

Cross-Cache Attacks for the Linux Kernel via PCP Messaging

Claudio Migliorelli^{1,2}, Andrea Mambretti¹, Alessandro Sorniotti¹,
Vittorio Zaccaria³, Anil Kurmus¹

¹IBM Research Europe – Zurich

²EPFL

³Politecnico di Milano

Exploitation in the Linux Kernel

The Goal — what the attacker wants

Kernel memory (single pool)



Exploitation in the Linux Kernel

The Goal — what the attacker wants

Kernel memory (single pool)

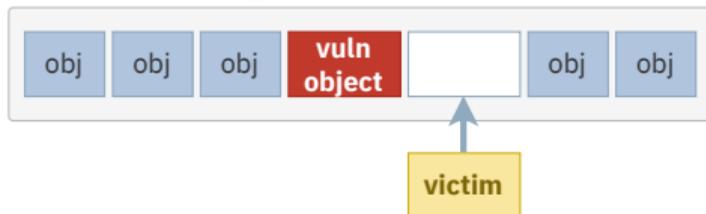


victim

Exploitation in the Linux Kernel

The Goal — what the attacker wants

Kernel memory (single pool)



Exploitation in the Linux Kernel

The Goal — what the attacker wants

Kernel memory (single pool)



Exploitation in the Linux Kernel

The Goal — what the attacker wants

Kernel memory (single pool)



Exploitation in the Linux Kernel

The Goal — what the attacker wants

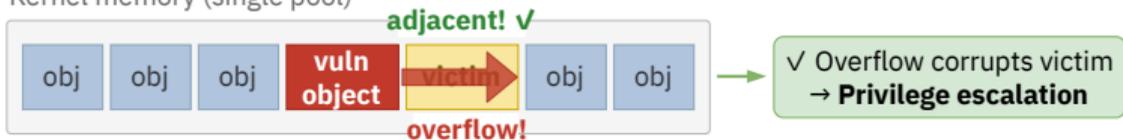
Kernel memory (single pool)



Exploitation in the Linux Kernel

The Goal — what the attacker wants

Kernel memory (single pool)



The Problem — what the kernel does

Pool A (object type X)



Pool B (object type Y)



?



Exploitation in the Linux Kernel

The Goal — what the attacker wants

Kernel memory (single pool)



The Problem — what the kernel does

Pool A (object type X)



?

Pool B (object type Y)



Exploitation in the Linux Kernel

The Goal – what the attacker wants

Kernel memory (single pool)



The Problem – what the kernel does

Pool A (object type X)



not adjacent! ✗

Pool B (object type Y)



✗ Overflow can't reach victim
→ **Attack fails**

Exploitation in the Linux Kernel

The Goal — what the attacker wants

Kernel memory (single pool)



The Problem — what the kernel does

Pool A (object type X)



Pool B (object type Y)



not adjacent! ✗

✗ Overflow can't reach victim
→ **Attack fails**

Cross-Cache Attack: force Pool A and Pool B to be adjacent
→ Overflow can cross the boundary

Previous Work

- Timing side-channels on SLUB: **PSPRAY** (USENIX 23)
- Cross-cache attacks for temporal vulnerabilities (page-recycling): **SLUBstick** (USENIX 24), **CROSS-X** (CCS 25)

Previous Work

- Timing side-channels on SLUB: **PSPRAY** (USENIX 23)
- Cross-cache attacks for temporal vulnerabilities (page-recycling): **SLUBstick** (USENIX 24), **CROSS-X** (CCS 25)

Previous work either focuses on SLUB or page-recycling cross-cache

This work:

