

Session 06

Modern Offensive and Defensive Solutions

Security of Information Systems (SIS)

Computer Science and Engineering Department

November 8, 2023

- ▶ HCFI: Hardware-enforced Control-Flow Integrity
- ▶ Losing Control: On the Effectiveness of Control-Flow Integrity under Stack Attacks

1 / 46

2 / 46

Attack and Defense

- ▶ attack: exploit vulnerabilities
- ▶ defense: prevent attacks, make attacks difficult, confine attacks
- ▶ attacker needs to find one security hole
- ▶ defender has to protect all security holes
- ▶ attacker invests time
- ▶ defense mechanisms incur overhead

3 / 46

4 / 46

Attacker Perspective and Mindset

- ▶ find one vulnerability and build from that
- ▶ look for something that is valuable
- ▶ do reconnaissance, look for weak spots
- ▶ create an attack chain
- ▶ use every trick in the book
- ▶ start from existing knowledge

Defender Perspective and Mindset

- ▶ protect all entry points
- ▶ users are vulnerable, as well as technology
- ▶ use multiple defensive layers
- ▶ monitor, be proactive
- ▶ discipline, best practices are worth more than skills
- ▶ invest more on valuable targets

5 / 46

6 / 46

Attacker Pros/Cons

- ▶ apart from ethical hackers, security researchers, it's a shady business
- ▶ you may not need skills, just a weak target and a database of attack vectors
- ▶ you may get caught
- ▶ you only need to find one spot
- ▶ possible great gains
- ▶ little time for fame (anonymous)
- ▶ the Internet gives you tons of targets
- ▶ but many targets give little more than fun

Defender Pros/Cons

- ▶ less resources (time) than an attacker
- ▶ must think of everything
- ▶ is being paid constructively
- ▶ you have a purpose: keep the system running
- ▶ it never ends

7 / 46

Honeypots

- ▶ baits
- ▶ a system appearing as vulnerable but closely monitored
- ▶ deflect, change attention and collect attacker information

8 / 46

- ▶ buffer overflows
- ▶ shellcodes
- ▶ memory protection (DEP, WX)
- ▶ memory randomization
- ▶ canaries
- ▶ code reuse
- ▶ CFI (Control Flow Integrity)
- ▶ memory safety, safe programming languages
- ▶ static and dynamic analysis
- ▶ hardware enhanced security

10 / 46

- ▶ <https://dl.acm.org/citation.cfm?id=2498135>
- ▶ issue with ASLR: memory disclosure / information leak
- ▶ one address leaked reveals all information
- ▶ do it at page level
- ▶ one leak may lead to other leaks that are chained together

11 / 46

SafeStack

- ▶ <https://clang.llvm.org/docs/SafeStack.html>
- ▶ part of the Code Pointer Integrity project:
<http://dslab.epfl.ch/proj/cpi/>
- ▶ moves sensitive information (such as return addresses) on a safe stack, leaving problematic ones on the unsafe stack
- ▶ reduced overhead, protects against stack buffer overflows
- ▶ microStache:
<https://www.springerprofessional.de/en/microstache-a-lightweight-execution-context-for-in-process-safe-16103742>

12 / 46

Address Sanitizer

- ▶ ASan
- ▶ <https://research.google.com/pubs/archive/37752.pdf>
- ▶ <https://github.com/google/sanitizers/>
- ▶ instruments code
- ▶ only useful in development
- ▶ detects out-of-bounds bugs, memory leaks

13 / 46

CFI/CPI

- ▶ <https://dl.acm.org/citation.cfm?id=1102165>
- ▶ <https://www.usenix.org/node/186160>
- ▶ <http://dslab.epfl.ch/proj/cpi/>
- ▶ coarse-grained CFI vs fine-grained CFI
- ▶ Control Flow Integrity, Code Pointer Integrity
- ▶ protect against control flow hijack attacks
- ▶ CPI is weaker than CFI but more practical (reduced overhead)
- ▶ CPI protects all code pointers, data based attacks may still happen
- ▶ CPS (Code Pointer Separation) is a weaker yet more practical for of CPI

14 / 46

Shellcodes

- ▶ difficult to inject due to DEP, small buffers and input validation
- ▶ preliminary parts of the attack may remap memory region
- ▶ shellcode may do stack pivoting and then load another shellcode
- ▶ alphanumeric shellcodes: still need a binary address to bootstrap

15 / 46

Code Reuse

- ▶ bypass DEP by using existing pieces of code
- ▶ code gadgets
- ▶ used in ROP (*Return-Oriented Programming*) and JOP (*Jump-Oriented programming*)

