



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Incident Management 101
Preparation & Initial Response (aka Identification)

GIAC Security Essentials
Certification (GSEC)
Practical Assignment
Version 1.4b

Option 1 - Research on Topics
in Information Security

Submitted by: Robin C. Dickerson on 09/07/04
Location: Denver, CO 2004

Paper Abstract:

According to SANS, there are six steps involved in properly handling a computer incident: Preparation, Identification, Containment, Eradication, Recovery, and Lessons Learned. Incident Management 101 provides guidelines, procedures, and tools designed to assist security specialists with the first two phases of Incident Management – Preparation and Initial Response (aka Identification phase). The intended audience is for incident handlers who are responding to suspicious activity (versus malicious code or DOS attacks) on both Unix and Windows systems. The guidelines, procedures and tools described are intended for business recovery, not for legal purposes such as preservation of evidence, forensic analysis, or prosecution.

Table of Contents

Paper Abstract:	1
Introduction	1
Section One	1
Preparation	1
Section Two	2
Initial Response.....	2
Conclusion	14
References.....	14

© SANS Institute 2005, Author retains full rights.

Introduction

Responding to hacking activity on computer systems is a stressful and critical responsibility. Many security specialists are expected to be able to monitor, respond, evaluate, advise, and mitigate suspicious activity on business critical systems. During an incident, security specialists are watched, evaluated, and are expected to be able to lead. You do not want to be ill prepared. This paper is intended to assist security specialists before an incident by providing instructions and guidelines on preparation, policy and training. It then covers initial response techniques such as creating a logbook, verifying an incident has occurred, conducting analysis, tools and tips, scope and impact, and next steps. By following these guidelines and customizing them for one's own unique environment, Incident Management can be an exciting experience.

Section One

Preparation

Preparing for an information system security incident begins with research, which comes long before being ready to react to alarms and events. First, determine if there is an established policy that details how security incidents are to be handled. The policy should identify who has authorization to conduct interviews, make requests, review sensitive data, and coordinate communications. For example, Human Resources and Legal may be the only departments authorized to respond to system misuse such as employees surfing pornographic website. Additionally, it should also contain of list of threats the organization intends to guard against and respond to. All stakeholders then need to become familiar with the policy and their roles and responsibilities. "Policies and supporting procedures that are documented, communicated, tested and enforced prepare you to respond to intrusions in a timely, controlled manner. They eliminate potential errors or omissions in advance of an intrusion," CERT paper titled "Establish policies and procedures for responding to intrusions." These policies should describe what actions the responder is allowed to take, such as when to take the impacted system offline, when to simply deny access to the intruder, under what circumstances the authorities should be engaged, when to contain, remove and restore the system, or when to simply continue monitoring for additional information (CERT. "Establish policies and procedures for responding to intrusions." p. 1). Lastly, the policy should state who needs to be notified, in what manner, and how often. Once the rules of engagement are understood, the investigation can proceed in an authorized manner.

What is an incident? CERT's definition is "The act of violating an explicit or implied security policy." SANS definition is "An adverse event in an information system, and/or network, or the threat of the occurrence of such an event. Incident

implies harm, or the intent to do harm". Examples of types of threats that may result in the Incident Management process being invoked may include:

- Attempts (either failed or successful) to gain unauthorized access to a system or its data.
 - Changes to system hardware, firmware, or software characteristics without the owner's knowledge, instruction, or consent.
 - Vandalism.
 - Theft of information.
 - Responding to high-rated Intrusion Detection Systems (IDS) alarms.
 - Responding to a customer or internal source identifying unusual system behavior.
 - Unwanted disruption or denial of service (not to be covered in this paper).
 - Virus outbreaks (not to be covered in this paper).
 - The unauthorized use of a system for the processing or storage of data.
- (Mandia, p.23)

Organizations and security specialists should be able to detect and respond to each of these types of threats and more. The best way to do this is by having documented procedures, education and testing. Conduct incident response tests on a periodic basis to identify weaknesses in the procedures and run the tools that will be used during a live investigation. Make sure your systems are logging and that people understand their roles and responsibilities. As with most things in life, practice makes perfect.

