



# | KernelSnitch

## Side-Channel Attacks on Kernel Data Structures

**Lukas Maar**   Jonas Juffinger   Thomas Steinbauer   Daniel Gruss   Stefan Mangard

Graz University of Technology, Austria

NDSS 2025

> [isec.tugraz.at](https://isec.tugraz.at)

## Novel OS side channel **KernelSnitch**



## Novel OS side channel **KernelSnitch**

- Timing side channel:

- ⚙ Different access timings of **kernel data structures**



## Novel OS side channel **KernelSnitch**

- Timing side channel:
  - ⚙ Different access timings of **kernel data structures**
- Amplification:
  - ⚙ Make timing difference **exploitable from user space**



## Novel OS side channel **KernelSnitch**

- Timing side channel:
  - ⚙ Different access timings of **kernel data structures**
- Amplification:
  - ⚙ Make timing difference **exploitable from user space**
- Attacks:
  - ⚙ Covert channel with up to  $580 \text{ kbit s}^{-1}$
  - ⚙ Website fingerprinting with  $F_1$  of 89.3 %
  - ⚙ Kernel heap pointer leak in under 65 s



## Novel OS side channel **KernelSnitch**

- Timing side channel:
  - ☞ Different access timings of **kernel data structures**
- Amplification:
  - ☞ Make timing difference **exploitable from user space**
- Attacks:
  - ☞ Covert channel with up to  $580 \text{ kbit s}^{-1}$
  - ☞ Website fingerprinting with  $F_1$  of 89.3 %
  - ☞ Kernel heap pointer leak in under 65 s
- Opensource our artifacts on:
  - 🔗 <https://github.com/IAIK/KernelSnitch>




# Motivation & Background

Hardware

Application



## Hardware



## Application

## Hardware



## Application

## Hardware




## Application



## Hardware




## Application






## Hardware



## Application






## Hardware



## Operating System



## Application






## Hardware



## Operating System



## Application





Empty list



List with  
2 elements



List with  
5 elements





Empty list



List with  
2 elements



List with  
5 elements



*Fast access*



Empty list



List with  
2 elements



List with  
5 elements



*Fast access*



*Medium fast access*



Empty list



List with  
2 elements



List with  
5 elements



*Fast access*




*Medium fast access*




*Slow access*


<sup>1</sup>This presentation only covers hash tables. For other data structures, see paper.



<sup>1</sup>This presentation only covers hash tables. For other data structures, see paper.

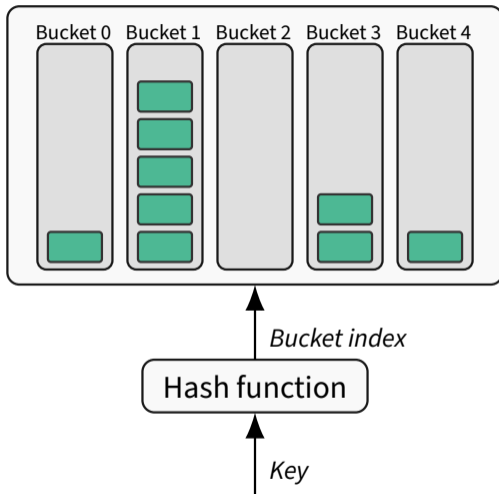


<sup>1</sup>This presentation only covers hash tables. For other data structures, see paper.



Hash function

<sup>1</sup>This presentation only covers hash tables. For other data structures, see paper.



<sup>1</sup>This presentation only covers hash tables. For other data structures, see paper.




