APPENDIX A

Answers to Questions

This appendix contains the answers to most of the questions at the end of each chapter. A few of the essay-style questions are left for the reader.

CHAPTER 1

1. What constitutes an incident?

   In NIST SP 800-61, an incident is defined as "violation or threat of violation of computer security policies, acceptable use policies, or standard security practices."

2. Who should define what constitutes an incident?

   Although the NIST publication provides a standard definition, every organization must further define what is considered an incident based on its own priorities and capabilities. An organization with a full-time IR team may have the capacity to handle any event that occurs. Smaller teams may need to restrict the definition of an "incident" to issues that can be effectively managed.

3. Is all malware OS specific?

   Malware, or anything with malicious content, is not necessarily data in compiled form. Interpreted code (such as Java and Python) can be used as malware and are not operating system specific.

4. In the first case study, what network architecture mistake did the victim company make?

   Network segmentation. The victim organization did not properly isolate the DMZ from the corporate environment.

5. How many phases are in the attack lifecycle?

   There are seven phases in the attack lifecycle: initial compromise, establish foothold, escalate privileges, internal reconnaissance, lateral movement, maintain presence, and complete mission.

6. Do all attacks encompass all phases of the attack lifecycle?

   No. The attack lifecycle was developed as a structured way to describe attack phases common to most incidents.

CHAPTER 2

1. List the groups within an organization that may be involved in an incident response. Explain why it is important to communicate with those groups before an incident occurs.

   - **Business line managers** These folks can help identify investigative priorities, as well as coordinate cooperation within their groups.
   - **Legal** Internal and external counsel can assist with internal coordination between groups, handle external communication, and provide a layer of support when external parties get involved.
   - **Human resources** HR can assist with internal personnel matters, should your investigation lead you down that path. HR can also help you reward or recognize team members or others in the company who have contributed to an IR. Remember that a high-profile IR is fast paced and can be very disruptive to both the company and employees' personal lives.
   - **Desktop IT support** These folks help with deployment of endpoint technology, and are a great resource for coordinating remediation events.
   - **Compliance** An organization's compliance officer can assist with navigating the various standards that may apply.
   - **Network infrastructure** If your organization has the ability to support a monitoring infrastructure, coordinate with the networking team to determine the ideal locations to collect traffic. If the configuration of the network and routing can be optimized to support security monitoring before an incident occurs, collection of network-based evidence will yield useful results earlier in the investigation.

2. Your organization receives a call from a federal law enforcement agency, informing you that they have information indicating a data breach occurred involving your environment. The agency provides a number of specific details, including the date and time when sensitive data was transferred out of your network, the IP address of the destination, and the nature of the content. Does this information match the characteristics of a good lead? Explain why or why not? What else might you ask for? How can you turn this information into an actionable lead?

   It is a good start. Good leads are relevant, detailed, and actionable. If the notice actually included the time of the incident, the destination address, and type of content, it may provide enough information for you to begin examining your environment for additional information. Most external notifications will not have sufficient information to identify the originating computer. Often lacking in specificity, external leads may be less actionable than other, internally generated leads. To begin generating actionable leads, you will need to match the provided information against logging and monitoring that may be in

place. Were any outgoing communications established with the destination IP address? What other traffic crossed the network boundary around the time of the event?

3. Explain the pros and cons of performing a live response evidence collection versus a forensic disk image. Why is a live response the most common method of evidence preservation during an IR?

   A live response is typically used for two purposes: to gather volatile evidence before a system is shut down for imaging, and as a "first look" at a system to determine whether it requires additional attention. In large enterprise investigations, you may find that most of your investigation is accomplished through performing live response. Many investigations involve several dozen computer systems, and most organizations lack the personnel or time to examine a significant number of forensic disk images. One significant reason to collect hard drive images rather than rely on live response (LR) is that the entire operating environment is preserved. Rarely do you know all of the questions that need to be answered at a single point in time, and repeating the LR every time a new data source is needed is a very disjointed means of collection. Furthermore, it is possible that the evidence that was once present on a system is overwritten or deleted by the time the question is asked.

4. During an investigation, you discover evidence of malware that is running on a system. Explain how you would respond, and why.

   How you respond depends on a number of factors. What stage is the investigation in? How much have you learned about the attackers and what they are after? Where was the malware discovered? On a random workstation, a domain server, or a point of sale terminal? Depending on the situation, your team may perform a live response, shut the system down for offline analysis, or gather the malware and allow the system to continue running in order to gather intelligence. In any situation, analysis of the malicious binary is an important step to understanding the incident. After collection, perform, at a minimum, static analysis of the binary to learn what allowed it to run on the system (compromised credentials or a system vulnerability, for example). Determine whether the malware has a communication or persistence component and use that information to improve your monitoring or enterprise-wide searches.

5. Explain why creating and searching for IOCs is a critical part of an investigation.

   Indicators of compromise are simply structured embodiments of suspicious activities, conditions, and markers of known malicious activity. In any investigation, you are looking for indicators in some form. It's not difficult to justify that you look for things you know are suspicious, as well as other markers of suspicious behavior. When you standardize on an IOC format, however, it provides a structured means to write, store, and share indicators. That becomes critical for any enterprise IR team.

6. When does the remediation process start, and why?

   The remediation process should start as early as possible in the investigation. The remediation team has a responsibility for short-term and long-term remediation planning. Short-term actions can be taken at any time, whether an emergency fix needs to be applied or the team is preparing for an eradication event. The planning of long-term remediation steps occurs throughout the investigation. As less optimal conditions are identified, the remediation team takes note and will research and prioritize the fix.

CHAPTER 3

1. Explain how a simple network service such as DHCP might become critical in an investigation.

   Many "simple" network services pass configurations or information down to workstations, such as in this example. Most generate log data that can help during an IR. In this example, a DHCP server component will log the IP address assignments passed out to DHCP clients. In some situations, this last hop is very important from an investigator's perspective. Take the example of an external notification of a system that has reached out and placed files on a remote drop site. That external notification may consist solely of an IP address of your web proxy and timestamp of the connection. If logging is enabled, you can search the proxy logs for that outbound connection and learn the internal IP address that issued the request. The next step is to correlate proxy logs with DHCP or DNS logs to determine the system that was assigned that address when the connection was established.

