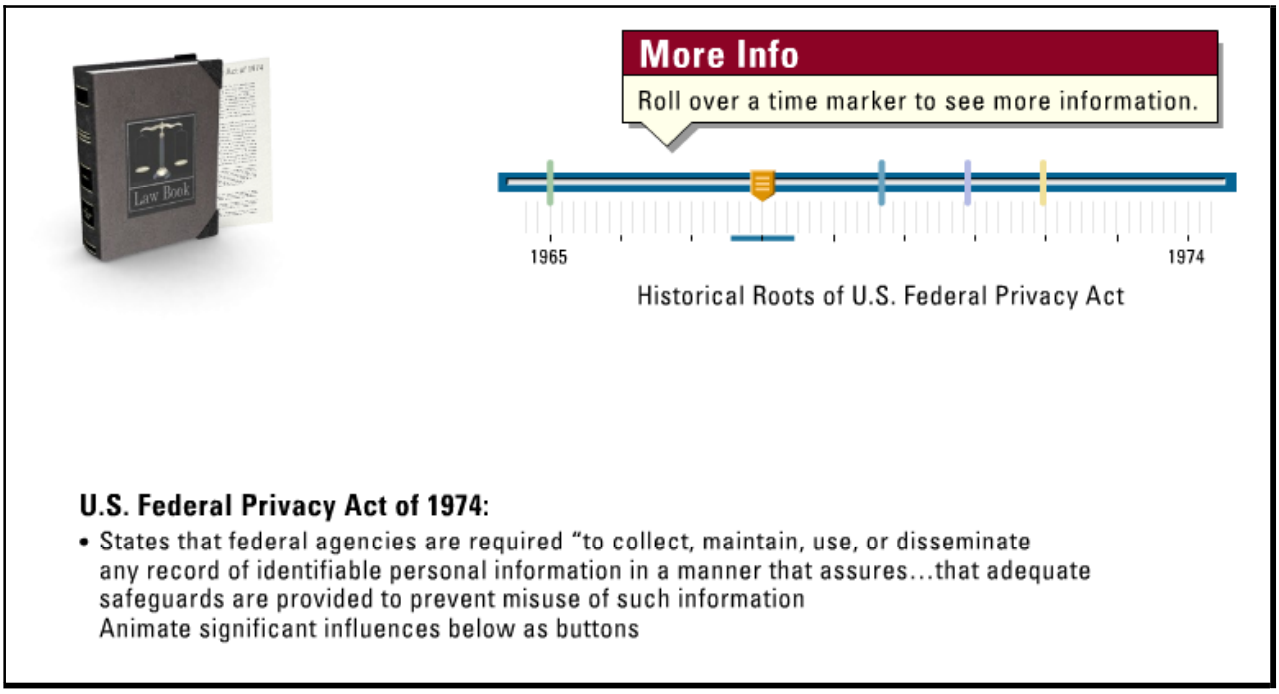

Roots of the Privacy Act of 1974 can be traced as far back as 1965 when hearings were held by the House of Representatives Special Subcommittee on Invasion of Privacy. Since then, several congressional committees have held numerous hearings and issued a number of reports on such topics as national data banks, commercial credit bureaus, and the effect of computers on personal privacy.

Significant influences in the Privacy Act include:

- "There must be no data record-keeping systems whose very existence is secret."
- "There must be a way for an individual to find out what information about him is in a record and how it is used."
- "There must be a way for an individual to prevent information about him obtained for one purpose from being used or made available for other purposes without his consent."
- "There must be a way for an individual to correct or amend a record of identifiable information about him."
- "Any organization creating, maintaining, using, or disseminating records of identifiable personal data must assure the reliability of the data for their intended use and must take reasonable precautions to prevent misuse of the data."

U.S. Federal Comprehensive Crime Control Act of 1984

This topic introduces the Comprehensive Crime Control Act of 1984.

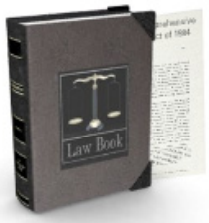

U.S. Federal Comprehensive Crime Control Act of 1984:

- Brought about major revisions to the law in many areas:
  - Bail
  - Sentencing
  - Criminal forfeiture
  - Youthful offenders
  - Treatment of offenders with mental disorders
  - The insanity defense
  - Immediate and long-range effects of officer's specific duties

The Comprehensive Crime Control Act of 1984 brought about major revisions to the law in many areas—including bail, sentencing, criminal forfeiture, youthful offenders, treatment of offenders with mental disorders, and the insanity defense—and had both immediate and long-range effects on officers' specific duties and on the overall scope of their jobs.

U.S. Medical Computer Crime Act of 1984

This topic introduces the U.S. Medical Computer Crime Act of 1984.

The U.S. Medical Computer Crime Act of 1984 was enacted to make it illegal for someone to alter any computerized medical record. Up to this point in time there was no law which prevented this act from occurring.

U.S. Computer Fraud and Abuse Act of 1986

While the development and possession of harmful computer code is not a criminal act, using the code can be. The Computer Fraud and Abuse Act (CFAA) [18 U.S.C. Section 1030] of 1986 (as amended in 1996) makes it illegal for anyone to distribute computer code or place it in the stream of commerce if the intent is to cause either damage or economic loss.

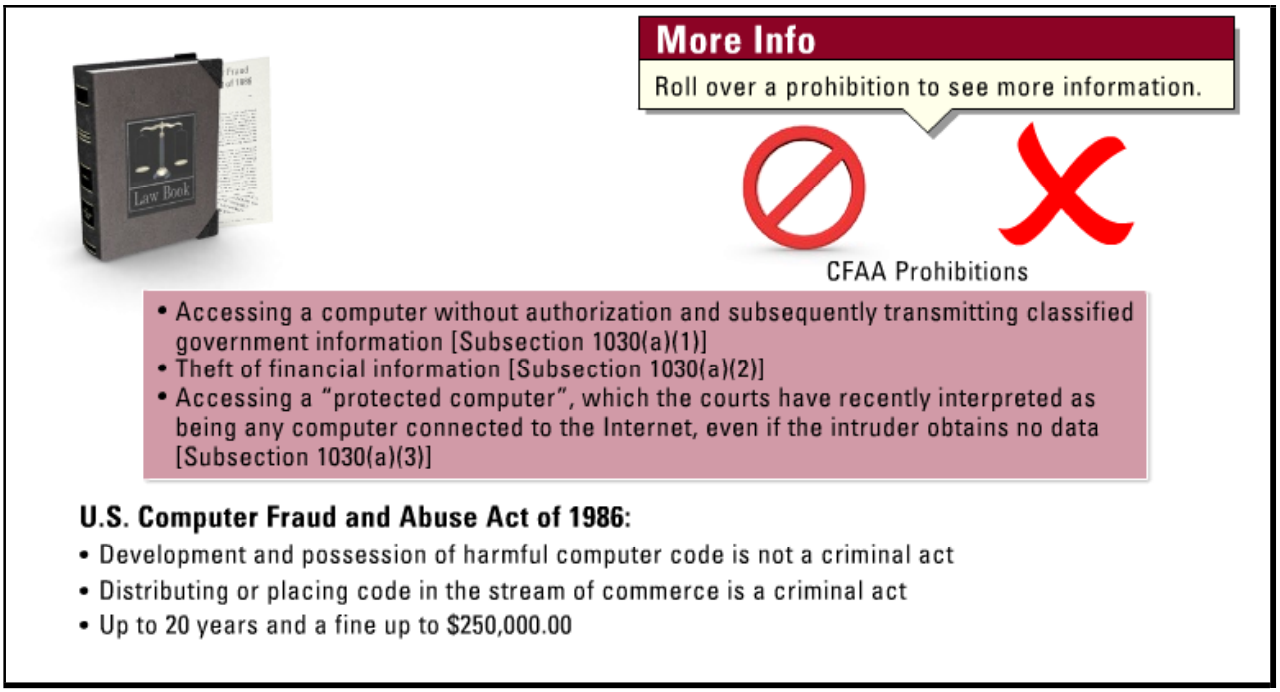

The CFAA focuses on a code's damage to computer systems and the subsequent economic losses, and it provides criminal penalties for either knowingly or recklessly releasing a computer virus into computers used in interstate commerce. Someone convicted under the CFAA could face a prison sentence as long as 20 years and a fine of up to $250,000.

When the CFAA was enacted in 1984 (as the Counterfeit Access Device and Computer Fraud and Abuse Act), it applied only to federal government computers and computers owned by large financial institutions. It was designed simply to give the Secret Service the jurisdiction to conduct investigations into computer crime. The first person prosecuted under the CFAA was Robert Morris, the Cornell University graduate student who released the first worm onto the Internet. Yet additional prosecutions were not immediately forthcoming: The unamended version of the 1984 CFAA resulted in only one prosecution. Since then, however, the act has been amended many times to counter new instances of computer crime.

In simple terms, the CFAA prohibits:

- Accessing a computer without authorization and subsequently transmitting classified government information [Subsection 1030(a)(1)]
- Theft of financial information [Subsection 1030(a)(2)]
- Accessing a "protected computer", which the courts have recently interpreted as being any computer connected to the Internet, even if the intruder obtains no data [Subsection 1030(a)(3)]
- Computer fraud [Subsection 1030(a)(4)]

- Transmitting code that causes damage to a computer system [Subsection 1030(a)(5)]
- Trafficking in computer passwords for the purpose of affecting interstate commerce or a government computer [Subsection 1030(a)(6)]
- Computer extortion [Subsection 1030(a)(7)]

Note Every state except Vermont has enacted a computer crime statute. Many of these statutes are based on the federal Computer Fraud and Abuse Act of 1986, but they vary widely in their definitions of computers, computer systems, computer networks, computer supplies, data, and other fundamental terms.

National Information Infrastructure Protection Act of 1996

The National Information Infrastructure Protection Act, which was signed into law by then-President Clinton in 1996, significantly amended the Computer Fraud and Abuse Act.

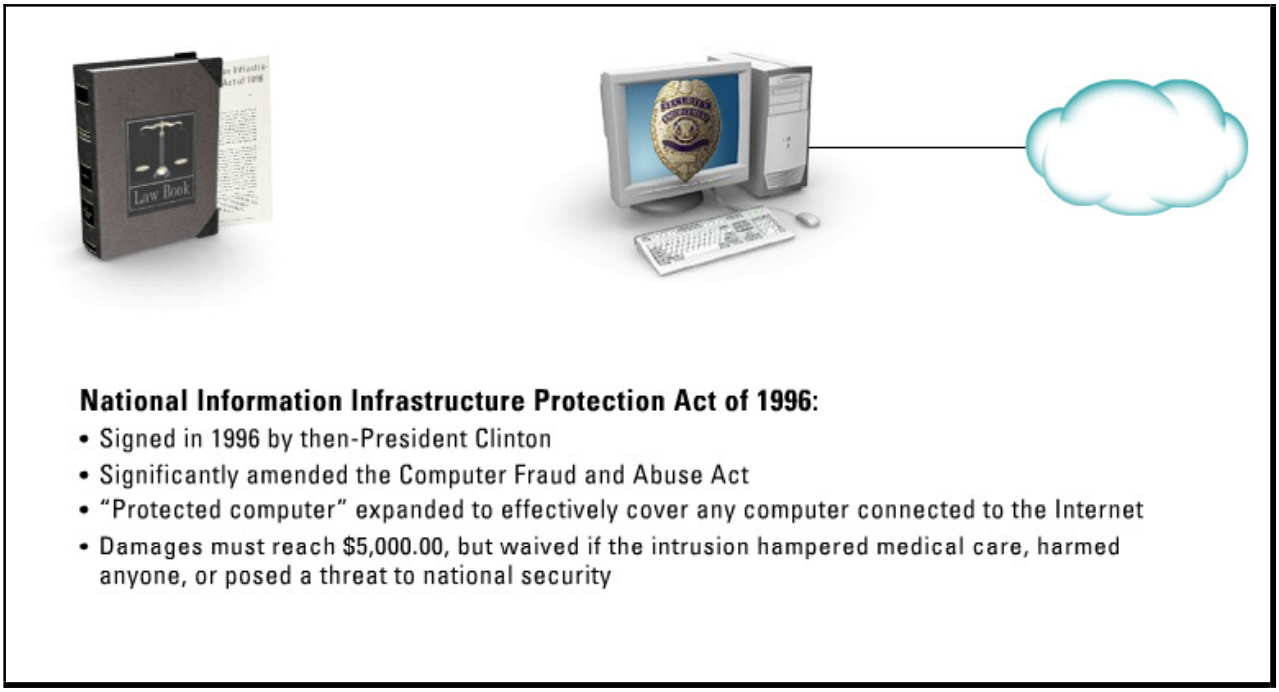

Its definition of a "protected computer" was expanded to effectively cover any computer connected to the Internet. Damages, as defined in the original act, must reach $5,000, but that requirement is waived if the intrusion hampered medical care, harmed anyone, or posed a threat to national security.

Federal Interest Computers

Federal interest computers are defined by law as two or more computers involved in a criminal offense, which are located in different states. Therefore, a commercial computer that is the victim of an intrusion coming from another state is a federal interest computer.

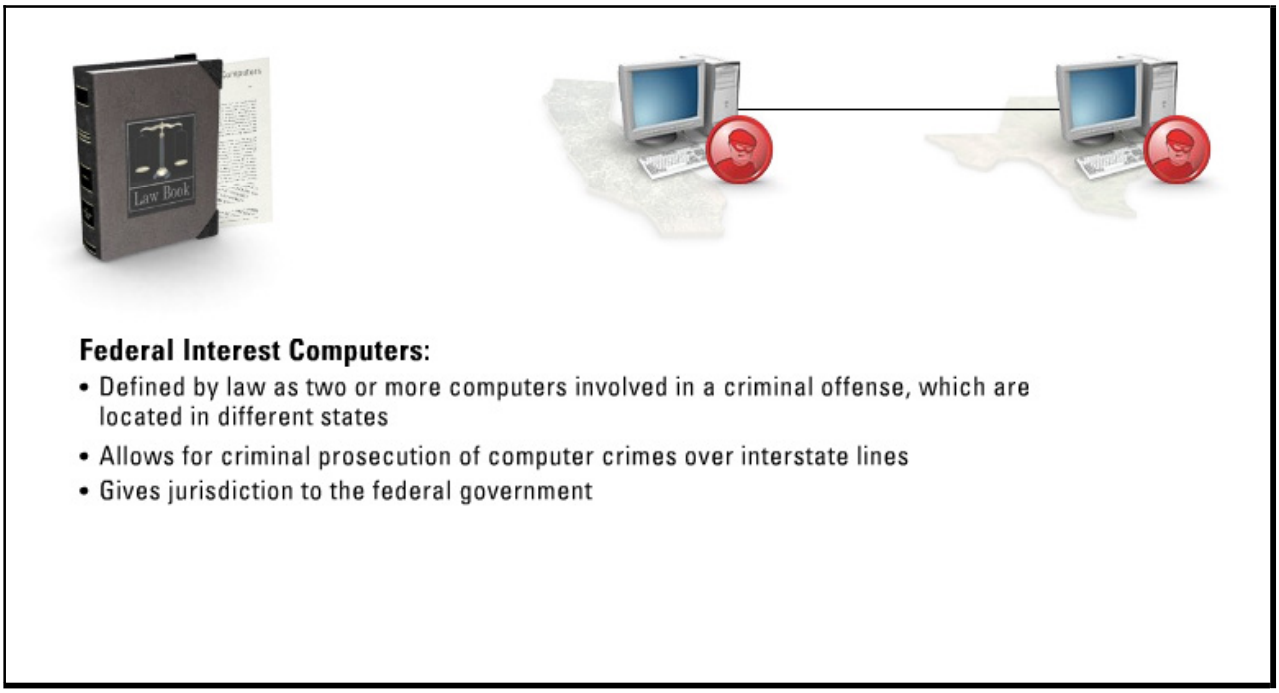

The FBI's National Computer Crime Squad (NCCS) investigates violations of the Computer Fraud and Abuse Act of 1986. These crimes cross multiple states or international boundaries. Violations of the Computer Fraud and Abuse Act include intrusions into government, financial, most medical, and federal interest computers.

Electronic Communications Privacy Act of 1986

The Electronic Communications Privacy Act (ECPA) [18 U.S.C. Sections 2510-2521, 2701-2710], which was signed into law in 1986, amended the Federal Wiretap Act to account for the increasing amount of communications and data transferred and stored on computer systems.

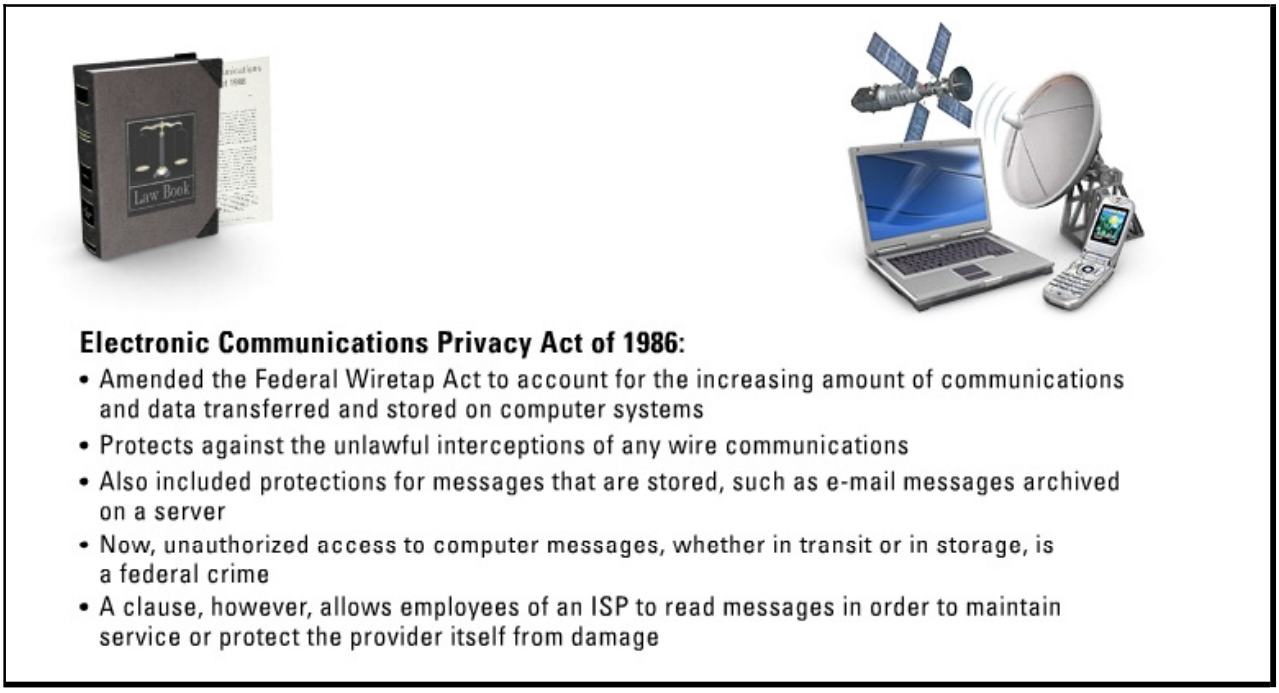

The ECPA protects against the unlawful interceptions of any wire communications—whether they are telephone or cell phone conversations, voicemail communications, e-mail communications, or any other communications in which data are sent over the wires. The ECPA also includes protections for messages that are stored, such as e-mail messages that are archived on servers. Now, under the law, unauthorized access to computer messages, whether in transit or in storage, is a federal crime.

There is a clause in the ECPA, however, that permits employees at an Internet service provider (ISP) to read the messages in order to maintain service or to protect the provider itself from damage. For example, if an ISP suspects that a virus is being disseminated via its systems, it has a right to intercept messages to determine whether its service is, indeed, a carrier of a virus.

Computer Security Act of 1987

This topic discusses the Computer Security Act of 1987.

Computer Security Act of 1987:

- Requires:
  - NIST to develop and promulgate standards and guidelines for protection of federal computer systems
  - Establishment of "security plans" for federal systems that contain "sensitive information"
  - Mandatory periodic training for all persons involved in management, use, or operation of federal systems that contain sensitive information
  - "Sensitive information" is that data that could adversely affect the national interest or the conduct of federal programs, or the privacy to which individuals are entitled under the Privacy Act

The Computer Security Act of 1987 (40 U.S. Code 759 and Public Law 100-235, Jan. 8, 1988) requires:

- National Institute of Standards and Technology (NIST) to "develop and promulgate" standards and guidelines for protection of federal computer systems
- Establishment of "security plans" for federal systems that contain "sensitive information"
- "Mandatory periodic training for all persons involved in management, use, or operation of federal systems that contain sensitive information"

The Computer Security Act also defines "sensitive information" as data that could "adversely affect the national interest or the conduct of federal programs, or the privacy to which individuals are entitled under...." the Privacy Act.

The Computer Security Act of 1987 provides the following key points:

- Requires the identification of systems that contain sensitive information
- Requires the establishment of security plans by all operators of federal computer systems that contain sensitive information
- Requires mandatory periodic training in computer security awareness and accepted computer security practices for all persons involved in management, use, or operation of federal computer systems that contain sensitive information
- Requires the NIST to establish a Computer Standards Program
  - The primary purpose of the program is to develop standards and guidelines to control loss and unauthorized modification or disclosure of sensitive information in systems and to prevent computer-related fraud and misuse

- Requires the establishment of a Computer System Security and Privacy Advisory Board within the Department of Commerce
    - The duties of the Board shall be: (1) to identify emerging managerial, technical, administrative, and physical safeguard issues relative to computer systems security and privacy; (2) to advise NIST and the Secretary of Commerce on security and privacy issues pertaining to federal computer systems; and (3) to report its findings to the Secretary of Commerce, the Director of the Office of Management and Budget, the Director of the National Security Agency, and the appropriate Committees of the Congress

---

Note The international nature of the Internet means that any attempt to deal with Internet-related crime will always be complicated by questions of jurisdiction. Laws vary from country to country and U.S. police have no power to intervene directly against criminal material on computers in another country or against criminals operating in another country.

Distinguishing Between NIST and NSA Controls

The NIST provides recommended security controls for federal information systems. It outlines electronic and physical controls for systems categorized under three levels of potential impacts, such as what would happen if someone steals information from a federal system and modifies the data or disrupts a government service.