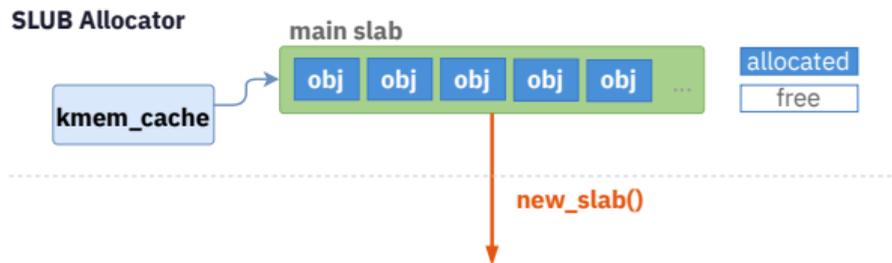
- PCPLost is a cross-cache memory massaging technique for spatial vulnerabilities (page-adjacency)
- Exploits overlooked interactions between SLUB and Page Frame Allocator
- ~ **90%** reliability in establishing adjacency, even under noise and mitigations

SLUB and Page Frame Allocator (PFA)

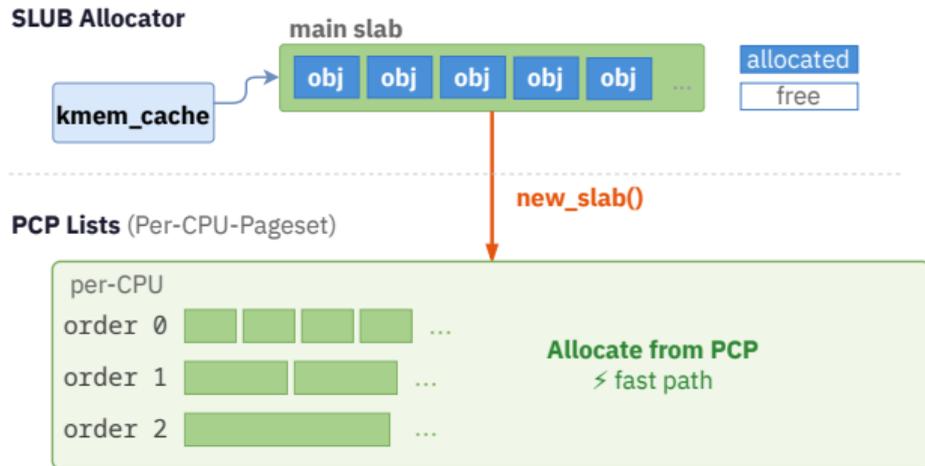
SLUB Allocator



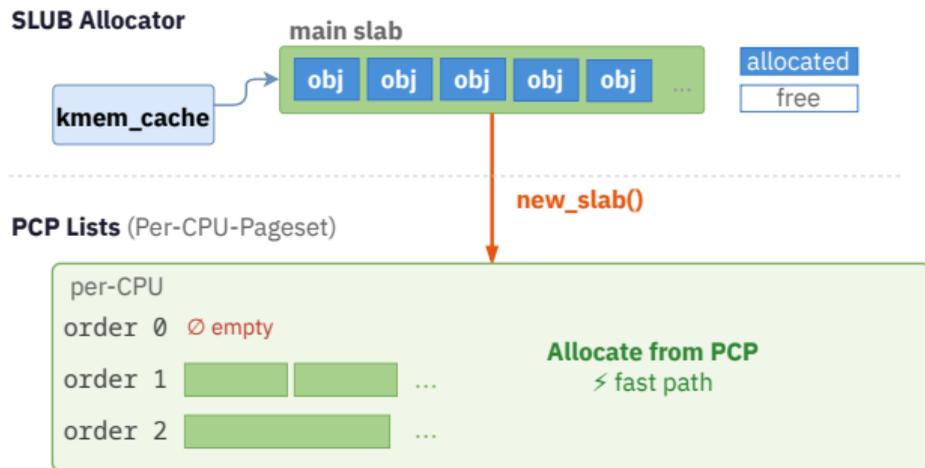
SLUB and Page Frame Allocator (PFA)



