

Introduction to Binary Exploitation

THS Hacking Workshops

Andrea Continella

06-12-2021

**UNIVERSITY
OF TWENTE.**

Assumptions

The following *concepts* apply, with proper modifications, to any machine architecture (e.g., ARM, x86), operating system (e.g., Windows, Linux, Darwin), and executable (e.g., Portable Executable (PE), Executable and Linkable Format (ELF))

For simplicity, we assume **ELFs** running on **Linux >= 2.6** processes on top of a **32-bit x86** machine.

Recap on calling convention

```
int foo(int a, int b) {
    int c = 14;
    return (a + b) * c;
}

int main(int argc, char * argv[]) {
    int avar, bvar, cvar;
    char * str;

    avar = atoi(argv[1]);
    bvar = atoi(argv[2]);
    cvar = foo(avar, bvar);

    gets(str);

    return 0;
}
```

```
./executable_file 10 20 <--
```

The foo() function receives two parameters by copy

- How does the CPU pass them to the function?
- Push them onto the stack!

80484d5:	83 c0 08
80484d8:	8b 00
80484da:	83 ec 0c
80484dd:	50
80484de:	e8 dd fe ff ff
80484e3:	83 c4 10
80484e6:	89 45 ec
80484e9:	83 ec 08
80484ec:	ff 75 ec
80484ef:	ff 75 e8
80484f2:	e8 8d ff ff ff
80484f7:	83 c4 10
80484fa:	89 45 f0
80484fd:	83 ec 0c
8048500:	ff 75 f4
8048503:	e8 78 fe ff ff
8048508:	83 c4 10
804850b:	83 ec 0c
804850e:	ff 75 f4
8048511:	e8 7a fe ff ff
8048516:	83 c4 10
8048519:	b8 10 86 04 08
804851e:	ff 75 f0

Push the second parameter, which happens to be on the stack, left there by previous instructions, at EBP-0x14.

```


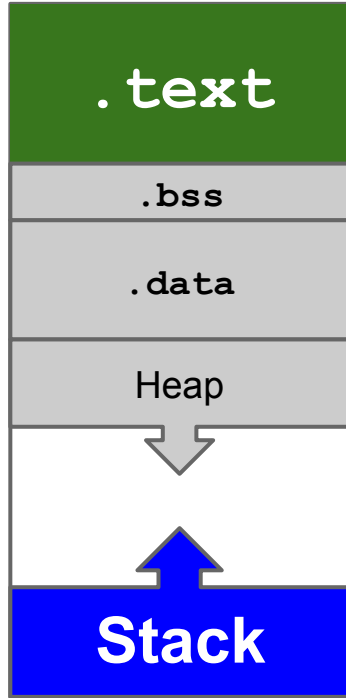
sub    $0xc,%esp
push  %eax
call   80483c0 <atoi@plt>
add    $0x10,%esp
mov    %eax,-0x14(%ebp)
sub    $0x8,%esp
pushl  -0x14(%ebp)
pushl  -0x18(%ebp)
call   8048484 <foo>
add    $0x10,%esp
mov    %eax,-0x10(%ebp)
sub    $0xc,%esp
pushl  -0xc(%ebp)
call   8048380 <gets@plt>
add    $0x10,%esp
sub    $0xc,%esp
pushl  -0xc(%ebp)
call   8048390 <puts@plt>
add    $0x10,%esp
mov    $0x8048610,%eax
pushl  -0x10(%ebp)

```

EIP (Instruction Pointer)



0xBFFDF000



<- ESP: stack pointer
(points to the top of the stack)

<- `EBP-0x18` (24 bytes below EBP)

<- `EBP-0x14` (20 bytes below EBP)

<- `EBP`

0xC0000000

80484d5:	83 c0 08
80484d8:	8b 00
80484da:	83 ec 0c
80484dd:	50
80484de:	e8 dd fe ff ff
80484e3:	83 c4 10
80484e6:	89 45 ec
80484e9:	83 ec 08
80484ec:	ff 75 ec
80484ef:	ff 75 e8
80484f2:	e8 8d ff ff ff
80484f7:	83 c4 10
80484fa:	89 45 f0
80484fd:	83 ec 0c
8048500:	ff 75 f4
8048503:	e8 78 fe ff ff
8048508:	83 c4 10
804850b:	83 ec 0c
804850e:	ff 75 f4
8048511:	e8 7a fe ff ff
8048516:	83 c4 10
8048519:	b8 10 86 04 08
804851e:	ff 75 f0

```



add    $0x8,%eax
mov    (%eax),%eax
sub    $0xc,%esp
push  %eax
call   80483c0 <atoi@plt>
add    $0x10,%esp
mov    %eax,-0x14(%ebp)
sub    $0x8,%esp
pushl  -0x14(%ebp)
pushl  -0x18(%ebp)
call   8048484 <foo>
add    $0x10,%esp
mov    %eax,-0x10(%ebp)
sub    $0xc,%esp
pushl  -0xc(%ebp)
call   8048380 <gets@plt>
add    $0x10,%esp
sub    $0xc,%esp
pushl  -0xc(%ebp)
call   8048390 <puts@plt>
add    $0x10,%esp
mov    $0x8048610,%eax
pushl  -0x10(%ebp)

```

Push the first parameter



0xBFFDF000



`<- ESP: stack pointer`
`(points to the top of the stack)`

`<- EBP-0x18` (24 bytes below EBP)

`<- EBP-0x14` (20 bytes below EBP)

`<- EBP` `0xC0000000`

80484d5:	83 c0 08
80484d8:	8b 00
80484da:	83 ec 0c
80484dd:	50
80484de:	e8 dd fe ff ff
80484e3:	83 c4 10
80484e6:	89 45 ec
80484e9:	83 ec 08
80484ec:	ff 75 ec
80484ef:	ff 75 e8
80484f2:	e8 8d ff ff ff
80484f7:	83 c4 10
80484fa:	89 45 f0
80484fd:	83 ec 0c
8048500:	ff 75 f4
8048503:	e8 78 fe ff ff
8048508:	83 c4 10
804850b:	83 ec 0c
804850e:	ff 75 f4
8048511:	e8 7a fe ff ff
8048516:	83 c4 10
8048519:	b8 10 86 04 08
804851e:	ff 75 f0