2. Discuss some of the challenges related to outsourced IT services as they pertain to conducting an investigation. How can they be dealt with?

   We have experienced several challenges when working with an outsourced IT provider, including the following:

   - **Perception of ownership** A contractor may not prioritize in the same way that an invested party would.
   - **Cost** Enterprise-wide changes, such as the installation of a software agent, can be cost prohibitive. Many service providers will charge a per-workstation fee. Paying the service provider for around-the-clock support and personnel to assist in an investigation adds up quickly.
   - **Responsiveness** The incident is your emergency, not theirs. Unless IR-related clauses are in a contract, short-notice changes may not be possible.
   - **Familiarity** Oftentimes, the service provider will be more aware of configurations and processes than the organization. This complicates IR planning and remediation execution.

3. Why is asset management a critical aspect of preparing for an incident?

   Knowing what software, services, and applications are in use by each part of the enterprise will allow you to understand where an attacker may move. When an organization loses track of servers and workstations, tracing activity to the endpoints becomes quite difficult.

4. What forensic tools are acceptable in court? Explain your answer.

   This is a bit of a trick question in that the process is evaluated, not necessarily the tool itself. This is why vendor claims about admissibility and suitability for meeting a legal standard are misleading. The vendors do, however, take advantage of a brand recognition and use that as a shield. For example, if two separate teams perform an analysis, one using a common automated forensic suite and the other using a manual process, the name recognition shortens the qualification process. This happens regardless of the actual quality of the result unless the other party catches the disparity, unfortunately. Over the past 15 years, the big forensic suites have had serious bugs that are not acknowledged publicly, but their recognizable brand results in an immediate pass through the system.

   The answer is that any tool may pass that standard, if it conforms to a sound methodology that can be tested, reviewed, is generally accepted by the relevant community, and is a repeatable process that has a known error rate.

5. Why is centralized logging important?

   Here are two reasons why centralized logging is important to an organization:

   - Log data is immediately exported to a system that may be outside of the control of an attacker. This serves as a trustworthy repository that cannot be easily altered.
   - Log data from many sources can be easily searched and sequenced for understanding how an attack has occurred (or is occurring).

6. What is a DNS blackhole? When would you use one? Write a BIND zone file that will accept all queries for a malicious domain and answer with an RFC1918 address.

   A DNS blackhole is a way to prevent connections or to redirect communications to a known malicious domain to a system that you control. Here is a sample BIND zone file for "malicious.com" that redirects all traffic to a web server that has been configured to log all requests. Naturally, this only affects malware that will resolve a hostname rather than communicate to a specific IP address.

```
;This zone redirects all queries for the domain malicious.com
to our internal blackhole web server.
$TTL 86400  ; 24 hours
@     IN    SOA    ns.company.com. admin.ns.company.com
                        201300001   ; Serial
```

```
                        28800 ; refresh  8 hours
                        7200        ; retry    2 hours
                        864000      ; expire  10 days
                        86400 )     ; min ttl  1 day
                  NS    ns.company.com.
            A     192.168.254.2     ; Web server with logging
*           IN    A     192.168.254.2
```

This zone file can be used for any number of malicious domains by using it as the master zone—for example, in the named configuration files:

```
zone "malicious.com" {type master; file "/var/named/zone/blackhole";}
```

CHAPTER 4

1. What are the five primary checklists you should use to collect initial information about an incident? List two key data points for each, and explain why they are important.

   - Incident Summary
     a. *Nature of the incident.* Know what you are investigating before allocating resources. Was the incident a spear phishing attempt? A series of failed logins at 4 A.M.?
     b. *Actions taken since detection.* Documenting what administrators or others have done before the IR began will help you understand what evidence may have been changed. Additionally, if any action caused a system to generate log data, you may need to deconflict the data sources so you don't report administrator actions as related to the incident.
   - Incident Detection
     a. *How the detection occurred.* This helps you understand the suspicious behavior and gives you a starting point for the investigation.
     b. *What were the time sources used during detection?* Ensure that the timestamp information is synchronized with a reliable source. If there is drift, document the offset and ensure that any analyst is aware that a drift was present.
   - Additional Details
     a. *Individual system details.* If your team performs a forensic image on a workstation, document completely that workstation's configuration. RAID parameters are as important as the make/model/serial number.
     b. *The responsible party for the server or workstation.* During the examination, the person or team responsible for a system's maintenance will be a

primary source of baseline information. These parties will be able to answer questions on what functions the application performed, the nature and location of stored data, and how to best implement certain remediation recommendations.

- Network Details
  a. *List of malicious IP addresses and host names.* This list will help you identify traffic that may need to be monitored and will be used in your remediation efforts.
  b. *Remediation steps taken by the network team.* This includes changes to perimeter firewalls or DNS. Some malware will be affected by these changes, preventing the investigators from identifying additional compromised systems.
- Malware Details
  a. *"Metadata" on malware collected in the enterprise.* Knowing exactly when and where malicious applications were installed helps define the scope of the incident. Note that the "where" includes a list of affected systems as well as logical path names.
  b. *The status of any analysis performed on the malicious applications.* Track who analyzed the files, internal to your organization as well as external. Note that submitting samples to sites such as VirusTotal means that the files are available to any company that has paid for access. Consider the dangers of submitting files to an antivirus (AV) company. First, the malware may have been written for your environment, so host names or user names may be exposed. Second, it becomes part of the AV company's intel feed. If your investigation lasts longer than a couple of AV update cycles, you'll find that AV may detect and remove files, completely subverting your own remediation plans. If you want to submit data, wait until your own remediation has been completed.

2. During an incident, you discover that an attacker conducted a brute force attack against a database server. You record the event in your incident management system. What time zone did you use? Explain why.

   In several chapters of this book, we address the importance of standardizing the representation of timestamps. The easiest way to keep timelines accurate and without the possibility of misinterpretation is to settle on a single time zone for documentation. What time zone you choose can depend on your organization. For example, if your company has one office in Chicago, no assets outside of U.S. Central time, and you aren't likely to collaborate with anyone, U.S. Central may work. Give the global nature of our investigations and the level of coordination, UTC is the best choice for us.

3. When maintaining a timeline of attacker activity, what four pieces of information should you record for each entry? Pick a file, such as win.ini on a

Microsoft Windows computer or /etc/hosts on a Unix-based host, and make as many time-related entries as you can for just that one file.

