# Chapter 14: Policies and Procedures

## *Policies and Procedures*

For most organizations, network and system security policies and procedures serve the purpose of ensuring information security. They achieve this by defining what constitutes information security, why it is important, and how to maintain it. In addition, the policies and procedures define the acceptable levels of information security. Before you can do so, however, you must first put in place a process that enables you to determine what is an adequate level of security for any given organization.

You should recall from the discussion in Chapter 1 that the elements of information security include confidentiality, integrity, availability, authentication, and access control. All five elements need to be addressed by whatever policies and procedures are implemented to address information security. In general terms, security policies are the set of rules and procedures that regulate how an organization manages, uses, protects, and distributes all information that directly or indirectly pertains to that organization.

### Policies Versus Procedures

Policies should always be developed before procedures. The development of procedures should flow from the policies. Policies should be concerned with what assets to protect and why they need to be protected. They are generally broad in their scope and are designed to set the tone and direction. In general, they can be thought of as the documents that spell out the *what* and *why* of information security for an organization. Procedures, on the other hand, must be much more precise and detailed. Procedures should be concerned with the specific measures necessary to protect the organization's assets. They can be thought of as the documents that spell the *who, when,* and *how* of information security within an organization.

### Information Security Policy Objectives

There are various reasons for an organization to develop network and system security policies and procedures. Some are obvious, while others are not so obvious. Some reasons concern the direct benefit that an organization gains from having policies and procedures, such as preventing or detecting fraud or deterring hackers. Other benefits are indirect in that the policies protect the organization from potential liability or save it from possible embarrassment. Below I have listed some of the objectives generally associated with network security policies.

- *Managing risk:* The primary goal of any policy concerning network and system security is to manage risk. It is almost impossible to completely secure an organization's information assets. As a result, an organization needs to identify the risks that its faces and develop measures to minimize the impact of those risks.
- *Ensuring business continuity:* The ongoing operation of the organization should be a fundamental goal of the policies developed by any organization. It is interesting to note how many organizations' policies tend to spell out what *cannot* be done in great detail but do a very poor job of addressing what *must* be done to ensure the operation of the organization. Organizational policies and procedures should ensure business resumption by outlining the appropriate actions necessary in response to an incident or disaster.
- *Defining responsibilities, expectations, and acceptable behaviors:* For any policy or procedure to be effective, those individuals subject to the policy or procedure must understand what is required of them to comply. Compliance to a policy cannot be achieved without reaching an understanding of what constitutes compliance. In addition, employees need to understand their responsibilities and how their

responsibilities may vary depending on the circumstances.

- *Discharging fiduciary duty and complying with any regulatory requirements:* Most organizations are subject to rules or regulations governing the responsibility of the corporate officers and regulating the operation of the organization. If a company is publicly traded, the corporate officers have a fiduciary duty to ensure the financial soundness of the organization. If they fail in that duty they can be held personally liable for the losses incurred. Most every organization is required to adhere to certain standards when it comes to accounting records and bookkeeping. Many organizations are also subject to federal, state, or local regulations that require certain measures be taken to protect the assets of the organization. Many organizations are subject to rules and regulations regarding the protection and disclosure of information pertaining to employees and customers. This is certainly true in the financial and health sectors. For many organizations, the absence of proper policies and procedures is considered automatic noncompliance.
- *Protecting the organization from liability:* The policies and procedures developed by an organization are often required to protect it from liability. In some cases, the existence of the policies and procedures are essential to demonstrate that an organization did not approve of an end user's actions or that an employee was or was not acting with the authorization of the organization.
- *Ensuring information integrity and confidentiality:* A key component of information security is protecting an organization's information assets. Ensuring the integrity and confidentiality of an organization's information is fundamental to that goal. Without information integrity, an organization cannot make sound business decisions. Without information confidentiality, an organization will lose its competitive edge through the loss of proprietary information regarding products, customers, and even partners and suppliers.

## Developing Security Policies

For an organization's information security policies and procedures to achieve the stated objectives, it is essential that certain elements be included in the policies and procedures. These elements can be thought of as key measures for the success for an organization's policy and procedures. The elements are the stepping stones in the development process. They are listed as follows:

- Identifying the organization's assets;
- Defining the risks:
- Defining how information assets are to be managed;
- Defining how information assets are to be accessed and what process will be used for authentication;
- Defining clearly and in detail what does and does not constitute appropriate use of company owned electronic media and services;
- Clearly defining what kind of information may be accessed and distributed and by what means;
- Defining what controls are to be put in place;
- Notifying users of monitoring and auditing procedures, information disclosure, and consequences for noncompliance;
- Identifying those responsible for security enforcement and how policies and procedures will be enforced;
- Developing steps to be taken in the event of noncompliance with policy, a security breach, or a disaster.

The first step is to determine responsibility for information security policy development. Too often, the IT unit is given sole responsibility for this task. However, if the policies and procedures are to be comprehensive, it will require the active participation of all business units. Development of information security policies must be a collaborative effort between the IT unit and the other business units within an organization. Any policy or procedure implemented without the active participation and "buy-in" of other business units faces an uphill battle.

The most critical factor in the success or failure of any information security policy is support from senior management: The policy developers must be empowered by senior management with the authority to implement the measures necessary to protect the organization's information assets. I cannot stress this fact strongly enough. Without the support of senior management any policies or procedures implemented are doomed to fail.

I have seen teams attempt to implement procedures only to see their effort undone by senior management's failure to back them when they encountered resistance. It was a no-win situation for the team. They were charged with the job of developing and implementing security procedures but not empowered with the authority necessary to succeed. As a result, they were ultimately seen as the bad guys and targeted by everyone's wrath. Basically, they had been set up to fail. If the group developing policies does not have the active support of senior management it is best not to even attempt the task.

Senior management needs to do more than just support the development and implementation of policies and procedures. Senior management needs to support a culture of information security within an organization. There needs to be recognition of the need for information security within every organization. Unfortunately, in most companies, information security is too often looked upon as something that can be dealt with after everything else. It is not recognized as a required core competency of the company. This fallacious mind-set can put an organization at risk.

Consider the following example: A student in one of my classes recounted a story indicative of the value that most companies place on information security. The student worked for a large software company that marketed a well-known database. During a cyclical downturn in business, the company went through a round of what was euphemistically called "rightsizing."

While most business units experienced moderate cuts in personnel, the information security and the business resumption planning groups were devastated. Essentially, both units were dissolved, and all personnel were laid off. Obviously, the company did not see information security and business resumption as a critical business activity.

As another example, at a company where I once worked, I submitted to senior management a recommendation that the company develop a policy to address "pretext calling." Pretext calling is a widespread practice used by information brokers to gain information on individuals from unsuspecting companies. Generally, an information broker poses as someone or some entity that is related to or associated with the individual with whom the targeted company does business. The targeted company could be a hospital, a financial institution, an insurance company, or even a school or government agency. The information broker usually gets a little bit of information from each contact. The information gathered is cumulative. With each contact the information broker gets more information, which in turn can be used to gain even more. Many companies are being hit by pretext calling. Even though the information broker lies and misrepresents himself or herself to the targeted company, this practice is not illegal. Companies are unwittingly giving out information on their employees, customers, and clients. It is not only bad for the customer, but it is bad for business. In addition, a company could find itself liable for how that information is used. It certainly would not instill customer confidence to know that a company was giving out customer information to anyone who calls. For that reason, I recommended that a policy and procedure be developed to address pretext calling.

Specifically, my recommendation was that the company should develop a general information privacy policy. Part of the implementation of that policy would include a training program to educate our staff on how to identify pretext calls. I argued that it would provide our company with a competitive advantage in that we could state to our customers that their information was safer with us than with our competitors. In addition, it would protect the company from possible liability. Finally, it would provide the company with a response to customers who contacted us with requests for information on how we handled this type of occurrence. Senior management thought it was a good idea but not a high priority, and that is where it ended. No one wanted to invest the time to develop the policy. Without the active support of senior management, it would have been impossible to develop a policy and attempt to impose it on the other business units.

Another common occurrence is for organizations to develop policies for no other reason than to say that the policies exist. The policies are really only for show. Once a year, when regulators or auditors are on site, the company can point to the policy manuals gathering dust in the corner and proudly proclaim that they have all the required policies covered. Of course, the fact that no one even knows what the policies state or whether or not the company is in compliance with the policies is not considered.

An essential part of security policy development is the risk assessment process. It is important
to go through a risk assessment process to determine what you want to protect, why you want to protect it, and from what you need to protect it. As described in Chapter 1, the steps associated with risk assessment include the following.

1. Identifying and prioritizing assets;
2. Identifying vulnerabilities;
3. Identifying threats and their probabilities;
4. Identifying countermeasures;
5. Developing a cost-benefit analysis;
6. Developing security policies.

The first step is to identify and prioritize assets and systems and then identify the vulnerabilities associated with those assets. When assessing vulnerabilities and the risks associated with them, it is important to weed out the possible threats from the probable ones. The process should be one of determining what threats are most likely and developing policies that address those threats and issues.

It is very important that the policies and procedures implemented within any organization should be real world-based. In other words, the policies and procedures should exist for the purpose of enhancing a preexisting process or function. As such, they should take into account the constraints of the real world and not try to achieve the apex of security. For example, it would be overkill to require all e-mail to be encrypted. You should not require passwords to be changed every week or require them to be 15 alphanumeric characters in length. While it might be very secure, it would not be logical to implement a hand scanner for biometric identification in an environment, such as a "clean room," where technicians wear special suits, including gloves.

As a rule, security policies and procedures that interfere with the operation of an organization are of little value. Those types of measures are usually ignored or circumvented by company personnel, so they tend to create security holes rather than plug them. If you make a process too arduous or annoying, people will ignore it. If you make the process of gaining access to a room too difficult, people will prop open the door. If you make passwords too hard to remember, people will write them down. All security measures (not just security policies), whenever possible, should complement the operational and business needs of an organization. The steps involved in information security policy implementation are fairly straightforward:

1. Developing a written security policies and procedures manual;
2. Developing an end user awareness and education program;
3. Developing a process for policy enforcement and procedure implementation;
4. Developing a process for the periodic review and updating of policies and procedures.

## Policy and Procedure Manuals

For a security policy to be practical, it must be documented. The plan must also be made available as a reference to all those subject to the policy. The policy and procedure manuals need to be kept current and updated with any necessary changes. Modifications to systems, personnel, business priorities, and other environmental factors must be reflected in the plan. That means regular and frequent reviews of the policy.

## Policy Format

There are many different ways in which one can format the policies. The type of format is relatively unimportant as long as the policy is understandable and achieves the desired results.

The most important thing is that policies are formalized and documented in some way. A policy should include, at a minimum, the following elements.

- *Policy statement:* This section should state the general policy, what the policy says, and what it entails. This section can be as short as a single sentence or as long as a page. If it goes beyond a page, perhaps you are attempting to cover in a single policy issues that should be covered by more than one policy.
- *Purpose:* This section should state why the policy is needed. Examples of the purpose for a policy could include something to the effect that the policy is to protect the company or employees, ensure the continued operation of the organization, or protect the financial health of the company.
- *Scope:* This section should cover how far the policy extends. The scope should spell out the circumstances under which the policy applies. It can also include the time frame, specific hardware or software, and/or events under which the policy is effective.
- *Compliance with policy:* This section should include a detailed explanation of what does and does not constitute compliance with the policy. The section can include examples, but be careful to word it in such a way that it allows you to include instances that may not be listed in your examples. The section should include wording to the effect "*examples include, but are not limited to ….*" Being too specific in detail may make the definition too narrow.
- *Penalties/consequences:* This section should explain the consequences for noncompliance with the policy. Specific punishments associated with noncompliance should be listed. If the consequences for noncompliance can include termination, then it should be clearly spelled out in this section of the policy. This section serves as a warning to employees and can protect an organization in the event that it finds itself in court as a result of terminating an employee for non-compliance with a policy. The fact that the organization had clearly warned all employees of the consequences can diminish any argument that an employee may have for termination without cause.

