HARDWARE ENFORCED SEPARATION IN EMBEDDED MULTICORE SOCS

GEOFF WATERS
SECURITY ARCHITECT
DIGITAL NETWORKING

NXP | SECURE CONNECTIONS FOR A SMARTER WORLD

Agenda

1. Types of Separation
2. Early adopters
3. Hardware enforcement mechanisms
4. Tamper proofing

Unsupervised Asymmetric Multiprocessing

- Security — no enforced isolation, cannot allow untrusted operating systems
- Requires cooperation among partitions
- How are global hardware resources managed?
  - Local access windows
  - Interrupt controller
  - Shared caches
  - IOMMU
- Boot sequence complexity
- Error management
- Resetting/rebooting partitions
- Debugging

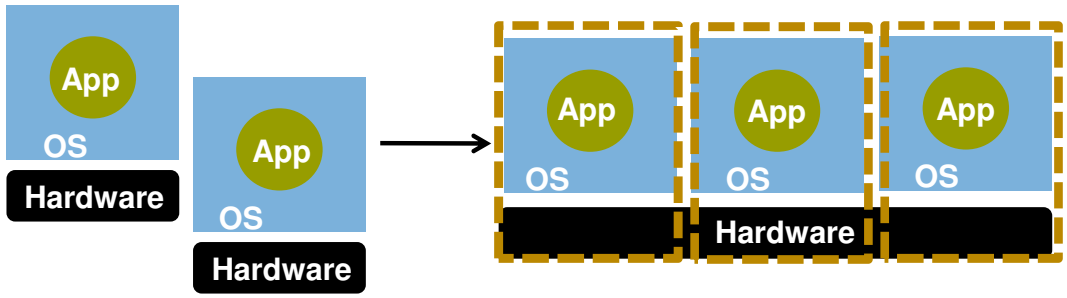

Partitioning with a Supervisor

- Supervisor: Layer of software more privileged than operating systems
- Provides:
  - Enforcement of system security, partition boundaries
  - Global resource management (e.g. interrupt controller)
  - Resource sharing and virtualization – CPUs, memory, I/O devices
  - Other services (e.g. debug)

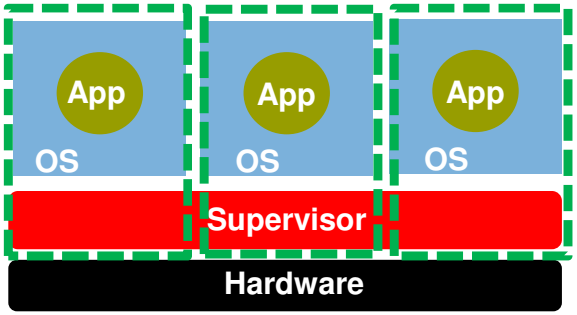

Operating Systems / Hypervisors / Privilege Levels

- Simple executive or an RTOS may use only one privilege level
- Traditional OS uses 2 privilege levels (kernel & user) to separate applications and OS kernel
- A hypervisor conceptually adds a new privilege level

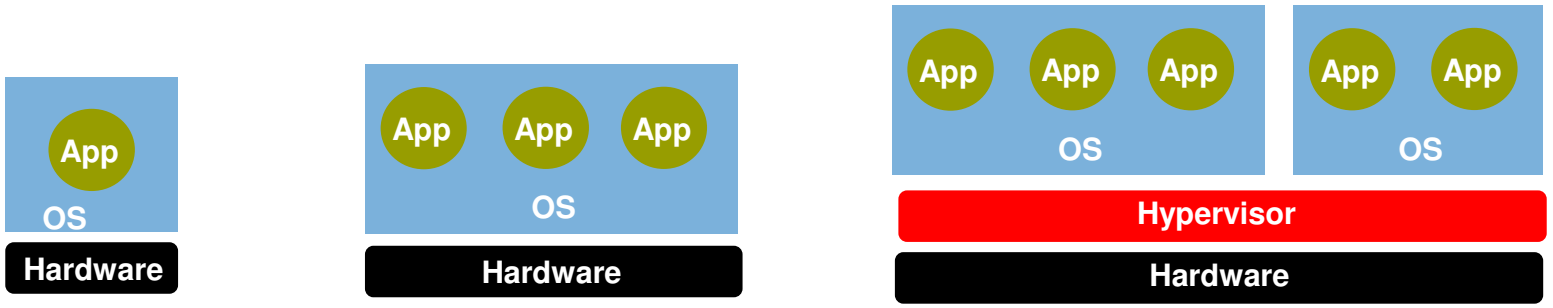

Motivations for Virtualization

- Efficiency: Consolidation onto fewer processors for higher hardware utilization
  - Oversubscription tolerated
- Ease of management
  - Create/destroy virtual machines as needed
  - Migrate running VM to different system
- Flexibility
  - Use different versions of the Linux kernel
  - Run legacy software or OS on HV
- Sandboxing– allows untrusted software to be added to a system (e.g. operator applications)

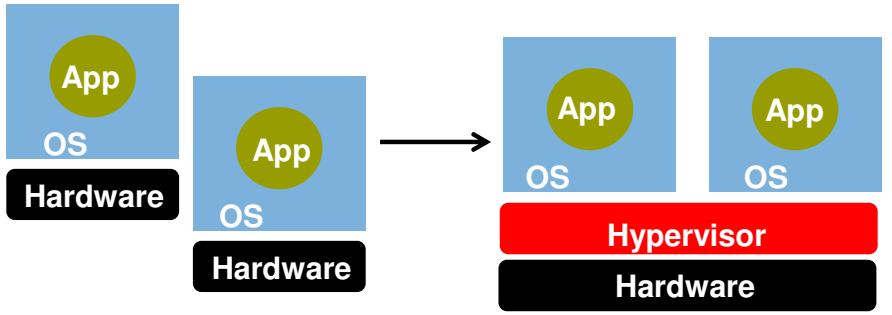

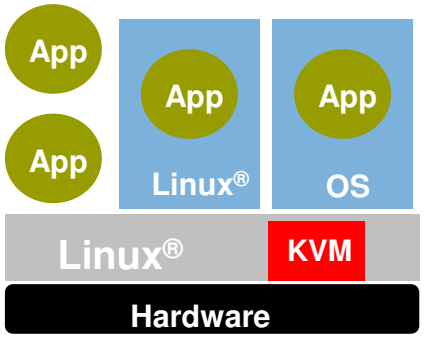

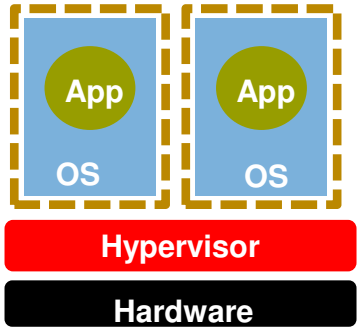

Virtual machines in sandboxes

Motivations for Separation

- Safety & Security
  - Sandboxing on steroids
  - Partition A cannot access Partition B's private resources
  - No communication between partitions except via explicit 1 way communications channels
  - Logical isolation for determinism
    - Partition A cannot effect the execution time of Partition B beyond established limits
- Migration — move to multicore, preserve investment in software
  - Run legacy software alongside new software
- In-service upgrade
- Efficiency: Consolidation onto fewer processors for higher hardware utilization

