WEEK 2 · SESSION 2 OF 2

Digital Evidence, Investigation & Legal Decision-Making

A security tool alert is not proof. An AI summary is not proof. Today we learn how evidence actually becomes a defensible conclusion — and where 1,500+ real 2026 court sanctions show what happens when nobody checks.

SESSION 1 → SESSION 2

From “Were you authorized?” to “Can you prove what happened?”

TODAY'S THROUGHLINE

Source evidence → analysis → conclusion

Every activity today comes back to this chain — and to what happens professionally when someone skips a link in it.

SESSION 2 FOCUS

Skills We'll Discuss Today:

Identify and evaluate relevant digital evidence

Across the many sources an incident can produce.

Explain why evidence integrity and documentation matter

Chain of custody isn't paperwork for its own sake.

Distinguish underlying evidence from analytical conclusions

The single most tested distinction today.

Evaluate AI-assisted legal research and digital investigation

Including a real, current, very expensive cautionary trend.

Develop and defend a professionally responsible course of action

In both the mock exercise and the mock hearing.

Determine when to stop, escalate, or seek more authority

Applied to a live insider-incident scenario.

SECTION 1

Why Digital Evidence Matters

Cybersecurity incidents create evidence — what an organization does with it matters

WHY EVIDENCE MATTERS

What Good Evidence Handling Makes Possible

- Understand what happened
- Determine the scope of an incident
- Contain an attack
- Recover systems
- Identify responsible accounts or devices
- Support disciplinary action
- Support insurance or regulatory requirements
- Support a legal investigation

TODAY'S GOVERNING PRINCIPLE

How evidence is identified, collected, preserved, analyzed, and documented matters.

SECTION 2

Sources of Digital Evidence

Depending on the incident, evidence can come from more places than you'd expect

SOURCES OF EVIDENCE

Where Evidence Can Come From

Authentication logs

Endpoint logs

Firewall logs

Cloud audit records

Email

Browser artifacts

File metadata

Network traffic

Memory

Mobile-device data

Security alerts

Photos / screenshots

Forensic images

SOURCES OF EVIDENCE

Not Everything Produced Is Underlying Evidence

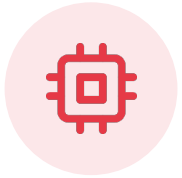

Interpretations

An analyst's summary is an interpretation.

An AI-generated incident report is an interpretation.

Both are useful — but both are a step removed from the underlying facts.

Underlying Evidence

The original logs, files, and records remain critical.

This is what a conclusion must ultimately trace back to — not a summary of a summary.

SECTION 3

Digital Forensics Foundations

Selected concepts from NIST SP 800-86 — not a forensics certification in one session

RESOURCE

NIST SP 800-86 — Guide to Integrating Forensic Techniques into Incident Response

csrc.nist.gov/pubs/sp/800/86/final

You will use selected concepts from this guide — not the entire document. The goal is understanding why professional investigations require systematic evidence handling, not becoming a certified forensic examiner in one session.

