

SoK: Automated Kernel Vulnerability Discovery and Exploit Generation

Anil Kurmus, **Andrea Mambretti**, Alessandro Sorniotti

IBM Research Europe – Zurich

Vincent Lenders, Damian Pfammatter, Bernhard Tellenbach

Armasuisse - Cyber Defense Campus

WOOT 2025

Once upon a time...

Once upon a time...



UNIVERSITY OF OXFORD
COMPUTING LABORATORY

MSC COMPUTER SCIENCE DISSERTATION

Automatic Generation of Control Flow Hijacking
Exploits for Software Vulnerabilities

Author:
Sean HEELAN

Supervisor:
Dr. Daniel KROENING

2009



armasuisse



Once upon a time...

2011


AEG: Automatic Exploit Generation

Thanassis Avgerinos, Sang Kil Cha, Brent Lim Tze Hao and David Brumley

Carnegie Mellon University, Pittsburgh, PA

{thanassis, sangkilc, brentlim, dbrumley}@cmu.edu

Once upon a time...



2016

Darpa - Cyber Grand Challenge

Darpa - Cyber Grand Challenge

Contributions:

First big push towards full end-2-end
AEG

Considerably advanced static and
dynamic analyses

Modern opensource tools (e.g., angr,
S2E)

Darpa - Cyber Grand Challenge

Contributions:

First big push towards full end-2-end
AEG

Considerably advanced static and
dynamic analyses

Modern opensource tools (e.g., angr,
S2E)

Limitations:

Simplified OS & Architecture (e.g., 7 syscalls)

Small size software

Binary code only

Proof-of-vulnerability only (e.g., IP control)

Darpa - Cyber Grand Challenge

Contributions:

First big push towards full end-2-end
AEG

Considerably advanced static and
dynamic analyses

Modern opensource tools (e.g., angr,
S2E)

Limitations:

Simplified OS & Architecture (e.g., 7 syscalls)

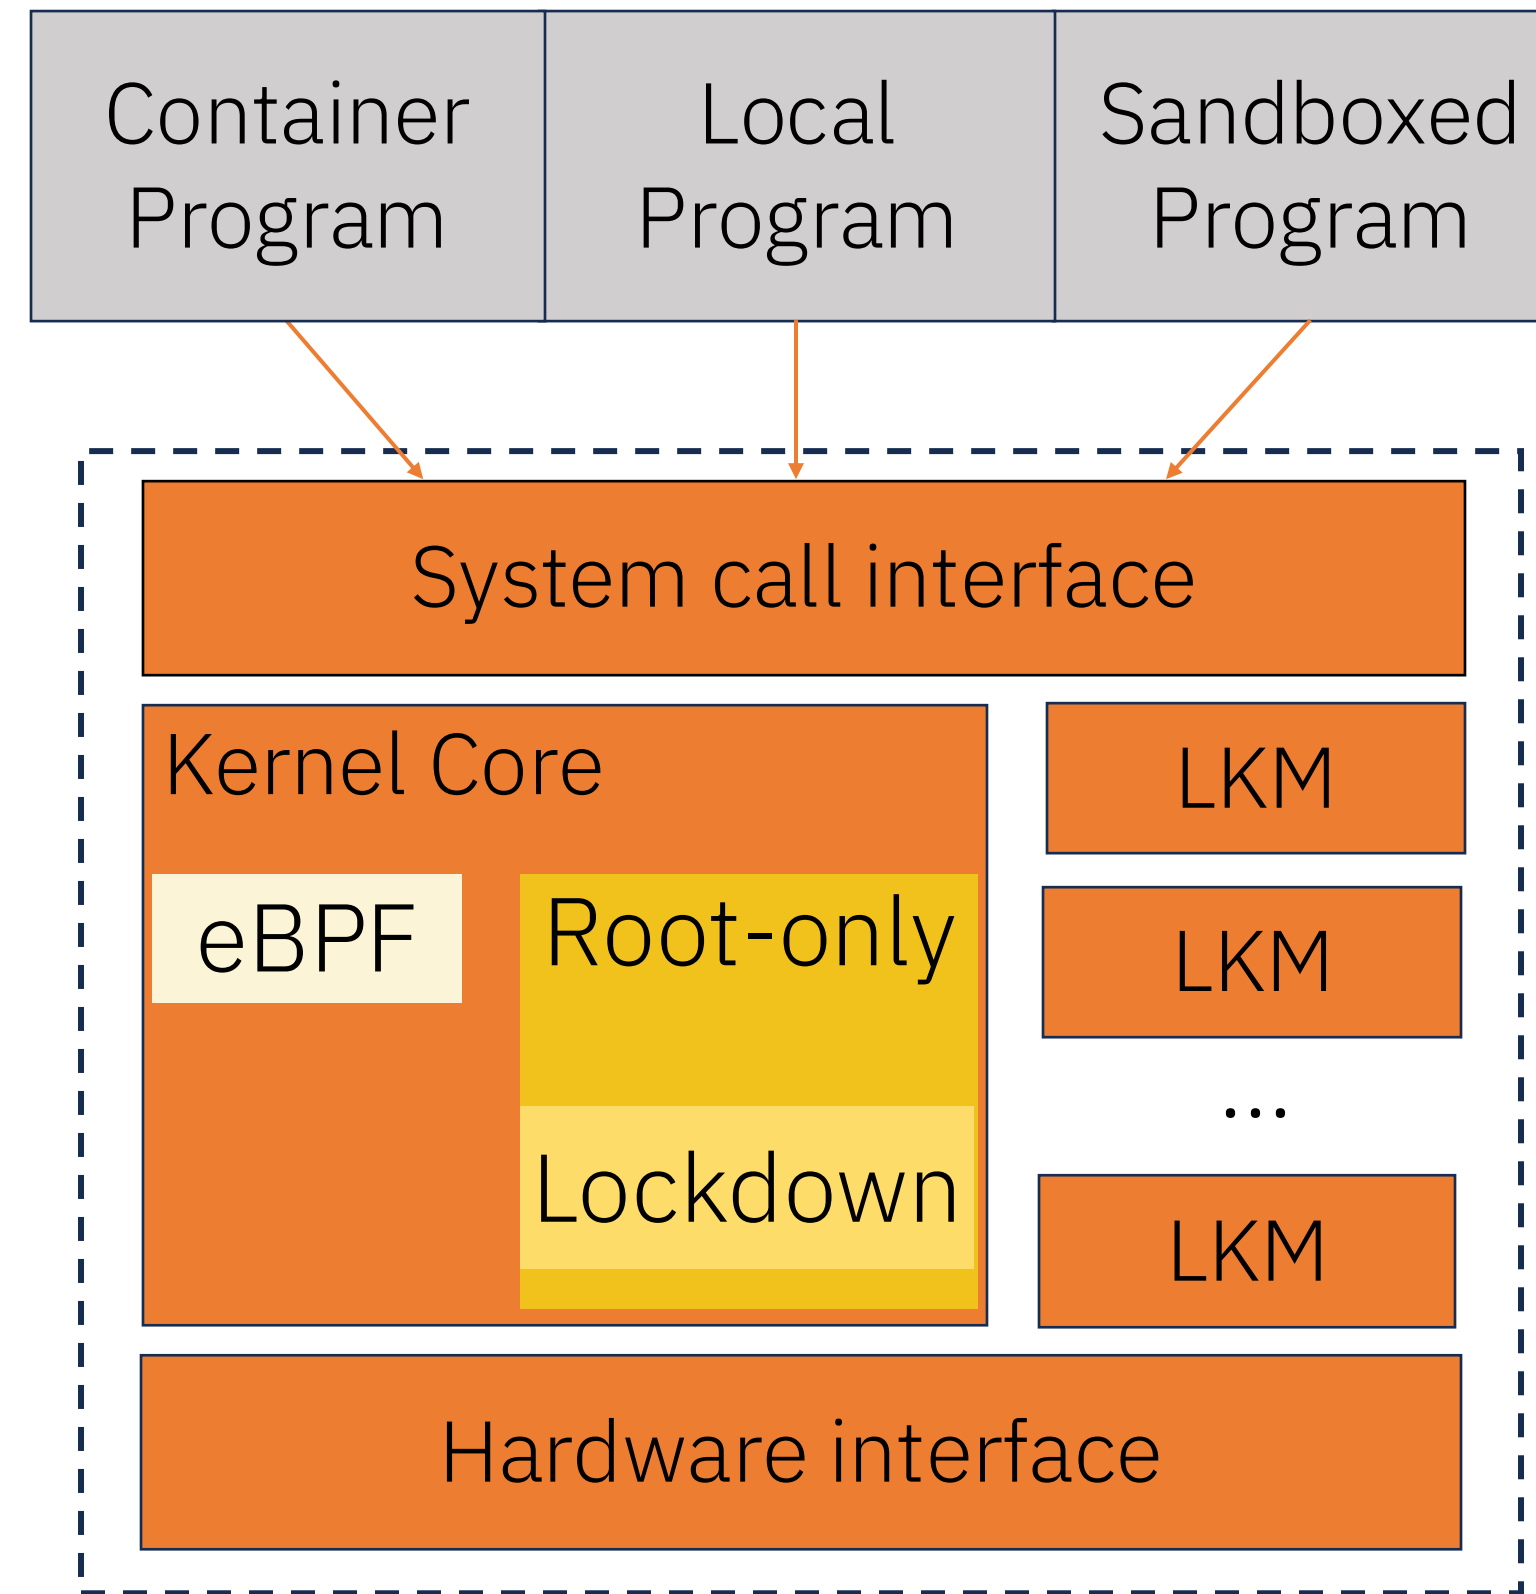

Small size software

Binary code only

Proof-of-vulnerability only (e.g., IP control)

~10 years to move it to the next level

The ultimate target





- Attacker
- Attack surface
- Running Kernel

The ultimate target

High Privilege

Large attack surface

Widely available and used



→ Attacker

■ Attack surface

----- Running Kernel

The ultimate target

High Privilege



Large attack surface

Widely available and used

40 Millions LoC

Several threat models

Several attacks/mitigations



- Attacker
- Attack surface
- Running Kernel

Threat models

Local

Local/User sandbox

Local/Container

Remote

Kernel sandbox (eBPF)

Lockdown

Physical

Vulnerability classes

PRECURSOR: might lead to a serious consequence only through a base vulnerability

HELPER: might only lead to an attacker achieving its goal in combination with another vulnerability

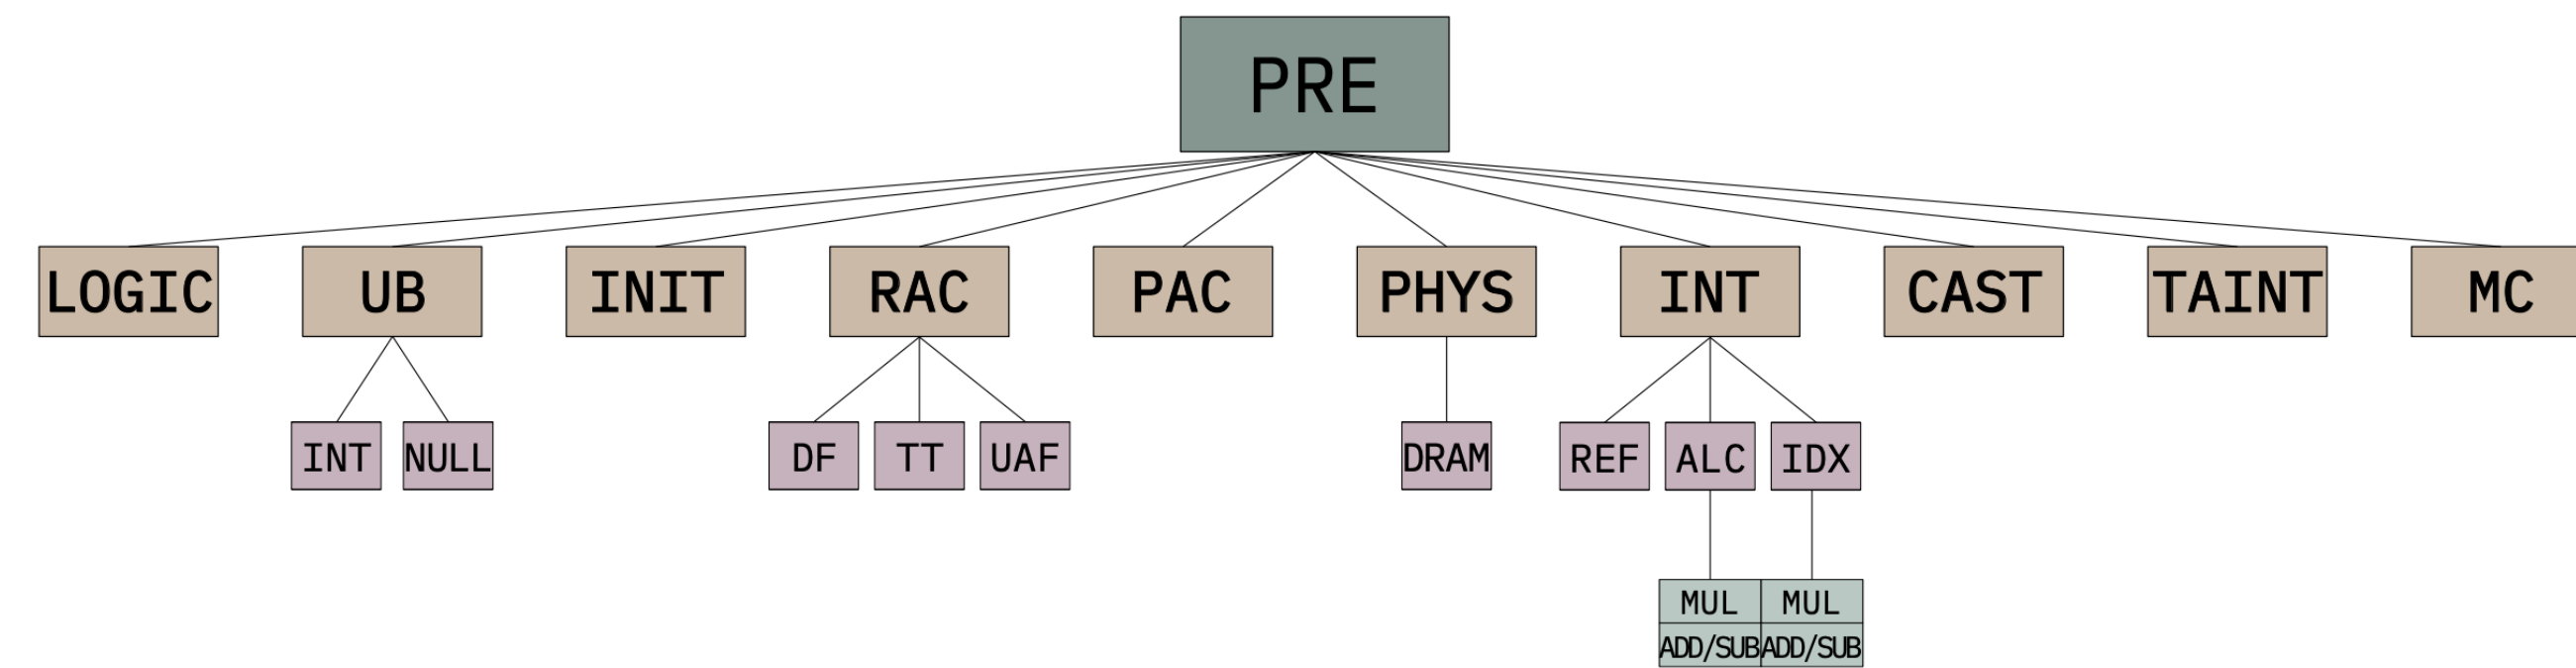
BASE: potentially, it allows attackers to directly reach the goal

Vulnerability classes

PRECURSOR: might lead to a serious consequence only through a base vulnerability

HELPER: might only lead to an attacker achieving its goal in combination with another vulnerability

BASE: potentially, it allows attackers to directly reach the goal

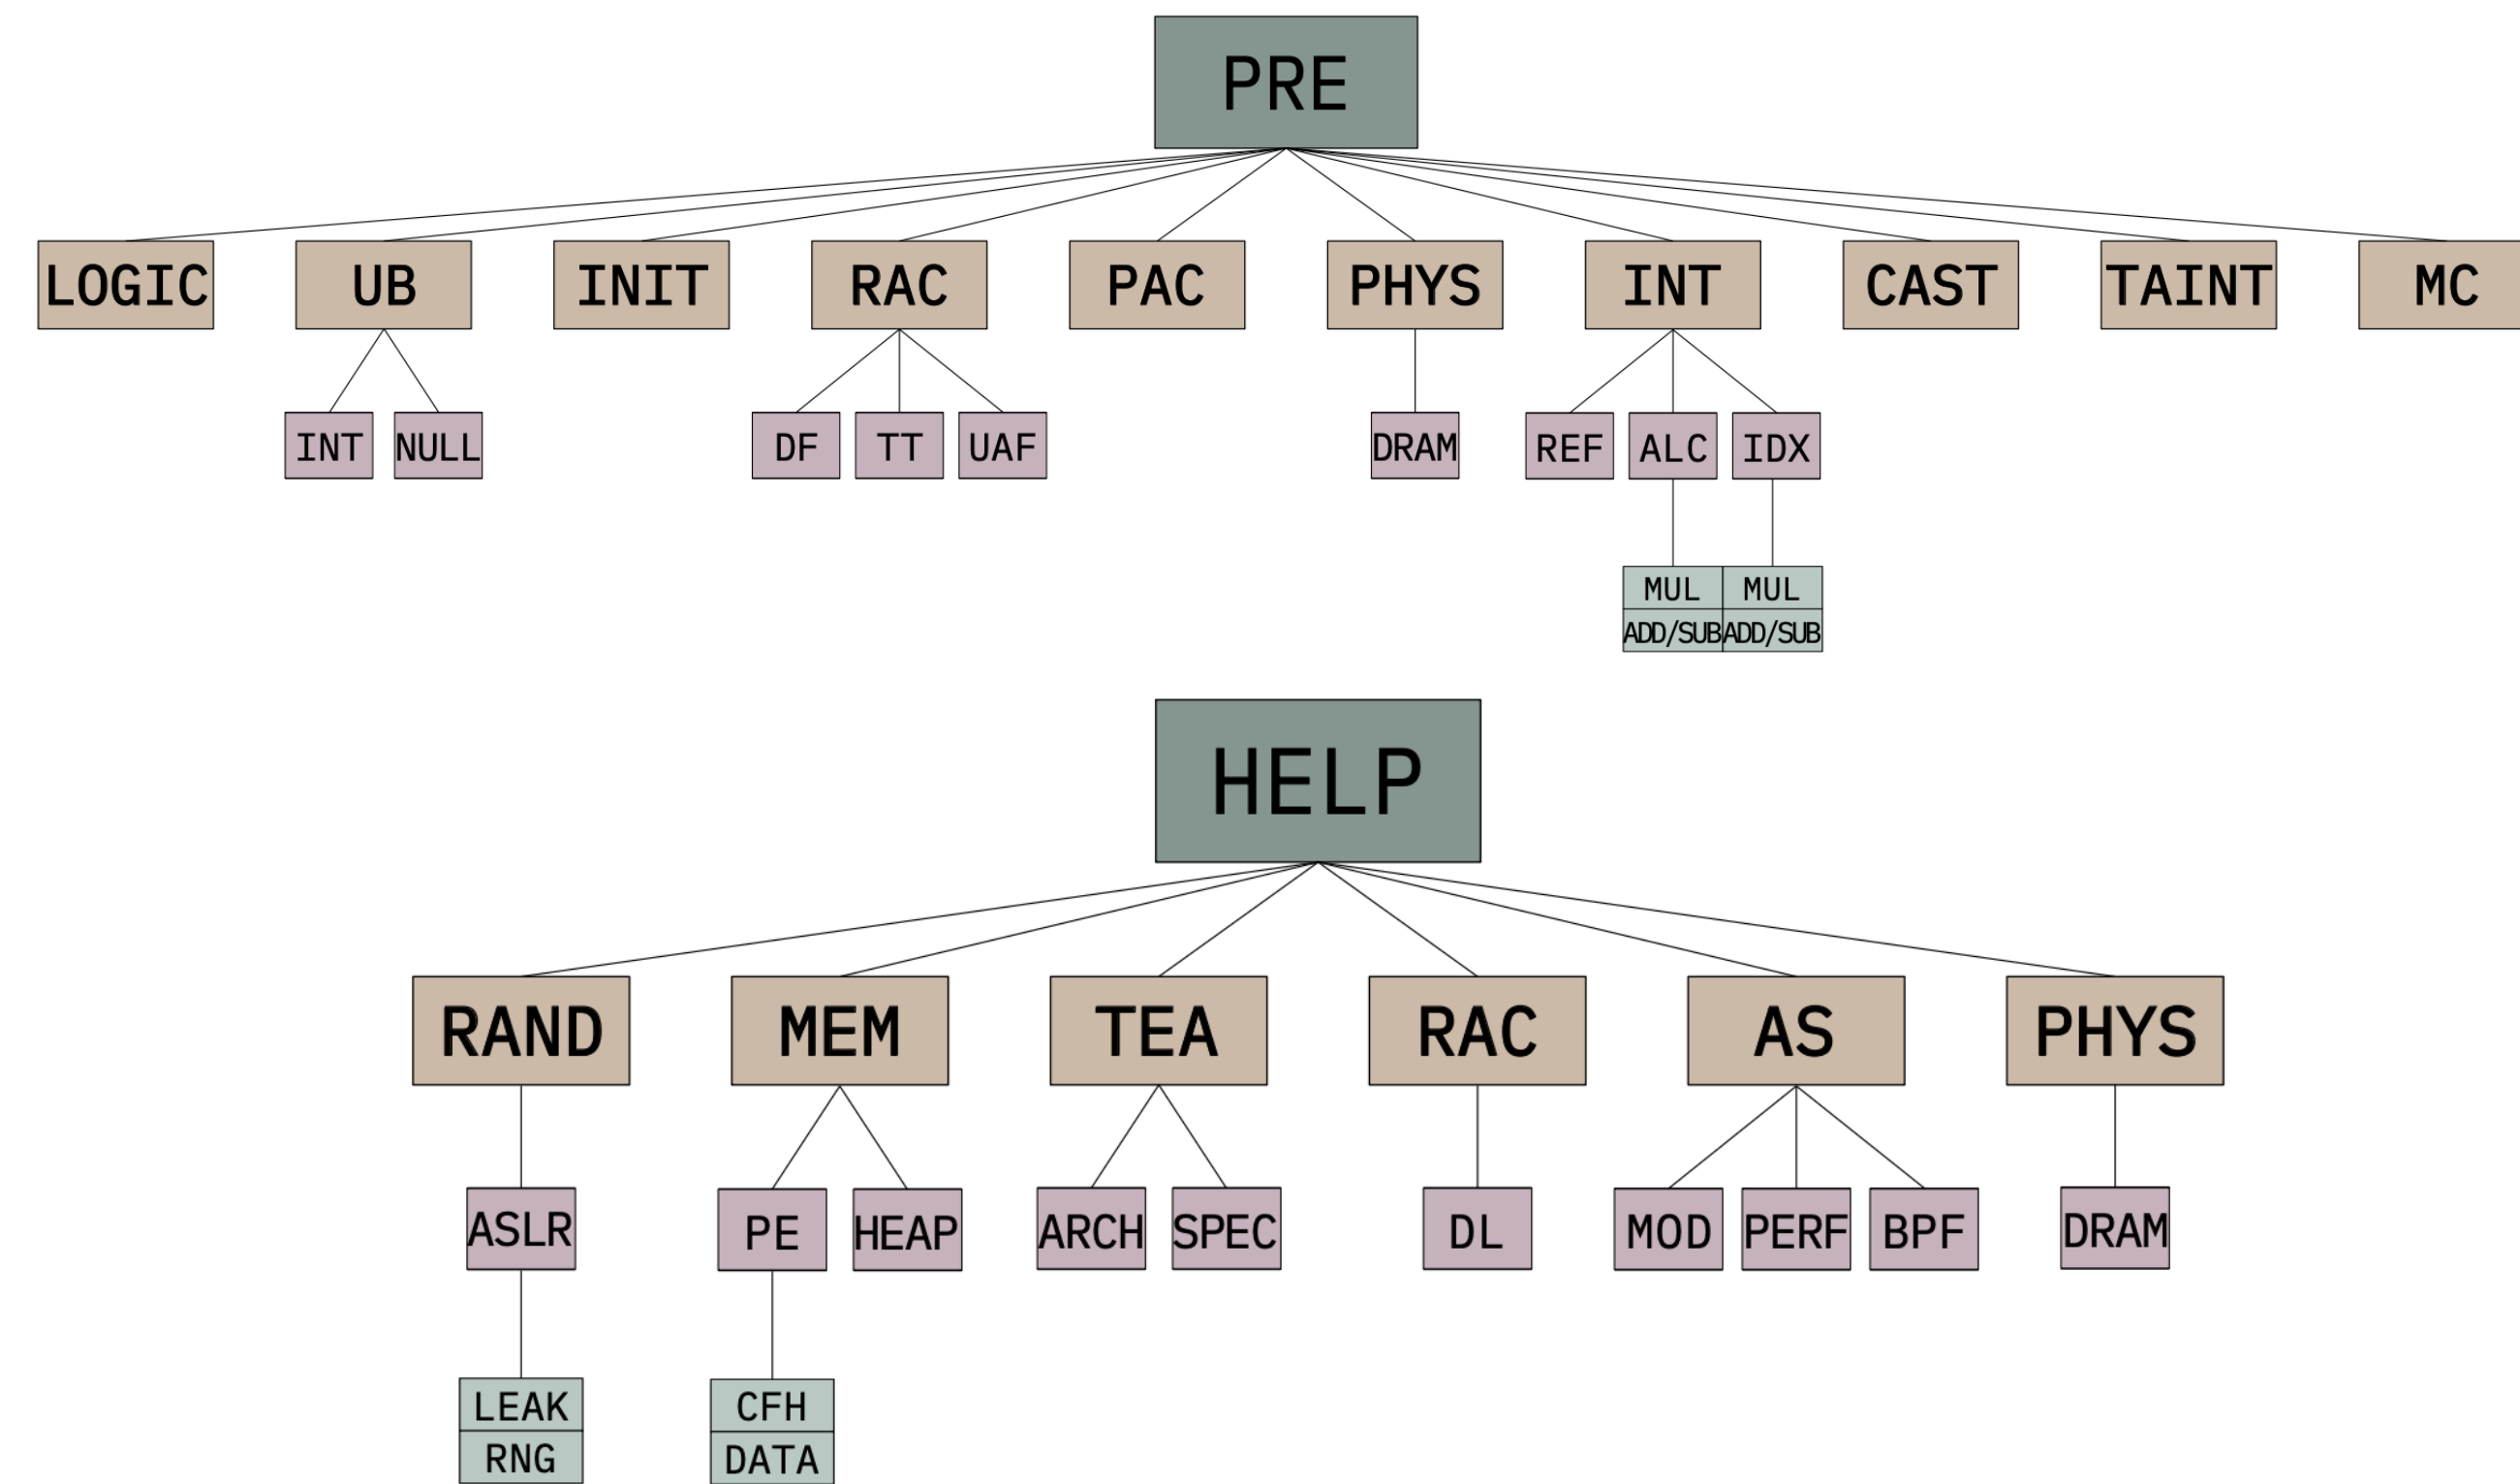


Vulnerability classes

PRECURSOR: might lead to a serious consequence only through a base vulnerability

HELPER: might only lead to an attacker achieving its goal in combination with another vulnerability

BASE: potentially, it allows attackers to directly reach the goal

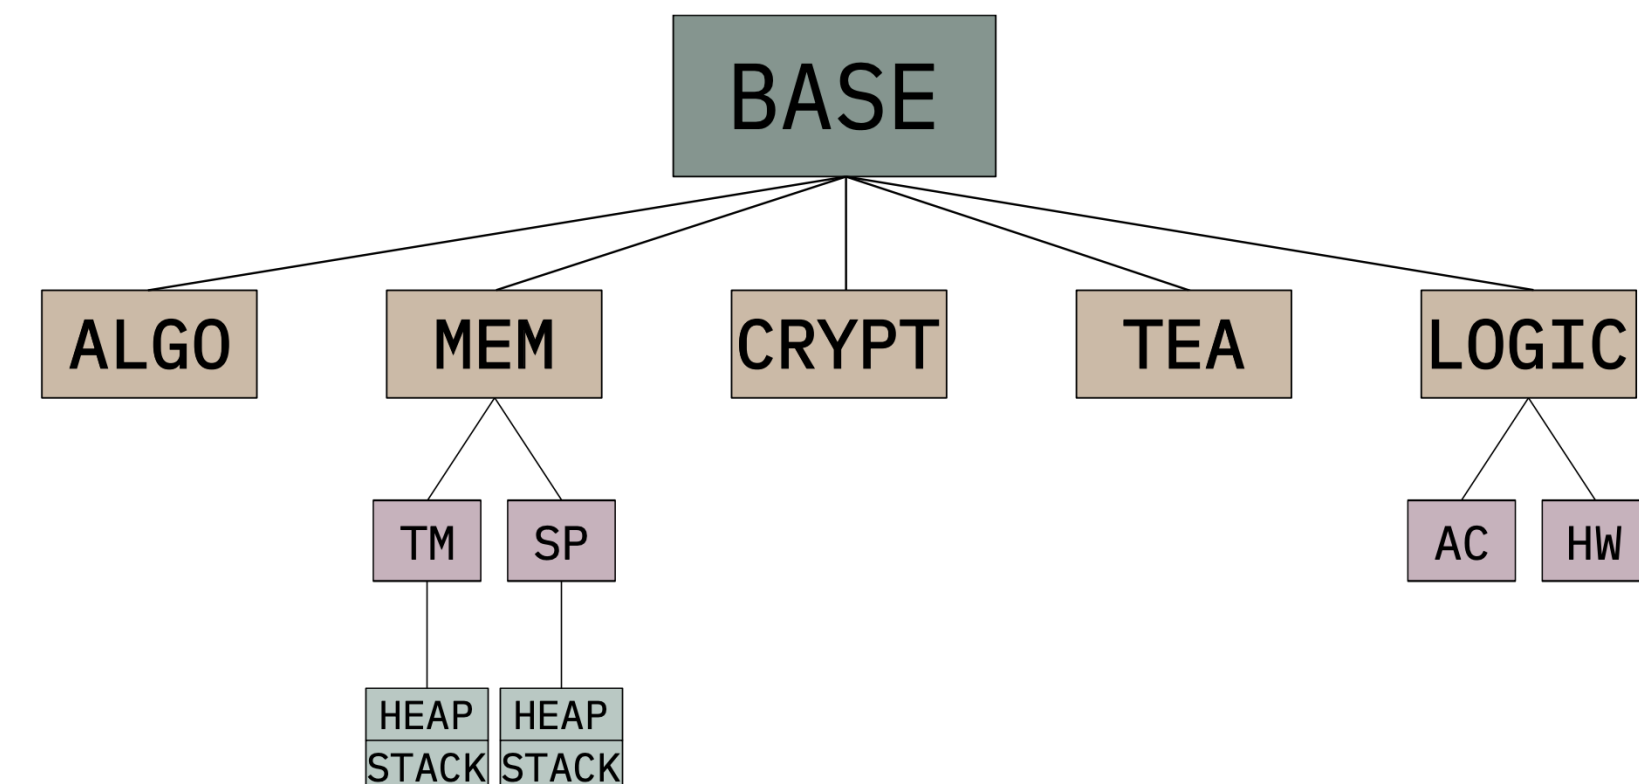
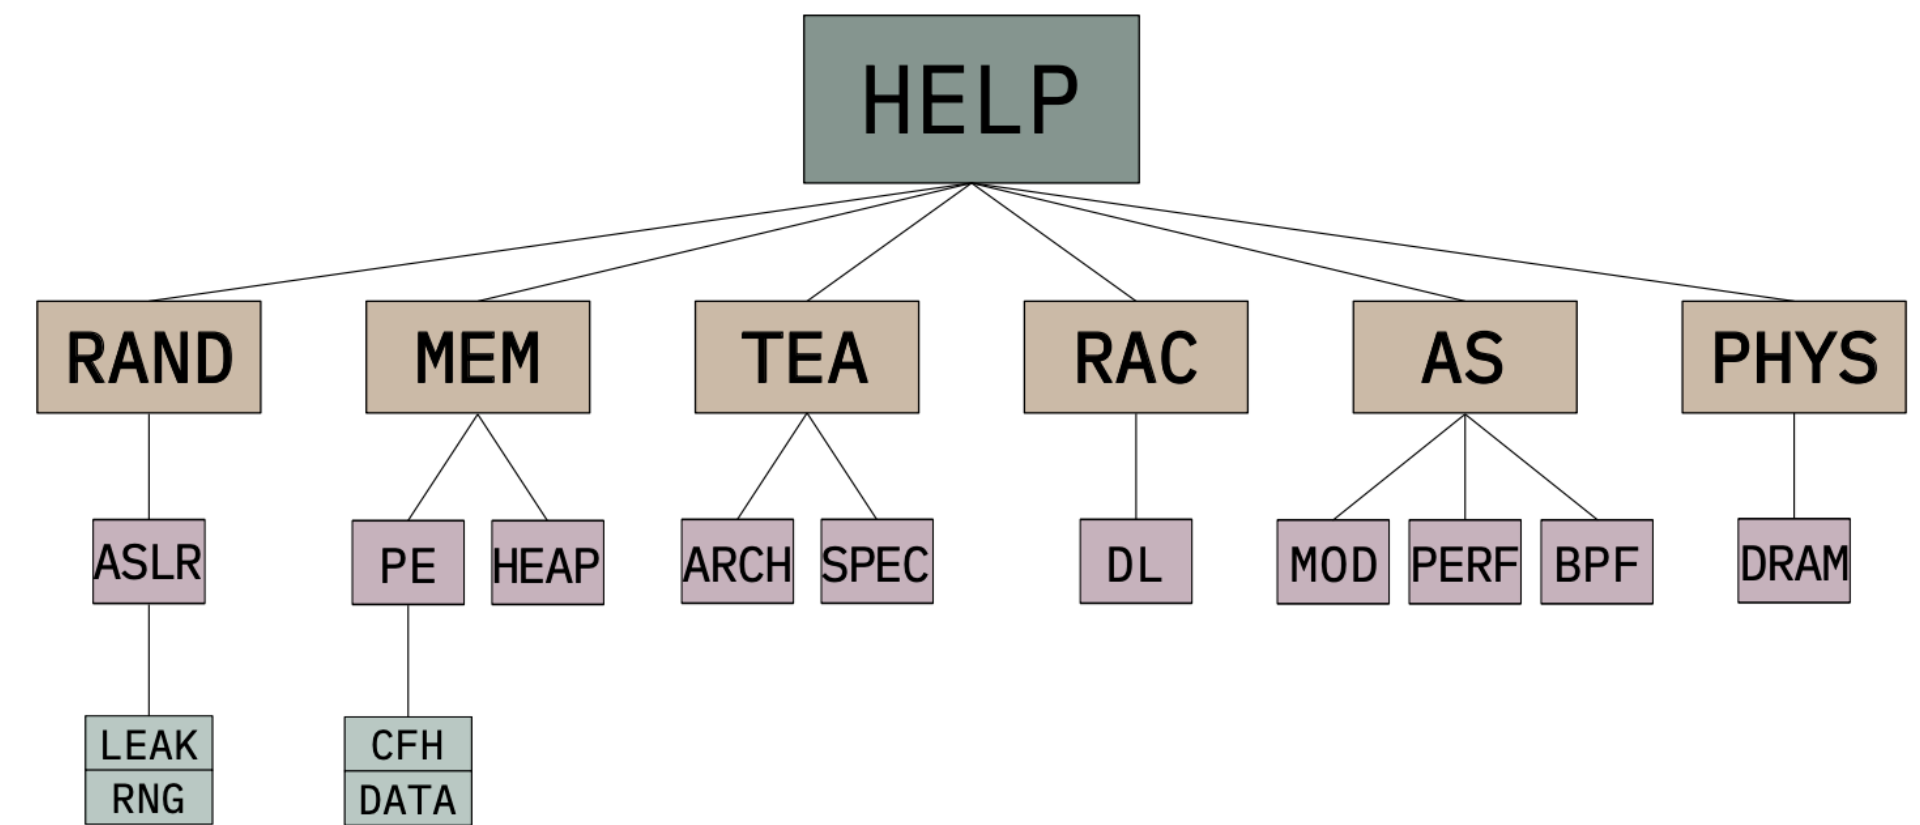