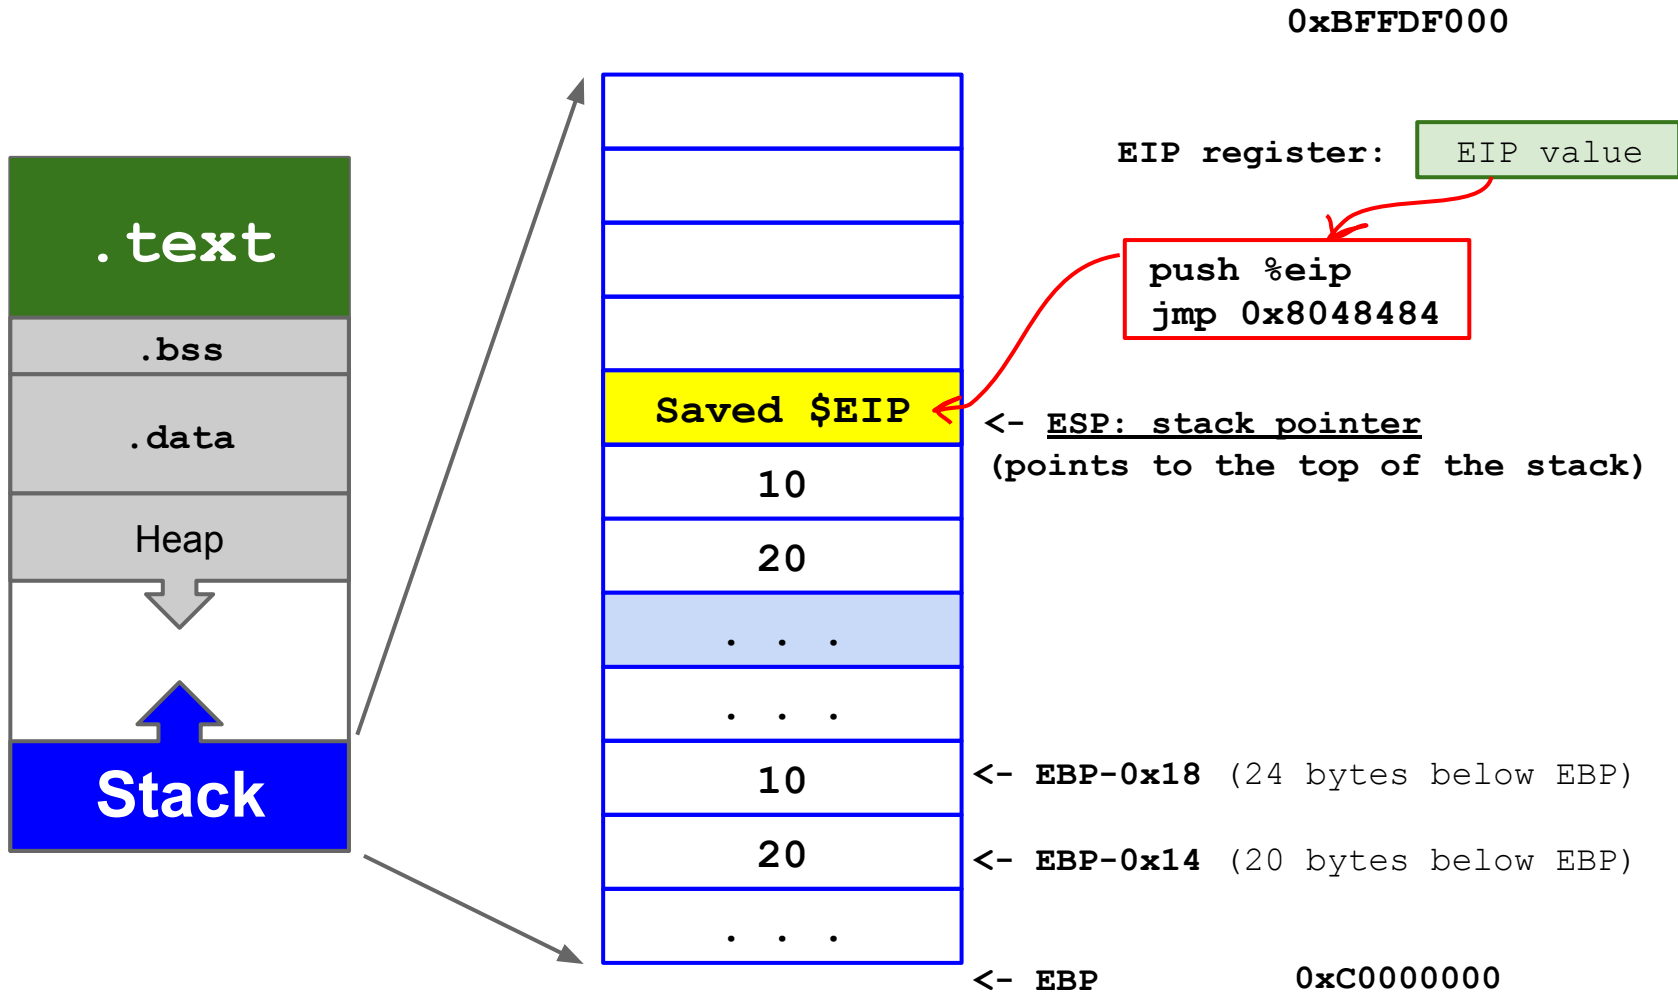
```
add    $0x8,%eax
mov    (%eax),%eax
sub    $0xc,%esp
push   %eax
call   80483c0 <atoi@plt>
add    $0x10,%esp
mov    %eax,-0x14(%ebp)
sub    $0x8,%esp
pushl  -0x14(%ebp)
pushl  -0x18(%ebp)
call   8048484 <foo>
add    $0x10,%esp
mov    %eax,-0x10(%ebp)
sub    $0xc,%esp
pushl  -0xc(%ebp)
call   8048380 <gets@plt>
add    $0x10,%esp
sub    $0xc,%esp
pushl  -0xc(%ebp)
call   8048390 <puts@plt>
add    $0x10,%esp
mov    $0x8048610,%eax
pushl  -0x10(%ebp)
```



