The University of Sydney
School of Computer Science
Dr Suranga Seneviratne
Senior Lecturer - Security

THE UNIVERSITY OF SYDNEY

INFO3616/CSEC3616/CSEC5616 - Tutorial - 1

1 Programming Satan's Computer

Ross Anderson and Roger Needham infamously described security engineering as "Programming Satan's computer". Explain the metaphor: why, in particular, is it "Satan's computer"? 1-2 sentences are sufficient.

The metaphor refers to a computer that acts intentionally maliciously, even to the degree that it reacts to the programmers' behaviour and tries to find the best moment to thwart them. As a result, Security Engineering must tolerate the presence of an intelligent and hostile opponent.

2 Safety/Reliability Engineering vs. Security Engineering

Explain the difference between Safety and Reliability Engineerings and Security Engineering.

Reliability Engineering or Safety Engineering try to address problems that we know at some point are going to happen. E.g. Equipment failures, disk failures, power failures. For some of these we can go even can calculate failure probabilities. E.g. What is the probability of a disk failure after a certain years of operation? Many products release this information. Due to the predictable nature and the fact that we are expecting them, makes planning for such issues much easier. On the other hand, it is difficult to predict and mitigate the threats coming from an active opponent which is handled by Security Engineering.

3 Security Goals - A

You may wish to revisit the lecture recording to answer the following questions.

a) True or false? Explain your answer in 1-2 sentences

- If you have the choice between secrecy and availability, and your system can only meet one of these goals simultaneously, always go for secrecy.
- Authentication is the process of obtaining permission to access a file.
- Privacy (in the context of internet) is my ability to conduct my online activities whilst being anonymous.

- **False** - There are situations where availability is more important than secrecy. An example can be the military in a particular combat situation: it may matter more that commanders can communicate at all than the communication being accessible to eavesdroppers.
- **False** - Authentication is about establishing the correctness of an identity. The process of obtaining permission to access to a resource is authorisation.
- **False** - Privacy is my ability to control what information I might share about me with other parties. What is described is anonymity, i.e. I want others to see and know what I am doing, I just don't want them know it was me.

b) Explain in 1-2 sentences.

- Why does non-repudiation require secure timestamping in many cases?
- Why is integrity of logfiles important to achieve accountability?
- Why is it very hard to estimate the complete attack surface of a given system?
- In the context of the last question, why is it important to remind yourself that security is a process?

- Because it is possible for someone to claim that their computer system had been compromised at some point and it was the attacker who carried out some activity later. The secure timestamping by a third party (e.g., TSA - Time Stamping Authority) provides a record what happened at which time, and hence makes such claims much harder.
- An attacker with access to a system can modify the log files to hide their traces. Hence, log files must be protected to ensure they provide a record of attacker activity (it is not always possible, however).
- This would require to correctly enumerate every possible way how a system can be breached. In practice, that is usually not possible: today's systems are too complex and one may easily overlook an attack vector.
- What is not possibly today, may be possible tomorrow: the environment changes, and hence security is not static, but a process.

c) Break a protocol

Danaerys tells Jon how he can enter the city. To enter the city, the guards will question him 'Who goes there?' And Jon must reply with the secret passphrase ('The Lord of Winterfell').
Give a reason why this does not meet the definition of a cryptographic protocol and show how it can be broken.

This method provides no freshness. Someone could overhear the password being spoken and use it to enter (i.e. replay attacks).

4 Security Goals-B

Someone tells you that you can "just encrypt data" and you can be sure that the following security goals are achieved: confidentiality, integrity, authenticity, and access control. For each goal, explain if the person is right or wrong.

The easiest way to solve this is to start with the last three goals.

- **Confidentiality: Right** - Once the data is encrypted confidentially is guaranteed. Without access
- **Integrity: Wrong** The attacker can make changes to the encrypted data, and it may still seem 'plausible enough' when decrypted.

- **Authenticity: Wrong** The attacker can make changes to the data, and there is no way to tell if they are authentic—see integrity. Note that authenticity is a special case of integrity which indicates that the data is exactly as it is when it was under the control of the originator. If a solution does not guarantee integrity it does not guarantee authenticity. Solutions that may guarantee integrity may still not provide authenticity.
- **Access Control: Can be argued either way** This security goal overlaps with others (or, alternatively stated, is more high-level). We tend to treat it as a security goal in its own right because it is all-important, even though it has many facets.

  Right We could argue that encryption controls access to the plain text data.

  Wrong We can equally say that encryption that there is nothing about encryption that will check an attacker's authorization to modify or delete the data. So, encryption provides only one aspect of access control.

