

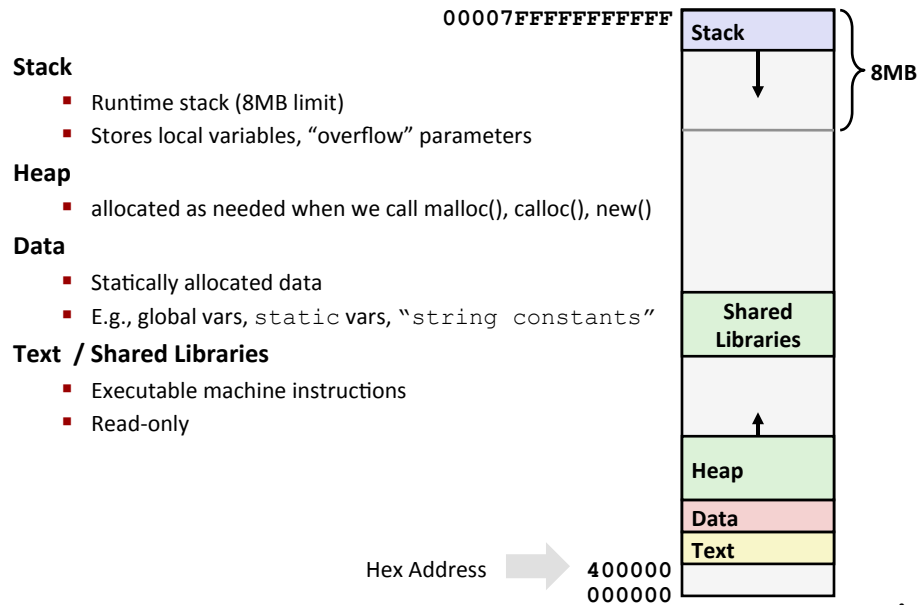
# CS 370: Computer Organization

## Overflow, Worms, and Viruses

Instructor:  
Professor Stephen P. Carl

### x86-64 Linux Memory Layout

*not drawn to scale*



## Memory Referencing Bug Example

```
typedef struct {
    int a[2];
    double d;
} struct_t;

double fun(int i) {
    volatile struct_t s;
    s.d = 3.14;
    s.a[i] = 1073741824; /* Possibly out of bounds */
    return s.d;
}
```

```
fun(0)  →  3.14
fun(1)  →  3.14
fun(2)  →  3.1399998664856
fun(3)  →  2.00000061035156
fun(4)  →  3.14
fun(6)  →  Segmentation fault
```

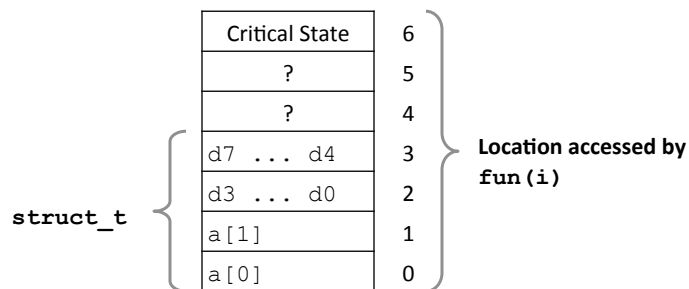
\*\*\* Result is system specific \*\*\*

## Memory Referencing Bug Example

```
typedef struct {
    int a[2];
    double d;
} struct_t;
```

```
fun(0)  →  3.14
fun(1)  →  3.14
fun(2)  →  3.1399998664856
fun(3)  →  2.00000061035156
fun(4)  →  3.14
fun(6)  →  Segmentation fault
```

### Explanation:



## Such problems are a BIG deal

- **Generally called a “buffer overflow”**
  - when exceeding the memory size allocated for an array
- **Why a big deal?**
  - It's the #1 technical cause of security vulnerabilities
  - #1 overall cause is social engineering / user ignorance
- **Common forms**
  - Unchecked lengths on string inputs
  - Bounded character arrays on the stack
- Buffer overruns on the stack sometimes referred to as **stack smashing**

7

## 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 library functions**
  - **strcpy, strcat**: Copy strings of arbitrary length
  - **scanf, fscanf, sscanf**, when given %s conversion specification

8

## Vulnerable Buffer Code

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

← btw, how big is big enough?

