

GhostRace: Exploiting and Mitigating Speculative Race Conditions

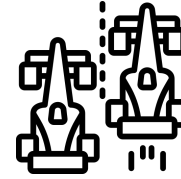
Hany Ragab, Andrea Mambretti, Anil Kurmus, and Cristiano Giuffrida



In This Talk

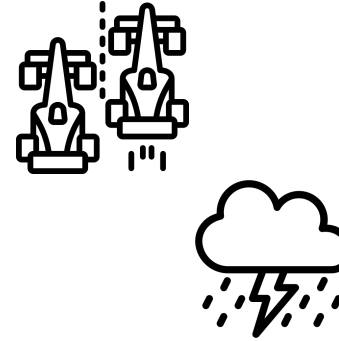
In This Talk

- Speculative Race Condition



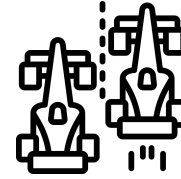
In This Talk

- Speculative Race Condition
- Inter-Process Interrupt Storming



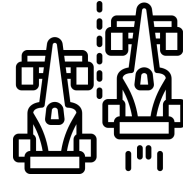
In This Talk

- Speculative Race Condition
- Inter-Process Interrupt Storming
- Gadget Scanner (1200+)



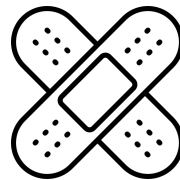
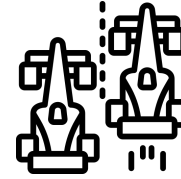
In This Talk

- Speculative Race Condition
- Inter-Process Interrupt Storming
- Gadget Scanner (1200+)
- PoC Exploit (12KB/s)



In This Talk

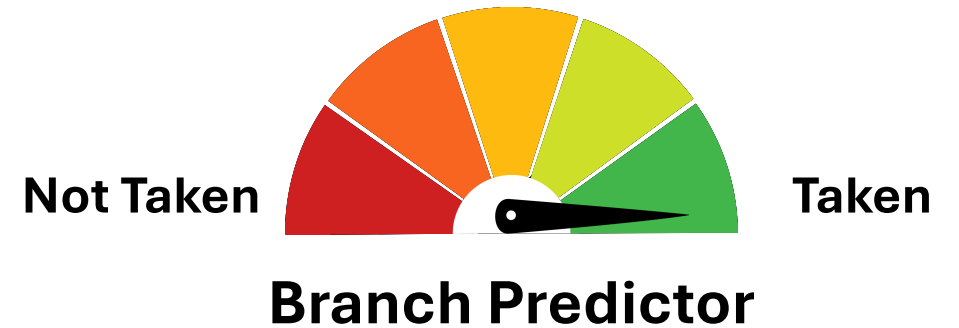
- Speculative Race Condition
- Inter-Process Interrupt Storming
- Gadget Scanner (1200+)
- PoC Exploit (12KB/s)
- Mitigation (5%)



Background

Spectre v1

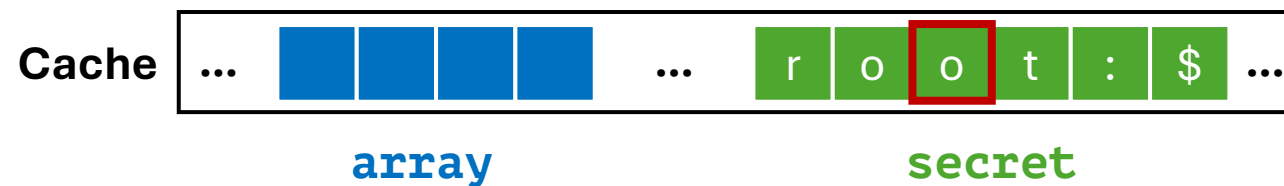
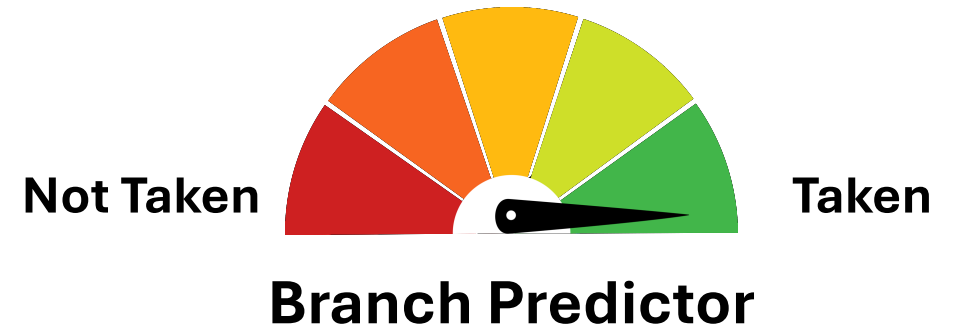
```
if (condition){  
  ...  
}
```



Background Spectre v1

Conditional Branch Misprediction

```
if (condition){  
  ...  
}
```



Background

Concurrency Bugs

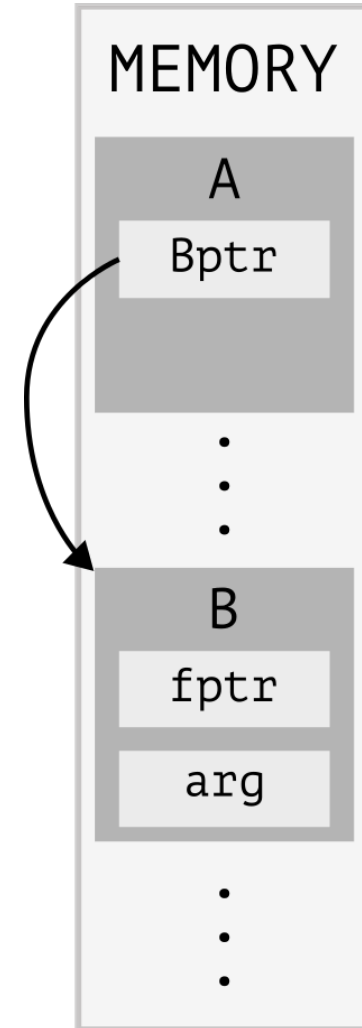
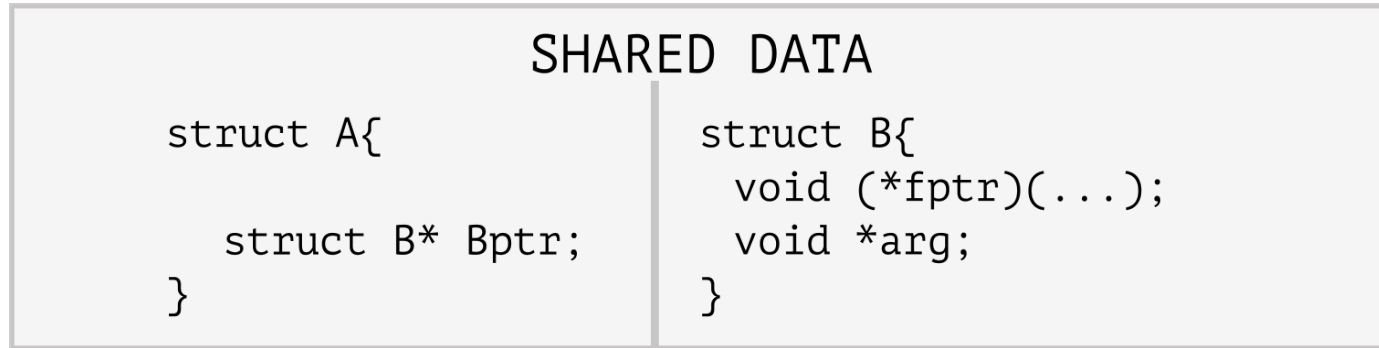
Background

Concurrency Bugs

```
                                SHARED DATA
struct A{
    struct B* Bptr;
}
struct B{
    void (*fptr)(...);
    void *arg;
}
```

