

CS 489 / 698: Software and Systems Security

Module 2: Program Security (Attacks)

memory errors

Meng Xu (*University of Waterloo*)


Spring 2023

Outline

- 1 Why studying memory errors?
- 2 Background: how a C program executes on a machine?
- 3 Textbook exploitation of a stack overflow vulnerability
- 4 A relatively formal definition to memory error
- 5 Case study: Heartbleed vulnerability
- 6 Concluding remarks

Memory errors are prevalent

% of memory safety vs. non-memory safety CVEs by patch year



Source: [BlackHat IL 2019 talk](#) by Matt Miller from Microsoft

Around 70% of all the vulnerabilities in Microsoft products addressed through a security update each year (2006 - 2018) are memory safety issues

Memory errors are prevalent

High+, impacting stable

Security-related assert

7.1%

Other


23.9%

Other memory unsafety

32.9%

Use-after-free

36.1%




Source: [Chromium Memory Safety Report](#) from Google.

Analysis based on 912 high or critical severity security bugs in Chromium reported in 2015 - 2020

Memory errors are prevalent

Memory Safety Vulnerabilities are Disproportionately Severe



Source: Blog post [Memory Safe Languages in Android 13](#) from Google.

Memory safety vulnerabilities disproportionately represent Android's most severe vulnerabilities

Memory errors can lead to severe consequences



Heartbleed Vulnerability
(CVE-2014-0610)

- A security bug in version 1.0.1 of OpenSSL, which is a widely used implementation of the Transport Layer Security (TLS) protocol
- It was introduced into OpenSSL in 2012 and publicly disclosed in April 2014
- At the time of disclosure, some 17% (around half a million) of the Internet's secure web servers certified by trusted authorities were believed to be vulnerable to the attack

Memory errors can lead to severe consequences



Heartbleed Vulnerability
(CVE-2014-0610)

- The Canada Revenue Agency (CRA) reported a theft of social insurance numbers belonging to 900 taxpayers, and said that they were accessed through an exploit of the bug during a 6-hour period on 8 April 2014.
- After the discovery of the attack, the agency shut down its website and extended the taxpayer filing deadline from 30 April to 5 May.
- On 16 April, the RCMP announced they had charged [a computer science student](#) in relation to the theft with unauthorized use of a computer and mischief in relation to data.

Heartbleed explanation

SERVER, ARE YOU STILL THERE?
IF SO, REPLY "POTATO" (6 LETTERS).



...this page about "books". User Eric's request
secure connection using key "4538538374224"
User Meg wants these 6 letters: POTATO. User
da wants pages about "irl games". Unlocking
secure records with master key 5130985733435
...to (http://www.heartbleed.com) this message: "P



POTATO

...this page about "books". User Eric's request
secure connection using key "4538538374224"
User Meg wants these 6 letters: **POTATO**. User
da wants pages about "irl games". Unlocking
secure records with master key 5130985733435
...to (http://www.heartbleed.com) this message: "P



Source: https://imgs.xkcd.com/comics/heartbleed_explanation.png

Heartbleed explanation

SERVER, ARE YOU STILL THERE?
IF SO, REPLY "BIRD" (4 LETTERS).



User Olivia from London wants pages about "h
ees in car why". Note: Files for IP 375.381.
.83.17 are in /tmp/files-3843. User Meg wants
these 4 letters: BIRD. There are currently 34
connections open. User Brendan uploaded the file
selfie.jpg (contents: 834ba962e2cb9ff89b43bf69



HMM...



