

Lecture 16: Security

CS343 – Operating Systems
Branden Ghena – Spring 2024

Some slides borrowed from:
Tyler Bletsch (NC State), Berkeley CS61C

Administrivia

- No Monday Office Hours
 - Memorial Day holiday
- Paging Lab due next week Thursday (May 30)
 - If you haven't yet, start putting some serious work into it

Today's Goals

- Introduce OS security considerations.
- Describe memory-based attacks and defenses.
- Explore speculative execution attacks and ramifications.

Why is computer security so important?

- Most public security happens at least in some portion on the honor system
 - Pretty easy to break a window
 - Keyed locks are easy to pick
 - Master keys can be determined and manufactured ([Matt Blaze attack](#))
 - Laws apply after you've done it



Early computers didn't have any security either

- Simple machines for doing computation do not have private files or contention
- Timeslicing machines meant there were multiple users, but all were employees of the same company
 - Permissions needed to be as secure as a file in a locked drawer on a desk




“The act of breaking into a computer system has to have the same social stigma as breaking into a neighbor's house. It should not matter that the neighbor's door is unlocked.”

- Ken Thompson, Turing Award Lecture, 1984

Connectivity of computers makes security a top concern

- Importantly, physical item security is dependent on the fact that one person can only steal one thing at a time
 - And it's usually obvious when theft occurs
- The internet changed all of this for computers
 - Usually not people breaking into computers manually, one at a time
 - Instead it is computers breaking into computers by means of scripting
 - And you can access a computer from anywhere on Earth
- Breaking into or controlling one car is a crime
 - Controlling 100,000 cars remotely is a problem for the manufacturer

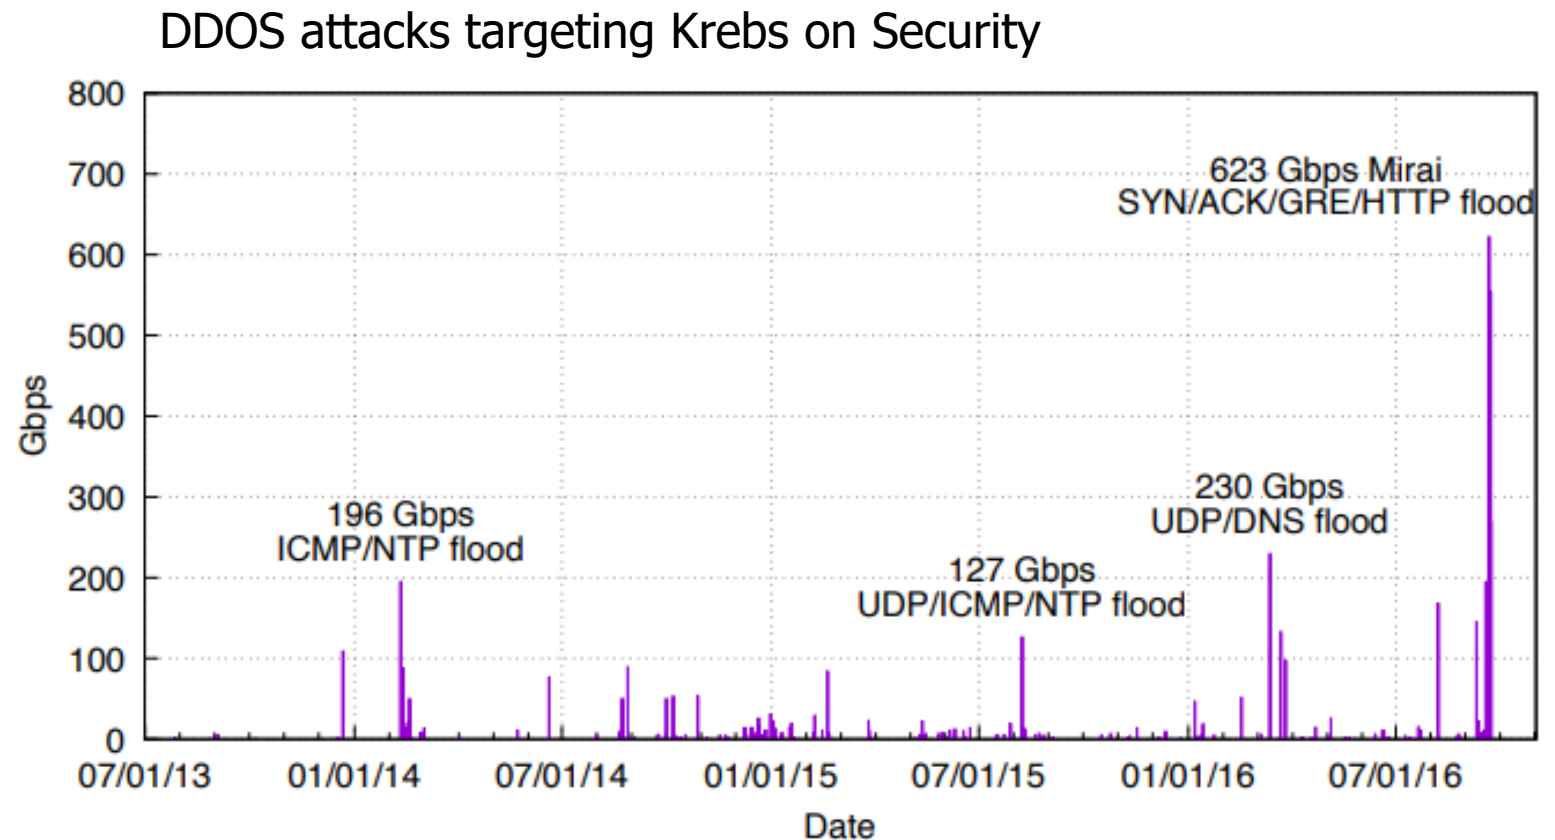
Mirai botnet (2016)



