Safety Analysis of NXP High Performance Layerscape Multicore Processors

Geoff Waters

Sr. Principal Engineer
NXP Digital Networking BL

June 2019 | Session #AMF-AUT-T3648

SECURE CONNECTIONS
FOR A SMARTER WORLD

Agenda

- NXP Multicore Processor Families
- Digital Networking Layerscape Products
- Safety Positioning and Preliminary Metrics
- Features Supporting Safety Goals
- Partitioning & Freedom from Interference
- Summary

NXP Automotive Microprocessors & Microcontrollers

BL DN

(Digital Networking)

High Performance Networking & Computing

- Highest networking & compute performance SoCs in NXP
- Experts in Linux, networking protocols, network security, virtualization
- #1 SoC Architecture in Mil/Aero

Products

QorIQ
Layerscape

BL Micros

Multimedia Processing

- HMI, Multimedia, Compute, Image Processing Leader
- GPUs with 1 to 16 Vec4 shaders, 8 to 256 GFLOPS
- With ML Framework
- Power efficiency, battery operation

Products

i.MX

BL AMP

ADAS (Advanced Driver Assistance Systems)	C&S (Connectivity & Security)	VDS (Vehicle Dynamics & Safety)	GPIS (General Purpose & Integrated Solutions)
Radar, LIDAR Vision Sensor Fusion	Gateway	Chassis & Safety Powertrain & Hybrid/EV	Body Electronics Edge Nodes
• #1 in Radar with strong IP and system knowledge • High performance low power accelerators • Scalable high performance roadmap for central processing	• #1 in Vehicle Networking with leading networking and security IP • #1 in Automotive HW Security with Strong IP and broad portfolio • End to end portfolio of networking devices (MCU/MPU, TX/RX)	• Long term Innovator in Chassis and Powertrain Control. • Significant Growth in Safety as Autonomous Control Drives Robust Fault Tolerant Systems	• 500+ customers • Broadest portfolio of integrated MCU+HV mixed-signal solutions • Complete Tools & Software enablement
Products	Products	Products	Products
S32R - Radar S32V - Vision	MPC564xB/C MPX574xG S32	MPC56xx MPC57xx S32S/P/H	S08/S12/PPC → ARM KEA – S32K S12 MagniV – S32M

Layerscape Portfolio

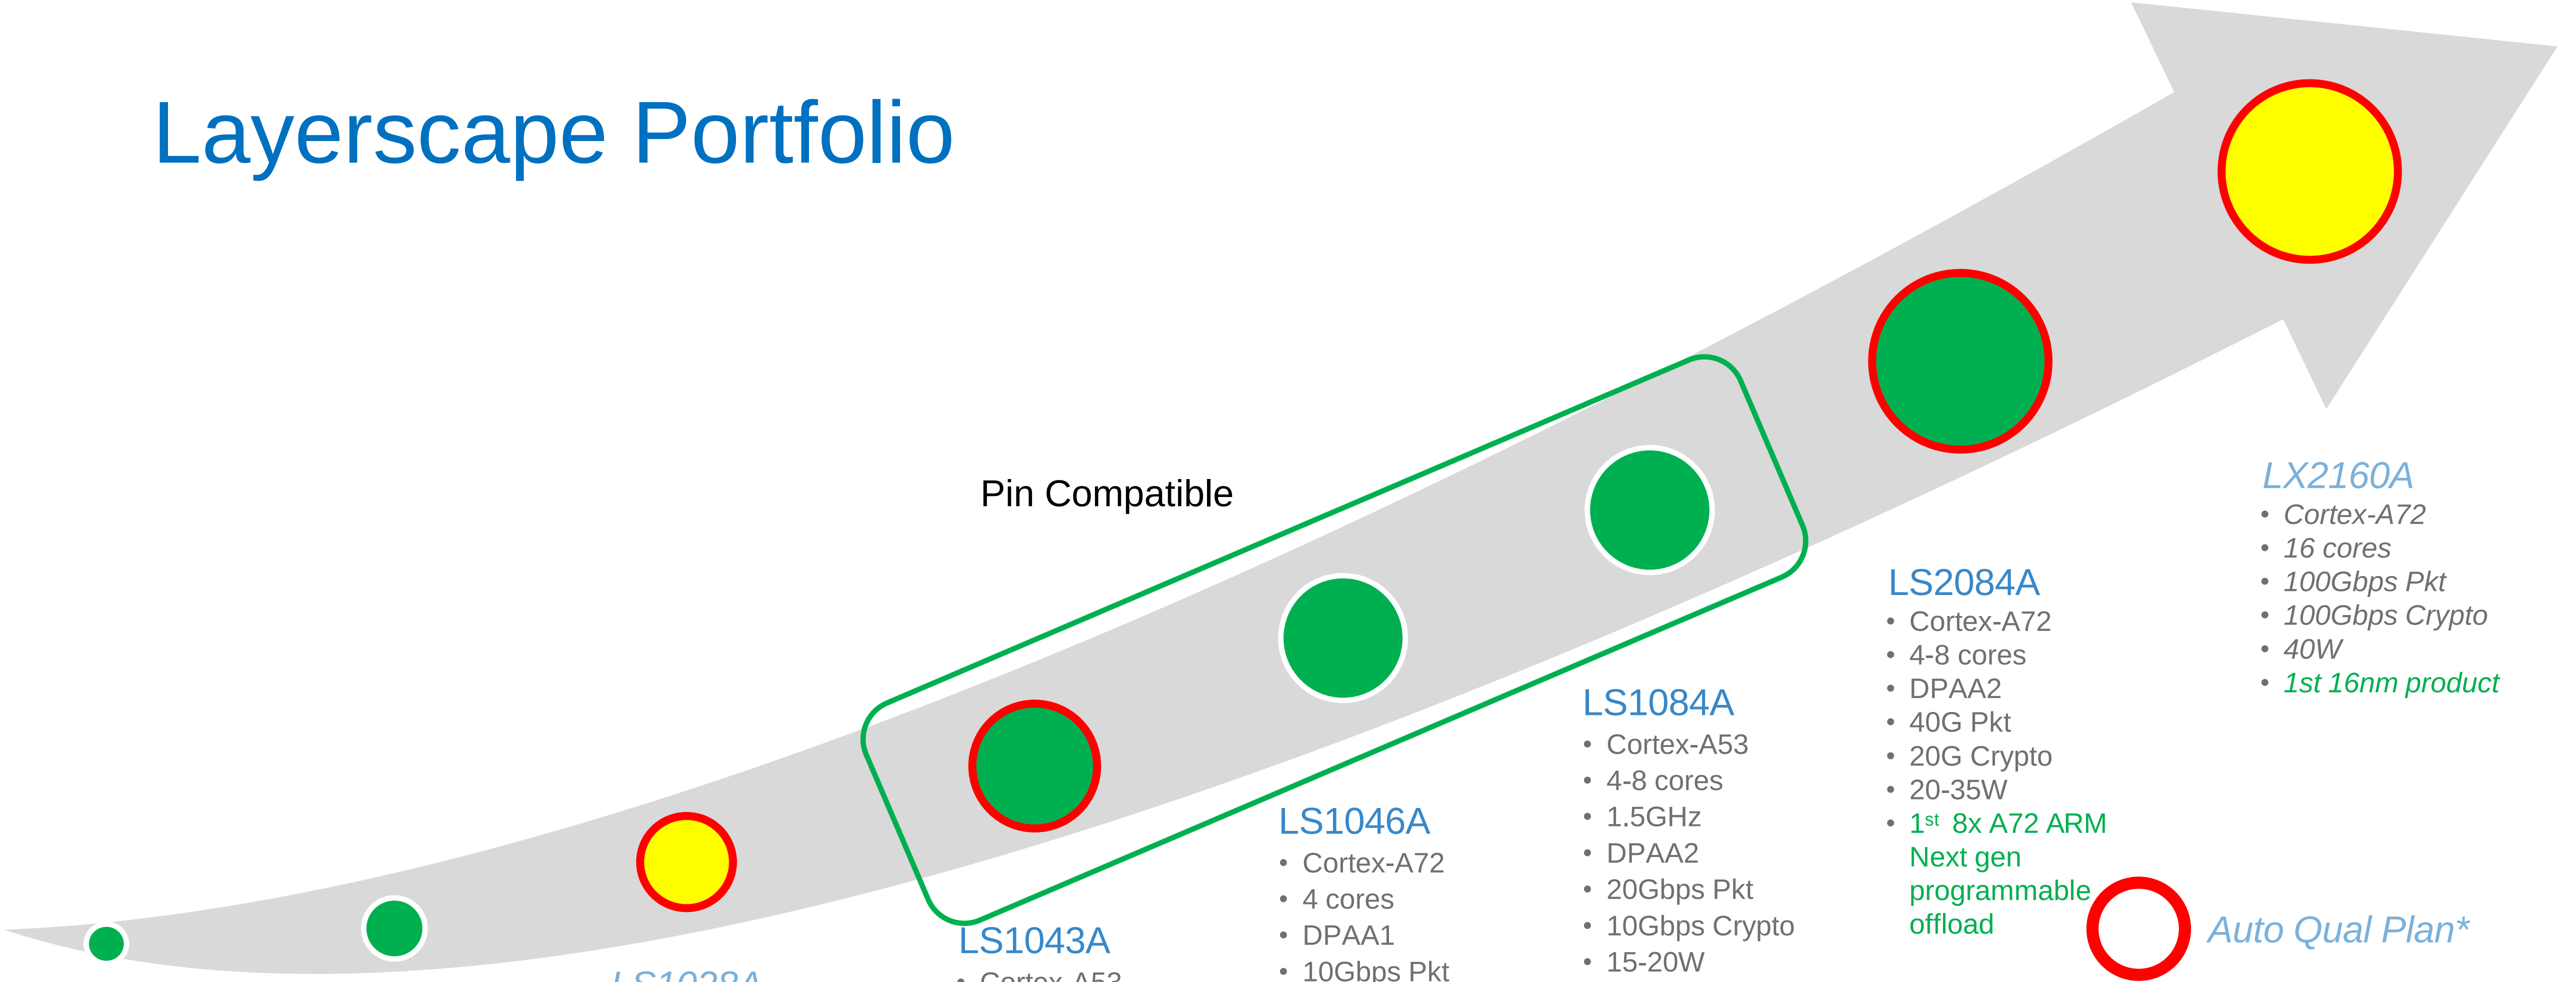

Pin Compatible

LS1012A
- Cortex-A53
- 800MHz
- 2Gbps Packet
- 1Gbps Crypto
- 1-2W
- Lowest power 64-bit ARM

LS1021A
- Cortex-A7
- 2 cores
- 1GHz
- 2Gbps Pkt
- 1Gbps Crypto
- 2W

LS1028A
- *Cortex-A72*
- *2 cores*
- *5Gbps Pkt*
- *5Gbps Crypto*
- 4-9W
- *1st with TSN switch*
- *Integrated GPU*

LS1043A
- Cortex-A53
- 2-4 cores
- 10Gbps Pkt
- 5Gbps Crypto
- 4-8W
- 1st 64-bit ARM processor

LS1046A
- Cortex-A72
- 4 cores
- DPAA1
- 10Gbps Pkt
- 10Gbps Crypto
- 8-15W
- 1st Value Tier A72 ARM

LS1084A
- Cortex-A53
- 4-8 cores
- 1.5GHz
- DPAA2
- 20Gbps Pkt
- 10Gbps Crypto
- 15-20W
- 1st 8x A53 ARM Next gen programmable offload

LS2084A
- Cortex-A72
- 4-8 cores
- DPAA2
- 40G Pkt
- 20G Crypto
- 20-35W
- 1st 8x A72 ARM Next gen programmable offload

LX2160A
- *Cortex-A72*
- *16 cores*
- *100Gbps Pkt*
- *100Gbps Crypto*
- *40W*
- *1st 16nm product*

⭕ *Auto Qual Plan**

🟡 *Sampling*

🟢 *Production*

* Additional products can be auto qualified with business justification

DN Processors in Mission Critical Applications

Aerospace

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

Military and Defense

Rocket navigation, Artillery Control Computer, IFF

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

Factory Automation

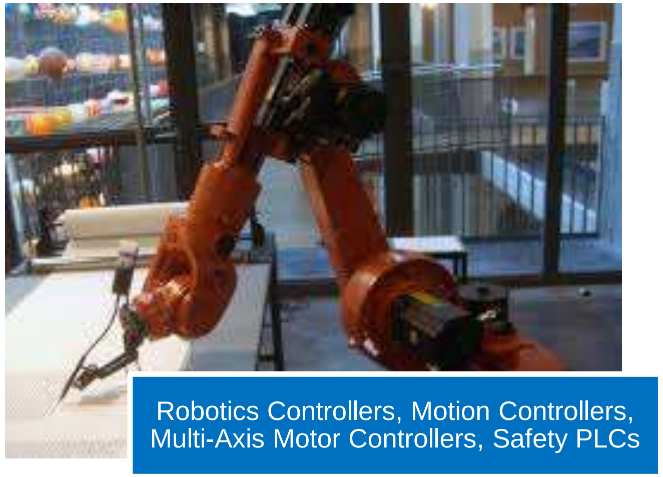

Robotics Controllers, Motion Controllers, Multi-Axis Motor Controllers, Safety PLCs

Railway

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

Power Grid

Power Distribution Relays, Smart Grid Communications

Multicore for Avionics Working Group F2F

Objectives

The Multicore for Avionics Working Group conference is a two-day, deep dive into technical training that targets skills development for engineers across a broad range of embedded technology solutions. In addition to live-demonstrations of the latest innovations from NXP and its partners, the event offers workshops and lectures over multiple markets allowing attendees to customize a schedule that is most relevant to their training needs.

