QorIQ T4240 DPAA Deep Dive

AMF-NET-T1282

Charlie Li

October 2013

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Three Generations of Acceleration

- Communication Processor Module (CPM)
  - Started with M683xx
  - PQI, PQII, and some PQIII
  - Ethernet, ATM and TDM
- Quicc Engine (QE)
  - Started with MPC8360
  - PQII Pro, some PQIII and some QorIQ
  - Ethernet, ATM and TDM
- Data Path Acceleration Architecture (DPAA)
  - Started with P4080 (DPAA 1.0); T4240 (DPAA 1.x)
  - Ethernet

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Enhancing Core Performance with Data Path Acceleration Architecture

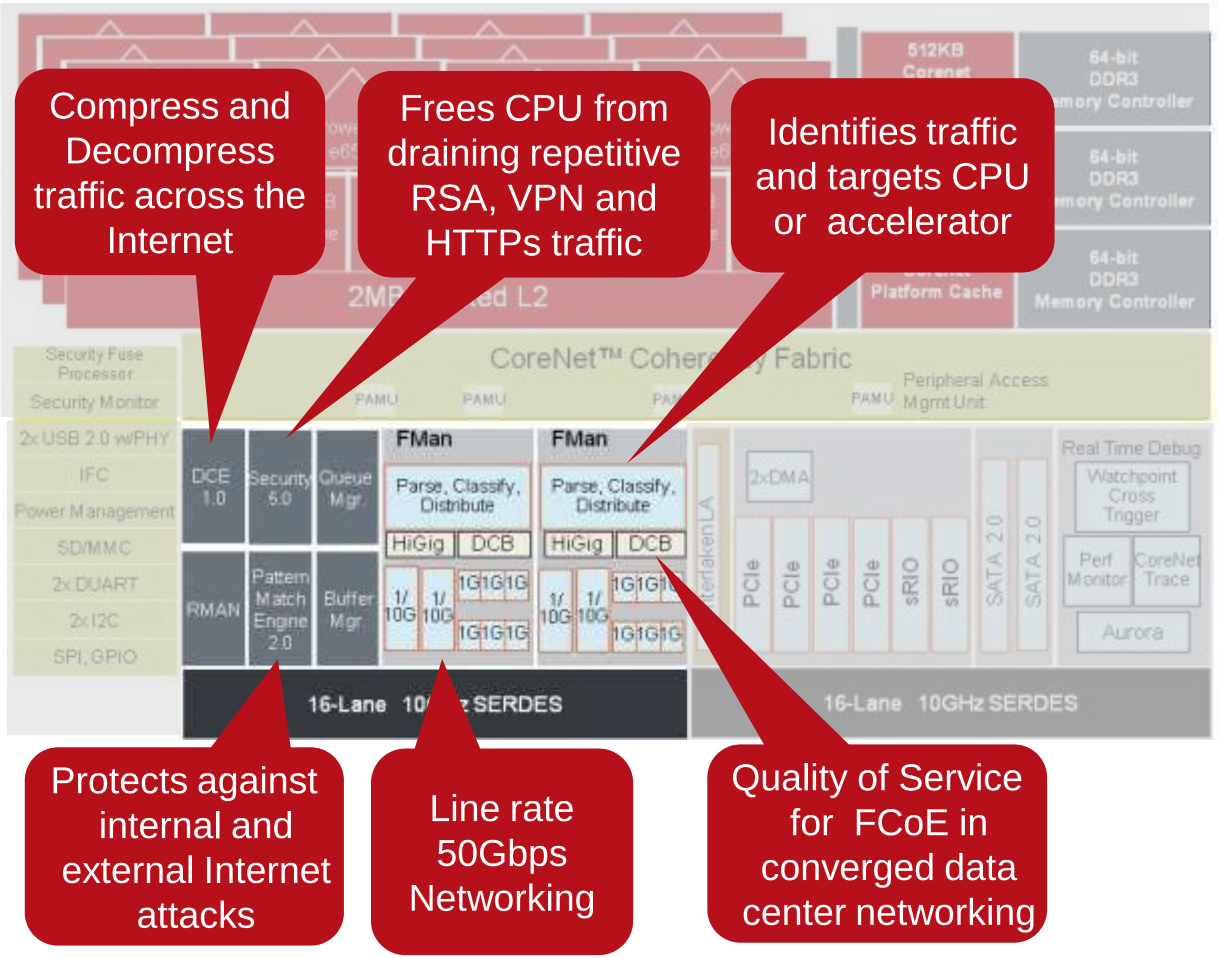

Compress and Decompress traffic across the Internet

Frees CPU from draining repetitive RSA, VPN and HTTPs traffic

Identifies traffic and targets CPU or accelerator

Protects against internal and external Internet attacks

Line rate 50Gbps Networking

Quality of Service for FCoE in converged data center networking

Hardware Accelerators	
FMAN Frame Manager	50 Gbps aggregate Parse, Classify, Distribute
BMAN Buffer Manager	64 buffer pools
QMAN Queue Manager	Up to 2^{24} queues
RMAN Rapid IO Manager	Seamless mapping sRIO to DPAA
SEC Security	40Gbps: IPSec, SSL Public Key 25K/s 1024b RSA
PME Pattern Matching	10Gbps aggregate
DCE Data Compression	20Gbps aggregate
Saving CPU Cycles for higher value work	

New	Enhanced

freescale™

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 Converge, QUICC Engine, Ready Play, SMARTMOS, Tower, TurboLink, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Data Path Acceleration Architecture Philosophy

- DPAA is design to balance the performance of multiple CPUs and Accelerators with seamless Integrations
  - ANY packet to ANY core to ANY accelerator or network interface **efficiently** WITHOUT locks or semaphores.
- “Infrastructure” components
  - Queue Manager (QMan)
  - Buffer Manager (BMan)
- “Accelerator” Components
  - Cores
  - Frame Manager (FMan)
  - RapidIO Message Manager (RMan)
  - Cryptographic accelerator (SEC)
  - Pattern matching engine (PME)
  - Decompression/Compression Engine (DCE)
  - DCB (Data Center Bridging)
- CoreNet
  - Provides the interconnect between the cores and the DPAA infrastructure as well as access to memory.

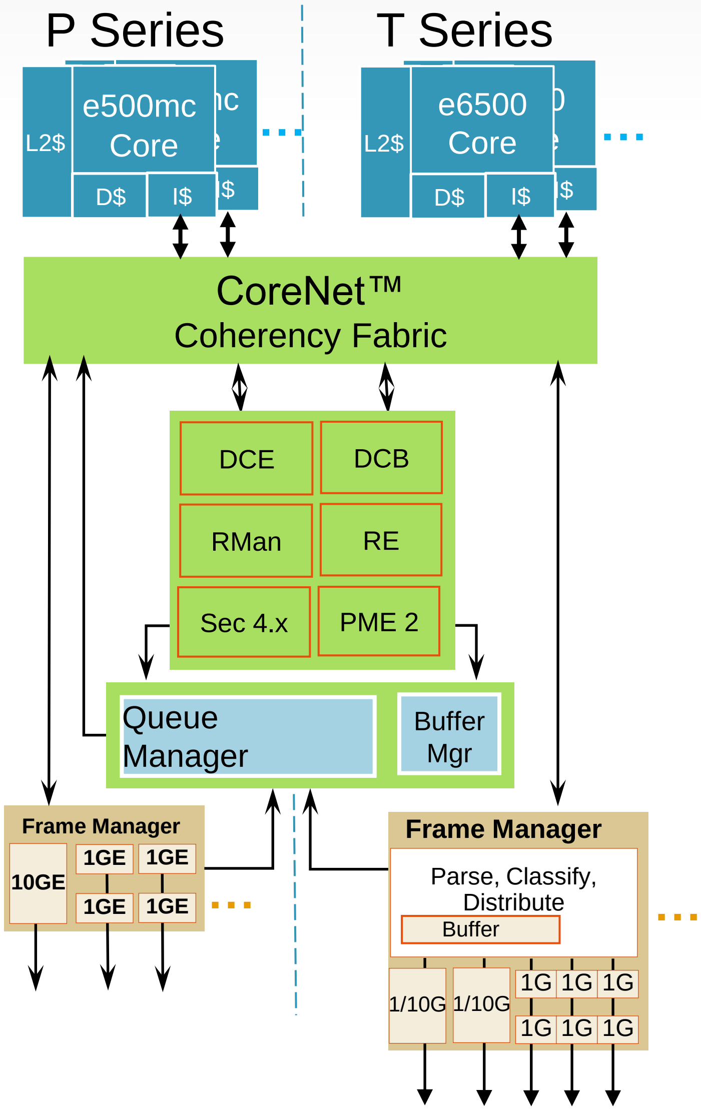

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

DPAA Terminology

- Buffer: Unit of contiguous memory, allocated by software
- **Frame**: Buffer(s) that hold a data element (generally a packet)
  - Frames can be single buffers or multiple buffers (scatter/gather lists)
    - A “simple frame” has one delimited data element
    - A “multi buffer frame” has two or more data elements
- **Frame Descriptor (FD)**: Proxy structure used to represent frames
- **Frame Queue**:
  - FIFO of related Frames Descriptor.(e.g. TCP session)
  - The basic queuing structure supported by QMan
- **Frame Queue Descriptor (FQD)**: Structure used to manage Frame Queues

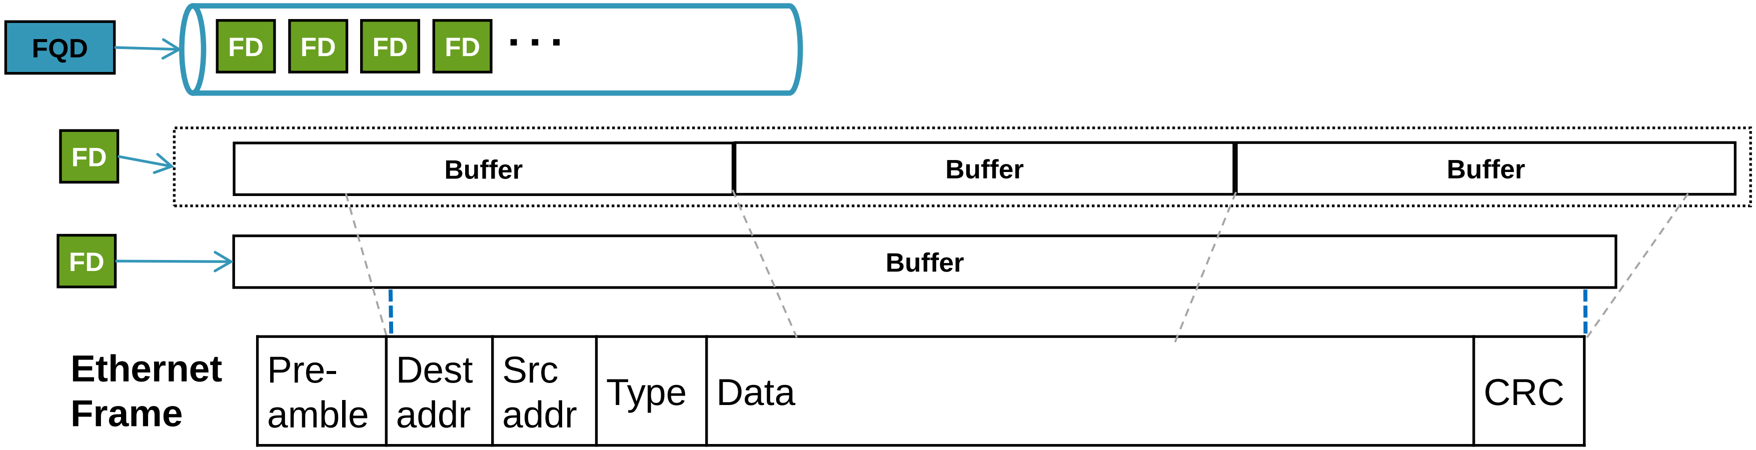

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

DPAA Building Block: Frame Descriptor (FD)

Simple Frame

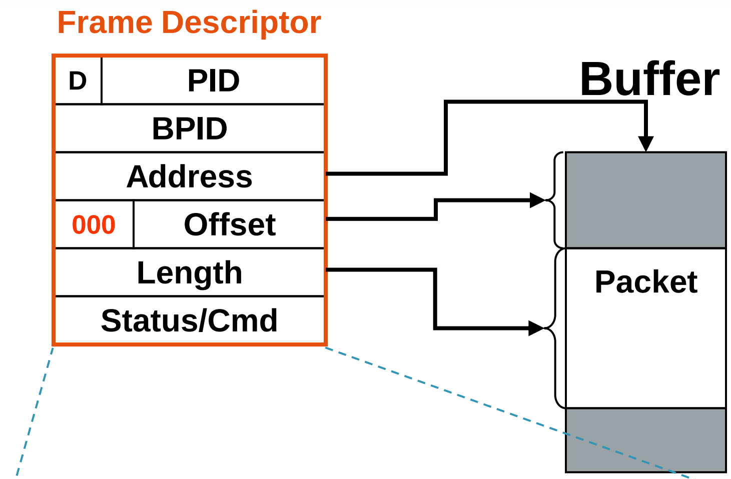

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
DD		LIODN offset						BPID								ELIODN offset				-	-	-	-	addr							
addr (cont)																															
Fmt			Offset									Length																			
STATUS/CMD																															

Multi-buffer Frame (Scatter/gather)

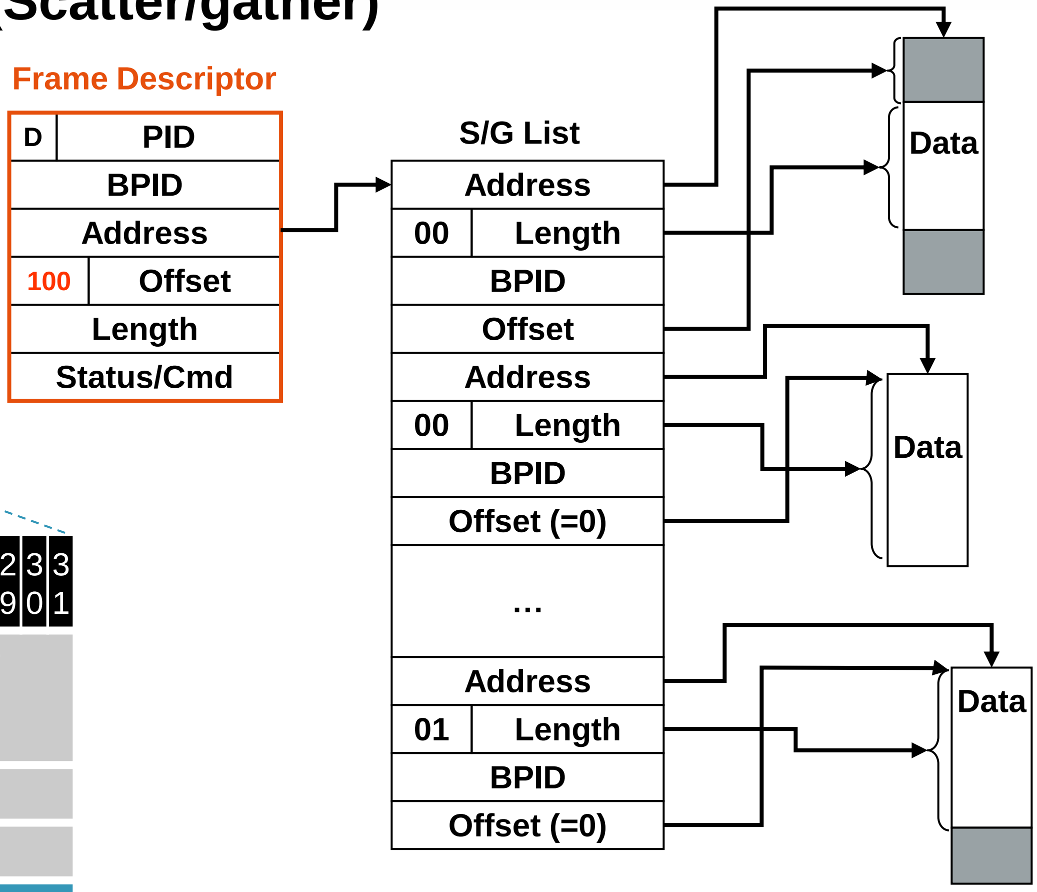

freescale™

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

DPAA Interaction: Compound Frame

- Compound frames allows related data to be passed in a single unit to the DPAA Accelerators.

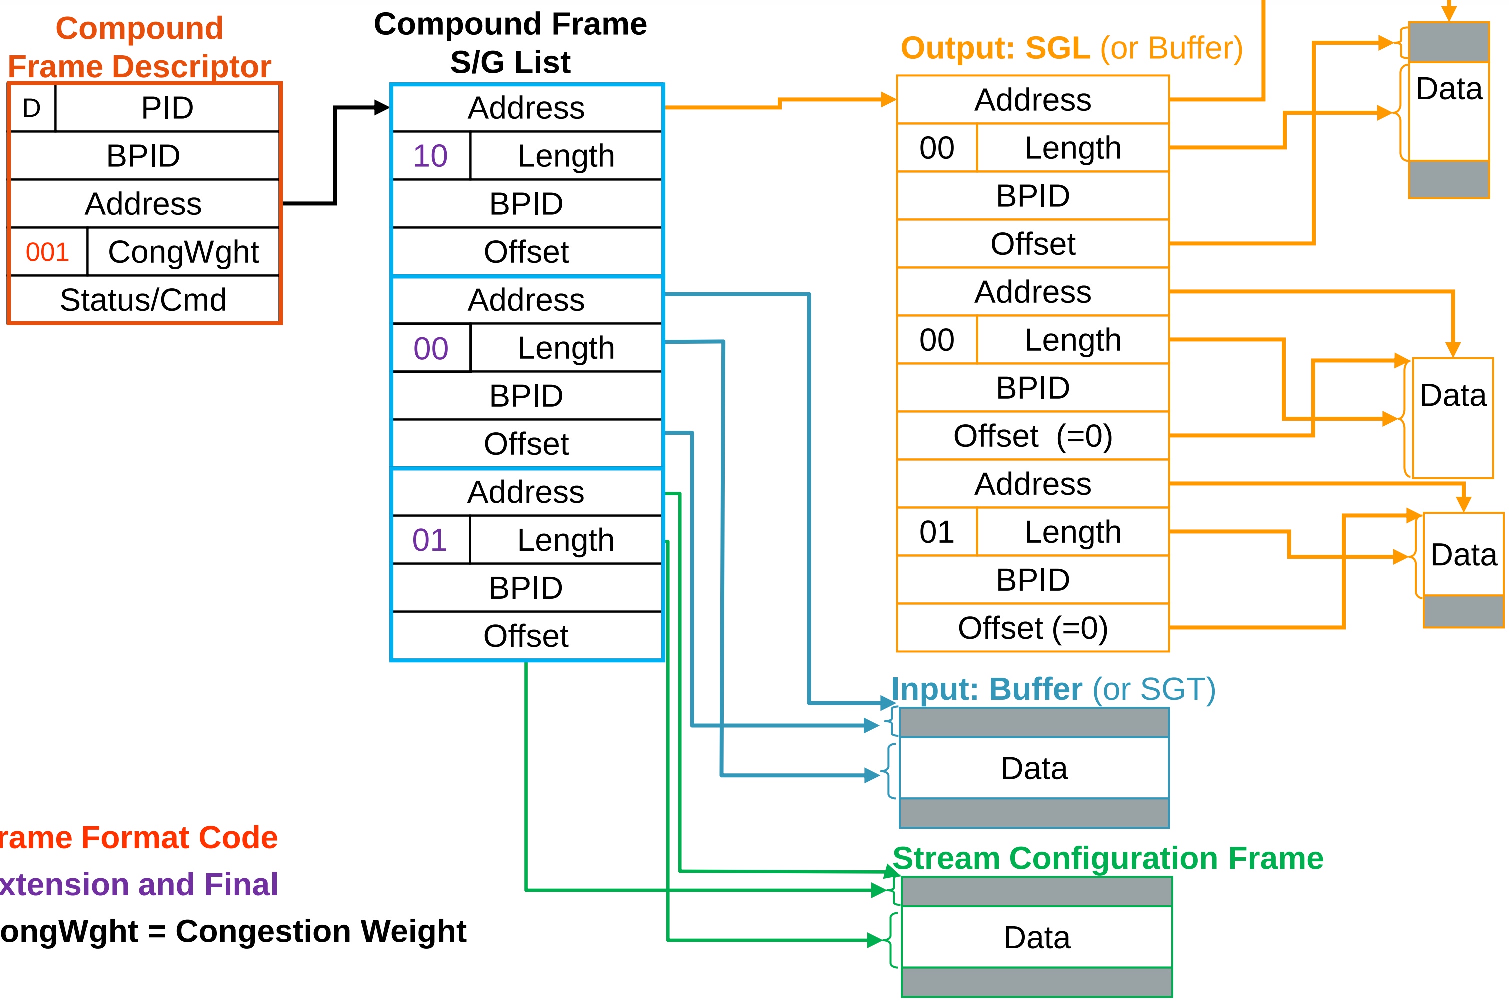

Compound Frame Descriptor

D	PID
BPID	
Address	
001	CongWght
Status/Cmd	

Compound Frame S/G List

Address	
10	Length
BPID	
Offset	
Address	
00	Length
BPID	
Offset	
Address	
01	Length
BPID	
Offset	

