

Lec13: Heap Exploitation

Taesoo Kim

Administrivia




- In-class CTF on **Dev 1** → Gathering at [Coda 9th, Atrium](#)
- Submit your team's challenge by **Nov 27**

NSA Codebreaker Challenge

- NSA Codebreaker Challenge → Due: Dec 08

Rank ▲	School	Task 0 ▼	Task 1 ▼	Task 2 ▼	Task 3 ▼	Task 4 ▼	Task 5 ▼	Task 6 ▼	Task 7 ▼	Task 8 ▼	Task 9 ▼	Score ▼
1	Georgia Institute of Technology	196	141	124	49	29	11	8	3	3	3	80,806.00
2	University of North Georgia	215	157	129	61	28	23	17	1	0	0	67,185.00
3	University of California, Santa Cruz	138	116	104	16	11	7	4	3	3	3	57,698.00
4	Dakota State University	158	99	82	22	13	10	8	2	1	1	43,048.00
5	Strayer University	142	65	49	20	14	12	8	1	0	0	30,292.00
6	SANS Technology Institute	63	51	41	19	16	10	6	2	0	0	28,673.00
7	Carnegie Mellon University	78	45	38	11	4	2	1	1	1	1	20,528.00

Trends of Vulnerability Classes



Classifying Heap Vulnerabilities

- Common: buffer overflow/underflow, out-of-bound read
 - *Much prevalent* (i.e., quality, complexity)
 - *Much critical* (i.e., larger attack surface)
- Heap-specific issues:
 - **Use-after-free** (e.g., dangled pointers)
 - Incorrect uses (e.g., double frees)

Simple High-level Interfaces

```
// allocate a memory region (an object)  
void *malloc(size_t size);
```

```
// free a memory region  
void free(void *ptr);
```

```
// allocate a memory region for an array  
void *calloc(size_t nmemb, size_t size);
```

```
// resize/reallocate a memory region  
void *realloc(void *ptr, size_t size);
```

```
// in C++
```

```
// new Type == malloc(sizeof(Type))
```

```
// new Type[size] == malloc(sizeof(Type)*size) -- Q. problem?
```

Review: Heap Allocation APIs

- Q0. `ptr = malloc(size); *ptr?`
- Q1. `ptr = malloc(0); ptr == NULL?`
- Q2. `ptr = malloc(-1); ptr == NULL?`
- Q3. `ptr = malloc(size); ptr == NULL` but valid? */* vaddr = 0? */*

- Q4. `free(ptr); ptr == NULL?`
- Q5. `free(ptr); *ptr?`
- Q6. `free(NULL)?`
- Q7. `free(ptr); free(ptr)?`

- Q8. `realloc(ptr, size); *ptr?`
- Q9. `realloc(NULL, size)?`
- Q10. `ptr = calloc(nmemb, size); *ptr?`

CS101: Common Goals of Heap Allocators

1. Performance
2. Memory fragmentation
3. Security

// either fast, secure, (external) fragmentation!

1. malloc() -> mmap()	& free() -> unmap()
2. malloc() -> brk()	& free() -> nop
3. malloc() -> base += size; return base	& free() -> nop

Memory Allocators

Allocators	B	I	C	Description (applications)
ptmalloc	✓	✓	✓	A default allocator in Linux
dlmalloc	✓	✓	✓	An allocator that ptmalloc is based on
jemalloc	✓		✓	A default allocator in FreeBSD
tcmalloc	✓	✓	✓	A high-performance allocator from Google
PartitionAlloc	✓		✓	A default allocator in Chromium
libumem	✓		✓	A default allocator in Solaris


- Recently, [Mimalloc](#) by Microsoft and [Scudo](#) by LLVM

Common Design Choices (Security-Related)