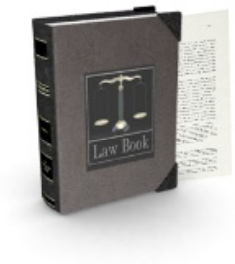

NIST:

- Provides recommended security controls for federal information systems
- Outlines electronic and physical controls for systems categorized under three levels of potential impacts (low, medium, and high impacts)
- Three classes of controls (management, operational, and technical)

NSA:

- The Reagan directive gave NSA control over all government computer systems containing "sensitive but unclassified" information
- Since enactment of the Computer Security Act, the NSA has sought to undercut NIST's authority

Low-, medium-, and high-impact levels are defined in the draft "Federal Information Processing Standard (FIPS) 199: Standards for Security Categorization of Federal Information and Information Systems."

Controls outlined in the publication fall into three classes—management, operational, and technical—and are then broken down further into families. For example, under the management class, families include security planning and acquisition of information systems and services. Operational class families focus on issues such as incident response and contingency planning and operations.

In 1987, the U.S. Congress, led by Rep. Jack Brooks, enacted a law reaffirming that the NIST, a division of the Department of Commerce, was responsible for the security of unclassified, non-military government computer systems. Under the law, the role of the National Security Agency (NSA) was limited to providing technical assistance in the civilian security realm. Congress rightly felt that it was inappropriate for a military intelligence agency to have control over the dissemination of unclassified information.

The law was enacted after President Reagan issued the controversial National Security Decision Directive (NSDD) 145 in 1984. The Reagan directive gave NSA control over all government computer systems containing "sensitive but unclassified" information. This directive was followed by a second directive issued by National Security Advisor John Poindexter that extended NSA authority over non-government computer systems.

Since the enactment of the Computer Security Act, the NSA has sought to undercut NIST's authority. In 1989, NSA signed a Memorandum of Understanding (MOU), which purported to transfer back to NSA

the authority given to NIST. The MOU created a NIST/NSA technical working group, which then developed the controversial Clipper Chip and Digital Signature Standard. The NSA has also worked in other ways to weaken the mandate of the Computer Security Act. In 1994, then-President Clinton issued the Presidential Decision Directive (PDD) 29. This directive created the Security Policy Board, which has recommended that all computer security functions for the government be merged under NSA control.

Canadian Criminal Code

Section 342.1 of the Criminal Code of Canada is part of a series of new "high tech" crimes that were introduced a few years back. The Criminal Code was also amended in order to expand the definition of "mischief" (see s. 430) to include anyone who willfully obstructs, interrupts, interferes with, alters, or destroys data.

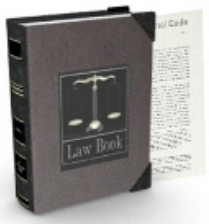

Canadian Criminal Code:

- Section 342.1 amended the Criminal Code in order to expand definition of "mischief" to include anyone who willfully obstructs, interrupts, interferes with, alters, or destroys data
- Purpose was to prohibit anyone from using a computer system "fraudulently and without colour of right"
- A law that attempts to control computer 'hackers' – But, the term hacker generally has two meanings:
    - Anyone who likes to fiddle around with computers and their software
    - A person who breaks into computer systems
- May have some very serious implications for Canadians

The purpose of Section 342.1 was, amongst other things, to prohibit anyone from using a computer system "fraudulently and without colour of right". In other words, if you use a computer system when you are not supposed to, or in a way that you know you are not supposed to, you may be committing an offense.

In creating a new crime category that prohibits the unauthorized use of a computer system, the Canadian government was presumably trying to pass a law that would allow the police to control computer hackers.

The term "hacker" is generally held to mean one of two different things:

- Anyone who likes to fiddle around (a technical term) with computers and their software; or
- A person who breaks into computer systems.

Unfortunately, Section 342.1 of the Criminal Code does not draw such a fine distinction. According to the law, if you use a computer system that you were not supposed to, and you know it, then you are guilty of an offense and could be liable for imprisonment "for a term not exceeding 10 years". But the law's clear-cut distinction between authorized and unauthorized use of a computer system may have some very serious implications for Canadians everywhere. That is because many of the service contracts that Canadians enter into nowadays contain language that limits their right to transfer or assign the use of the service to another person.

Section 1.1 defines that everyone commits mischief who willfully

- Destroys or alters data

- Renders data meaningless, useless, or ineffective
- Obstructs, interrupts, or interferes with the lawful use of data
- Obstructs, interrupts, or interferes with any person in the lawful use of data or denies access to data to any person who is entitled to access thereto

U.S. Copyright Act and Canada ’s Bill C-17

This topic discusses the U.S. Copyright Act and Canada’s Bill C-17.

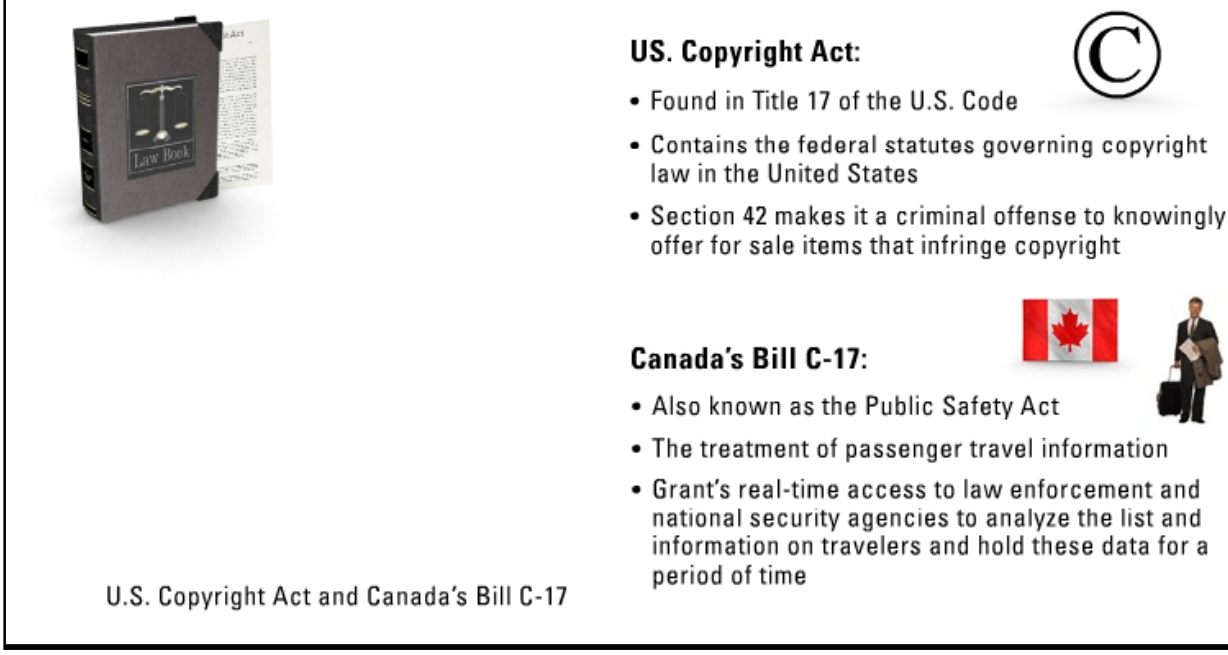

The U.S. Copyright Act is found in Title 17 of the U.S. Code and contains the federal statutes governing copyright law in the United States.

Section 42 addresses the importation of merchandise bearing trademarks that “copy or simulate” a U.S.-owned mark. However, it does not specify whether a foreign trademark legitimately applied to genuine merchandise “copies or simulates” the identical trademark in use in the United States. The rule to be deduced is that a U.S. trademark holder may rely on Section 42 to exclude parallel importations when the goods offered for sale abroad are materially different from those offered in the United States. In general terms, this section makes it a criminal offense to knowingly offer for sale items that infringe copyright.

The key issue on Canada’s Bill C-17, also known as the Public Safety Act, is the treatment of passenger travel information. The bill proposes to grant real-time access to law enforcement and national security agencies to analyze the lists and information on travelers and hold these data for a period of time. The Privacy Commissioner of Canada has protested continuously against these measures. In particular, the Commissioner argued that “the police have no business using this extraordinary access to personal information to search for people wanted on warrants for any offenses unrelated to terrorism.”

European Union Laws

The European Union has taken a number of steps to fight harmful and illegal content on the Internet. In April 1998, The European Commission presented to the European Council the results of a study on computer-related crime (the so-called COMCRIME study).

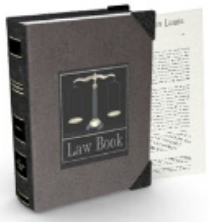

European Union Laws:

- The EU has taken a number of steps to fight harmful and illegal content on the Internet
- April 1998, the European Commission presented to the European Council a study on computer-related crime
- October 1999, the European Council concluded the high-tech crime should be included in the efforts to agree on common definitions and sanctions
- Recommendation No. R(89) 9 contains a minimum list of offenses necessary for a uniform criminal policy concerning computer-related crime
- November 23, 2001 ministers from 26 member countries together with Canada, Japan, South Africa, and the United States signed the treaty

In October 1999, the Tampere Summit of the European Council concluded that high-tech crime should be included in the efforts to agree on common definitions and sanctions. The European Parliament has also called for commonly acceptable definitions of computer-related offenses and for effective approximation of legislation, in particular in substantive criminal law. The Council of the European Union has adopted a Common Position on a number of initial elements as part of the Union's strategy against high-tech crime.

Recommendation No. R(89) 9 was adopted by the Council of Europe on September 13, 1989. It contains a minimum list of offenses necessary for a uniform criminal policy on legislation concerning computer-related crime.

The Council of Europe adopted another Recommendation concerning problems of procedural law connected with information technology on September 11, 1995.

The Council of Europe appointed a Committee of Experts on Crime in Cyberspace (PC-CY) in 1997 to identify and define new crimes, jurisdictional rights, and criminal liabilities due to communication on the Internet.

Canada, Japan, South Africa, and the United States were invited to meet with experts at the committee meetings and participate in the negotiations. The Ministers of Foreign Affairs finally adopted the Convention on November 8, 2001. It was open for signatures at a meeting in Budapest, Hungary, on November 23, 2001. Ministers or their representative from 26 member countries together with Canada, Japan, South Africa, and the United States signed the treaty.

Appendix to Recommendation No. R(95) 13

This topic discusses the appendix to Recommendation No. R(95) 13.

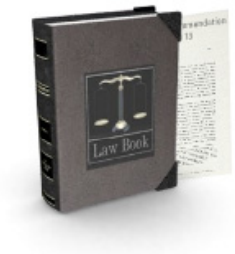

Appendix to Recommendation No. R(95) 13:

- Concerns problems of criminal procedure law connected with information technology, including:
  - Search and seizure
  - Technical surveillance
  - Obligations to cooperate with the investigating authorities
  - Electronic evidence
  - Use of encryption
  - Research, statistics, and training
  - International cooperation

The appendix to Recommendation No. R(95) 13, which concerns problems of criminal procedure law connected with information technology, includes the following:

I. Search and seizure

II. Technical surveillance

III. Obligations to cooperate with the investigating authorities

IV. Electronic evidence

V. Use of encryption

VI. Research, statistics, and training

VII. International cooperation

Foreign Corrupt Practices Act of 1997

The Foreign Corrupt Practices Act (FCPA) was enacted in 1977 and substantially revised in 1988. The provisions of the FCPA prohibit the bribery of foreign government officials by U.S. persons and prescribe accounting and record-keeping practices. The anti-bribery provisions of the FCPA are enforced by the Department of Justice (DOJ).

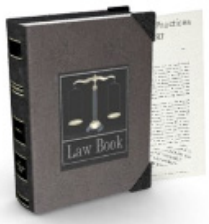

Foreign Corrupt Practices Act of 1997:

- Prohibit the bribery of foreign government officials by U.S. persons
  - Prescribe accounting and record-keeping practices
  - Anti-bribery provisions enforced by the Department of Justice
  - Apply to a U.S. person and make it illegal for U.S. persons to bribe a foreign government official for the purposes of obtaining or retaining business
  - Penalties vary based on whether the violator is a U.S. company or a U.S. individual
  - Companies can be fined up to $2 million
  - Individuals can be fined up to $100,000 and up to five years in prison

The anti-bribery provisions of the FCPA apply to any U.S. person and make it illegal for U.S. persons to bribe a foreign government official for the purpose of obtaining or retaining business. The wording of the FCPA is quite interesting and makes its scope rather clear.

Penalties for violating the anti-bribery provisions of the FCPA vary based on whether the violator is a U.S. company or a U.S. individual. U.S. companies can be fined up to $2 million while U.S. individuals (including officers and directors of companies that have willfully violated the FCPA) can be fined up to $100,000 and imprisoned for up to five years, or both. In addition, civil penalties may be imposed.

Summary

The key computer crime-related laws discussed in this lesson are:

- U.S. Privacy Laws
- U.S. Federal Privacy Act of 1974
- U.S. Federal Comprehensive Crime Control Act of 1984
- U.S. Medical Computer Crime Act of 1984
- U.S. Computer Fraud and Abuse Act of 1986
- National Information Infrastructure Protection Act of 1996
- Electronic Communications Privacy Act of 1986
- Computer Security Act of 1987
- Canadian Criminal Code
- U.S. Copyright Act and Canada's Bill C-17
- European Union Laws
- Appendix to Recommendation No. R(95) 13
- Foreign Corrupt Practices Act of 1997

Due Care

Overview

The most recognized definition of due care is “the care that a reasonable person would exercise under the circumstances; the standard for determining legal duty.” In the security context and for the security specialist who has a higher understanding of security attacks than the ordinary user, one can further define due care as “the cognizant exclusion or commission of actions that a reasonable person would expect will result in legal liability.” If you are a security professional in charge of caring for the sensitive data of your enterprise, you will be held at a much higher standard of due care.

Importance

It is important for the information security specialist to understand his or her responsibilities under the laws of due care.

Objectives

Upon completing this lesson, you will be able to:

- Define due care and due diligence
- Identify the legal obligations of a corporate director or officer
- Describe the evidence life cycle
- Define the common forms of admissible evidence
- Define Chain of Evidence
- Identify general guidelines for working with law enforcement

Outline

The lesson contains these topics:

- Characteristics of Due Care
- Liabilities of Corporate Officers
- The Evidence Life Cycle
- Characteristics of Admissible Evidence
- Chain of Evidence

- Contacting Law Enforcement about Suspected Computer Crimes

Characteristics of Due Care

Senior management of any company has the responsibility of protecting the company from a long list of activities that can negatively affect it, including protecting it from malicious code, natural disasters, intruders, and much more.

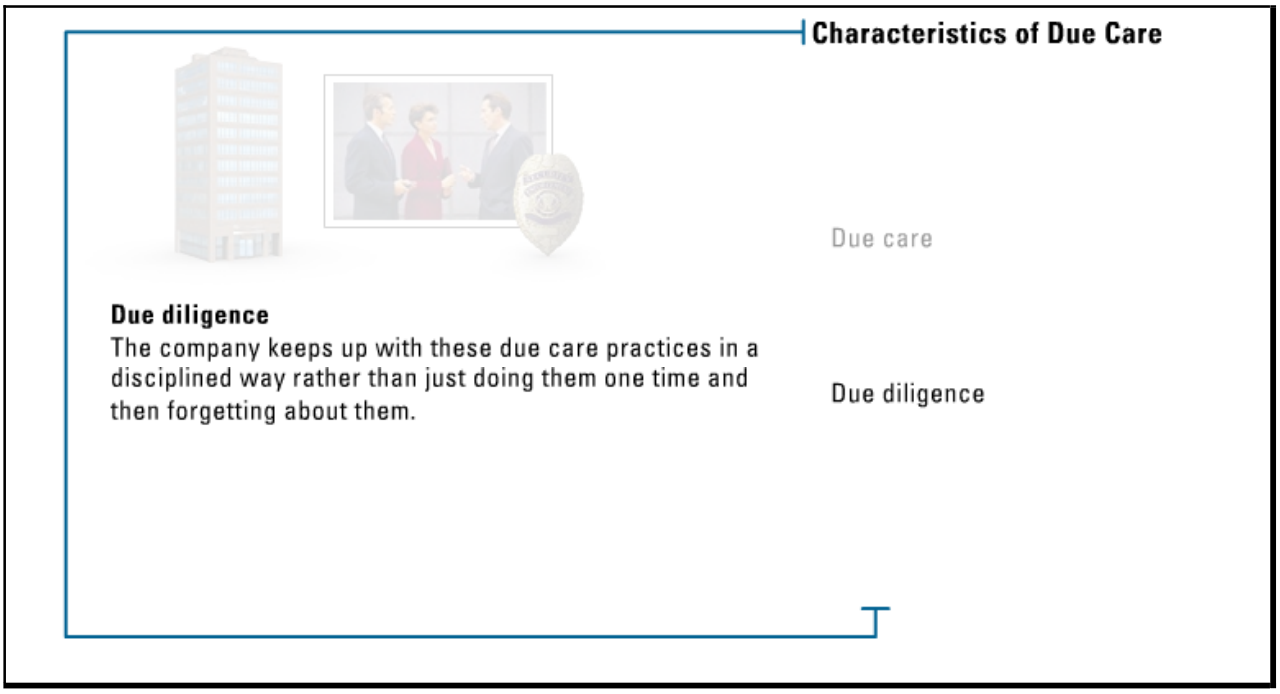

In order for senior management to hold to their responsibility, they must meet the following standards:

- Due care
- Due diligence

Due care basically states that a company practices common sense and prudent management concepts and acts responsibly. Due diligence basically states that the company keeps up with these due care practices in a disciplined way rather than just doing them one time and then forgetting about them.

In addition, management must follow the prudent man rule, which requires management to perform any duty that a prudent person would exercise in a similar circumstance. Management must perform these duties with diligence and care. The following is a list of actions required to show due care:

- Conducting background checks on potential employees
- Physical and logical access control
- Network security, which might require encryption
- Performing regular backups
- Disaster recovery and business continuity planning
- Conducting periodic reviews to strengthen disaster recovery plans
- Disseminating proper information to employees of expected behavior
- Developing security policies, standards, procedures, and guidelines

- Conducting security awareness training
- Keeping antivirus software up-to-date
- Conducting periodic penetration testing
- Keeping service level agreements (SLAs) current
- Implementing measures to ensure that software piracy is not taking place
- Ensuring proper auditing and log reviews are taking place

Caution	If a company or corporation does not practice due care pertaining to the security of its assets or customer information, it can be legally charged with negligence and held accountable for any ramifications of that negligence.

Liabilities of Corporate Officers

When a candidate chooses to take a job as a corporate director or officer at a particular company, his or her potential exposure to personal liability is significantly affected. After all, in recent years, personal liability for directors and officers has expanded into issues of wages, taxes, employee actions, environmental damage, regulatory matters, and disclosure or misrepresentation of company information.

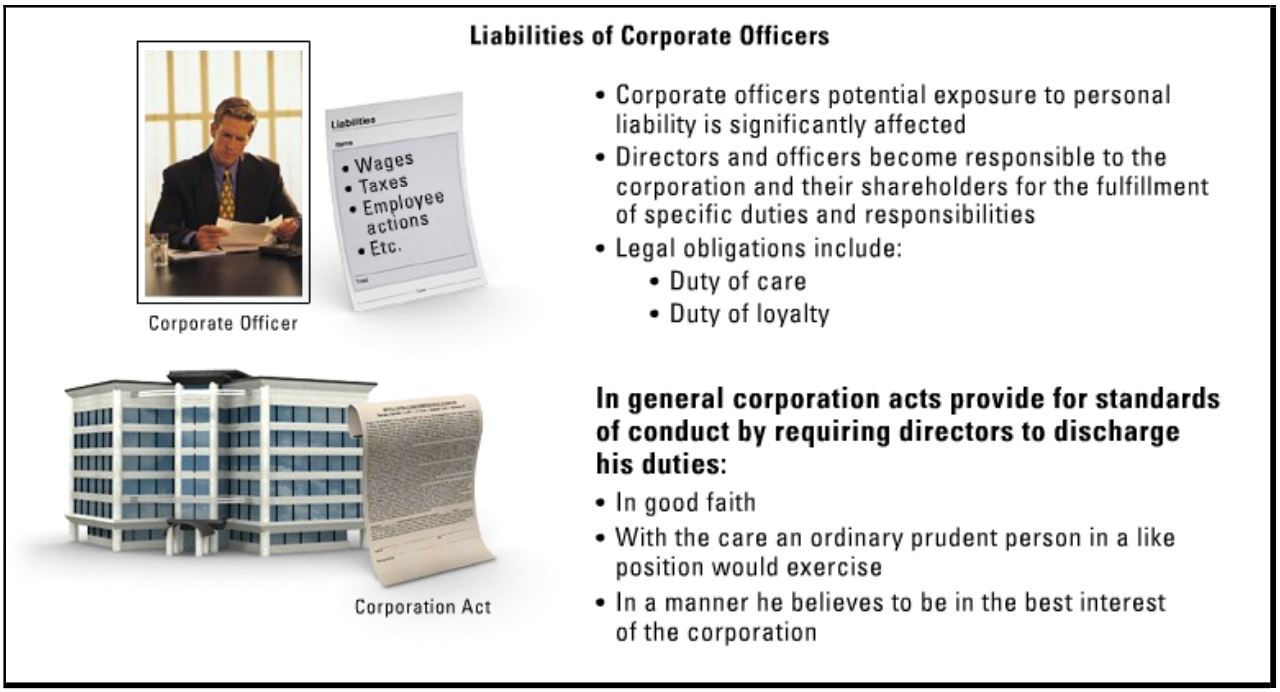

When a person agrees to serve as a director or an officer of a corporation, the person becomes responsible to the corporation and the shareholders for the fulfillment of specific duties and responsibilities. The legal obligations of a director or an officer fall into two broad categories: (1) the duty of care; and (2) the duty of loyalty.

In general, corporation acts provide for standards of conduct for directors by requiring that a director shall discharge his or her duties as a director:

- In good faith;
- With the care an ordinarily prudent person in a like position would exercise under similar circumstances; and
- In a manner he or she reasonably believes to be in the best interests of the corporation.