16 / 46

Return-Oriented Programming

- ▶ gadgets ending in `ret`
- ▶ may be chained together to form an attack
- ▶ Turing-complete language
- ▶ <http://www.suse.de/~krahmer/no-nx.pdf>
- ▶ <https://dl.acm.org/citation.cfm?id=2133377>
- ▶ most common way of creating runtime attack vectors
- ▶ JOP: <https://dl.acm.org/citation.cfm?id=1966919>
 - ▶ gadgets end up in an indirect branch not a `ret`

17 / 46

- ▶ prevent attacks
 - ▶ SafeStack
 - ▶ CFI/CPI, ASan
 - ▶ Microsoft CFG, RFG
- ▶ detect attacks
 - ▶ Microsoft EMET (*Enhanced Mitigation Experience Toolkit*), ProcessMitigations module

18 / 46

- ▶ <https://www.usenix.org/node/190963>
- ▶ <https://huhong-nus.github.io/advanced-DOP/>
- ▶ overwrites data, not code pointers
- ▶ bypasses CFI

19 / 46

Evolution of OS Security

- ▶ traditional main goals: functionality and reduced overhead
- ▶ recent focus on OS security: plethora of devices and use cases
- ▶ malware may easily take place among legitimate applications
- ▶ kernel exploits become more common
- ▶ OS virtualization, reduce TCB to hypervisor
- ▶ include hardware-enforced security features

21 / 46

Mandatory Access Control

- ▶ opposed to Discretionary Access Control, where owner controls permissions
- ▶ system-imposed settings
- ▶ increased, centralized security
- ▶ difficult to configure and maintain
- ▶ rigid, non-elastic
- ▶ Bell-LaPadula Model: <http://csrc.nist.gov/publications/history/bell176.pdf>
- ▶ SELinux, TrustedBSD, Mandatory Integrity Control

22 / 46

Role-Based Access Control

- ▶ <https://ieeexplore.ieee.org/abstract/document/485845>
- ▶ <https://dl.acm.org/citation.cfm?id=266751>
- ▶ aggregate permissions into roles
- ▶ role assignment, role authorization, permission authorization
- ▶ useful in organizations

23 / 46

Sandboxing

- ▶ assume application may be malware
- ▶ reduce potential damage
- ▶ confine access to a minimal set of allowed actions
- ▶ typically implemented at sandbox level (kernel enforced)
- ▶ iOS sandboxing, Linux seccomp

24 / 46

Application Signing

- ▶ ensure application is validated
- ▶ used by application stores and repositories: GooglePlay, Apple AppStore
- ▶ device may not run non-signed apps

25 / 46

iOS Jekyll Apps

- ▶ https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/wang_tielei
- ▶ apparently legitimate iOS app
- ▶ bypasses Apple vetting
- ▶ obfuscates calls to private libraries (part of the same address space, fixed from iOS 7)
- ▶ once installed turns out to be malware
- ▶ exfiltrates private data, exploits vulnerabilities

26 / 46

Jailbreaking/Rooting

- ▶ <https://dl.acm.org/citation.cfm?id=3196527>
- ▶ get root access on a device
- ▶ close to full control
- ▶ requires a critical vulnerability that gets root access
- ▶ tethered (requires re-jailbreaking after reboot) vs non-tethered
- ▶ essential for security researchers

27 / 46

Hardware-centric Attacks

- ▶ side channels
- ▶ undocumented hardware features
- ▶ imperfect hardware features that leak information
- ▶ proprietary features that get exploited
- ▶ hardware is part of TCB, reveals kernel memory

28 / 46

Sidechannel Attacks

- ▶ do not exploit vulnerabilities in applications or kernel code
- ▶ mostly use features such as

29 / 46

x86 Instruction Fuzzing

- ▶ <https://www.blackhat.com/docs/us-17/thursday/us-17-Domas-Breaking-The-x86-Instruction-Set-wp.pdf>
- ▶ <https://github.com/xoreaxeaxeax/sandsifter>
- ▶ <https://i.blackhat.com/us-18/Thu-August-9/us-18-Domas-God-Mode-Unlocked-Hardware-Backdoors-In-x86-CPUs.pdf>
- ▶ instruction of length N is placed at the end of the page
- ▶ creates a fuzzer for the x86 instruction set
- ▶ found glitches, hidden instructions

