CONTENTS

1 - Introduction

1.1 Information Security:

The information and the supporting processes, the computer systems and the networks, used for generating, storing and retrieving information and the human beings are important business assets of every Bank. The confidentiality, integrity and availability of information is essential for any financial Bank to maintain its competitive edge, cash-flow, profitability, legal compliance and commercial image. The application of Information Technology has brought about significant changes in the way the banking and the financial Banks process and store data. The telecommunication networks have played a catalytic role in the expansion and integration of the Information Systems, within and among the Banks, facilitating data accessibility to different users. This has made it imperative for each Bank to put in place adequate security controls to ensure data accessibility to all the authorized users, data inaccessibility to all the unauthorized users, maintenance of data integrity and implementation of safeguards against all security threats to guarantee information and information systems security across the Bank. This makes it necessary for each Bank to define, document, communicate, implement and audit Information Systems (IS) Security.

Many information systems in operation in our Bank may not have been designed to be sufficiently secure. Further, the level of security, which can be achieved through the application of technology, could also be limited and therefore, it requires to be supported by appropriate management policies and procedures. The selection of the security controls requires careful and detailed planning. The management of information and information systems security will require participation by all the employees in the Bank. It will also require participation from the third parties such as the suppliers, vendors, customers and shareholders. The Bank may also have to turn to specialist advice from Banks like IDRBT in the matter of information systems security, as and when required. The information systems security could be achieved by implementing a suitable set of controls, which consists of policies, practices, procedures, hardware and software functions. Each Bank has to establish these controls to ensure that its security requirements are met.

The Board of Directors has the responsibility for ensuring appropriate corporate policies, which set out the management responsibilities and the control practices for all the areas of information processing activities. A well-defined corporate Information Systems security policy which is being put in place and periodically reviewed and amended, as required, under the approval of the Board of Directors.

The management of risks is central to the Bank in the banking and financial sector. These Banks manage risks through prudent business practices, contractual arrangements with third parties, obtaining insurance coverage and use of appropriate security mechanisms. These Banks have now been increasingly dependent on the Information Technology (IT) for the efficient conduct of business, which necessitates growing levels of information systems security within the Banks. This document contains the guidelines / procedures for building up an "Information Systems Security environment" in the Bank.

1.2 What Are Information Security Policies

"Policies" are management instructions indicating a course of action, a guiding principle, or an appropriate procedure, which is expedient, prudent, or advantageous. Policies are high-level statements that provide guidance to workers who must make present and future decisions. It would also be correct to say that policies are generalized requirements that must be written down and communicated to certain groups of people inside, and in some cases outside, the organization. Policies can also be considered to be business rules. Although information security policies vary considerably by organization, they typically include general statements of goals, objectives, beliefs, ethics, controls, and worker responsibilities.

Policies are mandatory and can also be thought of as the equivalent of an organization-specific law. Special approval is required when a worker wishes to take a course of action that is not in compliance with policy. Because compliance is required, policies use definitive words like "do not..." "you must ...," or "you are obliged to ...". The words used to indicate policies must convey both certainty and indispensability. For simplicity and consistency, throughout this book, the word "must" has been employed, but equivalent words are certainly acceptable.

1.3 Difference Between Policy and Procedures & Guidelines:

Policies are distinct from and considerably higher-level than "procedures" (sometimes called "standard operating procedures"). A policy statement describes only the general means for addressing a specific problem; Procedures are specific operational steps or manual methods that workers must employ to achieve a certain goal. For instance, in many information technology departments there are specific procedures for performing back-ups of server hard drives. In this example, a policy could describe the need for back-ups, for storage off-site, and for safeguarding the back-up media (using encryption, physical security, etc.). A standard could define the software to be used to perform back-ups and how to configure this software. A procedure could describe how to use the back-up software, the timing for making back-ups, and other ways that humans interact with the back-up system (how to restore a file, how to securely transport storage media to an off-site repository, etc.).

Policies are higher-level requirement statements than "standards," although both types of management instructions require compliance. Policies provide general instructions, while standards provide specific technical requirements. Standards cover details such as implementation steps; systems design concepts, software interface mechanisms, software algorithms, and other specifics. The phrase "information security architecture" is gaining increasing acceptance as a collection of integrated information security standards. Standards would for example define the number of secret key bits required in an encryption algorithm such as SSL (Secure Sockets Layer, a widespread Internet encryption protocol). Policies, on the other hand, would simply define the need to use an approved encryption process when sensitive information is sent over public networks such as the Internet.

Separately, policies are generally aimed at a wider audience than standards. For example, a policy requiring the use of computer virus packages would apply to all personal computer users, but a standard requiring the use of public key digital certificates could be directed only at staff who conduct organizational business over the Internet.

2. Need for Information Security Policy and Guidelines

The business operations in the banking and the financial sector are becoming increasingly dependent on computerized information systems. It has now become impossible to separate technology from the business of the banks/financial Banks. The growing use of the personal computers and their networking in the financial sector has necessitated their integration in a Local Area or Wide Area Network environment. On account of the phenomenal growth in the use of IT and IT based applications by the Banks in its day-to-day operations, the need for putting in place the security controls for all the information systems has grown tremendously. The information systems security has, therefore, assumed great importance for the commercial success of the Banks, as the survival of the Banks depends on the speed, accuracy and reliability of the flow of information within the Banks vis-à-vis its customers.

The security controls are required to minimize the vulnerability to unauthorized use of the information and the information systems. However, such controls may have to be consistent with the degree of exposure of such systems and the information and the impact of loss to the Bank on account of unauthorized access and misuse, including accidental misuse, of such systems and information. The unauthorized including accidental misuse of the information may result in financial loss, competitive disadvantage, damaged reputation, improper disclosure, law suits and non-compliance with the regulatory provisions etc. Structured, well-defined and documented security policies, standards and guidelines lay the foundation for good information systems security.

No threat becomes obsolete. Further, new threats surface from time to time. The financial sector has witnessed rapid changes in the volume and the value of transactions and the introduction of the most modern and secured methods for the delivery of services to the customers. Still better information systems are being introduced at frequent intervals. Further, the banking and the financial sector is now poised to countenance various developments such as Internet banking, e-money, e-cheque, e-commerce etc., which have been made possible by the revolutionary researches and discoveries in Information Technology and its applications and the future promises to remain challenging. Constant developments of far reaching implications dictate constant vigilance and necessitate sound information systems security infrastructure.

Security policies are the foundation of your secure infrastructure. Your security policies serve as a guide and a reference point to numerous security tasks in your organization including:

- Securing applications
- Configuring user access controls
- Defining management duties and responsibilities
- Assuring standardization and consistency
- Retaining confidential and proprietary information
- Designing enterprise architecture
- Mitigating risk
- Responding to security incident investigations
- Disciplining employees for breach of policy
- Minimizing liabilities to customers and shareholders
- Assisting auditors in understanding security intentions
- Establishing a sense of awareness and training
- Avoiding disputes with different technical teams
- Expediting procurement and deployment of new systems

Without security policies, no enforcement of security configurations or standards can be made. By establishing a policy, you are implying that enforcement can or will follow. Without security policies, enforcement of them is not possible.

2.1 Implementation of Information Security Policy and Procedures :

At the Corporate level, the Chief Information Systems Security Officer (CISO) would be responsible for Information Systems Security. He will be assisted by a team of Officers comprising both Technical and Banking Officers to be responsible for the information systems security policies implementation in each of the offices/locations of the Bank. Information Systems Security Department in the Bank will address various issues such as the development of the Information Systems Security Policy, updation of the Information Systems Security Guidelines on an on-going basis, provision of consultancy and information on information systems security requirements, maintenance of centralised security functions etc. Further, the System Administration responsibilities should, among others, relate to the implementation of the security controls, compliance with the information systems security guidelines, management of day-today security functions etc. Information Systems Security department would be responsible for

- Identification of individuals to be responsible for the protection of information assets at each office/location of the Bank or as warranted
- Classification of information assets and specifications of the appropriate levels of security for each class of information assets
- Implementation of an awareness/education programme to ensure that the employees and the related third parties are aware of and observe their respective responsibilities for the maintenance and continuation of information systems security in the Bank
- Reporting of information systems security incidents and provision for their resolution
- Preparation of written (comprehensively documented) plans and procedures for business resumption/continuity following disasters
- Identification of the procedures and the processes for addressing exceptions or deviations from the information systems security policy document
- Co-ordination and co-operation among the various disciplines in the Bank such as technical, operational, audit, insurance and regulatory compliance
- Laying down precisely the responsibilities to ensure compliance with and to assess soundness and comprehensiveness of the information systems security policies on a continuous basis
- Review, updation and upgradation of the information systems security policies in the light of new threats and technology on a continuous basis and
- Preparation of the audit records, where necessary and the monitoring of the audit trails for the detection of uncharacteristic behavior of individuals and activities.

The Information Systems Security Managers/ Officers will serve as the supervisors and have to, therefore, monitor the successful implementation of the Information Systems Security Policy within their work-areas. They should be adequately trained on Information System Security standards like BS 7799 / ISO 17799 and should be encouraged to pursue courses on Information security/ Audit such as CISSP/ CISM/ CISA / CFE/ and other Internationally acclaimed certifications. This makes them key players in the implementation of information systems security policies.

2.2 Information Systems Security Administration :

Each business unit and information systems manager should lay down the need-to-know access privileges for the users within his business domain and communicate the same to the users. These access privileges should be documented. Further, these documented privileges should be reviewed periodically and changes should be made as and when deemed appropriate.

Each information access control system should have one or more Information Systems Security Administrator(s), appointed to ensure that the access control procedures are being monitored and enforced continuously. The Information Systems Security Administrator should:

a) Be responsible for maintaining accurate and complete access control privileges, based on the instructions from the information resource owner and in accordance with any applicable internal policies, directives and standards, laid down therefor;

b) Remain informed by the appropriate manager/s whenever the service of the employees is terminated or they are transferred or retire or are on leave or when have joint responsibilities, if any;

c) Monitor selected users with high-level access privileges and remove such privileges immediately, when such privileges are no longer required;

d) Monitor daily access activity to determine if any unusual activity has taken place such as repeated invalid access attempts that may threaten the integrity, confidentiality or the availability of information and the information systems. These unusual activities, whether intentional or accidental, must be brought to the attention of the information resource owner for investigation and resolution;

e) Ensure that each information system user be identified by a unique identification sequence (USERID), associated only with that user. The process should require that the user identity be authenticated prior to the user's gaining access to the information system by utilizing an established/properly chosen authentication technology;

f) Make periodic reviews on access to information systems by the users and report to the appropriate information resource owner; and

g) Ensure that the audit trail is collected, protected and available, whenever required.

The activities of the Information Systems Security Administrator/s (ISSA) have to be reviewed by an independent party such as Audit department, for the purpose, on a routine basis

IS Security Responsibilities

Information Owner

Data and records stored on systems are presumed to be the property of the Owner. For purposes of Policy, the "owner" of a collection of information will be responsible for the creation of that information or the primary user of that information. In this context, ownership does not signify proprietary interest, and ownership may be shared. The owner may delegate ownership responsibilities to another individual. Owners have the responsibility to:

- ✓ Know the information for which they are responsible.
- ✓ Determine a data retention period for their information, in accordance with the departmental records retention schedule (and any applicable state or federal laws or regulations).
- ✓ Ensure appropriate procedures are in effect to protect the integrity, confidentiality, and security of the information used or created within their area.
- ✓ Authorize access and assign custodianship.
- ✓ Specify controls and communicate the control requirements to the custodian and users of the information.
- ✓ Report promptly to the Information Security Officer the loss or misuse of information.
- ✓ Initiate appropriate actions when problems are identified.
- ✓ Promote education and awareness by utilizing programs administered by the Information Security Officer, where appropriate.

- ✓ Follow existing approval processes within their respective organizations for the selection, budgeting, purchase, and implementation of any computer system/software to manage information.

Custodian

The custodian of information is generally responsible for the processing and storage of the information. The custodian is responsible for the administration of controls as specified by the owner. Responsibilities may include:

- ✓ Providing and/or recommending physical safeguards.
- ✓ Providing and/or recommending procedural safeguards.
- ✓ Administering access to information.
- ✓ Evaluating the cost effectiveness of controls.
- ✓ Coordinating the maintenance of information security policies, procedures and standards as appropriate and in consultation with the Information Security Officer.
- ✓ Promoting education and awareness by utilizing programs administered by the Information Security Officer, where appropriate.
- ✓ Report promptly to the Information Security Officer the loss or misuse of information.
- ✓ Initiate appropriate actions when problems are identified.

User

The user is any person who has been authorized to read, enter, or update information. A user of information is expected to:

- ✓ Access information only in support of authorized job responsibilities or role.
- ✓ Comply with Information Security Policies and Standards and with all controls established by the owner and custodian.
- ✓ Refer all disclosures of confidential information to the Privacy Officer.
- ✓ Keep authentication devices (e.g. passwords, SecureCards, PINs, etc.) Confidential.
- ✓ Report promptly to the Information Security Officer the loss or misuse of any authentication device.
- ✓ Report promptly to the Information Security Officer the loss or misuse of information.
- ✓ Initiate appropriate actions when problems are identified.