The `call` instruction

- The CPU is about to call the `foo()` function
- When `foo()` will be over, where to jump?
- The CPU needs to **save the current EIP**
- **Where** does the CPU save the EIP?
 - On the **stack!**

```
call 0x8048484 <foo> == { push %eip  
                        jmp 0x8048484 ~> foo()
```



8048484:

8048485:

8048487:

804848a:

8048491:

8048494:

8048497:

8048499:

804849c:

804849f:

80484a2:

80484a5:

80484a6:

...

...

80484ec:

80484ef:

80484f2:

80484f7:

80484fa:

55

89 e5

83 ec 10

c7 45 fc 0e 00 00 00

8b 45 0c

8b 55 08

01 c2

8b 45 fc

0f af c2

89 45 fc

8b 45 fc

c9

c3

. . .

. . .

ff 75 ec

ff 75 e8

e8 8d ff ff ff

83 c4 10

89 45 f0

Function prologue

push %ebp

mov %esp,%ebp

sub \$0x4,%esp

movl \$0xe,-0x4(%ebp)

mov 0xc(%ebp),%eax

mov 0x8(%ebp),%edx

add %eax,%edx

mov -0x4(%ebp),%eax

imul %edx,%eax

mov %eax,-0x4(%ebp)

mov -0x4(%ebp),%eax

leave

ret

pushl -0x14(%ebp)

pushl -0x18(%ebp)

call 8048484 <foo>


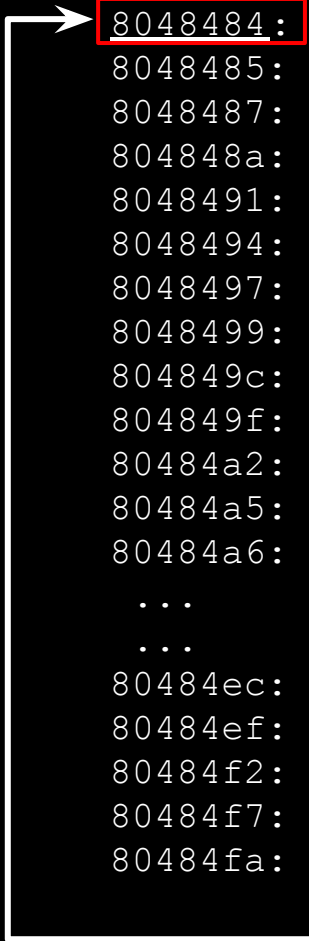
add \$0x10,%esp

mov %eax,-0x10(%ebp)

EIP (Instruction Pointer)

