CS 458 / 658: Computer Security and Privacy

Module 7 - Non-technical Aspects of Security and Privacy

Part 1 - Ethics and legal issues

Spring 2022

Outline

Outline

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

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!

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.

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.

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.

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

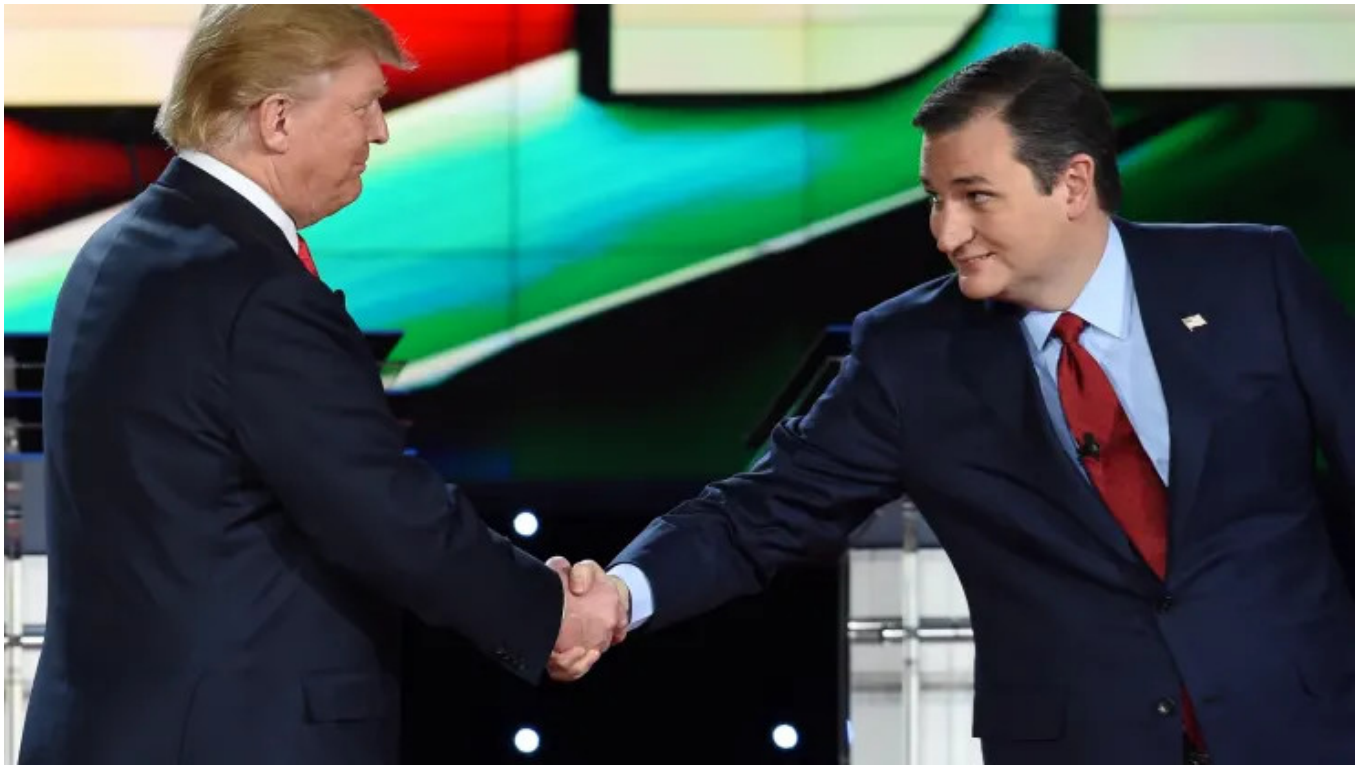

Donald Trump and Ted Cruz shake hands before the start of the Republican Presidential Debate (2015)

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).

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".

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

Linux kernel and the University of Minnesota

- You are an open-source enthusiast and make contributions to open-source projects regularly

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

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

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

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: How bad can this be?

