Introduction to intrusion tactics

- The case for securing networks
- Matt: A professional on dealing with attacks
- Denial of Service (DoS) attacks
- Read tcpdump logs
- Real-life DDoS attack
- Malicious packet sniffing
- IP Spoofing
- Overview of interception tactics
- Glossary terms from week 3 & wrap-up

The case for securing networks

Let's start by answering the question, why do we need secure networks?
As you've learned, networks are constantly at risk of attack from malicious hackers.
Attackers can infiltrate networks via malware, spoofing, or packet sniffing.
Network operations can also be disrupted by attacks such as packet flooding.
As we go along, you're going to learn about these and
other common network intrusion attacks in more detail.

Protecting a network from these types of attacks is important.
If even one of them happens, it could have a catastrophic impact on an organization.
Attacks can harm an organization by leaking valuable or
confidential information.
They can also be damaging to an organization's reputation and
impact customer retention.
Mitigating attacks may also cost the organization money and time.

Over the last few years,
there have been a number of examples of damage that cyber attacks can cause.
One notorious example was an attack against the American home-improvement
chain, Home Depot, in 2014.
A group of hackers compromised and infected Home Depot servers with malware.
By the time network administrators shut down the attack, the hackers had already
taken the credit and debit card information for over 56 million customers.

Now, you know why it's so important to secure a network.
But to keep a network secure,
you need to know what kinds of attacks to protect it from.
Coming up, you'll learn about some common network attacks.

How intrusions compromise your system

In this section of the course, you learned that every network has inherent vulnerabilities and could become the target of a network attack.

Attackers could have varying motivations for attacking your organization’s network. They may have financial, personal, or political motivations, or they may be a disgruntled employee or an activist who disagrees with the company's values and wants to harm an organization’s operations. Malicious actors can target any network. Security analysts must be constantly alert to potential

vulnerabilities in their organization’s network and take quick action to mitigate them.

In this reading, you’ll learn about network interception attacks and backdoor attacks, and the possible impacts these attacks could have on an organization.

Network interception attacks

Network interception attacks work by intercepting network traffic and stealing valuable information or interfering with the transmission in some way.

Malicious actors can use hardware or software tools to capture and inspect data in transit. This is referred to as **packet sniffing**. In addition to seeing information that they are not entitled to, malicious actors can also intercept network traffic and alter it. These attacks can cause damage to an organization’s network by inserting malicious code modifications or altering the message and interrupting network operations. For example, an attacker can intercept a bank transfer and change the account receiving the funds to one that the attacker controls.

Later in this course you will learn more about malicious packet sniffing, and other types of network interception attacks: on-path attacks and replay attacks.

Backdoor attacks

A **backdoor attack** is another type of attack you will need to be aware of as a security analyst. An organization may have a lot of security measures in place, including cameras, biometric scans and access codes to keep employees from entering and exiting without being seen. However, an employee might work around the security measures by finding a backdoor to the building that is not as heavily monitored, allowing them to sneak out for the afternoon without being seen.

In cybersecurity, backdoors are weaknesses intentionally left by programmers or system and network administrators that bypass normal access control mechanisms. Backdoors are intended to help programmers conduct troubleshooting or administrative tasks. However, backdoors can also be installed by attackers after they’ve compromised an organization to ensure they have persistent access.

Once the hacker has entered an insecure network through a backdoor, they can cause extensive damage: installing malware, performing a denial of service (DoS) attack, stealing private information or changing other security settings that leaves the system vulnerable to other attacks. A **DoS attack** is an attack that targets a network or server and floods it with network traffic.

Possible impacts on an organization

As you've learned already, network attacks can have a significant negative impact on an organization. Let's examine some potential consequences.

- **Financial**: When a system is taken offline with a DoS attack, or business operations are halted or slowed down by some other tactic, they prevent a company from performing the tasks that generate revenue. Depending on the size of an organization, interrupted operations can cost millions of dollars. In addition, if a malicious actor gets access to the personal information of the company's clients or customers, the company may face heavy litigation and settlement costs if customers seek legal recourse.
- **Reputation**: Attacks can also have a negative impact on the reputation of an organization. If it becomes public knowledge that a company has experienced a cyber attack, the public may become concerned about the security practices of the organization. They may stop trusting the company with their personal information and choose a competitor to fulfill their needs.
- **Public safety**: If an attack occurs on a government network, this can potentially impact the safety and welfare of the citizens of a country. In recent years, defense agencies across the globe are investing heavily in combating cyber warfare tactics. If a malicious actor gained access to a power grid, a public water system, or even a military defense communication system, the public could face physical harm due to a network intrusion attack.