push %eip

jmp 0x8048484



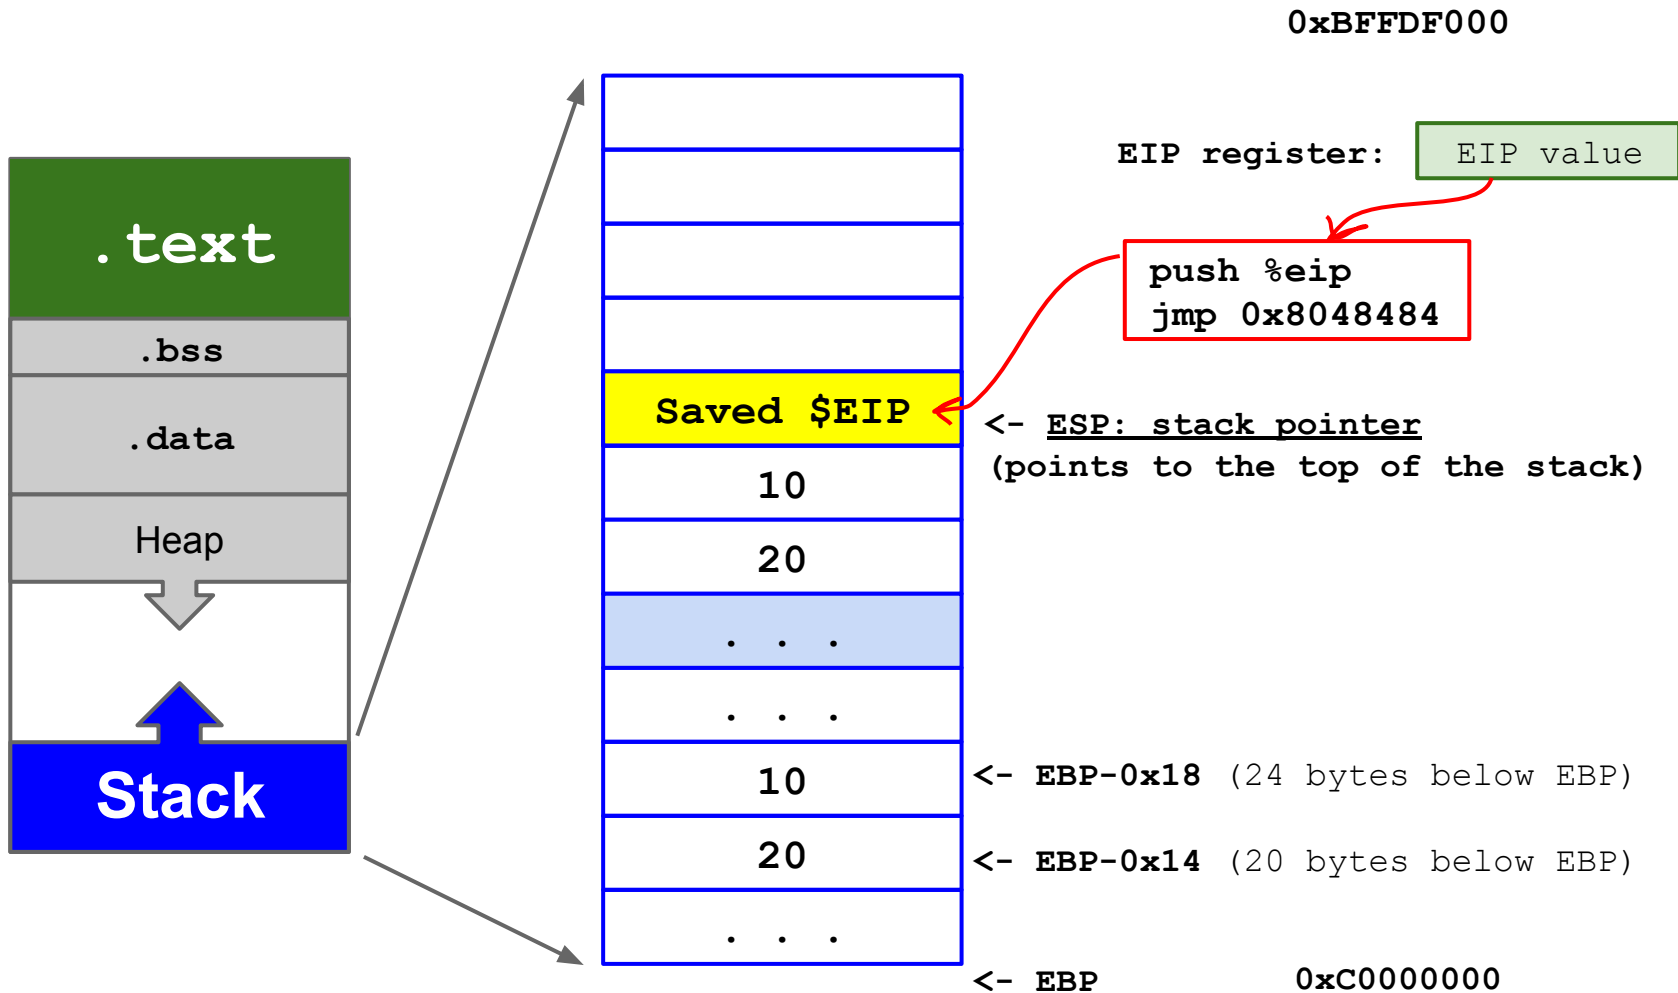
Function Prologue


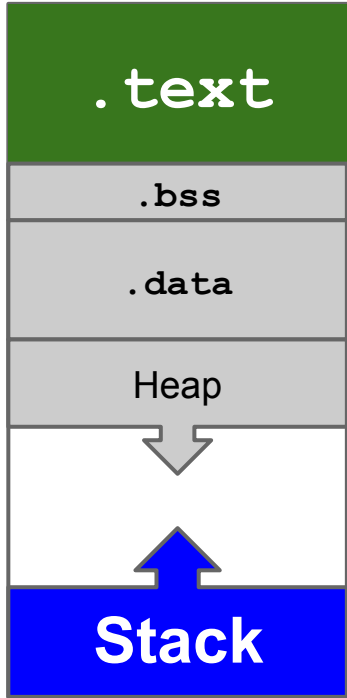
The CPU needs to remember where `main()`'s *frame* is located on the stack, so that it can be restored once `foo()`'s will be over.

The first 3 instructions of `foo()` take care of this.

<code>push %ebp</code>	save the current stack base address onto the stack
<code>mov %esp, %ebp</code>	the new base of the stack is the old top of the stack
<code>sub \$0x4, %esp</code>	allocate <code>0x4</code> bytes (32 bits integer) for <code>foo()</code> 's <u>local variables</u>

```
int foo(int a, int b) {  
    int c = 14;  
    c = (a + b) * c;  
    return c;  
}
```





0xBFFDF000

`<- ESP` ← 0x4 bytes subtraction

`<- EBP: base pointer address ESP`

```

Function prologue
push    %ebp
mov     %esp, %ebp
sub     $0x4, %esp

```

0xC0000000

8048484:	55
8048485:	89 e5
8048487:	83 ec 10
804848a:	c7 45 fc 0e 00 00 00
8048491:	8b 45 0c
8048494:	8b 55 08
8048497:	01 c2
8048499:	8b 45 fc
804849c:	0f af c2
804849f:	89 45 fc
80484a2:	8b 45 fc
80484a5:	c9
80484a6:	c3

...	...
...	...
80484ec:	ff 75 ec
80484ef:	ff 75 e8
80484f2:	e8 8d ff ff ff
80484f7:	83 c4 10
80484fa:	89 45 f0

```

push    %ebp
mov     %esp,%ebp
sub     $0x4,%esp
movl   $0xe,-0x4(%ebp)
mov     0xc(%ebp),%eax
mov     0x8(%ebp),%edx
add     %eax,%edx
mov     -0x4(%ebp),%eax
imul   %edx,%eax
mov     %eax,-0x4(%ebp)
mov     %eax,-0x4(%ebp)
leave  ),%eax
ret

```



Function Epilogue

The CPU needs to return back to **main()** 's execution flow.
The last 2 instructions of **foo()** take care of this.

these 2 instructions translate into these 3 instructions →

```
leave
ret
```

current base is the **new top** of the stack
restore the **saved EBP** to registry
pop the saved EIP and jump there

```
mov %ebp, %esp
pop %ebp
ret
```


Low addresses (0xBFFDF000)

EBP - "N*4" in hex

EBP-0x8

EBP-0x4

MEMORY ALLOCATION



MEMORY WRITING

```
{  
    ...  
    gets(var2);  
}
```

EBP

EBP+0x4

EBP+0x8

EBP+0x12

EBP + "(N+1)*4" in hex

```
foo(arg1, arg2,  
    ..., argN) {  
    var1;  
    var2;  
    ...  
    varN;  
}
```

High addresses (0xC0000000)

Now you know how the stack is used...