Section Two

Initial Response

The goals of initial response are to:

- Verify an incident has truly occurred.
- Determine what attacks were used to gain access.
- Identify which systems and data were accessed by the intruder.
- Determine what an intruder did after obtaining access (CERT. "Analyze all available information to characterize an intrusion." P. 1).

Disclaimer: Ideally, one should not conduct an investigation on a live system for many reasons. The primary reason is it may tip off a hacker that their activity has been detected. Some hackers booby trap the system and cause more harm or render a system unusable in an attempt to cover their tracks. Another reason not to conduct analysis on live machines is you may change the settings or tamper with the data in such a way that it can no longer be used as evidence in the event of prosecution. All security responders need to know when to proceed with an initial investigation and when to call the authorities. This information is typically communicated in an organization's security policy. That being said, this paper does not address how to conduct a full forensic analysis or preservation of evidence. This paper does however provide guidelines, procedures, and tools on reviewing information that will help a security specialist confirm an incident has occurred.

Figure 1.0 is a graphical representation of the procedural steps described in this document for conducting an Initial Response to a system security incident.

© SANS Institute 2005, Author retains full rights.

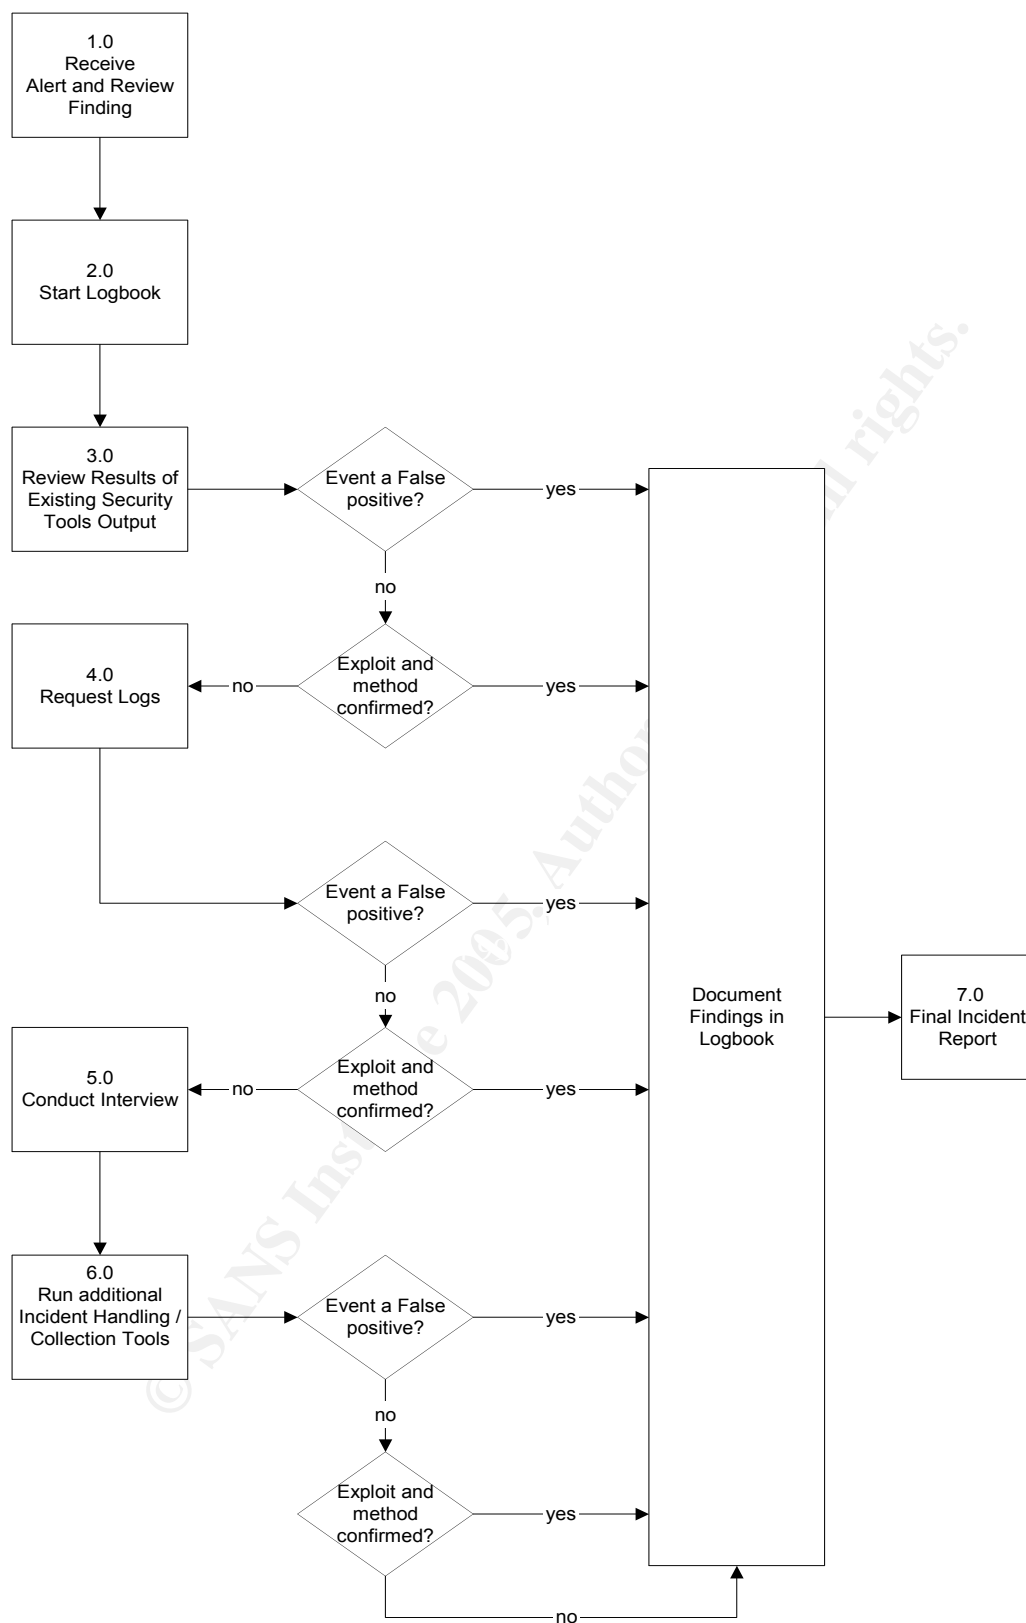


Figure 1.0 Initial Response flowchart

1.0 Receive Alert and Review Finding

