Chapter

1

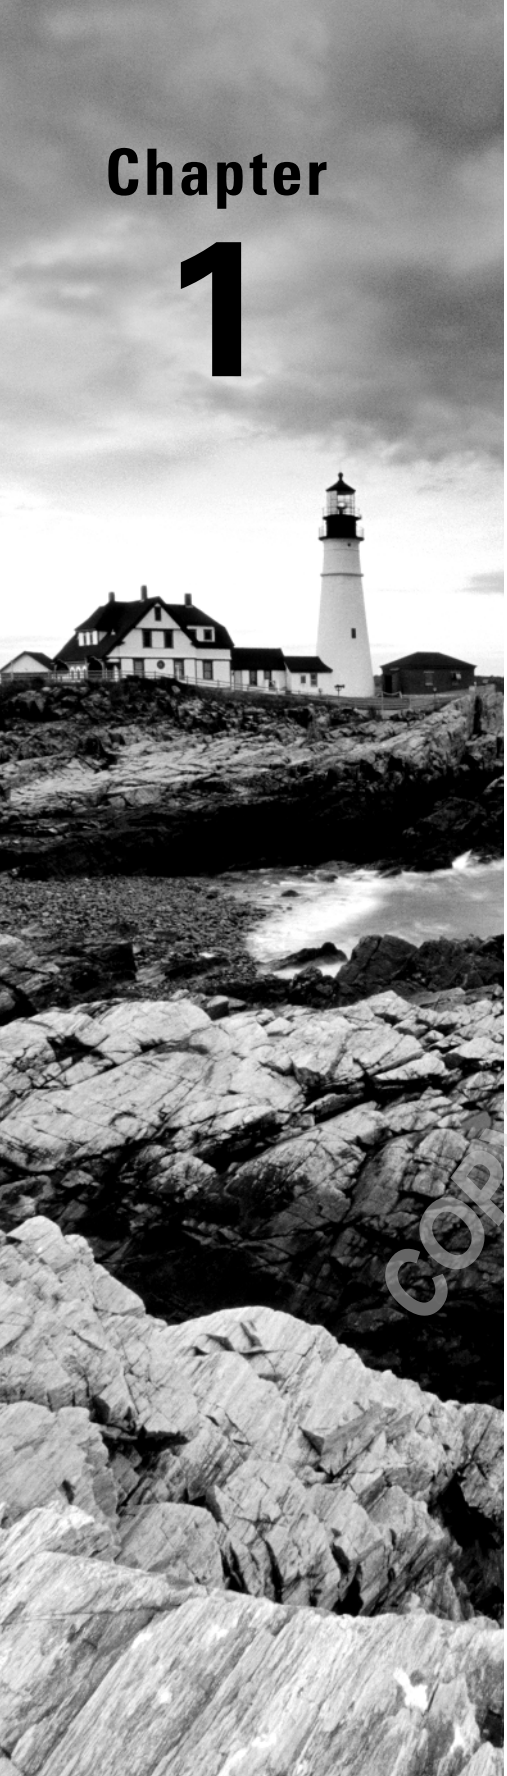

Introduction to Network Security

This chapter lays the foundation for future chapters covering the CCNA Security Exam objectives. In this chapter, I'll discuss the following topics:

- Threats to network security
- Network security objectives
- Classification of data
- Security controls
- Incident response
- Law and ethics

In Chapter 1, I'm going to explore some fundamental tenets of network security. It's important to have a foundation for the topics we will continue to explore in depth later on. I'll start with the threats to network security. Then I'll cover the objectives of network security and you'll learn some of the ways data is classified. Next, we will look at the types of controls that can be applied to help achieve network security.

In the two final sections, I will discuss the steps of incident response. Then we will briefly cover the law as it applies to computers and networks and the part ethics plays in network security.

Threats to Network Security

In today's world, the threat to computer networks is great. Because of the nature of the global communications network, threats can come from an increasing number of places and are constantly taking on new forms. An old joke in security circles goes, "What's the most secure computer in the world?" "The one that isn't turned on." It's even more true today than it was when it was first told.

The computer that isn't turned on isn't of much value to us, so we have to strike some balance between the usability of the computer network and the security controls that we apply.

Threats to network security are generally categorized into two basic types, external and internal. Along the lines of the joke I made earlier, the best way to secure the network from external threats is to disconnect it from external networks. Unfortunately, that isn't

an option in today's business climate. Virtually every business these days is connected to the Internet. Even some military computer networks are connected to the Internet. The days of maintaining complete physical separation of internal and external networks are over.

