

6.857 Rivest

L20.1 4/20/11

Admin: Quizzes & homeworks back at end of class

Today: Buffer overflow etc.

Buffer overflow

Basics

Details & Exploitation

Mitigation

Counterattacks

Format-string vulnerabilities

(?) Stuxnet

6.857 Rivest

L20.2 4/20/11

Buffer Overflow - Basics, Details, Exploitation

□ Slides

- Note escalation of privilege if program attacked is running with root privileges!

□ Take 6.858 for more on computer security!

Buffer Overflow - Mitigation

- Stackguard
(push canary value on stack when calling
check it before returning)
- Separate stack for return addresses
- ASLR - Address Space Layout Randomization
- Non-executable stack (NX)

Counterattacks

- Heap-based attack (malloc)
- Return to libc (no code executed on stack!)
 - each "return" executes snippet of code from libc
 - can construct "universal" vocabulary of such snippets

Videos:

David Brumley - Automatic Exploit Generation Overview

- AEG Exploits Demo ← 20 min total

security.ece.cmu.edu/aeg/

Note: now works with binary only!

6.857 Rivest

L20.4 4/20/11

Format String Vulnerabilities

(ref notes from 4/22/2009)

(Excellent notes from Jayant Krishnamurthy)

Code Injection Attacks

- Overview - why do them?
- Buffer Overflows (on the stack)
 - Program memory layout
 - Stack frames
 - A simple overflow
 - $\&$ Spawning Shells
 - ~~□ Heap based overflows and other~~
 - Putting it together
 - Defenses
 - other overflows: heap-based, return-to-libc
- Format String Exploits
 - C format strings
 - sketch of exploit
- XSS: Cross-Site Scripting
- SQL-injection

G.857 Spring 2009

~~Udacity~~ Lecture By

Jayant Krishnamurthi

L20, 4/22/2009

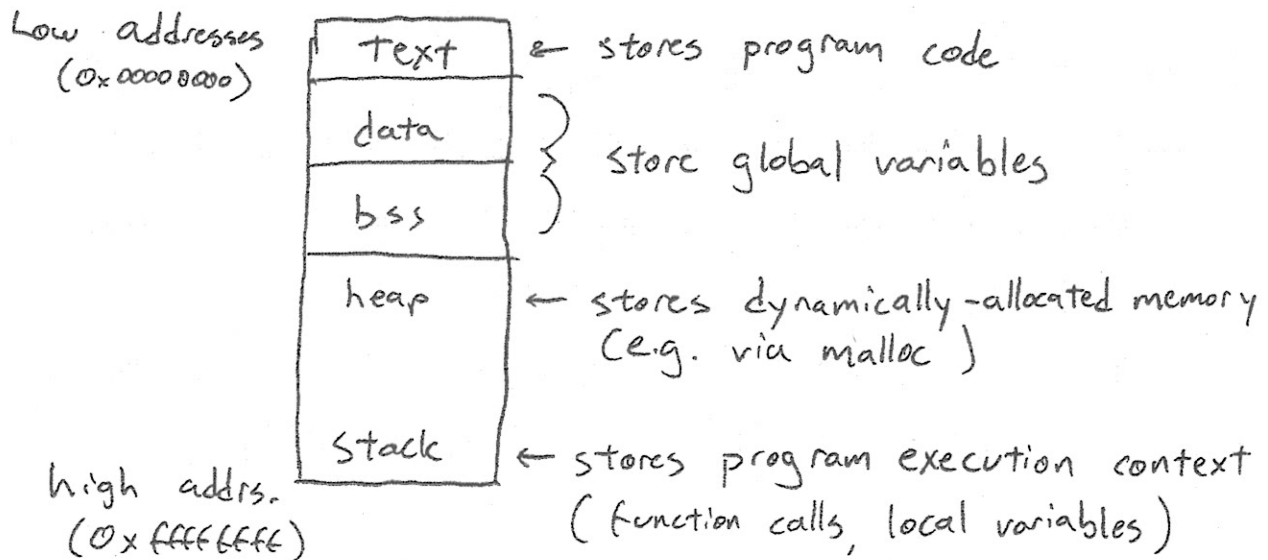
Why Code Injection:

1. Cause (potentially advantageous) incorrect behavior
2. Gain system privileges (root)
3. Gain access to a system
4. Steal information (XSS and SQL-injection)

Buffer Overflows:

- Common exploit that takes advantage of the fact that C does not perform boundary checks on arrays.
- Also exploits the layout of the program in memory

Basic Program Layout in Memory



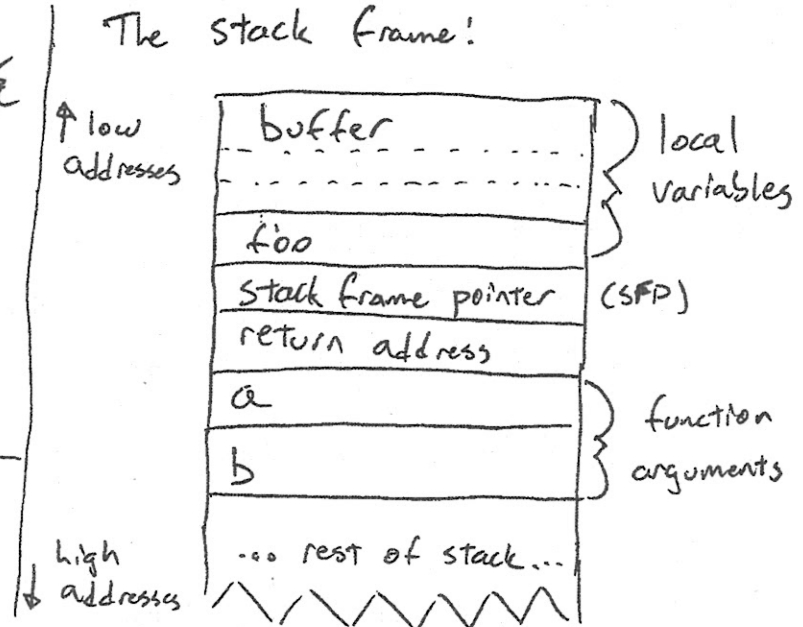
- Heap and stack are dynamic → their sizes change as the program runs.
- Heap grows up toward higher addresses, the stack grows down toward lower addresses
- Most common buffer overflow occurs on the stack

Stack Frames:

- whenever a function is called in a C program, a stack frame is created and added to the stack.

Example C program:

```
void test(int a, int b) {  
    char foo;  
    char buffer[10];  
}  
  
void main() {  
    test(1, 2);  
}
```



Buffering Buffers:

- C doesn't boundary check arrays
- strings are character arrays terminated with a null (0) byte.
 - functions like strcpy copy bytes until they reach a null byte.
- Putting too much data into a buffer is the basic mechanism of the buffer overflows (hence "overflow")

Uses of Buffer Overflows:

1. Cause crashes (as we've seen)
2. Overwrite variables with new values

(In the previous example, the value of flag was changed to $0x41414141$)

3. Execute arbitrary code

IDEA: change return address to a new, valid value.

A common location is the start of the buffer

itself, or an environment variable. Assembly code that is placed in the chosen location will be executed by the program.

the address of the buffer
can be found using a
debugger

Shellcode:

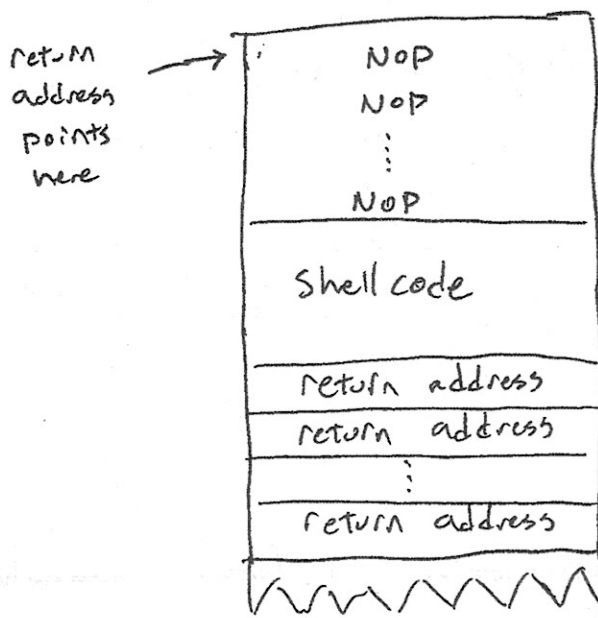
- Bytecode that opens up a shell. (use the `exec()` system call to execute a shell process)
- Somewhat tricky to make - typically have to avoid null bytes since they terminate C strings
- Can be as small as 31 bytes
- Can be all ASCII printable characters.

Let's say our overflow example declared buf as
`char * buf = argv[1];`

(meaning buf points to the 1st command line parameter).

Now what should we input to the program to cause an overflow?

Crafting an input buffer:



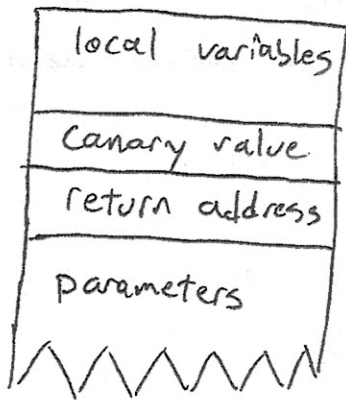
NOPs (short for No Operation) do nothing. This "NOP sled" lets us miss the exact address of the buffer by a little bit.

The return address is the start address of the buffer. Repeating the address several times lets us miss the exact position of the return address by a little bit.

- Ideally, when this is copied into buffer, we will overwrite the return address of the function call with the address of ~~the~~ buffer. This will cause execution ~~not~~ to jump to our custom code and spawn a shell.

Defense mechanisms:

- Address Space Randomization (ASR) - put the stack in a randomly chosen memory location so it's hard to guess the location of the buffer.
 - Safe functions - use strncpy instead of strcpy, since strncpy ~~hasna~~ lets you specify the maximum number of characters to copy
 - Non-Executable stacks - prevent memory locations on the stack from being interpreted as code. (Requires hardware support)
 - StackGuard - prevent the attacker from overwriting the return address by detecting changes and terminating the program.
- Change stack frames to look like:



The canary value is chosen randomly when the program starts. Before the function returns, it checks to make sure that the canary is still the same. It is difficult (though not impossible) to overwrite the return address without changing the canary.

Note: Only using safe functions can prevent all buffer overflows. The other mechanisms mainly ~~prevent~~ ^{prevent} the standard stack-based overflow that we saw earlier.

Heap Overflows: (or overflows in other program regions)


- Possible, though harder to find since the heap layout is not as transparent as the stack layout
- Can still execute arbitrary code by overwriting function pointers
- or just overwrite data ...

Return-to-libc Attacks:

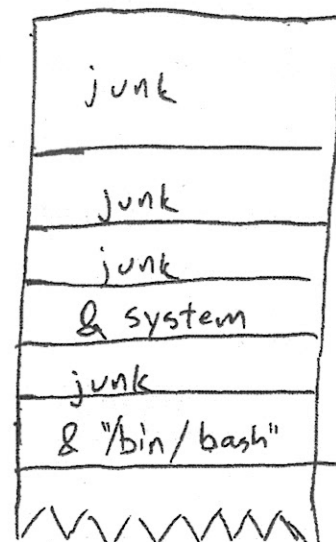
- Libc is a standard library including functions like printf(), exec(), etc.
- Basic idea: set up the stack to look like a function call to one (or more) functions in libc. It is possible to get a ^{root} shell by chaining several calls.
- Exploit works on non-executable stacks.

