1 Introduction to Network Forensics

In this chapter, you will learn about:

- What network forensics is
- Evidence handling standards
- Verification of evidence

Sitting in front of his laptop he stares at a collection of files and reflects on how easy it was to get them. He sent an e-mail to a sales manager at his target company—almost silly how obviously fake it was—and within minutes he knew that he had access to the sales manager's system. It took very little time for him to stage his next steps, which included installing a small rootkit to keep his actions from being noticed, and to ensure his continued presence on the system wouldn't be detected. It also provided him continued access without the sales manager needing open the e-mail message again. That had taken place weeks back and so far, there appeared to be no evidence that anyone had caught on to his presence not only on the system but, by extension, on the business network the sales manager's laptop was connected to.

It was this network that he was poring over now, looking at a collection of files related to the business's financial planning. There were also spreadsheets including lists of customer names, contact information, and sales projections to those customers. No really big score but definitely some interesting starting points. Fortunately, this user was well-connected with privileges in the enterprise network. This ended up giving him a lot of network shares to choose from, and for the last several weeks he has been busy looking for other systems on the network to take over. Getting access to the address book on this system was really helpful. It allowed him to send messages looking as though they came from this user, sending co-workers to a website that would compromise their systems with some client software, adding them to the growing botnet he had control over. File shares were also good places to not only get documents to make use of, but also to drop some more infected files. The key loggers that were installed have generated some interesting information and keeping an eye on all of that is an ongoing project.

Ultimately, this is becoming quite a little stronghold of systems. It's not exactly the best organization he's been in with respect to quality data from an intellectual property or large caches of credit card numbers or even health care information. However, having more systems to continue building the botnet is always good

COPYRIGHTED MATERIAL

and at some point months or even years down the road, more interesting information may show up. In the meantime, there may be vendors who have trust relationships with this network that could be exploited.

Once inside the network, he has so many potential places to go and places to probe. There is a lot of data to be found and even though it appears that disk encryption is being used fairly consistently across the organization, all of that data is accessible to him as an authenticated user on the network. Wiping logs in places where they were turned on was trivial. This little network was all his for the taking for apparently as long as he felt it would be useful.

Does this sound scary at all to you? In reality, this is far too common and although it's dramatized, it's not that far off from how networks become compromised. Not long ago, technical intrusions were more common than the type of attack just described. In a *technical intrusion*, attackers use software vulnerabilities to get into a system remotely. This type of attack targets servers sitting in a data center because those are exposed to the outside world. That's not the case anymore. As we continue to learn, attackers are using people to get into systems and networks. This was vividly illustrated in 2013 in Mandiant's report, "APT1: Exposing One of China's Cyber Espionage Units" (`https://www.fireeye.com/content/dam/fireeye-www/services/pdfs/mandiant-apt1-report.pdf`). Attackers send e-mail with malicious attachments, get someone to visit a website, or just simply park malicious software on a known website and wait for people to visit in order to infect their systems. Unfortunately, this is the world we now live in, a world where companies who haven't had systems compromised are becoming the minority rather than the majority.

This is one reason forensics is becoming such a hot skill to have. Well, that and the fact that the folks on various TV shows make it seem really cool, interesting, and easy. The reality is a different story, of course. Although the news and other media outlets make it seem as though attacks are carried out by solo *hackers* (an ambiguous and misleading word), the majority of outside attacks businesses are subject to today are perpetrated by well-funded and organized criminal enterprises. There is money to be made from these crimes; criminals are starting to use ransom and extortion to go directly for the money rather than trying to steal something to sell off later on.

The term *forensics* can be ambiguous. Because of that, it's helpful to have an understanding of what forensics currently is and isn't. Particularly when it comes to network forensics, it's more and more becoming part of incident response. Digital forensics practitioners have to be capable of more than locating images and deleted files that may be common for the large volume of child pornography cases that traditional law enforcement practitioners may be looking for. Sometimes, knowing how to extract files from a computer system isn't enough because information can be obscured and deleted very effectively. Certainly operating system forensics is important, but sometimes it takes more than just understanding what happened on the system itself.

