



# pointer subterfuge

```
void f2b(void *arg, size_t len) {  
    char buffer[100];  
    long val = ...; /* assume on stack */  
    long *ptr = ...; /* assume on stack */  
    memcpy(buff, arg, len); /* overwrite ptr? */  
    *ptr = val; /* arbitrary memory write! */  
}
```

# pointer subterfuge

```
void f2b(void *arg, size_t len) {  
    char buffer[100];  
    long val = ...; /* assume on stack */  
    long *ptr = ...; /* assume on stack */  
    memcpy(buff, arg, len); /* overwrite ptr? */  
    *ptr = val; /* arbitrary memory write! */  
}
```

# arbitrary memory write

bunch of scenarios that lead to *single arbitrary memory write*

typical result: arbitrary code execution

how?

# arbitrary memory write

bunch of scenarios that lead to *single arbitrary memory write*

typical result: arbitrary code execution

how?

overwrite existing machine code (insert jump?)

problem: usually not writable

overwrite return address directly

observation: don't care about stack canaries — skip them

overwrite other function pointer?

overwrite another data pointer — copy more?

# arbitrary memory write

bunch of scenarios that lead to *single arbitrary memory write*

typical result: arbitrary code execution

how?

*overwrite existing machine code (insert jump?)*

problem: usually not writable

overwrite return address directly

observation: don't care about stack canaries — skip them

overwrite other function pointer?

overwrite another data pointer — copy more?

# arbitrary memory write

bunch of scenarios that lead to *single arbitrary memory write*

typical result: arbitrary code execution

how?

overwrite existing machine code (insert jump?)

problem: usually not writable

*overwrite return address directly*

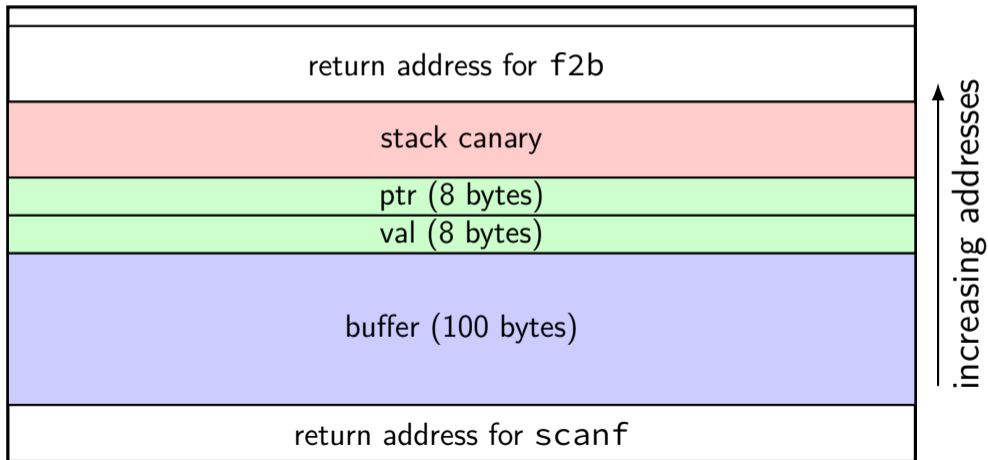
observation: don't care about stack canaries — skip them

overwrite other function pointer?

overwrite another data pointer — copy more?

# skipping the canary

highest address (stack started here)

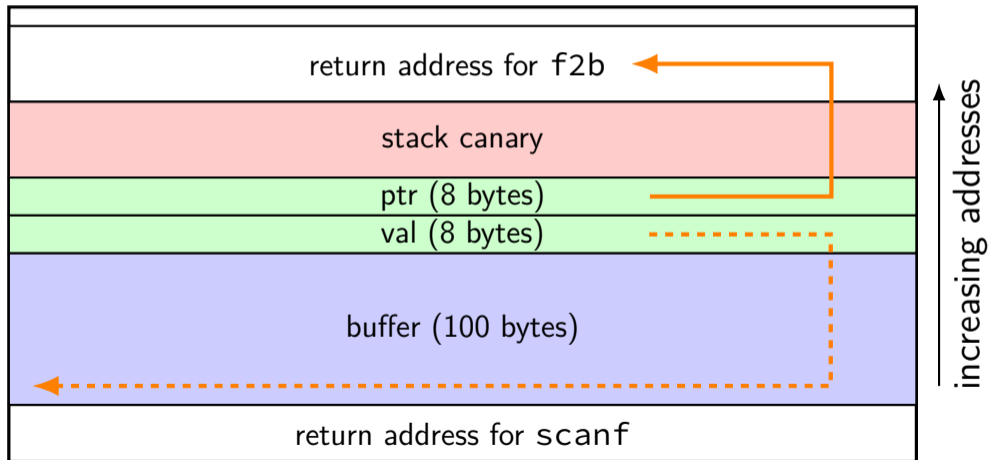


lowest address (stack grows here)