Output: SGL (or Buffer)

Address	
00	Length
BPID	
Offset	
Address	
00	Length
BPID	
Offset (=0)	
Address	
01	Length
BPID	
Offset (=0)	

Data

Data

Data

Input: Buffer (or SGT)

Data

Stream Configuration Frame

Data

Frame Format Code

Extension and Final

CongWght = Congestion Weight

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

DPAA Building Block: Frame Queue Descriptor

FQD Selected Field Description:

- FQD_LINK: Link to the next FQD in a queue of FQDs, used for Work Queues
- ORPRWS: ORP Restoration Window Size
- OA: ORP Auto Advance NESN Window Size
- ODP_SEQ: ODP Sequence Number
- ORP_NESN: ORP Next Expected Sequence Number.
- ORP_EA_HPTR, ORP_EA_TPTR: ORP Early Arrival Head and Tail Pointer
- PFDR_HPTR, PFDR_TPTR : PFDR Head and Tail Pointer
- **CONTEXT_A, CONTEXT_B**: Frame Queue Context A and B
- **STATE**: FQ State
- **DEST_WQ**: Destination Work Queue
- ICS_SURP: Intra-Class Scheduling Surplus or Deficit.
- IS: Intra-Class Scheduling Surplus or Deficit identifier
- ICS_CRED: Intra-Class Scheduling Credit
- CONG_ID: Congestion Group ID
- RA[1-2]_SFDR_PTR: SFDR Pointer for Recently Arrived frame # 1 and 2
- TD_MANT, TD_EXP : Tail Drop threshold Exponent and Mantissa
- C: FQD in external memory or in cache (Qman 1.1)
- X: XON or XOFF for flow control command (Qman1.1)

<table>
<tr><th>0</th><th>1</th><th>2</th><th>3</th><th>4</th><th>5</th><th>6</th><th>7</th><th>8</th><th>9</th><th>10</th><th>11</th><th>12</th><th>13</th><th>14</th><th>15</th><th>16</th><th>17</th><th>18</th><th>19</th><th>20</th><th>21</th><th>22</th><th>23</th><th>24</th><th>25</th><th>26</th><th>27</th><th>28</th><th>29</th><th>30</th><th>31</th></tr>
<tr><td>-</td><td>-</td><td>-</td><td>-</td><td>-</td><td>-</td><td>-</td><td>-</td><td colspan="24">FQD_LINK</td></tr>
<tr><td colspan="3">ORPRWS</td><td>OA</td><td colspan="14">ODP_SEQ</td><td colspan="14">ORP_NESN</td></tr>
<tr><td colspan="2">OLWS</td><td colspan="15">ORP_EA_HSEQ</td><td colspan="15">ORP_EA_TSEQ</td></tr>
<tr><td colspan="24">ORP_EA_HPTR</td><td colspan="8">ORP_EA_TPTR</td></tr>
<tr><td colspan="16">ORP_EA_TPTR</td><td colspan="16">PFDR_HPTR</td></tr>
<tr><td colspan="8">(cont…)</td><td colspan="24">PFDR_TPTR</td></tr>
<tr><td colspan="32">CONTEXT_A</td></tr>
<tr><td colspan="32">(cont…)</td></tr>
<tr><td colspan="32">CONTEXT_B</td></tr>
<tr><td colspan="11">FQ_CTRL</td><td>FE</td><td>R</td><td colspan="3">STATE</td><td colspan="16">DEST_WQ</td></tr>
<tr><td colspan="16">ICS_SURP</td><td>IS</td><td colspan="15">ICS_CRED</td></tr>
<tr><td colspan="32">BYTE_CNT</td></tr>
<tr><td colspan="8">CONG_ID</td><td colspan="24">FRM_CNT</td></tr>
<tr><td colspan="2">NRA</td><td colspan="2">OAC</td><td>C</td><td colspan="11">RA1_SFDR_PTR</td><td>X</td><td>IT</td><td>-</td><td>-</td><td>-</td><td colspan="11">RA2_SFDR_PTR</td></tr>
<tr><td colspan="2">NOD</td><td colspan="3">OAL</td><td colspan="11">OD1_SFDR_PTR</td><td colspan="5">OAL</td><td colspan="11">OD2_SFDR_PTR</td></tr>
<tr><td colspan="2">NPC</td><td>-</td><td>-</td><td>-</td><td colspan="11">OD3_SFDR_PTR</td><td>-</td><td>-</td><td>-</td><td colspan="8">TD_MANT</td><td colspan="5">TD_EXP</td></tr>
</table>

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Software Portal FQD Context_A Usage

- AE: Frame Annotation Stash Exclusive.
- DE: Frame Data Stash Exclusive
- CE: FQ Context Stash Exclusive
  - 0: Stash transaction issued as DIRECT0. PAMU translate this to LDEC
  - 1: Stash transaction issued as DIRECT1. PAMU translate this to LDECPE/LDECFE.
- AS: Frame Annotation Stashing Size
- DS: Frame Data Stashing Size
- CS: FQ Context Stashing Size
  - Number of 64 byte coherency granules (0, 1, 2, or 3) of Frame Annotation to be stashed.
- ADDR: FQ Context Address
  - the first 64 byte coherency granule containing the FQ context information to be stashed.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
-	-	-	-	-	AE	DE	CE	-	-	AS		DS		CS		-	-	-	-	-	-	-	-	ADDR							

32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
ADDR																										-	-	-	-	-	-

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Life of an Ingress Packet

- FMan receives packets
  - allocates internal buffers
  - retrieves data from MAC
- BMI
  - acquires a buffer from BMan
  - uses DMA to store data in it
- Parse+classify+keygen select a queue and policer profile
- Policer “colors” and optionally discards frame
- QMan applies active queue management and enqueues frame
- Frame is enqueued to one of a pool of cores
- Available core dequeue FD for processing

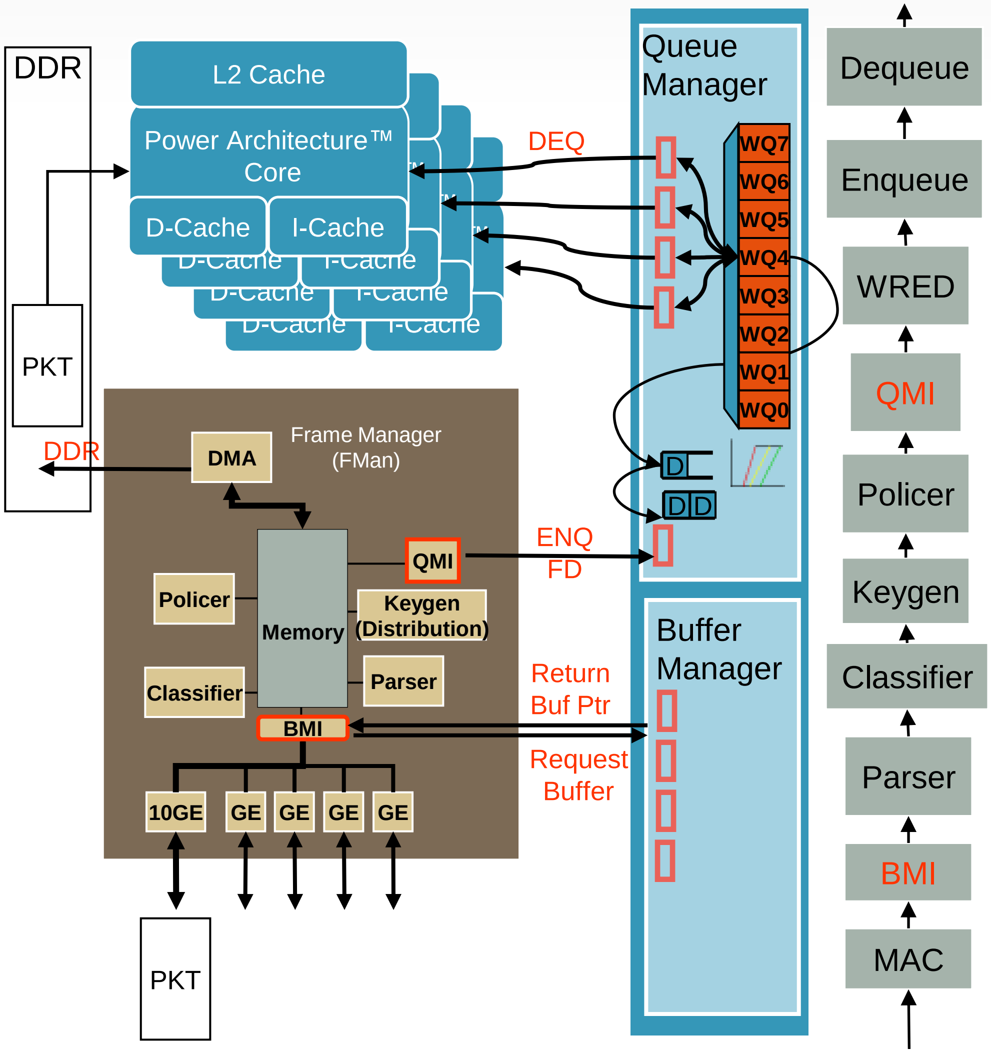

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Buffer Manager Functional Blocks

- Standardized command interface to SW and HW
  - Up to 66 Software portals for software: resolves any Multi Core race scenario
  - Up to 6 HW portal per HW block: simplified command for HW Accelerators
  - Up to 64 separate pools of free buffers
- BMan keeps a small per-pool stockpile of buffer pointers in internal memory
  - stockpile of 64 buffer pointers per pool, Maximum 2G buffer pointers
  - Absorbs bursts of acquire/release commands without external memory access
  - minimized access to memory for buffer pool management.
- Pools (buffer pointers) overflow into DRAM
- LIFO buffer allocation policy
  - A released buffer is immediately used for receiving new data, using cache lines previously allocated

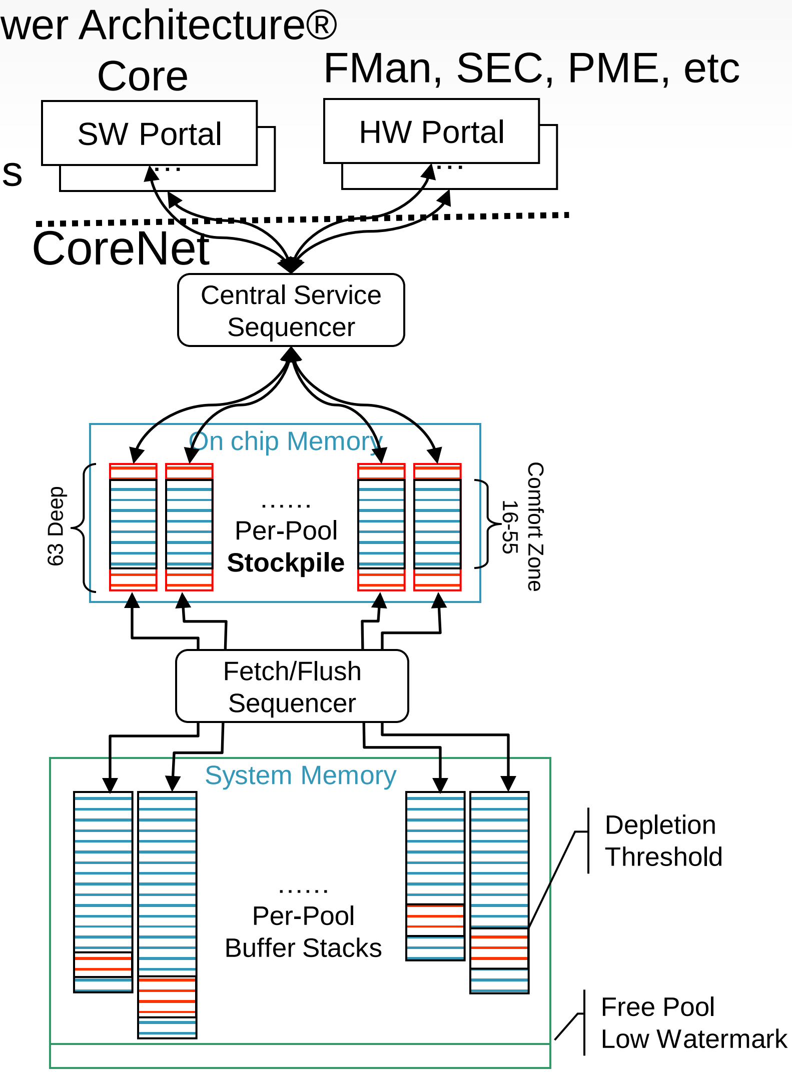

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

BMan SW Portal Components

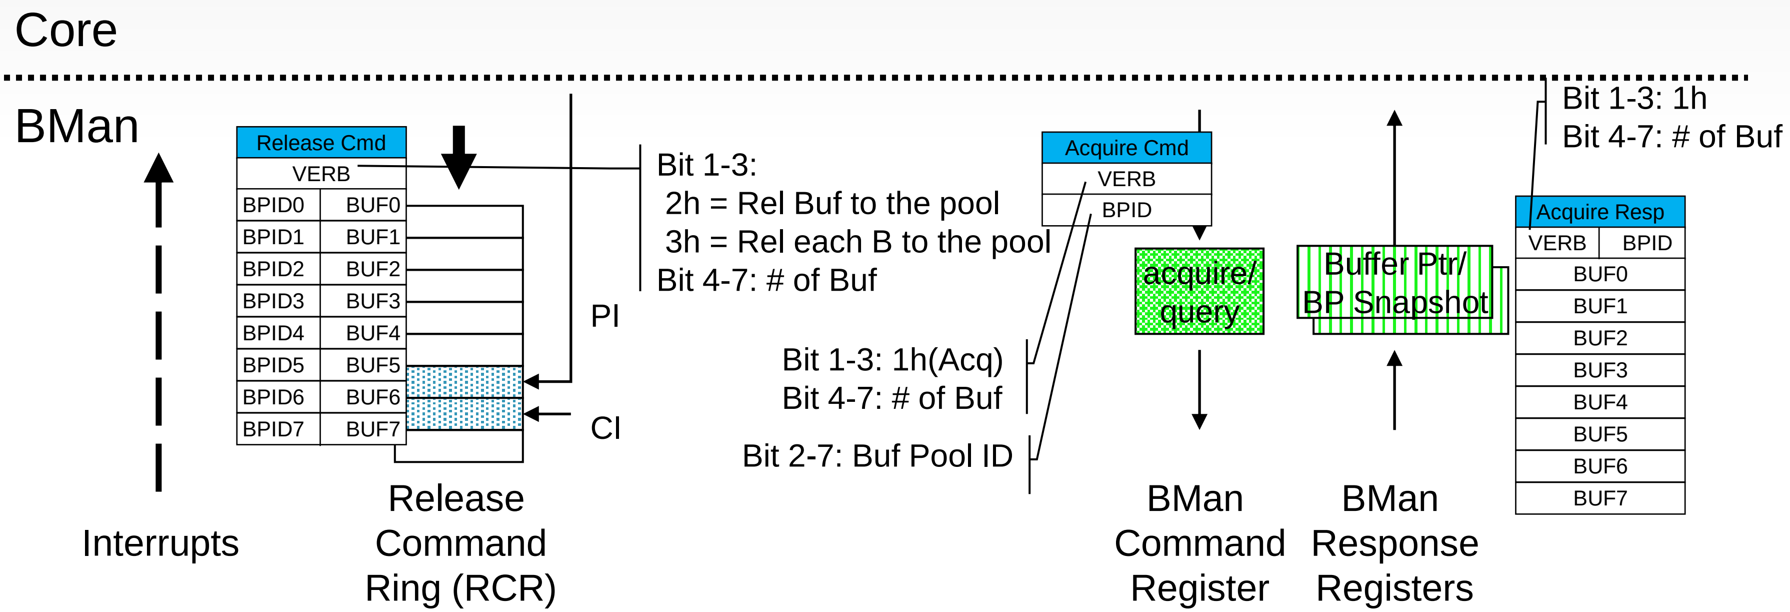

Core

BMan

Interrupts

Release Cmd	
VERB	
BPID0	BUF0
BPID1	BUF1
BPID2	BUF2
BPID3	BUF3
BPID4	BUF4
BPID5	BUF5
BPID6	BUF6
BPID7	BUF7

PI

CI

Release Command Ring (RCR)

Bit 1-3:
2h = Rel Buf to the pool
3h = Rel each B to the pool
Bit 4-7: # of Buf

Acquire Cmd
VERB
BPID

Bit 1-3: 1h(Acq)
Bit 4-7: # of Buf

Bit 2-7: Buf Pool ID

acquire/ query

BMan Command Register

Buffer Ptr/ BP Snapshot

BMan Response Registers

Bit 1-3: 1h
Bit 4-7: # of Buf

Acquire Resp	
VERB	BPID
BUF0	
BUF1	
BUF2	
BUF3	
BUF4	
BUF5	
BUF6	
BUF7	

- Software portals have 2 components
  - Management commands:
    - Command Registers (BCSPi_CR) : acquire 1-8 buffers OR query availability
    - Response Registers (BCSPi_RR0 / RR1): buffer address OR Buffer Pool Availability and Depletion state
  - Buffer Release:
    - Release Command Ring (RCR) Entry (BCSPi_RCRj): Circular FIFO
- Interrupts can be used to signal availability of space (in RCR) and that pools are depleted and require replenishment (RCR Interrupt Threshold Register)

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

BMan Command and Response

- BMan Command Type
  - BMan command registers (BCSPi_CR, RR and RCR) are 64B long.
  - Command Verb (1B) + Buffer Pool ID (1B)
    - Bit 0: Valid bit
    - Bit 1-3: Response Type. Valid encodings are:
      - 001 = Acquire buffers (Acquire)
      - 010= Release buffers to the pool identified in byte field 1 (Release)
      - 011= Release each buffer to the pool identified in byte field immediately preceding its buffer field (Release)
      - 100=Query buffer pool state, depletion and availability.
      - 110= Invalid command (Response)
      - 111= Stockpile ECC Error (Response)
    - Bit 4-7: Number of buffers associated with command type, maximum 8
      - 0h = Zero buffers. 1h = One buffer .... 8h = Eight buffers
    - Returns up to eight 48bit buffer addresses

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Verb x18	BPID 1	-	-	-	-	-	-	-	-	-	-	-	-	-	-
-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

New Frame Manager (FMan) Features

- FMan combines the Ethernet network interfaces with packet distribution logic to provide intelligent distribution and queuing decisions for incoming traffic at line rate.
- FMan key new features for QorIQ T4 processors.
  - Six 1G/2.5G multirate Ethernet MACs (mEMACs) per Frame Manager
  - Two 10G multirate Ethernet MACs (mEMACs) per Frame Manager
  - QMan interface: Supports priority based flow control message pass from Ethernet MAC to Qman
  - Comply with IEEE 803.3az (Energy efficient ethernet) and IEEE 802.1QBbb, in addition of IEEE Std 802.3®, IEEE 802.3u, IEEE 802.3x, IEEE 802.3z, IEEE 802.3ac, IEEE 802.3ab, and IEEE-1588 v2 (clock synchronization over Ethernet)
  - Port Virtualization: Virtual Storage profile (SPID) selection after classification or distribution function evaluation.
  - Rx port multicast support.
  - Egress Shaping.
  - Offline port: able to copy the frame into new buffers and enqueue back to the QMan.

