Software Engineering Institute
Carnegie Mellon University

Security of the Internet

Thomas A. Longstaff
James T. Ellis
Shawn V. Hernan
Howard F. Lipson
Robert D. McMillan
Linda Hutz Pesante
Derrick Simmel

CERT Division

[DISTRIBUTION STATEMENT A] Approved for public release and unlimited distribution.

http://www.sei.cmu.edu

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

The view, opinions, and/or findings contained in this material are those of the author(s) and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.

References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by Carnegie Mellon University or its Software Engineering Institute.

This report was prepared for the

SEI Administrative Agent
AFLCMC/AZS
5 Eglin Street
Hanscom AFB, MA 01731-2100

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

DM17-0052

Table of Contents

1 Overview of Internet Security

As of 1996, the Internet connected an estimated 13 million computers in 195 countries on every continent, even Antarctica (1). The Internet is not a single network, but a worldwide collection of loosely connected networks that are accessible by individual computer hosts in a variety of ways, including gateways, routers, dial-up connections, and Internet service providers. The Internet is easily accessible to anyone with a computer and a network connection. Individuals and organizations worldwide can reach any point on the network without regard to national or geographic boundaries or time of day.

However, along with the convenience and easy access to information come new risks. Among them are the risks that valuable information will be lost, stolen, corrupted, or misused and that the computer systems will be corrupted. If information is recorded electronically and is available on networked computers, it is more vulnerable than if the same information is printed on paper and locked in a file cabinet. Intruders do not need to enter an office or home, and may not even be in the same country. They can steal or tamper with information without touching a piece of paper or a photocopier. They can create new electronic files, run their own programs, and hide evidence of their unauthorized activity.

Basic Security Concepts

Three basic security concepts important to information on the Internet are confidentiality, integrity, and availability. Concepts relating to the people who use that information are authentication, authorization, and nonrepudiation.

When information is read or copied by someone not authorized to do so, the result is known as *loss of confidentiality*. For some types of information, confidentiality is a very important attribute. Examples include research data, medical and insurance records, new product specifications, and corporate investment strategies. In some locations, there may be a legal obligation to protect the privacy of individuals. This is particularly true for banks and loan companies; debt collectors; businesses that extend credit to their customers or issue credit cards; hospitals, doctors' offices, and medical testing laboratories; individuals or agencies that offer services such as psychological counseling or drug treatment; and agencies that collect taxes.

Information can be corrupted when it is available on an insecure network. When information is modified in unexpected ways, the result is known as *loss of integrity*. This means that unauthorized changes are made to information, whether by human error or intentional tampering. Integrity is particularly important for critical safety and financial data used for activities such as electronic funds transfers, air traffic control, and financial accounting.

Information can be erased or become inaccessible, resulting in *loss of availability*. This means that people who are authorized to get information cannot get what they need.

Availability is often the most important attribute in service-oriented businesses that depend on information (e.g., airline schedules and online inventory systems). Availability of the network itself is important to anyone whose business or education relies on a network connection. When a user cannot get access to the network or specific services provided on the network, they experience a *denial of service.*

To make information available to those who need it and who can be trusted with it, organizations use authentication and authorization. *Authentication* is proving that a user is whom he or she claims to be. That proof may involve something the user knows (such as a password), something the user has (such as a "smartcard"), or something about the user that proves the person's identity (such as a fingerprint). *Authorization* is the act of determining whether a particular user (or computer system) has the right to carry out a certain activity, such as reading a file or running a program. Authentication and authorization go hand in hand. Users must be authenticated before carrying out the activity they are authorized to perform. Security is strong when the means of authentication cannot later be refuted - the user cannot later deny that he or she performed the activity. This is known as *nonrepudiation.*

Why Care About Security?

It is remarkably easy to gain unauthorized access to information in an insecure networked environment, and it is hard to catch the intruders. Even if users have nothing stored on their computer that they consider important, that computer can be a "weak link", allowing unauthorized access to the organization's systems and information.

Seemingly innocuous information can expose a computer system to compromise. Information that intruders find useful includes which hardware and software are being used, system configuration, type of network connections, phone numbers, and access and authentication procedures. Security-related information can enable unauthorized individuals to get access to important files and programs, thus compromising the security of the system. Examples of important information are passwords, access control files and keys, personnel information, and encryption algorithms.

Judging from CERT® Coordination Center (CERT/CC) data and the computer abuse reported in the media, no one on the Internet is immune. Those affected include banks and financial companies, insurance companies, brokerage houses, consultants, government contractors, government agencies, hospitals and medical laboratories, network service providers, utility companies, the textile business, universities, and wholesale and retail trades.

The consequences of a break-in cover a broad range of possibilities: a minor loss of time in recovering from the problem, a decrease in productivity, a significant loss of money or staff-hours, a devastating loss of credibility or market opportunity, a business no longer able to compete, legal liability, and the loss of life.

2 History

The Internet began in 1969 as the ARPANET, a project funded by the Advanced Research Projects Agency (ARPA) of the U.S. Department of Defense. One of the original goals of the project was to create a network that would continue to function even if major sections of the network failed or were attacked. The ARPANET was designed to reroute network traffic automatically around problems in connecting systems or in passing along the necessary information to keep the network functioning. Thus, from the beginning, the Internet was designed to be robust against denial-of-service attacks, which are described in a section below on denial of service.

The ARPANET protocols (the rules of syntax that enable computers to communicate on a network) were originally designed for openness and flexibility, not for security. The ARPA researchers needed to share information easily, so everyone needed to be an unrestricted "insider" on the network. Although the approach was appropriate at the time, it is not one that lends itself to today's commercial and government use.

As more locations with computers (known as *sites* in Internet parlance) joined the ARPANET, the usefulness of the network grew. The ARPANET consisted primarily of university and government computers, and the applications supported on this network were simple: electronic mail (E-mail), electronic news groups, and remote connection to other computers. By 1971, the Internet linked about two dozen research and government sites, and researchers had begun to use it to exchange information not directly related to the ARPANET itself. The network was becoming an important tool for collaborative research.

During these years, researchers also played "practical jokes" on each other using the ARPANET. These jokes usually involved joke messages, annoying messages, and other minor security violations. Some of these are described in Steven Levy's *Hackers: Heroes of the Computer Revolution* (2). It was rare that a connection from a remote system was considered an attack, however, because ARPANET users comprised a small group of people who generally knew and trusted each other.

