The HITCHHIKER's Guide to High-Assurance System Observability Protection with Efficient Permission Switches

Chuqi Zhang, Jun Zeng, Yiming Zhang, Adil Ahmad, Fengwei Zhang, Hai Jin, and Zhenkai Liang ACM CCS, October 2024 Salt Lake City, U.S.A.









Observability captures a system's historical states



Logs are vulnerable when the OS is compromised



Logs are vulnerable when the OS is compromised



Logs are vulnerable when the OS is compromised





in-memory protection with eventual persistence





in-memory protection with eventual persistence





Low-assurance off-the-shelf environment





Low-assurance off-the-shelf environment







Problem 2: High synchronous protection overhead



Research questions we ask



Research questions we ask



Requirements for high-assurance secure environment

<u>Principled approach</u> to redesigning software components.









Memory isolation:

Stage-2 Page Table (S2PT) / Granule Protection Table (GPT)



Security monitor

Device isolation:

System Memory Management Unit (SMMU)

Memory isolation:

Stage-2 Page Table (S2PT) / Granule Protection Table (GPT)





Security monitor

























For security, log protection can not be unbounded asynchronous.



(Will take an unbounded delay Δt for async. protection)

For security, log protection can not be **unbounded asynchronous.**



For security, log protection can not be unbounded asynchronous.





2a. Unified log collection using eBPF



2a. Unified log collection using eBPF





Unified logs to be protected with *controlled delay*.



- Periodically triggers protections (timer interrupts).
- Prioritized and isolated timer interrupts.







We evaluated HITCHHIKER

Security environment's TCB and attack surface

Log protection effectiveness

Log protection efficiency

Secure environment

TCB: Reduced more than 10× compared to off-the-shelf environments.

Mitigated issues of off-the-shelf environment (TrustZone)

Issues	Issues
I01. SW drivers run in the TEE kernel space	I02. Wide interfaces between TEE system subcomponents
IO3. Excessively large TEE TCBs	I04. TAs can map physical memory in the NW
105. Information leaks to NW through debugging channel	I06. Absent or weak ASLR implementations
107. No stack cookies, guard pages, or execution protection	I08. Lack of software-independent TEE integrity reporting
I09. Ill-supported TA revocation	110. Validation bugs within the secure monitor
I11. Validation bugs within TAs	I12. Validation bugs within the trusted kernel
I13. Validation bugs in secure boot loader	I14. Bugs in memory protection
I15. Bugs in configuration of peripherals	I16. Bugs in security mechanisms
I17. Concurrency bugs (from multiple TAs)	I18. Software side-channels

16/18 TrustZone issues were mitigated.

Effectiveness on controlled short log protection

Protection Completion Ratio



Effectiveness on controlled short log protection

Protection Completion Ra Total time (ms) to compromise the full system



Performance efficiency on real-world applications

Runtime performance overhead



Takeaways

The state of the observability protection systems

• Low assurance in-memory environments and inefficiency under stress.

HitchHiker's design principles for observability protection

- Principled debloating strategies to designing high-assurance environments.
- Controlled protection delay using bounded timer and fast permission switches.

Code will (soon) be available at:



THANKS!