Chapter 1

Security Concepts

THE FOLLOWING CCST EXAM TOPICS ARE COVERED IN THIS CHAPTER:

✔ **1.0 Essential Security Principals**

- 1.1. Define essential security principles

  Vulnerabilities, threats, exploits, and risks; attack vectors; hardening; defense-in-depth; confidentiality, integrity, and availability (CIA); types of attackers; reasons for attacks; code of ethics

- 1.2. Explain common threats and vulnerabilities

  Malware, ransomware, denial of service, botnets, social engineering attacks (tailgating, spear phishing, phishing, vishing, smishing, etc.), physical attacks, man in the middle, IoT vulnerabilities, insider threats, Advanced Persistent Threat (APT)

It's true. . .you're not paranoid if they really are out to get you. Although "they" probably aren't after you personally, your network—no matter the size—is seriously vulnerable, so it's wise to be very concerned about keeping it secure. Unfortunately, it's also true that no matter how secure you think your network is, it's a good bet that there are still some very real threats out there that could breach its security and totally cripple your infrastructure!

I'm not trying to scare you; it's just that networks, by their very nature, are not secure environments. Think about it—the whole point of having a network is to make resources available to people who aren't at the same physical location as the network's resources.

Because of this, it follows that you've got to open access to those resources to users you may not be able to identify. One network administrator I know referred to a server running a much-maligned network operating system as "a perfectly secure server until you install the network interface card (NIC)." You can see the dilemma here, right?

Okay, with all this doom and gloom, what's a network administrator to do? Well, the first line of defense is to know about the types of threats out there, because you can't do anything to protect yourself from something you don't know about. But once you understand the threats, you can begin to design defenses to combat bad guys lurking in the depths of cyberspace just waiting for an opportunity to strike.

I'm going to introduce you to some of the more common security threats and teach you about the ways to mitigate them. I'll be honest—the information I'll be giving you in this chapter is definitely not exhaustive. Securing computers and networks is a huge task, and there are literally hundreds of books on this subject alone. To operate securely in a network environment, you must understand how to speak the language of security. As in any field, there is terminology.

In this chapter, you will learn the common types of attacks that all network professionals should understand to secure an enterprise network.

To find your included bonus material, as well as Todd Lammle videos, practice questions, and hands-on labs, please see `www.lammle.com/ccst`.

Technology-Based Attacks

Technology-based attacks are those that take advantage of weaknesses in software and the protocols that systems use to communicate with one another. This contrasts with attacks that target environmental or human weaknesses (covered later in this chapter). In this section, you'll learn about attacks that target technologies.

Denial of Service (DoS)/Distributed Denial of Service (DDoS)

A denial of service (DoS) attack does exactly what it sounds like it would do—it prevents users from accessing the network and/or its resources. Today, DoS attacks are commonly launched against a major company's intranet and especially its websites. "Joe the Hacker" (formerly a plumber) thinks that if he can make a mess of, say, Microsoft's or Amazon's website, he's done that company some serious damage. And you know what?

He's right!

Even though DoS attacks are nasty, strangely, hackers don't respect other hackers who execute them because they're really easy to deploy. It's true—even a pesky little 10-year-old can execute one and bring you to your knees. (That's just wrong!) This means that "real" bad guys have no respect for someone who uses DoS attacks, and they usually employ much more sophisticated methods of wreaking havoc on you instead. I guess it comes down to that "honor among thieves" thing. Still, know that even though a DoS-type attack won't gain the guilty party any esteemed status among "real" hackers, it's still not exactly a day at the beach to deal with.

Worse, DoS attacks come in a variety of flavors. Let's talk about some of them now.

The Ping of Death

Ping is primarily used to see whether a computer is responding to IP requests. Usually, when you ping a remote host, what you're really doing is sending four normal-sized Internet Control Message Protocol (ICMP) packets to the remote host to see whether it's available. But during a ping-of-death attack, a humongous ICMP packet is sent to the remote host victim, totally flooding the victim's buffer and causing the system to reboot or helplessly hang there, drowning. It's good to know that patches are available for most operating systems to prevent a ping-of-death attack from working.

Distributed DoS (DDoS)

Denial of service attacks can be made more effective if they can be amplified by recruiting helpers in the attack process. In the following sections, some terms and concepts that apply to a distributed denial of service attack are explained.

Botnet/Command and Control

A botnet is a group of programs connected on the Internet for the purpose of performing a task in a coordinated manner. Some botnets, such as those created to maintain control of Internet Relay Chat (IRC) channels, are legal, whereas others are illegally created to foist a DDoS. An attacker can recruit and build a botnet to help amplify a DoS attack, as illustrated in Figure 1.1.

The steps in the process of building a botnet are as follows:

1. A botnet operator sends out viruses or worms whose payloads are malicious applications, the bots, infecting ordinary users' computers.
2. The bots on the infected PCs log into a server called a command-and-control (C&C) server under the control of the attacker.

3. At the appropriate time, the attacker, through the C&C server, sends a command to all bots to attack the victim at the same time, thereby significantly amplifying the effect of the attack.

Traffic Spike

One of the hallmarks of a DDoS attack is a major spike in traffic in the network as bots that have been recruited mount the attack. For this reason, any major spike in traffic should be regarded with suspicion. A network intrusion detection system (IDS) can recognize these traffic spikes and may be able to prevent them from growing larger or in some cases prevent the traffic in the first place.

Some smaller organizations that cannot afford the pricier intrusion prevention systems (IPSs) or IDSs make use of features present on their load balancers. Many of these products include DDoS mitigation features such as the TCP SYN cookie option. It allows the load balancer to react when the number of SYN requests reaches a certain point. At that point, the device will start dropping requests when the SYN queue is full.

Coordinated Attack

Another unmistakable feature of a DDoS attack is the presence of a coordinated attack. As shown in Figure 1.1 and as just described in the section "Botnet/Command and Control," to properly amplify the attack the bots must attack the victim at the same time. The coordination of the bots is orchestrated by the command-and-control server, depicted in Figure 1.1. If all the bots can be instructed to attack at precisely the same second, the attack becomes much more overwhelming to the victim.

Friendly/Unintentional DoS

An unintentional DoS attack (also referred to as an attack from "friendly fire") is not caused by malicious individuals; instead, it's a spike in activity to a website or resource that overpowers its ability to respond. In many cases, it is the result of a relatively

FIGURE 1.1 Botnet.

