



Heapster

Analyzing the Security of Dynamic Allocators for Monolithic Firmware Images

Fabio Gritti, Fabio Pagani, Lukas Dresel, Ilya Grishchenko, Nilo Redini, Christopher Kruegel, and Giovanni Vigna

University of California, Santa Barbara

Security of Dynamic Allocators for Monolithic Firmware Images

Security of **Dynamic Allocators** for Monolithic Firmware Images

Dynamic Allocators

- Dynamic allocators are algorithms used to manage dynamic memory (i.e., heap memory) of programs.

Dynamic Allocators

- Dynamic allocators are algorithms used to manage dynamic memory (i.e., heap memory) of programs.
- A robust dynamic allocator algorithm is crucial for the performance of any software.

Dynamic Allocators

- Dynamic allocators are algorithms used to manage dynamic memory (i.e., heap memory) of programs.
- A robust dynamic allocator algorithm is crucial for the performance of any software.
- Composed by at least 2 primitives: malloc and free (i.e., **H**Heap **M**anagement **L**ibrary)

Dynamic Allocators

- Dynamic allocators are routinely abused as building blocks for complex exploits

(Pwn2Own) Zoom Heap based Buffer Overflow Remote Code Execution Vulnerability

ZDI-21-971

ZDI-CAN-13587

Dynamic Allocators

- Dynamic allocators are routinely abused as building blocks for complex exploits

(Pwn2Own) Zoom Heap based Buffer Overflow Remote Code Execution Vulnerability

ZDI-21-971

ZDI-CAN-13587



Zero Day Initiative
@thezdi

Our 1st [#Pwn2Own](#) [#AfterDark](#) concludes with the Mofoffensive Research Team combining a heap overflow and a stack-based buffer overflow to gain code execution on the LAN interface of the NETGEAR R6700 router. Their efforts earn \$5,000 and 1 Master of Pwn point. [#P2OAustin](#)

Dynamic Allocators

- Dynamic allocators are routinely abused as building blocks for complex exploits

(Pwn2Own) Zoom Heap based Buffer Overflow Remote Code Execution Vulnerability

ZDI-21-971

ZDI-CAN-13587



Zero Day Initiative
@thezdi

Our 1st [#Pwn2Own](#) [#AfterDark](#) concludes with the [Mofoffensive Research Team](#) combining a heap overflow and a stack-based buffer overflow to gain code execution on the LAN interface of the NETGEAR R6700 router. Their efforts earn \$5,000 and 1 Master of Pwn point. [#P2OAustin](#)

SECURITY RESEARCH

CVE-2021-43267: Remote Linux Kernel Heap Overflow | TIPC Module Allows Arbitrary Code Execution

MAX VAN AMERONGEN / NOVEMBER 4, 2021

Dynamic Allocators

- Dynamic allocators are routinely abused as building blocks for complex exploits

Current research is focused on allocators for “classic” systems



Our 1st [#Pwn2Own](#) [#AfterDark](#) concludes with the Mofoffensive Research Team combining a heap overflow and a stack-based buffer overflow to gain code execution on the LAN interface of the NETGEAR R6700 router. Their efforts earn \$5,000 and 1 Master of Pwn point. [#P2OAustin](#)

SECURITY RESEARCH

CVE-2021-43267: Remote Linux Kernel Heap Overflow | TIPC Module Allows Arbitrary Code Execution

MAX VAN AMERONGEN / NOVEMBER 4, 2021


Security of Dynamic Allocators for **Monolithic Firmware** Images

Monolithic Firmware Images

- Firmware images without an OS-abstraction

Monolithic Firmware Images


- Firmware images without an OS-abstraction



[Linux-based firmware image]

Monolithic Firmware Images

- Firmware images without an OS-abstraction



Monolithic Firmware Images

Empower a huge amount of diverse IoT devices

Monolithic Firmware Images

Empower a huge amount of diverse IoT devices



HUGE attack surface with a spectrum of different threat scenarios

Monolithic Firmware Images



Hack my toaster



Monolithic Firmware Images



Remote arbitrary write over
pacemaker...



Monolithic Firmware Images

- Very hard target for both static and dynamic analysis:
 - Binary ONLY
 - NO symbols
 - NO hardware
 - Scalability of re-hosting remains a challenge





Dynamic
allocators

Monolithic
firmware



Dynamic
allocators

Monolithic
firmware



Understand if the f/w binary is using a dynamic allocator?




Understand if the f/w binary is using a dynamic allocator?




Check if the f/w allocator is robust against attacks?

Research Scope

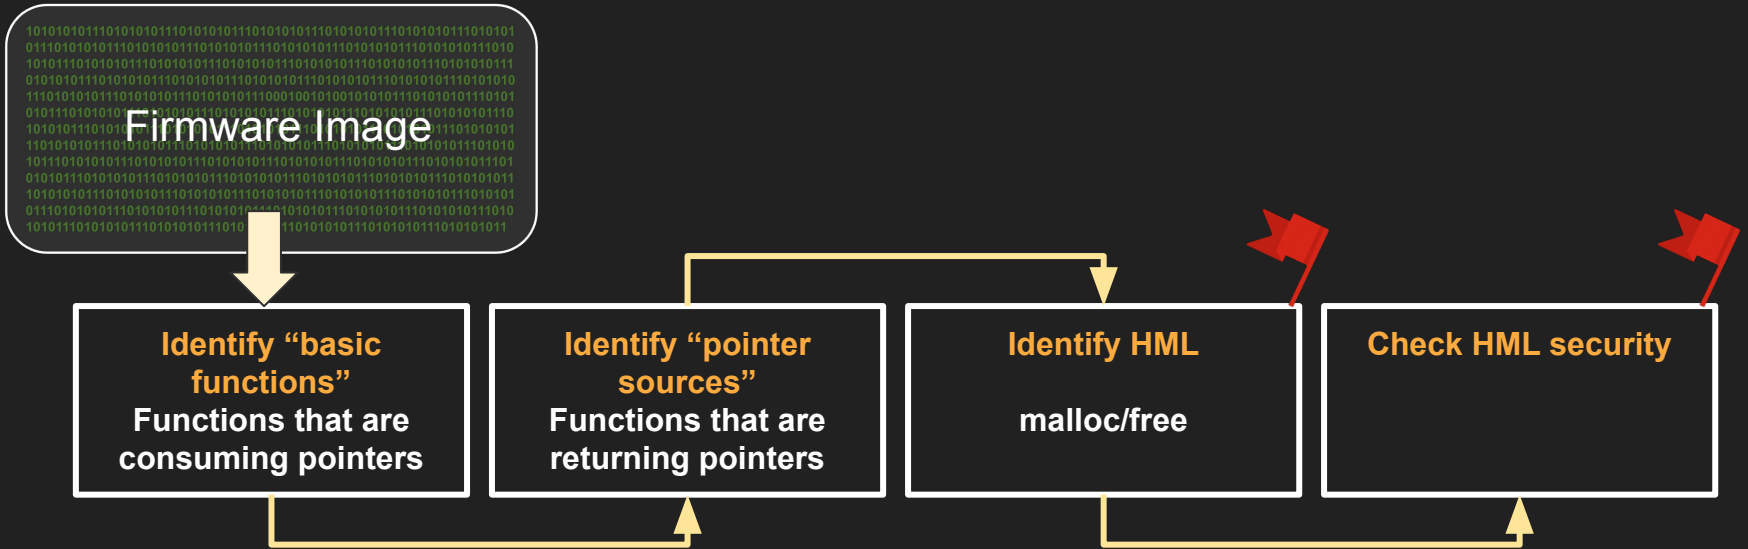
Monolithic Firmware Image




Research Scope

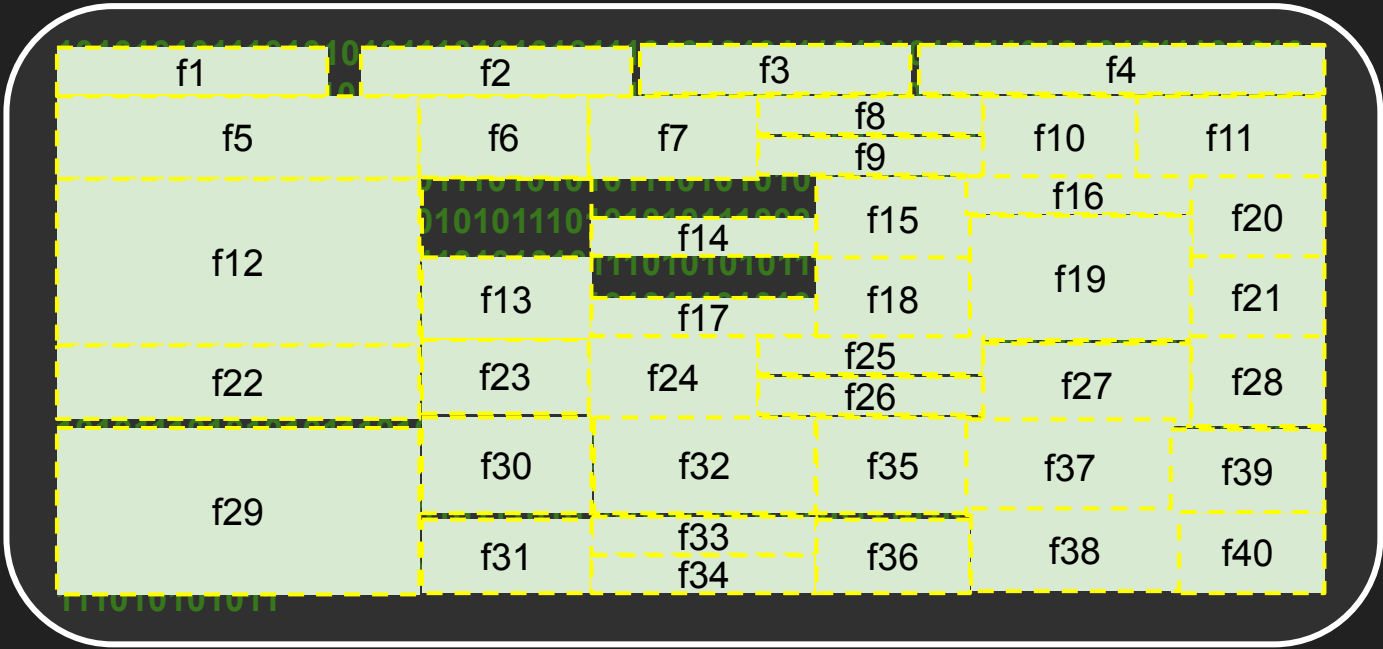



Heapster




10101010111010101011101010101110101010111010101011101010101110101011101010
1011101010101110101010111010101011101010101110101010111010101011101010101110
1010101110101010111010101011101010101110101010111010101011101010101110101010
1110101010111010101011101010101110101010111010101011101010101110101010111010
101011101010101110101010111010101011100010010100101010111010101011
10101010111010101011101010101110101010111010101011101010101110101011101010
1011101010101110101010111010101011101010101110101010111010101011101010101110
1010101110101010111010101011101010101110101010111010101011101010101110101010
1110101010111010101011101010101110101010111010101011101010101110101010111010
1010111010101011101010101110101010111010101011101010101110101010111010101011
1010101011101010101110101010111010101011101010101110101010111010101011101010
1011101010101110101010111010101011101010101110101010111010101011101010101110
1010101110101010111010101011101010101110101010111010101011101010101110101010
111010101011







~100/1000 functions





Approach



Memory allocators generate pointers

Approach



Memory allocators generate pointers



Pointers are eventually used to perform memory operations



Who uses memory pointers?

Identify Basic Functions

- These functions use pointers
 - memcpy(**addr1**,**addr2**,size)
 - memset(**addr1**,c,size)
 - memcmp(**addr1**,**addr2**,size)
- Simple to identify (i.e., “basic functions”)

Identify Basic Functions

Memcpy?

```
f8(X,Y,Z){  
    [ ... CODE ... ]  
}
```

X → b'wxyz\0wxyz'

Y → b'asdf\0asdf'

Z = 9

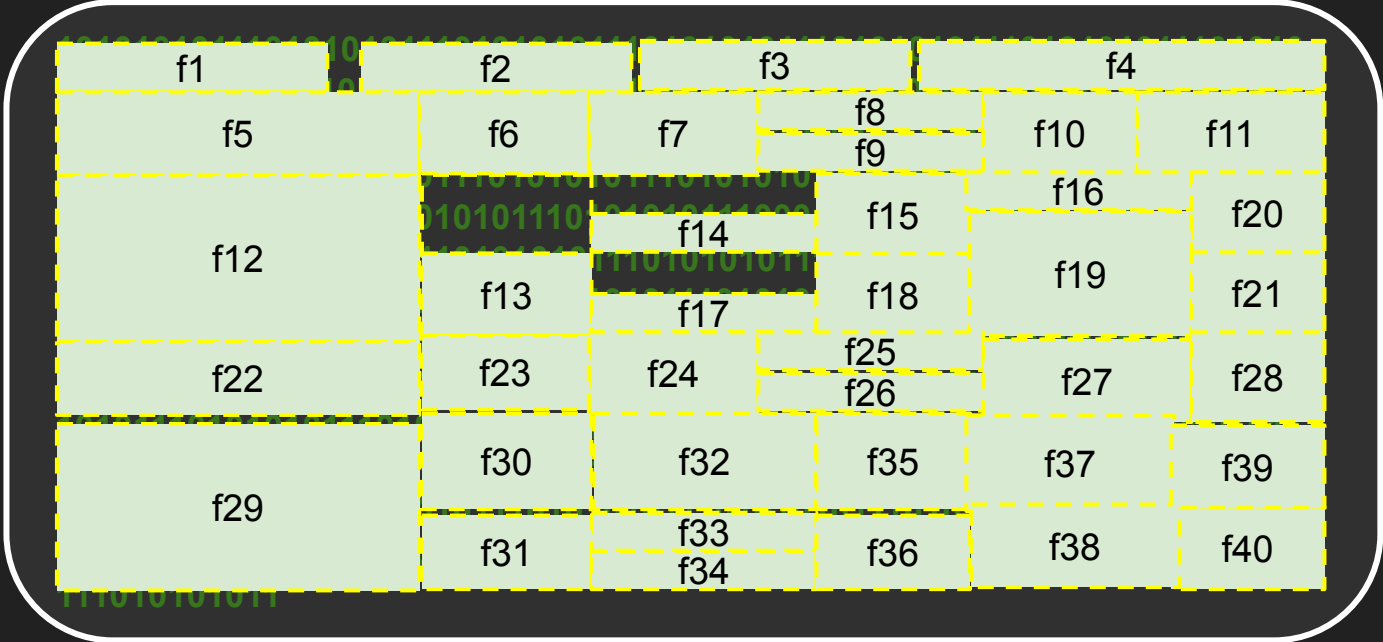
Identify Basic Functions

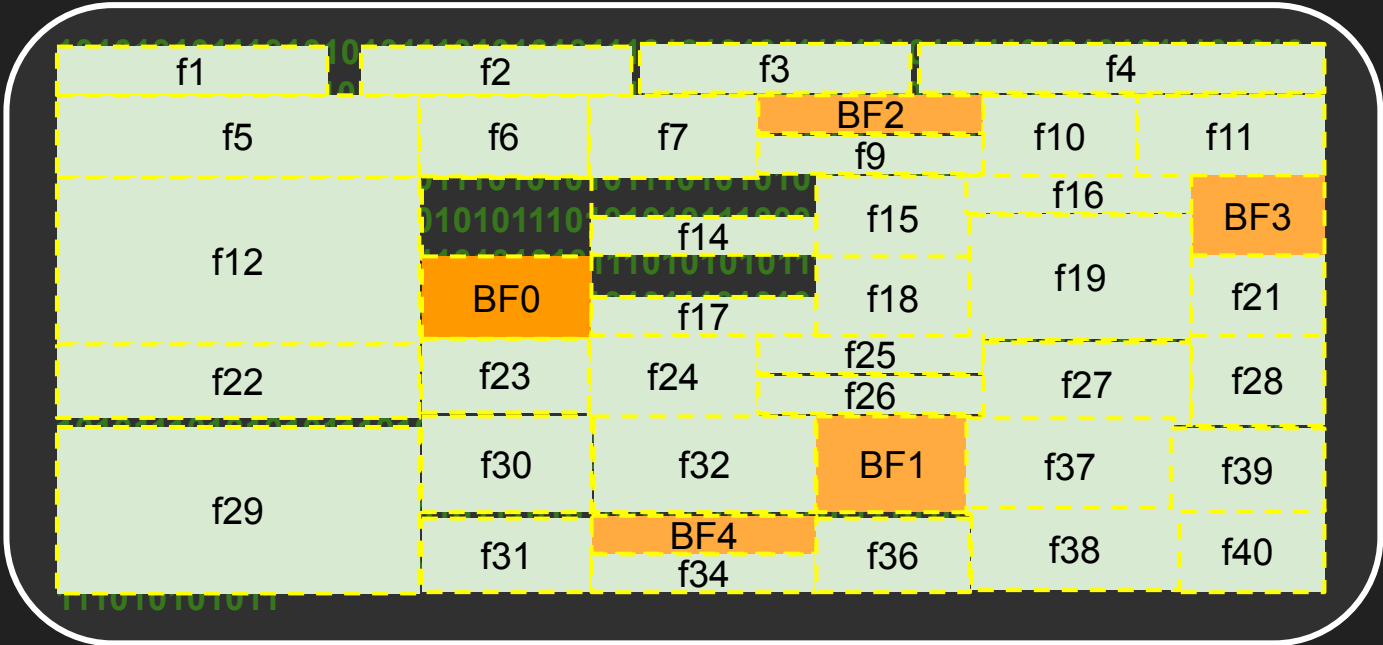
Memcpy?