Format

2 day event
Cities throughout America

20 hours of technical training sessions
From NXP and sponsors

General sessions, specialty tracks, sponsor demos, evening networking event

“Meet the Experts” opportunities

City/State	Date
Austin | Texas	October 22-23*

Target Date – 2019 Event in Planning

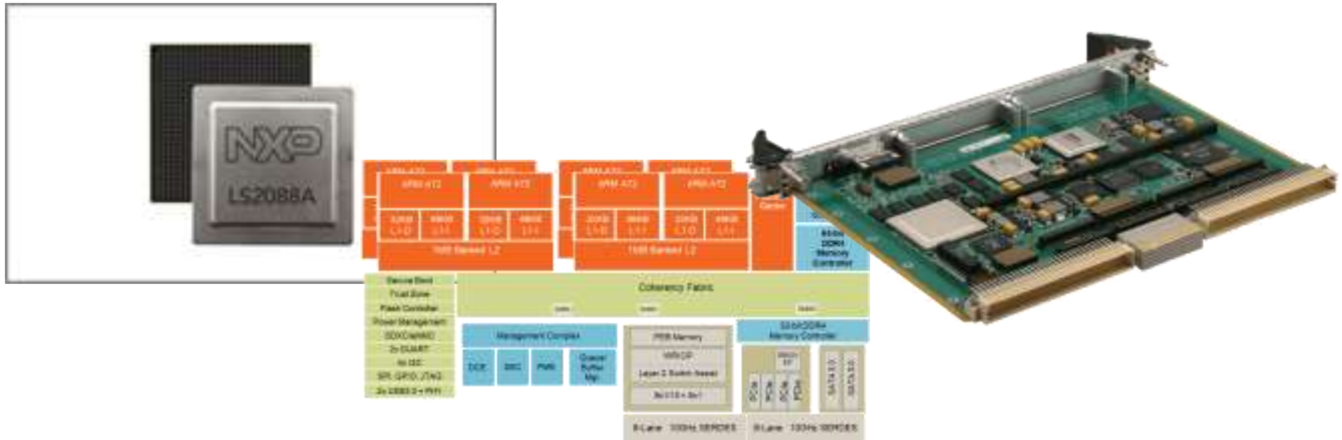

NXP Auto’s View of Domains and Functions

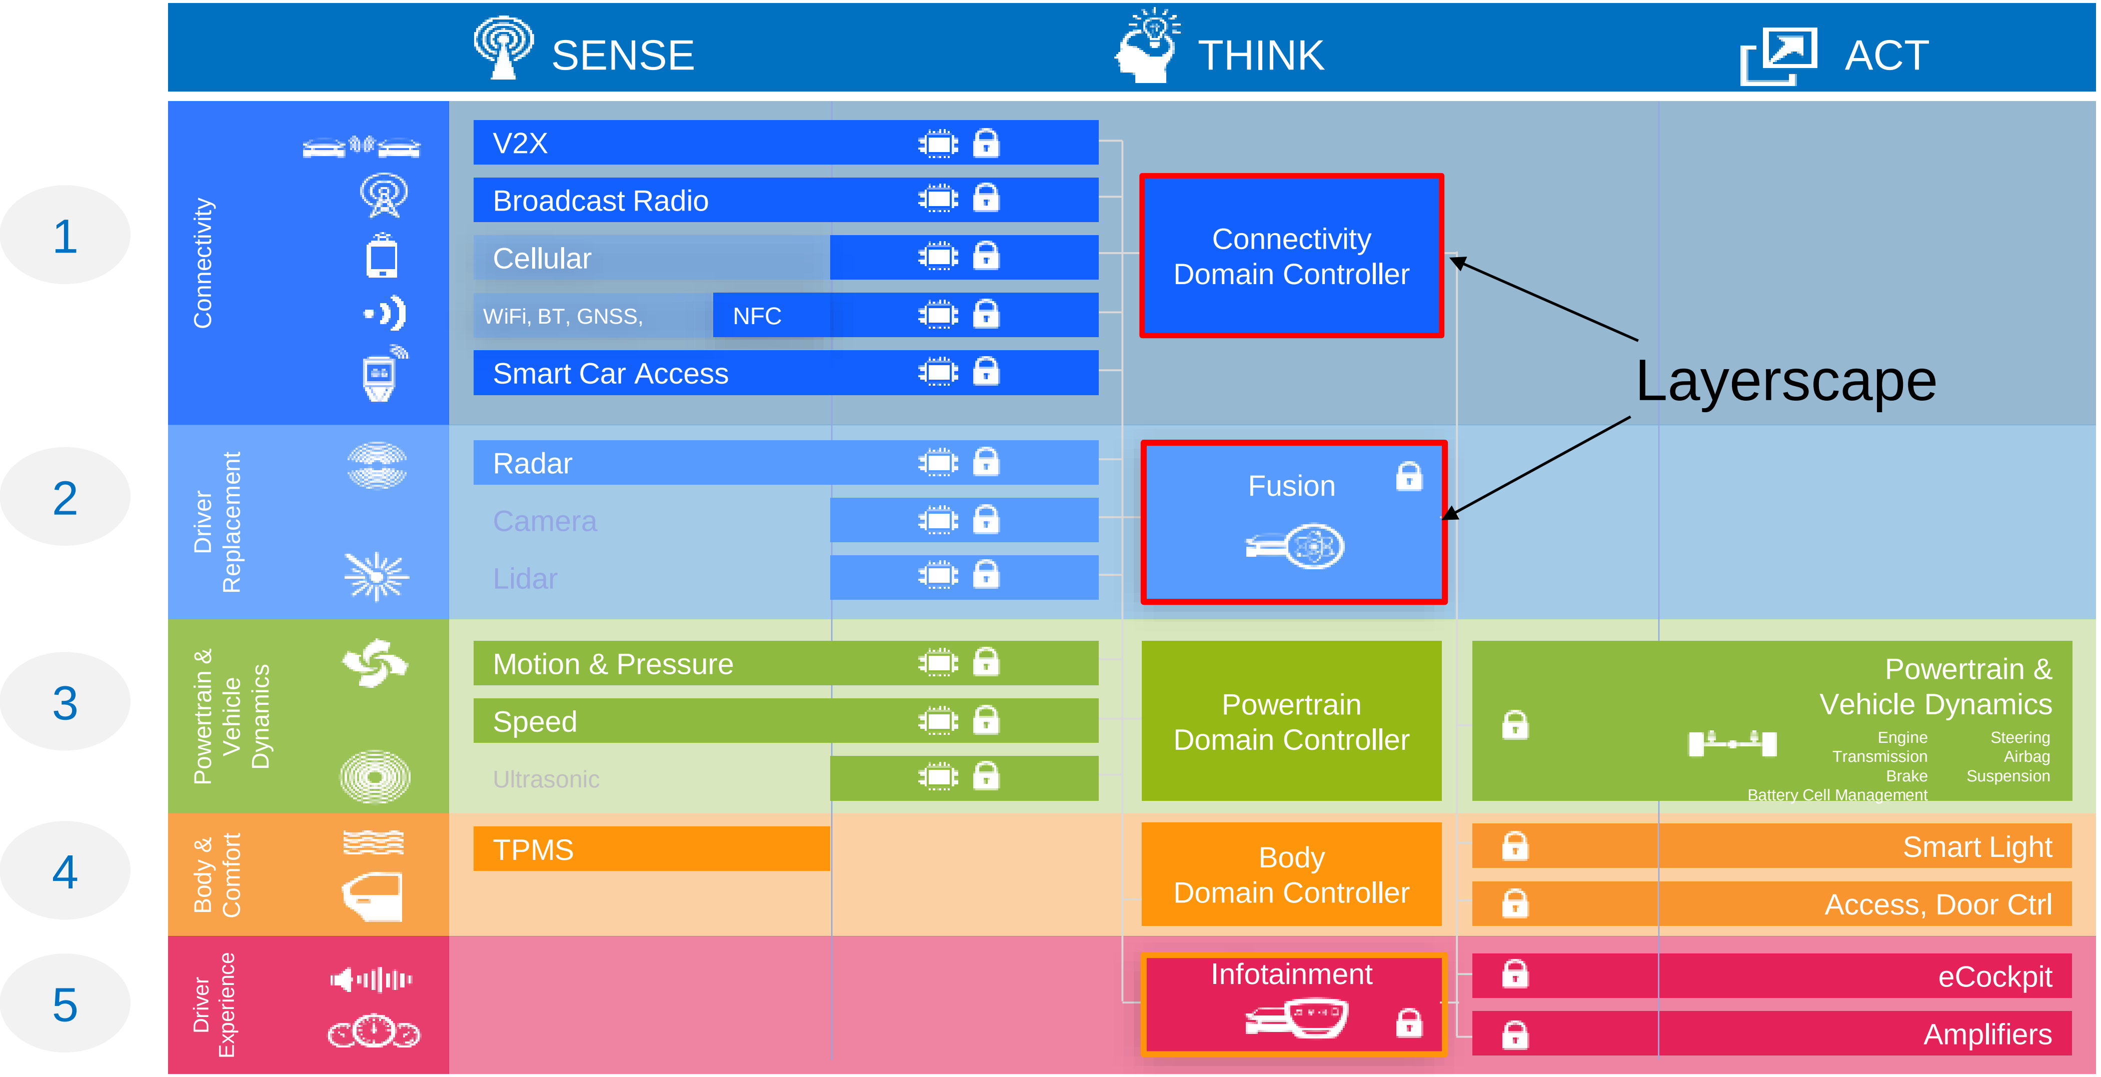

SENSE | THINK | ACT

1 Connectivity
- V2X
- Broadcast Radio
- Cellular
- WiFi, BT, GNSS, NFC
- Smart Car Access

Connectivity Domain Controller

2 Driver Replacement
- Radar
- Camera
- Lidar

Fusion

Layerscape

3 Powertrain & Vehicle Dynamics
- Motion & Pressure
- Speed
- Ultrasonic

Powertrain Domain Controller

Powertrain & Vehicle Dynamics: Engine, Transmission, Brake, Battery Cell Management, Steering, Airbag, Suspension