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Man Modular Architecture Processing Pipeline

Core dequeue & processing — **Core / SW**

QMan active queue mngt and scheduling (WRED) — **QMan**

QMI instructs QMan to enqueues FD
BMI releases internal buffer on completion — **QMI / BMI**
QMan / BMan

BMI instructs DMA to write frame IC and header to ext. Buffer — **BMI / DMA**

Per-group policing(RFC2698/4115) — **Policer**

Multiple stages: Parse / Classify / Distribute
→ Determine queue ID# — **PCD**

BMI instructs DMA to transfer frame to external buffer — **BMI / DMA**

Based upon Layer-2 packet size, BMI requests "right sized" buffer to BMan — **BMI / BMan**

BMI streams and allocate internal buffer for incoming frame + IC (Internal Context)
Calculate raw L4 checksum for parser — **BMI**

MAC Rx and validate — **MAC**

Rx

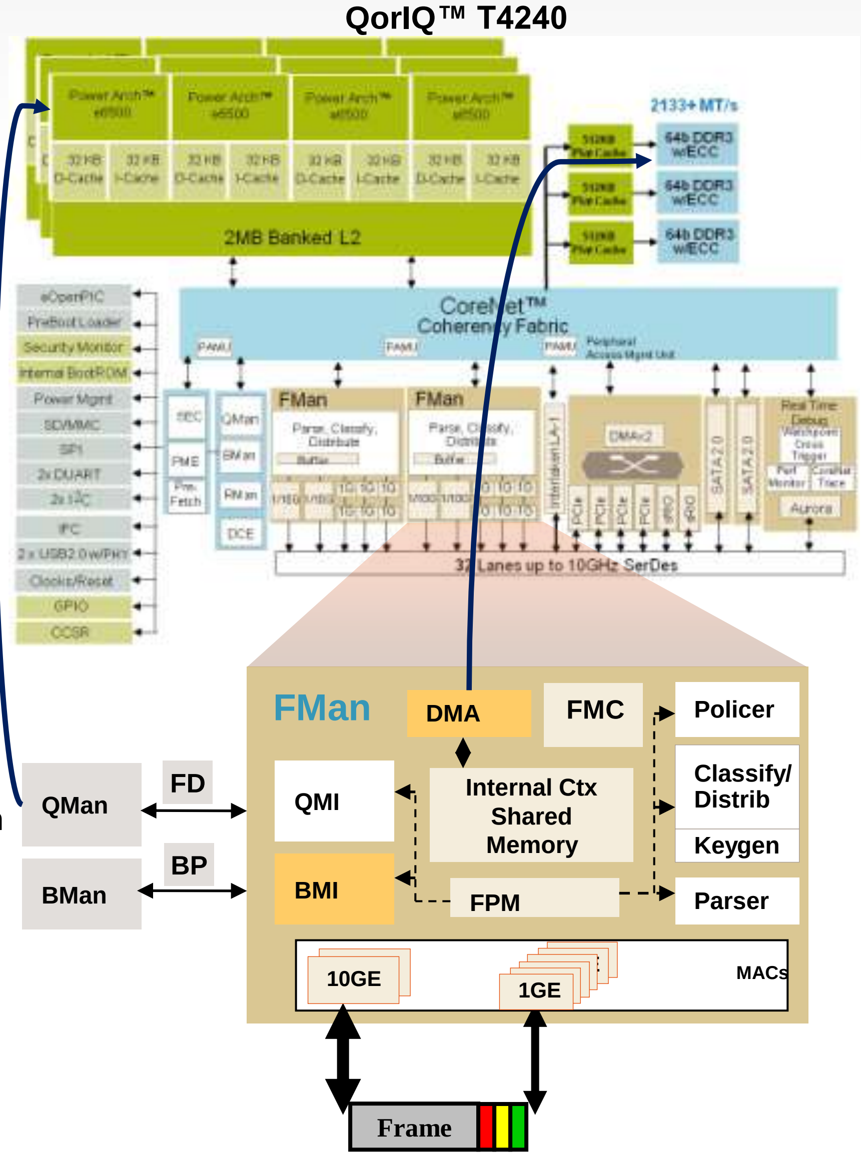

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

FMAN Ports Types

- Ethernet receive (Rx) and transmitter (Tx)
  - 1 Gbps/2.5Gbps/10Gbps
  - FMan_v3 some ports can be configured as HiGig.
  - Jumbo frames of up to 9.6 KB (add uboot bootargs "fsl_fm_max_frm=9600")
- Offline (O/H)
  - FMan_v3: 3.75Mpps (vs 1.5Mpps from the P series)
  - Able to dequeue and enqueue from/to a QMan queue. The FMan applies a Parse Classify Distribute (PCD) flow and (if configured to do so) enqueues the frame it back in a Qman queue. In FMan_v3 the FMan is able to copy the frame into new buffers and enqueue back to the QMan.
  - Use case: IP fragmentation and reassembly
- Host command
  - Able to dequeue host commands from a QMan queue. The FMan executes the host command (such as a table update) and enqueues a response to the QMan. The Host commands, require a dedicated PortID (one of the O/H ports).
  - The registers for Offline and Host commands are named O/H port registers.

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

1) Parser

- Performs parsing of common L2/L3/L4 headers, including tunneled protocols
- Can be augmented by the user to parse other standard protocols
- Can also parse proprietary, user-defined headers at any layer:
  - Self-describing, using standard fields such as proprietary Ethertype, Protocol ID, Next Header, etc.
  - Non-Self-Describing through configuration.
- Parse results, including proprietary fields, can be used by the classifier, and/or software.
- Soft parse can modify any field in parse results

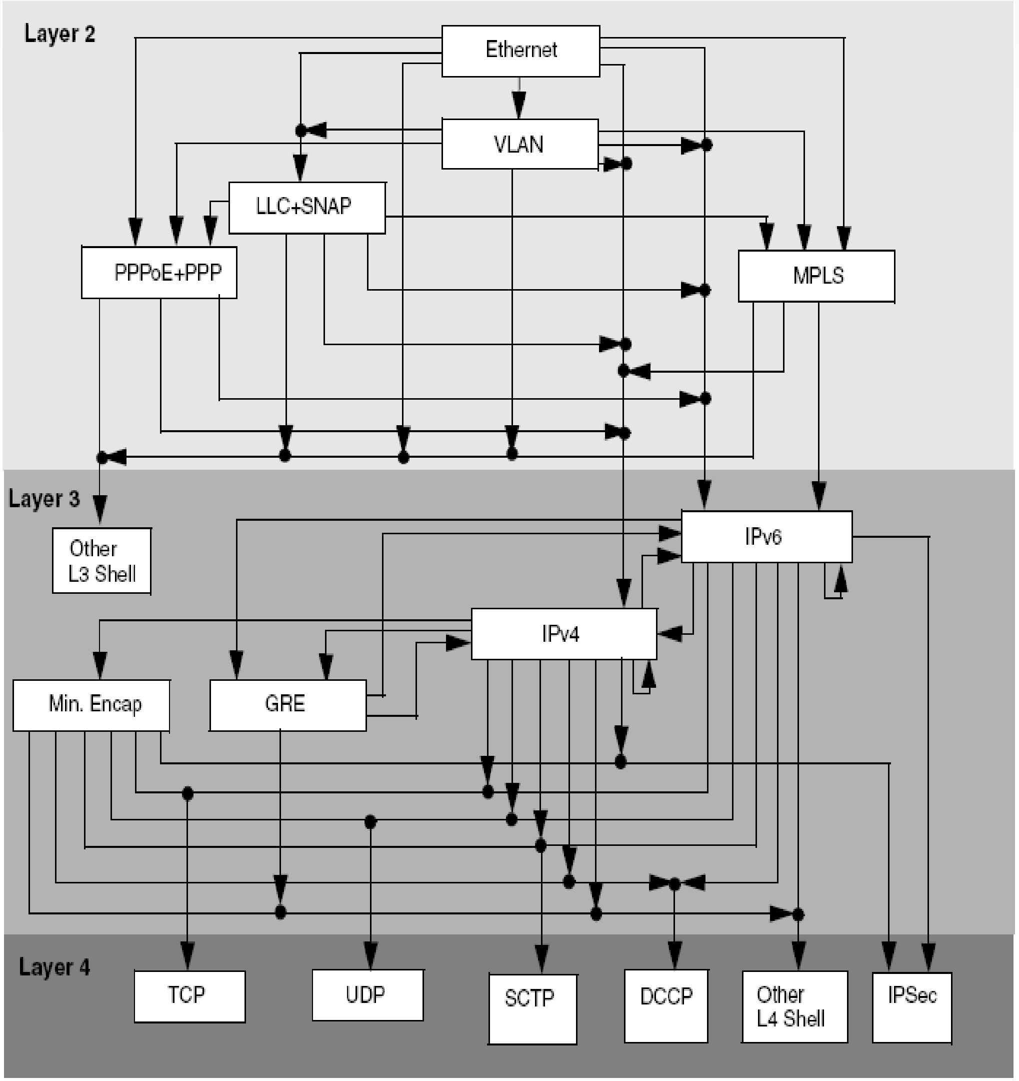

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Flexible Parsing of User Defined Fields (UDF)

The parser stores all parsing results in the Parse Array located inside the parser.

It is accessible through the Rx External Buffer Margins Register (FMBM_REBM(BSM))

Incoming frame

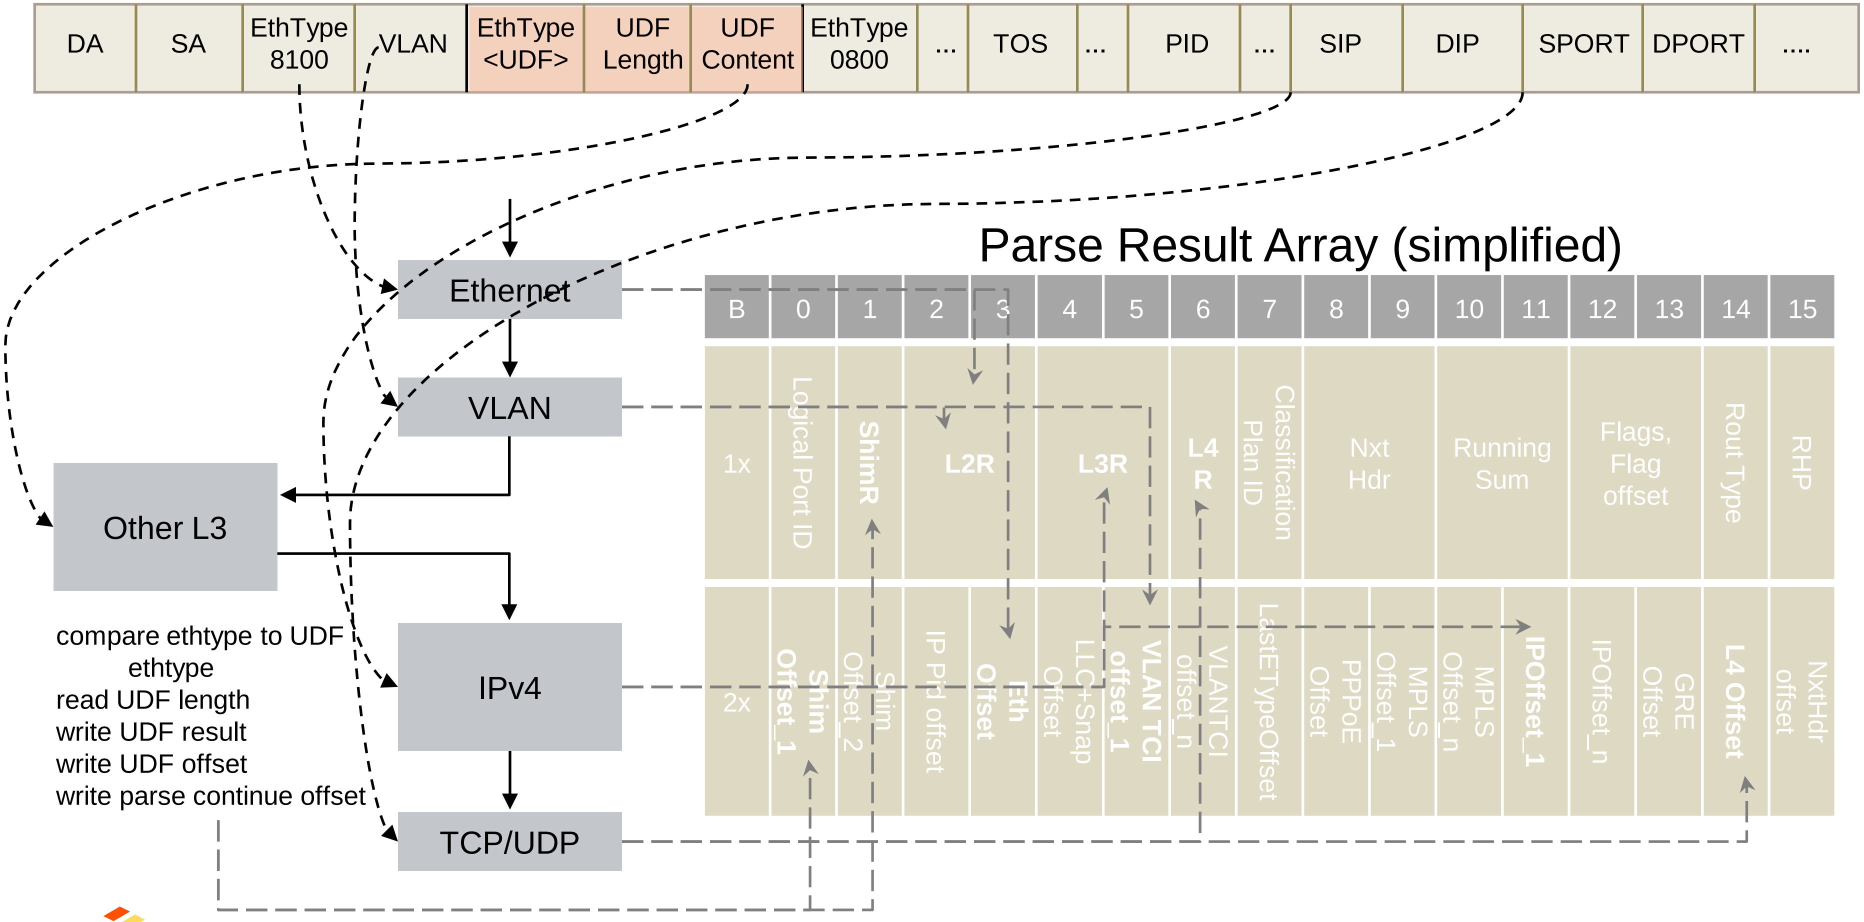

DA	SA	EthType 8100	VLAN	EthType <UDF>	UDF Length	UDF Content	EthType 0800	...	TOS	...	PID	...	SIP	DIP	SPORT	DPORT	....

compare ethtype to UDF ethtype
read UDF length
write UDF result
write UDF offset
write parse continue offset

Parse Result Array (simplified)

B	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1x	Logical Port ID	**ShimR**	**L2R**		**L3R**		**L4R**	Classification Plan ID	Nxt Hdr		Running Sum		Flags, Flag offset		Rout Type	RHP
2x	**Shim Offset_1**	Shim Offset_2	IP Pid offset	**Eth Offset**	LLC+Snap Offset	**VLAN TCI offset_1**	VLANTCI offset_n	LastETypeOffset	PPPoE Offset	MPLS Offset_1	MPLS Offset_n	**IPOffset_1**	IPOffset_n	GRE Offset	**L4 Offset**	NxtHdr offset

freescale™

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

2) KeyGen – Key Generator

- 256 Classification Plans
  - Indicates which fields of a parsed packet are of interest for key generation

- 32 Key Generation schemes
  - Direct Method
    - Used for port based or post coarse classification
  - Indirect Method
    - Based on the parsed protocol stack of the frame and the source port
    - The presence (or absence) of valid headers can direct the scheme used

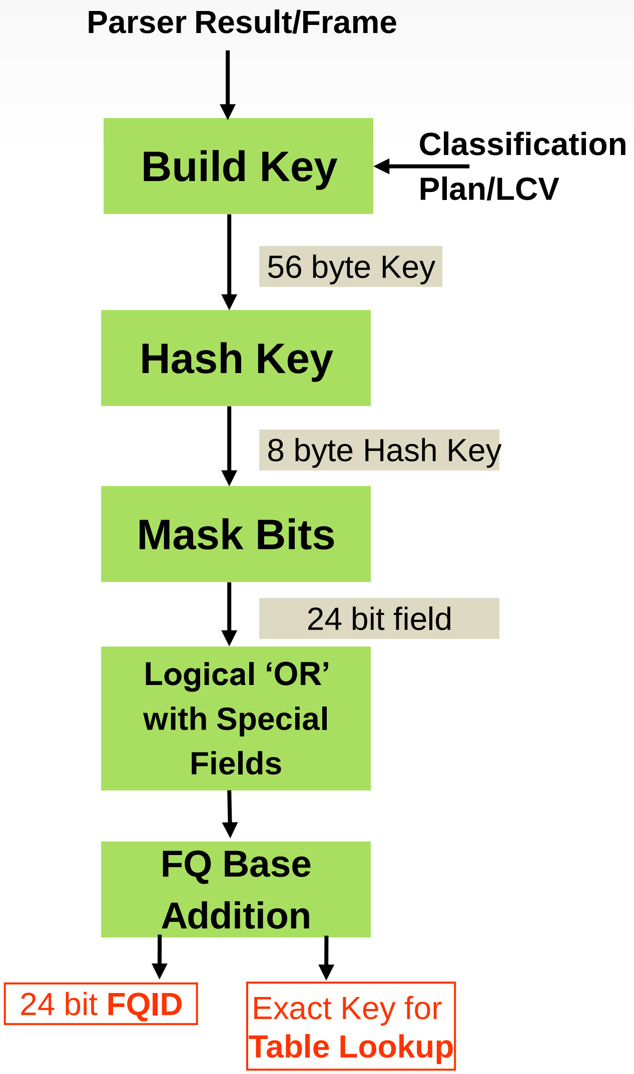

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

3) Classifier

- Up to three level tree search at line rate
  - Up to 256 bytes per tree level
  - Up to 512 bytes of total key data for the three level
- Each table is an ordered list of entries, returning the first match
  - Keys for the initial table are generated from parse results/KeyGen
  - Keys for subsequent tables are generated from parse and previous lookup results
  - Keys can be generated from any fields in the frame, including proprietary UDF
  - Each entry is individually maskable.
- Each table entry has Action description
  - Queue ID, next table, hash and distribute, or Drop.
- Output of classifier
  - Single queue
  - Set of queues and a KeyGen scheme if distributing
  - Policing Profile
- IPv4 Header Manipulation Use case

