June, 2010

Leveraging the QorIQ Data Path Acceleration Architecture (DPAA) for Wireless Applications

FTF-NET-F0704

Stephen Cole – Senior Systems Architect

Jonas Svennebring – Senior Software FAE

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale™ semiconductor

Agenda (time/presenter to be removed)

- Challenges in Wireless Industry
- QorIQ Application Mapping
- Data Path Acceleration – quick recap.
- Packet I/O
- PME – traffic monitoring and management
- SEC – Ciphering and packet formatting
- Use-case: Voice packet processing
- Use-case: HSPA, mobile broadband

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

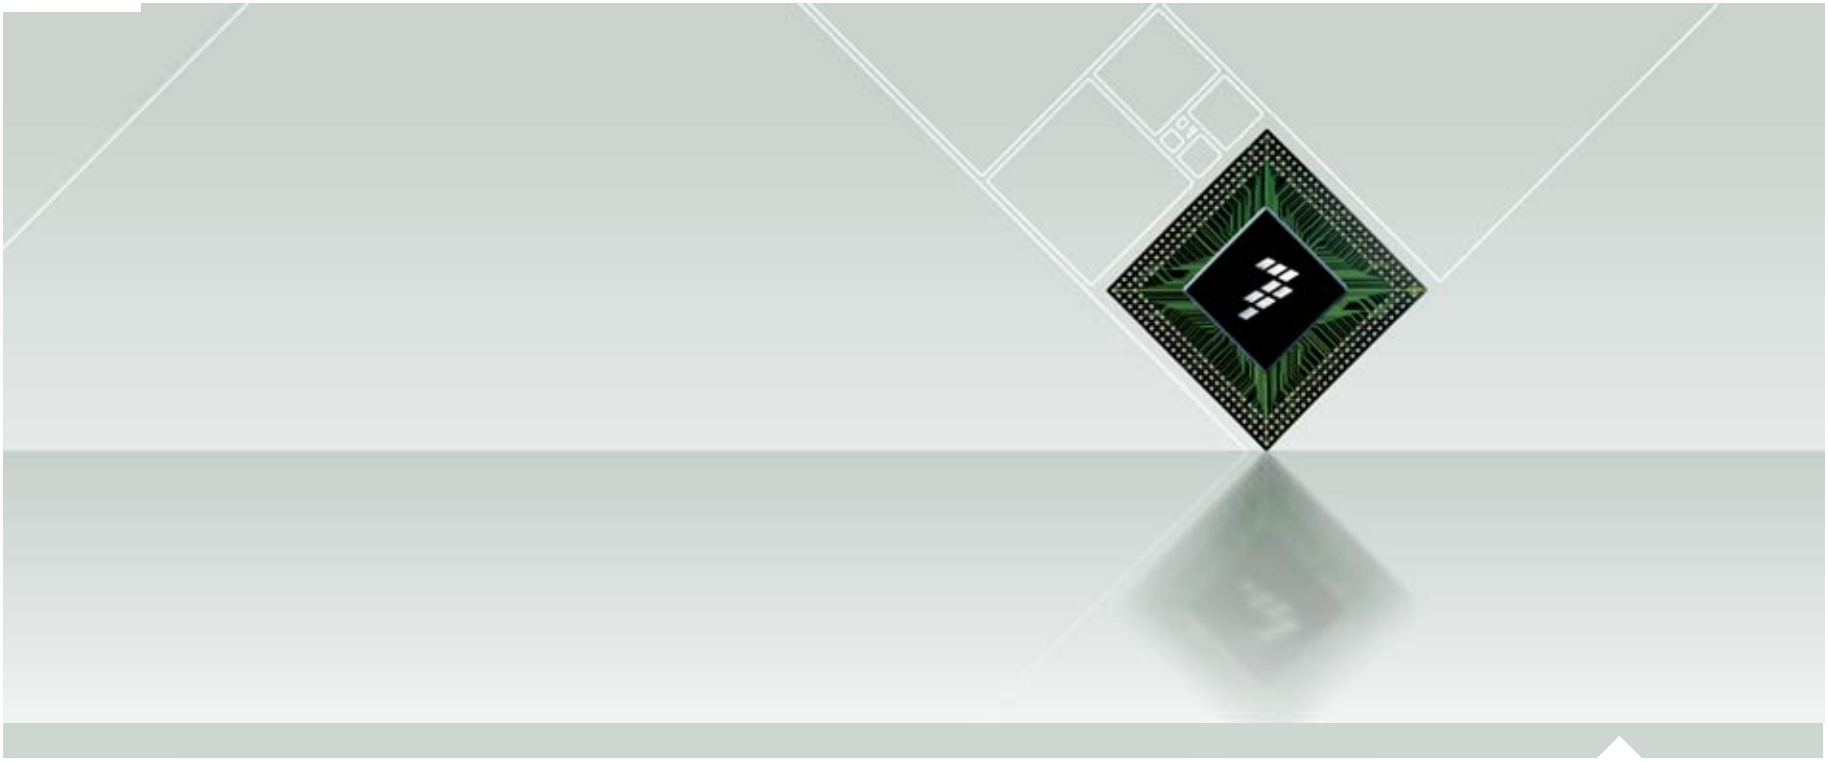

Challenges in Wireless Industry

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

Trends / Challenges

- Mobile broadband has taken off
  - High bandwidths
  - High user peak rates
  - Low network latency

- Smart Phones (iphone etc) require intensive signaling, i.e. control processing.

- WCDMA follows the LTE track in development
  - Competitive performance with HSPA+
  - Support for flat architecture

- Home base-station, will 2010/11 be the break through?
- All IP? Investments in ATM needs smooth migration, i.e. IP over ATM.
- Content Caches in field to offload main pipe to Internet?

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

Challenge: Mobile Broadband

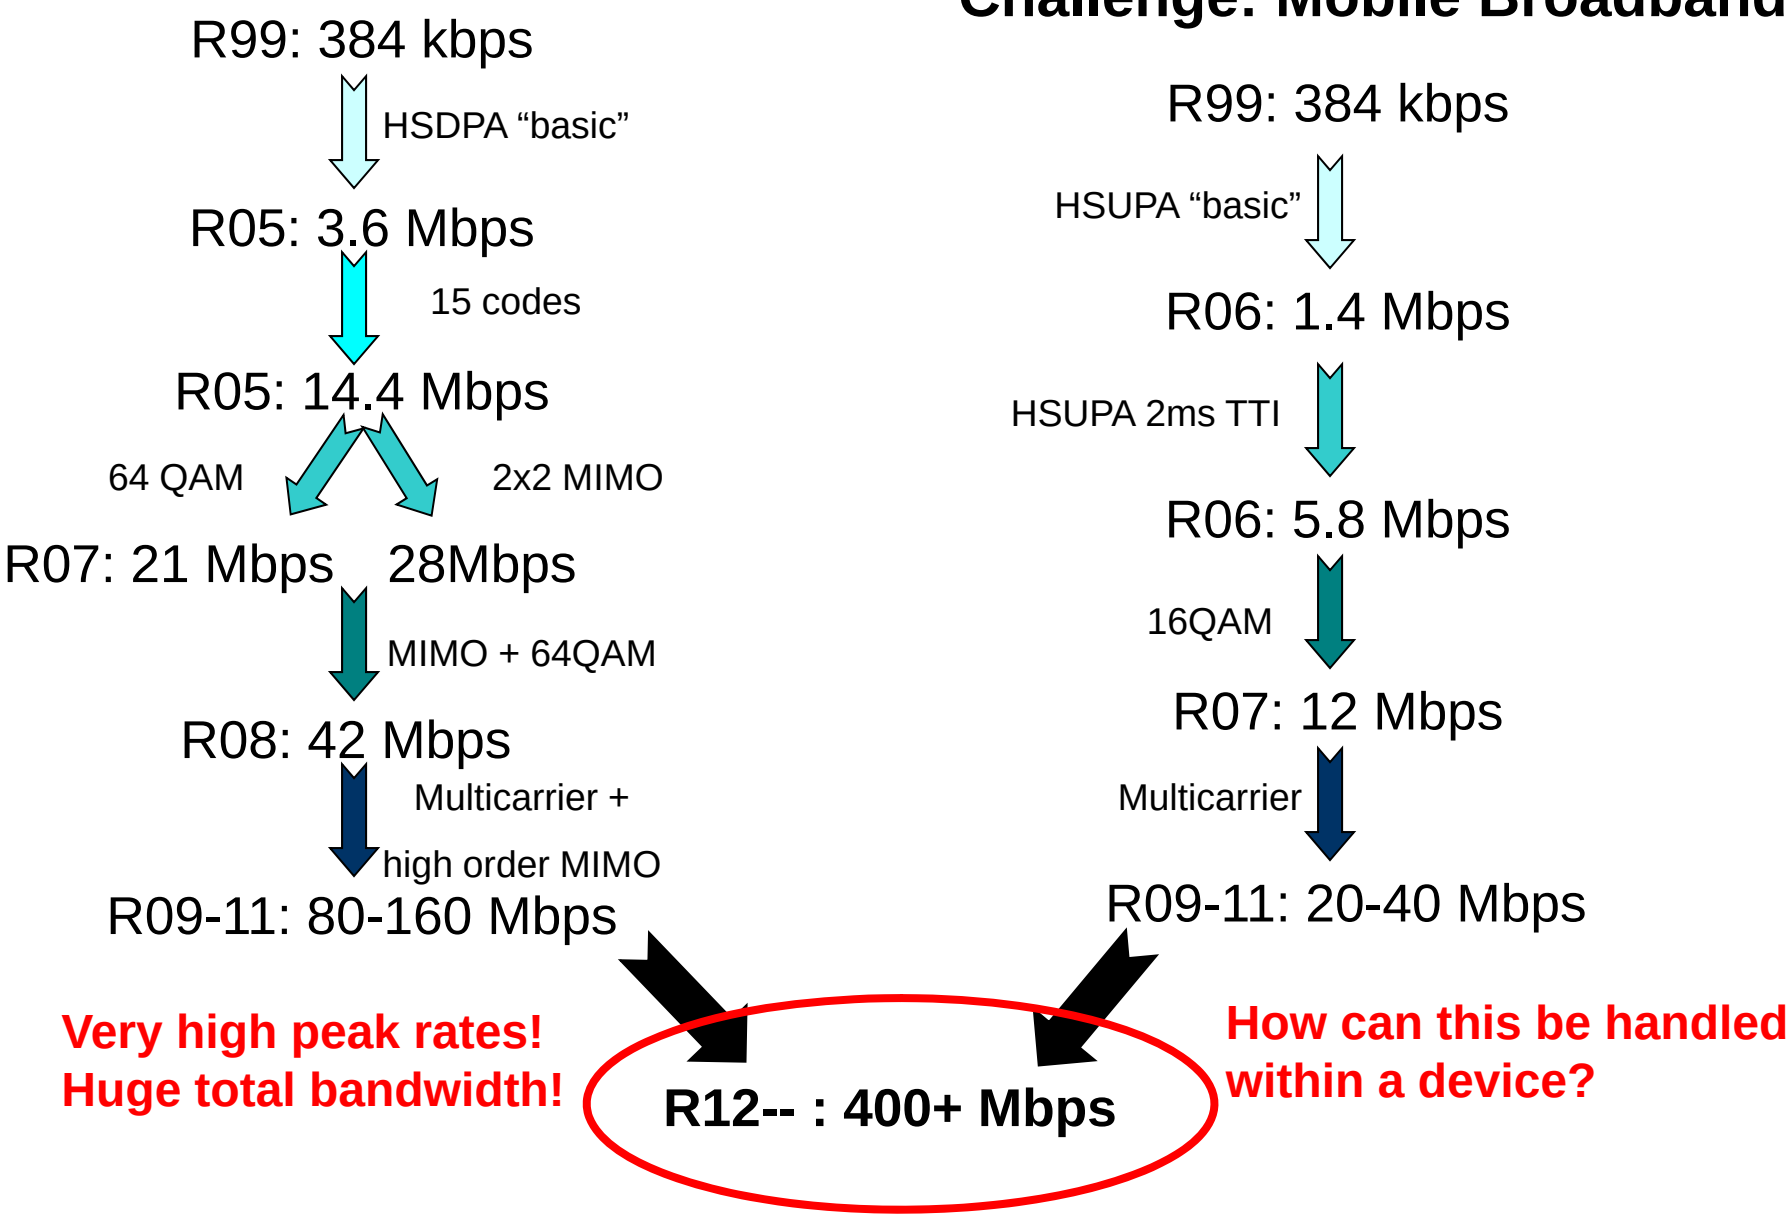

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

WCDMA and LTE Core Networks

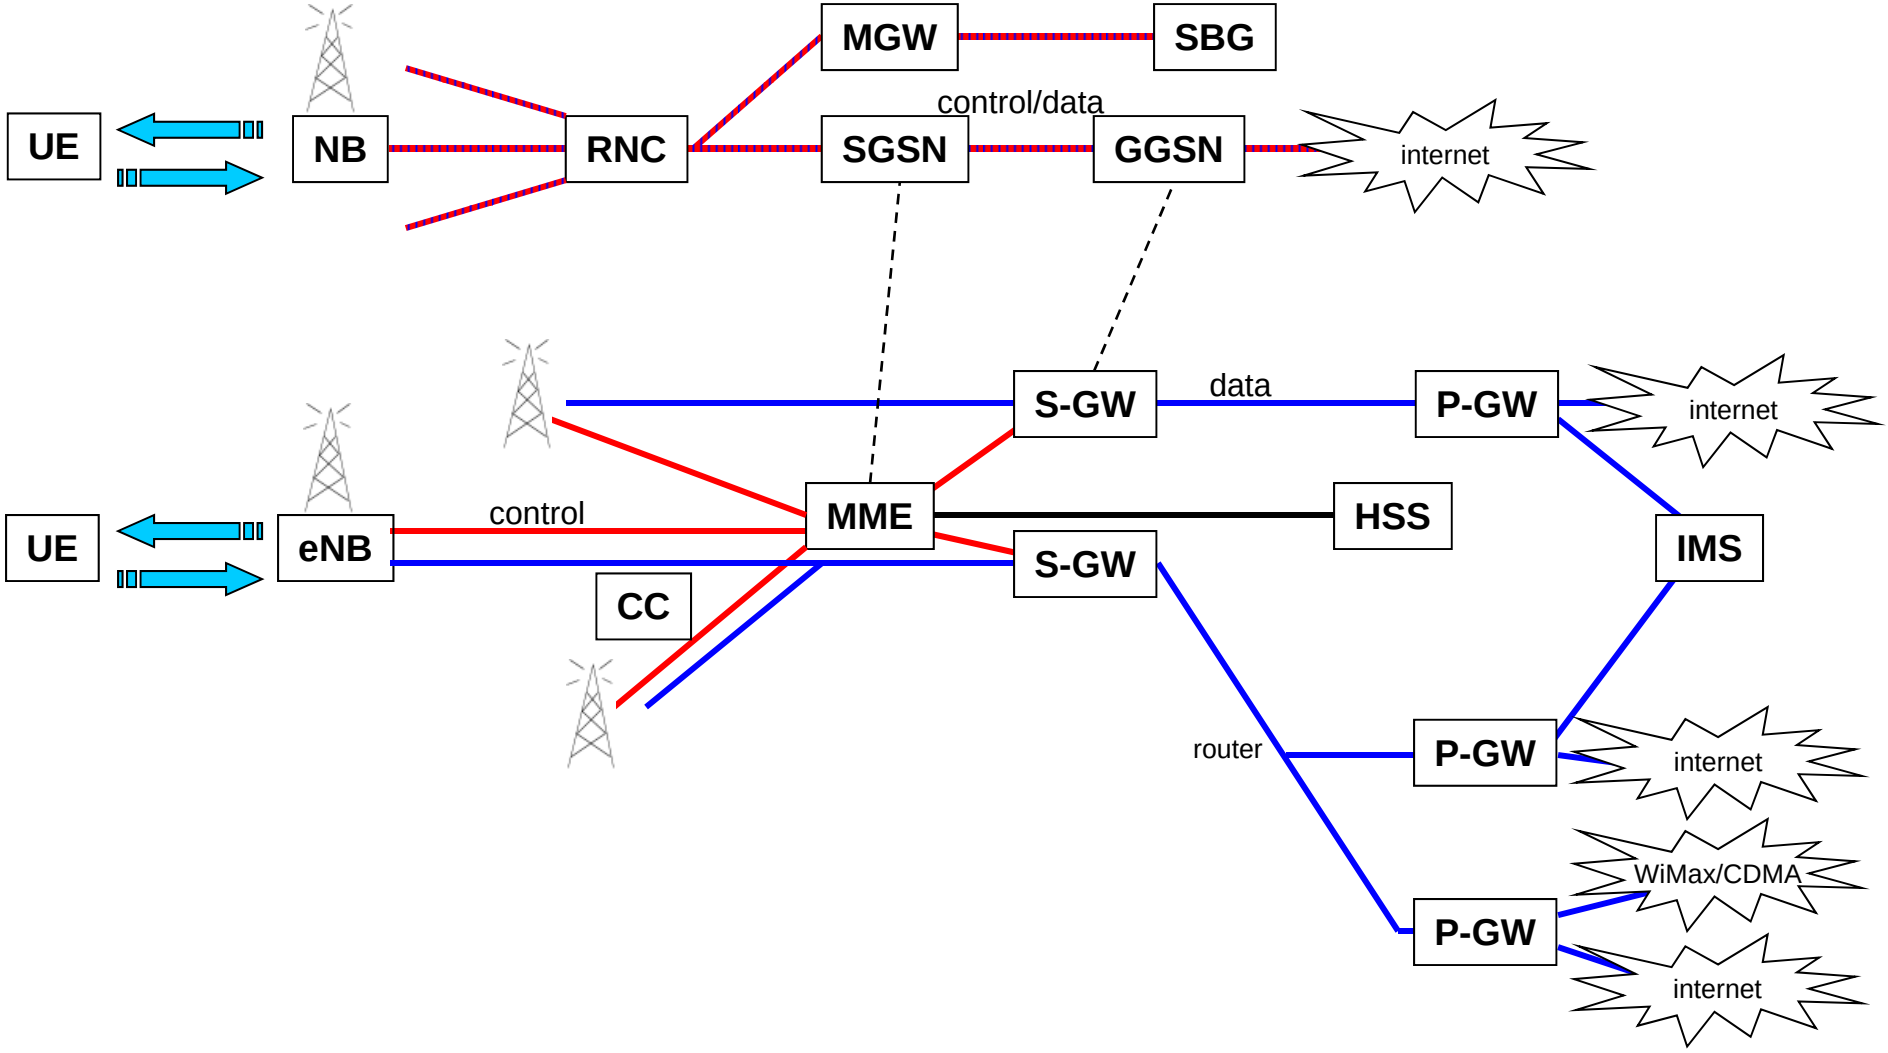

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

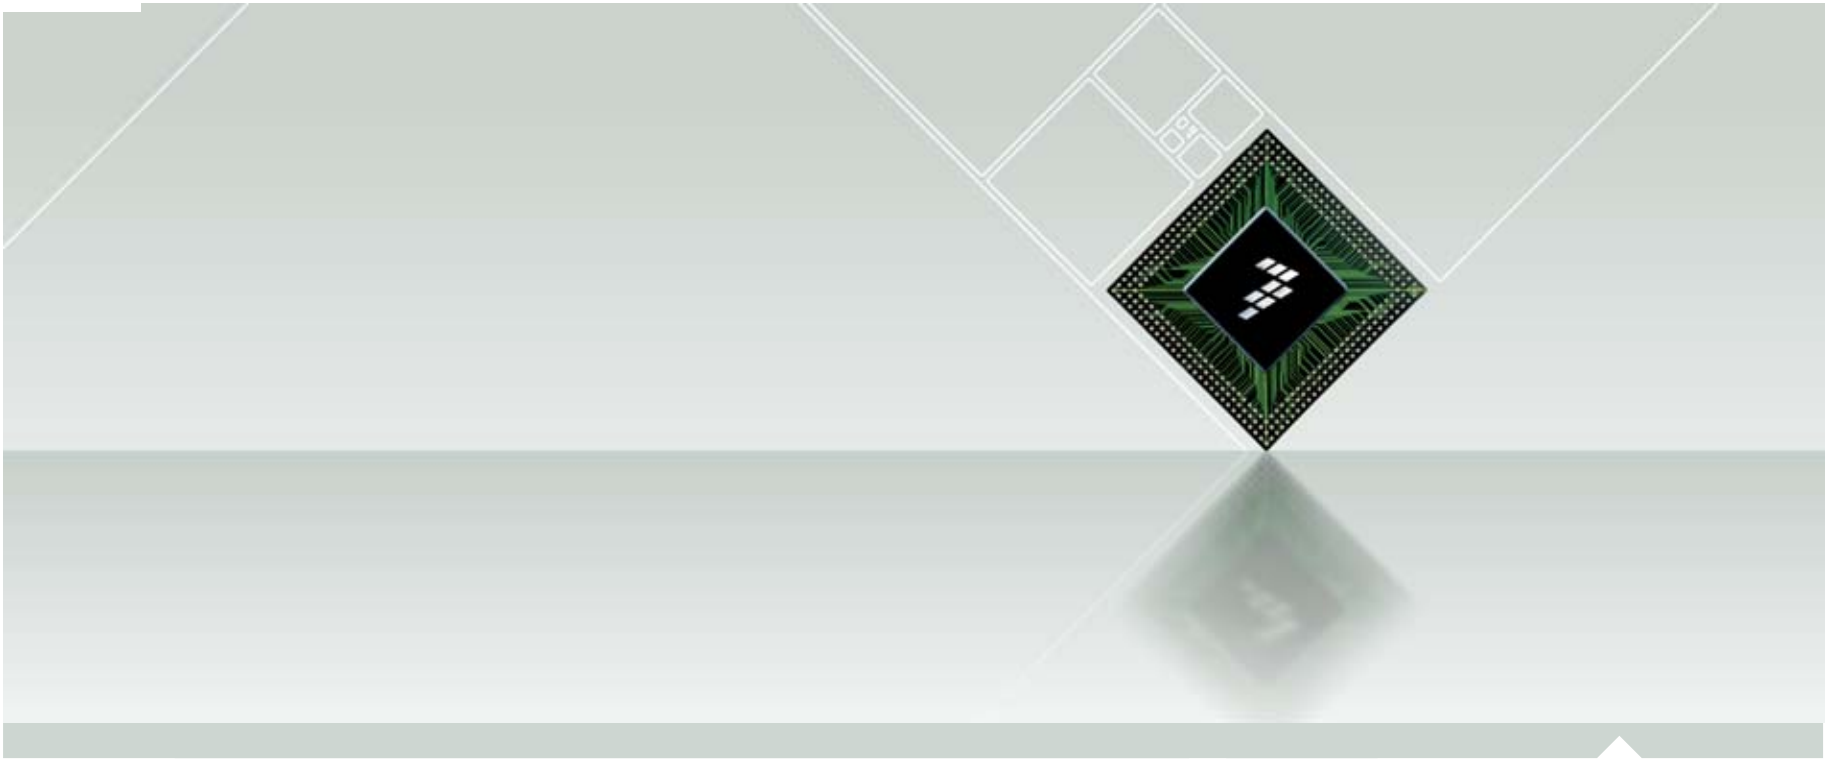

QorIQ Application Mapping

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

QorIQ™ Investment Continues Our Embedded Leadership Tradition

- **Our 15+ year Heritage:**
- **3rd Generation Data Path**
  - Gen-1: CPM MPC8260
  - Gen-2: QUICC Engine™ MPC8360
  - Gen-3: DPAA QorIQ P4080
- **Accelerating Connectivity**
  - eTSEC
  - SEC 4.0
  - PME 2.0
  - PCIe, Serial RapidIO®, XAUI
- **Power Architecture™ ISA**
  - e500 PowerQUICC® III
  - e500 QorIQ P1, P2 platforms
  - e500mc QorIQ P3, P4 platforms
- **Long life cycle management for industry**