4 Body & Comfort
- TPMS

Body Domain Controller

- Smart Light
- Access, Door Ctrl

5 Driver Experience

Infotainment

- eCockpit
- Amplifiers

Conceptual Vehicle Architecture

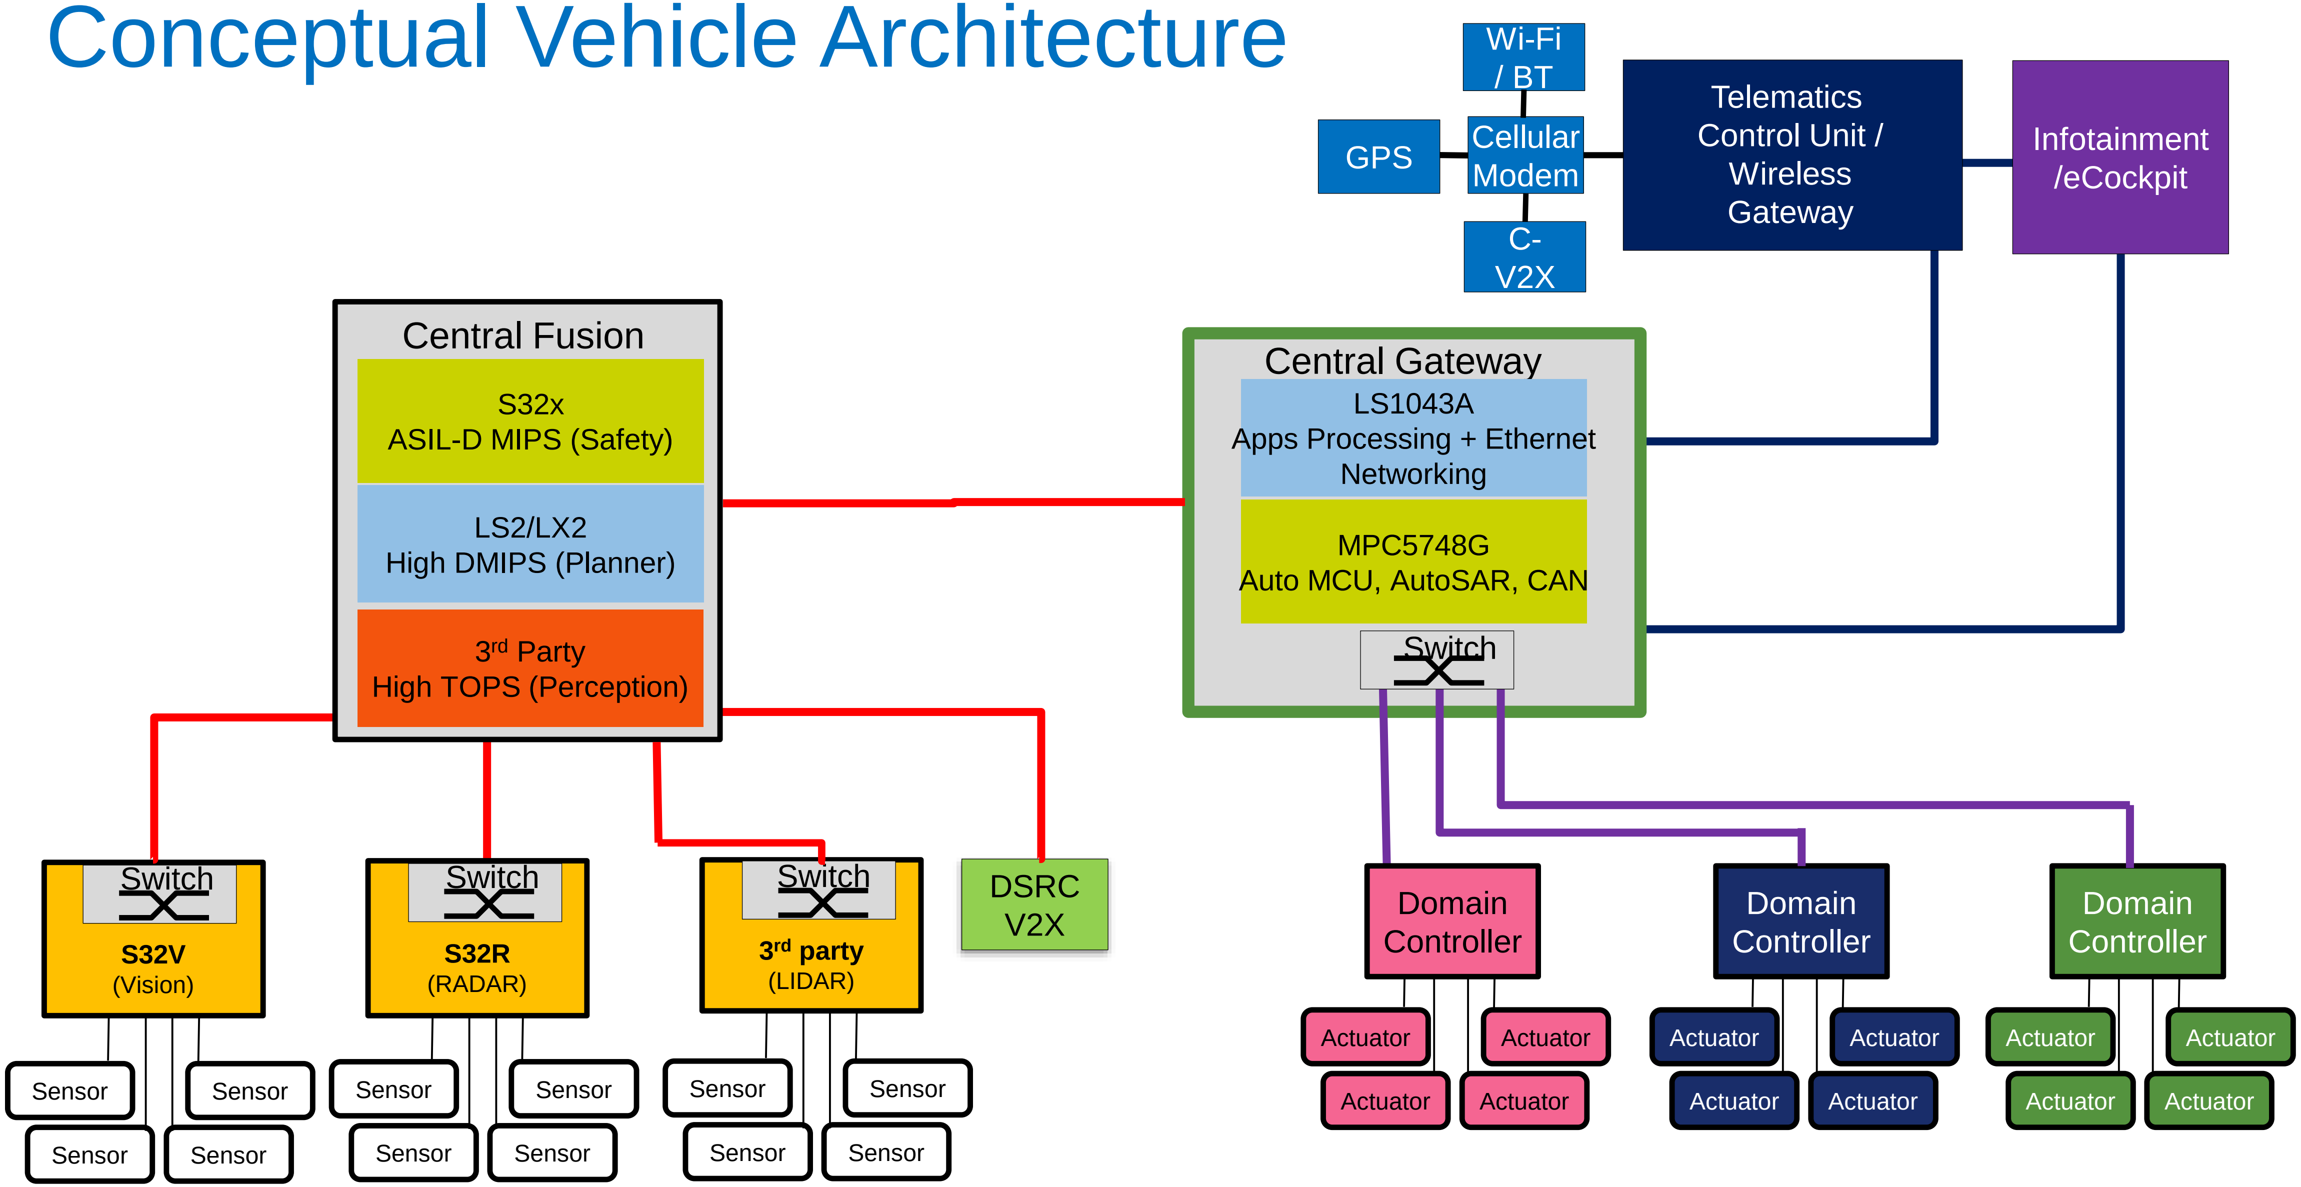

Layerscape Safety Positioning

- Layerscape SoCs were not designed specifically for the automotive market and do not offer certain features such as a Fault Collection and Control Unit (FCCU) that are normally provided in NXP purpose-built automotive products. However, Layerscape SoCs contain numerous reliability and security functions that can be leveraged as safety mechanisms

- Also, careful systems engineering at the board, software, and system level can compensate for some of the missing automotive-specific hardware features. Using this sort of holistic design approach, the high performance computing and network switching capabilities of Layerscape SoCs can be leveraged in a manner consistent with ISO 26262 ASIL B requirements

Layerscape LS2084A

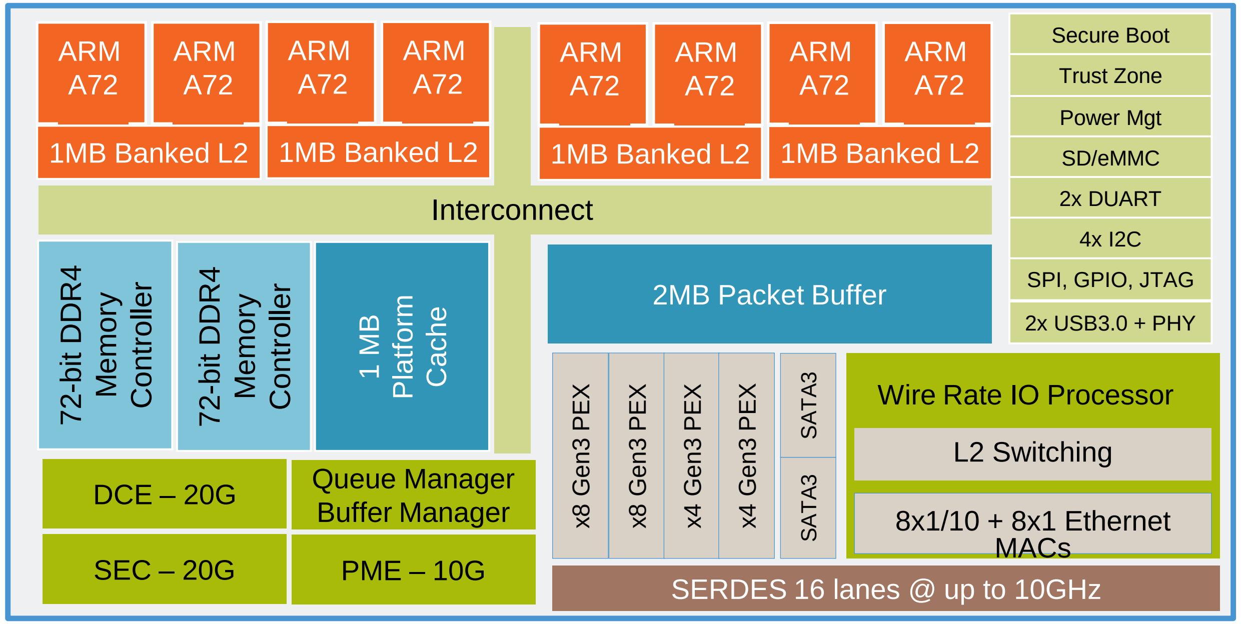

ARM A72 | ARM A72 | ARM A72 | ARM A72
1MB Banked L2 | 1MB Banked L2
ARM A72 | ARM A72 | ARM A72 | ARM A72
1MB Banked L2 | 1MB Banked L2

Secure Boot
Trust Zone
Power Mgt
SD/eMMC
2x DUART
4x I2C
SPI, GPIO, JTAG
2x USB3.0 + PHY

Interconnect

72-bit DDR4 Memory Controller
72-bit DDR4 Memory Controller
1 MB Platform Cache
2MB Packet Buffer

x8 Gen3 PEX
x8 Gen3 PEX
x4 Gen3 PEX
x4 Gen3 PEX
SATA3
SATA3

Wire Rate IO Processor
L2 Switching
8x1/10 + 8x1 Ethernet MACs

DCE – 20G
Queue Manager Buffer Manager
SEC – 20G
PME – 10G

SERDES 16 lanes @ up to 10GHz

Major Milestone	Schedule
Samples (Production Rev)	Dec 2017
Networking/Telecom Qualification	March 2018
AECQ100 grade 3 Qual on production rev	Nov 2018
PPAP Completion	Aug 2019

Auto Quality

- AEC Q100 Grade 3 (105C Tj)
- 15 years product longevity
- ZD-like approach to reduce risk of DPPM or Life failures
- Expected Operating Life fail rate <10 FIT
- Mission Profile: 10 years, 90C Tj-effective

Performance (Grade 3)

- ARM A72 x 8 @ 1.8 GHz
  - 86K DMIPS
  - SpecInt2k6 – 13.1, Rate -75.1
  - Neon SIMD in all CPUs
- 2x72b (w/ECC) DDR4 @ 1.8GT/s
  - 28.8GB/s memory BW
- High Speed IO
- Multiple PCIe Gen3 controllers
- Multiple Ethernet MACs (up to 10G)

Functional Safety

- Target ASIL-B*
- ECC protected memories
- Fault localization, containment and recovery
- Soft lockstep with determinism
- Excellent support for virtualization, containerization

Process & Package

- 28HPM, ~40W Thermal Max @ 105C
- 37.5 x 37.5 mm, lidded FCBGA, 1mm pitch, 1292 pins

Security

- 20Gbps Crypto Acceleration
- MACSEC, IPsec, SSL
- Trust Architecture
  - Secure Boot
  - Secure Debug
  - Secure Storage
  - Tamper Detection
  - HW Enforced Partitioning
  - ARM Trust Zone

LS2084A PMHF Spectrum

Safety Goal: Data Rx (Ethernet and PCIe), processing, data Tx is correct, else detected and ECU signals ‘Not Safe’

	Reliability Based[1]		IEC TR 62380 Based[2]	
	50% Fractional Safe Faults[3]	No Fractional Safe Faults	50% Fractional Safe Faults[3]	No Fractional Safe Faults
No Sys Safety Assumptions[4]	244.69	489.37	671.70	1343.40
Minimal Sys Safety Assumptions[5]	184.56	369.12	533.53	1067.17
Sys Safety Concept[6]	75.85	151.71	163.35	326.70

[1] Die and package failure rates based on Digital Networking 28nm reliability data, field returns
[2] Die and package failure rates based on IEC TR 62380
[3] Per ISO 26262 2018 Version Part 10, Clause 8.1.8
[4] Die level failure detection/correction; primarily ECC on internal RAMs
[5] Die level failure detection/correction, plus IO data corruption and IO/accelerator memory access violation detection
[6] Based on system safety concept from BlueBox vehicle platooning; safety MCU watchdog, external power & clock monitoring

