


QORIQ LAYERSCAPE LS2085/88 MULTICORE COMMUNICATIONS PROCESSORS FOR THE FUTURE OF NETWORKING



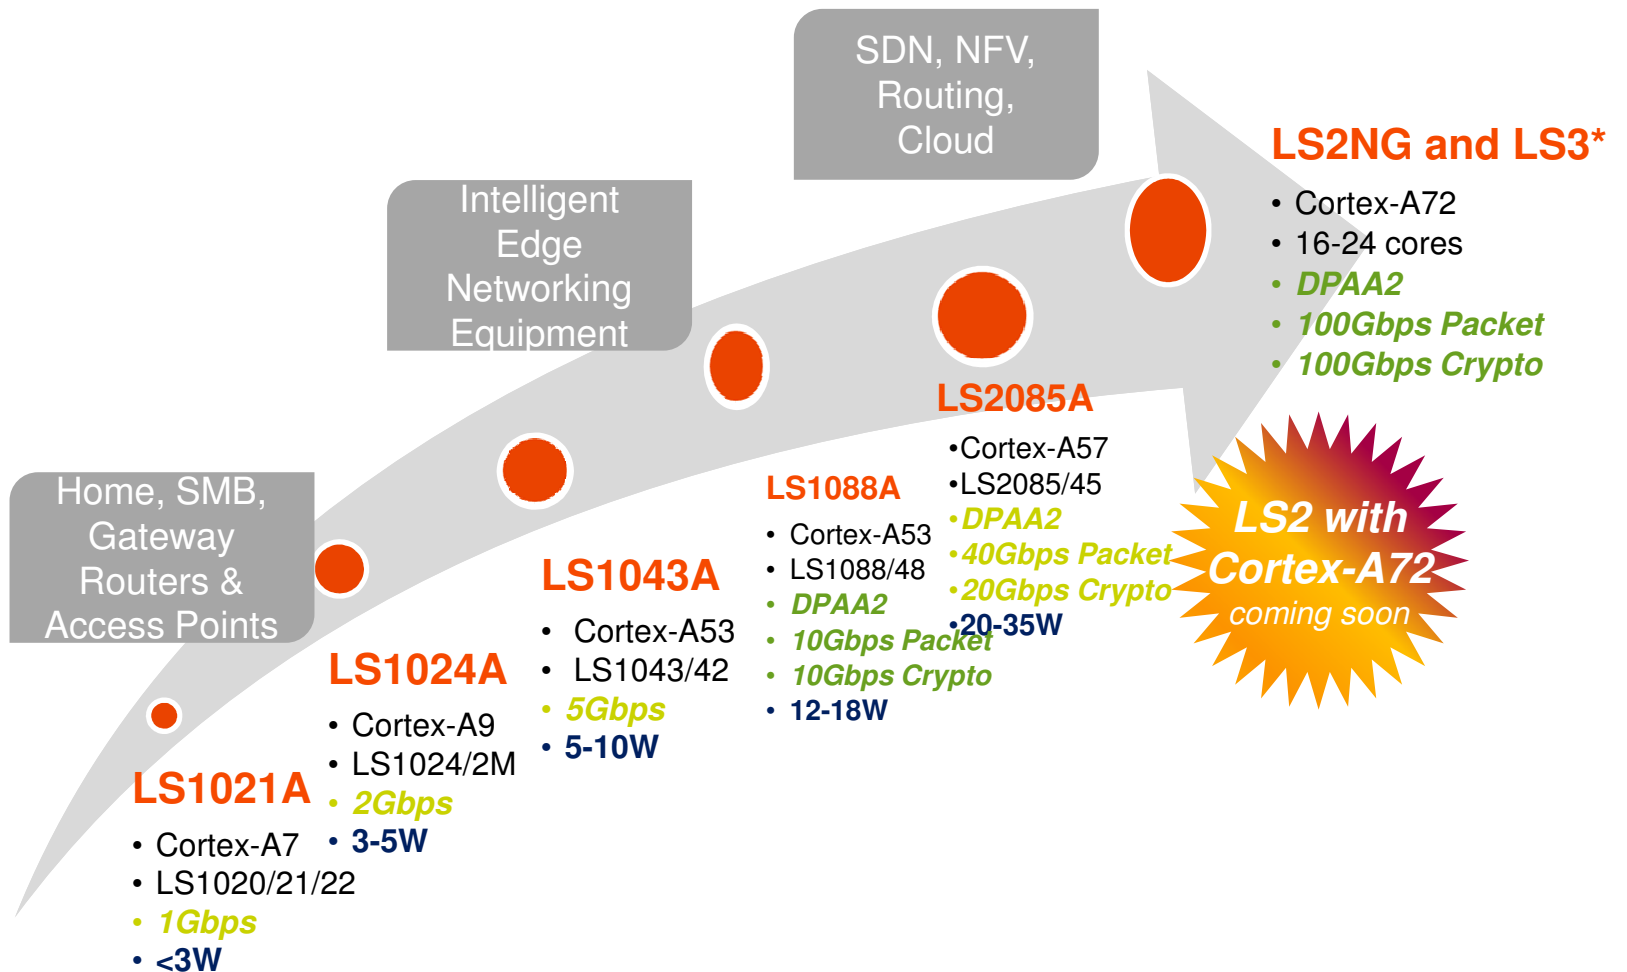
SECURE CONNECTIONS
FOR A SMARTER WORLD

Agenda



- Block diagram, new features
- ARM® related information
- Serdes information
- SoC infrastructure blocks
- LSDPAA introduction
- Use cases