Most entries in an attacker timeline include a minimum of four fields: the date the entry was added, the timestamp of the event, the source of the data, and a description of the event.

The following table is an example of how a modified /etc/hosts file from an HFS+ file system can be documented:

Date Added	Event Time (UTC)	Host	Description	Data Source
Jan 14, 2014	Oct 7 12:55:51 2013	planck	Last Access date	File system metadata
Jan 14, 2014	Oct 1 01:39:27 2013	planck	Last Modification date	File system metadata
Jan 14, 2014	Jul 29 01:39:27 2013	planck	File Create time	File system metadata
Jan 14, 2014	Jul 29 01:39:27 2013	planck	Inode Birth time	File system metadata

4. Why is it important to keep case notes? How often should you update them?

   Case notes serve to help you document progress, hypotheses, and findings as you perform examination on collected data. Its best to maintain notes in a central location as the investigation proceeds. Through seemingly endless cycles of collections, examination, research, taskings, and meetings, the maintenance of case notes will help data and tasks from being overlooked. Furthermore, when the investigation is over and time passes, the details not documented in reports may be forgotten. A historical reference for internal use or to answer external queries helps immensely.

5. Assume your web server farm was found to be operating as a "drop site" for attackers who were active in other companies. What would your investigative priorities be? How would you demonstrate that the incident was handled in a competent manner?

   This situation can be a bit tricky. Given time, one of the other companies will detect the incident and identify yours as a drop site. At that point, the matter may not be under your control. In this situation it is likely that the incident, from your organization's perspective, is motivated by access. Therefore, there is a low likelihood that you are a target. A good set of investigative priorities may be as follows:

   a. Identify how the external parties were able to use your site to store the data.

b. Validate the vulnerability or misconfiguration that led to the incident and review systems and logs for signs of activity unrelated to the storage of external data.

c. Identify the data placed on your servers and consider whether notifications should be performed.

d. Remediate the issue and monitor for additional activity.

During the entire process, the detail in your documentation lends credibility to your investigation. Ensure that timelines are complete and that all identified leads are followed, if possible. A common pitfall is to concentrate on one facet of the investigation (for example, malware or attempts to identify an "adversary").

CHAPTER 5

1. From the Practical Malware Analysis book (practicalmalwareanalysis.com/labs), generate host-based indicators for the binary file Lab03-03.exe.

   Dynamic analysis of this malware would reveal that it is a keylogger that performs process replacement on svchost.exe and creates a log file named practicalmalwareanalysis.log. An indicator of process replacement functionality includes the presence of common functions used for performing process replacement, such as CreateProcessA, WriteProcessMemory, and ResumeThread. Be sure to check that all of these functions are present together within a single file because just one or two is likely to generate false positives. You could also search for a process named svchost.exe that has no parent because that condition is unusual. Also, because the malware creates a file, you can look for the presence of a file named practicalmalwareanalysis.log. As a final check, you could examine all files on a system for one that has the same MD5 checksum as the malware.

2. In February 2013, Mandiant published a report (intelreport.mandiant.com) that detailed the group known as APT1. This report outlines the typical attack progression observed at numerous victim organizations. In the section titled "APT1: Attack Lifecycle," the typical process used for internal reconnaissance is described. Using the methodology described, generate a set of indicators that can help your organization identify this type of activity. Note that a methodology indicator does not necessarily identify malware. Consider both host-based and network-based indicators.

   The Mandiant report highlights a batch script that the APT1 attackers used to perform reconnaissance within a victim's environment. The script contained a sequence of commands that are built in to most versions of Microsoft Windows. Therefore, there is no "malware" to look for. Rather, we must

evaluate this activity to look for possible unique artifacts that would not typically be present on a system or the network due to normal user activity. Keep in mind that because every environment is different, some indicators that work well in one may not work well in another.

First, because you know the attacker is storing the output of their commands to a file in C:\WINNT\Debug\1.txt, you could examine the Debug directory on a handful of systems to see what files are commonly present. The results of that research would reveal that the presence of files with a .txt extension in the %systemroot%\Debug directory is unusual. You could create an indicator that looks for that condition.

Second, you could examine a handful of systems in your environment to see if users are commonly executing the commands listed in the APT1 report. For example, most users would never execute the commands ipconfig, tasklist, and netstat. You could write an indicator that examines the system to determine if those commands have been run. You can examine areas such as the Windows Prefetch, Shim Cache, event logs, and any local artifacts such as software metering logs to determine whether those commands have been run on a system.

Third, you could examine whether some of the queries, such as

```
net group 'domain admins' /domain
```

generate network or host indicators that are unique. For example, you might examine the event logs on domain controllers to see if this query generates an event that is unique enough to use as an indicator. The same goes for the local system event log for commands such as net user and net localgroup administrators. In this case, something to keep in mind is that normal users are unlikely to execute these commands, so any unique artifacts associated with them are potentially very good indicators.

3. From the *Practical Malware Analysis* book (practicalmalwareanalysis.com/labs), generate host-based and network-based indicators for the binary file Lab06-02.exe. Note that you may be able to generate extremely effective network signatures if you perform dynamic analysis on the binary and understand what the malware is looking for.

   Static or dynamic analysis would reveal two very good network-based indicators that you could look for. The program uses the HTTP user-agent "Internet Explorer 7.5/pma" and downloads the web page located at http://www.practicalmalwareanalysis.com/cc.htm. Because both of these are fairly unique values, any network traffic containing the user-agent or URL would be suspect. As a final check, you could also examine all files on a system for one that has the same MD5 checksum as the malware.

CHAPTER 6

1. If your evidence sources do not include multiple independent categories, what steps could you take to increase confidence in your conclusions?

   This depends on whether the incident is active during your response. If so, you can increase the amount of logging and detection measures in the environment to gather additional evidence of activity. Quite often, that will lead you to another source of evidence. There is rarely a situation where only one source of evidence, such as a registry key or a Prefetch file, is the sole means to determine that something happened. If so, you may need to address whether the incident meets your threshold for action.

2. In the ACH fraud scenario, think of another reasonable theory for how an attacker might have gained access to the CFO's computer. Explain how you would proceed with investigating that theory.

   Another reasonable theory that would require additional investigation is the use of valid credentials. Although malware may have been found on the computer, the presence does not mean that is the source of the activity. A complete explanation of the methodology is left to the reader.

