



**National Institute of
Standards and Technology**

Technology Administration
U.S. Department of Commerce

**Special Publication 800-61
(Draft)**

DRAFT Computer Security Incident Handling Guide

Recommendations of the National Institute of Standards and Technology

Tim Grance
Karen Kent
Brian Kim

NIST Special Publication 800-61
(Draft)

DRAFT Computer Security Incident Handling Guide

*Recommendations of the National
Institute of Standards and Technology*

Tim Grance, Karen Kent, Brian Kim

C O M P U T E R S E C U R I T Y

Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899-8930

September 2003



U.S. Department of Commerce

Donald L. Evans, Secretary

Technology Administration

Phillip J. Bond, Under Secretary for Technology

National Institute of Standards and Technology

Arden L. Bement, Jr., Director

Reports on Computer Systems Technology

The Information Technology Laboratory (ITL) at the National Institute of Standards and Technology (NIST) promotes the U.S. economy and public welfare by providing technical leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test methods, reference data, proof of concept implementations, and technical analysis to advance the development and productive use of information technology. ITL's responsibilities include the development of technical, physical, administrative, and management standards and guidelines for the cost-effective security and privacy of sensitive unclassified information in Federal computer systems. This Special Publication 800-series reports on ITL's research, guidance, and outreach efforts in computer security, and its collaborative activities with industry, government, and academic organizations.

**National Institute of Standards and Technology Special Publication 800-61
Natl. Inst. Stand. Technol. Spec. Publ. 800-61, ## pages (September 2003)**

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

THE PUBLIC COMMENT PERIOD FOR THIS DOCUMENT BEGINS ON
SEPTEMBER 15, 2003 AND ENDS ON OCTOBER 15, 2003.
COMMENTS MAY BE SUBMITTED TO THE COMPUTER SECURITY DIVISION, NIST,
VIA ELECTRONIC MAIL AT INCIDENTHANDLINGPUB800-61@NIST.GOV

OR VIA REGULAR MAIL AT

100 BUREAU DRIVE (MAIL STOP 8930)
GAITHERSBURG, MD 20899-8930

**U.S. GOVERNMENT PRINTING OFFICE
WASHINGTON: 2003**

For sale by the Superintendent of Documents, U.S. Government Printing Office
Internet: bookstore.gpo.gov — Phone: (202) 512-1800 — Fax: (202) 512-2250
Mail: Stop SSOP, Washington, DC 20402-0001

Acknowledgements

Acknowledgements will be added to the final version of the document.

Table of Contents

EXECUTIVE SUMMARY	1
1. INTRODUCTION.....	1-1
1.1 Authority.....	1-1
1.2 Purpose and Scope	1-1
1.3 Audience.....	1-1
1.4 Document Structure	1-1
2. ORGANIZING A COMPUTER SECURITY INCIDENT RESPONSE CAPABILITY	2-1
2.1 Events and Incidents.....	2-1
2.2 Need for Incident Response	2-1
2.3 Incident Response Policy and Procedure Creation	2-2
2.3.1 Policy and Procedure Elements	2-3
2.3.2 Sharing Information With Outside Parties	2-3
2.4 Incident Response Team Structure	2-7
2.4.1 Team Models.....	2-7
2.4.2 Team Model Selection.....	2-8
2.4.3 Incident Response Personnel	2-10
2.4.4 Dependencies Within Organizations	2-12
2.5 Incident Response Team Services	2-13
2.6 Recommendations	2-14
3. HANDLING AN INCIDENT	3-1
3.1 Preparation	3-1
3.1.1 Preparing to Handle Incidents	3-1
3.1.2 Preventing Incidents.....	3-3
3.2 Detection and Analysis	3-4
3.2.1 Incident Categories	3-4
3.2.2 Signs of an Incident.....	3-5
3.2.3 Sources of Precursors and Indications.....	3-6
3.2.4 Incident Analysis	3-8
3.2.5 Incident Documentation.....	3-12
3.2.6 Incident Prioritization	3-13
3.2.7 Incident Notification	3-15
3.3 Containment, Eradication, and Recovery	3-16
3.3.1 Choosing a Containment Strategy	3-16
3.3.2 Evidence Gathering and Handling.....	3-17
3.3.3 Identifying the Attacker.....	3-19
3.3.4 Eradication and Recovery	3-20
3.4 Post-Incident Activity.....	3-21
3.4.1 Lessons Learned.....	3-21
3.4.2 Using Collected Incident Data	3-22
3.4.3 Evidence Retention	3-24
3.5 Incident Handling Checklist.....	3-24
3.6 Recommendations	3-26
4. HANDLING DENIAL OF SERVICE INCIDENTS.....	4-1
4.1 Incident Definition and Examples.....	4-1

4.1.1	Reflector Attacks	4-1
4.1.2	Amplifier Attacks.....	4-2
4.1.3	Distributed Denial of Service Attacks	4-3
4.1.4	Synfloods.....	4-4
4.2	Preparation	4-6
4.2.1	Incident Handling Preparation	4-6
4.2.2	Incident Prevention.....	4-6
4.3	Detection and Analysis	4-7
4.4	Containment, Eradication, and Recovery	4-9
4.4.1	Choosing a Containment Strategy	4-9
4.4.2	Evidence Gathering and Handling.....	4-10
4.5	Checklist for Handling Denial of Service Incidents	4-11
4.6	Recommendations	4-12
5.	HANDLING MALICIOUS CODE INCIDENTS	5-1
5.1	Incident Definition and Examples.....	5-1
5.1.1	Viruses	5-1
5.1.2	Trojan Horses.....	5-2
5.1.3	Worms	5-2
5.1.4	Mobile Code	5-3
5.1.5	Blended	5-3
5.2	Preparation	5-3
5.2.1	Incident Handling Preparation	5-3
5.2.2	Incident Prevention.....	5-4
5.3	Detection and Analysis	5-5
5.4	Containment, Eradication, and Recovery	5-7
5.4.1	Choosing a Containment Strategy	5-7
5.4.2	Evidence Gathering and Handling.....	5-8
5.4.3	Eradication and Recovery	5-8
5.5	Checklist for Handling Malicious Code Incidents	5-9
5.6	Recommendations	5-9
6.	HANDLING UNAUTHORIZED ACCESS INCIDENTS	6-1
6.1	Incident Definition and Examples.....	6-1
6.2	Preparation	6-1
6.2.1	Incident Handling Preparation	6-1
6.2.2	Incident Prevention.....	6-2
6.3	Detection and Analysis	6-3
6.4	Containment, Eradication, and Recovery	6-5
6.4.1	Choosing a Containment Strategy	6-5
6.4.2	Evidence Gathering and Handling.....	6-6
6.4.3	Eradication and Recovery	6-7
6.5	Checklist for Handling Unauthorized Access Incidents.....	6-7
6.6	Recommendations	6-8
7.	HANDLING INAPPROPRIATE USAGE INCIDENTS	7-1
7.1	Incident Definition and Examples.....	7-1
7.2	Preparation	7-1
7.2.1	Incident Handling Preparation	7-1
7.2.2	Incident Prevention.....	7-2
7.3	Detection and Analysis	7-3

7.4 Containment, Eradication, and Recovery 7-5

7.5 Checklist for Handling Inappropriate Usage Incidents 7-5

7.6 Recommendations 7-5

8. HANDLING MULTIPLE COMPONENT INCIDENTS 8-1

8.1 Incident Definition and Examples 8-1

8.2 Preparation, Detection, and Analysis 8-1

8.3 Containment, Eradication, and Recovery 8-2

8.4 Checklist for Handling Multiple Component Incidents 8-2

8.5 Recommendations 8-3

List of Appendices

Appendix A— RECOMMENDATIONS A-1

A.1 Organizing a Computer Security Incident Response Capability A-1

A.2 Incident Preparation A-2

A.3 Incident Detection and Analysis A-4

A.4 Containment, Eradication, and Recovery A-6

A.5 Post-Incident Activity A-7

Appendix B— INCIDENT HANDLING SCENARIOS B-1

B.1 Scenario Questions B-1

B.2 Scenarios B-2

Appendix C— GLOSSARY C-1

Appendix D— ACRONYMS D-1

Appendix E— PRINT RESOURCES E-1

Appendix F— ONLINE TOOLS AND RESOURCES F-1

Appendix G— FREQUENTLY ASKED QUESTIONS G-1

Appendix H— CRISIS HANDLING STEPS H-1

List of Figures

Figure 2-1. Incident-Related Communications With Outside Parties	2-4
Figure 3-1. Incident Response Life Cycle	3-1
Figure 3-2. Incident Response Life Cycle (Detection and Analysis)	3-4
Figure 3-3. Incident Response Life Cycle (Containment, Eradication, and Recovery)	3-16
Figure 3-4. Incident Response Life Cycle (Post-Incident Activity)	3-21
Figure 4-1. Reflector Attack Using a DNS Server	4-2
Figure 4-2. Amplifier Attack.....	4-3
Figure 4-3. Distributed Denial of Service Attack	4-4
Figure 4-4. Synflood Attack.....	4-5
Figure 8-1. Example of a Multiple Component Incident	8-1

List of Tables

Table 3-1. Tools and Resources for Incident Handlers.....	3-2
Table 3-2. Common Sources of Precursors and Indications	3-7
Table 3-3. Excerpt of a Sample Diagnosis Matrix.....	3-12
Table 3-4. Sample Incident Response SLA Matrix	3-14
Table 3-5. Initial Incident Handling Checklist	3-25
Table 3-6. Generic Incident Handling Checklist for Uncategorized Incidents	3-25
Table 4-1. Denial of Service Precursors	4-8
Table 4-2. Denial of Service Indications	4-8
Table 4-3. Denial of Service Incident Handling Checklist	4-11
Table 5-1. Malicious Code Precursors	5-5
Table 5-2. Malicious Code Indications	5-6
Table 5-3. Malicious Code Incident Handling Checklist.....	5-9
Table 6-1. Actions to Prevent Unauthorized Access Incidents	6-2
Table 6-2. Unauthorized Access Precursors	6-3
Table 6-3. Unauthorized Access Indications.....	6-4
Table 6-4. Unauthorized Access Incident Handling Checklist	6-7
Table 7-1. Inappropriate Usage Indications	7-3
Table 7-2. Sample SLA for Inappropriate Usage Incidents.....	7-4
Table 7-3. Inappropriate Usage Incident Handling Checklist.....	7-5
Table 8-1. Multiple Component Incident Handling Checklist	8-3

EXECUTIVE SUMMARY

Computer security incident response—the mitigation of violations of security policies and recommended security practices—has become an important component of information technology (IT) programs. Security-related threats have become not only more numerous and diverse but also more damaging and disruptive. New types of security-related incidents emerge frequently. Preventative activities based on the results of risk assessments can lower the number of incidents, but not all incidents can be prevented. An incident response capability is therefore necessary for rapidly detecting incidents, minimizing loss and destruction, mitigating the weaknesses that were exploited, and restoring computing services. To that end, this publication provides guidelines for incident handling, particularly for analyzing incident-related data and determining the appropriate response to each incident. The guidelines can be followed independently of particular hardware platforms, operating systems, protocols, or applications.

Because performing incident response effectively is a complex undertaking, establishing a successful incident response capability requires substantial planning and resources. Continually monitoring threats through intrusion detection systems (IDS) and other mechanisms is essential. Establishing clear procedures for assessing the current and potential business impact of incidents is critical, as is implementing effective methods of collecting, analyzing, and reporting data. Building relationships and establishing suitable means of communication with other internal groups (e.g., human resources, legal) and with external groups (e.g., other incident response teams, law enforcement) are also vital.

This publication seeks to help both established and newly formed incident response teams. This document will assist organizations in establishing computer security incident response capabilities and handling incidents efficiently and effectively. More specifically, this document discusses the following items:

- + Organizing a computer security incident response capability
 - Establishing incident response policies and procedures
 - Structuring an incident response team, including outsourcing considerations
 - Recognizing which personnel groups may be called on to participate in incident response.
- + Handling incidents from initial preparation through the post-incident lessons learned phase
- + Handling specific types of incidents
 - **Denial of Service (DoS)**—an attack that prevents the usage of network, system, or application resources
 - **Malicious Code**—a virus, worm, Trojan horse, or other code-based malicious entity that infects a host
 - **Unauthorized Access**—a user gains access without permission to a network, system, application, data, or other resource
 - **Inappropriate Usage**—a user violates acceptable computing use policies
 - **Multiple Component**—a single incident that encompasses two or more incidents; for example, a malicious code infection leads to unauthorized access to a host, which is then used to gain unauthorized access to additional hosts.

Implementing the following requirements and recommendations should facilitate efficient and effective incident response for Federal departments and agencies.

Organizations must create, provision, and operate a formal incident response capability. Federal law requires Federal agencies to report incidents to the Federal Computer Incident Response Center (FedCIRC) office.

The Federal Information Security Management Act (FISMA) of 2002 requires Federal agencies to establish incident response capabilities. Each Federal civilian agency must designate a primary and secondary point of contact (POC) with FedCIRC, report all incidents, and internally document corrective actions and their impact. Each agency is responsible for determining specific ways in which these requirements are to be fulfilled.

Because computing has become so important, organizations should mitigate ever-increasing threats by responding quickly and efficiently to incidents. Doing so minimizes damage to resources and data and supports the availability of services. The concept of incident response has become widely accepted as a fundamental tenet of enterprise security. Incident response benefits organizations in other ways:

- + Information gained during incident handling can be valuable in preparing for future incident handling and in identifying systemic weaknesses that may indicate the need for stronger security controls.
- + A formal capability provides a logical source for collecting and reporting incident data, which is important because Federal civilian agencies are required to report incident information to a centralized Federal incident center.
- + Staff members will be prepared to handle incidents in a manner consistent with legal guidance, such as collecting and preserving evidence properly.

Establishing an incident response capability should include the following actions:

- + Creating an incident response policy
- + Developing procedures for performing incident handling and reporting, based on the incident response policy
- + Setting guidelines for communicating with outside parties regarding incidents
- + Selecting a team structure and staffing model
- + Establishing relationships between the incident response team and other groups, both internal (e.g., legal department) and external (e.g., law enforcement agencies)
- + Determining what services the incident response team should provide
- + Staffing and training the incident response team.

Organizations should reduce the frequency of incidents by effectively securing networks, systems, and applications.

Preventing problems is normally less costly and more effective than reacting to them after they occur. Thus, incident prevention is an important complement to an incident response capability. If security controls are insufficient, high volumes of incidents may occur, overwhelming the resources and capacity for response, which would result in delayed or incomplete recovery and possibly more extensive damage

and longer periods of service and data unavailability. Incident handling can be performed more effectively if organizations complement their incident response capability with adequate resources to actively maintain the security of networks, systems, and applications, freeing the incident response team to focus on handling serious incidents.

Organizations should document their guidelines for interactions with other organizations regarding incidents.

During incident handling, the organization may need to communicate with outside parties, including other incident response teams, law enforcement, the media, vendors, and external victims. Because such communications often need to occur quickly, organizations should predetermine communication guidelines so that only the appropriate information is shared with the right parties. If sensitive information is released inappropriately, it can lead to greater disruption and financial loss than the incident itself. Creating and maintaining a list of internal and external POCs, along with backups for each contact, should assist in making communications among parties easier and faster.

Organizations should emphasize the importance of incident detection and analysis throughout the organization.

Incidents can be detected primarily through the following means:

- + Security software, such as IDSs and antivirus scanners
- + Logs from operating systems, applications and network devices
- + Reports from people within and outside the organization.

In an organization, thousands or millions of possible signs of incidents may occur each day. Automation is needed to perform an initial analysis of the data and select events of interest for human review. Event correlation software and centralized logging can be of great value in automating the analysis process. However, the effectiveness of the process depends on the quality of the data that goes into it. Organizations should establish logging standards and procedures to ensure that adequate information is collected by logs and security software and that the data is reviewed regularly.

Organizations should create written guidelines for prioritizing incidents.

Prioritizing the handling of individual incidents is a critical decision point in the incident response process. Incidents should be prioritized based on the following:

- + Criticality of the affected resources (e.g., public Web server, user workstation)
- + Current and potential technical effect of the incident (e.g., root compromise, data destruction).

Combining the criticality of the affected resources and the current and potential technical effect of the incident determines the business impact of the incident—for example, data destruction on a user workstation might result in a minor loss of productivity, whereas root compromise of a public Web server might result in a major loss of revenue, productivity, access to services, and reputation, as well as the release of confidential data (e.g., credit card numbers, Social Security Numbers).

Incident handlers may be under great stress during incidents, so it is important to make the prioritization process clear. Organizations should decide how the incident response team should react under various circumstances, and then create a Service Level Agreement (SLA) that documents the appropriate actions and maximum response times. This documentation is particularly valuable for organizations that

outsource components of their incident response programs. Documenting the guidelines should facilitate faster and more consistent decision-making.

Organizations should use the lessons learned process to gain value from incidents.

After a major incident has been handled, the organization should hold a lessons learned meeting to review how effective the incident handling process was and identify necessary improvements to existing security controls and practices. Lessons learned meetings should also be held periodically for lesser incidents. The information accumulated from all lessons learned meetings should be used to identify systemic security weaknesses and deficiencies in policies and procedures. Follow-up reports generated for each resolved incident can be important not only for evidentiary purposes but also for reference in handling future incidents and in training new incident response team members. An incident database, with detailed information on each incident that occurs, can be another valuable source of information for incident handlers.

1. INTRODUCTION

1.1 Authority

The National Institute of Standards and Technology (NIST) developed this document in furtherance of its statutory responsibilities under the Federal Information Security Management Act (FISMA) of 2002, Public Law (P.L.) 107-347.

NIST is responsible for developing standards and guidelines, including minimum requirements, for providing adequate information security for all agency operations and assets; but such standards and guidelines shall not apply to national security systems. This guideline is consistent with the requirements of the Office of Management and Budget (OMB) Circular A-130, Section 8b(3), “Securing Agency Information Systems,” as analyzed in A-130, Appendix IV: Analysis of Key Sections. Supplemental information is provided in A-130, Appendix III.

This guideline has been prepared for use by Federal agencies. It may be used by nongovernmental organizations on a voluntary basis and is not subject to copyright, though attribution is desired.

Nothing in this document should be taken to contradict standards and guidelines made mandatory and binding on Federal agencies by the Secretary of Commerce under statutory authority, nor should these guidelines be interpreted as altering or superseding the existing authorities of the Secretary of Commerce, Director of the OMB, or any other Federal official.

1.2 Purpose and Scope

This publication seeks to assist organizations in mitigating the risks from information security incidents by providing practical guidance on responding to incidents effectively and efficiently. Agencies are encouraged to tailor the recommended guidelines and solutions to meet their specific security or business requirements. This guide replaces NIST Special Publication 800-3, *Establishing a Computer Security Incident Response Capability (CSIRC)*.

This document presents general incident response guidelines that are independent of particular hardware platforms, operating systems, and applications. Specifically, it includes guidance on establishing an effective incident response program, but the primary focus of the document is detecting, analyzing, prioritizing, and handling incidents.

1.3 Audience

This document has been created for computer security incident response teams (CSIRT), system and network administrators, security staff, technical support staff, chief information officers (CIO), chief security officers (CSO), and computer security program managers who are responsible for preparing for, or responding to, security incidents.

1.4 Document Structure

The remainder of this document is organized into seven major sections. Section 2 discusses the need for incident response, outlines possible incident response team structures, and highlights other groups within an organization that may participate in incident handling. Section 3 reviews the basic incident handling steps and provides advice for performing incident handling more effectively, particularly incident detection and analysis. Sections 4 through 8 provide specific recommendations for handling five types of

incidents: denial of service (DoS), malicious code, unauthorized access, inappropriate usage, and multiple component.

Appendix A contains a consolidated list of recommendations for incident response. Appendix B contains incident response scenarios and questions for use in incident response exercises. Appendices C and D contain a glossary and acronym list, respectively. Appendix E lists print resources, and Appendix F identifies online tools and resources that may be useful in planning and performing incident response. Appendix G covers frequently asked questions about incident response. Appendix H lists the major steps to follow when handling a security incident-related crisis.

2. ORGANIZING A COMPUTER SECURITY INCIDENT RESPONSE CAPABILITY

2.1 Events and Incidents

An *event* is any observable occurrence in a system or network. Events include a user connecting to a file share, a server receiving a request for a Web page, a user sending electronic mail (e-mail), and a firewall blocking a connection attempt. *Adverse events* are events with a negative consequence, such as system crashes, network packet floods, unauthorized use of system privileges, defacement of a Web page, and execution of malicious code that destroys data. This guide will address only adverse events that are computer security-related and excludes adverse events caused by sources such as natural disasters and power failures.

The definition of a *computer security incident* has evolved. In the past, a computer security incident was thought of as a security-related adverse event in which there was a loss of data confidentiality, disruption of data or system integrity, or disruption or denial of availability. New types of computer security incidents have emerged since then, necessitating an expanded definition of an incident.¹ An incident can be thought of as a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices.² Examples of today's incidents are as follows:

- + **Denial of Service**—an attacker sends specially crafted packets to a Web server, causing it to crash.
- + **Malicious Code**—a worm uses open file shares to quickly infect several hundred workstations within an organization.
- + **Malicious Code**—an organization receives a warning from an antivirus vendor that a new virus is spreading rapidly throughout the Internet.
- + **Unauthorized Access**—an attacker runs an exploit tool to gain access to a server's password file.
- + **Unauthorized Access (Extortion)**—a perpetrator obtains unauthorized administrator-level access to a system and then threatens the victim that the details of the break-in will be released to the press if the organization does not pay a designated sum of money.
- + **Inappropriate Usage**—a user provides illegal copies of software to others through peer-to-peer file sharing services.
- + **Inappropriate Usage (Harassment)**—a person threatens another person through e-mail.

2.2 Need for Incident Response

Incident response has become necessary because attacks frequently cause the compromise of personal and business data.³ Malicious code incidents such as the SQL Slammer worm,⁴ the Blaster worm,⁵ and the Love Letter worm⁶ have disrupted or damaged millions of systems and networks around the world.

¹ For the remainder of this document, the terms "incident" and "computer security incident" will be interchangeable.

² Violations of computer security policy and acceptable use policy are likely to be detected using the same means. In practice, incident response teams typically handle many acceptable use policy violations. Section 7 discusses this issue in more detail.

³ In the first 6 months of 2003, more than 76,000 cyber incidents were reported to CERT[®]/CC. Additional incident statistics are available at http://www.cert.org/stats/cert_stats.html.

⁴ "SQL" refers to Microsoft's Structured Query Language (SQL) Server product. More information on the worm is available at <http://www.cert.org/advisories/CA-2003-04.html>.

⁵ <http://www.cert.org/advisories/CA-2003-20.html>

⁶ <http://www.cert.org/advisories/CA-2000-04.html>

Heightened national security concerns are also raising awareness of the possible effects of computer-based attacks. These events—and many more—make the case daily for responding quickly and efficiently when computer security defenses are breached. To address these threats, the concept of computer security incident response has become widely accepted and implemented in the Federal Government, private sector, and academia.

The following are benefits of having an incident response capability:

- + Helping personnel to recover quickly and efficiently from security incidents, minimizing loss or theft of information, and disruption of services
- + Responding to incidents systematically so that the appropriate steps are taken
- + Using information gained during incident handling to better prepare for handling future incidents and to provide stronger protection for systems and data
- + Dealing properly with legal issues that may arise during incidents.

Besides the business reasons to establish an incident response capability, Federal departments and agencies must comply with law, regulations and policy directing a coordinated, effective defense against information security threats. Chief among these are the following:

- + OMB’s Circular No. A-130, Appendix III,⁷ which directs Federal agencies to “ensure that there is a capability to provide help to users when a security incident occurs in the system and to share information concerning common vulnerabilities and threats. This capability shall share information with other organizations ... and should assist the agency in pursuing appropriate legal action, consistent with Department of Justice guidance.”
- + FISMA,⁸ which requires agencies to have “procedures for detecting, reporting, and responding to security incidents” and establishes a centralized Federal information security incident center, in part to:
 - “Provide timely technical assistance to operators of agency information systems ... including guidance on detecting and handling information security incidents ...
 - Compile and analyze information about incidents that threaten information security ...
 - Inform operators of agency information systems about current and potential information security threats, and vulnerabilities”

2.3 Incident Response Policy and Procedure Creation

This section discusses policies and procedures related to incident response, with an emphasis on interactions with outside parties, such as the media, law enforcement agencies and incident reporting organizations.

⁷ <http://www.whitehouse.gov/omb/circulars/a130/a130trans4.html>

⁸ <http://csrc.nist.gov/policies/FISMA-final.pdf>

2.3.1 Policy and Procedure Elements

Policy governing incident response is highly individualized to the organization. However, most policies have the same key elements, regardless of whether the organization's incident response capability is indigenous or outsourced:⁹

- + Statement of management commitment
- + Purpose and objectives of the policy
- + Scope of the policy (to whom and what it applies and under what circumstances)
- + Definition of computer security incidents and their consequences within the context of the organization
- + Organizational structure and delineation of roles, responsibilities, and levels of authority; should include the authority of the incident response team to confiscate or disconnect equipment and to monitor suspicious activity, and the requirements for reporting certain types of incidents
- + Prioritization or severity ratings of incidents
- + Performance measures
- + Reporting and contact forms.

Procedures should be based on the incident response policy. Standard operating procedures (SOP) are a delineation of the specific technical processes, techniques, checklists, and forms used by the incident response team. SOPs should be comprehensive and detailed to ensure that the priorities of the organization are reflected in response operations. Also, following standardized responses should minimize errors, particularly those that might be caused by incident handling tempo and stress. SOPs should be tested to validate their accuracy and usefulness, then distributed to all team members. Training should be provided for SOP users; the SOP documents can be used as an instructional tool. Suggested SOP elements are presented throughout Sections 3 through 8.

2.3.2 Sharing Information With Outside Parties

The organization may need to communicate with outside parties regarding an incident. This includes reporting incidents to agencies such as the Federal Computer Incident Response Center (FedCIRC) and the CERT[®] Coordination Center (CERT[®]/CC), contacting law enforcement, and fielding inquiries from the media. Incident handlers may also need to discuss the incident with other involved parties, such as the organization's Internet Service Provider (ISP), the ISP that the attacker is using, the vendor of vulnerable software, or other incident response teams that may be familiar with unusual activity that the handler is trying to understand. An organization may want to—or be required to—communicate incident details with an outside organization for numerous reasons. The incident response team should discuss this at length with the organization's public affairs and legal advisors before an incident occurs to establish policies and procedures regarding information sharing. Otherwise, sensitive information regarding incidents may be provided to unauthorized parties; this action could lead to greater disruption and financial loss than the incident itself. The team should document all contacts and communications with outside parties for liability and evidentiary purposes.

⁹ Appendix F includes pointers to Web sites with sample incident response policies and procedural forms.

The following sections provide guidance on communicating with several types of outside parties regarding the handling of actual incidents, including the media, law enforcement, and incident reporting organizations. Figure 2-1 includes several of the outside parties with which the organization may need to communicate. The arrows indicate the direction of the communication—for example, the organization may initiate communications with software vendors. Double-headed arrows indicate that either party may initiate communications.



Figure 2-1. Incident-Related Communications With Outside Parties

2.3.2.1 The Media

Dealing with the media is an important part of incident response. The incident handling team should establish media communications procedures that are in compliance with the organization's policies on appropriate interaction with the media and information disclosure. Organizations often find it beneficial to designate a media point of contact and at least one backup contact. Ideally, all members of the incident response team should be prepared to interact with the media. The following actions should be considered for preparing those who may be communicating with the media:

- + Conduct training sessions on interacting with the media regarding incidents, which should include:
 - The importance of not revealing sensitive information, such as technical details of countermeasures (e.g., which protocols the firewall permits), which could assist other would-be attackers
 - The positive aspects of communicating important information to the public fully and effectively.
- + Establish procedures to brief media contacts on the issues and sensitivities regarding a particular incident before discussing it with the media.
- + Hold mock interviews and press conferences during incident handling exercises.

2.3.2.2 Law Enforcement

Many security-related incidents do not result in convictions because organizations do not properly contact law enforcement. Several levels of law enforcement are available to investigate incidents: Federal

investigatory agencies (e.g., the Federal Bureau of Investigation [FBI] and the Secret Service), district attorney offices, state law enforcement, and local (e.g., county) law enforcement. In addition, agencies have an Office of Inspector General (OIG) for investigation of violation of the law within each agency. The incident response team should become acquainted with its various law enforcement representatives before an incident occurs to discuss conditions under which incidents should be reported to them, and how the reporting should be performed. Note that law enforcement is usually interested only in particular types of incidents of a certain magnitude that involve criminal behavior.

Law enforcement should be contacted through designated individuals in a manner consistent with the requirements of the law and the organization's procedures. Many organizations prefer to appoint one incident response team member as the primary point of contact (POC) with law enforcement. This person should be familiar with the reporting procedures for all relevant law enforcement agencies and well prepared to recommend which agency, if any, should be contacted. Note that the organization should typically not contact multiple agencies because doing so might result in jurisdictional conflicts. The incident response team should understand what the jurisdictional issues are (e.g., physical location—an organization based in one state has a server located in another state attacked from a system in a third state, being used remotely by an attacker in a fourth state).

2.3.2.3 Incident Reporting Organizations

FISMA requires Federal agencies to report incidents to FedCIRC,¹⁰ which is a governmentwide incident response capability that assists Federal civilian agencies in their incident handling efforts. FedCIRC does not replace existing agency response teams; rather, it augments the efforts of Federal civilian agencies by serving as a focal point for dealing with incidents. FedCIRC analyzes the information provided by all agencies to identify trends and precursors of attacks, because these are easier to discern when reviewing data from many organizations than when reviewing the data of a single organization.

Each agency must designate a primary and secondary POC with FedCIRC, report all incidents, and internally document corrective actions and their impact. Each agency is responsible for determining specific ways in which these requirements are to be fulfilled. Organizations should create a policy that states who is designated to report incidents and how the incidents should be reported. FedCIRC allows agencies to report incidents online and provides long and short reporting forms.¹¹ The FedCIRC Web site provides a table listing reporting requirements, including types of incidents and the time frame for reporting each type to FedCIRC.¹² Examples of reporting requirements are a root compromise of a system containing classified information (immediate); a successful virus infection that provides unauthorized access to a unclassified system (within 2 hours of detection); and spam (monthly). All Federal agencies must ensure that their incident response procedures adhere to FedCIRC's reporting requirements and that the procedures are followed properly. This is mandatory for Federal agencies, but also beneficial for them because FedCIRC can provide better information to agencies if they receive better incident data from them.

Organizations other than Federal civilian agencies should not report incidents to FedCIRC. If an organization does not have its own CSIRT to contact, it can report incidents to other organizations, including—

¹⁰ <http://www.fedcirc.gov>

¹¹ <https://incidentreport.fedcirc.gov>

¹² <http://www.fedcirc.gov/incidentReporting/reportingRequirements.html>

- + **Information Analysis Infrastructure Protection (IAIP).**¹³ Because IAIP is part of the Department of Homeland Security (DHS), it is interested in any threats to critical U.S. infrastructures. Organizations can report incidents to IAIP using its incident report form.¹⁴
- + **CERT® Coordination Center (CERT®/CC).**¹⁵ CERT®/CC, previously known as CERT, is located at Carnegie Mellon University. This nongovernmental entity is interested in any computer security incidents that involve the Internet. CERT®/CC provides an incident reporting system for online incident reporting.¹⁶

2.3.2.4 Other Outside Parties

As previously mentioned, an organization may want to discuss incidents with several other groups, including:

- + **The Organization's ISP.** During a network-based DoS attack, an organization may need assistance from its ISP in blocking the attack or tracing its origin.
- + **Owners of Attacking Addresses.** If attacks are originating from an external organization's IP address space, incident handlers may want to talk to the designated security contacts for the organization, to alert them to the activity or to ask them to collect evidence. Handlers should be cautious if they are unfamiliar with the external organization because the owner of the address space could be the attacker or an associate of the attacker.
- + **Software Vendors.** Under some circumstances, incident handlers may want to speak to a software vendor about suspicious activity. This contact could include questions regarding the significance of certain log entries or known false positives for certain intrusion detection signatures, where minimal information regarding the incident may need to be revealed. More information may need to be provided in some cases—for example, if a server appears to have been compromised through an unknown software vulnerability. Incident handlers may have other questions for vendors, such as the availability of patches or fixes for new vulnerabilities.
- + **Other Incident Response Teams.** Groups such as the Forum of Incident Response and Security Teams (FIRST)¹⁷ promote information sharing among incident response teams. An organization may experience an unusual incident that is similar to ones handled by other teams; sharing information can facilitate more effective and efficient incident handling for all teams involved. An alternative to joining a formal group is to participate in incident-related mailing lists, anonymously providing nonsensitive information on an incident and asking for opinions.¹⁸
- + **Affected External Parties.** An incident may affect external parties directly; for example, an outside organization may contact the agency and claim that one of the agency's users is attacking it. This topic is discussed further in Section 7. Another way in which external parties may be affected is if an attacker gains access to sensitive information regarding them, such as credit card information. In some jurisdictions, organizations are required to notify all parties that are affected by such an incident. Regardless of the circumstances, it is preferable for the organization to notify affected external parties of an incident before the media or other external organizations do so. Handlers

¹³ Formerly known as the National Infrastructure Protection Center (NIPC), the IAIP Web site is located at <http://www.nipc.gov>.

¹⁴ <http://www.nipc.gov/incident/cirr.htm>

¹⁵ <http://www.cert.org>

¹⁶ <https://irf.cc.cert.org>

¹⁷ <http://www.first.org>

¹⁸ Examples of some popular mailing lists are provided in Appendix F.

should be careful to give out only the appropriate information—the affected parties may ask for details about internal investigations that should not be revealed publicly.

It is highly recommended that the incident response team discuss with its public affairs and legal representatives the circumstances under which each type of external organization can be contacted and the kind of information that can be provided. These procedures should be written, and all incident response team members should follow them.

2.4 Incident Response Team Structure

An incident response team is available for contact by anyone who discovers or suspects that an incident involving the organization has occurred. One or more team members, depending on the magnitude of the incident and availability of personnel, will then handle the incident. The incident handlers analyze the incident data, determine the impact of the incident, and act appropriately to limit the damage to the organization and restore normal services. Although the incident response team may have only a few members, the team's success is dependent on the participation and cooperation of individuals throughout the organization. This section identifies such individuals, discusses incident response team models, and provides guidance for selecting an appropriate model.

2.4.1 Team Models

Incident response team structure models fall into one of three categories:

- + **Central Incident Response Team.** A single incident response team handles incidents throughout the organization. This model is effective for small organizations and for large organizations with minimal geographic diversity in terms of computing resources.
- + **Distributed Incident Response Teams.** The organization has multiple incident response teams, each responsible for handling incidents for a particular logical or physical segment of the organization. This model is effective for large organizations (e.g., one team per division) and for organizations with major computing resources at distant locations (e.g., one team per geographic region, one team per major facility). However, the teams should be part of a single centralized entity so that the incident response process is consistent across the organization and information is shared among teams. This is particularly important because multiple teams may see components of the same incident or may handle similar incidents. Strong communication among teams and consistent practices should make incident handling more effective and efficient.
- + **Coordinating Team.** An incident response team provides guidance and advice to other teams without having authority over those teams—for example, a department-wide team may assist individual agencies' teams. This model can be thought of as a CSIRT for CSIRTs. Because the focus of this document is standard CSIRTs, the coordinating team model is not addressed in detail in this document.¹⁹

Incident response teams can also use any of three staffing models:

- + **Employees.** The organization performs all of its incident response work, with limited technical and administrative support from contractors.