## Policy Awareness and Education

A policy is of no value if no one knows what it states. End users and personnel must understand management's expectations and their responsibilities in regard to complying with an organization's policies. End users and employees must also understand the consequences for noncompliance. This aspect is very important for protecting the organization if litigation results from noncompliance.

The existence of a policy may be required to take punitive action against end users or employees who have acted in an unacceptable manner. Organizations that don't have a policy clearly defining unacceptable behavior may have no recourse.

Having a policy in place that prohibits certain types of behavior can also save an organization from liability for the actions of its end users or employees. The absence of a formal policy and an awareness process may make it difficult to hold an employee accountable in the event some inappropriate behavior on the part of the employee is discovered. With a written policy, an organization can demonstrate that any derogatory actions taken by an end user or employee were not in compliance with accepted behavior and were therefore not condoned by the organization.

Organizations should consider obtaining written acknowledgment from end users and employees stating that they have read and understand the organization's information security policy. This could be done as part of the general orientation for newly hired personnel or as part of the registration of new end users.

## Policy Enforcement

Compliance with policies needs to be enforced. The only way to ensure compliance is through monitoring and auditing. Those responsible for enforcing the IT security policies must have the support of senior management. If an organization's IT security policy is to be successful, it also needs the support of all business units within the organization.

### *Security Policy Suggestions*

Remember that the major emphasis of all policies and procedures is to prevent "bad things" from happening. It doesn't matter whether the bad thing is a mistake, disaster, or misdeed. Well-designed policies and procedures are flexible enough to address most "probable" threats.

That is why risk analysis is such an import part of the process.

Policies and procedures should also assume that the preventative measures will occasionally fail. As a result, they should include steps to detect "bad things." It is particularly important that the procedures spell out in detail what steps are to be taken in the event that all other measures have failed to prevent some "bad thing" from occurring. In other words, it should detail how the organization responds to an incident.

When developing procedures one needs to address the basic elements of network and system security covered in Chapter 2. They are listed as follows.

- Identification;
- Authentication;
- Access control (authorization);
- Availability;
- Confidentiality (secrecy);
- Integrity (accuracy);
- Accountability.

At the same time, you need to incorporate all of the various elements of security into all aspects of the operation of an organization and to address all probabilities. This includes procedures to address physical security and natural disasters as well as hardware and software security. You also need to address media controls and communication security. Most importantly, you need to address the human variable in your procedures in an effort to minimize temptation and stupidity and ensure compliance.

The framework required to adequately address the needs of a particular organization will largely depend on the type of organization. Large corporations require extensive policies that cover all the possibilities, while most small organizations, which may use technology to a more limited extent-or at least have less of it-will require a much less extensive set of policies.

Do not use a pound when an ounce will do the job. Overly complicated or detailed policies tend to create problems and are often ignored. Policies should be simple to understand and remember. The level of detail for each organization will vary, but the following sections provide some basic suggestions.

## Use of Company-Owned Electronic Media and Services

With the advent of new technologies, organizations are finding themselves relying increasingly on electronic modes of communication and information storage. Most employees in an organization have access to one or more forms of electronic media or service. They include but are not confined to the following:

- Computers (PCs, workstations, minicomputers, and mainframes);
- E-mail;
- Telephones and voice mail;
- Fax machines;
- LANs, intranets, and the Web.

Every organization that uses electronic media and services should have a policy that clearly defines the acceptable use of these media and services as company property. The policies should not only exist to protect the organization but also to protect the employees of the organization. The policy should specify the acceptable personal use of company-owned IT facilities and services. The policy should also cover when it is necessary to obtain management's permission and the process to do so. This policy should cover all technologies that could be exploited to receive and distribute information. Company systems and networks should not be used to generate or distribute material that is illegal or immoral or that contravenes the principles of the corporation. Such a policy ensures that appropriate measures are enacted to protect company assets and to educate employees of their responsibilities.

Often, an organization feels that developing a policy on the use of e-mail is all that is required. If the policy is to be truly effective, it must encompass more than just e-mail.

When it comes to developing such a policy, organizations can run the entire gamut from very liberal in their approach and loosely defined to very narrow definitions of what is acceptable use of company property with severe limitations on personal use. Each organization is different and the approach, and the philosophy that is brought to the task of developing a policy will vary greatly from company to company.

## What Does the Policy Cover?

It is very important that employees or end users understand what technologies or kinds of technologies the policy covers. Accordingly, organizations need to explain what the company's electronic media and services are and what they entail. It is to their benefit and the benefit of their employees that they understand that the policy covers more than just e-mail.

## Whose Property Is it?

A policy should state in clear terms that the electronic media and services are company property, not the employee's personal property. For example, employees very often become possessive about their PCs. They feel as if the PCs are their personal property and that no one has the right to access their PCs without first obtaining their (the employees') permission. It should be made clear that at any time, authorized personnel may review files on companyowned PCs, e-mail, or voice mail. This is not spying. Companies are, at times, obligated to perform such reviews to determine, among other things, whether there has been a breach of security, violation of company policy, or misuse of any company-owned media or services. Employees should be told that the company reserves the right to perform these reviews without prior notification of the employees. Make it clear to the employees that if they don't want the company to see something, they should not store it on company owned property.

## What Is Acceptable Use?

An organization has to determine for itself whether it will allow electronic media and services to be used for non-company-related purposes. The most reasonable approach is to allow limited, occasional use for personal, nonbusiness purposes (as is the case with personal phone calls). It is also important that policies be consistent with one another. It does not make sense for a policy to forbid the use of company e-mail for personal reasons while completely ignoring personal phone calls, voicemail, and faxes. Whatever an organization decides, the decision needs to be relayed to the employees in clear terms that spell out what the consequences are for violating the policy.

An organization should also protect itself by stating in writing that it is prohibited to use any of the company's electronic services for any purposes that violate state or federal laws. This includes requiring compliance with all copyright laws. If the company develops software, then the policy should also cover patents, trademarks, and intellectual property.

In addition, a policy should prohibit the use of company-owned electronic services to transmit, receive, or store information or data of a harassing or discriminatory nature or that is

derogatory to any group or individual. The policy should also prohibit any employee from using the company's electronic services to transmit, receive, or store information or data that is obscene or pornographic or that is defamatory or threatening in nature. This not only protects the organization; it protects the employees as well.

## Hacking

The policy should also prohibit attempts by employees or end users to "hack" other systems. It should be made clear that attempts to hack or access information without authorization will not be tolerated by an organization, and there should be severe consequences for doing so. This policy should not only apply to attempts to hacking company-owned systems; it should also apply to the hacking of outside systems using company-owned or -leased systems or services.

In addition, the policy should define the employees' responsibility to ensure that their logins and passwords remain confidential and the steps that they are required to take if they suspect that their passwords have been compromised. It should be made clear that these steps are not optional or suggested but are a required part of their job function and that failure to comply with the policy can result in adverse consequences.

## Unauthorized Software

Many organizations employ a cookie cutter approach to deploying desktop systems. Everyone

gets the same image of a specific suite of authorized software. While this can be aggravating to the end users, it is a sound management practice. At the very least, this approach reduces the costs associated with the installation of desktop systems. This is particularly true when employing a package such as Microsoft's System Management Server (SMS), which essentially pushes an image onto the desktop from the server. This approach can also reduce an organization's support costs by reducing the number of applications that the help desk supports.

In general, it is a good security practice to have a policy that prohibits end users from installing software on their desktop systems without authorization from the IT group. This can prevent malicious programs from being introduced to the network. When it comes to installing software on servers, there should not only be a policy in place that bars such activity, but the access control mechanisms should be in place to prevent such activity.

In many environments, it may be prudent to implement measures that prevent end users from installing software on their systems or in any way altering their desktop configuration. For example, Windows NT desktop systems can be installed with the local configuration capability disabled. Some programs designed to secure the desktop, such as Full Armor, Fortres 101, and Fool Proof, can be installed with Windows 3.X, 95, and 98. These systems provide some level of protection, but they can be circumvented and, in some cases, may actually pose risks.

## E-Mail

Employees should be made aware of the fact that e-mail is not a secure media. There is no guarantee that e-mail will remain private. They should also be made aware of the fact that email

transmitted on the Internet is particularly vulnerable to interception and disclosure. As such, information of an extremely sensitive or confidential nature should not be transmitted on the Internet unless the message is encrypted.

Every organization should reserve the right to review and disclose any employee's e-mail received or transmitted on or from company-owned electronic media or services. It should be made clear to every employee that this review and disclosure can be done without obtaining the employee's prior consent. This is not "big brother," it is common sense. A company has the right to protect itself. There have been a number of cases in the news media where the improper activities of an employee have landed an employer in court. The improper activities were later found to be detailed in the company e-mail. As a result, the company could be found libel for the employee's activities.

The results of a *Computerworld* survey regarding e-mail monitoring published in the magazine's October 1999 issue stated that 31% of the survey respondents had installed software that allowed for the active monitoring of e-mail and that another 21% were planning on installing software with that capability. Products such as Mailwall from Omniquad, MIMEsweeper from Content Technologies, and 2MA Messaging Manager from Re-Soft provide administrators with the ability to scan end users' e-mail for key words. These programs can scan both the e-mail subject and body for questionable, obscene, abusive, or illicit content.

## Identification

Any policy covering the acceptable use of company-owned electronic media and services should also deal with the issues of identity authentication and impersonation. Employees should be cautioned about relying on the stated identity of the sender of e-mail or any other type of transmission. E-mail messages in particular can easily be forged. Any policy should also prohibit employees from any attempt to hide their identity or to falsely represent themselves or attempt to represent themselves as someone else, when transmitting, receiving, or storing e-mail or other electronic communications.

## Communicate the Consequences

Finally, an organization should clearly define the consequences to any employee who knowingly violates the policy or allows another employee to violate the policy, so that there is

no possibility of misunderstanding.

For an organization to develop this type of policy, senior management, IT, and human resources need to work together, Developing a comprehensive policy governing the use of company-owned electronic media and services can protect an organization and save it from legal troubles down the road. For that policy to be effective, it must encompass more than just e-mail. An organization has to make the policy broad enough to incorporate all of the technologies that it is presently using. At the same time, individual elements of the policy have to be defined narrowly enough to make them meaningful and understandable. In addition, the policy needs to be periodically reviewed to ensure that it includes newly implemented technologies.

## Information Privacy

It is important that active steps be taken by all employees to ensure that information privacy is maintained. Corporate information pertaining to customers, employees, and company projects and products should be reviewed to determine their level of sensitivity. This is important from both a business and regulatory perspective. Disclosure of sensitive information can help competitors and scare away customers. In addition, a corporation may also be subject to regulatory requirements governing the disclosure of information. Web sites catering to children are subject to the Children's Online Privacy Protection Act, which is enforced by the Federal Trade Commission (FTC). Japan and most of the European nations have much stricter regulations than the United States governing the disclosure and sharing of personnel information by companies. As a result, a general policy is recommended. The policy should outline the requirements governing the actions of the organization for information privacy. Finally, the policy of organizations that have employees who frequently present at conferences or who are offered speaking engagements should cover what can and cannot be disclosed by the employee in his or her presentation. The policy can go so far as to include some type of review process by the management of the material being presented. This is to ensure that no sensitive proprietary or customer information is inadvertently disclosed.

## Information and Data Management

Depending on the environment in which you operate, you may want to consider classifying and prioritizing information by its level of importance or sensitivity. Correspondingly, the nature of the data will dictate the measures necessary to protect it. Determination of access levels should also be dictated by the sensitivity of the information or data.

Any policy should also define where information should reside and how it is to be moved, transported, or transmitted. The level of importance and sensitivity should be taken into account when these definitions are developed. For example, an organization may want to forbid information of critical importance from being copied to removable media such as floppies or tapes.

Information and data are valuable corporate assets and must be protected. Data can be defined as raw information, or information can be defined as meaningful data that has been organized in a coherent manner that allows for the reliable retrieval of data elements. One of the key components for the protection of information is to assign ownership. A policy on ownership should outline the responsibilities of the information guardian and the relationship with the custodian of the data.

Policies are also necessary to address the proactive management of information and data. Policies must address the availability of the data and ensure that the appropriate controls are in place and utilized. Development of these policies should entail the analysis of the risks and the establishment of appropriate classification and authorization standards for the data.