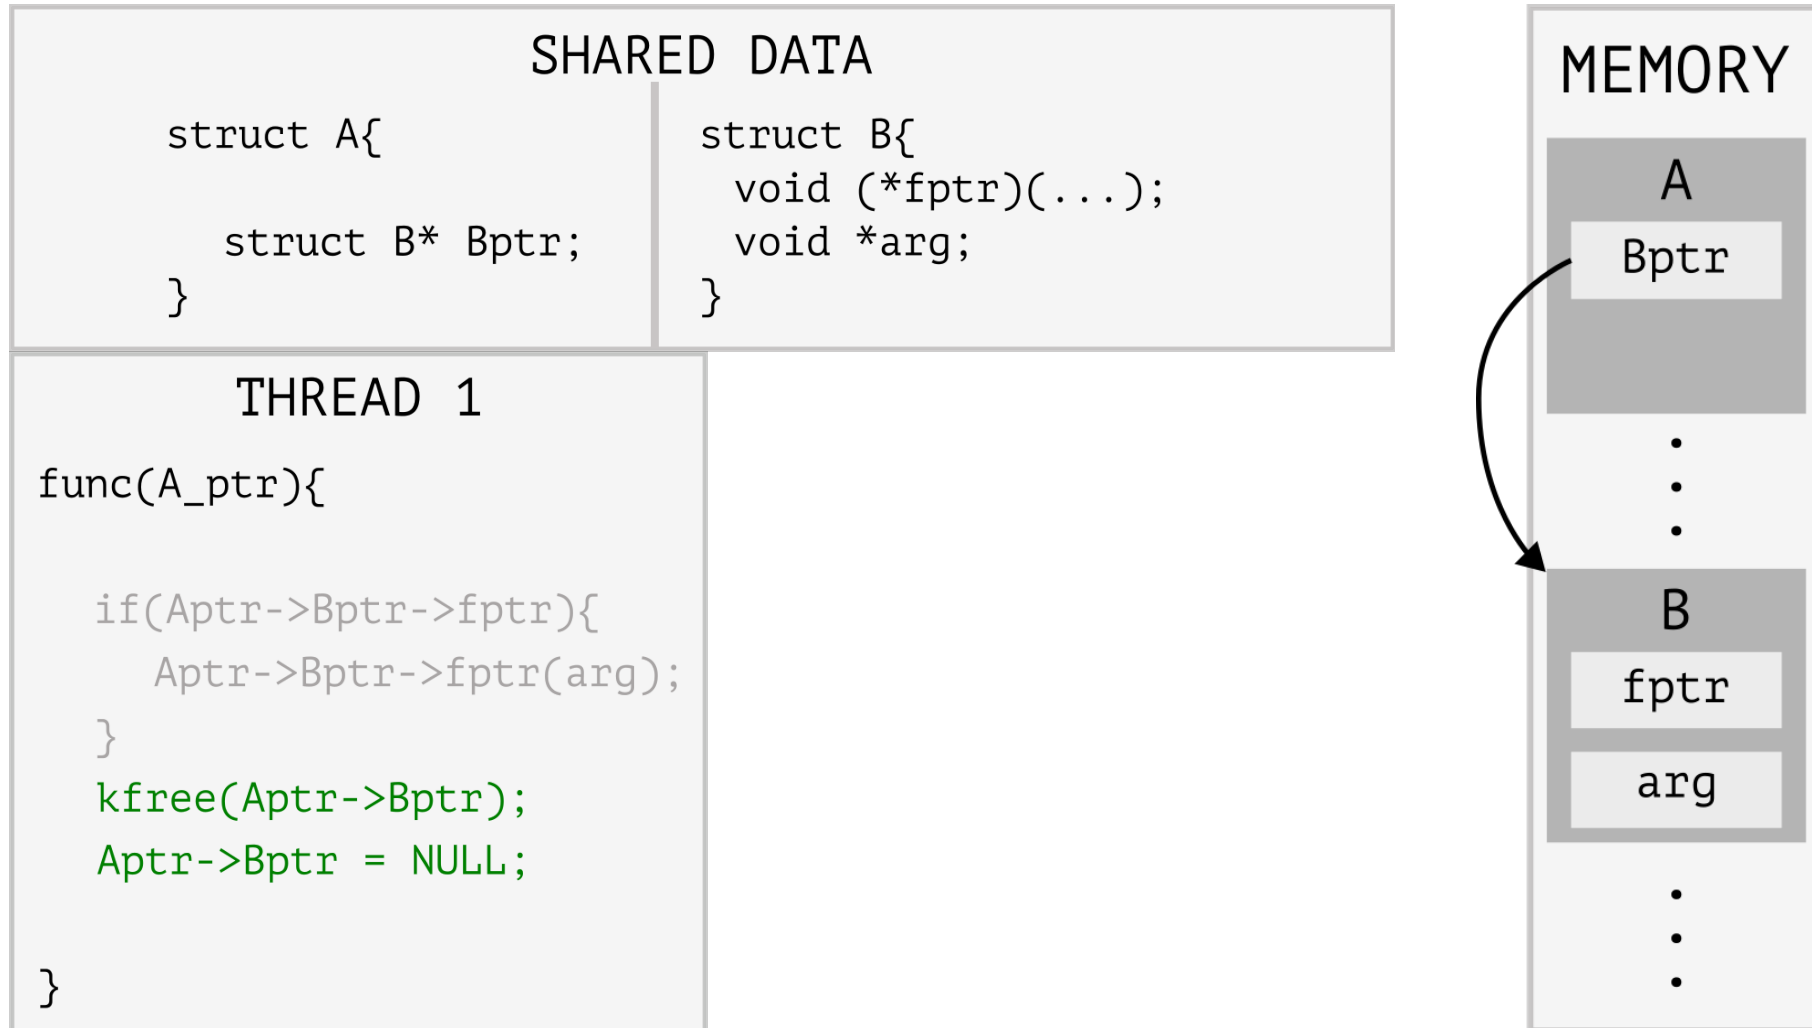
Background

Concurrency Bugs



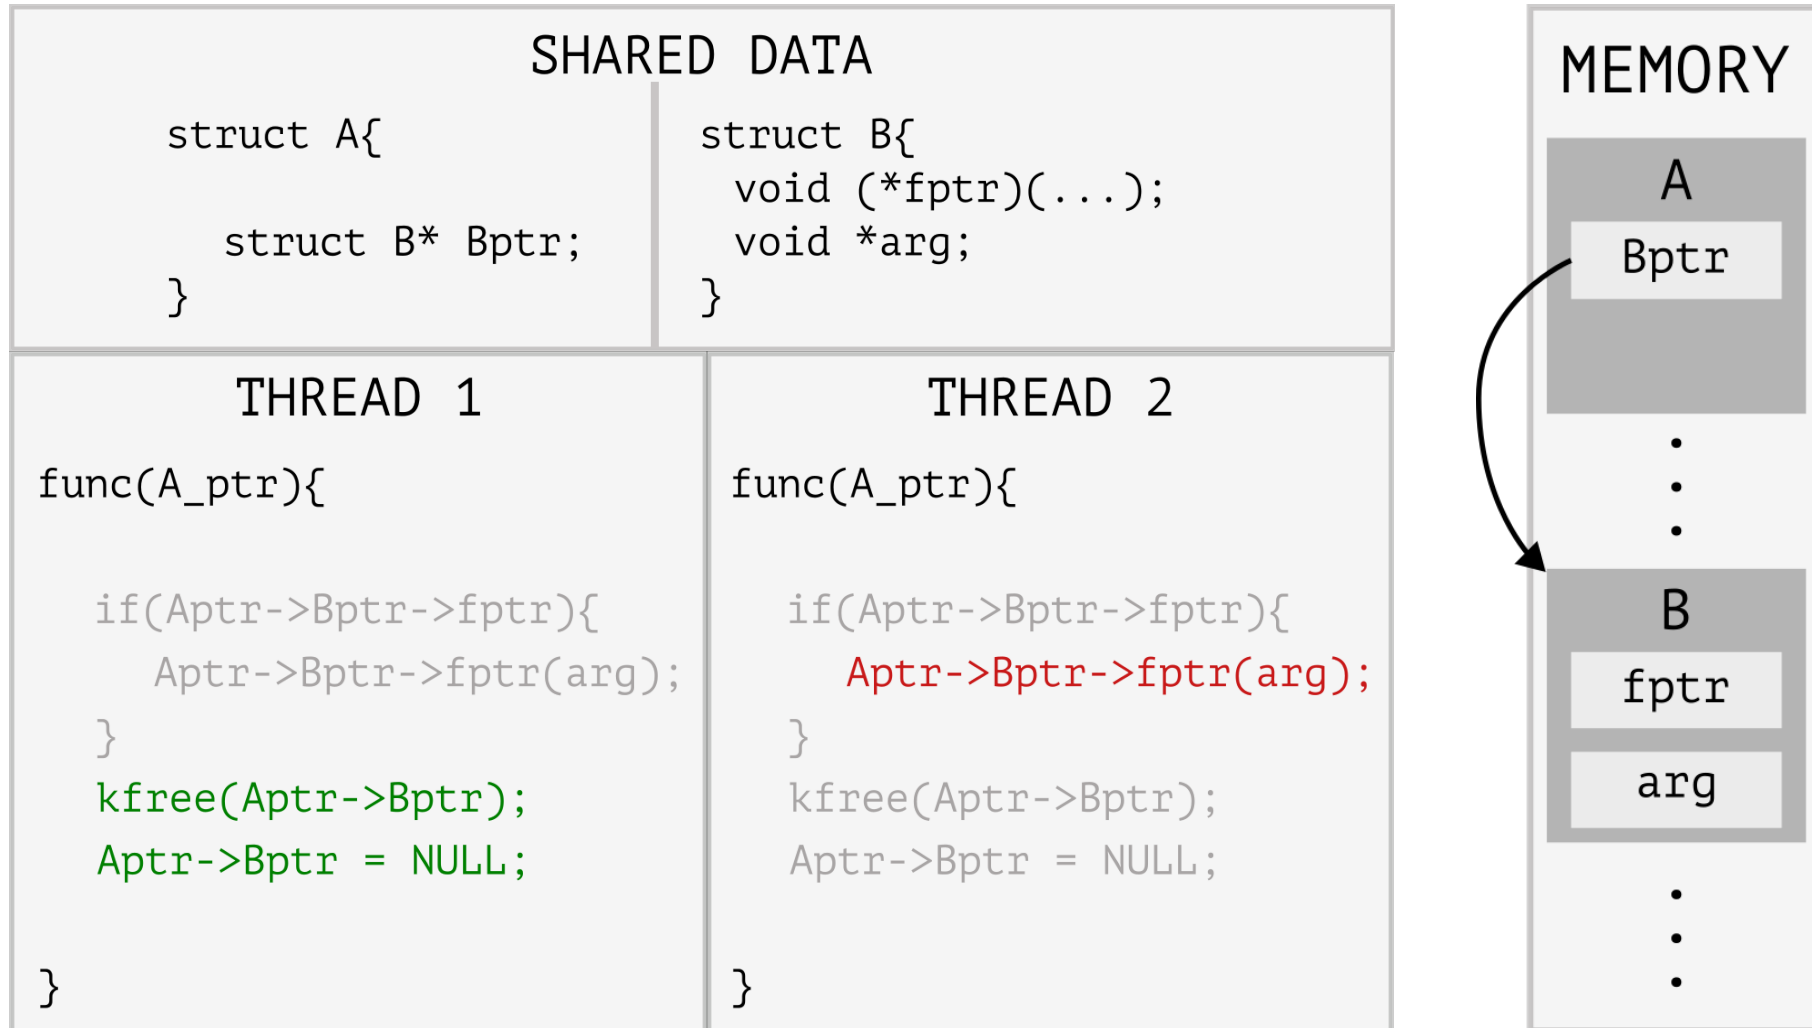
Background

Concurrency Bugs



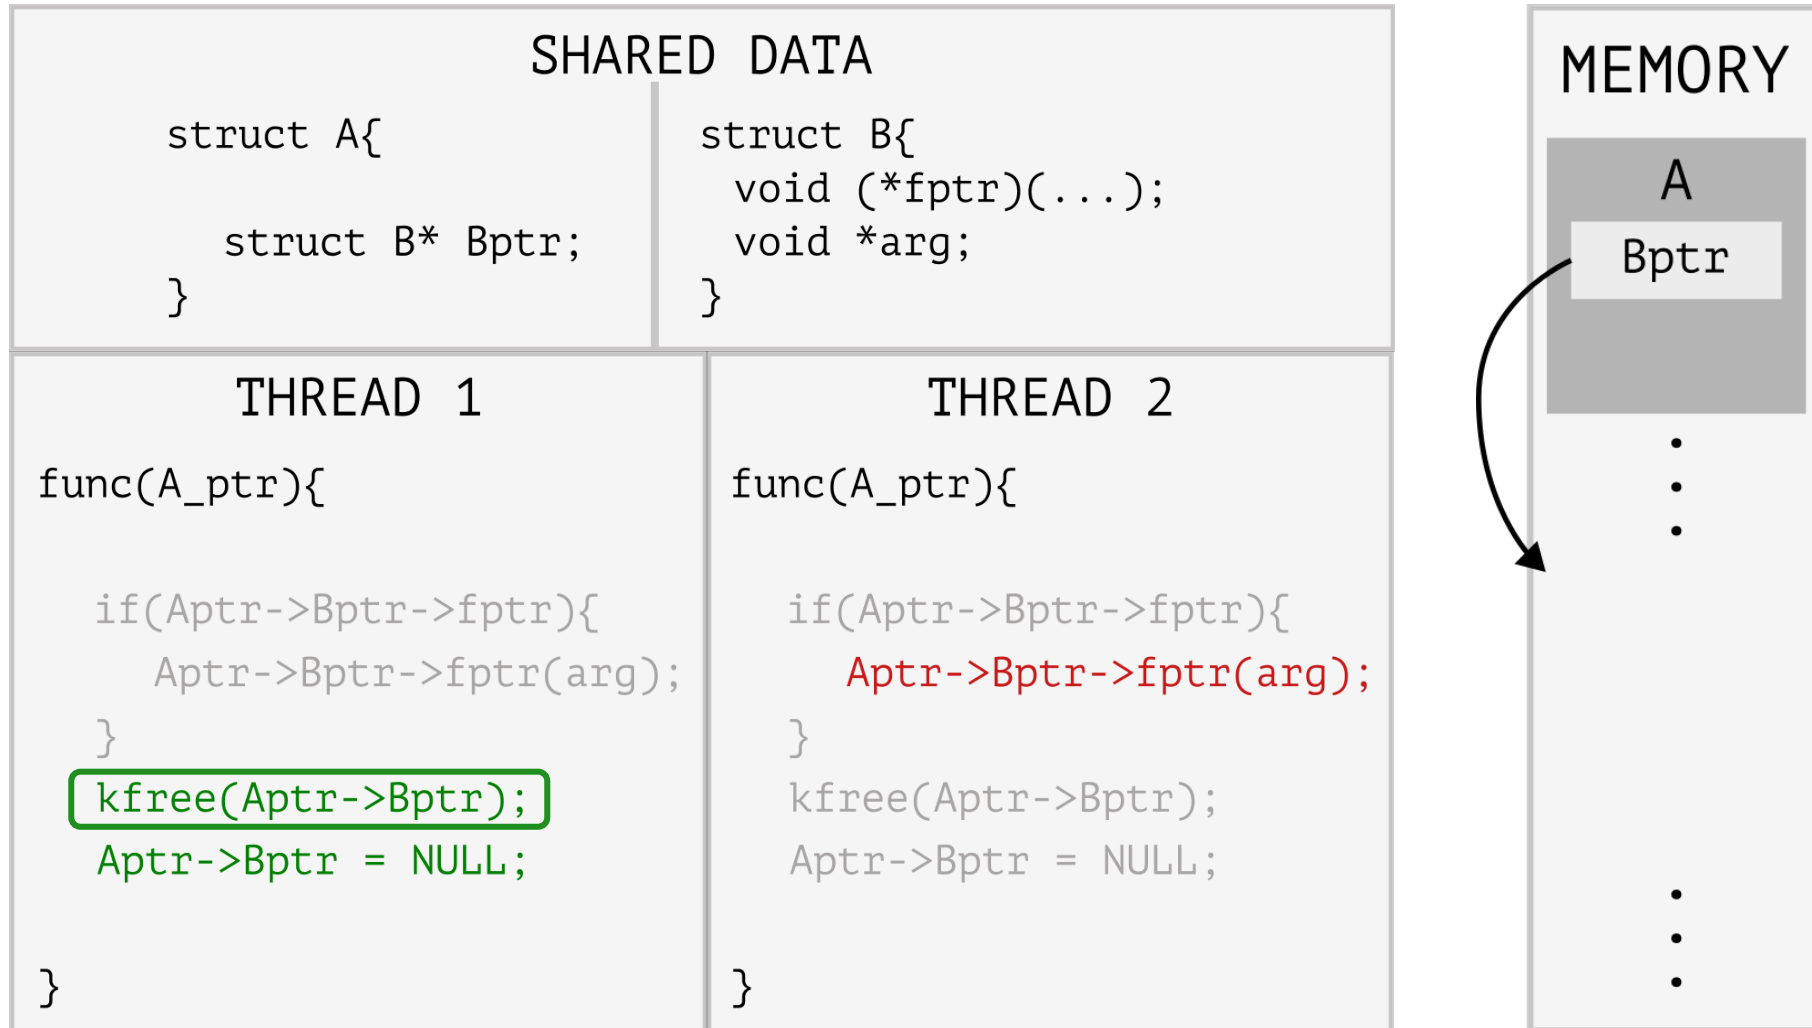
Background

Concurrency Bugs



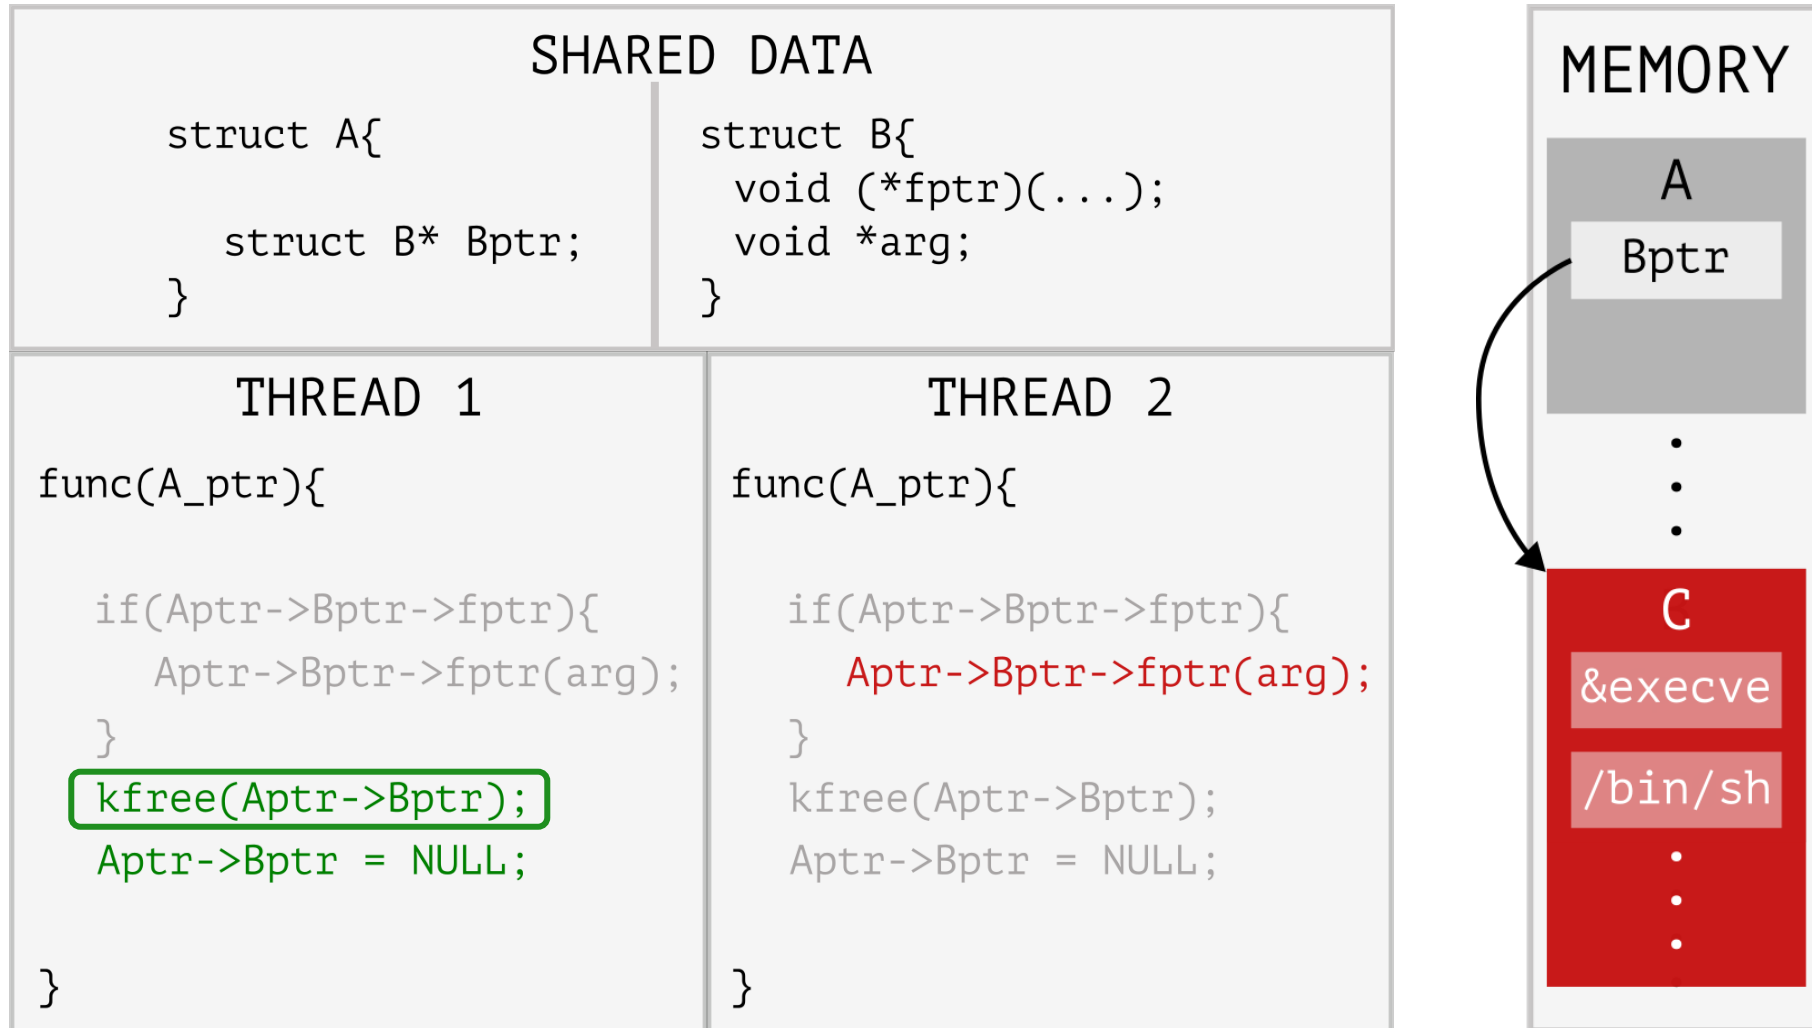
Background

Concurrency Bugs

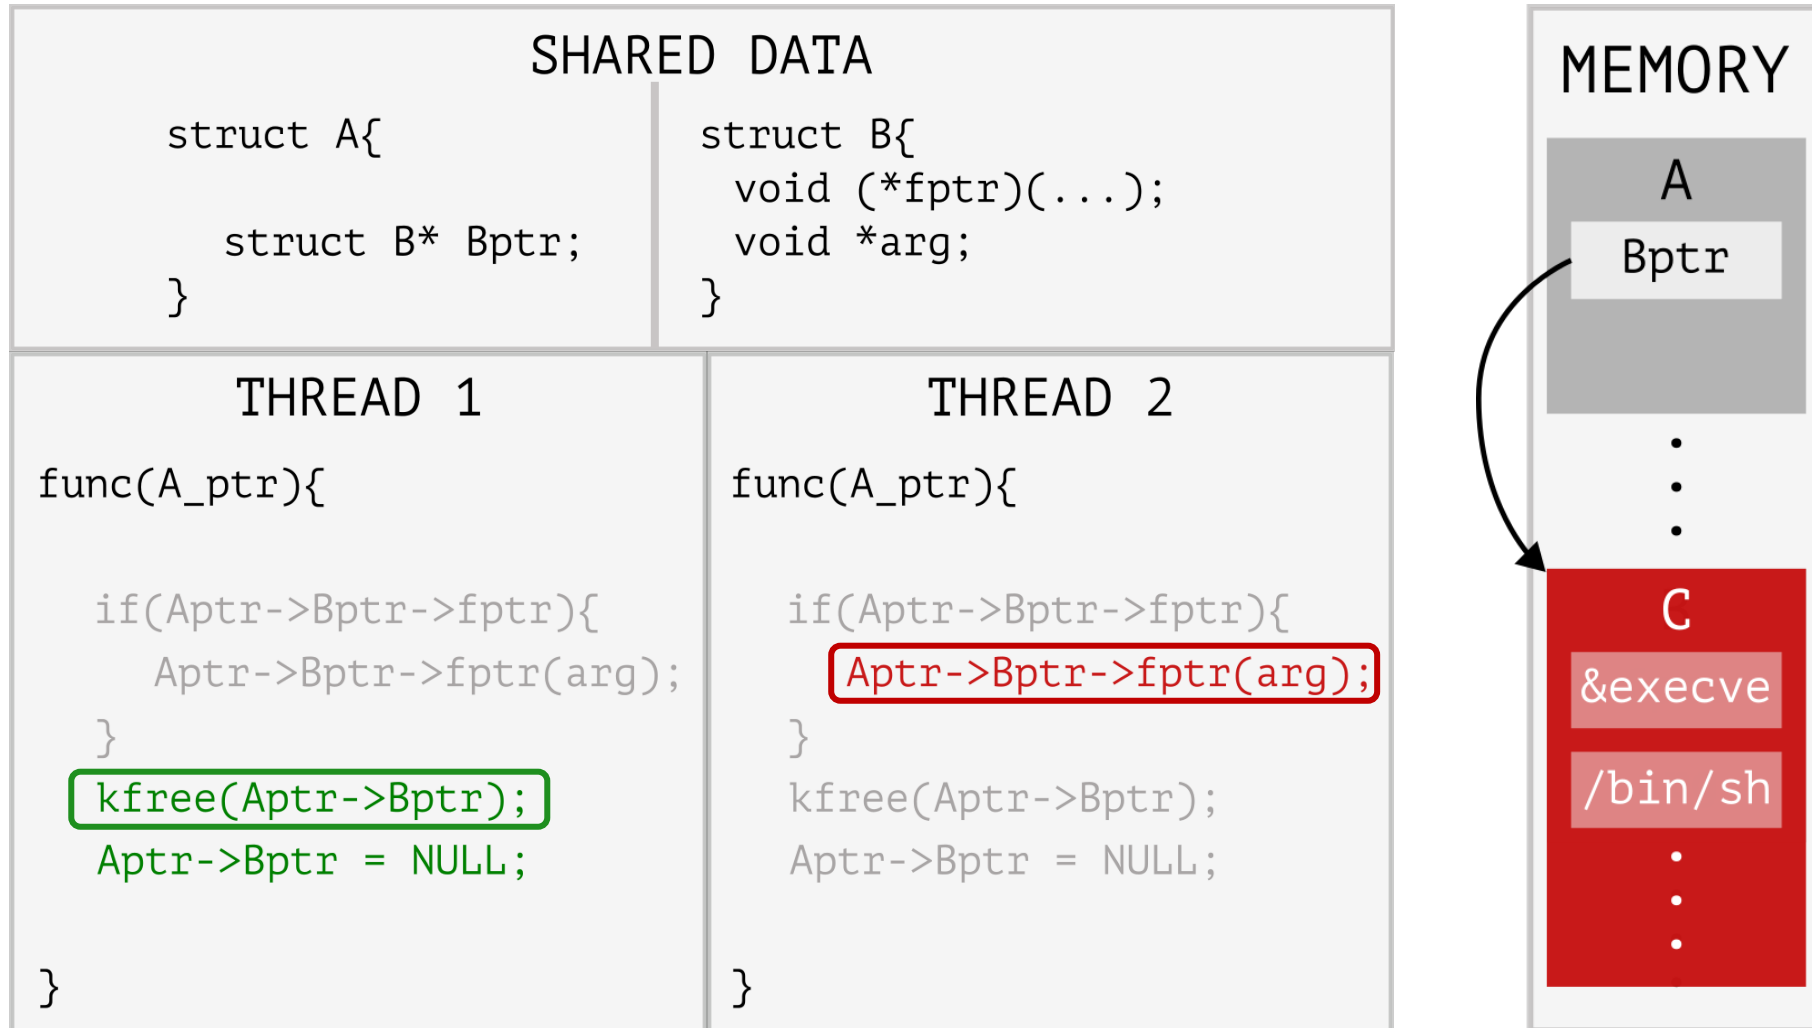


Background

Concurrency Bugs

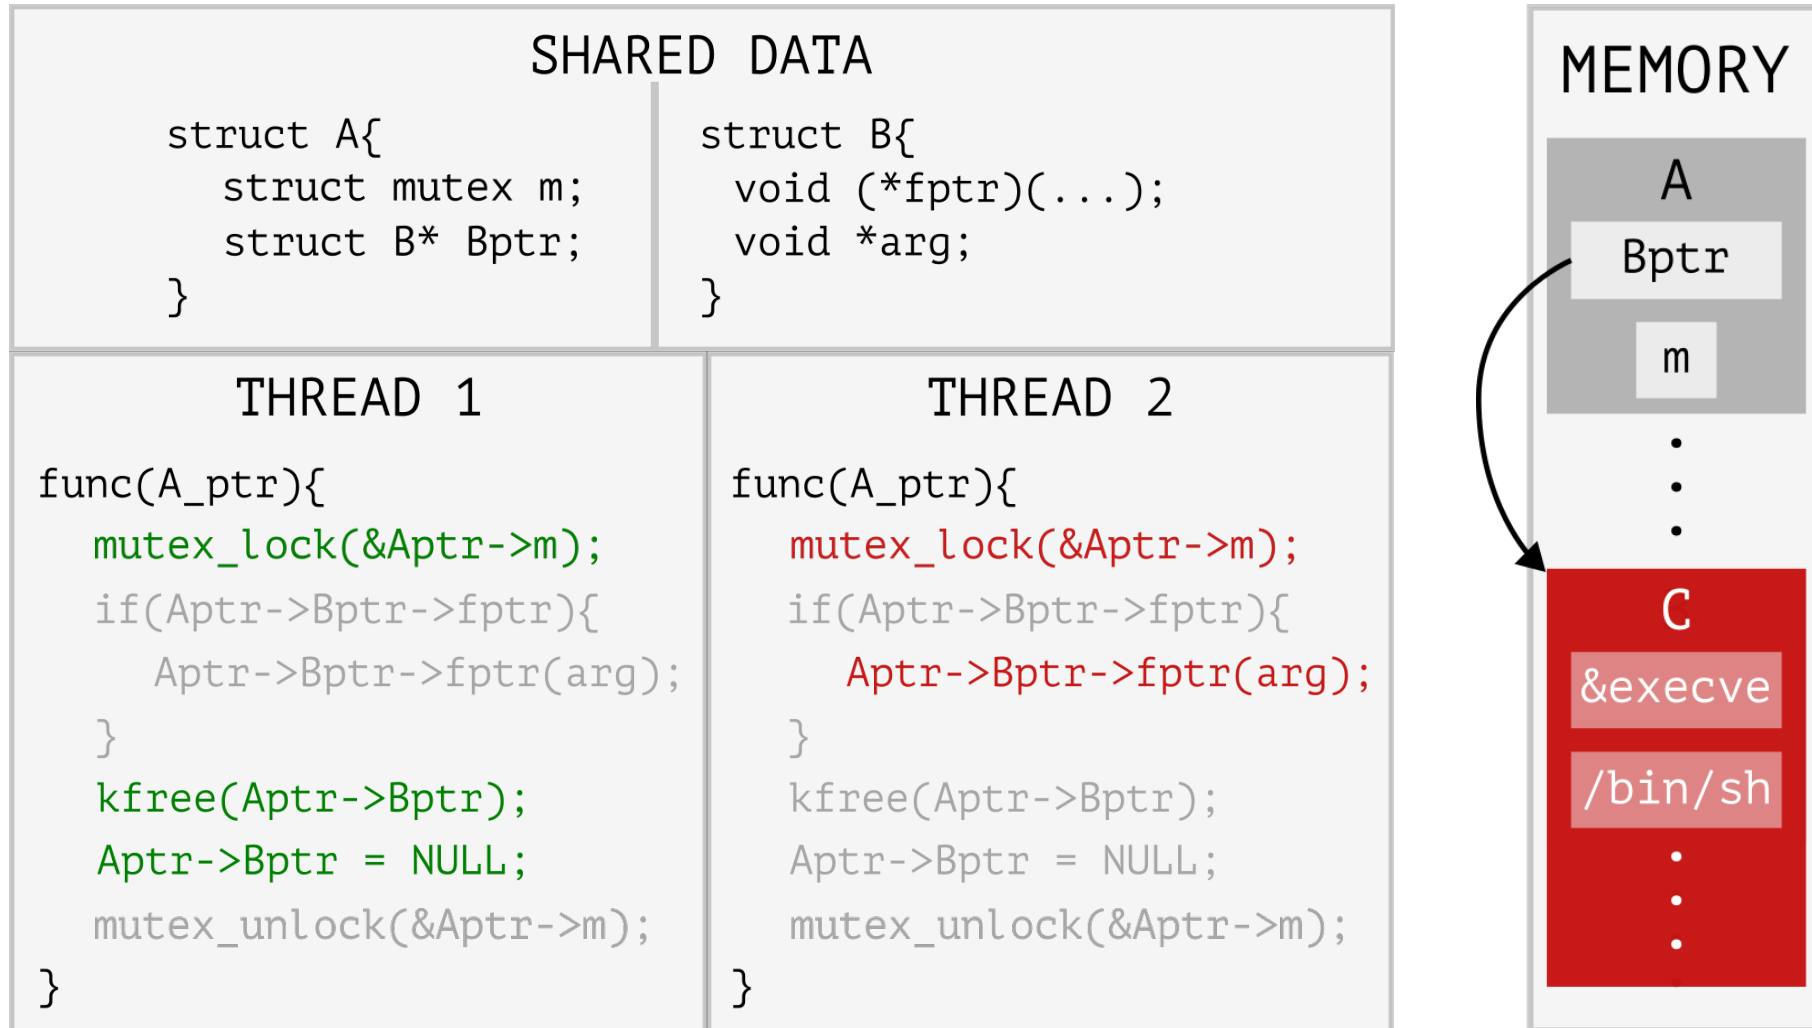


Concurrent Use-After-Free Attack



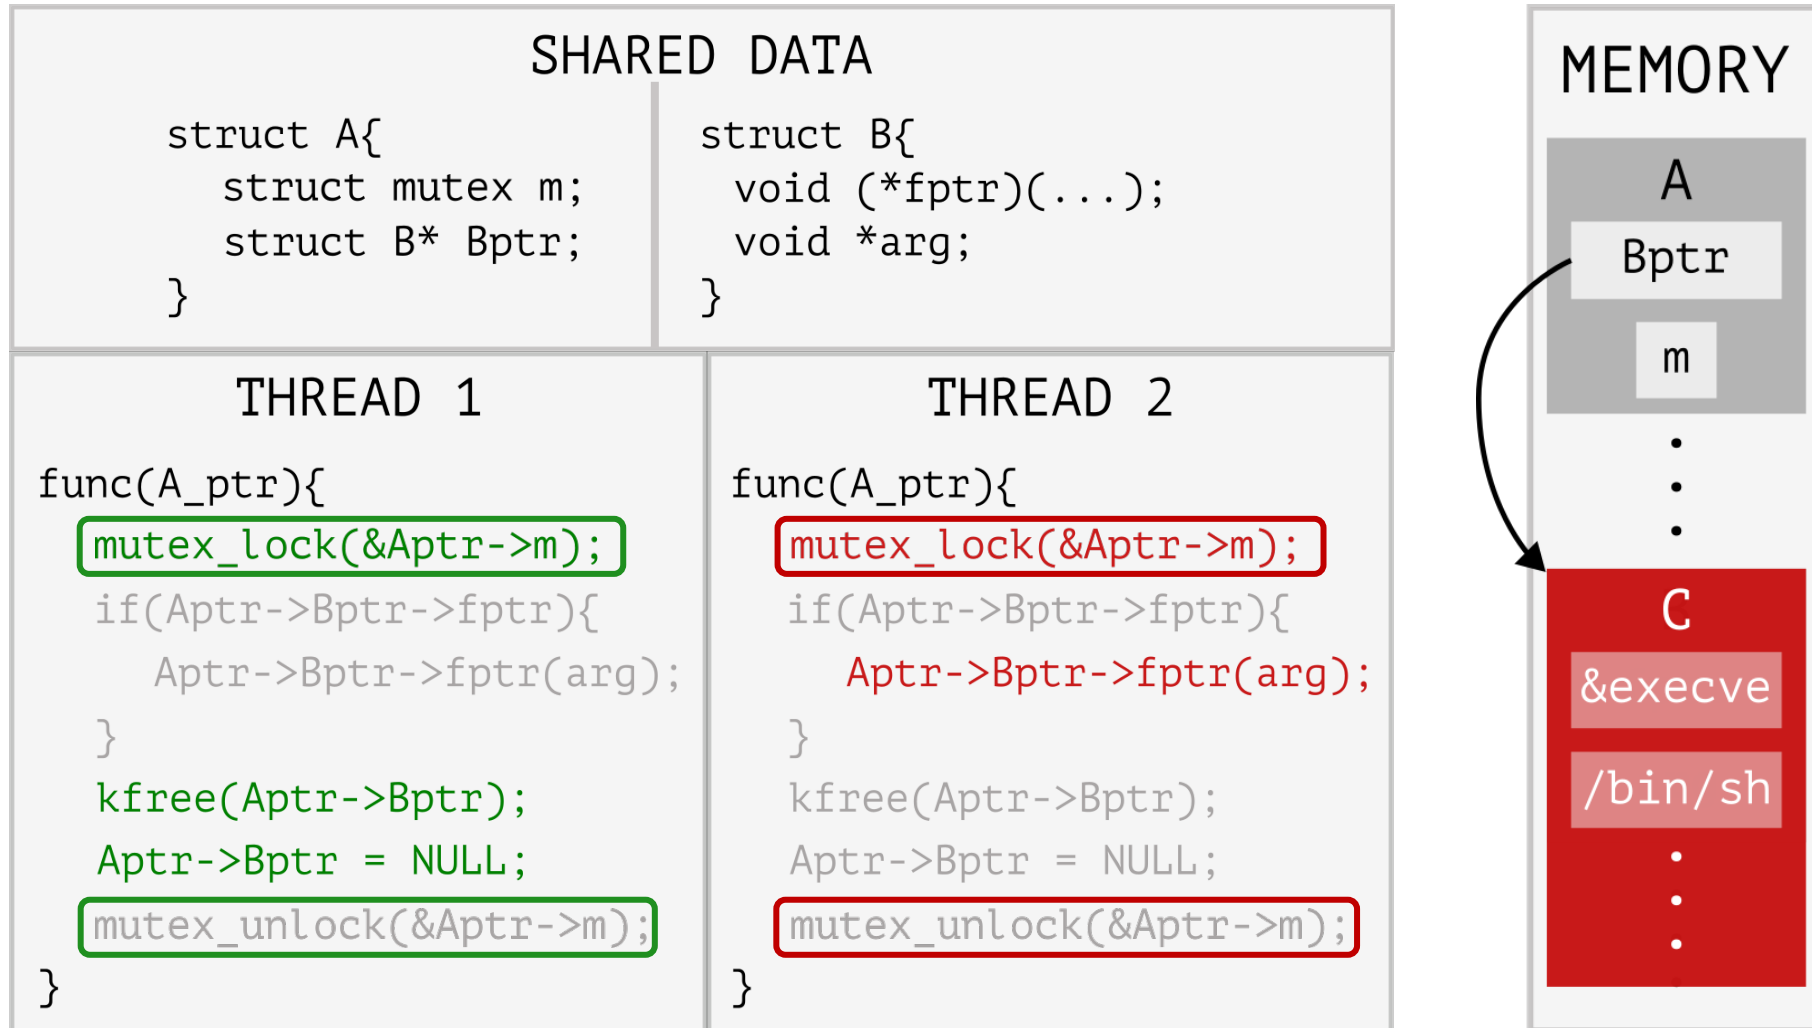
Background

Concurrency Bugs

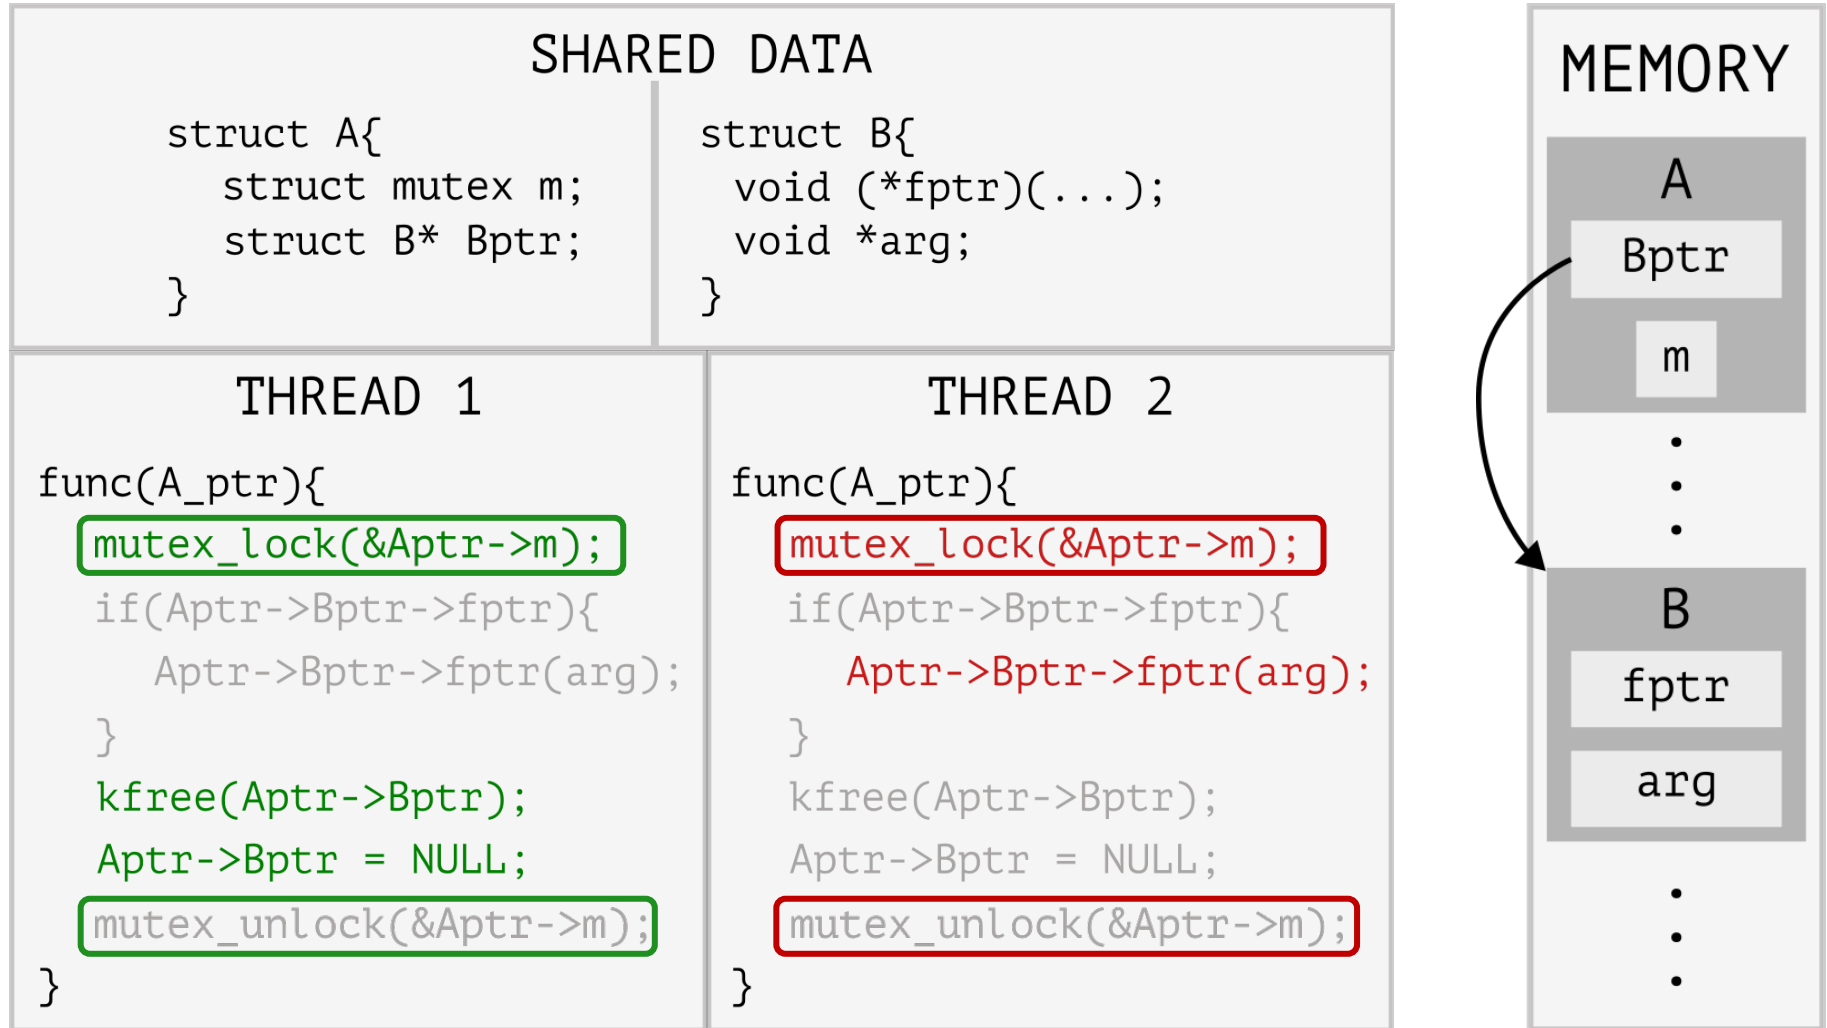


Background

Concurrency Bugs



Race Free Critical Region



Real Example

NFC Gadget

<pre>struct nfc_hci_dev { ... struct mutex msg_tx_mutex; ... struct hci_msg* cmd_pending_msg; ... }</pre>		SHARED DATA	<pre>struct hci_msg { ... void (*cb)(...); void *cb_context; ... }</pre>
THREAD 1	THREAD 2		
<pre>nfc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>	<pre>nfc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>		

**How do synchronization
primitives behave during
speculative execution?**

How do synchronization primitives behave during speculative execution?

What are the security implications for modern operating systems?

Objective

<pre>struct nfc_hci_dev { ... struct mutex msg_tx_mutex; ... struct hci_msg* cmd_pending_msg; ... }</pre>		SHARED DATA	<pre>struct hci_msg { ... void (*cb)(...); void *cb_context; ... }</pre>
THREAD 1	<pre>nfc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>	THREAD 2	<pre>nfc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>

Objective

Speculatively Hijack The Control-Flow

```
struct nfc_hci {  
    ...  
    struct mutex msg_tx_mutex;  
    ...  
    struct hci_dev *hdev;  
    ...  
};  
... {  
    ...);  
    context;
```

THREAD 1

```
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

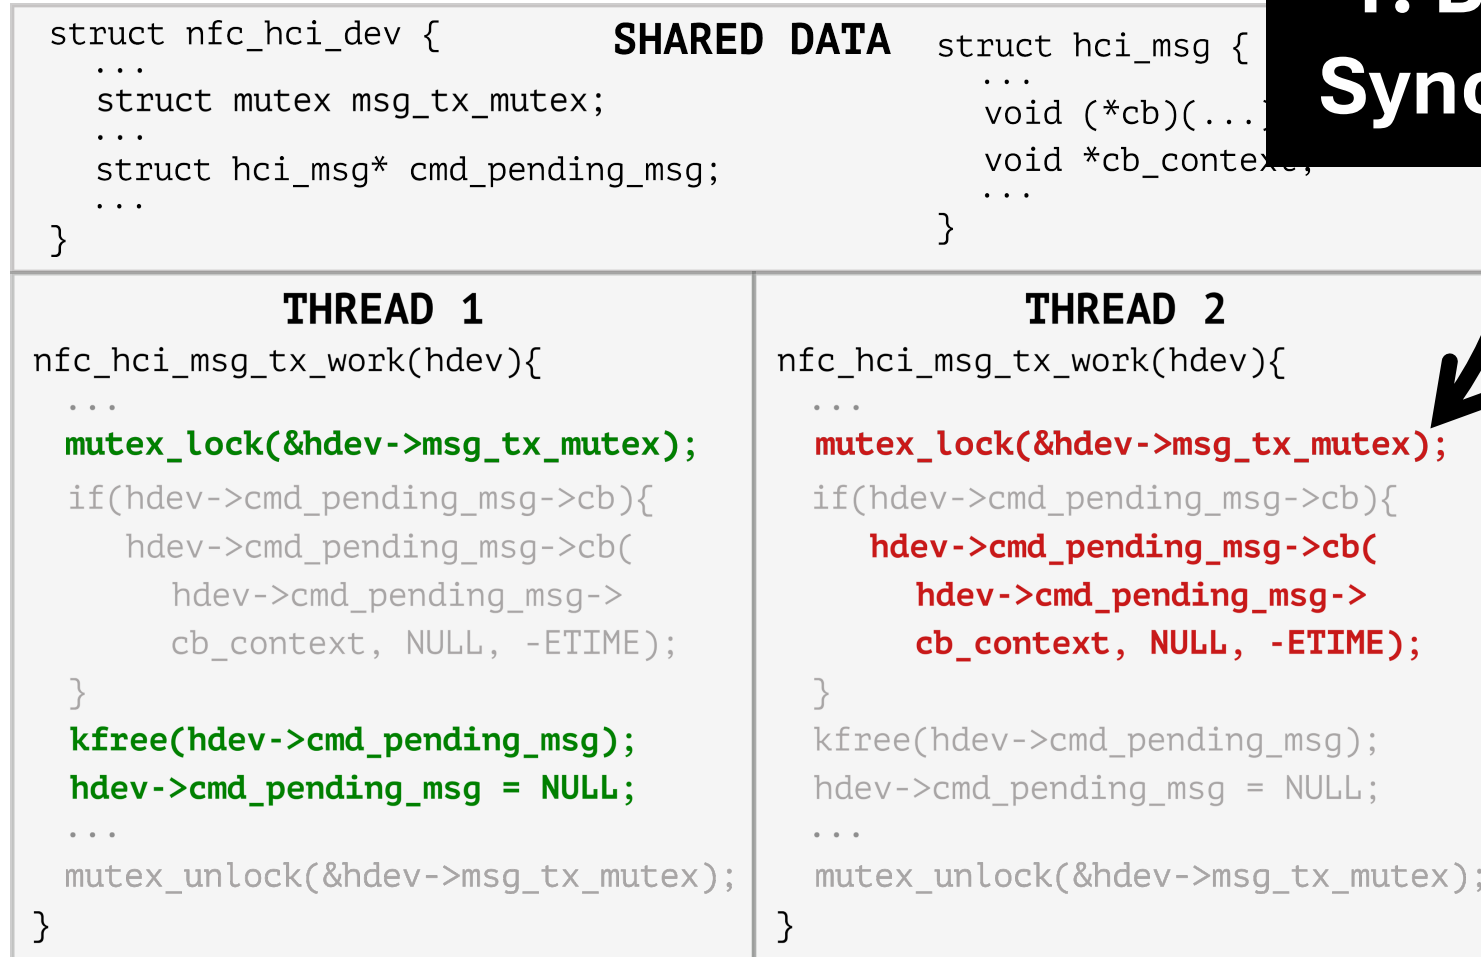
THREAD 2

```
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