Layerscape LX2160A

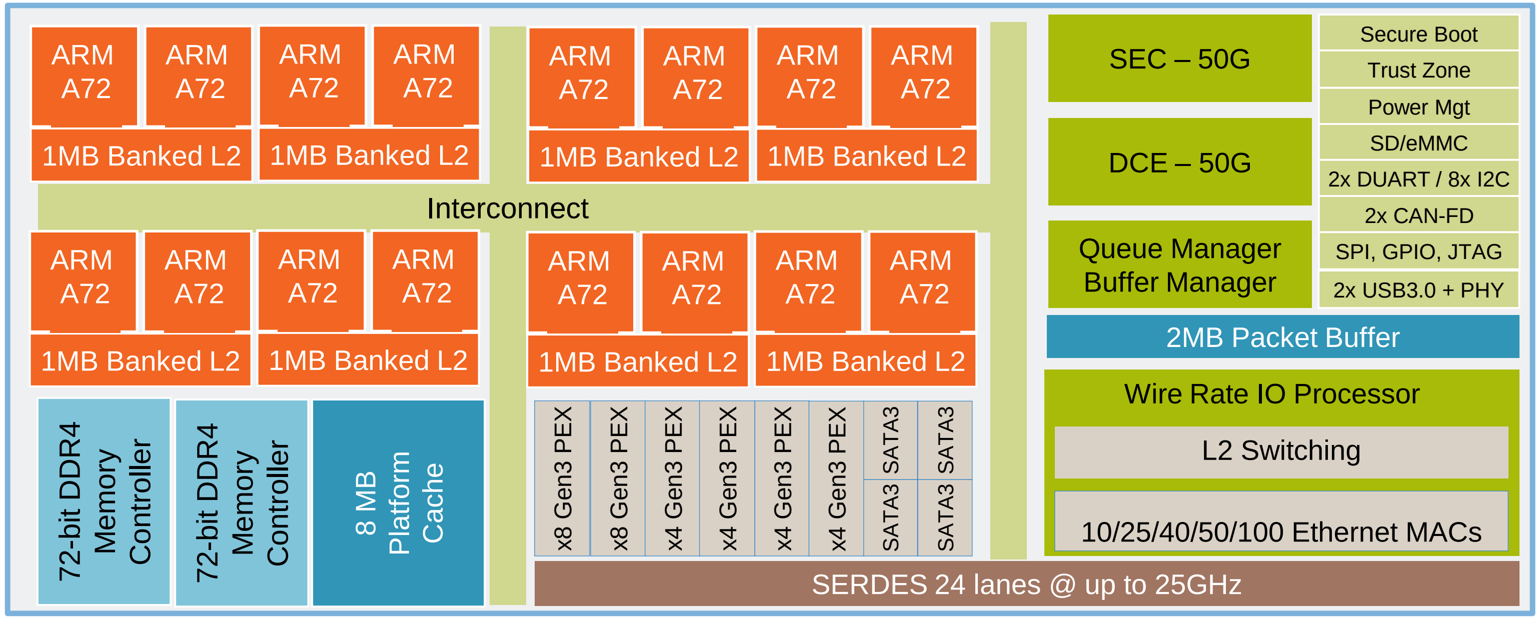

ARM A72 | ARM A72 | ARM A72 | ARM A72 | ARM A72 | ARM A72 | ARM A72 | ARM A72

1MB Banked L2 | 1MB Banked L2 | 1MB Banked L2 | 1MB Banked L2

Interconnect

ARM A72 | ARM A72 | ARM A72 | ARM A72 | ARM A72 | ARM A72 | ARM A72 | ARM A72

1MB Banked L2 | 1MB Banked L2 | 1MB Banked L2 | 1MB Banked L2

72-bit DDR4 Memory Controller | 72-bit DDR4 Memory Controller | 8 MB Platform Cache

x8 Gen3 PEX | x8 Gen3 PEX | x4 Gen3 PEX | x4 Gen3 PEX | x4 Gen3 PEX | x4 Gen3 PEX | SATA3 | SATA3 | SATA3 | SATA3

SEC – 50G

DCE – 50G

Queue Manager Buffer Manager

Secure Boot

Trust Zone

Power Mgt

SD/eMMC

2x DUART / 8x I2C

2x CAN-FD

SPI, GPIO, JTAG

2x USB3.0 + PHY

2MB Packet Buffer

Wire Rate IO Processor

L2 Switching

10/25/40/50/100 Ethernet MACs

SERDES 24 lanes @ up to 25GHz

Samples (Rev1):	Now
Samples (Rev2):	April 2020 (fully tested)
Telecom Production:	May 2020
Auto Grade 3 & PPAP:	Oct 2020

Auto Quality

- AEC Q100 Grade 3 (105 Tj)
- 15 years product longevity
- ZD-like approach to reduce risk of DPPM or Life failures
- Expected Operating Life fail rate <10 FIT
- Mission Profile: 10 years, 90C Tj-effective

Performance

- ARM A72 x 16 @ 2.2 GHz
  - ~201K DMIPS
  - SpecInt2k6 – 17.6, Rate -157
  - Neon SIMD in all CPUs
- 2x72b (including ECC) DDR4 up to 3.2GT/s
  - 51GB/s memory BW
- High Speed IO
- Multiple PCIe Gen3 controllers
- Multiple Ethernet MACs (up to 100G)

Functional Safety

- Target ASIL-B*
- ECC protected memories
- Fault localization, containment and recovery
- Soft lockstep with determinism
- Excellent support for virtualization, containerization

Process & Package

- 16FFC, ~25W Thermal Max @ 105C – 2.0GHz
- 40x40mm, Lidded FCBGA, 1mm pitch (1517 pins)

Security

- 50Gbps Crypto Acceleration
- MACSEC, IPsec, SSL
- Trust Architecture
  - Secure Boot
  - Secure Debug
  - Secure Storage
  - Tamper Detection
  - HW Enforced Partitioning
  - ARM Trust Zone

LX2160A PMHF Spectrum

Safety Goal: Data Rx (Ethernet and PCIe), processing, data Tx is correct, else detected and ECU signals ‘Not Safe’

	Reliability Based[1]		IEC TR 62380 Based[2]	
	50% Fractional Safe Faults[3]	No Fractional Safe Faults	50% Fractional Safe Faults[3]	No Fractional Safe Faults
No Sys Safety Assumptions[4]	59.45	118.90	441.03	882.05
Minimal Sys Safety Assumptions[5]	54.57	109.15	364.25	728.51
Sys Safety Concept[6]	38.16	76.33	116.33	232.66

[1] Die and package failure rates based on Digital Networking 28nm reliability data, field returns

[2] Die and package failure rates based on IEC TR 62380

[3] Per ISO 26262 2018 Version Part 10, Clause 8.1.8

[4] Die level failure detection/correction; primarily ECC on internal RAMs

[5] Die level failure detection/correction, plus IO data corruption and IO/accelerator memory access violation detection

[6] Based on system safety concept from BlueBox vehicle platooning; safety MCU watchdog, external power & clock monitoring

QorIQ Layerscape LS1043A

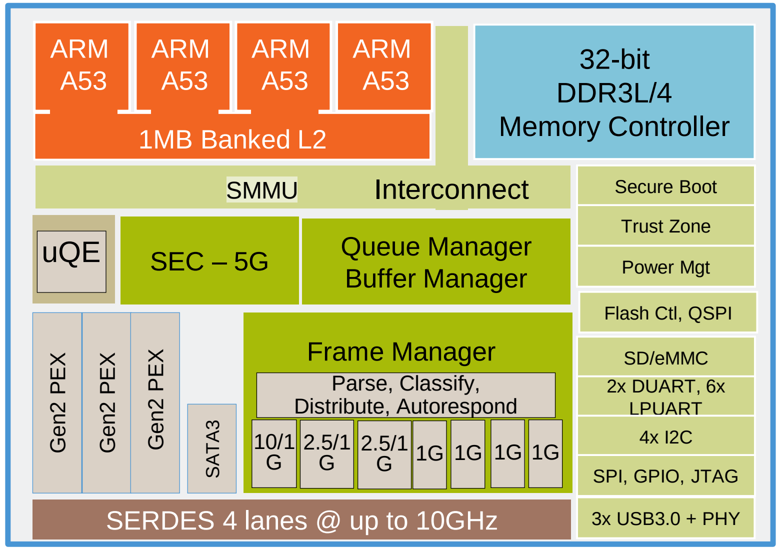

Major Milestone	Schedule
Engineering Samples Rev 1.1	Completed / October 4, 2016
Networking/Telecom Qualification	Completed / January 25, 2017
AECQ100 grade 3 Qual on Rev 1.1	Complete / Sept 12, 2017
PPAP Completion	June 2018 Updated PPAP (for new lidded package) Jan 2019

- AEC Q100 Grade 3 (105 Tj max)
- 15 years product longevity
- ZD-like approach to reduce risk of DPPM or Life failures
- Expected Operating Life fail rate <10 FIT
- Mission Profile: 10 years, 90C Tj-effective

Process & Package

- 28HPM, ~5-9W Thermal Max @ 105C
- 23x23mm, Lidded FCBGA, .8mm pitch (780 pins)

Performance

- ARM A53 x 4 @ up to 1.6GHz (LS1023A: 2 cores)
  - 19.5K DMIPS
  - SpecInt2k6 – 5.95, Rate -15
  - Neon SIMD in all CPUs
- 1x36b (including ECC) DDR3L/4 up to 1.6GT/s
  - 6.4GB/s memory BW
- High Speed IO
  - Multiple PCIe Gen2 controllers
  - Multiple Ethernet MACs (up to 10G)

Functional Safety

- Target ASIL-B*
- ECC protected memories
- Fault localization, containment and recovery
- Soft lockstep with determinism
- Excellent support for virtualization, containerization

Security

- 5Gbps Crypto Acceleration
- IPsec, SSL
- Trust Architecture
  - Secure Boot
  - Secure Debug
  - Secure Storage
  - Tamper Detection
  - HW Enforced Partitioning
  - ARM Trust Zone

Detecting Unsafe Hardware Operation

Layerscape SoCs are extensively verified pre-silicon, with additional post-silicon validation and qualification, to ensure proper operation.

Consequently, the most likely cause of unsafe hardware operations is operation of the SoC outside of a specified environmental parameters.

These environmental conditions include:

1. Min & Max operating temperature
2. Min & Max operating voltages
3. Min & Max clock rates and jitter

Operating outside of specified environmental parameters can lead to:

- Detected Correctable and Uncorrectable/Fatal Errors
- Undetected IO or processing errors (data corruption, incorrect results) aka glitches
- Undetected stoppage of operations, aka hangs

Detecting Out of Spec Environmental Parameters

- Detecting Out of Spec Temperature
- Layerscape SoCs incorporate a Thermal Management Unit (TMU) which can be polled by safety software to determine instantaneous and average temps
  - TMU can also generate interrupts when instantaneous of average threshold is exceeded
  - Safety software can notify Safety MCU of impending failure due to out of spec operation
- Detecting Out of Spec Power
- Layerscape SoCs do not incorporate fine grained power monitoring circuits
  - Voltage threshold detection is implemented at power on reset, coarse grained brown out/glitch detection available at runtime
  - External power monitoring can be provided by the recommended NXP System Basis Chip (SBC)/PMIC
  - Out of spec power condition must be reported to the Safety MCU
- Detecting Out of Spec Clocks
- Layerscape SoCs do not incorporate fine grained clock monitoring circuits
  - Input clocks are fed into PLLs which generate the clocks used by the cores, platform, and DDR controller(s).
  - PLL loss of lock detection is implemented
  - If SERDES PLL clock doesn’t lock on start up, SERDES block requests device reset.
  - Out of spec clock condition must be reported to the Safety MCU