Network forensics is becoming an extremely important set of skills when it comes to situations like the one described at the beginning of the chapter. Rather than relying on what the operating system and disks may be able to tell you, a network forensic investigator can go to the network itself and collect data of an attack in progress or look up historical information that may be available after

a company has suffered a security breach with someone taking up long-term residence, someone who has the ability to observe and piece together what they see into a coherent picture. This coherent picture may include information from other sources such as firewalls, application logs, antivirus logs, and a number of other sources.

One advantage to watching the network is that the network can't lie. Applications may not be doing what they are supposed to be doing. Logs may not be available or they may have been wiped. There may be root kits installed to obscure what is happening on a system. Once a network transmission is sent out on the wire, though, the bits are the bits.

Because of situations like the one described in the chapter-opening scenario, it's important to know exactly what forensics is as a practice as well as its role in incident response. Finally, there is a need for not only forensic practitioners in general because of the large number of incidents that occur in businesses around the world, but specifically, there is a need for network forensic practitioners.

What Is Forensics?

Before going further, let's define some terms.

The word *forensics* comes from the Latin *forens*, meaning belonging to the public. It is related to the word *forum*. If you have ever been involved in debate teams, you may be familiar with it as being related to debate and argumentation. If you are skilled in forensics, you may make a good lawyer. It is from this sense that the connotation of the word has come to mean something other than debate and argumentation. Investigating evidence, in the field or in the lab, to be used in a court case is the practice of forensics because the activity is related to the courts or trials.

This chapter expands on that by talking more specifically about *digital forensics*. Computer or *digital forensics* is the practice of investigating computers, digital media, and digital communications for potential artifacts. In this context, the word *artifact* indicates any object of interest. We wouldn't use the word *evidence* unless it's actually presented as part of a court case. You may say that an artifact is potential evidence. It may end up being nothing, but because it was extracted from the piles of data that may have been handed to the investigator, we need to refer to it in a way that makes clear the object is something of interest and potentially warrants additional investigation.

Because the word forensics is used in legal settings, you will often find that talk about forensics is involved with law enforcement. Traditionally, that has been the case. However, because many of the techniques and skills that are used by law enforcement are the same as those that may be practiced by an incident response specialist—someone who is investigating a suspicious event or set of events within a business setting—the word *forensics* also describes the process of identifying digital artifacts within a large collection of data, even in situations where law enforcement isn't involved.

For our purposes, the data we are talking about collecting is network information. This may be packet captures, which are bit-for-bit copies of all communication that has passed across a network

interface. The data collected may also come in the form of logs or aggregated data like network flow information.

Any time you handle information that could potentially be used in a court case, it's essential that it be maintained in its original condition, and that you can prove that it hasn't been tampered with. There are ways to ensure that you can demonstrate that the evidence hasn't been tampered with, including maintaining documentation demonstrating who handled it. Additionally, being able to have verifiable proof that the evidence you had at the end is the same as at the beginning is important. The reason for this is that in a course case , technical evidence, such as that from a digital forensic examination, is expected to adhere to an accepted set of standards.

Handling Evidence

The United States of America uses a *common law* legal system. This is at the federal as well as the state level, with the exception of the state of Louisiana, which uses a *civil law* system. The United Kingdom also uses a common law system. This means that legislatures enact laws and those laws are then interpreted by the courts for their applicability to specific circumstances. After a court has issued a ruling on a case, that case can then be used as a precedent in subsequent cases. This way every court doesn't have to make a wholly original interpretation of a law for every case. They build on previous cases to create a common interpretation of the law.

When it comes to addressing technical evidence in court cases, a couple of cases are worth understanding. The first case, Frye vs. United States, was a case in 1923 related to the admissibility of a polygraph test. As we continue to make technological advances, courts can have a hard time keeping up. The Frye standard was the one of the first attempts to codify a process that could help ensure that technical or scientific evidence being offered was standardized or accepted within the technical or scientific community. The courts needed a way to evaluate technical or scientific evidence to ensure that it was able to help the trier of facts determine the truth in a trial.

