



LS2085/8A Freescale's new QorIQ Layerscape Communications Processor

Zheng (John) Xu | Chief Architect, Freescale DNG


Aug. 2015



External Use

Freescale, the Freescale logo, Altivec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. Airfast, BeeKit, BeeStack, CoreNet, Flexis, Layerscape, MagniV, MXC, Platform in a Package, QorIQ Qonverge, QUICC Engine, Ready Play, SMARTMOS, Tower, TurboLink, UMEMS, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2014 Freescale Semiconductor, Inc.



- 
- Introduction
 - LS2085/8A Details
 - Core and Platform
 - Memory Subsystem
 - DPAA 2.0 Subsystem
 - Mgt Complex and LS Objects
 - WRIOP
 - AIOP
 - SEC, PME, DCE
 - QorIQ LS series Solutions
 - Example Applications
 - Performance

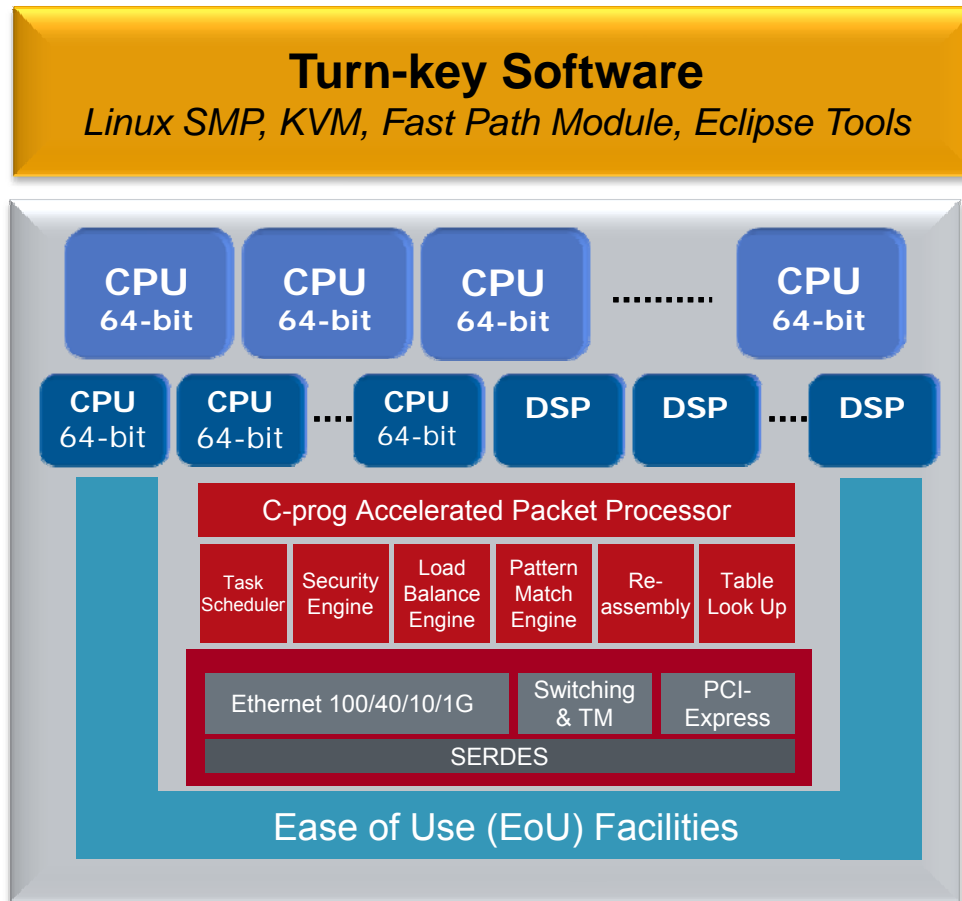
Introduction

- Traditional multicore processor approach is not sustainable for data plane processing due to performance, power and integration cost.
 - SDN demands high performance/low power data plane switches for reduced CAPEX and OPEX.
 - Scalable and deterministic performance under high stress conditions are required in data plane processing.
 - On-die hardware acceleration required to meet increasing network throughput requirements.
 - IPSec for security-on-wire
 - Data compression/decompression to reduce network throughput in data center
 - Deep Packet Inspection for Intrusion Detection
- Data plane processing needs to be programmable
 - SDN calls for highly programmable data plane switches to adapt to the ever-changing open networking standard (i.e. OpenFlow)
 - NFV requires flexible virtual switching for service chaining with evolving reference software implementation (i.e. OpenVSwitch)
 - Network virtualization overlay needs to support multiple protocols (i.e. VXLAN, NVGRE)



Freescale addresses SDN and NFV challenges with the new QorIQ LS family - an architecturally balanced multicore processing framework with highly programmable packet processing acceleration capabilities, focused on ease-of-use and performance-per-watt enhancements.

Advancing Multicore Design with Differentiated Solutions: QorIQ LS Series



1- Core Agnostic (ARM, Power Arch)

- ARM V8 Product Roadmap
- Small / Large footprints

2- Scalable Acceleration Elements

- Sized to Application Needs
- Turn key or C-programmable
- Wire rate I/O switching & TM

3- Ease of Use

- Real Time Monitoring / Debug
- SW Management utility
- I/O virtualization

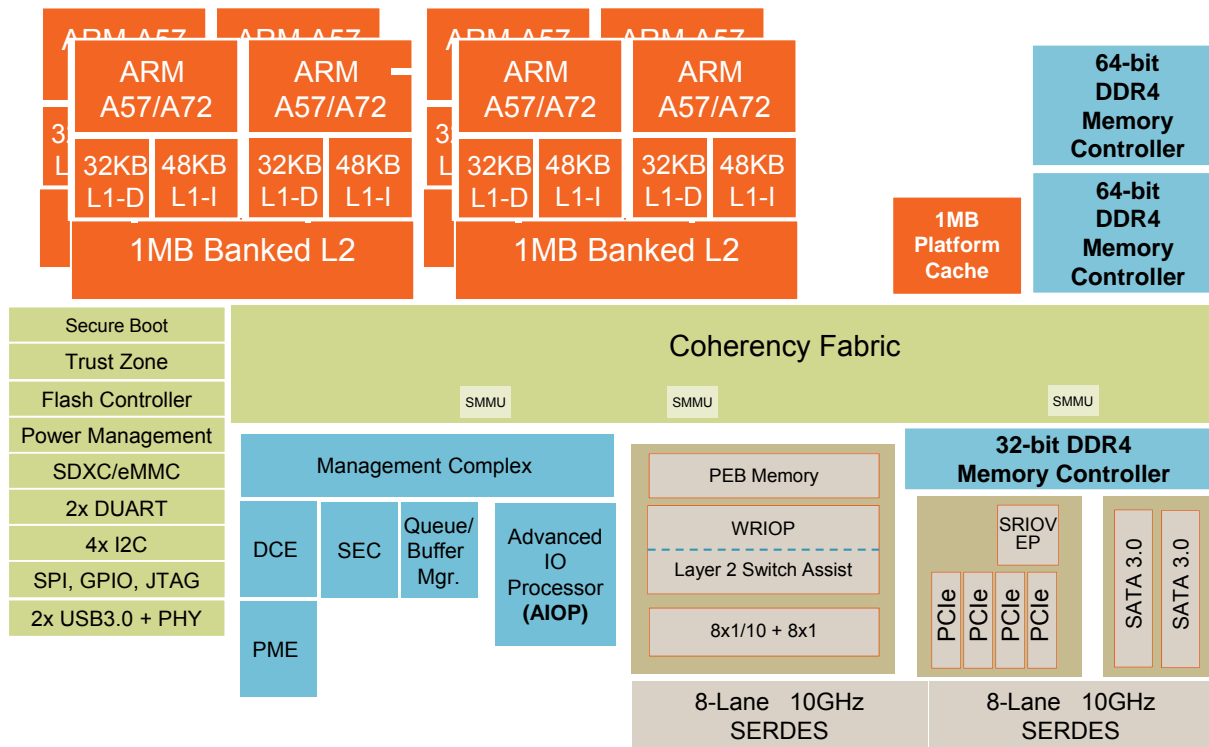
4- Turn-key Software

- Fast path modules
- Linux / BSP
- Hypervisor: KVM
- Eclipse-based tools

64-bit Multicore SoC Platform

- a) Industry standard tools & C-programmability
- b) Abstracts I/O and Acceleration
- c) Turn-key / Production-quality software

Introducing LS2085/8A, Current flagship of the QorIQ LS Family



General Purpose Processing

- 8x 64-bit ARMv8 A57 (LS2085A) or 8xA72 (LS2088A) CPUs up to 2.0GHz
- 1MB L2 cache in each 2xA57/A72 core cluster
- HW L1 & L2 Prefetch Engines
- Neon SIMD in all CPUs
- 1MB L3 platform cache
- 2x64b DDR4 up to 2.133GT/s

Accelerated Packet Processing

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

High Speed IO

- Supports 1x8, 4x4, 4x2, 4x1 PCIe Gen3 controllers
 - SR-IOV, End Point, 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/QSGMII
 - MACSec on up to 4x 1/10GbE
 - Layer 2 Switch Assist


