

Chapter 8 – Networking Threats, Assessments, and Defenses

What this section is about – what you should learn

In this chapter you will learn about another big category of attacks called Networking Threats, which are attacks that are made on the network infrastructure as opposed to attacks on an individual computer. The book then covers how to assess a network to see if it's vulnerable to these attacks, and ways to defend against them.

In some ways this chapter might seem like the junk drawer in your kitchen, as it has a bunch of seemingly unrelated stuff, like network protocols, Linux commands, and door locks. These are all important subjects, but unless you look at them from a specific perspective they might seem to be totally unrelated. The perspective you need to see how these are all related is to think about the infrastructure of the Internet, threats against it, and how to protect the servers and devices that provide the infrastructure.

The information in this chapter is different than what you've learned so far, as it requires some basic understanding of the infrastructure and how the Internet and TCP/IP function. And once again you're given a pretty large brain dump of information regarding all of the attacks. There are details on things like DNS and ARP poisoning attacks that you can only understand after you learn the details of network communication. Don't worry if some of the explanations don't make a lot of sense at this point. Remember this class is an introductory class, and the expectation is that you'll get a general sense of what network attacks are, as well as start to learn some of the terminology. To really understand how the attacks occur and the weaknesses they exploit requires taking several other classes where you'll learn about the underlying technologies, and months or years of experience. At this point you're just dipping your toes in the ocean, but soon you'll be in the water up to your neck, or over your head, or floating your boat ... however that saying goes. You'll learn all about the networking and the Internet infrastructure in later classes, but you'll need to know a little now to understand the threats talked about in this chapter. So, one of my major goals for this section is to help you learn about how the Internet works by looking "under the hood" at some of the Internet Infrastructure.

Internet Infrastructure

To understand what the Internet infrastructure is let's use two analogies, snail mail and phone calls. To start, let's look at what we're trying to accomplish. On the Internet, we want to deliver digital information from one computer to another. This information might be used for things like e-mail or web traffic, or login information, etc. It might be used for hundreds of different purposes; the point is that we need to get this data from one computer to another.

This is very similar to real life snail mail, where we send letters or packages from one person to another, or make a phone call to exchange information with another person. Let's look at the

infrastructure that makes this possible with snail mail first, as it's very similar to what happens on the Internet.

IP Addresses

The first thing to understand about snail mail and its infrastructure is that it's made possible by the fact every home has unique address. In the United States this address consists of a few parts which we all learned about in elementary school:

Name
Street Address (and possibly the apartment number)
City, State
ZIP Code

If we didn't have unique addresses you could send a letter or package to someone and hope that they get it. This happened way back in the historical past, before the postal system matured, where you might send a message addressed to someone like Sir Victor in Wessex, with the hope that it will be delivered. But now, using the addressing system developed over the years in the US, we can be relatively assured that each home or building has a unique address.

Also note that the components of a mailing address help us identify a specific location by first identifying a state, then identifying the city which is a location within the state, then identifying the address which is a location within the city. The designers of the system could have just assigned a single number to every building in the US, but think how crazy and difficult to use that system would be. It's much easier to use a hierarchical system that identifies a large zone like the state, followed by an intermediate zone (the city), and finally the building address.

Instead of home addresses, the Internet assigns each computer a numeric address commonly called an IP address. IP stands for Internet Protocol in this case, not Intellectual Property. An IP address is a set of 4 numbers that each range from 0-255, separated by a . or period. For example, the IP address for the web server at Columbia Basin College is 212.69.159.108. And much like mailing addresses, one set of numbers in the IP address will identify a network, and the remaining numbers will identify a specific host on the network. In the CBC network address the 212.69.159 portion of the IP address identifies the network used by the colleges in the state of Washington, while the 108 identifies the specific computer, the one hosting the college's web server.

DNS – Domain Name System

The next piece of Internet infrastructure that you need to learn about is the Domain Name System, or DNS it's more commonly known. This is the system that allows humans to use names instead of the IP addresses, when we want to communicate with another computer. You just learned that if you want to communicate with another computer on the Internet, you need to know that computer's IP address. The problem is that for most humans, IP numbers are hard to

remember; most humans are much better at remembering names. It's just like with phone numbers, it's much easier to remember someone's name than it is to remember their number. You might remember a few phone numbers, for people or places you call frequently, but most people would struggle to remember phone numbers or IP addresses. For example, if you wanted to use Google to do a search, is it easier for you to remember google.com or 172.217.14.238. The other nice thing about names, is that they remain fixed, while a person's address or phone number might change.