User Management

Staff who supervise users as defined above or who handle unit administrative responsibilities or as designated by head unit. User management is responsible for overseeing their user's access to information, including:

- ✓ Reviewing and approving all requests for access authorizations.
- ✓ Initiating security change requests to keep security record current so they accurately reflect the users role and required access.
- ✓ Promptly informing the appropriate Information Security Officer of employee terminations and transfers.
- ✓ Revoking physical access to terminated employees, i.e., confiscate keys, change combination locks, etc.
- ✓ Revoking physical access to students and others when access to information is no longer needed or appropriate.
- ✓ Providing the opportunity for training needed to properly use the computer systems.
- ✓ Reporting promptly to the Information Security Officer the loss or misuse of information.
- ✓ Initiating appropriate actions when problems are identified.
- ✓ Follow existing approval processes within their respective organization for the selection, budgeting, purchase, and implementation of any computer system/software to manage information.

Information Security Officer

The Information Security Officer is responsible for working with user management, owners, custodians, and users to develop and implement prudent security policies, procedures, and controls, subject to the approval of Counsel. Large departments with significant Confidential Information may have a departmental Information Security Liaison. Specific responsibilities include:

- ✓ Ensuring security policies, procedures, and standards are in place and adhered to.
- ✓ Providing basic security support for all systems and users.
- ✓ Advising owners in the identification and classification of computer resources.
- ✓ Advising systems development and application owners in the implementation of security controls for information on systems, from the point of system design, through testing and production implementation.

- ✓ Educating custodian and user management with comprehensive information about security controls affecting system users and application systems.
- ✓ Providing on-going security education.
- ✓ Performing security audits for compliance.
- ✓ Reporting regularly to the Security Officer and Privacy Officer on entity’s status with regard to information security.

3. Information Security Policy Standards

3.1 Introduction

What is an IT Security Policy?

An IT Security Policy is the most critical element of an IT security program. A security policy identifies the rules and procedures that all persons accessing computer resources must adhere to in order to ensure the confidentiality, integrity, and availability of data and resources. Furthermore, it puts into writing an organization's security posture, describes and assigns functions and responsibilities, grants authority to security professionals, and identifies the incident response processes and procedures.

Security Policy Basics

Security policies are high-level laws of the land regarding your security infrastructure. They are not procedures. (Procedures tell you how to implement security policies.) Upper management needs to hold someone accountable for drafting the security policies, overseeing their review, and implementing them. Without support from upper management, security policies often fall by the way side and never get written, understood, or implemented. The person being held responsible for security policies could be the Director of Information Security, the Chief Security Officer, the Director of Information Technology, the Chief Information Officer, or a knowledgeable employee appointed to be the information security officer.

Security is typically distributed, and security mechanisms should be built into all layers of the enterprise infrastructure. Security policies should describe the rules of the road for the following types of technology systems:

- Encryption mechanisms
- Access control devices
- Authentication systems
- Virtual Private Networks (VPNs)
- Firewalls
- Messaging systems
- Anti-virus systems
- Web sites
- Gateways

- Mission critical applications
- End-user desktops
- DNS servers
- Routers and switches

All security policies need to be written down. Policies that exist in someone's head are not really policies. When your organization has finished developing security policies, and right when you think you can breathe easy, it will be time to update your security policies. Since most IT organizations are deploying new technology continuously and retiring old systems, you will have to make sure your security policies still make sense for your new infrastructure. Similarly, when you are evaluating new equipment for possible procurement, you will want to make sure that the new equipment can properly be configured to meet your security requirements — if it can't, you may want to consider procuring alternative products. Security controls are mechanisms put into place to enforce security policies.

What Determines a Good IT Security Policy?

In general a good IT Security Policy does the following:

- Communicates clear and concise information and is realistic;
- Includes defined scope and applicability;
- Makes enforceability possible;
- Identifies the areas of responsibility for users, administrators, and management;
- Provides sufficient guidance for development of specific procedures
- Balances protection with productivity;
- Identifies how incidents will be handled; and
- Is enacted by a senior official (e.g., CEO)

Development of a security policy should be a collaborative effort with security officials, management, and those who have a thorough understanding of the business rules of the organization. A security policy should not impede an organization from meeting its mission and goals. However, a good policy will provide the organization with the assurance and the “acceptable” level of asset protection from external and internal threats.

3.2 Scope

What are the Components of a Security Policy?

A key point to consider is to develop a security policy that is flexible and adaptable as technology changes. Additionally, a security policy should be a living document routinely updated as new technology and procedures are established to support the mission of the organization.

The components of a security policy will change by organization based on size, services offered, technology, and available revenue. Here are some of the typical elements included in a security policy.

Security Definition – All security policies should include a well-defined security vision for the organization. The security vision should be clear and concise and convey to the readers the intent of the policy. In example:

“This security policy is intended to ensure the confidently, integrity, and availability of data and resources through the use of effective and established IT security processes and procedures.”

Further, the definition section should address why the security policy is being implemented and what the corresponding mission will entail. This is where you tie the policy to the mission and the business rules of the organization.

Enforcement – This section should clearly identify how the policy will be enforced and how security breaches and/or misconduct will be handled.

The Chief Information Officer (CIO) and the Information Systems Security Officer (ISSO) typically have the primary responsibility for implementing the policy and ensuring compliance. However, you should have a member of senior management, preferably the top official, implement and embrace the policy. This gives you the enforcement clout and much needed ‘buy-in’.

This section may also include procedures for requesting short-term exceptions to the policy. All exceptions to the policy should be reviewed and approved, or denied, by the Security Officer. Senior management should not be given the flexibility to overrule decisions. Otherwise, your security program will be full of exceptions that will lend themselves toward failure.

Security Profiles - A good security policy should also include information that identifies how security profiles will be applied uniformly across common devices (e.g., servers, workstations, routers, switches, firewalls, proxy servers, etc.).   The policy should reference applicable standards and procedures for locking down devices. Those standards may include security checklists to follow when adding and/or reconfiguring devices.

New devices come shipped with the default configuration for ease of deployment and it also ensures compatibility with most architectures. This is very convenient for the vendor, but a nightmare for security professionals. An assessment needs to be completed to determine what services are necessary on which devices to meet the organizational needs and requirements. All other services should be turned off and/or removed and documented in the corresponding standard operating procedure.

For example, if your agency does not have a need to host Internet or Intranet based applications then do not install Microsoft IIS. If you have a need to host HTML services, but do not have a requirement for allowing FTP, then disable it.

Auditing - All security programs should be audited on a routine and random basis to assess their effectiveness. The security officer must be given the authority, in writing, by the head of the organization to conduct audits of the program. If not, he or she could be subject to legal action for malicious conduct. Random and scheduled audits should be conducted and may include:

- Password auditing using password cracking utilities such as LC3 (Windows) and PWDump (Unix and Windows);
- Auditing user accounts database for active old accounts (persons who left the agency)
- Penetration testing to check for vulnerabilities using technical assessment tools such as ISS and Nessus;
- Social Engineering techniques to determine if you can get a username or password from a staff member;
- Simulate (off hours) network failure and evaluate your incident response team's performance and readiness;
- Test your back-up recovery procedures;
- Use Tripwire or similar product to monitor your critical binary files;
- Configure your Server OS to audit all events and monitor several times a day for suspicious activity;
- Use a port scanner (Nmap, Nessus, etc.) within your network to determine if your system administrators catch the traffic and take appropriate action.

These are just a few examples of the things to audit. The extent of your auditing will depend on the level of your security program.

Awareness Training - Security Awareness training for organizational staff must be performed to ensure a successful program. Training should be provided at different levels for staff, executives, system administrators, and security officers. Additionally, staff should be retrained on a periodic basis (e.g., every two years). A process should be in place for training newly hired staff within a certain time period. Staff completing training should be required to sign a written certification statement. This signed statement helps the security officer and management enforce the organization's security policies.

Trained staff can help alleviate some of the security burden from security officers. Trained staff can and often do provide advanced notification of suspicious events encountered on their machines which could prevent a worm or other Trojan from propagating throughout the entire network.

Administrative Policies vs. Technical Policies

Technical security policies describe how technology should be configured and used, and administrative security policies describe how people (end-users and management) should behave. The intended security rules for technology systems and data should be explicitly described in technical security policies. Technical security policies describe a rule or regulation pertaining to a piece of equipment, facility, or data.

Administrative security policies describe the intended behavior rules for people. Serving as a guide for both end-users and management, administrative policies should spell out the roles and responsibilities for all users of technology systems in the organization. It is very important to inform end-users and other management team members of administrative security policies. Users cannot be expected to follow policies if they do not know what they are. After reviewing the administrative policies, it is a good idea to get the user to sign the policy document attesting to the fact that they have read it, understand it and will abide by it.

Many organizations take the time to define technical security policies, while administrative security policies are often overlooked. While many technical security policies can be audited with online scanning tools, administrative security policies can only be audited with an in-person review. Auditors who review administrative policies will typically ask to see the actual formal policy document. Efficient auditors will also interview end-users and management to see if they understand their roles and responsibilities.

Administrative Security Policy

- Users must change their passwords each quarter.
- A designated employee should securely maintain a master list of passwords.
- End-users will update their virus signatures at least once a week.
- Users will not e-mail passwords across the corporate infrastructure.

- Users must use a card-key to enter the data center.
- Employees shall not use dial-out modems from their desktops.
- All users must read, sign and abide by the end-user security policies.
- Users will report suspicious network activity to the security officer.
- The security officer will manage and respond to all security incidents.
- Company information marked *Proprietary* must be used only for legitimate business purposes.

If your organization was being audited, here are some questions that an auditor might ask in regards to your administrative security policies:

1. Are employees informed about reporting security incidents? How would they know what to report, and to whom to report it? Where can employees turn for information to guide them on how to handle security incidents?
2. Are there security policies associated with change-management?
3. Who is responsible for risks associated with third-party vendors and partners?
4. Are there regular security reviews of IT systems? Are reports generated to capture the current security posture and make recommendations for corrective action? (You should assume that if an auditor asks if reports exist they will ask to see them.)
5. Is there a policy that defines acceptable use of the Internet?
6. What authorization is needed to change user IDs?
7. Are procedures for the disposal of media documented?
8. Who is responsible for enforcing policy breaches?
9. What are the reasons for allowing employees remote access?
10. How are employees made aware of security policies and procedures?

Technical Security Policy

- Servers will be configured to expire passwords once every quarter.
- Anti-virus software will be installed and properly configured on all user desktops.
- A card-key system will be installed at every data center entrance.
- All data and information must be assigned a custodial owner.
- All access control systems must be monitored for compliance.
- Online penetration tests should be conducted twice a year.

- Accounts must initiate a lock out after four unsuccessful attempts to login.
- Encryption is used to prevent access to sensitive and proprietary information.
- Incoming attachments must be scanned for viruses on the boundary server.
- Passwords must be at least 8 characters long and include upper and lower case characters and at least one numeric character.

If your organization was being audited, here are some questions that an auditor might ask in regards to your technical security policies:

1. Are stored passwords on the Web site encrypted? How?
2. How is the logical access to the Web site server controlled?
3. What controls are in place to protect audit log files?
4. Is there a master backup of router and firewall configuration files?
5. What outbound and inbound connections and services are being allowed through the firewall?
6. What is the process for authenticating firewall administrators?
7. Are Web servers protected from buffer overflow attacks? How?
8. What security controls exist to protect credit cards numbers? Are the credit card number encrypted?
9. How is the security of dial-up connections controlled?
10. Does the enterprise system architecture documentation include all physical and logical (VLAN) connections?

3.3 Standards for Information Control

Standards for Information Control

All involved systems and information are assets must be protected from misuse, unauthorized manipulation, and destruction. These protection measures may be physical and/or software based.

Unauthorized Use of Software: Computer software owned or licensed must not be copied for use at home or any other location, unless authorized by the Information Security Officer or as otherwise specified by the license agreement.

Installed Software: All software packages that reside on computers and networks must comply with applicable licensing agreements and restrictions and must comply with acquisition of software policies.

Virus Protection: Virus checking systems approved by the Information Security Officer and information services must be deployed. Wherever possible a multi-layered approach should be used (desktops, servers, gateways, etc.) that ensures all electronic files are appropriately scanned for viruses. Users are not authorized to turn off or disable virus-checking systems.

Access Controls: Physical and electronic access to Confidential Information and computing resources must be controlled. To ensure appropriate levels of access, a variety of security measures must be instituted as recommended by the Information Security Officer. Mechanisms to control access to Confidential Information include (but are not limited to) the following methods:

Authorization: Access will be granted on a "need to know" or "Minimum Necessary" basis and must be authorized by the immediate information owner or user management with the assistance of the Information Security Officer. Any of the following methods are acceptable for providing access under this policy:

<u>Context-based access:</u> An access control based on the context of a transaction (as opposed to being based on attributes of the initiator or target). The "external" factors might include time of day, location of the user, strength of user authentication, etc.
Role-based access: An alternative to traditional access control models (e.g., discretionary or non-discretionary access control policies) that permits the specification and enforcement of enterprise-specific security policies in a way that maps more naturally to an organization's structure and business activities. Each user is assigned to one or more predefined roles, each of which has been assigned the various privileges needed to perform that role.
<u>User-based access:</u> A security mechanism used to grant users of a system access based upon the identity of the user.

