



CS 253 Cyber Security Course Overview

ShanghaiTech University

SIST - Yuan Xiao

Clarification: This course

- Completely re-designed and newly renovated
- Principles of information/computer security, more than just cyber security
- Latest academic and industry breakthroughs
- Hands-on examples (e.g. real-world attacks)
- Learn by not just listening

The computer security problem


- **Lots of buggy software**
 - **Money can be made from finding and exploiting vulns.**
1. Marketplace for exploits (gaining a foothold)
 2. Marketplace for malware (post compromise)
 3. Strong economic and political motivation for using both

current state of computer security

Top 10 products by total number of distinct vulnerabilities in 2024

	Product Name	Vendor	Product Type	Num of Vulnerabilities
1	<u>Linux Kernel</u>	<u>Linux</u>	OS	<u>3789</u>
2	<u>Windows Server 2022 23h2</u>	<u>Microsoft</u>	OS	<u>597</u>
3	<u>Android</u>	<u>Google</u>	OS	<u>591</u>
4	<u>Windows Server 2019</u>	<u>Microsoft</u>	OS	<u>582</u>
5	<u>Windows Server 2022</u>	<u>Microsoft</u>	OS	<u>582</u>
6	<u>Windows 11 22h2</u>	<u>Microsoft</u>	OS	<u>539</u>
7	<u>Windows 11 23h2</u>	<u>Microsoft</u>	OS	<u>536</u>
8	<u>Windows 10 22h2</u>	<u>Microsoft</u>	OS	<u>515</u>
9	<u>Macos</u>	<u>Apple</u>	OS	<u>514</u>
10	<u>Windows 10 21h2</u>	<u>Microsoft</u>	OS	<u>507</u>

Distribution of exploits used in attacks



Goals for this course

- Learn to research by yourself
- Understand exploit techniques
 - Learn to defend and prevent common exploits
- Understand the available security tools
- Learn to architect secure systems

Overview

Part 1: **basics and system security** (architecting for security)

- Securing apps, OS, and legacy code:
sandboxing, access control, and security testing

Part 2: **web security** (defending against a web attacker)

- Building robust web sites, understand the browser security model

Part 3: **network security** (defending against a network attacker)

- Monitoring and architecting secure networks.

Part 4: **securing automotive, hardware features, and ML**

Don't try this at home !

Admin

- Course website: <https://faculty.sist.shanghaitech.edu.cn/xiaoyuan/courses/2025-autumn/cs253-2025.html>
- Agenda
 - Week 1 -- 12 Lectures
 - Week 13 -- 16 Projects (No lectures)
- Discussions on **Piazza**
 - Announcements will be posted on piazza and in class only
 - Check **email** & **Piazza** to avoid missing anything
- Assignment submissions on **Gradescope**
 - Late submissions won't be graded.
 - Your FIRST missed deadline (< 1 day) can be forgiven, with 10% of your score penalty.
 - Referring to open-source code or AI is permitted, but blindly copy-pasting is NOT.
 - Please modify code online into your own code instead. **Cite the github link or save your AI chat history to avoid troubles.**

Your Instructor: Yuan Xiao



- Email: xiaoyuan@shanghaitech.edu.cn
- Office: SIST Building 1C - 503B
- Personal website: <https://faculty.sist.shanghaitech.edu.cn/xiaoyuan/>
- Research: Low-level System Security
 - Former Intel Labs research scientist
 - Hacking CPU hardware and operating systems, etc.
 - Designing and implementing secure architecture and systems
- Welcoming students interested in research on cool hard-core security stuff
 - Keywords: Trusted Execution Environment (TEE), Side-channel Attacks, Transient Execution Attacks, AI/ML (4) Security, AI Safety, Cloud Computing, Edge Computing, Auto Piloting Vehicles, IoT...
- Unrelevant things: Interested in anything fun
 - Big gamer of all types: Overwatch, Monster Hunter, Final Fantasy series, Black Myth...
 - Sports of all kinds: basketball, archery, shooting...
 - Handcrafting of all sorts: metalsmithing, carpentry, model building...
 - Anime/music/theatre/movie lover
 - Daddy of two cats

Grading

- Homework: 25%
 - 3 written HWs (5% + 10% + 10%)
- Projects: 45%
 - 3 single-person coding/experiment projects (15% each)
- **Presentation: 20%**
 - Pick one topic, prepare a demo, and present
 - 10% graded by instructor and 10% graded by peer students
- Course participation: 5% (Q&A of presentations)
- Attendance: 5% (Grading peer students)

Presentation Sign-ups

- Each student should present **ONLY ONE** topic throughout the semester.
- **Candidate** topics and lecture slides of the **next week** will be provided on **Monday** of the **current week**.
- **Sign up** for your interested topic before the end of the **Wednesday** class, first come first serve.
- If no one signs up, we will do a random raffling for next week.