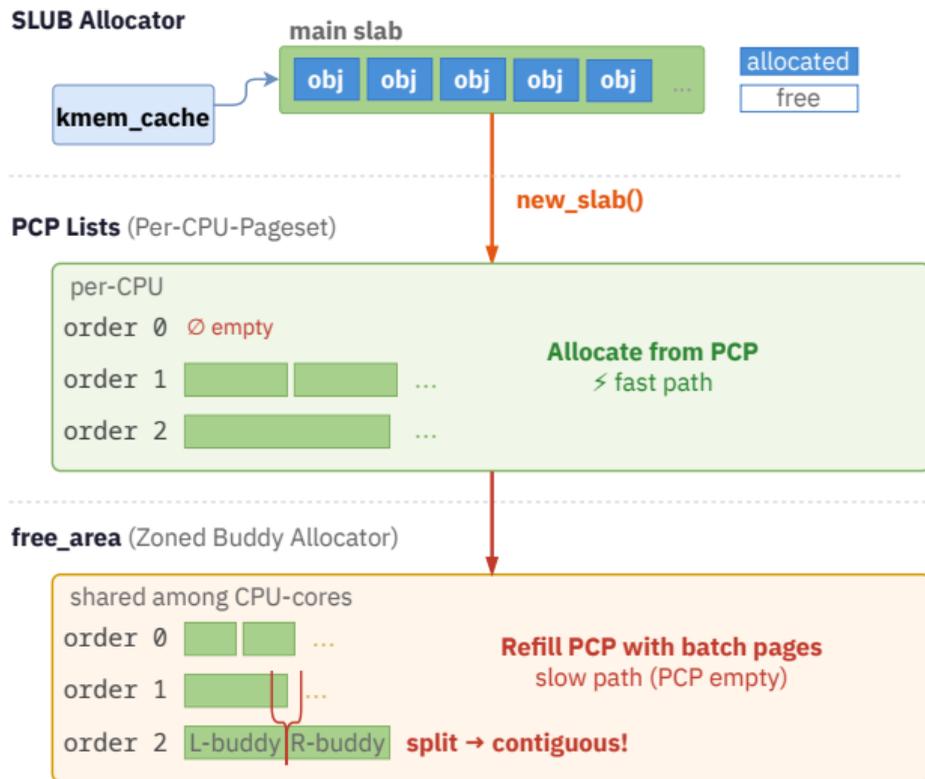
SLUB and Page Frame Allocator (PFA)



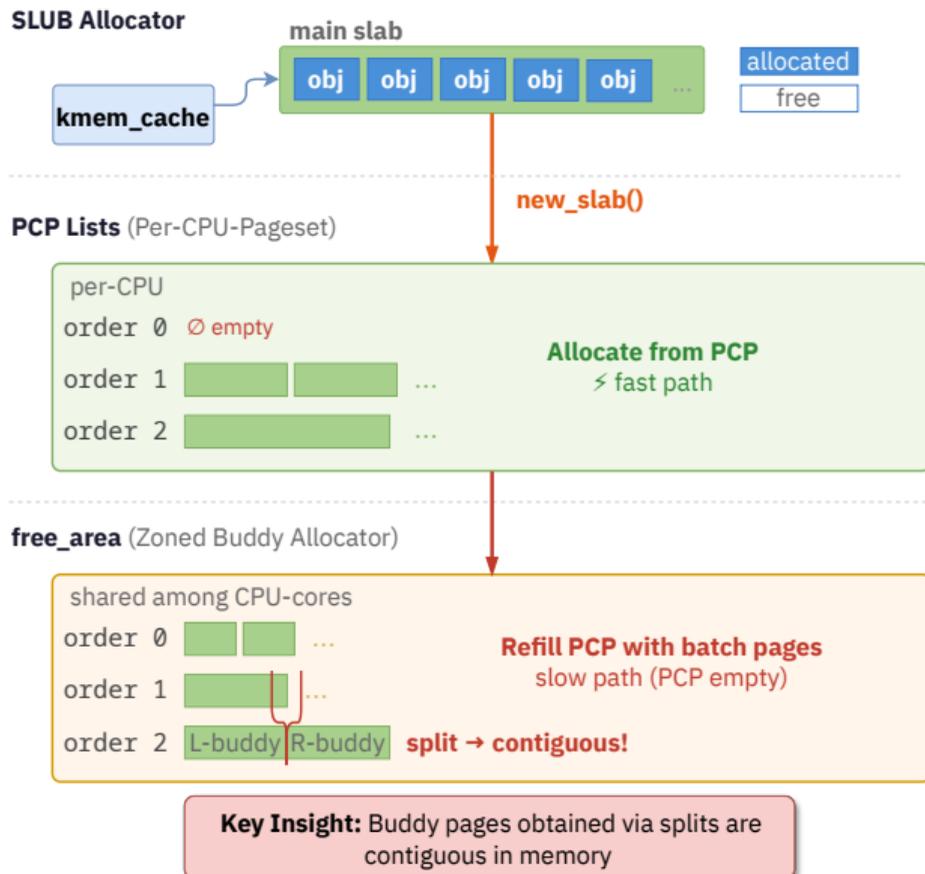
SLUB and Page Frame Allocator (PFA)