# skipping the canary

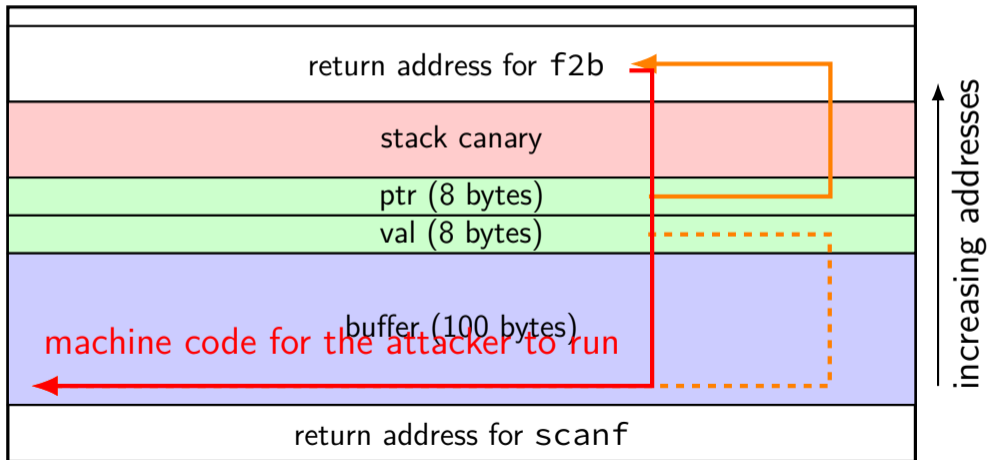
highest address (stack started here)



lowest address (stack grows here)

# skipping the canary

highest address (stack started here)



lowest address (stack grows here)

# exercise (1)

```
void vulnerable() {  
    int *array;  
    char buffer[100];  
    if (!Allocate(&array))  
        abort();  
    gets(buffer);  
    array[0] = atoi(buffer);  
    ...  
}
```

If return address is at 0x12345,  
where/how to place 0x12345 in input?

- A. beginning, as ASCII base-10 number
- B. beginning, as ASCII base-16 number
- C. 100 bytes into buffer, as bytes
- D. 104 bytes into buffer, as bytes
- E. 120 bytes into buffer, as bytes
- F. 136 bytes into buffer, as bytes
- G. none of these

```
vulnerable:  
    pushq %rbp  
    pushq %rbx  
    subq  $136, %rsp  
    movq  %fs:40, %rax  
    movq  %rax, 120(%rsp)  
    xorl  %eax, %eax  
    leaq  104(%rsp), %rdi  
    call  Allocate  
    testl %eax, %eax  
    je   call_abort  
    movq %rsp, %rdi  
    call gets  
    movq 104(%rsp), %rbp  
    movl  $10, %edx  
    movl  $0, %esi  
    movq %rsp, %rdi  
    call strtol  
    movl %eax, 0(%rbp)  
    ...
```

## exercise (2)

```
void vulnerable() {
    int *array;
    char buffer[100];
    if (!Allocate(&array))
        abort();
    gets(buffer);
    array[0] = atoi(buffer);
    ...
}
```

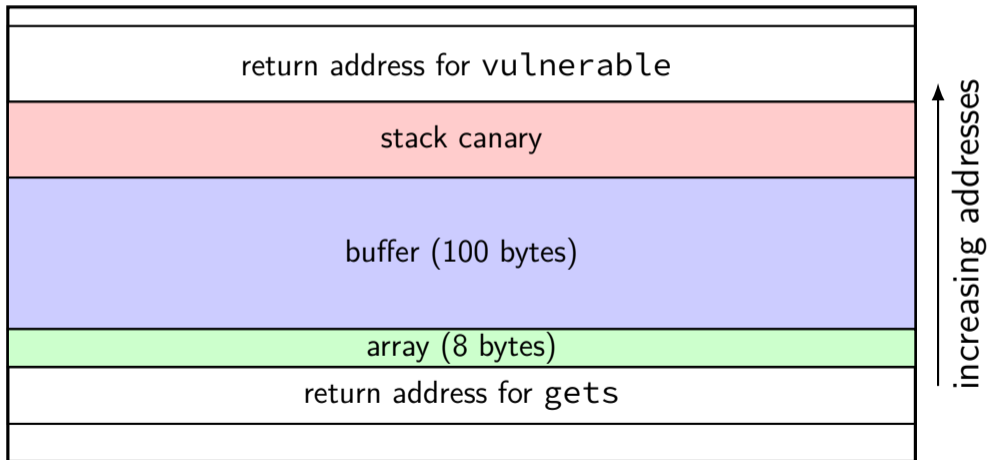
If we want to overwrite ret. addr. with 0x5678, where/how to place 0x5678 in input?