Datapath Acceleration

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

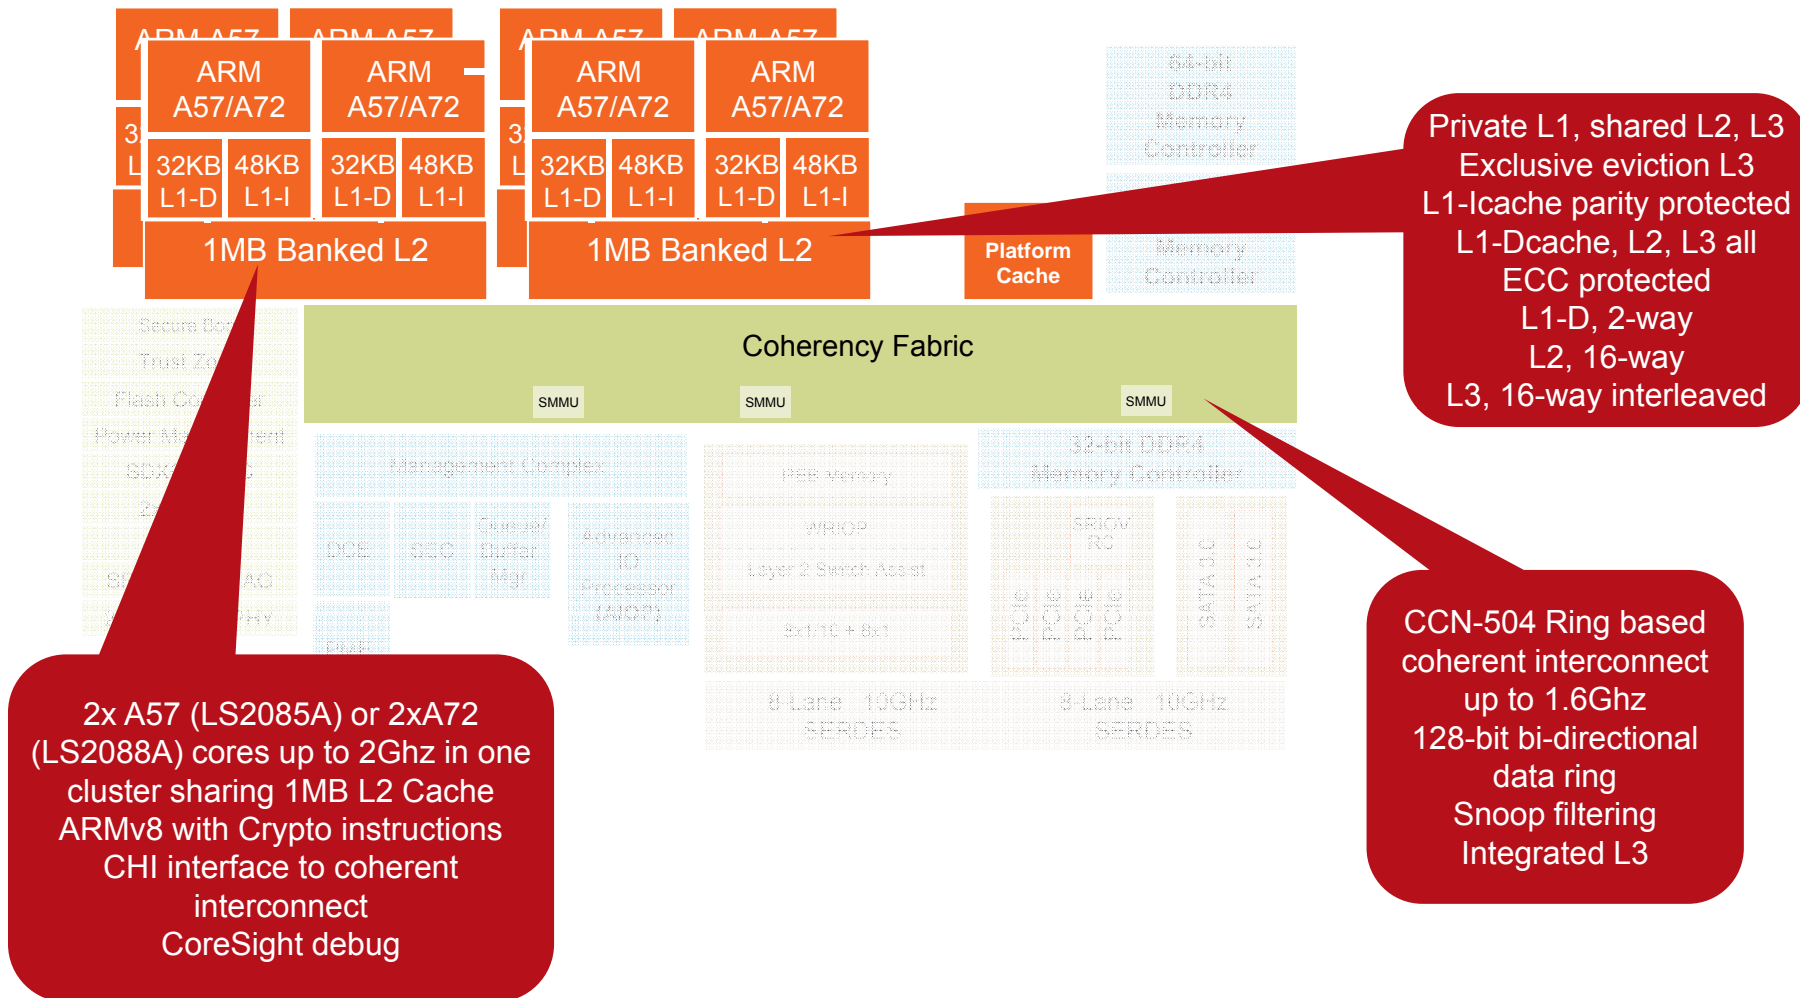
LS2085/8A Floorplan and Physical Metrics



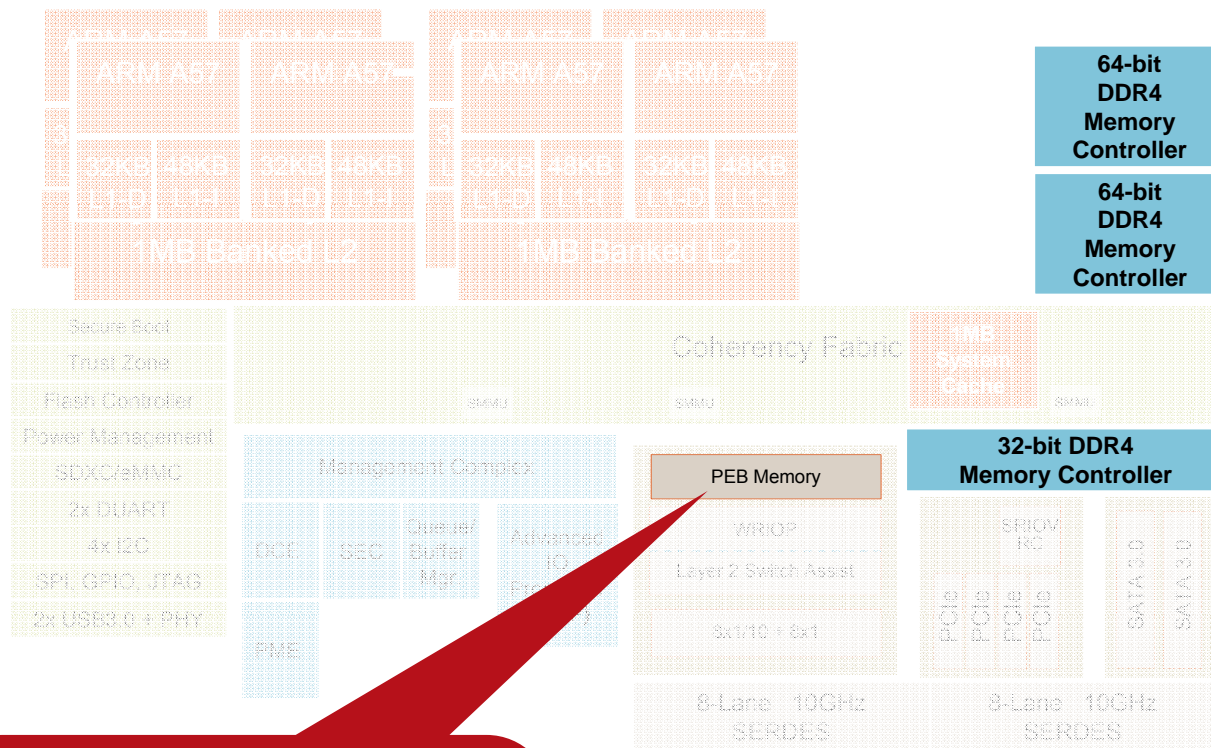
- 2.2 billion on-die transistors with TSMC 28nmHPM processing technology and 10 metal routing layers
- 1292-pin FC-PBGA 37.5x37.5mm package, 0.8mm pitch
- Operating junction temperature range, 0-105C
- Operating core logic DC power supply, Nominal 1V with +/-3% tolerance
- Preliminary power measurement
 - 40W TDP w/o AIOP @85C
 - 45W TDP w AIOP @85C



LS2085/8A Core, Platform and Cache Hierarchy



LS2085/8A Memory Subsystem



- 4MB SRAM non-coherent (reserved for DP) processing
- ECC protected
- Heavily bank interleaved SRAM design
- 230GB aggregated Rd/Wr bandwidth
- PEB buffers allocated by BMAN the same way as DRAM buffers

- 2x 64-bit DDR4 up to 2.133GT/s
ECC protected
- 2 memory channels with cache
line interleaving
- 34.1GB/s raw bandwidth
- x8, x4, x16 DIMM support
RDIMM, UDIMM
- 4CS with up to 128GB capacity
- Coherent memory space
- Enhanced order handling
mechanisms and number of
outstanding requests
- Improved bank hashing

- 1x 32-bit DDR4 up to 1.67GT/s
 - ECC protected
 - 6.4GB raw bandwidth
 - x8, x16 DIMM support
 - RDIMM and UDIMM
- 2CS with up to 64GB capacity
- Memory space non-coherent (reserved for DP processing)

LS2085/8A DPAA 2.0 Architecture

- Fully virtualized and isolated Data Path Acceleration Subsystem.
- Kernel bypassing and zero copy with user-space virtual address.
- Fully isolated and security provisioning for DPAA portal and memory accesses.
- Virtual switching allows convenient and isolated access to acceleration offload.