UAF Attack Challenges

<pre>struct nfc_hci_dev { ... struct mutex msg_tx_mutex; ... struct hci_msg* cmd_pending_msg; ... }</pre>		SHARED DATA	<pre>struct hci_msg { ... void (*cb)(...); void *cb_context; ... }</pre>
THREAD 1	THREAD 2		
<pre>nfc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>	<pre>nfc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>		

UAF Attack Challenges



1. Bypass The Sync. Primitive

UAF Attack Challenges

1. Bypass The Sync. Primitive

2. Create an Exploitation Window

```
struct nfc_hci_dev {  
    ...  
    struct mutex msg_tx_mutex;  
    ...  
    struct hci_msg* cmd_pending_msg;  
    ...  
}  
                                     SHARED DATA  
struct hci_msg {  
    ...  
    void (*cb)(...);  
    void *cb_context;  
    ...  
}
```

```
THREAD 1  
fc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

```
THREAD 2  
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

UAF Attack Challenges

1. Bypass The Sync. Primitive

3. Reallocate The Freed Memory With a Malicious Object

2. Create an Exploitation Window

```
struct nfc_hci {
    ...
    struct mutex msg_tx_mutex;
    ...
    struct hci_dev *hdev;
    ...
}
```

```
msg {
    ...
    void (*cb)(...);
    void *cb_context;
}
```

```
fc_hci_msg_tx_work(hdev){
    ...
    mutex_lock(&hdev->msg_tx_mutex);
    if(hdev->cmd_pending_msg->cb){
        hdev->cmd_pending_msg->cb(
            hdev->cmd_pending_msg->
            cb_context, NULL, -ETIME);
    }
    kfree(hdev->cmd_pending_msg);
    hdev->cmd_pending_msg = NULL;
    ...
    mutex_unlock(&hdev->msg_tx_mutex);
}
```

```
nfc_hci_msg_tx_work(hdev){
    ...
    mutex_lock(&hdev->msg_tx_mutex);
    if(hdev->cmd_pending_msg->cb){
        hdev->cmd_pending_msg->cb(
            hdev->cmd_pending_msg->
            cb_context, NULL, -ETIME);
    }
    kfree(hdev->cmd_pending_msg);
    hdev->cmd_pending_msg = NULL;
    ...
    mutex_unlock(&hdev->msg_tx_mutex);
}
```

