Chapter 1

The CERT® Guide to System and Network Security Practices

The Problem—In the Large[1]

Networks have become indispensable for conducting business in government, commercial, and academic organizations. Networked systems allow you to access needed information rapidly, improve communications while reducing their cost, collaborate with partners, provide better customer services, and conduct electronic commerce.

Many organizations have moved to distributed, client-server architectures where servers and workstations communicate through networks. At the same time, they are connecting their networks to the Internet to sustain a visible business presence with customers, partners, and suppliers. While computer networks have revolutionized the way companies do business, the risks they introduce can be devastating. Attacks on networks can lead to lost money, time, products, reputation, sensitive information, and even lives.

The 2000 Computer Security Institute/FBI Computer Crime and Security Survey (CSI 00) indicates that the number of computer crime and other information security breaches is still on the rise and that their cost is increasing. For example, 70 percent of the 585 respondents reported computer security breaches within the last 12 months—

1. This Problem description is directly quoted from (Allen 00a).

up from 62 percent in 1999. Furthermore, the financial losses for the 273 organizations that were able to quantify them totaled $265,586,240—more than double the 1999 figure of $123,779,000.

Engineering for ease of use is not being matched by engineering for ease of secure administration. Today's software products, workstations, and personal computers bring the power of the computer to increasing numbers of people who use that power to perform their work more effectively. Products are so easy to use that people with little technical knowledge or skill can install and operate them on their desktop computers. Unfortunately, it is difficult to configure and operate many of these products securely. This gap between the knowledge needed to operate a system and that needed to keep it secure is resulting in increasing numbers of vulnerable systems. (Pethia 00)

Technology evolves so rapidly that vendors concentrate on time to market, often minimizing that time by placing a low priority on security features. Until their customers demand products that are more secure, the situation is unlikely to change.

Users count on their systems being there when they need them and assume, to the extent that they think about it, that their Information Technology (IT) departments are operating all systems securely. But this may not be the case. System and network administrators typically have insufficient time, knowledge, and skill to address the wide range of demands required to keep today's complex systems and networks up and running. Additionally, evolving attack methods and emerging software vulnerabilities continually introduce new threats into an organization's installed technology and systems. Thus, even vigilant, security-conscious organizations discover that security starts to degrade almost immediately after fixes, workarounds, and new technology are installed. Inadequate security in the IT infrastructures can negatively affect the integrity, confidentiality, and availability of systems and data.

Who has this problem? The answer is, just about everyone. In fact, anyone who uses information technology infrastructures that are networked, distributed, and heterogeneous needs to care about improving the security of networked systems.

Whether you acknowledge it or not, your organization's networks and systems are vulnerable to both internal and external attack. Organizations cannot conduct business and build products without a robust IT infrastructure. And an IT infrastructure vulnerable to intruder attack cannot be robust. In addition, users have an organizational, ethical, and often legal responsibility to protect competitive and sensitive information. They must also preserve the reputation and image of their organizations and business partners. All of these can be severely compromised by successful intrusions.

As depicted in Figure 1.1, in the 1980s the intruders were system experts with a high level of expertise who personally constructed the methods for breaking into systems. Use of automated tools and exploit scripts was the exception rather than the rule. By the year 2000, due to the widespread and easy availability of intrusion tools and exploit

