Introduction
CS 236
Advanced Computer Security
Peter Reiher
April 1, 2008

Outline

- Subject of class
- Class topics and organization
- Reading material
- Class web page
- Grading
- Projects
- Office hours

Subject of Class

- Advanced topics in computer security
- Concentrating on unsolved problems and recent research
- Covering both networks and computers
  - Only real crypto research is out of scope
- Intended for students with serious research interest in security
- **Goal is to help such students learn how to do this kind of research**

Doing Research in Security

- A lot of bad research is done is security
  - Unimportant problems
  - Unrealistic approaches
  - Unverified conclusions
- The point of the class is to set you on the right road

Class Organization

- A little bit different
- Every Tuesday I will describe a problem area and a solution approach
- On Thursday, entire class will discuss that idea
  - Critiquing, designing, suggesting other alternatives
- More or less how a research group works

Tuesday Classes

- I will give a presentation
- Usually two parts
  1. Discussing problem and existing approaches
  2. Suggesting another approach
- Readings will be papers related to the area

In Between Classes

- I will assign students into groups
  - Probably of three students
- Each group should discuss the problem and idea among themselves
- In preparation for a more detailed discussion on Thursday
- Groups will change every week

Thursday Classes

- A general group discussion
  - Involving all students
- Maybe developing idea
- Maybe burying it
- Maybe coming up with something else

Associated Written Assignments

- Each group will produce a five page write-up
- Due before next Tuesday
- Describing their thoughts on the topic
- Will be graded

The Weekly Topics

- No topic the first week
  - Intro today, I won't be here Thursday
- No topic the last week
  - Students will present their projects in those sessions
- That leaves eight slots

Topics We Will Discuss

- Data flow in operating systems
  - Data tethers
- Botnet defenses
  - Infamy
- Securing web servers

Topics We Might Discuss

- Security for sensor networks
- Cyberwarfare and national scale cyber defense
- Data provenance issues
- Operating systems and TPM
- Ubiquitous computing security
- Worms, DDoS, IP spoofing
- Many other possibilities

Reading Material

- No textbook
- 2-4 papers for each class
- Papers will be made available on class web page
- In some cases, web pages may be used instead of papers

Class Web Page

- http://www.lasr.cs.ucla.edu/classes/236_1.spring08
- Will show class schedule
- And list papers for each class
  - With links to them
- Other useful information also there

Grading

- 40% weekly reports
- 10% class participation
- 50% project
- No final exam

Weekly Reports

- Done by small groups
- ~5 pages each
- Discussing/critiquing topic and approach for each week
- Due before the Tuesday of next week

Class Participation

- Not graded on brilliance
- But on involvement and ability to contribute to discussion
- If you can't regularly attend this class, you won't do well in it
- Also not a good class to sleep through
- Or to take if you don't care much about the subject

Class Projects

- Half of your grade
- Group projects (2-4 people)
- On some topic involving computer security
- Must be a research topic
  - Not just implementing known stuff
  - Need not be on topic covered in class

Project Proposals

- Project proposals due at end of 4th week of class (April 25)
- 1-page summary of what you want to do
- Can be submitted as hard copy or email
- Not graded, but required
- I'll approve and/or provide other feedback

Project Status Reports

- Due at end of 7th week of classes (May 16)
- 1-3 page summaries of the progress you've made to that date
  - Hint: there should be some
- Hard copy or email OK
- Not graded, but required

Project Presentation

- Last two class days reserved for project presentations
- In-class presentation of your project
  - Demo, if feasible
- Graded as part of project itself

Project Demonstration

- If not feasible to demo in class, arrange a separate demo with me
- Projects should (usually) produce something demonstrable
- Important that demo shows off something interesting about project
- Graded as part of project

Project Reports

- Written reports on project
- Due Monday of finals week (June 9)
- 15 pages is typical length
- Should:
  - Describe problem and approach
  - Cover difficulties and interesting points
  - Describe implementation
  - Show that you've learned something from it!

What Makes a Good Project?

- Probably requires coding
  - Hardware OK, if you can do it
  - Theoretical work acceptable, but you'll need real results
- Probably requires testing and/or measurement
- Should be research
  - Original work no one else has already done
  - Based on a promising idea
  - Ideally, this should be capable of being converted to a publishable research paper

Office Hours

- MW 2-3
- In 3532F Boelter Hall
- I'm around a lot, so other times can be arranged by appointment
- But I'll be away April 3
  - Possibly other days TBA

Prerequisites

- Should have taken CS 118 and 111
- Should have taken my CS 136 on Computer Security
  - Or similar class elsewhere
- I'm not going to check on this
- But I'll assume you know this material
  - I won't be presenting reviews of this material

Kinds of Security Things You Should Know About

- IPsec
- Security protocols
- Key exchange, certificates, certification hierarchies
- Basics of security threats and mechanisms
- Use of cryptography for authentication, privacy, and other purposes
- Basics of firewalls and virus protection systems
- Basics of viruses and worms

Kinds of Networking Things You Should Know About

- TCP/IP
- Routing protocols
- How DNS works
- Multicast protocols
- Basic ad hoc networking
- Basics of wireless networks
- Basic design and architecture of the Internet

Kinds of OS Things You Should Know About

- File systems
- Basic OS organization
- Important OS elements
  - E.g., booting and device drivers
- IPC and memory management

A Short Introduction

- What is this class really about?
- Learning how to do research in computer security
- Primarily by doing it
  - Partly the weekly discussions
  - Partly the projects

What's Worth Looking At?

- A matter of both opinion and perspective
- Basically,
  - Where are the big risks?
  - Where can we do better?
  - What technologies aren't good enough?

The IRC Hard Problems List

- The Infosec Research Council (IRC)
- Group of US government agencies that care a lot about security
  - Enough to fund research into it
- They are in the process of creating a "hard problems" list

What Are They After?