UAF Attack Challenges

1. Bypass The Sync. Primitive

3. Reallocate The Freed Memory With a Malicious Object

2. Create an Exploitation Window

```
struct nfc_hci {  
    ...  
    struct mutex msg_tx_mutex;  
    ...  
    struct hci_msg_tx_work msg_tx_work;  
    ...  
};
```

```
fc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

```
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

4. Hijack The Control-Flow

UAF Attack Challenges

```
struct nfc_hci_dev {  
    ...  
    struct mutex msg_tx_mutex;  
    ...  
    struct hci_msg* cmd_pending_msg;  
    ...  
}  
                                     SHARED DATA  
struct hci_msg {  
    ...  
    void (*cb)(...);  
    void *cb_context;  
    ...  
}
```

2. Create an Exploitation Window

THREAD 1

```
fc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

THREAD 2

```
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

1. Bypass The Sync. Primitive

UAF Attack Challenges

1. Bypass The Sync. Primitive



```
struct nfc_hci_dev {  
    ...  
    struct mutex msg_tx_mutex;  
    ...  
    struct hci_msg* cmd_pending_msg;  
    ...  
}  
                                     SHARED DATA  
struct hci_msg {  
    ...  
    void (*cb)(...);  
    void *cb_context;  
    ...  
}
```

THREAD 1

```
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

THREAD 2

```
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

Synchronization Primitives

Synchronization Primitives

Mutex lock on x86 arch:

- Conditional branch

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
```

Synchronization Primitives

Mutex lock on x86 arch:

- Conditional branch
- A series of inlined function calls

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
            ↳ arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)
                ↳ arch_atomic_try_cmpxchg_acquire(&lock, , ...)
                    ↳ arch_atomic_try_cmpxchg(&lock, , ...)
                        ↳ arch_try_cmpxchg(&lock, , ...)
                            ↳ __raw_try_cmpxchg(ptr, ...){
```

Synchronization Primitives

Mutex lock on x86 arch:

- Conditional branch
- A series of inlined function calls
- Atomic compare and exchange

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
            ↳arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)
                ↳arch_atomic_try_cmpxchg_acquire(&lock, , ...)
                    ↳arch_atomic_try_cmpxchg(&lock, , ...)
                        ↳arch_try_cmpxchg(&lock, , ...)
                            ↳__raw_try_cmpxchg(ptr, ...){
                                asm volatile(
                                    "lock cmpxchgq %2, %1"
                                    : "=a" (ret), "+m" (*ptr)
                                    : "r" (new), "0" (old)
                                    : "memory"
                                );
                            })
            return true;
    ...
}
```

Synchronization Primitives

Mutex lock on x86 arch:

- Conditional branch
- A series of inlined function calls
- Atomic compare and exchange

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
            ↳ arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)
                ↳ arch_atomic_try_cmpxchg_acquire(&lock, , ...)
                    ↳ arch_atomic_try_cmpxchg(&lock, , ...)
                        ↳ arch_try_cmpxchg(&lock, , ...)
                            ↳ __raw_try_cmpxchg(ptr, ...){
                                asm volatile(
                                    "lock cmpxchgq %2, %1"
                                    : "=a" (ret), "+m" (*ptr)
                                    : "r" (new), "0" (old)
                                    : "memory"
                                );
                            })
            return true;
    ...
}
```

Synchronization Primitives

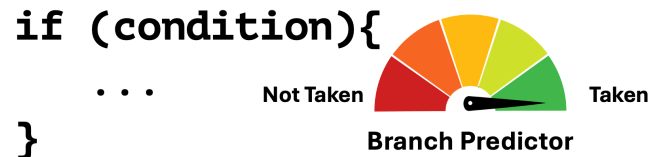
Mutex lock on x86 arch:

- Conditional branch
- A series of inlined function calls
- Atomic compare and exchange

Atomic != Serializing

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
            ↳ arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)
                ↳ arch_atomic_try_cmpxchg_acquire(&lock, , ...)
                    ↳ arch_atomic_try_cmpxchg(&lock, , ...)
                        ↳ arch_try_cmpxchg(&lock, , ...)
                            ↳ __raw_try_cmpxchg(ptr, ...){
                                asm volatile(
                                    "lock cmpxchgq %2, %1"
                                    : "=a" (ret), "+m" (*ptr)
                                    : "r" (new), "0" (old)
                                    : "memory"
                                );
                            }
            }
}
```

Synchronization Primitives



Mutex lock

- Conditional branch
- A series of inlined function calls
- Atomic compare and exchange

```
void mutex_lock(struct mutex *lock){  
    ...  
    if (!__mutex_trylock_fast(lock))  
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))  
            ↳ arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)  
                ↳ arch_atomic_try_cmpxchg_acquire(&lock, , ...)  
                    ↳ arch_atomic_try_cmpxchg(&lock, , ...)  
                        ↳ arch_try_cmpxchg(&lock, , ...)  
                            ↳ __raw_try_cmpxchg(ptr, ...){  
                                asm volatile(  
                                    "lock cmpxchgq %2, %1"  
                                    : "=a" (ret), "+m" (*ptr)  
                                    : "r" (new), "0" (old)  
                                    : "memory"  
                                );  
                            }  
            }  
}
```

Atomic != Serializing

Synchronization Primitives

`if (condition){`
...
`}`



Not Taken Taken
Branch Predictor

The diagram shows a semi-circular gauge with a needle pointing towards the 'Taken' side. The gauge is divided into segments of different colors (red, orange, yellow, green). The needle is positioned between the yellow and green segments, closer to the green segment.

Mutex lock

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
            ↳arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)
                ↳arch_atomic_try_cmpxchg_acquire(&lock, , ...)
                    ↳arch_atomic_try_cmpxchg(&lock, , ...)
                        ↳arch_try_cmpxchg(&lock, , ...)
                            ↳__raw_try_cmpxchg(ptr, ...){
                                asm volatile(
                                    "lock cmpxchgq %2, %1"
                                    : "=a" (ret), "+m" (*ptr)
                                    : "r" (new), "0" (old)
                                    : "memory"
                                );
                            }
            }
}
```

**Speculative
Race Condition**

UAF Attack Challenges

1. Bypass The Sync. Primitive:
Speculative Race Condition

<pre>struct nfc_hci_dev { ... struct mutex msg_tx_mutex; ... struct hci_msg* cmd_pending_msg; ... } SHARED DATA struct hci_r ... void (*cb ... void *cb_context, ... }</pre>	
<p>THREAD 1</p> <pre>nfc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>	<p>THREAD 2</p> <pre>nfc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>

UAF Attack Challenges

1. Bypass The Sync. Primitive:
Speculative Race Condition

2. Create an
Exploitation Window

SHARED DATA	
<pre>struct nfc_hci_dev { ... struct mutex msg_tx_mutex; ... struct hci_msg* cmd_pending_msg; ... }</pre>	<pre>struct hci_r ... void (*cb void *cb_ ... }</pre>
THREAD 1	THREAD 2
<pre>fc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>	<pre>nfc_hci_msg_tx_work(hdev){ ... mutex_lock(&hdev->msg_tx_mutex); if(hdev->cmd_pending_msg->cb){ hdev->cmd_pending_msg->cb(hdev->cmd_pending_msg-> cb_context, NULL, -ETIME); } kfree(hdev->cmd_pending_msg); hdev->cmd_pending_msg = NULL; ... mutex_unlock(&hdev->msg_tx_mutex); }</pre>

Creating UAF Exploitation Window

VICTIM CORE

----- User/Kernel Space Boundary

Creating UAF Exploitation Window

VICTIM CORE

1 `set_timer()`
`nfc_hci_msg_tx_work(hdev)`

----- User/Kernel Space Boundary

Creating UAF Exploitation Window

VICTIM CORE

1 `set_timer()`
`nfc_hci_msg_tx_work(hdev)`

```
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

----- User/Kernel Space Boundary

Creating UAF Exploitation Window

VICTIM CORE

1 `set_timer()`
`nfc_hci_msg_tx_work(hdev)`

```
nfc_hci_msg_tx_work(hdev){  
  ...  
  mutex_lock(&hdev->msg_tx_mutex);  
  if(hdev->cmd_pending_msg->cb){  
    hdev->cmd_pending_msg->cb(  
      hdev->cmd_pending_msg->  
      cb_context, NULL, -ETIME);  
  }  
  kfree(hdev->cmd_pending_msg); 2  
  hdev->cmd_pending_msg = NULL;  
  ...  
  mutex_unlock(&hdev->msg_tx_mutex);  
}
```

----- User/Kernel Space Boundary

Creating UAF Exploitation Window

VICTIM CORE

1 `set_timer()`
`nfc_hci_msg_tx_work(hdev)`

```
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

2

```
<nfc_hci_msg_tx_work>:  
...  
<kfree>:  
...  
    <__slab_free>:  
    ...  
    call <_raw_spin_lock_irqsave>  
    ...  
    call <list_del>  
    ...  
    call <_raw_spin_unlock_irqrestore>  
    lea rsp,[rbp-0x28]  
    pop rbx  
    pop r12  
    pop r13  
    pop r14  
    pop r15  
    pop rbp  
    ret  
...  
mov qword ptr [rax], 0  
...
```

----- User/Kernel Space Boundary

Creating UAF Exploitation Window

VICTIM CORE

1 `set_timer()`
`nfc_hci_msg_tx_work(hdev)`

```
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
            hdev->cmd_pending_msg->  
            cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

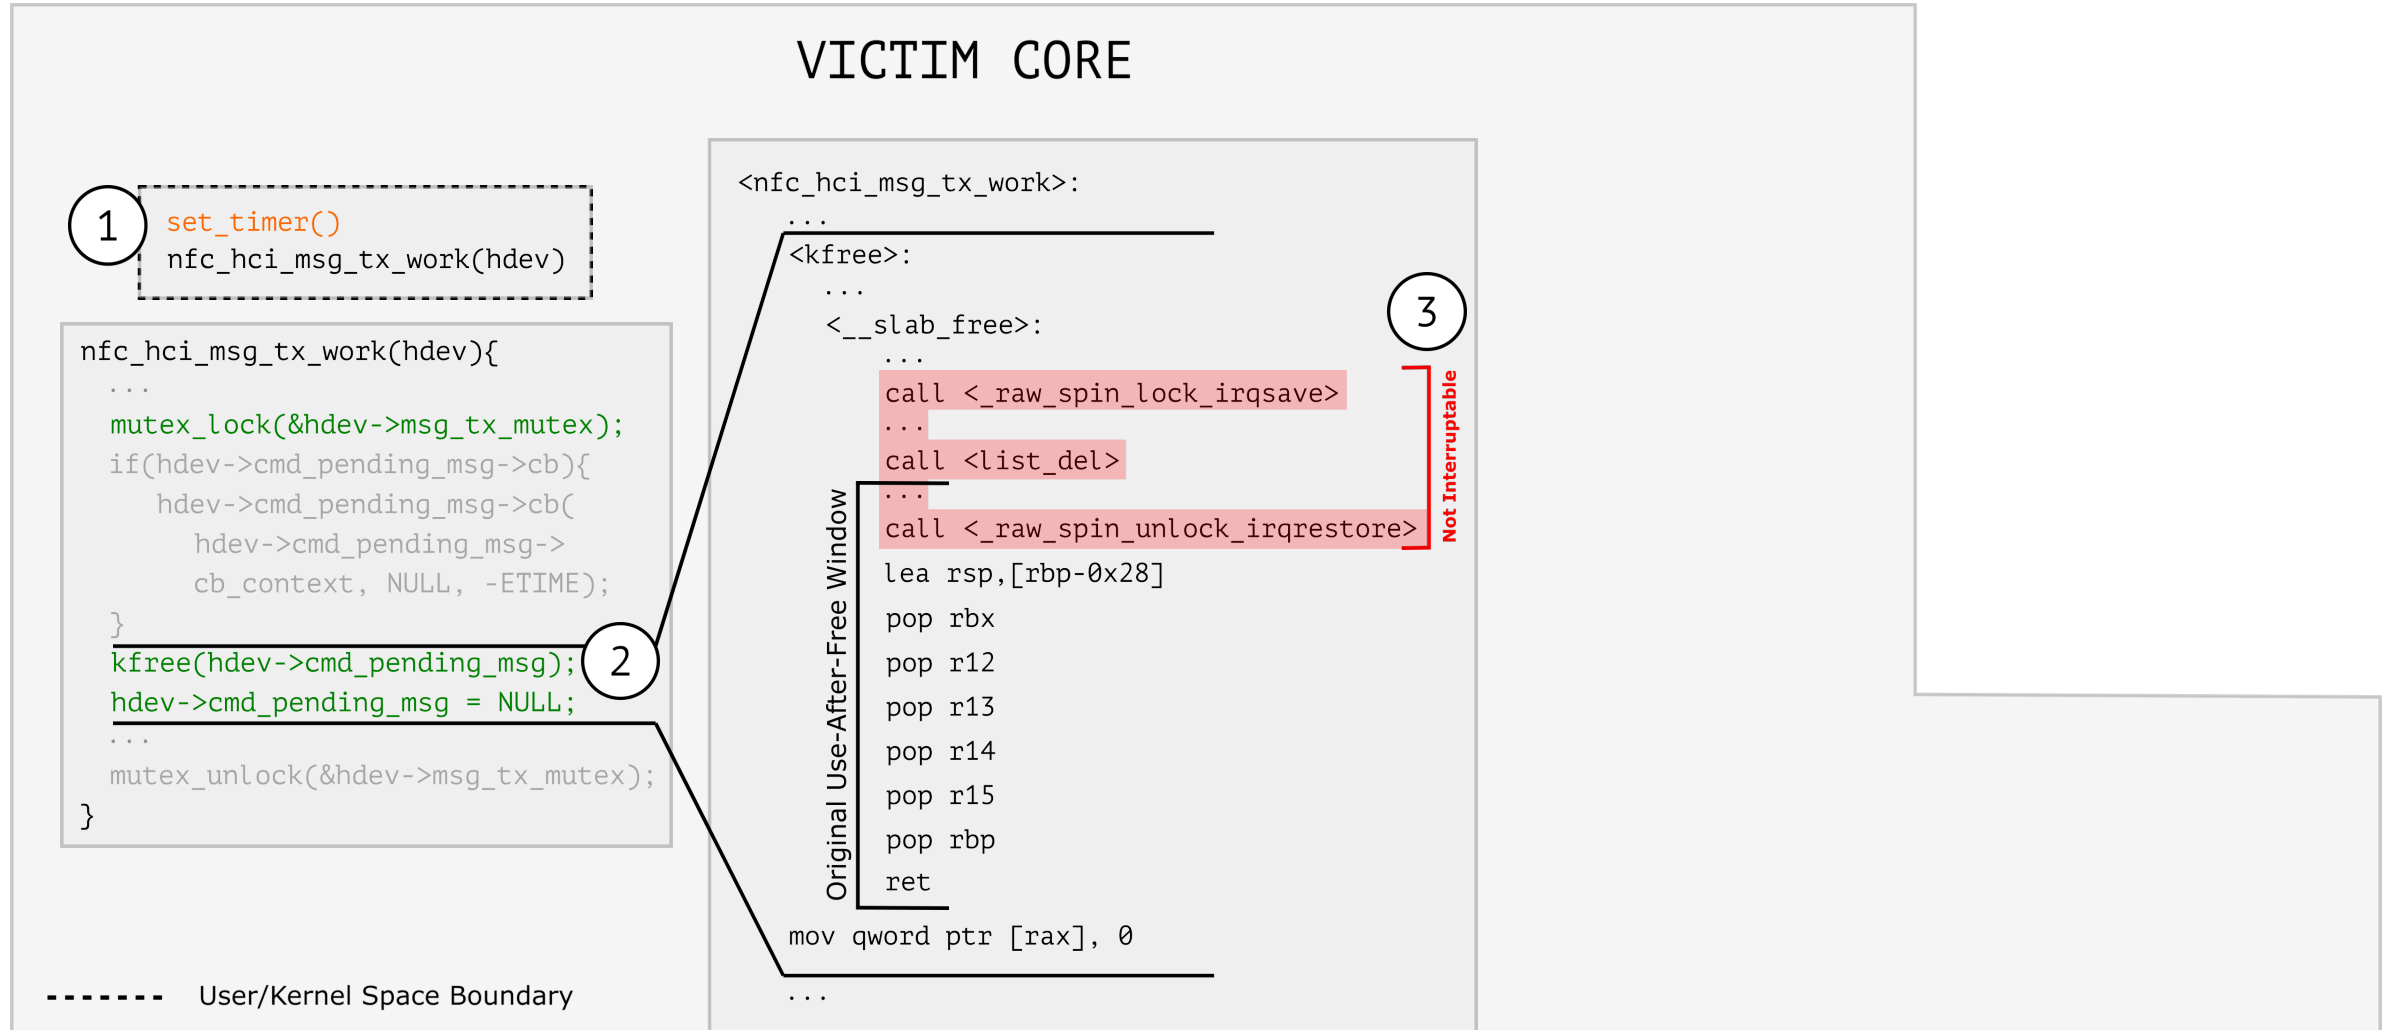
2