Linux kernel and the University of Minnesota

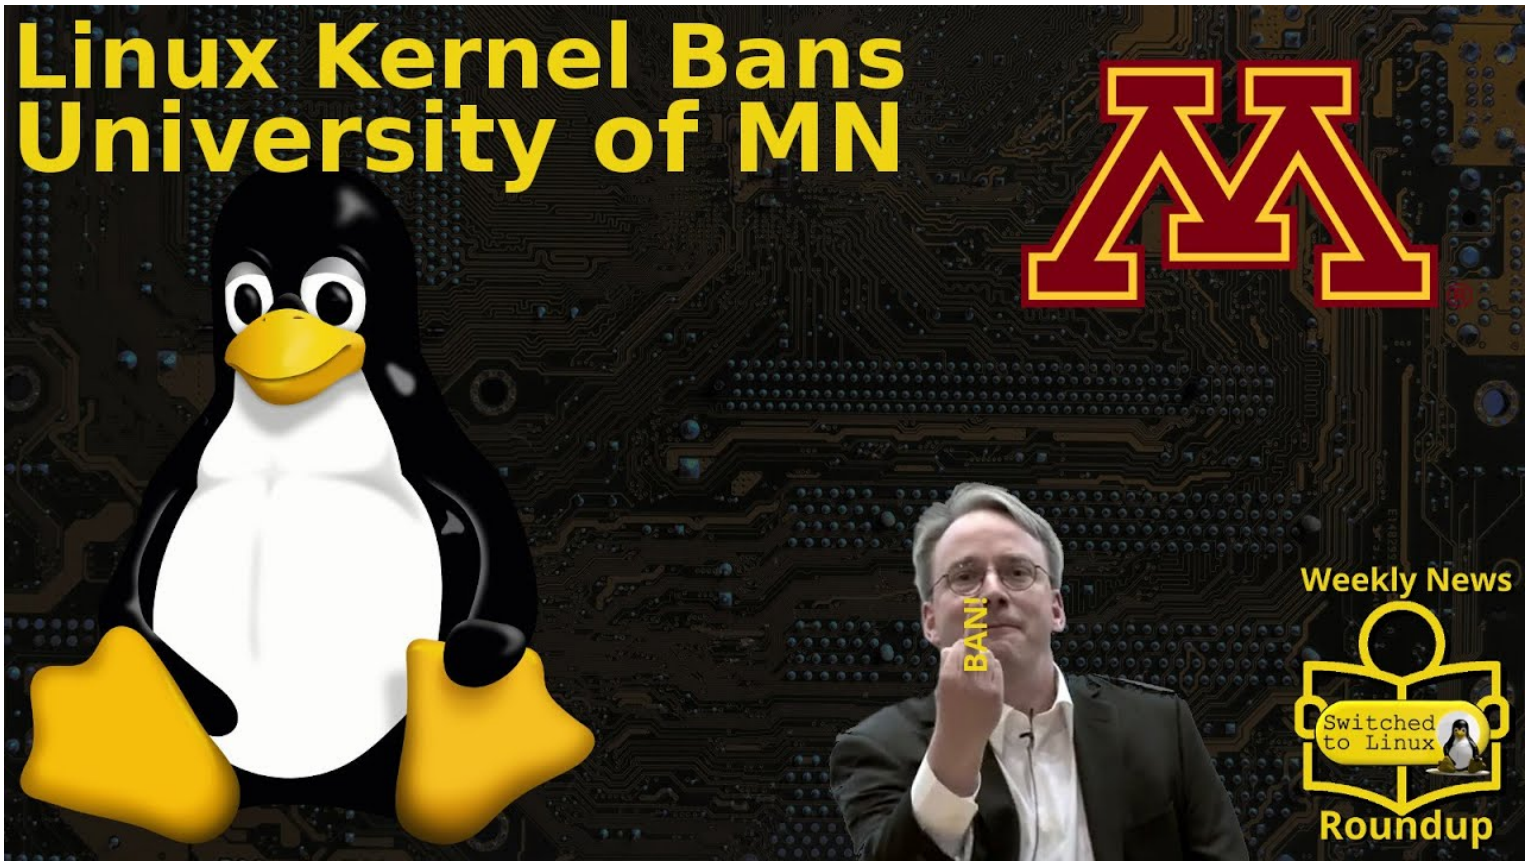

What we learned from the examples?

Be extremely cautious when human is involved in any form of activity, regardless of physical or virtual presence:

What we learned from the examples?

Be extremely cautious when human is involved in any form of activity, regardless of physical or virtual presence:

- In three words: Think before action

What we learned from the examples?

Be extremely cautious when human is involved in any form of activity, regardless of physical or virtual presence:

- In three words: Think before action
- In two words: Think twice

What we learned from the examples?

Be extremely cautious when human is involved in any form of activity, regardless of physical or virtual presence:

- In three words: Think before action
- In two words: Think twice
- In one word: Don't!

What we learned from the examples?

Be extremely cautious when human is involved in any form of activity, regardless of physical or virtual presence:

- In three words: Think before action
- In two words: Think twice
- In one word: Don't!

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

Outline

Laws, morality, and ethics

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

Laws, morality, and ethics

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

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

Laws, morality, and ethics

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

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?

What is law?

- Laws are a set of formal rules that governs how we behave as members of a society.

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.

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.

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?

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?