Key takeaways

Malicious actors are constantly looking for ways to exploit systems. They learn about new vulnerabilities as they arise and attempt to exploit every vulnerability in a system. Attackers leverage backdoor attack methods and network interception attacks to gain sensitive information they can use to exploit an organization or cause serious damage. These types of attacks can impact an organization financially, damage its reputation, and potentially put the public in danger. It is important that security analysts stay educated in order to maintain network safety and reduce the likelihood and impact of these types of attacks. Securing networks has never been more important.

Matt: A professional on dealing with attacks

My name's Matt, I'm a chaos specialist at Google.
They let us choose our own job titles to best describe what it is we do.
I spend a lot of my time planning for
how to take care of anything that might possibly be going wrong, and
when it does happen, putting a team in place to fix it as quickly as possible.
I had no intention of being in technology at all.
In high school, I was a lifeguard, first at public pools and then at a state beach.
Lifeguarding got me into really enjoying rescue.
So I got an EMT license, went through firefighter school. About halfway through
my college process, and well into when I was being a firefighter on a daily basis.
I was dealing with some burnout, some stress.
I needed a change in my life.
And a friend of mine who I'd been online gaming with since the early days
of online gaming, when it was all text based,
he said, I can tell you're burning out hard and you need a change.
My friends and I are going to San Francisco to start a startup.
Will you come with us?
And I said, you realize I am not a computer guy, right?
And he said, no, you're a computer guy, you just won't admit it.
The same thing that has drawn me into incident response in tech is what
originally drew me to medical response.
I really love being there for people on their worst day.
Being there when people really need you and
they don't know where else to turn to has always just fed this part of me, and
I'm lucky to find that same joy in DFIR, Digital Forensics and Incident Response.
What type of attacks have we faced at Google?
That's a hard question to answer,
because we face all of the kinds of attacks that most other companies face.
People after ransomware, people after industrial secrets,
other countries looking for intelligence information.
There was a really interesting attack that occurred a little while ago.
They were interested in getting a lot of information from technical companies,
specifically about vulnerabilities in software.
And they put in place a long running campaign to build personalities on
social media as though they were legitimate security researchers, and
then reach out to other security researchers in our field,
build relationships, and then just at the right moment, sneak in some malware.
Being under attack by an adversary who's made some progress is incredibly
stressful.
The first things you're thinking and

feeling are a little bit of a sense of panic.
Oh no, this is going to be a bad day.
How long am I going to be awake working on this?
What have they done?
What am I going to do?
And for me, the mantra that I repeat to myself is, as an incident responder,
I am here to help.
The things that are most important to having a good outcome
in an incident are what we call the 3Cs: Command, Control and
Communications. Meaning someone needs to be in charge of it affirmatively leading.
Someone needs to be exerting control over everyone involved so
that everyone's aligned, focused on the mission, and the biggest and
most important one of them all: proper communications.
If you have something to offer to the incident,
don't just go do it, Communicate to someone.
I think I could do this to help us make progress.
I think if we look over here, we'll find more data.
The advice that I would give somebody who wants to get into cybersecurity is if you
want it, you probably belong here.
The more people we have in here, who are passionate, curious question askers,
who want to know more, who want to build better, and who care about
making every thing more secure for the people who have to use technology,
those are people we want in the industry and I would want you here.

Denial of Service (DoS) attacks

Welcome back. In this video,
we're going to discuss denial of service attacks.
A denial of service attack is
an attack that targets a network or server
and floods it with network traffic.
The objective of a denial of service attack,
or a DoS attack, is to disrupt
normal business operations by
overloading an organization's network.
The goal of the attack is to send
so much information to a network device
that it crashes or is
unable to respond to legitimate users.
This means that the organization won't
be able to conduct their normal business operations,
which can cost them money and time.
A network crash can also leave them
vulnerable to other security threats and attacks.

