

CS 458 / 658: Computer Security and Privacy

Module 7 – Non-technical Aspects of Security and Privacy

Part 1 – Ethics and legal issues

Fall 2022

Module outline

- 1 Why studying ethics and laws?
- 2 Differences between laws, morality, and ethics
- 3 Ethical practices in security and privacy domain
- 4 Intellectual property
- 5 Other common legal issues in security and privacy domain

Motivation

- The course content includes a wide range of attacks.
- These attacks can have societal impacts and individual impacts.
- Your future work, being it research, industry, start-ups, software, security, ..., depends on your awareness of legal and ethical issues.

Cambridge Analytica




Facebook–Cambridge Analytica data scandal

A timeline of the Cambridge Analytica scandal

- In 2010, Facebook launched Open Graph. External developers can reach out to FB users and request access to not only their personal data, but also their **friends' personal data** too!
- In 2013, an app “thisisyourdigitallife” approached to almost 300,000 users and paid them to take a psychological test.
- In 2014, Facebook adapted its rules to limit a developer's access to user data, especially the friends' data.
- In 2015, The Guardian reported that Cambridge Analytica was helping Ted Cruz's presidential campaign. FB acknowledged the data leak and argued that they have legally pressured Cambridge Analytica to remove all of the data they had improperly acquired.
- In 2016, Cambridge Analytica was responsible for the “Defeat Crooked Hilary” video campaign on FB (assisting Trump's team).

A timeline of the Cambridge Analytica scandal



Donald Trump and Ted Cruz shake hands before the start of the Republican Presidential Debate (2015)

A timeline of the Cambridge Analytica scandal



Christopher Wylie, whistleblower of the Cambridge Analytica scandal

A timeline of the Cambridge Analytica scandal


- In March 2018, the scandal is exposed by The Guardian and The New York Times. The initial number is 50 million user profiles and later revised to 87 million (estimated by FB).
- In March 2018, Mark Zuckerberg first apologized for the situation, calling it an “issue”, a “mistake” and a “breach of trust”.
- In July 2018, United Kingdom’s Information Commissioner’s Office announced to fine FB £500,000 (\$663,000)
- In July 2019, the Federal Trade Commission announced to fine FB around \$5 billion to settle the data breach investigation
- In July 2019, the Securities and Exchange Commission announced to fine FB around \$100 million for misleading investors about the risks it faced from misuse of user data

Linux kernel and the University of Minnesota

- You are an open-source enthusiast and make contributions to open-source projects regularly
- You are deeply concerned that the Linux kernel might be vulnerable to supply chain attacks due to its loose review process
- You want to remind the code reviewers that they should tighten up their code review practices
- So you send an intentionally buggy piece of code to the Linux kernel reviewer and ask for it to be merged into the upstream

Q: If you were a developer, how would you feel about this?

Linux kernel and the University of Minnesota



What we learned from the examples?

Be extremely cautious when humans are involved in any form of activity, regardless of physical or virtual presence.

Fortunately, we have laws and ethics to guide us on making a right-or-wrong judgement call.

Module outline

- 1 Why studying ethics and laws?
- 2 Differences between laws, morality, and ethics**
- 3 Ethical practices in security and privacy domain
- 4 Intellectual property
- 5 Other common legal issues in security and privacy domain

Laws, morality, and ethics

Q: What do laws, morality, and ethics have in common?

A: They are all beliefs, claims, rules, and norms about how we should live and behave.

Q: What are the differences between laws, morality, and ethics?

What is law?

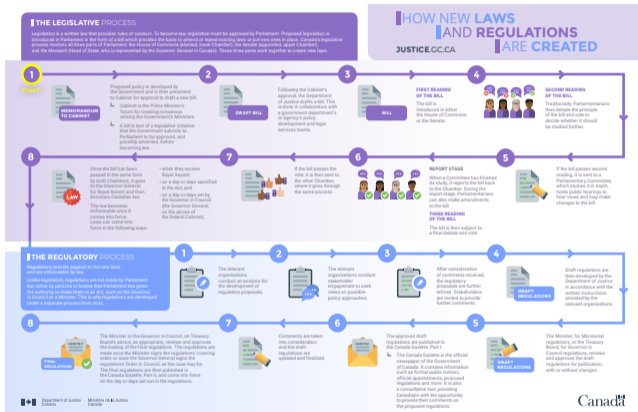
- Laws are a set of **formal rules** that governs how we behave as members of a society.
- The goal is to create a set of **basic and objectively enforceable** standard of behaviors.
- Specifies, in greater details, what we **MUST** do and more frequently, what we **MUST NOT** do.
- Laws are upheld and applied by a state-backed justice system.