A young girl with brown hair is in the foreground, looking towards a house on fire in the background. The house is engulfed in bright orange and yellow flames. Several people, including a firefighter in a dark uniform, are visible near the burning house. A yellow fire hose lies on the ground in the middle ground. The scene is set outdoors at dusk or dawn, with a dark sky.

WHAT IF WE OVERWRITE

OTHER DATA

Stack Smashing

1994 idea (well explained by aleph1)

- ["Smashing the stack for fun and profit"](#) (must read!)
- `foo()` allocates a buffer for a local variable, e.g., `char buf[8]`
- `buf` is filled **without size checking**
- Can easily happen in C:
 - `strcpy, strcat`
 - `fgets, gets`
 - `sprintf`
 - `scanf`

chall-0



A young girl with brown hair is in the foreground, looking towards a house that is on fire in the background. The house is engulfed in bright orange and yellow flames. Several people, including a firefighter in a dark uniform, are visible near the burning house. A yellow fire hose is laid out on the ground in the middle ground. The scene is set outdoors at night or dusk.

WHAT IF WE CHANGE

THE SAVED EIP

chall-1




Where to Jump?

Problem: We need to jump to a **valid memory location** that contains, or can be filled with, **valid executable machine code**

Solutions (i.e., exploitation techniques):

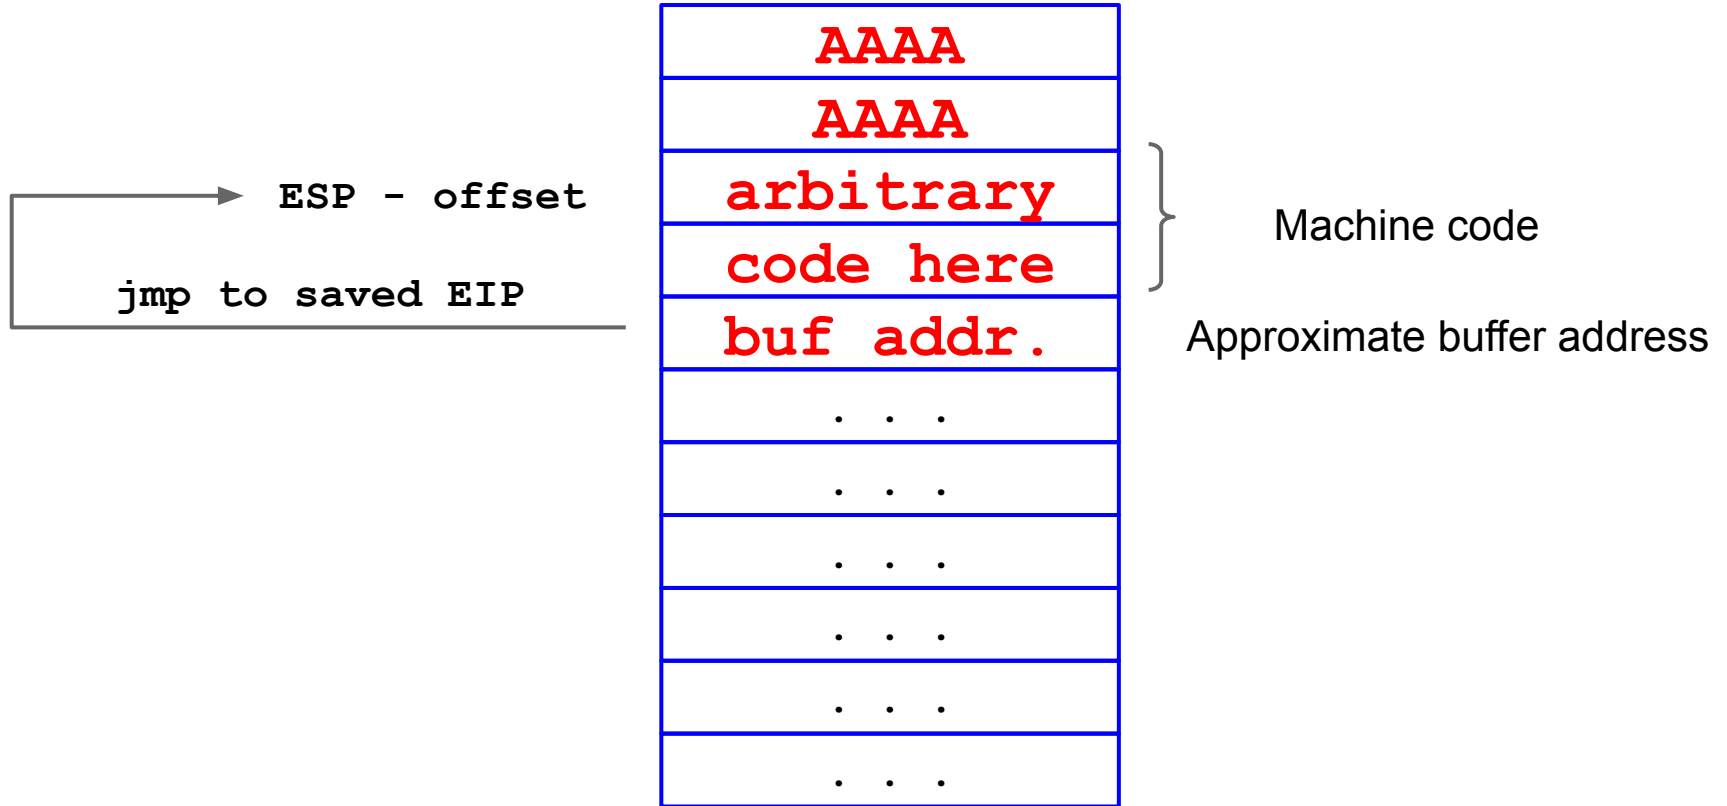
- Environment variable
- Built-in, existing functions
- Memory that we can control
 - The buffer itself <~ we will go with this
 - Some other variable

