Federal Information
Processing Standards Publication 191

November 9, 1994

Specifications for

Guideline for The Analysis Local Area Network Security

Contents

1 INTRODUCTION

1.1 Why LAN Security is Important

Local area networks (LANs) have become a major tool to many organizations in meeting data processing and data communication needs. Prior to the use of LANs, most processing and communications were centralized; the information and control of that information were centralized as well. Now LANs logically and physically extend data, processing and communication facilities across the organization

Security services that protect the data, processing and communication facilities must also be distributed throughout the LAN. For example, sending sensitive files that are protected with stringent access controls on one system, over a LAN to another system that has no access control protection, defeats the efforts made on the first system. Users must ensure that their data and the LAN itself are adequately protected. LAN security should be an integral part of the whole LAN, and should be important to all users.

Electronic mail (email), a major application provided by most LANs, replaces much of the interoffice and even interorganizational mail that is written on paper and placed in an envelope. This envelope provides some confidentiality between the sender and receiver, and it can even be argued that the integrity of the paper envelope provides the receiver with some degree of assurance that the message was not altered. Using electronic mail does not provide these assurances. Simple transfers on unprotected LANs of inadequately protected electronic mail messages can be captured and read or perhaps even altered. For some LANs, there can be no assurance that the message actually was sent from the named sender. Fortunately tools such as encryption, digital signatures, and message authentication codes help solve these problems and can help provide some assurance.

Understanding the necessity to provide security on a LAN and how to decide the appropriate security measures needed are major goals of this document.

1.2 Purpose

The intended readers of this document include organizational management, LAN administrators, system administrators, security officers, LAN users and others who have a responsibility for protecting information processed, stored or associated with a LAN. The purpose of this document is to help the reader understand the need for LAN security and to provide guidance in determining effective LAN security controls.

1.3 Overview of Document

Section 1 - Introduction - This section discusses the properties of a LAN, and the security concerns that result from those properties.

Section 2 - Threats, Vulnerabilities, Security Services & Mechanisms - This section describes threats, related vulnerabilities and the possible security services and mechanisms that could be used to protect the LAN from these threats.

Section 3 - Risk Management - This section describes the risk management process and how it can be used to plan and implement appropriate LAN security.

1.4 LAN Definition

The Institute of Electrical and Electronic Engineers (IEEE) has defined a LAN as "a datacomm system allowing a number of independent devices to communicate directly with each other, within a moderately sized geographic area over a physical communications channel of moderate rates" [MART89]. Typically, a LAN is owned, operated, and managed locally rather than by a common carrier. A LAN usually, through a common network operating system, connects servers, workstations, printers, and mass storage devices, enabling users to share the resources and functionality provided by a LAN.

According to [BARK89] the types of applications provided by a LAN include distributed file storing, remote computing, and messaging.

1.4.1 Distributed File Storing

Distributed file storing provides users transparent access to part of the mass storage of a remote server. Distributed file storing provides capabilities such as a remote filing and remote printing. Remote filing allows users to access, retrieve, and store files. Generally remote filing is provided by allowing a user to attach to part of a remote mass storage device (a file server) as though it were connected directly. This virtual disk is then used as though it were a disk drive local to the workstation. Remote printing allows users to print to any printer attached to any component on the LAN. Remote printing addresses two user needs: ongoing processing while printing, and shared use of expensive printers. LAN print servers can accept files immediately, allowing users to continue work on their local workstations, instead of waiting for the print job to be completed. Many users utilizing the same printer can justify the cost of high quality, fast printers.

1.4.2 Remote Computing

Remote computing refers to the concept of running an application or applications on remote components. Remote computing allows users to (1) remotely login to another component on the LAN, (2) remotely execute an application that resides on another component, or (3) remotely run an application on one or more components, while having the appearance, to the user, of running locally. Remote login allows users to login to a remote system (such as a multi-user system) as though the user were directly connected to the remote system. The ability to run an application on one or more components allows the user to utilize the processing power of the LAN as a whole.

1.4.3 Messaging

Messaging applications are associated with mail and conferencing capabilities. Electronic mail has been one of the most used capabilities available on computer systems and across networks. Mail servers act as local post offices, providing users the ability to send and receive messages across a LAN. A conferencing capability allows users to actively communicate with each other, analogous to the telephone.

1.5 The LAN Security Problem

The advantages of utilizing a LAN were briefly discussed in the previous section. With these advantages however, come additional risks that contribute to the LAN security problem.

1.5.1 Distributed File Storing - Concerns

File servers can control users' accesses to various parts of the file system. This is usually done by allowing a user to attach a certain file system (or directory) to the user's workstation, to be used as a local disk. This presents two potential problems. First, the server may only provide access protection to the directory level, so that a user granted access to a directory has access to all files contained in that directory. To minimize risk in this situation, proper structuring and management of the LAN file system is important. The second problem is caused by inadequate protection mechanisms on the local workstation. For example, a personal computer (PC) may provide minimal or no protection of the information stored on it. A user that copies a file from the server to the local drive on the PC loses the protection afforded the file when it was stored on the server. For some types of information this may be acceptable. However, other types of information may require more stringent protections. This requirement focuses on the need for controls in the PC environment.

1.5.2 Remote Computing - Concerns

Remote computing must be controlled so that only authorized users may access remote components and remote applications. Servers must be able to authenticate remote users who request services or applications. These requests may also call for the local and remote servers to authenticate to each other. The inability to authenticate can lead to unauthorized users being granted access to remote servers and applications. There must be some level of assurance regarding the integrity of applications utilized by many users over a LAN.

1.5.3 Topologies and Protocols - Concerns

The topologies and protocols used today demand that messages be made available to many nodes in reaching the desired destination. This is much cheaper and easier to maintain than providing a direct physical path from every machine to every machine. (In large LANs direct paths are infeasible.) The possible threats inherent include both active and passive wiretapping. Passive wiretapping includes not only information release but also traffic analysis (using addresses, other header data, message length, and message frequency). Active wiretapping includes message stream modifications (including modification, delay, duplication, deletion or counterfeiting).

1.5.4 Messaging Services - Concerns

Messaging services add additional risk to information that is stored on a server or in transit. Inadequately protected email can easily be captured, and perhaps altered and retransmitted, effecting both the confidentiality and integrity of the message.

1.5.5 Other LAN Security Concerns

Other LAN security problems include (1) inadequate LAN management and security policies, (2) lack of training for proper LAN usage and security, (3) inadequate protection mechanisms in the workstation environment, and (4) inadequate protection during transmission.

A weak security policy also contributes to the risk associated with a LAN. A formal security policy governing the use of LANs should be in place to demonstrate management's position on the importance of protecting valued assets. A security policy is a concise statement of top management's position on information values, protection responsibilities, and organizational commitment. A strong LAN security policy should be in place to provide direction and support from the highest levels of management. The policy should identify the role that each employee has in assuring that the LAN and the information it carries are adequately protected.

The LAN security policy should stress the importance of, and provide support for, LAN management. LAN management should be given the necessary funding, time, and resources. Poor LAN management may result in security lapses. The resulting problems could include security settings becoming too lax, security procedures not being performed correctly, or even the necessary security mechanisms not being implemented.

The use of PCs in the LAN environment can also contribute to the risk of the LAN. In general, PCs have a relative lack of control with regard to authenticating users, controlling access to files, auditing, etc. In most cases the protection afforded information that is stored and processed on a LAN server does not follow the information when it is sent locally to a PC.

Lack of user awareness regarding the security of the LAN can also add risk. Users who are not familiar with the security mechanisms, procedures, etc. may use them improperly and perhaps less securely. Responsibilities for implementing security mechanisms and procedures and following the policies regarding the use of the PC in a LAN environment usually fall to the user of the PC. Users must be given the proper guidance and training necessary to maintain an acceptable level of protection in the LAN environment.

1.6 Goals of LAN Security

The following goals should be considered to implement effective LAN security.

- Maintain the confidentiality of data as it is stored, processed or transmitted on a LAN;

- Maintain the integrity of data as it is stored, processed or transmitted on a LAN;

- Maintain the availability of data stored on a LAN, as well as the ability to process and transmit the data in a timely fashion;

- Ensure the identity of the sender and receiver of a message;

Adequate LAN security requires the proper combination of security policies and procedures, technical controls, user training and awareness, and contingency planning. While all of these areas are critical to provide adequate protection, the focus of this document is on the technical controls that can be utilized. The other areas of control mentioned above are discussed in the appendices.

2 THREATS, VULNERABILITIES, SERVICES & MECHANISMS

A threat can be any person, object, or event that, if realized, could potentially cause damage to the LAN. Threats can be malicious, such as the intentional modification of sensitive information, or can be accidental, such as an error in a calculation, or the accidental deletion of a file. Threats can also be acts of nature, i.e. flooding, wind, lightning, etc. *The immediate damage caused by a threat is referred to as an impact.*

Vulnerabilities are weaknesses in a LAN that can be exploited by a threat. For example, unauthorized access (the threat) to the LAN could occur by an outsider guessing an obvious password. The vulnerability exploited is the poor password choice made by a user. Reducing or eliminating the vulnerabilities of the LAN can reduce or eliminate the risk of threats to the LAN. For example, a tool that can help users choose robust passwords may reduce the chance that users will utilize poor passwords, and thus reduce the threat of unauthorized LAN access.

A security service is the collection of security mechanisms, supporting data files, and procedures that help protect the LAN from specific threats. For example, the identification and authentication service helps protect the LAN from unauthorized LAN access by requiring that a user identify himself, as well as verifying that identity. The security service is only as robust as the mechanisms, procedures, etc. that make up the service.

Security mechanisms are the controls implemented to provide the security services needed to protect the LAN. For example, a token based authentication system (which requires that the user be in possession of a required token) may be the mechanism implemented to provide the identification and authentication service. Other mechanisms that help maintain the confidentiality of the authentication information can also be considered as part of the identification and authentication service.

This section is composed of two parts. The first part discusses threats, impacts and related vulnerabilities. The threats are generally categorized based on the impact caused if the threat is realized. For each impact category there is a discussion regarding the threats that may cause the impact, potential losses from the threat, and the vulnerabilities that may be exploited by the threat. The second part of this section discusses LAN security services and the possible mechanisms that can be implemented to provide these services.

