
DATA SECURITY MANAGEMENT

INVESTIGATING A COMPUTER SECURITY INCIDENT

Peter Stephenson

INSIDE

Infrastructure Issues, Technologies Involved, Intrusion Detection, Forensic Analysis,
Back Tracing, Conducting an Investigation

You are a security professional. It is 3 A.M. on a Sunday. Your pager goes off and, upon calling the network control center, you find that there has been a serious intrusion into a critical system. What do you do?

It is late in the workday on a Friday. Your phone rings. A fraud examiner from internal auditing believes that a company computer was used to commit fraud in the payroll department. She wants your help. What do you do?

You receive a call from corporate security. They have received a complaint that someone from inside your company's network has broken into an outside system. Security wants your help. What do you do?

All of these scenarios have one thing in common: they are computer security-related incidents. The handling of such incidents must be methodical, preplanned, and consistent with practices that will stand up in a court of law should the necessity arise. That requires a lot of preparation, training, and implementation of the right tools, procedures, and policies. Over the course of the next few pages we will explore the implications of these and other types of incidents and provide you with a set of guidelines for managing them effectively.

INFRASTRUCTURE ISSUES

No investigation can be successful if the computing infrastructure will not support the basic requirements of

PAYOFF IDEA

When a security incident occurs and you are responsible for conducting the investigation, that is no time to begin determining how to perform an investigation. The procedures must already be established. This article provides a structured approach, with helpful checklists, to enable establishment of the policies and procedures necessary to conduct a security incident investigation effectively.

good information security. In the case of an intrusion, for example, complete system logs may be the key. Very often either there are no system logs or the logs are incomplete. If there are no policies and procedures in place for routinely gathering logging information, instituting an investigation, and following explicit guidelines for investigation and recovery, it is likely that your investigation will lead nowhere.

Most computer security incidents do not result in the capture and successful prosecution of the perpetrator. The FBI has estimated that fewer than ten percent of all computer incidents get reported, fewer than ten percent of those get investigated, fewer than ten percent of those result in prosecution, and fewer than ten percent of the prosecutions result in conviction and punishment. That means the computer criminal has a one in ten thousand chance of going to jail for a computer-related crime — great odds for any endeavor. Solid investigation can change those odds materially.

Additionally, the FBI and the Computer Security Institute, in their annual survey on computer crime and information security, gathered the following disturbing facts in 1998:

- 64 percent of respondents reported a security breach in 1998 — up 16 percent from the previous year
- security breaches cost the respondents who could quantify losses a total of \$136,822,000 — up 35 percent over the previous year
- 18 percent of respondents had no idea whether or not they had been hacked
- only 38 percent of respondents had a written intrusion policy, and only 22 percent had an evidence handling policy
- 74 percent of respondents reported attacks from inside their networks, and 70 percent reported attacks initiated from outside
- disgruntled employees accounted for attacks reported by 89 percent of the respondents, while outside hackers accounted for 79 percent (all respondents reported attacks from multiple sources)

It is clear from these statistics that there is a real problem. Fortunately, there are solutions.

In order to set the stage for incident response, there are a number of things that must be done. Here is a quick checklist before we go into more detail.

1. Implement appropriate policies, standards, and practices.
 2. Ensure that legal issues (such as privacy and ownership of company information) are documented in appropriate policies.
 3. Implement, equip, and thoroughly train a computer incident response team (CIRT).
 4. Implement appropriate access controls.
-

-
5. Implement appropriate vulnerability testing.
 6. Implement realtime intrusion detection and logging.
 7. Institute periodic incident response rehearsals and drills.
 8. Institute and maintain relationships with local law enforcement agencies.

One structured approach that allows implementation of the checklist is intrusion management. Although this is a topic for its own article, here is a brief description. Intrusion management is a four-level methodology that helps secure information assets on a large network. The definition of intrusion management is:

Limiting the possibility of a successful intrusion through effective preventive, quality management and detective processes, and facilitating successful investigation of an intrusion should one occur.

