


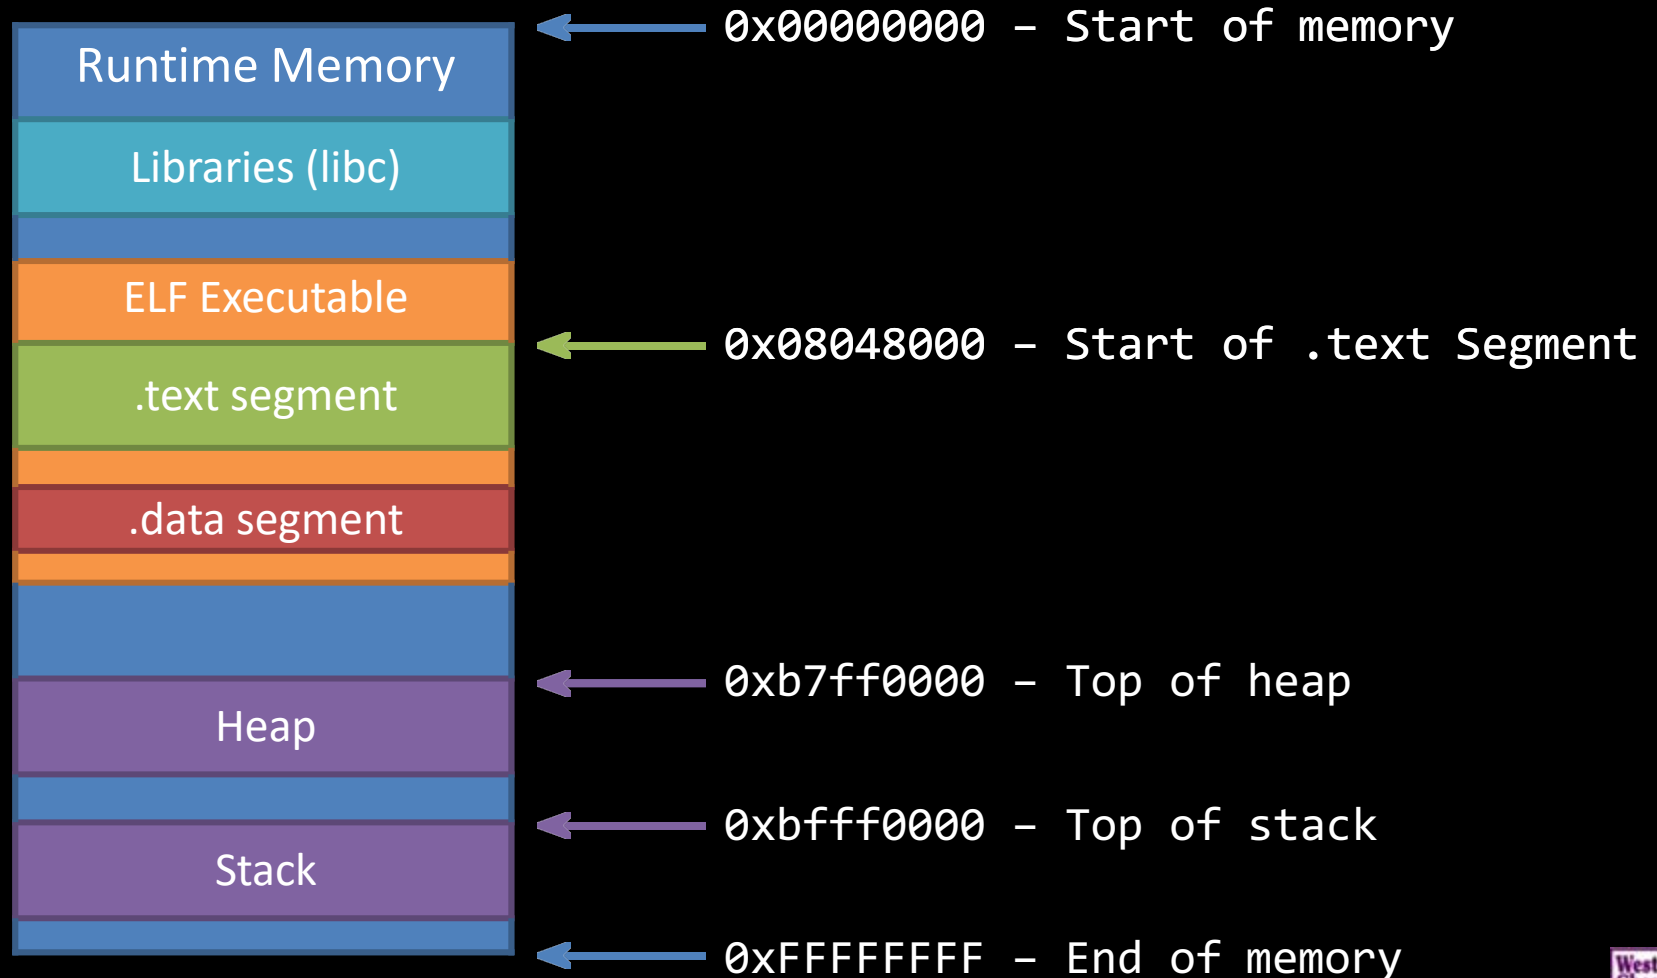
CSC 472 Software Security

Heap Exploitation (3): House of Force

Dr. Si Chen (schen@wcupa.edu)



Pseudo Memory Map



Heap in Linux (GNU C Library – glibc)

ptmalloc2



System call:

brk()

mmap()

DESCRIPTION

brk() and **sbrk()** change the location of the program break, which defines the end of the process's data segment (i.e., the program break is the first location after the end of the uninitialized data segment). Increasing the program break has the effect of allocating memory to the process; decreasing the break deallocates memory.

brk() sets the end of the data segment to the value specified by addr, when that value is reasonable, the system has enough memory, and the process does not exceed its maximum data size (see **setrlimit(2)**).

sbrk() increments the program's data space by increment bytes. Calling **sbrk()** with an increment of 0 can be used to find the current location of the program break.