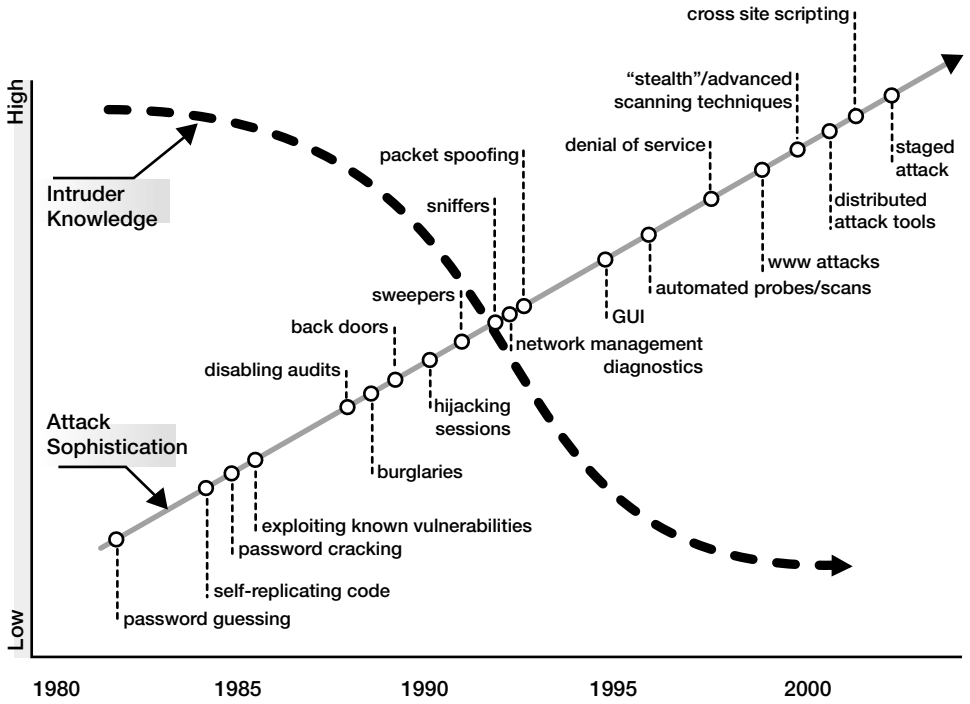

Figure 1.1 Attack sophistication versus intruder technical knowledge

scripts that can easily duplicate known methods of attack, absolutely anyone could attack a network. While experienced intruders are getting smarter, as demonstrated by increasingly sophisticated types of attacks, novice intruders require correspondingly decreasing knowledge to copy and launch known methods of attack. Meanwhile, as evidenced by distributed denial-of-service (DoS) attacks[2] and variants of the Love Letter Worm, both the severity and scope of attack methods are increasing.

In the early to mid-1980s, intruders manually entering commands on a personal computer could access tens to hundreds of systems; 20 years later they could use automated tools to access thousands to tens of thousands of systems. In the 1980s, it was also relatively simple to determine if an intruder had penetrated your systems and discover what he or she had done. By the year 2000, however, intruders could totally conceal their presence by, for example, disabling commonly used services and reinstalling their own versions, erasing their tracks in audit and log files. In the 1980s and early 1990s, DoS

2. Refer to the CERT tech tip *Denial of Service Attacks* (available at *http://www.cert.org/tech_tips/denial_of_service.html*) and CERT advisories on this subject.

attacks were infrequent and not considered serious. Today, a successful DoS attack on an Internet service provider that conducts its business electronically can put that provider out of business. Unfortunately, these types of attacks occur more frequently each year.

Because of the explosion of Internet use, the demand for competent system administrators with the necessary technical experience far exceeds the supply of individuals either graduating from formal degree programs or with knowledge and skills acquired through hands-on experience. As a result, people who are not properly qualified are being hired or promoted from within to do the job. This trend is exacerbated by the fact that some skilled, experienced system administrators change jobs frequently to increase their salaries or leave the job market because of burnout.

Today's audit and evaluation products typically focus on the underlying system and network technologies without considering the organizational concerns (e.g., policies, procedures) and human aspects (e.g., management, culture, knowledge and skills, incentives) that can dramatically affect the security posture of IT infrastructures. As a result, companies often implement incomplete or narrow solutions with the expectation that these will completely solve the problem.

The Problem—As Viewed by Administrators

Systems, networks, and sensitive information can be compromised by malicious and inadvertent actions despite an administrator's best efforts. Even when administrators know what to do, they often don't have the time to do it; operational day-to-day concerns and the need to keep systems functioning take priority over securing those systems. Administrators choose how to protect assets, but when managers are unable to identify which assets are the most critical and the nature of the threats against them (as part of a business strategy for managing information security risk), the protections an administrator offers are likely to be arbitrary at best. Unfortunately, managers often fail to understand that securing assets is an ongoing process, not a one-shot deal, and, as a result, they do not consider this factor when allocating administrator time and resources. Even if an organization decides to outsource security services, it will probably continue to be responsible for the establishment and maintenance of secure configurations and the secure operations of critical assets.