- Takes control of up to 600,000 insecure connected devices
 - IP-attached cameras, DVRs, routers, printers

Botnets can be directed towards denial-of-service attacks

- Mirai is used for DDOS attacks on various websites
 - Krebs on Security blog gets 623 Gbps of traffic during one attack



Outline

- **Design for security**
- Memory attacks and defenses
 - Buffer Overflows
 - Return-Oriented Programming
- Speculative execution attacks
 - Meltdown
 - Spectre

Trusted Computing Base (TCB)

- Trusted Computing Base is everything the OS relies on to enforce security
 - If everything outside of the TCB is “evil”, the TCB can still be trusted
 - Important to be a clear, minimum set of components
- TCB includes
 - Scheduler, Memory Management, Parts of file system, Parts of device drivers
- Anything else must be assumed malicious
 - Processes memory accesses, System call arguments, Received packets

Modern code bases are enormous

Program/Use Case	Millions of Lines of Code
Unix v1.0	0.01
Average iPhone app	0.04
Space Shuttle	0.4
Windows 3.1	2.5
Mars Curiosity Rover	5
Firefox (2015)	9.7
F-35 Fighter jet	24
Microsoft Office 2001	25
Windows 7	40
Facebook (2015)	62
Debian 5.0 codebase	68

- For many projects, no one person has read and understood all of it
- TCB needs to be agreed upon by everyone working on the project
 - And needs to be enforced by everyone in the project

<https://www.informationisbeautiful.net/visualizations/million-lines-of-code/>

Can we even trust the Trusted Computing Base?

- Can you trust the OS with your password? (or anything, really)
 - How do you know that the OS you're running hasn't already been taken over or modified in some way?
- Particularly large concern for server operators
 - Thousands of computers
 - All operated remotely without explicit users
 - Need to ensure that they aren't taken over
- Really malicious code might modify the OS if it has access
 - That way even if the computer reboots, the malicious part remains
 - Or modify the boot software (UEFI) to compromise *everything*

Hardware Root-of-Trust

- Idea: software can be tampered with, but hardware is MUCH more difficult
 - Requires physical access, at which point all bets are off anyways...
- When a server starts:
 1. Root-of-Trust chip boots first and hardware automatically checks the authentication of its code before starting it
 2. Root-of-Trust code checks authentication of OS code before booting the OS on the actual CPU
 3. OS actually starts running on the CPU
- Now the code running on the server can be trusted to be authentic

Writing auditable code

- Code style and semantics really do matter!!
 - If you want code to be secure, it needs to be read AND understood by many people
 - This is why I focus so much on semantics in Intro to C/C++
- Bad code style/semantics builds up cognitive load of the reader making them less likely to notice when something is wrong
 - 0 versus NULL
 - `&buf[0]` versus `&(buf[0])`
 - `int x, y, z;` versus `int x; int y; int z;`

Apple "goto fail" SSL bug

Spacing intentional. This code mixes tabs and spaces and has random extra line breaks.

It is actually decently commented overall, just not in this particular section.

...

```
if ((err = SSLFreeBuffer(&hashCtx)) != 0)
    goto fail;
```

```
if ((err = ReadyHash(&SSLHashSHA1, &hashCtx)) != 0)
    goto fail;
```

```
if ((err = SSLHashSHA1.update(&hashCtx, &clientRandom)) != 0)
    goto fail;
```

```
if ((err = SSLHashSHA1.update(&hashCtx, &serverRandom)) != 0)
    goto fail;
```

```
if ((err = SSLHashSHA1.update(&hashCtx, &signedParams)) != 0)
    goto fail;
    goto fail;
```

```
if ((err = SSLHashSHA1.final(&hashCtx, &hashOut)) != 0)
    goto fail;
```

...

Apple "goto fail" SSL bug

Spacing intentional. This code mixes tabs and spaces and has random extra line breaks.

It is actually decently commented overall, just not in this particular section.

...

```
if ((err = SSLFreeBuffer(&hashCtx)) != 0)
    goto fail;
```

```
if ((err = ReadyHash(&SSLHashSHA1, &hashCtx)) != 0)
    goto fail;
```

```
if ((err = SSLHashSHA1.update(&hashCtx, &clientRandom)) != 0)
    goto fail;
```

```
if ((err = SSLHashSHA1.update(&hashCtx, &serverRandom)) != 0)
    goto fail;
```

```
if ((err = SSLHashSHA1.update(&hashCtx, &signedParams)) != 0)
    goto fail;
```

```
goto fail;
```

```
if ((err = SSLHashSHA1.final(&hashCtx, &hashOut)) != 0)
    goto fail;
```

...

Apple "goto fail" SSL bug

Spacing intentional. This code mixes tabs and spaces and has random extra line breaks.

It is actually decently commented overall, just not in this particular section.

...

```
if ((err = SSLFreeBuffer(&hashCtx)) != 0)
    goto fail;
```

```
if ((err = ReadyHash(&SSLHashSHA1, &hashCtx)) != 0)
    goto fail;
```

```
if ((err = SSLHashSHA1.update(&hashCtx, &clientRandom)) != 0)
    goto fail;
```

```
if ((err = SSLHashSHA1.update(&hashCtx, &serverRandom)) != 0)
    goto fail;
```

```
if ((err = SSLHashSHA1.update(&hashCtx, &signedParams)) != 0)
    goto fail;
```

```
goto fail;
```

Outside of IF statement!! Always runs.

```
if ((err = SSLHashSHA1.final(&hashCtx, &hashOut)) != 0)
    goto fail;
```

...

