

```
29 #include <balloon.h>
30
31 int struct {
32     ScePspFVector3 mode;
33     ScePspFVector3 pos;
34     int sbuf[3];
35     float scnt;
36     t;
37 } BALLOONDAT;
38
39 static BALLOONDAT balloon;
40 static ScePspFVector3 sphere[28];
41 static ScePspFVector3 pole[28];
42
43 extern void DrawSphere(ScePspFVector3 *array, float r);
44 extern void DrawPole(ScePspFVector3 *array, float r);
45
46 void init_balloon(void)
47 {
48     int i;
49
50     balloon.mode=MODE
51     balloon.pos.x= 0.
52     balloon.pos.y=-8.
53     balloon.pos.z= 0.
54     balloon.t=0.0f;
55     balloon.scnt=2;
56
57     for (i=0; i<3; i++)
58         balloon.sbuf[i]=0;
59
60 }
61
62 void draw_balloon(void)
63 {
64     ScePspFVectors vec;
65     glTranslatef(SCEGU_TEXTURE);
66     glTranslatef(balloon.pos);
67 }
```

# Operating Systems and C Fall 2022, Security-Track 6. Stack-Based Exploits

exciting part of the course!

**performance-track lecture nr. 1:**  
*what you'll need for perflab.*  
array layout, what it means for performance, cache hierarchy, how associativity is organized in the cache. (important stuff for anyone)

## perflab:

- have two matrix multiplication procedures (rotate, smooth), have to rewrite it.
- about optimization techniques; blocking, loop unrolling, etc.

**performance-track lecture nr. 2:**  
*what you'll need for perflab.*  
*optimizations.* how to write code so compiler can derive performant code. manual transformations, blocking, loop unrolling

## attacklab:

- you have an executable, have to attack it.
- about the stack; code injection (smashing the stack), return-oriented programming (find interesting code in other programs).

**security-track lecture nr. 1:**  
*what you'll need for attacklab.*

## Brief review of assembly

## Arrays, strings and structs in assembly

## Structure of an .s program

## Stack-based exploits and counter-measures

basically  
what you need  
for the **attacklab**

(next week: not  
important for attacklab  
(culture-stuff) )

