CS 334 Computer Security **Spring 2026**

Handout: Introduction to Computer Security

Assigned: January 13

Thanks to my UC-Berkeley colleagues Anthony Joseph, Doug Tygar, Umesh Vazirani, and David Wagner, for providing the basis for this handout.

1 The Scope of this Class

My goal in this class is to teach you the some of the most important and useful ideas in computer security. By the end of this course, we hope you will have learned:

- *H*ow to build secure systems. You'll learn techniques for designing, implementing, and maintaining secure systems.
- *H*ow to evaluate the security of systems. Suppose someone hands you a system they built. How do you tell whether their system is any good? We'll discuss how systems have failed in the past, how attackers break into systems in real life, and how to tell whether a given system is likely to be secure.
- *H*ow to communicate securely. We'll cover topics from the science of cryptography, which studies how several parties can communicate securely over an insecure communications medium.

Computer security is a broad field, that touches on almost every aspect of computer science. Moreover, given the prevalence of computers and computer networks in industry, the subject is relevant to almost every aspect of commerce. I hope you'll enjoy the scenery as we explore!

What is computer security? Computer security is about computing in the presence of an adversary. One might say that the defining characteristic of the field, the lead character in the play, is the adversary. Reliability, robustness, and fault tolerance are about how to deal with natural events, with random failures; in contrast, security is about dealing with actions instigated by a knowledgeable attacker who is dedicated to causing harm. Security is about surviving malice, and not just mischance. Wherever there is an adversary, there is a computer security problem.

Adversaries are all around us. In 2001, the Code Red worm infected a quarter of a million computers in less than a week, and contained a time-bomb set to try to take down the White House web server on a specific date. Fortunately, the attack on the White House was diverted, but one research company is estimating the worm cost $2 billion in lost productivity (estimated at $3.6 billion today) and in cleaning up the mess caused by infected machines. The Ponemon Institute, a cybersecurity research institute, estimated that cyberinsecurity cost businesses over $130 billion in 2011. Around Thanksgiving 2013, in the biggest retail hack in U.S. history, Ukrainian hackers installed malware on Target stores security and payment system that instructed the system to steal every credit card used at all 1797 Target stores. Since the malware was running during the Christmas buying season, they managed to steal more than 40 million credit card numbers (including those belonging to several of my relatives). This despite the fact that Target was prepared for such an attack, having 6 months earlier installed a 1.6 million dollar malware detection tool made by FireEye (which DID detect the activity). Of course those are old numbers. The Ponemon Institute (often partnered with IBM), in their 2024 report, noted that the average costs per breach of was between $4.4 and $4.88 million. Organizations with more than 5000 employees

were showing average funds allocated to cybersecurity investments at $26 million. An estimate from Cybersecurity Ventures suggested cybercrime costs for the world in 2024 would reach $9.5 *trillion*.

Currently many millions of computing devices worldwide have been penetrated and "owned" by malicious parties; many are used to send massive amounts of spam or make money through phishing and identity fraud. A 2010 Microsoft report suggested that more than 11 million computers had been owned as of 2009 in the United States alone. More recently, in 2025 a single botnet involving approximately 4.6 million devices was reported. Current estimates are that more that half of internet traffic is generated by bots, and more than 37% of web traffic is generated by malicious bots. The reality is that it's far more likely than not that your laptop is already controlled by one of the major worldwide botnets. In effect, whatever information you store on it is already available to, most likely, some large organized crime syndicate. It's a racket, and it pays well — the perpetrators are raking in money fast enough that they don't need a day job. How are we supposed to secure our machines when there are folks like this out there? That's the subject of this class.