NAME

mmap, munmap - map or unmap files or devices into memory

SYNOPSIS

```
#include <sys/mman.h>
```

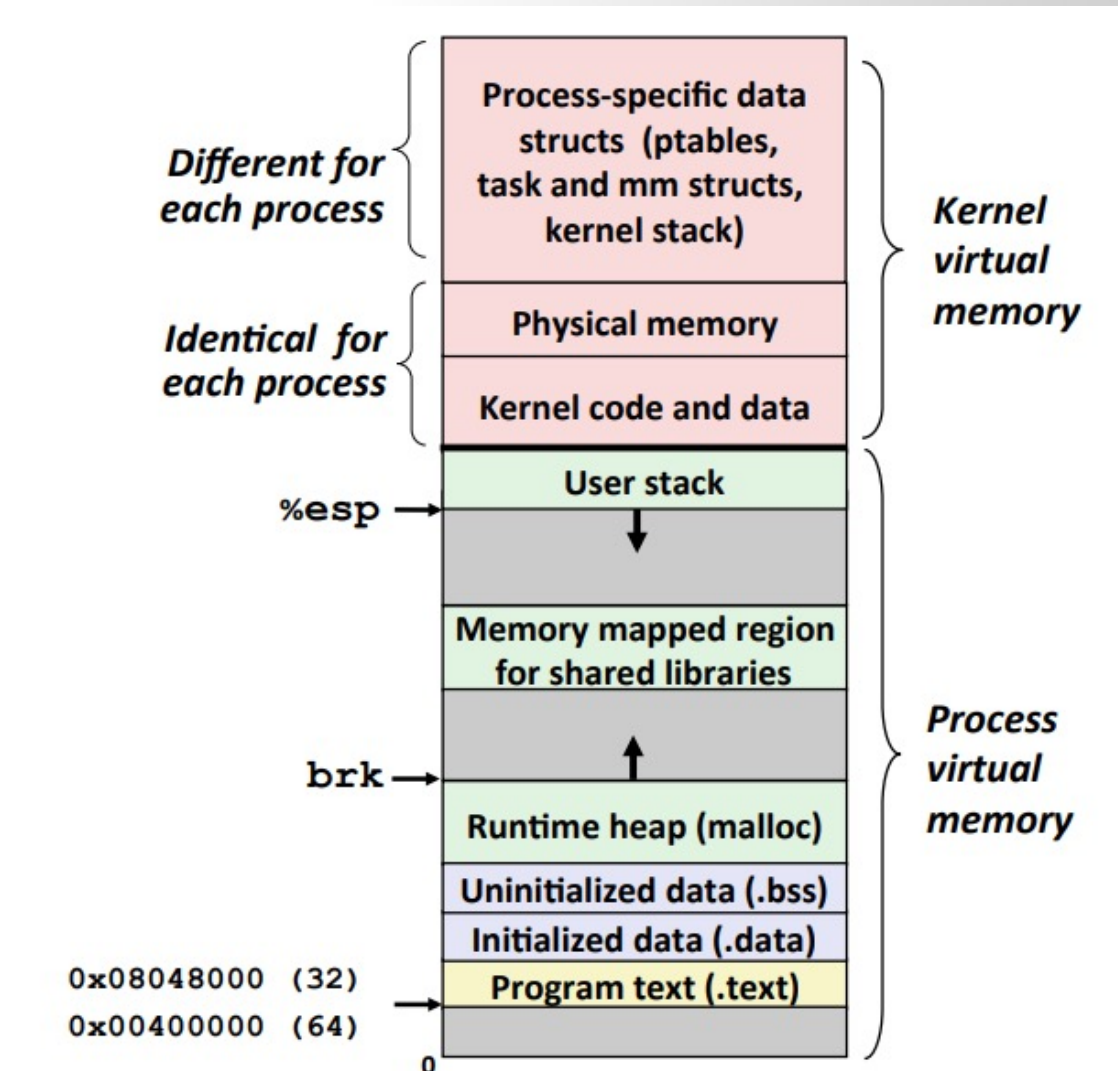
```
void *mmap(void *addr, size_t length, int prot, int flags,  
           int fd, off_t offset);  
int munmap(void *addr, size_t length);
```

See NOTES for information on feature test macro requirements.