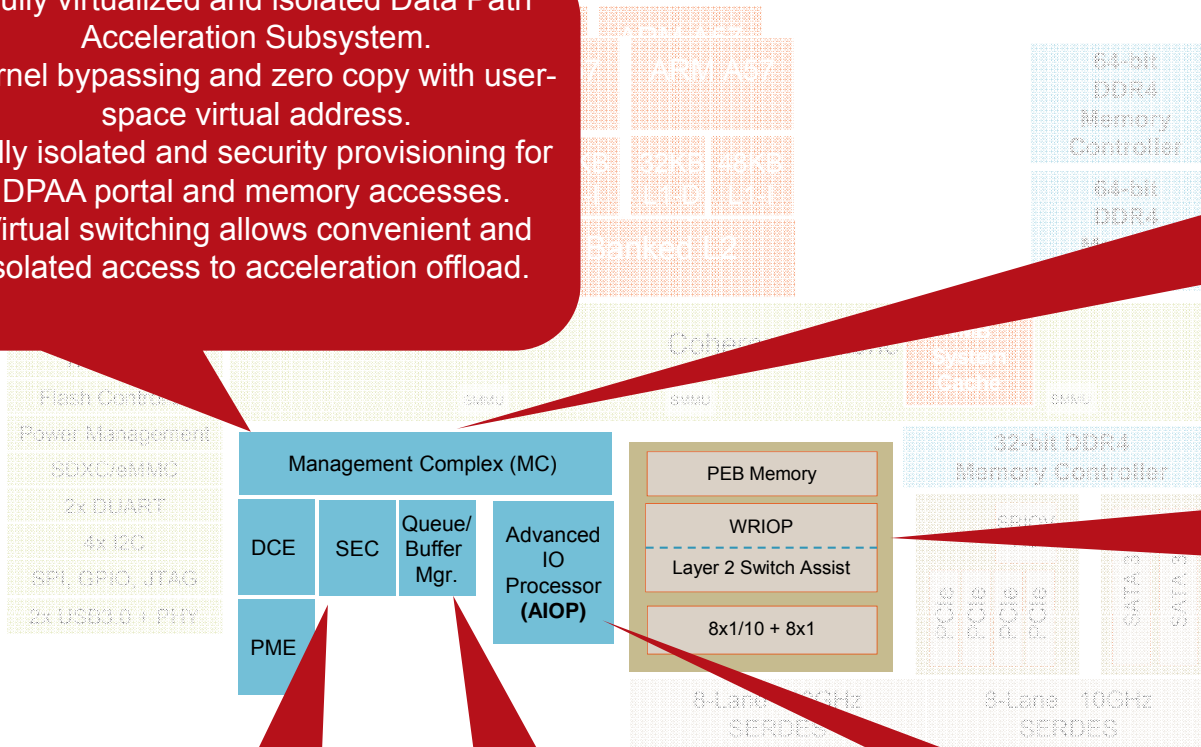
MC
Provides Layerscape object
abstraction
And allows application
software easy access of
DPAA 2.0 features.

WRIOP
Line rate 88Gbps Networking
Intelligent distribution, queuing
and drop decisions
Interface profiles
Embedded L2 switch

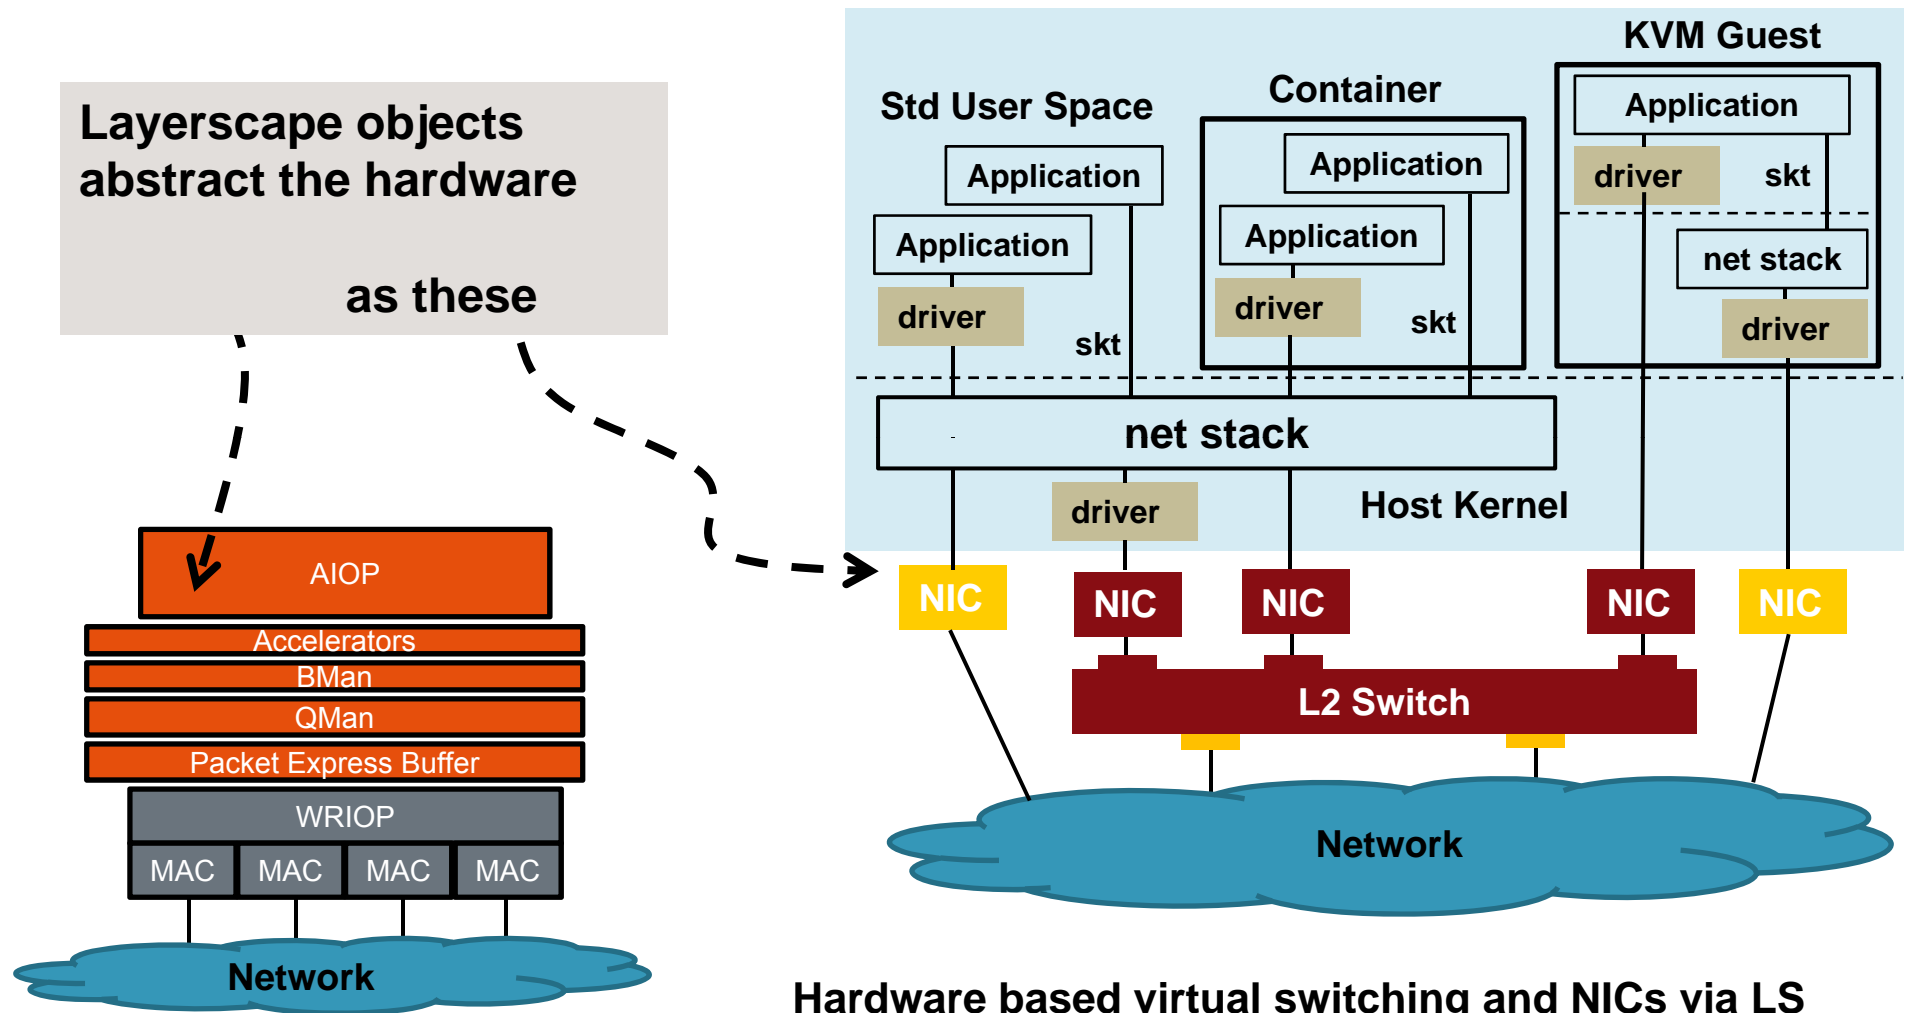
SEC, 20Gbps bulk crypto acceleration and numerous crypto algorithms
3rd Gen PME delivering 10Gbps Pattern Match/RegEx performance
2nd Gen DCE, 20Gbps Compression/Decompression aggregated