- A. beginning, as ASCII base-10 number
- B. beginning, as ASCII base-16 number
- C. 100 bytes into buffer, as bytes
- D. 104 bytes into buffer, as bytes
- E. 120 bytes into buffer, as bytes
- F. 136 bytes into buffer, as bytes
- G. none of these

```
vulnerable:
    pushq %rbp
    pushq %rbx
    subq $136, %rsp
    movq %fs:40, %rax
    movq %rax, 120(%rsp)
    xorl %eax, %eax
    leaq 104(%rsp), %rdi
    call Allocate
    testl %eax, %eax
    je call_abort
    movq %rsp, %rdi
    call gets
    movq 104(%rsp), %rbp
    movl $10, %edx
    movl $0, %esi
    movq %rsp, %rdi
    call strtol
    movl %eax, 0(%rbp)
    ...
```

# laying out stack to avoid subterfuge (1)

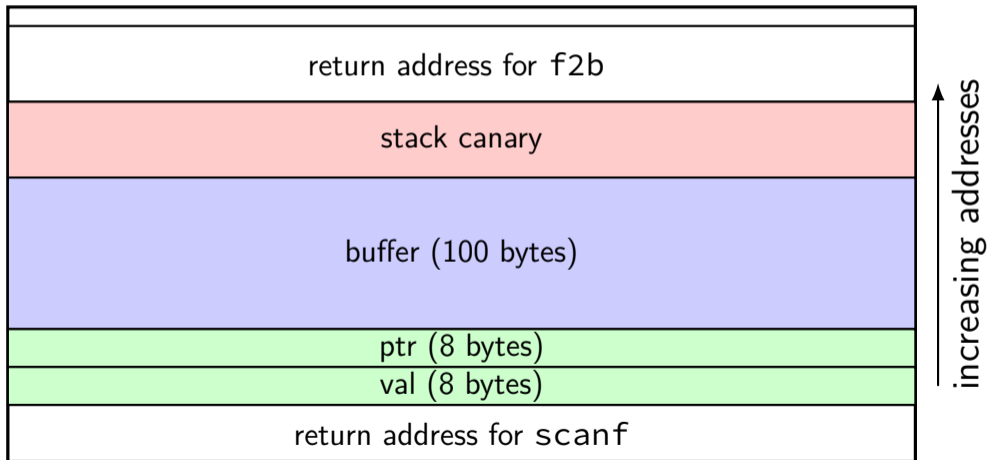
highest address (stack started here)



lowest address (stack grows here)

# laying out stack to avoid subterfuge (2)

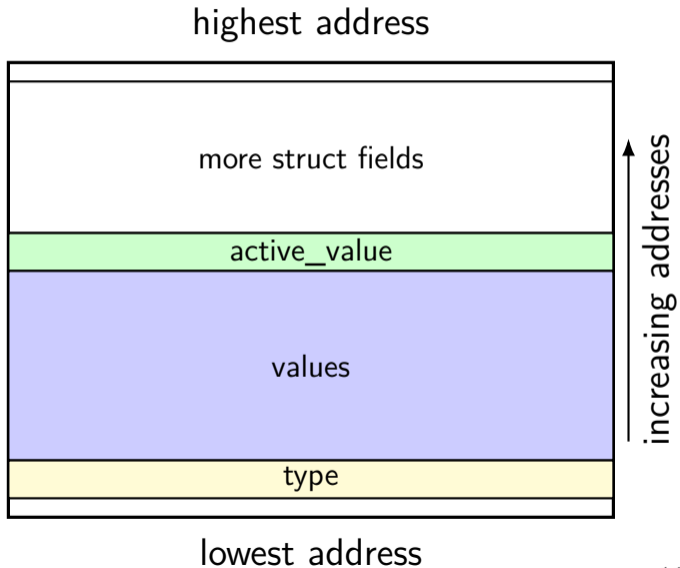
highest address (stack started here)



lowest address (stack grows here)