5 Sydney Opal Card

You are likely familiar with the Sydney Opal Card system: a swipe card is used to 'pay' and allow you to enter and leave the platforms via gates. Let us assume you were in charge when the system was planned. Using our reference framework of Policy, Incentives, Mechanism, and Assurance, please:

- State two policies to achieve. One sentence per policy is fine.
- State two incentives that work against those policies.
- State one attack based on each incentive (see note below). It's fine if the attack is described fairly generically (no need to refer to concrete technology that it attacks).
- Give two mechanisms you can think of that will work against the attack based on the incentive.
- State what assurance you place on each mechanism you have chosen.

This question is intentionally open: you are free to choose policy and incentives, and you are free to make assumptions how the transport system works. You are free to choose the attack based on the incentive (the one that your mechanism protects against). Policies, incentives, mechanism, and assurance do not have to be based on the electronic part of the system—it can also be entirely physical.

There are many possible answers. The following is an example.

- A high-level policy could be: *Customers must be billed correctly for their travel.* If desired, that can be broken down into security goals relating to determining the correct amount to pay (integrity, accountability, authenticity etc.), but that was not necessary here as an answer (but a good exercise if you did it).

  Another policy could be: *The billing system must be available whenever a customer wants to travel.*

  or

  Users' travel data must be confidential and must be accessible only to authorised NSW Transport staff

- Incentives could be:
  - Attacker wants to free-ride.
  - Attacker wants to disturb the train system in Sydney.
  - Attacker wants to access some target users' travel information
- Attacks could be:
  - Attacker forges an Opal card in someone else's name, which is also linked to that person's credit card. Alternatively: attacker breaks into online account and resets customer debt to 0.
  - Attacker stages a massive Denial-of-Service attack against the data centre that houses the accounting.
  - Attacker tries a phishing campaign to target some of the NSW transport employees hoping to harvest credentials for the customer database.
- Mechanisms could be:
  - Forgery is difficult to defend against in an absolute sense. One needs Opal cards with chips that are hard to tamper with and, ideally, have a cryptographic mechanism that makes the forgery detectable. History has shown that such systems are still hackable: it is a matter of effort. If the effort outweighs the benefit, the attacker is unlikely to try it. Note how easily this balance is shaken if, e.g., the attacker finds a way to forge many cards **and** sell them! Defending against an attacker trying to break in means to have intrusion detection in place, log suspicious events, good passwords etc.—really almost everything we will cover throughout the course.
  - Denial-of-Service attacks are extremely hard to defend against: the natural asymmetry of Internet-based services works against anything one can come up with. It is possible to make them less effective, however: for example, by having massive redundancy in place that can be spun up when needed (companies exist that provide exactly this security service!). One can alternatively disconnect the Internet interface from the internal Opal accounting, e.g. by having the internal accounting be able to run autonomously for a few days without synchronising with the Internet-facing systems or checking if Opal cards are charged. That would give the defenders time while keeping the system available. The drawback may be that since it is not possible to check if Opal Cards are still charged, one may induce a certain potential loss from people who take the opportunity to run up a debt and then just toss the Opal card. It may still be preferable to a wider outage, however.
  - Phishing is a common attack vector. We will discuss it more during our Usability lecture. Mechanisms to defend against can be deploying an email security solution that can flag suspicious emails or conducting user training and awareness programs so that the employees can identify phishing emails. Again and again what we see from various security incidents is that these two approaches are not perfect. Many high profile network breaches such as iCloud celebrity pictures, Sony hack, Bangladesh Bank heist, and eBay data breach have been traced back to attackers infiltrating the network using a well researched and crafted phishing attacks misguiding employees to enter their credentials.

6 Designing a New Student Experience

You have recently become the Head of the Higher Education Portfolio at your home university. Your first move is to cancel the subscription to the old "Whiteboard" system used to support lectures, enrolments, and grading. You replace it with something of your own design, called CAMPUS.

In the following, refer to the framework for discussing security aspects that we introduced in the lecture.

a) Incentives

For each of the following actors, give an example *incentive to attack* or *protect the system*. Each incentive must be different.

- The vice Chancellor of the University
- A student enrolled in a course
- A Tutor of the same course
- A Security guard in the CAMPUS data centre

There are multiple correct answers. Some examples are listed below.

- The Vice Chancellor is concerned about the university remaining operational and its reputation (and hence all data must be protected) and therefore has an incentive to protect the system.
- Students may be motivated to attack the system because they may wish to change their grades (or those of other students) or make the site unavailable so that they don't have to do the quiz etc.)
- Tutors need to enter the marks in order to do their job and get paid (so the site must be available). As such, they have an incentive to protect the system.
- Similarly, the guard wants to be paid and keep their job.

b) Policies

For each of the following actors, give a possible security policy they might want to design and a security goal that relates to it.