A distributed denial of service attack, or DDoS,
is a kind of DoS attack that uses
multiple devices or servers in
different locations to flood
the target network with unwanted traffic.
Use of numerous devices makes it more likely that
the total amount of traffic
sent will overwhelm the target server.
Remember, DoS stands for denial of service.
So it doesn't matter what part
of the network the attacker overloads;
if they overload anything, they win.
An unfortunate example I've
seen is an attacker who crafted
a very careful packet that caused
a router to spend extra time processing the request.
The overall traffic volume didn't overload the router;
the specifics within the packet did.

Now we'll discuss network level DoS attacks
that target network bandwidth to slow traffic.
Let's learn about three common network level DoS attacks.
The first is called a SYN flood attack.
A SYN flood attack is a type of

DoS attack that simulates the TCP connection
and floods the server with SYN packets.
Let's break this definition down a bit
more by taking a closer look at
the handshake process that is used to establish
a TCP connection between a device and a server.
The first step in the handshake is for the device to send
a SYN, or synchronize, request to the server.
Then, the server responds with
a SYN/ACK packet to
acknowledge the receipt of the device's request
and leaves a port open for
the final step of the handshake.
Once the server receives
the final ACK packet from the device,
a TCP connection is established.
Malicious actors can take advantage of the protocol by
flooding a server with
SYN packet requests for the first part of the handshake.
But if the number of SYN requests is
larger than the number of available ports on the server,
then the server will be overwhelmed
and become unable to function.

Let's discuss two other common DoS attacks
that use another protocol called ICMP.
ICMP stands for Internet Control Message Protocol.
ICMP is an internet protocol used by devices
to tell each other about
data transmission errors across the network.
Think of ICMP like a request
for a status update from a device.
The device will return
error messages if there is a network concern.
You can think of this like
the ICMP request checking in with
the device to make sure that all is well.
An ICMP flood attack
is a type of DoS attack performed by an attacker
repeatedly sending ICMP packets to a network server.
This forces the server to send an ICMP packet.
This eventually uses up all the bandwidth for incoming
and outgoing traffic and causes the server to crash.
Both of the attacks we've discussed so far,
SYN flood and ICMP flood,
take advantage of communication protocols
by sending an overwhelming number of requests.

There are also attacks that can overwhelm
the server with one big request.
One example that we'll discuss
is called the ping of death.

A ping of death attack is
a type of DoS attack that is caused when a hacker
pings a system by sending it
an oversized ICMP packet
that is bigger than 64 kilobytes,
the maximum size for a correctly formed ICMP packet.
Pinging a vulnerable network server with
an oversized ICMP packet
will overload the system and cause it to crash.
Think of this like dropping a rock on a small anthill.
Each individual ant can carry a certain amount of
weight while transporting food to and from the anthill.
But if a large rock is dropped on the anthill,
then many ants will be crushed, and the colony is unable to
function until it rebuilds its operations elsewhere.

Now that's it for DoS and DDoS attacks.
Coming up, we'll continue to
discuss common network attacks.

Read tcpdump logs

A **network protocol analyzer**, sometimes called a packet sniffer or a packet analyzer, is a tool designed to capture and analyze data traffic within a network. They are commonly used as investigative tools to monitor networks and identify suspicious activity. There are a wide variety of network protocol analyzers available, but some of the most common analyzers include:

- SolarWinds NetFlow Traffic Analyzer
- ManageEngine OpManager
- Azure Network Watcher
- Wireshark
- tcpdump

This reading will focus exclusively on tcpdump, though you can apply what you learn here to many of the other network protocol analyzers you'll use as a cybersecurity analyst to defend against any network intrusions. In an upcoming activity, you'll review a tcpdump data traffic log and identify a DoS attack to practice these skills.

tcpdump

tcpdump is a command-line network protocol analyzer. It is popular, lightweight–meaning it uses little memory and has a low CPU usage–and uses the open-source libpcap library. tcpdump is text based, meaning all commands in tcpdump are executed in the terminal. It can also be installed on other Unix-based operating systems, such as macOS®. It is preinstalled on many Linux distributions.

tcpdump provides a brief packet analysis and converts key information about network traffic into formats easily read by humans. It prints information about each packet directly into your terminal. tcpdump also displays the source IP address, destination IP addresses, and the port numbers being used in the communications.

