


Lecture 6B

Machine-Level Programming V: Miscellaneous Topics

Topics

- Linux Memory Layout
- Understanding Pointers
- Buffer Overflow



Linux Memory Layout

Stack

- Runtime stack (8MB limit)

Heap

- Dynamically allocated storage
- When call malloc, calloc, new

DLLs

- Dynamically Linked Libraries
- Library routines (e.g., printf, malloc)
- Linked into object code when first executed

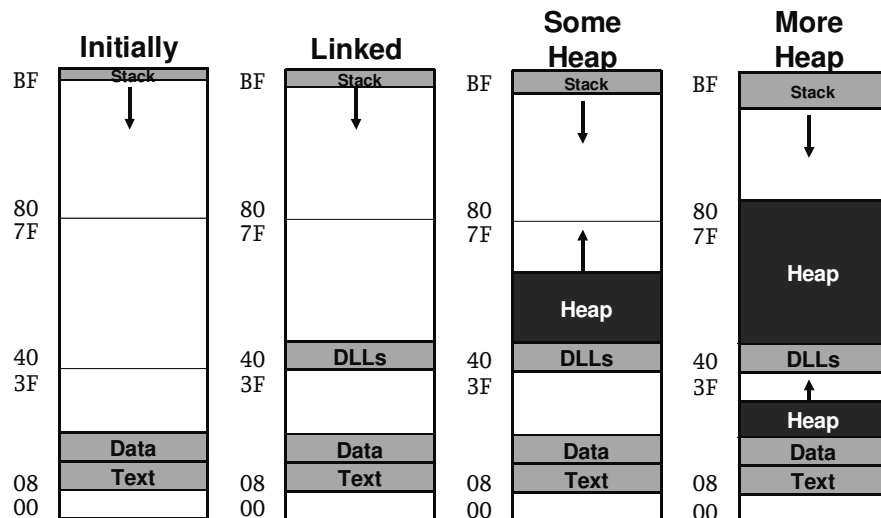
Data

- Statically allocated data
- E.g., arrays & strings declared in code

Text

- Executable machine instructions
- Read-only

Linux Memory Allocation



Text & Stack Example


```
(gdb) break main
(gdb) run
Breakpoint 1, 0x804856f in main ()
(gdb) print $esp
$3 = (void *) 0xbffffc78
```

Main

- Address 0x804856f should be read 0x0804856f

Stack

- Address 0xbffffc78



Dynamic Linking Example


```
(gdb) print malloc
$1 = {<text variable, no debug info>
0x8048454 <malloc>}
(gdb) run
Program exited normally.
(gdb) print malloc
$2 = {void *(unsigned int)}
0x40006240 <malloc>
```

Initially

- Code in text segment that invokes dynamic linker
- Address 0x8048454 should be read 0x40006240

Final

- Code in DLL region



Memory Allocation Example


```
char big_array[1<<24]; /* 16 MB */
char huge_array[1<<28]; /* 256 MB */

int beyond;
char *p1, *p2, *p3, *p4;

int useless() { return 0; }

int main()
{
    p1 = malloc(1 <<28); /* 256 MB */
    p2 = malloc(1 << 8); /* 256 B */
    p3 = malloc(1 <<28); /* 256 MB */
    p4 = malloc(1 << 8); /* 256 B */
    /* Some print statements ... */
}
```

Example Addresses



C operators

Operators

```
() [] -> .
! ~ ++ -- + - * & (type) sizeof
* / %
+ -
<< >>
< <= > >=
== !=
&
^
|
&&
||
?:
= += -= *= /= %= &= ^= != <<= >>=
,
```

Associativity

```
left to right
right to left
left to right
left to right
left to right
left to right
left to right
left to right
left to right
left to right
right to left
right to left
left to right
```

Note: Unary +, -, and * have higher precedence than binary forms

C pointer declarations

<code>int *p</code>	p is a pointer to int
<code>int *p[13]</code>	p is an array[13] of pointer to int
<code>int *(p[13])</code>	p is an array[13] of pointer to int
<code>int **p</code>	p is a pointer to a pointer to an int
<code>int (*p)[13]</code>	p is a pointer to an array[13] of int
<code>int *f()</code>	f is a function returning a pointer to int
<code>int (*f)()</code>	f is a pointer to a function returning int
<code>int (*(*f()) [13])()</code>	f is a function returning ptr to an array[13] of pointers to functions returning int
<code>int (*(*x[3])()) [5]</code>	x is an array[3] of pointers to functions returning pointers to array[5] of ints


Internet Worm and IM War

November, 1988

- Internet Worm attacks thousands of Internet hosts.
- How did it happen?

