ETH zürich

System Security Group

Arm CCA-based Normal-world Enclaves with Device Isolation

Edouard Michelin

January 27, 2025

Semester Project

Motivation

Motivation

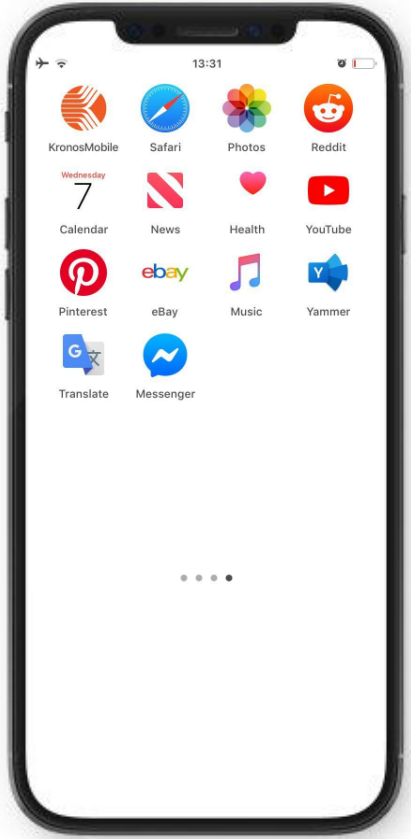
13:31
KronosMobile
Safari
Photos
Reddit
Wednesday
7
Calendar
News
Health
YouTube
Pinterest
eBay
Music
Yammer
Translate
Messenger

Motivation

Motivation

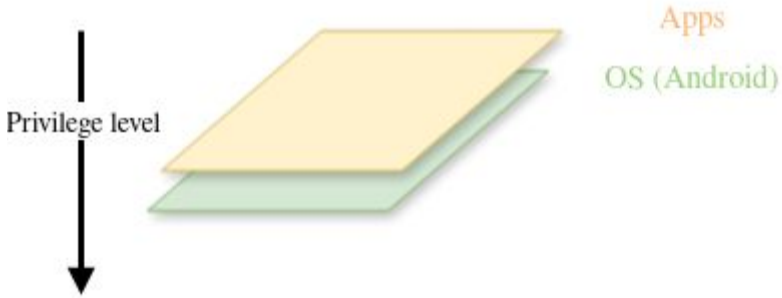
Privilege level
Apps
OS (Android)


13:31
KronosMobile
Safari
Photos
Reddit

Motivation

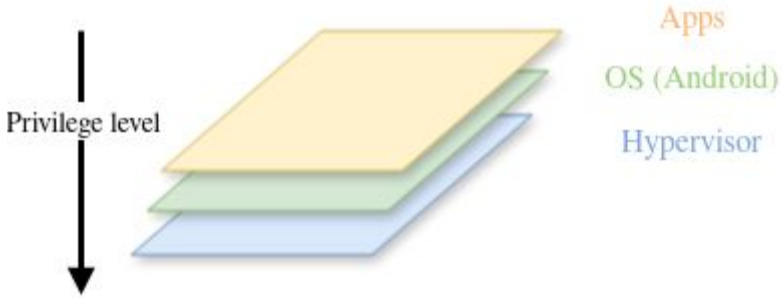
Privilege level
Apps
OS (Android)
Hypervisor


13:31
KronosMobile
Safari
Photos
Reddit

Motivation

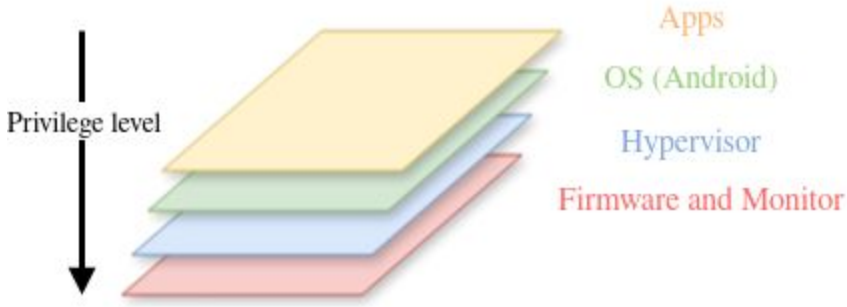
Privilege level
Apps
OS (Android)
Hypervisor
Firmware and Monitor


13:31
KronosMobile
Safari
Photos
Reddit

Motivation

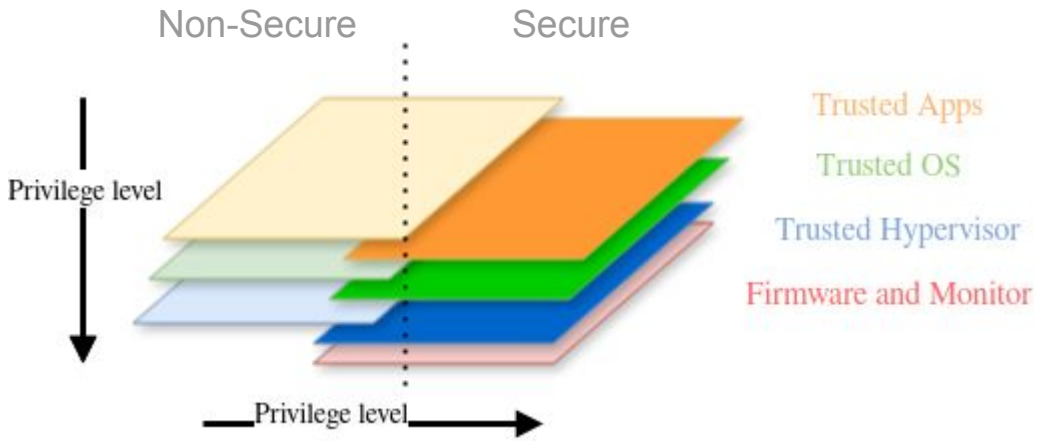
Non-Secure
Secure
Privilege level
Privilege level
Trusted Apps
Trusted OS
Trusted Hypervisor
Firmware and Monitor


13:31
KronosMobile
Safari
Photos
Reddit

Motivation

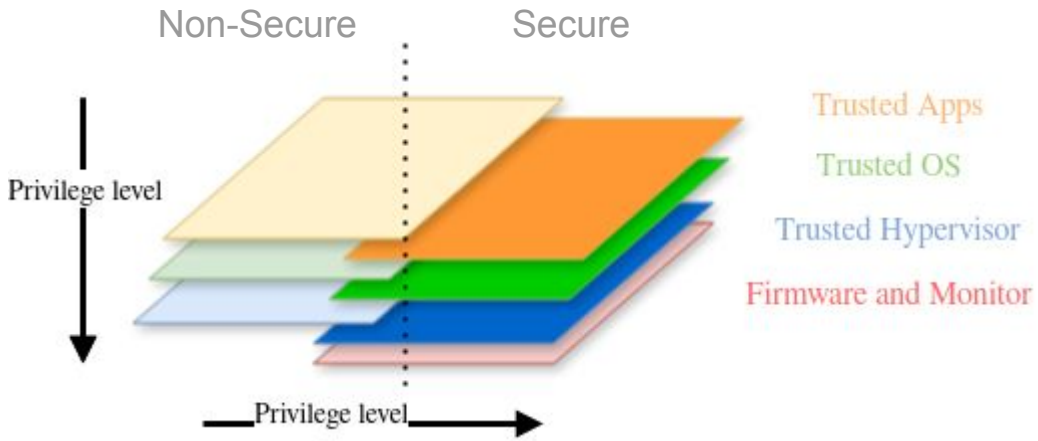

- Control divided among manufacturers, OS vendors and users.
- Users have the least control

  ⇒ Have to trust every stakeholder

Motivation

Non-Secure Secure

Privacy policies?

- GDPR (General Data Privacy Regulation)
- FADP (Federal Act on Data Protection)

- Control divide
- Users have th

⇒ Have to trust every stakeholder

Motivation

AP Exclusive: Google tracks your movements, like it or not

BY RYAN NAKASHIMA
Published 12:15 AM GMT+1, August 14, 2018

Share

SAN FRANCISCO (AP) — Google wants to know where you go so badly that it records your movements even when you explicitly tell it not to.

An Associated Press investigation found that many Google services on Android devices and iPhones store your location data even if you've used a privacy setting that says it will prevent Google from doing so.

Computer-science researchers at Princeton confirmed these findings at the AP's request.

For the most part, Google is upfront about asking permission to use your location information. An app like Google Maps will remind you to allow access to location if you use it for navigating. If you agree to let it record your location over time, Google Maps will display that history for you in a "timeline" that maps out your daily movements.

- Control div
- Users have

⇒ Have to

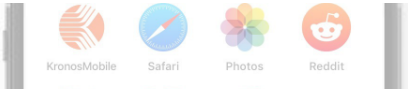

What if users could truly own their phone?