Most system and network administrators have developed their knowledge of how to protect and secure systems from experience and word of mouth, not by consulting a published set of procedures that serve as de facto standards generally accepted by the administrator community; no such standards currently exist. For this reason and those stated above, administrators are sorely in need of security practices that are easy to access, understand, and implement. The practices in this book are intended to meet these needs.

We recognize that it may not be practical to implement all steps within a given practice or even all practices. Business objectives, priorities, and an organization's ability to manage and tolerate risk dictate where IT resources are expended and determine the trade-offs among security and function, operational capability, and capacity. However, we believe that by adopting these practices, an administrator can act now to protect against today's threats, mitigate future threats, and improve the overall security of the organization's networked systems.

How to Use This Book

The most effective way to use this book is as a reference. We have attempted to provide adequate cross-referencing from one practice to other, related practices; and we have deliberately included some repetition from practice to practice to allow each to stand alone.

All practices assume the existence of the following information:

- Business objectives and goals from which security requirements derive. These may require periodically conducting an information security risk analysis and assessment to help set priorities and formulate protection strategies (see Key Definitions below).
- Organization-level and site-level security policies that can be traced to the above business objectives, goals, and security requirements. If such policies do not currently exist, the development of such policies is recognized as essential and is under way. Charles Cresson Wood (Wood 00), among others, has prepared an extensive reference guide describing all elements of a security policy along with sample policy language. Each practice in this book contains a closing section describing the security policy language that must be considered to ensure successful implementation of the practice. This language will likely need to be tailored to reflect the specific business objectives and security requirements of your organization and its computing environment. Appendix B lists all policy-related language and guidance presented in this book.

> Security policies define the rules that regulate how your organization manages and protects its information and computing resources to achieve security objectives. Security policies and procedures that are documented, well known, and visibly enforced establish expected user behavior and serve to
>
> *continued*

inform users of their obligations for protecting computing assets. Users include all those who access, administer, and manage your systems and have authorized accounts on your systems. They play a vital role in implementing your security policies.

A policy must be enforceable to achieve its objectives. In most organizations, the system administrators responsible for the technological aspects of information security do not have the authority to enforce security policies. It is therefore necessary to educate your management about security issues and the need for policies in specific topic areas such as acceptable use (refer to Section 2.15), and then to obtain a commitment to support the development, rollout, and enforcement of those policies.

Designate an individual in your organization to have responsibility for the development, maintenance, and enforcement of all security policies. The person who fills this role must have enough authority to enforce these policies. In many large organizations, the chief information officer (CIO) is the appropriate choice. While the CIO will probably delegate the tasks of writing and maintaining the policy, he or she must retain the responsibility and authority to enforce it.

As a general rule, policies are more successful if they are developed in cooperation with the people to whom they apply. Users, for example, are in the best position to evaluate how various policy statements might affect how they perform their work. Although middle- or high-level managers may be responsible for setting overall information security policies, they need to collaborate with system administrators, operations staff, security staff, and users in order to define reasonable technological and procedural protection measures for information resources.

When a new policy is first adopted in an established organization, not everyone will want to make the behavioral changes to comply with it. The responsible executive must be sure to explain the motivation for the policy. Peers, including those who participated in the development of the policy, can help accomplish this.

Train new employees about the policy as part of their initial orientation and inform all employees whenever the policy changes, retraining them if necessary. Make sure they understand the consequences of noncompliance.

To ensure user acceptance of any policies that require their compliance, require each user to sign a statement acknowledging that he or she understands the policy and agrees to follow it.

The practices in Part I provide a strong foundation through establishing secure configurations of computing assets. If these are set up correctly *and maintained,* many of the common vulnerabilities typically exploited by intruders will be eliminated. Following these practices can thus greatly reduce the impact of a significant number of known, recurring attacks. Part II assumes that the practices in Part I have been implemented and provides guidance on what to do if something suspicious, unexpected, or unusual occurs. The practices presented in Parts I and II are technology-neutral, that is, independent of any specific operating system or version. Appendix A presents examples of practice implementations that are operating-system-specific.