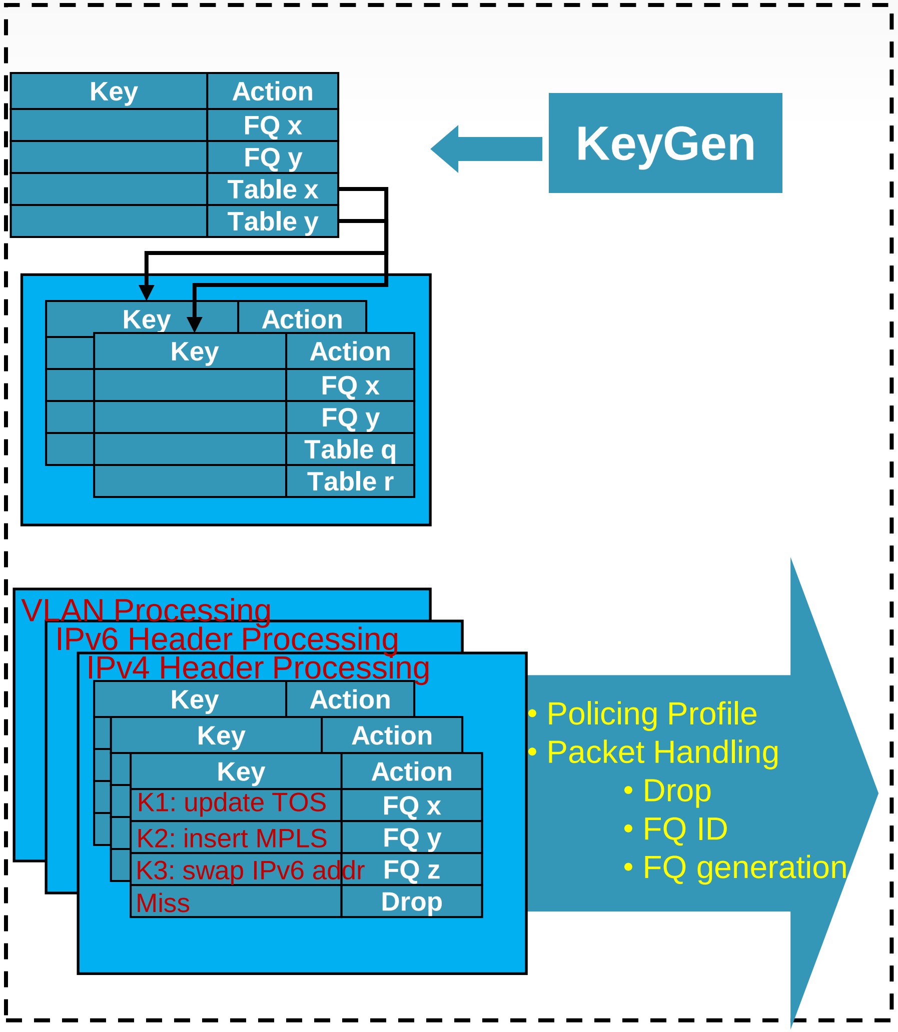

Key	Action
	FQ x
	FQ y
	Table x
	Table y

Key	Action
	FQ x
	FQ y
	Table q
	Table r

Key	Action
K1: update TOS	FQ x
K2: insert MPLS	FQ y
K3: swap IPv6 addr	FQ z
Miss	Drop

- Policing Profile
- Packet Handling
  - Drop
  - FQ ID
  - FQ generation

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Case Study: Spread Control and User Plane Traffic

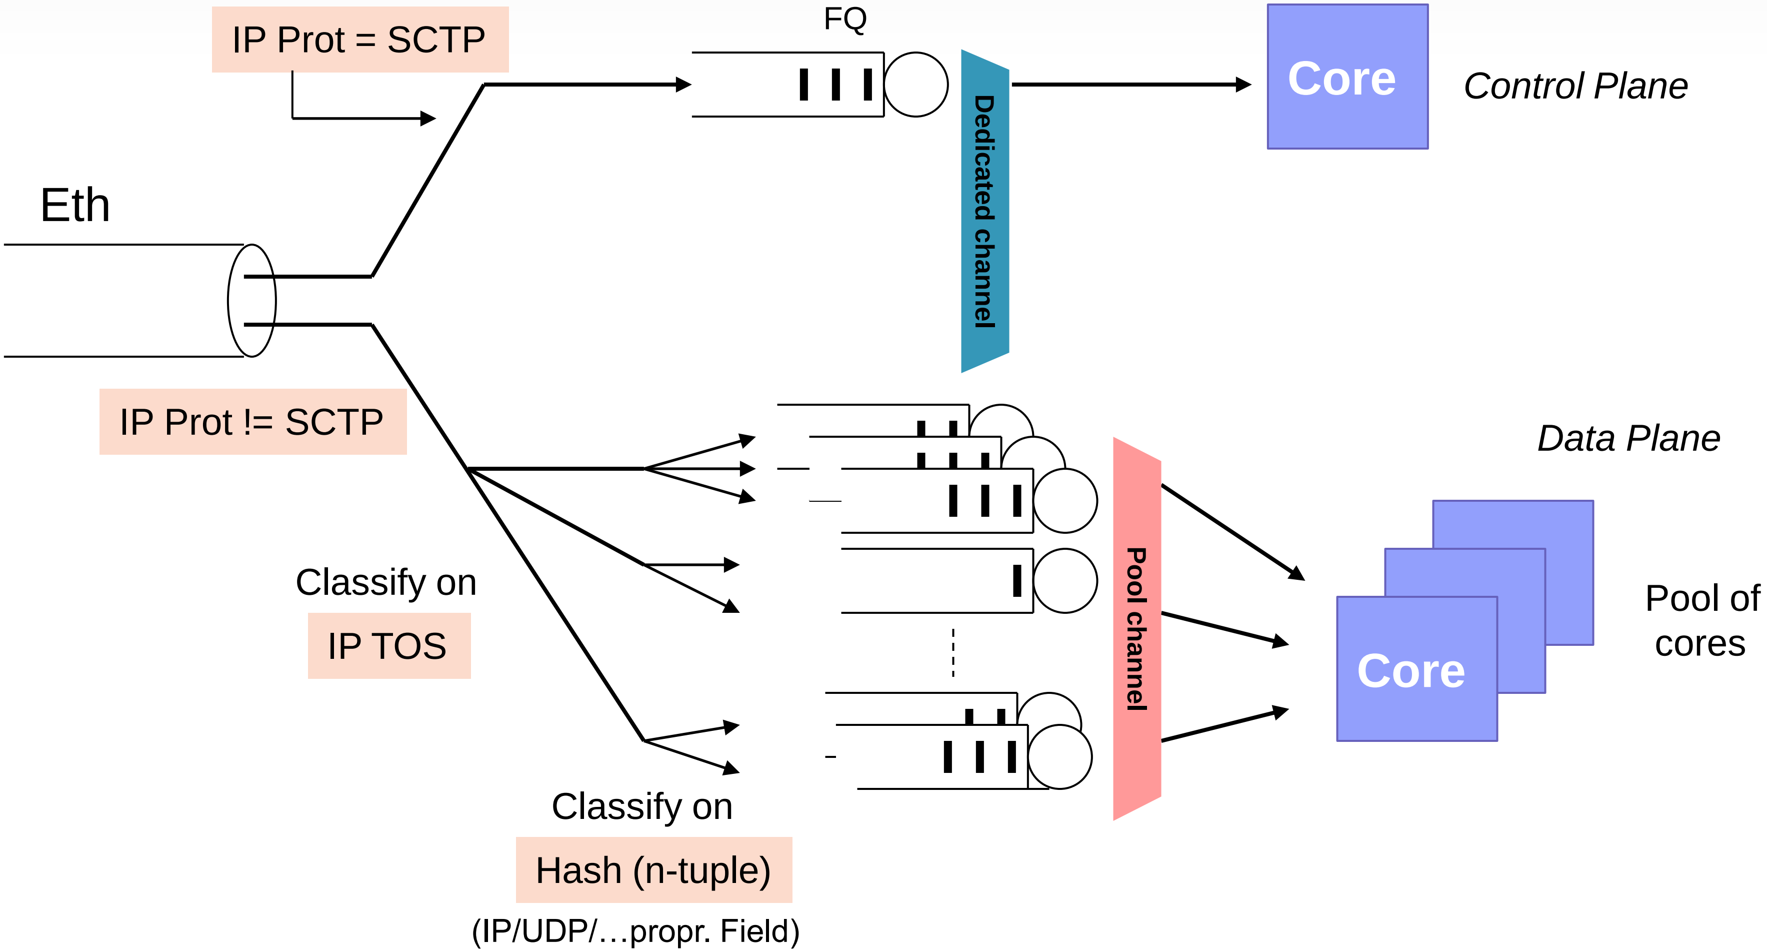

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Case Study: High Level Mapping to DPAA

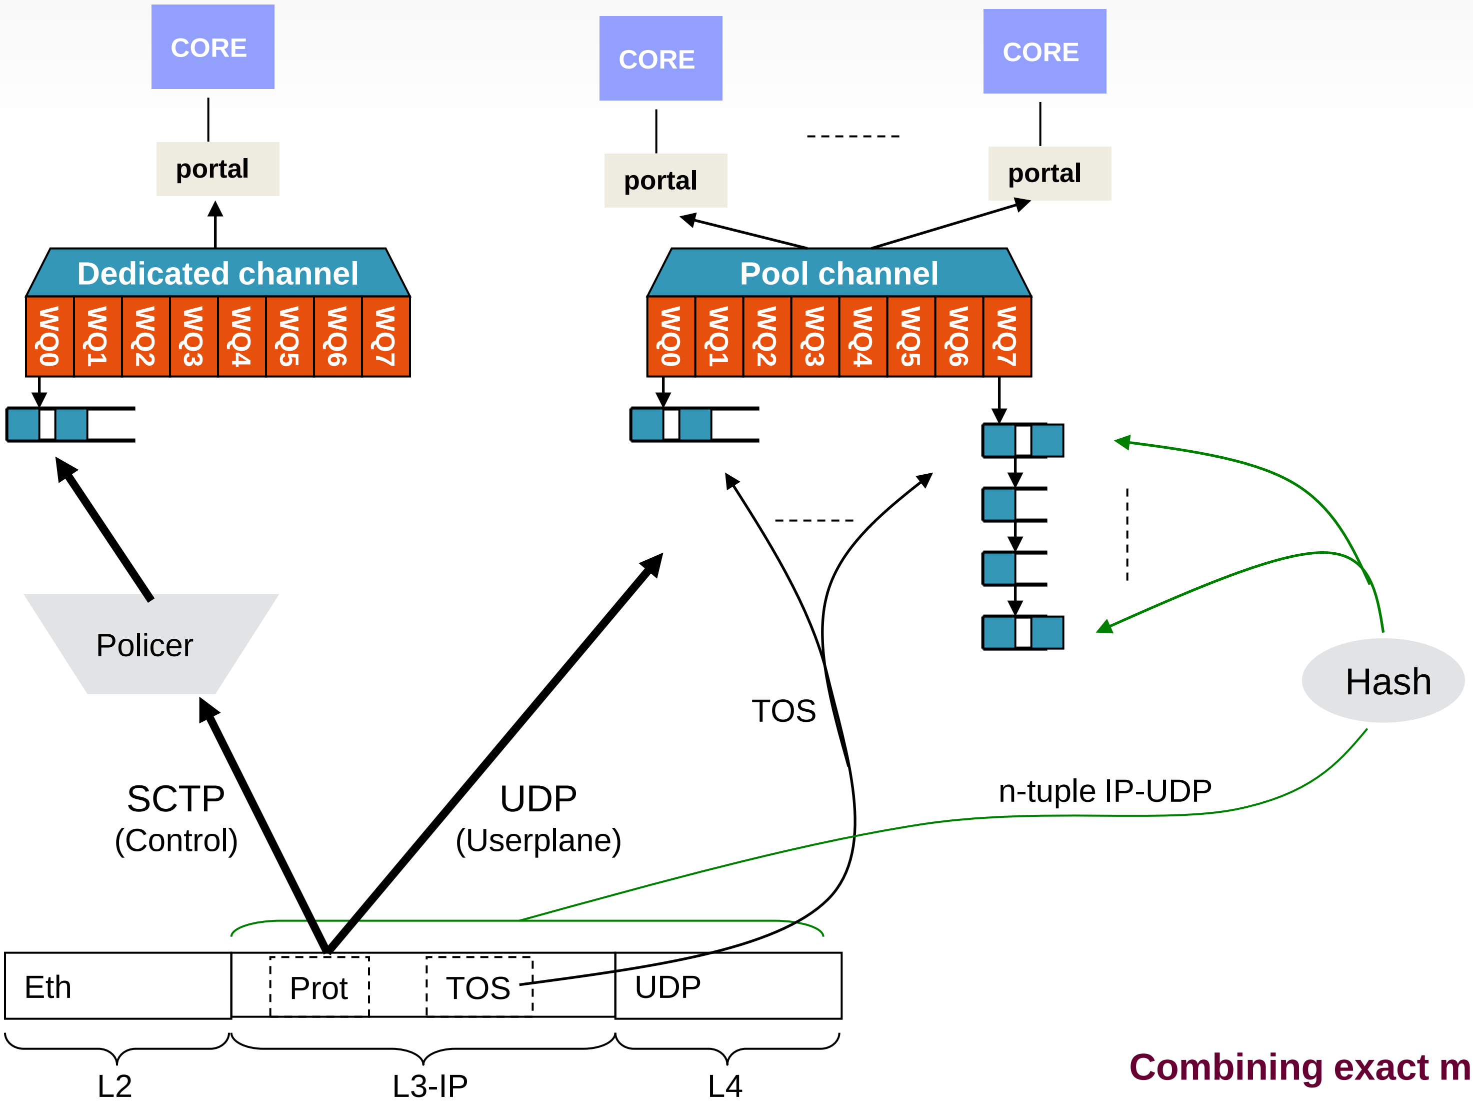

Combining exact match + Hash

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Case Study: Configure FMan with FMC Tool

- rootfs/etc/fmc/config/8c-128fq-p.xml

```
…
<distribution name="ipv4eth0">
        <queue count="128" base="0x3800"/>
        <key>
                <fieldref name="ipv4.src"/>
                <fieldref name="ipv4.dst"/>
                <fieldref name="ipv4.tos"/>
        </key>
        </distribution>

        <distribution name="ipv4eth1">
        <queue count="128" base="0x3880"/>
        <key>
                <fieldref name="ipv4.src"/>
                <fieldref name="ipv4.dst"/>
                <fieldref name="ipv4.tos"/>
        </key>
        </distribution>
…
```

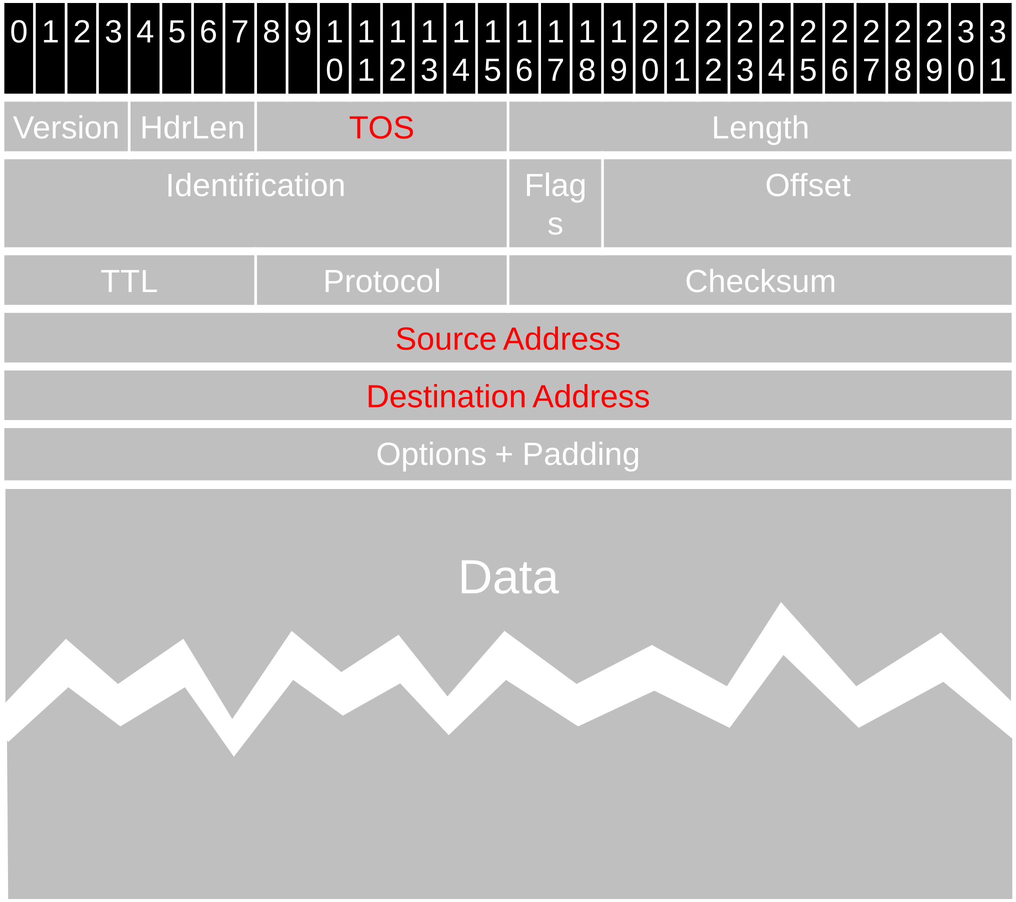

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

Version | HdrLen | TOS | Length

Identification | Flags | Offset

TTL | Protocol | Checksum

Source Address

Destination Address

Options + Padding

Data

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Virtual Storage Profiles

- Virtual storage profiles are not supported in P1023, P4080, P3041, P5020, P5040 and P2041.
  - In these devices the storage profile is configurable only on a per-port basis (named ‘hardware port storage profile’).
- The virtual storage profile mechanism allows the virtualization of the buffer pool selection for frame storage from the physical hardware ports.
  - The virtual storage profile ID (SPID) is selected as a result of the classification on the frame headers.
  - The same Storage Profile ID (SPID) values from the classification on different physical ports, may yield to different storage profile selection.

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

New Features for T4240 Queue Manager (QMan)

- T4xxx has a total of 50 Software Portals (SP), increase from 10 SP found in the P class processors.
- Supports Customer Edge Egress Traffic Management (CEETM) that provides hierarchical class based scheduling and traffic shaping:
  - Available as an alternate to FQ/WQ scheduling mode on the egress side of specific direct connect portals
  - Enhanced class based scheduling supporting 16 class queues per channel
  - Token bucket based dual rate shaping representing Committed Rate (CR) and Excess Rate (ER)
  - Congestion avoidance mechanism equivalent to that provided by FQ congestion groups
- A total of 48 algorithmic sequencers are provided, allowing multiple enqueue/dequeue operations to execute simultaneously.
- Support up to 295M enqueue/dequeue operations per second.