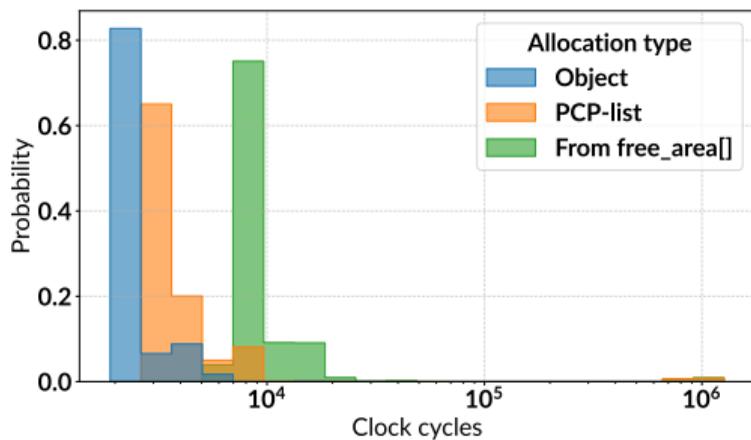
SLUB and Page Frame Allocator (PFA)



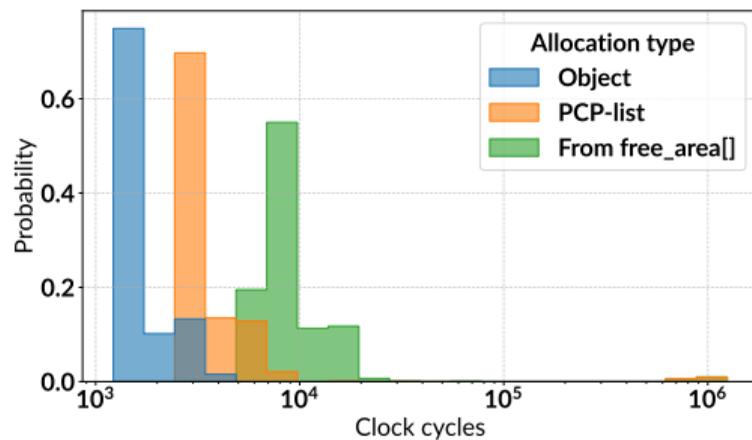
SLUB and Page Frame Allocator (PFA)



Timing side-channel on SLUB/PFA

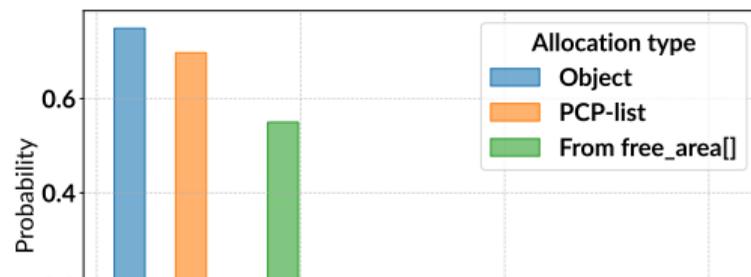
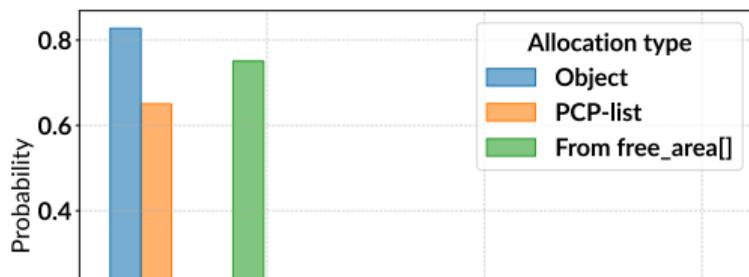