QMAN ,provides efficient, isolated, high-bandwidth event machine connections between separate control / management planes, data plane, services / acceleration functions, and physical network.

AIOP
17MPPS complex forwarding
30MPPS simple forwarding
Programmable engines,
accelerators



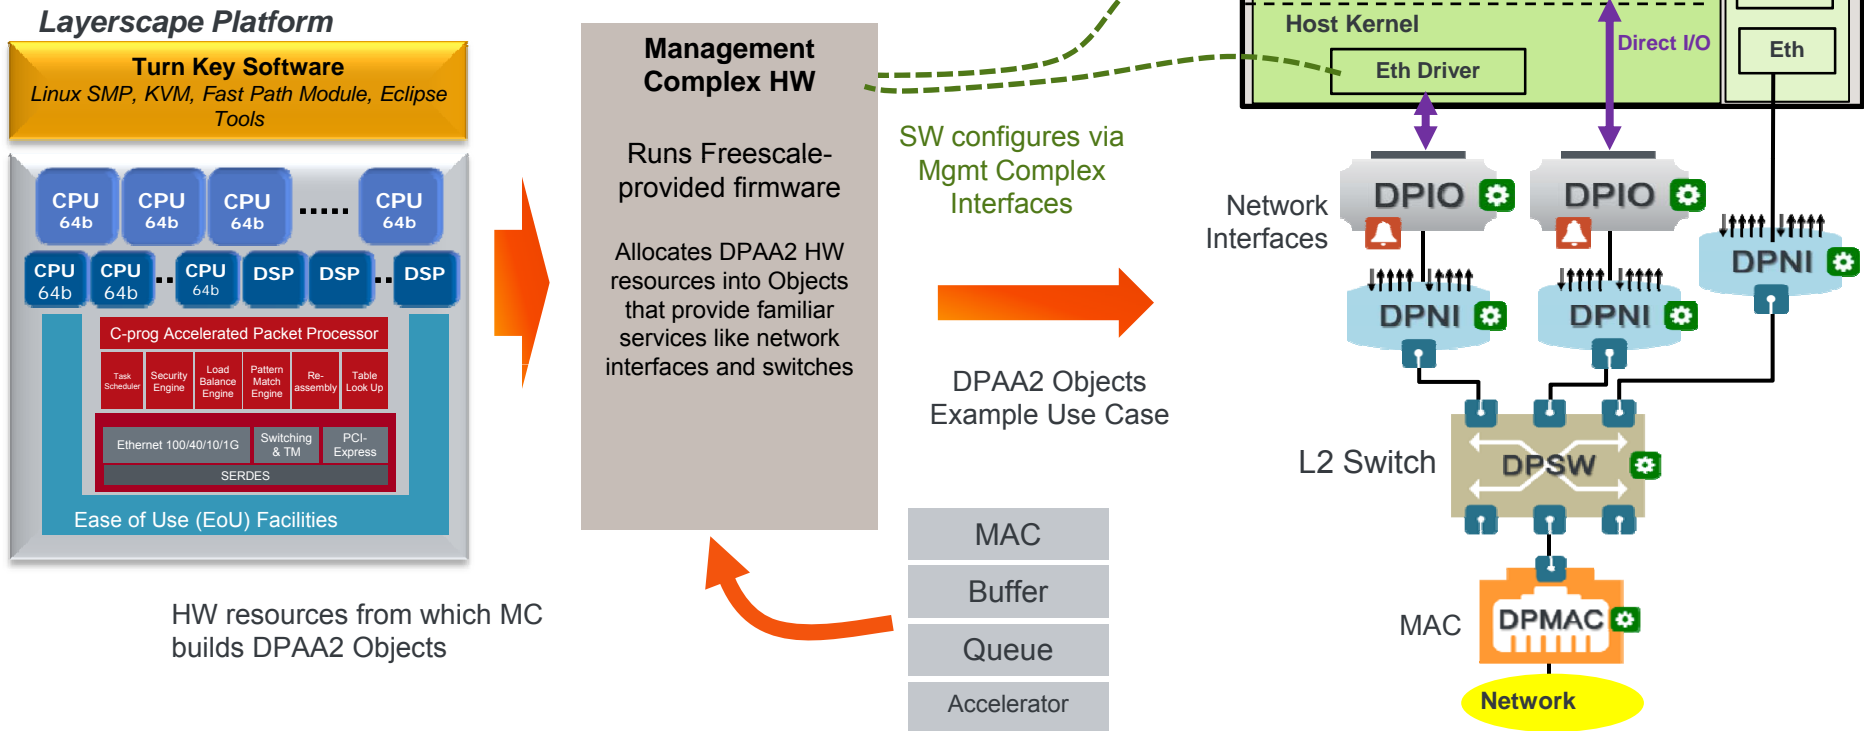
Management Complex Provides Hardware Abstraction



Hardware based virtual switching and NICs via LS Objects: SW mediated encapsulations of HW resources for a specific task.

Hardware Abstraction: Software Developer's View

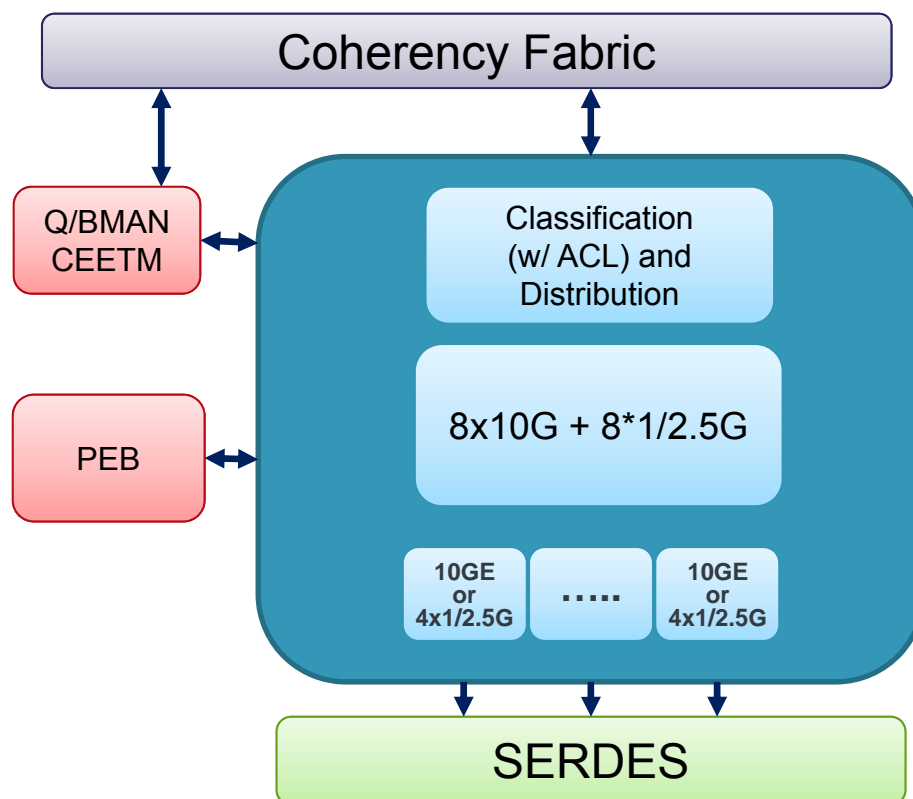
- Ease of use through DPAA2 Object Abstraction
- Performance because objects are HW that can be directly assigned to the software that uses them.