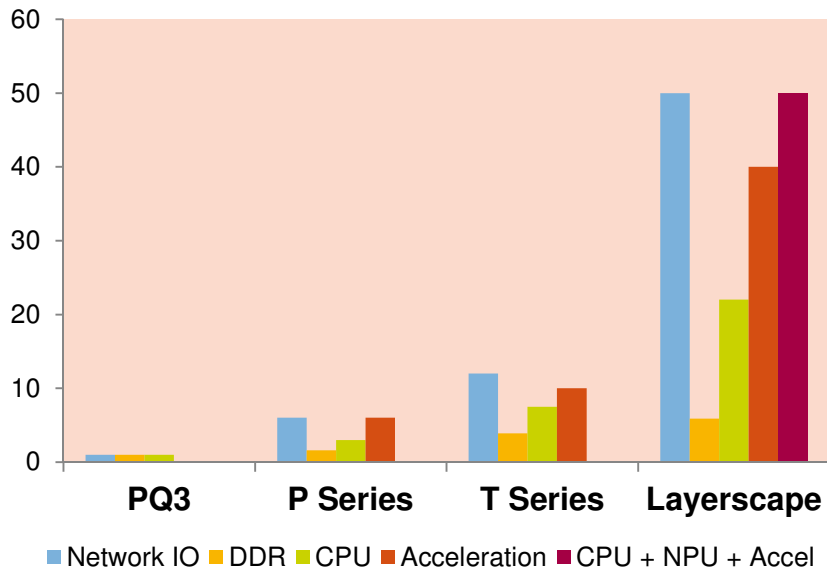
Continued Momentum: SoC Optimization



First to announce the extremely low-power Cortex[®]-A72 core in a networking processor



Layerscape : A New Architecture for a New Network



Many-core processor approach is not sustainable due to power, software complexity and integration costs



Need to provide right mix of high performance and programmability

MUST HAVE:

Advance Packet Processing

- Tightly coupled accelerators called as C functions
- H/W preloaded task state, headers, stack frame
- Customer programmable
- Run-to-completion model using standard C (C99)



3-4x Performance
over general purpose cores
in a lower power envelope



LS2085A's (AIOP) Packet Acceleration Benefits

1

Enterprise

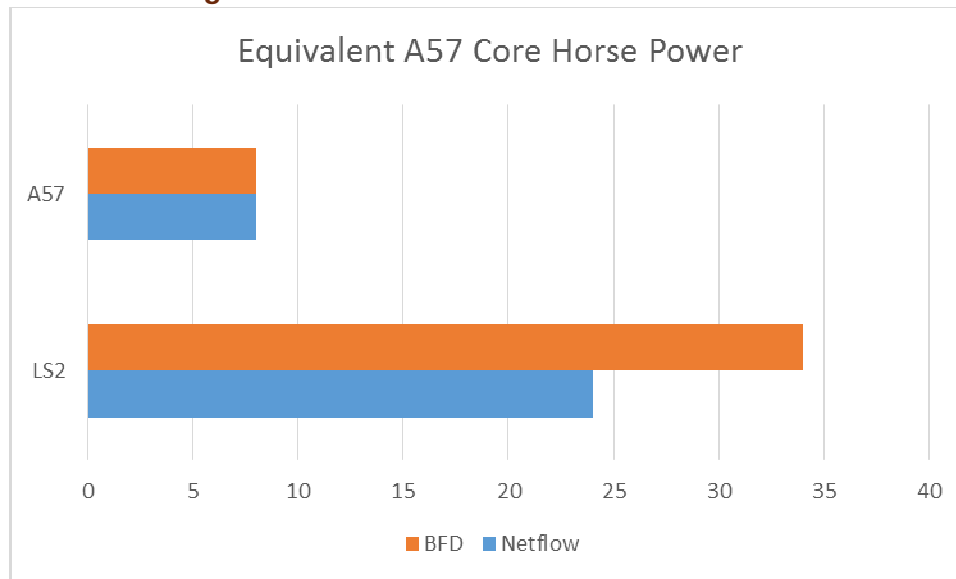
- SDN Switching
- NFV Appliances
- Enterprise Switching

2

Industrial and Control Applications

3

Storage



LS2085A BFD Performance	
GPP Performance	3Gbps
AIOP Performance (3.3ms on 25K sessions)	9.7Gbps
Equivalent A57 Core performance	34 Cores!

