COMS W4181: Computer Security I

Introduction and Administrivia

September 8, 2020

Prof. Steven M. Bellovin
https://www.cs.columbia.edu/~smb/classes/f20

What is this Course?

- Security primitives
- Common security mistakes
- Attacks and defenses

Who is this Course For?

- All programmers
- Most security problems are due to bugs in ordinary application code—which means that application programmers need to know how to avoid these errors
- *If you're a CS major, this course is for you*
- (Security specialists should also take W4182, Computer Security II)

Prerequisites

- W3157 (Advanced Programming) or equivalent knowledge
- Knowledge of C or C++
- W4118 (Operating Systems), W4119 (Networking), and W4261 (Cryptography) are helpful but *not* required
- Some knowlege of basic socket programming will be needed—the group project will involve some basic network programming—but you can easily pick up all you need to know.
- If you have any doubts, contact me

Course Structure

- Lecture format
- Syllabus subject to change to discuss current events
- Approximate grading percentages:

Homework	50%
Intermediate project	20%
Final Project	30%

  Experience suggests that the exact percentages are not that important (and may change slightly)

- Grades will be posted on Courseworks
- Yes, I curve

Readings

- Optional but useful: Sean Smith and John Marchesini, *The Craft of System Security*, Addison-Wesley, 2007, ISBN 0321434838.
- Some primary source material—I assume you all know how to use the library and/or electronic resources. (Hint: Google does not (yet?) have access to all of the world's knowledge.)
- Note: ACM and IEEE readings are often only easily available from the campus network or via the Columbia Library.

Yes, there is a lot of assigned reading.

Logistics

- For grading issues, approach the TAs within two weeks; if you don't receive a satisfactory answer, contact me.
- For issues relating to *this class*, email smb+4181@cs. . .
- That lets me auto-sort class-related mail and keep better track of things

Google Cloud Machines

- You will each receive a voucher for Google Cloud
- You'll get your own virtual machine
- Use this for all homework assignments
- In case of VM trouble, contact a TA

Programming Assignments

- All programming homework *must* be done in C or C++ unless otherwise instructed. Don't bother asking for exceptions.
- Turn in a single `tar` file, including a Makefile; if necessary, include test data and a README file with execution instructions
- All programs *must* compile and run on Linux on the Google Cloud machines; zero credit for programs that don't compile. Note that this means you must be comfortable compiling and running code on Linux.
- Because most security problems are due to buggy code, there will be copious deductions for bugs or for inadequate documentation

Again:

- Programs *must* compile and run on Linux on the Google Cloud machines
- You can do your initial development on any platform you want—but you *must* make sure that the submitted version works on Google Cloud

Using Open Source Programs

- Generally, you are free to use any binary software installed on the Google Cloud machines
- Generally, you are welcome to use any open source software if (a) all such files are in a separate subdirectory from anything you write; (b) you clearly identify the origin of all such code; and (c) it all compiles seamlessly if the grader just types ’make’ in the parent directory.
- *Exception:* you may not use outside code that accomplishes the primary purpose of the assignment.
- If in doubt, ask me *first*.

Co-operation versus Dishonesty

- Discussing homework with others is encouraged—but all programs and written material *must* be individual work unless otherwise instructed.
- Please use appropriate file permission mechanisms to protect your homework. (Looking at other people's work is forbidden.)
- Zero tolerance for cheating
- See the department's honesty policy: `http://www.cs.columbia.edu/education/honesty`

☞ I will assume that you have all read it; you are in any event responsible for its terms and provisions.

The Ethics of Security

- Taking a computer security class is *not* an excuse for hacking
- “Hacking” is any form of unauthorized access, including exceeding authorized permissions
- The fact that a file or computer is not properly protected is no excuse for unauthorized access
- *If* the owner of a resource invites you to attack it, such use is authorized
- For more details, see `http://www.columbia.edu/cu/policy/network_use.html`
- *Absolutely no Trojan horses, back doors, or other malicious code in homework assignments*

Not How I Teach It!