30 / 46

IME

- ▶ *Intel Management Engine*
- ▶ AMD Secure Technology
- ▶ hardware features and highly proprietary firmware
- ▶ <https://www.theverge.com/2018/1/3/16844630/intel-processor-security-flaw-bug-kernel-windows-linux>
- ▶ user space app could access kernel space memory access
- ▶ accused of being a backdoor to the system

31 / 46

rowhammer Attack

- ▶ <https://users.ece.cmu.edu/~yoonguk/papers/kim-isca14.pdf>
- ▶ <https://www.vusec.net/projects/drammer/>
- ▶ <https://googleprojectzero.blogspot.com/2015/03/exploiting-dram-rowhammer-bug-to-gain.html>
- ▶ <https://www.blackhat.com/docs/us-15/materials/us-15-Seaborn-Exploiting-The-DRAM-Rowhammer-Bug-To-Gain-Kernel.pdf>
- ▶ hardware fault in DRAM chips
- ▶ constant bit flip pattern in certain rows can cause a flip in another row (not belonging to the current process)
- ▶ may be exploited to get root access

32 / 46

Spectre and Meltdown

- ▶ <https://meltdownattack.com>
- ▶ <https://www.usenix.org/system/files/conference/usenixsecurity18/sec18-lipp.pdf>
- ▶ application may access data from another application
- ▶ Meltdown exploits a hardware race condition allowing an unprivileged process to read privileged data
- ▶ Spectre does a side channel attack on speculative execution features of modern CPUs
- ▶ hardware fixes by Intel, software solutions

33 / 46

KPTI

- ▶ *Kernel Page Table Isolation*
- ▶ <https://lwn.net/Articles/741878/>
- ▶ place kernel in separate address space
- ▶ mitigation against hardware-centric attacks

34 / 46

Hypervisor Attacks

- ▶ <https://dl.acm.org/citation.cfm?id=2484406>
- ▶ attack/compromise hypervisor, get control of VMs
- ▶ may exploit a vulnerability in the hypercall interface or may exploit a hardware bug
- ▶ hyperjacking

35 / 46

Evolution of Web Security

- ▶ path traversals, misconfigurations
- ▶ injections
- ▶ XSS
- ▶ misconfiguration
- ▶ unsafe communication
- ▶ application/language bugs

37 / 46

Secure Communication

- ▶ provide secure communication between client and server
- ▶ HTTPS everywhere
- ▶ Secure Cookie
- ▶ strong encryption, strong protocols

38 / 46

Attacks on Security Protocols

- ▶ <https://tools.ietf.org/html/rfc7457>
- ▶ <https://www.mitls.org/pages/attacks>
- ▶ flaws in protocol logic
- ▶ cryptographic design flaws
- ▶ implementation flaws

39 / 46

Connection Downgrade

- ▶ part of man-in-the-middle attack
- ▶ negotiate a connection with weaker protocol features than the current one
- ▶ ideally drop HTTPS altogether
- ▶ POODLE (*Padding Oracle On Downgraded Legacy Encryption*)
- ▶ <https://www.openssl.org/~bodo/ssl-poodle.pdf>

40 / 46

Advanced Injection Attacks

- ▶ LDAP, XPath injection
- ▶ blind SQL injection: content-based and time-based
- ▶ https://www.owasp.org/images/7/74/Advanced_SQL_Injection.ppt
- ▶ <https://nvisium.com/blog/2015/06/17/advanced-sql-injection.html>

41 / 46

Language Bugs

- ▶ bugs/vulnerabilities in frameworks
- ▶ bugs/vulnerabilities in web modules or language interpreter

42 / 46

Modern Offensive and Defensive Techniques

- ▶ attacks focus on low-level aspects of a system: hide features, exploit hardware, side channels, protocol design
- ▶ assume better/improved applications but imperfect system/protocol/configuration design
- ▶ defense takes more time and incurs significant overhead
- ▶ battle rages on

44 / 46

Keywords

- ▶ honeypot
- ▶ fine-grained ASLR
- ▶ SafeStack
- ▶ AddressSanitizer
- ▶ CFI/CPI
- ▶ code reuse
- ▶ ROP, JOP
- ▶ data-oriented attacks
- ▶ MAC, RBAC
- ▶ sandboxing
- ▶ Jekyll apps
- ▶ jailbreak, rooting
- ▶ side channel attacks
- ▶ IME attack
- ▶ Meltdown, Spectre
- ▶ KPTI
- ▶ rowhammer
- ▶ connection downgrade
- ▶ POODLE
- ▶ blind SQL injection

Resources

- ▶ see URLs across slides