Today we protect our networks by using firewalls, intrusion prevention systems, access lists, layer 2 controls, Application layer firewalls, proxies, and other systems to provide what is called *defense-in-depth*. That is to say, we layer on many types of defenses so that if an intruder defeats one security measure, there are many more that will continue to protect the network and its sensitive data. With defense-in-depth, intruders must work a lot harder to achieve their objective of penetrating the network.

External Threats

External threats come from anyone outside of your network who is trying to get in. They can take many different forms and are attempted by people with varying motives and skill sets, and there are new ways being discovered and exploited every day. The following list describes some of the means that are used to attempt to breach the perimeter of someone's network:

Social engineering While not necessarily a technology-based attack, social engineering can definitely be categorized as an attack against confidentiality. Typically, the attacker will pose as someone from the technical staff at a company either where the victim works or where they might have an account, such as a bank. The attacker will try to glean personal information from the victim in order to exploit something. For example, the attacker might try to gain access to an online bank account or find out the password to the victim's computer account at the office.

Denial of service attacks This is an attack that can disable or cripple the use of a system. It could just be a huge load of packets directed to the target, it could exploit a vulnerability in the operating system or code, or it could involve the use of incorrectly formatted data. There are several specific types of denial of service attacks:

SYN flood A SYN floods is the result of using the TCP three-way handshake to cause a denial of service. The attacker sends many crafted packets with spoofed source IP addresses to a host or device, such as a firewall or router. Because the source IP is spoofed, when the router or firewall responds with `SYN ACK`, it is never received. The router or firewall holds that SYN conversation open, waiting for the ACK response from the sending host, which never comes. Multiply this many times and it chews up a lot of resources, which is the intent of the attack. The attack exhausts all of the resources of the router or firewall in order to effect a denial of service.

Smurf attack You may be having visions of little blue people, but that's not what I'm referring to. In a smurf attack, a single host, using a spoofed source IP, sends a ping flood to a broadcast address. This is a form of denial of service attack. What happens is

that every IP that receives the ping request in that broadcast domain sends a response in an echo reply to the spoofed IP, which is a real IP and the target of the denial of service attack. This attack is well known and was a common type of attack in the '90s. It can still be used today, but it is less common because routers can be configured to not forward broadcast requests through the use of the `no ip directed-broadcast` command. Also, hosts are generally configured to not respond to broadcasts.

Distributed denial of service (DDOS) attacks Botnets are the weapon of choice in today's world when you want to conduct a DDOS attack. Heck, you can even rent one if you want to. But really, any group of computers under an attacker's direct or indirect control can be used to launch an attack using any of the aforementioned methods.

A botnet is a group of computers that have been compromised in such a way that they can be controlled by a third party, sometimes known as a *bot herder*. A compromised computer's owner seldom knows that it has been compromised, which is why the computers are sometimes referred to as *zombies*. Zombies are usually infected via specially crafted viruses, worms, or Trojan horses. They are usually controlled via Internet Relay Chat (IRC) by a command and control computer, usually a server under the attackers' control. Botnets can comprise tens of thousands of computers; a few have as many as a million. They can be used for spreading malware, denial of service attacks, and so on. The owners of such bots sometimes lease them to others.

Man-in-the-middle (MITM) attack This is a classic attack in which the attackers insert themselves into the middle of a conversation where they can observe traffic that goes from one computer to another. They can simply observe or manipulate traffic as it flows from source to destination and back.

Session hijacking This is similar to a man-in-the-middle attack, but the attacker is able to take over a session that the victim has launched. This is done by observing, by guessing, or by some other means of gaining access to the session ID that is in use while a victim is connected. In some cases, once the victim has been granted a session ID, the attacker assumes the role of the victim by causing a denial of service against the victim so the victim is out of the picture as far as the session goes. Then the attacker continues on with the session, posing as the victim.

Brute force attack A brute force attack involves trying all possible combinations until the correct one is found. This applies to passwords, encryption keys, and the like. An example of this type of attack occurs in the wireless world. It is now known that the Wired Equivalency Protocol (WEP) is subject to an attack in which the attacker breaks the WEP key. Once this is done, all traffic that is encrypted is available to the attacker, who is posing as another host on the wireless network. In a nutshell, this is done by exploiting a vulnerability in the WEP and capturing enough packets to run the exploit. This is just one type of brute force attack; there are many others.

These are but a few of the increasing number of ways that a network, computer, or application can be exploited.

Internal Threats

Internal threats to a network — in which network intrusion or data theft is carried out by people inside your organization — are generally considered the most dangerous type of threat because people inside your organization have more access and more knowledge about the network and the devices on the network. An internal threat could be something as simple as someone installing a keylogger on a coworker's computer and stealing a password to a system they want access to.

