Towards Safe Kernel Extensibility With eBPF

sooyee@cs.ubc.ca

University of British Columbia

Soo Yee Lim

eBPF Summit

@ebpfsummit

eBPF Lacks Run-time Memory Safety

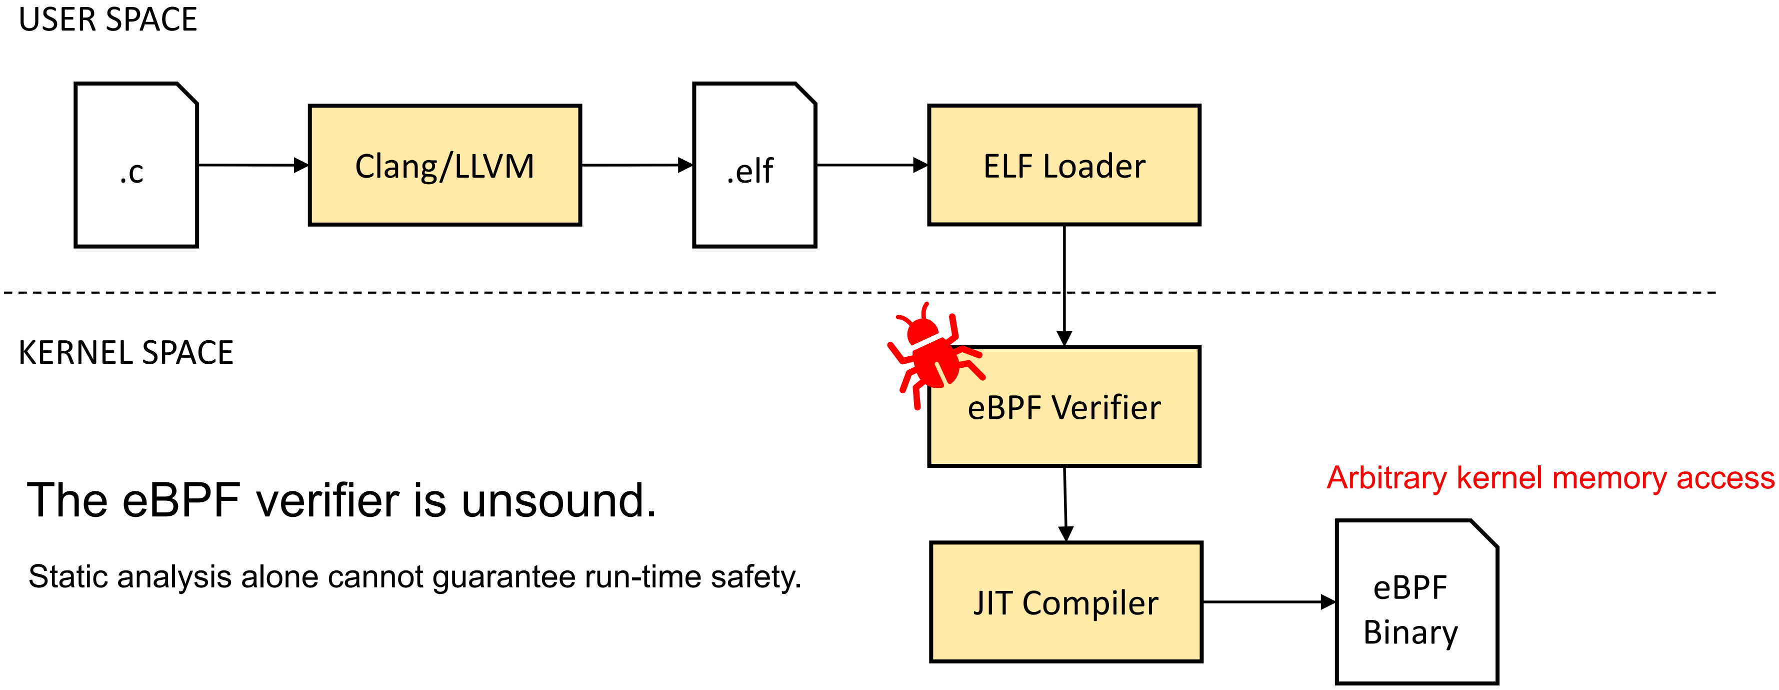

Arbitrary kernel memory access

! The eBPF verifier is unsound.

Static analysis alone cannot guarantee run-time safety.

SafeBPF

Dynamic Sandboxing of eBPF Programs with Software Fault Isolation and Hardware-Assisted Memory Tagging to enforce spatial memory safety at runtime.

Lim, Soo Yee, et al. "SafeBPF: Hardware-assisted Defense-in-depth for eBPF Kernel Extensions." *Proceedings of the 2024 on Cloud Computing Security Workshop*. 2024.