Interpreting output

tcpdump prints the output of the command as the sniffed packets in the command line, and optionally to a log file, after a command is executed. The output of a packet capture contains many pieces of important information about the network traffic.

types of information presented in a tcpdump packet capture.

Some information you receive from a packet capture includes:

- **Timestamp**: The output begins with the timestamp, formatted as hours, minutes, seconds, and fractions of a second.
- **Source IP**: The packet’s origin is provided by its source IP address.
- **Source port**: This port number is where the packet originated.
- **Destination IP**: The destination IP address is where the packet is being transmitted to.
- **Destination port**: This port number is where the packet is being transmitted to.

Note: By default, tcpdump will attempt to resolve host addresses to hostnames. It'll also replace port numbers with commonly associated services that use these ports.

Common uses

tcpdump and other network protocol analyzers are commonly used to capture and view network communications and to collect statistics about the network, such as troubleshooting network performance issues. They can also be used to:

- Establish a baseline for network traffic patterns and network utilization metrics.
- Detect and identify malicious traffic
- Create customized alerts to send the right notifications when network issues or security threats arise.
- Locate unauthorized instant messaging (IM), traffic, or wireless access points.

However, attackers can also use network protocol analyzers maliciously to gain information about a specific network. For example, attackers can capture data packets that contain sensitive information, such as account usernames and passwords. As a cybersecurity analyst, It’s important to understand the purpose and uses of network protocol analyzers.

Key takeaways

Network protocol analyzers, like tcpdump, are common tools that can be used to monitor network traffic patterns and investigate suspicious activity. tcpdump is a command-line network protocol analyzer that is compatible with Linux/Unix and macOS®. When you run a tcpdump command, the tool will output packet routing information, like the timestamp, source IP address and port number, and the destination IP address and port number. Unfortunately, attackers can also use network protocol analyzers to capture data packets that contain sensitive information, such as account usernames and passwords.

Real-life DDoS attack

Previously, you were introduced to Denial of Service (DoS) attacks. You also learned that volumetric distributed DoS (DDoS) attacks overwhelm a network by sending unwanted data packets in such large quantities that the servers become unable to service normal users. This can be detrimental to an organization. When systems fail, organizations cannot meet their customers' needs. They often lose money, and in some cases, incur other losses. An organization’s reputation may also suffer if news of a successful DDoS attack reaches consumers, who then question the security of the organization.

In this reading you’ll learn about a 2016 DDoS attack against DNS servers that caused major outages at multiple organizations that have millions of daily users.

A DDoS targeting a widely used DNS server

In previous videos, you learned about the function of a DNS server. As a review, DNS servers translate website domain names into the IP address of the system that contains the information for the website. For instance, if a user were to type in a website URL, a DNS server would translate that into a numeric IP address that directs network traffic to the location of the website’s server.

On the day of the DDoS attack we are studying, many large companies were using a DNS service provider. The service provider was hosting the DNS system for these companies. This meant that when internet users typed in the URL of the website they wanted to access, their devices would be directed to the right place. On October 21, 2016, the service provider was the victim of a DDoS attack.

Leading up to the attack

Before the attack on the service provider, a group of university students created a botnet. A **botnet** is a collection of computers infected by malware that are under the control of a single threat actor, known as the “bot-herder." Each computer in the botnet can be remotely controlled to send a data packet to a target system. In a botnet attack, cyber criminals instruct all the bots on the botnet to send data packets to the target system at the same time, resulting in a DDoS attack.

The group of university students posted the code for the botnet online so that it would be accessible to thousands of internet users and authorities wouldn’t be able to trace the botnet back to the students. In doing so, they made it possible for other malicious actors to learn the code to the botnet and control it remotely. This included the cyber criminals who attacked the DNS service

provider.

The day of attack

At 7:00 a.m. on the day of the attack, the botnet sent tens of millions of DNS requests to the service provider. This overwhelmed the system and the DNS service shut down. This meant that all of the websites that used the service provider could not be reached. When users tried to access various websites that used the service provider, they were not directed to the website they typed in their browser. Outages for each web service occurred all over North America and Europe.

The service provider’s systems were restored after only two hours of downtime. Although the cyber criminals sent subsequent waves of botnet attacks, the DNS company was prepared and able to mitigate the impact.