alignment; needed in  
assignment

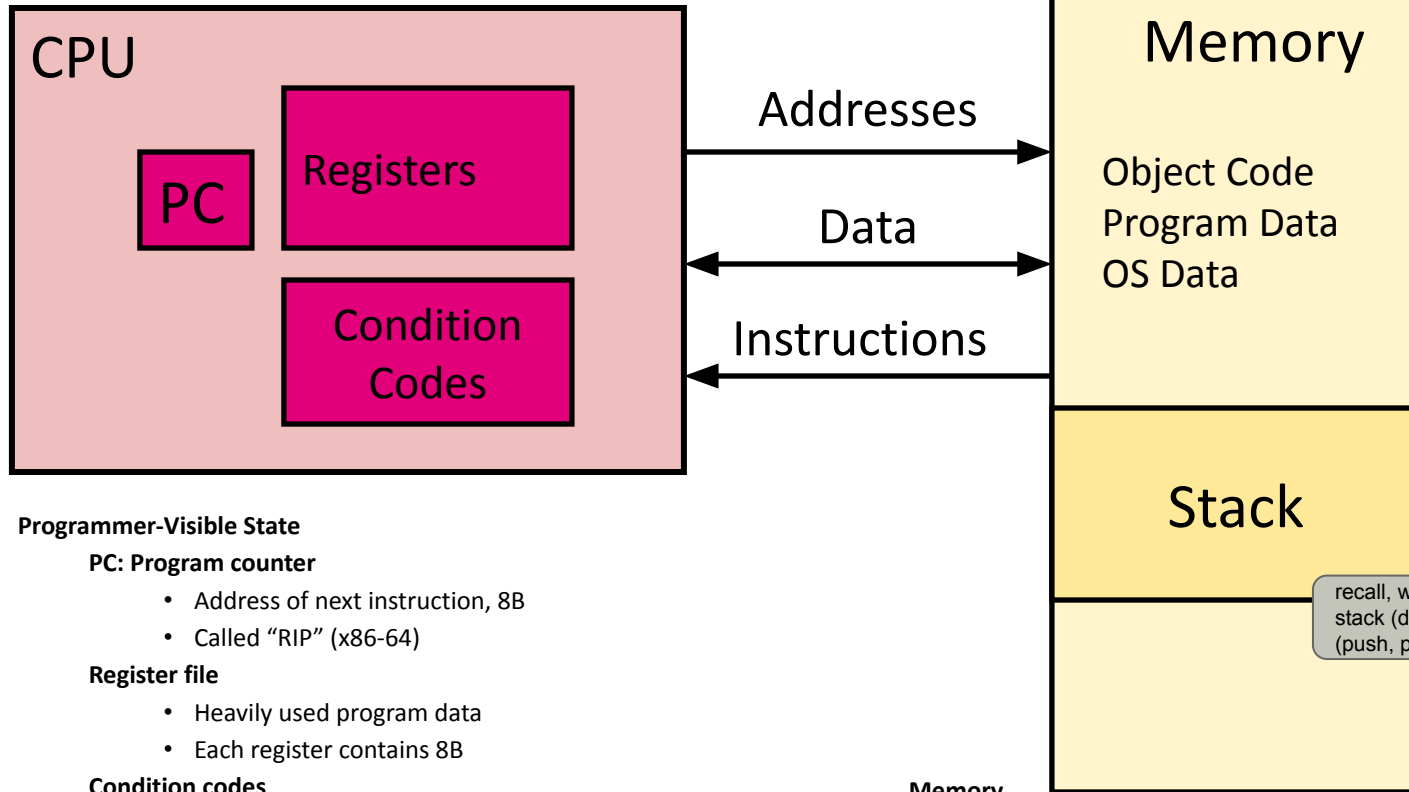
you'll need to  
disassemble a binary

# X86-64 Assembly

von Neumann Architecture

instruction either

- op on registers (state), OR
- transfers data to/from mem



## Programmer-Visible State

### PC: Program counter

- Address of next instruction, 8B
- Called "RIP" (x86-64)

### Register file

- Heavily used program data
- Each register contains 8B

### Condition codes

- Store status information about most recent arithmetic operation
- Used for conditional branching

## Memory

- Byte addressable array
- Code, user data, (some) OS data
- Includes stack used to support procedures

recall, what is a stack (data structure)? (push, pop)

aka. **program counter (pc)**

The **%rip** register is the current **instruction pointer**.  
Contains address of next instruction to be executed.

most instructions implicitly increment it.  
explicitly updated  $\Rightarrow$  change in control flow.

There are **16 general purpose registers** in x86-64.  
Additional registers for floating point, SIMD, ...  
16 registers: r0, r1, ..., r15

“register file”

# General-Purpose Registers

For historical reasons, r0-r7 are called **original registers**.

They have the following names:

- ax: register a
- bx: register b
- cx: register c
- dx: register d
  
- bp: register **b**ase **p**ointer (start of stack)
- sp: register **s**tack **p**ointer (current location in stack, grow downwards)
  
- si: register source index (source for data copies)
- di: register destination index (destination for data copies)

top & bottom  
for a given frame

usually first  
parameters of  
functions  
(if not arrays)

# General-Purpose Registers

memory is  
byte-addressable

why: e.g.  
C short is 2B

Register values can be accessed at different levels of granularity:

- **8B:**
  - original registers: **prefix r** rax, rsp, rsi
  - other registers: no suffix r8, r15
- **4B:**
  - original registers: **prefix e** eax, esp, esi
  - other registers: **suffix d** r8d, r15d
- **2B:**
  - original registers: **no prefix** ax, sp, si
  - other registers: **suffix w** r8w, r15w
- **1B (high byte):**
  - original registers (bits 8-15 from ax-dx) ah, bh, ch, dh
- **1B (low byte):**
  - original registers (bits 0-7 from ax-dx) al, bl, cl, dl
  - other registers: suffix b r8b, r15b

# Instructions

Three classes of instructions:

## 1. Transfer between memory and register

- Load/store data: register  $\leftrightarrow$  memory (e.g. mov)
- Push/pop: register  $\leftrightarrow$  stack

## 2. Arithmetic and comparison functions

we didn't talk much about these, except we walked about how these are built from logic gates

## 3. Transfer control

- Jumps to/from procedures
- Conditional branches

can loop

can if/then



The GNU tools (gcc, gdb) use AT&T Syntax for assembly.

example: `movq %rsp, %rbp`

Syntax is of the form

**OPERATOR source, destination**

never more than 2 operands.  
when there are 2, this is the form.

Register names are prefixed with %


The **alternative** is the **Intel syntax** (on windows): `MOVQ EBP, ESP` – no %

Look for % in the assembly code, if they are present you are dealing with AT&T syntax

# Procedure Call Example

```
804854e: e8 3d 06 00 00  call 8048b90 <main>
8048553: 50                pushl %eax
```

return address  
= address of next instruction




top of stack

`%rsp` 0x108

`%rip` 0x80854e

`call 8048b90`




instruction pointer  
updated to callee

`%rsp` 0x100

`%rip` 0x8048b90

# Procedure Call Example

```
8048591: c3                ret
```



**Q:** what happens if user input overwrites stack pointer, or return address?

**A:** user has full control over control flow. (exploit)

(garbage)

just a pop

# Stack Frame

## stack frame

- caller's saved registers<sup>(1)</sup>
- local vars
- args
- return address<sup>(2)</sup>

<sup>(1)</sup>: to restore caller state on return (bottom frame has no caller).

<sup>(2)</sup>: pushes it when it calls (top frame does not need this)

## Caller:

- Arguments
  - pushed by program (if needed)
- Return address
  - pushed by call

## Callee:


- Previous frame pointer (%rbp)
- Other callee-save registers (%rbx, %r12-15)
- Space for local variables
- Arguments for next function (when about to call another function)

extremely important for assignment; you'll spend hours seeing what stack looks like.

Caller Frame

Frame pointer  
→  
%rbp

Stack pointer  
→  
%rsp



Brief review of assembly

**Arrays, strings and structs in assembly**

Structure of an .s program

Stack-based exploits and counter-measures

# Array Allocation

## Basic Principle

compiler is going to reserve space for the array.