Sandboxing approach to untrusted code

- What if you don't know if you can trust some running code?
 - Or you *know* you actively don't trust it, but still want to run it
 - Example, PDF interpretation is actually a Turing-complete language
 - Lots of possibly buggy or abusable things going on in there
 - But we do still want to interpret PDFs!
- Sandboxing: running code with restricted access to other parts of the system
 - Reduces the possible attacks the code might make on your system

iOS “BlastDoor” Sandbox

- iOS uses BlastDoor to sandbox arriving iMessage data
 - Anyone can send *anything* over iMessage
 - Data needs to be decompressed and interpreted with various image file types supported
 - LOTS of attack surface: various targeted “zero-click attacks”
- BlastDoor limits possible interactions
 - No file system access
 - No network access
 - No interaction with other processes
 - On a crash, restarts with exponential delay

Principle of Least Privilege

- Only provide access to resources that are necessary for a legitimate purpose
- That way malicious behavior, that you aren't even aware of yet, has a limited amount of damage it can inflict

Security properties OS should enforce

- Confidentiality
 - Private information should remain private
 - Example: processes can't read memory in another process
- Integrity
 - Mechanisms should not be modified without permission
 - Example: OS data structures can't be modified by processes
- Availability
 - Resources on the computer should be able to be fairly accessed
 - Example: network access is shared among processes

