

CS361: Introduction to Computer Security

Introduction

Dr. Bill Young
Department of Computer Sciences
University of Texas at Austin

Last updated: January 22, 2020 at 07:46

What Does Security Mean?



“Security” is an extremely expansive term.

- Personal security
- Physical security
- Corporate security
- Homeland security
- Operations security
- Personnel security
- Communication security
- Computer security
- Network security
- System security

What do these have in common? What does “security” mean?

What Does Security Mean?

In the most general terms, *security* seems to mean something like “protection of assets against attack.”

But what *assets*? What kind of *attack*? What does *protection* mean? Doesn't the meaning of “protection” vary depending on the nature of the threat?



Some examples of threats:

- Interruption:** an asset becomes unusable, unavailable, or lost.
- Interception:** an unauthorized party gains access to an asset.
- Modification:** an unauthorized party tampers with an asset.
- Fabrication:** an asset has been counterfeit.

What is the Security Problem

“Security” is an increasingly prevalent problem in computer science. *Why do you suppose that is?*

What is the Security Problem

“Security” is an increasingly prevalent problem in computer science. *Why do you suppose that is?*



- Increased connectivity;
- Large number of valuable assets online;
- Low threshold to access;
- Sophisticated attack tools and strategies available;
- Others?

What do each of these mean? Why are they relevant?

America's failure to protect cyberspace is one of the most urgent national security problems facing the new administration... It is a battle we are losing. Losing this struggle will wreak serious damage on the economic health and national security of the United States. –CSIS report on *Securing Cyberspace for the 44th Presidency*

A top FBI official warned today that many cyber-adversaries of the U.S. have the ability to access virtually any computer system, posing a risk that's so great it could "challenge our country's very existence." –Computerworld

Cyber threats are asymmetric, surreptitious, and constantly evolving—a single individual or a small group anywhere in the world can inexpensively and secretly attempt to penetrate systems containing vital information or mount damaging attacks on critical infrastructures. Attack tools and resources are readily available on the Internet and new vulnerabilities are constantly discovered and exploited. Moreover, the pervasive interconnectivity of the IT infrastructure makes cyber attacks an increasingly attractive prospect for adversaries that include terrorists as well as malicious hackers and criminals. –Federal Plan for Cyber Security and Information Assurance Research and Development

Our general concern is IT security. IT security is commonly defined to have two aspects:

- The aim of *computer security* is to preserve computing resources against abuse and unauthorized use, as well as to protect data, from accidental or deliberate damage, disclosure or modification.
- The aim of *communication security* is to protect data that represents and codes information during its transmission in computer networks and distributed systems. *Network security* is a synonym.

We'll cover both areas. Prof. Shacham's class in Network Security and Privacy is largely about communication security.

Thought experiment: Banks clearly have a “security problem.” They have significant assets to protect. *How do they do it?*

Thought experiment: Banks clearly have a “security problem.” They have significant assets to protect. *How do they do it?*

- vaults
- guards
- alarm systems
- traceable assets
- procedural safeguards (against insider attacks)



What are the analogous safeguards use to protect information assets?

The Difficulty in Security

Most areas of computer science are concerned with ensuring that something good happens. In contrast, security is all about ensuring that *bad things never happen*.

Not only do you have to find *bugs* that make the system behave differently than expected, you have to identify any features of the system that are susceptible to misuse and abuse.

You have to defeat an *actively malicious adversary*. Security expert Ross Anderson characterizes this as "*Programming Satan's Computer*."



The Difficulty in Security

Thus, the hardest thing about security is convincing yourself that you've thought of all possible attack scenarios, before the attacker thinks of them.



“A good attack is one that the engineers never thought of.” –Bruce Schneier

The Difficulty in Security

Perfect security is probably impossible in any useful system.


Often, this means that a *tradeoff* is necessary between security and other important software project goals including:



- functionality,
- usability,
- efficiency,
- time-to-market,
- simplicity.

Is It Getting Better?

“The three golden rules to ensure computer security are: do not own a computer; do not power it on; and do not use it.” –Robert H. Morris (mid 1980’s), former chief scientist of the National Computer Security Center



“Unfortunately the only way to really protect [your computer] right now is to turn it off, disconnect it from the Internet, encase it in cement and bury it 100 feet below the ground.” –Prof. Fred Chang (2009), former director of research at NSA

“A plausible worst-case worm could cause \$50 billion or more in direct economic damage by attacking widely used services in Microsoft Windows and carrying a highly destructive payload.”