LS2085A Netflow Performance	
GPP Performance	10Gbps
AIOP Performance	20Gbps
Equivalent A57 Core performance	24 Cores!

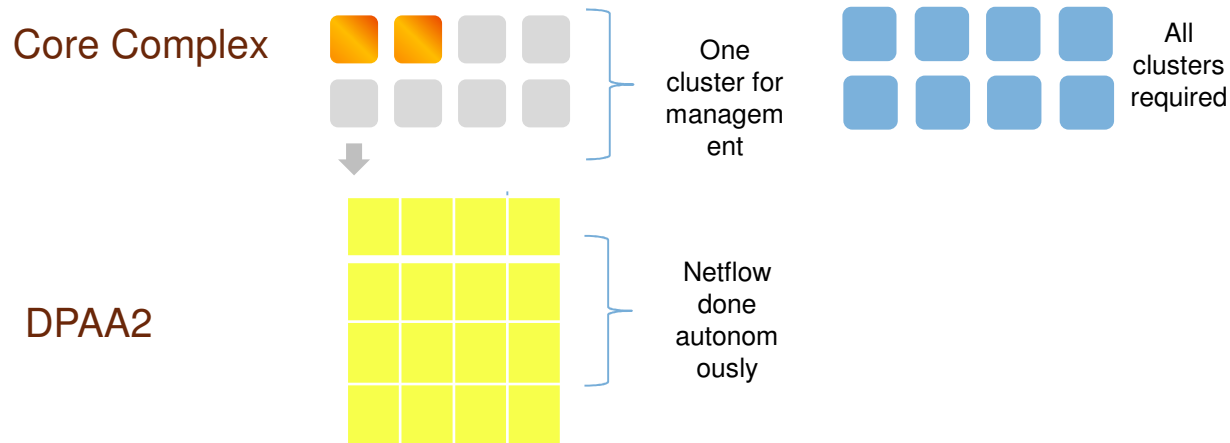
The LS2 device featuring AIOP behaves much better than an 8-core device (without the power dissipation penalty)!



Value of DPAA2: Power Efficiency

	LS2	A57-based solutions
Cores	8 @ 2GHz	8 @ 2.0-2.4GHz
Packet Acceleration	DPAA2 (16xe200)	Targeted HW acceleration

Use Case Netflow 20Gbps

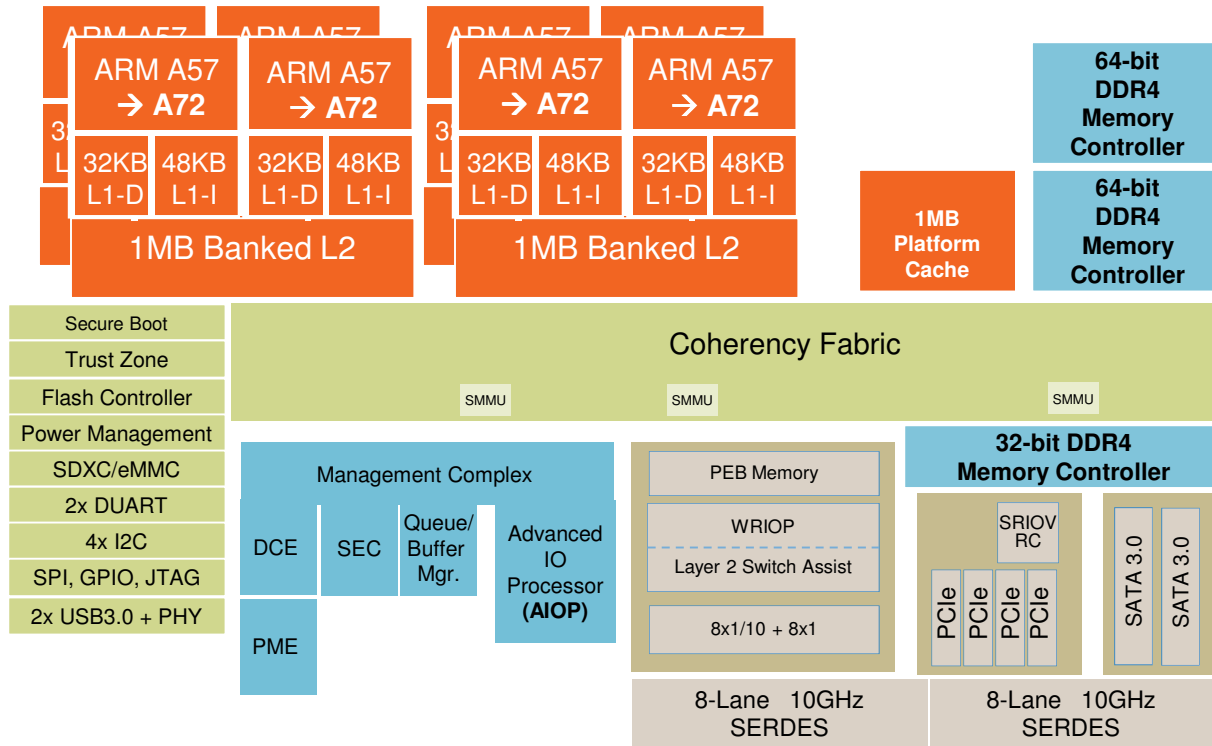


Power (GPP)	10W	40W
AIOP Power	10W	0W

...yields a >2x power advantage!