Vulnerability classes

PRECURSOR: might lead to a serious consequence only through a base vulnerability

HELPER: might only lead to an attacker achieving its goal in combination with another vulnerability

BASE: potentially, it allows attackers to directly reach the goal



Let's now consider CVE-2021-4083

TREAD 1

```
init_skb()
```

```
fd = acquire_fd() {
```

```
/* window of 12 insts  
between getting fd and  
refcnt++ */
```

```
}
```

```
recvmsg(fd)
```

USD GC

```
if (refcnt == my_refcnt){
```

```
/* I own skb */
```

```
...
```

```
delete(skb)
```

```
/* without locking */
```

```
}
```



Let's now consider CVE-2021-4083

TREAD 1

```
init_skb()
```

```
fd = acquire_fd() {
```

```
/* window of 12 insts  
between getting fd and  
refcnt++ */
```

```
}
```

```
recvmsg(fd)
```

USD GC

```
if (refcnt == my_refcnt){
```

```
/* I own skb */
```

```
...
```

```
delete(skb)
```

```
/* without locking */
```

```
}
```



Let's now consider CVE-2021-4083

TREAD 1

```
init_skb()
```

```
fd = acquire_fd() {
```

```
/* window of 12 insts  
between getting fd and  
refcnt++ */
```

```
}
```

```
recvmsg(fd)
```

SKB

fd
refcnt

USD GC

```
if (refcnt == my_refcnt){
```

```
/* I own skb */
```

```
...
```

```
delete(skb)
```

```
/* without locking */
```

```
}
```



Let's now consider CVE-2021-4083

TREAD 1

`init_skb()`

```
fd = acquire_fd() {  
    /* window of 12 insts  
    between getting fd and  
    refcnt++ */  
}
```

`recvmsg(fd)`

SKB

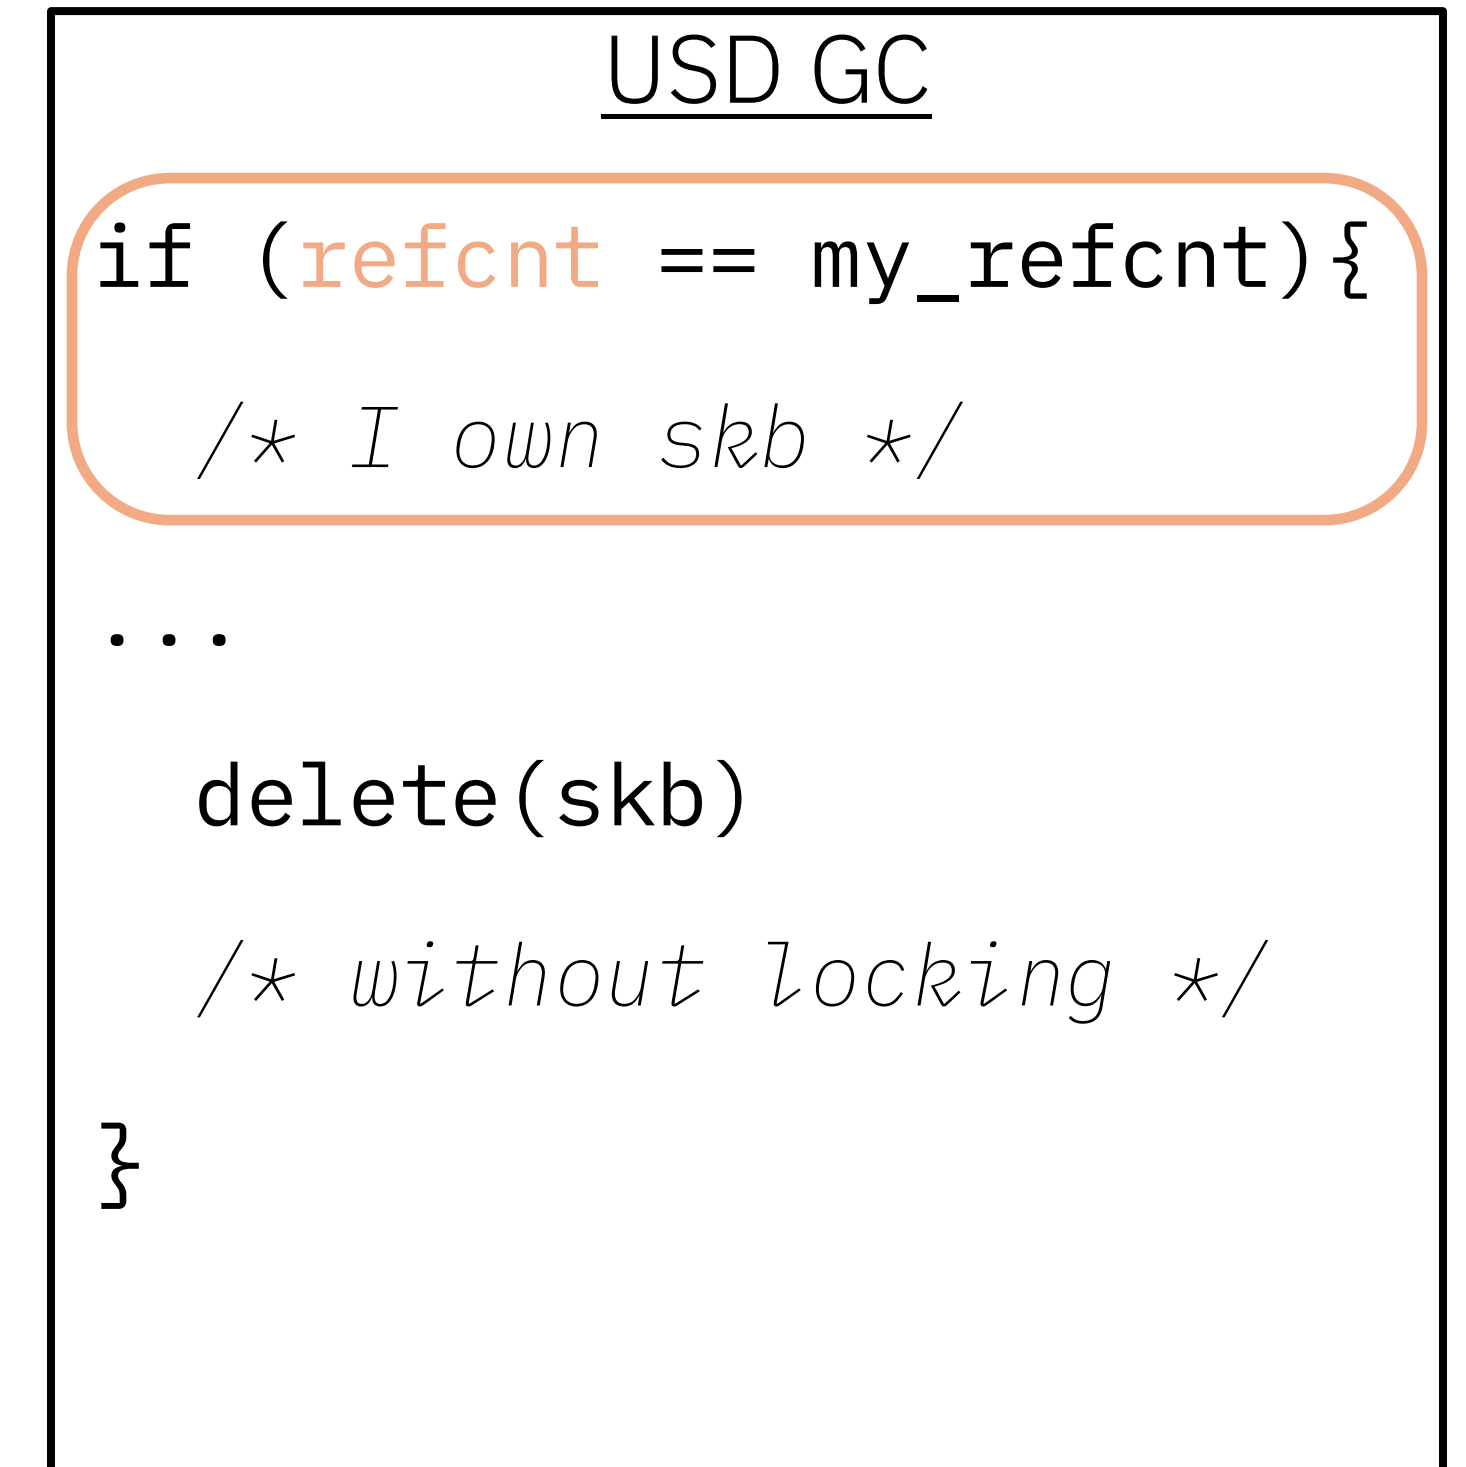
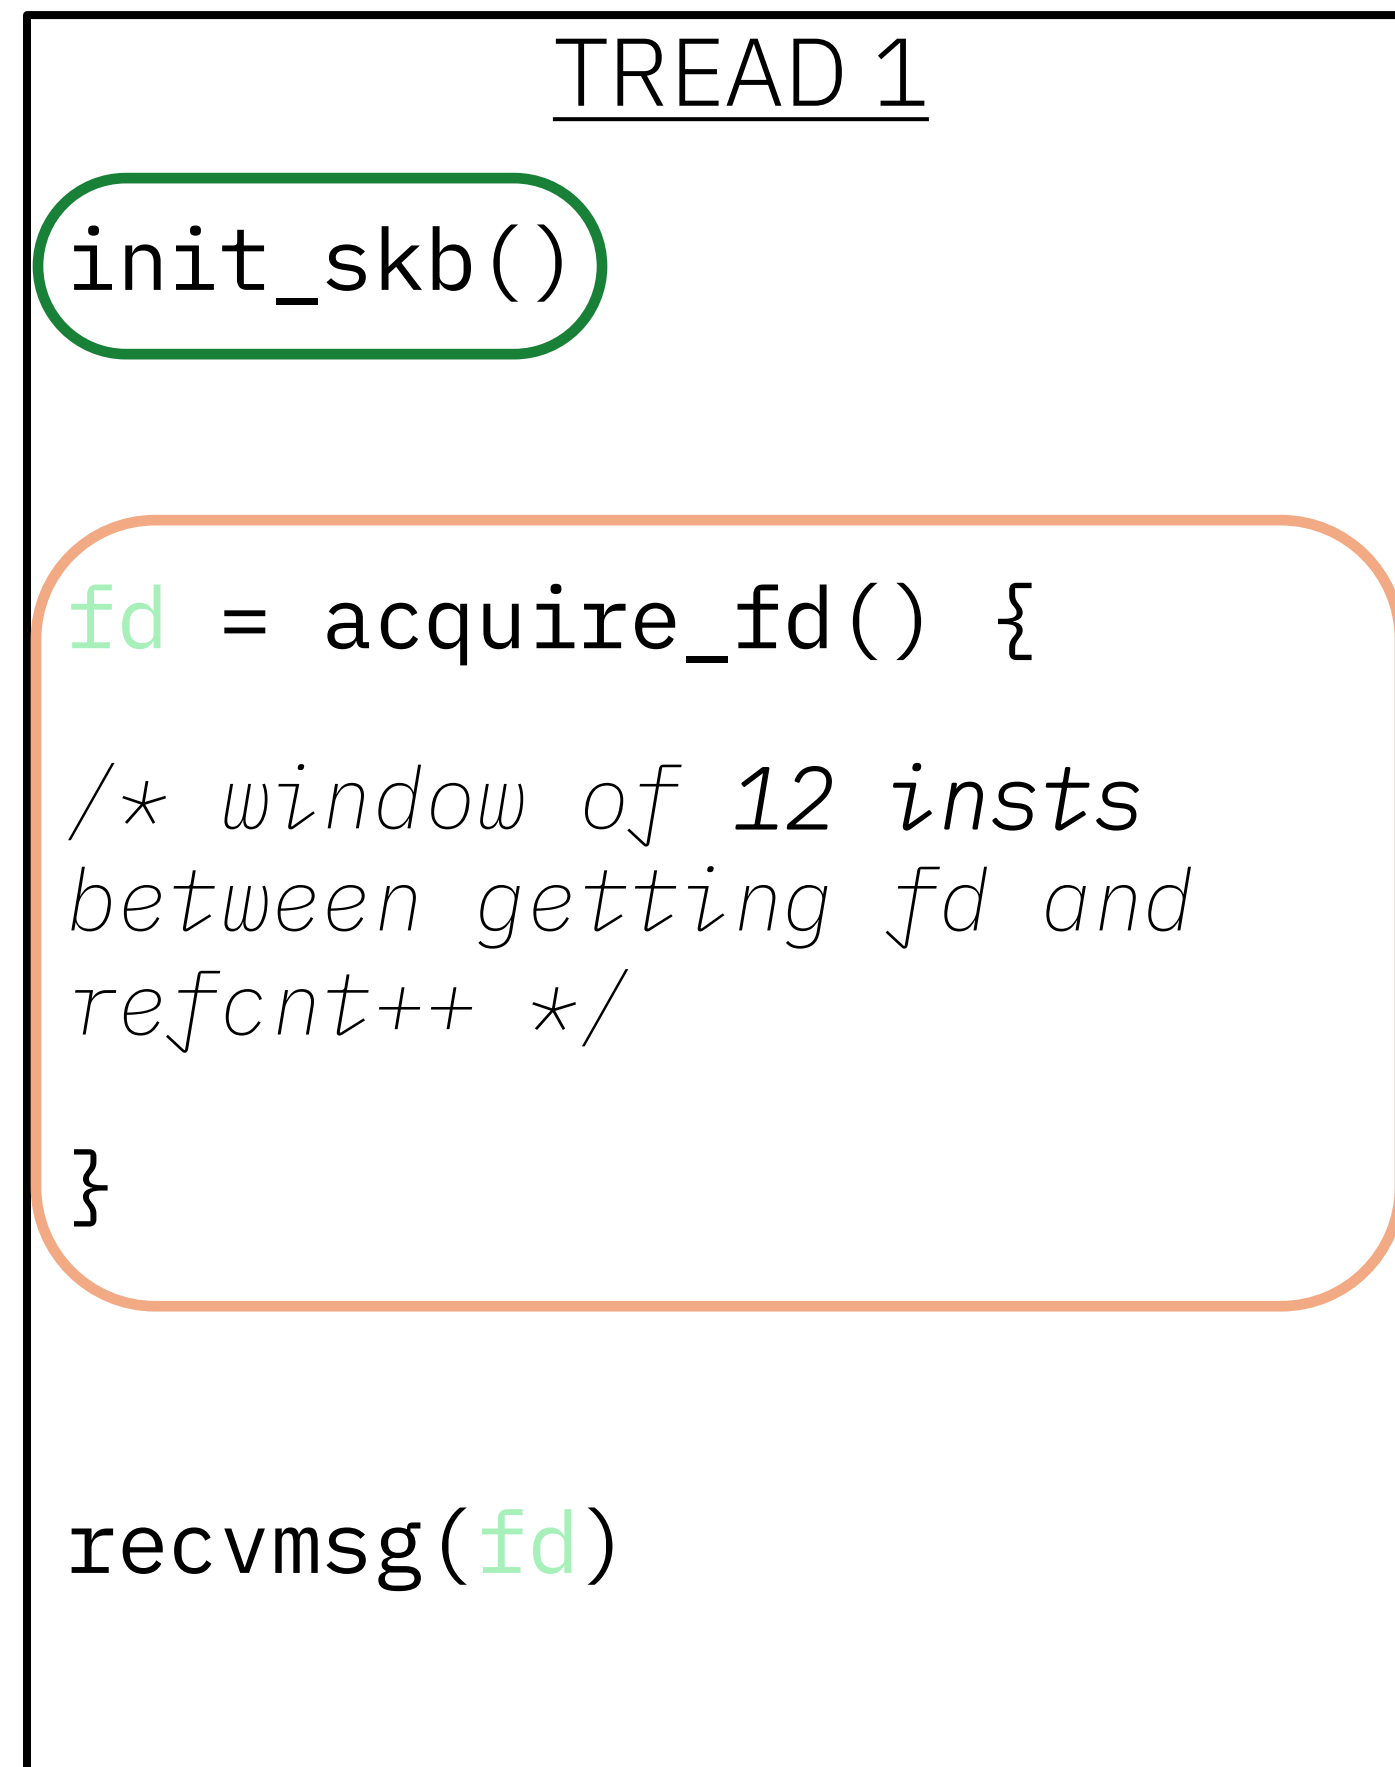
`fd`
`refcnt`

USD GC

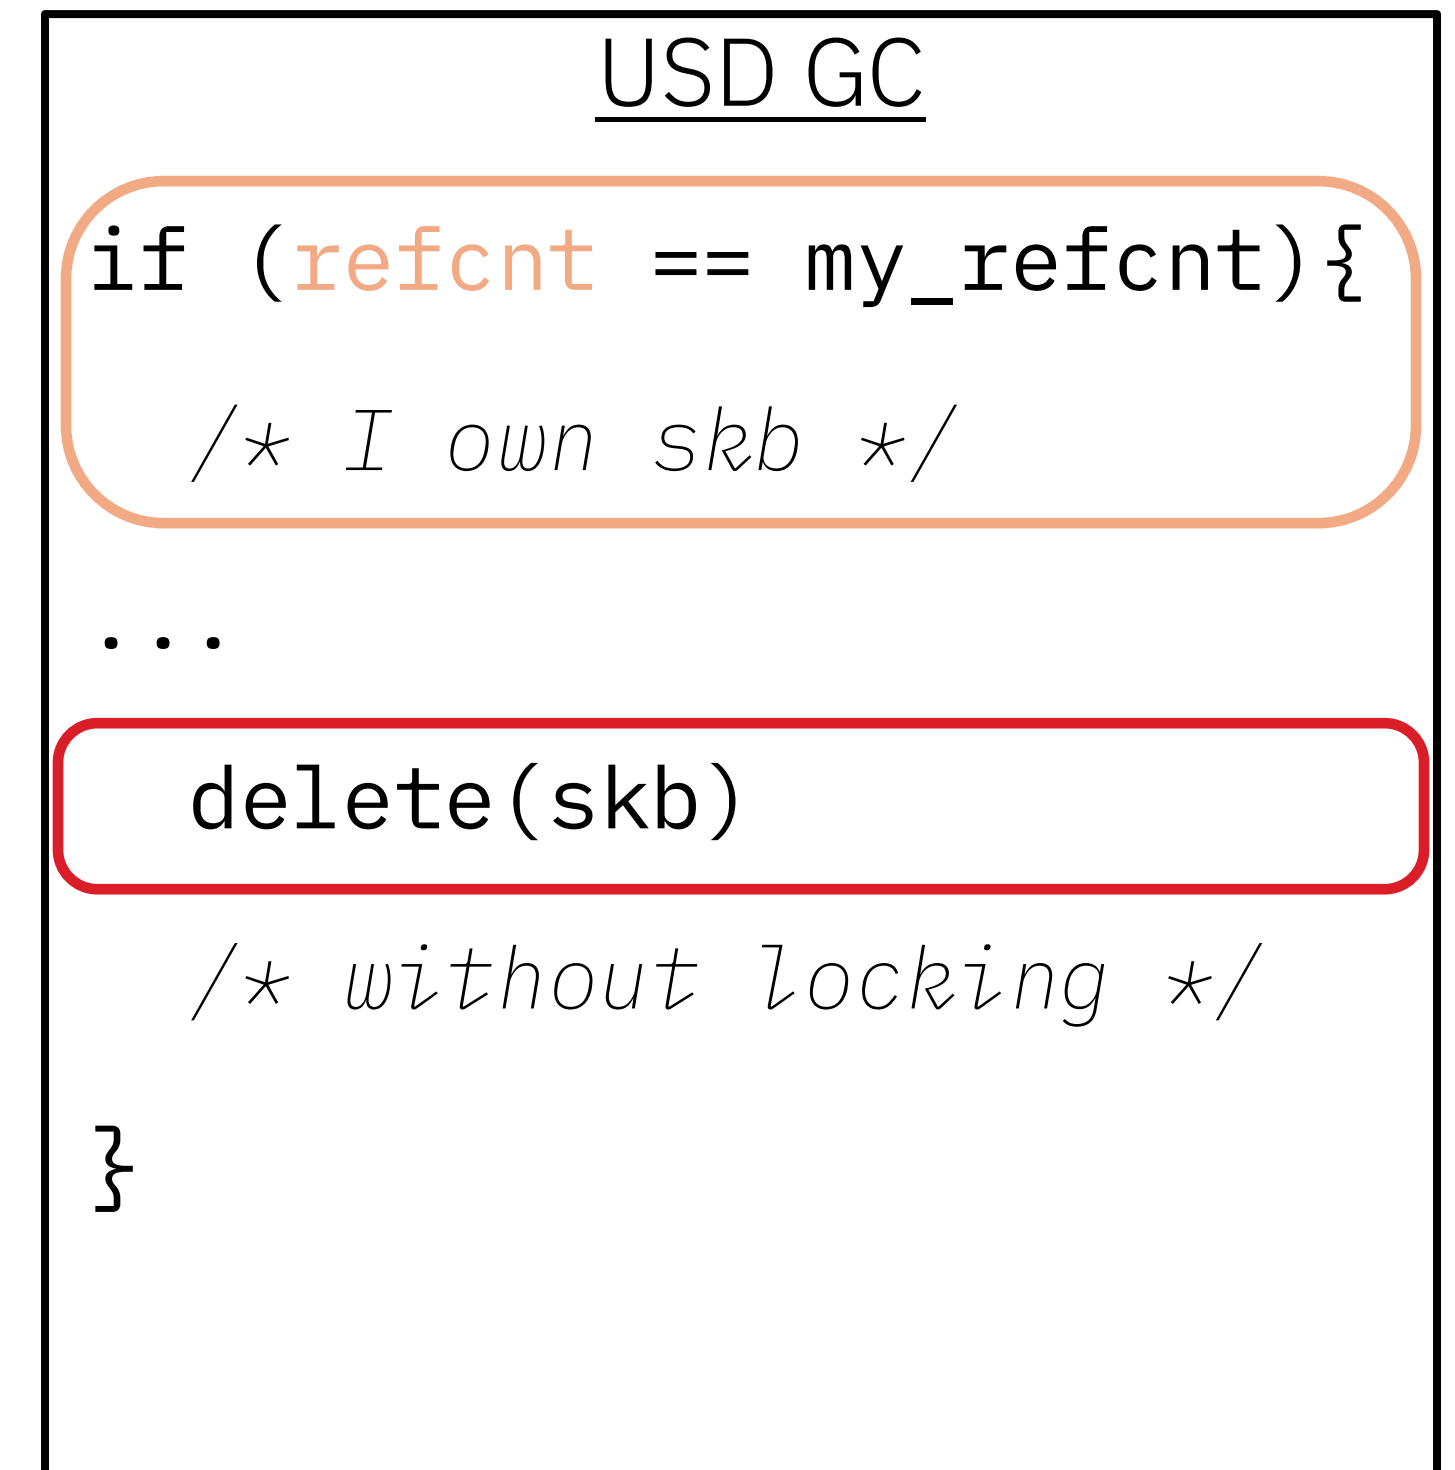
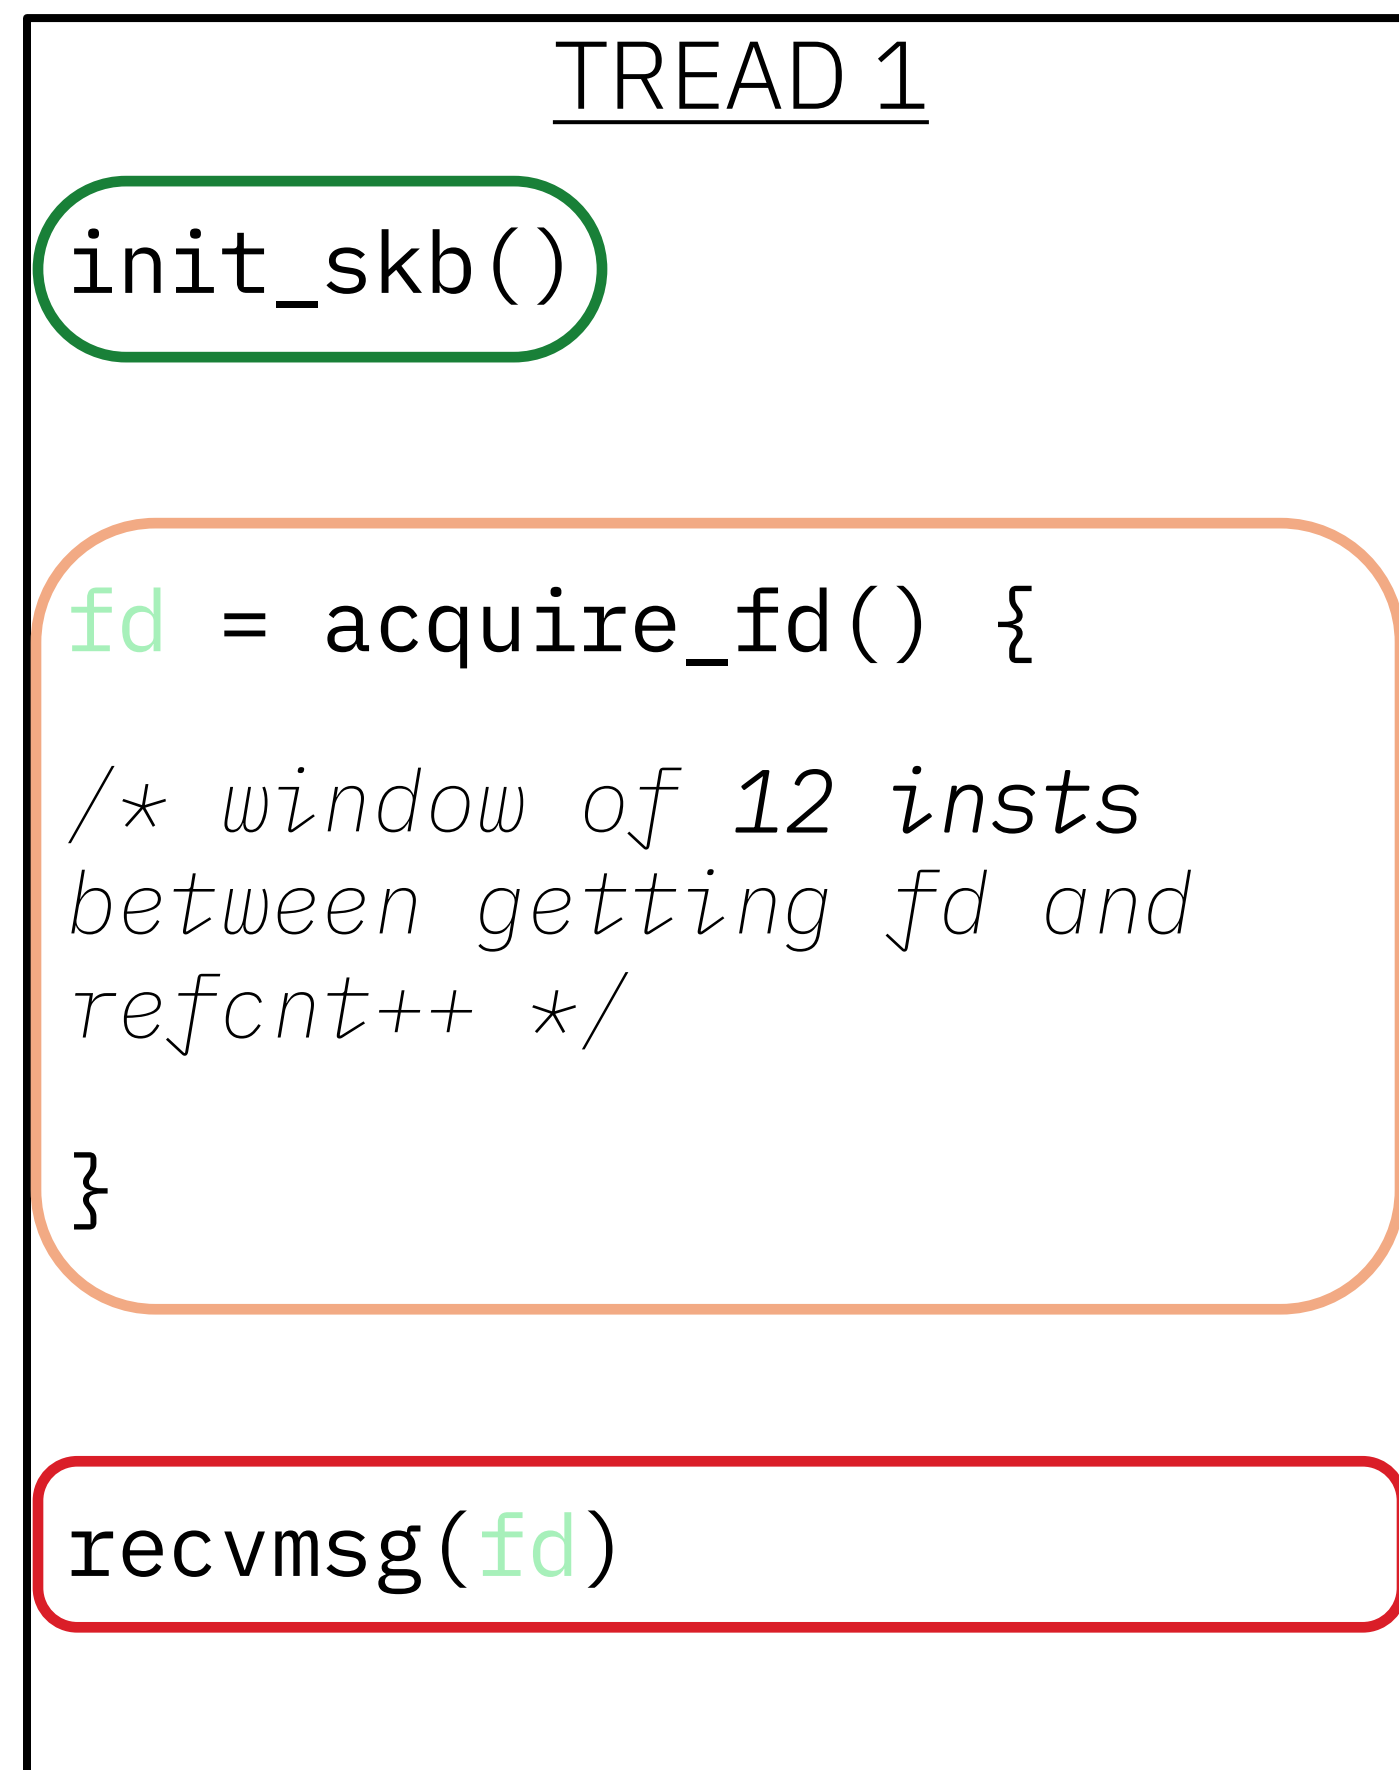
```
if (refcnt == my_refcnt){  
    /* I own skb */  
    ...  
    delete(skb)  
    /* without locking */  
}
```