Four Broad Activities

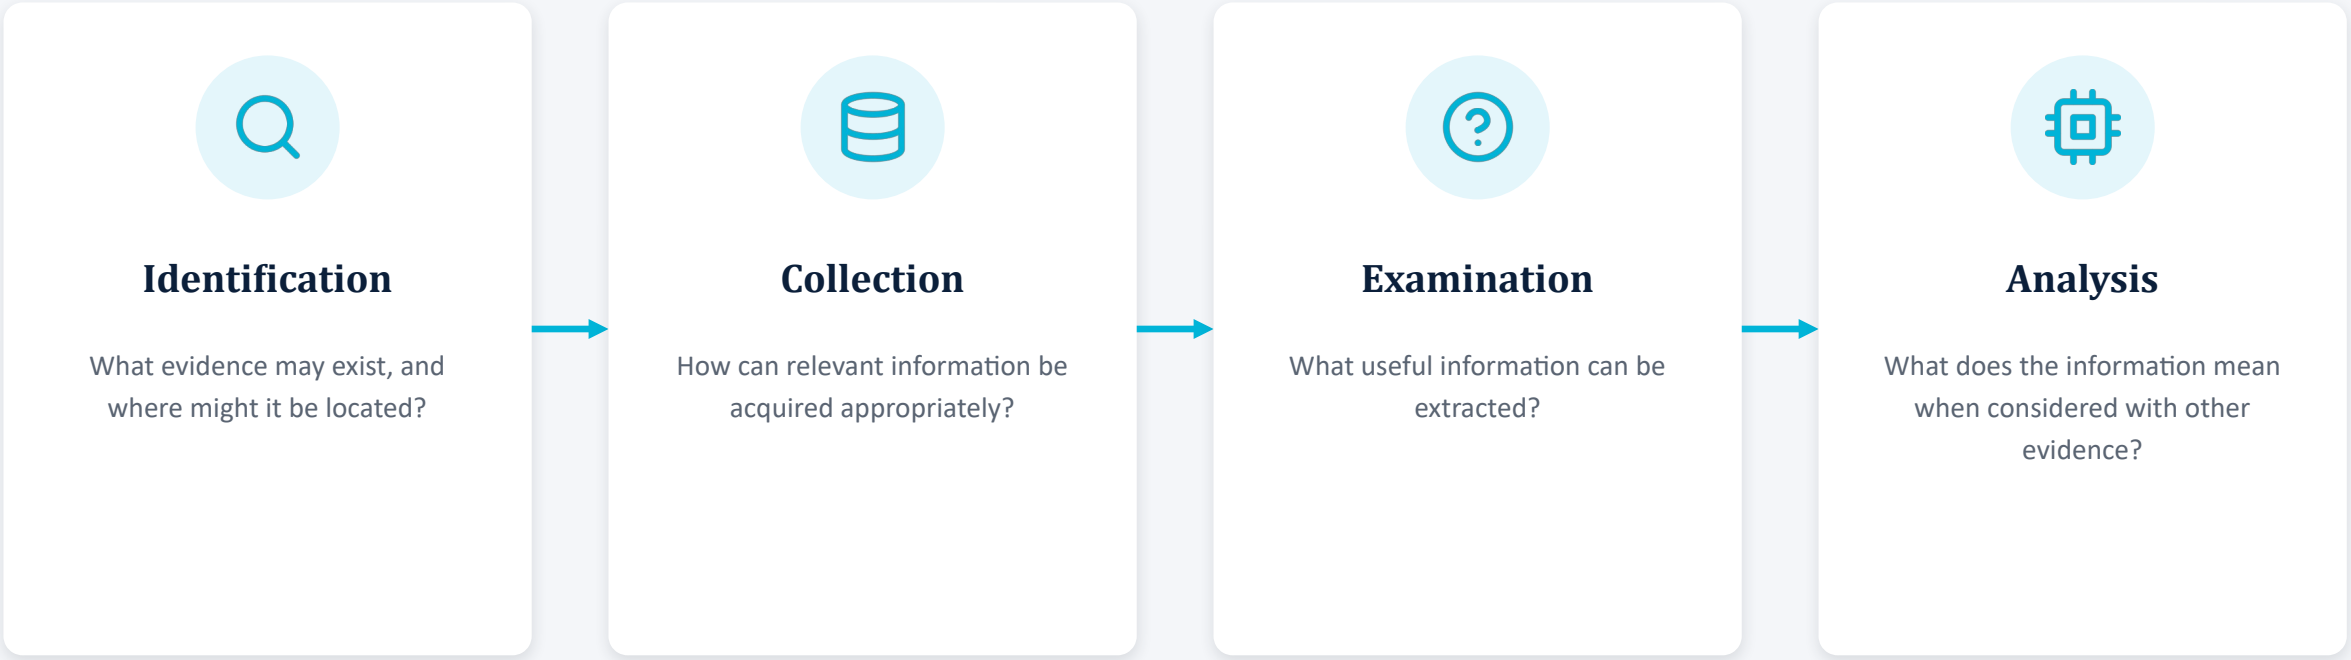

What This Session Is — and Isn't

This Is Not

- A digital-forensics certification
- Training to personally image drives or extract memory dumps
- Comprehensive coverage of every NIST SP 800-86 technique

This Is

- Understanding why systematic evidence handling matters professionally
- Building the judgment to know what's authorized to examine, and when to escalate to a specialist

SECTION 4

Evidence Integrity and Chain of Custody

A technically interesting file is less useful if nobody can explain where it came from

CHAIN OF CUSTODY

What Professionals Document

- What was collected
- Who collected it
- When it was collected
- Where it came from
- How it was collected
- Integrity information (hashes)
- Who subsequently handled it

“A technically interesting file is less useful if nobody can explain where it came from or whether it changed.”