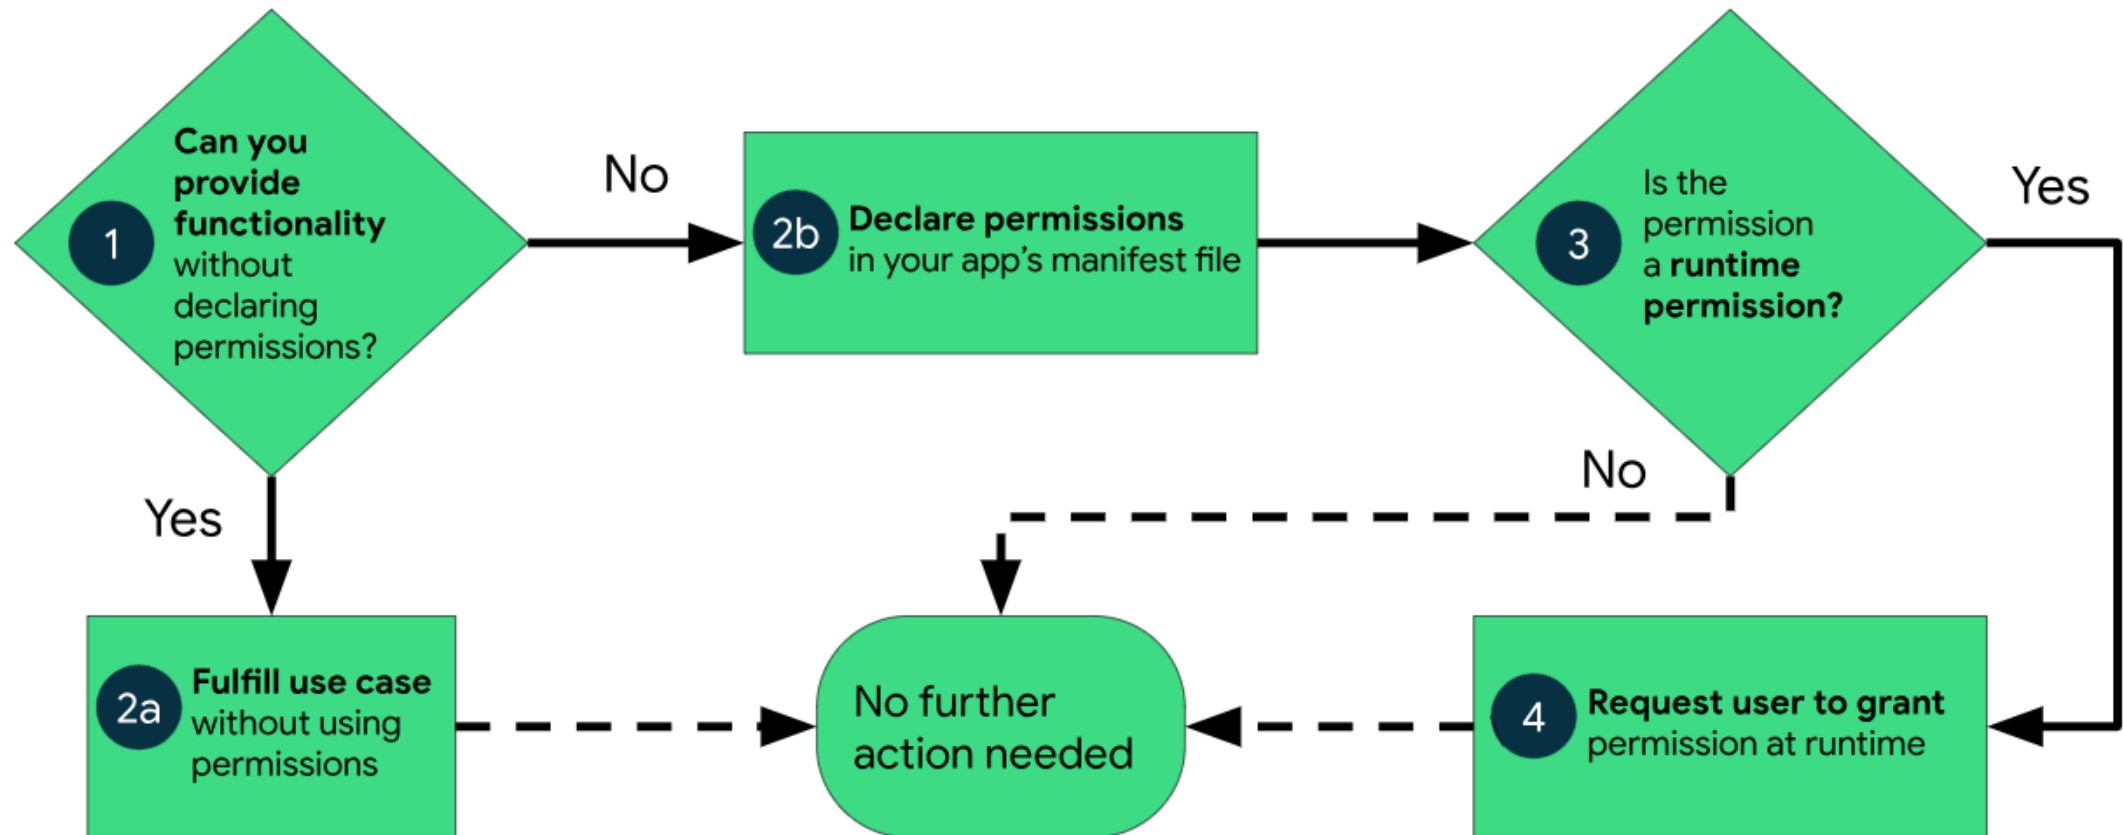
OS security concerns

- Processor access
 - Integrity: User versus kernel mode
 - Availability: Timeslicing
- Memory access
 - Confidentiality and Integrity: Virtual memory (and permissions)
 - Availability: Swapping
- File access
 - Confidentiality: Permissions (user and group)
 - Integrity: only accessible through system calls

What about devices?

- Device access
 - Confidentiality: User permissions... sort of?
- This gets complicated
 - Should any app I run be able to activate my webcam or microphone?
 - When should Uber be able to access my location?
- Still figuring this one out
 - Smartphones are at the forefront

Android access control model



- Ask the user to approve
 - Either at install time or at runtime

Authentication

- Act of proving some information, such as the identity of a computer system user
 - Often the responsibility of the kernel as a trusted entity
- Many actions are limited based on identity
 - File access privileges
 - Ability to install new programs
 - Access to certain hardware devices or mechanisms
- Kernel versus user process is one identity separation
 - Servers might have many different users

Identifying users

- Three overarching methods:
 1. Authentication based on “what you know”
 - Passwords, Security questions
 2. Authentication based on “what you have”
 - Security key, Cell phone
 3. Authentication based on “what you are”
 - Biometrics: fingerprint, face ID, retinal scan
- Multi-Factor Authentication (MFA) requires multiple different categories from the above



Break + xkcd



Outline

- Design for security
- **Memory attacks and defenses**
 - **Buffer Overflows**
 - **Return-Oriented Programming**
- Speculative execution attacks
 - Meltdown
 - Spectre

What's wrong with this code?

```
#include <stdlib.h>
#include <stdio.h>

int main() {
    char name[1024];
    printf("What is your name? ");
    scanf("%s", name);
    printf("%s is cool.\n", name);

    return 0;
}
```

Buffer overflow potential with “nice” input

```
tkblets@davros:~/jop/examples/code-injection $ ./cool
What is your name? Tyler
Tyler is cool.
tkblets@davros:~/jop/examples/code-injection $ █
```


Buffer Overflow

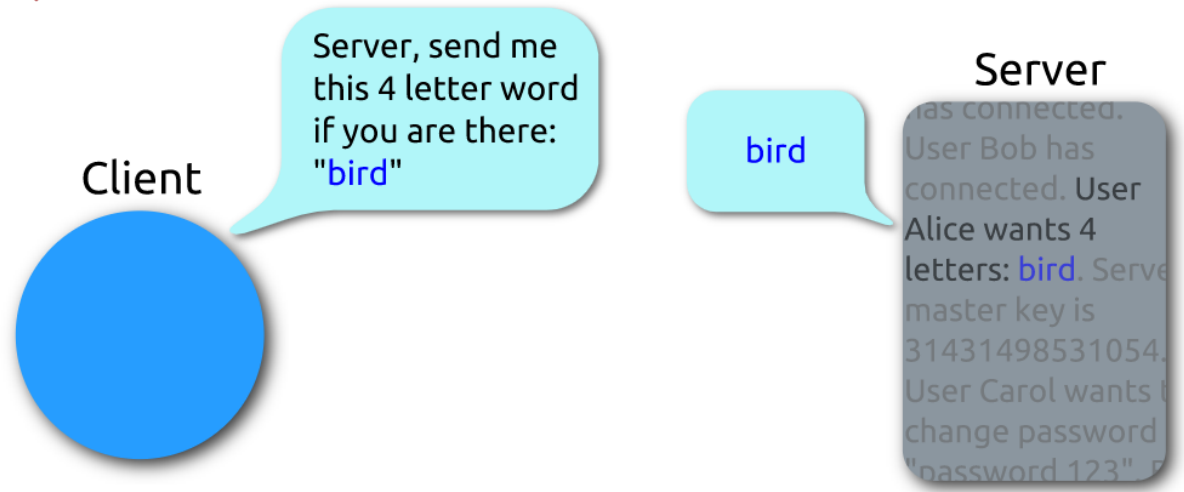
- Arrays (buffers) in C are not bounds checked
 - Can keep writing past the end of the array
 - Overwrites either data section or stack section
- Still an incredibly common problem in C
- **Key problem**
 - Trusting input from an untrustworthy source
 - Users are not part of the trusted computing base
 - Certainly not arbitrary inputs they can make

Heartbleed attack

- Vulnerability in OpenSSL
 - 2014
- Started the trend of vulnerabilities with cool names and logos