Intrusion management is a four-step process. The steps are avoidance, assurance, detection, and investigation. We define these steps, or levels as follows:

- Avoidance: using policies, standards, best practices, and tools such as firewalls, access control, and encryption to deflect attacks against information assets
- Assurance: vulnerability testing and system audits measure compliance with policies
- Detection: realtime logging and interception of intrusion or abuse attempts
- Investigation: tracing intrusions and abuses in a manner that facilitates appropriate responses. Lessons learned feed back into Avoidance

If, as the core of an infrastructure preparation a formal intrusion management program is implemented, it will have provided an advantage towards being able to bring the investigation of a computer security incident to a successful conclusion. Compare, for example, the items in the quick checklist above and the four levels of intrusion management. It will be found that the checklist fits well in the I/M process.

A few words regarding policies, standards, and practices are in order here. Many organizations have been lax in keeping policies, standards, and practices current with the state of their networks and the state of the art in terms of technology, business requirements, and legal issues. There are several specific areas, other than consistency with your business needs and network infrastructure, that should be considered. Some are:

- Privacy Issues. Generally the courts will side with the individual against the organization in matters of privacy if the organization does not have specific policies to protect it. One suggested policy is that
-

employees have no expectation of privacy in the workplace and that all data in any form is subject to scrutiny by the company at its own will and pleasure.

- Search and Seizure. All computers being used on company property are subject to seizure in the event of an investigation. This includes both user-owned computers and corporate-owned computers.
- Investigation Process and Authority. There should be a policy that creates a CIRT and vests it with appropriate authority. Additionally, there should be a detailed standard practice that dictates how the CIRT functions, including procedures, equipment and software to be used, mandatory training, and periodic drills.
- Logging. There should be a standard practice (supported by policy, of course) that mandates logging, defines logged events, and mandates a log retention period (six months at minimum).

TECHNOLOGIES INVOLVED

There are some specific enabling technologies that can help ensure a successful investigation. Some of these are implemented in advance (as part of the infrastructure) and some relate to investigative tools.

Intrusion Detection

Intrusion detection (implemented as part of level 3 of intrusion management) is a set of technologies, tools, and techniques intended to intercept efforts to abuse computers, data, or communications channels. The IDS (intrusion detection systems) FAQ (frequently asked questions) by Robert David Graham defines intrusion detection as follows (with our corrections to the inevitable typing errors in Internet FAQs):

An intrusion is somebody (aka, hacker or cracker) attempting to break into or misuse your system. The word "misuse" is broad, and can reflect something as severe as stealing confidential data to something minor such as misusing your e-mail system for spam (though for many of us, that is a major issue!).

An intrusion detection system (IDS) is a system for detecting such intrusions. For the purposes of this FAQ, IDS can be broken down into the following categories:

1. Network intrusion detection systems (NIDS) monitor packets on the network wire and attempts to discover if a hacker/cracker is attempting to break into a system (or cause a denial of service attack). A typical example is a system that watches for a large number of TCP connection requests (SYN) to many different ports on a target machine, thus discovering if someone is attempting a TCP port scan. An NIDS may run either on the target machine that watches its own traf-
-

fic (usually integrated with the stack and services themselves), or on an independent machine promiscuously watching all network traffic (hub, router, probe).

2. System integrity verifiers (SIV) monitor system files to find when an intruder changes them (thereby leaving behind a backdoor). The most famous of such systems is Tripwire.
3. Log file monitors (LFM) monitor log files generated by network services. In a similar manner to NIDS, these systems look for patterns in the log files that suggest an intruder is attacking. A typical example would be a parser for HTTP server log files that is looking for intruders who try well-known security holes, such as the phf attack. Note that a network IDS monitors many machines, whereas the others monitor only a single machine (the one they are installed on).