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

```
unix> ./bufdemo-nsp
Type a string: 012345678901234567890123
012345678901234567890123
```

```
unix> ./bufdemo-nsp
Type a string: 0123456789012345678901234
Segmentation Fault
```

9

## Buffer Overflow Disassembly

echo:

```
00000000004006cf <echo>:
4006cf: 48 83 ec 18      sub    $0x18,%rsp
4006d3: 48 89 e7         mov    %rsp,%rdi
4006d6: e8 a5 ff ff ff  callq 400680 <gets>
4006db: 48 89 e7         mov    %rsp,%rdi
4006de: e8 3d fe ff ff  callq 400520 <puts@plt>
4006e3: 48 83 c4 18     add    $0x18,%rsp
4006e7: c3              retq
```

call\_echo:

```
4006e8: 48 83 ec 08     sub    $0x8,%rsp
4006ec: b8 00 00 00 00  mov    $0x0,%eax
4006f1: e8 d9 ff ff ff  callq 4006cf <echo>
4006f6: 48 83 c4 08     add    $0x8,%rsp
4006fa: c3              retq
```

10

Sewanee: The University of the South

## Buffer Overflow Stack

*Before call to gets*

Stack Frame for call\_echo

---

Return Address (8 bytes)

---

20 bytes unused

---

[3]	[2]	[1]	[0]
-----	-----	-----	-----

buf ← %rsp

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

```
echo:
    subq $24, %rsp
    movq %rsp, %rdi
    call gets
    . . .
```

11

Sewanee: The University of the South

## Buffer Overflow Stack Example

*Before call to gets*

Stack Frame for call\_echo

---

00	00	00	00
00	40	06	f6

---

20 bytes unused

---

[3]	[2]	[1]	[0]
-----	-----	-----	-----

buf ← %rsp

```
void echo()
{
    char buf[4];
    gets(buf);
    . . .
}
```

```
echo:
    subq $24, %rsp
    movq %rsp, %rdi
    call gets
    . . .
```

**call echo:**

```
. . .
4006f1: callq 4006cf <echo>
4006f6: add $0x8,%rsp
. . .
```

12

## Buffer Overflow Stack Example #1

After call to gets

Stack Frame  
for call\_echo

00	00	00	00
00	40	06	f6
00	32	31	30
39	38	37	36
35	34	33	32
31	30	39	38
37	36	35	34
33	32	31	30

buf ← %rsp

```
void echo()
{
    char buf[4];
    gets(buf);
    . . .
}
```

```
echo:
    subq $24, %rsp
    movq %rsp, %rdi
    call gets
    . . .
```

call\_echo:

```
. . .
4006f1: callq 4006cf <echo>
4006f6: add $0x8,%rsp
. . .
```

```
unix> ./bufdemo-nsp
Type a string: 01234567890123456789012
01234567890123456789012
```

Overflowed buffer, but did not corrupt state

## Buffer Overflow Stack Example #2

After call to gets

Stack Frame  
for call\_echo

00	00	00	00
00	40	00	34
33	32	31	30
39	38	37	36
35	34	33	32
31	30	39	38
37	36	35	34
33	32	31	30

buf ← %rsp

```
void echo()
{
    char buf[4];
    gets(buf);
    . . .
}
```

```
echo:
    subq $24, %rsp
    movq %rsp, %rdi
    call gets
    . . .
```

call\_echo:

```
. . .
4006f1: callq 4006cf <echo>
4006f6: add $0x8,%rsp
. . .
```

```
unix> ./bufdemo-nsp
Type a string: 0123456789012345678901234
Segmentation Fault
```