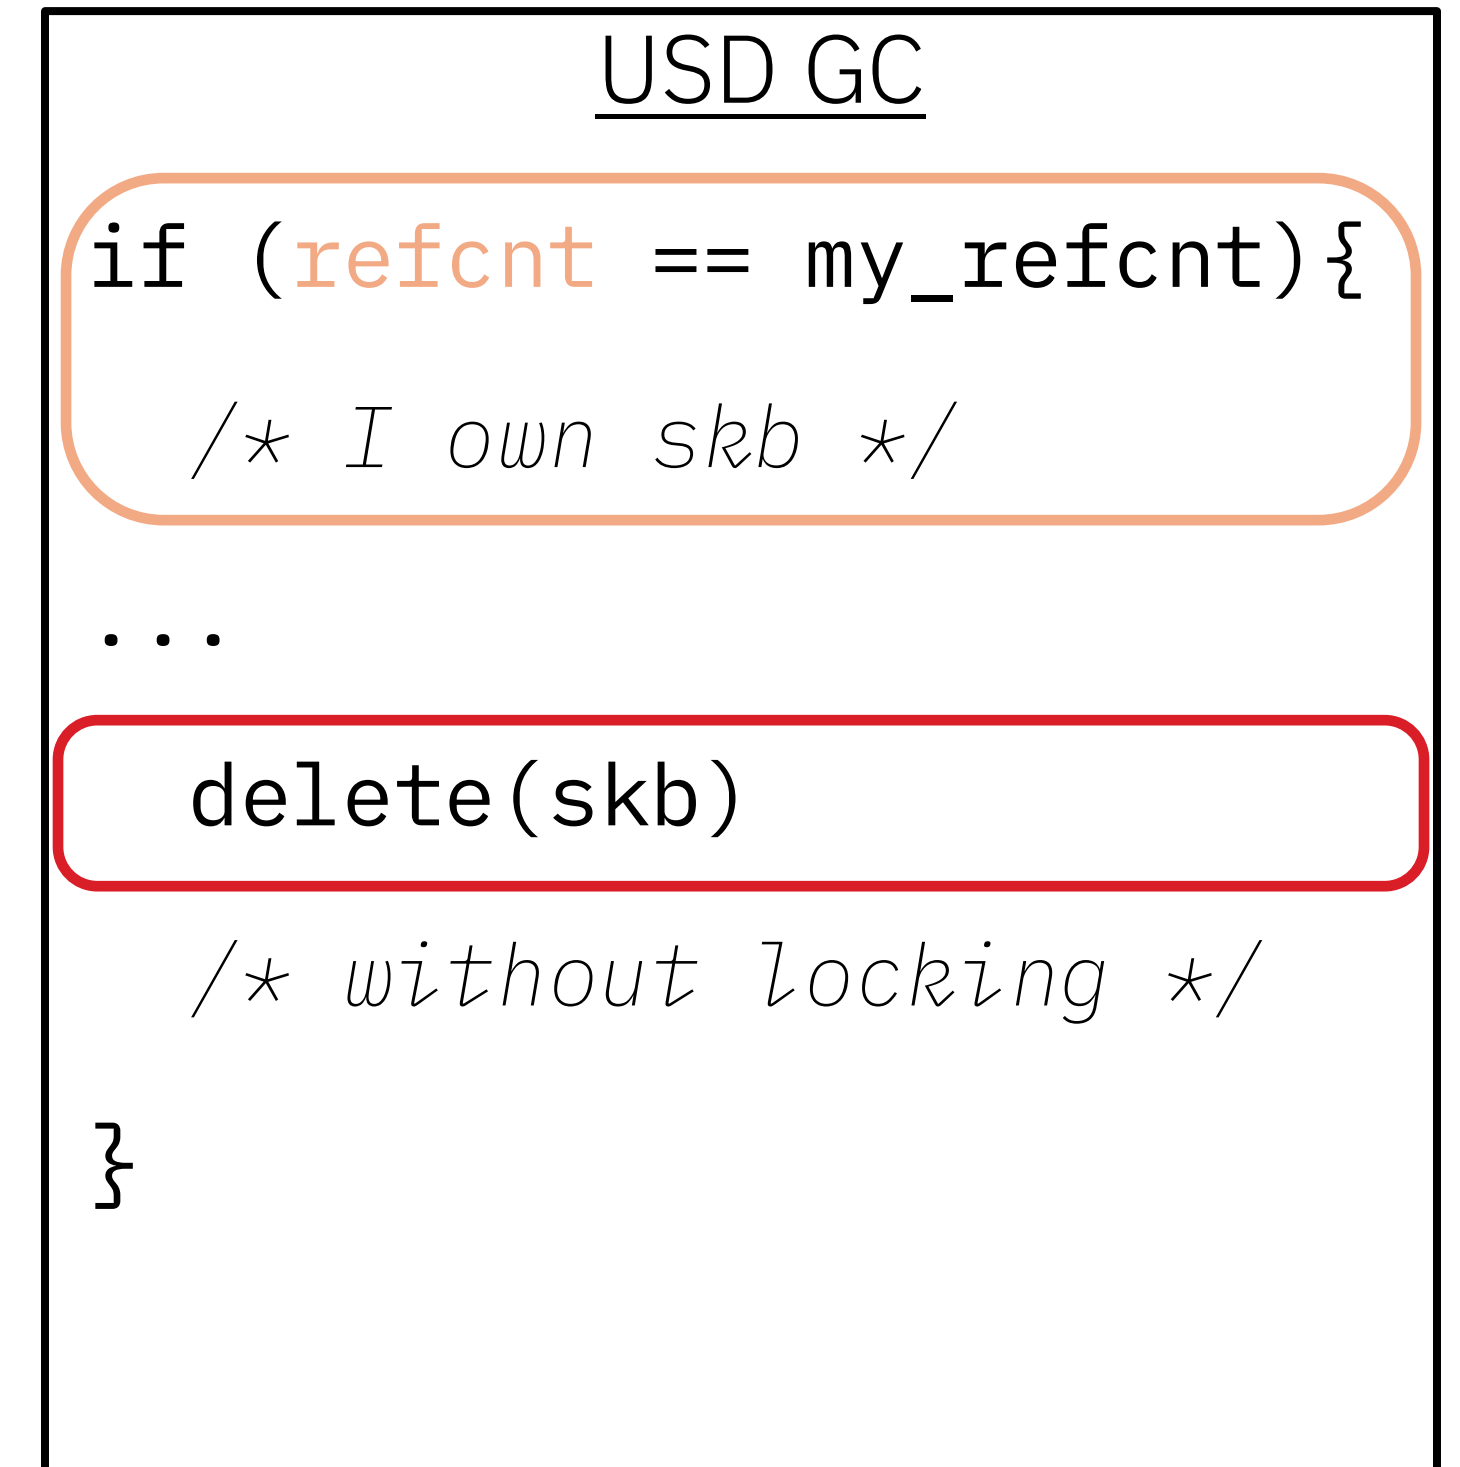
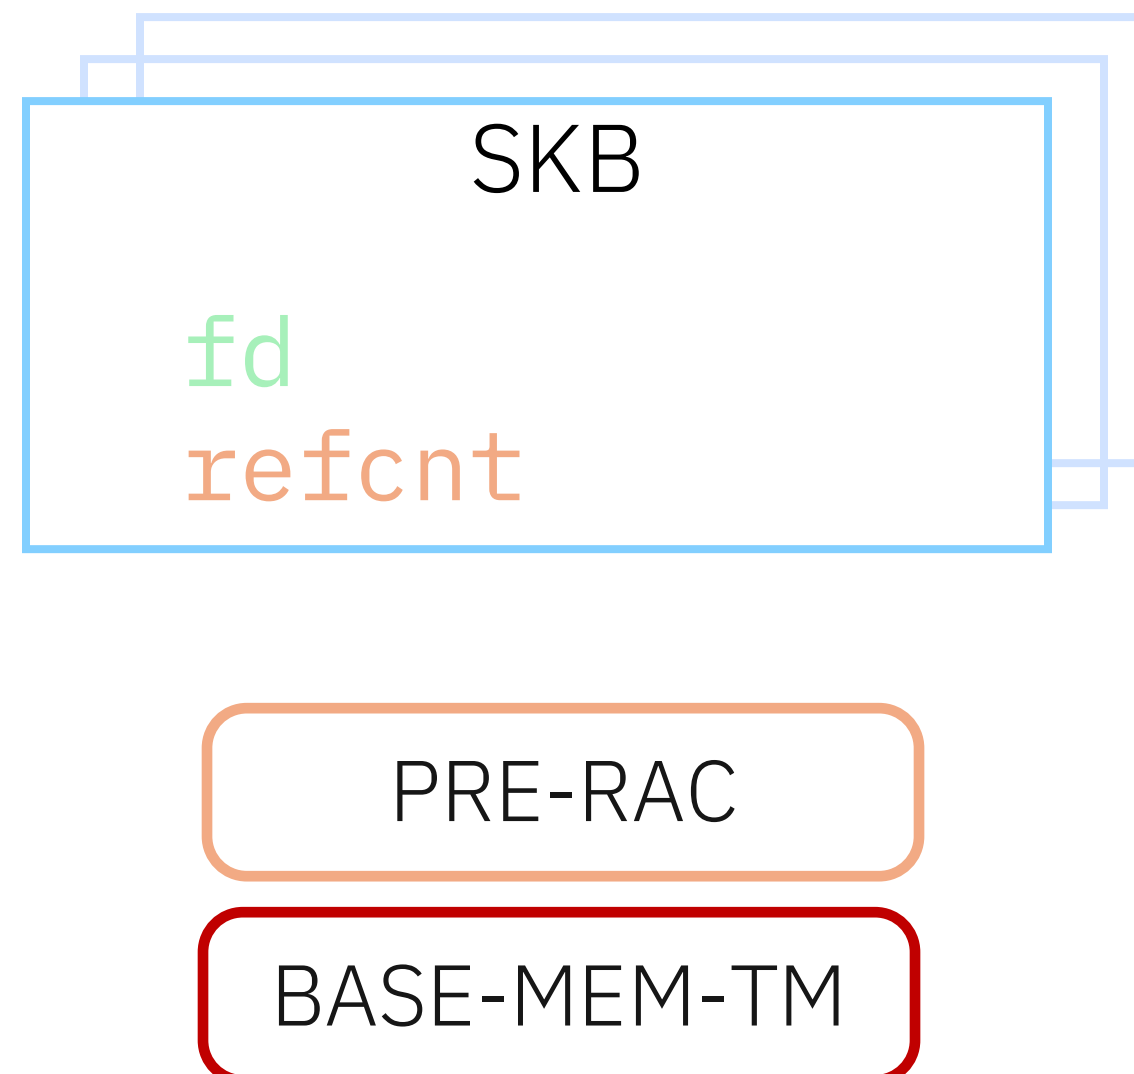
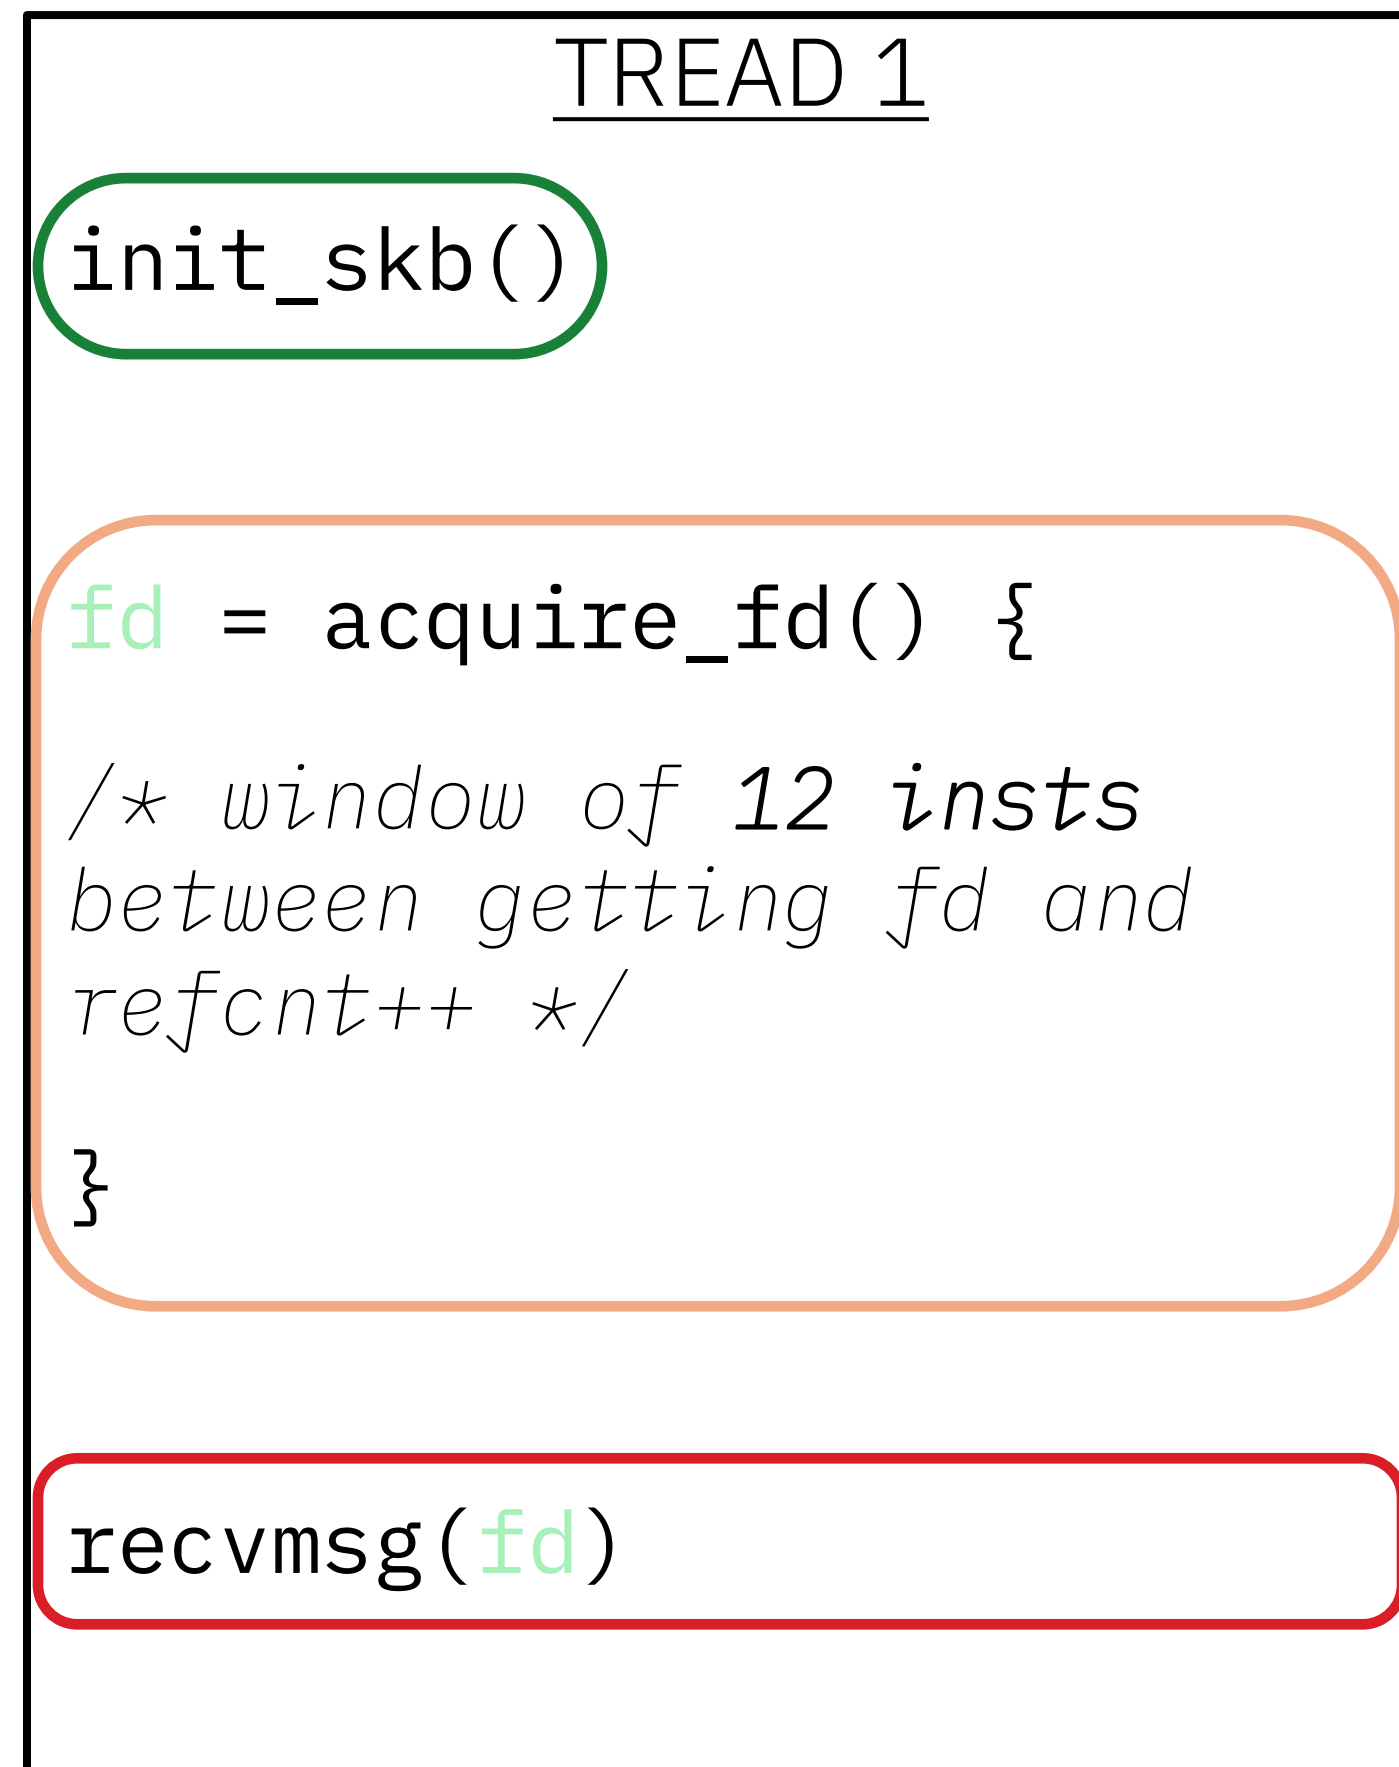
Let's now consider CVE-2021-4083