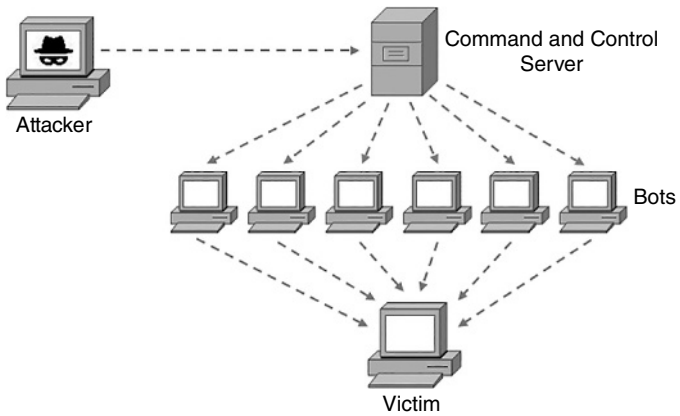

unknown URL suddenly being shared in a larger medium such as a popular TV or news show. For example, when Michael Jackson died, the amount of Twitter and Google traffic spiked so much that at first it was thought that an automated attack was under way.

Physical Attack

Physical attacks are those that cause hardware damage to a device. These attacks can be mitigated, but not eliminated, by preventing physical access to the device. Routers, switches, firewalls, servers, and other infrastructure devices should be locked away and protected by strong physical access controls. Otherwise, you may be confronted with a permanent DoS, covered in the next section.

Permanent DoS

A permanent DoS (PDoS) attack is one in which the device is damaged and must be replaced. It requires physical access to the device, or does it? Actually, it doesn't! An attack called phlashing attacks the firmware located in many systems. Using tools that fuzz (introduce errors into) the firmware, attackers cause the device to be unusable.

Another approach is to introduce a firmware image containing a Trojan or other type of malware.

Smurf

Smurfs are happy little blue creatures that like to sing and dance, but a Smurf attack is far more nefarious. It's a version of a DoS attack that floods its victim with spoofed broadcast ping messages. I'll talk about spoofing in more detail later; for now, understand that it basically involves stealing someone else's IP address.

Here's how it works: the bad guy spoofs the intended victim's IP address and then sends many pings (IP echo requests) to IP broadcast addresses. The receiving router responds by delivering the broadcast to all hosts in the subnet, and all the hosts respond with an IP echo reply—all of them at the same time. On a network with hundreds of hosts, this results in major network gridlock because all the machines are kept busy responding to each echo request. The situation is even worse if the routers have not been configured to keep these types of broadcasts confined to the local subnet (which thankfully they are by default!). Figure 1.2 shows a Smurf attack in progress.

Fortunately, Smurf attacks aren't very common anymore because most routers are configured in a way that prevents them from forwarding broadcast packets to other networks. Plus, it's really easy to configure routers and hosts so they won't respond to ping requests directed toward broadcast addresses.

SYN Flood

A SYN flood is also a DoS attack that inundates the receiving machine with lots of packets that cause the victim to waste resources by holding connections open. In normal communications, a workstation that wants to open a Transmission Control Protocol/Internet Protocol (TCP/IP) communication with a server sends a TCP/IP packet with the SYN flag set to 1, as

FIGURE 1.2 Smurf attack.

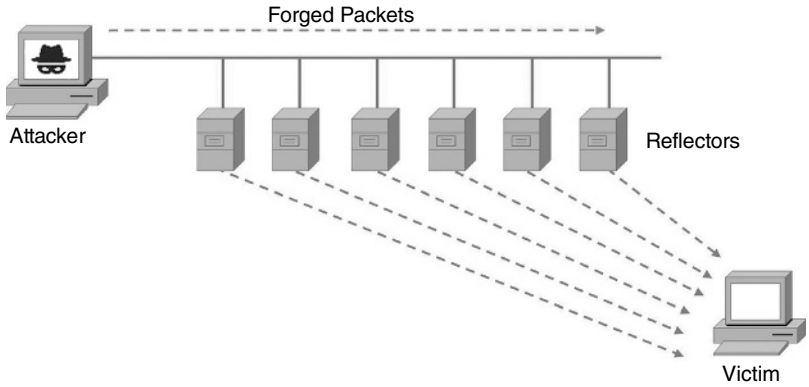

FIGURE 1.3 SYN flood.

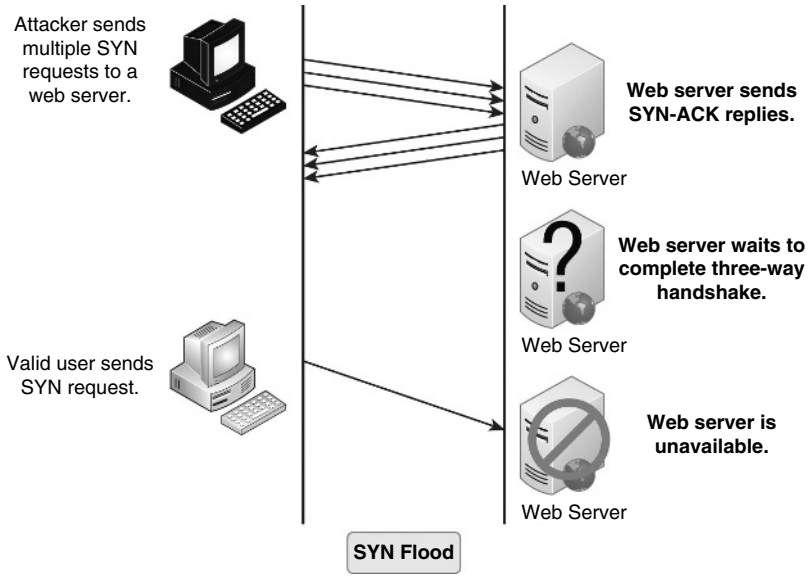

part of the three-way handshake process. The server automatically responds to the request, indicating that it's ready to start communicating with a SYN-ACK. In the SYN flood, the attacker sends a SYN, the victim sends back a SYN-ACK, and the attacker leaves the victim waiting for the final ACK. While the server is waiting for the response, a small part of memory is reserved for it. As the SYNs continue to arrive, memory is gradually consumed. Figure 1.3 shows an example of a simple DoS/SYN flood attack.

