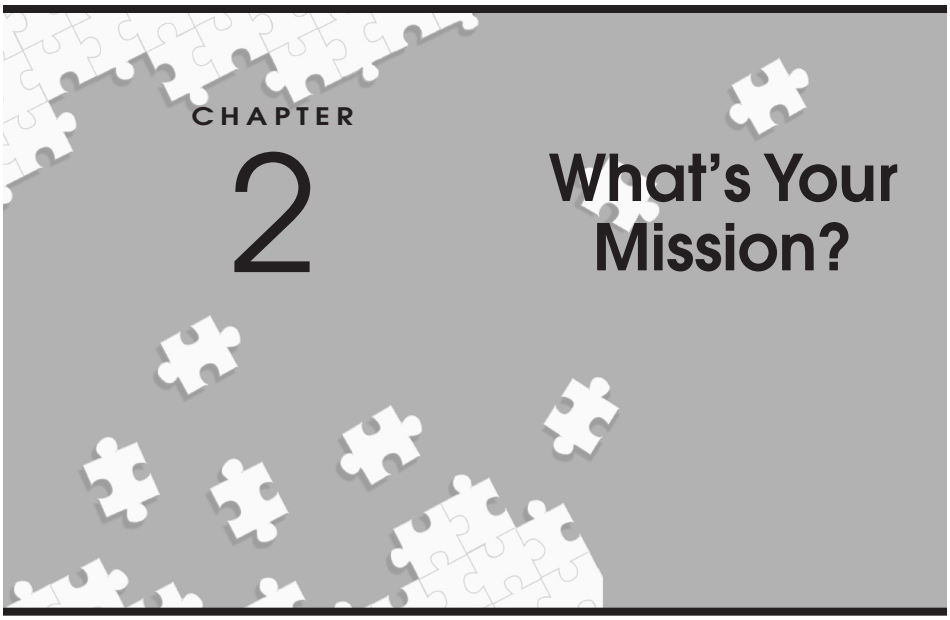

CHAPTER 2

# What's Your Mission?

Many people who enjoy putting jigsaw puzzles together take the same approach. That is, they find the pieces that outline the picture and put the frame together first, then proceed to identify and place the inner pieces. By having the border in place, it often becomes easier to get a sense of the overall picture, with the border serving as the foundation for completing the task at hand.

This chapter focuses on identifying those border pieces of the CIRT puzzle by providing numerous considerations that should be addressed during the formation of a team. By addressing these items first, a solid foundation can be laid for the other decisions that must be made toward the completion of the CIRT. The chapter first defines the focus and mission of the team, then moves into various operational aspects that may be considered.

## Focus and Scope

Management courses frequently highlight the importance of mission statements. It is often emphasized that all members of a team or organization should understand the group's mission so they can take

1. Who is included in the team's constituency? In other words, who owns the computers the team will be responsible for monitoring and responding to incidents on if they are attacked?
2. How dispersed is the constituency? If it is spread out, are there regional information technology resources that may be called upon during a crisis?
3. What is the ideal manner for the team to respond to an incident: on-site, remotely, or through some combination of these methods?
4. Will the team work with law enforcement either directly or indirectly?
5. Will the team's strategy be strictly reactive, or both proactive and reactive?
6. What services or functions will the team provide?
7. What type of activity is considered to be an incident (i.e., to what sort of activity will the team respond)?
8. How will incident reports be stored and tracked?
9. How will the incidents be counted?
10. What statistics need to be provided to reflect the team's activities? Who will need these statistics and how often? How granular should the statistics be?

**Figure 2–1** Questions to Help Identify the Team's Mission

ownership of the tasks at hand. The formation of a CIRT team is no different with respect to having a clearly identified mission. Establishing a mission up front by identifying the scope and focus of the team will ease the decision-making process for many later issues. The mission will also have a direct impact on the number and types of resources that need to be allocated to the team. To aid with this task, start by answering a few basic questions. (These questions are listed in Figure 2–1.)

## Know Who You're Protecting: Defining Your Constituency

First, who is included in the team's constituency? In other words, who are the people who own, use, operate, and are responsible for maintaining the computers that the team will be monitoring and responding to incidents on if they are attacked? In some cases, this question is

very easy to answer. For example, a university team would usually consider the students, staff, and any other authorized users to be its constituency. In other cases, the answer may not be so clear. For example, a hardware or software vendor may have a team established to address internal issues only, to address issues pertaining to its products only for its customers, or possibly to address incidents on its products and other vendors' products through a consulting arrangement. If the vendor has a partnership with another company, will the team have any responsibilities for addressing incidents as part of the partnership? Deciding who the team will support up front helps set the stage for the next set of questions that need to be addressed.

The distribution of the team's constituency or customer base will also have a major effect on many issues that must be addressed when forming a team. Will all of the team's clients be located in one building, one city, one state, one geographic region, one country, or worldwide? A large distribution will provide several obstacles or challenges when on-site support is needed and when support must be provided 24 hours per day.

## Defining Response

Once the constituency is identified, the manner in which the team will respond to incidents should be addressed. If the constituency is located in one central location, this issue should not be a problem. If the computers are dispersed over a wide geographic area or worldwide, however, on-site response can be a challenge. Depending on the technical expertise of the organization's system administrators, a team may be able to remotely provide assistance by walking a local system administrator through response procedures over the phone, rather than having someone from the team fly to the location and provide on-site support. This approach should be taken whenever possible, as the reaction time will generally be quicker (thus limiting potential damage) and the cost of responding to the incident will be lower.

If on-site support is required, then a minimum of two people should be sent to the location with specific checklists and tools at

their disposal. A two-person team is recommended for the following reasons:

- If problems arise during the response process, the second person can provide backup support to the primary person in completing a division of the tasks to be performed.
- If personnel are to be interviewed, having a second person can expedite the overall process: One person can conduct the interviews while the other performs the technical response.
- If the level of expertise varies between the individuals, the knowledge and experience of one may be supplemented by the knowledge and experience of the other. This is especially true for personnel who are relatively new to the incident response role.
- A two-person team can provide verification that procedures were correctly followed in case any questions arise about the tasks completed.
- One person can help with documenting the steps taken, while the other is performing the tasks at hand.

Keep in mind that the above checklist should be considered a guide to help with the incident response, but cannot cover every possible scenario that the incident handlers may encounter.

For teams that cover a large geographic region (including those with worldwide responsibility), the decision may be made to create several teams and have them geographically dispersed. This approach can help with covering various time zones and limit the potential response time for on-site support. On the downside, the possibility of detecting a wide-scale attack may be inhibited because varying resources analyze the incident data. If multiple teams are used, it is a good idea to have one central headquarters receive and evaluate the incident data from the subteams to provide wide-scale correlation. (As our tools progress in this realm, this correlation should become easier.)

## Working with Law Enforcement

Another issue that should be addressed up front is how the team will interact with the law enforcement community. Will the team work with law enforcement either directly or indirectly? In addressing this question, consider the following: The computer is like any other invention that has had a major impact on our daily lives in that it has both benefits and negative consequences. A good analogy is the automobile. When the car first came along, it enabled people to do things more quickly and effectively than they could prior to its invention. People were able to travel farther, travel together more efficiently, and travel in much more comfort, thereby doing more things in a day by covering more places. Cars also had negative effects—namely, automobile accidents and automobile theft. It took people some time to learn how to deal with these negative effects, design and implement safety and security features, and establish laws to address the problems and challenges of the new invention.

Similarly, the computer is an invention that has enabled people to do many more tasks in an increasingly efficient manner and has literally changed the way we live. In fact, even an auto mechanic cannot work on a car now without a computer to run diagnostic tests. Although the positive results have rapidly expanded the use of computers, some negative uses still need to be addressed—namely, computer crime. It will take us some time to catch up to the invention by developing and widely implementing security safeguards to help protect systems better and to establish case law. With this fact in mind, enabling an incident response team to work with law enforcement can have a very positive effect on the overall safeguards that are ultimately established. In fact, the best approach when investigating any incident that potentially involves a crime uses three experts: an attorney who is familiar with high-tech crime laws, the law enforcement agent, and the technical expert. Each has valuable knowledge and insight that can be vital when taking a case to trial.

One major advantage to working with a law enforcement agency is the benefit of extended networking. Specifically, many larger law enforcement agencies have developed relationships with other law

enforcement groups that may provide an added advantage in tracking an external perpetrator who has broken into a system or successfully launched a denial-of-service attack. A team trying to track an attacker on its own without any law enforcement involvement will typically find this task much more difficult, especially when the incident crosses international boundaries. It is far better (and many times easier) to provide the information to law enforcement officials and let them work with their contacts and resources to help track an intruder.

Working with law enforcement may also have some disadvantages. Often, when an organization is asked why it did not bring in law enforcement, the organization states that it didn't want the company name in the newspaper. It didn't want the publicity of a "hacking case." Although in most cases the media isn't interested in such things, this bad publicity is a real concern. There are ways to keep the company out of the news, such as using an attorney to keep the case as private as the laws will allow. Also, the organization may want to turn the case over to a prosecutor to pursue in conjunction with local, state, or federal law enforcement agents to pursue as a crime against the state or federal government, instead of naming the company in the proceedings.

An additional disadvantage of working with law enforcement may be the threat of losing control of the case. Inviting law enforcement to investigate a case may require that the case be fully investigated, even if your organization decides to stop its pursuit of the attacker. Although situations where a case cannot be stopped are very rare (all of the authors' dealings with law enforcement have been very cooperative), this outcome may be a possibility.

Even if your organization chooses to not include law enforcement as a regular part of its investigative team, it is a good idea to contact city, county, state, and federal (FBI especially) law enforcement agencies to introduce yourself and to get an idea of their services, contact information, capabilities, evidence requirements, and reporting procedures. It's always good to be prepared, even if you don't plan on using them.

Of course, legal considerations must also be addressed as they relate to privacy laws, company policies, and other issues that determine what information is shared with law enforcement. If the organi-

zation is considering the inclusion of law enforcement officials directly on the team, the first step would be to discuss the possibilities, concerns, and limitations with the appropriate legal organization as well as the management team.

If computer crime laws are to evolve so that they will better protect our information, they must be tried and tested in the criminal justice system. Organizations reporting to and working with law enforcement will facilitate this evolution. Even if an organization decides not to work directly with law enforcement, inevitably the team will encounter an incident where a law has been broken. Having contacts established up front with local, regional, state, and federal law enforcement agencies will help expedite the reporting process when this need arises. Groups such as InfraGard can help to establish these contacts.