In discharging his or her duties, a director is entitled to rely on information, opinions, reports, or statements, including financial statements and data, if prepared or presented by a person the director reasonably believes to be reliable and competent.

The “good faith” requirement demands the director always discharge duties honestly, conscientiously, fairly, and with undivided loyalty to the corporation. This requirement also demands the director manage the corporation in the best interests of all the shareholders. The duty of good faith is embodied in the duty of care and in the duty of loyalty and overlays an honesty requirement on both.

The Federal Sentencing Guidelines were extended in 1997 to cover company computer crimes and specified that senior corporate officers could be personally responsible to pay up to $290 million in fines if their company did not comply with the due care laws set out for them.

The Evidence Life Cycle

To begin the evidence life cycle, you need to start with the discovery and collection of evidence.

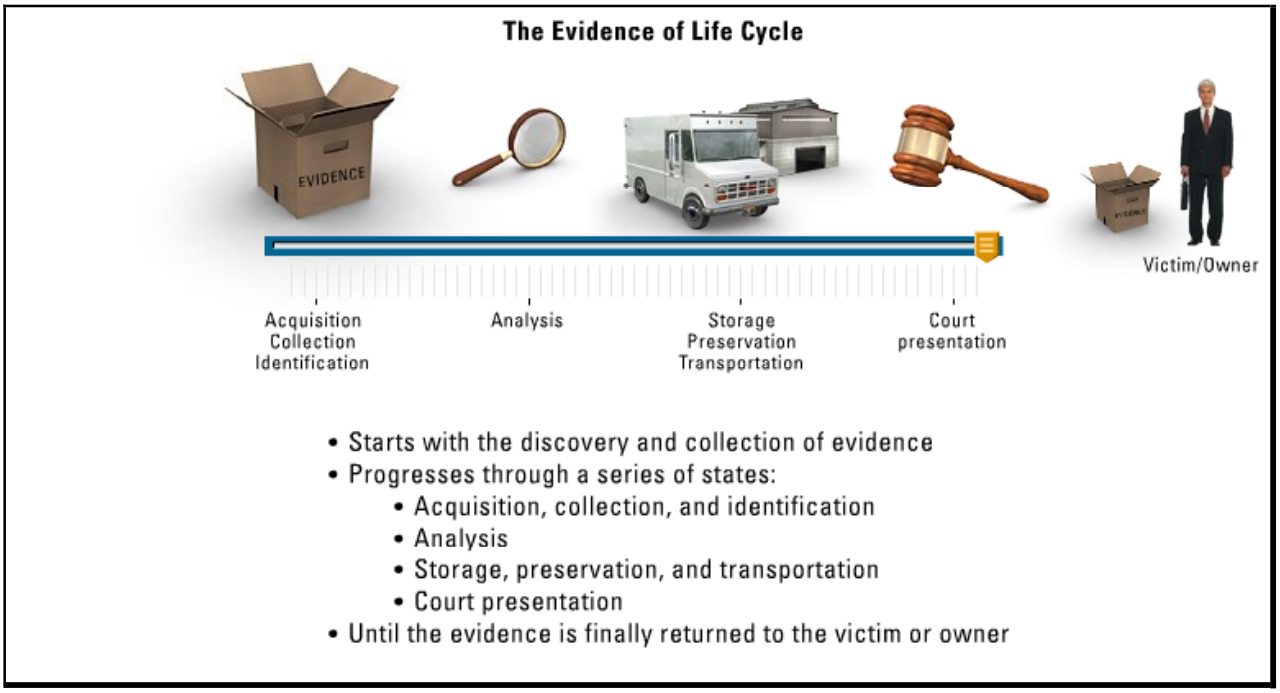

You will then progress through the following series of states until the evidence is finally returned to the victim or owner:

Acquisition, Collection, and Identification

- Properly mark the evidence you obtained or collected
  - Be extremely careful not to damage evidence when marking it
- Record evidence in a logbook
- Seal the evidence
  - Seal the evidence in a container with evidence tape

Analysis

- Comprehensively examine the evidence
- Can yield quality evidence that can be considered reliable in a court of law

Storage, Preservation, and Transportation

- Pack and preserve evidence to prevent contamination
- If not properly protected, the person or agency could be held liable
- Transport evidence to a stored and locked location

Court Presentation

- Present each piece of evidence in court
- Each time you transport evidence to and from the courthouse, you must handle it with proper care

- Continue to follow the Chain of Evidence

Return to Victim/Owner

- Once the trial is over, return the evidence to the victim/owner

Characteristics of Admissible Evidence

In a court of law, there are many types of evidence that can be offered to prove the truth or falsity of a given fact.

Characteristics of Admissible Evidence

In a court of law, many types of evidence can be offered to prove the truth of falsity of a given fact

The most common forms of evidence include:

- Direct evidence
- Real evidence
- Documentary evidence
- Demonstrative evidence

Hearsay rule:

- Prohibits officials from charging someone with a crime based on what someone who wishes to remain anonymous and will not come forward to testify told the official

The most common forms of evidence include: direct, real, documentary, and demonstrative. Direct testimony is that evidence given by an oral witness. Direct evidence is usually called to prove a specific act. Real evidence is made up of tangible objects that prove or disprove guilt. Physical crime evidence includes such things as the tools used in a crime, fruits of the crime, perishable evidence capable of reproduction, etc. Physical evidence is usually used to link the suspect to the scene of the crime. Documentary evidence is usually presented to a court of law in the form of business records, manuals, printouts, etc. Demonstrative evidence is usually used to aid the jury. It may be in the form of a model, an experiment, a chart, or an illustration.

Tip	Much of the evidence provided in computer crimes comes in the form of documentary evidence.

The hearsay rule prohibits a law enforcement official from charging someone with a crime based on what someone who wishes to remain anonymous and will not come forward to testify told the official. Hearsay evidence is secondhand evidence, which is not normally admissible in a court of law, unless it has firsthand evidence that can be used to provide the evidence's accuracy, trustworthiness, or reliability.

Chain of Evidence

Chain of Evidence provides a means of accountability and must be adhered to by law enforcement or other investigators when conducting any type of criminal investigation, including a computer crime investigation.

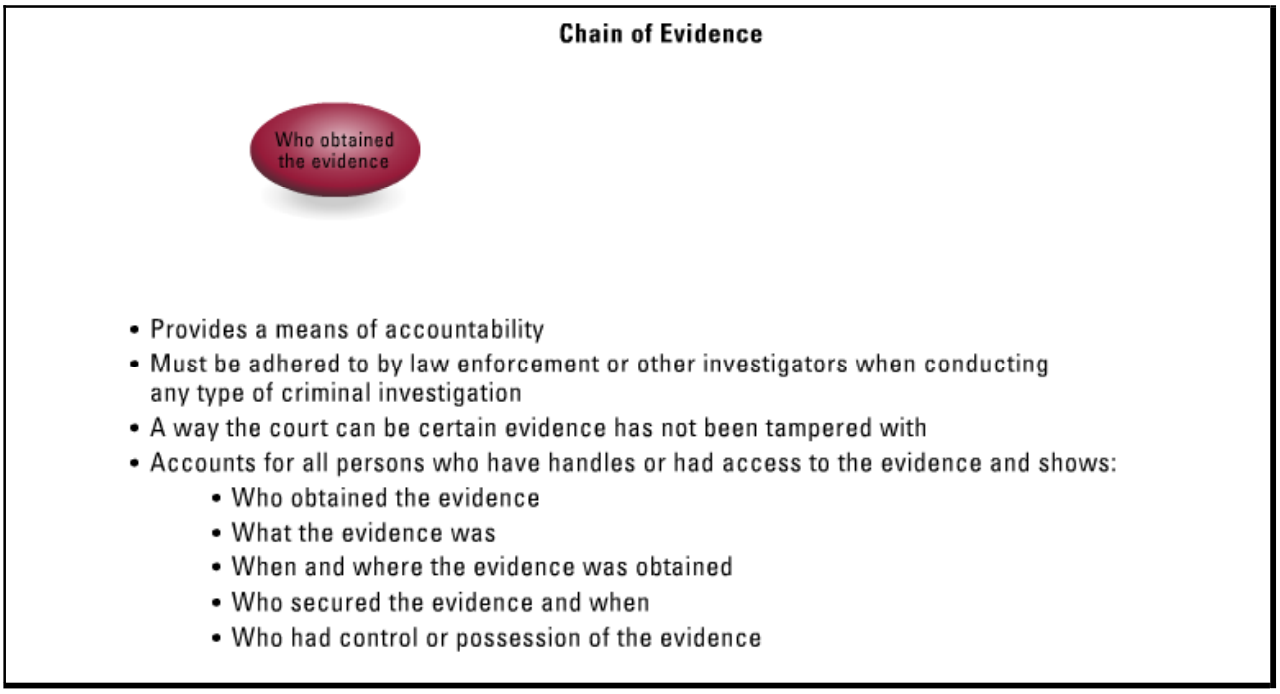

The Chain of Evidence is a way the court can be certain that the evidence has not been tampered with. The Chain of Evidence must account for all persons who have handled or had access to the evidence and shows:

- Who obtained the evidence
- What the evidence was
- When and where the evidence was obtained
- Who secured the evidence and when
- Who had control or possession of the evidence

Contacting Law Enforcement about Suspected Computer Crimes

Corporate counsel or management generally makes the decision to contact law enforcement, but the information security professional should handle the actual relationship with law enforcement.

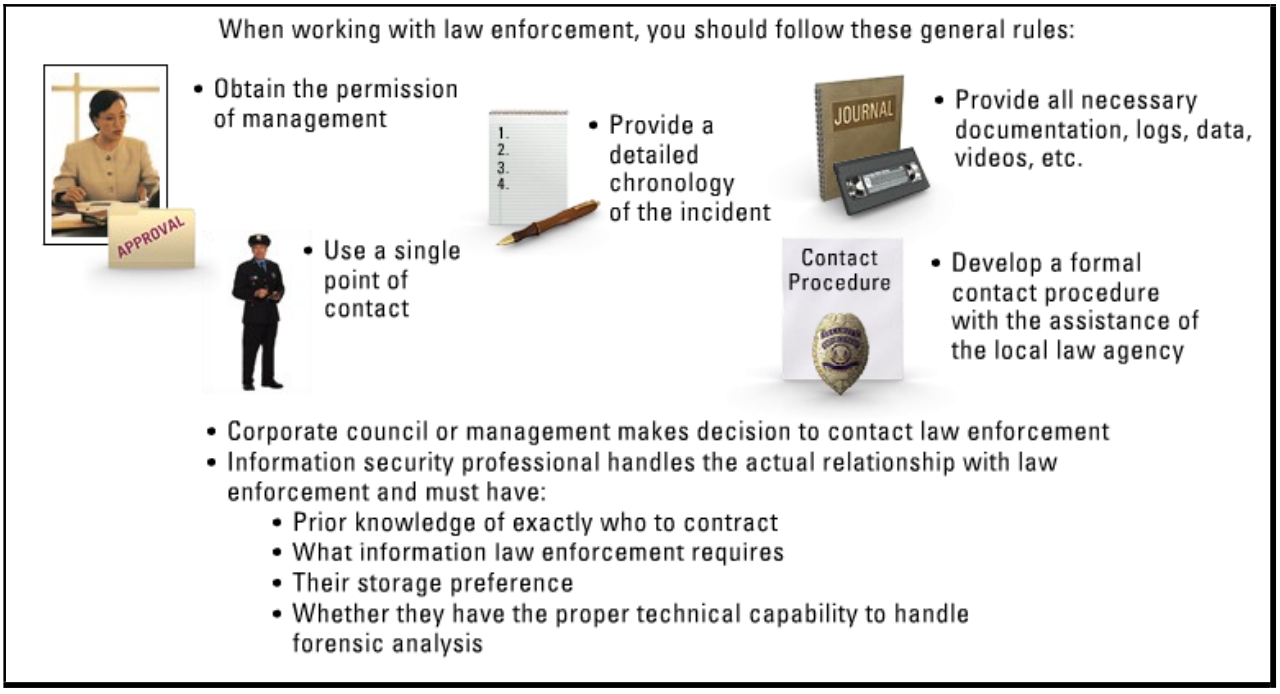

The information security professional must have prior knowledge of exactly who to contact, what information law enforcement requires, their media storage preferences, and whether they have the proper technical capability to handle forensic analysis. Having this prior knowledge can greatly smooth out the entire process and eliminate unnecessary delays.

When working with law enforcement, you should follow these general rules:

- Obtain the permission of management
- Use a single point of contact
- Provide a detailed chronology of the incident
- Provide all necessary documentation, logs, data, videos, etc.
- Develop a formal contact procedure with the assistance of the local law agency

Summary

The key points discussed in this lesson are:

- Due care basically states that a company practices common sense and prudent management concepts and acts responsibly. Due diligence basically states that the company keeps up with these due care practices in a disciplined way rather than just doing them one time and then forgetting about them.
- The legal obligations of a director or an officer fall into two broad categories: (1) the duty of care; and (2) the duty of loyalty.
- To begin the evidence life cycle, you need to start with the discovery and collection of evidence. You will then progress through a series of states until the evidence is finally returned to the victim or owner.
- The most common forms of evidence include: direct, real, documentary, and demonstrative.
- The Chain of Evidence is a way the court can be certain that the evidence has not been tampered with. The Chain of Evidence must account for all persons who have handled or had access to the evidence.
- The information security professional must have prior knowledge of exactly who to contact, what information law enforcement requires, their media storage preferences, and whether they have the proper technical capability to handle forensic analysis. Having this prior knowledge can greatly smooth out the entire process and eliminate unnecessary delays.

Investigation and Ethics

Overview

Investigating computer crimes can sometimes be a very difficult task to accomplish. Determining what you can do ethically and lawfully to determine who is perpetrating a crime can be very complex at times. This lesson will discuss how to properly investigate computer crime and the code of ethics you must follow when doing so.

Importance

It is important for the information security specialist to understand investigation techniques and the ethical bounds that he or she must follow when attempting to determine who is committing a computer crime.

Objectives

Upon completing this lesson, you will be able to:

- Identify the benefits of surveillance mechanisms
- Describe the U.S. warrant process
- Define enticement and entrapment
- Identify search and seizure rules and procedures
- Identify characteristics of good security incident procedures
- Identify guidelines for minimizing threats to data integrity
- Define computer ethics
- Identify the ethical considerations of the International Information Systems Security Certification Consortium and the Internet Architecture Board
- Identify common ethical fallacies in the computing world
- Identify the responsibilities of the Internet Architecture Board
- Define competitive intelligence and industrial espionage

Outline

The lesson contains these topics:

- Types of Surveillance
- The U.S. Warrant Process
- Enticement vs. Entrapment
- Search and Seizure Rules and Procedures
- What Constitutes a Security Incident?
- Guidelines for Data Integrity and Retention
- Computer Ethics
- Professional Code of Ethics
- Common Ethical Fallacies
- Internet Architecture Board
- Competitive Intelligence vs. Industrial Espionage

Types of Surveillance

Intruders will do whatever they can to not be seen when perpetrating their crimes. To protect a certain location or space, you can use surveillance mechanisms.

Types of Surveillance:

- Intruders do whatever they can not to be seen when perpetrating their crime
- High visible environments increase the likelihood of criminals being observed and reported
- Maximizing visibility can discourage crime:
  - Implementing physical features
  - Placing personnel in ways that maximize the ability to see
  - Using formal surveillance:
    - CCTV
    - Security patrols

Environments that are highly visible can increase the likelihood of criminal acts being observed and reported. Implementing physical features and placing personnel in ways that maximize the ability to see what is occurring can discourage crime.

You can also use formal surveillance, such as closed circuit television (CCTV) or organized security patrols, as an additional deterrent.

The U.S. Warrant Process

Court officials generally issue warrants and give them to law enforcement personnel to enter into databases. (However, there are several municipal, associate, and circuit courts that do not use law enforcement personnel to enter warrants.) Thus, the law enforcement agency is given the responsibility of locating the person and bringing that person to court.

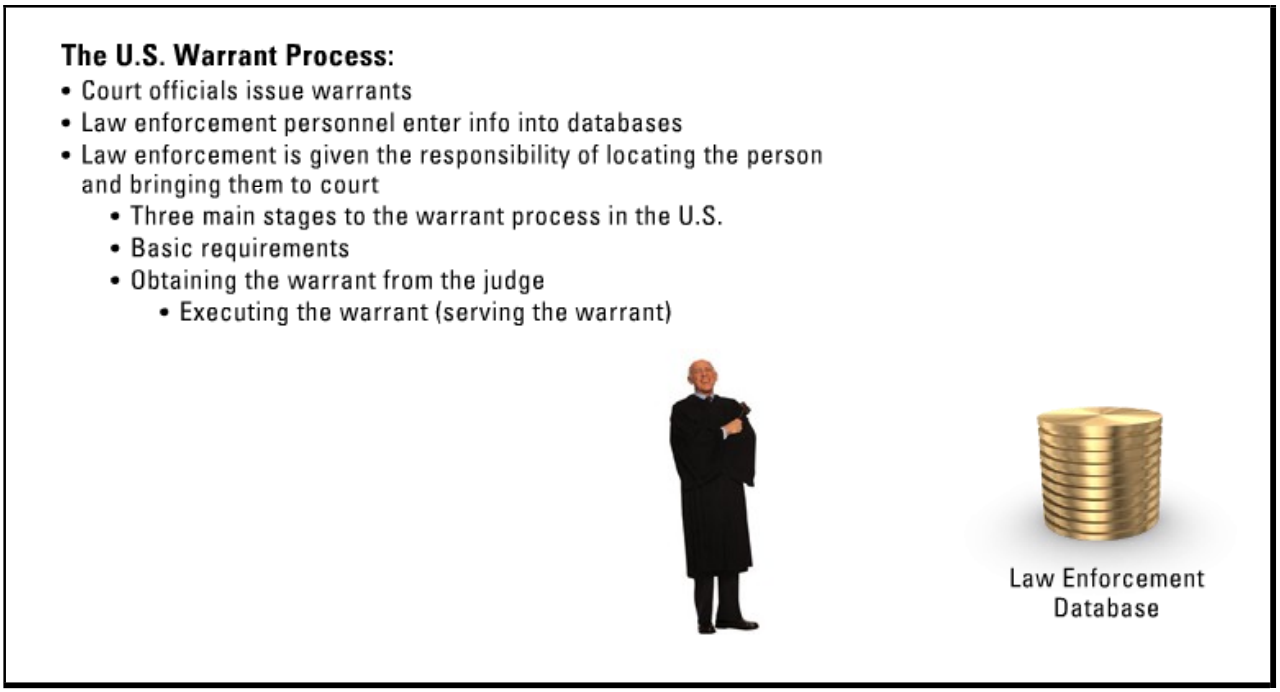

The warrant process in the United States includes three main stages:

- Basic requirements
- Obtaining the warrant from the judge
- Executing the warrant (serving the warrant)

During the basic requirements stage, it is determined if a warrant should in fact be issued. Government agencies issue a warrant to apprehend persons for:

- Committing crimes
- Failing to appear in a court
- Fleeing prosecution
- Violating probation
- Escaping prison

Enticement vs. Entrapment

This topic defines enticement and entrapment.

Entrapment encourages someone to commit a crime that the individual may have had no intention of committing. Conversely, **enticement** lures someone toward some evidence after that individual has already committed a crime.

Enticement is legal and ethical, whereas entrapment is neither legal nor ethical. To make matters more confusing, there is a fine line in the distinction between the two.

Note	Enticement is not necessarily illegal but does raise ethical arguments and may not be admissible in court.

Although entrapment is a defense to criminal charges, research has revealed no case holding that there is a federal constitutional right to be free from entrapment.

Search and Seizure Rules and Procedures

The Fourth Amendment protects against unreasonable searches or seizures.

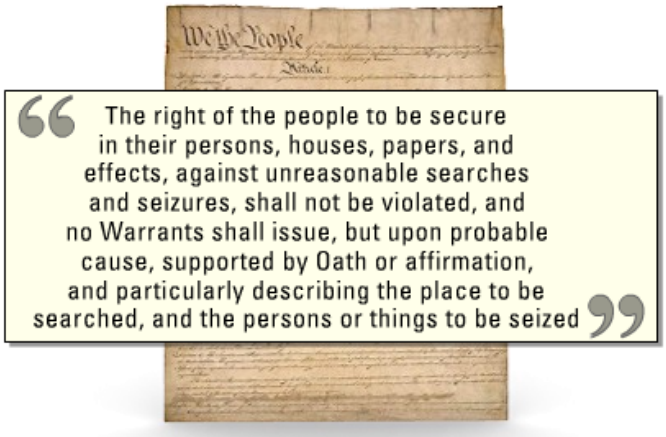

The Fourth Amendment:

- Protects against unreasonable searches and seizures
- Provides the right of people to be secure in their persons, houses, papers, and effects
- One of the liberties protected against State infringement

The flip side:

- Does permit searches and seizures that are considered reasonable
- Police may override your privacy concerns if:
  - They have probable cause to believe they can find evidence that you committed a crime and a judge issues a warrant
  - Particular circumstances justify the search without a warrant first being issued

The Fourth Amendment provides that the right of the people to be secure in their persons, houses, papers, and effects against unreasonable searches and seizures shall not be violated, and no warrants shall issue but upon probable cause, supported by oath or affirmation, and particularly describing the place to be searched, and the person or things to be seized. This protection is one of the liberties protected against State infringement by the Fourteenth Amendment.

The flip side is that the Fourth Amendment does permit searches and seizures that are considered reasonable. In practice, this means that the police may override your privacy concerns and conduct a search of your home, barn, car, boat, office, personal or business documents, bank account records, trash barrel or whatever, if:

- The police have probable cause to believe they can find evidence that you committed a crime, and a judge issues a search warrant, or
- The particular circumstances justify the search without a warrant first being issued.

What Constitutes a Security Incident?

Your organization should develop a list of what constitutes a security incident in the context of your business operations as you define your risk assessment and risk management procedures and privacy standards.