How This Book Is Organized

Figure 1.2 serves as one top-level depiction of how to secure and protect information assets. It includes steps to harden/secure, prepare, detect, respond, and improve.

Harden/Secure

Systems shipped by vendors are very usable but unfortunately often contain many weaknesses when viewed from a security perspective.[3] Vendors seek to sell systems that are ready to be installed and used by their customers. The systems perform as advertised, and they come with most, if not all, services enabled by default. Vendors apparently want to minimize telephone calls to their support organizations and generally adopt a "one size fits all" philosophy in relation to the systems they distribute. First, therefore, an administrator needs to redefine the system configuration to match the organization's security requirements and policy for that system.

Taking this step will yield a hardened (secure) system configuration and an operational environment that protects against known attacks for which there are designated mitigation strategies. To complete this step, follow the instructions below in the order listed:

1. Install only the minimum essential operating system configuration, that is, only those packages containing files and directories that are needed to operate the computer.

3. Refer to the CERT vulnerabilities database (at *http://www.kb.cert.org/kb/*), CERT vulnerability notes (at *http://www.cert.org/vul_notes*), and the Common Vulnerabilities and Exposures (CVE) site at *http://cve.mitre.org* for detailed vulnerability information.

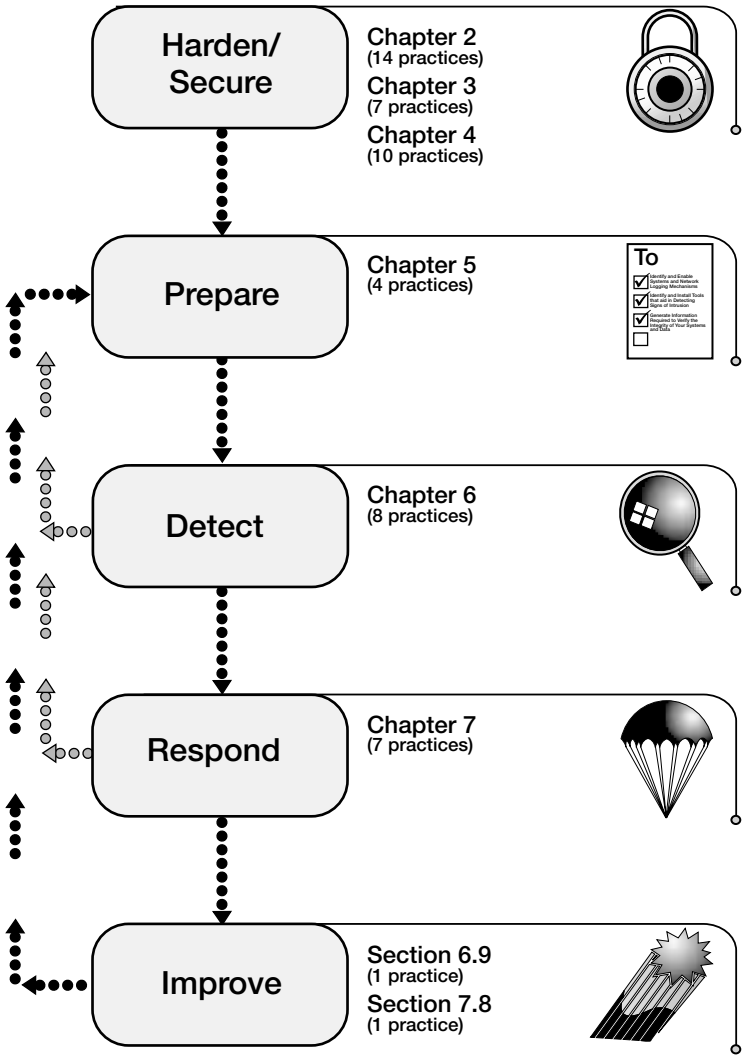

Figure 1.2 Securing information assets

2. Install patches to correct known deficiencies and vulnerabilities. Installing patches should be considered an essential part of installing the operating system but is usually conducted as a separate step.
3. Install the most secure and up-to-date versions of system applications. It is essential that all installations be performed before step 4, as any installation performed after privileges are removed can undo such removal and result in, for example, changed mode bits or added accounts.
4. Remove all privilege and access and then grant (add back in) privilege and access only as needed, following the principle "deny first, then allow."
5. Enable as much system logging as possible to have access to detailed information (needed in the case of in-depth analysis of an intrusion)