NXP Power Solution of Safety ADAS Module with LX2160A

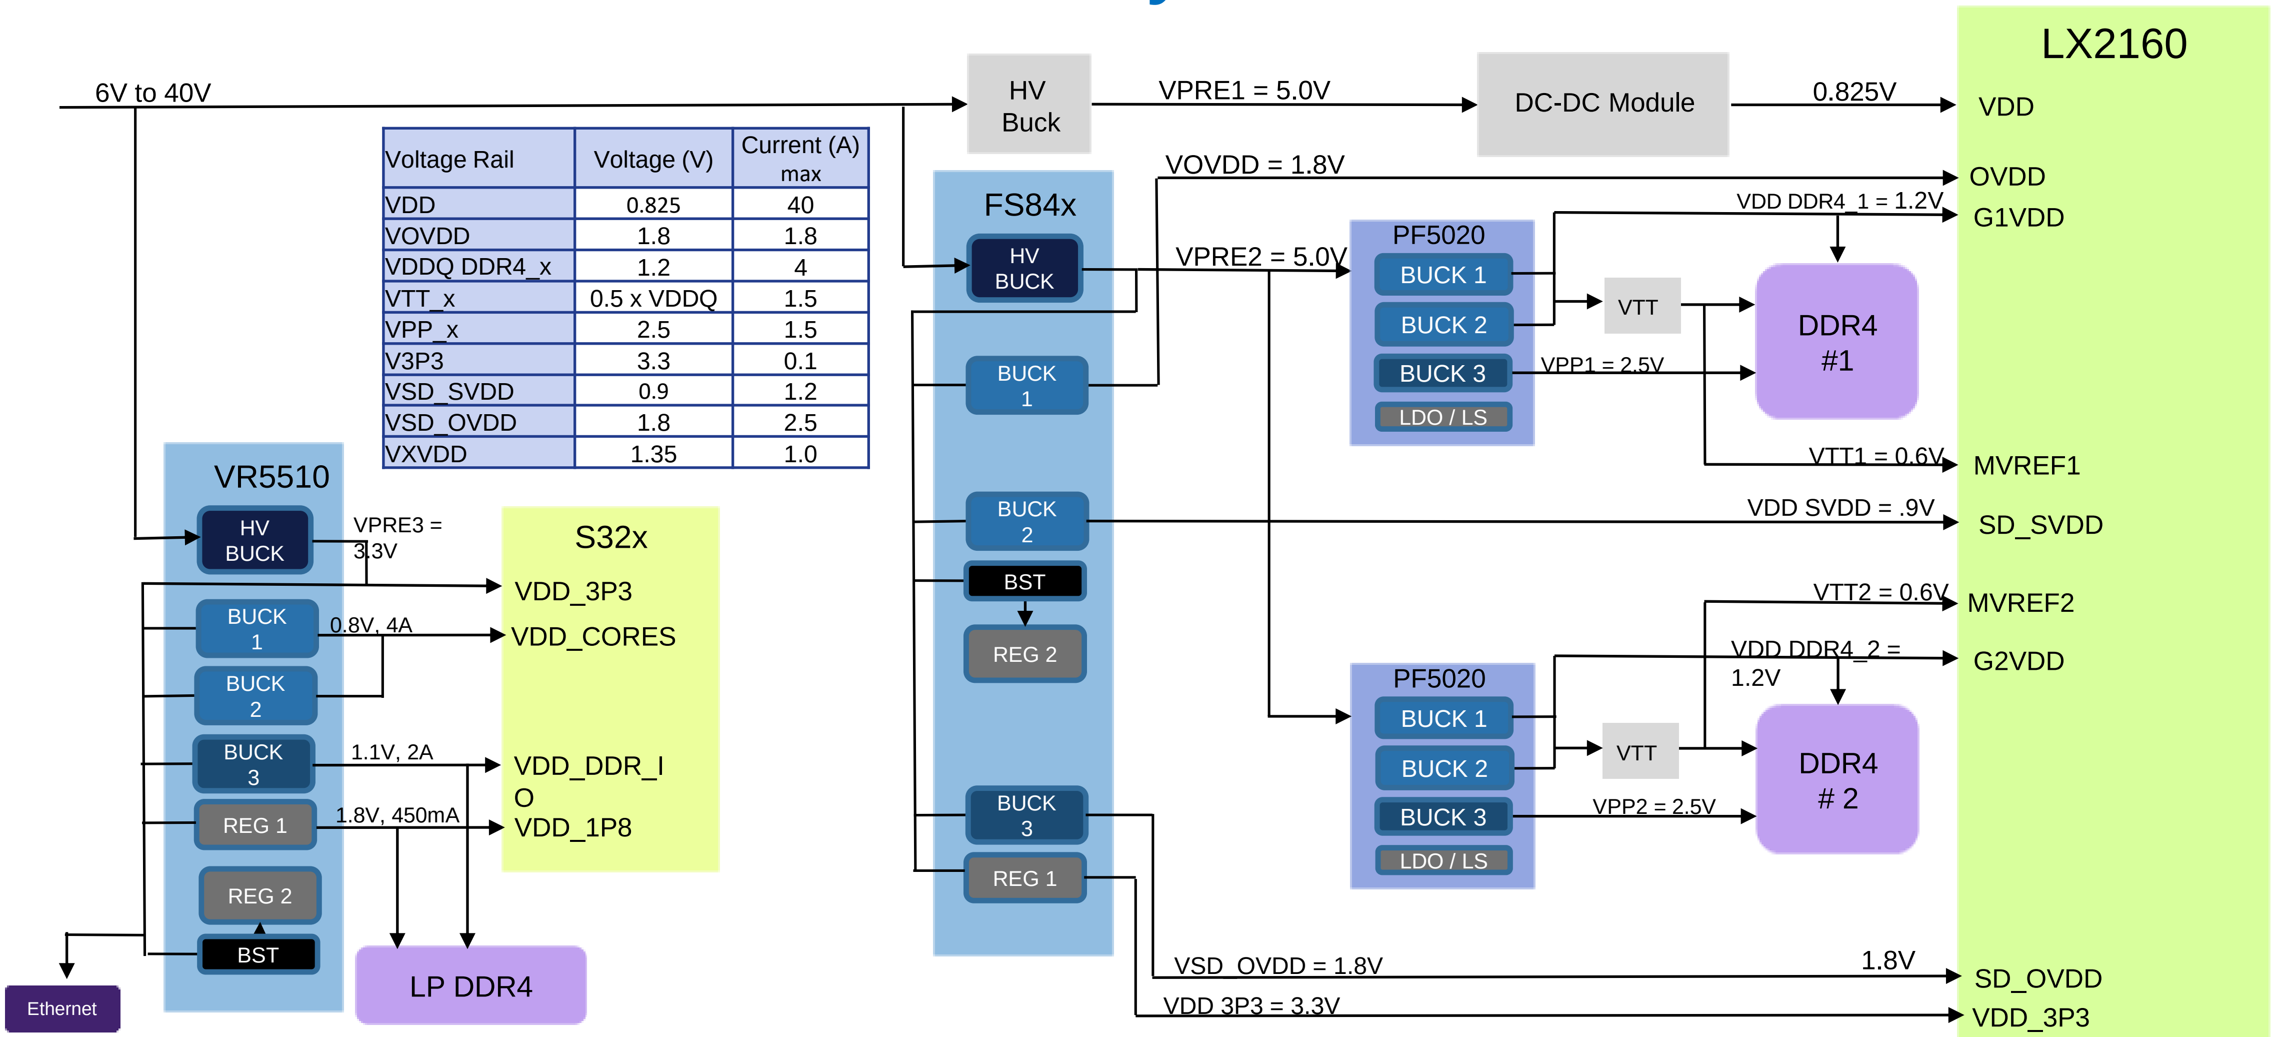

Voltage Rail	Voltage (V)	Current (A) max
VDD	0.825	40
VOVDD	1.8	1.8
VDDQ DDR4_x	1.2	4
VTT_x	0.5 x VDDQ	1.5
VPP_x	2.5	1.5
V3P3	3.3	0.1
VSD_SVDD	0.9	1.2
VSD_OVDD	1.8	2.5
VXVDD	1.35	1.0

S32x ASIL-D MCU along with VR5510 SBC would provide ASIL-D System Solution

Safety Signals

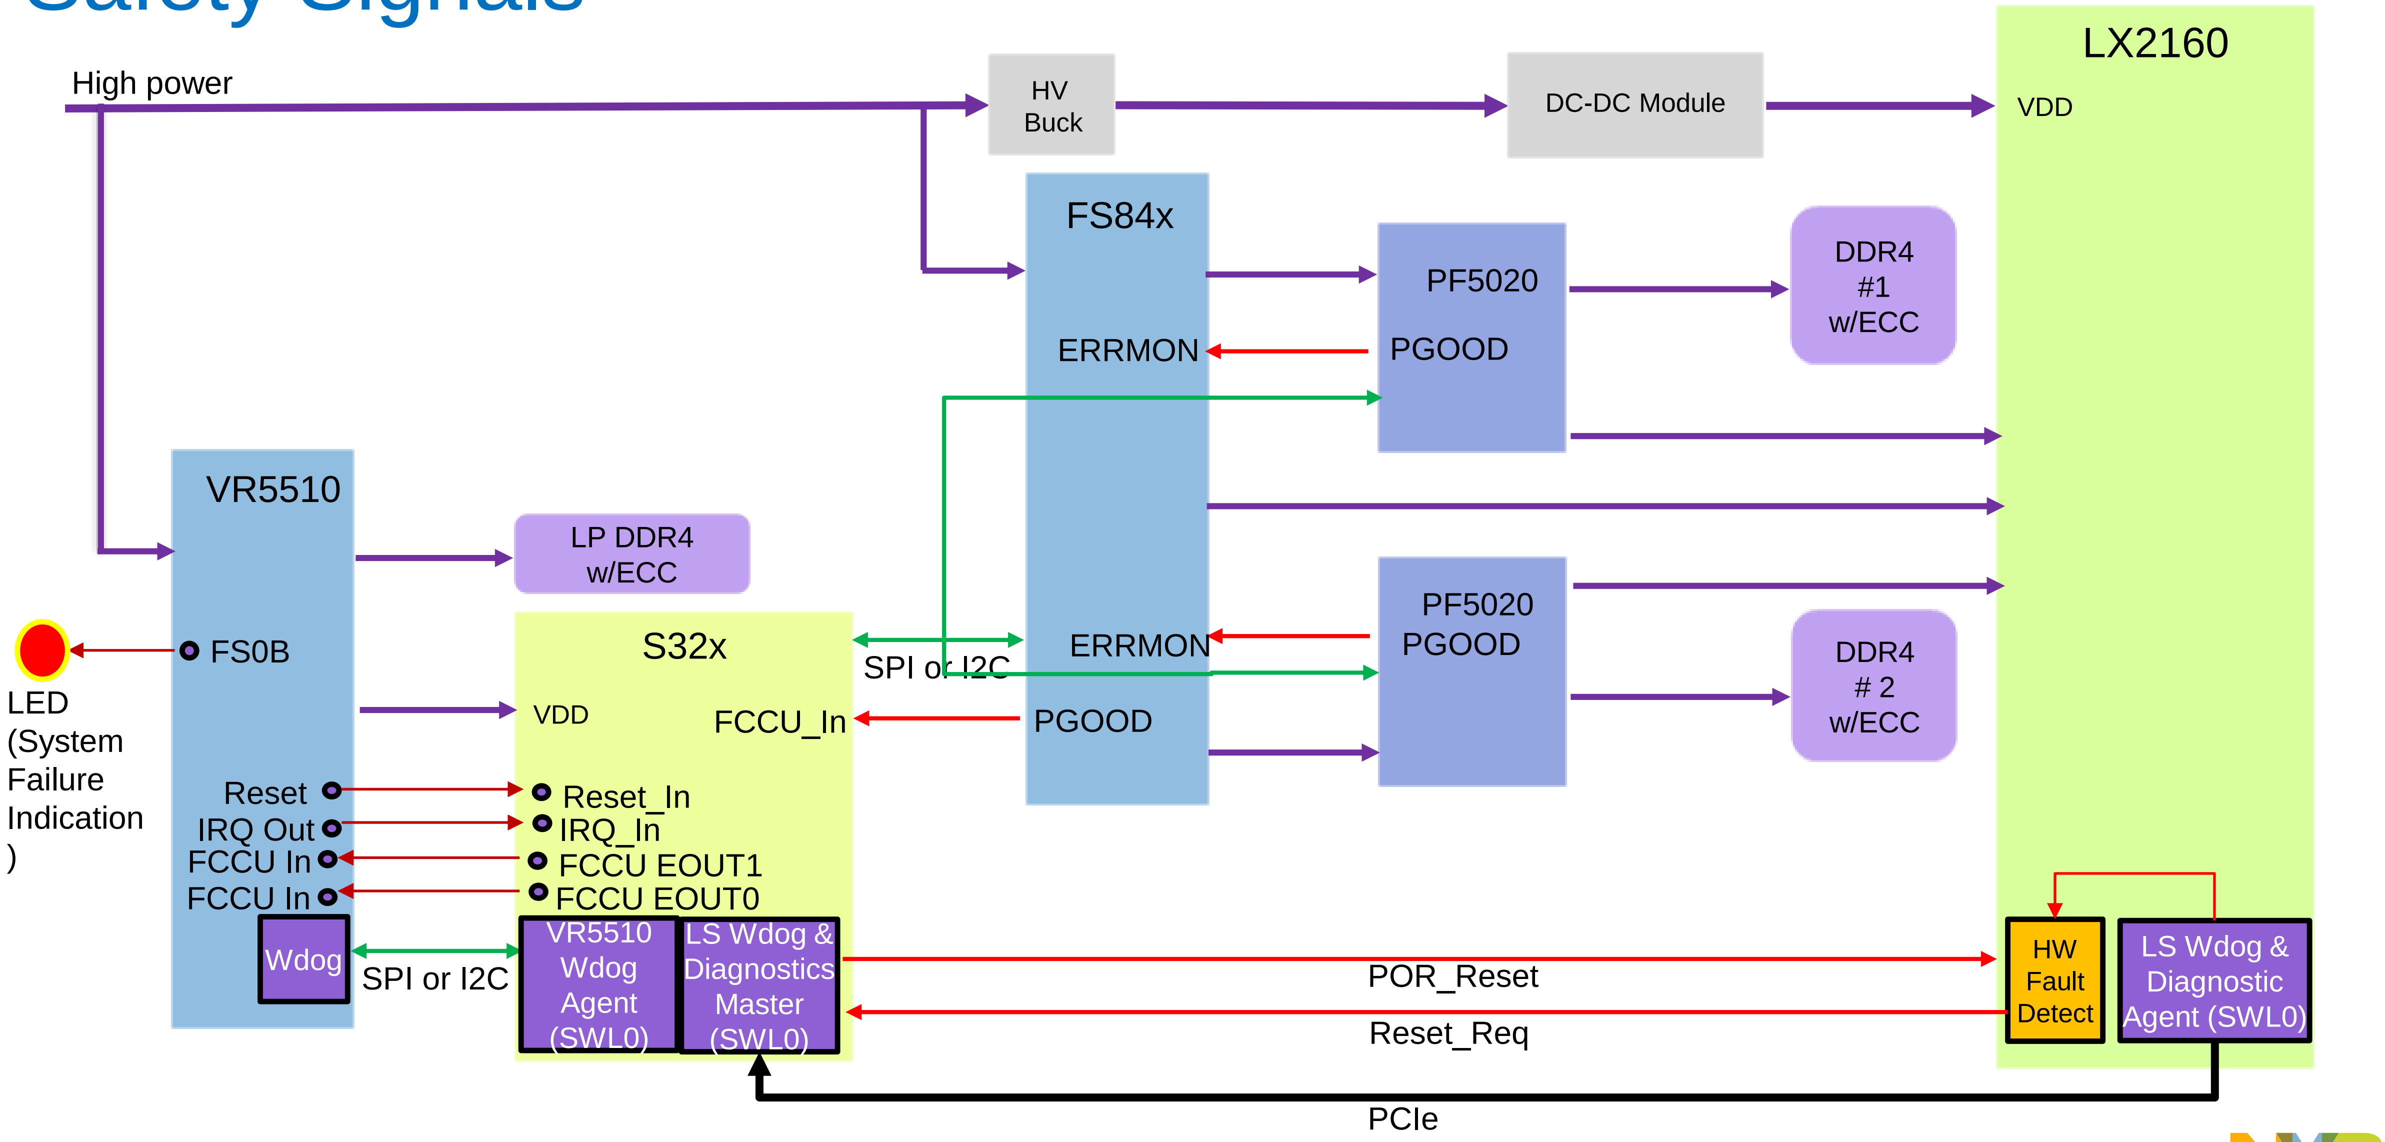

Detectable Correctable Errors

- The main correctable error detected by Layerscape SoCs is single bit flips in the internal and external memories
- These errors are detected by hardware on reads, including the address of the word where the bit flip occurred
- Software can maintain a count of total corrected single bit flips, as well as track the location of the corrections
- Regions of memory which are seldom read in operation should be protected by a ‘memory scrubber’ routine, which periodically reads the memory region to trigger single bit error corrections before fatal multi-bit errors accumulate
- Other detectable errors, which could be correctable depending on software, include;
  - Mis-directed reads/writes which are blocked by memory access controls
  - Hardware time-outs in certain non-CPU bus masters

Detectable Uncorrectable Errors

Layerscape internally detects many types of unrecoverable errors during SoC initialization, and at runtime. Detection can occur in hardware, firmware, or safety software.