DESCRIPTION

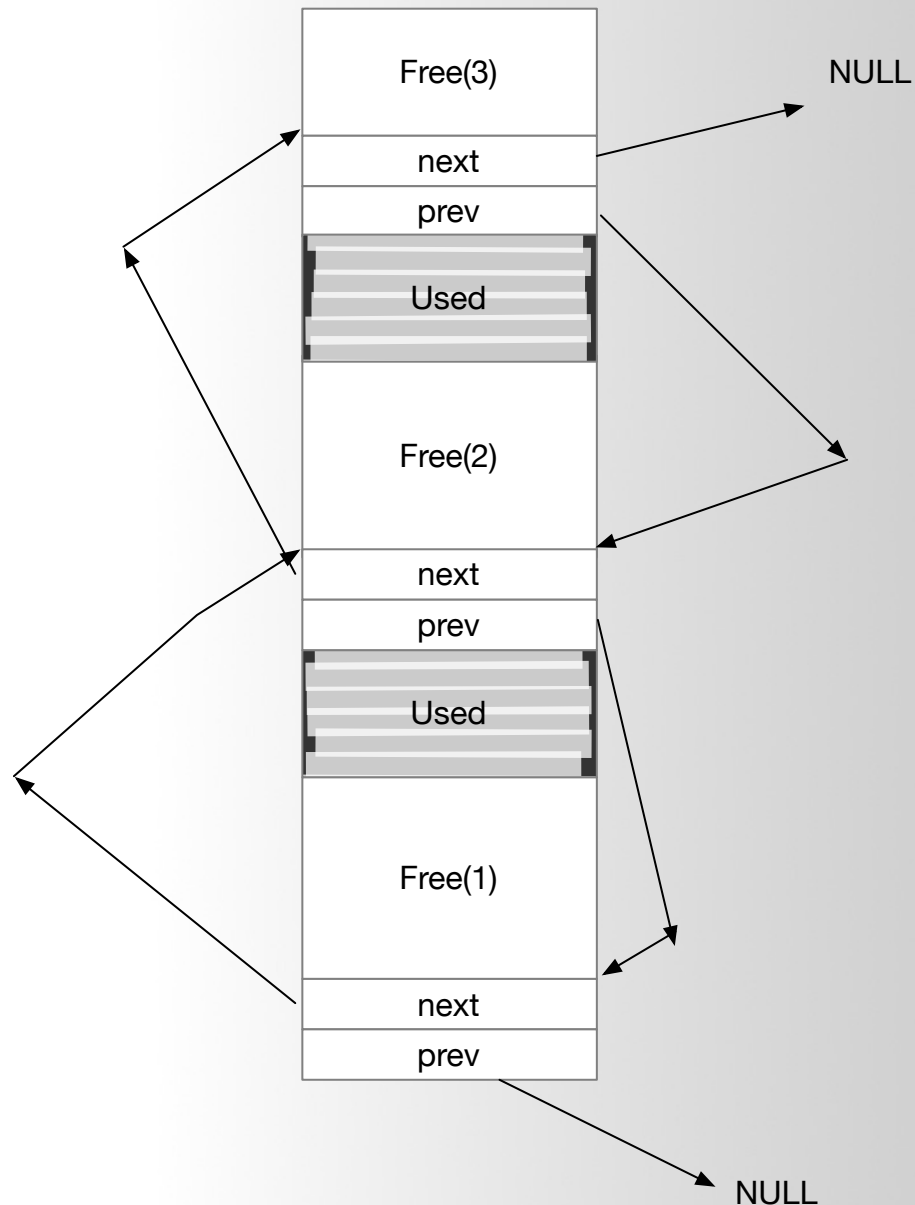
mmap() creates a new mapping in the virtual address space of the calling process. The starting address for the new mapping is specified in addr. The length argument specifies the length of the mapping (which must be greater than 0).

The Heap



Design your own Heap management system


- Linked List



Design your own Heap management system

- bitmap

H: header --> 11
B: Body --> 10
F: Free → 00



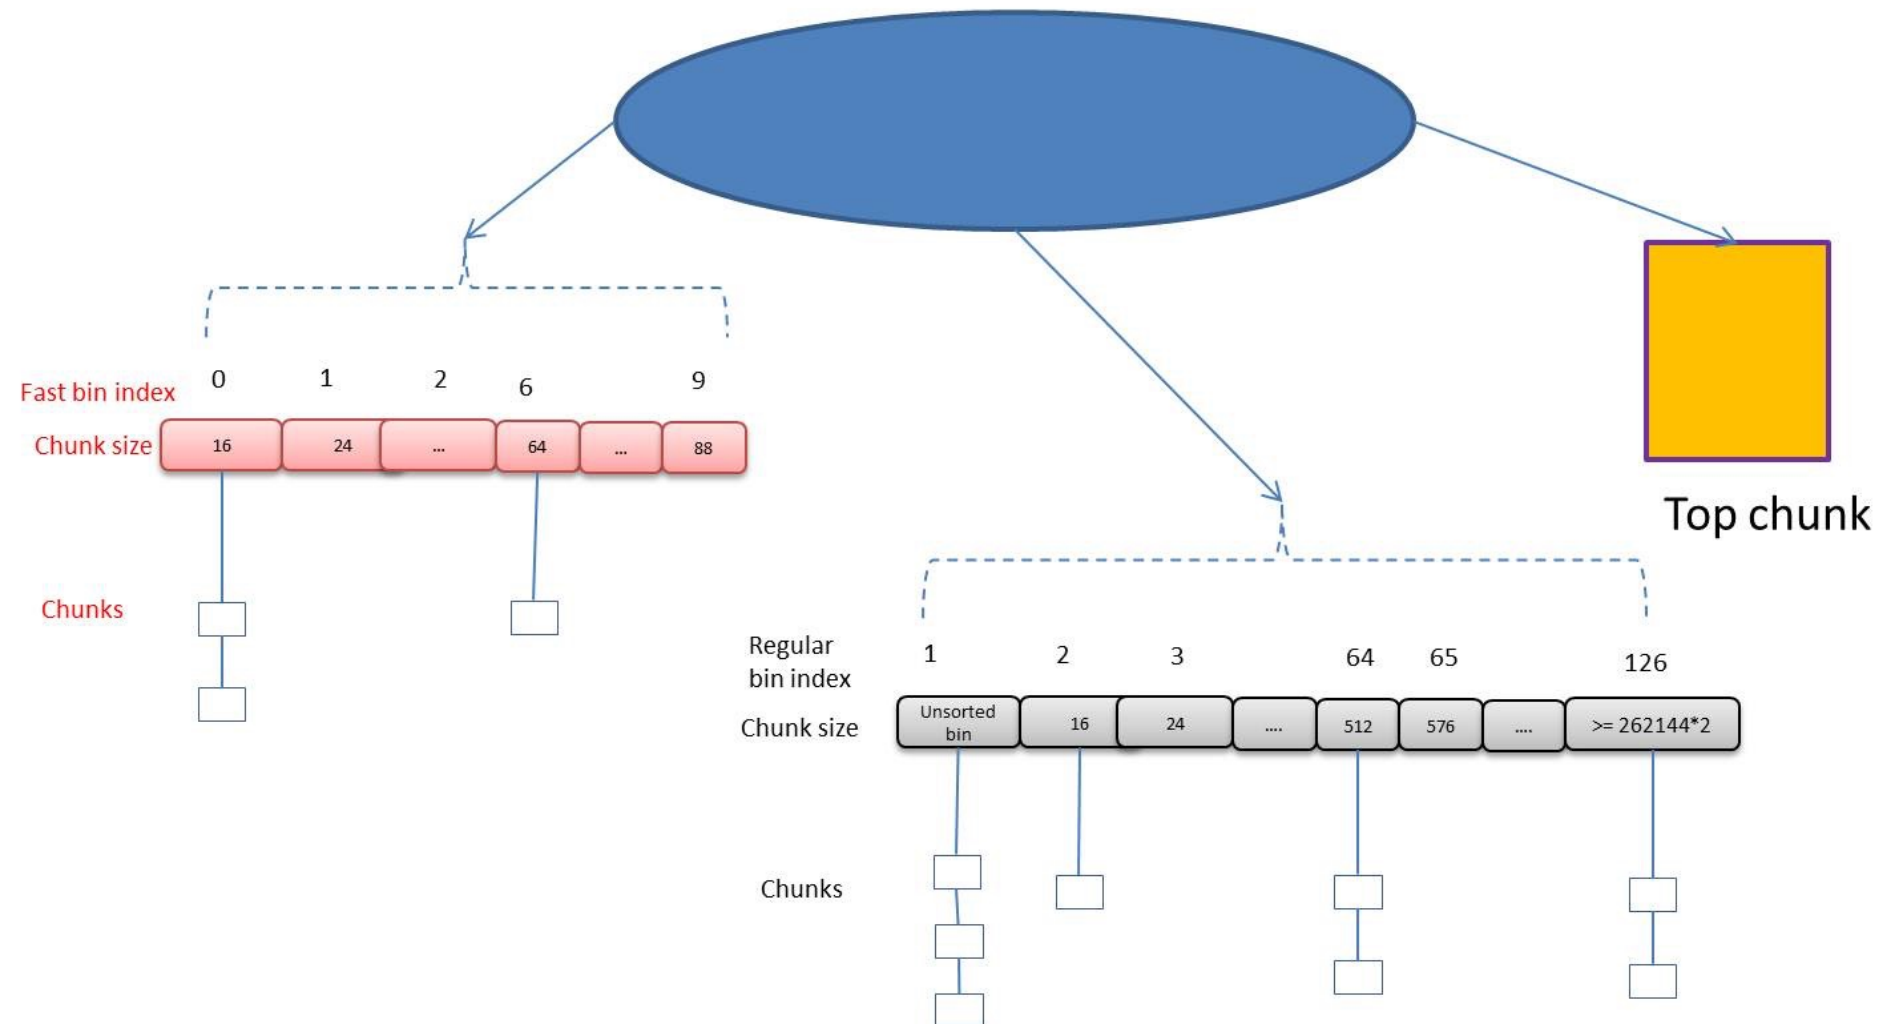
Bitmap representation:

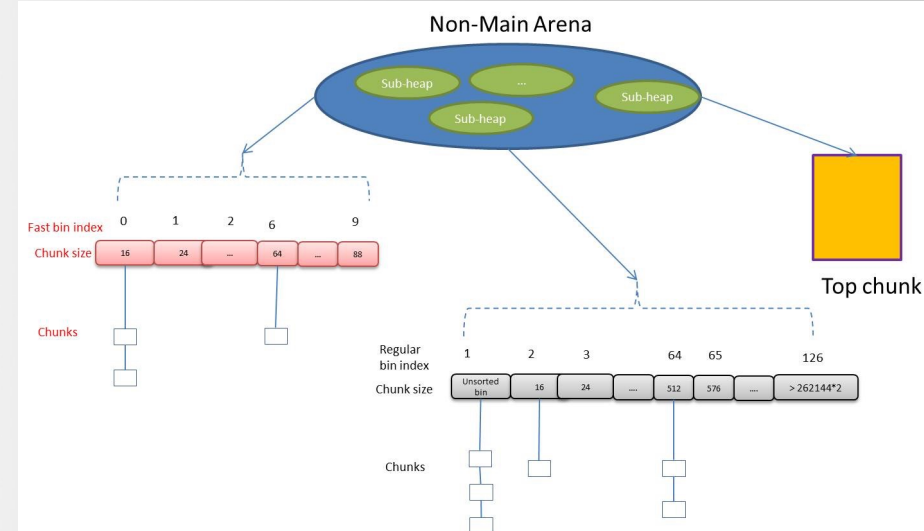
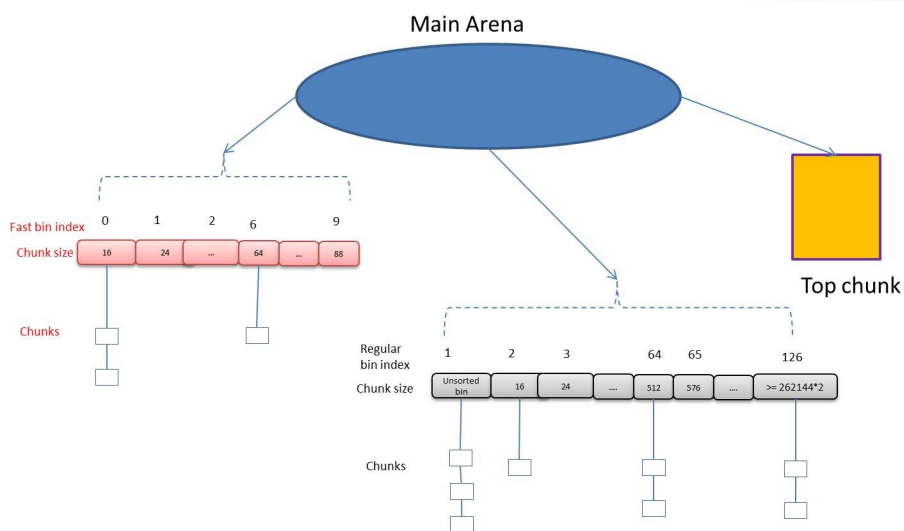
(HIGH) 11 00 00 10 10 10 11 00 00 00 00 00 00 00 10 11 (LOW)