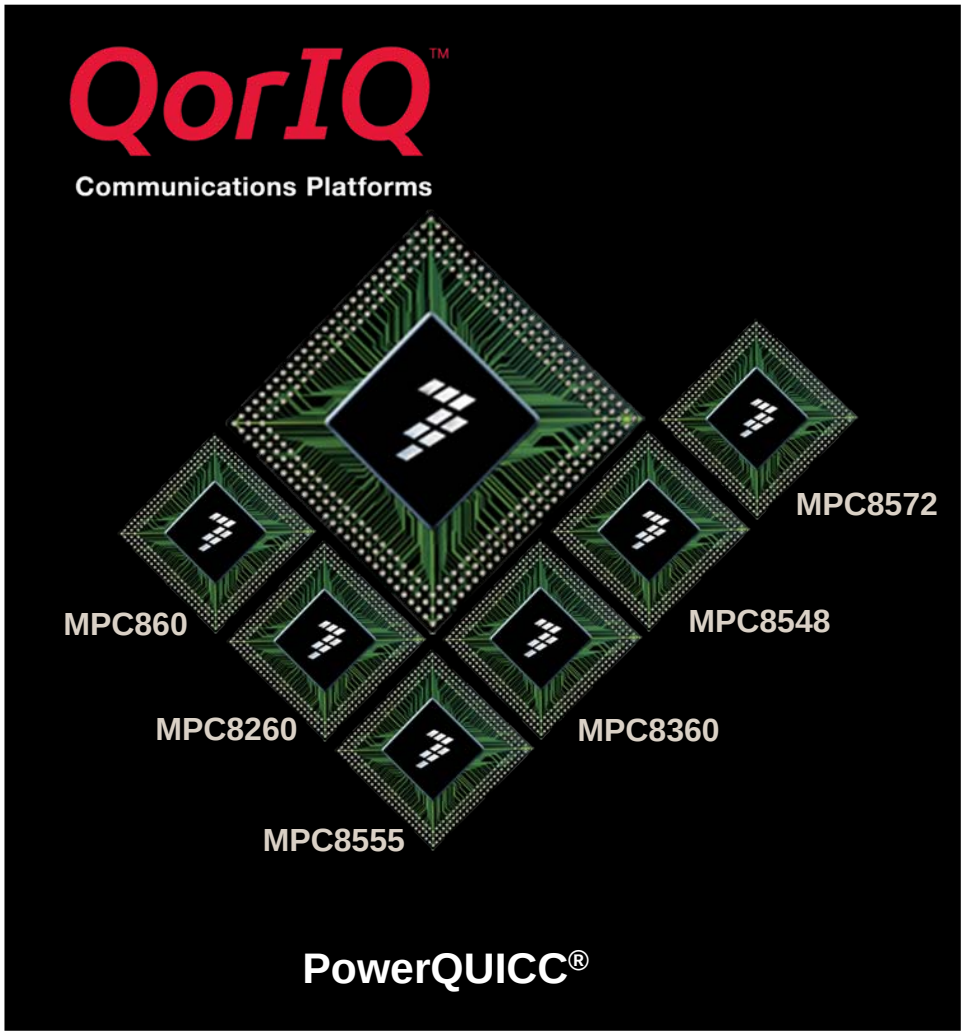

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale™ semiconductor

QorIQ™ Platform Levels matching Wireless Applications

45nm PLATFORMS / PRODUCTS

QorIQ P5
Highest-performing embedded processors
PRODUCTS: P5010 / P5020

QorIQ P4
Octal and Quad core family
PRODUCTS: P4080
P4040

QorIQ P3
Quad core family – entry level MC
PRODUCTS: P3040 / P3041

QorIQ P2
High performance Dual and Single Core
PRODUCTS:
P2020 and P2010

QorIQ P1
Low Power Dual and Single Core
PRODUCTS: P1020 , P1011, P1010


MME


SGSN


IS Gateway


Radio Network Control


Access Gateway


Content Caches


Basestation


Network Interface And Line cards

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor ™

QorIQ P4 Series P4080 Block Diagram

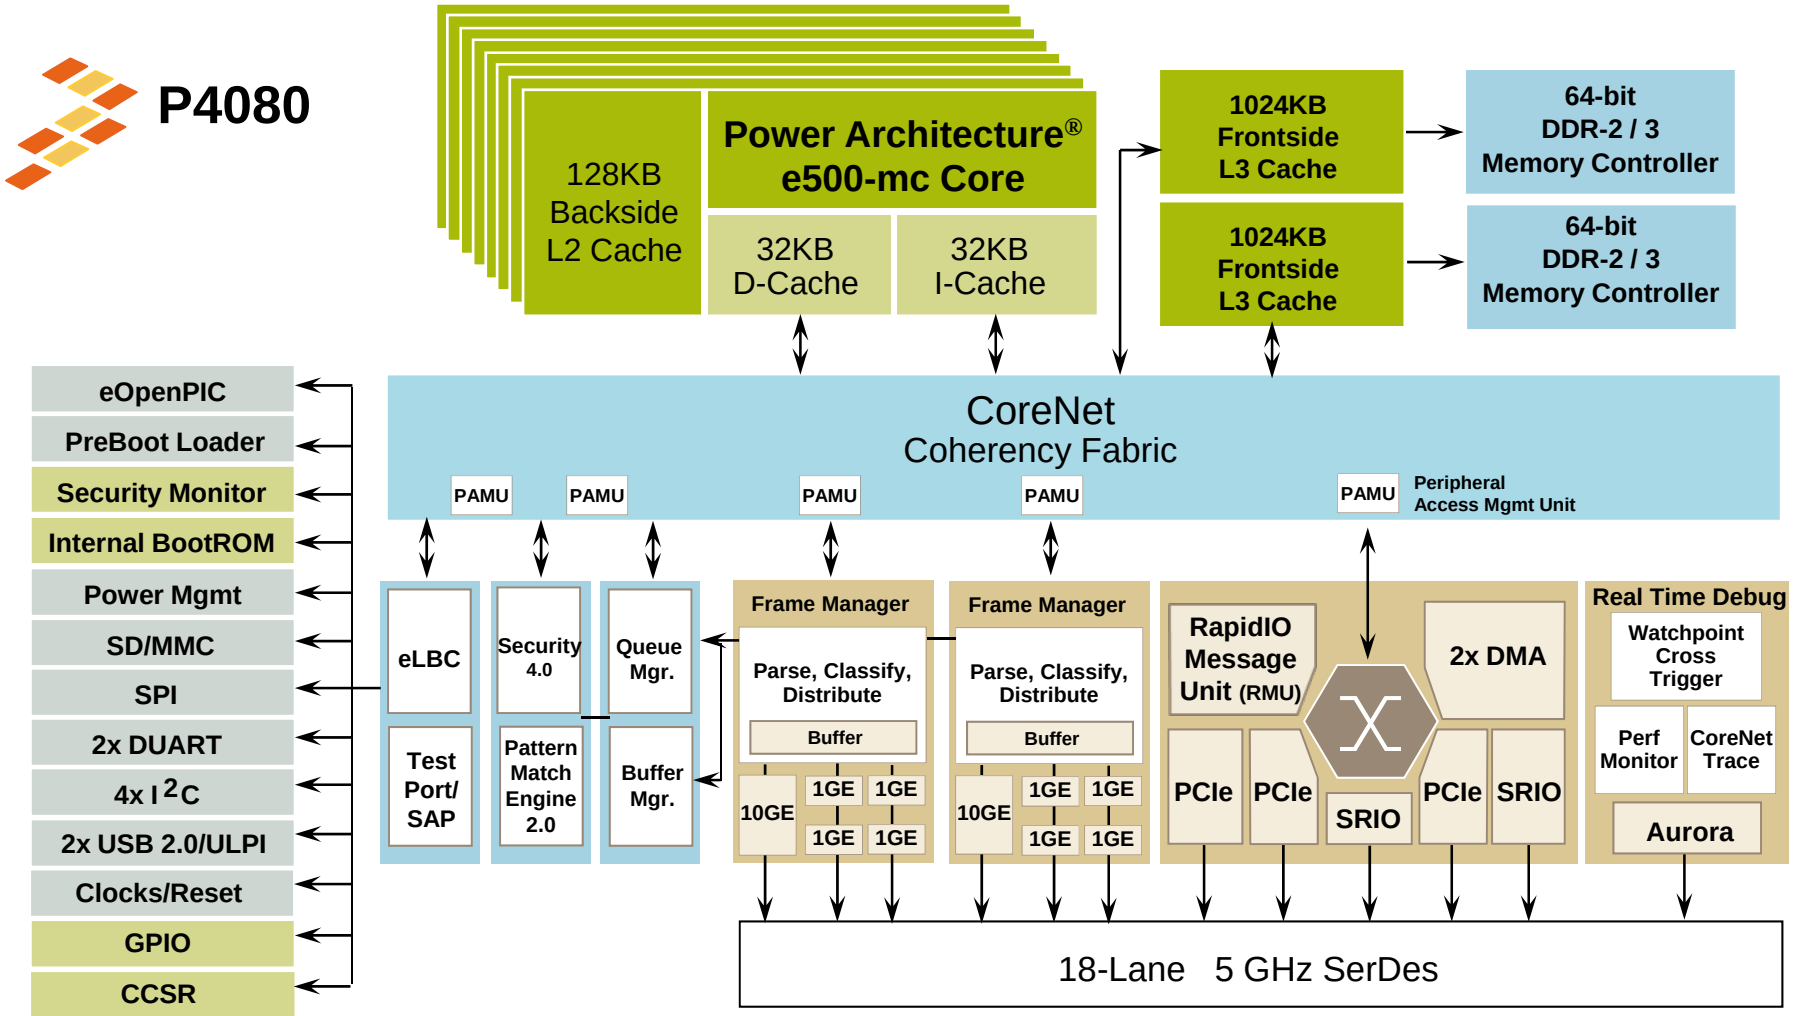

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

QorIQ P5 Series P5020 Block Diagram

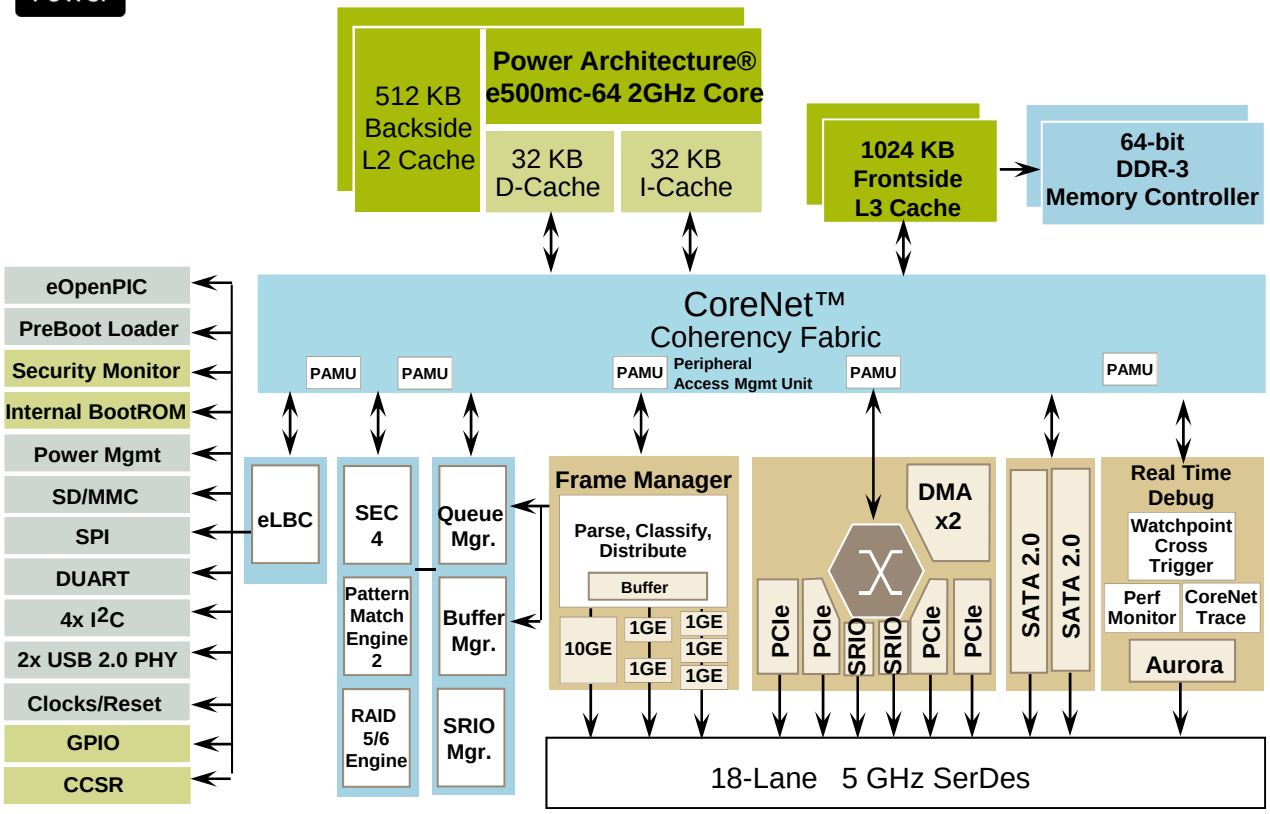

- **Pin Compatible to P4080, P4040 & P3041**
- **45nm SOI Process**

- **2x e500mc-64 core, built on Power Architecture® technology**
  - 2x 64-bit cores (up to 2+ GHz) with 512 KB backside L2 cache
  - Dual 1MB Shared L3 cache w/ECC
- **Memory Controller**
  - Dual 32/64-bit DDR3/3L w/ECC up to 1.3 GHz
- **Ethernet**
  - 5 x 10/100/1000 Ethernet controllers
  - 1 x 10GE controller (XAUI)
- **High Speed Interconnect**
  - 4 PCI Express® 2.0 controllers
  - 2 Serial RapidIO® 1.3 + 2.0 controllers
  - 2 SATA 3Gb/s
  - 2 USB 2.0 with PHY
- **CoreNet Switch Fabric**
- **Trusted Architecture**
- **Data Path Acceleration Architecture**
  - Security Engine (SEC)
  - Pattern Matching Engine (PME)
  - RAID 5/6 Engine
  - Enhanced RapidIO Messaging (Rman)

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

QorIQ P3 Series P3041 Block Diagram

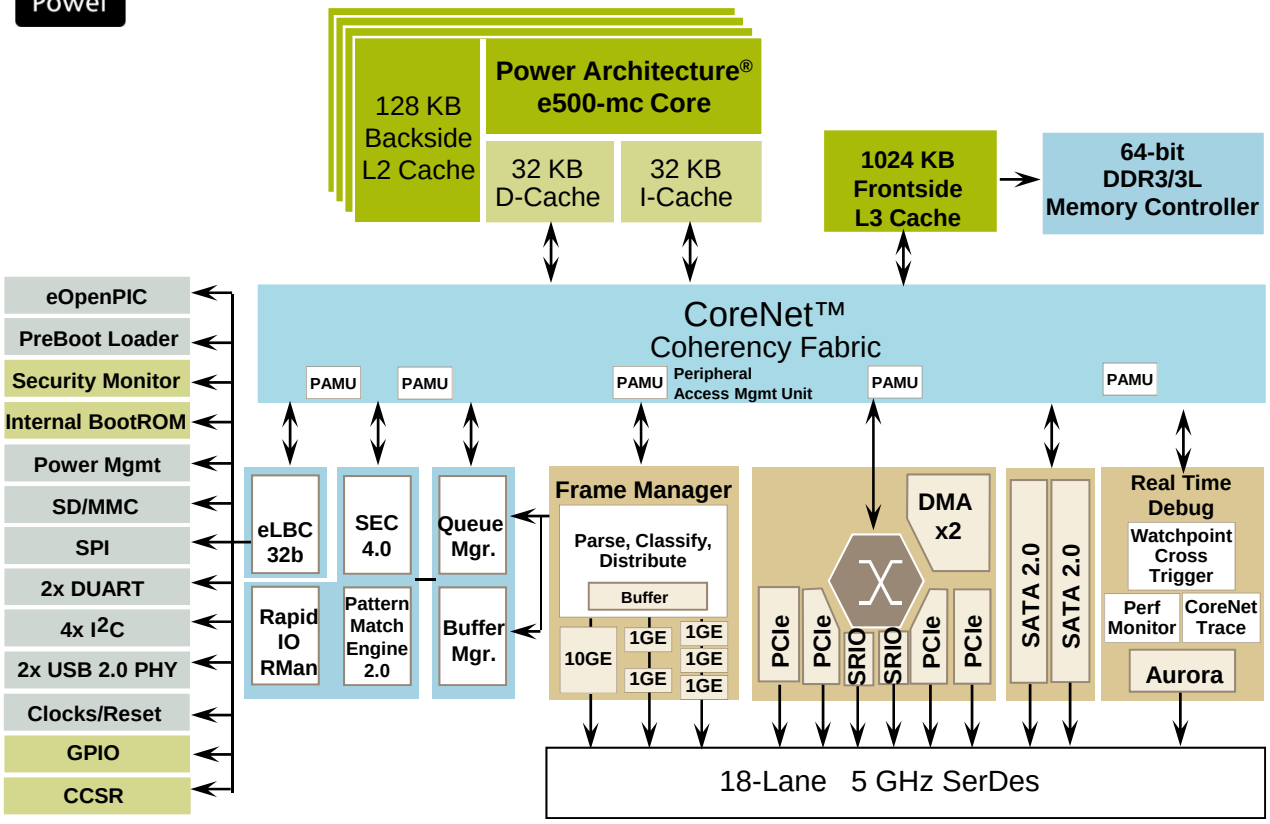

- **Pin Compatible to P4080, P4040, P5020 & P5010**
- **45nm SOI Process**

Quad e500mc Power Architecture®
- 4 cores (up to 1.5GHz) with 128KB backside L2 cache
- 1MB Shared L3 Cache w/ECC

Memory Controller
- 32/64bit DDR3/3L w/ECC up to 1.3 GHz

Ethernet
- 5 x 10/100/1000 Ethernet Controllers
- 1 x 10GE Controllers

High Speed Interconnect
- 4 PCIe 2.0 Controllers
- 2 SRIO 1.3 + 2.0 Controllers
- 2 SATA 2.0
- 2 USB 2.0 w/PHY

CoreNet Switch Fabric

Trusted Architecture

Datapath Acceleration Architecture
- Security Engine (SEC)
- Pattern Matching Engine (PME)
- Enhanced RapidIO Messaging (Rman)

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

MPC8569 PowerQUICC III
Bridging the Gap to the All-IP Network

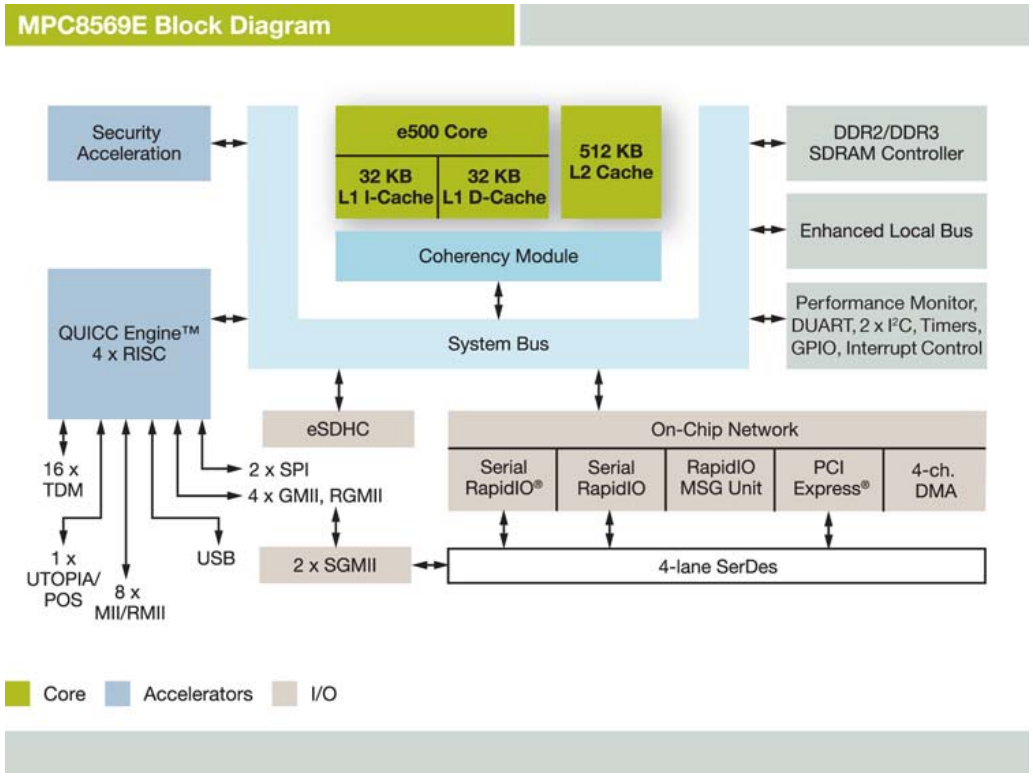

- **e500 PowerPC** from 800MHz to 1.33 GHz
  - 512KB L2 Cache w/ ECC
  - 36bit physical addressing
  - Double Precision Floating Point
- **System Interfaces**
  - 64b or 2x32-bit DDR2/3 w/ ECC
  - 800 Mbps/pin data rate
  - 16-bit Local Bus for SRAM/Flash
  - Timers, DUART, 2xI²C, GPIO, SPI
  - USB 2.0 Full Speed
- **High Speed Serial Interfaces**
  - Dual SGMII
  - Dual x1 Serial RapidIO or PCI-Ex
- **QUICC Engine**
  - 4 RISCs up to 667 MHz
  - Maximum of 8 Ethernet interfaces, one per UCC:
    - 4 x Gigabit Eth (up to 2 w/SGMII)
    - Up to 8 x 10/100 Ethernet
  - Multi-PHY UTOPIA/POS-PHY L2 (16-bit)
  - IEEE1588 Support v2
  - 16 x T1/E1 (512 x 64kbps channels)
- **Security Engine (SEC3.0)**
  - ARC4, 3DES, AES, RSA/ECC, RNG, XOR, Single pass SSL/TLS, Kasumi, SNOW
- **Four-channel DMA**
- **45nm SOI process technology**
- **Target <7W Power (@ 800MHz e500)**

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

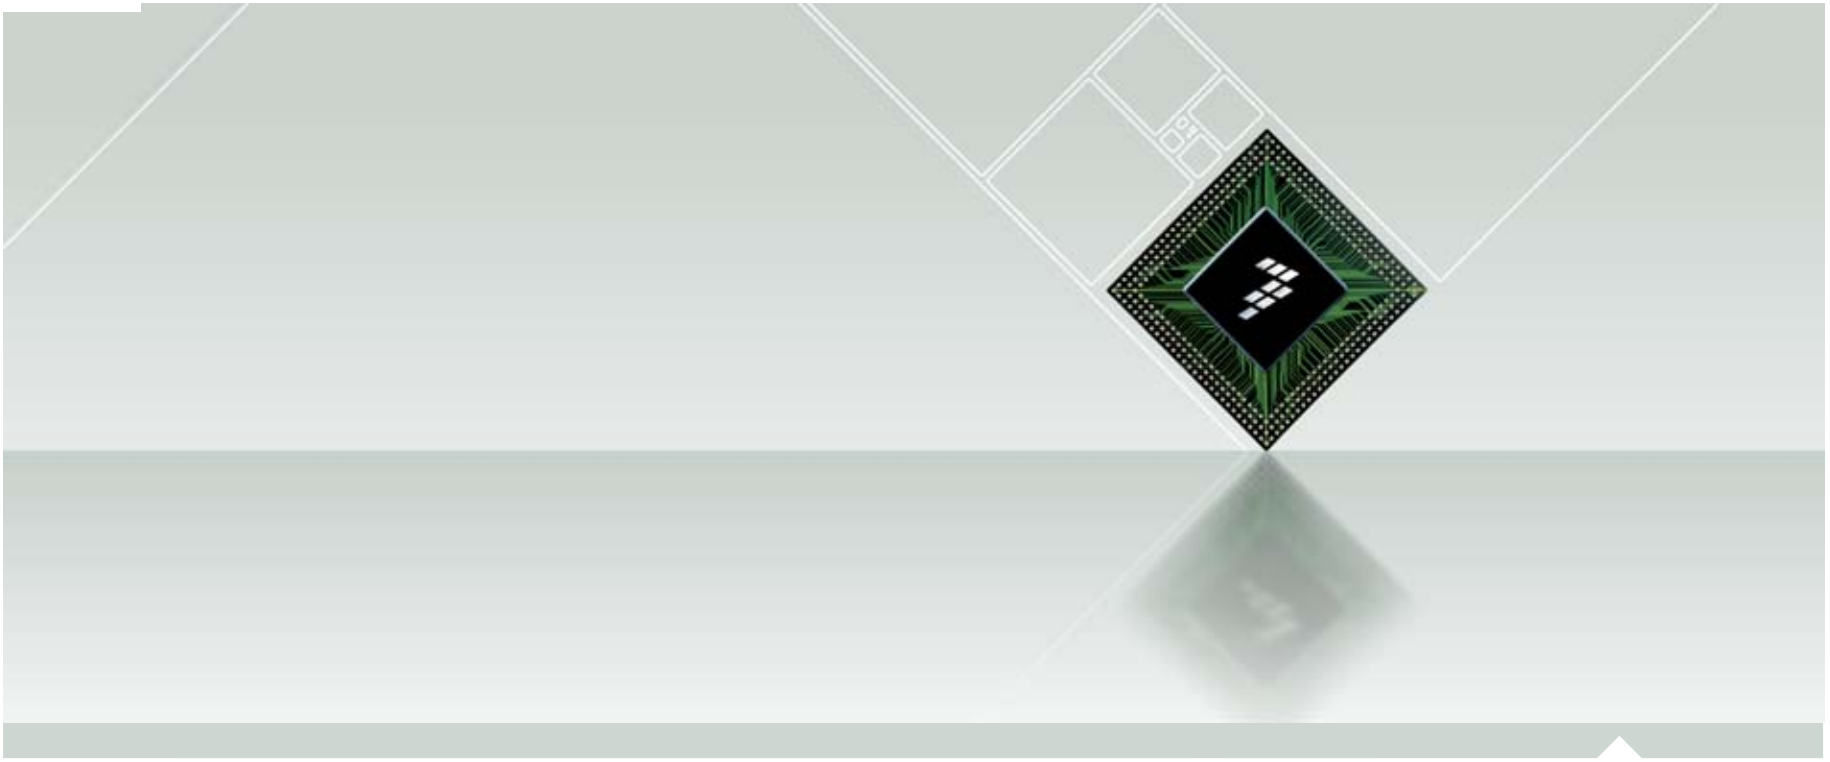

Data Path Acceleration: Quick Recap

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

DPAA – Data Path Acceleration Architecture

- **Frame Manager:**
- Parse header information
- Classify to what destination the packet is targeted
  ex. MAC, IP, UDP destination
- Distribute packets to allow load balancing between cores
- Police flows to avoid DoS attacks