July, 1999

- Microsoft launches MSN Messenger (instant messaging system).
- Messenger clients can access popular AOL Instant Messaging Service (AIM) servers



Internet Worm and IM War (cont.)

August 1999

- Mysteriously, Messenger clients can no longer access AIM servers.
- Microsoft and AOL begin the IM war:
 - AOL changes server to disallow Messenger clients
 - Microsoft makes changes to clients to defeat AOL changes.
 - At least 13 such skirmishes.
- How did it happen?

The Internet Worm and AOL/Microsoft War were both based on *stack buffer overflow* exploits!

- many Unix functions do not check argument sizes.
- allows target buffers to overflow.

String Library Code

- Implementation of Unix function gets
 - No way to specify limit on number of characters to read

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

- Similar problems with other Unix functions
 - strcpy: Copies string of arbitrary length
 - scanf, fscanf, sscanf, when given %s conversion specification

Vulnerable Buffer Code

```

/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    gets(buf);
    puts(buf);
}

```

```

int main()
{
    printf("Type a string:");
    echo();
    return 0;
}

```

Buffer Overflow Executions

```

unix> ./bufdemo
Type a string:123
123

```

```

unix> ./bufdemo
Type a string:12345
Segmentation Fault


```

```

unix> ./bufdemo
Type a string:12345678
Segmentation Fault

```

Buffer Overflow Stack



```

echo:
    pushl %ebp          # Save %ebp on stack
    movl %esp,%ebp
    subl $20,%esp      # Allocate space on stack
    pushl %ebx         # Save %ebx
    addl $-12,%esp     # Allocate space on stack
    leal -4(%ebp),%ebx # Compute buf as %ebp-4
    pushl %ebx         # Push buf on stack
    call gets          # Call gets
    . . .


```

Buffer Overflow Stack Example

```

unix> gdb bufdemo
(gdb) break echo
Breakpoint 1 at 0x8048583
(gdb) run
Breakpoint 1, 0x8048583 in echo ()
(gdb) print /x *(unsigned *)$ebp
$1 = 0xbffff8f8
(gdb) print /x *((unsigned *)$ebp + 1)
$3 = 0x804864d

```

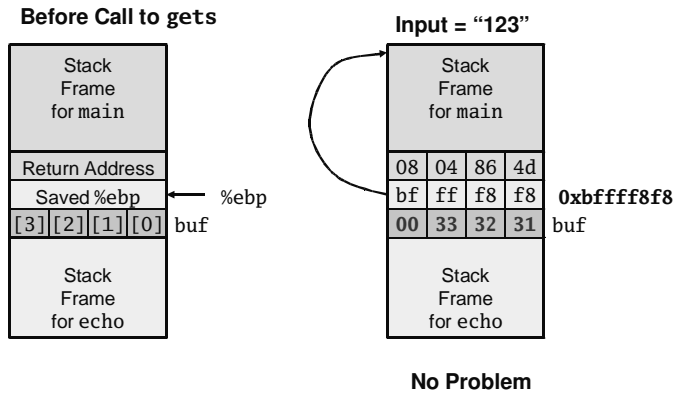


```

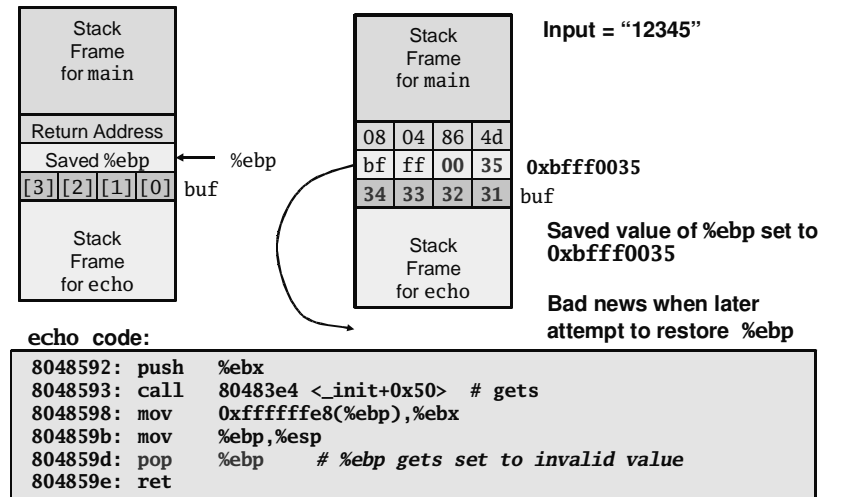
8048648: call 804857c <echo>
804864d: mov 0xfffffe8(%ebp),%ebx # Return Point