## InfraGard

Several groups have been formed in recent years to provide an avenue for networking and resource sharing between law enforcement officials and the technical community. One of these groups is InfraGard. "The National InfraGard Program began as a pilot project in 1996, when the Cleveland FBI Field Office asked local computer professionals to assist the FBI in determining how to better protect critical information systems in the public and private sectors. From this new partnership, the first InfraGard Chapter was formed to address both cyber and physical threats."[1] InfraGard's government component is staffed by the FBI and the Department of Homeland Security's National Information Protection Center (NIPC) and includes numerous chapters throughout the United States. In fact, all 56 FBI field offices have now opened local chapters with hundreds of members across the nation. InfraGard is seen as a cooperative effort between law enforcement and the private sector, with the participants being dedicated to increasing the security of critical infrastructures within the United States. From the beginning, the FBI has stated that it is not an FBI-run program, but rather a community program in which

1. http://www.infragard.net/history_main_pg.html.

the FBI and many other government agencies participate. The InfraGard program is often likened to a "neighborhood watch" program, in which businesses and agencies with similar interests share information and experiences to help reduce the risks of a networked community.

InfraGard was formed as a national organization, with individually governed and managed chapters. Working together within and between chapters, the group members strive to better protect critical information assets by enabling the flow of information between the technical community, corporate policy makers, the owners of the critical infrastructure, law enforcement, and lawmakers. Becoming involved with a local chapter provides an excellent avenue for meeting law enforcement officials, legal experts, and other technical resources in the area that may be contacted when the need arises. More information may be obtained by contacting the closest FBI office or from the national InfraGard Web page (http://www.infragard.net).

## Operational Strategy

At this point, we have addressed some basic factors in identifying the focus and scope of the team. The next major section of the puzzle to consider is the team's operational strategy. Specifically, will the team be strictly reactive or both proactive and reactive in nature? If it is reactive only, the team would strictly respond to computer incidents as they are detected or reported from the constituency. Tools such as intrusion detection systems (IDSs) may be used to monitor for and detect unauthorized activity as it happens. (Some IDSs can also be configured to help stop an attack in progress.)

It is very difficult, and sometimes ineffective, for a team to remain completely reactive. For example, if an incident affects a number of systems, it may be a prudent measure to proactively examine other systems that do not appear to have been affected. This activity ensures that no damage was inflicted and confirms that proper safeguards are in place to prevent the systems from being affected in the future. Although prevention programs are difficult to justify because it is nearly impossible to quantify the number of incidents prevented,

it stands to reason that an incident that is prevented consumes fewer resources than one that must be investigated.[2]

If the team is to be both proactive and reactive in nature, then services such as risk assessments, vulnerability analysis, and training may be included in the capabilities offered by the team. The number of tasks assigned to the team will have a direct impact on the number of personnel and tools required for the team to be completely operational. Keep in mind that some of these services may be outsourced or scaled appropriately for the organization, depending on the projected return on the investment for the services that are offered. More detail on the types of services that the team may consider offering is provided later in this chapter.

## Defining an Incident

Regardless of the strategy taken, the type of activity that is considered to be an incident should be clearly decided up front. It is strongly recommended that a clear, concise definition be developed for the "incidents" a team will address. Generic or vague definitions such as "unauthorized activity" leave too much room for interpretation and may negatively affect operations. For example, the number of personnel assigned to the team may prove insufficient for the volume of "unauthorized activity" reported and problems may be encountered in trying to enter and track the incident data in a database or trouble ticket system.

This question leads into a discussion on the very important topic of terminology. As this topic can be quite expansive, a separate chapter has been dedicated to this one piece of the puzzle. Refer to Chapter 3 for this discussion.

## Tracking an Incident

Although this is still very early in the book, you should consider your incident tracking strategy to be a part of the foundation of the

2. Committee on Institutional Cooperation. *Final Report: Incident Cost Analysis and Modeling Project,* University of Michigan, 1998.

incident response team. The ability to track incidents is important for several reasons:

- *Management Reporting.* The incident tracking system may be used as a tool to justify and report on the performance of the investment that has been made in the team and security tools. Reporting the type and severity of incidents encountered also helps upper management to gain a clear picture of the threats being realized in their environment. This information can then be used to focus efforts on specific areas, such as user awareness, antivirus software, and increased levels of defense, so as to improve the overall security posture of the organization.
- *Incident Triage.* A mechanism for grouping incidents can lead to streamlining incident investigation activities, setting priorities for incidents, and correlating incident data.
- *Customer Service.* The incident response team's reputation among its constituency is critical to its success. A team that permits incidents to "fall through the cracks" will lose favor with its constituency, and eventually management support. Management will naturally expect the team to handle incident information and evidence in a responsible, prudent manner. In addition to incident accountability, a team must have a well-planned method of keeping the affected constituency abreast of the status of the incident.

A dedicated database for tracking incidents, correlating activity data, and reporting statistics to upper management can be the most valuable tool that an incident response organization has at its disposal. Tracking processes and mechanisms must be established to follow the incident from the time it is reported to the time it is considered closed. Otherwise, if the team is very busy (as most are in today's environment), the chance for an incident to "fall through the cracks" is too great. Two or more tools may actually be required for this function, such as a triage tool that feeds data into a central database. The triage tool may be used for tracking open tickets, while the database is used for reviewing specific details about the activity.

Once the team is operational, the sources for reporting a computer incident to the team will typically have multiple origins. A team may accept incident reports via the telephone, e-mail, IDS, paging systems,

law enforcement, or other sources. Not all information that is reported to a team as an "incident" will, in fact, turn out to be a real incident, just as a call to the emergency 911 telephone system may not be an emergency. (*Note:* 911 is the code used within the United States for emergency calls to police, fire, and medical officials. Other codes are used in different countries, such as 0, 411, and 119.)

As activity is reported from various sites or systems, mechanisms for correlating various aspects of it need to be available. For example, how often does a specific IP address show up? Have any other attacks on a specific target been reported? How many scans of a particular port number have been detected?

Finally, statistics can be a wonderful tool for any team. A database can help to quickly identify useful data, such as the number of reports received, number of open incidents, number of false reports processed, number of times that a particular Internet service provider's (ISP) addresses have attacked systems in the team's constituency, and much more. These statistics can, in turn, help to identify specific threats to the systems that the team is trying to protect as well as justify the need for the team or additional resources moving forward. (Granted, statistics can be manipulated in many ways, but it is very difficult to dispute facts that are very well defined with a good deal of granularity.)

## Counting Incidents

A similar issue that must be addressed is how incidents will be counted. It is not uncommon for an attacker to gain unauthorized access to multiple sites within the same time frame or deny services to more than one system at once. A good example was the distributed denial-of-service attacks launched in February 2000. Over the course of three days, computer systems at Yahoo!, buy.com, eBay, CNN.com, amazon.com, ZDNet.com, E*Trade, Datek online, and Excite were all affected by the attack.[3] Depending on the accounting procedures implemented by an incident response team, these attacks could have

3. Levy, Steven, and Brad Stone. "Hunting the Hackers." *Newsweek,* February 21, 2000, pp. 39–44.

been combined into one incident or tracked separately as nine incidents with similar patterns.

Combining the activity that appears to be attributable to the same person or group of persons can help with the correlation of the activity. Conversely, by grouping the activity together, important statistics may be lost. Most teams will combine the activity that appears to come from the same source. In that case, it is recommended that the database or means of tracking statistics have enough granularity built in to depict the number of targets affected by the attack.

## Services Offered

Earlier in the chapter, we touched on the decision of whether the team should be proactive or reactive. If the team will have a proactive component, a new challenge arises: What additional services will the team provide? Some of the services that teams may choose to offer are user enrollment, vulnerability assessment, penetration testing, risk assessment, and architectural review. This list certainly isn't exhaustive and will need to be adjusted to best fit the organization's requirements. The following services are presented for consideration:

### *User Enrollment*

The process of creating, modifying, and removing user accounts and privileges on the computer systems. It also includes the definition of the authorizations, group memberships, and access profiles for users.

### *Vulnerability Assessment*

The process of searching for possible susceptibility for a system to be accessed in an unauthorized way or to have authorized access denied. Many commercial and free vulnerability assessment tools can help streamline this process, although these tools do require a certain amount of experience to use them effectively. There are many opinions regarding the frequency with which these assessments should be conducted, but nearly all security professionals agree that they're not done often enough.

### *Penetration Testing*

The process of attempting to gain unauthorized access to a computer system or facility. This focused attempt to break into a system or facility is usually conducted from the perspective of a "hostile" entity and attempts to measure how much effort must be expended to gain access. The network operations group or other entity that monitors the computer resources will typically not know ahead of time that the testing will be conducted. Therefore, the capability to detect and respond to an attack can be measured while searching for potential vulnerabilities.

### *Risk Assessment*

The process of rating and evaluating vulnerabilities, threats, value, and safeguards. It takes the results of a vulnerability assessment and adds in an analysis of threats, the value of the information, and the safeguards used to protect the information. Its purpose is to help make informed decisions based on the best balance between the risk that is posed to the organization's information and the benefit of protecting it. Although several different methodologies are used for this process, it can be a very valuable tool in making decisions about how much effort to protect a system is enough.

### *Architectural Review*

The process of evaluating the hardware, software, network, policy, and management of a system or group of systems to ensure they do what is intended, and do not do what is not intended. This service mirrors the thought that security should be part of everyone's effort—throughout the complete life cycle of information, from concept to disposal.

### *User Awareness Training*

The process of teaching and reinforcing knowledge of policies, procedures, and strategies while maintaining a computing environment. In terms of effectiveness, user awareness training can be one of the most valuable services that a team offers. Conversely, a lack of user

awareness can represent a significant threat to any network. In this forum, users can learn effective password management, the organization's information-handling policies, procedures for sharing information, virus risks, tactics to defeat social engineering, and more. As Richard Baker noted in his book *Network Security: How to Plan for It and Achieve It:*

> One of the biggest obstacles to effective computer security today is an epidemic of misplaced emphasis. Corporations spend a lot of time and money buying and installing elaborate computer security systems to protect themselves from well-publicized outsiders like youthful invaders and virus carriers. They do next to nothing to train employees to make regular backups or to avoid that stereotype of computer insecurity: the password on a sticky note attached to the monitor.[4]

### *Advisory Notification*

The process where security notifications are distributed to the constituency. Many teams keep an inventory of the computing platforms, software, network infrastructure, and services that their organization employs. As a manufacturer, software vendor, or other source of information publishes a notification that there may be a vulnerability in the product, a virus, or a security-related upgrade, the team would verify that the advisory is authentic and then forward it to the affected members of the constituency. Other teams (such as CERT CC and vendor teams) will be responsible for writing the advisory that is initially distributed.

### *Research and Development*

The process of creating, evaluating, testing, and integrating new products, policies, procedures, and strategies. A lab environment in which to test and discover methods of breaking into systems, techniques for securing systems, and ways to improve and implement the enforce-

---

4. Baker, Richard H. *Network Security: How to Plan for It and Achieve It.* New York: McGraw-Hill, 1995, p. 9.

ment of company policies—and as a proving ground that enforcement can be done—can be a very valuable service to the organization.