3. In the customer data loss scenario, a number of steps were taken to verify the customers' complaints. List at least two other useful steps you could have taken to help verify the customers' complaints or isolate the source of the data loss.

   In some situations, it would be worthwhile to immediately task someone to review web service logs and database server logs for activity during certain time periods. The success of this option would depend on the amount and quality of logs that are retained, as well as the accuracy of the time of the reported loss. Although we noted that there are challenges in involving an affected customer, we have worked with companies that have employed this practice successfully. In one situation, it was determined that a common Trojan was collecting e-mail and contact lists from many end users. The coordination was carefully executed, and the end user offered to allow a third party to image and review their computer.

CHAPTER 7

1. What investigative questions is a live data collection likely to help answer?

   Potentially, a live collection may answer the majority of questions you may have during the initial stages of an investigation. It depends on the level of detail that you collect during a live response. A small collection consisting of users, processes, select registry keys, and network state can help you determine if there are signs of malicious activity. A comprehensive collection that includes data

sources, such as browsing history and the NTFS master file table, can reveal far more. Naturally, there are reasons for and against voluminous collections.

2. Should you perform a live data collection on each system you suspect is compromised? Explain your answer.

   Generally, it is advisable to collect data from any system that you suspect has been compromised or is otherwise involved in an incident. If your investigation is at the point where a significant amount of data has been learned about the incident and you need to identify the scope, performing a targeted (with respect to indicators) LR over a large population of systems can answer that investigative question. It can also give you enough information to determine whether a subset requires a detailed examination.

3. In what situations would collecting an image of memory be most useful to the investigation?

   We have found memory images to be most useful in two situations. First, when malware is primarily memory-resident and leaves little trace evidence on storage. Second, when attackers use encryption. We've gained access to many a password-protected RAR file through the examination of memory images.

4. During an investigation, you identify a suspect system that runs an operating system you have not dealt with before. You are tasked with performing a live data collection. Explain how you should proceed to gather live data.

   In the chapter, we discussed the primary types of data that should be collected in an LR. As a baseline, research how that type of information can be retrieved from the environment. Are built-in commands available? Perhaps specific GNU tools can be recompiled on the unfamiliar platform? After the baseline is established, determine other types of information that is quickly obtainable and is of potential value. For example, showprods can list software that has been installed on a given day in IRIX. Finally, determine whether a trusted environment is necessary. If so, you will need to find another workstation that is not compromised and build a complete LR toolkit, including binaries for the commands you identified, the associated libraries, and any support files.

CHAPTER 8

1. A new hard drive duplication and imaging product called "Cyber Imager Pro" is released on the market. Your boss wants you to evaluate the product for use in your organization to perform forensic duplications. Describe what you would do to evaluate the product and determine whether or not it is acceptable.

   First, ensure that the tool performs the function that you are expecting. An example of where this has gone wrong is when IT administrators used

Norton Ghost to preserve hard drives. It was a well-maintained, reliable tool with a good history in the IT space. It did not, however, perform a complete disk image. After the functionality has been defined, test the tool in a series of situations, including when the source material is damaged. Follow the testing procedures used by NIST, at the URL listed in Chapter 8. Thoroughly document the process and findings.

2. If you have connected evidence hard drives to a system for imaging, do you need to use a write blocker if you are going to boot to a Linux-based forensic CD? Explain why or why not.

   The short answer is that you should always use a write blocker, if one is available. Depending on the forensic CD you use, the state of the source material, the volume definition, and the file system in use, there may be a chance that the simple read-only flags you pass to mount commands are not sufficient. This is a situation where familiarity with common file systems and partitioning schemes is essential.

3. You are tasked with establishing the standard hard drive imaging tool and procedure for your organization. What considerations and steps will you take to create the standard?

   Considerations should include the following items:

   - The level of experience of the persons who will perform the imaging
   - The native environment used by the investigators
   - The forensic tools that will be used to review the data
   - The various types of storage used by the organization
   - The available tools and whether verification has been performed

   When creating the standard disk imaging procedures for an organization, research the considerations listed here. Generate thorough documentation, including step-by-step guides that can be distributed to system administrators if necessary.

4. You are attempting to image a hard drive, but the process seems to randomly fail part of the way through. The imaging software reports that it is unable to read from the source drive. How would you approach troubleshooting the issue?

   If the software handles the errors gracefully, let the process complete and document the errors that were reported. If the tool fails to generate a valid image, you may want to experiment with dd_rescue. We have found that dd_rescue can be helpful in recovering a drive with errors by changing the input buffer as they occur. When the exact sector or sectors that return errors are identified, dd_rescue can skip ahead, replacing the bad sectors with zero-filled buffers. In some circumstances, working backward from the end of the drive is effective. Another option is to find a company that can do data recovery.

CHAPTER 9

1. When architecting a new network monitoring system, what types of questions should be asked of the IT or network staff? How can you help ensure complete visibility into traffic egressing your networks?

   Three initial questions help you understand the effort required to perform reasonably comprehensive network monitoring:

   - How many network egress points are there?
   - What is the average volume of traffic crossing the egress points?
   - Does hardware at the egress points support SPAN interfaces?

   Complete visibility requires that all traffic traverses a connection being monitored, the sensor can keep up with the data transfer rate, and the method used reliably applies the rules or signatures to the network data.

2. How can your team detect the following suspicious activities through statistical network monitoring?

   a. Installation of dropper malware

   b. Malware that retrieves commands from a remote site

   c. Potential data theft

   Statistical network monitoring can help identify suspicious activities through analysis of transfer size and frequency. The following are examples that may detect anomalous activity, depending on the situation:

   a. **Installation of dropper malware** HTTP GET requests over a certain size, issued immediately after e-mail is accessed.

   b. **Malware-retrieving commands** Connections occurring at a predictable frequency, where the originating system retrieves a small amount of data.

   c. **Potential data theft** Abnormal amounts of data exiting the network.

3. What is perfect forward secrecy and how does it affect your ability to decrypt SSL traffic? How might you decrypt SSL traffic that is encrypted using a forward secrecy algorithm?

   Perfect forward secrecy is a key-exchange process that generates a pair of session keys for a protected communication. The process is designed to ensure that the long-term secret keys cannot be compromised should the session keys be determined. Decryption of sessions that employ PFS is difficult because each session needs to be attacked independently.

