




Part 1

Understanding and Exploiting Windows Networks

- ◆ Chapter 1: Network Investigation Overview
- ◆ Chapter 2: The Microsoft Network Structure
- ◆ Chapter 3: Beyond the Windows GUI
- ◆ Chapter 4: Windows Password Issues
- ◆ Chapter 5: Windows Ports and Services



Chapter 1

Network Investigation Overview

As mentioned in the introduction, this chapter provides background information to those readers who do not have a great deal of experience in conducting network investigations. Since much of this book will focus on the techniques used to conduct these investigations, a basic working knowledge of the steps required to use them is essential to getting the most out of this text. Those who have an extensive amount of experience in this area will probably be able to skim this chapter and proceed to Chapter 2, “The Microsoft Network Structure.”

With that disclaimer out of the way, we’ll now cover the steps generally involved in conducting an investigation of a network intrusion or similar network-related incident. It is important to note that this section will deal with broad generalities. Every investigation is unique, and it is the responsibility of the investigator to analyze each situation to determine the appropriate investigative approach. Making these decisions and implementing the associated techniques require a great deal of subject matter expertise, and the remainder of this book is designed to provide you with the information and techniques that you will need to be an effective Windows network investigator.

In this chapter, you will learn to

- ◆ Gather important information from the victim of a network incident
- ◆ Identify potential sources of evidence in a network investigation
- ◆ Understand types of information to look for during analysis of collected evidence

Performing the Initial Vetting

The vast majority of intrusion investigations begin with a phone call. Someone, somewhere has encountered something that makes them suspect that they are the victim of a computer hacker. The first thing any investigator must learn is that many of the people who pick up a phone to report an incident are *not* victims. It is important to conduct an initial assessment of any report and determine its legitimacy in order to avoid unnecessary and unproductive false starts.

WHEN YOU ARE THE VICTIM

This section largely deals with situations where you are working in the capacity of an outside consultant or law enforcement officer, but the questions and techniques discussed still apply to internal corporate security departments or similar groups. All too often, IT administrators, users, and even Security Operations Center (SOC) monitoring analysts leap too quickly to the conclusion that the sky is falling. It is the responsibility of the highly trained security professional (that would be you) to cut to the heart of the matter, provide a reasonable triage of the situation, and either begin the necessary investigation or restore peace and tranquility to the world by telling the people involved, “It will all be OK.”

Since most cases begin with a phone call, it makes sense to perform your initial investigation while on the phone. This saves a great deal of time by allowing you to get preliminary information to determine exactly what resources (if any) you will need to bring to bear to conduct an appropriate investigation into the incident being reported. Obviously, if the reported incident involves classified or otherwise sensitive information, you will need to factor operation-security concerns into your approach. In such cases, you may need to perform even your initial vetting in person at an appropriately secure facility. While each situation will be unique, the following list of questions will provide you with a good starting point for performing your initial inquiries:

What makes you believe that you are the victim of a computer crime? This simple, open-ended question provides you with a lot of information about both the incident and your reporting party. Allow the reporting person to provide you with the story in his own words for a while. Listen for things that indicate the experience and knowledge level of the reporting person. In addition, start assessing the likelihood that an incident has actually occurred. Responses to this question will range from “Our security team was conducting a routine audit of our IDS (intrusion detection system) logs and noticed some anomalies that we found suspicious,” a good sign, to “I received an email and my virus-scanning thing said it was infected,” a not-so-good sign. If the response has anything to do with aluminum foil and alien mind rays, simply refer the caller to the appropriate counseling service—or to your favorite rival agency (you know the drill).

What systems are involved, what data do they store, and were they damaged? Here you are looking to determine whether or not any alleged incident falls within your territorial and subject-matter jurisdictions or your assigned area of responsibility. If all of the computers are located in Spokane and you are a local police officer in Denver, you probably need to end this call with a referral to another agency. Likewise, if you are assigned to a Computer Emergency Response Team (CERT) for a large company and the caller is asking about their mother’s home PC, not their company computer, then perhaps you should provide them with a number for a local IT security firm. Check to ensure that you are the appropriate person to address the alleged incident.

When did the attack occur? While this seems like a fairly simple question, you may be surprised at some of the answers it can generate. It is not at all uncommon for an organization to wait many weeks or months before notifying law enforcement of an incident. Internal politics involving Legal, Public Relations, and other departments can stretch out for long periods of time while the pros and cons of reporting the incident to outside people are debated. This

question will give you an idea of how stale the case may be and how long the victim organization has had to unintentionally lose and delete important evidence.

How was the attack discovered, and who knows about the discovery? This question gives you an idea of how likely it is that the offender knows that his activities have been detected. If the victim organization detected a few anomalies that suggest an attack and immediately called you, then you may have the advantage of catching the attacker unaware. If, on the other hand, the attack was discovered because all systems reported `U h4v3 b33n H4x0red` at bootup, it is a fair guess that the attacker already knows that the victim is aware of the incident. An additional consideration here is that a large percentage of computer incidents are perpetrated by inside users of the impacted systems. Thus, if the victim organization has already circulated emails announcing that they have detected an attack, it is a fair guess that your as-of-yet-unidentified suspect has also been made aware of the discovery.

Did the attacker seem to have familiarity with the network or systems impacted? This question can be used to begin gauging the competency of the attacker, as well as to try to determine whether you are dealing with a rogue insider or an outside attacker. If the attacker gained access to the system using an old administrator account and in one command line copied a file from `C:\files\secret stuff\my special projects\stuff I never told anyone else about\project X\plans.doc`, then you can bet that either the attacker had inside information or the attacker has been to this system before and this is simply the first time that the victim has noticed.

After you have an idea of what has transpired, you will be in a position to make suggestions to the caller to help preserve any evidence that may exist. The instructions that you give in this regard will depend on the specifics of the case, and by the end of this book you will have the knowledge necessary to make that determination. In many cases, the best advice is simply to suggest that the computer be left powered on and that only the network cable be disconnected if necessary to prevent further damage. Again, there will be situations where this is *not* the best idea, but each case must be analyzed independently.