Key takeaways

As demonstrated in the above example, DDoS attacks can be very damaging to an organization. As a security analyst, it’s important to acknowledge the seriousness of such an attack so that you’re aware of opportunities to protect the network from them. If your network has important operations distributed across hosts that can be dynamically scaled, then operations can continue if the baseline host infrastructure goes offline. DDoS attacks are damaging, but there are concrete actions that security analysts can take to help protect their organizations. Keep going through this course and you will learn about common mitigation strategies to protect against DDoS attacks.

Malicious packet sniffing

In this video, we'll discuss
packet sniffing, with a focus on how
threat actors may use this technique to
gain unauthorized access to information.
Previously, you learned about
the information and data packets
that travel across the network.
Packets include a header which contains
the sender's and receiver's IP addresses.
Packets also contain a body, which
may contain valuable information like names,
date of birth, personal messages,
financial information, and credit card numbers.

Packet sniffing is the practice of using
software tools to observe
data as it moves across a network.
As a security analyst,
you may use packet sniffing to analyze and capture
packets when investigating ongoing incidents
or debugging network issues.
Later in this certificate program,
you'll gain hands-on practice
with some packet sniffing software.
However, malicious actors may also use
packet sniffing to look at
data that has not been sent to them.
This is a little bit like opening somebody else's mail.
It's important for you to learn about how
threat actors use packet sniffing
with harmful intent so you can be
prepared to protect against these malicious acts.
Malicious actors may insert themselves in the middle of
an authorized connection between two devices.
Then they can use packet sniffing to spy on
every data packet as it comes across their device.
The goal is to find valuable information in
the data packets that they can
then use to their advantage.
Attackers can use software applications
or a hardware device to look into data packets.
Malicious actors can access a network packet with
a packet sniffer and make changes to the data.

They may change the information
in the body of the packet,
like altering a recipient's bank account number.

Packet sniffing can be passive or active.
Passive packet sniffing is a type of
attack where data packets are read in transit.
Since all the traffic on a network
is visible to any host on the hub,
malicious actors can view
all the information going
in and out of the device they are targeting.
Thinking back to the example of a letter being delivered,
we can compare a passive packet sniffing
attack to a postal delivery person
maliciously reading somebody's mail.
The postal worker, or packet
sniffer, has the right to deliver the mail,
but not the right to read the information inside.
Active packet sniffing is a type of
attack where data packets are manipulated in transit.
This may include injecting
internet protocols to redirect the packets to
an unintended port or
changing the information the packet contains.
Active packet sniffing attack would
be like a neighbor telling the delivery person
"I'll deliver that mail for you," and then reading the mail
or changing the letter before putting it in your mailbox.
Even though your neighbor knows you
and even if they deliver it to the correct house,
they are actively going out of
their way to engage in malicious behavior.

The good news is that
malicious packet sniffing can be prevented.
Let's look at a few ways
the network security professional
can prevent these attacks.
One way to protect against
malicious packet sniffing is to use
a VPN to encrypt and protect
data as it travels across the network.
If you don't remember how VPNs work,
you can revisit the video about
this topic in the previous section of the program.
When you use a VPN,

hackers might interfere with your traffic,
but they won't be able to decode it
to read it and read your private information.
Another way to add a layer of
protection against packet sniffing is to make sure
that websites you have use
HTTPS at the beginning of the domain address.
Previously, we discussed how HTTPS uses SSL/TLS to
encrypt data and prevent eavesdropping
when malicious actors spy on network transmissions.
One final way to help protect yourself against
malicious packet sniffing is to
avoid using unprotected WiFi.
You usually find unprotected WiFi in
public places like coffee shops,
restaurants, or airports.
These networks don't use encryption.
This means that anyone on the network can access
all of the data traveling to and from your device.
One precaution you can take is avoiding
free public WiFi unless you have
a VPN service already installed on your device.

Now you know how threat actors may use
packet sniffing and how to
protect a network from these attacks.
Let's move on to discuss other network intrusions.

IP Spoofing

Next, let's learn about
another kind of network attack called IP spoofing.
IP spoofing is a network attack
performed when an attacker changes the source IP of
a data packet to impersonate
an authorized system and gain access to a network.
In this kind of attack,
the hacker is pretending to be someone they are
not so they can communicate over
the network with the target computer and get
past firewall rules that may prevent outside traffic.
Some common IP spoofing attacks are on-path attacks,
replay attacks, and smurf attacks.
Let's discuss these one at a time.