You are notified of a suspected security event. Begin by reviewing the details of the event. What type of alarm is it? If an automatic tool has alerted you, become familiar with the event and determine what it is triggering on. For example, if it is a Network or Host-based Intrusion Detection System (NIDS or HIDS) event, you may want to check the vendor's website to get a description of what the signature is; or perhaps there is a central repository or documentation that describes how the alarm is triggered and what the significance is. You may even be able to tell at this point from the description of the alarm that it is not applicable to your environment. For instance, if you receive an Apache vulnerability, but you are only running Microsoft's IIS, this can be ruled out as a false positive. Know what you are dealing with and know your environment.

2.0 Start Logbook

The next step is to start a logbook. The logbook is used to document *everything*; all people interviewed, what happened, which systems were involved, what action was taken, what tools and commands were used, and what the results were. The first entry in your logbook should include an Incident Notification Checklist:

- Who is calling/paging?
 - Date/time
 - Phone
- Nature of Incident (virus, DOS, defacement, theft, unusual activity)
- When did the incident occur?
- How was the incident detected?
- Who discovered the incident?
- When was the incident detected?
- What is the immediate and future impact to client?
 - Is it a business critical machine?
- Targeted computer(s):
 - Hostname
 - OS
 - IP address(s)
 - Location
- Attacking computer(s):
 - IP address(s)
- What security precautions are currently taken on the system?
 - Real time intrusion detection, tcp wrappers, Tripwire, vulnerability scans conducted regularly? (Mandia, p.77)

Before you delve into an initial investigation, check to see if there are any authorized activities happening on the machine such as scheduled changes, tools, or scans. There is nothing more frustrating than finding out three hours into an investigation that the activity was planned.