In essence, the Frye standard says that any scientific or technical evidence that is presented before the court must be generally accepted by a meaningful portion of the community of those responsible for the process, principle, or technique being presented. Acceptance by only a small number of colleagues who are also working in a related area doesn't necessarily rise to the standard of general acceptance by the community. Scientific evidence such as that resulting from DNA testing or blood type testing has passed this standard of reliability and veracity and is therefore allowed to be presented in a trial.

The federal court system and most U.S. states have moved past the Frye standard. Instead, they rely on the case Daubert vs. Merrell Dow Pharmaceuticals, Inc. Essentially, the standard of determining whether scientific or technical evidence is relevant hasn't changed substantially. What the majority opinion in the Daubert case argued was that because the Federal Rules of Evidence (FRE) were passed in 1975, those should supersede Frye, which was older. The Supreme Court ruled that in cases where the FRE was in conflict with common laws, such as the standard set by Frye, the FRE had precedence.

The intention of the continuing progress of case law related to technical evidence is to ensure that the evidence presented can be used to assist the trier of facts. The role of the trier of facts in a court case is to come to the truth of the situation. Frye was used to make sure technical evidence was accepted by a community of experts before it could be considered admissible in court. Daubert said that because the Federal Rules of Evidence came later than Frye, it should become the standard in cases of technical evidence. While expert witnesses are used to explain the evidence, the expert witness alone is not sufficient. The witness is a stand-in at trial for the evidence. A witness can be questioned and can provide clarifying information that the evidence directly cannot.

When it comes to digital evidence, we have to consider issues related to the appropriate handling of the data because it can be easily manipulated. For that reason, there's a risk that digital evidence could be considered hearsay if it's mishandled because of the FRE requirements regarding hearsay evidence. Hearsay is relevant here because *hearsay* is any evidence that is not direct, meaning that it doesn't come from a primary source that can be questioned by the opposition. In short, because there isn't someone sitting on the stand indicating what they saw, it's potentially hearsay unless it is a recording of regular business activities. Of course, the legal aspects are much more complicated than this short discussion might imply, but those are the essentials for those of us without law degrees.

All of this is to say that we have to handle potential evidence carefully so it cannot be questioned as being inauthentic and an inaccurate representation of the events. Fortunately, there are ways that we can not only demonstrate that nothing has changed but also demonstrate a complete record of who has handled the evidence. It is essential that when evidence has been acquired that it be documented clearly from the point of acquisition using the techniques outlined in the following sections.

Cryptographic Hashes

The best way to demonstrate that evidence has not changed from the point of acquisition is to use a cryptographic hash. Let's say, for example, that you have an image of a disk drive that you are going to investigate. Or, for our purposes, what may be more relevant is to say that we have a file that contains all of the network communications from a particular period of time. In order to have something we can check against later, we would generate a cryptographic hash of those files. The cryptographic hash is the result of a mathematical process that, when given a particular data set as input, generates a fixed-length value output. That fixed-length value can be verified later on with other hashes taken from the same evidence. Because hashing a file will always generate the same value (that is, output), as long as the file (the input data) hasn't changed, courts have accepted cryptographic hashes (of sufficient complexity) as a reliable test of authenticity when it comes to demonstrating that the evidence has not changed over a period of time and repeated interactions.

Two separate sets of data creating the same hash value is called a *collision*. The problem of determining the collision rate of a particular algorithm falls under a particular probability theory called the *birthday paradox*. The birthday paradox says that in order to get a 50% probability that two people in a given room have the same birthday, month and day, all you need is to have 23 people in the room. In order to get to 100% probability, you would need 367 people in the room. There is a very slim potential

for having 366 people in a room who all have a different birthday. To guarantee that you would have a duplicate, you would need to have 367 (365 + 1 for leap day + 1 to get the duplicate). This particular mathematical problem has the potential to open doors for attacks against the hash algorithm.