Internal threats have long been considered serious, but statistics are starting to show that this may not be the case. The CSI Computer Crime and Security Survey keeps track of losses attributed to computer crime and security breaches. In Figure 1.1, a graph from this study shows the loss due to insiders based on a survey size of 420 respondents.

FIGURE 1.1 Percentage of losses due to insiders

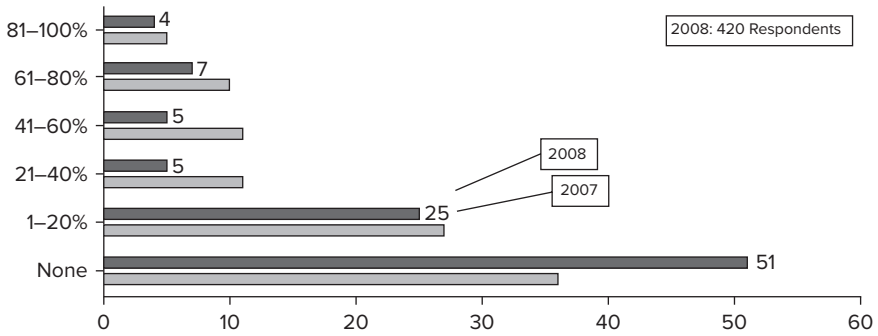

You can sign up to receive a free copy of the CSI Computer Crime and Security Survey at the following Web page: `http://www.gocsi.com/forms/csi_survey.jhtml`.

Internal threats are successful for a number of reasons:

- Improper patching
- Default configurations and default passwords
- Vendor best practices not followed

- Insecure programming practices
- Lax administration procedures

Application Security

Application attacks are attempts to gain control of a network or computer through a specific application, often a web browser or email application. Application attacks deserve special mention because they are said to account for 75 percent of all attacks carried out in today's networking environment. They are frequently successful due to a number of factors:

- Secure programming practices are not widespread. Programmers are usually under the gun to deliver code as quickly as possible, so security is often an afterthought. The best way to combat this is to use the Systems Development Life Cycle (SDLC) as a process to ensure that security is "built in" and not an afterthought.
- There's no firewall to crack through. With public facing web services, the application is already accessible.
- There are a growing number of security tools and applications that are built specifically for web attacks. And they are relatively easy to use.

Network Security Objectives

In a discussion about the tenets of network security, inevitably the CIA triad comes up (see Figure 1.2). CIA stands for confidentiality, integrity, and availability. These are collectively known as the objectives of network security.

FIGURE 1.2 The CIA triad

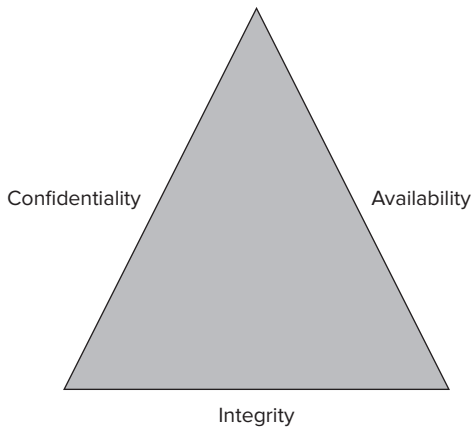

Let's discuss each of them briefly:

Confidentiality Whether data is at rest or in motion, we are concerned about its confidentiality. Encryption is one way to protect data confidentiality. (Encryption will be covered in Chapter 9, "Understanding Cryptographic Solutions.") For example, when you make a purchase on the Web and you put in your credit card data, you should be concerned about the possibility of an unauthorized person seeing your personal details. That's why you should look for the use of HTTPS and/or the little lock symbol (depending on your browser) somewhere on your browser screen to indicate a secure session.

Integrity Integrity refers to efforts made to ensure that data (whatever it may be) arrives at its destination unchanged. One method of checking integrity is to perform a check on the data using hash functions. A cryptographic hash function is defined as taking a string of variable length as input and producing a fixed-length hash. These checks help determine that the data that was originally sent is the same as the data that arrived.

Availability Is data available? For example, if you are unable to access your Gmail account, which happed to me one morning recently, you have an availability problem. We have become a data-oriented society, and it's upsetting if our data isn't available.

So why do I mention the three objectives of information security? Think about the methodology a hacker uses to carry out an attack. When an attack takes place, it is against one of these three objectives. A denial of service attacks the availability of data. What about a simple defacement of a website? Wouldn't that be a simple attack against integrity? Let's take that example a bit further. Suppose someone adds false information to a Wikipedia entry and then someone else reads it and believes it to be true? Without getting into whether you should or should not believe everything you read on the Internet, such an action could be potentially damaging, and an attack against integrity. Finally, if someone breaks into a database and steals a load of credit card numbers, an attack against confidentiality has occurred.