However, the problem with names is that they're not what's required to make a phone call or get network data where it need to go. With addresses and phone numbers, when we knew a name but not a number, we used to use something called a phone book. The phone book listed names in alphabetical order, along with that person's address and phone number. If you're old enough you might remember the phone book, but with the advent of cell phones and the Internet the phone book has gone the way of the pay phone and is no longer with us. There was a phone book for each city or town, or maybe a few towns in you lived in a sparsely populated area.

The Internet started out just using IP addresses, but this was simple when there were only 4 or 5 computers connected to the Internet. This was back in the late 1970s. But as the Internet started to grow a few people recognized the need for a system to allow the use of names, and a way to convert the names to IP addresses. In 1983 the DNS system was introduced, and just like the phone book, it provides a way to convert DNS names to IP addresses. DNS is part of the Internet infrastructure, and if you are in charge of a network that you want to connect to the Internet one of the things you must do is run a DNS server that has the names and IP addresses of the computers on your network.

As you read the chapter, you'll see that there are some attacks that have been directed at the DNS system itself, and other attacks that make use of DNS to try and break into other devices. The attacks directed at DNS have the goal of getting users to go the wrong computer, one that's set up to look like the real computer, and then provide their secure information. This would be similar to an attacker being able to change the phone book, and publish a fake street address for a business. The attacker would try and change the address for a business like a credit card company, where people would send bank checks to make their monthly payments. I don't think this phone book spoofing actually happened, but it should help explain what's happening with the DNS attacks.

Network Routers

Another key part of the Internet infrastructure are the routers that connect the different smaller networks, and move IP traffic from one network to another. You may be familiar with network routers if you have a wireless router in your home. Your home router communicates with all the devices in your home, allowing them to "talk" to each other, but the router's main purpose is to allow your devices to communicate with other computers on the Internet. Your router doesn't connect to other web servers or devices directly, but instead talks to another router run by your ISP. This first router at your ISP also talks to other routers, building a

network of networks. The network routers all talk with each other, building routes or paths from one network to another. When your computer wants to communicate with another device on the Internet, your router, and several other routers will each forward your network packets along the path, until they reach their final destination. You'll learn all about the process the routers use to build the routes and paths in later classes. For now, you need to realize that there will be several routers that handle your network packets as they get forwarded along the path to their final destination.

A good way to visualize what the network routers do, and how this makes your network communications vulnerable, is to go back at looking at how the US Postal service delivers snail mail. Let's say you live in Kennewick Washington, and you want to send a letter to your brother Gustavo, who lives in Bradenton Florida. You leave the letter in your mailbox, where it's picked up by your neighborhood's postal worker. This worker won't deliver the letter to Florida for you, but they will start it on its way.

Your neighborhood postal worker is like your home router, they don't deliver things from end to end, but they look at the destination address, and then decide where to send the message next. In the case of the snail mail, the letter will first be sent by truck to another US Postal processing center in Spokane Washington. The Spokane processing center will look at the destination, and decide that to get the letter to Florida it first needs to go to Seattle Washington. Spokane will bundle your letter with all of the mail that needs to go to the processing center in Seattle, put it in a truck, and drive it to Seattle. The Seattle processing center will look at the letter's final destination, and decide that it needs to next go to a processing center in Denver. The Denver processing center will look at the destination, and decide that the letter next needs to go to Miami. When the letter arrives in Miami, their processing center will look at the address in the destination, and route the letter to Bradenton. When the letter arrives in Bradenton, it will be given to one of local carriers, who will look at the address, and deliver the letter to your brother.

You don't need to memorize the route the letter took from Kennewick to Bradenton, not because it's probably wrong, but because the specific route isn't really important. What's important is that you see that several people and processing centers had to look at your letter, and decide the next intermediate destination for the letter along its path to the final destination. A very similar process happens with the routers on the Internet. There will be several routers that have to handle any network packets you send to another computer, and each of them will look at the final destination address to determine the next intermediate destination for the network packets.