¹⁹ Information on the Coordinating team model, as well as extensive guidance on other team models, will be available in an upcoming CERT[®]/CC document, tentatively titled *Organizational Models for Computer Security Incident Response Teams (CSIRTs)*.

- + **Partially Outsourced.** The organization outsources portions of its incident response work. Section 2.4.2 discusses the major factors that should be considered with outsourcing. Although incident response duties can be divided among the organization and one or more outsourcers in many ways, a few arrangements have become commonplace:
 - The most prevalent arrangement is for the organization to outsource 24x7 monitoring of intrusion detection sensors, firewalls, and other security devices to an offsite managed security services provider (MSSP). The MSSP identifies and analyzes suspicious activity and reports each detected incident to the organization's incident response team. Because the MSSP's employees can monitor multiple customers at once, this model may provide a 24x7 monitoring and response capability at a skill and cost level that may be superior to a comparable internal team.
 - Some organizations perform basic incident response work in-house and call on contractors to assist with handling incidents, particularly those that are more serious or widespread. The services most often performed by the contractors include computer forensics, advanced incident analysis, incident containment and eradication, and vulnerability mitigation.
- + **Fully Outsourced.** The organization completely outsources its incident response work, typically to an onsite contractor. This model is most likely to be used when the organization needs a full-time, onsite incident response team but does not have enough available, qualified employees.

2.4.2 Team Model Selection

When selecting appropriate structure and staffing models for an incident response team, organizations should consider the following factors:

- + **The Need for 24x7 Availability.** Larger organizations, as well as smaller ones that support critical infrastructures, usually need incident response staff to be available 24 hours a day, 7 days a week. This typically means that incident handlers can be contacted at any time by phone or pager, but it can also mean that an onsite presence is required at all times. Real-time availability is the best for incident response because the longer an incident lasts, the more potential there is for damage and loss. Real-time contact is often needed when working with other agencies and organizations—for example, tracing spoofed traffic back to its source through router hops. An incident response team that can react quickly to investigate, contain, and mitigate incidents should be genuinely useful to the organization.
- + **Full-Time Versus Part-Time Team Members.** Organizations with limited funding, staffing, or incident response needs may have only part-time incident response team members. In this case, the incident response team can be thought of as a volunteer fire department. When an emergency occurs, the team members are contacted rapidly, and those who can assist do so. An existing group such as the IT help desk can act as a first point of contact for incident reporting. The help desk members can be trained to perform the initial investigation and data gathering and then alert the incident response team if it appears that a serious incident has occurred. Organizations with part-time team members should ensure that they maintain their incident response skills and knowledge.
- + **Employee Morale.** Incident response work is very stressful, as are the on-call responsibilities of most team members. This combination makes it easy for incident response team members to become overly stressed. Many organizations will also struggle to find willing, available, experienced, and properly skilled people to participate, particularly in 24-hour support.
- + **Cost.** Cost is a major factor, especially if employees are required to be onsite 24 hours a day. Organizations may fail to include incident response-specific costs in budgets. For example, most

organizations do not allocate sufficient funding for training and maintaining skills. Because the incident response team works with so many facets of IT, its members need much broader knowledge than most IT staff members. They must also understand how to use the tools of incident response, such as computer forensics software. The organization should also provide funding for regular team exercises so the team can gain practical experience and improve its performance. Other costs that may be overlooked include physical security for the team's work areas and communications mechanisms.

- + **Staff Expertise.** Incident handling requires specialized knowledge and experience in several technical areas. Outsourcers may possess deeper knowledge of intrusion detection, vulnerabilities, exploits, and other aspects of security than employees of the organization. Also, managed security service providers may be able to correlate events among customers so that they can identify new threats more quickly than any individual customer could. Section 2.4.3 contains additional information on recommended team member skills.
- + **Organizational Structures.** If an organization has three departments that function independently, incident response may be more effective if each department has its own incident response team. The main organization can host a centralized incident response entity that facilitates standard practices and communications among the teams.

When considering outsourcing, organizations should keep these issues in mind:²⁰

- + **Current and Future Quality of Work.** The quality of the outsourcer's work remains a very important consideration. Organizations should consider not only the current quality of work, but also the outsourcer's efforts to ensure the quality of future work—for example, minimizing turnover and burnout and providing a solid training program for new employees.
- + **Division of Responsibilities.** Organizations are usually unwilling to give an outsourcer authority to make operational decisions for the environment (e.g., disconnecting a Web server). It is important to decide the point at which the outsourcer hands off the incident response to the organization. One partially outsourced model addresses this issue by having the outsourcer provide incident data to the organization's internal team, along with recommendations for further handling the incident. The internal team ultimately makes the operational decisions.
- + **Sensitive Information Revealed to the Contractor.** Dividing incident response responsibilities and restricting access to sensitive information can limit this. For example, a contractor may determine what user ID was used in an incident but not know what person is associated with the user ID. The contractor can report to the organization that user ID 123456 is apparently being used to download pirated software without knowing who 123456 is. Trusted employees within the organization can then take over the investigation.
- + **Lack of Organization-Specific Knowledge.** Accurate analysis and prioritization of incidents are dependent on specific knowledge of the organization's environment. The organization should provide the outsourcer regularly updated documents that define what incidents they are concerned about, which resources are critical, and what the level of response should be under various sets of circumstances. The organization should also report all changes and updates made to its IT infrastructure, network configuration, and systems. Otherwise, the contractor has to make a best guess as to how each incident should be handled, inevitably leading to mishandled incidents and

²⁰ NIST Special Publication 800-36, *DRAFT Guide to Selecting IT Security Products*, provides guidance on deciding whether to outsource services and selecting an outsourcer. The draft is available at http://csrc.nist.gov/publications/drafts/Products_PC_100802.pdf. Another good resource is the CERT[®]/CC publication *Outsourcing Managed Security Services*, available at <http://www.cert.org/security-improvement/modules/omss/>.

frustration on both sides. Lack of organization-specific knowledge can also be a problem when incident response is not outsourced, if communications are weak among teams or if the organization simply does not collect the necessary information.

- + **Lack of Correlation.** Correlation among multiple data sources is very important. If the intrusion detection system records an attempted attack against a Web server, but the outsourcer has no access to the Web logs, it may not be able to determine whether the attack was successful. To be efficient, the outsourcer will require administrative privileges to critical systems and security device logs remotely over a secure channel. This will increase administration costs, introduce additional access entry points, and increase the risk of unauthorized disclosure of sensitive information.
- + **Handling Incidents at Multiple Locations.** Effective incident response work will sometimes require a physical presence at the organization's facilities. If the outsourcer is offsite, consider where the outsourcer is located, how quickly it can have an incident response team at any facility, and how much this will cost. Consider onsite visits; perhaps there are certain facilities or areas where the outsourcer should not be permitted to work.
- + **Maintaining Incident Response Skills In House.** Organizations that completely outsource incident response should strive to maintain basic incident response skills in house. Situations may arise in which the outsourcer is unavailable (e.g., a new worm attacks thousands of organizations simultaneously, or a natural disaster or national flight stoppage occurs). The organization should be prepared to perform its own incident handling if the outsourcer is unable to. The organization's technical staff must also be able to understand the significance, technical implications, and impact of the outsourcer's recommendations.

2.4.3 Incident Response Personnel

Regardless of which incident response model an organization chooses, a single employee should be in charge of incident response.²¹ In a fully outsourced model, this person is responsible for overseeing and evaluating the outsourcer's work. In all other models, this responsibility is generally achieved by having a team manager and a deputy team manager who assumes authority in the absence of the team manager. The managers typically perform a variety of tasks, including acting as a liaison with upper management and other teams and organizations, defusing crisis situations, and ensuring that the team has the necessary personnel, resources, and skills. Managers should also be technically adept and have excellent communication skills, particularly an ability to communicate to a range of audiences. Finally, team managers should be able to maintain positive working relationships with other groups, even under times of high pressure.

In addition to the team manager and deputy team manager, some teams also have a technical lead—a person with strong technical skills and incident response experience who assumes oversight of and final responsibility for the quality of the technical work that the entire incident response team undertakes. The position of technical lead should not be confused with the position of incident lead. Larger teams often assign an incident lead as the primary POC for handling a specific incident. Depending on the size of the incident response team and the magnitude of the incident, the incident lead may not actually perform any actual incident handling, such as data analysis or evidence acquisition. Instead, the incident lead may be coordinating the handlers' activities, gathering information from the handlers, and providing updates regarding the incident to other groups.

²¹ Of course, there should be at least one other person designated as an alternate, to oversee the incident response capability when the primary person is unavailable.

Members of the incident response team should have excellent technical skills because they are critical to the team's success. Unless the team members command a high level of technical respect across the organization, people will not turn to them for assistance. Technical inaccuracy in functions such as issuing advisories can undermine the team's credibility, and poor technical judgment can cause incidents to worsen. Critical technical skills include system administration, network administration, programming, technical support, and intrusion detection. Every team member should have good problem solving skills; there is no substitute for real-world troubleshooting experience, such as dealing with operational outages. It is not necessary for every team member to be a technical expert—to a large degree, practical and funding considerations will dictate this—but having at least one highly proficient person in each major area of technology (e.g., particular operating systems, Web servers, e-mail servers) is a necessity.

It is important to counteract staff burnout by providing opportunities for learning and growth. Suggestions for building and maintaining skills are as follows:

- + Budget enough funding to maintain, enhance, and expand proficiency in technical areas and security disciplines. Consider sending each full-time team member to at least two technical conferences per year and each part-time team member to at least one.
- + Ensure the availability of books, magazines, and other technical references that promote deeper technical knowledge.
- + Give team members opportunities to perform other tasks, such as creating educational materials, conducting security awareness workshops, writing software tools to assist system administrators in detecting incidents, and conducting research.
- + Consider rotating staff members in and out of the incident response team.
- + Maintain sufficient staffing so that team members can have uninterrupted time off work (e.g., vacations).
- + Create a mentoring program to enable senior technical staff to help less experienced staff learn incident handling.
- + Participate in exchanges in which team members temporarily trade places with others (e.g., network administrators) to gain new technical skills.
- + Occasionally bring in contractors with deep technical knowledge in needed areas, as funding permits.
- + Develop incident handling scenarios and have the team members discuss how they would handle them. Appendix B contains a set of scenarios and a list of questions to be used during scenario discussions.
- + Conduct simulated incident handling exercises for the team. Exercises are particularly important because they not only improve the performance of the incident handlers, but also identify issues with policies and procedures, and with communication.

Incident response team members should have other skills in addition to technical expertise. Teamwork skills are of fundamental importance because cooperation and coordination are necessary for successful incident response. Every team member should also have good communication skills. Speaking skills are particularly important because the team will interact with a wide variety of people, including incident victims, managers, system administrators, human resources, public relations, and law enforcement. Writing skills are important when team members are preparing advisories and procedures. Although not everyone within a team needs to have strong writing and speaking skills, at least a few people within

every team should possess them so the team can represent itself well in front of senior management, users, and the public at large.

2.4.4 Dependencies Within Organizations

It is important to identify other groups within the organization that may be needed to participate in incident handling so that their cooperation can be solicited before it is needed. Every incident response team relies on the expertise and judgment of others, including—

- + **Management.** Management invariably plays an important role in incident response. In the most fundamental sense, management establishes incident response policy, budget, and staffing. Ultimately, management is held responsible for coordinating incident response among various stakeholders, minimizing damage, and reporting to Congress, OMB, the General Accounting Office (GAO), and other parties.
- + **Information Security.** According to FISMA, information security is a responsibility of the CIO. Thus, the reporting and analysis of incidents, usually to the information security office, should be an organizational policy requirement. Also, information security staff members may be needed during incidents—for example, altering network security controls (e.g., firewall rulesets) to contain an incident.
- + **Telecommunications.** Some incidents involve unauthorized access to telephone lines, such as dialing into unsecured modems. Private Branch Exchange (PBX) compromises often are intertwined with break-ins into other systems. The telecommunications staff is aware of current capabilities, and POCs and procedures for working with telecommunications carriers.
- + **IT Support.** IT technical experts, such as system administrators, network administrators, and software developers, not only have the needed technical skills to assist during an incident but also usually have the best understanding of the technology with which they deal on a daily basis. This understanding can facilitate decisions such as whether to disconnect an attacked system from the network.
- + **General Counsel/Legal Department.** Legal experts should review incident response policies and procedures to ensure their compliance with law and Federal guidance, including the right to privacy. In addition, the guidance of the general counsel or legal department should be sought if there is reason to believe that an incident may have legal ramifications, including evidence collection, prosecution of a suspect, or a lawsuit.
- + **Public Affairs and Media Relations.** Depending on the nature and impact of an incident, a need may exist to inform the media and, by extension, the public (within the constraints imposed by security and law enforcement interests). More information on this is provided in Section 2.3.2.
- + **Human Resources.** When an employee is the apparent target of an incident or is suspected of causing an incident, the human resources department often becomes involved—for example, in assisting with disciplinary proceedings or employee counseling.
- + **Business Continuity Planning.** Computer security incidents undermine the business resilience of an organization and act as a barometer of its level of vulnerabilities and the inherent risks. Business continuity planning professionals should be made aware of incidents and their impacts so they can fine-tune business impact assessments, risk assessments, and continuity of operations plans. Also, because business continuity planners have extensive expertise in minimizing operational disruption during severe circumstances, they may be valuable in planning responses to certain types of incidents,

such as a DoS. Organizations should also ensure that incident response policies and procedures do not conflict with business continuity processes.

- + **Physical Security and Facilities Management.** Some computer security incidents occur through breaches of physical security or involve coordinated logical and physical attacks. Threats made against the organization may not indicate whether logical or physical resources are being targeted. Also, the incident response team may need access to facilities during incident handling—for example, to acquire a compromised workstation from a locked office. Thus, close coordination between physical security and facilities management and the incident response team is important.

2.5 Incident Response Team Services

The main focus of an incident response team is performing incident response; however, it is fairly rare for a team to perform incident response only. The following are examples of additional services that an incident response team might offer:²²

- + **Advisory Distribution.** A team may issue advisories that describe new vulnerabilities in operating systems and applications and provide information on mitigating the vulnerabilities.²³ Promptly releasing such information is a high priority because of the direct link between vulnerabilities and incidents. Distributing information about current incidents can also be useful in helping others identify signs of such incidents. It is recommended that only a single team within the organization distribute computer security advisories, to avoid duplication of effort and the spread of conflicting information.
- + **Vulnerability Assessment.** An incident response team can examine networks, systems, and applications for security-related vulnerabilities; determine how they can be exploited and what the risks are; and recommend how the risks can be mitigated.²⁴ These responsibilities can be extended so that the team performs auditing or penetration testing, perhaps visiting sites unannounced to perform on-the-spot assessments. Incident handlers are well suited to performing vulnerability assessments because they routinely see all kinds of incidents and have first-hand knowledge of vulnerabilities and how they are exploited. However, because the availability of incident handlers is unpredictable, organizations should typically give primary responsibility for vulnerability assessments to another team and use incident handlers as a supplemental resource.
- + **Intrusion Detection.** An incident response team may assume responsibility for intrusion detection because others within the organization do not have sufficient time, resources, or expertise. The team generally benefits because it should be poised to analyze incidents more quickly and accurately, based on the knowledge it gains of intrusion detection technologies. Ideally, however, primary responsibility for intrusion detection should be assigned to another team, with members of the incident response team participating in intrusion detection as their availability permits.
- + **Education and Awareness.** Education and awareness are resource multipliers—the more the users and technical staff know about detecting, reporting, and responding to incidents, the less drain there

²² CERT®/CC provides a more detailed list of potential team services at <http://www.cert.org/csirts/services.html>.

²³ Teams should word advisories so that they do not blame any person or organization for security issues. Teams should meet with legal advisors to discuss the possible need for a disclaimer in advisories, stating that the team and organization has no liability in regard to the accuracy of the advisory. This is most pertinent when advisories may be sent to contractors, vendors, and other nonemployees who are users of the organization's computing resources.

²⁴ NIST Special Publication 800-42, *DRAFT Guideline on Network Security Testing*, provides guidance on performing vulnerability assessments and penetration testing. The draft is available at <http://csrc.nist.gov/publications/drafts/security-testing.pdf>.

should be on the incident response team. This information can be communicated through many means: workshops and seminars, Web sites, newsletters, posters, and even stickers on monitors.

- + **Technology Watch.** A team can perform a technology watch function, which means that it looks for new trends in information security threats. Examples of this are monitoring security-related mailing lists, analyzing intrusion detection data to identify an increase in worm activity, researching new rootkits that are publicly available, and monitoring honeypots.²⁵ The team should then make recommendations for improving security controls based on the trends that they identify. A team that performs a technology watch function should also be better prepared to handle new types of incidents.
- + **Patch Management.** Giving the incident response team the responsibility for *patch management* (e.g., acquiring, testing, and distributing patches to the appropriate administrators and users throughout the organization) is generally not recommended. Patch management is a time-intensive, challenging task that cannot be delayed every time an incident needs to be handled. In fact, patch management services are often needed most when attempting to contain, eradicate, and recover from large-scale incidents. Having the same people responsible for incident response and patch management is generally not effective.

2.6 Recommendations

The items below summarize the key recommendations presented in this section for organizing a computer security incident handling capability.

- + **Establish a formal incident response capability.** Organizations should be prepared to respond quickly and effectively when computer security defenses are breached. FISMA requires Federal agencies to establish incident response capabilities.
- + **Create an incident response policy and use it as the basis for incident response procedures.** The incident response policy is the foundation of the incident response program. It defines which events are considered incidents, establishes the organizational structure for incident response, defines roles and responsibilities, and lists the requirements for reporting incidents, among other items.
- + **Establish policies and procedures regarding incident-related information sharing.** The organization will want or be required to communicate incident details with outside parties, such as the media, law enforcement agencies and incident reporting organizations. The incident response team should discuss this requirement at length with the organization's public affairs and legal advisors to establish policies and procedures regarding information sharing. The team should comply with existing organization policy on interacting with the media and other outside parties.
- + **Provide pertinent information on incidents to the appropriate incident reporting organization.** Federal civilian agencies are required to report incidents to FedCIRC. Reporting benefits the agencies because the incident reporting organizations use the reported data to provide information to the agencies regarding new threats and incident trends.
- + **Consider the relevant factors when selecting an appropriate incident response team model.** Organizations should carefully weigh the advantages and disadvantages of each possible team structure model and staffing model in the context of the organization's needs and available resources.
- + **Select people with appropriate skills for the incident response team.** The credibility and proficiency of the team depend to a large extent on the technical skills of its members. Poor technical judgment can undermine the team's credibility and cause incidents to worsen. Critical technical skills

²⁵ Section 3.2.3 discusses honeypots in more depth, including the risks they can pose.

include system administration, network administration, programming, technical support, and intrusion detection. Teamwork and communications skills are also needed for effective incident handling.

- + **Identify other groups within the organization that may need to participate in incident handling.** Every incident response team relies on the expertise and judgment of other teams, including management, information security, IT support, legal, public affairs, and facilities management.
- + **Determine which services the team should offer.** Although the main focus of the team is incident response, most teams perform additional functions. Examples include distributing security advisories, performing vulnerability assessments, educating users on security, and monitoring intrusion detection sensors.

3. HANDLING AN INCIDENT

The incident response process has several phases, from initial preparation through post-incident analysis. During preparation, the organization attempts to limit the number of incidents that will occur by selecting and implementing a set of controls based on the results of risk assessments. However, residual risk will inevitably persist after controls are implemented; furthermore, no control is foolproof. Detection of security breaches is thus necessary to alert the organization whenever incidents occur. In keeping with the severity of the incident, the organization can act to mitigate the impact of the incident by containing it and ultimately recovering from it. After the incident is adequately handled, the organization issues a report that details the cause and cost of the incident and the steps the organization should take to prevent future incidents. The major phases of the incident response process—preparation, detection and analysis, containment/eradication/recovery and post-incident activity—are described in detail throughout this section. Figure 3-1 presents a diagram of the incident response life cycle.




Figure 3-1. Incident Response Life Cycle

3.1 Preparation

Incident response methodologies typically emphasize preparation—not only establishing an incident response capability so that the organization is ready to respond to incidents, but also preventing incidents by ensuring that systems, networks, and applications are sufficiently secure. Although the incident response team is not typically responsible for incident prevention, it is so important that it has come to be considered a fundamental component of incident response programs. Also, the incident response team’s expertise should be valuable in establishing recommendations for securing systems. This section provides basic guidance on preparing to handle incidents and on preventing incidents.

3.1.1 Preparing to Handle Incidents

Table 3-1 lists tools and resources available that may be of value during incident handling. Please see Appendix F for information about specific software that may be useful for incident analysis, and for a list of Web sites that contain valuable information regarding incident response. Section 3.2 provides information about detecting incidents through the use of IDSs, centralized logging, and other mechanisms.

Table 3-1. Tools and Resources for Incident Handlers

Acquired	Tool / Resource
Incident Handler Communications and Facilities	
	Contact information for team members and others within and outside the organization (primary and backup contacts), such as law enforcement and other incident response teams; information may include phone numbers, e-mail addresses, public encryption keys (in accordance with the encryption software described below), and instructions for verifying the contact's identity
	On-call information for other teams within the organization, including escalation information (see Section 3.2.6 for more information on escalation)
	Incident reporting mechanisms , such as phone numbers, e-mail addresses, and online forms that users can use to report suspected incidents; at least one mechanism should permit people to report incidents anonymously
	Pagers or cell phones to be carried by team members for off-hours support, onsite communications
	Encryption software to be used for communications among team members, within the organization and with external parties; software must use a Federal Information Processing Standards (FIPS) 140-2 validated encryption algorithm ²⁶
	War room for central communication and coordination; if a permanent war room is not necessary, the team should create a procedure for procuring a temporary war room when needed
	Secure storage facility for securing evidence and other sensitive materials
Incident Analysis Hardware and Software	
	Computer forensic workstations²⁷ and/or backup devices to create disk images, preserve log files, and save other relevant incident data
	Laptops , which provide easily portable workstations for activities such as analyzing data, sniffing packets, and writing reports
	Spare workstations, servers, and networking equipment , which may be used for many purposes, such as restoring backups and trying out malicious code; if the team cannot justify the expense of additional equipment, perhaps equipment in an existing test lab could be used, or a virtual lab could be established using OS emulation software
	Blank media , such as floppy diskettes, CD-Rs, and DVD-Rs
	Easily portable printer to print copies of log files and other evidence from non-networked systems
	Packet sniffers and protocol analyzers to capture and analyze network traffic that may contain evidence of an incident
	Computer forensic software to analyze disk images for evidence of an incident
	Floppies and CDs with trusted versions of programs to be used to gather evidence from systems
Incident Analysis Resources	
	Port lists , including commonly used ports, Trojan horse ports
	Documentation for OSs, applications, protocols, intrusion detection signatures, antivirus signatures
	Network diagrams and lists of critical assets , such as Web, e-mail, and FTP servers
	Baselines of expected network, system and application activity
	Cryptographic hashes of critical files to speed the analysis, verification, and eradication of incidents
Incident Mitigation Software	
	Media , including OS boot disks and CD-ROMs, OS media, and application media

²⁶ FIPS 140-2, *Security Requirements for Cryptographic Modules*, is available at <http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf>.

²⁷ A computer forensic workstation is specially designed to assist incident handlers in acquiring and analyzing data. These workstations typically contain a set of removable hard drives, which can be used for evidence storage.

Acquired	Tool / Resource
	Security patches from OS and application vendors
	Backup images of OS, applications, and data stored on secondary media

Many incident response teams create a *jump kit*, which is a portable bag or case that contains materials that an incident handler may likely need during an offsite investigation. The jump kit is ready to go at all times so that when a serious incident occurs, incident handlers can grab the jump kit and go. Jump kits contain many of the same items listed in Table 3-1. For example, each jump kit typically includes a laptop, loaded with appropriate software (e.g., packet sniffers, computer forensics). Other important materials include backup devices, blank media, basic networking equipment and cables, and operating system and application media and patches. Because the purpose of having a jump kit is to facilitate faster responses, the team should refrain from borrowing items from the jump kit. It is also important to keep the jump kit current at all times (e.g., installing security patches on laptops, updating operating system media). Organizations should balance the cost of creating and maintaining jump kits with the savings from containing incidents more quickly and effectively.

3.1.2 Preventing Incidents

Keeping the number of incidents reasonably low is very important to protect the business processes of the organization. If security controls are insufficient, high volumes of incidents may occur, overwhelming the incident response team. This can lead to slow and incomplete responses, which translate to a larger negative business impact (e.g., more extensive damage, longer periods of service and data unavailability). A sound approach to improving the organization's security posture and preventing incidents is to conduct periodic risk assessments of systems and applications. These assessments should determine what risks are posed by combinations of threats and vulnerabilities.²⁸ Each risk should be prioritized, and the risks can be mitigated, transferred, or accepted until a reasonable overall level of risk is reached. Incorporating or at least examining the control strategies of responsible peer organizations can provide reasonable assurance that what works for others should work for the organization.²⁹

Another benefit of conducting risk assessments regularly is that critical resources are identified, allowing staff to emphasize monitoring and response activities for those resources.³⁰ This should not be interpreted as a justification for organizations to ignore the security of resources that are deemed to be less than critical, because the organization is only as secure as its weakest link. Note that regardless of how effective a risk assessment is, it reflects only the current risk. New threats and vulnerabilities are constantly emerging, and computer security is an ongoing process that requires diligence to be effective.

It is outside the scope of this document to provide specific advice on securing networks, systems, and applications. Incident response teams are generally not responsible for securing resources; however, they can be advocates of good security practices. Also, other documents already provide good advice on general security concepts and operating system and application-specific guidance.³¹ The following text,

²⁸ Guidance on risk management is available in NIST Special Publication 800-30, *Risk Management Guide for Information Technology Systems*, at <http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf>. Another useful resource is the Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) approach, discussed at <http://www.cert.org/octave>.

²⁹ <http://csrc.nist.gov/publications/nistpubs/> provides links to the NIST Special Publications on computer security, which include documents on operating system and application security baselines.

³⁰ Information on identifying critical resources is discussed in Federal Information Processing Standard (FIPS) 199, *Standards for Security Categorization of Federal Information and Information Systems*, at <http://csrc.nist.gov/publications/drafts/FIPS-PUB-199-ipd.pdf>.

³¹ <http://csrc.nist.gov> is a good starting point for documents on network, host, and application security.

however, provides a brief overview of some of the main recommended practices for securing networks, systems, and applications:

- + **Patch Management.** Many information security experts agree that a large percentage of incidents involve exploitation of a relatively small number of vulnerabilities in systems and applications.³² Large organizations should implement a patch management program to assist system administrators in identifying, acquiring, testing, and deploying patches.³³
- + **Host Security.** All hosts should be hardened appropriately. Besides keeping each host properly patched, hosts should be configured to provide only the minimum services to only the appropriate users and hosts—the principle of least privilege. Insecure default settings (e.g., default passwords, unsecured shares) should be changed. Warning banners should be displayed whenever a user attempts to gain access to a secured resource. Hosts should have auditing enabled and should log significant security-related events.
- + **Network Security.** The network perimeter should be configured to deny all activity that is not expressly permitted. Only activity necessary for the proper functioning of the organization should be permitted. This includes securing all connection points, such as modems, virtual private networks (VPN), and dedicated connections to other organizations.
- + **Malicious Code Prevention.** Software to detect and stop malicious code, such as viruses, worms, and Trojan horses, should be deployed throughout the organization. Malicious code protection should be deployed at the host level (e.g., server and workstation operating systems), the application server level (e.g., e-mail server, Web proxies) and the application client level (e.g., e-mail clients, instant messaging clients). Section 5 examines malicious code prevention in more detail.
- + **User Awareness and Education.** Users should be made aware of policies and procedures regarding appropriate use of networks, systems, and applications. This should reduce the frequency of incidents involving malicious code and incidents caused by violations of acceptable use policies. IT staff should be trained so that they can maintain their networks, systems, and applications in accordance with the organization’s security standards.

3.2 Detection and Analysis



Figure 3-2. Incident Response Life Cycle (Detection and Analysis)

3.2.1 Incident Categories

Incidents can occur in countless ways, so it is impractical to develop comprehensive procedures with step-by-step instructions for handling every incident. The best that the organization can do is to prepare

³² The *SANS/FBI Top 20 List – Twenty Most Critical Internet Security Vulnerabilities* identifies some of the most commonly exploited vulnerabilities. It is available from <http://www.sans.org/top20>.

³³ NIST Special Publication 800-40, *Procedures for Handling Security Patches*, is recommended for information on techniques for patching systems. <http://csrc.nist.gov/publications/nistpubs/800-40/sp800-40.pdf>

generally to handle any type of incident and more specifically to handle common incident types. The incident categories listed below are neither comprehensive nor intended to provide definitive classification for incidents; rather, they simply give a basis for providing guidance on how to handle incidents based on their primary category:

- + **Denial of Service**—an attack that prevents the usage of network, system, or application resources
- + **Malicious Code**—a virus, worm, Trojan horse, or other code-based malicious entity that infects a host
- + **Unauthorized Access**—a person gains logical or physical access without permission to a network, system, application, data, or other resource
- + **Inappropriate Usage**—a person violates acceptable computing use policies
- + **Multiple Component**—a single incident that encompasses two or more incidents.

Some incidents fit into more than one category. A team should categorize incidents by the transmission mechanism—for example:

- + A virus that creates a backdoor should be handled as a malicious code incident, not an unauthorized access incident, because the malicious code was the transmission mechanism.
- + A virus that creates a backdoor that has been used to gain unauthorized access should be treated as a multiple component incident, because there are two transmission mechanisms.

This section focuses on recommended practices for handling any type of incident. Sections 4 through 8 give more specific advice based on the incident categories.

3.2.2 Signs of an Incident

For many organizations, the most challenging part of the incident response process is accurately detecting and assessing possible incidents—determining whether an incident has occurred and, if so, the type, extent, and magnitude of the problem. What makes this so challenging is a combination of three factors:

- + Incidents may be detected through many different means, with varying levels of detail and fidelity. Automated detection capabilities include network-based and host-based intrusion detection software (IDS), antivirus software, and log analyzers. Incidents may also be detected through manual means, such as problems reported by users. Some incidents have overt signs that can be easily detected, whereas others are virtually impossible to detect without automation.
- + The volume of potential signs of incidents is typically high; for example, it is not uncommon for an organization to receive thousands or even millions of intrusion detection sensor alerts per day.³⁴
- + Deep, specialized technical knowledge and extensive experience are necessary for proper and efficient analysis of incident-related data. In most organizations, the few people with this level of knowledge are probably assigned to other tasks.

Signs of an incident fall into one of two categories: indications and precursors. A *precursor* is a sign that an incident may occur in the future. An *indication* is a sign that an incident may have occurred or may be

³⁴ For example, a single Web vulnerability scan can generate hundreds of alerts on both a network-based IDS and the Web server's host-based IDS product. An attacker performing such a scan on 10 Web servers could generate several thousand IDS alerts.

occurring now. There are too many types of indications to exhaustively list them, but here are some examples:

- + The network intrusion detection sensor alerts when a buffer overflow attempt occurs against an FTP server.
- + The antivirus software alerts when it detects that a host is infected with a worm.
- + The Web server crashes.
- + Users complain of slow access to the Internet.
- + The system administrator sees a filename with unusual characters.
- + The user calls the help desk to report a threatening e-mail message.
- + The host records an auditing configuration change in its log.
- + The application records multiple failed login attempts from an unfamiliar remote system in its log.
- + The e-mail administrator sees a large number of bounced e-mails with suspicious content.
- + The network administrator notices an unusual deviation from typical network traffic flows.

One should not think of incident detection as being strictly reactive. In some cases, the organization can detect activities that are likely to precede an incident. For example, a network IDS sensor may record unusual port scan activity targeted at a group of hosts, which occurs shortly before a DoS attack is launched against one of the same hosts. The intrusion detection alerts regarding the scanning activity serve as a *precursor* of the subsequent DoS incident. Other examples of precursors are as follows:

- + Web server log entries that show the usage of a Web vulnerability scanner
- + An announcement of a new exploit that targets a vulnerability of the organization's mail server
- + A threat from a hacktivist group stating that the group will attack the organization.

Not every attack can be detected through precursors. Some attacks have no precursors, whereas other attacks generate precursors that the organization fails to detect. If precursors are detected, the organization may have an opportunity to prevent the incident by altering its security posture through automated or manual means to save a target from attack.³⁵ In the most serious cases, the organization may decide to act as if an incident is already occurring, so that the risk is mitigated quickly. At a minimum, the organization can monitor certain activity more closely—perhaps connection attempts to a particular host or a certain type of network traffic.

3.2.3 Sources of Precursors and Indications

Precursors and indications are identified using many different sources, with the most common being computer security software alerts, logs, publicly available information, and people. Table 3-2 lists common sources of precursors and indications for each category.

³⁵ An example of an automated security change is intrusion prevention software, which may detect unusual reconnaissance activity and block subsequent related activity. An example of a manual security change is an administrator creating a new firewall rule to block connection attempts to a particular host.

Table 3-2. Common Sources of Precursors and Indications

Precursor or Indication Source	Description
Computer Security Software Alerts	
Network and host-based IDS	IDS products are designed to identify suspicious events and record pertinent data regarding them, including the date and time the attack was detected, the type of attack, the source and destination Internet Protocol (IP) addresses, and the username (if applicable and known). IDS software often produces <i>false positives</i> —alerts that indicate malicious activity is occurring, when in fact there has been none. Analysts should manually validate IDS alerts, either by closely reviewing the recorded supporting data or by getting related data from other sources.
Antivirus software	When it detects malicious code, antivirus software typically sends alerts to the affected host and a centralized antivirus console. Current antivirus products are very effective at detecting and eradicating or isolating malicious code if their signatures are kept up to date. This updating task can be overwhelming in large organizations. One way of addressing it is to configure centralized antivirus software to push signature updates to individual hosts, rather than rely on hosts to be configured to pull updates. Because detection varies among antivirus products, some organizations use products from multiple vendors to provide better coverage and higher accuracy. Antivirus software should be deployed in at least two levels: at the network perimeter (e.g., firewalls, e-mail servers ³⁶) and at the host level (e.g., workstations, file servers, client software).
File integrity checking software	Incidents may cause changes to important files; file integrity checking software can detect such changes. It works by using a hashing algorithm to obtain a cryptographic checksum for each designated file. If the file is altered and the checksum is recalculated, an extremely high probability exists that the new checksum will not match the old checksum. By recalculating checksums regularly, changes to files can be detected.
Third-party monitoring service	Some organizations pay a third party to monitor their publicly accessible services, such as Web, Domain Name System (DNS) and FTP servers. The third party automatically attempts to access each service every x minutes. If the service cannot be accessed, the third party alerts the organization using the methods specified by the organization, such as phone calls, pages, and e-mails. Some monitoring services can also detect and alert on changes in certain resources—for example, a Web page. Although a monitoring service is mainly useful from an operational standpoint, it can also provide an indication of a denial of service attack or server compromise.
Logs	
Operating system, service and application logs	Logs from operating systems, services, and applications (particularly audit-related data) are frequently of great value when an incident occurs. They can provide a wealth of information, such as which accounts were accessed and what actions were performed. Unfortunately, in many incidents, the logs contain no evidence because logging was not enabled or configured properly on the host. To facilitate effective incident handling, organizations should require a baseline level of logging on all systems, and a higher baseline level of logging on critical systems. All systems should have auditing turned on and should log audit events, particularly administrative-level activity. All systems should be checked periodically to verify that logging is functioning properly and adheres to the logging standards.
Network device logs	Logs from network devices such as firewalls and routers are not typically used as a primary source of precursors or indications. Although these devices are usually configured to log blocked connection attempts, they provide little information about the nature of the activity. Still, they can be valuable in identifying trends (e.g., a significantly increased number of attempts to access a particular port) and in correlating events detected by other devices.