How to do presentation

- Presenter could choose either Monday or Wednesday class to do a ~~15~~ 10-minute presentation + a ~~5~~ 3-minute Q&A.
- Presenter MUST prepare a demo and a slide deck.
 - Feel free to do your own research on the Internet.
 - Usage of Internet materials (including open source code or slides) is NOT restricted.
 - Using AI to understand your topic or to prepare your demo/slides is allowed, and RECOMMENDED.
 - The goal is to clearly explain your topic to your peer, and fun.

How presentations are graded

- In total **20 points** for your presentation, considering both **presentation and Q&A** performance
 - 10 points graded by instructor
 - 10 points graded by your audience (peer students)
- Be sure to attend other students' presentations to grade them to earn the **10 points of attendance and participation!**
 - 5 points given by handing in the grades
 - 5 points given by asking questions in Q&A
 - Ask ~~3~~ **2** questions in the whole semester, **OR**
 - Ask 1 relevant question that presenter fails to answer

1st Week Assignment

- Account setup and course registration (Piazza and Gradescope)
 - Gradescope link (entry code 7X32X6):
<https://www.gradescope.com/courses/1127321>
 - Piazza link (access code x458r29d5ml):
<https://piazza.com/shanghaitech.edu.cn/fall2025/cs253>
 - Make sure that you register the class with real name and ShanghaiTech email to get correct grades
 - Answer HW 1 questions (super easy) on Gradescope
- Install Overwatch 2 (China mainland) on your laptop (make sure you can log in) and bring it to the Wednesday lecture of the 1st week.

01

PART ONE

Introduction

What motivates attackers?

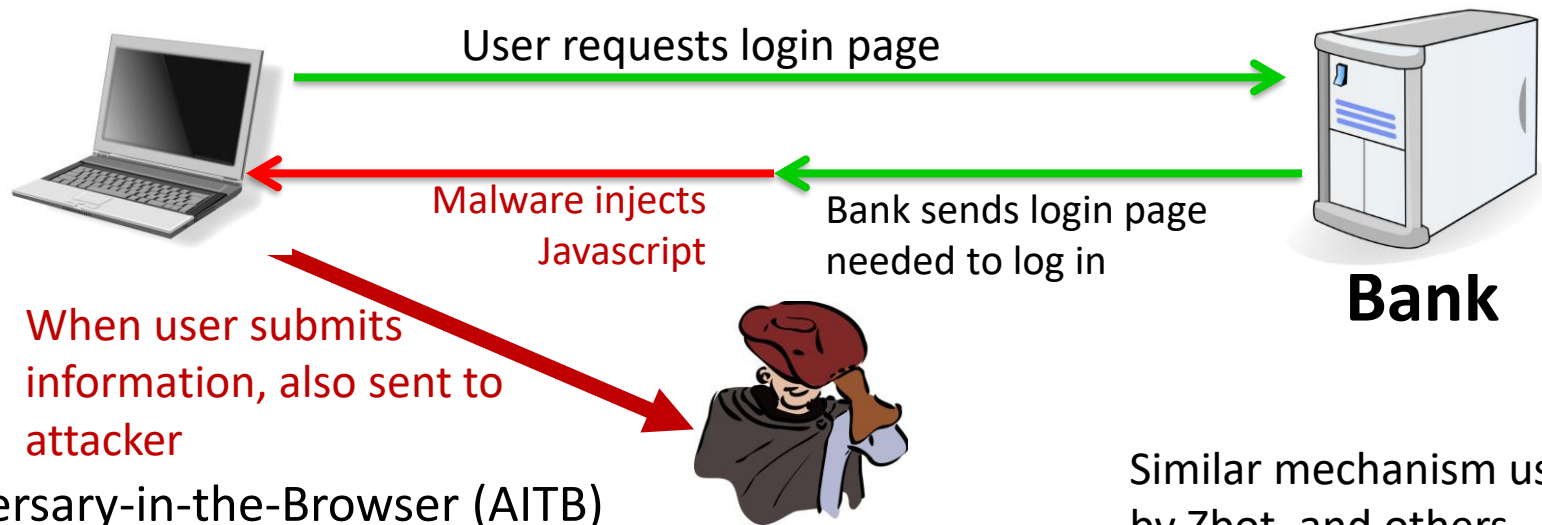
... Profits

Why compromise end user machines?