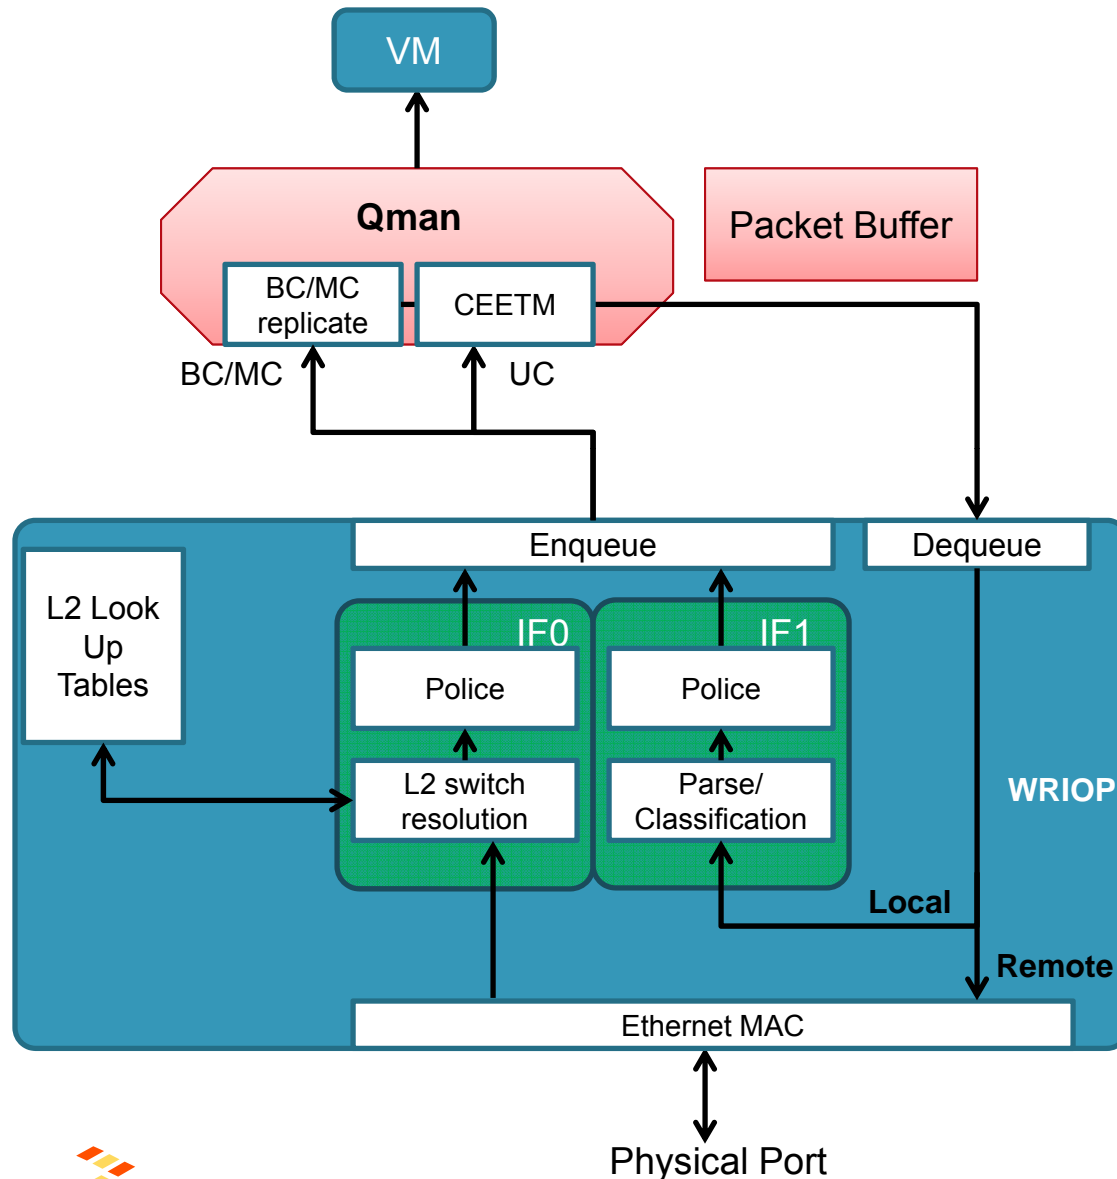
LS2085/8A Wire-Rate IOP (WRIOP)

- **16 physical ports and 2 recycle ports of the following types**
 - Ethernet MACs at 1/2.5/10Gbps
- **CRCs, IEEE 1588v2, MACSEC (802.1AE)**
- **Broadcast/Multicast**
- **Higher layer protocol offloads**
 - TCP/IP Checksum check/generations
- **Parse & Classify**
 - Parsing of standard headers and non-standard headers
 - 3 user defined tags that include 3 fields to parse at 32b
 - Classification to determine Interface Profile ID
 - Determines action to be taking on packet
- **Embedded Virtualization Features**
 - Virtual MACs, Interface Profiles (up to 1K)
 - Virtual device MIBs associated with interface profile
 - Independent physical & virtual port reset/reconfiguration
- **Policing and Flow Control**
 - Dual rate 3 color policing/marking
 - Policing on up to 256 profiles (RFC4115, RFC2968) based on classification result
 - Transmit pause-frame on buffer depletions or congestion on queues
 - PFC IEEE802.1Qbb, 802.3x based on queue congestion
 - PFC mapping to traffic class, statistics

- **Statistics**
 - 802.3 basic and mandatory managed objects statistic counters
 - IETF Management Information Database (MIB) package (RFC2665)
 - Remote Network Monitoring (RMON) counters
 - IFP statistics

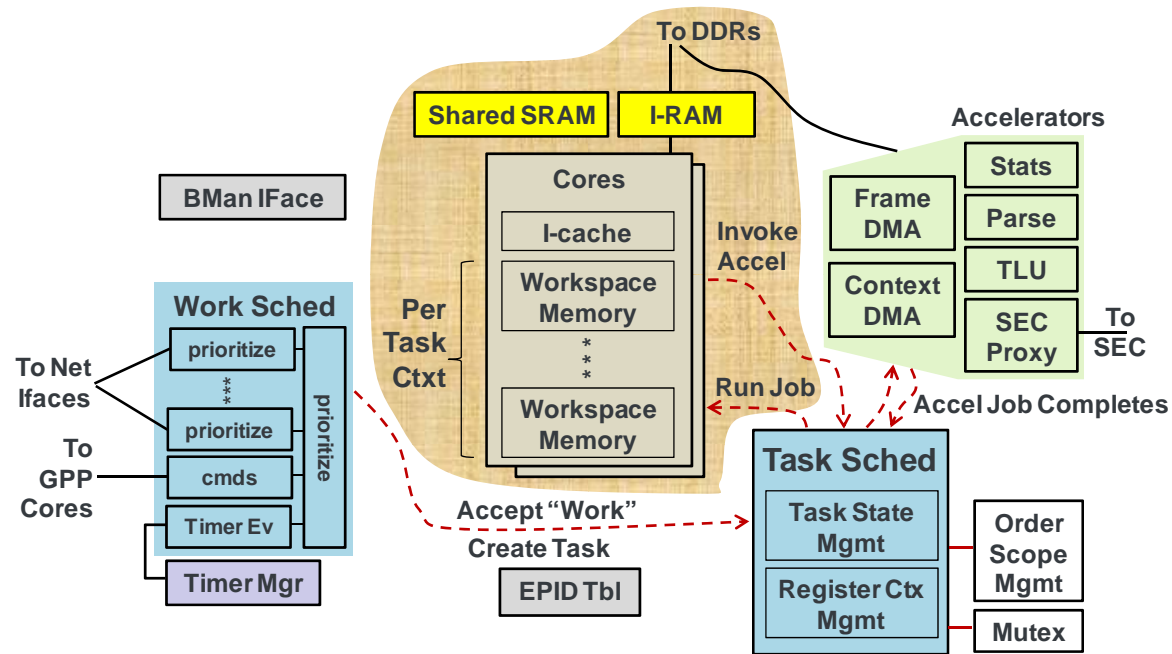


Embedded L2 Switch (WRIOP and QBMAN)



- L2 switching (including virtual switching) between 16 physical ports and 2 recycling ports (for VMs)
- L2 Lookup Tables include ACL, VLAN/Port, MAC address, ACL (TCAM) lookups
- VLAN aware bridging and Ethernet multicast according to 802.1Q
- QOS and Traffic Shaping according to 802.1Q
- Strict Priority/weighted round robin scheduling into 8 queues per port along with port-level CR/ER rate shaping for Egress traffic
- MSTP, RSTP and GMRP according to 802.1Q
- Hardware address learning
- MAC address aging
- Ingress policing
- Broadcast/Multicast packets

LS2085/8A, Advanced IO Processor (AIOP)



Fast path data plane/packet processor (16xe200 @800 supporting 256 concurrent task))

Hardware task scheduler

Minimized context switching overheads

C programmable

Packet processing accelerators

- Table lookup (EM/LPM/ACL)
- Packet infrastructure (BQMan, DMA,...)
- Parser, SEC, timer etc..