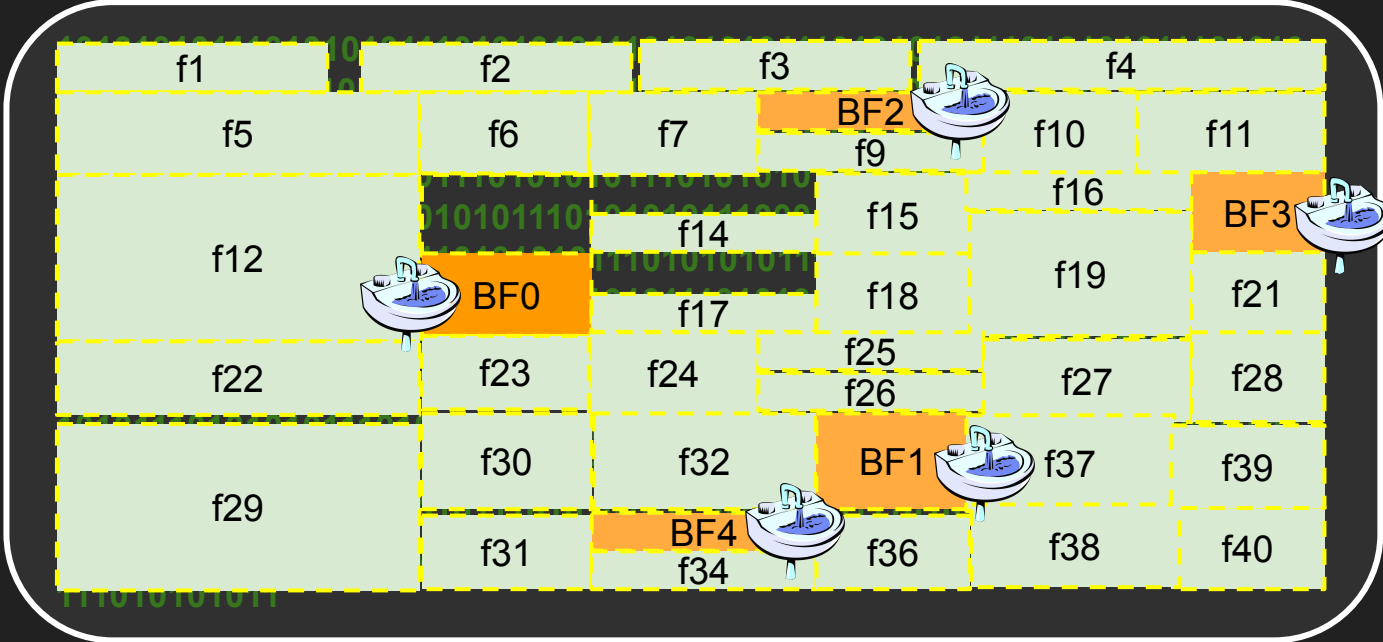
```
f8(X,Y,Z){  
    [ ... CODE ... ]  
}
```

```
X → b'wxyz\0wxyz'  
Y → b'asdf\0asdf'  
Z = 9
```

→ Does buffer at **X** contain exactly 9 bytes `b'asdf\0asdf'`?
→ Is buffer at **Y** unchanged?











Where are the sources?




Identify Pointer Sources

- Functions that provide arguments to the basic functions



```
v25 = (char *) f19(x);  
v26 = v25;  
if(v25){  
    v25[28] = v27;  
    v25[29] = 1;  
    memcpy(v25 + 12, v21 + 2, 16)  
}
```



Identify Pointer Sources



```
v25 = (char *) f19(x);  
v26 = v25;  
if(v25){  
    v25[28] = v27;  
    v25[29] = 1;  
    memcpy(v25 + 12, v21 + 2, 16)  
}
```



- Use static taint engine (Reaching Definition)
- Collect *all* the functions that are returning values that define basic functions' arguments

Identify Pointer Sources



```
v25 = (char *) f19(x);  
v26 = v25;  
if(v25){  
    v25[28] = v27;  
    v25[29] = 1;  
    memcpy(v25 + 12, v21 + 2, 16)  
}
```



```
v25 defs:  
{<f19_RETURN>}
```

- Use static taint engine (Reaching Definition)
- Collect *all* the functions that are returning values that define basic functions' arguments

Identify Pointer Sources



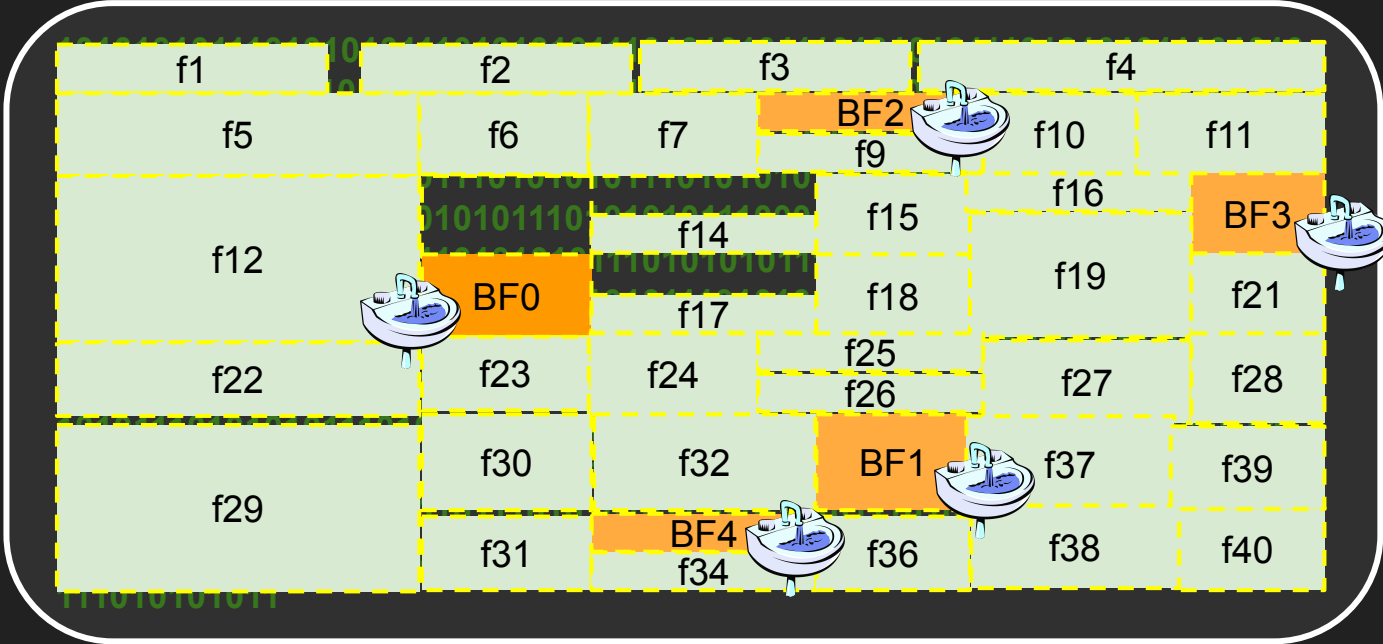
```
v25 = (char *) PS(x);  
v26 = v25;  
if(v25){  
    v25[28] = v27;  
    v25[29] = 1;  
    memcpy(v25 + 12, v21 + 2, 16)
```

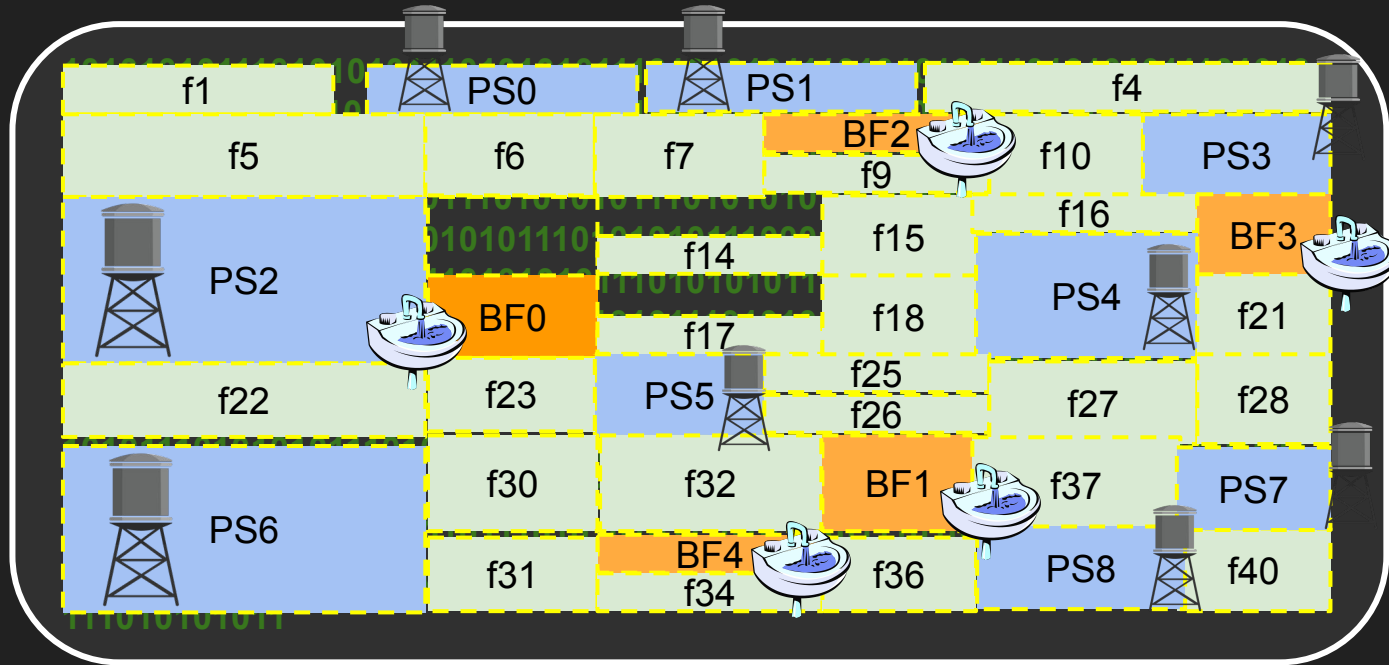



v25 defs:
{<f19_RETURN>}

f19 is a Pointer
Source!












Memory allocator is a pointer source





Many pointer sources, how to identify malloc?



Identify Malloc



Malloc returns pointers inside heap region

(CortexM: 0x20000000 -> 0x40000000)

Identify Malloc



Malloc returns pointers inside heap region

(CortexM: 0x20000000 -> 0x40000000)



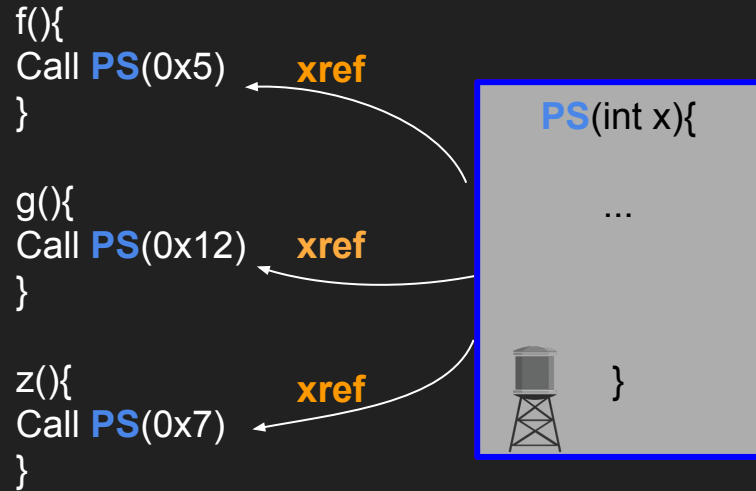
Malloc returns different addresses to subsequent invocations

(Serve every request with a different memory block)