   The only way to decrypt SSL traffic that uses forward secrecy through Diffie-Hellman (DH or ECDHE) is to gain access to the session keys, which are not accessible in a packet capture. A server-side attack where session keys are compromised would be necessary.

4. How would you quickly identify whether a large .pcap file with thousands of sessions contained FTP activity? How would you extract the transferred files?

Review the capture files for FTP and FTP-data connections. Once the sessions have been identified, you can reconstruct the files in Wireshark or similar tools.

CHAPTER 10

1. What are some common ways you might discover services or applications that are helpful to an investigation?

   The most expeditious method to understand the data sources available to an investigation is to discuss the tools used by IT systems management and networking staff. The tools that they install for management, compliance, business continuity, and availability gather and store significant amounts of information.

2. Are network services such as DHCP critical to incident response? If DHCP logs are not available, how might you still determine what system an IP address was assigned to?

   Yes, network services such as DHCP provide essential information to the investigator. In the event DHCP logs are not retained, review other services that may show a connection between an IP address and another property. For example, SAMBA logs may capture the system and user name, and web proxy logs may connect an address with a host name.

3. Your company recently deployed a new application across the enterprise. The application is a new disaster recovery tool that automatically backs up data from users' systems to a central server. The backup data is encrypted, but the backup software maintains a local log file in plain text. How might this service be of use to incident responders?

   If you are this far into the book, the answer is likely to be obvious. If the backup process runs at a short interval, the log (and the files stored remotely) may capture attacker activity before they have the opportunity to clean up. The point we made in the answer to Question 1 applies to this situation. If you schedule regular meetings with IT staff to understand the tools they use, your team can test the tools on virtual machines to understand how they identify and store data prior to their deployment. If the new data source fills a gap in your LR or examination plans, add it to the collection and analysis process. In this situation, you would want to retrieve the disaster recovery log file during a live response.

4. In this chapter we mentioned "booting a copy of a forensic image." What are we talking about? Describe the tools and methods that exist to perform that task.

   In certain circumstances an examiner may want to let an image run, typically to extract run-time data out of memory, review data that is difficult to interpret

on a forensic workstation, or to observe the behavior of malware in a unique environment. When this becomes necessary, a few methods are available that will ensure the integrity of the original evidence. The first option is to let a copy of the image boot in a virtual machine. Nearly all virtualization applications provide a means to convert a true disk image to their native format. Depending on the desired outcome and the evidence's operating system, using a virtual machine may not work. If the original system is available, an examiner can restore the hard drive image to another of equivalent size, and allow the restored image to start. Another option is to use specialized forensic hardware that leaves the restored drive in a read-only state while caching all writes to another medium. When examination is complete, the system is powered down, and the restored drive remains identical to the original image. Any processes should be thoroughly documented in your investigator's notes.

CHAPTER 11

1. Based on the information presented in this chapter, what do you think is the most challenging part of the analysis process? Explain why.

   This answer can be subjective, depending on your experience in the field. For some, the preparation for analysis is difficult. Getting access to data and structures that are not easily processed by the common forensic suites can be a challenge. In some situations, an analyst must consult a number of references, including developer documentation or open-source projects, to understand, parse, and search the data. In other situations, it is data minimization and learning the ability to identify when sampling and statistical analysis are appropriate techniques that do not reduce the accuracy of the findings.

2. Given the following scenario, explain how you would proceed. An investigation stakeholder tells you that one of the most critical objectives is to prove that a file with a specific MD5 hash was not present on a system at the time of analysis. You have a recent forensic disk image for the system.

   The first step is to thoroughly understand what the stakeholder is attempting to answer. The simple answer is to generate MD5 hashes of every file on the file system and then perform a search. The result may not be accurate, though. Determine where the given hash originated as well as the reliability of that source. For example, if the hash was from a dropper Trojan identified by an IDS product, it may not make sense to perform a hash search because the executable may remove itself post-execution. Perhaps the hash came from an analysis of the Trojan, and belonged to a second-stage loader that was never successfully retrieved in your environment. In both situations, the malicious situation detected by the IDS would be present, but a simple MD5 hash search would have cleared the system.

3. List four common types of text encoding. Explain how you would perform an effective keyword search if your source data was encoded with them.

   This question asks for text-encoding (or data-encoding) examples. Note that ASCII, UTF, and similar schemes apply to character encoding. When searching, regardless of encoding, always perform validation tests to ensure the tool and encoding scheme are performing as you expect.

   - **Base16** This type of encoding is simply the hexadecimal representation of data. To perform a search for strings that may be stored in an alphanumeric representation, perform a character-by-character conversion of your search term.
   - **UUencoding** The uuencode function is similar to a stream encoder, where three of the original bytes are taken at a time, split into four 6-byte groups, and shifted into the printable ASCII range. Because of the stream-like nature of the encoder, determining the alignment of the original text may be difficult. When performing searches, it is best to identify blocks of uuencoded data, decode them, and perform a search as you normally would. Uuencoded data is very easy to identify.
   - **Base64** As with uuencode, this function is similar to a stream encoder, and the most effective strategy is to decode blocks of data and search the result.
   - **URL encoding** Also known as "percent encoding," this encoding scheme is most commonly found in environments that use the MIME type application/x-www-form-urlencoded, such as in web server logs and network captures. RFC 3986 is the primary specification that includes this encoding scheme. It is a character-based scheme, so a simple conversion of the terms prior to a search is typically sufficient to yield expected results.

4. A manager at another office lets you know to expect a disk image via courier within the next day. You are tasked with recovering deleted files. What questions would you ask before the image arrives? Why?

   One of the main themes of this chapter is that, as an examiner, context is essential for any task that you are asked to complete. If you are leading the investigation, never leave doubt in your analysts' minds on the core elements of the inquiry. The first questions should pertain to the source of the data and how it has been handled. Ensure that evidence control has been established and you are aware how the "disk image" was generated. A few questions are listed here:

   - When was the image gathered?
   - When was the matter or incident discovered?
   - What actions were taken by the people who first responded to the system?
   - When did the suspected deletions take place?
   - What was deleted?

- Are other versions of the deleted files available?
- If the data is nonstandard, are samples of similar data available?

CHAPTER 12

1. Which attribute within the Master File Table (MFT) contains timestamps that cannot be directly manipulated through the Windows API?