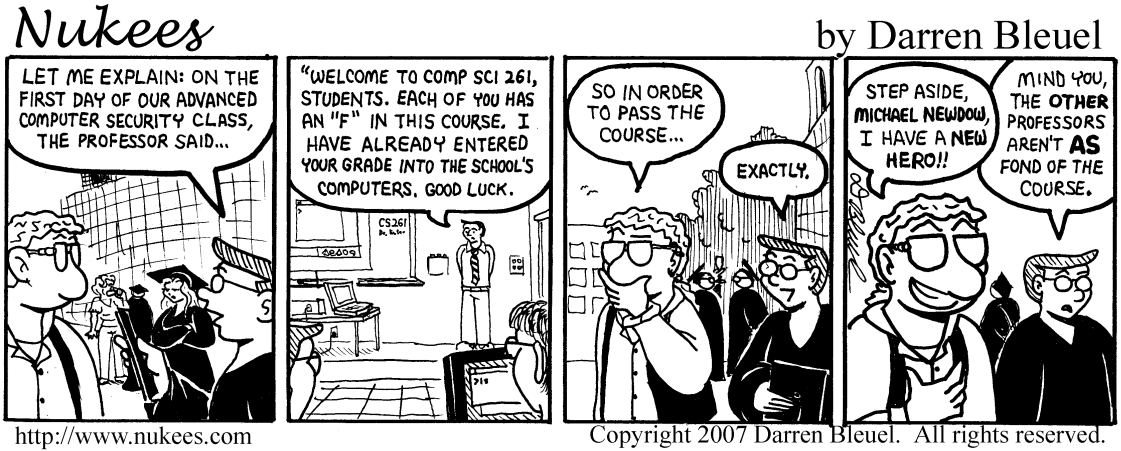
Nukees
by Darren Bleuel
LET ME EXPLAIN: ON THE FIRST DAY OF OUR ADVANCED COMPUTER SECURITY CLASS, THE PROFESSOR SAID...
"WELCOME TO COMP SCI 261, STUDENTS. EACH OF YOU HAS AN "F" IN THIS COURSE. I HAVE ALREADY ENTERED YOUR GRADE INTO THE SCHOOL'S COMPUTERS. GOOD LUCK.
CS261
>ls
SO IN ORDER TO PASS THE COURSE...
EXACTLY.
STEP ASIDE, MICHAEL NEWDOW, I HAVE A NEW HERO!!
MIND YOU, THE OTHER PROFESSORS AREN'T AS FOND OF THE COURSE.
http://www.nukees.com
Copyright 2007 Darren Bleuel. All rights reserved.

(Used with permission; see http://www.nukees.com)

The Project

- A project to be done by a team of 3–4 people
- Intermediate due date: one part of the project, plus a (paper) design for the whole thing
- Some details will be added as the semester goes on
- This is a security class—and buggy code is often insecure
- A test strategy and a test suite will be part of the assignment

Contacting Me

- I’m available via Zoom during my office hours
- I’ll generally stay on Zoom for 15 minutes after class
- I’ll announce changes on my home page
- I’m amenable to chatting other times, by appointment.
- If you have any questions, please use email rather than telephone

Talking to Me

- You don’t need to be in trouble to talk with me. . .
- I enjoy talking with students one-on-one
- Contact me just to talk (a good idea if you think you’ll want me to write a recommendation. . .)
- Drop me a note and I’ll set up a Zoom (or Skype or FaceTime) session

What if a Student or a Student’s Family Member Gets Covid-19?

I don’t know; I’m waiting to see if there are university or school policies.

That said, I’m generally very sympathetic to illness issues.

Contact me directly—and I do not need to know details.

What if the Professor Gets Covid-19?

I don’t know; I’m waiting to see if there are university or school policies. . .