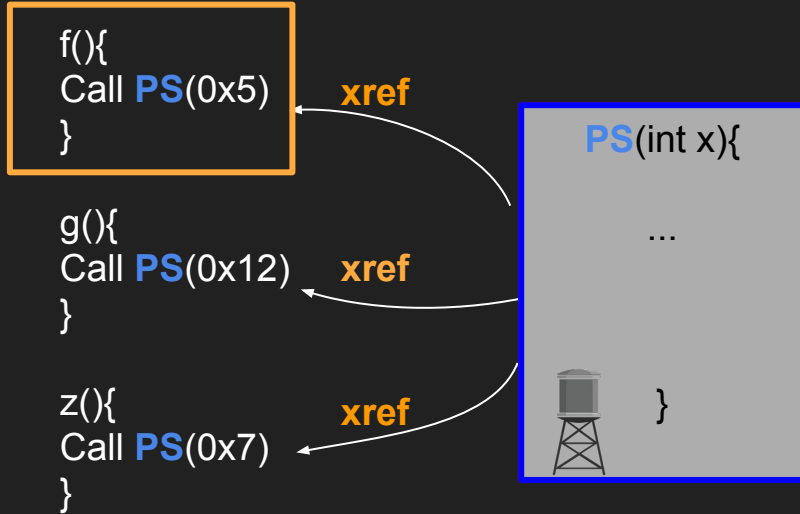
Identify Malloc



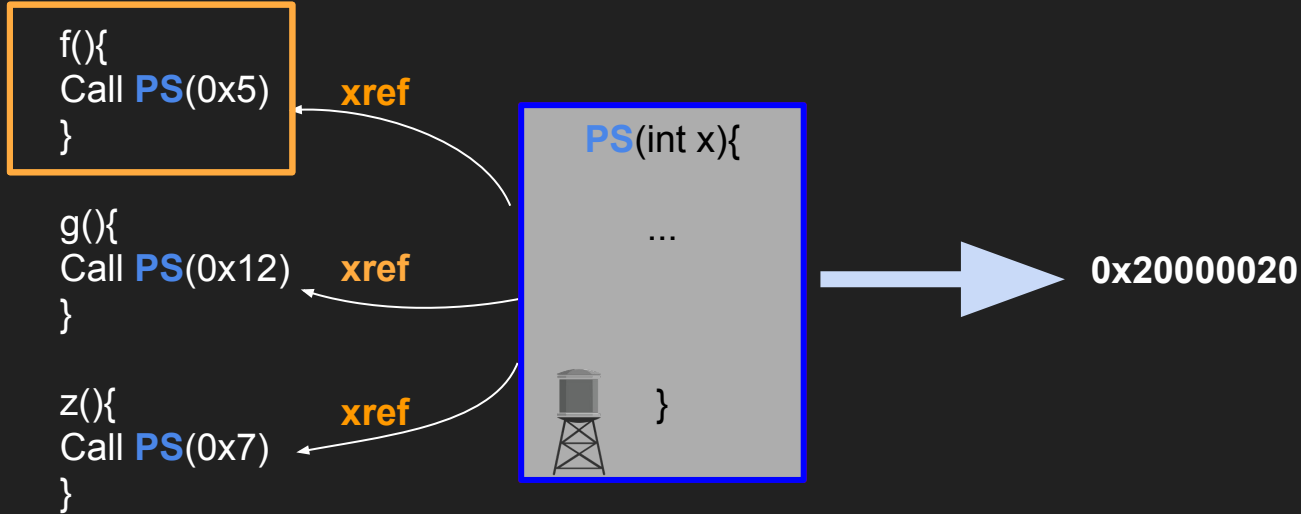
Identify Malloc



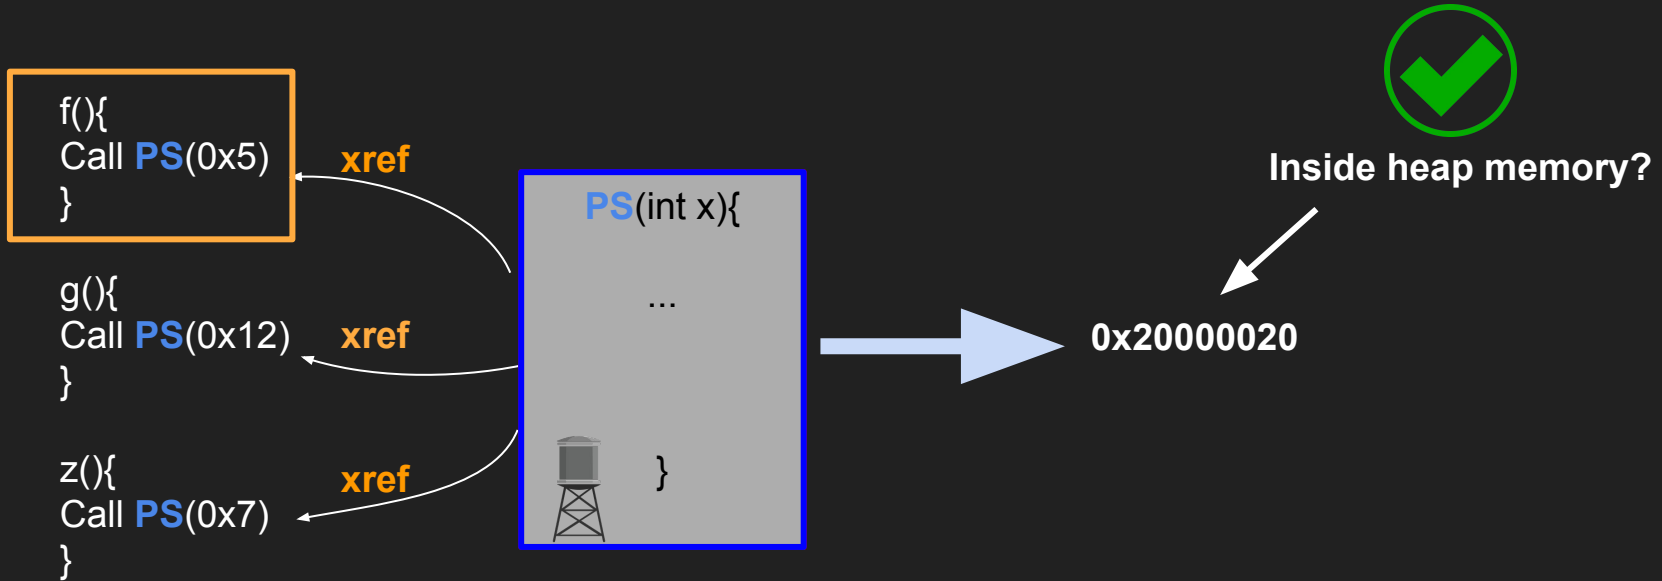
Identify Malloc



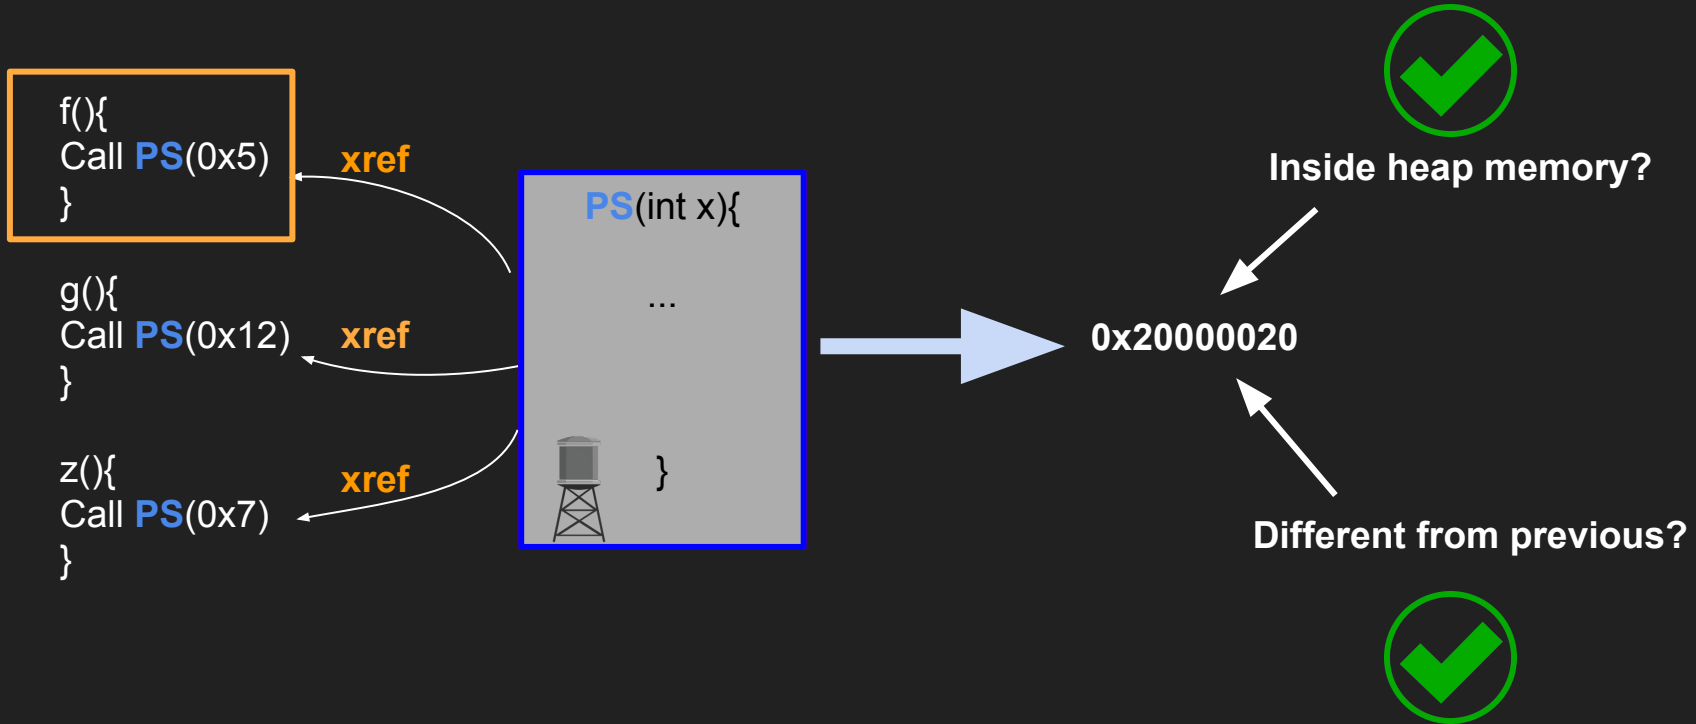
Identify Malloc



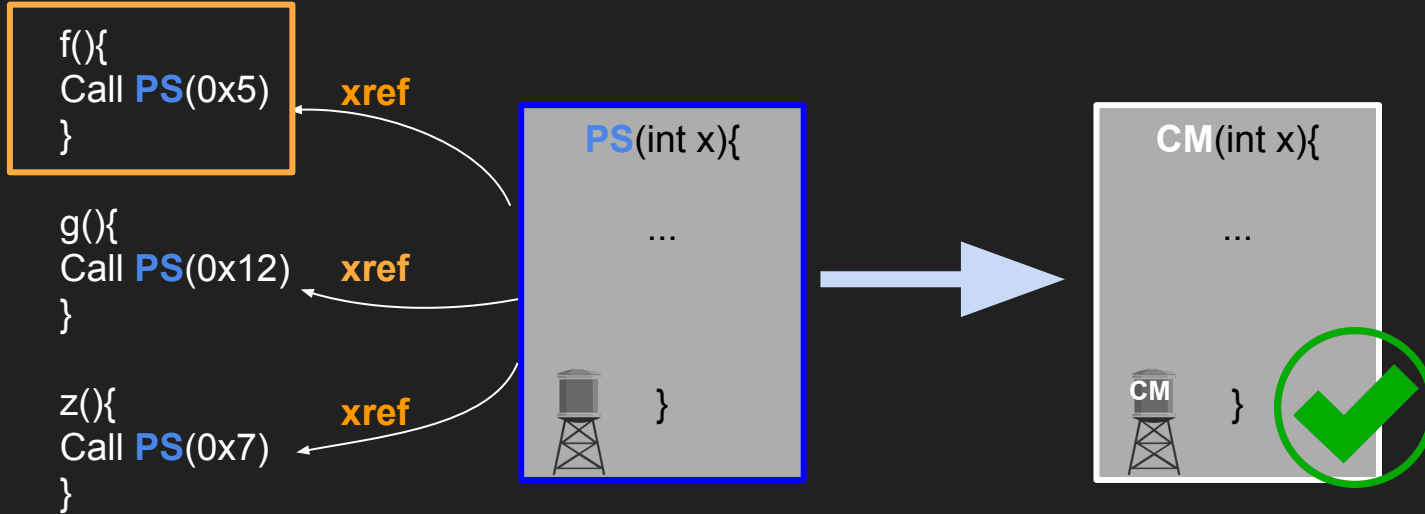
Identify Malloc



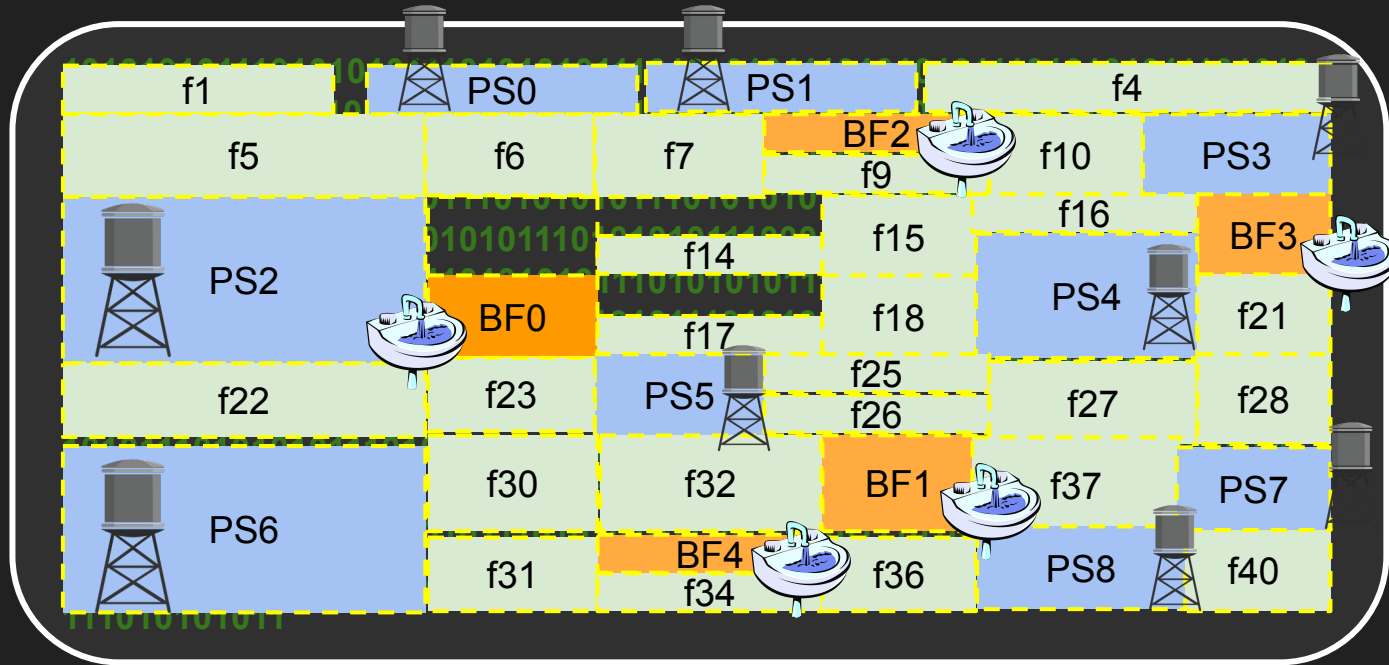
Identify Malloc

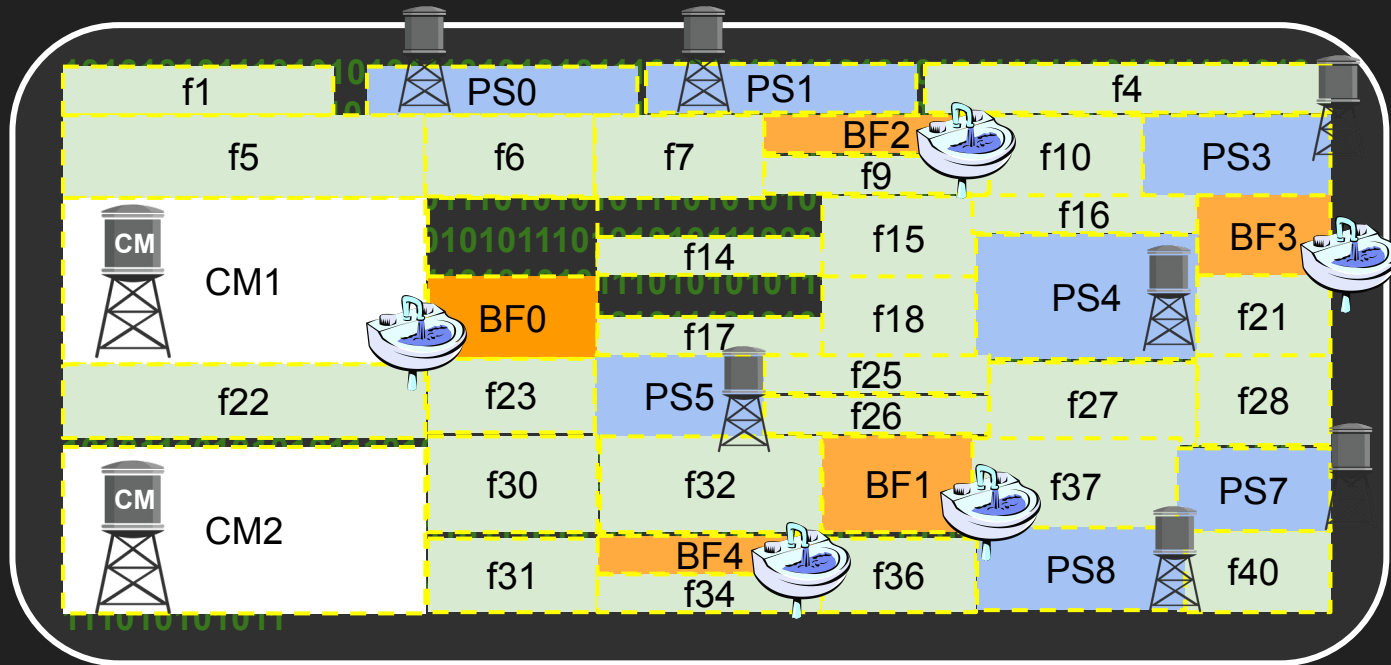


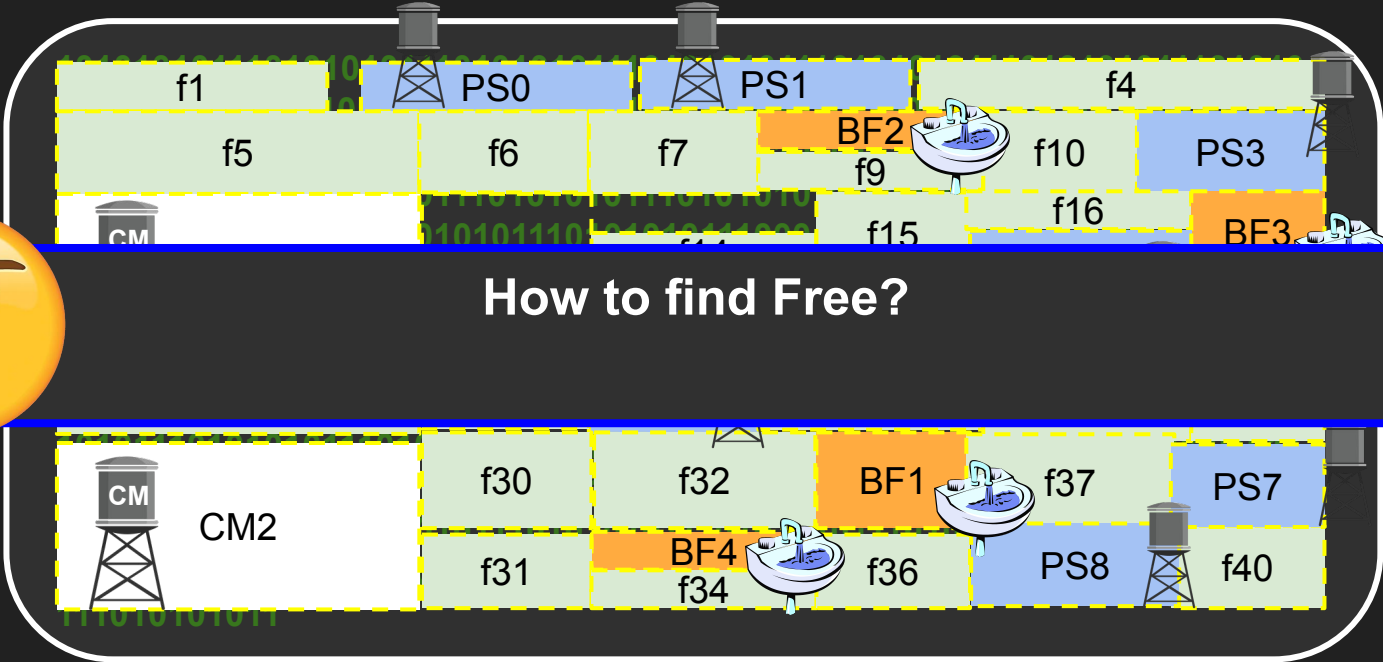
Identify Malloc



PS is a candidate malloc!







