

CS 458 / 658
Computer Security and Privacy

Module 7
Non-technical Aspects

Winter 2014

Module outline

- ① Administering security
- ② Legal and ethical issues

Module outline

- ① Administering security
- ② Legal and ethical issues

Administering security

- So far in this course, we've talked about a lot of things you can do **technically** to protect your programs, operating systems, networks, databases, and Internet applications
- But there's more to security and privacy than just these technical solutions
- Next, we will look at several **non-technical** aspects of administering security

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
 - What do they do today, now that the files are all electronic?
- Many users do not appreciate the security and privacy risks in using computers
- A **security plan** is a document put together by an organization that explains what the security goals are, how they are to be met, and how they'll **stay met**
 - Employees can use this document to inform their actions

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
- It has seven parts, which we will look at in turn:
 - Policy
 - Current state
 - Requirements
 - Recommended controls
 - Accountability
 - Timetable
 - Continuing attention

Policy

- 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's?
A security manager? A security group in IT?
 - Commitment
 - Institutionally, who provides security support for staff?
Where does security fit into the org chart?

Current state

- The security plan should contain a risk analysis (see later) describing the current status of the system
 - What assets and controls are there? What might go wrong? What vulnerabilities are currently exposed?
- What should you do if new assets are added or new vulnerabilities are discovered?
- List the limits of security responsibility
 - Who is responsible for the security of the Internet uplink router to the company's ISP?
- How is people's privacy affected? Perform a Privacy Impact Assessment (PIA)

Requirements

- What needs does the organization have?
 - **Who** is allowed/not allowed to do **what**?
 - What audit logs should be kept?
 - Do you need to be able to measure the ongoing effectiveness of the security controls?
- **Not** anything to do with **mechanism**
 - The policy statement doesn't say anything about **how** to accomplish the listed goals
 - It should be technology-neutral
 - For example, it 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.

Recommended controls

- Here's where you list mechanisms to control vulnerabilities identified in the "Current state" section, to satisfy the needs in the "Requirements" section, taking into account the priorities in the "Policy" section.
- They may be any of the security controls we've talked about in this course, or other similar ones
 - Program, OS, Network, Internet application, Database, etc.

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?
- Probably different people will be accountable for different pieces of the plan

Timetable

- Any reasonably sized security 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
 - What order, noting dependencies
- Include milestones to track progress along the way

Continuing attention

- 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 writes 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

Business continuity plans

- The Business Continuity Plan (BCP) is another kind of security plan
 - Focus is on Availability
- What will your organization do if it encounters a situation that is:
 - 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

Catastrophic failures

- Some examples of such failures:
 - Fire / earthquake destroys your data centre
 - A utility (phone, network, electricity, etc.) fails or goes out of business
 - Flood prevents operations staff from being able to reach your offices
 - Pandemic outbreak of avian flu keeps 1/3 of your staff home sick
 - See IST's pandemic plan (listed as a reading)
- What do you do?
 - Consult your **business continuity plan**

Don't blame “the computer”

- If your business can't go on because some computer isn't working right, that's **not** the computer's fault; it's **yours**, for not having a backup contingency
 - Some (physical) stores can't sell you goods if their computers are down
 - Better stores have a fallback procedure where they keep track of sales on paper until the computer comes back up and the accounts can be reconciled

Advance planning

- You need to write an actual plan, which should include things like:
 - Who is in charge when a catastrophe occurs
 - This person will also be the one to declare when the emergency is over and things can get back to normal
 - See uWaterloo's Emergency Response policy (listed as a reading)
 - What needs to be done
 - To deal with **keeping the business going**, not with dealing with the emergency itself; someone else will do things like call the fire department
 - Who will do it

Advance planning

- But writing the plan isn't 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

Incident response plans

- You notice that your company's home page has been defaced
- What do you do?
- Follow your company's **incident response plan**
 - "Incident" in this case refers to a security breach

Incident response plans

- The incident response plan needs to consider a number of things
 - Legal issues
 - The incident has legal ramifications. Under what circumstances should law enforcement get involved?
 - Preserving evidence
 - How can you quickly recover from the incident while maintaining as much **forensic evidence** as possible?
 - Records
 - Keep careful track of everything you do once you notice the breach
 - Public Relations
 - Speak with one voice

After the incident

- Once you have recovered from the incident, hold a review to ask:
- Is any security control action to be taken?
 - How did the breach occur? Have you patched that particular hole? Have you established procedures so that other similar problems are less likely to happen in the future? Was lack of user training an issue?
- Did the incident response plan work?
 - Did everyone know whom to notify? Did the response team have the needed resources? Was the response fast enough? What should be done differently next time?