Smashing the Stack

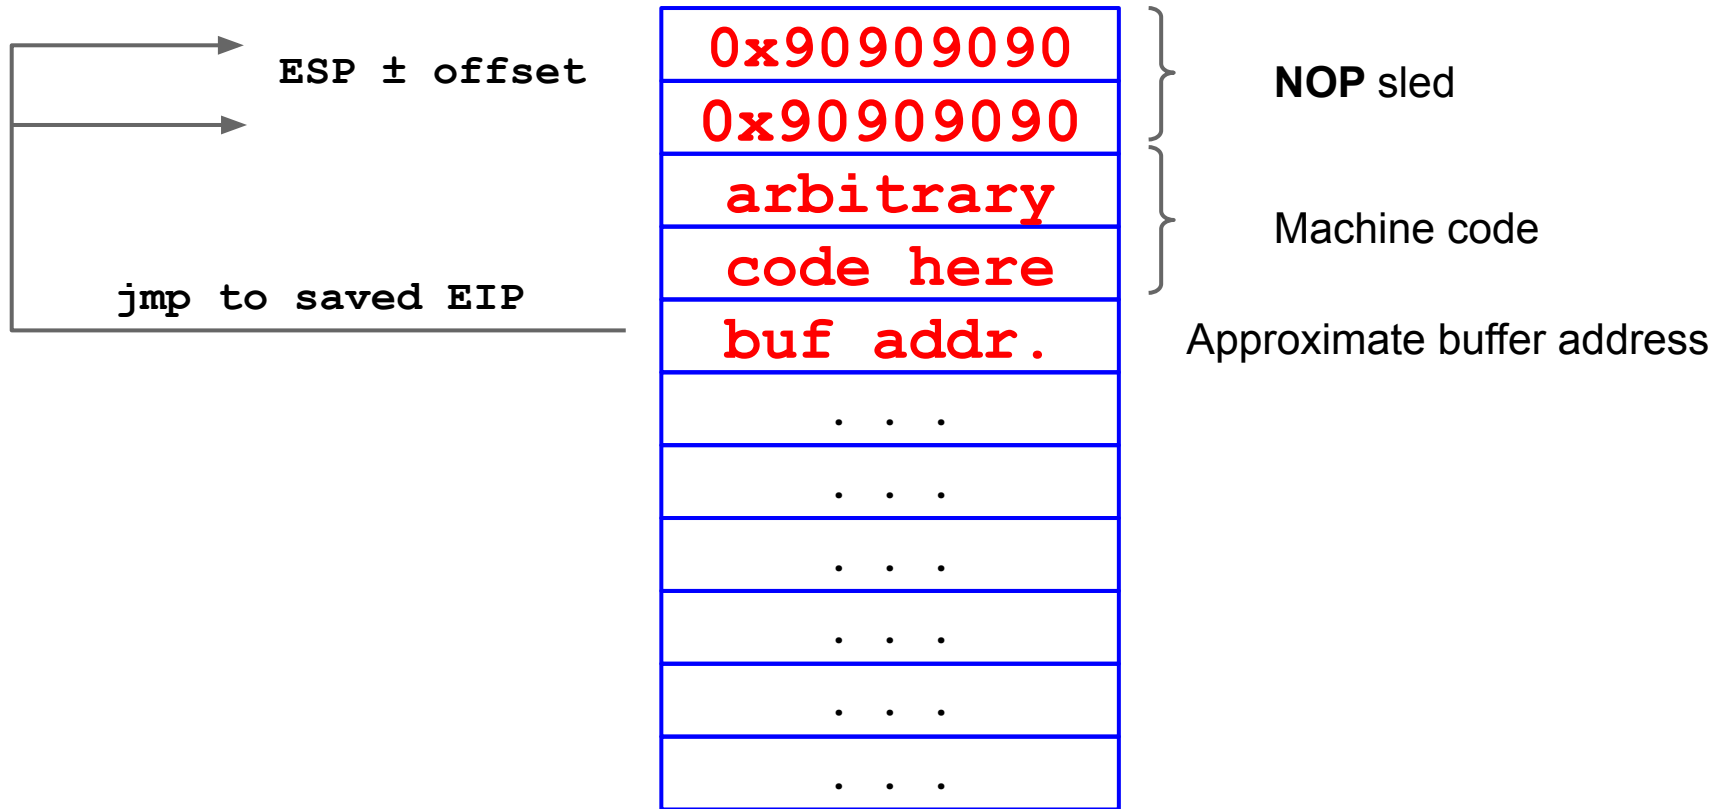


Where to jump?

Smashing the Stack




NOP (0x90) Sled to the Resc



What to Execute? 5h311c0d3

Historically, goal of the attacker: to spawn a (privileged) **shell** (on a local/remote machine)



(Shell)code: sequence of machine instructions (that are needed to open a shell)

In general, a shellcode may do just anything (e.g., open a TCP connection, launch a VPN server, a reverse shell).