Threat model & Assumptions

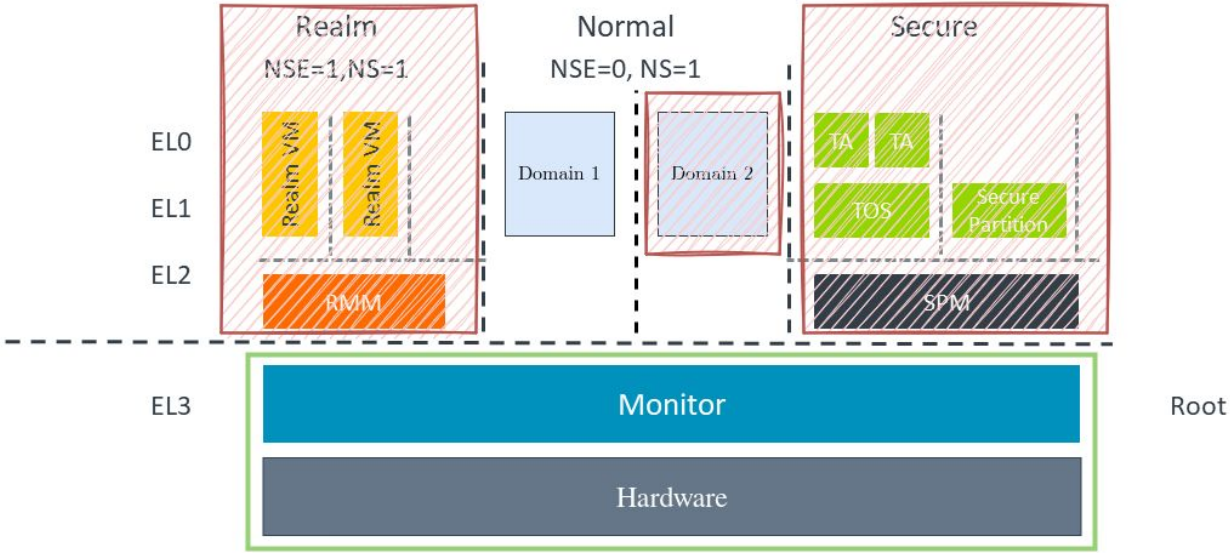

- Goal: **confidentiality** and **integrity** of **code**, **data**, and **peripheral interaction** – with a **small TCB**.
- Availability and side-channel attacks out of scope.

Background - Arm **TrustZone**

- 2 worlds: Secure & Non-Secure (a.k.a. Normal)
- Isolation enforced by Address Space Controllers (ASC)

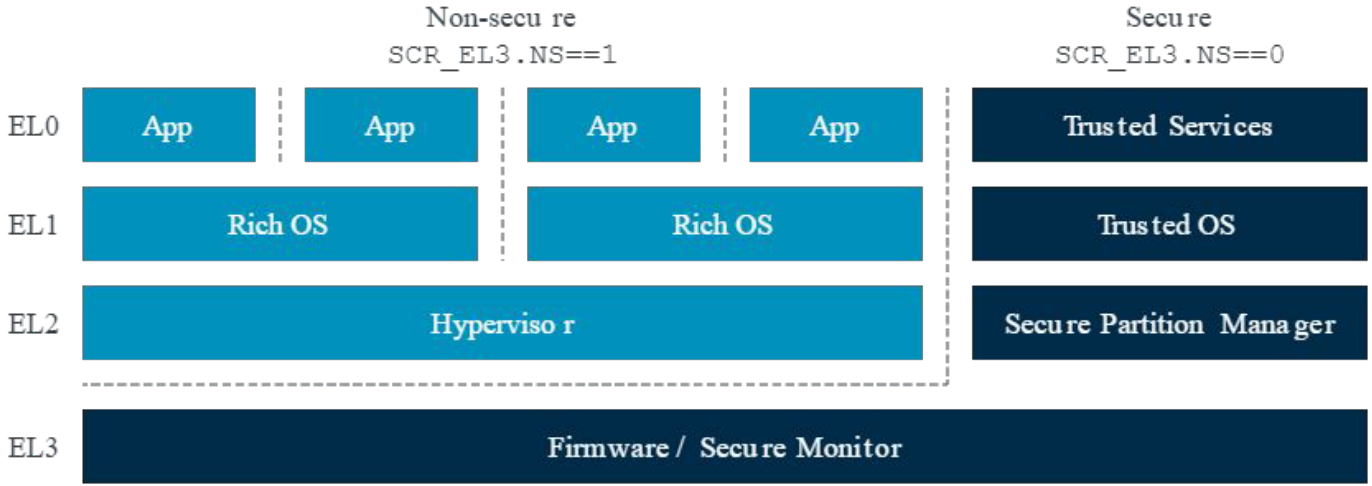

https://developer.arm.com/documentation/102418/0102/TrustZone-in-the-processor/Security-States

Background - Arm **TrustZone**

- 2 worlds: Secure & Non-Secure (a.k.a. Normal)
- Isolation enforced by Address Space Controllers (ASC)

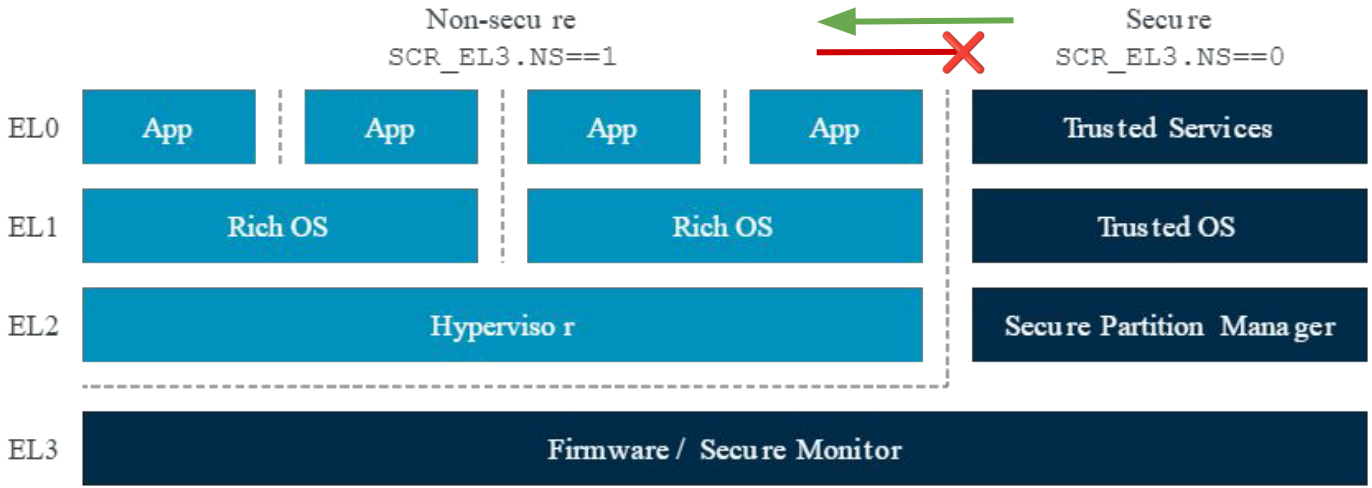

https://developer.arm.com/documentation/102418/0102/TrustZone-in-the-processor/Security-States

Background - Arm **TrustZone**

- 2 worlds: Secure & Non-Secure (a.k.a. Normal)
- Isolation enforced by Address Space Controllers (ASC)

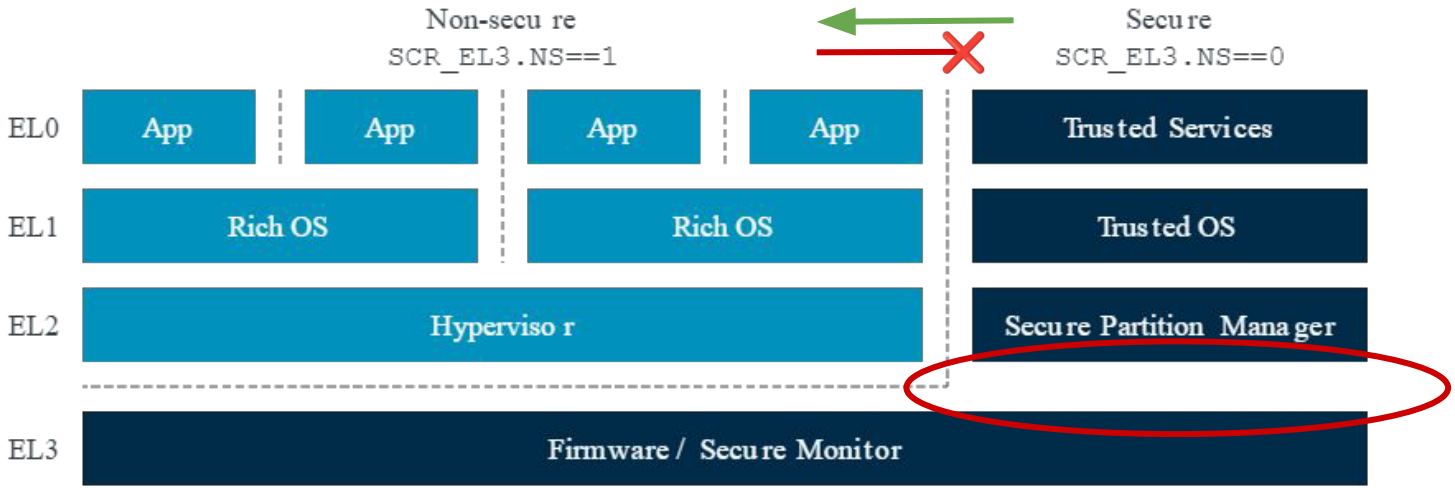

https://developer.arm.com/documentation/102418/0102/TrustZone-in-the-processor/Security-States