A hypothetical attack to execute system("/bin/sh"); and ~~gain~~ get a shell using the vulnerable program from before:

~~Attack~~ Stack Frame



overwrite
into →



- The return address is now the address of the system function.
- The argument to system is the address ~~to~~ ^{of} the string "/bin/bash" stored somewhere else in memory (e.g., in an environment variable).
- This will execute a shell, but it won't maintain the privileges of the executing program because system() drops privileges.

Format String Exploits!

- Format strings are arguments to printf containing special ~~characters~~ escape sequences that begin with "%"
- If programmers call printf() incorrectly, we can cause all kinds of trouble. (We can write arbitrary memory locations)

Notable Escape Sequences:

- %x - print a value in hexadecimal
- %s - interpret the argument as a pointer ~~then~~ to a char buffer (a string). Print the string.
- %n - save the number of bytes written so far to the ~~address~~ location pointed to by the argument

Some printf Examples:

`printf("%x", 16);` → prints "10"

```
char*foo = "abcd";  
int a = 10;
```


`printf("%x, %s %n", foo, foo, &a);` → prints the address of "abcd", then "abcd", then saves 13 in a (8 chars for the address, 1 space, 4 chars in "abcd").

`printf(argv[1]);` ← the wrong way to print a string.
Note that escape characters in `argv[1]` will be interpreted by `printf()`.

`printf("%s", argv[1]);` ← the right way to print a string.

- Format String Exploits occur when people use `printf()` ~~the~~ incorrectly to print strings.
- By including escape characters in the string, (especially `%n`), we can write arbitrary addresses.
- The arguments for the escape sequences ~~some~~ are calculated by adding an offset to the stack pointer

Normal printf call stack:



The location of the i th parameter P_i is ~~very~~ computed by adding to ESP, even if P_i wasn't provided in the call.

`printf("%x");` → prints ^{the} ~~some~~ hexadecimal value

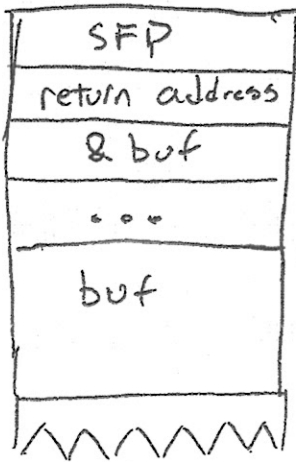
- If the format string is also allocated on the stack, we can control the arguments to the escape sequences as well.

~~Example~~

A Vulnerable Program: (ignore the ^{potential} buffer overflow...)

```
void main (int argc, char* int argv[]) {
    char buf [100];
    strcpy (buf, argv[1]);
    printf (buf);
}
```

- In the printf() call, the stack will look like:



Since buf is below the printf call, at some point printf will start using its contents as the arguments to the escape sequences. Relatively easy to find which escape sequence first reads its argument from buf.

- Can now use the %n sequence ~~to write~~ and control its argument (the address to write) => can write to arbitrary memory locations, and set them to values of our choosing.

The exploit string!

- Say we figure out that the k th `printf` ^{escape sequence} argument ~~usually~~ uses the first word of buffer as its argument.
- ~~Say~~ To write to `<address>`, our string looks like
"`<address> %x %x ... %x %n`"
 $\underbrace{\hspace{10em}}_{k-1 \text{ "%x"s"}}$
- This writes something like $4 + 8(k-1)$ to `<address>`; $4 + 8(k-1)$ is (probably) the length of the printed string.
- By using several `%n`'s, we can write any value we want.

Cross-Site Scripting (XSS):

- Attacks on websites to run some code on the client viewing the website
- Can be used to steal login information, cookies

Simple script (in PHP...)