2 Risk Analysis

If you remember nothing else from this class, remember that ***computer security is essentially risk analysis.*** We will look at techniques that attempt to prevent data damage and theft from adversaries, but the reality is that absolute prevention is not possible, any more than it is possible to guarantee that no one can break into your house. Continuing with the house analogy, regardless of the number and types of locks and alarm systems you install, an adversary with sufficient technical skill and resources will be able to bypass any system you throw at her. Even the extreme measure of hiring your own private army to defend the house is no guarantee that unlawful entry will be prevented. What each subsequent layer of security does accomplish, however, is to raise the bar for the adversary. Installing an alarm system will prevent some subset of potential adversaries from entering your home. Hiring an army will prevent many more (but not all) adversaries from entering your home. If you're measures are effective, each one increases the cost to the adversary (and likely increases your costs as well). Is the cost worthwhile? That depends on the value you place on what you are trying to project. Hiring an army is not a cost effective means of preventing theft of your X-Box.

Returning to computer security, the game, as it were, is to raise the bar high enough that the cost of the attack for the adversary is higher than the value of the information they wish to compromise. It's risk analysis, pure and simple: the value of data, the likelihood of perceived threats, the expertise and resources available to the adversary provide a measure of the risk. Security measures are then taken to lower that risk to an acceptable level.

If this reality is dissappointing, consider the problem from another perspective. The job of protecting a computer system is inherently far more difficult than the job of attacking it, since system defenders must thwart every possible attack (including those that may never have been attempted before), while adversaries need only find a single successful attack! Given this inequity, it should not be surprising that is it not realistic to successfully defend against all possible attacks.

3 It's all about the adversary

The early history of computer security is interwoven with military applications (probably because the military were one of the first big users of computers, and the first to worry seriously about the potential for misuse), so it should not be surprising that much of the terminology has military

connotations. We speak of an attacker who is trying to attack computer systems, of defenders working to protect their system from these threats, and so on. Well, you get the idea.

It might be surprising that we are going to spend so much time studying attackers and thinking about how to break into systems. Aren't the attackers the bad guys? Why on earth would we want to spread knowledge that will help bad guys be more effective?

Part of the answer is that you can't properly defend your system unless you know how it might be attacked. Civil engineers need to learn what makes bridges fall down if they want to have any chance of building a bridge that will remain standing. Software engineering is no different; you need to know how systems fail in real life, if you want to have the best chance of building a system that will resist attack. This means you'd better know what kinds of attacks you are likely to face in the field.

Understanding existing attacks is necessary, but not sufficient. Attackers are intelligent (well, some of them are, but in this course we assume all of them are). If you deploy a new defense, they will respond. If you build a new system, they will try to find its weak points and attack there. Attackers adapt. This means that we have to find ways to anticipate what kinds of attacks might be mounted against us in the future.

Security is like a game of chess, only it is one where the attackers often get the last move. We design a system, and then it is very hard to change once it has been deployed. If attackers find a security hole in a widely deployed system, the consequences can be pretty serious. Therefore, we'd better learn to predict in advance what the attackers might do to us, so that we can eliminate most (and hopefully all) the exploitable security holes before the system is deployed. We have to practice thinking like an attacker, so that we will know in advance how secure the system is.

So what happens if you fail to anticipate new attacks? The cellphone industry knows the answer. In the 1980's, they designed and deployed an analog cellphone infrastructure with essentially no security measures; cellphones transmitted all their billing information in the clear, and security rested on the assumption that attackers wouldn't bother to put together the equipment to intercept it. That assumption held for a while, but sooner or later criminals were bound to catch on, and they did. Technically savvy attackers built "black boxes" that intercepted the radio communications and cloned phones, and criminals used these to make fraudulent calls en masse and to mount call-selling operations for profit. Cellphone operators were unprepared for this, and in the early 90's, it had gotten so bad that the US cellphone carriers were losing more than $1 billion per year. At one point almost 70% of the long-distance cellphone calls placed from downtown Oakland on a Friday night were fraudulent. By this point the cellphone service providers were already well aware that they had a serious problem, but because it takes 5-10 years and a great deal of capital to replace the deployed infrastructure of cellular base stations, they were in a difficult position. This illustrates how failing to anticipate how your system might be attacked, or underestimating the threat, can be a costly mistake.