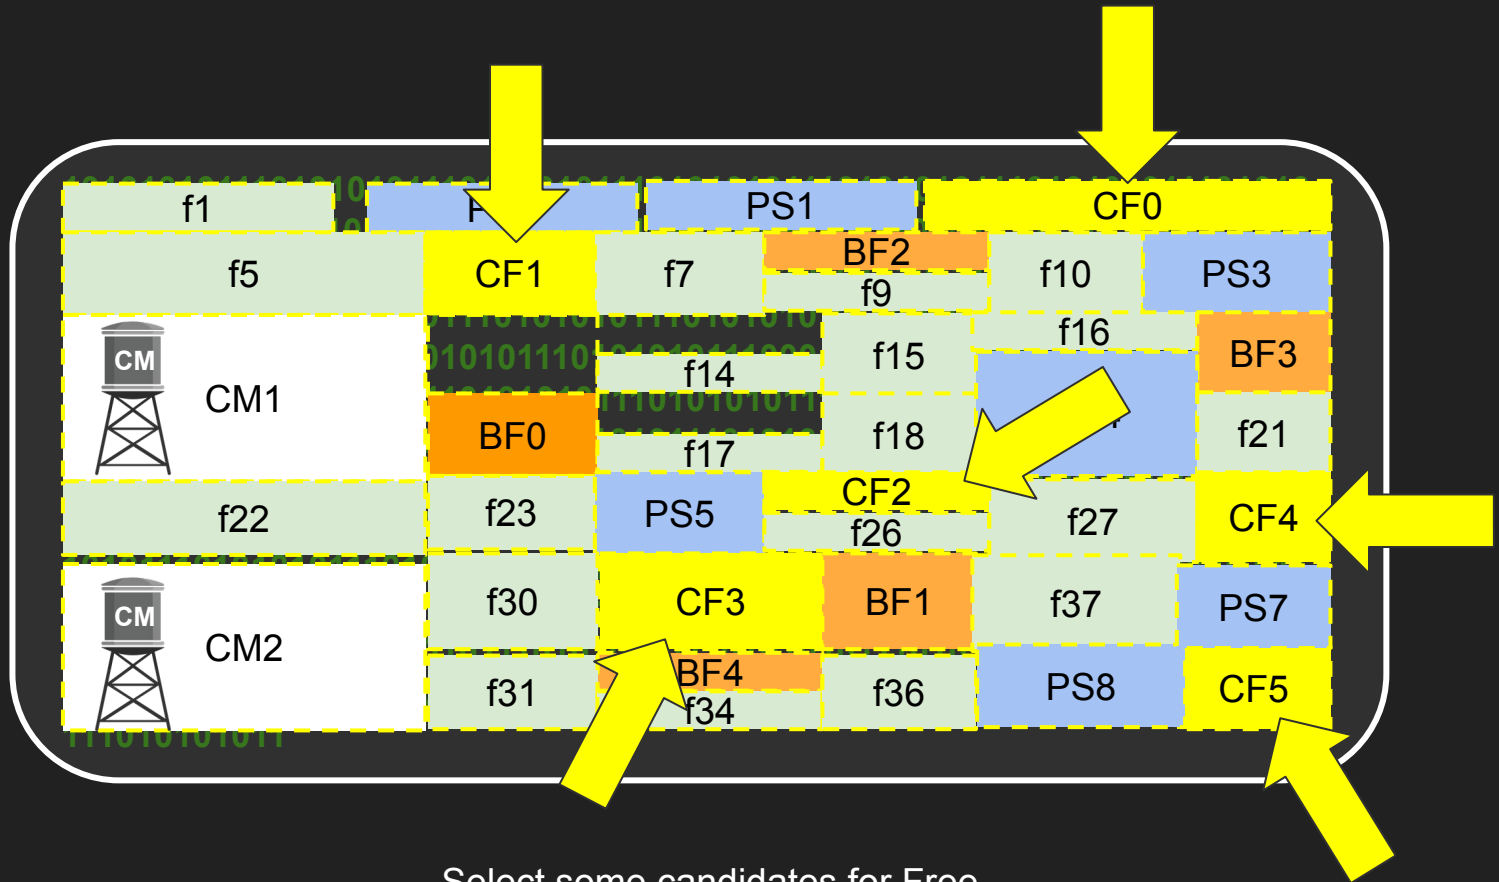
How to find Free?





Identify Free



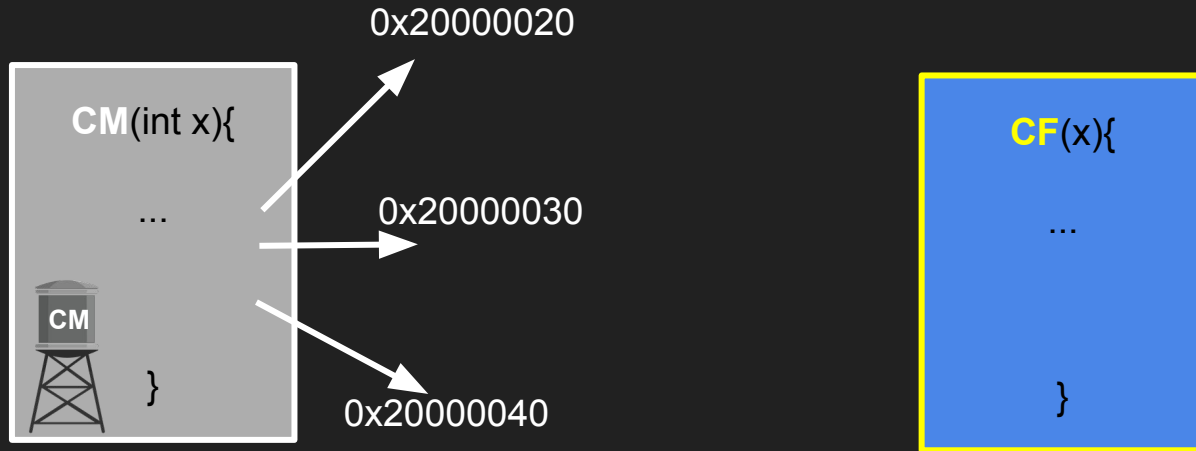
Malloc and Free are LIFO
(malloc returns the last freed pointer)



Identify Free

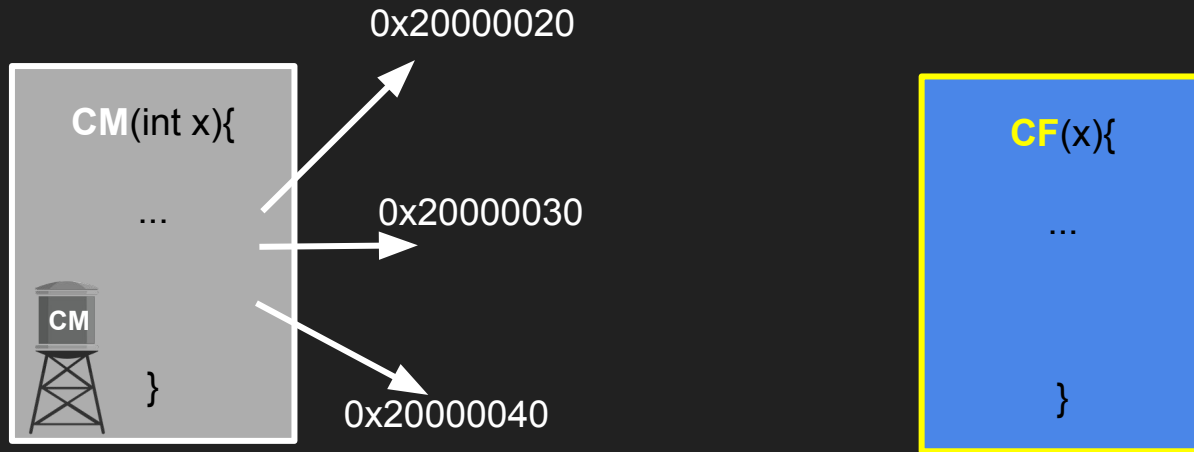


Identify Free




Call **CM** an X amount of times

Identify Free



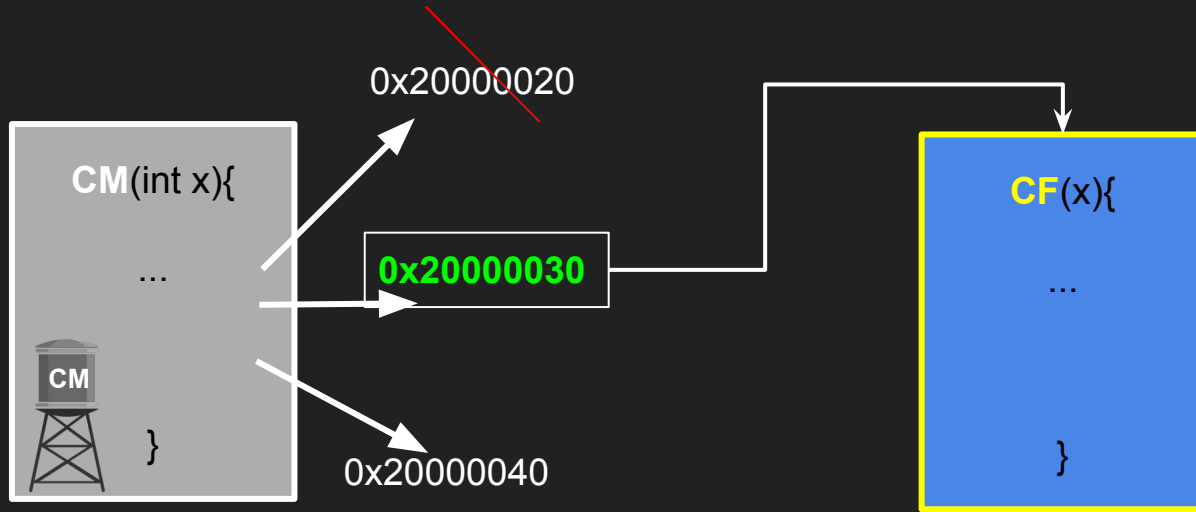
Allocated addresses= [0x20000020, 0x20000030, 0x20000040]

Identify Free




Allocated addresses= [0x20000020, 0x20000030, 0x20000040]

Identify Free




Allocated addresses= [0x20000020, 0x20000030, 0x20000040]

Identify Free



Allocated addresses= [0x20000020, 0x20000030, 0x20000040]


Identify Free



Call **CM** again

Allocated addresses= [**0x20000020**, 0x20000030, 0x20000040]


Identify Free

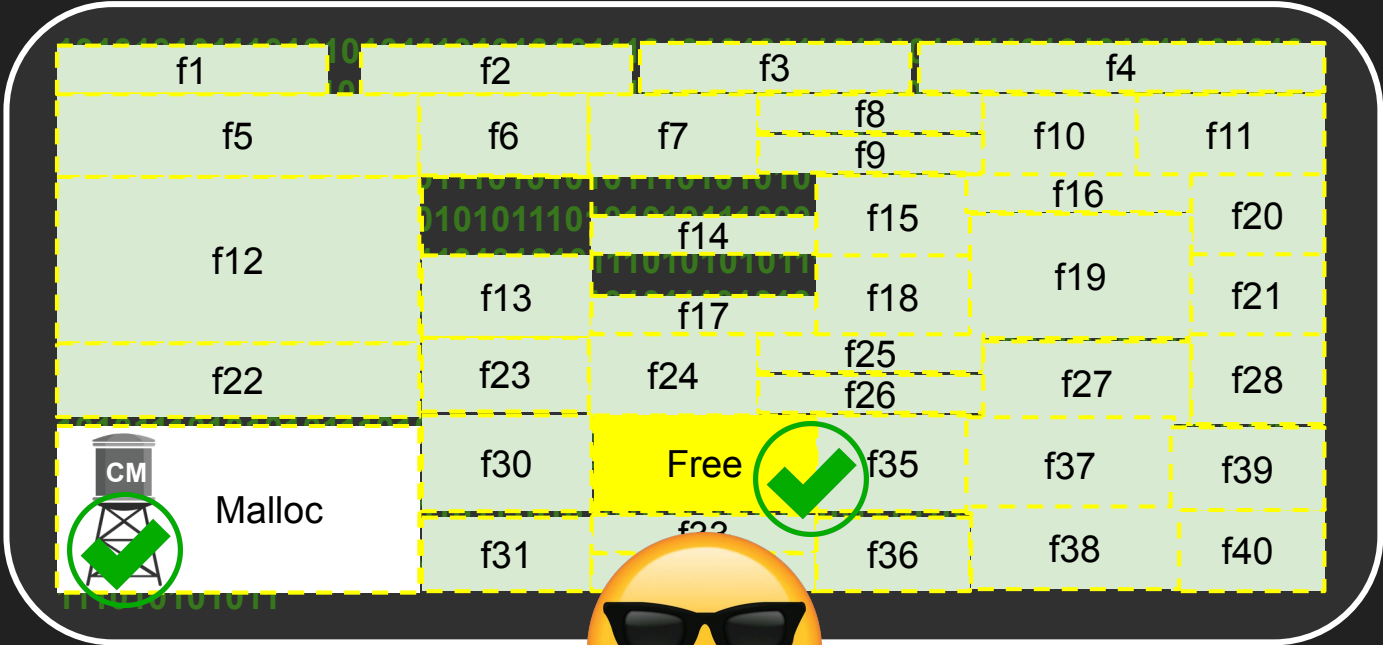


Identify Free

Heap Management Library (i.e., HML)







Evaluation

20 monolithic firmware images (ground truth)

P²IM: Scalable and Hardware-independent Firmware Testing via Automatic Peripheral Interface Modeling

Bo Feng *Northeastern University* Alejandro Mera *Northeastern University* Long Lu *Northeastern University*

HALucinator: Firmware Re-hosting Through Abstraction Layer Emulation

Abraham A. Clements^{*1}, Eric Gustafson^{*1,2},
Tobias Scharnowski³, Paul Grosen², David Fritz¹, Christopher Kruegel²,
Giovanni Vigna², Saurabh Bagchi⁴, and Mathias Payer⁵
¹Sandia National Laboratories, ²UC Santa Barbara, ³Ruhr-Universität Bochum,
⁴Purdue University, ⁵École Polytechnique Fédérale de Lausanne
{aacleme, djfritz}@sandia.gov, tobias.scharnowski@rub.de,
{edg, pcgrosen, chris, vigna}@cs.ucsb.edu,
sbagchi@purdue.edu, mathias.payer@epfl.ch



What You Corrupt Is Not What You Crash: Challenges in Fuzzing Embedded Devices

Toward the Analysis of Embedded Firmware through Automated Re-hosting

Eric Gustafson^{1,2}, Marius Muench³, Chad Spensky¹, Nilo Redini¹, Aravind Machiry¹, Yanick Fratantonio³,
Aurélien Francillon³, Davide Balzarotti³, Yung Ryn Choe², Christopher Kruegel¹, and Giovanni Vigna¹
¹University of California, Santa Barbara
{edg, cspensky, nredini, machiry, chris, vigna}@cs.ucsb.edu
²Sandia National Laboratories
{edgusta, yrchoe}@sandia.gov
³EURECOM
{marius.muench, francill, yanick.fratantonio, balzarot}@eurecom.fr

Evaluation

20 monolithic firmware images
(ground truth)



**P²IM: Scalable and Hardware-independent Firmware Instantiation through
Automatic Peripheral Interface Modeling**

Bo Feng *Northeastern University*, Alejandro Mera *Northeastern University*, Long Chen *Northeastern University*

**HALucinator: Firmware Re-hosting
Through Abstraction Layer Emulation**

Abraham A. Clements^{*1}, Eric Gustafson^{*1,2},
Tobias Scharnowski³, Paul Grosen², David Fritz¹, Christopher Kruegel²,
Giovanni Vigna², Saurabh Bagchi⁴, and Mathias Payer⁵
¹Sandia National Laboratories, ²UC Santa Barbara, ³Ruhr-Universität Bochum,
⁴Purdue University, ⁵École Polytechnique Fédérale de Lausanne
{aacleme, djfritz}@sandia.gov, tobias.scharnowski@rub.de,
{edg, pcgrosen, chris, vigna}@cs.ucsb.edu,
sbagchi@purdue.edu, mathias.payer@epfl.ch



**What You Compute Is Not What You Crash:
Challenges in Fuzzing Embedded Devices**

Automated Re-hosting for the Analysis of Embedded Firmware through Automated Re-hosting

Eric Gustafson^{1,2}, Marius Muench³, Chad Spensky¹, Nilo Redini¹, Aravind Machiry¹, Yanick Fratantonio³,
Gilles Francillon³, Davide Balzarotti³, Yung Ryn Choe², Christopher Kruegel¹, and Giovanni Vigna¹
¹University of California, Santa Barbara
{edg, cspensky, nredini, machiry, chris, vigna}@cs.ucsb.edu
²Sandia National Laboratories
{edgusta, yrchoe}@sandia.gov
³EURECOM
{marius.muench, francill, yanick.fratantonio, balzarot}@eurecom.fr