Background - Arm **TrustZone**

- 2 worlds: Secure & Non-Secure (a.k.a. Normal)
- Isolation enforced by Address Space Controllers (ASC)

Limitations

- Secure world over-privileged
- Monitor not fully isolated
- **Core-wise partitioning capabilities board-dependant**
- **Asynchronous aborts**

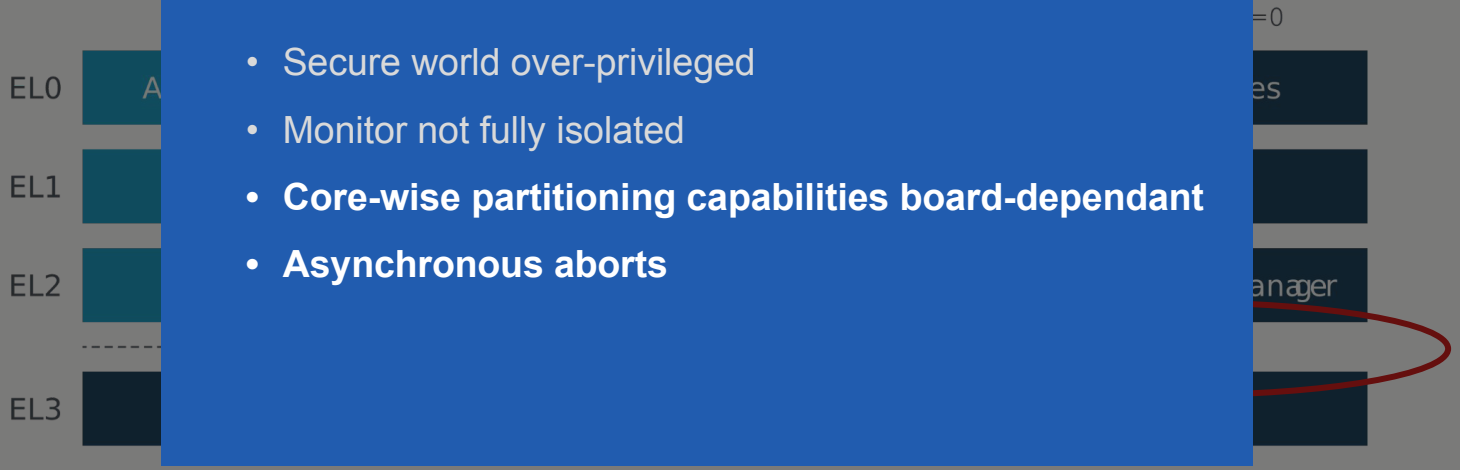

https://developer.arm.com/documentation/102418/0102/TrustZone-in-the-processor/Security-States

Background - Arm **CCA**

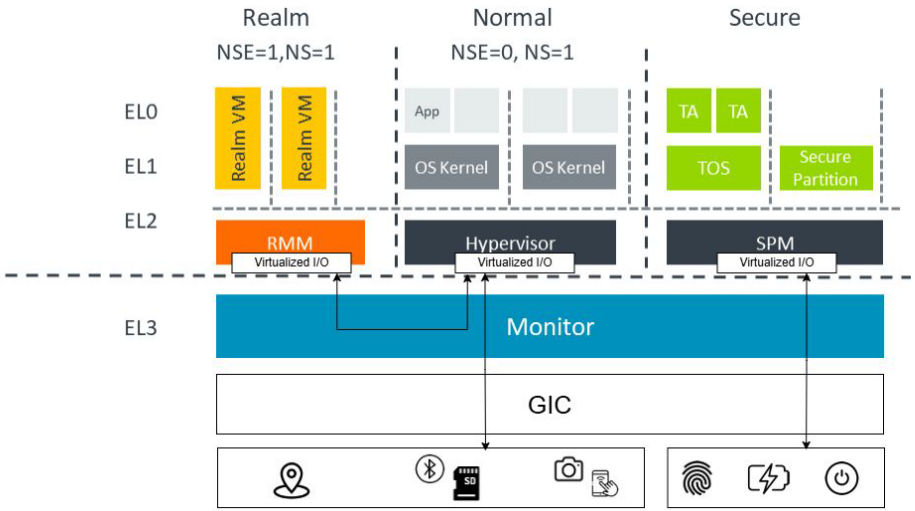

- 2 new worlds:
  - ⇒ Root (Monitor)
  - ⇒ Realm (confidentials VMs)
- Isolation enforced by Granule Protection Check (GPC) during address translation
- GPC checks assignments of regions in Granule Protection Table (GPT)

NS = Non-Secure – NSE = Non-Secure Extension
RMM = Realm Management Monitor
TA = Trusted App – TOS = Trusted OS
SPM = Secure Partition Manager
GIC = Generic Interrupt Controller

https://developer.arm.com/documentation/den0125/0300/Arm-CCA-Hardware-Architecture

Background - Arm **CCA**

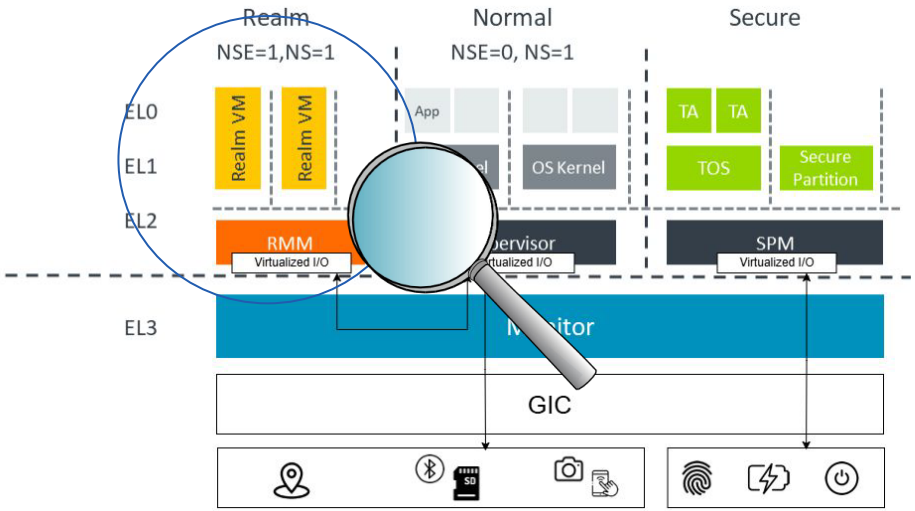

- 2 new worlds:
  - ⇒ Root (Monitor)
  - ⇒ Realm (confidentials VMs)
- Isolation enforced by Granule Protection Check (GPC) during address translation
- GPC checks assignments of regions in Granule Protection Table (GPT)

NS = Non-Secure – NSE = Non-Secure Extension
RMM = Realm Management Monitor
TA = Trusted App – TOS = Trusted OS
SPM = Secure Partition Manager
GIC = Generic Interrupt Controller

https://developer.arm.com/documentation/den0125/0300/Arm-CCA-Hardware-Architecture

Background - Arm **CCA**

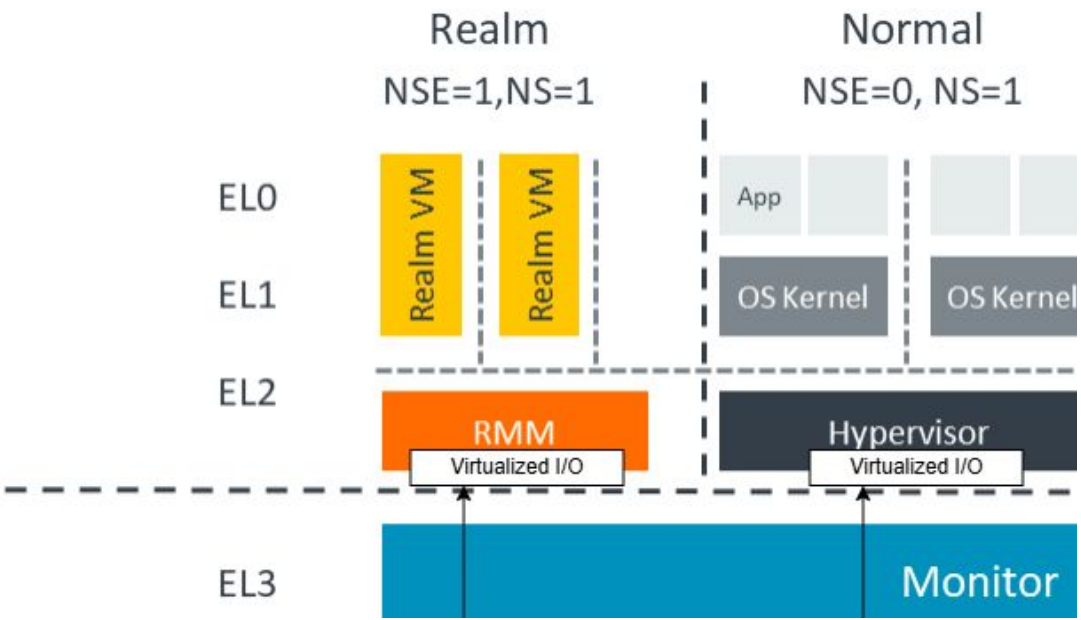

https://developer.arm.com/documentation/den0125/0300/Arm-CCA-Hardware-Architecture