For example, say you've made a connection from your home computer to the web server at the University of Washington in Seattle. I'm going to assume that you get your Internet through Charter, and I'm also going to simplify the route your packets take and make up some router and network names. You'll learn how to see the actual networks your packets go through in a video I've made for you, but for now I want to simplify things to make them understandable. Ok, in this example the first place the network packets from your computer will go is to your

home router. Your router will look at the IP address of the destination for the network packets, see they're bound for the UW web server, and then send them to the Charter router for your neighborhood. This router will check the destination IP address, see they're bound for the UW web server, and send them to a router in Yakima, which will forward the packets to a router in Seattle, which will send the packets to the main router for UW, which will send the packets to the UW's web server.

Once again, it isn't critical that you know the specific route the network packets take. The important thing is that you understand that any time your computer communicates with another computer on the Internet your network packets will be seen by several different routers and networks.

If you want to see the route your network packets take, and all the routers that have access to your network packets, open a Command Windows and use the `tracert` command. If you type `tracert destination`, where *destination* is the name or IP address of the computer you want to communicate with, you'll see a list of the networks your packets traverse. The following figure shows the route packets from my home computer take between my computer and one of the computers at Google.com.

```
C:\WINDOWS\system32>tracert google.com

Tracing route to google.com [172.217.14.238]
over a maximum of 30 hops:

  0  0 ms  0 ms  0 ms  192.168.1.1
  1  2 ms  1 ms  1 ms  192.168.1.1
  2  6 ms  4 ms  3 ms  fdr01.knwc.wa.nwestnet.net [50.46.181.118]
  3  3 ms  4 ms  5 ms  cr1-knwcwaxa-b-be500.bb.as20055.net [64.52.98.2]
  4  9 ms  8 ms  10 ms  cr1-knwcwaxa-b-be-13-bb.as20055.net [137.83.80.144]
  5  11 ms  9 ms  10 ms  cr1-yakmwafp-b-be-12.bb.as20055.net [107.191.236.84]
  6  13 ms  9 ms  10 ms  cr2-sttlwawb-b-be-14.bb.as20055.net [107.191.236.48]
  7  9 ms  10 ms  12 ms  cr2-sttlwawb-b-n-be-10.bb.as20055.net [107.191.236.127]
  8  10 ms  10 ms  9 ms  pr1-sttlwawb-b-be-10.bb.as20055.net [137.83.80.99]
  9  10 ms  9 ms  11 ms  google-sttlwawb-b.pni.as20055.net [107.191.239.2]
 10  9 ms  8 ms  9 ms  74.125.243.193
 11  9 ms  11 ms  9 ms  209.85.254.249
 12  9 ms  9 ms  12 ms  sea30s02-in-f14.1e100.net [172.217.14.238]

Trace complete.
```

The router names might look non-sensical, but if you live in Southeast Washington you can probably deduce some information about the router locations from their names. For example, routers with `knwc` in the name are probably somewhere in Kennewick. The router with `yakma` is probably in Yakima, and the routers with `sttlwa` in the name are probably in Seattle.

Once again, the specific routers don't matter. What you need to take away from this is that any time your computer communicates with another computer on the Internet the network packets

will be handled by several different routers, each of which can read and change the data in the packets. When it comes to network attacks and threats, you might be able to see the potential problems that could arise when you have to trust multiple intermediate devices to deliver your messages. With snail mail, any of the postal workers or processing centers could open your mail, read the message, possibly change it, and then send it along its way. The same threat exists for network communications, where anyone with access to one of the intermediate routers can look at your network communications, possibly change them, and then send them along their way. This is what happens in the Man In The Middle Attack.

Hopefully you can see that these types of network attacks are a major concern. But, before you get so concerned that you quit using the Internet here's some good news. The good news is that most network connections are now encrypted. This means that while it's possible for an attacker to see your network packets, the data in the packet will be scrambled and there's no way for the attacker to read or change it.

MAC Addresses and ARP Poisoning

In addition to IP addresses, there's another type of network address to learn about. This is something called a Media Access Control (MAC) layer address. MAC addresses look much different than IP addresses, and have the form `xx-xx-xx-xx-xx-xx`, where each `xx` is a hex number between 00 and FF. You might also see them written with the colon character as a delimiter instead of the `-`.