- Organization develops a list of what constitutes a security incident
- Based on:
  - Risk assessment
  - Risk management procedures
  - Privacy standards

"The act of violating an explicit or implied security policy"

- Attempts to gain unauthorized access to a system or its data
- Unwanted disruption or denial of service
- The unauthorized use of a system
- Changes to system hardware, firmware, or software without the owner's knowledge, instruction, or consent

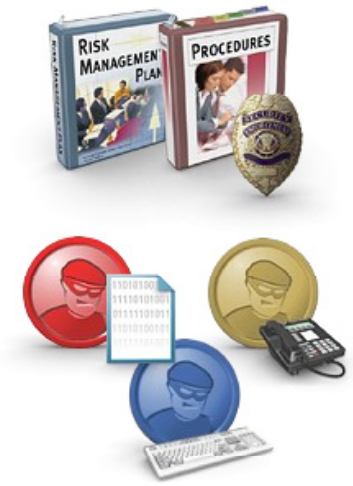

Your organization will need to implement accurate and current security incident procedures for those items that have been identified as incidents. The procedures will need to be formal, documented, report-and-response procedures. These security incident procedures relate to internal reporting of security incidents and do not specifically require you to report the incident to any outside entity.

A good but fairly general definition of an incident is "the act of violating an explicit or implied security policy". Unfortunately, this definition relies on the existence of a security policy that, while generally understood, varies between organizations. The types of activities that are generally considered to be in violation of a typical security policy include:

- Attempts (either failed or successful) to gain unauthorized access to a system or its data
- Unwanted disruption or denial of service
- The unauthorized use of a system for the processing or storage of data
- Changes to system hardware, firmware, or software characteristics without the owner's knowledge, instruction, or consent

Guidelines for Data Integrity and Retention

This topic discusses guidelines for minimizing threats to data integrity.

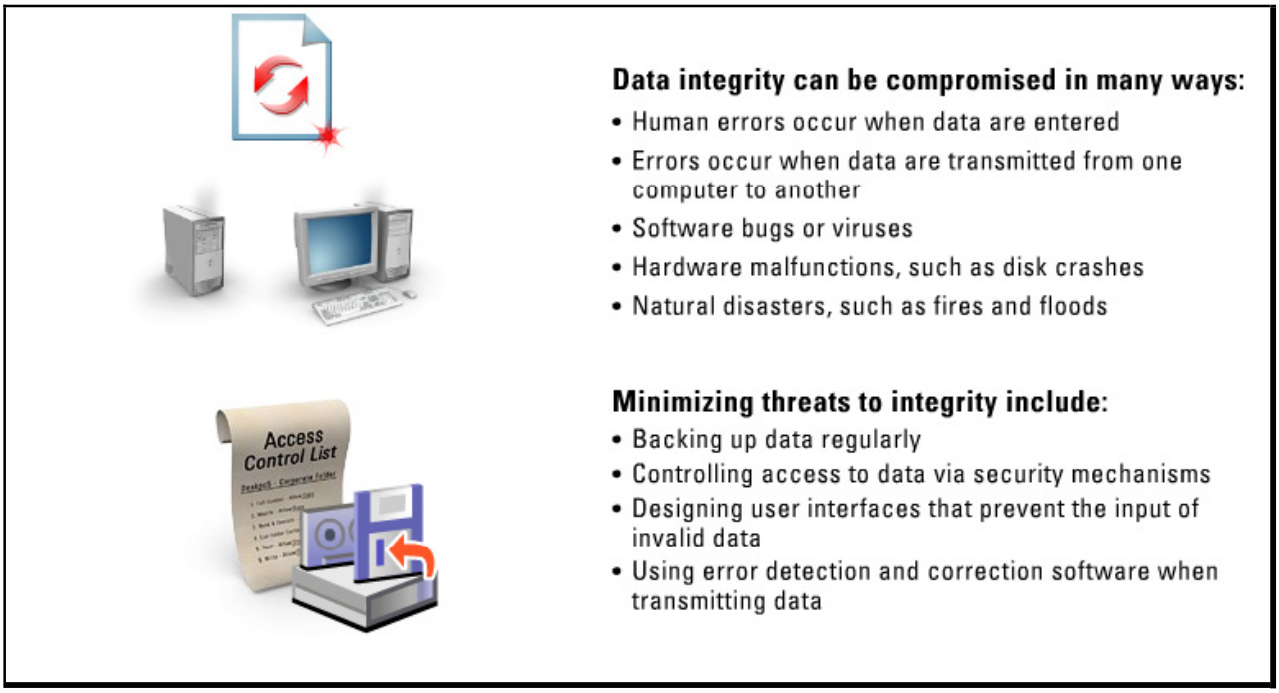

Data integrity can be compromised in a number of ways:

- Human errors occur when data are entered
- Errors occur when data are transmitted from one computer to another
- Software bugs or viruses
- Hardware malfunctions, such as disk crashes
- Natural disasters, such as fires and floods

There are many ways to minimize these threats to data integrity. They include:

- Backing up data regularly
- Controlling access to data via security mechanisms
- Designing user interfaces that prevent the input of invalid data
- Using error detection and correction software when transmitting data

Computer Ethics

Ethics are the principles of right and wrong that individuals acting as free moral agents can use to make choices to guide their behavior. **Computer ethics** is a new branch of ethics that is growing and changing rapidly as computer technology also grows and develops. The term computer ethics is open to interpretations both broad and narrow.

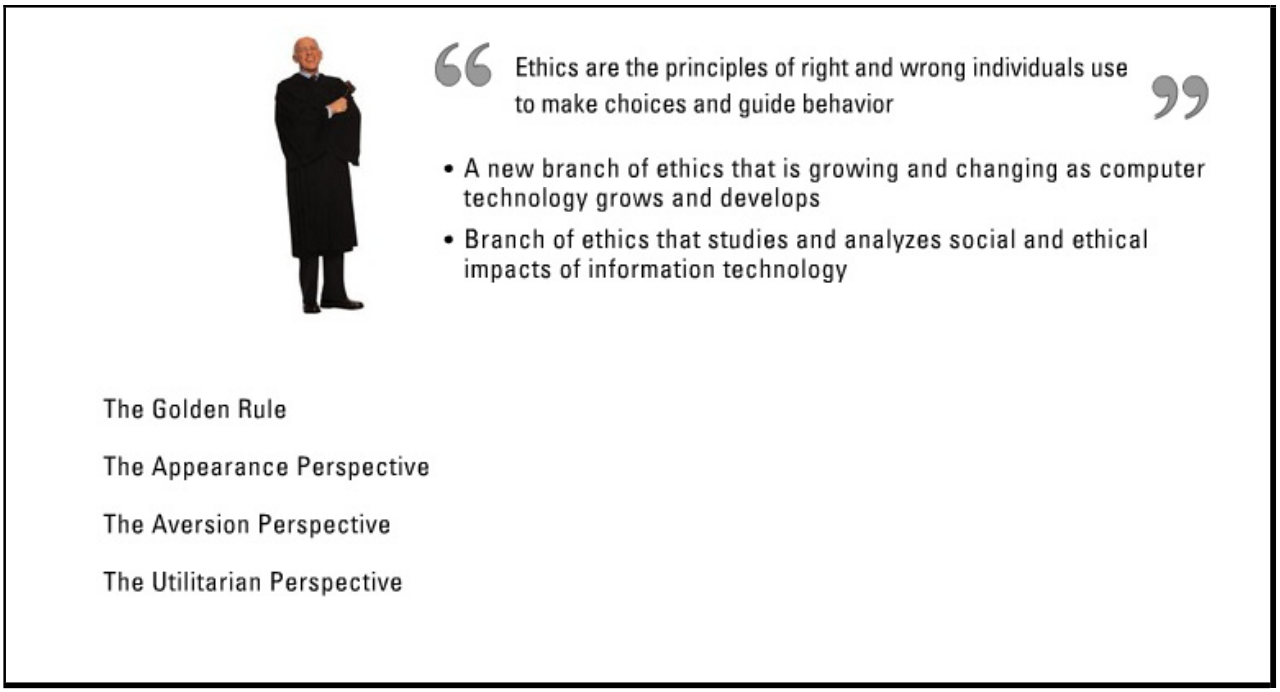

Various principles have been put forward to describe ethics. These include:

- The Golden Rule - Do unto others as you would have them do unto you
- The Utilitarian Perspective - Take the action that achieves the higher or greater value
- The Aversion Perspective - Take the action that produces the least harm, or the least potential cost
- The Appearance Perspective - The appearance of unethical behavior can do as much harm as actual unethical behavior

In the industrialized nations of the world, the "information revolution" already has significantly altered many aspects of life—in banking and commerce, work and employment, medical care, national defense, transportation, and entertainment. Consequently, information technology has begun to affect (in both good and bad ways) community life, family life, human relationships, education, freedom, democracy, and so on. Computer ethics in the broadest sense can be understood as that branch of applied ethics that studies and analyzes such social and ethical impacts of information technology.

Professional Code of Ethics

This topic discusses the ethical considerations of the International Information Systems Security Certification Consortium (ISC^2) and the Internet Architecture Board (IAB).

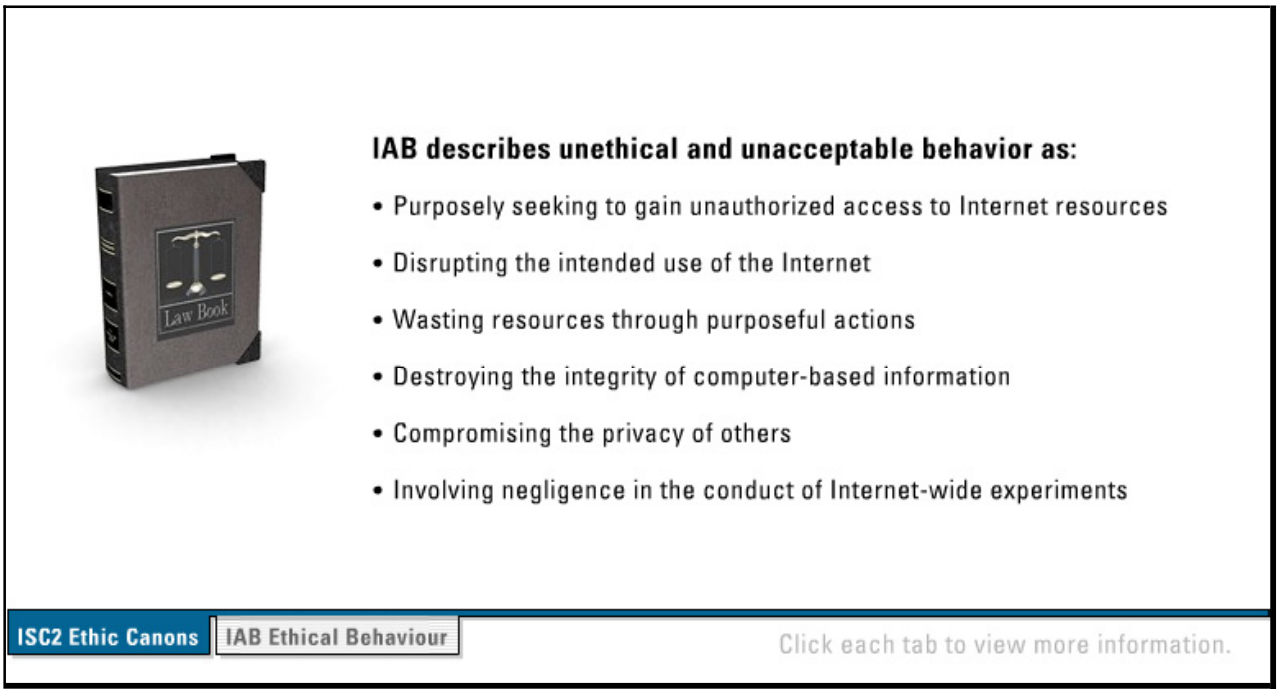

The ISC^2 provides these code of ethics canons:

- Protect society, the commonwealth, and the infrastructure
- Act honorably, honestly, justly, responsibly, and legally
- Encourage the growth of research
- Discourage unnecessary fear or doubt
- Discourage unsafe practices
- Observe and abide by all contracts, expressed or implied
- Avoid any conflicts of interest
- Provide diligent and competent service to principals
- Advance and protect the profession

IAB describes unethical and unacceptable behavior as:

- Purposely seeking to gain unauthorized access to Internet resources
- Disrupting the intended use of the Internet
- Wasting resources through purposeful actions
- Destroying the integrity of computer-based information
- Compromising the privacy of others
- Involving negligence in the conduct of Internet-wide experiments

Generally Accepted System Security Principles (GASSP): GASSP are developed and maintained with guidance from security professionals, IT product developers, information owners, and other organizations having extensive experience in defining and stating the principles of information security.

Common Ethical Fallacies

The computing world has supported many common ethical fallacies. These fallacies exist because people read and interpret rules differently.

Common Ethical Fallacies:

- Hackers only want to learn "how things work". These people are not making a profit off of their deeds. Thus, they cannot be illegal or unethical
- Hacking does not actually hurt anyone
- The first amendment protects a person's right to write a virus
- Information should be shared freely and openly, therefore sharing confidential information should be legal and ethical

These fallacies exist because people read and interpret rules differently

Common ethical fallacies include:

- Hackers only want to learn "how things work". These people are not making a profit off of their deeds. Thus, they cannot be illegal or unethical.
- Hacking does not actually hurt anyone.
- The first amendment protects a person's right to write a virus.
- Information should be shared freely and openly, therefore sharing confidential information should be legal and ethical.

Internet Architecture Board

The **Internet Architecture Board (IAB)**, formally called the Internet Activities Board, is the technical body that oversees the development of the Internet suite of protocols (commonly referred to as Transmission Control Protocol/Internet Protocol [TCP/IP]).

Internet Architecture Board:

- Formerly called the Internet Activities board
- The technical body that oversees the development of the Internet suite of protocols (TCP/IP)
- Has two task forces
- The Internet Research Task Force (IRTF)
- The Internet Engineering Task Force (IETF)

The IAB has two task forces, each of which are charged with investigating a particular area:

- The Internet Research Task Force (IRTF)
- The Internet Engineering Task Force (IETF)

The IAB's responsibilities include architectural oversight of IETF activities, Internet Standards Process oversight and appeal, and the appointment of the RFC Editor. The IAB is also responsible for the management of the IETF protocol parameter registries as well as:

- **IESG Confirmation**: The IAB confirms the IETF Chair and IESG Area Directors from nominations provided by the IETF Nominating Committee.
- **Architectural Oversight**: The IAB provides oversight of, and occasional commentary on, aspects of the architecture for the protocols and procedures used by the Internet.
- **Standards Process Oversight and Appeal**: The IAB oversees the process used to create Internet standards. The IAB serves as an appeal board for complaints of improper execution of the standards process through acting as an appeal body in respect of an IESG standards decision.
- **RFC Series and IANA**: The IAB is responsible for editorial management and publication of the Request for Comments (RFC) document series, and for administration of the assignment of IETF protocol parameter values by the IETF Internet Assigned Numbers Authority (IANA).
- **External Liaison**: The IAB acts as representative of the interests of the IETF in liaison relationships with other organizations concerned with standards and other technical and organizational issues relevant to the Internet.

- **Advice to ISOC**: The IAB acts as a source of advice and guidance to the Board of Trustees and Officers of the Internet Society concerning technical, architectural, procedural, and (where appropriate) policy matters pertaining to the Internet and its enabling technologies.
- **IRTF Chair**: The IAB selects a chair of the IRTF for a renewable two-year term.

Competitive Intelligence vs. Industrial Espionage

This topic defines competitive intelligence and industrial espionage.

Competitive intelligence is "a formalized, yet continuously evolving process by which the management team assesses the evolution of its industry and the capabilities and behavior of its current and potential competitors to assist in maintaining or developing a competitive advantage" (Prescott and Gibbons 1993). Competitive intelligence tries to ensure that the organization has accurate, current information about its competitors and a plan for using that information to its advantage. Competitive intelligence uses public information—information that can be legally and ethically identified and accessed.

Industrial espionage is defined as spying on one's competitors to gain a competitive advantage and is illegal.

Competitive intelligence differs from industrial espionage in that competitive intelligence uses public sources to find and develop information on competition, competitors, and the market environment, unlike industrial espionage, which develops information by illegal means like "cracking".

Summary

The key points discussed in this lesson are:

- Intruders will do whatever they can to not be seen when perpetrating their crimes. To protect a certain location or space, you can use surveillance mechanisms.
- The warrant process in the United States includes three main stages:
  - Basic requirements
  - Obtaining the warrant from the judge
  - Executing the warrant (serving the warrant)
- Entrapment encourages someone to commit a crime that the individual may have had no intention of committing. Conversely, enticement lures someone toward some evidence after that individual has already committed a crime.
- The Fourth Amendment protects against unreasonable searches or seizures. The flip side is that the Fourth Amendment does permit searches and seizures that are considered reasonable.
- Your organization will need to implement accurate and current security incident procedures for those items that have been identified as incidents. The procedures will need to be formal, documented, report-and-response procedures.
- There are many ways to minimize threats to data integrity. They include:
  - Backing up data regularly
  - Controlling access to data via security mechanisms
  - Designing user interfaces that prevent the input of invalid data
  - Using error detection and correction software when transmitting data
- Computer ethics in the broadest sense can be understood as that branch of applied ethics that studies and analyzes such social and ethical impacts of information technology.
- You should follow the ethical considerations set forth by the ISC2 and the IAB.
- The computing world has supported many common ethical fallacies. These fallacies exist because people read and interpret rules differently.
- The IAB is the technical body that oversees the development of the Internet suite of protocols.
- Competitive intelligence differs from industrial espionage in that competitive intelligence uses public sources to find and develop information on competition, competitors, and the market environment, unlike industrial espionage, which develops information by illegal means like “hacking”.

Physical Security

Overview

You must protect servers and other systems storing critical or sensitive information not only from unauthorized technological access, but also from unauthorized physical access. After all, if the attacker can gain physical access to a server, he or she can circumvent just about all security devices in use on the network.

Objectives

Upon completing this module, you will be able to:

- Define physical security and list its components
- Explain methods used to secure the perimeter of a facility
- List ways to ensure security of power, HVAC, water, and gas systems, as well as fire controls
- Describe the pros and cons of various intrusion detection systems
- Explain methods used to secure compartmentalized areas

Outline

The module contains these lessons:

- Introduction to Physical Security
- The Perimeter
- Inside the Building
- Intrusion Detection Systems
- Compartmentalized Areas

Introduction to Physical Security

Overview

This lesson will discuss how to properly provide physical security in the enterprise.

Importance

It is vital for the information security specialist to understand how to properly implement physical security.

Objectives

Upon completing this lesson, you will be able to:

- Define physical security
- List the major sources of physical loss
- Identify the elements that make up the physical security in an organization
- Identify the objectives of the layered defense

Outline

The lesson contains these topics:

- What Is Physical Security?
- Sources of Physical Loss
- Physical Security Elements
- The Layered Defense

What Is Physical Security?

You set **physical security** in place to address threats, vulnerabilities, and countermeasures. Essentially, it physically protects an enterprise's resources and sensitive information, including people, facilities, data, equipment, support systems, media, and supplies.

Physical security is an important consideration when choosing a secure site and its design and configuration. You should also consider physical security when choosing the methods for protecting against unauthorized access, theft of equipment and information, and environmental and safety threats, as well as when choosing the measures needed to protect people, the facility, and its resources.

Threats to physical security include:

- Emergencies
  - Fire and smoke
  - Building collapse
  - Utility loss
  - Water damage
  - Toxic materials
- Natural disasters
  - Earthquakes
  - Storm damage
- Human intervention
  - Sabotage
  - Vandalism

- War
- Strikes

Sources of Physical Loss

This topic lists major sources of physical loss.

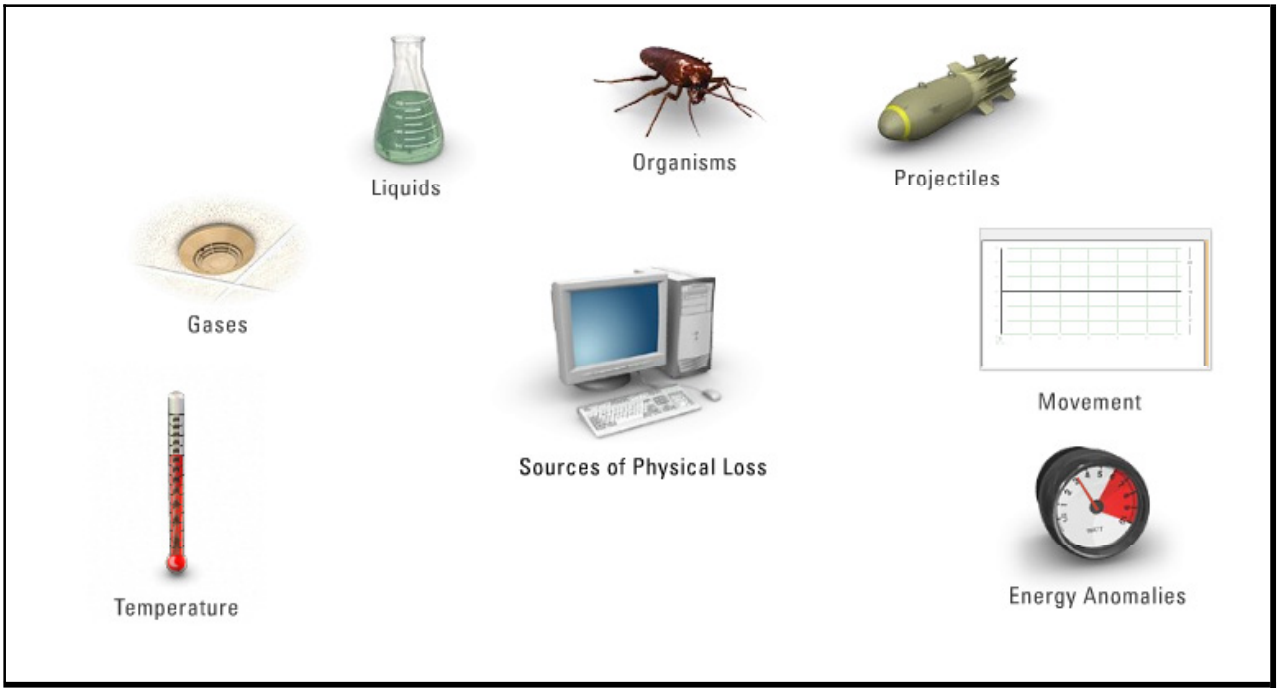

Seven major sources of physical loss include:

- **Temperature** - Extreme variations in heat or cold
- **Gases** - War gases, commercial vapors, humidity, dry air, and fuel vapors
- **Liquids** - Water and chemicals
- **Organisms** - Viruses, bacteria, people, animals, and insects
- **Projectiles** - Meteorites, falling objects, cars, trucks, bullets, and aircraft
- **Movement** - Collapse, shearing, shaking, vibration, and slides
- **Energy Anomalies** - Clectric surges, magnetism, static electricity, radio waves, and microwaves

Source: *Fighting Computer Crime* by Donn B. Parker

Physical Security Elements

The actual facility design your organization implements can affect the level of physical security available.