³⁶ NIST Special Publication 800-45, *Guidelines on Electronic Mail Security*, includes guidance on deploying antivirus software to e-mail servers. It is available at <http://csrc.nist.gov/publications/nistpubs/800-45/sp800-45.pdf>.

Precursor or Indication Source	Description
Honeypot logs	Some organizations are sufficiently concerned with detecting precursors of incidents that they have deployed deceptive measures such as honeypots so that they can collect better data on precursors. <i>Honeypots</i> are hosts that have no authorized users other than the honeypot administrators because they serve no business function; all activity directed at them is considered suspicious. Attackers will scan and attack honeypots, giving administrators data on new trends and attack tools, particularly malicious code. However, honeypots are a supplement to, not a replacement for, properly securing networks, systems and applications. If honeypots are to be used by an organization, qualified incident handlers and intrusion detection analysts should manage them. The legality of honeypots has not been clearly established, so organizations should carefully study the legal ramifications before planning any honeypot deployments.
Publicly Available Information	
Information on new vulnerabilities and exploits	Keeping up with new vulnerabilities and exploits can prevent some incidents from occurring and assist in the detection and analysis of new attacks. Several organizations, such as FedCIRC, CERT [®] /CC, IAIP ³⁷ and the Department of Energy's Computer Incident Advisory Capability (CIAC), ³⁸ periodically provide threat update information through briefings, Web postings, and mailing lists.
Information on incidents at other organizations	Reports of incidents that have occurred at other organizations can provide a wealth of information. There are Web sites and mailing lists where incident response teams and security professionals can share information regarding reconnaissance and attacks that they have seen. In addition, some organizations acquire, consolidate, and analyze logs and intrusion detection alerts from many other organizations. ³⁹
People	
People from within the organization	Users, system administrators, network administrators, security staff, and others from within the organization may report signs of incidents. It is important to validate all such reports. Not only do users generally lack the knowledge to determine if an incident is occurring, but also even the best-trained technical experts make mistakes. One approach is to ask people who provide such information how confident they are of the accuracy of the information. Recording this estimate along with the information provided can help considerably during incident analysis, particularly when conflicting data is discovered.
People from other organizations	Although few reports of incidents will originate from people at other organizations, they should be taken very seriously. A classic example is a hacker who identifies a serious vulnerability in a system and either informs the organization directly or publicly announces the issue. Another possibility is that the organization might be contacted by an external party that claims someone at the organization is attacking it. External users may also report other indications, such as a defaced Web page or an unavailable service. Other incident response teams also may report incidents. It is important to have mechanisms in place for external parties to report indications and for trained staff to monitor those mechanisms carefully; this may be as simple as setting up a phone number and e-mail address, configured to forward messages to the help desk.

3.2.4 Incident Analysis

Incident detection and analysis would be easy if every precursor or indication were guaranteed to be accurate; unfortunately, this is not the case. For example, user-provided indications such as a complaint of a server being unavailable are often incorrect. Intrusion detection systems are notorious for producing large numbers of false positives—incorrect indications. These examples demonstrate what makes incident detection and analysis so difficult: each indication should be evaluated to determine if it is legitimate. Making matters worse, the total number of indications from human and automated sources

³⁷ <http://www.nipc.gov>

³⁸ <http://www.ciac.org/ciac>

³⁹ The Internet Storm Center (<http://isc.incidents.org>) is a free source of trend information.

may be thousands or millions a day. Finding the few real security incidents that occurred out of all the indications can be a daunting task.

Even if an indication is accurate, it does not necessarily mean that an incident has occurred. Some indications, such as a Web server crash or modification of critical files, could happen for several reasons other than a security incident, including human error. Given the occurrence of indications, however, it is reasonable to suspect that an incident might be occurring and to act accordingly. In general, incident handlers should assume that an incident is occurring until they have determined that it is not.

Determining whether a particular event is actually an incident is sometimes a matter of judgment. It may be necessary to collaborate with other technical and information security personnel to make a decision. In many instances, an incident should be handled the same way regardless of whether it is security related. For example, if an organization is losing Internet connectivity every 12 hours and no one is certain of the cause, the staff would want to resolve the problem just as quickly and would use the same resources to diagnose the problem, regardless of its cause.

Some incidents are easy to detect, such as an obviously defaced Web page. However, many incidents are not associated with such clear symptoms. Skilled attackers are careful to cover their tracks, and even unskilled attackers are becoming more difficult to detect because the tools that they use are more sophisticated and stealthy. Small signs such as one change in one system configuration file may be the only indications that an incident has occurred. In incident handling, detection may be the most difficult task. Incident handlers are responsible for analyzing ambiguous, contradictory, and incomplete symptoms to determine what has happened. Although technical solutions exist that can make detection somewhat easier, the best remedy is to build a team of highly experienced and proficient staff members who can analyze the precursors and indications effectively and efficiently and take appropriate actions. Without a well-trained and capable staff, incident detection and analysis will be conducted inefficiently; and costly mistakes will be made.

The incident response team should work quickly to analyze and validate each incident, documenting each step taken. When the team believes that an incident has occurred, the team should rapidly perform an initial analysis to determine the incident's scope, such as which networks, systems, or applications are affected; who or what originated the incident; and how the incident is occurring. The initial analysis should provide enough information for the team to prioritize subsequent activities, such as containment of the incident and deeper analysis of the effects of the incident. When in doubt, incident handlers should assume the worst until additional analysis indicates otherwise.

Performing the initial analysis and validation is challenging. The following are recommendations for making incident analysis easier and more effective:

- + **Profile Networks and Systems.** Profiling is one of the best technical measures for aiding in incident analysis. *Profiling* is measuring the characteristics of expected activity so that changes to it can be more easily identified. Examples of profiling are running file integrity checking software on hosts to derive checksums for critical files and monitoring network bandwidth usage and host resource usage to determine what the average and peak usage levels are on various days and times. If the profiling process is automated, changes to activity can be detected and reported to administrators quickly. In practice, it is difficult to detect incidents accurately using most profiling techniques; organizations should use profiling as one of several detection and analysis techniques.
- + **Understand Normal Behaviors.** Incident response team members should study networks, systems, and applications to gain a solid understanding of what their normal behavior is so that abnormal behavior can be recognized more easily. Many intrusion detection analysts have been told at some point to identify unusual occurrences. Without a solid understanding of what "usual" is, "unusual" is

very difficult to define. No incident handler will have a comprehensive knowledge of all behavior throughout the environment, but handlers should know which experts could fill in the gaps.

One way to gain this knowledge is through reviewing log entries and security alerts. This may be tedious if filtering is not used to condense the logs to a reasonable size. As handlers become more familiar with the logs and alerts, they should be able to focus on unexplained entries, which are usually more important to investigate, as well as more interesting. Conducting frequent log reviews should keep the knowledge fresh, and the analyst should be able to notice trends and changes over time. The reviews also give the analyst an indication of the reliability of each source. Reviewing logs and investigating interesting entries are also good preparation for handling incidents, which requires these skills.

- + **Use Centralized Logging and Create a Log Retention Policy.** Information regarding an incident may be recorded in several places, such as firewall, router, network IDS, host IDS, and application logs. Organizations should deploy one or more centralized logging servers and configure logging devices throughout the organization to send duplicates of their log entries to the centralized logging servers. Incident handlers benefit by having all pertinent log entries available together, and this consolidation also provides secure storage for logs, which reduces the impact of attackers disabling logging or modifying logs on individual hosts that they compromise. In addition, creating and implementing a log retention policy that specifies how long log data should be maintained may be extremely helpful in analysis, as older log entries may show reconnaissance activity or previous instances of similar attacks. The length of time to maintain log data is dependent on several factors, including the organization's data retention policies and the volume of data. Generally, log data should be retained for at least a few weeks, preferably for at least a few months.
- + **Perform Event Correlation.** Evidence of an incident may be captured in several logs. Each log may contain different types of data regarding the incident—a firewall log may have the source IP address that was used, whereas an application log may contain a username. A network intrusion detection sensor may detect that an attack was launched against a particular host, but it may not know if the attack was successful. The analyst may need to examine the host's logs in order to determine that. Correlating events among multiple indication sources can be invaluable in validating whether a particular incident occurred, as well as rapidly consolidating the pieces of data. Using centralized logging makes event correlation easier and faster, as it pulls together data from networks, hosts, services, applications and security devices.
- + **Keep All Host Clocks Synchronized.** Protocols such as the Network Time Protocol (NTP) synchronize clocks among hosts. This is important for incident response because event correlation will be more difficult if the devices reporting events have inconsistent clock settings. From an evidentiary standpoint, it is highly preferable to have consistent timestamps in logs—for example, to have three logs that show an attack occurred at 12:07:01 a.m., rather than logs that list the attack as occurring at 12:07:01, 12:10:35, and 11:07:06.
- + **Maintain and Use a Knowledge Base of Information.** The knowledge base should include information that handlers need to reference quickly during incident analysis. Although it is possible to build a knowledge base with a complex structure, a simple approach can be effective. Text documents, spreadsheets, and relatively simple databases provide effective and flexible mechanisms for sharing data among team members. Appendix F provides pointers to documents that may be of use during protocol analysis, such as commonly used port numbers. The knowledge base should also contain other information, including the following:
 - Links to malicious code and hoax information; the most comprehensive and up-to-date sources are typically the major antivirus software vendors

- Links to lists of domains that have been blacklisted for sending spam
- Explanations of the significance and validity of precursors and indications, such as intrusion detection alerts, operating system log entries, and application error codes.
- + **Use Internet Search Engines for Research.** Comprehensive Internet search engines such as Google and AltaVista can help analysts find information on unusual activity, particularly scanning. For example, an analyst may see some unusual scans targeting TCP port 22912. Performing a search on the terms “TCP,” “port,” and “22912” may return some hits that contain logs of similar activity or even an explanation of the significance of the port number. Because most public mailing lists related to incident response or intrusion detection have Web-based archives, Internet search engines will include list archives in their searches. Handlers may want to search private mailing lists and forums that they can access as well as contact other CSIRTs to ask them if they have seen such activity.
- + **Run Packet Sniffers to Collect Additional Data.** Sometimes the indications do not record enough detail to permit the handler to understand what is occurring. If an incident is occurring over a network, the fastest way to collect the necessary data may be to have a packet sniffer capture network traffic. Configuring the sniffer to record traffic that matches specified criteria should keep the volume of data manageable and minimize the inadvertent capture of other information. Sniffers can provide the purest and most complete data about network-based attacks. Some incidents are very difficult to resolve without using a sniffer.
- + **Consider Filtering the Data.** In many organizations, there is simply not enough time to review and analyze all the indications. When presented with large volumes of data, it is human nature to be overwhelmed and, in many cases, simply ignore the data. To promote effective incident detection, it is necessary to overcome that reaction and ensure that at least the most suspicious activity is investigated. One effective strategy is to filter indications so that categories of indications that tend to be insignificant are not shown to the indication analyst. Another strategy is to filter indications so that only categories of indications that are of the highest significance are shown to the analyst. This approach is dangerous, however, because new malicious activity may not fall into one of the chosen indication categories. Nevertheless, this approach is better than not reviewing the indications at all.
- + **Consider Experience as Being Irreplaceable.** Determining the intent of activity is often challenging. Imagine that a handler sees some unusual activity involving a DNS server—not an attack, but some unusual traffic patterns and port numbers. Is this reconnaissance for an impending attack against DNS server—or against another server, using the DNS server as an intermediary? Or could it be benign traffic created by a load balancer? There are several possible explanations for the data, and handlers may lack sufficiently detailed information to conclusively determine which explanation is correct. The best way to determine the intent of suspicious activity is to gain as much incident handling experience as possible. Incident analysis is a demanding technical activity, but also a bit of an art form. An experienced handler can review the data and quickly get an intuitive sense of the significance of the incident.
- + **Create a Diagnosis Matrix for Less Experienced Staff.** Such a matrix may be most helpful for help desk staff, system administrators, and others who perform their own analysis of precursors and indications. It may also be helpful for new intrusion detection analysts and incident response team members. A small excerpt of a matrix is shown in Table 3-3. It lists potential symptoms on the left side and incident categories across the top. The boxes within the matrix indicate which symptoms are typically associated with each incident category and how strongly that symptom is associated with the category. The strength can be listed in any way that will be helpful—from “yes” or “no” to a percentage. The diagnosis matrix provides guidance for less experienced staff members who may see the symptoms but cannot identify the likely underlying cause, and can also be used as a training tool.

The matrix should be even more valuable if it also has supporting text, such as a brief justification of each matrix entry and advice on how to validate each type of incident.

Table 3-3. Excerpt of a Sample Diagnosis Matrix

Symptom	Denial of Service	Malicious Code	Unauthorized Access	Inappropriate Usage
Files, critical, access attempts	Low	Medium	High	Low
Files, inappropriate content	Low	Medium	Low	High
Host crashes	Medium	Low	Medium	Low
Port scans, incoming, unusual	High	Low	Medium	Low
Port scans, outgoing, unusual	Low	High	Medium	Low
Utilization, bandwidth, high	High	Medium	Low	Medium
Utilization, e-mail, high	Low	High	Medium	Medium

- + **Seek Assistance From Others.** Occasionally, the team will be unable to determine the full cause and nature of an incident. If the team lacks sufficient information to contain and eradicate the incident, then it should consult with internal resources (e.g., information security staff) and external resources (e.g., FedCIRC, other CSIRTs, contractors with incident response expertise) for analysis, containment, and eradication assistance. It is important to accurately determine the cause of each incident so that it can be fully contained and the exploited vulnerabilities can be mitigated to prevent similar incidents from occurring.

3.2.5 Incident Documentation

As soon as an incident response team suspects that an incident is occurring or has occurred, it is important to immediately start recording all facts regarding the incident.⁴⁰ A logbook is an effective and simple medium for this but personal digital assistants (PDA), laptops, audio recorders, and digital cameras can also serve this purpose.⁴¹ Documenting system events, telephone conversations, observed changes in files, and so on can lead to a more efficient and systematic, and less error-prone handling of the problem. Every step taken from the time the incident was detected to its final resolution should be documented and timestamped. Every document regarding the incident should be dated and signed by the incident handler. Information of this nature can also serve as evidence in a court of law if legal prosecution is pursued. Whenever possible, handlers should work in teams of at least two, so that one person can record and log events while the other person performs the technical tasks. More information on evidence is presented in Section 3.3.2.

The incident response team should maintain records about the status of incidents, along with other pertinent information.⁴² Using an application or a database for this purpose is necessary to ensure that incidents are handled and resolved in a timely manner. For example, an incident handler may receive an urgent call pertaining to an incident that was handled the previous day by a handler who has just left on vacation. The handler can quickly become familiar with the incident by accessing the incident database, which should contain information on the following:

⁴⁰ Incident handlers should log only the facts regarding the incident, not personal opinions or conclusions. Subjective material should be presented in incident reports, not recorded as evidence.

⁴¹ Any devices that are potential sources of evidence should not themselves be used to record other evidence.

⁴² The organization's incident response policy should state how long incident information should be retained.

- + The current status of the incident
- + A summary of the incident
- + Actions taken by all incident handlers on this incident
- + Contact information for other involved parties (e.g., system owners)
- + A list of evidence gathered during the incident investigation
- + Comments from incident handlers
- + Next steps to be taken (e.g., waiting for a system administrator to patch an application).⁴³

The incident response team should take care to safeguard data related to incidents because it often contains sensitive information—for example, data on exploited vulnerabilities, recent security breaches, and users that may have performed inappropriate actions. To reduce the risk of sensitive information being released inappropriately, the team should ensure that access to incident data is restricted properly. For example, only authorized personnel should have access to the incident database. E-mails regarding an incident, as well as documents such as incident reports, should be encrypted so that only the sender and intended recipients can read them.

3.2.6 Incident Prioritization

Prioritizing the handling of the incident is perhaps the most critical decision point in the incident handling process. Incidents should not be handled on a first-come, first-served basis as a result of resource limitations. Instead, handling should be prioritized based on two factors:

- + **Current and Potential Technical Effect of the Incident.** Incident handlers should consider not only the current negative technical effect of the incident (e.g., unauthorized user-level access to data), but also the likely future technical effect of the incident if it is not immediately contained (e.g., root compromise). For example, if a worm is spreading among workstations, the current impact may be minor, but within a few hours the worm traffic may cause a major network outage.
- + **Criticality of the Affected Resources.** Resources affected by an incident, such as firewalls, Web servers, Internet connectivity, user workstations, and applications, have different significance to the organization. The criticality of a resource is based primarily on its data or services, users, trust relationships and interdependencies with other resources, and visibility (e.g., a public Web server vs. an internal department Web server). Many organizations have already determined resource criticality through their business continuity planning efforts or their Service Level Agreements (SLA), which state the maximum time for restoring each key resource. When possible, the incident response team should acquire and reuse existing valid data on resource criticality.⁴⁴

Combining the criticality of the affected resources and the current and potential technical effect of the incident determines the business impact of the incident—for example, data destruction on a user

⁴³ The Trans-European Research and Education Networking Association (TERENA) has developed RFC 3067, *TERENA's Incident Object Description and Exchange Format Requirements* (<http://www.ietf.org/rfc/rfc3067.txt>). The document provides recommendations for what information should be collected for each incident. The IETF Extended Incident Handling (inch) Working Group (<http://www.ietf.org/html.charters/inch-charter.html>) has created two Internet-Draft documents that expand on TERENA's work—*Requirements for Format for INcident Report Exchange (FINE)* (<http://www.ietf.org/internet-drafts/draft-ietf-inch-requirements-01.txt>) and *The Incident Data Exchange Format Data Model and XML Implementation* (<http://www.ietf.org/internet-drafts/draft-ietf-inch-iodef-01.txt>)

⁴⁴ A fundamental concept of business continuity planning is Business Impact Analysis (BIA), which refers to determining the impact of particular events. BIA information for an organization may be directly applicable to incident prioritization.

workstation could result in a minor loss of productivity, whereas root compromise of a public Web server could result in a major loss of revenue, productivity, access to services, reputation, and the release of confidential data (e.g., credit card numbers, Social Security Numbers). The team should prioritize the response to each incident based on its estimate of the business impact caused by the incident. For example, inappropriate usage incidents that are not security-related typically do not need to be handled as quickly as other types of incidents because their business impact is relatively low. (Section 7 provides guidance on prioritizing such incidents.)

Organizations should document prioritization guidelines in a format such as the sample matrix shown in Table 3-4. The column headings list the criticality tiers, and the row headings list the impact categories. Each value within the matrix specifies the maximum time that the incident response team has to begin their response to the incident. This can be thought of as an SLA for incident response. The SLA is based on initial response time, not the conclusion of the incident, because the length of time needed to handle an incident varies widely. Organizations should customize the matrix based on their own needs and their approach to identifying resource criticality. For example, an organization may have several criticality classifications. Also, minor incidents, such as a virus infection that does no damage, may best be handled by the local IT staff instead of the incident response team. It may also be desirable to establish two versions of the matrix—one for incidents that occur during the standard workday, and another for incidents that occur off-hours.

Table 3-4. Sample Incident Response SLA Matrix

Current impact or likely future impact of the incident	Resources currently impacted or likely to be impacted by the incident		
	Tier 1 (e.g., Internet connectivity, public Web servers, firewalls, customer data)	Tier 2 (e.g., system administrator workstations, file and print servers, XYZ application data)	Tier 3 (e.g., user workstations)
Root-level access	15 minutes	30 minutes	1 hour
Unauthorized data modification	15 minutes	30 minutes	2 hours
Unauthorized access to sensitive data	15 minutes	1 hour	1 hour
Unauthorized user-level access	30 minutes	2 hours	4 hours
Services unavailable	30 minutes	2 hours	4 hours
Annoyance	30 minutes	Local IT staff	Local IT staff

More than one matrix entry may apply to an incident if it affects multiple resources (e.g., systems, applications, data). The incident handler can identify all applicable matrix entries and follow the most urgent action first. For example, if malicious code has provided unauthorized user-level data access of a tier 1 resource (30-minute response) and system compromise of a tier 3 resource (1-hour response), the handlers should address the issue on the tier 1 resource first and then on the tier 3 resource. Of course, the handlers may want to look at the tier 3 resource sooner than the designated 1-hour maximum, particularly if the team thinks that it may have information that would be helpful in handling the incident with the tier 1 resource.

The matrix approach encourages organizations to consider carefully how the incident response team should react under various circumstances. By providing a framework for making incident handling decisions, the matrices save incident handlers time. During an incident, handlers are often under great

stress, and decision-making can be challenging. Incident handlers should also have discretion to deviate from the matrix based on their judgment, especially when unforeseen or unusual circumstances occur.

Organizations should also establish an escalation process for those instances when the team does not respond to an incident within the designated time. This can happen for many reasons: for example, cell phones and pagers may fail, people may have personal emergencies, or an incident handler may fall back to sleep after answering a call in the middle of the night. The escalation process should state how long a person should wait for a response and what the person should do if no response occurs. Generally, the first step is to duplicate the initial contact, such as calling the same cell phone numbers. After waiting for a brief time—perhaps 15 minutes—the caller should escalate the incident to a higher level, such as the incident response team manager. If that person does not respond within a certain time, then the incident should be escalated again to a higher level of management. This process should be repeated until someone responds.

3.2.7 Incident Notification

When an incident is analyzed and prioritized, the incident response team needs to notify the appropriate individuals within the organization and, in some cases, other organizations.⁴⁵ Timely reporting and notification enable all those who need to be involved to play their roles. Given the magnitude and complexity of today's information security threats, cooperative incident response is likely the most effective approach. Incident response policies should include provisions concerning incident reporting—at a minimum, what must be reported to whom and at what times (e.g., initial notification, regular status updates). The exact reporting requirements vary among agencies, but parties that are typically notified include—

- + CIO
- + Head of information security
- + Local information security officer
- + Other incident response teams within the organization
- + System owner
- + Human resources (for cases involving employees, such as harassment through e-mail)
- + Public affairs (for incidents that may generate publicity)
- + Legal department (for incidents with potential legal ramifications).

During the handling of an incident, the team may need to notify certain parties frequently of the current status of the incident. In some cases, such as a major malicious code infection, the team may need to send organization-wide updates. The team should plan and prepare several communication methods, and select the methods that are appropriate for a particular incident. For example, if the e-mail server has been overwhelmed by malicious code, the team should not send incident updates by e-mail. Possible communication methods include—

- + E-mail
- + Web site (Intranet-based)

⁴⁵ Section 2.3.2 provides more information on communicating with external parties.

- + Telephone calls
- + In person (e.g., daily briefings)
- + Voice mailbox greeting (e.g., set up a separate voice mailbox for incident updates, and update the greeting message to reflect the current incident status)
- + Paper (e.g., post notices on bulletin boards and doors, hand out notices at all entrance points).

3.3 Containment, Eradication, and Recovery



Figure 3-3. Incident Response Life Cycle (Containment, Eradication, and Recovery)

3.3.1 Choosing a Containment Strategy

When an incident has been detected and analyzed, it is important to contain it before the spread of the incident overwhelms resources or the damage increases. Most incidents require containment, so it is important to consider it early in the course of handling each incident. An essential part of containment is decision-making (e.g., shut down a system, disconnect it from a wired or wireless network, disconnect its modem cable, disable certain functions). Such decisions are much easier to make if strategies and procedures for containing the incident have been predetermined. Organizations should define acceptable risks in dealing with incidents and develop strategies accordingly.

Containment strategies vary based on the type of incident. For example, the strategy for containing a malicious code infection is quite different from that of a distributed denial of service attack. Sections 4 through 8 of this document provide specific guidance on containing several types of incidents. It is highly recommended that organizations create separate containment strategies for each major type of incident. The criteria should be documented clearly to facilitate quick and effective decision-making. Criteria for determining the appropriate strategy include—

- + Potential damage to and theft of resources
- + Need for evidence preservation
- + Service availability (e.g., network connectivity, services provided to external parties)
- + Time and resources needed to implement the strategy
- + Effectiveness of the strategy (e.g., partially contains the incident, fully contains the incident)
- + Duration of the solution (e.g., emergency workaround to be removed in 4 hours, temporary workaround to be removed in 2 weeks, permanent solution).

In certain cases, some organizations delay the containment of an incident so that they can monitor the attacker's activity, usually to gather additional evidence. The incident response team should discuss

delayed containment with its legal department to determine if it is feasible. If an organization knows that a system has been compromised and allows the compromise to continue, it may be liable if the attacker uses the compromised system to attack other systems. The delayed containment strategy is dangerous because an attacker could escalate unauthorized access or compromise other systems in a fraction of a second. This strategy should be attempted by only a highly experienced incident response team that can monitor all of the attacker's actions and disconnect the attacker in a matter of seconds. Even then, the value of delayed containment is usually not worth the high risk that it poses.

Another potential issue regarding containment is that some attacks may cause additional damage when they are contained. For example, a compromised host may run a process that pings another host periodically. When the incident handler attempts to contain the incident by disconnecting the compromised host from the network, the subsequent pings will fail. Because of the failure, the process may overwrite all the data on the host's hard drive. Handlers should not assume that just because a host has been disconnected from the network, further damage to the host has been prevented.

3.3.2 Evidence Gathering and Handling

It is important to clearly document how all evidence, including compromised systems, has been preserved. All evidence should be collected according to procedures that meet all applicable laws and regulations, developed from previous discussions with legal staff and appropriate law enforcement agencies.⁴⁶ Also, evidence should be accounted for at all times; whenever evidence is transferred from person to person, chain of custody forms should detail the transfer and include each party's signature. A detailed log should be kept for all evidence, including the following:

- + Identifying information (e.g., the location, serial number, model number, hostname, and IP address of a computer)
- + Name, title, and phone number of each individual who collected or handled the evidence during the investigation
- + Time and date (including time zone) of each occurrence of evidence handling
- + Locations where the evidence was stored.

Collecting evidence from computing resources presents some challenges. It is generally desirable to acquire evidence from a system of interest as soon as one suspects that an incident may have occurred. Many incidents cause a dynamic chain of events to occur; an initial system snapshot may do more good in identifying the problem and its source than most other actions that can be taken at this stage. Also, from an evidentiary standpoint, it is much better to get a snapshot of the system as-is rather than doing so after incident handlers, system administrators, and others have inadvertently altered the state of the machine during the investigation. Users and system administrators should be made aware of the steps that they should take to preserve evidence.

Before copying the files from the affected host, it is often desirable to capture information that may not be recorded in a file system or image backup, such as current network connections, processes, login sessions, open files, network interface configurations, and the contents of memory. This data may hold clues as to the attacker's identity or the attack methods that were used. It is also valuable to document how far the

⁴⁶ *Searching and Seizing Computers and Obtaining Electronic Evidence in Criminal Investigations*, from the Computer Crime and Intellectual Property Section (CCIPS) of the Department of Justice (DOJ), provides legal guidance on evidence gathering. The document is available at <http://www.cybercrime.gov/s&smanual2002.htm>. Another helpful document is *Best Practices for Seizing Electronic Evidence*, available from the U. S. Secret Service at http://www.secretservice.gov/electronic_evidence.shtml.

local clock deviates from the actual time. However, risks are associated with acquiring information from the live system. Any action performed on the host will alter the state of the machine to some extent. Also, the attacker may currently be on the system and notice the handler's activity, which could have disastrous consequences.

A well-trained and careful incident handler should be able to issue only the minimum commands needed to acquire the dynamic evidence without inadvertently altering other evidence. A single poorly chosen command can irrevocably destroy evidence; for example, simply displaying the contents of a directory can alter the last access time on each listed file. Furthermore, running commands from the affected host is dangerous because they may have been altered or replaced (e.g., Trojan horses, rootkits) to conceal information or cause additional damage. Incident handlers should use a write-protected floppy or a CD that contains trusted commands and all dependent files so that all necessary commands can be run without using the affected host's commands. Incident handlers can also use *write blocker* programs that prevent the host from writing to its hard drives.

After acquiring volatile data, an incident handler with computer forensics training should immediately make a full disk image to sanitized write-protectable or write-once media. A disk image preserves all data on the disk, including deleted files and file fragments. If it is possible that evidence may be needed for prosecution or internal disciplinary actions, the handlers should make at least two full images, label them properly, and securely store one of the images to be used strictly as evidence. (All evidence, not just disk images, should be tagged and stored in a secure location.) In some cases, handlers may acquire and secure the original disk as evidence; the second image can then be restored to another disk as part of system recovery.

Obtaining a disk image is superior to a standard file system backup for computer forensic purposes because it records more data. Imaging is also preferable because it is much safer to analyze an image than it is to perform analysis on the original resource—the analysis may inadvertently alter or damage the original. If the business impact of taking down the system outweighs the risk of keeping the system operational, disk imaging may not be possible. A standard file system backup can capture information on existing files, which may be sufficient for handling many incidents, particularly those that are not expected to lead to prosecution. Both disk imaging and file system backups are valuable regardless of whether the attacker will be prosecuted, because they permit the target to be restored while the investigation continues using the image or backup.

Computer forensic software is valuable not only for acquiring disk images, but also for automating much of the analysis process, such as:

- + Identifying and recovering file fragments and hidden and deleted files and directories from any location (e.g., used space, free space, slack space)
- + Examining file structures, headers, and other characteristics to determine what type of data each file contains, instead of relying on file extensions (e.g., .doc, .jpg, .mp3)
- + Displaying the contents of all graphics files
- + Performing complex searches
- + Graphically displaying the acquired drive's directory structure
- + Generating reports.

Organizations should use only programs that have been validated as acceptable for computer forensics purposes. The validation that the software functions properly and accurately may have been given by a court, or occurred through independent testing and verification. The use of other software may introduce serious errors into the forensics process, such as inadvertently altering the data that is being copied. Using software that affects the integrity of the data not only complicates analysis, but also reduces the likelihood that the data can be used as evidence.

During evidence acquisition, it is often prudent to acquire copies of supporting log files from other resources—for example, firewall logs that show what IP address an attacker used. As with hard drive and other media acquisition, logs should be copied to sanitized write-protectable or write-once media. One copy of the logs should be stored as evidence, whereas a second copy could be restored to another system for further analysis. Many incident handlers create a *message digest* for log files and other pieces of digital evidence; this refers to generating a cryptographic checksum for a file, so that if the file is modified and the checksum is recalculated, there is only an infinitesimal chance that the checksums will be the same. Message digests should be generated using software and a message digest algorithm that are FIPS 140-2 and FIPS 180-2 validated.⁴⁷ (Message digests are also useful for other computer forensic purposes—for example, when acquiring media, handlers can generate checksums of the original media and the duplicates to show that integrity was maintained during imaging.) Incident handlers should also document the local clock time on each logging host and what deviation, if any, there is from the actual time.

To assist in incident analysis, handlers may want to duplicate an aspect of an incident that was not adequately recorded. For example, a user visited a malicious Web site, which then compromised the workstation. The workstation contains no record of the attack. A handler may be able to determine what happened by setting up another workstation and contacting the same Web site while using packet sniffers and host-based security software to record and analyze the activity. Handlers should be very careful when duplicating such attacks.

Another example where incident duplication may occur is when an internal user is suspected of downloading inappropriate files. If the firewall has recorded which FTP servers the user visited, an incident handler may decide to access the same FTP servers to see what types of materials they contain and determine whether the filenames on the user's workstation correspond to filenames on the FTP servers. Handlers should only consider accessing external services if they are available to the public (e.g., FTP server that permits anonymous logons). Although it may be acceptable to monitor network traffic to determine what FTP account and password a user provided, it is generally not acceptable to reuse that information to gain access to the FTP server.

3.3.3 Identifying the Attacker

During incident handling, system owners and others typically want to know who committed the attack. Although this can be important, particularly if the organization wants to prosecute the attacker, incident handlers should generally stay focused on containment, eradication, and recovery. Identifying the attacker can be a time-consuming and futile process that can keep a team from achieving their primary goal—minimizing the business impact. The following items describe the most commonly performed activities for attacker identification:

⁴⁷ FIPS 180-2 is titled the *Secure Hash Standard*. The Secure Hash Algorithm (SHA-1) is an example of a popular message digest algorithm that is FIPS 140-2 and FIPS 180-2 validated. See <http://csrc.nist.gov/cryptval/140-1/1401val.htm> for information on FIPS 140-2 validation, and <http://csrc.nist.gov/publications/fips/fips180-2/fips180-2.pdf> for the FIPS 180-2 standard.

- + **Validating the Attacker's IP Address.** New incident handlers often focus on the attacker's IP address. The handler may attempt to validate that the address was not spoofed by using pings, traceroutes, or other methods of verifying connectivity. However, this is not helpful, because at best it indicates that a host at that address responds to the requests. A failure to respond does not mean the address is not real—for example, a host may be configured to ignore pings and traceroutes. The attacker may have received a dynamic address (e.g., from a dial-up modem pool) that has already been assigned to someone else. More importantly, if the IP address is real and the team pings it, the attacker may be tipped off that the organization has detected their activity. If this occurs before the incident has been fully contained, the attacker could cause additional damage, such as wiping out hard drives with evidence of the attack. The team should consider acquiring and using IP addresses from another organization (e.g., an ISP) when performing actions such as address validation, so that the true origin of the activity is concealed from the attacker.
- + **Scanning the Attacker's System.** Some incident handlers do more than perform pings and traceroutes to check an attacking IP address—they may run port scanners, vulnerability scanners and other tools to attempt to gather more information on the attacker. For example, the scans may indicate that Trojan horses are listening on the system, which implies that the attacking host itself has been compromised. Incident handlers should discuss this with legal representatives before performing such scans, because they may violate organization policies or even break the law.
- + **Researching the Attacker Through Search Engines.** In most attacks, incident handlers will have at least a few pieces of data regarding the possible identity of the attacker, such as a source IP address, an e-mail address, or an IRC nickname. Performing an Internet search using this data may lead to more information on the attacker—for example, a mailing list message regarding a similar attack, or even the attacker's Web site. Research such as this generally does not need to be performed before the incident has been fully contained.
- + **Using Incident Databases.** Several groups collect and consolidate intrusion detection and firewall log data from various organizations into incident databases.⁴⁸ Some of these databases allow people to search for records corresponding to a particular IP address. Incident handlers could use the databases to see if other organizations are reporting suspicious activity from the same source. Of course, the organization can also check its own incident tracking system or database for related activity.

