CS329E: Elements of Security

Intro to Security

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

Last updated: February 2, 2015 at 10:42

What Does Security Mean?

"Security" is an extremely expansive term.

- Personal security
- Physical security
- Corporate security
- National (homeland) security
- Operations security
- Personnel security
- Communication security
- Computer security
- Network security
- System security

What do they all have in common? I.e., 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?

What Does Computer Security Mean?

Some examples of threats in *computer security*:

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.

Can you think of examples of each of these? Have any of them happened to you?

What is the Security Problem

"Security" is an increasingly prevalent problem for computer users.
Why do think that is?

What is the Security Problem

"Security" is an increasingly prevalent problem for computer users.
Why do think that is?

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

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

Nature of the Threat

America's failure to protect cyberspace is one of the most urgent national security problems facing the new administration that will take office in January 2009. ... 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, Dec. 2008*

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, March 24, 2010*

Nature of the Threat

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 (2006)*

Nature of the Threat

A dozen determined computer programmers can, if they find a vulnerability to exploit, threaten the United States' global logistics network, steal its operational plans, blind its intelligence capabilities or hinder its ability to deliver weapons on target. – *William J. Lynn, U.S. Deputy Secy of Defense, Foreign Affairs (2010)*

IT Security

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

- *Computer security* protects computing resources against abuse and unauthorized use, as well as to protect data, from accidental or deliberate damage, disclosure or modification. (Baker, 1991; Amoroso, 1994)
- *Communication security* protects data that represents and codes information during its transmission in computer networks and distributed systems. *Network security* is a synonym. (Davies and Price, 1984; Devargas, 1993)

Defining Security

Thought experiment: Banks clearly have a "security problem." They have significant assets to protect. What measures do banks use to protect their assets?

Defining Security

Thought experiment: Banks clearly have a "security problem." They have significant assets to protect. What measures do banks use to protect their assets?

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

Can we envision analogous safeguards 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*. 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

Security Isn't the Point

Security is often treated as an afterthought. *No-one builds a digital system for the purpose of being secure*. They build digital systems to do something useful.

Security mechanisms may be viewed as a nuisance to be subverted, bypassed, or disabled.

It's often hard to convince management to allocate extra resources to prevent attacks that may never occur.

The More Things Change ...

"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

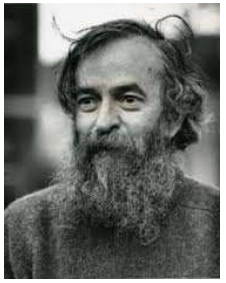

"Unfortunately the only way to really protect [your computer] right now is to turn it off, disconnect it from the Internet, encase it in cement and bury it 100 feet below the ground." –Prof. Fred Chang (2009), former director of research at NSA

Perfect Security: It Ain't Happening

Perfect security is unachievable in any useful system.

We trade-off security with other important goals: functionality, usability, efficiency, time-to-market, and simplicity.

A Security Paradox

"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, 6/14/04

Nevertheless, organizations often choose not to investigate or prosecute intruders (hackers). Why might that be?

A Security Paradox

"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, 6/14/04

Nevertheless, organizations often choose not to investigate or prosecute intruders (hackers). Why might that be?

- 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.

A Security Paradox

"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

Secure OS Projects

Numerous "secure" operating systems have been developed over the years:

- KSOS (Kernelized Secure OS, 1979),
- PSOS (Provably Secure OS, 1980),
- SCOMP (Secure Communications Processor, 1983),
- LOCK (Logical Co-processor Kernel, 1987),
- Secure Xenix (1987),
- Greenhills Integrity OS (2000)

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.

You Can't Defend Everything

Modern information management systems are a complex, "target-rich" environment comprising: hardware, software, storage media, peripheral devices, data, and people.

He who defends everything defends nothing. –old military adage

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.

Compromised system software such as the compiler could introduce arbitrary functionality into the object code.

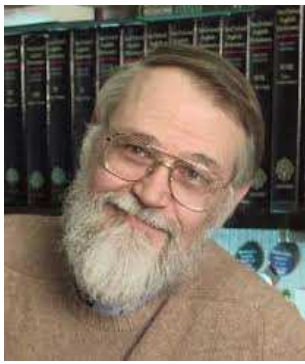

Complete assurance would require scrutinizing not only the source program, but also the compiler, linker, loader, assembler, microcode, and even hardware.

Where is the Risk?

"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 Risk Management

Viega and McGraw, *Building Secure Software* assert that software and system security is "all about managing risk."

Risk is the possibility that a particular threat will adversely impact an information system by exploiting a particular vulnerability. The assessment of risk must take into account the consequences of an exploit.

Risk management is a process for an organization to identify and address the risks in their environment.