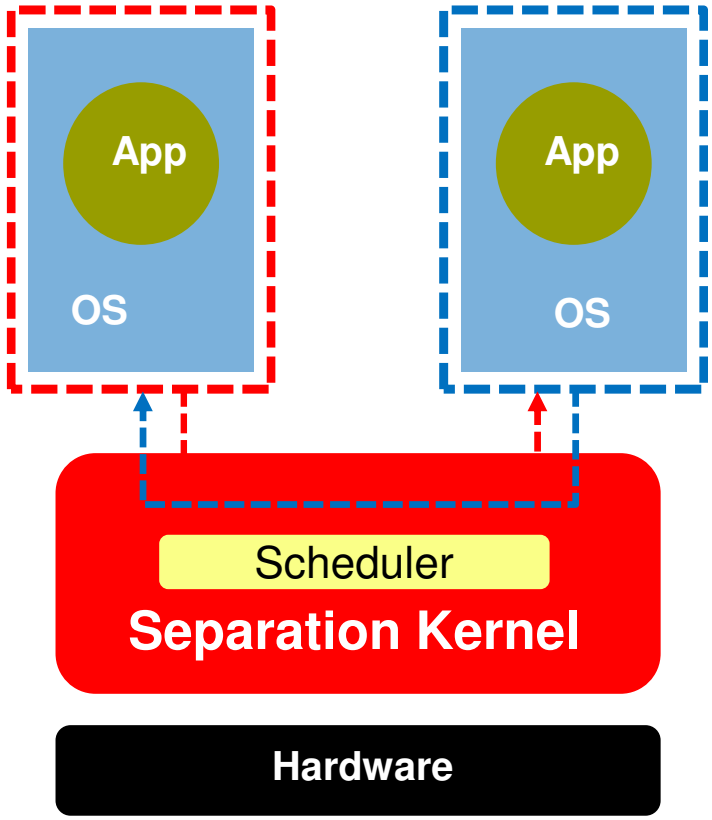

Agenda

1. Types of Separation
2. Early adopters
3. Hardware enforcement mechanisms
4. Tamper proofing

Early Adopters and Next Wave

Aerospace

Main Flight Control, Secondary Flight Control, Aircraft Engine Management, Cockpit Display

Military and Defense

Rocket navigation, Artillery Control Computer, IFF

UAV Flight Computer, Defense Airborne Computer, Weapon Navigation System, Ground Control System

Factory Automation

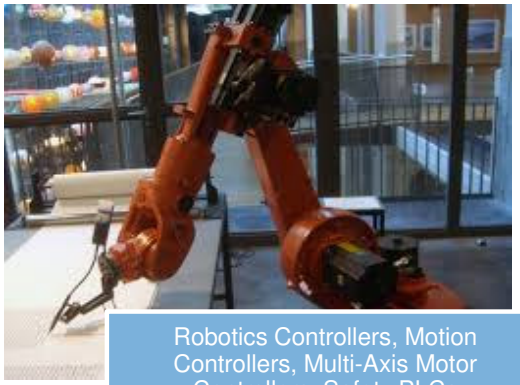

Robotics Controllers, Motion Controllers, Multi-Axis Motor Controllers, Safety PLCs

Railway

Traction Control, Railway Signaling Controller, Railway Communications, Brake Controller

Power Grid

Power Distribution Relays, Smart Grid Communications

Agenda

1. Types of Separation
2. Early adopters
3. Hardware enforcement mechanisms
   A. CPU Extensions & MMU
   B. IO MMU
   C. ‘Devices’
4. Tamper proofing