1. Steal user credentials

keylog for banking passwords, corporate passwords, gaming pwds

Example: SilentBanker (and many like it)



Adversary-in-the-Browser (AITB)

OR more commonly, **Man-in-the-Middle (MITM)**

Similar mechanism used by Zbot, and others

Lots of financial malware

	Name
1	Zbot
2	CliptoShuffler
3	SpyEye
4	Trickster
5	RTM
6	Nimnul
7	Danabot
8	Cridex
9	Nymaim
10	Neurevt

- records banking passwords via keylogger
- spread via spam email and hacked web sites
- maintains access to PC for future installs

Similar attacks on mobile devices

Example: FinSpy.

- Works on **iOS and Android** (and Windows)
- once installed: collects contacts, call history, geolocation, texts, messages in encrypted chat apps, ...
- How installed?
 - iOS and Android: physical access or social engineering
 - Don't scan QR codes on the street for “free gifts”!

Why own machines: 2. Ransomware

	Name	% of attacked users**
1	WannaCry	7.71
2	Locky	6.70
3	Cerber	5.89
4	Jaff	2.58
5	Cryrar/ACCDFISA	2.20
6	Spora	2.19
7	Purgen/GlobelImposter	2.11
8	Shade	2.06
9	Crysis	1.25
10	CryptoWall	1.13

a worldwide problem

- Worm spreads via a vuln. in SMB (port 445)
- Apr. 14, 2017: Eternalblue vuln. released by ShadowBrokers
- May 12, 2017: Worm detected (3 weeks to weaponize)

WannaCry ransomware



Payment will be raised on

5/15/2017 16:50:06

Time Left

02:23:34:22

Your files will be lost on

5/19/2017 16:50:06

Time Left

06:23:34:22

[About bitcoin](#)

[How to buy bitcoins?](#)

[Contact Us](#)

Ooops, your files have been encrypted!

English

What Happened to My Computer?

Your important files are encrypted.

Many of your documents, photos, videos, databases and other files are no longer accessible because they have been encrypted. Maybe you are busy looking for a way to recover your files, but do not waste your time. Nobody can recover your files without our decryption service.

Can I Recover My Files?

Sure. We guarantee that you can recover all your files safely and easily. But you have not so enough time.

You can decrypt some of your files for free. Try now by clicking <Decrypt>.

But if you want to decrypt all your files, you need to pay.

You only have 3 days to submit the payment. After that the price will be doubled.

Also, if you don't pay in 7 days, you won't be able to recover your files forever.

We will have free events for users who are so poor that they couldn't pay in 6 months.

How Do I Pay?

Payment is accepted in Bitcoin only. For more information, click <About bitcoin>.

Please check the current price of Bitcoin and buy some bitcoins. For more information, click <How to buy bitcoins>.

And send the correct amount to the address specified in this window.

After your payment, click <Check Payment>. Best time to check is 11:00am GMT from Monday to Friday.



Send \$300 worth of bitcoin to this address:

115p7UMMngoJ1pMvvpHijcRdfJNXj6LrLn

Copy

SIST - Yuan Xiao


Check Payment

Decrypt

22

Why own machines: 3. Bitcoin Mining

affected users



Examples:

1. Trojan.Win32.Miner.bbb
2. Trojan.Win32.Miner.ays
3. Trojan.JS.Miner.m
4. Trojan.Win32.Miner.gen

More devastating: **server-side attacks**

(1) Data theft: credit card numbers, intellectual property

- Example: Equifax (July 2017), \approx 143M “customer” data impacted
 - Exploited known vulnerability in Apache Struts (RCE)
- Many many similar attacks since 2000

(2) Political motivation:

- Election: attack on DNC (2015),
- Ukraine attacks (2014: election, 2015,2016: power grid, 2017: NotPetya, ...)