Software Fault Isolation

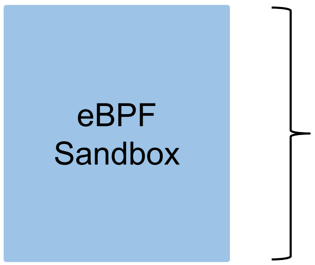

eBPF
Sandbox

addr = 0xA000
len = 4096 bytes

Consider an invalid memory access at address 0xB123

and_mask = 0xFFF; or_mask = 0xA000

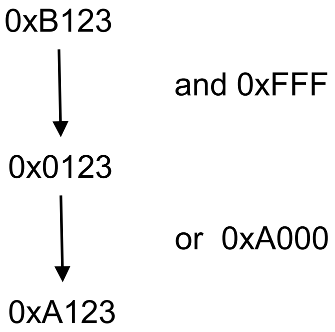

0xB123

and 0xFFF

0x0123

or 0xA000

0xA123

All memory accesses always fall within the eBPF sandbox.

Memory Tagging with ARM MTE

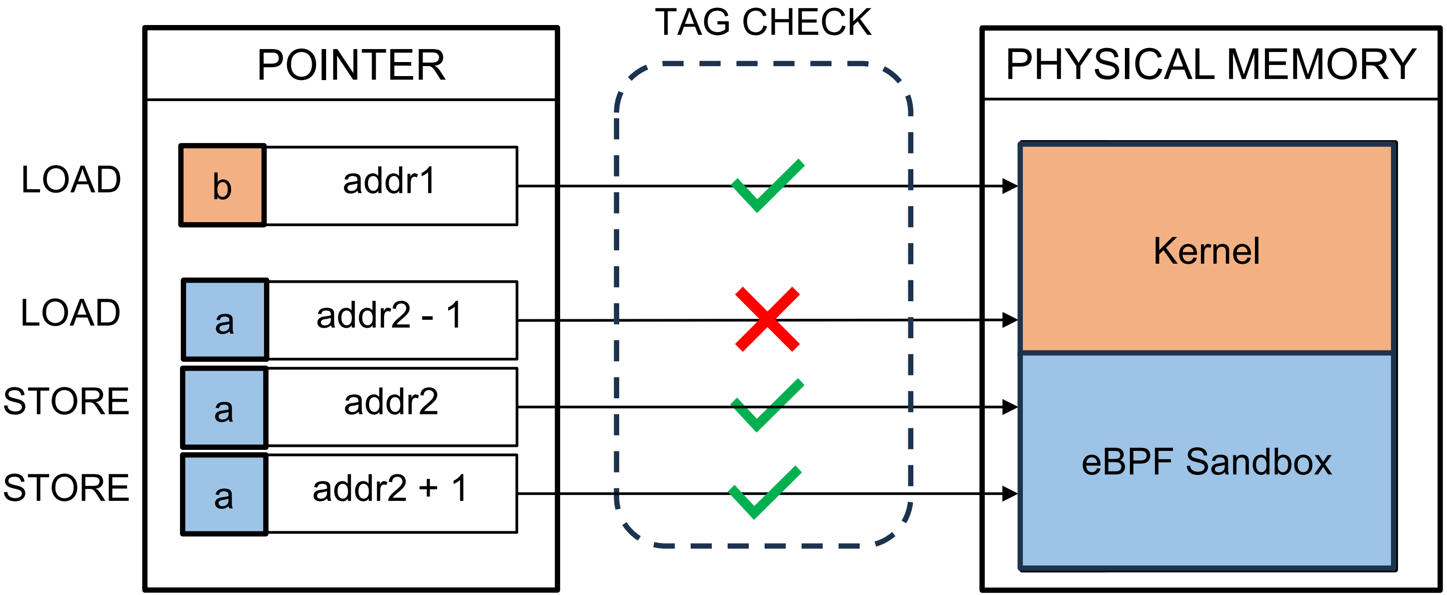

Evaluation Results

- SafeBPF successfully prevents 7 high-severity vulnerabilities.
- SafeBPF incurs 0% - 4% overhead on webserver macrobenchmarks.

Email: sooyee@cs.ubc.ca
Website: https://s00y33.github.io/

Proceedings will be available at the the Cloud Computing Security Workshop 2024.

Other eBPF Work

Lim, Soo Yee, Bogdan Stelea, Xueyuan Han, and Thomas Pasquier. "Secure namespaced kernel audit for containers." In *Proceedings of the ACM Symposium on Cloud Computing*. 2021.

Cao, Xuechun, Shaurya Patel, **Soo Yee Lim**, Xueyuan Han, and Thomas Pasquier. "FetchBPF: Customizable Prefetching Policies in Linux with eBPF." In *USENIX Annual Technical Conference*. 2024.