WEEK 2 · SESSION 1 OF 2

Cybercrime, Access & Legal Boundaries

Where is the line between legitimate cybersecurity work and unauthorized activity? Today we find out — through scope, the CFAA, vulnerability research, and two cases making headlines right now.

LAST WEEK → THIS WEEK

From “Should you?” to “Are you authorized to?”

THIS WEEK'S BIG QUESTION

“Where is the line between legitimate cybersecurity work and unauthorized activity?”

The same tools used by penetration testers, analysts, and researchers can just as easily be used for unauthorized access. The difference is rarely the tool.

SESSION 1 FOCUS

What You'll Be Able to Do By the End of Today

Explain why authorization is fundamental

The foundation for every other objective this week.

Explain foundational CFAA concepts

Access without authorization vs. exceeding authorized access.

Distinguish authorized testing from unauthorized activity

Using real scope documents and real discoveries.

Evaluate the legal and professional boundaries of vulnerability research

Using a live 2026 case as the test.

Recognize the importance of written scope

And know what a scope document should actually contain.

Evaluate AI-assisted legal research

As a tool to identify questions — never as the final authority.

Identify legal issues in unauthorized access

Grounded in the Computer Fraud and Abuse Act.

Determine when to stop, escalate, or seek more authority

The single most important professional habit this week builds.

SECTION 1

Authorization Comes Before Testing

The tool doesn't know whether you have permission — you have to know

AUTHORIZATION

The Technology Is Identical. The Authorization Is Not.

A vulnerability scanner can be used during:

- An authorized penetration test
- A cybersecurity lab exercise
- A CTF competition
- A vulnerability disclosure program
- A bug bounty program
- **An unauthorized scan of someone else's systems**

The technology may be identical. The authorization is not.

AUTHORIZATION

Before Testing, Know Exactly What You Have Permission to Do

- Which systems are in scope
- Which systems are excluded
- When testing may occur
- Which techniques are permitted
- Which techniques are prohibited
- How discovered data should be handled
- How findings should be reported
- When testing must stop
- Who to contact if something unexpected happens

THE PROFESSIONAL DEFAULT

When scope is unclear, the answer is not to guess.

Clarify the authorization before proceeding.

SECTION 2

Cybercrime and the Computer Fraud and Abuse Act

The federal law that defines unauthorized computer access — and its real limits

THE CFAA

Two Core Concepts to Recognize

Access Without Authorization

Never had permission to access the system at all — no account, no credentials, no invitation, under any circumstances.

Exceeding Authorized Access

Had some legitimate access, but used it to obtain or alter information off-limits to that access — going beyond what was granted.

Also central to CFAA analysis: knowledge and intent — did the person know their access was unauthorized?

Avoid Oversimplifying: Not Every ToS Violation Is a Federal Crime

DOJ charging policy is explicit that these are NOT sufficient, by themselves, to warrant federal charges:

- Embellishing an online dating profile against a site's terms of service
- Creating a fictional account on a hiring, housing, or rental website
- Using a pseudonym on a platform that prohibits them
- Checking sports scores or paying bills at work

Read the exact language of legal material — don't collapse nuanced legal concepts into oversimplified rules.

2026 CASE IN THE NEWS

Clear-Cut “Without Authorization”: The Snowflake Breach

Connor Riley Moucka, 26 (Kitchener, Ontario) — pleaded guilty August 5, 2026

Stolen credentials, not a technical exploit

Moucka and co-conspirators used stolen login credentials — no clever hack, just valid-looking access using access nobody authorized.

165+ victim organizations, billions of records

Compromised cloud-hosted data belonging to at least 165 customers of a single SaaS provider, including AT&T, LendingTree, and Ticketmaster.

Extortion, not just access

Stolen data included banking information, driver's license and passport numbers, and Social Security numbers — used to extort victims for over $2.5 million.