In 1986, the first well-publicized international security incident was identified by Cliff Stoll, then of Lawrence Berkeley National Laboratory in northern California. A simple accounting error in the computer records of systems connected to the ARPANET led Stoll to uncover an international effort, using the network, to connect to computers in the United States and copy information from them. These U.S. computers were not only at universities, but at military and government sites all over the country. When Stoll published his experience in a 1989 book, *The Cuckoo's Egg* (3), he raised awareness that the ARPANET could be used for destructive purposes.

In 1988, the ARPANET had its first automated network security incident, usually referred to as "the Morris worm" (4). A student at Cornell University (Ithaca, NY), Robert T. Morris, wrote a program that would connect to another computer, find and use one of several vulnerabilities to copy itself to that second computer, and begin to run the copy of itself at the new location. Both the original code and the copy would then repeat these actions in an infinite loop to other computers on the ARPANET. This "self-replicating automated network attack tool" caused a geometric explosion of copies to be started at computers all around the ARPANET. The worm used so many system resources that the attacked computers could no longer function. As a result, 10% of the U.S. computers connected to the ARPANET effectively stopped at about the same time.

By that time, the ARPANET had grown to more than 88,000 computers and was the primary means of communication among network security experts. With the ARPANET effectively down, it was difficult to coordinate a response to the worm. Many sites removed themselves from the ARPANET altogether, further hampering communication and the transmission of the solution that would stop the worm.

The Morris worm prompted the Defense Advanced Research Projects Agency (DARPA, the new name for ARPA) to fund a computer emergency response team, now the CERT® Coordination Center, to give experts a central point for coordinating responses to network emergencies. Other teams quickly sprang up to address computer security incidents in specific organizations or geographic regions. Within a year of their formation, these incident response teams created an informal organization now known as the Forum of Incident Response and Security Teams (FIRST). These teams and the FIRST organization exist to coordinate responses to computer security incidents, assist sites in handling attacks, and educate network users about computer security threats and preventive practices.

In 1989, the ARPANET officially became the Internet and moved from a government research project to an operational network; by then it had grown to more than 100,000 computers. Security problems continued, with both aggressive and defensive technologies becoming more sophisticated. Among the major security incidents (5) were the 1989 WANK/OILZ worm, an automated attack on VMS systems attached to the Internet, and exploitation of vulnerabilities in widely distributed programs such as the sendmail program, a complicated program commonly found on UNIX-based systems for sending and receiving electronic mail. In 1994, intruder tools were created to "sniff" packets from the network easily, resulting in the widespread disclosure of user names and password information. In 1995, the method that Internet computers use to name and authenticate each other was exploited by a new set of attack tools that allowed widespread Internet attacks on computers that have trust relationships (see the section on exploitation of trust, below) with any other computer, even one in the same room. Today the use of the World Wide Web and Web-related programming languages create new opportunities for network attacks.

Although the Internet was originally conceived of and designed as a research and education network, usage patterns have radically changed. The Internet has become a home for private and commercial communication, and at this writing it is still expanding into important areas of commerce, medicine, and public service. Increased reliance on the Internet is expected over the next five years, along with increased attention to its security.

3 Network Security Incidents

A *network security incident* is any network-related activity with negative security implications. This usually means that the activity violates an explicit or implicit security policy (see the section on security policy). Incidents come in all shapes and sizes. They can come from anywhere on the Internet, although some attacks must be launched from specific systems or networks and some require access to special accounts. An intrusion may be a comparatively minor event involving a single site or a major event in which tens of thousands of sites are compromised. (When reading accounts of incidents, note that different groups may use different criteria for determining the bounds of an incident.)

A typical attack pattern consists of gaining access to a user's account, gaining privileged access, and using the victim's system as a launch platform for attacks on other sites. It is possible to accomplish all these steps manually in as little as 45 seconds; with automation, the time decreases further.

Sources of Incidents

It is difficult to characterize the people who cause incidents. An intruder may be an adolescent who is curious about what he or she can do on the Internet, a college student who has created a new software tool, an individual seeking personal gain, or a paid "spy" seeking information for the economic advantage of a corporation or foreign country. An incident may also be caused by a disgruntled former employee or a consultant who gained network information while working with a company. An intruder may seek entertainment, intellectual challenge, a sense of power, political attention, or financial gain.

One characteristic of the intruder community as a whole is its communication. There are electronic newsgroups and print publications on the latest intrusion techniques, as well as conferences on the topic. Intruders identify and publicize misconfigured systems; they use those systems to exchange pirated software, credit card numbers, exploitation programs, and the identity of sites that have been compromised, including account names and passwords. By sharing knowledge and easy-to-use software tools, successful intruders increase their number and their impact.

Types of Incidents

Incidents can be broadly classified into several kinds: the probe, scan, account compromise, root compromise, packet sniffer, denial of service, exploitation of trust, malicious code, and Internet infrastructure attacks.

Probe

A probe is characterized by unusual attempts to gain access to a system or to discover information about the system. One example is an attempt to log in to an unused account. Probing is the electronic equivalent of testing doorknobs to find an unlocked door for easy entry. Probes are sometimes followed by a more serious security event, but they are often the result of curiosity or confusion.

Scan

A scan is simply a large number of probes done using an automated tool. Scans can sometimes be the result of a misconfiguration or other error, but they are often a prelude to a more directed attack on systems that the intruder has found to be vulnerable.

Account Compromise

An account compromise is the unauthorized use of a computer account by someone other than the account owner, without involving system-level or root-level privileges (privileges a system administrator or network manager has). An account compromise might expose the victim to serious data loss, data theft, or theft of services. The lack of root-level access means that the damage can usually be contained, but a user-level account is often an entry point for greater access to the system.

Root Compromise

A root compromise is similar to an account compromise, except that the account that has been compromised has special privileges on the system. The term *root* is derived from an account on UNIX systems that typically has unlimited, or "superuser", privileges. Intruders who succeed in a root compromise can do just about anything on the victim's system, including run their own programs, change how the system works, and hide traces of their intrusion.

Packet Sniffer

A packet sniffer is a program that captures data from information packets as they travel over the network. That data may include user names, passwords, and proprietary information that travels over the network in clear text. With perhaps hundreds or thousands of passwords captured by the sniffer, intruders can launch widespread attacks on systems. Installing a packet sniffer does not necessarily require privileged access. For most multi-user systems, however, the presence of a packet sniffer implies there has been a root compromise.