Identification/Authentication: Unique user identification (user id) and authentication is required for all systems that maintain or access Confidential Information. Users will be held accountable for all actions performed on the system with their user identification. At least one of the following authentication methods must be implemented:
strictly controlled passwords, biometric identification, and/ortokens in conjunction with a PIN. The user must provide authentication that is known only to that user and a designated security manager. An automatic timeout re-authentication must be required after a certain period of no activity. The maximum 15 minutes unless the department user(s) has a business reason for a longer period, as approved by the Information Security Officer. The user must log off or secure the system when leaving it.

Data Authentication: The Organization must be able to provide corroboration that confidential information has not been altered or destroyed in an unauthorized manner, to include, but not limited to: the use of check sums, double keying, message authentication codes, or digital signatures.

Remote Access: Confidential information that is stored or accessed remotely must maintain the same level of protections as information stored and accessed within the network.

Physical Access: Access to areas in which Confidential Information is accessed or stored must be controlled. The following physical controls must be in place: Mainframe computer systems must be installed in an access-controlled area. The area in and around the computer facility must afford protection against fire, water damage, and other environmental hazards, such as power outages and extreme temperature situations. File servers containing Confidential Information must be installed in a secure area to prevent theft, destruction, or access by unauthorized individuals. Workstations or personal computers (PCs) must be secured against use by unauthorized individuals. Local policies must be developed on secure and appropriate workstation use and physical safeguards which must include procedures that will: Position workstations to minimize unauthorized viewing of Confidential Information. Grant workstation access only to those who need it in order to perform their job function. Establish workstation location criteria to eliminate or minimize the possibility of unauthorized access to Confidential Information. Employ physical safeguards as determined by risk analysis, such as locating workstations in controlled access areas or installing covers or enclosures to preclude passerby access to Confidential Health Information. Use automatic screen savers with passwords to protect unattended machines that are turned on. Power down unused systems at end of day.

Emergency Access: Each department is required to establish a mechanism to provide emergency access to systems and applications in the event that the assigned custodian or owner is unavailable , or when operating under an emergency mode. Procedures must be documented to address: Authorization, Implementation, and Revocation.

Information Disposal: Information disposal must be consistent with established departmental records retention schedules. The disposal of information must ensure the continued protection of Confidential Information. Hard copy (paper and microfilm/fiche) must be shredded before being discarded or confidentially recycled. Magnetic media (floppy disks, hard drives, zip disks, etc.) must be erased with a degaussing device or disk "wiping" software before being discarded or reused. CD ROM Disks must be defaced or broken in half before being discarded.

Data Transfer/Printing:
Electronic Mass Data Transfers: Downloading and uploading Confidential Information between systems must be strictly controlled. Individually identifiable health information may only be transferred in accordance with the Privacy Policy.

Other Electronic Data Transfers and Printing: Confidential Information must be stored in a manner inaccessible to unauthorized individuals. Confidential Information must not be downloaded, copied or printed indiscriminately or left unattended and open to compromise. Medical information that is downloaded for educational purposes must, where possible, be de-identified before use.

Other Media: Confidential Information stored on any external media (diskettes, cd-roms, portable storage, memory sticks, etc…) must be protected from theft and unauthorized access. Such medium must be appropriately labeled so as to identify it as Confidential Information. Further, external medium containing Confidential Information must never be left unattended in unsecured areas. . Use of campus mail to deliver Confidential

Information requires the use of sealed envelopes, marked confidential, and the inclusion of both a delivery and return address.

Confidential Information must never be stored on mobile computing devices (laptops, personal digital assistants (PDA), smart phones, tablet PC's, etc.) unless they have the following minimum security requirements implemented:

- Power-on passwords
- Auto logoff or screen saver with password, and
- Encryption of stored data or other acceptable safeguards approved by Information Security Officer.

Further, mobile computing devices must never be left unattended in unsecured areas. If Confidential Information is stored on external medium or mobile computing devices and there is a breach of confidentiality as a result, then the owner of the medium/device will be held personally accountable and is subject to the terms and conditions of Information Security Policies and Confidentiality Statement signed as a condition of employment or affiliation.

Audit Controls: The audit processes be implemented to examine logged information in order to identify questionable data access activities, investigate breaches, respond to potential weaknesses, and assess the security program.

Communications/Network Controls: Technical security mechanisms must be put in place to guard against unauthorized access to data that is transmitted over a communications network. The following features must be implemented:

Backup & Recovery: Controls must recovery from any damage to computer equipment or files within a reasonable period of time. Each entity is required to develop and maintain a plan for responding to a system emergency that includes performing backups, preparing critical facilities that can be used to facilitate continuity of operations in the event of an emergency, and recovering from a disaster. An applications and data criticality analysis must be developed and documented to assess the sensitivity, vulnerabilities, and security of its programs and information it receives, manipulates, stores, and/or transmits. A data backup plan must be documented and routinely updated to create and maintain, for a specific period of time, retrievable exact copies of information. A disaster recovery plan must be developed and documented which contains a process enabling the entity to restore any loss of data in the event of fire, vandalism, natural disaster, or system failure. An emergency mode operation plan must be developed and documented which contains a process enabling the entity to continue to operate in the event of fire, vandalism, natural disaster, or system failure. Testing and revision procedures should be developed and documented requiring periodic testing of written contingency plans to discover weaknesses and the subsequent process of revising the documentation, if necessary. Backup data must be stored in an off-site location and protected from physical damage. Backup data must be afforded the same level of protection as the original data.

Equipment Control: Procedures must be in place which ensure proper equipment control. Documented security procedures must be established for bringing hardware and software into and out of a facility and for maintaining a record of that equipment. This includes, but is not limited to, the marking, handling, and disposal of hardware and storage media.

Compliance

The failure to comply with Information Security Policies and Standards may result in disciplinary action, up to and including dismissal, in accordance with applicable procedures, or, in the case of outside affiliates, termination of the terms of affiliation. Failure to comply with Information Security Policies and Standards by students may constitute grounds for corrective action in accordance with procedures. Possible disciplinary/corrective action may be instituted for, but is not limited to, the following:

- ✓ Unauthorized disclosure of Confidential Information.
- ✓ Unauthorized disclosure of a sign-on code (user id), password, or PIN.
- ✓ Attempting to obtain a sign-on code or password that belongs to another person.
- ✓ Using or attempting to use another person's sign-on code or password.

- ✓ Unauthorized use of an authorized password to invade patient privacy by examining records or information for which there has been no request for review.

- ✓ Installing or using unlicensed software on computers.

- ✓ The intentional unauthorized destruction of information.

Attempting to get access to sign-on codes for purposes other than official business, including completing fraudulent documentation to gain access.

4. Information Security Procedures and Guidelines

Procedures are specific operational steps or manual methods that workers must employ to achieve a certain goal. For instance, in many information technology departments there are specific procedures for performing back-ups of server hard drives. In this example, a policy could describe the need for back-ups, for storage off-site, and for safeguarding the back-up media (using encryption, physical security, etc.). A standard could define the software to be used to perform back-ups and how to configure this software. A procedure could describe how to use the back-up software, the timing for making back-ups, and other ways that humans interact with the back-up system (how to restore a file, how to securely transport storage media to an off-site repository, etc.).

4.1 User accounts and passwords

Passwords - Passwords are a critical element in protecting the infrastructure. Remember, your security policy is only as good as the weakest link. If you have weak passwords then you are at a higher risk for compromise not only by external threats, but also from insiders. If a password is compromised through social engineering or password cracking techniques, an intruder now has access to your resources. The result will mean that, you have just lost confidentiality and possibly the integrity of the data, and availability may have been compromised or in progress.

The policy should clearly state the requirements imposed on users for passwords. Passwords should not be any of the following:

- Same as the username;
- Password;
- Any personal information that a hacker may be able to obtain (e.g., street address, social security number, names of children, parents, cars, boats, etc.);
- A dictionary word; or
- Telephone number

These are some examples of passwords not to use. You should force users through automated password policy techniques to require a minimum of eight characters, use of a combination of symbols, alpha charters, and numerals, and a mixture of uppercase and lowercase. Users should be required to change their password at least quarterly. Previous passwords should not be authorized. Lastly, an account lockout policy should be implemented after a predetermined number of unsuccessful logon attempts.

Another tip to consider is that you should be logging all successful and failed logon attempts. A hacker may be trying several accounts to logon to your network. If you see several ‘failed’ logon attempts in a row and then no activity; does this mean the hacker gave up or did he “successfully” logon?

4.2 Access Control

User Access to Computer Resources - This section should identify the roles and responsibilities of users accessing resources on the organization’s network. This should include information such as:

- Procedures for obtaining network access and resource level permission;
- Policies prohibiting personal use of organizational computer systems;
- Passwords;
- Procedures for using removal media devices;
- Procedures for identifying applicable e-mail standards of conduct;
- Specifications for both acceptable and prohibited Internet usage;
- Guidelines for applications;
- Restrictions on installing applications and hardware;
- Procedures for Remote Access;
- Guidelines for use of personal machines to access resources (remote access);
- Procedures for account termination;
- Procedures for routine auditing;
- Procedures for threat notification; and
- Security awareness training;

Depending on the size of an organization’s network, a more detailed listing may be required for the connected Wide Area Networks (WAN), other Local Area Networks (LAN), Extranets, and Virtual Private Networks (VPN).

Some organizations may require that other connected (via LAN, WAN, VPN) or trusted agency’s meet the terms and conditions identified in the organization’s security policy before they are granted access. This is done for the simple reason that your security policy is only as good as the weakest link. For example, If Company ‘A’ has a rigid security policy and Company ‘B’ has a substandard policy and wants to partner with Company ‘A’, Company ‘B’ may request to have a network connection to Company ‘A’ (behind the firewall). If Company’ A’ allows this without validating Company ‘B’s’ security policy then Company ‘A’ can now be compromised by exploits launched from Company ‘B’.

When developing a security policy one should take situations such as this one very serious and develop standards that must be met in order for other organizations to be granted access. One method is to require the requesting organization to meet, at a minimum, your policy and guidelines.

4.3 Hardware acquisition

Technology is increasingly creating value for the customer but at the same time, technology is also destroying profits for the service providers like the banks. Does this mean that the banks can ignore technology? The choice before the banks is between annihilation by not adopting technology and bankruptcy because of the extreme competition and lowering profit margins.

Banking is very far from being a homogenous activity. In fact it is a collection of more than 150 specific products / market services. Lines of business differ by customer whether retail, corporate or other financial institutions. They differ by distribution channel, whether by branch, by direct salesman or by mail. They differ by product route, whether lending, deposit gathering or payment product.

What should be the overall strategy for the bankers and financial institutions in using information technology? The following four points can act as guiding factor:

- Link business strategy more effective with technological reality
- Adopt simultaneous tight-loose policies to manage systems investments
- Treat routine automation differently from distinctive automation
- Consider industry systems capacity when making product decisions

With the above background, it is very clear that though technology is a necessity, the cost of acquiring it and its perceived benefits should weigh more in the minds of Bankers.

4.4 Operating System Security

From the moment you boot your computer (after the BIOS loads), you are interacting with the *operating system*. This integral piece of software defines what you can do with the computer system and how you do it. Whether you're interacting with the file system or chatting with someone on an instant messenger program, the operating system is working behind the scenes to provide you with a (hopefully) flawless experience as it interprets your actions and converts them into something your computer can process.

While operating systems vary on many levels, the most common operating systems provide much more than a simple interface between user and machine. Included are programs that provide the user with numerous extras, from simple screen savers to complex file-encryption schemes. However, it's important to understand that these programs are extras that are added to the OS and are not necessary for the computer to operate.

Many users become intimately familiar with the operating system's accessories (such as Solitaire), but forget about the security features that are included to help the user maintain a safe and reliable operating environment. As a result, many information systems exist in an insecure state that leaves the system at risk to a virus infection or a complete compromise by an attacker.

This section is dedicated to operating system security issues. From setting up a secure home network to creating strong passwords, it's important to understand the details of using an operating system in a safe and secure manner. In today's connected world, it's irresponsible to set up a computer without regard to security. It takes only one virus or Trojan horse to create a ripple effect of infected computers and compromised systems.

Before delving into the security side of an operating system (OS), it's important that you know where the OS begins and ends. This brief overview describes the functionality and purpose of the operating system and how it's used to create the computing experience.

"Any computer system includes a basic set of programs called the *operating system*. The most important program in the set is called the *kernel*. It is loaded into [memory] when the system boots and contains many critical procedures that are needed for the system to operate. The other programs are less crucial utilities; they can provide a wide variety of interactive experiences for the user—as well as doing all the jobs the user bought the computer for—but the essential shape and capabilities of the system are determined by the kernel."

OS Functions

In short, the OS must provide two main functions:

- It must manage the resources available to the computer system.
- It must provide a reliable, stable, secure, and consistent interface for applications to access the computer's resources.

The first function is critical to the OS because it defines how applications access the system's resources. By controlling the various aspects of how hardware and software are used, the OS ensures that every application gets a chance to use the processor. The second, related function defines the methods by which an application can access these resources. Because the OS often acts as a buffer between an executing program and the hardware, it needs to provide some means of allowing applications to access resources without needing to know the details of each and every unique computer system.

OS Types

There are four main types of operating systems, classified according to the types of programs they support and way these programs interact with users:

- **Real-time operating system.** This OS is most often found in robotic machinery and scientific devices. It doesn't provide much room for user operation, with the exception of some configuration changes. Typically, this OS contains highly polished timing mechanisms due to the impact even the slightest error could have in automated production or measurements.