128 byte per chunk
 $1\text{MB} / 128 = 8\text{k}$

```
→ ~ $ git clone git://sourceware.org/git/glibc.git
```

Main Arena





- **Arena: the top level memory management entity.**
- There are **two types** of arenas.
 - Main arena covers the traditional heap area: the space between **start_brk** and **brk** for a process from kernel point of view, only one main arena exists for a process.
 - Non-main arena manages the memory fetched from kernel via **mmap()** system call, there could be 0 to 2*(number of cpu cores) such arenas based on process threads usage.


```
struct malloc_state
{
    /* Serialize access. */
    mutex_t mutex;

    /* Flags (formerly in max_fast). */
    int flags;

#ifdef THREAD_STATS
    /* Statistics for locking. Only used if THREAD_STATS is defined. */
    long stat_lock_direct, stat_lock_loop, stat_lock_wait;
#endif

    /* Fastbins */
    mfastbinptr fastbins[NFASTBINS];

    /* Base of the topmost chunk -- not otherwise kept in a bin */
    mchunkptr top;

    /* The remainder from the most recent split of a small request */
    mchunkptr last_remainder;

    /* Normal bins packed as described above */
    mchunkptr bins[NBINS * 2 - 2];

    /* Bitmap of bins */
    unsigned int binmap[BINMAPSIZE];

    /* Linked list */
    struct malloc_state *next;

    /* Linked list for free arenas. */
    struct malloc_state *next_free;


    /* Memory allocated from the system in this arena. */
    INTERNAL_SIZE_T system_mem;
    INTERNAL_SIZE_T max_system_mem;
};
```

Pa

University

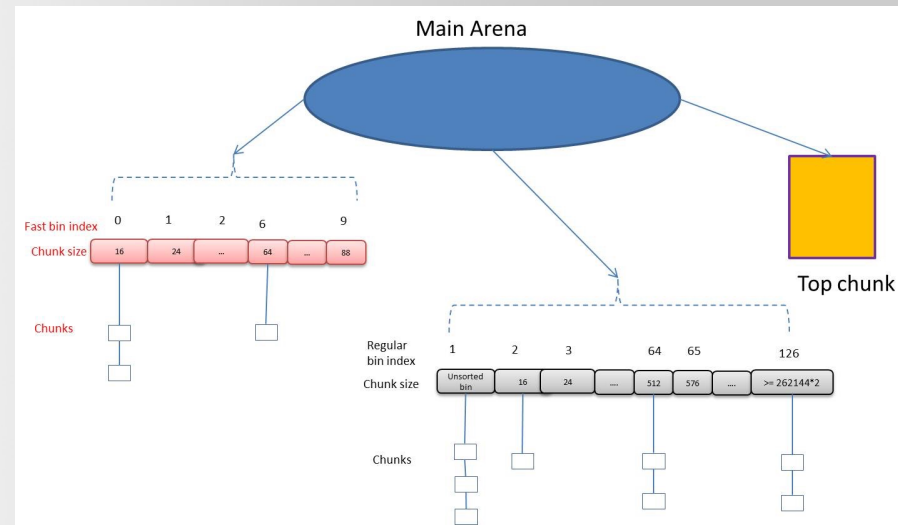
```
→ ~ $ git clone git://sourceware.org/git/glibc.git
```