We'll go into much more detail and share some ideas on how to implement these services later in the book.

## The Importance of Credibility

The importance of the team's credibility cannot be emphasized strongly enough. One bad report, advisory, or action can be detrimental to the entire team's credibility and the damage can take years to repair. Team credibility is a result of a combination of effectiveness, integrity, professionalism, timeliness, ethics, consistency, and the ability to deal with incidents discreetly, among several other factors. If the team's credibility is damaged, members of the team's constituency will lose faith in the team and stop relying on those personnel for support in responding to incidents. For this reason, advice should not be provided unless vulnerabilities and patches have been fully tested and verified. It can almost go without saying—the incident response team must practice what it preaches. In other words, the incident response team's own tools must be patched, maintained, and well managed so as to prevent incidents.

The CERT Coordination Center provides an excellent example of protecting team credibility. The center will not issue an advisory without fully testing the vulnerabilities and patches first in its lab to verify the steps that are recommended. Given the number of computer vulnerabilities that continue to appear, a team might potentially do no other activity than issue vulnerability alerts or advisories on a daily basis. In an effort to underscore the seriousness of a vulnerability discussed in a CERT CC advisory, the coordination center will issue alerts only on vulnerabilities that are considered very serious and may potentially affect multiple computer systems. When CERT CC does issue an advisory, it will include an MD5 checksum for patches that are recommended for downloading. This step is also part of integrity protection, providing an added check for administrators who are downloading a patch to ensure that it is the same patch that is intended to be downloaded.

## Summary

At this point, we have identified several key pieces that should provide a border to our incident response puzzle. Although these pieces may not be clearly defined as yet, the intent was to begin the consideration process that should be given to each issue. Every question listed in Figure 2–1 and discussed in this chapter will directly affect the type and number of resources needed for the team to succeed. Many of the following chapters will return to this basic set of questions as well. It is strongly recommended that the entire book be read before attempting to answer these questions. Numerous points yet to be made will influence each of these considerations.

The chapter concluded by touching on the importance of the team's credibility. If the issues described in this chapter form the border to the puzzle, the team's credibility could be thought of as the foundation on which the puzzle is built. After all, if the integrity of the team suffers in some way, the continuity of the picture we are trying to build will suffer as well. One large corner piece of this puzzle remains to be considered—terminology. Keep reading to complete the border.

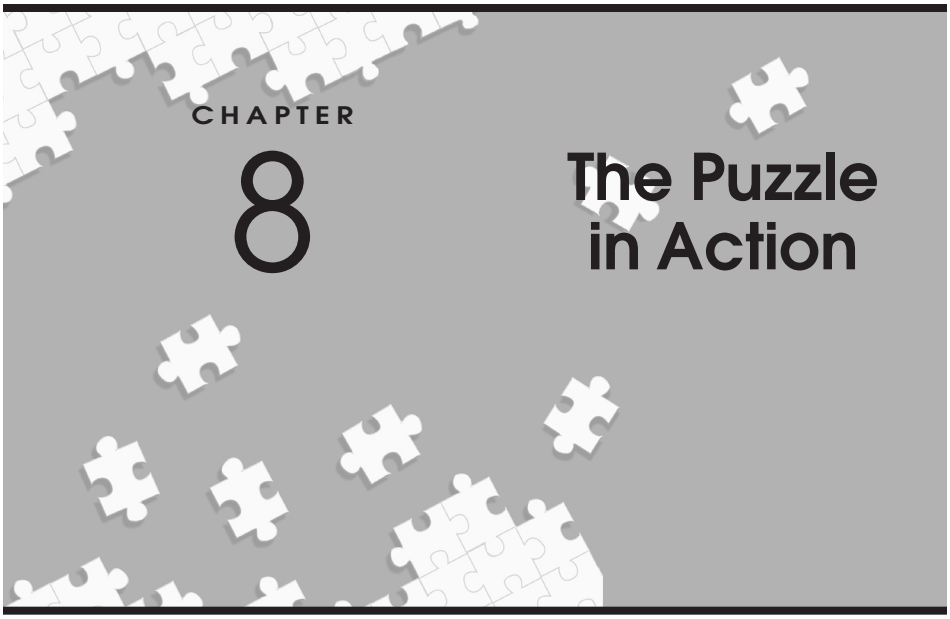


In the simplest form, everything with computers can be broken down into 1's and 0's. Similarly, computer security initiatives should always be able to be broken down into their simplest form, policies. Policies identify what is authorized and what is not, assign organizational responsibilities, communicate acceptable levels of risks, and much more. The policies may be expanded in the form of procedures, which provide the step-by-step guidelines for putting the policies into action. From there, it's a matter of implementing and configuring systems appropriately, purchasing and adding security tools to monitor and safeguard the systems, and training and authorizing end users to use the resources appropriately.

When the policies and procedures are violated, then a computer incident (e.g., unauthorized access, denial of service) may have occurred. To detect and respond to these violations of the organization's security policies, incident response policies and procedures should be in place. These policies may be in the form of stand-alone documentation, or they may be incorporated into other documentation such as company security policies or disaster recovery plans.

**NOTE** Unfortunately, not all organizations have existing computer security policies. Many people view the writing of a security policy as a huge undertaking that is nearly impossible to accomplish. Depending on the level of support from upper management, the task may be more daunting to complete in some organizations as compared to others. In the ideal situation, the organization has a security policy and is serious about covering all facets of the security equation. If the organization does not have existing policies, however, this omission should not stop the development of a CIRT. Ideally the organization will develop security policies in the near future or simultaneously as the CIRT is developed, but policies should not be viewed as a mandatory requirement for the formation of a CIRT.

This chapter focuses on the operational aspects of computer incident response. Considerations that should be given to specific incident-handling procedures will be described in detail, as will the life cycle of an incident. The information provided in this chapter can, in turn, be used to write computer incident policies and procedures. Together, these policies and procedures complete the incident response puzzle by filling in the center piece. Because computer security begins with policies, what better place to envision this piece of the puzzle than in the center where it belongs.

## The Life Cycle of an Incident

The best way to determine the policies required for incident response is to examine the typical life cycle of computer incident response. Figure 8–1 provides a flowchart outlining the major phases of the incident response life cycle; each phase is described in detail in the sections that follow. Not all incidents are identical, of course: Many have unique attributes. Therefore, the steps outlined in this section will address the typical case. The incident handler, however, must always be prepared for the unexpected.

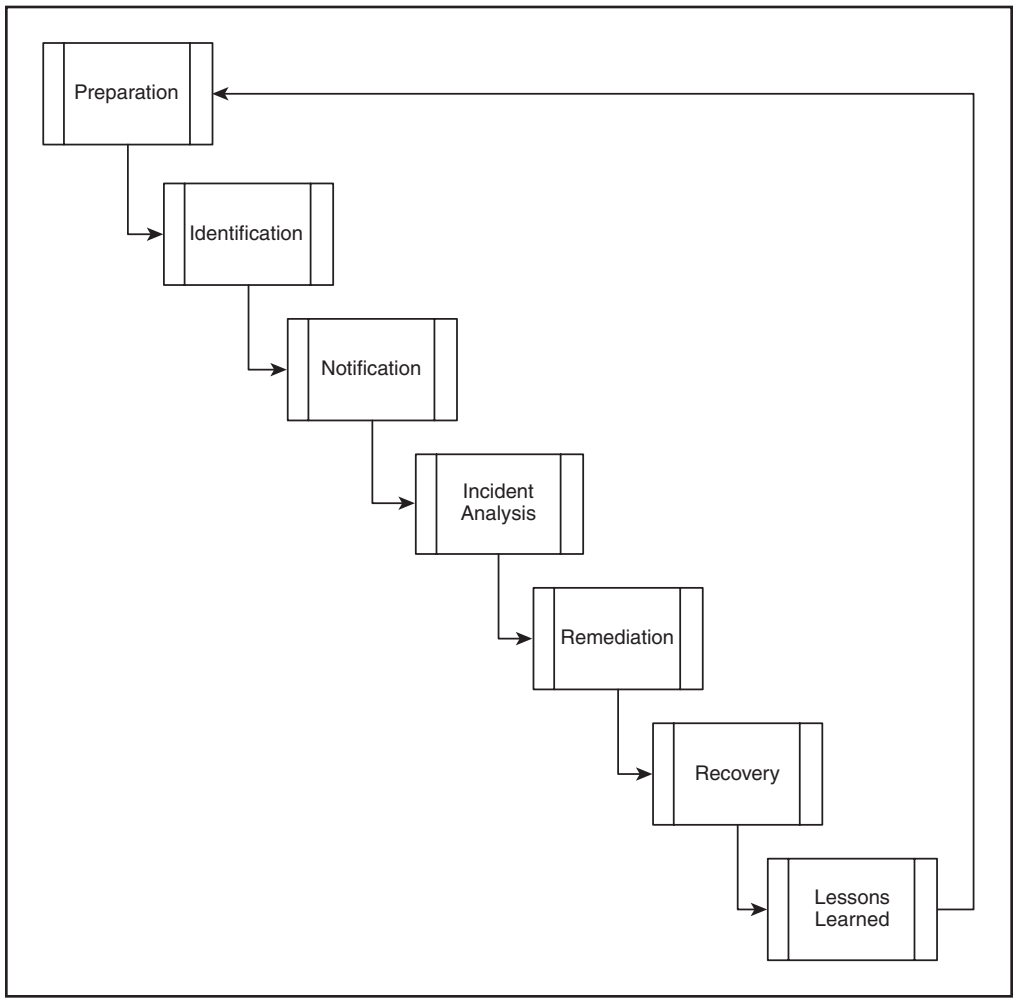


**Figure 8–1** The Incident Response Life Cycle

The terms used to describe the various stages of incident response may vary somewhat from publication to publication. Despite the differing terminology used to describe the various phases, most agree that the process is cyclical and addresses many of the issues outlined in the different phases described in this chapter. The terms used to describe each phase and the number of phases may vary, but the basic elements will generally remain very similar.

The level and detail of response required in each phase depend on the type and severity of the incident. In the sections that follow, the worst-case scenario of a compromise has been used as the main focus for discussion points. Actual steps taken, however, may vary according to the type of incident, and some steps may be skipped or significantly condensed. In addition, if the reported activity turns out to not be an incident, the process may be aborted at any time, at which point the team resumes the preparation phase.

## Step One: Preparation (Preparing for Compromise)

Incident response always begins with the steps taken to protect the organization's information resources before an incident takes place. These steps may be realized through the documentation of specific policies and procedures, end user awareness training, hardening of operating systems, installation of security tools, and the like. Just as security affects everyone in the organization at all levels, so should security safeguards be implemented at all levels.