```
<nfc_hci_msg_tx_work>:  
...  
<kfree>:  
...  
    <__slab_free>:  
    ...  
        call <_raw_spin_lock_irqsave>  
    ...  
        call <list_del>  
    ...  
        call <_raw_spin_unlock_irqrestore>  
    lea rsp,[rbp-0x28]  
    pop rbx  
    pop r12  
    pop r13  
    pop r14  
    pop r15  
    pop rbp  
    ret  
...  
mov qword ptr [rax], 0  
...
```

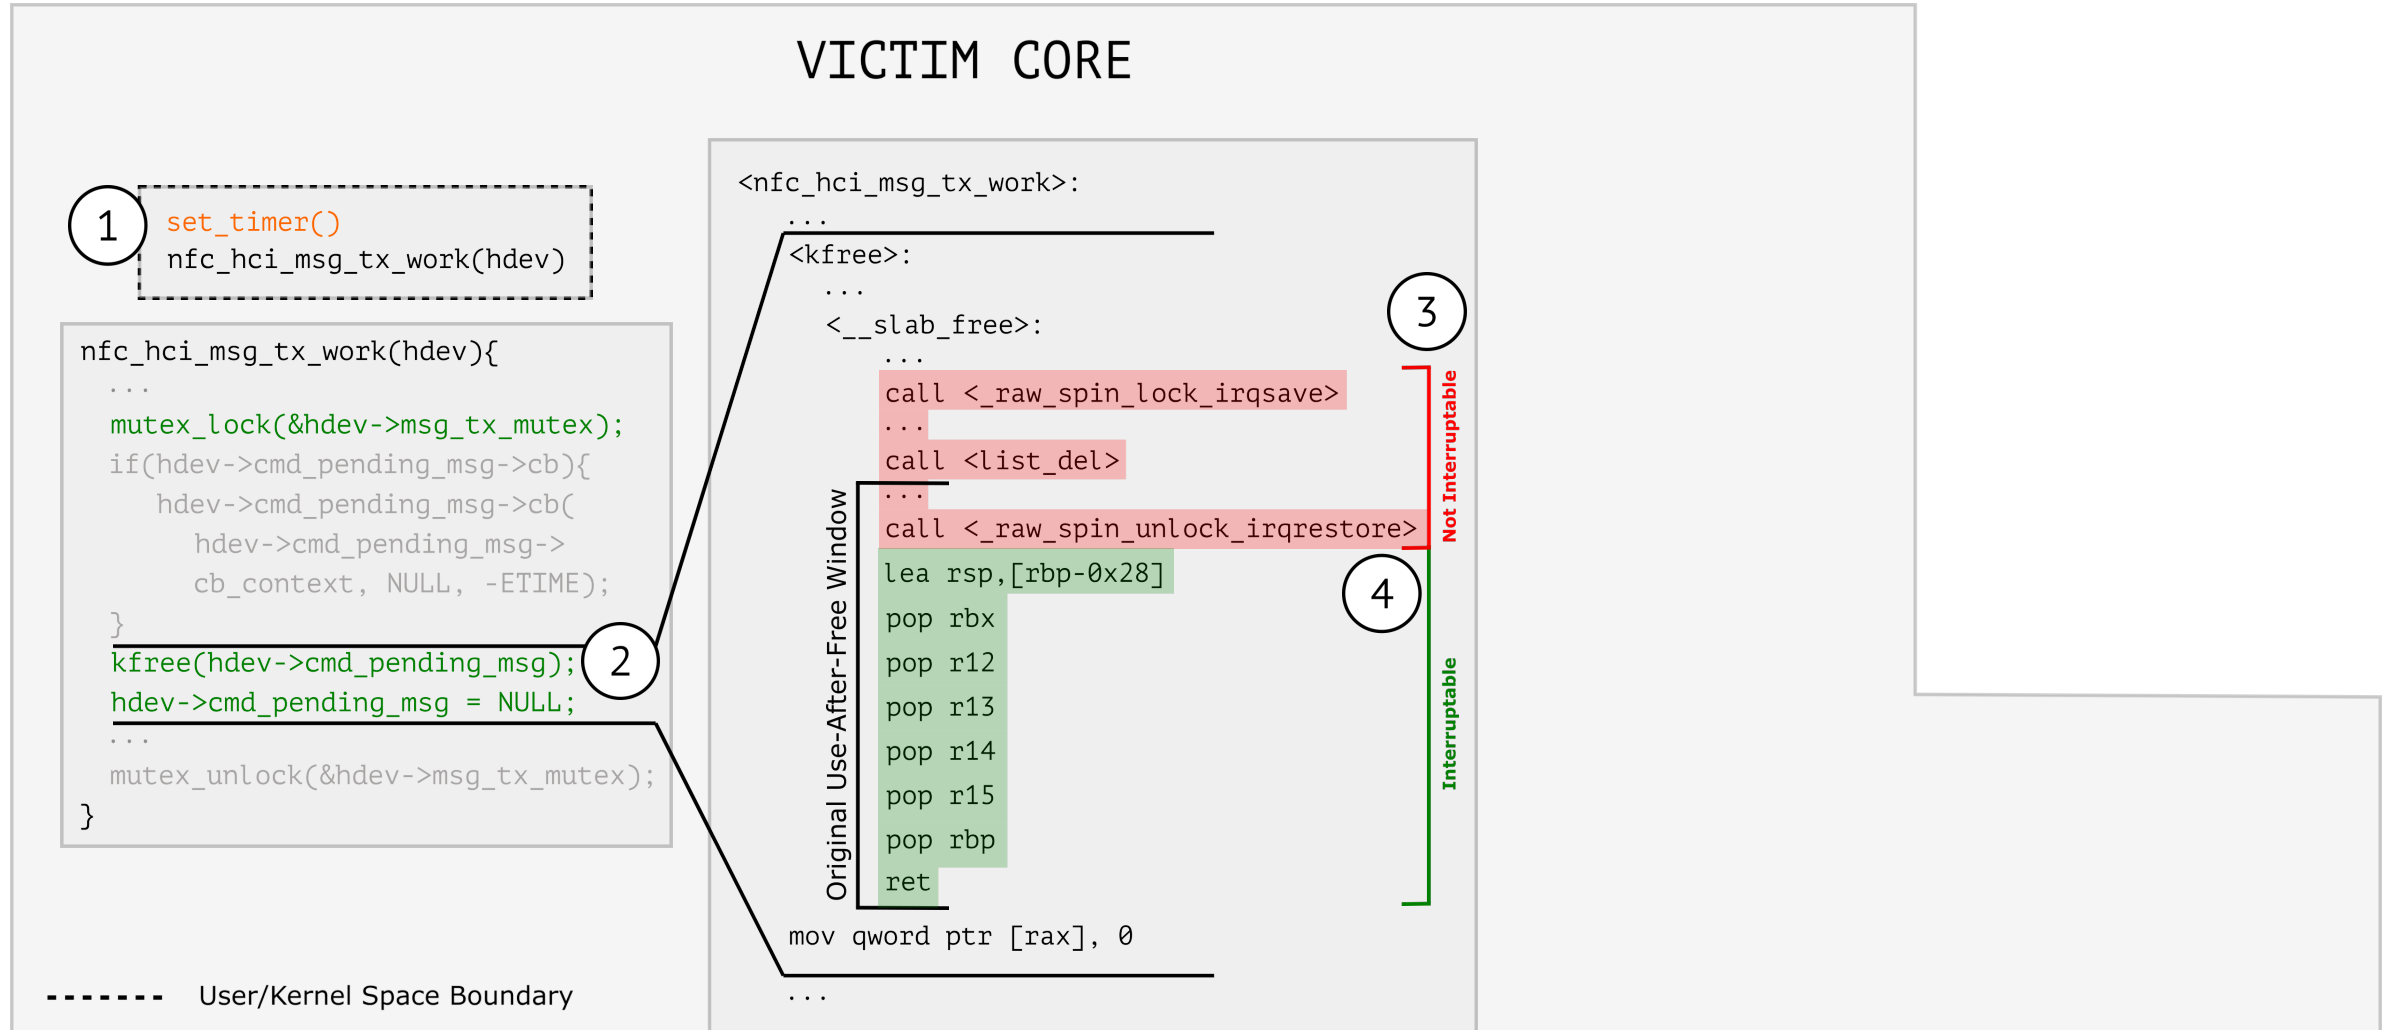
Original Use-After-Free Window