Heartbeat – Normal usage

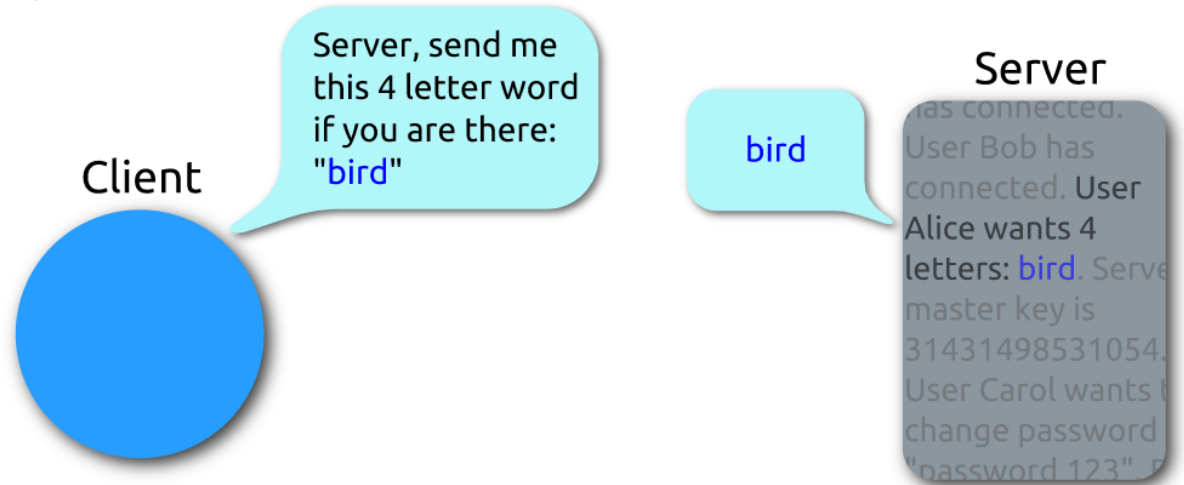


Heartbleed attack

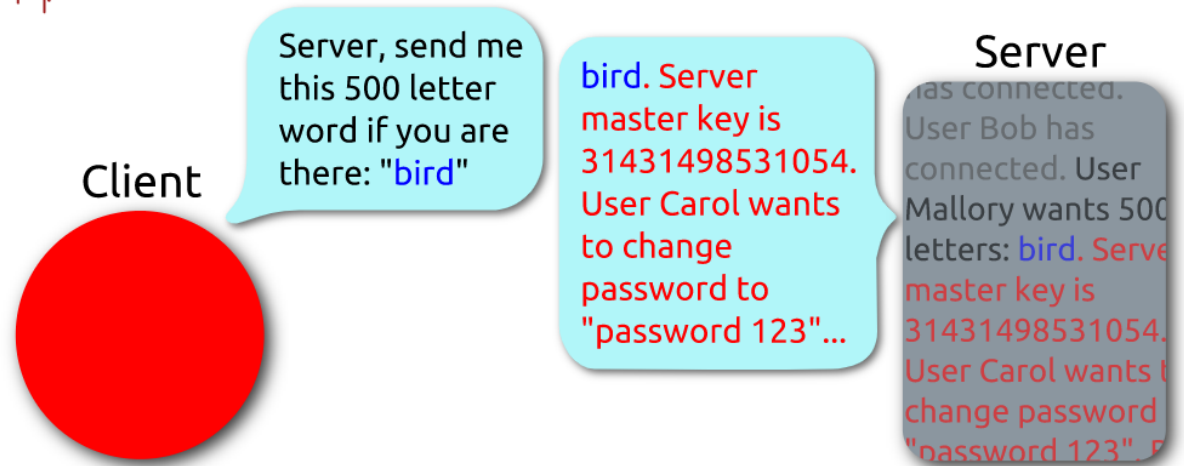
- Vulnerability in OpenSSL
 - 2014
- Started the trend of vulnerabilities with cool names and logos



Heartbeat – Normal usage



Heartbeat – Malicious usage



Return addresses constantly live on the stack

- Recall: When a function is called...
 - parameters are pushed on stack
 - return address pushed on stack
 - called function puts local variables on the stack
- Memory layout



- C's calling convention means arbitrary execution could happen anywhere!

What do you do with arbitrary execution?

- Open a shell that can run anything...
- Top: C code
- Middle: position-independent x86 assembly
- Bottom: machine code hex

```
int main(int argc, char *argv[])
{
    char *sh;
    char *args[2];

    sh = "/bin/sh";
    args[0] = sh;
    args[1] = NULL;
    execve(sh, args, NULL);
}
```

(a) Desired shellcode code in C

```
    nop
    nop                // end of nop sled
    jmp  find          // jump to end of code
cont: pop  %esi        // pop address of sh off stack into %esi
    xor  %eax,%eax     // zero contents of EAX
    mov  %al,0x7(%esi) // copy zero byte to end of string sh (%esi)
    lea  (%esi),%ebx   // load address of sh (%esi) into %ebx
    mov  %ebx,0x8(%esi) // save address of sh in args[0] (%esi+8)
    mov  %eax,0xc(%esi) // copy zero to args[1] (%esi+c)
    mov  $0xb,%al     // copy execve syscall number (11) to AL
    mov  %esi,%ebx    // copy address of sh (%esi) to %ebx
    lea  0x8(%esi),%ecx // copy address of args (%esi+8) to %ecx
    lea  0xc(%esi),%edx // copy address of args[1] (%esi+c) to %edx
    int  $0x80        // software interrupt to execute syscall
find: call cont        // call cont which saves next address on stack
sh:  .string "/bin/sh " // string constant
args: .long 0          // space used for args array
     .long 0          // args[1] and also NULL for env array
```

(b) Equivalent position-independent x86 assembly code

```
90 90 eb 1a 5e 31 c0 88 46 07 8d 1e 89 5e 08 89
46 0c b0 0b 89 f3 8d 4e 08 8d 56 0c cd 80 e8 e1
ff ff ff 2f 62 69 6e 2f 73 68 20 20 20 20 20 20
```

(c) Hexadecimal values for compiled x86 machine code

Morris Worm

- November 02, 1988
 - Roughly 88,000 computers on internet at the time
- Worm
 - Invading program that installs itself on additional computers
- Infected several thousand computers, taking down internet for several days



How the worm entered computers: three methods

1. Debug vulnerability in *sendmail* – an email sending service
 - Connect, enter debug mode, send arbitrary code to execute
2. Buffer overflow in *finger* – a command to list user details
 - Send request with more than 512 bytes of arguments
 - Execute `/bin/sh`
3. Guess passwords
 - Get list of users for the machine worm is already running in
 - Guess username, reverse username, 400 “popular” words, entire dictionary

Effects of Morris Worm


- Morris Worm created too many copies of itself
 - Checked if there was already a worm on the computer before running
 - 1 out of 7 of the executables just ran anyways (too high a default)
- Computers ended up with many processes running
 - **Check your understanding:** How are too many processes harmful?