- The Vice Chancellor of the University
- A student enrolled in a course
- A tutor of the same course

Again several answers are possible.

- The Vice Chancellor may ask that all data remain correct, secret, the privacy of all users and staff is guaranteed etc. The related security goals would be secrecy, integrity, and privacy etc.
- The student may want the grading to be trackable to understand whether someone manipulated their marks. This would relate to accountability.
- The tutor would like to work online and get the payments in a timely manner. The security goal is availability.

7 Terminology

a) Person-in-the-middle-attack
What is a Person-in-the-middle aka Man-in-the-middle attack?

An attack on the network (or some transmission bus) where the attacker manages to be between two communicating parties, usually carried out to eavesdrop on, or tamper with, the data. Examples are attacks on TLS (with broken authentication) or sniffing devices clipped into a data bus. We will discuss this more when we discuss asymmetric cryptography later in this course.

b) Distributed Denial-of-Service
What is a Denial-of-Service attack? And a distributed one?

An attack where someone over the network attempts to exhaust the resources of the victim. DoD attacks can be carried out by a form of malicious crowdsourcing (as demonstrated by **LOIC**), or by botnets. What usually happens is bots (compromised computers) or a group of people start sending large volumes of requests towards a web service and when the service can't handle such large load is is not going to be available even for the legitimate users. The term "distributed" comes when this attack is co-ordinated in a manner from different geo locations, networks, and devices so that it is extremely difficult to seperate attack traffic from benign traffic. "Amplification" is a common tactic often used with DDoS attacks, *e.g.,* DNS amplification: the response is larger than the query and may even propagate to other DNS servers. A summary of notable DDoS attacks in history can be found **here.** We will cover more on DDoS attacks during our web security lecture.

c) Supply Chain Attack
What is a Supply Chain Attack?

A supply chain attack is an attack in which the attacker inserts a malicious code or even a malicious component into a trusted piece of software or hardware. By doing so the attacker can hijack the victims's distribution systems to turn any application they sell, any software update they push out, even the physical equipment they ship to customers, into Trojan horses. With one well-placed intrusion, the attackers can create a springboard to the networks of a supplier's customers - sometimes numbering hundreds or even thousands of victims. Supply chain attacks are increasing because of their economies of scale. Recent examples for supply chain attacks includes **SolarWinds** and **Kaseya** attacks. We will discuss about supply chain attacks during our software security lecture.

d) Security Goals
Tick all that are **not** classic security goals. Explain your choice.

[ ] Authenticity

[ ] Integrity

[ ] Deterrence

[ ] Authorisation

[ ] Confidentiality

[ ] Denial-of-service prevention

[ ] Authenticity

[ ] Integrity

[X] Deterrence

[ ] Authorisation

[ ] Confidentiality

[X] Denial-of-service prevention

The security goals we discussed during the lecture are; *Confidentiality, Integrity, Authenticity, Authorisation, Accountability, Non-repudiation, Deniability, Availability, and Privacy.* "Denial-of-service prevention" is only one aspect of the goal "Availability". "Deterrence" (the idea threat of punishment will deter people from committing crime and reduce the probability and/or level of offending in society) is not a security goal.

8 Homework - ACSC Essential Eight

Essential Eight by the Australian Cyber Security Centre (ACSC) is a series of baseline mitigation strategies recommended for organisations to improve their security posture. Implementing these strategies as a minimum makes it much harder for adversaries to compromise systems.

- Find information about the Australian Cyber Security Centre (ACSC) Essential Eight maturity model.

  ACSC Essential eight defines eight aspects of security w.r.t Windows-based systems.

  `https://www.data3.com/solutions/security/acsc-essential-eight/`

  They are:

  - Application Control
  - Application Patching
  - Restrict Administrative Privileges
  - Patch Operating Systems
  - Configure Microsoft Office Macro Settings
  - Using Application Hardening
  - Multi-Factor Authentication
  - Regular Backups

  When implementing the Essential Eight, organisations should first identify a target maturity level that is suitable for their environment. Organisations should then progressively implement each maturity level until that target is achieved.

  `https://www.cyber.gov.au/acsc/view-all-content/publications/essential-eight-maturity-model`

- Discuss how each strategy is applicable to a new superannuation provider.

  Being a high-risk company, the company must aim to go to maturity level three as quickly as possible.