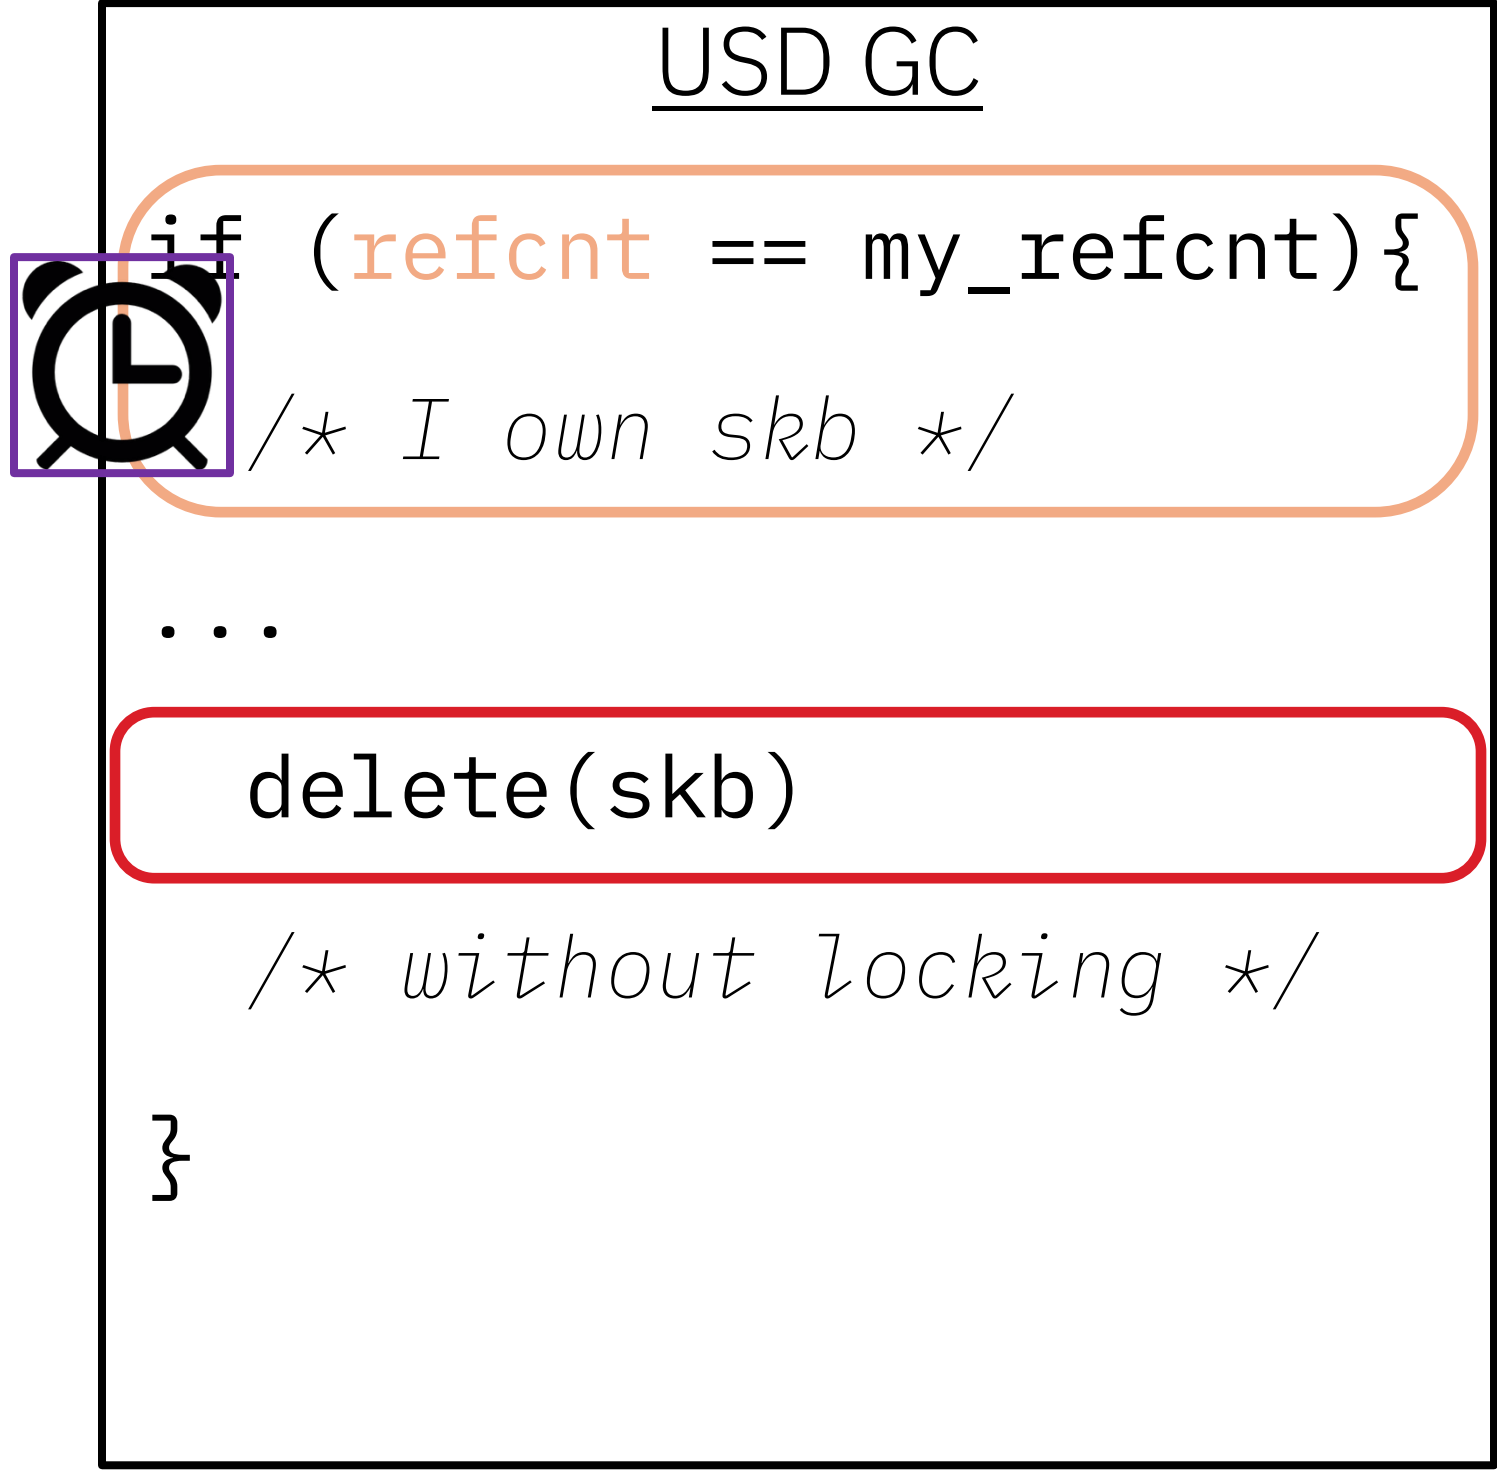

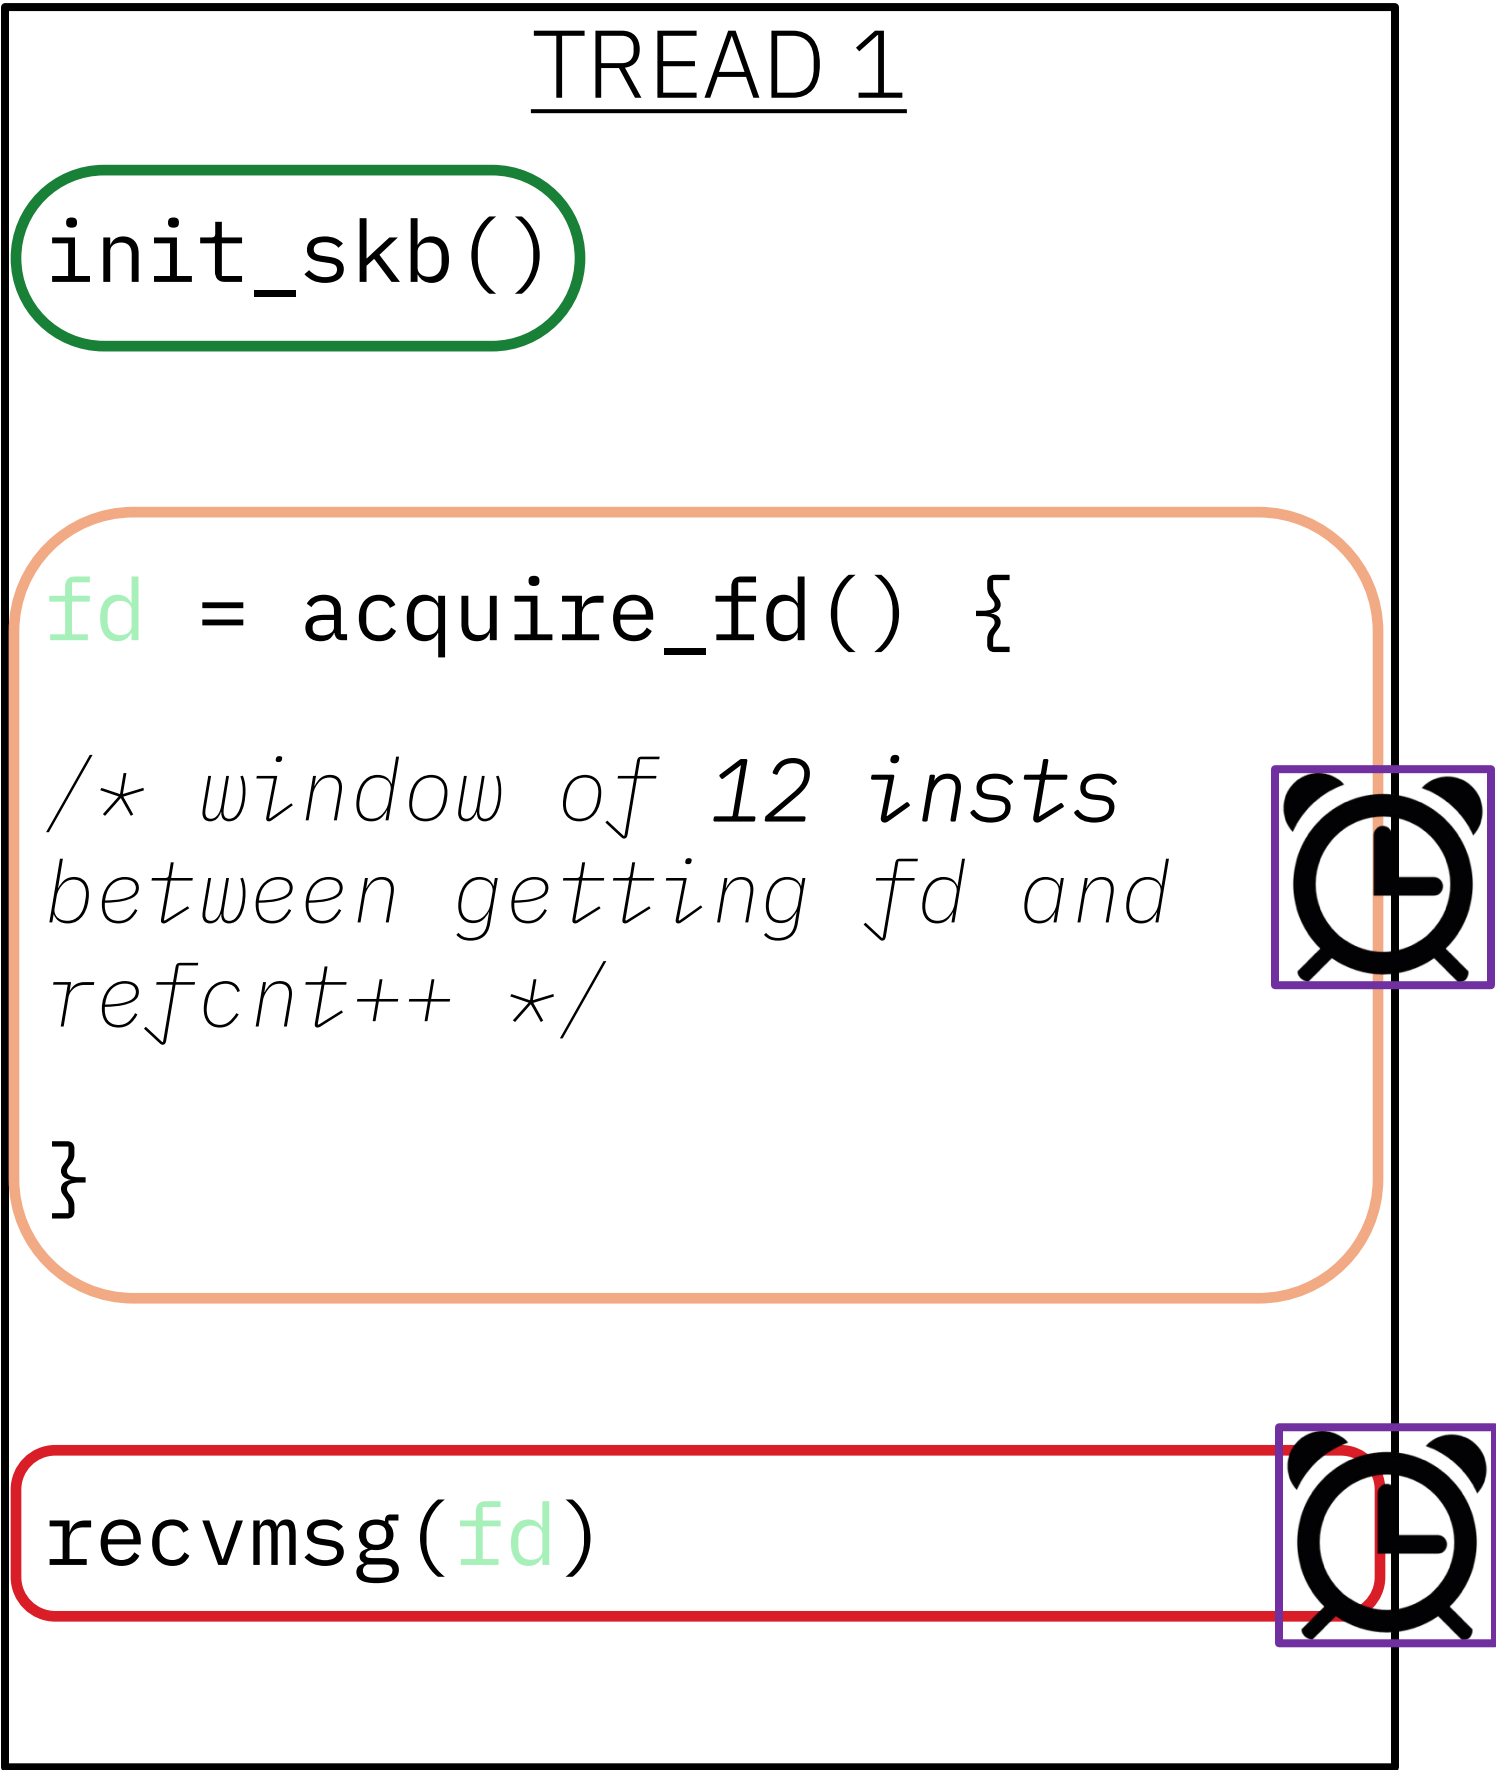
Let's now consider CVE-2021-4083



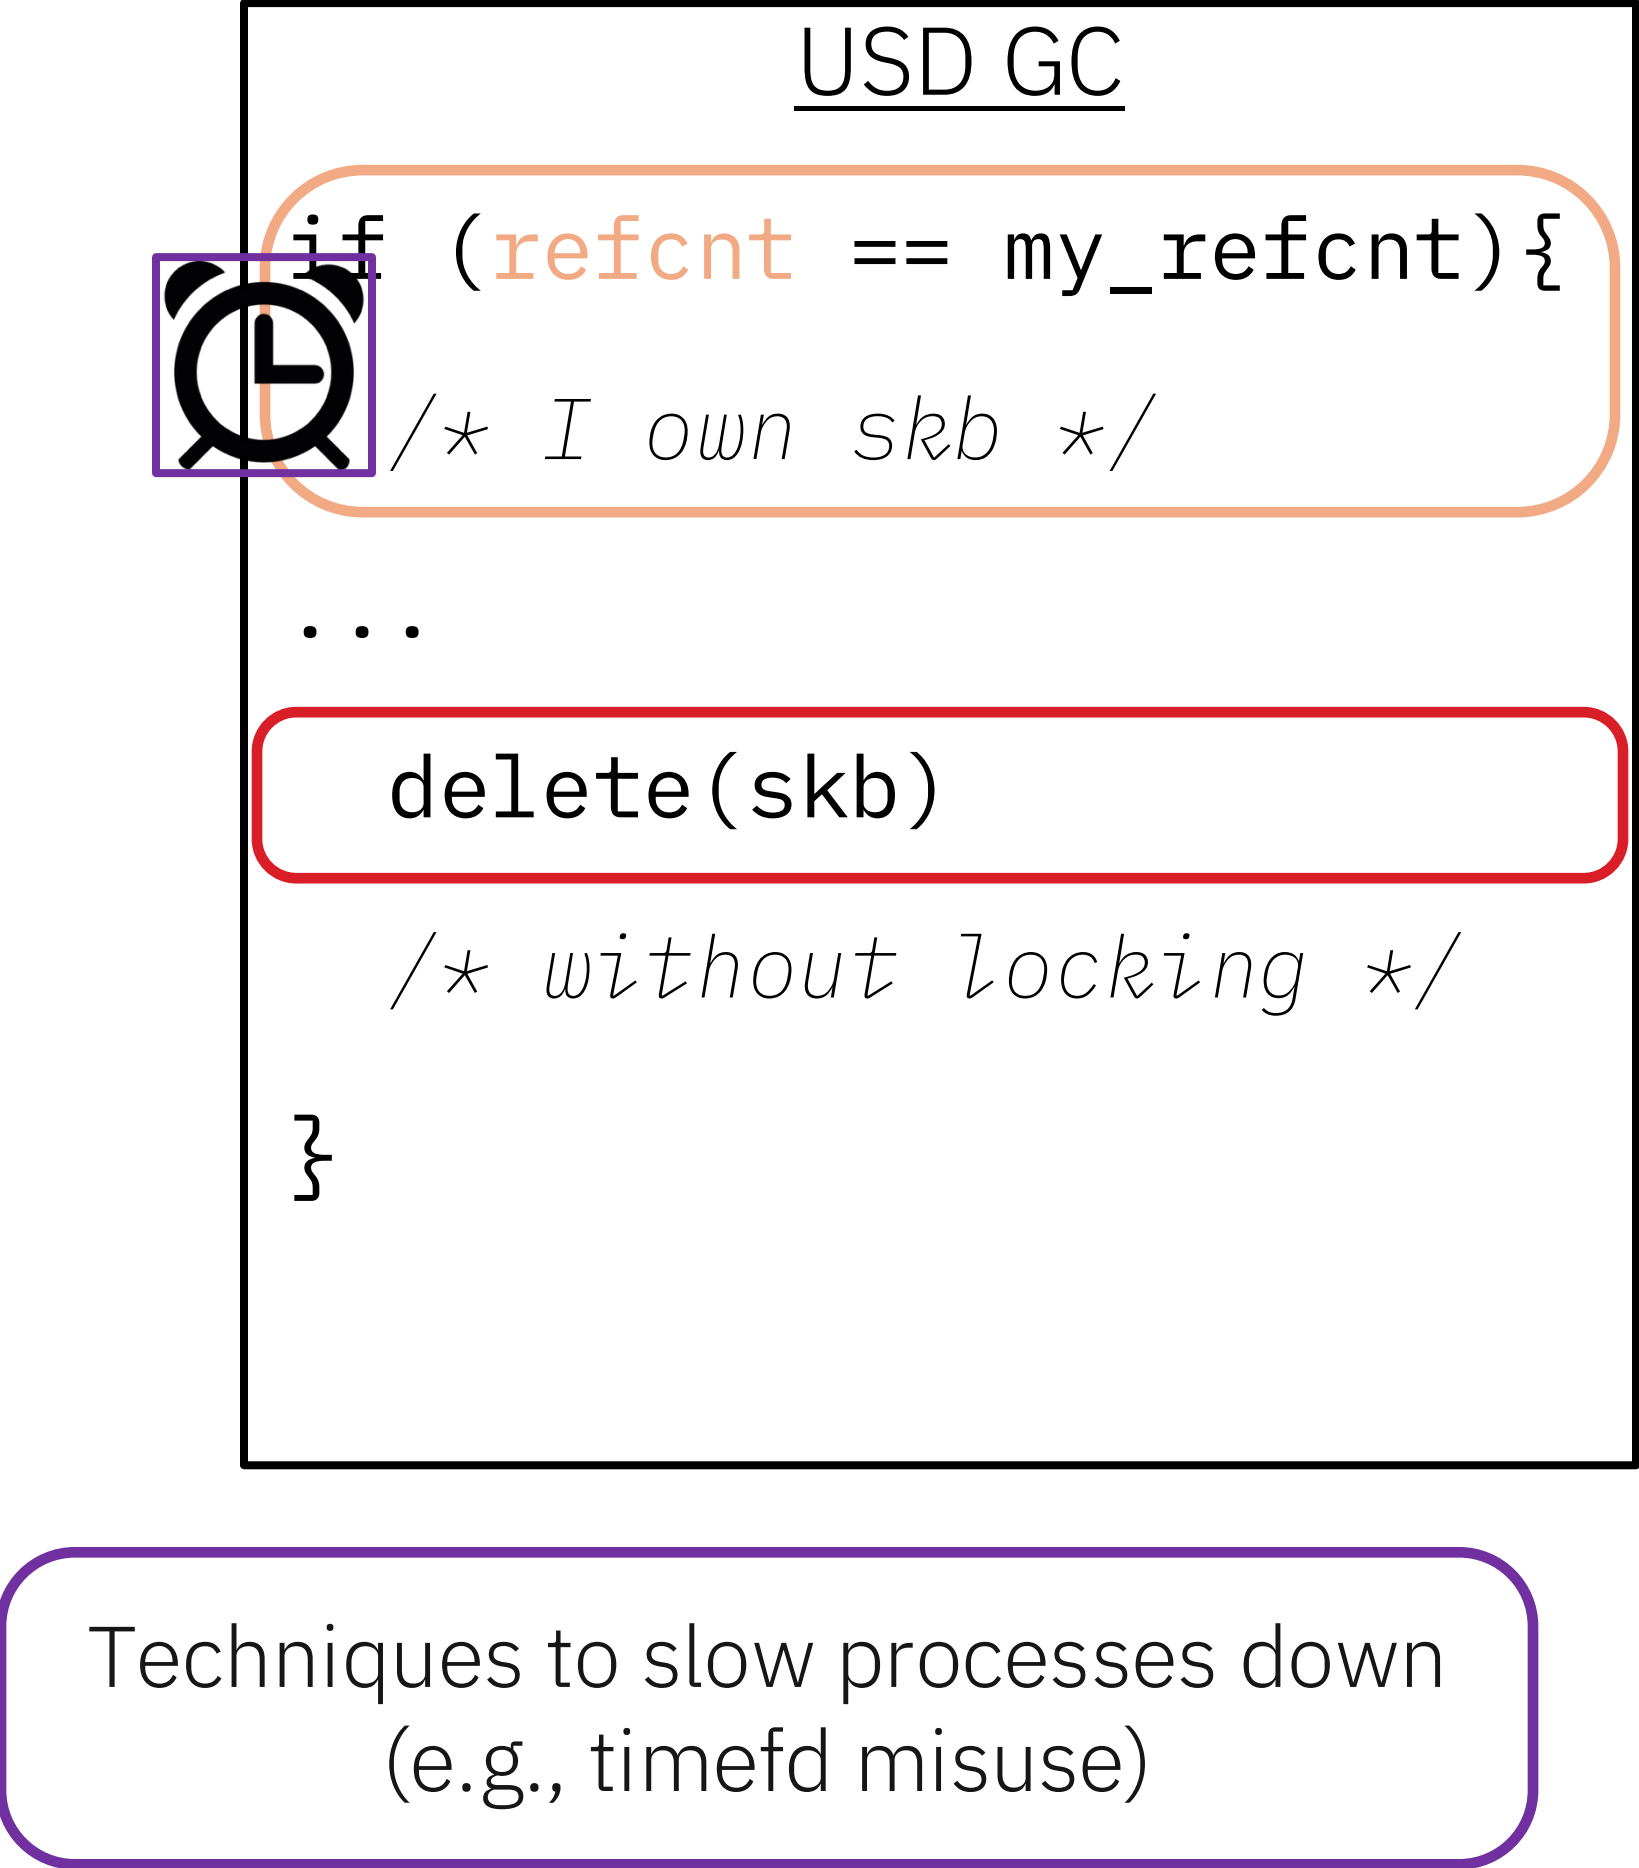


Let's now consider CVE-2021-4083



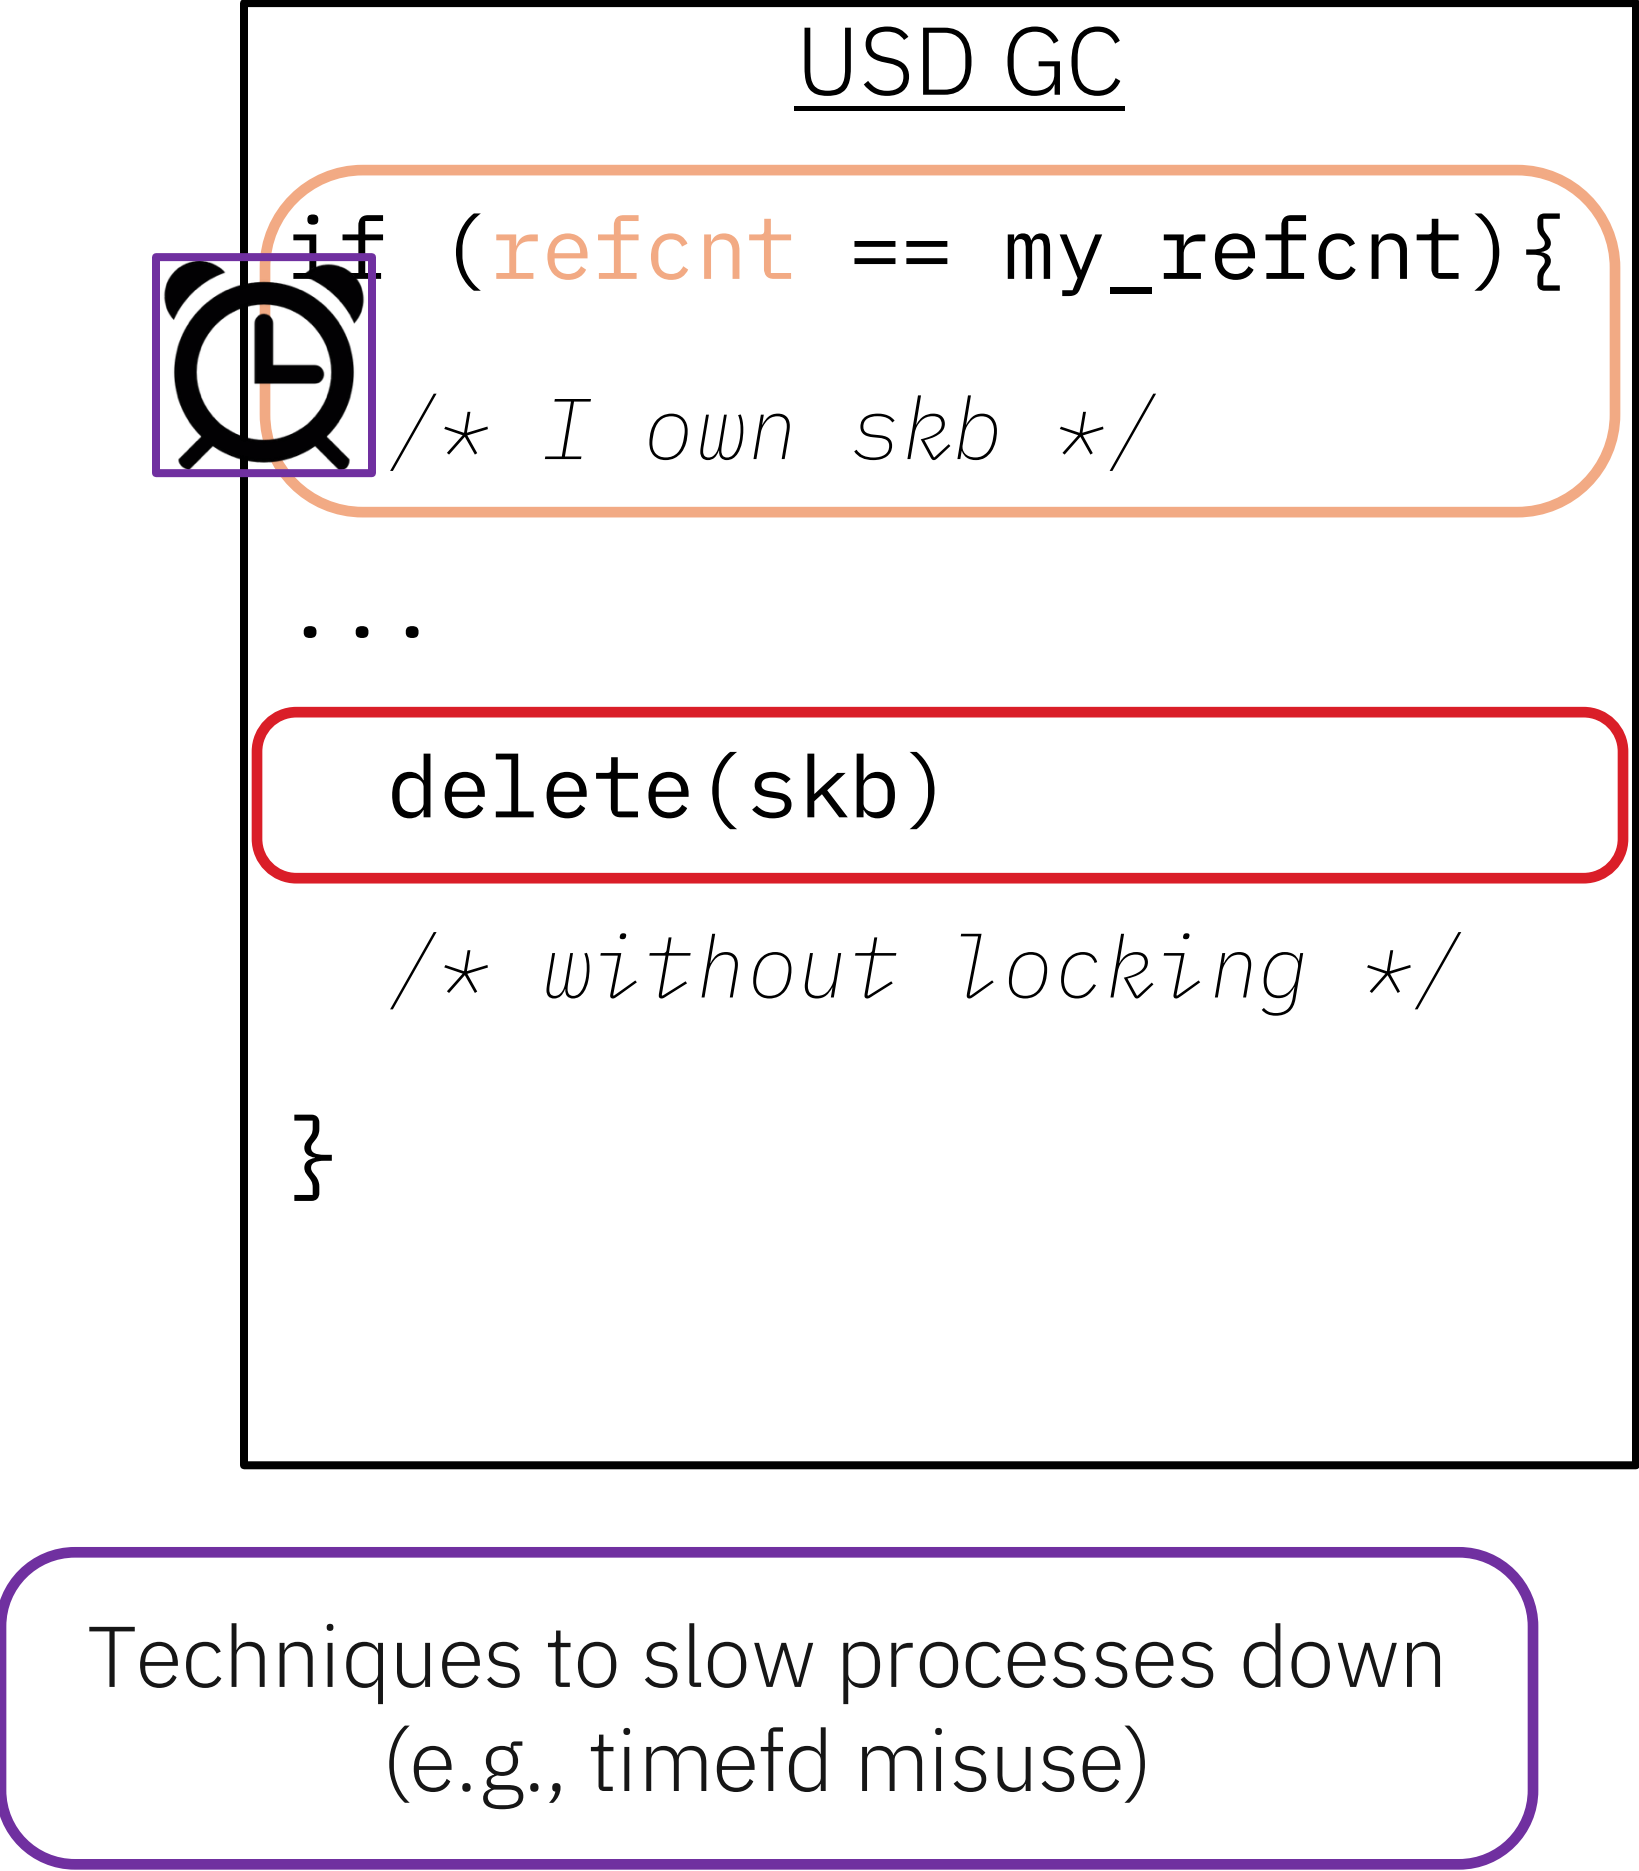
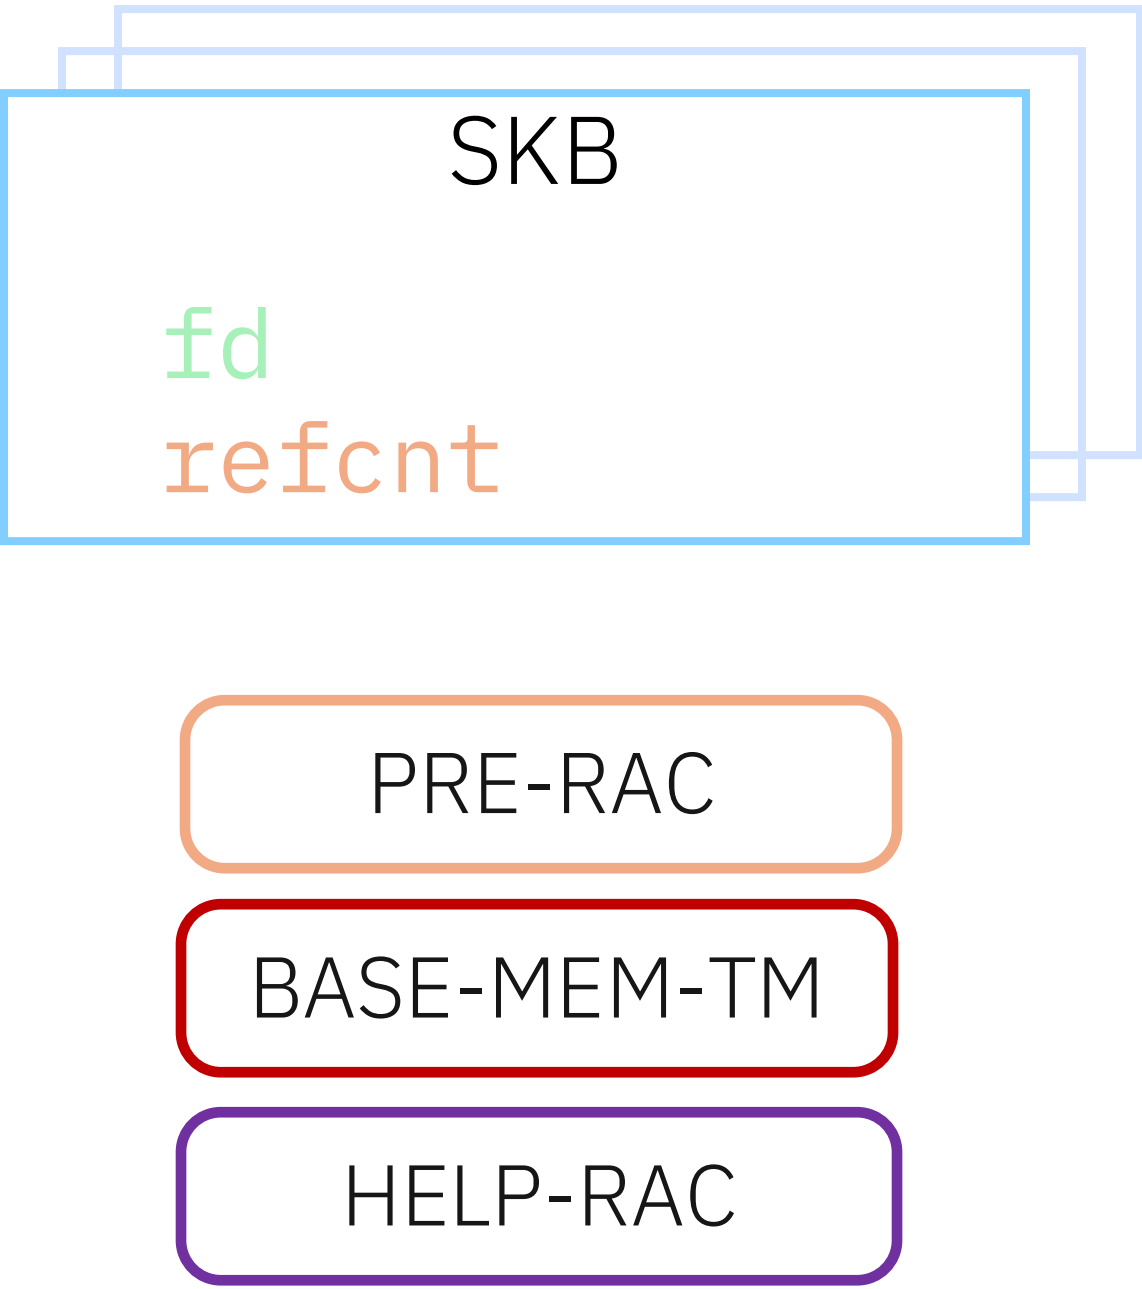

Let's now consider CVE-2021-4083



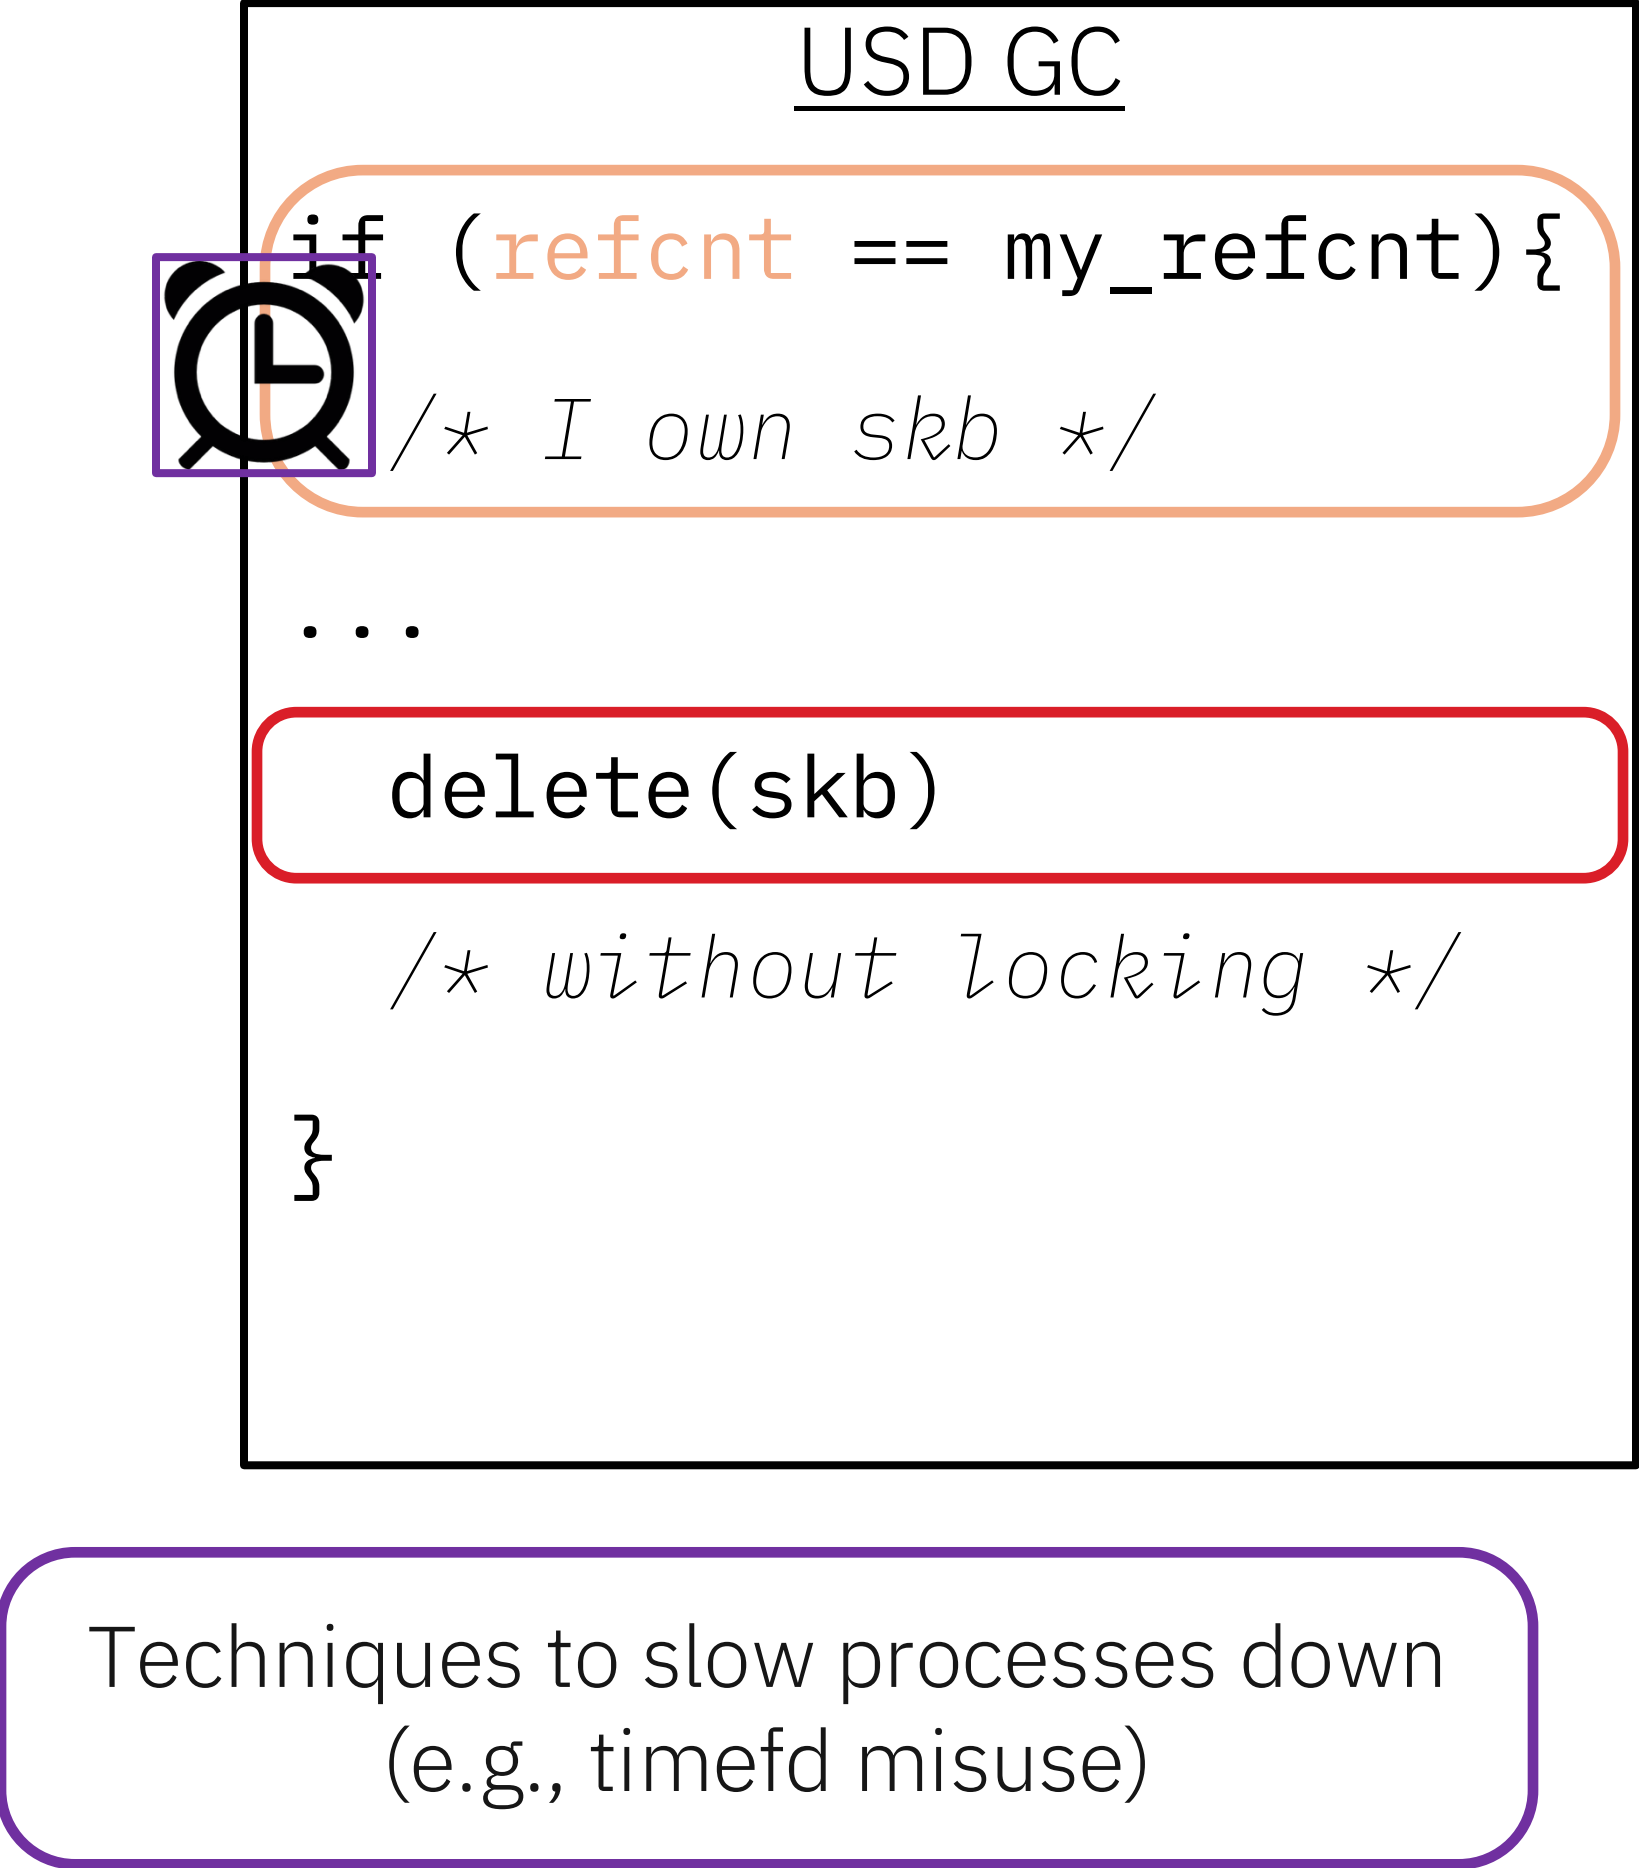
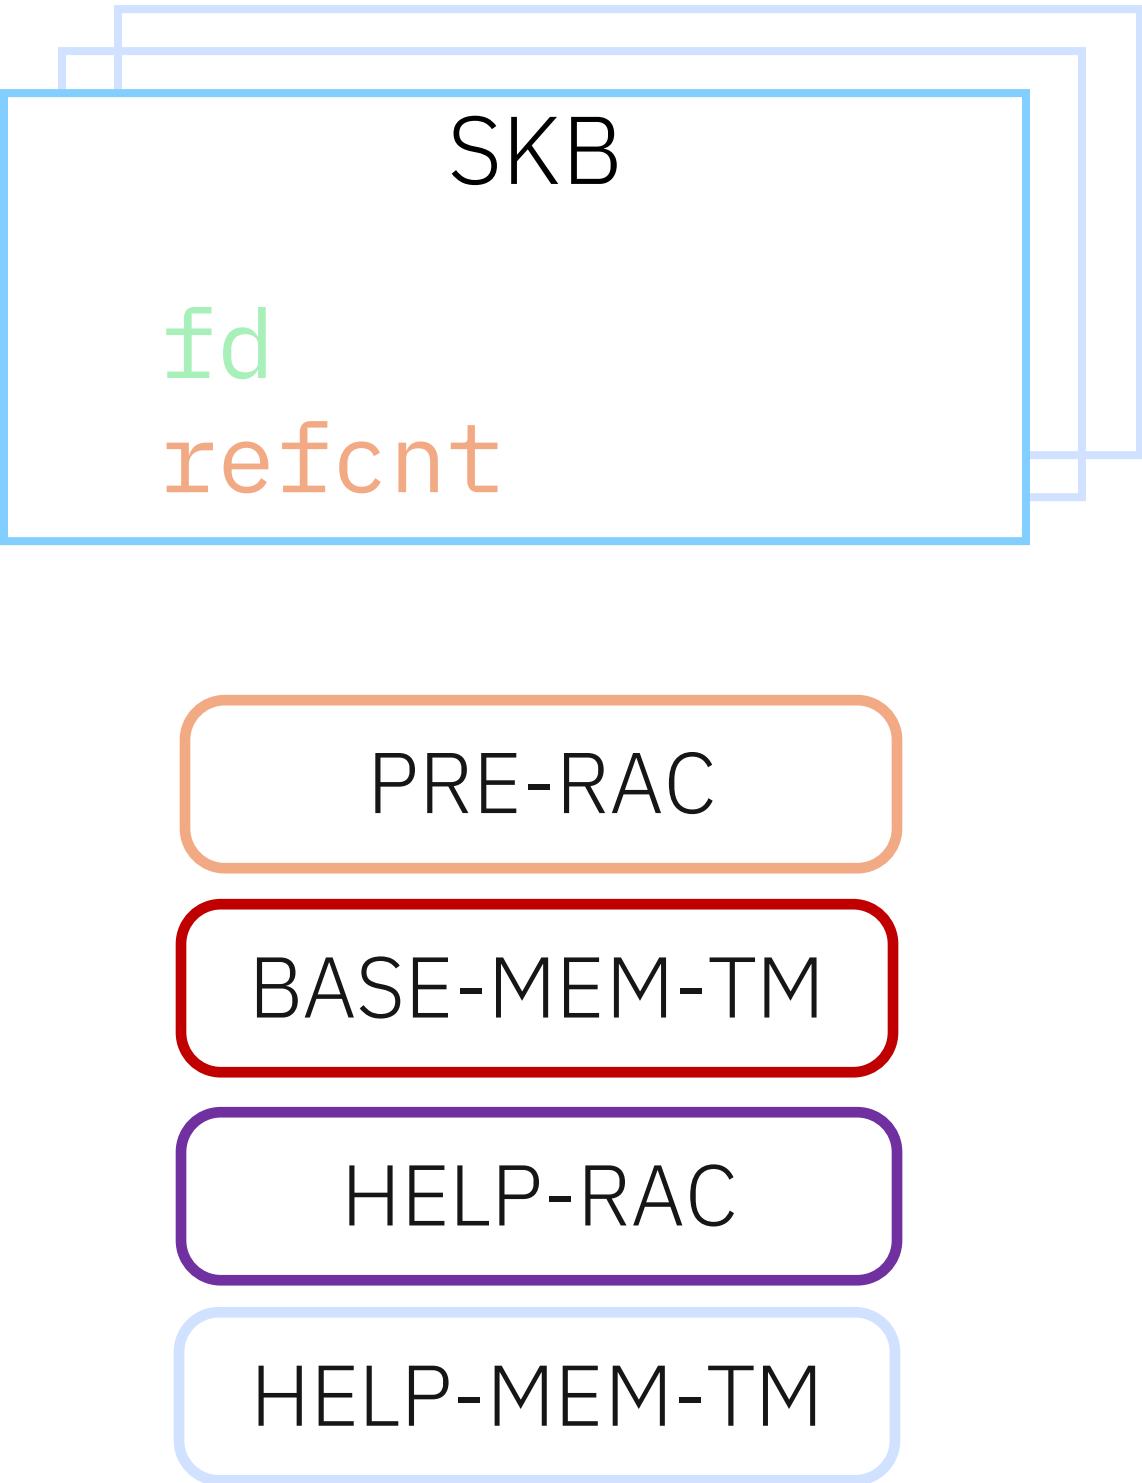