```

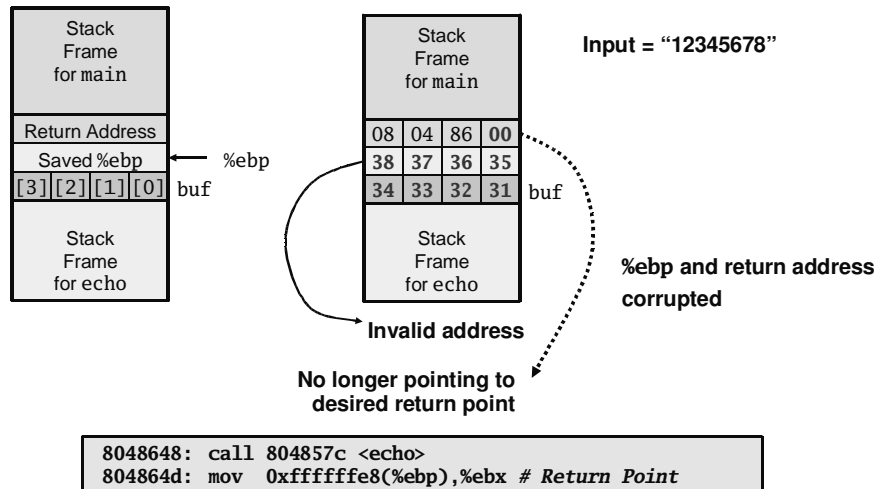
Buffer Overflow Example #1



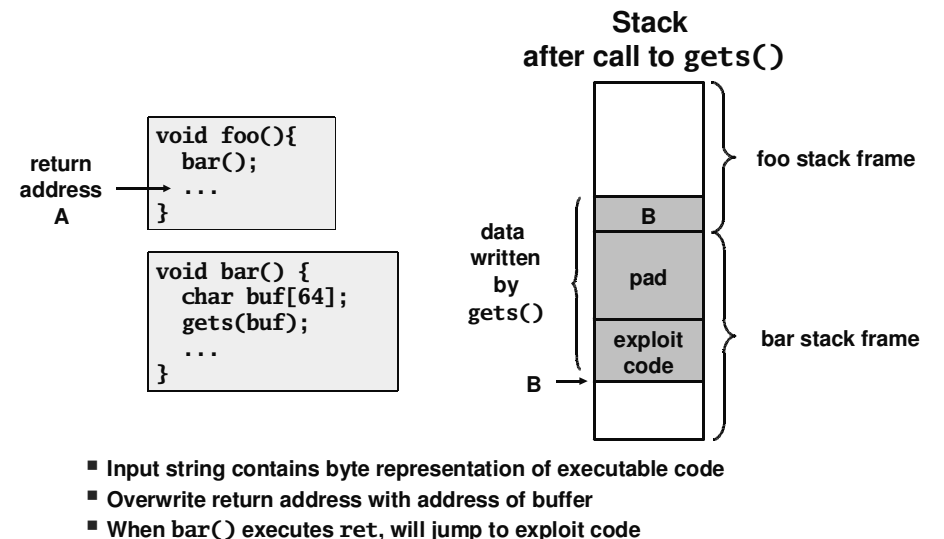
Buffer Overflow Stack Example #2



Buffer Overflow Stack Example #3



Malicious Use of Buffer Overflow



Exploits Based on Buffer Overflows

Buffer overflow bugs allow remote machines to execute arbitrary code on victim machines.

Internet worm

- 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 by sending phony argument:
 - *finger "exploit-code padding new-return-address"*
 - exploit code: executed a root shell on the victim machine with a direct TCP connection to the attacker.

Avoiding Overflow Vulnerability

```
/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    fgets(buf, 4, stdin);
    puts(buf);
}
```

Use Library Routines that Limit String Lengths

- fgets instead of gets
- strncpy instead of strcpy
- Don't use scanf with %s conversion specification
 - Use fgets to read the string

Final Observations

Memory Layout

- OS/machine dependent (including kernel version)
- Basic partitioning: stack/data/text/heap/DLL found in most machines

Type Declarations in C

- Notation obscure, but very systematic

Working with Strange Code

- Important to analyze nonstandard cases
 - E.g., what happens when stack corrupted due to buffer overflow
- Helps to step through with GDB