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

Lengthy legislative process

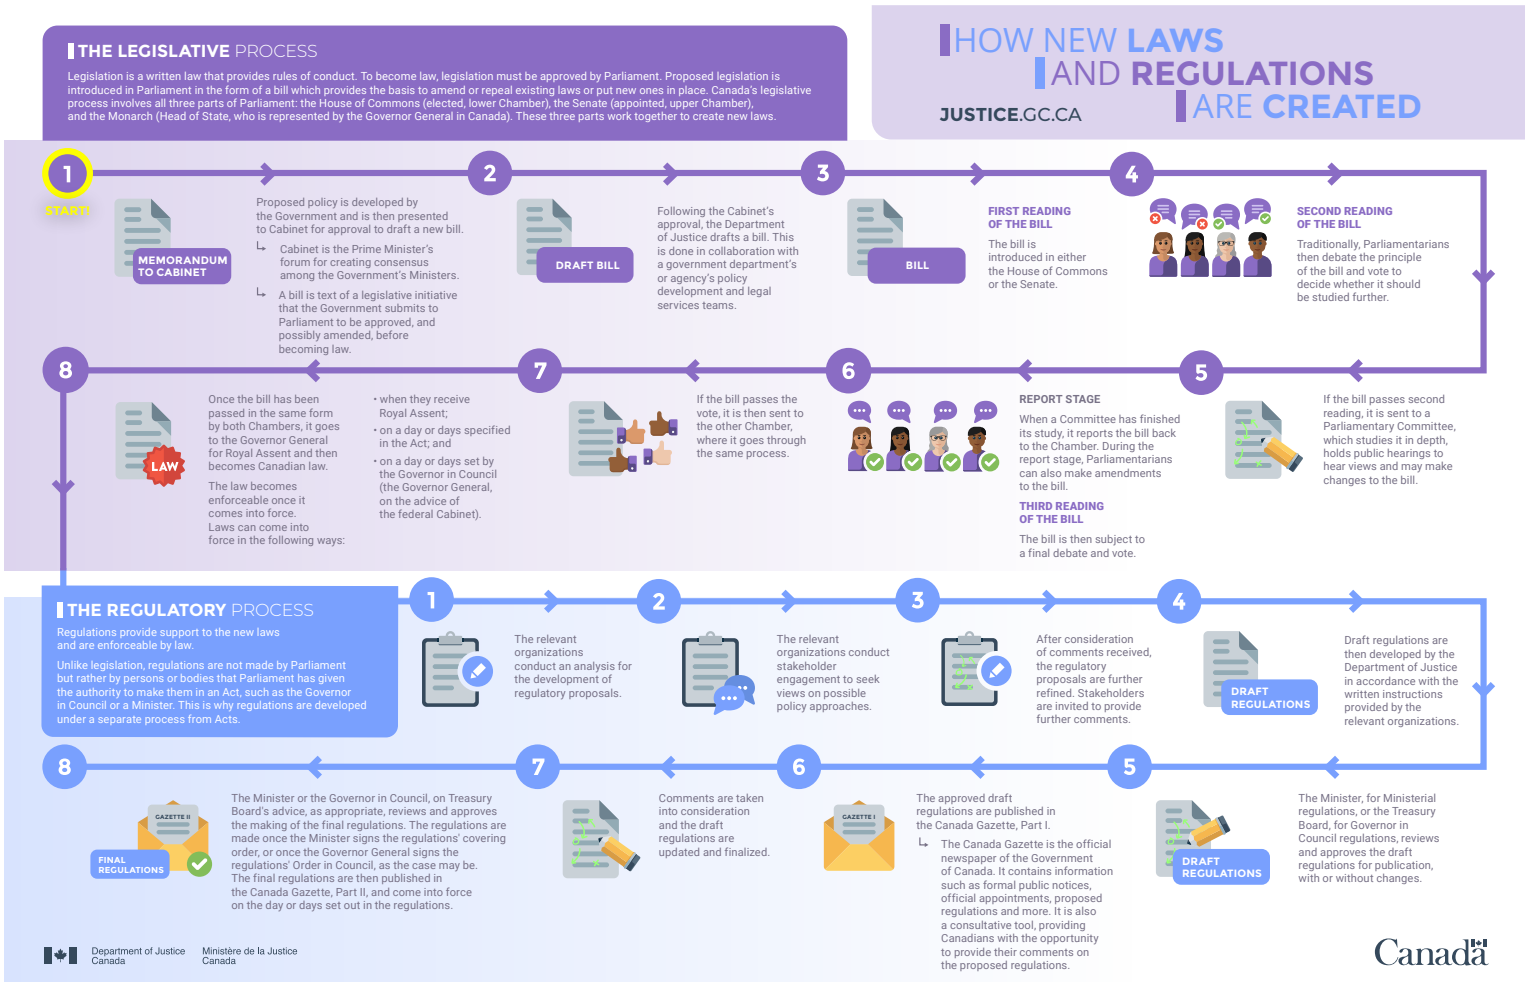

The legislative process in Canada

Non-violations of law

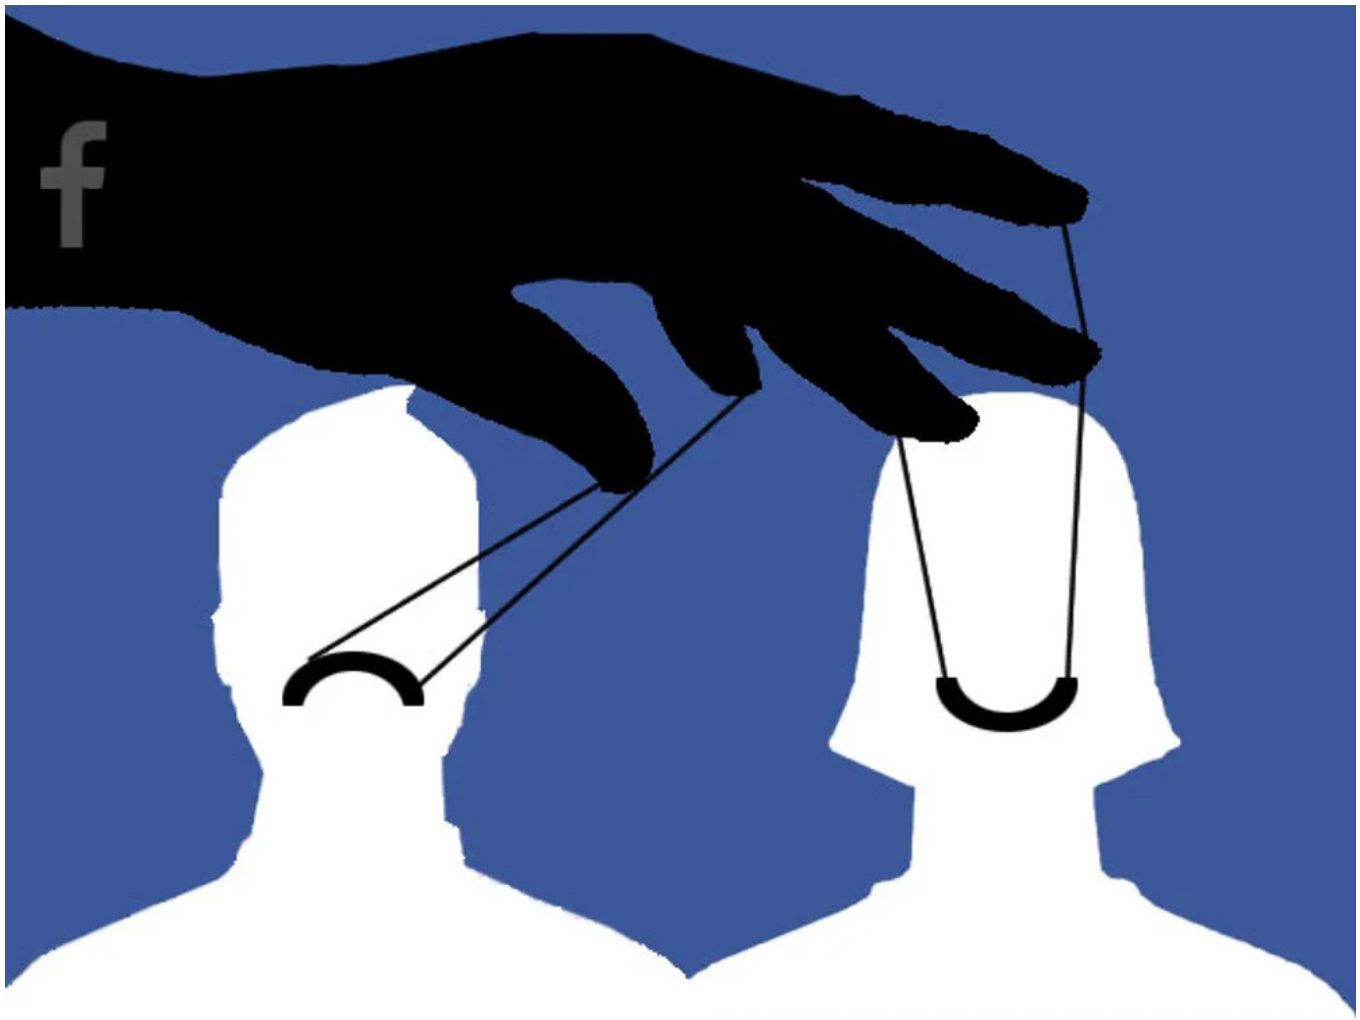

The (secret) mood manipulation study by Facebook in 2012

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.

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?

Right and wrong

Q: Why there are no legal violations here?

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.*

Right and wrong

Q: Why there are no legal violations here?

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?

What is morality?

- Morality refers to an informal framework of values, principles, beliefs, customs, ways of living.

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.

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.

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.

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.

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?

From morality to ethics

- There is rarely a moral authority agreed by every individual
- *"The unexamined life is not worth living"* — Socrates

What is ethics?

What is ethics?

- Ethics is a branch of philosophy that answers a basic question: what should I do? (out of all possibilities)

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.

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

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

In an ideal world, our ethical reflections shape the laws and moral systems a society will develop.

Outline

The science of building consensus on ethics

Ethics is about solving shared practical problems by building consensus through rigorous and logical argument.

The science of building consensus on ethics

Ethics is about solving shared practical problems by building consensus through rigorous and logical argument.

- Consequentialist theories
  - concerned with the ethical consequences of particular actions
- Duty-based theories
  - concerned with the intentions of the person making ethical decisions about particular actions
- Agent-centered theories
  - concerned with the overall ethical status of individuals and less concerned with particular actions

Consequentialist theories

The Utilitarian Approach

Consequentialist theories

The Utilitarian Approach

- "The best life is one that produces the least pain and distress"

Consequentialist theories

The Utilitarian Approach

- "The best life is one that produces the least pain and distress"
- Actions could be described as good or bad depending upon the amount and degree of happiness and/or pain they would produce

Consequentialist theories

The Utilitarian Approach

- "The best life is one that produces the least pain and distress"
- Actions could be described as good or bad depending upon the amount and degree of happiness and/or pain they would produce

Utilitarianism is one of the most common approaches to making and justifying ethical decisions, especially decisions with consequences that concern large groups of people.