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Queuing Structure

- **Frame Descriptor (FD)**
  - The basic queue Element that describe a frame
  - Usually a single IO packet will use a single frame
  - Other scenarios: commands with no buffer
- **Frame Queue Descriptor (FQD)**
  - A linked list of FD's
  - Usually a frame queue is associated with a **flow or interface**
  - Enqueue operation must include the target FQ as a parameter
  - Dequeue operation may use FQ as a parameter for operation
  - Head of frame queue can be associated to ODP
- **Work Queue Structure**
  - Linked list of FQD
  - Hold flows of the same priority and designation
  - Dequeue operation may use WQ as a parameter for operation
- **Channel**
  - Set of eight WQ channel served by a single type of entity
  - Dequeue from channel can be configured to be:
    - strict priority
    - round robin (Simple, Weighted or Deficit)
  - Dequeue may use channel as a parameter for operation

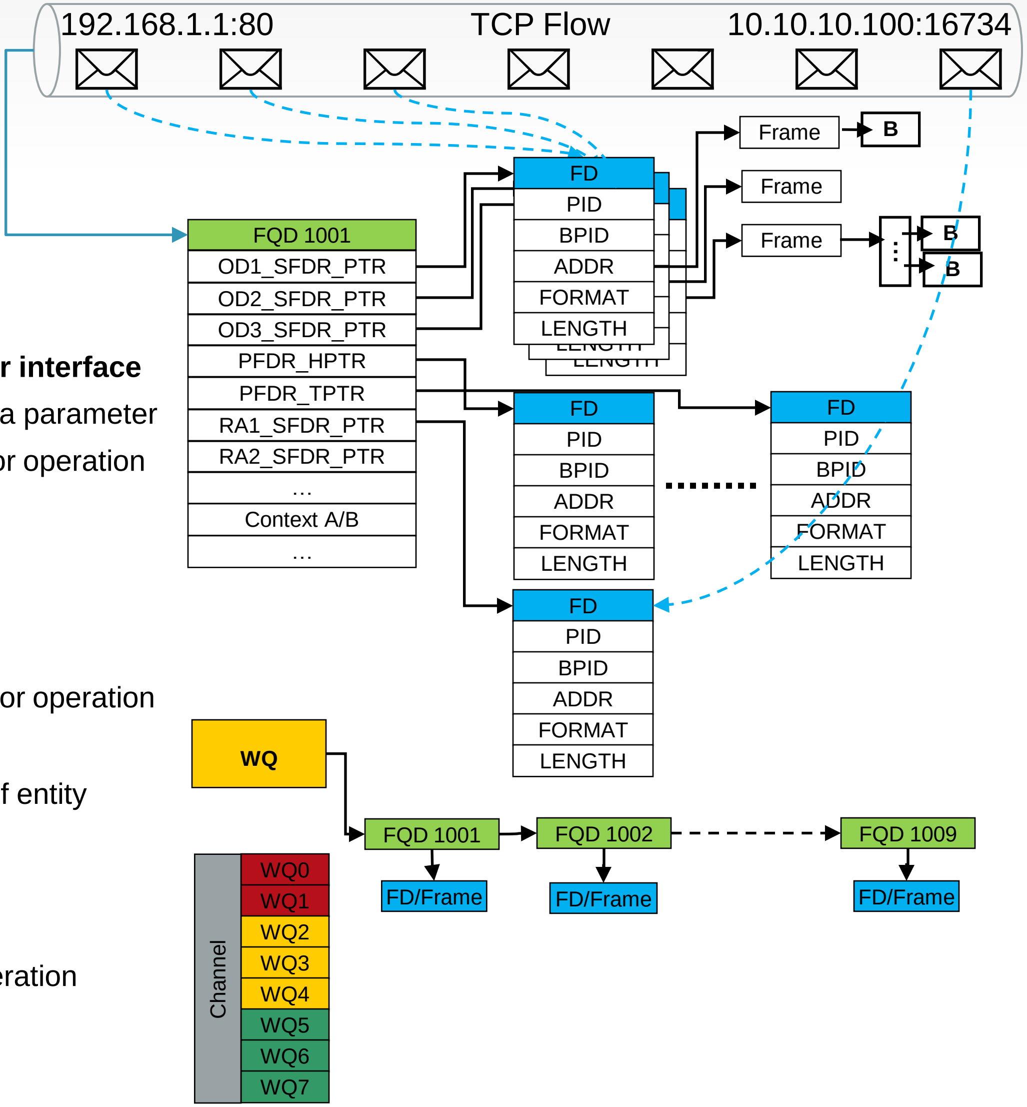

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

QMan Communicates with Portals

- Portals are the interface between QMan and the accelerator which use them
  - Direct Connect Portals has direct connect signals to Dedicated Channel
    - Dedicated Channels are always serviced by a single entity, e.g. FMan, etc.
  - Software Portals use CoreNet as the physical interconnect to the processor core.
    - Each Software Portal serves a Dedicated Channel, and optionally services Pool Channels.
    - Software and QMan interact by “reading” and “writing” data across CoreNet
- Each channel consists of 8 WQs, and thus there are 8 possible priorities.
- QMan contains a total of 110 channels for T4240
  - 16 Dedicated Channels (Sub Portals) per Frame Manager’s Direct Connect Portal
  - 1 for SEC (8 SPs), 1 for PME, 1 for DCE, 1 for RMan
  - Up to 50 CoreNet dedicated channels for software portals
  - 15 for CoreNet pool channels which are shared by all software portals

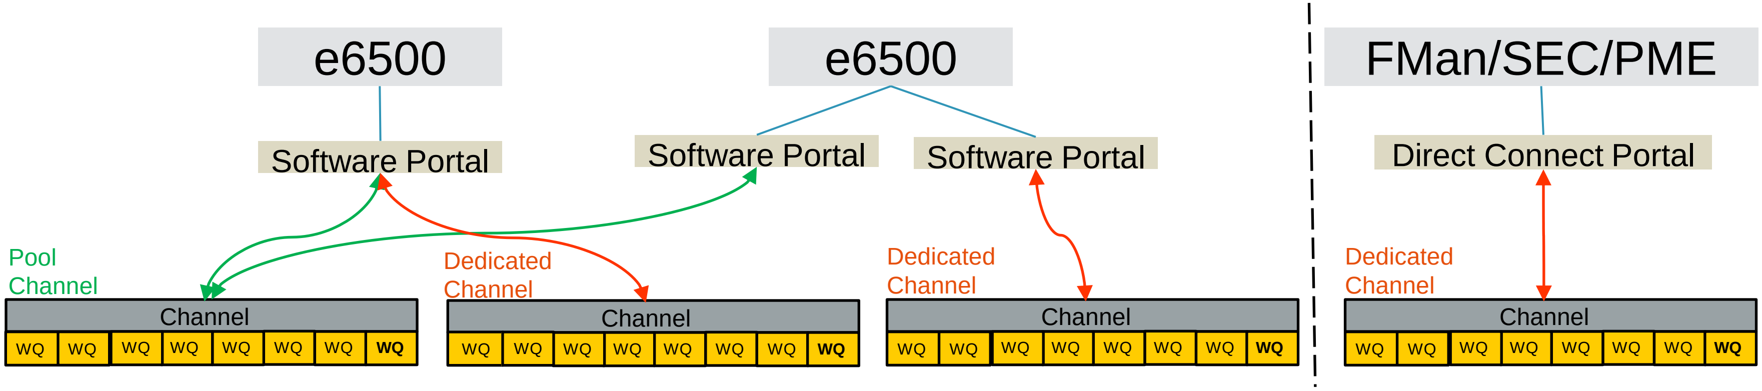

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Queue Management

- QMan provides a way to inter-connect DPAA components
  - Cores (including IPC)
  - Hardware offload accelerators
  - Network interfaces – Frame Manager
- Queue management
  - High performance interfaces (“portals”) for enqueue/dequeue
  - Internal buffering of queue/frame data to enhance performance
- Congestion avoidance and management
  - RED/WRED
  - Tail drop for single queues and aggregates of queues
  - Congestion notification for “loss-less” flow control
- Load spreading across processing engines (cores, HW accelerators)
  - Order restoration
  - Order preservation/atomicity
- Delivery to cache/HW accelerators of per queue context information with the data (Frames)
  - This is an important offload for software using hardware accelerators

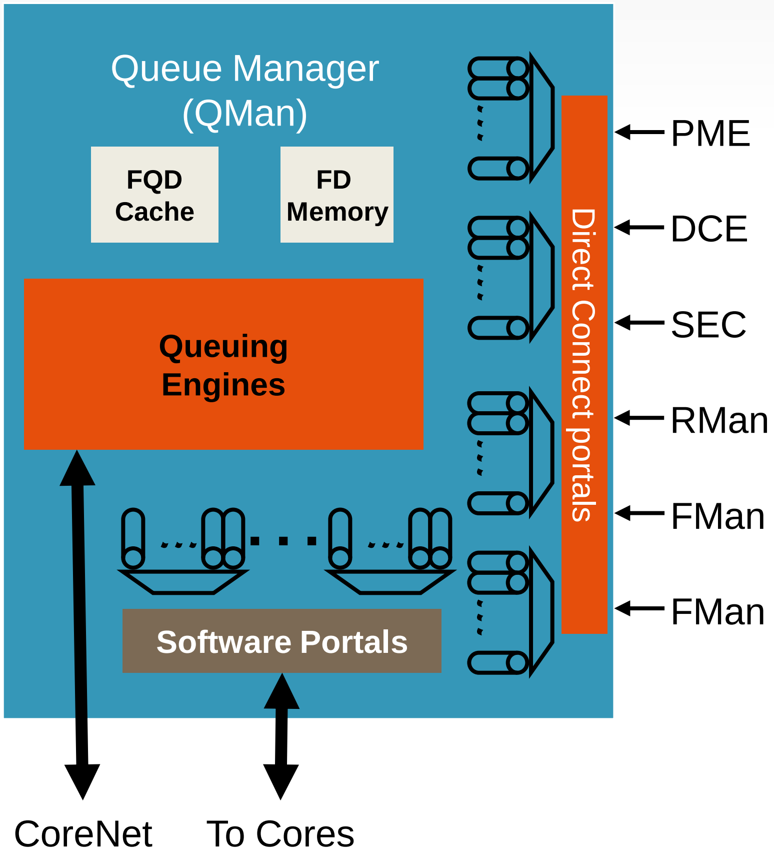

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

QMan Software Portal Components

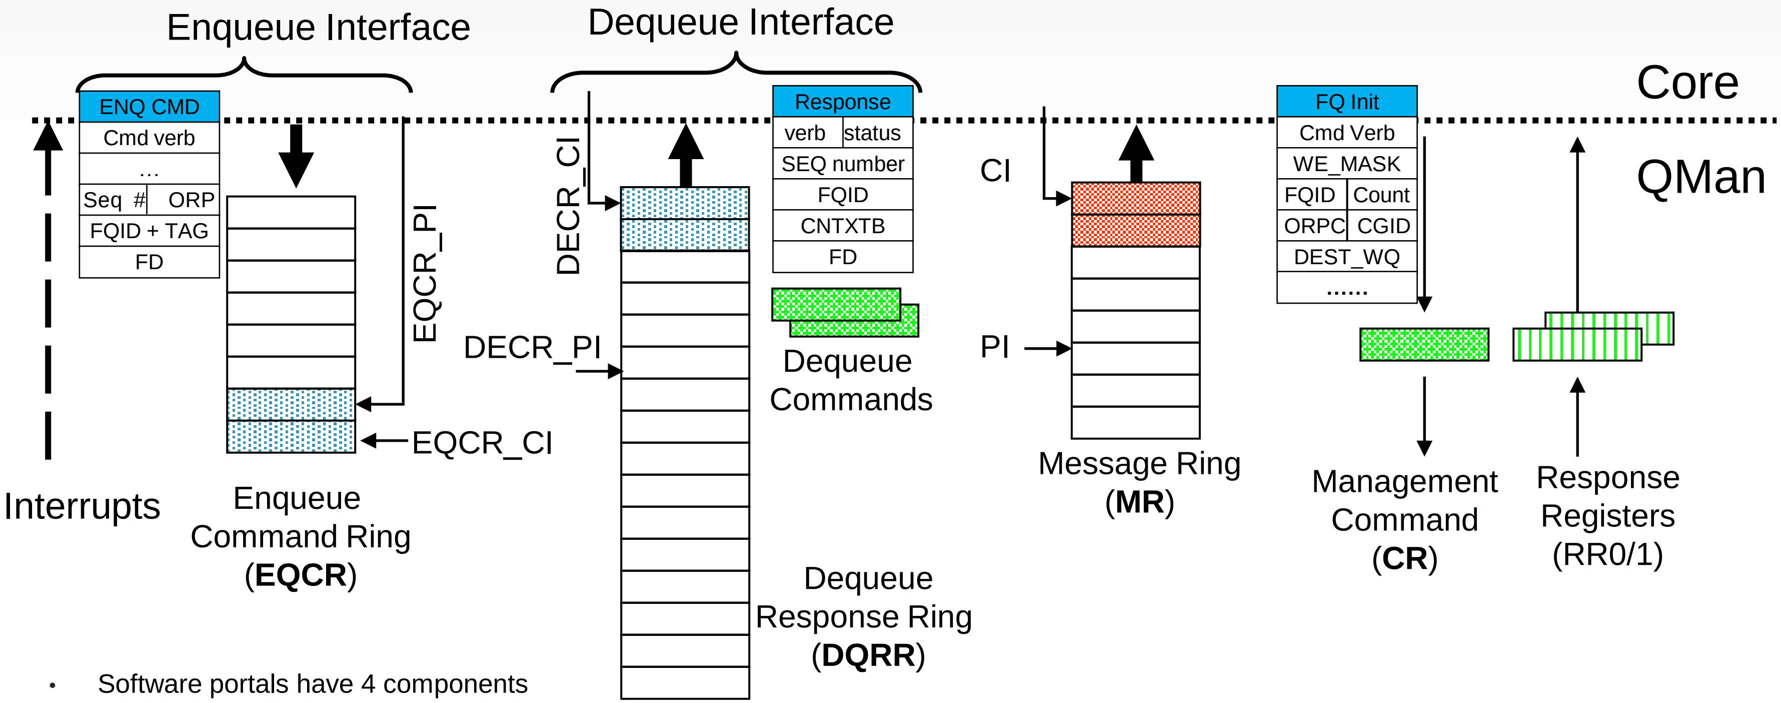

- Software portals have 4 components
  - Enqueue: EQCR
  - Dequeue: Command registers + DQRR
  - Messages: MR (e.g. enqueue rejections)
  - Management commands: command/response registers
- Interrupts can be used to signal availability of data or space (in EQCR)
- Rings provide finite size FIFOs
  - Up to 16 entries for DQRR, 8 entries for EQCR and MR
- Portal components are implemented inside QMan to reduce access latency
  - Unlike traditional BD rings which are in “memory” and “registers”
- QMan can “push” (stash) DQRR entries across CoreNet into the appropriate core’s cache
- PI and CI are the basic mechanisms used with rings but other forms of notification of data availability and data consumption are supported
- When these other mechanisms are used QMan maintains PI/CI

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

QMan Cache Warming

- In addition to stashing DQRR entries into cache, QMan’s software portals can also “warm” a core’s (L1 or L2) cache with frame and queue related data
  - Actual frame data for single buffer frames
  - Scatter gather list for multi-buffer frames
  - Frame “annotations”
    - Data between Address and Offset at start of frame
    - Used to pass additional information about the frame which is not “frame data” e.g. FM parse results
  - Data referenced by FQ Context
- Stashing options can be configured on a per FQ basis
- Cache warming are performed at the time that the frame is dequeued (i.e., DQRR entry is created)

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Core Load Distribution

- Static flow-based distribution (Dedicated Channel)
  - Set of WQs with different FQs directed statically to different cores
  - Distribution of frames (selection of FQ) is based on hash keys, ensuring that packets from the same traffic flow will always go to the same cores
  - Static not dynamic, doesn't react to core load, assign work to the cores in a static or fixed manner
- Adaptive load balancing (Pool Channel)
  - Load spread the packets (or the Frame Queue) to the cores based on actual core availability/readiness
  - QMan provides two mechanisms to deal with out of order packets:
    - Order preservation: ensure that related packets are processed in order (and typically one at a time); can also provide “atomicity” – atomic access to data
    - Order restoration: allows frames to be processed out of order and then restores their order later on before they are transmitted

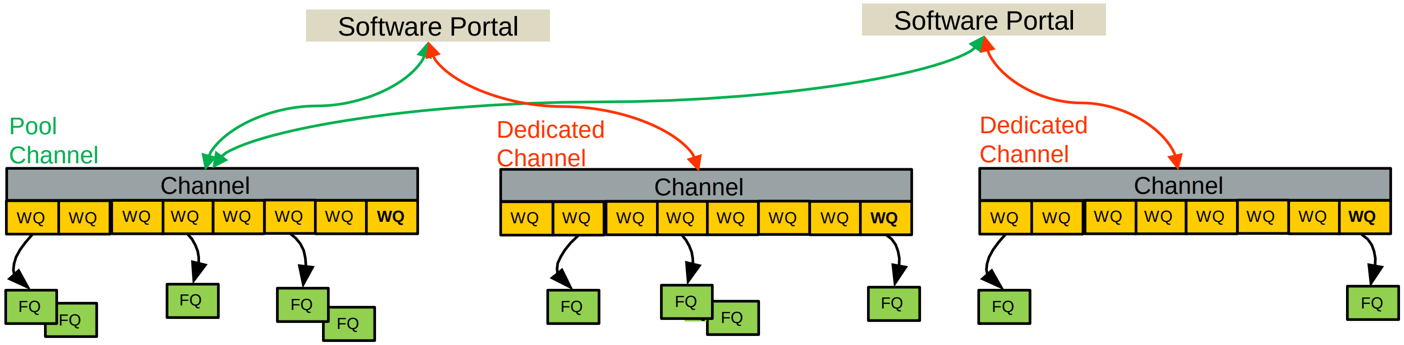

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Dequeue Modes

- QMan supports 2 modes on software portals:
  - Push Mode:
    - Qman continues to push entries into DQRR in attempt to keep it “full”
    - QMan provides 2 command registers
      - One register is “static” and QMan repeatedly executes this command
      - One register is “volatile” and QMan executes that command a limited number of times
    - Push mode is “just like” a BD ring
  - Pull Mode:
    - QMan provides a single command register
    - Software must issue a new command for each dequeue operation

- Push mode is the most common mode
- Pull mode offers more control to applications

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Dequeue Scheduling

- Class Scheduler schedules **WQ**
  - A Class scheduler per channel
  - Two levels of scheduling:
    - Use Weighted Interleaved Round Robin to schedule within the medium priority queues (2 to 4) and low priority queues (5 to 7)
    - Strict priority of all 8 WQs with programmable elevation (CS_ELEV) of the low priority tier over the medium priority tier
  - Maintain active FQ states transition
    - Keeps track of the last RR winner, selection counters, elevation counter
- Intra-Class scheduling schedules **FQ**
  - Schedule a frame queue within a work queue
  - Use Modified Deficit Round Robin with ICS_CRED + ICS_SURP

15-bit Credit	**16-bit Surplus**

  - First Dequeue **surplus = surplus + credit**
  - Dequeue 1-3 frames,
  - subtract frame length(s) bytes from surplus
  - If surplus > 0, dequeue 1-3 frames more
  - If surplus <=0, reschedule Frame Queue