$T \ A[L];$

A is an Array of data type  $T$  and length  $L$

Contiguously allocated region of  $L * \mathbf{sizeof}(T)$  bytes

examples follow

# Array Allocation



8 bytes = 64 bits  
(address size =  
word size)

# Array Access

```
int A[5] = {0, 1, 2, 3, 4};
```

Array of data type *int* and length 5

Identifier **A** can be used as a pointer to array element 0: Type *int\**

val from  
previous slide

Reference	Type	Value
<code>val[4]</code>	<code>int</code>	<code>4</code>
<code>val</code>	<code>int *</code>	<code>x</code>
<code>val+1</code>	<code>int *</code>	<code>x + 4</code>
<code>&amp;val[2]</code>	<code>int *</code>	<code>x + 8</code>
<code>val[5]</code>	<code>int</code>	<code>??</code>
<code>*(val+1)</code>	<code>int</code>	<code>1</code>
<code>val + i</code>	<code>int *</code>	<code>x + 4 i</code>

undefined  
behavior




# Array Example

type alias

```
#define ZLEN 5
typedef int zip_dig[ZLEN];

zip_dig cmu = { 1, 5, 2, 1, 3 };
zip_dig mit = { 0, 2, 1, 3, 9 };
zip_dig ucb = { 9, 4, 7, 2, 0 };
```

(ucb = berkeley)



20 = 5\*4

Declaration “`zip_dig cmu`” equivalent to “`int cmu[5]`”  
Example arrays were allocated in successive 20 byte blocks

Not guaranteed to happen in general

# Array Example

```
void zincr(zip_dig z) {  
    int i;  
    for (i = 0; i < ZLEN; i++)  
        z[i]++;  
}
```

addressing mechanism:  
 $z + 4*i$

```
# rdx = z  
movq    $0, %rax          # %rax = i  
.L4:      # loop:  
addq    $1, (%rdx,%rax,4) # z[i]++  
addq    $1, %rax          # i++  
cmpl    $5, %rax         # i:ZLEN (ZLEN==5)  
jne     .L4              # if !=, goto loop
```

\$1 is  
constant 1

%r is  
register r

# Strings

Declared as read-only data.


```
.section      .rodata  
.Label:  
.string      "String constant\n"
```

1 byte per character,  
terminated by \0

we won't spend more time  
on strings, because you  
won't need them in the  
assignment.

# Structures & Alignment

## Unaligned Data



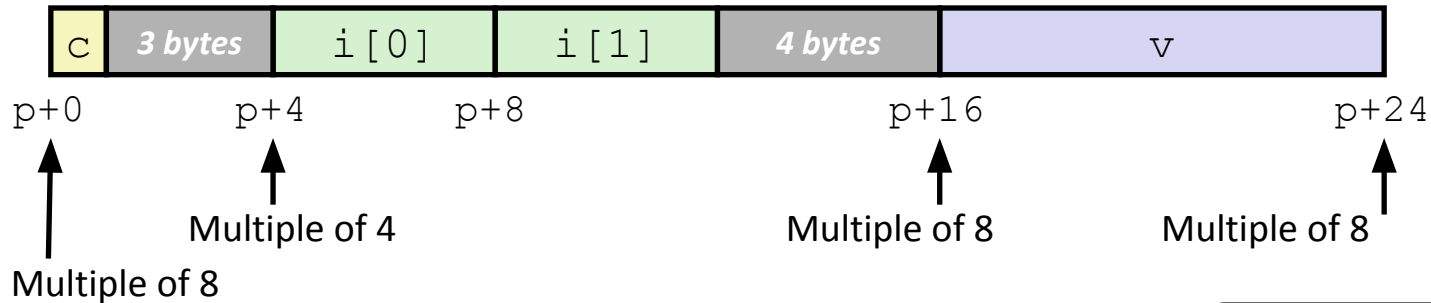
not a power of 2.  
moving data in memory,  
and to/from register,  
is a mess. instead:

```
struct S1 {  
    char c;  
    int i[2];  
    double v;  
} *p;
```

## Aligned Data

Primitive data type requires  $K$  bytes  $\Rightarrow$

Its address must be multiple of  $K$



starting address must be  
multiple of  $K$  (1, 2, 4, 8, 16, ...)

whole data structure is  
padded up to multiple of its  
largest  $K$

gcc stack is 16-byte aligned.  
remember that.

# Specific Cases of Alignment (x86-64)


- 1 byte: **char**, ...  
no restrictions on address
- 2 bytes: **short**, ...  
lowest 1 bit of address must be  $0_2$
- 4 bytes: **int**, **float**, ...  
lowest 2 bits of address must be  $00_2$
- 8 bytes: **double**, **char \***, ...  
lowest 3 bits of address must be  $000_2$

**Q:** but Willard, that's wasteful.  
**A:** when we talk dynamic memory allocation (heap), we'll see these bits used for something else (tags).

# Meeting Overall Alignment Requirement

For largest alignment requirement  $K$   
Overall structure must be multiple of  $K$

```
struct S2 {  
    double v;  
    int i[2];  
    char c;  
} *p;
```

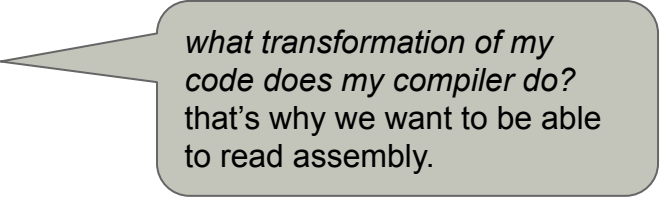


order in descending order of size makes it easier for compiler to store your struct more compactly.