Background - Arm **CCA**

1. No physical access to devices

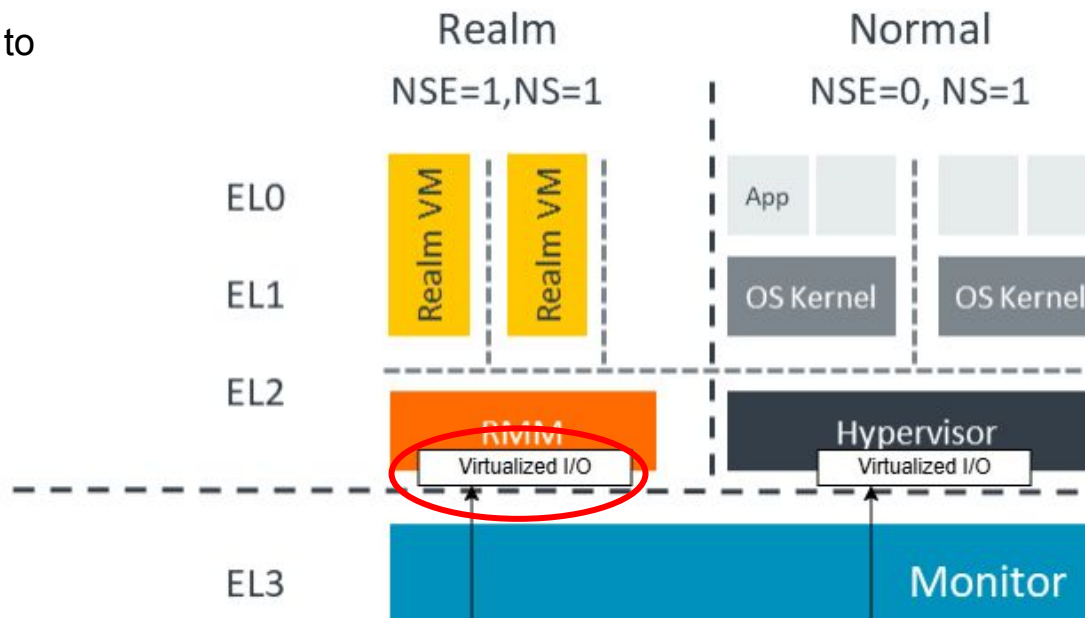

https://developer.arm.com/documentation/den0125/0300/Arm-CCA-Hardware-Architecture

Background - Arm **CCA**

1. No physical access to devices
2. Resources virtualized
   ⇒ RMM in TCB

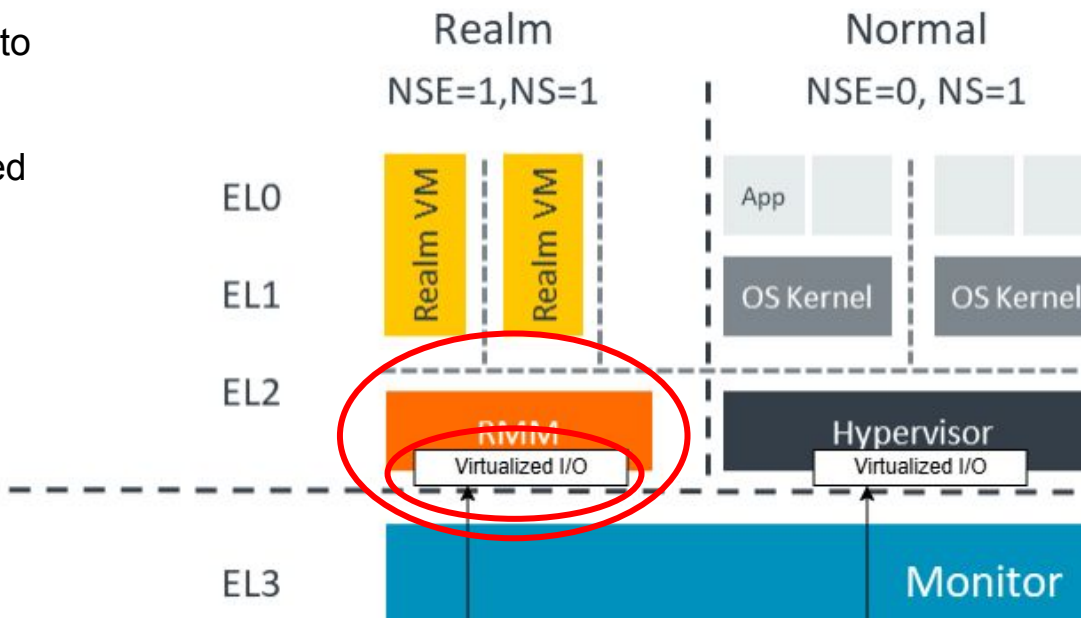

https://developer.arm.com/documentation/den0125/0300/Arm-CCA-Hardware-Architecture

Background - Arm **CCA**

1. No physical access to devices
2. Resources virtuali
   ⇒ RMM in TCB

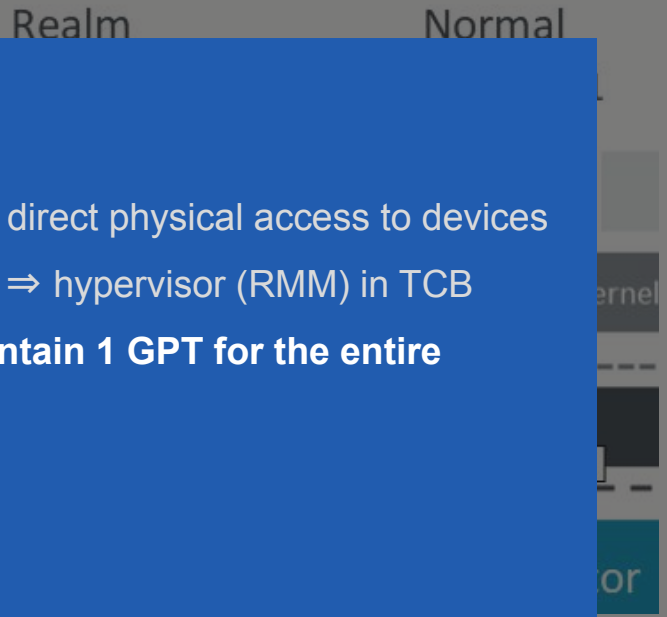

Limitations

- Realms do not have direct physical access to devices
- Realms are for VMs ⇒ hypervisor (RMM) in TCB
- **Default setups maintain 1 GPT for the entire system**

https://developer.arm.com/documentation/den0125/0300/Arm-CCA-Hardware-Architecture

Background - **Interrupts on Arm**

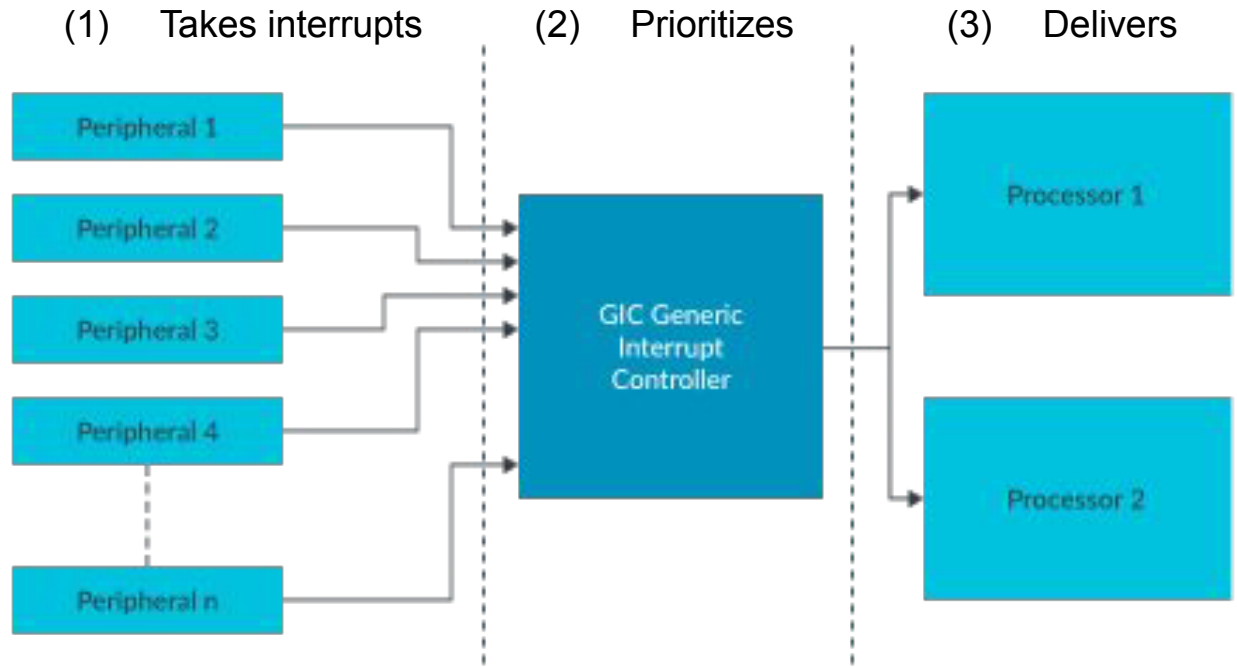

https://developer.arm.com/documentation/198123/0302/What-is-a-Generic-Interrupt-Controller-