Q: Why are laws not enough in the context of computer security and privacy?

A:

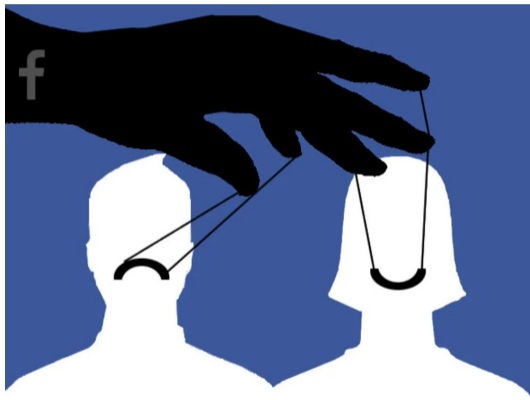
- The **lengthy legislative process** does not match with the fast-pacing tech industry.
- Laws usually have a very **narrow** focus.

Lengthy legislative process



The legislative process in Canada

Non-violations of law



The (secret) mood manipulation study by Facebook in 2012

Some facts about the mood manipulation study

For one week in January 2012, data scientists in Facebook skewed the content of News Feed for 689,003 users.

- Some people were shown content with more positive words
- Others were shown content analyzed as sadder than average.

Finding 1: More negative News Feeds led to more negative status messages, as more positive News Feeds led to positive statuses.

Finding 2: Omitting (either positive or negative) emotional content reduced the amount of words the person subsequently produced.

Right and wrong

Q: Why there are no legal violations here?

A: Quoted from [Facebook Terms of Service](#):

Product research and development: We use the information we have to develop, test and improve our Products, including by conducting surveys and research, and testing and troubleshooting new products and features.

However, we might still have some upset feelings here.

What is morality?

- Morality refers to an **informal** framework of values, principles, beliefs, customs, ways of living.
- Morality is usually not enforced by the state, but by **social pressure** to conform to moral norms.
- An individual who is strongly bounded to a moral system may even consider questioning the moral system as wrong.
- Usually, the process of moral formation is **unconscious**, e.g., by family, by community, or by culture.
- The application of morality is almost a habit without an explicit thinking and reasoning process.

Q: A legal + moral system to classify the rights and wrongs?

A: There is rarely a moral authority agreed by every individual

From morality to ethics

“The unexamined life is not worth living” — Socrates

What is ethics?

- Ethics is a branch of philosophy that answers a basic question: **what should I do?** (out of all possibilities).
- Usually, the process of making an ethical decision is a **conscious** reasoning process based on each individual's values, principles, and purpose — do something that is good, right, and meaningful.
- Ethics is the framework to reason about issues that the **laws** cannot or do not address.
- Ethics is the framework to examine a **moral system** to see whether the principles and rules there make sense.

General theories of ethics

There are three approaches to ethical analysis:

- Consequentialism
 - concerned with the ethical **consequences** of particular actions
- Deontology (duty-based)
 - concerned with the **actions** themselves, and not the consequences
 - following one's **duty**
- Virtue ethics (agent-centered)
 - concerned about **agents** (individuals) and less concerned with particular actions or consequences
 - focus on virtues like courage, temperance, wisdom, justice, generosity, etc.

Ethical Theories: summary

Deliberative process

What is the **outcome** I should (try) to produce?

What is my **obligation** (duty) in this situation?

What kind of **person** should I be (or try to be)?



Ethical conduct

I will do the action that will achieve the best **consequences**

I will do the right thing, I will never fail to do my **duty**

I will do whatever a fully **virtuous** person would do in the circumstances

Motivation

Consequentialist
I aim to produce the most good

Deontology
I aim to perform the right action

Virtue ethics
I aim to develop my character

Module outline

- 1 Why studying ethics and laws?
- 2 Differences between laws, morality, and ethics
- 3 Ethical practices in security and privacy domain**
- 4 Intellectual property
- 5 Other common legal issues in security and privacy domain

Ethical practices in security and privacy

We will see a few ethical practices:

- ① Responsible disclosure
- ② Build ethically
- ③ Talk to non-experts
- ④ Talk to independent experts
- ⑤ Codes of professional ethics

Responsible disclosure

Q: You have found a security vulnerability, what should you do?