- **Single-user, single task system.** This type of OS is used by devices such as a PDA or other miniature computers. It basically allows one user to operate one program at a time. If another program is needed, the user must close the currently executing application.
- **Single-user, multitasking system.** This type of OS is most familiar because it includes most Microsoft Windows systems. In this OS, a user can open multiple programs and jump back and forth between applications as required. In fact, there is much debate that although Windows Server OSes appear to be multiuser systems, they're actually single-user, multitasking Oses (with the exception of Terminal Services).
- **Multi-user system.** A true multi-user operating system allows many users to access the computer's resources simultaneously. A common example of this type of OS is Linux. In this type of system, the OS manages requests from numerous users, and maintains rigorous control over the resources to ensure that one user doesn't affect any other user.

OS Tasks

The OS is responsible for various tasks within the computing environment. These tasks are often what make one OS more reliable or easier to use than another, and also determine the power of the OS:

- **Processor management.** The OS needs to ensure that each application gets a share of the processor's time, and that the processor is used efficiently to accomplish real work.
- **Memory management.** This defines the methods by which the OS allocates memory to applications and OS functions.
- **Device management.** Because a computer system is composed of various hardware components (hard drive, monitor, mouse, keyboard, and so on), the OS must be able to manage how these components interact with each other.
- **Storage management.** The OS not only controls active resources, but defines how files and data are stored in a reliable fashion.

- **Application interface.** An OS is really a bridge between applications and the computer's resources, which means that it must provide application-programming interfaces (APIs) for applications to connect.
- **User interface.** Whether this is a command line or a graphical user interface (GUI), the OS is responsible for interacting with the end user.

This is a very brief summary of the major tasks that an OS should handle. The following sections describe security-related issues that the OS must also deal with to maintain confidentiality, integrity, and availability of system resources.

OS Security Weaknesses

Now that you have had a brief overview of what an operating system should provide for a user with regard to functionality, let's take a look at the security aspects of some favorite operating systems. In this section, we discuss the two most common operating system families and the security features they include. We also examine methods by which these security features can be attacked and/or bypassed, and how to protect against these types of attacks.

4.5 Data Classification

Classification is used to promote proper controls for safeguarding the confidentiality of information. Regardless of classification, the integrity and accuracy of all classifications of information must be protected. The classification assigned, and the related controls applied, are dependent on the sensitivity of the information. Information must be classified according to the most sensitive detail it includes. Information recorded in several formats (e.g., source document, electronic record, report) must have the same classification regardless of format. The following levels are to be used when classifying information:

Confidential Information

Confidential Information is information that is protected by State or Federal Statute. It represents very important and highly sensitive material. This information is private or otherwise sensitive in nature and must be restricted to those with a legitimate business need for access.

Examples of Confidential Information may include: personnel information, proprietary information of commercial research sponsors, information concerning select agents, individual health information, student records, system access passwords, and information file encryption keys.

Unauthorized disclosure of this information to people without a business need for access may be against laws and regulations, or may cause significant problems for the organization, or its business associates. Decisions about the provision of access to this information must always be cleared through the information owner.

Public Information

Public information includes all information that is not confidential.
Examples of Public Information may include Organization's Events Calendars, class schedules, minutes of meetings, and material posted to the Organization's web pages.

4.6 Incident Handling

Incident Handling refers to those practices, technologies and/or services used to respond to suspected or known breaches to security safeguards. Once a suspected intrusion activity has been qualified as a security breach (i.e., incident), it is imperative that the incident be contained as soon as possible, and then eradicated so that any damage and risk exposure to the Agency and the Commonwealth are avoided or minimized. Information technology security incidents refer to deliberate, malicious acts which may be technical (e.g., creation of viruses, system hacking) or non-technical (e.g., theft, property abuse, service disruption). In several cases, if the incident is left "unchecked" (i.e., not contained), then the damage resulting from these incidents continues to spread within, and across, Agencies.

Handling incidents can be logistically complex, and may require information and assistance from sources outside the Agency (e.g., technical specialists, law enforcement entities such as state police and the public affairs office). Industry best practices suggest that organizations who adopt both proactive and reactive means to address incident handling are better able to limit the negative implications of incidents. Examples of proactive activities include establishing communication mechanisms to report incidents and to disseminate incident alerts; and identifying technical experts who can provide emergency assistance if needed. Examples of reactive activity include blocking or aborting computer processes; temporarily denying user access; and deploying inoculation software.

Information security technology traditionally focuses on protecting the perimeter to keep the bad guys and the bad code out of the enterprise. But as every CIO knows, information security breaches in large enterprises are inevitable. Hackers will penetrate the network, or — in what many believe are more frequent occurrences — insiders will compromise customer and company data. With such compromises a certainty, enterprises are left scrambling to manage these proliferating incidents.

Recently, intrusion detection systems (IDS) have emerged to help detect perimeter breaches and intrusions. However, organizations that install these systems quickly find that in order to be effective, they need a process to immediately respond to those alarms and confirm the incident to enable sound incident management, containment and mitigation. According to Gartner research, “intrusion detection sounds like a good idea but alerts you only that something is going on. It is not always so effective to just see the alarms going off" and not have the tools to address the problem. Too often, however, the IDS bells ring, but with no effective means to respond and sort out all the false positives, the IDS becomes white noise, and, ultimately, shelfware.

So how are enterprises responding to and addressing all these, alarms, intrusions and compromises? After all, any effective security process – whether it be home security, industrial security, or national security – needs an effective response mechanism that enables effective incident handling and containment. By analogy, home security systems that do not provide an armed response do little more than annoy the neighbors.

Proper incident response is a critical component of the information security framework, as dictated by necessity, and mandated by industry best practices and regulations. For too many organizations, however, the incident response process is either non-existent or merely consists of attempting to patch systems and restore them to their believed previous state, often days or weeks after the incident has occurred and the damage proliferated. The typical incident response is often *ad hoc* and haphazard, involving tedious analysis of log files, invasive analysis of compromised systems using disjointed tools, and the disruption of mission critical systems. When and if the incident is finally confirmed, the problem is kept under wraps with organizations content to move on without preserving evidence and properly documenting the incident for reporting, prosecution or internal "lessons learned" analysis. However, failing to adopt a formal incident response, or simply employing a so-called "patch and proceed" approach is no longer viable under the current state of industry best practices for several reasons, which are set forth below.

Effective Containment and Mitigation Requires Immediate Response

Computer security incidents are like cancer—early intervention and containment are critical in order to prevent the spread of the problem. In many cases, however, the incident response process involves weeks-long delays while the organization's computer incident response team (CIRT) or outside consultants travel to the site of the compromised systems, followed by painstakingly inefficient analysis of log files and other data with stand-alone utilities. Meanwhile, either the damage proliferates while the CIRT tries to confirm whether the incident did in fact take place, or mission-critical systems are taken off-line, causing disruption of operations and substantial monetary loss. In incident response investigations, the analysis must be as rapid as possible to mitigate the loss and increase the likelihood of identifying the culprit.

For instance, hackers and malevolent insiders often cover their tracks by deleting event log and system files, hiding their installed malware by renaming it with innocuous file extensions, cloaking created backdoors, and other similar techniques. Deleted file recovery is a particularly crucial incident response function, as in addition to erasing log files to mask the incident; perpetrators will also maliciously delete system files and other critical company data. This deleted information must be quickly restored before it is overwritten and lost. Network-enabled computer forensics tools can quickly undelete files, locate hidden malware (even if renamed) through file signature and hash analysis, find backdoors and other evidence, and make complete bit-stream image backups of drives housing compromised data. Additionally, network-enabled computer forensics operate in a live environment, which allows a very rapid response without taking any systems off-line and thus disrupting operations. Additionally, as the target systems are not taken off-line, the key live data of the compromised system (open ports, live registry, RAM dumps) can be easily captured and preserved. Rapid deleted file recovery, disk imaging, file signature and hash analysis, and live data capture are only some of the key functions that network-enabled computer forensic software provide for effective incident response.

Ironically, incident response teams that simply seek to “patch and proceed” without understanding the cause of the intrusion and the extent of the compromise either fail to fully contain the damage, or at a minimum leave systems vulnerable to further compromise. Far from being a hindrance to quickly and safely proceeding with business, rapid enterprise incident response with network-based computer forensics tools enable prompt and detailed incident identification, management and recovery, all while preserving data for subsequent post-mortem analysis. Among many other benefits, this allows the CIRT to fully understand the scope and nature of the incident for proper and thorough recovery and remediation.

Effective Incident Detection Requires an Integrated Response Process_

In June 2003, Gartner created a major stir in the information security industry when it issued a research report calling into question the effectiveness of intrusion detection systems. The report listed several problems associated with the IDS process, including a high rate of false positives and negatives, the burden associated with the need full-time monitoring by information systems staff, and “a taxing incident response process.” The report ultimately questions the value of future IDS investments and recommends that organizations spend their resources on perimeter protection such as firewalls.

However, network-enabled computer forensics analysis represents a natural extension of an IDS alert by providing a very rapid and thorough incident response process, anywhere on the network. An enterprise-wide computer forensics system can be utilized to quickly and concisely confirm whether or not the system is truly compromised. Instead of manually pouring through log files to attempt to confirm an incident, network-enabled computer forensics tools provide a means to thoroughly analyze, from one central location, any potentially compromised system anywhere on the network. This capability greatly enhances the effectiveness of IDS by providing an accurate confirmation and response mechanism to intrusion alerts.

Further, this rapid response and confirmation ability greatly reduces the administrative burden in monitoring and responding to IDS alerts. The latest generation of network-enabled computer forensics systems represents a quantum leap in the power, efficiency and accuracy of the incident response process. Integrating IDS monitoring with a networked computer forensics system is a crucial step that enables organizations to greatly improve the effectiveness of their IDS, while providing a broad and highly effective enterprise-wide investigation capability.

The Insider Threat Requires Response and Investigation

Many computer security professionals tend to think of computer security incidents as problems that originate from outside the perimeter, such as denial of service attacks, worm infections, and website defacements. As a result, many incident response teams fail to recognize and prepare for security compromises perpetrated by insiders. This is a serious mistake that results in substantial losses and costs. As reflected by recent surveys, many incidents, particularly those resulting in significant financial harm, are the work of rogue employees and other trusted individuals.

The insider threat takes many forms, whether it is unauthorized access to customer privacy information, theft of intellectual property and trade secrets, financial fraud, improper deletion of computer files (as in the case of Arthur Andersen) or various employee policy violations such as email harassment and Internet pornography. Notably, industry regulations and best practices do not differentiate between computer incidents with internal or external origins. As such, the incident response process needs to be just as effective in addressing the internal theft of intellectual property as with denial of service attacks.

In terms of defined industry best practices, ISO 17799 provides very detailed requirements for incident response, internal investigations, and preservation and analysis of computer evidence consistent with best practices and computer forensics protocols. An enterprise's overall security framework must, under ISO 17799, include an effective incident response approach "to ensure a quick, effective and orderly response to security incidents." An ISO 17799-compliant enterprise should employ the best methods and tools available to respond to breaches or suspected breaches of its information security, and must collect and preserve the resulting evidence in a forensically sound manner for investigation and reporting purposes.

The leading international financial standards-setting institution, the Basel Committee on Banking Supervision (the "Basel Committee") has promulgated important new standards for electronic banking. In a report entitled "Risk Management for Electronic Banking," the Basel Committee addresses several components of information security, including a strong focus on the necessity of incident response processes. In the report, the Basel Committee establishes that " effective incident response mechanisms are . . . critical to minimize operational, legal and reputation risks arising from internal and external attacks." As a result, banks should "develop appropriate incident response plans . . . that ensure business continuity, control reputation risk and limit liability associated with disruptions in their e-banking services."

In order to implement this Risk Management Principle, the Basel Committee highlighted eight specific actions that should be undertaken by banks, including the following four functions or capabilities that banks should develop:

- Incident response plans to address recovery of e-banking systems and services under various scenarios, business and geographic locations . . .

- Mechanisms to identify [a] crisis as soon as it occurs, assess its materiality, and control the reputation risk associated with any disruption in service.

- Incident response teams with the authority to act in an emergency and sufficiently rained in analyzing incident detection/response systems and interpreting the significance of related output.

- A process for collecting and preserving forensic evidence to facilitate appropriate post-mortem reviews of any e-banking incidents as well as to assist in the prosecution of attackers.

Thus, a key factor for banks is to be able to quickly and thoroughly respond to security incidents. In order to manage risks adequately, banks must have contingency plans in place to address incidents as they occur, and those plans "should set out a process for restoring or replacing e-banking processing capabilities, reconstructing supporting transaction information, and include measures to be taken to resume availability of critical e-banking systems and applications in the event of a business disruption."

Under these compelling regulations and defined best practices, organizations must employ the best methods and tools available to respond to breaches or suspected breaches of its information security, and must collect and preserve the resulting evidence in a forensically sound manner. For incident response, "best practices" is embodied by network-enabled incident response and computer forensics systems for computer security incidents that occur throughout an organization's network.