----- User/Kernel Space Boundary

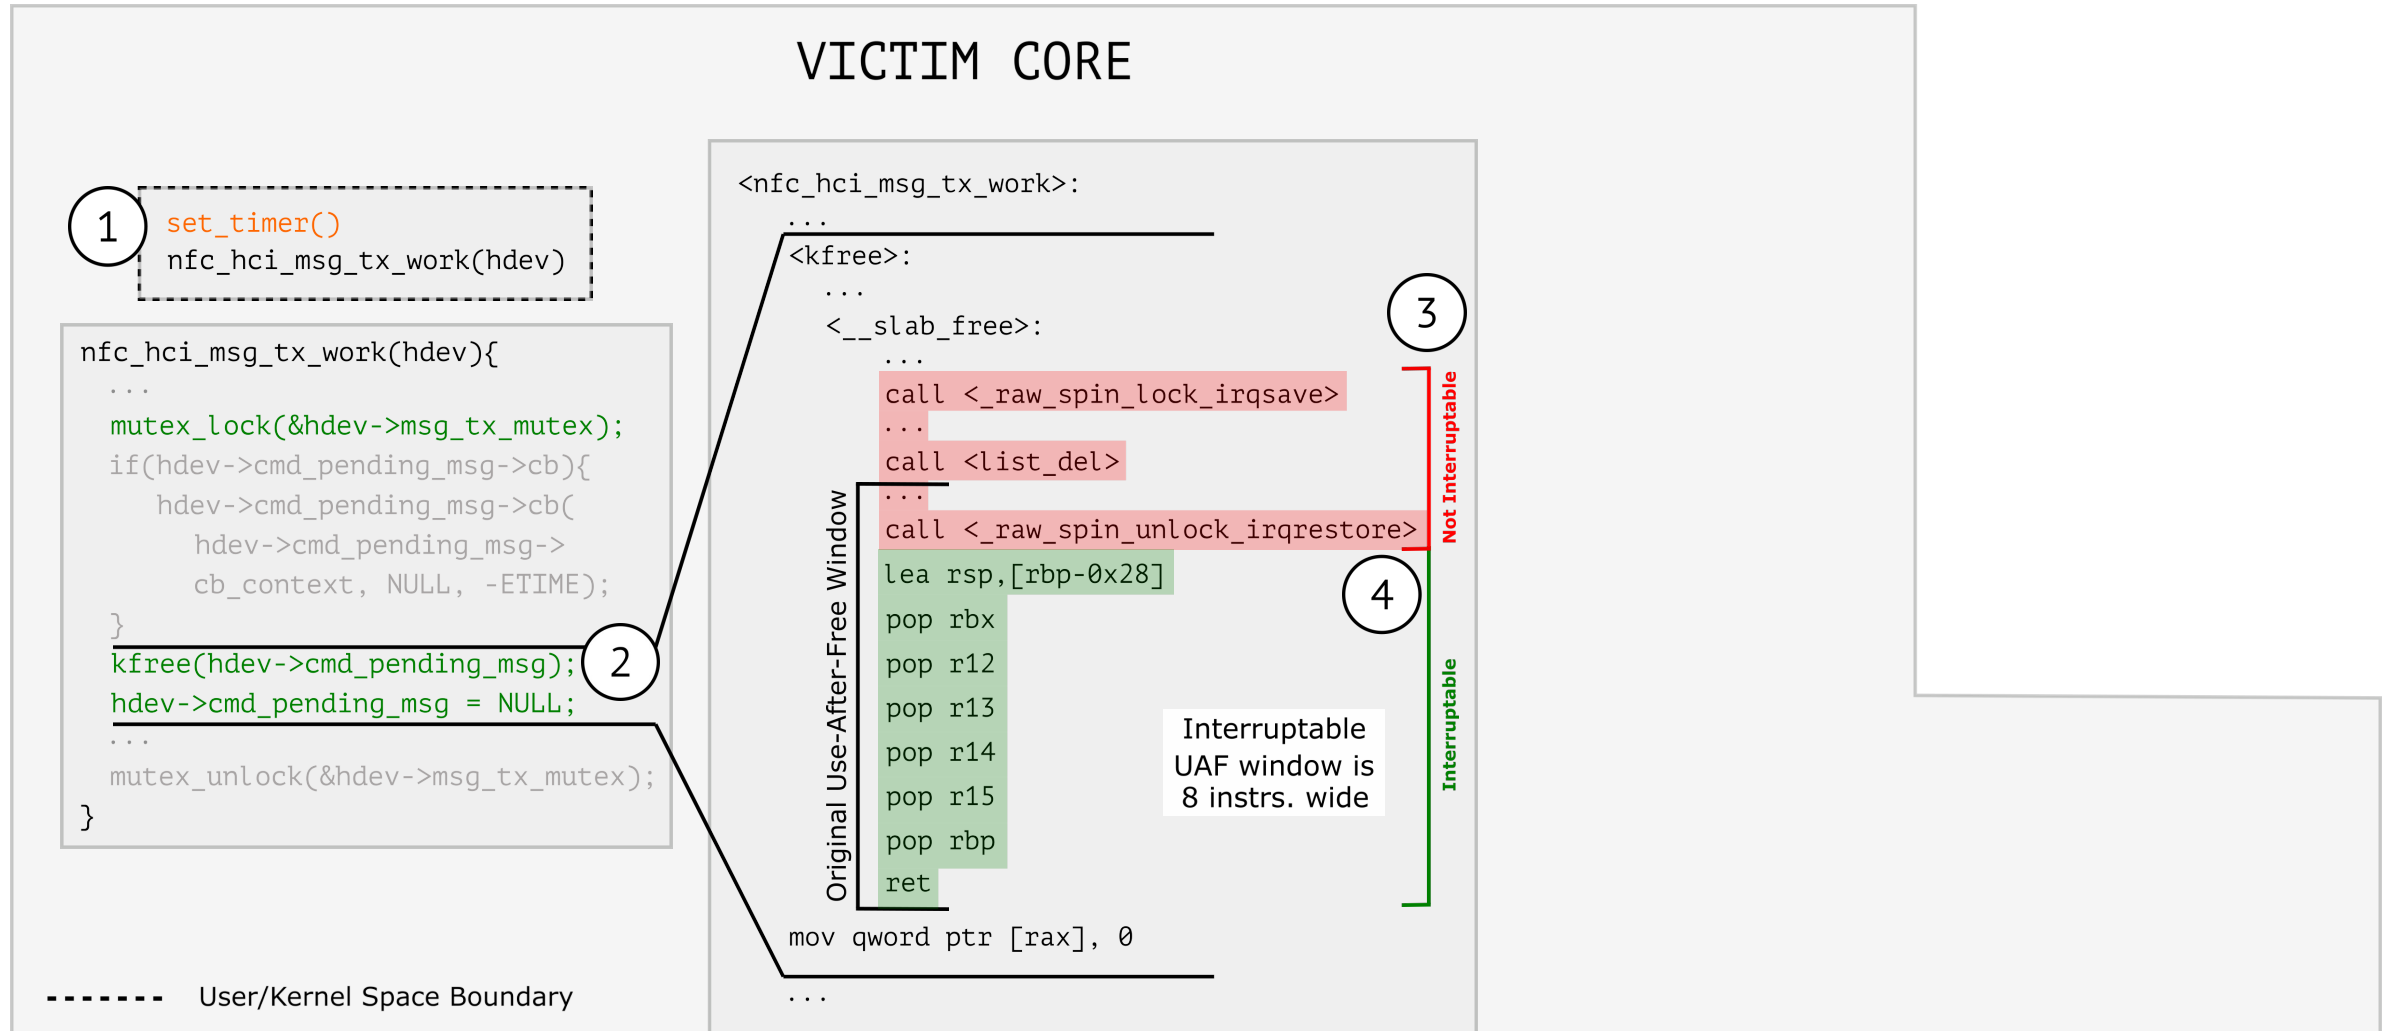
Creating UAF Exploitation Window



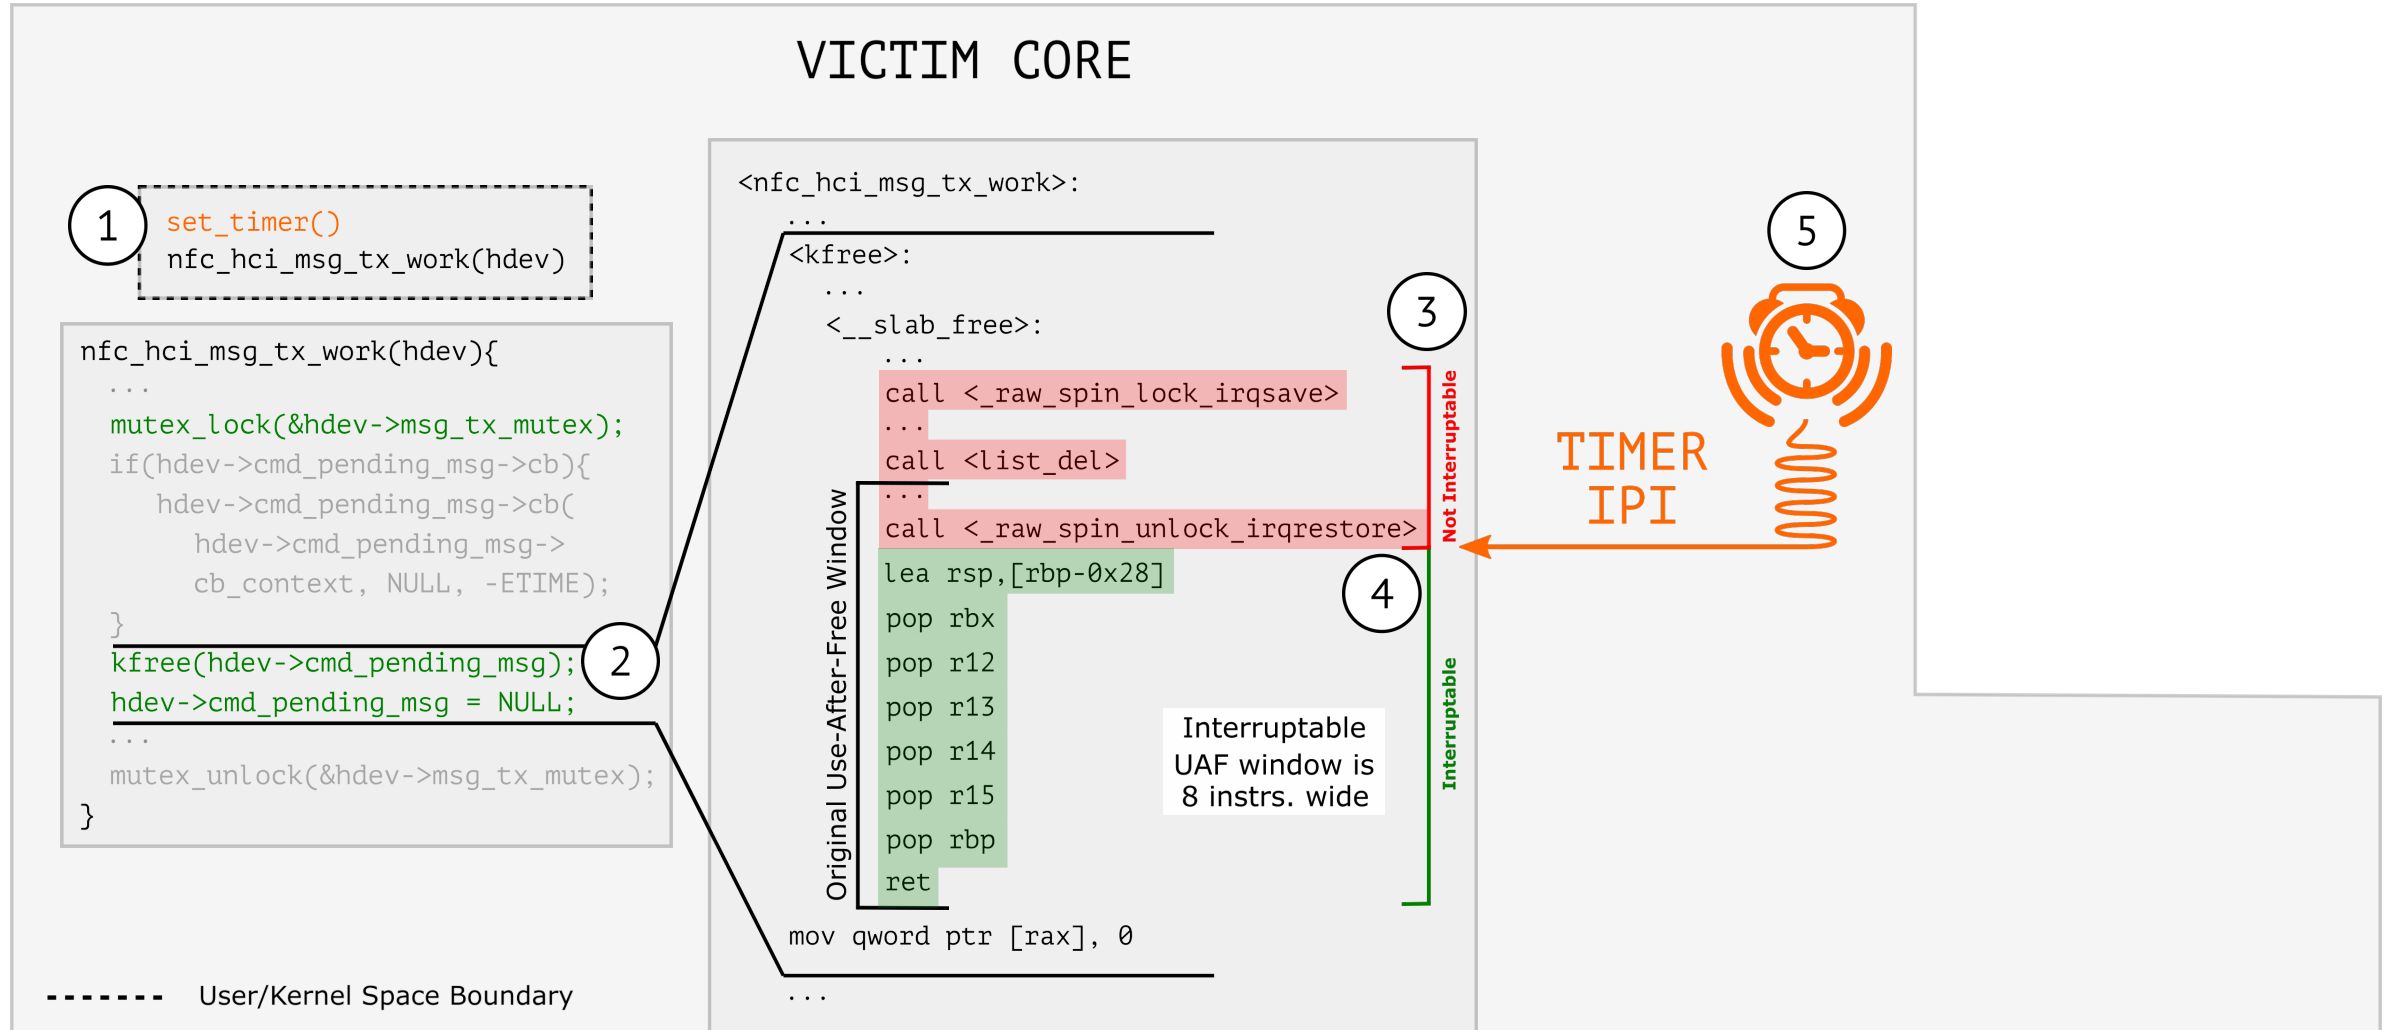
Creating UAF Exploitation Window



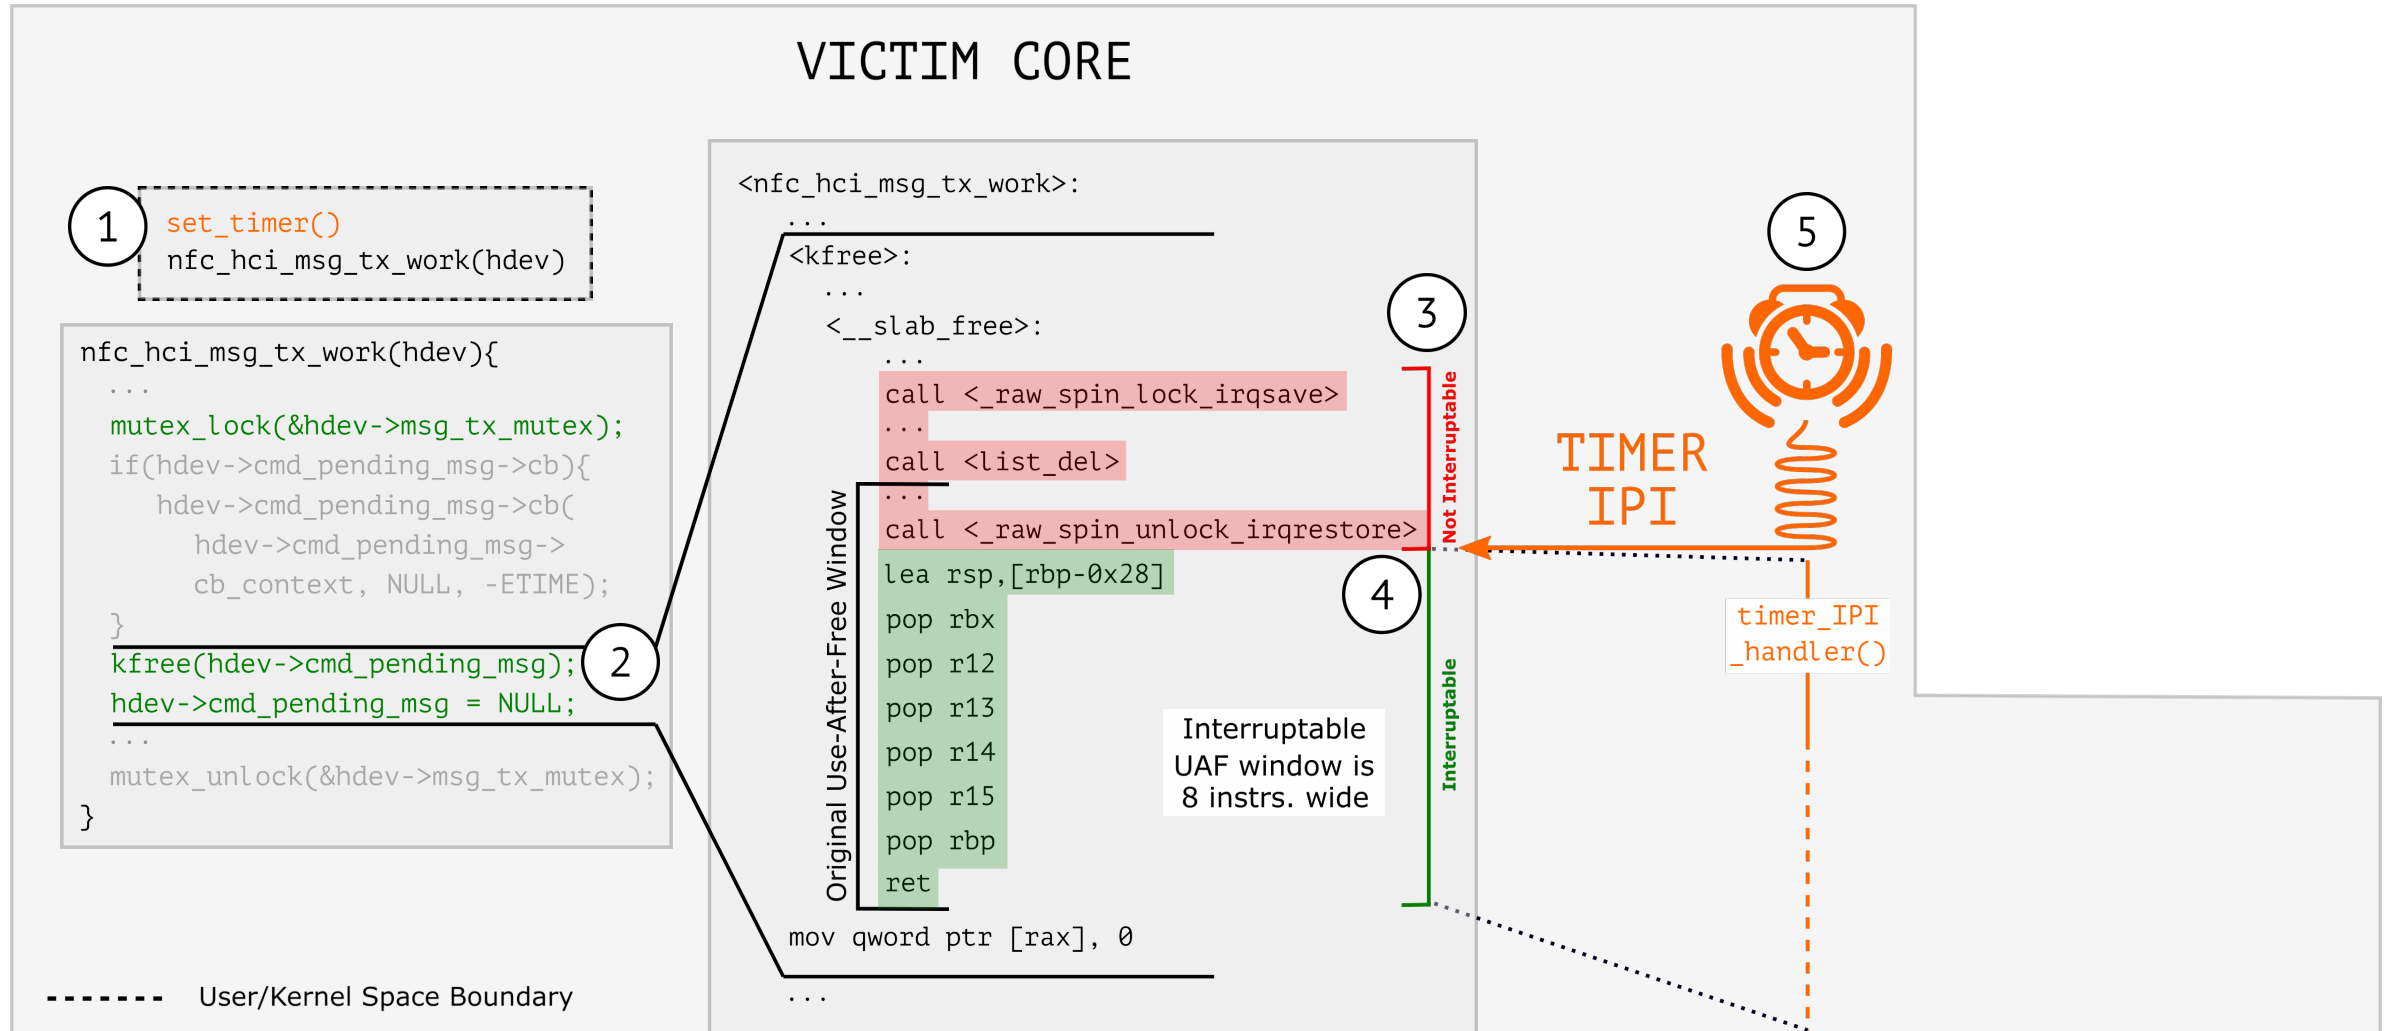
Creating UAF Exploitation Window



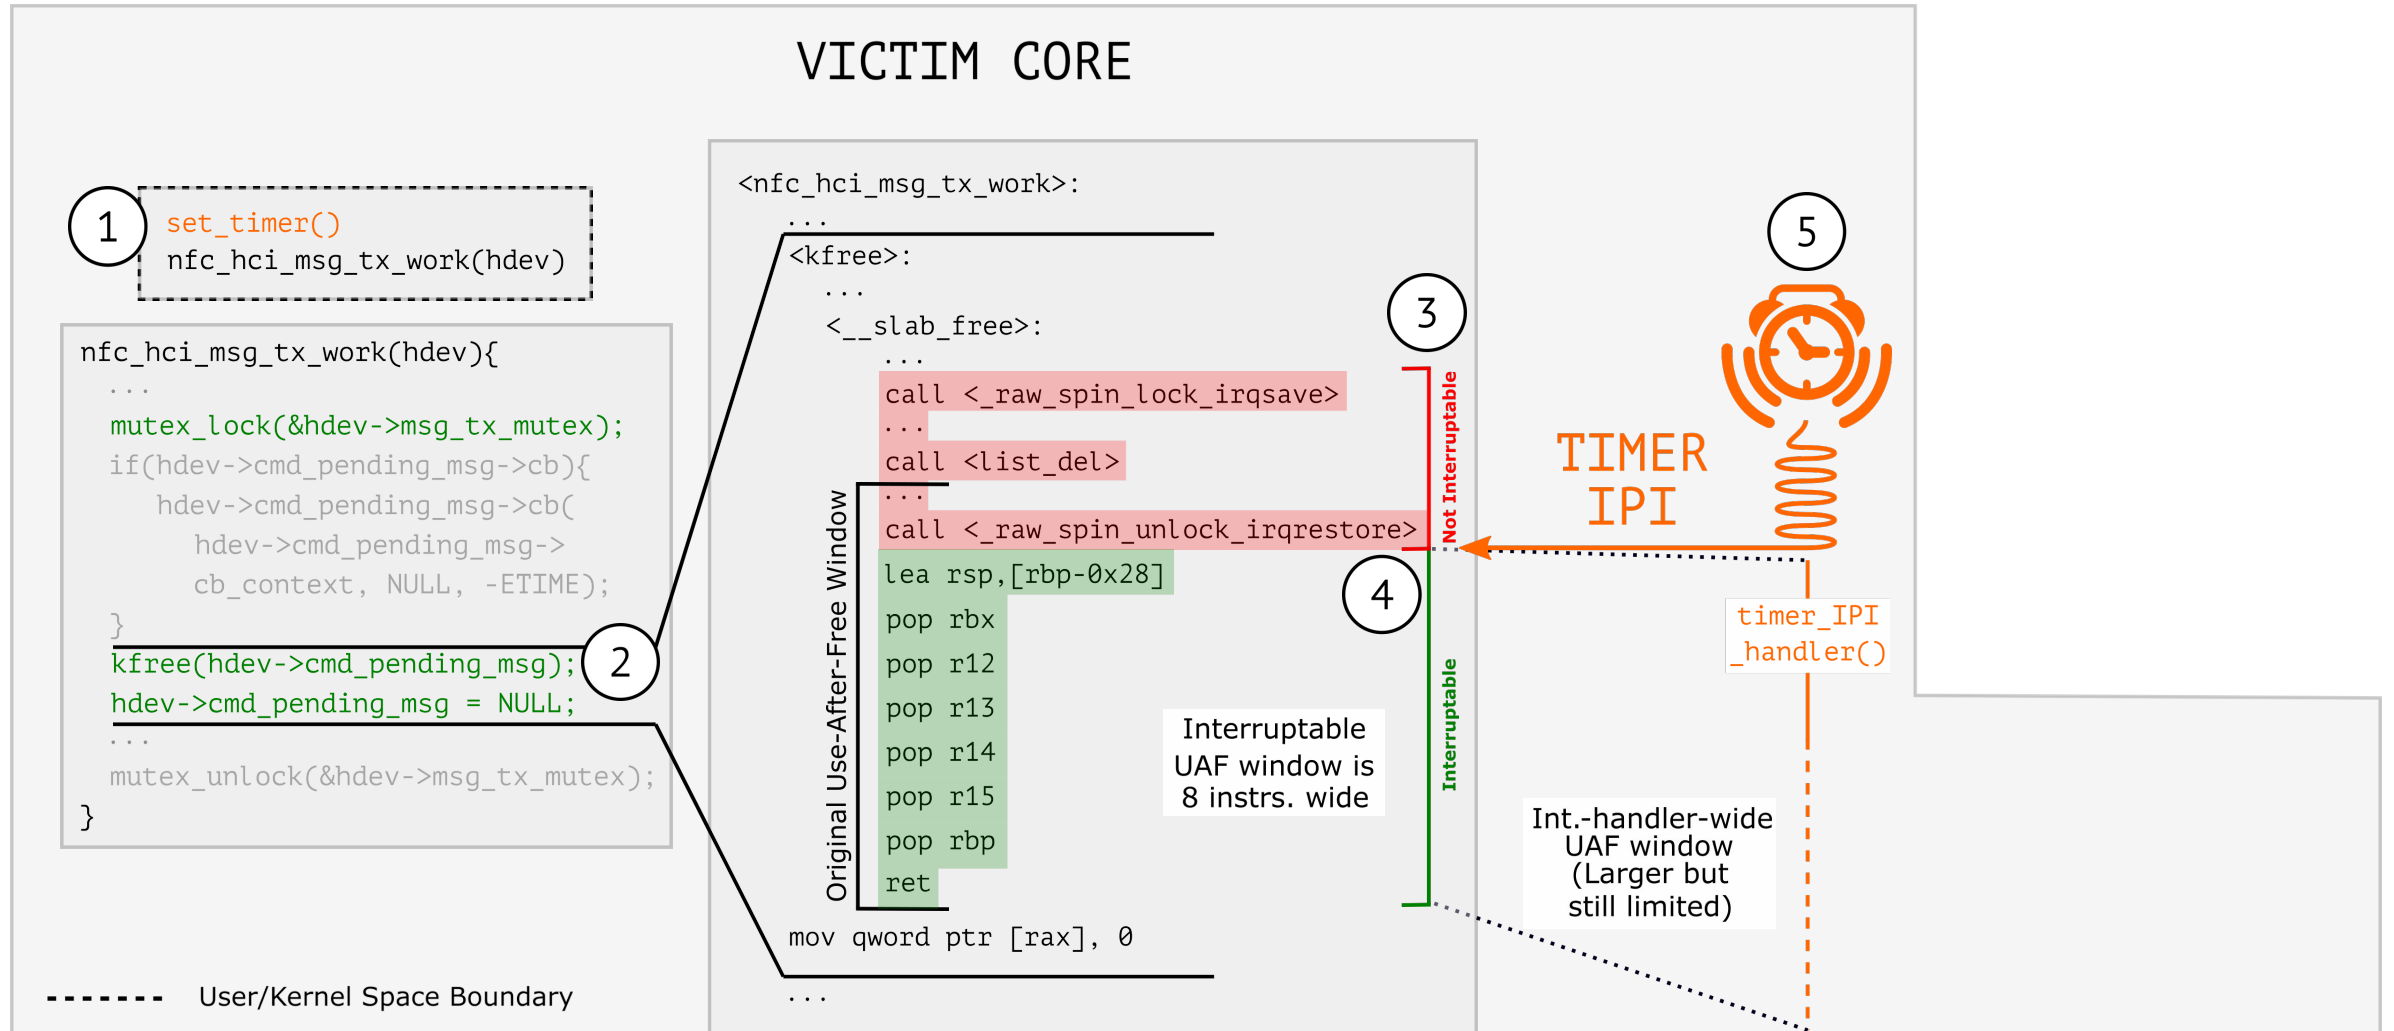
Creating UAF Exploitation Window



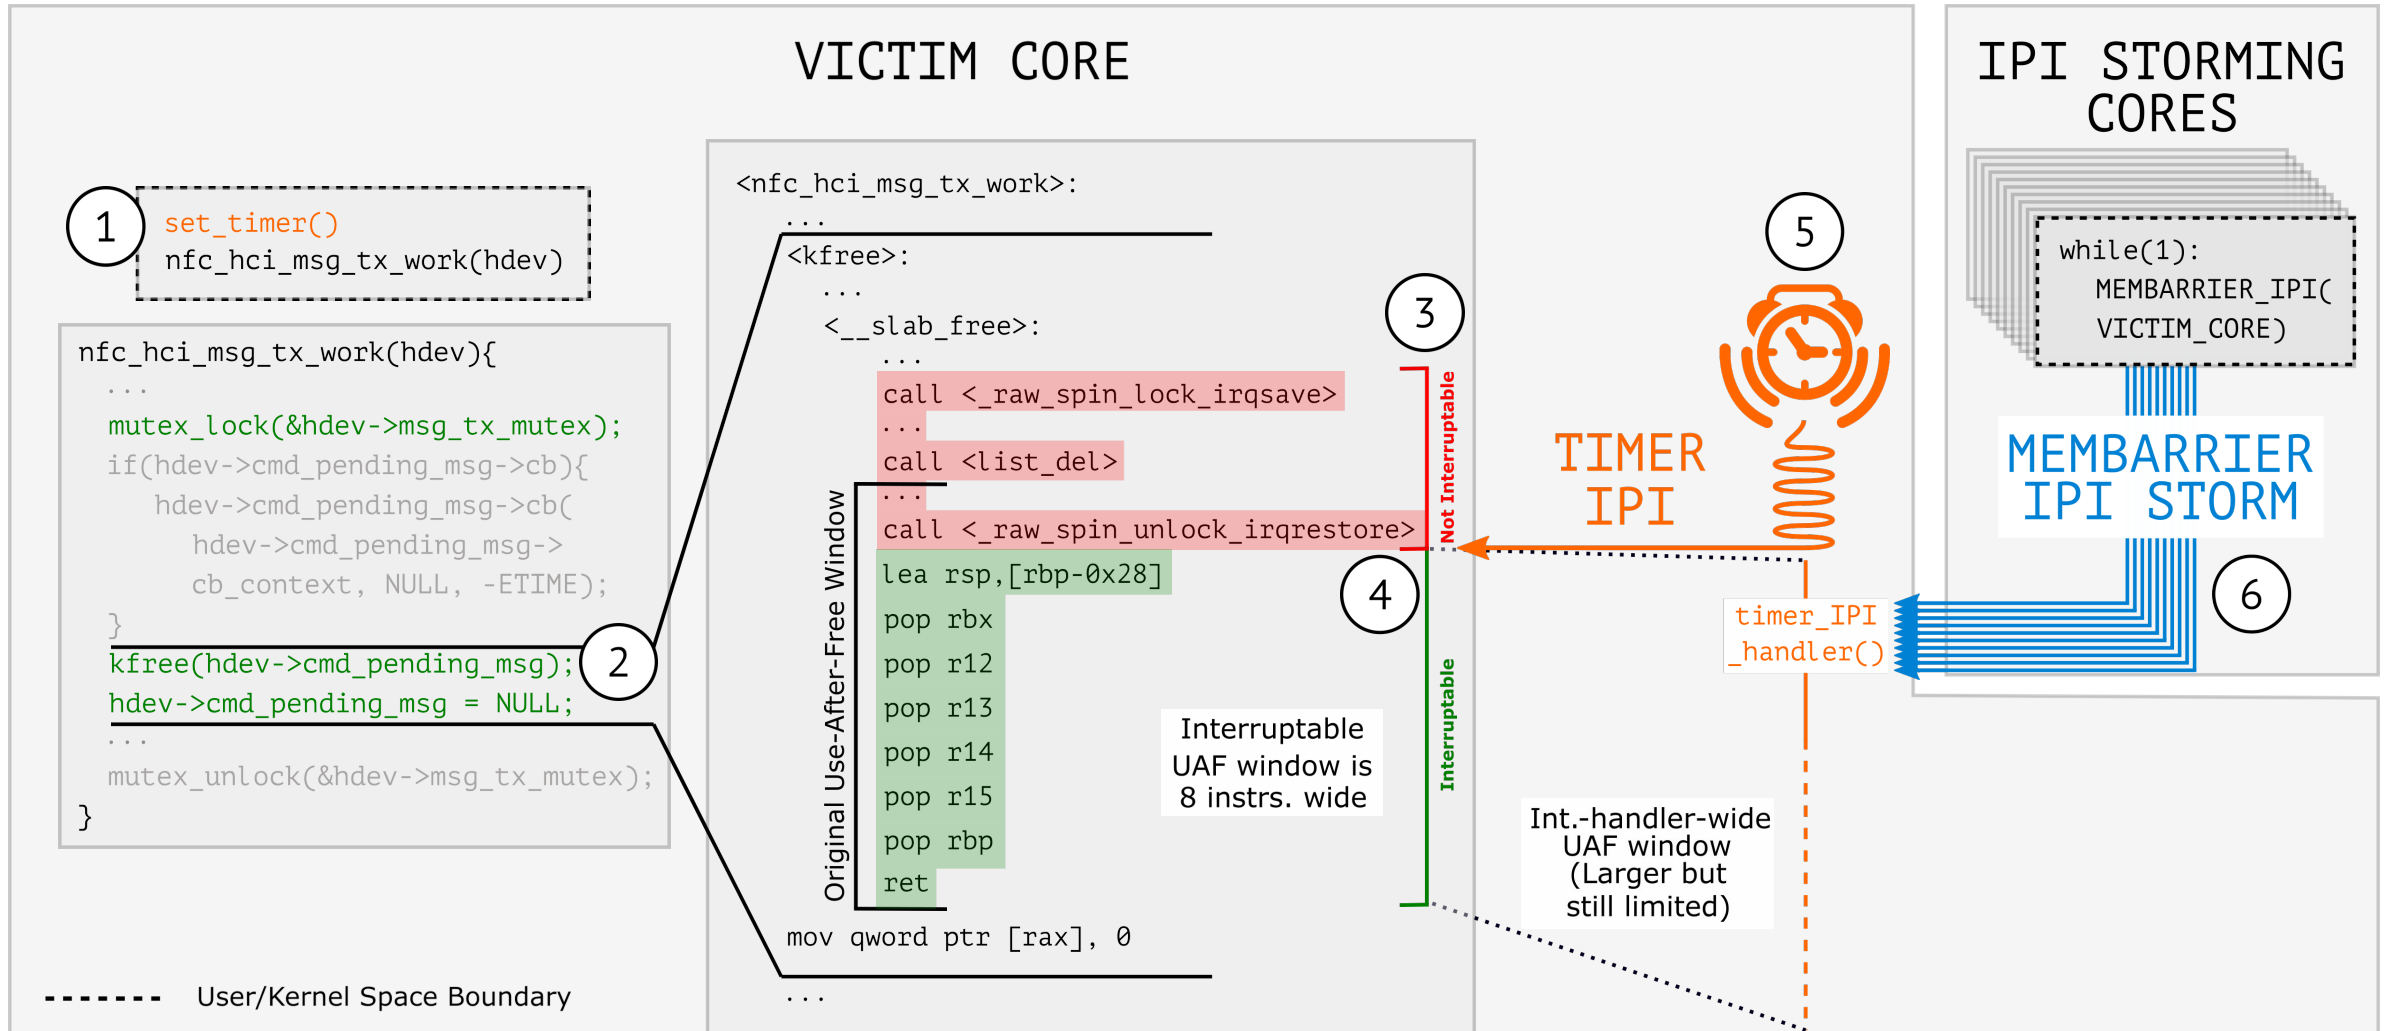
Creating UAF Exploitation Window



