

Chapter

1

Ethical Hacking

THE FOLLOWING CEH EXAM TOPICS ARE COVERED IN THIS CHAPTER:

- ✓ Professional code of conduct
- ✓ Appropriateness of hacking



COPYRIGHTED MATERIAL



Welcome to the exciting world of information security and, specifically, the important world of what is referred to as *ethical hacking*. You're here because you want to take the exam that will get you the Certified Ethical Hacker (CEH) certification. Perhaps you have done the training from EC-Council, the organization that manages the CEH, and you want a resource with a different perspective to help you as you prepare for the exam. Or you've decided to go the self-study route and you have enough experience to qualify for the exam. One way or another, you're here now, and this book will help you improve your understanding of the material to prepare for the exam.

The exam covers a wide range of topics, often at a deeply technical level, so you really need to have a solid understanding of the material. This is especially true if you choose to go on to the practical exam. This chapter, however, will be your starting point, and there is nothing technical here. In it, you'll get a chance to understand the foundations of the entire exam. First, you'll learn just what ethical hacking is, as well as what it isn't. The important part of the term *ethical hacking* is the *ethical* part. When you take the exam, you will be expected to abide by a code. It's essential to understand that code so you can live by it throughout your entire career.

Finally, you'll learn what EC-Council is, as well as the format and other details of the exam that will be useful to you. While some of it may seem trivial, it can be helpful to get a broader context for why the exam was created and learn about the organization that runs it. Personally, I find it useful to understand what's underneath something rather than experience it at a superficial level. As a result, you'll get the macro explanation, and you can choose to use it or not, depending on whether you find it helpful. It won't be part of the exam, but it may help you understand what's behind the exam so you understand the overall intentions.

Overview of Ethics

Before we start talking about ethical hacking, I will cover the most important aspect of that, which is ethics. You'll notice it's not referred to as "hacking ethically." It's ethical hacking. The important part is in the front. Ethics can be a challenging subject because you will find that it is not universal. Different people have different views of what is ethical and what is not ethical. It's essential, though, that you understand what ethics are and what is considered ethical and unethical from the perspective of the Certified Ethical Hacker certification. This is a critical part of the exam and the certification. After all, you are being entrusted with access to sensitive information and critical systems. To keep yourself viable as a professional,

you need to behave and perform your work in an ethical manner. Not only will you be expected to behave ethically, you will be expected to adhere to a code of ethics.

As part of the code of ethics, you will be sworn to keep information you obtain as part of your work private, paying particular attention to protecting the information and intellectual property of employers and clients. When you are attacking systems that belong to other people, you could be provided with internal information that is sensitive. You could also come across some critical information vital to the organization for which you are working. Failing to protect any of that data violates the code of ethics by compromising the confidentiality of that information.

You are expected to disclose information that needs to be disclosed to the people who have engaged your services. This includes any issues that you have identified. You are also expected to disclose potential conflicts of interest that you may have. It's important to be transparent in your dealings and also do the right thing when it comes to protecting your clients, employers, and their business interests. Additionally, if you come across something that could have an impact on a large number of people across the Internet, you are expected to disclose it in a responsible manner. This doesn't mean disclosing it in a public forum. It means working with your employer, any vendor that may be involved, and any computer emergency response team (CERT) that may have jurisdiction over your findings.

The first-time responsible disclosure was identified and documented in the 1990s, well over 20 years ago. The security researcher Rain Forest Puppy developed a full disclosure policy, sometimes called the Rain Forest Puppy Policy (RFP or RFPolicy). It advocated working closely with vendors to ensure they had time to fix issues before announcing them. At the time, there was a tendency for so-called researchers to just publish findings to the public to make a name for themselves without regard to the possibility of exposing innocent people when the vulnerabilities they found were exploited by attackers.

On the other side, companies that developed software hadn't caught up with the idea that they needed to be on top of vulnerabilities, and the slow months- or years-long pace of software development wasn't possible any longer with word of vulnerabilities getting out within minutes around the world. Hackers may have tried to notify a company only to have that company ignore the contact. Other companies may have acknowledged the bug but then dragged their feet about getting fixes out to their customers. The RFPolicy was an attempt to ensure that those who found vulnerabilities didn't just announce them indiscriminately but also had the ability to make the announcement if the company started to drag their feet. Wide acceptance of this policy within the security community dramatically increased the collaboration between those who were looking for vulnerabilities and those companies who had to be conscious of their consumers who may be exposed to attack if vulnerabilities were announced.