Conversely, the "patch and proceed" methodology is not compliant with these regulations and standards for two reasons. First, with the growing standardization of network-enabled computer forensics tools, "patch and proceed" is simply no longer consistent with best practice. Secondly, without the proper response, collection and preservation of evidence, the internal and regulatory incident reporting requirements under these regulations and standards cannot be met.

The Defense and Prosecution of Claims Require Preservation of Evidence

In addition to information security regulatory requirements, enterprises face an increasing risk of various cyber liability claims stemming from such security breaches as theft of customer privacy data, denial of service attacks that are launched from a company's compromised systems, misappropriation of intellectual property, insider financial fraud, and the destruction of computer data. These legal claims take many forms, including government enforcement actions, class action suits from customers, shareholder suits (for lack of internal controls to investigate and stem insider fraud), and other claims.

Organizations that fail to properly respond to a myriad of potential liability-causing incidents will find themselves unable to defend their interests in court, subjecting the enterprise to significant legal exposure.

More than just being able to defend its interests, companies are increasingly seeking to prosecute perpetrators, particularly in cases of industrial espionage, financial fraud, and theft of intellectual property and trade secrets. However, a company may not be able to prosecute or obtain a civil injunction against an employee who leaves the company with source code or the customer list if the digital evidence trail is not properly collected and preserved. Courts mandate that computer evidence be collected and handled in a manner consistent with best practices. Additionally, claims under cyber-insurance policies, which are growing increasingly popular, often require proper incident handling and preservation of evidence (usually under the provisions of ISO 17799).

4.7 Change Management

Change Control Procedures

In order to minimize the corruption of information systems there should be strict control over the implementation of changes. Formal change control procedures should be enforced. They should ensure that security and control procedures are not compromised, that support programmers are given access only to those parts of the system necessary for their work, and that formal agreement and approval for any change is obtained. Changing application software can impact the operational environment. Wherever practicable, application and operational change control procedures should be integrated. In particular following controls should be considered.

- Identification and recording of significant changes
- Assessment of the potential impact of such changes
- Formal approval procedure for proposed changes
- Communication of change details to all relevant persons
- Procedures identifying responsibilities for aborting and recovering from unsuccessful changes

Change control process should include:

- Maintaining a record of agreed authorization levels
- Ensuring changes are submitted by authorized users
- Reviewing controls and integrity procedures to ensure that they will not be compromised by the changes
- Identifying all computer software, information database entities and hardware that require amendment;
- Obtaining formal approval for detailed proposals before work commences

- ensuring that the authorized user accepts changes prior to any implementation
- Ensuring that implementation is carried out to minimize business disruption
- Ensuring that the system documentation set is updated on the completion of each change and that old documentation is archived or disposed of
- Maintaining a version control for all software updates
- Maintaining an audit trail of all change requests
- Ensuring that operating documentation and user procedures are changed as necessary to be appropriate
- Ensuring that the implementation of changes takes place at the right time and is not disturbing the business processes involved.

Many organizations maintain an environment in which users test new software and which is segregated from development and production environments. This provides a means of having control over new software and allowing additional protection of operational information that is used for testing purposes.

Technical review of operating system changes

Periodically it is necessary to change the operating system e.g. to install a newly supplied software release or patches. When changes occur, the application systems should be reviewed and tested to ensure that there is no adverse impact on operation or security. This process should cover:

- Review of application control and integrity procedures to ensure that they have not been compromised by the operating system changes;
- Ensuring that the annual support plan and budget will cover reviews and system testing resulting from operating system changes;
- Ensuring that notification of operating system changes is provided in time to allow appropriate reviews to take place before implementation;
- Ensuring that appropriate changes are made to the business continuity plans.

Restrictions on changes to software packages

Modifications to software packages should be discouraged as far as possible and vendor supplied software packages should be used without modification. Where it is deemed essential to modify the software package, the following points should be considered:

- The risk of built-in controls and integrity process being compromised
- Whether the consent of the vendor should be obtained
- The possibility of obtaining the required changes from the vendor as standard program updates
- The impact if the organization becomes responsible for the future maintenance of the software as a result of changes.

If changes are deemed essential the original software should be retained and the changes applied to a clearly identified copy. All changes should be fully tested and documented, so that they can be reapplied if necessary to future software upgrades.

Covert Channels and Trojan code

A covert channel can expose information by some indirect and obscure means. It may be activated by changing a parameter accessible by both secure and insecure elements of a computing system, or by embedding information into a data stream. Trojan code is designed to affect a system in a way that is not authorized and not readily noticed and not required by the recipient or user of the program. Covert channels and Trojan code rarely occur by accident. Where covert channels and Trojan code rarely occur by accident. Where covert or Trojan code are a concern, the following should be considered

- Buying programs only from a reputable source
- Buying programs in source code so the code may be verified
- Using evaluated procedure
- Inspecting all source code before operational use
- Controlling access to and modification of code once installed
- Use staff of proven trust to work on key systems

Outsourced software development

Where software development is outsourced, the following points should be considered:

- Licensing arrangements, code ownership and intellectual property rights
- Certification of the quality and accuracy of the work carried out
- Escrow arrangements in the event of failure of the third party
- Rights of access for audit of the quality and accuracy of work done

- Contractual requirements for quality of code
- Testing before installation to detect Trojan code

4.8 Virus Control Procedures

Virus Control procedures

The following procedure was developed to insure network integrity as well as performance

Prevention

Each local area network should follow a regularly scheduled anti-virus screening procedure. This procedure should include the following items.

- Prior to contact with any systems/networks, and using Anti-virus software, scan all floppy disks on a system(s) specifically designated for this purpose.
- Any system that has been serviced should be scanned before the service request is signed off and before any re-introduction of the system to the network.
- If possible, anti-virus software should be memory resident on all workstations. On all local area networks, anti-virus software should run as an NLM on the Novell file server. If this is not possible, the workstations and file server should be scanned at the close of the business day or at a minimum weekly.
- Downloads from public bulletin boards or networks should be discouraged.

Virus Isolation and Eradication

As part of any troubleshooting procedure, especially if the immediate analysis offers no explainable cause of the problem in question, or if the profile fits a virus, a virus scan should be executed from a verified clean, write protected, bootable, floppy disk.

Standalone Virus Procedure

If a virus is suspected of being present and a virus scan can positively identify the virus, use the following standalone virus procedure.

- Inform the user not to use their PC or any floppy disks in their possession.
- Scan the PC using the Anti-virus Software.
- Scan all diskettes and any other resources that may have come in contact with the infected PC.
- Notify the Information Services Department of the virus type and any contacts that may potentially be infected. The Help Desk will act as the coordinating office for information regarding the infection. The Help Desk will notify all LAN Administrators of the potential problem.
- Once the virus has been identified and eradicated, notify the Help Desk for dissemination of information. If the threat can be positively identified as a stand alone virus, there is no need to disconnect the PC from the LAN else, treat the infection/problem as a traveling virus and follow its procedure.

Traveling Virus Procedure

If a virus is suspected of being present and the virus cannot be identified or is identified as a transmittable or "traveling" virus, the following procedure should be followed.

- The personal computer should be physically disconnected from the LAN.
- The LAN should be physically isolated from other LANs and NET.
- All access and mail servers on the LAN should be shutdown.
- The user should be informed not to use their PC or any floppy disks in their possession.
- The Information Services Department Help Desk must be notified of the "possibility" of the presence of a "traveling" virus on a LAN connected to NET. The Help Desk will electronically isolate the LAN from NET. The Help Desk should then notify all NET LAN

Administrators of the potential problem and act as the coordinating office for information regarding the problem.

- All resources on the LAN should be scanned using Anti-virus software. Don't forget to scan diskettes and standalone PCs that have shared resources or any contact with the infected LAN.
- All file servers should be scanned on an hourly basis until all scanning procedures are complete and a clean bill of health is issued by the LAN Administrator.

- The Information Services Department Help Desk should be contacted to determine if any further incidence of infection is still present on other LANs on the wide area network. If all indications from the Help Desk are that NET is virus free, they will electronically enable LAN access back to NET. The LAN Administrator may then reconnect the NET router and restore the access and mail servers.

4.9 Database Security

One of the more recent evolutions in network security has been the movement away from protecting the perimeter of the network to protecting data at the source. The reason behind this change has been that perimeter security no longer works in today's environment. Today, more than just your employees need access to data. Essentially, partners and customers must have access to this data as well meaning that your database cannot simply be hidden behind a firewall.

Of course, as your databases become more exposed to Internet, it is imperative that you properly secure them from attacks from the outside world. Securing your databases involves not only establishing a strong policy, but also establishing adequate access controls. In this paper, we will cover various ways databases are attacked, and how to prevent them from being “hacked”.

Current Database Security Environment

It is very easy in the security community to create an air of fear, uncertainty, and doubt (FUD). As both security and database professionals, it is important to see through the FUD, determine the actual risks, and investigate what can be done about the situation. The truth is most databases are configured in a way they can be broken into relatively easily. However, this is not to say that databases cannot be made properly secured. It is the information to properly lock down these databases that has not been made available, and that the proper lockdown procedures have not been taken. On the other hand, the number of databases compromised so far has not been nearly on the scale that we have seen web servers being attacked and compromised. The reasons for this are several:

- There are less databases than web servers
- Knowledge of database security has been limited
- Getting a version of enterprise databases to learn and test on was difficult
- Databases were traditionally behind a firewall

All of the above has changed significantly over the past few years. First, there is an increasing interest for databases in the Black Hat hacker community. The number of talks on database security has grown significantly over the past few years. The number of database exploits reported on security news groups has increased significantly.

Next, downloading database software has also become much simpler. All of the latest versions are available for download from each vendor's respective websites for anyone with a fast enough Internet connection. Also, for each of the major platforms profiled within this white paper, the installation process has become increasingly simple in design. Lastly, the increasing number of threats against databases is not going to cause the end of the world. However, it is necessary for us to start taking database security seriously by taking a proactive approach to understand the risks and lockdown procedures.

Understanding Vulnerabilities

In order to understand vulnerabilities, we should start by describing the various classes of vulnerabilities:

- Vendor bugs
- Poor architecture
- Misconfigurations
- Incorrect usage

Vendor Bugs

Vendor bugs are buffer overflows and other programming errors that result in users executing the commands they are allowed to execute. Downloading and applying patches usually fix vendor bugs. To ensure you are not vulnerable to one of these problems, you must stay aware of the patches, and install them immediately when they are released.

Poor Architecture

Poor architecture is the result of not properly factoring security into the design of how an application works. These vulnerabilities are typically the hardest to fix because they require a major rework by the vendor. An example of poor architecture would be when a vendor utilizes a weak form of encryption.

Misconfigurations

Misconfigurations are caused by not properly locking down databases. Many of the configuration options of databases can be set in a way that compromises security. Some of these parameters are set insecurely by default. Most are not a problem unless you unsuspectingly change the configuration. An example of this in Oracle is the REMOTE_OS_AUTHENT parameter. By setting REMOTE_OS_AUTHENT to true, you are allowing unauthenticated users to connect to your database.

Incorrect Usage

Incorrect usage refers to building applications utilizing developer tools in ways that can be used to break into a system. SQL INJECTION is an example of incorrect usage.

Database Worms

In the recent past, a new set of threats have emerged – worms that propagate through vulnerabilities in databases rather than through more traditional operating system or web server holes. Despite their lack of sophistication, these worms have been somewhat successful because of the poor state of database security. Security in databases has generally been ignored and the threat management of these applications has been non-existent.

The damage caused by a worm is dependent on several factors:

1) The number of targets for the worm

2) The success rate of infection

3) The resilience of the worm

A critical factor in the effect of a worm is the quantity of potential targets. Most people assume that databases are always behind firewalls. Unfortunately this is not always the case. Databases are a critical piece of an organizations infrastructure and cannot always be hidden behind a firewall. The success rate of infection is critical to whether or not the worm is able to spread through to other systems.

The Spida worm was effective because a large number of Microsoft SQL Server databases have blank “sa” passwords. Those databases with non-blank passwords were not infected.

The last factor is the resilience of the worm. Is it hard to detect? Does it leave a back door? Are there any bugs that cause the infection to fail on certain databases? A well-developed worm is much harder to fight than a “noisy” and “sloppy” worm that is easy to remove from the system.

The truth is, there are not many resources out there to keep up with database security. There are a few simple tasks that can be performed to reduce your security risk at a reasonable level.

- Stay patched
- Stay aware of database security holes.
- To ask questions on database security, check out:
- Explore possible third-party solutions

Provide multiple levels of security:

- Perform audits and pen tests on your databases regularly
- Encryption of data in motion
- Encryption of data at rest within the database
- Monitor your log files
- Implement intrusion detection

Databases are extremely complex beasts, and generic auditing, vulnerability assessment, and IDS solutions just don't cut it. It is strongly recommended that you find a vendor that caters directly to these specific applications or find a strong partner that understands databases. Databases are your most valuable assets, and you should place significantly more effort towards securing them.

By staying informed and aware of security vulnerabilities to you databases, you should be able to keep the risks to a minimum.

4.10 Network Security

A basic understanding of computer networks is requisite in order to understand the principles of network security. In this section, we'll cover some of the foundations of computer networking, then move on to an overview of some popular networks. Following that, we'll take a more in-depth look at TCP/IP, the network protocol suite that is used to run the Internet and many intranets.

Once we've covered this, we'll go back and discuss some of the threats that managers and administrators of computer networks need to confront, and then some tools that can be used to reduce the exposure to the risks of network computing.