Denial of Service

The goal of denial-of-service attacks is not to gain unauthorized access to machines or data, but to prevent legitimate users of a service from using it. A denial-of-service attack can come in many forms. Attackers may "flood" a network with large volumes of data or deliberately consume a scarce or limited resource, such as process control blocks or pending network connections. They may also disrupt physical components of the network or manipulate data in transit, including encrypted data.

Exploitation of Trust

Computers on networks often have trust relationships with one another. For example, before executing some commands, the computer checks a set of files that specify which other computers on the network are permitted to use those commands. If attackers can forge their identity, appearing to be using the trusted computer, they may be able to gain unauthorized access to other computers.

Malicious Code

Malicious code is a general term for programs that, when executed, would cause undesired results on a system. Users of the system usually are not aware of the program until they discover the damage. Malicious code includes Trojan horses, viruses, and worms. Trojan horses and viruses are usually hidden in legitimate programs or files that attackers have altered to do more than what is expected. Worms are self-replicating programs that spread with no human intervention after they are started. Viruses are also self-replicating programs, but usually require some action on the part of the user to spread inadvertently to other programs or systems. These sorts of programs can lead to serious data loss, downtime, denial of service, and other types of security incidents.

Internet Infrastructure Attacks

These rare but serious attacks involve key components of the Internet infrastructure rather than specific systems on the Internet. Examples are network name servers, network access providers, and large archive sites on which many users depend. Widespread automated attacks can also threaten the infrastructure. Infrastructure attacks affect a large portion of the Internet and can seriously hinder the day-to-day operation of many sites.

Incidents and Internet Growth

Since the CERT® Coordination Center began operating in 1988, the number of security incidents reported to the center has grown dramatically, from less than 100 in 1988 to almost 2,500 in 1995, the last year for which complete statistics are available as of this writing. Through 1994, the increase in incident reports roughly parallels the growth of the size of the Internet during that time. Figure 1 shows the growth of the Internet and the corresponding growth of reported security incidents.

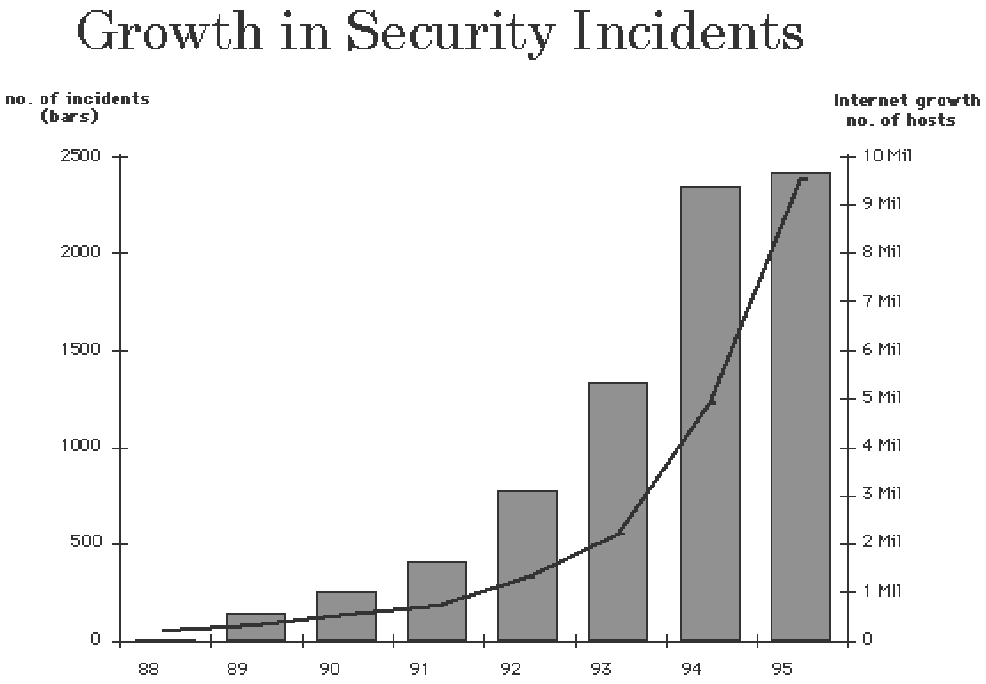

Figure 1

The data for 1995 and partial data for 1996 show a slowing of the rate at which incidents are reported to the CERT/CC (perhaps because of sites' increased security efforts or the significant increase in other response teams formed to handle incidents). However, the rate continues to increase for serious incidents, such as root compromises, services outages, and packet sniffers.

Incident Trends

In the late 1980s and early 1990s, the typical intrusion was fairly straightforward. Intruders most often exploited relatively simple weaknesses, such as poor passwords and misconfigured systems, that allowed greater access to the system than was intended. Once on a system, the intruders exploited one or another well-known, but usually unfixed, vulnerability to gain privileged access, enabling them to use the system as they wished.

There was little need to be more sophisticated because these simple techniques were effective. Vendors delivered systems with default settings that made it easy to break into systems. Configuring systems in a secure manner was not straightforward, and many system administrators did not have the time, expertise, or tools to monitor their systems adequately for intruder activity.

Unfortunately, all these activities continue in 1996; however, more sophisticated intrusions are now common. In eight years of operation, the CERT Coordination Center has seen intruders demonstrate increased technical knowledge, develop new ways to exploit system vulnerabilities, and create software tools to automate attacks. At the same time, intruders with little technical knowledge are becoming more effective as the sophisticated intruders share their knowledge and tools.

Intruders' Technical Knowledge

Intruders are demonstrating increased understanding of network topology, operations, and protocols, resulting in the infrastructure attacks described in the previous section on Internet infrastructure attacks.

Instead of simply exploiting well-known vulnerabilities, intruders examine source code to discover weaknesses in certain programs, such as those used for electronic mail. Much source code is easy to obtain from programmers who make their work freely available on the Internet. Programs written for research purposes (with little thought for security) or written by naive programmers become widely used, with source code available to all. Moreover, the targets of many computer intrusions are organizations that maintain copies of proprietary source code (often the source code to computer operating systems or key software utilities). Once intruders gain access, they can examine this code to discover weaknesses.

Intruders keep up with new technology. For example, intruders now exploit vulnerabilities associated with the World Wide Web to gain unauthorized access to systems.