Duty-based theories

The Duty-Based Approach a.k.a., deontological ethics

Duty-based theories

The Duty-Based Approach a.k.a., deontological ethics

- Doing what is right is not about the consequences of our actions (something over which we ultimately have no control) but about having the proper intention in performing the action

Duty-based theories

The Duty-Based Approach a.k.a., deontological ethics

- Doing what is right is not about the consequences of our actions (something over which we ultimately have no control) but about having the proper intention in performing the action
- The ethical action is one taken from duty, i.e., it is done precisely because it is our obligation to perform the action

Duty-based theories

The Duty-Based Approach a.k.a., deontological ethics

- Doing what is right is not about the consequences of our actions (something over which we ultimately have no control) but about having the proper intention in performing the action
- The ethical action is one taken from duty, i.e., it is done precisely because it is our obligation to perform the action
- Use categorical imperative to define ethical obligation: *"act only in accordance with that maxim through which you can at the same time will that it become a universal law"*

Duty-based theories

The Duty-Based Approach a.k.a., deontological ethics

- Doing what is right is not about the consequences of our actions (something over which we ultimately have no control) but about having the proper intention in performing the action
- The ethical action is one taken from duty, i.e., it is done precisely because it is our obligation to perform the action
- Use categorical imperative to define ethical obligation: *"act only in accordance with that maxim through which you can at the same time will that it become a universal law"*

Q: How does this approach apply to the mood manipulation study?

Agent-centered theories

The Virtue Approach

Agent-centered theories

The Virtue Approach

- Let's forget about the question *"what is a good action"*, instead, focus on *"what makes a good person"*

Agent-centered theories

The Virtue Approach

- Let's forget about the question *"what is a good action"*, instead, focus on *"what makes a good person"*
- In order to become a good person, one has to cultivate the right set of character traits (i.e., virtues)

Agent-centered theories

The Virtue Approach

- Let's forget about the question *"what is a good action"*, instead, focus on *"what makes a good person"*
- In order to become a good person, one has to cultivate the right set of character traits (i.e., virtues)
- Classic set of virtues include courage, wisdom, justice, honesty, etc — everything you would expect from a Hollywood movie

Virtue ethics is perhaps more useful in the long-term and less useful when applied to a specific context. Virtue ethics emphasizes the importance of role models and education to ethical behavior, which sometimes is merely reinforcing current cultural norms.

Ethical theories summary

	Consequentialist	**Duty-based**	**Agent-centered**
Deliberative process	What kind of outcomes should I produce (or try to produce)?	What are my obligations in this situation, and what are the things I should never do?	What kind of person should I be (or try to be), and what will my actions show about my character?
Ethical conduct	Ethical conduct is the action that will achieve the best consequences.	Ethical conduct involves always doing the right thing: never failing to do one's duty.	Ethical conduct is whatever a fully virtuous person would do in the circumstances.
Motivation	Aim is to produce the most good.	Aim is to perform the right action.	Aim is to develop one's character.

Outline

Responsible disclosure

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

Responsible disclosure

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

Responsible vulnerability disclosure

Responsible disclosure

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

Responsible vulnerability disclosure

- A private full disclosure to all responsible parties (e.g., software vendors for most software bugs)

Responsible disclosure

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

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)

Responsible disclosure

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

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

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

Talk to non-experts

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...

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...

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)

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: 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

Outline

Legal protections

Q: 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

Legal protections

Q: 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 defences, we can also use legal defences

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:

- Trade secrets,
- Trademarks,
- Patents, and
- Copyrights

Types of intellectual property

Four kinds of IP here:

- Trade secrets,
- Trademarks,
- Patents, and
- 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

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
- To get one, make a legal filing showing that you are using the name in commerce
  - This lets you sue others who use that name in a confusing manner
- Domain names are often protected under trademark law

Trademarks

- Trademarks protect names, brands, logos
- To get one, make a legal filing showing that you are using the name in commerce
  - This lets you sue others who use that name in a confusing manner
- Domain names are often protected under trademark law

Example: Microsoft vs MikeRoweSoft
Mike Rowe used to own the domain name "MikeRoweSoft.com"

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

Cryptography patents

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!
- 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

Legal copying

Even the rights granted to the copyright holder aren't absolute,
Anyone can copy a work without permission in certain circumstances

- In the US, these circumstances are broad, but loosely defined
- In Canada, there are very specific circumstances

Fair use in the U.S.

In the US, these exceptions are called fair use

- For purposes such as criticism, comment, news reporting, teaching, scholarship, or research
- Four tests:
  - the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes;
  - the nature of the copyrighted work;
  - the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and
  - the effect of the use upon the potential market for or value of the copyrighted work

Fair dealing in Canada

In Canada, the fair dealing exception to copyright law is defined more narrowly:

- It applies to private study, research, criticism, review, news reporting, education, parody, and satire
  - This is an exhaustive list!
- Time shifting, format shifting, backup copies, copying for private purposes, and mash-ups ("YouTube exception") are also legal

Paracopyright

In 1998, the US passed the Digital Millennium Copyright Act (DMCA):

- It didn't make any additional acts of making copies illegal; rather, it made illegal the circumvention of a technological copy protection mechanism that might be in place
- It also made illegal the manufacture, selling, or "traffic" of devices that might help you circumvent such mechanisms

Paracopyright in Canada

Canada, as of 2012, has restrictions similar to the DMCA. The rules are known as "digital lock rules" and they override fair dealing rights

- The rules prohibit circumventing a "technological measure" (e.g., encryption) to access data on DVDs, software, ebooks, etc.
- *Even if* the copyright in the underlying work has expired, or if the license on the work allows such use

There are very limited exceptions that apply, e.g., for law enforcement or encryption research.