3.0 Review Results of Existing Security Tools Output

Determine if your organization runs any security checking applications or tools on a regular basis. See if there were any recent reports or scans ran within the last few days or week. For example, if you conduct vulnerability scans or use security health checking tools on a regular basis, check the results for the server in question. Review the security tool's findings to see if the vulnerability that is being triggered exists or is even applicable. For example, if you see an IIS web attack AND your security-checking tool checks for this vulnerability, you can possibly assume the server is not vulnerable. Another example may be that someone is trying to log in as root. Perhaps the security checking tools that run regularly can confirm that that remote root login has been disabled and tested, thus you can again identify this as an unsuccessful event. Or worse, you find in your security-tool results that the system does have the vulnerability that is being exploited. In either case, start with what is readily available to help determine the system and security health of your system, if such data exists.

Document any findings you uncover from reviewing the latest security-tool results. In the logbook, note what tools or application findings were reviewed. Note the date of the last report or scan. Include any relevant information that would cause you to either conclude it is a false positive, non-issue on the system because it is protected, or result that shows a vulnerability exists on the device in question.

4.0 Request Logs

Attackers usually leave trace information on the systems they are attempting to access or have successfully gained access to. Trace information can often be found in log and audit files. This trace information can be used to search for source IPs of the offending system or other events or connections originating from that system that were previously unnoticed (CERT. "Analyze all available information to characterize an intrusion" p.3).

Start with a broad search and as you gather more evidence, start to narrow the search down. If you have real-time IDS (Network or Host-based Intrusion Detection Sensors), you already have a good beginning point. The IDS logs identify suspicious behavior coming from IP addresses. Use past IDS logs to look for reconnaissance activity from the offender. Intruders typically attempt a number of different attacks or scans before they are successful, so looking for historical data may be a good start (CERT. "Analyze all available information to characterize an intrusion" p.3). Also the IDS logs can be used to narrow down which logs you will want to look at from the compromised system or network logs based on dates, times, source IPs, and keywords.

Server Logs: Windows Systems

Windows systems have three types of logs: System, Application, and Security logs.

“System logs capture information such as system processes and device driver activities... Activities related to user programs and commercial off-the-shelf applications populate the Application log... System auditing and the security processes used by NT are found in the Security log. Security events that are audited by Windows systems include:

- Changes in user privileges
- Changes in the audit policy
- File and directory access
- Printer activity
- System logins and logouts” (Mandia, p.260)

Tip: for Security Specialists that do not have personal Unix workstations to conduct reviews of large log files, one recommended tool is UltraEdit. UltraEdit has powerful Unix grep-like search capabilities. This tool can save hours of work and your eyesight.

During an investigation, the Windows Security Logs will probably be the best resource for information. Note that only the system administrator can access the Security Logs, which can be found by going to Start, Programs, Administrative Tools, Event Viewer. It is in the Event Viewer that all three logs are housed. In the Security Logs, look at the event IDs (a number in the event column); each number represents a specific type of system event. For example, 624 means a new account was added, 517 means the audit log was cleared. The Incident Response: Investigating Computer Crime book recommends filtering on event ID 624 (addition of a new account, 626 (user account enabled), 636 (changing an account group) and 642 (user account changed) (Mandia, p.262). For a list of Security Event IDs, please visit the Microsoft website or visit <http://is-it-true.org/nt/atips/atips155.shtml>.

Another good source of information is in the Microsoft Registry. It is here that the security specialist and system administrator can identify software and applications that were installed on a system and then were manually deleted (Mandia, p.278). Look for unusual or suspicious looking applications such as password cracking tools, sniffer programs, keyboard capturing tools, etc.