Product Definition: LS208X Platform



General Purpose Processing

- 8x ARM® A57 CPUs, 64b, 2.0GHz
 - 1MB L2 cache
- **Rev2: 8x ARM A72 CPUs**
- HW L1 & L2 Prefetch Engines
- Neon SIMD in all CPUs
- 1MB L3 platform cache w/ECC
- 4MB Coherent Cache
- 2x64b DDR4 up to 2.4GT/s

Accelerated Packet Processing

- 20Gbps SEC- crypto acceleration
- 10Gbps Pattern Match/RegEx
- 20Gbps Data Compression Engine

Express Packet IO

- Supports 1x8, 4x4, 4x2, 4x1 PCIe Gen3 controllers
 - SR-IOV support, Root Complex
- 2 x SATA 3.0, 2 x USB 3.0 with PHY

Network IO

- Wire Rate IO Processor:
 - 8x1/10GbE + 8x1G
 - XAUI/XFI/KR and SGMII
 - MACSec on up to 4x 1/10GbE
 - Layer 2 Switch Assist

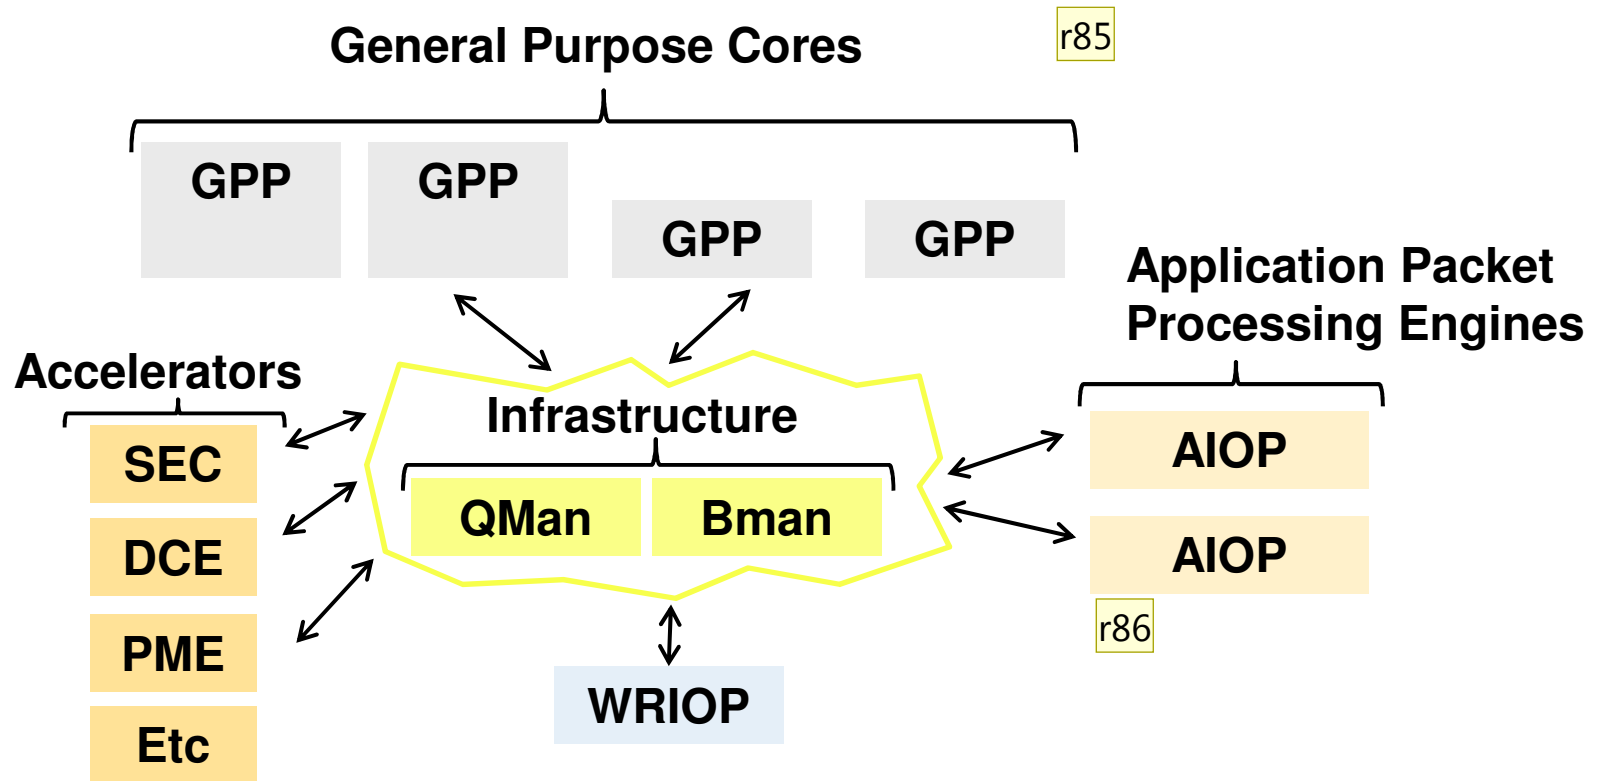
Other Parametrics

- 37.5x37.5 Flipchip
- 1mm Pitch
- 1292pins

Datapath Acceleration

- **SEC**- crypto acceleration
- **DCE** - Data Compression Engine
- **PME** – Pattern Matching Engine
- **L2 Switching** -- via Datapath Acceleration Hardware
- **Management Complex** – Configuration Abstraction