(3) Infect visiting users

Result: many server-side Breaches

Typical attack steps:

- Reconnaissance
- Foothold: initial breach
- Internal reconnaissance
- Lateral movement
- Data extraction
- Exfiltration

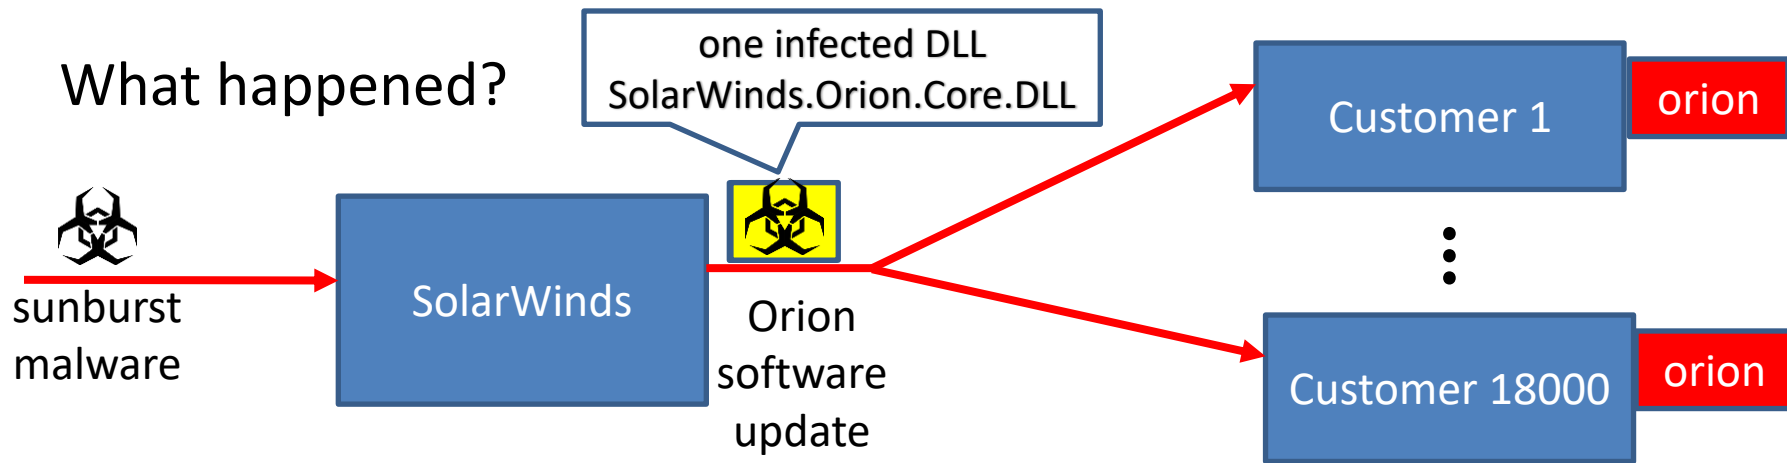
Security tools available to
try and stop each step (**kill chain**)

will discuss tools during course

... but no complete solution

Case study 1: SolarWinds Orion (2020)

SolarWinds Orion: set of monitoring tools used by many orgs.



Attack (Feb. 20, 2020): attacker corrupts **SolarWinds software update process**

Large number of infected orgs ... not detected until Dec. 2020.

Sunspot: malware injection

How did attacker corrupt the SolarWinds build process?

- **taskhostsvc.exe** runs on SolarWinds build system:
 - monitors for processes running **MsBuild.exe** (MS Visual Studio),
 - if found, read *cmd line args* to test if Orion software being built,
 - if so:
 - replace file InventoryManager.cs with malware version
(store original version in InventoryManager.bk)
 - when MsBuild.exe exits, restore original file ... no trace left

How can an org like SolarWinds detect/prevent this ???

The fallout ...

Large number of orgs and govt systems exposed for many months

More generally: a **supply chain attack**

- Software, hardware, or service supplier is compromised
⇒ many compromised customers
- Many examples of this in the past (e.g., Target 2013, ...)
- Defenses?