Overflowed buffer and corrupted return pointer

Sewanee: The University of the South

## Buffer Overflow Stack Example #3

*After call to gets*

Stack Frame  
for call\_echo

00	00	00	00
00	40	06	00
33	32	31	30
39	38	37	36
35	34	33	32
31	30	39	38
37	36	35	34
33	32	31	30

buf ← %rsp

```
void echo()
{
    char buf[4];
    gets(buf);
    . . .
}
```

```
echo:
    subq $24, %rsp
    movq %rsp, %rdi
    call gets
    . . .
```

**call\_echo:**

```
. . .
4006f1: callq 4006cf <echo>
4006f6: add $0x8,%rsp
. . .
```

```
unix> ./bufdemo-nsp
Type a string: 012345678901234567890123
012345678901234567890123
```

**Overflowed buffer, corrupted return pointer, but program seems to work!**

15

Sewanee: The University of the South

## Buffer Overflow Stack Example #3 Explained

*After call to gets*

Stack Frame  
for call\_echo

00	00	00	00
00	40	06	00
33	32	31	30
39	38	37	36
35	34	33	32
31	30	39	38
37	36	35	34
33	32	31	30

buf ← %rsp

**register\_tm\_clones:**

```
. . .
400600: mov %rsp,%rbp
400603: mov %rax,%rdx
400606: shr $0x3f,%rdx
40060a: add %rdx,%rax
40060d: sar %rax
400610: jne 400614
400612: pop %rbp
400613: retq
```

**“Returns” to unrelated code**  
**Lots of things happen, without modifying critical state**  
**Eventually executes retq back to main**

16


## Code Injection Attacks

```

void P() {
    Q();
    ...
}

int Q() {
    char buf[64];
    gets(buf);
    ...
    return ...;
}
    
```

return address A



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

17

## Exploits Based on Buffer Overflows

- *Buffer overflow bugs can allow remote machines to execute arbitrary code on victim machines*
- **Distressingly common in real programs**
  - Programmers keep making the same mistakes ☹
  - Recent measures make these attacks much more difficult
- **Examples across the decades**
  - Original "Internet worm" (1988)
  - "IM wars" (1999)
  - Twilight hack on Wii (2000s)
  - ... and many, many more

18



## Worms and Viruses

- **Worm: A program that**
  - Can run by itself
  - Can propagate a fully working version of itself to other computers
  
- **Virus: Code that**
  - Adds itself to other programs
  - Does not run independently
  
- **Both are (usually) designed to spread among computers and to wreak havoc**

19

## 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 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.
- **Once on a machine, scanned for other machines to attack**
  - invaded ~6000 computers in hours (10% of the Internet in 1988!)
    - see June 1989 article in *Comm. of the ACM*
  - the young author of the worm was prosecuted...
  - and CERT was formed... still housed at CMU
  - CERT = Computer Emergency Response Team

20

## Example 2: IM War

July, 1999

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



21

## 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
- What was really happening?
  - AOL had discovered a buffer overflow bug in their own AIM clients
  - They exploited it to detect and block Microsoft: the exploit code returned a 4-byte signature (the bytes at some location in the AIM client) to server
  - When Microsoft changed code to match signature, AOL changed signature location

22

## How to avoid buffer overflow attacks

- Avoid overflow vulnerabilities
- Employ system-level protections
- Have compiler use “stack canaries”