DPAA 1.x Strong Partitioning/Secure Virtualization

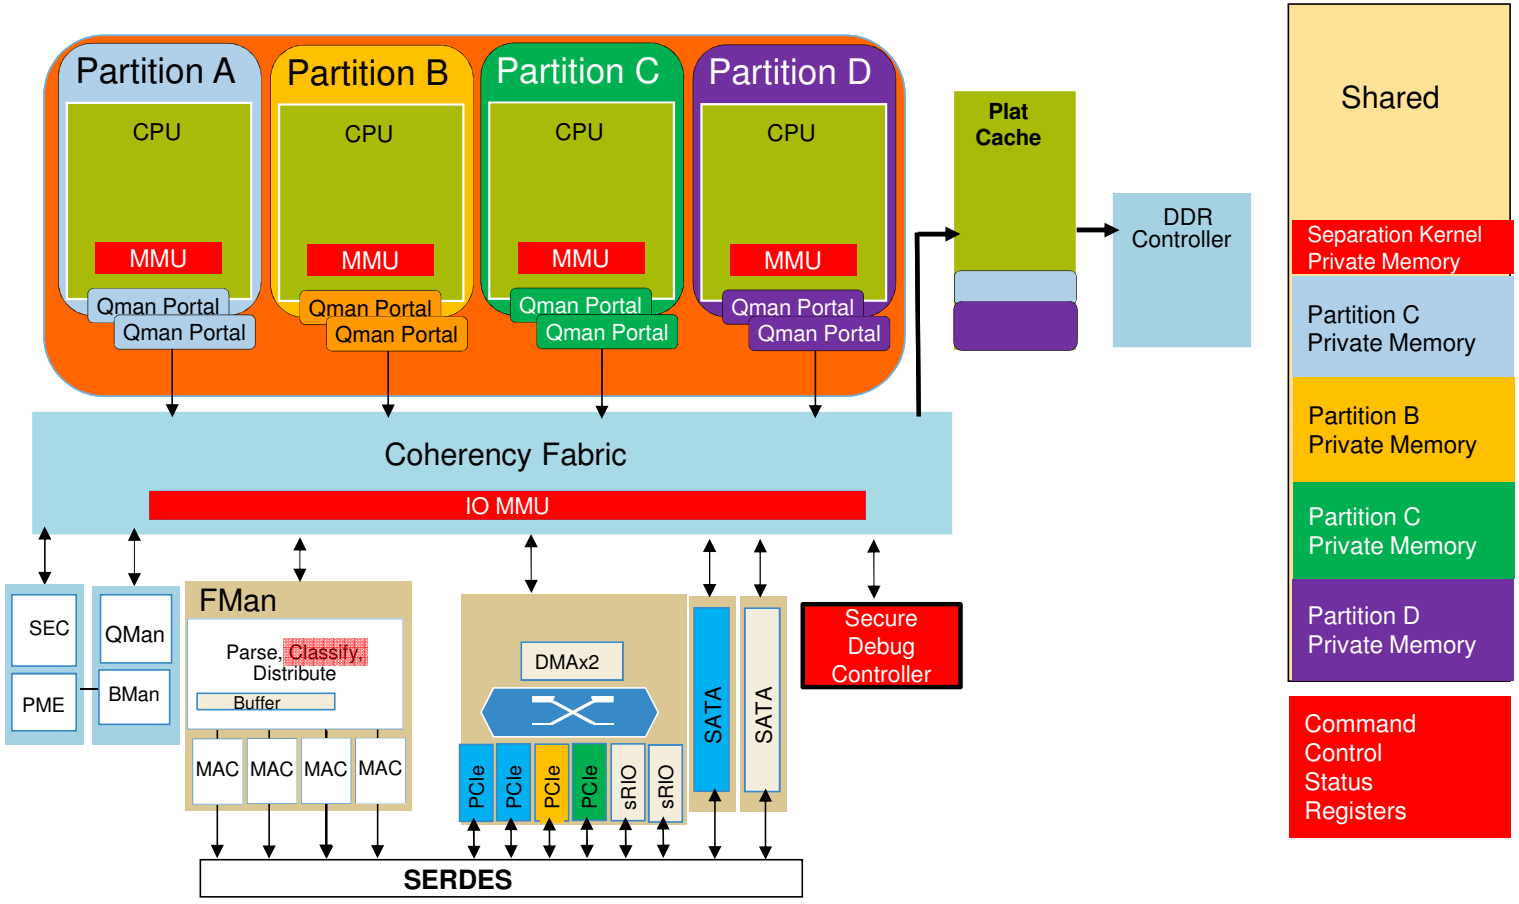

DPAA 2.0 Strong Partitioning/Secure Virtualization

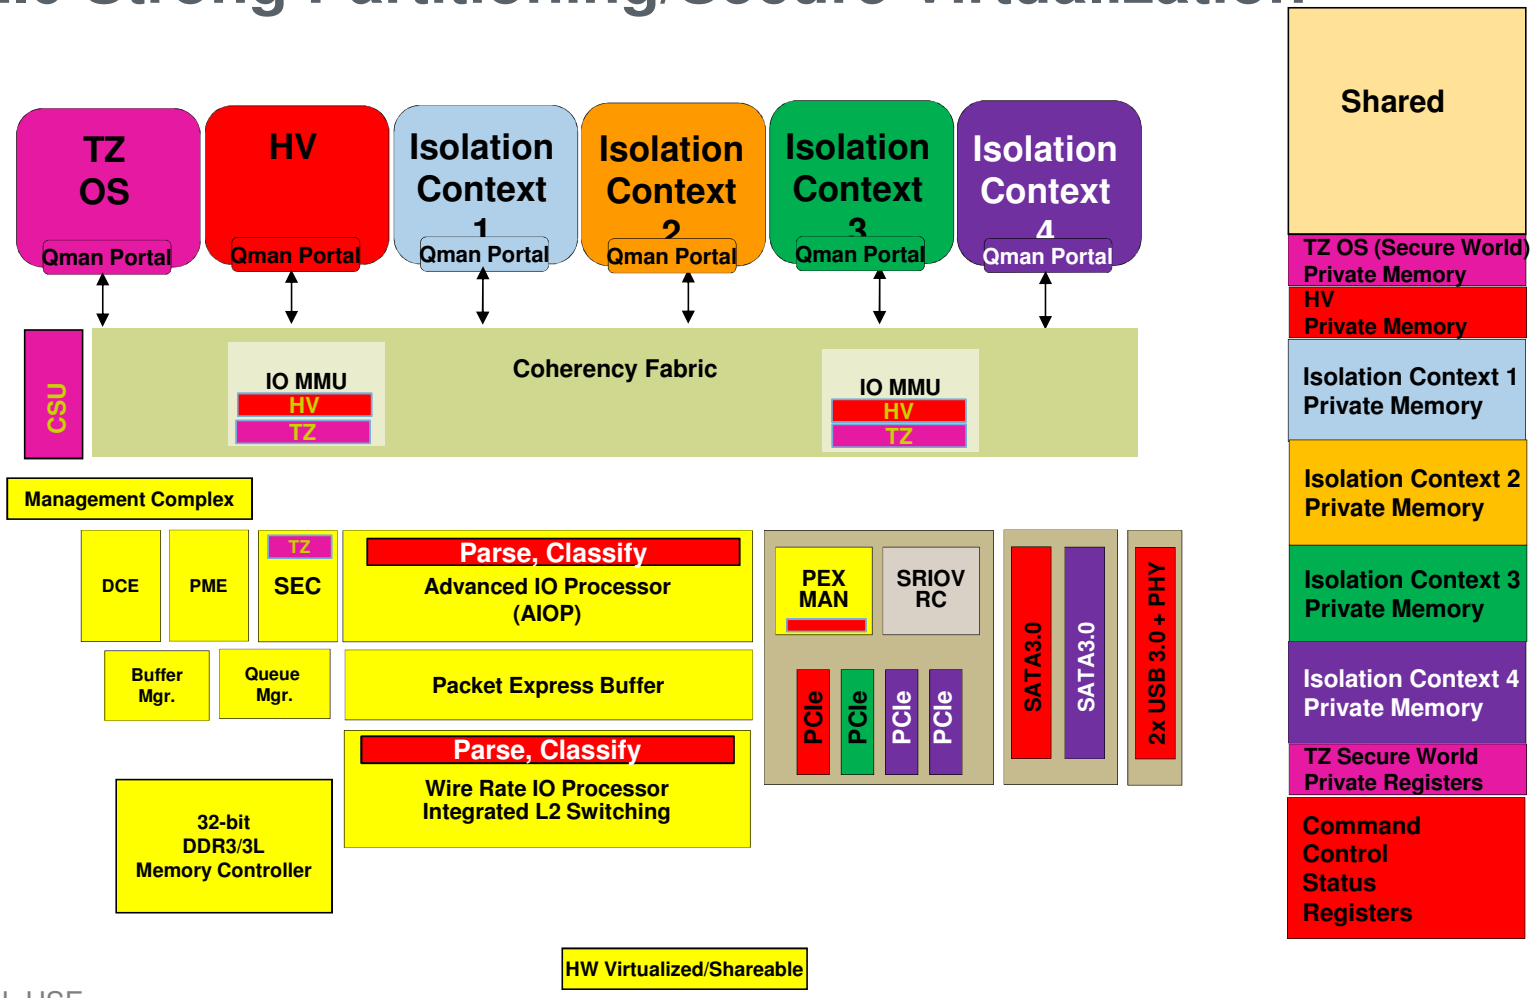

Comparison of Processor Virtualization Capabilities

ARM, Power, x86 architectures all support similar mechanisms to support virtualization.