An on-path attack is
an attack where the malicious actor places themselves in
the middle of an authorized connection
and intercepts or alters the data in transit.
On-path attackers gain access to
the network and put themselves between two devices,
like a web browser and a web server.
Then they sniff the packet
information to learn the IP and
MAC addresses to devices
that are communicating with each other.
After they have this information,
they can pretend to be either of these devices.

Another type of attack is a replay attack.
A replay attack is
a network attack performed when
a malicious actor intercepts
a data packet in transit and delays
it or repeats it at another time.
A delayed packet can cause
connection issues between target computers,
or a malicious actor may take
a network transmission that was sent by
an authorized user and repeat it at
a later time to impersonate the authorized user.

A smurf attack is a combination of

a DDoS attack and an IP spoofing attack.
The attacker sniffs an authorized user's IP address
and floods it with packets.
This overwhelms the target computer and can
bring down a server or the entire network.

Now that you've learned about
different kinds of IP spoofing,
let's talk about how you can protect
the network from this kind of attack.
As you previously learned,
encryption should always be implemented so that the data
in your network transfers can't
be read by malicious actors.
Firewalls can be configured
to protect against IP spoofing.
IP spoofing makes it seem like
the malicious actor is an authorized user
by changing the sender's address of
the data packet to match the target network's address.
So if a firewall receives a data packet from the internet
where the sender's IP address
is the same as the private network,
then the firewall will deny the transmission
since all the devices with that IP address
should already be on the local network.
You can make sure that your firewalls
configure correctly by creating a rule to
reject all incoming traffic that has
the same IP address as the local network.

That's it for IP spoofing.
You've learned how IP spoofing is used in
some common attacks like on-path attacks,
replay attacks, and smurf attacks.

Overview of interception tactics

In the previous course items, you learned how packet sniffing and IP spoofing are used in network attacks. Because these attacks intercept data packets as they travel across the network, they are called interception attacks.

This reading will introduce you to some specific attacks that use packet sniffing and IP spoofing. You will learn how hackers use these tactics and how security analysts can counter the threat of interception attacks.

A closer review of packet sniffing

As you learned in a previous video, **packet sniffing** is the practice of capturing and inspecting data packets across a network. On a private network, data packets are directed to the matching destination device on the network.

The device's **Network Interface Card (NIC)** is a piece of hardware that connects the device to a network. The NIC reads the data transmission, and if it contains the device's MAC address, it accepts the packet and sends it to the device to process the information based on the protocol. This occurs in all standard network operations. However, a NIC can be set to promiscuous mode, which means that it accepts all traffic on the network, even the packets that aren't addressed to the NIC's device. You'll learn more about NIC's later in the program. Malicious actors might use software like Wireshark to capture the data on a private network and store it for later use. They can then use the personal information to their own advantage. Alternatively, they might use the IP and MAC addresses of authorized users of the private network to perform IP spoofing.

A closer review of IP spoofing

After a malicious actor has sniffed packets on the network, they can impersonate the IP and MAC addresses of authorized devices to perform an IP spoofing attack. Firewalls can prevent IP spoofing attacks by configuring it to refuse unauthorized IP packets and suspicious traffic. Next, you'll examine a few common IP spoofing attacks that are important to be familiar with as a security analyst.

On-path attack

An **on-path attack** happens when a hacker intercepts the communication between two devices or servers that have a trusted relationship. The transmission between these two trusted network

devices could contain valuable information like usernames and passwords that the malicious actor can collect. An on-path attack is sometimes referred to as a **meddler-in-the middle attack** because the hacker is hiding in the middle of communications between two trusted parties.

Or, it could be that the intercepted transmission contains a DNS system look-up. You'll recall from an earlier video that a DNS server translates website domain names into IP addresses. If a malicious actor intercepts a transmission containing a DNS lookup, they could spoof the DNS response from the server and redirect a domain name to a different IP address, perhaps one that contains malicious code or other threats. The most important way to protect against an on-path attack is to encrypt your data in transit, e.g. using TLS.