<http://shell-storm.org/shellcode/>

Basically: execute `execve("/bin/sh")`

Writing a Shellcode in a Nutshell

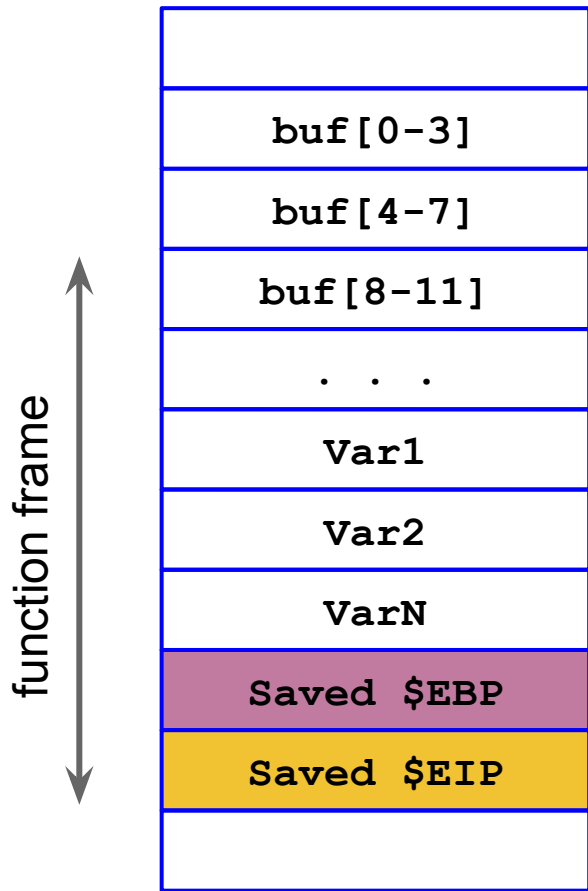
General idea: Translate `execve("/bin/sh", NULL, NULL)` into machine instructions

Invoke system call by executing a software interrupt through the `int` instruction

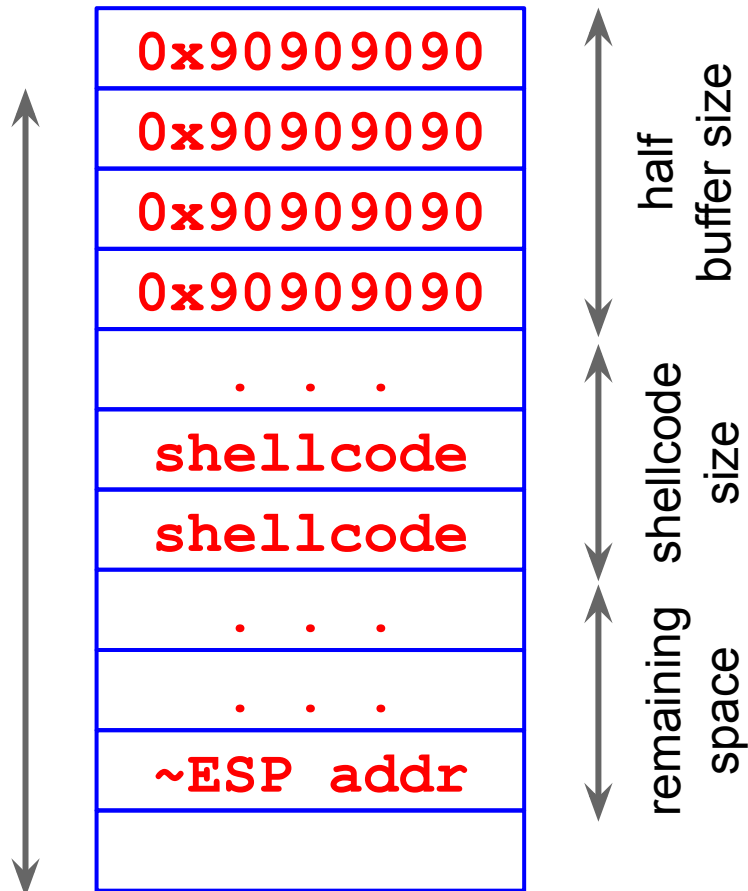
Trick: Addresses should be relative

Problem: Avoid zeros (replace semantically equivalent instructions)

```
shellcode = "\x31\xc0\x50\x89\xe2\x68\x2f\x2f\x73\x68\x68  
\x2f\x62\x69\x6e\x89\xe3\x50\x53\x89\xe1\xb0\x0b\xcd\x80"
```



Vulnerable program's memory layout (function frame).



Memory layout of a possible exploit.