User Olivia from London wants pages about "h
ees in car why". Note: Files for IP 375.381.
.83.17 are in /tmp/files-3843. User Meg wants
these 4 letters: **BIRD**. There are currently 34
connections open. User Brendan uploaded the file
selfie.jpg (contents: 834ba962e2cb9ff89b43bf69

BIRD



Source: https://imgs.xkcd.com/comics/heartbleed_explanation.png

Heartbleed explanation

SERVER, ARE YOU STILL THERE?
IF SO, REPLY "HAT" (500 LETTERS).



a connection. Jake requested pictures of deer.
User Meg wants these 500 letters: **HAT**. Lucas requests the "missed connections" page. Eve (administrator) wants to set server's master key to "14835038534". Isabel wants pages about snakes but not too long". User Karen wants to change account password to "Coffee&Tea". User



HAT. Lucas requests the "missed connections" page. Eve (administrator) wants to set server's master key to "14835038534". Isabel wants pages about snakes but not too long". User Karen wants to change account password to "Coffee&Tea". User

a connection. Jake requested pictures of deer.
User Meg wants these 500 letters: **HAT**. Lucas requests the "missed connections" page. Eve (administrator) wants to set server's master key to "14835038534". Isabel wants pages about snakes but not too long". User Karen wants to change account password to "Coffee&Tea". User



Source: https://imgs.xkcd.com/comics/heartbleed_explanation.png

Outline

- 1 Why studying memory errors?
- 2 Background: how a C program executes on a machine?**
- 3 Textbook exploitation of a stack overflow vulnerability
- 4 A relatively formal definition to memory error
- 5 Case study: Heartbleed vulnerability
- 6 Concluding remarks

A simple C program


```
1 #include <stdio.h>
2 #include <string.h>
3
4 int main(void) {
5     char buff[8];
6     int pass = 0;
7
8     printf("Enter the password: ");
9     gets(buff);
10
11     if(strcmp(buff, "warriors")) {
12         printf("Wrong password\n");
13     } else {
14         printf("Correct password\n");
15         pass = 1;
16     }
17
18     if(pass) {
19         printf ("Root privileges granted\n");
20     }
21     return 0;
22 }
```

Try with

```
gcc -m64 -fno-stack-protector
```

And password "golden-hawks"

Von Neumann architecture



Implications of the Von Neumann architecture

- Code and data reside in the same memory space and can be addressed in a unified way
 - If you manage to get the PC register to point to a memory address contains your logic, you have effectively hijacked the control flow.
- There is only one unified memory, it is the job of the compiler / programming language / runtime to find a way to utilize the memory efficiently.
 - Variables declared in a program (e.g., `int i = 0;`) needs to be mapped to an address in the memory, and the mapping logic needs to be (ideally) consistent on the same architecture.


Definition: memory

Q: What is a conventional way of dividing up the “memory”?

A: Four types of memory on a conceptual level:

- Text (where program code is initially loaded to)
- Stack
- Heap
- Global (a.k.a., static)

Memory layout (Linux x86-64 convention)



Example

```
1 #include <stdlib.c>
2
3 /// where is this variable hosted?
4 const char *HELLO = "hello";
5
6 /// where is this variable hosted?
7 long counter;
8
9 void main() {
10     /// where is this variable hosted?
11     int val;
12
13     /// where is this variable hosted?
14     /// where is its content allocated?
15     char *msg = malloc(120);
16
17     /// what is freed here?
18     free(msg);
19
20     /// what is freed here (at end of function)?
21 }
22
23 /// what is freed here (at end of execution)?
```

Example (and answers)

```
1 #include <stdlib.c>
2
3 // this is in the data section
4 const char *HELLO = "hello";
5
6 // this is in the BSS section
7 long counter;
8
9 void main() {
10     // this is in the stack memory
11     int val;
12
13     // the msg pointer is in the stack memory
14     // the msg content is in the heap memory
15     char *msg = malloc(120);
16
17     // msg content is explicitly freed here
18     free(msg);
19
20     // the val and msg pointer is implicitly freed here
21 }
22
23 // the global memory is only destroyed on program exit
```

Outline

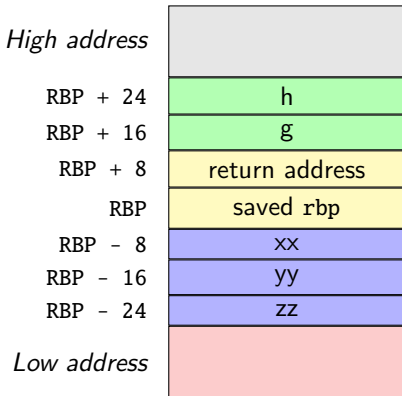
- 1 Why studying memory errors?
- 2 Background: how a C program executes on a machine?
- 3 Textbook exploitation of a stack overflow vulnerability**
- 4 A relatively formal definition to memory error
- 5 Case study: Heartbleed vulnerability
- 6 Concluding remarks