Reminder: Threat model & Assumptions

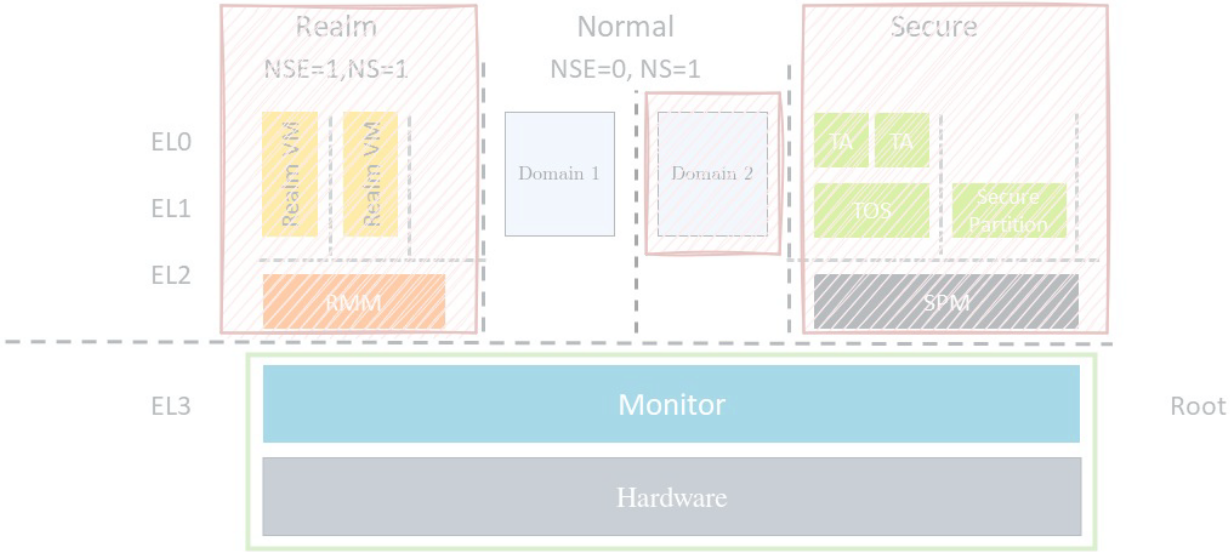

- Goal: **confidentiality** and **integrity** of **code**, **data**, and **peripheral interaction** – with a **small TCB**.
- Availability and side-channel attacks out of scope.

Background - Related works - **TEEtime**

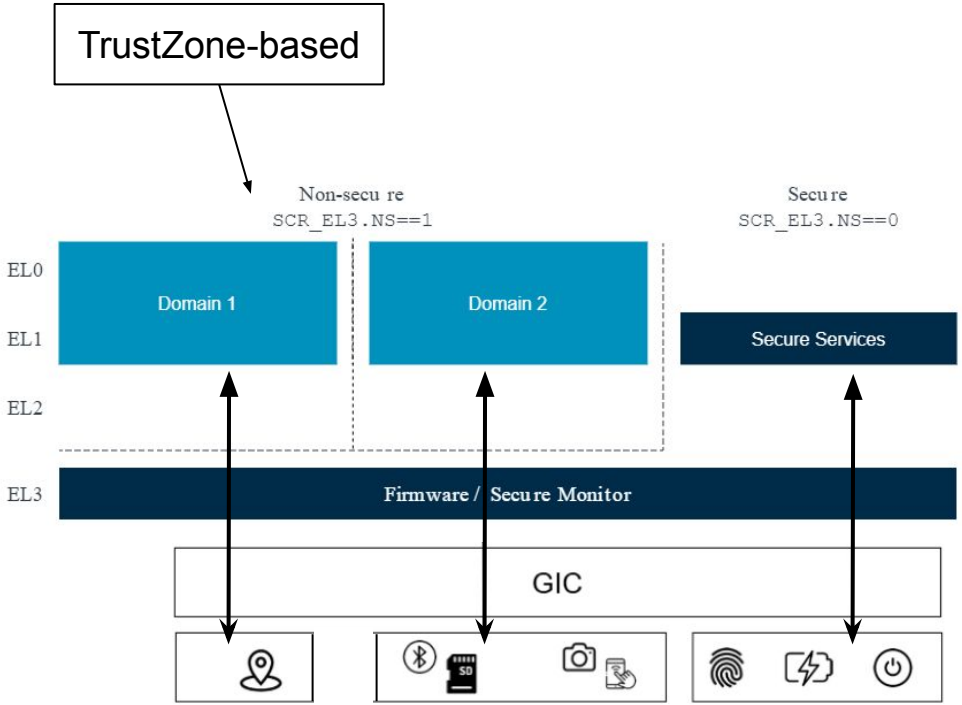

- All software domains <u>equally privileged</u>.
- Temporal sharing vs Spatial sharing (multi-core).
- Domains own peripheral interaction.
- Interrupt isolation (via GIC):
  - ⇒ Configuration/handling only by owner domain.
  - ⇒ Interrupts only trigger in owner domain

Background - Related works - **SHELTER**

- Per-Enclave GPT ⇒ 1 Address Space per Core
- Interrupt isolation out of scope
- EL0 only – Syscalls to (untrusted) host OS

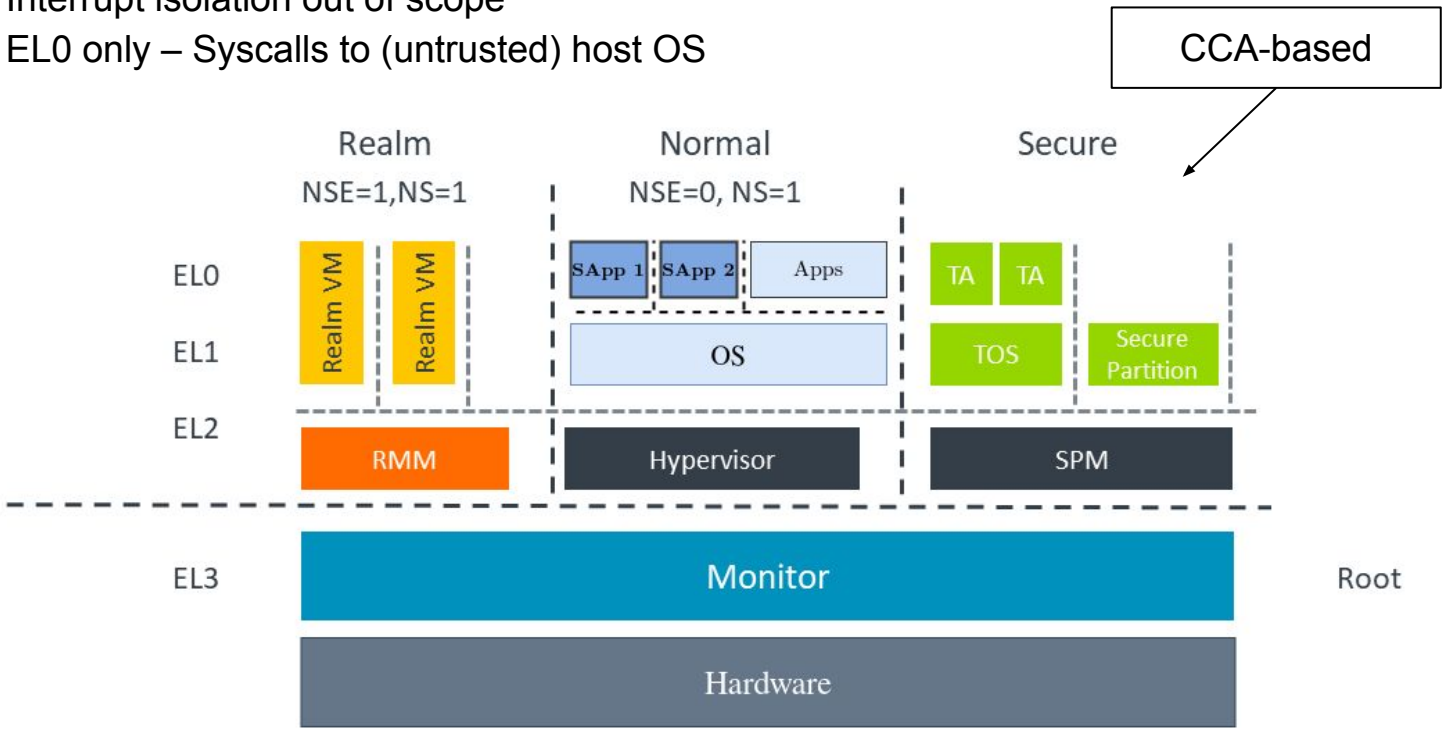

Background - Related works - **DevLore**

- Integrated device isolation for Realm VMs
  - MMIO ✔ (GPT, S2 tables)
  - DMA ✔ (GPT, S2 tables)
  - Interrupts ✔ (RMM)
- GIC to Root memory (config binding)
- Interrupts routing relies on trustful RMM

Problem statement

Design	Core-wise partitioning	Memory isolation	Device/interrupt isolation	Hypervisor-free design	Arm CCA-enabled
CCA Realms	N/A	✔	❌	❌	✔
TEEtime	Board-dependant	✔	✔	✔	❌
SHELTER	✔	✔	❌	❌	✔
DevLore	✔	✔	✔	❌	✔
This Work	✔	✔	✔	✔	✔