Let's now consider CVE-2021-4083



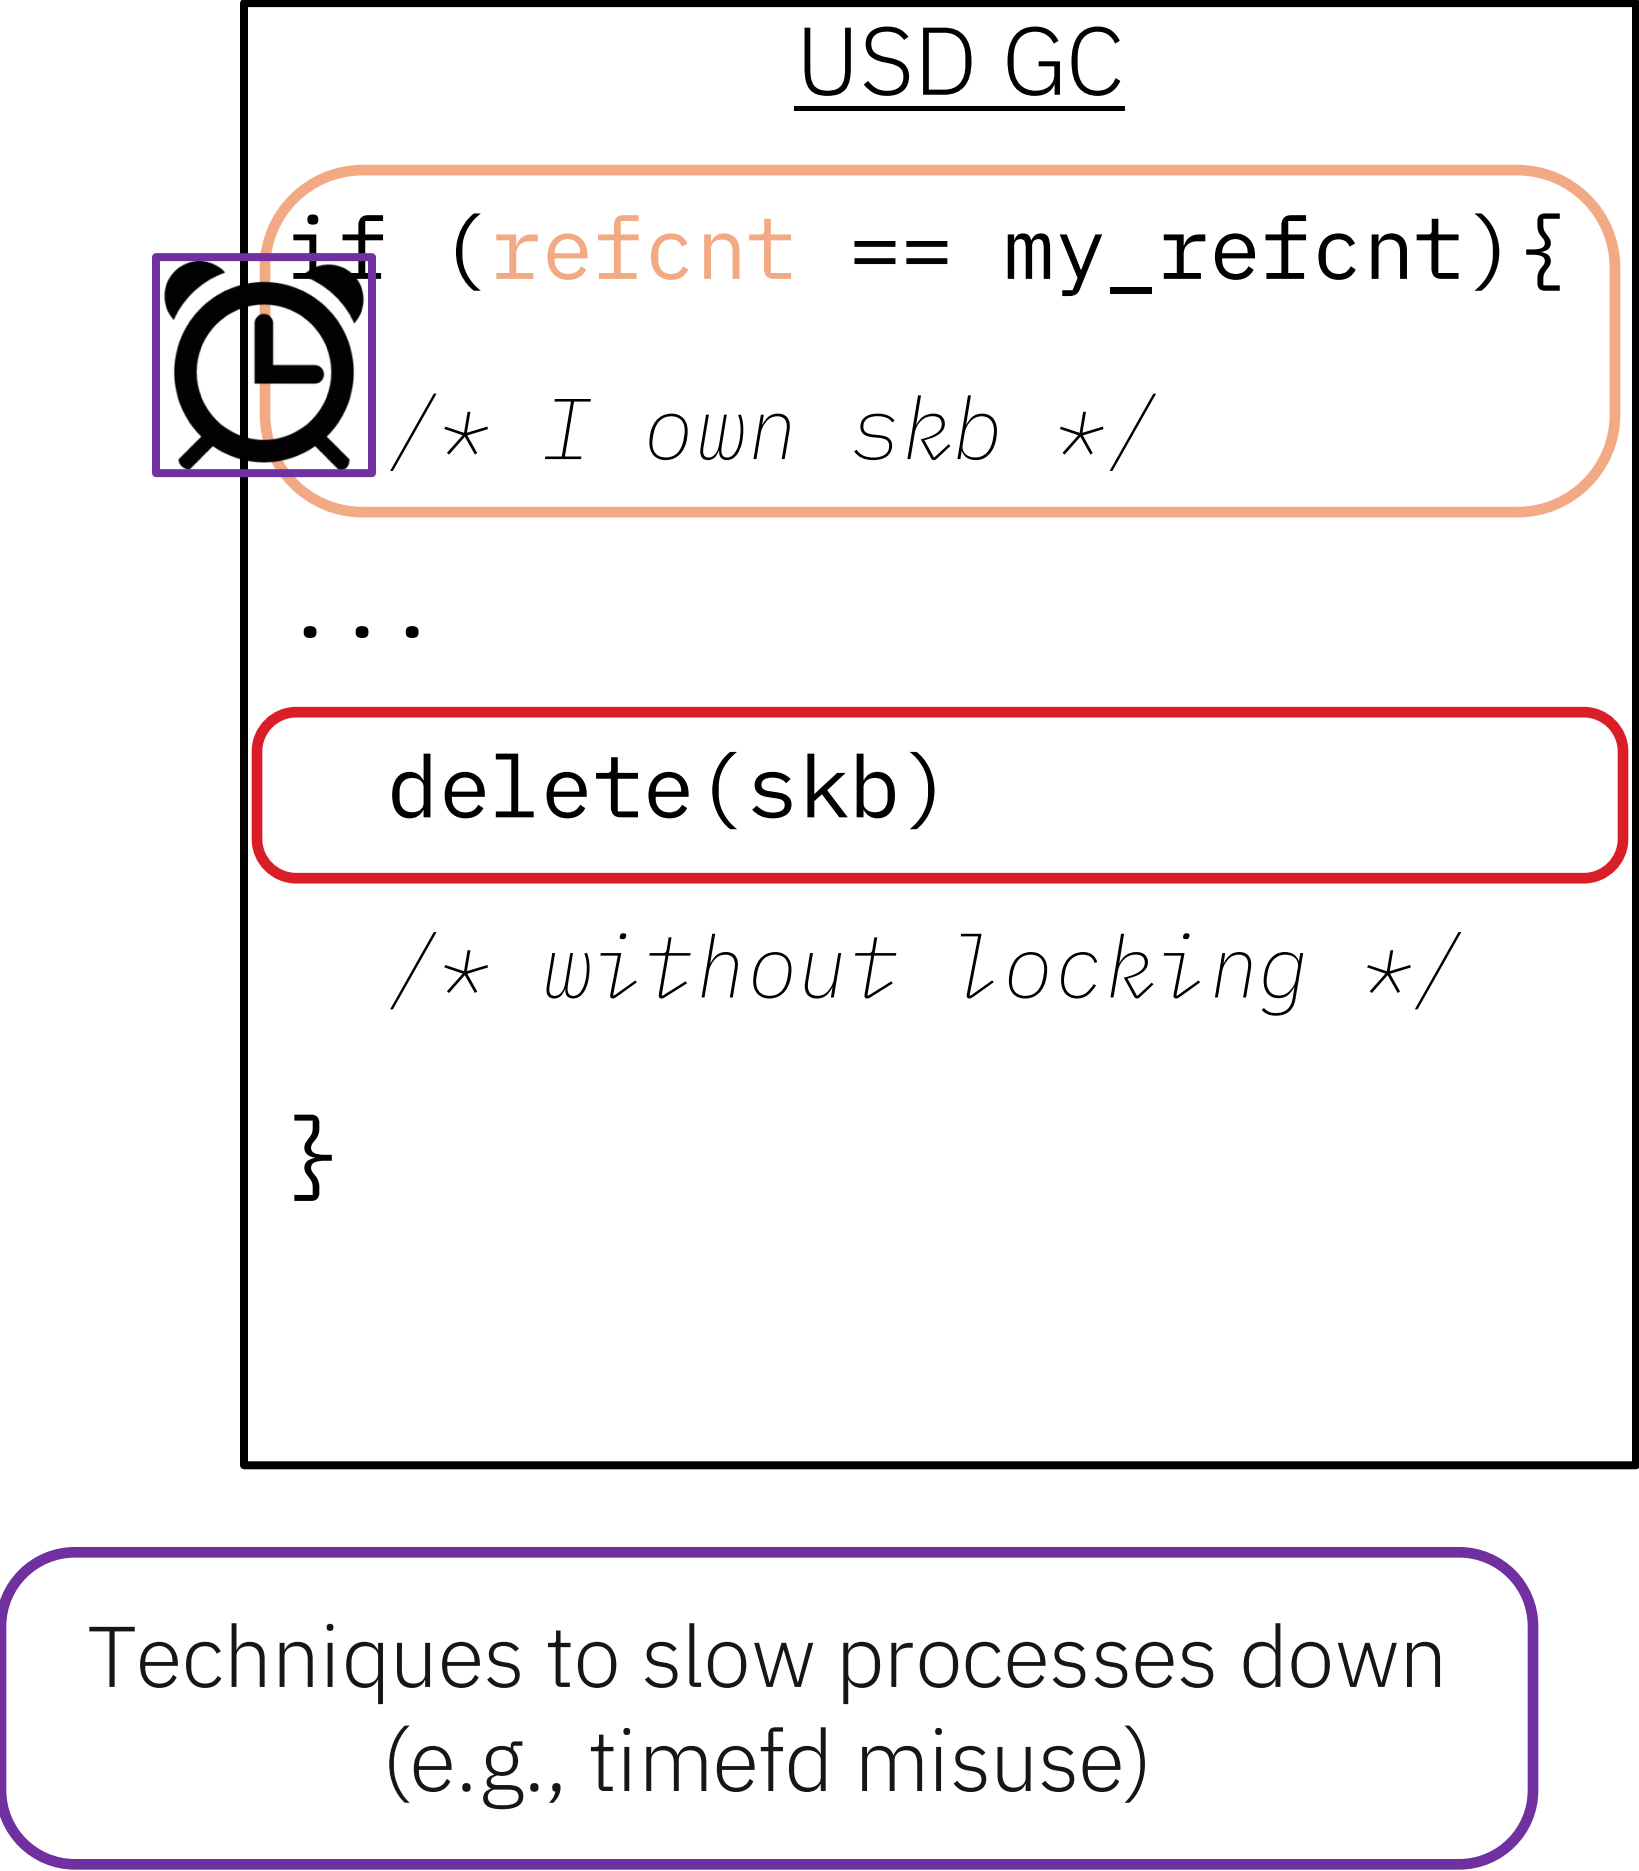
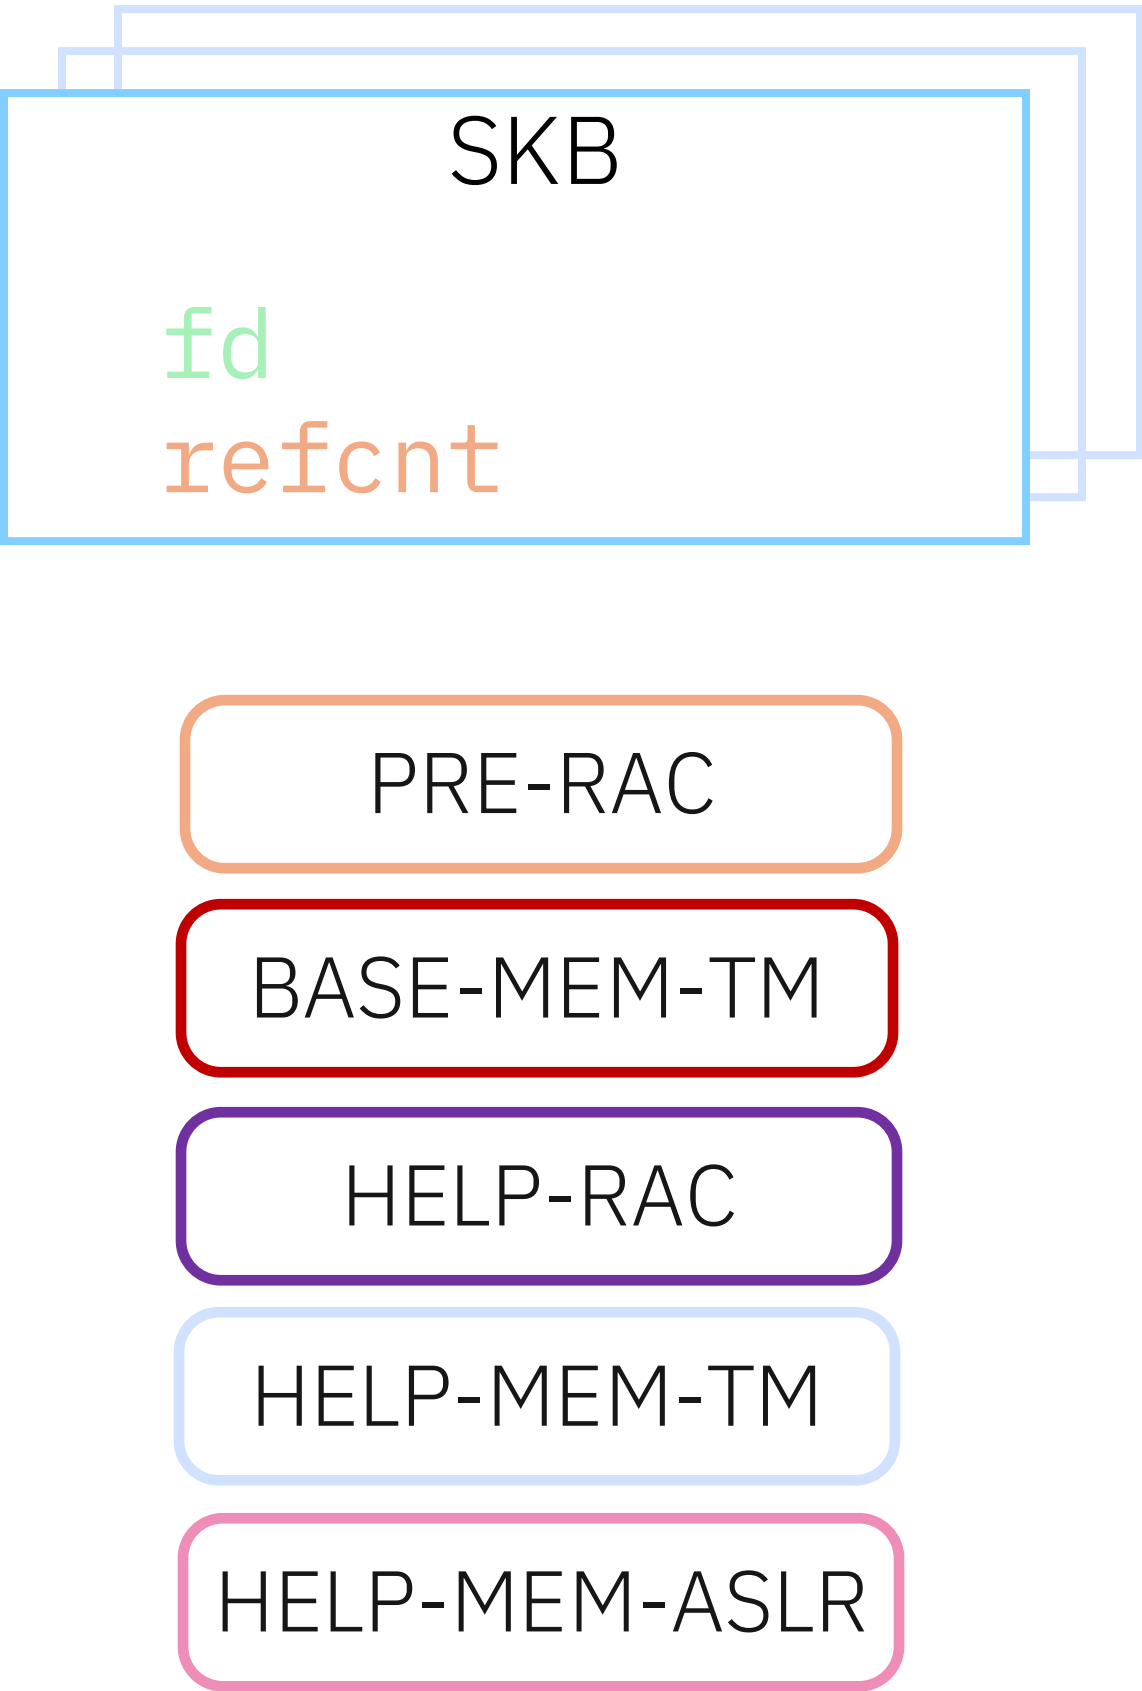

Let's now consider CVE-2021-4083



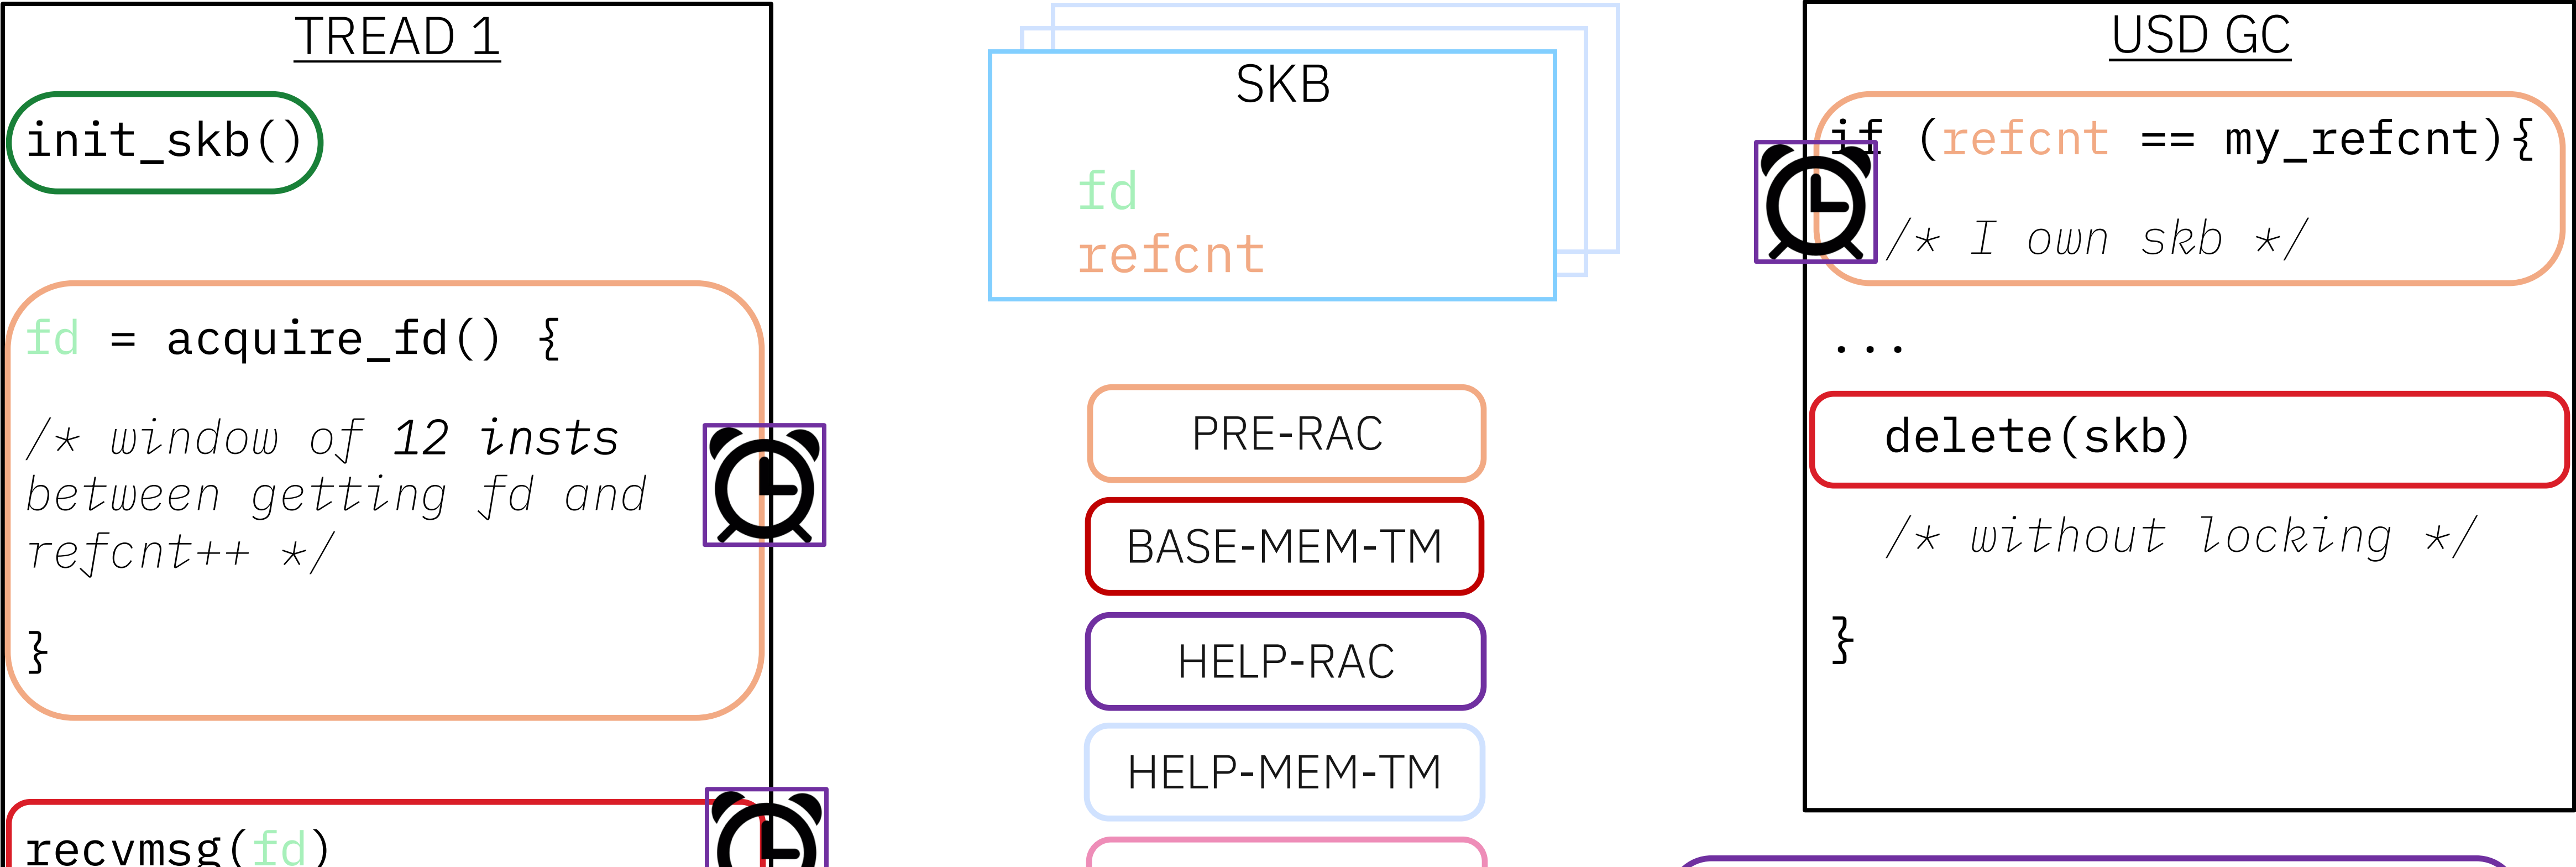
Let's now consider CVE-2021-4083




Let's now consider CVE-2021-4083



Let's now consider CVE-2021-4083



Improves way of thinking about complex exploits




Exploitation Steps

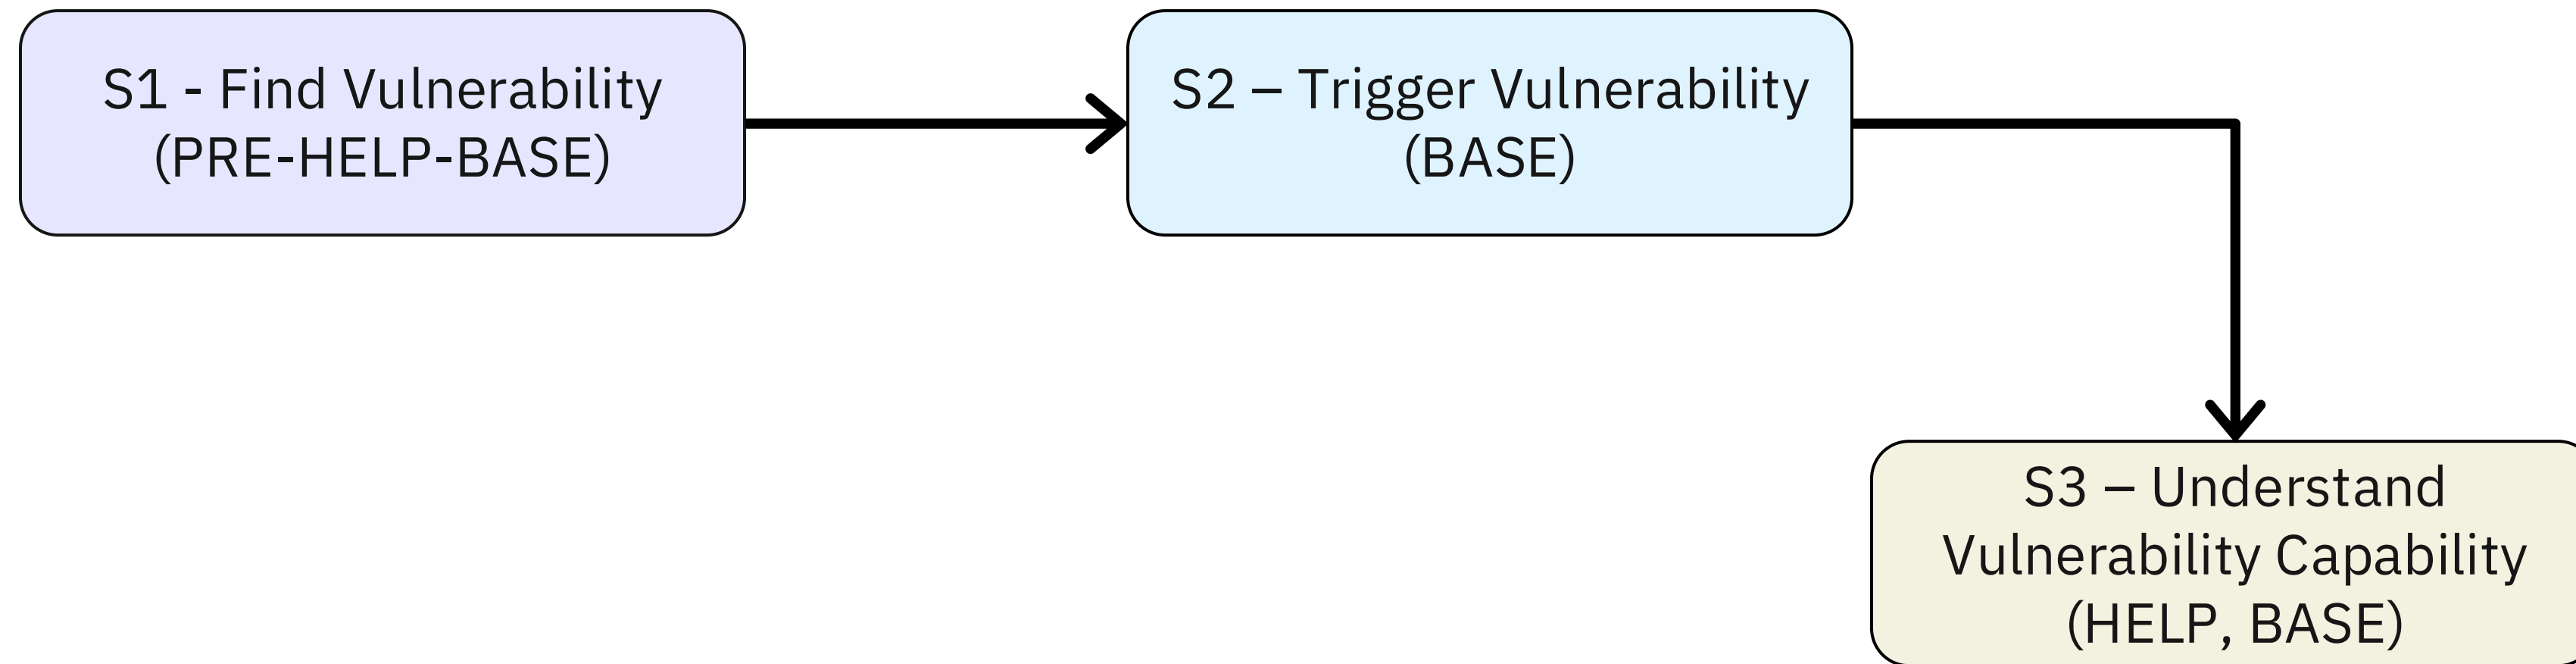
Exploitation Steps

S1 - Find Vulnerability
(PRE-HELP-BASE)