3.3.4 Eradication and Recovery

After an incident has been contained, eradication may be necessary to eliminate components of the incident, such as deleting malicious code and disabling breached user accounts.⁴⁹ For some incidents, eradication is either not necessary or taken care of during recovery. In recovery, administrators restore systems to normal operation and (if applicable) harden the security of the systems to prevent similar incidents. Recovery may involve such actions as restoring systems from clean backups, rebuilding systems from scratch, replacing compromised files with clean versions, installing patches, changing passwords, and tightening network perimeter security (e.g., firewall rulesets, boundary router access control lists). It is also often desirable to employ higher levels of system logging or network monitoring as part of the recovery process. Once a resource is successfully attacked, it is often attacked again; or other resources within the organization are attacked in a similar manner. Many valuable resources are available on the Internet for recovering and securing systems.⁵⁰ Because eradication and recovery actions

⁴⁸ The Distributed Intrusion Detection System (DSHield), located at <http://www.dshield.org>, is a free, publicly available incident database.

⁴⁹ Sections 4 through 8 present information on eradication for each incident category.

⁵⁰ <http://csrc.nist.gov/publications/nistpubs/> provides links to the NIST Special Publications on computer security. CERT®/CC also provides useful documents on securing systems and recovering from incidents at http://www.cert.org/tech_tips/.

are typically operating system or application-specific, detailed recommendations and guidance regarding them are outside the scope of this document.

3.4 Post-Incident Activity



Figure 3-4. Incident Response Life Cycle (Post-Incident Activity)

3.4.1 Lessons Learned

One of the most important parts of incident response is also the most often omitted: learning and improving. Each incident response team should evolve to reflect new threats, improved technology, and lessons learned. Many organizations have found that holding a “lessons learned” meeting with all involved parties after a major incident, and periodically after lesser incidents, is extremely helpful in improving security measures and the incident handling process itself.⁵¹ This meeting provides a chance to achieve closure with respect to an incident by reviewing what occurred, what was done to intervene, and how well intervention worked. The meeting should be held within several days of the end of the incident. Questions to be answered in the lessons learned meeting include—

- + Exactly what happened, and at what times?
- + How well did staff and management perform in dealing with the incident? Were the documented procedures followed? Were they adequate?
- + What information was needed sooner?
- + Were any steps or actions taken that might have inhibited the recovery?
- + What would the staff and management do differently the next time a similar incident occurs?
- + What corrective actions can prevent similar incidents in the future?
- + What additional tools or resources are needed to detect, analyze, and mitigate future incidents?

Small incidents generally need limited post-incident analysis, with the exception of incidents performed through new attack methods that are of widespread concern and interest. After serious attacks have occurred, it is usually worthwhile to hold post-mortem meetings that cross team and organizational boundaries, to provide a mechanism for information sharing. The primary consideration in holding such meetings is ensuring that the right people are involved. Not only is it important to invite people who have been involved in the incident that is being analyzed, but also it is wise to consider who should be invited for the purpose of facilitating future cooperation.

The success of such meetings also depends on the agenda. Collecting input about expectations and needs (including suggested topics to cover) from participants before the meeting increases the likelihood that the

⁵¹ Multiple incidents can be covered in a single lessons learned meeting.

needs of the participants will be met. In addition, establishing rules of order before or during the start of a meeting can minimize confusion and discord. Having one or more moderators who are skilled in group facilitation can yield a high payoff. Finally, it is also important to document the major points of agreement and action items, and communicate them to parties who could not attend the meeting.

Lessons learned meetings provide other benefits. Reports from these meetings serve as good material for training new team members by showing them how more experienced team members respond to incidents. Updating incident response policies and procedures is another important part of the lessons learned process. Post-mortem analysis of the way an incident was handled will often reveal a missing step or an inaccuracy in a procedure, providing impetus for change. Because of the changing nature of information technology and changes in personnel, the incident response team should review all related documentation and procedures for handling incidents at designated intervals.

Another important post-incident activity is creating a follow-up report for each incident, which can be quite valuable for future use. First, it provides a reference that can be used to assist in handling similar incidents. Also, creating a formal chronology of events (including timestamped information such as log data from systems) is important for legal reasons, as is creating a monetary estimate of the amount of damage the incident caused in terms of any loss of software and files, hardware damage, and staffing costs (including restoring services). This estimate may become the basis for subsequent prosecution activity by entities such as the FBI or the U.S. Attorney General's office.

3.4.2 Using Collected Incident Data

Lessons learned activities should produce a set of objective and subjective data regarding each incident. Over time, the collected incident data should be useful in several capacities. The data, particularly the total hours of involvement and the cost, may be used to provide justification for additional funding of the incident response team. A study of incident characteristics may indicate systemic security weaknesses and threats, as well as changes in incident trends. This data can be put back into the risk assessment process, ultimately leading to the selection and implementation of additional controls. Another good use of the data is measuring the success of the incident response team. If incident data is collected and stored properly, it should provide several measures of the success (or at least the activities) of the incident response team. Also, organizations (e.g., Federal agencies) that are required to report incident information will need to collect the necessary data to meet their requirements.

Organizations should focus on collecting data that is actionable, rather than collecting data simply because it is available. For example, counting the number of precursor port scans that occur each week and producing a chart at the end of the year that shows port scans increased by 8% is not very helpful and may be quite time-consuming. Absolute numbers are not informative—understanding how they represent threats to the business processes of the organization is what matters. Organizations should decide what incident data to collect based on reporting requirements and on the expected return on investment from the data (e.g., identifying a new threat and mitigating the related vulnerabilities before they can be exploited.) Possible metrics for incident-related data include—

- + **Number of Incidents Handled.** Handling more incidents is not necessarily better—for example, the number of incidents handled may decrease because of better network and host security controls, not because of negligence by the incident response team. The number of incidents handled is best taken as a measure of the relative amount of work that the incident response team had to perform. It is more effective to produce separate incident counts for each incident category (e.g., unauthorized access). Subcategories can be used as well to provide more information. For example, a growing number of unauthorized access incidents performed by insiders could prompt stronger policy provisions concerning background investigations for personnel and misuse of computing resources

and stronger security controls on internal networks (e.g., deploying intrusion detection software to more internal networks and hosts).

- + **Time Per Incident.** For each incident, time can be measured in several ways:
 - Total amount of labor spent working on the incident
 - Elapsed time from the beginning of the incident to its resolution
 - Elapsed time for each stage of the incident handling process (e.g., containment, recovery)
 - How long it took the incident response team to respond to the initial report of the incident.
- + **Objective Assessment of Each Incident.** The response to an incident that has been resolved can be analyzed to determine how effective it was. The following are examples of performing an objective assessment of an incident:
 - Reviewing logs, forms, reports, and other incident documentation for adherence to established incident response policies and procedures
 - Identifying which precursors and indications of the incident were recorded to determine how effectively the incident was logged
 - Determining if the incident caused damage before it was detected
 - Determining if the actual cause of the incident was identified
 - Calculating the estimated monetary damage from the incident
 - Identifying which measures, if any, could have prevented the incident.
- + **Subjective Assessment of Each Incident.** Incident response team members may be asked to assess their own performance, as well as that of other team members and of the team as a whole. Another valuable source of input is the owner of a resource that was attacked—to determine if the owner thinks the incident was handled efficiently and if the outcome was satisfactory.

Besides using these metrics to measure the team's success, organizations may also find it useful to audit their incident response programs periodically. Audits will identify problems and deficiencies that can then be corrected. At a minimum, an incident response audit should evaluate the following items against applicable regulations, policies, and best practices:

- + Incident response policies and procedures
- + Tools and resources
- + Team model and structure
- + Incident handler training and education
- + Incident documentation and reports
- + The measures of success discussed earlier in this section.

3.4.3 Evidence Retention

Organizations should establish policy for how long evidence from an incident should be retained. Most organizations choose to retain all evidence for at least a few months after the incident ends. The following factors should be considered during the policy creation:

- + **Prosecution.** If it is possible that the attacker will be prosecuted, evidence may need to be retained until all legal actions have been completed. In some cases, this may take several years. Also, evidence that does not seem significant now may become more important in the future. For example, if an attacker is able to use knowledge gathered in one attack to perform a more severe attack later, evidence from the first attack may be key to explaining how the second attack was accomplished.
- + **Data Retention.** Most organizations have data retention policies that state how long certain types of data may be kept. For example, an organization may state that e-mail messages should be retained for only 180 days. If a disk image contains thousands of e-mails, the organization may not want the image to be kept for more than 180 days unless it is absolutely necessary.
- + **Cost.** Original hardware (e.g., hard drives, compromised systems) that is stored as evidence, as well as hard drives and other devices that are used to hold disk images, are individually inexpensive for most organizations. However, if an organization stores dozens of such components for years, the cost can be substantial. Also, the organization must retain functional computers that can use the stored hardware (e.g., hard drives) and media (e.g., backup tapes); if the stored evidence can no longer be read, it has lost its value.

3.5 Incident Handling Checklist

The checklist in Table 3-5 provides the major steps to be performed in the initial handling of an incident. The items address only the detection and analysis of an incident; after that has been completed, incident handlers should use checklists that are geared toward a particular type of incident. Sections 4 through 8 contain handling checklists for each of the five incident categories. A generic checklist is provided in Table 3-6 for handling incidents that do not fit into any of the categories.

Note that the actual steps performed may vary based on the type of incident being handled and the nature of individual incidents. For example, if the handler knows exactly what has happened based on analysis of indications (Table 3-5, step 1.1), there may be no need to perform steps 1.2 or 1.3 to further research the activity. The checklists are intended to provide guidance to handlers on the major steps that should be performed, not to dictate the exact sequence of steps that should always be followed.

Table 3-5. Initial Incident Handling Checklist

	Action	Completed
Detection and Analysis		
1.	Determine whether an incident has occurred	
1.1	Analyze the precursors and indications	
1.2	Look for correlating information	
1.3	Perform research (e.g., search engines, knowledge base)	
1.4	As soon as the handler believes an incident has occurred, begin documenting the investigation and gathering evidence	
2.	Classify the incident using the categories presented in Section 3.2.1 (e.g., denial of service, malicious code, unauthorized access, inappropriate usage, multiple component)	
3.	Follow the appropriate incident category checklist; if the incident does not fit into any of the categories, follow the generic checklist	

Table 3-6. Generic Incident Handling Checklist for Uncategorized Incidents

	Action	Completed
Detection and Analysis		
1.	Prioritize handling the incident based on the business impact	
1.1	Identify which resources have been affected and forecast which resources will be affected	
1.2	Estimate the current and potential technical effect of the incident	
1.3	Find the appropriate cell(s) in the prioritization matrix, based on the technical effect and affected resources	
2.	Report the incident to the appropriate internal personnel and external organizations	
Containment, Eradication, and Recovery		
3.	Acquire, preserve, secure, and document evidence	
4.	Contain the incident	
5.	Eradicate the incident	
5.1	Identify and mitigate all vulnerabilities that were exploited	
5.2	Remove malicious code, inappropriate materials, and other components	
6.	Recover from the incident	
6.1	Return affected systems to an operationally ready state	
6.2	Confirm that the affected systems are functioning normally	
6.3	If necessary, implement additional monitoring to look for future related activity	
Post-Incident Activity		
7.	Create a follow-up report	
8.	Hold a lessons learned meeting	

3.6 Recommendations

The items below summarize the key recommendations presented in this section for handling incidents.

- + **Acquire tools and resources that may be of value during incident handling.** The team will be more efficient at handling incidents if various tools and resources are already available to them. Examples include contact lists, encryption software, network diagrams, backup devices, computer forensic software, port lists, and security patches.
- + **Prevent incidents from occurring by ensuring that networks, systems, and applications are sufficiently secure.** Of course, preventing incidents is beneficial to the organization and also reduces the workload of the incident response team. Performing periodic risk assessments and reducing the identified risks to an acceptable level are effective in reducing the number of incidents. User, IT staff, and management awareness of security policies and procedures is also very important.
- + **Identify precursors and indications through alerts generated by several types of computer security software.** Network and host-based intrusion detection systems, antivirus software, and file integrity checking software are valuable for detecting signs of incidents. Each type of software may detect incidents that the other types of software cannot, so the use of several types of computer security software is highly recommended. Third-party monitoring services can also be helpful.
- + **Establish mechanisms for outside parties to report incidents.** Outside parties may want to report incidents to the organization; for example, they may believe that one of the organization's users is attacking them. Organizations should publish a phone number and e-mail address that outside parties can use to report such incidents.
- + **Require a baseline level of logging and auditing on all systems, and a higher baseline level on all critical systems.** Logs from operating systems, services, and applications frequently provide value during incident analysis, particularly if auditing was enabled. The logs can provide information such as which accounts were accessed and what actions were performed.
- + **Profile networks and systems.** Profiling measures the characteristics of expected activity levels so that changes in patterns can be more easily identified. If the profiling process is automated, deviations from expected activity levels can be detected and reported to administrators quickly, leading to faster detection of incidents and operational issues.
- + **Understand the normal behaviors of networks, systems, and applications.** Team members who understand what normal behavior is should be able to recognize abnormal behavior more easily. This knowledge can best be gained by reviewing log entries and security alerts; the handlers should become familiar with the typical data and can investigate the unusual entries to gain more knowledge.
- + **Use centralized logging and create a log retention policy.** Information regarding an incident may be recorded in several places. Organizations should deploy centralized logging servers and configure devices to send duplicates of their log entries to the centralized servers. The team benefits because it can access all log entries at once; also, changes made to logs on individual hosts will not affect the data already sent to the centralized servers. A log retention policy is important because older log entries may show previous instances of similar or related activity.
- + **Perform event correlation.** Indications of an incident may be captured in several logs. Correlating events among multiple sources can be invaluable in collecting all the available information for an incident and validating whether the incident occurred. Centralized logging makes event correlation easier and faster.

- + **Keep all host clocks synchronized.** If the devices reporting events have inconsistent clock settings, event correlation will be more difficult. Clock discrepancies may also cause issues from an evidentiary standpoint.
- + **Maintain and use a knowledge base of information.** Handlers need to reference information quickly during incident analysis; a centralized knowledge base provides a consistent, maintainable source of information. The knowledge base should include general information, such as commonly used port numbers and links to virus information, as well as data on precursors and indications of previous incidents.
- + **Create a diagnosis matrix for less experienced staff.** Help desk staff, system administrators, and new incident response team members may need assistance in determining what type of incident may be occurring. A diagnosis matrix that lists incident categories and the symptoms associated with each category can provide guidance as to what type of incident is occurring and how the incident can be validated.
- + **Start recording all information as soon as the team suspects that an incident has occurred.** Every step taken, from the time the incident was detected to its final resolution, should be documented and timestamped. Information of this nature can serve as evidence in a court of law if legal prosecution is pursued. Recording the steps performed can also lead to a more efficient and systematic, and less error-prone handling of the problem.
- + **Safeguard incident data.** It often contains sensitive information regarding such things as vulnerabilities, security breaches, and users that may have performed inappropriate actions. The team should ensure that access to incident data is restricted properly, both logically and physically.
- + **Prioritize incidents by business impact, based on the criticality of the affected resources and the technical impact of the incident.** Because of resource limitations, incidents should not be handled on a first-come, first-served basis. Instead, organizations should establish written guidelines that outline how quickly the team must respond to the incident and what actions should be performed, based on the incident's current and potential business impact. This guidance saves time for the incident handlers and provides a justification to management and system owners for their actions. Organizations should also establish an escalation process for those instances when the team does not respond to an incident within the designated time.
- + **Include provisions regarding incident reporting in the organization's incident response policy.** Organizations should specify which incidents must be reported, when they must be reported, and to whom. The parties most commonly notified are the CIO, head of information security, local information security officer, other incident response teams within the organization, and system owners.
- + **Establish strategies and procedures for containing incidents.** It is important to contain incidents quickly and effectively to limit their business impact. Organizations should define acceptable risks in containing incidents and develop strategies and procedures accordingly. Containment strategies should vary based on the type of incident.
- + **Follow established procedures for evidence gathering and handling.** The team should clearly document how all evidence has been preserved. Evidence should be accounted for at all times. The team should meet with legal staff and law enforcement agencies to discuss evidence handling, then develop procedures based on those discussions.
- + **Capture volatile data from systems as evidence.** This includes lists of network connections, processes, login sessions, open files, network interface configurations, and the contents of memory.

Running carefully chosen commands from trusted media can collect the necessary information without damaging the system's evidence.

- + **Obtain system snapshots through full forensic disk images, not file system backups.** Disk images should be made to sanitized write-protectable or write-once media. This process is superior to a file system backup for investigatory and evidentiary purposes. Imaging is also valuable in that it is much safer to analyze an image than it is to perform analysis on the original system because the analysis may inadvertently alter the original.
- + **Hold lessons learned meetings after major incidents.** Lessons learned meetings are extremely helpful in improving security measures and the incident handling process itself.

4. HANDLING DENIAL OF SERVICE INCIDENTS

4.1 Incident Definition and Examples

A *denial of service* (DoS) is an action that prevents or impairs the authorized use of networks, systems, or applications by exhausting resources such as CPU, memory, bandwidth, and disk space. Successful DoS attacks conducted against well-known Web sites, making the Web sites unavailable to users, have received extensive media attention since the late 1990's.⁵² Other examples of DoS attacks include—

- + Using all available network bandwidth by generating unusually large volumes of traffic
- + Sending malformed TCP/IP packets to a server so that its operating system will crash
- + Sending illegal requests to an application in order to crash it
- + Establishing many simultaneous login sessions to a server so that other users cannot connect to it
- + Consuming all available disk space by creating many large files.

Four types of DoS attacks merit a more detailed description: reflector attacks, amplifier attacks, distributed DoS attacks, and synfloods.

4.1.1 Reflector Attacks

In a *reflector attack*, a host sends requests with a spoofed source address to a service on an intermediate host.⁵³ (The utilized service is typically UDP-based, so it is easier to carry out a DoS attack by spoofing the address.) That host generates a reply to each request and sends these replies to the spoofed address. Because the intermediate host unwittingly performs the attack, that host is known as a *reflector*. During a reflector attack, a DoS could occur to the host at the spoofed address, the reflector itself, or both hosts. Examples of commonly used reflector services include echo (port 7), chargen (port 19), DNS (port 53), Simple Network Management Protocol (SNMP) (port 161) and Internet Security Association and Key Management Protocol (ISAKMP) (port 500).

In some cases, two reflectors can be used to create a self-contained DoS. An attacker can send requests to a reflector using the spoofed source address of another reflector. So when the first reflector generates its replies, it will send them to the second reflector. If the combination of reflectors is chosen properly, a loop can occur between the two reflectors. The earliest reflector attacks established loops between the echo and chargen services, but current attacks create loops among various combinations of reflector services. Most reflector attacks can be prevented through network-based and host-based firewall rulesets that reject suspicious combinations of source and destination ports.

⁵² Guidance for Web server security is available from NIST Special Publication 800-44, *Guidelines on Securing Public Web Servers*, at <http://csrc.nist.gov/publications/nistpubs/800-44/sp800-44.pdf>.

⁵³ For more information, read *An Analysis of Using Reflectors for Distributed Denial-of-Service Attacks* by Vern Paxson, (<http://www.icir.org/vern/papers/reflectors.CCR.01/reflectors.html>).

In Figure 4-1, the first diagram shows the network traffic pattern of a normal DNS query and response. The DNS client sends a query from its UDP port 1792 to the server's DNS port, 53. The DNS server responds to the query by sending a UDP packet to the client's UDP port 1792. The diagram in the second box shows a reflector attack that uses the same DNS server.⁵⁴ The attacker sends a packet to the DNS server; however, it crafts the packet so it uses a spoofed source address, in this case j.k.l.m. The packet also uses an unusual source port—7. Port 7 is usually associated with echo, a reflector service. When the DNS server receives the packet, it generates and sends a response to the victim at the spoofed address, j.k.l.m. If the victim is offering the echo service, it may create a packet that echoes the received data back to the DNS server. This can cause a loop between the DNS server and the victim if the DNS server responds to the packets sent by the victim.

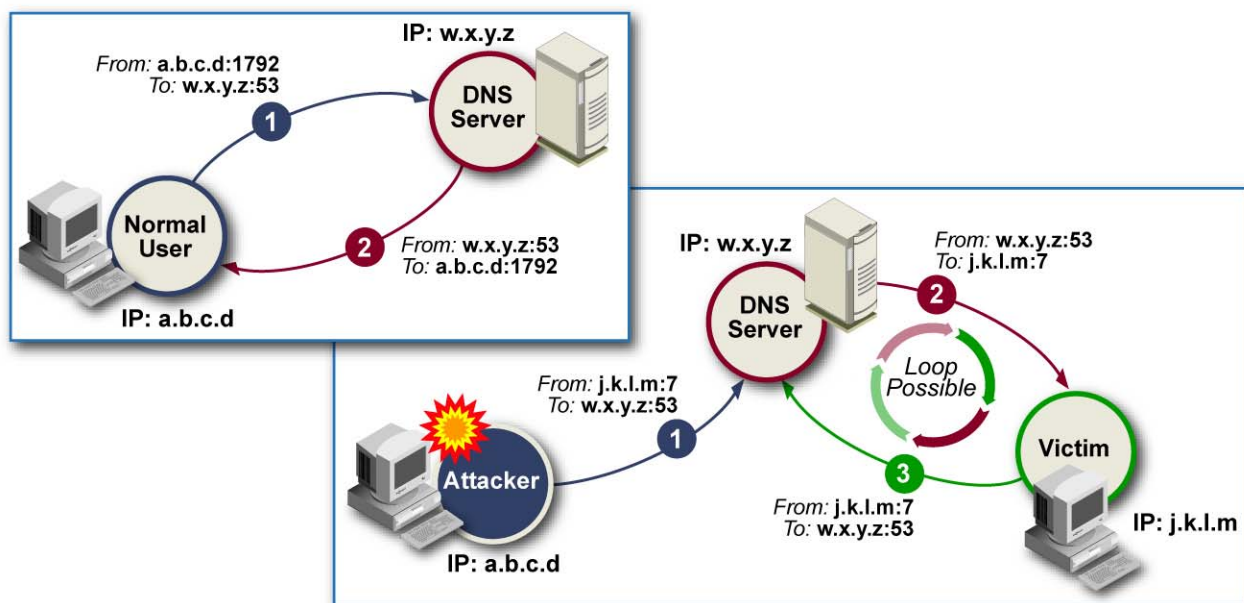


Figure 4-1. Reflector Attack Using a DNS Server

4.1.2 Amplifier Attacks

Like a reflector attack, an amplifier attack involves sending requests with a spoofed source address to an intermediate host. However, an amplifier attack does not use a single intermediate host—instead, its goal is to use a whole network of intermediate hosts. It attempts to do this by sending an ICMP or UDP request to an expected broadcast address, hoping that many hosts will receive the broadcast and respond to it.⁵⁵ Because the attacker's request uses a spoofed source address, the responses are all sent to the spoofed address, which may cause a DoS for that host or the host's network. Most environments block amplifier attacks by configuring border routers not to forward directed broadcasts, but some still permit

⁵⁴ The starburst elements in the graphics in this document mark the malicious actions performed by attackers.

⁵⁵ The request used for an amplifier attack is usually an ICMP message, but UDP can also be used, particularly with the same reflector services mentioned previously. TCP cannot be used; broadcasts are connectionless, and TCP is a connection-oriented protocol, so the concept of a broadcast does not exist for TCP.

them. This has led to the creation of *smurf amplifier lists*, which identify networks that forward directed broadcasts and attempt to educate network administrators on the dangers of this.⁵⁶

Figure 4-2 illustrates a simple ICMP-based amplifier attack. The attacker sends an ICMP ping to a target address. The ping packet has two noteworthy characteristics:

- + The source IP address of the packet is spoofed as j.k.l.m, the IP address of the intended victim.
- + The destination IP address of the packet is w.x.y.255, which is likely to be a valid broadcast address.

The figure shows that the ping penetrated the w.x.y network perimeter and was broadcast to w.x.y hosts. This activity generated numerous ping replies, which simultaneously reach the intended victim at IP address j.k.l.m. Attackers are likely to send crafted pings to many networks at once so that the victim receives a larger number of simultaneous ping replies.

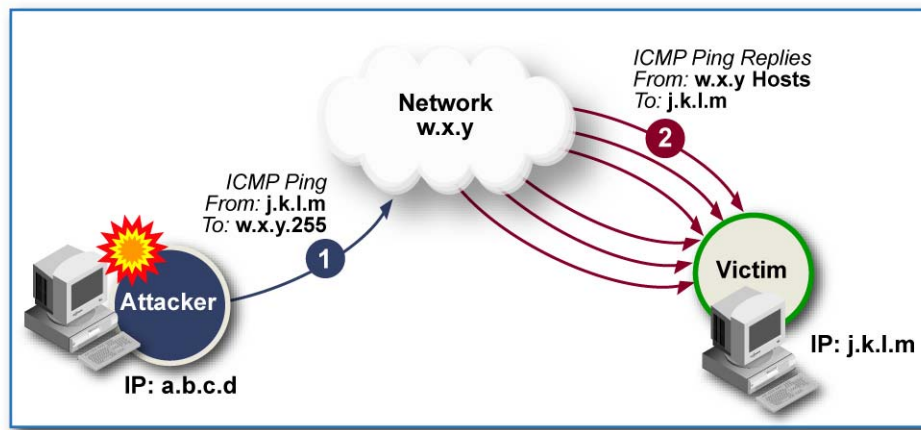


Figure 4-2. Amplifier Attack

4.1.3 Distributed Denial of Service Attacks

Network bandwidth is so large for most organizations that a single attacking machine cannot cause a network DoS. However, attackers now perform *distributed denial of service* (DDoS) attacks, which coordinate an attack among many computers.⁵⁷ If enough hosts are used, the total volume of generated network traffic can exhaust not only the resources of a targeted host but also the available bandwidth for nearly any organization. DDoS attacks have become an increasingly severe threat, and the lack of availability of computing and network services now translates to significant disruption and major financial loss. No organization can completely protect itself from DDoS attacks, but implementing the guidance presented in Section 4.2.2 can reduce the threat of such attacks.

⁵⁶ *Smurf* is a type of amplifier attack that uses ICMP messages. One example of a smurf amplifier list is available at <http://www.netscan.org>.

⁵⁷ For more information on distributed denial of service attacks, see Dave Dittrich's Web page on DDoS attacks and tools (<http://staff.washington.edu/dittrich/misc/ddos>) or read *A Taxonomy of DDoS Attacks and DDoS Defense Mechanisms*, from Jelena Mirkovic, Janice Martin, and Peter Reiher of the University of California, Los Angeles (http://lasr.cs.ucla.edu/ddos/ucla_tech_report_020018.pdf).

DDoS attacks typically use two types of components: *agents*, which run on compromised hosts and perform the actual attacks; and a *handler*, which is a program that controls the agents, telling them when to attack, what to attack, and how to attack it.⁵⁸ In some cases, a handler program is not used; the attacker may communicate with the agents through another means, such as an Internet Relay Chat (IRC) channel; or the agents may be preprogrammed with attack instructions. Attackers often use hundreds or thousands of agents when performing a DDoS attack. Figure 4-3 shows the three steps of a DDoS attack. First, attackers compromise hosts and deploy DDoS agents to them. Second, an attacker uses a handler to instruct the DDoS agents on what to attack, when to attack and how to attack. Finally, the agents follow the instructions and attack the targeted victims.

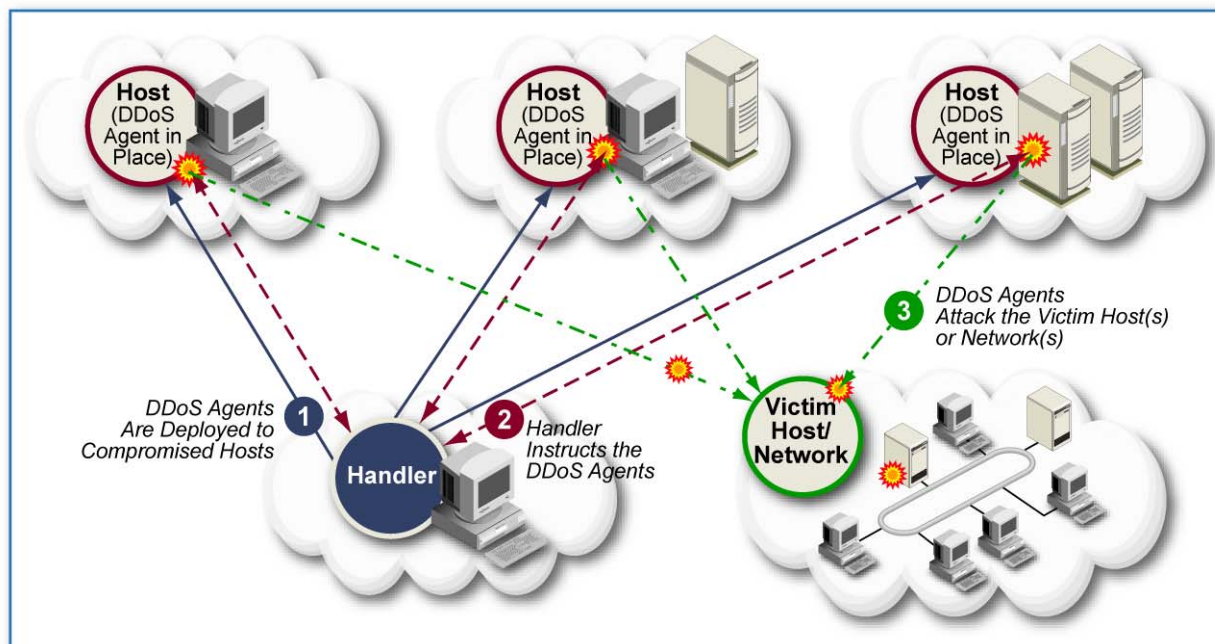


Figure 4-3. Distributed Denial of Service Attack

4.1.4 Synfloods

A synflood occurs when an attacker initiates many TCP connections in a short time (by sending SYN packets) but does not complete the TCP three-way handshakes necessary to fully establish each connection. At one time, many operating systems permitted only a small number of connections to be pending for each service at any moment. If an attacker initiated 100 TCP connections to a particular service port and did not complete any of them, the operating system would have no resources available for another connection to be initiated to that service port until the old connections started timing out, generally after about a minute. This would cause a temporary DoS for the targeted service; if the attacker continued to send SYN packets, the DoS would be extended. Current operating systems are far more resistant to synfloods, and many firewalls now offer protection against synfloods, so they have become less of a threat. Still, synfloods can occur if attackers initiate many thousands of TCP connections in a short time; DDoS tools make this type of attack possible.

⁵⁸ Agents are also known as slaves or daemons, and handlers are also known as masters.

Figure 4-4 shows the difference between normal TCP connections and an attempted synflood. The left side of the diagram shows the three-way TCP handshake, composed of SYN, SYN/ACK and ACK packets. The hosts do not consider the TCP connection to be established fully until the ACK has been received. The right side of the diagram shows a synflood attempt. The attacker has started several three-way TCP handshakes with the server, but the connections are not completed. Because the attacker has spoofed the source address of the connection attempts, the server tries to continue the handshakes by contacting the spoofed host. If the attacker spoofs an IP address that is used by an unavailable host, the server will not receive any responses to the handshakes, effectively causing the DoS.

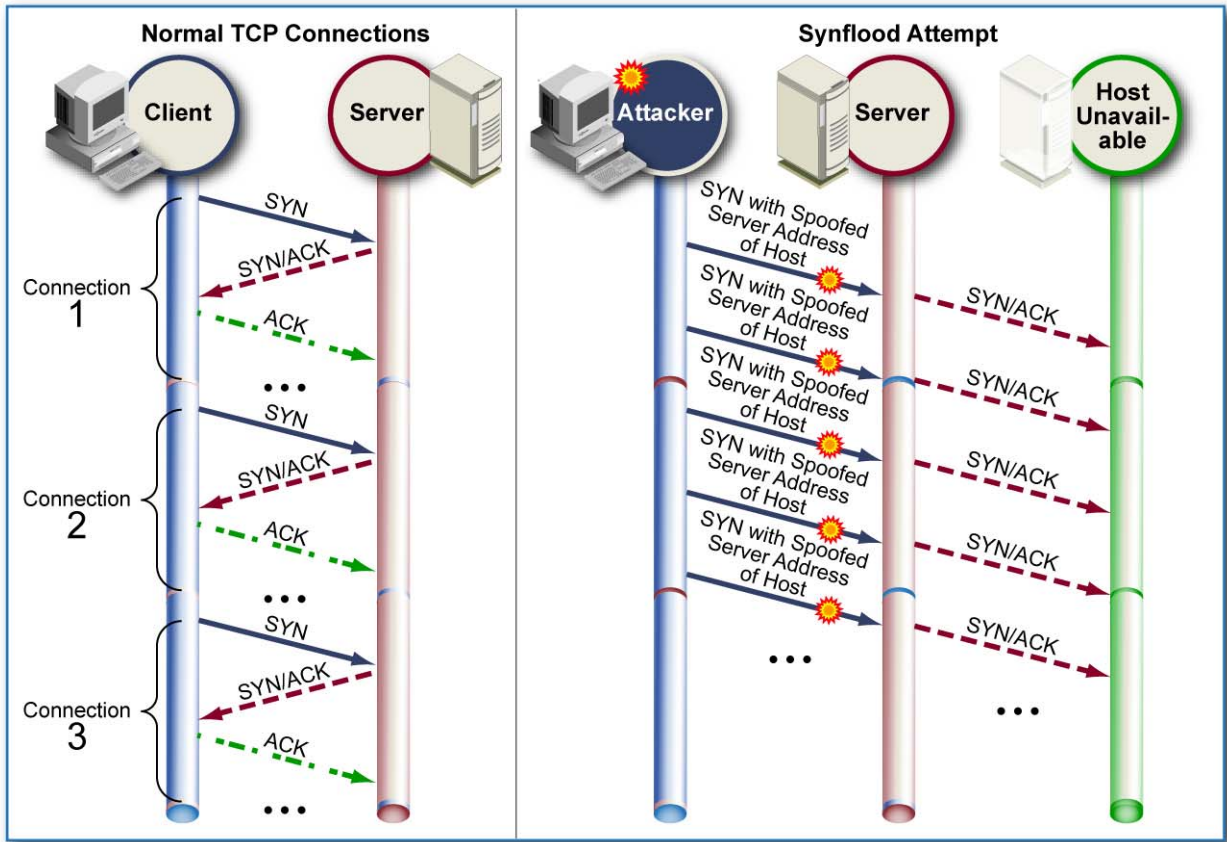


Figure 4-4. Synflood Attack

4.2 Preparation

This section provides guidance on preparing to handle DoS incidents and on preventing DoS incidents.