   The $FILE_NAME attribute timestamps cannot be directly manipulated through the Windows API, whereas the $STANDARD_INFORMATION timestamps can be. However, note that Windows can still update the $FILE_NAME timestamps upon certain conditions, such as when files are moved or copied. The chapter discusses how attackers can use "double time modification" to exploit these conditions.

2. What NTFS artifact, specific to a directory, may record metadata for deleted files?

   Slack space within a nonresident $INDX attribute for a directory may contain the size and Modify/Access/Create timestamps for deleted files previously stored in this path.

3. What criteria must be met for a file to be resident within the MFT?

   The size of the file at the time of its initial creation must be under 800 bytes—small enough to fit in the $DATA attribute within the MFT record itself.

4. How could an attacker load and execute malicious code from an Alternate Data Stream (ADS) on a Windows 7 system?

   Although Windows 7 prevents the direct execution of code stored within an ADS, it is still possible to load batch files, Visual Basic, and PowerShell scripts, from an ADS.

5. What filenames can be obtained by parsing the contents of a Prefetch file for a given application?

   The Prefetch file can contain a listing of files written to or loaded by the application within the first ten seconds of execution. This typically includes the application's own executable file, supporting libraries, and most input or output files.

6. An attacker connects to an unprotected WinVNC screen-sharing service on a workstation. The current user is logged out, so the attacker must supply valid Windows credentials to proceed. What will the logon type be for this authentication?

   The log would contain an entry for a Logon Type 2 (interactive). Due to the screen-sharing application, the logon would appear to originate from the "console."

7. An attacker mounts the ADMIN$ share on a remote server in order to remotely schedule a task via the "at" command. What logon type would this activity generate?

   The log would contain an entry for a Logon Type 3 (network), recorded on the target system. If the attacker supplied a domain account's credentials, the authoritative domain controller would also record the logon event.

8. What source of evidence records the username responsible for creating a scheduled task?

   The Task Scheduler Operational event log "Microsoft-Windows-TaskScheduler/Operational.evtx" records an Event ID 106 event upon the registration of a scheduled task. This event includes the username that created the task.

9. What distinguishes evidence in the Shim Cache from Prefetch files?

   The Shim Cache can record the presence of executable files that have not executed on a system. In addition, it can record the presence of non-executable files, such as Visual Basic scripts, that the operating system also processes for the Shim Cache. Finally, the cache may record the file's Standard Information Last Modified timestamp, depending on the version of Windows.

10. What registry keys can record the directories accessed by a user via Explorer during an interactive session? What metadata is stored within the values of these keys?

    The ShellBag keys, stored in HKEY_USERS\{SID}_Classes\Local Settings\Software\Microsoft\Windows\Shell\ on Windows Vista and later, can track directories accessed by a user through Explorer during an interactive session. By parsing the contents of BagMRU values, an analyst can obtain the Standard Information Modified, Accessed, and Created timestamps for each path tracked by the ShellBags.

11. What distinguishes the evidence in UserAssist keys from that stored in MUICache keys?

    Both UserAssist and MUICache track application executed by a user through Windows Explorer. However, UserAssist includes different metadata—the number of times each program ran and a Last Executed timestamp. The MUICache only tracks the FileDescription data from a Portable Executable file's Version Information resources section.

12. Name two persistence mechanisms that do not require the use of the registry.

    DLL load order hijacking, recurring scheduled tasks, and user startup folders (such as C:\ProgramData\Microsoft\Windows\Start Menu\Programs\Startup) all provide the ability to automatically execute malicious code without changes to the registry.

CHAPTER 13

1. In an investigation, a network administrator notifies you that a Mac OS X system appears to be generating significant traffic that doesn't correlate with normal activity. How would you identify the source of the traffic on the system?

   The first step would be to obtain packet traces from the network administrator. If a full packet trace cannot be captured, obtain the connection information. Assuming the system is live, you may first start by reviewing the network sockets that are in use at the time of collection with the lsof utility. If the ports from the connection information are not in use, determine whether a user may have been present at the time of the traffic. Review the log data in /var/log/system.log, and if they are enabled, the launchd logs in /var/log. Either data source may have additional leads.

2. What would be the top files you would want to obtain during a live response to augment the data sources discussed in Chapter 7?

   The top files that we collect during an IR include data in the following locations. These files contain significant information on system activity and system configuration. Depending on the scope of the live response, directory listings of startup folders, additional log files, and software installation receipts may be useful.

   - /var/log/system/log
   - /var/log/authd.log
   - /var/log/wifi.log
   - /var/log/asl/*
   - /var/db/dslocal.*

3. The Unix touch command allows one to update a file's access and modification times. What other sources of time-based data could one use to establish whether manipulation occurred?

   The first location to examine would be the file system for additional timestamps generated and maintained by automated processes. Has the file been indexed by Spotlight? Has it been tagged as a version in Managed Storage? The next location to review would be within the file itself. Are internal data structures present that provide reliable timestamps?

4. Name a few persistence mechanisms that an attacker can use to maintain access to a compromised Mac OS X system. What live response commands can you use to automate a review for those mechanisms?

   - Plist files in LaunchAgents or LaunchDaemon directories

During a live response, gather all files in the LaunchAgent and LaunchDaemon directories under /System/Library, /Library, and the Library directories for each user account.

- Plist files in StartupItems directories

  During a live response, gather all files in the StartupItems directories under /System/Library and /Library.

- Traditional cron jobs

  Review the contents of /usr/lib/cron.

- Kernel extensions

  Review the kext bundles in /Library/Extensions. Detection of this persistence mechanism is a bit complicated to automate unless a baseline has been established for the system under review.

- Login items

  Review the com.apple.loginitems.plist files in each user's Library/Preferences directory.

Patrick Wardle (@patrickwardle), a security researcher, has compiled a great synopsis of current OS X persistence techniques.

CHAPTER 14

1. During an investigation, a database administrator discovers that hundreds of SQL queries with a long execution time were run from a single workstation. The database did not record the specific query—just the date, time, source, and query execution time. You are tasked to determine what queries were executed. How would you proceed?