(a) `kmalloc-4k`



(b) `kmalloc-2k`

Timing side-channel on SLUB/PFA



Key Insight

PCP List vs. free_area allocations are detectable via timing side-channel

PCPLost attack core steps

1. Repeatedly drain the PCP list for the target order (Probe & Drain)
 - Forces PFA page splits to create contiguous layouts in the PCP list

PCPLost attack core steps

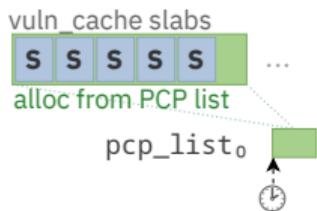
1. Repeatedly drain the PCP list for the target order (Probe & Drain)
 - Forces PFA page splits to create contiguous layouts in the PCP list
2. Allocate vulnerable and target cache from the PCP list

PCPLost attack core steps

1. Repeatedly drain the PCP list for the target order (Probe & Drain)
 - Forces PFA page splits to create contiguous layouts in the PCP list
2. Allocate vulnerable and target cache from the PCP list
3. Trigger OOB write

Same-order PCP messaging

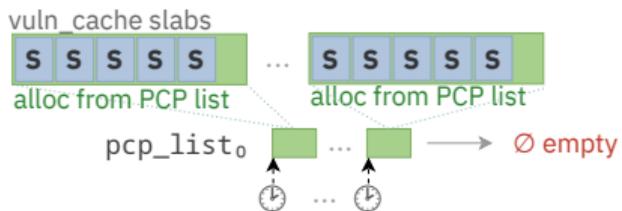
1 **Probe & Drain** — spray objects to exhaust PCP list (order = 0)



■
spray obj

Same-order PCP messaging

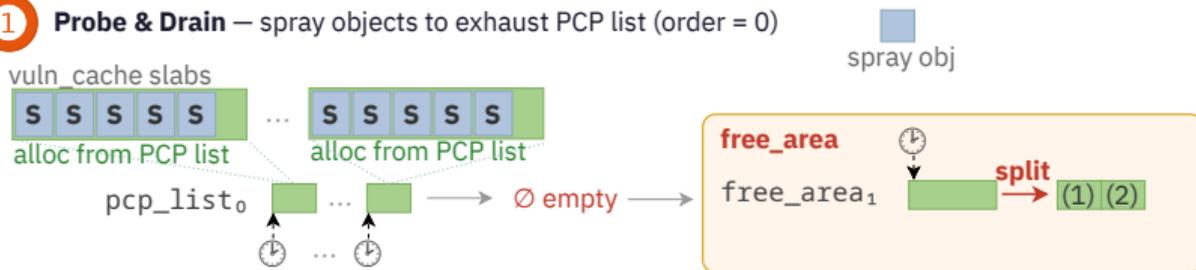
1 **Probe & Drain** — spray objects to exhaust PCP list (order = 0)