Table 1.1 lists the network security objectives and some of the types of attack strategy used against each.

TABLE 1.1 Examples of Attack Types

Attack Category	Attack Strategy
Confidentiality	Man-in-the-middle, packet capture, port mapping
Integrity	Malicious code/data diddling, keylogger, proxy
Availability	DDOS, SYN flood, smurf attack

Classification of Data

Data classification is one of the first tasks an organization needs to embrace in order to adequately protect the data. If you don't know the value of the data that you are trying to protect, how can you know if you are assigning the appropriate controls? In many ways, it is enlightening to an organization to invest the time to determine what data they have, who owns the data, and how the data should be protected.

The following lists a number of important tenets surrounding data classification:

- All data is not equal.
- Some data might be embarrassing or damaging if made public.
- Some compliance measures require classification.
- You should focus your efforts on protecting the most critical data instead of trying to protect all data equally.

So what are the most important points to consider when classifying data?

First, if the organization goes to the trouble of classifying data, it sends a strong message about how dedicated it is to information security.

Second, the owners of the data get to determine what its sensitivity is. This helps system administrators and network security professionals determine the appropriate controls to safeguard the data.

Last, some types of data are subject to regulatory controls. Regulations may determine what controls should be used to safeguard the data.

Classification levels differ from organization to organization. Many firms try to follow a classification system similar to that of the many military organizations. Their classification system is shown in Table 1.2.

TABLE 1.2 Military Data Classifications

Classification	Description of Data
Unclassified	Data that has no confidentiality, integrity, or availability requirements. Not sensitive to the organization, so no need to secure.
Sensitive but unclassified (SBU)	Data that could provide some embarrassment to the organization if revealed, but no major security restrictions.
Confidential	Data that has the least-restrictive protection within the classified realm. This data must have confidentiality protection.

Classification	Description of Data
Secret	Data that must be secured at significant cost and effort. This data should be more restrictive that Confidential but less restrictive than Top Secret.
Top Secret	Data that must be secured with the most effort and cost, if necessary. Data at this level is usually available only to those who have been cleared at this level and who have a legitimate need to know.

Some organizations don't feel the need to have five levels of data classifications. Many organizations come up with their own scheme that fits their organizational needs. A common data classification scheme used in the private sector is shown in Table 1.3.

TABLE 1.3 Private Sector Data Classifications

Classification	Description of Data
Public	Data that is in the public domain, such as white papers, stock information, marketing information. No protection is required.
Sensitive	Very similar to the Sensitive But Unclassified (SBU) classification in the military model. An example might be a time-sensitive piece of data that could cause embarrassment if revealed early but would not cause a security breach.
Private	Data that is important to the organization and is protected accordingly. An employee directory might fall into this classification level. This information is not intended to be revealed to the outside world.
Confidential	The highest level within a private sector organization and afforded the highest level of security controls and expense. A trade secret, a formula, or non-public financial data are examples.

When an organization takes on the task of classification, it usually uses criteria like the following:

Value This is the cornerstone of classifying data. What is the data worth to the organization? If you can't determine what it is worth, you have no basis on which to decide how much to spend to secure it.

Age Ever notice that even the Pentagon declassifies certain data after a specific period of time? That's because it is no longer relevant from a security standpoint. An example might be a technology that was once secret but has now advanced well beyond its original form, so there is no reason to maintain the secrecy.

Usefulness There's no reason to protect information no one cares about. For example, a three-year-old-financial report isn't really useful any more.

Personal data This is protected by many different laws, so by definition, it has to be classified.

Once you decide on a scheme, you need to do some more tasks to finish the job of classifying data. Generally you need to identify who owns the data, how to classify it, and how best to secure it. Here's a checklist to help determine how to proceed:

1. Identify the custodian of the data.
2. Determine how the data is classified and labeled.
3. Identify the data owner and classify accordingly.
4. Identify any exceptions.
5. Identify specific controls to be used with each classification, if necessary.
6. Define the rules for declassifying data and disposal, if necessary.
7. Create a security awareness program that addresses information classification.

Are there times when data must be divulged, even though it is classified? As with everything else, there are certain gray areas or exceptions.

Court orders are one of the first things that come to mind in a discussion about divulging sensitive data. It is sometimes ordered that certain data be revealed to the court, including to court personnel. However, there is some expectation that the data would be divulged only to those with a need to know and that it would be protected while in the court's custody.