When you hear *cryptographic*, you may think *encryption*. We are not talking about encrypting the evidence. Instead, we are talking about passing the evidence through a very complicated mathematical function in order to get a single output value. Hashing algorithms used for this purpose are sometimes called *one-way functions* because there is no way to get the original data back from just the hash value. Similarly, for a hash algorithm to be acceptable for verifying integrity, there should be no way to have two files with different contents generate the same hash value. This means that we can be highly confident that if we have one hash value each time we test a file, the content of that file hasn't changed because it shouldn't be possible to make any change to the content of the file such that the original hash value is returned. The only way to get the original hash value is for the data to remain unaltered.

NOTE A cryptographic hash takes into consideration only the data that resides within the file. It does not use any of the metadata like the filename or dates. As a result, you can change the name of the file and the hash value for that file will remain the same.

NOTE Cryptography is really just about secret writing, which isn't necessarily the same as encryption. Hashes are used in encryption processes as a general rule because they are so good at determining whether something has changed. If you have encrypted something, you want to make sure it hasn't been tampered with in any fashion. You want to know that what you receive is exactly what was sent. The same is true when we are talking about forensic evidence.

For many years, the cryptographic hash standard used by most digital forensic practitioners and tools was Message Digest 5 (MD5). MD5 was created in 1992 and it generates a 128-bit value that is typically represented using hexadecimal numbering because it is shorter and more representative than other methods like printing out all 128 binary bits. To demonstrate the process of hashing, I placed the following text into a file:

Hi, this is some text. It is being placed in this file in order to get a hash value from the file.

The MD5 hash value for that file is 2583a3fab8faaba111a567b1e44c2fa4. No matter how many times I run the MD5 hash utility against that file, I will get the same value back. The MD5 hash algorithm is non-linear, however. This means that a change to the file of a single bit will yield an entirely different result, and not just a result that is one bit different from the original hash. Every bit in the file will make a difference to the calculation. If you have an extra space or an end of line where there wasn't one in the original input, the value will be different. To demonstrate this, changing the first letter of the text file from an H to a G is a single-bit difference in how it is stored on the computer since the value for H is 72 and the value for G is 71 on the ASCII table. The hash value resulting from this

altered file is 2a9739d833abe855112dc86f53780908. This is a substantive change, demonstrating the complexity of the hashing function.

NOTE MD5 is the algorithm but there are countless implementations of that algorithm. Every program that can generate an MD5 hash value contains an implementation of the MD5 algorithm.

One of the problems with the MD5 algorithm, though, is that it is only 128 bits. This isn't an especially large space in which to be generating values, leading it to be vulnerable to collisions. As a result, for many purposes, the MD5 hash has been superseded by the Secure Hash Algorithm 1 (SHA-1) hash. The SHA-1 hash generates a 160-bit value, which can be rendered using 40 hexadecimal digits. Even this isn't always considered large enough. As a result, the SHA-2 standard for cryptographic hashing has several alternatives that generate longer values. One that you may run into, particularly in the encryption space, is SHA-256, which generates a 256-bit value. Where the 128-bit MD5 hash algorithm has the potential to generate roughly 3.4 × 10^38 unique values, the SHA-256 hash algorithm can yield 1.15 × 10^77 unique values. It boggles the mind to think about how large those numbers are, frankly. Generating a SHA-1 hash against our original text file gives us a value of 286f55360324d42bcb1231ef5706a9774ed0969e. The SHA-256 hash value of our original file is 3ebcc1766a03b456517d10e315623b88bf41541595b5e9f60f8bd48e06bcb7ba. These are all different values that were generated against the same input file.

One thing to keep in mind is that any change at all to the data in the source file will generate a completely different value. Adding or removing a line break, for example, would constitute removing an entire character from the file. If that were done, the file may look identical to your eyes but the hash values would be completely different. To see the difference, you would have to view the file using something like a hexadecimal editor to see how it is truly represented in storage and not just how it is displayed.