- 1 Make it public.
- 2 Tell the responsible parties first.

Responsible vulnerability disclosure

- A **private full disclosure** to all responsible parties (e.g., software vendors for most software bugs).
- Wait for either a patch from the responsible parties or a specific amount of time (e.g., 90 days or 120 days).
- A **public partial disclosure** if you want to further pressure the responsible parties; or a **public full disclosure** in the interests of potential victims.

Build ethically

Tips for incorporate ethical decisions when building something new:

- Get as many **dissenting** voices as possible.
- Explain how something works, what is possible to go wrong, and how bad actors can take advantage to a **non-expert**.
- The privacy and data protection norms and cultural values **vary by region and country**.
- **Consult** other experts (e.g. ethics, religions, advocates, activists).

Talk to non-experts



What if the tool works as intended?

- Who does this affect?
- Does this data need to be collected?

What if the tool does not work as intended?

- Failure modes? Abuse cases?
- Who does this affect?

Talk to independent experts

Institutional review board (IRB), a.k.a., independent ethics committee (IEC), ethical review board (ERB), or research ethics board (REB), etc...

IRB is a committee that applies research ethics by reviewing the methods proposed for research to ensure that they are ethical.

Codes of professional ethics

You will probably be a member of one or more professional societies:

- Association for Computing Machinery (ACM).
- Institute of Electrical and Electronics Engineers (IEEE).
- Canadian Information Processing Society (CIPS).

These organizations have **codes of professional ethics**.

Example: Canadian Information Processing Society (CIPS)

These are the high-level bullets from CIPS' code:

- Protect Public Interest and Maintain Integrity.
- Demonstrate Competence and Quality of Service.
- Maintain Confidential Information and Privacy.
- Avoid Conflicts of Interest.
- Uphold Responsibility to the IT Profession.

Module outline

- 1 Why studying ethics and laws?
- 2 Differences between laws, morality, and ethics
- 3 Ethical practices in security and privacy domain
- 4 Intellectual property**
- 5 Other common legal issues in security and privacy domain

Legal protections

How can we defend against a threat?

- Prevent it: block the attack
- Deter it: make the attack harder or more expensive
- Deflect it: make yourself less attractive to attacker
- Detect it: notice that attack is occurring (or has occurred)
- Recover from it: mitigate the effects of the attack

In addition to (sometimes instead of, unfortunately) using technological defenses, we can also use **legal** defenses.

Overview of intellectual property

In contrast to real property, so-called “intellectual property” (IP) differs in important ways:

- It is **non-depletable**.
- It is **replicable**.
- It has **minimal marginal cost**.

So the laws for IP differ from the laws for real property, and indeed are much more complicated.

Types of intellectual property

Four kinds of IP here:

- 1 Trade secrets
- 2 Trademarks
- 3 Patents
- 4 Copyrights

These four kinds of IP:

- Cover different intangibles
- Convey different rights
- Have different durations
- Use different registration process

Trade secrets

- This is the simplest kind of IP
- You want to protect some secret information
 - The formula for Coca-Cola
 - The method for computing how many airline seats to oversell
 - Your new $O(n)$ sorting algorithm
- Just don't tell anyone, and call it a trade secret
 - Unfortunately, you have to tell **someone**, or it's not useful
 - **You get legal protection if that person passes it on**

Trade secrets: reverse engineering

- **Reverse engineering** is the process of taking a finished product and taking it apart to figure out how it works
 - If someone successfully does this and published the results, you've effectively **lost your trade secret** protection
 - General rule for trade secrets: **it has to be a secret**
- A similar rule applies to software, with some caveats we'll see later
- RC4 was originally a trade secret (a violation of Kerckhoffs's principle), but it was reverse engineered in 1994

Trademarks™

- Trademarks protect **names, brands, logos** and also **domain names**
- To get one, make *some* legal filing showing that you are using the name in commerce
 - This lets you sue others who use that name in a confusing manner
- Trademarks last while they are used (they have to be renewed)