Logging. In order for some types of IDS to work, logs must be comprehensive and timely (e.g., they should be created and made available to the LFM in realtime). However, in addition, they should be retained for a period of no less than six months. The usual problem with logs is that they are incomplete (i.e., they do not contain enough useful information) and they are not available for analysis when the event is discovered. The combination of good intrusion detection and a log retention policy can mitigate both of those problems.

The usual excuses for not creating comprehensive logs are size of the logs and performance hits due to their creation. The former is not a problem if a log parser such as AXENT Technology's Intruder Alert (IDA) is used as an LFM. The IDA watches logs, no matter how detailed, and reports important (as defined by you) events immediately, thus removing the requirement for human analysis on an ongoing basis. The latter is a genuine problem and requires care in designing the logging system to avoid.

Logs should be offloaded from the machine where they are logging and stored on a log host specifically dedicated to preserving and protecting logs. This preserves resources on the system being logged, protects the logs themselves, and provides a mechanism for adjusting network and host performance. Finally, it allows products such as the ITA to work on groups of logs with virtually no performance effects against the machine being logged.

Forensic Analysis

Forensic computer analysis is the science of collecting clues or leads from a computer involved in a security event. It requires specialized tools, such as noninvasive bit stream backup software (such as SafeBack from Sydex) or IPFilter (from NTI in Gresham, OR) which extracts e-mail

addresses from a disk. These tools have, as their purpose, gathering evidence and leads from the disks of computers involved in an incident without disturbing the integrity of the disk or the data on it. They are specifically designed to produce information that can be used in a court of law. Their use requires specific and significant training and experience. A discussion of the details of forensic computer analysis is well beyond the scope of this article.

Back Tracing

Back tracing is the technique of tracing an intrusion to its source. There are two general types of back tracing: network and telephone. Network back tracing involves tracing an intruder backwards through routers and hosts on a large network. Intruders will generally jump from system to system, using precracked computers as intermediate platforms from which to launch attacks. By so doing they mask their true origin. Network back tracing requires cooperation from the administrators of intermediate systems.

Telephone back tracing always requires the assistance of the telephone company (and, thus, a court order). Skilled intruders use a technique called phreaking to break into telephone switches and jump from switch (or PBX) to switch, masking the phone number from which they are actually calling.

CONDUCTING AN INVESTIGATION

There are several approaches to managing an intrusion and conducting an investigation. The SANS Institute publishes a ten-step approach to responding to a security incident. While these steps do not, explicitly, address investigation, they are a very good starting point.

- Step 1. Remain calm
- Step 2. Take good notes
- Step 3. Notify the right people and get help
- Step 4. Enforce a "need to know" policy
- Step 5. Use out-of-band communications
- Step 6. Contain the problem
- Step 7. Make a backup of the affected system(s) as soon as practicable
- Step 8. Get rid of the problem
- Step 9. Get back in business
- Step 10. Learn from this experience

Kenneth Rosenblatt, an assistant district attorney in Santa Clara County, CA, is a well-known prosecutor of computer-related crime. In his book, *High Technology Crime — Investigating Cases Involving Computers*, (Ken-

neth S. Rosenblatt, KSK Publications, San Jose, CA), the author defines six specific goals of an investigation of a computer security incident:

1. To understand how the intruder is entering the system
2. To obtain the information you need to justify a trap and trace of the phone line the intruder is using
3. To discover why the intruder has chosen the victim's computer
4. To gather as much evidence of the intrusion as possible
5. To obtain information that may narrow your list of suspects, or at least confirm that the intruder is not a current employee
6. To document the damage to the victim caused by the intruder, including the time and effort spent by the victim in investigating the incident and determining the amount of damage to its computer

In this author's upcoming book on the topic, corporate investigation teams are offered a specific set of seven steps for meeting those goals:

1. Eliminate the obvious
2. Hypothesize the attack
3. Reconstruct the crime
4. Perform a traceback to the suspected source computer
5. Analyze the source, target, and intermediate computers
6. Collect evidence including, possibly, the computers themselves
7. Turn your findings and evidentiary material over to corporate investigators or law enforcement for follow-up

One should begin the investigation by getting the lay of the land. Two things must be done immediately. First, preserve the crime scene and second, gather basic witness information to get an idea of what happened and when.