When Layerscape HW detects an uncorrectable error, it asserts the Reset_Req signal to tell external logic that it is in an unrecoverable state and in need of reset. Uncorrectable errors detected in hardware include:

Sources:

- SERDES (PLL lock failure)
- Run Control Power Mgt (RCPM) Unit time-out
- POR BIST
- Multi-bit ECC Error
- Interconnect Misc Node
- Secure Debug Controller
- Security Monitor
- Service Processor
- Management Complex
- Integrated Flash Controller
- TrustZone Watchdog Timer
- Per CPU Watchdog Timers
- Any software with write access to Reset_Ctl Register

Layerscape Interconnects

- LS1043A uses CCI-400
- LS2 uses CCN-504
- LX2 uses CCN-508
- These Arm interconnects route transactions across the interconnect (from ‘node’ to ‘node’) in ‘packets’.
- All nodes perform the following error detection and reporting;
  - classifying the error as either correctable or uncorrectable fatal
  - logging the relevant error information in dedicated Error Syndrome registers that are mapped into the configuration address space and accessed over the block’s configuration bus
  - Signaling the error to the Misc Node (MN)
- Examples of errors include;
  - Correctable – single bit ECC error in L3. Corrected and reported if threshold count is reached
  - Uncorrectable – double bit ECC error in L3. Reported immediately.
- When the MN receives an error signal, the signal is sticky and is only cleared by the error handler reading the Error Signal Valid registers in the MN.
- Layerscape’s implementations of these interconnects also includes parity checking over each packet, allowing corrupted transactions to be detected and reported to the Misc Node (MN), thereby triggering the Reset_Req.

Layerscape Memories

Layerscape devices have 3 classes of memory

- Internal SRAM
- External DRAM (DDR main memory)
- External NVRAM (multiple types and interfaces supported)

Layerscape SoCs have extensive ECC on internal SRAMs. Some IO buffers are parity protected only due to short duration of data residence.

- Exact error reporting pathways and reactions depend on the block the SRAM is located in.
- Error injection for self test also supported

Customers must provision boards with wider DRAM memories (x36 or x72) for the DDR controller to perform ECC.

NVRAM ECC support depends on the specific NVRAM interface. Managed flash (typically with serial interfaces) includes error detection, often at the block level. The Layerscape Integrated Flash Controller (IFC) supports ECC similarly to external DRAM; the system is configured with extra data bits to store ECC data calculated by the IFC.

Watchdog Timers

- Layerscape SoCs are provisioned with a number of timers, each of which can generate multiple timer interrupts at configured intervals.
- Each Armv8 core has a dedicated watchdog timer. When enabled, if the timer expires without software reaction, the core (and the software running on it) is considered non-responsive.
  - Individual core watchdogs are capable of triggering the Reset_Req.
  - Individual core watchdog Reset_Reqs can be masked.
    - If the system is virtualized, restarting a hung VM is more appropriate than triggering a SoC reset.
    - If the core hardware itself is corrupted, the hung VM won't restart, and the hypervisor or safety software can trigger the Reset_Req
- TrustZone Secure World has a dedicated watchdog timer.
  - Safety software should run in TZ Secure World, and this is the logical master watchdog for triggering (in hardware) the Reset_Req.
  - The TZ Watchdog Reset_Req is only maskable when the SoC is in debug mode.
- Many hardware blocks also include timers, where if the block's transaction isn't completed within the expected period of time (or internal processing is hung), the timer triggers the block to generate a catastrophic error interrupt.
  - These blocks can be independently reset, in some cases allowing for block recovery
  - Safety software may determine restarting the block and unwinding its incomplete operations is too complicated
  - Even in this case, the safety software can notify the Safety MCU of an impending restart

Memory Access Controls

- Layerscape SoCs have excellent support for virtualization, making concepts like ECU consolidation easy to adopt.
- This virtualization support also includes strong, hardware enforced partitioning.
  - Accidental or malicious attempts by partition 1 to access partition 2’s private resources are detected and blocked.
- Memory access controls act as a whitelist for software partitions and the IOs/accelerators working for them. Corrupted addresses, including from misprogrammed or corrupted descriptors, are highly likely to miss the whitelist. This results in a blocked transaction and an error interrupt.
- Some access violations can be configured to trigger the Reset_Req (via the Security Monitor)
- Hardware providing this enforcement includes;
  - CPU MMUs
  - IO MMUs (called SMMU in Layerscape)
  - Datapath Acceleration Architecture hardware Queue Manager and Buffer Managers
  - Generic Interrupt Controller (GIC)
  - TrustZone Secure World/Non-Secure World partitioning IPs;
    - TrustZone Address Space Controller (TZASC)
    - TrustZone Protection Controller (TZPC)
    - TrustZone Memory Access (TZMA)

Memory Management Unit (MMU)

- MMUs translate virtual addresses into a physical address which are put onto the system bus
- Armv8 CPUs used in Layerscape SoCs offer two stage address translation
  - Virtual address (VA) -> Intermediate physical (IPA) -> Physical (PA)
- Important concepts; Process ID, Page Table, Translation Lookaside Buffer
  - The process running on the CPU is identified by Process ID (PID) Registers (updated by privileged software each time it schedules that process to run
  - The process can’t spoof its PID
  - PID is fed to MMU along with virtual address; MMU accesses page tables specific to that process
  - Translation Lookaside Buffer (TLB)
    - A PID aware cache of the page table entries of recently translated addresses
  - Page table is data structure containing mapping from VA→ PA
    - PID aware, also contains access permissions for the page (see example from ARMv8)
- MMU configuration creates a PID whitelist
  - A given PID can be blocked from accessing a 64KB page, given read access only, etc.
  - Pages can also be marked as No Execute.
  - Careful whitelist configuration reduces the probability of a corrupted transaction completing

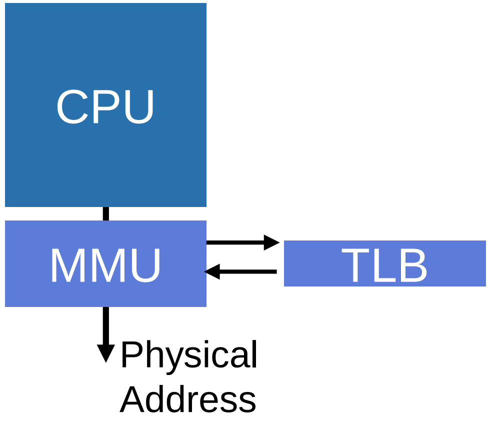

IO MMU

- IO MMUs simplify software development
  - 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, which the Layerscape SMMU supports

- 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, performing an important security function

- Access protection is a security function, but it is also a safety function
  - While software errors are assumed to be tested out of existence in ASIL systems, the reality is there will be more software bugs than hardware bugs. Software programming DMAs with incorrect source or destination addresses will trigger IO MMU errors.
  - Misconfigured non-CPU bus masters may try to access memory they aren't meant to access, triggering an memory access violation interrupt
  - Multibit corruption of packets on the interconnect may not be caught by parity. A corrupted address is highly unlikely to fall in a legal access window.

TrustZone Secure World Partitioning IP

- Creating hardware enforced separation between TrustZone Non-Secure World (the rich execution environment) and Secure World (trusted execution environment) is mainly a security feature, but like the MMUs and IO MMUs, permission based access to memory mapped regions also acts as a safety check.

- Most software will execute in Non-Secure World, and if any of it tries to access memory ranges owned by Secure World, the TrustZone partitioning Ips will block the attempt.

- Most CPU, platform, and IP block configuration registers are TZ SecureWorld access by default.

- The TrustZone Address Space Controller (TZASC) sits in front of the DDR controller(s), blocking Non-Secure World access to configured regions of DDR memory.

Datapath Acceleration Architecture

Software Portals & Isolation

- All Layerscape datapath resources (accelerators, QDMA, and Ethernet) are accessed through Queue Manager and Buffer Manager 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

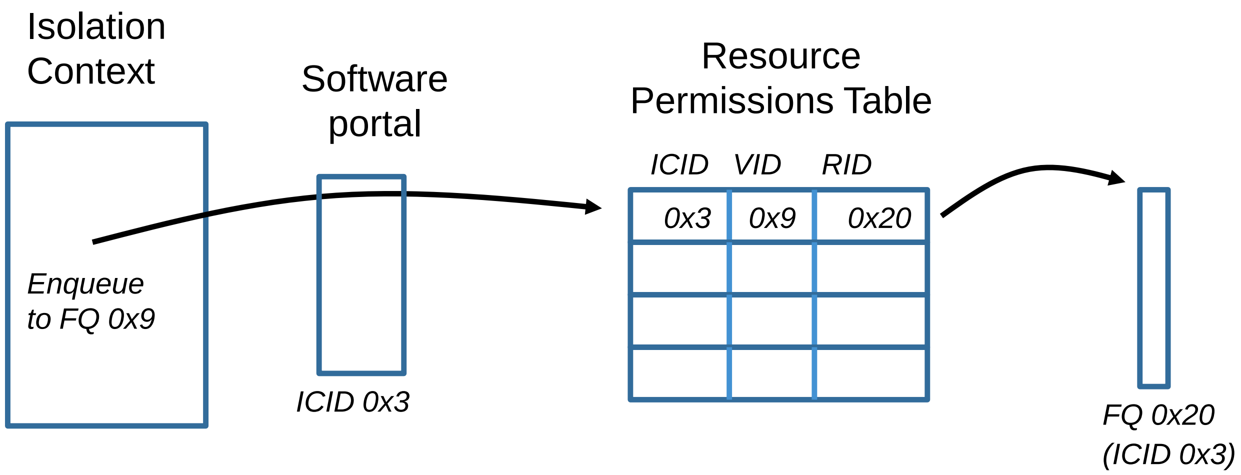

ICID	VID	RID
0x3	0x9	0x20

- This allows the datapath to virtualize (in hardware) all network interfaces and accelerators
- The network interfaces and accelerators temporarily (and unspoofably) take on the access permissions of the software partition/VM that generated the request
- Deliberate or accidental attempts to access Qman FQs of Bman buffer pools will be blocked and trigger an error

Network Interfaces

- Layerscape devices implement multiple Ethernet MACs, typically integrated into a large networking engine (Fman, WRIOP).
- The Ethernet MACs themselves include standard Ethernet frame checking (CRC) and in some cases, cryptographic data integrity, encryption, and replay detection (MACSEC).
- The Fman/WRIOP the MACs are embedded within provides network processing offloads, including;
  - Parsing & classification, with ingress policing and egress shaping
  - Interface virtualization
  - Protocol offload
- The Fman/WRIOP supports self test capability, including loop-back and link training tests.
- Large SRAMs in these engines (look-up/classification tables) are ECC protected. Smaller buffer memories holding transient frame data may be parity protected only.
- The Fman/WRIOP can generate interrupts for specific virtual interfaces. There is a single interrupt for signaling an unrecoverable error to safety software. The Fman/WRIOP does not directly trigger a Reset_Req.
- Ethernet ports can be configured to strip Ethernet headers & CRC, or deliver full L2 frame to software
  - Delivering full frame provides end to end CRC protection.(at the expense of software CRC checking)
  - Ports can also be configured to transmit a frame with software generated CRC for outbound end to end data integrity.

Network Protocol Usage Considerations

- Ethernet has CRCs to detect corrupted frames, however when a CRC error is detected, the frames must be discarded.
  - Statistics are maintained on the number of discarded frames, thresholds can be set for generating interrupts if too many frames arrive corrupted.
- IP (OSI layer 3) is also an unreliable protocol.
  - IPsec can be used to add cryptographic data integrity, encryption, and replay detection.
- OSI layer 4 options include UDP and TCP.
  - UDP/IP/Ethernet should be used where some packet loss is acceptable.
  - TCP/IP/Ethernet should be used where reliable transmission is required. If a portion of TCP data isn’t delivered due to Ethernet frame corruption, the sending TCP stack will retransmit the missing data with sequence information, allowing the receiving TCP stack to reassemble the complete message.
- Application layer communication can include a range of reliability features, including forward error correction, to make the make the loss of individual frames tolerable.
- A heartbeat protocol can run over an unreliable protocol, however the interval after which a missing heartbeat is considered an indication the system is not safe should be set large enough (and the interval between sending heartbeat messages small enough) that an occasional lost frame is tolerated.

Silent Data Corruption

- Silent data corruption is the undetectable, uncorrectable portion of the failures that can occur in a device that appears to otherwise be operating normally with in spec clock and power.
- These are essentially soft errors in sequential logic, and the probability of such silent data corruption events is documented in the LX2 FMEDA.

Freedom From Interference

Interference Channels and Resource Usage

MCP_Resource_Usage_4: The applicant has identified the available resources of the MCP and of its interconnect in the intended final configuration, has allocated the resources of the MCP to the software applications hosted on the MCP and has verified that the demands for the resources of the MCP and of the interconnect do not exceed the available resources when all the hosted software is executing on the target processor.