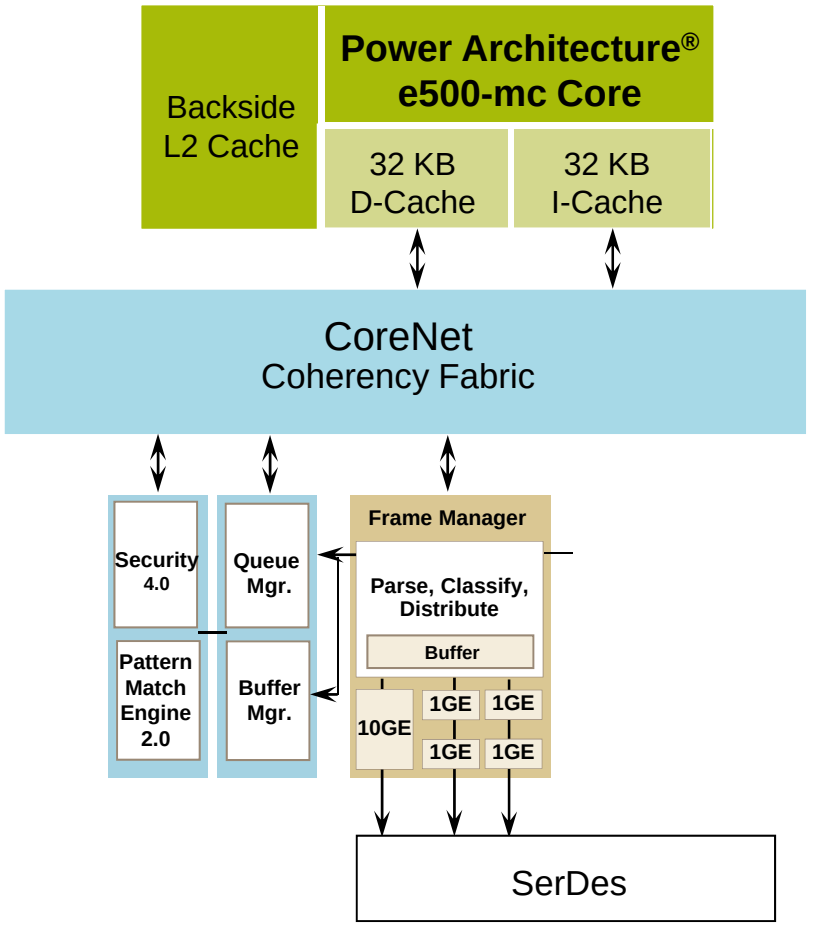

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

QMan Software Portals

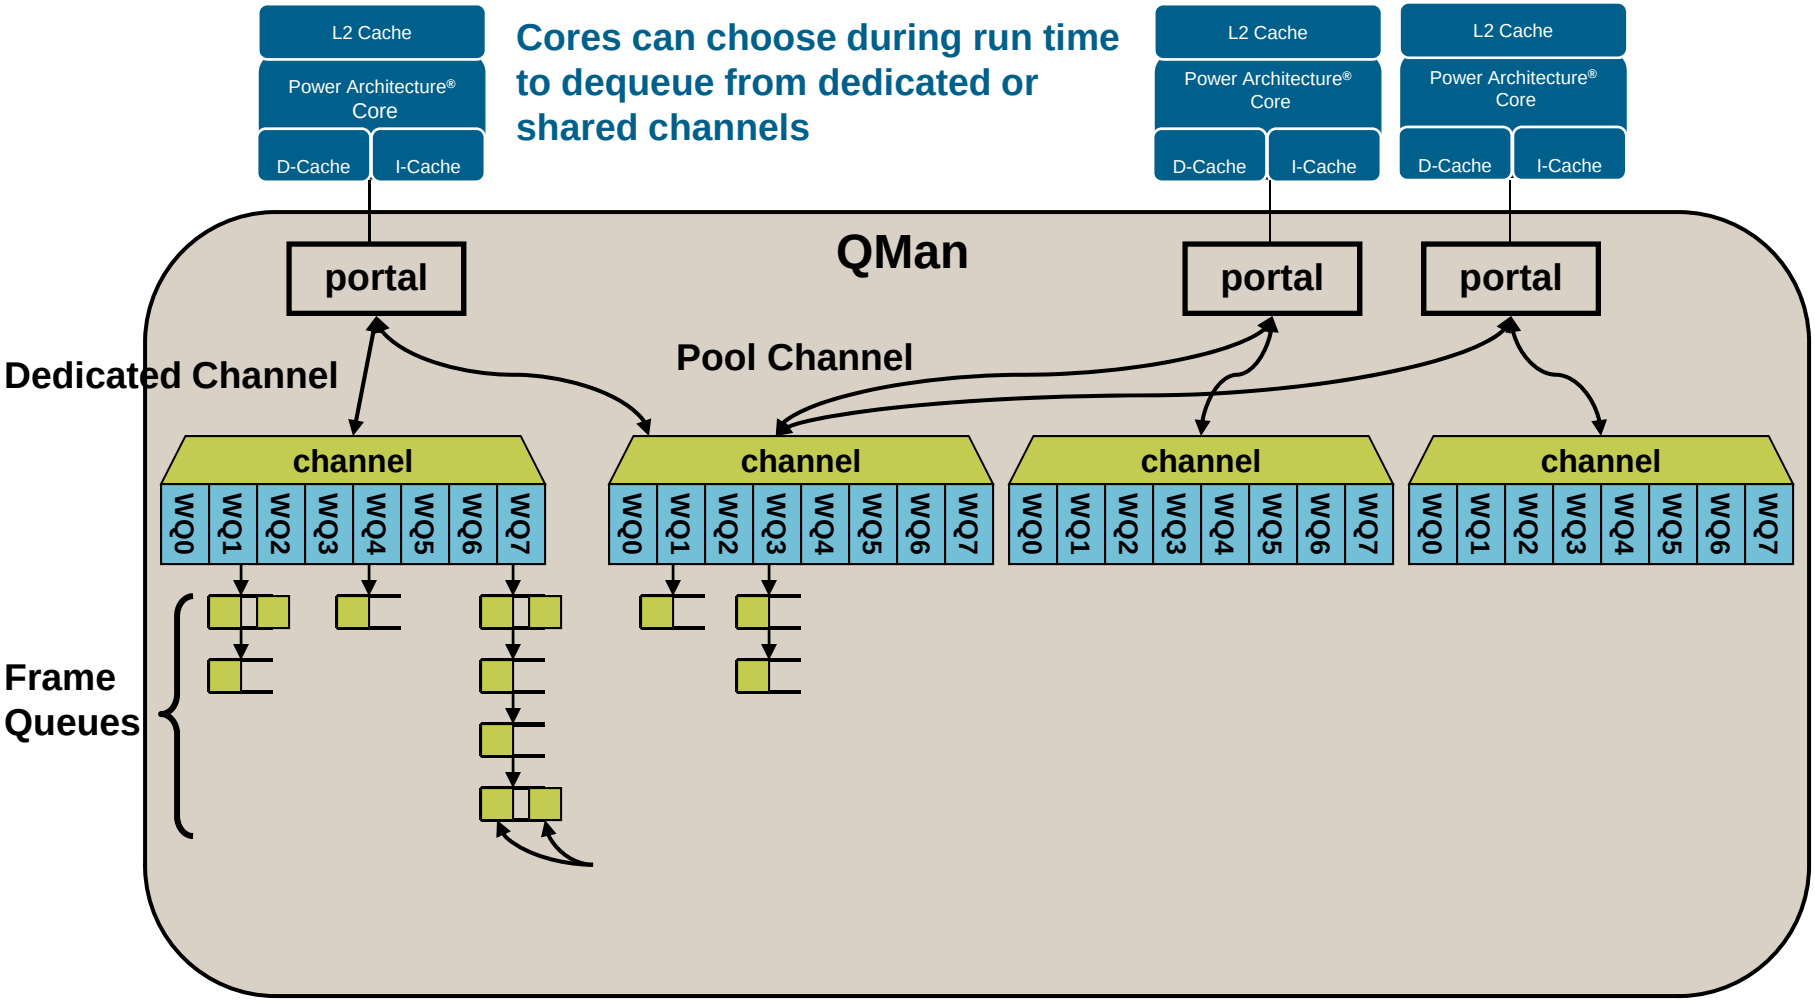

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

QoS supported through Channel Scheduler

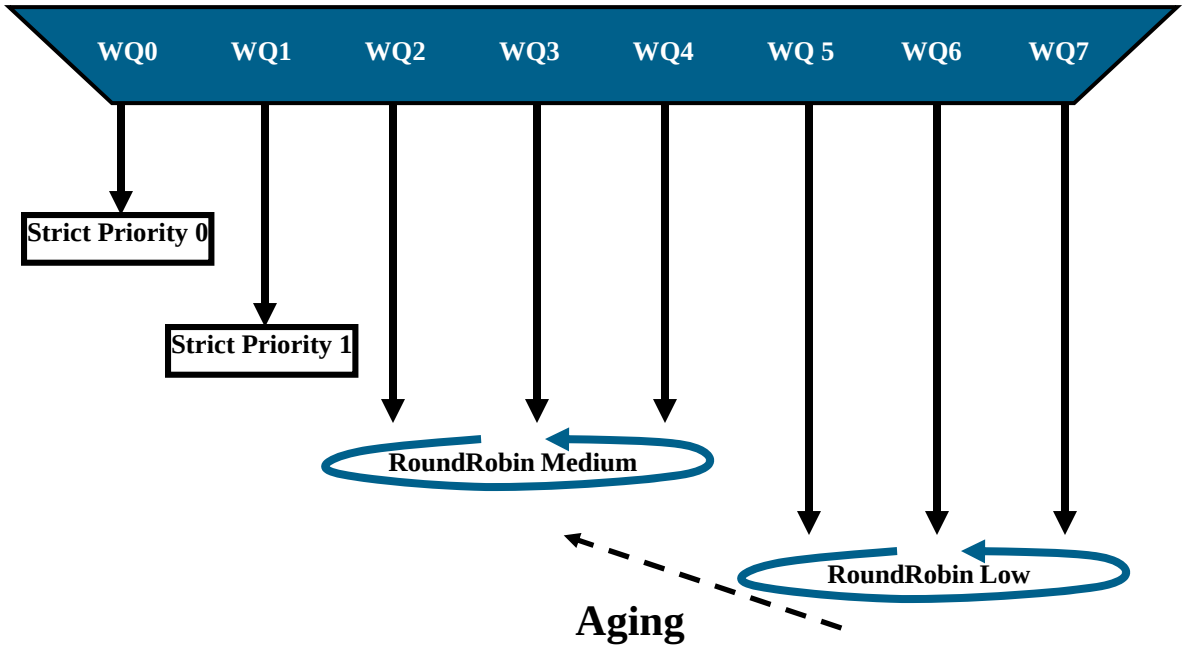

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

Queue Manager (QMan)

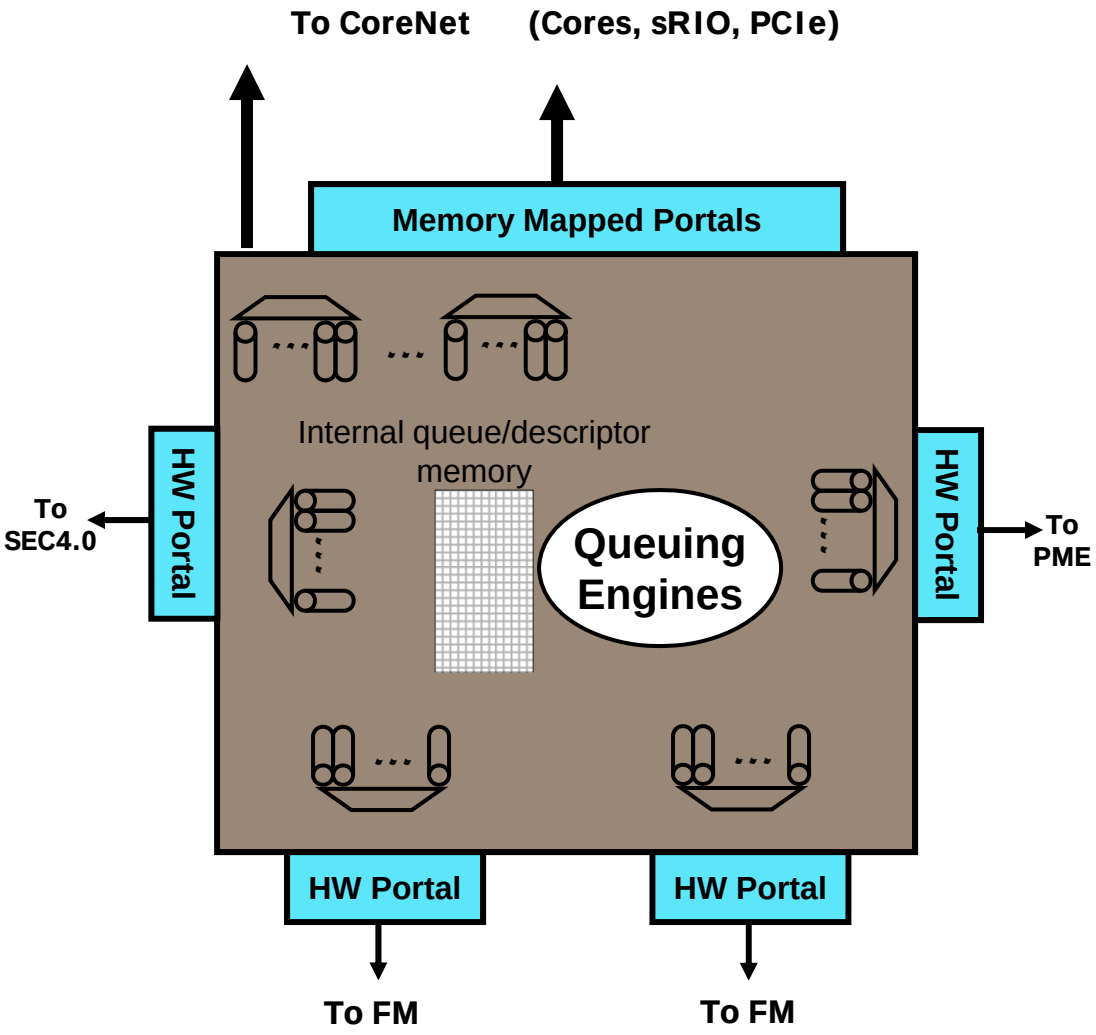

- QMan provides a set of building blocks
  - Frame queues (FQ) which are enqueued onto
    - …Work queues (WQ) which are organized into
      - …Channels with prioritized class scheduling between WQs
        - … which can be used to build connections between blocks (cores, network I/Fs, HW accelerators)
- Frame queues are “logical” queues
  - Actual data is stored in memory buffers, QMan queues “ frame descriptors”
- QMan also supports active queue management:
  - Tail drop on FQs
  - WRED on groups of FQs
  - Tail drop on groups of FQs
- Channels can be shared between consumers which facilitates load spreading
- Channels may also be dedicated to a single consumer