1. **Binning**: size-base groups/operations
 - e.g., caching the same size objects together
2. **In-place metadata**: metadata before/after or even inside
 - e.g., putting metadata inside the freed region
3. **Cardinal metadata**: no encoding, direct pointers and sizes
 - e.g., using raw pointers for linked lists

ptmalloc in Linux: Memory Allocation

```
ptr = malloc(size);
```

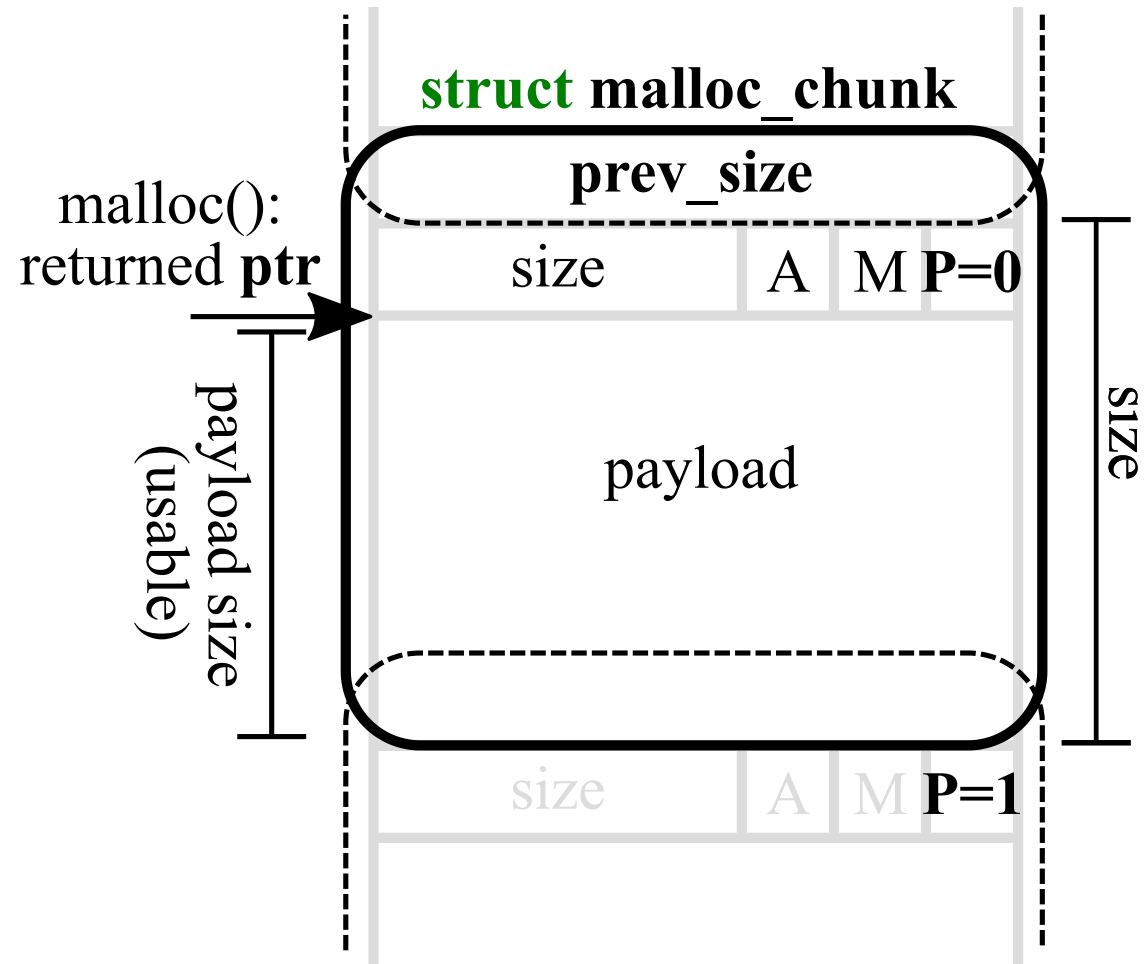


ptmalloc in Linux: Data Structure (glibc)

```
1  struct malloc_chunk {
2      // size of "previous" chunk
3      // (only valid when the previous chunk is freed, P=0)
4      size_t prev_size;
5
6      // size in bytes (aligned by double words): lower bits
7      // indicate various states of the current/previous chunk
8      //   A: allocated in a non-main arena
9      //   M: mmapped
10     //   P: "previous" in use (i.e., P=0 means freed)
11     size_t size;
12
13     [...]
14 };
```

- Q. How to know if the current chunk is in-use or freed?

