CS 3IS3 Fall 2007

01 Basic Information Security

William M. Farmer

Department of Computing and Software
McMaster University

23 September 2007

The Information Age

- Information drives commerce and culture.
- Made possible by modern computing and communication technology.
- Key infrastructure: Internet.

The New Information World (1)

- World Wide Web.
  - The Web has become a universal library.
  - Example: The Wikipedia has eclipsed the Encyclopaedia Britannica as the most important encyclopedic source of information.
  - Essentially all new public information is put on the Web.
  - There are several projects to put vast amounts of old information on the Web.
  - Example: Google's agreements with five major libraries (Harvard, University of Michigan, New York Public Library, Oxford, and Stanford).
- Commerce.
  - Information is now a major commodity.
  - Information systems are a major tool of commerce.

The New Information World (2)

- Digital Property.
  - Much property is now digital.
  - Examples include books, articles, news, music, video, and software.
  - Digital property can be reproduced almost instantaneously at extremely low cost.
- Information Ownership
  - Who should own intellectual and digital property?
  - Example: Myriad Genetics is fighting to keep hold of a patent for two breast cancer genes.
  - Who should own the metadata about intellectual and digital property?
- Privacy.
  - Privacy is threatened by the new technology.
  - Example: Identity theft.
  - Encryption may enable some privacy to be preserved.

The New Information World (3)

- Risks.
  - Much of the economy depends on computer networks and software.
  - Many systems of our economy are tied together.
- Cybercrime.
  - Crime via the computer and communication networks is a new development of major concern.
  - Example: Original design of the Internet infrastructure is inadequate.
  - Crime can be perpetrated electronically from a distance.
  - National borders are no longer a major obstacle to crime.
- Information Warfare.
  - Warfare may now include attacks on information systems (possibly instead of on military resources).
  - Example: May 2007 cyberattack on Estonia.
  - Small countries and groups can attack large countries.

Information Security

- Concerned with the protection of:
  - Electronically stored and manipulated information.
  - The systems used to store and manipulate information.
- Growing, dynamic field.
  - Has major importance in the information age.
  - Network security is an important subfield.
- Closely related to the problem of software reliability.
  - Information systems and security mechanisms are heavily based on software.
  - Software is difficult to develop and maintain and very often unreliable.

Why is Information Security Unique?

- Concerned with misuse instead of proper use.
- Hard to engineer.
  - Involves most components of an information system.
  - Information security requirements clash with many other system requirements.
  - Cuts across component boundaries and levels of abstraction.
  - Hard to separate from other concerns.
- A system is only as secure as its weakest component.

What Needs to be Protected?

1. Data.
    - Confidentiality.
    - Integrity.
    - Availability.
2. Information systems.
    - System confidentiality.
    - System integrity.
    - Availability of services.
    - System resources (disk storage, CPU cycles, etc.).
    - Monitoring mechanisms.
    - Security mechanisms.
3. Your personal and organization's reputation.

Confidentiality

- Confidentiality (also called privacy) is the state in which information or resources are concealed.
- Confidentiality also applies to metadata about information and resources.
  - Examples include existence, location, protection, etc.
- Confidentiality is achieved by following the need to know principle, a special case of the principle of least privilege.
- Military interest in keeping information secret was the main driving force behind the development of mechanisms to achieve confidentiality in the years between World War II and the advent of the Internet.

Integrity

- Integrity is the state in which data or resources have not been accidently or maliciously modified or destroyed.
- Integrity also applies to metadata about information and resources.
  - Examples include origin, provenance, access history, etc.
- An integrity violation reduces the trustworthiness of the data or resources.
- There are two approaches to maintain integrity:
  - Prevention of unauthorized attempts to modify the data or resources.
  - Detection of integrity violations or unauthorized modifications.
- The banking industry has been a major player in the development of mechanisms to achieve integrity.

Availability

- Availability is the state in which information or resources can be used as needed.
- Availability is an important aspect of reliability.
- Denial of service attacks are attempts to block availability.
  - They are difficult to detect because they can look like legitimate, but possibly atypical, attempts to access information and resources.