Outline

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:

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?

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!

CS 458 / 658: Computer Security and Privacy

Module 7 - Non-technical Aspects of Security and Privacy

Part 2 - Administering security and privacy

Spring 2022

Outline

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.

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.

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

Red alert!

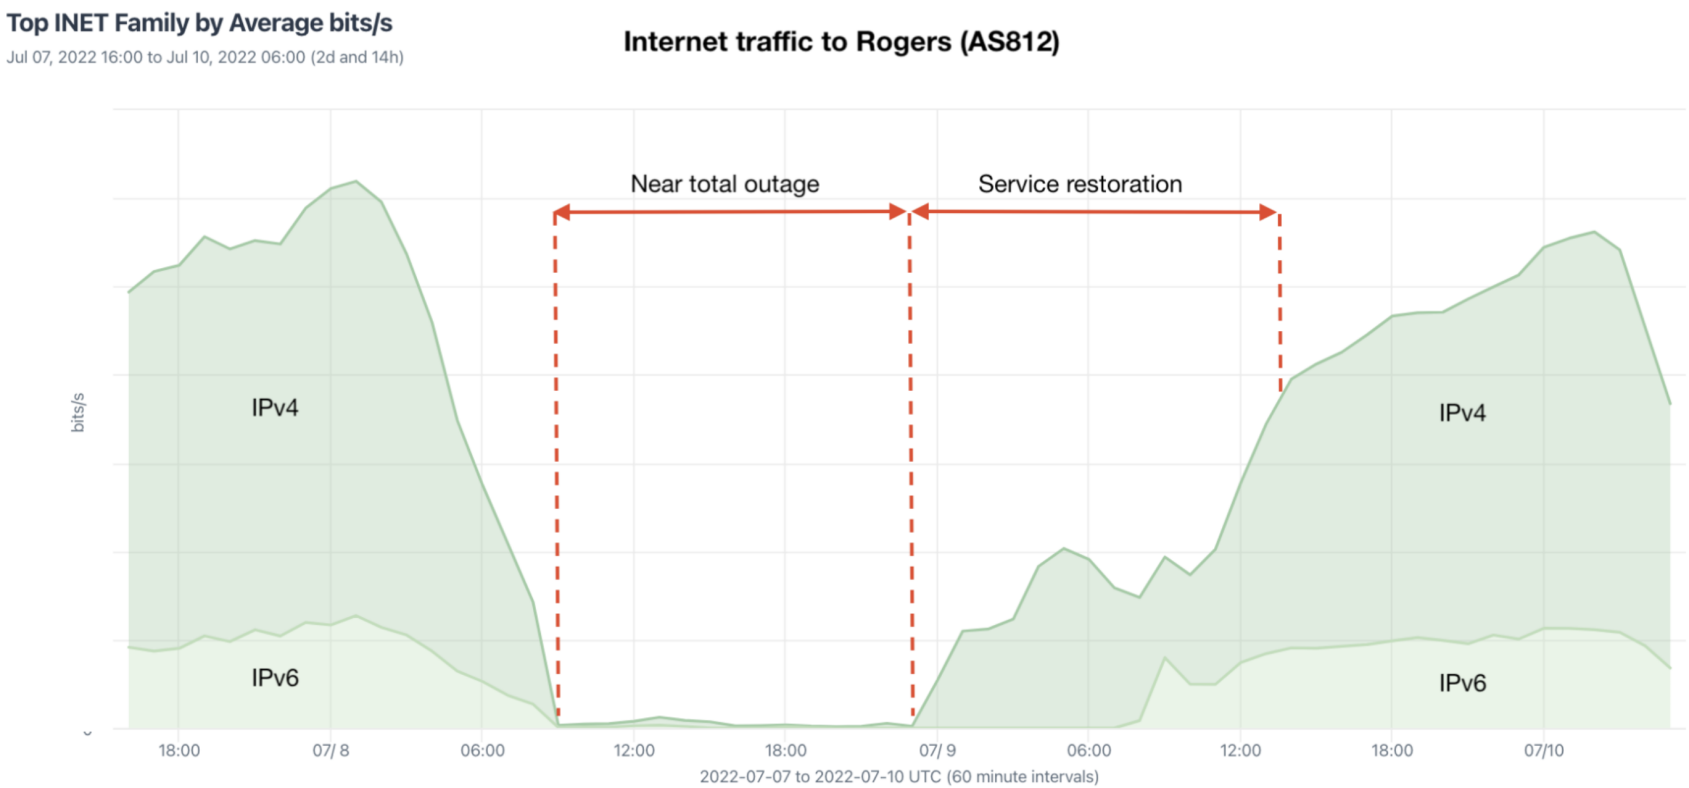

Q: What should you do as a Rogers' employee at 6:00 UTC?

Goals of security planning

A security plan is a document that explains

- what the security goals are
- how they are to be met
- how they'll stay met

Employees can use this document to inform their actions

Goals of security planning

A security plan is a document that explains

- what the security goals are
- how they are to be met
- how they'll stay met

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.

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:

- Policy
- Current state
- Requirements
- Recommended controls
- Accountability
- Timetable
- Continuing attention

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

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

3/ Requirements

What are the security and privacy needs of the organization

- 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.