ptmalloc in Linux: Memory Allocation



(a) allocated chunk

Remarks: Memory Allocation

- Given an allocated ptr,
 1. Immediately lookup its size (SIZE)
 2. Check if the **previous** object is allocated/freed ($P = 0$ or 1)
 3. Check if the **next** object is allocated/freed (Q. how?)
 4. Iterate to the next object ($\text{ptr} + \text{SIZE}$)
 5. Iterate to the prev object **if freed** ($\text{ptr} - \text{PREV_SIZE}$)
 6. Not possible to iterate to the previous object if allocated


ptmalloc in Linux: Data Structure

```
1  struct malloc_chunk {
2    [...]
3    // double links for free chunks in small/large bins
4    // (only valid when this chunk is freed)
5    struct malloc_chunk* fd;
6    struct malloc_chunk* bk;
7
8    // double links for next larger/smaller size in
9    // (only valid when this chunk is freed)
10   struct malloc_chunk* fd_nextsize;
11   struct malloc_chunk* bk_nextsize;
12  };
```

rgebins

- Q. What if we access fd/bk of the allocated chunk?

ptmalloc in Linux: Memory Free



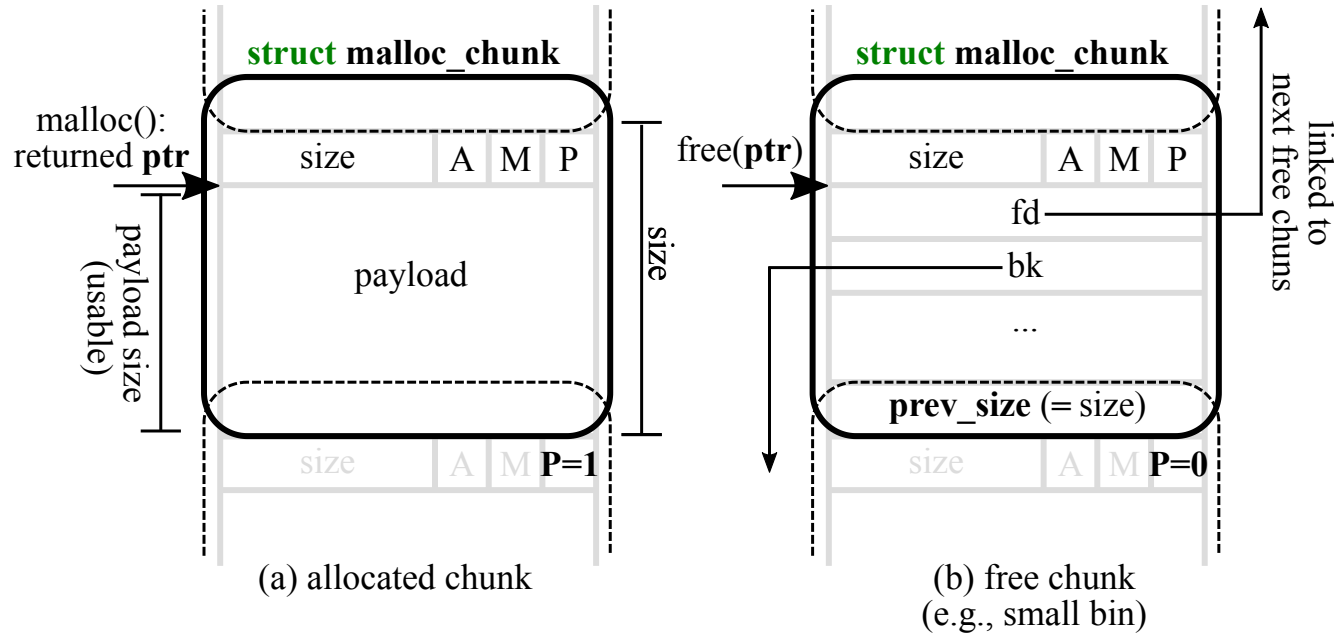
(b) free chunk

Remarks: Memory Free

- Given a free-ed ptr,
 1. All benefits as an allocated ptr (previous remarks)
 2. Iterate b/w **free** objects via fd/bk links
- Invariant: **no two adjacent** free objects
 1. When free() invoked, it is always consolidated to adjacent (i.e., fd/bk) objects!