■
spray obj

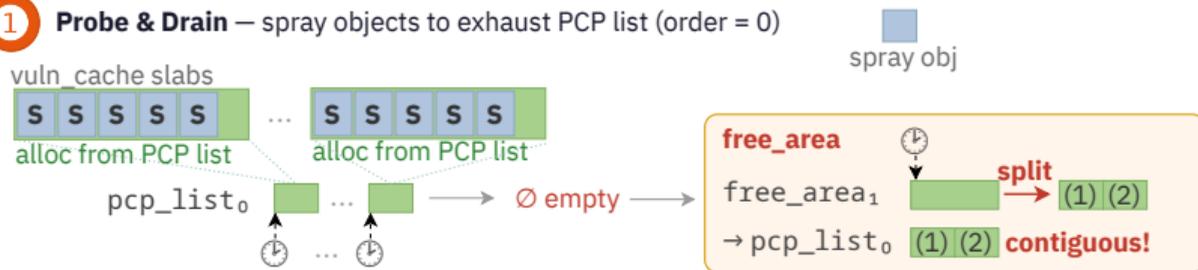
Same-order PCP messaging

1 **Probe & Drain** — spray objects to exhaust PCP list (order = 0)



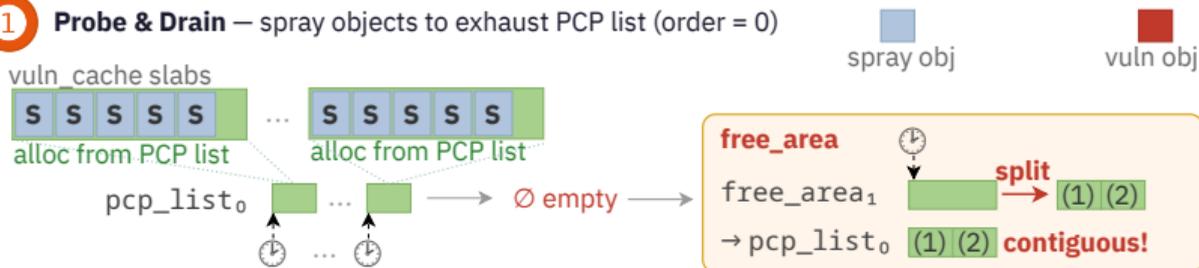
Same-order PCP messaging

1 **Probe & Drain** — spray objects to exhaust PCP list (order = 0)

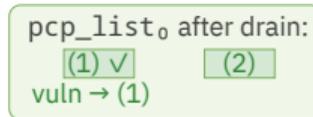


Same-order PCP messaging

- 1 **Probe & Drain** — spray objects to exhaust PCP list (order = 0)



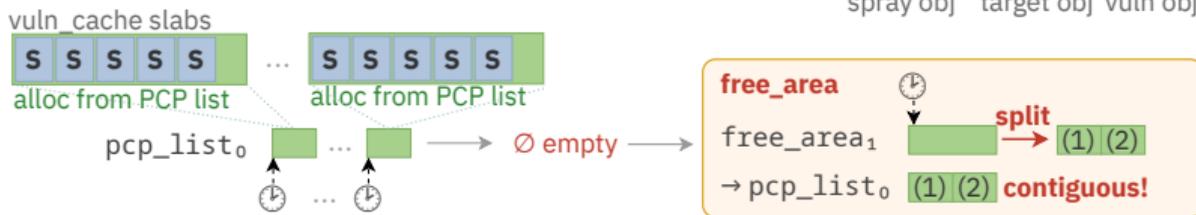
- 2 **Allocate** — vulnerable and target slabs from PCP list (contiguous pages)



Same-order PCP messaging

1 **Probe & Drain** — spray objects to exhaust PCP list (order = 0)

■ spray obj ■ target obj ■ vuln obj



2 **Allocate** — vulnerable and target slabs from PCP list (contiguous pages)

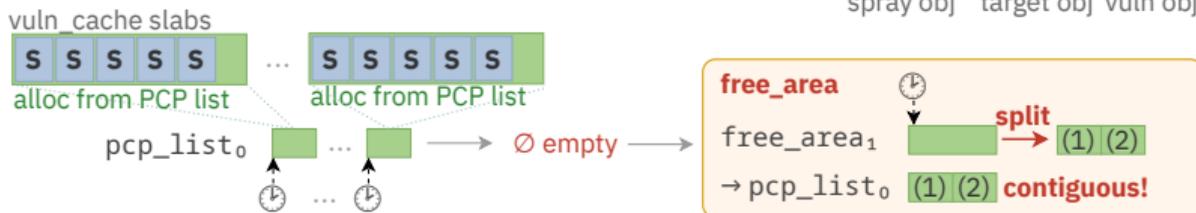


pcp_list₀ after drain:
(1) ✓ (2) ✓
vuln → (1) target → (2)