SG Buffer Management in hardware

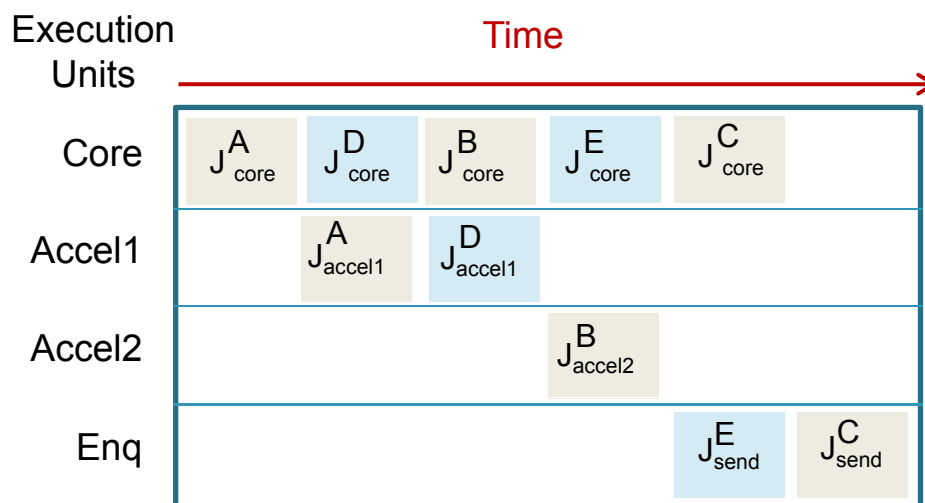
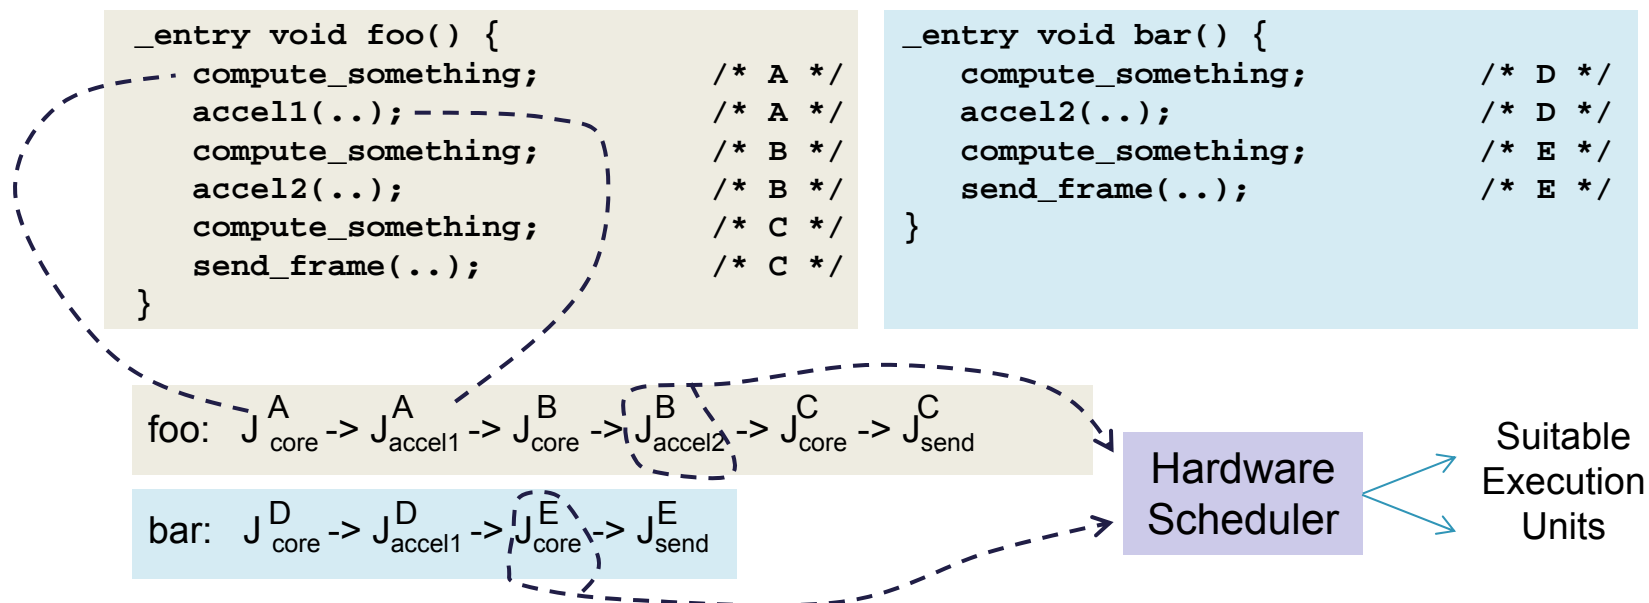
Packet order maintenance & synchronization in hardware