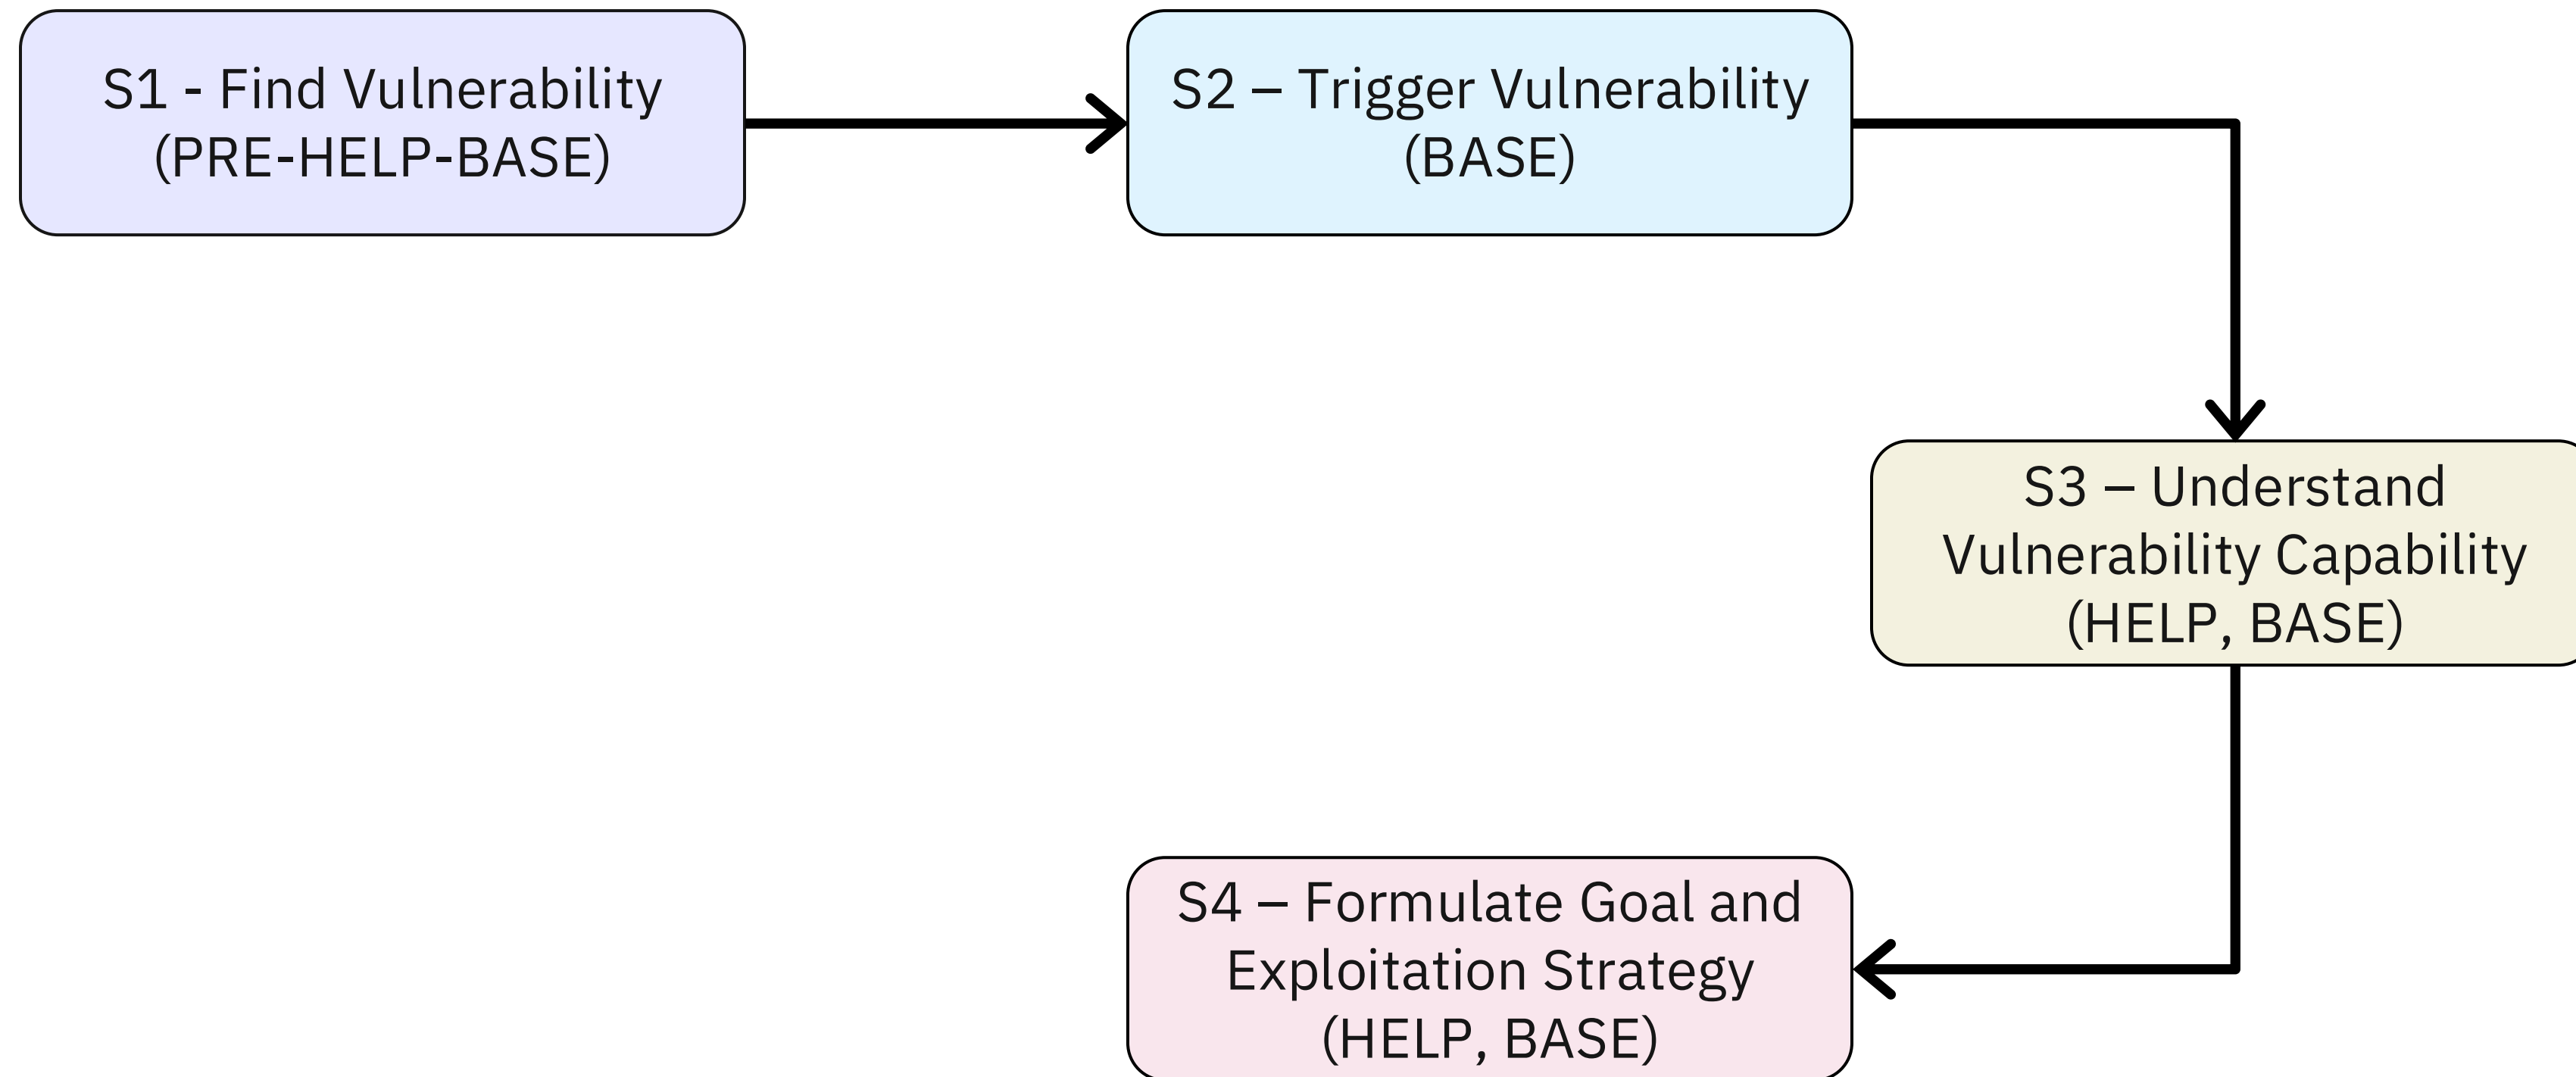
Exploitation Steps



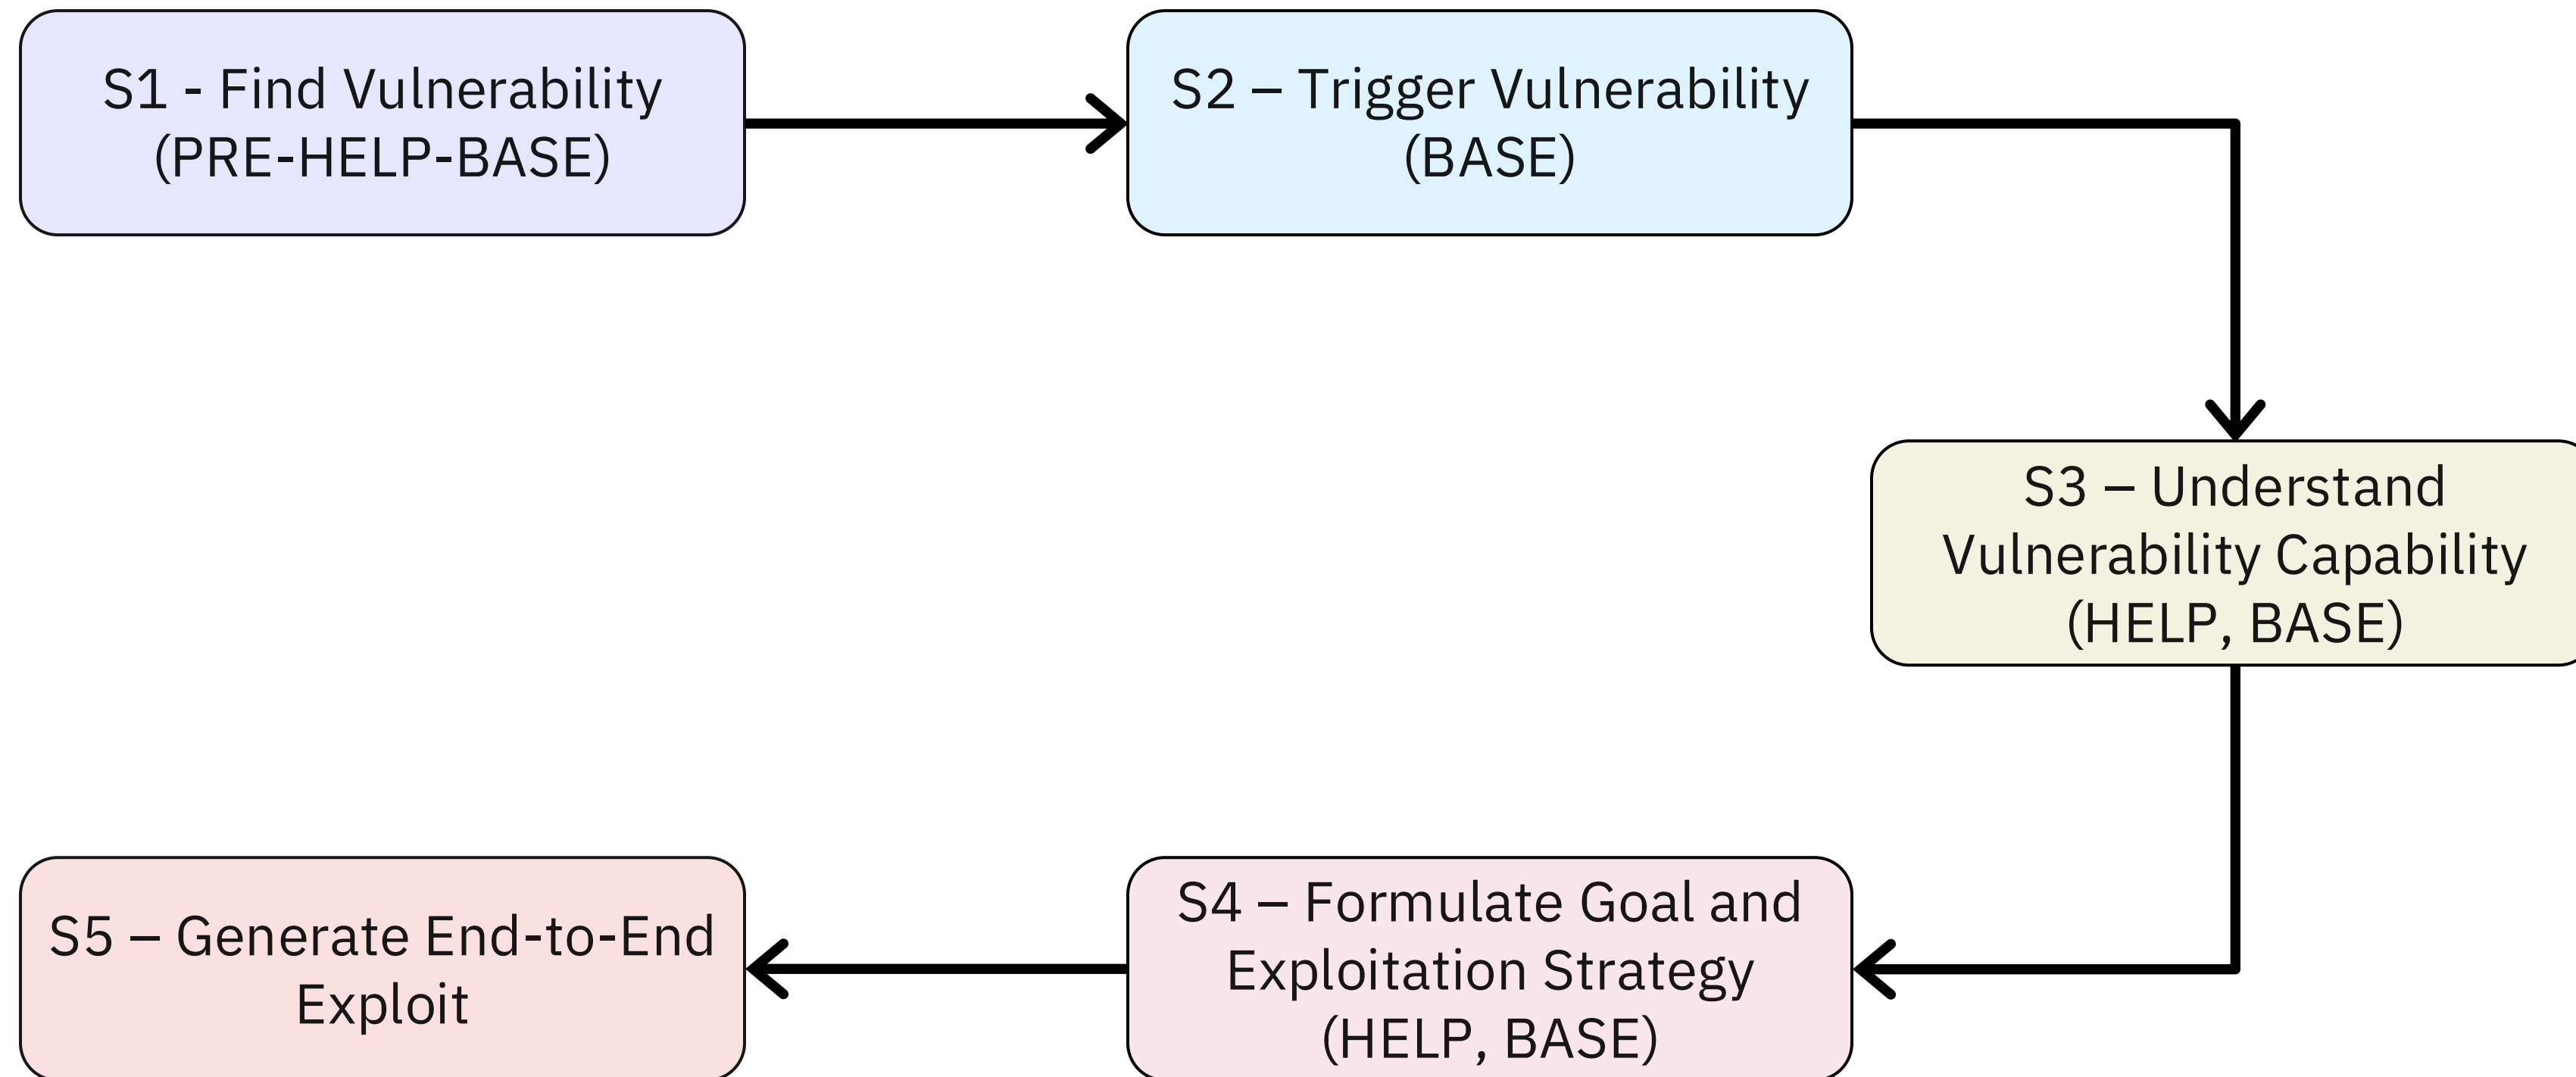
Exploitation Steps



Exploitation Steps



Exploitation Steps



SoK – methodologies & statistics

First: paper from top tier venues from 2013 to 2023.

Second: collected cited work not included in first selection.

Filtered for new techniques or automation of some steps of kernel exploitation

SoK – methodologies & statistics

First: paper from top tier venues from 2013 to 2023.

Second: collected cited work not included in first selection.

Filtered for new techniques or automation of some steps of kernel exploitation



41 TOTAL



armasuisse




SoK – methodologies & statistics

First: paper from top tier venues from 2013 to 2023.

Second: collected cited work not included in first selection.

Filtered for new techniques or automation of some steps of kernel exploitation




SoK – methodologies & statistics

First: paper from top tier venues from 2013 to 2023.

Second: collected cited work not included in first selection.

Filtered for new techniques or automation of some steps of kernel exploitation

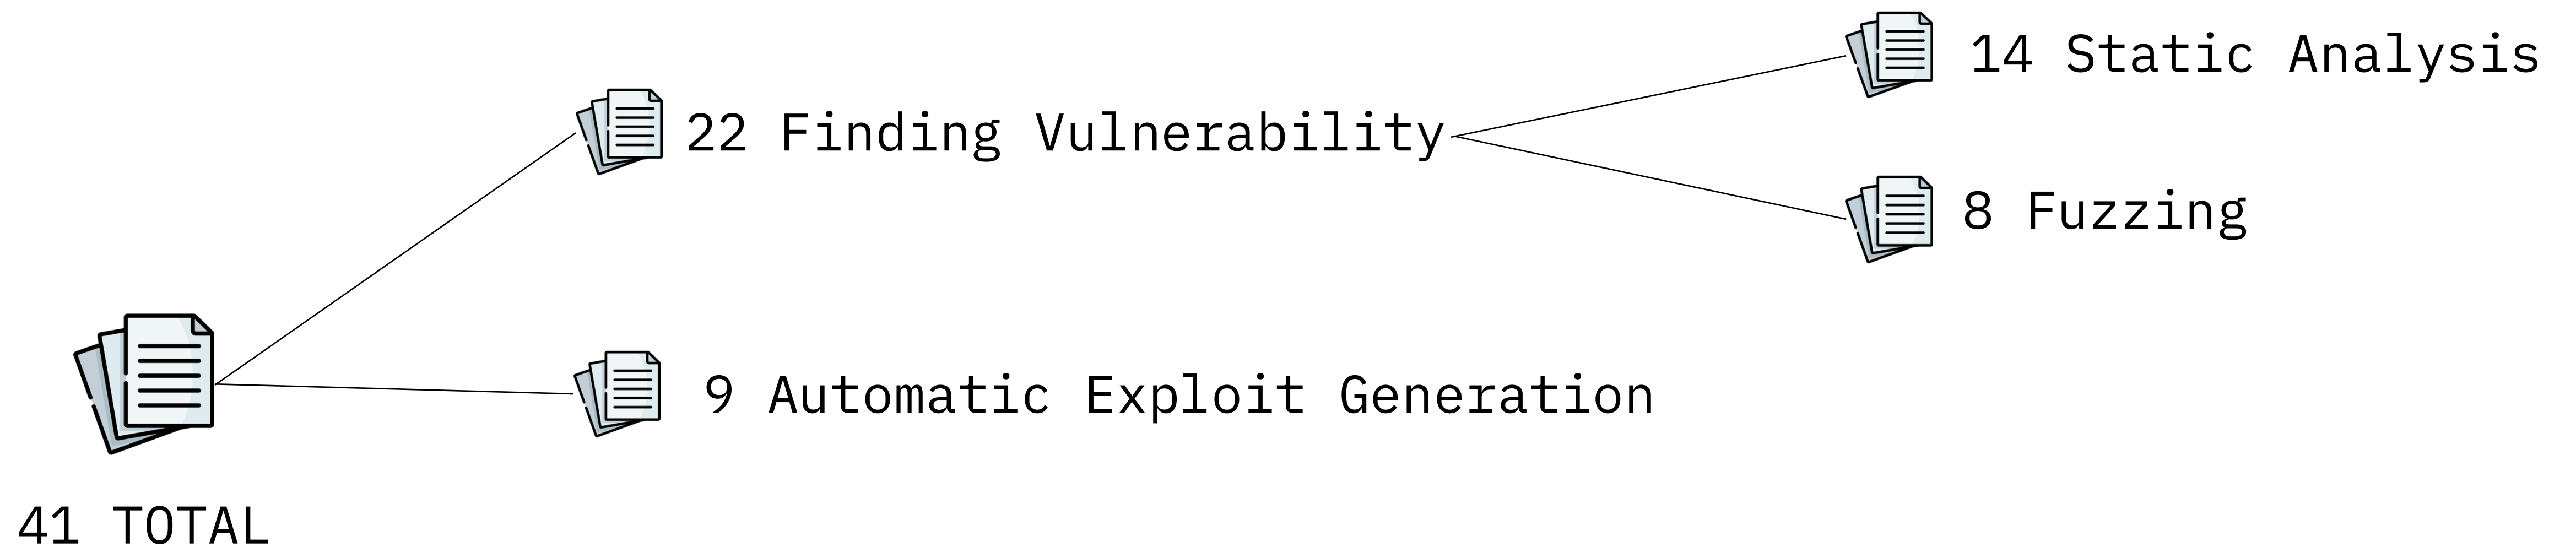


SoK – methodologies & statistics

First: paper from top tier venues from 2013 to 2023.

Second: collected cited work not included in first selection.

Filtered for new techniques or automation of some steps of kernel exploitation

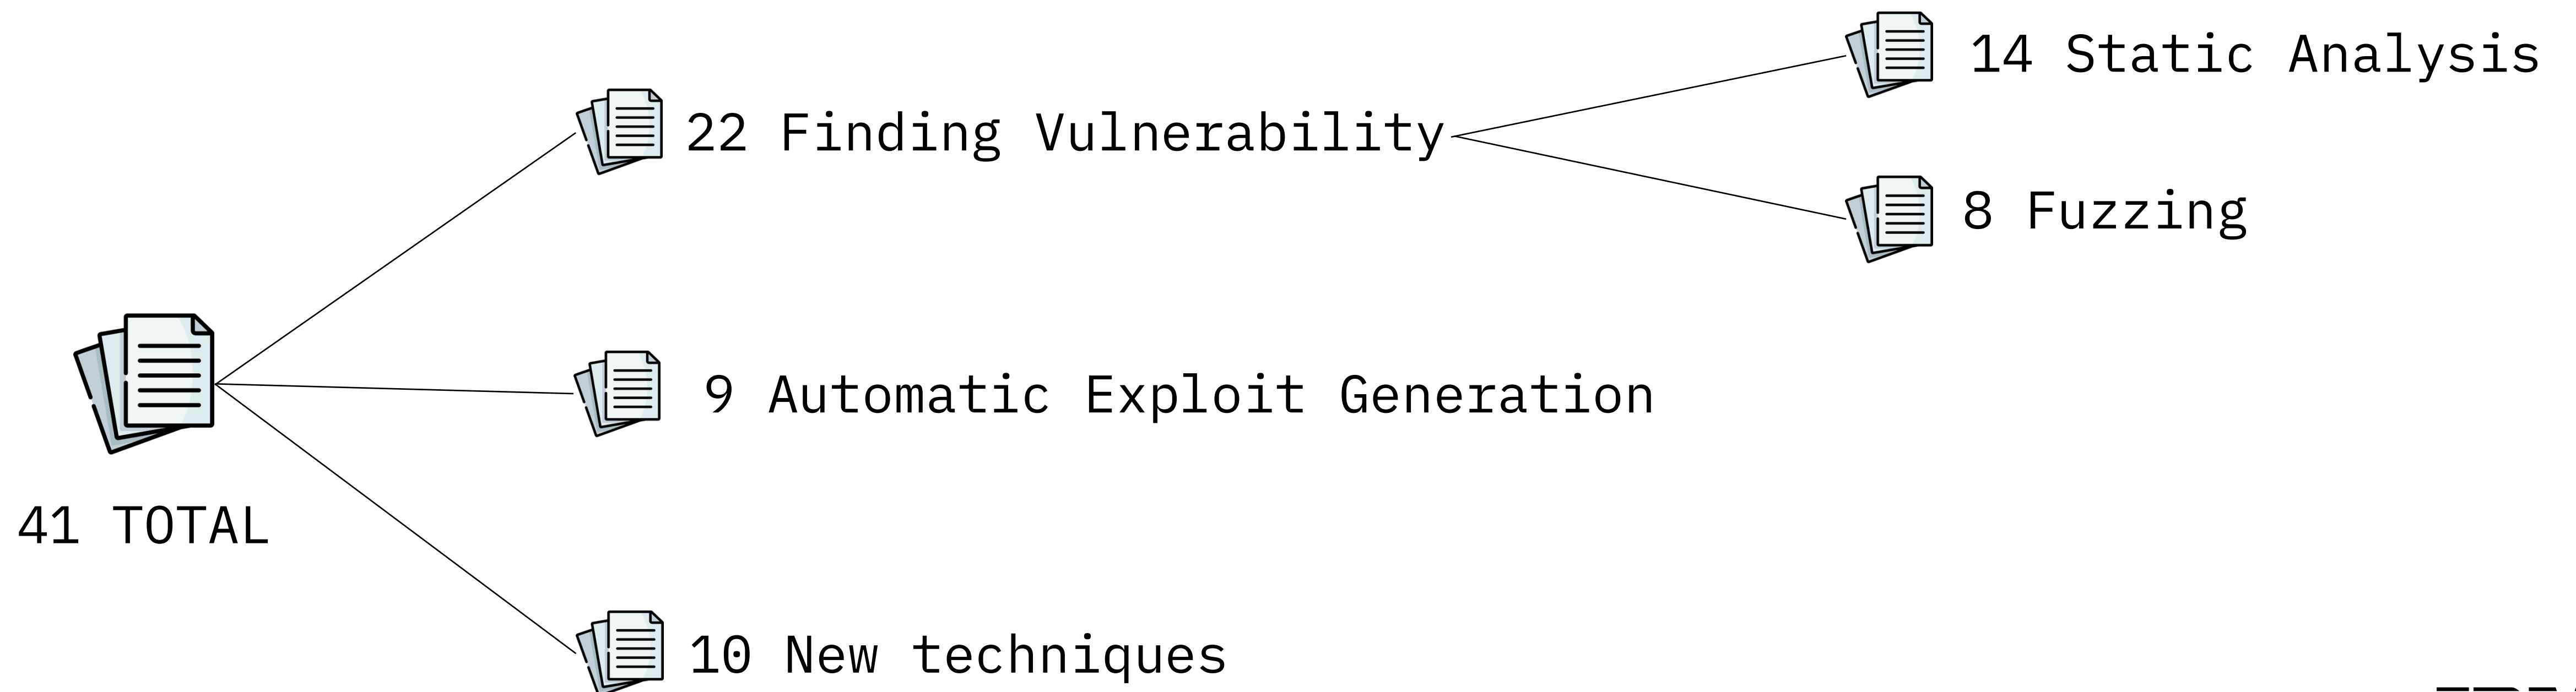


SoK – methodologies & statistics

First: paper from top tier venues from 2013 to 2023.

Second: collected cited work not included in first selection.

Filtered for new techniques or automation of some steps of kernel exploitation



Open problems – Attacker Model and Goals

Open problems – Attacker Model and Goals

Open problems – Attacker Model and Goals

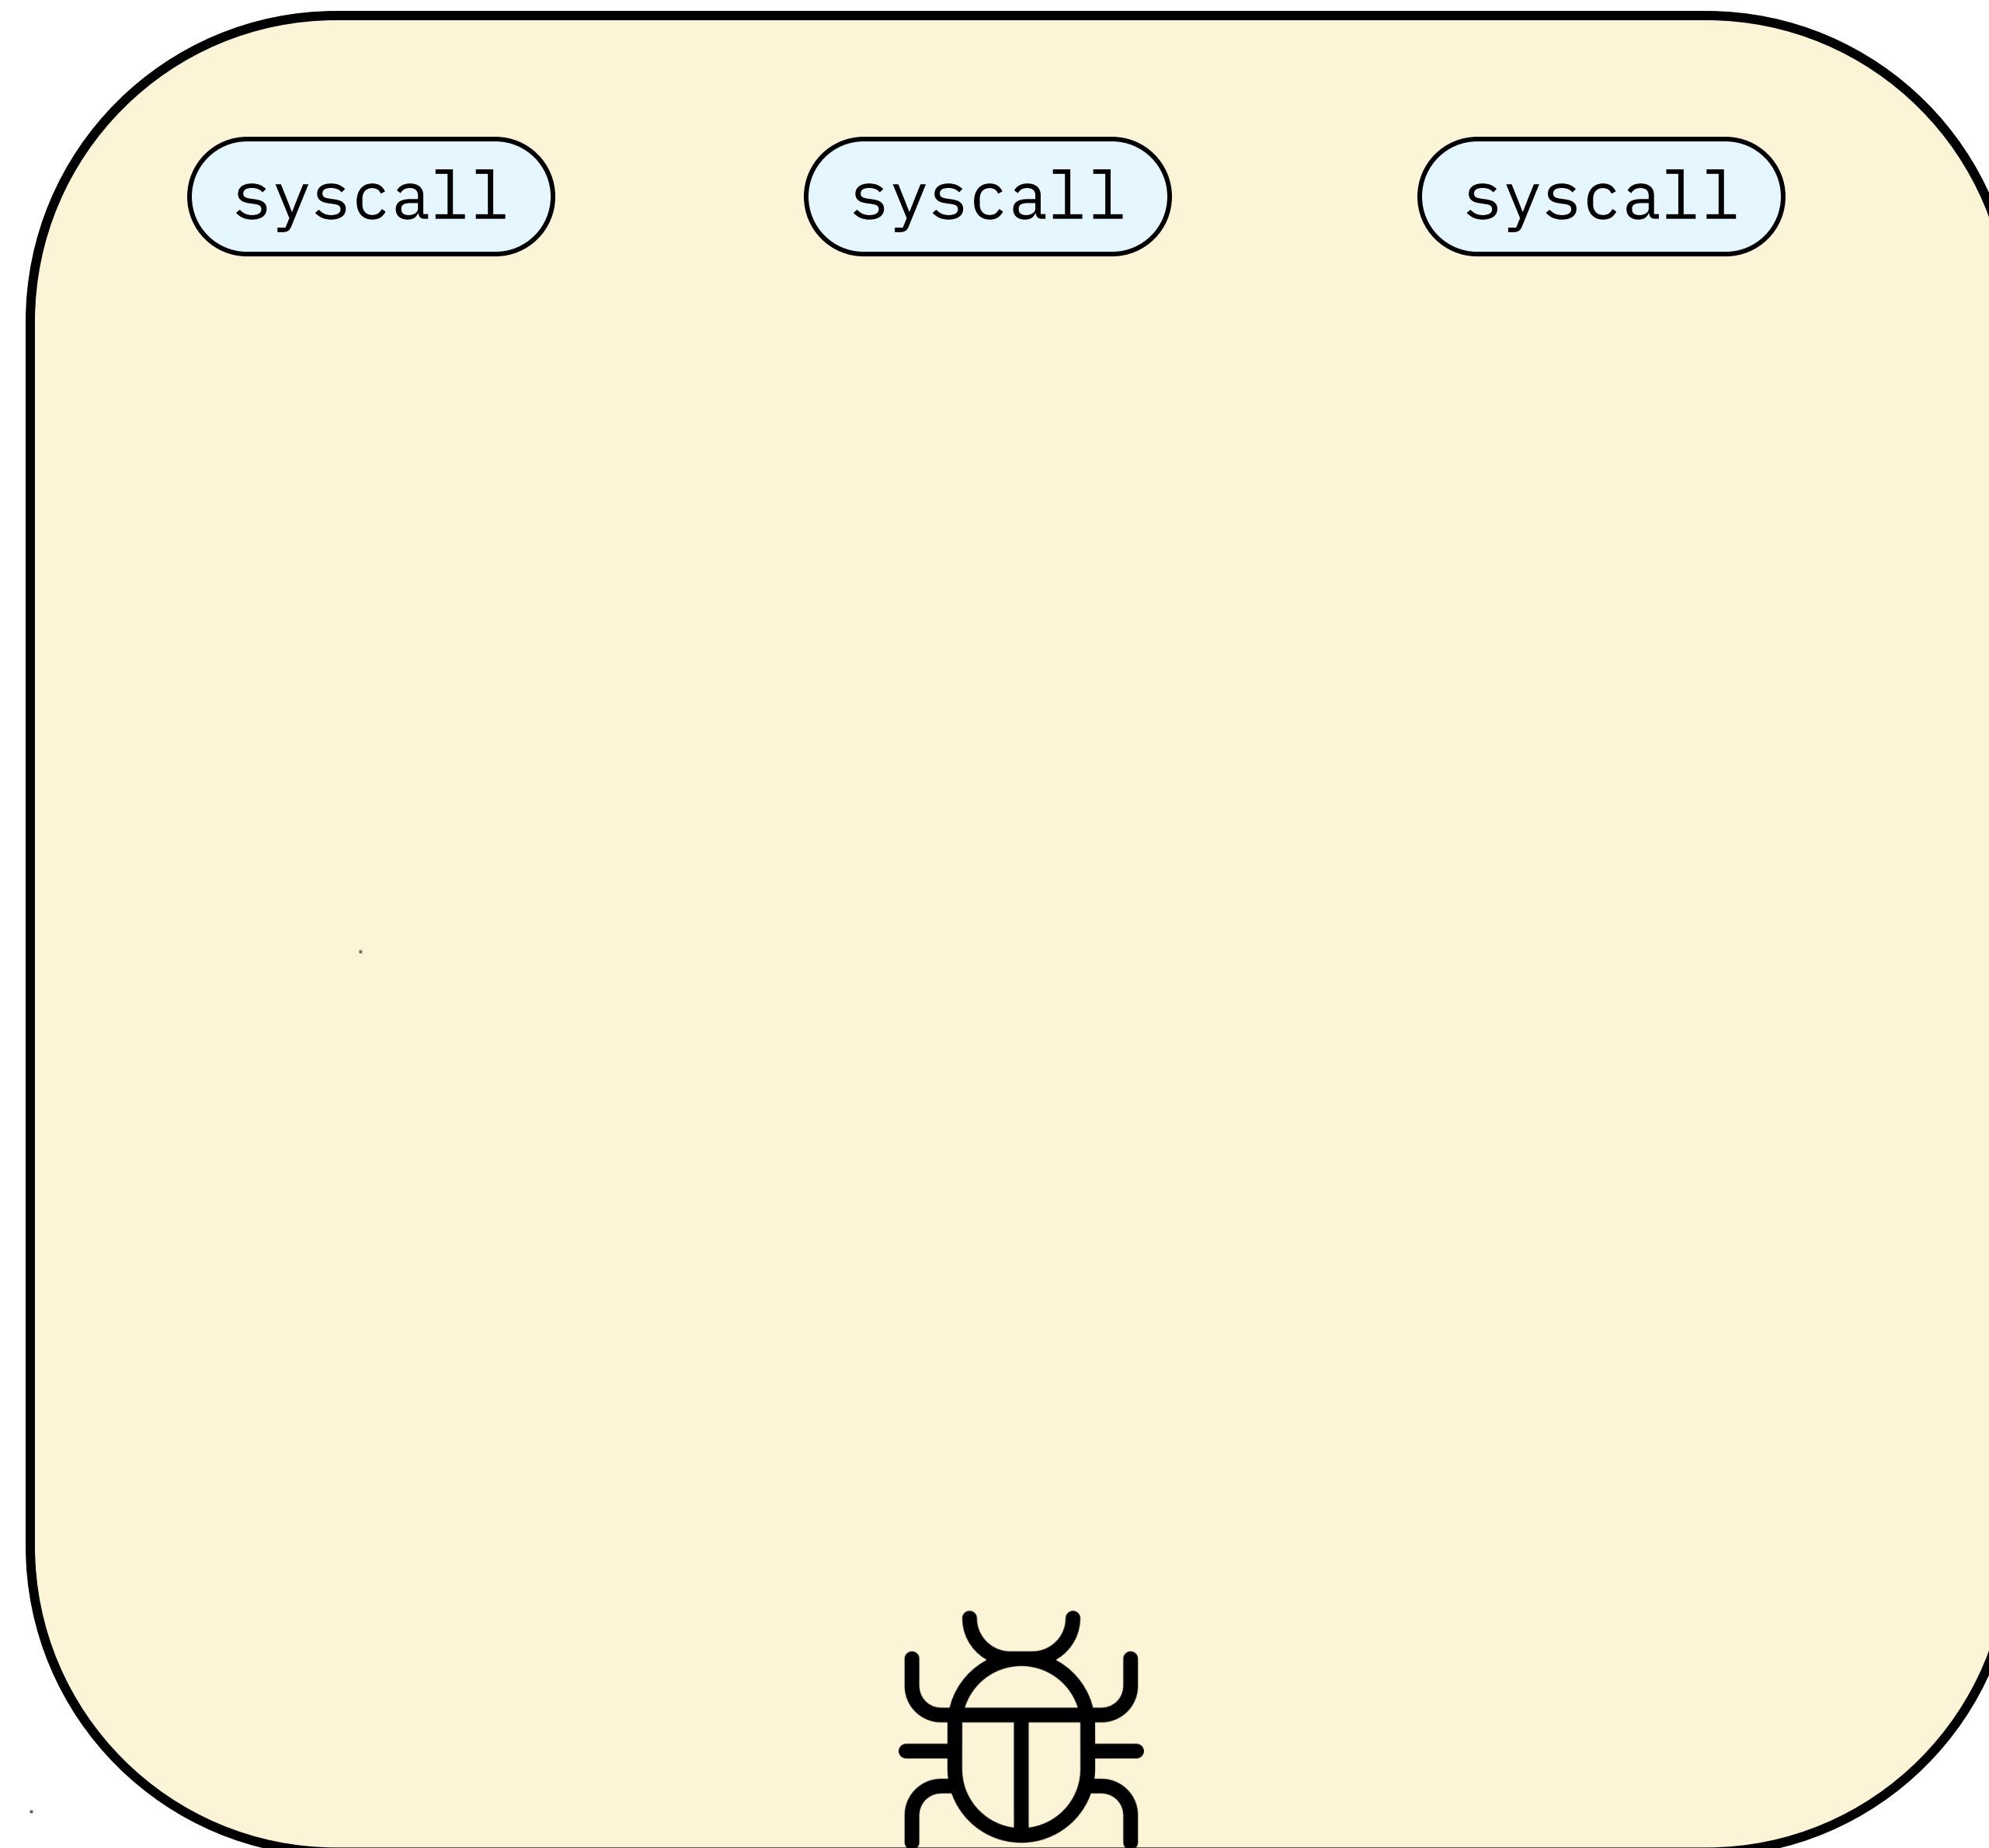
Most works are either **threat model agnostic** or focus **solely** to the **local attacker** scenario.

Open problems – Attacker Model and Goals


Most works are either **threat model agnostic** or focus **solely** to the **local attacker** scenario.