Threats and Attacks

- A threat is a potential violation of confidentiality, integrity, or availability.
- An attack is an attempt to violate confidentiality, integrity, or availability.

Kinds of Threats

- Snooping.
- Modification.
- Spoofing.
- Repudiation of origin.
- Denial of receipt.
- Delay.
- Denial of service.

Where do the Threats Come From?

- Hardware failures.
- Software failures.
- Configuration mistakes.
- Operational mistakes.
- Insiders.
- Hackers.
- Malicious code (such as viruses).
- Criminals, vandals, and terrorists.
- Natural disasters.

Security Policies

- A security policy is a document that states what services and behavior are allowed and disallowed.
- The security policy for an organization defines what is meant by “security” within the organization.
- A security policy should be a written document available to all members of the organization.
- The composition of security policies is a concern: security vulnerabilities can occur if the security policies conflict.

Security Mechanisms

- A security mechanism is a method, tool, or procedure for enforcing a security policy (Bishop).
- An organization's security posture is the collection of security mechanisms that the organization has in place.
- A security system is a collection of coherent security mechanisms intended to enforce a security policy.

Security Strategies

- A security strategy is an approach to enforcing a security policy.
- Goals of a security strategy:

  1. Prevention: Prevent an attack from occurring.
  2. Detection: Detect an attack.
  3. Recovery: Stop an attack and then recover, or function as best as possible under an attack.

Assumptions

- Underlying every security policy are certain assumptions.
- Two principal assumptions are:

  1. The policy correctly and unambiguously partitions the set of system states into secure and nonsecure states.
  2. There exists a set of security mechanisms that will enforce the security policy.

Security Mechanisms

- Let s be a system such that P is the set of its possible states and $Q \subseteq P$ is the set of its secure states.
- Let m be a security mechanism and $R_m \subseteq P$ be the set of states to which s is restricted by m.
- m is secure if $R_m \subseteq Q$.
- m is precise if $R_m = Q$.
- m is broad if R_m is not secure.
- The goal of a security system is to behave as a single precise security mechanism.
- In practice security mechanisms are usually broad, allowing the system to enter some nonsecure states.

Trust and Assurance

- Trust is a measure of the confidence that a system satisfies its requirements.
- Assurance is evidence that a system satisfies its requirements.
- The more assurance a system has the more it is trusted.
- Assurance is established in three major steps:
  1. A specification of the requirements of the system is produced.
  2. A design is produced that satisfies the requirements specification of the system.
  3. An implementation is produced that satisfies the design of the system.

Operational Issues

- The selection of a security policy or security system requires a cost-benefit analysis.
  - The costs include the cost of developing security measures as well as the cost of security breaches.
- A risk analysis is needed to assess the likelihood of specific threats and the level of damage they would cause.
  - Risk is a function of environment and time.
  - Some risks may be considered acceptable.
- Security measures are constrained by both law and custom.

Human Issues

- Human issues play a major role in information security.
- Threats that have never happened are often hard for people to take seriously.
- Those who are responsible for security must be given the power needed to implement adequate security measures.
- Security needs adequate human and material resources.
- Personnel need adequate training and must understand the importance of security measures.
- Insiders should be carefully monitored.
- Measures should be taken to counter social engineering attacks.
- Human error should be expected.

Security Life Cycle

1. Threats
2. Policy
3. Security system development
    3.1 Requirements specification
    3.2 Design
    3.3 Implementation
    3.4 Operation and maintenance

Protection Systems

- A protection system describes the conditions under which a system is secure (Bishop).
  - The protection system is a subsystem of the system itself.
  - As the system changes, the conditions that determine what is meant by being secure changes as well. That is, what is meant by security changes over time.
- A protection state is a state of the protection system.
  - It is a substate of a corresponding state of the system.
  - It determines what security means in the corresponding state of the system.