Server Logs: Unix Systems

Unix systems have a myriad of logs that store powerful information such as logons, startups, and shutdowns (Mandia, p.322); all of which are probably in a different location depending on the flavor of Unix. For obtaining and viewing Unix logs, this paper will only give some very general information. Each flavor of Unix has its own standards and commands. It is advised that you request the help of the system administrator/or consult the man pages (Unix help pages) to determine exactly where each type of log is stored. “Most Unix flavors keep their log files in /var/adm or /var/log subdirectories,” according to the Incident Response: Investigating Computer Crime book. The Incident Response: Investigating Computer Crime book recommends reviewing the following logs:

Log	Description
Su Command	Shows when a user switches to another userID. Sometimes used to access the root account. Hackers often try to elevate their system privileges.
Logged-on User – utmp and wtmp	Provides information about who is currently logged in. Often can be found by executing the w, who, command.
Logon Attempts	Successful and unsuccessful logon attempts.
Cron Logs	Scheduled programs for future execution.
Lastlog	Shows the login-name, port, and last login time

For Unix systems, use the grep and find command to do key word or string searches in the logs.

Web logs

When dealing with attacks on a webserver/website, the web logs will provide the most fruitful information. At the time of writing this paper, the two most popular web applications are Microsoft Windows Internet Information Service (IIS) and Apache. It is these logs that will help solidify or dispel if an incident has truly occurred.

IIS logs, if stored in their default location, can be found in the C:\WINNT\System32\LogFiles\W3SVC1 directory. New IIS logs are created daily. (Mandia, p.371)

Apache logs, if stored in their default location, can be found in the /usr/local/apache/logs location. The specific log you want is the access_log. (Mandia, p.372)

In reviewing these logs, use either the grep command or a control F on Windows systems (or UltraEdit's Find in Files if you have it) and look for the source IP address of the attacker, if known. You will then be able to see what all he or she has attempted to do. To see if the attacker's request was successful, search for HTTP Status Codes in the 200 ranges. Codes in the 200 ranges indicate the request was successfully fulfilled. Using a Search Engine such as Google can identify other codes.

Please note that if the target system has been successfully exploited, the integrity of ALL the logs cannot be guaranteed. The attacker very well may have removed his or her trace evidence in the logs. If you are unable to find anything in the above logs, it is time to look at networks logs.

Network Logs

According to CERT's paper titled "Analyze all available information to characterize an intrusion", "Firewall, network monitor, and router logs often remain intact and contain information on hacker activity, even if an attacker gains system-level access." It is unlikely that the intruder has also compromised network equipment. Use any and all information you have to look for more evidence of an intrusion and methodology of the attack. Search for items such as IP addresses (source and destination), protocols, date and times.

Tip: Make sure you know what time zone your logs are in when doing correlation analysis. Microsoft IIS web logs are in GMT whereas normal OS logs are probably on system time.

Throughout this phase of the investigation, be sure to update your logbook entry with findings, observations, commands, system times, and logs.

5.0 Conduct Interview(s)

If you have reason to believe that further investigation is needed, start talking to people, especially the system administrator. Ask them at least the following questions and document their responses in the logbook entry:

- Have you noticed any recent unusual activity?
- How many people have administrative access to the system?
- Have there been any changes (new software, OS upgrades) made recently?
- What applications provide remote access on the system?
- What are the logging capabilities on the system? (Mandia, p.81)

Signs that help security and system specialists recognize that a real incident is occurring are IDS alerts, system error reports, system performance statistics straying from a known baseline, such as a taxed CPU, memory or excessive disk utilization, unexpected shutdowns or restarts, and file system warnings. Other big tip offs are unexpected, unusual, or suspicious process behavior. Please note that being able to identify unusual process behaviors may require engaging technical expertise from system administrators or application owners. As a general guideline, look for missing processes, extra processes, unusual process behavior or processes that have unusual user identification associated with them. Specifically look for:

- Processes running at unexpected times.
- Processes terminating prematurely.
- New, unexpected, or previously disabled processes or services.
- Inactive user accounts that are spawning processes and using CPU resources.