4.2.1 Incident Handling Preparation

In addition to the general guidance presented in Sections 3.1.1 and 3.2.3, other actions should be performed while preparing to handle DoS incidents:

- + Talk with the organization's ISPs and their second-tier providers to determine how they can assist in handling network-based DoS attacks. This may include filtering or limiting traffic (such as blocking a particular source IP address or setting a maximum limit for incoming ICMP traffic), providing logs of DoS traffic, and tracing attacks back to their source.⁵⁹ Clearly establish what procedures the organization should follow when requesting the assistance of the ISPs or second-tier providers, including 24x7 primary and backup POCs, multiple communication channels, and methods for the ISP to authenticate requests from the organization.
- + Deploy and configure intrusion detection software to detect DoS and DDoS traffic. Network intrusion detection software typically has signatures that identify various types of DoS and DDoS attacks; signatures should be enabled and tuned appropriately to detect attacks with reasonable accuracy. To detect incoming network-based attacks, the organization may need to deploy an additional network IDS sensor outside the external firewall. Organizations should also discuss intrusion detection with their ISPs, because some attacks may overwhelm ISP resources and not even reach the organization's border routers. The ISP may be able to perform traffic monitoring that can detect major DoS attacks occurring over the ISP's networks.
- + Perform ongoing resource monitoring to establish baselines of network bandwidth utilization and critical host resource utilization, and log or alert when there is a significant deviation from the baselines.
- + Identify Web sites that provide statistics on latency between various Internet Service Providers and between various physical locations. This is often referred to as *Internet health monitoring*.⁶⁰ When a network-based DoS occurs, incident handlers could utilize such Web sites to attempt to determine if similar attacks are currently affecting other organizations (e.g., a worm causing regional disruptions).
- + Meet with network infrastructure administrators to discuss how they can assist in analyzing and containing network-based DoS and DDoS attacks. For example, administrators may be able to adjust logging during an attack (e.g., to collect more information on a particular type of activity). Administrators may also be useful in protecting evidence, such as assisting with the acquisition of copies of logs.
- + Maintain local copies (electronic and/or paper) of any computer-based information that may be valuable in handling DoS incidents in case the organization's Internet or internal network connectivity is lost during an incident.

4.2.2 Incident Prevention

Section 3.1.2 has guidelines and pointers to resources on incident prevention. The following items provide additional recommendations for preventing DoS incidents:

⁵⁹ Many ISPs will not provide an attacked organization with logs of the attack unless ordered to do so by a court.

⁶⁰ An example of an Internet health monitoring Web site is the Internet Health Report, located at <http://www.internetpulse.net>.

- + Configure the network perimeter to deny all incoming and outgoing traffic that is not expressly permitted.⁶¹ This should include—
 - Blocking the usage of services, such as echo and chargen, that no longer serve a legitimate purpose and are used in DoS attacks.
 - Performing egress and ingress filtering to block obviously spoofed packets.⁶²
 - Blocking traffic from unassigned IP address ranges, known as bogon lists.⁶³ Attack tools that spoof IP addresses may utilize addresses that have not yet been assigned for Internet usage.
 - Writing and sequencing firewall rules and router access control lists to block traffic properly.⁶⁴
 - Configuring border routers not to forward directed broadcasts.
 - Limiting incoming and outgoing ICMP traffic to only the necessary types and codes.
 - Blocking outgoing connections to common IRC and instant messaging ports if the usage of such services is not permitted.
- + Implement rate limiting for certain protocols, such as ICMP, so that they can only consume a designated percentage of the total bandwidth. Rate limiting can be implemented at the organization’s network perimeter (e.g., border routers, firewalls) and by the organization’s ISPs.
- + On Internet-accessible hosts, disable all unneeded services, and restrict the use of services that may be used in DoS attacks.
- + Implement DoS prevention software. The software can study network traffic patterns, detect major deviations, and block traffic accordingly. Unfortunately, DoS anomaly detection is not completely accurate, so the software may block legitimate traffic or fail to block a true DoS attack.⁶⁵
- + Implement redundancy for key functions (e.g., multiple ISPs, firewalls, Web servers).
- + Ensure that networks and systems are not running near maximum capacity, or it could be easy for a minor DoS attack to take up the remaining resources.

4.3 Detection and Analysis

DoS attacks can be detected through particular precursors and indications, primarily those listed in the following tables. Table 4-1, Denial of Service Precursors, lists possible precursors of a DoS attack, explains the reason why each action might be performed, and provides a recommended response to potentially prevent a related incident from occurring. Table 4-2, Denial of Service Indications, lists

⁶¹ A good source of information on blocking DDoS attacks is *Strategies to Protect Against Distributed Denial of Service (DDoS) Attacks*, available at <http://www.cisco.com/warp/public/707/newsflash.html>.

⁶² *Ingress filtering* is the process of blocking incoming packets from false IP addresses (e.g., packets with reserved source addresses), while *egress filtering* blocks outgoing packets from false IP addresses (e.g., packets with internal network source addresses accidentally leaving the organization and entering the Internet). See RFC 2267, *Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing*, for more information (<http://www.ietf.org/rfc/rfc2267.txt>).

⁶³ <http://www.cymru.com/Documents/bogon-list.html>

⁶⁴ Suppose that the first rule permits queries to reach the public DNS server, and for the DNS server to reply to the queries, using UDP port 53. A second rule may forbid the usage of the echo service, UDP port 7. A clever attacker may use the DNS server as a reflector by sending it packets that use UDP source port 53 and destination port 7. Because the ruleset is evaluated sequentially, the traffic will be permitted because it matches the first rule.

⁶⁵ An example of a commercial DoS prevention product that uses anomaly-based methods is Mazu Networks’ Enforcer.

malicious actions such as a network-based DDoS, a DoS against an operating system, and a DoS against an application, and possible indications of each action. These tables can easily be customized by the organization to include environment-specific precursors and indications, which should facilitate a more efficient and effective incident handling process.

Table 4-1. Denial of Service Precursors

Precursor	Response
DoS attacks are often preceded by reconnaissance activity—generally, a low volume of the traffic that will be used in the actual attack—to determine which attacks may be effective.	If handlers detect unusual activity that appears to be preparation for a DoS attack, the organization may be able to block the attack by quickly altering its security posture—for example, altering firewall rulesets to block a particular protocol from being used or protect a vulnerable host.
A newly released DoS tool could pose a significant threat to the organization.	Investigate the new tool and, if possible, alter security controls so that the tool should not be effective against the organization.

Table 4-2. Denial of Service Indications

Malicious Action	Possible Indications
Network-based DoS against a particular host	<ul style="list-style-type: none"> • User reports of system unavailability • Unexplained connection losses • Network intrusion detection alerts • Host intrusion detection alerts (until the host is overwhelmed) • Increased network bandwidth utilization • Large number of connections to a single host • Asymmetric network traffic pattern (large amount of traffic going to the host, little traffic coming from the host) • Firewall and router log entries • Packets with unusual source addresses
Network-based DoS against a network	<ul style="list-style-type: none"> • User reports of system and network unavailability • Unexplained connection losses • Network intrusion detection alerts • Increased network bandwidth utilization • Asymmetric network traffic pattern (large amount of traffic entering the network, little traffic leaving the network) • Firewall and router log entries • Packets with unusual source addresses • Packets with nonexistent destination addresses
DoS against the operating system of a particular host	<ul style="list-style-type: none"> • User reports of system and application unavailability • Network and host intrusion detection alerts • Operating system log entries • Packets with unusual source addresses
DoS against an application on a particular host	<ul style="list-style-type: none"> • User reports of application unavailability • Network and host intrusion detection alerts • Application log entries • Packets with unusual source addresses

Although these tables may be helpful in analyzing incidents, they are missing an important component—the indications that are associated with benign activities. Benign and malicious events may present similar symptoms, which can make it difficult for analysts to promptly determine if an incident has occurred. Extending the indications table to include benign activities should assist in distinguishing benign from malicious activity. For example, if the organization loses Internet connectivity, many—but

not all—of the symptoms may be similar to a network-based DDoS. This information should be added to the table, with information as to how the benign activity may be distinguished from its malicious counterpart.

DoS attacks pose some additional challenges in terms of incident analysis:

- + DoS attacks often use connectionless protocols (UDP and ICMP) or use a connection-oriented protocol in such a way that full connections are not established (e.g., sending TCP SYN packets to create a synflood attack). Therefore, it is relatively easy for attackers to use spoofed source IP addresses, making it difficult to trace the source of attacks. ISPs may be able to assist in tracing the activity, but it is often more effective to review logs for previous reconnaissance activity that appears to be related. Because the attacker would want to receive the results of the reconnaissance, such activity is unlikely to use a spoofed address, so it may indicate the location of the attacker.
- + DDoS attacks often use hundreds or thousands of workstations that are controlled by a single handler (or no handler at all). The victim site will not see the IP of the handler, and even if it could, it is likely that it is just another host that the attacker has compromised.
- + Network-based DoS attacks are difficult for IDS sensors to detect with a high degree of accuracy. For example, synflood alerts are one of the most common false positives in network IDS products. If an attacker performs a rapid SYN scan, many IDS products will report it as a synflood, even though the activity is sending only one request to each port. If a server crashes, hosts trying to reconnect to it may keep sending SYN packets. Sometimes many legitimate connections in a short time (e.g., retrieving many elements of a Web page) will also cause a synflood alert to be triggered.
- + When an outage occurs, no one may realize that a DoS attack caused it. For example, a Web server may crash occasionally as a result of operating system instability, requiring a reboot for its functionality to be restored. If an attacker sends some specially crafted packets to the Web server that cause it to crash, system administrators may assume the crash resulted from the operating system's instability and not realize that an attack took place.

4.4 Containment, Eradication, and Recovery

In addition to the general guidance presented in Section 3.3, this section gives specific recommendations for performing containment, and gathering and handling evidence for DoS incidents.

4.4.1 Choosing a Containment Strategy

Containment for a DoS incident usually consists of stopping the DoS. Sometimes this is easy; usually it is not. Often the first thought is to block all traffic from the source of the activity. However, as previously mentioned, such attacks often have spoofed source addresses or use hundreds or thousands of compromised hosts—in either case, making it difficult or impossible to implement effective filtering based on source IP addresses. Even if the organization can block the source addresses that are being used, the attacker can simply move to other IP addresses. Other possible solutions for containing a DoS are as follows:

- + **Correct the Vulnerability or Weakness That Is Being Exploited.** For example, if the attack can occur because packet filters do not block packets using UDP port 7 (echo) and a publicly accessible host is accidentally running echo, the filters should be altered to block packets destined for the echo port; and the host's configuration should be changed so that it no longer offers the echo service. If an unpatched operating system is susceptible to a DoS from specially crafted packets, patch the

operating system. The host may need to be temporarily disconnected from the network to halt the DoS while the host is strengthened.

- + **Implement Filtering Based on the Characteristics of the Attack.** For example, if the attack is using ICMP echo requests, one could alter the perimeter security to temporarily block such requests from entering the network. Unfortunately, this is not always practical—if an attacker is sending a SYN flood to a Web server’s HTTP port, blocking SYN packets destined for that port will itself cause a DoS for users. Also, DoS attack tools are increasingly versatile, so if one attack method is blocked, attackers can easily switch to another method. Another strategy is rate filtering—permitting only a certain number of packets per second to use a specific protocol or contact a certain host. Although filtering techniques can be valuable in containing incidents, they can introduce additional problems. For example, adding new rules to a router or firewall may have a substantial negative impact on the device’s performance, causing network slowdowns or even a denial of service. Organizations should carefully consider where filtering should be implemented (e.g., border router, firewall) and should be prepared to upgrade networking devices if necessary to facilitate filtering of long-term attacks.
- + **Have the ISP Implement Filtering.** A network-based DoS from external hosts may overwhelm the organization’s Internet routers. The organization must rely on its ISPs to implement filtering that blocks the activity.
- + **Relocate the Target.** If a particular host is being targeted and other containment strategies are not working, the host could be moved to a different IP address. This is deemed “security through obscurity”, because the attacker may locate the moved target and attack it again. The targeted service could be transferred to a different host—of course, one without the same vulnerability.
- + **Attack the Attackers.** For example, administrators may use programs that are designed to remotely shut off attacking DDoS agents, or they may modify network or server configurations to bounce attack traffic back to its source. However, if the source address is spoofed, or the source address is legitimate but shared (e.g., proxying firewall), these techniques may inadvertently attack an innocent party. Such *hack back* techniques should not be used.

As previously mentioned in Section 3, the decision-making process for containing a DoS incident should be easier if recommended actions are predetermined. This can be done with a matrix or other written guidelines for when each potential solution should be implemented, if ever. The containment strategy may include several solutions in sequence, for example:

1. Implement filtering based on the characteristics of the attack
2. Correct the vulnerability or weakness that is being exploited
3. Have the ISP implement filtering
4. Relocate the target.

4.4.2 Evidence Gathering and Handling

Gathering evidence on DoS attacks is often challenging and time consuming, for any of several reasons:

- + **Identifying the Source of Attacks From Observed Traffic.** IP source addresses are frequently spoofed. DDoS attacks may use hundreds or thousands of hosts, each of which may use multiple spoofed addresses. Even if hosts use their actual addresses, these are the intermediate boxes generating the attack traffic, not the system that orchestrated the overall attack.

- + **Tracing Attacks Back Through ISPs.** Incident handlers may have to contact several ISPs in turn to trace an attack back to its source, assuming that it is possible, technically and logistically. For example, some ISPs may not be willing or able to cooperate with requests without subpoenas, which will take time to obtain. Also, it is much more difficult to trace back an attack that has ended, rather than an ongoing attack. Because of the amount of time that it can take to get assistance from each ISP, it is likely that the attack will end before it can be traced back.
- + **Learning How the Attacking Hosts Were Compromised.** In a DDoS, the agent hosts may have been compromised using dozens of methods over a long period of time. The attacker may not have even been responsible for any of the compromises; instead, the attacker is simply utilizing hosts that others had previously compromised.
- + **Reviewing a Large Number of Log Entries.** Because most DoS attacks work by overwhelming resources, it follows that they may generate unusually large numbers of log entries. Depending on logging standards and practices, log entries may be overwritten, eliminating evidence. Also, it may take a long time to review all of the log entries and extract the pertinent information.

4.5 Checklist for Handling Denial of Service Incidents

The checklist in Table 4-3 provides the major steps to be performed in handling a DoS incident. This checklist is a continuation of the Initial Incident Handling Checklist presented in Table 3-5. Note that the exact sequence of steps may vary based on the nature of individual incidents, and on the strategies chosen by the organization for halting DoS attacks that are in progress.

Table 4-3. Denial of Service Incident Handling Checklist

	Action	Completed
Detection and Analysis		
1.	Prioritize handling the incident based on the business impact	
1.1	Identify which resources have been affected and forecast which resources will be affected	
1.2	Estimate the current and potential technical effect of the incident	
1.3	Find the appropriate cell(s) in the prioritization matrix based on the technical effect and affected resources	
2.	Report the incident to the appropriate internal personnel and external organizations	
Containment, Eradication, and Recovery		
3.	Acquire, preserve, secure, and document evidence	
4.	Contain the incident—halt the denial of service if it has not already stopped	
4.1	Identify and mitigate all vulnerabilities that were used	
4.2	If not yet contained, implement filtering based on the characteristics of the attack, if feasible	
4.3	If not yet contained, contact the ISP for assistance in filtering the attack	
4.4	If not yet contained, relocate the target	
5.	Eradicate the incident; if step 4.1 was not performed, identify and mitigate all vulnerabilities that were used	
6.	Recover from the incident	
6.1	Return affected systems to an operationally ready state	
6.2	Confirm that the affected systems are functioning normally	

	Action	Completed
6.3	If necessary and feasible, implement additional monitoring to look for future related activity	
Post-Incident Activity		
7.	Create a follow-up report	
8.	Hold a lessons learned meeting	

4.6 Recommendations

The items below summarize the key recommendations presented in this section for handling DoS incidents.

- + **Configure firewall rulesets to prevent reflector attacks.** Most reflector attacks can be stopped through network-based and host-based firewall rulesets that reject suspicious combinations of source and destination ports.
- + **Configure border routers to prevent amplifier attacks.** Amplifier attacks can be blocked by configuring border routers not to forward directed broadcasts.
- + **Determine how the organization's ISPs and second-tier providers can assist in handling network-based DoS attacks.** ISPs can often filter or limit certain types of traffic, slowing or halting a DoS attack. They can also provide logs of DoS traffic and may be able to assist in tracing the source of the attack. The organization should meet with the ISPs in advance to establish procedures for requesting such assistance.
- + **Configure security software to detect DoS attacks.** Intrusion detection software can detect many types of DoS activity. Establishing network and system activity baselines, and monitoring for significant deviations from those baselines, can also be useful in detecting attacks.
- + **Configure the network perimeter to deny all incoming and outgoing traffic that is not expressly permitted.** By restricting the types of traffic that can enter and leave the environment, the organization will limit the methods that attackers can use to perform DoS attacks.
- + **Create a containment strategy that includes several solutions in sequence.** The decision-making process for containing DoS incidents is easier if recommended solutions are predetermined. Because the effectiveness of each possible solution will vary among incidents, organizations should select several solutions and determine in which order the solutions should be attempted.

5. HANDLING MALICIOUS CODE INCIDENTS

5.1 Incident Definition and Examples

Malicious code refers to a program that is covertly inserted into another program with the intent to destroy data, run destructive or intrusive programs, or otherwise compromise the security or integrity of the victim's data. Generally, malicious code is designed to perform these nefarious functions without the system's user knowing. Malicious code attacks can be divided into five categories: viruses, Trojan horses, worms, mobile code, and blended.

5.1.1 Viruses

A *virus* is designed to self-replicate—make copies of itself—and distribute the copies to other files, programs, or computers. Viruses insert themselves into host programs and propagate when the infected program is executed, generally by user interaction (e.g., opening a file, running a program, clicking on a file attachment). Viruses have many purposes—some are designed to play annoying tricks, whereas others have destructive intent. Some viruses present themselves as jokes while performing secret destructive functions. Categories of viruses include file infector viruses, boot sector viruses, macro viruses, and virus hoaxes.

- + **File Infector Viruses.** *File infector viruses* attach themselves to executable programs, such as word processors, spreadsheet applications and computer games. When they have infected a program, they propagate to infect other programs on the system and on other systems that use a shared infected program. The virus may also reside in the system's memory, so that each time a new program is executed, the virus infects the program. Another method of file infector execution involves the virus modifying the manner in which the computer opens a file, rather than modifying the actual program running the file. In this scenario, the virus executes first, and then the program is run. Jerusalem and Cascade are two of the best-known file infector viruses.⁶⁶
- + **Boot Sector Viruses.** A *boot sector virus* infects the master boot record (MBR) of a hard drive or the boot sector of removable media, such as floppy diskettes. The boot sector is an area at the beginning of a drive or disk where information about its structure is stored. Boot sectors contain boot programs that are run at host startup to boot the operating system. The MBR of a hard drive is a unique location on the disk where a computer's basic input/output system (BIOS) can locate and load the boot program. Removable media such as floppy disks need not be bootable to infect the system; if an infected disk is in the drive when the computer boots, the virus could be executed. Boot sector viruses are easily concealed, have a high rate of success, and can harm a computer to the point of complete inoperability. Symptoms of a boot sector virus infection include a computer that displays an error message during booting or cannot boot. Form, Michelangelo, and Stoned are examples of boot sector viruses.
- + **Macro Viruses.** Macro viruses are the most prevalent and successful type of virus. *Macro viruses* attach themselves to documents such as word processing files and spreadsheets. As the name implies, a macro virus uses an application's macro programming language to execute and propagate. Many popular software packages, such as Microsoft Office, use macro programming languages in their products to automate complex or repetitive tasks. Attackers have taken advantage of macro programming capabilities to distribute malicious code. Macro viruses tend to spread quickly because users frequently share documents from applications with macro capabilities. Also, when a macro

⁶⁶ More information on viruses is available from most antivirus vendor Web sites, including Sophos (<http://www.sophos.com/virusinfo>), Symantec (<http://securityresponse.symantec.com/avcenter/vinfodb.html>), and Trend Micro (<http://www.trendmicro.com/vinfo/virusencyclo>).

virus infection occurs, the virus also infects the template that the program uses to create and open files. Consequently, every document that is created or opened with the infected template is also infected. The Concept, Marker, and Melissa viruses are well-known examples of macro viruses.

- + **Virus Hoaxes.** As the name implies, *virus hoaxes* are false virus warnings. The phony viruses are usually described as being of devastating magnitude and requiring immediate action to adequately protect computer resources from infection. Despite the illegitimacy of their messages, virus hoaxes are just as prevalent in the digital world as actual viruses. Virus hoaxes are circulated by innocent end users who believe they are helping by distributing these warnings to the Internet community. In general, the hoaxes cause little damage, although some malicious virus hoaxes direct users to alter operating systems settings or delete files, which may cause security or operational problems. Virus hoaxes can be time-consuming, as many of the hoax recipients may contact technical support staff to warn them of the new threat or to ask for guidance. Examples of widely distributed virus hoaxes include Good Times and Bud Frogs.

5.1.2 Trojan Horses

Named after the wooden horse from Greek mythology, *Trojan horses* are programs that appear to be benign but actually have a hidden malicious purpose. Some Trojan horses completely replace an existing file on a system with a malicious version, whereas other Trojan horses are presented as one thing (e.g., a game available for download) but are really another (e.g., a game and a password sniffer). Trojan horses are often difficult to detect because they appear to be performing a useful function. Trojan horses tend to fit into one of three models:

- + Continuing to perform the function of the original program and additionally performing a separate malicious activity
- + Continuing to perform the function of the original program but modifying the function to perform malicious activity (e.g., a Trojan horse version of a login program that collects passwords) or to disguise other malicious activity (e.g., a Trojan horse version of a process listing program that does not display certain processes that are malicious)
- + Performing a malicious function that completely replaces the function of the original program.

The objective of most Trojan horses is to allow a remote user to gain access to, and full control of, the victim computer. To accomplish this, Trojan horses consist of a client component and a server component. The client resides on the intruder's remote computer and will seek to establish a connection with the server, which resides on the infected host. When a connection between client and server is established, the remote intruder can execute commands on the infected computer and transfer or modify files. Another common purpose of Trojan horses is to act as DDoS agents.

Other Trojan horses do not provide remote access to the victim system. Many Trojan horses exist simply to hide evidence of an incident, such as suppressing the names of malicious processes from being displayed in process lists. Other Trojan horses collect information, such as passwords; a Trojan horse could also be configured to mail the password list to a particular external e-mail account on a daily basis.

5.1.3 Worms

Worms are self-replicating programs that are completely self-contained, meaning they do not require a host program to infect a victim. Worms are self-propagating; unlike viruses, they can create fully functional copies and execute themselves without user intervention. Worms take advantage of known vulnerabilities and configuration weaknesses, such as open Windows shares. The nature of worms allows

them to span the globe in a matter of minutes. Examples include the Blaster worm and the SQL Slammer worm.⁶⁷

5.1.4 Mobile Code

Mobile code is software that is transmitted from a remote system to a local system and then executed on the local system without the user's explicit instruction.⁶⁸ Mobile code often acts as a mechanism for a virus, worm, or Trojan horse to be transmitted to the user's workstation. In other cases, mobile code takes advantage of vulnerabilities to perform its own exploits, such as unauthorized data access or root compromise. Popular vehicles for mobile code include Java applets, ActiveX, JavaScript, and VBScript.

5.1.5 Blended

A *blended attack* is an instance of malicious code that uses multiple methods to spread. The well-known Nimda "worm" is actually an example of a blended attack.⁶⁹ It uses four distribution methods:

- + **E-mail.** A user on a vulnerable host opens an infected e-mail attachment; Nimda looks for e-mail addresses on the host and then sends copies of itself to those addresses.
- + **Windows shares.** Nimda scans hosts for unsecured Windows file shares; it can then use NetBIOS⁷⁰ as a transport mechanism to infect files on that host in the hopes that a user will run an infected file, which will activate Nimda on that host.
- + **Web servers.** Nimda scans Web servers, looking for known vulnerabilities in Microsoft IIS. If it finds a vulnerable server, it attempts to transfer a copy of itself to the server and infect it and its files.
- + **Web clients.** If a vulnerable Web client visits a Web server that has been infected by Nimda, the client's workstation will become infected.

In addition to using these methods, blended attacks spread through other services, such as instant messaging and peer-to-peer file sharing. People tend to refer to most instances of blended attacks as worms. Although most people refer to Nimda as a worm, technically it has virus, worm, and mobile code characteristics.

5.2 Preparation

This section provides guidance on preparing to handle malicious code incidents and on preventing malicious code incidents.

5.2.1 Incident Handling Preparation

In addition to the following general guidance presented in Sections 3.1.1 and 3.2.3, other actions should be taken in preparation for handling malicious code incidents:

⁶⁷ More information on the Blaster worm is available at <http://www.cert.org/advisories/CA-2003-20.html>; Slammer worm information is located at <http://www.cert.org/advisories/CA-2003-04.html>.

⁶⁸ NIST Special Publication 800-28: *Guidelines on Active Content and Mobile Code* contains more information on mobile code. It is available at <http://csrc.nist.gov/publications/nistpubs/800-28/sp800-28.pdf>.

⁶⁹ The CERT®/CC advisory on the Nimda worm is available at <http://www.cert.org/advisories/CA-2001-26.html>.

⁷⁰ NetBIOS is a networking service that is used by many Windows hosts.

- + **Make Users Aware of Malicious Code Issues.** This information should include a basic review of the methods that malicious code uses to propagate and the symptoms of infections. Holding regular user education sessions helps to ensure that users are aware of the risks that malicious code poses.
- + **Read Antivirus Vendor Bulletins.** Sign up for mailing lists from antivirus vendors that provide timely information on new malicious code threats.
- + **Deploy Host-Based Intrusion Detection Systems to Critical Hosts.** Host-based IDS software can detect signs of malicious code incidents, such as configuration changes and system executable modifications. File integrity checkers are useful in identifying the affected components of a system.

Some organizations configure their network perimeters to block connections to specific common Trojan horse ports, with the goal of preventing Trojan horse client and server component communications. However, this approach is generally ineffective. Known Trojan horses use hundreds of different port numbers, and many Trojan horses can be configured to use any port number. Also, some Trojan horses use the same port numbers that legitimate services use, so their communications cannot be blocked by port number. Some organizations also implement port blocking incorrectly, so legitimate connections are sometimes blocked.⁷¹ Implementing filtering rules for each Trojan horse port will also increase the demands placed on the filtering device. Generally, a Trojan horse port should be blocked only if the organization has a serious Trojan horse infestation.

5.2.2 Incident Prevention

Section 3.1.2 offered general guidelines and resources useful in preventing incidents from occurring. The following paragraphs provide specific advice on preventing malicious code incidents:

- + **Use Antivirus Software.** Antivirus software is a necessity to combat the threat of malicious code and limit damage. The software should be running on all hosts throughout the organization, and all copies should be kept current with the latest virus signatures so that the newest threats can be thwarted.⁷² Antivirus software should also be used for applications used to transfer malicious code, such as e-mail, file transfer, and instant messaging software. The software should be configured to perform periodic scans of the system as well as real-time scans of each file as it is downloaded, opened, or executed. The antivirus software should also be configured to disinfect and quarantine infected files.⁷³ Some antivirus products not only look for viruses, worms, and Trojan horses, but they also examine HTML, ActiveX, JavaScript, and other types of mobile code for malicious content.
- + **Block Suspicious Files.** Configure e-mail servers and clients to block attachments with file extensions that are associated with malicious code (e.g., .pif, .vbs), and suspicious file extension combinations (e.g., .txt.vbs, .htm.exe).
- + **Limit the Use of Nonessential Programs With File Transfer Capabilities.** Examples include peer-to-peer file and music sharing programs, instant messaging software, and IRC clients and servers. These programs are frequently used to spread malicious code among users.
- + **Educate Users on the Safe Handling of E-mail Attachments.** Antivirus software should be configured to scan each attachment before opening it. Users should not open suspicious attachments

⁷¹ A filtering device may be misconfigured to block all connections with a port of 27374, whether 27374 is the source or destination port. 27374 is used by SubSeven, but it could also legitimately be used as a client port by operating systems.

⁷² Some organizations prefer to configure antivirus software to automatically download and install signature updates, whereas other organizations want to test a new signature update before permitting users to utilize it in case it inadvertently breaks system functionality. Each approach has benefits and downsides; neither is clearly preferable over the other.

⁷³ If a file cannot be disinfected, it is generally preferable from an incident handling and evidentiary standpoint to quarantine it instead of deleting it.

or attachments from unknown sources. Users should also not assume that if the sender is known, the attachment is not infected. Senders may not know that their systems are infected with malicious code that can extract e-mail addresses from files and send copies of the malicious code to those addresses. This activity creates the impression that the e-mails are coming from a trusted person even though the person is not aware that they have been sent. Users can also be educated on file types that they should never open (e.g., *.bat*, *.com*, *.exe*, *.pif*, *.vbs*). Although user awareness of good practices should lessen the number and severity of malicious code incidents, organizations should assume that users will make mistakes and infect systems.

- + **Eliminate Open Windows Shares.** Many worms spread through unsecured shares on hosts running Windows. If one host in the organization is infected with a worm, it could rapidly spread to hundreds or thousands of other hosts within the organization through their unsecured shares. Organizations should routinely check all hosts for open shares and direct the system owners to secure the shares properly. Also, the network perimeter should be configured to prevent traffic that uses NetBIOS ports from entering or leaving the organization’s networks. This should not only prevent external hosts from directly infecting internal hosts through open shares but should also prevent internal worm infections from spreading to other organizations through open shares.
- + **Use Web Browser Security to Limit Mobile Code.** All Web browsers should have their security settings configured so as to prevent unsigned ActiveX and other mobile code vehicles from unknowingly being downloaded to and executed on local systems. Organizations should consider establishing an Internet security policy that specifies which types of mobile code may be used from various sources (e.g., internal servers, external servers).
- + **Configure E-mail Clients to Act More Securely.** E-mail clients throughout the organization should be configured to avoid actions that may inadvertently permit infections to occur. For example, e-mail clients should not automatically execute attachments.

5.3 Detection and Analysis

Because malicious code incidents can take many forms, they may be detected via a number of precursors and indications. Table 5-1, Malicious Code Precursors, lists possible precursors of a malicious code attack, explains the reason why each action might be performed, and provides a recommended response to potentially prevent a subsequent incident. Table 5-2, Malicious Code Indications, lists malicious code actions, such as virus, worm, and Trojan horse infections, and provides possible indications of each action. Blended attacks are not listed because they are detected according to the individual methods that they use, such as virus, worm, and mobile code techniques. The tables in this section can easily be customized by the organization to include environment-specific precursors and indications, which should facilitate a more efficient and effective incident handling process.

Table 5-1. Malicious Code Precursors

Precursor	Response
An alert warns of new malicious code that targets software that the organization uses.	Research the new virus to determine whether it is real or a hoax. This can be done through antivirus vendor Web sites and virus hoax sites. ⁷⁴ If the malicious code is confirmed as authentic, ensure that antivirus software is updated with virus signatures for the new malicious code. If a virus signature is not yet available, and the threat is serious and imminent, the activity might be blocked through other means, such as configuring e-mail servers or clients to block e-mails matching characteristics of the new malicious code. The team might also want to notify

⁷⁴ A good source of hoax information, the CIAC Hoaxbusters, is available at <http://hoaxbusters.ciac.org>.

Precursor	Response
	antivirus vendors of the new virus.
Antivirus software detects and successfully disinfects or quarantines a newly received infected file.	Determine how the malicious code entered the system and what vulnerability or weakness it was attempting to exploit. If the malicious code might pose a significant risk to other users and hosts, mitigate the weaknesses that the malicious code used to reach the system and would have used to infect the target host.

Table 5-2. Malicious Code Indications

Malicious Action	Possible Indications
A virus that spreads through e-mail infects a host.	<ul style="list-style-type: none"> • Antivirus software alerts of infected files • Sudden increase in the number of e-mails being sent and received • Changes to templates for word processing documents, spreadsheets, etc. • Deleted, corrupted, or inaccessible files • Unusual items on the screen, such as odd messages and graphics • Programs start slowly, run slowly, or do not run at all • System instability and crashes • If the virus achieves root-level access, see the indications for “Root compromise of a host” as listed in Table 6-3, Unauthorized Access Indications
A worm that spreads through a vulnerable service infects a host.	<ul style="list-style-type: none"> • Antivirus software alerts of infected files • Port scans and failed connection attempts targeted at the vulnerable service (e.g., open Windows shares, HTTP) • Increased network usage • Programs start slowly, run slowly, or do not run at all • System instability and crashes • If the worm achieves root-level access, see the indications for “Root compromise of a host” as listed in Table 6-3, Unauthorized Access Indications
A Trojan horse is installed and running on a host.	<ul style="list-style-type: none"> • Antivirus software alerts of Trojan horse versions of files • Network intrusion detection alerts of Trojan horse client-server communications • Firewall and router log entries for Trojan horse client-server communications • Network connections between the host and unknown remote systems • Unusual and unexpected ports open • Unknown processes running • High amounts of network traffic generated by the host, particularly if directed at external host(s) • Programs start slowly, run slowly, or do not run at all • System instability and crashes • If the Trojan horse achieves root-level access, see the indications for “Root compromise of a host” as listed in Table 6-3, Unauthorized Access Indications
Malicious mobile code on a Web site is used to infect a host with a virus, worm or Trojan horse.	<ul style="list-style-type: none"> • Indications listed above for the pertinent type of malicious code • Unexpected dialog boxes, requesting permission to do something • Unusual graphics, such as overlapping or overlaid message boxes
Malicious mobile code on a Web site exploits vulnerabilities on a host.	<ul style="list-style-type: none"> • Unexpected dialog boxes, requesting permission to do something • Unusual graphics, such as overlapping or overlaid message boxes • Sudden increase in the number of e-mails being sent and received • Network connections between the host and unknown remote systems • If the mobile code achieves root-level access, see the indications for “Root compromise of a host” as listed in Table 6-3, Unauthorized Access Indications
A user receives a virus hoax message.	<ul style="list-style-type: none"> • Original source of the message is not an authoritative computer security group, but a government agency or an important official person • No links to outside sources • Tone and terminology attempt to invoke panic or a sense of urgency • Urges recipients to forward the message to other people

Prioritizing malicious code incidents properly is important because of their tendency to spread to other systems. In most cases, a basic analysis of the incident should identify which malicious code was used. It is then relatively easy to determine the likely impact of the incident. The incident handler may not yet be aware of all systems that have been infected during the incident; but in most cases, it should be apparent whether the incident involves only a few systems or thousands of servers and workstations across the organization.

5.4 Containment, Eradication, and Recovery

In addition to the general guidance presented in Section 3.3, this section gives specific recommendations for performing containment, and gathering and handling evidence for malicious code incidents.