**KernelSnitch**

User space


Process 1

Process 2


Kernel space



*Data  
container  
structure*




Process 1 → syscall accesses the data structure



Process 1 → syscall accesses the data structure

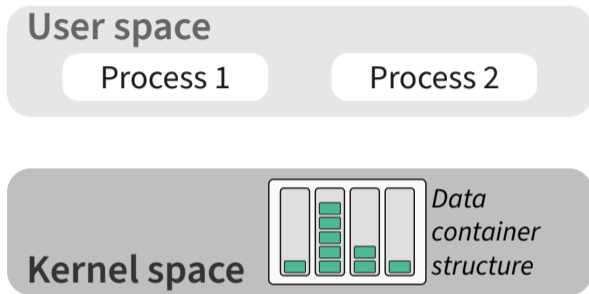
Process 2 → syscall appends data



Process 1 → syscall accesses the data structure

Process 2 → syscall appends data

Process 1 → syscall re-accesses the data structure



Process 1 → syscall accesses the data structure

Process 2 → syscall appends data

Process 1 → syscall re-accesses the data structure

**Deduce security-critical information**



**Access primitive**

**Append/remove primitive**



## Access primitive

- Syscalls that access structures

## Append/remove primitive

## Access primitive

- Syscalls that access structures
- Such as:

### Get POSIX time

```
1 sys_clock_gettime(id):  
2     sign = current.signal  
3     h = timer_hash(id, sign)  
4     hbucket =  
5         posix_timers_hhtable[h]  
6     for tim in hbucket:  
7         if tim.sign == sign and  
8             tim.id == id:  
9             return tim.get_time()  
10    return ERROR
```

## Append/remove primitive

## Access primitive

- Syscalls that access structures
- Such as:

### Get POSIX time

```
1 sys_clock_gettime(id):  
2     sign = current.signal  
3     h = timer_hash(id, sign)  
4     hbucket =  
5         posix_timers_hhtable[h]  
6     for tim in hbucket:  
7         if tim.sign == sign and  
8             tim.id == id:  
9             return tim.get_time()  
10    return ERROR
```

## Append/remove primitive

Hash

## Access primitive

- Syscalls that access structures
- Such as:

### Get POSIX time

```
1 sys_clock_gettime(id):  
2     sign = current.signal  
3     h = timer_hash(id, sign)  
4     hbucket =  
5         posix_timers_hhtable[h]  
6     for tim in hbucket:  
7         if tim.sign == sign and  
8             tim.id == id:  
9             return tim.get_time()  
10    return ERROR
```

Hash  
Bucket

## Append/remove primitive

## Access primitive

- Syscalls that access structures
- Such as:

### Get POSIX time

```
1 sys_clock_gettime(id):  
2     sign = current.signal  
3     h = timer_hash(id, sign)  
4     hbucket =  
5         posix_timers_hhtable[h]  
6     for tim in hbucket:  
7         if tim.sign == sign and  
8             tim.id == id:  
9             return tim.get_time()  
10    return ERROR
```

Hash  
Bucket  
Timing leak

## Append/remove primitive

## Access primitive

- Syscalls that access structures
- Such as:

### Get POSIX time

```
1 sys_clock_gettime(id):  
2     sign = current.signal  
3     h = timer_hash(id, sign)  
4     hbucket =  
5         posix_timers_hhtable[h]  
6     for tim in hbucket:  
7         if tim.sign == sign and  
8             tim.id == id:  
9             return tim.get_time()  
10    return ERROR
```

Hash  
Bucket  
Timing leak

## Append/remove primitive

- Syscalls that modify structures

## Access primitive

- Syscalls that access structures
- Such as:

### Get POSIX time

```
1 sys_clock_gettime(id):  
2     sign = current.signal  
3     h = timer_hash(id, sign)  
4     hbucket =  
5         posix_timers_hhtable[h]  
6     for tim in hbucket:  
7         if tim.sign == sign and  
8             tim.id == id:  
9             return tim.get_time()  
10    return ERROR
```

Hash  
Bucket  
Timing leak

## Append/remove primitive

- Syscalls that modify structures
- Such as:

### Create new POSIX timer

```
1 sys_timer_create():  
2     sign = current.signal  
3     id = sign.next_id  
4     h = timer_hash(id, sign)  
5     hbucket =  
6         posix_timers_hhtable[h]  
7     tim = k_itimer(sign, id)  
8     hbucket.append(tim)  
9     return id
```