You can use a number of utilities to generate these values. The preceding values were generated using the built-in, command-line utilities on a Mac OS system. Linux has similar command-line utilities available. On Microsoft Windows, you can download a number of programs, though Microsoft doesn't include any by default. Microsoft does, however, have a utility that you can download that will generate the different hash values for you. The name of the utility is File Checksum Identity Verifier (FCIV).

Any time you obtain a file such as a packet capture or a log file, you should immediately generate a hash value for that file. MD5 hash values are considered acceptable in court cases as of the time of this writing, though an investigation would be more durable if algorithms like SHA-1 or SHA-256, which generate longer values, were to be used. MD5 continues to demonstrate flaws the longer it is used and those flaws may eventually make evidence verification from MD5 hashes suspect in a court case.

Over the course of looking at packet captures in Chapter 4, we will talk about some other values that perform similar functions. One of those is the cyclic redundancy check (CRC), which is also mathematically computed and is often used to validate that data hasn't been altered. These sorts of values, though, are commonly called checksums rather than hashes.

Chain of Custody

Sometimes it seems as though TV shows like *NCIS*, *CSI*, *Bones*, and others that portray forensics simultaneously advance and set back the field of forensics. Although some of the technical aspects of forensics, including the language, are ridiculous, these shows do sometimes get things right. This was especially true in the early days of *NCIS*, as an example, where everything they collected was bagged and tagged. If evidence is handed off from one person to another, it must be documented. This documentation is the *chain of custody*. Evidence should be kept in a protected and locked location if you are going to be presenting any of it in court. Though this may be less necessary if you are involved in investigating an incident on a corporate network, it's still a good habit. For a start, as noted earlier in this chapter, you never know when the event you are investigating may turn from a localized incident to something where legal proceedings are required. As an example, the very first well-known distributed denial of service (DDoS) attack in February 2000 appeared as a number of separate incidents to the companies involved. However, when it came time to prosecute Michael Calce, known as Mafiaboy, the FBI would have needed evidence and that evidence would have come from the individual companies who were targets of the attacks—Yahoo, Dell, Amazon, and so on.

Even in the case of investigating a network incident in a business setting, documenting the chain of custody is a good strategy. This ensures that you know who was handling the potential evidence at any given time. It provides for accountability and a history. If anything were to go wrong at any point, including loss of or damage to the evidence, you would have a historical record of who was handling the evidence and why they had it.

Keeping a record of the date and time for handing off the evidence as well as who is taking responsibility for it and what they intend to do with it is a good chain-of-custody plan. It doesn't take a lot of time and it can be very important. As always, planning can be the key to success, just as lack of planning can be the doorway to failure. The first time you lose a disk drive or have it corrupted and that drive had been handed around to multiple people, you will recognize the importance of audit logs like chain-of-custody documentation. Ideally, you would perform a hash when you first obtain the evidence to ensure that what you are getting is exactly what you expect it to be. You should have a hash value documented so you will have something to compare your hash to in order to demonstrate that no changes have occurred.

Incident Response

Incident response may be harder to get your head around if you are a forensic practitioner. If you are a system or network administrator trying to get your hands around the idea of forensics, incident response should be old hat to you. When networks belonging to businesses or other organizations (schools, non-profits, governmental, and so on) are subject to a malware infestation, as an example, that would probably trigger an incident response team to get the incident under control as well as

investigate the cause of the incident. Depending on who you talk to you may get different answers, but the process of incident response can be boiled down to four stages: preparation; detection and analysis; containment, eradication, and recovery; and post-incident activity.

What exactly is an incident? How does an incident differ from an event? This is another area where you may find that you get differing opinions depending on whom you talk to. Rather than getting into a deep discussion here, let's go with simple. An *event* is a change that has been detected in a system. This could be something as simple as plugging an external drive into a system. That will trigger a log message in most cases. That would be an event. Someone attempting to ping a system behind a firewall where the messages are blocked and logged may be an event. An event may even be updating system software, as in the case with a hot fix or a service pack.