Information and data integrity is not just concerned with protecting information content. Integrity must also address the accuracy of the data elements. A policy concerning data integrity should identify the requirements for secure data storage and mechanisms for the backup of data, and the requirements for the procedures to preserve and test the accuracy of the data. In the appropriate environment, data integrity also includes data entry standards to ensure that information is entered in a consistent and uniform format.

To ensure data integrity, a policy should be enacted governing proper procedures to protect against the potential threat from computer viruses. The policy should cover requirements for virus scans and copying files from outside sources to company-owned systems.

Information and data management policies should also state that all files that reside on company-owned devices or media, such as PCs, removable disks, and tapes, are the property of the company. As such, the policy should prohibit employees from removing company information from the premises without authorization. This policy, while difficult to enforce, may be a useful legality to have in place. In addition, as a precaution, a company should reserve the right to examine, access, use, and disclose any or all information or data, transmitted, received, or stored on any electronic media, device, or service owned or paid for by the company.

## Systems Administration

One of the biggest challenges in devising proper security procedures is determining how to deal with the control and monitoring of the administrators of the organization's various systems. For example, many organizations operate in an environment where an individual or individuals have access to or responsibility for all aspects of system administration. The organization may have a small IT unit where using delineation of responsibility and segregation of duties as a control procedure is not practical. How do you segregate duties when there is only one person in the department?

However, whenever possible, segregation of duties should be implemented. The individual or individuals responsible for the day-to-day administration should not also be the individual or individuals responsible for creating new accounts. In addition, the individual or individuals who create new accounts should not be responsible for determining the level of access given to those accounts. All new accounts should be reviewed by an individual not responsible for creating accounts. If possible, a distinction should be made between system administration and security adminisration. System administration functions should be audited at least annually.

All system changes and daily jobs performed by administrators and operators should be recorded in a log or schedule and should be reviewed daily. All system backups should be recorded and logged and the logs reviewed and retained. Backups should also be tested periodically, at least weekly. All security access changes should be documented, reviewed, and filed. A policy will also stipulate the records retention schedule and destruction of logs, schedules, and other documentation.

In addition, systems should be classified according to their confidentiality and criticality to the operation of the organization to determine appropriate security measures. System classification is also required for disaster recovery planning.

System auditing and validation should be addressed in some manner through policies. They can either be incorporated into existing polices or be in a separate policy. Chapter 15 discusses auditing in more detail.

## Remote Network Access

Many organizations have requirements for remote network access. Sales staff, field engineers,

and even delivery personnel and drivers often require access to an organization's network. In addition, with the growth in telecommuting, many employees are now working from home, rather than coming into the office. As a result, more employees require access to the company's systems from outside the corporate network. Any remote access to the corporate network should be tightly controlled and subject to stringent security measures. A policy for remote access should address issues associated with authentication and access control. At a minimum, the policy should require any connection to utilize some kind of secure ID procedure. Refer to the discussion in Chapter 7 regarding modems for more detail.

Another consideration is third-party access to the corporate network. Many organizations have vendors, partners, customers, or joint ventures that require access to the corporate network. Policies need to be developed to ensure that proper controls are implemented, maintained, and monitored for all third-party access to an organization's network.

### Security of Telecommunications

Related to remote access are the issues associated with secure communications. Different modes of telecommunications are subject to different potential for disclosure. A policy should detail what measures should be taken when using each of the different modes of electronic communications, based upon the sensitivity of the information. Refer to Chapter 9 for a more detailed discussion of the issues associated with the different modes of telecommunications.

### Physical Security

Physical access to IT facilities should be restricted to only those authorized personnel who need access to perform their job functions. A policy should define who are the appropriate individuals and what processes and safeguards should be enacted. Where appropriate, the policy should also cover company IT assets while in the possession of employees. The types of issues addressed should include computer room or network center fire suppression systems and environmental controls. For example, if the computer room is not monitored constantly, is there an automated system in place that pages someone in the event that the fire suppression system is triggered or the environmental controls fail?

### Use of Standards

Policies should be developed that dictate a standard platform or common operating environment that is deployed throughout the organization. Adherence to the platform should be mandatory. In addition to reducing costs and administration requirements for an organization, standards can also protect data and infrastructure. Standards also aid in the interoperability and portability of applications in a distributed computing environment. Standards should even be considered for the look of the desktop. Nothing is more annoying to someone from the IT group than to sit down at the system and find that all of the icons have been changed to nonstandard images such as flowers, bumblebees, and smiley faces. Nonstandard icons impede the support process. In addition, installing these nonstandard icons is a security risk in that it is an introduction of unknown files into the network. Consideration should be given to restricting local administration of all desktops to ensure that standards are maintained.

### Reporting Noncompliance

Frequently, organizations educate employees and end users on their responsibility to report noncompliance but never put in place a mechanism to provide that capability. There are times when an employee may not feel comfortable reporting an incident of noncompliance. If the noncompliance involves a supervisor, systems administrator, or real criminal activity the individual may be apprehensive to report the occurrence for fear of reprisal. In this type of circumstance you need to be able to provide a way to report issues of noncompliance anonymously. Consider setting up a hotline for reporting such matters. To ensure the caller's anonymity, consider using an outside service or third party for this function.

## Personnel-Related Policies

### Introduction

There are several personnel-related policies that should be implemented which impact procedures in the areas of employee hiring, termination, and attendence.

### Pre-Employment Screening

Before hiring someone for the IT group, check all references. Never assume simply that

because a person gave references that the reference will be good. Talk to the references and question them about the candidate. Also consider a credit check. A bad credit rating or a history of bankruptcy on a potential employee's record may indicate someone who is not responsible or who is financially strapped. This is an indication that the person could be a potential risk, especially if he or she will be involved with systems that process financial transactions. Financial institutions, in particular, like to have employees who are financially responsible. If possible, consider drug screening. Many people consider drug screening to be an invasion of privacy, but when hiring new employees it can be a useful tool to weed our questionable applications. The drug test can also be required if you wish to have the employee
bounded.

## Mandatory Vacation Policy

Every employee in the IT unit should be required to take at least five consecutive business days off each year. Also consider rotating job functions and responsibilities. This is something
that should be considered for most every position in an organization. It is unfortunate, but most often the employee that embezzles from a company or is caught committing some fraud is almost always considered to be a model employee up until the crime is discovered. Such employees are often considered to be hard workers, because they almost never take any time off. The reason they never take time off is because if someone filled in for them while they were out, the irregularities would be discovered. As a result, they come into work everyday, sick or healthy, without fail, and always do their job.

For this reason, I recommend that organizations adopt a policy that requires each and every employee to take off five consecutive business days, so that someone else can perform their job function for that period. This may not identify every instance of employee misappropriation, but it will at least catch some.

## New Account Policy

When an account is created for a new end user or employee the system administrator should not be the one to determine what level of authorization and access to assign the account. This policy would cover the process of notifying the system administrator of the new end user accounts and the level of access required. There should also be a follow-up review process. The follow-up review should be performed by someone other than the system administrator to
ensure that the access level assigned to the new account was the authorized level.

## Security Access Change Request

When an access level change is requested for an existing account, such changes should be documented and authorized by someone other than the requesting parties. When the access level change is effected, it should be reviewed by someone other than the system administrator who made the change.

## Employee Termination Checklist

When an employee is terminated, either voluntarily or involuntarily, all access to systems should be deleted, and all keys, badges, files, or equipment should be recovered. If the employee was a system administrator all passwords should be changed, if possible. If it is not possible to change all passwords, then at least change the passwords for privileged accounts. If the employee had dial-in access, consider changing telephone numbers. Someone other than
the system administrator should review that the employee's access to all systems has been removed.

Procedures also need to be developed to handle occasions when turnover takes place for critical IT employees. Specifications need to be developed to address hardware and software training for firewalls and network operations. The policy should seek to avoid allowing a situation to develop where one person has all the knowledge, and there is no succession plan in place.

### Information Protection Team

Any corporate information security policy should include the formation of an information protection team. This team should be responsible for reviewing, monitoring, and enhancing policies and standards for the organization in regards to information security. The team's charter should include reviewing the security implications of any new major system prior to its implementation. The information protection team and its authority should be codified in an organization's policies and procedures.

### Crisis Management Planning

Every organization's planning and procedures should include some kind of crisis management planning. Most organizations will need at least two sections to any crisis management procedure: One section should deal with disaster recovery planning, and the other section should cover computer security incident response planning. Chapter 16 discusses crisis management in more detail.

# Chapter 15: Auditing, Monitoring, and Intrusion Detection

## *Overview*

It is stating the obvious to say that, in this day and age and for some time to come, organizations will rely heavily on computers and networks for their existence. As a result, the accuracy of those systems and people who use and maintain them are crucial to an organization's survival. Moreover, since people make mistakes and since some people can be dishonest or malicious, organizations need to regularly audit and monitor their computers and networks.

With the introduction of computers and networks, the concept of an audit has expanded to have multiple meanings. Historically, audits have had the effect of reducing reliance on administrative and procedural controls. In other words, the controls were not built into the process but were in the verifications that took place afterward. This does not mean that an audit negated the need for procedures and controls, but it would only catch any deviations from those procedures and controls after the fact. The residual risk was deemed acceptable to the operation of the organization, so that an audit was only required periodically. This orientation toward residual risk may have been acceptable in the past, but it is very dangerous in today's environment. As a result, in some contexts "audit" has become synonymous with "monitor."

This chapter covers three separate aspects of "auditing." First is the traditional electronic data processing (EDP) audit to which most IT departments are subject and which is usually performed with the assistance of an outside firm or by an internal audit department. EDP audits can review issues such as the controls for application development, records retention, copyright requirements, and general operational issues, but this chapter focuses on security audits. The second aspect of auditing examined in this chapter is system auditing and tools that are available to periodically check the integrity of either an individual system or a network in general. The third is intrusion detection, which is a process of ongoing auditing or monitoring of the security and integrity of an organization's systems and networks.

## *What Is an Audit?*

Traditionally, an audit is an independent review of a given subject. Its purpose is to report on conformance to required standards. One of the functions that an EDP audit serves is to verify compliance to company policies and to ensure that required security procedures and practices are being followed. In addition, an EDP audit usually entails the process of monitoring and analyzing systems, networks, and end-user activity.

In addition to reviewing compliance to policies and procedures, an audit is concerned with risk assessment. An EDP audit assesses the risks to and associated with systems and networks to determine if the existing controls are adequate to protect the organization's assets. Some of the areas that a security audit would review include the following.

- Ensuring that desk manuals and procedures are up to date;
- Ensuring proper segregation of duties with proper reviews of work;
- Ensuring that adequate physical controls are in place;
- Ensuring that user authentication controls are adequate;
- Ensuring that audit trails are maintained;
- Ensuring that disaster recovery/business resumption plans are in place and tested regularly;
- Ensuring proper controls for application development and implementation;
- Ensuring that data integrity is monitored and maintained;
- Ensuring that general policies and procedures are followed.

An audit can be an opportunity to validate an organization's security policies and can provide IT with a chance to have an outside party test the security measures that have been implemented. It is not uncommon to employ a "tiger team" or "white hat hackers," as they are sometimes called, to test security measures. These are network security experts who test system and network defenses by attempting to "hack" into them. This hacking is done with the knowledge and consent of the organization that owns the network or systems that they are attempting to penetrate. Such individuals are usually hired consultants, but some organizations employ internal staff for tiger teams.

If an organization does business through a partner or a third party, then the organization's IT unit may need to audit that partner's or third party's security measures. This is particularly true

if an organization uses a portal, colocation, or ASP vendor to provide Internet-enabled or branded Internet services to customers. It would be extremely risky for an organization to enter into an agreement with an ASP without first certifying all aspects of the ASP's computer operation, including security. When using a colocation or ASP service, a company can find itself the indirect victim of a denial-of-service attack directed at another subscriber of the service.

As mentioned above, there are many areas reviewed during an audit. Consequently, for large installations, it may be necessary to categorize the functions and audit the functions separately. For example, the functions can be categorized under the following headings:

- Operational audits;
- System audits;
- Activity and usage audits.