Note: The need to use Worst Case scenarios is implicit in this objective.

Shared resources within Layerscape create the potential for interference channels (vs idealized system), however in the hardware, these interference channels will exist at the nanosecond to microsecond level.

At the scale of a software function, interference channels are dictated by software scheduling.

Interference channels capable of causing side channel information leakage have been demonstrated with Spectre series of attacks.

Background Issue #1: Hardware Resource Contention

- Multicore Processors can execute several software applications at the same time because they have two or more processing cores that can each host and execute software applications. Several applications may therefore attempt to access the same shared resources of the MCP (such as memory, cache and external interfaces) at the same time, causing contention for those resources.
- Most MCPs have internal mechanisms such as “interconnects” to handle and arbitrate the demands for MCP resources, but the contention for shared resources between applications usually causes delays in access to the resources. These delays are a form of time interference between applications, which can cause applications to take much longer to execute than when executing on their own.

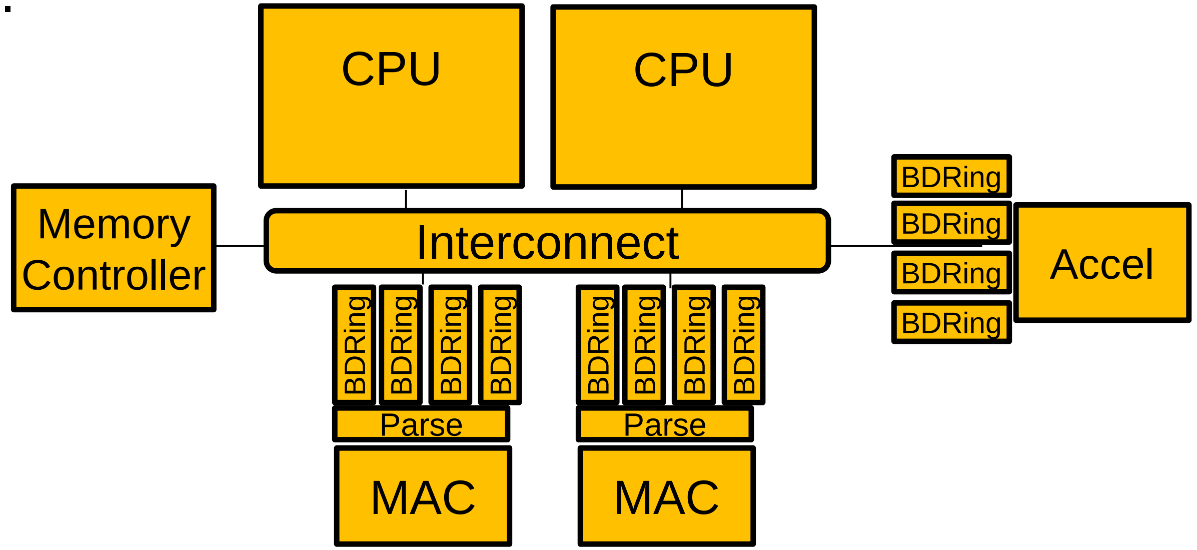

Background Issue #1: Hardware Resource Contention

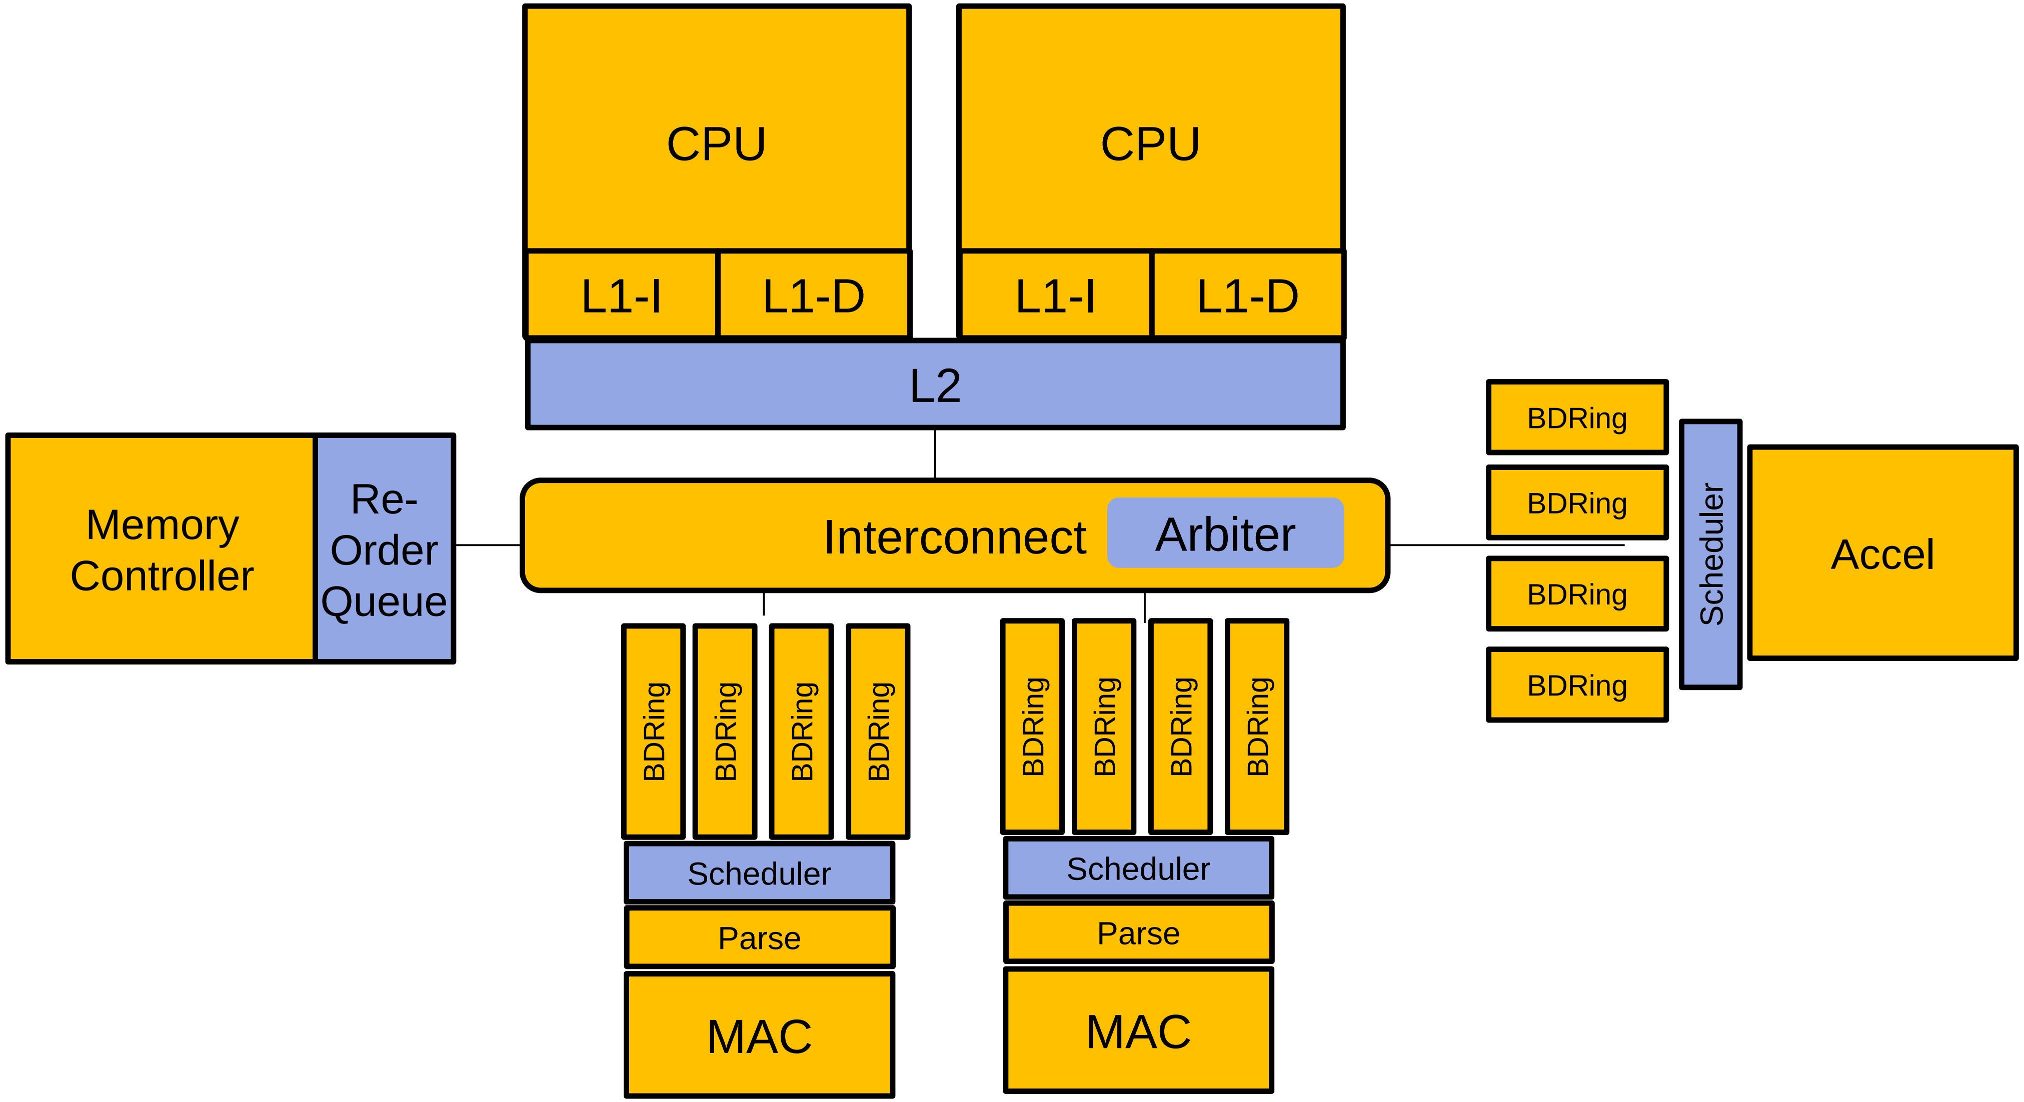

Background Issue #2: Software Resource Contention

- There could also be functional interference between applications via MCP mechanisms. Interference could also occur due to software components installed on the MCP, such as operating systems or software hypervisors.
- Interference between software applications executing on an MCP could cause safety critical software applications to behave in a non-deterministic or unsafe manner, or could prevent them from having sufficient time to complete the execution of their safety-critical functionality.

Not a Layerscape HW design issue. Resolving this issue is a core competency of some of our sponsors.