Meeting with the Victim Organization

Once you have gathered enough information to determine that some type of incident occurred and that you are the appropriate person or agency to respond to that incident, it is time to get your investigation under way. At this stage, it is best to arrange a meeting with the reporting person and anyone else who has relevant information about the incident.

MEETINGS ABOUT MEETINGS

It may be in your best interest to also schedule a one-on-one meeting with the reporting person prior to including anyone else in the conversation. This gives you an opportunity to question that person in a little more detail before moving into a setting where his peers and bosses will be watching. If at this private meeting he realizes that a mistake has been made (such as, “Oops, we weren’t hacked; I accidentally deleted those files”), then he can get out and call the whole thing off. If such a realization is made in front of a roomful of people assembled to discuss the big incident that has been discovered, the reporting person’s fight-or-flight instincts may kick in and lead him to provide you (and everyone else) with false or misleading information to save face.

If possible, the first face-to-face meeting with the victim organization should take place in a quiet meeting room with at least one whiteboard available. After the initial introductions, have the reporting person explain what is known about the incident in very broad terms. During this meeting, there are some very specific pieces of information that you will need to obtain, so don't let the initial overview get into too much detail. After everyone agrees on a very general view of what you are all gathered to discuss, take control of the meeting and begin to gather information in a systematic manner. The following sections will give you some ideas on information that you need to ascertain, but keep in mind that no two investigations will be exactly alike.

THE BIG MEETING

Once word gets out that law enforcement or security consultants are coming to interview staff about a possible computer crime incident, things can spiral out of control within the victim organization very quickly. Everyone who thinks they are important will insist on attending, and the initial introductions will sound like a job fair as everyone explains what their unit does and how important they are to the overall mission of the organization. You will likely encounter representatives from the Human Resources department, senior managers, chief information officers, company lawyers, computer incident response teams, outside consultants, and all other imaginable players. Just take it all in and note who the key players really are. This is your opportunity to once again size up the people with whom you are dealing. Also, never forget that many computer crimes are committed by people within the victim organization. Don't reveal too much about your thoughts, techniques, or plans in these types of meetings, because the perpetrator may be sitting in the room.

Understanding the Victim Network Information

Before you can even begin a serious discussion of any incident, you must first establish a baseline understanding of the network environment in which the incident took place. This is no different than performing an initial assessment of the scene of a burglary or any other crime. Just as an investigator of a physical crime must identify possible points of entry or exit, location of valuables, items that may be missing or moved, and so on, the same concepts apply when conducting a computer-related investigation.

FOR MORE INFORMATION

Remember that this chapter is only a high-level summary of the issues involved in responding to a reported computer intrusion. The remainder of this book will discuss issues specific to conducting network investigations in a Windows environment, but for readers who feel they need additional background information on intrusion response in general, we recommend *Incident Response and Computer Forensics, Second Edition* by Prosise, Mandia, and Pepe (Osborne, 2003) to supplement your existing knowledge.

One of the first things that you will need to get clear in your own mind is the topology of the victim network. The topology refers both to the physical location of the various pieces of hardware, media, and so on that constitute the network and to the way that data logically flows through that network. You should have a clear understanding of any connections that lead to outside networks such as partner organizations or the Internet. Identify which security controls, such as firewalls, IDSs, and filtering routers are in place at possible entry or exit points to the network and within the core of the network. Obtaining a current network diagram (if available) or using a whiteboard to sketch out the network visually at this point can be very helpful. Start trying to identify possible sources of evidence within the network, such as devices that generate logs and/or monitor network communications. Gain an understanding of any proprietary technologies or systems with which you are not familiar by asking specific and detailed questions to clarify the network's design and function.

DID LEIA ATTACK FRODO OR WAS IT PICARD?

Keep in mind that the administrators and other people whom you will be interviewing work on the victim network day in and day out. They will know much of it like the back of their hands, and they will often speak to you as if you should as well, referring to computers by their internally assigned names (such as Frodo, Leia, or Picard) and speaking in organization-specific acronyms. When conducting initial interviews, make sure that you understand everything clearly. Nobody is fully versed in all current aspects of network technology, every proprietary vendor's product, and the implementation details of these items in every network. You must ask questions—lots of questions. This is not the time to allow your ego to interfere with your interview. If you don't know something, ask the interviewee to explain the technology in question and how it impacts the network's function.

Get a sense of how the network is used and what normal patterns of usage might be. By understanding what type of activity is typical, you will be in a better position when analyzing evidence for activity that may be abnormal and malicious. Here are some questions that will help you determine normal usage patterns:

- ◆ Do you have employees who log in from remote locations?
- ◆ Do partner organizations have access to any of your systems?
- ◆ During what times do your employees normally access the network?
- ◆ Do remote connections normally last for long periods of time (such as interactive user logons), short periods of times (such as automated transactions or updates), or variable amounts of time?
- ◆ Which systems house sensitive data, and which users should have access to those systems?
- ◆ Are all of your systems located in this facility, or are you using remote data centers or cloud service providers?

By asking these and similar questions, you will be able to understand both how the network is structured and how it is used by legitimate users. Without this information, it is virtually impossible to perform a successful network investigation.

SILVER LINING?

When you start asking the important questions, the fact is that the victim organization may not know many of the answers. Many organizations lack adequate data governance and information rights management; many simply do not know where their sensitive information is stored or who should be accessing it. If you are working for an internal security team, these meetings are often a good time to point out that extra budget should be allocated to systems and processes designed to identify ways to keep ahead of incidents in the future.

Understanding the Incident

Now that you have had a chance to get acquainted with the electronic crime scene, let's get into the details of the incident itself. You've already given the reporting person two opportunities (once in the initial vetting and once at the beginning of the face-to-face meeting) to give you the highlights of what has occurred, so you should have a fair idea of what has happened that raised concern. At this stage, you should direct the conversation and get all the detailed information that you can about the timeline, methods, scope, and outcome of the incident. Don't allow the interviewees to rush ahead of you. Make sure that you understand all of the necessary details of each step before allowing the conversation to move forward.