Here's another example. Suppose your company decided to outsource the HR department. The company would almost certainly need to turn over personnel records to the outsourcing company performing the HR function. This kind of data is generally protected legally through contractual obligations. This kind of situation calls for due diligence in that the outsourcing company must build in safeguards and determine effective controls. A good contract would specify the right to unannounced audits of security controls at any time.

Last, the roles that are played in an information security classification scheme are as follows.

Owner The owner of the data is ultimately the person responsible for it, and it's usually a high-ranking member of management. The owner is generally different than the custodian.

Custodian The custodian is the person who takes care of the data, usually someone within the IT department. This is the person who maintains the server, the database, the backup tape, and so on.

User The user is someone who uses the data as part of their day-to-day duties. Users are generally governed by policies as to what they can or can't do with the data and how to handle it.

In the next section, we will look at the types of security controls that can be employed to protect the data after it's classified.

Security Controls

Security controls are measures applied to manage and reduce risk to your data. The strength of security controls and the type applied depend on the data being safeguarded. Security controls are generally lumped into three categories: administrative, technical, and physical.

Once data has been categorized and the owner has been identified, it is the custodian's responsibility to apply security controls to safeguard the data.

Security Controls by Type

Let's begin by examining the three basic types of security controls.

Administrative Administrative controls are typically associated with policies and procedures. Here are some examples of administrative controls:

- Security policies and procedures
- Awareness training
- Audits
- Change control procedures
- Background checks of employees
- Prudent hiring practices
- Job rotation
- Separation of duties

Administrative controls form the management layer of the control structure. If you don't have some policies and procedures, then the other two types of security controls are probably not going to work effectively.

Technical Technical controls are usually hardware and software based, as are the following examples:

- Firewalls
- Intrusion prevention systems
- Router access Controls Lists
- Virtual private network (VPN) devices
- Identity management systems such as Terminal Access Controller Access-Control System Plus (TACACS+) and Remote Authentication Dial In User Service (RADIUS)
- Network admission control systems
- Tokens and smartcards

As with administrative controls, all of the technical controls in the world are worthless if they are the only line of defense. The combination of all three types helps to achieve data security.

Physical Physical controls are typically mechanical in nature. The following are examples of physical controls:

- Locks
- Uninterruptible power supplies
- Diesel generators
- Security guards
- Motion sensors
- Alarm systems
- Safes
- Fire suppression systems

It must be said that even though there is a need for physical controls, human safety must come above all else. Doors and locks must have safeguards to allow exit in the event of an emergency.

Security Controls by Purpose

There is another way to categorize a security control — by its purpose. Each type of control mentioned earlier can also be categorized as follows:

- Preventative control
- Deterrent control
- Detective control

Sometimes a control can be in more than one of the preceding categories. For example, a security camera can be both a deterrent control (causing a would-be intruder to have second thoughts about breaking in) and a detective control (it can make a recording of what did happen so that security personnel have clues to an intruder's identity).

An example of a preventative control is proper hiring practices. When they are in place, human resource issues are prevented from arising because due diligence is put into the hiring process.

A security guard posted in the lobby of a bank is a great example of a deterrent control. Their main purpose is to make would-be robbers have second thoughts.

Finally, can you think of a detective control that is used in an IT environment? The first thing that pops into my mind is an Intrusion Detection System (IDS). This system alerts you after the fact that some anomalous condition occurred.

Incident Response

Responding to a security incident is an important part of security operations. How you respond is even more important. Maintaining one's composure and following a specific set of procedures is tantamount to success. There are typically six phases to incident response. You can remember these six phases by using the mnemonic word PICERL, as illustrated here:

- P - Preparation
- I - Identification
- C - Containment
- E - Eradication
- R - Restoration
- L - Lessons learned

We will explore each of these phases and then discuss some of the important facets of cybercrime.

Preparation

When we talk about the preparation phase from an incident handling perspective, we are talking about the steps required to get a team of people ready to handle incidents. It is through preparation that the elements come together. The following sections are not necessarily an exhaustive discussion but are meant to convey some of the key elements.

Policy

To prepare for future security incidents, there must be a firm policy in place and a list of procedures for what will happen when an incident occurs (as it inevitably will). Decisions must be made ahead of time about what will be done. It's of the utmost importance to have management buy in and sign off on both the policy and the procedures for handling an incident. Because an incident could potentially involve law enforcement, the legal team should be involved early on along with management to decide when, or even if, law enforcement should be called upon.

What happens when an incident takes place? Who gets notified and how? Is it a good idea to communicate over normal channels, or should there be an out-of-band method? These are some of the questions that need answers prior to jumping in. Usually there is a call tree of specific management personnel that need to be notified in the event of an incident. Because security incidents may compromise or break systems that are normally used to communicate, a secure, out-of-band method for communicating is strongly recommended.