Effects of Morris Worm

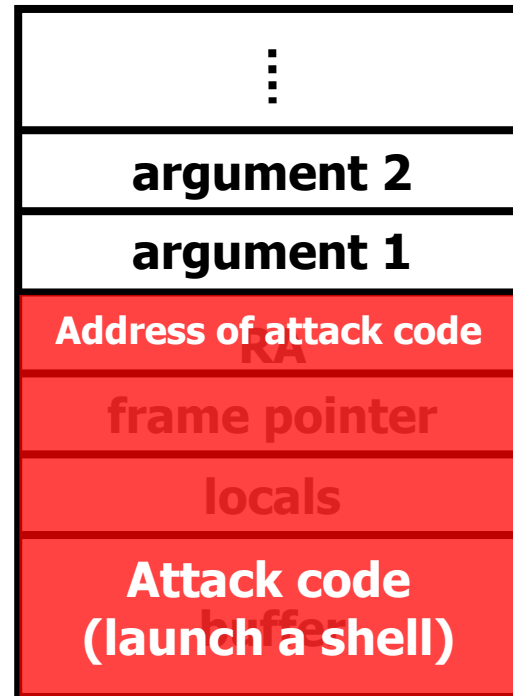
- Morris Worm created too many copies of itself
 - Checked if there was already a worm on the computer before running
 - 1 out of 7 of the executables just ran anyways (too high a default)
- Computers ended up with many processes running
 - Long response time due to so many processes
 - Thrashing due to too much memory pressure
 - Slowed computers to a halt
- Outcomes:
 - Invaded ~6000 computers in hours (**10% of the Internet** at the time)
 - CERT was created to manage software security
 - First Computer Fraud and Abuse Act (CFAA) prosecution

Overcoming no-execute

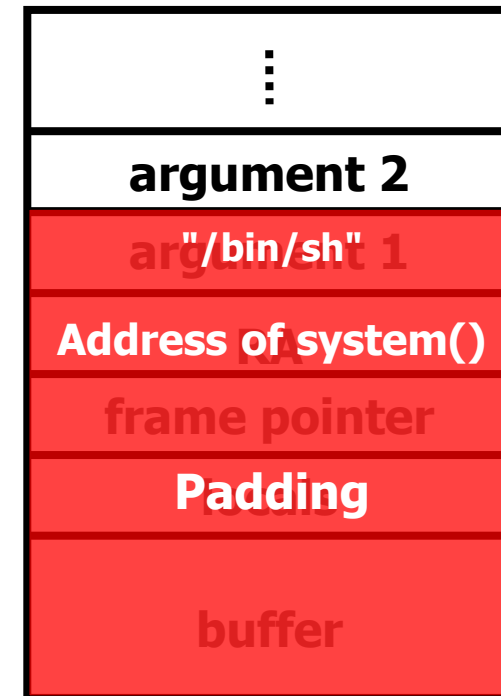
- Do we need malicious code to have malicious behavior? **No**



Default Stack



Code injection




Code reuse (!)

"Return-into-libc" attack


Return-oriented programming

- More general process to enable arbitrary execution without code rewrite
- Look through assembly instructions followed by a return
 - Known as "gadgets"
- Chain these gadget together to make working code
 - By placing addresses on stack




Gadgets can create a Turing-complete programming environment


- Loading constants




- Control flow



- Arithmetic




- Memory



Address-space layout randomization (ASLR)

- Randomize memory region locations in virtual memory
 - Already spread throughout physical memory
- Move locations of libraries and code relative to each other
 - Arbitrary address for attacker to send code to gets harder to predict!
- Implemented 2005-2007
 - Linux, MacOS, and Windows
 - 2011 for Android and iOS



Overcoming ASLR

- ASLR is a probabilistic approach, merely increases attacker's expected work
 - Each failed attempt results in crash; at restart, randomization is different
- Counters:
 - Information leakage
 - Program reveals a pointer? Game over.
 - De-randomization attack
 - Just keep trying! (carefully)
 - 32-bit ASLR defeated in 216 seconds
 - BlastDoor sandbox has delay after crash for exactly this scenario
 - Under certain scenarios is less effective
 - Poor source of randomness

Break + Question

- The Common Vulnerabilities and Exposures (CVE) system documents publicly released software vulnerabilities.
- How long has it been since the last CVE due to a buffer overflow?

Break + Question

- The Common Vulnerabilities and Exposures (CVE) system documents publicly released software vulnerabilities.
- How long has it been since the last CVE due to a buffer overflow?
 - Today is Thursday (May 23rd, 2024)

Discovered Wednesday, May 22nd

[CVE-2024-4453](#)

GStreamer EXIF Metadata Parsing Integer Overflow Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of GStreamer. Interaction with this library is required to exploit this vulnerability but attack vectors may vary depending on the implementation. The specific flaw exists within the parsing of EXIF metadata. The issue results from the lack of proper validation of user-supplied data, which can result in an integer overflow before allocating a buffer. An attacker can leverage this Source: Zero Day Initiative

Max CVSS	7.8
EPSS Score	N/A
Published	2024-05-22
Updated	2024-05-22

Break + Question

- The Common Vulnerabilities and Exposures (CVE) system documents publicly released software vulnerabilities.
- How long has it been since the last CVE due to a buffer overflow?