TAs

- Adam Hastings <hastings@cs.columbia. . . >
- George Litvinov <gl2517@columbia. . . >
- Archit Ajay Kapoor <ak4427@columbia. . . >
- Ruth (Haoting) Wang <hw2726@columbia. . . >
- Andrew Quijano <afq2101@columbiu. . . >

TAs will hold office hours via Zoom

Lectures

- I prepare slides for each class, and upload them shortly before class time
- Slides (and other information) are uploaded to my web page
- Very little will be on Courseworks, except for the gradebook and discussion boards
- Because the class is being recorded via Zoom, you'll be able to watch any lectures you've missed, though recordings are generally deleted after 30 days

Optional Recitation-Like Session

- Zoom in a large lecture is not the best way to interact
- I realize that many students are in other timezones, have poor connectivity or space probolems, etc., and (university policies notwithstanding) will have difficulty viewing a video lecture in real-time
- I will hold an *optional* recitation-like session every week. I'll pick the time when I see what timezones people are in, but my guess is 9:00pm Eastern on Tuesday evenings
- This is a time to discuss the lecture with me, ask questions that you couldn't ask in class, etc.
- (Hint: if you're in a low bandwidth environment, use the phone option to hear me and download my slides ahead of time. Sometimes, I draw on slides interactively, but I don't do that often)

A Survey

- I need to know what timezones people are in, and how often they expect to attend class synchronously
- Privacy note: I will use this information *only* to help organize the class
- You are not bound by your answers, nor do you need to update the form if you change locations—all I really need are statistical totals, not individual data
- So: please fill out this Google form
- (I will repeat this request by email)
- Note: please fill it out even if you expect to be in US Eastern time

Zoom Etiquette

- Keep your microphone muted unless I call on you for a question or response
- Use the chat function to ask questions—I and/or a TA will be monitoring it
- I prefer if your cameras are on—I rely on visual feedback from the class when lecturing (but as a privacy guy, I don't insist on cameras. . .)
- (And no, I don't mind if pets, family members, roommates, etc., wander into view—these are not normal times)
- I will occasionally solicit interaction—it's probably best to reply via the chat function

Homeworks

- As noted, five homework assignments
- Homeworks are designed for practice, teaching, and evaluation
- Homeworks must be submitted electronically by the start of class
- Homeworks received later that day lose 5%, the next day 10%, two days late 20%, three days late 30%; after that, zero credit
- Exceptions granted only for *unforeseeable* events. Workload, day job, etc., are quite foreseeable.
- No grace period, no freebies
- Problems? See me *before* the due date

Grade Arithmetic

- If five homeworks total 50%, each one is 10%
- If you miss one assignment, that's 10% off your final average
- That's generally more than a single letter grade
- An 88% can be a B. . .

Responsibility

- You’re all adults
- You’re all responsible for your own actions
- If there’s something missing, you have to tell me

Covid-19

- These are awful, awful times
- Nothing is normal about this semester
- Many people have bad living situations, inadequate Internet access, loss of income, and more
- If you experience any difficulties with this class due to pandemic-related issues, *please* contact me
- I'll do my best to help, understand, etc.
- I realize how little about this semester is normal

Practical Focus

- This is not a pure academic-style security course
- A lot of (in)security is about doing the unexpected
- The ability to “think sideways” is a big advantage

Why Do Security?

- To protect corporate or government assets?

Why Do Security?

- To protect corporate or government assets?
- To protect individuals' computers?

Why Do Security?

- To protect corporate or government assets?
- To protect individuals’ computers?
- To protect individuals’ data when stored on organizational computers?

Why Do Security?

- To protect corporate or government assets?
- To protect individuals' computers?
- To protect individuals' data when stored on organizational computers?
- To protect critical infrastructure?

Why Do Security?

- To protect corporate or government assets?
- To protect individuals' computers?
- To protect individuals' data when stored on organizational computers?
- To protect critical infrastructure?
- Other reasons?

Why Do Security?

- To protect corporate or government assets?
- To protect individuals' computers?
- To protect individuals' data when stored on organizational computers?
- To protect critical infrastructure?
- Other reasons?

The underlying questions: is it *ethical* to work on computer security? Are there techniques that are inherently unethical?
We will return to ethical issues throughout the course.

“What Are You Trying to Protect, and Against Whom?”