–Nicholas Weaver and Vern Paxson

Nevertheless, organizations often choose not to investigate or prosecute intruders (hackers). Why do you suppose that is?


“A plausible worst-case worm could cause \$50 billion or more in direct economic damage by attacking widely used services in Microsoft Windows and carrying a highly destructive payload.”

–Nicholas Weaver and Vern Paxson

Nevertheless, organizations often choose not to investigate or prosecute intruders (hackers). Why do you suppose that is?

- They don't want to expose vulnerabilities in their systems.
- They want to protect their public image.
- Intruders are sometimes viewed as mere pranksters.
- Sometimes electronic assets are not viewed as valuable.

“You can’t understand a company’s network security by looking at public events—that’s a bad metric. All the public events tell you are, these are attacks that were successful enough to steal data, but were unsuccessful in covering their tracks.” –Bruce Schneier



Numerous “secure” operating systems have been developed over the years:

- PSOS (Provably Secure OS)
- KSOS (Kernelized Secure OS)
- LOCK (Logical Co-processor Kernel)
- SCOMP (Secure Communications Processor)
- Secure Xenix
- Greenhills Integrity RTOS

Which of these do you use? Which have you even heard of? *Why do you suppose that is?*

Some Rules of Thumb

Maybe it's because those systems violated one or more of these rules:

- Security is often inversely related to utility.
- Security expenditure should be relative to the threat.
- Security should be considered from an overall systems point of view.
- Security should be affordable and cost effective.
- Security should be as simple as possible.



Identifying the Risk

He who defends everything defends nothing. —old military adage



Security planning is really about *risk management*—how much to spend on protection against what.

This should be done within the broader framework of managing non-IT risks.

Network Security eSummit, January 23, 2020


KEYNOTE: The Top Critical Success Factor to Security is
Managing Risk

Mark Thomas, President, Escoute Consulting

Where is the Risk?

Ken Thompson (Turing Award lecture, 1983) said that *not even complete control over source code is sufficient to ensure the absence of malicious functionality.*

“You can't trust code that you did not completely create yourself. (Especially code from companies that employ people like me.) No amount of source-level verification or scrutiny will protect you from using untrusted code. [...] As the level of program gets lower, these bugs will be harder and harder to detect. A well-installed microcode bug will be almost impossible to detect.” –Ken Thompson



Security is About Managing Risk

In *Building Secure Software*, Viega and McGraw assert that software and system security is “all about managing risk.” This can be done through:

Risk acceptance: some risks are simply tolerated by the organization.

Risk avoidance: not performing an activity that would incur risk.

Risk mitigation: taking actions to reduce the losses due to a risk.

Risk transfer: shift the risk to someone else.

There is generally much more money in a bank than in a convenience store; but which is more likely to be robbed? Why?

Annualized Loss Expectancy

One common tool for risk assessment is *annualized loss expectancy* (ALE), which is a table of possible losses, their likelihood, and potential cost. **Example:** consider a bank with the following ALE. Where should they put their security dollars?

Loss type	Amount	Incidence	ALE
*SWIFT fraud	\$50,000,000	.005	\$250,000
ATM fraud (large)	\$250,000	.2	\$50,000
ATM fraud (small)	\$20,000	.5	\$10,000
Teller theft	\$3,240	200	\$648,000

* The Society for Worldwide Interbank Financial Telecommunication ("SWIFT") operates a worldwide financial messaging network, allowing large scale transfer of funds.

Is ALE the Right Model?

Annualized Loss Expectancy effectively computes the “expected value” of any security expenditure. *Is that the right risk model?*

$$EV = \sum P(X_i) * X_i$$

Consider the following two scenarios:

- 1 I give you a dollar.
- 2 We flip a coin. Heads: I give you \$100. Tails: you give me \$98.

The expected values are the same, but the risks seem quite different. What if it were \$1000 at risk instead of \$100?

What factors go into assessing risk in real life? In computer security?

Aspects of Computer Security

Historically, computer security has been defined to encompass:

Confidentiality: (also called secrecy/privacy)

who can *read* information?

Integrity: who can *write*, modify or generate information?

Availability: are resources available when needed?



Which do you think is the most important?


Aspects of Computer Security

Some experts (e.g., NSA) add to this list:

Authentication: how do we establish identity?

Non-repudiation: can I deny my actions?