For examples of responsible disclosure, look at the work of Dan Kaminsky. He has found serious flaws in the implementations of the Domain Name System (DNS), which impacts everyone on the Internet. He worked responsibly with vendors to ensure that they had time to fix their implementations and remediate the vulnerabilities before he disclosed them. In the end, he did disclose the vulnerabilities in a very public manner, but only after vendors had time to fix the issue. This meant he wasn't putting people in the path of compromise

and potential information disclosure. Even though he was using the software in a way that it wasn't intended to be used, he was using an ethical approach by attempting to address an issue before someone could make use of the issue in a malicious way.

As you perform work, you will be given access to resources provided by the client or company. Under the code of ethics you will need to agree to, you cannot misuse any of the equipment. You can't damage anything you have access to as part of your employment or contract. There will be times when the testing you are performing may cause damage to a service provided by the infrastructure of the company you are working for or with. As long as this is unintentional or agreed to be acceptable by the company, this is okay. One way to alleviate this concern is to keep lines of communication open at all times. If it happens that an unexpected outage occurs, ensuring that the right people know so it can be remedied is essential.

Perhaps it goes without saying, but you are not allowed to engage in any illegal actions. Similarly, you cannot have been convicted of any felony. Along the same lines, though it's not directly illegal, you can't be involved with any group that may be considered "black hat," meaning they are engaged in potentially illegal activities, such as attacking computer systems for malicious purposes.

Colorful Terminology

You may regularly hear the terms *white hat*, *black hat*, and *gray hat*. White-hat hackers are people who always do their work for good. Black-hat hackers, probably not surprisingly, are people who do bad things, generally actions that are against the law. Gray-hat hackers, though, fall in the middle. They are working for good, but they are using the techniques of black-hat hackers.

Communication is also important when you embark on an engagement, regardless of whether you are working on contract or are a full-time employee. When you are taking on a new engagement, it's essential to be clear about the expectations for your services. If you have the scope of your services in writing, everything is clear and documented. As long as what you are being asked to do is not illegal and the scope of activities falls within systems run by the company you are working for, your work would be considered ethical. If you stray outside of the scope of systems, networks, and services, your actions would be considered unethical.

When you keep your interactions professional and ensure that it's completely clear to your employer what you are doing, as long as your actions are against systems belonging to your employer, you should be on safe ground ethically.

Overview of Ethical Hacking

These days, it's hard to look at any source of news without seeing something about data theft, Internet-based crime, or various other attacks against people and businesses. What we see in the news, actually, are the big issues, with large numbers of records compromised or big companies breached. What you don't see is the number of system compromises where the target of the attack is someone's personal computer or other device. Consider, for example, the Mirai botnet, which infected smaller, special-purpose devices running an embedded implementation of Linux. The number of devices thought to have been compromised and made part of that botnet is well over 100,000, with the possibility of there being more than one million.

Each year, millions of new pieces of malware are created, often making use of new vulnerabilities that have been discovered. Since 2005, there has not been a year without at least 10 million data records compromised. In the year 2017, nearly 200 million records were compromised. These numbers are just from the United States. To put this into perspective, there are only about 250 million adults in the United States, so it's safe to say that every adult has had their information compromised numerous times. To be clear, the data records that we're talking about belong to individual people and not to businesses. There is minimal accounting of the total value of intellectual property that may have been stolen, but it's clear that the compromise has been ongoing for a long time.

All of this is to say there is an urgent need to improve how information security is handled. It's believed that to protect against attacks, you have to be able to understand those attacks. Ideally, you need to replicate the attacks. If businesses are testing attacks against their own infrastructure early and often, those businesses could be in a better position to improve their defenses and keep the real attackers out.

This type of testing is what ethical hacking really is. It is all about ferreting out problems with a goal of improving the overall security posture of the target. This may be for a company in terms of their infrastructure or even desktop systems. It may also be performing testing against software to identify bugs that can be used to compromise the software and, subsequently, the system where the software is running. The aim is not to be malicious but to be on the "good" side to make the situation better. This is something you could be hired or contracted to perform for a business. They may have a set of systems or web applications they want tested. You could also have software that needs to be tested. There are a lot of people who perform testing on software—both commercial and open source.

Ethical hacking can be done under many different names. You may not always see the term *ethical hacking*, especially when you are looking at job titles. Instead, you will see the term *penetration testing*. It's essentially the same thing. The idea of a penetration test is to attempt to penetrate the defenses of an organization. That may also be the goal of an ethical hacker. You may also see the term *red teaming*, which is generally considered a specific type

of penetration test where the testers are adversarial to the organization and network under test. A red teamer would actually act like an attacker, meaning they would try to be stealthy so as not to be detected.