Operational security audits seek to ensure that proper controls have been established to identify deviations from established standards and policies. This type of audit is designed to mitigate vulnerabilities introduced by poor management.

There are several objectives for system security auditing. The first is to validate the system configuration. System security audits also seek to analyze the system configuration to mitigate

vulnerabilities introduced by the faulty implementation of a system, network, or application. The types of things a system audit reviews or looks for includes, among other things:

- *Accounts without passwords:* It happens more often than you would think.
- *Adherence to and enforcement of password policies:* How easy is it to crack the

passwords?

- *Shared accounts:* Are there accounts to which more than one person has the password?
- *Dormant accounts:* These accounts are often used by hackers and should be deleted.
- *Files with no owner:* These files are open to abuse, because anyone can take possession of them.
- *Files with inappropriate access rights:* These files are also open to abuse. It is very important that critical system files have the proper access rights.
- *Separation of duties:* Is there a process of checks and balances in place with proper reviews, or does one or two individuals have all the controls?

Even a secure system that is properly configured is vulnerable to attack, and auditing provides an excellent way of determining whether and how such attacks may take place.

Another reason for a system security audit is to monitor for attempted probes, attacks, and other unusual occurrences. Auditing a system can also assist in setting baselines for system usage, which are used to identify abnormal activity.

System monitoring relies heavily on system audit logs or event logs. General system log files record particular events including the following:

- Logins or attempted logins;
- Logouts;
- Remote system access;
- File opens, closes, renames, and deletions;
- Changes in privileges or security attributes;
- Changes in access control levels.

These log files are usually maintained on the server's or system's local disk drives and as such are vulnerable to alteration. It is generally a good practice to either move the log files to another server on a daily basis or simply print out the pertinent log entries to ensure a hardcopy record that can not be altered.

There are several software tools available to aid in the process of auditing a system. Two of the best known open source freeware programs are COPS and SATAN, which are discussed in Chapter 7. There are also a number of commercial products available from vendors such as Internet Security Systems (ISS), Secure Networks, Cisco, and Netective, just to name a few.

The key to an activity and usage audit is the establishment of baseline metrics to assist in identifying potential security problems. System activity audits seek to analyze deviations from the normal patterns of usage and other unusual activities. Baseline metrics should be established to assist in identifying potential security problems. The purpose is to identify abnormal usage and to identify possible attacks in progress. Exceeding these metrics or thresholds should trigger alarms in the system or should cause action to be initiated by whoever reviews the reports or logs. For example, baseline metric may reveal that a particular employee accesses data 20 times more than any other employee does? Why? What causes him or her to deviate from the normal behavior pattern?

Baseline metrics are also a key element of certain types of IDSs, which are discussed in more detail later in the chapter. However, this type of auditing does not require an IDS to be effective. Auditing of this nature is geared toward the application level, and most IDSs are generally geared toward the operating system level. It is usually not difficult to generate simple queries that identify anomalies in activity for a given period. Sometimes all it takes is a

daily eyeball review of standard reports. Identifying anomalies for overall system activity on a
daily basis is relatively easy. It is more difficult to identify anomalous behavior for a given operator over an extended period. An employee who accesses files 20 times more in a given day than any other employee would be easy to spot, but identifying anomalous behavior for a given employee over a period of time is more difficult because it requires comparisons to historical data. For example, a user account that logs in at an unusual hour for a given period might warrant investigation. However, one would not know if the hour that the login occurred was unusual unless there was historical data that indicated that the behavior was not normal.
Many applications have built-in logging capabilities that can be utilized for routine reviews. In general, it is advisable to log, monitor, and review transactions that require any kind of an override of system parameters. For example, in the financial industry it is normal for controls to be put in place for each employee that specify an upper dollar limit on transactions. These controls or system parameters may require a supervisor override if a transaction is over the specified amount. Most systems of this type generate a daily supervisor override report that is reviewed by a third party to ensure the legitimacy of the transactions and to serve as a precaution against collusion.

These types of reports are usually one of the first things examined by external auditors. External auditors examine them to ensure that the reports are, in fact, reviewed on a daily basis. Auditors will sometimes sample the transactions to ensure that they are legitimate. Auditors may also review whether the trigger mechanism that flags the transaction is adequate
for the function.

Another auditing function that is frequently a standard feature of many system applications is data change tracking. This is the ability to identify the last operator that made a change to a data element, record, or file. The minimum characteristics that should be recorded include the following:

- Identification of the operator making the change;
- Type of change;
- File and data element;
- Date and time of a change;
- Whether the change was successful;
- What the element was before and after the change.

This capability can be particularly crucial when investigating fraudulent activity involving insiders, such as employees.

## Audit Mistakes

Ideally, an audit should be seen as an opportunity to improve processes. Unfortunately, the reality is sometimes one of finger-pointing and recrimination. Based on personal experience, some of the more common mistakes that contribute to a difficult EDP audit are described as follows:

- *Not consulting with IT in the scheduling or planning process:* Nothing will ensure a difficult EDP audit like scheduling one during a period when the IT division is stretched to the limit working on projects. This results in the IT division feeling imposed upon and resentful of the untimely intrusion. The IT division's resources may already be stretched to the breaking point when they start getting requests to provide all sorts of information and reports for the auditors. On the other hand, the auditors feel that IT is not cooperating, because IT is not responding in a timely manner to the requests for information. This makes for strained relationships and almost ensures that a process that should be one of open communications becomes painful and difficult.
- *Auditors not properly trained to perform an EDP audit:* I've been involved in EDP

audits where the auditors did not have the technical background necessary to adequately perform the audit. In these instances the results were mixed. In some cases, the auditors simply accepted everything they were told by the IT group to be factual and accurate. There was no process of independent verification. While this might make the process easier on the IT group, it is not a true audit and does not serve the needs of the organization as a whole. In other cases I've seen the lack of technical knowledge on the part of auditors make then insecure about information with which they are provided. In some cases I've seen it border on paranoia. Since the auditors had no way of independently verifying information with which they were provided, they doubted everything.

- *Leaving it up to IT to enforce unilateral changes within the organization:* It is not unusual for deficiencies in procedures to be identified, over which the IT unit has no control. For instance, access levels within applications may be administered by the IT unit, but those who determine the actual level of access may reside within other business unit. As an example, the ultimate authority as to who has what access to the HRMS is the director of human resources. The IT group supports the HRMS package, but it is human resources who owns it, and it is they who determine who will have access to what information. On more than one occasion, I have seen audit findings in a final report regarding issues over which IT had no control or say in the process. However, the items were still cited as deficiencies in the audit. The IT group is left to correct the deficiency, over the objections of another business group.
- *Doing it by the book:* Auditors sometimes fail to recognize one of the cardinal rules of network security, which is that security measures and procedures that interfere with the operation of an organization are of little value. Those types of measures are usually ignored or circumvented by company personnel, so they tend to create security holes rather than plug them. Whenever possible, security measures should compliment the operational and business needs of an organization. Some auditors have a tendency to site any deviation from standard recommended practices, even if the deviation makes sense operationally for an organization. Security is a balancing process—balancing the security needs with the business needs and the probable with the possible. Too often auditors concentrate on the possible and not the probable.
- *Audit report does a hatchet job on IT:* It is not uncommon for the final audit report to be unnecessarily harsh on the IT unit. This is often a result of the mistakes listed above. Misunderstandings, lack of communication, and general distrust often lead to harsh findings. This is very unfortunate, since the security audit is actually an opportunity to test, learn, and improve an organization's security. As such, it should be welcomed, but too often it is met with dread. The IT unit and the audit group need to work together in developing the final report, so that it is comprehensive and practical. It needs to be comprehensive in that no area is glossed over. It needs to be practical in that no audit recommendations should constrict or interfere with the operation of the organization.
- *Lack of management support to implement audit recommendations:* The surest way to ensure that an audit is a failure is for management to fail to support the implementation of the audit recommendations. Management support is critical when implementing policy changes, particularly when those changes meet with resistance. In some cases it may simply be a matter of management not allocating the resources necessary to implement the recommendations. Most organizations have projects with deadlines and commitments that existed before the audit. Implementing the audit recommendations is always something that is given low priority. Ultimately, the recommendations are never implemented, and the same findings are usually cited at

the next audit.

## Deficiencies of Traditional Audit Techniques

The unfortunate reality is that it is not possible to build a completely secure system or network. Procedures are sometimes ignored. Passwords are vulnerable, and technologies fail or are subverted. Even in an environment where everything functions according to plan, the systems are still vulnerable to abuse by privileged insiders, such as system administrators.

The ultimate goal of a network security scheme is to prevent successful attacks on a network. Traditionally, the primary tool for ensuring network security has been the firewall. However, firewalls are almost useless for monitoring activity on the internal network. Organizations are beginning to recognize the need to audit or monitor their internal networks simply because the

majority of all attacks and losses involve insiders.

While traditional security audits may identify weakness in security measures or even expose security breaches, it is usually after the fact. Audit tools, such as COPS or SATAN, will only identify weaknesses in the configuration or implementation of systems or networks. Neither one of these approaches identifies problems as they occur; instead, they are concerned with residual risk. Traditionally, the residual risk was deemed acceptable to the operation of the organization, so that an audit was only required periodically. In today's Internet-connected environment the paradigm of residual risk is no longer valid. As a result, more proactive methods are required to audit or monitor networks and systems. Today there are new tools available that provide administrators with the ability to monitor network and system security on-line in real time.

## Intrusion Detection

Competent system administrators have always monitored their systems for intrusions. The process usually entailed reviewing logs on a daily basis. Intrusions were sufficiently rare that after-the-fact reviews were usually adequate to address any possible problems. Unfortunately, times have changed drastically. After-the-fact reviews are no longer adequate; real-time or near real-time responses to intrusions are necessary. In addition, the volume of activity on the networks today dwarfs what was the norm 10–15 years ago. As a result, it is not humanly possible to review the amount of information in today's log files without some automated process. Without the automation of the review and monitoring process, it could be weeks before a system administrator knows about an intrusion to his or her system.

In general terms an "intrusion" can be defined as an unauthorized attempt or achievement to access, alter, render unavailable, or destroy information on a system or the system itself. Basically, an intrusion is somebody attempting to break into or misuse a system. Some observers differentiate misuse and intrusion. The term intrusion is usually used in reference to attacks that originate from outside an organization. Misuse is usually used to describe an attack that originates from the internal network. However, not everyone makes this differentiation.

Intrusion detection is the art of detecting unauthorized, inappropriate, or anomalous activity. The art of intrusion detection has been practiced by system and network administrators for years. However, intrusion detection has recently received more attention in the media largely due to the fact that so many companies are now marketing IDSs. Supposedly, these new IDSs can identify attacks in progress, generate real-time alerts, and even launch countermeasures or

reconfigure routers or firewalls to counter an attack.

## Intrusion Detection Systems (IDSs)

IDSs act much like security guards or sentries. They constantly scan network traffic or host audit logs. While the present batch of IDS products provide useful tools to augment an organization's network security, it is necessary to get past the marketing hype to evaluate a system's effectiveness. Presently, no single system provides truly effective end-to-end intrusion detection capability. In addition, IDSs are not a new concept. In Chapter 7, we discussed the TCPWrapper, a UNIX-based freeware IDS that has been around for many years.

Generally, IDSs fall into one of two categories:

- Network-based IDSs;
- Host-based IDSs.

While there are merits to both approaches neither method by itself is sufficient to monitor all threats. As a result, the current trend in the industry is to combine the two approaches.

### *Host-Based Intrusion Detection Systems*

Host-based products reside on the host and are capable of automatically monitoring and denying services if suspicious activity is detected. They monitor activity on the individual host as opposed to monitoring activity on the network. Host-based IDSs still rely on system audit logs, much the same way system administrators do, but IDSs automate the process. Typically a host-based IDS monitors system, event, and security logs on Windows NT and the

syslog file for UNIX. The host-based IDS uses system log files and the system's own auditing agents to monitor the system.