- An unusually large number of processes. (CERT. "Monitor and inspect network activities for unexpected behavior. ")

Another area to be suspect of is unusual user behavior. Log files and commands can assist a security specialist looking for suspicious activity. Again, being able to identify unusual user behavior may require the technical expertise from a system administrators or application developer. Specifically look for:

- Repeated failed login attempts especially to privileged accounts.
- Logins from unusual locations or times.
- Unusual attempts to change user identity.
- Unusual processes run by users.
- Attempts to access restricted files.
- Users logged in for an abnormal length of time (both short and long).
- Users executing an unexpected command. (CERT. "Monitor and inspect network activities for unexpected behavior. " and CERT. "Inspect files and directories for unexpected changes.")

Again, document all information that you are able to glean from the system administrator in the logbook entry, even if it seems trivial. What you document now may be very relevant later.

6.0 Run additional Incident Handling / Collection Tools

Being prepared to handle and investigate a computer incident also requires selecting, installing, and becoming familiar with tools that will assist you in the initial response process (CERT. "Prepare to respond to intrusions."). Security tools identification and testing needs to occur long before an incident occurs. You don't want to be caught trying to install a new tool and learn how to read the results during a live incident. The security specialist should assist with the tool selection by identifying the greatest risks and exposures in the environment. Some things to consider when selecting security tools are ease of use, value of output, licensing and funding, technical expertise, reliability, scalability, and vendor support. There are so many good tools available, both free and fee-based, that this paper couldn't possibly touch on all of them. Instead, some basic tools will be briefly presented along with a list of other URLs for security tools.

Network vulnerability scanning tool - "Nmap ("Network Mapper") is a free open source utility for network exploration or security auditing. It was designed to rapidly scan large networks, although it works fine against single hosts. Nmap uses raw IP packets in novel ways to determine what hosts are available on the network, what services (application name and version) those hosts are offering, what operating systems (and OS versions) they are running, what type of packet filters/firewalls are in use, and dozens of other characteristics. Nmap runs on most types of computers and both console and graphical versions are available. Nmap is free software, available with full source code under the terms of the GNU GPL." <http://www.insecure.org/nmap/index.html>

One feature about this tool that is noteworthy is it compatible on both Unix and Windows machines. It can quickly identify what ports and vulnerabilities may exist on your server. Note on installation; you must also install WinPcap_3_0.exe first on your Windows system.

Nessus is another great vulnerability scan tool, but it can only be run from a Unix system. Nessus is available at <http://www.nessus.org/>

Sniffer – “Ethereal is a sniffer/protocol analyzer that runs on a variety of platforms,” SANS Cookbook. Good to use when you need to monitor the type of traffic flowing through your environment. Ethereal is easy to install and obtain results from. It also has a decent GUI. Ethereal is available at <http://www.ethereal.com/>

Windows Security – DumpSec. “Dumpsec is a security auditing program for Microsoft Windows® NT/2000. It dumps the permissions (DACLS) and audit settings (SACLs) for the file system, registry, printers and shares in a concise, readable format, so that holes in system security are readily apparent. DumpSec also dumps user, group and replication information. DumpSec is a must-have product for Windows NT systems administrators and computer security auditors.” <http://www.systemtools.com/somarsoft/>

Key information is to look at the Report section, and run scans on the following:

- Permissions for file system
- Permissions for shares
- Permissions for users

Missing Patches on Windows Systems - HFNetChk.exe. “HFNetChk.exe is the multi-threaded command-line tool you can use to assess a computer or selected group of computers for the absence of security patches. You can use HFNetChk to assess patch status for the Windows NT 4.0, Windows NT Terminal Server, Windows 2000, Windows XP operating systems, as well as hotfixes and service packs for IIS 4.0, IIS 5.0, SQL Server 7.0, SQL Server 2000 (including MSDE), Exchange Server 5.5, Exchange Server 2000, Windows Media Player, Front Page Server Extensions, Microsoft Java Virtual Machine, Microsoft Data Access Components (MDAC), and Internet Explorer 5.01 or later.” <http://hfnetchk.shavlik.com/>

Who is logged in (Windows) – PsLoggedOn “You can determine who is using resources on your local computer with the “net” command (“net session”), however, there is no built-in way to determine who is using the resources of a remote computer. In addition, NT comes with no tools to see who is logged onto a computer, either locally or remotely. *PsLoggedOn* is an applet that displays both the locally logged on users and users logged on via resources for either the local computer, or a remote one. If you specify a user name instead of a computer, *PsLoggedOn* searches the computers in the network neighborhood

and tells you if the user is currently logged on. Full source code is included.”