Human Factors

When dealing with a potential security breach, it is easy to get worked up and make mistakes. The first order of business is to remain calm and think everything through. You should keep a notebook and take notes of everything that is observed and done. This will be valuable if an incident leads to criminal charges. It's important to keep a chronological record of what is happening so you don't have to recall it from memory while on the witness stand! If you take very good notes, the facts will be there in black and white for you to recall, complete with time and dates.

Also in the realm of human factors is the task of building a great team to react when an incident happens. It's best to pick a team from many disciplines so you have representation across the enterprise and the ability to look at any given issue with some level of expertise.

You should establish an organization that can support a response time service-level agreement of, say, 15 minutes to 30 minutes. It is largely up to management to determine their comfort level and how much money and manpower they want to allocate. If you have an incident at a remote location, either at another branch of your organization or at a company you work with, you need to decide if someone there can help you or if you need to send someone.

Additional Recommendations

A few other elements are key in the preparation phase. The following list lays out a number of recommendations:

- Establish a command post or war room from which to work.
- Establish a communications plan.
- Establish training and test runs.
- Coordinate with support desks and system administrators.
- Have emergency passwords and crypto keys available.
- Have a bag available with tools, drives, cables, USB keys, extra batteries, and so on.
- Have bootable media available with security tools and a laptop (or two).
- Have available cell phones and emergency call lists.

Identification

If something happens to your network, you need to decide how serious it is. In other words, you need to identify whether you have an *incident* (a serious breach) or an *event* (something happened, but it's not serious). This is analogous to someone running into the back of your car while you are stopped at a stoplight. You know something happened because you felt a jolt. But you don't know if it's a big deal until you get out and look at the bumper. You might find that there's no damage and you both go on your merry ways. On the other hand, you might know instantly that things are bad and you should call someone. It's the same for incident handling. For some things, it's intuitively obvious that there's a problem. For others, you might not know until you do some investigation.

In most cases, it's not wise to let just anyone declare that an incident has occurred. There should be established procedures about who can declare an incident. The responsible party should be someone who has been specifically trained and can quickly recognize if something is serious or not. That's not to say there is always someone who can recognize this. Keep in mind that it's better to declare an incident that turns out not to be one than it is not to declare one and have something serious happen and maybe cause damage.

How do you know if an incident has happened? Usually you have observed one or more suspicious events. For example, you observe a system log where a logon to the system has been attempted a hundred times in a row in a short period. That raises your eyebrows. First, no sane human being would try to log in to a system a hundred times. Second, it would be physically impossible for a human being to attempt the logon that many times in a short period of time. Something doesn't add up here. This would be a suspicious event, for sure. Other examples are unexplained new files or file modifications or missing logs.

So what happens after you declare an incident? You have gone through the preparation phase, so you know who to contact and who should do what; in other words, you follow the playbook.

Real World Scenario

A DNS Hack

While working at a transportation company several years ago, I ran into a situation one Friday afternoon. (Why do all these things happen on Friday?) The person who was responsible for administering the DNS servers called me and said that he couldn't log in as root anymore. Being the inquisitive type, I used my personal account to log in and started looking at the system. I noticed that the `etc/password` file had been modified 20 minutes earlier. I asked if there was any chance that he changed the password by accident, and the answer, of course, was no. Was there anyone else who had access? No, again. Well, it was time to call the boss and relay the bad news. It turned out the servers were scheduled to be updated with a patch over the weekend to close a security hole. Looked like the update was scheduled too late. The shelf life of vulnerabilities is getting shorter all the time.

Because there was no incident response team, the prudent decision at that time was to remove the servers from the network. The company was lucky that in the grand scheme of things, this action wasn't that critical. Imagine what might have happened if this had been a revenue-generation server, like an e-commerce server or credit card processing server. After much deliberation, the decision was made to rebuild the DNS servers over the weekend and be back in business on Monday morning.

The person who hacked in wasn't very smart. They could have created an account that wouldn't have been noticed, like an account with a variation of someone else's username. They could have lurked about for days or months without being noticed. They could have used their newfound access to do surveillance or launch attacks on other systems. The lessons learned here were to be cognizant of security patches and get them installed quickly. From an identification point of view, I was able to see that there was a serious problem almost immediately and quick action was taken.

In "A DNS Hack," identification was made by looking at a system. This is just one of the areas where identification is performed. It is also performed at the network perimeter and using firewall logs, IDS/IPS alerts, and router logs. There are also host file integrity products as well as antivirus products that can alert you to a problem.