One thing to keep in mind is that the victim may have already developed a theory of the crime that might or might not bear any similarity to reality. They may even have put together a very fancy, post-incident response report and believe that they are handing you a gift-wrapped case ready for prosecution. While we have received many such reports, we have also never seen one that was 100 percent accurate. As the investigator, it is your job to review any information that you receive and check it for factual accuracy.

After you have determined exactly what the alleged attacker did that caused such upset, it is time to ask one of the most important questions of the interview: "What have you done in response to the incident?" This can be a very telling question. First, you can further gauge the competency of your victims by listening to the steps that they took and analyzing the appropriateness of their response. Second, you get a good idea at this point how much evidence might still be available to you.

For example, if you ask your victim what they did in response to the incident and receive an answer of, "We screamed in sheer panic for 30 seconds and then immediately called you," then you know two things: these may not be the most technically proficient people, and your evidence is likely right where the attacker left it. If on the other hand you receive a response such as, "We immediately downed the affected systems, did a bit-level zeroing of all media contained within them, reinstalled from known-good media, and restored the network to full functionality," you know you are dealing with a fairly technically competent crew who has stomped all over your evidence and your chances of working a successful case.



Real World Scenario

TRUST NO ONE

At the outset of one intrusion investigation, we were presented with a very nice report from a highly paid security contractor who analyzed the logs from the victim system and came to a conclusion about the crime. His report indicated that the initial attack occurred on November 15 and that it consisted of a series of failed attempts to intrude upon the box that eventually led to a successful attack. The report concluded that the attacker was unfamiliar with the system and that this was the first attempted attack against the victim system.

In performing our own analysis of the same logs, we noted that the attack on November 15 had been successful on the first attempt. In fact, we noticed that it exploited a piece of code that had been written by the victim organization and that had never been disseminated to any other group. While the logs did show some experimentation by the attacker, they were indicative of attempts to further increase the attacker's control of the system after already exploiting the box. The method of attack was fairly complicated and suggested a good deal of familiarity with the system. This attack was clearly either the work of someone with inside information or the work of someone who had exploited this system before and who was returning to enter the system once again.

The contractor who created the report was also responsible for keeping the victim system secure. He believed that *his* systems were secure and that, prior to this incident, nobody had broken into them. Whether out of malice or simply as the result of preconceived notions, reports by people who work closely with the victim systems are bound to contain some type of bias. Be certain to review them carefully and come to your own, independent opinion that is based on the facts at hand rather than unsubstantiated beliefs.

Identifying and Preserving Evidence

There are many possible sources of electronic evidence that you can use during the course of your investigation. One of the biggest challenges of dealing with electronic evidence is that it, even more than physical evidence, needs to be collected promptly and correctly. Much of this book will talk about the proper ways to collect digital evidence from memory and from disk, but first you must identify where that evidence may be.

These days, you should first ask the victim about their network topology, both logical and physical. It is important to understand whether you are dealing with a network that is primarily supported by an onsite server room, an offsite dedicated data center, or a cloud service provider in another country. Understanding whether critical systems are running as an installed OS on dedicated servers or as virtual machines can also impact how you collect evidence and the order in which it should be collected. Chapter 17, "The Challenges of Cloud Computing and Virtualization," addresses some of these considerations in more detail.

One of the most useful sources of evidence in any network investigation will be the logs generated automatically by various devices throughout the network. Since it is created by

an automated process during the normal course of doing business, log evidence falls under an exception to the hearsay rule under the U.S. Federal Rules of Evidence and is admissible as evidence at trial. Log evidence can provide the most thorough and accurate account of what transpired on a network. Teaching you to identify, collect, and analyze these logs from Windows computers is the subject of a number of chapters in this book.

In addition to logs that are generated by Windows-based computers, many other devices also generate valuable logs of evidence. It is important to identify what logs are kept, where they are kept, and for how long they are kept. This bears repeating: *Identify what logs are kept, where, and for how long.* This is an extremely simple question that often is very difficult to get correctly answered. In many IT shops, logs get configured at initial installation and then are never seen again. Many generations of IT workers may have come and gone between the time logging was enabled and the time you arrive asking to see the logs. Many organizations will automate the logging subsystems to rotate and archive logs with no human intervention, creating an “out of sight, out of mind” situation. You may need to dig, poke, and prod to arrive at an accurate accounting of which devices within the network are configured to log, where those logs are stored, where they are archived, and for how long they are kept before being deleted and overwritten.

DUMB AS A LOG

You will find that in many organizations, logs are not on the top of the administrators’ daily chore lists. We have frequently asked administrators if they back up and preserve their logs and have been told that they do *not* do so; however, we cannot stop our inquiry at that point. Next, we ask them if they back up the data on the victim computer. To which almost all administrators will quickly volunteer that they perform full system backups and that they archive those backups in a grandfather-father-son or some other common rotation. Logs are simply data, and when it comes to Windows logs, they are almost always stored on the main system drive of the computer. If the administrators are backing up the system and archiving those backups, then they are also backing up and preserving their logs as well, whether they intended it (or even realized it) or not.

In addition to identifying the log evidence that may be floating around the victim organization, you should also inquire about backups of any system that was impacted or that might have been impacted by the incident. Frequently, backup tapes or other media are overwritten in a set rotation. You want to ensure that any backups that may prove useful in your investigation are pulled out of that rotation and seized as evidence as soon as possible to avoid their inadvertent destruction. Also, you will need to identify any possible sources of evidence that may exist outside your victim organization, such as logs at an Internet service provider or data held at a partner organization. U.S. law enforcement officers will want to issue a preservation letter to secure that evidence immediately to avoid losing it and any benefit that you may get from it.

2703(F) ORDERS

U.S. federal law, in the Electronic Communication Transactional Records Act, Title 18 U.S.C. 2703(f), states the following about the requirement to preserve evidence:

“(1) In general.—A provider of wire or electronic communication services or a remote computing service, upon the request of a governmental entity, shall take all necessary steps to preserve records and other evidence in its possession pending the issuance of a court order or other process.

“(2) Period of retention.—Records referred to in paragraph (1) shall be retained for a period of 90 days, which shall be extended for an additional 90-day period upon a renewed request by the governmental entity.”