Main Arena



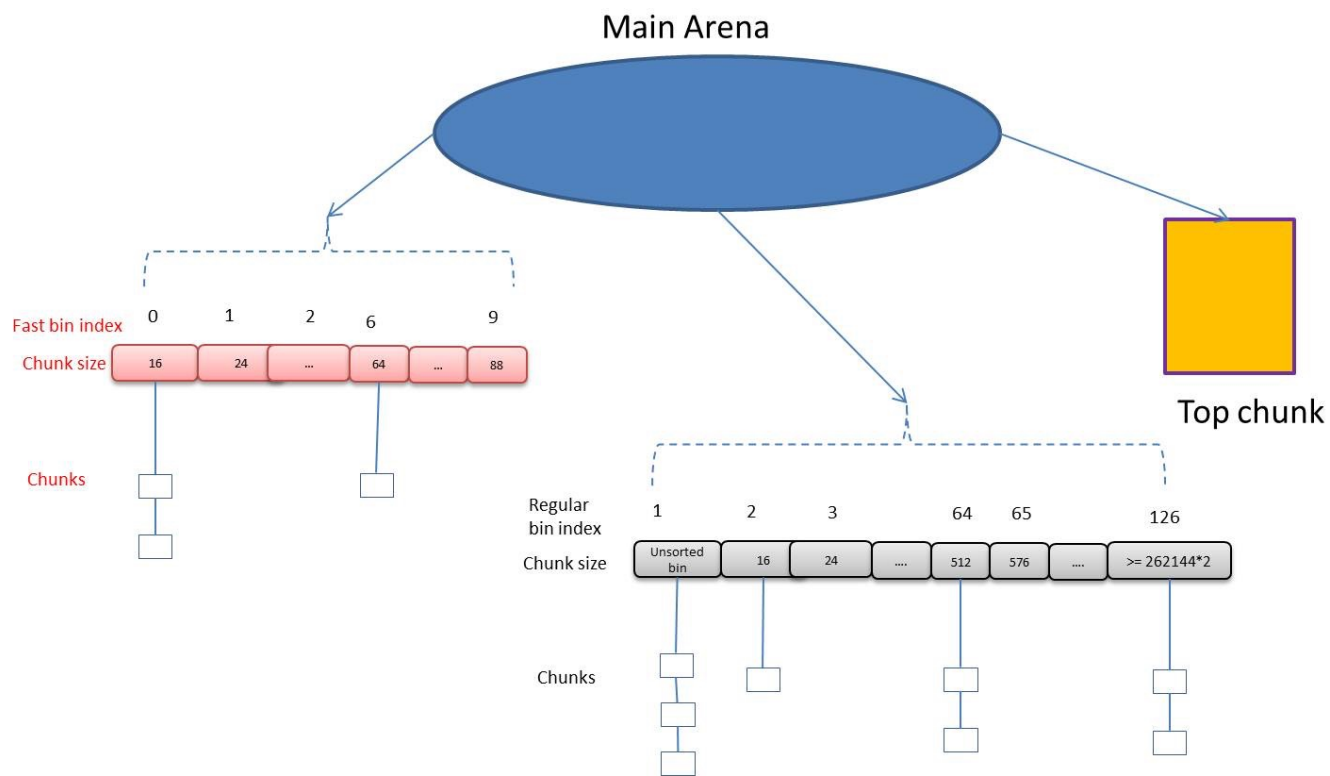
Bins and Chunks

- Internally, the heap manager needs to keep track of freed chunks so that `malloc` can reuse them during allocation requests. In a naive implementation, the heap manager could do this by simply storing all freed chunks together on some enormous linked list. This would work, but it would make [*malloc*](#) slow.
- Since *malloc* is a high-utilization component of most programs, this slowness would have a huge impact on the overall performance of programs running on the system.
- To improve performance, **the heap manager instead maintains a series of lists called “bins”,** which are designed to maximize speed of allocations and frees.



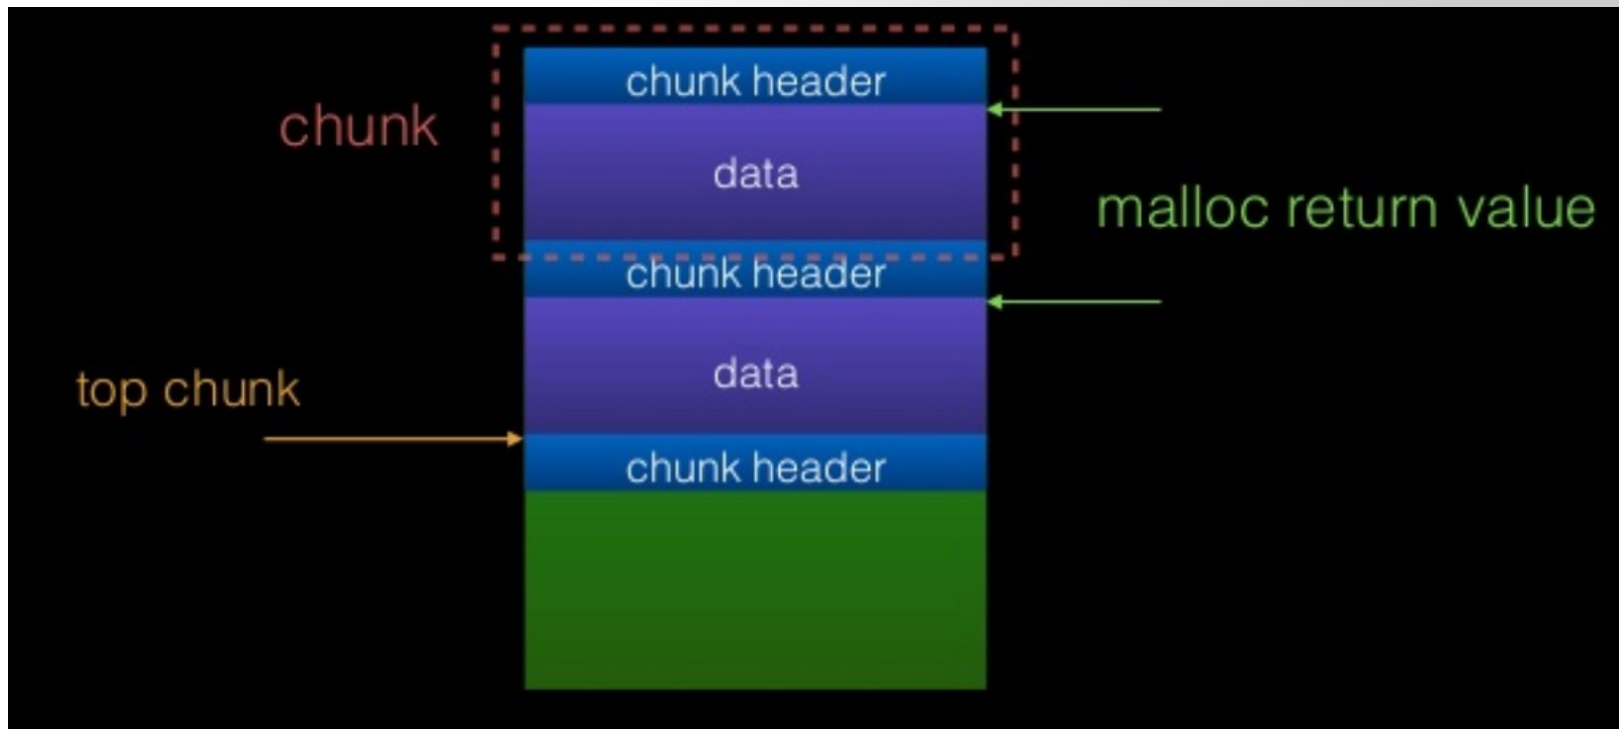
Bins and Chunks

- A bin is a list (doubly or singly linked list) of **free (non-allocated) chunks**. Bins are differentiated based on the size of chunks they contain:
 - Fast bin (16 ~ 80 bytes)
 - Unsorted bin
 - Small bin (< 512 bytes)
 - Large bin (> 512 bytes)



Mechanism of glibc malloc

- Allocated chunk
- Free chunk
- **Top chunk**



Top Chunk

- Top Chunk: Chunk which is at the top border of an arena is called top chunk. It doesn't belong to any bin. Top chunk is used to service user request when there is NO free blocks, in any of the bins. If top chunk size is greater than user requested size top chunk is split into two:
 - User chunk (of user requested size)
 - Remainder chunk (of remaining size)
- The remainder chunk becomes the new top. If top chunk size is lesser than user requested size, top chunk is extended using sbrk (main arena) or mmap (thread arena) syscall.

