Buffer Overflows

CSE 351 Winter 2021

Instructor:

Mark Wyse

Teaching Assistants:

Kyrie Dowling Catherine Guevara Ian Hsiao Jim Limprasert Armin Magness Allie Pfleger Cosmo Wang Ronald Widjaja



Alt text: I looked at some of the data dumps from vulnerable sites, and it was ... bad. I saw emails, passwords, password hints. SSL keys and session cookies. Important servers brimming with visitor IPs. Attack ships on fire off the shoulder of Orion, c-beams glittering in the dark near the Tannhäuser Gate. I should probably patch OpenSSL.

http://xkcd.com/1353/

Administrivia

- Lab 2 due today!
- hw13 due Wednesday (2/10)
- hw14 due Friday (2/12)
- Lab 3 released Wednesday, due Monday (2/22)
 - You will have everything you need by the end of this lecture

Buffer Overflows

- Address space layout review
- Input buffers on the stack
- Overflowing buffers and injecting code
- Defenses against buffer overflows

not drawn to scale

Review: General Memory Layout

- Stack
 - Local variables (procedure context)
- Heap
 - Dynamically allocated as needed
 - new,malloc(),calloc(),...
- Statically-allocated Data
 - Read/write: global variables (Static Data)
 - Read-only: string literals (Literals)
- Code/Instructions
 - Executable machine instructions
 - Read-only



not drawn to scale

Memory Allocation Example

```
char big array[1L<<24]; /* 16 MB */
int global = 0;
int useless() { return 0; }
int main() {
 void *p1, *p2;
 int local = 0;
 p1 = malloc(1L << 28); /* 256 MB */
 p2 = malloc(1L << 8); /* 256 B */
 /* Some print statements ... */
```

Where does everything go?



not drawn to scale

Memory Allocation Example



What Is a Buffer?

- A buffer is an array used to temporarily store data
- You've probably seen "video buffering..."
 - The video is being written into a buffer before being played
- Buffers can also store user input





Reminder: x86-64/Linux Stack Frame

- Caller's Stack Frame
 - Arguments (if > 6 args) for this call
- Current/ Callee Stack Frame
 - Return address
 - Pushed by call instruction
 - Old frame pointer (optional)
 - Caller-saved pushed before setting up arguments for a function call
 - Callee-saved pushed before using long-term registers
 - Local variables (if can't be kept in registers)
 - "Argument build" area
 (Need to call a function with >6 arguments? Put them here)



- C does not check array bounds
 - Many Unix/Linux/C functions don't check argument sizes
 - Allows overflowing (writing past the end) of buffers (arrays)
- * "Buffer Overflow" = Writing past the end of an array
- Characteristics of the traditional Linux memory layout provide opportunities for malicious programs
 - Stack grows "backwards" in memory
 - Data and instructions both stored in the same memory

- Stack grows *down* towards lower addresses
- Buffer grows *up* towards higher addresses
- If we write past the end of the array, we overwrite data on the stack!

Enter input: hello

No overflow 😊



- Stack grows *down* towards lower addresses
- Buffer grows *up* towards higher addresses
- If we write past the end of the array, we overwrite data on the stack!

Enter input: helloabcdef



- Stack grows *down* towards lower addresses
- Buffer grows *up* towards higher addresses
- If we write past the end of the array, we overwrite data on the stack!

Enter input: helloabcdef

Buffer overflow! 🛞



- Buffer overflows on the stack can overwrite "interesting" data
 - Attackers just choose the right inputs
- Simplest form (sometimes called "stack smashing")
 - Unchecked length on string input into bounded array causes overwriting of stack data
 - Try to change the return address of the current procedure
- Why is this a big deal?
 - It was the #1 technical cause of security vulnerabilities
 - #1 overall cause is social engineering / user ignorance

String Library Code

* Implementation of Unix function gets()



What could go wrong in this code?
Doit has before size

String Library Code

* Implementation of Unix function gets()

```
/* Get string from stdin */
char* gets(char* dest) {
    int c = getchar();
    char* p = dest;
    while (c != EOF && c != '\n') {
        *p++ = c;
        c = getchar();
    }
    *p = '\0';
    return dest;
}
```

