

CS 453 / 698: Software and Systems Security

Module: Non-technical Aspects of Security

Lecture: administering security

Meng Xu (*University of Waterloo*)

Fall 2024

Outline


- 1 Security planning
- 2 Risk analysis
- 3 Closing remarks

Security planning



Fire!

Top OTT Service by Average bits/s | **Internet Traffic served by Facebook**
 Oct 04, 2021 06:00 to Oct 05, 2021 00:00 (18h) | **Global outage 4-Oct-2021**



silent night, holy night, all is calm, all is bright ...

Q: What should you do as Facebook employee at 16:00 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: high-level goals and priorities
- Current state: risk analysis, anticipation of new situations
- Requirements: **what** are the security and privacy needs
- Recommended controls: **how** to provide those needs
- Accountability: **who** is responsible for what
- Timetable: **when** the elements of the plan will be performed
- Continuing attention: how often the plan should be updated

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

- 1 Security planning
- 2 Risk analysis
- 3 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

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?
- 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

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

A concrete example

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

Q: What is the saving here?

A concrete example

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

Q: What is the saving here?

A: $10,000 \times 0.01 - 100 = 0$

A concrete example

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

Q: What is the saving here?

A: $10,000 \times 0.01 - 100 = 0$

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

Q: What is the saving here?

A: $10,000 \times 0.01 - 100 = 0$

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

Q: What is the saving here?

A: $10,000 \times 0.01 - 100 = 0$

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

- 1 Security planning
- 2 Risk analysis
- 3 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



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 thefts after?

- Hardware?
- Software?
- Data?

Thieves

Q: What are most thefts after?

- Hardware?
- Software?
- Data?

Q: How do we secure a 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?

〈 End 〉