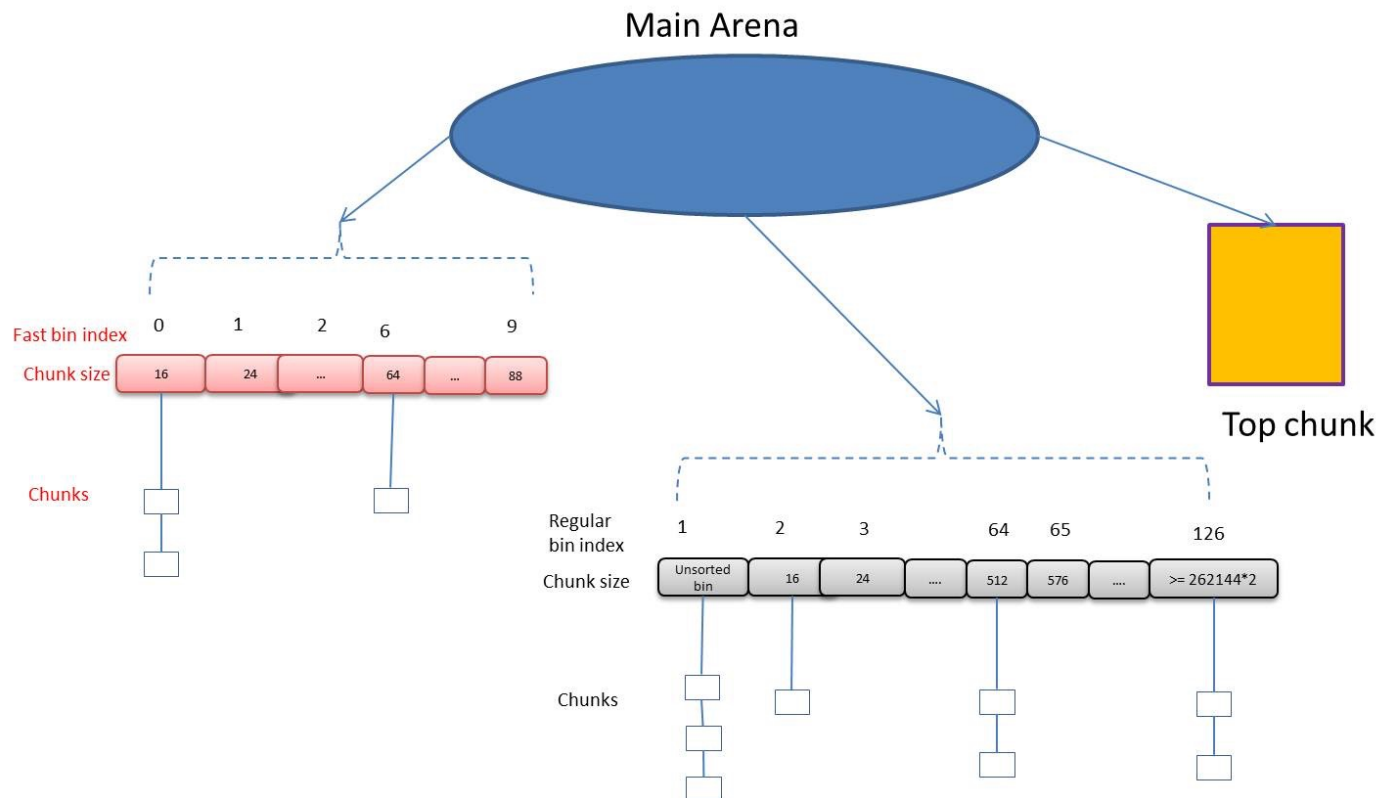
The Malloc Maleficarum (2004)

■ The Malloc Maleficarum

- In late 2001, "Vudo Malloc Tricks" and "Once Upon A free()" defined the exploitation of overflowed dynamic memory chunks on Linux.
- In late 2004, a series of patches to GNU libc malloc implemented over a dozen mandatory integrity assertions, effectively rendering the existing techniques obsolete.
- It is for this reason, a small suggestion of impossibility, that they present the Malloc Maleficarum:
 - The House of Prime
 - The House of Mind
 - The House of Force
 - The House of Lore
 - The House of Spirit
 - The House of Chaos

House of Force

- **House of Force:** In this technique, attacker abuses top chunk size and tricks 'glibc malloc' to service a very large memory request (greater than heap system memory size) using top chunk. Now when a new malloc request is made, GOT entry of free would be overwritten with shellcode address. Hence from now on whenever free is called, shellcode gets executed!!




The Malloc Maleficarum (2004)

■ The Malloc Maleficarum

- In late 2001, "Vudo Malloc Tricks" and "Once Upon A free()" defined the exploitation of overflowed dynamic memory chunks on Linux.
- In late 2004, a series of patches to GNU libc malloc implemented over a dozen mandatory integrity assertions, effectively rendering the existing techniques obsolete.
- It is for this reason, a small suggestion of impossibility, that they present the Malloc Maleficarum:
 - ~~The House of Prime~~
 - ~~The House of Mind~~
 - **The House of Force**
 - ~~The House of Lore~~
 - **The House of Spirit**
 - ~~The House of Chaos~~