Attacker goals considered are only to PE or IL. **DoS** is tackled **by only one paper**

Open problems – Full Exploitability Assessment



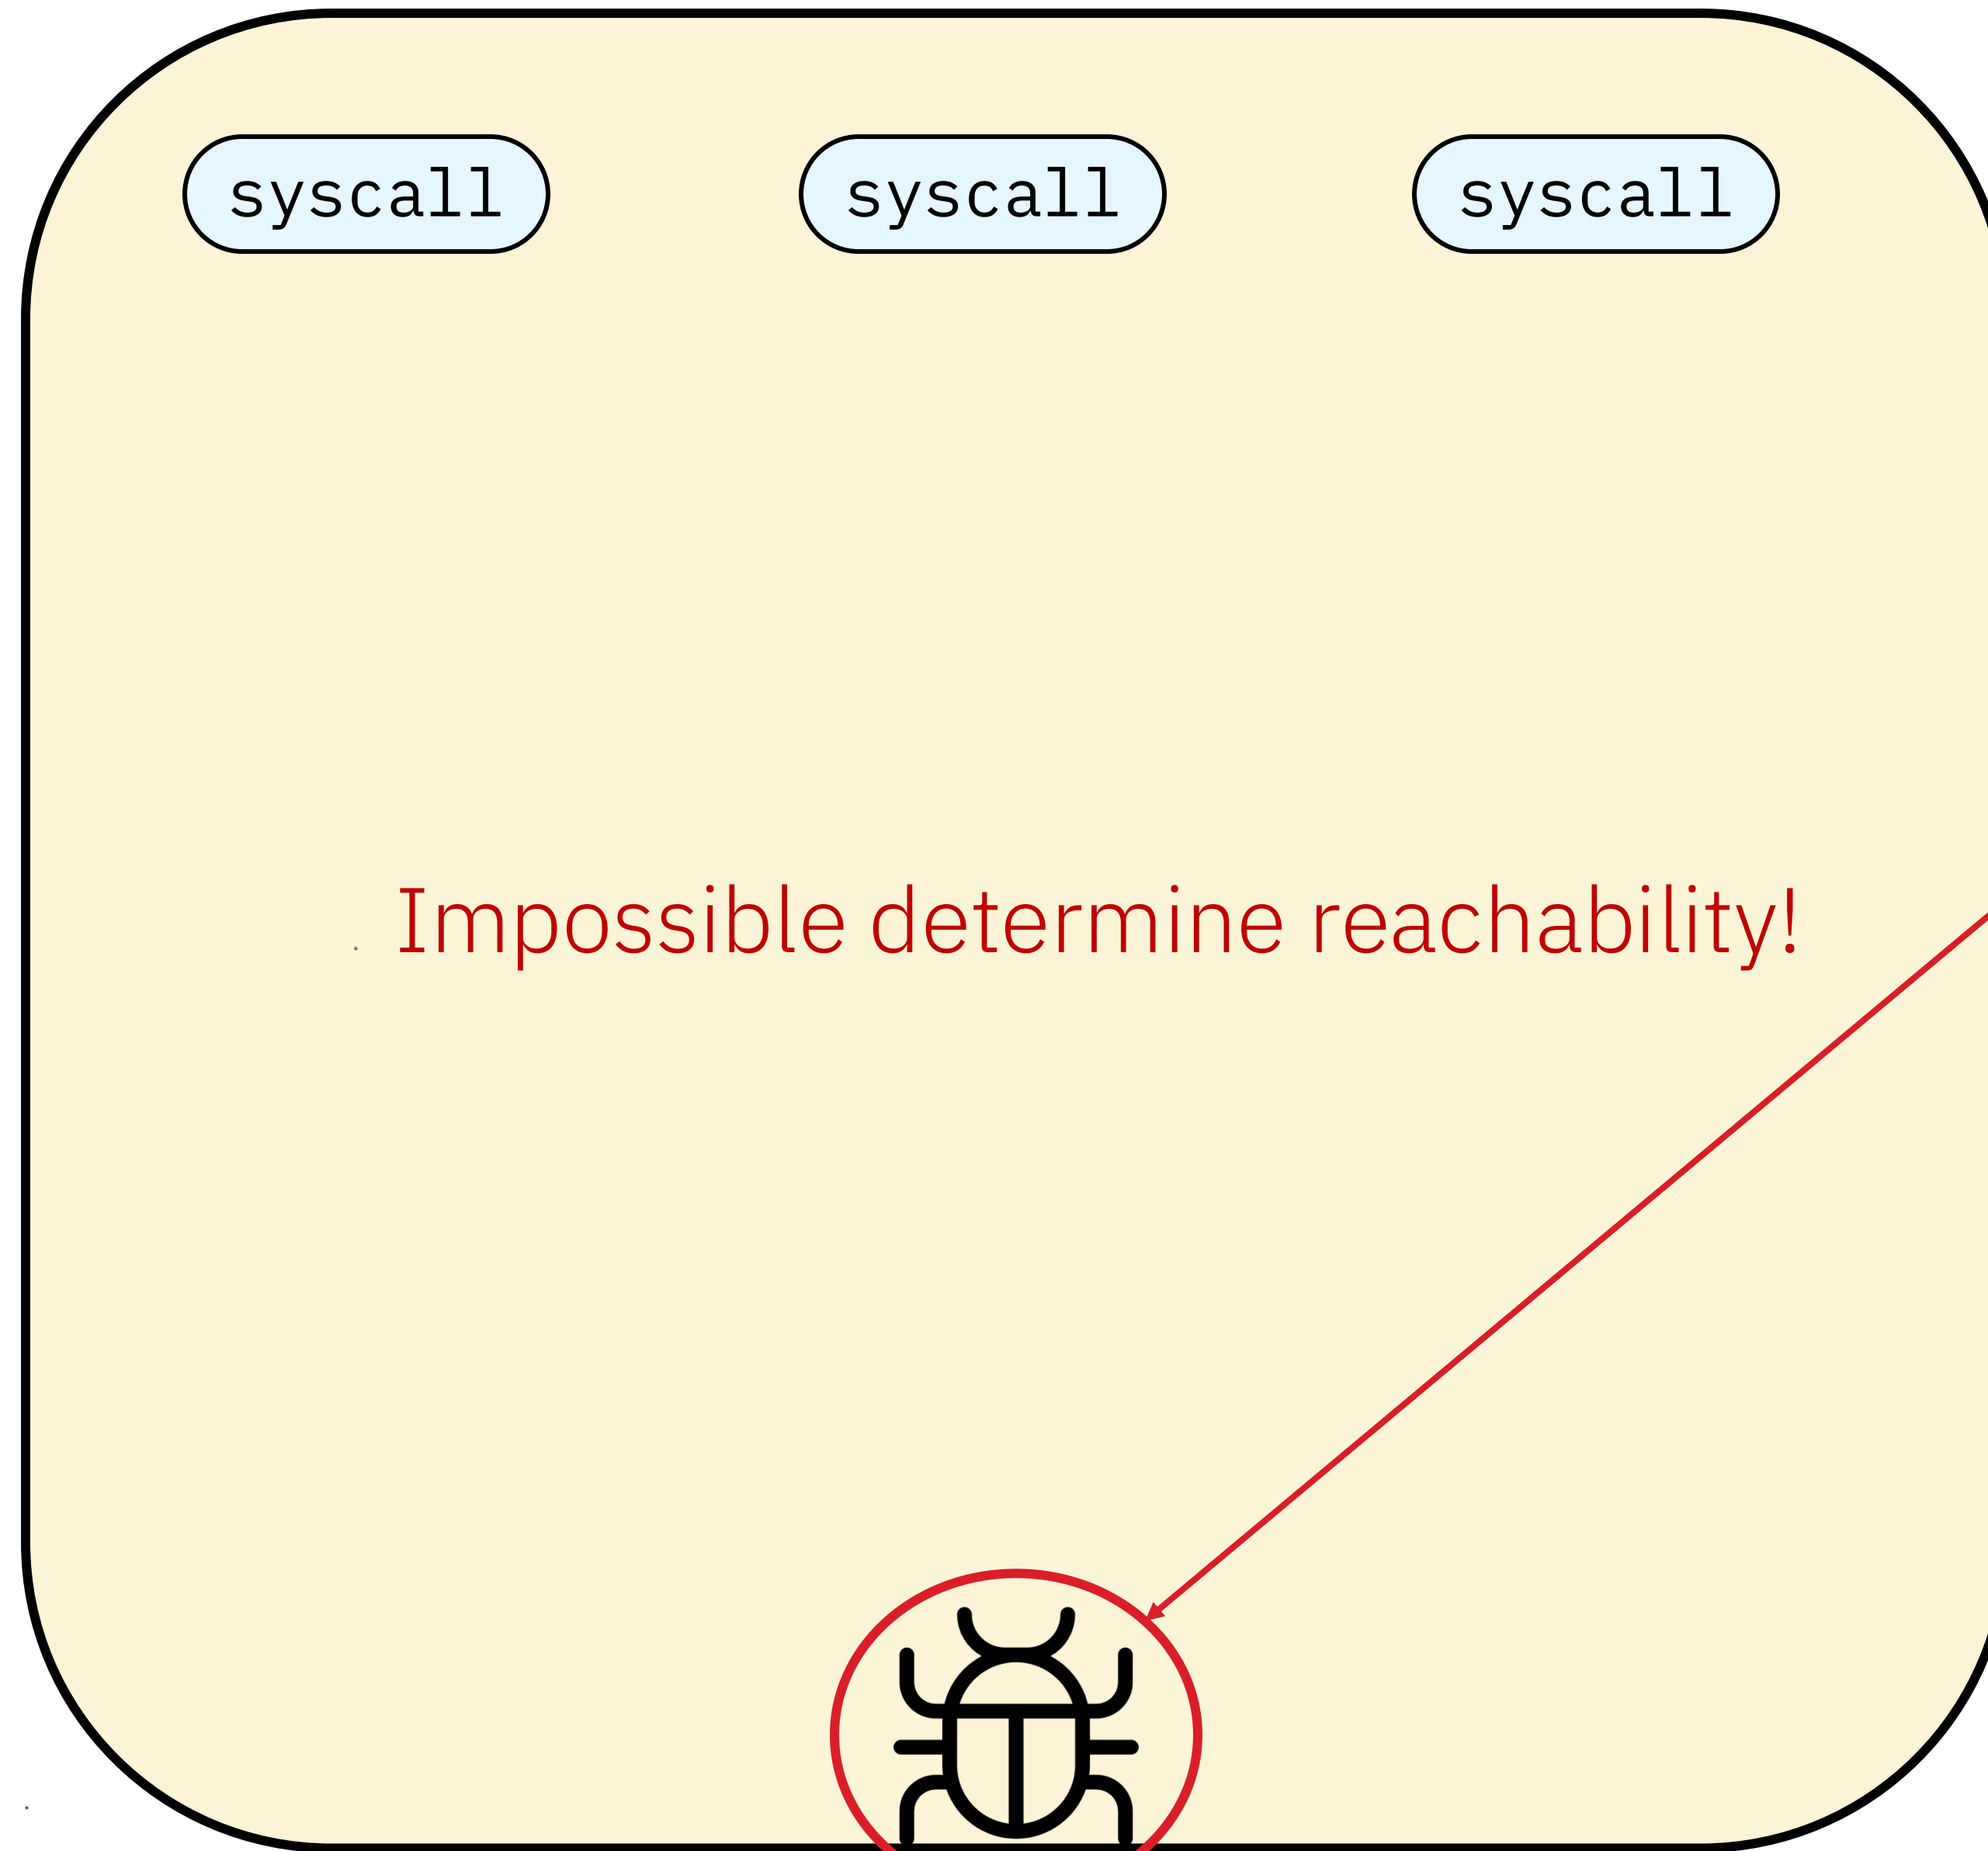
Open problems – Full Exploitability Assessment