# other subterfuge cases (1)

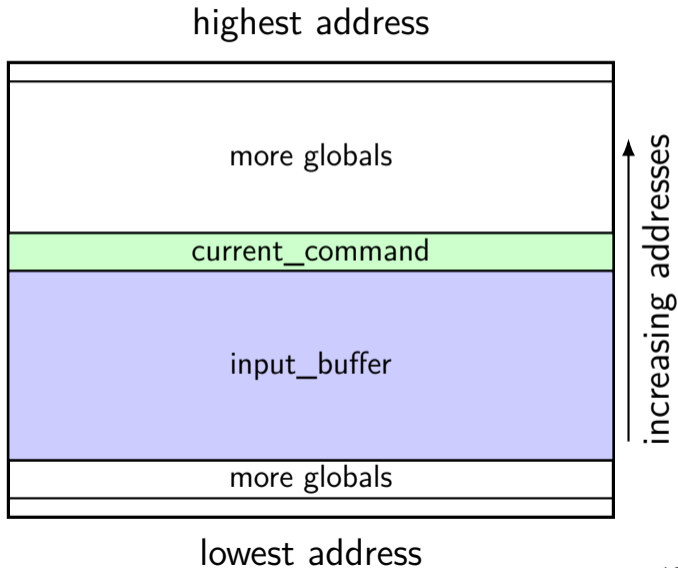
```
struct Command {  
    CommandType type;  
    int values[MAX_VALUES];  
    int *active_value;  
    ...  
};
```



## other subterfuge cases (2)

```
Command *current_command;
char input_buffer[4096];

void run_next_command() {
    if (!current_command) {
        current_command =
            getNext();
    }
    current_command-> ...
    ...
}
```





# beyond return addresses

pointer subterfuge let us overwrite anything

my example: showed return address

but return address is tricky to locate exactly

but there are usually *much easier options!*

# arbitrary memory write

bunch of scenarios that lead to *single arbitrary memory write*

typical result: arbitrary code execution

how?

overwrite existing machine code (insert jump?)

problem: usually not writable

overwrite return address directly

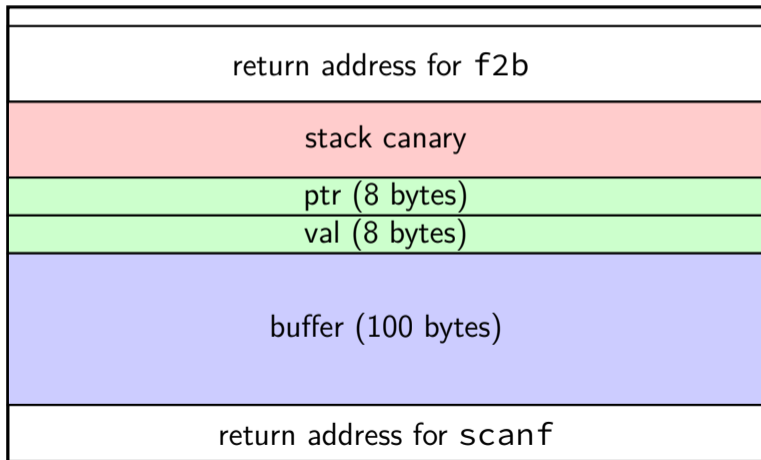
observation: don't care about stack canaries — skip them

*overwrite other function pointer?*

overwrite another data pointer — copy more?

# attacking the GOT

highest address (stack started here)



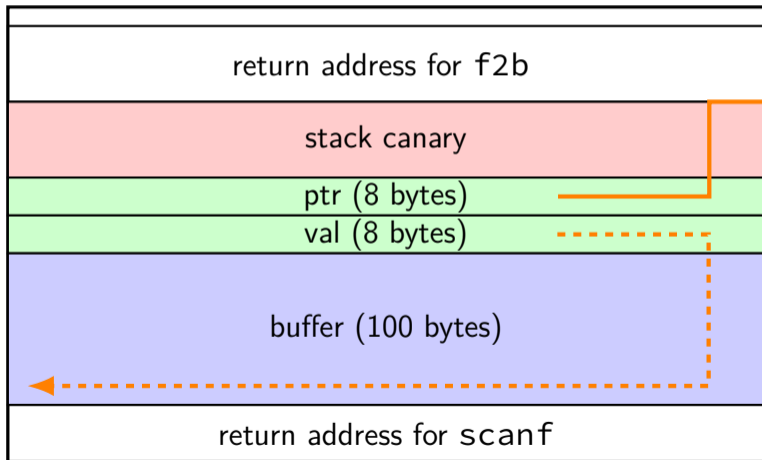
lowest address (stack grows here)

global offset table

GOT entry: printf
GOT entry: fopen
GOT entry: exit

# attacking the GOT

highest address (stack started here)