2.1 Threats and Vulnerabilities

Identifying threats requires one to look at the impact and consequence of the threat if it is realized. The impact of the threat, which usually points to the immediate near-term problems, results in disclosure, modification, destruction, or denial of service. The more significant long-term consequences of the threat being realized are the result of lost business, violation of privacy, civil law suits, fines, loss of human life or other long term effects. Consequences of threats will be discussed in Section 3, *Risk Management*. The approach taken here is to categorize the types of impacts that can occur on a LAN so that specific technical threats can be grouped by the impacts and examined in a meaningful manner. For example, the technical threats that can lead to the impact 'LAN traffic compromise' in general can be distinguished from those threats that can lead to the impact 'disruption of LAN functionalities'. It should be recognized that many threats may result in more than one impact; however, for this discussion a particular threat will be discussed only in conjunction with one impact. The impacts that will be used to categorize and discuss the threats to a LAN environment are:

- **Unauthorized LAN access** - results from an unauthorized individual gaining access to the LAN.
- **Inappropriate access to LAN resources** - results from an individual, authorized or unauthorized, gaining access to LAN resources in an unauthorized manner.
- **Disclosure of data** - results from an individual accessing or reading information and possibly revealing the information in an accidental or unauthorized intentional manner.
- **Unauthorized Modification to data and software** - results from an individual modifying, deleting or destroying LAN data and software in an unauthorized or accidental manner.
- **Disclosure of LAN traffic** - results from an individual accessing or reading information and possibly revealing the information in an accidental or unauthorized intentional manner as it moves through the LAN.
- **Spoofing of LAN traffic** - results when a message appears to have been sent from a legitimate, named sender, when actually the message had not been.
- **Disruption of LAN functions** - results from threats that block LAN resources from being available in a timely manner.

2.1.1 Unauthorized LAN Access

LANs provide file sharing, printer sharing, file storage sharing, etc. Because resources are shared and not used solely by one individual there is need for control of the resources and accountability for use of the resources. *Unauthorized LAN access occurs when someone, who is not authorized to use the LAN, gains access to the LAN (usually by acting as a legitimate user of LAN).* Three common methods used to gain unauthorized access are password sharing, general password guessing and password capturing. Password sharing allows an unauthorized user to have the LAN access and privileges of a legitimate user; with the legitimate user's knowledge and acceptance. General password guessing is not a new means of unauthorized access. Password capturing is a process in which a legitimate user unknowingly reveals the user's login id and password. This may be done through the use of a trojan horse program that appears to the user as a legitimate login program; however, the trojan horse program is designed to capture passwords. Capturing a login id and password as it is transmitted across the LAN unencrypted is another method used to ultimately gain access. The methods to capture cleartext LAN traffic, including passwords, is readily available today. Unauthorized LAN access can occur by exploiting the following types of vulnerabilities:

- lack of, or insufficient, identification and authentication scheme,
- password sharing,
- poor password management or easy to guess passwords,
- using known system holes and vulnerabilities that have not been patched,
- single-user PCs that are not password protected at boot time,
- underutilized use of PC locking mechanisms,
- LAN access passwords that are stored in batch files on PCs,
- poor physical control of network devices,
- unprotected modems,
- lack of a time-out for login time period and log of attempts,
- lack of disconnect for multiple login failures and log of attempts,
- lack of 'last successful login date/time' and 'unsuccessful login attempt' notification and log,
- lack of real-time user verification (to detect masquerading).

2.1.2 Inappropriate Access to LAN Resources

One of the benefits of using a LAN is that many resources are readily available to many users, rather than each user having limited dedicated resources. These resources may include file stores, applications, printers, data, etc. However, not all resources need to be made available to each user. To prevent compromising the security of the resource (i.e. corrupting the resource, or lessening the availability of the resource), only those who require the use of the resource should be permitted to utilize that resource. *Unauthorized access occurs when a user, legitimate or unauthorized, accesses a resource that the user is not permitted to use.* Unauthorized access may occur simply because the access rights assigned to the resource are not assigned properly. However, unauthorized access may also occur because the access control mechanism or the privilege mechanism is not granular enough. In these cases, the only way to grant the user the needed access rights or privileges to perform a specific function is to grant the user more access than is needed, or more privileges than are needed. Unauthorized access to LAN resources can occur by exploiting the following types of vulnerabilities:

- use of system default permission settings that are too permissive to users,
- improper use of administrator or LAN manager privileges,
- data that is stored with an inadequate level or no protection assigned,
- lack of or the improper use of the privilege mechanism for users,
- PCs that utilize no access control on a file level basis.

2.1.3 Disclosure of Data

As LANs are utilized throughout an agency or department, some of the data stored or processed on a LAN may require some level of confidentiality. *The disclosure of LAN data or software occurs when the data or software is accessed, read and possibly released to an individual who is not authorized for the data.* This can occur by someone gaining access to information that is not encrypted, or by viewing monitors or printouts of the information. The compromise of LAN data can occur by exploiting the following types of vulnerabilities:

- improper access control settings,
- data, that has been deemed sensitive enough to warrant encryption, stored in unencrypted form,
- application source code stored in unencrypted form,
- monitors viewable in high traffic areas,
- printer stations placed in high traffic areas,
- data and software backup copies stored in open areas.

2.1.4 Unauthorized Modification of Data and Software

Because LAN users share data and applications, changes to those resources must be controlled. *Unauthorized modification of data or software occurs when unauthorized changes (additions, deletions or modifications) are made to a file or program.*

When undetected modifications to data are present for long periods of time, the modified data may be spread through the LAN, possibly corrupting databases, spreadsheet calculations, and other various application data. This can damage the integrity of most application information.

When undetected software changes are made, all system software can become suspect, warranting a thorough review (and perhaps reinstallation) of all related software and applications. These unauthorized changes can be made in simple command programs (for example in PC batch files), in utility programs used on multi-user systems, in major application programs, or any other type of software. They can be made by unauthorized outsiders, as well as those who are authorized to make software changes (although the changes they make are not authorized). These changes can divert information (or copies of the information) to other destinations, corrupt the data as it is processed, or harm the availability of system or LAN services.

PC viruses can be a nuisance to any organization that does not choose to provide LAN users the tools to effectively detect and prevent virus introduction to the LAN. Currently viruses have been limited to corrupting PCs, and generally do not corrupt LAN servers (although viruses can use the LAN to infect PCs). [WACK89] provides guidance on detecting and preventing viruses.

The unauthorized modification of data and software can occur by exploiting the following types of vulnerabilities:

- write permission granted to users who only require read permission to access,
- undetected changes made to software, including the addition of code to create a trojan horse program,
- lack of a cryptographic checksum on sensitive data,
- privilege mechanism that allow unnecessary write permission,
- lack of virus protection and detection tools.

2.1.5 Disclosure of LAN Traffic

The disclosure of LAN traffic occurs when someone who is unauthorized reads, or otherwise obtains, information as it is moved through the LAN. LAN traffic can be compromised by listening and capturing traffic transmitted over the LAN transport media (tapping into a network cable, listening to traffic transmitted over the air, misusing a provided network connection by attaching an analysis device, etc.). Many users realize the importance of confidential information when it is stored on their workstations or servers; however, it is also important to maintain that confidentiality as the information travels through the LAN. Information that can be compromised in this way includes system and user names, passwords, electronic mail messages, application data, etc. For example, even though passwords may be in an encrypted form when stored on a system, they can be captured in plaintext as they are sent from a workstation or PC to a file server. Electronic mail message files, which usually have very strict access rights when stored on a system, are often sent in plaintext across a wire, making them an easy target for capturing. The compromise of LAN traffic can occur by exploiting the following types of vulnerabilities:

- inadequate physical protection of LAN devices and medium,
- transmitting plaintext data using broadcast protocols,
- transmitting plaintext data (unencrypted) over the LAN medium,

2.1.6 Spoofing of LAN Traffic

Data that is transmitted over a LAN should not be altered in an unauthorized manner as a result of that transmission, either by the LAN itself, or by an intruder. LAN users should be able to have a reasonable expectation that the message sent, is received unmodified. *A modification occurs when an intentional or unintentional change is made to any part of the message including the contents and addressing information.*

Messages transmitted over the LAN need to contain some sort of addressing information that reports the sending address of the message and the receiving address of the message (along with other pieces of information). *Spoofing of LAN traffic involves (1) the ability to receive a message by masquerading as the legitimate receiving destination, or (2) masquerading as the sending machine and sending a message to a destination.* To masquerade as a receiving machine, the LAN must be persuaded into believing that the destination address is the legitimate address of the machine. (Receiving LAN traffic can also be done by listening to messages as they are broadcast to all nodes.) Masquerading as the sending machine to deceive a receiver into believing the message was legitimately sent can be done by masquerading the address, or by means of a playback. A playback involves capturing a session between a sender and receiver, and then retransmitting that message (either with the header only, and new message contents, or the whole message). The spoofing of LAN traffic or the modification of LAN traffic can occur by exploiting the following types of vulnerabilities:

Vulnerabilities

- transmitting LAN traffic in plaintext,
- lack of a date/time stamp (showing sending time and receiving time),
- lack of message authentication code mechanism or digital signature,
- lack of real-time verification mechanism (to use against playback).

2.1.7 Disruption of LAN Functions

A LAN is a tool, used by an organization, to share information and transmit it from one location to another. This need is satisfied by LAN functionalities such those described in Section 1.4, *LAN Definition*. *A disruption of functionality occurs when the LAN cannot provide the needed functionality in an acceptable, timely manner.* A disruption can interrupt one type of functionality or many. A disruption of LAN functionalities can occur by exploiting the following types of vulnerabilities:

Vulnerabilities

- inability to detect unusual traffic patterns (i.e. intentional flooding),
- inability to reroute traffic, handle hardware failures, etc,
- configuration of LAN that allows for a single point of failure,
- unauthorized changes made to hardware components (reconfiguring addresses on workstations, modifying router or hub configurations, etc.),
- improper maintenance of LAN hardware,
- improper physical security of LAN hardware.

2.2 Security Services and Mechanisms

A security service is the collection of mechanisms, procedures and other controls that are implemented to help reduce the risk associated with threat. For example, the identification and authentication service helps reduce the risk of the unauthorized user threat. Some services provide protection from threats, while other services provide for detection of the threat occurrence. An example of this would be a logging or monitoring service. The following services will be discussed in this section:

- **Identification and authentication** - is the security service that helps ensure that the LAN is accessed by only authorized individuals.
- **Access control** - is the security service that helps ensure that LAN resources are being utilized in an authorized manner.
- **Data and message confidentiality** - is the security service that helps ensure that LAN data, software and messages are not disclosed to unauthorized parties.
- **Data and message integrity** - is the security service that helps ensure that LAN data, software and messages are not modified by unauthorized parties.
- **Non-repudiation** - is the security service by which the entities involved in a communication cannot deny having participated. Specifically the sending entity cannot deny having sent a message (non-repudiation with proof of origin) and the receiving entity cannot deny having received a message (non-repudiation with proof of delivery).
- **Logging and Monitoring** - is the security service by which uses of LAN resources can be traced throughout the LAN.

The mechanisms, procedures and guidance provided in this section should not be considered as mandatory requirements in this document. This FIPS Guideline is voluntary, and the controls listed here should be considered as potential solutions, and not required solutions. Determining the appropriate controls and procedures to use in any LAN environment is the responsibility of those in each organization charged with providing adequate LAN protection.