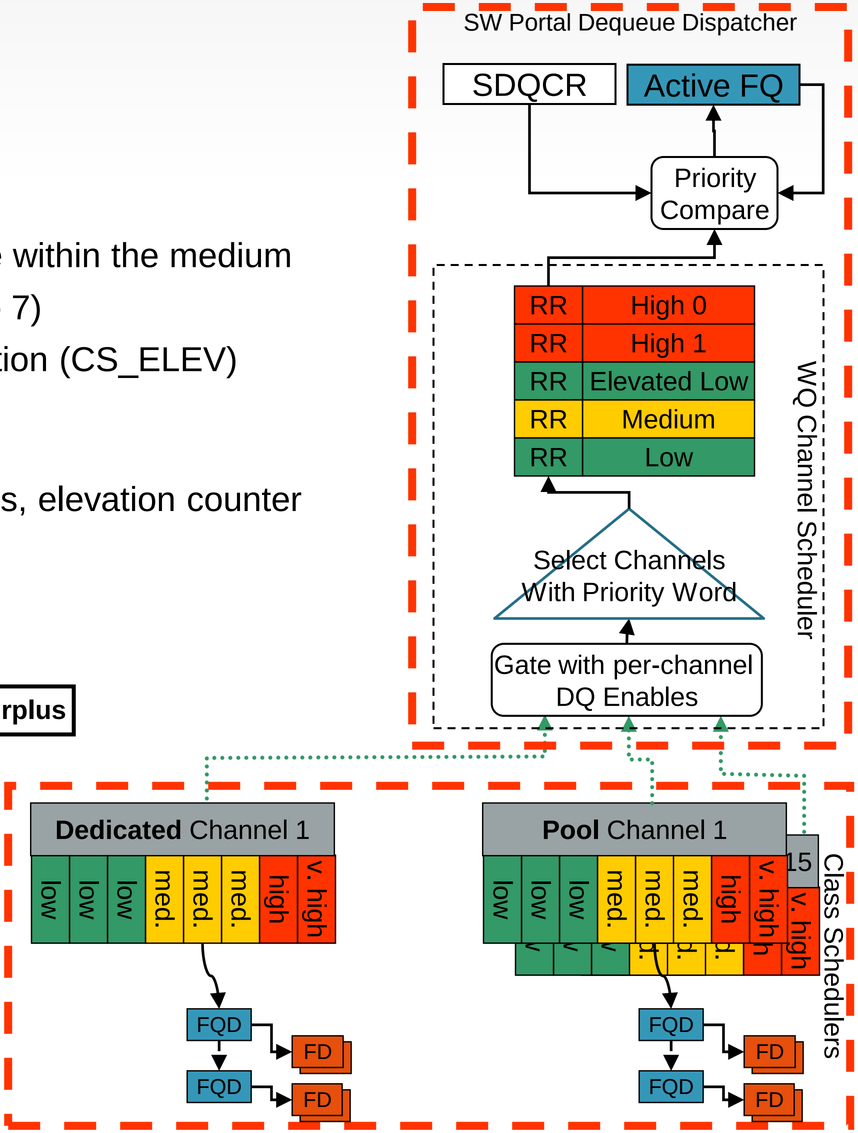

SW Portal Dequeue Dispatcher

SDQCR

Active FQ

Priority Compare

RR	High 0
RR	High 1
RR	Elevated Low
RR	Medium
RR	Low

Select Channels With Priority Word

Gate with per-channel DQ Enables

WQ Channel Scheduler

Dedicated Channel 1

low | low | low | med. | med. | med. | high | v. high

Pool Channel 1

15

low | low | low | med. | med. | med. | high | v. high | v. high

FQD

FD

Class Schedulers

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Static Dequeue: Who Is on First?

e6500 core X

Push Mode

SDQCR						
-	S S	F C	D P	-	-	DCT
Token						
DQ SRC						
DQ SRC (cont)						

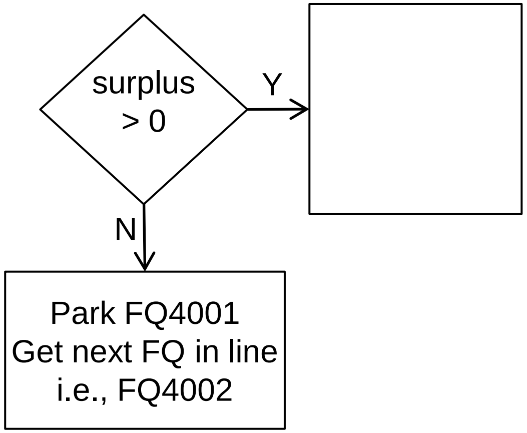

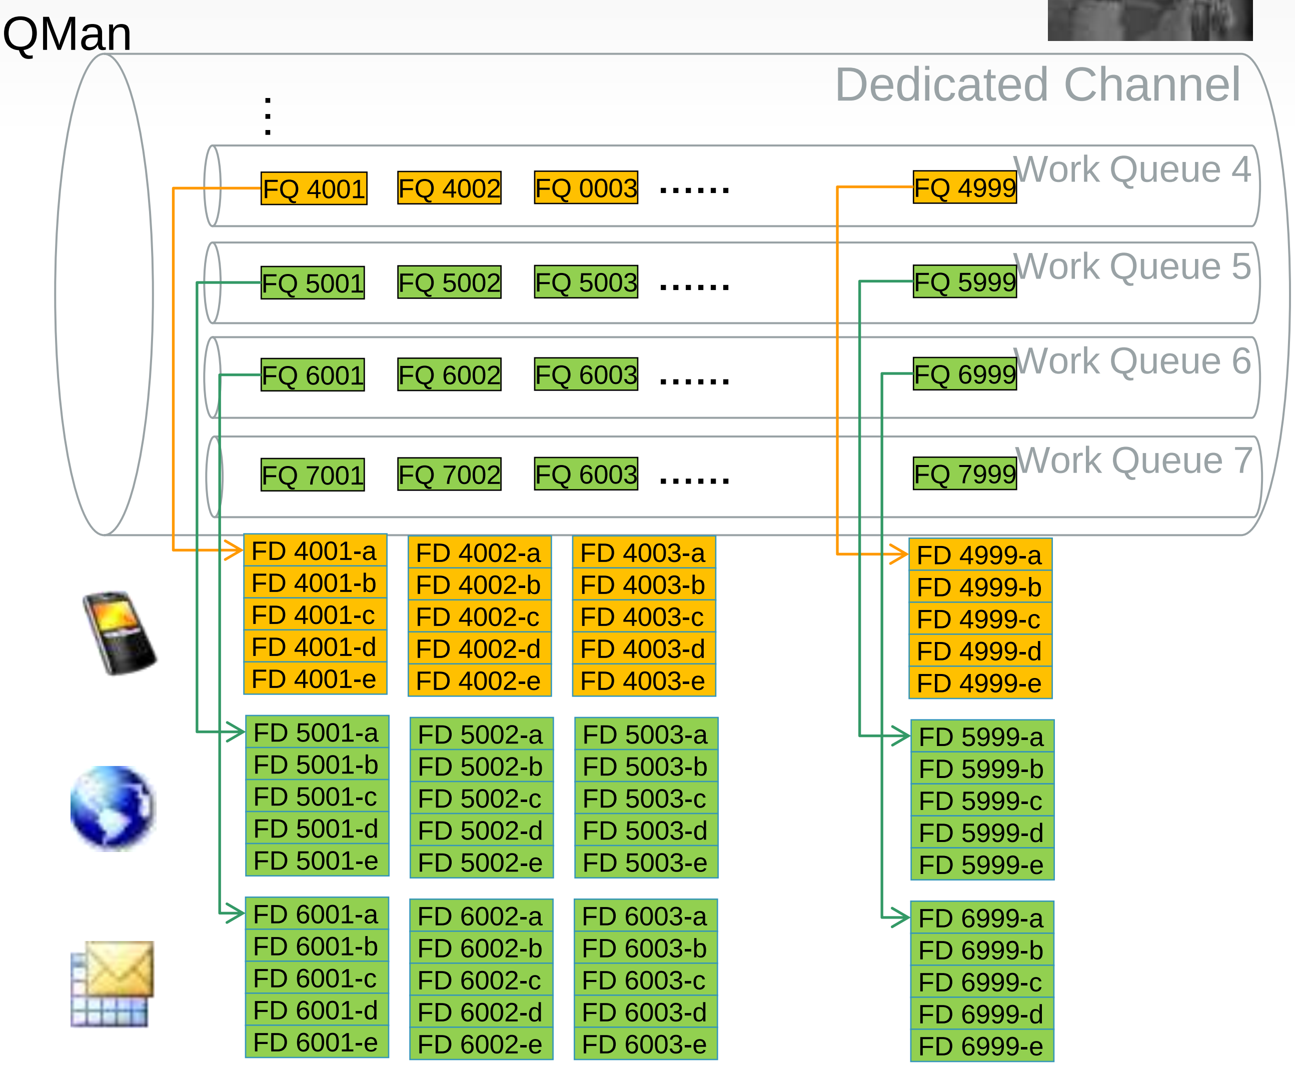

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Volatile Dequeue: Who Is on First?

e6500 core X | QMan

Push Mode

VDQCR		
P	E	Number of Frames
FQID=FQ6002		
……		
……		

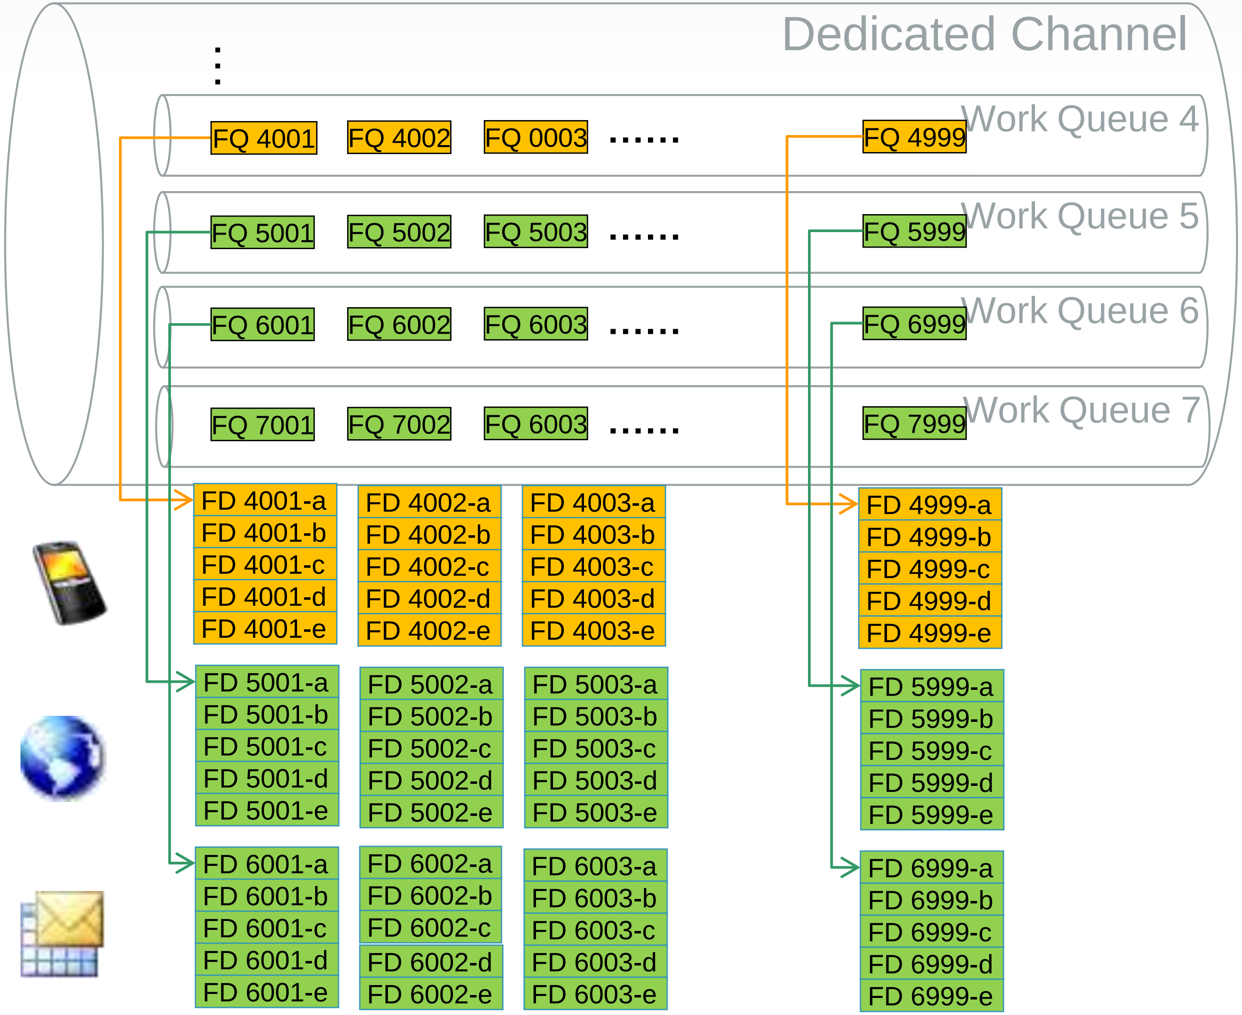

Dedicated Channel

⋮

FQ 4001 FQ 4002 FQ 0003 …… FQ 4999 Work Queue 4

FQ 5001 FQ 5002 FQ 5003 …… FQ 5999 Work Queue 5

FQ 6001 FQ 6002 FQ 6003 …… FQ 6999 Work Queue 6

FQ 7001 FQ 7002 FQ 6003 …… FQ 7999 Work Queue 7

FD 4001-a FD 4001-b FD 4001-c FD 4001-d FD 4001-e	FD 4002-a FD 4002-b FD 4002-c FD 4002-d FD 4002-e	FD 4003-a FD 4003-b FD 4003-c FD 4003-d FD 4003-e	FD 4999-a FD 4999-b FD 4999-c FD 4999-d FD 4999-e
FD 5001-a FD 5001-b FD 5001-c FD 5001-d FD 5001-e	FD 5002-a FD 5002-b FD 5002-c FD 5002-d FD 5002-e	FD 5003-a FD 5003-b FD 5003-c FD 5003-d FD 5003-e	FD 5999-a FD 5999-b FD 5999-c FD 5999-d FD 5999-e
FD 6001-a FD 6001-b FD 6001-c FD 6001-d FD 6001-e	FD 6002-a FD 6002-b FD 6002-c FD 6002-d FD 6002-e	FD 6003-a FD 6003-b FD 6003-c FD 6003-d FD 6003-e	FD 6999-a FD 6999-b FD 6999-c FD 6999-d FD 6999-e

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Customer Edge Egress Traffic Management (CEETM)

- QMan 1.2 (i.e. QorIQ T42xx) supports egress traffic management by provides hierarchical class based scheduling and traffic shaping.
- On a specific QMan Direct Connect Portals (DCPs), each sub-portal that supports CEETM can be configured to use either the <u>regular FQ/WQ scheduling mode</u> OR <u>CEETM scheduling mode</u>.
  - A given sub-portal can switch between FQ/WQ and CEETM scheduling mode.
  - CEETM supports up to 8 logical network interfaces (LNI) that can each be mapped to a DCP sub-portal, whereas a DCP can support up to 16 sub-portals.
  - CEETM is supported on only a subset of the DCP portals, not on all DCP portals.
  - A single instance of the CEETM is associated with a single DCP and therefore a single egress I/O module.
- CEETM maintains the following functionality equivalent of regular FQ/WQ scheduling mode:
  - Congestion management capabilities including WRED
  - Dequeued frame context (Context_A and Context_B)
  - Priority or traffic class flow control
- FQID xF00000 – xFFFFFF are reserved for CEETM as Logical FQIDs (LFQID).
  - xF00000 – xF00FFF (4k) for DCP portal 0 (i.e. FMAN0)
  - xF10000 – xF10FFF (4k) for DCP portal 1 (i.e. FMAN1)

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

CEETM Scheduling Hierarchy (QMAN 1.2)

- Logics
  - **Green** denotes logic units and signal paths that relate to the request and fulfillment of *Committed Rate (CR)* packet transmission opportunities.
  - **Yellow** denotes the same for *Excess Rate (ER).*
  - **Black** denotes logic units and signal paths that are used for unshaped opportunities or that operate consistently whether used for CR or ER opportunities.
- Scheduler
  - Channel Scheduler: channels are selected to send frame from Class Queues.
  - Class scheduler: frames are selected from Class Queues . Class 0 has highest priority.
- Algorithm
  - Strict Priority (SP)
  - Weighted Scheduling
  - Shaped Aware Fair Scheduling (SAFS)
  - Weighted Bandwidth Fair Scheduling (WBFS)

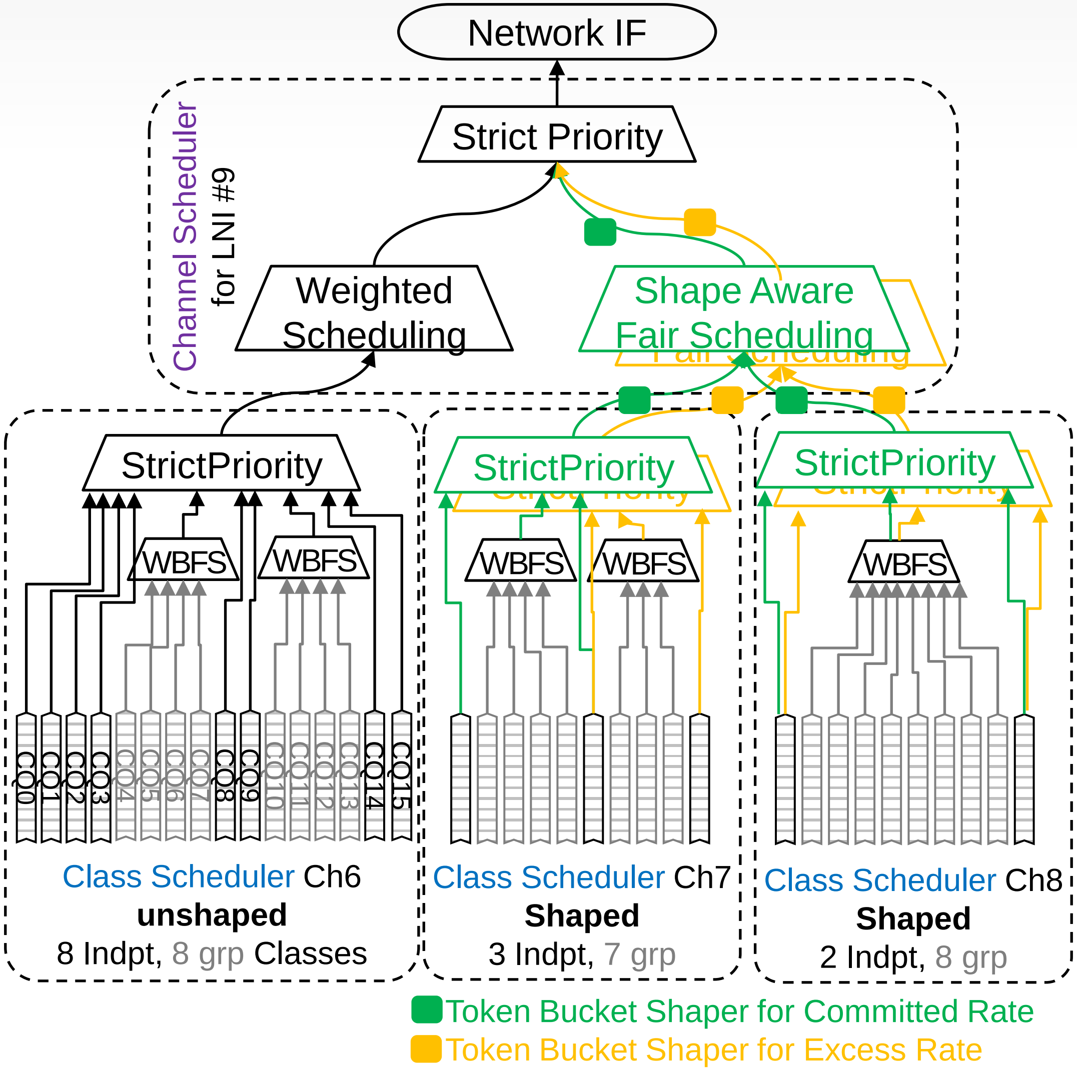

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Weighted Bandwidth Fair Scheduling (WBFS)

- Weighted Bandwidth Fair Scheduling (WBFS) is used to schedule packets from queues within a priority group (A or B group) such that each gets a “fair” amount of bandwidth made available to that priority group.
- The premises for fairness for algorithm is:
  - available bandwidth is divided and offered equally to all classes.
  - offered bandwidth in excess of a class’s demand is to be re-offered equally to classes with unmet demand.