*Policies and Procedures.* Specific policies and procedures that should be documented for the organization in preparation for an incident include the following:

- Computer security policy or policies
- Incident response procedures
- Recall procedures
- Backup and recovery procedures

As previously noted, successful computer security begins with policies. The policies provide the foundation from which a security program is built, and provide "reference points for a wide variety of information security activities including: designing controls into applications, establishing user access controls, performing risk analyses, conducting computer crime investigations, and disciplining workers for security violations."[1] To be completely successful, the policies

1. Wood, Charles Cresson. *Information Security Policies Made Easy,* Version 7. Sausalito, CA: Baseline Software, 1999, p. 1.

must be clearly and concisely written, and enforced by management. To document policies but then not enforce them through human resource and legal action diminishes the writing of the policies to simply a paperwork drill.

Developing computer security policies involves identifying key business resources and supporting policies, defining specific roles in the organization, and determining the capabilities and functionality of those roles. One inclusive policy may be written for an organization or a shorter, overarching policy may be documented with smaller supporting policies written separately to address specific concerns. Examples of policies that could be addressed include the following: user account policy, remote access policy, acceptable use policy, firewall management policy, consent to monitoring policy, and special access policy.

The organization's security policy not only is important to communicate to the employees or members of the group what is authorized activity, but also may prove valuable should an authorized user intentionally abuse his or her privileges. Supporting documentation such as end-user agreements can prove quite useful for prosecuting or addressing an insider threat. Providing a copy of the security policy to end users and having them sign end-user agreements after they have read the policy is the approach taken by many proactive organizations that use such agreements. The security policy may also be summarized in an information packet or bulletin that provides a ready reference for end users in shorter fashion, thereby reinforcing the larger document, which may not be closely read. Some organizations choose to promote awareness of existing policies even further, by requiring employees to attend a "Security 101" course when they begin employment. The more ways in which the policy can be consistently communicated and reinforced, the better the chances for a successful implementation of the document.

Security policies should indicate management support for the computer incident response program, by identifying the incident response team as a key business resource. Security practices such as routinely changing passwords and using unique passwords should be specified in the appropriate policy documentation. The policies should also indicate the responsibility of the incident response team to perform the services assigned to it, such as vulnerability assessments, reporting requirements, and monitoring of systems.

Many security policies are supported through additional documentation in the form of written procedures. The procedures are intended to provide step-by-step guidelines for enforcing the policy. Every incident response team should have documented procedures available for immediate access by the team members. The procedures should identify the roles and responsibilities of the team, as well as offer step-by-step instructions for performing the assigned tasks. Flowcharts can be extremely useful tools for incident-handling procedures and can aid with the clarification of steps to be taken during a crisis situation. The procedures should address every service or responsibility of the team in detail, from start to finish. Examples of flowcharts and processes that may be addressed in the incident response procedures include the following:

- Responding to a "virus warning" inquiry or other request for information
- Monitoring intrusion detection systems
- Processing each type of event or incident that is reported (e.g., successful intrusion, attempted intrusion, denial-of-service attack, probes or scans)
- Eradicating a computer virus
- Entering information into the trouble ticket system or database
- Reporting incidents to law enforcement
- Reporting incidents to other teams
- Conducting penetration tests
- Responding to reported activity

Outlining processes in the form of a flowchart can be a quite valuable exercise in documenting the procedures, as it will force each step to be examined in detail. The simplified flowchart in Figure 8–2 shows how the flowchart can assist with outlining the procedures to be followed. Most flowcharts will not be so simple, however, and may actually require multiple pages to document. The simple version here is included to reinforce the point of using flowcharts. Note that flowcharts should also be accompanied by supporting verbiage and not used as the sole method of procedural documentation.

Incident-handling procedures should prioritize how incidents are to be managed when more than one response is needed. Depending on

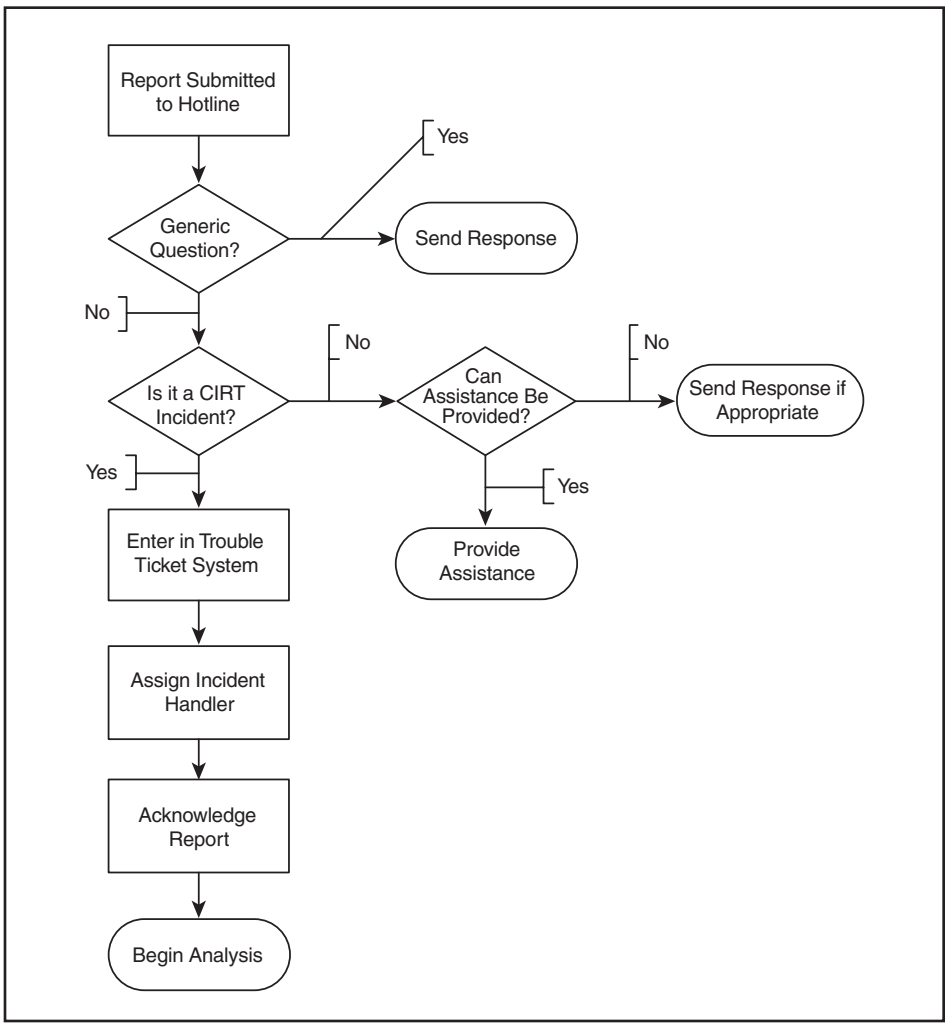


**Figure 8–2** Sample Flowchart

the size of the team's constituency, it may not be uncommon for multiple incidents to be handled simultaneously. Prioritizing the order in which incidents are tackled can assist with resource assignment issues during peak periods. For example, the team may decide to prioritize the assignment of resources to incidents as follows:

1. Ongoing attack (intruder is currently in the system or resources are being denied on an increasing or large scale)

2. Successful attack on systems identified as critical to the business or operations
3. Root compromise
4. Level of severity of the attack (i.e., number of systems affected, type of attack)
5. External tip is received that requires investigation

The actual order of priorities chosen will vary between organizations, and each group must determine up front the best priorities for response for its particular constituency. Typically, the priorities assigned will reflect the location of the system, the type of data maintained on the system, and the potential impact of the loss of that system on overall operations.

*Establish Response Guidelines.* In addition to determining the priorities of response, the team should discuss various response guidelines with upper management to determine actions that are deemed acceptable and those that may need consultation prior to the action being taken. For example, if a Web site is compromised, can the CIRT authorize the system being taken off-line without higher approvals?

One of the best methods to work out the acceptable response guidelines is to discuss various scenarios with upper management and walk through the possible response procedures that may be followed. Conducting tabletop exercises of various theoretical situations can be an excellent tool for developing response guidelines, as well as for training at various levels. Former New York City Mayor Rudolph Giuliani discusses how the use of tabletop exercises with his staff enabled him to remain calm during the attacks on the World Trade Center and thereafter in September 2001. Stated Giuliani, "We conducted tabletop exercises designed to rehearse our response to a wide variety of contingencies. We'd blueprint what each person in each agency would do if the city faced, say, a chemical attack or a biochemical attack. . . . The goal was to build a rational construct for myself, and for the people around me. I wanted them ready to make decisions when they couldn't check with me."[2] The middle of an incident is not the time to discuss

2. Giuliani, Rudolph. *Leadership.* New York: Miramax Books, 2002, pp. 62–63.

with management what options for response are available, unless it is a unique incident deemed critical to the business.

*Establishing Contacts.* Points of contact outside of the immediate team and recall information for all team members should be readily available and kept current for when the need arises to make an immediate notification. The recall list should identify all team members and give a priority ranking for contacting them in an emergency situation should they need to be recalled to work. Co-locating this list with the written incident-handling procedures can assist with locating the information quickly in a stressful situation. Contact information for law enforcement agencies and emergency personnel (e.g., fire, police) should also be included in the outside contacts list.

> **NOTE** Some groups place the contact information within the procedures document. This approach is not recommended as it is more difficult to update the information when a change in the team occurs. By keeping the names and contact information separate, it is much easier to keep the list current. Additionally, the contacts are easier to locate in one central list as opposed to searching through the entire procedures document.

*Backup Procedures.* Backup procedures should be documented and strictly followed. The importance of a good backup plan is never quite as evident as when a major incident has occurred and reliable data are needed quickly to restore operations. Multiple copies of backups should be maintained and stored off-site for added protection. The backup procedures and media should be periodically tested to help ensure their integrity.

*Evaluate System Security.* Computer systems should be routinely evaluated for their overall level of security. Operating systems should be hardened to help protect against well-known vulnerabilities being exploited. ("Hardening an operating system" refers to locking down a system to ensure that it is not providing too much access or running unnecessary services.) Vulnerability assessments and penetration tests are the most common methods used to evaluate the security of a

system. These tests can be performed by internal personnel or outsourced to a consulting firm.

*Warning Banners.* Warning banners should be posted on systems at the "points of access." The purpose of the warning banner can vary, but it is normally used to indicate that the system is private and that use of the system is subject to monitoring. Figure 8–3 provides a sample warning banner. It is strongly recommended that the organization's legal counsel review the wording used in the banner to ensure that the goals for its use are met.