An *incident*, on the other hand, is commonly something that is attributable to human interaction and is often malicious. An incident is always an event, because every incident would result in some sort of observable change to the system. If all of your web servers were infected by malware, that malware would be observable on the system. It would result in events on all of the systems and you would have an incident on your hands. A single system being infected with malware would be an event but wouldn't be enough to rise to a level where you would call an incident response team.

A forensic practitioner would obviously be necessary at the detection and analysis phase but they would typically be involved in the preparation stage as well. Over the course of the book, we will be going over some items that you may want to make sure are in place as an organization goes through preparation stages. Preparation is a very large category of activities, including planning, but from the standpoint of a forensic investigator, it is primarily when you make sure you will have what you need when it comes to doing an analysis. There may also be activity when it comes to eradication, to ensure that the source of the incident has been completely removed. Finally, a forensic investigator would be involved in post-incident activities for lessons learned and process improvement.

In most cases, you would have an incident response team, even if it is small and ad hoc, to deal with incidents because handling incidents is a process. The larger the organization and the more systems involved, the larger the incident response team would likely be. Creating a team up front would be another important activity when it comes to planning. Your organization, as part of the creation of security policies, standards, and processes, should create an incident response team or at least have documentation for how to handle an incident, should one occur. Considering that it's widely believed that a significant proportion of companies in the United States have been breached, meaning they have had attackers compromise systems to gain unauthorized access, "should one occur" is a bit euphemistic. In reality, I should say *when* an incident occurs. If you haven't had to deal with an incident, it may simply be a result of lack of appropriate detection capabilities.

Forensic practitioners are definitely needed as part of the incident response effort. They need not be full-time forensic practitioners, but simply people already employed at the company who happen to have the knowledge and skills necessary to perform a forensic investigation. They can get to the root cause of an incident, and that requires someone who can dig through filesystems and logs and look in other places within the operating system on the affected hosts.

Without understanding the root cause, it would be difficult to say whether the incident is under control. It would also be difficult to know whether you have found all of the systems that may be impacted because incidents, like unauthorized system access or malware infestations, will commonly impact multiple devices across a network. This is especially true when there is a large commonality in system deployments. In other words, if all systems are created from the same source image, they will all be vulnerable in the same way. Once an attacker finds a way into one, all of the others that have been built using the same image are easy targets.

The forensic investigator will need to be focused on identifying the source of the attack, whether it's a system compromise or a malware infection, to determine what may need to be addressed to make sure a subsequent, similar attack isn't successful. They will also need to be focused on finding any evidence that the attacker attempted to compromise or infect other hosts on the local network. If there is evidence of attempts against systems not on the organization's network, the incident response team should have the capability to reach out to other organizations, including a computer emergency response team (CERT) that may be able to coordinate attacks across multiple organizations.

This is where you may run into the need for the collected artifacts in a larger investigation and potential criminal action. Coordinating with law enforcement will help you, as a forensic investigator, determine your best course of action if there is evidence of either substantial damage or evidence that the attack involves multiple organizations. This is another area where planning is helpful—determining points of contact for local and federal law enforcement ahead of time for when an incident occurs.

The Need for Network Forensic Practitioners

In early 2016, a task force was assembled to talk about how to best approach educating more professionals who are capable of filling thousands of jobs that are expected to be available in the coming years. While this is generally referred to as a need for cybersecurity workers, the term *cybersecurity* is fairly vague and covers a significant amount of ground. The federal government alone is planning on large spending around making sure they can support a growing need for skilled and/or knowledgeable people to prevent attacks, defend against attacks, and then respond when an attack has been detected. The initial plan was to spend $3.1 billion to modernize and if the plan is implemented properly, there will continue to be a need for people who are capable of responding to incidents.