DPAA2 Hardware



Slide 8

r85

Use different slide or eliminate big/little, use GPP

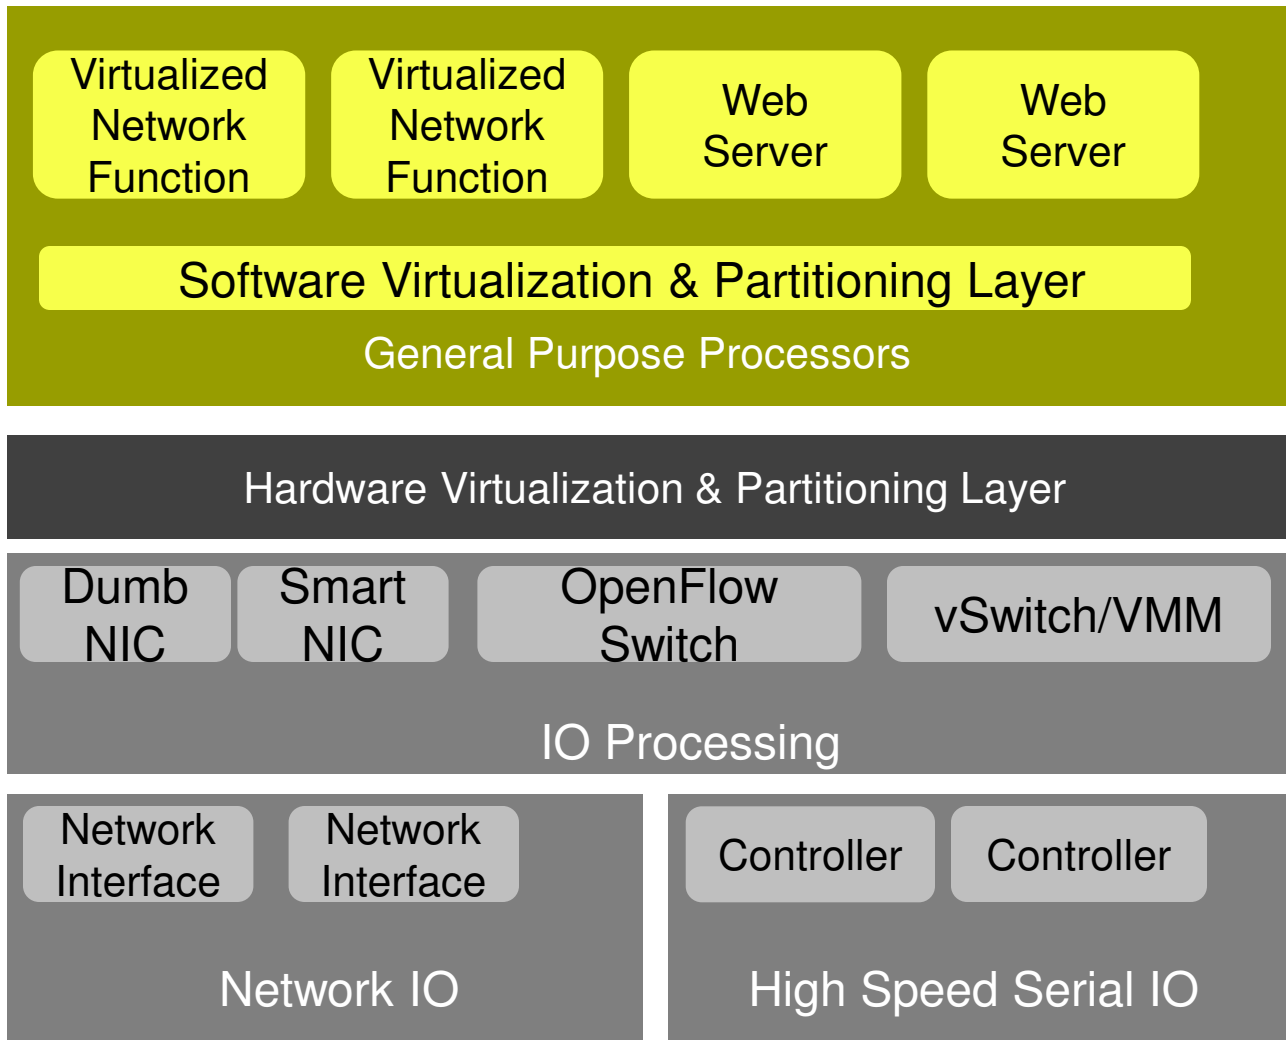
ra6722; 02/12/2014

r86

If not spelling out Advanced IO Processor, call these C-programmable packet processors

ra6722; 02/12/2014

QorIQ Layerscape architecture Concept: Functional View



High Performance General Purpose Processors with HW enforced/accelerated virtualization capabilities.

VMs in the GPPs perceive direct access to private IO and acceleration resources; IOs and accelerators adopt the permissions of the VMs using them.

Virtualized IOs & accelerators instantiated as network objects useful to the VMs; such as normal and smart NICs, vSwitch.




Major New Features Summary

- 64-bit ARM® core
8 x A57 cores, 2GHz
Rev2: 8 x A72
- DPAA2
Major paradigm/software model shift
- DDR4 only
- SATA x2 gen3
- USB x2 SuperSpeed/High Speed with PHY
- PCIe Gen 3
- QSPI, QDMA

ARM cores

- 8 ARM cores: 4 clusters, 2 cores/clusers
- Up to 2GHz.
 - Each cluster can run at different speeds
- General purpose processor ARMv8 64-bit architecture
- Superscalar, dual symmetric single-cycle ALU/shift
- Advanced SIMD and floating-point unit
- 64-byte cache line size
- L1 cache: 48KB I-Cache, 32KB D-Cache