Code Validation: how can I be sure my implementation is correct?



Other Aspects of Security

There are lots of other topics often mentioned when discussing computer security:

- authorization,
- access control,
- firewalls,
- passwords,
- certificates,
- cryptography,
- digital signatures, etc.



We'll talk about all of these, but these are *mechanisms* for protecting one or more of the major aspects such as confidentiality or integrity.

How do I protect my information from unauthorized disclosure?

Historically, this was the first computer security concern, and remains extremely important in military and commercial settings.

- How do you group and categorize data?
- How do you characterize who is authorized to see what?
- Can authorizations change over time?
- How are the permissions administered and checked?
According to what rules?
- How do you control the flow of “permissions” in the system?
I.e., can I authorize others to view data that I am authorized to view?

How do I protect my information from unauthorized modification?

Integrity is a fuzzier notion than confidentiality and more context dependent. But for many commercial applications it is *more important* than confidentiality. Can you give some examples?

- Who is authorized to supply or modify data?
- How do you separate and protect assets?
- Can you detect and/or correct erroneous or unauthorized changes to data?
- Can authorizations change over time?

How do I ensure that my information/system is available when I need it?

Threats to availability are often called *denial of service* (DoS) attacks.

- Are resources provided in a timely fashion?
- Are resources allocated fairly by the system?
- Is the system so difficult/tedius to use as to be useless?
- If faults occur, can the system compensate/recover?
- How is concurrency controlled by the system?

Many virus and worm attacks are DoS attacks. The MyDoom worm cost businesses an estimated \$38.5 billion, according to some estimates.

Characteristics of Attackers

Malicious attackers must have:

- *method*: the skills, knowledge and tools to carry out the attack;
- *opportunity*: the time and access needed to attack;
- *motive*: a reason to attempt penetration of the system.



Knowledge is widely available. *Keeping your security mechanism secret usually is not a good security approach.* Experts call that “security by obscurity.”

Characteristics of Computer Intrusion

“If one overlooks the basement windows while assessing the risks to one’s house, it does not matter how many alarms are put on the doors and upstairs windows.” –Melissa Danforth

A computing system comprises: hardware, software, storage media, data, people. An intruder may target any one or any combination of these.

Principle of Easiest Penetration: an intruder will use any available means to subvert the security of a system.



This implies that security analysis must be thorough, comprehensive and on-going.

"Why put steel doors in paper walls?" –Rich DeMillo

Most Leaks Don't Involve Technology

In 1996, news of possible signs of life in a Martian meteorite called ALH84001 leaked out ahead of a press conference that had been scheduled by NASA. *This was partly because a high-ranking White House official told a prostitute about the meteorite.* NASA had to scramble to reschedule its press conference to an earlier date to satisfy the growing demand for information from the press and the public.

–www.newscientist.com (Aug 1, 2006)



Why You Need a Systems Point of View

Some random vulnerabilities:

Program	Effect
zLib	DoS affecting many programs, including those that display PNG files.
Internet Explorer	Malicious PNG file can be used to execute arbitrary code when displayed in IE.
libPNG	DoS affecting users of Firefox, Opera, Safari, and many others.
MS GDI+	JPG-rendering code enables the remote execution of arbitrary code. Affects IE, MS Office, and other MS products.
zLib	Potential for remote code execution. Affects programs that display or manipulate PNG files.
Windows Graphics Rendering Engine	Rendering of WMF files enables remote execution of arbitrary code. Exploited through IE.
Java 2 Platform	Rendering of GIF image allows remote execution of arbitrary code through hostile applet.

Taking a Systems Point of View

Notice that none of the programs (in the table) were “security features” of the relevant systems. They were all related to displaying images.



Yet each exploit meant almost total compromise of the security of the system.

What does that mean for the security professional/community?

How can a system experience loss or harm?

- A *vulnerability* is a weakness in security.
- A *threat* is a set of circumstances that has the potential to cause harm.
- An *attack* is an (attempted) exploitation of some vulnerability in the system.
- A *control* or *countermeasure* is a means of removing or reducing a vulnerability.

Typically, when you buy a book on “computer security” it will have as subject matter one of two broad topics:

- technical and theoretical aspects of security;
- physical, procedural, and operational solutions to protecting information (the more applied end of the spectrum).

This course concentrates on the first of these.

The second is also extremely important for the security professional, but not currently covered in any UT CS courses.