CY 2550 Foundations of Cybersecurity

Exploits and Patches

Alina Oprea Associate Professor, Khoury College Northeastern University March 30 2020

Announcements

- Forensics project released today, due on April 4
- Exploit project is the last one, due on April 17
- Final exam
 - Take home
 - Released on April 13 at 11:45am EST, due on April 14 at noon
 - Submitted through Gradescope
 - Questions on the material to test general understanding
 - Might include questions from the "Countdown to Zero Day" book

Focus on Attacks

- Software is notorious for having bugs
 - Functionality that doesn't work as intended, or at all
 - Crashes that cause unreliability, data loss
- To an attacker, software bugs are opportunities
- Exploits
 - Weaponized software bugs
 - Use programming errors to an attacker's advantage
- Typical uses
 - Bypass authentication and authorization checks
 - Elevate privileges (to admin or root)
 - Hijack programs to execute unintended, arbitrary code
 - Enable unauthorized, persistent access to systems

Outline

- Program Execution Basics
- Buffer Overflows Attacks
 - C examples
- Mitigations
- Web-based attacks
- SQL Basics
- SQL Injection
- Patches

Program Execution

Code and Data Memory

Program Execution

The Stack

Compilers



- Computers don't execute source code
- Instead, they execute machine code
- Compilers translate source code to machine code
- Assembly is human-readable machine code

	00000000000040052 <u>d <main>:</main></u>						
C Source Code			55		push	rbp	
	×84-64	machine	48 89	e5	mov	rbp,rsp	
			48 83	ec 20	sub	rsp,0x20	
<pre>#include <stdio.h></stdio.h></pre> Code in hexadecin		exadecimai	89 7d	ec	mov	DWORD PTR [rbp-0x14],edi	
			48 89	75 e0	mov	QWORD PTR [rbp-0x20],rsi	
		40053c:	83 7d	ec 01	cmp	DWORD PTR [rbp-0x14],0x1	
<pre>int main(int argc, char** argv) {</pre>		400540:	7e 36		jle	400578 <main+0x4b></main+0x4b>	
		400542:	c7 45	fc 01 00 00 00	mov	DWORD PTR [rbp-0x4],0x1	
int i;		400549:	eb 23		jmp	40056e <main+0x41></main+0x41>	
<pre>if (argc > 1) { for (i = 1; i < argc; ++i) {</pre>		40054b:	8b 45	fc	mov	eax,DWORD PTR [rbp-0x4]	
		40054e:	48 98		cdqe		
		400550:	48 8d	14 c5 00 00 00	lea	rdx,[rax*8+0x0]	
print+("%s	printf("%s\n",argv[i]); 400557:		00				
}		400558:	48 8b	45 eØ	mov	rax,QWORD PTR [rbp-0x20]	
}		40055c:	48 01	d0	add	rax,rdx	
		40055f:	48 8b	00	mov	rax,QWORD PTR [rax]	
else {	else {		48 89	с7	mov	rdi,rax	
printf("%s\	<pre>n", "Hello world");</pre>	400565:	e8 a6	fe ff ff	call	400410 <puts@plt></puts@plt>	
1		40056a:	83 45	fc 01	add	DWORD PTR [rbp-0x4],0x1	
}		40056e:	8b 45 fc		mov	eax,DWORD PTR [rbp-0x4]	
return 1;		400571:	.: 3b 45 ec		cmp	eax,DWORD PTR [rbp-0x14]	
}		400574:	7c d5		jl	40054b <main+0x1e></main+0x1e>	
		400576:	eb 0a		jmp	400582 <main+0x55></main+0x55>	
		400578:	bf 14	x86-64	↓ mo∨	edi,0x400614	
40057d: 400582: 400587: 400588:		40057d:	e8 8e	assembly	call	400410 <puts@plt></puts@plt>	
		400582:	b8 01		mov	eax,0x1	
		400587:	c9 c3		leave		
		400588:			ret	7	

Computer Memory

- Running programs exists in memory
 - Program memory the code for the program
 - Data memory variables, constants, and a few other things, necessary for the program
 - OS memory always available for system calls
 - E.g. to open a file, print to the screen, etc.







The Stack

- Data memory is laid out using a specific data structure
 - The stack
- Every function gets a frame on the stack
 - Frame created when a function is called
 - Contains local, in scope variables
 - Frame destroyed when the function exits
- The stack grows downward
- Stack frames also contain control flow information
 - More on this in a bit...