You can see that the preyed-upon machine can't respond to any other requests because its buffers are already overloaded, and it therefore rejects all packets requesting connections,

even valid ones, which is the idea behind the attack. Crafting a firewall rule to prevent flooding of SYN packets might help. However, if the threat actor can coordinate a large number of SYN packets from distributed hosts, reverse proxies can mitigate the attack.

Reflective/Amplified Attacks

Reflected or amplified attacks increase the effectiveness of a DoS attack. Two of the more effective of these types of attacks involve leveraging two functions that almost all networks use, DNS and NTP. In the next two sections, these attacks are described.

DNS

A DNS amplification attack is a form of reflection attack in that the attacker delivers traffic to the victim by reflecting it off a third party. Reflection conceals the source of the attack. It relies on the exploitation of publicly accessible open DNS servers to deluge victims with DNS response traffic.

The attacker sends a small DNS message to an open resolver using the victim's IP address as the source. The type of request used returns all known information about the DNS zone, which allows for the maximum level of response amplification directed to the victim's server. The attack is magnified by recruiting a botnet to send the small messages to a large list of open resolvers (DNS servers). The response from the DNS server overwhelms the victim, as shown in Figure 1.4.

NTP

Although NTP reflection attacks use the same process of recruiting bots to aid the attack, the attacks are not reflected off DNS servers; they are instead reflected off Network Time Protocol (NTP) servers. These servers are used to maintain time synchronization between devices in a network.

The attacker (and accompanying bots) sends a small spoofed 8-byte UDP packet to vulnerable NTP servers that requests a large amount of data (megabytes worth of traffic) be

FIGURE 1.4 DNS amplification attack.

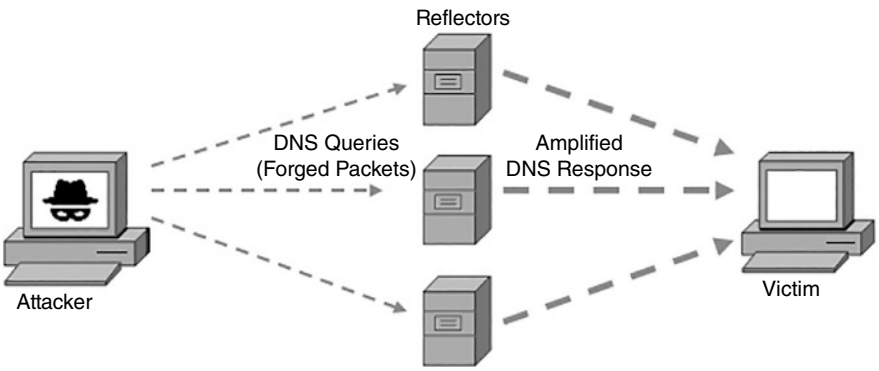

sent to the DDoS's target IP address. The attackers use the `monlist` command, a remote command in older versions of NTP, which sends the requester a list of the last 600 hosts that have connected to that server. This attack can be prevented by using at least NTP version 4.2.7 (which was released in 2010).

On-Path Attack (Previously Known as Man-in-the-Middle Attack)

Interception! But it's not a football, it's a bunch of your network's packets—your precious data. An on-path attack (previously known as a man-in-the-middle attack) happens when someone intercepts packets intended for one computer and reads the data. A common guilty party could be someone working for your very own ISP using a packet sniffer and augmenting it with routing and transport protocols. Rogue ATM machines and even credit-card swipers are tools that are also increasingly used for this type of attack. Figure 1.5 shows an on-path/man-in-the-middle attack.

DNS Poisoning

DNS clients send requests for name to IP address resolution (called queries) to a DNS server. The search for the IP address that goes with a computer or domain name usually starts with a local DNS server that is not authoritative for the DNS domain in which the requested computer or website resides. When this occurs, the local DNS server makes a request of the DNS server that does hold the record in question. After the local DNS server receives the answer, it returns it to the local DNS client. After this, the local DNS server maintains that record in its DNS cache for a period called the Time to Live (TTL), which is usually an hour but can vary.

FIGURE 1.5 On-path/man-in-the-middle attack.

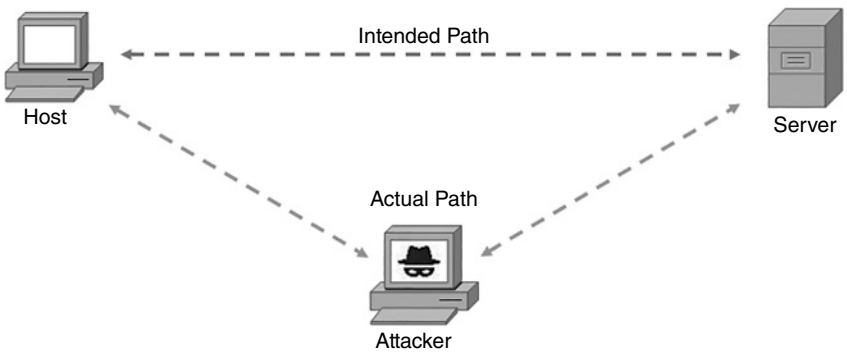

In a DNS cache poisoning attack, the attacker attempts to refresh or update that record when it expires with a different address than the correct address. If the attacker can convince the DNS server to accept this refresh, the local DNS server will then be responding to client requests for that computer with the address inserted by the attacker. Typically, the address they now receive is for a fake website that appears to look in every way like the site the client is requesting. The hacker can then harvest all the name and password combinations entered on their fake site. This type of attack is really a combination of a DNS cache poisoning attack and a phishing attempt.

To prevent this type of attack, the DNS servers should be limited in the updates they accept. In most DNS software, you can restrict the DNS servers from which a server will accept updates. This can help prevent the server from accepting these false updates.

VLAN Hopping

VLANs, or virtual LANs, are Layer 2 segmentation in a switched network. A VLAN may also span multiple switches. When devices are segregated into VLANs, access control lists (ACLs) can be used in a router to control access between VLANs in the same way it is done between real LANs. When VLANs span switches, the connection between the switches is called a trunk link, and it carries the traffic of multiple VLANs. Trunk links are also used for the connection from the switch to the router.