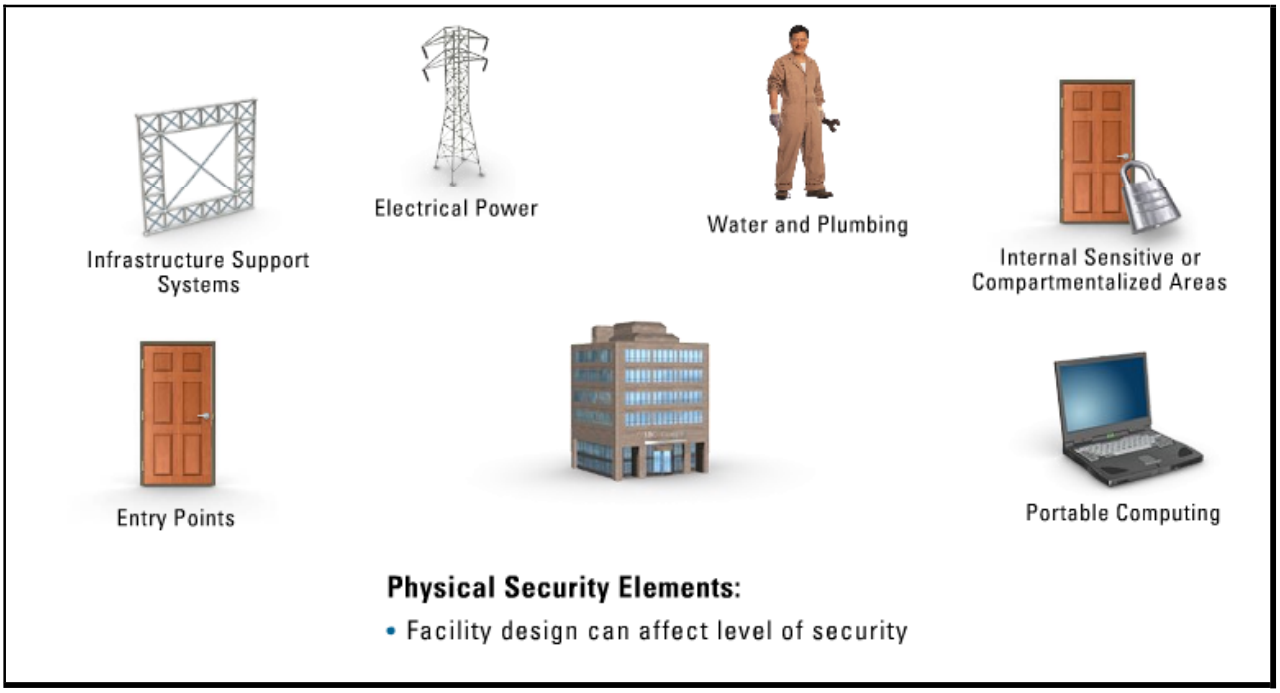

The following elements make up the physical security in an organization:

- **Entry points**
  - Doors and windows
  - Roof access
  - Service or delivery doors
  - Fire escapes
- **Infrastructure support systems**
  - Power
  - Water and plumbing
  - Heating, ventilation, and air conditioning (HVAC)
  - Physical infrastructure support
  - Cabling, plugs, and sockets
  - Loose wires and exposed cabling
- **Electrical Power** - A disruption in the electrical power supply can have a serious business impact. Continuous power is vital to the well-being of all major enterprises. As such, you should work with utility providers to create a plan for identifying and configuring a protection strategy for information systems that will meet the power demands of your company.
- **Water and Plumbing** - Water damage can cause irreparable damage to computer systems. Common sources of water damage include broken pipes, fire-suppression systems, evaporative coolers, and condensers.

- **Internal Sensitive or Compartmentalized Areas** - These are areas where highly sensitive information may be processed and stored, including the data center or server room, communications centers, switching centers, or end-user areas.
- **Portable Computing** - Remote computing and wireless connectivity require special implementations to provide for security.

The Layered Defense

Every facility should provide a **layered defense** that starts with perimeter security, continues through the building grounds and building entry points, and finally continues to each room or department.

The Layered Defense:

- Objective: To prevent or deter unauthorized or illegal events from occurring

The objectives of the layered defense are to prevent or deter unauthorized or illegal events from occurring, and if they do occur, to detect the event and delay the activity to allow for a timely response. The physical security points in the layered defense are:

- The perimeter
- Building entry points
- Building floors, office suites, and offices
- Compartmentalized areas

Summary

The key points discussed in this lesson are:

- Physical security physically protects an enterprise's resources and sensitive information, including people, facilities, data, equipment, support systems, media, and supplies.
- Seven major sources of physical loss include:
  - Temperature
  - Gases
  - Liquids
  - Organisms
  - Projectiles
  - Movement
  - Energy Anomalies
- The following elements make up the physical security in an organization:
  - Entry Points
  - Infrastructure Support Systems
  - Electrical Power
  - Water and Plumbing
  - Internal Sensitive or Compartmentalized Areas
  - Portable Computing
- The objectives of the layered defense are to prevent or deter unauthorized or illegal events from occurring, and if they do occur, to detect the event and delay the activity to allow for a timely response.

The Perimeter

Overview

In a layered defense, the perimeter is the first line of defense that intruders must overcome. This lesson will discuss the security features that can mitigate the possibility of an attacker successfully penetrating the perimeter of a facility.

Importance

Understanding perimeter weaknesses and how to strengthen them is an important first step in protecting the facility.

Objectives

Upon completing this lesson, you will be able to:

- List possible protective barriers on the perimeter
- Identify common types of protective lighting systems
- Describe the logistics of perimeter intrusion detection sensors
- List surveillance devices used to protect the perimeter
- Identify the levels and main components of a CCTV system
- Identify the basic types of building construction
- Identify methods attackers use to gain entry through a locked door
- Define mantraps
- Identify security controls for windows
- Identify types of locking devices
- Identify considerations regarding guard stations

Outline

The lesson contains these topics:

- Perimeter Security
- Lighting

- Perimeter Intrusion Detection
- Surveillance Devices
- Closed-Circuit Television
- Building Materials
- Doors
- Mantraps
- Windows
- Locks
- Guard Stations

Perimeter Security

The first line of defense for any enterprise is perimeter security, which is usually located as far away as possible from the main building.

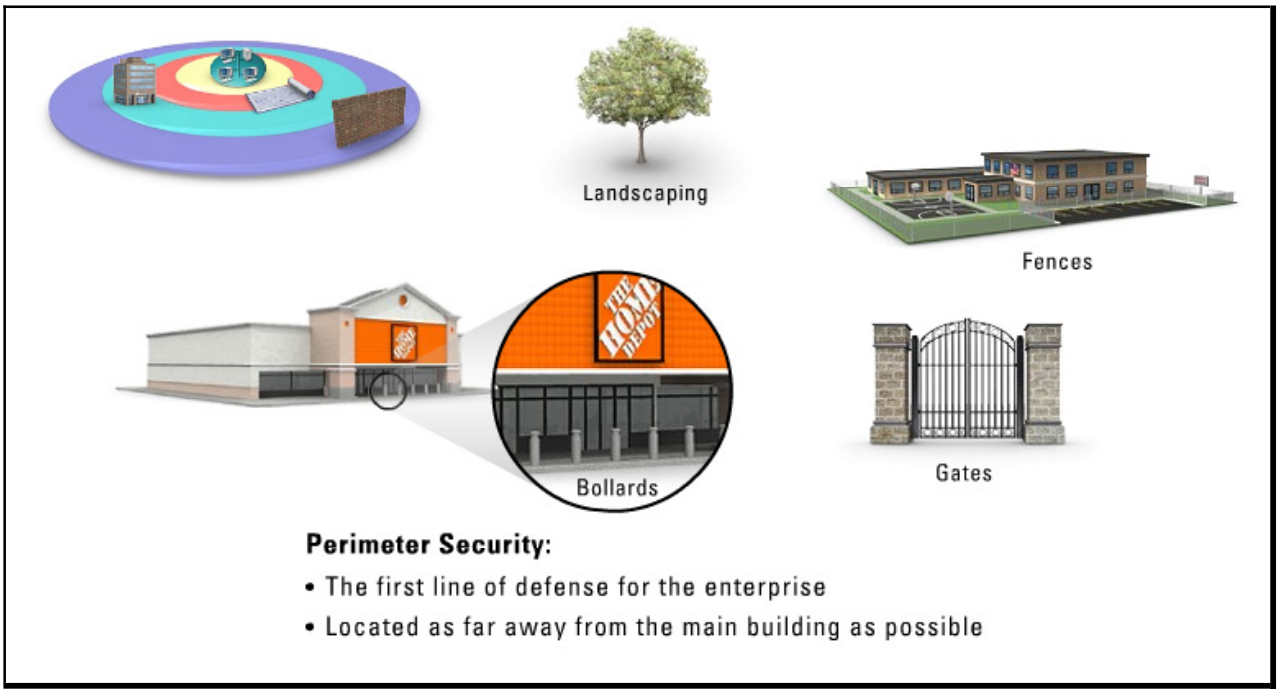

Protective barriers on the perimeter can include:

- **Landscaping** - Shrubs and trees can provide a barrier point
- **Fences** - Fences or wall structures can secure an area and are usually supplemented with security patrol
- **Gates** - Moving barriers, such as swinging, sliding, lowering, or rolling gates can be used to control the entrance of persons or vehicles; gates are separated into classes:
  - Class I: Residential gate operation
  - Class II: Commercial, such as parking lot or garage
  - Class III: Industrial/limited access (warehouse, factory, loading dock, etc.)
  - Class IV: Restricted access operation that requires supervisory control (prison, airport, etc.)
- **Bollards** - Heavy duty rising posts can be used to provide traffic control and property protection

Lighting

Good lighting can provide excellent protection in any environment. Because those who perpetrate unlawful acts prefer to conceal themselves in the cover of darkness, shedding light on an area deters trouble.

5 Protective Lighting Systems

Lighting:

- Good lighting provides excellent protection
- Shedding light on an area deters trouble
- Have light dispersed around bordering areas

You can achieve protective lighting by having even light dispersed around bordering areas, glaring lights in the eyes of an intruder, and showing a small amount of light on areas patrolled by security. Common types of protective lighting systems include the following:

- **Continuous Lighting** - The most common type of lighting, which consists of a series of fixed luminaries arranged to flood a given area continuously during times of darkness
- **Trip Lighting** - Lighting activated by some trigger, such an intruder crossing a sensor
- **Standby Lighting** - Lighting activated when suspicious activity is suspected
- **Emergency Lighting** - Lighting used for limited times when power fails
- **Gaseous Lighting** - Lighting provided by lamps, such as high-pressure sodium and mercury vapor lamps; these lamps have an inherent security weakness in that they can take several minutes to re-ignite

Perimeter Intrusion Detection

Perimeter intrusion detection sensors are devices that can detect intrusion across or under a land boundary or through a physical barrier, such as a chain-link fence.

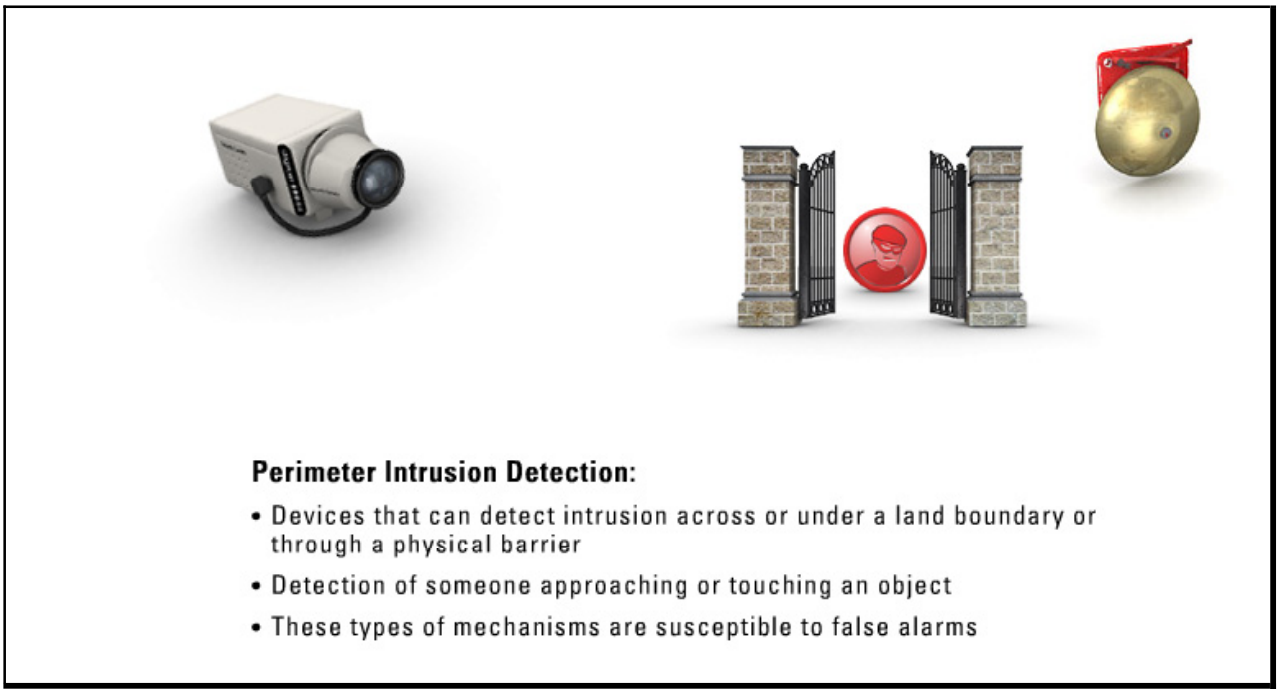

Other perimeter intrusion detection mechanisms can detect someone approaching or touching an object, such as a door or a vehicle. A problem with these types of mechanisms lies in the fact that they are susceptible to false alarms caused by non-intruders, such as animals, flying debris, or the weather.

Surveillance Devices

To provide for increased protection for locked buildings, you might consider using surveillance devices.

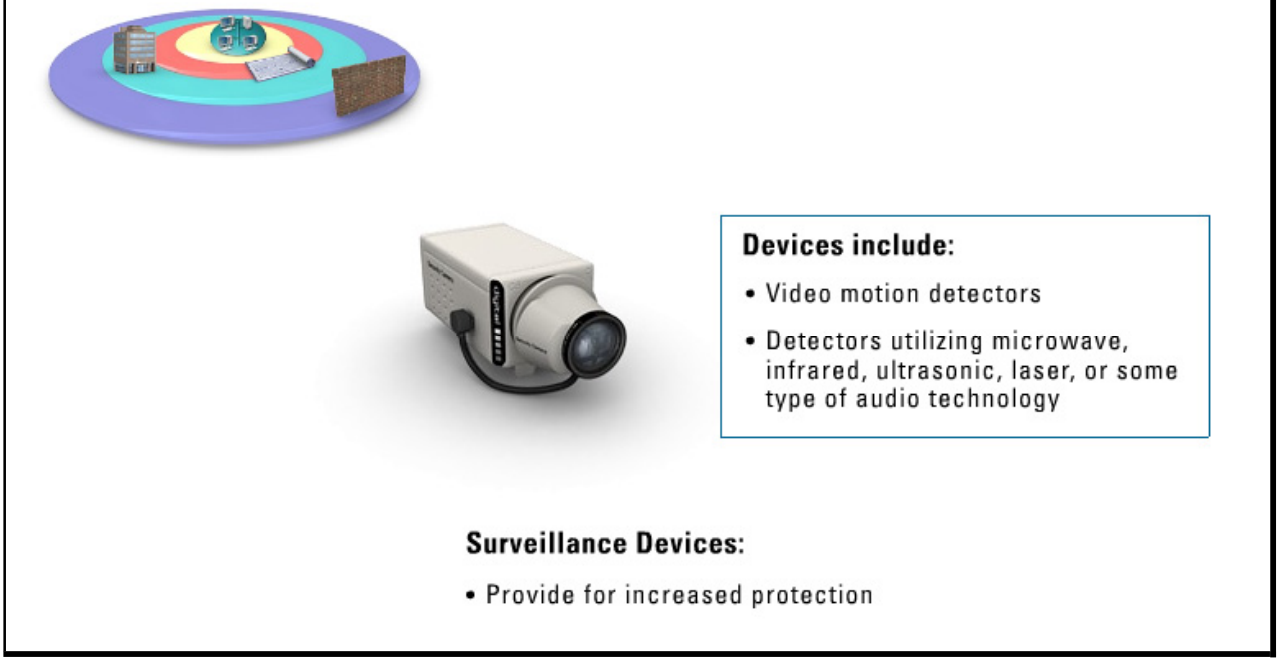

Surveillance devices include video motion detectors or detectors utilizing microwave, infrared, ultrasonic, laser, or some type of audio technology.

Closed-Circuit Television

Closed-circuit television (CCTV) is a system that uses video cameras to capture video transmissions and display them on connected monitors.

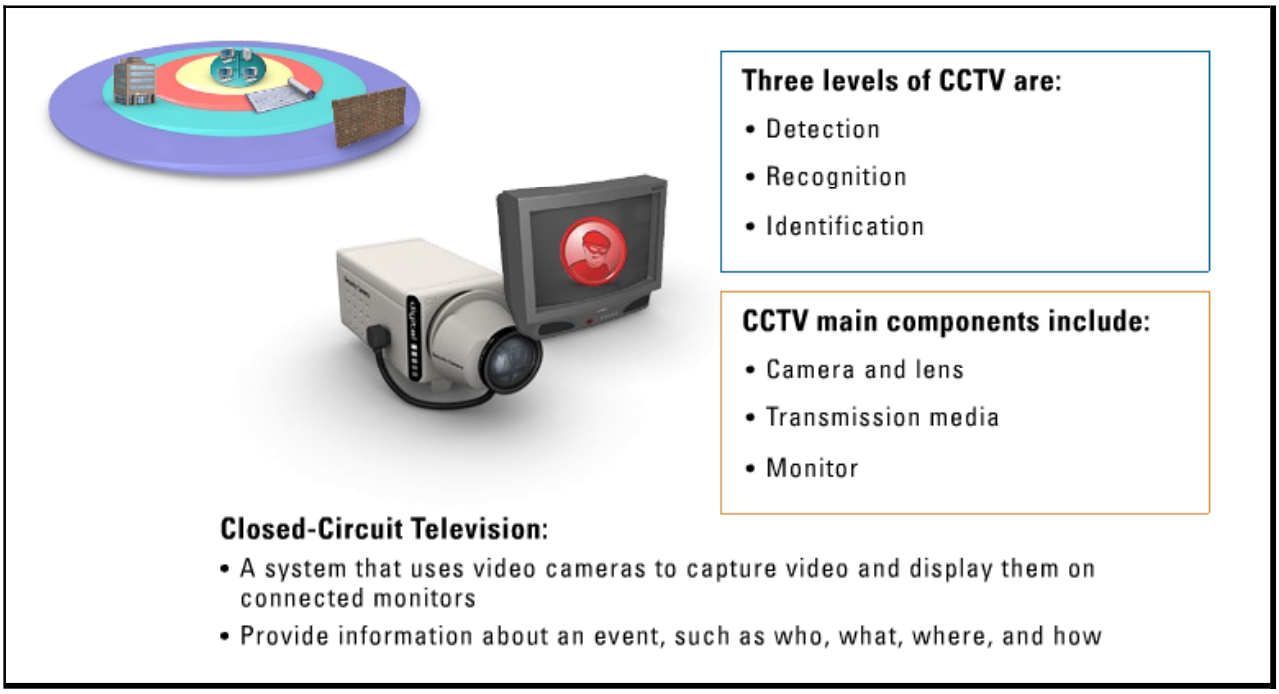

The CCTV system can provide information about an event, such as who, what, where, and how. The three levels of CCTV are:

- **Detection** - The ability to detect the presence of an object
- **Recognition** - The ability to determine the type of an object
- **Identification** - The ability to determine object details

---

Note Original domestic television systems were called broadcast, because anyone who selected the correct station could receive the transmission; these systems were often called open systems. CCTV systems do not use the broadcast medium, but instead have some type of hard-wired direct connection. Thus, CCTV systems are considered closed systems.

The main components of a CCTV system include:

- Camera and lens
  - Depth-of-field: The area between the nearest and farthest points that appear to be in focus
  - Field-of-view: The entire area that can be captured by the camera lens
- Transmission media
- Monitor

Building Materials

Security requirements for the outside building structure include the mandatory building codes.