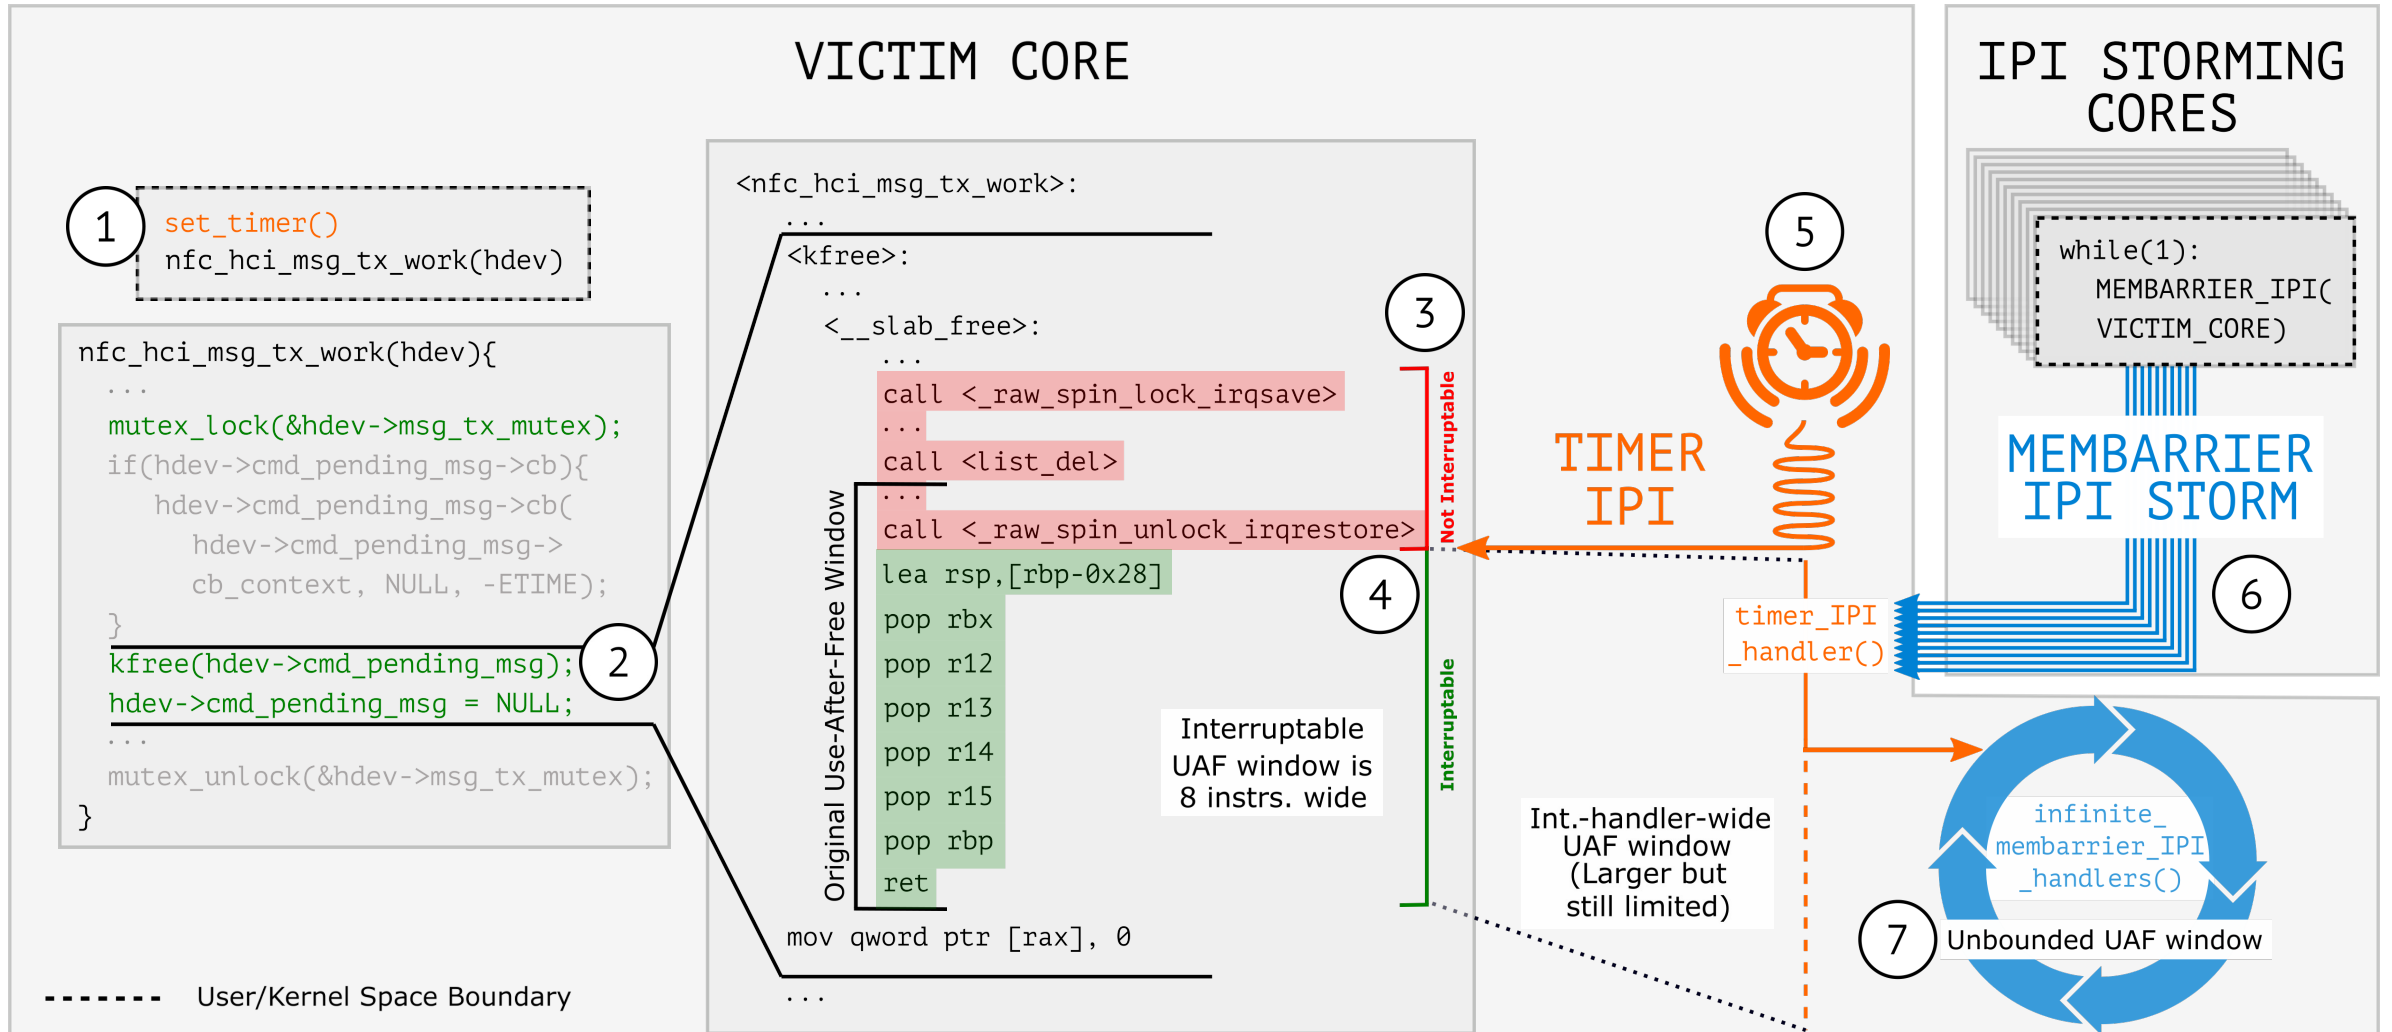
Creating UAF Exploitation Window



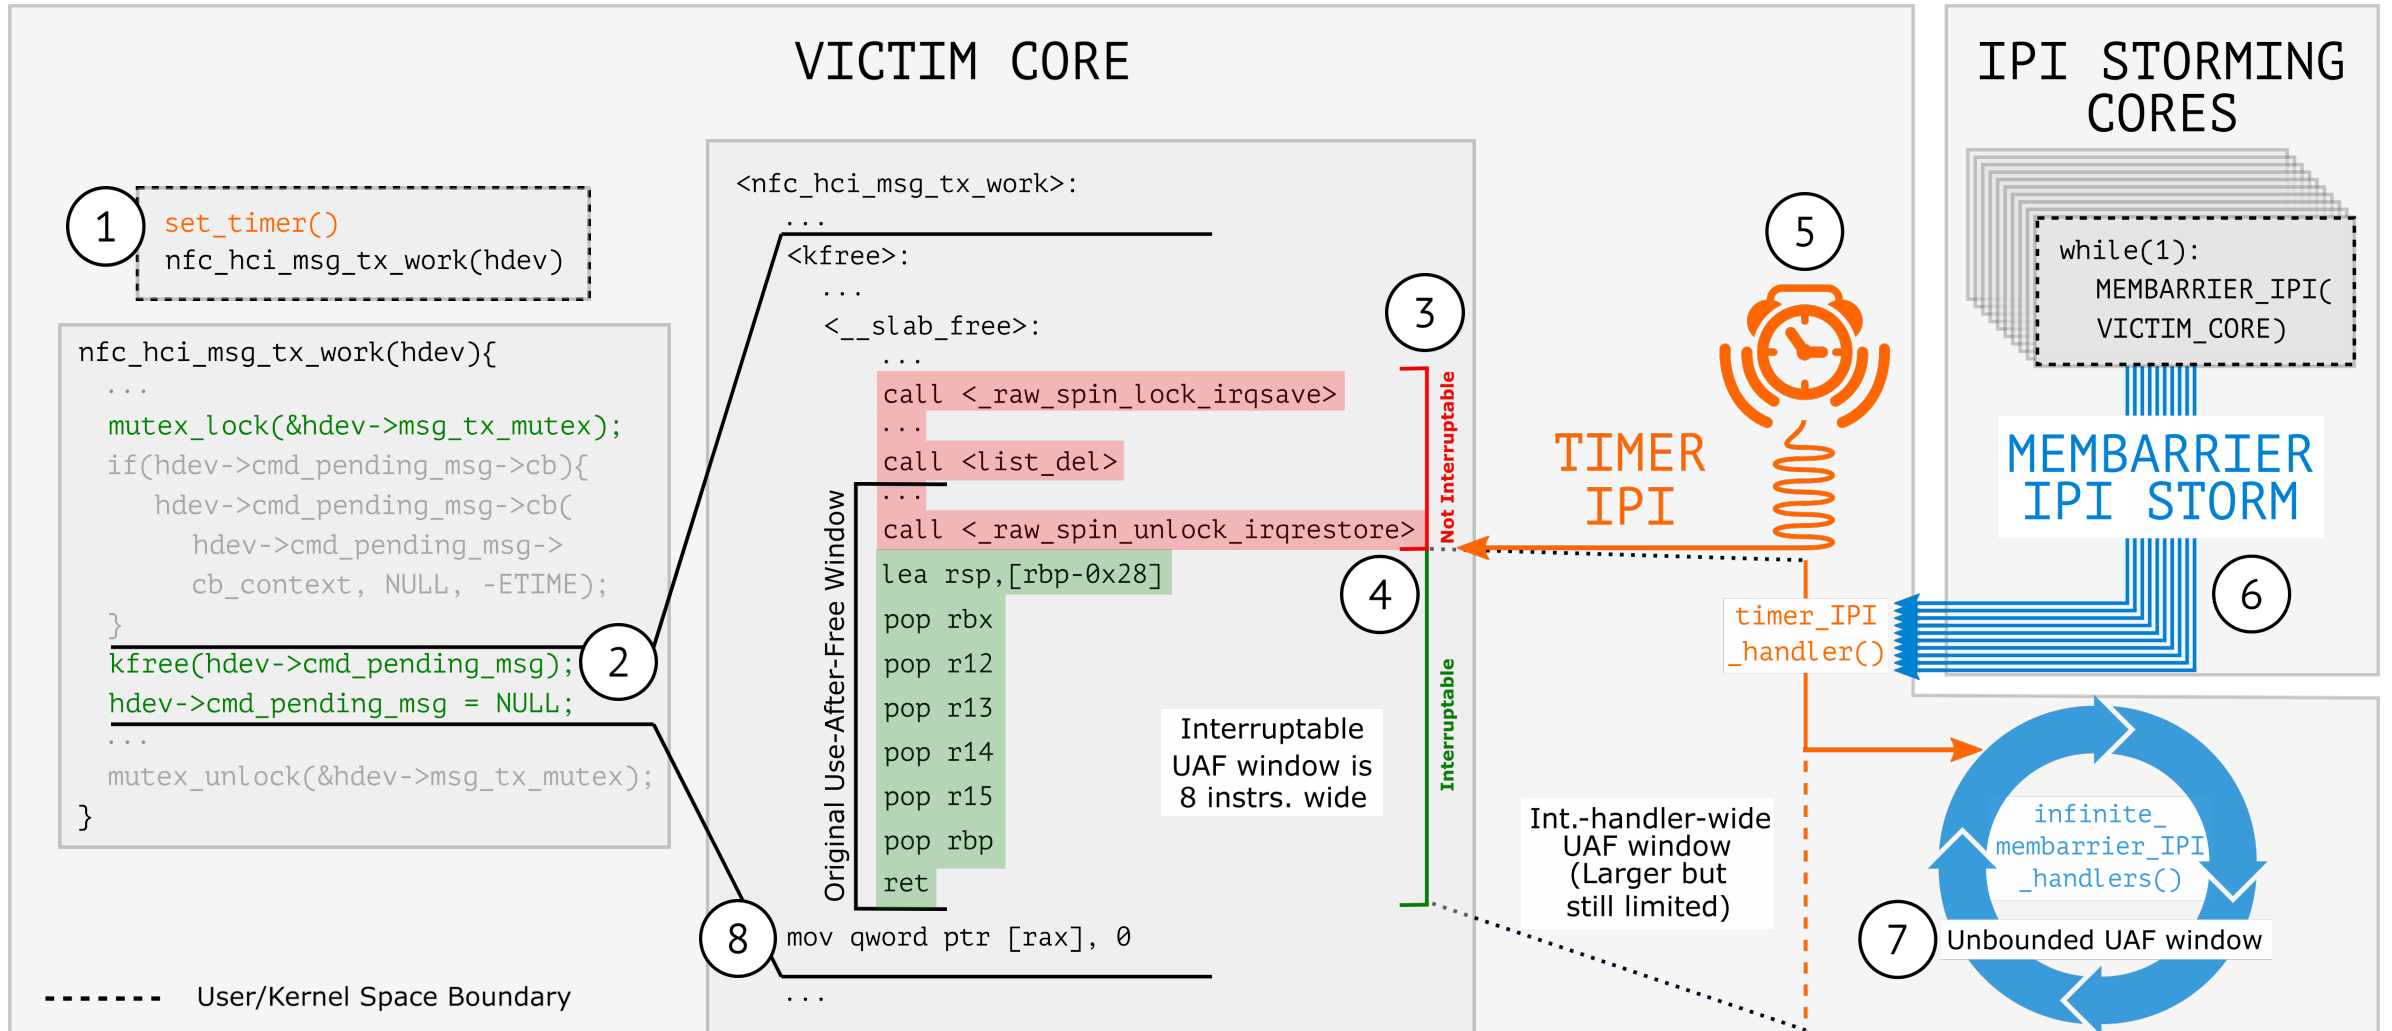
Creating UAF Exploitation Window



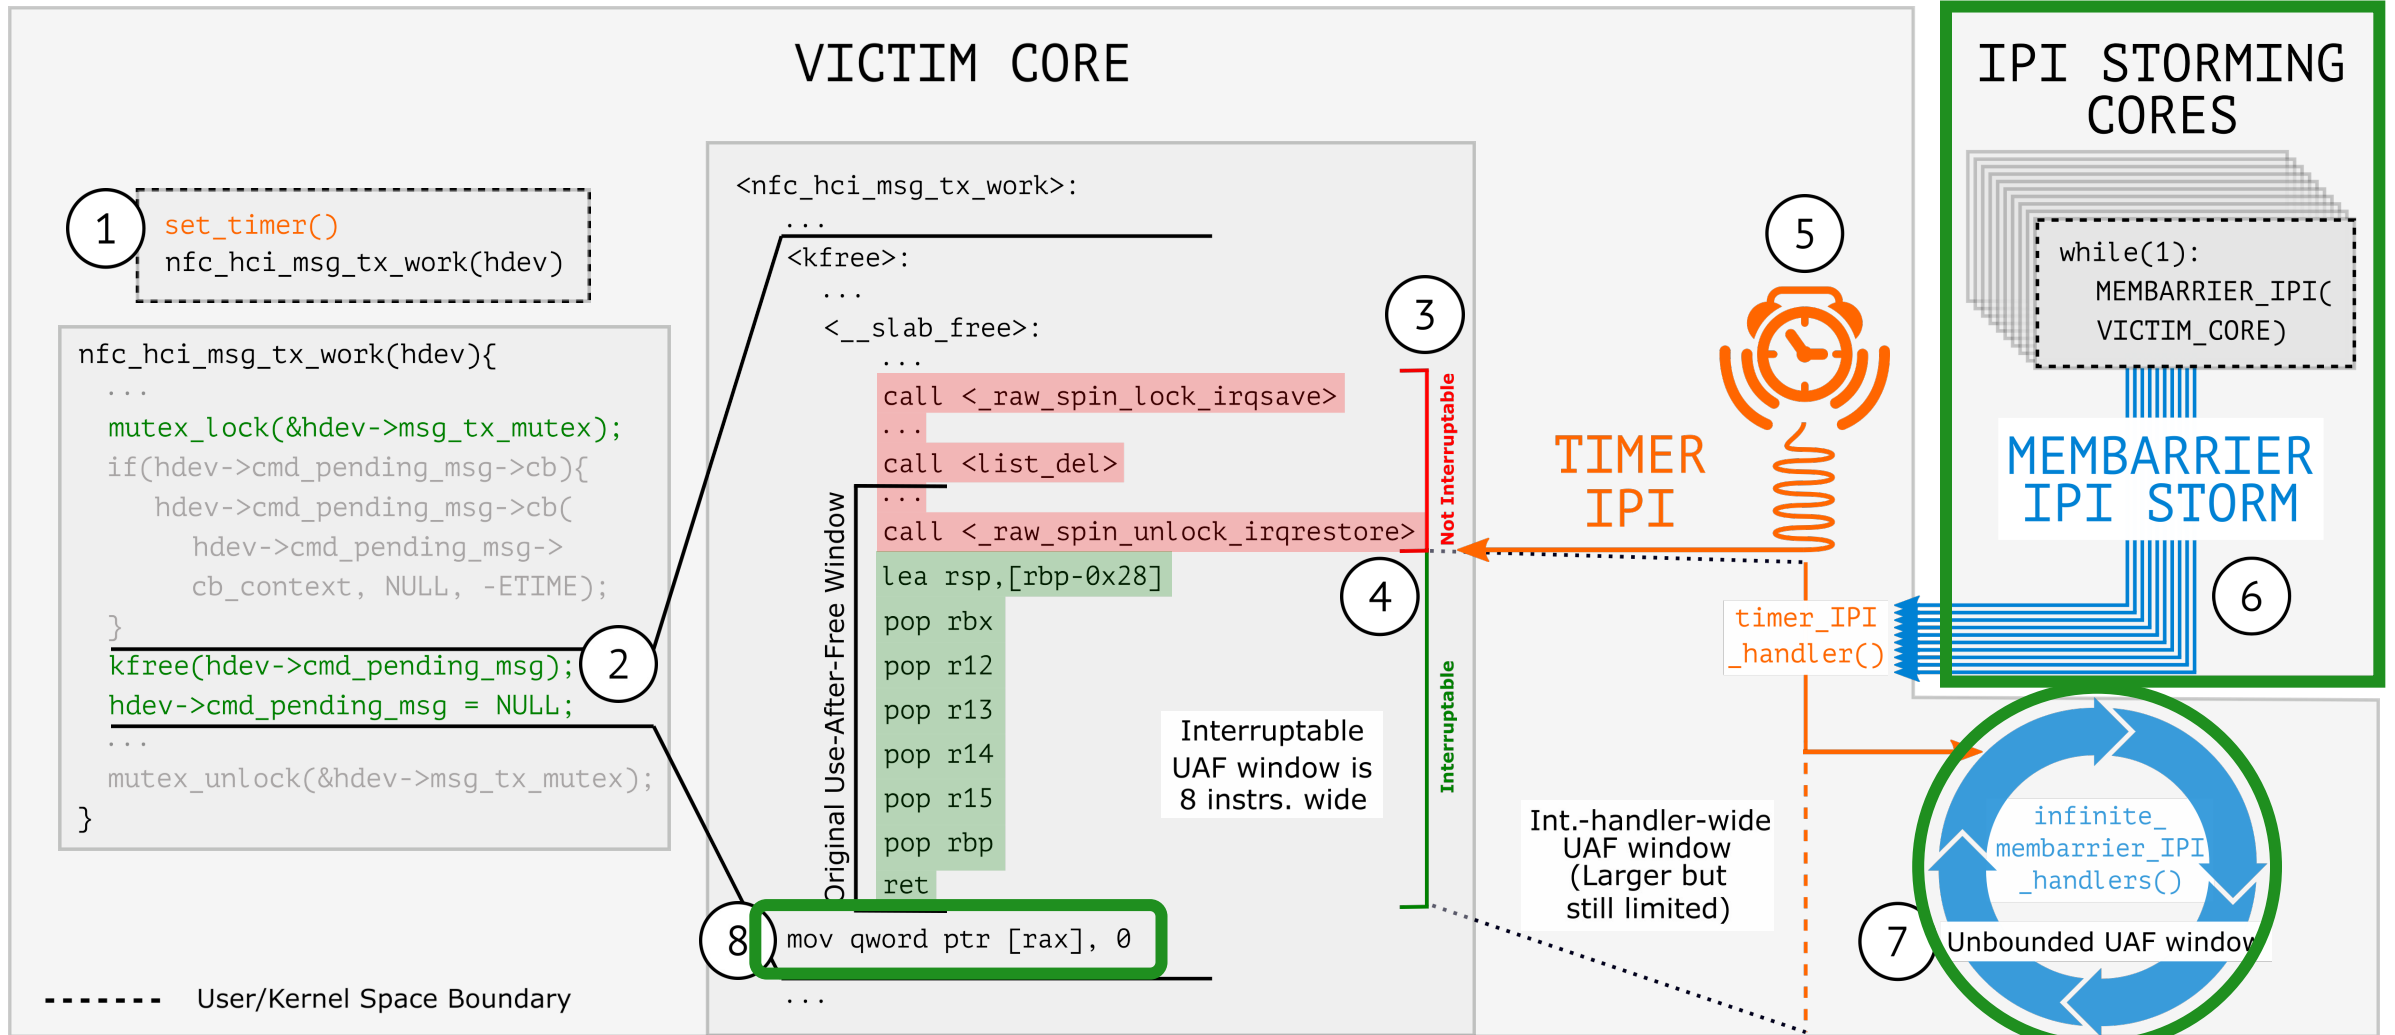
Creating UAF Exploitation Window



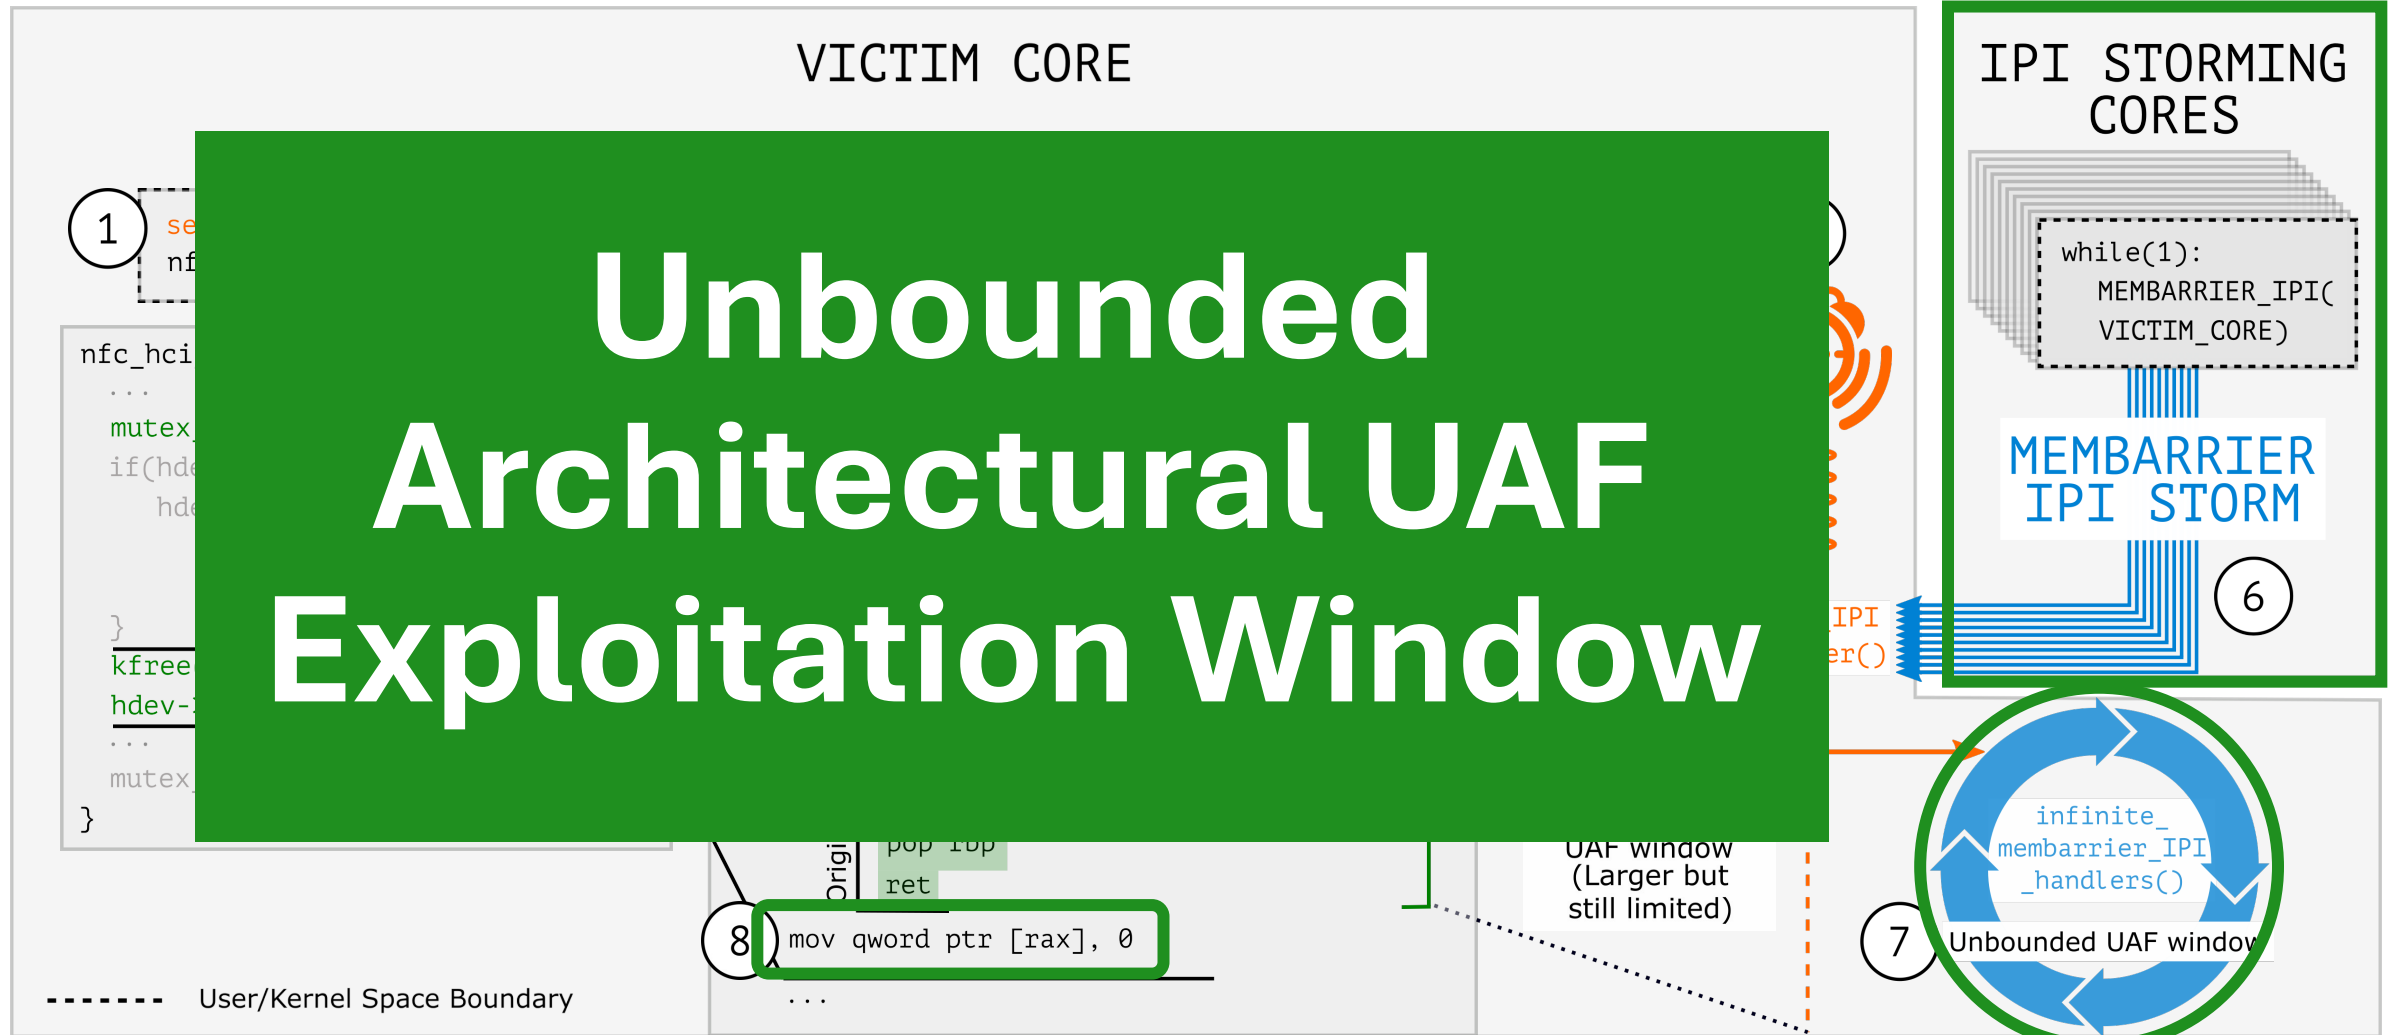
Creating UAF Exploitation Window



Creating UAF Exploitation Window



Creating UAF Exploitation Window



UAF Attack Challenges

1. Bypass The Sync. Primitive:
Speculative Race Condition