One of the challenging aspects of this sort of activity is having to think like an attacker. Testing of this nature is often challenging and requires a different way of thinking. When doing any sort of testing, including ethical hacking, a methodology is important, as it helps ensure that your actions are both repeatable and verifiable. There are a number of methodologies you may come across. Professionals who have been doing this type of work for a while may have developed their own style. However, they will often follow common steps, such as the ones I am going to illustrate as we move through the chapter.

EC-Council helps to ensure that this work is done ethically by requiring anyone who has obtained the Certified Ethical Hacker certification to agree to a code of conduct. This code of conduct holds those who have their CEH certification to a set of standards ensuring that they behave ethically, in service to their employers. They are expected to not do harm and to work toward improving the security posture rather than doing damage to that posture.

Methodologies

As with so many things, using a methodology is valuable when it comes to ethical hacking or security testing. Methodologies can help with consistency, repeatability, and process improvement. Consistency is important because you want to run the same sets of tests or probes no matter who you are testing against. Let's say you are working with a company that keeps asking you back. Without consistency, you may miss some findings from one test to another, which may let them think they got better, or the finding doesn't exist any longer. This would be a bad impression to leave a company with. Similarly, repeatability gives you the ability to do the same tests every time you run the assessment. In fact, if you are working with a team, every one of you should be able to run the sets of tests. Again, you want to be sure that any organization you are assessing will have the same perspective on their security posture, no matter how many times they come to you and no matter who the organization is.

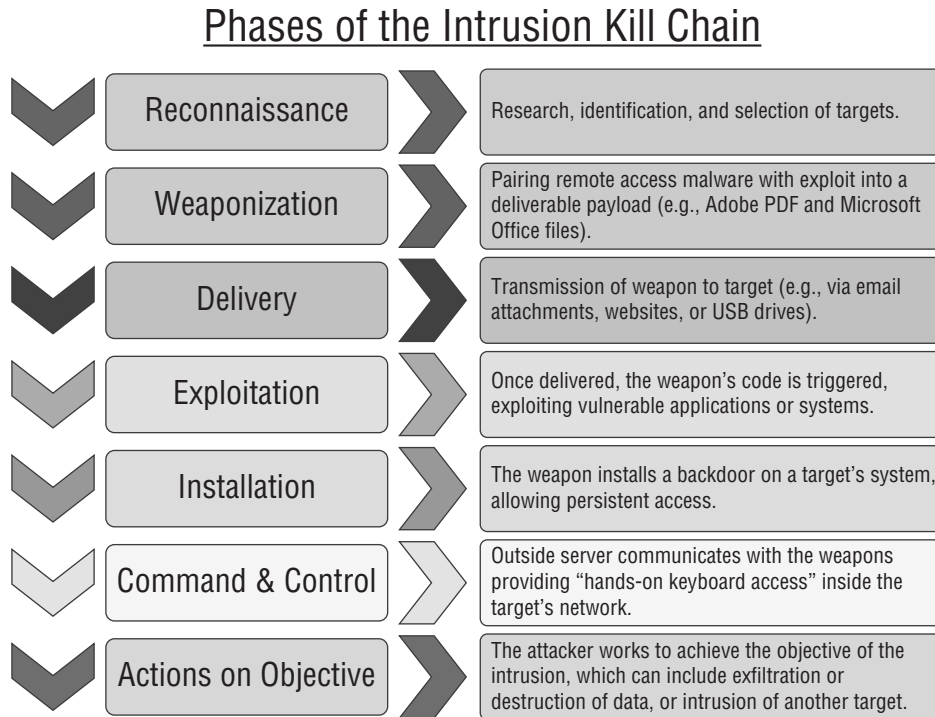
There are some testing or assessment methodologies that get used throughout the industry. They are typically built around expectations of what an attacker would do. While this may or may not be realistic, there are some other ways of viewing how attackers operate that are useful to understand. The first is the cyber kill chain, and the other is the attacker lifecycle. When we get to look at the methodology of ethical hacking, you will see the similarities.

Cyber Kill Chain

A commonly referred-to framework in the information security space is the cyber kill chain. A kill chain is a military concept of the structure of an attack. The idea of a kill chain is that you can identify where the attacker is in their process so you can adapt your own response

tactics. Lockheed Martin, a defense contractor, adapted the military concept of a kill chain to the information security (or cybersecurity) space. Figure 1.1 shows the cyber kill chain, as developed by Lockheed Martin.