No way to specify limit on number of characters to read

- Similar problems with other Unix functions:
 - Strcpy: Copies string of arbitrary length to a dst
 - Scanf, fscanf, sscanf, when given %s specifier

Vulnerable Buffer Code

```
/* Echo Line */
void echo() {
    char buf[8]; /* Way too small! */
    gets(buf);
    puts(buf);
}
```

```
void call_echo() {
    echo();
}
```

unix> ./buf-nsp Enter string: 123456789012345 123456789012345

unix> ./buf-nsp Enter string: 1234567890123456 Segmentation fault (core dumped)

Buffer Overflow Disassembly (buf-nsp)

echo:

000000000	040114	6 <ec< th=""><th>cho>:</th><th>24 by thes</th></ec<>	cho>:	24 by thes
401146:	48 83	ec 1	. 8	sub \$0x18,%rsp
• • •				calls printf
401159:	48 8d	7c 2	24 08	lea 0x8(%rsp),%rdi
40115e:	b8 00	00 0	00 00	mov \$0x0,%eax
401163:	e8 e8	fe f	ff ff	callq 401050 <gets@plt></gets@plt>
401168:	48 8d	7c 2	24 08	lea 0x8(%rsp),%rdi
40116d:	e8 be	fe f	ff ff	callq 401030 <puts@plt></puts@plt>
401172:	48 83	c4 1	. 8	add \$0x18,%rsp
401176:	с3			retq
call_echo:				
0000000000	401177	<cal< td=""><td>l ech</td><td>>:</td></cal<>	l ech	>:
401177:	48 83	ec O	8	sub \$0x8,%rsp
40117b:	b8 00	00 0	00 00	mov \$0x0,%eax
401180:	e8 c1	ff f	f ff	callq 401146 <echo></echo>
⊊ 401185:	48 83	c4 0	8	add \$0x8,%rsp
401189:	с3			retq

Buffer Overflow Stack

Before call to gets



Buffer Overflow Example

Before call to gets



Buffer Overflow Example #1





Overflowed buffer, but did not corrupt state

Buffer Overflow Example #2

After call to gets



unix> ./buf-nsp
Enter string: 1234567890123456
Segmentation fault (core dumped)

Overflowed buffer and corrupted return pointer

Buffer Overflow Example #2 Explained

After return from echo



"Returns" to a valid instruction, but bad indirect jump so program signals SIGSEGV, Segmentation fault

Malicious Use of Buffer Overflow: Code Injection Attacks <u>Stack after call to gets ()</u>



- Input string contains byte representation of executable code
- Overwrite return address A with address of buffer B
- When bar() executes ret, will jump to exploit code

Practice Question

- smash_me is vulnerable to stack smashing!
- What is the minimum number of characters that gets must read in order for us to change the return address to a stack address?



Exploits Based on Buffer Overflows

Buffer overflow bugs can allow attackers to execute arbitrary code on victim machines

- Distressingly common in real programs
 - Programmers keep making the same mistakes S
 - Recent measures make these attacks much more difficult
- Examples across the decades
 - Original "Internet worm" (1988)
 - Heartbleed (2014, affected 17% of servers)
 - Similar issue in Cloudbleed (2017)
 - Hacking embedded devices
 - Cars, Smart homes, Planes

Example: the original Internet worm (1988)

- Exploited a few vulnerabilities to spread
 - Early versions of the finger server (fingerd) used gets() to read the argument sent by the client:
 - finger droh@cs.cmu.edu
 - Worm attacked fingerd server with phony argument:
 - finger "exploit-code padding new-return-addr"
 - Exploit code: executed a root shell on the victim machine with a direct connection to the attacker
- Scanned for other machines to attack
 - Invaded ~6000 computers in hours (10% of the Internet)
 - see June 1989 article in Comm. of the ACM
 - The author of the worm (Robert Morris*) was prosecuted...

Example: Heartbleed (2014)

HOW THE HEARTBLEED BUG WORKS:



Example: Heartbleed (2014)



Example: Heartbleed (2014)



Heartbleed Details

- Buffer over-read in OpenSSL
 - Open source security library
 - Bug in a small range of versions
- "Heartbeat" packet
 - Specifies length of message
 - Server echoes it back
 - Library just "trusted" this length
 - Allowed attackers to read contents of memory anywhere they wanted
- Est. 17% of Internet affected
 - "Catastrophic"
 - Github, Yahoo, Stack Overflow, Amazon AWS, ...