```
struct nfc_hci_dev {  
    ...  
    struct mutex msg_tx_mutex;  
    ...  
    struct hci_msg* cmd_pending_msg;  
    ...  
}  
                                     SHARED DATA  
struct hci_r  
    ...  
    void (*cb  
    ...  
    void *cb_context,  
    ...  
}
```

2. Create an Exploitation
Window:
IPI Storming

```
THREAD 1  
nfc_hci_msg_tx_work(hdev){  
    mutex_lock(&hdev->msg_tx_mutex);  
    hdev->cmd_pending_msg->cb(  
    hdev->cmd_pending_msg->cb(  
    hdev->cmd_pending_msg->  
    cb_context, NULL, -ETIME);  
}  
kfree(hdev->cmd_pending_msg);  
hdev->cmd_pending_msg = NULL;  
...  
mutex_unlock(&hdev->msg_tx_mutex);  
}
```

```
THREAD 2  
nfc_hci_msg_tx_work(hdev){  
    ...  
    mutex_lock(&hdev->msg_tx_mutex);  
    if(hdev->cmd_pending_msg->cb){  
        hdev->cmd_pending_msg->cb(  
        hdev->cmd_pending_msg->  
        cb_context, NULL, -ETIME);  
    }  
    kfree(hdev->cmd_pending_msg);  
    hdev->cmd_pending_msg = NULL;  
    ...  
    mutex_unlock(&hdev->msg_tx_mutex);  
}
```

UAF Attack Challenges

1. Bypass The Sync. Primitive:
Speculative Race Condition

3. Reallocate The Freed Memory:
msgsnd syscall

2. Create an Exploitation Window:
IPI Storming

4. Hijack The Control-Flow:
Speculatively Execute A Controlled Disclosure Gadget

```
struct nfc ...  
...  
struct m ...  
...  
struct h ...  
...  
}
```

THREAD 1

```
nfc_hci_msg_tx_work(hdev){  
...  
mutex_lock(&hdev->msg_tx_mutex);  
hdev->cmd_pending_msg->cb(  
hdev->cmd_pending_msg->cb(  
hdev->cmd_pending_msg->  
cb_context, NULL, -ETIME);  
}  
kfree(hdev->cmd_pending_msg);  
hdev->cmd_pending_msg = NULL;  
...  
mutex_unlock(&hdev->msg_tx_mutex);  
}
```

THREAD 2

```
nfc_hci_msg_tx_work(hdev){  
...  
mutex_lock(&hdev->msg_tx_mutex);  
if(hdev->cmd_pending_msg->cb){  
hdev->cmd_pending_msg->cb(  
hdev->cmd_pending_msg->  
cb_context, NULL, -ETIME);  
}
```

Speculative Concurrent Use-After-Free Attack

1. Bypass The Sync. Primitive:
Speculative Race Condition

3. Reallocate The Freed Memory:
msgsnd syscall

2. Create an Exploitation Window:
IPI Storming

4. Hijack The Control-Flow:
Speculatively Execute A Controlled Disclosure Gadget

```
struct nfc ...
struct m ...
struct h ...
...
}
```

THREAD 1

```
hci_msg_tx_work(hdev){
    mutex_lock(&hdev->msg_tx_mutex);
    hdev->cmd_pending_msg->cb(
    hdev->cmd_pending_msg->cb(
    hdev->cmd_pending_msg->
    cb_context, NULL, -ETIME);
}
kfree(hdev->cmd_pending_msg);
hdev->cmd_pending_msg = NULL;
...
mutex_unlock(&hdev->msg_tx_
}
```

THREAD 2

```
nfc_hci_msg_tx_work(hdev){
    ...
    mutex_lock(&hdev->msg_tx_mutex);
    if(hdev->cmd_pending_msg->cb){
        hdev->cmd_pending_msg->cb(
        hdev->cmd_pending_msg->
        cb_context, NULL, -ETIME);
    }
```

Linux Kernel Attack Surface

Linux Kernel Attack Surface

- Gadget scanner using the Coccinelle patching engine

Linux Kernel Attack Surface

- Gadget scanner using the Coccinelle patching engine
- Statically analyzed Linux v5.15.83

Linux Kernel Attack Surface

- Gadget scanner using the Coccinelle patching engine
- Statically analyzed Linux v5.15.83

**Identified 1283 Speculative
Concurrent Use-After-Free
Gadgets**

Mitigating Speculative Race Conditions

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
            ↳arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)
                ↳arch_atomic_try_cmpxchg_acquire(&lock, , ...)
                    ↳arch_atomic_try_cmpxchg(&lock, , ...)
                        ↳arch_try_cmpxchg(&lock, , ...)
                            ↳__raw_try_cmpxchg(ptr, ...){
                                asm volatile(
                                    "lock cmpxchgq %2, %1"
                                    : "=a" (ret), "+m" (*ptr)
                                    : "r" (new), "0" (old)
                                    : "memory"
                                );
                            })
        return true;
    ...
}
```

Mitigating Speculative Race Conditions

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
            ↳ arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)
                ↳ arch_atomic_try_cmpxchg_acquire(&lock, , ...)
                    ↳ arch_atomic_try_cmpxchg(&lock, , ...)
                        ↳ arch_try_cmpxchg(&lock, , ...)
                            ↳ __raw_try_cmpxchg(ptr, ...){
                                asm volatile(
                                    "lock cmpxchgq %2, %1"
                                    : "=a" (ret), "+m" (*ptr)
                                    : "r" (new), "0" (old)
                                    : "memory"
                                );
                            }}
            return true;
    ...
}
```

Mitigating Speculative Race Conditions

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
            ↳ arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)
                ↳ arch_atomic_try_cmpxchg_acquire(&lock, , ...)
                    ↳ arch_atomic_try_cmpxchg(&lock, , ...)
                        ↳ arch_try_cmpxchg(&lock, , ...)
                            ↳ __raw_try_cmpxchg(ptr, ...){
                                asm volatile(
                                    "lock cmpxchgq %2, %1"
                                );
                            }
            return true;
    ...
}
```

lfence

Mitigating Speculative Race Conditions

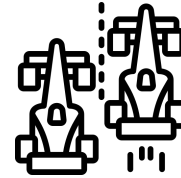
LMBench:
Geomean
≈5%

```
void mutex_lock(struct mutex *lock){
    ...
    if (!__mutex_trylock_fast(lock))
        if (atomic_long_try_cmpxchg_acquire(&lock, ...))
            ↳ arch_atomic_long_try_cmpxchg_acquire(&lock, , ...)
              ↳ arch_atomic_try_cmpxchg_acquire(&lock, , ...)
                ↳ arch_atomic_try_cmpxchg(&lock, , ...)
                  ↳ arch_try_cmpxchg(&lock, , ...)
                    ↳ __raw_try_cmpxchg(ptr, ...){
                      asm volatile(
                        "lock cmpxchgq %2, %1"
                        ↳ fence
                      });
                }
            return true;
        ...
    }
```

GhostRace: Exploiting and Mitigating Speculative Race Conditions

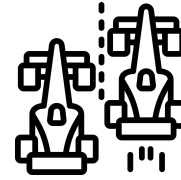
GhostRace: Exploiting and Mitigating Speculative Race Conditions

- Speculative Race Condition



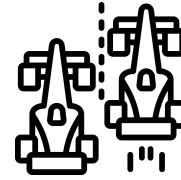
GhostRace: Exploiting and Mitigating Speculative Race Conditions

- Speculative Race Condition
- Inter-Process Interrupt Storming



GhostRace: Exploiting and Mitigating Speculative Race Conditions

- Speculative Race Condition



- Inter-Process Interrupt Storming

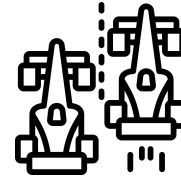


- Gadget Scanner (1200+)



GhostRace: Exploiting and Mitigating Speculative Race Conditions

- Speculative Race Condition



- Inter-Process Interrupt Storming



- Gadget Scanner (1200+)

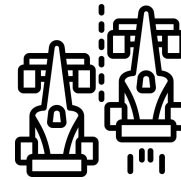


- PoC Exploit (12KB/s)



GhostRace: Exploiting and Mitigating Speculative Race Conditions

- Speculative Race Condition



- Inter-Process Interrupt Storming



- Gadget Scanner (1200+)



- PoC Exploit (12KB/s)

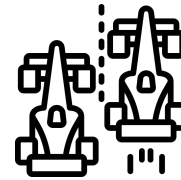


- Mitigation (5%)



GhostRace: Exploiting and Mitigating Speculative Race Conditions

- Speculative Race Condition
- Inter-Process Interrupt Storming
- Gadget Scanner (1200+)
- PoC Exploit (12KB/s)
- Mitigation (5%)



IBM Research | Zurich