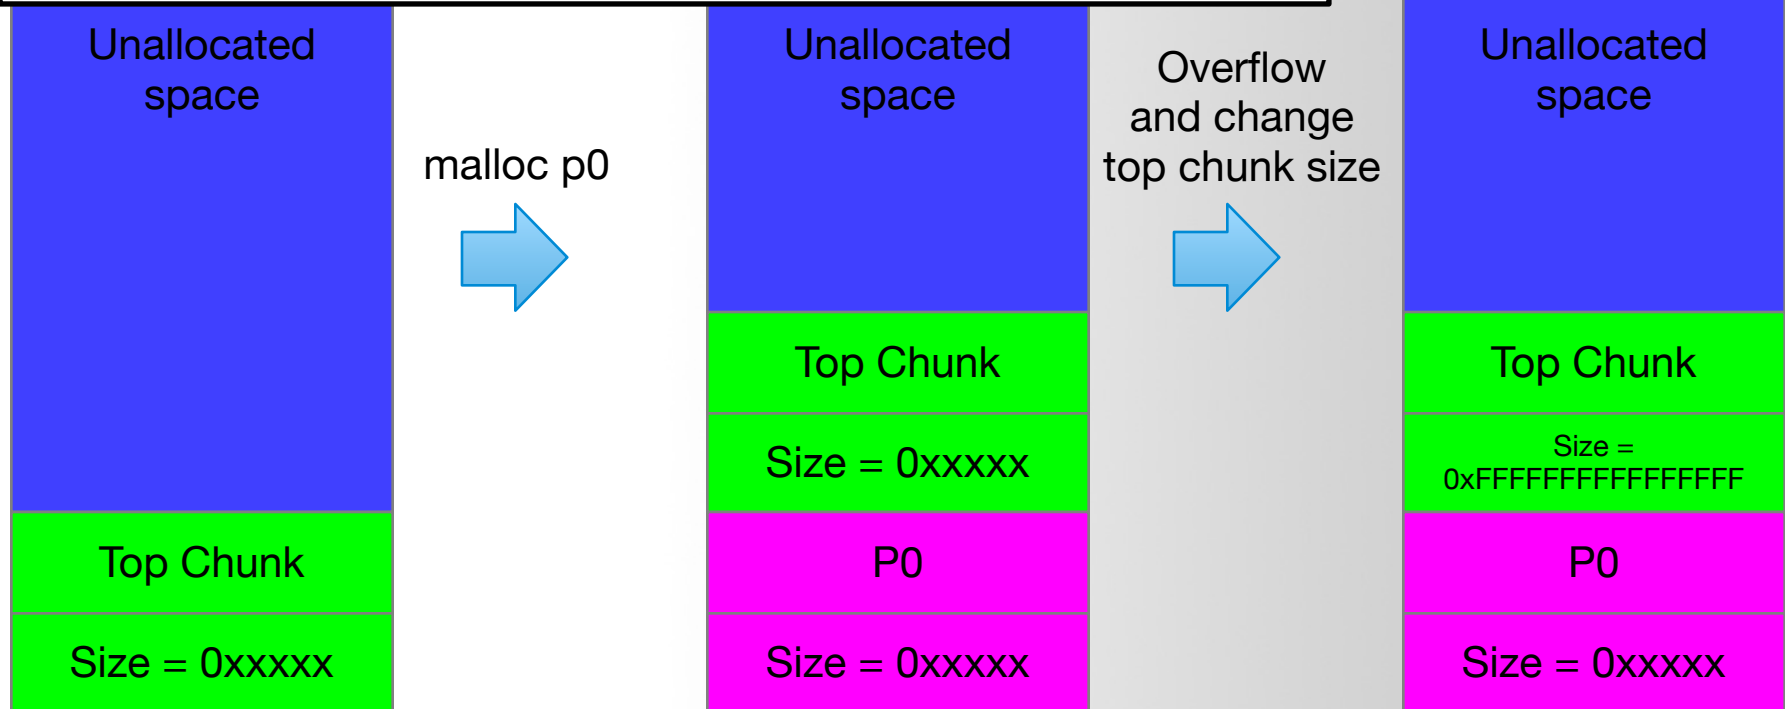
<https://dl.packetstormsecurity.net/papers/attack/MallocMaleficarum.txt>

House of Force



House of Force

- This attack assumes an overflow into the top chunk's header. The size is modified to a very large value (-1 in this example).
- This ensures that all initial requests will be serviced using the top chunk, instead of relying on mmap.
- On a 64 bit system, -1 evaluates to 0xFFFFFFFFFFFFFFFF.
- A chunk with this size can cover the entire memory space of the program.



House of Force

Let us assume that the attacker wishes 'malloc' to return address P. Now, any malloc call with the size of: $\&\text{top_chunk} - P$ will be serviced using the top chunk. Note that P can be after or before the top_chunk.



House of Force

E.g. `top_chunk=0x601200`

`malloc(0xffe00030)`

`0xffe00030 < top_chunk_size`

`0xffe00030+0x601200=0x100401230`

`top_chunk=0x401230`



House of Force

- Prerequisites: Three malloc calls are required to successfully apply house of force as listed below:
 - Malloc 1: Attacker should be able to control the size of top chunk. Hence heap overflow should be possible on this allocated chunk which is physically located previous to top chunk.
 - Malloc 2: Attacker should be able to control the size of this malloc request.
 - Malloc 3: User input should be copied to this allocated chunk.

House of Force - example

```
#!/gcc hof.c -o hof -no-pie
#include<stdio.h>
#include<stdlib.h>

char bss_var[] = "This is a string that we want to overwrite.";

int main()
{
    unsigned long *p1 = malloc(0xf0);
    unsigned long *top = (unsigned long*)( (unsigned char*)p1 + 0xf0);

    //overwrite the size of top chunk
    *(top+1) = -1;

    unsigned long evil_size = (unsigned long)bss_var - sizeof(long)*2 - (unsigned long)top - sizeof(long) * 2;
    malloc(evil_size);

    unsigned long *victim = malloc(0x100);

    printf("The victim pointer: %p\n", victim);

    return 0;
}
```


Exercise: BambooBox

Source Code:

- https://github.com/ctf-wiki/ctf-challenges/blob/master/pwn/heap/house-of-force/hitcontraning_lab11/bamboobox.c

Solution:

https://github.com/ctf-wiki/ctf-challenges/blob/master/pwn/heap/house-of-force/hitcontraning_lab11/exp.py

Q & A