<http://www.sysinternals.com/ntw2k/freeware/psloggedon.shtml>.

PsLoggedOn can be run on Windows NT 4.0, 2000, and XP systems.

Open ports and listening applications (Windows) – Fport. “Fport reports all open TCP/IP and UDP ports and maps them to the owning application. This is the same information you would see using the 'netstat -an' command, but it also maps those ports to running processes with the PID, process name and path. Fport can be used to quickly identify unknown open ports and their associated applications.”

<http://www.foundstone.com/index.htm?subnav=resources/navigation.htm&subcontent=/resources/proddesc/fport.htm>.

Fport can be run on Windows NT 4.0, 2000, and XP systems.

Note: Fport is similar to netstat, but if fport shows a rogue process listening for connections, and netstat shows current connections to that process, you may want to kill that process. Use them together! Fport shows currently listening ports, not which are currently servicing remote systems. Use netstat to show all listening ports and all current connections to those ports (Mandia, p.237).

Running Processes (Windows) – PsList. “Most UNIX operating systems ship with a command-line tool called "ps" (or something equivalent) that administrators use to view detailed information about process CPU and memory usage. Windows NT/2K comes with no such tool natively, but you can obtain similar tools with the Windows NT Workstation or Server Resource Kits. The tools in the Resource Kits, *pstat* and *pmon*, show you different types of information, and will only display data regarding the processes on the system on which you run the tools.

PsList is utility that shows you a combination of the information obtainable individually with *pmon* and *pstat*. You can view process CPU and memory information, or thread statistics. What makes *PsList* more powerful than the Resource Kit tools is that you can view process and thread statistics on a remote computer.” <http://www.sysinternals.com/ntw2k/freeware/pslist.shtml>

Knoppix – “KNOPPIX is a bootable CD with a collection of [GNU/Linux](#) software, automatic hardware detection, and support for many graphics cards, sound cards, SCSI and USB devices and other peripherals. KNOPPIX can be used as a Linux demo, educational CD, rescue system, or adapted and used as a platform for commercial software product demos. It is not necessary to install anything on a hard disk. Due to on-the-fly decompression, the CD can have up to 2 GB of executable software installed on it.” <http://www.knopper.net/knoppix/index-en.html>

Unix Security – “The Coroner's Toolkit (TCT) is a collection of tools that are either oriented towards gathering or analyzing forensic data on a Unix system.”

<http://www.fish.com/tct/>

Rootkits:

chkrootkit. "chkrootkit is a tool to locally check for signs of a rootkit."

<http://www.chkrootkit.org/>

Rootkit Hunter – "Rootkit scanner is scanning tool to ensure you for about 99.9% you're clean of nasty tools. This tool scans for rootkits, backdoors and local exploits by running tests like:

- MD5 hash compare
- Look for default files used by rootkits
- Wrong file permissions for binaries
- Look for suspected strings in LKM and KLD modules
- Look for hidden files
- Optional scan within plaintext and binary files

Rootkit Hunter is released as GPL licensed project and free for everyone to use." http://www.rootkit.nl/projects/rootkit_hunter.html

Tiger – "This package identifies common security and configuration problems. It also checks for common signs of intrusion." (CERT. "Identifying tools that aid in detecting signs of intrusion"

<ftp://coast.cs.purdue.edu/pub/tools/unix/scanners/tiger/>

Some useful Unix commands (again may vary depending on flavor of Unix)

W – who is logged in

PS – processes status

Netstat – good interface and network information; has many different flags

Lsof - Open ports and listening applications

Make sure you are using a trusted tool source as a hacker could have replaced the systems binaries with her or her own malicious commands. You could think you are executing a PS command and next thing you know, your system is deleting the hard drive.

Lastly, continue to update your logbook entry with the tools you run, the date and times you ran them, and their output.