A VLAN hopping attack results in traffic from one VLAN being sent to the wrong VLAN. Normally, this is prevented by the trunking protocol placing a VLAN tag in the packet to identify the VLAN to which the traffic belongs.

This process is shown in Figure 1.6.

The attacker can circumvent this by a process called double tagging, which is placing a fake VLAN tag into the packet along with the real tag. When the frame goes through multiple switches, the real tag is taken off by the first switch, leaving the fake tag.

When the frame reaches the second switch, the fake tag is read, and the frame is sent to the VLAN to which the hacker intended the frame to go.

FIGURE 1.6 VLAN hopping.

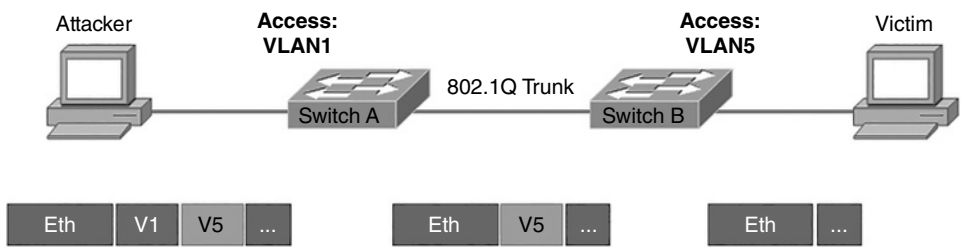

ARP Spoofing

ARP spoofing is the process of adopting another system's MAC address for the purpose of receiving data meant for that system. It usually also entails ARP cache poisoning. ARP cache poisoning is usually a part of an on-path/man-in-the middle attack. The ARP cache contains IP address to MAC address mappings that a device has learned through the ARP process. One of the ways this cache can be poisoned is by pinging a device with a spoofed IP address. In this way, an attacker can force the victim to insert an incorrect IP address to MAC address mapping into its ARP cache. If the attacker can accomplish this with two computers having a conversation, they can effectively be placed in the middle of the transmission. After the ARP cache is poisoned on both machines, they will be sending data packets to the attacker, all the while thinking they are sending them to the other member of the conversation.

Rogue DHCP

Dynamic Host Configuration Protocol (DHCP) is used to automate the process of assigning IP configurations to hosts. When configured properly, it reduces administrative overload, reduces the human error inherent in manual assignment, and enhances device mobility. But it introduces a vulnerability that when leveraged by a malicious individual can result in an inability of hosts to communicate (constituting a DoS attack) and peer-to-peer attacks.

When an illegitimate DHCP server (called a rogue DHCP server) is introduced to the network, unsuspecting hosts may accept DHCP Offer packets from the illegitimate DHCP server rather than the legitimate DHCP server. When this occurs, the rogue DHCP server will not only issue the host an incorrect IP address, subnet mask, and default gateway address (which makes a peer-to-peer attack possible); it can also issue an incorrect DNS server address, which will lead to the host relying on the attacker's DNS server for the IP addresses of websites (such as major banks) that lead to phishing attacks. An example of how this can occur is shown in Figure 1.7.

FIGURE 1.7 Rogue DHCP.

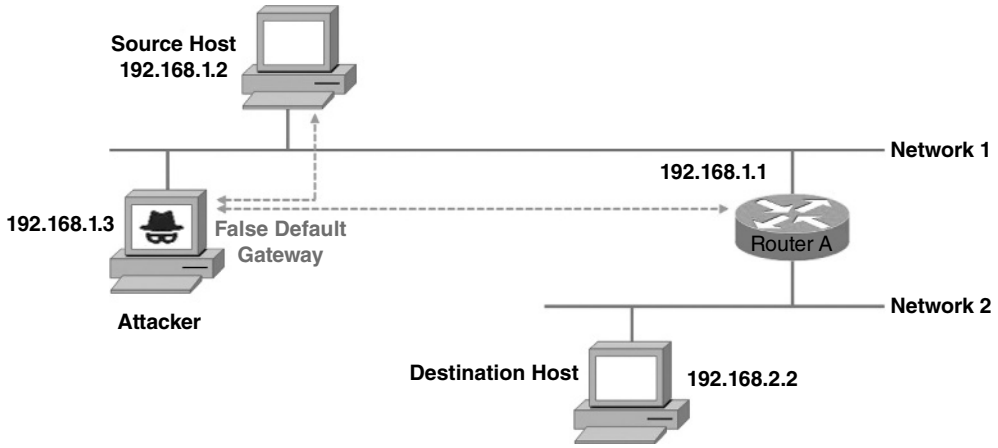

In Figure 1.7, after receiving an incorrect IP address, subnet mask, default gateway, and DNS server address from the rogue DHCP server, the DHCP client uses the attacker's DNS server to obtain the IP address of their bank. This leads the client to unwittingly connect to the attacker's copy of the bank's website. When the client enters their credentials to log in, the attacker now has the client's bank credentials and can proceed to empty out their account.

IoT Vulnerabilities

Thanks to the Internet of Things (IoT), we can finally have our fridge in the kitchen connected to the Internet! But the downside of that magic is that IoT devices tend to be made cheaply, so they don't have a lot of features you can tune, and they usually aren't built with security in mind.

Because of that, IoT devices have become quite the popular attack vector because they are usually directly connected to your network. The best way to protect yourself is to bury your Amazon Echo in your backyard, but if you would prefer to use Alexa, then the next best thing is to make sure your IoT devices are on an isolated wireless network that uses strong wireless security. That way, if someone does hack your fridge, they can't breach your actual network.

There is an enterprise version of IoT called Operational Technology (OT). OT is the other side of the IT coin that is usually staffed by nontechnical engineers who manage the infrastructure used in places such as power companies, oil and gas, and various other plants. Securing OT is the next big gold rush for IT folk because its critical infrastructure typically lacks solid network and security design.

Rogue Access Point (AP)

These are APs that have been connected to your wired infrastructure without your knowledge. The rogue may have been placed there by a determined hacker who snuck into your facility and put it in an out-of-the-way location or, more innocently, by an employee who just wants wireless access and doesn't get just how dangerous doing this is. Either way, it's just like placing an open Ethernet port out in the parking lot with a sign that says "Corporate LAN access here—no password required!"