Brief review (x86-64 + ex 3.67)

Arrays, strings and structs in assembly

**Structure of an .s program**



*what transformation of my code does my compiler do?  
that's why we want to be able to read assembly.*

Stack-based exploits and counter-measures

# Virtual Memory

you have different **sections** in your assembly program, which corresponds to different regions in virtual memory.

4 sections

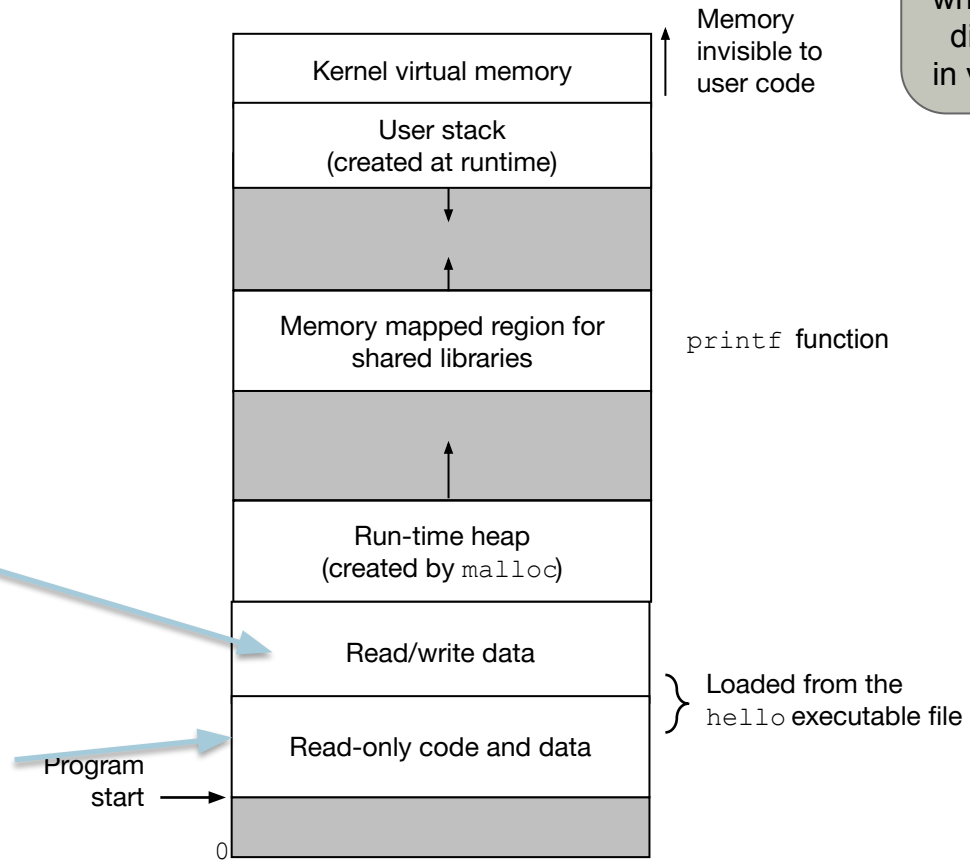
defined, but not initialized

(uninitialised) **.bss:**

(initialised read-write) **.data:**

(initialised read-only) **.rodata:**

(program) **.text:**





# Object file

-h (headers) reveals structure of object file (the sections, their size, etc.).

```
% gcc -c ex17.c
% objdump -h ex17.o

ex17.o:      file format elf64-x86-64

Sections:
Idx Name          Size      VMA           LMA           File off  Algn
  0 .text          00000718  0000000000000000  0000000000000000  00000040  2**0
             CONTENTS, ALLOC, LOAD, RELOC, READONLY, CODE
  1 .data          00000000  0000000000000000  0000000000000000  00000758  2**0
             CONTENTS, ALLOC, LOAD, DATA
  2 .bss           00000000  0000000000000000  0000000000000000  00000758  2**0
             ALLOC
  3 .rodata        00000238  0000000000000000  0000000000000000  00000758  2**3
             CONTENTS, ALLOC, LOAD, RELOC, READONLY, DATA
  4 .comment       00000035  0000000000000000  0000000000000000  00000990  2**0
             CONTENTS, READONLY
  5 .note.GNU-stack 00000000  0000000000000000  0000000000000000  000009c5  2**0
             CONTENTS, READONLY
  6 .eh_frame      00000198  0000000000000000  0000000000000000  000009c8  2**3
             CONTENTS, ALLOC, LOAD, RELOC, READONLY, DATA
```