FIGURE 1.1 Cyber kill chain



The first stage of the cyber kill chain is *reconnaissance*. This is where the attacker identifies their target as well as potential points of attack. This may include identifying vulnerabilities that could be exploited. There may be a lot of information about the target gathered in this phase, which will be useful later in the attack process.

Once the attacker has identified a target, they need to determine how to attack the target. This is where *weaponization* comes in. The attacker may create a custom piece of malware, for instance, that is specific to the target. They may just use a piece of common off-the-shelf (COTS) malware, though this has the potential to be discovered by antivirus software installed in the victim's environment. The attacker may decide this doesn't matter by sending out more malicious software to more individuals.

Delivery is how you get the weapon (the malware or the link to a rogue website) into the victim's environment. This could be a network-based attack, meaning there is an exposed

service that may be vulnerable to exploit remotely. This could be sending an attachment in via email, or it could be that the malicious software is hosted on a web server the victim is expected to visit and get infected when they hit the website. *Exploitation* is when the malicious software infects the victim's system.

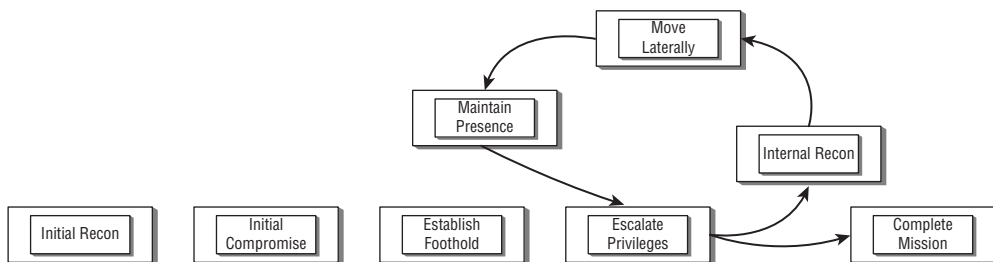
Exploitation leads to *installation*. The attacker will install additional software to maintain access to the system and perhaps give them remote access to the system. Once installation is complete, the attacker moves to *command and control*. You will sometimes see this referred to as C2. The command-and-control phase gives attackers remote access to the infected system. This may involve installation of additional software, or it may involve sending directives to the infected system. The attacker may be trying to get information from the infected system or have the system perform actions like participating in a large-scale denial-of-service attack.

These actions are called *actions on objectives*. Each attacker may have different objectives they are trying to achieve. Attackers who are criminally oriented are probably looking for ways to monetize the infected systems by stealing information that could be stolen or by selling off access to another organization. So-called nation-state actors may be looking to gain access to intellectual property. No matter what the organization is, they have objectives they are trying to achieve. They will keep going until they achieve those objectives, so there is a lot of activity that happens in this phase of the kill chain.

Attack Lifecycle

The security technology and consulting company FireEye Mandiant often refers to a different methodology called the *attack lifecycle*. This is different from the cyber kill chain, though there are some similarities. Rather than a theoretical exercise or one with a military focus, the attack lifecycle describes exactly how attackers have operated for as far back as there have been attacks against computing infrastructure. If you go back and look at how the Chaos Computer Club operated in the 1980s or Kevin Mitnick and his contemporaries operated in the late 1970s into the 1980s and beyond, you can map their actions directly into the attack lifecycle. Figure 1.2 shows how the attack lifecycle looks.

FIGURE 1.2 Attack lifecycle



One significant difference between the attack lifecycle is a recognition that the attack is not one and done. There is a loop that happens in the middle. Attackers don't keep launching attacks from outside the network. Once they get into the environment, they use the compromised systems as launch points for additional compromises within the environment. Attackers will gain access to a system and use that system and anything discovered there, like credentials, to move off to another system in the network. Before we get there, though, an attacker identifies a victim and potential attack possibilities in the *initial recon* stage. The attacker is doing reconnaissance, including identifying names and titles using open source intelligence, meaning they use public sources like social network sites, to generate attacks. To gain access, they launch attacks—commonly, these would be phishing attacks. This is the *initial compromise* stage.

Once they have compromised a system, the attacker will work to *establish footholds*. This includes making sure they retain access to the system so they can get back in when they need to. It's perhaps important to recognize that these attacks don't happen in a bang-bang fashion. It may take days or weeks to move from one phase of the attack lifecycle to another. This depends on the organization performing the attacks. These are not individuals. They are organizations, so there may be different employees working on different stages.