Opportunity for a new architecture leveraging Arm CCA

Gathering the best of TEEtime, SHELTER and DevLore

Arm CCA

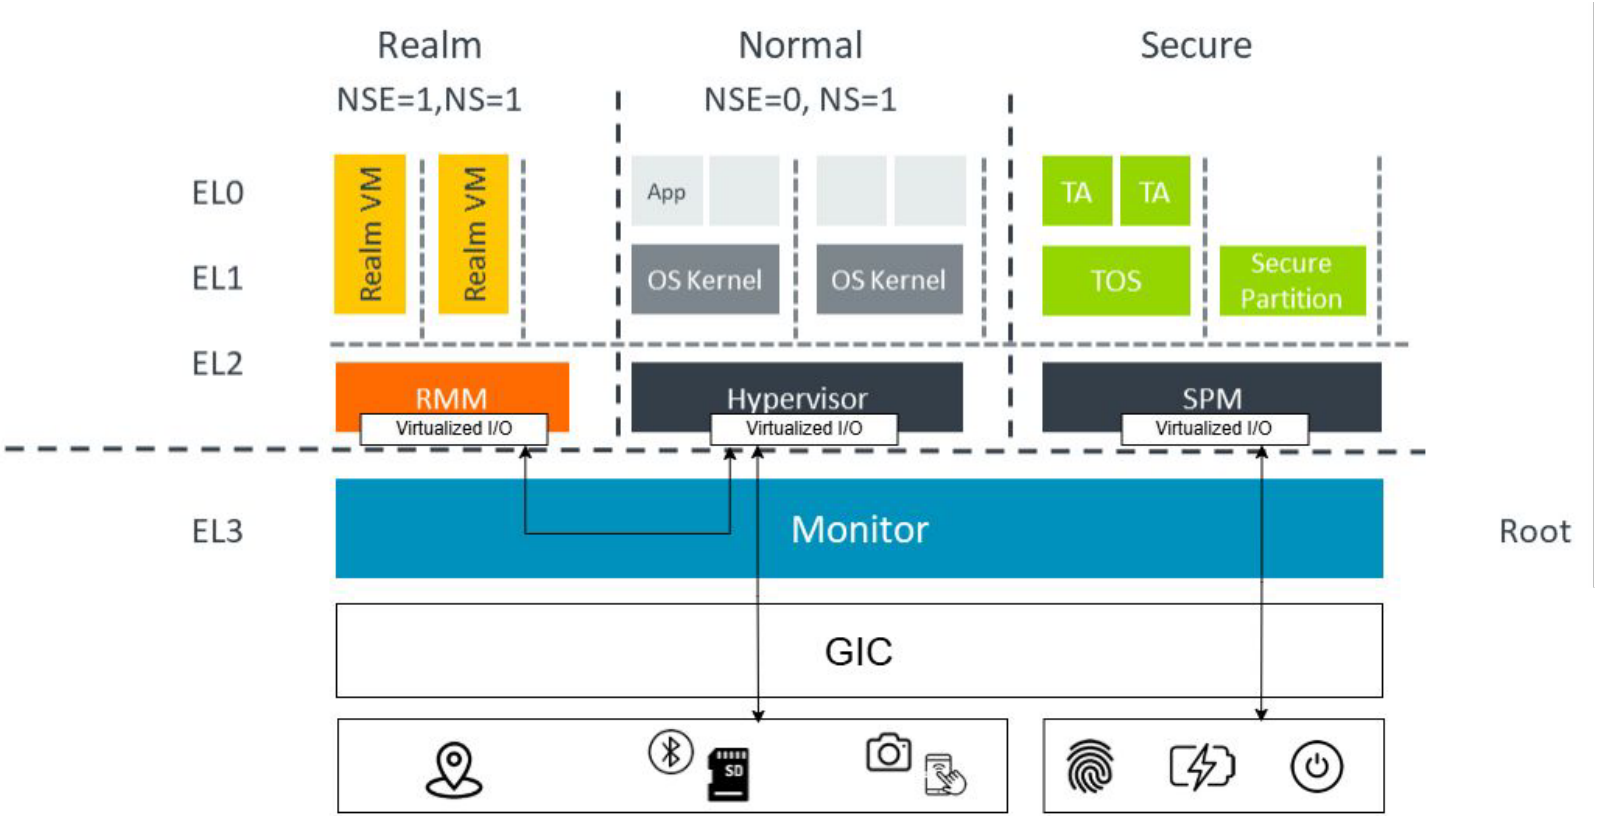

https://developer.arm.com/documentation/den0125/0300/Arm-CCA-Hardware-Architecture

This work

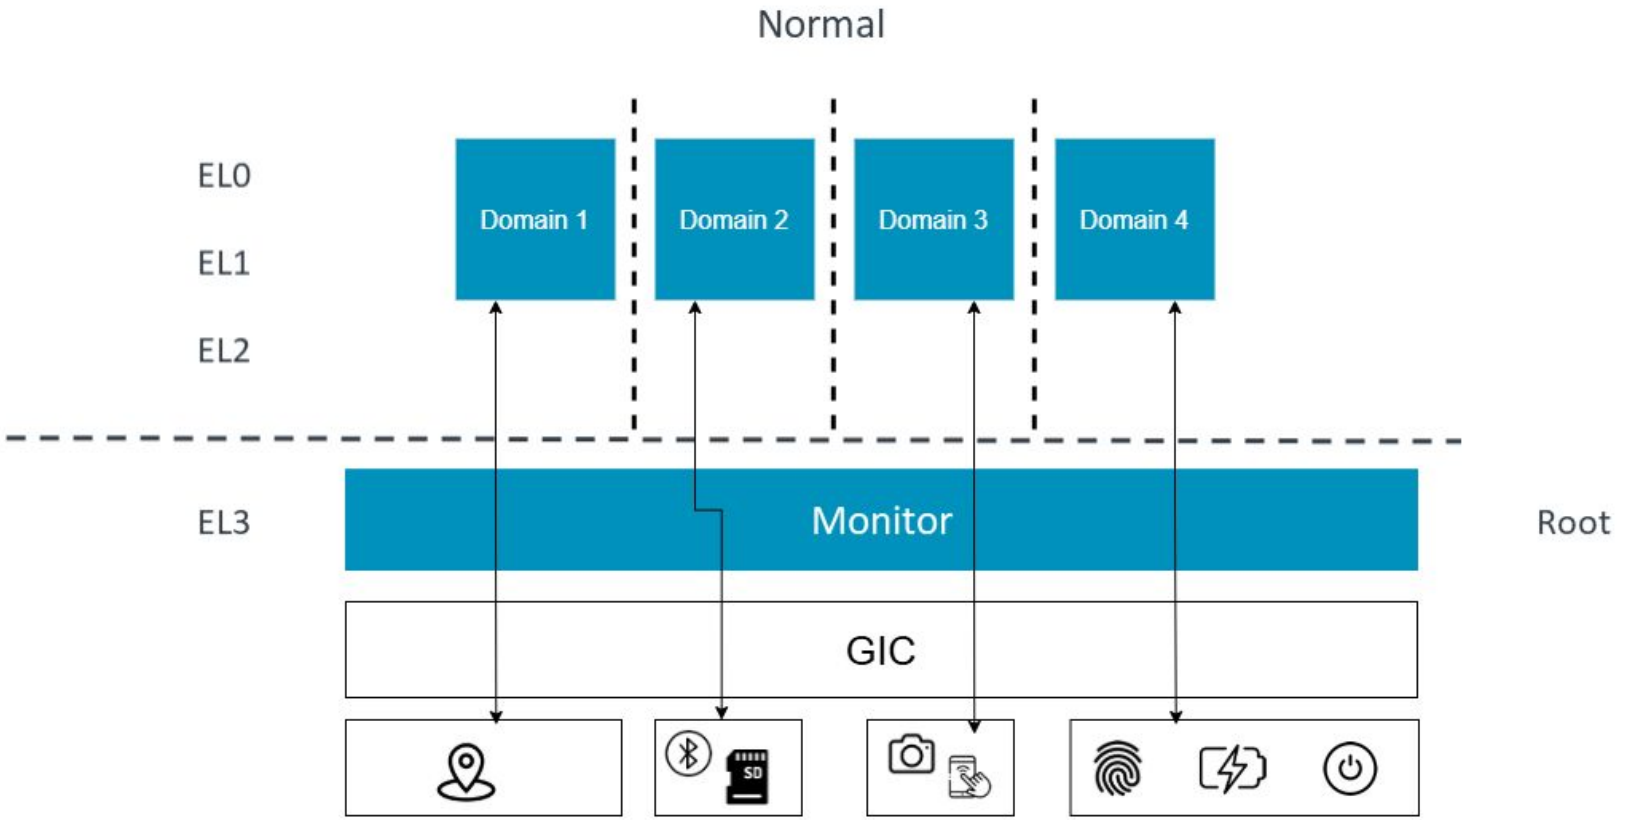
Normal
EL0
EL1
EL2
EL3
Domain 1
Domain 2
Domain 3
Domain 4
Monitor
Root
GIC
SD