Chapter 2 contains practices for hardening and securing general-purpose servers and workstations. These include configuring, minimizing deployed services, authenticating users, controlling access, performing backups, and performing remote administration in a secure manner. Additional hardening details can be found in the CERT implementation *Installing and Securing Solaris 2.6 Servers.*[4] Chapter 3 addresses more specific details for securing public web servers, such as web server placement, security implications of external programs, and using encryption. Chapter 4 provides guidance on deploying firewall systems, including firewall architecture and design, packet filtering, alert mechanisms, and phasing new firewalls into operation. The practices in Chapters 3 and 4 build upon and assume previous configuration of a secure general-purpose server as described in Chapter 2. This relationship is shown in Figure 1.3.

Prepare

The philosophy of the preparation step hinges on the recognition that a collection of vulnerabilities exists that are yet to be identified, requiring an administrator to be in a position to recognize when these vulnerabilities are being exploited. To support such recognition, it is vitally important to characterize a system so that an administrator can understand how it works in a production setting. Through a thorough examination and recording of a known baseline state and of expected changes at the network, system (including kernel), process, user, file, directory, and hardware levels, the administrator and his or her manager learns the expected behavior of an information asset. In addition, the administrator must develop policies and procedures to identify, install, and

4. Available at *http://www.cert.org/security-improvement* under UNIX implementations.

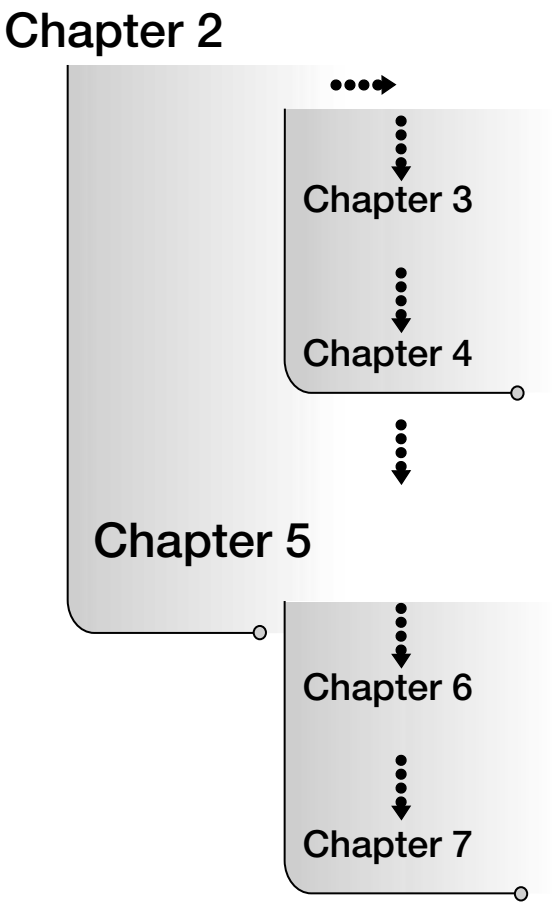

Figure 1.3 Practice dependencies

understand tools for detecting and responding to intrusions well before such policies, procedures, and tools need to be invoked.

One way to think about the distinction between the hardening and securing step and the characterization part of preparing is that hardening attempts to solve *known* problems by applying known solutions, whereas characterization helps identify new problems and formulate new solutions. In the case of characterization, the problems are identified through anomaly-based detection techniques, that is, departures from normal behavior, so that new solutions can be formulated and applied.

Chapter 5 contains practices for characterizing information assets, preparing to detect signs of intrusion, and preparing to respond to intrusions. As shown in Figure

1.3, the practices in this chapter are necessary precursors to those in Chapters 6 and 7. In addition, the practices in Part I (Chapters 2, 3, and 4) should be viewed as precursors to those described in Part II (Chapters 5, 6, and 7). This higher-level relationship is also depicted in Figure 1.3.