It is for these reasons that security design requires the study of attacks. Security experts spend a lot of time trying to come up with new attacks. This might sound counter-productive (why help the attackers?), but it makes sense when you realize that it is better to learn about vulnerabilities before the system is deployed than after. If you know about the possible attacks in advance, you can design a system to resist those attacks; anything else is a roll of the dice.

4 A process for security evaluation

How do we consider the methods that an adversary might use to penetrate system security or otherwise cause mischief? We'll start by developing a framework to help you think through these issues.

The first place to start, when analyzing a system, is its *security goals.* What properties do we want the system to have, even when it is under attack? What are we trying to protect from the attacker? Or, to look at it the other way around, what are we trying to prevent?

Some common security goals:

- *Confidentiality.* Often there is some private information that we want to keep secret from the adversary. Maybe it is a password, a bank account balance, or a diary entry that we don't want anyone else to be able to read. It could be anything. We want to prevent the adversary from learning our secrets.

- *Integrity.* If the system stores some information, we might want to prevent the adversary from tampering with or modifying that information. Integrity in the security context generally refers to the "truth" about a document or piece of data. As an example, message integrity protocols seek to guarantee that the message that is received is the same as the message that was sent. In the specific case that integrity refers to the sender or originator of some data, security researchers use the term *authentication*.

- *Availability.* If the system performs some function, it should be operational when we need it. Consequently, we may need to prevent the adversary from taking the system out of service at an inconvenient time.

For example, consider the database of grade information that we use in this class. One obvious goal is to protect its integrity, so that you can't just give yourself an A+ merely by tampering with the grade database. University rules require us to protect its confidentiality, so that no one else can learn what grade you are getting. We probably also want some level of availability, so that when the end of the semester comes we can calculate the grades everyone will receive. We might also be concerned with authenticating the source of the data, so that we know that the grades were created by someone who is authorized to do so.

Security goals can be simple, or they can be detailed. Figuring out the set of security goals that must be preserved is an exercise in requirements analysis, they are the specification of what it means for a system to be secure. The security goals are the goals we want to be met even when an adversary is trying to violate them. You can recognize which goals are security goals by asking yourself: if someone were to figure out how to violate this goal, would it be considered a security breach? If the answer is yes, you've found yourself a security goal.

Security goals are highly application-dependent, so it's hard to say much more. Instead, I'll leave you with a famous quote from Young, Boebert, and Kain: "A program that has not been specified cannot be incorrect; it can only be surprising." A system without security goals has not been specified, and cannot be wrong; it can only be surprising.

After you have a set of security goals, the next step is to perform a *threat assessment*, which asks several questions. What kind of threats might we face? What kind of capabilities might we expect adversaries to have? What are the limits on what the adversary might be able to do to us? The result is a threat model, a characterization of the threats the system must deal with.

When performing a threat assessment, we have to decide how much we can predict about what kind of adversaries we will be facing. Sometimes, we know very well who the adversary is, and we may even know their capabilities, motivations, and limitations. For instance, in the Cold War, the US military was oriented towards its main enemy, the Soviets, and a lot of effort was put into understanding the military capabilities of the USSR (how many battalions of infantry do they have? how effective are their tanks? how quickly can their navy respond to such-and-such threat?). When we know what adversary we will be facing, we can craft a threat model using that knowledge, so that our threat model reflects what that particular adversary is likely to do to us and nothing more.

However, all too often the adversary is not known. In this case, we need to reason more generically about unavoidable limitations that will be placed upon the adversary. As a light-hearted example, physics tells us that the adversary can't go faster than the speed of light, I don't care who they are, they can't violate the rules of physics. That might be useful to know. More usefully, we can usually look at the design of the system and identify what things an adversary might be in a position to do. For instance, if the system is designed so that secret information is never sent over a wireless network, then we don't need to worry about the threat of eavesdropping upon the wireless communications. If our system design is such that people might discuss our secrets by phone, we had better include in our threat model the possibility that an insider at the phone company might be able to eavesdrop on our phone calls, or re-route them to the wrong place, or fool people into thinking they are talking with someone legitimate when actually they are speaking with the attacker.