Demo

Demo

Stack layout (Linux x86-64 convention)

```
1 long foo(  
2     long a, long b, long c,  
3     long d, long e, long f,  
4     long g, long h)  
5 {  
6     long xx = a * b * c;  
7     long yy = d + e + f;  
8     long zz = bar(xx, yy, g + h);  
9     return zz + 20;  
10 }
```



Argument a to f passed by registers.

What is heap and why do we need it?

In C/C++, the **heap** is used to manually allocate (and free) **new regions of process memory** during program execution.

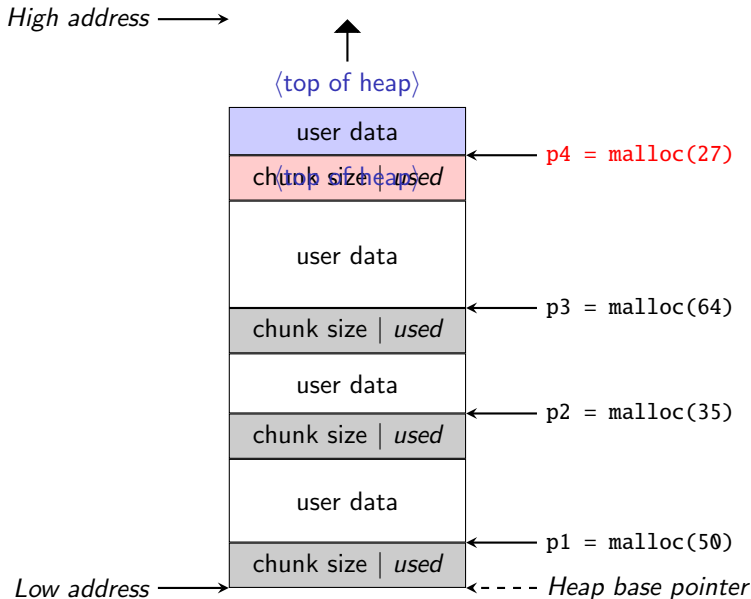
Heap vs stack

```
1 typedef struct Response {
2     int status;
3     char message[40];
4 } response_t;
5
6 response_t *say_hello() {
7     response_t* res =
8         malloc(sizeof(response_t));
9     if (res != NULL) {
10        res->status = 200;
11        strncpy(res->message, "hello", 6);
12    }
13    return res;
14 }
15 void send_back(response_t *res) {
16     // implementation omitted
17 }
18 void process() {
19     response_t *res = say_hello();
20     send_back(res);
21     free(res);
22 }
```

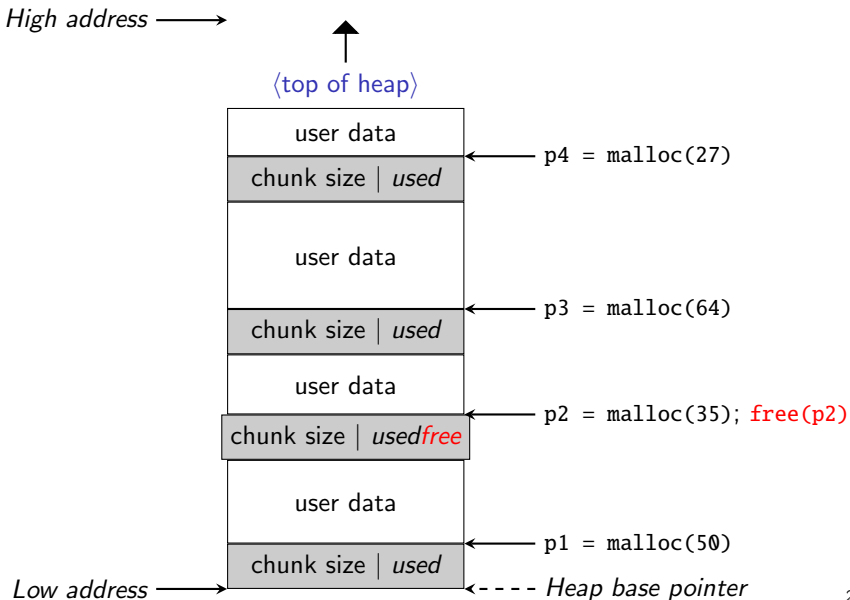
```
1 typedef struct Response {
2     int status;
3     char message[40];
4 } response_t;
5
6 void say_hello(response_t *res) {
7     res->status = 200;
8     strncpy(res->message, "hello", 6);
9 }
10 void send_back(response_t *res) {
11     // implementation omitted
12 }
13 void process() {
14     struct Response res;
15     say_hello(&res);
16     send_back(&res);
17 }
```