Category	Feature	QorIQ e500mc, e5500	QorIQ e6500	LS-A Cortex A53/A57	X86
Privilege	3rd privilege level	Yes		Yes	Yes
	Direct register access	Yes		*	
	Direct system calls	Yes		Yes	
MMU	Domain separation	Yes		Yes	Yes
	Extended Address space	Yes		Yes	Yes
	Hardware guest physical address translation	No*	Yes (LRAT)	Yes	Yes (EPT/NPT)
Interrupts	Direct guest interrupt management	Yes		Yes	Yes (x2 APIC)

Privilege Levels with Hypervisors

- Power Architecture virtualization and ARMv8 virtualization extensions add new privilege levels / processor modes

Power Book ‘E’
(e500v2)

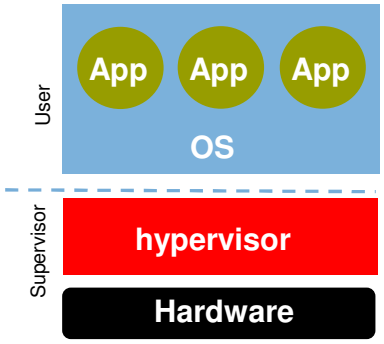

Power Book ‘E’
Virtualization
(e500-mc)

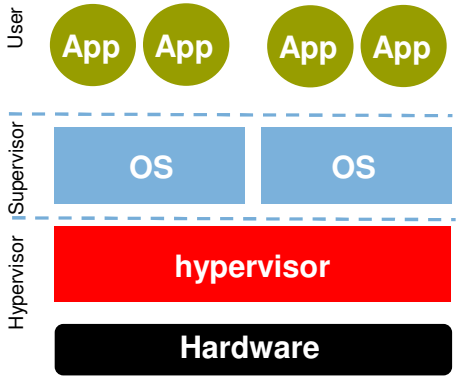

ARMv8
(Cortex A53/A57)

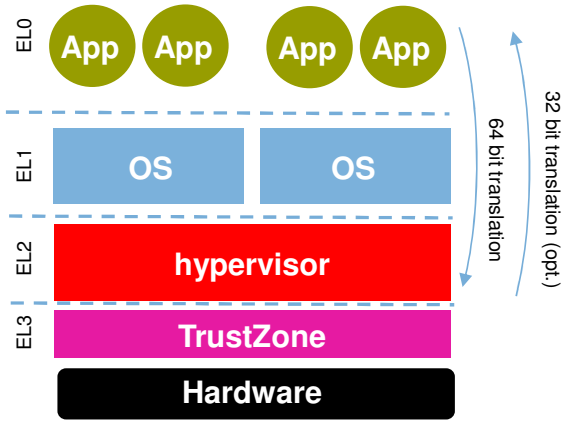

ARM Execution Levels

In ARMv8 architecture, the higher the level number, the higher the privilege level. Higher privilege levels have more direct access hardware resources.

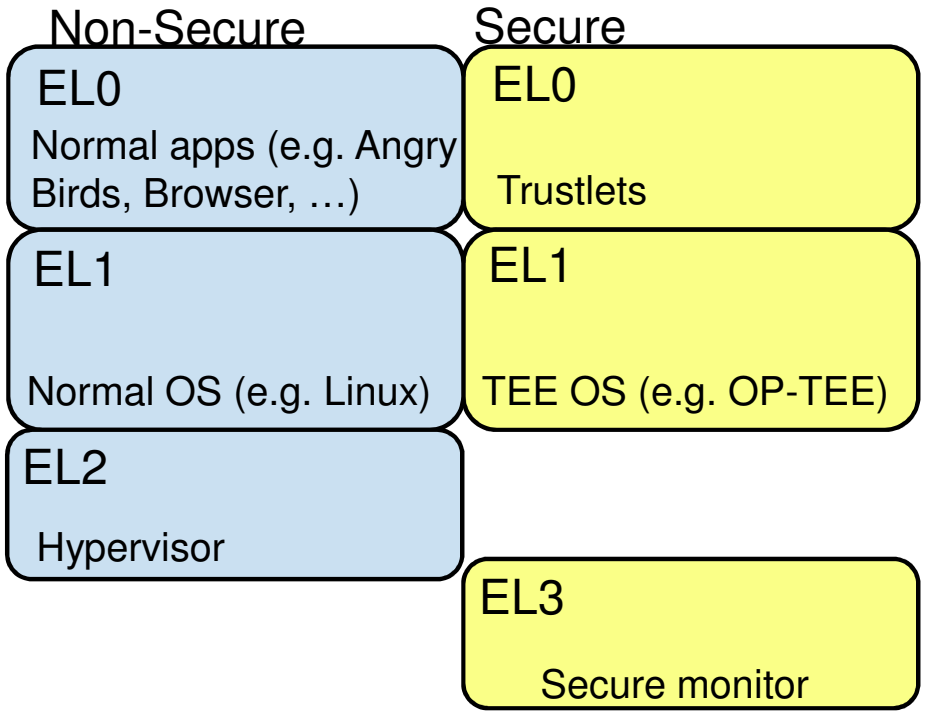

Memory Management Unit (MMU)

- MMUs translate virtual addresses into a physical address which are put onto the system bus
- Older/simpler CPUs may have only a single translation stage
  - Virtual Address (VA) → Physical Address (PA)
- Newer CPUs (including Power e6500 & ARMv8) offer two stage address translation
  - Virtual address (VA) -> Intermediate physical (IPA)
  - Intermediate (IPA) -> Physical (PA)
- Important concepts; Process ID, Page Table, Translation Lookaside Buffer
  - The process running on the CPU is identified by Process ID (PID) Regis (updated by privileged software each time it schedules that process to ru
  - The process can't spoof its PID
  - PID is fed to MMU along with virtual address; MMU accesses page table to that process
  - Translation Lookaside Buffer (TLB)
    - A PID aware cache of the page table entries of recently translated adc
  - Page table is data structure containing mapping from VA→ PA
    - Also contains access permissions for the page (see example from AR
    - Also PID aware
- This is a general description, different CPUs have different levels of hierarc page tables, different content in the MMU's 'cache'