Detect

This step occurs during the monitoring of transactions performed by some asset (such as looking at the logs produced by a firewall system or a public web server). The administrator notices some unusual, unexpected, or suspicious behavior, learns something new about the asset's characteristics (see Section 5.3), or receives information from an external stimulus (a user report, a call from another organization, a security advisory or bulletin). These indicate either that something needs to be analyzed further or that something on the system has changed or needs to change (a new patch needs to be applied, a new tool version needs to be installed, etc.). Analysis includes investigating unexpected, suspicious behavior that may be the result of an intrusion and drawing some initial conclusions, which are further refined during the **Respond** step. Possible changes include a number of improvement actions (see **Improve** below):

- Installing a patch (rehardening)
- Updating the configuration of a logging, data collection, or alert mechanism
- Updating a characterization baseline to add unexpected but now acceptable behavior or remove no longer acceptable behavior
- Installing a new tool

Chapter 6 contains practices for detecting signs of intrusion in the following information assets:

- Detection tools
- Networks
- Systems (including processes and user behavior)
- Network and system performance
- Files and directories
- Hardware
- Access to physical resources

Chapter 6 practices assume that those described in Chapter 5 have been implemented.

Respond

For the purposes of this book, response includes recovery. In this step, an administrator further analyzes the effects of, scope of, and damage caused by an intrusion, contains these effects as far as possible, works to eliminate future intruder access, and returns information assets to a known, operational state—possibly while continuing analysis. Other parties that may be affected are notified, and evidence is collected and protected in the event of legal proceedings against the intruder.

Chapter 7 addresses response practices and assumes that the relevant practices in Chapter 5 have been implemented.

Improve

Improvement actions, described in Part II, Sections 6.9 and 7.8, typically occur following a detection or response activity. In addition to those noted under **Detect** above, improvement actions may involve the following steps:

- Holding a post-mortem review meeting to discuss lessons learned
- Updating policies and procedures
- Updating tool configurations and selecting new tools
- Collecting measures of resources required to deal with the intrusion and other security business case information

Improvement actions may cause you to revisit **Harden/Secure, Prepare,** and **Detect** practices.

Chapter Structure

Each chapter includes an overview of the problem being addressed and an introduction to possible solutions. The major sections within each chapter provide detailed descriptions of solutions that serve as security practices. Each security practice consists of an introduction, a series of practical steps presented in the order of recommended implementation, and a section covering policy considerations that complements these steps and helps ensure that they will be deployed effectively.

The recommended steps are addressed directly to a mid-level system or network administrator with several years of experience.[5] In some cases (policy considerations,

5. Refer to the SAGE (System Administrators Guild) job description for intermediate system administrators, available at *http://www.usenix.org/sage/jobs/jobs-descriptions.html#Intermediate.*

deployment plans, etc.), the person addressed is the manager responsible for system and network administration.

Each chapter closes with a checklist summarizing all of the practices and steps within each one. The checklist also serves as a table of contents that can be reviewed prior to reading each chapter.

Key Definitions

The following definitions are used throughout the book:

Assets generally include information, hardware, software, and people. Asset values are determined based on the impact to the organization if the asset is lost. Critical assets are those that are essential to meeting an organization's mission and business objectives. (Alberts 00) For the purposes of this book, assets include the information, hardware, and software that make up the information technology infrastructure of an organization.

Threat is defined as anything that may compromise an asset. This could be a person, such as an employee or a hacker, or it could be a competitor or anyone else with deliberate intent to compromise an asset. Threats also include anything that results in accidental disruption to an asset (such as a natural disaster), the means of access to do so, or any outcome or consequence that results in an unwanted effect such as disclosure, modification, destruction, loss, or interruption. (Alberts 00) Threats include vulnerabilities and risks of exposure.