Risk

- 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** x **impact**
- Note that both probability and impact of a given risk will change over time, so continual review is needed

Risk analysis

- It is impossible to completely eliminate risk
 - No system is absolutely secure
 - Even in your daily life, there are risks all around you
 - Crossing the street?
 - 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?
- As you can see, 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

Risk analysis

- In our setting, a risk analysis usually comprises the following steps:
 - Identify assets
 - Determine vulnerabilities
 - Estimate likelihood of exploitation
 - Compute expected loss
 - Survey applicable controls
 - Project savings due to control

Identify assets

- Way back in lecture 1, we identified three main assets we would want to protect:
 - Hardware, software, data
- Here, we add three more
 - People
 - Skills to run the system, network, or specific programs
 - Documentation
 - On hardware and software, but also the security plan, business continuity plan, and incident response plan
 - Supplies
 - Paper, forms, printer toner, etc. that play a supporting role

Determine vulnerabilities

- This step is where you best apply the knowledge obtained in this course
- “Think like an attacker” and be very creative
 - Even outlandish; this part can be fun!
- Come up with as many attacks on your own systems as you can, both technical and non-technical, against assets in each of the six categories
 - Confidentiality, integrity, availability
 - Don't forget privacy issues as well

Estimate likelihood of exploitation

- This is the hardest step, and there are experts trained in doing it
- 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
 - How often has this risk been a problem in the past?
 - Distinguish something that might happen once a year from something that might happen once a month
- Take into account existing controls and their own probabilities of failure

Compute expected loss

- Identify the impact of the risk
- Also a tricky step (even though estimates are usually good enough)
- Some examples:
 - 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
 - Different from value of data to competitor
 - Cost of delaying or outsourcing data processing if your systems are unavailable

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

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 (for example, 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
 - Hopefully, this is positive

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

Physical threats

- There are two major classes of physical threats:
 - Nature, e.g.:
 - Fire
 - Flood
 - Blackouts
 - Human, e.g.:
 - Vandals
 - Thieves
 - Targeted attackers

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

Physical controls against humans

- Earlier, we looked at being able to recover from natural disasters
 - Many of these techniques will also be useful against thefts, etc.
- Now, we will discuss what additional measures are necessary to protect against humans
 - Need to not only recover from the loss, but also deal with the release of potentially sensitive data

Vandals

- Some human attacks aren't actually after the data
- Sir George Williams (later Concordia U)
"Computer Centre Incident" of 1969 — the largest student uprising in Canadian history



- How would you control this kind of threat?

Thieves

- Most thefts are after what?
 - Hardware?
 - Software?
 - Data?
- We've already talked about controls against theft of software and data
- What about hardware?

Targeted attackers

- What if the thieves are actually targeting you?
- Now what are they most likely to be after?
 - Hardware?
 - Software?
 - Data?

Protecting offline data

- We have a good sense of how to protect data on an active machine hooked up to a network
- What about data sitting on a shelf?
 - Backup tapes / disks
 - Printouts / reports
- What happens after they're on the shelf?
- Why is offline data like this attractive to attackers?

Protecting offline data


- It's obviously harder for a network-based attacker to get at that kind of data
- But what about a physical attacker?
 - Thief
 - Insider
- How do you safely dispose of data?
 - Paper
 - Magnetic media
 - Optical media

Putting it together

- So now we know how to protect:
 - Programs
 - Operating Systems
 - Networks
 - Internet applications
 - Databases
 - Physical computers and data
- How can we test if we've done it right?

Tiger teams

- Tiger teams are teams of security professionals
- You can hire them to try to break into your site, systems, networks, etc.
 - And tell you what's wrong



Legal protections

- Remember this from module 1:
- 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

Legal protections

- The most obvious legal protections are against threats to hardware
- If someone steals a laptop, it's completely straightforward that he can be charged with a crime
- What if someone copies the laptop's hard disk, but leaves the laptop where it is?
- This is much newer law, and is often less clear
 - Caveat: IANAL; this course does not constitute formal legal advice. :-)

Module outline

- ① Administering security
- ② Legal and ethical issues

Overview of IP

- 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
- Four kinds of IP concern us:
 - Trade secrets, trademarks, patents, and copyrights

Overview of IP

- These four kinds of IP:
 - Cover different kinds of intangibles
 - Convey different rights
 - Have different durations
 - Have different registration requirements
 - (But are nonetheless often confused for each other!)
- Note: IP law is similar, but not identical, in Canada and the US; we will make note of the most important differences

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, you've 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, but it was reverse engineered in 1994

Trademarks

- Even though the RC4 algorithm was no longer protected, its **name** was!
- 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

Patents

- Applies to **inventions**, 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
- Notably:
 - 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 example of each type of crypto

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
 - It's sometimes not obvious when they apply
- In Canada, there are very specific circumstances (details on later slides)

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!
- In addition, there is a similar set of tests as in the US.
- Time shifting, format shifting, backup copies, copying for private purposes, and mash-ups (“YouTube exception”) are also legal

Private copying of sound recordings

- Private copying as mentioned on the previous slide explicitly does not cover the private copying of sound recordings:
 - However, you are allowed to copy a sound recording “onto an audio recording medium for the private use of the person who makes the copy”
 - In exchange, everyone pays a levy (about 21 cents) on blank audio recording media like tapes and blank CDs
- Some people argue this makes the downloading of songs over a P2P network legal in Canada
 - But uploading still probably isn't!