Design - **Execution Isolation & Overview**

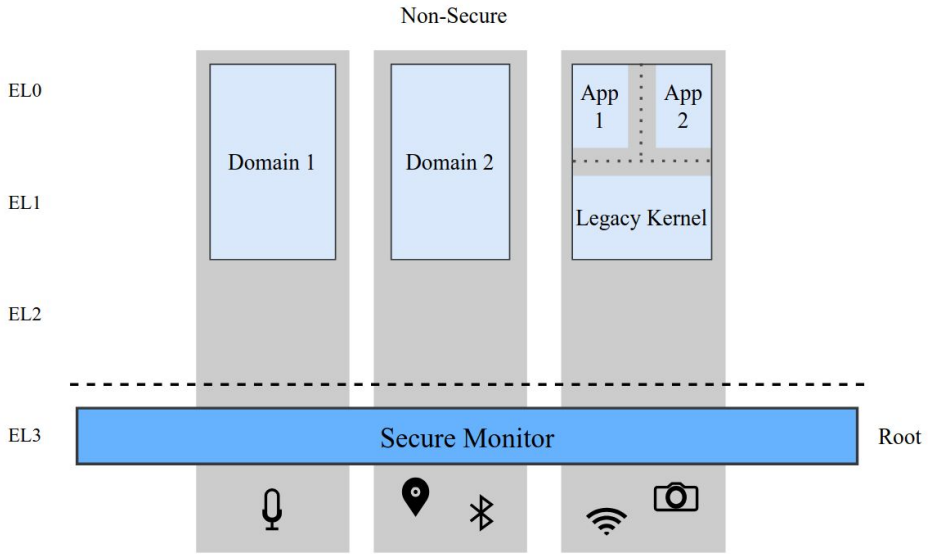

- Secure Monitor ⇒ security operations
- Legacy OS ⇒ non-security: scheduling (in spatial)
- 3 stages via calls to Monitor
  1. Setup
  2. Run
  3. Yield
- Insight from TEEtime

Design - **Memory Isolation**

- Leveraging GPTs.
- Per-domain GPT design, insight from SHELTER.

Design - Memory Isolation - **Per-domain GPTs**

- Monitor maintains 1 GPT per domain

  ⇒ Each domain has its own mapping of **[**memory region ⇛ security state**]**.

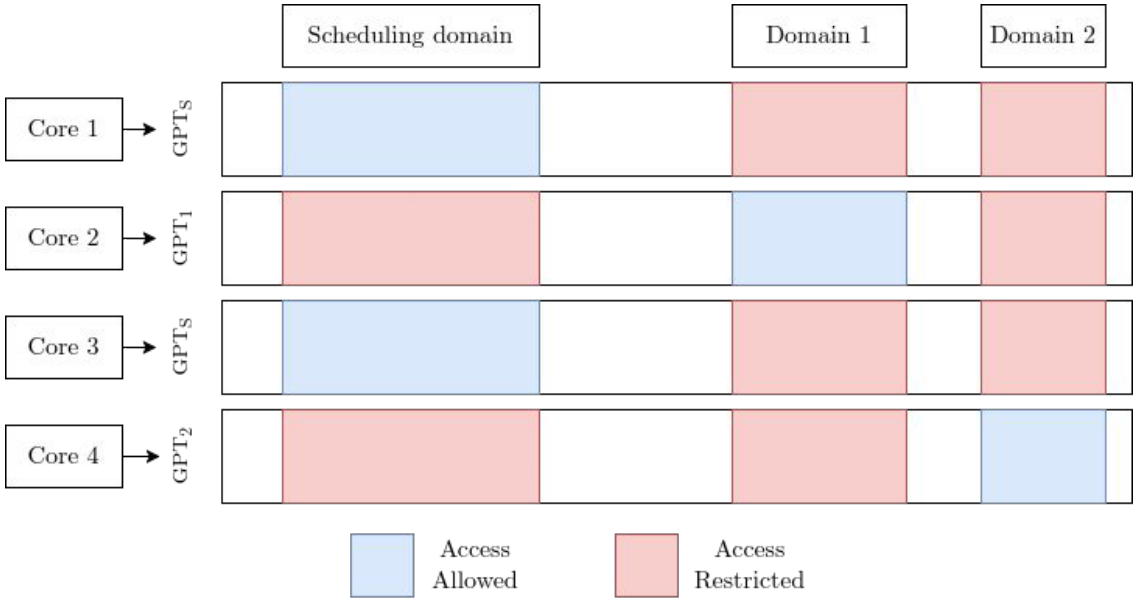

Design - **Memory Isolation**

- Leveraging GPTs.
- New per-domain GPT design.
- Assignment of memory regions to domains.
- Access control enforced by GPC.
- Core assigned a domain ⇒ assigned its GPT.

Design - GPT Management - **At T_0**

Design - GPT Management - **Allocate Domain Region**

Design - GPT Management - **Restrict Access to New Domain**

Design - GPT Management - **Swap GPT of Scheduled Core**

Design - **Interrupt Isolation**

- GIC memory marked as Root ⇒ configuration only by Monitor.
- GPTs for MMIO devices.
- Affinity to route interrupts.
- Insight from TEEtime and DevLore.
- (DMA for future work)

Design - Interrupt Isolation - **Affinity**

- Logical address of cores – hierarchical format

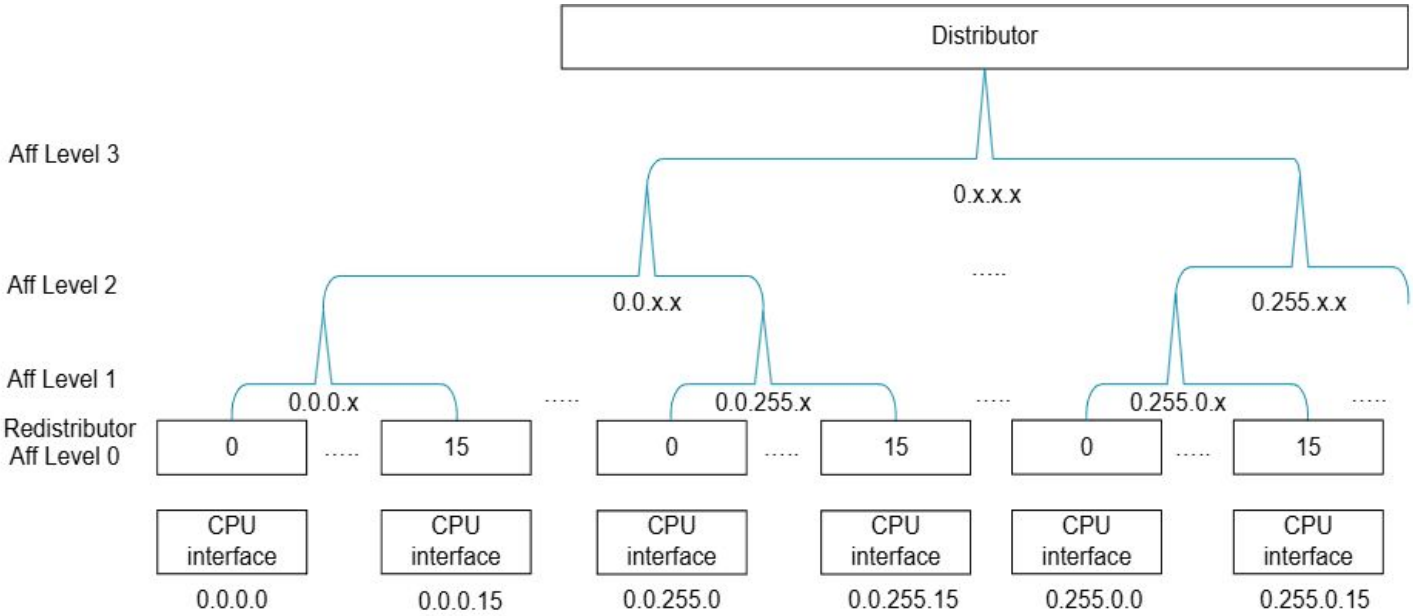

https://developer.arm.com/documentation/101206/0003/Operation/Affinity-routing-and-assignment

Design - Interrupt Isolation - **Affinity**

- Allows routing interrupts to specific cores

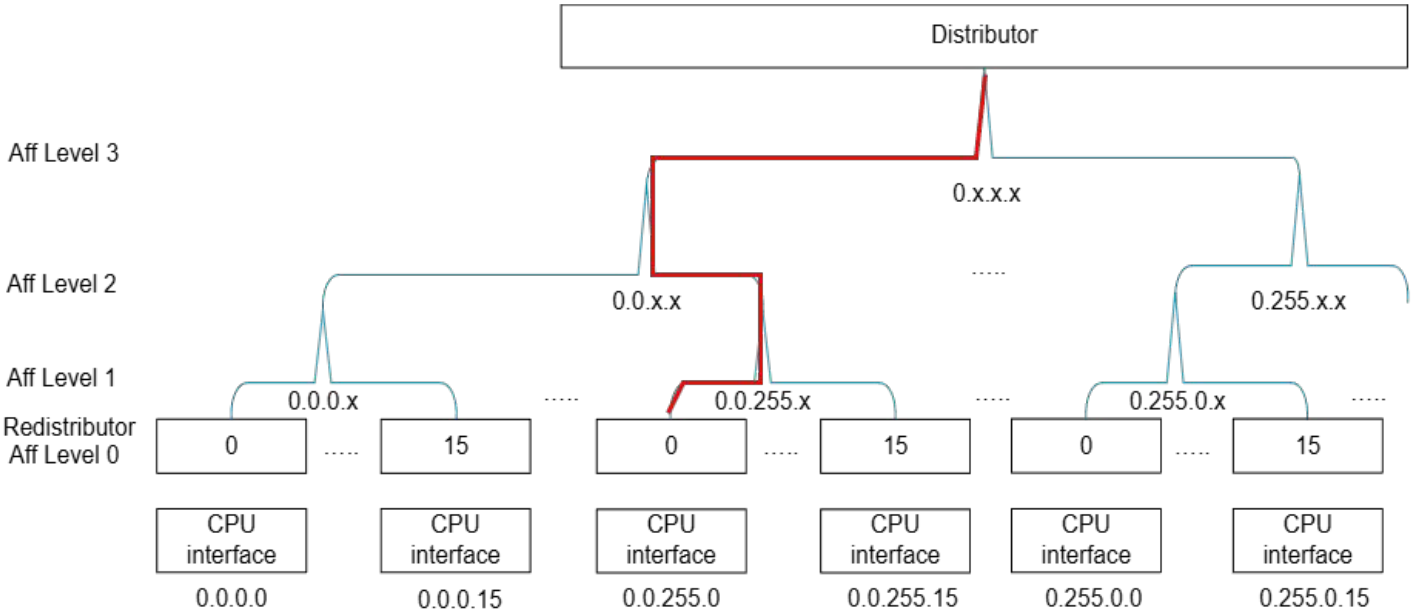

https://developer.arm.com/documentation/101206/0003/Operation/Affinity-routing-and-assignment