2.2.1 Identification and Authentication

The first step toward securing the resources of a LAN is the ability to verify the identities of users [BNOV91]. The process of verifying a user's identity is referred to as authentication. Authentication provides the basis for the effectiveness of other controls used on the LAN. For example the logging mechanism provides usage information based on the userid. The access control mechanism permits access to LAN resources based on the userid. Both these controls are only effective under the assumption that the requestor of a LAN service is the valid user assigned to that specific userid.

Identification requires the user to be known by the LAN in some manner. This is usually based on an assigned userid. However the LAN cannot trust the validity that the user is in fact, who the user claims to be, without being authenticated. The authentication is done by having the user supply something that only the user has, such as a token, something that only the user knows, such as a password, or something that makes the user unique, such as a fingerprint. The more of these that the user has to supply, the less risk in someone masquerading as the legitimate user.

A requirement specifying the need for authentication should exist in most LAN policies. The requirement may be directed implicitly in a program level policy stressing the need to effectively control access to information and LAN resources, or may be explicitly stated in a LAN specific policy that states that all users must be uniquely identified and authenticated.

On most LANs, the identification and authentication mechanism is a userid/password scheme. [BNOV91] states that "password systems can be effective if managed properly [FIPS112], but seldom are. Authentication which relies solely on passwords has often failed to provide adequate protection for systems for a number of reasons. Users tend to create passwords that are easy to remember and hence easy to guess. On the other hand users that must use passwords generated from random characters, while difficult to guess, are also difficult to be remembered by users. This forces the user to write the password down, most likely in an area easy accessible in the work area". Research work such as [KLEIN] detail the ease at which passwords can be guessed. Proper password selection (striking a balance between being easy-to-remember for the user but difficult-to-guess for everyone else) has always been an issue. Password generators that produce passwords consisting of pronounceable syllables have more potential of being remembered than generators that produce purely random characters. [FIPS180] specifies an algorithm that can be used to produce random pronounceable passwords. Password checkers are programs that enable a user to determine whether a new passwords is considered easy-to-guess, and thus unacceptable.

Password-only mechanisms, especially those that transmit the password in the clear (in an unencrypted form) are susceptible to being monitored and captured. This can become a serious problem if the LAN has any uncontrolled connections to outside networks. Agencies that are considering connecting their LANs to outside networks, particularly the Internet, should examine [BJUL93] before doing so. If, after considering all authentication options, LAN policy determines that password-only systems are acceptable, the proper management of password creation, storage, expiration and destruction become all the more important. [FIPS 112] provides guidance on password management. [NCSC85] provides additional guidance that may be considered appropriate.

Because of the vulnerabilities that still exist with the use of password-only mechanisms, more robust mechanisms can be used. [BNOV91] discusses advances that have been made in the areas of token-based authentication and the use of biometrics. A smartcard based or token based mechanism requires that a user be in possession of the token and additionally may require the user to know a PIN or password. These devices then perform a challenge/response authentication scheme using realtime parameters. Using realtime parameters helps prevent an intruder from gaining unauthorized access through a login session playback. These devices may also encrypt the authentication session, preventing the compromise of the authentication information through monitoring and capturing.

Locking mechanisms for LAN devices, workstations, or PCs that require user authentication to unlock can be useful to users who must leave their work areas frequently. These locks allow users to remain logged into the LAN and leave their work areas (for an acceptable short period of time) without exposing an entry point into the LAN.

Modems that provide users with LAN access may require additional protection. An intruder that can access the modem may gain access by successfully guessing a user password. The availability of modem use to legitimate users may also become an issue if an intruder is allowed continual access to the modem.

Mechanisms that provide a user with his or her account usage information may alert the user that the account was used in an abnormal manner (e.g. multiple login failures). These mechanisms include notifications such as date, time, and location of last successful login, and number of previous login failures. The type of security mechanisms that could be implemented to provide the identification and authentication service are listed below.

Mechanisms

- password based mechanism,
- smartcards/smart tokens based mechanism,
- biometrics based mechanism,
- password generator,
- password locking,
- keyboard locking,
- PC or workstation locking,

- termination of connection after multiple failed logins
- user notification of 'last successful login' and 'number of login failures',
- real-time user verification mechanism,
- cryptography with unique user keys.

2.2.2 Access Control

This service protects against the unauthorized use of LAN resources, and can be provided by the use of access control mechanisms and privilege mechanisms. Most file servers and multi-user workstations provide this service to some extent. However, PCs which mount drives from the file servers usually do not. Users must recognize that files used locally from a mounted drive are under the access control of the PC. For this reason it may be important to incorporate access control, confidentiality and integrity services on PCs to whatever extent possible. Appendix C highlights some of the concerns that are inherent in the use of PCs.

According to [NCSC87], access control can be achieved by using discretionary access control or mandatory access control. Discretionary access control is the most common type of access control used by LANs. The basis of this kind of security is that an individual user, or program operating on the user's behalf is allowed to specify explicitly the types of access other users (or programs executing on their behalf) may have to information under the user's control. Discretionary security differs from mandatory security in that it implements the access control decisions of the user. Mandatory controls are driven by the results of a comparison between the user's trust level or clearance and the sensitivity designation of the information.

Access control mechanisms exist that support access granularity for acknowledging an owner, a specified group of users, and the world (all other authorized users). This allows the owner of the file (or directory) to have different access rights than all other users, and allows the owner to specify different access rights for a specified group of people, and also for the world. Generally access rights allow read access, write access, and execute access. Some LAN operating systems provide additional access rights that allow updates, append only, etc.

A LAN operating system may implement user profiles, capability lists or access control lists to specify access rights for many individual users and many different groups. Using these mechanisms allows more flexibility in granting different access rights to different users, which may provide more stringent access control for the file (or directory). (These more flexible mechanisms prevent having to give a user more access than necessary, a common problem with the three level approach.) Access control lists assign the access rights of named users and named groups to a file or directory. Capability lists and user profiles assign the files and directories that can be accessed by a named user.

User access may exist at the directory level, or the file level. Access control at the directory level places the same access rights on all the files in the directory. For example, a user that has read access to the directory can read (and perhaps copy) any file in that directory. Directory access rights may also provide an explicit negative access that prevents the user from any access to the files in the directory.

Some LAN implementations control how a file can be accessed. (This is in addition to controlling who can access the file.) Implementations may provide a parameter that allows an owner to mark a file sharable, or locked. Sharable files accept multiple accesses to the file at the same time. A locked file will permit only one user to access it. If a file is a read only file, making it sharable allows many users to read it at the same time.

These access controls can also be used to restrict usage between servers on the LAN. Many LAN operating systems can restrict the type of traffic sent between servers. There may be no restrictions, which implies that all users may be able to access resources on all servers (depending on the users access rights on a particular server). Some restrictions may be in place that allow only certain types of traffic, for example only electronic mail messages, and further restrictions may allow no exchange of traffic from server to server. The LAN policy should determine what types of information need to be exchanged between servers. Information that is not necessary to be shared between servers should then be restricted.

Privilege mechanisms enable authorized users to override the access permissions, or in some manner legally bypass controls to perform a function, access a file, etc. A privilege mechanism should incorporate the concept of least privilege. [ROBA91] defines least privilege as "a principle where each subject in a system be granted the most restrictive set or privileges needed for the performance of an authorized task." For example, the principle of least privilege should be implemented to perform the backup function. A user who is authorized to perform the backup function needs to have read access to all files in order to copy them to the backup media. (However the user should not be given read access to all files through the access control mechanism.) The user is granted a 'privilege' to override the read restrictions (enforced by the access control mechanism) on all files in order to perform the backup function. The more granular the privileges that can be granted, the more control there is not having to grant excessive privilege to perform an authorized function. For example, the user who has to perform the backup function does not need to have a write override privilege, but for privilege mechanisms that are less granular, this may occur. The types of security mechanisms that could be implemented to provide the access control service are listed below.

Mechanisms

- access control mechanism using access rights (defining owner, group, world permissions),
- access control mechanism using access control lists, user profiles, capability lists,
- access control using mandatory access control mechanisms (labels),

- granular privilege mechanism,

2.2.3 Data and Message Confidentiality

The data and message confidentiality service can be used when the secrecy of information is necessary. As a front line protection, this service may incorporate mechanisms associated with the access control service, but can also rely on encryption to provide further secrecy protection. Encrypting information converts it to an unintelligible form called ciphertext, decrypting converts the information back to its original form. Sensitive information can be stored in the encrypted, ciphertext, form. In this way if the access control service is circumvented, the file may be accessed but the information is still protected by being in encrypted form. (The use of encryption may be critical on PCs that do not provide an access control service as a front line protection.)

It is very difficult to control unauthorized access to LAN traffic as it is moved through the LAN. For most LAN users, this is a realized and accepted problem. The use of encryption reduces the risk of someone capturing and reading LAN messages in transit by making the message unreadable to those who may capture it. Only the authorized user who has the correct key can decrypt the message once it is received.

A strong policy statement should dictate to users the types of information that are deemed sensitive enough to warrant encryption. A program level policy may dictate the broad categories of information that need to be stringently protected, while a system level policy may detail the specific types of information and the specific environments that warrant encryption protection. At whatever level the policy is dictated, the decision to use encryption should be made by the authority within the organization charged with ensuring protection of sensitive information. If a strong policy does not exist that defines what information to encrypt, then the data owner should ultimately make this decision.

Cryptography can be categorized as either secret key or public key. Secret key cryptography is based on the use of a single cryptographic key shared between two parties . The same key is used to encrypt and decrypt data. This key is kept secret by the two parties. If encryption of sensitive but unclassified information (except Warner Amendment information) is needed, the use of the *Data Encryption Standard* (DES), FIPS 46-2, is required unless a waiver is granted by the head of the federal agency. The DES is a secret key algorithm used in a cryptographic system that can provide confidentiality. FIPS 46-2 provides for the implementation of the DES algorithm in hardware, software, firmware or some combination. This is a change from 46-1 which only provided for the use of hardware implementations. For an overview of DES, information addressing the applicability of DES, and waiver procedures see [NCSL90].

Public key cryptography is a form of cryptography which make use of two keys: a public key and a private key. The two keys are related but have the property that, given the public key, it is computationally infeasible to derive the private key [FIPS 140-1]. In a public key cryptosystem, each party has its own public/private key pair. The public key can be known by anyone; the private key is kept secret. An example for providing confidentiality is as follows: two users, Scott and Jeff, wish to exchange sensitive information, and maintain the confidentiality of that information. Scott can encrypt the information with Jeff's public key. The confidentiality of the information is maintained since only Jeff can decrypt the information using his private key. There is currently no FIPS approved public-key encryption algorithm for confidentiality. Agencies must waive FIPS 46-2 to use a public-key encryption algorithm for confidentiality. Public key technology, in the form of digital signatures, can also provide integrity and non-repudiation. This will be discussed in Section 2.2.4, *Data Integrity*.