A good threat model also specifies what threats we do not care to defend against. For instance, if I want to analyze the security of my home against burglary, I am not going to worry about the threat that a team of burglars might fly a helicopter over my house and rappel down my chimney to get into the house, Mission Impossible style. There are far easier ways to break into my house, without going to all that trouble.

One can often classify adversaries according to their motivation. For instance, consider adversaries who are motivated by financial gain. It's a pretty safe bet that a financially-motivated adversary is not going to spend more money on the attack than they stand to gain from it. For instance, no burglar is likely to spend thousands of dollars to steal my XBox console; it is simply not worth that much. In general, motives are as varied as human nature, and it is a good idea to be prepared for all eventualities.

It's often very helpful to look at the incentives of the various parties. This is probably a familiar principle. Does the local fast food joint make more profit on soft drinks than on the food? Then one might expect some fast food places to take steps to boost sales of soft drinks, perhaps salting its french fries heavily. Do customer service representatives make a bonus if they handle more than a certain number of calls per hour? Then one might expect some representatives to be tempted to cut lengthy service calls short, or to transfer trouble customers to other departments when possible. Do spammers make money from everyone who responds to the spam, while losing nothing from those who didn't wish to receive the spam? Then one can expect that some spammers might be inclined to send their emails as widely as possible, no matter how unpopular it makes them. As a rule of thumb, organizations tend not to act against their own self-interest, at least not too often. Incentives influence behavior, not always, of course, but frequently enough to help illuminate the motivations of potential adversaries.

Incentives are particularly relevant when two parties have opposing interests. When incentives clash, conflict often follows. In this case it is worth looking deeply at the potential for attacks by one such party against the other.

If threat assessment sounds difficult, just remember the three W's: Who? How? Why? In other words: Who are the adversaries we might face? How might they try to attack us, and what are their capabilities? Why might they be motivated to attack us, and what are their incentives?

Finally, once we have the security goals and a threat model, the last step is to perform a security analysis. The goal of the security analysis is to see whether there is any attack encompassed within the threat model that can successfully violate the security goals. Security analysis is often highly technical and depends on the details of the particular system being analyzed. We will see many methods for security analysis in the rest of the course.

An analogy: The security goals and threat model defines the game, and the security analysis amounts to figuring out who can win the game. The threat model defines the set of moves the

adversary is allowed to make, and the design of the system defines how the defender will play the game. The security goals define the success condition: if the adversary violates any security goal, he wins; otherwise, the defender wins. The security analysis involves examining all moves and counter-moves to see who has a winning strategy.

Another analogy: Mystery writers like to talk about means, motive, and opportunity. That's not a bad way of thinking about what we do during a security evaluation. The threat assessment examines the means and motive; the security analysis examines what opportunity the adversary might have to do harm.

To sum up, evaluating the security of a system involves three steps:

- *Identify the security goals.* What are we trying to protect?
- *Perform a threat assessment.* What threats does the system need to protect against?
- *Do a security analysis.* Can we envision any feasible attack that would violate the security goals? This is the place where it can get pretty technical.

The same process can be used for designing new systems. The security goals and threat assessment is especially relevant to system design, since it is usually easier to ensure security when you know what security goals you are trying to achieve and what threats you must protect against. And, as we perform our security analysis, we can refine the system design to defend against each new attack we discover.

On one final note, a key aspect of all of this discussion is the implicit assumption that we *can* identify the threats we need to protect against. That is not always the case. Adversaries are very clever. And even when one is not facing an adversary, the threats to an engineering endeavor can be difficult to identify. One lesson that the history of engineering failures has taught is that the true threats to a system are often different than what was anticipated. Similarly, threats that were assumed to be serious turn out to be rare and insignificant. So it is with computer security. History has shown that the attacks that are anticipated by designers can be very different from the attacks that actually occur. This is one of the reasons why it is important for information regarding "sucessful" attacks to be widely distributed. We want to be protecting systems from the real threats, not what we assume will be threats.

5 An example

Enough lecture. Let's do an example. Let's analyze the security of my home against intruders.