Unified DPAA Frame Formats

Multi-buffer (Scatter/gather) Frame

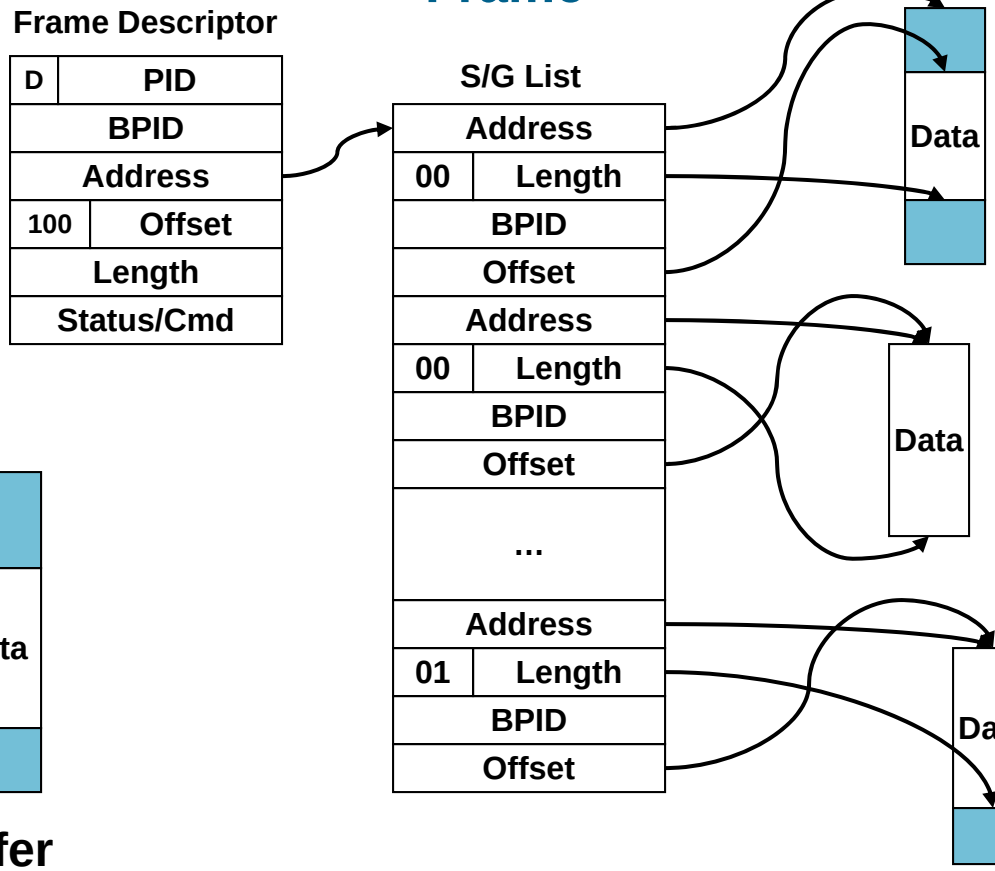

Simple Frame

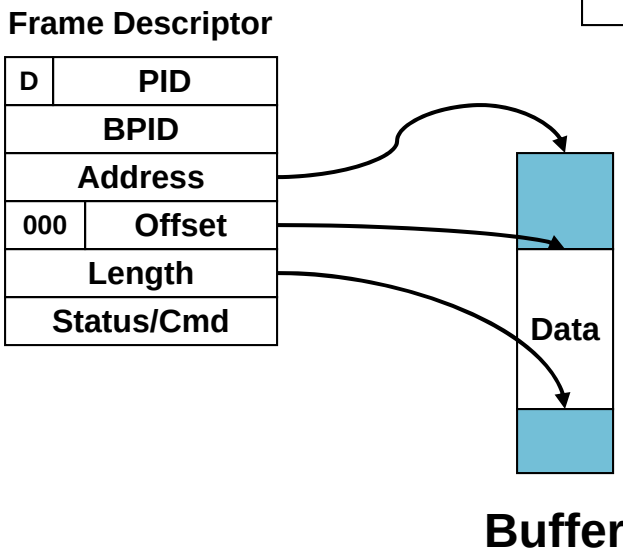

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

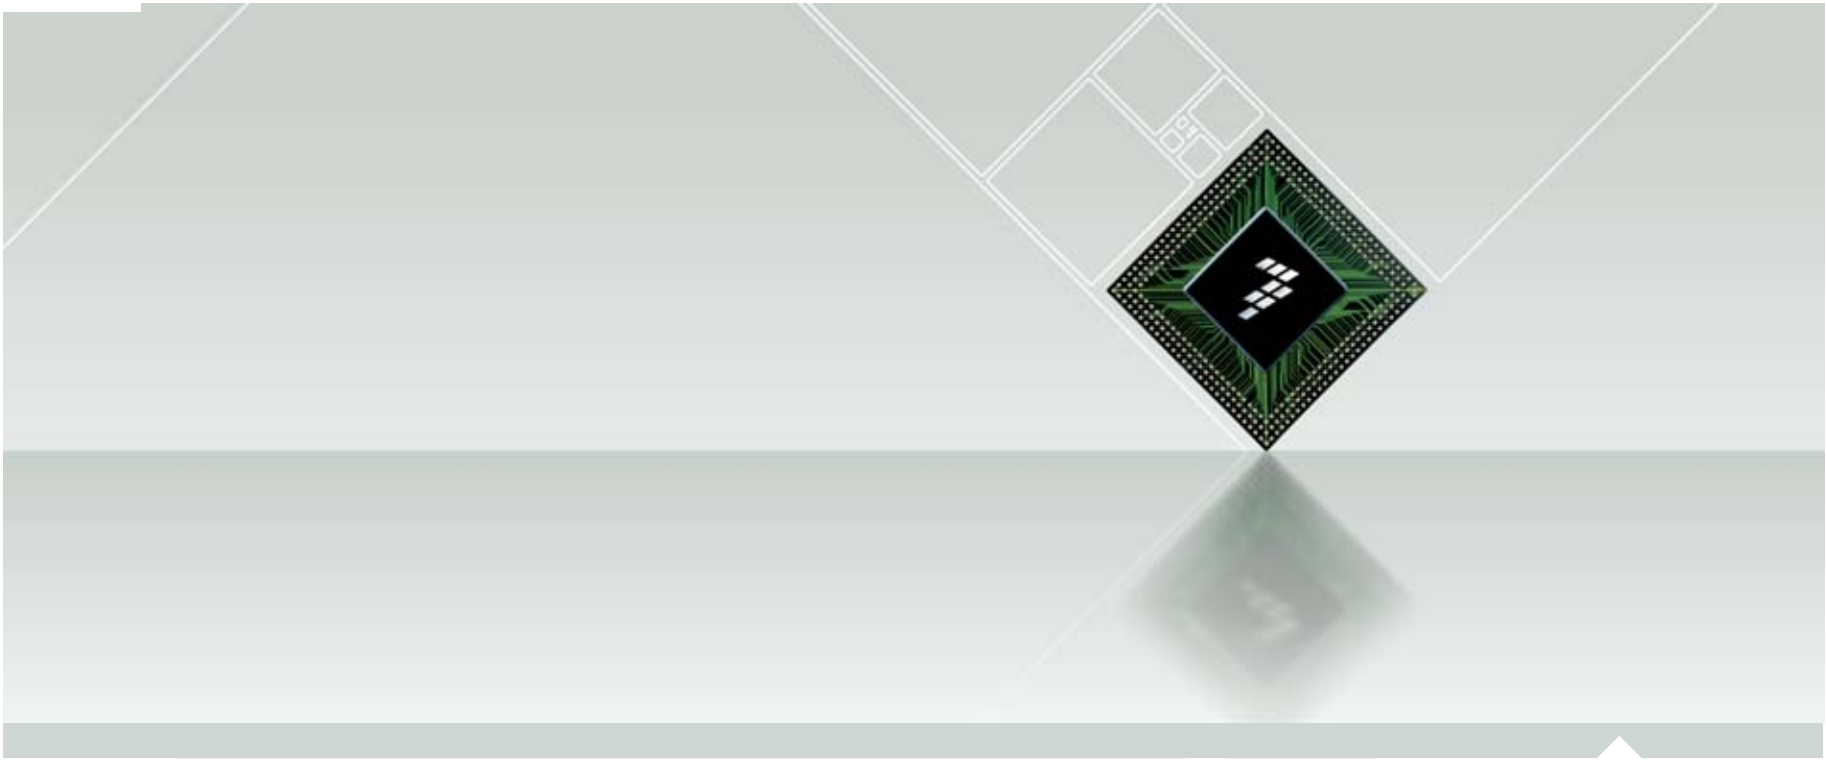

Packet I/O

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

FMan/QMan Ingress Packet Processing

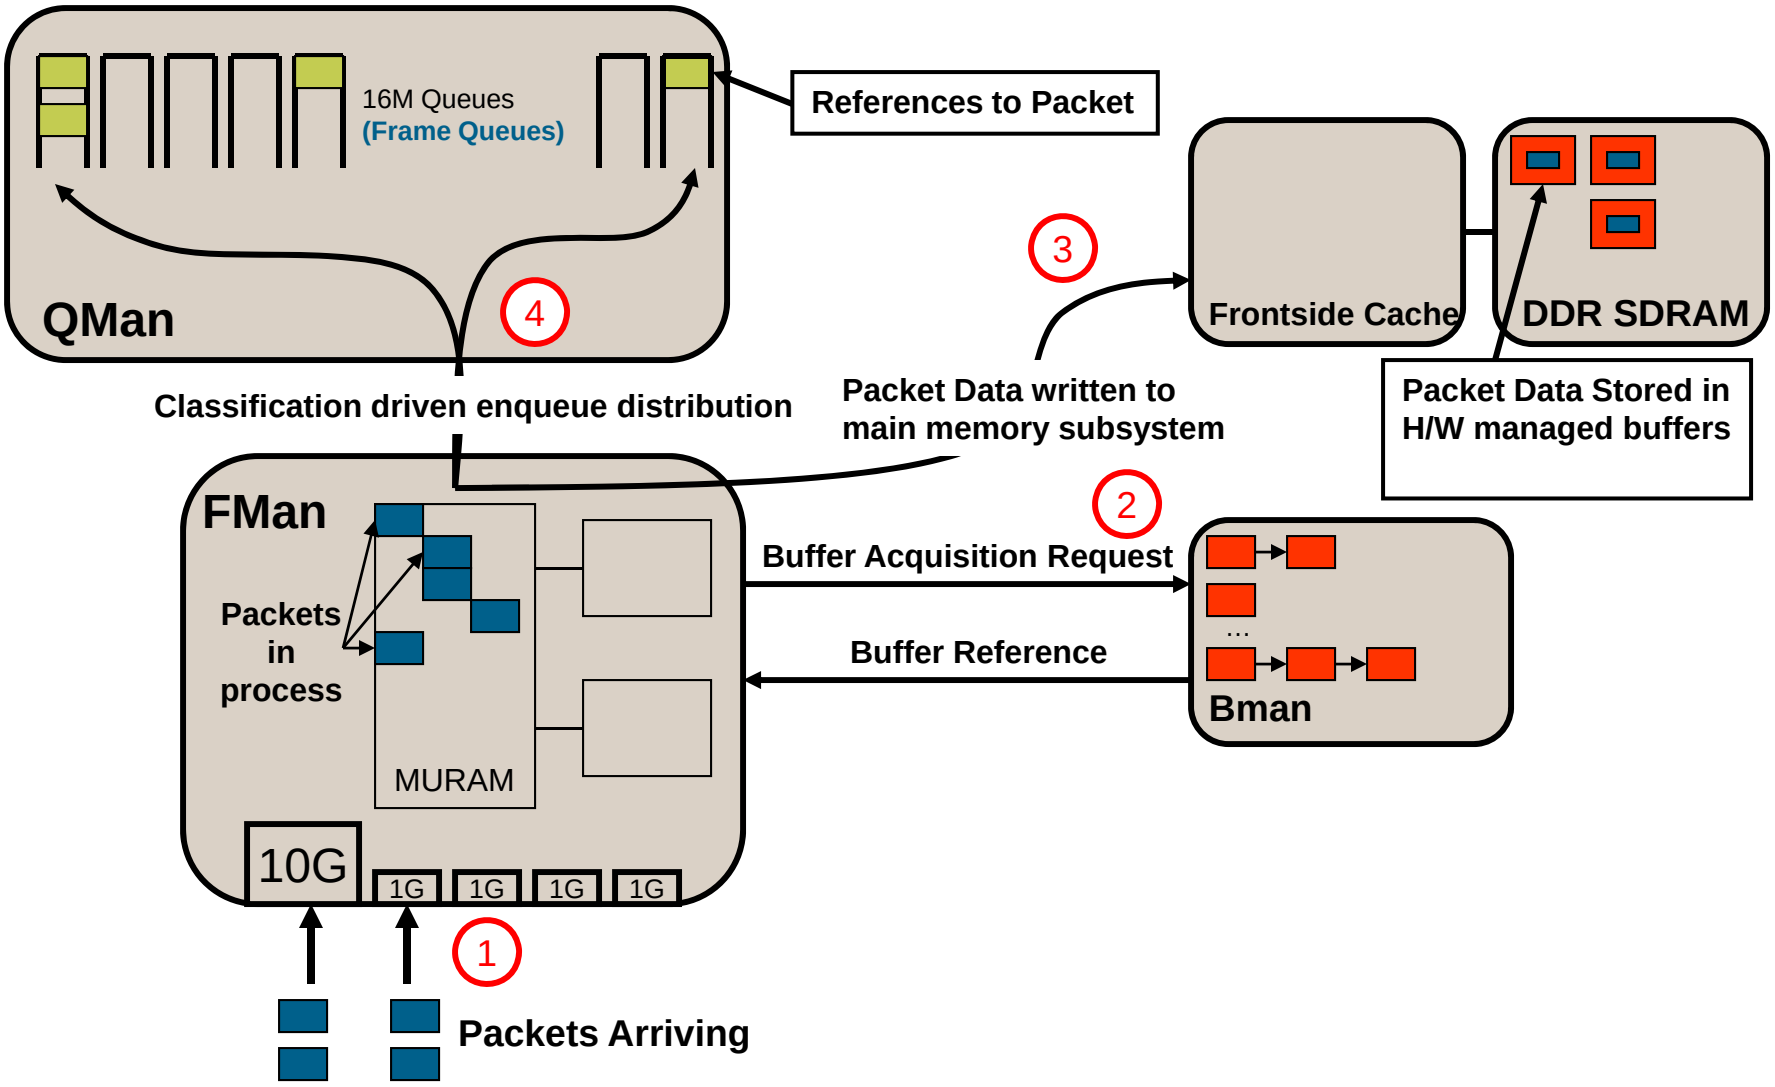

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

Example Location of NAPT Functionality

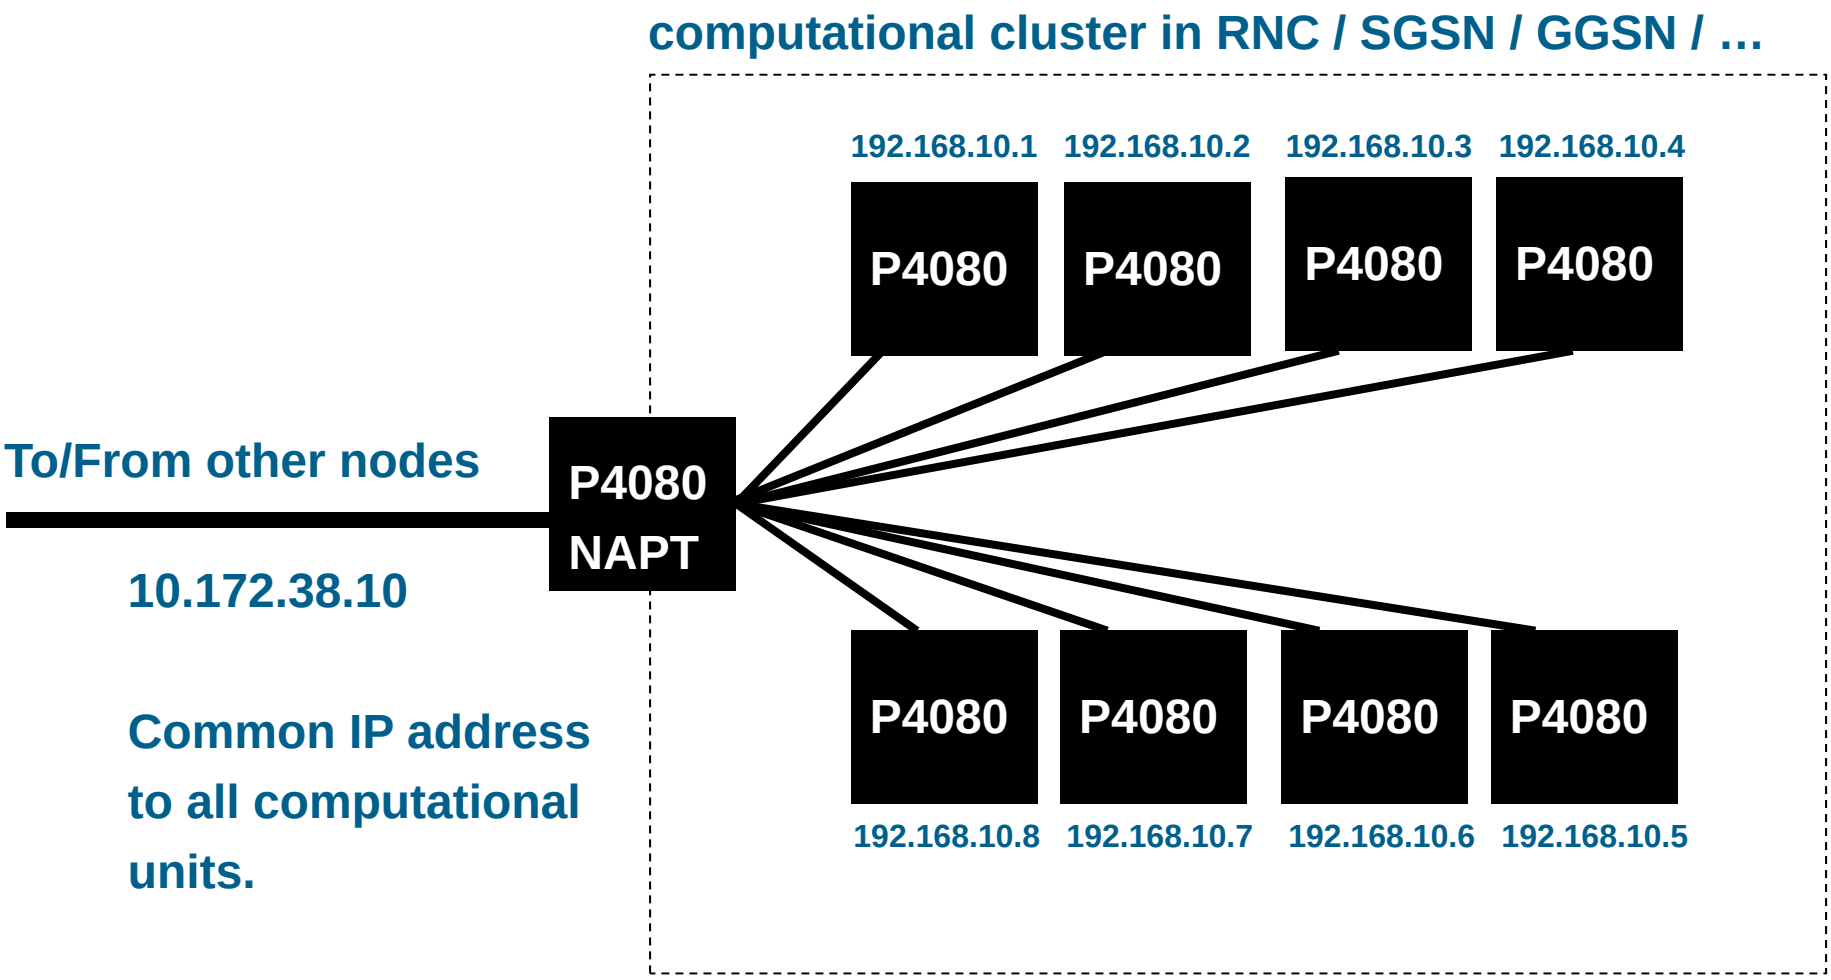

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

NAPT Basic Principle

```
typedef struct SHLL_node
{
SHLL_node           *Next;
uint32_t            IpFrom;
uint32_t            IpTo;
uint16_t            PortFrom;
uint16_t            PortTo;
union mac_address_t EthTo;
uint64_t            StatBytes;
uint64_t            StatPackets;
}
```

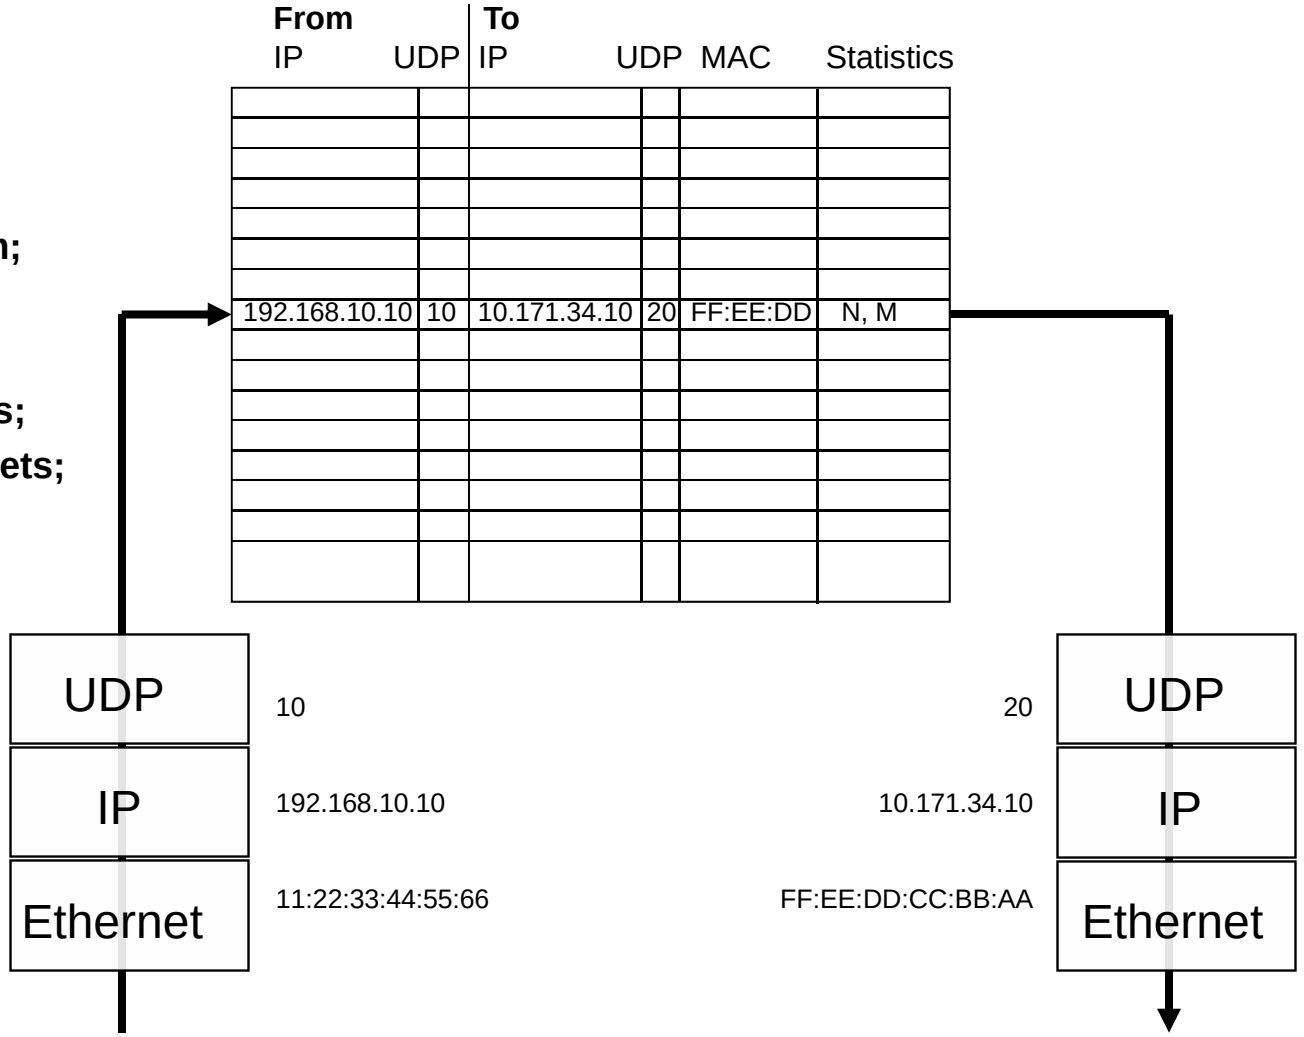

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

System Configuration: Linux® OS and Bare-metal

- Demo implementation based on separated slow- and fastpath
- Linux running on 1 core with control software on top
- 7 cores bare-metal application processing packets
- Hypervisor protects the partitions from each other

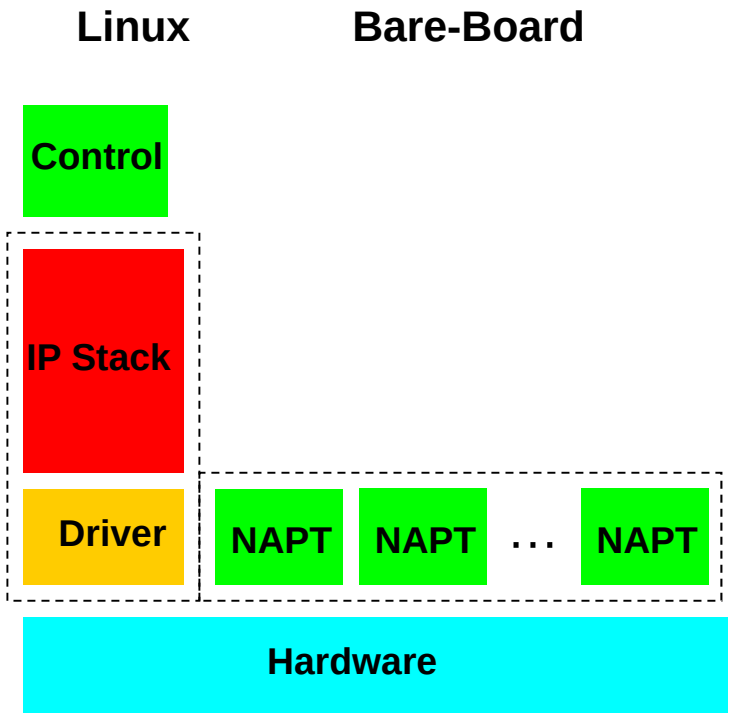

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

Core / Accelerator Partition

- Frame Manager Ingress
  1. Buffer allocation in DDR
  2. Parser / classification / policing
  3. IP checksum
  4. UDP checksum
  5. Enqueue to core

Accelerator

- Software
  6. Dequeue from Frame Manager
  7. Translation lookup
  8. Enqueue to Frame Manager

Core

- Frame Manager Egress
  9. Dequeue from core
  10. Add UDP checksum
  11. Add IP checksum
  12. Buffer deallocation from DDR

Accelerator

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

RMan Datapath

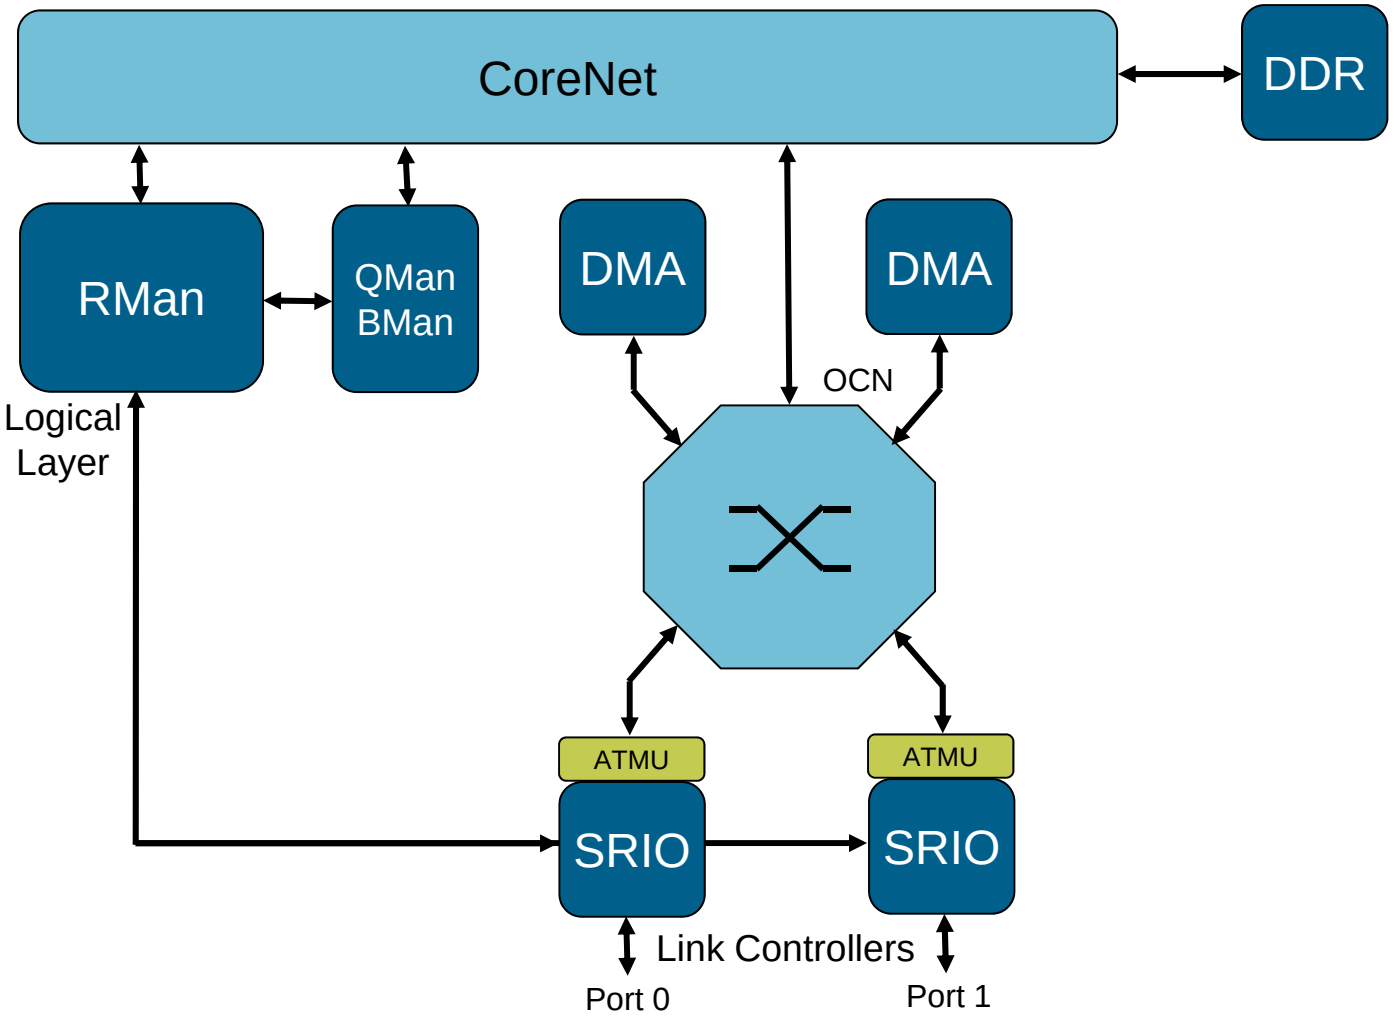

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

RMan for QorIQ: Greater Performance and Functionality

- Many queues allow multiple inbound/outbound queues per core
  - Hardware queue management via QorIQ Datapath 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