The ISO/OSI Reference Model

The *International Standards Organization* (ISO) *Open Systems Interconnect* (OSI) Reference Model defines seven layers of communications types, and the interfaces among them. Each layer depends on the services provided by the layer below it, all the way down to the physical network hardware, such as the computer's network interface card, and the wires that connect the cards together.

An easy way to look at this is to compare this model with something we use daily: the telephone. In order for you and I to talk when we're out of earshot, we need a device like a telephone. (In the ISO/OSI model, this is at the application layer.) The telephones, of course, are useless unless they have the ability to translate the sound into electronic pulses that can be transferred over wire and back again. (These functions are provided in layers below the application layer.) Finally, we get down to the physical connection: both must be plugged into an outlet that is connected to a switch that's part of the telephone system's network of switches.

If I place a call to you, I pick up the receiver, and dial your number. This number specifies which central office to which to send my request, and then which phone from that central office to ring. Once you answer the phone, we begin talking, and our session has begun. Conceptually, computer networks function exactly the same way.

It isn't important for you to memorize the ISO/OSI Reference Model's layers; but it's useful to know that they exist, and that each layer cannot work without the services provided by the layer below it.

Figure 1: The ISO/OSI Reference Model

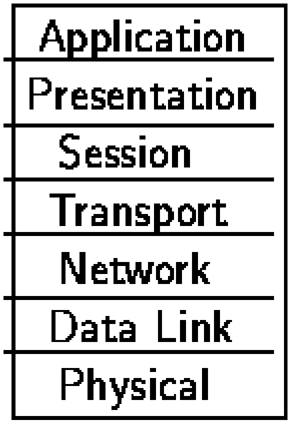

Application
Presentation
Session
Transport
Network
Data Link
Physical

Types And Sources Of Network Threats

Now, we've covered enough background information on networking that we can actually get into the security aspects of all of this. First of all, we'll get into the types of threats there are against networked computers, and then some things that can be done to protect yourself against various threats.

Denial-of-Service

DoS (Denial-of-Service) attacks are probably the nastiest, and most difficult to address. These are the nastiest, because they're very easy to launch, difficult (sometimes impossible) to track, and it isn't easy to refuse the requests of the attacker, without also refusing legitimate requests for service.

The premise of a DoS attack is simple: send more requests to the machine than it can handle. There are toolkits available in the underground community that make this a simple matter of running a program and telling it which host to blast with requests. The attacker's program simply makes a connection on some service port, perhaps forging the packet's header information that says where the packet came from, and then dropping the connection. If the host is able to answer 20 requests per second, and the attacker is sending 50 per second, obviously the host will be unable to service all of the attacker's requests, much less any legitimate requests (hits on the web site running there, for example).

Such attacks were fairly common in late 1996 and early 1997, but are now becoming less popular.

Some things that can be done to reduce the risk of being stung by a denial of service attack include

Not running your visible-to-the-world servers at a level too close to capacity

Using packet filtering to prevent obviously forged packets from entering into your network address space.

Obviously forged packets would include those that claim to come from your own hosts, addresses reserved for private networks as defined in RFC 1918, and the *loopback* network (127.0.0.0).

Keeping up-to-date on security-related patches for your hosts' operating systems.

Unauthorized Access

``Unauthorized access'' is a very high-level term that can refer to a number of different sorts of attacks. The goal of these attacks is to access some resource that your machine should not provide the attacker. For example, a host might be a web server, and should provide anyone with requested web pages. However, that host should not provide command shell access without being sure that the person making such a request is someone who should get it, such as a local administrator.

Executing Commands Illicitly

It's obviously undesirable for an unknown and untrusted person to be able to execute commands on your server machines. There are two main classifications of the severity of this problem: normal user access, and administrator access. A normal user can do a number of things on a system (such as read files, mail them to other people, etc.) that an attacker should not be able to do. This might, then, be all the access that an attacker needs. On the other hand, an attacker might wish to make configuration changes to a host (perhaps changing its IP address, putting a start-up script in place to cause the machine to shut down every time it's started, or something similar). In this case, the attacker will need to gain administrator privileges on the host.

Confidentiality Breaches

We need to examine the threat model: what is it that you're trying to protect yourself against? There is certain information that could be quite damaging if it fell into the hands of a competitor, an enemy, or the public. In these cases, it's possible that compromise of a normal user's account on the machine can be enough to cause damage (perhaps in the form of PR, or obtaining information that can be used against the company, etc.)

While many of the perpetrators of these sorts of break-ins are merely thrill-seekers interested in nothing more than to see a shell prompt for your computer on their screen, there are those who are more malicious, as we'll consider next. (Additionally, keep in mind that it's possible that someone who is normally interested in nothing more than the thrill could be persuaded to do more: perhaps an unscrupulous competitor is willing to hire such a person to hurt you.)

Destructive Behavior

Among the destructive sorts of break-ins and attacks, there are two major categories.

Data Diddling.

The data diddler is likely the worst sort, since the fact of a break-in might not be immediately obvious. Perhaps he's toying with the numbers in your spreadsheets, or changing the dates in your projections and plans. Maybe he's changing the account numbers for the auto-deposit of certain paychecks. In any case, rare is the case when you'll come in to work one day, and simply know that something is wrong. An accounting procedure might turn up a discrepancy in the books three or four months after the fact. Trying to track the problem down will certainly be difficult, and once *that* problem is discovered, how can any of your numbers from that time period be trusted? How far back do you have to go before you think that your data is safe?

Data Destruction.

Some of those perpetrate attacks are simply twisted jerks who like to delete things. In these cases, the impact on your computing capability -- and consequently your business -- can be nothing less than if a fire or other disaster caused your computing equipment to be completely destroyed.

Firewalls

As we've seen in our discussion of the Internet and similar networks, connecting an organization to the Internet provides a two-way flow of traffic. This is clearly undesirable in many organizations, as proprietary information is often displayed freely within a corporate *intranet* (that is, a TCP/IP network, modeled after the Internet that only works within the organization).

In order to provide some level of separation between an organization's intranet and the Internet, *firewalls* have been employed. A firewall is simply a group of components that collectively form a barrier between two networks.

A number of terms specific to firewalls and networking are going to be used throughout this section, so let's introduce them all together.

Bastion host.

A general-purpose computer used to control access between the internal (private) network (intranet) and the Internet (or any other untrusted network). Typically, these are hosts running a flavor of the Unix operating system that has been customized in order to reduce its functionality to only what is necessary in order to support its functions. Many of the general-purpose features have been turned off, and in many cases, completely removed, in order to improve the security of the machine.

Router.

A special purpose computer for connecting networks together. Routers also handle certain functions, such as *routing* , or managing the traffic on the networks they connect.

Access Control List (ACL).

Many routers now have the ability to selectively perform their duties, based on a number of facts about a packet that comes to it. This includes things like origination address, destination address, destination service port, and so on. These can be employed to limit the sorts of packets that are allowed to come in and go out of a given network.

Demilitarized Zone (DMZ).

The DMZ is a critical part of a firewall: it is a network that is neither part of the untrusted network, nor part of the trusted network. But, this is a network that connects the untrusted to the trusted. The importance of a DMZ is tremendous: someone who breaks into your network from the Internet should have to get through several layers in order to successfully do so. Those layers are provided by various components within the DMZ.

Proxy.

This is the process of having one host act in behalf of another. A host that has the ability to fetch documents from the Internet might be configured as a *proxy server*, and host on the intranet might be configured to be *proxy clients*. In this situation, when a host on the intranet wishes to fetch the web page, for example, the browser will make a connection to the proxy server, and request the given URL. The proxy server will fetch the document, and return the result to the client. In this way, all hosts on the intranet are able to access resources on the Internet without having the ability to direct talk to the Internet.

Types of Firewalls

There are three basic types of firewalls, and we'll consider each of them.

Application Gateways

The first firewalls were application gateways, and are sometimes known as proxy gateways. These are made up of bastion hosts that run special software to act as a proxy server. This software runs at the *Application Layer* of our old friend the ISO/OSI Reference Model, hence the name. Clients behind the firewall must be *proxitized* (that is, must know how to use the proxy, and be configured to do so) in order to use Internet services. Traditionally, these have been the most secure, because they don't allow anything to pass by default, but need to have the programs written and turned on in order to begin passing traffic.

Figure 5: A sample application gateway

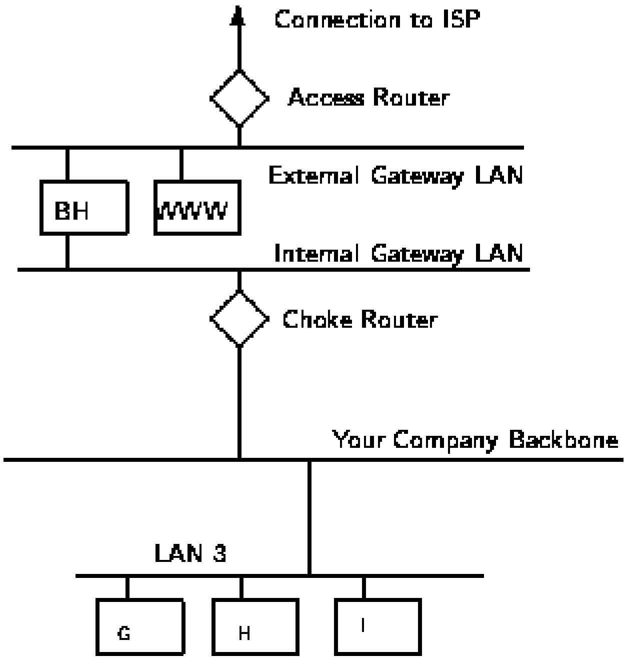

These are also typically the slowest, because more processes need to be started in order to have a request serviced.

Packet Filtering

Packet filtering is a technique whereby routers have *ACLs* (Access Control Lists) turned on. By default, a router will pass all traffic sent it, and will do so without any sort of restrictions. Employing ACLs is a method for enforcing your security policy with regard to what sorts of access you allow the outside world to have to your internal network, and vice versa.

There is less overhead in packet filtering than with an application gateway, because the feature of access control is performed at a lower ISO/OSI layer (typically, the transport or session layer). Due to the lower overhead and the fact that packet filtering is done with routers, which are specialized computers optimized for tasks related to networking, a packet filtering gateway is often much faster than its application layer cousins.

Because we're working at a lower level, supporting new applications either comes automatically, or is a simple matter of allowing a specific packet type to pass through the gateway. (Not that the *possibility* of something automatically makes it a good idea; opening things up this way might very well compromise your level of security below what your policy allows.)

There are problems with this method, though. Remember, TCP/IP has absolutely no means of guaranteeing that the source address is really what it claims to be. As a result, we have to use layers of packet filters in order to localize the traffic. We can't get all the way down to the actual host, but with two layers of packet filters, we can differentiate between a packet that came from the Internet and one that came from our internal network. We can identify which network the packet came from with certainty, but we can't get more specific than that.

Hybrid Systems

In an attempt to marry the security of the application layer gateways with the flexibility and speed of packet filtering, some vendors have created systems that use the principles of both.

Figure 6: A sample packet filtering gateway

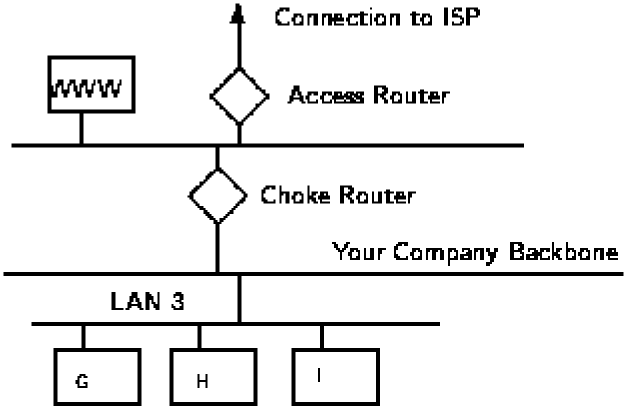

In some of these systems, new connections must be authenticated and approved at the application layer. Once this has been done, the remainder of the connection is passed down to the session layer, where packet filters watch the connection to ensure that only packets that are part of an ongoing (already authenticated and approved) conversation are being passed.

Other possibilities include using both packet filtering and application layer proxies. The benefits here include providing a measure of protection against your machines that provide services to the Internet (such as a public web server), as well as provide the security of an application layer gateway to the internal network. Additionally, using this method, an attacker, in order to get to services on the internal network, will have to break through the access router, the bastion host, and the choke router.

Lots of options are available, and it makes sense to spend some time with an expert, either in-house, or an experienced consultant who can take the time to understand your organization's security policy, and can design and build a firewall architecture that best implements that policy. Other issues like services required, convenience, and scalability might factor in to the final design.

The business of building firewalls is in the process of becoming a commodity market. Along with commodity markets come lots of folks who are looking for a way to make a buck without necessarily knowing what they're doing. Additionally, vendors compete with each other to try and claim the greatest security, the easiest to administer, and the least visible to end users. In order to try to quantify the potential security of firewalls, some organizations have taken to firewall certifications. The certification of a firewall means nothing more than the fact that it *can* be configured in such a way that it can pass a series of tests. Similarly, claims about meeting or exceeding U.S. Department of Defense ``Orange Book'' standards, C-2, B-1, and such all simply mean that an organization was able to configure a machine to pass a series of tests. This doesn't mean that it was loaded with the vendor's software at the time, or that the machine was even usable. In fact, one vendor has been claiming their operating system is ``C-2 Certified'' didn't make mention of the fact that their operating system only passed the C-2 tests without being connected to any sort of network devices.