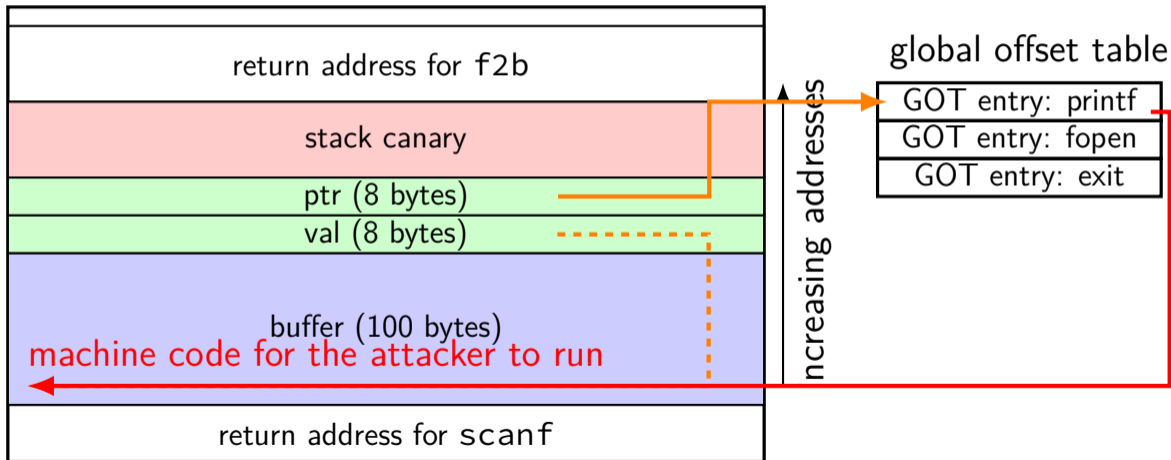
global offset table

GOT entry: printf
GOT entry: fopen
GOT entry: exit

lowest address (stack grows here)

# attacking the GOT

highest address (stack started here)



lowest address (stack grows here)

# arbitrary memory write

bunch of scenarios that lead to *single arbitrary memory write*

typical result: arbitrary code execution

how?

overwrite existing machine code (insert jump?)

problem: usually not writable

overwrite return address directly

observation: don't care about stack canaries — skip them

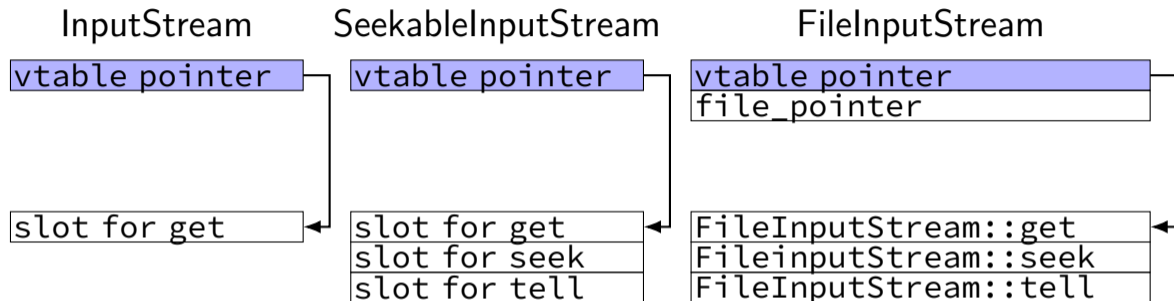
*overwrite other function pointer?*

overwrite another data pointer — copy more?

# C++ inheritance

```
class InputStream {
public:
    virtual int get() = 0;
    // Java: abstract int get();
    ...
};
class SeekableInputStream : public InputStream {
public:
    virtual void seek(int offset) = 0;
    virtual int tell() = 0;
};
class FileInputStream : public SeekableInputStream {
public:
    virtual int get();
    virtual void seek(int offset);
    virtual int tell();
    ...
};
```

# C++ inheritance: approx memory layout





# C++ implementation (pseudo-code)

```
struct InputStream_vtable {  
    int (*get)(InputStream* this);  
};
```

```
struct InputStream {  
    InputStream_vtable *vtable;  
};
```