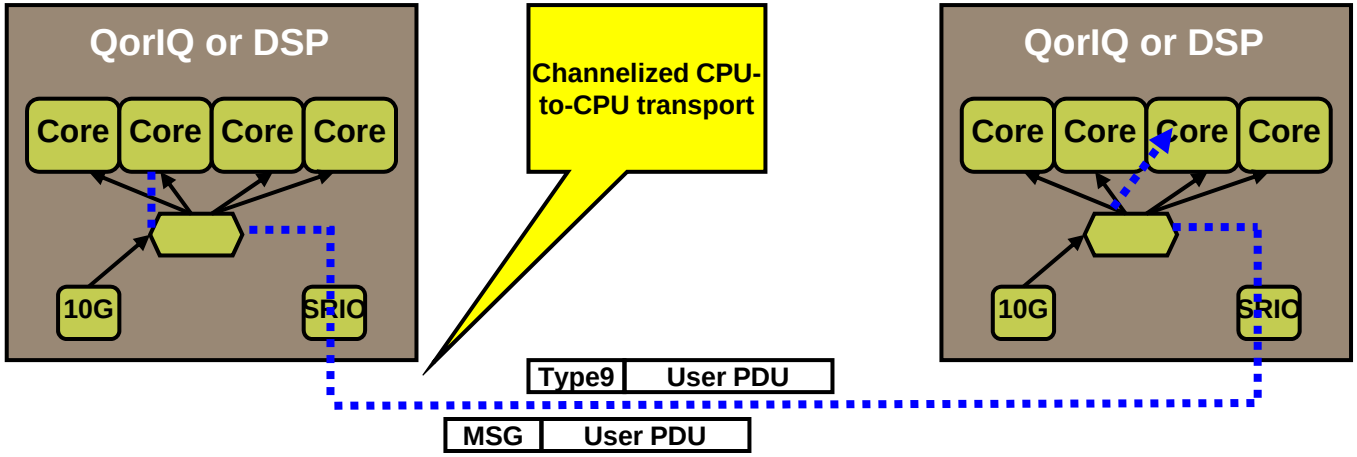

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

RMan Enables New Zero-CPU Overhead Use Models

Scalable Multicore System

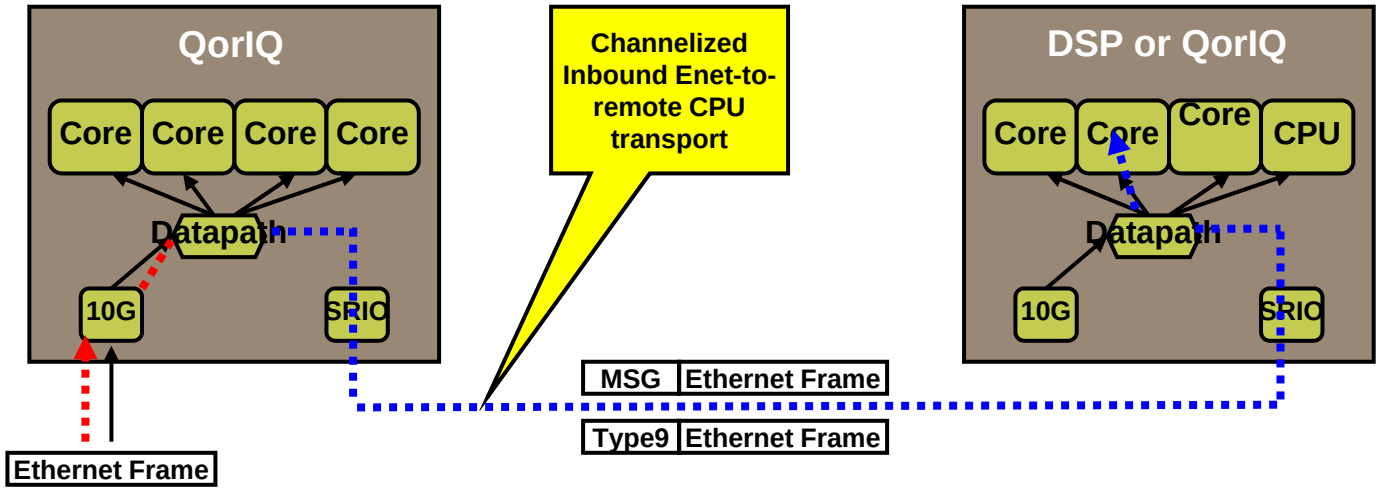

Ethernet Bridging

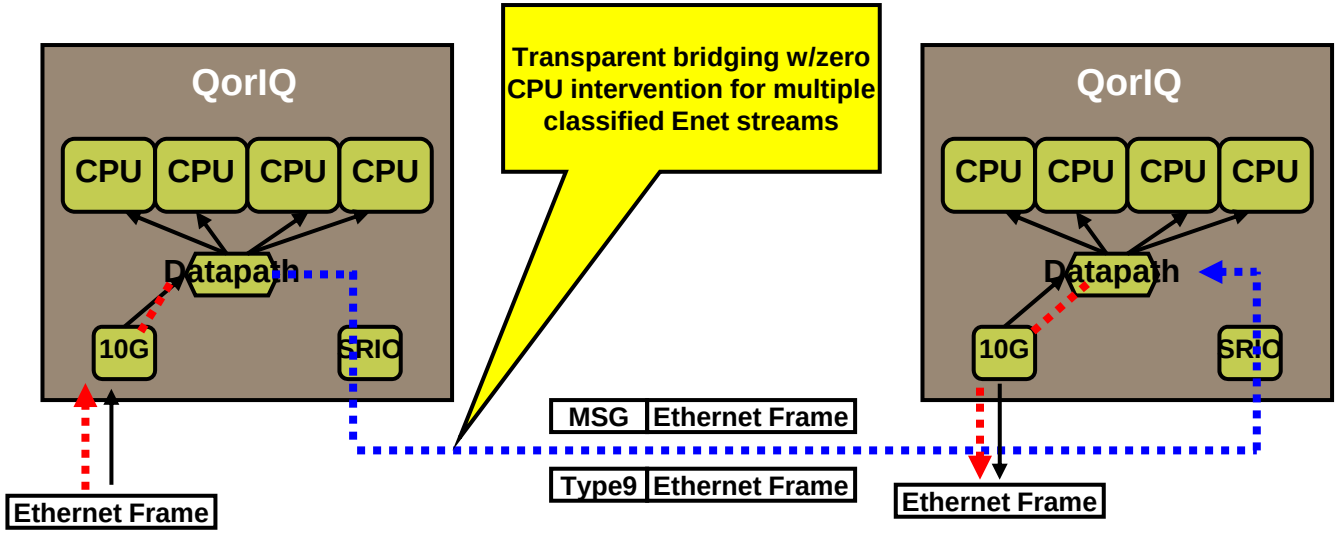

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

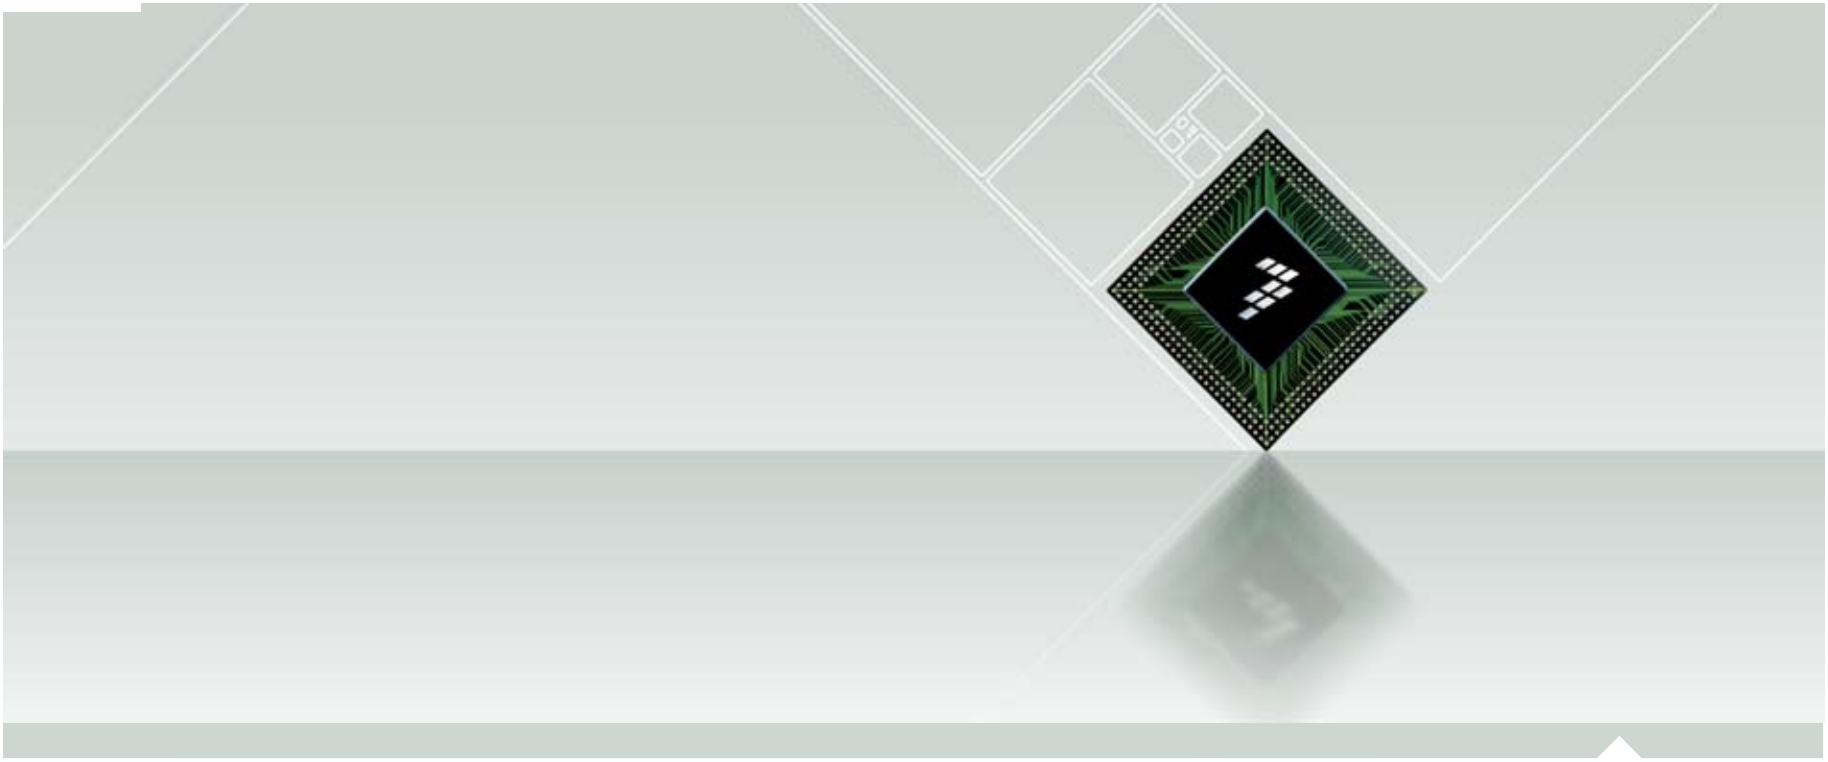

PME: Traffic Monitoring and Management

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

Deploy Deep Packet Inspection and Policy Control

- DPI enables companies to:
  - Understand the network traffic and pattern
  - Gather business intelligence
  - Identify trends and adapt to those trends
- Policy Control enables:
  - A smarter pipe requires fine-tuned network controls
  - Control and manage growing usage
  - Fair usage to all network users

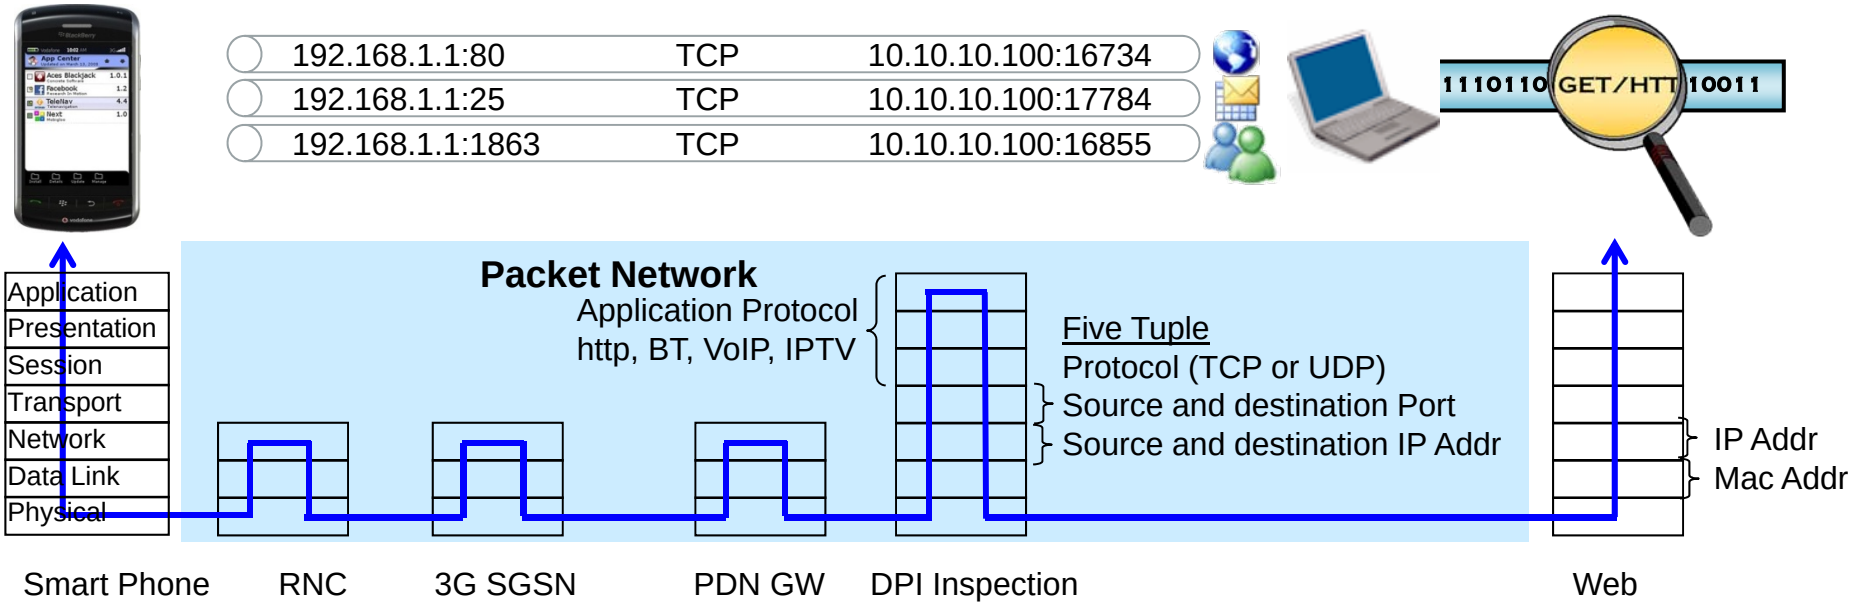

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

Policy Control: DPI with Pattern Matcher

- 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
- Stateful rules operate on a per session basis
  - User-defined logic reacts to pattern matches detected by the DXE
  - Can be used to further qualify the pattern match
  - Protocol state tracking (e.g. track the “normal” transitions of SMTP)

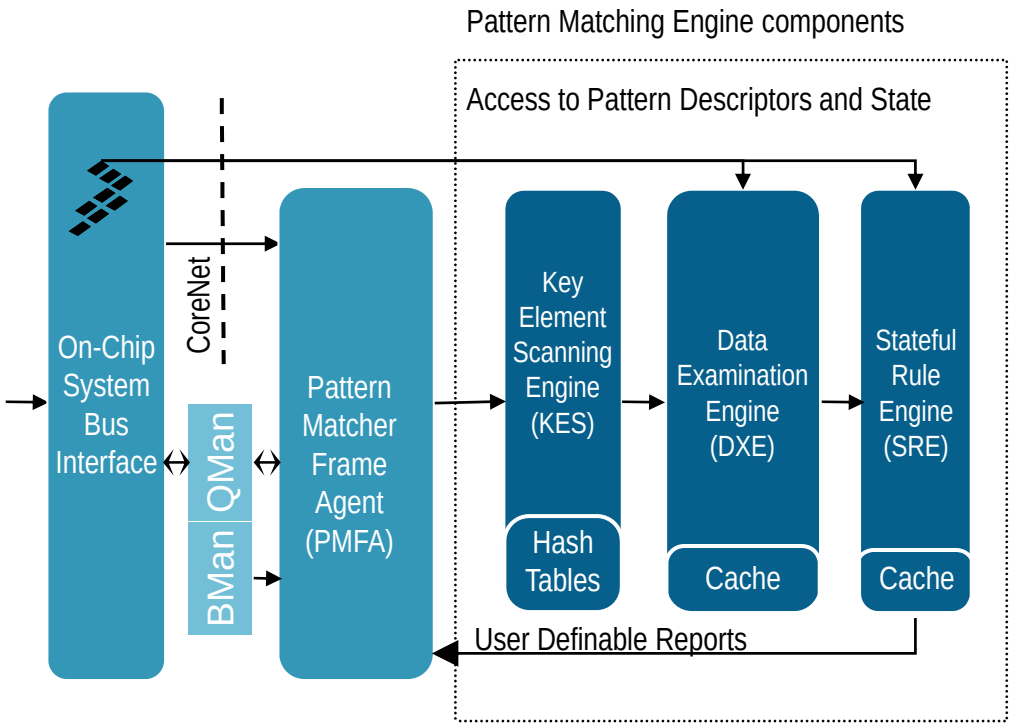

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

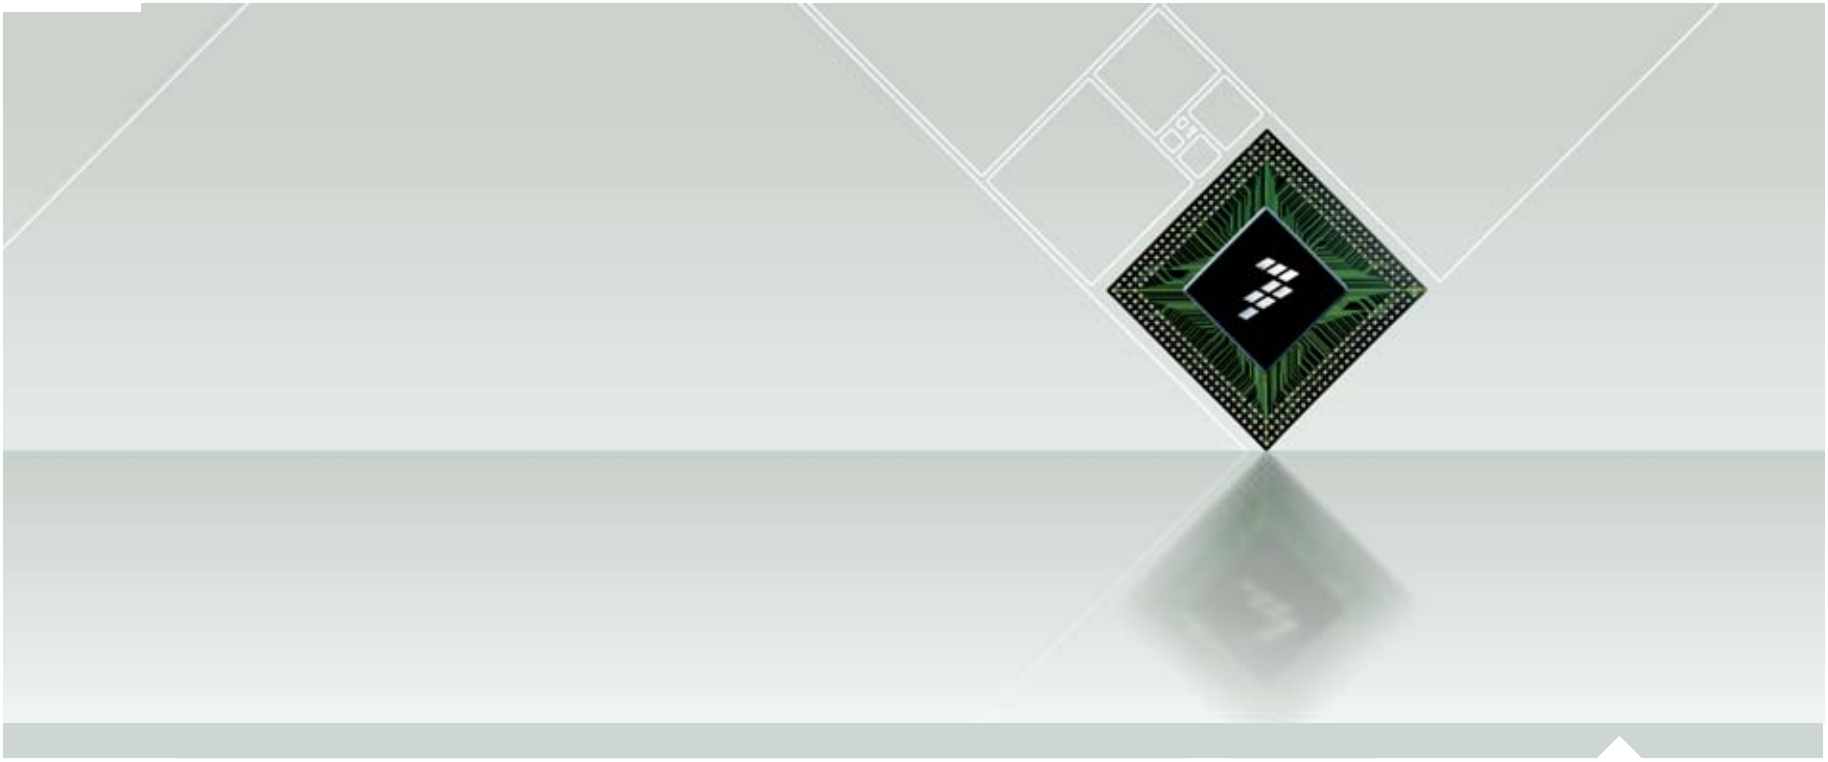

SEC: Ciphering and Packet Formatting

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

P4080 RLC AM Downlink Processing

The SEC dequeues from its highest priority FQ, an action which gets it the Shared Descriptor, the FD, SG List, and initial data.

The Shared Descriptor is only fetched once per FQ dequeue, and a configurable max number of FDs can be processed from that FQ.