- L2 Cache: 1MB with ECC
- 40-bit physical addressing
- Hardware page table walk



r12

General Purpose Processors: ARM v8

ra6722; 02/12/2014

Rev2: Highest Performance Cortex® A72 ARM® v8 Processor

- **Highest Single-Threaded Performance**

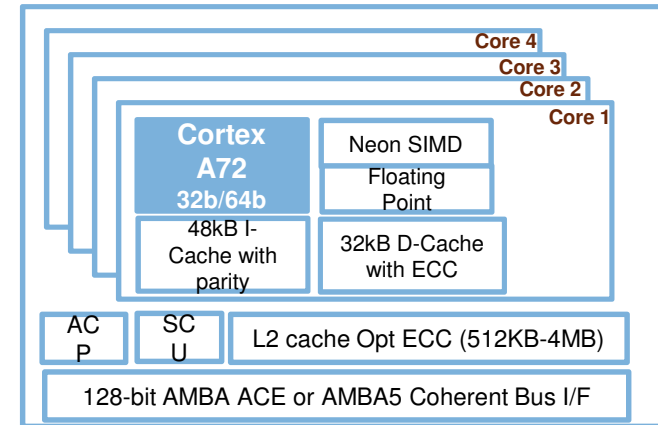
- Lower power enabling maximum performance in thermal limit
- Large performance increase across integer, memory-streaming, float

- **Significant Advancements in Power Efficiency**

- 17.4% power reduction from Cortex-A57
- 6% ~ 10% cluster area reduction lowers static power


- **Enhanced Multicore Scalability**

- Larger L2 for optimizing SDN/NFV applications



Performance Improvement A72 relative to A57

L2 cache in LS2088A set to 512KB



Improved performance in floating point, branch prediction and data pre-fetch)

High Speed Serial Interface: SERDES

- Two x 8-lane 10Gbps SERDES: total 16 lanes
Shared by
 - 16 x Ethernet Controllers
 - 4 x PCIe Controllers
 - 2 x SATA controllers (Gen3, 6Gbps)
- Ethernet Interfaces:
 - 10Gbps: 8 x XFI or 2 x XAUI
 - 1Gbps: 16 x SGMII, 4 x QSGMII(16 ports)
- SGMII base line at 1.25Gbaud for 1Gbps Ethernet
SGMII can also runs at 3.125Gbaud for 2.5Gbps Ethernet
QSGMII runs at 5Gbaud