Other aspects of the new sophistication of intruders include the targeting of the network infrastructure (such as network routers and firewalls) and the ability to cloak their behavior. Intruders use Trojan horses to hide their activity from network administrators; for example, intruders alter authentication and logging programs so that they can log in without the activity showing up in the system logs. Intruders also encrypt output from their activity, such as the information captured by packet sniffers. Even if the victim finds the sniffer logs, it is difficult or impossible to determine what information was compromised.

Techniques to Exploit Vulnerabilities

As intruders become more sophisticated, they identify new and increasingly complex methods of attack. For example, intruders are developing sophisticated techniques to monitor the Internet for new connections. Newly connected systems are often not fully configured from a security perspective and are, therefore, vulnerable to attacks.

The most widely publicized of the newer types of intrusion is the use of the packet sniffers described in the section above on packet sniffers. Other tools are used to construct packets with forged addresses; one use of these tools is to mount a denial-of-service attack in a way that obscures the source of the attack. Intruders also "spoof" computer addresses, masking their real identity and successfully making connections that would not otherwise be permitted. In this way, they exploit trust relationships between computers.

With their sophisticated technical knowledge and understanding of the network, intruders are increasingly exploiting network interconnections. They move through the Internet infrastructure, attacking areas on which many people and systems depend. Infrastructure attacks are even more threatening because legitimate network managers and administrators typically think about protecting systems and parts of the infrastructure rather than the infrastructure as a whole.

In the first quarter of 1996, 7.5% of 346 incidents handled by the CERT Coordination Center involved these new and sophisticated methods, including packet sniffers, spoofing, and infrastructure attacks. A full 20% involved the total compromise of systems, in which intruders gain system-level, or root, privileges. This represents a significant increase in such attacks over previous years' attacks, and the numbers are still rising. Of 341 incidents in the third quarter of 1996, nearly 9% involved sophisticated attacks, and root compromises accounted for 33%.

Intruders' Use of Software Tools

The tools available to launch an attack have become more effective, easier to use, and more accessible to people without an in-depth knowledge of computer systems. Often a sophisticated intruder embeds an attack procedure in a program and widely distributes it to the intruder community. Thus, people who have the desire but not the technical skill are able to break into systems. Indeed, there have been instances of intruders breaking into a UNIX system using a relatively sophisticated attack and then attempting to run DOS commands (commands that apply to an entirely different operating system).

Tools are available to examine programs for vulnerabilities even in the absence of source code. Though these tools can help system administrators identify problems, they also help intruders find new ways to break into systems.

As in many areas of computing, the tools used by intruders have become more automated, allowing intruders to gather information about thousands of Internet hosts quickly and with minimum effort. These tools can scan entire networks from a remote location and identify individual hosts with specific weaknesses. Intruders may catalog the information for later exploitation, share or trade with other intruders, or attack immediately. The increased availability and usability of scanning tools means that even technically naive, would-be intruders can find new sites and particular vulnerabilities.

Some tools automate multiphase attacks in which several small components are combined to achieve a particular end. For example, intruders can use a tool to mount a denial-of-service attack on a machine and spoof that machine's address to subvert the intended victim's machine. A second example is using a packet sniffer to get router or firewall passwords, logging in to the firewall to disable filters, then using a network file service to read data on an otherwise secure server.

The trend toward automation can be seen in the distribution of software packages containing a variety of tools to exploit vulnerabilities. These packages are often maintained by competent programmers and are distributed complete with version numbers and documentation.

A typical tool package might include the following:

- network scanner
- password cracking tool and large dictionaries
- packet sniffer
- variety of Trojan horse programs and libraries
- tools for selectively modifying system log files

- tools to conceal current activity
- tools for automatically modifying system configuration files
- tools for reporting bogus checksums

4 Internet Vulnerabilities

A vulnerability is a weakness that a person can exploit to accomplish something that is not authorized or intended as legitimate use of a network or system. When a vulnerability is exploited to compromise the security of systems or information on those systems, the result is a security incident. Vulnerabilities may be caused by engineering or design errors, or faulty implementation.

Why the Internet Is Vulnerable

Many early network protocols that now form part of the Internet infrastructure were designed without security in mind. Without a fundamentally secure infrastructure, network defense becomes more difficult. Furthermore, the Internet is an extremely dynamic environment, in terms of both topology and emerging technology.

Because of the inherent openness of the Internet and the original design of the protocols, Internet attacks in general are quick, easy, inexpensive, and may be hard to detect or trace. An attacker does not have to be physically present to carry out the attack. In fact, many attacks can be launched readily from anywhere in the world - and the location of the attacker can easily be hidden. Nor is it always necessary to "break in" to a site (gain privileges on it) to compromise confidentiality, integrity, or availability of its information or service.

Even so, many sites place unwarranted trust in the Internet. It is common for sites to be unaware of the risks or unconcerned about the amount of trust they place in the Internet. They may not be aware of what can happen to their information and systems. They may believe that their site will not be a target or that precautions they have taken are sufficient. Because the technology is constantly changing and intruders are constantly developing new tools and techniques, solutions do not remain effective indefinitely.

Since much of the traffic on the Internet is not encrypted, confidentiality and integrity are difficult to achieve. This situation undermines not only applications (such as financial applications that are network-based) but also more fundamental mechanisms such as authentication and nonrepudiation (see the section on basic security concepts for definitions). As a result, sites may be affected by a security compromise at another site over which they have no control. An example of this is a packet sniffer that is installed at one site but allows the intruder to gather information about other domains (possibly in other countries).

Another factor that contributes to the vulnerability of the Internet is the rapid growth and use of the network, accompanied by rapid deployment of network services involving complex applications. Often, these services are not designed, configured, or maintained securely. In the rush to get new products to market, developers do not adequately ensure that they do not repeat previous mistakes or introduce new vulnerabilities.

Compounding the problem, operating system security is rarely a purchase criterion. Commercial operating system vendors often report that sales are driven by customer demand for performance, price, ease of use, maintenance, and support. As a result, off-the-shelf operating systems are shipped in an easy-to-use but insecure configuration that allows sites to use the system soon after installation. These hosts/sites are often not fully configured from a security perspective before connecting. This lack of secure configuration makes them vulnerable to attacks, which sometimes occur within minutes of connection.