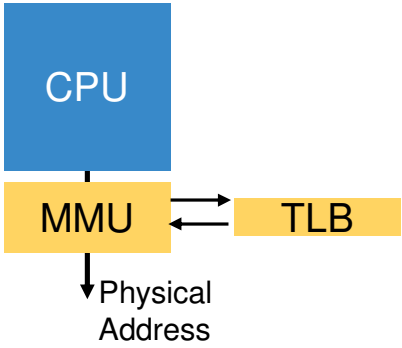

Table D4-33 Access permissions for instruction execution for stage 1 of the EL1&0 translation regime

UXN	PXN	AP[2:1]	SCTLR_EL1.WXN	Access from EL1	Access from EL0
0	0	00	0	R, W, Executable	Executable
			1	R, W, Not executable[a]	Executable
		01	0	R, W, Not executable[b]	R, W, Executable
			1	R, W, Not executable	R, W, Not executable[c]
		10	x	R, Executable	Executable
		11	x	R, Executable	R, Executable
0	1	00	x	R, W, Not executable	Executable
		01	0	R, W, Not executable	R, W, Executable
			1	R, W, Not executable	R, W, Not executable[c]
		10	x	R, Not executable	Executable
		11	x	R, Not executable	R, Executable

QorIQ Virtualized Interrupts

- QorIQ Power “Book E” Embedded Virtualization
  - Allows a number of interrupts to be vectored directly to guest (GIVOR)
  - Others trap to the hypervisor

- QorIQ Layerscape ARMv8
  - Route to one of:
    - Guest OS (current or different)
    - Hypervisor
    - TrustZone
  - Basic model – all interrupts virtualized:
    - Physical go to Hypervisor initially
    - Interrupts directed to guests:
      - Hypervisor maps “virtual” interrupt for that Guest OS

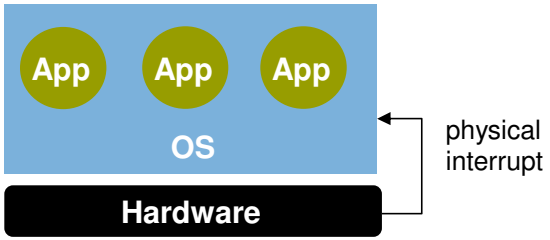

System without Virtualization

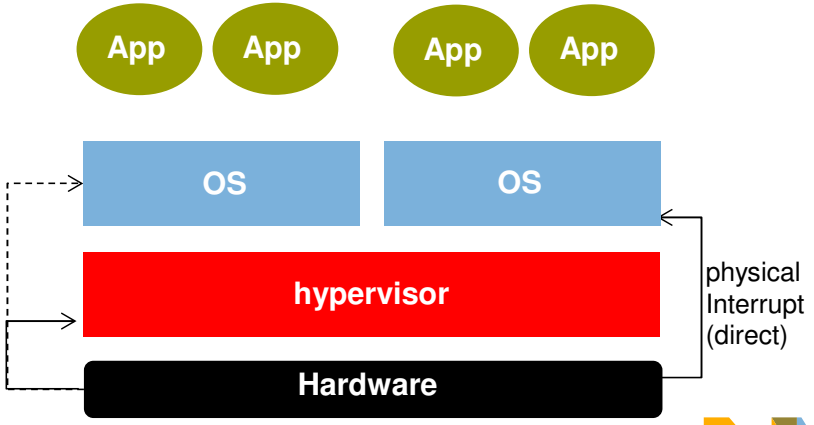

System with Virtualization

Agenda

1. Types of Separation
2. Early adopters
3. Hardware enforcement mechanisms
   A. CPU Extensions & MMU
   B. IO MMU
   C. ‘Devices’
4. Tamper proofing

Comparison of Processor Virtualization Capabilities

ARM, Power, x86 architectures all support similar mechanisms to support virtualization.

Category	Feature	QorIQ e500mc, e5500	QorIQ e6500	LS-A Cortex A53/A57	X86
IOMMU		Yes (PAMU)		Yes (SMMU)	Yes (VT-d)

- IO MMUs exist principally for convenience
  - Allow the guest OS to use unmodified device drivers
  - OS will program descriptors with Intermediate Physical Address, IO MMU with translate to Physical Address
  - Note; Applications using user space device drivers will program descriptors with VA, requiring 2 stage translation
  - QorIQ PAMU offers single stage translation, ARMv8 SMMU supports 2-stage translation
- Like MMUs, IO MMUs can include access permissions look-up in the translation
  - If partition A is blocked from directly accessing partition B's memory by the MMU, it could try programming a hardware block with DMA capability to access partition B's memory on partition A's behalf.
  - A properly configured IO MMU will block this

Agenda

1. Types of Separation
2. Early adopters
3. Hardware enforcement mechanisms
   A. CPU Extensions & MMU
   B. IO MMU
   C. ‘Devices’
4. Tamper proofing

Datapath Acceleration Architecture

- Datapath Infrastructure
  - Queue Manager
    - Work Distribution, congestion control
    - Order preservation and restoration
    - Prioritization, shaping
  - Buffer Manager
- ‘Processors’
  - ARM or Power ISA GPPs
  - Other programmable Engines
  - Network interfaces (WRIOP)
    - Parse, Classify, Police
- Accelerators
  - Crypto Acceleration (SEC)
  - Pattern Matching Acceleration
  - Data Compression Engine