static analysis



Open problems – Full Exploitability Assessment



Open problems – Full Exploitability Assessment



Open problems – Full Exploitability Assessment




dynamic analysis

syscall

syscall

syscall



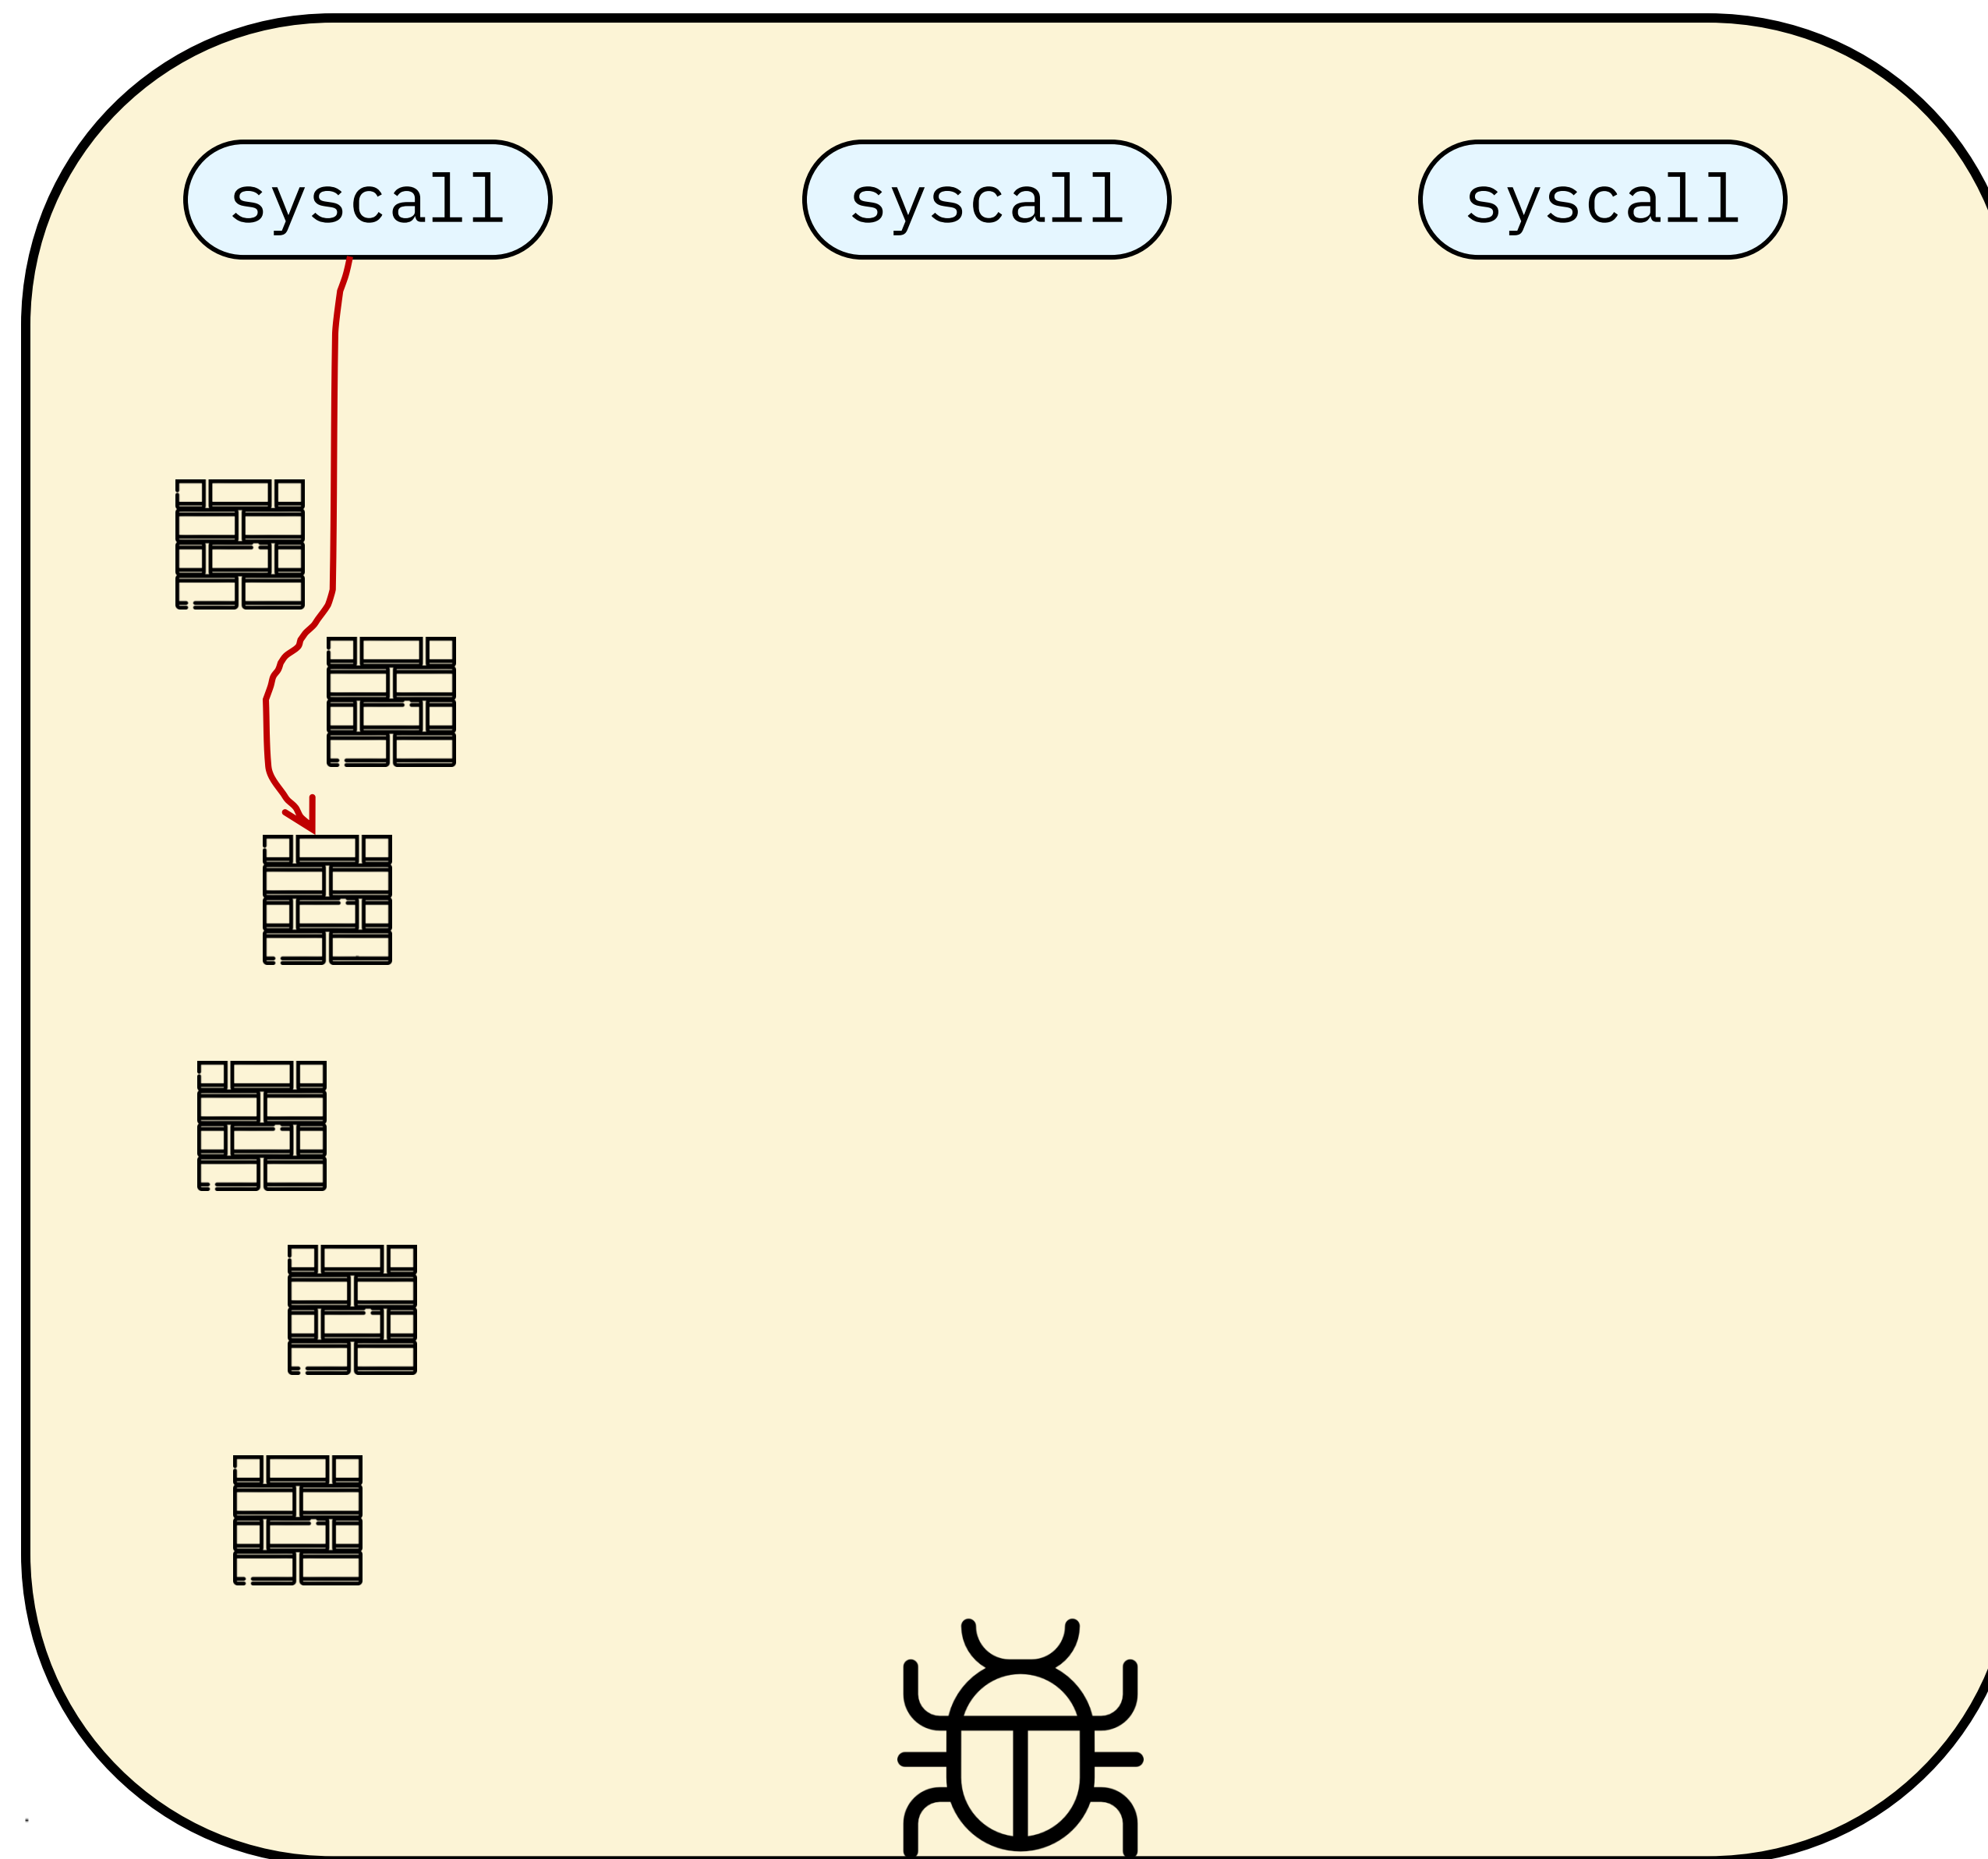
armasuisse



Open problems – Full Exploitability Assessment



dynamic analysis



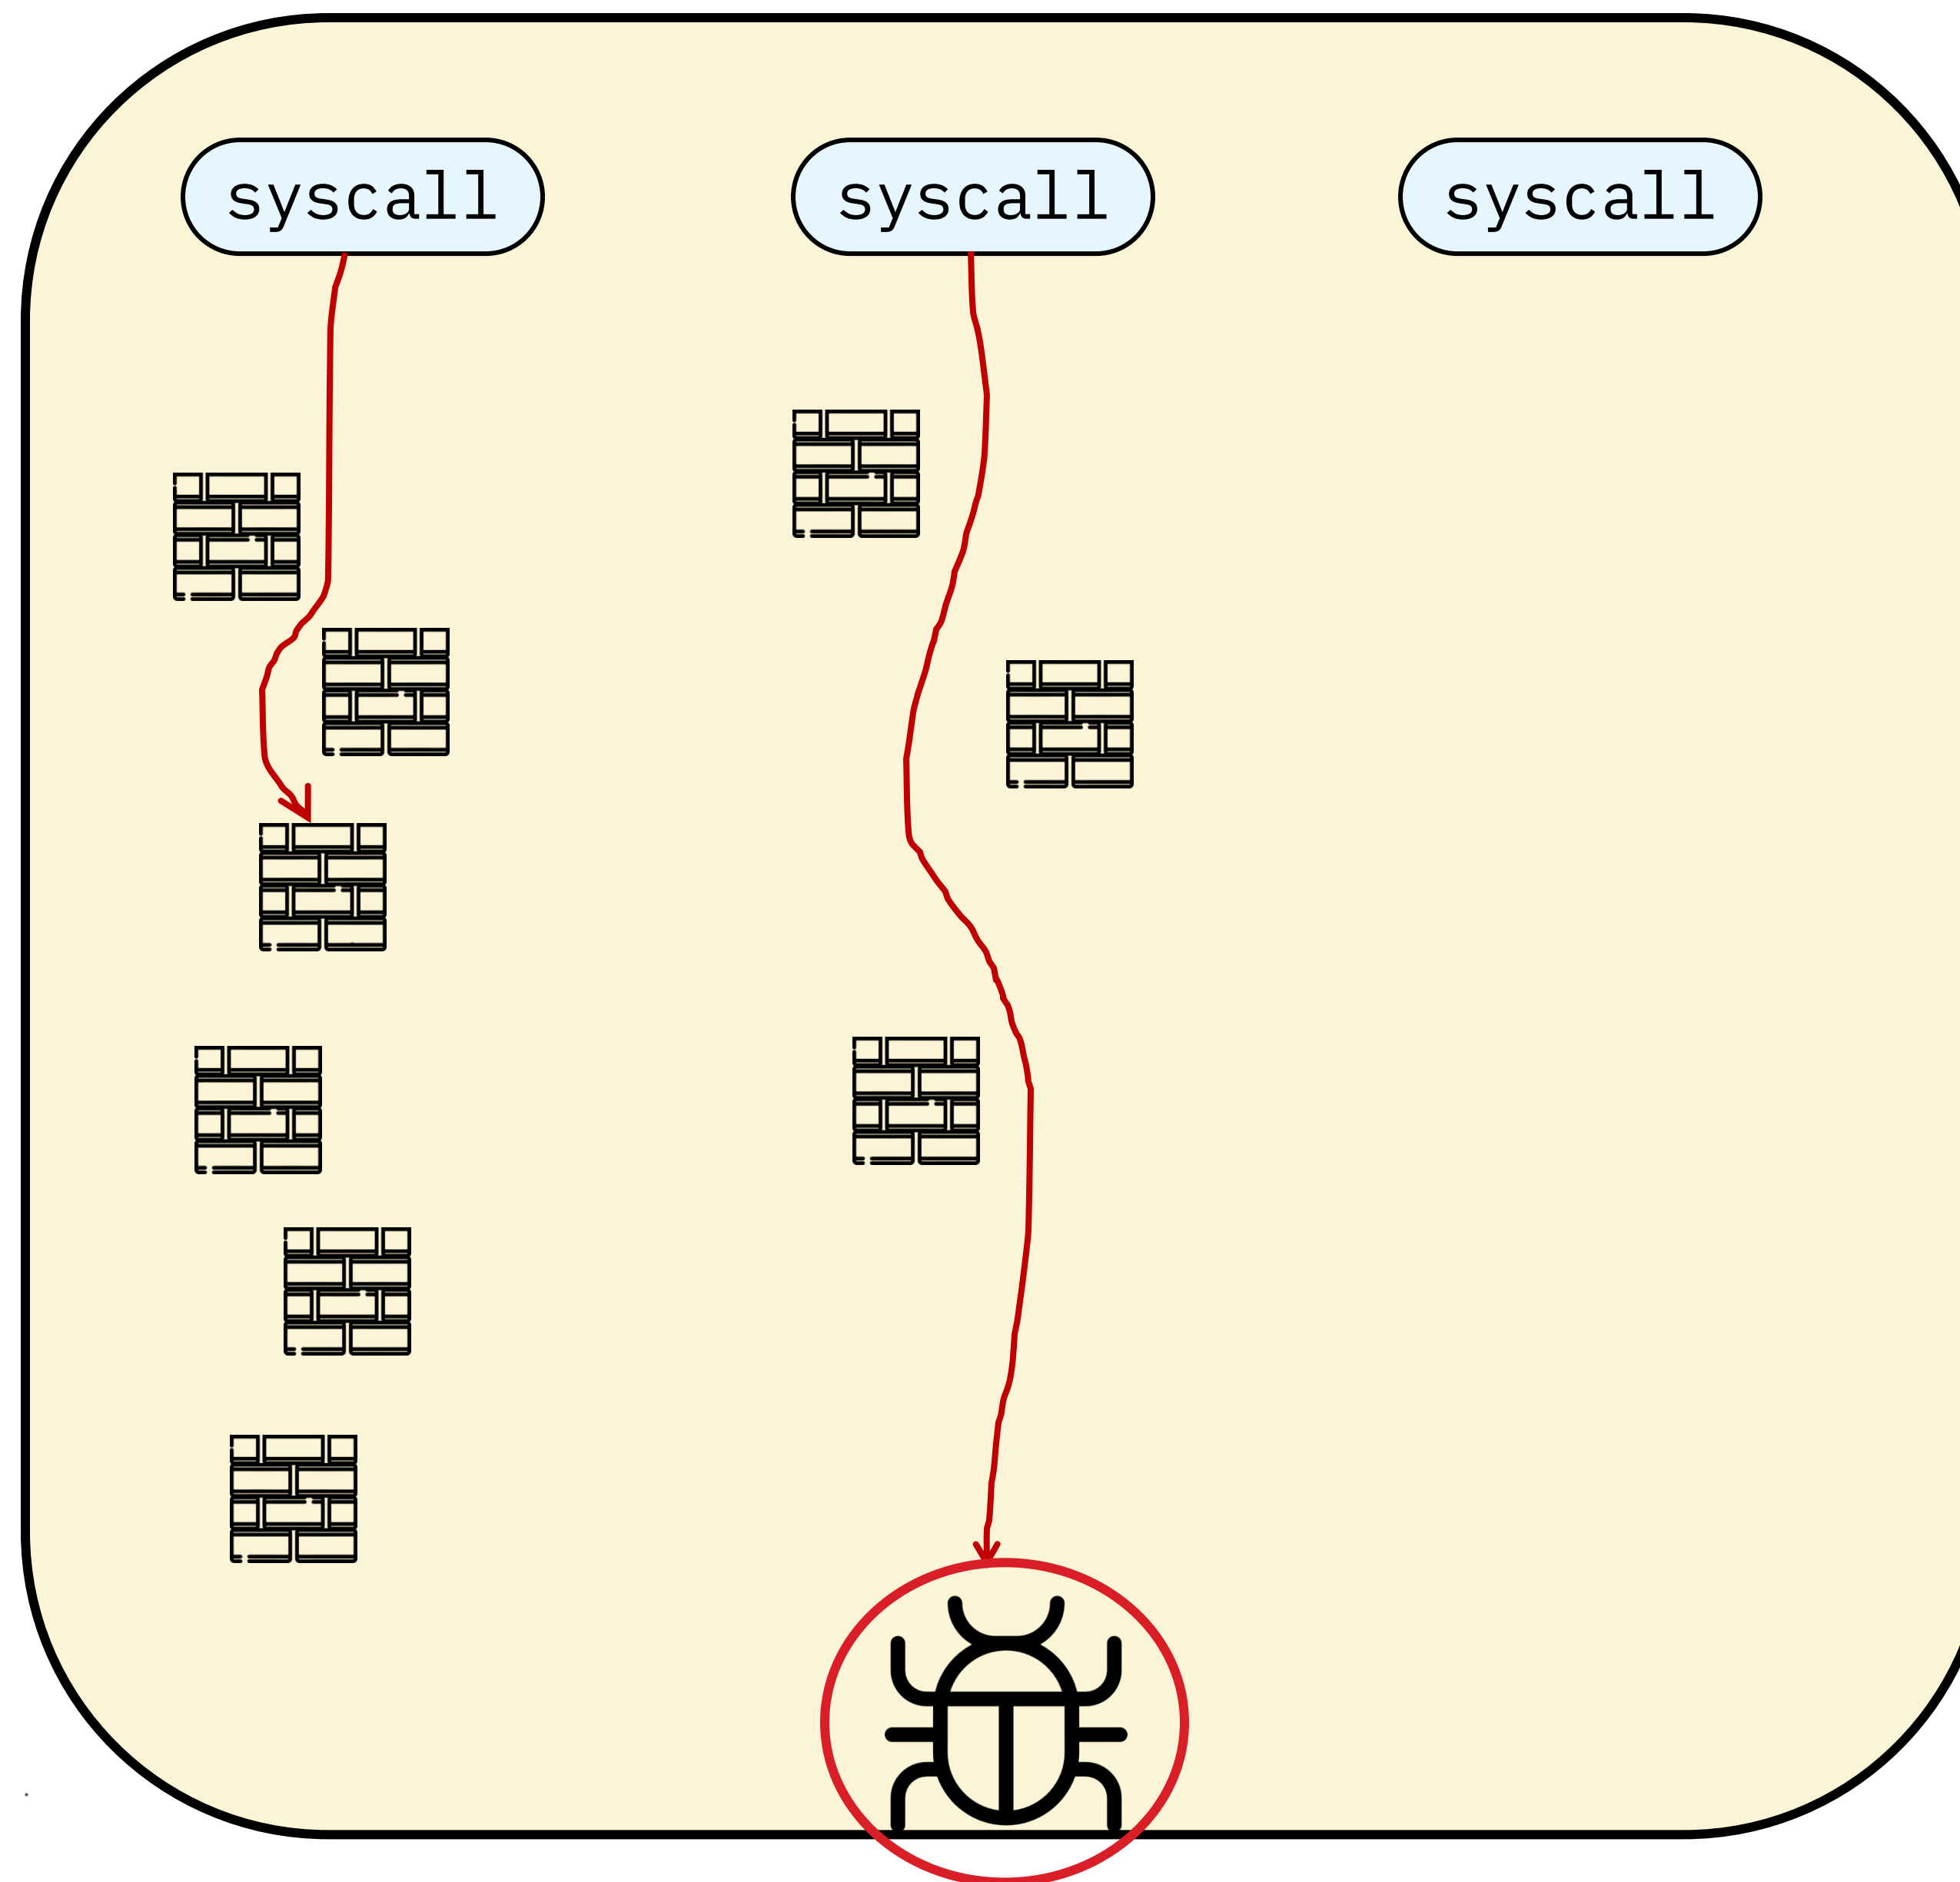
armasuisse




Open problems – Full Exploitability Assessment



dynamic analysis



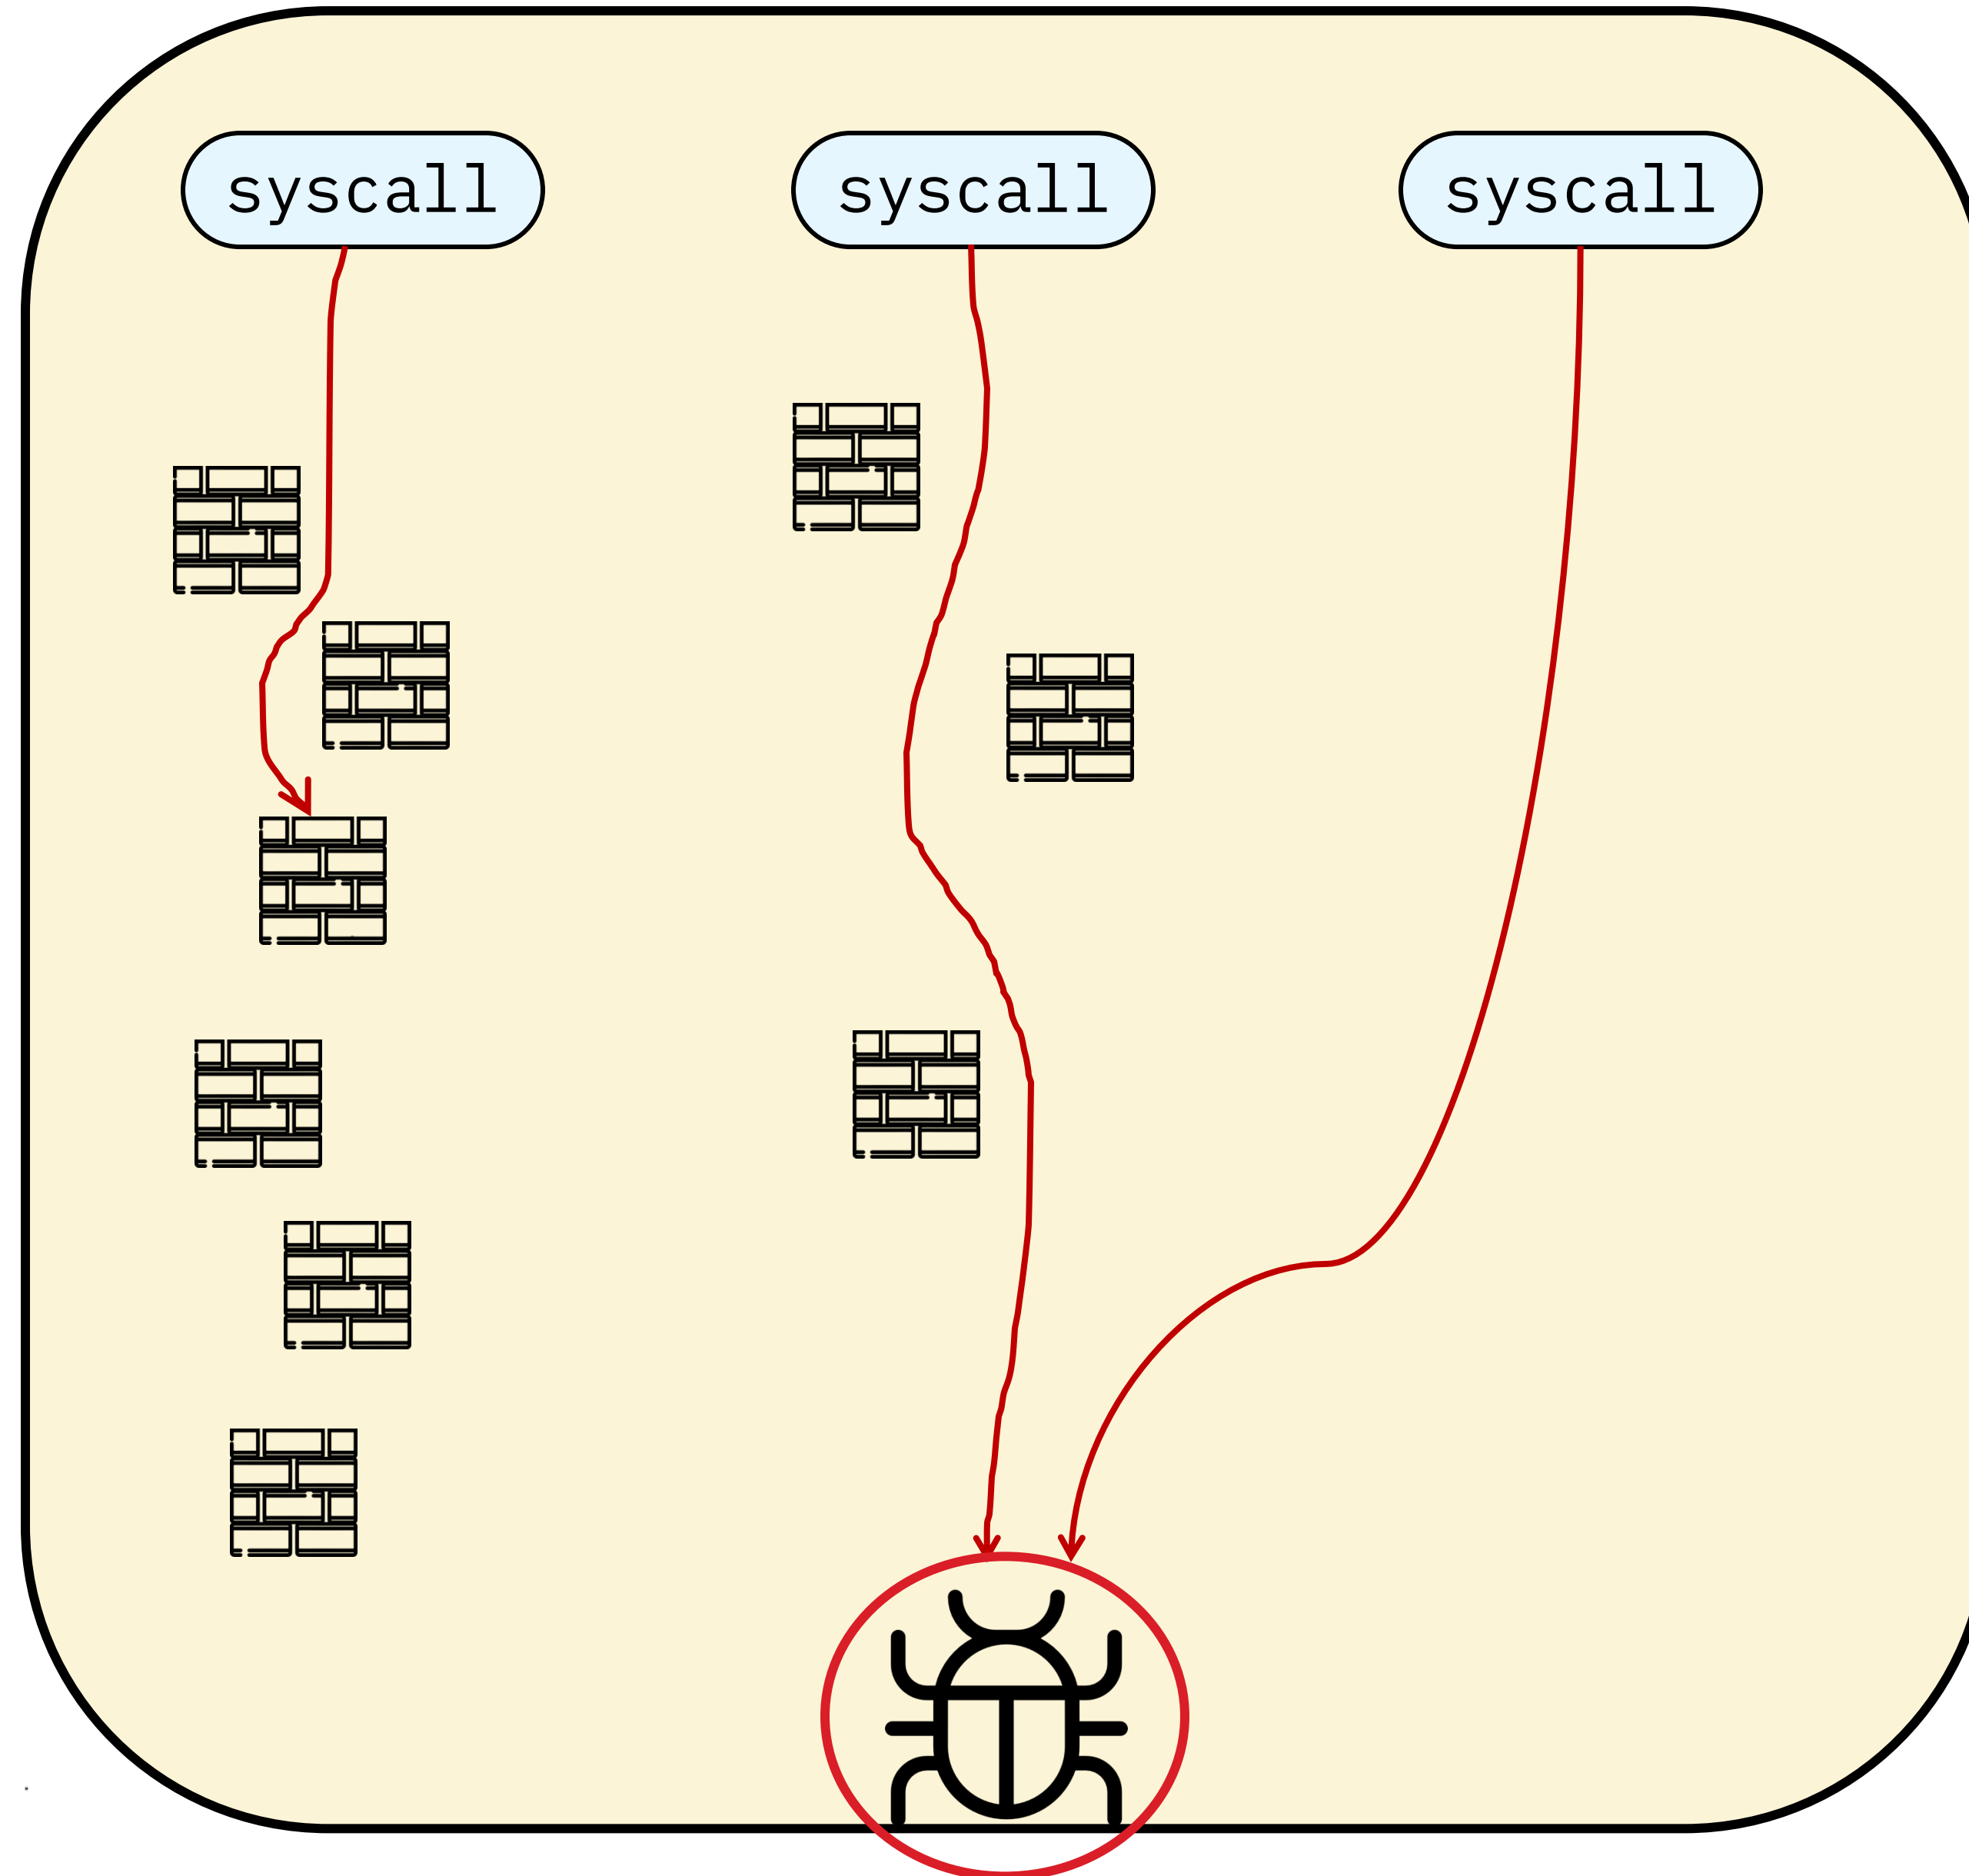
armasuisse



Open problems – Full Exploitability Assessment



dynamic analysis



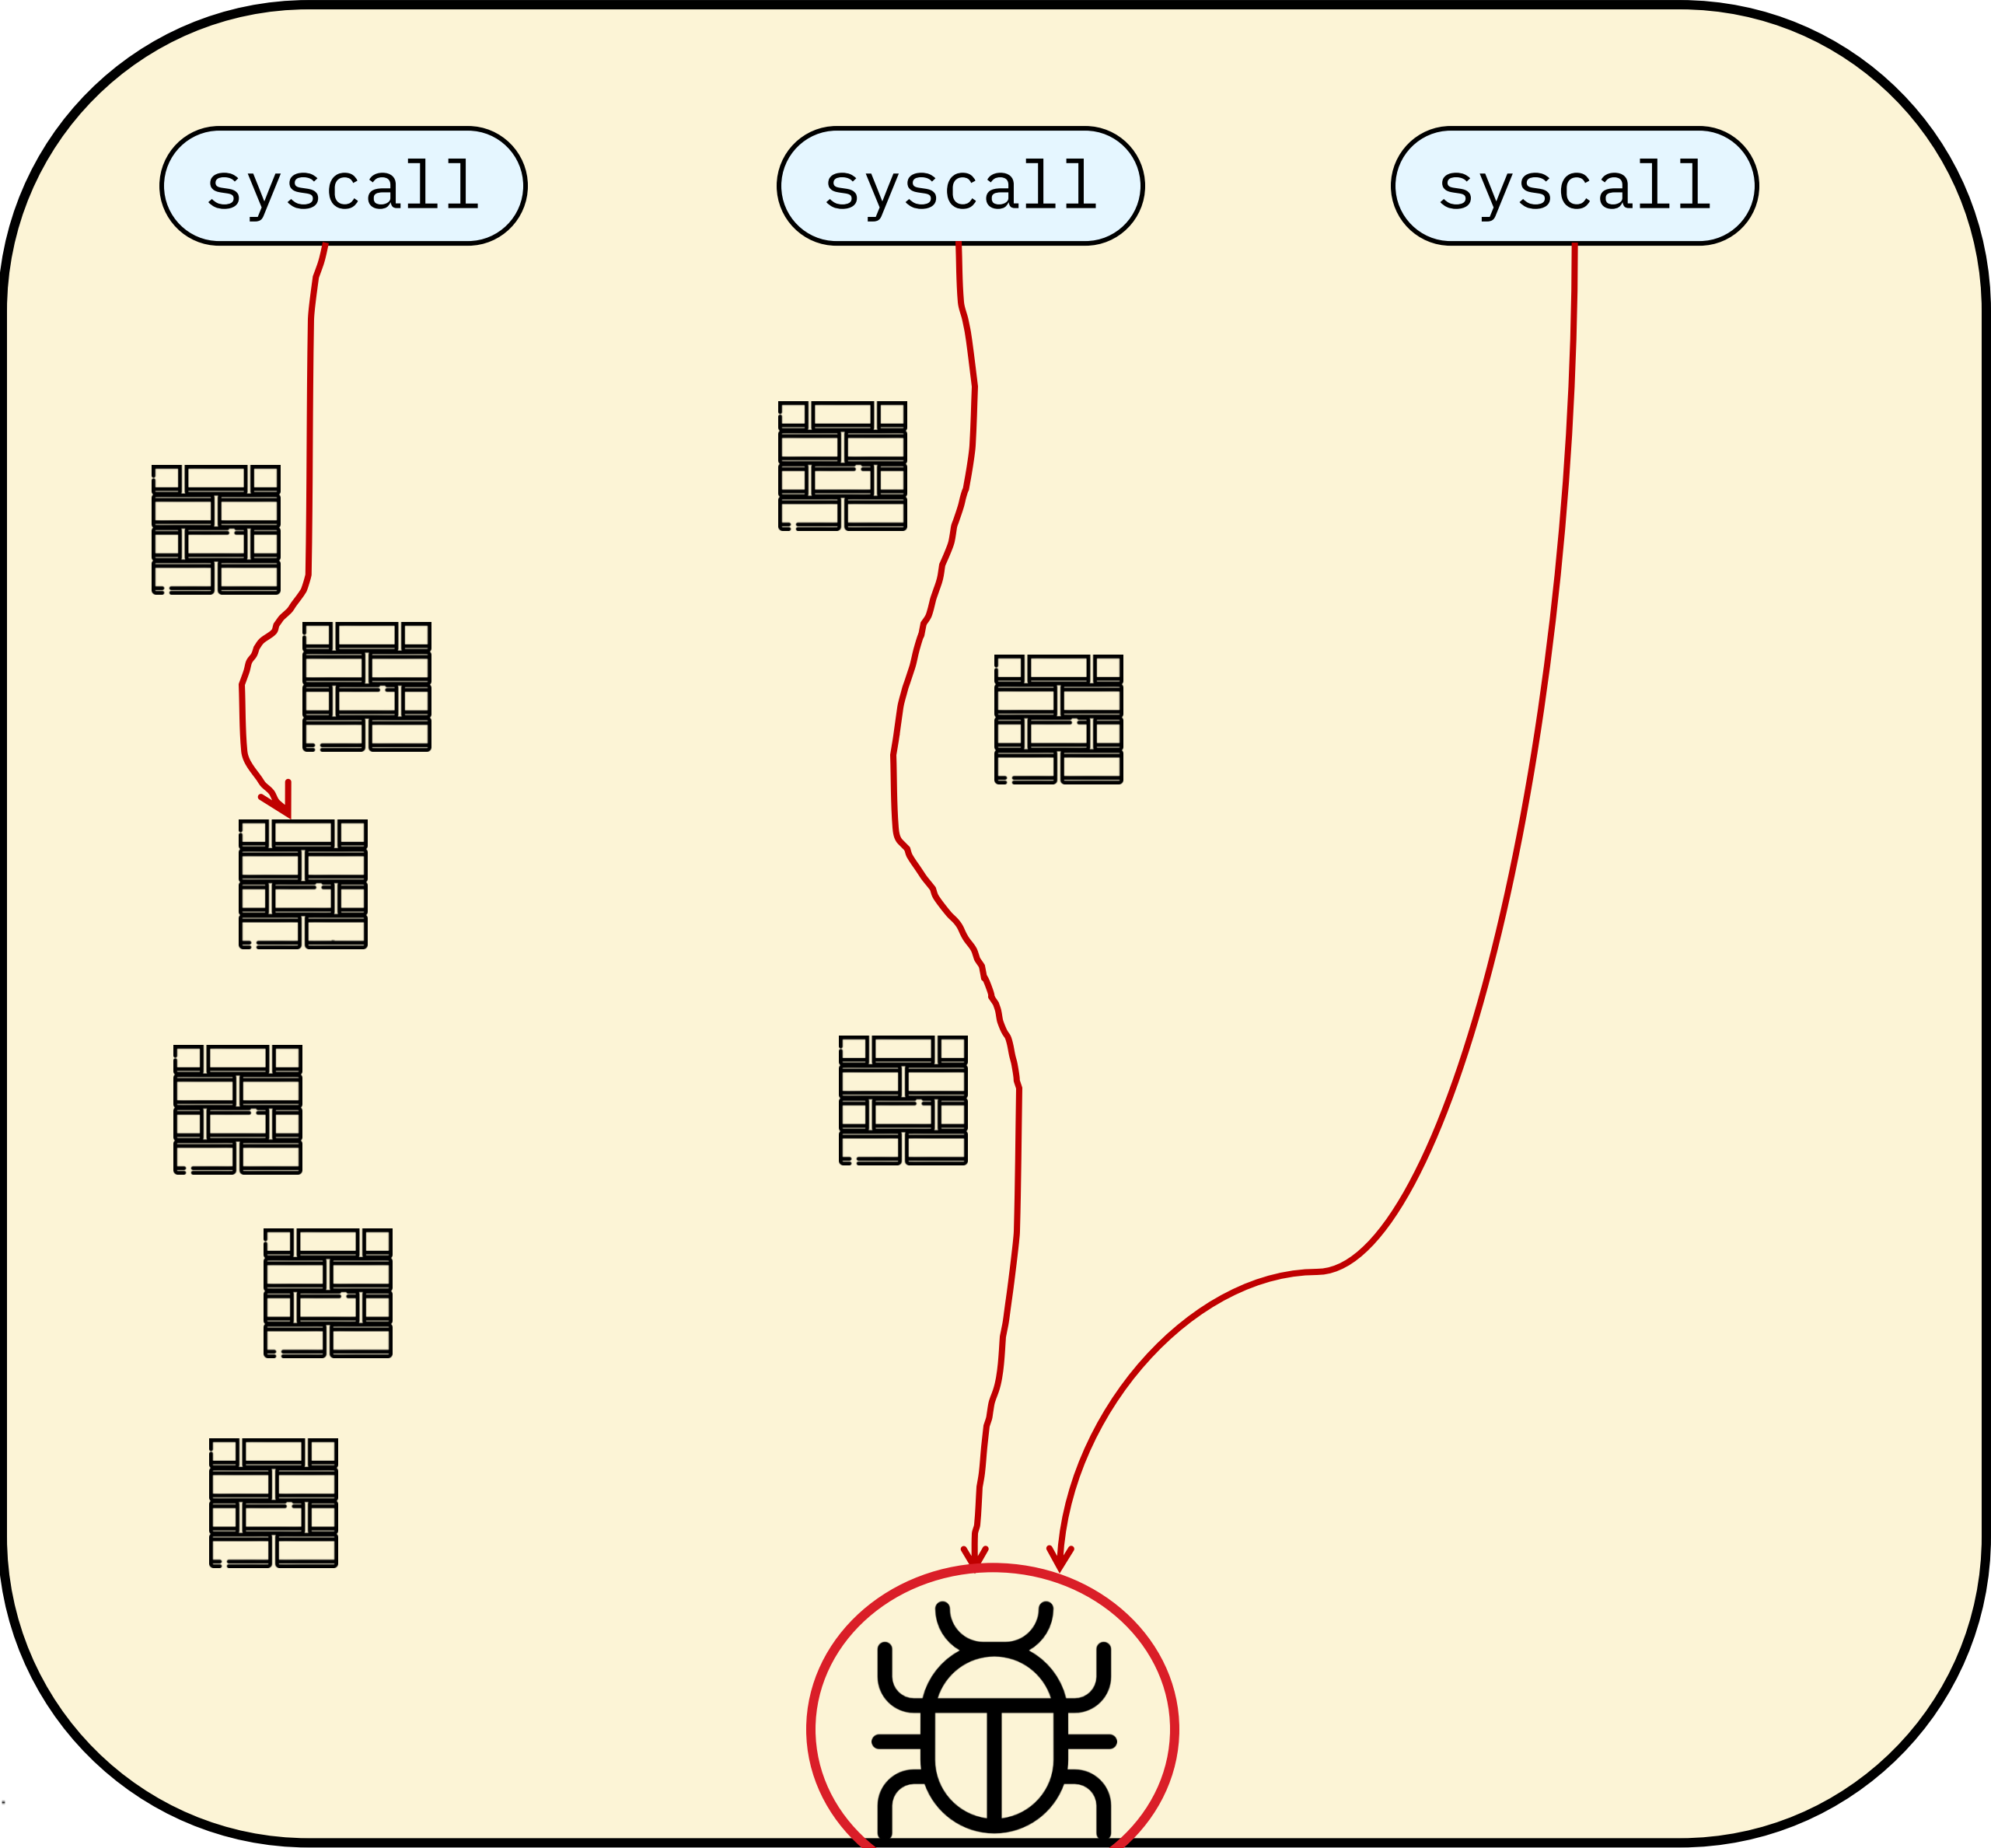
armasuisse




Open problems – Full Exploitability Assessment



dynamic analysis

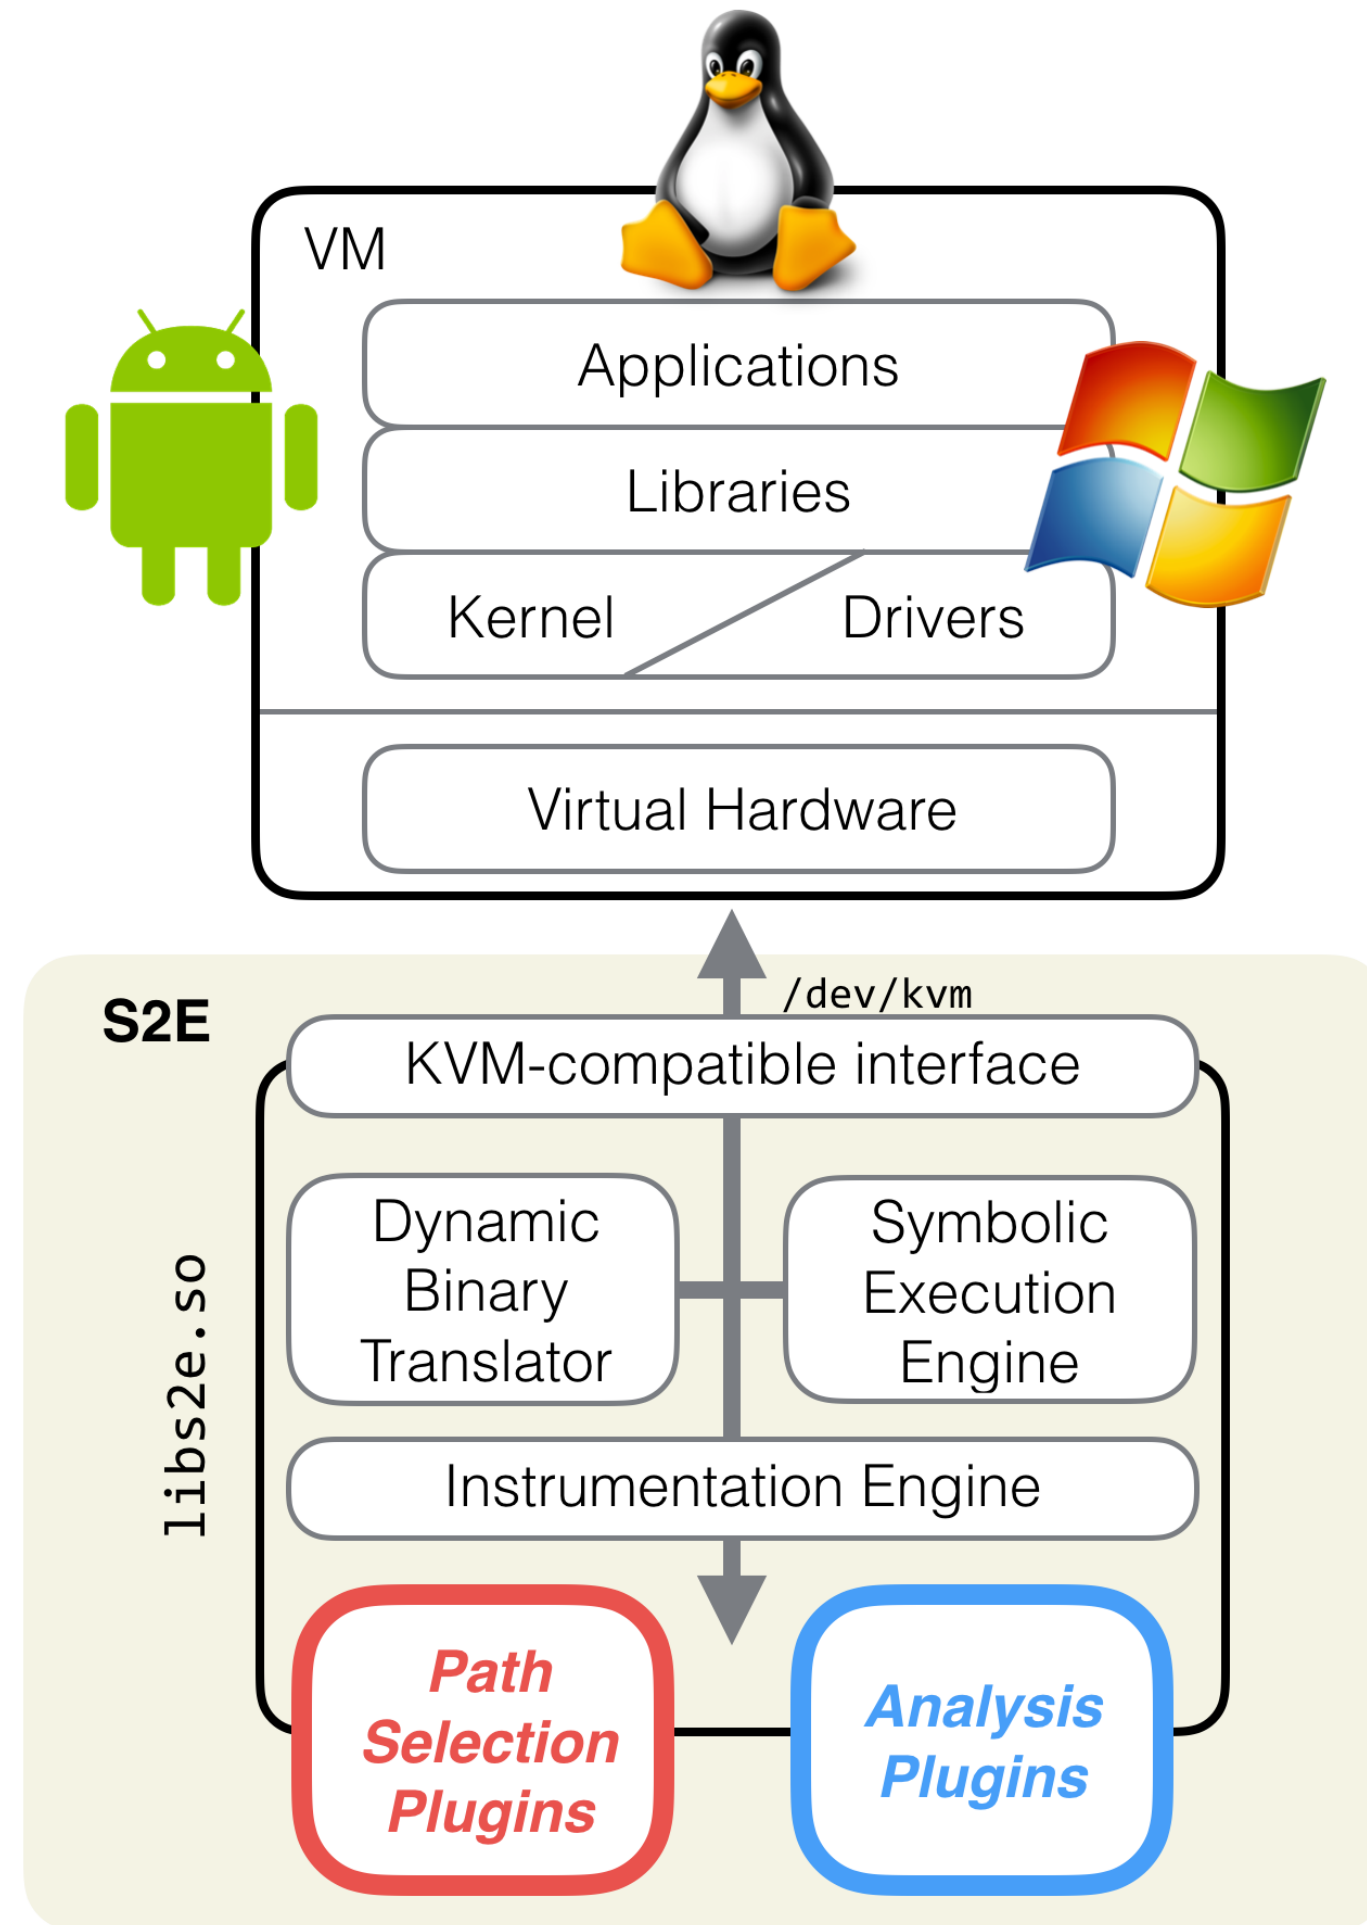


Only GREBE & SyzScope explore multiple/all paths to the same bug.



Open problems – Portability

Let's consider the widely used S²E as an example



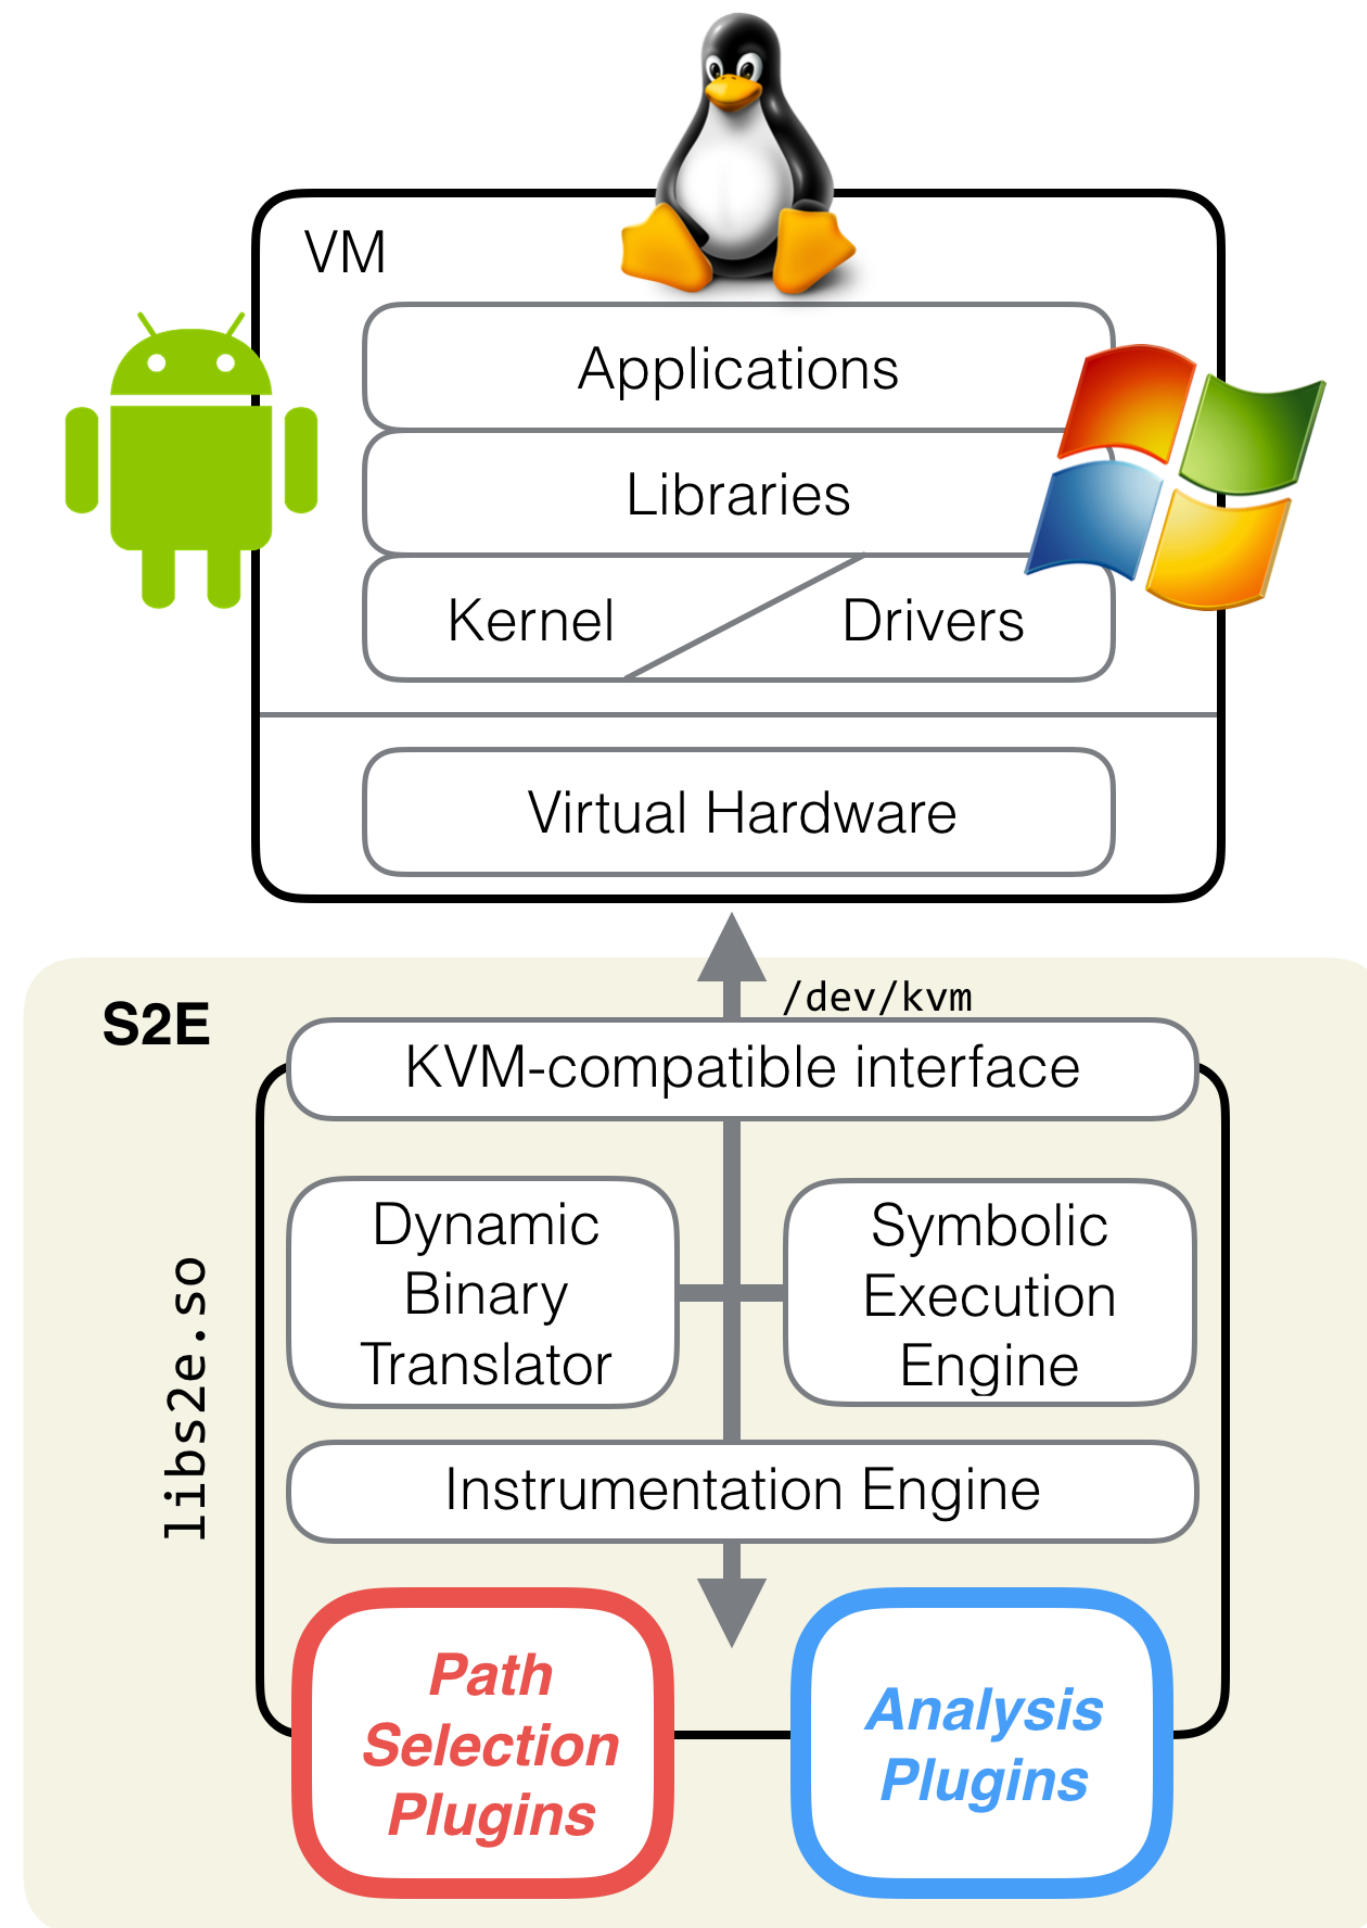
Depends on:

1. Qemu v3.0.0 (mainline is at v10)
2. LLVM lifter
3. Klee
4. Instrumented kernel (version 4.9.3 until last year)
5. Custom KVM implementation



Open problems – Portability

Let's consider the widely used S²E as an example



Depends on:

1. Qemu v3.0.0 (mainline is at v10)
2. LLVM lifter
3. Klee
4. Instrumented kernel (version 4.9.3 until last year)
5. Custom KVM implementation



Open problems – Reproducibility

Reproducing and comparing with previous work is hard and requires expert knowledge

Each paper select their favorite set of bugs and kernel version

The community could agree on one research Linux version and bug dataset

The known bugs could be ported to the supported kernel automatically (difficult!)

Open problems – Tool reuse

We found that Syzkaller, LLVM, S2E and angr are overall highly used - kudos!

However, **papers** using these tools **re-implemented** within them **similar features** which are never merged into the tool mainline

Often, **no code is provided**

Methodologies from different **papers** often **cannot be combined** due to **incompatible or brittle codebases**

Conclusions - Status Kernel AEG

No work has yet achieved full AEG for the kernel at scale.

Exploitation steps are often tackled in isolation.

Focus is given to only the most prominent threat model and attacker goals.

Very limited re-use of state of the art beyond major analysis framework.