Finally, the explosive growth of the Internet has expanded the need for well-trained and experienced people to engineer and manage the network in a secure manner. Because the need for network security experts far exceeds the supply, inexperienced people are called upon to secure systems, opening windows of opportunity for the intruder community.

Types of Technical Vulnerabilities

The following taxonomy is useful in understanding the technical causes behind successful intrusion techniques, and helps experts identify general solutions for addressing each type of problem.

Flaws in Software or Protocol Designs

Protocols define the rules and conventions for computers to communicate on a network. If a protocol has a fundamental design flaw, it is vulnerable to exploitation no matter how well it is implemented. An example of this is the Network File System (NFS), which allows systems to share files. This protocol does not include a provision for authentication; that is, there is no way of verifying that a person logging in really is whom he or she claims to be. NFS servers are targets for the intruder community.

When software is designed or specified, often security is left out of the initial description and is later "added on" to the system. Because the additional components were not part of the original design, the software may not behave as planned and unexpected vulnerabilities may be present.

Weaknesses in How Protocols and Software Are Implemented

Even when a protocol is well designed, it can be vulnerable because of the way it is implemented. For example, a protocol for electronic mail may be implemented in a way that permits intruders to connect to the mail port of the victim's machine and fool the machine into performing a task not intended by the service. If intruders supply certain data for the "To:" field instead of a correct E-mail address, they may be able to fool the machine into sending them user and password information or granting them access to the victim's machine with privileges to read protected files or run programs on the system. This type of vulnerability enables intruders to attack the victim's machine from remote sites without access to an account on the victim's system. This type of attack often is just a first step, leading to the exploitation of flaws in system or application software.

Software may be vulnerable because of flaws that were not identified before the software was released. This type of vulnerability has a wide range of subclasses, which intruders often exploit using their own attack tools. For readers who are familiar with software design, the following examples of subclasses are included:

- race conditions in file access
- non-existent checking of data content and size
- non-existent checking for success or failure
- inability to adapt to resource exhaustion
- incomplete checking of operating environment
- inappropriate use of system calls
- re-use of software modules for purposes other than their intended ones

By exploiting program weaknesses, intruders at a remote site can gain access to a victim's system. Even if they have access to a nonprivileged user account on the victim's system, they can often gain additional, unauthorized privileges.

Weaknesses in System and Network Configurations

Vulnerabilities in the category of system and network configurations are not caused by problems inherent in protocols or software programs. Rather, the vulnerabilities are a result of the way these components are set up and used. Products may be delivered with default settings that intruders can exploit. System administrators and users may neglect to change the default settings, or they may simply set up their system to operate in a way that leaves the network vulnerable.

An example of a faulty configuration that has been exploited is anonymous File Transfer Protocol (FTP) service. Secure configuration guidelines for this service stress the need to ensure that the password file, archive tree, and ancillary software are separate from the rest of the operating system, and that the operating system cannot be reached from this staging area. When sites misconfigure their anonymous FTP archives, unauthorized users can get authentication information and use it to compromise the system.

5 Improving Security

In the face of the vulnerabilities and incident trends discussed above, a robust defense requires a flexible strategy that allows adaptation to the changing environment, well-defined policies and procedures, the use of robust tools, and constant vigilance.

It is helpful to begin a security improvement program by determining the current state of security at the site. Methods for making this determination in a reliable way are becoming available. Integral to a security program are documented policies and procedures, and technology that supports their implementation.

Security Policy, Procedures, and Practices

Security Policy

A policy is a documented high-level plan for organization-wide computer and information security. It provides a framework for making specific decisions, such as which defense mechanisms to use and how to configure services, and is the basis for developing secure programming guidelines and procedures for users and system administrators to follow. Because a security policy is a long-term document, the contents avoid technology-specific issues.

A security policy covers the following (among other topics appropriate to the organization):

- high-level description of the technical environment of the site, the legal environment (governing laws), the authority of the policy, and the basic philosophy to be used when interpreting the policy
- risk analysis that identifies the site's assets, the threats that exist against those assets, and the costs of asset loss
- guidelines for system administrators on how to manage systems
- definition of acceptable use for users
- guidelines for reacting to a site compromise (e.g., how to deal with the media and law enforcement, and whether to trace the intruder or shutdown and rebuild the system)

Factors that contribute to the success of a security policy include management commitment, technological support for enforcing the policy, effective dissemination of the policy, and the security awareness of all users. Management assigns responsibility for security, provides training for security personnel, and allocates funds to security. Technological support for the security policy moves some responsibility for enforcement from individuals to technology. The result is an automatic and consistent enforcement of policies, such as those for access and authentication. Technical options that support policy include (but are not limited to)

- challenge/response systems for authentication
- auditing systems for accountability and event reconstruction

- encryption systems for the confidential storage and transmission of data
- network tools such as firewalls and proxy servers

There are many books and papers devoted to site security policies, including requests for comments RFC 1244 (6) and RFC 1281 (7), guidelines written by the Internet Engineering Task Force.

Security-Related Procedures

Procedures are specific steps to follow that are based on the computer security policy. Procedures address such topics as retrieving programs from the network, connecting to the site's system from home or while traveling, using encryption, authentication for issuing accounts, configuration, and monitoring.

Security Practices

System administration practices play a key role in network security. Checklists and general advice on good security practices are readily available. Below are examples of commonly recommended practices:

- Ensure all accounts have a password and that the passwords are difficult to guess. A one-time password system is preferable.
- Use tools such as MD5 checksums (8), a strong cryptographic technique, to ensure the integrity of system software on a regular basis.
- Use secure programming techniques when writing software. These can be found at security-related sites on the World Wide Web.
- Be vigilant in network use and configuration, making changes as vulnerabilities become known.
- Regularly check with vendors for the latest available fixes and keep systems current with upgrades and patches.
- Regularly check on-line security archives, such as those maintained by incident response teams, for security alerts and technical advice.
- Audit systems and networks, and regularly check logs. Many sites that suffer computer security incidents report that insufficient audit data is collected, so detecting and tracing an intrusion is difficult.

Security Technology

A variety of technologies have been developed to help organizations secure their systems and information against intruders. These technologies help protect systems and information against attacks, detect unusual or suspicious activities, and respond to events that affect security. In this section, the focus is on two core areas: operational technology and cryptography. The purpose of operational technology is to maintain and defend the availability of data resources in a secure manner. The purpose of cryptography is to secure the confidentiality, integrity, and authenticity of data resources.