DDR controllers

- Supports DDR4 only
- Two 64-bit DDR controllers with ECC, 2.1Ghz data rate
- One 32-bit DDR controller with ECC, 1.6Ghz data rate
 - Intended for AIOP only.
 - Can be power gated if not needed
- Support x4, x8, and x16 memory widths
 - Programmable support for single, dual, and quad ranked devices and modules
 - Support for both unbuffered and registered DIMMs
 - 4 chip-selects per controller
 - Max amount of DDR memory supported
 - 16Gb is highest density approved by JEDEC
 - x4, we can put $64/4=16$ device for one CS
 - 16Gb*16/8=32Gbytes/CS
 - 4 Chip select =128Gbytes
 - Caution: For 2166Mhz, we might limit to 2 CS. We need to wait until silicon.
 - And we have two controllers for main memory
- Support self-refresh mode for implementation of battery backed main memory

r17

Slide 14

r17

- Two 64b Main DDR controllers with ECC, etc
 - Used by whole SoC
- One 32b DPAA2 DDR controller with ECC, etc.
 - Optional. When enabled, for use by AIOP only

Don't need the CS details. Main DDR capacity is up to 256GB. DPAA2 DDR capacity is up to ?

ra6722; 02/12/2014

PCI Express

- Four Gen 3 PCI Express controllers
- Power-on configuration options allow root complex or endpoint
- The physical layer operates at 2.5, 5, or 8 Gbaud data rate per lane.
- Link Width
 - PEX1, PEX2, PEX4: x1/x2/x4 r18
 - PEX3: x1/x2/x4/x8
- Both 32-bit and 64-bit addressing, 256-byte maximum payload size
- Inbound INTx transactions
- Message Signaled Interrupt (MSI and MSI-X) transactions
- PCI Express controller 3 supports end-point SR-IOV
 - Two physical functions
 - 32 virtual functions per physical function
 - Eight MSI-X per virtual function

Slide 15

r18

This can be abstracted a bit...

All controllers support

x

y

z

3 controllers support up to x4 link width

1 controller supports up to x8 link width

- This controller supports end point SRIOV

- 2 PFs, 32 VFs per PF, 8 MSI-X per VF

ra6722; 02/12/2014

USB r26

- Two USB controllers with integrated PHY
- Complies with USB specification, Rev. 3.0
- Supports super-speed (5 Gbps), high-speed (480 Mbps), full-speed (12 Mbps), and low-speed (1.5 Mbps) operations
- Both controllers support operation as a stand-alone USB host controller
 - Supports USB root hub with one downstream-facing port
 - Enhanced host controller interface (EHCI)-compatible
- Both controllers supports operation as a stand-alone USB device
 - Supports one upstream-facing port
 - Supports six programmable USB endpoints

- SYSCLK is also used as reference clock for USB PHY

SYSCLK requirements for USB PHY

Parameter/condition	Symbol	Min	Typ	Max	Unit	Notes
SYSCLK frequency for USB 3.0 superspeed only	f_{SYSCLK}	—	100, 125	—	MHz	—
SYSCLK frequency for USB 2.0 only or mixed 2.0/3.0	f_{SYSCLK}	—	100	—	MHz	—

Slide 16

r26

Either make this overview level or training level.

ra6722; 02/12/2014

SATA

- 2xSATA controllers support 1.5Gbps, 3.0Gbps, 6.0Gbps
- AHCI 1.3 Compliant
 - Define both hardware behavior and programming Interface
 - Use standard driver with minimal Freescale specific code
- 2 x AHCI controller
 - In AHCI spec, it supports 1-32 ports/controller.
 - In LS2, we support one port/AHCI controller
- pp2c, pp3c, pp4c, pp5c should be configured based on the speed of SYSCLK
 - These registers set OOB timing parameters

r20

Slide 17

r20

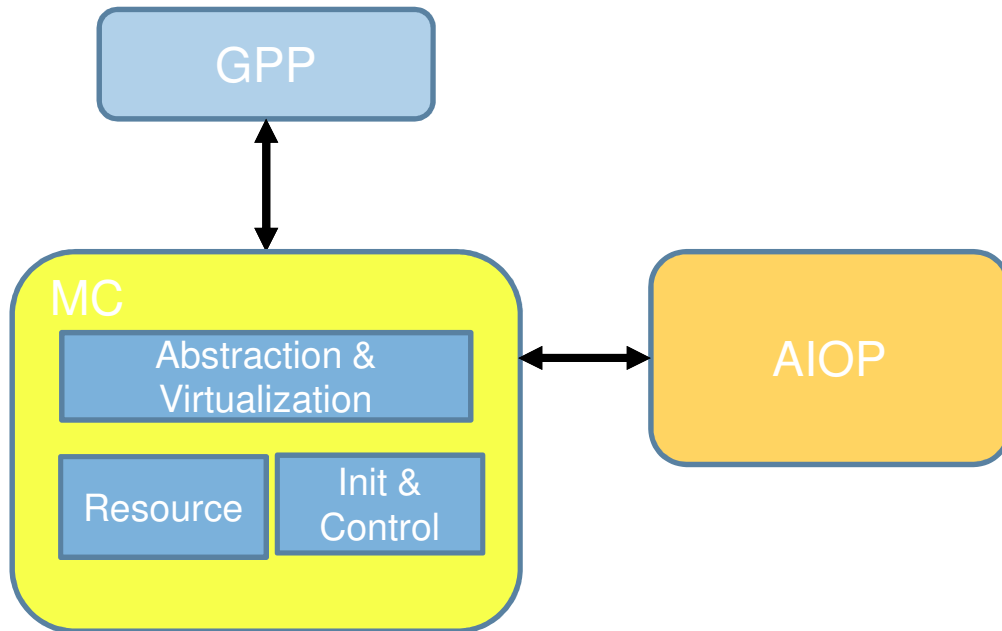
If this is an overview slide, drop the register/timing info. If this is a training/config slide, add the rest of the details.