The easiest way to explain MAC addresses is to compare them to the IMEI (International Mobile Equipment Identity) number on your cell phone. The IMEI is a unique number identifying your phone, separate from the phone number. That is, you can change your phone number, but the IMEI is burned on the phone and cannot be changed. It's not on the SIM card either, so swapping SIMs will not change a phone's IMEI.

Or, in case you don't know about phones and IMEI numbers, here's another analogy that may help. For this analogy let's consider a person's identity. Each person is a unique individual, regardless of their phone number or mailing address. You can move, or change your phone, but you'll still be you. In this analogy, your identity is like a computer's MAC address. And, if your name is Mac and you have a Mac computer it would be Mac's Mac's MAC address. ☺

All silliness aside, a computer's network card has a hard-coded address called the MAC address. While you can change the IP address on a network card or computer, you can't change the MAC address. Changing the IP address is done all the time as computers are moved from place to place. For example, say you have a laptop that you take with you everywhere. You use your laptop at home, at school, at the mall, at your relative's houses, etc. Each time you connect your laptop to a new network, there will be a device on that network that assigns your computer a new IP address associated with that network. However, even though the computer is getting a new IP address on each network, the computer's MAC address remains constant.

Going back to the analogy where we compare the MAC address to your identity, let's say you go on a trip around the US, staying in a different hotel each night. If someone wants to call you at the hotel, you'll have a different number each night, but you'll still be you. If the phone in the room rings, one of the things you'll want to do is make sure that the caller wants to talk to you, and not someone in a different room, or a previous occupant of the room you're in. This verification is pretty simple to accomplish, you and the caller will do a little verbal handshake where you ask each other to verify your names. Something like, "Hi, is this Amina?". "Yes, who is calling please?".

With networks, the MAC addresses are used between any two devices on the same network. For example, going back to the example where you are using UW's web server, when your computer sends a packet to your home router it will send it to your router's MAC address. That is, the final destination will be given by the IP address, but the first "hop" in that journey is between your computer and the network router. So, the network packets will have two destinations, the final destination which is specified by the IP address, and the next direct hop destination which is specified by the MAC address. When your router receives the network packets it will leave the IP address for the final destination unchanged, but it will change the MAC address for the next direct hop to the MAC address for the Charter router for your neighborhood. The Charter router for your neighborhood knows that it should look at the packet because it was sent to that router's MAC address. Each router along the path will repeat this process for any packets sent to its MAC address. That is, it will look at the IP address of the final destination, then change the destination MAC address to the device which is the next hop in the path.

Here's another analogy that might help. When you send a letter to Florida the final destination address is Florida. But the next hop in the path will be Spokane. The Post Office will change the next hop address to Spokane, and send the letter there. Spokane will look at the final destination, and see that the next hop in the path should be Seattle, so it will change the next hop address to Seattle, and so on, until the letter finally reaches your brother in Florida.

There's one last piece of information you need before we get back to the network attacks. You also need to know what the Address Resolution Protocol (ARP) does. This is the way that devices on a network find out the MAC address for other devices. They do this by doing something called a broadcast, which is like shouting in a room so that everyone will hear. In this case, your computer will know the IP address for your home router, but it won't know the router's MAC address. So, your computer will broadcast on the network, asking the device with IP address of the router to send back its MAC address. Once again, this is like shouting in room, and asking for the person named Cory Wong to give you their phone number. The difference is that your router will respond, but I'm pretty sure you won't get any phone numbers using the shouting method.

When your computer gets the MAC address for the router it stores it in a table in memory. This way your computer will have quick access to MAC address the next time it needs to send network packets to the router, and it won't have to do another broadcast to get it. The table

that stores the MAC addresses on your computer is called the ARP table. If you want to view the ARP table on your computer, open a Command Window and type: `arp -a`

The following table shows an example ARP table.

```
C:\WINDOWS\system32>arp -a

Interface: 192.168.1.45 --- 0x4
 Internet Address      Physical Address      Type
 192.168.1.1           3c-37-86-23-d0-c2    dynamic
 192.168.1.2           00-15-99-11-41-49    dynamic
 192.168.1.5           90-a8-22-93-99-6c    dynamic
```

ARP poisoning is when an attacker changes the MAC address in one of the ARP tables, so network packets are sent to a different device, a device of the attacker's choosing. This won't make a difference on your home network, but it could be a problem on a public network like those at an airport or coffee shop where there might be multiple routers. It's also a problem on the Internet, where the routers connecting multiple networks might be tricked into sending packets to the wrong router.