3/ Requirements

What are the security and privacy needs of the organization

- 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

- to control issues identified in the "Current state" section,
- to satisfy the needs in the "Requirements" section, and
- taking into account the priorities in the "Policy" section.

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

- Program, OS, Network, Internet application, 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.

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?

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 plans

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 BCP

Outline

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 $\times$ 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

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?

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:

- Identify assets
- Determine vulnerabilities
- Estimate likelihood of exploitation
- Compute risk exposure
- Survey applicable controls
- Project savings due to control

1/ Identify assets

The main assets we would want to protect:

- Hardware
- Software
- Data

1/ Identify assets

The main assets we would want to protect:

- Hardware
- Software
- Data

What else?

1/ Identify assets

The main assets we would want to protect:

- Hardware
- Software
- Data

- Documentation
- Procedures
- Reputation

2/ Determine vulnerabilities

This step is where you apply the knowledge obtained in this course

- Also called threat modeling
- "Think like an attacker" and be very creative, even outlandish
- Come up with as many attacks on your own systems as you can, both technical and non-technical, against assets identified before
- Confidentiality, integrity, availability, privacy, etc.

3/ Estimate likelihood of exploitation

This is the hardest step, and there are experts trained in doing it — this is called *actuarial science*

- It's difficult to estimate the probability of each risk
  - Especially if it's so unlikely that it's never happened before
  - Otherwise, frequency analysis can be useful
- Take into account existing controls and their own effectiveness

3/ Estimate likelihood of exploitation

This is the hardest step, and there are experts trained in doing it — this is called *actuarial science*

- It's difficult to estimate the probability of each risk
  - Especially if it's so unlikely that it's never happened before
  - Otherwise, frequency analysis can be useful
- Take into account existing controls and their own effectiveness

Q: What is the chance that a buffer overflow bug can cause arbitrary code execution? With stack canaries? With ASLR?

4/ Compute risk exposure

Identify 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

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
  - Note that a control that protects against one vulnerability might make another one worse!
  - Also watch out for interactions among different controls

6/ Project savings due to control

- The expected cost of not controlling the risk is just the risk exposure, as computed earlier
- 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
- Savings = Risk exposure − Cost of control − New risk exposure

6/ Project savings due to control

- The expected cost of not controlling the risk is just the risk exposure, as computed earlier
- 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
- Savings = Risk exposure − Cost of control − New risk exposure

Q: If savings = 0, should we apply the control?

A concrete example

	No exploit	**Exploited**
Data breach (1% chance)	\$0	\$10,000
With control mechanisms	\$100	\$100

e.g., a firewall that completely removes the chance of exploitation

A concrete example

	No exploit	**Exploited**
Data breach (1% chance)	\$0	\$10,000
With control mechanisms	\$100	\$100

e.g., a firewall that completely removes the chance of exploitation

Q: What is the saving here?

A concrete example

	No exploit	**Exploited**
Data breach (1% chance)	\$0	\$10,000
With control mechanisms	\$100	\$100

e.g., a firewall that completely removes the chance of exploitation

Q: What is the saving here?
Savings = Risk exposure − Cost of control − New risk exposure
A: $10{,}000 \times 0.01 - 100 - 0 = 0$

A concrete example

	No exploit	**Exploited**
Data breach (1% chance)	\$0	\$10,000
With control mechanisms	\$100	\$100

e.g., a firewall that completely removes the chance of exploitation

Q: What is the saving here?
Savings = Risk exposure − Cost of control − New risk exposure
A: $10{,}000 \times 0.01 - 100 - 0 = 0$

Q: Do you want to use this control mechanism?

A concrete example

	No exploit	**Exploited**
Data breach (1% chance)	\$0	\$10,000
With control mechanisms	\$100	\$100

e.g., a firewall that completely removes the chance of exploitation

Q: What is the saving here?
Savings = Risk exposure − Cost of control − New risk exposure
A: $10{,}000 \times 0.01 - 100 - 0 = 0$

Q: Do you want to use this control mechanism?
A: Yes assuming risk aversion

A concrete example

	No exploit	**Exploited**
Data breach (1% chance)	\$0	\$10,000
With control mechanisms	\$100	\$100

e.g., a firewall that completely removes the chance of exploitation

Q: What is the saving here?
Savings = Risk exposure − Cost of control − New risk exposure
A: $10{,}000 \times 0.01 - 100 - 0 = 0$

Q: Do you want to use this control mechanism?
A: Yes assuming risk aversion

Q: What does this remind you?

Cybersecurity insurance

	No exploit	**Exploited**
Data breach (1% chance)	\$0	\$10,000
With insurance cost	\$100	\$100

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

Cybersecurity insurance

Frankly, I don't think we or anybody else really knows what they're doing when writing cyber. People who say they have a firm grasp on the risk are kidding themselves.

— Warren Buffet, 2018

Outline

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

Physical threats

There are two major classes of physical threats:

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

Q: What are the major differences in the security controls needed to protect against these two classes?

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

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

Q: What are most thieves after?

- Hardware?
- Software?
- Data?

Thieves

Q: What are most thieves after?

- Hardware?
- Software?
- Data?

Q: How do we secure hardware?
A: Guards, lockdown equipments, Apple AirTag?

Targeted attackers

What if the thieves are actually targeting you?