Such gauges as market share, certification, and the like are no guarantees of security or quality. Taking a little bit of time to talk to some knowledgeable folks can go a long way in providing you a comfortable level of security between your private network and the big, bad Internet.

Additionally, it's important to note that many consultants these days have become much less the advocate of their clients, and more of an extension of the vendor. Ask any consultants you talk to about their vendor affiliations, certifications, and whatnot. Ask what difference it makes to them whether you choose one product over another, and vice versa. And then ask yourself if a consultant who is certified in technology XYZ is going to provide you with competing technology ABC, even if ABC best fits your needs.

Single Points of Failure

Many ``firewalls'' are sold as a single component: a bastion host, or some other black box that you plug your networks into and get a warm-fuzzy, feeling safe and secure. *The term ``firewall'' refers to a number of components that collectively provide the security of the system.* Any time there is only one component paying attention to what's going on between the internal and external networks, an attacker has only one thing to break (or fool!) in order to gain complete access to your internal networks.

Secure Network Devices

It's important to remember that the firewall only one entry point to your network. Modems, if you allow them to answer incoming calls, can provide an easy means for an attacker to sneak *around* (rather than *through*) your front door (or, firewall). Just as castles weren't built with moats only in the front, your network needs to be protected at all of its entry points.

Secure Modems; Dial-Back Systems

If modem access is to be provided, this should be guarded carefully. The *terminal server*, or network device that provides dial-up access to your network needs to be actively administered, and its logs need to be examined for strange behavior. Its password need to be strong -- not ones that can be guessed. Accounts that aren't actively used should be disabled. In short, it's the easiest way to get into your network from remote: guard it carefully.

There are some remote access systems that have the feature of a two-part procedure to establish a connection. The first part is the remote user dialing into the system, and providing the correct user-id and password. The system will then drop the connection, and call the authenticated user back at a known telephone number. Once the remote user's system answers that call, the connection is established, and the user is on the network. This works well for folks working at home, but can be problematic for users wishing to dial in from hotel rooms and such when on business trips.

Other possibilities include one-time password schemes, where the user enters his userid, and is presented with a ``challenge,'' a string of between six and eight numbers. He types this challenge into a small device that he carries with him that looks like a calculator. He then presses enter, and a ``response'' is displayed on the LCD screen.

The user types the response, and if all is correct, he login will proceed. These are useful devices for solving the problem of good passwords, without requiring dial-back access. However, these have their own problems, as they require the user to carry them, and they must be tracked, much like building and office keys.

No doubt many other schemes exist. Take a look at your options, and find out how what the vendors have to offer will help you *enforce your security policy effectively.*

Crypto-Capable Routers

A feature that is being built into some routers is the ability to session encryption between specified routers. Because traffic traveling across the Internet can be seen by people in the middle who have the resources (and time) to snoop around, these are advantageous for providing connectivity between two sites, such that there can be secure routes.

Virtual Private Networks

Given the ubiquity of the Internet, and the considerable expense in private leased lines, many organizations have been building *VPNs* (Virtual Private Networks). Traditionally, for an organization to provide connectivity between a main office and a satellite one, an expensive data line had to be leased in order to provide direct connectivity between the two offices. Now, a solution that is often more economical is to provide both offices connectivity to the Internet. Then, using the Internet as the medium, the two offices can communicate.

The danger in doing this, of course, is that there is no privacy on this channel, and it's difficult to provide the other office access to ``internal'' resources without providing those resources to everyone on the Internet.

VPNs provide the ability for two offices to communicate with each other in such a way that it looks like they're directly connected over a private leased line. The session between them, although going over the Internet, is private (because the link is encrypted), and the link is convenient, because each can see each others' internal resources without showing them off to the entire world.

A number of firewall vendors are including the ability to build VPNs in their offerings, either directly with their base product, or as an add-on. If you have need to connect several offices together, this might very well be the best way to do it.

Security is a very difficult topic. Everyone has a different idea of what ``security'' is, and what levels of risk are acceptable. The key for building a secure network is to *define what security means to your organization* . Once that has been defined, everything that goes on with the network can be evaluated with respect to that policy. Projects and systems can then be broken down into their components, and it becomes much simpler to decide whether what is proposed will conflict with your security policies and practices.

Many people pay great amounts of lip service to security, but do not want to be bothered with it when it gets in their way. It's important to build systems and networks in such a way that the user is not constantly reminded of the security system around him. Users who find security policies and systems too restrictive will find ways around them. It's important to get their feedback to understand what can be improved, and it's important to let them know *why* what's been done has been, the sorts of risks that are deemed unacceptable, and what has been done to minimize the organization's exposure to them.

Security is everybody's business, and only with everyone's cooperation, an intelligent policy, and consistent practices, will it be achievable

4.11 Business continuity and Disaster recovery

An important part of an Information Systems Security Programme is a comprehensively documented plan to ensure the continuation of the critical business operations of the Bank in the event of disruption. A disaster recovery plan outlines the roles and responsibilities under such situations and the system/procedures to be adopted for business continuity.

The disaster recovery is that part of the business resumption plan which ensures that the information and the information processing facilities are restored to their normal operating conditions as soon as possible after disruption. The disaster recovery plan should include the following:

a) Listing of business activities which are considered critical, preferably with priority rankings, including the time frame, adequate to meet business commitments;

b) Identification of the range of disasters that must be protected against;

c) Identification of the processing resources and locations, available to replace those supporting critical activities;

d) Identification of personnel to operate information processing resources at the disaster recovery site;

e) Identification of information to be backed up and the location for storage, as well as the requirement for the information to be saved for back-up purpose on a stated schedule and compliance therewith;

f) Information back-up systems being capable of locating and retrieving critical information in a timely fashion; and

g) Agreements entered into with the service providers/contractors/ vendors for priority resumption of services under the terms and conditions specified therefor therein.

The disaster recovery plan will have to be tested as frequently as necessary, as per the terms and conditions specified therefor in the agreement/ s with the service providers/contractors/vendors, to find problems, if any in the execution of the plan as also to keep the personnel trained therefor and in the operation of the back-up system. The record of each of these exercises should be documented, submitted to higher-ups and preserved. A periodic re-evaluation of the disaster recovery plan, to ascertain that it still serves the purpose, will have to be undertaken. A minimal frequency for both the testing of the disaster recovery plan and the re-evaluation exercise of its appropriateness/suitability will require to be specified by the Bank. The agreement/s with the service providers/contractors/vendors will have to include the terms and conditions for switch-over to the primary system on the resolution of the problems thereat. Further, if the implementation of the disaster recovery plan requires close co-ordination among various service providers/contractors/vendors, the terms and conditions, warranting close co-ordination & co-operation among them, will have to be specified in each relevant agreement, setting out the obligations to be met by each of the service providers/contractors/vendors including the penalties/punitive measures in case of non-compliance.

Important aspects of Business continuity Management

Objective is to counteract interruptions to business activities and to protect critical business processes from the effects of major failures or disasters

- A business continuity management process should be implemented to reduce the disruption caused by disasters and security failures which may result of for example, natural disasters, accidents, equipment failures and deliberate actions to an acceptable level through a combination of preventive and recovery controls.

- The consequences of disasters, security failures and loss of service should be analysed. Contingency plans should be developed and implemented to ensure that business processes can be restored with in the required time scales. Such plans should be maintained and practiced to become an integral part of all other management processes.

- Business continuity management should include controls to identify and reduce risks, limit the consequences of damaging incidents and ensure the timely resumption of essential operations.

Business continuity Management process

There should be a managed process in place for developing and maintaining business continuity throughout the organization. It should bring together the following key elements of business continuity management:

- Understanding the risks the organisation is facing in terms of their likelihood and their impact, including an identification and prioritization of critical business processes;
- Understanding the impact which interruptions are likely to have on the business and establishing the business objectives of information processing facilities
- Considering the purchase of suitable insurance which may form part of the business continuity process
- Formulating and documenting a business continuity plans in line with the agreed strategy
- Regular testing and updating of the plans and processes put in place

- Ensuring that the management of business continuity is incorporated in the organisation's processes and structure. Responsibility for coordinating the business continuity management process should be assigned at an appropriate level within the organization e.g. at the information security forum.

Business continuity and Impact analysis

Business continuity should begin by identifying events that can cause interruptions to business processes, e.g. equipment failure, flood and fire. This should be followed by a risk assessment to determine the impact of those interruptions, both in terms of damage scale and recovery period. Both of these activities should be carried out with full involvement from owners of business resources and processes. This assessment considers all business processes and is not limited to the information processing facilities.

Depending on the results of the risk assessment, a strategy plan should be developed to determine the overall approach to business continuity. Once this plan has been created, it should be endorsed by management.

Writing and implementing continuity plans

Plans should be developed to maintain or restore business operations in the required time scales following interruption to, or failure of, critical business processes. The business continuity planning process should consider the following

- Implementation of emergency procedures to allow recovery and restoration in required time scales. Particular attention needs to be given to the assessment of external business dependencies and the contracts in place.
- Documentation of agreed procedures and processes
- Appropriate education of staff in the agreed emergency procedures and processes including crisis management
- Testing and updating of the plans

The planning process should focus on the required business objectives, e.g. restoring of specific services to customers in an acceptable amount of time. The services and resources that will enable this to occur should be considered, including staffing, noninformation processing resources as well as fallback arrangements for information processing facilities.

Business continuity planning framework

A single framework of business continuity plans should be maintained to ensure that all plans are consistent, and to identify priorities for testing and maintenance. Each business continuity plan should specify clearly the conditions for its activation, as well as the individuals responsible for executing each component of the plan. When new requirements are identified, established emergency procedures, e.g. evacuation plans or any existing fallback arrangements should be amended as appropriate.

A business continuity planning framework should consider the following:

- The conditions for activating the plans which describe the process to be followed before each plan is activated
- Emergency procedures, which describe the actions to be taken following an incident, which jeopardizes business operations. This should include arrangements for public relations management and for effective liaison with appropriate public authorities e.g. police, fire service and local government
- Fall back procedures which describe the actions to be taken to move essential business activities or support services to alternative temporary locations, and to bring business processes back into operation in the required time scales.
- Resumption procedures which describe the actions to be taken to return to normal business operations
- A maintenance schedule which specifies how and when the plan will be tested and the process for maintaining the plan
- Awareness and education activities which are designed to create understanding of the business continuity processes and ensure that the processes continue to be effective.
- The responsibilities of the individuals, describing who is responsible for executing which component of the plan. Alternatives should be nominated as required.

Each plan should have a specific owner. Emergency procedures, manual fallback plans and resumption plans should be within the responsibility of the owners of the appropriate business resources or processes involved. Fall back arrangements for alternative technical services, such as information processing and communication facilities, should usually be the responsibility of the service providers

Testing the plans

Business continuity plans may fail on being tested, often because of incorrect assumptions, oversights or changes in equipment or personnel. They should therefore be tested regularly to ensure that they are up to date and effective. Such tests should also ensure that all members of the recovery team and other relevant staff are aware of plans.

The test schedule for business continuity plan(s) should indicate how and when each element of the plan should be tested. It is recommended to test the individual components of the plan(s) frequently. A variety of techniques should be used in order to provide assurance that the plan(s) will operate in real life. These should include:

- Table top testing of various scenarios
- Simulations particularly for training people in their post incident crisis management roles
- Technical recovery testing i.e. ensuring information systems can be restored effectively
- Testing recovery at an alternate site running business processes in parallel with recovery operations away from the main site
- Tests of supplier facilities and services ensuring externally provided services and products will meet the contracted commitment
- Complete rehearsals i.e. testing that the organization, personnel, equipment, facilities and processes can cope with interruptions

The techniques can be used by any organization and should reflect the nature of the specific recovery plan.

Maintaining and reassessing the plans

Business continuity plans should be maintained by regular reviews and updates to ensure their continuing effectiveness. Procedures should be included within the organisation's change management programme to ensure that business continuity matters are appropriately addressed.

Responsibility should be assigned for regular reviews of each business continuity plan; the identification of changes in business arrangements not yet reflected in the business continuity plans should be followed by an appropriate update of the plan. This formal change control process should ensure that the updated plans are distributed and reinforced by regular reviews of the complete plan.

Examples of situations that might necessitate updating plans include the acquisition of new equipment, or upgrading of operational systems and changes in:

Personnel
Addresses or telephone numbers
Business strategy
Location, facilities and resources
Legislation
Processes or new/withdrawn ones
Risk (operational and finacial)

4.12 Internet and email

E-mail – An email usage policy is a must. Several viruses, Trojans, and malware use email as the vehicle to propagate themselves throughout the Internet. A few of the more recent worms were Code Red, Nimda, and Gonner. These types of exploits prey on the unsuspecting user to double click on the attachment thereby infecting the machine and launching propagation throughout the entire network. This could cause several hours and/or days of downtime while remedial efforts are taken.