SECTION 5

AI-Assisted Digital Investigation

Genuinely useful — and, right now, generating real courtroom sanctions when trusted blindly

Genuine Capability, Genuine Risk

Potential Uses

- Summarizing logs
- Identifying patterns
- Developing investigative questions
- Organizing timelines
- Explaining unfamiliar events
- Correlating information
- Helping organize reports

Real Risks

- Inventing events
- Misinterpreting timestamps
- Overlooking context
- Creating false correlations
- Omitting relevant evidence
- Exposing confidential information
- Generating conclusions that can't be reproduced

AI-ASSISTED INVESTIGATION

A Concrete Example

“An AI-generated incident summary says an employee exfiltrated data at 2:14 a.m.”

The professional response is NOT:

- Copying that statement directly into the incident report

The professional response IS:

- Returning to the underlying evidence and verifying it

2026 CASE IN THE NEWS

This Exact Failure Is Playing Out in Courts Right Now

AI-generated fake citations: 1,598 documented U.S. and international court cases as of June 2026

The trend is accelerating, not slowing

719 documented cases by January 2026, 1,227 by early April, 1,598 by June 9 — nearly 8 new cases a day and climbing.

Penalties are real money and real careers

A combined ~$109,700 in sanctions and fees in one Oregon case; a Nebraska attorney suspended after denying, then admitting, AI use in a brief with 57 of 63 citations defective.

Courts converge on one standard

Multiple appellate courts now use nearly identical language: no filing should contain a citation the responsible professional has not personally read and verified.

THE PATTERN THAT MATTERS

A lawyer citing a fake case and an analyst copying an unverified AI timestamp are the same professional failure.

Both skip verification. Both treat AI confidence as a substitute for evidence.

Both are entirely preventable with one habit: return to the source before you conclude.