Clearly, the worst type of rogue AP is the one some hacker has cleverly slipped into your network. It's particularly nasty because the bad guy probably didn't do it to simply gain access to your network. Nope—the hacker likely did it to entice your wireless clients to disastrously associate with their rogue AP instead! This ugly trick is achieved by placing their AP on a different channel from your legitimate APs and then setting its SSID in accordance with your SSID. Wireless clients identify the network by the SSID, not the MAC address of the AP or the IP address of the AP, so jamming the channel that your AP is on will cause your stations to roam to the bad guy's AP instead. With the proper DHCP scope configured on the AP, the hacker can issue the client an address, and once that's been done, the bad guy has basically "kidnapped" your client over to their network and can freely perform a peer-to-peer attack. Believe it or not, this can all be achieved from a laptop while Mr. Hacker simply sits in your parking lot, because there are many types of AP applications that will run on a laptop—yikes!

But you're not helpless—one way to keep rogue APs out of the wireless network is to employ a wireless LAN controller (WLC) to manage your APs. This is a nice mitigation technique because APs and controllers communicate using Lightweight Access Point Protocol (LWAPP) or the newer Control and Provisioning of Wireless Access Points (CAPWAP), and it just so happens that one of the message types they share is called Radio Resource Management (RRM). Basically, your APs monitor all channels by momentarily switching from their configured channel and by collecting packets to check for rogue activity. If an AP is detected that isn't usually managed by the controller, it's classified as a rogue, and if a wireless control system is in use, that rogue can be plotted on a floor plan and located. Another great benefit of this mitigation approach is that it enables your APs to also prevent workstations from associating with the newly exposed rogue.

Evil Twin

An evil twin is an AP that is not under your control but is used to perform a hijacking attack. A hijacking attack is one in which the hacker connects one or more of your users' computers to their network for the purpose of a peer-to-peer attack.

The attack begins with the introduction of an access point (AP) that is under the hacker's control. This access point will be set to use the same network name or SSID your network uses, and it will be set to require no authentication (creating what is called an open network).

Moreover, this access point will be set to use a different channel than the access point under your control.

To understand how the attack works, you must understand how wireless stations (laptops, tablets, and so on) choose an access point with which to connect. It is done by SSID and not by channel. The hacker will "jam" the channel on which your access point is transmitting. When a station gets disconnected from an access point, it scans the area for another access point with the same SSID. The stations will find the hacker's access point and connect to it.

Once the station is connected to the hacker's access point, it will receive an IP address from a DHCP server running on the access point and the user will now be located on the same network as the hacker. At this point, the hacker is free to commence a peer-to-peer attack.

Ransomware

Ransomware is a class of malware that prevents or limits users from accessing their information or systems. In many cases the data is encrypted, and the decryption key is only made available to the user when the ransom has been paid.

Password Attacks

Password attacks are among the most common attacks there are. Cracked or disclosed passwords can lead to severe data breaches. The end game of a phishing attack is often to

learn a password. In this section, you'll learn about the two major approaches to cracking a password.

Brute-Force

A brute-force attack is a form of password cracking. The attacker attempts every possible combination of numbers and letters that could be in a password. Theoretically, given enough time and processing power, any password can be cracked. When long, complex passwords are used, however, it can take years.

Setting an account lockout policy is the simplest mitigation technique to defeat brute-force attacks. With such a policy applied, the account becomes locked after a set number of failed attempts.

Dictionary

Similar to a brute-force attack, a dictionary attack uses all the words in a dictionary until a key is discovered that successfully decrypts the ciphertext. Manual brute-force cracking requires considerable time and processing power, and it is very difficult to complete. A dictionary attack uses common passwords or passphrases to try in the password attack, and this reduces the time and resources—that is, if it works! It also requires a comprehensive dictionary of words.

An automated program uses the hash of the dictionary word and compares this hash value to entries in the system password file. Although the program comes with a dictionary, attackers also use extra dictionaries that are found on the Internet.

You should implement a security rule that says that a password must *not* be a word found in the dictionary to protect against these attacks. You can also implement an account lockout policy so that an account is locked out after a certain number of invalid login attempts.

Advanced Persistent Threat

Probably the scariest attack in this chapter is the Advanced Persistent threat (APT), which is when a group such as a government or state actor breaches a system for a long time. These attacks tend to have a goal, so they are well planned and funded, and they are in it for the long haul.

One of the most famous examples is Stuxnet, where US and Israeli intelligence agencies built a worm that lay dormant until it found its way onto computers that ran Iran's secret nuclear program infrastructure. Then it caused havoc.

Hopefully you don't find yourself against such an entity in your day-to-day job, but all you can do is make sure your company's security posture is as strong as possible.

Hardening Techniques

There are many different hardening techniques we can employ to secure our networks from compromise. When evaluating the techniques to be employed in your network, you should keep a few things in mind: evaluate your risk, evaluate the overhead the hardening

introduces, and prioritize your list of hardening techniques to be implemented. Many of these hardening techniques are low-hanging fruit and should be employed, such as changing default passwords on network appliances and operating systems. Just make sure you have a system in place so complex passwords are not forgotten and are kept safe. Other techniques might require much more effort, such as patch management and firmware changes. In the following sections, I will introduce you to a few hardening techniques that can be used to secure your organization.

Changing Default Credentials

When installing a network device, the very first thing you must do is log into the device. There is often a standardized default username and password for each vendor or each vendor's product line. Most devices make you change the default password upon login to the device.

Changing the default password to a complex password is a good start to hardening the device. However, changing the username will also ensure that a brute-force attack cannot be performed against the default username. There are many different websites dedicated to listing the default credentials for network devices, so it doesn't take tremendous skill to obtain the default username and password of a device.

Avoiding Common Passwords

Avoiding common passwords is another simple measure to harden a device or operating system. There are several dictionaries that you can find on the Internet that will include common passwords. Some dictionaries are even collections of compromised passwords that have been made public.

When creating a password, it is always best practice to make the password at least 12 to 18 characters, based on the sensitivity of its use. You should always include symbols, numbers, and upper- and lowercase alpha characters. You should also resist substituting characters for symbols that look like the character. This substitution is often called "leet speak," and it is in every downloadable dictionary on the Internet. An example of a "leet speak" password is *p@$$word*. Another common pitfall in creating passwords is the use of words; passwords should be random and complex. An example of a complex password is *GLtNjXu#W6*qkqGkS$*. You can find random password generators on the Internet, such as `https://passwordsgenerator.net/`.

