CYBR-2100 Cyber Ethics & Cyber Law

Week 3 - Privacy, Surveillance & Digital Identity

Course outline captured from Blackboard

Welcome to Week 3.

Last week, we examined authorization and the legal boundaries of cybersecurity activity.

This week, we turn to the information cybersecurity professionals collect, monitor, protect, analyze, and increasingly use to make decisions about people.

Cybersecurity and privacy frequently overlap, but they are not the same thing.

A system can be secure and still create significant privacy concerns.

This Week's Big Question

How much information should an organization collect, monitor, infer, and retain in the name of security?

Modern cybersecurity systems can collect enormous amounts of information about individuals.

They can also use AI and analytics to create new information through inference.

A system might collect login activity but infer insider-threat risk.

It might collect device activity but infer productivity.

It might capture an image but use that image to determine identity.

This week, you will examine what happens when cybersecurity, privacy, surveillance, digital identity, and AI intersect.

What You Will Learn

By the end of Week 3, you should be able to:

- Distinguish privacy risk from cybersecurity risk.
- Identify how organizations collect, use, share, retain, and dispose of information.
- Distinguish collected data from inferred information.
- Evaluate privacy risks created by cybersecurity monitoring.
- Analyze surveillance for necessity and proportionality.
- Identify stakeholder rights, responsibilities, risks, and potential harms.
- Evaluate digital identity technologies.
- Identify privacy and security concerns associated with biometrics.
- Determine what information is needed before claiming that a privacy law or regulation applies.

- Evaluate AI-enabled monitoring, profiling, scoring, and decision-making.
- Recommend practical privacy and cybersecurity safeguards.

Privacy and Cybersecurity Are Not the Same Thing

Cybersecurity often focuses on protecting systems and information from risks such as:

- unauthorized access;
- disclosure;
- modification;
- disruption;
- compromise;
- destruction.

Privacy asks additional questions about how information about individuals is processed.

For example:

- Should this information be collected?
- Why is it needed?
- What will it be used for?
- Who can access it?
- Will it be combined with other information?
- What can be inferred from it?
- How long will it be retained?
- Can it be used for another purpose?
- Can an individual challenge inaccurate information?
- What happens when a decision is made from it?

A database can be extremely well protected from attackers while the organization's authorized use of the data still creates privacy problems.

The Data Life Cycle

When evaluating privacy, look beyond initial collection.

Consider the entire life of the information.

Collection

What information is gathered?

Generation and Inference

What additional information is created from the collected data?

Use

Why is the information being processed?

Sharing

Who receives it?

Storage

Where is it maintained and protected?

Retention

How long is it kept?

Disposal

What eventually happens to it?

Collected Data vs. Inferred Data

This distinction becomes increasingly important as organizations use AI.

An organization might collect:

- login times;
- IP addresses;
- location information;
- browsing activity;
- file activity;
- facial images;
- device information.

An AI-enabled system might then infer:

- insider-threat risk;
- identity confidence;
- productivity;
- suspicious behavior;
- likelihood of fraud;
- unusual behavior.

The inference may have significant consequences even though the organization never directly collected that particular conclusion from the individual.

Class Meeting 1

Privacy, Data Collection & Surveillance

We begin with an employee-monitoring scenario.

A fictional organization deploys an endpoint security platform that collects:

- login activity;
- browsing information;
- application activity;
- file activity;
- USB-device use;
- location information;
- screenshots;
- typing and mouse activity.

The platform also produces AI-generated productivity and insider-risk scores.

You will determine:

- which information supports legitimate cybersecurity objectives;
- which collection may be excessive;
- who should have access;
- whether cybersecurity data should be reused by Human Resources;
- what employees should know;
- how long the data should be retained;
- what happens when an AI-generated score is wrong.

Privacy Data Map

You will create a simple data map for the monitoring system.

For different types of information, your team will identify:

- what is collected;
- why it is collected;
- where it goes;
- who can access it;
- what decisions may be based on it;
- how long it should be retained;
- possible harms;
- information AI may infer.

The purpose is to see how quickly a cybersecurity system can become a significant privacy system.