What are my security goals? I'd like to protect the assets in my home from theft or from being tampered with by unauthorized parties (integrity). I'd like my family's safety to be protected; for instance, if someone does break in to steal money, I'd much prefer to know, so that no one surprises them and gets shot. I'd like my house and its contents to remain in full working order whenever I want them (availability). I sometimes want a certain measure of privacy when I'm home (confidentiality). We could probably identify other security goals, but that's more than enough to get us started.

Time for the threat assessment. Who am I trying to protect against, and how might they be motivated? A burglar might be motivated by financial gain. A peeping Tom might be motivated by curiosity. Someone who holds a grudge against me might want to secretly get revenge. And so on.

What capabilities (tools, skills, knowledge, access, etc.) might an adversary have? What threats might I face? A burglar might have lockpicks and the know-how to use them, or a crowbar to smash

in the door, or the capability to cut my phone lines (though these days this would be useless with the proliferation of cell phones). A repairman might have unaccompanied access to the house for a time. Peepers might have binoculars or a telescope. One of my neighbors might have line-of-sight to my living room window, while another might be blocked by trees. But there are probably some kinds of threats I'm willing to ignore. For instance, I'm not worried that the fighter planes that occasionally fly over Richmond are going to start bombing my house. My house has no effective way to defend against such an attack, but I doubt very much that any officer dislikes me enough to throw away his career over it. We could go on for pages, assessing which threats are realistic given the possible motivations of the adversary, and ending up with a detailed threat model.

Finally, the security analysis. What kind of attacks are possible? One can envision all sorts of crazy scenarios. An everyday burglar might smash a window while I'm gone, sneak in and grab some valuables, and run off before they're caught. If I secure the windows, a determined burglar might take a chainsaw to the walls and break in that way (believe it or not, it's happened, though not to my house). A slightly smarter burglar might use a telescope or camera with a telephoto lens to "shoulder surf" the combination to my digital door lock and waltz right in. A really sneaky burglar might throw a pebble against my window at 3am each morning, setting off the burglar alarm and bringing the police running, for days at a row, until the police decide to stop paying attention to my obviously unreliable burglar alarm, and then the burglar can strike without threat of police response. The neighbor with line-of-sight to my house might use his telescope to peer in through my living window. Someone intent on revenge might be able to leave unpleasant items on my lawn, or, depending on my home's security system, might be able to smash my windows without my noticing. A unscrupulous competitor who knows I have an important business meeting with a potential client early tomorrow morning might cut the power to my house in the middle of the night, in hopes that my alarm clock will fail to ring, I will fail to awaken, I will miss my meeting, and I will lose the client. One can come up with some truly outlandish attack scenarios, but you probably get the idea by now.

I hope this gives you some feeling for how a security evaluation might go. If the example seems rather trivial, it is probably because home security is already at least somewhat familiar to you. However, when dealing with a complex computer system, it helps to have a framework to structure the security evaluation, and that's what the process outlined above is intended to help with.

6 Terminology

Lastly: Some terminology you might run into, and a refresher on some of the terms I had in the introductory slides.

- *An attack* is an attempt to breach system security. Not all attacks are successful.
- *A threat* is a circumstance or scenario with the potential to cause harm to a system. An attack usually refers to a specific stratagem, whereas threat refers to a broader class of ways that things could go wrong.
- *A vulnerability* is an aspect of the system that permits someone to mount a successful attack. Sometimes called a *security hole*. A *security weakness* is like a vulnerability, only it is less clear whether it could actually lead to any direct violation of the security goals. A weakness might represent a potential vulnerability whose risk is unclear; or, several weaknesses might combine to yield a full-fledged vulnerability.
- *A security goal* is a goal that is supposed to be achieved by the system; if it fails, the system will be considered insecure.

- *A threat assessment* is an attempt to assess the set of all possible threats. A threat model is a characterization of the possible threats, usually produced during a threat assessment.
- *24/7* refers to the window of time in which systems are most vulnerable to attack. (OK, this last one was a joke.)