A stack-based implementation of
(roughly) the same functionality

Heap: what happens after malloc()?



Heap: what happens after `free()`?



Real-world heap manager

For implementation details of the `glibc`¹ memory allocator, refer to the article from [Azeria Labs](#).

¹GNU C library

For exploitation of memory errors

Smashing The Stack For Fun And Profit

How2Heap — Educational Heap Exploitation

Outline

- ① Why studying memory errors?
- ② Background: how a C program executes on a machine?
- ③ Textbook exploitation of a stack overflow vulnerability
- ④ A relatively formal definition to memory error**
- ⑤ Case study: Heartbleed vulnerability
- ⑥ Concluding remarks

A quick recap

This presentation is about **memory corruption**, a.k.a.,

- **memory errors**, or
- **violations of memory safety properties**, or
- **unsafe programs**

A program is **memory safe** if it is free of memory errors.

Definition: safety

Q: What is “safe” in memory safety?

Observation 1: At runtime, memory is a pool of **objects**

Observation 2: Each **object** has known and limited **size** and **lifetime**

Observation 3: Once allocated, the size of an **object** never changes

Observation 4: A memory access is always **object-oriented**, i.e.

- Memory read: (object_id, offset, length)
- Memory write: (object_id, offset, length, value)

Wait..., in C/C++, pointers are just 32/64-bit integers. I can do:
`int *p = 0xdeadbeef; int v = *p;` Which object I refer to here?

Definition: safety

Q: What is “safety” in memory safety?

At any point of time during the program execution,
for any object in memory, we know its
(object_id, size [int], alive [bool])

At the same time, for each memory access, we know:

- Memory read: (object_id, offset [int], length [int])
- Memory write: (object_id, offset [int], length [int], _)

Definition: spatial safety

At any point of time during the program execution,
for any object in memory, we know its
(**object_id**, **size [int]**, alive [bool])

At the same time, for each memory access, we know:

- Memory read: (**object_id**, **offset [int]**, **length [int]**)
- Memory write: (**object_id**, **offset [int]**, **length [int]**, **_**)

It is a violation of spatial safety if:

- $\text{offset} + \text{length} \geq \text{size}$ or
- $\text{offset} < 0$

Example: spatial safety violations

```
1 int foo(int x) {  
2     int arr[16] = {0};  
3     return arr[x];  
4 }
```

```
1 long foo() {  
2     int a = 0;  
3     return *(long *)(&a);  
4 }
```

Definition: NULL-pointer dereference

```
1 int foo(int *p) {  
2     // it is possible that p == NULL  
3     return *p + 42;  
4 }
```

NULL-pointer dereference is sometimes considered as **undefined behavior** — meaning, its behavior is not given in the C language specification, although most operating systems chooses to panic the program on such behavior.