NIST Privacy Framework

We will begin using the NIST Privacy Framework to organize privacy-risk thinking.

NIST Privacy Framework:

https://www.nist.gov/privacy-framework

You do not need to memorize the entire framework.

We will use its major functions as practical questions.

Identify-P

What data processing occurs, and what privacy risks does it create?

Govern-P

Who is responsible for privacy decisions?

Control-P

What controls and data-management capabilities are available?

Communicate-P

What do individuals and organizations understand about the processing?

Protect-P

How is the information protected?

AI Privacy Analysis Lab

You will use ChatGPT Edu to evaluate the fictional employee-monitoring system.

The goal is not to have AI make the privacy decision for you.

You will use it to identify potential issues and then evaluate the quality of its analysis.

Look for whether the AI:

- identifies secondary use;
- distinguishes collected from inferred information;
- identifies excessive collection;
- questions retention;
- considers inaccurate risk or productivity scores;
- identifies affected stakeholders;
- considers power differences;
- assumes that particular laws apply without enough information;
- recommends practical safeguards.

You will compare its analysis with your own.

Between Class Meetings

Privacy Observation Activity

Choose one technology that collects, generates, or infers information about individuals.

Examples include:

- smartphone;
- learning platform;
- social media platform;
- fitness device;
- streaming service;
- vehicle;
- smart speaker;
- workplace system;
- facial-recognition system;
- AI assistant.

Do not submit private personal information.

Come to Class Meeting 2 prepared to explain:

- What data the technology collects.
- What information it may infer.
- Why the organization wants the information.
- One potential privacy benefit.
- One potential privacy harm.
- One question you would ask the organization about its use of the data.

Class Meeting 2

Digital Identity, Regulation & AI-Enabled Surveillance

The second class meeting examines what happens when information is used to establish, verify, monitor, classify, or make decisions about identity.

Digital Identity

Digital identity is more than a username.

Identity systems may involve:

- identifiers;

- credentials;
- authentication factors;
- identity attributes;
- devices;
- biometrics;
- behavioral characteristics;
- identity proofing;
- federation.

We will use current NIST Digital Identity Guidelines.

NIST SP 800-63-4 Digital Identity Guidelines:

https://csrc.nist.gov/pubs/sp/800/63/4/final

Identity Proofing

Identity proofing asks:

Who are you, and how does the organization establish that?

Authentication

Authentication asks:

Can you demonstrate control of an authenticator associated with that identity?

Federation

Federation allows identity information established by one trusted system to be used by another service.

Digital Identity Case

You will evaluate a fictional college identity-verification system intended to reduce account sharing during online examinations.

The proposed system requires students to:

- photograph government identification;
- capture a live facial image;
- complete facial comparison;
- permit behavioral monitoring;
- allow AI analysis of gaze, movement, background noise, and device behavior.

The vendor retains some information to improve its algorithms.

You will examine:

- cybersecurity and academic-integrity benefits;
- data collection;

- data inference;
- necessity;
- proportionality;
- false matches;
- accessibility;
- retention;
- vendor use;
- AI accuracy;
- stakeholder impact.

Biometrics

Biometric technologies may use:

- face;
- fingerprints;
- voice;
- iris;
- gait;
- behavioral characteristics.

Consider an important security difference:

A compromised password can usually be changed.

A person's face or fingerprint cannot simply be replaced.

AI-enabled biometric systems also raise questions involving:

- accuracy;
- false matches;
- false non-matches;
- demographic performance;
- accessibility;
- consent;
- surveillance;
- data retention;
- secondary use.

Privacy Laws and Regulatory Obligations

Privacy regulation depends on context.

Before deciding that a particular law or regulation applies, ask:

- What organization is involved?
- What information is involved?
- Whose information is involved?
- What industry is involved?
- What jurisdiction applies?
- What is being done with the information?

Do not assume that every situation involving health information automatically means HIPAA applies.

Do not assume every educational situation automatically means FERPA resolves the issue.

Do not assume that every privacy problem is automatically a regulatory violation.

First determine applicability.