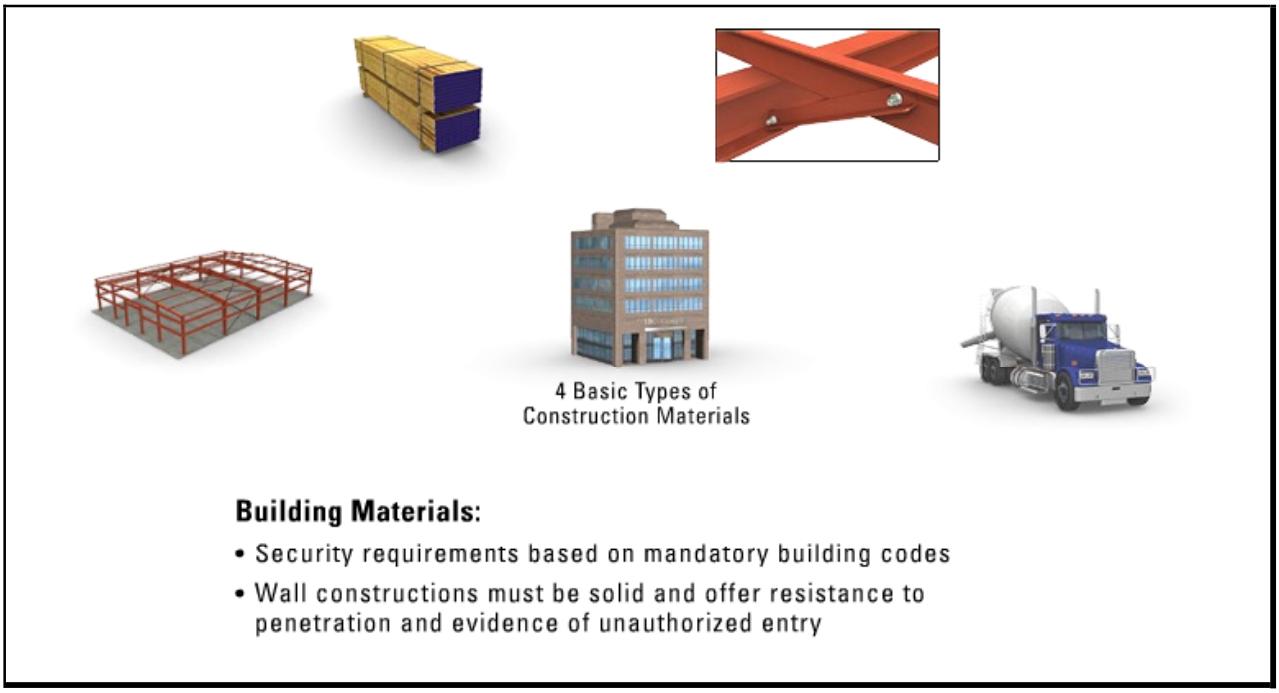

Basic types of building construction include:

- **Light Frame** - Typical of most homes; fire survival ability is rated at 30 minutes
- **Heavy Timber** - Structures with a minimum thickness of four inches; fire survival ability rated at one hour
- **Incombustible** - Structures consisting of steel; fire survival will weaken infrastructure (eventually causing collapse)
- **Fire Resistant** - Structures consisting of incombustible elements, usually encased in concrete

Building materials used in wall constructions must be solid and offer resistance to penetration and evidence of unauthorized entry.

Doors

The quality of a door and its frame, hinges, lock, and installation method are all critical factors in protecting against unauthorized entry.

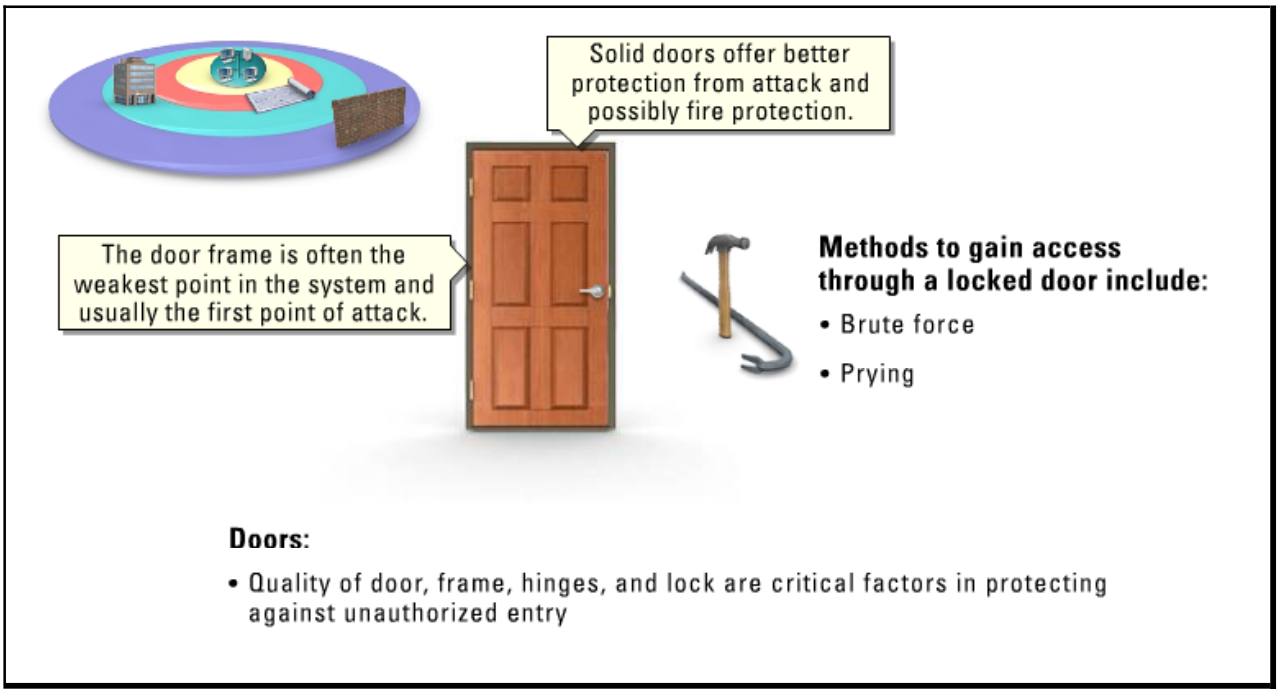

Methods attackers use to gain entry through a locked door include:

- **Brute Force** - Such as kicking in the door
- **Prying** - Forcing the door open with a tool, such as a crowbar

Hollow-core doors are just slightly less resistant to attack than drywall, which attackers can kick in, cut, or saw with minimal effort. Solid-core doors offer better protection to attack and possibly fire protection.

Caution	The doorframe is often the weakest point in the entire entry system and therefore is usually the first point of attack.

Mantraps

A **mantrap** is a system that routes personnel through two interlock-controlled doors into the facility.

Mantraps:

- A system that routes personnel through two interlock-controlled doors into the facility
- Only one door can be opened at a time
- Inner door will not unlock if outer door is open (and vice versa)
- Authentication required to enter second door

One of the key characteristics of the mantrap is that only one door can be opened at a time. To accomplish this, the design specifies that the inner door will not unlock if the outer door is open, or vice versa. In most cases, a person must produce some type of authentication to enter the second door; if the person cannot produce the proper authentication, the person will be trapped between the two doors and cannot enter or leave until a security person intervenes.

Windows

The materials of a window and the method used to construct it determine its security.

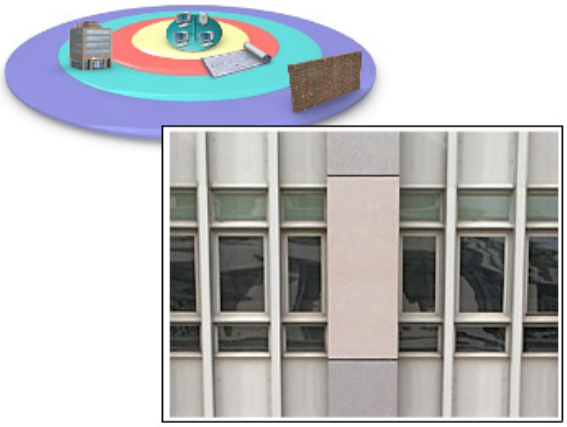

Windows

In residential homes, the most common type of window is made of standard plate glass. This type of glass is easy to cut and shatters into dangerous shards when broken. A more protective type of window contains tempered glass, which is five to seven times more break -resistant and, when cut or pierced, shatters into many small fragments, similar to a car's windshield.

Acrylics are also a type of window material. Standard acrylics are usually not as strong as polycarbonate acrylics, but they are more break-resistant than standard plate glass. The concern with acrylics is that when burned, they produce toxic fumes that may harm personnel. Polycarbonate windows are lightweight plastics that are 20 times stronger than standard acrylic of the same thickness; glass-clad polycarbonates combine the best qualities of glass and acrylics. In fact, windows made from glass-clad polycarbonate are resistant to abrasion, chemicals, and fires, and are even anti-ballistic. Because they are very expensive, they are usually only used in very high security areas.

Security controls for windows include:

- **Laminated Glass** - Laminated glass is made by bonding a plastic inner layer between two outer layers of glass.
- **Wired Glass** - Wired glass is made by embedding a continuous sheet of wire mesh between two layers of ordinary glass.
- **Solar Window Films** - Solar window films are materials affixed to windows to offer efficiency for heating and cooling and to filter out the sun's damaging ultraviolet rays. They also protect the glass from shattering.
- **Window Security Film** - Window security film is similar to solar film, but it is designed from a security perspective. It protects from unwanted entries, storm damage, and in some cases, bomb blasts. It is the most effective way to improve the integrity of plate or tempered glass.

- **Glass Breakage Sensor** - Glass breakage sensors are specially designed microphones tuned to the frequency of breaking glass.

Locks

The most accepted and used security device is the lock.

Locks:

- The most accepted and used security device
- Basic tools of lock picking include:
  - Tension wrench
  - Pick
- Another lock defeating technique is called raking
  - Performed with a pick that has a wider tip
- Locks are considered delay devices
- Types of locks include:
  - Key lock
  - Combination lock
  - Electronic combination lock
  - Deadbolt lock
  - Keyless lock
  - Smart lock

The basic tools of lock picking are the tension wrench and the pick; you can buy either from a locksmith supply house. The tension wrench imparts a rotary motion to the key plug of the lock and aids in finding the lock tumblers, while the pick is used to move the binding tumblers, one at a time, to the shear line. When all tumblers are aligned properly with the shear line, the lock opens.

Another technique intruders use to defeat locks is called raking. Racking is performed with a pick that has a wider tip. It is inserted all the way to the back of the plug. The pick is then pulled out quickly, and when this happens all the pins are bounced up. As the rake exits, you turn the plug using a tension wrench. Some of the upper pins will happen to fall on the ledge created by the turning pins. An intruder can easily pick the remaining pins.

Note	Locks are not a foolproof way to bar entry and are thus considered delay devices. They keep honest people out, but cannot keep out determined intruders as most locks are easily bypassed and most keys are readily duplicated.

The types of locking devices include:

- **Key Lock** - A lock that requires a key to open
  - **Warded Lock** - Wards are obstructions to the keyhole that prevent all but the properly cut key from entering
  - **Wafer/Tumbler Lock** - Wafers under spring tension are located in the core or plug of the lock and protrude outside the diameter of the plug into a shell formed by the body of the lock
  - **Pin Tumbler Lock** - The key moves pins so that a shear line can be obtained, thus allowing the key to turn the plug and operate the lock; more secure than warded and wafer/tumbler locks

    - **Interchangeable Core** - A lock with a core that can be removed and replaced using a special-change key
- **Combination Lock** - A sequence of numbers in proper order are required to open the lock
- **Electronic Combination Lock** - Uses digital readouts and obtains its power from the energy created when the dials are turned; offers higher security than combination locks, but is much more expensive
- **Deadbolt Lock** - A bolt inserted into the frame of the door for added security
- **Keyless Lock** - A push button lock that has buttons that are pushed in sequence to open the door; sometimes called a cipher lock
- **Smart Lock** - An inexpensive plastic card that is pre-authenticated to open a door; smart locks are used in most hotels

Guard Stations

In enterprises with very high security concerns, guard stations can provide excellent deterrence and a flexible security and safety response in the event of an unauthorized intrusion or a security-related incident.

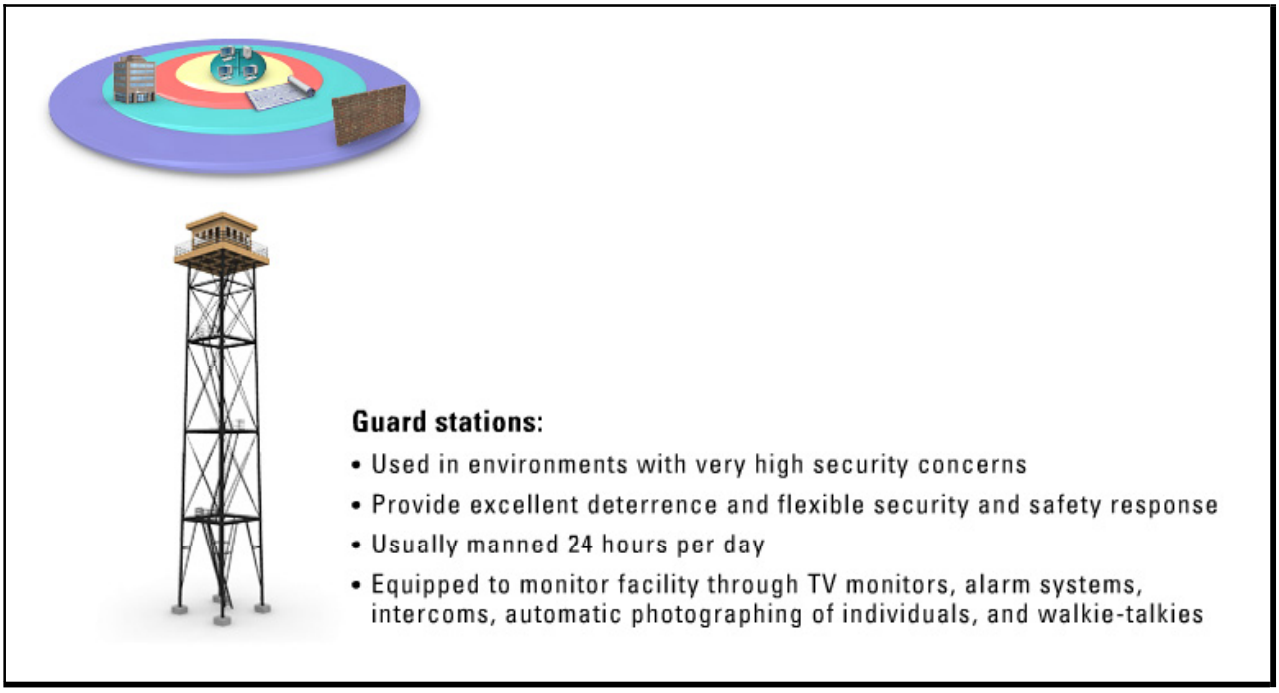

Guard stations are specially constructed enclosures that are usually manned 24 hours per day. They are usually equipped to monitor the facility through TV monitors, alarm systems, intercoms, automatic photographing of individuals entering the facility, and walkie-talkies for emergency communication.

When deciding on whether or not to implement a guard station, ask management the following questions:

- Is hiring or contracting more cost-effective?
- Will the guards require certification or licensing?
- Will the guards be armed?
- Are there any special union considerations?
- How will the guards be screened for the position?
- Will bonding be necessary?
- What specific training will be required?
- What is the impact on insurance policies?

Summary

The key points discussed in this lesson are:

- The first line of defense for any enterprise is perimeter security, which is usually located as far away as possible from the main building.
- Good lighting can provide excellent protection in any environment. Because those who perpetrate unlawful acts prefer to conceal themselves in the cover of darkness, shedding light on an area deters trouble.
- Perimeter intrusion detection sensors are devices that can detect intrusion across or under a land boundary or through a physical barrier, such as a chain-link fence.
- Surveillance devices include video motion detectors or detectors utilizing microwave, infrared, ultrasonic, laser, or some type of audio technology.
- CCTV is a system that uses video cameras to capture video transmissions and display them on connected monitors.
- Building materials used in wall constructions must be solid and offer resistance to penetration and evidence of unauthorized entry.
- The quality of a door and its frame, hinges, lock, and installation method are all critical factors in protecting against unauthorized entry.
- A mantrap is a system that routes personnel through two interlock-controlled doors into the facility.
- The materials of a window and the method used to construct it determine its security.
- The most accepted and used security device is the lock. However, locks are not a foolproof way to bar entry and are thus considered delay devices.
- In enterprises with very high security concerns, guard stations can provide excellent deterrence and a flexible security and safety response in the event of an unauthorized intrusion or a security-related incident.

Inside the Building

Overview

Keeping the inside of the building safe in the event of a disaster is the next layer of security that you must address. This lesson will discuss the security measures for utilities, heating, ventilation, and air conditioning (HVAC), and fire systems.

Importance

Understanding the security mechanisms for protecting assets inside the building is a necessary skill for the information security professional.

Objectives

Upon completing this lesson, you will be able to:

- Identify security controls for electrical power systems
- Identify security controls for HVAC, water, and gas systems
- Identify the fire controls needed for building and personnel security

Outline

The lesson contains these topics:

- Power Controls
- HVAC, Water, and Gas Controls
- Fire Controls

Power Controls

The first element of security within a building is securing the power supply. Because most technology-based security mechanisms rely on power, the security controls for electrical power require specific technical training and expert consultants.

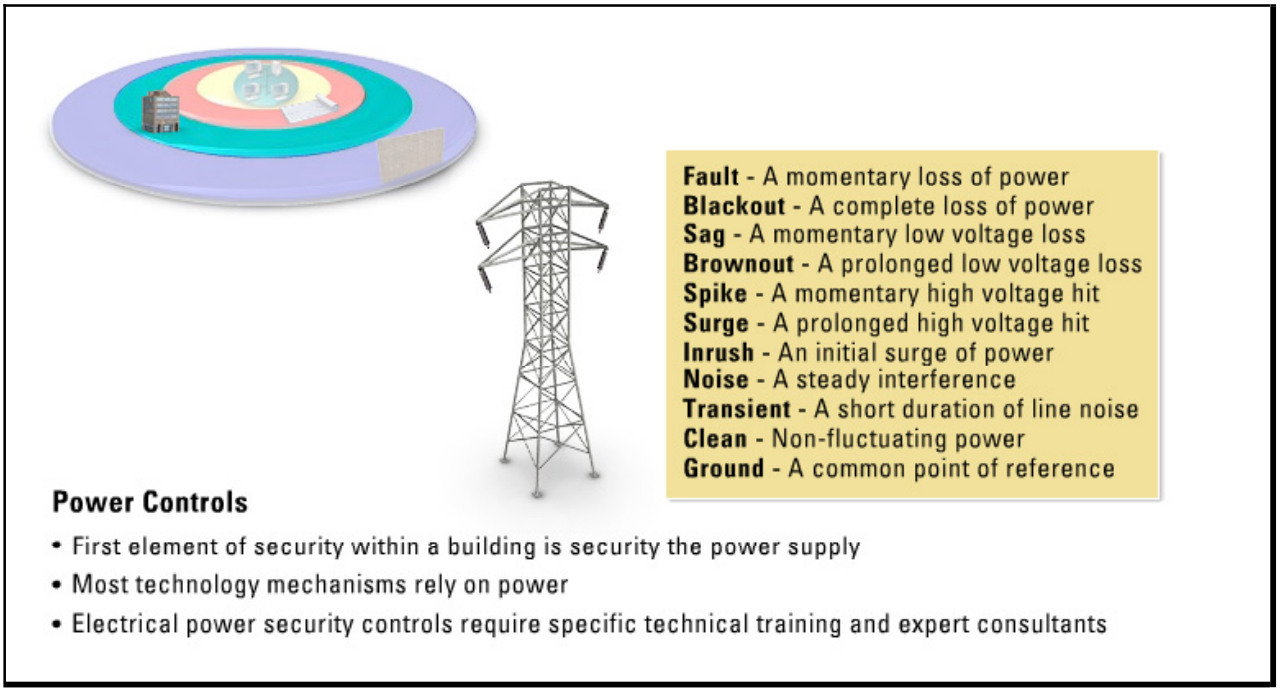

Computer systems require a clean steady supply of power to operate and provide a long life. The most common threats to computer systems include noise, brownouts, and humidity. Here are some terms used when discussing electrical power issues:

- **Fault** - A momentary loss of power
- **Blackout** - A complete loss of power
- **Sag** - A momentary low voltage loss
- **Brownout** - A prolonged low voltage loss
- **Spike** - A momentary high voltage hit
- **Surge** - A prolonged high voltage hit
- **Inrush** - An initial surge of power
- **Noise** - A steady interference
- **Transient** - A short duration of line noise
- **Clean** - Non-fluctuating power
- **Ground** - A common point of reference

To control interference, you should install a single-socket, stand-alone, power line filter suppressor. Some industrial equipment may require twist-type, extreme duty, or other high-current filter suppressors. Remember to always install filter-suppressants on problematic equipment.

An uninterruptible power supply (UPS) is an essential power control for today's computing environment. A UPS is basically a battery that will supply clean power to a system or systems in the event of a power shutdown or lengthy interruption from the primary source of electrical power.

Static electricity, although low in current, can be very high in voltage and has the ability to damage sensitive integrated circuits when the electrical current touches the metal contacts. For this reason, you should minimize static in the areas where people use computer equipment. To do this, provide an atmosphere where static is difficult to build up; you can use higher humidity to defeat static build-up. You can also use anti-static carpets, anti-static mats, or anti-static sprays to control static around computer equipment.

HVAC, Water, and Gas Controls

Separating air conditioning controls for a data center or server room from other controls used throughout the building is a good idea. Doing so allows for greater control over humidity and temperature in areas that contain sensitive equipment. Also, if possible, the server room should have a separate air conditioning system that is independent from the rest of the building.