Synchronous programming model

Deterministic performance

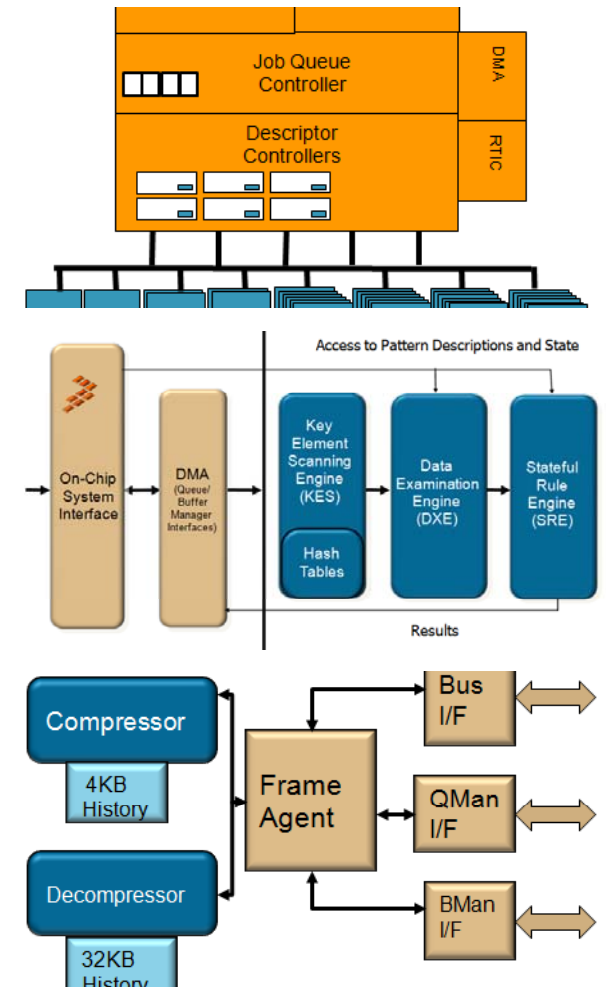
Synchronous Programming Model

Ease of Use and Parallelism

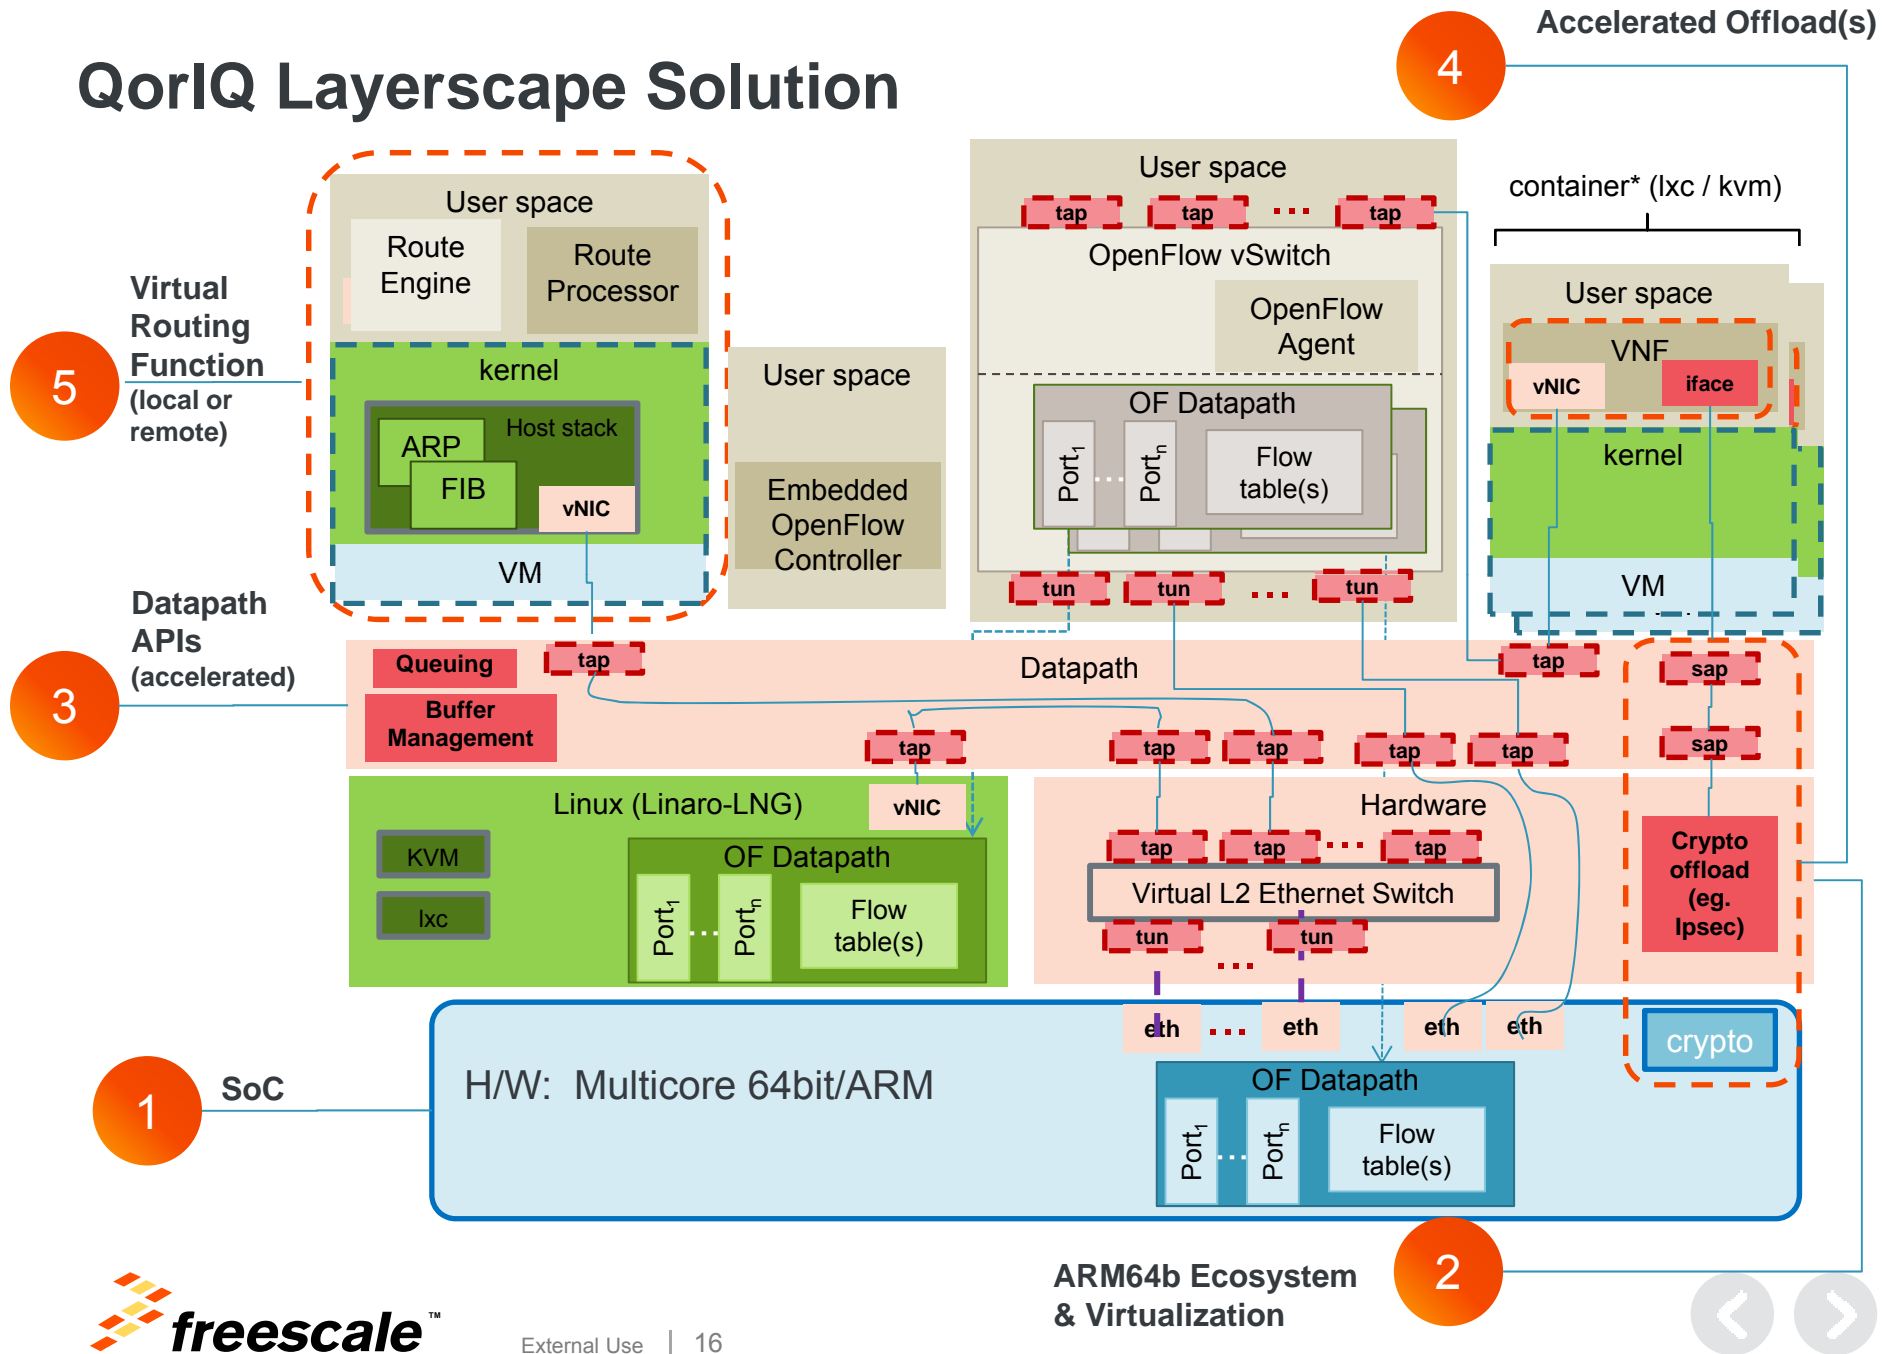


LS2085/8A SEC, PME and DCE Acceleration Engines

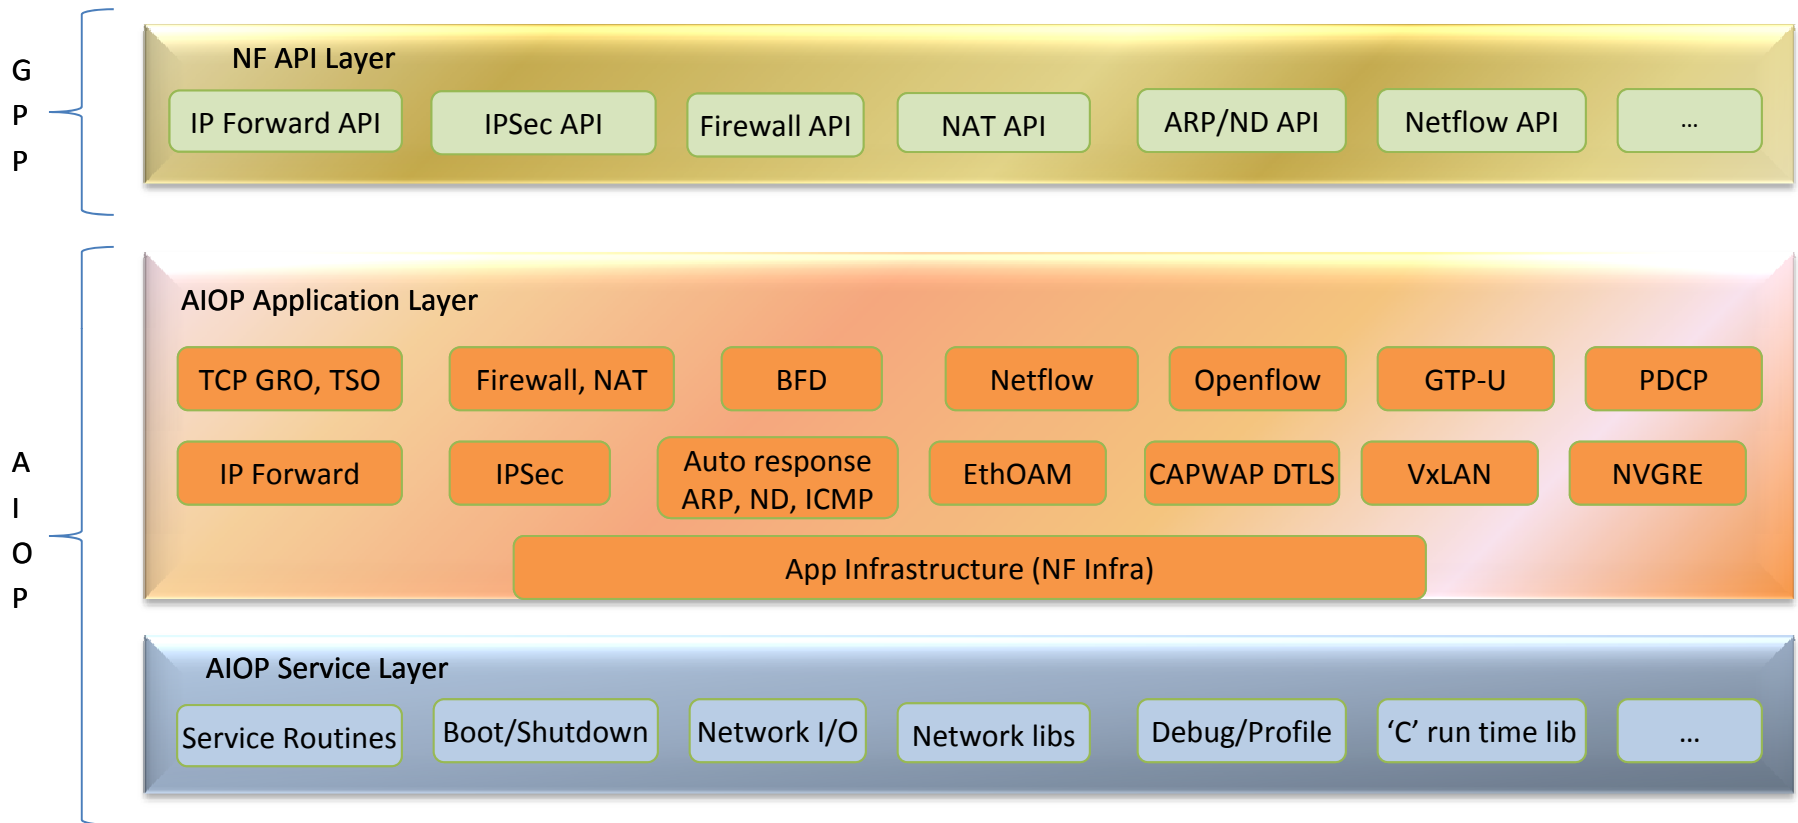
- SEC, 5th generation
 - AES (128, 192, 256-bit), EIA, 23.8Gbps AES-256
 - DES/3DES, 14.3Gbps 3DES
 - SHA-1,2,256,384,512 digest, 34.1Gbps SHA-256
 - MD5 128-bit digest
 - Header & Trailer off-load for following Security Protocols:
 - IPSec, SSL/TLS, 3G RLC, PDCP, SRTP, Wifi, MACSEC
 - 23.8 Gbps IPSec (AES-HMAC-SHA-2)
 - NIST certified Random Number Generators
 - 3G acceleration for Snow, ZUC, Kasumi
 - CRC32, CRC32C
- PME, 3rd generation
 - Perl meta-characters including wildcards, repeats, ranges, anchors, etc.
 - Stateful rules with user-defined instruction reacting to pattern match events, can be used to correlate patterns, qualify matches or track protocol state changes
- DCE, 2nd generation
 - Deflate as specified as in RFC1951
 - GZIP as specified in RFC1952
 - Zlib as specified in RFC1950
 - 12Gbps comp/decomp or 20Gbps combined