Two important things to note in this statute are that the request can be made by any governmental entity and that the receiving organization must preserve the evidence for 90 days while a court order or other process is prepared. This gives broad authority to rattle off a request citing 18 U.S.C. 2703(f), requesting the immediate preservation of logs or other evidence. Such a request is generally referred to as a *preservation request* or *preservation letter*. You can then further develop your case, consult your prosecutor or legal advisor, and obtain the appropriate process required under the Electronic Communications Privacy Act or other applicable law to retrieve whatever evidence you are seeking.

Outside the United States, countries that are signatories to the Council of Europe’s Convention on Cybercrime must enact similar legislation that allows authorities to expeditiously preserve digital evidence. More information about this convention is available from the Council of Europe’s website, <http://conventions.coe.int/treaty/en/treaties/html/185.htm/>.

Establishing Expectations and Responsibilities

Before ending your meeting, you will want to determine exactly what the victim organization is expecting from you. A victim who tells a private security firm that they do not want anyone to ever know about the incident and simply want to identify and repair any damages may meet with a receptive audience. That same victim making that same request of law enforcement may not be so fortunate. For most law enforcement agencies, the goal is generally to identify and prosecute an offender. This early in an investigation it would be inappropriate for criminal or other investigators to make any promises of future confidentiality; such decisions are at the discretion of prosecutors and judges (at least in the United States; in other countries the rules may be very different). In short, such promises are not usually within the authority of a U.S. criminal investigator to make. It is important that all parties understand what can and cannot be promised at all stages of the investigation; everyone should be kept well-informed so their expectations are kept reasonable. Failure to ensure this can have disastrous effects later in the investigation.

You might need various members of the victim organization to assist you in your investigation. You will likely need to schedule follow-up meetings with specific administrators to further elaborate on the workings of specific systems so that you fully understand the environment in which your investigation is taking place. You might need to ask someone to locate old records

that indicate how log rotations were initially configured and exactly what types of events are being audited by those logs. You might also need to establish parameters for contacting you and responding to any further incidents or anomalies. Make sure that all of these types of issues are resolved and that everyone understands what their responsibilities are before ending the meeting.

It is important that you remember that at this stage you might not have identified the full scope of the incident or the location of all possible sources of evidence. Make sure that you keep open lines of communication with all of the involved players to ensure that you have up-to-date and accurate information. Also, never forget that many incidents are perpetrated by inside employees, so stress the importance of keeping the incident a secret to anyone who must be involved. Ask for complete secrecy from all parties, but assume that each of them has told everyone they know.

Collecting the Evidence

Once you have met with and interviewed the relevant members of the victim organization, it is time to take the information that you have learned and proceed with collecting evidence. Again, many of the techniques used to collect that evidence will be discussed later in this book, but in general terms you must collect evidence in a way that preserves its value in a criminal proceeding. This means that you do not substantively alter the evidence during collection and that you maintain an accurate chain of custody for each piece of evidence that you collect. Evidence in a network investigation can consist of many different things, and we will look at some of the different types of evidence that you may want to collect.

COLLECTING THE CLOUD

If your victim network takes services from a cloud service provider, the legal agreements between the victim and the provider can become critical to your ability to access the logs or other data you require to conduct your investigation. You should get the legal eagles involved early in the process if a cloud service provider is involved.

One of the more obvious places to look for evidence is the logs of devices designed for network security. Items such as network or host-based intrusion detection systems (IDS) can provide a wealth of information about successful and attempted attacks performed within and against a network. An IDS monitors communications that come into, out of, or through a network or specific host (depending on the type of IDS and its configuration). These communications are then analyzed on the fly by the IDS, which looks for signatures of known attacks and other anomalies that might indicate malicious or prohibited activity. The response of the IDS to a suspected problem can range from noting its findings in a log to storing a copy of the suspect traffic, sending an alert to a specific user, or even taking active countermeasures against the perceived threat (which technically would make the device an intrusion prevention system, or IPS). The information and logs created by IDSs can be a great starting point for an investigation since they can provide a summary of detected malicious activity within the network.

DIGITAL SOURCES OF EVIDENCE

Just as with a physical crime scene, a digital crime scene can be rich with potential evidence. Don't let the fact that you are now investigating a digital crime distract you. Digital crimes are investigated using the same basic principles as any other criminal offense. If you were investigating a bank robbery, you would undoubtedly survey the crime scene, interview witnesses, canvass the area for discarded items, round up security tapes from nearby establishments, and so on. The same logic applies at a digital crime scene. You determine the topology and normal usage of the network, speak to the system administrators, examine logs of impacted and related systems, and examine the output from IDSs and other network security devices. Investigating a network incident does involve specialized knowledge and methods, but don't let that fact distract you. Digital crime is still crime, and its investigation follows the same general route as any other investigation.

Other devices can also generate security-related logs. Firewalls, which are devices configured to permit certain network traffic while blocking other types of connection attempts, can be configured to stand as sentinels at the entry to a network, between subnets, or on specific hosts. Firewalls are often configured to log the packets that did not meet the criteria established for allowable communications and were thus blocked. Proxy servers, or Application layer firewalls, can provide even more specific log data regarding the activities of the various users and systems within the network. Even routers are often configured as a first line of defense by dropping certain types of communication as soon as they try to enter or exit the network. These *screening routers* can also be configured to log the packets that they drop, although such logs are much more common in firewalls and proxy servers.

Devices that are designed to accept or authenticate inbound network connections will frequently perform a great deal of security-related logging as well. Remote access servers, RADIUS (Remote Authentication Dial In User Service) servers, wireless access points, VPN (virtual private network) concentrators, and other methods of connecting or authenticating to a network are usually configured to log any attempted and/or successful connections. As possible entry points to a network, these devices should be familiar to the victim organization's administrators, but legacy, redundant, or backup systems are often overlooked by administrators but specifically targeted by intruders. Don't forget to analyze any network diagrams and other information for indications of ways into a network that may have been omitted in previous discussions with the administrators.