FTC Privacy and Security Resources

The Federal Trade Commission provides privacy and data-security guidance that we may use during this week's work.

https://www.ftc.gov/business-guidance/privacy-security

Surveillance Proportionality

Security monitoring exists on a spectrum.

An organization responding to a security problem might consider:

- Reviewing authentication logs.
- Increasing endpoint logging.
- Monitoring browsing activity.
- Capturing screenshots.
- Activating webcam monitoring.
- Using AI to evaluate behavior.
- Continuously tracking location.

More monitoring does not automatically mean better security.

You will evaluate when surveillance becomes disproportionate to the problem being addressed.

An important question is:

Could a less intrusive control achieve the same security objective?

AI Policy Reviewer

You will also use ChatGPT Edu to review a fictional monitoring policy or technology proposal.

You will ask AI to identify possible:

- unclear purposes;
- excessive collection;

- secondary uses;
- retention problems;
- access risks;
- AI-generated inferences;
- bias concerns;
- transparency gaps;
- governance weaknesses.

Then you will evaluate whether the AI critique is actually correct.

Your Week 3 Graded Work

Privacy & Surveillance Analysis

Points: 40

Category: Case Studies, Simulations & Interactive Activities

You will analyze a privacy or surveillance scenario.

Your analysis will address:

- data collection;
- data use;
- data inference;
- stakeholders;
- privacy risks;
- cybersecurity benefits;
- necessity and proportionality;
- legal and regulatory questions;
- ethical considerations;
- AI or automation;
- bias and equity;
- safeguards;
- professional recommendation.

You will use authoritative resources to support significant claims.

Policy/Compliance Activity

Points: 15

Category: Knowledge Checks & Applied Analysis

Instead of a traditional quiz, you will review a fictional organizational policy or proposed technology deployment.

You will identify:

- strengths;
- weaknesses;
- privacy concerns;
- cybersecurity concerns;
- legal or regulatory questions;
- AI and automated-decision risks;
- missing governance controls;
- recommended improvements.

You will be evaluated on your ability to identify and explain meaningful issues rather than memorize regulatory terminology.

Preparation & Professional Engagement

Points: 10

Category: Professionalism & Engagement

This week's engagement includes:

- Privacy Observation preparation;
- privacy data mapping;
- monitoring discussions;
- NIST Privacy Framework activities;
- AI privacy analysis;
- digital identity case;
- biometric analysis;
- regulatory applicability discussion;
- surveillance proportionality exercise;
- AI policy review.

Week 3 Checklist

Before completing Week 3:

- Participate in Class Meeting 1.
- Participate in the employee-monitoring analysis.
- Complete the Privacy Data Map activity.
- Work with the introductory NIST Privacy Framework concepts.
- Participate in the AI Privacy Analysis Lab.
- Complete the Privacy Observation Activity before Class Meeting 2.
- Participate in the digital identity discussion.
- Evaluate the digital identity case.

- Participate in the biometric and equity discussion.
- Participate in the regulatory applicability activity.
- Complete the Surveillance Proportionality exercise.
- Participate in the AI Policy Reviewer activity.
- Submit the Privacy & Surveillance Analysis.
- Submit the Policy/Compliance Activity.
- Complete required professional engagement activities.

Before You Leave Week 3

You should be able to explain why these statements matter:

Secure does not automatically mean private.

Collecting information and inferring information are not the same thing.

A legitimate security objective does not automatically justify unlimited surveillance.

More data does not automatically produce better decisions.

An AI-generated risk score is an inference, not an unquestionable fact.

Regulatory applicability must be established rather than assumed.

Privacy decisions affect people, not simply databases.

What to Expect Next Week

Week 4 moves into intellectual property, data ownership, software and content rights, responsible disclosure, vulnerability research, and professional conduct.

We will examine questions such as:

Who owns code?

Who owns data?

What can you legally and ethically reuse?

What happens when AI generates code or content based on existing material?

What should a cybersecurity researcher do after discovering a vulnerability?

When does responsible disclosure become irresponsible disclosure?

And what happens when the interests of researchers, software vendors, customers, and the public conflict?