# Assembly file

a function

a function

```
file "ex17.c"
2 .section ".rodata"
3 .LC0:
4 .string "ERROR: %s\n"
5 .text
6 .globl die
7 .type die, @function
8 die:
9 .LFB2:
10 .cfi_startproc
11 pushq %rbp
12 .cfi_def_cfa_offset 16
13 .cfi_offset 6, -16
14 movq %rsp, %rbp
15 .cfi_def_cfa_register 6
16 subq $16, %rsp
17 movq %rdi, -8(%rbp)
18 call _errno_location
19 movl (%rax), %eax
20 testl %eax, %eax
21 je .L2
22 movq -8(%rbp), %rax
23 movq %rax, %rdi
24 call perror
25 jmp .L3
26 .L2:
27 movq -8(%rbp), %rax
28 movq %rax, %rsi
29 movl $.LC0, %edi
30 movl $0, %eax
31 call printf
32 .L3:
33 movl $1, %edi
34 call exit
35 .cfi_endproc
36 .LFE2:
37 .size die, .-die
38 .section ".rodata"
39 .LC1:
40 .string "%d %s %s\n"
41 .text
42 .globl Address_print
43 .type Address_print, @function
44 Address_print:
45 .LFB3:
46 .cfi_startproc
47 pushq %rbp
48 .cfi_def_cfa_offset 16
49 .cfi_offset 6, -16
50 movq %rsp, %rbp
51 .cfi_def_cfa_register 6
52 subq $16, %rsp
53 movq %rdi, -8(%rbp)
54 movq -8(%rbp), %rax
55 leaq 520(%rax), %rcx
56 movq -8(%rbp), %rax
57 leaq 8(%rax), %rdx
58 movq -8(%rbp), %rax
59 movl (%rax), %eax
60 movl %eax, %esi
61 movl $.LC1, %edi
```

read-only data

program

read-only data

program

# Object file

disassemble all sections,  
not just those containing instructions

```
...  
% objdump -D ex17 >> ex17.d  
% vi ex17.d
```

reveals the actual  
address of each instruction

```
653  
654 000000000040098b <Address print>:  
655 40098b: 55 push %rbp  
656 40098c: 48 89 e5 mov %rsp,%rbp  
657 40098f: 48 83 ec 10 sub $0x10,%rsp  
658 400993: 48 89 7d f8 mov %rdi,-0x8(%rbp)  
659 400997: 48 8b 45 f8 mov -0x8(%rbp),%rax  
660 40099b: 48 8d 88 08 02 00 00 lea 0x208(%rax),%rcx  
661 4009a2: 48 8b 45 f8 mov -0x8(%rbp),%rax  
662 4009a6: 48 8d 50 08 lea 0x8(%rax),%rdx  
663 4009aa: 48 8b 45 f8 mov -0x8(%rbp),%rax  
664 4009ae: 8b 00 mov (%rax),%eax  
665 4009b0: 89 c6 mov %eax,%esi  
666 4009b2: bf f3 10 40 00 mov $0x4010f3,%edi  
667 4009b7: b8 00 00 00 00 mov $0x0,%eax  
668 4009bc: e8 df fd ff ff callq 4007a0 <printf@plt>  
669 4009c1: 90 nop  
670 4009c2: c9 leaveq  
671 4009c3: c3 retq  
672
```

last 2 phases of  
assignment: need  
to use this to find  
patterns in the code

Brief review (x86-64 + ex 3.67)

Arrays, strings and structs in assembly

Structure of an .s program

**Stack-based exploits and counter-measures**

# Vulnerable C code

example of vulnerable C code

```
echo.c
1 #include <stdio.h>
2
3 void echo()
4 {
5     char buf[4];
6     gets(buf);
7     puts(buf);
8
9 }
10
11 int main()
12 {
13     echo();
14
15     return 0;
16 }
```

disable stack protection  
(more on that later today)  
(so it's not really a problem  
today; compilers prevent  
problem by default)

```
% gcc -fno-stack-protector echo.c -o echonp
```

```
% ./echonp
12345678
12345678
% ./echonp
123456789
123456789
% ./echonp
12345678901234567890123
12345678901234567890123
% ./echonp
123456789012345678901234567890123
123456789012345678901234567890123
[1] 19108 segmentation fault (core dumped) ./echonp
```

process tried accessing  
something outside itself.  
what's happening?  
let's investigate.

# Vulnerable C code

tip: for assignment, always use objdump. with it, you get assembly, and the object, and the address where it is

```
% objdump -D echonp >> echonp.d
```

The stack is 16B aligned  
=>  
rsp is decreased with 16B  
sub \$0x10, %rsp

```
322  
323 0000000000400566 <echo>:  
324 400566: 55 push %rbp  
325 400567: 48 89 e5 mov %rsp,%rbp  
326 40056a: 48 83 ec 10 sub $0x10,%rsp  
327 40056e: 48 8d 45 f0 lea -0x10(%rbp),%rax  
328 400572: 48 89 c7 mov %rax,%rdi  
329 400575: b8 00 00 00 00 mov $0x0,%eax  
330 40057a: e8 d1 fe ff ff callq 400450 <gets@plt>  
331 40057f: 48 8d 45 f0 lea -0x10(%rbp),%rax  
332 400583: 48 89 c7 mov %rax,%rdi  
333 400586: e8 a5 fe ff ff callq 400430 <puts@plt>  
334 40058b: 90 nop  
335 40058c: c9 leaveq  
336 40058d: c3 retq  
337  
338 000000000040058e <main>:  
339 40058e: 55 push %rbp  
340 40058f: 48 89 e5 mov %rsp,%rbp  
341 400592: b8 00 00 00 00 mov $0x0,%eax  
342 400597: e8 ca ff ff ff callq 400566 <echo>  
343 40059c: b8 00 00 00 00 mov $0x0,%eax  
344 4005a1: 5d pop %rbp  
345 4005a2: c3 retq  
346 4005a3: 66 2e 0f 1f 84 00 00 nopw %cs:0x0(%rax,%rax,1)  
347 4005aa: 00 00 00  
348 4005ad: 0f 1f 00 nopl (%rax)  
349
```