DHCP Snooping

An attack called DHCP spoofing is carried out by an attacker running a rogue DHCP server on your LAN. The rogue DHCP server has particular options set such as DNS or the default gateway in an attempt to redirect valid traffic to an exploited website to further compromise the clients. A rogue DHCP server can also be introduced to create an interruption of service.

DHCP snooping is a feature on Cisco switches as well as other vendor switches. It prevents a rogue DHCP server, also called a spurious DHCP server, from sending DHCP messages to clients. When DHCP snooping is configured on switches, all switchports are

considered untrusted ports. Only a trusted port can forward DHCP messages; all untrusted ports are filtered for DHCP messages.

Change Native VLAN

When data is transmitted on a trunk link that is not tagged with a VLAN ID, the data will default to the native VLAN. The native VLAN is also the VLAN used for switch management. The default native VLAN is VLAN 1 on all unconfigured switches from the factory; it is also the default membership for all ports on an unconfigured switch. This creates a potential security issue; if a new switch is plugged into the network and the default VLAN is not changed on the switch ports, a user will have direct access to the management network.

If the native VLAN is not changed from the default, an attacker can use the native VLAN to launch an attack using the Dynamic Trunking Protocol (DTP) or launch a VLAN hopping attack. A DTP attack is performed by plugging a rogue Cisco switch into a port that is set to default trunking negotiation. An attacker can then use the untagged packets to move data on the native VLAN of the trunk. A VLAN hopping attack is when the attacker tags the frame twice. The intended VLAN is tagged, and then the default VLAN of 1 is tagged. When the switch receives the frame on an access link, the first tag is stripped off and then the frame is switched onto the intended VLAN.

You can mitigate the risk of users being exposed to the management VLAN and VLAN hopping by changing the native VLAN to another VLAN number. It is common to change the native VLAN on trunks to VLAN 999 or another unused VLAN. Then do not use VLAN 1 for any device management and create another VLAN for that purpose.

Patching and Updates

When operating systems are installed, they are usually point-in-time snapshots of the current build of the operating system. From the time of the build to the time of install, several vulnerabilities can be published for the operating system. When an operating system is installed, you should patch it before placing it into service. Patches remediate the vulnerabilities found in the operating system and fixed by the vendor. Updates add new features not included with the current build. However, some vendors may include vulnerability patches in updates. Network devices also have patches and updates that should be installed before placing them into service.

After the initial installation of the device or operating system and the initial patches and updates are installed, you are not done! Vendors continually release patches and updates to improve security and functionality, usually every month and sometimes outside of the normal release cycle. When patches are released outside of the normal release cycle, they are called out-of-band patches and are often in response to a critical vulnerability.

Microsoft products are patched and updated through the Windows Update functionality of the operating system. However, when an administrator is required to patch and update an entire network, Windows Server Update Services (WSUS) can be implemented. A WSUS server enables the administrator to centrally manage patches and updates. The administrator can also report on which systems still need to be patched or updated.

Upgrading Firmware

When you purchase a network device, you don't know how long it's been sitting on the shelf of a warehouse. In that time, several exploits could have been created for vulnerabilities discovered. It is always recommended that a device's *firmware* be upgraded before the device is configured and put into service.

Most hardware vendors will allow downloading of current firmware. However, some vendors require the device to be covered under a maintenance contract before firmware can be downloaded. It is also best practice to read through a vendor's changelog to understand the changes that have been made from version to version of firmware.

Defense in Depth

The concept of defense in depth is not the implementation of a single security measure but the combination of several different security measures. This creates a depth to our defense strategy to make is tougher for a bad actor to launch a successful attack and gain access to sensitive data. An example of a defense-in-depth strategy is shown in Figure 1.8 as it applies to a malicious email. The malicious email should be filtered by the spam filter, but if it makes it through the spam filter, then the antivirus and antimalware software should pick up the attempt. If none of the prior detections work at stopping the email, the end-user training should prevent the malicious link in the email from being clicked. However, if it is clicked, then the content filtering software should prevent the exploit from calling back to the server hosting the malware.

FIGURE 1.8 Example of defense in depth.

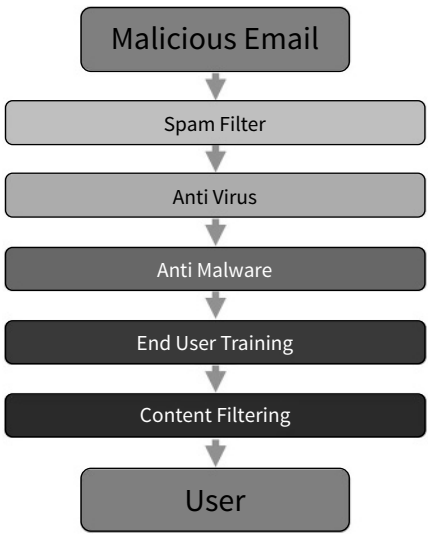

Security in this example is not left to a single mechanism. The collection of these mechanisms makes it tougher for a bad actor to succeed. The following sections cover several different security measures according to the CompTIA Network+ exam objectives, but they in no way cover all of the different security measures out there. It is always best to assess what you are trying to secure first and then identify the weaknesses. Once you understand the weaknesses, security measures can be implemented to address each weakness identified. The result will be a very diverse set of security measures and a defense-in-depth strategy.

Social-Based Attacks

Despite what we see in the movies, most breaches aren't due to a bad guy pressing a few keys on their keyboard and suddenly hacking into the Pentagon. They usually just rely on some of these low-tech attacks.

Social Engineering

The sad reality is that no matter how much money we spend on advanced security solutions to prevent bad actors from messing with the business, the average user can easily be tricked into doing something that can compromise everything.

Social engineering is really just a fancy term for being a conman, where the attacker simply talks their way into getting what they need. This comes in many forms, but essentially the bad actor will pretend to be someone in authority to gain access to information they shouldn't have access to. A common example of this is the hacker phoning a user and claiming to be their tech support so they can try to get the user's account information.