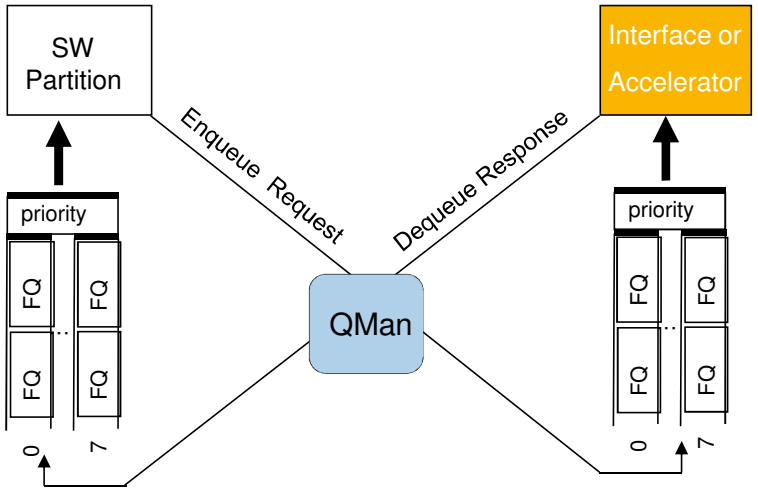

Datapath Acceleration Architecture: Software Portals & Isolation

- All QorIQ Layerscape datapath resources are accessed through software portals
- Portals can be put in an isolated mode where DPAA resource IDs are virtual
- A resource permissions table maps virtual ID to real ID

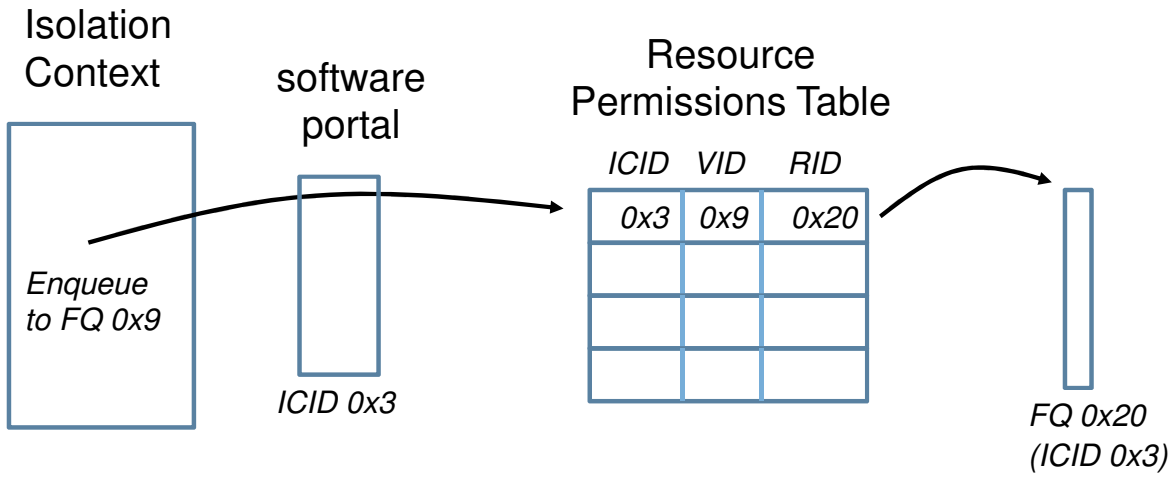

- If table lookup fails, it's an access violation

DPAA 2.0 Separation Mechanics

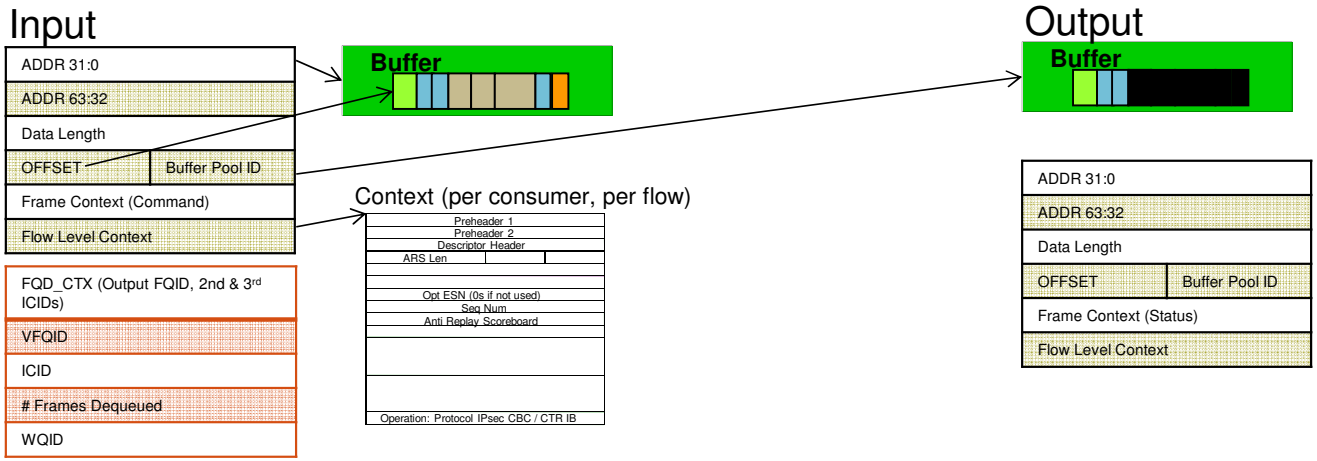

- Bman filters buffer release with ICID
- IO MMU filters buffer and context reads with ICIDs

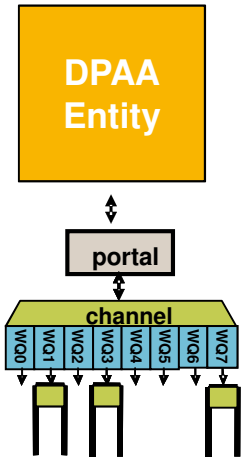

- Bman filters buffer request with ICID
  - Output BPID in FLC
- QMan filters enqueue with ICID
  - Output FQID in Dequeue Summary Info