Example: Microsoft vs MikeRoweSoft

- Mike Rowe used to own the domain name “MikeRoweSoft.com”
- Settlement reached: Rowe granted Microsoft ownership of MikeRoweSoft domain in exchange for an Xbox (and additional compensation)

Patents

- Applies to **inventions** (including algorithms), which must be:
 - Novel
 - Useful
 - Non-obvious
- The bargain is that:
 - You tell everyone how your invention works
 - In exchange, you get to have a monopoly over it for 20 years
- The most difficult form of IP to obtain

Patents in cryptography

Many cryptographic algorithms are (or were) patented

- Diffie-Hellman (expired 1997)
- RSA (expired 2000)
- IDEA (block cipher used in early PGP, expired 2012)
- Lots of patents on elliptic curve cryptography

Since 2000, you could pick a good unpatented candidate of each type of crypto

NOTE: unlike trade secrets, this is not against Kerckhoffs's principle.

Copyright[©]

- Copyright is the most well-known kind of IP
- Protects expressions of ideas in a tangible medium
 - But not ideas themselves!
- Example: you design a sorting algorithm.
 - Copyright protects copies of your implementation, but not other implementations of your idea.
- No filing requirement
 - But you can get additional benefits if you do file
- Lasts a “limited time”
 - Currently: life+70 years in the US, life+50 in Canada
- The copyright holder has monopoly rights over certain uses of the work; primarily, making copies

Let's fill this table

Type	Covers	Example	Filing?	Duration
Trade secrets				
Trademarks				
Patents				
Copyright				

Q: Covers: 1) inventions; 2) brands, names, logos; 3) expressions; 4) secrets

Examples: 1) a secret formula; 2) a melody; 3) an engine design; 4) a domain name

Filing: 1) none; 2) some; 3) not necessary (but additional benefits); 4) complicated filing

Duration: 1) 20 years; 2) as long as it's used/defended; 3) as long as it's secret; 4) lifetime + 50 (CA)/70 (US) years

Solution

A:

Type	Covers	Example	Filing?	Duration
Trade secrets	secrets	secret formula	none	as long as it's secret
Trademarks	brands, names, logos	a domain name	some filing	as long as it's defended
Patents	inventions	an engine design	complicated filing	20 years
Copyright	expressions	a melody, a movie, a videogame, an implementation	not necessary, but helps	lifetime + 50 (CA)/ 70 (US) years

Module outline

- 1 Why studying ethics and laws?
- 2 Differences between laws, morality, and ethics
- 3 Ethical practices in security and privacy domain
- 4 Intellectual property
- 5 Other common legal issues in security and privacy domain

Cyber crime

- We saw that laws regarding intellectual property differ from those about real property
- Similarly, laws about unauthorized access to computers, networks, or services differ from those about physical trespass
 - But until those new laws came about, courts had to make really stretched analogies to handle such events

Cyber crime

- Early on, there were bizarre rulings:
 - The value of stolen data was the value of the paper it was printed on
 - The value of a stolen manual was the value of the equipment it was intended for
- Things seem to have settled down somewhat
 - GDPR:
General Data Protection Regulation
 - PIPEDA:
The Personal Information Protection and Electronic Documents Act
 - HIPAA:
Health Insurance Portability and Accountability Act
- But there are still many recent and active issues!
 - If your ISP keeps a copy of your incoming email, is that wiretapping?

Rules of evidence

Another problem with prosecuting computer crime is producing evidence admissible in court:

- Should the log files of the machine that was broken into be admissible?
- How should you preserve electronic evidence from the time of the intrusion to the time of a possible trial?

Computer forensics replace regular forensics

Cybercrime treaty

- Worse, computer crime is often international (two or more jurisdictions)
- Rules of evidence, police powers, etc. in one country don't usually carry over to another
- The Council of Europe cybercrime treaty (to which Canada and the US are also signatories) stipulates that member countries should pass laws making it easier for law enforcement to access telecommunications traffic (including voice, data, and Internet)

Bill C-13 (“Cyberbullying Law”)

Full name: Protecting Canadians from Online Crime Act:

- Really a “lawful access” law
- Passed in December 2014
- Any “public officer” (not just the police) can demand that any computer data in a person’s control not be deleted (until a production order can be obtained)
- Lowers standard for seizing of computer data, transmission data, and tracking data to “reasonable grounds for suspicion”
- Provides immunity to ISPs that “voluntarily” hand over customer data to government
 - Even though the Supreme Court had recently ruled that unconstitutional!