Denial of Service (DoS) and Distributed Denial of Service Attacks (DDoS)

Another type of network attack discussed in this chapter are the Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks. As the book says, these are nuisance attacks that prevent a site from processing valid traffic. A good analogy for a DoS attack would be your phone. Say someone keeps calling you and hanging up when you answer. This will be annoying, but the bigger problem is it will also prevent you from answering valid calls. If you get attacked on your phone this way, you can simply block the number and the attack will stop. To get around this, the attacker can recruit thousands of other people to also make nuisance calls to your number, or in other words distribute the attack. This is very similar to what happens with DoS and DDoS attacks. But instead of making a phone call, in a DoS or DDoS attack the attacking computer starts the process of opening a network connection with the target computer. The target computer will respond, but the attacking computer never completes the network connection. The target computer only has a limited number of network connections it can open at any time, so by tying them up, even briefly, the attacker is able to prevent valid users from connecting. Like a phone call from a single phone number, a DoS attack is relatively easy to block, which stops the attack. This is why the attackers evolved to using thousands of computers to all attack at the same time in a Distributed Denial of Service attack. In a DDoS attack it's very difficult, maybe even impossible, to block all of the devices involved in the attack.

Hopefully this explanation makes sense. But the main point of this is, and the thing that should be of interest to you and anyone who owns a computer or smart device, is where do these thousands of attacking devices come from? The devices used in the attack are devices that are

owned by unsuspecting members of the general public, people like you and me. We don't know our devices are being used, and certainly wouldn't explicitly agree to let them be used as part of an attack. But if you don't take some basic security precautions it's quite possible for an attacker to plant some code on your device. This code allows the attacker to take control of your device at any time, and have your device participate in the attack. You might not even notice when this occurs as each individual device only generates a relatively small amount of network traffic during the DDoS attack.

If an attacker is able to compromise a device and plant the DDoS code, the device is called a zombie. The network they make up is called a zombie net and it is controlled by the attacker. The DDoS code lies dormant until it receives a message from the controller, telling it which computer to attack. At this point all of the zombies perform a DoS attack, typically overwhelming the target.

Another interesting fact about DDoS attacks is that any smart device connected to the Internet can become a zombie. This has been a problem because many smart devices aren't secured, and there might not even be a way to secure some of the more poorly designed devices. The people making the smart devices want to make them easy to use, and often don't even take security into consideration during early development. Sooner or later they realize that they'll need to add some security, but it's much easier to make sure things are secure from the start and can be much more difficult to retrofit once a product is built.

So, how do you protect your devices from becoming zombies? Just do the normal basic security items listed above, stay away from sketchy web sites, don't reply to texts from people you don't know, etc. If you have any smart devices or routers, make sure their firmware is up to date.

As far as defending yourself from DDoS attacks, you probably don't have to worry about your home devices being attacked. The big worry is for corporate, government and military networks. The problem in protecting these networks and devices is that if you are attacked with a DoS or DDoS attack there's not much you can do at the time of the attack. You might find articles on the Internet that promote methods for protecting networks from DDoS attacks, but it's not really possible at this time.

The only real defense is to physically distribute your network servers and spread them around different parts of the Internet. Take Amazon for example. Amazon doesn't use a single server to host amazon.com. Instead, they have dozens of servers located in different data centers around the country. Using multiple servers allows Amazon to distribute the load which might overwhelm a single server. It also allows users to connect to servers that may be physically closer to them on the Internet, that is a server that takes fewer network hops to reach. These servers are all linked together to keep their databases synched, so users don't know or see that there are multiple servers. And a DDoS attack launched against one Amazon server won't impact the other servers. Or, if a DDoS attack is launched against all the Amazon servers the impact will be diluted. While this defense can help mitigate against DDoS attacks, replicating

and syncing servers across the Internet is going to be expensive. Plus, it has to be set up ahead of time. It isn't something you can decide to do at the last minute if your network comes under attack.