push old base pointer on stack

new base pointer := old stack pointer

new stack pointer := old stack pointer - 16

= %rsp

address of buff into rax then into 1st arg initialize return

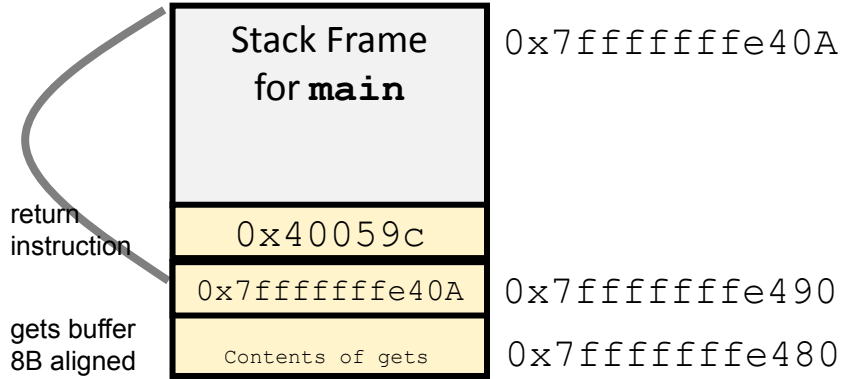
evict stack frame

enough space for buf

# Vulnerable code

so, why vulnerable?

*Before call to gets*



```
% gdb ./echonp
GNU gdb (Ubuntu 7.11.1-0ubuntu1~16.5) 7.11.1
Copyright (C) 2016 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 "x86_64-linux-gnu".
Type "show configuration" for configuration details.
For bug reporting instructions, please see:
<http://www.gnu.org/software/gdb/bugs/>.
Find the GDB manual and other documentation resources online at:
<http://www.gnu.org/software/gdb/documentation/>.
For help, type "help".
Type "apropos word" to search for commands related to "word"...
Reading symbols from ./echonp...done.
(gdb) break echo
Breakpoint 1 at 0x40056e: file echo.c, line 6.
(gdb) r
Starting program: /home/phbo/Class/C/TMP/echonp

Breakpoint 1, echo () at echo.c:6
6      gets(buf);
(gdb) p/x $rbp
$1 = 0x7fffffffef490
(gdb) p/x *((unsigned long*)$rbp)
$2 = 0x7fffffffef4a0
(gdb) p/x *((unsigned long*)$rbp+1)
$3 = 0x40059c
```

`%rbp+1` is return instruction

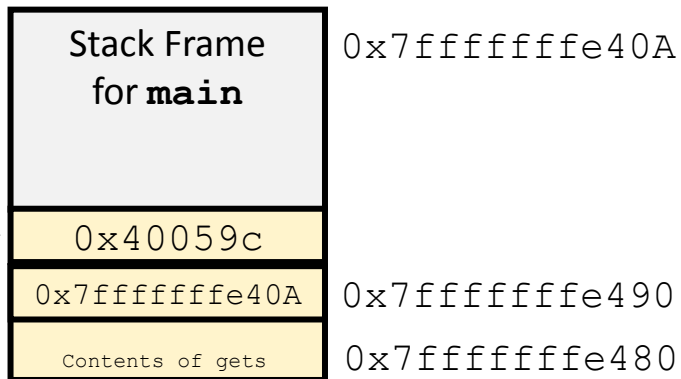
```
$5 = 0x40059c
(gdb) p/x $rsp
$4 = 0x7fffffffef480
```

if I write more than 16 bytes,  
then I overwrite return address.  
I take over the machine then  
(or provoke segmentation fault)

```
342 400597: e8 ca ff ff ff      callq 400566 <echo>
343 40059c: b8 00 00 00 00     mov $0x0,%eax
```