Operational Technology

Intruders actively seek ways to access networks and hosts. Armed with knowledge about specific vulnerabilities, social engineering techniques, and tools to automate information gathering and systems infiltration, intruders can often gain entry into systems with disconcerting ease. System administrators face the dilemma of maximizing the availability of system services to valid users while minimizing the susceptibility of complex network infrastructures to attack. Unfortunately, services often depend on the same characteristics of systems and network protocols that make them susceptible to compromise by intruders. In response, technologies have evolved to reduce the impact of such threats. No single technology addresses all the problems. Nevertheless, organizations can significantly improve their resistance to attack by carefully preparing and strategically deploying personnel and operational technologies. Data resources and assets can be protected, suspicious activity can be detected and assessed, and appropriate responses can be made to security events as they occur.

One-Time Passwords. Intruders often install packet sniffers to capture passwords as they traverse networks during remote log-in processes. Therefore, all passwords should at least be encrypted as they traverse networks. A better solution is to use one-time passwords because there are times when a password is required to initiate a connection before confidentiality can be protected.

One common example occurs in remote dial-up connections. Remote users, such as those traveling on business, dial in to their organization's modem pool to access network and data resources. To identify and authenticate themselves to the dial-up server, they must enter a user ID and password. Because this initial exchange between the user and server may be monitored by intruders, it is essential that the passwords are not reusable. In other words, intruders should not be able to gain access by masquerading as a legitimate user using a password they have captured.

One-time password technologies address this problem. Remote users carry a device synchronized with software and hardware on the dial-up server. The device displays random passwords, each of which remains in effect for a limited time period (typically 60 seconds). These passwords are never repeated and are valid only for a specific user during the period that each is displayed. In addition, users are often limited to one successful use of any given password. One-time password technologies significantly reduce unauthorized entry at gateways requiring an initial password.

Firewalls. Intruders often attempt to gain access to networked systems by pretending to initiate connections from trusted hosts. They squash the emissions of the genuine host using a denial-of-service attack and then attempt to connect to a target system using the address of the genuine host. To counter these address-spoofing attacks and enforce limitations on authorized connections into the organization’s network, it is necessary to filter all incoming and outgoing network traffic.

A firewall is a collection of hardware and software designed to examine a stream of network traffic and service requests. Its purpose is to eliminate from the stream those packets or requests that fail to meet the security criteria established by the organization. A simple firewall may consist of a filtering router, configured to discard packets that arrive from unauthorized addresses or that represent attempts to connect to unauthorized service ports. More sophisticated implementations may include bastion hosts, on which proxy mechanisms operate on behalf of services. These mechanisms authenticate requests, verify their form and content, and relay approved service requests to the appropriate service hosts. Because firewalls are typically the first line of defense against intruders, their configuration must be carefully implemented and tested before connections are established between internal networks and the Internet.

Monitoring Tools. Continuous monitoring of network activity is required if a site is to maintain confidence in the security of its network and data resources. Network monitors may be installed at strategic locations to collect and examine information continuously that may indicate suspicious activity. It is possible to have automatic notifications alert system administrators when the monitor detects anomalous readings, such as a burst of activity that may indicate a denial-of-service attempt. Such notifications may use a variety of channels, including electronic mail and mobile paging. Sophisticated systems capable of reacting to questionable network activity may be implemented to disconnect and block suspect connections, limit or disable affected services, isolate affected systems, and collect evidence for subsequent analysis.

Tools to scan, monitor, and eradicate viruses can identify and destroy malicious programs that may have inadvertently been transmitted onto host systems. The damage potential of viruses ranges from mere annoyance (e.g., an unexpected "Happy Holidays" jingle without further effect) to the obliteration of critical data resources. To ensure continued protection, the virus identification data on which such tools depend must be kept up to date. Most virus tool vendors provide subscription services or other distribution facilities to help customers keep up to date with the latest viral strains.

Security Analysis Tools. Because of the increasing sophistication of intruder methods and the vulnerabilities present in commonly used applications, it is essential to assess periodically network susceptibility to compromise. A variety of vulnerability identification tools are available, which have garnered both praise and criticism. System administrators find these tools useful in identifying weaknesses in their systems. Critics argue that such tools, especially those freely available to the Internet community, pose a threat if acquired and misused by intruders.

Cryptography

One of the primary reasons that intruders can be successful is that most of the information they acquire from a system is in a form that they can read and comprehend. When you consider the millions of electronic messages that traverse the Internet each day, it is easy to see how a well-placed network sniffer might capture a wealth of information that users would not like to have disclosed to unintended readers. Intruders may reveal the information to others, modify it to misrepresent an individual or organization, or use it to launch an attack. One solution to this problem is, through the use of cryptography, to prevent intruders from being able to use the information that they capture.

Encryption is the process of translating information from its original form (called *plaintext*) into an encoded, incomprehensible form (called *ciphertext*). Decryption refers to the process of taking ciphertext and translating it back into plaintext. Any type of data may be encrypted, including digitized images and sounds.

Cryptography secures information by protecting its confidentiality. Cryptography can also be used to protect information about the integrity and authenticity of data. For example, checksums are often used to verify the integrity of a block of information. A checksum, which is a number calculated from the contents of a file, can be used to determine if the contents are correct. An intruder, however, may be able to forge the checksum after modifying the block of information. Unless the checksum is protected, such modification might not be detected. Cryptographic checksums (also called message digests) help prevent undetected modification of information by encrypting the checksum in a way that makes the checksum unique.

The authenticity of data can be protected in a similar way. For example, to transmit information to a colleague by E-mail, the sender first encrypts the information to protect its confidentiality and then attaches an encrypted digital signature to the message. When the colleague receives the message, he or she checks the origin of the message by using a key to verify the sender's digital signature and decrypts the information using the corresponding decryption key. To protect against the chance of intruders modifying or forging the information in transit, digital signatures are formed by encrypting a combination of a checksum of the information and the author's unique private key. A side effect of such authentication is the concept of nonrepudiation. A person who places their cryptographic digital signature on an electronic document cannot later claim that they did not sign it, since in theory they are the only one who could have created the correct signature.

Current laws in several countries, including the United States, restrict cryptographic technology from export or import across national borders. In the era of the Internet, it is particularly important to be aware of all applicable local and foreign regulations governing the use of cryptography.