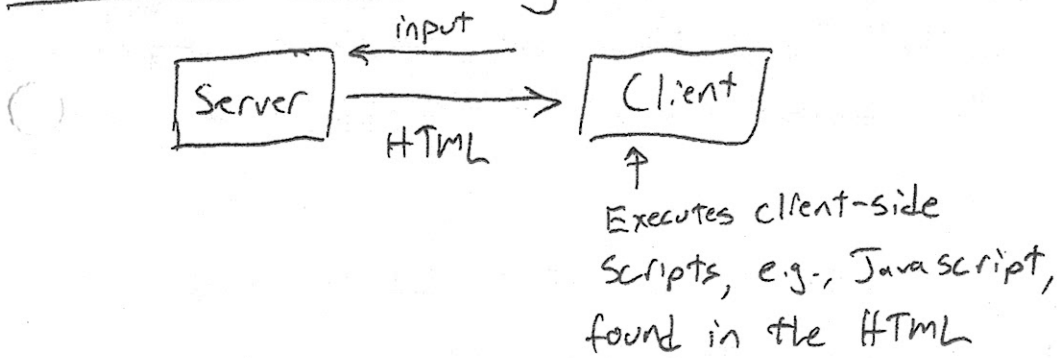
```
<html>
  <body> Hi
  <? echo $_GET["name"]; ?>
</body>
</html>
```

script.php → Hi

script.php?name=bob → Hi bob

Notes on
next page →

Cross-Site Scripting (XSS)



- XSS modifies the generated HTML to include ^{extra} scripts ~~that~~ on the page seen by the client.
- These extra scripts can steal login information and cookies from the client and send them to an adversary.

A simple vulnerable program:

```
<html>
```

```
Hi, <? echo $_GET["name"]; ?>
```

GET parameters are specified in the url after the ? character

```
</html>
```

Sample Run

```
script.php? name = Jayant
```

```
→ <html>
```

```
Hi, Jayant
```

```
</html>
```

```
script.php? name = <script> alert("hello"); </script>
```

```
→ <html>
```

```
Hi, <script> alert("hello"); </script>
```

```
</html>
```

By setting the name parameter, I can include Javascript on the page.

- Useful for phishing attacks: the attacker crafts a URL, then sends it to other people. For example, if my bank had an XSS vulnerability, ~~It~~ an attacker could generate a URL that included a script to steal login information, then send the link out in fake emails from the bank. When clients click on the link and log in to their accounts, the adversary's script will run and send the client's login information to the attacker.
- XSS attacks can occur whenever user-controlled values are printed out as part of an HTML document. ~~Unescaped~~

Prevention:

- Escape all variables before printing them out
 - Replace < with <, etc.
- (Note that the problem here is analogous to the problem that ~~causes~~^{enables} SQL injection attacks.)

SQL Injection Attacks

```
SELECT * FROM users  
WHERE username = '<username>'  
and password = '<password>';
```

Program generates a SQL query by replacing <username> and <password> with the specific values

Input:
username = admin

password = ' OR 't' = 't

⇒ SELECT * FROM users
WHERE username = 'admin'
AND password = '' OR 't' = 't';

user can log in as an admin without knowing the password!

More insidious:

password = ' ; DROP TABLE users ;

← Runs 2 SQL queries instead of 1, and the 2nd query deletes all users!

Prevention:

- Escape strings before putting them in SQL queries.

→ Replaces ' with \', which is not interpreted as a '.

- Most languages have libraries for this.

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

void test(char * buf) {
    char foo[12];
    int i;

    // Here's the target overflow
    strcpy(foo, buf);

    // Do a memory dump of the 10 words on the stack below foo
    for (i = 0; i < 10; i++) {
        printf("%08x: %08x \n", &((int *) foo)[i], ((int *) foo)[i]);
    }
}

int main(int argc, char * argv[]) {
    if (argc >= 1) {
        test(argv[1]);
        printf("%s \n", argv[1]);
    }

    printf("We can skip this print statement\n");
    printf("This is the end, my friend \n");
    return 0;
}
```