There are a couple of approaches that host-based intrusion detection software can employ. One is to employ a wrapper, like TCPWrapper. This approach wraps the various host network services in an extra layer or shell that interprets network packet requests to the various services. The other approach employs agents that run as separate processes and monitor the requests to host. Both approaches are effective at detecting anomalous activity or misuse of host systems.

One advantage to host-based agents is that they can monitor changes to critical system files and changes in user privileges. When a key system file changes, the IDS compares the files properties with known attack signatures to see if there is a match. One popular method for detecting intrusions involves verifying key system files and executables via checksums at regular intervals for unexpected changes. For example, Chapter 7 discusses using MD5 to monitor changes to system files and the Tripwire IDS, which also provides this function. The first time one of these systems is run, it generates a snapshot of the file attributes, including file sizes and access rights. This information is stored in a database. Each subsequent run of the IDS compares the attributes of the files on the disk to the attributes stored in its database. If the attributes have changed then an alarm is sounded.

Some host-based IDSs monitor TCP port activity and notify system administrators when specific ports are accessed or scanned. They can also monitor and record when physical ports are accessed. This can be useful if the port has a modem connected to it.

Perhaps the biggest drawback to host-based IDSs, such as TCPWrapper and Tripwire, is that the intrusion detection process is not real-time. Host-based intrusion detection programs, regardless of whether they use some wrapper or agent, generally identify intrusion attempts after they have been attempted or succeeded. The lag between the intrusion and its discovery can be substantial. By then it can be too late. This is a weakness with host-based IDSs in general. Another general weakness with host-based IDSs, like TCPWrapper and Tripwire, is that they don't have any capability to proactively react to an intrusion. Nor do they allow the system administrator to be proactive.

Another drawback to the host-based approach is that to secure the entire network, it is necessary to load the IDS on every computer. However, this aspect of host-based IDSs can also be a benefit. If you only desire to monitor one system, the cost of host-based IDSs is often lower than those for their network-based counterparts. As we have already discussed,

there are freeware versions of host-based IDSs available on the Internet. In addition, hostbased
IDSs usually require no additional hardware, since they run on the system itself.
Network-based IDSs very often require a dedicated system or device to function. This too increases the cost.

Another advantage to a host-based IDS is that it monitors specific systems and can identify non-network-based attacks. The host-based IDSs can monitor system file integrity, file permissions, and other file system parameters that a network-based IDS does not monitor. In addition, host-based IDSs can monitor terminal connections that bypass the network, and they can also monitor the specific activities of someone logged into the host and accessing files. In addition, a host-based IDS can monitor the activities of applications and processes running on the host. A network-based IDS can only monitor the network, not what is occurring on a specific host.

### *Network-Based Intrusion Detection Systems*

Netwrok-based IDS products run on the network and monitor activity analyzing patterns and reporting on suspicious activity. A network-based IDS usually employs a dedicated network server or device with a network adapter configured for promiscuous mode to monitor and analyze all traffic in real time as it travels across the network. The network-based IDS monitors packets on the network wire and attempts to discern the legitimate traffic from the malicious. Some vendors state that a dedicated server is not necessary for the functioning of their network-based IDS. However, in reality it would not be advisable to run an IDS on a general-purpose application server. Would you want your network's IDS running on the company's payroll server?

When compared to host-based IDSs, network-based IDSs have advantages and disadvantages.

Depending on the system, a network-based IDS may be less expensive to implement. This is due to the fact that a network-based IDS is operating system-independent and is not required to be loaded on all hosts on a network to be effective.

In addition, host-based IDSs will miss many network-based attacks. Host-based IDSs do not examine packet headers, so they cannot detect denial-of-service attacks. Network-based IDSs are also much more stealthy than host-based IDSs. With a host-based IDS, if the system is compromised a hacker can readily see if there is an IDS present. It would be very difficult to determine if a network-based IDS was on a network simply by examining the wire. About the only thing a hacker could determine is that there is a device on the network running in promiscuous mode. A network-based IDS can also provide superior controls on event logs. With many host-based IDSs, the audit logs reside on the system locally. As a result, if the system is compromised, a hacker can manipulate the log files to hide his or her tracks.

Another weakness of network-based IDSs is the fact that they become less effective as network traffic increases. They work very well on an empty network, but as the number of packets increase, their effectiveness decreases to the point where they cannot identify any intrusions. This is a major weakness considering today's high transaction volume and the growth of fast Ethernet and switched Ethernet.

### *Knowledge-Based Intrusion Detection Systems*

There are two general approaches employed for identifying hostile intrusions. One is knowledge-based, and the other is statistical-based. The two approaches are very different and
employ different technologies.

Most of the IDSs deployed today are knowledge-based. Knowledge-based IDSs are sometimes referred to as misuse detection systems, expert systems, or model- or signaturebased

IDSs.

Knowledge-based IDSs rely on the ability to recognize known attacks. A knowledge-based IDS recognizes known intrusion scenarios and attack patterns. The knowledge-based IDS relies on a database of attack "signatures" or "patterns" that can be changed for different systems. For example, a host-based, knowledge-based IDS may monitor keystrokes for attack patterns. The IDS has a database of known keystroke patterns that are known to be a threat.

Knowledge-based IDSs employ many different techniques to identify intrusion patterns or signatures. For a host-based, knowledge-based IDS the process can involve monitoring keystrokes, reviewing files for changes and monitoring ports. The review of files can function much the same way as a virus scanner on a PC. The scan searches for known patterns or changes that have been made to critical files since the last scan. String signatures look for text strings that indicates a possible attack. An example of a string that might raise a red flag for a UNIX system would be someone examining the contents of the password file or hosts file using "cat /passwd" or "cat /hosts." You should always be suspicious of someone who wants to examine the password file or review what other hosts are on the network. When monitoring ports, a host-based, knowledge-based IDS can compare audit logs to the signatures of common techniques. As an example, a significant number of failed TCP connections to wellknown ports may be an indication that someone is scanning ports, or a large number of unacknowledged SYN-ACK packets is probably an indication that the system is under a SYN flooding attack.

A network-based, knowledge-based IDS examines packets on the network. Packets are considered suspect if they match a known signature, string, or pattern. A network-based, knowledge-based IDS can examine the protocol stack for suspicious invalid or fragmented packets that violate the TCP/IP protocol. The ping-of-death with its oversized ICMP packets would be an example of a known signature. A network-based, knowledge-based IDS can also examine packet headers for dangerous or illogical combinations in packet headers. Another well-known header signature is a TCP packet with both the SYN and FIN flags set, signifying that the originator wishes to start and stop a connection at the same time. This can be an indication that a system is being probed by an intruder.

Knowledge-based systems that employ pattern matching simply translate known intrusions into patterns that are then matched against the system or network activity. The IDS attempts to match activity to the patterns representing intrusion scenarios. The IDS monitors the activity, accumulating more and more evidence for an intrusion attempt until a threshold is crossed. The basic approach underlying pattern matching is that if it looks like a duck, walks like a duck, and quacks like a duck, then it must be a duck. However, for pattern matching to work the patterns must be easily recognizable, and they must be distinguishing. In other words, they must not look like any other normal or legitimate activity.

The advantages of knowledge-based IDSs is that they usually have low false alarm rates. This is due to the fact that they usually watch for very specific signatures, strings, and patterns. In addition, because they watch for specific events they are able to report with some detail and certainty on the threat being faced, which makes it easier to determine the appropriate course of action.

The major disadvantage to knowledge-based IDSs is that they are only effective against threats with which they are already familiar. As a result, they are useless against new techniques for which they have no signature or pattern in the knowledge base. In addition, it is not a simple matter to create a signature or pattern for an attack. It is not easy to translate known attack scenarios into patterns that can be used by a knowledge-based IDS. It requires keeping the IDS up-to-date with new vulnerabilities and environments. Further, it requires time-consuming analysis of each new vulnerability to update the IDS's knowledge base. As a result, vendors don't update their databases as often as they should.

Another common weakness of knowledge-based IDSs is that they are ineffective against passive attacks, such as network sniffing and wiretaps. They are also ineffective against IP or sequence number spoofing, DNS-based attacks, session hijacking, and redirects. In addition, a knowledge-based IDS will not detect the fraudulent or malicious activity of a privileged insider if the activity does not match a known pattern or signature. This is particularly true if the activity is performed through an application. For example, fraudulently transferring funds from one account to another will not be flagged, since it would be within the normal parameters of the system. Some of the better known network-based IDS products are from AXENT, Cisco, and Internet Security Systems (ISS).

### *Statistical-Based Intrusion Detection Systems*

Statistical-based IDSs identify intrusions by developing base-line measurements for "normal" activity and assuming that anything that deviates significantly from the norm is an intrusion. In other words, intrusions are recognized by identifying deviations from normal or expected behavior of the system or the users. Statistical-based IDSs are also referred to as behaviorbased IDS or just simply anomaly detection systems. The underlying philosophy is predicated on the concept that anything new, different, or unknown must represent a threat to the security of the system or network.

A statistical-based IDS (SIDS) develops a model for "normal" patterns of activity and behavior by collecting information from various sources. The SIDS learns what is normal by knowing what the patterns have been historically. It requires large quantities of information and data to develop an accurate and useful model. The more information the IDS can acquire, the more it learns and the more accurate the model. These models can be developed at the system, user, or application level. A model can be developed for any kind of activity that needs to be monitored. The models, which are based on historical information, are used to compare and validate the ongoing activity for a system, user, or application. When a statistically significant deviation is observed, an alarm is generated. In other words, anything that does not conform to previously learned pattern or behavior is deemed to be suspicious.

The major advantage of a SIDS versus a knowledge-based system is that a SIDS does not rely on a predefined set of known attack patterns or signatures. As a result, the SIDS can detect attempts to exploit new vulnerabilities. At least theoretically it can. SIDSs are also less dependent on operating system–specific mechanisms.

Another advantage to a SIDS is that it can detect the fraudulent or malicious activity of a privileged insider. For example, the fraudulent transfer of funds from one account to another could set off alarms if the user did not normally access that account, or if the dollar amount was over what was normal for the individual, or if it was done at an unusual time. In other words, the alarm would go off if the transfer were statistically significantly different from the user's normal activity or behavior.

There aren't many SIDSs on the market today—at least not for the standard vanilla corporate computer system or network. This is due to a number of factors. First, they tend to have a high number of false alarms. This is due to the fact that SIDSs consider almost any activity that is new or different to be a threat. Very few networks are static. In addition, developing user profiles may be difficult, especially in an environment where users work irregular schedules or there is a high turnover. As a result, it would be very difficult to implement a SIDS in an environment where changes to the users, network topology, servers, or applications are the norm.

SIDSs must also be flexible enough to modify their model as a user's or network's activity patterns change. However, this flexibility can actually be exploited to an intruder's advantage. If the incremental changes are minor and performed over an extended period, an intruder can "teach" the SIDS to accept fraudulent or malicious activity as "normal."

Another weakness of the SID approach is that regardless of whether the SIDS is taught over time that an intrusive activity is normal or whether it is simply an oversight on the part of the system, a SIDS will not recognize attacks of any kind that conform to behavior that it has learned to be normal. In other words, if it's not abnormal, then the SIDS will assume it is not an intrusion. This logic is often fallacious.

Another concern with SIDSs is how to determine what components to monitor for creating the models. The possibilities are endless and include file access, access time, network connections, volume of packets, and CPU utilization. In addition, determining when a deviation from the norm becomes statistically significant can be difficult. Like most things in network security, it is a balancing process.

Many of the SIDSs in use today utilize neural networks. A neural network is a type of artificial intelligence that can be trained to learn. The training usually involves feeding the neural network large amounts of data and programming a complex set of rules about data relationships. Once set in place, the rules can be adjusted by the neural network based on additional input. Neural networks "learn" from examples and additional input. A neural network is capable of learning from examples to find patterns in data from a representative data sample. The more examples or input the network receives the more it learns. Neural networks are able to predict future events based on past performance.

A neural network usually involves large parallel processing systems employing the concept of fuzzy logic. Neural networks are sometimes described in terms of knowledge layers, with more complex networks having more layers. These systems examine the inputted data and make determinations based on the complex set of rules and past examples.