- A list of the problems that most need solving
  - From US government perspective
- Particularly those that require substantial research
- With an eye towards creating a roadmap for future security research

Who Is the IRC?

- Representatives from most relevant agencies
  - IARPA – IC Advanced Research and Development Activity
  - CIA - Central Intelligence Agency
  - DOD - Department of Defense (including the Air Force, Army, Defense Advanced Research Projects Agency, National Reconnaissance Office, National Security Agency, Navy, and Office of the Secretary of Defense)
  - DOE - Department of Energy
  - DHS - Department of Homeland Security
  - FAA - Federal Aviation Administration
  - NASA - National Aeronautics and Space Administration
  - NIH - National Institutes of Health
  - NIST - National Institute of Standards and Technology
  - NSF - National Science Foundation
  - TSWG - Technical Support Working Group

Where Did Their List Come From?

- Much internal expertise
  - E.g., Doug Maughan, Carl Landweir, Karl Levitt
- Also outside experts
  - Steve Bellovin, Marc Donner, Joan Feigenbaum, James R Gosler , Steve Kent, Peter G. Neumann, Fred Schneider

What's On the List?

- Nine broad topics
- Covering wide range of privacy and security issues
- Not only of concern to US government
  - Or just to government at all
- Best opinion of top security experts of where research is needed

Why Should You Care?

- Revised list will be used to guide government research priorities
  - Intended as tool to get more research funding from Congress
- A lot of the great research of next few years will be in these areas
- If experts are right, you should be focusing attention here

The List

1. GLOBAL SCALE IDENTITY MANAGEMENT
2. INSIDER THREATS
3. AVAILABILITY OF TIME-CRITICAL SYSTEMS
4. BUILDING SCALABLE SECURE SYSTEMS
5. ATTACK ATTRIBUTION AND SITUATIONAL UNDERSTANDING
6. INFORMATION PROVENANCE
7. SECURITY WITH PRIVACY
8. ENTERPRISE LEVEL SECURITY METRICS
9. COPING WITH MALWARE

1. Global Scale Identity Management

- Scope: Identification, authentication, authorization, requisite key infrastructure
- Motivation: Need for seamless IAA across many systems, costs of divergent IAA systems, limits of current PKI, quantum
- Challenges: Scale, churn, anonymity, federation
- Goal: allow seamless identity management in all systems

2. Insider Threats

- Motivation: Frequency and severity of incidents historically, increasing potential
- Challenges: Not unauthorized access, inside knowledge of defenses, "help" from outsiders with substantial resources
- Approaches: Connections to hard problem #1, pervasive auditing, and redundancy
- Goal: Mitigate the insider threat in cyber space so far as it is in physical space

3. Availability of Time-Critical Systems

- Motivation: SCADA, military, home-land security first responders
- Value availability over secrecy
- Work in lossy, ad hoc wireless environments
- Challenges:
  - Limited resources
  - Computational processing power
  - Service quality guarantees given dynamics
- Distributed systems compound problem

4. Building Scalable Secure Systems

- Motivation: High Consequence Systems
- Challenges:
  - Today's systems are huge
  - Catastrophic bugs can be tiny
  - Some developers may be working against us
- Components, subsystems, architectures
- Approaches:
  - Help formal verification to scale
  - Development and formal V&V environments
  - Means of correctly composing formal models
- Goal: E.g., fully verified truly trustworthy TCB

5. Attack Attribution and Situational Understanding

- Motivation: Respond to the unpreventable
- Challenges:
  - Some attacks may be acts of war, others the work of teens, others nations posing as teens.
  - Hostile networks, anonymizers, recordless public access such as wi-fi and internet cafes.
- Big picture and appropriate response
  - Response selection: E.g., degradation of mission instead of total failure
- Attribution: ID of adversaries despite measures to conceal identification

6. Information Provenance

- Motivation: Life-critical and releasability decisions both require pedigree of data
- Challenges:
  - Volume
  - Degree of automated processing and transformation
  - Provenance vs. privacy
- Goal: Track pedigree for every byte of information in exabyte scale systems transforming terabytes of data per day

7. Security With Privacy

- Motivation: More of our interactions and transactions are occurring in cyberspace. Data mining poses risks to privacy and identity theft poses risks to security.
- Challenges: Current strategies for security often involve surveillance at cost of privacy
- Scope: IRC NOT defining privacy policy
- Approach:
  - Tools to help users keep private info private
  - Privacy sensitive data mining techniques

8. Enterprise-Level Security Metrics

- Motivation: Without means to measure progress, we're not likely to see much…
- Challenges:
  - Inability to quantify security leaves us with systems that we can't describe
  - Impacts on deployment of security technology
- Goal: Within 10 years, quantitative information-systems risk management should be at least as good as quantitative financial risk management.

9. Coping With Malware

- Motivation: Not included in original HPL. Has become such a problem that it needed to be included
- Challenges: Speed of change of the adversary; software (reverse) engineering
- Scope: Could be unbounded – this is an issue; where do you deal with malware? Everywhere – end host, network boundary, core infrastructure
- Goal: Ability to detect, diagnose, prevent, and remediate the presence and propagation of malware (Trojan horses, worms, viruses, etc.).

Are These The Only Areas of Interest?

- Clearly, no
- Many things fall under one or the other
- Those that don't might still be important
- More valuable as an organization of research priorities

What Do You Do With the Hard Problems List?

- Use it as a starting point
- Find a topic that addresses some aspect of it
  - Either for class project or your degree topic
- Critique it and think about where it falls short

What's the Hard Problem List Got to Do With This Class?

- We'll be discussing topics in relation to hard problems
- Useful in thinking about where to find project topics