Now what are they most likely to be after?

- Hardware?
- Software?
- Data?

Putting it together

So now we know how to protect:

- Programs
- Operating Systems
- Networks
- Internet applications
- Databases
- Physical computers and data

Putting it together

So now we know how to protect:

- Programs
- Operating Systems
- Networks
- Internet applications
- Databases
- Physical computers and data

Q: How can we test if we've done it right?

Red and blue teaming

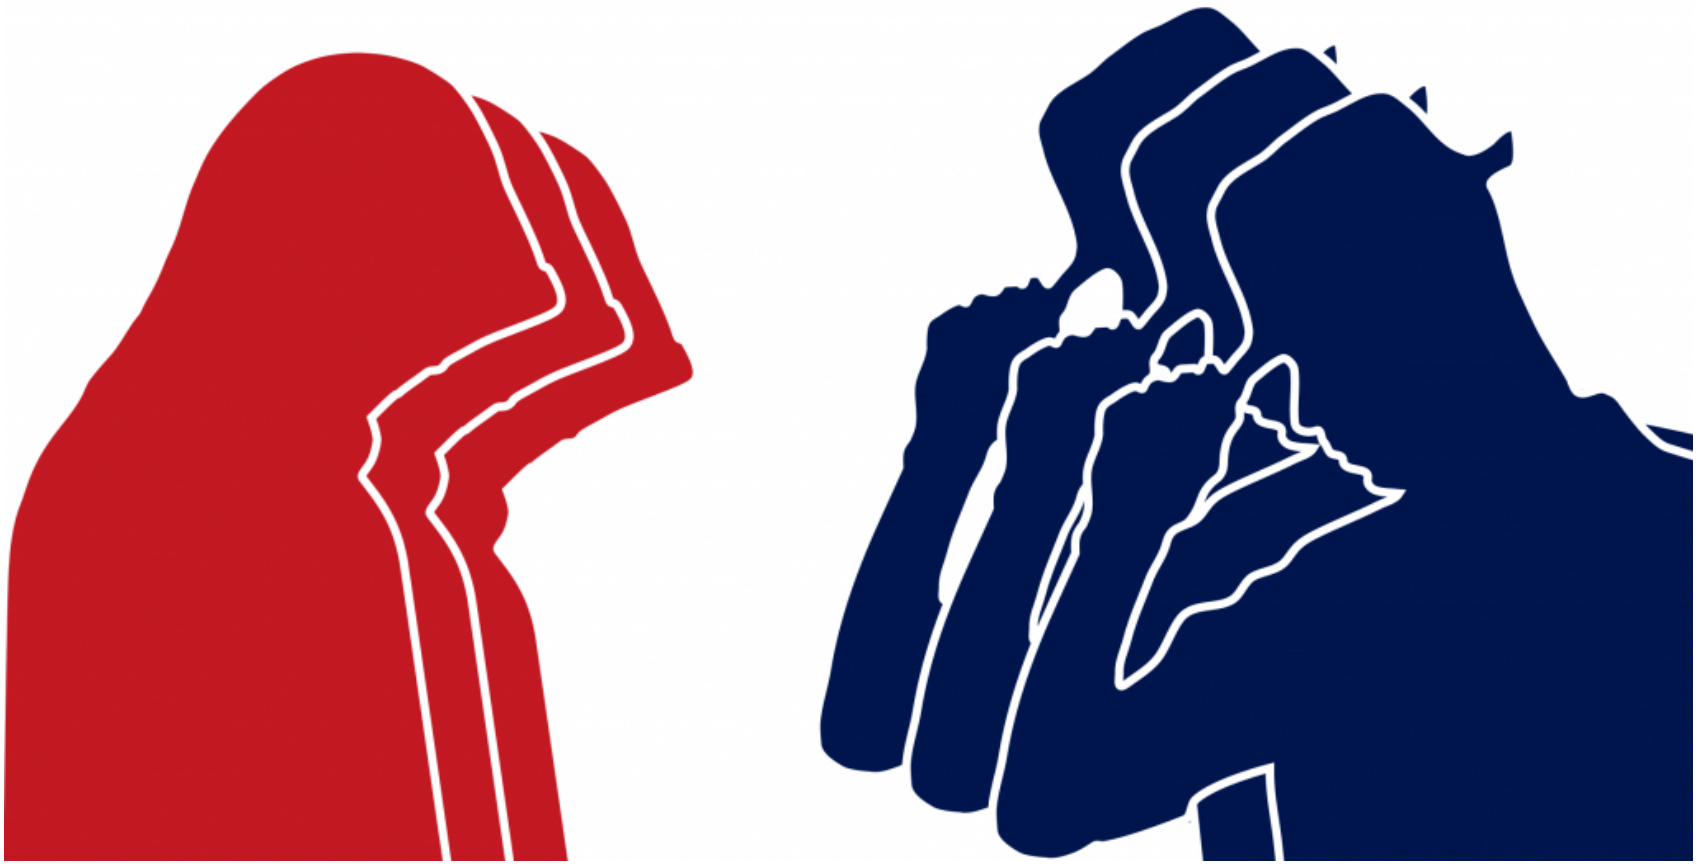

Concluding remarks

- Security is science
- Security is art
- Security is a mindset
- Security is a practice