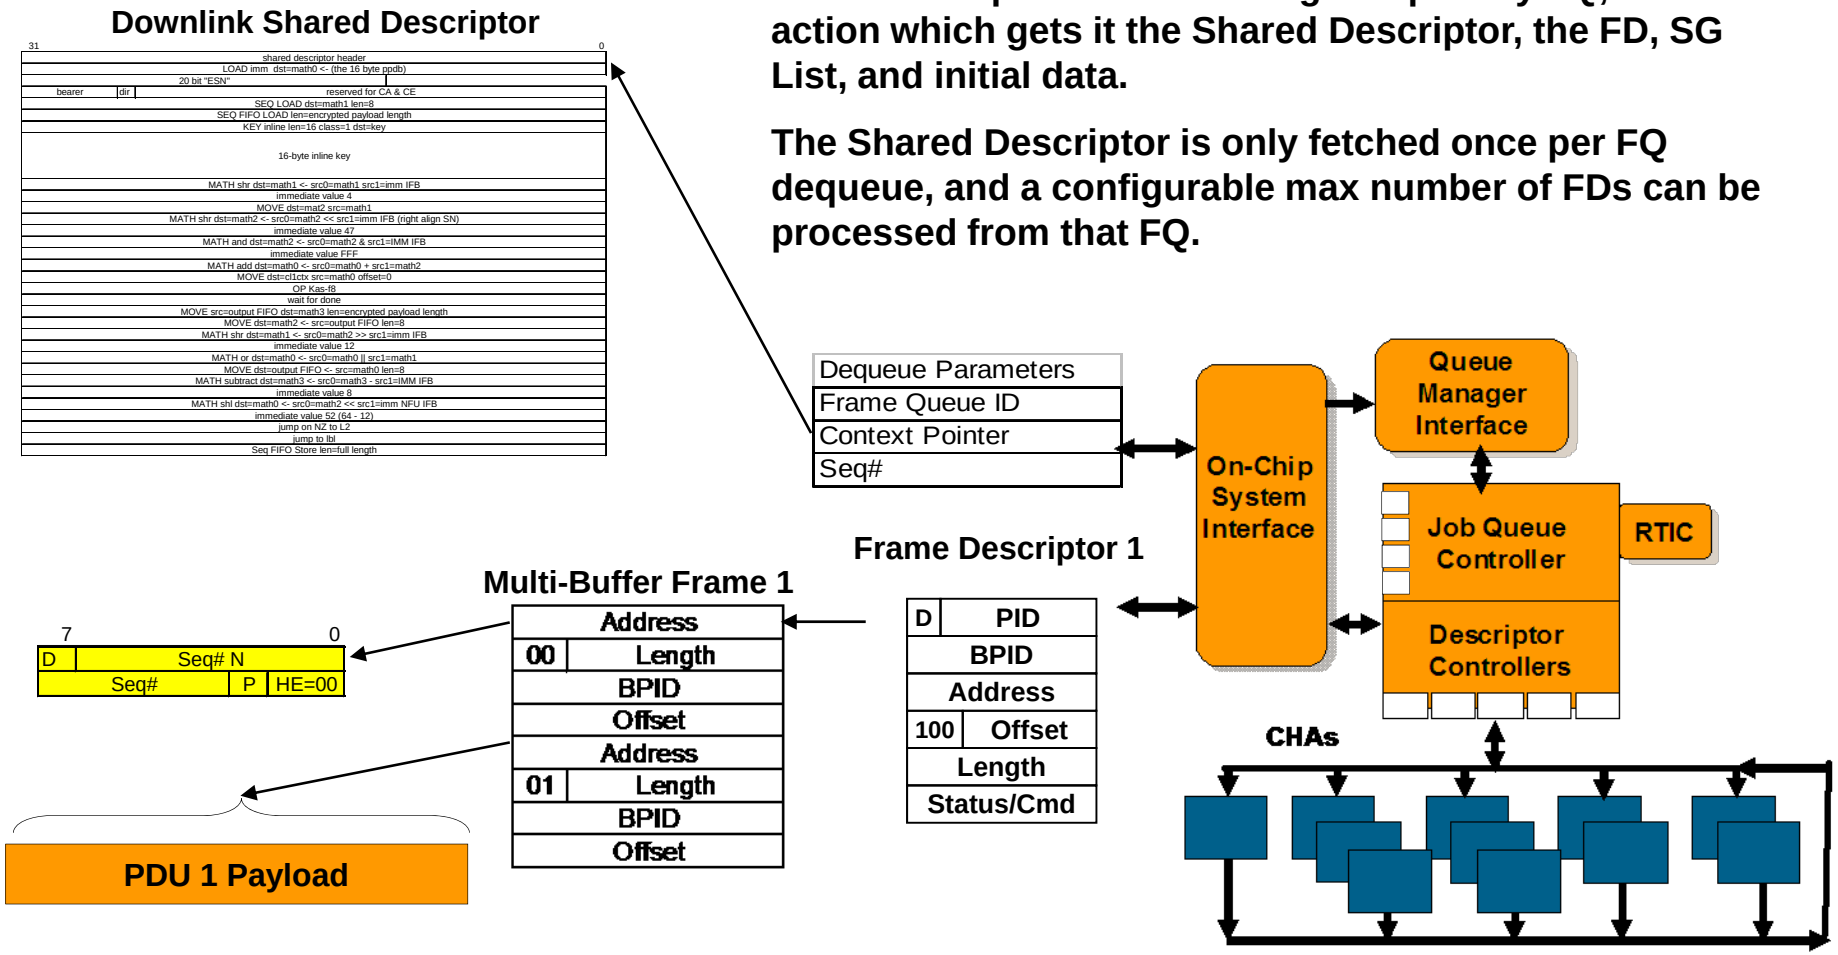

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

SEC 4.0 for P4080 rev2

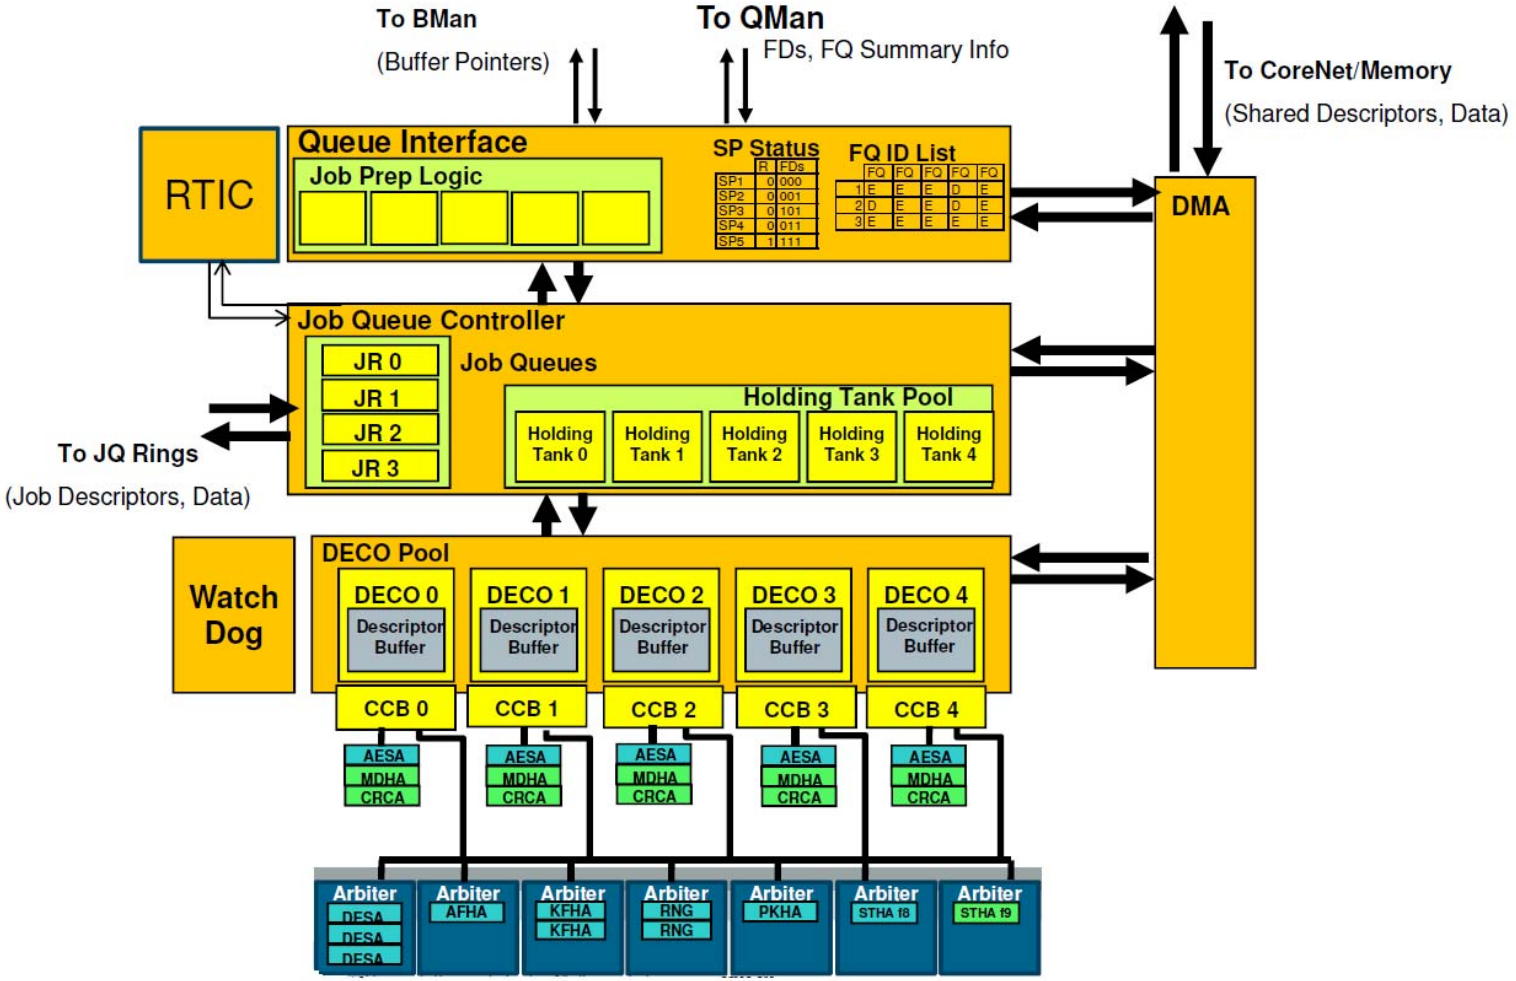

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

DECO Commands

- Load Commands:
  - [SEQ] KEY – Load cipher keys
  - [SEQ] LOAD – Load register
  - [SEQ] FIFO LOAD – Loads In/Out FIFO
- Store Commands:
  - [SEQ] STORE – Store register
  - [SEQ] FIFO STORE – Store In/Out FIFO

- MOVE – Move data between SEC registers.

- MATH – Perform arithmetic operation (sum, and, or, xor, bitshift, …)

- OPERATION – Execute cipher operation (Kasumi, Snow, AES, …)
- JUMP – Branch in the descriptor
- SEQ IN/OUT PTR – Sets pointer for in/out sequence data.
- JD/SD HEADER – First word in the descriptor.
- SIGNATURE – Last word in trusted descriptor.

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor ™

P4080 RLC AM Downlink Processing

An output Frame Descriptor is sent out from the SEC that in turn points to the generated output data. Input data is either left in memory but can also be freed and the buffer returned to BMan.

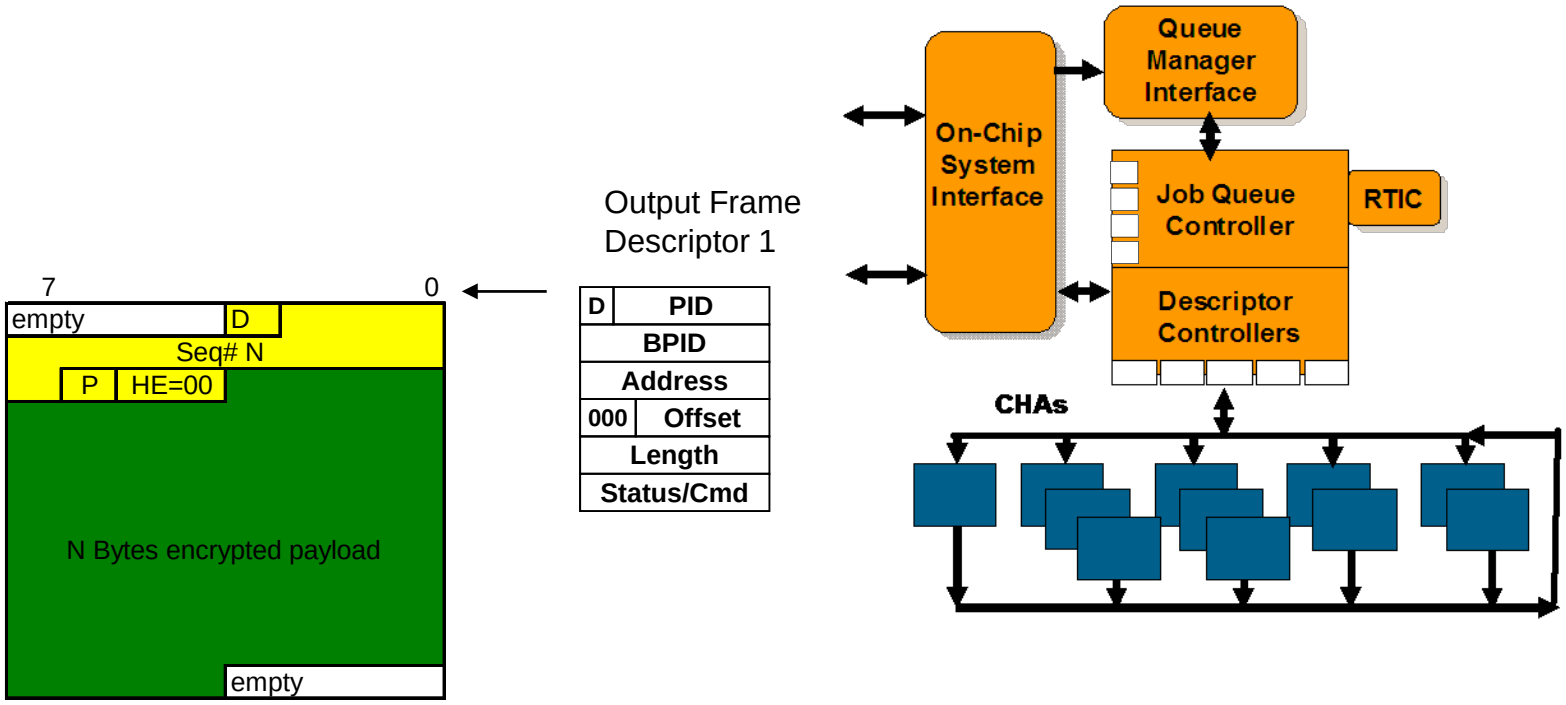

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

P4080 RLC AM Downlink Processing

The frame descriptors is enqueued into a work queue / channel back to a core for further processing.

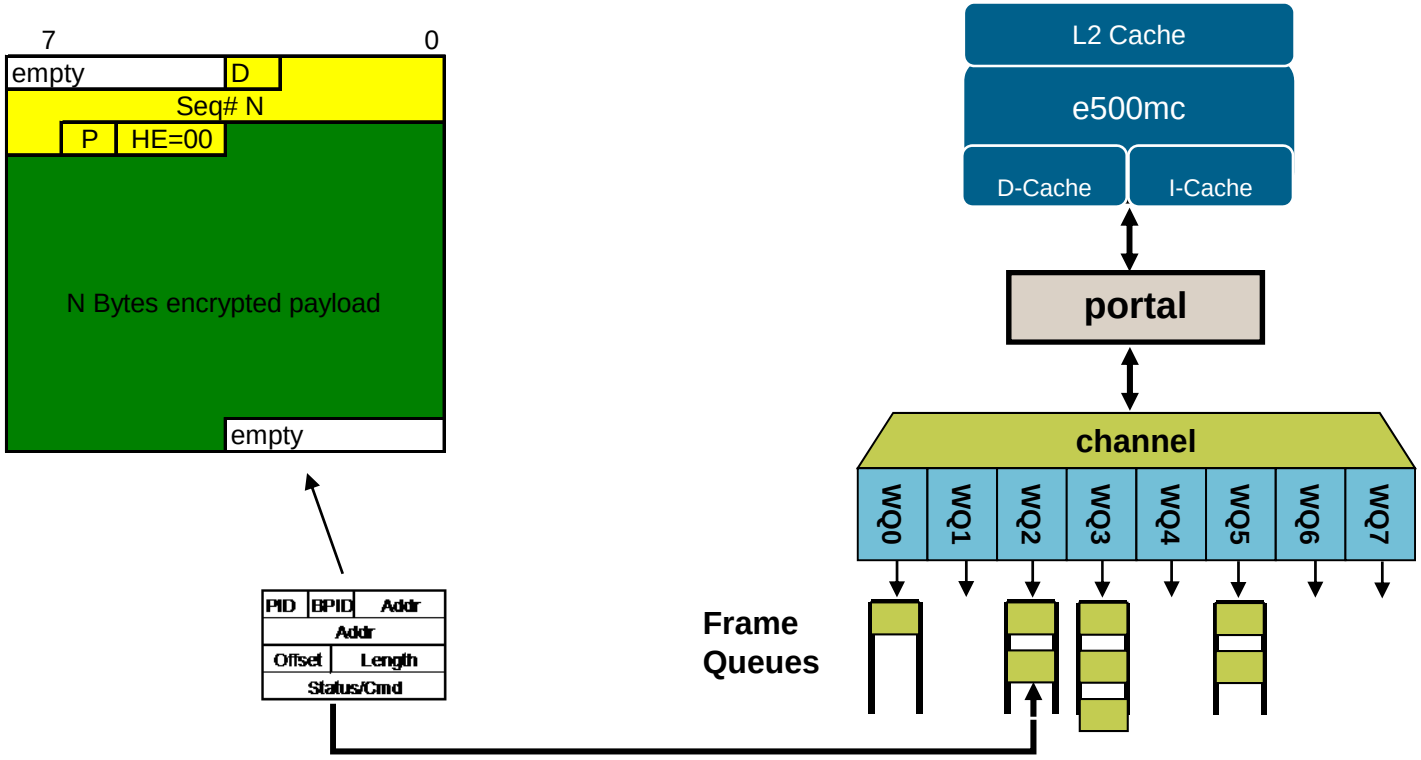

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

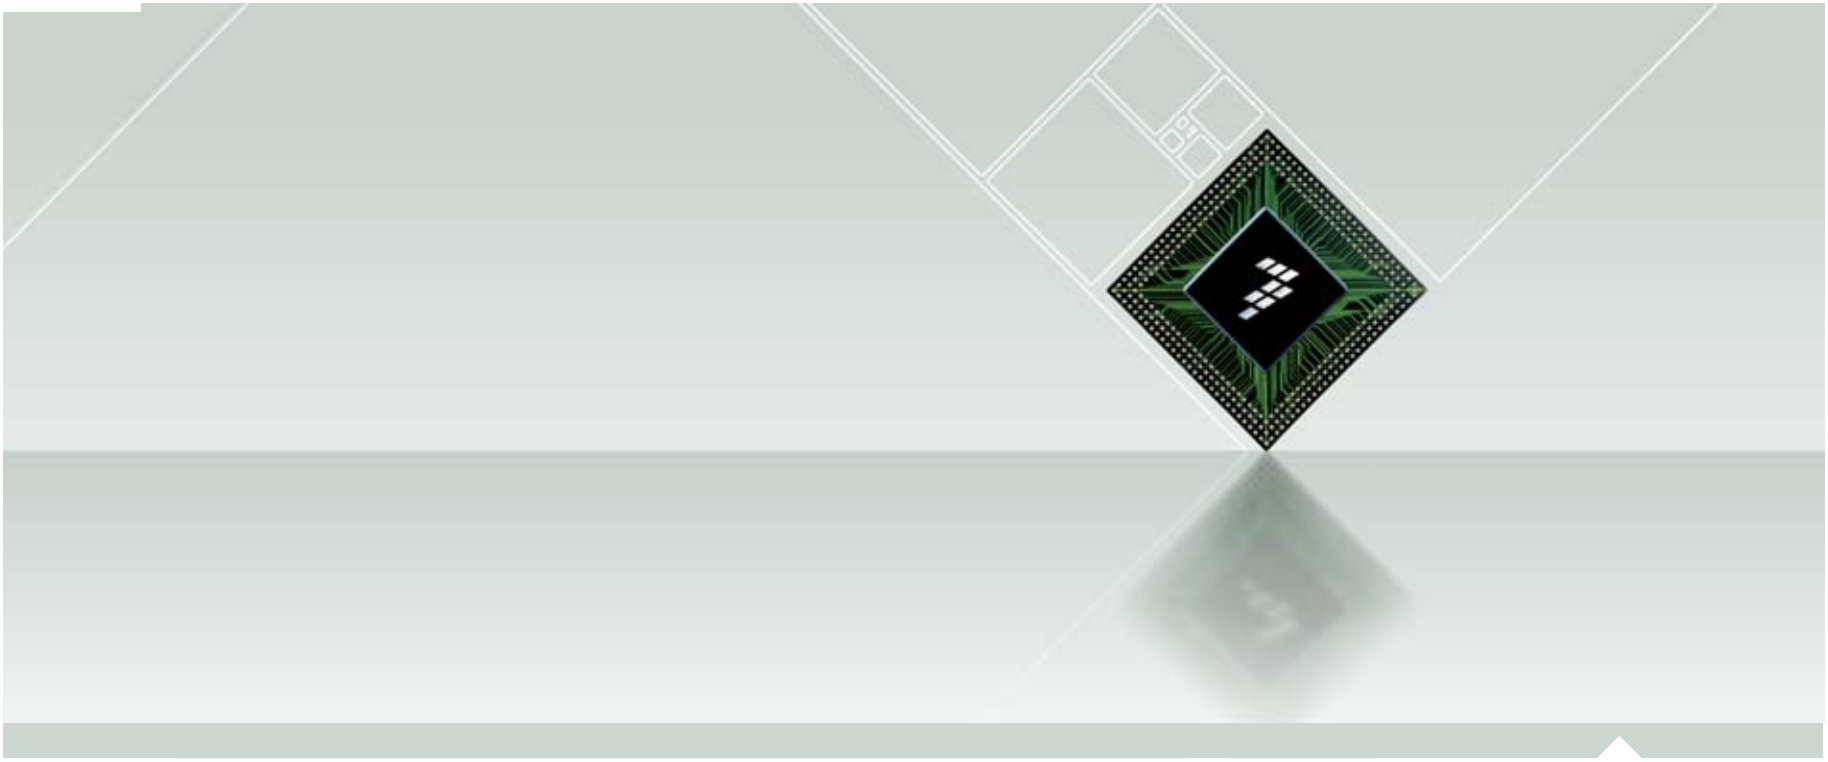

Use Case: Voice Packet Processing

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

Challenge: Voice Packets

- Voice Packets are ciphered/formatted in RNC (3G) and (e)NB (LTE and 3G flat architecture).
- Voice packets are large to the number but small to the size (~20 bytes).
- Typically each active flow has a packet per ~20ms, i.e. it is not possible to group packets of same flow or latency will increase and phone conversation will not work (lag severely).
- Hardware Acceleration is a big benefit but overhead/byte of sending to accelerator has been a big problem since packets are so small.