5.4.1 Choosing a Containment Strategy

Because malicious code works surreptitiously and can propagate to other systems rapidly, early containment of a malicious code incident is needed to stop it from spreading and causing further damage. If the infected system is not critical, disconnecting it from the network immediately is strongly recommended. If the system performs critical functions, it should remain on the network only if the damage to the organization from the services being unavailable is greater than the security risks posed by not immediately disconnecting the system. Other actions that may need to be performed when containing a malicious code incident are as follows:

- + **Any of the Actions Listed in Section 5.2.2.** If one host has been infected, it is highly likely that other systems will be infected, so the containment process includes the prevention of the spread of the malicious code to other systems.
- + **Identifying and Isolating Other Infected Hosts.** Antivirus alert messages are a good source of information, but not every infection will be detected by antivirus software. Incident handlers may need to search for indications of infection through other means, such as—
 - Performing port scans to detect hosts listening on a known Trojan horse or backdoor port
 - Using antivirus scanning and clean-up tools released to combat a specific instance of malicious code
 - Reviewing logs from e-mail servers, firewalls, and other systems that the malicious code may have passed through, as well as individual host logs
 - Configuring network and host intrusion detection software to identify activity associated with infections
 - Auditing the processes running on systems to confirm that they are all legitimate.
- + **Sending Unknown Malicious Code to Antivirus Vendors.** Occasionally, malicious code that cannot be definitively identified by antivirus software may enter the environment. Eradicating the malicious code from systems and preventing additional infections may be difficult or impossible without having updated antivirus signatures from the vendor. Incident handlers should be familiar with the procedures for submitting copies of unknown malicious code to the organization's antivirus vendors.
- + **Configuring E-mail Servers and Clients to Block E-mails.** Many e-mail programs can be configured manually to block e-mails by particular subjects, attachment names, or other criteria that

correspond to the malicious code. This is neither a foolproof nor an efficient solution, but it may be the best option available if an imminent threat exists and antivirus signatures are not yet available.

- + **Blocking Outbound Access.** If the malicious code attempts to generate outbound e-mails or connections, handlers should consider blocking access to IP addresses or services to which the infected system may be attempting to connect.
- + **Shutting Down E-mail Servers.** During the most severe malicious code incidents, with hundreds or thousands of internal hosts infected, e-mail servers may become completely overwhelmed by viruses trying to spread via e-mail. It may be necessary to shut down an e-mail server to halt the spread of e-mail-borne viruses.
- + **Isolating Networks From the Internet.** Networks may become overwhelmed with worm traffic when a severe worm infestation occurs. Occasionally a worm will generate so much traffic throughout the Internet that network perimeters are completely overwhelmed. It may be better to disconnect the organization from the Internet, particularly if the organization's Internet access is essentially useless as a result of the volume of worm traffic. This protects the organization's systems from being attacked by external worms; should the organization's systems already be infected, prevents them from attacking other systems and adding to the traffic congestion.

Identifying infected hosts and vulnerable hosts is made quite complicated by the dynamic nature of computing. If all hosts were powered on and connected to the network at all times, malicious code cleanup would be relatively easy. The actual situation is that hosts may be infected and powered off, moved to other networks, or left on while the system owner is out of the office. Vulnerable hosts that are shut off while their owners are on vacation may quickly become infected when they are powered back on. The identification of vulnerable hosts and infected hosts should not rely solely on user participation. However, organizations often lack the personnel and time to track down each machine manually, particularly when there are substantial numbers of mobile users and telecommuters. Organizations should carefully consider these issues before a large-scale malicious code incident occurs, so that they are prepared to implement effective containment strategies.

5.4.2 Evidence Gathering and Handling

Although it is certainly possible to gather evidence on malicious code incidents, it is often futile, because malicious code is either transmitted automatically or is accidentally transmitted by infected users. It is therefore very difficult and time-consuming to identify the source of malicious code. However, collecting a sample of malicious code for further examination might be useful in some cases.

5.4.3 Eradication and Recovery

Antivirus software effectively identifies and removes malicious code infections; however, some infected files cannot be disinfected. (Files can be deleted and replaced with clean backup copies; in the case of an application, the affected application can be reinstalled.) If the malicious code provided attackers with root-level access, it may not be possible to determine what other actions the attackers may have performed.⁷⁵ In such cases, the system should either be restored from a previous, uninfected backup or be rebuilt from scratch. Of course, the system should then be secured so that it will not be susceptible to another infection from the same malicious code.

⁷⁵ Section 6 provides additional guidance on handling and recovering from root compromises.

5.5 Checklist for Handling Malicious Code Incidents

The checklist in Table 5-3 provides the major steps to be performed in handling a malicious code incident. This checklist is a continuation of the Initial Incident Handling Checklist presented in Table 3-5. Note that the exact sequence of steps may vary based on the nature of individual incidents and on the strategies chosen by the organization for containing incidents.

Table 5-3. Malicious Code Incident Handling Checklist

	Action	Completed
Detection and Analysis		
1.	Prioritize the handling of the incident based on its business impact	
1.1	Identify which resources have been affected and forecast which resources will be affected	
1.2	Estimate the current and potential technical effect of the incident	
1.3	Find the appropriate cell(s) in the prioritization matrix, based on the technical effect and affected resources	
2.	Report the incident to the appropriate internal personnel and external organizations	
Containment, Eradication, and Recovery		
3.	Contain the incident	
3.1	Identify infected systems	
3.2	Disconnect infected systems from the network	
3.3	Mitigate vulnerabilities that were exploited by the malicious code	
3.4	If necessary, block the transmission mechanisms for the malicious code	
4.	Eradicate the incident	
4.1	Disinfect, quarantine, delete, and replace infected files	
4.2	Mitigate the exploited vulnerabilities for other hosts within the organization	
5.	Recover from the incident	
5.1	Confirm that the affected systems are functioning normally	
5.2	If necessary, implement additional monitoring to look for future related activity	
Post-Incident Activity		
6.	Create a follow-up report	
7.	Hold a lessons learned meeting	

5.6 Recommendations

The items below summarize the key recommendations presented in this section for handling malicious code incidents.

- + **Make users aware of malicious code issues.** Users should be familiar with the methods that malicious code uses to propagate and the symptoms of infections. Holding regular user education sessions helps to ensure that users are aware of the risks that malicious code poses. Teaching users how to safely handle e-mail attachments should reduce the number of infections that occur.

- + **Read antivirus bulletins.** Bulletins regarding new malicious code threats provide timely information to incident handlers.
- + **Deploy host-based intrusion detection systems, including file integrity checkers, to critical hosts.** Host-based IDS software, particularly file integrity checkers, can detect signs of malicious code incidents, such as configuration changes and modifications to executables.
- + **Use antivirus software, and keep it updated with the latest virus signatures.** Antivirus software should be deployed to all hosts and all applications that may be used to transfer malicious code. The software should be configured to detect and disinfect or quarantine malicious code infections. All antivirus software should be kept current with the latest virus signatures so the newest threats can be detected.
- + **Configure software to block suspicious files.** Files that are very likely to be malicious should be blocked from the environment, such as those with file extensions that are usually associated with malicious code, as well as files with suspicious combinations of file extensions.
- + **Eliminate open Windows shares.** Many worms spread through unsecured shares on hosts running Windows. A single infection may rapidly spread to hundreds or thousands of hosts through unsecured shares.
- + **Contain malicious code incidents as quickly as possible.** Because malicious code works surreptitiously and can propagate to other systems rapidly, early containment of a malicious code incident is needed to stop it from spreading and causing further damage. Infected systems should be disconnected from the network immediately. Organizations may need to block malicious code at the e-mail server level, or even temporarily suspend e-mail services to gain control over serious e-mail-borne malicious code incidents.

6. HANDLING UNAUTHORIZED ACCESS INCIDENTS

6.1 Incident Definition and Examples

An *unauthorized access* incident occurs when a user gains access to resources that the user was not intended to have. Unauthorized access is typically gained through the exploitation of operating system or application vulnerabilities, the acquisition of usernames and passwords, or social engineering. Attackers may acquire limited access through one vulnerability and use that access to attack more vulnerabilities, eventually gaining higher levels of access. Examples of unauthorized access incidents include—

- + Performing a remote root compromise of an e-mail server
- + Defacing a Web server
- + Guessing and cracking passwords
- + Copying a database containing credit card numbers
- + Running a packet sniffer on a workstation to capture usernames and passwords
- + Using a permission error on an anonymous FTP server to distribute pirated software and music files
- + Dialing into an unsecured modem and gaining internal network access
- + Posing as an executive, calling the help desk, resetting the executive's e-mail password, and learning the new password
- + Using an unattended, logged-in workstation without permission.

6.2 Preparation

This section provides guidance on preparing to handle unauthorized access incidents and on preventing unauthorized access incidents.

6.2.1 Incident Handling Preparation

In addition to following the general guidance presented in Sections 3.1.1 and 3.2.3, other actions should be performed while preparing to handle unauthorized access incidents:

- + Configure network-based and host-based IDS software (such as file integrity checkers and log monitors) to identify and alert on attempts to gain unauthorized access. Each type of intrusion detection software may detect attacks that others are not able to detect.
- + Use centralized log servers so pertinent information from hosts across the organization is stored in a single secured location.
- + Establish procedures to be followed when all users of an application, system, trust domain, or organization should change their passwords because of a password compromise. The procedures should adhere to the organization's password policy.
- + Discuss unauthorized access incidents with system administrators so that they understand their roles in the incident handling process.

6.2.2 Incident Prevention

If the general guidance presented in Section 3.1.2 on incident prevention is applied, the number of unauthorized access incidents should be effectively reduced. In general, employing a strong layered defense strategy, with several security layers between unauthorized users and the resources they are attempting to exploit, is the recommended practice for reducing incidents. Additional steps that support a layered defense strategy are listed in Table 6-1.

Table 6-1. Actions to Prevent Unauthorized Access Incidents

Category	Specific Actions
Network Security	<ul style="list-style-type: none"> • Configure the network perimeter to deny all incoming traffic that is not expressly permitted. • Properly secure all remote access methods, including modems and VPNs. An unsecured modem can provide easily attainable unauthorized access to internal systems and networks. War dialing is the most efficient technique for identifying improperly secured modems.⁷⁶ When securing remote access, carefully consider the trustworthiness of the clients; if they are outside the organization's control, they should be given as little access to resources as possible, and their actions should be closely monitored. • Put all publicly accessible services on secured demilitarized zone (DMZ) network segments. The network perimeter can then be configured so that external hosts can establish connections only to hosts on the DMZ, not internal network segments. • Use private IP addresses for all hosts on internal networks. This will severely restrict the ability of attackers to establish direct connections to internal hosts.
Host Security	<ul style="list-style-type: none"> • Perform regular vulnerability assessments to identify serious risks and mitigate the risks to an acceptable level. • Disable all unneeded services on hosts. Separate critical services so they run on different hosts. If an attacker then compromises a host, immediate access should be gained only to a single service. • Run services with the least privileges possible, to reduce the immediate impact of successful exploits. • Use host-based firewall software to limit individual hosts' exposure to attacks. • Limit unauthorized physical access to logged-in systems by requiring hosts to lock idle screens automatically and asking users to log off before leaving the office. • Regularly verify the permission settings for critical resources, including password files, sensitive databases and public Web pages. This process can easily be automated to report changes in permissions on a regular basis.
Authentication and Authorization	<ul style="list-style-type: none"> • Create a password policy that requires the use of complex, difficult-to-guess passwords, forbids password sharing, and directs users to use different passwords on different systems, especially external hosts and applications. • Require sufficiently strong authentication, particularly for accessing critical resources. • Create authentication and authorization standards for employees and contractors to follow when developing software. For example, passwords should be strongly encrypted using a FIPS 140-2 validated algorithm when they are transmitted or stored. • Establish procedures for provisioning and deprovisioning user accounts. These should include an approval process for new account requests and a process for periodically disabling or deleting accounts that are no longer needed.
Physical Security	<ul style="list-style-type: none"> • Implement physical security measures that restrict access to critical resources.

⁷⁶ War dialing is the process of dialing blocks of phone numbers to identify modems that are listening, then attempting to gain access to the host that the modem is connected to. A brief on war dialing is available from @Stake, Inc. at http://www.atstake.com/research/reports/acrobat/wardialing_brief.pdf.

6.3 Detection and Analysis

Because unauthorized access incidents can occur in many forms, they can be detected through dozens of types of precursors and indications. Table 6-2, Unauthorized Access Precursors, lists possible precursors of an unauthorized access attack, explains the reason why each action might be performed, and provides a recommended response to potentially prevent a subsequent incident from occurring. Table 6-3, Unauthorized Access Indications, lists malicious actions such as root compromise, data modification, and unauthorized account usage. For each such action, the table lists possible indications of the action. These tables can easily be customized by the organization to include environment-specific precursors and indications, which should facilitate a more efficient and effective incident handling process.

Table 6-2. Unauthorized Access Precursors

Precursor	Response
Unauthorized access incidents are often preceded by reconnaissance activity to map hosts and services and to identify vulnerabilities. Activity may include port scans, host scans, vulnerability scans, pings, traceroutes, DNS zone transfers, OS fingerprinting, and banner grabbing. Such activity is detected primarily through IDS software, secondarily through log analysis.	Incident handlers should look for distinct changes in reconnaissance patterns—for example, a sudden interest in a particular port number or host. If this activity points out a vulnerability that could be exploited, the organization may have time to block future attacks by mitigating the vulnerability (e.g., patching a host, disabling an unused service, modifying firewall rules).
A new exploit for gaining unauthorized access is released publicly, and it poses a significant threat to the organization.	The organization should investigate the new exploit and, if possible, alter security controls to minimize the potential impact of the exploit for the organization.
Users report possible <i>social engineering</i> attempts—attackers trying to trick them into revealing sensitive information, such as passwords, or encouraging them to download or run programs and file attachments.	The incident response team should send a bulletin to users with guidance on handling the social engineering attempts. The team should determine what resources the attacker was interested in and look for corresponding log-based precursors, as it is likely that the social engineering is only part of the reconnaissance.
A person or system may observe a failed physical access attempt (e.g., outsider attempting to open a locked wiring closet door, unknown individual using a cancelled ID badge).	If possible, security should detain the person. The purpose of the activity should be determined, and it should be verified that the physical and computer security controls are strong enough to block the apparent threat. (An attacker who cannot gain physical access may perform remote computing-based attacks instead.) Physical and computer security controls should be strengthened if necessary.

Table 6-3. Unauthorized Access Indications

Malicious Action	Possible Indications
Root compromise of a host	<ul style="list-style-type: none"> • Existence of unauthorized security-related tools or exploits • Unusual traffic to and from the host (e.g., attacker may use the host to attack other systems) • System configuration changes, including— <ul style="list-style-type: none"> ○ Process/service modifications or additions ○ Unexpected open ports ○ System status changes (restarts, shutdowns) ○ Changes to log and audit policies and data ○ Network interface card set to promiscuous mode (packet sniffing) ○ New administrative-level user account or group • Modifications of critical files, timestamps and privileges, including executable programs, OS kernels, system libraries, and configuration and data files • Unexplained account usage (e.g., idle account in use, account in use from multiple locations at once, unexpected commands from a particular user, large number of locked-out accounts) • Significant changes in expected resource usage (e.g., CPU, network activity, full logs or file systems) • User reports of system unavailability • Network and host intrusion detection alerts • New files or directories with unusual names (e.g., binary characters, leading spaces, leading dots) • Highly unusual operating system and application log messages • Attacker contacts the organization to say that he has compromised a host
Unauthorized data modification (e.g., Web server defacement, FTP warez server ⁷⁷)	<ul style="list-style-type: none"> • Network intrusion detection alerts • Increased resource utilization • User reports of the data modification (e.g., defaced Web site) • Modifications to critical files (e.g., Web pages) • New files or directories with unusual names (e.g., binary characters, leading spaces, leading dots) • Significant changes in expected resource usage (e.g., CPU, network activity, full logs or file systems)
Unauthorized usage of standard user account	<ul style="list-style-type: none"> • Access attempts to critical files (e.g., password files) • Unexplained account usage (e.g., idle account in use, account in use from multiple locations at once, commands that are unexpected from a particular user, large number of locked-out accounts) • Web proxy log entries showing the download of hacker tools
Physical intruder	<ul style="list-style-type: none"> • User reports of network or system unavailability • System status changes (restarts, shutdowns) • Hardware is completely or partially missing (i.e., a system was opened and a particular component removed) • Unauthorized new hardware (e.g., attacker connects a packet sniffing laptop to a network or a modem to a host)
Unauthorized data access (e.g., database of customer information, password files)	<ul style="list-style-type: none"> • Intrusion detection alerts of attempts to gain access to the data through FTP, HTTP, and other protocols • Host-recorded access attempts to critical files

Unauthorized access incidents differ from other types of incidents in that they tend to occur in several steps. Typically, attackers will perform multiple reconnaissance activities to map networks; identify hosts; determine what operating system, services, and applications each host runs; and find vulnerabilities that may be remotely exploitable. Reconnaissance has become so commonplace that organizations often

⁷⁷ A *warez server* is a file server that is used to distribute illegal content. Originally the term “warez” referred to pirated software, but it now also includes other illegal content such as copies of copyrighted songs and movies. Attackers often exploit vulnerabilities in FTP servers to gain unauthorized access so they can use the server to distribute their warez files.

ignore it because of time and resource limitations. However, it is important for organizations to review reconnaissance activity at least minimally, to get a sense of the risks they currently face.

After the reconnaissance steps have been completed, attackers begin taking actions to acquire unauthorized access to systems. Many vulnerabilities permit privileged access to be gained in a single step, while other vulnerabilities only provide user-level access. Ultimately, most attackers are seeking administrator-level access to systems, so they generally look first for vulnerabilities that can grant privileged access. If such a vulnerability cannot be found or successfully exploited, attackers may attempt to find and exploit vulnerabilities that can provide user-level access and then carry out additional attacks to elevate the access level. Because this process may take a considerable amount of time, the attack may be detected at an intermediate step, when some access has been granted but additional access is being pursued. The incident response team should endeavor to detect, validate, and halt such incidents before full administrator access is gained. If the attacker gains administrator-level access, the attacker is likely to install rootkits and establish backdoors so that the attacker can access the system remotely with administrative privileges in the future.

During unauthorized access incidents, it is often difficult to distinguish benign activity from malicious activity. Indications such as system shutdowns, audit configuration changes, and executable modifications are probably caused by routine system administration, not attacks. The key to determining the source of the activity is the organization's change management process. If a system is scheduled for maintenance, such as an operating system upgrade, this information should be provided to the staff that monitor and analyze precursors and indications (including outsourcers or contractors). When suspicious indications are detected, the analyst can quickly verify that they are caused by a planned maintenance activity. Of course, if the change management process does not accurately capture the necessary information, then incident handlers may need to contact system administrators to confirm that they have performed the activity. Malicious activity may escalate to a full root compromise by the time the handler can reach the system administrator.

When prioritizing unauthorized access incidents, determining the current and likely future impact of the incident can be very difficult. Because attackers want to elevate user-level access privileges to administrator-level access, ongoing incidents could potentially result in root-level access. The current impact of the incident may be difficult to judge until extensive analysis has been done, and the incident may need to be prioritized before the analysis is complete. Therefore, it is best to prioritize unauthorized access incidents based on an estimate of the current impact, with the assumption that the impact will become more severe without intervention. Time frames can then be assigned to each impact category by the criticality of the resources that have been accessed without authorization.

6.4 Containment, Eradication, and Recovery

In addition to the general guidance presented in Section 3.3, this section gives specific recommendations for performing containment, and gathering and handling evidence for unauthorized access incidents.

6.4.1 Choosing a Containment Strategy

Response time is critical when attempting to contain an unauthorized access incident. Extensive analysis may be required to determine exactly what has happened; and in the case of an active attack, the state of things may be changing rapidly. In most cases, it is advisable to perform an initial analysis of the incident, prioritize the incident, implement initial containment measures, and then perform further analysis to determine if the containment measures were sufficient. For example, it may not be immediately apparent if an attacker has copied a system's password file. Disconnecting the system from the network while incident handlers determine if the password file was compromised should prevent an

attacker from using those passwords—assuming that the passwords are unique to that system. In most environments, however, they are not. Because of trust relationships between systems and users providing the same or similar passwords for many systems, stolen passwords are often used to access systems other than the one they were stolen from. An incident can quickly expand from a single host to dozens of hosts within minutes.

Incident handlers walk a fine line when choosing containment strategies because if they assume the worst, the containment strategy could be to shut all networks and systems down. Incident handlers should consider more moderate solutions that focus on mitigating the risks to the extent practical, rather than shutting down the whole environment for days at a time (unless, of course, the extent of the malicious activity is so great that a complete shutdown is merited). An appropriate combination of the following actions should be effective in the initial or final containment of an unauthorized access incident:

- + **Isolate the Affected Systems.** This is the simplest technique for containing an unauthorized access incident—disconnect each affected system from the network. This prevents the affected systems from being further compromised. However, it can be challenging to identify all affected systems. Attackers often use one compromised system as the source of attacks against other internal systems. Handlers should examine other systems for signs of successful attacks and contain those components of the incident as well. If many systems need to be checked, automated methods could be used, such as performing port scans for backdoors.
- + **Disable the Affected Service.** If an attacker is using a particular service to gain unauthorized access, containing the incident may include temporarily or permanently disabling the service. For example, if the attacker is exploiting an FTP vulnerability and the unauthorized access is limited to the FTP data files, the incident could be contained by temporarily disabling the FTP service. If the server was inadvertently running FTP, then FTP should be disabled permanently.
- + **Eliminate the Attacker’s Route into the Environment.** If possible, prevent the attacker from accessing nearby resources that might be the next targets while minimizing disruption of services to authorized users. Examples include temporarily blocking incoming connections to a particular network segment or disconnecting a remote access server.
- + **Disable User Accounts That May Have Been Used in the Attack.** The same accounts and passwords that were acquired from one system may work on other systems; therefore, the accounts may need to be disabled across the enterprise. Handlers should also look for new user accounts that may have been created by the attacker.
- + **Enhance Physical Security Measures.** If an unauthorized access incident involves a breach of physical security, additional containment strategies should be followed. For example, if an outsider is suspected of gaining access to a server room, not only should the server room be secured more strongly, but also the physical security staff or law enforcement may need to search the facility to confirm that the intruder is not still present. Other security changes may be merited; if the attacker can breach security in one instance, other opportunities may present themselves.

6.4.2 Evidence Gathering and Handling

When handlers suspect that unauthorized access has been gained to a system, they should make a full image backup of the system. Other relevant data, including host and application logs, intrusion detection alerts, and firewall logs, may provide correlating evidence of the unauthorized access. If a physical security breach occurred during the incident, additional evidence may be available through physical security system logs, security camera tapes, and eyewitness accounts. Unauthorized access incidents are more likely than most other incidents to lead to prosecution, so it is important to follow established

evidence gathering and handling procedures, and to contact law enforcement if the situation merits their involvement.

6.4.3 Eradication and Recovery

Successful attackers frequently install rootkits, which modify or replace dozens or hundreds of files, including system binaries. Rootkits hide much of what they do, making it tricky to identify what was changed.⁷⁸ Therefore, if an attacker appears to have gained root access to a system, handlers cannot trust the operating system software. Typically, the best solution is to restore the system from a known good backup or reinstall the operating system and applications from scratch, and then secure the system properly. Changing all passwords on the system, and possibly on all systems that have trust relationships with the victim system, is also highly recommended. Some unauthorized access incidents involve the exploitation of multiple vulnerabilities, so it is important for handlers to identify all vulnerabilities that were used and to determine strategies for correcting or mitigating each vulnerability. Other vulnerabilities that are present should be mitigated as well, or an attacker may use them instead.

If an attacker only gains a lesser level of access than administrator-level, eradication and recovery actions should be based on the extent to which the attacker gained access. Vulnerabilities that were used to gain access should be mitigated appropriately. Additional actions should be performed as merited to identify and address weaknesses systemically. For example, if an attacker gained user-level access by guessing a weak password, then not only should that account’s password be changed to a stronger password, but also the system administrator and owner should consider enforcing stronger password requirements. If the system was in compliance with the organization’s password policies, the organization should consider revising its password policies.

6.5 Checklist for Handling Unauthorized Access Incidents

The checklist in Table 6-4 provides the major steps to be performed in handling an unauthorized access incident. This checklist is a continuation of the Initial Incident Handling Checklist in Table 3-5. Note that the exact sequence of steps may vary based on the nature of individual incidents and on the strategies chosen by the organization for containing incidents.

Table 6-4. Unauthorized Access Incident Handling Checklist

	Action	Completed
Detection and Analysis		
1.	Prioritize handling the incident based on its business impact	
1.1	Identify which resources have been affected and forecast which resources will be affected	
1.2	Estimate the current technical effect of the incident	
1.3	Find the appropriate cell(s) in the prioritization matrix, based on the technical effect and affected resources	
2.	Report the incident to the appropriate internal personnel and external organizations	
Containment, Eradication, and Recovery		
3.	Perform an initial containment of the incident	
4.	Acquire, preserve, secure, and document evidence	

⁷⁸ Tools such as chkrootkit (<http://www.chkrootkit.org/>) can be useful in determining whether a rootkit has been installed on a system.

	Action	Completed
5.	Confirm the containment of the incident	
5.1	Further analyze the incident and determine if containment was sufficient (including checking other systems for signs of intrusion)	
5.2	Implement additional containment measures if necessary	
6.	Eradicate the incident	
6.1	Identify and mitigate all vulnerabilities that were exploited	
6.2	Remove components of the incident from systems	
7.	Recover from the incident	
7.1	Return affected systems to an operationally ready state	
7.2	Confirm that the affected systems are functioning normally	
7.3	If necessary, implement additional monitoring to look for future related activity	
Post-Incident Activity		
8.	Create a follow-up report	
9.	Hold a lessons learned meeting	

6.6 Recommendations

The items below summarize the key recommendations presented in this section for handling unauthorized access incidents.

- + **Configure intrusion detection software to alert on attempts to gain unauthorized access.** Network and host-based intrusion detection software (including file integrity checking software) is valuable for detecting attempts to gain unauthorized access. Each type of software may detect incidents that the other types of software cannot, so the use of multiple types of computer security software is highly recommended.
- + **Configure all hosts to use centralized logging.** Incidents are easier to detect if data from all hosts across the organization is stored in a centralized, secured location.
- + **Establish procedures for having all users change their passwords.** A password compromise may force the organization to require all users of an application, system, or trust domain—or perhaps the entire organization—to change their passwords.
- + **Configure the network perimeter to deny all incoming traffic that is not expressly permitted.** By limiting the types of incoming traffic, attackers should be able to reach fewer targets and should be able to reach the targets using designated protocols only. This should reduce the number of unauthorized access incidents.
- + **Secure all remote access methods, including modems and VPNs.** Unsecured modems provide easily attainable unauthorized access to internal systems and networks. Remote access clients are often outside the organization's control, so granting them access to resources increases risk.
- + **Put all publicly accessible services on secured DMZ network segments.** This permits the organization to allow external hosts to initiate connections to hosts on the DMZ segments only, not to hosts on internal network segments. This should reduce the number of unauthorized access incidents.
- + **Disable all unneeded services on hosts and separate critical services.** Every service that is running presents another potential opportunity for compromise. Separating critical services is

important because if an attacker compromises a host that is running a critical service, immediate access should be gained only to that one service.

- + **Use host-based firewall software to limit individual hosts' exposure to attacks.** Deploying host-based firewall software to individual hosts and configuring it to deny all activity that is not expressly permitted should further reduce the likelihood of unauthorized access incidents.
- + **Create and implement a password policy.** The password policy should require the use of complex, difficult-to-guess passwords, as well as ensure that authentication methods are sufficiently strong for accessing critical resources. Weak and default passwords are likely to be guessed or cracked, leading to unauthorized access.
- + **Provide change management information to the incident response team.** Indications such as system shutdowns, audit configuration changes, and executable modifications are probably caused by routine system administration, rather than attacks. When such indications are detected, the team should be able to use change management information to verify that the indications are caused by authorized activity.
- + **Select containment strategies that balance mitigating risks and maintaining services.** Incident handlers should consider moderate containment solutions that focus on mitigating the risks as much as is practical while maintaining unaffected services.
- + **Restore or reinstall systems that appear to have suffered a root compromise.** The effects of root compromises are often difficult to identify completely. The system should be restored from a known good backup, or the operating system and applications should be reinstalled from scratch. The system should then be secured properly so the incident cannot recur.

7. HANDLING INAPPROPRIATE USAGE INCIDENTS

7.1 Incident Definition and Examples

An *inappropriate usage* incident occurs when a user performs actions that violate acceptable computing use policies. Although such incidents are often not security-related, handling them is very similar to handling security-related incidents. Therefore, it has become commonplace for incident response teams to handle many inappropriate usage incidents. Examples of incidents a team might handle include users who—

- + Download password cracking tools or pornography
- + Send spam promoting a personal business
- + E-mail harassing messages to coworkers
- + Set up an unauthorized Web site on one of the organization's computers
- + Use file or music sharing services to acquire or distribute pirated materials
- + Transfer sensitive materials from the organization to external locations.

Certain inappropriate usage incidents are more challenging to handle because they are targeted at outside parties. Of course, this raises liability concerns. What makes these incidents particularly interesting is that in some cases, the organization is not actually the source of the attacks—but it appears to outside parties that the organization attacked them. The handlers should work quickly to investigate the activity, collect evidence, and determine if the activity originated from the organization's networks or systems. Examples of inappropriate usage incidents directed at outside parties include—

- + An internal user defacing another organization's public Web site
- + An internal user purchasing items from online retailers with stolen credit card numbers
- + A third party sending spam e-mails with spoofed source e-mail addresses that appear to belong to the organization
- + A third party performing a DoS against an organization by generating packets with spoofed source IP addresses that belong to the organization.

7.2 Preparation

This section provides guidance on preparing to handle inappropriate usage incidents and on preventing inappropriate usage incidents.

7.2.1 Incident Handling Preparation

In addition to the general guidance presented in Sections 3.1.1 and 3.2.3, other actions should be performed while preparing to handle inappropriate usage incidents:

- + Meet with representatives of the organization's human resources and legal departments to discuss the handling of inappropriate usage incidents. Monitoring and logging of user activities should comply with the organization's policies. Also, the incident response team should understand the intricacies of handling incidents that directly involve employees. For example, early indications may imply that a

particular employee has been downloading pornography; further analysis shows that someone else used the employee's account. Discretion and confidentiality should be incorporated into incident response procedures.

- + Discuss liability issues with the organization's public affairs and legal departments, particularly for incidents that are targeted at outside parties. It is critical for incident handlers to understand when they should discuss incidents with the allegedly attacked party and what information they should reveal.
- + Configure network-based IDS software to identify certain types of activity, including—
 - The use of unauthorized services, such as peer-to-peer file and music sharing
 - Spam (e.g., mail relaying attempts)
 - Any file activity (e.g., e-mail attachments, FTP transfers, Web requests) with suspicious words in the filename (e.g., “confidential,” sexually explicit terms)
 - Outbound reconnaissance activity and attacks.
- + Log user activities such as FTP commands, Web requests, and e-mail headers. This may be done through proxies and application logs or by some network-based IDS sensors. The goal is to log the basic information on such activities without storing sensitive content, such as e-mail text and file attachments. Organizations should balance privacy considerations with the value of such information for investigative and evidentiary purposes.

7.2.2 Incident Prevention

Generally, little can be done to prevent inappropriate usage incidents from occurring, other than requiring users to read and sign an acceptable use policy and informing users that their activities are regularly monitored. However, there are suggested actions that may be of some help in reducing certain types of inappropriate usage incidents:

- + Configure network devices to prevent the use of services that violate organization policies, such as peer-to-peer file sharing and music sharing services. However, blocking this traffic may be impractical because such services may use millions of workstations around the world and be tunneled over various ports.
- + Configure the organization's e-mail servers so that they cannot be used for unauthorized mail relaying, a common way to send spam.⁷⁹
- + Implement spam filtering software on all e-mail servers. This should not only block much of the spam sent to users from external parties but also help to prevent internal users from sending spam.
- + Implement uniform resource locator (URL) filtering to prevent access to inappropriate Web sites. Of course, this is only effective if users are forced to use it. This can be accomplished by creating Web proxy servers that run the URL filtering software and configuring network firewalls so that all outgoing Web requests must be made by the proxy servers. Users must go through one of the proxy servers to access external Web sites.

⁷⁹ The Mail Abuse Prevention System (MAPS) Web site (<http://mail-abuse.org/tsi/ar-fix.html>) provides advice on preventing mail relaying for dozens of e-mail programs.

- + Consider limiting outbound connections that use encrypted protocols, such as Secure Shell (SSH), HTTP Secure (HTTPS) and IP Security Protocol (IPSec). Permitting unnecessary encrypted connections may allow users to perform actions that security controls cannot monitor. For example, a user could establish an SSH connection to an external server and download illegal materials; because the connection is encrypted, network security controls could not determine the nature of the activity. Possible methods for limiting the traffic include firewall rulesets and URL filtering (e.g., blocking access to public HTTPS proxy servers).

7.3 Detection and Analysis

Inappropriate usage incidents are most often detected through user reports, such as seeing inappropriate material on a user’s screen or receiving a threatening e-mail. There are usually no precursors of inappropriate usage.⁸⁰ Table 7-1, Inappropriate Usage Indications, lists actions such as unauthorized service usage and access to inappropriate materials. For each such action, the table lists possible indications of the action. The table can easily be customized by the organization to include environment-specific precursors and indications, which should facilitate a more efficient and effective incident handling process.

Table 7-1. Inappropriate Usage Indications

Inappropriate Action	Possible Indications
Unauthorized service usage (e.g., Web server, file sharing, music sharing)	<ul style="list-style-type: none"> • Network intrusion detection alerts • Unusual traffic to and from the host • New process/software installed and running on a host • New files or directories with unusual names (e.g., “warez” server style names) • Increased resource utilization (e.g., CPU, file storage, network activity) • User reports • Application log entries (e.g., Web proxies, FTP servers, e-mail servers)
Access to inappropriate materials (e.g., downloading pornography, sending spam)	<ul style="list-style-type: none"> • Network intrusion detection alerts • User reports • Application log entries (e.g., Web proxies, FTP servers, e-mail servers) • Inappropriate files on workstations, servers, or removable media
Attack against external party	<ul style="list-style-type: none"> • Network intrusion detection alerts • Outside party reports • Network, host, and application log entries⁸¹

The incident response team should be cautious about assisting with reports of inappropriate usage that are not clearly incidents. For example, a manager might report that an employee seems to be wasting a lot of time with computers, but the manager does not have any idea what the employee is doing. The manager may ask the team to monitor the employee’s Internet usage and analyze the user’s hard drive. The team should not assist with such requests unless appropriate management and human resources personnel provide written approval.

Analyzing inappropriate usage incidents is typically straightforward, except for incidents that have been reported by outside parties. The key to analyzing such incidents is determining whether the organization was really the source of the attack or if spoofing has simply created that appearance. This should be fairly easy to determine if proper logging is being performed and the organization has good security controls in

⁸⁰ One exception is when an internal user performs reconnaissance before attacking an external network or host. In that case, the same precursors that apply to denial of service or unauthorized access incidents are applicable.

⁸¹ Examples of indications include e-mail server log entries of bounced e-mails with forged source addresses and firewall log entries of TCP RST packets that do not have a corresponding SYN (i.e., backscatter from spoofed packets).

place. For example, if firewalls are configured to log all connections and connection attempts, the firewall logs should record any outgoing attacks—assuming that all outgoing communications pass through a firewall. If the organization’s perimeter is unsecured, it will be more difficult to determine whether the attack came from the organization because the attacker could have taken another route that circumvents the firewall.