   The evidence may reside in three locations, depending on how well the environment was configured before the incident. On the system that issued the SQL query, examine swap files, output files, shell or command history files, any history retained by the application used to connect to the database, and if the event was recent, system memory itself. On the server, assuming that the standard logging facility only captured the information just listed, the only hope for recovery of the queries may be debug or development logs. Finally, if network monitoring is in place, the SQL query may be present in network captures.

2. You begin investigating a Windows 7 computer, and quickly notice that there is no C:\Users directory. Assume the data was not deleted. Provide a realistic explanation of why there is no C:\Users directory, and how you can get access to the user data.

One common reason for the absence of data in C:\Users is that the user or administrator moved the user's home directories to another volume. It may occur when the profile requires more drive space, for performance reasons, or for simpler OS image management for the administrators. When a profile directory is moved, you can locate it by examining the ProfileImagePath registry value in the following subkey:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\ProfileList\(user SID)

3. As you are investigating a Windows XP system, you find operating system artifacts that indicate the attacker ran heidisql.exe multiple times. Assuming the filename is the original name of the program, what do you suspect the application was? What registry key or keys would you examine to determine additional information about what the attacker did?

   HeidiSQL is an open-source SQL tool. Once installed, it maintains configuration information in the registry. When you are faced with a tool that is unfamiliar, it is wise to examine it in a VM, and then monitor the interaction with the file system and registry to understand how it stores configurations and recently accessed data. HeidiSQL stores server profiles and a query history for each under the registry subkey HKEY_CURRENT_USER\Software\HeidiSQL.

4. You are investigating what artifacts an application creates on a Windows 8 system. You notice that under the AppData directory, the application is storing information in multiple directories: Local, LocalLow, and Roaming. Why would the application do that? What is the purpose of the different AppData directories (Local, LocalLow, and Roaming)?

   The purpose of each AppData directory is described in the Microsoft Roaming User Deployment Guide. Data that remains on the local system is stored in Local or LocalLow. LocalLow is a special case of Local, where data stored under a lower integrity level is placed. Roaming contains a copy of the data fetched from the server during logon.

CHAPTER 15

1. What are the primary differences between static and dynamic analysis techniques? What are the advantages of each?

   The primary difference between static and dynamic analysis is whether the program being examined is allowed to execute. When static analysis is being performed, the potentially malicious code is not run, leaving the interpretation of basic information and instructions to the examiner. Dynamic analysis has the potential to reveal additional important information because the code is allowed to interact with local and potentially remote systems.

2. In building a safe triage environment, what are some common configuration choices you should make? If you were to create such an environment, what utilities would you prepare for use in both a static and a dynamic analysis environment?

   A safe triage environment is sufficiently isolated to prevent malicious code from affecting or communicating with systems outside of the testing environment, so the first choice would be to use virtualization software. To perform static analysis, utilities that extract a binary file's strings and functions are required. It can also be advantageous to make use of AV scanners to determine whether the unknown binary has already been examined. For the advanced examiners, a disassembler and good OS library reference guides are required. Dynamic analysis is performed with debuggers as well as file, process, network, and registry monitors.

3. On a system you are investigating, you discover a file that you suspect is malware. How would you proceed to safely obtain a copy of the malware and place it in your triage environment? Address investigating both a live system and a forensic image of a hard drive.

   In both situations, the goal is identical. You need to obtain a copy of the suspect binary and transfer it into your analysis environment without putting any other environment at risk. Whether you extract the binary from a forensic image or copy it from a live system, ensure that it cannot be accidentally run. This may involve changing the file's extension, clearing execute bits, or immediately compressing the binary in a TAR or ZIP file. We've found that it is best to use encryption, if it is part of the compression suite. This will prevent AV or mail quarantine systems from removing the data. Finally, transfer the compressed file to the analysis virtual machine. Typically, the process of a drag-and-drop copy into an analysis VM is best because network shares are not exposed to the analysis VM.

4. Given the following information, describe what you would do next to determine more about a file you are looking at:

   a. An MD5 hash lookup in major databases comes up negative.

   b. Antivirus scans do not indicate a threat.

   c. The file contains few legible strings, most of which are uninformative.

   With very little information learned from static analysis techniques, begin to examine the binary in a dynamic environment. In a virtual machine, allow the executable to run while monitoring all file and network activity. Note the libraries that it imports and the functions called and parameters that are passed. The next step would be to examine the program flow in a debugger. Depending on the author of the executable, you may need to identify and work around code that is designed to frustrate reverse-engineering efforts.

5. You are performing dynamic analysis of a file in Windows, and the output from Process Monitor shows:

   a. The malware calls the SetWindowsHookEx function.

   b. A file named bpk.dat is created in the user's application data temporary folder.

   c. The file is written to at a seemingly random interval, but appears to happen more often when the system is in use.

   You inspect the file content and find unreadable binary data. What theories might you propose about the file's functionality?

   If the process registers a callback to SetWindowsHookEx, you can expect that the potential malware is performing some function on an event. Given the output file that grows over time, as the system is in use, it is likely to be capturing keystrokes. To validate that theory, examine the behavior of the potential malware running in a debugger, or examine it in a disassembler.

CHAPTER 16

1. You analyze an image of a hard drive to locate deleted RAR files. You complete your analysis and do not find any deleted RAR files. Your boss tells you not to write a report. What would you do and why?

   Keep in mind that the decisions made during an investigation or an IR may have unintended consequences at a later time. They could be motivated by a sense of urgency, desire by management for a particular outcome, or worse. During the formation of your incident response unit, generate several reporting templates that can be used for large and small incidents alike and mandate their use. If that process fails due to external pressure, however, you can ensure that the integrity of your work is preserved by completing case notes and compiling them into a single location for backup purposes.

2. Explain why active voice is the preferred writing style of technical reports. Provide at least three clear examples that illustrate active voice versus passive voice.

   One of the primary reasons for using active voice is that it helps to structure sentences in a concise way. The subject clearly performs an action.

   Examples of passive voice:

   - The hard drive image was reviewed by Mandiant for evidence of data theft.
   - The connection to the remote web server was initiated by the backdoor Trojan at 23:41 on August 10, 2013.
   - During the periods of activity in October, the SQL server was queried 452 times.

Examples of active voice:

- Mandiant reviewed the hard drive image for evidence of data theft.
- The backdoor Trojan initiated a connection to the remote web server at 23:41 on August 10, 2013.
- The attackers queried the SQL server 452 times during the month of October.