...

```
InputStream *s = ...;  
int c = (s->vtable->get)(s);
```

# C++ implementation (pseudo-code)

```
struct SeekableInputStream_vtable {  
    struct InputStream_vtable as_InputStream;  
    void (*seek)(SeekableInputStream* this, int offset);  
    int (*tell)(SeekableInputStream* this);  
};
```

```
struct FileInputStream {  
    SeekableInputStream_vtable *vtable;  
    FILE *file_pointer;  
};
```

...

```
FileInputStream file_in = { the_FileInputStream_vtable, ... };  
InputStream *s = (InputStream*) &file_in;
```

# C++ implementation (pseudo-code)

```
SeekableInputStream_vtable the_FileInputStream_vtable = {  
    &FileInputStream_get,  
    &FileInputStream_seek,  
    &FileInputStream_tell,  
};
```

...

```
FileInputStream file_in = { the_FileInputStream_vtable, ... };  
InputStream *s = (InputStream*) &file_in;
```

# calling virtual method

```
SeekableInputStream *in = ...; // 8(%rsp)
in->get();
in->seek(10);
```

---

```
# in->get();
movq    8(%rsp), %rdi // rdi <- in
movq    (%rdi), %rax // rax <- vtable
call    *(%rax)      // call vtable[0]
```

```
# in->seek(10);
movq    8(%rsp), %rdi // rdi <- in
movl    $10, %esi    // esi <- 10
movq    (%rdi), %rax // rax <- vtable
call    *8(%rax)     // call vtable[1]
```

# FileInputStream assembly (1)

```
_ZN15FileInputStreamC2Ev: // constructor
    # rdi == this
    movq    $_ZTV15FileInputStream+16, (%rdi)
    ...
    ret
# VTable for FileInputStream
_ZTV15FileInputStream:
    # offset (for multiple inheritance)
    .quad  0
    # info for typeid() operator
    .quad  _ZTI15FileInputStream
    # VTable pointer points here
    # FileInputStream::get
    .quad  _ZN15FileInputStream3getEv
    # FileInputStream::seek
    .quad  _ZN15FileInputStream4seekEi
    # FileInputStream::tell
    .quad  _ZN15FileInputStream4tellEv
```

# attacking function pointer tables

option 1: overwrite table entry directly

required/easy for Global Offset Table — fixed location  
usually not possible for VTables — read-only memory


option 2: create table in buffer (big list of pointers to shellcode),  
point to buffer

useful when table pointer next to buffer  
(e.g. C++ object on stack next to buffer)

option 3: find suitable pointer elsewhere

e.g. point to wrong part of vtable to run different function

## exercise



```
class VulnerableClass {  
public:  
    char buffer[100];  
    virtual void foo();  
    virtual void bar();  
};  
VulnerableClass objs[10];
```

Assume `gets(objs[0].buffer)` is run and eventually `ptr->foo()` will be run where `ptr == &objs[1]`.

input start: \_\_\_\_\_

input+50 bytes: \_\_\_\_\_

input+100 bytes: \_\_\_\_\_

A. shellcode            B. address of `objs[0].buffer[0]`

C. address of `objs[0].buffer[50]`

D. address of original vtable

E. address of `objs[0]`'s vtable

F. address of `objs[1]`'s vtable pointer

# arbitrary memory write

bunch of scenarios that lead to *single arbitrary memory write*

typical result: arbitrary code execution

how?

overwrite existing machine code (insert jump?)

problem: usually not writable

overwrite return address directly

observation: don't care about stack canaries — skip them

overwrite other function pointer?

*overwrite another data pointer — copy more?*



## write to write

```
struct A {
    char name[100];
    long irrelevant;
    ...
    struct B* other_thing;
    ...
};
struct B {
    char name[100];
    ...
}
...
    gets(a_object->name);
    gets(a_object->other_thing->name);
...
```