Another factor that can complicate incident analysis is when the identity of the person causing the incident cannot be determined by reviewing existing data. Increasing the monitoring of computing or physical resources is effective in identifying the individual if the activity continues. A more challenging alternative is to generate a profile of the suspected perpetrator’s usage characteristics and goals, and to work with human resources on expanding the profile. This technique is effective only in certain cases. Proactive monitoring is the recommended method of identifying who is causing incidents.

Inappropriate usage incidents are generally easy to prioritize. Unless a crime is involved or the public will learn what happened, these incidents do not need to be handled with the same urgency as other incidents. Table 7-2 shows a sample matrix for prioritizing inappropriate usage incidents. Although it is constructed differently from the matrix presented in Section 3.2.6, both matrices prioritize responses based on the business impact of the incident. Table 7-2 defines the business impact by two factors: (1) whether the activity is criminal or not, and (2) which individuals or groups are aware of the incident or are likely to become aware of it. Criminal activity dictates a faster response for evidentiary reasons. Incidents that will become public knowledge should generally be handled more quickly than those that will be known only to a few individuals within the organization.

Table 7-2. Sample SLA for Inappropriate Usage Incidents

Current or likely awareness of the incident	Nature of the incident	
	Criminal activity	Non-criminal activity
Public (including an outside party that says it was attacked by the organization)	15 minutes Within 1 hour, team contacts public affairs, human resources, legal department, and law enforcement	1 hour Within 2 hours, team contacts public affairs and human resources
Some groups within the organization	2 hours Within 4 hours, team contacts human resources, legal department, and law enforcement	4 hours Within 8 hours, team contacts public affairs and human resources
Few individuals within the organization	4 hours Within 8 hours, team contacts human resources, legal department, and law enforcement	1 day Within 2 days, team contacts public affairs and human resources

There is one caveat to prioritizing inappropriate usage incidents. A sizable number of these incidents are actually follow-on activities to previous incidents, such as a root compromise of a host or a successful malicious code infection. Prioritizing an incident that encompasses two or more incidents will be discussed in Section 8.

7.4 Containment, Eradication, and Recovery

Inappropriate usage incidents typically require no containment, eradication, or recovery actions, other than possibly deleting objectionable materials or uninstalling unauthorized software.⁸² For most inappropriate usage incidents, evidence acquisition is important. Evidence may be needed for prosecuting or disciplining an individual and for limiting liability by demonstrating that the organization did its best to prevent, detect, and halt the activity. Evidence storage is particularly important, as internal users have physical access to many facilities. Addressing the threat of having evidence altered or destroyed may require coordination with the organization’s physical security staff.

7.5 Checklist for Handling Inappropriate Usage Incidents

The checklist in Table 7-3 provides the major steps to be performed in handling an inappropriate usage incident. This checklist is a continuation of the Initial Incident Handling Checklist in Table 3-5. Note that the sequence of steps may vary based on the nature of individual incidents.

Table 7-3. Inappropriate Usage Incident Handling Checklist

	Action	Completed
Detection and Analysis		
1.	Prioritize the handling of the incident based on its business impact	
1.1	Determine whether the activity seems criminal in nature	
1.2	Forecast who will likely become aware of the activity	
1.3	Find the appropriate cell(s) in the prioritization matrix, based on the criminality and public awareness	
2.	Report the incident to the appropriate internal personnel and external organizations	
Containment, Eradication, and Recovery		
3.	Acquire, preserve, secure, and document evidence	
4.	If necessary, contain and eradicate the incident (e.g., remove inappropriate materials)	
Post-Incident Activity		
5.	Create a follow-up report	
6.	Hold a lessons learned meeting	

7.6 Recommendations

The items below summarize the key recommendations presented in this section for handling inappropriate usage incidents.

- + **Discuss the handling of inappropriate usage incidents with the organization’s human resources and legal departments.** Processes for monitoring and logging user activities should comply with the organization’s policies and all applicable laws. Procedures for handling incidents that directly involve employees should incorporate discretion and confidentiality.

⁸² An exception to this is when a user is attacking another organization; such incidents should be contained as quickly as possible to prevent additional damage to others’ systems and to limit potential liability. Another possible exception is that the incident response team may want to inform an external organization that one of their hosts contains illegal materials.

- + **Discuss liability issues with the organization's legal departments.** Liability issues may arise during inappropriate usage incidents, particularly for incidents that are targeted at outside parties. Incident handlers should understand when they should discuss incidents with the allegedly attacked party and what information they should reveal.
- + **Configure network-based intrusion detection software to detect certain types of inappropriate usage.** Intrusion detection software has built-in capabilities to detect certain inappropriate usage incidents, such as the use of unauthorized services, outbound reconnaissance activity and attacks, and improper mail relay usage (e.g., sending spam).
- + **Log basic information on user activities.** Basic information on user activities such as FTP commands, Web requests, and e-mail headers may be valuable for investigative and evidentiary purposes.
- + **Configure all e-mail servers so they cannot be used for unauthorized mail relaying.** Mail relaying is commonly used to send spam.
- + **Implement spam filtering software on all e-mail servers.** Spam filtering software can block much of the spam sent by external parties to the organization's users, as well as spam that is sent by internal users.
- + **Implement URL filtering software.** URL filtering software prevents access to many inappropriate Web sites. Users should be required to use the software, typically by preventing access to external Web sites unless the traffic passes through a server that performs URL filtering.

8. HANDLING MULTIPLE COMPONENT INCIDENTS

8.1 Incident Definition and Examples

A *multiple component incident* is a single incident that encompasses two or more incidents. Figure 8-1 provides an example of the steps that could comprise a multiple component incident:

1. Malicious code spread through e-mail compromises an internal workstation.
2. An attacker (who may or may not be the one who sent the malicious code) uses the infected workstation to compromise additional workstations and servers.
3. An attacker (who may or may not have been involved in steps 1 or 2) uses one of the compromised hosts to launch a DDoS attack against another organization.

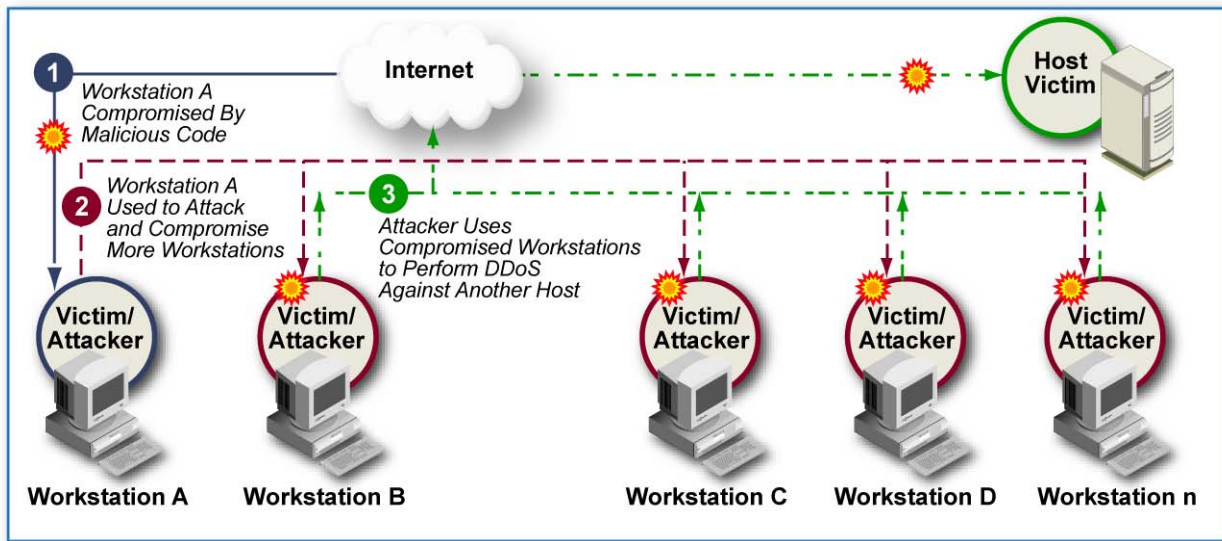


Figure 8-1. Example of a Multiple Component Incident

This multiple component incident consists of a malicious code incident, several unauthorized access incidents, and a DoS incident. If the organization is made aware of the incident by a complaint from the target of the DoS attack, then an inappropriate usage incident has also occurred. As shown by the example, a multiple component incident may involve several incidents performed by different perpetrators. This factor complicates the incident analysis process.

8.2 Preparation, Detection, and Analysis

Multiple component incidents tend to be hard to analyze. Incident handlers may know about one portion of the incident only and may not realize that the incident is composed of several stages. Handlers may also be aware of multiple incidents but not realize that they are related. Further complicating the analysis is that the incident's stages may occur over a period of weeks or months. Unless the organization has excellent logging and log archiving processes in place, the evidence of earlier stages of the incident may be gone. Even if the data is available, it can be challenging for the analyst to determine which indications are related among all the data.

The main preparation for handling multiple component incidents is the same as that previously noted for each individual incident category. Another helpful activity is to conduct exercises in which the incident response team reviews scenarios involving multiple component incidents. The use of centralized logging and correlation software has already been recommended for facilitating more efficient incident analysis. This is particularly true for analyzing multiple component incidents, which typically have several precursor and indicator sources. Incident handlers should diagnose an incident as having multiple components more quickly if all precursors and indicators are accessible from a single point of view.

8.3 Containment, Eradication, and Recovery

Handlers should not become fixated on immediately determining all components of an incident. Every incident that is detected could be a multiple component incident, but it could take an extended period of time for a handler to authoritatively determine that an incident has only a single component. Meanwhile, the initial incident has not been contained. It is generally better to contain the initial incident and then search for signs of other components. Experienced handlers should be able to make an educated guess as to whether an incident has other components. It can be generally assumed that unauthorized access incidents are more likely to have multiple components, and other types of incidents are less likely to have multiple components.

Handlers who are aware of multiple components of an incident should separately prioritize the handling of each component, because not enough resources will likely be available to handle all components simultaneously. If the organization has created a prioritization matrix for each incident category, the handler can look up each component in the proper prioritization matrix and follow the most urgent instructions first. Another factor to consider is how current each component is—a DoS attack in progress should usually be addressed more quickly than a malicious code infection that occurred 6 weeks ago. Also, if one component creates a path for attackers to reach targets, the handler may be able to contain the whole incident by containing just that one component. (Of course, the other components will still need to be handled, just not as urgently.) Handlers should be cautious, though, as attackers may have created or discovered additional paths to the targets.

8.4 Checklist for Handling Multiple Component Incidents

The checklist in Table 8-1 provides the major steps to be performed in handling a multiple component incident. This checklist is a continuation of the Initial Incident Handling Checklist in Table 3-5. Note that the sequence of steps may vary based on the nature of individual incidents and on the strategies chosen by the organization for containing incidents.

Table 8-1. Multiple Component Incident Handling Checklist

	Action	Completed
Detection and Analysis		
1.	Prioritize handling the incident based on its business impact	
1.1	Follow the step 1 instructions for each applicable incident category	
1.2	Determine the proper course of action for each incident component	
2.	Report the incident to the appropriate internal personnel and external organizations	
Containment, Eradication, and Recovery		
3.	Follow the Containment, Eradication, and Recovery steps for each component, based on the results of step 1	
Post-Incident Activity		
4.	Create a follow-up report	
5.	Hold a lessons learned meeting	

8.5 Recommendations

The items below summarize the key recommendations presented in this section for handling multiple component incidents.

- + **Use centralized logging and event correlation software.** Incident handlers should identify an incident as having multiple components more quickly if all precursors and indications are accessible from a single point of view.
- + **Contain the initial incident and then search for signs of other incident components.** It can take an extended period of time for a handler to authoritatively determine that an incident has only a single component; meanwhile, the initial incident has not been contained. It is generally better to contain the initial incident first.
- + **Separately prioritize the handling of each incident component.** Resources are probably too limited to handle all incident components simultaneously. Components should be prioritized based on response guidelines for each component and how current each component is.

Appendix A—RECOMMENDATIONS

Appendix A lists the major recommendations presented in sections 2 through 8 of this document. The first group of recommendations applies to organizing an incident response capability. The remaining recommendations have been grouped by the phases of the incident response lifecycle—preparation; detection and analysis; containment, eradication, and recovery; and post-incident activity. Each group contains general recommendations for its incident response phase, as well as any applicable recommendations for handling particular categories of incidents (e.g., denial of service) during the phase.

A.1 Organizing a Computer Security Incident Response Capability

- + **Establish a formal incident response capability.** Organizations should be prepared to respond quickly and effectively when computer security defenses are breached. FISMA requires Federal agencies to establish incident response capabilities.

A.1.1 Incident Response Policy and Procedure Creation

- + **Create an incident response policy and use it as the basis for incident response procedures.** The incident response policy is the foundation of the incident response program. It defines which events are considered incidents, establishes the organizational structure for incident response, defines roles and responsibilities, and lists the requirements for reporting incidents, among other items.
- + **Establish policies and procedures regarding incident-related information sharing.** The organization will want or be required to communicate incident details with outside parties, such as the media, law enforcement agencies and incident reporting organizations. The incident response team should discuss this requirement at length with the organization's public affairs and legal advisors to establish policies and procedures regarding information sharing. The team should comply with existing organization policy on interacting with the media and other outside parties.
- + **Provide pertinent information on incidents to the appropriate incident reporting organization.** Federal civilian agencies are required to report incidents to FedCIRC. Reporting benefits the agencies because the incident reporting organizations use the reported data to provide information to the agencies regarding new threats and incident trends.

A.1.2 Incident Response Team Structure and Services

- + **Consider the relevant factors when selecting an appropriate incident response team model.** Organizations should carefully weigh the advantages and disadvantages of each possible team structure model and staffing model in the context of the organization's needs and available resources.
- + **Select people with appropriate skills for the incident response team.** The credibility and proficiency of the team depend to a large extent on the technical skills of its members. Poor technical judgment can undermine the team's credibility and cause incidents to worsen. Critical technical skills include system administration, network administration, programming, technical support, and intrusion detection. Teamwork and communications skills are also needed for effective incident handling.
- + **Identify other groups within the organization that may need to participate in incident handling.** Every incident response team relies on the expertise and judgment of other teams, including management, information security, IT support, legal, public affairs, and facilities management.
- + **Determine which services the team should offer.** Although the main focus of the team is incident response, most teams perform additional functions. Examples include distributing security

advisories, performing vulnerability assessments, educating users on security, and monitoring intrusion detection sensors.

A.2 Incident Preparation

- + **Acquire tools and resources that may be of value during incident handling.** The team will be more efficient at handling incidents if various tools and resources are already available to them. Examples include contact lists, encryption software, network diagrams, backup devices, computer forensic software, port lists, and security patches.
- + **Prevent incidents from occurring by ensuring that networks, systems, and applications are sufficiently secure.** Of course, preventing incidents is beneficial to the organization and also reduces the workload of the incident response team. Performing periodic risk assessments and reducing the identified risks to an acceptable level are effective in reducing the number of incidents. User, IT staff, and management awareness of security policies and procedures is also very important.

A.2.1 Denial of Service Incidents

- + **Configure firewall rulesets to prevent reflector attacks.** Most reflector attacks can be stopped through network-based and host-based firewall rulesets that reject suspicious combinations of source and destination ports.
- + **Configure border routers to prevent amplifier attacks.** Amplifier attacks can be blocked by configuring border routers not to forward directed broadcasts.
- + **Determine how the organization's ISPs and second-tier providers can assist in handling network-based DoS attacks.** ISPs can often filter or limit certain types of traffic, slowing or halting a DoS attack. They can also provide logs of DoS traffic and may be able to assist in tracing the source of the attack. The organization should meet with the ISPs in advance to establish procedures for requesting such assistance.
- + **Configure security software to detect DoS attacks.** Intrusion detection software can detect many types of DoS activity. Establishing network and system activity baselines, and monitoring for significant deviations from those baselines, can also be useful in detecting attacks.
- + **Configure the network perimeter to deny all incoming and outgoing traffic that is not expressly permitted.** By restricting the types of traffic that can enter and leave the environment, the organization will limit the methods that attackers can use to perform DoS attacks.

A.2.2 Malicious Code Incidents

- + **Make users aware of malicious code issues.** Users should be familiar with the methods that malicious code uses to propagate and the symptoms of infections. Holding regular user education sessions helps to ensure that users are aware of the risks that malicious code poses. Teaching users how to safely handle e-mail attachments should reduce the number of infections that occur.
- + **Read antivirus bulletins.** Bulletins regarding new malicious code threats provide timely information to incident handlers.
- + **Deploy host-based intrusion detection systems, including file integrity checkers, to critical hosts.** Host-based IDS software, particularly file integrity checkers, can detect signs of malicious code incidents, such as configuration changes and modifications to executables.

- + **Use antivirus software, and keep it updated with the latest virus signatures.** Antivirus software should be deployed to all hosts and all applications that may be used to transfer malicious code. The software should be configured to detect and disinfect or quarantine malicious code infections. All antivirus software should be kept current with the latest virus signatures so the newest threats can be detected.
- + **Configure software to block suspicious files.** Files that are very likely to be malicious should be blocked from the environment, such as those with file extensions that are usually associated with malicious code, as well as files with suspicious combinations of file extensions.
- + **Eliminate open Windows shares.** Many worms spread through unsecured shares on hosts running Windows. A single infection may rapidly spread to hundreds or thousands of hosts through unsecured shares.

A.2.3 Unauthorized Access Incidents

- + **Configure intrusion detection software to alert on attempts to gain unauthorized access.** Network and host-based intrusion detection software (including file integrity checking software) is valuable for detecting attempts to gain unauthorized access. Each type of software may detect incidents that the other types of software cannot, so the use of multiple types of computer security software is highly recommended.
- + **Configure all hosts to use centralized logging.** Incidents are easier to detect if data from all hosts across the organization is stored in a centralized, secured location.
- + **Establish procedures for having all users change their passwords.** A password compromise may force the organization to require all users of an application, system, or trust domain—or perhaps the entire organization—to change their passwords.
- + **Configure the network perimeter to deny all incoming traffic that is not expressly permitted.** By limiting the types of incoming traffic, attackers should be able to reach fewer targets and should be able to reach the targets using designated protocols only. This should reduce the number of unauthorized access incidents.
- + **Secure all remote access methods, including modems and VPNs.** Unsecured modems provide easily attainable unauthorized access to internal systems and networks. Remote access clients are often outside the organization's control, so granting them access to resources increases risk.
- + **Put all publicly accessible services on secured DMZ network segments.** This permits the organization to allow external hosts to initiate connections to hosts on the DMZ segments only, not to hosts on internal network segments. This should reduce the number of unauthorized access incidents.
- + **Disable all unneeded services on hosts and separate critical services.** Every service that is running presents another potential opportunity for compromise. Separating critical services is important because if an attacker compromises a host that is running a critical service, immediate access should be gained only to that one service.
- + **Use host-based firewall software to limit individual hosts' exposure to attacks.** Deploying host-based firewall software to individual hosts and configuring it to deny all activity that is not expressly permitted should further reduce the likelihood of unauthorized access incidents.
- + **Create and implement a password policy.** The password policy should require the use of complex, difficult-to-guess passwords, as well as ensure that authentication methods are sufficiently strong for

accessing critical resources. Weak and default passwords are likely to be guessed or cracked, leading to unauthorized access.

A.2.4 Inappropriate Usage Incidents

- + **Discuss the handling of inappropriate usage incidents with the organization's human resources and legal departments.** Processes for monitoring and logging user activities should comply with the organization's policies and all applicable laws. Procedures for handling incidents that directly involve employees should incorporate discretion and confidentiality.
- + **Discuss liability issues with the organization's legal departments.** Liability issues may arise during inappropriate usage incidents, particularly for incidents that are targeted at outside parties. Incident handlers should understand when they should discuss incidents with the allegedly attacked party and what information they should reveal.
- + **Configure network-based intrusion detection software to detect certain types of inappropriate usage.** Intrusion detection software has built-in capabilities to detect certain inappropriate usage incidents, such as the use of unauthorized services, outbound reconnaissance activity and attacks, and improper mail relay usage (e.g., sending spam).
- + **Log basic information on user activities.** Basic information on user activities such as FTP commands, Web requests, and e-mail headers may be valuable for investigative and evidentiary purposes.
- + **Configure all e-mail servers so they cannot be used for unauthorized mail relaying.** Mail relaying is commonly used to send spam.
- + **Implement spam filtering software on all e-mail servers.** Spam filtering software can block much of the spam sent by external parties to the organization's users, as well as spam that is sent by internal users.
- + **Implement URL filtering software.** URL filtering software prevents access to many inappropriate Web sites. Users should be required to use the software, typically by preventing access to external Web sites unless the traffic passes through a server that performs URL filtering.

A.2.5 Multiple Component Incidents

- + **Use centralized logging and event correlation software.** Incident handlers should identify an incident as having multiple components more quickly if all precursors and indications are accessible from a single point of view.

A.3 Incident Detection and Analysis

- + **Identify precursors and indications through alerts generated by several types of computer security software.** Network and host-based intrusion detection systems, antivirus software, and file integrity checking software are valuable for detecting signs of incidents. Each type of software may detect incidents that the other types of software cannot, so the use of several types of computer security software is highly recommended. Third-party monitoring services can also be helpful.
- + **Establish mechanisms for outside parties to report incidents.** Outside parties may want to report incidents to the organization; for example, they may believe that one of the organization's users is

attacking them. Organizations should publish a phone number and e-mail address that outside parties can use to report such incidents.

- + **Require a baseline level of logging and auditing on all systems, and a higher baseline level on all critical systems.** Logs from operating systems, services, and applications frequently provide value during incident analysis, particularly if auditing was enabled. The logs can provide information such as which accounts were accessed and what actions were performed.
- + **Profile networks and systems.** Profiling measures the characteristics of expected activity levels so that changes in patterns can be more easily identified. If the profiling process is automated, deviations from expected activity levels can be detected and reported to administrators quickly, leading to faster detection of incidents and operational issues.
- + **Understand the normal behaviors of networks, systems, and applications.** Team members who understand what normal behavior is should be able to recognize abnormal behavior more easily. This knowledge can best be gained by reviewing log entries and security alerts; the handlers should become familiar with the typical data and can investigate the unusual entries to gain more knowledge.
- + **Use centralized logging and create a log retention policy.** Information regarding an incident may be recorded in several places. Organizations should deploy centralized logging servers and configure devices to send duplicates of their log entries to the centralized servers. The team benefits because it can access all log entries at once; also, changes made to logs on individual hosts will not affect the data already sent to the centralized servers. A log retention policy is important because older log entries may show previous instances of similar or related activity.
- + **Perform event correlation.** Indications of an incident may be captured in several logs. Correlating events among multiple sources can be invaluable in collecting all the available information for an incident and validating whether the incident occurred. Centralized logging makes event correlation easier and faster.
- + **Keep all host clocks synchronized.** If the devices reporting events have inconsistent clock settings, event correlation will be more difficult. Clock discrepancies may also cause issues from an evidentiary standpoint.
- + **Maintain and use a knowledge base of information.** Handlers need to reference information quickly during incident analysis; a centralized knowledge base provides a consistent, maintainable source of information. The knowledge base should include general information, such as commonly used port numbers and links to virus information, as well as data on precursors and indications of previous incidents.
- + **Create a diagnosis matrix for less experienced staff.** Help desk staff, system administrators, and new incident response team members may need assistance in determining what type of incident may be occurring. A diagnosis matrix that lists incident categories and the symptoms associated with each category can provide guidance as to what type of incident is occurring and how the incident can be validated.
- + **Start recording all information as soon as the team suspects that an incident has occurred.** Every step taken, from the time the incident was detected to its final resolution, should be documented and timestamped. Information of this nature can serve as evidence in a court of law if legal prosecution is pursued. Recording the steps performed can also lead to a more efficient and systematic, and less error-prone handling of the problem.

- + **Safeguard incident data.** It often contains sensitive information regarding such things as vulnerabilities, security breaches, and users that may have performed inappropriate actions. The team should ensure that access to incident data is restricted properly, both logically and physically.
- + **Prioritize incidents by business impact, based on the criticality of the affected resources and the technical impact of the incident.** Because of resource limitations, incidents should not be handled on a first-come, first-served basis. Instead, organizations should establish written guidelines that outline how quickly the team must respond to the incident and what actions should be performed, based on the incident's current and potential business impact. This guidance saves time for the incident handlers and provides a justification to management and system owners for their actions. Organizations should also establish an escalation process for those instances when the team does not respond to an incident within the designated time.
- + **Include provisions regarding incident reporting in the organization's incident response policy.** Organizations should specify which incidents must be reported, when they must be reported, and to whom. The parties most commonly notified are the CIO, head of information security, local information security officer, other incident response teams within the organization, and system owners.

A.4 Containment, Eradication, and Recovery

- + **Establish strategies and procedures for containing incidents.** It is important to contain incidents quickly and effectively to limit their business impact. Organizations should define acceptable risks in containing incidents and develop strategies and procedures accordingly. Containment strategies should vary based on the type of incident.
- + **Follow established procedures for evidence gathering and handling.** The team should clearly document how all evidence has been preserved. Evidence should be accounted for at all times. The team should meet with legal staff and law enforcement agencies to discuss evidence handling, then develop procedures based on those discussions.
- + **Capture volatile data from systems as evidence.** This includes lists of network connections, processes, login sessions, open files, network interface configurations, and the contents of memory. Running carefully chosen commands from trusted media can collect the necessary information without damaging the system's evidence.
- + **Obtain system snapshots through full forensic disk images, not file system backups.** Disk images should be made to sanitized write-protectable or write-once media. This process is superior to a file system backup for investigatory and evidentiary purposes. Imaging is also valuable in that it is much safer to analyze an image than it is to perform analysis on the original system because the analysis may inadvertently alter the original.

A.4.1 Denial of Service Incidents

- + **Create a containment strategy that includes several solutions in sequence.** The decision-making process for containing DoS incidents is easier if recommended solutions are predetermined. Because the effectiveness of each possible solution will vary among incidents, organizations should select several solutions and determine in which order the solutions should be attempted.

A.4.2 Malicious Code Incidents

- + **Contain malicious code incidents as quickly as possible.** Because malicious code works surreptitiously and can propagate to other systems rapidly, early containment of a malicious code incident is needed to stop it from spreading and causing further damage. Infected systems should be disconnected from the network immediately. Organizations may need to block malicious code at the e-mail server level, or even temporarily suspend e-mail services to gain control over serious e-mail-borne malicious code incidents.

A.4.3 Unauthorized Access Incidents

- + **Provide change management information to the incident response team.** Indications such as system shutdowns, audit configuration changes, and executable modifications are probably caused by routine system administration, rather than attacks. When such indications are detected, the team should be able to use change management information to verify that the indications are caused by authorized activity.
- + **Select containment strategies that balance mitigating risks and maintaining services.** Incident handlers should consider moderate containment solutions that focus on mitigating the risks as much as is practical while maintaining unaffected services.
- + **Restore or reinstall systems that appear to have suffered a root compromise.** The effects of root compromises are often difficult to identify completely. The system should be restored from a known good backup, or the operating system and applications should be reinstalled from scratch. The system should then be secured properly so the incident cannot recur.

A.4.4 Multiple Component Incidents

- + **Contain the initial incident and then search for signs of other incident components.** It can take an extended period of time for a handler to authoritatively determine that an incident has only a single component; meanwhile, the initial incident has not been contained. It is generally better to contain the initial incident first.

A.5 Post-Incident Activity

- + **Hold lessons learned meetings after major incidents.** Lessons learned meetings are extremely helpful in improving security measures and the incident handling process itself.

A.5.1 Unauthorized Access Incidents

- + **Separately prioritize the handling of each incident component.** Resources are probably too limited to handle all incident components simultaneously. Components should be prioritized based on response guidelines for each component and how current each component is.

Appendix B—INCIDENT HANDLING SCENARIOS

Exercises involving incident handling scenarios provide an inexpensive and effective way to build incident response skills and identify potential issues with incident response processes. The incident response team or individual team members are presented with a brief scenario and a list of questions related to the scenario. The team then discusses each question and determines the most likely answer. The goal is to determine what the team would really do and to compare that with policies, procedures, and generally good practices to identify any discrepancies or deficiencies. For example, the answer to one question may indicate that the response would be delayed because the team lacks a particular piece of software or because another team within the organization does not provide off-hours support.

The questions listed below are applicable to nearly any scenario. Each question is followed by a reference to the related section(s) of the document. After the questions are scenarios, which each have additional incident-specific questions. Organizations are strongly encouraged to adapt these questions and scenarios for use in their own incident response exercises.

B.1 Scenario Questions

Preparation:

1. Would the organization consider this activity to be an incident? If so, which of the organization's policies does this activity violate? (*Section 2.1*)
2. What measures are in place to attempt to prevent this type of incident from occurring or to limit its impact? (*Section 3.1.2*)

Detection and Analysis:

3. What precursors of the incident, if any, might the organization detect? Would any precursors cause the organization to attempt to take action before the incident occurred? (*Sections 3.2.2, 3.2.3*)
4. What indications of the incident might the organization detect? Which indications would cause someone to think that an incident might have occurred? (*Sections 3.2.2, 3.2.3*)
5. How would the incident response team analyze and validate this incident? (*Section 3.2.4*)
6. To which people and groups within the organization would the team report the incident? (*Section 3.2.7*)
7. How would the incident response team prioritize the handling of this incident? (*Section 3.2.6*)

Containment, Eradication, and Recovery:

8. What strategy should the organization take to contain the incident? Why is this strategy preferable to others? (*Section 3.3.1*)
9. What could happen if the incident is not contained? (*Section 3.3.1*)

10. What sources of evidence, if any, should the organization acquire? How would the evidence be acquired? Where would it be stored? How long should it be retained? (*Sections 3.2.5, 3.3.2, 3.4.3*)

Post-Incident Activity:

11. Who would attend the lessons learned meeting regarding this incident? (*Section 3.4.1*)
12. What could be done to prevent similar incidents from occurring in the future? (*Section 3.1.2*)
13. What could be done to improve detection of similar incidents? (*Section 3.1.2*)

General Questions:

14. How many incident response team members would participate in handling this incident? (*Section 2.4.3*)
15. Besides the incident response team, what groups within the organization would be involved in handling this incident? (*Section 2.4.4*)
16. To which external parties would the team report the incident? When would each report occur? How would each report be made? (*Section 2.3.2*)
17. What other communications with external parties may occur? (*Section 2.3.2*)
18. What tools and resources would the team use in handling this incident? (*Section 3.1.1*)
19. What aspects of the handling would have been different if the incident had occurred at a different day and time (on-hours versus off-hours)? (*Section 2.4.2*)
20. What aspects of the handling would have been different if the incident had occurred at a different physical location (onsite versus offsite)? (*Section 2.4.2*)

B.2 Scenarios

Scenario 1: DNS Server Denial of Service

On a Saturday afternoon, external users start having problems accessing the organization's public Web sites. Over the next hour, the problem worsens to the point where nearly every attempt to access any of the organization's public Web sites fails. Meanwhile, a member of the organization's networking staff responds to automatically generated alerts from the Internet border router and determines that much of the organization's Internet bandwidth is being consumed by an unusually large volume of UDP packets to and from both of the organization's public DNS servers. An analysis of the traffic shows that the DNS servers are receiving high volumes of requests from a single external IP address. The networking administrator also notices that all of the DNS requests from that address have a source port of either UDP 7 or UDP 19. While this analysis is taking place, the organization's network intrusion detection sensors record suspicious activity related to the echo and chargen services.

The following are additional questions for this scenario:

1. Whom should the organization contact regarding the external IP address used in all of the packets?

2. Suppose that after the initial containment measures were put in place, the network administrators detected that nine internal hosts were attempting the same unusual requests to the DNS server. How would that affect the handling of this incident?
3. Suppose that two of the nine internal hosts left the network before their system owners were contacted. How would the system owners be identified?

Scenario 2: Internally Generated Spam

On a Thursday morning, the organization's "abuse" e-mail account receives a complaint from a person about receiving spam from the organization. The message contains a copy of a spam (with full e-mail headers) that is promoting a get-rich-quick scheme. A security administrator who monitors the abuse account reviews the e-mail and determines that the headers appear to show that the spam was generated using the organization's mail server. The security administrator forwards the e-mail to the incident response team's address, along with a brief note about the activity. A member of the incident response team analyzes the activity and confirms that the spam headers are genuine and that it was sent from the organization's mail server.

The following are additional questions for this scenario:

1. How would the incident response team validate the origin of the spam?
2. How would the organization respond to complaints regarding the spam?

Scenario 3: Worm and DDoS Agent Infestation

On a Tuesday morning, a new worm is released on the Internet. The worm exploits a Microsoft Windows vulnerability that was publicly announced 2 weeks before, at which time patches were released. The worm spreads itself through two methods: (1) e-mailing itself to all addresses that it can locate on an infected host and (2) identifying and sending itself to hosts with open Windows shares. The worm is designed to generate a different attachment name for each copy that it mails; each attachment has a randomly generated filename that uses one of over a dozen file extensions. The worm also chooses from more than 100 e-mail subjects and a similar number of e-mail bodies. When the worm infects a host, it gains administrative rights and attempts to download a DDoS agent from different IP addresses using FTP. (The number of IP addresses providing the agent is unknown.) Although the antivirus vendors quickly post warnings about the worm, it spreads very quickly, before any of the vendors have released signatures. The organization has already incurred widespread infections before antivirus signatures become available 3 hours after the worm started to spread.

The following are additional questions for this scenario:

1. How would the incident response team identify all infected hosts?
2. How would the organization attempt to prevent the worm from entering the organization before antivirus signatures were released?
3. How would the organization attempt to prevent the worm from being spread by infected hosts before antivirus signatures were released?
4. Would the organization attempt to patch all vulnerable machines? If so, how would this be done?

5. How would the handling of this incident change if infected hosts that had received the DDoS agent had been configured to attack another organization's Web site the next morning?
6. How would the incident response team keep the organization's users informed about the status of the incident? What if e-mail services were overloaded or unavailable due to the worm?
7. What additional measures, if any, would the team take to take care of hosts that are not currently connected to the network (e.g., staff members on vacation, off-site employees that dial in occasionally)?

Scenario 4: Use of Stolen Credit Card Numbers

On a Monday morning, the organization's legal department receives a call from the FBI regarding some suspicious activity originating from the organization's network. Later that day, an FBI agent meets with the organization's CIO and two members of the legal department to discuss the activity. The FBI has been investigating activity involving online purchases made with several stolen credit card numbers, and more than 30 of the transactions during the past week have been traced to one of the organization's IP addresses. The agent asks for the organization's assistance, and in turn, the legal department and the CIO ask for the incident response team's assistance in acquiring the necessary evidence. It is vitally important that this matter be kept confidential.

The following are additional questions for this scenario:

1. From what sources might the incident response team gather evidence?
2. How would the team identify which host is currently using the specified IP address? How would the team demonstrate which host had been using the specified IP address a week ago?
3. What would the team do to keep the investigation confidential?