REQUIRED RESOURCE

DOJ Justice Manual — CFAA Guidance (9-48.000)

justice.gov/jm/jm-9-48000-computer-fraud

Focus your reading on: access without authorization — and — exceeding authorized access. Come to Class Meeting 2 ready with one distinction, one surprise, one implication for practitioners, and one open question.

SECTION 3

Security Research & Vulnerability Disclosure

Discovering a flaw is not the same as being authorized to act on it

What Are You Authorized to Do With What You Find?

Discovering a vulnerability does not automatically authorize you to exploit it further.

Professional security researchers need to understand:

Vulnerability disclosure policies	Authorized testing scope	Prohibited activities	Privacy requirements
Data-handling expectations	Reporting procedures	Disclosure expectations	Safe-harbor language

RESOURCE

CISA Vulnerability Disclosure Policy Template

cisa.gov/vulnerability-disclosure-policy-template

A model for how organizations set clear expectations and authorization boundaries for outside researchers — exactly the kind of document Microsoft and Nightmare Eclipse disagreed about the meaning of.

2026 CASE IN THE NEWS

When Disclosure Breaks Down: Microsoft vs. “Nightmare Eclipse”

An independent researcher publicly disclosed unpatched Microsoft bugs — with working exploit code — in May 2026

Four named vulnerabilities

BlueHammer, RedSun, UnDefend, and YellowKey — affecting Windows Defender and BitLocker, published on GitHub and GitLab with exploit code.

The researcher says Microsoft cut them off first

Nightmare Eclipse claims MSRC revoked their vulnerability-reporting account access, leaving them feeling they had no private channel left.

Some bugs were later used in real attacks

Microsoft and CISA both confirmed some of the disclosed flaws were subsequently exploited by real-world attackers before a patch existed.

2026 CASE IN THE NEWS

Two Arguments, Both With Real Weight

Microsoft's Position

The researcher never attempted to report the bugs so Microsoft could fix them first — that would have been “responsible.” Publishing exploit code before a patch existed may have directly aided the attackers who went on to use it. Microsoft's Digital Crimes Unit says it will pursue cases against “these actors and those that enable their criminal activity.”

The Security Community's Response

Veteran researcher Katie Moussouris (who pioneered Microsoft's own bug bounty program) called the threat of prosecution “over the top” and warned it will create a chilling effect — fewer researchers reporting bugs at all. Others noted disclosure norms exist partly to protect the vendor, not just the customer, and that a company revoking a researcher's reporting access first complicates the story.

CLASS DISCUSSION

Applying This Week's Concepts to a Live Case

Does this case involve unauthorized access at all — or is it a dispute about disclosure conduct, a different question entirely?

If a company revokes a researcher's private reporting channel, does that change the researcher's professional obligations?

Does “we warned them and got no response” excuse publishing working exploit code before a patch exists?

What would a well-written vulnerability disclosure policy have needed to prevent this dispute from happening at all?

SECTION 4

Activity 1: Where Should Research Stop?

A cybersecurity student finds a vulnerable system — and has to decide what happens next

ACTIVITY 1

The Scenario

SCENARIO

A cybersecurity student notices that a local bakery's customer WiFi portal has an exposed admin login page. Out of curiosity, they try a few common default credentials — and one of them works. Now logged into the admin panel, they can see customer names, order history, and stored payment metadata. They never contacted the bakery owner first. They take a screenshot “just in case” and start drafting an email to let the owner know.

The question is not simply whether the student meant well.

Where should professional security research stop when authorization does not exist?

ACTIVITY 1

Discussion Structure

Pinpoint the line

As a group, identify the exact moment the student's actions crossed from noticing a problem into unauthorized access.

Apply the CFAA concepts

Was this 'access without authorization,' or could it be read differently? What about knowledge and intent?

Decide what happens next

The student is already logged in and has a screenshot. What should they do now — and does that choice matter separately from what they already did?