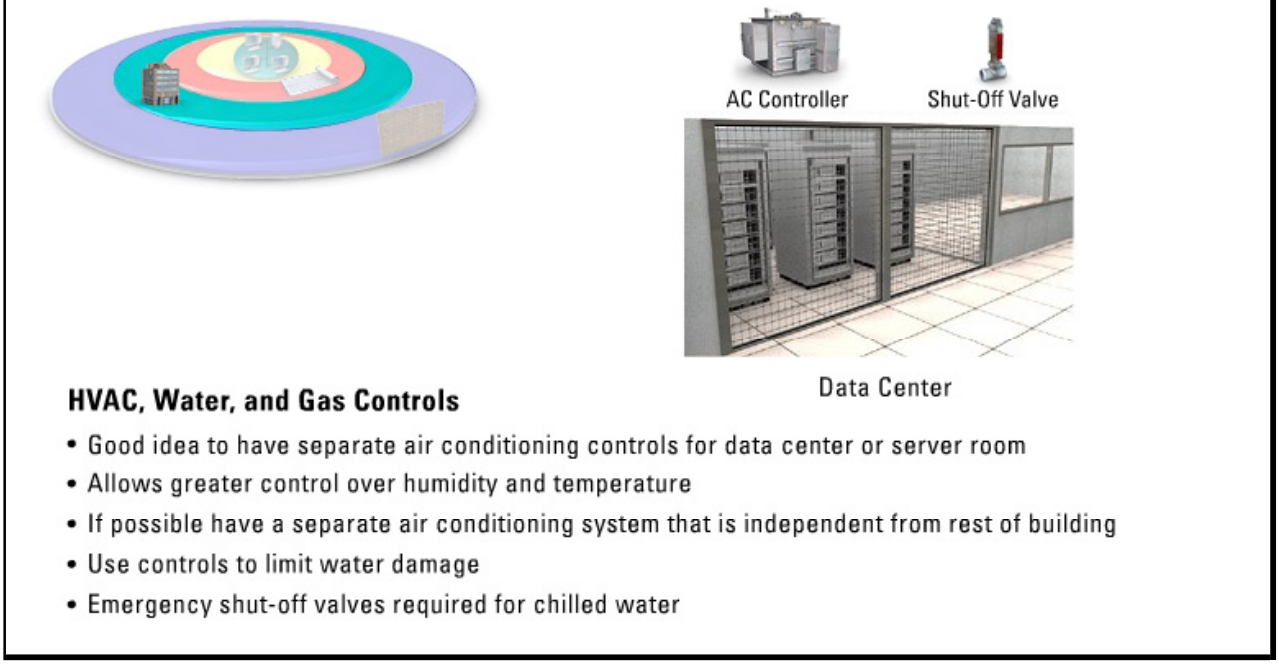

Because water can damage any computer system beyond the point of repair, you should have controls in effect to limit damage of this type. Do not place a computer room next to or directly below an area that can possibly flood. Make sure you check for pipes under raised floors that could damage equipment if they leak.

Emergency shut-off valves should be a requirement for chilled water. If an emergency occurs, the shut-off valve would automatically stop the flood of water into the pipeline; this would ensure the safety of the water supply and possibly prevent outside contaminants from entering the water supply.

If your facility uses natural gas, you should follow these safety precautions:

- Identify the location of the main shut-off valve
- Verify the shut-off valves work correctly
- Attach a wrench to a cord near the shut-off valve (if required)
- Communicate the locations of the shut-off valves to the local fire department
- Paint shut-off valves with white or fluorescent paint to increase their visibility
- Secure the main gas line in a fenced and locked area and use appropriate access controls
- Know the gas piping throughout the facility

Fire Controls

Fire prevention refers to three main areas of fire control: fire prevention, fire detection and containment, and fire suppression.

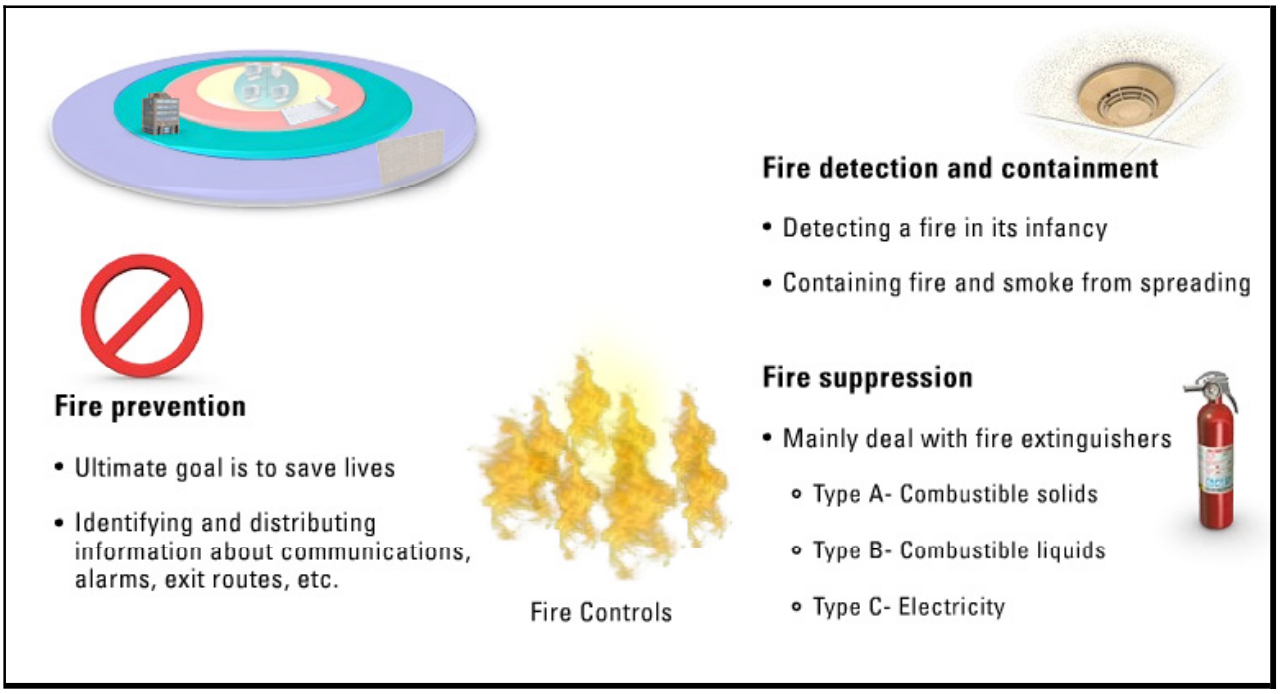

The goal of fire detection is to identify a fire in its infancy, while it is small and controllable. Fire containment and suppression involve how to deal with a fire in the event of an emergency. Containment involves the use of fire barriers (firewalls), vents, dampers, and the HVAC system to keep fire and smoke from spreading. Suppression methods involve using a fixed or portable extinguishing system. The ultimate goal of fire protection is to save lives. With this goal in mind, you must identify and distribute information about communications, alarms, exit routes, and refuge areas to all personnel in the facility.

Many fire detection systems are readily available and quite effective. Ionization-type smoke detectors react to the charged particles in smoke, whereas photoelectric detectors react to a light blockage caused by smoke. Also, heat detectors react to the heat of a fire, and combination detectors offer several of these functions in a single unit.

Fire suppression techniques mainly deal with extinguishers. You should have portable fire extinguishers near all electrical equipment. For computers, you should use type A, B, or C extinguishers.

- **Type A** - Combustible solids; suppression via water or soda acid
- **Type B** - Combustible liquids; suppression via gas (halon), CO_2, or soda acid
- **Type C** - Electricity; suppression via gas (halon) or CO_2
- **Type D** - Combustible metals; suppression via dry powder

Remember that the primary purpose of a fire extinguisher is to provide an escape route that people can follow to exit the building. If the fire is small enough, then trained personnel can use the extinguisher to extinguish the fire; but only after everyone else has been evacuated.

Summary

The key points discussed in this lesson are:

- Computer systems require a clean steady supply of power to operate and provide a long life. The most common threats to computer systems include noise, brownouts, and humidity.
- Separating air conditioning controls for a data center or server room from other controls used throughout the building is a good idea.
- Do not place a computer room next to or directly below an area that can possibly flood. Make sure you check for pipes under raised floors that could damage equipment if they leak.
- If your facility uses natural gas, you should follow several safety precautions involving the shut-off valves.
- The ultimate goal of fire protection is to save lives. With this goal in mind, you must identify and distribute information about communications, alarms, exit routes, and refuge areas to all personnel in the facility.

Intrusion Detection Systems

Overview

Intrusion detection is the ability to identify hostile acts against the enterprise. For security protection, you can install penetration sensors on windows, doors, ceilings, walls, or any other entry point in the facility. This lesson will discuss the various methods of intrusion detection in the enterprise.

Importance

It is vital for the information security specialist to understand the mechanics of both electronic and physical intrusion detection systems.

Objectives

Upon completing this lesson, you will be able to:

- Identify the pros and cons of an intrusion detection system that involves breaking a circuit
- Identify the pros and cons of an intrusion detection system that involves interrupting a light source
- Explain the logistics of passive infrared detection systems
- Identify the pros and cons of using sound detection systems
- Identify the pros and cons of using vibration detection systems
- Explain the logistics of motion detection systems

Outline

The lesson contains these topics:

- Breaking a Circuit
- Interrupting a Light Source
- Passive Infrared Detection
- Sound Detection
- Vibration Detection
- Motion Detection

Breaking a Circuit

To protect entry points into a building, floors, or offices, you can install an electrically sensitized strip of metal foil or wire that will indicate when someone opens a door or window. Any action that breaks the foil or wire will break the electrical circuit and activate an alarm.

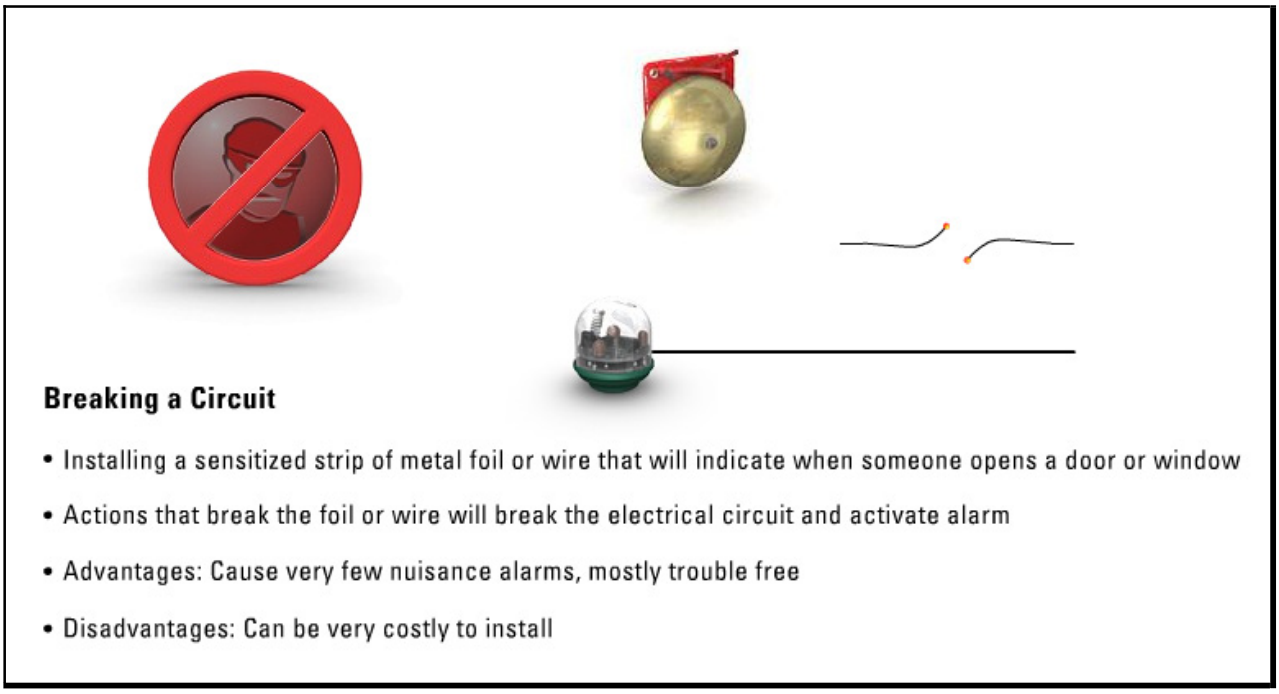

The advantages of this type of system are they cause very few nuisance alarms and are mostly trouble free. If an alarm is triggered, more than likely an unauthorized entry is in progress. The downside to these types of systems is the fact that they can be very costly to install if you have many points of entry.

Interrupting a Light Source

Another method for detecting intrusion is using a device that senses an interruption in a light source.

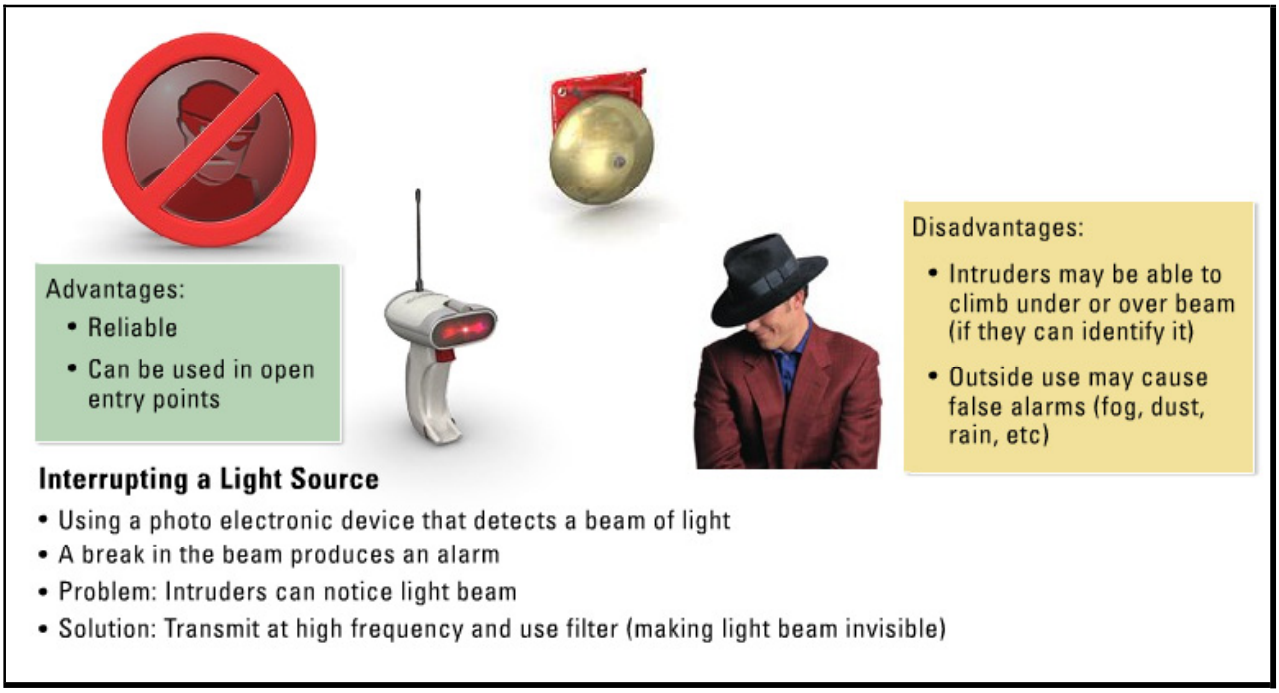

If you have a photoelectric device that detects a beam of light, it can also detect a break in the light source. The problem is intruders can notice the light, especially if the room is dark, and can easily avoid breaking the light path. To counter this problem, the light from the light beam can be transmitted at a frequency of several thousand vibrations per second. This frequency offers an infrared filter to cover the light source to make the beam invisible. The beam can then be crisscrossed with mirrors until it reaches the light receiver. If an intruder crosses any portion of the beam, the contact will be broken and an alarm will activate.

Advantages include:

- Reliable source of detection
- Can be used in open entry points where obstructions cannot be used

Disadvantages include:

- Intruders may be able to climb under or over the beam if they can identify it
- If used outside, fog, dust, rain, bright sunlight, or smoke can cause false alarms

Passive Infrared Detection

A **passive infrared detector** (PIR) is a device that emits a certain level of infrared energy in the form of light energy.

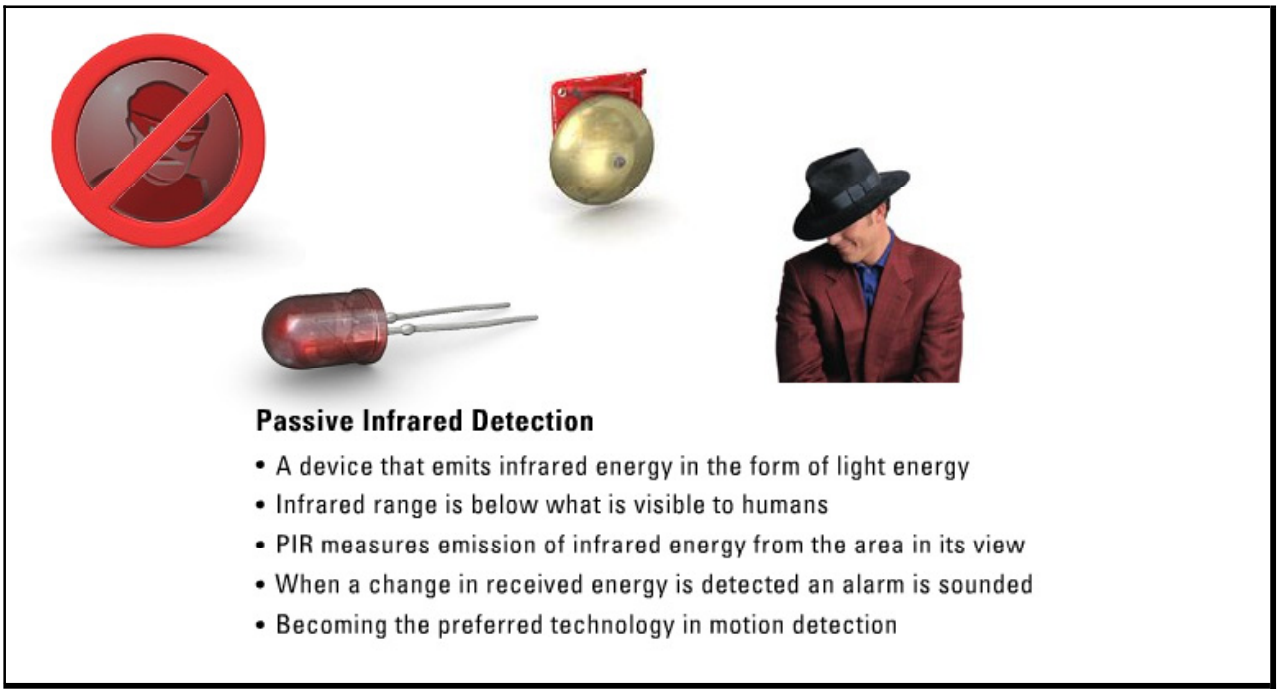

This light energy in the infrared range is below what is visible to humans. The PIR measures the emission of infrared energy from the area in its view. Lenses or mirrors might also be used to focus the received energy on the measuring element. When a change in received energy is detected, the PIR will go into an alarm condition. PIRs can detect radiations of heat (body heat) as well as movement (blocking the received infrared energy).

The PIR is becoming the preferred technology in motion detection because of its great ability and flexibility to control the area viewed through a variety of precision lenses and mirrored optics.

Sound Detection

Sound detectors have a microphone that is sensitive enough to detect even minute changes in sound volume; this technology uses sound as an alarm condition.

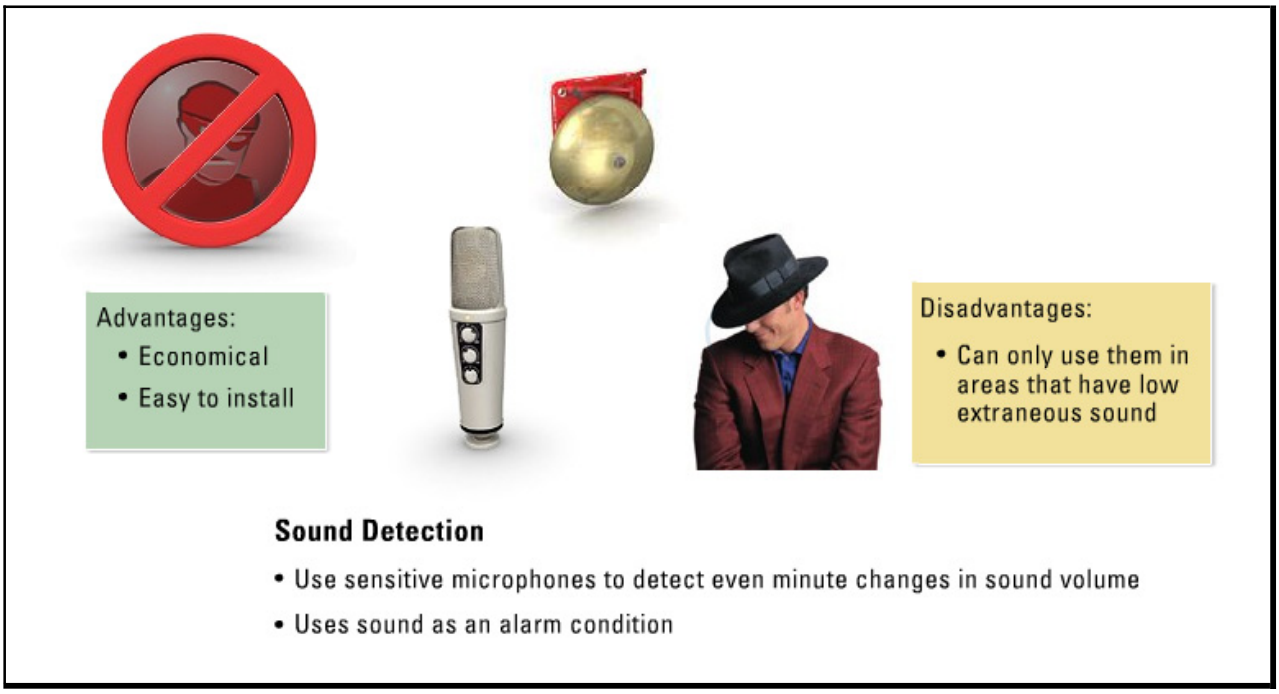

The advantages of sound detection are they are economical and you can easily install them. The disadvantage is the fact that you can only use them in areas that have low extraneous sound.

Vibration Detection

Vibration sensors detect any vibration caused by an intruder touching, jiggling, or opening a door, window, and so on.