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor ™

Per Flow Ciphering

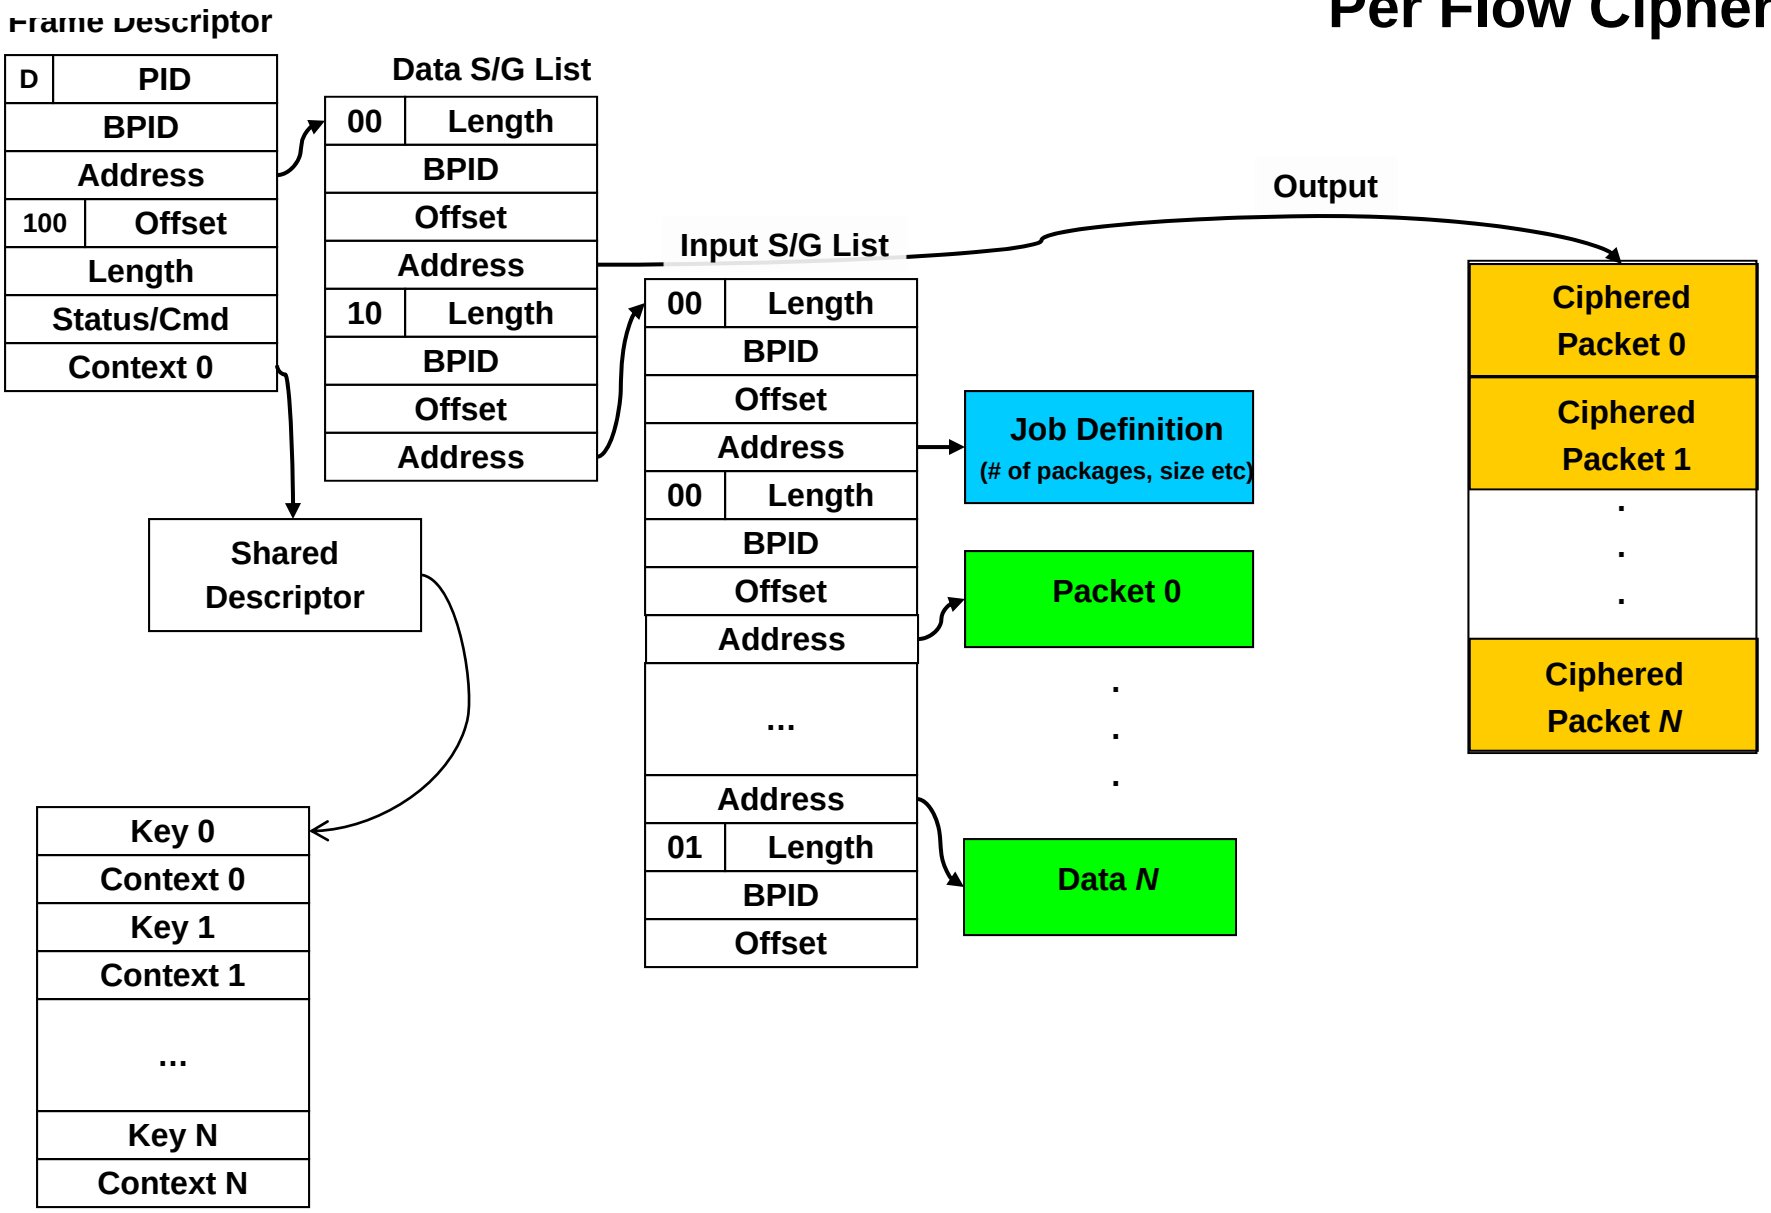

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

S/G output – buffers at random location

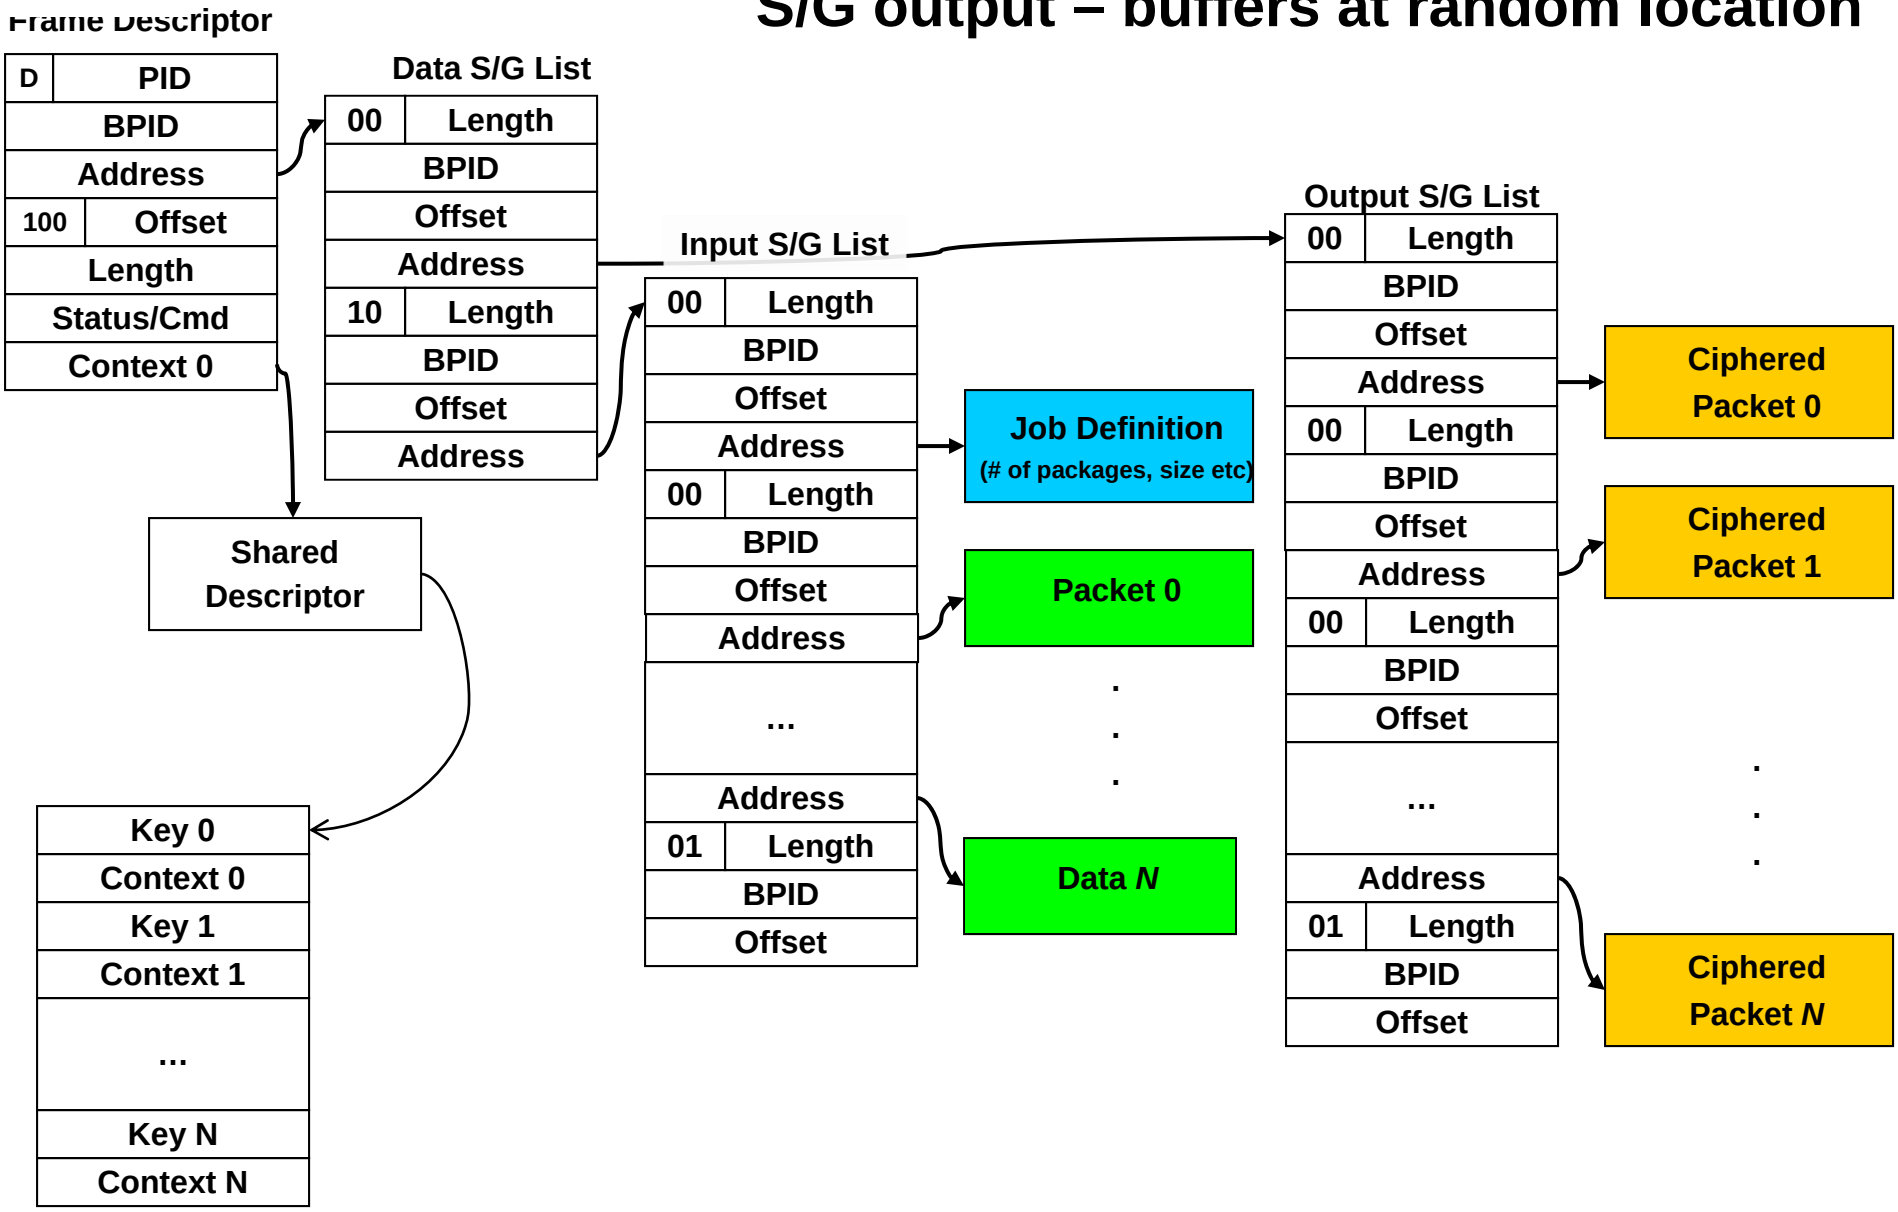

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

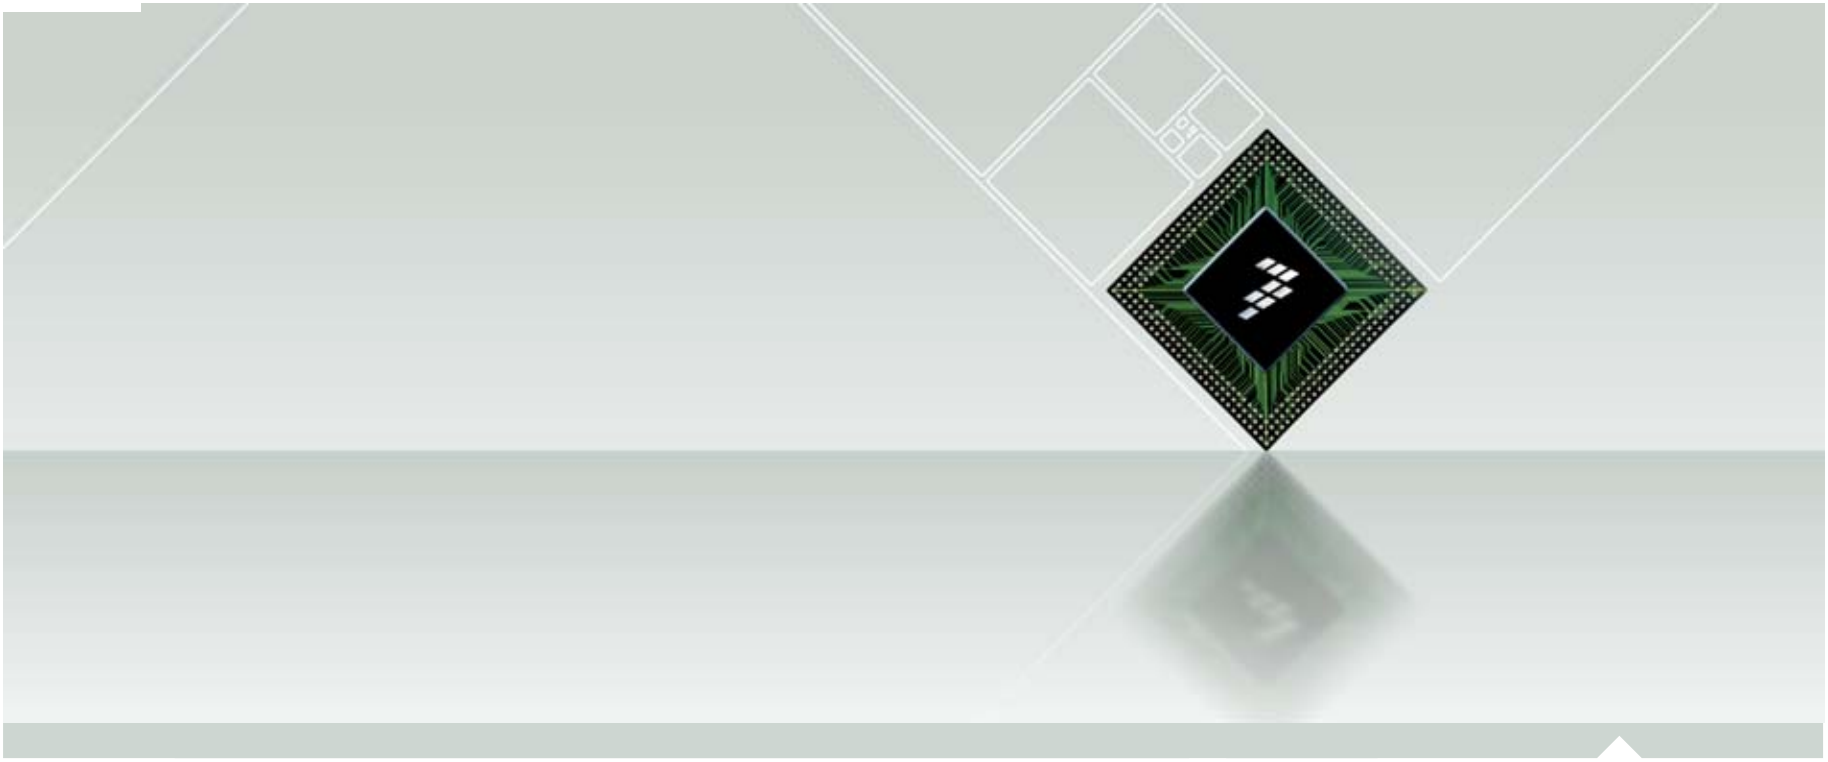

Use Case: HSPA Mobile Broadband

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

Challenge: Data Packets

- Huge total bandwidth!
- Very high individual user peak rate.
- Complex protocol processing.

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor ™

Overall Flow HS

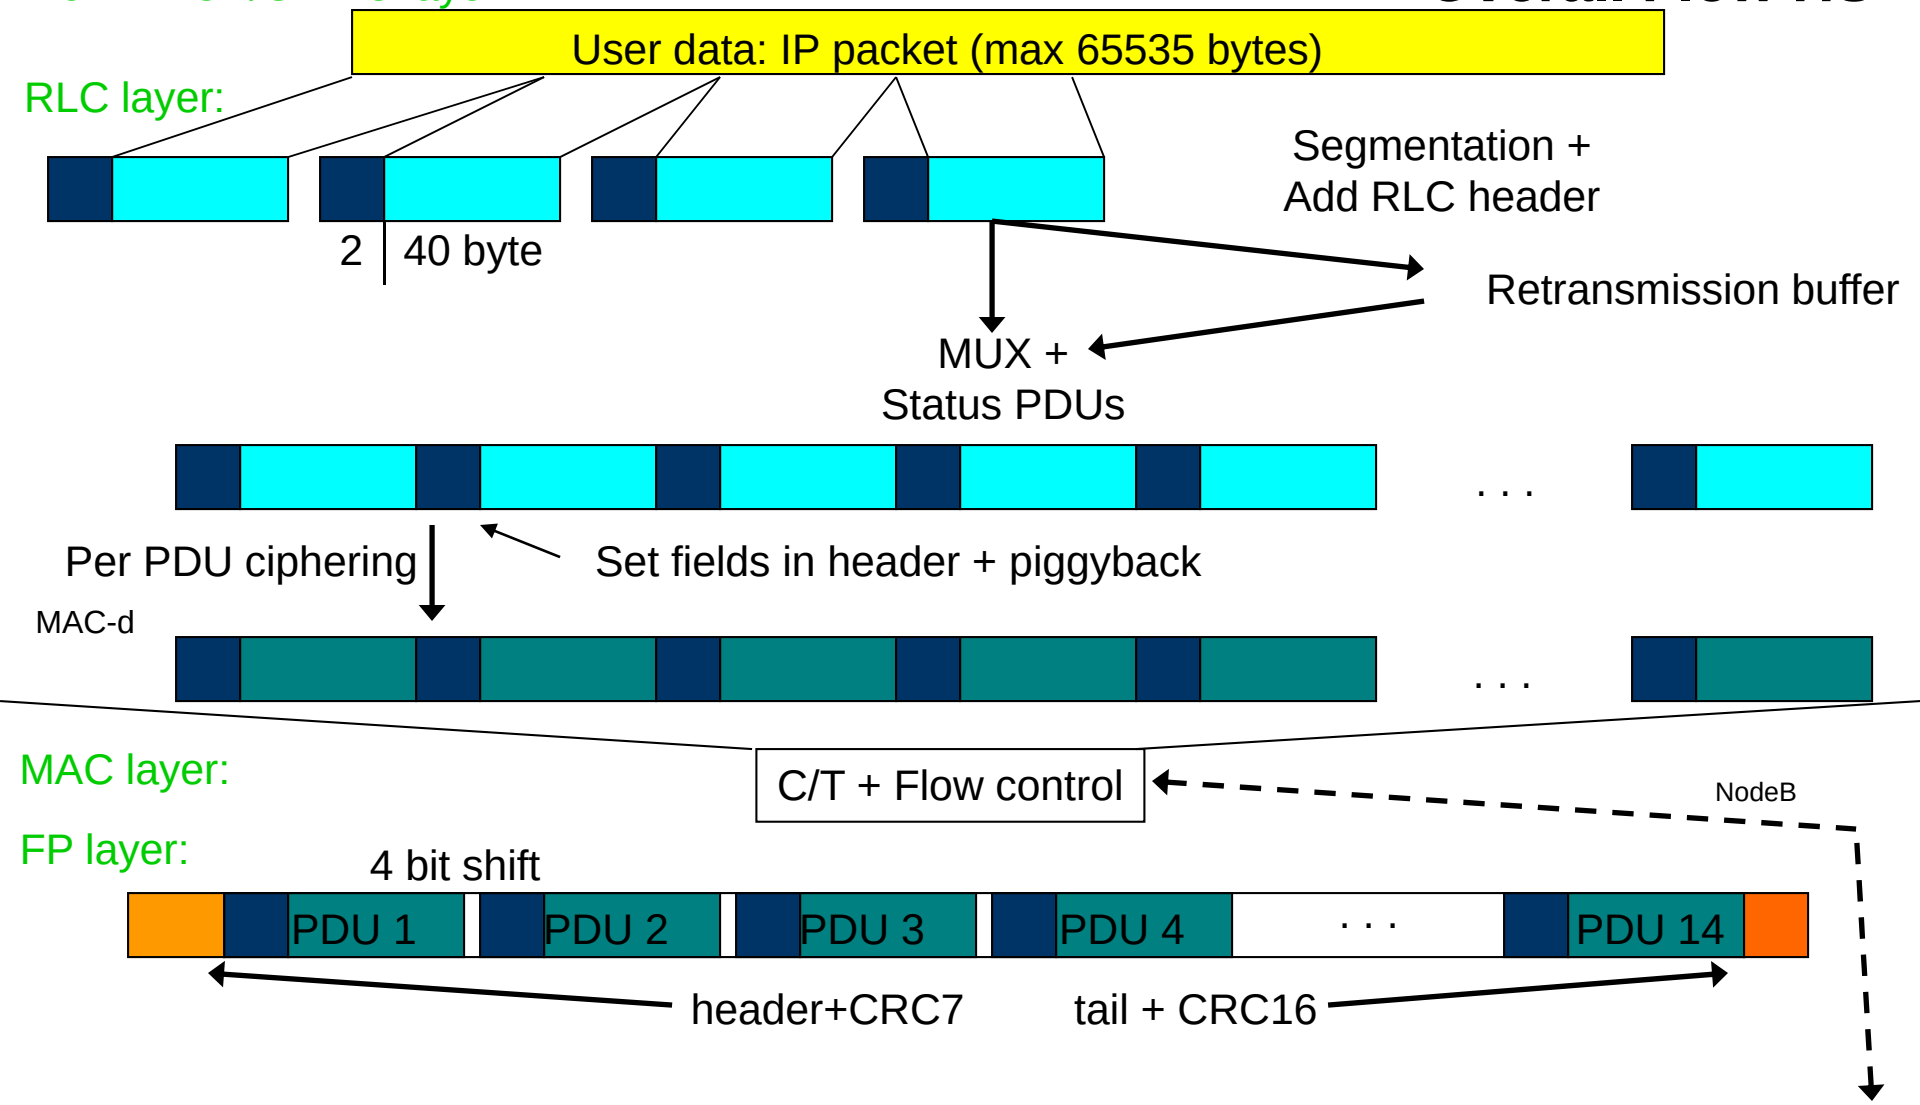

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

Frame Type 1 with MAC and RLC PDU – R7

Gray = Depends on configuration and version

Blue = Empty

Example of special cases:

E bit can have “alternative interpretation”