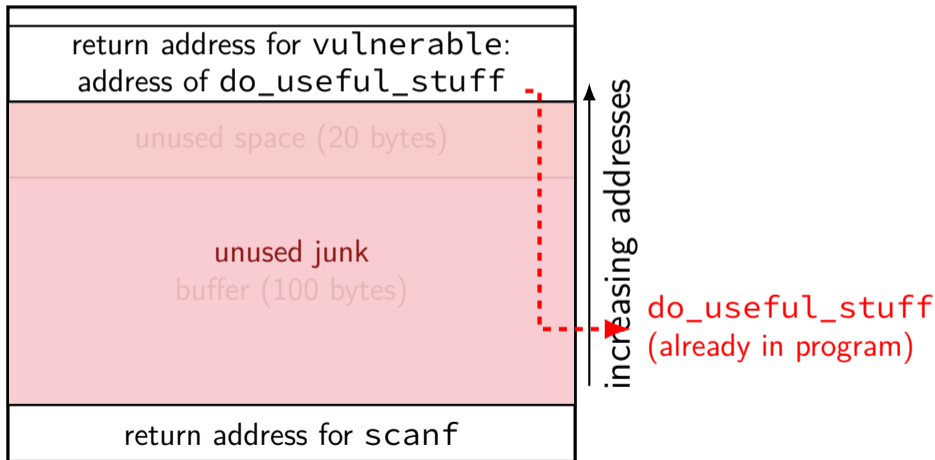
## so far overwrites

once we found a way to overwrite function pointer  
easiest solution seems to be: direct to our code

...but alterante places to direct it to

# return-to-somewhere

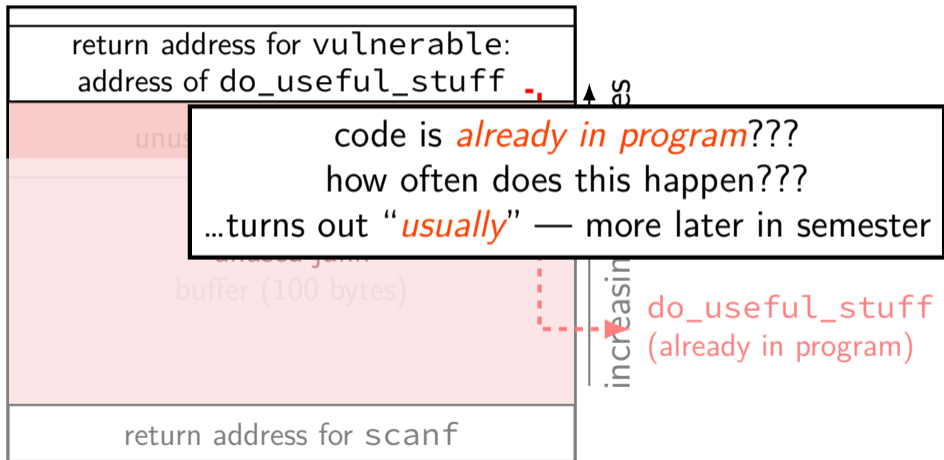
highest address (stack started here)



lowest address (stack grows here)

# return-to-somewhere

highest address (stack started here)



lowest address (stack grows here)

# example: system()

## NAME

`system` - execute a shell command

## SYNOPSIS

```
#include <stdlib.h>
```

```
int system(const char *command);
```

part of C standard library

in any program that dynamically links to libc

challenge: need to hope argument register (rdi) set usefully

# locating system() Linux

```
$ ldd /bin/ls
linux-vdso.so.1 (0x00002aaaaade000)
libselinux.so.1 => /lib/x86_64-linux-gnu/libselinux.so.1 (0x00002aaaaab3a000)
libc.so.6 => /lib/x86_64-linux-gnu/libc.so.6 (0x00002aaaaab65000)
libpcre2-8.so.0 => /usr/lib/x86_64-linux-gnu/libpcre2-8.so.0 (0x00002aaaaad57000)
libdl.so.2 => /lib/x86_64-linux-gnu/libdl.so.2 (0x00002aaaaade7000)
/lib64/ld-linux-x86-64.so.2 (0x00002aaaaaab000)
libpthread.so.0 => /lib/x86_64-linux-gnu/libpthread.so.0 (0x00002aaaaaded000)
$ objdump --dynamic-syms /lib/x86_64-linux-gnu/libc.so.6 | grep system
0000000000156a80 g DF .text 0000000000000067 GLIBC_2.2.5 svcerr_systemerr
0000000000055410 g DF .text 000000000000002d GLIBC_PRIVATE __libc_system
0000000000055410 w DF .text 000000000000002d GLIBC_2.2.5 system
```

if address randomization disabled:

address should be  $0x00002aaaaab650 + 0x55410$

`ldd` — “what libraries does this load and where?”  
similar tools for other OSes

# case study (simplified)

bug in NTPd (Network Time Protocol Daemon)