So, while you probably don't have to worry about your personal devices being attacked you should worry about preventing your computer or any other devices you own from being part of the attack. I have a friend, Troy Thompson, who used to be the Director of Cyber Security at PNNL/Battelle, and this was his idea. He thought that if your device was converted to a zombie and used in a DDoS attack you should have to pay a fine. After all it's pretty simple to protect your devices, all you need to do is run anti-virus software and keep your OS and applications up to date. So, if you're not taking the basic security precautions and your device becomes a zombie and is used in an attack you should pay a fine. There are already similar laws for other devices that set the precedent. For example, several states have laws that allow you to be ticketed if you leave your car running and unattended. Part of the reason for these laws is environmental, but another reason for the law is it makes it too easy for someone to steal the car.

<https://www.goodhousekeeping.com/life/news/a41993/car-idling-laws/>

In any case, what do you think? Should you be fined if a device you own is used in DDoS attack, and you failed to perform even the basic security tasks? You don't have to answer this formally, or upload your answer anywhere. It's just food for thought.

There's one last thing I want to mention about DDoS attacks, but before I do I need to give you a warning. This warning is one that will be repeated throughout your cyber security classes as you learn how to do things that can be illegal and get you in serious legal trouble if you use them incorrectly or on systems where you do not have prior written authorization. The information you are about to read or view will teach you how to use tools or processes that can be used to perform actions that are only legal if you conduct them against systems and in situations where you have clear, written, prior authorization to perform the action(s) from a person who is legally authorized to grant the permission. For example, you may learn how to crack passwords, launch Denial of Service or DDoS attacks, redirect network packets, encrypt data, plus many others. Performing these actions without prior authorization is a legal violation and may result in local, state or federal criminal charges, fines, and/or jail time. You should also note that many of the actions leave clues that can easily be traced, making it relatively easy for authorities to catch the person responsible for the attack.

To avoid legal jeopardy, you must only use these tools on or against your personal systems, or systems set up by the college for testing where you have explicit permission to run the tests. Note that your personal systems mean those devices owned and operated by you, but do not include those of your other family members.

Ok, enough of the scary warnings, let's get back to DDoS attacks. One of the scary things about DoS and DDoS attacks is how easy they are to launch. If you're serious about attacking a large

organization you can rent a Botnet to launch a large-scale attack¹, and the cost is surprisingly cheap. Some of the sites that do this pretend that they're legitimate by assuming that you're using the tools to stress test a site or network that you own, but this is usually just a pretense to try and protect themselves from legal charges prosecution. There are DDoS attack tools that you can download and use to launch an attack. These tools are easy to use, all you need to do is provide a target DNS name or IP address, and then press a button².

The Junk Drawer

The remainder of the chapter covers what seems to me to a dozen unrelated subjects, ranging from PowerShell scripts, to a handful of Linux commands, to robots and locks for your laptop. These are all loosely related by the fact that they're either used in network attacks, or used to protect servers that host an organization's Internet infrastructure.

My suggestion is to read through this material and get what you can out of it. You should be able to easily understand the information on Physical Security Controls, but don't stress about trying to memorize anything else as you will have entire classes on programming in Python, Linux, and the network commands such as ping and netstat, analyzing network packets, and network reconnaissance tools.

Applying What You've Learned to Your Home System(s)

This brings us to looking at practical things you can do to protect your own devices. You might think there's nothing that you can apply to your home devices since these attacks are aimed at the Internet infrastructure, but there are a couple of things you can do.

The first set of things you can do at home is prevent your devices from becoming zombies and being used in a DDoS attack. The specific actions you take to protect your devices are the same actions you're hopefully already doing such as:

1. Ensuring that your anti-virus and anti-malware programs are running and up to date.
2. Ensuring that all OS and application patches and updates have been installed.
3. Not downloading attachments in emails from unknown senders, or visiting sketchy web site links.