HE valid values differs over revisions

Length field can be 7 or 15 bits

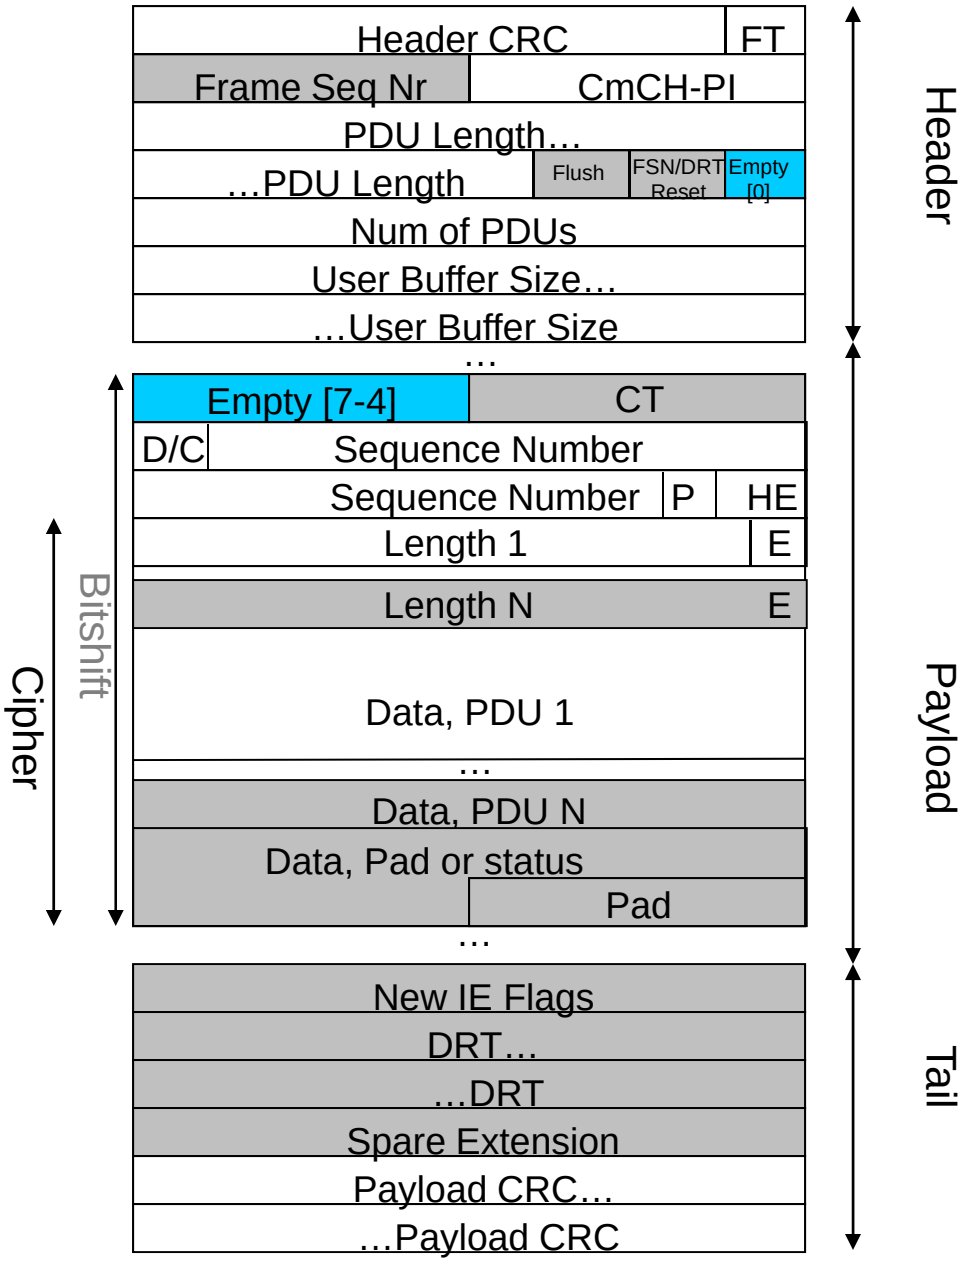

Frame Type 2 with MAC and RLC PDU – R7

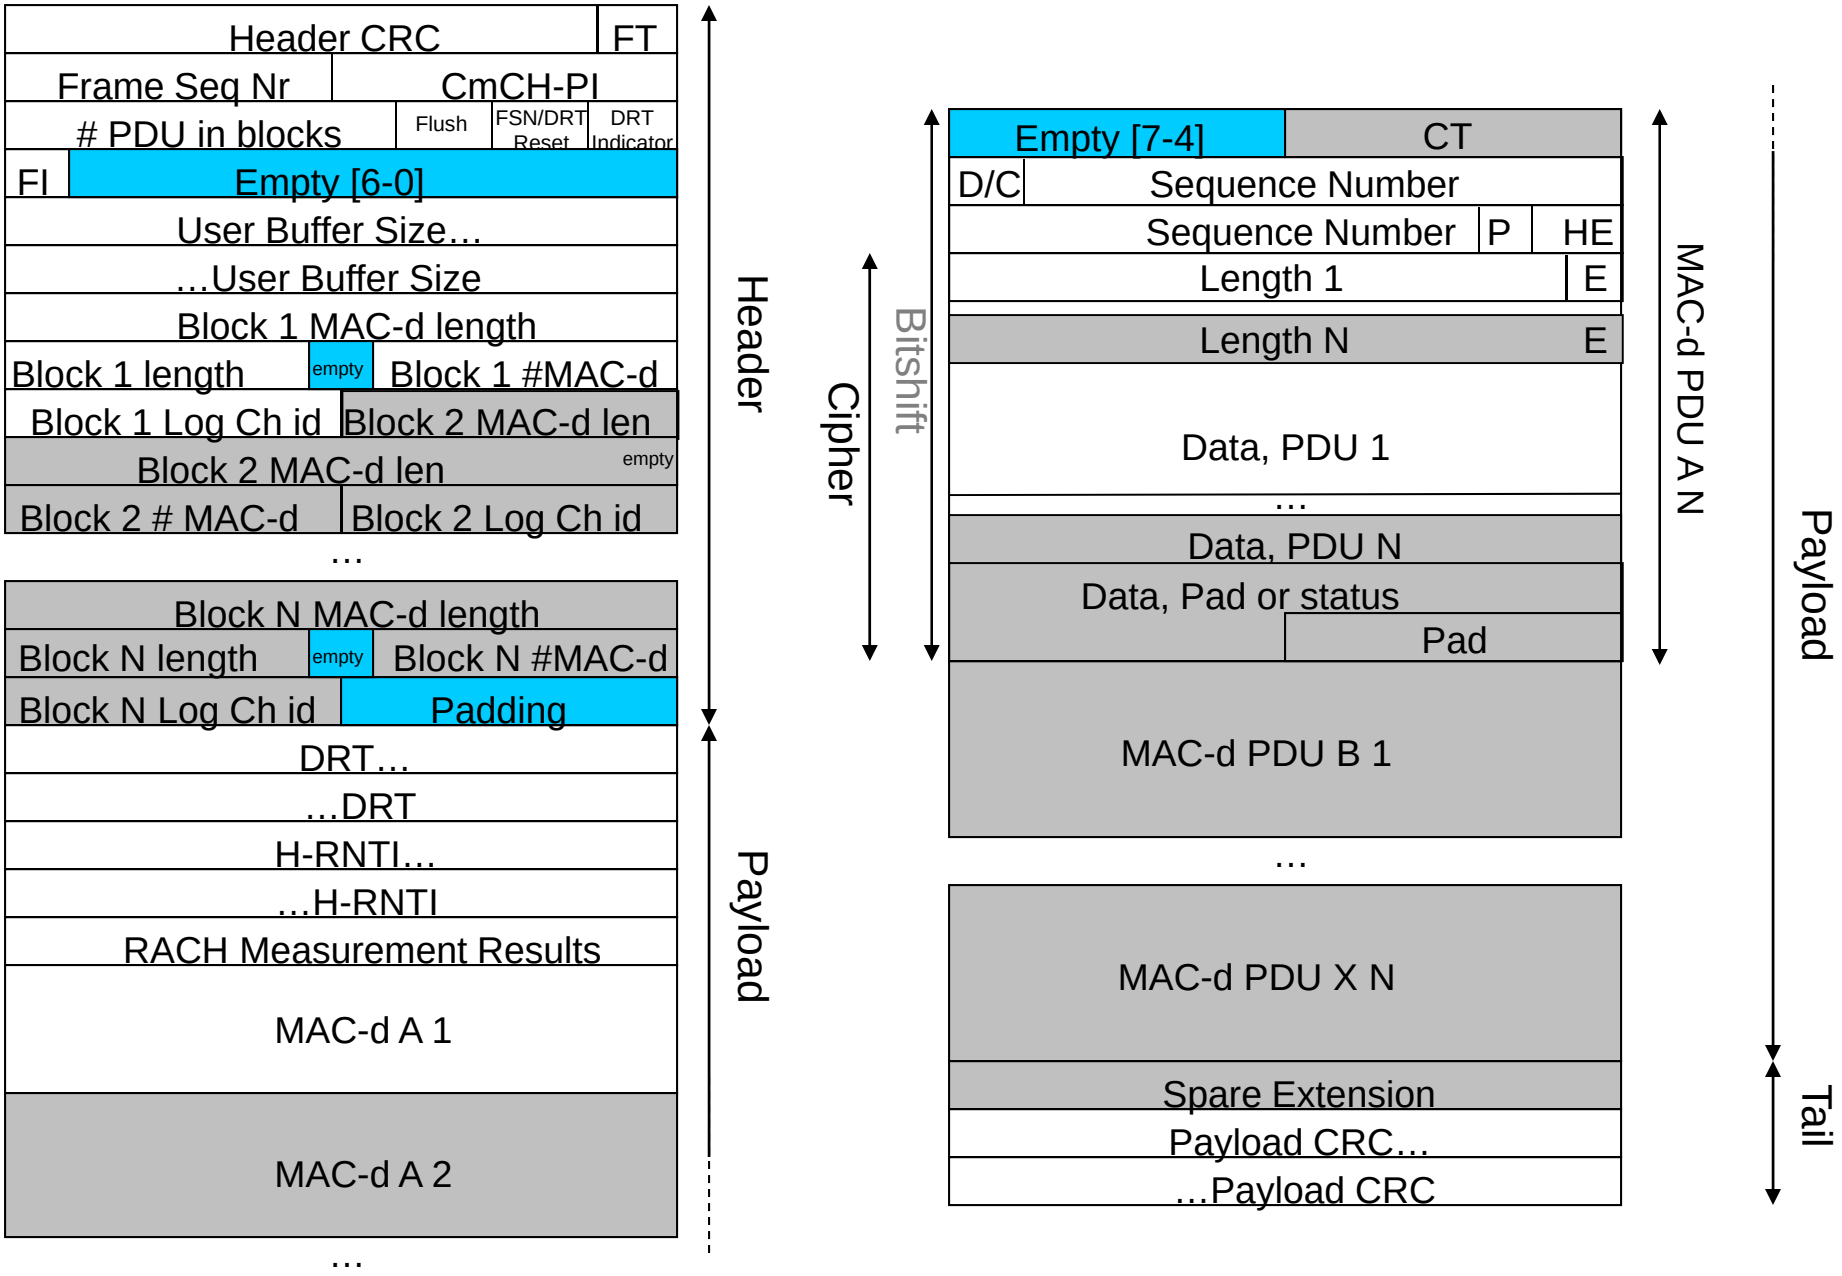

HSPA: AM RLC/MAC/FP Downlink

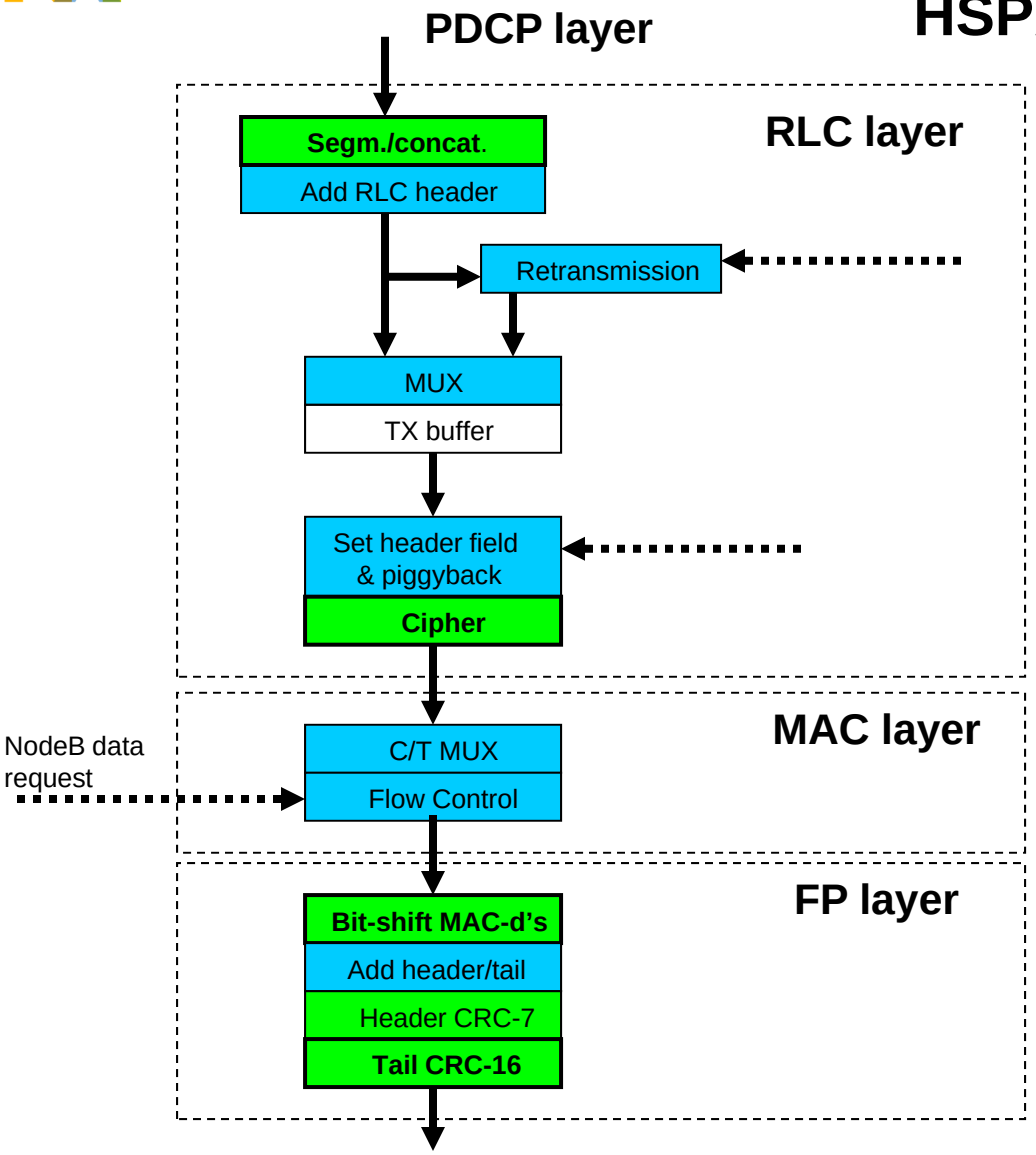

Bold frame indicates intensive work.

Blue = Core
Green = Accelerator

Features for mobile broadband:

- Cipher blocks of packets
- Auto-update sequence number
- Bit-shifting
- CRC-7 / CRC-11 / CRC-16
- Dual CRC (header/payload)

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor

RLC / MAC / FP flow on P4080: Example A

- User data assembled in memory, from GTP-U.

- When MAC flow control indicates that we can send data:

Core:

1. Checks retransmission buffer.
2. Checks status transmission buffer.
3. Creates RLC headers with status.
4. Checks if C/T MUX is used.
5. Adds pointers to segments and RLC headers.
6. Concatenation/piggyback/pad. of last user-data packet
7. Saves pointers for retransmission buffer
8. Creates FP header/tail.

SEC: segments, adds header, bitshift and cipher. Adds FP header/tail and computes CRC-7 and CRC-16.

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor ™

Type-1 MAC-d processing

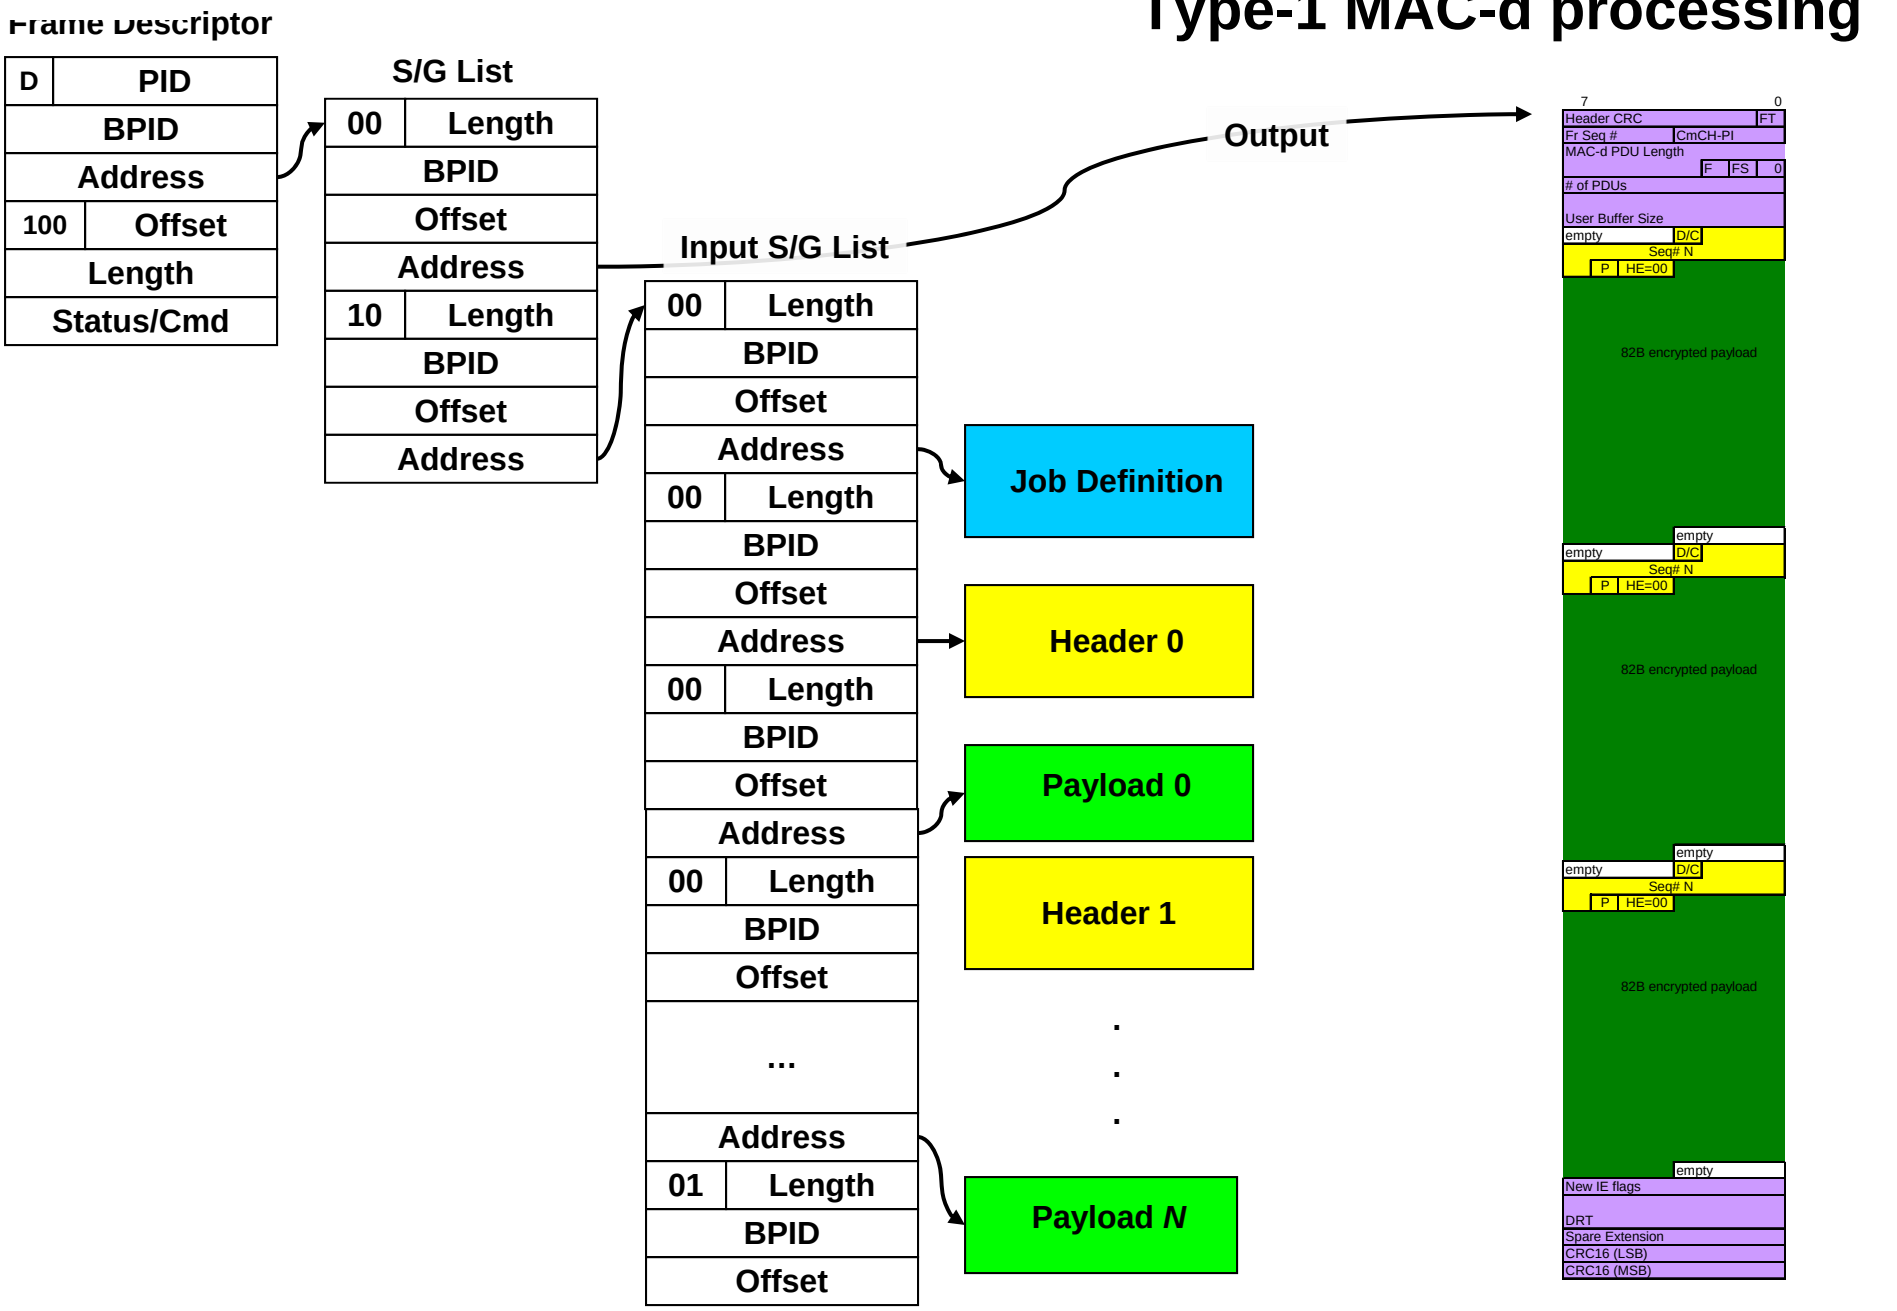

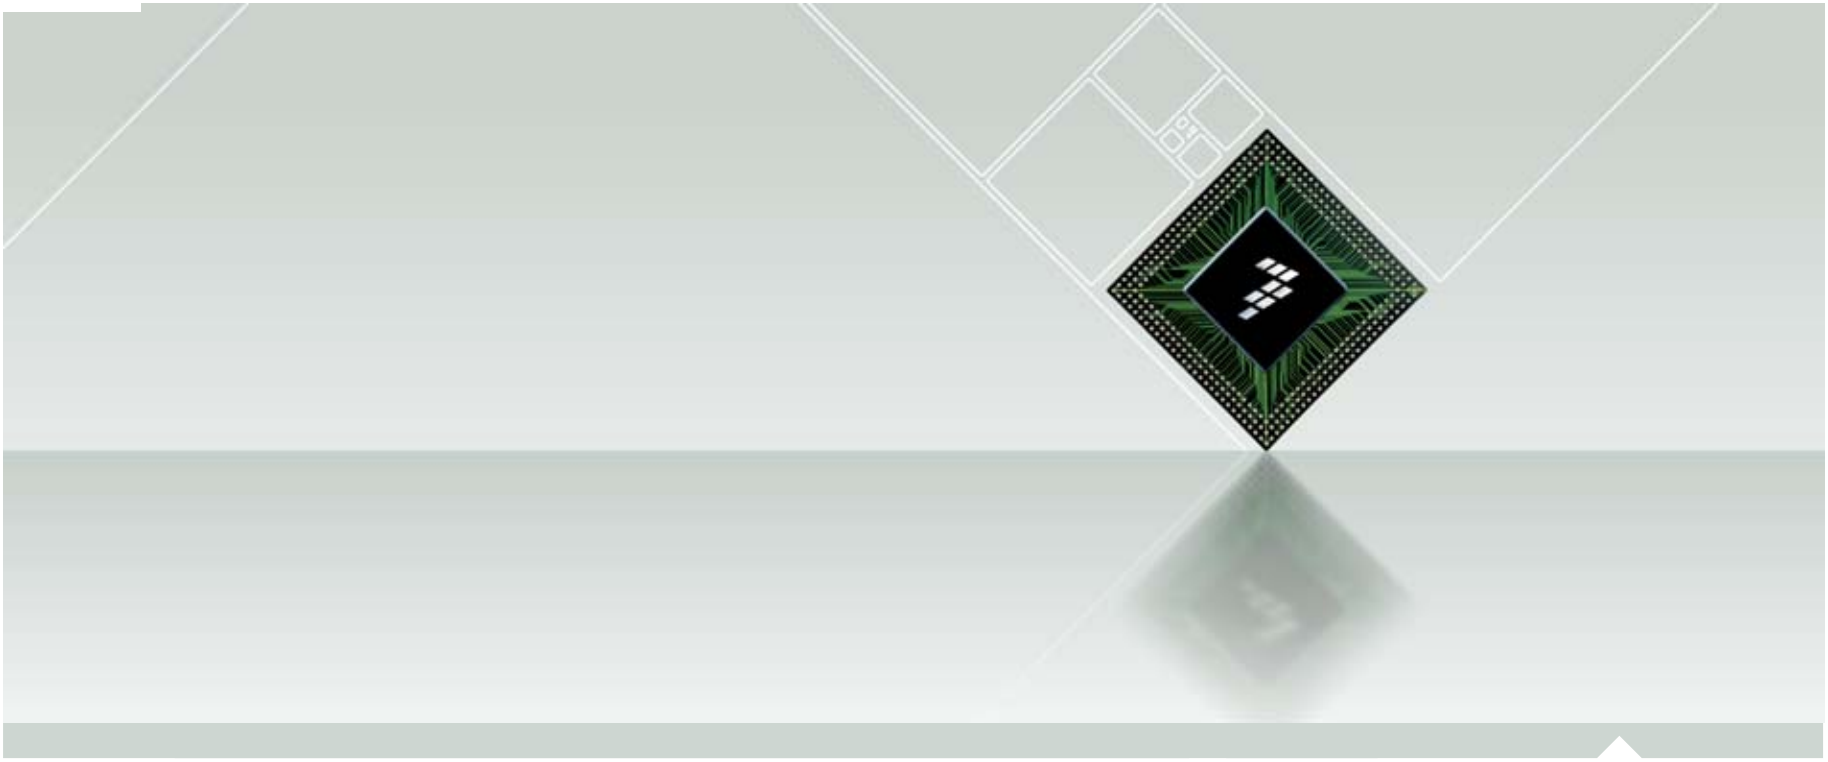

DEMO

Freescale, the Freescale logo, AltiVec, C-5, CodeTEST, CodeWarrior, ColdFire, C-Ware, mobileGT, PowerQUICC, StarCore, and Symphony are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off. BeeKit, BeeStack, CoreNet, the Energy Efficient Solutions logo, Flexis, MXC, Platform in a Package, Processor Expert, QorIQ, QUICC Engine, SMARTMOS, TurboLink and VortiQa are trademarks of Freescale Semiconductor, Inc. All other product or service names are the property of their respective owners. © 2010 Freescale Semiconductor, Inc.

freescale semiconductor


freescale™
semiconductor