Agenda

1. Types of Separation
2. Early adopters
3. Hardware enforcement mechanisms
4. Tamper proofing

Generic Trust Architecture SoC

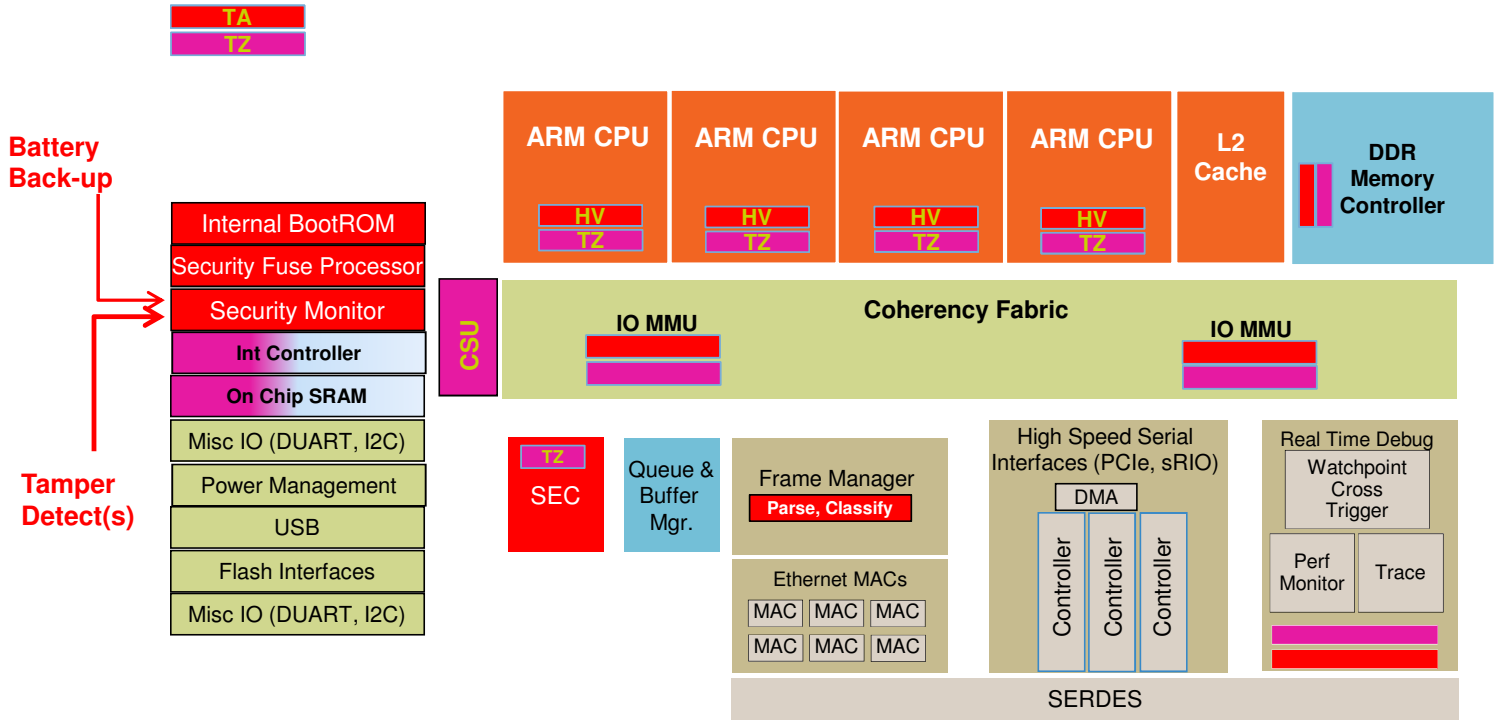

Secure Boot: Verifying Code before Execution

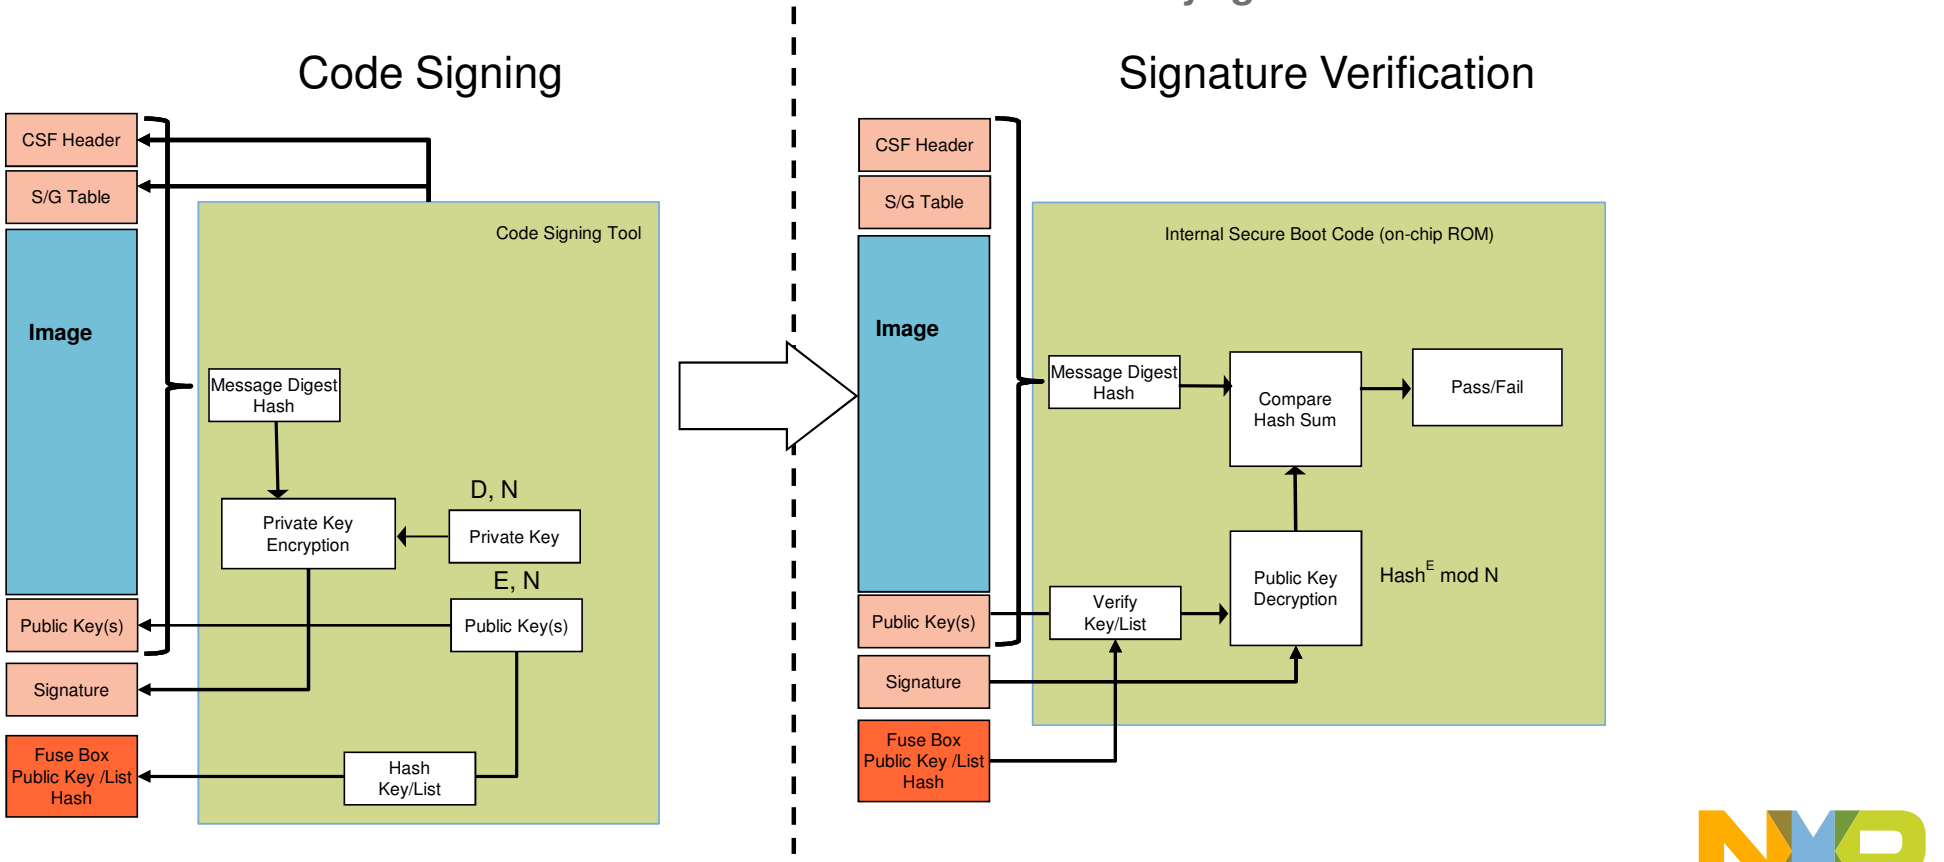

Tamper Detection Sources

- Hardware:
  - External Tamper Detection via TMP_DETECT and LP_TMP_DETECT
  - Secure Debug Controller (if set to Conditionally Closed with Notification)
  - Run Time Integrity Checker (in SEC)
  - Security Fuse Processor (if fuse array read fails, including hamming code check)
  - Security Monitor (OTPMK and ZMK hamming code check)
  - All sensitive flops upon detection of scan entry and exit (expert mode debug)
  - Power Glitch
  - In Trust 2.0:
    - Monotonic counter roll-over
- Software:
  - ISBC (Boot 0)