😥 Heartbeat – Normal usage



By FenixFeather - Own work, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=32276981

Hacking Cars (2010)

- UW CSE research demonstrated wirelessly hacking a car using buffer overflow
 - http://www.autosec.org/pubs/cars-oakland2010.pdf
- Overwrote the onboard control system's code
 - Disable brakes, unlock doors, turn engine on/off



Hacking DNA Sequencing Tech (2017)

Computer Security and Privacy in DNA Sequencing

Paul G. Allen School of Computer Science & Engineering, University of Washington

- Potential for malicious code to be encoded in DNA!
- Attacker can gain control of DNA sequencing machine when malicious DNA is read
- Ney et al. (2017): <u>https://dnasec.cs.washington.edu/</u>





Figure 1: Our synthesized DNA exploit

Dealing with buffer overflow attacks

- 1) Employ system-level protections
- 2) Avoid overflow vulnerabilities
- 3) Have compiler use "stack canaries"

1) System-Level Protections



Any attempt to execute this code will fail

1) System-Level Protections

- Non-executable code segments
 - Wait, doesn't this fix everything?
- Works well, but can't always use it
- Many embedded devices *do not* have this protection
 - *e.g.*, cars, smart homes, pacemakers
- Some exploits still work!
 - Return-oriented programming
 - Return to libc attack
 - JIT-spray attack



Any attempt to execute this code will fail

ASLR

1) System-Level Protections

- At start of program, allocate random amount of space on stack
- Shifts stack addresses for entire program
 - Addresses will vary from one run to another
- Makes it difficult for hacker to predict beginning of inserted code
- Example: Address of variable local for when Slide 5 code executed 3 times:
 - 0x7ffd19d3f8ac
 - 0x7ffe8a462c2c
 - 0x7ffe927c905c
 - Stack repositioned each time program executes

9:20m



2) Avoid Overflow Vulnerabilities in Code

```
/* Echo Line */
void echo()
{
    char buf[8]; /* Way too small! */
    fgets(buf, <u>8</u>, stdin);
    puts(buf);
}
```

- Use library routines that limit string lengths
 - fgets instead of gets (2nd argument to fgets sets limit)
 - strncpy instead of strcpy
 - Don't use scanf with %s conversion specification
 - Use fgets to read the string
 - Or use <code>%ns</code> where <code>n</code> is a suitable integer

2) Avoid Overflow Vulnerabilities in Code

- Alternatively, don't use C use a language that does array index bounds check
 - Buffer overflow is impossible in Java
 - ArrayIndexOutOfBoundsException
 - Rust language was designed with security in mind
 - Panics on index out of bounds, plus more protections

3) Stack Canaries

- Basic Idea: place special value ("canary") on stack just beyond buffer
 - Secret value that is randomized before main()
 - Placed between buffer and return address
 - Check for corruption before exiting function
- GCC implementation
 - -fstack-protector

unix>./buf Enter string: **12345678** 12345678 unix> ./buf
Enter string: 123456789
*** stack smashing detected ***

This is extra

(non-testable)

material

Setting Up Canary

Before call to gets



This is extra

(non-testable)

material

Checking Canary

After call to gets



Summary of Prevention Measures

- 1) Employ system-level protections
 - Code on the Stack is not executable
 - Randomized Stack offsets
- 2) Avoid overflow vulnerabilities
 - Use library routines that limit string lengths
 - Use a language that makes them impossible
- 3) Have compiler use "stack canaries"

Think this is cool?

- * You'll love Lab 3 🔅
 - Released Wednesday, due next Friday (11/13)
 - Some parts *must* be run through GDB to disable certain security features
- Take CSE 484 (Security)
 - Several different kinds of buffer overflow exploits
 - Many ways to counter them
- Nintendo fun!
 - Using glitches to rewrite code: <u>https://www.youtube.com/watch?v=TqK-2jUQBUY</u>
 - Flappy Bird in Mario: <u>https://www.youtube.com/watch?v=hB6eY73sLV0</u>