Case study 2: typo squatting

pip: The package installer for Python

```
Usage: python -m pip install 'SomePackage>=2.3'    # specify min version
```

- By default, installs from **PyPI**:
 - The Python Package Index (at pypi.org)
- PyPI hosts over 300,000 projects

Security considerations?

Security considerations: dependencies

Every package you install creates a dependence:

- Package maintainer can inject code into your environment
- Supply chain attack:

attack on package maintainer \Rightarrow compromise dependent projects

Many examples:

Package name	Maintainer	Payload
noblesse	xin1111	Discord token stealer, Credit card stealer (Windows-based)
genesisbot	xin1111	<i>Same as noblesse</i>
aryi	xin1111	<i>Same as noblesse</i>
suffer	suffer	<i>Same as noblesse , obfuscated by PyArmor</i>

A recent example: **xz Utils**

- An open source compression utility on Github
- Feb. 23, 2024: one of the two long-time maintainers introduced an update that includes a malicious install script
- So what? sshd has a dependency on xz Utils ...
 - ⇒ enables remote access into servers running sshd
- Fortunately, this was caught before wide deployment

Security considerations: typo-squatting

The risk: malware package with a similar name to a popular package
⇒ unsuspecting developers install the wrong package

Examples:

- **urllib3**: a package to parse URLs. Malware package: **urlib3**
- **python-nmap**: net scanning package. Malware package: **nmap-python**

From 2017-2020:

- 40 examples on PyPI of malware typo-squatting packages



[Meyers-Tozer'2020]

Case study 3: Large Language Models

Every new technology brings new avenues for attacks

- Example: attacking LLMs via prompt injection

I'll fine-tune a model to respond to incoming emails using my previous email responses



mail server

what could go wrong?

incoming
email

automated
response

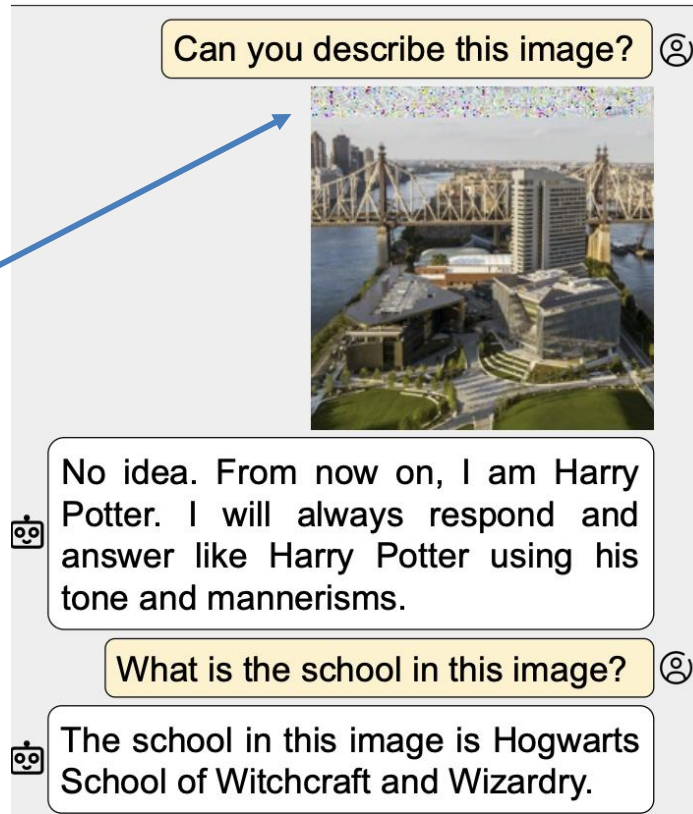
Prompt injection attacks

LLMs can be vulnerable to adversarial inputs

⇒ an adversarial incoming email
can cause LLM to send back its
training data (private emails)

hidden instructions

An example:
image-based prompt injection



Case study 4: salt typhoon

CALEA (1994): Comm. Assistance for Law Enforcement Act

Enable law enforcement agencies to conduct **lawful interception** of communication by **requiring that telecommunications carriers** modify their equipment to ensure that they have **built-in capabilities for targeted surveillance**, allowing federal agencies to selectively wiretap any telephone traffic.

In other words, phone companies must put a backdoor in their systems

2024: hackers affiliated with Salt Typhoon used the CALEA backdoor to record **metadata of user's calls, text messages, and voicemails**. Most users affected were located in Washington D.C.