Access Control Matrix Model: Definitions

- The access control matrix model is a framework for specifying a protection system.
- Let O be the set of objects that need protection, S be the set of subjects that can access the objects in O, and R be the set of rights subjects can be granted for accessing objects.
- An access control matrix is a matrix

  $$A : S \times O \rightarrow \mathcal{P}(R)$$

  that assigns each subject-object pair $(s, o) \in S \times O$ a set $A(s, o) \subseteq R$ of rights.
- A protection state is represented by (S, O, A).
- A state s of a system is secure if it conforms to the protection state corresponding to s.

Access Control Matrix Examples

1. Process-file example.
   - The objects are files.
   - The subjects are processes.
   - The rights are read, write, execute, append, own.
2. Unix file system.
   - The objects are files, directories, and devices.
   - The subjects are Unix users.
   - The rights are read (r), write (w), and execute (x).
3. Programming language.
   - The objects are variables and procedures.
   - The subjects are procedures.
   - The rights are read, write, and call.

Commands (1)

- As a system changes, the protection state changes.
  - The initial protection state is denoted by (S_0, O_0, A_0).
- Modifications of the protection state can be modeled as applications of the following primitive commands:
  1. Create a subject: **create subject** s.
  2. Create an object: **create object** o.
  3. Enter a right into the matrix: **enter** r **into** $A(s, o)$.
  4. Delete a right from the matrix: **delete** r **from** $A(s, o)$.
  5. Delete a subject: **destroy subject** s.
  6. Delete an object: **destroy object** o.
- The primitive commands can be combined into defined commands.
  - Defined commands may include conditions.

Commands (2)

- A defined command is mono-operational if it only invokes a single primitive command.
- A defined command is monoconditional if it only includes a single condition.
- A protection system based on the access control matrix model is represented by a tuple (S_0, O_0, A_0, R, C) where:
  - (S_0, O_0, A_0) is an initial protection state.
  - R is a finite set of rights.
  - C is a finite set of defined commands.

Access Control Matrix Model: Summary

- The access control matrix is a simple mechanism for representing a protection system.
- Access control matrices are not an efficient mechanism for implementing a protection system.
- There are 6 primitive commands for modifying an access matrix mechanism.
- A protection system based on the access control matrix model includes an initial protection state, a finite set of rights, and a finite set of defined commands.

Verifying Security

- Fundamental question: How can we verify that an information system is secure?
- Approaches:
  1. Informal: Show that the information system enforces its security policy.
  2. Mathematical: Develop a mathematical model S of the information system and a mathematical model P of its security policy and then mathematically prove that S enforces P.
  3. Formal: Develop a formal mathematical model S of the information system and a formal mathematical model P of its security policy and then prove in a formal proof system that S enforces P.
  4. Computational: Develop a generic algorithm for determining whether an arbitrary information system enforces an arbitrary security policy and then apply it to the information system and its security policy.

The Safety Problem

- Let S be a protection system based on the access control matrix model.
- A right r that is added by S to an element of an access control matrix not containing r is said to be leaked by S.
- If S can never leak the right r, S is said to be safe with respect to the right r. If S can leak the right r, S is said to be unsafe with respect to the right r.
- The safety problem is to determine whether or not a given protection system based on the access control matrix model with an initial state s_0 is safe with respect to a given right r.

Basic Results

- **Theorem 1** (Harrison, Ruzzo, Ullman). The safety problem is undecidable.
  - Proof: Reduce the halting problem to the safety problem by showing how a access control matrix can simulate a Turing machine.
- **Theorem 2** (Harrison, Ruzzo, Ullman). The safety problem for mono-operational protection systems is decidable.
  - Proof: Without loss of generality, we may assume that the length of the shortest sequence of commands that leaks a given right is less than or equal to

$$|R|(|S_0|+1)(|O_0|+1)+1.$$

- **Theorem 3** (Harrison, Ruzzo). The safety problem for monoconditional protection systems with **create**, **enter**, and **delete** primitive commands (but no **destroy** primitive commands) is decidable.