Smurf attack

A **smurf attack** is a network attack that is performed when an attacker sniffs an authorized user's IP address and floods it with packets. Once the spoofed packet reaches the broadcast address, it is sent to all of the devices and servers on the network.

In a smurf attack, IP spoofing is combined with another denial of service (DoS) technique to flood the network with unwanted traffic. For example, the spoofed packet could include an Internet Control Message Protocol (ICMP) ping. As you learned earlier, ICMP is used to troubleshoot a network. But if too many ICMP messages are transmitted, the ICMP echo responses overwhelm the servers on the network and they shut down. This creates a denial of service and can bring an organization's operations to a halt.

An important way to protect against a smurf attack is to use an advanced firewall that can monitor any unusual traffic on the network. Most next generation firewalls (NGFW) include features that detect network anomalies to ensure that oversized broadcasts are detected before they have a chance to bring down the network.

DoS attack

As you've learned, once the malicious actor has sniffed the network traffic, they can impersonate an authorized user. A **Denial of Service attack** is a class of attacks where the attacker prevents the compromised system from performing legitimate activity or responding to legitimate traffic. Unlike IP spoofing, however, the attacker will not receive a response from the targeted host. Everything about the data packet is authorized including the IP address in the header of the packet. In IP spoofing attacks, the malicious actor uses IP packets containing fake IP addresses. The attackers keep sending IP packets containing fake IP addresses until the network server crashes.

Pro Tip: Remember the principle of defense-in-depth. There isn't one perfect strategy for stopping each kind of attack. You can layer your defense by using multiple strategies. In this case, using industry standard encryption will strengthen your security and help you defend from DoS attacks on more than one level.

Key takeaways

This reading covered several types of common IP spoofing attacks. You learned about how packet sniffing is performed and how gathering information from intercepting data transmissions can give malicious actors opportunities for IP spoofing. Whether it is an on-path attack, IP spoofing attack, or a smurf attack, analysts need to ensure that mitigation strategies are in place to limit the threat and prevent security breaches.

Glossary terms from week 3 & wrap-up

Nice job finishing this section!
Let's review what you've learned so far.
We discussed how to secure networks.
We also learned about network intrusion tactics
like malicious packet sniffing and IP spoofing.
Finally, we discussed how
a security analyst can protect against these attacks.
You've learned about DoS and DDoS
attacks like ICMP flooding, SYN attacks,
and the ping of death,
which try to overwhelming a network by
flooding it with unwanted data packets.

Now, just think about everything you
know already about network attacks.
What you've learned in these videos will be
essential in your work as a security analyst.
Coming up, you'll learn
about how security analysts can protect
the network using various security hardening techniques.

Terms and definitions from Course 3, Week 3

Active packet sniffing: A type of attack where data packets are manipulated in transit

Botnet: A collection of computers infected by malware that are under the control of a single threat actor, known as the “bot-herder"

Denial of service (DoS) attack: An attack that targets a network or server and floods it with network traffic

Distributed denial of service (DDoS) attack: A type of denial or service attack that uses multiple devices or servers located in different locations to flood the target network with unwanted traffic

Internet Control Message Protocol (ICMP): An internet protocol used by devices to tell each other about data transmission errors across the network

Internet Control Message Protocol (ICMP) flood: A type of DoS attack performed by an attacker repeatedly sending ICMP request packets to a network server

IP spoofing: A network attack performed when an attacker changes the source IP of a data packet to impersonate an authorized system and gain access to a network

Network Interface Card (NIC): Hardware that connects computers to a network

On-path attack: An attack where a malicious actor places themselves in the middle of an authorized connection and intercepts or alters the data in transit

Packet sniffing: The practice of capturing and inspecting data packets across a network

Passive packet sniffing: A type of attack where a malicious actor connects to a network hub and looks at all traffic on the network

Ping of death: A type of DoS attack caused when a hacker pings a system by sending it an oversized ICMP packet that is bigger than 64KB

Replay attack: A network attack performed when a malicious actor intercepts a data packet in transit and delays it or repeats it at another time

Smurf attack: A network attack performed when an attacker sniffs an authorized user’s IP address and floods it with ICMP packets

Synchronize (SYN) flood attack: A type of DoS attack that simulates a TCP/IP connection and floods a server with SYN packets