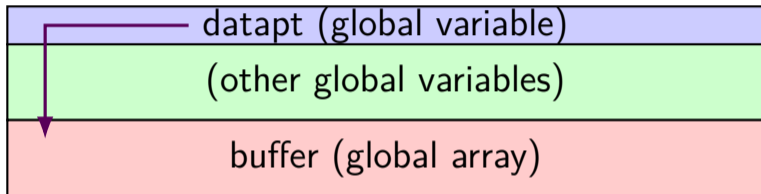
via Stephen Röttger, “Finding and exploiting ntpd vulnerabilities”

<https://googleprojectzero.blogspot.com/2015/01/finding-and-exploiting-ntpd.html>

```
static void
ctl_putdata(
    const char *dp,
    unsigned int dlen,
    int bin    /* set to 1 when data is binary */
) {
    ...
    memmove((char *)datapt, dp, (unsigned)dlen);
    datapt += dlen;
    datalinenelen += dlen;
```

# the target

```
memmove((char *)datapt, dp, (unsigned)dlen);
```



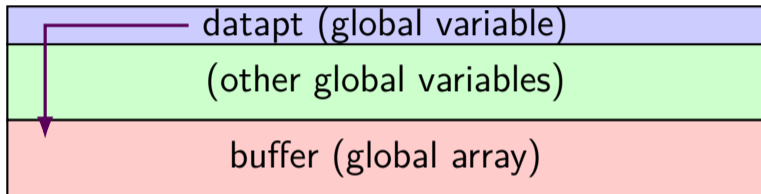


## more context

```
memmove((char *)datapt, dp, (unsigned)dlen);  
...  
...  
strlen(some_user_supplied_string)  
/* calls strlen@plt  
   looks up global offset table entry! */
```

# the target

```
memcpy((char *)datap, dp, (unsigned)dlen);
```



strlen GOT entry

## overall exploit

overwrite `datapt` to point to `strlen` GOT entry

overwrite value of `strlen` GOT entry


example target: `system` function

executes command-line command specified by argument

supply string to provide argument to “`strlen`”

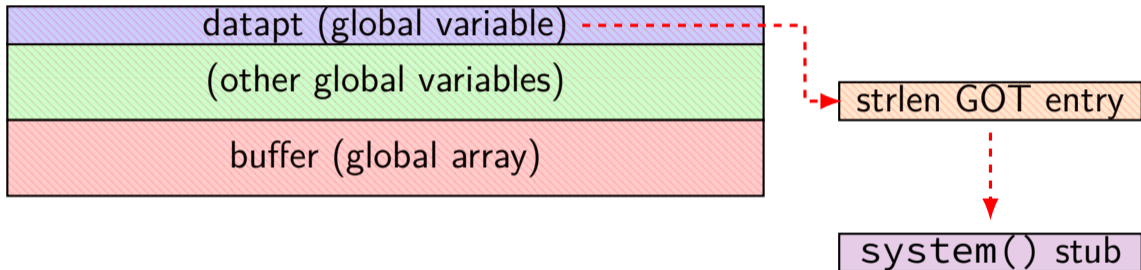
# the target

```
memcpy((char *)datap, dp, (unsigned)dlen);
```



# the target

```
memcpy((char *)datap, dp, (unsigned)dlen);
```



## overall exploit: reality

real exploit was more complicated

needed to defeat more mitigations

needed to deal with not being able to write \0

actually tricky to send things that trigger buffer write  
(meant to be local-only)

# subterfuge exercise

```
struct Student {
    char email[128];
    struct Assignment *assignments[16];
    ...
};
struct Assignment {
    char submission_file[128];
    char regrade_request[1024];
    ...
};
void SetEmail(Student *s, char *new_email) { strcpy(s->email, new_email); }
void AddRegradeRequest(Student *s, int index, char *request) {
    strcpy(s->assignments[index]->regrade_request, request);
}
void vulnerable(char *STRING1, char *STRING2) {
    SetEmail(s, STRING1); AddRegradeRequest(s, 0, STRING2);
}
```

exercise: to set `0x1020304050` to `0xAABBCCDD`, what should `STRING1`, `STRING2` be?

(assume 64-bit pointers, no padding in structs, little-endian)

# subterfuge exercise solution

```
struct Student { char email[128]; struct Assignment *assignments[16]; ... };  
struct Assignment { char submission_file[128]; char regrade_request[1024]; ... };
```

STRING1 (email) controls *what address to overwrite* (want 0x1020304050)

`&s->assignments[0] == &email[128]`

make bytes 128-128+8 be pointer to fake assignment

want fake assignment->regrade\_request address to be 0x1020304050

fake assignment address needs to be at 0x1020304050 - 128

STRING2 (regrade\_reqest) controls *what value to set* (want 0xAABBCCDD)



# backup slides