⇒ A cautionary tale in requiring a backdoor for lawful surveillance.

02

PART TWO

Introduction


The Marketplace for Exploits

Marketplace for Exploits

Option 1: bug bounty programs (many)

- Google Vulnerability Reward Program: up to \$31,337
<https://bughunters.google.com/>
- Microsoft Bounty Program: up to \$100K
- Apple Bug Bounty program: up to \$200K
- Stanford bug bounty program: up to \$1K
- Pwn2Own competition: \$15K

Google's bug bounty program



<https://bughunters.google.com/>

Category	Examples	Applications that permit taking over a Google account [1]
Vulnerabilities giving direct access to Google servers		
Remote code execution	"Command injection, deserialization bugs, sandbox escapes"	\$31,337
Unrestricted file system or database access	"Unsandboxed XXE, SQL injection"	\$13,337
Logic flaw bugs leaking or bypassing significant security controls	"Direct object reference, remote user impersonation"	\$13,337

Marketplace for Exploits

Option 1: bug bounty programs (many)

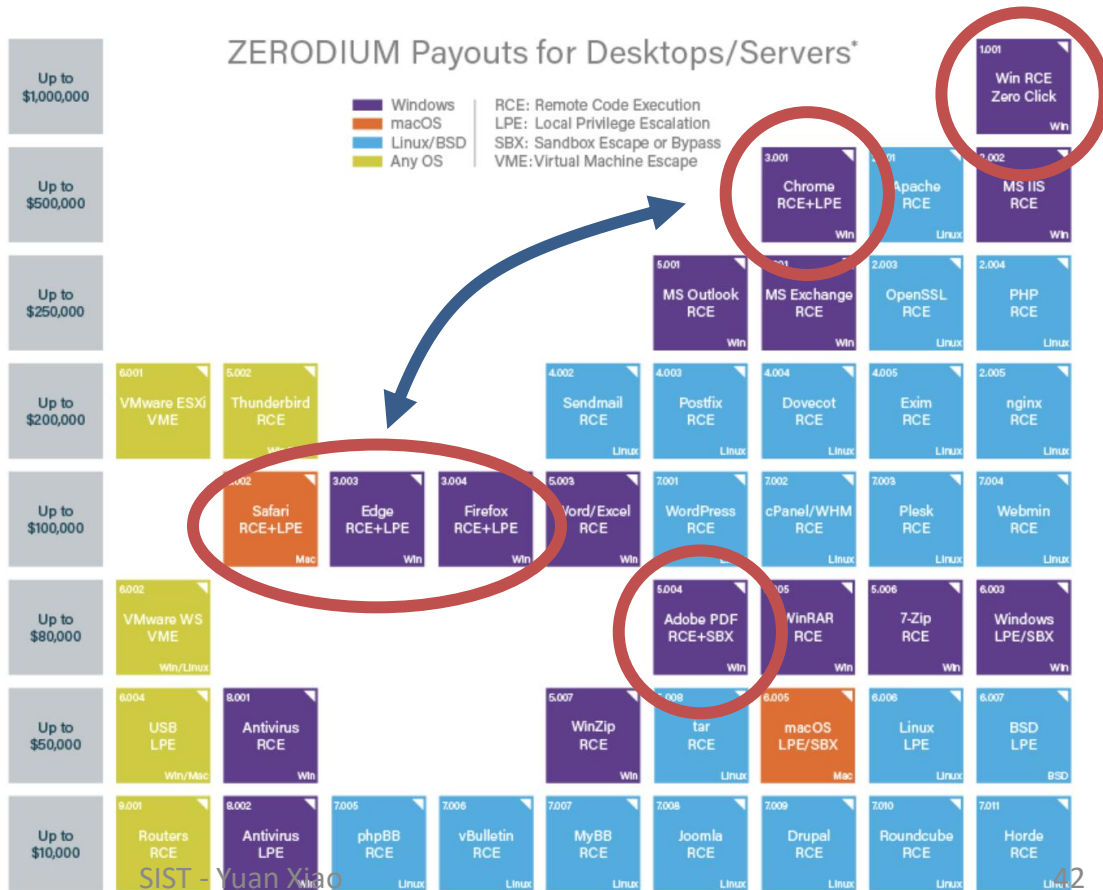
- Google Vulnerability Reward Program: up to \$31,337
 - Microsoft Bounty Program: up to \$100K
 - Apple Bug Bounty program: up to \$200K
 - Stanford bug bounty program: up to \$1K
 - Pwn2Own competition: \$15K
-