The placement of warning banners has been a topic of much debate between the security and legal communities for some time. The U.S. Department of Justice contends that the warning banner must be seen by the intruder upon entering the system for it to be recognized as "off limits" or subject to monitoring. With this idea in mind, the banner should be displayed on every point of access to the system (e.g., all open ports). The counter-argument cites the "no trespassing" sign analogy. If a "no trespassing" sign is posted on a fence, does it have to be posted at the very spot where an intruder jumps the fence

*Users (authorized or unauthorized) have no explicit or implicit expectation of privacy.*

Any use of this system and the files maintained or processed on this system may be intercepted, monitored, recorded, copied, audited, inspected, and disclosed to authorized management and law enforcement personnel, as well as authorized officials of other agencies. *By using this system, the user consents to such interception, monitoring, recording, copying, auditing, inspection, and disclosure at the discretion of management.*

*Unauthorized or improper use of this system may result in administrative disciplinary action and civil and criminal penalties. By continuing to use this system you indicate your awareness of and consent to these terms and conditions of use. LOG OFF IMMEDIATELY if you do not agree to the conditions stated in this warning.*

**Figure 8–3** Sample Warning Banner

for him or her to be cited with trespassing? This sticking point, as with many security issues, will remain for the courts to iron out as more case law is established. In the meantime, it is recommended that warning banners be used and posted at as many points of entry to the network as possible.

*Security Tools.* In addition to setting up the safeguards already mentioned, the preparation phase may include the implementation of specific security tools. Firewalls, intrusion detection systems, secure identification devices, biometrics, encryption programs, and other such tools may increase the overall security of the infrastructure. Completing a risk analysis as described earlier in this book can help with identifying the right tools for the organization. Ideally, several layers of security should be incorporated in the infrastructure to provide "defense in depth." The more layers a perpetrator has to transcend to reach the most sensitive information, the less likely he or she is to succeed.

*Training.* Training requirements need to be considered and implemented for all levels of the organization. Computer security statistics from numerous surveys have helped to substantiate the threat posed by a lack of end-user awareness. End users need to be made aware of not only the basics of computer security, but also the presence of the incident response team and the need to notify the team of suspicious activity. Training is a security countermeasure that is estimated to require up to a 10 percent investment in resources and effort of the security/incident response team but will typically have a 90 percent return on investment if done correctly. Awareness training is a powerful tool that should not be overlooked.

## Step Two: Incident Identification

Incident identification normally begins with someone or something noticing activity that appears suspicious. The following are examples of how this might occur:

- An end user notices that the system indicates an incorrect time for when he or she last logged into the system.

- A system administrator notices that an authorized user has higher privileges than were assigned.
- An end user notices that a file has been modified, yet no one else should have access to it.
- An outside organization notices probing activity against its site stemming from your site.
- System performance begins to unexpectedly slow down and the cause is not readily apparent.
- An intrusion detection system sends an alarm to the management console, drawing attention to a violation of the signature file.
- Firewall logs have a gap or period of time with no activity accounted for during normal operating hours.
- The organization's Web page is listed on a well-known "hacker Web site" as having been compromised.
- An end user reports additional files in a personal directory that he or she did not create or store there.

These are just a few examples of how suspicious activity may be noted. Once it is detected, the activity should be reported to the computer incident response team and investigated to determine whether an incident has occurred.

Every response team should have a report form that identifies the information needed to investigate and track an incident. The report form should be available to all members of the constituency and may be posted on an intranet, documented in security policies or procedures, or provided as a separate file or form. When suspicious activity is noted, the form should be completed and submitted to the team. Some of the information requested on the form may not be immediately available, but that omission should not hinder the reporting of the activity. Constituents should be encouraged to report the activity as quickly as possible so that the proper level of response can be initiated. Routinely, follow-up communications between the team and the person submitting the report will take place to gather further details.

Most incident response teams also provide a phone number for reporting suspicious activity verbally. Teams with a dispersed constituency should try to provide a toll-free number for such reporting. The e-mail address and phone numbers (local and toll-free) should try to follow the same naming convention, if possible. Using an easy-to-

remember address and phone number will help the team to be more accessible during a crisis situation.

The information requested on the incident report form will vary from team to team. The CERT CC report form is available on its Web site at http://www.cert.org and may be used as a model for developing other forms. Appendix A also provides a sample incident report form that may be used or edited according to the requirements of the organization. The information requested in the form should mirror specific data fields in the incident tracking database or trouble ticket system. Whenever possible, excerpts of audit logs, copies of suspicious e-mails (including header information), and other supporting documentation should be submitted with the report form to help with the investigation.

**NOTE** Extreme care should be given to not use the system that is being attacked to report the incident. Doing so may tip off the attacker that he or she has been discovered. Whenever possible, some form of out-of-band communications should be used. In other words, a different system, the phone, a fax machine, or a mode of communication other than the attacked system should be used to make the report.

The incident report form may be completed by someone external to the team and sent in, or it may be completed by a team member who receives a report via the phone or notices the suspicious activity directly. As soon as the report is received, it should be entered into the trouble ticket system and assigned a unique tracking number, and the team should acknowledge to the person sending the report that it has been received. The reporting party should be given the tracking number to use in case any further information or activity is detected. (The importance of acknowledging the receipt of reports will be addressed in more detail later in this chapter.)

At this point, the report should be reviewed to try and determine the circumstances surrounding the suspicious activity. If more information is needed to gain a clearer understanding of the events, then the appropriate party should be contacted and a request for information made. Depending on the nature of the report, the appropriate

party may be an end user who submitted the report, someone from system administration who may have access to audit logs or other supporting evidence, or even an outside party such as the incident response team for an Internet service provider (ISP). Again, care should be taken to not divulge more information than is absolutely necessary. Sometimes the actual attacker may become involved in the incident communications without the incident handler realizing that he or she is talking to the attacker. It is always better to err on the side of caution throughout the incident-handling process.

Suspicious activity does not always equate to a computer incident (e.g., a misconfigured router reported as probing or scanning a network). As the information reported and obtained is reviewed, the incident handler should be able to determine whether an incident has occurred. If it appears that no incident has taken place, then the trouble ticket should be closed and the outcome should be communicated to the reporting party. If an incident has occurred, however, the trouble ticket should remain open until no further action is required on the part of the incident response team. Furthermore, if a successful attack (e.g., unauthorized access, denial of service) has occurred, the following steps should be taken:

1. A complete system backup should be made and stored in a safe location for further investigation and use.
2. All observations noted and steps taken in the course of investigating the incident should be logged as the analysis proceeds.
3. The appropriate notifications should be made.

### Step Three: Notification

If an incident has occurred, the proper authorities need to be notified. Proper authorities may include upper management, law enforcement, another incident response team, or others as identified in the incident response procedures. It is extremely important that escalation procedures be documented up front, before an incident occurs. This will help to eliminate or reduce the room for error during a crisis situation. The procedures should address who is to be notified for each type of incident and at what point. The incident response team leader must

always apply a certain amount of subjectivity, but the documentation of guidelines will cover the majority of situations. For the cases not covered, the best rule of thumb is to report the activity to upper management when in doubt.

As previously discussed, the integrity of the team can never be ignored. The level of trust that a team enjoys with its constituency and others will directly affect the team's level of success. Wrongful disclosure of incident information can quickly reduce the level of trust between the team and the people whom the team supports, as well as between the team and other teams and organizations. The incident response procedures should provide clear guidance on how and when information concerning an incident may be disclosed. This guidance should take into account any restrictions imposed on the team by upper management or outside organizations, and consider requests the team may receive for such information. For example, a government team should address Freedom of Information Act requests, a team in the health care industry should address requirements imposed by the Healthcare Information Portability and Accountability Act (HIPAA), and teams in financial organizations should address any requirements resulting from the Gramm-Leach-Bliley Act.

Teams should always hold the information reported to them in the strictest confidence. Many teams will not share the information beyond the immediate team members. Sometimes, however, information must be provided outside of the team (e.g., warning a site that it has been the victim of an attack or contacting law enforcement when it appears a computer crime has been committed or human life may be in danger, such as when an emergency response system has been disabled). When information does need to be shared, the purpose for disclosing the information, the requesting party, and the category of the information should be considered to determine if and how the information should be provided. More specifically:

- Disclosure of information concerning the incident should always be guarded and limited only to those with a valid "need to know." Even if people are aware of the incident and ask questions out of curiosity, the information should not be shared with them unless they have a valid reason for knowing the facts. This limitation

even applies to incident response team members who may not have a role in the incident validating their awareness.

- The entity that is to receive the information will govern the amount or type of information to be provided. For example, a response to an inquiry received from the news media about a specific incident will drastically differ from the level of detail provided to upper management. Likewise, if the incident involves law enforcement, much more detail will be provided to the officer or agent working the case than possibly even the owner of the system, especially when an insider threat is suspected.
- The category of the information that is to be provided will also determine the extent and type of data to be disclosed. For example, information deemed necessary for public release will be more generic with less detail than that released internally to the organization.

In addition, it should always be remembered that once the information is shared or disclosed, control over that data is effectively lost. Therefore the information can easily be disclosed to other parties, and that spread can come back to haunt the incident response team.

The timing of when the information is provided should also be taken into consideration. Validating reports and facts to gain the full picture of what has taken place normally takes time. It would be nice to be able to delay reporting until the full picture is clear, but this option is not always prudent. Sometimes immediate notification must be made, such as when a potential threat to other systems or even human life is at stake. In these cases, an initial report should provide a "heads up" to the activity and annotate the fact that the investigation is still proceeding to gain further information.

How the information is reported or disclosed will vary as well. In some cases, the notification may occur through the completion of a report. If a specific report format must be used, that format should be identified in the incident response procedures with directions on what information is to be given in each section. The basic questions of *who, what, why, when,* and *how* can provide a simple format for reporting that covers the major elements. As with the initial report made to the team, care should be taken to not use the attacked system as the

medium for submitting the report. If the report must be transmitted over a nonsecure medium such as the Internet, extra security precautions such as the use of encryption or a virtual private network (VPN) should be used to further protect against disclosure of the information. The report may also be made via a phone call. Again, if specific pieces of information will be provided in the phone notification, these elements should be documented and explained in the incident-handling procedures.

Although notification is listed as step 3 in this incident life cycle, with serious or large incidents this process will normally continue through the remaining steps. Successful attacks will typically require an initial notification as well as follow-up reports as the situation is investigated further. If multiple reports are provided, a method should be invoked for linking sequential reports together so that the flow of information can be followed and the potential for confusion is eliminated. Other papers or resources may include this phase with other phases in incident handling. It is broken out here to give it the attention it soundly deserves, as the notification and communication process can directly affect the success or failure of a response effort.