	Initial Distribution		First ReDistribution		Second Redistribution		Total BW Attained
BW available	10G		1.5G		.2G		0G
Number of classes with unmet demand	5		3		2		
Bandwidth to be offer to each class	2G		.5G		.1G		
	Demand	Offered & Retained	Unmet Demand	Offered & Retained	Unmet Demand	Offered & Retained	
Class 0	.5G	.5G	0				.5G
Class 1	2G	2G	0				2G
Class 2	2.3G	2G	.3G	.3G	0		2.3G
Class 3	3G	2G	1G	.5G	.5G	.1G	2.6G
Class 4	4G	2G	2G	.5G	1.5G	.1G	2.6G
Total Consumption	**11.8G**	**8.5G**		**1.3G**		**.2G**	**10G**

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Addressing Ordering Requirements

- There are two basic approaches to addressing this requirement:
  - Order restoration
    - Take note of the correct order (or sequence) of packets before processing starts and restore the packets to that order before they are transmitted
  - Order preservation
    - Ensure that related packets are processed in order (and typically one at a time)
    - Order preservation can also provide “atomicity” – atomic access to data used in processing the frame

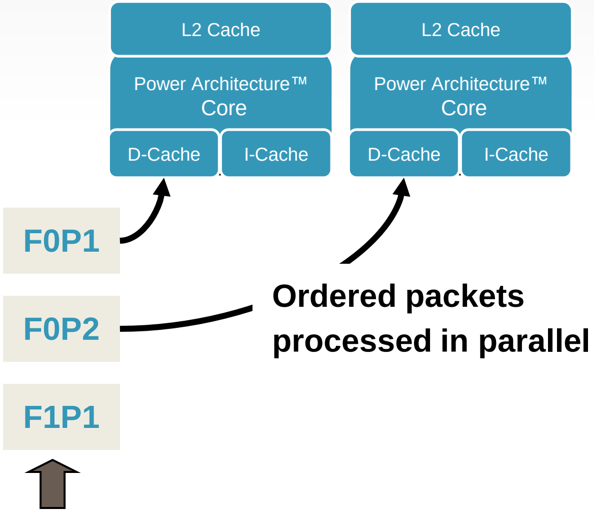

Ordered packets processed in parallel

- QMan requires that related frames (which must be transmitted in order) be placed on the same frame queue for both of these approaches
  - This does not mean that only related frames are placed on a given FQ
  - Many sets of related frames can be placed on an FQ
  - Frame Manager is responsible for achieving this

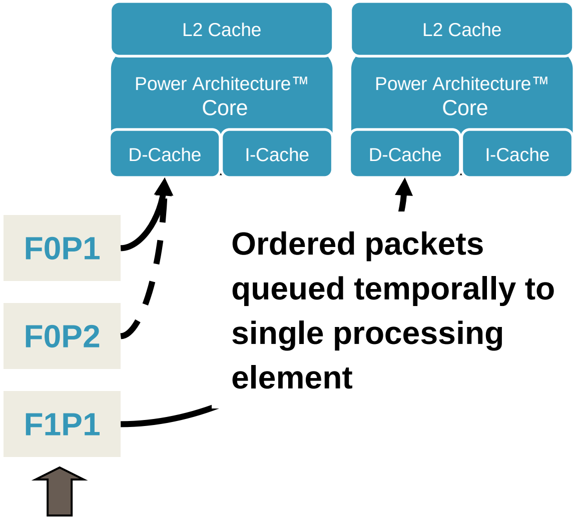

Ordered packets queued temporally to single processing element

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

QMan Order Restoration

- QMan’s order restoration support has two components:
- Order Definition Point (ODP)
  - A point defined relative sequence to each Frames’ pass
  - ODP id is associated to FQ-ID
  - assigning a monotonically increasing 14 bits sequence number to a series of frames
  - QM supports single ODP on head of queue
  - ODP can be made anywhere in the system i.e. SW can be an ODP
- Order Restoration Point (ORP)
  - A point relative sequence associated to single ODP is restored
  - Allows frames insertion into the flow (single sequence number with more/last indication)
- Behavior highlights
  - Configurable number of “in-flight” packets per ORP
  - resection of ORP is part of enqueue command but Queue tail is not associated to ORP i.e., enqueue to single destination queue can respect many ORP’s
  - Note: Frames are not “marked”

ODP-A

ODP-B

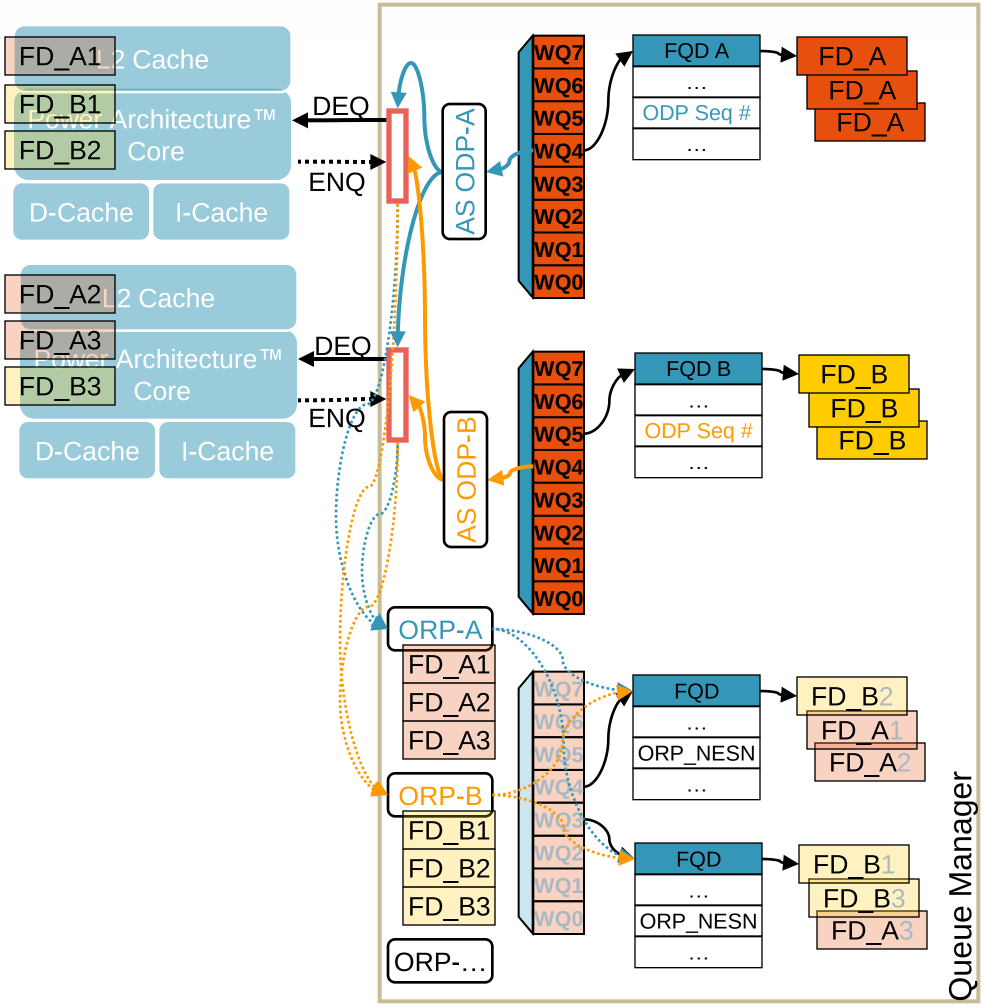

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Order Restoration Configuration

- Treatment of a FD is determined by:
  - ORP Restoration Window Size (ORPRWS)
  - ORP Auto Advance NESN Window Toggle (OA)
  - ODP Sequence Number (SEQ)
  - ORP Next Expected Sequence Number (NESN)
  - ORP Acceptable Late Arrival Window Size (OLWS)
  - ORP Early Arrival Head Sequence Number (EA_HSEQ)
  - ORP Early Arrival Tail Sequence Number (EA_TSEQ)
  - ORP Early Arrival Head Pointer (EA_HPTR)
  - ORP Early Arrival Tail Pointer (EA_HPTR)

Note: SEQ is 14 bit wide. Create a sliding (8k - 1) sliding windows before
and after NESN

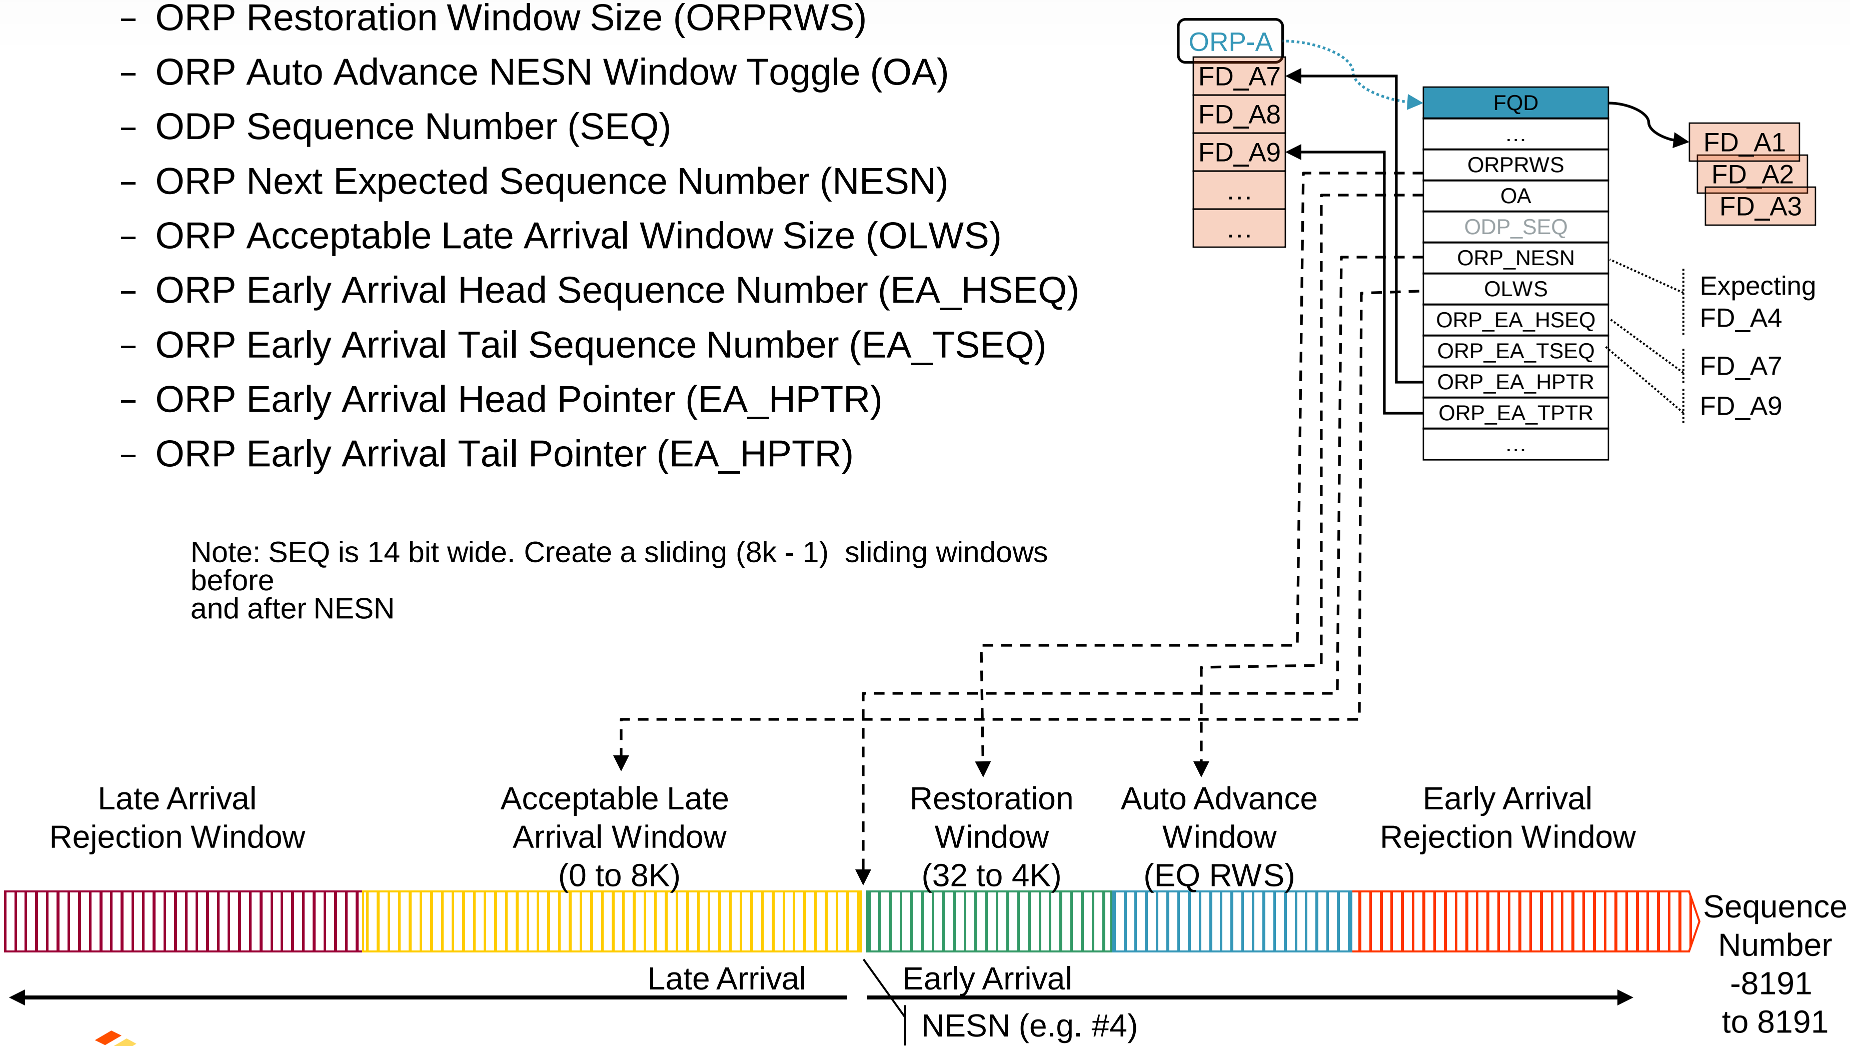
ORP-A
FD_A7
FD_A8
FD_A9
…
…
FQD
…
ORPRWS
OA
ODP_SEQ
ORP_NESN
OLWS
ORP_EA_HSEQ
ORP_EA_TSEQ
ORP_EA_HPTR
ORP_EA_TPTR
…
FD_A1
FD_A2
FD_A3
Expecting FD_A4
FD_A7
FD_A9
Late Arrival Rejection Window
Acceptable Late Arrival Window (0 to 8K)
Restoration Window (32 to 4K)
Auto Advance Window (EQ RWS)
Early Arrival Rejection Window
Sequence Number -8191 to 8191
Late Arrival
Early Arrival
NESN (e.g. #4)

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Congestion Management and Avoidance

- Both FM and QMan are involved in supporting congestion management and avoidance
- QMan provides the following support
  - Congestion management **(loss-less flow control, threshold/tail drop)**
  - Congestion avoidance **(RED/WRED)**
- Congestion Groups (CG) define granularity
  - Every frame queue has a tail drop threshold and configured to a congestion group
  - Congestion calculations are done on groups of queues
  - Congestion avoidance/management calculations configured in the CG
- Congestion Group Details
  - 256 Congestion Groups
  - Time aware weighted average queue depth of all queues in the CG
  - 3 color configurable WRED curves
  - Enqueued packet may be rejected (discarded) due to WRED policy
  - Instantaneous CG depth +/- hysteresis value can initiate congestion state messages to enqueue sources
    - Lossless flow control on interfaces supporting PAUSE semantics

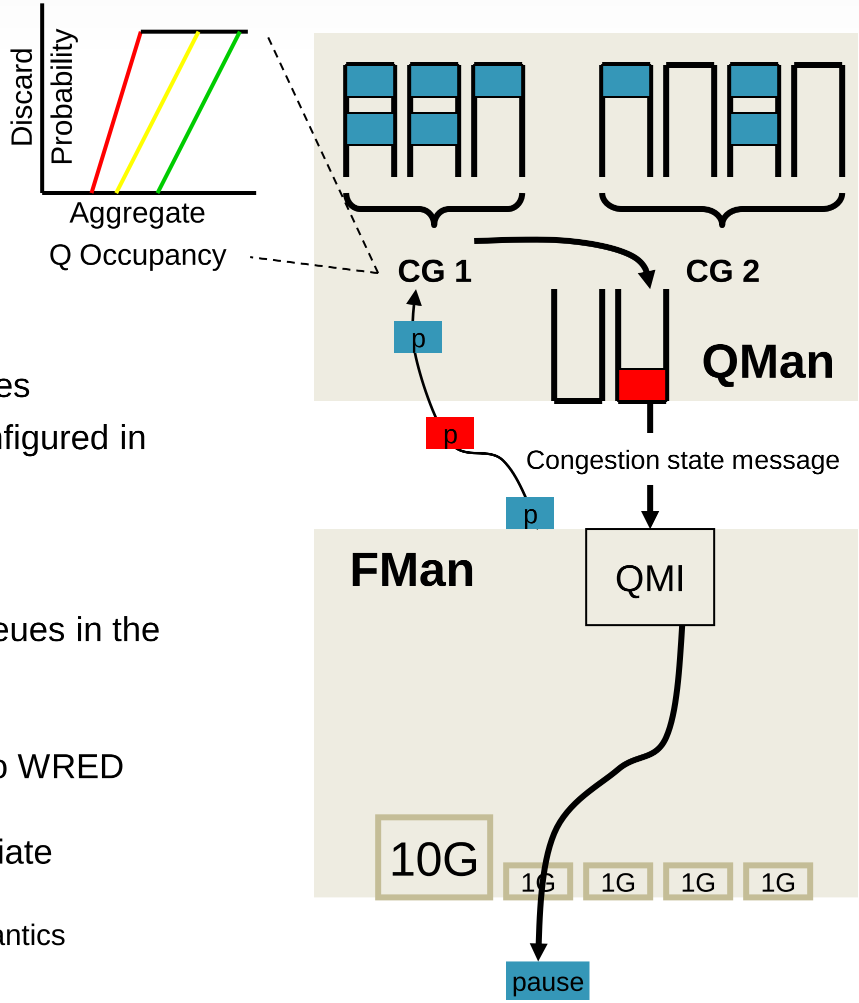

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Setup Congestion Group Record

- On enqueue:
  - The color for the frame being enqueued is used to select a probability curve
  - The frame may be selected for random discard