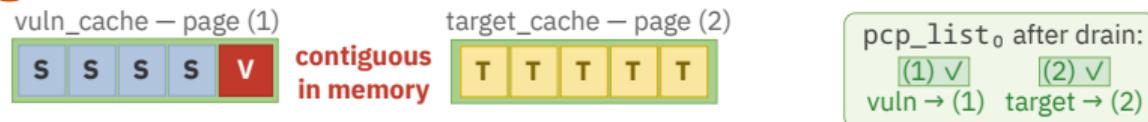
Same-order PCP messaging

1 **Probe & Drain** — spray objects to exhaust PCP list (order = 0)

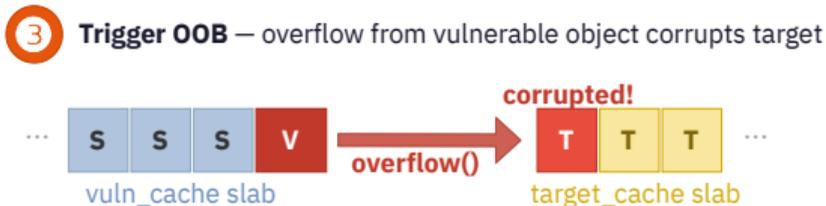
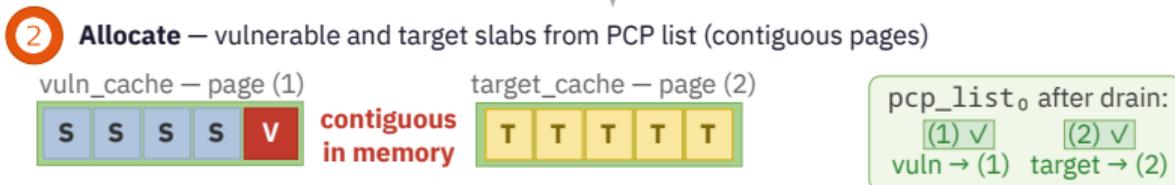
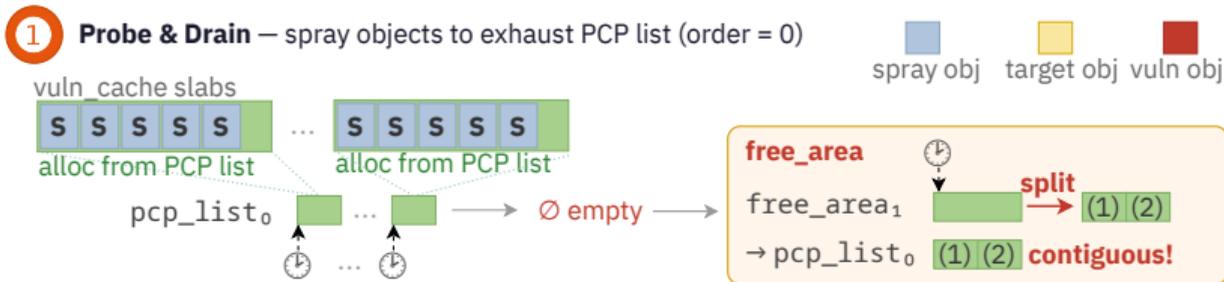
■ spray obj ■ target obj ■ vuln obj



2 **Allocate** — vulnerable and target slabs from PCP list (contiguous pages)



Same-order PCP messaging



Evaluation

Reliability evaluation

Vuln. cache	Target cache	Idle (%)	Noise (%)
<i>Same-order ($n_v = n_t$)</i>			
kmalloc-16	kmalloc-64	98.83	97.42
kmalloc-64	kmalloc-128	96.80	90.94
kmalloc-128	kmalloc-16	99.84	98.28
<i>Cross-order ($n_v > n_t$)</i>			
kmalloc-2k	kmalloc-16	94.30	87.58
kmalloc-4k	kmalloc-64	93.20	92.42
kmalloc-4k	kmalloc-128	90.31	90.08

With CPU pinning. 800 runs per scenario.