URLs of other security tools:

<http://www.insecure.org/tools.html>

<http://www.sysinternals.com/>

<http://www.cert.org/security-improvement/implementations/i042.07.html>

<http://chiht.dfn-cert.de/>

Some useful URLs:

<http://www.doshelp.com/trojanports.htm> - list of well-known trojan ports.

<http://www.arin.net/> - look up IP addresses to identify where they originate

7.0 Final Incident Report

Hopefully by now you have identified how and if the attacker successfully compromised your system. If you have successfully identified the attack vector, the next action should be to see if any similar systems suffer from the same problem. Similar systems can include systems of the same OS, systems in the same IP address range, systems in the same trusted domain, systems that have at least one network service in common (FTP, Telnet, SMTP), systems that share the same file system (CERT. "Analyze all available information to characterize an intrusion."). Once all systems have been checked for the same vulnerabilities, write up your findings in the final incident report. The final incident report should provide details on how the incident was recognized, who was involved, what evidence was recovered, what vulnerabilities exist on the system, and how wide spread the damage is. Once this is documented per your companies' security policy, you are ready for the next phases of Incident Handling as recommended by SANS – Containment, Eradication, Recovery, and Lessons Learned.

If you have not been able to confidently identify how the system was compromised, this too should be included in the final incident report. This allows management to determine the next best course of action, whether that means hiring a security forensics specialist team, engaging law enforcement, or simply deciding to take no further action.

Conclusion

In summary, this paper has provided guidelines, a procedure, and suggested tools to use during the Preparation and Initial Response phases of Incident Management. With preparation, practice, and procedures, the sound of your pager going off doesn't have to invoke feelings of dread; instead you are ready to respond to security incidents in a professional and confident manner.

List of References

Mandia, Kevin and Proise, Chris. Incident Response: Investigating Computer Crime. Berkeley: McGraw-Hill, 2001.

Rosenblatt, Kenneth S. High-Technology Crime. Investigating Cases Involving Computers. San Jose: KSK Publications, 1995. 84 – 96.

Casey, Eoghan. Digital Evidence and Computer Crime. San Diego: Academic Press, 2001. 59 – 171.

Cole, Eric. SANS Security Essentials Cookbook Version 2.2. SANS Institute, 2003.

Cole, Eric and Fossen, Jason and Northcutt, Stephen and Pomeranz, Hal. SANS Security Essentials & the CISSP 10 Domains Version 2.2 Book 2. SANS Institute, 2004.

CERT. "Prepare to respond to intrusions." URL: <http://www.cert.org/security-improvement/practices/p045.html> (21 Aug. 2004).

CERT. "Responding to Intrusions." URL: <http://www.cert.org/security-improvement/modules/m06.html> (21 Aug. 2004).

CERT. "Analyze all available information to characterize an intrusion." URL: <http://www.cert.org/security-improvement/practices/p046.html> (21 Aug. 2004).

CERT. "Identify data that characterize systems and aid in detecting signs of suspicious behavior." URL: <http://www.cert.org/security-improvement/practices/p091.html> (21 Aug. 2004).

CERT. "Monitor and inspect network activities for unexpected behavior." URL: <http://www.cert.org/security-improvement/practices/p094.html> (21 Aug. 2004).

CERT. "Monitor and inspect system activities for unexpected behavior." URL: <http://www.cert.org/security-improvement/practices/p095.html> (21 Aug. 2004).

CERT. "Inspect files and directories for unexpected changes." URL: <http://www.cert.org/security-improvement/practices/p096.html> (21 Aug. 2004).

CERT. "Identifying tools that aid in detecting signs of intrusion" URL: <http://www.cert.org/security-improvement/implementations/i042.07.html> (21 Aug. 2004).

CERT. "Establish a policy and procedures that prepare your organization to detect signs of intrusion." URL: <http://www.cert.org/security-improvement/practices/p090.html> (21 Aug. 2004).

CERT. "Establish policies and procedures for responding to intrusions." URL: <http://www.cert.org/security-improvement/practices/p044.html> (21 Aug. 2004)