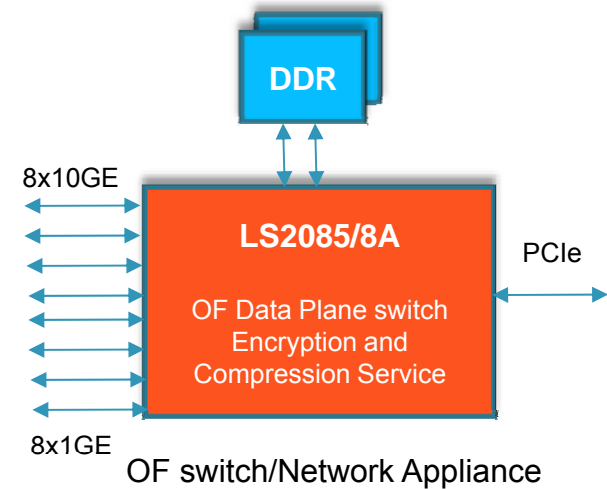
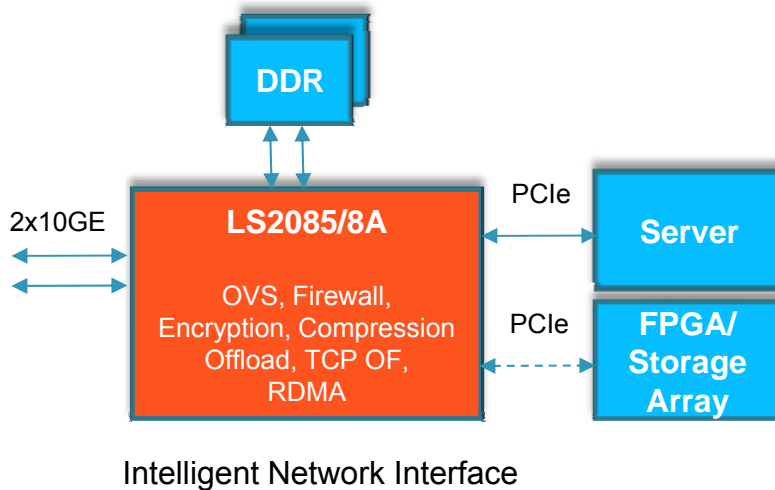
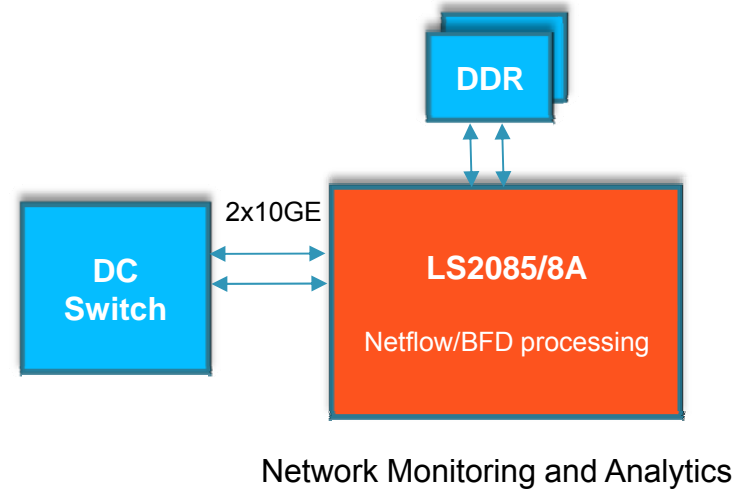
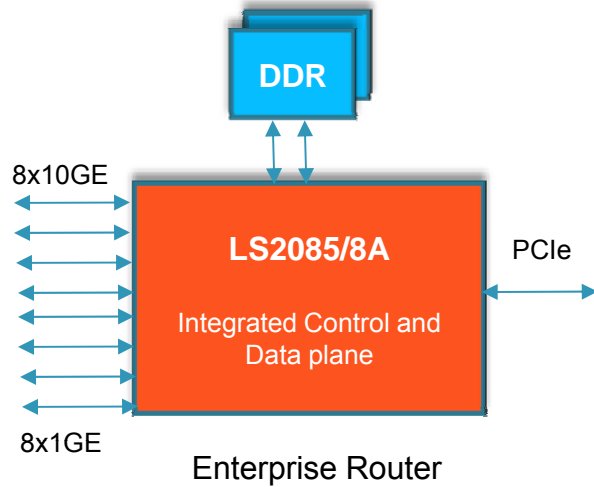
QorIQ Layerscape Solution



AIOP Software: Building Blocks



Example Applications



Core and Platform Performance Results


- Operating condition
 - 8xA57 (LS2085A) or 8xA72 (LS2088A) @2Ghz
 - CCN-504 @1.6Ghz
 - Platform 800Mhz
 - DDR 2.133GT/s
- Dhrystone
 - LS2085A, 89600 composite, 5.6 DMIPS/Mhz
- CoreMark
 - LS2085A, 69900 composite, 4.37 CoreMark/Mhz
- SpecInt2006
 - 74 for SpecINT2006-Rate or 12 for SpecInt2006 Single core
 - LS2088A improves SpecInt performance by ~12% with 15% core power reduction
- LMBench latency
 - LS2085A, L1, 4 cycles; L2, 18 cycles; DDR, 208 cycles
- Stream bandwidth for two 64-bit DDR memory channels
 - LS2085A, 19.7GB/s achieved out of 34.1GB/s theoretical

AIOP Packet Processing Performance Results

Use cases / Benchmarks	Results
Complex Fwd Packet Processing Proof point <ul style="list-style-type: none">• 10K Algorithmic Access Control List (AACL) Rules• 5 classification stages per frame<ol style="list-style-type: none">1. Logical port (index/exact match)2. Policy Based Routing – Access Control List (ACL)3. Longest Prefix Match (LPM) Routing4. IP SA spoof check (exact match)5. ARP (exact match)	20Gbps @128B Packet Size 17MPPS
Netflow (IPFIX) Packet Processing	20Gbps @ 128B
Simple IPSec Fwd	15Gbps@390B
L2 Switch – Physical & Virtual	120Gbps

Advanced IO Processor (AIOP) Benefits

Fast path: Complex IP Forwarding Efficiency



- 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)



4-6x Power Performance
over general purpose cores
in a lower power envelope



AIOP assisted standard SW packages / “objects”, black box or white box

Smart NIC
LRO/GSO

OpenFlow Switch

Open vSwitch
VxLAN
NVGRE
VM Manager


Virtual Appliance
Offload
(IP FWD, FW, Ipsec,
QoS, SLB/ADC)

Switch
Supplement
(BFD, Eth-OAM,
Netflow, sFlow)

Example use case,
Netflow with AIOP and
2xA57 cores, 3x
performance @50% power
compared to all 8 cores.



LS2085/8A Reference Design Board



- LS2085/8A FC-PBGA processor
- Two ports of 72-bits DDR4 (including ECC) up to 2.133GT/s
 - Each port supports two DIMM connectors.
 - Each DIMM connector supports single/dual rank DDR4 module.
- One port of 40-bits DDR4 (including ECC) up to 1.67GT/s with one DIMM connector and two CS.
- Four RJ45 connectors for 10GE support
- Four SFP+ cages for XFI support
- Two PCIe connectors supporting
 - PCIe card (x4/x8)
 - PCIe card (x4)
- Two SATA connectors
- Two USB 3.0 ports
- One SD/MMC card slot
- NOR/NAND flash interface
- 64MB high speed flash with SPI interface

Disclaimer

- The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.
- The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Freescale assumes no obligation to update or otherwise correct or revise this information. Freescale reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of Freescale to notify any person of such revisions or changes.
- Freescale makes no representations or warranties with respect to the contents hereof and assumes no responsibility for any inaccuracies, errors or omissions that may appear in this information.



www.Freescale.com