Neural networks are being employed for credit risk analysis, predicting market trends, weather forecasting, and fraud detection. For example, VISA and Mastercard use neural networks to identify fraudulent activity. The neural networks comb the millions of daily transactions to identify anomalies in activity based on each "individual" cardholder's past patterns. This is an impressive accomplishment, considering the volume of transactions and the number of cardholders each company has in its customer base.

## Defense In-Depth Approach

Like a firewall, an IDS should be seen as just one more tool in a defense indepth approach. Security measures should be multitiered, and IDSs can serve as another layer of security. Before you deploy an IDS, however, make sure that you weigh the pros and cons and be sure that the vendor you pick has the system that best meets your needs. Some of the pros of IDSs are listed as follows:

- Can detect some abuses and intrusions;
- Can identify where attacks are occurring;
- Can be useful for collecting evidence;
- Can alert administrators that someone is probing;
- Can take corrective action against certain types of abuses or intrusions.

Some IDS cons are listed as follows:

- Misses many types of abuses and intrusions;
- Do not work well no high-speed or heavy-volume networks;
- Generates false alarms.

An IDS can add depth to your overall security, helping to identify possible intrusions and abuses, but an IDS by itself does not ensure security. IDSs have a long way to go before they are as effective as much of the marketing hype would have you believe. Network-based IDSs' inability to function effectively on noisy, high-speed, or high-volume networks is just one example of the limitations that IDSs have to overcome before they become truly effective. Even when they are functioning correctly, all IDSs still miss many specific and harmful types of attacks. The most effective approach to intrusion detection is to use a combination of network-based and host-based detection.

### *Future Directions*

Intrusions or abuses usually are not confined to a single system or network segment. We now work in an environment where information is distributed over large networks that are centrally administered. As a result, it would be useful to have intrusion detection tools that employed a distributed approach where the host-based IDS communicated with the networkbased IDS, and both notified central administration of any anomalies.

To this end, the IETF has formed a working group to study intrusion detection. According to the working group's charter, the purpose of the IETF intrusion detection working group is: "… to define data fromats and exchange procedures for sharing information of interest to intrusion detection and response systems and to management systems [that] may need to interact with them."

One can only hope that the IETF's efforts lead to integrated end-to-end IDSs that are able to monitor for and react to intrusions and abuses for the entire enterprise.

# Chapter 16: Crisis Management

This chapter describes the planning process that every organization should go through to prepare for an event that threatens the operation or viability of the organization. Disaster recovery and computer security incident response planning can be thought of as two sides of a coin. The two topics are closely related and share some common methodologies and goals. Both are concerned with ensuring the availability and integrity of an organization's networks and systems.

### *Disaster Recovery Planning*

From time to time, many businesses face a catastrophic event that can threaten the viability of the organization. Accordingly, every organization should formulate a set of procedures that details actions to be taken in anticipation of a catastrophic event. The procedures should be designed as if the catastrophic event is inevitable and is going to take place tomorrow. This type of plan is referred to as a disaster recovery plan. In some organizations, disaster recovery planning is called contingency planning or business resumption planning.

Some organizations believe that having hot site recovery services is the same as having a disaster recovery plan. A hot site is a facility that is designed to be activated in the event that an organization's computers or computer facilities are rendered inoperable. A hot site is preconfigured with the power, environmental controls, telecommunications, and computers necessary for an organization to resume computer operations with a minimal disruption in service.

In response to questions about their disaster recovery plans, colleagues have told me that they have contracted for hot site services or maintained redundant systems at another facility, as if all they need to care about was ensuring that the systems were covered. In these cases, the emphasis was on the hardware and software and not on the business and people. A disaster recovery plan is about the resumption of business operations, not just network and computer operations.

The requirements for a disaster recovery plan vary for each organization. However, for most organizations, the minimum objectives of a disaster recovery plan is to provide the information and procedures necessary to do the following:

1. Respond to the occurrence of a disaster;
2. Notify the necessary personnel;
3. Assemble disaster recovery teams;
4. Recover data that may have been lost as a result of the occurrence;
5. Resume processing as quickly as possible to ensure minimal disruption of an organization's operations;
6. Comply with any regulatory requirements that dictate the existence of a disaster recovery plan for the organization.

One of the key factors in the success of a business resumption plan is proper planning for the IT group. Most organizations today rely heavily on computers, networks, telecommunications, and IT in general. As a result, IT plays a key role in most organizations' disaster recovery planning. Usually, the IT unit develops its own separate disaster recovery plan, which details the actions necessary to minimize system downtime, thereby minimizing the disruption of the organization's operation. The IT plan is integrated with an organization's overall disaster recovery plan.

The topic of IT disaster recovery planning is expansive enough that it could easily fill a book. In fact many books have been written on the topic. This chapter aims to discuss IT disaster recovery planning from a business perspective and demonstrate how it ties into network security. It is important to remember that one of the key elements of information security is "availability." This refers to the availability of the information that is located on the systems and networks of the organization. Proper planning is necessary to ensure the availability of mission-critical systems. It is crucial in the planning process to determine what is an adequate level of preparation and what is a mission-critical system.

## What Level of Preparation?

The extent to which an organization is willing to invest resources into IT disaster recovery planning should be directly related to the business of the organization. Different organizations have different recovery needs, with regard to IT. As a result, the plans developed by different organizations should reflect their needs. For example, a nonprofit organization that relies on fundraising for income could probably survive several, days if not weeks, of downtime. A bank, on the other hand, could find itself out of business if its systems were down for that period. Most banks could accept a few hours to a day or two of downtime as a result of a catastrophic occurrence, while a stock brokerage firm that trades on the NYSE or NASDAQ could find itself in financial ruin if its systems were down for a few hours and it was unable to

trade. The amount of resources that go into IT disaster recovery preparedness is dictated by the operational needs of the organization and the organization's ability, or lack thereof, to survive downtime.

While it would be nice if every organization had unlimited resources to prepare for the immediate resumption of business after a disaster, like everything else in business, expenditures on disaster recovery planning must be justified through cost analysis. In the case of the nonprofit organization, there is little financial loss associated with the downtime itself. In other words, the inability to do business for a day or two has relatively minor financial impact on the organization. In such circumstances, it would be difficult to justify the cost of extensive disaster recovery preparation that included such things as redundant systems and telecommunications and hot sites. In contrast, the brokerage firm could most likely demonstrate that the inability to function, for even a short period of time, could potentially cost the company a significant amount of money. As a result, the brokerage firm could justify substantial expenditures for IT-related disaster recovery preparation and planning.

Disaster recovery planning decisions have to be made like most every other business decision.

The cost of the recovery has to be weighed against the losses incurred as a result of any downtime that may occur. When estimating the cost of the downtime, it is important to include the soft costs as well as the hard costs. The hard costs, such as lost revenue directly related to the downtime, are the easy ones to quantify. The soft costs are the hard-to-quantify items such as customer good will, level of service, and satisfaction or consumer confidence. It takes a thorough knowledge of an organization's business to be able to estimate the soft costs. Consequently, it may be difficult for an IT unit to estimate these costs alone. Therefore, participation from other business units within an organization is vital to the process of determining the costs associated with any downtime.

## What to Restore First?

Just as different organizations have different recovery needs, different functions within an organization have varying levels of priority for recovery. Any IT disaster recovery plan should assign levels of importance to each system to ascertain which systems will be given priority when restoring services. Mission-critical functions need to be identified prior to the occurrence of a disaster, so that when a disaster does occur, IT does not waste time restoring superfluous systems instead of those that are truly required. Once again, this takes a thorough knowledge of the organization's business—and input from outside of the IT unit. One approach is to gather this information through an assessment team headed by IT but with the participation of management and staff knowledgeable in the functioning of the organization and familiar with the various systems and applications. This process should include formally validating the teams' understanding of the organization's business.

Another approach would be to identify or assign ownership for each application and obtain the owners' input. By working your way up the hierarchy of applications and systems, you should be able to assign priority to each. This process should also include obtaining management's perspective on how critical each application is to the conduct of business. The process of prioritizing systems should be quantified, by performing a detailed analysis of each application, to determine what the cost would be to the organization to lose access to a particular function.

## Review and Test

From a cost analysis perspective, successful disaster recovery preparation is proportionate to the potential loss. From an operational perspective, a successful disaster recovery plan is responsive to the business needs of the organization.

From a general management perspective, a disaster recovery plan must be kept current and updated with any necessary changes. Modifications to systems, personnel, business priorities, and other environmental factors must be reflected in the plan. That means regular and frequent reviews of the disaster recovery plan. For most organizations, the shelf life of a disaster recovery plan is about three to four months. In other words, that is how long it will take for the plan to become out-of-date and need revision. During the three to four month period, personnel will turn over, technologies will be introduced and/or retired, new products will be released, and business priorities will change. As a result, the plan will need to be revised to reflect these changes.

There must also be regular and comprehensive tests of any disaster recovery plan, and the results of any tests must be incorporated into the plan. Moreover, key personnel must understand their roles and responsibilities in the plan. If they do not, then the best-written plan in the world will be of little value to an organization when a diasster strikes.

## Disaster Recovery Planning Case Study

As an example of why it is important to update such plans, a financial institution with which I had a business relationship undertook a comprehensive review of its then existing disaster recovery preparations. At that time, the business resumption plan was about eight years old. Since the time the plan was originally developed, individual sections of the plan have been updated to reflect changes in such things as personnel and technology, but the plan in its entirety has never been reviewed to determine whether it is still adequate for the financial institution's business model.

The purpose of the review was to ensure that the company was properly prepared to deal with a disaster, either limited in scope or large in scale, that interrupted normal business operations.

For the review, a team was assembled to evaluate and revise the existing disaster recovery plan. The team included a cross-section of business units within the organization. Most of the team participants were also key members of the business resumption team. The idea was for the team to review and revise the plan as necessary to address the needs of the organization. As with any disaster recovery plan, the goal of the existing plan was the swift resumption of the operations of organization in the event of a disaster.

The first task was to review the existing plan from start to finish. The review of the existing plan only confirmed what many of the team already knew. That was that the organization's business model had changed significantly since the plan was first conceived, but that none of those changes was reflected in the plan. After a thorough review of the existing plan, the team came to the conclusion that simply revising the plan would not address the needs of the organization. The plan was so outdated, and there were so many deficiencies in it that the plan needed to be entirely rewritten. Some of the major shortcomings of the organization's disaster recovery plan are detailed as follows.

- *Remote delivery systems:* When the financial institution's disaster recovery plan was first developed years earlier, well over 50% of all business with customers was performed in the branches by tellers. When the plan was originally developed, the financial institution had few of the electronic delivery systems that were in place at the time of the review. Those that had been deployed years earlier were not used as extensively by the financial institution's customer base as those that were in place during the time of the review. As a result, the primary emphasis of the organization's disaster recovery plan, as it existed, was the resumption of branch operations. A review of how the financial institution was operating revealed that over 85% of all transactions were performed through one of the electronic delivery systems. As a result, the financial institution's disaster recovery plan needed to be modified to reflect a primary emphasis on the restoration of the systems that were providing services to customers. The original plan did not adequately address restoring the organization's call center operations, Internet banking, or banking-by-phone systems. As a result, it was necessary to drastically rewrite the existing plan to ensure the swift resumption of these electronic delivery services in the event of a disaster.
- *Geographic range:* Another problem with the plan was that when it was first developed years earlier, both the overwhelming majority of the customer base and all of the financial institutions offices were located in the San Francisco Bay Area. One of the basic assumptions of the plan was that any event large enough to disrupt the operation of the financial institution, such as an earthquake, would also affect the customers. This assumption was used to prioritize systems. For example, if our telephones were affected, then the assumption was that the customers' telephones

would be affected. In this case, restoration of a system such as the bank-by-phone system would not be given a high priority. In addition, it was assumed that our customers, who would also be subject to the disaster, would understand if it required time to restore normal service. At the time the plan was conceived, the company was thinking like a small, community financial institution, and the plan reflected that mind-set. It was not necessarily a wrong mind-set, because at the time the plan was conceived, the organization was a small community financial institution. However, in the intervening years, the financial institution grew and expanded outside the San Francisco Bay Area. At the time of the review, it had offices throughout California, Texas, Oregon, New Jersey, and Arizona. In fact, over 40% of its customer base lived and work outside the San Francisco Bay Area. As a result, it could no longer be assumed that the customers would be subject to the same disaster that struck the financial institution. In addition, it could not and should not be assumed that the customers would be understanding if there were a major disaster in the San Francisco Bay Area. Customers outside the Bay Area would not care if there was an earthquake in San Francisco; they would still want access to the services that they expected the financial institution to provide. As a result, the financial institution needed to devise a disaster recovery plan that ensured the swift resumption of operations for all events, even if there was a major disaster in the San Francisco Bay Area.