VARIETY IS THE SPICE OF LIFE

You will find that the amount of available log evidence varies dramatically from one investigation to another. The largest factor in this equation is the victim organization. If your victim is a government agency handling sensitive information, you will probably have more logs than you can imagine detailing all aspects of the incident. If, on the other hand, your victim is a small mom-and-pop company whose system administrator is the family's youngest son, then you may be wishing that you had more evidence on which to proceed. You can only work with the evidence that is available.

If your victim's network employs a Security Information and Event Management (SIEM) product, it can be a great source of information to you. A SIEM tool acts as a repository for logs from various sources and can be used to perform advanced analytics on the log files to identify potential problems. SOCs that monitor networks for security or other adverse events often rely on SIEM tools to detect incidents, and many network security incidents are first brought to light by these tools. Keep in mind that a SIEM is only as good as its configuration allows it to be; the garbage-in, garbage-out rule certainly applies here. A SIEM does not conduct your investigation for you, but it can be an excellent guide to point you toward potential sources of additional evidence.

Another exceptional source of evidence, if you are lucky enough to be working in an environment where they are used, is network monitoring tools. Products like NetWitness (www.netwitness.com), Solera (www.soleranetworks.com), and AccessData's SilentRunner products (<http://accessdata.com/products/>) record network traffic, extract metadata about all communications, and enable a wealth of network forensic data that can make your life much easier. These tools allow investigators to query historical network communications to look for file transfers, exploit use, malicious software downloads, authentication attempts, and a whole host of other information, all from one console. If these systems are implemented in your victim's network, they can help speed up your investigation and point you toward systems and logs that can hold more evidence related to your case.

Data that is stored in the memory of a running system can be of great evidentiary value. By determining which processes are running, which ports are listening, which connections are active, which users are logged in, and other information about a running system, you can generate a good picture of what that computer was doing at the moment you seized it. We will examine this concept in more detail in Chapter 5, "Windows Ports and Services." Such information can be of extreme importance in a network investigation, and the methods of gathering this type of evidence will be discussed in Chapter 6, "Live-Analysis Techniques."

The logs from individual victim computers are usually vital pieces of information in any network investigation. The logs should be collected and analyzed offline to avoid modifying or altering their content to the point that it jeopardizes their evidentiary value. In "Part III: Analyzing the Logs" (Chapter 11–Chapter 17) of this book, we detail methods of performing Windows log analysis and outline the information that can be gained about an incident from that analysis. Some network administrators have taken additional steps to preserve logs in centralized locations for easier analysis. Typically, shops that use such log-aggregating techniques are more security conscious, and the administrator will be able to guide you to the logs and explain how they are stored. Keep in mind that the logs from computers that are not known victims can also be important. Evidence of failed attacks might be present on these systems, which can lead to additional charges against the perpetrator and provide you with additional information about her methods and techniques. Also, other computers might have been involved in authentication of compromised accounts or other aspects of network activity and might contain log evidence of that activity despite never having been compromised themselves.

The data stored on victim systems is also critical to a successful investigation. Any files that are present on a victim system may later be found on the suspect's computer, allowing you to further tie the suspect machine to the incident. Tools left behind by the attacker can be analyzed to determine additional information about the attacker and the attacker's techniques (see Chapter 10, "Introduction to Malware," for details on suspect tool analysis). Evidence contained within the registry or elsewhere on the system can be of critical importance in locating and prosecuting the offender. In "Part II: Analyzing the Computer" (Chapter 6–Chapter 10), we

outline many of the types of evidence that can be found on both victim and suspect computers to help further an investigation and solidify a criminal prosecution.

HONEYPOTS

Keep in mind that if the attacker is not yet aware that the incident has been discovered, you may have the option of setting up monitoring equipment to watch for future illegal activity. By configuring proper data-capture tools, you can monitor the victim box to gather more evidence about your attacker as more attacks are made. This can be a great way to identify other compromised machines, aid in identifying your attacker, learn more about the attacker's methods, and gain more evidence to use in criminal prosecution. You will have to weigh the risks versus the rewards with the victim organization based on the sensitivity of the information being exposed to further attacks and the willingness of the victim organization to accept that risk. Because of the legal and risk-management sensitivities of such an approach, it is normally necessary to talk to your prosecutor or legal advisor before performing any network monitoring to ensure compliance with all applicable laws.

Analyzing the Evidence

Now that you have identified and collected the evidence, the real work can begin. Obviously, after the evidence has been properly collected, you should make working copies of all digital evidence and use these copies when performing your analysis. While this phase of your investigation is more static and controlled than evidence collection, it is still a time-sensitive process. Keep in mind that you have secured and preserved all of the evidence of which you are currently aware; however, it is very common that your analysis of that evidence will lead you to uncover more sources of evidence. Digital evidence can be easily destroyed, whether maliciously by the attacker, accidentally through hardware failure, or systematically through log rotations. This creates an urgency to complete your analysis as quickly as possible in order to follow any logical investigative steps that your analysis may suggest.

You can perform many types of analysis on the evidence collected from the victim organization, and this text explains tools and techniques for doing so in a Windows network. When you perform your analysis, many facts can be of assistance to your investigation. The specifics of each case determine what is and is not useful, but when you consider that you might have literally terabytes of data to sort through, it might help to know what types of needles you are searching for in that digital haystack. We provide some examples of data that are frequently of investigative interest and suggest some techniques for locating that data both in this section and throughout this book. For now, let's focus on patterns and data that are frequently helpful to the investigator.

One of the simplest places to start is to focus on activity that occurred around the time of the first known incident. You will often collect vast amounts of data during the evidence-collection phase of your investigation, and limiting the scope of your initial search to a finite time period can expedite your discovery of relevant data. For example, you might focus your initial log analysis efforts on the date and time of the first malicious activity that the victim noticed, or you

might perform a forensic analysis of all files added or modified during the time that the attacker was first known to have accessed the system. You can always expand the scope of your analysis to prior events to look for previously undetected intrusions or intrusion attempts after you have a better idea of what occurred during the reported incident. Chapter 12, “Windows Event Logs,” discusses ways to make your searches more effective by filtering based on time and date ranges, as well as other criteria, to expedite your analysis times.