Evaluation

799 monolithic firmware images
(wild dataset)


FIRMXRAY: Detecting Bluetooth Link Layer Vulnerabilities From Bare-Metal Firmware

Haohuang Wen
wen.423@osu.edu
The Ohio State University

Zhiqiang Lin
zlin@cse.ohio-state.edu
The Ohio State University

Yinqian Zhang
yinqian@cse.ohio-state.edu
The Ohio State University

Evaluation





Evaluation

- 799 monolithic firmware images
- 340 use a dynamic memory allocator (~42%)

Evaluation

- **799** monolithic firmware images
- **340** use a dynamic memory allocator (~42%)
- **10** different HML families in **32** different variations





SECURITY TESTING




HML Security Testing



Heap exploitation primitive

- Heap overflow
- Use-after-free
- Double-free
- Fake-free




Heap vulnerable state

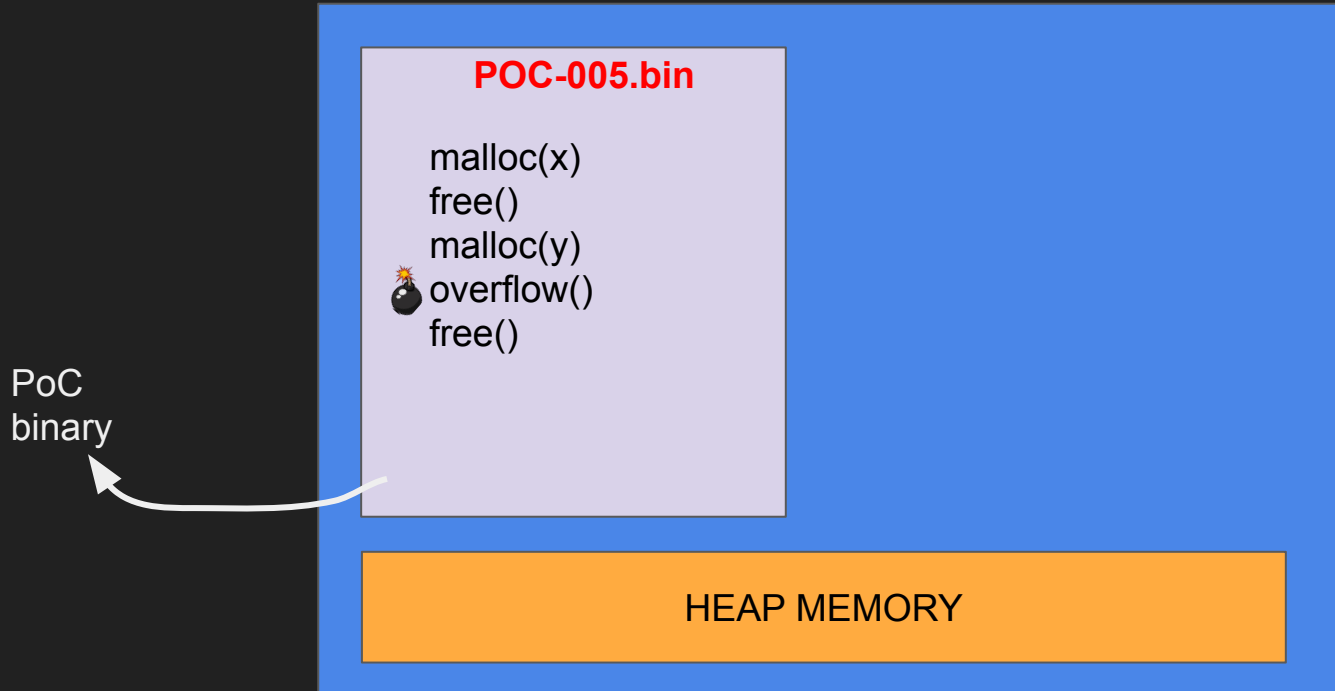
- Overlapped chunk
- Out-of-heap allocation
- Restricted write
- Arbitrary write

HML Security Testing

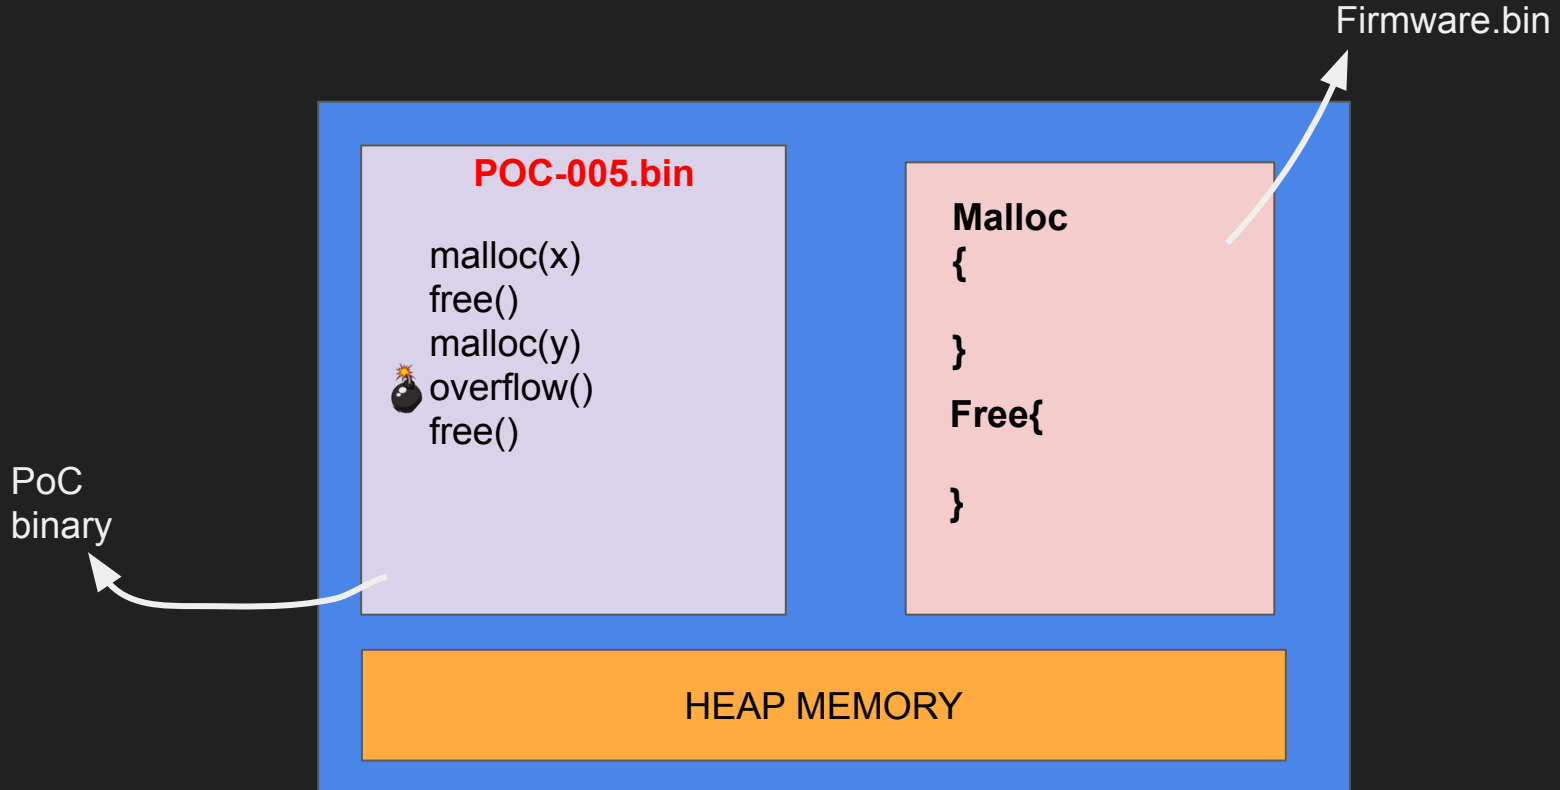
- Generation of PoC(s) that will be symbolically traced