Current IT Security Metrics

"If you can't measure something, you can't understand it. If you can't understand it, you can't control it. If you can't control it, you can't improve it." H. James Harrington

Some Popular IT Security Metrics:

- Security risk assessment matrices
- Security vulnerabilities and incident statistics
- Annualized loss expectancy (ALE)
- Return on investment (ROI)
- Total cost of ownership (TCO)

Current Metrics are Flawed

Typical Security Risk Assessment:

		Likelihood of Event		
		High	Medium	Low
Severity of Impact	High	"We're Doomed!"	Bad	Outlier
	Medium	Bad	Not Good	Error
	Low	Annoyance	Typical	"Whatever..."

Problem: This doesn't measure security risk; it measures human judgments about risk. *That can be useful, assuming you understand what you're getting.*

Counting Vulnerabilities or Incidents

The number of vulnerabilities discovered or security-related "incidents" are often used as general indicators of the level of security.

But these depend critically on:

- How thorough are your scans?
- How many systems are scanned?
- What severity is assigned to what vulnerabilities/incidents?
- How many applications are deployed?
- How does your number compare with peers?

Annualized Loss Expectancy

ALE is a common tool for risk assessment. Given potential threats, where do you put your security dollars?

Risks in a large bank:

Loss type	Amount	Incidence	ALE
SWIFT fraud	$50,000,000	0.005	$250,000
ATM fraud (large)	$250,000	0.2	$100,000
ATM fraud (small)	$20,000	0.5	$10,000
Teller theft	$3,240	200.0	$648,000

But this is really nothing more than expected value! Is that the right way to compute risk?

* 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?

Consider the following two scenarios:

1. I give you a dollar.
2. We flip a coin. Heads: I give you $1000. Tails: you give me $998.

The expected values are exactly the same, but the risks seem quite different.

Often ALE deals in opinions and expectations because IT security does not have data to define actual probabilities.

ROI and TCO

Return on Investment (ROI) attempts to calculate how much benefit will be gained from an investment.

- How do current security expenditures affect future losses? This is very hard to estimate.
- Traditionally, ROI (in a financial setting) involves profit or rate of return. These don't apply well for a security investment.

Total Cost of Ownership (TCO) seeks to quantify the cost over the entire lifecycle of the investment.

- Only really applies to security purchases, not to measurement of the IT security process.
- Data for adequate comparison is lacking.

Lessons from Other Industries

We typically don't have very good data for estimating IT security risk!

Other industries—insurance, manufacturing, design—have a long history of dealing with risk.

Some lessons:

1. Metrics and process will improve as the ability to collect, analyze and understand data improves.
2. *Security is a business process.* You must measure the business process to measure security.
3. *Security results from human activities.* You must understand people as well as technology.

Are Goals Appropriate? Cisco Example:

The security goals must align with the mission of the enterprise.

Operating System Count & Urgent Vulnerabilities

	# Hosts	# Vuln.	# Hosts w/ Vuln.	% Hosts w/ Vuln.
Windows	4212	593	347	8.2
Linux	8026	62	41	< 1
Solaris	2733	216	143	5.2
Cisco	4626	6	6	< 1

Presented with these findings, Cisco management responded that since Linux and Cisco OS teams "had less than 1 percent of their hosts with high severity vulnerabilities, those teams must be spending too much time, effort and resources patching their hosts." (Hayden, *IT Security Metrics*, p. 84)

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?

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

Authentication: how do we establish identity?

Non-repudiation: can I deny my actions?

Which among these is the most important?

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.

Confidentiality: Questions

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?

Integrity: Questions

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?

Availability: Questions

How do I ensure that my information / system is there 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?

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

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

"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

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

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

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

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, who then sold the information to a tabloid. 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 (8/1/06)

"Life if tough, but it's tougher if you're stupid." –John Wayne in *Sands of Iwo Jima*

Why You Need a Systems Point of View

Several severe system breaches:

Date	Program	Effect
March 2002	zLib	DoS affecting many programs, including those that display PNG files.
Nov. 2002	Internet Explorer	Malicious PNG file can be used to execute arbitrary code when displayed in IE.
Aug. 2004	libPNG	DoS affecting users of Firefox, Opera, Safari, and many others.
Sep. 2004	MS GDI+	JPG-rendering code enables the remote execution of arbitrary code. Affects IE, MS Office, and other MS products.
July 2005	zLib	Potential for remote code execution. Affects programs that display or manipulate PNG files.
Dec. 2005	Windows Graphics Rendering Engine	Rendering of WMF files enables remote execution of arbitrary code. Expoited through IE.
Jan. 2007	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?

Attacks: Some Terminology

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.

Parts of Computer Security

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).

Both are important. We'll tend to concentrate on the first, but also mention the second frequently.