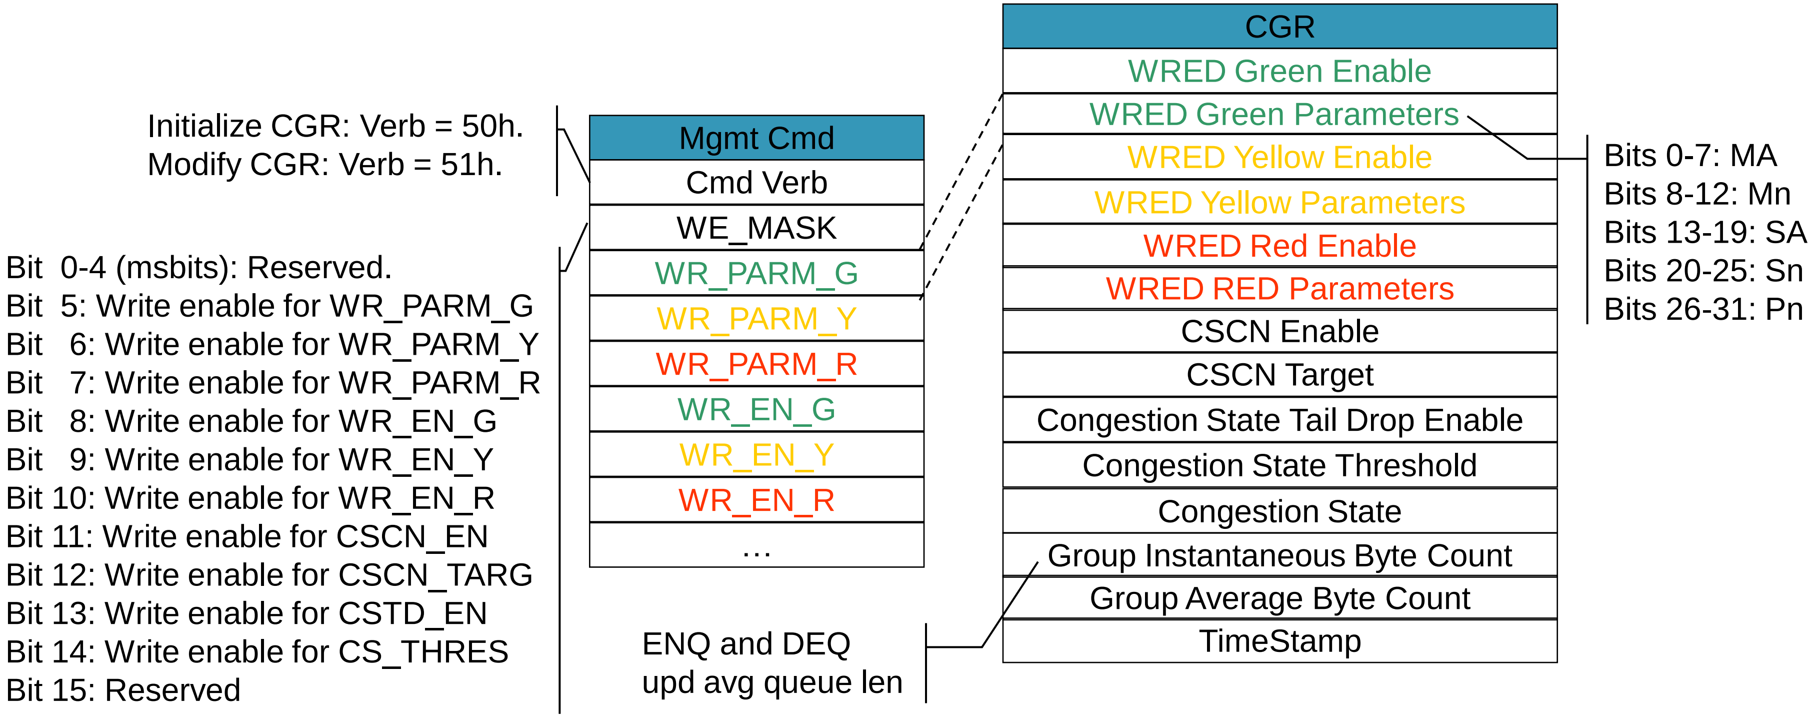

Initialize CGR: Verb = 50h.
Modify CGR: Verb = 51h.

Bit 0-4 (msbits): Reserved.
Bit 5: Write enable for WR_PARM_G
Bit 6: Write enable for WR_PARM_Y
Bit 7: Write enable for WR_PARM_R
Bit 8: Write enable for WR_EN_G
Bit 9: Write enable for WR_EN_Y
Bit 10: Write enable for WR_EN_R
Bit 11: Write enable for CSCN_EN
Bit 12: Write enable for CSCN_TARG
Bit 13: Write enable for CSTD_EN
Bit 14: Write enable for CS_THRES
Bit 15: Reserved

Mgmt Cmd
Cmd Verb
WE_MASK
WR_PARM_G
WR_PARM_Y
WR_PARM_R
WR_EN_G
WR_EN_Y
WR_EN_R
…

ENQ and DEQ
upd avg queue len

CGR
WRED Green Enable
WRED Green Parameters
WRED Yellow Enable
WRED Yellow Parameters
WRED Red Enable
WRED RED Parameters
CSCN Enable
CSCN Target
Congestion State Tail Drop Enable
Congestion State Threshold
Congestion State
Group Instantaneous Byte Count
Group Average Byte Count
TimeStamp

Bits 0-7: MA
Bits 8-12: Mn
Bits 13-19: SA
Bits 20-25: Sn
Bits 26-31: Pn

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Security (SEC) 5.0 Overview

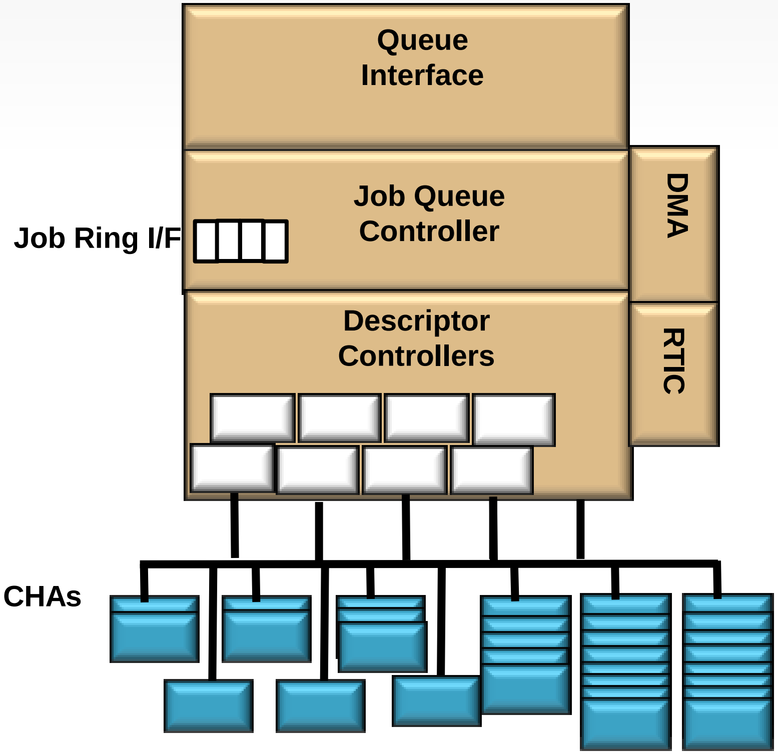

Supports protocol processing for the following:

- IPSec
- 802.1ae (MACSEC)
- SSL/TLS/DTLS
- 3GPP RLC
- LTE PDCP
- SRTP
- 802.11i (Wi-Fi)
- 802.16e (WiMax)

Public Key Hardware Accelerators (PKHA) ~25K RSA Ops/sec (1024b)
RSA and Diffie-Hellman (to 4096b)
Elliptic curve cryptography (1023b)
Data Encryption Standard Accelerators (DESA) ~15Gbps
DES, 3DES (2K, 3K)
ECB, CBC, OFB modes
Advanced Encryption Standard Accelerators (AESA) ~40Gbps
Key lengths of 128-, 192-, and 256-bit
ECB, CBC, CTR, CCM, GCM, CMAC, OFB, CFB, and XTS

ARC Four Hardware Accelerators (AFHA) ~7.5Gbps

Compatible with RC4 algorithm
Message Digest Hardware Accelerators (MDHA) ~40Gbps
SHA-1, SHA-2 256,384,512-bit digests
MD5 128-bit digest
HMAC with all algorithms
Kasumi/F8 Hardware Accelerators (KFHA) ~9Gbps
F8 , F9 as required for 3GPP
A5/3 for GSM and EDGE
GEA-3 for GPRS
Snow 3G Hardware Accelerators (STHA) ~12Gbps
Implements Snow 3.0
ZUC Hardware Accelerators (ZHA) ~14Gbps
Implements 128-EEA3 & 128-EIA3
CRC Unit~40Gbps
Standard and user defined polynomials
Random Number Generator, random IV generation

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Pattern Matching Engine (PME) 2.X Overview

- Regex support plus significant extensions:
  - Patterns can be split into 256 sets each of which can contain 16 subsets
  - 32K patterns of up to 128B length
  - 9.6 Gbps raw performance
- Combined hash/NFA technology
  - No “explosion” in number of patterns due to wildcards
  - Low system memory utilization
  - Fast pattern database compiles and incremental updates
- Matching across “work units”
  - Finds patterns in streamed data
- Pipeline of processing
  - PME offers pipeline of filtering, matching, and behavior base engine for complete pattern matching solution

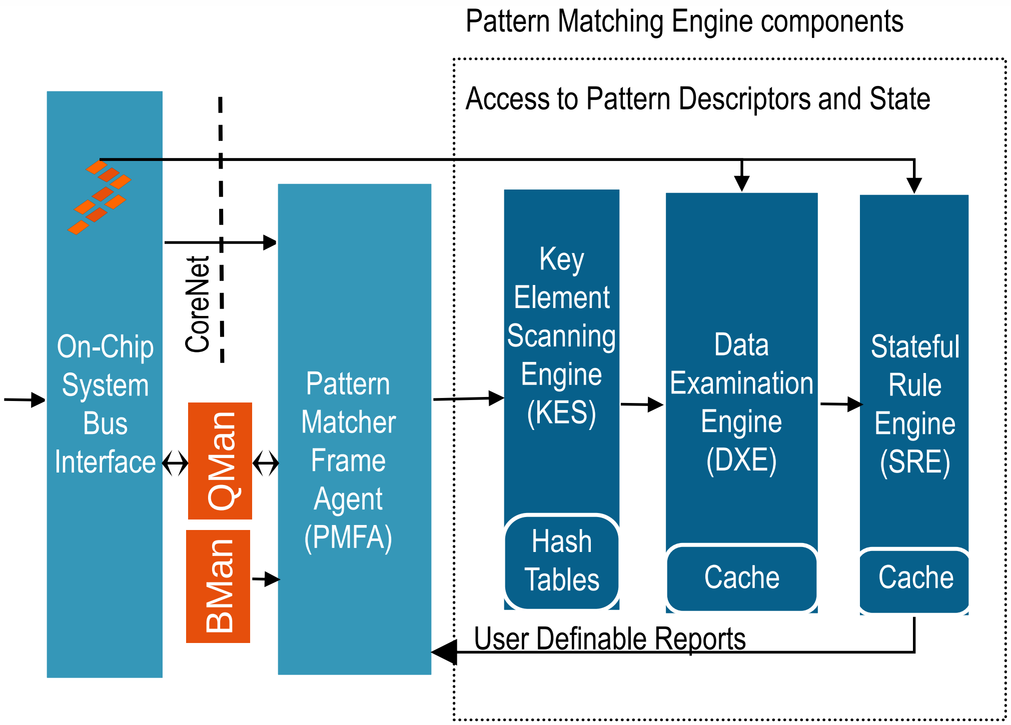

Pattern Matching Engine components
Access to Pattern Descriptors and State
On-Chip System Bus Interface
CoreNet
QMan
BMan
Pattern Matcher Frame Agent (PMFA)
Key Element Scanning Engine (KES)
Hash Tables
Data Examination Engine (DXE)
Cache
Stateful Rule Engine (SRE)
Cache
User Definable Reports

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

RapidIO Message Manager (RMan) Overview

- Many queues allow multiple inbound/outbound queues per core
  - Hardware queue management via QorIQ Data Path Architecture (DPAA)
- Supports all messaging-style transaction types
  - Type 11 Messaging
  - Type 10 Doorbells
  - Type 9 Data Streaming
- Enables low overhead direct core-to-core communication

Device-to-Device Transport

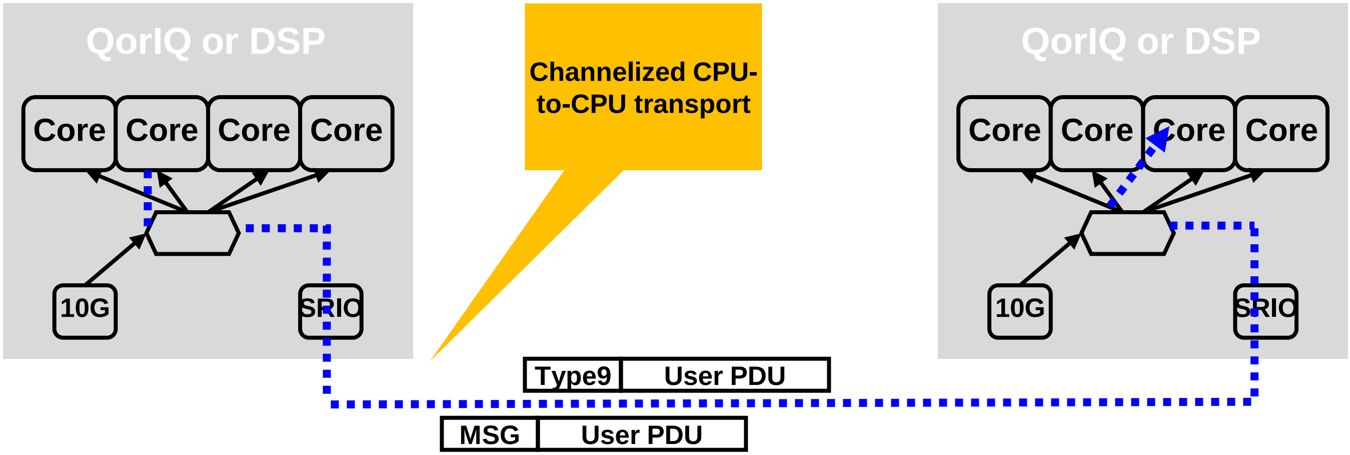

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Decompression/Compression Engine (DCE) Overview

- Deflate
  - As specified as in RFC1951
- GZIP
  - As specified in RFC1952
- Zlib
  - As specified in RFC1950
  - Interoperable with the zlib 1.2.5 compression library
- Encoding
  - supports Base 64 encoding and decoding (RFC4648).
- Operate up to 600Mhz
  - 10Gbps Compress
  - 10Gbps Decompress
  - 20Gbps Aggregate

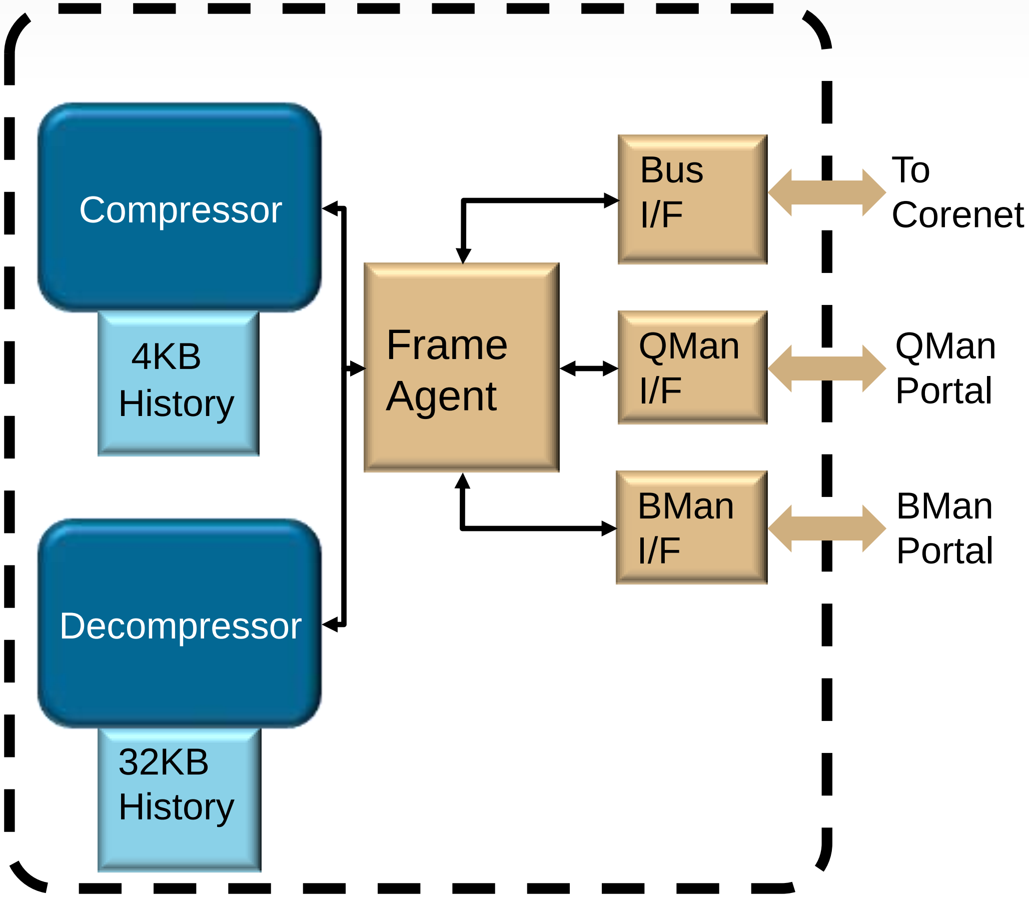

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, ColdFire+, C-Ware, the Energy Efficient Solutions logo, Kinetis, mobileGT, PEG, PowerQUICC, Processor Expert, QorIQ, Qorivva, SafeAssure, the SafeAssure logo, StarCore, Symphony and VortiQa are trademarks of Freescale Semiconductor, 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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Data Center Bridging (DCB) Overview

- QMan 1.2 (e.g. QorIQ T42xx) supports Data Center Bridging (DCB).
- DCB refers to a series of inter-related IEEE specifications collectively designed to enhance Ethernet LAN traffic prioritization and congestion management.
- DCB can be used in:
  - Between data center network nodes:
  - LAN/network traffic
  - Storage Area Network (SAN) (e.g. Fiber Channel (loss sensitive)) and,
  - IPC traffic (e.g. Infiniband (low latency))
- The DPAA is compliant with the following DCB specifications (traffic management related) :
  - IEEE Std. 802.1Qbb: Priority-based flow control (PFC)
    - To avoid frame loss, a PFC Pause frames can be sent autonomously by HW.
  - IEEE Std. 802.1Qaz: Enhanced transmission selection (ETS)
    - Support Weighted bandwidth fairness.
  - IEEE 802.1Qau: Quantized Congestion Notification (QCN)
    - end-to end congestion control mechanism.

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Agenda

- DPAA Overview
- BMan Enablement
- FMan Enablement
- QMan Enablement
  - QMan Building Blocks
  - QMan Functions
  - QMan Scheduling
  - Order Restoration, Order Preservation and Atomicity
  - Congestion Management and Avoidance
- Accelerator Overview
- Conclusion

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

Conclusion

- DPAA is about seamless Integration of multiple CPUs and Accelerators for high speed Interfaces.
- The Data Path Acceleration Architecture **Accelerators** included:
  - Queue Manager
  - Buffer Manager
  - Frame Manager
  - Hardware accelerators such as SEC, PME, DCE, and RMan
  - Power Architecture Cores
- Seamless **Integration** of these components address multicore requirements:
  - Load spreading
  - Packet ordering
  - Device virtualization
  - Inter-core communication
  - HW buffer management

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, Vybrid and Xtrinsic are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2013 Freescale Semiconductor, Inc.

freescale™