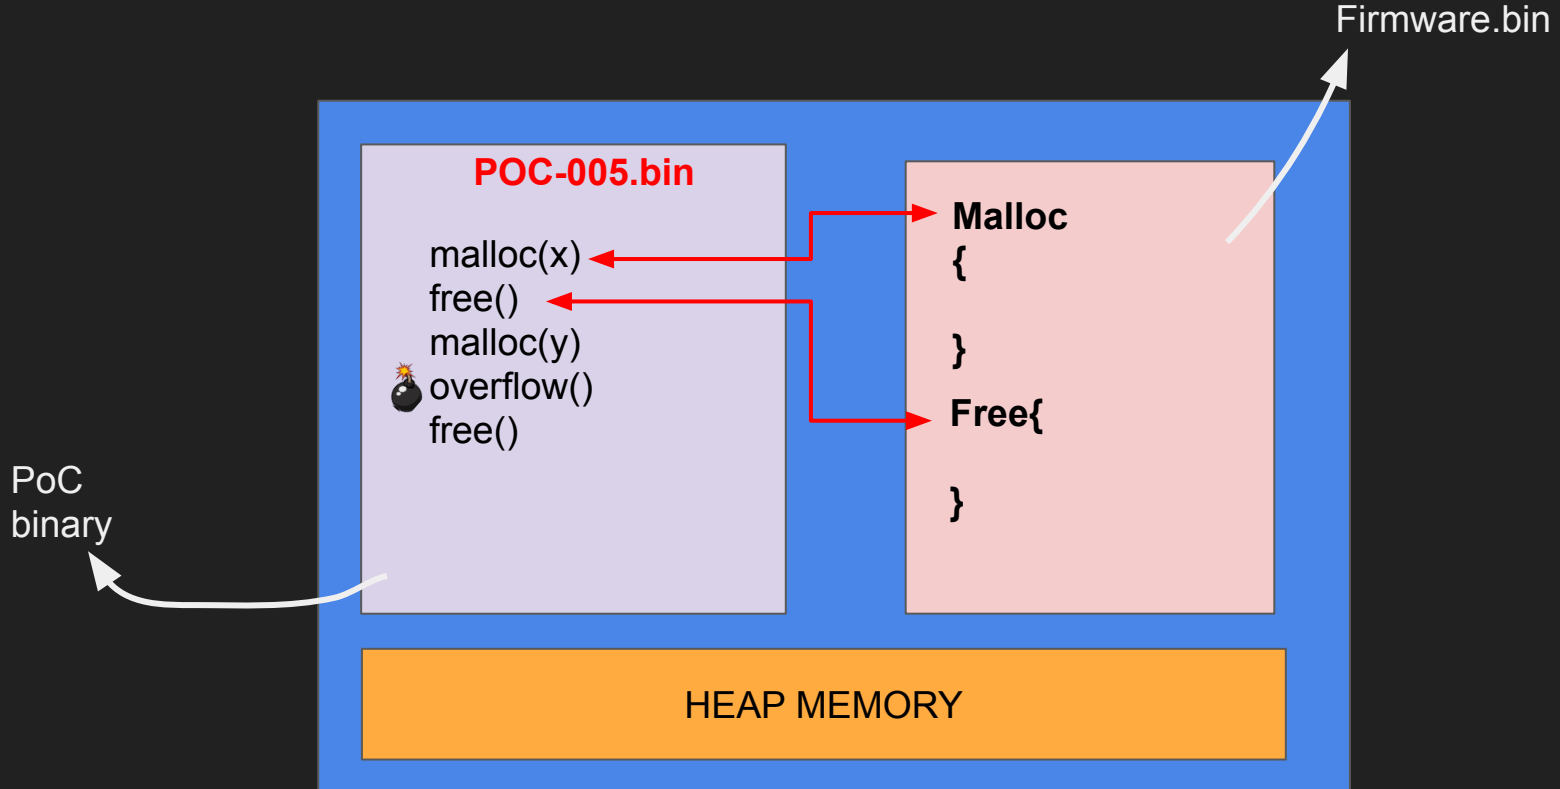
HML Security Testing



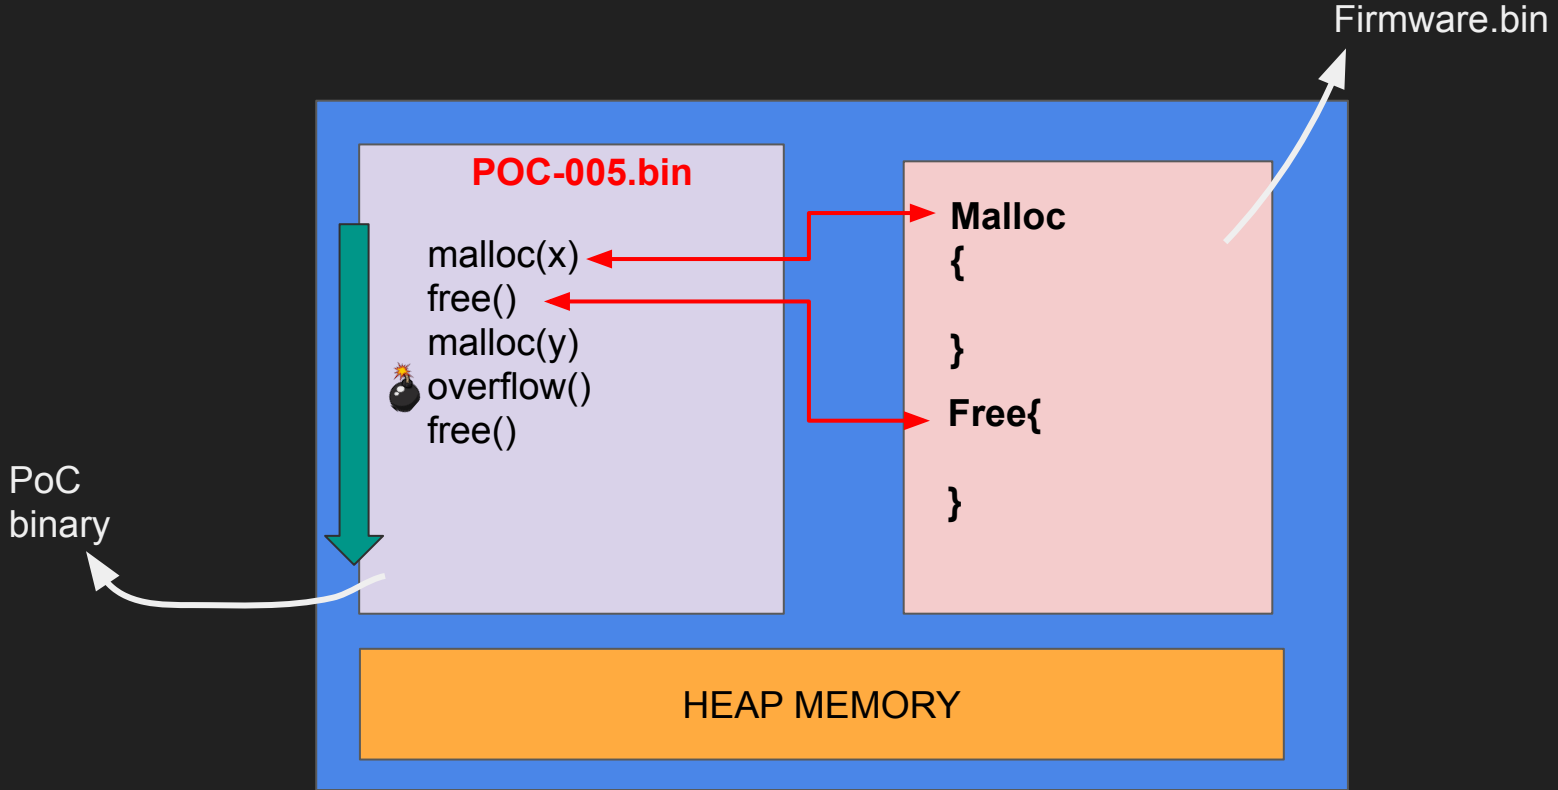
HML Security Testing



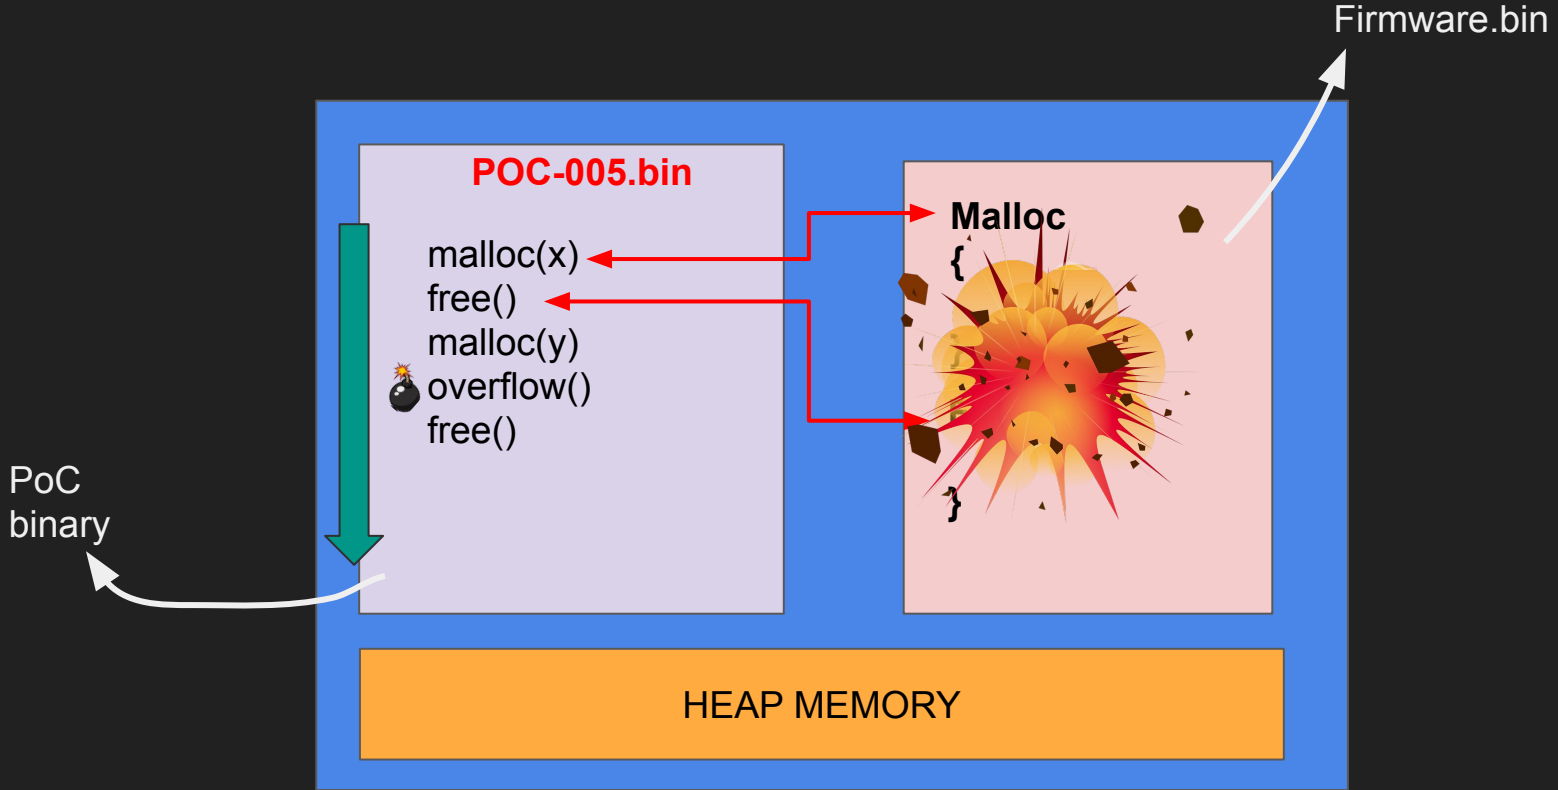
HML Security Testing




HML Security Testing




HML Security Testing



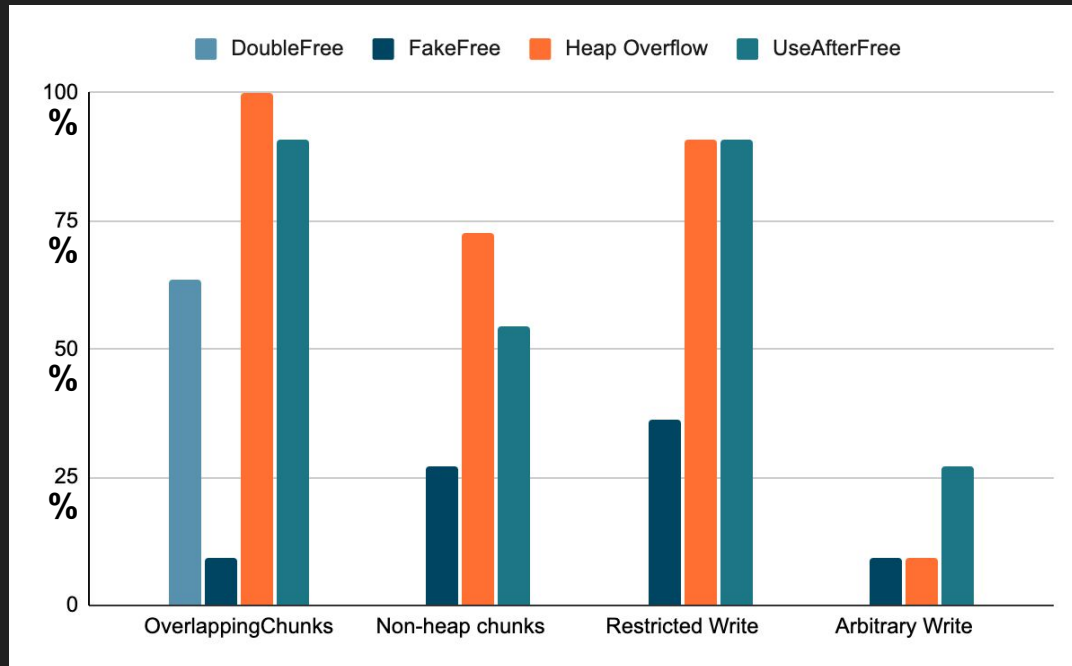


32 unique HML representatives




Security Test Results

- All the tested HML were vulnerable to at least one heap exploitation primitive




Hardware Example



Developed application that uses malloc/free

Hardware Example




Developed application that uses malloc/free



Unknown HML included in the firmware

Hardware Example



Developed application that uses malloc/free




Unknown HML included in the firmware



Heapster detected possible attacks

Hardware Example



Developed application that uses malloc/free



Unknown HML included in the firmware



Heapster detected possible attacks



Attacks confirmed on the board



Conclusions

Takeaways

- Dynamic memory allocators are often used in monolithic firmware

Takeaways

- Dynamic memory allocators are often used in monolithic firmware
- Different and unique implementations in different variants

Takeaways

- Dynamic memory allocators are often used in monolithic firmware
- Different and unique implementations in different variants
- *Every* tested HML was vulnerable to at least one heap exploitation technique



- Open source
 - github.com/ucsb-seclab/heapster
- Support
 - <https://angr.io/invite/>
 - Ping me @degrigis

Thanks!

 degrigis@ucsb.edu

 [@degrigis](https://twitter.com/degrigis)

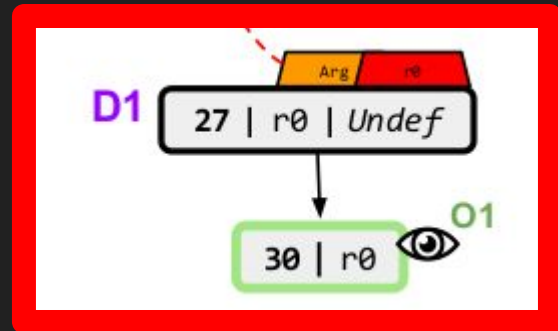
Extra

Static Taint Engine

```
1 last = 0x0
2 void mem_init(){
3     last = 0x2000C000;
4 }
5
6 int malloc(x){
7     chunk = last
8     last = last + x
9     return chunk;
10 }
11
12 int baz(){
13     void *x = 0x2000;
14     return x;
15 }
16
17 void bar(x,y){
18     if(y==0)
19         v1 = baz();
20     else
21         v1 = malloc(y);
22     foo(v1);
23 }
24
25 void foo(a){
26     int b[10];
27     memcmp(a, b, 10);
28 }
```

```
1 malloc:
2     mov r0, <arg_0>
3     ldr r1, [last]
4     add r0, r0, last
5     str r0, [last]
6     mov r0, r1
7     ret ;[04]
8
9 baz:
10    mov r0, 0x2000
11    ret ;[03]
12
13 bar:
14    mov r0, <arg0>
15    mov r1, <arg1>
16    cmp r1, 0
17    bne tag
18    call baz
19    b return
20 tag:
21    call malloc
22 return:
23    call foo ;[02]
24    ret
```

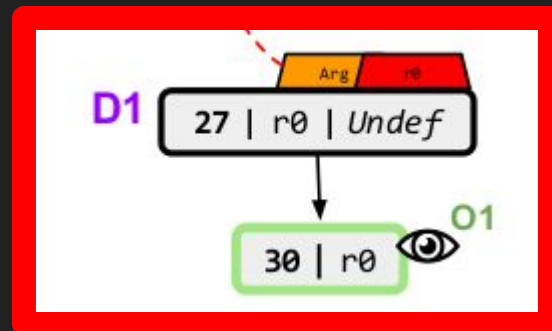
```
26 foo:
27     mov r0, <arg0>
28     mov r1, var_b
29     mov r2, 0x10
30     call memcmp ;[01]
31     ret
```



```

1 last = 0x0
2 void mem_init(){
3     last = 0x2000C000;
4 }
5
6 int malloc(x){
7     chunk = last
8     last = last + x
9     return chunk;
10 }
11
12 int baz(){
13     void *x = 0x2000;
14     return x;
15 }
16 void bar(x,y){
17     if(y==0)
18         v1 = baz();
19     else
20         v1 = malloc(y);
21     foo(v1);
22 }
23
24 void foo(a){
25     int b[10];
26     memcmp(a, b, 10);
27 }
28
29 malloc:
30     mov r0, <arg_0>
31     ldr r1, [last]
32     add r0, r0, last
33     str r0, [last]
34     mov r0, r1
35     ret ;[04]
36
37 baz:
38     mov r0, 0x2000
39     ret ;[03]
40
41 bar:
42     mov r0, <arg0>
43     mov r1, <arg1>
44     cmp r1, 0
45     bne tag
46     call baz
47     b return
48 tag:
49     call malloc
50     return:
51     call foo ;[02]
52     ret
53
54 foo:
55     mov r0, <arg0>
56     mov r1, var_b
57     mov r2, 0x10
58     call memcmp ;[01]
59     ret


```




```

1 last = 0x0
2 void mem_init(){
3     last = 0x2000C000;
4 }
5
6 int malloc(x){
7     chunk = last
8     last = last + x
9     return chunk;
10 }
11
12 int baz(){
13     void *x = 0x2000;
14     return x;
15 }
16
17 void bar(x,y){
18     if(y==0)
19         v1 = baz();
20     else
21         v1 = malloc(y);
22     foo(v1);
23 }
24
25 void foo(a){
26     int b[10];
27     memcmp(a, b, 10);
28 }
29
30 malloc:
31     mov r0, <arg_0>
32     ldr r1, [last]
33     add r0, r0, last
34     str r0, [last]
35     mov r0, r1
36     ret ;[O4]
37
38 baz:
39     mov r0, 0x2000
40     ret ;[O3]
41
42 bar:
43     mov r0, <arg0>
44     mov r1, <arg1>
45     cmp r1, 0
46     bne tag
47     call baz
48     b return
49 tag:
50     call malloc
51     return:
52     call foo ;[O2]
53     ret
54
55 foo:
56     mov r0, <arg0>
57     mov r1, var_b
58     mov r2, 0x10
59     call memcmp ;[O1]
60     ret


```



```

1 last = 0x0
2 void mem_init(){
3     last = 0x2000C000;
4 }
5
6 int malloc(x){
7     chunk = last
8     last = last + x
9     return chunk;
10 }
11
12 int baz(){
13     void *x = 0x2000;
14     return x;
15 }
16
17 void bar(x,y){
18     if(y==0)
19         v1 = baz();
20     else
21         v1 = malloc(y);
22     foo(v1);
23 }
24
25 void foo(a){
26     int b[10];
27     memcmp(a, b, 10);
28 }
29
30 malloc:
31     mov r0, <arg_0>
32     ldr r1, [last]
33     add r0, r0, last
34     str r0, [last]
35     mov r0, r1
36     ret ;[04]
37
38 baz:
39     mov r0, 0x2000
40     ret ;[03]
41
42 bar:
43     mov r0, <arg0>
44     mov r1, <arg1>
45     cmp r1, 0
46     bne tag
47     call baz
48     b return
49 tag:
50     call malloc
51     return:
52     call foo ;[02]
53     ret
54
55 foo:
56     mov r0, <arg0>
57     mov r1, var_b
58     mov r2, 0x10
59     call memcmp ;[01]
60     ret

```



```

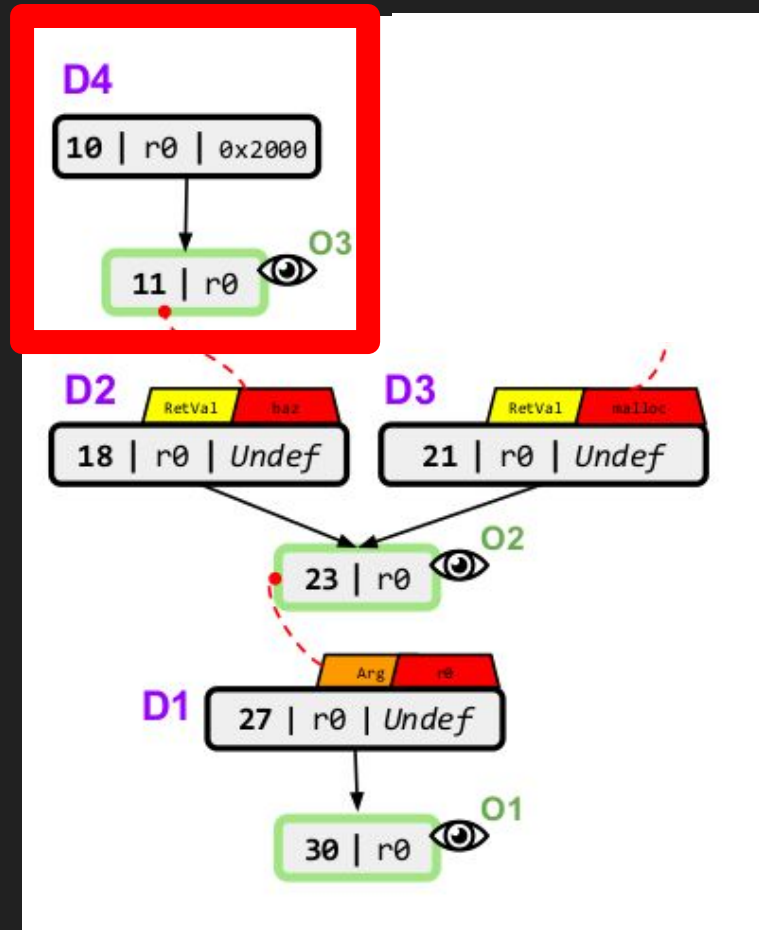
1 last = 0x0
2 void mem_init(){
3   last = 0x2000C000;
4 }
5
6 int malloc(x){
7   chunk = last
8   last = last + x
9   return chunk;
10 }
11
12 int baz(){
13   void *x = 0x2000;
14   return x;
15 }
16
17 void bar(x,y){
18   if(y==0)
19     v1 = baz();
20   else
21     v1 = malloc(y);
22   foo(v1);
23 }
24
25 void foo(a){
26   int b[10];
27   memcmp(a, b, 10);
28 }

```

```

1 malloc:
2   mov r0, <arg_0>
3   ldr r1, [last]
4   add r0, r0, last
5   str r0, [last]
6   mov r0, r1
7   ret ;[O4]
8
9 baz:
10  mov r0, 0x2000
11  ret ;[O3]
12
13 bar:
14  mov r0, <arg0>
15  mov r1, <arg1>
16  cmp r1, 0
17  bne tag
18  call baz
19  b return
20 tag:
21  call malloc
22 return:
23  call foo ;[O2]
24  ret
25
26 foo:
27  mov r0, <arg0>
28  mov r1, var_b
29  mov r2, 0x10
30  call memcmp ;[O1]
31  ret

```



```

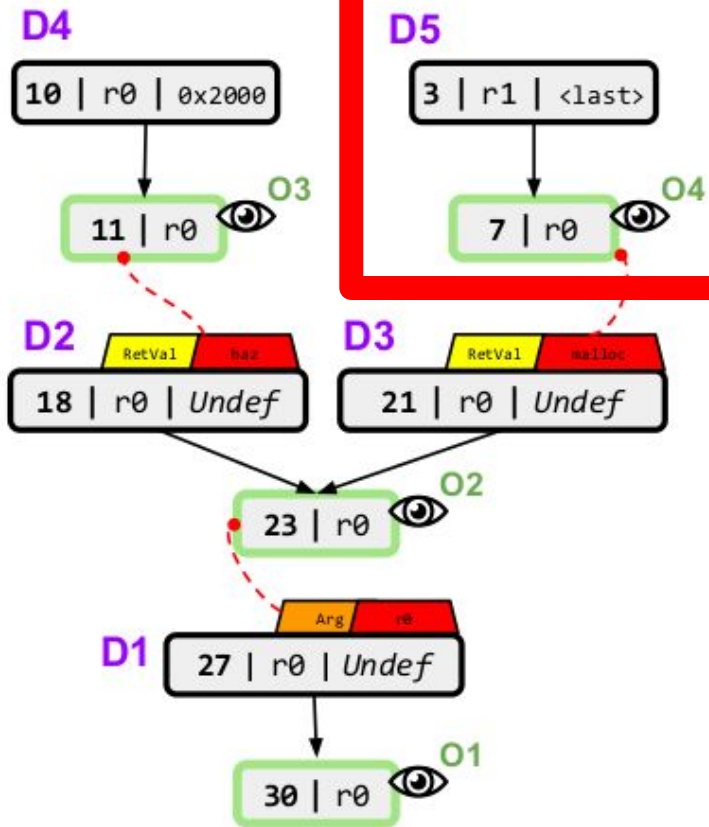
1 last = 0x0
2 void mem_init(){
3   last = 0x2000C000;
4 }
5
6 int malloc(x){
7   chunk = last
8   last = last + x
9   return chunk;
10 }
11
12 int baz(){
13   void *x = 0x2000;
14   return x;
15 }
16
17 void bar(x,y){
18   if(y==0)
19     v1 = baz();
20   else
21     v1 = malloc(y);
22   foo(v1);
23 }
24
25 void foo(a){
26   int b[10];
27   memcmp(a, b, 10);
28 }