Scenario 5: Compromised Database Server

On a Tuesday night, a database administrator performs some off-hours maintenance on several production database servers. The administrator notices some unfamiliar and unusual directory names on one of the servers. After reviewing the directory listings and viewing some of the files, the administrator concludes that the server has been attacked and calls the incident response team for assistance. The team's investigation determines that the attacker successfully gained root access to the server 6 weeks ago.

The following are additional questions for this scenario:

1. What sources might the team use to determine when the compromise had occurred?
2. How would the handling of this incident change if the team found that the database server had been running a packet sniffer and capturing passwords from the network?
3. How would the handling of this incident change if the team found that the server was running a process that would copy a database containing sensitive customer information each night and e-mail it to an external address?

Scenario 6: Virus Hoax

On a Wednesday afternoon, a user forwards an e-mail to the help desk about a terrible new virus. The user received the e-mail from a friend at another organization. The e-mail states that the new virus cannot be detected by antivirus software and that users should look for and delete three particular virus files from their hard drives. A help desk agent researches the message, determines that it is a virus hoax, and responds to the user by e-mail.

Meanwhile, other users receive the same virus warning e-mail from other outside parties and forward it to others inside and outside the organization. By Thursday afternoon, the help desk has received several calls that appear to be related to individuals deleting the three “virus files” from their hard drives; these files are actually legitimate files that certain applications use. The lead help desk agent asks for the incident response team’s assistance.

The following is an additional question for this scenario:

1. Would the organization proactively identify hosts that are missing the three files? If so, how would this be done? If not, what negative effect would this have?

Scenario 7: Unauthorized Materials on the FTP Server

While creating a weekly usage report, a network administrator notices that off-hours bandwidth utilization on the organization’s DMZ segment has been significantly higher than usual. The administrator configures network monitoring software to collect more detailed statistics on DMZ bandwidth usage. The next day, the administrator sees that an unusually large percentage of the activity involves the organization’s FTP server. The network administrator contacts the FTP server administrator, who has just returned from vacation, regarding the increase in activity. The FTP administrator quickly determines that the server is hosting unauthorized materials, which appear to include pirated software, songs, and movies. The administrator contacts the incident response team regarding the activity.

The following are additional questions for this scenario:

1. Would the team attempt to identify all individuals who had uploaded illegal materials to the FTP server? If so, how would this be done?
2. Would the team attempt to verify that the unauthorized materials were illegal? If so, how would this be done?

Scenario 8: Outbound DDoS Attack

On a Sunday night, the organization’s network intrusion detection sensor alerts on suspected outbound DDoS activity involving a high volume of ICMP pings. The intrusion analyst reviews the alerts; although the analyst cannot confirm that the alerts are accurate, they do not match any known false positives. The analyst contacts the incident response team so that it can investigate the activity further. Because the DDoS activity uses spoofed source IP addresses, it takes considerable time and effort to determine which host or hosts within the organization are producing it; meanwhile, the DDoS activity continues. The investigation shows that five servers appear to be generating the DDoS traffic. Analysis of the five servers shows that each contains signs of a DDoS rootkit. In addition, three of the servers appear to have been used to attack other internal hosts, and one appears to have been used for attacking external hosts as well.

The following are additional questions for this scenario:

1. How would the team determine which hosts within the organization were producing the traffic? Which other teams might assist the incident response team?
2. Would the organization contact the owners of the IP addresses that the DDoS attack had targeted? If so, who would contact them, and how would the contact be performed?
3. If the incident response team determined that the initial compromise had been performed through a modem in one of the servers, how would the team further investigate this activity?

Scenario 9: Unauthorized Access to Payroll Records

On a Wednesday evening, the organization's physical security team receives a call from a payroll administrator who caught an unknown person leaving her office. The administrator saw the person run down the hallway and enter a staircase that leads to a building exit. The administrator had left her workstation unlocked and unattended for just a few minutes. The payroll program is still logged in and on the main menu, as it was when she left it, but the administrator notices that the mouse appears to have been moved. The incident response team has been asked to acquire evidence related to the incident and to determine what actions were performed (e.g., payroll data access or modification, Trojan horse delivery).

The following are additional questions for this scenario:

1. How would the team determine what actions had been performed?
2. How would the handling of this incident differ if the payroll administrator had recognized the person leaving her office as a former payroll department employee?
3. How would the handling of this incident differ if the physical security team determined that the person had used social engineering techniques to gain physical access to the building and the payroll department?
4. How would the handling of this incident differ if the team had reason to believe that the person was a current employee?
5. How would the handling of this incident differ if remote access logs from the previous week showed an unusually large number of failed login attempts using the payroll administrator's user ID?

Scenario 10: Hacking Tool Download

On a Friday afternoon, a network intrusion detection sensor records some suspicious FTP activity involving an internal user downloading files from an external FTP server. The intrusion analyst reviews the alerts and notices that the alerts are false positives. Although the alerts indicate that an attack has occurred, the supporting data recorded by the sensor show no signs of an attack. However, the data raise other concerns, because they show that the user is downloading executables from a suspicious directory structure containing repeated spaces and periods, as well as characters that are not usually seen in FTP directory names. The intrusion analyst uses an Internet search engine to look for more information on the executable names, and several of them match the names of hacking tools. The analyst contacts the incident response team to perform further analysis and determine how this activity should be handled.

The following are additional questions for this scenario:

1. How would the team determine what files the user had downloaded?
2. How would the team confirm that the files that had been downloaded were hacking tools?
3. How would the handling of this incident differ if the user suspected of downloading the tools was a member of the organization's information security team?
4. How would the handling of this incident differ if the user suspected of downloading the tools was a member of the incident response team?
5. How would the handling of this incident differ if the user suspected of downloading the tools was a contractor that had just found out that his contract was not being renewed?

Scenario 11: Disappearing Host

On a Thursday afternoon, a network intrusion detection sensor records vulnerability scanning activity directed at internal hosts that is being generated by an internal IP address. Because the intrusion detection analyst is unaware of any authorized, scheduled vulnerability scanning activity, she reports the activity to the incident response team. When the team begins the analysis, it discovers that the activity has stopped and that there is no longer a host using the IP address.

The following are additional questions for this scenario:

1. What data sources might contain information regarding the identity of the vulnerability scanning host?
2. How would the team identify who had been performing the vulnerability scans?
3. How would the team confirm that the files that had been downloaded were hacking tools?
4. How would the handling of this incident differ if the vulnerability scanning were directed at the organization's most critical hosts?
5. How would the handling of this incident differ if the vulnerability scanning were directed at external hosts?
6. How would the handling of this incident differ if the physical security staff discovered that someone had broken into the facility half an hour before the vulnerability scanning occurred?

Scenario 12: Telecommuting Compromise

On a Saturday night, network intrusion detection software records some probes and scans originating from an internal IP address. Host intrusion detection software on a few servers also records some of the probes and scans. The intrusion detection analyst determines that the internal IP address belongs to the organization's VPN server and contacts the incident response team. The team reviews the intrusion detection, firewall, and VPN server logs and identifies the external IP address that is generating the activity, the user ID that was authenticated for the session, and the name of the user associated with the user ID.

The following are additional questions for this scenario:

1. What should the team's next step be (e.g., calling the user at home, disabling the user ID, disconnecting the VPN session)? Why should this step be performed first? What step should be performed second?
2. Suppose that the identified user's personal computer had become compromised by a game containing a Trojan horse that was downloaded by a family member. How would this affect the team's analysis of the incident? How would this affect evidence gathering and handling?
3. How would the handling of this incident differ if the external IP address belonged to the identified user, but the IP address was a firewall device performing Network Address Translation (NAT) for 8 hosts behind it?
4. What should the team do in terms of eradicating the incident from the user's personal computer?
5. Suppose that the user installed antivirus software and determined that the Trojan horse had included a keystroke logger. How would this affect the handling of the incident? How would this affect the handling of the incident if the user were a system administrator? How would this affect the handling of the incident if the user were a high-ranking executive in the organization?
6. How would the handling of this incident differ if the external IP address were not being used by the identified user?
7. How would the handling of this incident differ if the user reinstalled the operating system on the affected host before the team could perform any analysis on the host?

Appendix C—GLOSSARY

Selected terms used in the *Computer Security Incident Handling Guide* are defined below.

Agent: A program used in DDoS attacks that sends malicious traffic to hosts based on the instructions of a handler.

Baselining: Monitoring resources to determine typical utilization patterns so that significant deviations can be detected.

Blended Attack: Malicious code that uses multiple methods to spread.

Boot Sector Virus: A virus that plants itself in a system's boot sector and infects the master boot record.

Computer Forensics: The practice of gathering and retaining computer-related data in a manner that makes the data admissible in a court of law.

Computer Security Incident: See "incident."

Computer Security Incident Response Team (CSIRT): A capability set up for the purpose of assisting in responding to computer security-related incidents; also called a Computer Incident Response Team (CIRT) or a CIRC (Computer Incident Response Center, Computer Incident Response Capability).

Denial of Service (DoS): An attack that prevents or impairs the usage of network, system, or application resources.

Distributed Denial of Service (DDoS): A denial of service technique that uses numerous hosts.

Egress Filtering: The process of blocking outgoing packets that use obviously false IP addresses, such as source addresses from internal networks.

Event: Any observable occurrence in a network or system.

False Positive: An alert that incorrectly indicates that malicious activity is occurring.

File Infector Virus: A virus that attaches itself to a program file, such as a word processor, spreadsheet application, or game.

File Integrity Checker: Software that generates, stores, and compares message digests for files to detect changes to the files.

Forensics: See "computer forensics."

Handler: A type of program used in DDoS attacks to control agents distributed throughout a network. Also refers to an incident handler, which refers to a person that performs incident response work.

Honeypot: A host that is designed to collect data on suspicious activity and has no authorized users other than its administrators.

Inappropriate Usage: A user violates acceptable computing use policies.

Incident: A violation or imminent threat of violation of computer security policies, acceptable use policies, or standard computer security practices.

Incident Handling: The mitigation of violations of security policies and recommended practices.

Incident Response: See “incident handling.”

Indication: A sign that an incident may have occurred or may be currently occurring.

Ingress Filtering: The process of blocking incoming packets that use obviously false IP addresses, such as reserved source addresses.

Intrusion Detection System: Software that looks for suspicious activity and alerts administrators.

Macro Virus: A virus that attaches itself to documents and uses the macro programming capabilities of the document’s application to execute and propagate.

Malicious Code: A virus, worm, Trojan horse, or other code-based entity that infects a host, typically with malicious intent.

Message Digest: A cryptographic checksum, typically generated for a file, that can be used to detect changes to the file; SHA-1 is an example of a message digest algorithm.

Mobile Code: Software that is transmitted from a remote system to a local system, then executed on the local system without the user’s explicit instruction; examples of mobile code software are Java, JavaScript, VBScript, and ActiveX.

Multiple Component Incident: A single incident that encompasses two or more incidents.

Packet Sniffer: Software that observes and records network traffic.

Patch Management: The process of acquiring, testing, and distributing patches to the appropriate administrators and users throughout the organization.

Port Scanning: Using a program to remotely determine which ports on a system are open (e.g., whether systems allow connections through those ports).

Precursor: A sign that an attacker may be preparing to cause an incident.

Profiling: Measuring the characteristics of expected activity so that changes to it can be more easily identified.

Risk: The probability that one or more adverse events will occur.

Scanning: Sending packets or requests to another system to gain information to be used in a subsequent attack.

Signature: A recognizable, distinguishing pattern associated with an attack, such as a binary string in a virus or a particular set of keystrokes used to gain unauthorized access to a system.

Social Engineering: An attempt to trick someone into revealing information (e.g., a password) that can be used to attack systems or networks.

Threat: The potential source of an adverse event.

Trojan Horse: A non-self-replicating program that seems to have a useful purpose, but in reality has a different, malicious purpose.

Unauthorized Access: A user gains access without permission to a network, system, application, data, or other resource.

Victim: A machine that is attacked.

Virus: A self-replicating program that runs and spreads by modifying other programs or files.

Virus Hoax: An urgent warning message about a nonexistent virus.

Vulnerability: A weakness in a system, application, or network that is subject to exploitation or misuse.

Worm: A self-replicating, self-propagating, self-contained program that uses networking mechanisms to spread itself.

Appendix D—ACRONYMS

Selected acronyms used in the *Computer Security Incident Handling Guide* are defined below.

BIA	Business Impact Analysis
BIOS	Basic Input/Output System
CCIPS	Computer Crime and Intellectual Property Section
CERT[®]/CC	CERT [®] Coordination Center
CIAC	Computer Incident Advisory Capability
CIO	Chief Information Officer
CIRC	Computer Incident Response Center, Computer Incident Response Capability
CIRT	Computer Incident Response Team
CPU	Central Processing Unit
CSIRC	Computer Security Incident Response Capability
CSIRT	Computer Security Incident Response Team
CSO	Chief Security Officer
DHS	Department of Homeland Security
DMZ	Demilitarized Zone
DNS	Domain Name System
DOJ	Department of Justice
DoS	Denial of Service
DDoS	Distributed Denial of Service
FAQ	Frequently Asked Questions
FBI	Federal Bureau of Investigation
FedCIRC	Federal Computer Incident Response Center
FINE	Format for INcident Report Exchange
FIPS	Federal Information Processing Standards
FIRST	Forum of Incident Response and Security Teams
FISMA	Federal Information Security Management Act
FTP	File Transfer Protocol
GAO	General Accounting Office
HTTP	HyperText Transfer Protocol
HTTPS	HyperText Transfer Protocol Secure
IAIP	Information Analysis Infrastructure Protection
IANA	Internet Assigned Numbers Authority
ICMP	Internet Control Message Protocol
IDS	Intrusion Detection System
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPSec	IP Security Protocol
IRC	Internet Relay Chat
ISAKMP	Internet Security Association and Key Management Protocol
ISP	Internet Service Provider
IT	Information Technology

MBR	Master Boot Record
MSSP	Managed Security Services Provider
NAT	Network Address Translation
NIPC	National Infrastructure Protection Center
NIST	National Institute of Standards and Technology
NTP	Network Time Protocol
OCTAVE	Operationally Critical Threat, Asset, and Vulnerability Evaluation
OIG	Office of Inspector General
OMB	Office of Management and Budget
OS	Operating System
PBX	Private Branch Exchange
PDA	Personal Digital Assistant
PGP	Pretty Good Privacy
P.L.	Public Law
POC	Point of Contact
RFC	Request for Comments
SHA	Secure Hash Algorithm
SLA	Service Level Agreement
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SOP	Standard Operating Procedure
SQL	Structured Query Language
SSH	Secure Shell
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TERENA	Trans-European Research and Education Networking Association
UDP	User Datagram Protocol
URL	Uniform Resource Locator
VPN	Virtual Private Network

Appendix E—PRINT RESOURCES

- Allen, Julia H, *The CERT® Guide to System and Network Security Practices*, Addison-Wesley, 2001.
- Bace, Rebecca Gurley, *Intrusion Detection*, Que, 1999.
- Casey, Eoghan, *Digital Evidence and Computer Crime*, Academic Press, 2000.
- Kruse II, Warren and Heiser, Jay, *Computer Forensics: Incident Response Essentials*, Addison-Wesley, 2001.
- Mandia, Kevin and Prosis, Chris, *Incident Response: Investigating Computer Crime*, McGraw-Hill, 2001.
- Marcella Jr., Albert J. (Editor) and Greenfield, Robert S. (Editor), *Cyber Forensics: A Field Manual for Collecting, Examining, and Preserving Evidence of Computer Crimes*, Auerbach Publications, 2002.
- Northcutt, Stephen, *Computer Security Incident Handling: Step-by-Step*, SANS Press, 2003.
- Northcutt, Stephen and Novak, Judy, *Network Intrusion Detection (Third Edition)*, Que, 2002.
- Northcutt, Stephen, et al., *Inside Network Perimeter Security: The Definitive Guide to Firewalls, Virtual Private Networks (VPNs), Routers, and Intrusion Detection Systems*, Que, 2002.
- Schiffman, Mike, *Hacker's Challenge: Test Your Incident Response Skills Using 20 Scenarios*, McGraw-Hill, 2001.
- Schiffman, Mike, et al., *Hacker's Challenge 2: Test Your Network Security and Forensic Skills*, McGraw-Hill, 2002.
- Schultz, E. Eugene and Shumway, Russell, *Incident Response: A Strategic Guide to Handling System and Network Security Breaches*, Que, 2002.
- Schweitzer, Douglas, *Incident Response: Computer Forensics Toolkit*, John Wiley and Sons, 2003.
- van Wyk, Kenneth and Forno, Richard, *Incident Response*, O'Reilly and Associates, 2001.

Appendix F—ONLINE TOOLS AND RESOURCES

The lists below provide examples of online tools and resources that may be helpful in establishing and maintaining an incident response capability.

Incident Response Organizations

Organization	URL
AusCERT—Australian Computer Emergency Response Team	http://www.auscert.org.au
CCIPS—Computer Crime and Intellectual Property Section, U.S. Department of Justice	http://www.cybercrime.gov/
CERT [®] /CC—CERT [®] Coordination Center, Carnegie Mellon University	http://www.cert.org
CERT [®] /CC Incident Reporting System	https://irf.cc.cert.org
CIAC—Computer Incident Advisory Capability, U.S. Department of Energy	http://www.ciac.org/ciac
DOD-CERT—U.S. Department of Defense Computer Emergency Response Team	http://www.cert.mil
FedCIRC—Federal Computer Incident Response Center	http://www.fedcirc.gov
FedCIRC Incident Reporting System	https://incidentreport.fedcirc.gov
FIRST—Forum of Incident Response and Security Teams	http://www.first.org
HTCIA—High Technology Crime Investigation Association	http://www.htcia.org
IAIP—Information Analysis Infrastructure Protection, U.S. Department of Homeland Security	http://www.nipc.gov
IAIP Incident Report form	http://www.nipc.gov/incident/cirr.htm
IETF Extended Incident Handling (inch) Working Group	http://www.ietf.org/html.charters/inch-charter.html
InfraGard	http://www.infragard.net
ISC—Internet Storm Center	http://isc.incidents.org

Incident Response-Related Mailing Lists

Mailing List Name	Archive Location
Bugtraq	http://www.securityfocus.com/archive/1
Dshield	http://www.dshield.org/pipermail/list
Focus-IDS	http://www.securityfocus.com/archive/96
Forensics	http://www.securityfocus.com/archive/104
Incidents	http://www.securityfocus.com/archive/75
Intrusions	http://cert.uni-stuttgart.de/archive/intrusions
LogAnalysis	http://sisyphus.iocaine.com/pipermail/loganalysis

Technical Resource Sites

Resource Name	URL
CHIHT (Clearing House for Incident Handling Tools)	http://chiht.dfn-cert.de
COAST (Computer Operations, Audit and Security Technology) Intrusion Detection Pages	http://www.cerias.purdue.edu/coast/intrusion-detection/welcome.html
CSIRT Development, CERT [®] /CC	http://www.cert.org/csirts
CSRC—Computer Security Resource Center, National Institute of Standards and Technology	http://csrc.nist.gov
DShield (Distributed Intrusion Detection System)	http://www.dshield.org
Incident Handling Links and Documents	http://www.honeypots.net/incidents/links
Intrusion Detection FAQ, SANS Institute	http://www.sans.org/resources/idfaq
Intrusion Detection Links and Documents	http://www.honeypots.net/ids/links
Loganalysis.org	http://www.loganalysis.org
SANS Institute Reading Room	http://www.sans.org/rr
SecurityFocus	http://www.securityfocus.com
The Electronic Evidence Information Center	http://www.e-evidence.info

Vulnerability and Exploit Information Resources

Resource Name	URL
CERT [®] /CC Advisories	http://www.cert.org/advisories
CERT [®] /CC Incident Notes	http://www.cert.org/incident_notes
CERT [®] /CC Vulnerability Notes Database	http://www.kb.cert.org/vuls
CIAC Bulletins and Advisories	http://www.ciac.org/cgi-bin/index/bulletins
Common Vulnerabilities and Exposures (CVE)	http://www.cve.mitre.org
ICAT Vulnerability Metabase	http://icat.nist.gov
Information Analysis Infrastructure Protection (IAIP)	http://www.nipic.gov
Packet Storm	http://www.packetstormsecurity.com
SANS/FBI Top 20 List	http://www.sans.org/top20
SecurityFocus Vulnerabilities Database	http://www.securityfocus.com/bid

Training Resources

Training Resource Name	Types of Training	URL
CERT [®] /CC	Incident response	http://www.cert.org/training
Computer Forensic Services	Computer forensics	http://www.computer-forensic.com/training.html
Foundstone	Incident response, computer forensics	http://www.foundstone.com
MIS Training Institute (MISTI)	Incident response, intrusion detection, computer forensics	http://www.misti.com
SANS Institute	Incident response, intrusion detection, computer forensics	http://www.sans.org

NIST Technical Resource Documents

Resource Name	URL
NIST Special Publication 800-28—Guidelines on Active Content and Mobile Code	http://csrc.nist.gov/publications/nistpubs/800-28/sp800-28.pdf
NIST Special Publication 800-30—Risk Management Guide for Information Technology Systems	http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf
NIST Special Publication 800-31—Intrusion Detection Systems	http://csrc.nist.gov/publications/nistpubs/800-31/sp800-31.pdf
NIST Special Publication 800-36—DRAFT Guide to Selecting IT Security Products	http://csrc.nist.gov/publications/drafts/Products_PC_100802.pdf
NIST Special Publication 800-40—Procedures for Handling Security Patches	http://csrc.nist.gov/publications/nistpubs/800-40/sp800-40.pdf
NIST Special Publication 800-41—Guidelines on Firewalls and Firewall Policy	http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf
NIST Special Publication 800-42—DRAFT Guideline on Network Security Testing	http://csrc.nist.gov/publications/drafts/security-testing.pdf
NIST Special Publication 800-43—Systems Administration Guidance for Windows 2000 Professional	http://csrc.nist.gov/itsec/guidance_W2Kpro.html
NIST Special Publication 800-44—Guidelines on Securing Public Web Servers	http://csrc.nist.gov/publications/nistpubs/800-44/sp800-44.pdf
NIST Special Publication 800-45—Guidelines on Electronic Mail Security	http://csrc.nist.gov/publications/nistpubs/800-45/sp800-45.pdf
NIST Special Publication 800-48—Wireless Network Security: 802.11, Bluetooth, and Handheld Devices	http://csrc.nist.gov/publications/nistpubs/800-48/NIST_SP_800-48.pdf

Other Technical Resource Documents

Resource Name	URL
CERT [®] /CC Security Improvement Modules	http://www.cert.org/security-improvement
Computer Security Incident Response Team (CSIRT) Frequently Asked Questions (FAQ)	http://www.cert.org/csirts/csirt_faq.html
Distributed Denial of Service Defense Tactics	http://razor.bindview.com/publish/papers/DDSA_Defense.html
FAQ: Network Intrusion Detection Systems, Robert Graham	http://www.robertgraham.com/pubs/network-intrusion-detection.html
Handbook for Computer Security Incident Response Teams (CSIRTs)	http://www.cert.org/archive/pdf/csirt-handbook.pdf
How to Design a Useful Incident Response Policy	http://www.securityfocus.com/infocus/1467
Incident Response: Managing Security at Microsoft	http://www.microsoft.com/technet/treeview/default.asp?url=/technet/itsolutions/msit/security/msirsec.asp
Managing the Threat of Denial-of-Service Attacks	http://www.cert.org/archive/pdf/Managing_DoS.pdf
Responding to Intrusions	http://www.sei.cmu.edu/pub/documents/sims/pdf/sim006.pdf
RFC 2350: Expectations for Computer Security Incident Response	http://www.ietf.org/rfc/rfc2350.txt
RFC 3067: TERENA's Incident Object Description and Exchange Format Requirements	http://www.ietf.org/rfc/rfc3067.txt

Resource Name	URL
RFC 3227: Guidelines for Evidence Collection and Archiving	http://www.ietf.org/rfc/rfc3227.txt
Sample Incident Handling Forms, SANS Institute	http://www.sans.org/incidentforms
System Security	http://www.linux-mag.com/2000-10/security_01.html
A Taxonomy of DDoS Attacks and DDoS Defense Mechanisms	http://lasr.cs.ucla.edu/ddos/ucla_tech_report_020018.pdf

Knowledge Base Resources

Resource Name	URL
IANA Protocol Numbers and Assignment Services	http://www.iana.org/numbers.html
Domain Name System Parameters	http://www.iana.org/assignments/dns-parameters
ICMP Type Numbers	http://www.iana.org/assignments/icmp-parameters
Internet Multicast Addresses	http://www.iana.org/assignments/multicast-addresses
Internet Protocol V4 Address Space	http://www.iana.org/assignments/ipv4-address-space
IP Protocol Numbers	http://www.iana.org/assignments/protocol-numbers
IP Version Numbers	http://www.iana.org/assignments/version-numbers
Port Numbers	http://www.iana.org/assignments/port-numbers
Syslog Parameters	http://www.iana.org/assignments/syslog-parameters
TCP Header Flags	http://www.iana.org/assignments/tcp-header-flags
TCP Option Numbers	http://www.iana.org/assignments/tcp-parameters
IETF RFCs for Common Protocols (DNS, FTP, HTTP and SMTP)	http://www.ietf.org/rfc.html
RFC 959: File Transfer Protocol (FTP)	http://www.ietf.org/rfc/rfc0959.txt
RFC 1034: Domain Names—Concepts and Facilities	http://www.ietf.org/rfc/rfc1034.txt
RFC 1035: Domain Names—Implementation and Specification	http://www.ietf.org/rfc/rfc1035.txt
RFC 2065: Domain Name System Security Extensions	http://www.ietf.org/rfc/rfc2065.txt
RFC 2228: FTP Security Extensions	http://www.ietf.org/rfc/rfc2228.txt
RFC 2616: Hypertext Transfer Protocol—HTTP/1.1	http://www.ietf.org/rfc/rfc2616.txt
RFC 2617: HTTP Authentication: Basic and Digest Access Authentication	http://www.ietf.org/rfc/rfc2617.txt
RFC 2821: Simple Mail Transfer Protocol	http://www.ietf.org/rfc/rfc2821.txt
RFC 2822: Internet Message Format	http://www.ietf.org/rfc/rfc2822.txt
RFC 2965: HTTP State Management Mechanism	http://www.ietf.org/rfc/rfc2965.txt
Neohapsis Ports List	http://www.neohapsis.com/neolabs/neo-ports/
Network ICE Port Knowledgebase	http://www.iss.net/security_center/advice/Exploits/Ports/
The Internet Ports Database	http://www.portsdb.org

Free and Open Source Incident Detection and Analysis Software

Software Name	Software Type	URL
Chkrootkit	Rootkit detection	http://www.chkrootkit.org
Ethereal	Protocol analyzer	http://www.ethereal.com
IPTraf	Network statistics utility	http://iptraf.seul.org
Snort	Network-based intrusion detection	http://www.snort.org
Swatch	Log file monitoring	http://swatch.sourceforge.net
TCPDump	Packet sniffer	http://www.tcpdump.org
TCPTrace	Tcpdump file analysis	http://www.tcptrace.org/index.html
The Coroner's Toolkit	Computer forensics	http://www.porcupine.org/forensics/tct.html
The Sleuth Kit	Computer forensics	http://www.sleuthkit.org/sleuthkit/index.php
WinDump	Packet sniffer	http://windump.polito.it

Appendix G—FREQUENTLY ASKED QUESTIONS

Users, system administrators, information security staff members, and others within organizations may have questions about incident response. The following are frequently asked questions (FAQ) regarding incident response. Organizations are encouraged to customize this FAQ and make it available to their user community.

1. What is an incident?

In general, an incident is a violation of computer security policies, acceptable use policies, or standard computer security practices. Examples of incidents are—

- + A distributed denial of service attack against a public Web server
- + A worm that infects hundreds of workstations on a network and effectively shuts the network down
- + An attacker who gains remote administrator-level access to an e-mail server
- + A user who downloads password cracking tools
- + A user who defaces another organization’s public Web site.

2. What is incident handling?

Incident handling is the process of detecting and analyzing incidents and limiting the incident’s effect. For example, if an attacker breaks into a system through the Internet, the incident handling process should detect the security breach. Incident handlers will then analyze the data and determine how serious the attack is. The incident will be prioritized, and the incident handlers will take action to ensure that the progress of the incident is halted and that the affected systems return to normal operation as soon as possible.

3. What is incident response?

The terms “incident handling” and “incident response” are synonymous in this document.⁸³

4. What is an incident response team?

An incident response team (also known as a Computer Security Incident Response Team [CSIRT]) is responsible for providing incident response services to part or all of an organization. The team receives information on possible incidents, investigates them, and takes action to ensure that the damage caused by the incidents is minimized. In some organizations, the incident response team is a formalized, full-time group; in others, incident response team members are pulled from other job functions as needed. Some organizations have no incident response team because they outsource incident response duties.

⁸³ Definitions of “incident handling” and “incident response” vary widely. For example, CERT®/CC uses “incident handling” to refer to the overall process of incident detection, reporting, analysis, and response, whereas “incident response” refers specifically to incident containment, recovery, and notification of others. See http://www.cert.org/csirts/csirt_faq.html for more information.

5. What services does the incident response team provide?

The particular services that incident response teams offer vary widely among organizations. Besides performing incident handling, a team typically distributes advisories regarding new threats, and educates and raises the awareness of users and technical staff on their roles in incident prevention and handling. Many teams also assume responsibility for intrusion detection system monitoring and management. Some teams perform additional services, such as auditing and penetration testing.

6. To whom should incidents be reported?

Organizations should establish clear POCs for reporting incidents internally. Some organizations will structure their incident response capability so that all incidents are reported directly to the incident response team, whereas others will use existing support structures, such as the IT help desk, for an initial point of contact. The organization should recognize that external parties, such as other incident response teams, would report some incidents. Federal agencies are required under the law to report all incidents to the Federal Computer Incident Response Center (FedCIRC).

7. How are incidents reported?

Most organizations have multiple methods for reporting an incident. Different reporting methods may be preferable as a result of variations in the skills of the person reporting the activity, the urgency of the incident, and the sensitivity of the incident. A phone number or pager number should be established to report emergencies. An e-mail address may be provided for informal incident reporting, whereas a Web-based form may be useful in formal incident reporting. Sensitive information can be provided to the team by sending a fax to a machine in a secured area, or by using a public key published by the team to encrypt the material.

8. What information should be provided when reporting an incident?

The more precise the information is, the better. For example, if a workstation appears to have been infected by malicious code, the incident report should include the following data:

- + The user's name, user ID, and contact information (e.g., phone number, e-mail address)
- + The workstation's location, model number, serial number, hostname, and IP address
- + The date and time that the incident occurred
- + A step-by-step explanation of what happened, including what was done to the workstation after the infection was discovered. This explanation should be detailed, including the exact wording of messages, such as those displayed by the malicious code or by antivirus software alerts.

9. How quickly does the incident response team respond to an incident report?

The response time depends on several factors, such as the type of incident, the criticality of the resources and data that are affected, the severity of the incident, existing Service Level Agreements (SLA) for affected resources, the time and day of the week, and other incidents that the team is handling. Generally, the highest priority is handling incidents that are likely to cause the most damage to the organization or to other organizations.

10. When should a person involved with an incident contact law enforcement?

Communications with law enforcement agencies should be initiated by the incident response team members, the Chief Information Officer (CIO) or other designated official—users, system administrators, system owners, and other involved parties should not initiate contact.

11. What should someone do who discovers that a system has been attacked?

The person should immediately stop using the system and contact the incident response team. The person may need to assist in the initial handling of the incident—for instance, disconnecting the network cable from the attacked system or physically monitoring the system until incident handlers arrive to protect evidence on the system.

12. What should someone do who receives spam?

The person should forward the spam to the e-mail address that the organization has designated for reporting spam. Statistics compiled on spam may be used to justify additional antispam measures. The statistics will also be provided to incident reporting organizations that study trends in computer security incidents. The person usually should not reply to the spam message in any way, including asking to be removed from a mailing list, as this would validate to the sender that the e-mail address is valid and actively used.

13. What should someone do who receives a warning from a friend about a new virus?

The person should check a virus hoax Web site to see if the new virus is legitimate or a hoax. Many of the virus warnings distributed through e-mail are hoaxes, and some of the instructions provided in the hoaxes may cause damage to systems if they are followed. Antivirus vendor Web sites often contain virus hoax information; the Computer Incident Advisory Capability (CIAC) Hoaxbusters site is another good source.⁸⁴ A person who is still in doubt about the authenticity of a virus warning should contact the help desk for further assistance.

14. What should someone do who is contacted by the media regarding an incident?

A person who has been part of the incident response may answer the media's questions in accordance with the organization's policy regarding incidents and outside parties. If the person is not qualified to represent the organization in terms of discussing the incident, the person should make no comment regarding the incident, other than to refer the caller to the organization's public affairs office. This will allow the public affairs office to provide accurate and consistent information to the media and the public.

⁸⁴ <http://hoaxbusters.ciac.org>

Appendix H—CRISIS HANDLING STEPS

Appendix H lists the major steps that should be performed when a technical professional believes that a serious incident has occurred and the organization does not have an incident response capability available. This serves as a basic reference of what to do for someone who is faced with a crisis and does not have time to read through this entire document.

1. **Document everything.** This includes every action that is performed, every piece of evidence, every conversation with users, system owners, and others regarding the incident.
2. **Find a coworker who can provide assistance.** Handling the incident will be much easier if two or more people work together. For example, one can perform actions while the second documents them.
3. **Analyze the evidence to confirm that an incident has occurred.** Perform additional research as necessary (e.g., Internet search engines, software documentation) to better understand the evidence. Reach out to other technical professionals within the organization for additional help.
4. **Notify the appropriate people within the organization** if it appears that an incident has occurred. This should include the CIO, the head of information security, and the local security manager. Use discretion when discussing details of the incident with others; tell only the people who need to know and use communication mechanisms that are reasonably secure. (If the attacker has compromised e-mail services, do not send e-mails about the incident.)
5. **Stop the incident if it is still in progress.** The most common way to do this is to disconnect affected systems from the network. In some cases, firewall and router configurations may need to be modified to stop network traffic that is part of an incident, such as a DoS attack.
6. **Preserve evidence from the incident.** Make backups (preferably disk image backups, not file system backups) of affected systems. Make copies of log files that contain evidence related to the incident.
7. **Wipe out all effects of the incident.** This includes malicious code infections, inappropriate materials (e.g., pirated software), Trojan horse files, and any other changes made to systems by incidents. If a system has been fully compromised, rebuild it from scratch or restore it from a known good backup.
8. **Identify and mitigate all vulnerabilities that were exploited.** The incident probably occurred by taking advantage of vulnerabilities in operating systems or applications. It is critical to identify such vulnerabilities and eliminate or otherwise mitigate them so that the incident does not recur.
9. **Confirm that operations have been restored to normal.** Make sure that data, applications, and other services affected by the incident have been returned to normal operations.
10. **Create a final report.** This report should detail the incident handling process, and provide an executive summary of what happened and how a formal incident response capability would have helped to handle the situation, mitigate the risk, and limit the damage more quickly.