- *Disaster recovery services:* In reviewing the changes that needed to be incorporated into the financial institution's disaster recovery plan, it was determined that the financial institution had outgrown the company with which it had contracted for "hot site" disaster recovery services. There were a number of reasons for choosing to review alternatives to the present service provider:

o *Multiple hot sites (local versus remote):* The service provider that the financial institution was using at that time could only offer a single hot site, which was located out of state. In the event the financial institution was forced to activate its disaster recovery plan, it would have to transport personnel, media, and supplies to the out-of-state location. This would add 24–48 hours to the time it would take to resume computer operations. While this scenario might be acceptable in the event of a major disaster, such as an earthquake, it would not be acceptable if the financial institution experienced a very localized event such as a fire in the computer room or simply a failure of a major system. Under the scenario of a limited disaster, the financial institution would want the option of activating a hot site that would be accessible locally. Since the travel time would be minimal, a local hot site would substantially reduce the amount of time it would take the company to get systems back on-line. Ideally, the best service provider could offer multiple hot sites, with a choice of local and remote sites.

o *Possible contention for service:* Another issue was the fact that the company that the financial institution used for disaster recovery services had contracts with several clients in the San Francisco Bay Area to provide disaster recovery services. In the event of a major disaster, such as an earthquake, the financial institution would have to compete with the other Bay Area clients for time and resources at the single disaster recovery center.

o *Disaster recovery center's capacity:* After reviewing the facilities available at the disaster recovery center, it was determined that the service provider did not have adequate resources to handle all of its San Francisco Bay Area customers. In addition, the financial institution was utilizing ISDN circuits to communicate with the service provider's disaster recovery center. As a result, it required that there be a one-to-one relationship between our branch offices and the service provider's ISDN ports at the disaster recovery center. At the time the contract was originally signed for disaster recovery services, the financial institution only had a limited number of locations that required connectivity to the disaster recovery center in the event of a disaster. In the intervening years, the financial institution had grown and was continuing to grow. It was therefore concluded that at that time the existing disaster recovery center did not have adequate capacity to handle all of the financial institution's branch offices simultaneously. The situation would only become worse, since new offices were opening at a rate of two or three a year.

After a review of existing disaster recover preparations, it was determined that the financial institution needed a disaster recovery service provider with abundant capacity and the ability to provide multiple hot sites. Ideally, any service provider chosen would be able to offer both local hot sites, which can be activated for a limited event, and remote hot sites in the event the

San Francisco Bay Area is affected by a major disaster.

Another conclusion of the existing plan review was that the present telecommunications configuration was designed for the resumption of branch operations and, as such, was not adequate for the electronic delivery systems. In addition, it was not flexible enough to address all contingencies. Nor was the present configuration as easy to implement as it should be in the event of a disaster. In the event of a disaster, the desire was to minimize the amount of intervention required to implement the backup telecommunications.

The financial institution needed to deploy a telecommunications configuration that encompassed the electronic delivery systems, that was flexible enough to address all contingencies, and that was relatively easy to implement. Accomplishing this goal would require the expenditure of funds to purchase new equipment and services.

## Outsourcing Plan Development and Maintenance

Many organizations don't have the time, resources, and expertise to put together a comprehensive disaster recovery plan. Under these circumstances, an organization should consider outsourcing the process of developing a disaster recovery plan—and even the plan's maintenance. Generally, out-sourcing the development of a business resumption plan includes the following

- Project planning and orientation;
- Reviewing recovery strategies;
- Defining recovery plans and supporting documentation;
- Developing test programs;
- Developing and implementing plan maintenance procedures.

Any consultant hired to develop a plan should provide disaster recovery education to selected company personnel to enhance their ability to understand and respond to emergency outages and to prepare them to participate in the development of the organization's overall recovery capability.

The recovery plans developed by the consultant should define the detailed actions that the company must take to declare a disaster, notify appropriate personnel of the disaster, activate the recovery plans, and execute a timely restoration and recovery. The plan should also include testing programs that define primary and secondary objectives of the testing and the frequency of testing. In other words, each test can have a different objective. One can test telecommunications while another tests operational procedures.

Any disaster recovery plan should also include a maintenance program to ensure that the recovery plan remains up-to-date. The maintenance procedures should include periodic reviews of the technology platforms.

Typically, a plan is developed by collecting information through interviews, workshops, teleconferences, and questionnaires, as deemed appropriate by the consultant. The plan development process should also make use of existing documentation where applicable. Information collected is used to evaluate the ability of company's disaster recovery plan to meet the organization's business requirements. Upon completion and documentation of the disaster recovery plan, the company should validate its contents by performing a detailed and thorough walk-through.

## *Computer Security Incident Response Plan*

Another aspect of crisis management planning is computer and network security incident response planning. Every crisis management plan should include a computer security incident response plan (CSIRP). This plan outlines actions that the company must take when there is fraud or misuse of company-owned electronic media or services, a theft or destruction of company information, or a penetration of, or attack on, company-owned systems and networks. The plan should address such things as what constitutes a security incident, identifying key personnel, a communication and notification process, as well as an escalation process. Needless to say, the information protection team should be a key component of any CSIRP.

Since security incidents occur with more frequency than disasters, organizations are finding that security incident response planning in some respects is more important than disaster recovery planning. In general, organizations experience few real disasters but deal with many security incidents. Denial-of-service attacks and virus outbreaks are becoming common. Most companies are willing to admit denial-of-service or virus incidents, but few are willing to disclose when their networks or systems are truly compromised.

### General Recommendations

Just as with disaster recovery planning, there is no single CSIRP that fits all organizations. There is no universal CSIRP template that can be applied to an organization. Each organization's security requirements and needs are unique. However, the following sections outline some general recommendations for CSIRP.

### Legal Counsel

The first step is to identify a computer security incident response team (CSIRT). This team can be different from the information protection team (IPT), but the IPT should have some representation on this group. In addition, since it may be necessary to take legal action against

the parties responsible for the incident, it is good idea to either have legal counsel on the CSIRT or at least have it readily accessible. The legal counsel may be necessary to determine whether it is possible to terminate or prosecute the individual or individuals responsible for the incident.

### Liability

Legal counsel may also be required to assist in assessing an organization's liability for any computer security incident. The liability can come in many forms. An organization may be liable for the direct loss resulting from a fraud or destruction of company assets. An organization may find itself financially liable for the disclosure of information regarding customers, employees, or partners. An organization may also need to assess its liability as a result of a customer, employee, partner, or hacker using the organization's systems to launch an attack on another company's system.

I have read accounts of system administrators tracking the activities of hackers on their systems to gather more information on the hacker. Rather than shutting down the hacker completely, the administrator limits the damage and monitors the activity to gather evidence on the crime and to identify the perpetrator. Clifford Stoll's account in his book *The Cuckoo's*

*Egg* is one example.

This kind of action or inaction can have risks, not the least of which would be an organization's liability should the hacker damage, steal, or misuse another organization's systems or information. If the hacker damages another company's systems, the question may arise, why didn't the first company stop the hacker when it had the chance? As a result, I recommend against this type of approach and suggest that if a hacker is detected, shut him or her down immediately. However, do gather as much evidence in the process as possible: Save audit and system logs to identify the origins and time of attack. Print out all logs to avoid having them altered or overwritten. Take detailed notes about what occurred, when it occurred, and any actions taken as a result. In addition, avoid using the compromised system or network for communications regarding the incident. It is possible that the hacker or hackers could intercept messages.

## Retaliation

A CSIRP should not include any measures for retaliation. It can be tempting to retaliate against a spammer or to trace a hacker back to his or her system of origin. This kind of cybervigilantism is illegal and can result in additional liability. I have even heard some accounts where system administrators tracked down hackers, physically went to the hacker's locations, and threatened them with bodily harm. While this approach may provide a certain amount of satisfaction, I strongly recommend against it.

Another reason to avoid retaliation is that hackers often use the systems of other innocent victims from which to launch attacks. This masks their location and makes it difficult to trace the attack back to its true source. In this case, retaliating would only create another victim. In many cases, this is actually the true intention of the hacker. The hacker hopes for a retaliation to be directed against the system or network from which he or she is launching the attack. If an organization does retaliate it can find itself not a victim, but a perpetrator liable for its actions.

## Triage

A CSIRP should include a triage process to handle all information regarding incidents. This triage process should provide the initial assessment and analysis and determine what if any escalation is necessary. Triage should act as the focal point for all information and funnel that information to the appropriate groups.

### *Sources for Information on CSIRP*

As stated previously, it is not practical to use a cookie cutter approach to developing a CSIRP. However, there are many sources of information available that can provide some general guidelines to assist in the development of a plan.

The CERT coordination Web site has the "Handbook for Computer Security Incident Response Teams (CSIRTs)," which can be downloaded free of charge at the URL http://www.cert.org/nav/reports.html. Another useful document is "Expectations for Computer Security Incident Response," which is available from the IETF at the URL http://www.ietf.org/rfc/rfc2350.txt?number=2350. The SANS Institute also has publications with detailed recommendations for dealing with the various phases of an computer security incident. However, there is a cost associated with the publications, and they are not inexpensive. Information regarding the SANS publications is available at the URL http://www.sans.org/newlook/publications/incident_handling.htm. These and other available sources go into much more detail than I can in this limited space.

Most importantly, organizations need to spend some time planning what to do in the event of an attack, security breach, or fraud before it occurs. The time to start thinking about what to do is not during the crisis, but before.

# Chapter 17: Cookies, Cache, and AutoComplete

Today, millions of people use the Web every day for shopping, banking, education, business, and entertainment. An essential component of that process is the Web browser. Browsers such as Netscape's Navigator and Microsoft's Internet Explorer are the end user's interface to the Web. Normally when Web surfers think of browser security, if they think of it at all, it is in reference to SSL. People are usually most concerned about the interception of information—such as credit card numbers—as it traverses the network. However, there is exposure associated with files that reside on the local disk drive of the Web surfer's PC. These files are created, accessed, and manipulated by the Web browser and various Web servers every time a Web surfer uses his or her Web browser software. Very few people are aware of the potential risks associated with these files. At the very least they raise privacy issues. At the worst they expose the Web surfer to fraud from malicious Web sites or from the fact that the files can be accessed long after one has shut down the browser and logged off the Web.

This chapter will discuss some of the basic security issues associated with Web browser software. Specifically, we will discuss the internal functions of Navigator and Internet Explorer. We will look at how they work, what to look for, and ways to protect yourself when surfing the Web.

## *Cookies*

Much has been written about cookie files and their possible uses and abuses. Essentially, cookies are text files that are stored on a Web surfer's disk drive by Web browser software, such as Navigator and Internet Explorer. Cookies are an invention of Netscape Navigator but were copied by Microsoft's Internet Explorer. The two browser software employ slightly different approaches when storing cookies. Navigator stores all cookies in a single file aptly named COOKIE.TXT. Internet Explorer creates a file for each cookie, but stores them in the WINDOWS\Cookies directory. A cookie is actually created by a Web server and passed to the browser. The cookie contains information pertaining to the Web site being visited that is stored on a PC's hard drive as a .txt file. By storing the information on the Web surfer's disk drive, Web sites avoid having to store information on their servers. Basically, cookies are used to track what sites you've visited. A cookie can be used to track the number of times you visit a site, to store the personal perference setting for a particular site, or to hold your authentication for a particular site.