Evaluation

Reliability evaluation

Vuln. cache	Target cache	Idle (%)	Noise (%)
<i>Same-order ($n_v = n_t$)</i>			
kmalloc-16	kmalloc-64	98.83	97.42
kmalloc-64	kmalloc-128	96.80	90.94
kmalloc-128	kmalloc-16	99.84	98.28
<i>Cross-order ($n_v > n_t$)</i>			
kmalloc-2k	kmalloc-16	94.30	87.58
kmalloc-4k	kmalloc-64	93.20	92.42
kmalloc-4k	kmalloc-128	90.31	90.08

With CPU pinning. 800 runs per scenario.

Real-world CVEs

- PCPLost validated on **6 CVEs**
 - **Spatial** bugs
 - **Temporal** bugs via pivoting
- Bypassing mainline mitigations
SLAB_FREELIST_RANDOM
SLAB_FREELIST_HARDENED
- Bypassing experimental SLAB_VIRTUAL

Evaluation

Reliability evaluation

Vuln. cache	Target cache	Idle (%)	Noise (%)
<i>Same-order ($n_v = n_t$)</i>			
kmalloc-16	kmalloc-64	98.83	97.42
kmalloc-64	kmalloc-128	96.80	90.94
kmalloc-128	kmalloc-16	99.84	98.28
<i>Cross-order ($n_v > n_t$)</i>			
kmalloc-2k	kmalloc-16	94.30	87.58
kmalloc-4k	kmalloc-64	93.20	92.42
kmalloc-4k	kmalloc-128	90.31	90.08

With CPU pinning. 800 runs per scenario.

Real-world CVEs

- PCP/lost validated on **6 CVEs**
 - **Spatial** bugs
 - **Temporal** bugs via pivoting
- Bypassing mainline mitigations
SLAB_FREELIST_RANDOM
SLAB_FREELIST_HARDENED
- Bypassing experimental SLAB_VIRTUAL
Mitigated by SLAB_VIRTUAL_GP (with guard pages)

Conclusion

- **PCP list internals** expose a **timing side-channel**

Conclusion

- **PCP list internals** expose a **timing side-channel**
- We use the side-channel to mount PCPLost exploiting **SLUB/PFA interactions**

Conclusion

- **PCP list internals** expose a **timing side-channel**
- We use the side-channel to mount PCPLost exploiting **SLUB/PFA interactions**
- Mainline kernel defenses stop at SLUB
 - **The Page Allocator remains unguarded**
 - Cross-cache security requires mitigations at that level

Conclusion

- **PCP list internals** expose a **timing side-channel**
- We use the side-channel to mount PCPLost exploiting **SLUB/PFA interactions**
- Mainline kernel defenses stop at SLUB
 - **The Page Allocator remains unguarded**
 - Cross-cache security requires mitigations at that level
- We disclosed PCPLost to kernel maintainers

Evaluation with real-world CVEs

CVE	Category	Vulnerable cache	Target cache	Mitgations bypassed
CVE-2024-53141	OOB	kmalloc-cg-1k	kmalloc-cg-2k	SV,FH
CVE-2021-22555	OOB/UAF	kmalloc-cg-4k	kmalloc-cg-1k	FR,SV,FH
CVE-2022-0185	OOB	kmalloc-4k	kmalloc-cg-4k	SV,FH
CVE-2023-0461	UAF	kmalloc-512	kmalloc-cg-1k	FR,SV,FH
CVE-2021-3715	UAF	kmalloc-192	kmalloc-32	FR,FH
CVE-2022-27666	OOB	-	kmalloc-4k	FR,FH

SLAB_FREELIST_RANDOM bypass