Recap: Ethics and Legal Issues

- Laws, morality, ethics:
 - Laws: set of formal rules enforced by a justice system
 - Morality: informal framework of values, principles, beliefs, etc., many times unconscious, and enforced by social pressure
 - Ethics: branch of philosophy that questions what to do following a reasoning process.
- Ethical theories: consequentialist (focused on the outcome), deontology (focused on the duty), virtue ethics (focused on developing your character).
- Ethical practices in security and privacy: responsible disclosure, build ethically, talk to non-experts, talk to independent experts, follow codes of professional ethics.
- Intellectual property: very different from physical property.
 - 1 Trade secret: protects secrets as long as they are secret. No need to file.
 - 2 Trademarks: protects names, brands, logos, domains. Easy to file, holds while in use.
 - 3 Patents: protects inventions (you tell everyone how they work). Hard to obtain, holds for 20 years.
 - 4 Copyright: does not protect the idea, but the expression of the idea in a tangible medium. No filing required, last for the author's lifetime +50/70 years.
- Legal issues: cyber crimes vs. physical crimes. Bizarre rulings at first, now we have better laws, international treaties, etc. But still a lot of work to be done!

CS 458 / 658: Computer Security and Privacy

Module 7 – Non-technical Aspects of Security and Privacy

Part 2 – Administering security and privacy

Fall 2022

Module outline

- 6 Security planning
- 7 Risk analysis
- 8 Closing remarks

Security planning

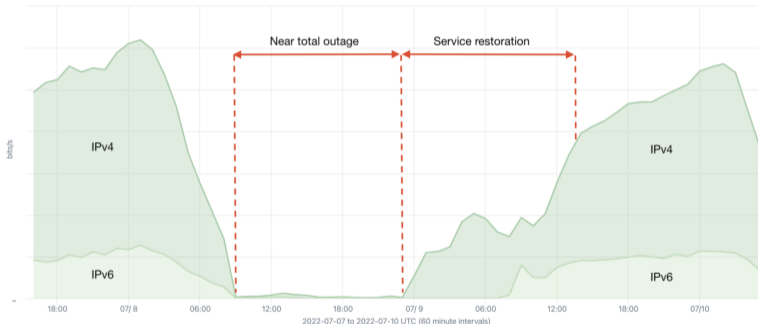
It used to be that employees understood that when you went home for the day, you locked up all your files in your filing cabinet.

How do you achieve the same effect today now that files are all electronic?

Red alert!

Top INET Family by Average bits/s
 Jul 07, 2022 16:00 to Jul 10, 2022 06:00 (2d and 14h)

Internet traffic to Rogers (AS812)



Q: What should you do as a Rogers' employee at 6:00 UTC to provide availability?

Do not “think” about what to do, simply follow the security plan

Goals of security planning

A **security plan** is a document that explains

- what the security goals are (e.g., availability, confidentiality)
- how they are to be met (e.g., what kind of guidelines should be followed)
- how they'll **stay** met (e.g., how to keep providing these in the future)

Employees can use this document to inform their actions

Analogy: Go to a construction site and ask the manager-in-charge, what is your safety plan here?

Contents of a security plan

A security plan is both a description of the current state of the security of an organization, as well as a plan for improvement.

Usually, a security plan has seven parts:

- 1 Policy (high-level goals and priorities)
- 2 Current state (risk analysis, anticipation of new situations)
- 3 Requirements (*what* are the security and privacy needs)
- 4 Recommended controls (*how* we provide those needs)
- 5 Accountability (*who* is responsible for what)
- 6 Timetable (*when* the elements of the plan will be performed)
- 7 Continuing attention (*how often* the plan should be updated)

We will examine them in turn.

1/ Policy

A **policy** is a high-level statement of purpose and intent.

The policy statement should specify:

- Goals
 - Relative importance of confidentiality, integrity, availability
 - Which has higher priority: securing data or serving customers?
- Responsibility
 - Whose job is getting security right? Every employee? A security manager? A security group in IT?
 - What are their roles respectively?
- Commitment
 - Institutionally, who provides security support for staff?
 - Where does security fit into the organization chart?