## Access primitive

- Syscalls that access structures
- Such as:

### Get POSIX time

```
1 sys_clock_gettime(id):  
2     sign = current.signal  
3     h = timer_hash(id, sign)  
4     hbucket =  
5         posix_timers_hhtable[h]  
6     for tim in hbucket:  
7         if tim.sign == sign and  
8             tim.id == id:  
9             return tim.get_time()  
10    return ERROR
```

Hash  
Bucket  
Timing leak

## Append/remove primitive

- Syscalls that modify structures
- Such as:

### Create new POSIX timer

```
1 sys_timer_create():  
2     sign = current.signal  
3     id = sign.next_id  
4     h = timer_hash(id, sign)  
5     hbucket =  
6         posix_timers_hhtable[h]  
7     tim = k_itimer(sign, id)  
8     hbucket.append(tim)  
9     return id
```



## Access primitive

- Syscalls that access structures
- Such as:

### Get POSIX time

```
1 sys_clock_gettime(id):  
2     sign = current.signal  
3     h = timer_hash(id, sign)  
4     hbucket =  
5         posix_timers_hhtable[h]  
6     for tim in hbucket:  
7         if tim.sign == sign and  
8             tim.id == id:  
9             return tim.get_time()  
10    return ERROR
```

Hash  
Bucket  
Timing leak

## Append/remove primitive

- Syscalls that modify structures
- Such as:

### Create new POSIX timer

```
1 sys_timer_create():  
2     sign = current.signal  
3     id = sign.next_id  
4     h = timer_hash(id, sign)  
5     hbucket =  
6         posix_timers_hhtable[h]  
7     tim = k_itimer(sign, id)  
8     hbucket.append(tim)  
9     return id
```

## Access primitive

- Syscalls that access structures
- Such as:

### Get POSIX time

```
1 sys_clock_gettime(id):  
2     sign = current.signal  
3     h = timer_hash(id, sign)  
4     hbucket =  
5         posix_timers_hhtable[h]  
6     for tim in hbucket:  
7         if tim.sign == sign and  
8             tim.id == id:  
9             return tim.get_time()  
10    return ERROR
```

Hash  
Bucket  
Timing leak  
Append

## Append/remove primitive

- Syscalls that modify structures
- Such as:

### Create new POSIX timer

```
1 sys_timer_create():  
2     sign = current.signal  
3     id = sign.next_id  
4     h = timer_hash(id, sign)  
5     hbucket =  
6         posix_timers_hhtable[h]  
7     tim = k_itimer(sign, id)  
8     hbucket.append(tim)  
9     return id
```

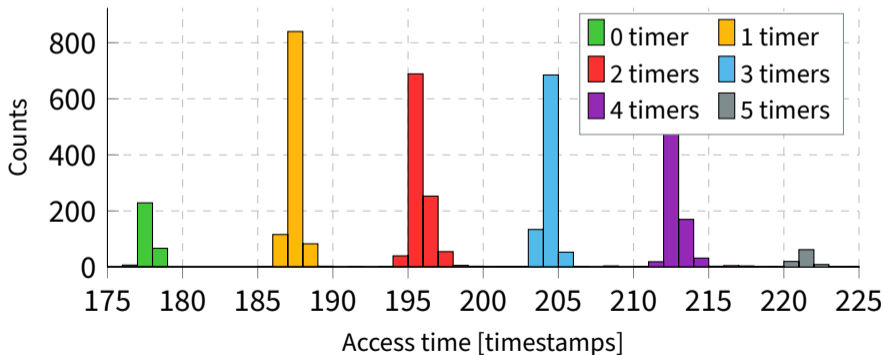
Access randomly selected buckets<sup>2</sup>

- Repeat 4096 accesses
- Buckets between 0 and 5 timers

<sup>2</sup> Applied information leakage amplification, see paper.