SECTION 5

Activity 2: The Scope Challenge

A real-looking pentest authorization — and discoveries that don't fit neatly inside it

ACTIVITY 2

Your Team's Written Authorization

PENETRATION TEST AUTHORIZATION — CLIENT: RIDGEVIEW RETAIL GROUP

Authorized network range	203.0.113.0/24 (external) and 10.20.0.0/16 (internal)
Testing window	Monday–Friday, 9:00 PM–5:00 AM local time only
Prohibited techniques	No denial-of-service testing • No social engineering of employees • No physical access attempts
Excluded production systems	Payment processing system (PCI scope) • Production customer database server • Primary email server
Reporting	All findings reported to the named client contact within 24 hours of discovery

ACTIVITY 2

During the Assessment, Your Team Discovers…

1. A staging server responding on the same authorized 203.0.113.0/24 range
2. A third-party AWS S3 storage bucket referenced in an API response, hosted outside the client's own infrastructure
3. Hardcoded database credentials inside a config file on an in-scope web server
4. A forgotten legacy subdomain hosting an old customer portal — not in the authorized range and not mentioned anywhere in the paperwork
5. An employee's personal social media profile, found while researching the organization online

For each discovery: clearly in scope — clearly out of scope — or unclear and requiring additional authorization?

ACTIVITY 2

Debrief — Reasoning Through Each Discovery

1. Staging server	**In scope**	Same authorized IP range — covered by the written authorization as written.
2. Third-party S3 bucket	**Unclear**	Hosted outside the client's own infrastructure — the client cannot authorize testing on a third party's system. Needs separate authorization or must be excluded and reported only.
3. Hardcoded credentials found	**Nuanced**	Discovering them on an in-scope system is fine to report. Using them to log into the excluded production database is not — discovery and use are two different authorization questions.
4. Legacy subdomain	**Unclear**	Not in the authorized range, not mentioned in the paperwork. Stop and clarify before touching it further — do not assume it's fair game just because it's related.
5. Employee's personal profile	**Out of scope**	Not a technical system, and researching individuals crosses into the prohibited social-engineering territory named in the authorization.

SECTION 6

AI-Assisted Legal Research

A confident AI answer is not the same thing as an accurate legal claim

Useful Assistant, Not a Legal Authority

Where AI Helps

- Explaining unfamiliar terminology
- Identifying questions that need further research
- Summarizing a difficult passage — after you've read it yourself
- Identifying competing interpretations
- Challenging your understanding
- Suggesting authoritative sources to investigate

Where AI Falls Short

- AI is not the legal authority
- A confident answer does not make a legal claim accurate
- AI can miss context, invent details, or misstate how a rule actually applies
- AI cannot substitute for reading the primary source yourself