FIPS 140-1, *Security Requirements for Cryptographic Modules*, should be used by agencies to specify the security requirements needed to protect the equipment that is used encryption. This standard specifies requirements such as authentication, physical controls and proper key management for all equipment that is used for encryption. Systems that implement encryption in software have additional requirements placed on them by FIPS 140-1. LAN servers, PCs, encryption boards, encryption modems, and all other LAN and data communication equipment that has an encryption capability should conform to the requirements of FIPS 140-1. The types of security mechanisms that could be implemented to provide the message and data confidentiality service are listed below.

Mechanisms

- file and message encryption technology,
- protection for backup copies on tapes, diskettes, etc,
- physical protection of physical LAN medium and devices,
- use of routers that provide filtering to limit broadcasting (either by blocking or by masking message contents).

2.2.4 Data and Message Integrity

The data and message integrity service helps to protect data and software on workstations, file servers, and other LAN components from unauthorized modification. The unauthorized modification can be intentional or accidental. This service can be provided by the use of cryptographic checksums, and very granular access control and privilege mechanisms. The more granular the access control or privilege mechanism, the less likely an unauthorized or accidental modification can occur.

The data and message integrity service also helps to ensure that a message is not altered, deleted or added to in any manner during transmission. (The inadvertent modification of a message packet is handled through the media access control implemented within the LAN protocol.) Most of the security techniques available today cannot prevent the modification of a message, but they can detect the modification of a message (unless the message is deleted altogether).

The use of checksums provide a modification detection capability. A Message Authentication Code (MAC), a type of cryptographic checksum, can protect against both accidental and intentional, but unauthorized, data modification. A MAC is initially calculated by applying a cryptographic algorithm and a secret value, called the key, to the data. The initial MAC is retained. The data is later verified by applying the cryptographic algorithm and the same secret key to the data to produce another MAC; this MAC is then compared to the initial MAC. If the two MACs are equal, then the data is considered authentic. Otherwise, an unauthorized modification is assumed. Any party trying to modify the data without knowing the key would not know how to calculate the appropriate MAC corresponding to the altered data. FIPS 113, *Computer Data Authentication*, defines the Data Authentication Algorithm, based on the DES, which is used to calculate the MAC. See [SMID88] for more information regarding the use of MACs.

The use of electronic signatures can also be used to detect the modification of data or messages. An electronic signature can be generated using public key or private key cryptography. Using a public key system, documents in a computer system are electronically signed by applying the originator's private key to the document. The resulting digital signature and document can then be stored or transmitted. The signature can be verified using the public key of the originator. If the signature verifies properly, the receiver has confidence that the document was signed using the private key of the originator and that the message had not been altered after it was signed. Because private keys are known only to their owner, it may also possible to verify the originator of the information to a third party. A digital signature, therefore, provides two distinct services: nonrepudiation and message integrity. FIPS PUB 186, Digital Signature Standard, specifies a digital signature algorithm that should be used when message and data integrity are required.

The message authentication code (MAC) described above can also be used to provide an electronic signature capability. The MAC is calculated based on the contents of the message. After transmission another MAC is calculated on the contents of the received message. If the MAC associated with the message that was sent is not the same as the MAC associated with the message that was received, then there is proof that the message received does not exactly match the message sent. A MAC can be used to identify the signer of the information to the receiver. However, the implementations of this technology do not inherently provide nonrepudiation because both the sender of the information and the receiver of the information share the same key. The types of security mechanisms that could be implemented to provide the data and message integrity service are listed below.

Mechanisms

- message authentication codes used for software or files,
- use of secret key based electronic signature,
- use of public key digital signature,
- granular privilege mechanism,
- appropriate access control settings (i.e. no unnecessary write permissions),
- virus detection software,
- workstations with no local storage (to prevent local storage of software and files),
- workstations with no diskette drive/tape drive to prevent introduction of suspect software.
- use of public key digital signatures.

2.2.5 Non-repudiation

Non-repudiation helps ensure that the entities in a communication cannot deny having participated in all or part of the communication. When a major function of the LAN is electronic mail, this service becomes very important. Non-repudiation with proof of origin gives the receiver some confidence that the message indeed came from the named originator. The nonrepudiation service can be provided through the use of public key cryptographic techniques using digital signatures. See Section 2.2.4 *Data and Message Integrity* for a description and use of digital signatures. The security mechanism that could be implemented to provide the non-repudiation service is listed below.

Mechanisms

- use of public key digital signatures.

2.2.6 Logging and Monitoring

This service performs two functions. The first is the detection of the occurrence of a threat. (However, the detection does not occur in real time unless some type of real-time monitoring capability is utilized.) Depending on the extensiveness of the logging, the detected event should be traceable throughout the system. For example, when an intruder breaks into the system, the log should indicate who was logged on to the system at the time, all sensitive files that had failed accesses, all programs that had attempted executions, etc. It should also indicate sensitive files and programs that were successfully accessed in this time period. It may be appropriate that some areas of the LAN (workstations, fileservers, etc.) have some type of logging service.

The second function of this service is to provide system and network managers with statistics that indicate that systems and the network as a whole are functioning properly. This can be done by an audit mechanism that uses the log file as input and processes the file into meaningful information regarding system usage and security. A monitoring capability can also be used to detect LAN availability problems as they develop. The types of security mechanisms that could be used to provide the logging and monitoring service are listed below.

Mechanisms

- logging of I&A information (including source machine, modem, etc.),
- logging of changes to access control information,
- logging of use of sensitive files,
- logging of modifications made to critical software,
- utilizing LAN traffic management tools,
- use of auditing tools.

3 RISK MANAGEMENT

A systematic approach should be used to determine appropriate LAN security measures. Deciding how to address security, where to implement security on the LAN, and the type and strength of the security controls requires considerable thought. This section will address the issues involving risk management of a LAN. The elements that are common to most risk management processes will be examined in terms of the unique properties of a LAN that may require special considerations beyond the risk process of a centralized system or application. In presenting this information, a simple risk management methodology will be introduced that may be considered as a candidate among the different methodologies and techniques that are currently available.

It is the reader's task to determine the appropriate level of protection required for his or her LAN. This is accomplished through risk management. [KATZ92] defines risk management as the process of:

- estimating potential losses due to the use of or dependence upon automated information system technology,
- analyzing potential threats and system vulnerabilities that contribute to loss estimates, and
- selecting cost effective safeguards that reduce risk to an acceptable level.

There are many risk management methodologies that an organization may use. However all should incorporate the process defined above.

3.1 Current Approaches

One of the most important considerations in choosing a methodology or technique is that the results obtained from the risk assessment be useful in providing LAN security. If the methodology is too complicated to use, if it requires input data that is too detailed, or if it produces results that are too intricate to infer what the risk to the LAN actually is, the methodology will not be useful and will not lead to effective LAN security. On the other hand, if the methodology does not allow for reasonable granularity in its definition of variables such as loss, likelihood and cost, the results produced may be too simple and may not reflect the true risk to the LAN. Those responsible within the organization should adopt the risk assessment approach that provides a technique that is understandable, easily used, and produces results that helps the organization to effectively secure its LANs.

In 1979, NIST published FIPS 65 [FIPS65] which described a quantitative method for performing risk analysis. This document was issued as a guideline and not a standard. Therefore the use of FIPS 65 is **not** mandatory for performing risk analysis. [KATZ92] points out that its primary use was for the risk analysis of large data centers. [FIPS65] describes how an estimate of risk (i.e. Annual Loss Expectancy) could be obtained by estimating, for each application data file: (1) the frequency of occurrence of harmful impact (i.e., destruction, modification, disclosure or unavailability of the data file) and (2) the consequences (in dollars) that could result from each of the impacts [KATZ92]. [KATZ92] explains that "recognizing the lack of empirical data on frequency of occurrence of impacts and the related consequences, FIPS 65 suggested an 'order of magnitude approach' to approximating these values. That this concept was not well understood by users of that method has been illustrated by numerous attempts to be too precise in quantifying the input data to FIPS 65 and, by the same token, interpreting the results as having more precision than they actually had. " FIPS 65 may be used for a risk assessment of a LAN; however agencies may choose other methodologies and techniques if the agency finds them to be more appropriate and effective.

Automated risk analysis tools are available that are tailored specifically to the LAN environment. [GILB89] points out the many benefits of using automated risk analysis tools. However there is a concern in using automated risk analysis tools. There are many techniques available to calculate risk. While most depend on a loss variable and a likelihood or probability variable, the manner in which these variables are represented, the calculations that are used on these variables, and the manner in which the risk value is represented is not always made available to the user. This disadvantage is compounded because there is currently no standard method or agreed upon approach for performing risk analysis. While there exists a proposed standard framework [KATZ92] for risk analysis that provides vendors with some guidance in developing these tools, there are no agreed upon methods for representing the necessary variables to perform a risk analysis, and there are no agreed upon methods for calculating risk using these variables. Because of this lack of consistent agreement with the risk community, coupled with the proprietary nature of the tools, determining the effectiveness of any particular method may be difficult. On the other hand, if the methodology used by the tool is understood and deemed acceptable for the user, then the tool may prove to be quite adequate. The underlying question in determining if a tool will be effective for a particular environment should be, "What is the automated risk analysis tool measuring, and are the results produced by it useful for providing appropriate LAN security?" [GILB89] discusses the use of automated risk analysis tools, and examines criteria that can be considered in the automated tools selection process.

Another approach for performing risk analyses is to develop sets of baseline security controls needed for predefined levels of risk. The predefined levels of risk may be based on the asset alone (e.g. data is considered sensitive due to an agency policy or federal mandate), the consequence that would result from the loss of the asset (e.g. the agency may not be able to meet its mission) or other factors. This allows data owners and those responsible for ensuring the security of the LAN to determine the level of risk for specific assets, and follow the guidance and implement the controls that have been deemed appropriate. This approach may provide an agency with the benefit of having consistent protection for specified types of assets. This approach has been implemented in [DOE89], [HHS91], [NASA90]. A benefit of this approach is that the user is not only provided with a risk analysis methodology, but also with an awareness and understanding of the agency policies that have derived the baseline controls. In organizations where the responsibility for security resides with someone who is not a security practitioner, this approach may provide enough knowledge and direction to provide effective security.

Other methodologies and approaches are available. Some require a manual process; others are implemented in software. Whatever risk analysis method is chosen by an organization, it must be effective in helping to implement effective LAN security and thus reduce the risk to the LAN.

3.2 Participants

LAN security should address the concerns and needs of the organization as a whole. This perspective can only be obtained by including representatives from relevant areas of the organization. Minimally this should include:

- **LAN Management** is responsible for the operation of the LAN. LAN Management can provide the risk assessment group the correct LAN configurations, including hardware, software, data, and functionality mapping. LAN Management can also determine the immediate impacts that can occur if a threat is realized.

- **Organizational Management** is responsible for supporting the LAN security policy by providing funding to implement required security services and making a commitment to ensure compliance with policy goals. Organizational management has the proper perspective in assessing the longterm consequences to the organization if a threat is realized.

- **Security Personnel** are responsible for ensuring that organizational security policies are developed and adhered to.

- **Data and Application Owners** are responsible for ensuring that their data and applications are adequately protected and are available to authorized users.