chall-2

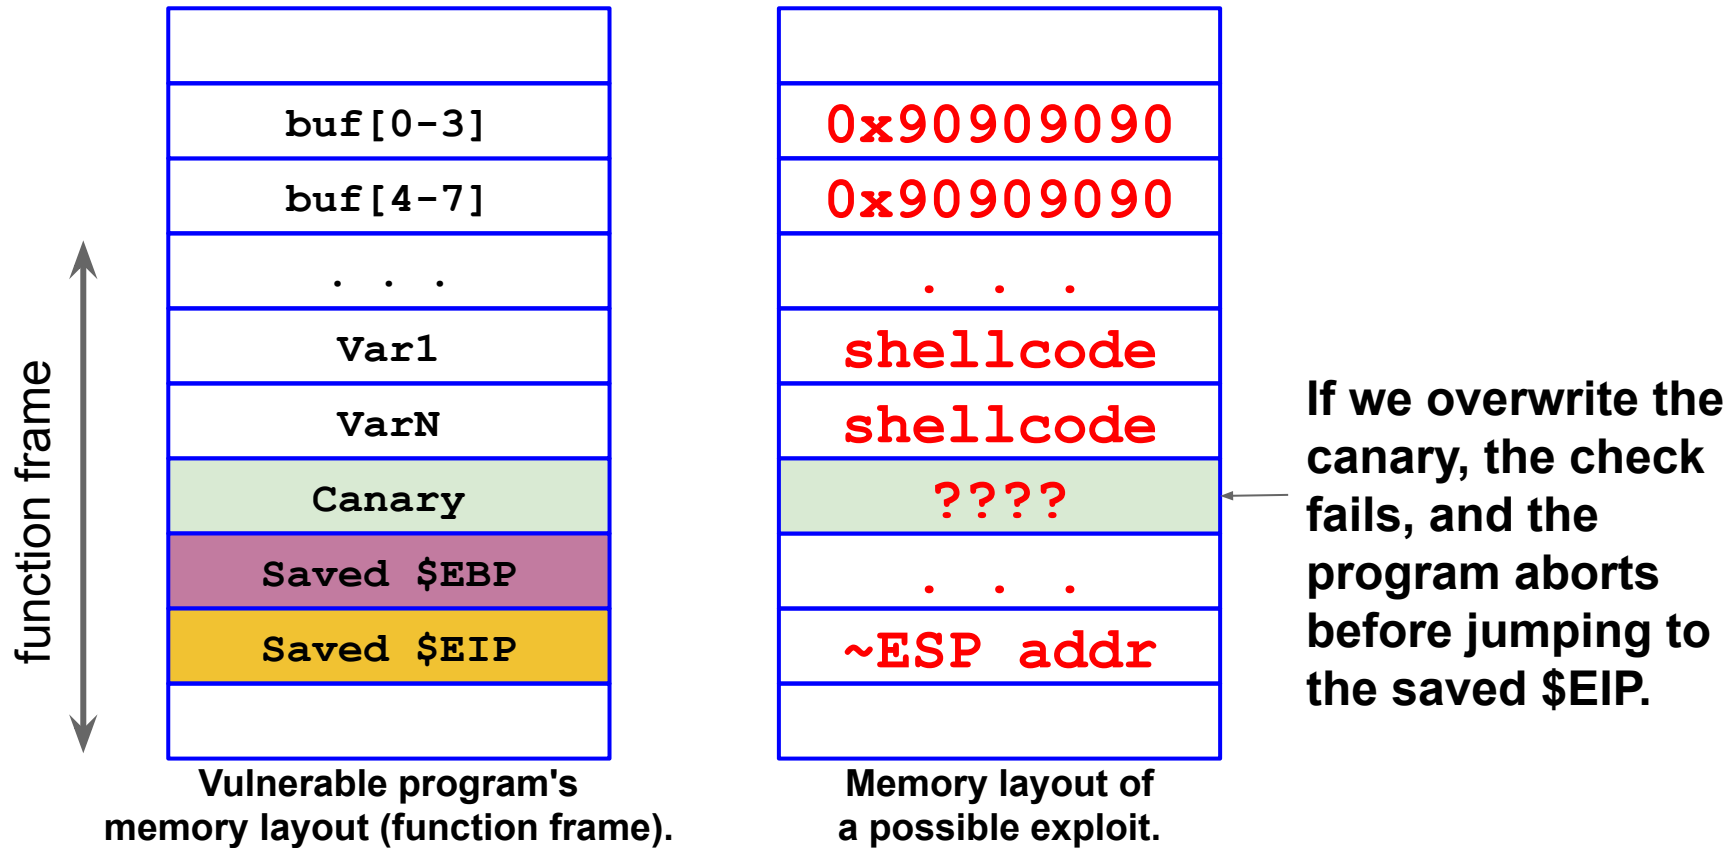
```
echo 0 > /proc/sys/kernel/randomize_va_space
```



Protection Mechanisms

- **Stack Canaries**
- **ASLR**: Address Space Layout Randomization
- **NX** or **DEP**: No section is executable & writable
- **RELRO**: Read-Only after Relocation
- **PIE**: Position Independent Executable
- **CFI**: Control-Flow Integrity

Stack Canaries



Non-executable Stack

Non-executable stack (data != code)

- No stack smashing or local variables
 - Issue: some programs (e.g., JVM older versions) actually need to execute code on the stack.
- The hardware **NX bit** mechanism is used
 - Implementations: **DEP**, since Windows XP SP2; OpenBSD **W^X**; **ExecShield** in Linux
- **Bypass**: don't inject code, but point the return address to existing machine instructions (code-reuse attacks)
 - C library functions: “return to libc” (ret2libc)
 - Generalization: return oriented programming (ROP)

Built-in, Existing Functions

The address of a system library or function (e.g., return to libc attack).

PROS:

- Works remotely and reliably

- No need for executable stack

- A function is executable usually :-)

CONS:

- Need to prepare the stack frame carefully

Address Space Layout Randomization

Repositioning the stack, among other things, at each execution at random; impossible to guess return addresses correctly

Active by default in Linux > 2.6.12, randomization range 8MB

`/proc/sys/kernel/randomize_va_space`

chall-3

```
echo 2 > /proc/sys/kernel/randomize_va_space
```





Bonus: chall-4



Further Reading

[textbook] Chris Anley et al., "*The Shellcoder's Handbook. Discovering and Exploiting Security Holes*", 2007
<https://www.wiley.com/en-it/The+Shellcoder's+Handbook:+Discovering+and+Exploiting+Security+Holes,+2nd+Edition-p-9780470080238>

Aleph One, "Smashing The Stack For Fun And Profit", Phrack, 1996, <http://phrack.org/issues/49/14.html#article>
V. Van der Veen et al., "*Memory Errors: The Past, the Present and the Future*". RAID 2012
https://dx.doi.org/10.1007/978-3-642-33338-5_5 (short history about mitigations against memory corruption exploitation)

C. Cowan et al., "*StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks*", USENIX Security 1998 https://www.usenix.org/legacy/publications/library/proceedings/sec98/full_papers/cowan/cowan.pdf
(introduces stack canaries)

H. Shacham. "*The Geometry of Innocent Flesh on the Bone: Return-into-libc without Function Calls (on the x86)*". CCS 2007 <https://acmccs.github.io/papers/geometry-ccs07.pdf> (introduces the concept of return oriented programming)

[course] Modern Binary Exploitation - CSCI 4968 <http://security.cs.rpi.edu/courses/binexp-spring2015/>
/ <https://github.com/RPISEC/MBE>

Thanks!

Andrea Continella

a.continella@utwente.nl

<https://conand.me>



@_conand