  - ESBC/Trusted-Uboot (Boot 1)
  - Any SW with write access to the Security Monitor can declare a security violation.

Runtime Integrity Checking

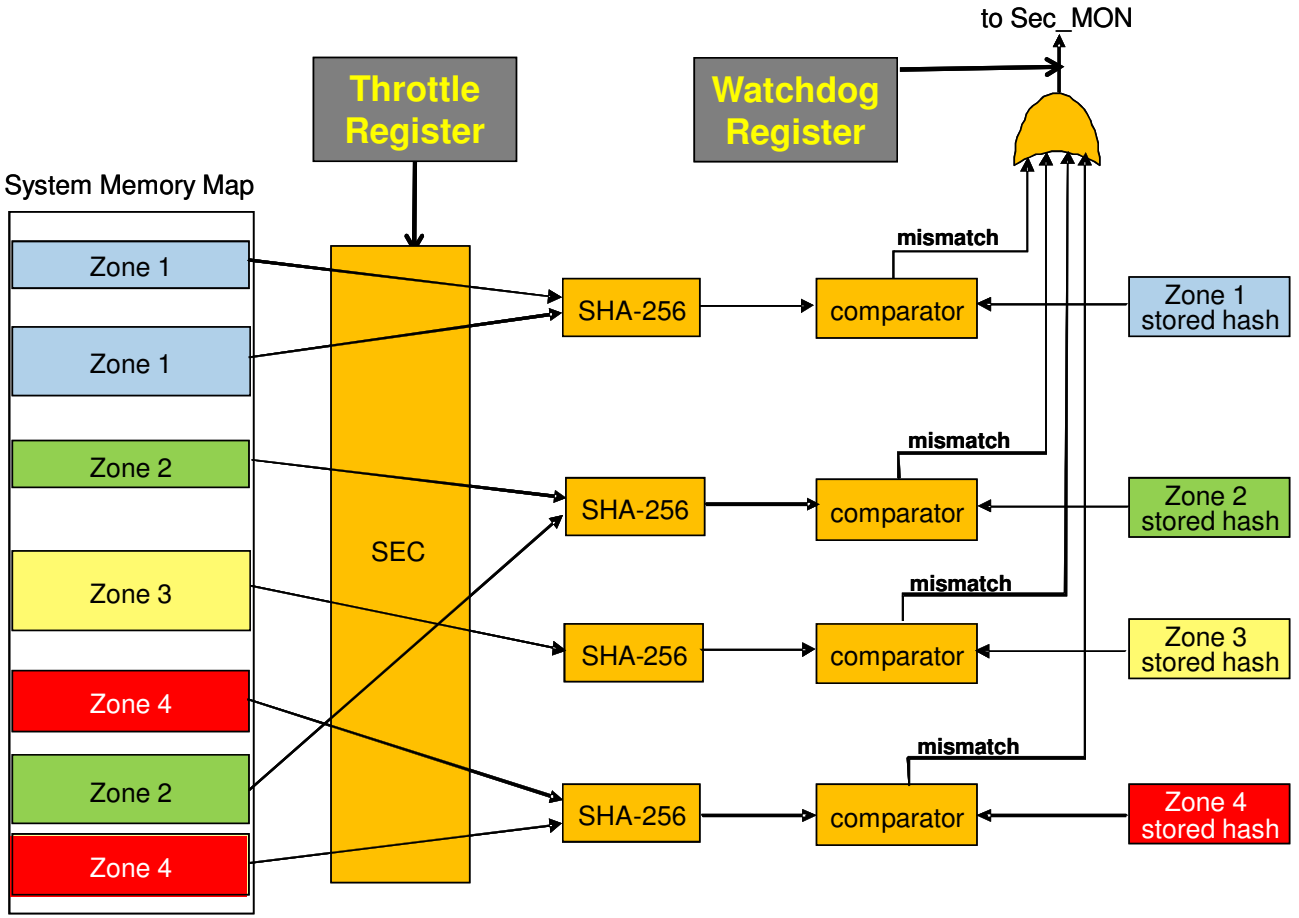

Summary

- The MILS architecture is seeing increasing adoption in other industries dealing with cyber-physical devices.
- Support for safely and securely combining mixed criticality functions and providing a means for compositional security certification should make MILS the architecture of choice for a variety of industrial applications, including
  - power generation & distribution
  - factory & service robotics
  - intelligent transportation networks
- NXP QorIQ Processors are well suited for MILS systems, and we look forward to on-going cooperation with the CertMILS program.

SECURE CONNECTIONS
FOR A SMARTER WORLD