- **LAN Users** are responsible for providing accurate information about their applications, data and LAN usage.

The above list generally represents those individuals involved in the risk analysis of most computer systems and applications (with the exception of LAN management if there is no network). What is unique about this list with regard to forming a team to assess LAN risks is that each group listed above may be multiplied to account for each part of an organization the LAN serves, each application that is processed on the LAN, and for the different requirements and mandates that are in place throughout the organization. The requirements of the "LAN owner" in addition to the needs of many data and application owners have to all be considered.

The ultimate goal of effective **overall** LAN security may not be met if strong team leadership is not in place from the beginning. For example, organizations that lack strong centralized management of the LAN may have a difficult time assessing needs in any hierarchical manner, since each local manager or application owner may view his needs as a priority over other local managers and application owners, regardless of what the risk analysis results indicate.

Initially, those within the organization charged with performing the risk analysis need to make some determination regarding the proposed scope and boundary of the risk analysis. With this information, the necessary participants in the risk process can be chosen.

3.3 Elements of Risk Management

Operation of a LAN involves risk. The term *risk management* is commonly used to define the process of determining risk, applying controls to reduce the risk, and then determining if the residual risk is acceptable. Risk management supports two goals: measure risk (risk assessment) and selecting appropriate controls that will reduce risk to an acceptable level (risk mitigation). Issues that should be addressed when assessing LAN security include:

1. Assets - What should be protected?

2. Threats - From what do the assets need protection and what is the likelihood that a threat will occur?

3. Impacts - What are the immediate damages if the threat is realized (e.g. disclosure of information, modification of data)?

4. Consequences - What are the long-term effects of the threat being realized (e.g. damage to reputation of organization, loss of business)?

5. Controls - What are the effective security measures (security services and mechanisms) needed to protect the assets?

6. Risk - After implementation of the security controls, is the remaining risk acceptable?

The goal of risk assessment is to determine the risk to the LAN. The risk assessment process is conducted in two steps. The first step defines the boundary of the environment, determines the scope of the assessment and selects the appropriate methodology to use. In step two the risk analysis is conducted. The risk analysis can be broken down into asset identification, threat and vulnerability identification, likelihood assessment, and risk measure.

The goal of risk mitigation is to apply effective security controls such that the residual risk to the LAN is acceptable. Risk mitigation involves three steps: determining those areas where risk is unacceptable; selecting effective safeguards, and evaluating the controls and determining if the residual risk to the LAN is acceptable.

Organizations can select from a variety of risk management methodologies. The goal is for an organization to choose the most effective approach for the organization. The methodology discussed here consists of seven processes (outlined in Figure 3.1).

Figure 3.1 - Risk Management Process

1. Define the Scope and Boundary and Methodology

2. Identify and Value Assets,

3. Identify Threats and Determine Likelihood,

4. Measure Risk,

5. Select Appropriate Safeguards,

6. Implement and Test Safeguards,

7. Accept Residual Risk.

3.4 Risk Assessment

3.4.1 Process 1 - Define the Scope and Boundary, and Methodology

This process determines the direction that the risk management effort will take. It defines how much of the LAN (the boundary) and in how much detail (the scope) the risk management process should entail. The boundary will define those parts of the LAN that will be considered. The boundary may include the LAN as a whole or parts of the LAN, such as the data communications function, the server function, the applications, etc. Factors that determine the boundary may be based on LAN ownership, management or control. Placing the boundary around a part of the LAN controlled elsewhere may result in cooperation problems that may lead to inaccurate results. This problem stresses the need for cooperation among those involved with the ownership and management of the different parts of the LAN, as well as the applications and information processed on it.

The scope of the risk management effort must also be defined. The scope can be thought of as a logical outline showing, within the boundary, the depth of the risk management process. The scope distinguishes the different areas of the LAN (within the boundary) and the different levels of detail used during the risk management process. For example some areas may be considered at a higher or broader level, while other areas may be treated in depth and with a narrow focus. For smaller LANs, the boundary may be the LAN as a whole, and the scope may define a consistent level of detail throughout the LAN. For larger LANs, an organization may decide to place the boundary around those areas that it controls and to define the scope to consider all areas within the boundary. However the focus on data communications, external connections, and certain applications might be more narrow. Changes in the LAN configuration, the addition of external connections, or updates or upgrades to LAN software or applications may influence the scope.

The appropriate risk management methodology for the LAN may have been determined prior to defining the boundary and scope. If the methodology has already been determined, then it may be useful to scrutinize the chosen methodology given the defined boundary and scope. If a methodology has not been chosen, the boundary and scope information may be useful in selecting a methodology that produces the most effective results.

3.4.2 Process 2 - Identify and Value Assets

Asset valuation identifies and assigns value to the assets of the LAN. All parts of the LAN have value although some assets are definitely more valuable than others. This step gives the first indication of those areas where focus should be placed. For LANs that produce large amounts of information that cannot be reasonably analyzed, initial screening may need to be done. Defining and valuing assets may allow the organization to initially decide those areas that can be filtered downward and those areas that should be flagged as a high priority.

Different methods can be used to identify and value assets. The risk methodology that an organization chooses may provide guidance in identifying assets and should provide a technique for valuing assets. Generally assets can be valued based on the impact and consequence to the organization. This would include not only the replacement cost of the asset, but also the effect on the organization if the asset is disclosed, modified, destroyed or misused in any other way.

Figure 3.2 - Simple Asset Valuation

The value of the asset can be represented in terms of the potential loss. This loss can be based on the replacement value, the immediate impact of the loss, and the consequence. One of the simplest valuing techniques to indicate the loss of an asset is to use a qualitative ranking of high, medium and low. Assigning values to these rankings (3=high, 2=medium, and 1=low) can assist in the risk measure process.

Because the value of an asset should be based on more than just the replacement cost, valuing assets is one of the most subjective of the processes. However, if asset valuation is done with the goal of the process in mind, that is, to define assets in terms of a hierarchy of importance or criticality, the relativeness of the assets becomes more important than placing the "correct" value on them.

The risk assessment methodology should define the representation of the asset values. Purely quantitative methodologies such as FIPS 65 may use dollar values. However having to place a dollar value on some of the consequences that may occur in today's environments may be sufficient to change the perception of the risk management process from being challenging to being unreasonable.

Many risk assessment methodologies in use today require asset valuation in more qualitative terms. While this type of valuation may be considered more subjective than a quantitative approach, if the scale used to value assets is utilized consistently throughout the risk management process, the results produced should be useful. Figure 3.2 shows one of the simplest methods for valuing assets.

Throughout this discussion of the risk management process, a simple technique for valuing assets (as shown in Figure 3.2), determining risk measure, estimating safeguard cost, and determining risk mitigation will be presented. This technique is a simple, yet valid technique; it is being used here to show the relationship between the processes involved in risk management. The technique is not very granular and may not be appropriate for environments where replacement costs, sensitivities of information and consequences vary widely.

One of the implicit outcomes of this process is that a detailed configuration of the LAN, as well as its uses is produced. This configuration should indicate the hardware incorporated, major software applications used, significant information processed on the LAN, as well as how that information flows through the LAN. The degree of knowledge of the LAN configuration will depend on the defined boundary and scope. Figure 3.3 exemplifies some of the areas that should be included.

Figure 3.3 - Defining the LAN Configuration

Hardware configuration - includes servers, workstations, PCs, peripheral devices, external connections, cabling maps, bridges or gateway connections, etc.

Software configuration - includes server operating systems, workstation and PC operating systems, the LAN operating system, major application software, software tools, LAN management tools, and software under development. This should also include the location of the software on the LAN and from where it is commonly accessed.

Data - Includes a meaningful typing of the data processed and communicated through the LAN, as well as the types of users who generally access the data. Indications of where the data is accessed, stored and processed on the LAN is important. Attention to the sensitivity of the data should also be considered.

After the LAN configuration is completed, and the assets are determined and valued, the organization should have a reasonably correct view of what the LAN consists of and what areas of the LAN need to be protected.

3.4.3 Process 3 - Identify Threats and Determine Likelihood

The outcome of this process should be a strong indication of the adverse actions that could harm the LAN, the likelihood that these actions could occur, and the weaknesses of the LAN that can be exploited to cause the adverse action. To reach this outcome, threats and vulnerabilities need to be identified and the likelihood that a threat will occur needs to be determined.

Large amounts of information on various threats and vulnerabilities exist. The Reference and Further Reading Sections of this document provide some information on LAN threats and vulnerabilities. Some risk management methodologies also provide information on potential threats and vulnerabilities. User experience and LAN management experience also provide insight into threats and vulnerabilities.

The degree to which threats are considered will depend on the defined boundary and scope defined for the risk management process. A high level analysis may point to threats and vulnerabilities in general terms; a more focused analysis may tie a threat to a specific component or usage of the LAN. For example a high level analysis may indicate that the consequence due to loss of data confidentiality through disclosure of information on the LAN is too great a risk. A more narrowly focused analysis may indicate that the consequence due to disclosure of personnel data captured and read through LAN transmission is too great a risk. More than likely, the generality of the threats produced in the high level analysis, will, in the end, produce safeguard recommendations that will also be high level. This is acceptable if the risk assessment was scoped at a high level. The more narrowly focused assessment will produce a safeguard that can specifically reduce a given risk, such as the disclosure of personnel data.

The threats and vulnerabilities discussed in Section 2 may be used as a starting point, with other sources included where appropriate. New threats and vulnerabilities should be addressed when they are encountered. Any asset of the LAN that was determined to be important enough (i.e., was not filtered through the screening process) should be examined to determine those threats that could potentially harm it. For more focused assessments, particular attention should be paid to detailing the ways that these threats could occur. For example, methods of attack that result in unauthorized access may be from a login session playback, password cracking, the attachment of unauthorized equipment to the LAN, etc. These specifics provide more information in determining LAN vulnerabilities and will provide more information for proposing safeguards.

This process may uncover some vulnerabilities that can be corrected by improving LAN management and operational controls immediately. These improved controls will usually reduce the risk of the threat by some degree, until such time that more thorough improvements are planned and implemented. For example, increasing the length and composition of the password for authentication may be one way to reduce a vulnerability to guessing passwords. Using more robust passwords is a measure that can be quickly implemented to increases the security of the LAN. Concurrently, the planning and implementation of a more advanced authentication mechanism can occur.

Figure 3.4 Assigning Likelihood Measure

> The likelihood of the threat occurring can be normalized as a value that ranges from 1 to 3. A 1 will indicate a low likelihood, a 2 will indicate a moderate likelihood and a 3 will indicate a high likelihood.

Existing LAN security controls should be analyzed to determine if they are currently providing adequate protection. These controls may be technical, procedural, etc. If a control is not providing adequate protection, it can be considered a vulnerability. For example, a LAN operating system may provide access control to the directory level, rather than the file level. For some users, the threat of compromise of information may be too great not to have file level protection. In this example, the lack of granularity in the access control could be considered a vulnerability.