The Workflow You'll Practice Today

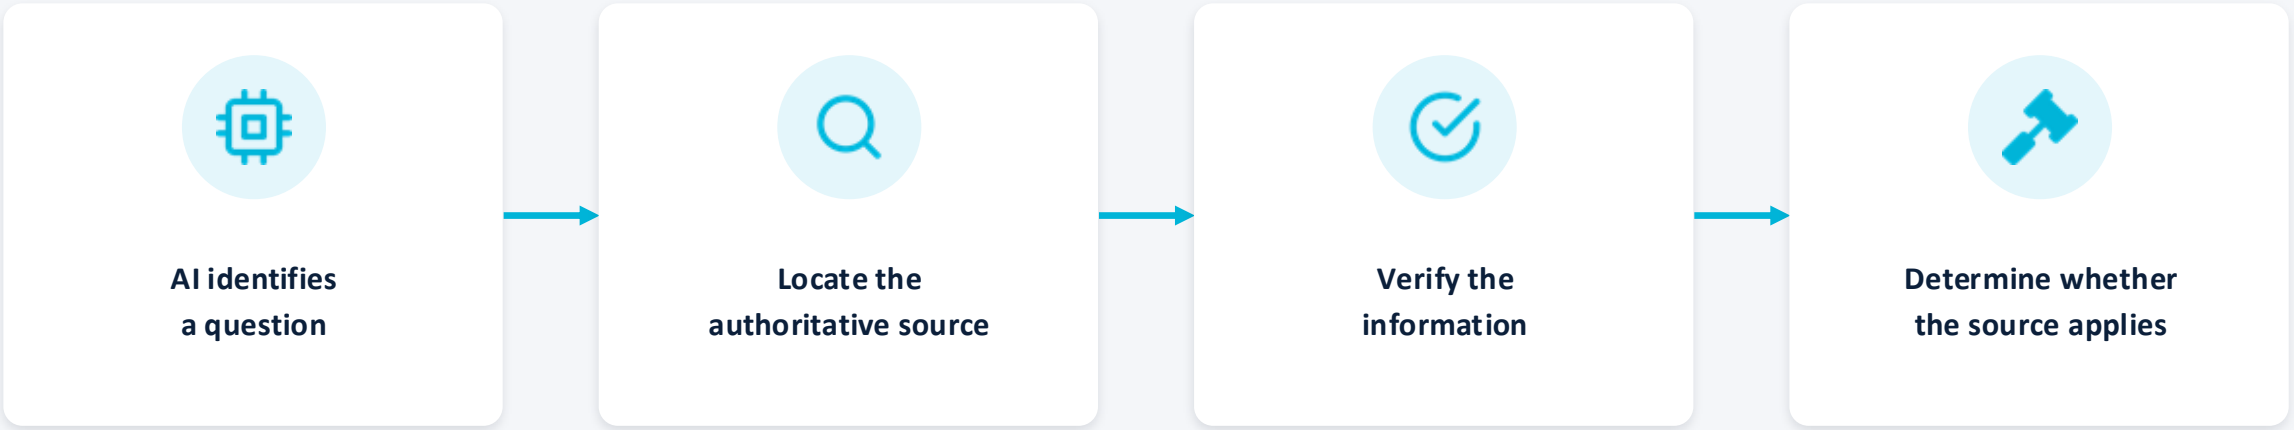

You may use ChatGPT Edu to clarify difficult language in the DOJ reading — but always compare any AI explanation against the original source.

Practice: Try the Workflow Yourself

In pairs, pick one question raised during today's Microsoft/Nightmare Eclipse discussion. Use ChatGPT Edu to help identify what to research — then locate and verify an actual authoritative source before drawing any conclusion.

- If the AI gives you a confident-sounding legal answer, that confidence is not evidence it's correct
- A source is only authoritative if you can identify who published it and why they're positioned to know
- Be ready to say what you verified and what you couldn't confirm — an honest gap beats a confident guess

SECTION 7

Between Class Meetings

What to complete before Class Meeting 2

BEFORE CLASS MEETING 2

DOJ Authorization Research — Required Prep

Review the DOJ CFAA guidance: justice.gov/jm/jm-9-48000-computer-fraud — focus on access without authorization and exceeding authorized access

Come to Class Meeting 2 prepared to discuss:

1. One important distinction between the two concepts
2. One part of the guidance you found surprising or particularly important
3. One implication for penetration testers or security researchers
4. One question you still have

Read the actual DOJ material first. Use ChatGPT Edu afterward only to clarify difficult language — and compare any AI explanation with the original source.

BEFORE YOU LEAVE TODAY

You Should Be Able to Explain Why These Statements Matter

- Good intentions do not create authorization.
- Finding a vulnerability does not automatically authorize you to exploit it.
- When scope is unclear, stop and clarify.
- Technical access does not automatically establish legal authority.
- AI can help identify legal questions, but authoritative sources must be verified.

CLASS MEETING 2

Digital Evidence, Investigation & Legal Decision-Making

Bring your four DOJ discussion items. We'll move from authorization questions into how evidence is identified, collected, and analyzed — and where AI-generated conclusions can go wrong.