Access randomly selected buckets<sup>2</sup>

- Repeat 4096 accesses
- Buckets between 0 and 5 timers



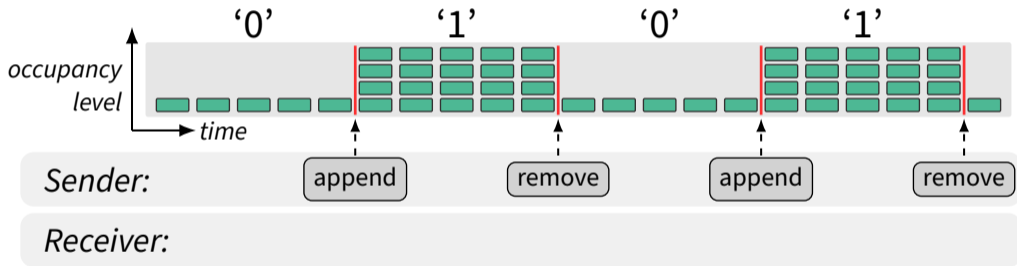
<sup>2</sup> Applied information leakage amplification, see paper.

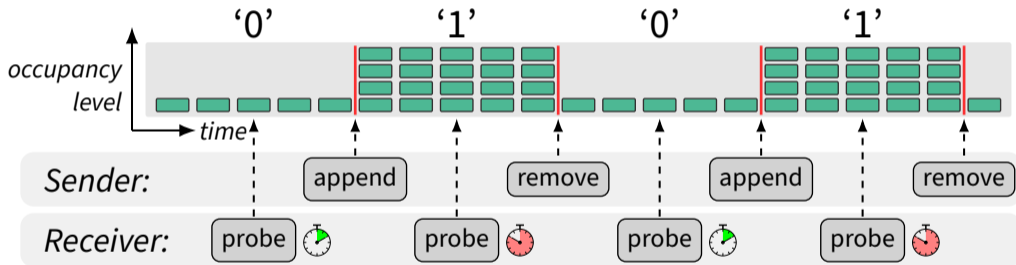
# Attacks



*Sender:*

*Receiver:*





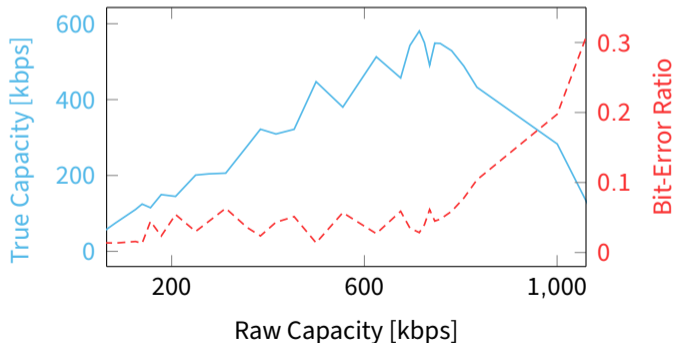


Vary time slice

- Determine raw capacity and BER
- Calculate true capacity

## Vary time slice

- Determine raw capacity and BER
- Calculate true capacity
- True capacity up to  $580 \text{ kbit s}^{-1}$






*Browser:*

*Spy:*




*Browser:* Appends/removes when accessing a website ...

*Spy:*



*Browser:* Appends/removes when accessing a website for all buckets.


*Spy:*



*Browser:* Appends/removes when accessing a website for all buckets.

*Spy:* Probes all buckets.