*Lets talk about each...*

24

## 1. Avoid Overflow Vulnerabilities in Code (!)

```

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

### ALWAYS use library routines that limit string lengths

- `fgets` instead of `gets`
- `strncpy` instead of `strcpy`
- Don't use `scanf` with `%s` conversion specification. Instead:
  - Use `fgets` to read the string
  - Or use `%ns` where `n` is a suitable integer

25

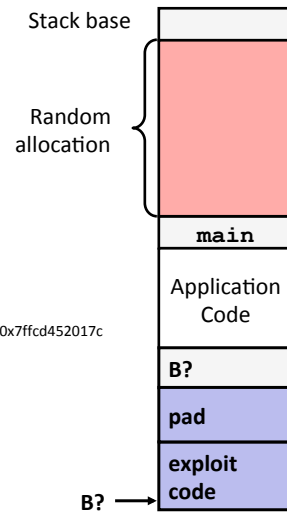
## 2. System-Level Protections

### Randomized stack offsets

- At start of program, allocate random amount of space on stack
- Shifts stack addresses for entire program
- Makes it difficult for hacker to predict beginning of inserted code
- 5 example executions of memory allocation code:

local      0x7ffe4d3be87c    0x7fff75a4f9fc    0x7ffeadb7c80c    0x7ffeaea2fdac    0x7ffcd452017c

- Stack repositioned each time program executes




26

## 2. System-Level Protections

### Nonexecutable code segments

- IA-32 and older code can mark region of memory as either "read-only" or "writeable"
  - Can execute anything readable
- X86-64 added explicit "execute" permission
- Stack marked as non-executable



Any attempt to execute this code will fail

27

### 3. Stack Canaries

#### Idea:

- Place special value (“canary”) on stack just beyond buffer
- Check for corruption before exiting function

#### gcc Implementation

- Uses option `-fstack-protector`
- This option is now the default (previously, was disabled)

```
unix> ./bufdemo-sp
Type a string: 0123456
0123456
```

```
unix> ./bufdemo-sp
Type a string: 01234567
*** stack smashing detected ***
```

28

### Protected Buffer Disassembly


#### echo:

```
40072f: sub    $0x18,%rsp
400733: mov    %fs:0x28,%rax
40073c: mov    %rax,0x8(%rsp)
400741: xor    %eax,%eax
400743: mov    %rsp,%rdi
400746: callq 4006e0 <gets>
40074b: mov    %rsp,%rdi
40074e: callq 400570 <puts@plt>
400753: mov    0x8(%rsp),%rax
400758: xor    %fs:0x28,%rax
400761: je     400768 <echo+0x39>
400763: callq 400580 <__stack_chk_fail@plt>
400768: add    $0x18,%rsp
40076c: retq
```

29

## Setting Up Canary

*Before call to gets*




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

```
echo:
    . . .
    movq    %fs:40, %rax # Get canary
    movq    %rax, 8(%rsp) # Place on stack
    xorl    %eax, %eax # Erase canary
    . . .
```

30

## Checking Canary

*After call to gets*



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

Input: 0123456

```
echo:
    . . .
    movq    8(%rsp), %rax # Retrieve from stack
    xorq    %fs:40, %rax # Compare to canary
    je     .L6 # If same, OK
    call   __stack_chk_fail # FAIL
.L6:
    . . .
```

## Return-Oriented Programming Attacks

- **Challenge (for hackers)**
  - Stack randomization makes it hard to predict buffer location
  - Marking stack nonexecutable makes it hard to insert binary code
- **Alternative Strategy**
  - Use existing code
    - E.g., library code from `stdlib`
  - String together fragments to achieve overall desired outcome
  - (Note: *does not overcome stack canaries*)
- **Construct program from *gadgets***
  - Gadget is a sequence of instructions ending in `ret`
    - Encoded by single byte `0xc3`
  - Code positions fixed from run to run
  - Code is executable

32

## Gadget Example #1

```
long ab_plus_c
(long a, long b, long c)
{
    return a*b + c;
}
```

```
0000000004004d0 <ab_plus_c>:
4004d0: 48 0f af fe  imul %rsi,%rdi
4004d4: 48 8d 04 17  lea (%rdi,%rdx,1),%rax
4004d8: c3           retq
```

$rax \leftarrow rdi + rdx$


Gadget address = `0x4004d4`

- Use tail end of existing functions

33


## Gadget Example #2

```
void setval(unsigned *p) {
    *p = 3347663060u;
}
```



- Repurpose byte codes

## ROP Execution



- Trigger with `ret` instruction
  - Starts executing Gadget 1
- Final `ret` in each gadget will start next one