Internet – The World Wide Web was the greatest invention, but the worst nightmare from a security standpoint. The Internet is the pathway in which vulnerabilities are manifested. The black-hat community typically launches their 'zero day' and old exploits on the Internet via IRC chat rooms, through Instant Messengers, and free Internet email providers (hotmail, yahoo, etc.). Therefore, the Internet usage policy should restrict access to these types of sites and should clearly identify what, if any, personal use is authorized.

Moreover, software should be employed to filter out many of the forbidden sites that include pornographic, chat rooms, free web-based email services (hotmail, Yahoo, etc.), personals, etc. There are several Internet content filtering applications available that maintain a comprehensive database of forbidden URLs.

The Internet is the world's largest network *of networks* . When you want to access the resources offered by the Internet, you don't really connect to *the* Internet; you connect to a network that is eventually connected to the *Internet backbone* , a network of extremely fast (and incredibly overloaded!) network components. This is an important point: the Internet is a network of *networks* -- not a network of hosts.

A simple network can be constructed using the same protocols and such that the Internet uses without actually *connecting* it to anything else.

Figure 3: A Simple Local Area Network

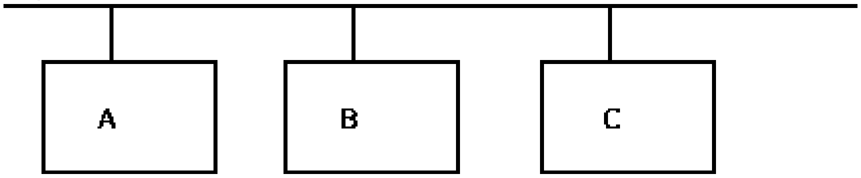

I might be allowed to put one of my hosts on one of my employer's networks. We have a number of networks, which are all connected together on a *backbone* , that is a network of our networks. Our backbone is then connected to other networks, one of which is to an *Internet Service Provider* (ISP) whose backbone is connected to other networks, one of which is the Internet backbone.

If you have a connection ``to the Internet'' through a local ISP, you are actually connecting your computer to one of their networks, which is connected to another, and so on. To use a service from my host, such as a web server, you would tell your web browser to connect to my host. Underlying services and protocols would send *packets* (small datagrams) with your query to your ISP's network, and then a network they're connected to, and so on, until it found a path to my employer's backbone, and to the exact network my host is on. My host would then respond appropriately, and the same would happen in reverse: packets would traverse all of the connections until they found their way back to your computer, and you were looking at my web page. This shows how the hosts on that network are provided connectivity to other hosts on the same LAN, within the same company, outside of the company, but in the same ISP *cloud* , and then from another ISP somewhere on the Internet.

Figure 4: A Wider View of Internet-connected Networks

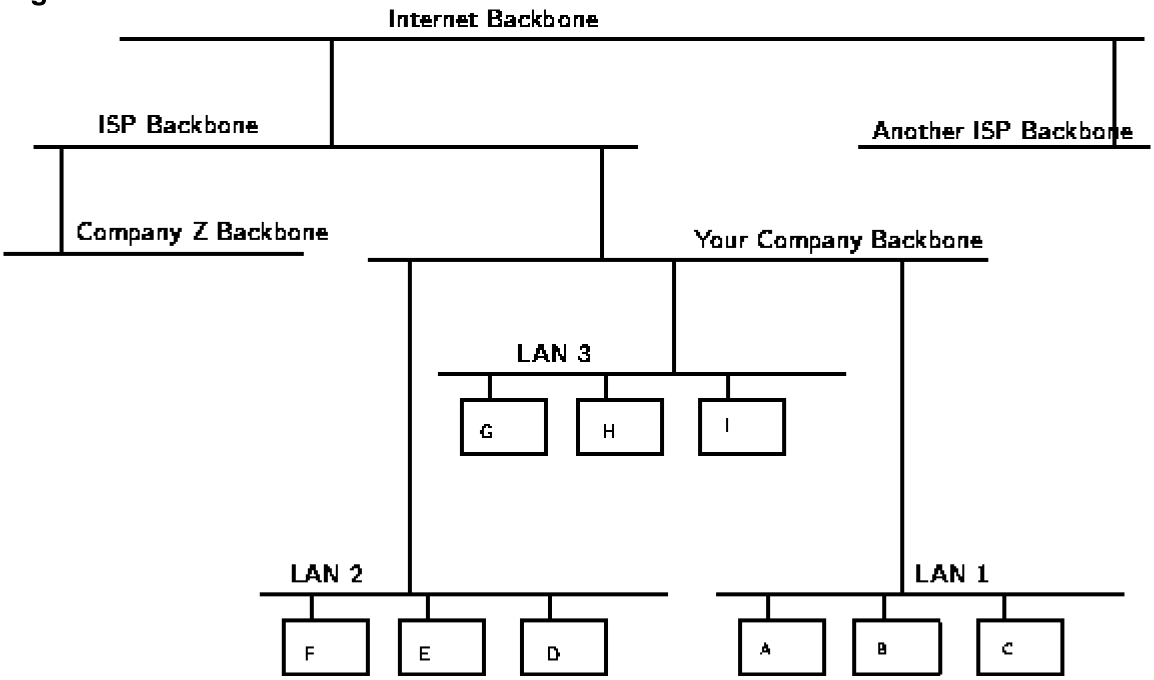

The Internet is made up of a wide variety of hosts, from supercomputers to personal computers, including every imaginable type of hardware and software. How do all of these computers understand each other and work together?

TCP/IP: The Language of the Internet

TCP/IP (Transport Control Protocol/Internet Protocol) is the ``language'' of the Internet. Anything that can learn to ``speak TCP/IP'' can play on the Internet. This is functionality that occurs at the Network (IP) and Transport (TCP) layers in the ISO/OSI Reference Model. Consequently, a host that has TCP/IP functionality (such as Unix, OS/2, MacOS, or Windows NT) can easily support applications (such as Netscape's Navigator) that uses the network.

Open Design

One of the most important features of TCP/IP isn't a technological one: The protocol is an ``open'' protocol, and anyone who wishes to implement it may do so freely. Engineers and scientists from all over the world participate in the *IETF* (Internet Engineering Task Force) working groups that design the protocols that make the Internet work. Their time is typically donated by their companies, and the result is work that benefits everyone.

IP

As noted, IP is a ``network layer'' protocol. This is the layer that allows the hosts to actually ``talk'' to each other. Such things as carrying datagrams, mapping the Internet address (such as 10.2.3.4) to a physical network address (such as 08:00:69:0a:ca:8f), and routing, which takes care of making sure that all of the devices that have Internet connectivity can find the way to each other.

Understanding IP

IP has a number of very important features which make it an extremely robust and flexible protocol. For our purposes, though, we're going to focus on the security of IP, or more specifically, the lack thereof.

Attacks Against IP

A number of attacks against IP are possible. Typically, these exploit the fact that IP does not perform a robust mechanism for *authentication* , which is proving that a packet came from where it claims it did. A packet simply claims to originate from a given address, and there isn't a way to be sure that the host that sent the packet is telling the truth. This isn't necessarily a weakness, *per se* , but it is an important point, because it means that the facility of host authentication has to be provided at a higher layer on the ISO/OSI Reference Model. Today, applications that require strong host authentication (such as cryptographic applications) do this at the application layer.

IP Spoofing.

This is where one host claims to have the IP address of another. Since many systems (such as router access control lists) define which packets may and which packets may not pass based on the sender's IP address, this is a useful technique to an attacker: he can send packets to a host, perhaps causing it to take some sort of action.

Additionally, some applications allow login based on the IP address of the person making the request (such as the Berkeley *r-commands*)[2]. These are both good examples how trusting untrustable layers can provide security that is -- at best -- weak.

IP Session Hijacking.

This is a relatively sophisticated attack, first described by Steve Bellovin. This is very dangerous, however, because there are now toolkits available in the underground community that allow otherwise unskilled bad-guy-wannabes to perpetrate this attack. IP Session Hijacking is an attack whereby a user's session is taken over, being in the control of the attacker. If the user was in the middle of email, the attacker is looking at the email, and then can execute any commands he wishes as the attacked user. The attacked user simply sees his session dropped, and may simply login again, perhaps not even noticing that the attacker is still logged in and doing things.

For the description of the attack, let's return to our large network of networks. In this attack, a user on host A is carrying on a session with host G. Perhaps this is a telnet session, where the user is reading his email, or using a Unix shell account from home. Somewhere in the network between A and B sits host H which is run by a naughty person. The naughty person on host H watches the traffic between A and G, and runs a tool which starts to impersonate A to G, and at the same time tells A to shut up, perhaps trying to convince it that G is no longer on the net (which might happen in the event of a crash, or major network outage). After a few seconds of this, if the attack is successful, naughty person has ``hijacked'' the session of our user. Anything that the user can do legitimately can now be done by the attacker, illegitimately. As far as G knows, nothing has happened.

This can be solved by replacing standard telnet-type applications with encrypted versions of the same thing. In this case, the attacker can still take over the session, but he'll see only ``gibberish'' because the session is encrypted. The attacker will not have the needed cryptographic key(s) to decrypt the data stream from G, and will, therefore, be unable to do anything with the session.

TCP

TCP is a transport-layer protocol. It needs to sit on top of a network-layer protocol, and was designed to ride atop IP. (Just as IP was designed to carry, among other things, TCP packets.) Because TCP and IP were designed together and wherever you have one, you typically have the other, the entire suite of Internet protocols are known collectively as ``TCP/IP.'' TCP itself has a number of important features that we'll cover briefly.

Guaranteed Packet Delivery

Probably the most important is guaranteed packet delivery. Host A sending packets to host B expects to get acknowledgments back for each packet. If B does not send an acknowledgment within a specified amount of time, A will resend the packet.

Applications on host B will expect a data stream from a TCP session to be complete, and in order. As noted, if a packet is missing, it will be resent by A, and if packets arrive out of order, B will arrange them in proper order before passing the data to the requesting application.

This is suited well toward a number of applications, such as a telnet session. A user wants to be sure every keystroke is received by the remote host, and that it gets every packet sent back, even if this means occasional slight delays in responsiveness while a lost packet is resent, or while out-of-order packets are rearranged.

It is not suited well toward other applications, such as streaming audio or video, however. In these, it doesn't really matter if a packet is lost (a lost packet in a stream of 100 won't be distinguishable) but it *does* matter if they arrive late (i.e., because of a host resending a packet presumed lost), since the data stream will be paused while the lost packet is being resent. Once the lost packet is received, it will be put in the proper slot in the data stream, and then passed up to the application.

UDP

UDP (User Datagram Protocol) is a simple transport-layer protocol. It does not provide the same features as TCP, and is thus considered ``unreliable.'' Again, although this is unsuitable for some applications, it does have much more applicability in other applications than the more reliable and robust TCP.

Lower Overhead than TCP

One of the things that makes UDP nice is its simplicity. Because it doesn't need to keep track of the sequence of packets, whether they ever made it to their destination, etc., it has lower overhead than TCP. This is another reason why it's more suited to streaming-data applications: there's less screwing around that needs to be done with making sure all the packets are there, in the right order, and that sort of thing.

4.13 Backup restoration

Some of those perpetrate attacks are simply twisted jerks who like to delete things. In these cases, the impact on your computing capability -- and consequently your business -- can be nothing less than if a fire or other disaster caused your computing equipment to be completely destroyed.

Where Do They Come From?

How, though, does an attacker gain access to your equipment? *Through any connection that you have to the outside world.* This includes Internet connections, dial-up modems, and even physical access. (How do you know that one of the temps that you've brought in to help with the data entry isn't really a system cracker looking for passwords, data phone numbers, vulnerabilities and anything else that can get him access to your equipment?)

In order to be able to adequately address security, all possible avenues of entry must be identified and evaluated. The security of that entry point must be consistent with your stated policy on acceptable risk levels.

Hope you have backups

This isn't just a good idea from a security point of view. Operational requirements should dictate the backup policy, and this should be closely coordinated with a disaster recovery plan, such that if an airplane crashes into your building one night, you'll be able to carry on your business from another location. Similarly, these can be useful in recovering your data in the event of an electronic disaster: a hardware failure, or a breakin that changes or otherwise damages your data.

Don't put data where it doesn't need to be

Although this *should* go without saying, this doesn't occur to lots of folks. As a result, information that doesn't need to be accessible from the outside world sometimes is, and this can needlessly increase the severity of a break-in dramatically.

Avoid systems with single points of failure

Any security system that can be broken by breaking through any one component isn't really very strong. In security, a degree of redundancy is good, and can help you protect your organization from a minor security breach becoming a catastrophe.

Stay current with relevant operating system patches

Be sure that someone who knows what you've got is watching the vendors' security advisories. Exploiting old bugs is still one of the most common (and most effective!) means of breaking into systems.

Watch for relevant security advisories

In addition to watching what the vendors are saying, keep a close watch on groups like CERT and CIAC. Make sure that at least one person (preferably more) is subscribed to these mailing lists

Have someone on staff be familiar with security practices

Having at least one person who is charged with keeping abreast of security developments is a good idea. This need not be a technical wizard, but could be someone who is simply able to read advisories issued by various incident response teams, and keep track of various problems that arise. Such a person would then be a wise one to consult with on security related issues, as he'll be the one who knows if web server software version such-and-such has any known problems, etc.

This person should also know the ``dos'' and ``don'ts'' of security, from reading such things as the ``Site Security Handbook.''