2/ Current state

- A risk analysis (see later) of the current status of the system
 - What assets and controls are there?
 - What might go wrong?
 - What vulnerabilities are currently exposed?
- A reasonable anticipation of new situations:
 - What to do if new assets are added?
 - How to respond to the discovery of new vulnerabilities?
- A fair comparison with state-of-the-art and state-of-the-practice
 - What are the recent incidents?
 - Why they happen in the first place?
 - What can we learn from their experience?

3/ Requirements

What are the security and privacy needs of the organization (aligned with the goals in step 1), e.g.,

- **Who** is allowed/not allowed to do **what**?
- What audit logs should be kept?

Be careful that requirements statement should **not** say anything about **how**, i.e., the concrete “mechanism” to satisfy the needs. The requirements statement should be technology-neutral.

Example: a requirements statement might say that employees should be allowed to access their email while travelling; it should not say any of the words VPN, ssh, TLS, IPsec, etc.

4/ Recommended controls

Here's where you list mechanisms

- taking into account the priorities in the “Policy” section.
- to control issues identified in the “Current state” section.
- to satisfy the needs in the “Requirements” section.

They are essentially the security topics we've explored in this course:

- Program, OS, Network, Internet applications, Database, etc.

5/ Accountability

Who is accountable if the security controls aren't implemented, aren't implemented properly, or fail?

- Desktop users?
- Project leaders?
- Managers?
- Database admins?
- Information officers?
- Human resources?

6/ Timetable

Any reasonably sized plan will be too big to implement all at once

- Obtaining new hardware / software
- Configuring / installing it
- Training users

The timetable section of a security plan lists how and when the elements of the plan will be performed, in what order and dependency relationships.

Important: Include milestones to track progress along the way

7/ Continuing attention

"The Only Constant in Life Is Change" — Heraclitus

- The state of the organization isn't static
- The state of the world isn't static
- There will be new vulnerabilities
- Existing controls will become ineffectual

The security plan should list a process for periodic review and updating of the plan itself

Who develops the security plan?

Who performs the security analysis, makes recommendations, and writes the security plan?

The **security planning team** should have representation from a number of different constituencies:

- Upper management / CTO / CIO (setting policy)
- IT (hardware group, sysadmins)
- Systems and application programmers, DB admins
- Data entry personnel
- Physical security personnel
- Representative users
- External consulting / advisory board

Business Continuity Plan

The Business Continuity Plan (BCP) is another kind of security plan, with a sheer focus on **availability**.

It aims to lay down a way out for situations that are:

- Catastrophic: a large part (or all) of a computing capability is suddenly unavailable
- Long duration: the outage is expected to last for so long that business would suffer if left unattended

Taking actions after planning

Writing the plan is far from enough!

Before something occurs, you need to:

- Acquire redundant equipment
- Arrange for regular data backups
- Stockpile supplies
- Train employees so that they know how to react
 - This may also involve **live testing** of the Business Continuity Plan (BCP)

Module outline

- 6 Security planning
- 7 Risk analysis**
- 8 Closing remarks

Risk

Definition: A **risk** is a **potential problem** that a system or its users may experience.

Risks have two important characteristics:

- Probability: what is the probability (between 0 and 1) that the risk will occur? (That is, the **risk** will turn into a **problem**)
- Impact: if the risk occurs, what harm will happen? This is usually measured in terms of money (cost to clean up, direct losses, PR damage to the company, etc.)

The **risk exposure** = **probability** × **impact**

Motivations for risk analysis

- It is impossible to completely eliminate risk
 - No system is absolutely secure
 - The bug-free software is the software not-written
- We perform risk analysis to determine if the benefits of some action outweigh the risks
 - If not, is there anything we can do to reduce the risk exposure, either by controlling the probability or reducing the impact?
- Risk analysis is not specific to security and privacy issues
 - But bringing risk analysis to those issues is a relatively new, and extremely useful, phenomenon

Procedures for risk analysis

A risk analysis usually comprises the following steps:

- ① Identify assets
- ② Determine vulnerabilities
- ③ Estimate likelihood of exploitation
- ④ Compute risk exposure
- ⑤ Survey applicable controls
- ⑥ Project savings due to control

We're gonna see all these with an example; consider Alice, who owns a small company



1/ Identify assets

The main assets we would want to protect:

- Hardware
- Software
- Data
- **What else?**
- Reputation

Example: Alice's company



Asset	Value
Computer	\$2 000
Software	\$1 000
Client data	\$10 000

2/ Determine vulnerabilities (threat modeling)

- This step is where you apply the knowledge obtained in this course
- Also called threat modelling
- Come with as many attacks as you can, both technical and non-technical, against the assets identified before

Example: Alice's company



Asset	Value
Computer	\$2 000
Software	\$1 000
Client data	\$10 000

Threat	Affected asset
Data Breach	Steals "Client data"
...	...