Another popular social engineering attack is the attacker pretending to be a customer and bluffing their way through tech support so they can have the password reset. A good example of this is when a large energy company, Suncor, was crippled because a hacker convinced tech support to reset the Azure Global Administrator's password (they also reset the multi-factor authentication on the account), which gave the hacker complete access to the company's Microsoft infrastructure! Within a couple minutes, the hacker locked out all the IT staff and was able to wreak havoc on all the computers because they were managed by Azure.
I certainly wouldn't want to be that tech support person!

The best way to combat this threat is to continuously train staff to be on the lookout for social engineering attempts. For example, there is never a reason for IT to ask a user for their credentials. So, if someone who claims to be in IT asks for that kind of information, users should be taught to hang up and report the incident to the team that handles security.

Insider Threats

Sometimes the bad actor actually works for the company, which makes life harder because they are already on the network. Because of this, a disgruntled worker might take advantage of their existing access to the company systems, or even their position, to do things that can cause all kinds of problems.

A common example of this is when IT staff are laid off: they might leave some going-away presents such as changing passwords or otherwise setting up a time bomb that will take a system offline if not found in time. Hopefully it goes without saying, but never do this! Losing your job is not worth going to prison or getting a large fine.

The best way to reduce the risk of malicious insider threats is to enforce strict least privilege wherever possible; by ensuring that people can't access things they don't need for their day-to-day job, there will be less overall risk. It is also a good idea to have some kind of security monitoring so you can tell if users are misbehaving.

A different version of an insider threat is a user who doesn't mean to be malicious but can compromise your security anyway. Imagine a user brings their personal laptop into the office and decides to directly connect the laptop to the Ethernet jack in their office. Well, if the laptop has malware in it, then it's possible that it can spread across the network and infect all the systems.

My favorite one is when a user wants better wireless in their office area and decides to buy a cheap home router like a D-Link, instead of asking IT to install a proper access point. If they plug the D-Link into the Ethernet jack, then they just created a rogue DHCP server that will accidentally knock all the endpoints in the VLAN offline due to them getting the wrong IP information. This is because the cheap router will almost always be closer to the endpoints than the corporate DHCP server, so it will answer DHCP requests before the proper one can respond.

The best way to handle accidental insider threats so to define strict policies that specify what an employee is and isn't allowed to do; this is typically called an acceptable use policy (AUP). The idea is by explicitly saying that the employee can't use non-IT issued equipment, they can be disciplined if they break the rules. Beyond that, solutions such as Network Access Control (NAC) can help prevent unauthorized devices from being added to the network.

Phishing

Phishing comes in many different forms, but the main idea is that the attacker will send the victim an email that asks the user to click a link that will typically either download some malware or take them to a login page that will steal their credentials if they enter their username and password.

To make things more believable, the attacker will make the email look somewhat official and will pretend to be something urgent like a message from "Netflix" saying your account is delinquent and will be shut down if you don't click the link immediately. You can usually tell a phishing email from the real thing very easily because they tend to be lazy representations that are full of typos and sometimes just flatout use wrong logos or colors.

A more sure-fire way of identifying a phishing email is to look at the email address you received the mail from. The address is usually spoofed by the attacker to make it appear more legitimate. Although it may say something like support@netflix.com..., if you hover your mouse over it, you might see the real address is thisisafakenetflixsupportemail@gmail.com. Likewise, just like social engineering, your bank is never going to ask you for your banking password. So any requests like that can be immediately discarded.

There are also a couple of spin-off versions, such as "smishing," which is the same thing except the attacker texts it to your phone instead. I currently get about 10 texts a day that range from saying my packages can't be delivered, to my online services being shut off for non-pay, to the government deciding to give me money! In fact, I looked at my phone while writing this and found this message:

> *Netflix:We couldn't process your monthly payment, so your account is temporarily on hold. Please update your billing information here: https: Netflix-Canada-update.com*

We also have "vishing" attacks, where you are called on the phone by the attacker instead of them sending you a text or an email. If you have ever received a spam call from "Microsoft" saying they detected a problem with your computer and you need to give them remote access to your computer so they can fix it, they tried to social engineer you into compromising your computer. If you fall for it, they will probably install some lovely malware for you!

The last type of phishing that we will talk about is called spear phishing, and it's basically just a targeted form of phishing. Instead of sending out random emails to see if someone falls for it, the attacker will do more research and send a targeted message to achieve a result.

A frequent example that I see is when the attacker spoofs an email from the company's CEO that is sent to the finance department at 10:50 a.m. The message will say something like "Please pay this $500,000 invoice immediately as the deal needs to close by 11:00 AM. – Mr. CEO"; because of the short deadline, finance might pay the invoice without doing any verification, and the money really goes to the attacker.

Just as with social engineering, the best way to combat the threat is to continuously train users on how to detect phishing attempts and how to deal with them. Security teams frequently will send their own phishing emails to their employees in order to assess them, and those who click the link are treated to some mandatory security training.

Vishing

A vishing attack is similar to a phishing attack, with a small twist. The threat actor will use their voice to attempt to steal credentials, in lieu of a well-crafted email. Typically, the threat actor will call a helpdesk and pretend to be an employee. They will plead with the help desk person to help them change their password, or they could lose their job—something along those lines to pull the heartstrings (so to speak) of the help desk person. These types of attacks are really more social engineering than anything else.

Real Life

In September 2023, a really large cyberattack was carried out on MGM Resorts International. This cyberattack was aided by a successful vishing attempt on MGM's IT helpdesk.

Smishing

The reason these are called smishing attacks is that the Short Message Service (SMS), also known as texting, is the delivery method. A smishing attack involves well-crafted text messages to fool a victim into logging into a phishing site. Other messages might coax them into purchasing and sending gift cards. There are many different variations on this method of attack. If you can motivate a person to carry out an act, it's probably been done via smishing. This is also another good reason to never publish your cellphone number on social media sites such as Facebook or LinkedIn.

Spear Phishing

Another form of phishing is spear phishing. Spear phishing is when the attacker uses information that the target would be less likely to question because it appears to be coming from a trusted source. Suppose, for example, that you receive a message that appears to be from your spouse that says to click to see that video of your children from last Christmas. Because it appears far more likely to be a legitimate message, it cuts through your standard defenses like a spear, and the likelihood that you would click this link is higher. Generating the attack requires much more work on the part of the attacker, and it often involves using information from contact lists, friend lists from social media sites, and so on.