# Buffer Overflow Stack Example

*Before call to gets*

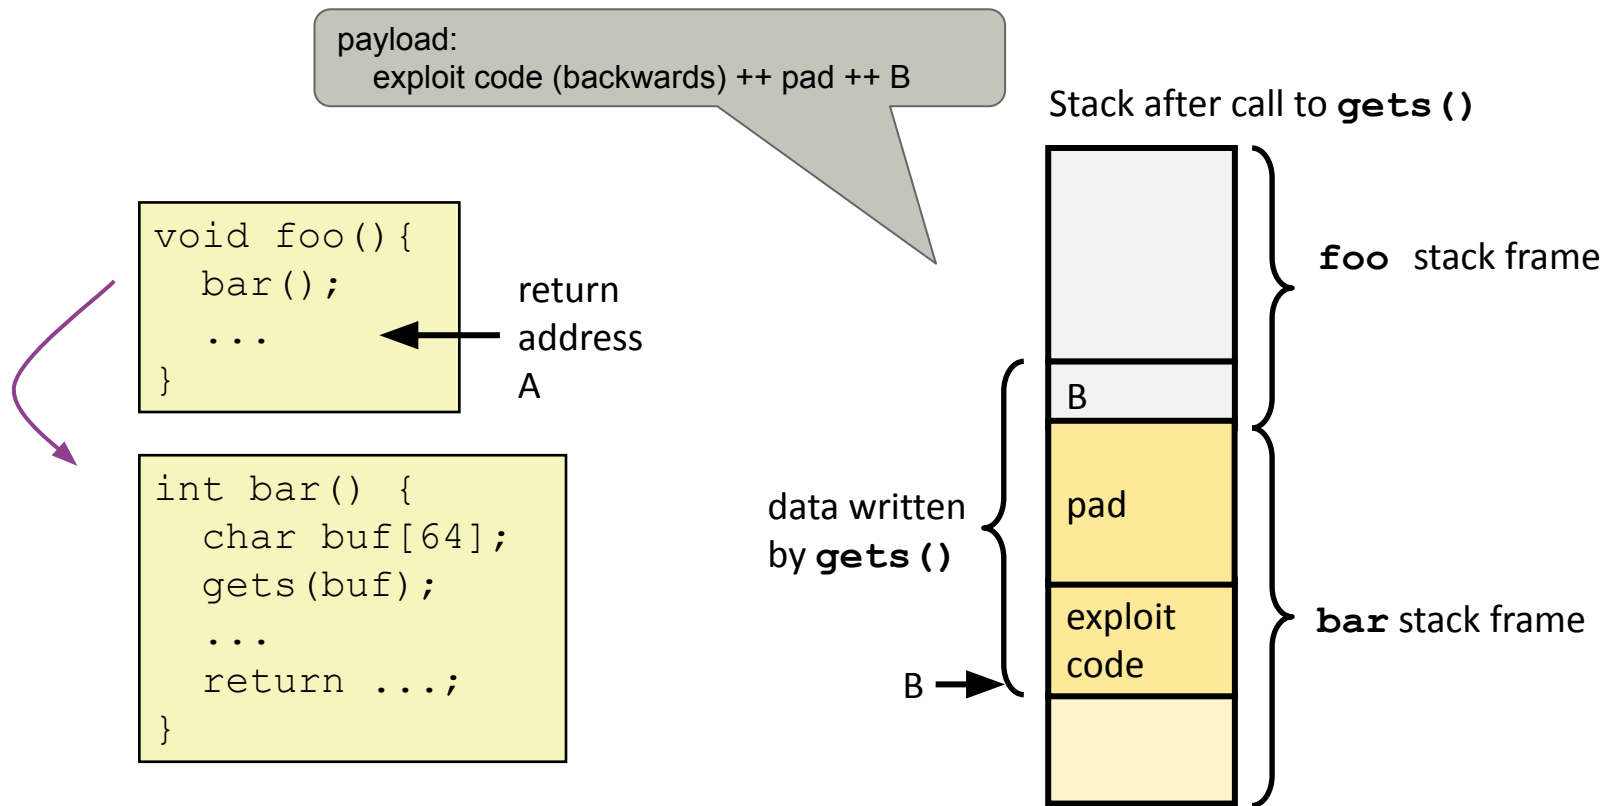


depending on input,  
different termination behavior.

```
quit anyway. (y or n) y
% ./echonp
1234567890123456789012345
1234567890123456789012345
[1] 23189 segmentation fault (core dumped) ./echonp
% ./echonp
123456789012345678901234
123456789012345678901234
[1] 23255 illegal hardware instruction (core dumped) ./echonp
% ./echonp
12345678901234567890123
12345678901234567890123
```



# Malicious Use of Buffer Overflow



Input string contains **byte representation of executable code**  
Overwrite return address with address of exploit code  
When `bar()` executes `ret`, will jump to exploit code

# 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**

Morris worm, 1988

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.

how to stop this:

- want to make it impossible to execute code on the stack.
- want to make it hard to figure out where exploit code starts.
- want to protect return address.

# Avoiding Overflow Vulnerability

use a function that checks how much you read.

```
/* 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
- Or use **%ns** where **n** is a suitable integer

# System-Level Protections

## Randomized stack offsets

At start of program, allocate random amount of space on stack  
Makes it difficult for hacker to predict beginning of inserted code

## Nonexecutable code segments

In traditional x86, can mark region of memory as either  
“read-only” or “writeable”

- Can execute anything readable

X86-64 added explicit “execute” permission

# Stack Canaries

## Idea

Place special value (“canary”) on stack just beyond buffer

Check for corruption before exiting function

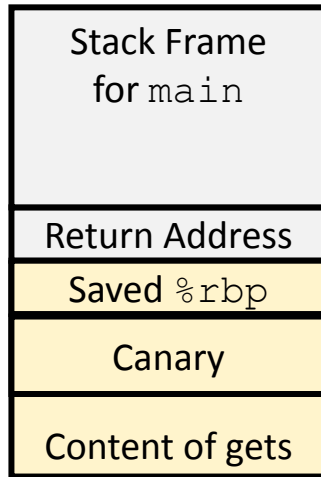
## GCC Implementation

**-fstack-protector**

**-fstack-protector-all (default)**

# Setting Up Canary

*Before call to gets*



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