3/ Estimate likelihood of exploitation

- Estimate the probability of each risk/threat
- This is the hardest step, and there are experts trained in doing it — this is called *actuarial science*
- **Frequency analysis** can be useful if the risk has happened before

Example: Alice's company



Asset	Value
Computer	\$2 000
Software	\$1 000
Client data	\$10 000

Threat	Affected asset	Prob
Data Breach	Steals "Client data"	1%
...

4/ Compute risk exposure

Recall: **risk exposure** = **probability** (of the risk) × **impact** (of the risk)

Identifying the **impact** of the risk is also a tricky step (even though estimates are usually good enough)

Some examples include:

- Legal obligations to conserve confidentiality or integrity
- Penalties for failing to provide a service
- Could release of data cause harm to a person?
- Value of keeping data out of competitor's hands
- Cost of delaying or outsourcing data processing if your systems are unavailable

4/ Compute risk exposure

Recall: **risk exposure** = **probability** (of the risk) × **impact** (of the risk)

Identifying the **impact** of the risk is also a tricky step (even though estimates are usually good enough)

Example: Alice's company



Asset	Value
Computer	\$2 000
Software	\$1 000
Client data	\$10 000

Threat	Affected asset	Prob	Impact
Data Breach	Steals "Client data"	1%	\$10 000
...

The impact of "Data Breach" is estimated to be \$10 000.

Risk exposure = $0.01 \times \$10\,000 = \100 .

5/ Survey applicable controls

- For each risk, think of different ways to control the vulnerability
 - Again, both technical and non-technical means
- Classify each control as to how well it protects against each vulnerability
 - A control that protects against one vulnerability might make another one worse!
 - Also watch out for interactions among different controls

Example: Alice's company



Threat	Prob	Impact
Data Breach	1%	\$10 000
...

Control	Cost	Protects	P. Fail
Firewall	\$100	Data Breach	0%
...

6/ Project savings due to control

- The expected cost of not controlling the risk is just the risk exposure, as computed earlier: $\text{risk exposure} = \text{probability} \times \text{impact}$
- For each control, the cost of the control is its direct cost (e.g., buying the network monitoring equipment, training, etc.), plus the exposure of the **controlled risk**
 - Most controls aren't perfect: even with the control, there will still be a (smaller, hopefully) probability of a problem
- $\text{Savings} = \text{Risk exposure} - \text{Cost of control} - \text{New risk exposure}$

If savings = 0, in many cases it might still be worth it to apply the control (risk aversion)

A concrete example

Example: Alice's company



Threat	Prob	Impact
Data Breach	1%	\$10 000
...

Control	Cost	Protects	P. Fail
Firewall	\$100	Data Breach	0%
...

Savings = Risk exposure – Cost of control – New risk exposure

Q: What are the savings here? Do we buy the firewall?

A: Risk exposure of “Data breach”: \$100

Cost of control: \$100

New risk exposure: $0.01 \cdot 0 \cdot \$10\,000 = \0

Savings: $\$100 - \$100 - \$0 = \0

We might still buy it, assuming **risk aversion**

Cybersecurity insurance

The “gamble” of not paying for the security control vs. paying for it is similar to having an insurance policy:

What do you prefer?

- Pay \$100 to have insurance,
- Not pay for insurance, but then we have a 1% probability of paying \$10 000?