To do much else, the attacker will need to *escalate privileges*. They need to have administrative privileges to move into the loop that happens as they continue to move through the environment, gathering additional systems along the way. They will probably be gathering credentials from memory or disk here. They will also be investigating connections the system is known to have had with other systems in the network. This is a form of *internal recon*. They may also be trying to identify other credentials that are known to the system.

The reconnaissance is necessary to be able to *move laterally*. This is sometimes known as east-west movement. If you think about the network diagram, the connection to the outside world is quite often on the top. On a map, this would be north, so moving into and out of the network is known as north-south. Any movement within the organization is side to side or lateral movement. On a map, side to side would be east-west. To make those lateral movements, attackers need to know what systems there are. It may be servers, since individual systems are likely to know a lot of servers they communicate with, but it may also be individual workstations. In an enterprise network, it may be possible to authenticate using captured credentials against other workstations, which may have access to different sets of servers.

With every system the attacker gets access to, they need to *maintain presence*. This means some form of persistence, so any malware that is allowing the attacker access remains running. You might use the Windows registry, scheduled tasks, or other types of persistence to keep any malware running so the attacker can keep getting back in when they want.

The last phase of the attack lifecycle, though leaving it until the end is misleading, is *complete mission*. Again, attacks tend not to be one and done. Once an attacker is in your environment, they will likely be continuing to revisit to see if there is anything else you need. They may be continuing to get a broader reach within the organization. The complete mission phase is where data may be exfiltrated from the environment. This, again, may not

be a onetime thing. The attacker may continue to find additional targets in the environment to exploit, which would likely mean additional exfiltration. This means there would be continuous returns to this phase. After all, if you are planning to take up years-long residence, you don't want to wait years before getting data out because you can't, as an attacker, ever know when something may change and you lose access.

Methodology of Ethical Hacking

The basic methodology is meant to reproduce what real-life attackers would do. You will see similarities here to both the cyber kill chain and the attack lifecycle. Companies can shore up their security postures using information that comes from each stage covered here. One thing to keep in mind when it comes to information security is that not everything is about protection or prevention. You need to be able to detect all of these attacker activities.

Reconnaissance and Footprinting

Reconnaissance is where you gather information about your target. You want to understand the scope of your endeavor up front, of course. This will help you narrow your actions so you aren't engaging in anything that could be unethical. You'll have some sense of who your target is, but you may not have all the details. Gathering the details of your target is one of the reasons for performing reconnaissance. Another reason is that while there is a lot of information that has to be public just because of the nature of the Internet and the need to do business there, you may find information leaked to the rest of the world that the organization you are working for would do better to lock down.

The objective of reconnaissance and footprinting is determining the size and scope of your test. *Footprinting* is just getting an idea of the "footprint" of the organization, meaning the size and appearance. This means trying to identify network blocks, hosts, locations, and people. The information gathered here will be used later as you progress through additional stages.

Keep in mind that while you are looking for details about your target, you will find not only network blocks, which may exist within enterprise networks, but also potentially single hosts, which may belong to systems that are hosted with a service provider. As these systems will run services that may provide entry points or just house sensitive data, it's necessary to keep track of everything you gather and not limit yourself to information available about network blocks that the company may have.

In the process of doing this work, you may also turn up personal information belonging to employees at your target. This will be useful when it comes to social engineering attacks. These sorts of attacks are commonplace. In fact, some estimates suggest that 80 to 90 percent of infiltrations are a result of these social engineering attacks. They are not the only means of accessing networks, but they are commonly the easiest way in.

Scanning and Enumeration

Once you have network blocks identified, you will want to identify systems that are accessible within those network blocks; this is the scanning and enumeration stage. More important, however, you will want to identify services running on any available host. Ultimately, these services will be used as entry points. The objective is to gain access, and that may be possible through exposed network services. This includes not only a list of all open ports, which will be useful information, but also the identity of the service and software running behind each open port.

This may also result in gathering information that different services provide. This includes the software providing the service, such as nginx, Apache, or IIS for a web server. Additionally, there are services that may provide a lot of details about not only the software but the internals of the organization. This may be usernames, for instance. Some Simple Mail Transfer Protocol (SMTP) servers will give up valid usernames if they are queried correctly. Windows servers using the Server Message Block (SMB) protocol or the Common Internet File System (CIFS) protocol can be asked for information. You can get details like the directories being shared, usernames, and even some policy information. The objective of this phase is to gather as much information as you can to have starting points for when you move into the next phase. This phase can be time-consuming, especially as the size of the network and enterprise you are working with grows. The more details you can gather here, the easier the next stage will be for you.