For example, if the cookie file is tracking the number of times you visit a particular site, the cookie would be opened each time you browse the site and the counter stored in the cookie would be increased by one. This type of information can be useful for marketers and Webmasters to determine the number of return visitors a site receives. When customizing your preferences for sites like Yahoo, a cookie file would be created and stored on your PC's hard drive. The cookie information would be used to customize banner information that addressed your particular areas of interest or to determine whether to use frames or no-frames.

Each time you visited the Web site, the Web server would open the cookie file and modify the information displayed according to the preferences stored in the cookie. When a Web surfer uses a site that requires a registration process, such as the NY Times or Amazon.com Web sites, the authentication information is written to a cookie and stored on his or her PC. With some cookies this can include usernames and passwords. The information stored in the cookie
file is read by the Web server each time the Web surfer visits the site. The purpose is to save the end user the trouble of having to enter in the information each time he or she visits the

Web site. Theoretically, only the Web site that created the cookie information can read or modify that information. In addition, not every Web site employs cookies, so you will not automatically receive a cookie when visiting a site.

Figure 17.1 shows the various cookies stored on my PC's disk drive under the WINDOWS\Cookies directory. In this case the files were stored by IE5. The origins of many of them are easy to guess by their names. Others are more obscure. I have highlighted the cookie issued by the Amazon.com Web site. This cookie is used to authenticate my identity to the Amazon.com server during each visit to its Web site.

Figure 17.1: IE5 cookie files.

For example, in the past Amazon.com has used cookies to identify return visitors to its Web site. Essentially, the purpose of the cookie was to identify a returning customer without requiring the entry of information. If you were a first time visitor to the Amazon.com Web site, the page would display "Hello! New to Amazon.com?" If however you were a returning customer to the Amazon site it could identify you. For example, as a returning customer, I would simply type in the URL www.amazon.com and up would pop the Amazon.com home page saying "Hello, John E. Canavan." The site was able to identify me by the information that it had stored in the cookie on my disk drive at the time I originally signed up with Amazon.com. If I were to delete the cookie, the Amazon.com Web site would treat me as a first-time visitor. I had hoped to provide examples of the Amazon.com Web pages in action, but Amazon.com wouldn't allow me to use images of their Web pages.

Figure 17.2 shows the content of the COOKIES.TXT file that resides on my PC's disk drive. This is the file that is created by Netscape's Navigator. Navigator stores all of its cookies in a single file. Looking at Figure 17.2 you can see the sites that I have visited using Navigator include the NY Times, Double Click, Yahoo, and Netscape. The other information stored in the cookie files will vary from site to site.

Figure 17.2: Navigator's cookie.txt file.

The benefit of cookies files to Web surfers is dubious at best. Cookie files were developed for marketing and tracking purposes. Even for that purpose, their value is questionable because end users can access, change, delete, or block cookie information. However, the average Web surfer is neither aware of the existence of cookie files nor of the ability to manipulate them.

The risks associated with cookie files are obvious. Cookies can be an invasion of one's privacy. They basically let people know where you've been on the Web. If someone is able to access your PC they can view the cookies. If you've been to the Playboy Web site and if it issues cookies (I really don't know. Honest!), then playboy.com would clearly show in your cookies. I guess you could always say that you visited the site to read the articles.

The fact that cookies reveal this kind of information can be viewed as a good thing or a bad thing. Parents can use cookies to ensure that children aren't accessing inappropriate sites, and companies can use them to check for policy violations. However, in either case it would be much more effective to put preventative measures in place rather then check after the fact.

In addition to the risks associated with someone having direct access to your PC there are risks from malicious Web sites. I stated earlier that "theoretically" only the Web site that created or issued the cookie should be able to access it. However, the reality is that from time to time, vulnerabilities are found in Navigator and Internet Explorer that allow for malicious Web sites to read, alter, and delete cookies issued by other Web sites. As a result, it is very dangerous for Web sites to store sensitive information in cookie files.

Imagine the risk associated with a Web server that uses a cookie to store credit card information for on-line transactions. Several free Web-based e-mail services use cookies for authentication. As a result, the operator of a malicious Web site has the potential to access the cookie, determine the e-mail service provider, steal the password, and access the e-mail account. At the very least, someone would be able to identify the owner of the cookie by visiting the issuing Web site. Referring back to the example of Amazon.com, if a hacker Web site had access to the cookie issued to me by Amazon, they could go to the site and up would pop my name.

There is also risk associated with using a site that simply allows the option of using cookies for authentication. I know of at least one on-line stock trading service that gives customers the
option to use cookies as part of the authentication. I have seen first hand an instance where an end user went to the trading site and up popped someone else's name and account information.

It is quite possible that the individual whose information was erroneously displayed did the correct thing and did not enable the cookies option, but the trading Web site read some information in the end user's cookie files that caused the incorrect identification. One may believe that he or she is protected by not using cookies only to find the poor design of a Web server application circumventing all precautions. In this case, the best protection is not to frequent Web sites that use cookies for authentication.

There are several options available to end users for controlling cookies. First, end users can configure your browser to prompt you before downloading a cookie. This gives end users the chance to accept or deny cookies depending on the site. They can also disable cookies altogether, so that your browser denies all cookies. This can be done with Navigator under Edit/Preferences and Advanced. For IE5 it can be found under Tools/Internet options and then clicking on the Security tab to custom level. Disabling cookies completely can cause problems, since some sites require cookies to be enabled to access the site. Unfortunately, the error message end users receive does not specifically tell them that they need to enable cookies. It is usually some nondescript message about the browser not being supported.

Another option to consider to protect yourself against this type of threat is disabling or prompting for Java, JavaScript, and ActiveX. Again, this will at least give you the option to accept or deny based upon your level of comfort with the site you are visiting. Keep in mind, however, that if you configure your browser to prompt you before accepting cookies, Java, JavaScript, and ActiveX, you will be prompted constantly.

There are also a number of utilities that can be used to monitor and control cookies. Programs such as CookiePal, CookieJar, CookieCop, McAfee's Internet Guard Dog, and Norton's Internet Security can all be employed for controlling cookies.

### Cache Files

Cookie files are not the only files created on the disk drive of a Web surfer's PC when visiting a Web site. Both Netscape's Navigator and Microsoft's Internet Explorer also "cache" or store files that have recently been accessed at Web sites. This is referred to as caching, and the files stored on a PC's disk drive are referred to as cache files. The advantage to caching comes into play when reloading Web pages that have recently been viewed. The browser is able to load a page from the files cached on the local disk drive rather than having to reload it from the server. Loading files from the local drive is much faster than loading them over the network from the server. The cached files can include html files, graphic or image files such as gifs and jpgs, and text files.

The risks associated with cached files are similar to those associated with cookie files. Basically, they tell someone where you've been. However, viewing the cache files cannot only
tell someone where you've been browsing, but they can actually allow them to view the files you browsed, since the html files along with their associated graphic or image files are cached
on your hard drive when you browse most Web pages.

Figure 17.3 is an example of the files cached on my disk drive by Netscape's Navigator. The files are cached in a subdirectory aptly named cache and have nondescript names, but the extensions identify the file types. The files are easily viewed by simply clicking on them. Internet Explorer also caches files, but IE stores cache files under the temporary Internet files directory, which is a subdirectory of the Windows directory.

Figure 17.3: IE5 cache files.

The fact that browsers cache the pages viewed on-line should be remembered when using the Internet to access bank and brokerage accounts for financial transactions. For example, all of the pages viewed when using an Internet banking system are stored in the PC's cached files. Those cache files can include account numbers and balances. If you use your Internet banking system on a PC that is shared by others, say at work, then anyone with physical access to that PC can potentially view information regarding your accounts.

Using Windows Explorer to access the directory where the files reside you can retrieve the cache files. In addition, Internet Explorer provides the ability to view cached files. Under Tools/Internet Options under the General tab you will find temporary Internet files. There are two buttons: One is delete files and the other is settings. This is depicted in Figure 17.4. If you click on settings, there is an option to view files.

Figure 17.4: IE5 Internet options.

Figure 17.5 shows the Internet Explorer cache files that reside on my PC's disk drive. It shows the file types and the location from which they originated. Figure 17.5 shows html, gif, and jpg files from the Yahoo and NASA Web sites.

Figure 17.5: Cache files.

In addition to the risk from individuals with physical access to your PC being able to retrieve information from your cache files, there is also the risk of someone doing the same over the network. One way this can be accomplished is if you are sharing your disk drive. Another would be through vulnerabilities that have from time to time been identified with both Navigator and Explorer. These include the Cache-Cow vulnerability, which affected Netscape Communicator 4.05, and the more recently discovered vulnerability for both IE and Navigator that involves using cookies to run JavaScript that can grab cache files and even html bookmarks. These kinds of vulnerabilities could allow a malicious Web site to grab or view information in the cache files that reside on your PC's disk drive.

When using Internet Explorer you can mitigate the risk of exposing confidential information in cache files by configuring the browser not to cache secure pages. In other words, if the page uses SSL with HTTPS then the browser will not cache the files to the PC's disk drive. This is accomplished by going into Tools/Internet Options and clicking on the Advanced tab. Scroll down until you find "Do not save encrypted pages to disk" and click on the box next to the option. This will prevent SSL pages from being cached to your hard drive. Figure 17.6 shows the Internet options screen with the appropriate option highlighted.

Figure 17.6: IE5 Internet options.

Another option is to simply delete the cached files before exiting Navigator or Internet Explorer. For Internet Explorer this can be done under Internet option. Referring back to Figure 17.4 you can see that there is a delete files button under the temporary Internet files section. Simply clicking on this button deletes all cached files. With Netscape the cache files are deleted by going into Edit/Preferences and clicking on the button clear disk cache. This screen is depicted in Figure 17.7.

Figure 17.7: IE5 preferences.

With Internet Explorer, you also have the option of configuring the browser to delete cache files automatically when exiting the software. Referring to Figure 17.6, the option directly under "Do not save encrypted pages to disk" is the option "Delete temporary Internet files folder when browser is closed." By checking this option IE will delete the files cached to the hard drive during the browser session at the time the session ends.

Since most of the vulnerabilities associated with grabbing or viewing cached information over
the network involves JavaScript, you should consider disabling the function. However, many reputable Web sites employ JavaScript for legitimate purposes, so there is a tradeoff with this option.

## *AutoComplete*

Another concern with using Internet Explore 5.0 in an environment where a PC might be shared with or used by others is the option of AutoComplete. The AutoComplete option can be found under Tools/Internet Options by clicking on the content tab. Figure 17.8 illustrates the AutoComplete settings box. The concern here is the option "user names and passwords on forms." If that box is checked, IE5 will store sign-on account information for those sites that require authentication, including Internet banking and brokerage Web sites. The risk associated with this is obvious.

Figure 17.8: IE5 AutoComplete settings.

Figure 17.9 illustrates how AutoComplete functions. The example again employs the Internet banking system of the fictitious Anybank Corporation. In Figure 17.9, IE5 has been configured to autocomplete user names and passwords. In addition, I have already logged into
the on-line banking system using two different account numbers and have shut down and reopened IE5 to demonstrate the effects of AutoComplete. When I logged into the Anybank Internet banking system, IE5 recorded the account information. During subsequent attempts to
log in to the Internet banking system I need only enter the first number of an account, and IE5 will display the entire account number for all accounts beginning with the number entered.

Figure 17.9: An example of AutoComplete.

In Figure 17.9 when the number 5 is entered into the account number field a drop-down menu appears displaying two account numbers, 55000037390 and 59001260504. These are the two account numbers that I previously had used to login to the Anybank Internet banking system. IE5 stored these account numbers and passwords and associated them with the Anybank Internet URL. When IE5 is configured in this manner, the end user need only highlight the account number, and the password is automatically entered.

Obviously, AutoComplete for usernames and passwords should not be enabled. This is especially true when working in an environment where PCs may be shared. I would also recommend against using the AutoComplete function on a laptop. If the laptop is lost or stolen, it could be used to access on-line accounts. If AutoComplete for usernames and passwords is enabled then access to accounts can be easily obtained. When IE5 was first released we found that AutoComplete was enabled by default. I believe that has since changed.