As specific threats and related vulnerabilities are identified, a likelihood measure needs to be associated with the threat/vulnerability pair (i.e. What is the likelihood that a threat will be realized, given that the vulnerability is exploited?). The risk methodology chosen by the organization should provide the technique used to measure likelihood. Along with asset valuation, assigning likelihood measures can also be a subjective process. Threat data for traditional threats (mostly physical threats) does exist and may aid in determining likelihood. However experience regarding the technical aspects of the LAN and knowledge of operational aspects of the organization may prove more valuable to decide likelihood measure. Figure 3.4 defines a simple likelihood measure. This likelihood measure coincides with the asset valuation measure defined in Figure 3.1. Although the asset valuation and the likelihood measures provided in this example appear to be weighted equally for each threat/vulnerability pair, it is a user determination regarding which measure should be emphasized during the risk measurement process.

3.4.4 Process 4 - Measure Risk

In its broadest sense the risk measure can be considered the representation of the kinds of adverse actions that may happen to a system or organization and the degree of likelihood that these actions may occur. The outcome of this process should indicate to the organization the degree of risk associated with the defined assets. This outcome is important because it its the basis for making safeguard selection and risk mitigation decisions.

There are many ways to measure and represent risk. [KATZ92] points out that depending on the particular methodology or approach, the measure could be defined in qualitative terms, quantitative terms, one dimensional, multidimensional, or some combination of these. The risk measure process should be consistent with (and more than likely defined by) the risk assessment methodology being used by the organization. Quantitative approaches are often associated with measuring risk in terms of dollar losses (e.g. FIPS 65). Qualitative approaches are often associated with measuring risk in terms

Figure 3.5 - One Dimensional Approach to Calculate Risk

The risk associated with a threat can be considered as a function of the relative likelihood that the threat can occur, and the expected loss incurred given that the threat occurred. The risk is calculated as follows:

risk = likelihood of threat occurring (given the specific vulnerability) x loss incurred

The value estimated for loss is determined to be a value that ranges from 1 to 3. Therefor risk may be calculated as a number ranging from 1 to 9 meaning a risk of 1 or 2 is considered a low risk, a risk of 3 or 4 would be a moderate risk, and a risk of 6 or 9 would be considered a high risk.

LIKELIHOOD	LOSS	RISK
1	1	1 - LOW
1	2	2 - LOW
1	3	3 - MODERATE
2	1	2 - LOW
2	2	4 - MODERATE
2	3	6 - HIGH
3	1	3 - MODERATE
3	2	6 - HIGH
3	3	9 - HIGH

of quality as indicated through a scale or ranking. One dimensional approaches consider only limited components (e.g. risk = magnitude of loss X frequency of loss). Multidimensional approaches consider additional components in the risk measurement such as reliability, safety, or performance. One of the most important aspects of risk measure is that the representation be understandable and meaningful to those who need to make the safeguard selection and risk mitigation decisions.

Figure 3.6 - Calculating Cost Measure

> In this example cost measure, the cost of the safeguard is the amount needed to purchase or develop and implement each of the mechanisms. The cost can be normalized in the same manner as was the value for potential loss incurred. A 1 will indicate a mechanism with a low cost, a 2 will indicate a mechanism with a moderate cost, and a 3 will indicate a mechanism with a high cost.

Figure 3.5 provides an example of a one dimensional approach for calculating risk. In this example, the levels of risk are now normalized (i.e. low, medium and high) and can be used to compare risks associated with each threat. The comparison of risk measures should factor in the criticality of the components used to determine the risk measure. For simple methodologies that only look at loss and likelihood, a risk measure that was derived from a high loss and low likelihood may result in the same risk measure as one that resulted from a low loss and high likelihood. In these cases, the user needs to decide which risk measure to consider more critical, even though the risk measures may be equal. In this case, a user may decide that the risk measure derived from the high loss is more critical than the risk measure derived from the high likelihood.

With a list of potential threats, vulnerabilities and related risks, an assessment of the current security situation for the LAN can be determined. Areas that have adequate protection will not surface as contributing to the risk of the LAN (since adequate protection should lead to low likelihood) whereas those areas that have weaker protection do surface as needing attention.

3.5 Risk Mitigation

3.5.1 Process 5 - Select Appropriate Safeguards

The purpose of this process is to select appropriate safeguards. This process can be done using risk acceptance testing.

Risk acceptance testing is described by [KATZ92] as an activity that compares the current risk measure with acceptance criteria and results in a determination of whether the current risk level is acceptable. While effective security and cost considerations are important factors, there may be other factors to consider such as: organizational policy, legislation and regulation, safety and reliability requirements, performance requirements, and technical requirements.

The relationship between risk acceptance testing and safeguard selection can be iterative. Initially, the organization needs to order the different risk levels that were determined during the risk assessment. Along with this the organization needs to decide the amount of residual risk that it will be willing to accept after the selected safeguards are implemented. These initial risk acceptance decisions can be factored into the safeguard selection equation. When the properties of the candidate safeguards are known, the organization can reexamine the risk acceptance test measures and determine if the residual risk is achieved, or alter the risk acceptance decisions to reflect the known properties of the safeguards. For example there may be risks that are determined to be too high. However after reviewing the available safeguards, it may be realized that the currently offered solutions are very costly and cannot be easily implemented into the current configuration and network software. This may force the organization into either expending the resources to reduce the risk, or deciding through risk acceptance that the risk will have to be accepted because it is currently too costly to mitigate.

Many sources exist that can provide information on potential safeguards (See the Reference and Further Reading Sections). The methodology discussed here defines safeguards in terms of security services and mechanisms. A security service is the sum of mechanisms, procedures, etc. that are implemented on the LAN to provide protection. The security services (and mechanisms) provided in Section 2 can be used as a starting point. The security services should be related to the threats defined in the risk assessment.

In most cases the need for a specific service should be readily apparent. If the risk acceptance results indicate that a risk is acceptable, (i.e., existing mechanisms are adequate) then there is no need to apply additional mechanisms to the service that already exists.

Figure 3.7 - Comparing Risk and Cost

> To calculate risk/cost relationships use the risk measure and the cost measure associated with each threat/mechanism relationship and create a ratio of the risk to the cost (i.e., risk/cost). A ratio that is less than 1 will indicate that the cost of the mechanism is greater than the risk associated with the threat. This is generally not an acceptable situation (and may be hard to justify) but should not be automatically dismissed. Consider that the risk value is a function of both the loss measure and the likelihood measure. One or both of these may represent something so critical about the asset that the costly mechanism is justified. This situation may occur when using simple methodologies such as this one.

After the needed security services are determined, consider the list of security mechanisms for each service. For each security service selected, determine the candidate mechanisms that would best provide that service. Using the threat/vulnerability/risk relationships developed in the previous processes, choose those mechanisms that could potentially reduce or eliminate the vulnerability and thus reduce the risk of the threat. In many cases, a threat/vulnerability relationship will yield more than one candidate mechanism. For example the vulnerability of using weak passwords could be reduced by using a password generator mechanism, by using a token based mechanism, etc. Choosing the candidate mechanisms is a subjective process that will vary from one LAN implementation to another. Not every mechanism presented in Section 2 is feasible for use in every LAN. In order for this process to be beneficial, some filtering of the mechanisms presented needs to be made during this step.

Selecting appropriate safeguards is a subjective process. When considering the cost measure of the mechanism, it is important that the cost of the safeguard be related to the risk measure to determine if the safeguard will be cost-effective. The methodology chosen by the organization should provide a measure for representing costs that is consistent with the measures used for representing the other variables determined so far. Figure 3.6 shows a cost measure that is consistent with the other measuring examples presented. This cost measuring method, while appearing to only consider the cost of the safeguard, can have the other factors mentioned above factored in.

When a measure (or cost) is assigned to the safeguard, it can be compared to the other measures in the process. The safeguard measure can be compared to the risk measure (if it consists of one value, as shown in Figure 3.7) or the components of the risk measure. There are different ways to compare the safeguard measure to the risk measure. The risk management methodology chosen by the organization should provide a method to select those effective safeguards that will reduce the risk to the LAN to an acceptable level.

3.5.2 Process 6 - Implement And Test Safeguards

The implementation and testing of safeguards should be done in a structured manner. The goal of this process is to ensure that the safeguards are implemented correctly, are compatible with other LAN functionalities and safeguards, and provide expected protection.

This process begins by developing a plan to implement the safeguards. This plan should consider factors such as available funding, users' learning curve, etc. A testing schedule for each safeguard should be incorporated into this plan. This schedule should show how each safeguard interacts or effects other safeguards (or mechanisms of some other functionality). The expected results (or the assumption of no conflict) of the interaction should be detailed. It should be recognized that not only is it important that the safeguard perform functionally as expected and provide the expected protections, but that the safeguard does not contribute to the risk of the LAN through a conflict with some other safeguard or functionality.

Each safeguard should first be tested independently of other safeguards to ensure that it provides the expected protection. This may not be relevant to do if the safeguard is designed to interwork with other safeguards. After testing the safeguard independently, the safeguard should be tested with other safeguards to ensure that it does not disrupt the normal functioning of those existing safeguards. The implementation plan should account for all these tests and should reflect any problems or special conditions as a result of the testing.

3.5.3 Process 7 - Accept Residual Risk

After all safeguards are implemented, tested and found acceptable, the results of the risk acceptance test should be reexamined. The risk associated with the threat/vulnerability relationships should now be reduced to an acceptable level or eliminated. If this is not the case, then the decisions made in the previous steps should be reconsidered to determine what the proper protections should be.

Appendix A - LAN Security Policy

A computer security policy is a concise statement of top management's position on information values, protection responsibilities, and organizational commitment. This policy is one of the key components of an overall computer systems security program. It is this policy statement that can drive the initial security requirements for a LAN. However it may be appropriate to address LAN security goals, responsibilities, etc. with a separate policy to be used in conjunction with the existing broader policy. This section discusses establishing a security policy that could be applied to a LAN. It also presents one example of a LAN security policy. This example policy is for example purposes only. It is not intended to be used, as is, by an Agency. The purpose of this example policy is to highlight the issues that should be considered in developing a LAN security policy.

The LAN security policy should be issued by the appropriate level of organizational management, i.e, the person in the organization to whom employees covered by this policy ultimately report. The policy should be created by a team of individuals that may include top management, security officers, and LAN management. The policy should state:

• Information value - Management's position on the value of information;

• Responsibilities - Who is responsible for protecting the information on the LAN;

• Commitment - The organization's commitment to protecting information and the LAN;

• Applicability - What constitutes the LAN environment and what parts, if any, are exempted.

The LAN security policy should be written such that modifications are rarely required. The need for changes may indicate that it is too specific. For example, requiring that a specific virus detection package be used and including the name of the package in the policy may be too specific, considering the rapid pace that virus software packages are developed. It may be more reasonable to merely state that virus detection software should exist on LAN PCs, servers, etc. and let LAN management specify the product.