Implementation - **Setup**

- Trusted Firmware-A v2.9 **+ 1.7k SLoC** ⇒ Firmware & Secure Monitor implementation
- Functional prototype implemented on an RME-enabled Arm FVP
- Linux v6.5 as the scheduling domain

Implementation - Some Features - **For Memory Isolation**

- Size of SRAM increased for L0 GPTs
- Extension of GPT library for multi-GPT support
- Initialization of GPTs during boot stages
- Correct GPT applied on core’s warm reset
- Multi-threaded synchronization primitives
- TLB invalidation when GPT swapped or modified
- Ensure no sharing of cached GPT entries in TLB across cores

Implementation - Some Features - **For Interrupt Isolation**

- Ensure no overlapping in peripheral assignments
- GIC moved to Root world

  ⇒ GIC updates in Legacy OS hooked to Monitor
- Set affinity of interrupts for binding of routing

Implementation - Some Features - **For Loading Binaries**

- User-level loader
- Kernel module for SMC

Results - **Security Evaluation**

Summary of attacks

Design	Attack	Successfully mitigated?
TEEtime	Privilege escalation to EL3 (Monitor)	✔
SHELTER	Iago attacks	✔
DevLore	Interrupt starvation	✔

Results - Security Evaluation - **Privilege escalation in TEEtime**

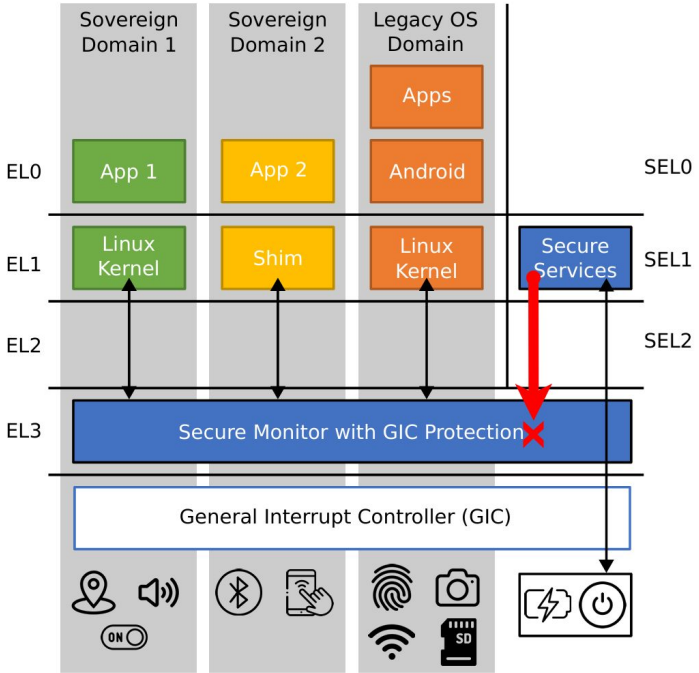

F. Groschupp, M. Kuhne, M. Schneider, I. Puddu, S. Shinde, and S. Capkun, I"t's TEEtime: A new architecture bringing sovereignty to smartphones", 2023.

- EL3 cache-based code injection from SEL1[1]
  ⇒ Secure Monitor compromised
- CCA: Monitor in Root world ✔
- Attack mitigated & secure services out of TCB

1. D. Cerdeira, J. Martins, N. Santos, and S. Pinto, "ReZone: Disarming TrustZone with TEE privilege reduction", 2022

Results - Security Evaluation - **Iago attacks in SHELTER**

- SHELTER mitigates memory-based Iago attacks
- No checks against syscalls not related to SApp memory (`getpid()`, `time()`) [1]
- Many Iago attacks still possible, e.g., connection-replay Iago attack [2,3]
- This work: no syscalls to host (legacy) OS ✔

1. Y. Zhang et al., “SHELTER: Extending arm CCA with isolation in user space”, 2023
2. Stephen Checkoway and Hovav Shacham, “Iago attacks: why the system call API is a bad untrusted RPC interface", 2013
3. Thomas Ristenpart and Scott Yilek, “When Good Randomness Goes Bad: Virtual Machine Reset Vulnerabilities and Hedging Deployed Cryptography”, 2003

Results - Security Evaluation - **Guarantee of int. delivery in DevLore**

- Realm VMs trust the RMM to let interrupts pass
- Compromised RMM can starve VMs
- This work: domains have direct physical access to peripherals ✔

Results - **Performance Evaluation**

Operation	Time (μs)
Setup	
Setup clean state	83
Grant peripheral access	264
GPT transition	165
Core scheduling	258
Yield	
Memory & cache cleanup	54
GPT transition + swap	173
Withdraw peripheral access	175

Breakdown of life cycle operations for a small program

Results - **Performance Evaluation**

Operation	Time (μs)
Setup	
Setup clean state	83
Grant peripheral access	264
GPT transition	165
Core scheduling	258
Yield	
Memory & cache cleanup	54
GPT transition + swap	173
Withdraw peripheral access	175

Breakdown of life cycle operations for a small program

- What if the domain is larger?

Results - **Performance Evaluation**

Operation	Time (μs)
Setup	
Setup clean state	83
Grant peripheral access	264
GPT transition	165
Core scheduling	258
Yield	
Memory & cache cleanup	54
GPT transition + swap	173
Withdraw peripheral access	175

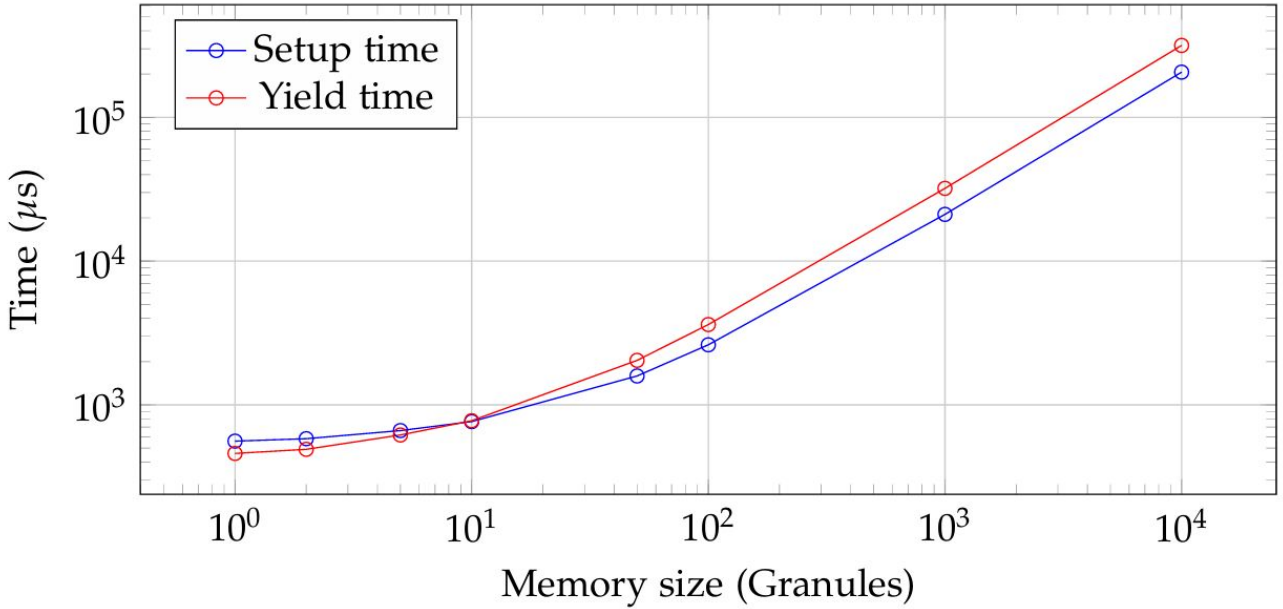

Setup and yield time with different domain sizes

- Area of improvement: granule transition

Summary of contributions

- Multi-GPT design

- Normal world enclaves with (small TCB and) isolation across
  - Execution
  - Memory
  - Peripheral interaction (no DMA)

- Small TCB: Hardware, Firmware, Secure Monitor

Thank you for listening!