Containment

Once you've identified the problem, it needs to be contained. You could do as I did and remove the system from the network. Safe move, but what if you wanted to observe the behavior? What if you need to do forensics? There are many factors to weigh when looking at an incident. Is it over or is it ongoing? How critical are the systems involved? Can you

secure the area and still preserve evidence if you need to? Everyone handles these things differently, so be prepared. Can you filter the interloper at the perimeter with a router access control list (ACL)? Is the system stable enough to keep running? Lots of questions and the answers will depend on a lot of people's opinions, including management.

One thing that is almost always necessary is to do a backup as soon as possible. Usually a hard drive needs to be imaged multiple times. From a forensics standpoint, the original is kept for evidence. Backup copy 1 is then kept for possible production on a newly built system. Any other copies are made from backup copy 1 and could be used for forensics purposes.

Eradication

The next step is to decide how to get rid of the problem and get back to business. What's the first order of business when starting to get rid of the problem? If you guessed "find the last known clean backup," you hit the nail on the head. In today's environment, given the sophistication of many virus and worm writers, the chances that you are going to be rebuilding the system are pretty good. If that's the case, then you are definitely going to want the last good backup, because you're going to have to restore the data.

If you know how the attack happened, you will have to determine the best course of action to fix the system so the incident isn't repeated. If you don't know how it happened, maybe it's time to do a vulnerability assessment on the system to see if there are any big holes.

Recovery

Recovery is, simply put, getting the system back online and into production. The following steps are performed in the recovery phase:

- Restore the system.
- Validate the system.
- Put the system back into production.
- Monitor the system.

Lessons Learned

The final aspect of the incident response process is taking stock of the lessons learned. The goal here is to apply the knowledge learned during an incident to improve capabilities and hopefully the security.

This process should include a formal meeting at which a report should be presented explaining what happened, what was done, and how the situation can be improved to prevent a future occurrence.

Law and Ethics

In this section, we will explore the law as it relates to network security. We will cover the major types of law and how they might apply to a given situation. Intellectual property is a critical element in network security — I will briefly define the basic types. Ethics issues make big headlines in the news these days, and it plays an important role in how network security personnel work and how they are treated. In fact, it could be argued that the ethical standards of information security personnel need to be above and beyond those of other personnel. I'll wrap up this section with a discussion on major governance topics and how compliance with regulations fits into network security.

Legal Matters

The law is a driving factor behind policies at most organizations. To be more specific, compliance with certain laws and regulations is a motivating factor. Various regulatory agencies and the associated laws that they enforce make it necessary for anyone involved with information security to have at least a cursory knowledge of relevant aspects of the law. Depending on your job within information security, you might need to have more than just a cursory knowledge about the law. In particular, forensics specialists need to be conversant with rules of evidence, chain of custody, and other topics.

Let's begin by discussing the three basic types of law. Criminal law, civil law (also known as tort law), and administrative law are the three types typically found in most countries. That doesn't mean that they are all administered the same way or observed the same way in each country. This makes it more difficult to deal with a cybercriminal because more often than not, you are dealing with someone who might be operating in several countries.

Criminal Law

Criminal law involves transgressions of a jurisdiction's criminal code — laws against theft, murder, illegal drugs and other antisocial behavior. In the United States, these laws are enforced on local, state, and federal levels. Cop and lawyer shows on TV often portray the criminal law process, typically beginning when someone has committed a crime and then portraying the process as they travel through the justice system. Penalties can be probation, prison time, or fines and sometimes a combination of one or more.

As the computer/network industry continues to mature, more and more cases of criminal activity are starting to make headlines.

Generally speaking, prosecutors need to establish just three things to prove a criminal case in court:

Means Were they capable of committing the crime?

Motivation Was there a reason the culprit committed the crime?

Opportunity Were they able to commit the crime?

You might ask why this is important. It ties directly into network security because evidence must be kept and a chain of custody of the evidence must be established. Evidence in a network security case could include, for example, physical hard drives, firewall logs, and intrusion detection messages.

Civil Law

Civil law deals with a perceived wrongdoing that is not necessarily criminal. It covers such matters as personal injuries, medical malpractice, slander, and business dealings. Civil law is usually about receiving money, although sometimes the court may require that an action be taken (or undone, in some cases). Prison is not an option in civil law.

Administrative Law

Administrative law usually involves the regulations and actions of government agencies. For instance, if a company was found to be negligent in paying its employees, a state or federal agency could force the company to correct the issue monetarily and could also levy an administrative fine as a punitive step.

Intellectual Property

Intellectual property is the largest asset of many companies. Therefore, it must be treated as confidential information and protected accordingly. Intellectual property may take one of the following forms:

Trademarks A trademark is a word, a name, or a symbol that is used in trade and indicates where goods come from. It distinguishes one firm's products from another's.