Last MAJOR overflow vulnerability: Tuesday, May 21st

[CVE-2023-3943](#)

Stack-based Buffer Overflow vulnerability in ZkTeco-based OEM devices allows, in some cases, the execution of arbitrary code. Due to the lack of protection mechanisms such as stack canaries and PIE, it is possible to successfully execute code even under restrictive conditions. This issue affects ZkTeco-based OEM devices (ZkTeco ProFace X, Smartec ST-FR043, Smartec ST-FR041ME and possibly others) with firmware ZAM170-NF-1.8.25-7354-Ver1.0.0 and possibly others.

Source: Kaspersky Labs

Max CVSS	10.0
EPSS Score	0.04%
Published	2024-05-21
Updated	2024-05-21

Outline


- Design for security
- Memory attacks and defenses
 - Buffer Overflows
 - Return-Oriented Programming
- **Speculative execution attacks**
 - **Meltdown**
 - **Spectre**

First, some background knowledge

- To understand Speculative Execution Attacks you really need to understand low-level software and hardware
- A few pieces of background knowledge will be useful:
 - Timing Side Channels
 - Speculative Execution
 - Keeping the kernel in Virtual Memory

Background: Side channel attacks

- Important for understanding speculative execution attacks
- Many physical systems have properties that may leak information about internal state
 - Determine RSA key bits based on power use during a decrypt operation
 - Determine length of password by how long it takes to check it



Timing attacks are one side channel

- Timing attacks can be overcome with constant-time algorithms which always take as long as the worst-case execution time
 - But this means reducing performance
- Caches are essentially one big timing attack
 - Speeds up access to data if it is present in the cache
 - This was the goal!!
 - An attack can know which data was accessed recently
 - But that seems harmless, right?

Background: Speculative Execution

Modern processors want to always be doing something

- What if we're going to branch based on a memory load?
- What if we just guess what the result will be and start executing early!!

So they are often "speculatively executing" instructions

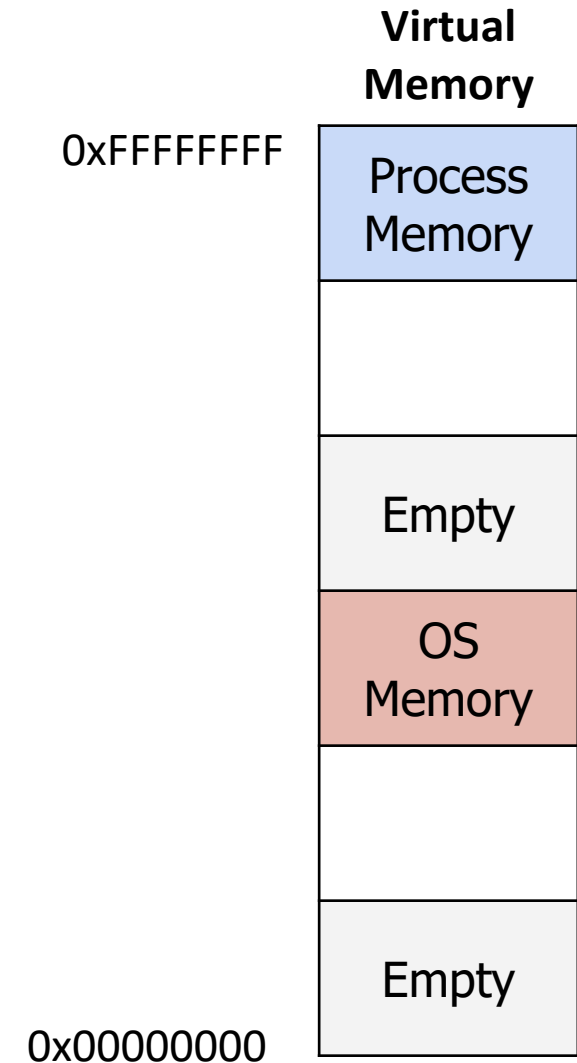
- Perform the operation and throw out the result if we shouldn't actually do it
- For example, branch prediction

Optimization: Kernel Mapped in Virtual Memory

Page tables map virtual memory to physical memory for a process

But actually, we often leave the OS memory in the page table too...

- Each page is marked as no-read, no-write
- Faster to switch back to the OS
 - No need to TLB flush or page table swap if the OS intends to go right back to process
- Also allows the kernel to swap out parts of its own memory if necessary
 - Such as page tables themselves



Meltdown

Security vulnerability in all modern processors that allows arbitrary reads from memory

Disclosed in January 2018 by: (told Intel in June 2017)

- Jann Horn ([Google Project Zero](#)),
- Werner Haas, Thomas Prescher ([Cyberus Technology](#)),
- Daniel Gruss, Moritz Lipp, Stefan Mangard, Michael Schwarz ([Graz University of Technology](#))

Details:

- <https://hackernoon.com/a-simplified-explanation-of-the-meltdown-cpu-vulnerability-ad316cd0f0de>
- <https://meltdownattack.com/meltdown.pdf>



Step 1: Read from a kernel address

```
mov $KERNEL_ADDRESS_OF_SECRET, %r12  
mov (%r12), %eax
```

`%eax` now holds a byte of memory that we shouldn't be able to access

- This will be an invalid page fault!
- Once the instruction actually hits the end of the pipeline...
- For now, it loads that value into `%r12` right away and continues executing speculatively

Step 2: Read based on secret

```
mov $KERNEL_ADDRESS_OF_SECRET, %r12  
mov (%r12), %eax  
mov MY_ARRAY(%eax), %edx
```

`%edx` is a valid read from our own memory

- This is never going to finish either because the process will have an exception from the prior instruction, but it will start executing...
- `MY_ARRAY` here is a 256-byte array which is not in the cache

Step 3: Handle the Exception

```
mov $KERNEL_ADDRESS_OF_SECRET, %r12
```

```
mov (%r12), %eax
```

```
mov MY_ARRAY(%eax), %edx
```

The processor realizes you tried to read from memory you didn't have access to and generates an exception

- You can catch these and recover
- The invalid instruction and ones after it are rolled back as if they never happened

Everything's still safe right?