### Step Four: Incident Analysis

The process for deciding how to handle an incident includes several aspects: characterizing the incident; considering the prevailing circumstances of the information asset or environment affected by it; and weighing the costs, merits, and drawbacks of various restoration or response options. The result of this process is a course of action that represents the incident response team's best judgment about how the incident should be addressed given the circumstances and options at the time.

Without a good understanding of the cause of an incident, it is extremely difficult to select a course of action that will effectively correct and securely restore the affected information resources. To diagnose an incident or attack, incident handlers attempt to determine whether the characteristics of the incident and circumstances surrounding it have a known or previously observed cause. Diagnosis

can involve matching characteristics and circumstances with known conditions, or it could be a process of eliminating unlikely causes. Several factors influence the selection of the course of action to follow:

- The impact on and circumstances in the information environment at the time the suspicious activity occurred
- The criticality of affected information assets for business operations
- The real or potential damage caused by the incident
- The location of the system targeted

Evidence for the analysis may be provided through a variety of media, including alarms, logs, and remnant files. *Alarms* refer to the auditing data from security programs or tools that are programmed to trigger an alert when a specific event takes place. For example, the files provided by an intrusion detection system may be extremely valuable in determining the avenue used to gain access to a system. *Logs* refer to auditing files that track specific events as they occur. Both normal activity and malicious activity may be depicted in the audit logs. For example, many audit logs are set to record when users log onto the network or system. The activity of authorized users will not be of great importance to the incident analysis, unless an authorized user's account has been compromised and is used to gain access to the system. All too often, incident handling is analogous to "looking for the needle in the haystack." An audit log that notes an authorized user's account being accessed during a time period when the user is not at work or on vacation may be a clue to finding that "needle." Finally, *remnant files* refer to files or programs that an intruder may have left on the system once access was gained. Examples of remnant files include sniffer logs, password files, source code for programs, and exploit scripts that the intruder may have used to store his or her "goods" or to target other systems.

It is very common for intruders to install programs, create accounts, or open ports to allow for alternative points of access to the network while they are in the system. These changes enable later access should the intruder's initial avenue be shut off. Part of the incident analysis should include a vulnerability assessment of the tar-

geted system to help identify avenues through which the attack may have been launched and any new vulnerabilities that may have been introduced as the result of the attack. Even if the incident appears to be clear on the surface, a vulnerability assessment should always be performed as part of the incident response.

In some cases, particularly those resulting from the actions of an insider or those through which physical access to the system was obtained, the evidence or clues to the case may be found in the surrounding area and may not be limited to computer media. For example, a torn printout in the trash can or a Post-it note with a user's account name and password may reveal how account information was used to gain access to the system. The main consideration when visiting the targeted system is to not limit your view to the system itself, but rather to take into consideration the surrounding environment and any clues it may provide.

If the incident involved an insider to the organization or any kind of illegal activity, a forensic analysis should be performed, preferably by personnel trained in computer forensics. This analysis will normally begin with photographing the "crime scene" from all directions (i.e., the computer system and surrounding area) and include imaging the computer media available in the environment. The use of forensic tools that have already been scrutinized in the legal system will increase the strength of the evidence that is preserved in this way if the case goes to trial. Additionally, the chain of custody of all evidence obtained must be strictly enforced and documented. (More information concerning computer forensics is provided in Chapter 11.)

Once the cause of the attack has been determined, the recommended course of action to remedy the situation may be determined and followed. Depending on the severity level and nature of the response, the course of action may first need approval from upper management.

### Step Five: Remediation

The course of action and remediation phase normally begins by containing the incident (if applicable) and then removing the cause of the

activity. To contain an incident means to limit the exposure or spread of the event. For example:

- If a system has been compromised or accessed by an unauthorized person, containment will mean either kicking the person out of the system or limiting the intruder's reach within the system to monitor his or her activity.
- In the case of a denial-of-service attack, containment may mean limiting the systems affected by blocking ports used in the attack or isolating access between affected and unaffected segments.
- If a virus or worm is affecting operations and the antivirus software is not limiting its spread, the containment may include disconnecting infected segments or systems from the network or disconnecting points of access to the network to minimize the possible spread.
- If an authorized user is suspected of using the system for unauthorized purposes, containment may require the employee to be placed on paid administrative leave until the situation can be fully investigated.

The method(s) of containment will typically be the focus of the course of action selected by the incident response team. Because isolating or removing a computer from a network or organization can directly affect the business operations, the recommended course of action should be explained to the owner or manager of the business unit prior to steps being taken if guidelines for that particular response have not previously been agreed to by upper management. The owner or manager should be able to judge how the action will further alter operations or hurt the business and if the recommendations are acceptable.

The steps that are taken on the network or system will normally be performed by the system administrator or IT staff responsible for the system, with the incident handler providing support. Therefore, once the course of action is determined and approved, the appropriate technical resource should be identified and requested to be on-site, if not already available. Working with the actual owners of the system(s) to perform the containment (and recovery) steps will normally

be better received than having the incident response team show up and "take over" the system. Additionally, working with the technical staff or system administrators may yield additional clues about the incident. For example, they may help to identify configuration settings that may have resulted in a vulnerability that has been exploited or describe suspicious activity that may have been previously noted.

The most common method of containment is realized by completely removing a compromised system (or systems) from the network for further investigation. If resources allow, the system should be taken to a safe location to perform a root-cause analysis of the incident so as to validate the suspected source of the incident. A replacement system may be installed on the network, enabling operations to be restored while the team analyzes the compromised system. If a complete backup system is not available, then the same result may be accomplished by swapping the hard drive of the affected computer with a new hard drive.

Care must be taken with the new system to ensure that the intruder does not immediately gain access or deny services to the replacement. This consideration is particularly important if the exact cause of the incident was not fully determined in the analysis phase. The following steps may help to protect the replacement system:

- Move the host to a different IP address
- Require all users to change their passwords
- Install missing patches to guard against known vulnerabilities
- Examine trust relationships with other hosts for possible avenues to the targeted system

Containment may also mean isolating the affected computer system(s) in an effort to prevent further compromise, stop the spread of a virus or other form of malicious logic, or limit a denial-of-service attack. The isolation tactic may be taken when further evidence is desired to identify (and possibly prosecute) the intruder. There are multiple ways in which isolation of network segments may be realized. For example, adding a router or firewall to block access to other segments of a network may limit the reach of an intruder. Shutting down e-mail servers may help to stop the spread of a virus infection.

Disabling specific services or blocking ports on computer systems may help to slow down a denial-of-service attack.

Many people assume that the first step to take in the case of a successful compromise is to immediately "kick the intruder off the system or network." In reality, this is not always the case. If a computer crime appears to have been committed, the organization may decide to involve law enforcement and try to obtain further evidence that may later be used in a court of law to prosecute the intruder. Possibly a Title III (wiretap) order will be required and an intrusion detection system will be added to the compromised system to monitor for further activity. Access beyond the compromised system should be removed or restricted to limit potential exposure to other systems on the network while the monitoring takes place.

Passwords on compromised systems and systems that regularly interact with the compromised system should also be changed as part of the remediation phase. If, however, the decision has been made to "fishbowl" or monitor the intruder to gain further evidence, this step may be partially left for the recovery stage so as to not alert the intruder to the fact that he or she has been discovered. It is important to keep as low a profile as possible throughout this part of the incident life cycle. Additionally, any code that may have been compromised during the incident should be avoided while the system is still on-line, even if it has been isolated. The compromised code may include Trojan horses, which may either spread the incident further or alert the intruder to their discovery.

Once the situation has been contained and the cause of the incident is verified, then steps to remove the vulnerability should be taken. This removal may be done by changing system configurations, updating antivirus software signature files, blocking specific ports or services that are no longer needed, or undertaking more substantial work such as changes to application programs. If the incident resulted from a vulnerability for which a patch is available from the software manufacturer, then systems should be patched whenever possible before being placed back in production. Care should be taken during the patch process to ensure that the patch has not been altered (i.e., check MD5 checksums provided by the vendor prior to installation), and that it does not negatively affect the operation of applications that may be running on the system. Typically, it is best to test the

installation of the patch on one or a few systems first, prior to large-scale patching. Once the vulnerability has been removed from the equation, the incident life cycle then leads into the recovery phase.

### Step Six: System Restoration

In the worst-case scenario, restoring operations involves taking the system or systems off-line and rebuilding them. Rebuilding them may mean loading operating systems and programs from scratch or simply reloading files from backup sources. Depending on the amount of time that an intruder stayed in the system, the latest backup may not be the best one to use. The backup tape or other medium should be reviewed to determine its integrity and to decide whether it is the best source from which to restore operations. All too often, additional restoration work will be needed for the accounts or files that were affected during the incident. Nevertheless, the backup tape/medium should provide a foundation to begin that process.

In the simplest form, the typical recovery procedures can be broken down as follows:

1. Install an operating system from media that is known to be authentic, preferably from the vendor's original media.
2. Disable unnecessary services and apply secure configuration changes to applications.
3. Install appropriate vendor security patches to the operating system and all of the applications on the system.
4. Consult advisories, vendor bulletins, and security documentation.
5. Change passwords.
6. Reconnect the system.

Depending on the type of incident that has occurred, system restoration may be much easier to achieve. For example, a virus incident may just require antivirus signature files to be updated to achieve the system restoration. A scanning incident may not require any changes, if the potential vulnerability for which the person appears to have been searching is not present.

Once the system has been restored, configuration settings should be checked to ensure that they are equivalent to the initial state of the

system prior to the incident. If the incident exploited a certain configuration setting, that setting should be changed accordingly to prevent a repeat occurrence. If not done before, all passwords on the compromised system should be changed, as well as any passwords on systems that regularly interact with that system. If a patch or fix exists for the vulnerability that was exploited to launch the attack, then the patch or fix should be made to the system. Finally, the system should be checked to ensure that it is operating normally.

In some cases, it may be appropriate to not restore the system immediately. For example, if the system is no longer needed, an upgrade to the system will be made in the near future, or the identified vulnerability cannot be readily fixed, management may decide to delay the recovery procedures. Restoring a system with the same vulnerability is an invitation to repeat the incident. Recovery should take place only when the reoccurrence of the incident can be prevented and/or the security of the system is strengthened. Once the system is restored, it should be closely monitored for repeat attempts or attacks, and for vulnerabilities that may not have been discovered during the incident analysis.

### Step Seven: Lessons Learned

The final phase of the incident life cycle is always the "lessons learned" phase. Every incident that requires a response effort should be analyzed for lessons learned. The incident response team members should discuss the steps taken and address any concerns or problems encountered along the way. The review should focus on the facts and not place blame for steps that did not go well. Both the positive and the negative aspects of the incident response should be discussed. The following questions may be considered during this review:

- Were the response efforts provided appropriate? Did the selected course of action work?
- Was there enough information available to analyze the incident? If not, what else would have helped and how could that have been obtained?
- Were all appropriate parties kept informed of the status of the incident response? Was the information flow sufficient?