Patents A patent for an invention grants property rights to the inventor for a period of 20 years from the date the application was filed in the United States. U.S. patents are only valid inside the United States and its territories. A patent grants the right to prohibit others from making, using, or selling the invention.

Trade secrets A trade secret is intellectual property that is not in the public domain. For example, the formula for Coca-Cola is a trade secret that is heavily protected. One of the tests to determine whether information is a trade secret is whether due care is taken to protect it.

Copyrights A copyright protects original works, which include musical, literary, dramatic, and artistic works, whether they are published or unpublished. The Copyright Act of 1976 gives the owner of a copyright specific rights, including the right to control reproduction, distribution, and public performance or display.

For further information on intellectual property, visit the U.S. Patent and Trademark Office website at `http://www.uspto.gov/web/offices/pac/doc/general/whatis.htm`.

Ethics

Ethics refers to a set of standards and principles that are deemed to be higher than the law. Ethics involve morals and standards that are considered to be proper behavior in a given situation. Certain industry groups and professions usually have a code of ethics that they expect their members or constituents to abide by. Within the information security community, there are a number of ethics codes developed by industry organizations:

- International Information Systems Security Certification Consortium $(ISC)^2$ Code of Ethics
- Global Information Assurance Certification (GIAC) Code of Ethics
- Information Systems Security Association (ISSA) Code of Ethics
- Information Systems Audit and Control Association (ISACA) Code of Professional Ethics
- Internet Architecture Board (IAB)
- Generally Accepted System Security Principles (GAASP)
- Computer Ethics Institute

Review Questions

1. Which type of law is also called tort law?
   - **A.** Administrative
   - **B.** Criminal
   - **C.** Judicial
   - **D.** Governmental
   - **E.** Civil

2. What is the guiding principle behind ethics instead of the law?
   - **A.** Common sense
   - **B.** The Golden Rule
   - **C.** Morals
   - **D.** Creed

3. What type of attack constitutes 75 percent of all attacks on today's networks?
   - **A.** Internal
   - **B.** External
   - **C.** Application
   - **D.** Password cracking
   - **E.** Network

4. Which one of the following can be described as the "getting back to business" aspect of incident handling?
   - **A.** Preparation
   - **B.** Investigation
   - **C.** Recovery
   - **D.** Lessons Learned
   - **E.** Containment

5. Which is *not* one of the things a prosecutor needs to prove in a cybercrime case?
   - **A.** Means
   - **B.** Money
   - **C.** Opportunity
   - **D.** Motivation

6. Which of the following is an external threat?
   - **A.** An unpatched system
   - **B.** Port scanning
   - **C.** Default configurations
   - **D.** Insecure programming practices

7. Which of the following would be part of the containment process?
   - **A.** Investigating an incident
   - **B.** Performing a backup
   - **C.** Building a new server
   - **D.** Formatting the hard drive

8. Which of the following is described as data that doesn't have any confidentiality, integrity, or availability requirements?
   - **A.** Classified
   - **B.** Secret
   - **C.** Sensitive but unclassified
   - **D.** Unclassified
   - **E.** Confidential

9. Which of the following grants protections for 20 years in the United States?
   - **A.** Trademark
   - **B.** Patent
   - **C.** Copyright
   - **D.** Intellectual property

10. Which of the following is *not* one of the phases of incident response?
    - **A.** Preparation
    - **B.** Containment
    - **C.** Identification
    - **D.** Lessons learned
    - **E.** Detection

P - Preparation
I - Identification
C - Containment
E - **Eradication**
R - Recovery
L - Lessons learned

Answers to Review Questions

1. E. Civil law is also known as tort law. This type of law involves the recovery of monetary damages.
2. C. Ethics can involve common sense, the Golden Rule, and creeds, but the best answer is morals, which are considered to be higher than the law.
3. C. Application attacks make up approximately 75 percent of all network attacks today.
4. C. Recovery describes the process of getting a server back into production and returning a business back to normal.
5. B. A prosecutor needs to prove means, motive, and opportunity.
6. B. Port scanning is usually done from outside, so that is your first clue that an attack is coming from outside the organization.
7. B. Performing a backup is part of the containment process because you may need to preserve evidence.
8. D. In the military data classification system, unclassified is the label for data that doesn't have any confidentiality, integrity, or availability requirements.
9. B. A U.S. patent is protected by law for a period of 20 years. Other types of intellectual property don't have the same protections.
10. E. Detection is not one of the phases of incident response. Remember the phases of incident response with the following mnemonic: PICERL.