ANALYZE ACCURATELY BUT QUICKLY

Keep in mind that many attackers do not directly attack their victim organization from their home computer (those who do are the low-hanging fruit that make for quick-and-easy cases). More sophisticated attackers compromise a series of computers and bounce their commands through them in order to obscure their actual location. As a result, the analysis of the evidence seized at the victim organization often leads not directly to the attacker but rather to another victim. It is important that you perform your analysis quickly so that you may contact the other victim location and obtain their logs and other sources of evidence before so much time has elapsed that their logs have rotated and been lost forever. A clever attacker will make you repeat this process several times before you manage to find her actual IP address and location, so make certain that you perform each step of your analysis as quickly as you can accurately do so. Also, don’t forget to issue preservation letters as soon as possible to keep any evidence intact while you arrange to collect it.

As part of your initial interviews, you learned a great deal about the network and its normal usage. You should look for connections to the network that break from these normal usage patterns. For example, in a network that has many users in the northeastern United States, connections from Brazil might be suspicious. Similarly, a company that is staffed from 9:00 a.m. to 5:00 p.m., Monday through Friday, might not have a great deal of legitimate network activity at 1:00 a.m. on a Sunday. The existence of such activity is, again, what we call a clue.

Attackers will frequently create or modify accounts in order to ensure that their control of the system can be maintained. Look for accounts that have been modified since the date of the incident. Also, check all accounts that have increased privileges on the system and confirm that each one is a legitimate account. Chapter 2 addresses this issue in more detail and outlines how to identify accounts with elevated privileges; Chapter 14, “Other Audit Events,” discusses the log entries that Windows generates when accounts are created or modified.

Similarly, hackers also usurp accounts that have previously been inactive or disabled, so look for accounts that are suddenly being used after long periods of inactivity. Chapter 4, “Windows Password Issues,” discusses many of the ways that hackers obtain passwords for valid accounts in order to disguise their activity as normal network traffic. In addition, rogue insiders may already have an account on the system that they are utilizing to perform unauthorized acts. If you do identify an account that is being used maliciously, be certain to document as much of that account’s activity as possible. Chapter 13, “Logon and Account Logon Events,” details the Windows auditing capability for user logons and shows how to use those logs to track user activity on the system.

Many attackers attempt to hide the evidence of their presence by altering or deleting logs. These alterations may result in large gaps in log files that in themselves can be evidence of a crime. Sometimes the event that tipped off the victim organization to the presence of an incident is the deletion of all of their system logs or the disabling of the logging functions on victim machines. Chapter 13 shows you how to correlate logs from various systems so that if one system is not logging or has had its logs altered or erased by the attacker, other evidence can still be located and used to document the attacker's actions. Chapter 15, "Forensic Analysis of Event Logs," demonstrates how to recover deleted log records.

If a computer has been intruded upon, the hacker may have targeted that machine for a specific purpose, especially if the computer in question stores particularly sensitive or valuable data. Focus on files that are known to be compromised, suspected to be compromised, or likely to be targeted. Analyze which users have accessed those files and whether each access was legitimate. Chapter 14 discusses the Windows file access audit capability and how to use it to perform such an analysis.

Many network services are required to successfully perform work within the network, and many of the computers that provide these services will generate logs of their activities. By examining these logs, you can gain valuable insight into activities throughout the network. Chapter 11, "Analyzing the Logs," deals with services such as DHCP that may generate logs as soon as a computer connects to the network. Chapter 13 details the role of the domain controllers in granting access to many network resources and illustrates the logs that they create.

Many intruders will install software on a victim system that performs unauthorized functions and/or reports back to the intruder. It is important to know what malicious software (also called *malware*) has been installed on any victim system. Identifying malware and its function can help you learn more about the attacker, gain insight into his purpose, identify other compromised systems, and lead to other sources of evidence or security concerns. For example, identifying that a piece of malware left by an attacker on one system is being used to capture valid usernames and passwords that are usable throughout the network would greatly impact the scope of your investigation. Often the attacker will install the malware in a way that ensures that it will restart whenever the system is rebooted. Chapter 9, "Registry Evidence," shows you how to analyze the registry and other locations to help locate installed malware.

Malware can do many different things to the victim system, from monitoring network communications to providing a backdoor through which the intruder can reenter, but it can be difficult to detect. Chapter 5 discusses many of the common Windows services and ports to help you recognize the software that is supposed to be on the victim machine and to help you better identify the malware that is not supposed to be there. Chapter 6 deals extensively with techniques used to query the RAM of a running system to identify malware and document the effects that malware may be having on the system at that moment. Finally, Chapter 3, "Beyond the Windows GUI," talks at length about a special category of malware known as *rootkits*, which have the ability to hide their presence on the victim system while exerting a great deal of control over that system.

There are many different types of information that can further any network investigation and many different techniques to identify, collect, analyze, and understand that information. This book focuses on the elements that are unique to a Windows environment, while leaving more general sources of evidence such as IDS logs, firewall rules, and the like to be discussed by others. It is vital that you never lose sight of the fact that any network—be it based primarily on

Microsoft, open source, or other platforms—will have many different pieces of evidence available within it, and it is your job to know how to properly handle all of it.

Analyzing the Suspect's Computers

After analyzing the evidence from the victim network, you will hopefully have developed enough information to spur your investigation in the correct direction. Law enforcement will serve subpoenas for outside IP addresses that were used by the attacker, possibly leading you to other victim networks and even more evidence to be analyzed. At the end of this process, you will (hopefully) arrive at an IP address being used directly by your attacker, obtain a subpoena for the provider to whom that address is assigned, and identify the computer that your attacker was using to perform the evil deeds that spawned the investigation in the first place.

At this point you have discovered another valuable source of evidence: the suspect's equipment. When searching the suspect's home or office, be aware of the many possible pieces of useful information that you might find. Obviously, you will want to seize the suspect's computer (and in a forensically sound manner), but don't forget the many other potential sources of evidence. A savvy attacker will often store incriminating files on removable media, physically hidden somewhere they are hard to find. When you consider the wide array of digital media on the market today, there is virtually no place that cannot hide some form of storage device. Make certain that any search warrant that you obtain contains appropriate language to allow you to search for any electronic, magnetic, optical, or other storage media so that you may perform a thorough search of the area. Attackers also frequently have printouts, scraps of paper, or other notes lying around that contain usernames, passwords, IP addresses, computer names, and so on. Ensure that your warrant contains appropriate language to allow you to seize this very valuable evidence.