The second thing you can do at home is protect your router, if you have one, by ensuring that you've done the following:

4. Ensure that the firmware on your home router is up to date, if you have a router.

¹ <https://www.imperva.com/learn/ddos/booters-stressers-ddosers/>

² <https://www.softwaretestinghelp.com/ddos-attack-tools/>

5. Ensure that the default admin password for the router has been changed, and the new password is a strong password.
6. If you have a wireless router, make sure you are using a strong authentication protocol. WEP, which is the original protocol, is very easy to crack. It's been replaced by WPA-2, which is what you should select if given a choice. The issues with WEP have been known for years, so the chances of your router using it are very small, unless you have an older router. (Note – you'll learn all about this in the CSIA 330 Wireless Security class.)

Besides the basic actions, there are a couple of new things you can do:

7. Don't log in to any sensitive sites when you're using free wireless to prevent Man in the Middle attacks, especially if you do any mobile computing. For example, don't log in to your bank account when you're in Starbucks or travelling and waiting around in some airport. You never know who is really controlling the router you're using in your connection. I'm sure you've heard this advice before. And I know it can be easy advice to forget or ignore. In our always connected society it might feel really hard to be offline. And what could it hurt to connect to the "FreeStarbucks" wireless network?

Another thing you can do is to check the log on your home router, if you have one. This won't improve your security, but it will provide an eye-opening demonstration of how your home network defenses are almost always being probed. Your router log can show you all of the attempted connections coming from outside of your network, most of which will be an attacker probing your network. Usually you initiate every network connection that passes through your home router from inside your network. This happens when you do things like asking for a web page or open a connection to play an online game, at which time a network request will be initiated on your home computer, inside your network. The router makes the connection to the web server for you, and knows that the data being sent back is in response to your request. So, any requests from outside your network, that aren't initiated by a device inside your network, are an indication that someone is trying to gain access to your network. These attempts are typically blocked by the router to protect you, as the router also acts as a firewall. The attempts are also typically logged by your router.

Of course, it's possible that you have a device inside your network that people can access. Maybe you've set up your own VPN, or maybe you or your kids are running a Minecraft server. In this case, you must typically configure the firewall portion of the router to allow these requests into your network. You'll learn all about routers, firewall, and network traffic in a later class. But for now, I thought it would be interesting for you to check your router to see how many attacks are being performed against your home router. Every time I've checked my router the log is full of evidence of attempted network attacks, and I imagine yours is probably the same.

The specific steps for checking your router log will vary, depending on the make and model of router, but the general steps will be the same, and they include:

- A. Start your web browser.
- B. Log in to your router.
- C. Find the command in the router's Control Panel to display the logs

I've made a video that demonstrates the process, and has some guidance on figuring out what to do if the steps for your specific router are different.

Here is the link to the video, plus links to a couple other videos that show you the same thing, and links to some of the sites referenced in the video:

<https://tonysako.com/home/cs150-introduction-to-computer-security/intro-to-computer-security-network-threats-basic-home-router-security/> - Tips on checking the firmware on a home router, setting the admin password, and viewing router logs for evidence of attacks.

<https://tonysako.com/cs150-viewing-wireless-router-logs/> - An older version of the previous video. This video has tips on discovering the IP address for your home router, and recovering or resetting the admin password.

<https://www.techwalla.com/articles/how-to-check-a-routers-log> - Techwalla's video on checking router logs

<https://www.howtogeek.com/233952/how-to-find-your-routers-ip-address-on-any-computer-smartphone-or-tablet/>

<https://www.quora.com/How-can-you-know-router-admin-password-without-reset>

Ways to Check Your Comprehension

The test over this chapter won't be for a few weeks, so if you want check your comprehension you can use the review questions at the end of the chapter or the Practice Test. I don't have the review questions loaded in Canvas, so you'll have to just read them and figure out your answers on your own. Or you could try and connect with some other students in the class and drill each other using these questions.

If you want to use the Practice Test, look in the Test 2 Canvas Module for the link to a Practice Test. You can take the Practice Test as few or as many times as you want. You're not required to take the practice test, but it's also a good way to check your comprehension and prepare for the "real" test.

The Activities for This Section

There are three sets of activities for this section, a Hands-On assignment from the book, an additional Hands-On assignment that's not in the book, and a Case Project or writing

assignment. I've also created an optional activity that you might find of interest and want to do if you have a home router.

Before submitting your Hands-On Projects, add all of the information for your submission into a single document. Make sure that this document has the proper header information (your name, project number, date) and well as the project and step number for each item in the document. That is, if you are submitting a screen shot for Project 2-1, step 16, make sure and add some text that says "Project 2-1 #16", or something to that effect.