```
$ gcc -fno-stack-protector -o overflow overflow.c
$ ./overflow AAA
bffff518: 00414141
bffff51c: b7eeclae
bffff520: b7f91b19
bffff524: 00000003
bffff528: bffff548
bffff52c: 0804845e
bffff530: bffff73b
bffff534: 080496a0
bffff538: bffff558
bffff53c: 080484c9
AAA
```

We can skip this print statement
This is the end, my friend

```
$ ./overflow AAAABBBBCCCCDDDEEEFF
```

```
bffff4f8: 41414141
bffff4fc: 42424242
bffff500: 43434343
bffff504: 00000003
bffff508: 45454545
bffff50c: 00464646
bffff510: bffff727
bffff514: 080496a0
bffff518: bffff538
bffff51c: 080484c9
```

Segmentation fault

```
$ gdb overflow
```

```
GNU gdb 6.8-debian
```

```
Copyright (C) 2008 Free Software Foundation, Inc.
```

```
License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>
```

```
This is free software: you are free to change and redistribute it.
```

```
There is NO WARRANTY, to the extent permitted by law. Type "show copying"  
and "show warranty" for details.
```

```
This GDB was configured as "i486-linux-gnu"...
```

```
(gdb) disas main
```

```
Dump of assembler code for function main:
```

```
0x0804842f <main+0>: lea 0x4(%esp),%ecx
0x08048433 <main+4>: and $0xffffffff0,%esp
0x08048436 <main+7>: pushl -0x4(%ecx)
0x08048439 <main+10>: push %ebp
0x0804843a <main+11>: mov %esp,%ebp
0x0804843c <main+13>: push %ecx
0x0804843d <main+14>: sub $0x14,%esp
0x08048440 <main+17>: mov %ecx,-0x8(%ebp)
0x08048443 <main+20>: mov -0x8(%ebp),%eax
0x08048446 <main+23>: cmpl $0x0,(%eax)
0x08048449 <main+26>: jle 0x8048479 <main+74>
0x0804844b <main+28>: mov -0x8(%ebp),%edx
0x0804844e <main+31>: mov 0x4(%edx),%eax
0x08048451 <main+34>: add $0x4,%eax
0x08048454 <main+37>: mov (%eax),%eax
0x08048456 <main+39>: mov %eax,(%esp)
0x08048459 <main+42>: call 0x80483d4 <test>
0x0804845e <main+47>: mov -0x8(%ebp),%edx
0x08048461 <main+50>: mov 0x4(%edx),%eax
0x08048464 <main+53>: add $0x4,%eax
0x08048467 <main+56>: mov (%eax),%eax
0x08048469 <main+58>: mov %eax,0x4(%esp)
0x0804846d <main+62>: movl $0x804856d,(%esp)
0x08048474 <main+69>: call 0x8048330 <printf@plt>
0x08048479 <main+74>: movl $0x8048574,(%esp)
0x08048480 <main+81>: call 0x8048340 <puts@plt>
0x08048485 <main+86>: movl $0x8048595,(%esp)
0x0804848c <main+93>: call 0x8048340 <puts@plt>
0x08048491 <main+98>: mov $0x0,%eax
0x08048496 <main+103>: add $0x14,%esp
0x08048499 <main+106>: pop %ecx
0x0804849a <main+107>: pop %ebp
0x0804849b <main+108>: lea -0x4(%ecx),%esp
0x0804849e <main+111>: ret
```

```
End of assembler dump.
```

```
(gdb) quit
```

```
$ hexdump exploit
```

```
00000000 4141 4141 4141 4141 4141 4141 4141 4141
```

```
00000010 4141 4141 8485 0804
```

```
00000018
```

```
$ ./overflow `cat exploit`b
```

```
ffff4f8: 41414141
bffff4fc: 41414141
bffff500: 41414141
bffff504: 00000003
bffff508: 41414141
bffff50c: 08048485
bffff510: bffff700
bffff514: 080496a0
bffff518: bffff538
bffff51c: 080484c9
This is the end, my friend
$
```