Low 13

Stack Frame Example

IP

```
argv
   int fcount(char s[], char c) {
0:
                                                                                argc
      int cnt;
      int pos;
                                                                               IP = ...
                                                                main()
      for (pos = 0; pos < strlen(s); pos = pos + 1) {</pre>
1:
                                                                                 "†"
        if (s[pos] == c)
2:
                                                                             "testing"
             cnt = cnt + 1;
3:
       }
                                                                fcount()
                                                                               IP = 9
4:
      return cnt;
5: }
                                                                                 cnt
6:
                                                                                pos
7: void main(int argc, char* argv[]) {
      int cnt = fcount("testing", "t"); // should return 2
8:
9: <sup>}</sup>
```

14

High

Memory

Two Call Example

0: int fcount(char s[], char c) {
 integer cnt;
 integer pos;
1-4: ...

```
5: }
```

IP

- 6: void main(int argc, char* argv[]) {
 - 7: fcount("testing", "t"); // should return 2
 - 8: fcount("elevate", "e"); // should return 3
 - 9: }



Recursion Example

- 0: int r(int n) {
- 1: if (n > 0) r(n 1);
- 2: return n;
- 3: }
- 4: void main(int argc, char* argv[]) {
 5: r(4); // should return 4
 6: }



Review

- Running programs exist in memory (RAM)
- Code is in program memory
 - CPU keeps track of current instruction in the IP register
- Data memory is structured as a stack of frames
 - Each function invocation adds a frame to the stack
 - Each frame contains
 - Local variables that are in scope
 - Saved IP to return to

Fun Fact

- What is a stack overflow?
- Memory is finite
 - If recursion goes too deep, memory is exhausted
 - Program crashes
 - Called a stack overflow

Buffer Overflows

A Vulnerable Program

Smashing the Stack

Shellcode

NOP Sleds

Memory Corruption

- Programs often contain bugs that corrupt stack memory
- Usually, this just causes a program crash
 - The infamous "segmentation" or "page" fault
- To an attacker, every bug is an opportunity
 - Try to modify program data in very specific ways
- Vulnerability stems from several factors
 - Low-level languages are not memory-safe
 - Control information is stored inline with user data on the stack

Threat Model

- Attacker's goal:
 - Inject malicious code into a program and execute it
 - Gain all privileges and capabilities of the target program (e.g., setuid)
- System's goal: prevent code injection
 - Integrity program should execute faithfully, as programmer intended
 - Crashes should be handled gracefully
- Attacker's capability: submit arbitrary input to the program
 - Environment variables
 - Command line parameters
 - Contents of files
 - Network data

A Vulnerable Program



A Normal Exampl

What if the data in string s is longer than 32 characters?

- 0: void func_print(char s[]) {
 - // only holds 32 characters, max

char buffer[32];

1: strcpy(buffer, s);

2: printf("%s\n",buffer);

3: }

IP

strcpy() does not check the length of the input!

4: void main(int argc, char* argv[]) {

- 5: for (int i=1; i < argc; i++) {
- 6: func_print(argv[i]);

8: }





Demo: Program crash

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

```
void func_print(char in_str[]){
    char buffer[5];
    strcpy(buffer,in_str);
    printf("Buffer is:%s\n",buffer);
```

```
int main(int argc, char* argv[]) {
    int i;
    for (i=1; i<argc; i++){
        printf("Argument:%s\n",argv[i]);
        func_print(argv[i]);
    }
}</pre>
```

printf("Function returned successfully\n");

```
return(1);
```

```
-bash-4.2$
-bash-4.2$ gcc -o print.0 print.c -m32
-bash-4.2$ ./print.o 123
Argument:123
Buffer is:123
Function returned successfully
```

```
[-bash-4.2$ ./print.o 123 abcde
Argument:123
Buffer is:123
Argument:abcde
Buffer is:abcde
Function returned successfully
```

Demo: Program crash

[-bash-4.2\$ objdump -d print.o

ebp: Base of stack esp: top of stack eax: function return value



Demo: Program crash

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

```
void func_print(char in_str[]){
    char buffer[5];
    strcpy(buffer,in_str);
    printf("Buffer is:%s\n",buffer);
```

```
int main(int argc, char* argv[]) {
    int i;
    for (i=1; i<argc; i++){
        printf("Argument:%s\n",argv[i]);
        func_print(argv[i]);
    }
</pre>
```

```
printf("Function returned successfully\n");
```

```
return(1);
```

[-bash-4.2\$./print.o 1234567890 Argument:1234567890 Buffer is:1234567890 Function returned successfully

- Buffer is of size 5, but it is allocated for 13 characters
- One character is reserved for \n (end of string)

[-bash-4.2\$./print.o 1234567890123 Argument:1234567890123 Buffer is:1234567890123 Segmentation fault

Smashing the Stack

- Buffer overflow bugs can overwrite saved instruction pointers
 - Usually, this causes the program to crash
- Key idea: replace the saved instruction pointer
 - Can point anywhere the attacker wants
 - But where?
- Key idea: fill the buffer with malicious code
 - Remember: machine code is just a string of bytes
 - Change IP to point to the malicious code on the stack

Exploit v1

- 0: void func_print(char s[]) {
 // only holds 32 characters, max
 char buffer[32];
- 1: strcpy(buffer, s);
- 2: printf("%s\n",buffer);
- 3: }