Understanding Modern Heap Allocators

- Maximize memory usage: reusing free memory regions!
- Data structure to minimize fragmentation (i.e., fd/bk consolidation)
- Data structure to maximize performance (i.e., $O(1)$ in free/malloc)



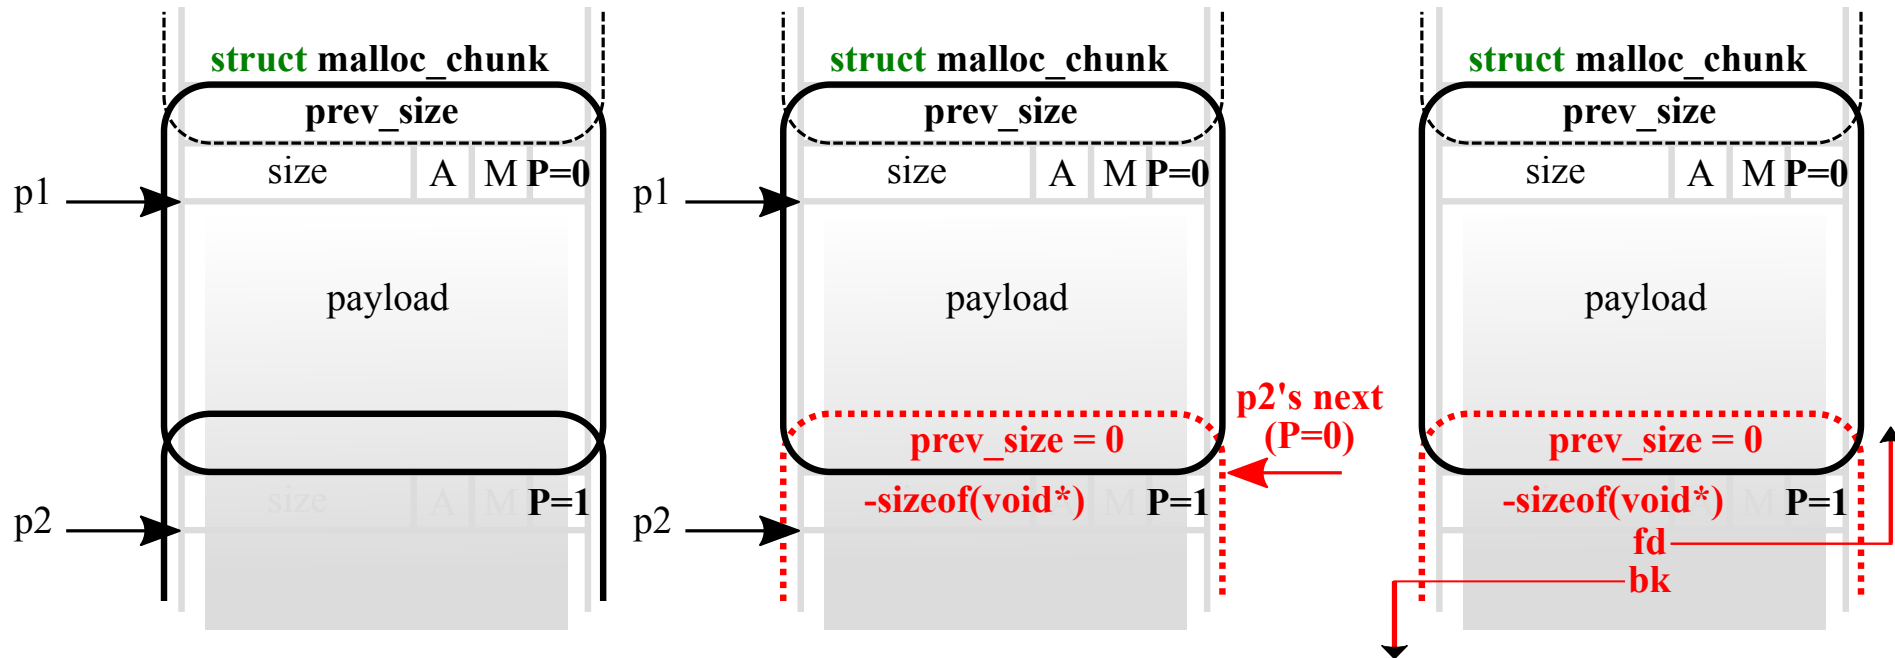
Security Implication of Heap Overflows

- A heap overflow can overwrite the heap metadata
- Incorrect API invocation would destroy the consistency of the metadata
- Allocated/freed objects can be easily crafted for benefits (and fun!)

```
1  void *p1 = malloc(sz);
2  void *p2 = malloc(sz);
3
4  /* overflow on p1 */
5
6  free(p1);
```

Example: Unsafe Unlink (< glibc 2.3.3)