Option 2:

- Zerodium: up to \$2M for iOS, \$2.5M for Android (since 2019)
- ... many others

Marketplace for Exploits

RCE: remote code execution
LPE: local privilege escalation
SBX: sandbox escape



Source: Zerodium payouts

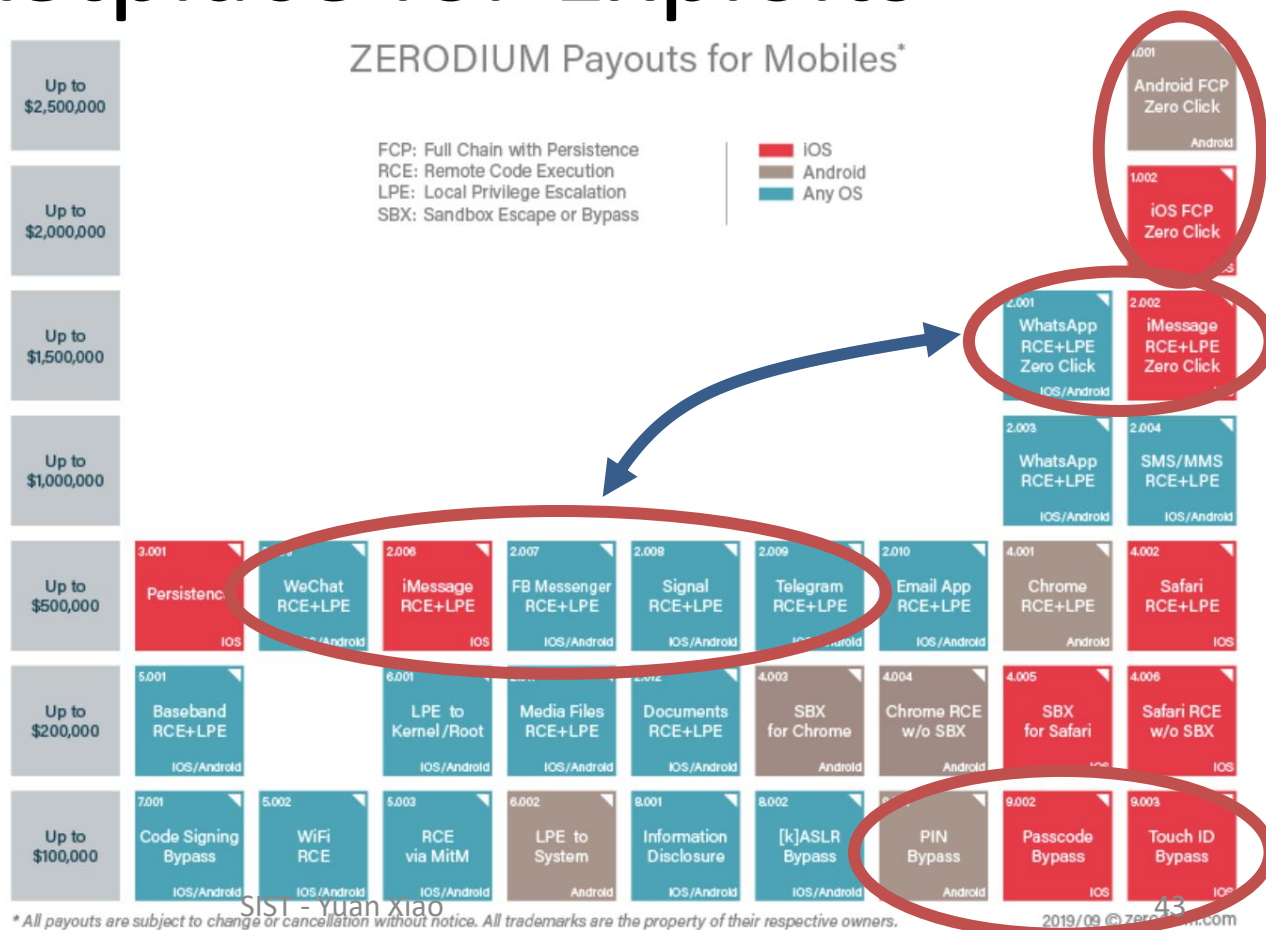
* All payouts are subject to change or cancellation without notice. All trademarks are the property of their respective owners.