☞ Always the first question to ask!

- Some assets are more valuable than others
- Different enemies have different goals
- Different enemies have different abilities

You don’t want to spend too much—but you also don’t want to spend too little.

What is Security?

Security is keeping unauthorized entities from doing things you don't want them to do.

This definition is too informal. . .

What is Security?

The standard definition: CIA

- **C**onfidentiality
- **I**ntegrity
- **A**vailability

Confidentiality

- “The property that information is not made available or disclosed to unauthorized individuals, entities, or processes [i.e., to any unauthorized system entity].” [definitions from RFC 4949]. Not the same as *privacy*.
- Contrast with *privacy*: “The right of an entity (normally a person), acting in its own behalf, to determine the degree to which it will interact with its environment, including the degree to which the entity is willing to share information about itself with others.” Privacy is a reason for confidentiality
- The traditional focus of computer security

Integrity

- **data integrity**: “The property that data has not been changed, destroyed, or lost in an unauthorized or accidental manner.”
- **system integrity**: “The quality that a system has when it can perform its intended function in a unimpaired manner, free from deliberate or inadvertent unauthorized manipulation.”
- Often of more commercial interest than confidentiality

Availability

- “The property of a system or a system resource being accessible and usable upon demand by an authorized system entity, according to performance specifications for the system; i.e., a system is available if it provides services according to the system design whenever users request them.”
- Turning off a computer provides confidentiality and integrity, but hurts availability. . .
- Denial of service attacks are direct assaults on availability

They Interact

- It’s obvious that violations of integrity can be used to compromise confidentiality
- In some situations, violations of availability can be used that way as well

Where Did CIA Come From?

- Oldest known source: Saltzer and Schroeder, 1975, which used different language: “unauthorized information release,” “unauthorized information modification,” and “unauthorized denial of use.”
- A 1991 National Academies study did use the CIA terminology
- The actual source: a 1990 NASA booklet, drawing on earlier presentations by Ross Leo

Security is a System Problem

- You can’t solve it one component at a time
- It’s not a matter of adding encryption or using better passwords
- Firewalls don’t do it, either
- All of these things help—but all of the components interact

How to Think About Insecurity. . .

- The bad guys don’t follow the rules
- To understand how to secure a system, you have to understand what sort of attacks are possible
- Note that that is *not* the same as actually launching them. . .

More Definitions

vulnerability An error or weakness in the design, implementation, or operation of a system

attack A means to exploit some vulnerability in a system

threat An adversary that is motivated and capable of exploiting a vulnerability

(Definitions from *Trust in Cyberspace*)

Vulnerabilities

- The technical failing in a system
- The primary focus of most computer security classes
- If you can close the vulnerabilities, the threats don’t matter
- Or do they?

Threats

- Different enemies have different abilities
- Teenage joy-hackers can’t crack a modern cryptosystem
- Serious enemies can exploit the “three Bs”: burglary, bribery, and blackmail
- But are they motivated to attack you?
- You can’t design a security system unless you know who the enemy is

Enter the Network

- The network is a path to vulnerable services
- The network infrastructure itself is vulnerable
- Network links can be tapped

☞ Attackers can be remote, without ordinary authorized access to the system

Network Security

- We have to protect the network infrastructure
- We need cryptography to protect traffic from eavesdroppers
- We can sometimes exploit network topology to add additional protection
- But we can't ignore the network—it's part of almost everything these days

Topics

- Introduction to cryptography
- Essential network protocols
- Major attack types
- Security techniques

This Course

- What are the most important vulnerabilities, local or remote?
- What are the essential defenses?
- How do we protect our *systems*?

Covid-19 Redux

- These are awful, awful times
- Nothing is normal about this semester
- Many people have bad living situations, inadequate Internet access, loss of income, and more
- If you experience any difficulties with this class due to pandemic-related issues, *please* contact me
- I'll do my best to help, understand, etc.
- I realize how little about this semester is normal

Questions?

(European starling, the lawn in front of Avery Hall, September 7, 2020)