Cybersecurity insurance

Cyber insurance products may cover the following first-party and post-breach expenses:

- Privacy attorney
- IT forensic investigation
- Compliance with state notification laws
- Credit monitoring for breached individuals
- Public relation firm to manage the crisis
- Regulatory fines
- Class action lawsuits resulting from the breach

Risk analysis practice

We identify the following risks in a company:

No.	Threat	Loss/Impact	Annual Prob. of Occurrence	Risk exposure (annual)
T1	Power loss	\$300 000	1%	
T2	Data theft	\$1 500 000	5%	
T3	Data breach	\$2 000 000	8%	
T4	Unauthorized access	\$200 000	4%	

We have these controls

No.	Control	Protects against	Buy cost	Annual cost	Annual Prob. of Failure (%)
C1	Firewalls	T2, T3	\$25 000	\$7 000	10%
C2	Security audit	T2, T3, T4	\$5 000	\$80 000	15%

Q: Compute the expected or average annual loss without protection for each threat — the risk exposure

A: Annual loss = Loss if it occurs × probability of occurrence, e.g.
 For T1: $300\,000 \cdot 0.01 = 3\,000$ dollars annual loss

Risk analysis practice

We identify the following risks in a company:

No.	Threat	Loss/Impact	Annual Prob. of Occurrence	Risk exposure (annual)
T1	Power loss	\$300 000	1%	\$3 000
T2	Data theft	\$1 500 000	5%	\$75 000
T3	Data breach	\$2 000 000	8%	\$160 000
T4	Unauthorized access	2\$00 000	4%	\$8 000

We have these controls

No.	Control	Protects against	Buy cost	Annual cost	Annual Prob. of Failure (%)
C1	Firewalls	T2, T3	\$25 000	\$7 000	10%
C2	Security audit	T2, T3, T4	\$5 000	\$80 000	15%

Q: If we buy C1, what are the total savings after n years?

$$\text{Total Savings} = [\text{Risk Exposure w/o control}] - [\text{Cost control}] - [\text{Risk Exposure w/ control}] = R - C - R'$$

$$\text{A: } R = (3\,000 + 75\,000 + 160\,000 + 8\,000) \cdot n = 246\,000 \cdot n \qquad C = 25\,000 + 7\,000 \cdot n$$

$$R' = (3\,000 + 75\,000 \cdot 0.1 + 160\,000 \cdot 0.1 + 8\,000) \cdot n = 34\,500 \cdot n$$

$$\text{Total Savings} = R - C - R' = 204\,500 \cdot n - 25\,000$$

Risk analysis practice

We identify the following risks in a company:

No.	Threat	Loss/Impact	Annual Prob. of Occurrence	Risk exposure (annual)
T1	Power loss	\$300 000	1%	\$3 000
T2	Data theft	\$1 500 000	5%	\$75 000
T3	Data breach	\$2 000 000	8%	\$160 000
T4	Unauthorized access	2\$00 000	4%	\$8 000

We have these controls

No.	Control	Protects against	Buy cost	Annual cost	Annual Prob. of Failure (%)
C1	Firewalls	T2, T3	\$25 000	\$7 000	10%
C2	Security audit	T2, T3, T4	\$5 000	\$80 000	15%

Other possible questions:

Q: Which of the two security controls is better after n years?

Which of the two security controls is better (in general)? (how would you answer this)

If savings were 0, is it worth purchasing the control?

...

Module outline

- 6 Security planning
- 7 Risk analysis
- 8 Closing remarks**

Physical security

All the firewalls in the world won't help you defend against an attacker who **physically** steals your laptop off your desk.

See databreaches.net for **many** examples of personal information being lost in incidents just like this.

We need to protect the physical machines, as well as the software and data on those machines.

Physical threats

There are two major classes of physical threats:

- Nature, e.g.:
 - Fire
 - Flood
 - Blackouts
- Human, e.g.:
 - Vandals
 - Thieves
 - Targeted attackers

Vandals

Some human attacks aren't actually after the data.

Example: Sir George Williams University (later Concordia University) “Computer Centre Incident” of 1969 — the largest student uprising in Canadian history



Q: How would you control this kind of threat?

Thieves and Targeted Attackers

Most thieves are after:

- Hardware
- Software
- Data

We could also have targeted attackers that are after something in particular.

There are many ways to secure the hardware:

- Guards
- Lockdown equipments
- Apple AirTag?
- ...

Putting it together

So now we know how to protect:

- Programs (M2)
- Operating Systems (M3)
- Networks (M4)
- Internet applications (M5)
- Databases (M6)
- Physical computers and data (M7)

To test if we've done this right: **red** and **blue** teaming.

- Red team: usually hired hackers that try to break into your system ("ethical hacking")



- Blue team: usually the company's security experts that try to defend



And we are doooooone!