- Which steps went well? Which steps could be improved?
- Did the incident-handling procedures cover all needed steps or requirements? What documentation was the most helpful? Least helpful?
- Have steps been taken to prevent a reoccurrence of the incident?
- Should anyone else be alerted to the vulnerability exploited, such as a vendor?
- Might other systems within the constituency be vulnerable to the same attack? If so, what steps can be taken to mitigate the risks to those systems?
- Can the vulnerability exploited be addressed in organizational security policies? Do any policies need to be rewritten?
- Are there any other lessons learned that should be documented or acted upon?

Documentation changes identified during the review should be addressed as quickly as possible. The review may identify weaknesses in the organization's security policy or specific procedures that need to be addressed. It may not be the responsibility of the incident response team to correct these problems, but it is the responsibility of the team to notify the appropriate party of the deficiency. Providing recommended changes may help to speed the change being completed.

Another outcome of the review may be the identification of deficiencies in the incident-handling procedures. Those steps that worked well during the incident response and those that did not may be used to write improved procedures for future responses.

It may be determined that distributing an advisory is warranted to prevent a reoccurrence of the attack. This advisory may take the form of notifying a vendor of a newly discovered vulnerability, notifying the constituency of a specific vulnerability or threat, or notifying another organization of the incident or problem (e.g., CERT CC). The pertinent information should be obtained, documented clearly, and distributed as quickly as possible so that the problem may be addressed with all due speed.

Finally, a post-incident report should thoroughly document what took place and how the organization was affected by the event. The cost of the incident to the organization should be quantified (if possible), and any intangible damage or costs should be noted. These

costs should be included in the incident database for tracking statistics and generic report generation. Copies of the report should be filed in a safe location for future reference, and submitted to management as required.

Upon completion of the report, the incident should be closed in any trouble ticket system, as no further action should be required of the incident response team. Of course, just because an incident is closed, it does not mean it cannot be reopened. On numerous occasions, incidents have been reopened due to new evidence, new activity, or a reoccurrence of activity that appears to originate from the same source. At this point, the incident life cycle begins again.

## Sample Incidents

Building on the discussion of the incident life cycle just presented, let's examine the phases of this life cycle through a couple of hypothetical incidents. Because the preparation step will remain the same for all incidents, we will begin this discussion with step 2, identification. As previously noted, the discussion of each step focused on the system compromise, so our hypothetical cases will use other types of incidents for further review.

### *Unauthorized Use Example*

1. *Identification:* An employee reports to her manager that a coworker is spending a great deal of time surfing the Internet instead of doing his job. The employee further states that some of the sites the coworker visits are offensive to her. The manager contacts human resources, which in turn asks the CIRT to investigate the situation further.
2. *Notification:* In this case, the manager is already aware of the report and does not need to be notified. If the employee is suspected of having or accessing child pornography, however, then upper management should be made aware of the potential violation.
3. *Analysis:* This type of incident will require a forensic examination of the employee's computer system and any other computer media he uses. The manager may provide the system to the CIRT for the

imaging and review, or the incident handlers may need to image the drives at the person's desk—possibly after hours, so as to not tip him off. If no unauthorized programs or files are found from the analysis, then the case may be closed and the team returns to the preparation phase.

4. *Remediation:* If unauthorized files or images are found on the employee's system, the focus would be on the human factor—that is, discussing the violation with the employee and human resources or management taking the proper steps to address the situation from their angle. For example, the employee may be given a written warning to not use the business computer for this sort of activity in the future.
5. *System Restoration:* System restoration in this case would require the hard drive or any other corrupt media to be cleaned and authorized software to be reinstalled.
6. *Lessons Learned:* The remediation steps taken depend on the proper policies (human resources or computer security) being in place. If a policy is lacking, then a lesson learned may be to strengthen the documented policy. If lack of awareness is an issue, then the lesson learned may be to enhance the awareness of existing policies.

### *Attempted Unauthorized Access*

1. *Identification:* An employee reports to the CIRT that every morning when she comes into the office, her system is turned on and another user's ID appears in the login screen. She does not recognize the user's ID and she "swears" that she turned the computer off the previous night.
2. *Notification:* As this activity is not definitively an incident, the decision may be made to notify management when more information is obtained. Therefore, we will skip this phase and move right into the analysis.
3. *Analysis:* The employee's system may be examined for signs of attempted or unauthorized access. Any audit logs that are available should be checked for signs of suspicious activity. If nothing stands out and the suspicious activity continues, a video camera may be used to watch for someone accessing or attempting to

access the system after hours. If the activity stops or no substantial evidence of wrongdoing is found, the case may be closed and the team returns to preparing for the next incident.

4. *Remediation:* If an unauthorized individual is discovered as attempting to access the employee's system, physical security or other measures may be employed to keep the person out of the area in the future. For example, if the janitor is discovered to be the perpetrator, he or she may be relieved of these duties or moved to another location without computer systems. If the perpetrator is another employee, management may need to interview the individual to ascertain why the other person's system is used and counsel the worker if necessary to stop the activity.
5. *System Restoration:* As the system does not appear to have been successfully accessed, no restoration activity is needed.
6. *Lessons Learned:* If physical access to the system was not protected, some physical security measures may be implemented to bolster this protection. If awareness of authorized users is an issue, then training or counseling may be required.

### *Attempted Denial-of-Service*

1. *Identification:* The intrusion detection system (IDS) alarm on suspicious activity sounds, indicating that a denial-of-service attack has been attempted. The attack appears to have come from an account at a large ISP. No degradation of system performance is noted. The IDS team reports the alarm to the CIRT.
2. *Notification:* Because the activity does not appear to have been successful, further notification of upper management may not be immediately required. Rather, the incident may be included in a weekly or monthly incident summary.
3. *Analysis:* The incident handlers should work with the operations personnel to check system logs in an attempt to verify the alarm as valid. If it appears to be a "false positive" or "false alarm," then the case should be closed and no further activity is required.
4. *Remediation:* If further information indicates that a denial-of-service attack was really attempted, then the team may consider contacting the ISP and reporting the activity to it. Many proactive and security-conscious ISPs will address abuses by their end

users. The amount of feedback received by the CIRT will vary, but the contact is worth a try. Most problems can be reported to the address of "abuse@" followed by the ISP's domain name.

5. *System Restoration:* As the attack does not appear to have been successful, no restoration activity is needed.
6. *Lessons Learned:* There may or may not be lessons learned in this hypothetical case. If contact with the ISP was successful, that fact may be noted for future incidents involving the ISP's users. Likewise, if the contact was not successful or included problems, that fact could be noted for future reference.

## Incident Reporting

As previously noted in this chapter, every response team should have a report form that identifies the information it requires to investigate and track an incident. The information requested on the incident report form will vary from team to team, and should mirror specific data fields in the incident tracking database and/or trouble ticket system. Appendix A provides a sample form that may be used as a guide for an organization's own incident report form. The CERT CC report form is also available at http://www.cert.org and provides additional fields that may be considered.

The desired information should be clearly requested in the report form and allow little, if any, room for ambiguity. Pick lists or selection options can provide a tremendous advantage by eliminating confusion for both the reporting party and the incident handler. Pick lists are also extremely valuable in the incident database and trouble ticket systems. Without such lists, it's amazing how many different ways information may be defined and reported. Caution should be taken, however, to ensure that an "other" category is included for those cases that do not match any of the selections. The "other" selection should be accompanied by an area where specific comments may be made to expand upon the entry. Instructions for completing the report form should explain how to handle unclear areas or questions as they arise.

The report form should also ask for contact information for the reporting individual. In addition, the form should contain a place to indicate whether the e-mail address provided should not be used for

communication regarding the incident. The ability to remain anonymous should also be considered, as some people may be reluctant to make a report if they have to provide their identity.

## Feedback

Once the report form is completed and submitted to the incident response team, it should be acknowledged in some fashion. Lack of acknowledgment often leads to a feeling of helplessness and frustration on the part of the constituency. If the reporting party took the time to complete the report and submit it, then he or she should receive the satisfaction of a response indicating that someone was interested in the activity. The acknowledgment should include a unique incident or event tracking number and any pertinent information that needs to be passed on at that time. If additional information is needed, a request for it may accompany the response. At the same time, care should be taken to not disclose any information about the incident in the process of responding. The following provides a sample acknowledgment that may be used to respond to a report:

> Thank you for your incident report dated June 15, 2002. We are analyzing the information reported to ascertain what may have caused the suspicious activity noted, and we are tracking this activity as event #0601-25.
>
> We appreciate the inclusion of audit log excerpts with the report, but need to clarify the time zone for which these logs were recording activity. Please respond with the time zone used and reference the tracking number so we can update our records appropriately.
>
> If you discover any additional information concerning this activity or have any questions, please feel free to contact us at 800-123-4567.
>
> Regards,
> (Person responding or team signature)

In some cases little, if anything, can be done regarding the activity reported (e.g., a scan of network addresses). An acknowledgment or response to the report is still warranted, however. Without it, the person completing the report may be reluctant to continue reporting in

the future. In addition, despite the limited response capability that may be provided to that report, the information may still be vital to the monitoring of the organization's information security. The following is a sample acknowledgment to this type of report:

> Thank you for your incident report dated December 1, 2002. We are tracking this event as incident #1201-01. Please reference this number should you discover any additional information concerning this incident.
>
> While we may not be able to take action against this source based on the information you provided, your reports, along with those of other system administrators, will help us to understand the scope and frequency of these problems. Your information will be added to our database so we can correlate it with past activity.
>
> If you discover additional information concerning this activity or have any questions, please feel free to contact us at 800-123-4567.
>
> Regards,
> (Person responding or team signature)

If activity is reported that turns out to not be an incident, then a response should be provided indicating that fact to the person who submitted the report. Depending on how the incident tracking numbers are assigned, the acknowledgment may include a tracking number. If applicable, the acknowledgment should be an awareness education opportunity in which feedback is provided describing why the activity was not considered to be an incident. If people are reporting nonincident activity, the education provided in the feedback can help to cut down on the number of unneeded reports.

Regardless of the type of acknowledgment sent, the response should be signed either by the incident handler who is responding to the report or with a team signature. Including an individual name can add to a sense of uniqueness to the process and detract from the feeling of a "canned" response, but it can also hinder follow-on communications regarding the incident. Specifically, if the person reporting discovers further information on the event and wishes to report it, that individual may believe that he or she can speak only with the person who signed the initial response. If that person is out for a few days or working a different shift, the unavailability can add to

frustration for the reporting party and slow the communication. The person reporting should be encouraged to communicate with other members of the team as well, should they need to follow up on an incident.