1. Overwriting to p2's size to `-sizeof(void*)`, treating now as if p2 is free
2. When `free(p1)`, attempt to consolidate it with p2 as p2 is free



Example: Unsafe Unlink (< glibc 2.3.3)

- To consolidate, perform unlink on p2 (removing p2 from the linked list)
- Crafted `fd/bk` when `unlink()` result in an arbitrary write!

```
1 // unlink(P):
2   FD = P->fd;
3   BK = P->bk;
4   FD->bk = BK; // NOTE. let's abuse this write!
5   BK->fd = FD;
6
7   p2's fd = dst - offsetof(struct malloc_chunk, bk);
8   p2's bk = val;
9
10  => *dst = val (arbitrary write!) // NOTE. any catch?
```

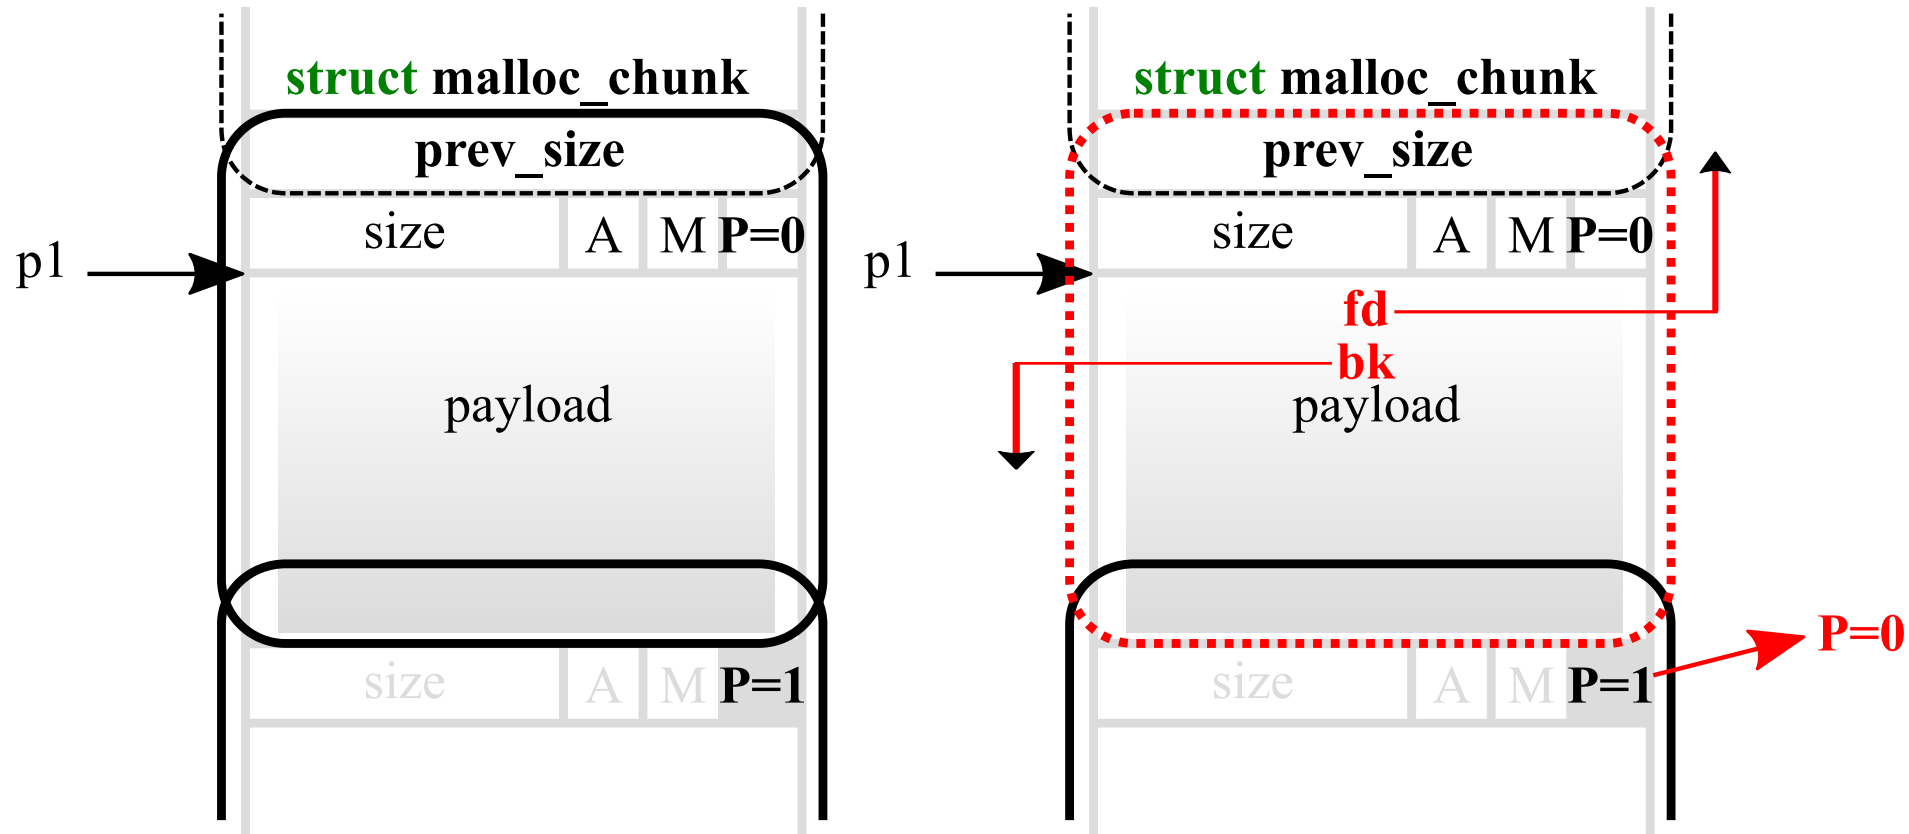
- Q. How to prevent this exploit technique?