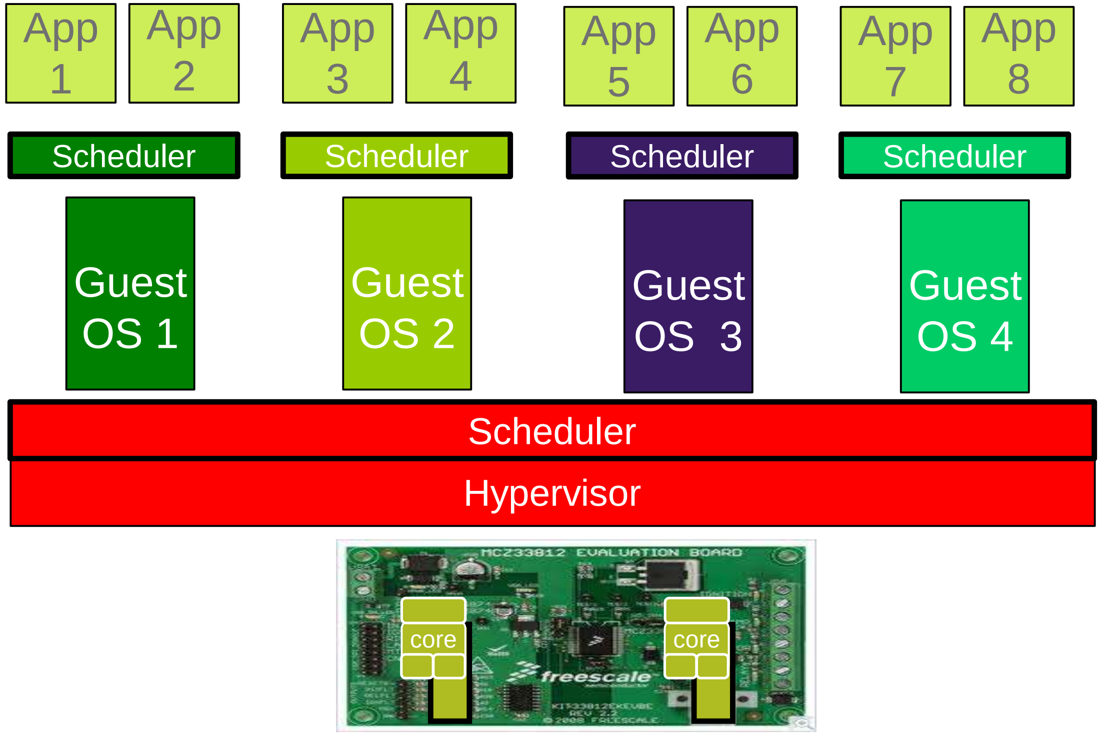

Unsupervised Asymmetric Multi-Processing

- 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

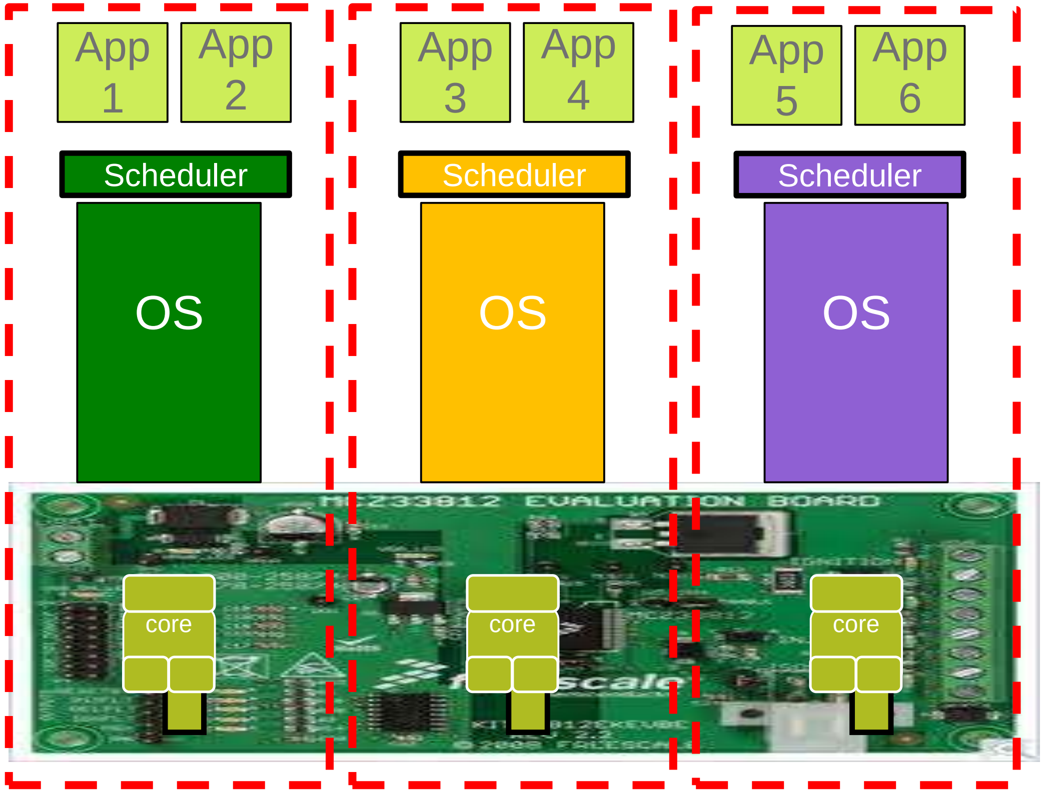

Symmetric Multi-processing (SMP)

An MCP software architecture in which a single operating system controls the execution of the processes on all the cores and may dynamically allocate sections of processes to run in parallel on separate cores.

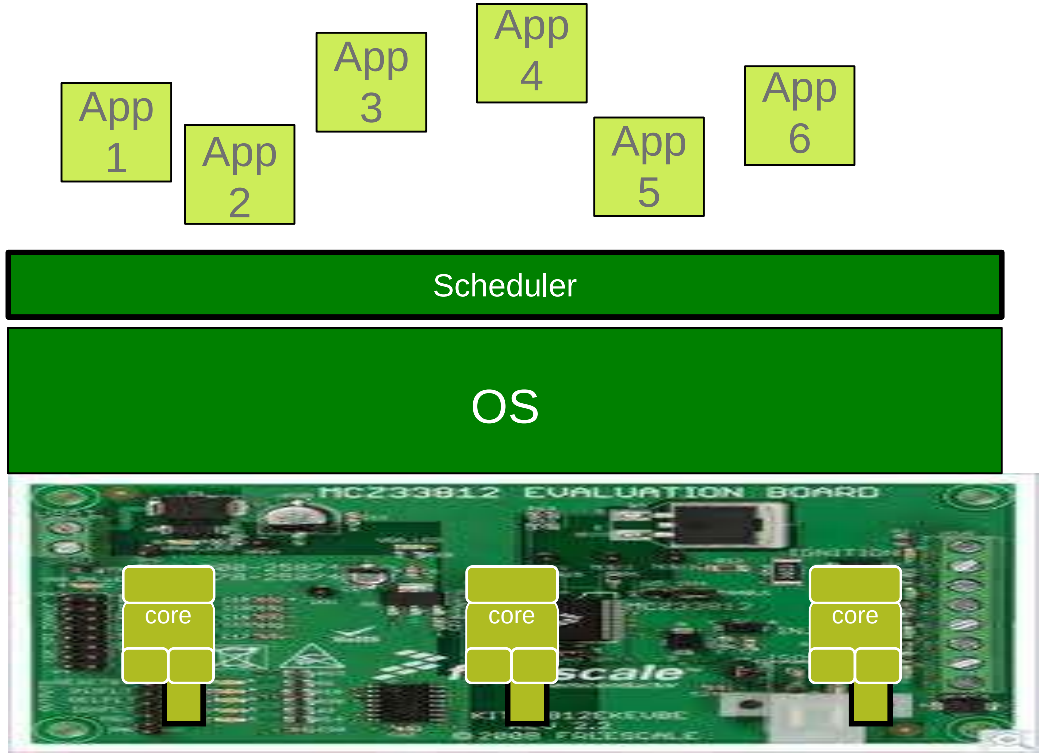

Bound MultiProcessing (aka Process Affinity)

- Processes aren’t dynamically scheduled on cores, they are bound to specific dedicated cores.
- When the process wants to run, it doesn’t have to wait for another process to yield the core.

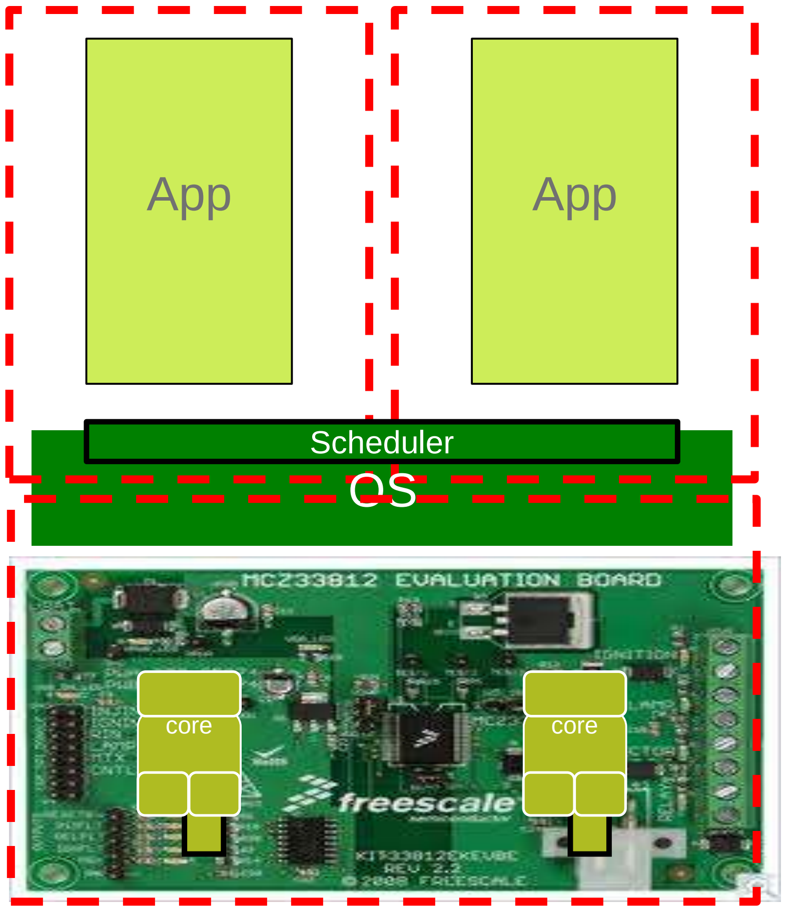

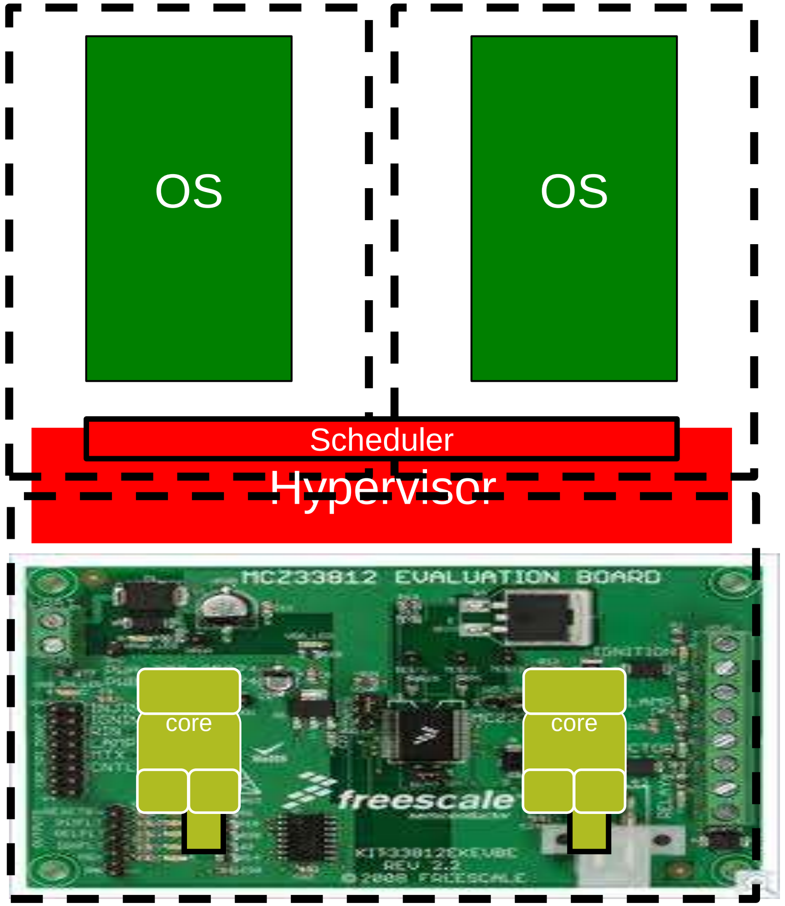

Determinism

- Determinism / deterministic: The ability to produce a predictable outcome generally based on the preceding operations and data. The outcome occurs in a specific period of time with repeatability.
- (From DO-297/ED-124).

Contributors to non-determinism in Layerscape based systems:

- SW
  - OS/HV scheduling
  - Separation kernels perform time sensitive scheduling
- HW
  - Branch prediction
  - Prefetcher
  - Caches
  - Load-on-store collisions
  - Snoops
  - Interrupt latency
  - DDR page hits/misses/collisions
  - Contention

- Presentation on Layerscape Determinism @
- MCFA Face to Face Workshop Presentations & Meeting Minutes : 2017_MCFA_Presentations

Critical Configuration Settings

- Those configuration settings that the applicant has determined to be necessary for the deterministic execution of the software or any settings that, if inadvertently altered, could change the behavior of the processor so as to cause the hosted software to no longer comply with its requirements. (See objectives MCP_Resource_Usage_1 and MCP_Resource_Usage_2.).

- Configuration registers are in CCSR (Configuration, Control, and Status Register) and DCSR (Debug Control and Status Register) space. DCSR space is not included in public documentation.

Layerscape Critical Configuration Registers include registers controlling;

- Pin muxing
- Clock control & generation
- Power Mgt
- Initiator & Target Identification, Arbitration
- Memory Space Access control
- Security violations

- Very few registers in Layerscape can be physically locked, however the majority of critical configuration registers are accessible only by TrustZone Secure World software.

Robust Partitioning: Resource

- Robust Resource Partitioning (adapted from DO-248C / ED-94C and DO-297 / ED-124)

Achieved when:

- Software partitions cannot contaminate the storage areas for the code, I/O or data of other partitions.
- Software partitions cannot consume more than their allocations of shared resources.
- Failures of hardware unique to a software partition cannot cause adverse effects on other software partitions.

- Note: Software that provides partitioning should have at least the same DAL as the highest DAL of the software that it partitions.

NXP refers to this as Logical Partitioning, relying on memory access controls

- CPU MMUs
- Platform IO MMU
- DPAA resource partitioning
  - Virtual network interfaces
  - Virtual accelerators
- Layerscape SoCs, particularly DPAA2 SoCs, have excellent support for Robust Resource Partitioning

Robust Partitioning: Timing

- Robust Time Partitioning (on an MCP) is achieved when, as a result of mitigating the time interference between partitions hosted on different cores, no software partition consumes more than its allocation of execution time on the core(s) on which it executes, irrespective of whether partitions are executing on none of the other active cores or on all of the other active cores.

- Robust Time Partitioning is largely the domain of the kernel
- Ecosystem partners offering Separation Kernels support these requirements

Layerscape in Automotive

- Highest CPU and IO performance SoCs in NXP
- Scalability – 1-16 ARM core SoCs
- Quality & Longevity – Best quality available in high performance processing. Many devices already on 15 year longevity program.
- Safety – We've demonstrated safety for mil/aero and other critical infrastructure applications. Working to prove ASIL-B equivalence with auto-centric collateral (FMEDA, Safety Manual).
- Security – Secure Boot, Secure Debug, Hardware Enforced Partitioning & Virtualization
- Software – SDKs with a very PC-like look & feel. Broad support in Linux, history of working with WindRiver, GHS, and QNX.

SECURE CONNECTIONS
FOR A SMARTER WORLD