The LAN security policy should clearly define and establish responsibility for the protection of information that is processed, stored and transmitted on the LAN, and for the LAN itself. Primary responsibility may be with the data owner, i.e., the manager of the organizational component that creates the data, processes it, etc. Secondary responsibility may then be with the users and end users, i.e. those persons within the organization given access to the information by those with primary responsibility. LAN management should clearly define the role of the individuals responsible for maintaining the availability of the LAN. The example LAN security policy below defines responsibilities for functional managers (who may have primary responsibility), users (who may have secondary responsibility), LAN managers (who are responsible for implementing and maintaining LAN security and availability), and local administrators (who are responsible for maintaining security in their part of the LAN environment). Local administrators are usually responsible for one or a subset of the servers and workstations on a LAN. These responsibilities were compiled from [OLDE92], [COMM91], [WACK91], and [X9F292].

An Example LAN Security Policy

Purpose

The information residing on the XYZ Agency local area network (LAN) is mission critical. The size and complexity of the LAN within XYZ has increased and now processes sensitive information. Because of this specific security measures and procedures must be implemented to protect the information being processed on the XYZ LAN. The XYZ LAN facilitates sharing of information and programs by multiple users. This environment increases security risk and requires more stringent protection mechanisms than would be needed for a standalone microcomputer (PC) operation. These expanding security requirements in the XYZ computing environment are recognized by this policy which addresses the use of the XYZ LAN.

This policy statement has two purposes. This first is to emphasize for all XYZ employees the importance of security in the XYZ LAN environment and their role in maintaining that security. The second is to assign specific responsibilities for the provision of data and information security, and for the security of the XYZ LAN itself.

Scope

All automated information assets and services that are utilized by the XYZ Agency Local Area Network (LAN) are covered by this policy. It applies equally to LAN servers, peripheral equipment, workstations, and personal computers (PCs) within the XYZ LAN environment. XYZ LAN resources include data, information, software, hardware, facilities, and telecommunications. The policy is applicable to all those associated with the XYZ LAN, including all XYZ employees, vendors, and contractors utilizing the XYZ LAN.

Goals

The goals of the XYZ information security program are to ensure the integrity, availability and confidentiality of data which are sufficiently complete, accurate, and timely to meet the needs of XYZ without sacrificing the underlying principles described in this policy statement. Specifically the goals are as follows:

- Ensure that the XYZ LAN environment has appropriate security commensurate with sensitivity, criticality, etc.;

- Ensure that security is cost-effective based on a cost versus risk ratio, or that is necessary to meet with applicable mandates;

- Ensure that appropriate support for the security of data in each functional area is provided for;

- Ensure individual accountability for data, information, and other computing resources to which individuals have access;

- Ensure auditibility of the XYZ LAN environment;

- Ensure that employees are provided sufficient guidance for the discharge of responsibilities regarding automated information security;

- Ensure that all critical functions of the XYZ LAN have appropriate contingency plans or disaster recovery plans to provide continuity of operation;

- Ensure that all applicable federal department and organizational policies, mandates, etc. are applied and adhered to.

Responsibilities

The following groups are responsible for implementing and maintaining security goals set forth in this policy. Detailed responsibilities are presented in *Responsibilities for Ensuring XYZ LAN Security.*

1. Functional Management (FM) - those employees who have a program or functional responsibility (not in the area of computer security) within XYZ. *Functional Management is responsible for informing staff about this policy, assuring that each person has a copy, and interacting with each employee on security issues.*

2. LAN Management Division (LM) - employees who are involved with the daily management and operations of the XYZ LAN. They are responsible for ensuring the continued operation of the LAN. *The LAN Management Division is responsible for implementing appropriate LAN*

security measures in order to comply with the XYZ LAN security policy.

3. Local Administrators (LA) - employees who are responsible for ensuring that end users have access to needed LAN resources that reside on their respective servers. *Local administrators are responsible for ensuring that the security of their respective servers is in accordance with the XYZ LAN security policy.*

4. End Users (U) - are any employees who have access to the XYZ LAN. They are responsible for using the LAN in accordance with the LAN security policy. *All users of data are responsible for complying with security policy established by those with the primary responsibility for the security of the data, and for reporting to management any suspected breach of security.*

Enforcement

The failure to comply with this policy may expose XYZ information to the unacceptable risk of the loss of confidentiality, integrity or availability while stored, processed or transmitted on the XYZ LAN. Violations of standards, procedures or guidelines in support of this policy will be brought to the attention of management for action and could result in disciplinary action up to and including termination of employment.

GENERAL POLICIES OF THE LAN

GP1. Every personal computer should have an "owner" or "system manager" who is responsible for the maintenance and security of the computer, and for following all policies and procedures associated with the use of the computer. The primary user of the computer may fill this role. These users should be trained and given guidance so that they can adequately follow all policies and procedures.

GP2. In order to prevent unauthorized access to LAN data, software, and other resources residing on a LAN server, all security mechanisms of the LAN server must be under the exclusive control of the local administrator and the relevant personnel of the LAN Management Division.

GP3. In order to prevent the spread of malicious software and to help enforce program license agreements, users must ensure that their software is properly licensed and safe.

GP4. All software changes and backups on the servers will be the responsibility of the LAN Management Division.

GP5. Each user must be assigned a unique USERID and initial password (or other identification information and authentication data), only after the proper documentation has been completed. Users must not share their assigned USERIDs.

GP6. Users must be authenticated to the LAN before accessing LAN resources.

GP7. USERIDs must be suspended after a consecutive period of non-use.

GP8. Use of LAN hardware such as traffic monitors/recorders and routers must be authorized and monitored by the LAN Management Division.

GP9. The Computer Security Act of 1987 (P.L. 100-235) states that "Each agency shall provide for the mandatory periodic training in computer security awareness and accepted computer practices of all employees who are involved with the management, use, or operation of each Federal computer system within or under the supervision of that agency".

- Employees responsible for the management, operations and use of the XYZ LAN must receive training in computer security awareness and acceptable computer practices.

- Computer security training should be implemented into existing training programs such as orientation programs for new employees, and training courses involved with information technology systems equipment and software packages.

GP10. Security reports must be generated and reviewed on a daily basis.

SPECIFIC RESPONSIBILITIES FOR ENSURING XYZ LAN SECURITY

1. Users

Users are expected to be knowledgeable about and adhere to XYZ Agency security policies, and other applicable laws, policies, mandates and procedures. Users are ultimately responsible for their own behavior. Specifically users are responsible for the following:

U1. Responsible for understanding and respecting relevant Federal laws, Department policies and procedures, XYZ policies and procedures, and other applicable security policies and associated practices for the XYZ LAN.

U2. Responsible for employing available security mechanisms for protecting the confidentiality and integrity of their own information when required.

U2.1. Follow site procedures for security of sensitive data as well as for the XYZ LAN itself. Use file protection mechanisms to maintain appropriate file access control.

U2.2. Select and maintain good passwords. Use FIPS 112, Password

Usage for guidance on good password selection. Do not write passwords down, or disclose them to others. Do not share accounts.

U3. Responsible for advising others who fail to properly employ available security mechanisms. Help to protect the property of other individuals. Notify them of resources (e.g. files, accounts) left unprotected.

U4. Responsible for notifying the local administrator or management if a security violation or failure is observed or detected.

U5. Responsible for not exploiting system weaknesses.

U5.1. Do not intentionally modify, destroy, read or transfer information in an unauthorized manner: do not intentionally deny others authorized access to or use of LAN resources and information.

U5.2. Provide the correct identity and authentication information when requested and not attempt to assume another party's identity.

U6. Responsible for ensuring that backups of the data and software on their own workstation's fixed disk drive are performed.

U7. Responsible for being familiar with how malicious software operates, methods by which it is introduced and spread, and the vulnerabilities that are exploited by malicious software and unauthorized users.

U8. Responsible for knowing and utilizing appropriate policies and procedures for the prevention, detection, and removal of malicious software.

U9. Responsible for knowing how to monitor specific systems and software to detect signs of abnormal activity, and what to do or whom to contact for more information.

U10. Responsible for utilizing the technical controls that have been made available to protect systems from malicious software.

U11. Responsible for knowing and utilizing contingency procedures for containing and recovering from potential incidents.

2. Functional Managers

Functional managers (and higher-level management) are responsible for the development and implementation of effective security policies that reflect specific XYZ LAN objectives. They are ultimately responsible for ensuring that information and communications security is, and remains, a highly visible and critical objective of day-to-day operations. Specifically functional managers are responsible for the following:

FM1. Responsible for implementing effective risk management in order to provide a basis for the formulation of a meaningful policy. Risk management requires identifying the assets to be protected, assessing the vulnerabilities, analyzing risk of exploitation, and implementing cost-effective safeguards.

FM2. Responsible for ensuring that each user receive, at a minimum, a copy of the security policy and site handbook (if any) prior to establishing an account for the user.

FM3. Responsible for implementing a security awareness program for users to ensure knowledge of the site security policy and expected practices.

FM4. Responsible for ensuring that all personnel within the operating unit are made aware of this policy and responsible for incorporating it into computer security briefings and training programs.

FM4. Responsible for informing the local administrator and the LAN Management Division of the change in status of any employee who utilizes the XYZ LAN. This status change includes an interagency position change, interdivision position change, or a termination from XYZ employment.

FM5. Responsible for ensuring that users understand the nature of malicious software, how it is generally spread, and the technical controls to use for protection.

3. Local Area Network (LAN) Management Division

The LAN Management Division (or designated personnel) is expected to enforce (to the extent possible) local security policies as they relate to technical controls in hardware and software, to archive critical programs and data, and to control access and protect LAN physical facilities. Specifically, LAN management is responsible for the following:

NM1. Responsible for rigorously applying available security mechanisms for enforcement of local security policies.

NM2. Responsible for advising management on the workability of the existing policies and any technical considerations that might lead to improved practices.

NM3. Responsible for securing the LAN environment within the site and interfaces to outside networks.

NM4. Responsible for responding to emergency events in a timely and effective manner.

> NM4.1. Notify local administrators if a penetration is in progress, assist other local administrators in responding to security violations.
>
> NM4.2. Cooperate with local administrators in locating violators and assist in enforcement efforts.

NM5. Responsible for employing generally approved and available auditing tools to aid in the detection of security violations.

NM6. Responsible for conducting timely audits of LAN server logs.

NM7. Responsible for remaining informed on outside policies and recommended practices and when appropriate, informing local users and advising management of changes or new developments.

NM8. Responsible for judiciously exercising the *extraordinary* powers and privileges that are inherent in their duties. Privacy of users should always be a major consideration.

NM9. Responsible for developing appropriate procedures and issuing instructions for the prevention, detection, and removal of malicious software consistent with the guidelines contained herein.

NM10. Responsible for backing up all data and software on the LAN servers on a timely basis.

NM11. Responsible for identifying and recommending software packages for the detection and removal of malicious software.