The crime scene can be preserved by getting people away from any computers or devices considered part of the virtual crime scene, disconnecting any communications links to those computers, and performing a physical, or bit stream backup for the purpose of preserving evidence. Then conduct preliminary witness interviews, very informally at this point, to get a rough picture of what happened. The news reporter's five Ws (Who, What, Where, When, Why) are a good guideline for this questioning.

Once this point has been reached, obviously wrong explanations for the event can be eliminated and one can begin to hypothesize how the attack, or other incident, occurred. Next, a copy of the backup will be used (there should be two: one for evidence and one as a working copy) to create a mirror of the affected computer(s). Never work on the original computer or the evidence copy of backup. On the mirror, begin to reconstruct the crime and test the hypotheses. A replica of the network may

need to be created in the lab; at this point, one may be ready to start the process of tracing back to the suspected origin of the event.

Back tracing is very difficult and, in many cases, it cannot be done unless the intruder is online. However, if the intruder was careless, there may be footprints left on intermediate machines. Amateurs may not have used intermediate systems and will be very easy to back trace.

Next, begin using forensic techniques to extract clues from the computers involved. The tools must be specifically for this activity and they must meet several criteria:

- they must not alter the data as a side effect of the collection process
- they must collect all of the data we want and only the data we want
- we must be able to establish that they worked properly, e.g., as advertised
- they must be accepted, generally, by the computer forensic investigative community
- the results produced must be repeatable

Our next step is evidence preservation and it will certainly include all bit stream backups, related floppy disks, and, perhaps, the involved computers themselves. Remember not to turn on the computers unless there is a bootable floppy in the A: drive, to prevent the computer from booting from its hard drive, and perhaps, writing information to the hard disk.

Finally, the findings will be analyzed, conclusions drawn, and final report prepared. The report presents the conclusions, the evidentiary material to support them, recommendations, and lessons learned. It is not one's place to take action unless directed by management to do so.

During the course of the investigation, it may be necessary to get back online with the systems that were included in the investigation. In fact, there may be systems (servers, for example) which cannot be taken out of service at all. That complicates the work, but it is important to remember that business needs drive security (including investigations), not the other way around. Business needs are never subordinated to the investigation unless forced to by outside influences (damage to the system, intervention by law enforcement, etc.).

SUMMARY

The investigation of a computer security incident is complex. Security professionals must balance the need for a "clean" investigation with the overriding need to keep the business running smoothly and implementing mitigating controls for the future. Most of investigative work is done before the event even occurs. That work is in the form of infrastructure preparation, training of the CIRT, and rehearsals of mock incidents. If prepared for an event before it ever occurs, the intrusive portions of in-

vestigations can be conducted rapidly and the organization can get back to the business of doing business.

A most important lesson to be taken from this discussion is that the investigation of computer security incidents cannot exist in a vacuum. Investigation is not a stand-alone effort. It is, rather, an integral part of a comprehensive intrusion management program. A successful investigation is supported and enabled by good preparation in the form of policies, standards, and practices, as well as a well-trained and equipped CIRT, intrusion detection, and logging mechanisms, and a generally secure network.

Investigation of computer security incidents is, as stated above, very complex. The surface has merely been touched here. However, enough has been provided to get things started to prepare for what most respondents to the FBI/CSI 1998 study found: intrusions are becoming an inevitable fact of business life.

Peter Stephenson is currently President of InfoSEC Technologies, Inc., an information security consultancy. He has 32 years of experience in technology fields, including 17 as a consultant. He has written or contributed to 12 books, written numerous training courses, lectured worldwide on security, and has published many articles in industry publications. He can be reached via e-mail at pstephen@versalink.com.