Definition: NULL-pointer dereference

At any point of time during the program execution,
for any object in memory, we know its
(**object_id** \neq 0, size [int], alive [bool])

At the same time, for each memory access, we know:

- Memory read: (**object_id**, offset [int], length [int])
- Memory write: (**object_id**, offset [int], length [int], _)

It is a NULL-pointer dereference if

- `object_id == 0`

Definition: temporal safety

At any point of time during the program execution,
for any object in memory, we know its
(**object_id**, size [int], **alive** [bool])

At the same time, for each memory access, we know:

- Memory read: (**object_id**, offset [int], length [int])
- Memory write: (**object_id**, offset [int], length [int], _)
- Memory free: (**object_id**)

It is a violation of temporal safety if:

- !alive

Example: temporal safety violations

```
1 int foo() {
2     int *p = malloc(sizeof(int));
3     *p = 42;
4     free(p);
5     return *p;
6 }
```

```
1 int *ptr;
2
3 void foo() {
4     int p = 100;
5     ptr = &p;
6 }
7 int bar() {
8     return *ptr;
9 }
```

```
1 int foo() {
2     int *p = malloc(sizeof(int));
3     *p = 42;
4     free(p);
5     free(p);
6     return *p;
7 }
```

Definition: temporal safety (revisited)

At any point of time during the program execution,
for any object in memory, we know its
(**object_id**, size [int], **status** [alloc|init|dead])

At the same time, for each memory access, we know:

- Memory read: (**object_id**, offset [int], length [int])
- Memory write: (**object_id**, offset [int], length [int], _)
- Memory free: (**object_id**)

It is a violation of temporal safety if:

- Read: **status** != **init**
- Write: **status** == **dead**
- Free: **status** == **dead**

Example: temporal safety violations

```
1 int foo() {  
2     int p;  
3     return p;  
4     // what is the value returned?  
5 }
```

```
1 int foo() {  
2     int *p = malloc(sizeof(int));  
3     return *p;  
4     // what is the value returned?  
5 }
```

Definition: memory leak

At any point of time during the program execution,
for any object in memory, we know its
(**object_id**, size [int], **status** [alloc|init|dead])

At the same time, for each memory access, we know:

- Memory read: (**object_id**, offset [int], length [int])
- Memory write: (**object_id**, offset [int], length [int], _)
- Memory free: (**object_id**)

It is a memory leak if **exists one object_id** whose:

- status != dead

Example: memory leak

```
1 int foo() {
2     int *p = malloc(sizeof(int));
3     int *q = malloc(sizeof(int));
4     *p = 42;
5     free(q);
6     return *p;
7 }
```

Outline

- 1 Why studying memory errors?
- 2 Background: how a C program executes on a machine?
- 3 Textbook exploitation of a stack overflow vulnerability
- 4 A relatively formal definition to memory error
- 5 Case study: Heartbleed vulnerability**
- 6 Concluding remarks

Heartbleed vulnerability I

```
1 int dtls1_process_heartbeat(SSL *s) {
2     unsigned char *p = &s->s3->rrec.data[0], *pl;
3     unsigned short hbtype;
4     unsigned int payload;
5     unsigned int padding = 16; /* Use minimum padding */
6
7     /* Read type and payload length first */
8     hbtype = *p++;
9     n2s(p, payload);
10    pl = p;
11
12    /* ... redacted ... */
13
14    if (hbtype == TLS1_HB_REQUEST) {
15        unsigned char *buffer, *bp;
16
17        /* Allocate memory for the response */
18        buffer = OPENSSL_malloc(1 + 2 + payload + padding);
19        bp = buffer;
20
21        /* Enter response type, length and copy payload */
22        *bp++ = TLS1_HB_RESPONSE;
```

Heartbleed vulnerability II

```
23     s2n(payload, bp);
24     memcpy(bp, pl, payload);
25
26     /* Random padding */
27     RAND_pseudo_bytes(bp, padding);
28
29     /* Send out the response */
30     r = dtls1_write_bytes(
31         s, TLS1_RT_HEARTBEAT, buffer, 3 + payload + padding
32     );
33
34     /* ... redacted ... */
35
36     /* Clean-up used resources */
37     OPENSSL_free(buffer);
38     return r;
39 }
40
41 else { /* ... redacted ... */ }
42 }
```