Example: Mitigation on Unlink (glibc 2.27)

```
1  #define unlink(AV, P, BK, FD)
2      /* (1) checking if size == the next chunk's prev_size
3  +   if (chunksize(P) != prev_size(next_chunk(P)))
4  +   malloc_printerr("corrupted size vs. prev_size");
5   FD = P->fd;
6   BK = P->bk;
7   /* (2) checking if prev/next chunks correctly point to
*/
8  +   if (FD->bk != P || BK->fd != P)
9  +   malloc_printerr("corrupted double-linked list");
10 +   else {
11     FD->bk = BK;
12     BK->fd = FD;
13     ...
14 + }
```

- Q. Does it prevent the exploit completely?

Security Implication of NULL Overflow in Heap



Heap Exploitation Techniques!

Fast bin dup	House of einherjar
Fast bin dup into stack	House of orange
Fast bin dup consolidate	Tcache dup
Unsafe unlink	Tcache house of spirit
House of spirit	Tcache poisoning
Poison null byte	Tcache overlapping chunks
House of lore	*Unsorted bin into stack
Overlapping chunks 1	*Fast bin into other bin
Overlapping chunks 2	*Overlapping small chunks
House of force	*Unaligned double free
Unsorted bin attack	*House of unsorted einherjar

NOTE. * are what our group recently found and reported!