The processor never saved any results from the invalid accesses to memory in registers

- So there's no problem, right?

We forgot about the cache

The load affected the cache!!!

```
mov $KERNEL_ADDRESS_OF_SECRET, %r12  
mov (%r12), %eax  
mov MY_ARRAY(%eax), %edx
```

The value at address **MY_ARRAY+%eax** was saved in our cache

Step 4: Time loads from memory

```
for (int i=0; i<255; i++){  
    start_time = time();  
    int temp = MY_ARRAY[i*CACHE_BLOCKSIZE];  
    stop_time = time();  
  
    if ((stop_time-start_time) <= SHORT_TIME){  
        secret = i;  
    }  
}
```

The cache speeds up the access to the one memory address that was cached due to speculative execution

Step 5: Repeat and Profit

- Now we know the value of a single byte
- But we can repeat this process over and over to read arbitrary memory
 - Read from memory at ~ 500 kbps
- Incredible part is how relatively simple this attack is
 - Does require systems knowledge of multiple domains
 - Computer architecture, OS, and security

How do we fix this?


1. Stop speculatively executing
 - Already in the hardware
 - Would slow all computers down a lot
2. Stop caching speculative loads
 - Already in the hardware
 - Would slow all computers down a lot
3. Stop leaving OS memory in the page table ✓
 - Would slow all computers down somewhat
 - Kernel Page Table Isolation
 - Estimated 5-30% performance loss
 - Improved by use of PCID bit in TLB

Sidebar: how long were we vulnerable to Meltdown

- From the authors, every Intel processor implementing out-of-order execution is potentially affected
 - Which is roughly every processor from 1995-2018 (20+ years)
 - Some non-Intel processors are affected as well around the same time range
- New processors can be designed to avoid the vulnerability

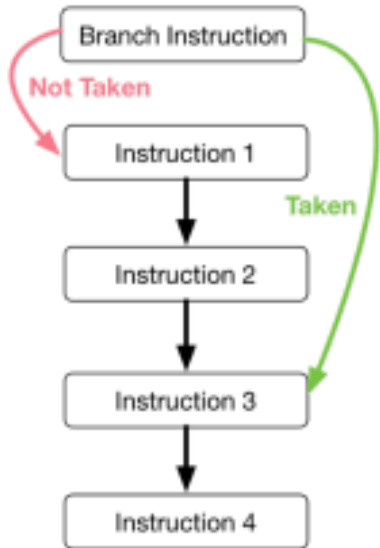
Spectre

- Speculative execution targeting branch prediction
- Disclosed in January 2018 by
 - [Jann Horn](#) (Google Project Zero) and
 - [Paul Kocher](#) in collaboration with, in alphabetical order, [Daniel Genkin](#) (University of Pennsylvania and University of Maryland), [Mike Hamburg](#) (Rambus), [Moritz Lipp](#) (Graz University of Technology), and [Yuval Yarom](#) (University of Adelaide and Data61)



Background: Branch Prediction


Sample Program



A sample five instruction program used to demonstrate effects of branch prediction.

No Branch Prediction


Time (not to scale)



Without branch prediction, the majority of a branch instruction is spent determining whether the branch condition is true (take the branch) or false (do not take the branch).

Branch Prediction (Correct Guess)


Time (not to scale)



This is the time saved by guessing the branch condition instead of waiting to compute it. Modern processors guess correctly more than 96% of the time on normal workloads, resulting in a significant speed boost.

Branch Prediction (Wrong Guess)

Time (not to scale)



The time wasted by incorrectly predicting the branch destination is called the misprediction penalty. During that time, the processor speculatively executes instructions (Instruction 1 and Instruction 2 in this example). These instructions are unexecuted once the processor realizes it made a mistake.

Incredibly accurate in modern day computers >95%

Spectre v1

- Repeat meltdown-style attack using conditional branches
 - Conditional branches are especially prevalent for bounds checks in software virtual machines (like the Javascript runtime)
1. Train conditional branch predictor that bounds check branch always succeeds
 2. Make an invalid bounds-checked read, affecting cache state
 3. Use cache timing analysis to determine value of read byte

Spectre v2

- Combine indirect branch prediction and in-kernel ROP gadgets
 - Indirect branch predictors try loading a guessed address
1. Train indirect branch predictor to go to a particular address
 2. Make a system call requesting something
 3. Within the system call, a branch mis-prediction then runs the targeted gadget, affecting cache state
 - Note: the gadget runs with kernel permission on physical memory
 4. Use cache timing attack to determine result

Spectre fallout

- Spectre allows code inside a process to access all memory of the process
 - Bypassing any security mechanisms or containerization
 - Example: Javascript running inside a web browser
 - Led to increased push for “one website per process”
- Spectre is harder to fix too. Can't just change page tables
 - No one simple thing can fix all of these problems
 - Stopping branch prediction helps, but we don't want to stop it everywhere
 - Active research on targeted branch prediction disabling

Ramifications of speculative execution attacks

- Particularly big deals in the era of cloud computing
 - Anyone can run a program on an AWS server
 - And now can maybe read data from the other running programs...
- Speculative execution attacks are a new era for computer security
 - Hardware is still being actively developed to address attacks
 - Websites can be fixed in hours, Programs in days, OSes in weeks, and Hardware takes years
 - Attacks are still being developed
- Role of the OS: mitigate hardware issues as best possible

Security is an arms race

- There is no single fix for system security
 - New attacks are constantly being discovered
 - New solutions are constantly being applied
- 1. Find a vulnerability and how it can be exploited
- 2. Fix vulnerability
- 3. Go back to 1
- But if the OS is designed with security in mind, it's hopefully harder to find vulnerabilities in the first place

Outline

- Design for security
- Memory attacks and defenses
 - Buffer Overflows
 - Return-Oriented Programming
- Speculative execution attacks
 - Meltdown
 - Spectre