Gaining Access

Gaining access is what many people consider to be the most important part of a penetration test, and for many, it's the most interesting. This is where you can demonstrate that some services are potentially vulnerable. You do that by exploiting the service. There are no theoretical or false positives when you have compromised a system or stolen data and you can prove it. This highlights one of the important aspects of any ethical hacking: documentation. Just saying, "Hey, I did this" isn't going to be sufficient. You will need to demonstrate or prove in some way that you did manage to compromise the system.

Technical attacks, like those looking for vulnerabilities in listening network services, are sometimes thought of as how systems get compromised, but the reality is that social engineering attacks are far more likely to be the way attackers gain access to systems. This is one of the reasons why enumeration is important—because you need targets for social engineering attacks. There are a number of ways to perform social engineering attacks, including using email to either infect a machine with malware or get the user to provide information that can be used in other ways. This may be the username and password, for instance.

Another mechanism for gathering information from users is to get them to visit a website. This may be a website that you, as the attacker, have loaded with malicious software that will infect their systems. Or, as before, you may be asking them for information. You've seen malware mentioned twice here. Understanding how malware works and where it can be used can be an important part of gaining access.

You will not always be asked to perform social engineering attacks. Companies may be handling security awareness, which commonly includes awareness of social engineering attacks, in other ways and not want or expect you to do phishing attacks or web-based attacks. Therefore, you shouldn't rely on using these techniques, in spite of the comparative ease of doing so, to get access to systems.

Maintaining Access

Once you are in, emulating common attack patterns means that you should maintain access. If you've managed to compromise a user's system, when the user shuts the system down, you will lose access. This may mean that you will need to re-compromise the system. Since exploits are not always guaranteed to be effective, you may well not get in the next time you attempt the compromise. Beyond that, you may have used a compromise that relied on a vulnerability that was fixed. Your next attempt may fail because the vulnerability is no longer there. You need to give yourself other means to get into the system so you can make sure you retain the ability to see what is happening on that system and potentially the enterprise network overall.

This is another stage where malware can be beneficial. You may need to install a rootkit, for example, that can provide you with a backdoor as well as the means to obscure your actions and existence on the system. You may need to install additional software on the system to maintain access. This may require copying the software onto your target system once you have done the initial compromise.

Therefore, this stage isn't as simple as perhaps it seems. There may be a number of factors that get in the way of ensuring that you maintain access. There are, though, a number of ways of maintaining access. Different operating systems allow for different techniques, but each operating system version or update can make different techniques harder. Ethical hacking is dependent on the circumstances, which is part of what makes it challenging. There are no single answers or straightforward approaches. One Windows 10 system may be easily compromised because there are patches that are available but missing. Another Windows 10 system may be difficult to get into because it is up-to-date, and it has been locked down with permissions and other settings.

Maintaining access is often called *persistence*. This is where any access mechanism is installed to persist on a system. No matter whether a user logs out or reboots a system, the attacker can continue to get in. This is commonly done by installing software that reaches out, or beacons, to systems on the Internet somewhere. The reason for this is because inbound access is often blocked by a firewall. Outbound access is often allowed from the inside of a network in a completely unrestricted manner.

Covering Tracks

Covering your tracks is where you hide or delete any evidence to which you managed to get access. Additionally, you should cover up your continued access. This can be accomplished with malware that ensures that your actions aren't logged or perhaps misreports system information, like network connections.

One thing to keep in mind when you are trying to cover your tracks is that sometimes your actions may also provide evidence of your work. One example is that wiping logs on a Windows system will leave a log entry indicating that the logs have been wiped. This may be an indication to anyone watching the logs that someone tried to erase evidence. It's not a guarantee that the log wipe was malicious, but it may be enough to prompt someone to investigate further. Because of this, covering tracks can be challenging. This may, though, be exactly what you've been asked to do—challenge and test the response capabilities of the operations team. As a result, it's always important to keep in mind the objectives of your engagement.

Summary

It's hard to overstate the importance of ethics. You will be expected to adhere to a code of ethics when you sign up for your CEH certification and pass your exam. You'll need to act in a professional manner at all times with your clients and employers. You will need to be a responsible custodian of any data entrusted to you.