NM12. Responsible for developing procedures that allow users to report computer viruses and other incidents and then responsible for notifying potentially affected parties of the possible threat.

NM13. Responsible for promptly notifying the appropriate security or incident response personnel of all computer security incidents including malicious software.

NM14. Responsible for providing assistance in determining the source of malicious software and the extent of contamination.

NM15. Responsible for providing assistance for the removal of malicious software.

NM16. Responsible for conducting periodic reviews to ensure that proper security procedures are followed, including those designed to protect against malicious software.

4. Local Administrators

Local administrators (or designated personnel) are expected to utilize, on their assigned server, the available LAN security services and mechanisms to support and enforce applicable security policies and procedures. Specifically local administrators are responsible for the following:

LA1. Responsible for managing all users' access privileges to data, programs and functions.

LA2. Responsible for monitoring all security-related events and the following-up on any actual or suspected violations where appropriate. When appropriate, responsible for notifying and coordinating with the LAN Management Division the monitoring or investigation of security-relevant events.

LA3. Responsible for maintaining and protecting LAN server software and relevant files using available security mechanisms and procedures.

LA4. Responsible for scanning the LAN server with anti-virus software at regular intervals to assure no virus becomes resident on the LAN server.

LA5. Responsible for assigning a unique USERID and initial password (or other identification information or authentication data) to each user only after proper documentation has been completed.

LA6. Responsible for promptly notifying the appropriate security or incident response personnel of all computer security incidents, including malicious software;

> LA6.1. Notify the LAN Management Division if a penetration is in progress, assist other local administrators in responding to security violations.
>
> LA6.2. Cooperate with other local administrators and the LAN Management Division in finding violators and assisting in enforcement efforts.

LA7. Responsible for providing assistance in determining the source of malicious software and the extent of contamination.

Appendix B - Personal Computer Considerations

Personal computers typically do not provide technical controls for user authentication, access control, or memory protection that differentiates between system memory and memory used for user applications. Because the lack of controls and the resultant freedom with which users can share and modify software, personal computers are more prone to attack by viruses, unauthorized users and related threats.

Virus prevention in the PC environment must rely on continual user awareness to adequately detect potential threats and then to contain and recover from the damage. Personal computer users are in essence personal computer managers, and must practice their management as a part of their general computing. Personal computers generally do not contain auditing features, thus a user needs to be aware at all times of the computer's performance, i.e., what is normal or abnormal activity. Ultimately, personal computer users need to understand some of the technical aspects of their computers in order to detect security problems, and to recover from those problems. Not all personal computer users are technically oriented, thus this poses some problems and places even more emphasis on user education and involvement in virus prevention.

Because of the dependence on user involvement, policies for LAN environments (and thus PC usage) are more difficult to implement than in a multi-user computer environment. However, emphasizing these policies as part of a user education program will help to ingrain them in users' behavior. Users should be shown via illustrated example what can happen if they do not follow the policies. An example where users share infected software and them spread the software throughout an organization would serve to effectively illustrate the point, thus making the purpose of the policy more clear and more likely to be followed. (It is not suggested that an organization actually enact this example, merely illustrate it). Another effective method for increasing user cooperation is to create a list of effective personal computer management practices specific to each personal computing environment. Creating such a list would save users the problem of determining how best to enact the policies, and would serve as a convenient checklist that users could reference as necessary.

For guidance on general protection of PCs see [STIE85]. For guidance on protecting against malicious software see [WACK89].

Appendix C - Contingency Planning for LANs

A computer security incident is any adverse event whereby some aspect of computer security could be threatened: loss of data confidentiality, loss of data or system integrity, or disruption or denial of availability. In a LAN environment the concept of a computer security incident can be extended to all areas of the LAN (hardware, software, data, transmissions, etc.) including the LAN itself. Contingency plans in a LAN environment should be developed so that any LAN security incident can be handled in a timely manner, with as minimal an impact as possible on the ability of the organization to process and transmit data. A contingency plan should consider (1) incident response, (2) back-up operations, and (3) recovery.

1. The purpose of incident response is to mitigate the potentially serious effects of a severe LAN security-related problem. It requires not only the capability to react to incidents, but the resources to alert and inform the users if necessary. It requires the cooperation of all users to ensure that incidents are reported and resolved and that future incidents are prevented [WACK91,5]. [WACK91] is recommended as guidance in developing an incident response capability.

2. Back-up Operations plans are prepared to ensure that essential tasks (as identified by a risk analysis) can be completed subsequent to disruption of the LAN environment and continuing until the LAN is sufficiently restored [NIST74,65].

3. Recovery plans are made to permit smooth, rapid restoration of the LAN environment following interruption of LAN usage [NIST74,65]. Supporting documents should be developed and maintained that will minimize the time required for recovery. Priority should be given to those applications, services, etc. that are deemed critical to the functioning of the organization. Back-up operation procedures should ensure that these critical services and applications are available to users.

Appendix D - Training and Awareness

The Computer Security Act of 1987 (P.L. 100-235) states that "Each agency shall provide for the mandatory periodic training in computer security awareness and accepted computer practices of all employees who are involved with the management, use, or operation of each Federal computer system within or under the supervision of that agency."

[TODD89] provides a framework for identifying computer security training requirements for a diversity of audiences who should receive some form of computer security training. It focuses on learning objectives based upon the extent to which computer security knowledge is required by an individual as it applies to his or her job function. For detailed discussion and guidance for general computer security training the reader is directed to [TODD89].

To maintain security in a LAN environment, training in certain areas of LAN operation and use should be received by LAN users. Security mechanisms, procedures, etc. may not be effective if they are used improperly. Training areas that should be considered are listed below for functional managers, LAN managers and general users. The training area for functional managers focuses on (1) the need to understand the importance of the security policy and (2) how that policy needs to be implemented into the LAN for it to be effective. The training area for LAN managers focuses on the need to understand how security is provided for operationally on the LAN. It also directs attention on the need for effective incident response. The training area for all users focuses on (1) recognizing the user role in the security policy and the responsibilities assigned there, (2) using the security services and mechanisms effectively to maintain security, and (3) understanding how to use the incident response procedures. Specifically these areas are discussed below.

Functional Managers

1. Recognize the importance of the LAN security policy and how this policy drives the decisions made regarding LAN security. Recognize the importance of determining adequate security for different types of information that the functional manager owns (or has responsibility for).

2. Recognize the LAN as a valuable resource to the organization and the need for protecting that resource. Recognize the importance of providing for adequate protection (through funding, personnel, etc.).

LAN Management

1. Understand how the LAN operates in all aspects. Ability to recognize normal operating behavior versus abnormal operating behavior.

2. Understand LAN management's role in implementing the security policy into the LAN.

3. Understand how the security services and mechanisms work. Ability to recognize improper use of the security mechanisms by users.

4. Understand how to use the incident response capability effectively.

LAN Users

1. Understand the security policy and the user responsibilities dictated there. Understand why maintaining LAN security is important.

2. Understand how to use the security services and mechanisms provided by the LAN to maintain the security of the LAN and protect critical information.

3. Understand how to use the incident response capability, how to report and incident, etc.

4. Recognize normal workstation or PC behavior versus abnormal behavior.

References

[MART89] Martin, James, and K. K. Chapman, The Arben Group, Inc.; Local Area Networks, Architectures and Implementations, Prentice Hall, 1989.

[BARK89] Barkley, John F., and K. Olsen; Introduction to Heterogenous Computing Environments, NIST Special Publication 500-176, November, 1989.

[NCSC87] A Guide to Understanding Discretionary Access Control in Trusted Systems, NCSC-TG-003, Version 1, September 30, 1987

[NCSL90] National Computer Systems Laboratory (NCSL) Bulletin, Data Encryption Standard, June, 1990.

[SMID88] Smid, Miles, E. Barker, D. Balenson, and M. Haykin; Message Authentication Code (MAC) Validation System: Requirements and Procedures, NIST Special Publication 500-156, May, 1988.

[OLDE92] Oldehoeft, Arthur E.; Foundations of a Security Policy for Use of the National Research and Educational Network, NIST Interagency Report, NISTIR 4734, February 1992.

[COMM91] U.S. Department of Commerce Information Technology Management Handbook, Attachment 13-D: Malicious Software Policy and Guidelines, November 8, 1991.

[WACK89] Wack, John P., and L. Carnahan; Computer Viruses and Related Threats: A Management Guide, NIST Special Publication 500-166, August 1989.

[X9F292] Information Security Guideline for Financial Institutions, X9/TG-5, Accredited Committee X9F2, March 1992.

[BJUL93] National Computer Systems Laboratory (NCSL) Bulletin, Connecting to the Internet: Security Considerations, July 1993.

[BNOV91] National Computer Systems Laboratory (NCSL) Bulletin, Advanced Authentication Technology, November 1991.

[KLEIN] Daniel V. Klein, "Foiling the Cracker: A Survey of, and Improvements to, Password Security", Software Engineering Institute. (This work was sponsored in part by the Department of Defense.)

[GILB89] Gilbert, Irene; Guide for Selecting Automated Risk Analysis Tools, NIST Special Publication 500-174, October, 1989.

[KATZ92] Katzke, Stuart W. ,Phd., "A Framework for Computer Security Risk Management", NIST, October, 1992.

[NCSC85] Department of Defense Password Management Guideline, National Computer Security Center, April, 1985.

[NIST85] Federal Information Processing Standard (FIPS PUB) 112, Password Usage, May, 1985.

[ROBA91] Roback Edward, NIST Coordinator, Glossary of Computer Security Terminology, NISTIR 4659, September, 1991.

[TODD89] Todd, Mary Anne and Constance Guitian, Computer Security Training Guidelines, NIST Special Publication 500-172, November, 1989.

[STIE85] Steinauer, Dennis D.; Security of Personal Computer Systems: A Management Guide, NBS Special Publication 500-120, January, 1985.

[WACK91] Wack, John P.; Establishing a Computer Security Incident Response Capability (CSIRC), NIST Special Publication 800-3, November, 1991.

[NIST74] Federal Information Processing Standard (FIPS PUB) 31, Guidelines for Automatic Data Processing Physical Security and Risk Management, June, 1974.

Further Reading

[1] Berson, T.A, and Beth, T. (Eds.); Local Area Network Security Workshop LANSEC '89 Proceedings, Springer-Verlag, Berlin, 1989.

[2] Federal Information Processing Standard Publication (FIPS PUB) 83, Guideline on User Authentication Techniques for Computer Network Access Control, September, 1980.

[3] Gahan, Chris; LAN Security, the Business Threat from Within, BICC Data Networks Limited, November, 1990.

[4] Muftic, Sead; Security Mechanisms for Computer Networks, Ellis Horwood Limited, West Sussex, England, 1989.

[5] National Research Council; Computers At Risk: Safe Computing in the Information Age, National Academy Press, Washington, D.C., 1991.

[6] Schweitzer, James A.; Protecting Information on Local Area Networks, Butterworth Publishers, Stoneham, MA, 1988.