```
echo:  
    . . .  
    mov %fs:20, %rax    # Get canary  
    mov %rax, -8(%rbp) # Put on stack  
    xor %eax, %eax     # Erase canary  
    . . .
```

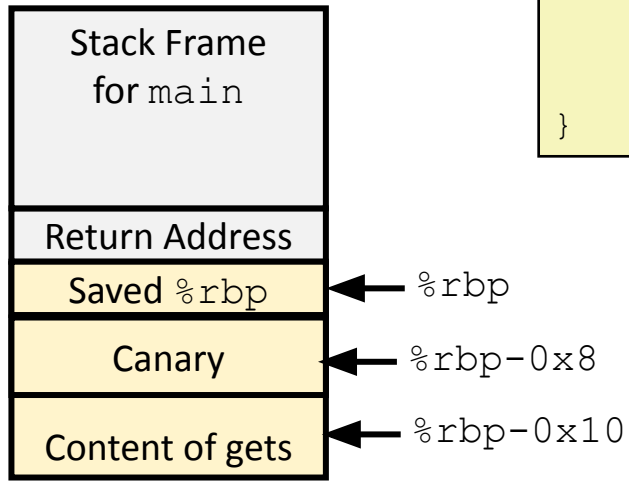
random part of memory into  
rax

put canary just below base  
pointer

must overwrite canary to  
overwrite return address

# Checking Canary

*Before call to gets*



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

```
echo:
    . . .
    mov -8(%rbp), %rax    # Retrieve from
stack
    xor %fs:20, %rax     # Compare with
Canary
    je .L24              # Same: skip ahead
    call __stack_chk_fail # ERROR
.L24:
    . . .
```

# Canary Example

```
361 00000000004005d6 <echo>:
362 4005d6: 55          push  %rbp
363 4005d7: 48 89 e5    mov   %rsp,%rbp
364 4005da: 48 83 ec 10 sub   $0x10,%rsp
365 4005de: 64 48 8b 04 25 28 00 mov   %fs:0x28,%rax
366 4005e5: 00 00
367 4005e7: 48 89 45 f8 mov   %rax,-0x8(%rbp)
368 4005eb: 31 c0       xor   %eax,%eax
369 4005ed: 48 8d 45 f0 lea   -0x10(%rbp),%rax
370 4005f1: 48 89 c7    mov   %rax,%rdi
371 4005f4: b8 00 00 00 00 mov   $0x0,%eax
372 4005f9: e8 c2 fe ff ff callq 4004c0 <gets@plt>
373 4005fe: 48 8d 45 f0 lea   -0x10(%rbp),%rax
374 400602: 48 89 c7    mov   %rax,%rdi
375 400605: e8 86 fe ff ff callq 400490 <puts@plt>
376 40060a: 90         nop
377 40060b: 48 8b 45 f8 mov   -0x8(%rbp),%rax
378 40060f: 64 48 33 04 25 28 00 xor   %fs:0x28,%rax
379 400616: 00 00
380 400618: 74 05      je    40061f <echo+0x49>
381 40061a: e8 81 fe ff ff callq 4004a0 <__stack_chk_fail@plt>
382 40061f: c9        leaveq
383 400620: c3        retq
384
385
386 0000000000400621 <main>:
387 400621: 55          push  %rbp
388 400622: 48 89 e5    mov   %rsp,%rbp
389 400625: b8 00 00 00 00 mov   $0x0,%eax
390 40062a: e8 a7 ff ff ff callq 4005d6 <echo>
391 40062f: b8 00 00 00 00 mov   $0x0,%eax
392 400634: 5d         pop   %rbp
393 400635: c3        retq
394 400636: 66 2e 0f 1f 84 00 00 nopw  %cs:0x0(%rax,%rax,1)
395 40063d: 00 00 00
396
```



# Canary Example

```
12345678901234567890123
% ./echo
12345678
12345678
% ./echo
123456789
123456789
*** stack smashing detected ***: ./echo terminated
[1] 23643 abort (core dumped) ./echo
```

when compiled w/  
protection

# GCC protections

## Compile-time

```
cc echo.c -o echo
echo.c: In function 'echo':
echo.c:6:2: warning: implicit declaration of function 'gets' [-Wimplicit-function-declaration]
  gets(buf);
  ^
/tmp/cchUozYT.o: In function `echo':
echo.c:(.text+0x24): warning: the `gets' function is dangerous and should not be used.
```

## Run-time

```
[1] 17102 abort (core dumped) ./echo
% ./echo
123456789
123456789
*** stack smashing detected ***: ./echo terminated
[1] 17139 abort (core dumped) ./echo
```

pretty clear that you shouldn't use **gets**.

shameless self-promotion

**goal:** tools that developers can use to write secure SW.

sample research (past supervisions):

- **analyze binaries** for information leaks
- reduce **timing leaks** in the **Linux kernel**
- automatically **fix vulnerabilities** in JavaScript
- automatically generate (i.e. **synthesize**) a secure program **from formal specification**
- assess **privacy risk** in **analytics programs** (data scientists; Google search for “Privugger”)

I like code, and I like proofs.


I created the “Applied Information Security” course.

I’m a barista in Analog.

Willard Rafnsson  
IT University of Copenhagen

wilr@itu.dk

<https://www.willardthor.com/>



# Take-Aways

Alignment is necessary when working with structs

Object files are sequences of hex codes, that can be mapped to assembly instructions

Buffer overflows cause vulnerabilities that can be exploited maliciously

Counter measures include (i) randomized stack addresses, (ii) non executable code on stack and (iii) canaries.