6 Information Warfare

Extensive and widespread dependence on the Internet has called new attention to the importance of information to national security. The term *information warfare* refers to the act of war against the information resources of an adversary. Like warfare on land or in the air, information warfare is one component of a range of attack strategies for dominating an adversary in order to gain or maintain an objective.

Information warfare is divided into two categories: offensive and defensive. The purpose of offensive information warfare is to attack the information resources of an adversary to gain dominance. Defensive information warfare is the protection of your information assets against attack.

Information assets can take many forms, from messages sent by courier in diplomatic bags to the computers used to analyze enemy positions based on satellite data. In computer security, information assets include digital information, the computers that process them, and the networks that transmit the digital information from place to place. Computer security is a key element for protecting the availability, integrity, and confidentiality of all these information assets.

Internet security protects information assets consisting of computers, information, and networks that are part of the Internet. Internet security is related to information warfare when the Internet contains information assets that are important to the information warfare objective. For example, if an adversary can use the Internet to access battle plans, the Internet is being used for information warfare.

Internet security is important to both offensive and defensive information warfare because the Internet is a global and dependable resource on which many countries rely. Historically, military networks and computers were unreachable by nonmilitary participants. The Internet, however, provides a cost-effective way for military and government units to communicate and participate in achieving objectives. Use of the Internet means that individuals, multinational companies, and terrorist organizations all can gain access to important information resources of governments and military forces. Thus, it is important to address Internet security concerns as a key component of defensive information warfare.

Because the Internet is global, it can be an avenue of attack for offensive information warfare by many governments. One of the battlefields for a future military offensive could very well involve the Internet. Intruder technology (as described in a separate section above) could be used by a government as a weapon against information resources, or used randomly by a terrorist organization against civilian targets.

In the study of information warfare, there are many new problems to solve that are not evident in other forms of warfare. These problems include identifying the enemy, responding without making your systems vulnerable to attack, and gathering intelligence on the Internet about preparations for a military exercise. These and other problems are likely to be the subject of discussion and investigation for some time to come.

7 The Future

Research and development efforts are underway to allow critical applications to operate in the future in a more secure environment than exists today.

Internetworking Protocols

Most of the network protocols currently in use have changed little since the early definitions of the ARPA research and education network when trust was the norm. To have a secure foundation for the critical Internet applications of the future, severe weaknesses must be addressed: lack of encryption to preserve privacy, lack of cryptographic authentication to identify the source of information, and lack of cryptographic checksums to preserve the integrity of data (and the integrity of the packet routing information itself). New internetworking protocols are under development which use cryptography to authenticate the originator of a packet and to protect the integrity and confidentiality of data.

The IETF (Internet Engineering Task Force) Proposed Standard for the Next Generation Internet Protocol (IPng) is being designed to cope with the vastly increased addressing and routing needs associated with the exponential growth of the Internet. IPng provides integral support for authenticating hosts and protecting the integrity and confidentiality of data.

The first release of IPng is officially termed IPv6 (Internet Protocol version 6). Since it is impractical to replace the existing protocol instantly and simultaneously throughout the Internet, IPv6 is designed to coexist with the current version of IP, allowing for a gradual transition over the course of years. Implementations of IPv6 for many routers and host operating systems are underway.

In the future, authentication protocols will increasingly be supported by technology that authenticates individuals (in the context of their organizational or personal roles) through the use of smart cards, fingerprint readers, voice recognition, retina scans, and so forth.

Protocol design, analysis, and implementation will be the subject of continued research. A primary goal is 100% verifiably secure protocols (that is, protocols as provably secure as the cryptographic algorithms supporting them), but researchers are nowhere near attaining this goal.

Intrusion Detection

Research is underway to improve the ability of networked systems and their managers to determine that they are, or have been, under attack. Intrusion detection is recognized as a problematic area of research that is still in its infancy. There are two major areas of research in intrusion detection: anomaly detection and pattern recognition.

Research in anomaly detection is based on determining patterns of "normal" behavior for networks, hosts, and users and then detecting behavior that is significantly different (anomalous). Patterns of normal behavior are frequently determined through data collection over a period of time sufficient to obtain a good sample of the typical behavior of authorized users and processes. The basic difficulty facing researchers is that normal behavior is highly variable based on a wide variety of innocuous factors. Many of the activities of intruders are indistinguishable from the benign actions of an authorized user.

The second major area of intrusion detection research is pattern recognition. The goal here is to detect patterns of network, host, and user activity that match known intruder attack scenarios. One problem with this approach is the variability that is possible within a single overall attack strategy. A second problem is that new attacks, with new attack patterns, cannot be detected by this approach.

Finally, to support the needs of the future Internet, intrusion detection tools and techniques that can identify coordinated distributed attacks are critically needed, as are better protocols to support traceability.

Software Engineering and System Survivability

Current software engineering methods and practice have had only limited success in managing the intellectual complexity of designing and implementing software. Moreover, in the design of software systems, security concerns are typically an afterthought (addressed through add-ons and software patches) rather than being an integral part of the overall design. This means that software systems of any significant size and complexity are likely to have exploitable security flaws. Because managing the intellectual complexity of software is difficult, up-front security design in products is rare, and detailed knowledge about systems is widespread, systems will be breached in spite of our best efforts to make them invulnerable. Therefore, the concept of information systems security must encompass the specification of systems that exhibit behaviors that contribute to survivability in spite of intrusions. Only then can systems be developed that are robust in the presence of attack and are able to survive attacks that cannot be completely repelled.

System survivability is the capacity of a system to continue performing critical functions in a timely manner even if significant portions of the system are incapacitated by attack or accident. We use the term *system* in the broadest possible sense, which includes networks and large-scale "systems of systems".

Although the concepts and practices associated with system survivability are embryonic, they include (but are not limited to) traditional areas of software engineering and computer science such as reliability, testing, dependability, fault tolerance, verification of correctness, performance, and information system security. Promising research in survivability encompasses a wide variety of research methods in software engineering. Inoculation tools may be developed that will automate the distribution of security fixes, throughout an entire network infrastructure, to provide comprehensive protection from a newly discovered security flaw. The concept of inoculation may be further generalized to encompass adaptive networks, which consist of distributed cooperative network elements that exchange information on security problems and actively change and adjust in response to security threats.