Patch for the Heartbleed vulnerability I

```
1 diff --git a/ssl/d1_both.c b/ssl/d1_both.c
2 index 7a5596a6b3..2e8cf681ed 100644
3 @@ -1459,26 +1459,36 @@ dtls1_process_heartbeat(SSL *s)
4     unsigned int payload;
5     unsigned int padding = 16; /* Use minimum padding */
6
7 - /* Read type and payload length first */
8 - hbtype = *p++;
9 - n2s(p, payload);
10 - pl = p;
11 -
12     if (s->msg_callback)
13         s->msg_callback(0, s->version, TLS1_RT_HEARTBEAT,
14             &s->s3->rrec.data[0], s->s3->rrec.length,
15             s, s->msg_callback_arg);
16
17 + /* Read type and payload length first */
18 + if (1 + 2 + 16 > s->s3->rrec.length)
19 +     return 0; /* silently discard */
20 + hbtype = *p++;
21 + n2s(p, payload);
22 +
```

Patch for the Heartbleed vulnerability II

```
23 +   if (1 + 2 + payload + 16 > s->s3->rrec.length)
24 +       return 0; /* silently discard per RFC 6520 sec. 4 */
25 +   pl = p;
26 +
27   if (hbtype == TLS1_HB_REQUEST)
28       {
29       unsigned char *buffer, *bp;
30 +     unsigned int write_length = 1 /* heartbeat type */ +
31 +         2 /* heartbeat length */ + payload + padding;
32       int r;
33
34 +     if (write_length > SSL3_RT_MAX_PLAIN_LENGTH)
35 +         return 0;
36 +
37       /* Allocate memory for the response, size is 1 byte
38        * message type, plus 2 bytes payload length, plus
39        * payload, plus padding
40        */
41 -     buffer = OPENSSL_malloc(1 + 2 + payload + padding);
42 +     buffer = OPENSSL_malloc(write_length);
43     bp = buffer;
```

Outline

- ① Why studying memory errors?
- ② Background: how a C program executes on a machine?
- ③ Textbook exploitation of a stack overflow vulnerability
- ④ A relatively formal definition to memory error
- ⑤ Case study: Heartbleed vulnerability
- ⑥ Concluding remarks**

Memory errors are prevalent

High+, impacting stable

Security-related assert

7.1%

Other


23.9%

Other memory unsafety

32.9%

Use-after-free

36.1%



Source: [Chromium Memory Safety Report](#) from Google.

Analysis based on 912 high or critical severity security bugs in Chromium reported in 2015 - 2020

Statistics can be misleading...


This is a personal note: one explanation why we have a **disproportionately** high number of memory errors reported amongst all security vulnerabilities is that — **we know memory errors too well.**

- Memory errors have **universally** accepted definitions (e.g., why the website is named **Stack Overflow?**)
 - Once you find a memory error, you do not need to diligently argue that this is a bug and not a feature
- Memory errors often lead to a set of known consequences that are **generally** considered severe (e.g., data leak or denial-of-service)
 - Once you find a memory error, you do not need to construct a working exploit to justify it
- Finding memory errors typically **do not require** program-specific domain knowledge (the bug is rooted in C/C++ language semantics instead of program logic)
 - If you have a technique that can find memory errors in one codebase, you can scale it up to millions of codebases developed in C/C++.

In fact, very few types of vulnerabilities meet these requirements.

Gradual adoption of memory-safe languages

Memory Safety Vulnerabilities Per Year




Source: Blog post [Memory Safe Languages in Android 13](#) from Google.

Number of memory safety vulnerabilities starts to decrease with the adoption of memory-safe languages

Gradual adoption of memory-safe languages

Memory unsafe code and Memory safety vulnerabilities




Source: Blog post [Memory Safe Languages in Android 13](#) from Google.

Number of memory safety vulnerabilities correlates to the portion of unsafe code

Gradual adoption of memory-safe languages

New Native Code



Source: Blog post [Memory Safe Languages in Android 13](#) from Google.

Rust on the rise in Android native implementations

⟨ **End** ⟩