# Confusion Matrix of Website Fingerprinting



Top 100 websites

- Use Convolutional Neural Network (CNN)
- $F_1$  score of 89.3%

Recall:

## Access primitive

```
1 index = hash_fn(kaddr, user_id)
2 bucket = hash_table[index]
3 for e in bucket
```



## Hash:

- Known: hash\_fn, user\_id
- Unkown: `kaddr`

## Recall:

### Access primitive

```
1 index = hash_fn(kaddr, user_id)
2 bucket = hash_table[index]
3 for e in bucket
```

Recall:

## Access primitive

```
1 index = hash_fn(kaddr, user_id)
2 bucket = hash_table[index]
3 for e in bucket
```

Hash:

- Known: hash\_fn, user\_id
- Unkown: kaddr

Exploit strategy (online phase):

- Same kaddr, but different user\_ids
- Use our side channel to detect hash collisions

Recall:

## Access primitive

```
1 index = hash_fn(kaddr, user_id)
2 bucket = hash_table[index]
3 for e in bucket
```

Hash:

- Known: hash\_fn, user\_id
- Unknown: kaddr

Exploit strategy (online phase):

- Same kaddr, but different user\_ids
- Use our side channel to detect hash collisions

Bruteforce (offline phase):

- Test all possible kaddrs
- Match hash collisions





- 👓 Exploit **futex hash table**
  - `user_id`: user address
  - `kaddr`: `mm_struct`



- 🔦 Exploit **futex hash table**
  - `user_id`: user address
  - `kaddr`: `mm_struct`
- 🔦 Consider `mm_struct` alignment
  - $\sim 2^{35.5}$  different kernel addresses



- 🔦 Exploit **futex hash table**
  - `user_id`: user address
  - `kaddr`: `mm_struct`
- 🔦 Consider `mm_struct` alignment
  - $\sim 2^{35.5}$  different kernel addresses
- 🔦 Less than 65 s

# Mitigations & Conclusion





 We disclosed KernelSnitch.





We disclosed KernelSnitch.



**Linux Torvalds:** "KASLR is broken for local accesses."

- No mitigation.



📄 We disclosed KernelSnitch.

🐧 **Linux Torvalds:** "KASLR is broken for local accesses."

- No mitigation.

🐧 **Kees Cook:** "They **leak heap KASLR**, not code KASLR. Heap KASLR is **hard** to expose, even locally."

- Presented patch.



📄 We disclosed KernelSnitch.

🐧 **Linux Torvalds:** "KASLR is broken for local accesses."

- No mitigation.

🐧 **Kees Cook:** "They **leak heap KASLR**, not code KASLR. Heap KASLR is **hard** to expose, even locally."

- Presented patch.

🐧 **Linus Torvalds:** "What voodoo programming is this?"

- Rejected the patch.



📄 We disclosed KernelSnitch.

🐧 **Linux Torvalds:** "KASLR is broken for local accesses."

- No mitigation.

🐧 **Kees Cook:** "They **leak heap KASLR**, not code KASLR. Heap KASLR is **hard** to expose, even locally."

- Presented patch.

🐧 **Linus Torvalds:** "What voodoo programming is this?"

- Rejected the patch.

👓 **Still exploitable to this day!**



- We presented KernelSnitch
  - 🕒 Timing side channel on data structures
  - 🕒 Software-induced







- **We presented KernelSnitch**
  - 🕒 Timing side channel on data structures
  - 🕒 Software-induced
- **We demonstrated leakage amplification**
  - 🕒 Exploitable from user space
  - 🕒 From untrusted and isolated users



- **We presented KernelSnitch**
  - 🕶️ Timing side channel on data structures
  - 🕶️ Software-induced
- **We demonstrated leakage amplification**
  - 🕶️ Exploitable from user space
  - 🕶️ From untrusted and isolated users
- **We showed three attacks**
  - 🕶️ Covert channel
  - 🕶️ Website fingerprinting
  - 🕶️ Kernel heap pointer leak



# | KernelSnitch

## Side-Channel Attacks on Kernel Data Structures

**Lukas Maar**   Jonas Juffinger   Thomas Steinbauer   Daniel Gruss   Stefan Mangard

Graz University of Technology, Austria

NDSS 2025

> [isec.tugraz.at](https://isec.tugraz.at)