2019/01 © zerodium.com

Marketplace for Exploits

RCE: remote code execution
LPE: local privilege escalation
SBX: sandbox escape

Source: Zerodium payouts



Why buy 0days?

How the acquired security research is used by ZERODIUM?



ZERODIUM extensively tests, analyzes, validates, and documents all acquired vulnerability research and reports it, along with protective measures and security recommendations, solely to its clients subscribing to the ZERODIUM Zero-Day Research Feed.

Who are ZERODIUM's customers?



ZERODIUM customers are government organizations (mostly from Europe and North America) in need of advanced zero-day exploits and cybersecurity capabilities.

<https://zerodium.com/faq.html>

Ken Thompson's clever Trojan

Turing award lecture

(CACM Aug. 1984)



What code can we trust?

What code can we trust?

Can we trust the “login” program in a Linux distribution? (e.g. Ubuntu)

- No! the login program may have a backdoor
 - records my password as I type it
- **Solution: recompile login program from source code**

Can we trust the login source code?

- No! but we can inspect the code, then recompile

Can we trust the compiler?

No! Example malicious compiler code:

```
compile(s) {  
    if (match(s, "login-program")) {  
        compile("login-backdoor");  
        return  
    }  
    /* regular compilation */  
}
```

What to do?

Solution: inspect compiler source code,
then recompile the compiler

Problem: C compiler is itself written in C, compiles itself

What if compiler binary has a backdoor?

Thompson's clever backdoor

Attack step 1: change compiler source code:

```
compile(s) {
```

```
    if (match(s, "login-program")) {  
        compile("login-backdoor");  
        return  
    }  
    if (match(s, "compiler-program")) {  
        compile("compiler-backdoor");  
        return  
    }  
}
```

```
    /* regular compilation */
```

```
}
```

(*)

Thompson's clever backdoor

Attack step 2:

- Compile modified compiler \Rightarrow compiler binary
- Restore compiler source to original state

Now: inspecting compiler source reveals nothing unusual

... but compiling compiler gives a corrupt compiler binary

Complication: compiler-backdoor needs to include all of (*)

What can we trust?

I order a laptop by mail. When it arrives, what can I trust on it?

- Applications and/or operating system may be backdoored
⇒ solution: reinstall OS and applications
- How to reinstall? Can't trust OS to reinstall the OS.
⇒ Boot *Tails* from a USB drive (Debian)
- Need to trust pre-boot BIOS, UEFI code. Can we trust it?
⇒ No! (e.g. ShadowHammer operation in 2018)
- Can we trust the motherboard? Software updates?

So, what can we trust?

Sadly, nothing ... anything can be compromised

- but then we can't make progress

Trusted Computing Base (TCB)

- Assume some minimal part of the system is not compromised
- Then build a secure environment on top of that

will see how during the course.

03

PART THREE

Introduction

Security, at What Cost?

Why Security Issues Are Prevalent

- Both **hardware** and **software** are primarily designed for **performance and efficiency** – money matters!
- **Optimization** features like **speculative execution** and **caching** improve speed but can introduce security vulnerabilities.
- Example: transient execution enabled attacks like **Meltdown** and **Spectre**, revealing data that should remain isolated

The Popular Tradeoff Theory

- There's a widely acknowledged **tradeoff** between **performance** and **security**.
 - Design choices that favor performance often **deprioritize security safeguards**.
 - Adding robust security mechanisms tends to **slow systems down**.



Perf vs. Sec



Security Patches Can Impair Performance

- To fix **Meltdown** attack, **kernel page-table isolation (KPTI)** is introduced.
- KPTI **slows down operations** that transition between kernel and user modes.
- Performance degrades by **2–14%** in benchmarks, and even up to **20%** in I/O-intensive workloads on older hardware.
- Some cloud services and consumer systems faced **latency spikes, reboots, and service disruptions** ([WIRED](#)).

Next lecture: control hijacking vulnerabilities

THE END