3. Design two metadata tables that were not discussed in this chapter. Explain why you chose the layout and the fields you included.

   Providing the metadata associated with a data source is an essential practice. On our teams, we mandate that if data is presented from a file on a system, a minimum amount of data will be provided to assist a reader in locating the source. For example, this table shows file metadata for a file on the EXT4 file system. The number of fields requires two lines per record.

File Path		Permissions	Owner	Group
MD5 Hash	Access Time	File Modified	Inode Change	File Size
/root/.bash_history		-rw-------	root	root
8d8724c73edf2dddb 457f36cad221457	04/09/13 09:12:07	04/09/13 09:12:07	04/09/13 09:12:07	1131

   The next table shows the metadata for a file on an NTFS file system. Typically, when the details are relevant to the situation that is being presented, we'll include another table that lists the ACLs and ownership information.

File Name	File Path				
MD5 Hash	File Created	Last Written	Last Accessed	Entry Modified	File Size
keys.db0	C:\Windows\System32				
e271828182845904 5235360287471352	01/23/11 19:25:33	01/23/11 19:25:33	03/08/13 13:11:27	03/08/13 13:11:27	23,805

4. During an analysis, you discover what appears to be credit card numbers in a file. You provide an excerpt of the file in a figure. Are there any special considerations you should take as to how the data is presented in the figure?

   During any investigation where sensitive information is found to be accessed by unauthorized parties, additional accidental exposure by the response team is the last thing anyone wants to do. In any situation where you have PCI, PII, HIPAA, or other data that is protected by regulation or policy, you will need

to define a scheme for representing a cleansed sample in a report as well as a means for securely storing and transmitting raw data.

CHAPTER 17

1. List at least five of the factors critical to remediation efforts.

   The list should contain five of the following seven factors: incident severity, remediation timing, the remediation team, technology, budget, management support, and public scrutiny.

2. List at least three of the five qualities essential for a strong remediation lead.

   The list should contain three of the following five qualities: in-depth understanding IT and security, focus on execution, understanding of internal politics, proven track record of building support for initiatives, and the ability to communicate with technical and nontechnical personnel.

3. Should the remediation team include members from the investigation team? Why or why not?

   Yes. A member from the investigation team will have a better understanding of the attacker's tools, tactics, and procedures as well as how the investigation into the activity is progressing. The member from the investigation team should also be able to provide solid posturing recommendations to improve the organization's logging and monitoring stance as it relates to the current incident.

4. Give at least five examples of personnel who should be members of the remediation team.

   Any five of the following personnel types are correct: investigation team member, system architects, network architects, application developers, subject matter experts, internal or external legal counsel, compliance officers, business line managers, human resources public relations, and executive management.

5. Define posturing actions. Define containment actions. What are some differences between the two types of actions?

   Posturing actions are taken during an ongoing incident and are designed to be implemented while having little impact on the attacker. *Containment actions* prevent the attacker from performing a specific action that the organization cannot allow to continue. Some differences are that posturing actions are specifically designed not to alert the attacker and to improve visibility and security of the environment. Containment actions are designed to be invasive (and often draconian in nature) and directly interfere with the attacker's ability to perform some type of action. In addition, posturing actions are often

used when gathering intelligence about the attacker and his activities is more important than stopping the attacker.

6. What are some common reasons eradication events fail?

   The single biggest reason eradication events fail is improper planning. Other common reasons that eradication events fail include the following: eradicating too soon, eradicating too late, failing to implement all planned eradication actions, failing to be comprehensive (such as not cleaning or rebuilding one compromised system), and failing to verify that all eradication actions were performed properly.

7. Explain the mean time to remediate concept. Is this a useful concept?

   The *mean time to remediate* is the time between the discovery of an incident and the complete eradication of the threat. This concept is extremely useful because it allows management to determine how successful its incident response process and team is. In general, the lower the mean time to remediate, the less damage an attacker is able to do to an environment.

8. Give an example of a remediation effort where a combined action approach would be more valuable to the incident response effort than either an immediate or delayed action approach.

   An example of any scenario where there is near real-time data theft, especially financial data or actual money, is correct. A combined action is the preferred approach in these cases because the organization needs to immediately prevent the attacker from stealing any more data/money, and the organization needs to perform a comprehensive investigation and remediation to ensure the attacker is properly eradicated from the environment.

9. Pick a recent public breach disclosure and create a containment plan based on the known information. What actions would you implement as part of the eradication effort that you would not implement as part of the containment plan? Explain your choices.

   This exercise is left to the reader.

10. Create a complete remediation plan, including posturing/containment, eradication, and strategic recommendations, for one of the two case studies presented in Chapter 1.

    This exercise is left to the reader.

CHAPTER 18

1. List at least three other strategic recommendations you think should be proposed for the victim organization from Case Study #1.

   This answer is subjective; there are many answers that can be considered correct. Any answer that improves the security of the environment and is not included in the ten strategic recommendations already listed should be considered correct. The intent of the exercise is to force the reader to consider additional ways to enhance the security of the environment.

2. Do you think implementing an immediate containment plan was the right course of action in this case study? Why or why not?

   Yes, implementing an immediate containment plan was the right course of action in this situation. An immediate containment plan was the right course of action because the victim organization needed to stop the attacker from continuing to steal card holder data as soon as possible. Removing the attacker from the overall environment was a secondary goal because the attacker was focused on stealing data from the restricted financial environment only; the corporate environment did not contain data the attacker was interested in.

3. What are three of the four main goals of an eradication event?

   Any three of the following four goals are correct: to remove the attacker's ability to access the environment; to deny the attacker access to compromised systems, accounts, and data; to remove the attack vector the attacker used to gain access to the environment; and to restore the organization's trust in its computer systems and user accounts.

4. Should nontechnical members of the remediation team provide input on technical recommendations?

   Nontechnical members of the remediation team should be allowed to provide input on technical recommendations because they will understand the nontechnical implications actions have for the business. For example, a nontechnical business line owner may not understand the technical details required in order to implement a change on one of his systems, but he will understand the impact to the business that the change will have. Therefore, it is critical that he comment on the proposed action to ensure that the business is okay with the change.

5. Build a complete remediation plan for Case Study #2. The remediation plan only needs to list the high-level action, not the detail associated with implementing each activity. Provide and explain each recommendation.

   This exercise is left to the reader.