```

```

1 malloc:
2   mov r0, <arg_0>
3   ldr r1, [last]
4   add r0, r0, last
5   str r0, [last]
6   mov r0, r1
7   ret ;[O4]
8
9 baz:
10  mov r0, 0x2000
11  ret ;[O3]
12
13 bar:
14  mov r0, <arg0>
15  mov r1, <arg1>
16  cmp r1, 0
17  bne tag
18  call baz
19  b return
20 tag:
21  call malloc
22 return:
23  call foo ;[O2]
24  ret
25
26 foo:
27  mov r0, <arg0>
28  mov r1, var_b
29  mov r2, 0x10
30  call memcmp ;[O1]
31  ret

```



Firmware Initialization

Firmware initialization (Cortex-M)

- **Reset Handler execution**

- Need to execute the compiler-injected stub that unpacks .bss and global data

```
void __noreturn Reset_Handler()  
{  
    int i; // r1  
    int v1; // r0  
  
    for ( i = 0; (unsigned int)&sdata + i < 0x200009CC; i += 4 )  
        *(void **) ((char *)&sdata + i) = *(void **) (i + 134248564);  
    v1 = SystemInit();  
    start(v1);  
    while ( 1 )  
        ;  
}
```

Firmware initialization

- **Heap Auxiliary Functions**
 - Responsible to write heap-specific data in memory

```
void mem_init()
{
    unsigned int v0; // [sp+4h] [bp+4h]

    ram = (unsigned int)&unk_1FFF3B57 & 0xFFFFFFFF; // stores 0x1fff3b54 in 0x1fff9368
    v0 = (unsigned int)&unk_1FFF3B57 & 0xFFFFFFFF;
    *(_WORD *)v0 = 0x5800;
    *(_WORD *) (v0 + 2) = 0;
    *(_BYTE *) (v0 + 4) = 0;
    ram_end = ram + 0x5800; // stores 0x1fff9354 in 0x1fff936c
    *(_BYTE *) (ram + 0x5804) = 1;
    *(_WORD *) ram_end = 0x5800;
    *(_WORD *) (ram_end + 2) = 0x5800;
    lfree = ram; // stores 0x1fff3b54 in 1fff9370
}
```


Firmware initialization

- Can be implemented in many ways...

```
int heap_init()
{
    memset_alt(byte_200050B8, 0x610u, 'H');
    memset_alt(byte_200056C8, 0xA20u, 'V');
    memset_alt(byte_200060E8, 0x1040u, 'L');
    memset_alt(byte_20007128, 0x1D18u, 'M');
    memset_alt(byte_20008E40, 0x1CB0u, 'm');
    memset_alt(byte_2000AAF0, 0xE10u, 's');
    memset_alt(byte_2000B900, 0x41A0u, 'T');
    memset_alt(byte_2000FAA0, 0xD20u, 't');
    bzero((int)word_20005068, 48);
    init_mem(0, (int)byte_200050B8);
    init_mem(1, (int)byte_200056C8);
    init_mem(2, (int)byte_200060E8);
    init_mem(3, (int)byte_20007128);
    init_mem(4, (int)byte_20008E40);
    init_mem(5, (int)byte_2000AAF0);
    init_mem(6, (int)byte_2000B900);
    return init_mem(7, (int)byte_2000FAA0);
}
```

pacemaker

```
unsigned int __fastcall init_mem(int a1, int dest_addr)
{
    int v2; // r6
    char *v3; // r4
    int v4; // r5
    unsigned int result; // r0
    int v6; // r3

    v2 = 5 * a1;
    v3 = (char *)&dword_8085D76 + 10 * a1;
    v4 = *((unsigned __int16 *)v3 + 2);
    byte_20005098[a1] = dest_addr;
    for ( result = 0; *((unsigned __int16 *)v3 + 4) > result; ++result )
    {
        *(_WORD *) (dest_addr + 4) = *(_WORD *) &dword_8085D76 + v2;
        *(_WORD *) (dest_addr + 6) = 0;
        *(_WORD *) (dest_addr + 8) = 0;
        *(_BYTE *) (dest_addr + 10) = 0;
        *(_BYTE *) (dest_addr + 11) = 0;
        if ( *((unsigned __int16 *)v3 + 4) - 1 == result )
            v6 = 0;
        else
            v6 = dest_addr + v4;
        *(_DWORD *) dest_addr = v6;
        dest_addr += v4;
    }
    return result;
}
```

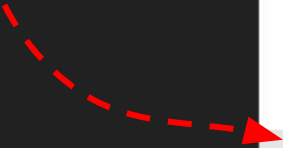

Firmware initialization

- Can be implemented in many ways...

```
int sub_11EFC()
{
    __BYTE v1[6]; // [sp+0h] [bp-D0h] BYREF
    __int16 v2; // [sp+6h] [bp-CAh] BYREF
    __int16 v3; // [sp+Ah] [bp-C6h] BYREF
    __int16 v4; // [sp+Eh] [bp-C2h] BYREF
    __int16 v5; // [sp+12h] [bp-BEh] BYREF
    __int16 v6; // [sp+2Ah] [bp-A6h] BYREF
    __int16 v7; // [sp+2Eh] [bp-A2h] BYREF
    __int16 v8; // [sp+32h] [bp-9Eh] BYREF
    __int16 v9; // [sp+36h] [bp-9Ah] BYREF

    sub_11440(119);
    sub_3B4(v1, &byte_2000288D, 6);
    sub_3B4(&v2, &unk_20004370, 4);
    sub_3B4(&v3, &unk_20004378, 4);
    sub_3B4(&v4, &unk_2000437C, 4);
    sub_402(&v5, 24);
    sub_3B4(&v6, &dword_200028A8, 4);
    sub_3B4(&v7, &dword_200028AC, 4);
    sub_3B4(&v8, &dword_200028A4, 4);
    sub_3B4(&v9, &unk_20004382, 96);
    return sub_114EC(487424, v1, 150);
}
```

Calls free() to insert first
Free chunk.



AC603_VIITA

HML Classification

HML Classifications


Malloc001

Malloc002

HML Classifications



HML Classifications



HML Classifications

Binary Similarity?

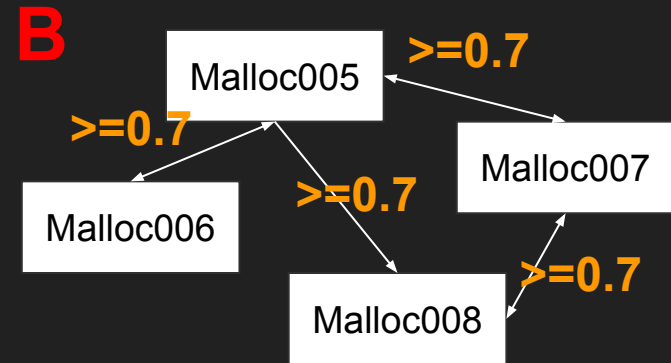
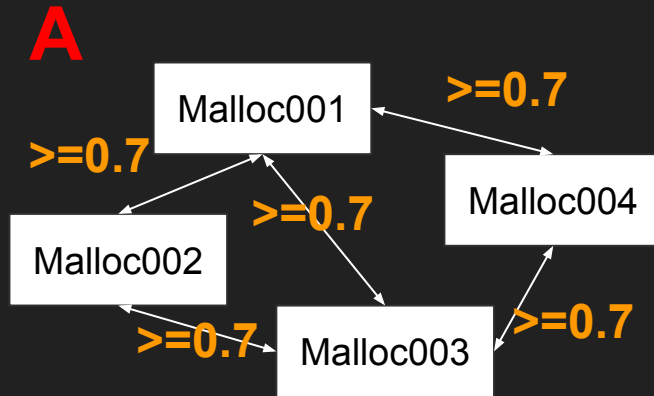


Same “family” if ≥ 0.7


Malloc001_3

Malloc002_3

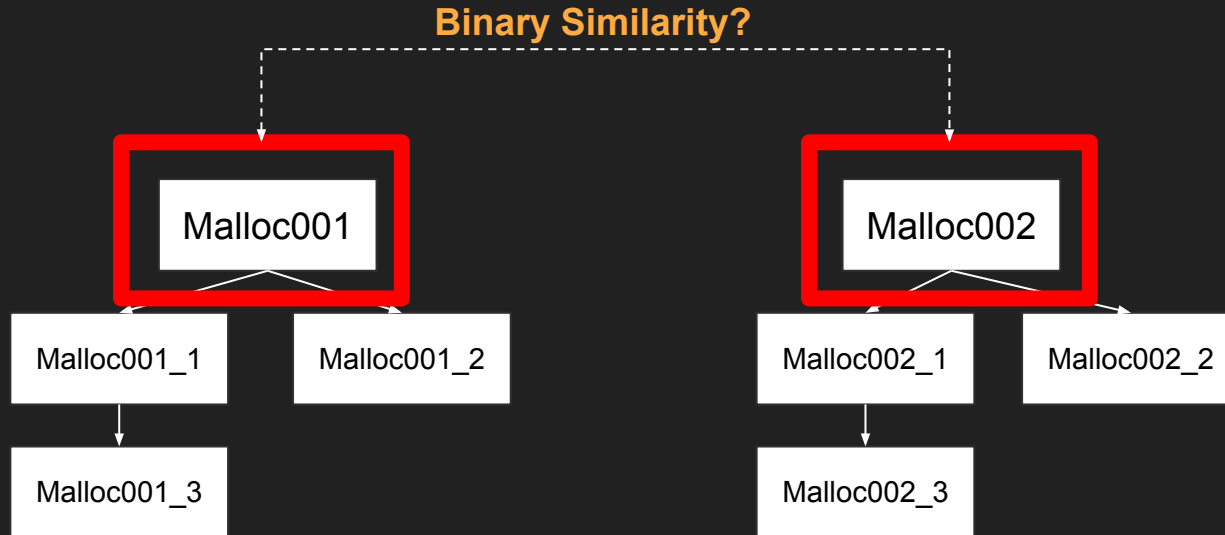
HML Classifications



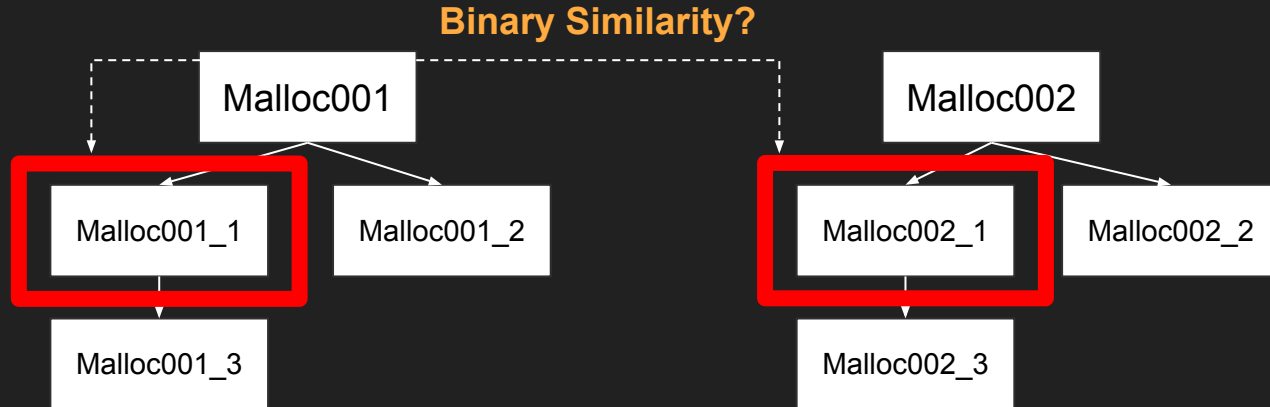
HML Classifications



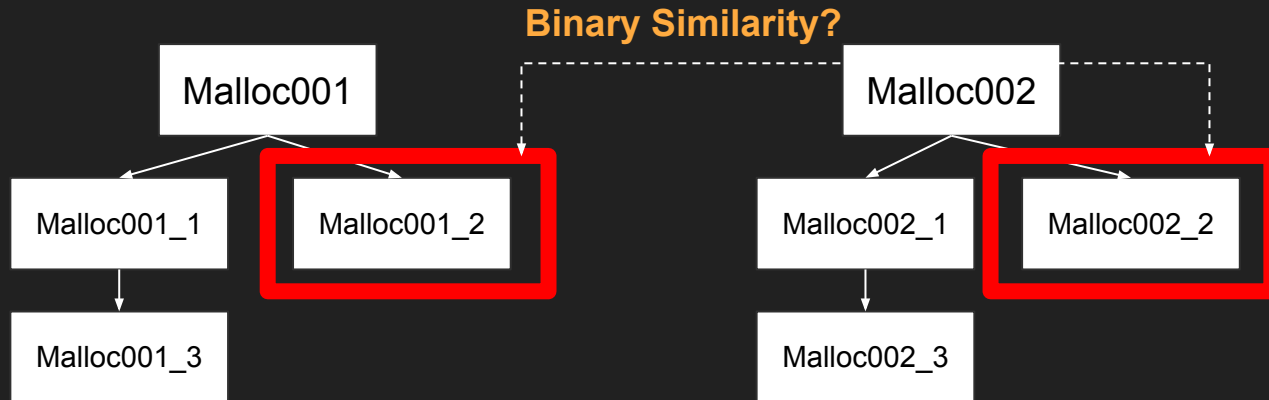
HML Classifications



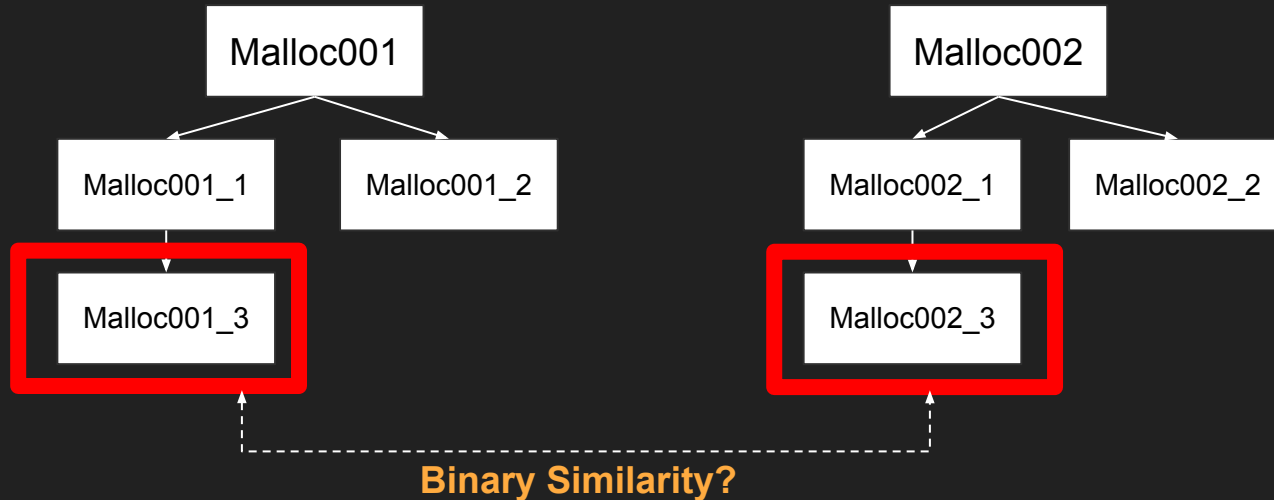
HML Classifications



HML Classifications



HML Classifications




HML Classifications

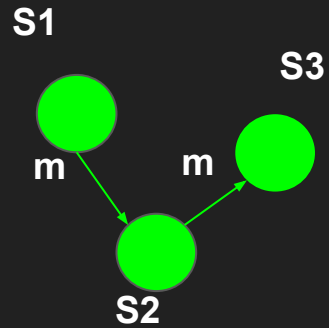


PoC Tracing

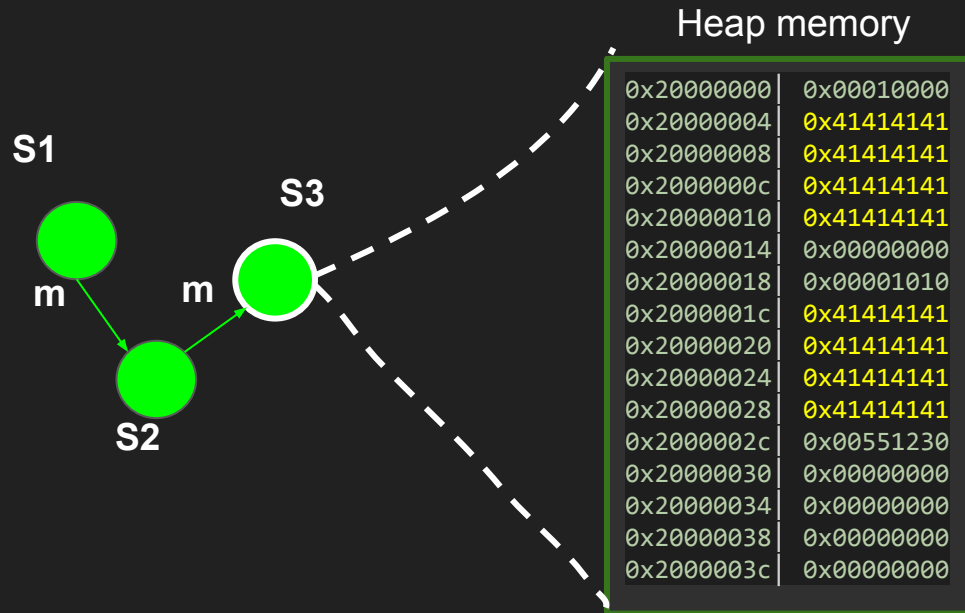
HML Security Testing (PoC tracing)