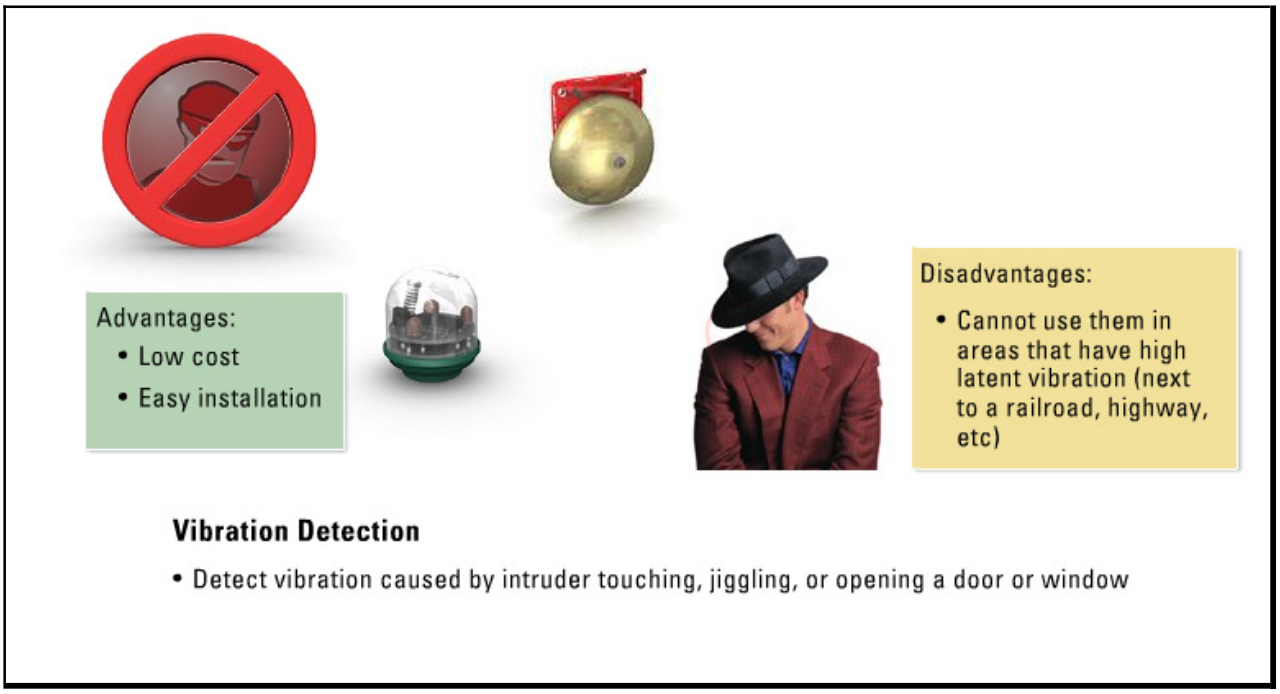

When you install the vibration sensor next to the entry point, no one can use that entry point without an alarm sounding. Advantages to vibration sensors include low cost and easy installation. The main disadvantage is the fact that you cannot use them in areas that have a high degree of latent vibration, such as next to a railroad or highway.

Motion Detection

Motion detection can occur via ultrasonic or microwave sensors and can be a very effective way to protect interior areas. Motion detection works by flooding the protected area with acoustic or microwave energy and detecting the Doppler shift in transmitted and received frequencies when motion occurs within the area.

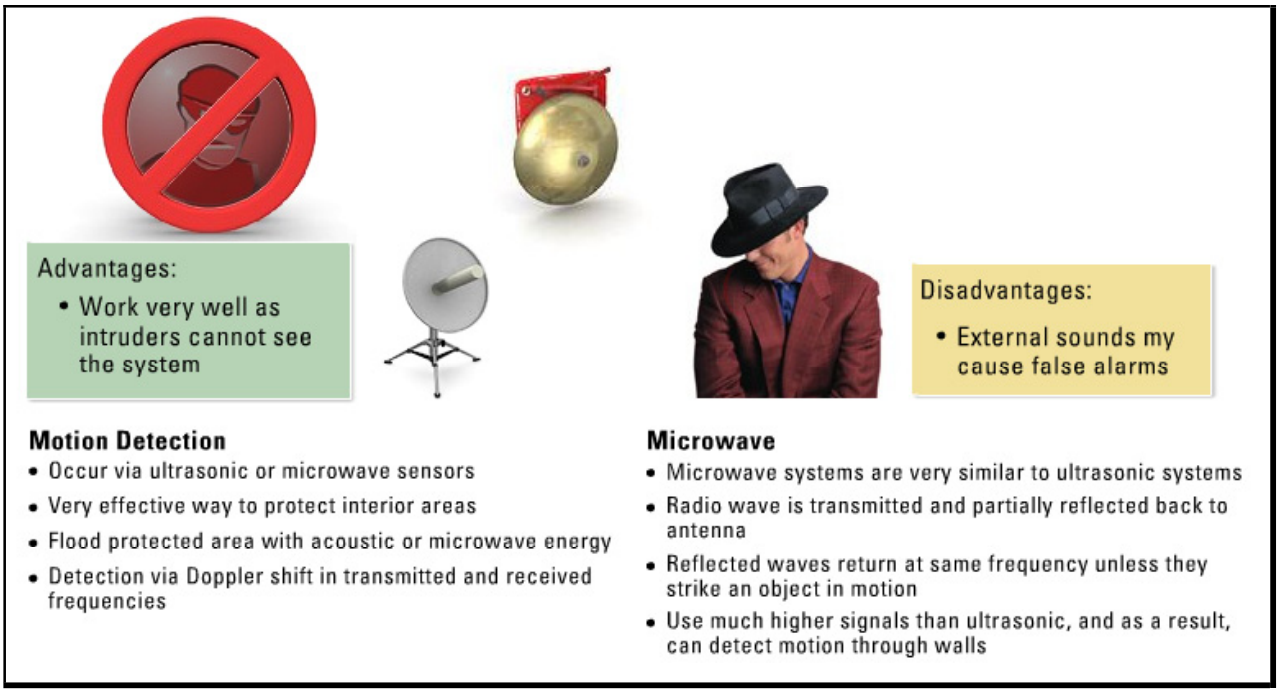

Ultrasonic systems are made up of the transceiver, an electronic unit, and a control unit. The transmitter is used to generate a pattern of acoustic energy that will fill the entire protected area (usually an enclosed room). When motion within the protected area occurs, the energy field is disturbed and an alarm is triggered. The main advantage of ultrasonic systems are the fact that they work very well as intruders cannot see the system. The main disadvantage is that external sounds may cause false alarms, which may cause the user to lower the sensitivity, and hence lower the effectiveness of the system.

Microwave systems are very similar to ultrasonic systems. In the microwave system, a pattern of radio waves is transmitted and partially reflected back to an antenna. If all objects in the protected area are stationary, the reflected waves return at the same frequency, but if they strike an object in motion, they return a different frequency, which will trigger an alarm. The advantage of the microwave system is that a large area can be covered if antennae are properly placed. The disadvantage of the microwave system is the fact that it is difficult to confine the energy transmitted. It is possible for the waves to penetrate thin walls and windows. If a person outside the protected area then crosses these waves, a false alarm can be generated.

Note	The principles of ultrasonic and microwave systems are the same, but microwave signals use a much higher frequency, and as a result, can detect motion through walls.

Summary

The key points discussed in this lesson are:

- To protect entry points into a building, floors, or offices, you can install an electrically sensitized strip of metal foil or wire that will indicate when someone opens a door or window. Any action that breaks the foil or wire will break the electrical circuit and activate an alarm.
- Another method for detecting intrusion is using a device that senses an interruption in a light source.
- A PIR is a device that emits a certain level of infrared energy in the form of light energy.
- Sound detectors have a microphone that is sensitive enough to detect even minute changes in sound volume; this technology uses sound as an alarm condition.
- Vibration sensors detect any vibration caused by an intruder touching, jiggling, or opening a door, window, and so on.
- Motion detection works by flooding the protected area with acoustic or microwave energy and detecting the Doppler shift in transmitted and received frequencies when motion occurs within the area.

Compartmentalized Areas

Overview

A compartmentalized area is one that you must protect at all times. This type of area contains sensitive equipment and information, and is usually a restricted-access area. These areas are typically found in government facilities, chemical laboratories, aerospace enterprises, and electronic-based organizations.

Importance

Understanding the security-based architecture of compartmentalized areas will help the information security specialist better protect important company assets.

Objectives

Upon completing this lesson, you will be able to:

- Describe the process for determining who has access to a compartmentalized area
- Identify good security practices for data centers/server rooms
- Describe methods for securing electronic equipment and data from theft
- Explain object-level security protection

Outline

The lesson contains these topics:

- Effective Access Control
- Data Center or Server Room
- Computer Equipment Protection
- Object Protection

Effective Access Control

To be effective, you must automatically control who can access a compartmentalized area and when.

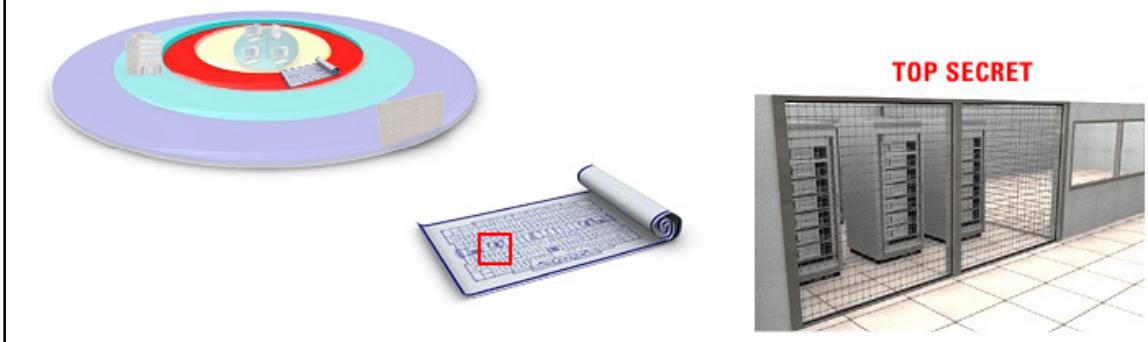

Effective Access Control

- A compartmentalized area is one that you must protect at all times
- Determining who has access to a compartmentalized area, an exhaustive background check is usually necessary
- Determine who has need-to-know clearance

To determine who has access to a compartmentalized area, you usually perform an exhaustive background investigation and determine who has need-to-know clearance. This process helps security officers verify who has access, when access is permitted, and for what purposes someone should have access.

Data Center or Server Room

The data center room or server room requires strict security clearance, as integrity in these rooms must remain extremely high.

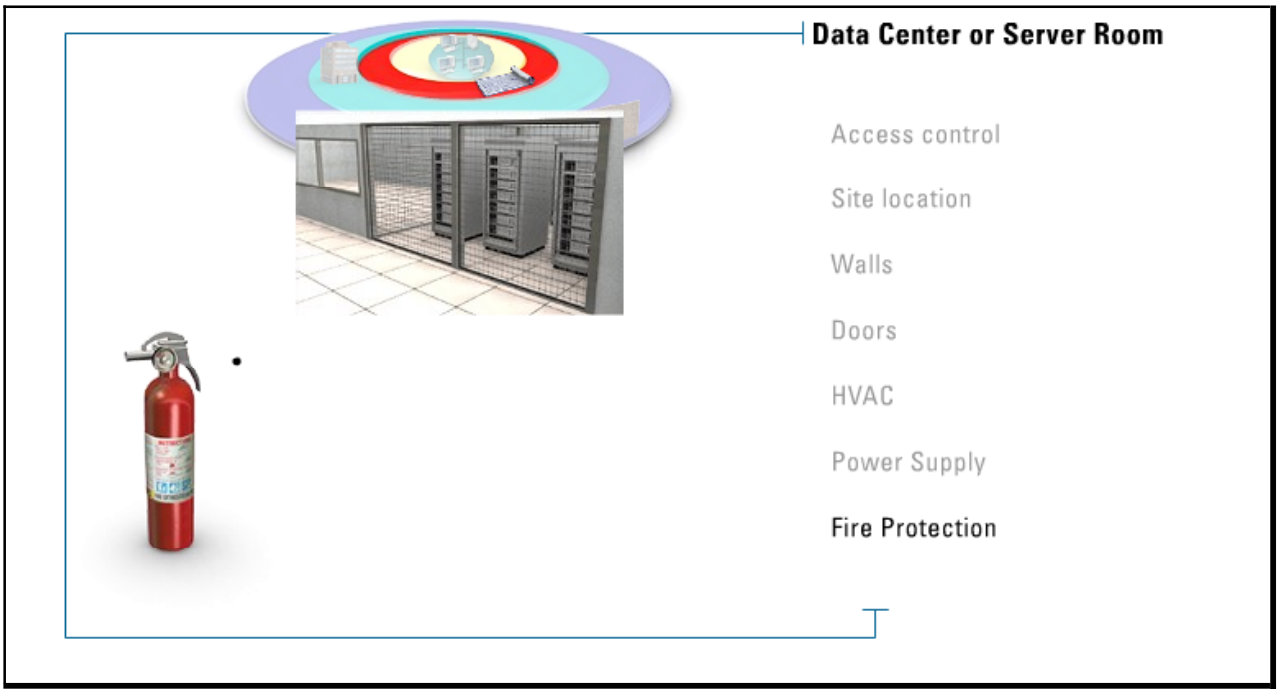

Good security practices in these rooms include:

- **Access Control**
  - Smart cards/badges required for access
  - Alarmed doors during non-working hours
  - Post an access control list stating who is allowed unescorted access
  - Document all visitors including why they were in the room
  - Create access control policies for working hours, non-working hours, and weekends
  - Enforce strict key control for locks
  - Change lock combination as appropriate
  - Use CCTV to view visitors
- **Site Location**
  - Location should not be easily reached by visitors or the general public
  - Location should be away from external windows or walls
  - Location should be away from water pipes or other support system facilities
- **Walls**
  - Room construction should be via a single unit
  - No wall should be part of the external facility structure

- Roof, floors, and walls should not adjoin an insecure area
- Use shatter-resistant glass if being used as wall barrier

- **Doors**
  - Doors should have a solid core
  - Doors should not open outward
  - Door hinges should be fixed to the frame
  - Door frames should be permanently fixed to wall studs
  - Review emergency door locking mechanisms
- **HVAC**
  - Should be a system separate from the rest of the building
  - Duct and vent size should not allow security breaches
  - Maintain positive pressure
  - Place barriers on ducts and vents if required
- **Power Supply**
  - A UPS or generator should exist for a certain amount of time defined by policy
  - Test backup systems regularly
  - Electrical facilities should be separate from the rest of the facility
  - Properly secure electrical closets
  - Properly secure and test cables and wiring
  - Provide emergency lighting
  - Emergency power-off switches should be located near all exits
  - Protect emergency power-off switches with protective plastic covers
- **Fire Protection**
  - Portable extinguishers should be located at exits and near equipment
  - Install fire sensors/detectors
  - Document and test emergency plans
  - Install water sensors under the raised floor

Computer Equipment Protection

The simplest method of protecting computers from theft is to bolt down the equipment. You can also use cables on servers and workstations with locks that require special keys, such as electronic tokens or smart cards.

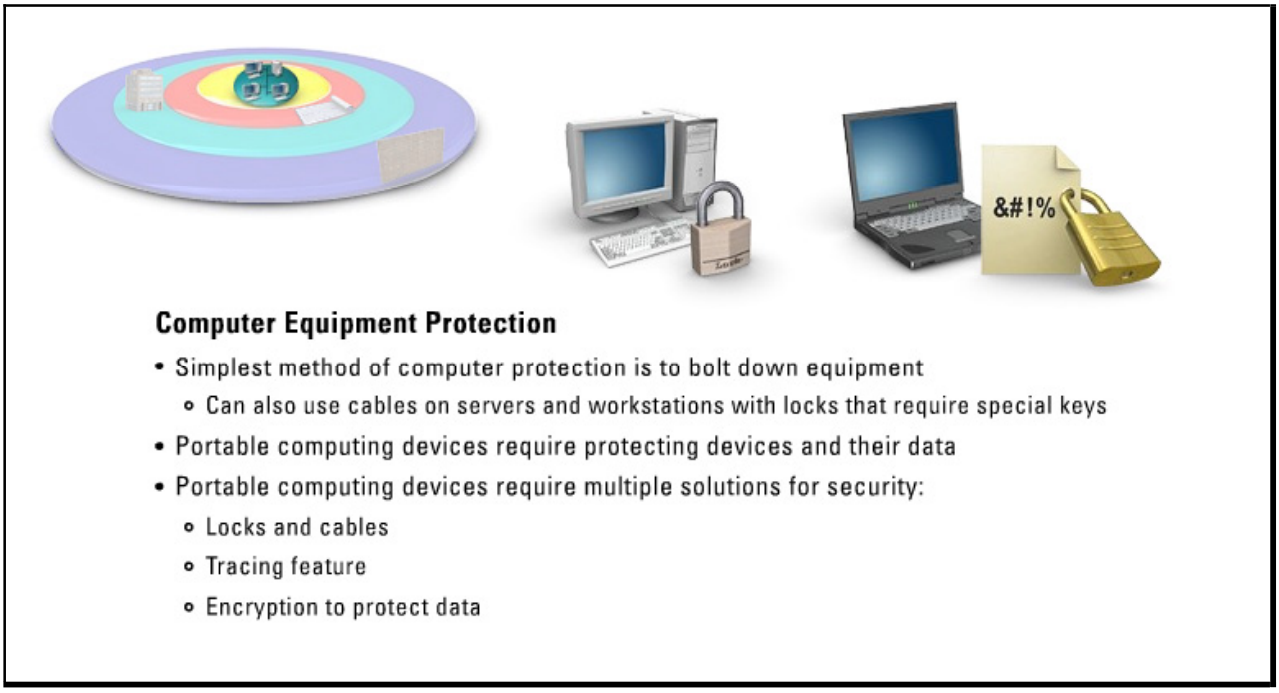

Portable computing devices, such as laptops, notebooks, and handheld devices like personal digital assistants (PDAs) require protecting the devices and their data, and keeping the security controls easy for the user. This means portable devices will require multiple solutions for security. For laptops and notebooks, docking stations should be used to secure the device, with locks and cables security the docking station.

Another possible security measure for these types of devices is a tracing feature. This is software that transmits a homing signal to a monitoring server. These signals can be transmitted over any Internet connection or phone line. If a portable device is missing, the user notifies the monitoring network. When the device is next connected, a message will be sent to the monitoring server indicating the device's location.

To protect data, you can install encryption software on the hard drive. All data on the hard drive are then encrypted and can only be accessed through some type of access control system. Solutions include entering a password, inserting a smart card, or using a biometric device.

Object Protection

The object level is the final layer of protection in the layered security scheme. These objects are items that you would place inside a security container, such as a safe, vault, or locking file cabinet. Safes are designed to be resistant to attack, but remember they are not attack proof. Given enough time and energy, any safe can be opened.

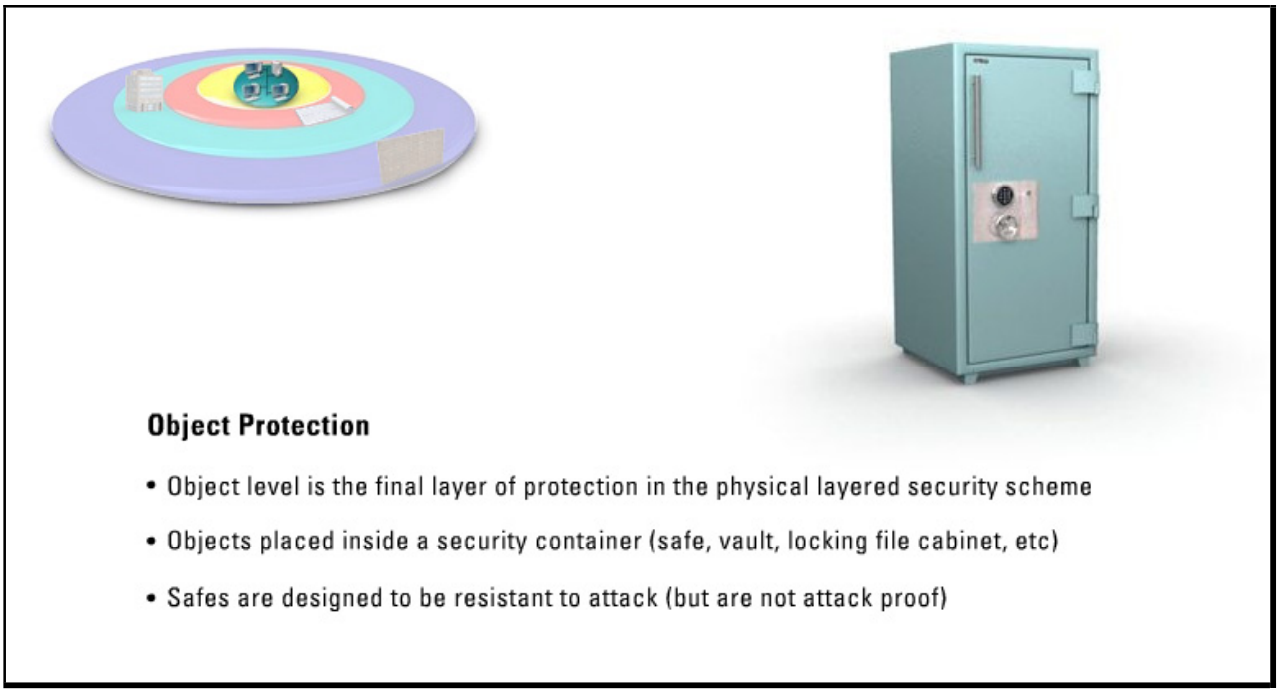

To determine what type of security container is required, it is important to first understand the security ramifications of the item that needs to be secured. For example, if the item is a sensitive paper object, a high security locking cabinet may be the way to go.

Combination dials and key locks are common for safes, but some safes have mechanical or electronic push-button locks. Security measures for locks include the following:

- If you have a lightweight safe, install it in reinforced concrete or anchor it to the building.
- Install the safe in a visible location. People should see who is attempting to access it.
- Use good combinations, and change the combinations often.
- Install a relocking device, which is a device inside the lock that will activate if a forced entry is attempted.

Summary

The key points discussed in this lesson are:

- To be effective, you must automatically control who can access a compartmentalized area and when.
- The data center room or server room requires strict security clearance, as integrity in these rooms must remain extremely high.
- The simplest method of protecting computers from theft is to bolt down the equipment. Another possible security measure for these types of devices is a tracing feature. To protect data, you can install encryption software on the hard drive.
- The object level is the final layer of protection in the layered security scheme. These objects are items that you would place inside a security container, such as a safe, vault, or locking file cabinet.