ra6722; 02/12/2014

MC (Management Complex)

- MC abstracts Layerscape architecture hardware into software-managed objects that are:
 - Simpler to use than directly managing the hardware
 - Customers do not need to program the hardware directly
 - Application-oriented in terminology and use
 - Based on concepts generally familiar to programmers and system architects

Value of the Management Complex: Ease of use “Agent” between the GPP and the AIOP



- Abstraction and Virtualization
 - Allocates and configure elements: Virtual NI, switches networking functions/accelerators
 - Dedicated vs. Virtual
- Resource Allocation
 - No more awareness of queues, buffers, pointers (QBMan, CTLU, Accelerators)
- Initialization and Control
 - Session establishment and tear-down
 - Device discovery
 - Policy elements

Reduce awareness and head-aches of working with AIOP

Slide 19

r28

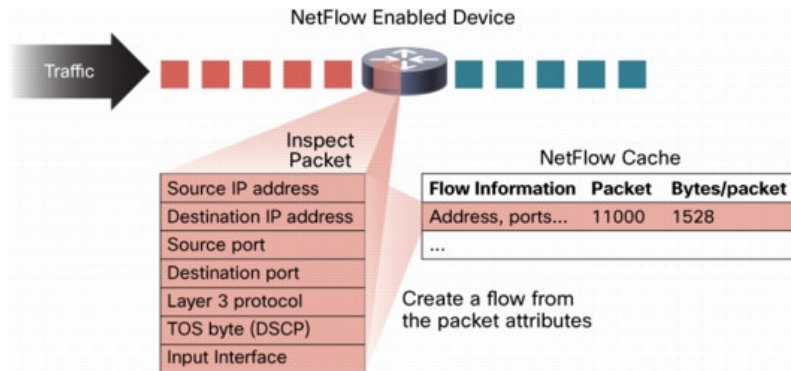
Duplicate.

ra6722; 02/12/2014

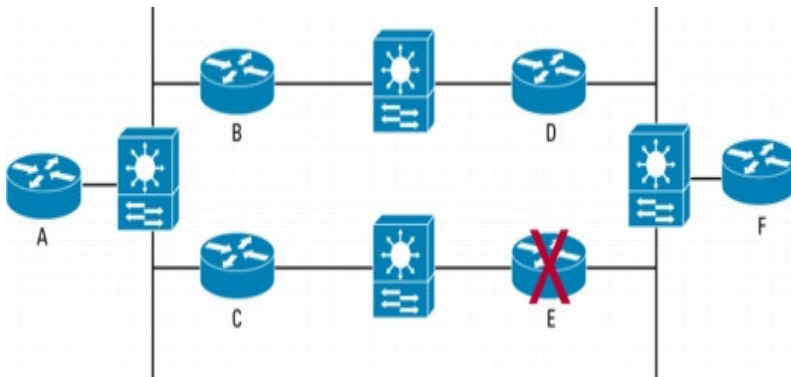
Next-Generation Data Centers Networking

LS208X eases the implementation of analytics that will optimize the use of the data center network infrastructure

Netflow Use Case



BFD



Implementation Details

- Key tool for network trouble shooting, capacity planning and anomaly detection
- Implementation in the LS208X saves the cost of expensive ASICs normally used to inspect all packets.
- LS208X's AIOP Performance is 20Gbps
- Implements Cisco netflow RFC 3954 in AIOP.
- Support for IPV4 and IPV6
- ACL support for selective monitoring.
- Support multiple observation domains and observation points.
- Support millions of flows (200K flows/s).

- Provide failure detection in less than 1 sec.
- Performance Implementation: 9.7Gbps at 3.3ms/25K flows detection rate (which is less than 50% the capacity of LS2085A).
- Confirms to RFC 5880/81/82/83.
- Support for IPV4/V6 single hop and multi hop peer failure detections.
- Support BFD Authentication.

Slide 20

r83

This slide should come earlier to set the stage for the LS2 value proposition. Too much redundant info with other Netflow/BFD slides.

ra6722; 02/12/2014



SECURE CONNECTIONS
FOR A SMARTER WORLD