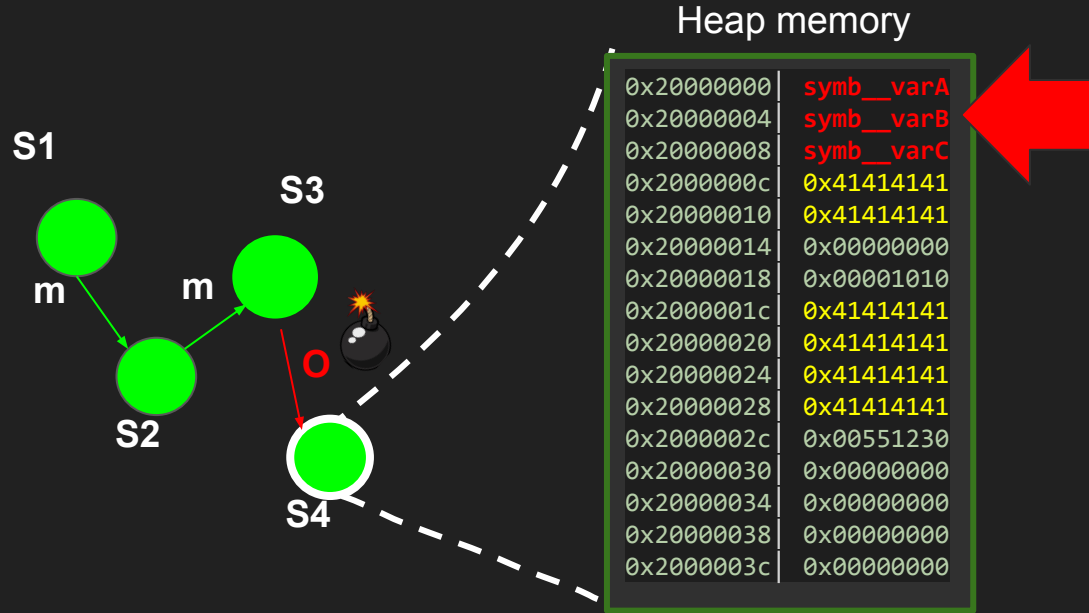
HML Security Testing (PoC tracing)




HML Security Testing (PoC tracing)



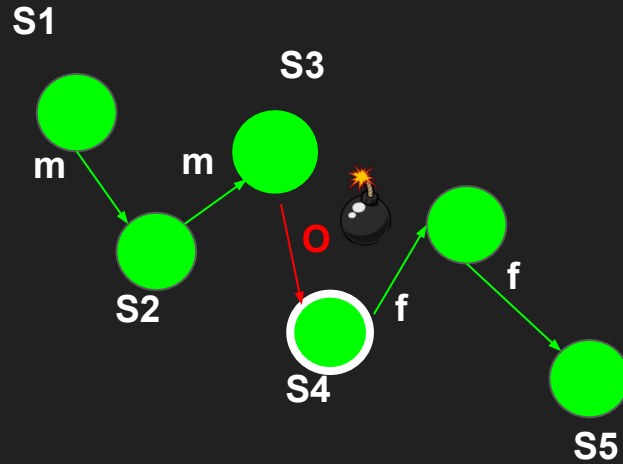
HML Security Testing (PoC tracing)



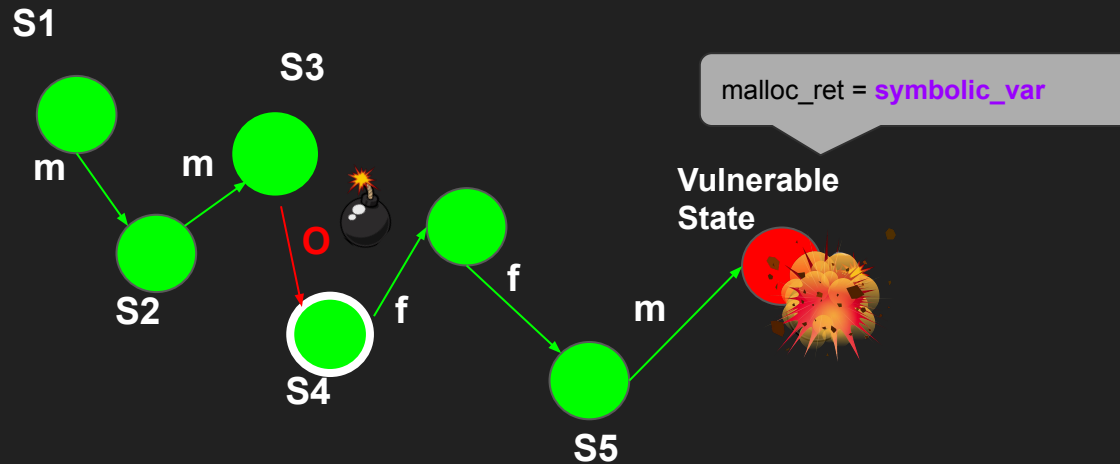
HML Security Testing (PoC tracing)




HML Security Testing (PoC tracing)



HML Security Testing (PoC tracing)

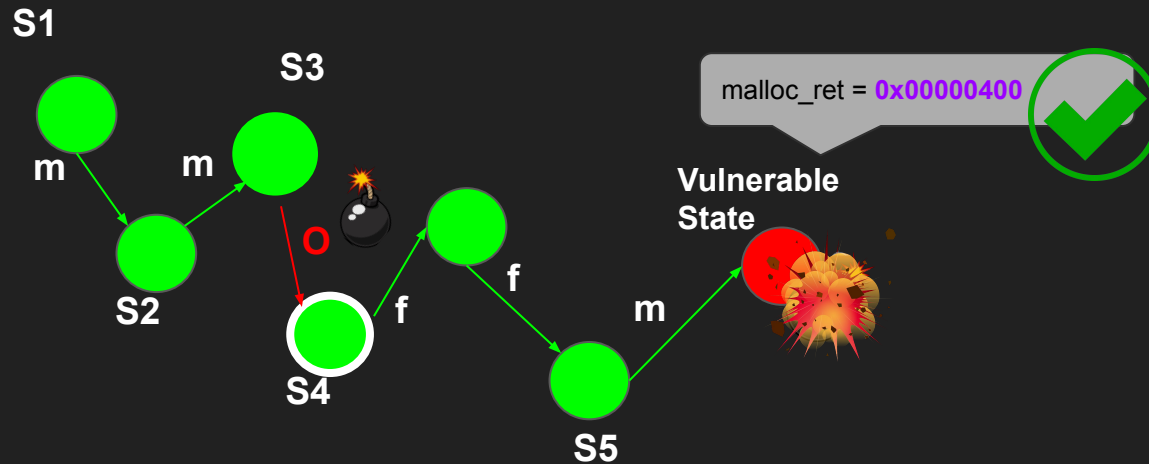


HML Security Testing (PoC tracing)




Can I concretize value to be out-of-heap?



HML Security Testing (PoC tracing)




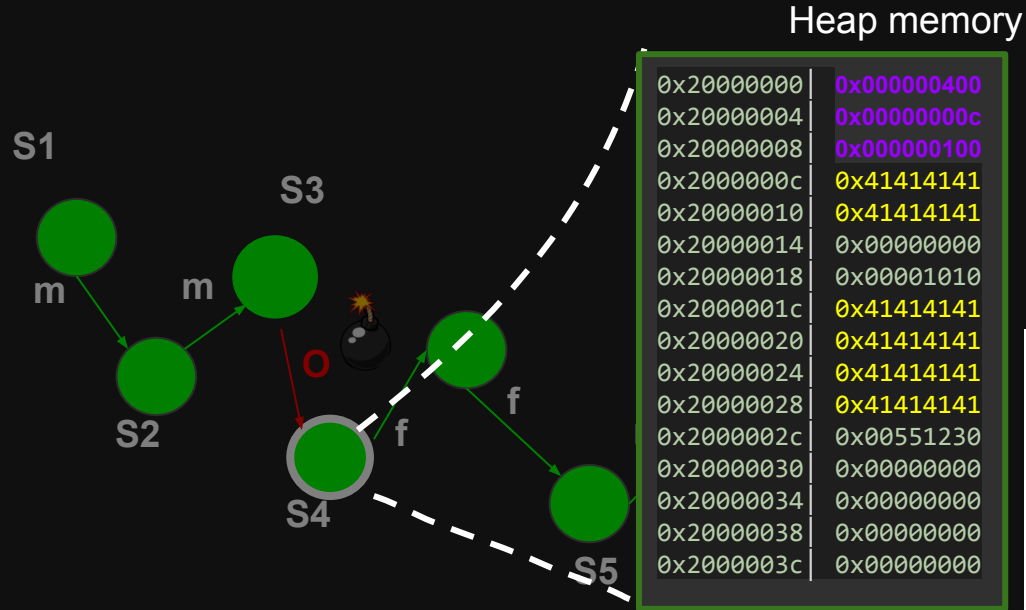
HML Security Testing (PoC tracing)



HML Security Testing (PoV generation)



HML Security Testing (PoV generation)






```
m1=malloc(X)
free(m1)
m2=malloc(Y)
overflow(m1)
free(m2)
```



POC-005.bin



```
m1=malloc(10)
free(m1)
m2=malloc(20)
overflow(m1)
free(m2)
```



POV-005.bin


PoV Feasibility

Security Test Results

What about the feasibility of PoVs attacks?

Security Test Results

What about

```
m1=malloc(10)
free(m1)
m2=malloc(20)
 overflow(m1)
free(m2)
```

of PoVs attacks?

POV-005.bin

Security Test Results

```
m1=malloc(10)
free(m1)
m2=malloc(20)
💣 overflow(m1)
free(m2)
```

m1=malloc(10)


free(m1)

m2=malloc(20)

💣 overflow(m1)

free(m2)

Security Test Results

```
m1=malloc(10)
free(m1)
m2=malloc(20)
 overflow(m1)
free(m2)
```

```
m1=malloc(10)
```

No re-hosting, extremely challenging...

```
overflow(m1)
```

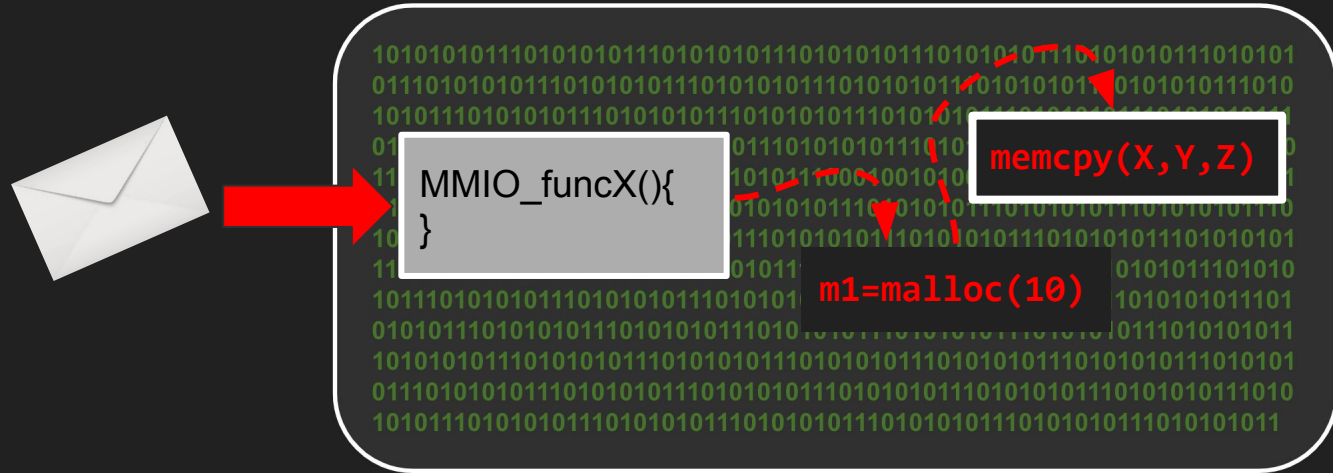
```
free(m2)
```


Security Test Results



P1: Can I reach a call to `malloc` from functions that read data from MMIO regions?

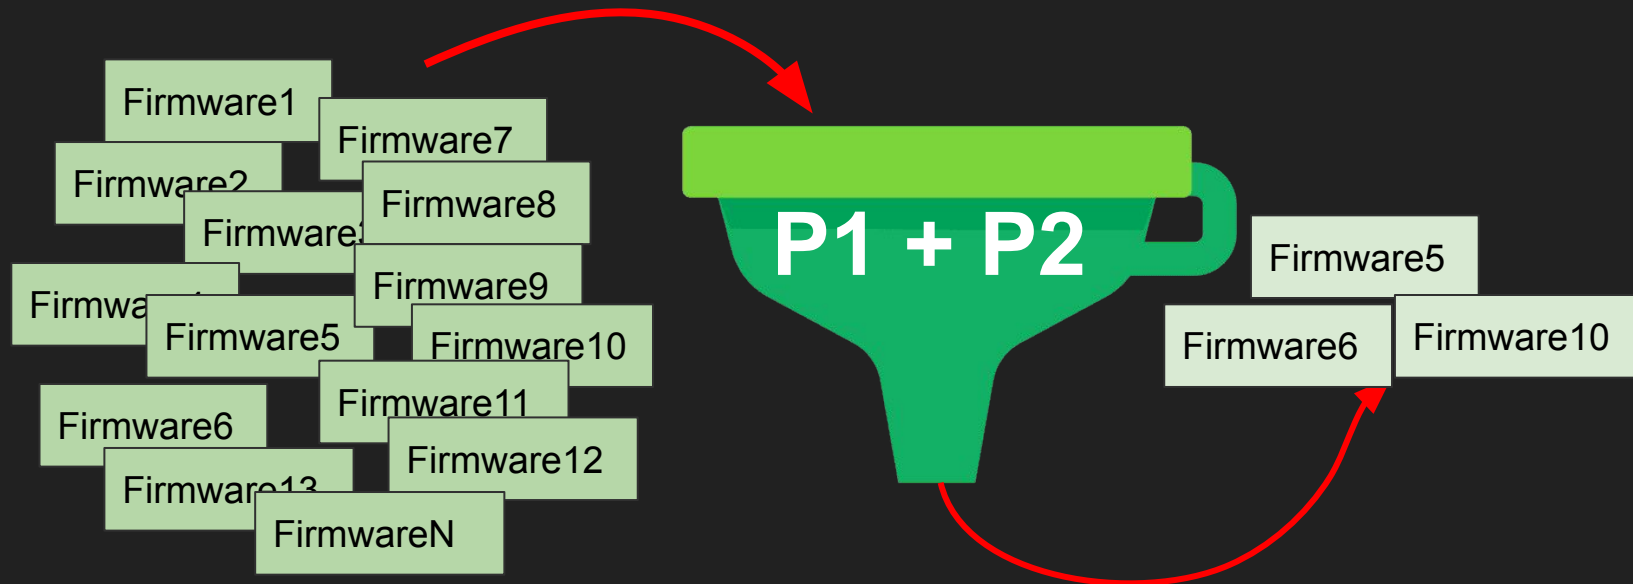
Security Test Results



P2: Can I observe a flow from malloc to a memcpy with not constant size?

Security Test Results

- Manually review of pre-selected images



Security Test Results

54 firmware selected for manual investigation

Security Test Results

54 firmware selected for manual investigation

4 contains valid exploitation primitives

Security Test Results

Additional research must be done!