Hands-On Project 8-1

The homework for this section includes Hands-On Project 8-1 **DNS Poisoning**. As you learned, most Internet traffic is based on TCP/IP. With TCP/IP every computer is assigned an IP address, which is a 4 part number similar to 104.56.72.78. To send packets to another device you must know the recipient's IP address. It's almost like phones, where every phone needs a number, and you need to know the recipient's number to call them. And just like with phone numbers, it's much easier for humans to remember names than it is to remember numbers. If you're old enough you might remember the phone book, it provided this service for phone numbers. That is, it's easier to remember a person's name than their phone number. But to call someone you need to know their phone number. With the white pages phone book, or the old 411 service, you could "resolve" a person's name to their phone number. DNS does this exact same thing for IP addresses. The difference is that DNS is built into most network applications. All you need to do is type a URL or email address with an easy to remember DNS name, and the network will automatically query the DNS system and resolve this to an IP address. But imagine what would happen if an attacker could change the DNS system so that when you asked for IP address of a legitimate site, it returned the IP address of their fake site.

You'll learn the details about the DNS system in a later networking class. And you'll see that the DNS system itself is pretty trustworthy as it's relatively difficult for someone to hijack the entire system. But a key part of the DNS configuration is done on each computer. This is the part where you have a list of DNS servers that will be queried any time you need an IP address, or a small local "phonebook" that can be used like the contacts list on your cell phone. If an attacker can change either of these two files, they can fool your computer into using a fake IP addresses, and redirect some of your network traffic to computers they control.

In this exercise, you'll learn how to edit the file on your computer that acts like your phone's contact list. This will give you an idea of what big DNS does, that is, you give it a name and it returns an IP address. Plus, you'll get an idea of how it can be attacked. While this might seem like a stupid computer trick you can play on your friends remember that if you doing this can be a crime. That is, changing the local DNS file on your computer so that your friend is redirected to their CBC Canvas account when they try to access Facebook would be a funny trick, as long as you are there to laugh and then help them fix it. But changing the file on anyone else's computer, where you don't have prior written authorization, is a federal crime.

Note that you're not going to be able to complete the exercise as it's written in the book. Some websites check to ensure that the name you're using to connect to them is the correct name, and will return an error if you use the wrong DNS name and IP address combination. To help you out I've made a video that demonstrates the problems you might encounter and shows how to complete the project.

What to submit for Project 8-1:

Read the following notes carefully, as they explain exactly what you need to submit to receive credit.

Create screenshots after steps #12 and #13, and answer the question.

Tracert Hands-On Project

Note that this project is not in the book. In this project you'll get experience running the tracert command, which will demonstrate how many routers and networks handle network packets sent from your computer.

Run the tracert command and use it to display the routers and path between your computer and the website sire-usa.com. You can run tracert in a Command Prompt window, or use Open Visual Tracert. If you decide to use Open Visual Tracert I suggest you install and run it in your VM.

What to submit for this project

When the tracert command has finished, create a screenshot showing the list of routers that handled the network packets.

Required Case Project Homework (Writing Assignment)

The writing assignment for this section requires you to do some research and write a paper. You can select any of the subjects described in Case Project 8-1, 8-2, 8-3, 8-4, or 8-5. You only need to write one paper, but it must be on one of these subjects. All of the papers require you to do some research, so make sure and keep track of the papers or web sites you use for research, and include them as references in your paper.

Hopefully you remember how your paper must be formatted, and the other guidelines for writing papers. But if not, you can refer to the Written Project Guidelines document for details on how your paper/report will be graded. You can find the document at:

<https://tonysako.com/writingprojectguidelines2021/>

Also, remember to check your TurnItIn score. If the score is higher than 30% your submission will NOT be graded. You will need to either edit your material and put more of it in your own words, or add more original material. Once you have made your changes, you can resubmit

your work. There is no way to check your TurnItIn score before submitting your work. But don't worry about making multiple submissions, everyone does it and it has no impact on your grade.

<https://tonysako.com/home/cs150-introduction-to-computer-security/cs150-checking-your-turnitin-score/> - How to check your TurnItIn score