Web-Related Programming and Scripting Languages

Downloading interesting, informative, or entertaining "content" from a remote site to a user's local machine is central to the activity of Web browsing (or "net surfing"). The content getting the most attention from Web users and the greatest concern from security experts is executable content, code to be executed on the local machine on download. This executable content may provide live audio of a conference in progress, a jazz tune, three-dimensional (3-D) animation effects, or hostile code that destroys the local file system. Executable code is authored using one or more Web-related programming or scripting languages designed specifically for the production of platform-independent executable content. Languages in this category include JAVA and ActiveX. Executable content is called an "applet" in JAVA and a "control panel" in ActiveX.

Web-related programming languages pose new security challenges and concerns because code is downloaded, installed, and run on a user's machine without a review of source code (the recommended practice for secure use of publicly available software). These activities can be triggered by following any hypertext link or opening any page while browsing. A user may not even be aware that code has been downloaded and executed. Some Web-related programming languages, most notably JAVA, have built-in security features, but security experts are concerned about the adequacy of these features.

As executable content makes Web browsing even more alluring, further research in software engineering and greater user awareness will be necessary to counter security risks. Presently, the security of executable content depends upon the correctness of multiple vendors' implementations, the inherent security of platform-independent "virtual machines," and the safety of the source code that is executed. In the foreseeable future, users need to be educated about the risks so they can make informed choices about where to place their trust.

Intelligent Autonomous Agents - A New Computing Paradigm

The future Internet environment is likely to be increasingly dependent on an agent-based model of computing, with significant implications for Internet security. Agents are executable software objects with executions that are not tied to any specific host or computing resource or to any geographical or logical network location. Agents perform computation and communication defined by a user, but the execution platforms are typically outside the user's administrative control (and outside the administrative control of the user's organization). The conceptual model of agent operation is one in which an intelligent agent, at the request of a user, goes to one or more remote hosts to perform a computation or gather information and then returns to the user with the result. An agent's mode of operation may range from partially to fully autonomous, and the degree to which an agent is autonomous may vary throughout the life of that agent.

A future agent-based computing environment may include features such as these:

- Agents share information and cooperate to complete the user's task.
- Agents protect themselves with intrinsic security mechanisms but also depend on some measure of extrinsic security provided by the infrastructure and cooperating agents.

- Since most of an agent's activity takes place outside the user's domain of administrative control (and hence outside any firewall designed to protect the user), the traditional firewall has little to contribute to security.
- Replication and agent diversity provide increased survivability while under attack and under conditions of degraded or uncertain infrastructure support.
- Agents communicate to enhance the detection of threats. Specialized sensor agents are specifically designed to detect particular types of threats, and groups of diverse sensor agents provide the entire agent "collective" with a comprehensive profile of current threats.
- The agent-supported infrastructure protects itself and takes defensive action without user intervention.

Acknowledgments: CERT is a registered trademark and service mark of Carnegie Mellon University. UNIX is a registered trademark of AT&T Bell Laboratories.

8 List of Acronyms

ARPA	Advanced Research Projects Agency
CERT/CC	CERT® Coordination Center
DARPA	Defense Advanced Research Projects Agency
FIRST	Forum of Incident Response and Security Teams
FTP	File Transfer Protocol
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPng	Next Generation Internet Protocol (official name is IPv6)
IPv6	Internet Protocol version 6 (also informally called IPng)
NFS	Network File System
RFC	Request for Comments

9 Bibliography

Caelli, W., Longley, D., and Shain, M., *Information Security Handbook,* Stockton Press, New York, 1991.

Chapman, D. B., and Zwicky, E. D. *Building Internet Firewalls,* O'Reilly & Associates, Sebastopol, CA, 1995.

Cheswick, W. R., and Bellovin, S. M., *Firewalls and Internet Security: Repelling the Wily Hacker,* Addison-Wesley, New York, 1994.

Comer, D. E., *Internetworking with TCP/IP,* 3 vols., Prentice-Hall, Englewood Cliffs, NJ, 1991 and 1993.

Garfinkel, S., *PGP: Pretty Good Privacy,* O'Reilly & Associates, Sebastopol, CA, 1995.

Garfinkel, S., and Spafford, G., *Practical UNIX and Internet Security,* 2nd ed., O'Reilly & Associates., Sebastopol, CA, 1996.

Kaufman, C., Perlman, R., and Speciner, M., *Network Security: Private Communication in a Public World,* PTR Prentice-Hall, Englewood Cliffs, NJ, 1995.

McGraw, G., and Felten, E. W., *Java Security,* John Wiley& Sons, New York, 1996.

National Research Council, *Computers at Risk: Safe Computing in the Information Age,* National Academy Press, Washington, D. C., 1991.

Schneier, B., *Applied Cryptography: Protocols, Algorithms, and Source Code in C,* 2d ed., John Wiley & Sons, New York, 1996.

10 References

1. Network Wizards. Data is available on line: http://www.isc.org/ds/.
2. Levy, S., *Hackers: Heroes of the Computer Revolution*, Anchor Press/Doubleday, Garden City, NY, 1984.
3. Stoll, C., *The Cuckoo's Egg: Tracking a Spy Through the Maze of Computer Espionage*, Doubleday, New York, 1989.
4. Denning, P. J., (ed.), *Computers Under Attack: Intruders, Worms, and Viruses*, ACM Press, Addison-Wesley, New York, 1990.
5. CERT Coordination Center, CERT* advisories and other security information, CERT/CC, Pittsburgh, PA. Available online: http://www.cert.org/.
6. Internet Engineering Task Force, Site Security Policy Handbook Working Group, *Site Security Handbook*, RFC 1244, available online from ftp://info.cert.org/pub/ietf/ssphwg/
7. Note (added Sept. 1, 1998): RFC 1244 has been replaced by RFC 2196, which is available online from http://www.ietf.org.rfc/rfc2196.txt.
8. Internet Engineering Task Force, Network Working Group, *Guidelines for the Secure Operation of the Internet*, RFC 1281. Available on line: ftp://info.cert.org/pub/ietf/ssphwg/
9. Note (added Sept. 1, 1998): The current location of RFC 1281 is http://www.ietf.org/rfc/rfc1281.txt.
10. Internet Engineering Task Force, Network Working Group, *The MD5 Message-Digest Algorithm*, RFC 1321. Available on line: http://www.ietf.org/rfc/rfc1321.txt.

Prepared for presentation on the web February 1998