Information security risk analysis and assessment methods help an organization identify important assets, threats against these assets, security requirements for these assets, and weaknesses or vulnerabilities in current practice that increase the likelihood of these assets being compromised. Refer to *Operationally Critical Threat, Assets, and Vulnerability Evaluation*SM *(OCTAVE*SM**)** *Framework, Version 1.0* (Alberts 99), *Survivable Network Analysis Method* (Mead 00), *Secure Computing* (Summers 97), *Network Intrusion Detection: An Analyst's Handbook* (Northcutt 99), and "Web of Worries" (Kessler 00) for more information on this subject. You can find additional guidance in a publication titled *Information Security Risk Assessment, Practices of Leading Organizations* (GAO/AIMD-00-33), published by the U.S. General Accounting Office (Washington, D.C., November 1999).

Attack connotes an action conducted by an adversary, the attacker, on a potential victim. From the perspective of the administrator responsible for maintaining a system, an attack is a set of one or more events that has one or more security consequences. From the perspective of a neutral observer, the attack can either be successful—an intrusion—or unsuccessful—an attempted or failed intrusion. From the perspective of an intruder, an attack is a mechanism to fulfill an objective. Intrusion implies forced

entry, while attack implies only the application of force. Information-gathering probes and scans conducted by an intruder are considered attacks for the purposes of this book. (Allen 99)

An **incident** is a collection of data representing one or more related attacks. Attacks may be related by attacker, type of attack, objectives, sites, or timing. (Allen 99)

An **intrusion** refers to an actual illegal or undesired entry into an information system. Intrusion includes the act of violating the security policy or legal protections that pertain to an information system. (Allen 99) Additionally, an intrusion represents a deliberate event as a result of an intruder gaining access that compromises the confidentiality, integrity, or availability of computers, networks, or the data residing on them. *Breach* is used as a synonym. While a DoS attack does not constitute actual "entry" or "access," it does compromise the availability of the denied asset and is thus considered an intrusion for the purposes of this book.

Sources for This Book

The CERT series of security improvement modules listed below, all of which are available on the CERT web site, served as the primary source documents for this book.

Securing Network Servers (Allen 00b)

Securing Desktop Workstations (Simmel 99)

Securing Public Web Servers (Kossakowski 00)

Deploying Firewalls (Fithen 99)

Detecting Signs of Intrusion (Allen 00c)

Responding to Intrusions (Kossakowski 99)

The scope of and topics addressed by each module, and the set of modules as a whole, were explicitly chosen to address 75–80 percent of the practices designed to solve the problems that are reported to CERT. The practices describe the steps necessary to protect systems and networks from malicious and inadvertent compromise. The practice level (technology-neutral) was intentionally chosen to be as specific as possible while remaining broadly applicable and ensuring that the practices retain their utility and shelf life longer than the most up-to-date operating system version. To keep the size of each module manageable and easy to digest in a short period of time, each module addresses an important but *relatively narrowly defined* problem in network and system security.

Complete reference information is available in the Bibliography.

Other Sources of Information

There are many excellent sources of information about emerging intruder trends, attack scenarios, security vulnerabilities, vulnerability detection, and ways to mitigate their effects. The most common sources, which are referred to frequently throughout this book and are recommended for administrators wishing to stay current, are listed below.

- Vendor web sites
- CERT current activity, advisories, summaries, incident notes, vulnerability notes, and tech tips available at the CERT web site (see sidebar).
- Web sites of computer and network security organizations (see sidebar).
- Mailing lists, some of which are sponsored by vendors
- USENET news groups

Links to many of these sites can be found on the web sites for CERT and this book, as noted in the Preface.

Advisories address Internet security problems. They offer an explanation of the problem; information that helps a reader determine if his or her site has the problem; fixes or workarounds; and vendor information. All advisories published since 1988 are available from the CERT web site advisory archives. Advisories are available at *http://www.cert.org/advisories.*

Summaries are published each quarter. They contain information on the most frequent, high-impact types of security incidents and vulnerabilities that were reported to the CERT during the previous three months; they also provide pointers to more information. Summaries are available at *http://www.cert.org/summaries.*