Use-after-free

- Simple in concept, but difficult to spot in practice!
- Q. Why is it so critical in terms of security?

```
1 int *ptr = malloc(size);  
2 free(ptr);  
3  
4 *ptr; // BUG. use-after-free!
```

Use-after-free

1. What would be the `*ptr`? if nothing happened?
2. What if another part of code invoked `malloc(size)`?

```
1 int *ptr = malloc(size);  
2 free(ptr);  
3  
4 *ptr; // BUG. use-after-free!
```

Use-after-free: Security Implication

1. What would be the *ptr? if nothing happened?


- → Heap pointer leakage (e.g., fd/bk)

2. What if another part of code invoked malloc(size)?

- → Hijacking function pointers (e.g., handler)


```
1  struct msg { ... void (*handler)(); ... };
2
3  struct msg *ptr = malloc(size);
4  free(ptr);
5
6  // later ...
7
8  ptr->handler(); // BUG. use-after-free!
```

Use-after-free with Application Context




```
1 class Div: Element;
2 class Body: Element;
3 class Document { Element* child; };
```

Use-after-free with Application Context



```
1 class Div: Element;  
2 class Body: Element;  
3 class Document { Element* child; };  
4  
5 // (a) memory allocations  
6 Document *doc = new Document();  
7 Body *body = new Body();  
8 Div *div = new Div();
```

Dangled Pointers and Use-after-free



```
1 // (b) using memory: propagating pointers
2 doc->child = body;
3 body->child = div;
4
5 // (c) memory free: doc->child is now dangled
6 delete body;
7
8 // (d) use-after-free: dereference the dangled pointer
9 if (doc->child)
10     doc->child->getAlign();
```

API Misuse: Double Free

1. What happen when free two times?
2. What happen for following malloc()s?

```
1 char *ptr = malloc(size);
2 free(ptr);
3 free(ptr);    // BUG!
4
5 malloc(size); // Q. what does it likely return?
6 malloc(size); // Q. what does it likely return?
```

Binning and Security Implication

- e.g., size-based caching (e.g., fastbin)

(fastbin)

Bins

sz=16 [-]---->[fd]---->[fd]-->NULL

sz=24 [-]---->[fd]---->NULL

sz=32 [-]---->NULL

...

Double Free

- Bins after doing free() two times

```
1 char *ptr = malloc(sz=16);  
2 free(ptr);  
3 free(ptr); // BUG!
```

(fastbin)

```
      Bins  ptr      ptr  
sz=16 [ -]---->[XX]---->[XX]---->[fd]---->[fd]-->NULL  
sz=24 [ -]---->[fd]---->NULL  
sz=32 [ -]---->NULL  
...
```

Double Free: Security Implication

```

1 char *ptr = malloc(sz=16);
2 free(ptr);
3 free(ptr); // BUG!
4
5 ptr1 = malloc(sz=16) // hijacked!
6 ptr2 = malloc(sz=16) // hijacked! Q. why problematic?

```

(fastbin)

```

Bins          (1)  (2)
      +-----+-----+
      |           v     v
sz=16 [ -]---+ [XX]---->[XX] +--->[fd]---->[fd]-->NULL
sz=24 [ -]---->[fd]---->NULL
sz=32 [ -]---->NULL
...

```

Double Free: Mitigation

- Check if the bin contains the pointer that we'd like to free()

```
1 // @glibc/malloc/malloc.c
2
3 /* Check that the top of the bin is not the record we
e going to
4     add (i.e., double free). */
5 if (__builtin_expect (old == p, 0))
6     malloc_printerr ("double free or corruption
asttop)");
7     ...
```

- Q. How to bypass?

Summary

- Two classes of **heap**-related vulnerabilities
 - Traditional: buffer overflow/underflow, out-of-bound read
 - Specific: **use-after-free**, **dangled pointers**, double free
- Understand why they are security critical and non-trivial to eliminate!
- Mitigation approaches taken by allocators

Today's Tutorial

- In-class tutorial:
 - Exploring common techniques
 - Exploiting tcache (simple binning)

```
$ ssh lab09@54.88.195.85  
Password: <password>
```

```
$ cd tut09-advheap
```

References

- [CVE-2014-0160](#)
- [CVE-2018-11360](#)
- [CVE-2018-17182](#)
- [Vudo - An object superstitiously believed to embody magical powers](#)