July 2012 Supreme Court Decisions

- The Supreme Court of Canada determined that fair dealing provisions should be interpreted in a “broad and liberal manner” and affirmed a “technology-neutral” approach to fair dealing
- It specifically recognized that copying materials for teaching purposes was a legitimate use of fair dealing principles
- This is a major blow to Access Copyright, which was an expensive licencing regime to allow copying at universities (however uWaterloo had already withdrawn from Access Copyright in August 2011)

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
- Problem: this applies even when the copy protection mechanism is broken to make a "fair use" copy!
- 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, with even fewer exceptions
- 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 licence on the work allows such use
- There are very limited exceptions that apply, e.g., for law enforcement, encryption research, to unlock a cellphone, or to access content if the person has a perceptual disability.

Paracopyright in Canada (cont.)

- However, violating the digital lock rules does not carry significant penalties for individuals in non-commercial infringement
- Similar rules apply to unauthorized downloading, file sharing, etc.
- It is not an infringement to possess software that can circumvent digital locks.
- Unauthorized downloading will likely be dealt with through a “notice-and-notice” system, whereby rights holders send notices of infringements to ISPs, who then forward notices to subscribers.
- Currently unclear whether an ISP is required to disclose subscriber information to copyright holder

Computer crime

- We saw that laws regarding intellectual property differ from those about real property
- Similarly, laws about unauthorized access of 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

Computer 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 the manual for
- Things seem to have settled down somewhat
 - At least as long as (para)copyright doesn't get involved
- 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
 - “Chain of custody”
- 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?
 - And there’s usually no physical evidence at all to speak of!
 - **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-30 (“Protecting Children from Internet Predators Act”)

- Proposed in February 2012 to give more power to law enforcement for investigating computer crimes
- Would have required ISPs to perform deep packet inspection (DPI) on customers' traffic and to hand over customer information without a warrant
- Would have given the government the authority to install its own DPI hardware within private ISPs
- Bill was abandoned in February 2013 after wide criticism

Scope of “Lawful Access”

- Who will be covered?
 - Phone companies
 - ISPs
 - Web server operators?
 - Wireless Nomad?
 - Coffeeshops with free Wi-fi?
 - Anyone who has Wi-fi at home?
 - Will you have to have a Wi-fi router with a “back door”?
 - Operators of Tor nodes?

Dangers of building in “back doors”

- From a technical point of view, adding the ability for anyone to surreptitiously intercept communications is designing in a **weakness**
- These dangers have been known since the Clipper Chip days, back before the turn of the century
- Remember that the technology can't tell whether any oversight requirements are met

Dangers of building in “back doors”

- If there's a back door built into the system, the bad guys will find a way to use it
 - Technical means (hacking in)
 - Social engineering
- This isn't just a hypothetical concern!

Greek wiretapping scandal

- Many senior Greek government officials had copies of their calls routed to a bank of throwaway cell phones
- Unknown people used CALEA interception capabilities built into phone switches to comply with US law
 - But this was in Greece!

Harmonization

- Industry prefers to only build one thing, not one thing for each jurisdiction
 - “Any new technical requirements must be based on international standards” (Information Technology Association of Canada submission to Customer Name and Address consultation)
- So we tend to end up with the biggest “back doors” required anywhere

Regulation of Investigatory Powers

- Example of a big back door required elsewhere:
- In the UK, Part III of the RIP Act went into effect in October 2007
- You can be served notice to:
 - Decrypt data
 - Hand over decryption keys
 - Don't tell that this has happened

Redress for software failures

- If flaws are discovered in most products you buy, you can get a new one (with the flaw repaired), or at least a refund
 - Not so with software
- Why is that?
- Note that **embedded software** usually doesn't have this problem: flaws in embedded software (in things like cars, for example) are usually fixed by the manufacturers

Reporting flaws and failures

- What should you do if you discover a flaw or failure in a software product?
 - Especially a security flaw
- Vendors prefer that you tell them, and no one else
 - And then they can tell no one else, and the problem is solved?
 - Some vendors will even back up this preference by suing you (or having you arrested!) if you publicly disclose a security flaw in their products

Full disclosure

- Some people (but not usually vendors) prefer **full disclosure**
 - When you find a problem, post it to a full disclosure mailing list of security professionals (like Bugtraq)
 - The reasoning is that by the time you (the good guys) have found the problem, the bad guys probably have as well, and may be actively exploiting it
 - You need to plug the hole as quickly as you can, until the vendor comes up with an official fix
 - Further, without disclosure, vendors sometimes have little incentive to fix the problem at all

Responsible disclosure

- Vendors countered with **responsible disclosure**:
 - If you find a security flaw, tell the vendor
 - Tell no one else for at least 30 days
 - If the vendor hasn't announced the flaw, with credit to you, and hopefully with a fix, in that 30 days, you should contact a **coordinating centre** like CERT to decide what to do next
- There is ongoing debate (see link on LEARN) as to which way is best
 - Best for whom?

Codes of professional ethics

- As a computer security professional (or even not specifically in security), you will be expected to uphold certain ethical standards
 - Note: ethics != law
- 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**
 - Linked to on LEARN

Example: CIPS

- Most professional codes of ethics have similar flavours, with some difference in detail
- 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

Recap

- Administering security
- Legal and ethical issues