This is just at the level of the federal government. Large consulting companies like Mandiant and Verizon Business as well as the large accounting companies that are also involved in security consulting are hiring a lot of people who have skills or knowledge in the area of forensics. When companies suffer a large-scale incident, particularly smaller or medium-sized companies that can't afford full-time staff capable of handling a complete response, they often bring in a third party to help them out. This has several advantages. One of them is that a third party is less likely to make any assumptions because they have no pre-existing knowledge of the organization. This allows them to be thorough rather than potentially skipping something in the belief they know the answer because of the way "it's supposed to work." Hiring information technology people who are skilled in information security

and forensics can be really expensive. This is especially true for smaller companies that may just need someone who knows a little networking and some Windows administration.

Large companies will often have a staff of people who are responsible for investigations, including those related to digital evidence. This means that the federal government, consulting companies, and large companies are all looking for you, should you be interested in taking on work as a network forensic investigator. This will be challenging work, however, because in addition to an understanding of common forensic procedure and evidence handling, you also need a solid understanding of networking. This includes the TCP/IP suite of protocols as well as a number of application protocols. It also includes an understanding of some of the security technology that is commonly in place in enterprise networks like firewalls and intrusion detection systems.

Because there is currently no end in sight when it comes to computers being compromised by attackers around the world, there is no end in sight for the need for skilled forensics professionals. For forensic investigators without a foundation in network protocols and security technologies, this book intends to address that gap.

Summary

Businesses, government agencies, educational institutions, and non-profits are all subject to attack by skilled adversaries. These adversaries are, more and more, well-funded professional organizations. They may be some form of organized crime or they may be nation-states. The objectives of these two types of organizations may be significantly different but the end result is the same—they obtain some sort of unauthorized access to systems and once they are in place, they can be difficult to detect or extricate. This is where forensics professionals come in.

Forensics is a wide and varied field that has its basis in the legal world. Forensics, in a general sense, is anything to do with court proceedings. For our purposes, while the practice of digital forensics may have some foundation in law enforcement professionals performing investigations as part of criminal proceedings, the skills necessary to perform those investigations cross over to other areas. When it comes to investigations performed within an enterprise rather than by a law enforcement agency, the skills and techniques are the same but there may be differences in how artifacts and evidence are handled. That isn't always the case, of course, because even if you are just looking for the root cause, there is a possibility of what you find being necessary as part of a court case.

Because there is a possibility that artifacts and evidence may be used in court, it's generally a good idea to make use of cryptographic hashes as well as keeping a chain-of-custody document. These two activities will help you maintain accountability and a historical record of how the evidence and artifacts were handled. This is helpful if you have to refer to the events later on.

When it comes to working in an organization that isn't law enforcement, you may be asked to perform forensic investigations as part of an incident response. Incident response teams are becoming common practice at all sizes of organization. It's just how any organization has to operate to ensure that they can get back on their feet quickly and efficiently when an attack happens—whether it's

someone who has infiltrated the network by sending an infected e-mail or whether it's an attacker who has broken into the web server through a commonly known vulnerability.

Given the number of organizations around the world that have suffered these attacks, including several highly publicized attacks at Sony, Target, Home Depot, TJ Maxx, and countless others, there is a real need for forensics practitioners who can work with network data. This is because companies are using intrusion detection systems that will generate packet captures surrounding an incident and some organizations will actually perform a wire recording on a continuous basis simply in case an incident takes place. The network is the best place to capture what really happened because the network—the actual wire—can't lie.

References

Morgan, Steve. "Help Wanted: 1,000 Cybersecurity Jobs At OPM, Post-Hack Hiring Approved By DHS." (*Forbes*, January 13, 2016.) Retrieved June 22, 2016, from http://www.forbes.com/sites/stevemorgan/2016/01/31/help-wanted-1000-cybersecurity-jobs-at-opm-post-hack-hiring-approved-by-dhs/#3f10bfe12cd2.

Umberg, Tommy and Cherrie Warden. "Digital Evidence and Investigatory Protocols." *Digital Evidence and Electronic Signature Law Review*, 11 (2014). DEESLR, 11(0). doi:10.14296/deeslr.v11i0.2131.