Attackers generally perform a recon of their intended target to determine the structure of the network, locate potential vulnerabilities, and develop an idea of which machines are most valuable to the attacker. They will then exploit a vulnerable system and gain a foothold within the network from which they can perform further recon, launch further attacks, set up rogue sniffers, and perform other steps to increase their influence within the network. As the attackers gain control over more boxes, they will add rogue processes, install backdoor listeners, and otherwise embed into each system to ensure that the boxes remain under their control. When valuable data is discovered, the attackers will exfiltrate that data from the victim network to store or possibly to sell. Each of these steps has the potential of leaving evidence for you to find not only at the victim's location but also on the computers that the attackers are using.

Once you have located and properly collected the evidence from your suspect, you must analyze that evidence to try to tie the suspect to the incident. There are many types of evidence that you can use to accomplish this task, and we will explore some of the more common ones here.

The suspect may frequently have performed open-source intelligence gathering about the victim network. Most organizations offer entirely too much information about themselves and their networks to public access over the Internet, and many attackers will use this against the victim by culling through these pieces of information to assist in their target recon. Attackers typically map out as much information as possible about the victim organization and its network. Information about personnel can be used for social-engineering and spear phishing attacks. Information about projects being performed by various divisions or offices can be used to help the attacker focus the attack on the areas most likely to yield the information being

sought. Finally, information about the network's structure and uses can help the attacker find vulnerabilities through which to compromise the network. You should carefully search the suspect's data for any mention of the victim organization, its IP addresses, personnel, or network. All of this can be useful evidence if the case goes to trial. Chapter 9 discusses some of the places where this type of evidence might be located on the suspect's computers.



Real World Scenario

SEARCH, SEARCH, AND SEARCH AGAIN

When it comes to executing search warrants, a search for digital evidence can be one of the most difficult types of warrants to serve. When executing a search warrant for crack cocaine, you can search anywhere in the named property that a single rock of crack could be located—effectively anywhere in that property. The same applies with digital evidence. Modern removable media can be even smaller than a crack rock, and your search should be performed with that level of thoroughness. To illustrate the point, here are just a few of the places that we have found digital evidence:

- ◆ A Secure Digital (SD) card was hidden under the paper inside of a tin of Altoids mints.
- ◆ Digital evidence was stored in cell phones and digital cameras.
- ◆ A piece of Juicy Fruit gum was removed from a pack and replaced with a Sony Memory Stick, which was then wrapped in the original Juicy Fruit wrapper and placed back in the pack.
- ◆ A hard drive was placed in a plastic bag, hung on a coat hanger, and then hidden by a shirt hung on top of it on the hanger.
- ◆ In the middle of large collection of commercial audio CDs, one CD was removed from its case and replaced with a CD-R containing evidence.

There are many other possible places to store digital evidence. Watches, pens, Swiss army knives, and even dolls that contain USB flash drives are being marketed. Micro SD cards are about the size of a fingernail but able to store gigabytes of data. Video game machines and digital video recorders can be modified to store data and then connected to a home network to allow ready access to that data from any computer. In short, ensure that your searches are adequately thorough so that you don't miss a vital piece of evidence that would seal your case for the prosecutor.

You should search for any files that may have come from the victim organization's computers, since such evidence is as damning as stolen televisions or any other stolen property. By using hash analysis techniques, you can quickly scan the suspect's machine for any proprietary files that might have been taken from the victim systems. This can be a powerful source of evidence and give you a great deal of leverage with the suspect in subsequent interviews.

During your analysis of the evidence from the victim network, you likely performed tool analysis on any tools left behind by the attacker. If you find those same tools on the attacker's systems, that is obviously great evidence linking that computer to the crime. In addition, if your analysis determined that the attacker compromised a particular service on the victim machine, the presence of hacker tools capable of exploiting that service is also powerful evidence. Finally,

the presence of tools commonly used to recon and attack computers, such as scanners, sniffers, exploit scripts or toolsets, rootkits, and mass rooters, can also be evidence in your investigation.

TOOLS OF THE TRADE

Hackers rely on a wide array of tools to perform their evil deeds. For those who are not familiar with this terminology, here's a brief summary of some of the main categories of tools used by attackers:

Scanner/Port Scanner A scanner is a tool used for target recon that attempts connections to multiple different ports on multiple machines. A scanner can provide a great deal of information regarding the open ports and services on a target system, providing details such as the operating system used, the services offered, the ports to which the services are listening, and the versions of the OS and services. For further information, read about Nmap, one of the hacker scanners of choice, at www.insecure.org.

Sniffer This is a software package that uses the computer's existing network interface card to monitor the traffic that the computer is able to receive. Sniffers can be general-purpose sniffers, which are designed to capture any type of network communication, or they can be specialized sniffers that are configured to scan for particular types of information such as usernames, passwords, and so on. For further information read about Wireshark (www.wireshark.org) and Cain & Abel (www.oxid.it).

Trojan Any program that purports to have a useful function but instead performs a malicious function is generically called a Trojan horse, or simply Trojan. Hackers will frequently replace common system commands with Trojanized versions that perform a similar function to the real system tool but also conceal information from the user or perform some other malicious function. Rootkits (discussed in Chapter 3) frequently contain a number of Trojans.

Mass Rooter This multipurpose tool both scans for a known vulnerability and then actively exploits that vulnerability. Mass rooters can compromise numerous systems in a matter of minutes.

Exploit Any method of taking advantage of a vulnerability on a target system to gain unauthorized access to that system or its resources is generically called an exploit. Exploits can exist as source code, as compiled executables, or as modules for a more complex framework. For an example, read about the Metasploit Project (www.metasploit.com).

The suspect might have logs on her own systems showing connections to the victim organization. We have found history files on suspect machines detailing every command typed by the attacker and recorded perfectly for presentation in court on the suspect's own machine. Routers and wireless devices owned by the suspect can also maintain connection logs that can be used against an attacker. Also, remnants of the commands used to perform the attack might still exist in slack space or in the registry of the suspect's computers. Perform a thorough search of the suspect's computers for any ties to the victim organization, its IP addresses, and its machine names. Chapter 9 explores this issue in more detail.