Environmental

Some attacks become possible because we have created an environment that they can develop in and be acted out upon. An example of environmental behavior is the urge for people to pick up a wallet on the street. They might have malicious intent, or they might not. However, they will pick it up and look inside. The environment is a public place, and there is potential for value lost. The following are issues that are created by user behavior that can be carried out by malicious threat actors.

Tailgating

Tailgating is the term used for someone being so close to you when you enter a building that they are able to come in right behind you without needing to use a key, a card, or any other security device. Many social-engineering intruders who need physical access to a site will use this method of gaining entry. Educate users to beware of this and other social-engineering ploys and prevent them from happening.

Access control vestibules (mantraps) are a great way to stop tailgating. An access control vestibule is a series of two doors with a small room between them that helps prevent unauthorized people from entering a building.

Piggybacking

Piggybacking and tailgating are similar but not the same. Piggybacking is done with the authorization of the person with access. Tailgating is done when the attacker sneaks inside without the person with access knowing. This is why access control vestibules (mantraps) and turnstiles deter tailgating, and live guards and security training deters piggybacking.

Shoulder Surfing

Shoulder surfing involves nothing more than watching someone when they enter their sensitive data. They can see you entering a password, typing in a credit card number, or entering any other pertinent information. The best defense against this type of attack is simply to survey your environment before entering personal data. Privacy filters can be used that make the screen difficult to read unless you are directly in front of it.

Malware

Malware is a broad term describing any software with malicious intent. Although we use the terms *malware* and *virus* interchangeably, distinct differences exist between them. The lines have blurred because the delivery mechanism of malware and viruses is sometimes indistinguishable.

A virus is a specific type of malware, the purpose of which is to multiply, infect, and do harm. A virus distinguishes itself from other malware because it is self-replicating code that often injects its payload into documents and executables. This is done in an attempt to infect more users and systems. Viruses are so efficient in replicating that their code is often programmed to deactivate after a period of time, or they are programmed to only be active in a certain region of the world.

Malware can be found in a variety of other forms, such as covert cryptomining, web search redirection, adware, spyware, and even ransomware, and these are just a few. Today the largest threat of malware is ransomware because it's lucrative for criminals.

Ransomware

Ransomware is a type of malware that is becoming popular because of anonymous currency, such as Bitcoin. Ransomware is software that is often delivered through an unsuspecting random download. It takes control of a system and demands that a third party be paid. The "control" can be accomplished by encrypting the hard drive, by changing user password information, or via any of several other creative ways. Users are usually assured that by paying the extortion amount (the ransom), they will be given the code needed to revert their systems back to normal operations. CryptoLocker was a popular ransomware that made headlines across the world, as shown in Figure 1.9.

FIGURE 1.9 CryptoLocker.

You can protect yourself from ransomware by having antivirus/antimalware software with up-to-date definitions and by keeping current on patches.

Summary

In this chapter, you learned common attack types that you might expect on an enterprise network. These attack types can be categorized into technology-based attacks and those that are the result of human failure or of the network environment that exists.

Technology-based attacks include denial-of-service (DoS)/distributed denial of service (DDoS) attacks, on-path attacks, DNS poisoning, VLAN hopping, ARP spoofing, rogue DHCP, rogue access point (AP), evil twin, ransomware, and password attacks.

Human and environmental attacks include social engineering, phishing, tailgating, piggybacking, and shoulder surfing.

Exam Essentials

Explain common technology-based attacks. These include (DoS)/distributed denial-of-service (DDoS) attacks, on-path attacks, DNS poisoning, VLAN hopping, ARP spoofing, rogue DHCP, rogue access point (AP), evil twin, ransomware, and password attacks.

Describe (DoS)/distributed denial of service (DDoS) attacks. This includes the architecture and behavior of a botnet and of the role of the command-and-control-server.

Identify human and environmental attacks. These include social engineering, phishing, tailgating, piggybacking, and shoulder surfing.

Review Questions

The following questions are designed to test your understanding of this chapter's material. For more information on how to get additional questions, please see `www.lammle.com/ccst`.

You can find the answers to these questions in Appendix.

1. Which of the following is *not* a technology-based attack?
   - **A.** DoS
   - **B.** Ping of death
   - **C.** Shoulder surfing
   - **D.** Malware
2. A command-and-control server is a part of which of the following attacks?
   - **A.** DDoS
   - **B.** Ping of death
   - **C.** Shoulder surfing
   - **D.** Malware
3. Which of the following is a DoS attack that floods its victim with spoofed broadcast ping messages?
   - **A.** SYN flood
   - **B.** Smurf attack
   - **C.** Land attack
   - **D.** Ping of death
4. Which of the following is an attack that inundates the receiving machine with lots of packets that cause the victim to waste resources by holding connections open?
   - **A.** Ping of death
   - **B.** Zero day
   - **C.** Smurf attack
   - **D.** SYN flood
5. In which of the following does the attacker (and their bots) sends a small spoofed 8-byte UDP packet to vulnerable NTP servers that requests a large amount of data (megabytes worth of traffic) be sent to the DDoS's target IP address?
   - **A.** SYN flood
   - **B.** NTP amplification
   - **C.** Smurf attack
   - **D.** DNS amplification

6. Which of the following was previously known as a man-in-the-middle attack?
   - **A.** VLAN hopping
   - **B.** On-path attack
   - **C.** LAND attack
   - **D.** Smurf attack

7. Double tagging is a part of which of the following attacks?
   - **A.** VLAN hopping
   - **B.** Smurf attack
   - **C.** DDoS
   - **D.** Malware

8. Which of the following is the process of adopting another system's MAC address for the purpose of receiving data meant for that system?
   - **A.** Certificate spoofing
   - **B.** ARP spoofing
   - **C.** IP spoofing
   - **D.** URL spoofing

9. Which of the following is connected to your wired infrastructure without your knowledge?
   - **A.** Rogue AP
   - **B.** Command-and-control server
   - **C.** Zombies
   - **D.** Botnet

10. Which of the following uses the same SSID as your AP?
    - **A.** Rogue AP
    - **B.** Rogue DHCP
    - **C.** Evil twin
    - **D.** Zombie