IP

4: void main(int argc, char* argv[]) {

- 5: for (int i=1; i < argc; i++) {
- 6: func_print(argv[i]);
- 7: }
- 8: }



Malicious Code

- The classic attack when exploiting an overflow is to inject a payload
 - Sometimes called shellcode, since often the goal is to obtain a privileged shell
 - But not always!
- There are tools to help generate shellcode
 - Metasploit
- Example shellcode:

```
{
    // execute a shell with the privileges of the
    // vulnerable program
    exec("/bin/sh");
}
```

Hitting the Target

- Address of shellcode must be guessed exactly
 - Must jump to the precise start of the shellcode
- However, stack addresses often change
 - Change each time a program runs
- Challenge: how can we reliably guess the address of the shellcode?
 - Cheat!
 - Make the target even bigger so it's easier to hit ;)



Hit the Ski Slopes

- Most CPUs support no-op instructions
 - Simple, one byte instructions that don't do anything
 - On Intel x86, 0x90 is the NOP
- Key idea: build a NOP sled in front of the shellcode
 - Acts as a big ramp
 - If the instruction pointer lands anywhere on the ramp, it will execute NOPs until it hits the shellcode

Exploit v2

- 0: void print(string s) {
 // only holds 32 characters, max
 string buffer[128];
- 1: strcpy(buffer, s);
- 2: puts(buffer);

3: }

IP

4: void main(integer argc, strings argv) {

5: for (; argc >
$$0$$
; argc = argc - 1) {

6: print(argv[argc]);



Demo: Running Attack Code



[-bash-4.2\$ gcc vuln.c -o vuln.o -fno-stack-protector -m32	
[-bash-4.2\$./vuln.o	
Enter some text:	
[123	
You entered: 123	
[-bash-4.2\$./vuln.o	
Enter some text:	
[abc1234567	
You entered: abc1234567	
-bash-4 2\$ 🗍	

Goal: overwrite buffer in echo() function so that secretFunction() is called

Demo: Running Attack Code

-bash-4.2\$ objdump -d vuln.o

Start									
address of +	0804848d <secretfunction>:</secretfunction>								
function	804848d:	55	push %ebp						
lanetion	804848e:	89 e5	mov %esp,%ebp						
	8048490:	83 ec 18	sub \$0x18,%esp						
	8048493:	c7 04 24 94 85 04 08	movl \$0x8048594,(%esp)						
	804849a:	e8 b1 fe ff ff	call 8048350 <puts@plt></puts@plt>						
	804849f:	c7 04 24 a8 85 04 08	movl \$0x80485a8,(%esp)						
	80484a6:	e8 a5 fe ff ff	call 8048350 <puts@plt></puts@plt>						
	80484ab:	c9	leave						
	80484ac:	c3	ret						

Demo: Running Attack Code

-bash-4.2\$ objdump -d vuln.o

080484ad <ech 80484ad: 80484ae: 80484b0: 80484b3: 80484b3: 80484b4: 80484b4: 80484c2: 80484c6: 80484c4: 80484c4: 80484d5: 80484d9: 80484e0: 80484e0: 80484e5: 80484e6:</ech 	55 89 e5 83 ec 38 c7 04 24 d1 85 04 08 e8 91 fe ff ff 8d 45 e4 89 44 24 04 c7 04 24 e2 85 04 08 e8 9e fe ff ff 8d 45 e4 89 44 24 04 c7 04 24 e5 85 04 08 e8 5b fe ff ff c9 c3	<pre>push %ebp mov %csp,%ebp sub \$0x38,%esp movl \$0x00485d1,(%esp) call \$048350 <puts@plt> lea -0x1c(%ebp),%eax mov %eax,0x4(%esp) movl \$0x80485e2,(%esp) call 804830 <isoc99_scanf@plt> lea -0x1c %ebp),%eax mov %eax,(x4(%esp) movl \$0x80485e5,(%esp) call 8048340 <printf@plt> leave ret</printf@plt></isoc99_scanf@plt></puts@plt></pre>	Size of function 0x38 = 48+8 = 56 Size of buffer 0x1c = 16+12=28
0040400.		ebp: eax: size 4 for retu buffer (wan	4 bytes Irn value; right after It to override) ³⁶





Mitigations

- Stack canaries
 - Compiler adds special sentinel values onto the stack before each saved IP
 - Canary is set to a random value in each frame
 - At function exit, canary is checked
 - If expected number isn't found, program closes with an error
- Non-executable stacks
 - Modern CPUs set stack memory as read/write, but no eXecute
 - Prevents shellcode from being placed on the stack
- Address space layout randomization
 - Operating system feature
 - Randomizes the location of program and data memory each time a program executes



Review

- Programs are vulnerable to memory corruption
- Buffer overflow attacks
 - Make programs crash
 - Run malicious code
 - Use disassembly to learn address space of program and craft attack
 - More advanced attacks (return-to-libc)
- Mitigations: stack canaries, non-executable stacks, ASLR
 - Implemented in modern compilers
 - Still examples of vulnerabilities in the wild (HeartBleed)