## Tracking Incidents

The number of incidents processed and responded to by every team will vary according to the size of the constituency. If an organization is large enough to establish its own incident response team, then the chances are very good that the team will need a database or mechanism to store and track incident data. Many larger teams have both a database and a trouble ticket system. The trouble ticket system helps with tracking the reports as they are received as well as already open items. Typically, an incident will be in one of three states from the incident response team's perspective:

- *Open:* The incident has not been resolved and an action is required of the response team.
- *Waiting or pending:* The incident has not been resolved, but the team is waiting for further information, an investigation to be conducted, or a response from another person or group.
- *Closed:* No further action is required of the incident response team.

The trouble ticket system should annotate the current state of each incident and indicate what action is pending. It should also identify the flow of steps taken, so an incident handler responding for the first time to an incident can quickly see the history of action taken to date.

The trouble ticket system is the tool used to triage the requests and reports as they are received. Normally, the system would assist with identifying those incidents that are in the open or waiting state. Depending on the programs used for the system and the database, the data from the trouble ticket may be directly ported into the database. Likewise, if a closed incident is reopened, information previously entered into the database should be able to be recalled through the trouble ticket system if needed.

Working in combination, these systems can be extremely important tools for incident handling. Not only can they be used to track open incidents (ensuring that a report does not fall through the cracks) and store incident data, but they can also aid in the correlation of activity. The "correlation of activity" refers to the process of looking for patterns in activity that may be attributed to either the same source or the same type of attack. Running queries on the database can automate the steps taken to identify attacks stemming from the same source, targeting the same destination, using the same port or service, occurring during the same time period, involving the same "handle," or following any other pattern. Without the database, this task can be tremendously time-consuming and the human eye may overlook important incident elements that are not readily apparent.

The database should be configured with enough storage to expand based on the number of records stored. As the incident response team becomes widely known and succeeds, the number of reports will increase exponentially, and the database must be able to support this growth. The statistics in Table 8–1 were taken from the CERT CC Web site and give the number of incidents processed over the years by that team. Although CERT CC can be considered as having the largest constituency and routinely handles more incidents than many other teams, the statistics should indicate how quickly the number of incidents can grow with time.

Another important use of the database is statistics generation. Statistics regarding the number of inquiries and incidents responded to can be used to help justify the hiring of additional people for the team or the purchase of additional resources. Statistics on parts of the organization experiencing the most successful attacks may indicate an area with training deficiencies or other problems that need to be resolved. Statistics indicating a decrease in successful incidents may be used to gauge the success of security programs or particular tools (e.g., antivirus software). Depending on the specific reporting needs of the organization, the data maintained in the database can be very useful for justifying the team's existence as well as pinpointing strengths and weaknesses in particular security programs.

Although an incident response team may use many tools, the database should be considered one of the (if not the) most important tool. Therefore, due consideration should be given to the type of data-

**Table 8–1** Annual Number of Incidents Handled by the CERT CC

| Year | Number of Incidents Reported |
|---|---|
| 1988 | 6 |
| 1989 | 132 |
| 1990 | 252 |
| 1991 | 406 |
| 1992 | 773 |
| 1993 | 1,334 |
| 1994 | 2,340 |
| 1995 | 2,412 |
| 1996 | 2,573 |
| 1997 | 2,134 |
| 1998 | 3,734 |
| 1999 | 9,859 |
| 2000 | 21,756 |
| 2001 | 52,658 |
| 2002 | 82,094 |
| **Total** | **182,463** |

Source: http://www.cert.org/stats/cert_stats.html.

base purchased, the fields identified for tracking, and the hiring of a skilled database administrator for programming and maintaining the system. The database is not the area to cut costs or save dollars should funding problems emerge.

## Keeping Current

One of the biggest complaints or problems facing system administrators today is keeping current with the latest vulnerabilities. With so many identified and noted in various forums, how do system administrators know which ones to address? The same dilemma faces incident response teams. For incident handlers, the situation is often even more complex because they have more operating systems about which to be knowledgeable. How does the team keep current?

Several resources may be used to accomplish this goal. Advisories and alerts provided by other teams and vendors can be excellent tools with which to keep abreast of the latest holes and fixes. CERT CC advisories, in particular, should be monitored closely for alerts of serious problems. Advisories or postings from other teams and security groups can also provide valuable information regarding a new vulnerability.

Vendor advisories for specific systems covered by the team should be monitored as well. Patches or fixes available from the vendor should be identified in the announcement. Many vendors forward their alerts to the CERT CC, which also posts their announcements on the center's Web site. Some vendors include special notes concerning the fix that should be carefully considered before any action is taken. For example, a vendor may indicate a temporary fix is available that has not been thoroughly tested. In this case, the potential threat posed by the vulnerability must be balanced against the potential problems that may be encountered by installing a temporary fix to determine whether action should be taken immediately.

Some vendors now offer alert services that are tailored to the organization. For a fee, this service provides daily, weekly, or monthly updates through a subscription-based service geared to the operating systems and programs present in the organization. Typically the updates will rank the severity of the threats identified. Some organizations offer the service with a focus on intelligence gathering, drawing information from additional resources that may provide indications and warnings of potential threats. The type of information and spin given to each advisory or report varies between service providers and should be selected based on the specific needs of the team.

The following avenues may also be used to keep current:

- Mailing lists and newsgroups available on the Internet provide additional sources of information for keeping current. Some lists and newsgroups are better than others, and sometimes the team must sift through a lot of information to find the most applicable or valuable pieces.
- Technical groups such as the Forum of Incident Response and Security Teams (FIRST) and InfraGard (both of which were described earlier in this book) can be valuable sources for establishing contacts to provide guidance on specific issues as well as updates to the latest vulnerabilities.
- Conferences can be a valuable source of information on the latest tools, attacks, and responses. Some conferences, such as those sponsored by the SysAdmin, Audit, Network, Security (SANS) Institute, offer training that can lead to certifications.
- Training from both internal and external sources can provide updates on vulnerabilities, threats, and the latest developments for addressing those threats.
- Trade publications, books, and magazines may be useful for researching various subjects.

A team cannot afford to rely strictly on one source of information to keep current with vulnerabilities and countermeasures. The best approach is to utilize a combination of resources with time slotted for team members to conduct research. Despite the best efforts to stay up-to-date, remember to always be prepared for the unexpected.

## Writing Computer Security Advisories

Some teams decide to write their own advisories on vulnerabilities about which they want to alert their constituency. Other teams simply rely on forwarding advisories from other sources, such as those published by CERT CC or specific vendors. A few simple rules should be followed by teams chartered with writing their own advisories:

1. Keep it simple. The advisory should stick to the facts and avoid technical jargon.

2. Always include a fix or some steps to lessen the vulnerability. If a vulnerability does not have a readily available fix or countermeasure, the decision may be to not advertise the problem for further exploitation.
3. If a patch will be downloaded, include the MD5 checksum whenever it is available. The MD5 checksum is a digital signature or fingerprint of the patch and should be used to validate that the correct patch has been downloaded before it is installed on the system. Although the MD5 checksum is not foolproof, this step does significantly increase the overall security of the patch process. (Note: On some occasions, a patch has been compromised and modified to include a Trojan horse program when people are downloading it.)
4. Whenever possible, test the vulnerability and proposed fix in a lab environment to verify that the patch fixes the problem and doesn't inject other problems. This step may not always be possible, but it is a good security measure for protecting the team's integrity if it is an option.

---

**NOTE** For most teams that are internal to a specific company or organization, the testing of the patch is handled by the system administrators. The role of the CIRT is to identify and qualify the vulnerabilities, and then advise the appropriate entities of vulnerabilities, warnings, and informational bulletins. It is then the responsibility of the system administrators to test patches appropriately in their environment, troubleshoot any problems that arise, and be prepared with backups to restore the system if the patch goes bad.

---

The best format to follow with writing advisories has four parts:

1. *Problem:* Briefly describe the vulnerability—what it is and what can happen if it is exploited.
2. *Symptoms:* Identify any symptoms that may indicate the vulnerability has been exploited on a system.
3. *Fix:* Describe the steps that can be taken to prevent the

vulnerability from being exploited or to recover from an attack. Remember to include the MD5 checksum, if applicable.

4. *Point of contact information:* How can further information be obtained? Who can be contacted for questions or problems?

Advisories can be distributed to the constituency through several means. One of the most popular methods is through a list server, sending the advisory electronically through e-mail. Many teams also post copies of the advisory on their Web sites: on intranet sites, Web pages, or both. The final method is through paper versions, physically sending the report out to people or posting it on bulletin boards around corporate buildings. In this day of automation, the hard copy distribution is the method used the least, but it can be very valuable when a major event is taking place and people need to be made aware of it prior to turning computer systems on. For example, the Melissa virus made its debut on a Friday afternoon in March 1999. Offices that posted warnings on their doors before employees returned to work on the following Monday were able to give notice of the activity prior to computer systems being turned on. Very similar circumstances were experienced more recently when the Slammer worm spread in January 2003. The advisory steps taken by organizations in these cases may have helped to stop the further spread of the virus or worm by increasing the awareness of end users.

## Summary

This chapter focused on the center of the incident response puzzle, the policies and procedures. The operational aspects of computer incident response were discussed in detail, with particular attention being given to the response life cycle. The phases addressed in detail regarding the life cycle of incident response may be summarized as follows: preparation, incident identification, notification, incident analysis, remediation, and lessons learned.

Often the lessons learned will feed directly into improvements that can be made to strengthen the security of the infrastructure, thus beginning the life cycle again. This outline may be used as a starting point for documenting the procedures that an incident

response team should follow. It is reiterated here to reinforce the importance of these steps to the incident response puzzle.

This chapter also presented an overview of reporting criteria, addressing topics such as the report form, the importance of feedback, and the use of a trouble ticket system. The importance of a database for storing the incident data was discussed, pointing out the importance of this tool for incident correlation and statistic generation and tracking. Methods for keeping current with the latest vulnerabilities and trends once the team is formed were presented, as well as rules of thumb for writing and distributing advisories. Not every incident response team should write advisories. Several good sources of advisories are available that can just as easily be leveraged. Too many advisories can lead to the same problem experienced by many with respect to too many vulnerabilities: In time, people will tend to ignore the warnings if they are too frequent in distribution.

Combined, these topics outline many of the daily considerations that must be taken into account by the team in operation. They complete the overall picture of the incident response team puzzle. The remainder of the book provides more details on issues presented in this and earlier chapters. Although the elements described in Chapter 8 will assist with the task of writing policies and procedures for the response team, it is strongly recommended that you visit the remaining chapters before those procedures are documented or finalized. Additional details in the following chapters may provide further, more granular hints on developing your procedures.