Incident notes provide information about current intruder activity. Vulnerability notes provide high-quality, validated information about vulnerabilities. Because CERT's understanding of the scope of a vulnerability may change, information that originally appears in these notes may later become part of an advisory. Both contain information that might help to protect systems from intrusion, and both may be updated from time to time. Incident notes are available at *http://www.cert.org/incident_notes.* Vulnerability notes are available at *http://www.cert.org/vul_notes.*

continued

Tech tips contain information on a number of Internet security issues and guidance on specific topics to secure and protect UNIX and Windows NT systems. Tech tips are available at *http://www.cert.org/tech_tips.*

Implementations provide technology-specific guidance for carrying out steps in a practice. They are available for UNIX and Windows NT systems in specific topic areas. Implementations are available at *http://www.cert.org/security-improvement.*

Mailing lists and web sites appear, disappear, and change frequently. Be sure that the sources you consult are up-to-date and reliable. Links to many of these sources can be found on the book web site.

General security information. The following sources provide both broad and detailed information on a wide range of information, system, and network security topics:

AUSCERT (Australian Computer Emergency Response Team) at *http://www.auscert.org.au.*

Bugtraq and Security Focus at *http://www.securityfocus.com.* BugTraq is a full-disclosure, moderated mailing list providing detailed discussion and announcement of computer security vulnerabilities: what they are, how to exploit them, and how to fix them.

CERIAS (Center for Education, Research, and Information Assurance Security) at *http://www.cerias.purdue.edu* (formerly known as Computer Operations, Audit, and Security Team [COAST]).

CERT/CC (CERT Coordination Center) at *http://www.cert.org.*

CVE (Common Vulnerabilities and Exposures) at *http://cve.mitre.org.* CVE is a list or dictionary that provides common names for publicly known information security vulnerabilities and exposures.

CIAC (Computer Incident Advisory Capability) at *http://ciac.llnl.gov.*

CSI (Computer Security Institute) at *http://gocsi.com.*

DFNCERT (German Computer Emergency Response Team) at *http://www.cert.dfn.de/eng/.*

FIRST (Forum of Incident Response and Security Teams) at *http://www.first.org.* Contact information for FIRST teams can be obtained from *http://www.first.org/team-info.*

ICSA (Trusecure) at *http://www.trusecure.com.*

IETF (Internet Engineering Task Force) at *http://www.ietf.org.*

SANS Institute at *http://www.sans.org.*

Security Portal at *http://www.securityportal.com.*

USENIX Advanced Computing Systems Association at *http://www.usenix.org.*

Security fixes and patches. Monitor security fixes and patches that are produced by the vendors of your systems and obtain and install all that apply. A general index of vendor sites can be found at *http://www.cert.org/security-improvement/implementations/data/vendor_list.html.*

Advisories. Subscribe to advisories that are issued by various security incident response teams and update your systems against those threats that apply to your site's technology. Sites that publish such advisories include AUSCERT, CERT, and CIAC.

Mailing lists and USENET newsgroups. Read relevant mailing lists and subscribe to USENET newsgroups (*http://www.cert.org/othersources/usenet.html*) to keep up to date with the latest information being shared by fellow administrators.

Subscribers to mailing lists usually receive announcements about security problems and software updates soon after they are available. Web sites vary considerably in the timeliness of their announcements, so you need to determine how often to look for information there. Some news-oriented web sites are updated one or more times a day, so we recommend that you monitor these daily.

Security tools. It is important to review regularly sites that contain a wide range of useful and publicly available security tools. These include the following:

CERIAS at *ftp://ftp.cerias.purdue.edu/pub/tools/unix/*

CIAC at *http://ciac.llnl.gov/ciac/ToolsUnixSysMon.html*

Insecure.org at *http://www.insecure.org/tools.html*

TAMU (Computer and Information Services Network Group at Texas A&M University) at *http://www.net.tamu.edu/network/public.html*

Wietse Venema's site at *ftp://ftp.porcupine.org/pub/security/*

Summary

This chapter has set the stage for understanding and making effective use of the practices to follow by establishing the general context, describing the book's overall organization, and providing key definitions and information sources referenced throughout. The three chapters that constitute Part I address practices for establishing secure configurations of general-purpose servers, user workstations, public web servers, and firewall systems—necessary first steps before proceeding to the practices contained in Part II. The three chapters in Part II describe what preparation, detection, and response actions should be taken when an unexpected event or behavior occur.