SECTION 6

Evidence vs. Analysis

The distinction underneath everything else covered today

Keep This Distinction in Mind All Week

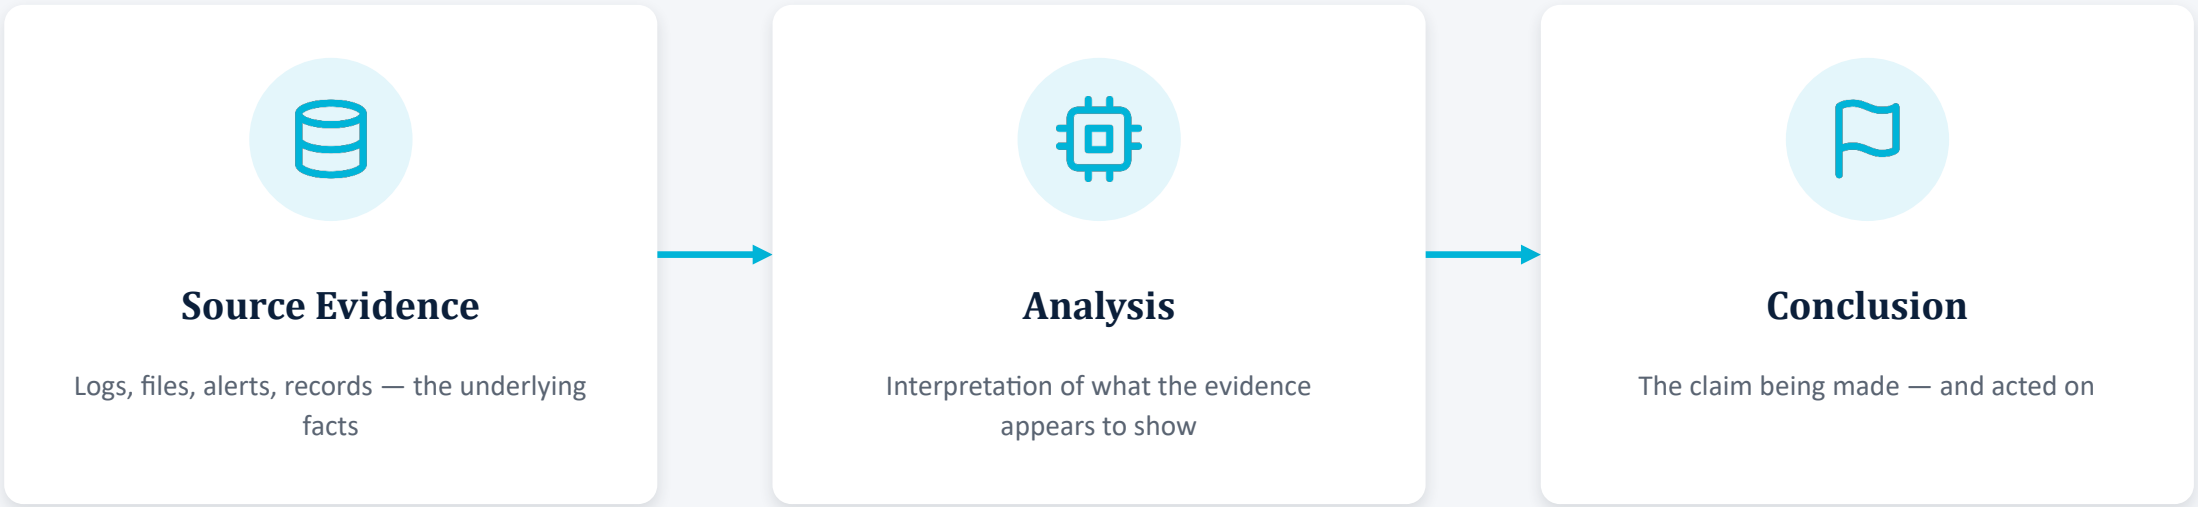

A security tool alert is not automatically proof. An AI-generated summary is not automatically proof. An analyst's theory is not automatically proof.

SECTION 7

Mock Legal/Ethical Decision Exercise

A suspected insider incident — with real pressure to skip steps

MOCK EXERCISE

The Scenario

SCENARIO

The organization believes an employee may have copied proprietary information. The employee uses a company laptop, but also has a personal cloud-storage account accessible through the browser. The security team may technically be able to access that personal account. Management wants the team to “check everything.” An AI incident-analysis tool suggests that company files were probably transferred to the personal account.

MOCK EXERCISE

A Real Legal Wrinkle: The Stored Communications Act

Accessing an employee's personal accounts — even ones left logged in on a company device — can create real federal liability, not just a policy question.

Benz v. PHB Realty Company

An employer accessed a former employee's personal email and social media after finding them logged in on a work computer. A federal court in Kansas allowed her Stored Communications Act claim to proceed — the case later settled.

Why this matters for today's scenario

Just because access is technically possible — an account left logged in, a saved password — does not mean examining it is authorized or legally safe.

MOCK EXERCISE

What You Must Determine — Part 1

What evidence actually exists

Separate from what management or the AI tool has assumed or suggested.

What remains an assumption

Including the AI tool's suggestion that files were transferred.

What the security team is authorized to examine

Company laptop and company accounts are not automatically the same as the personal cloud account.

What privacy concerns exist

For the employee, and for anyone else whose data might be visible in that personal account.

MOCK EXERCISE

What You Must Determine — Part 2

What evidence should be preserved

And how, given what was covered on chain of custody today.

Whether additional authority is required

Before anyone examines the personal cloud account.

How much weight the AI conclusion should receive

Given everything covered about AI-generated conclusions today.

What the security team should do next

A concrete, defensible course of action — not just a list of concerns.

SECTION 9

Wrap-Up

Connecting today back to the whole of Week 2

BEFORE YOU LEAVE TODAY

You Should Be Able to Explain Why These Statements Matter

- An analytical conclusion is not the same thing as the evidence supporting it.
- AI can help identify investigative questions, but underlying evidence must be verified.
- Technical access does not automatically establish legal authority — even to a personal account left logged in.
- A technically interesting finding is only as useful as the documentation behind it.
- 1,598 court sanctions in 2026 tell the same story cybersecurity investigators need to hear: verify before you conclude.

NEXT WEEK

Privacy, Surveillance & Digital Identity

Who gets to collect information about you? Who gets to monitor you? What happens to the data after it's collected? How much surveillance is justified in the name of security?

You'll examine these through privacy and surveillance cases, stakeholder analysis, regulatory requirements, and technologies that increasingly identify, track, profile, and evaluate people.