Attackers will frequently discuss their exploits with other people. Some like to brag about their technical accomplishments and how many systems they “own,” while others may be attempting to sell the compromised information to the highest bidder. Regardless of their intent, remnants of electronic communications made from the attacker to other individuals can frequently be found on the suspect’s computers. Check for emails, chat logs, website postings, and other sources of communication to see if your suspect made admissions to others that can be used against her in the interview room and in court.



Real World Scenario

TALKING THEMSELVES INTO A CORNER

We once investigated an intrusion into a government system in which the logs showed an attack being flawlessly executed against the victim machine from a particular IP address. Within 60 seconds of that attack, six more identical attacks were initiated from six different, geographically distributed IP addresses. This was the digital equivalent of, “Hey bud, watch this!” The suspect had gone into an IRC chat room to show his other hacker buddies the new attack that he had discovered, and each of them then tried the same attack. Our suspicions were confirmed after multiple simultaneous search warrants were executed, and the suspects confirmed our theory during their interviews.

Remember that any link that you can find on the suspect’s computers to the victim organization can be powerful evidence. In addition to making wonderful fodder for a jury, this type of evidence can also be used to provide the suspect and her attorney with an incentive to cooperate with your investigation. The suspect’s computer is also likely to give you additional leads into other machines that were compromised by the attacker, generating even more cases and charges. Frequently, when faced with overwhelming evidence in one attack, the suspect will cooperate by providing information about other attacks and allow you to identify and assist other victims.

Recognizing the Investigative Challenges of Microsoft Networks

Many excellent books have been written about responding to computer incidents, but the majority of these books discuss the topic in broad terms without addressing the specifics of any given platform. This book takes the next step in dealing directly with networks that rely primarily on Microsoft products to provide the majority of their core network functions.

The primary obstacle faced by security practitioners of Microsoft-based networks is the proprietary and closed nature of the source code. Unlike open-source alternatives, Microsoft’s products are distributed only as compiled executables without any accompanying source code. As a result, in order for anyone to determine how the product reacts to any given situation, the product must be set up in a test environment and subjected to that situation. With open-source options, the source code for the product can be analyzed to make determinations of how the

product is supposed to handle certain eventualities. This is not necessarily a security problem with Microsoft products. Indeed, it could be argued that protecting the source code actually enhances the security of the product since potential attackers are not able to parse through it to locate vulnerabilities. Despite the philosophical arguments that always accompany the open-source versus closed-source debate, it does limit the options available to those who investigate intrusions within a Microsoft network.

Examples of how this can hamper the investigative process can be found in the Microsoft log files. Most of the logs stored on a Linux/Unix platform are plain-text logs. They can be searched, grouped, or sorted using any text editor or other utility that can read text. By contrast, the system logs on Microsoft systems are stored in a proprietary, binary format that requires special software to even read them. Since source code is not available for the Microsoft operating systems, you cannot analyze it to determine how the OS will record particular system events.

Since you cannot do your own analysis of the code, you must rely heavily on documentation provided by Microsoft (or others) and on independent testing to accurately report how the OS will respond to certain events. Unfortunately in many cases, such as log analysis, the available literature is fairly sparse. This means that you must put a great deal of work into determining how the operating system records events before you can even begin to use those recorded logs to make our case. Fortunately, this book outlines the major functions of the operating systems of which you will need to be aware and will show you how to use those functions to conduct a productive investigation.

The Bottom Line

Gather important information from the victim of a network incident. It is important to properly vet any report of an incident to ensure that the appropriate people and resources are utilized to address every report. As the number of reported incidents continues to rise, this requirement becomes more and more important to ensure the most efficient utilization of limited agency resources.

We outlined various questions and considerations that any investigator responding to an incident should keep in mind when first interviewing the members of the victim organization. The steps you take at this stage can set the tone for the rest of your investigation and are vital to a rapid and effective response.

Master It You are called regarding a possible computer intrusion into a defense contractor's network. After performing an initial interview with the reporting person by phone, you feel confident that an incident has occurred and that you should continue your investigation. What steps would you next take to gather additional information to launch an investigation?

Identify potential sources of evidence in a network investigation. Evidence within a digital crime scene can be located in many different places. It is important to consider how data flows through a network to determine which network devices may have recorded information that can be of evidentiary value. In addition to logs that may be kept on the victim computer, explore logs generated by firewalls, IDSs, routers, wireless devices, authentication servers, and proxy servers that may have recorded information about the attack.

Master It You are called to a company where they suspect that a disgruntled system administrator has accessed the company's database from outside the company and

deleted multiple important records. The logs on the database server have been deleted, leaving no trace of the attack. What are some other possible sources of evidence for this incident?

Understand types of information to look for during analysis of collected evidence. After the evidence is properly secured, the analysis phase should be completed as quickly and accurately as possible to allow time to follow up on any other investigative leads that the analysis may suggest. The analysis should be thorough and may be time-consuming, but as new investigative leads are discovered, you should take immediate action to preserve that evidence for later collection.

Once suspects are located, a thorough search for digital evidence should ensue to gather all possible evidence of their involvement in the incident. As analysis of collected evidence occurs, you may uncover evidence that proves the reported incident along with evidence of crimes that were not previously known. Thorough analysis and interviewing may lead to the discovery of multiple other victims and other crimes.

Evidence to search for will depend on the specific investigation, but common items of interest include the following:

- ◆ Access around the time of the suspected incident
- ◆ Access at unusual times or from unusual locations
- ◆ Repeated failed access attempts
- ◆ Evidence of scanning or probing that preceded the incident
- ◆ Data transfers that occurred after the incident
- ◆ Evidence of the victim's files, IP addresses, and the like on the suspect's computers
- ◆ Detection of known malicious software or exploit methods

Master It While investigating an alleged attack against a local government finance server, you locate and seize a computer believed to have been used by the suspect. What are some types of evidence that you should look for on the suspect's computer?

