TC Offloads on DPAA Devices

Mădălin Bucur, NXP

Netdev 2.1

Montreal

April 7, 2017

drivers/net/ethernet/freescale?/dpaa

DPAA 1.x

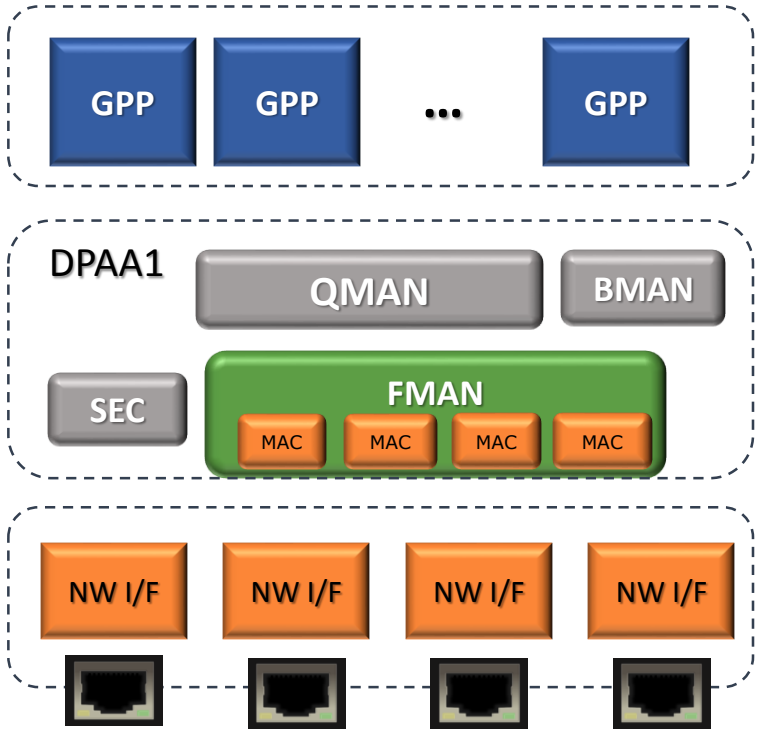

From 2008 on PowerPC
From 2015 on ARM

DPAA 2.x

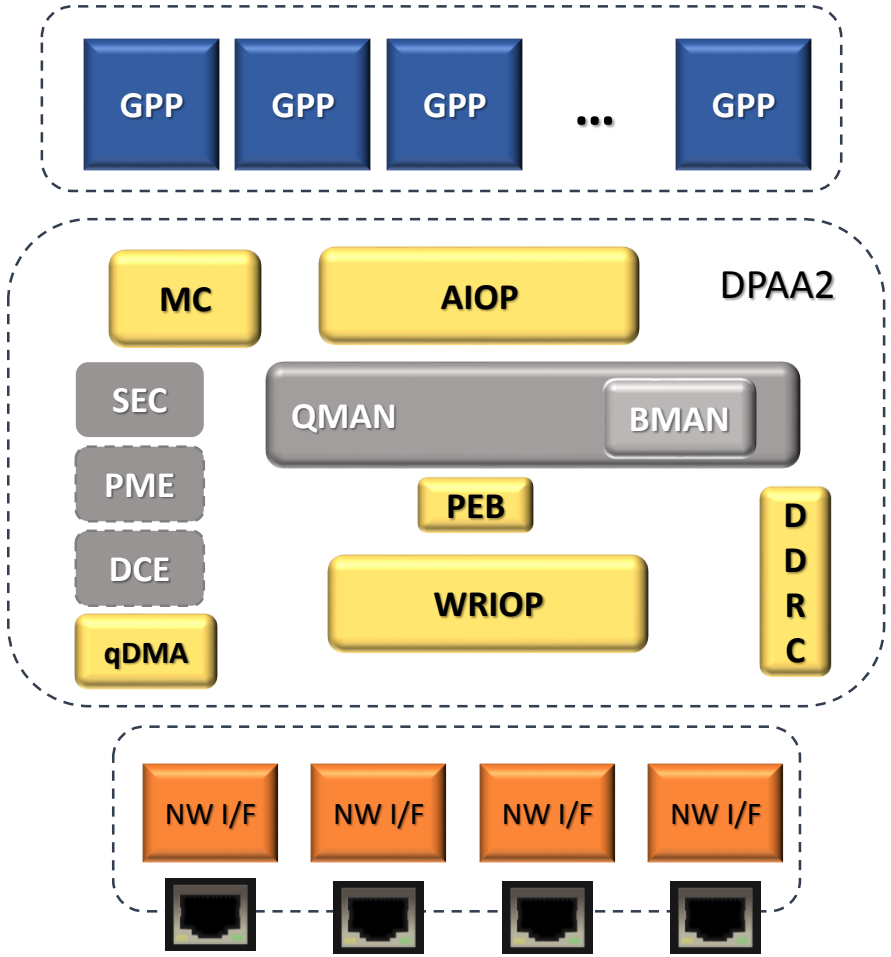

DPAA - Present

- DPAA 1.x PowerPC support in vanilla kernel
- DPAA 1.x ARM support in review (QMan, device tree)
- DPAA 2.x ARM support in preparation

- Supporting Rx & Tx checksum offload

- Supporting MQPRIO Offload
  - 4 priority levels
  - mapped to HW prioritized workqueues

DPAA 1.x

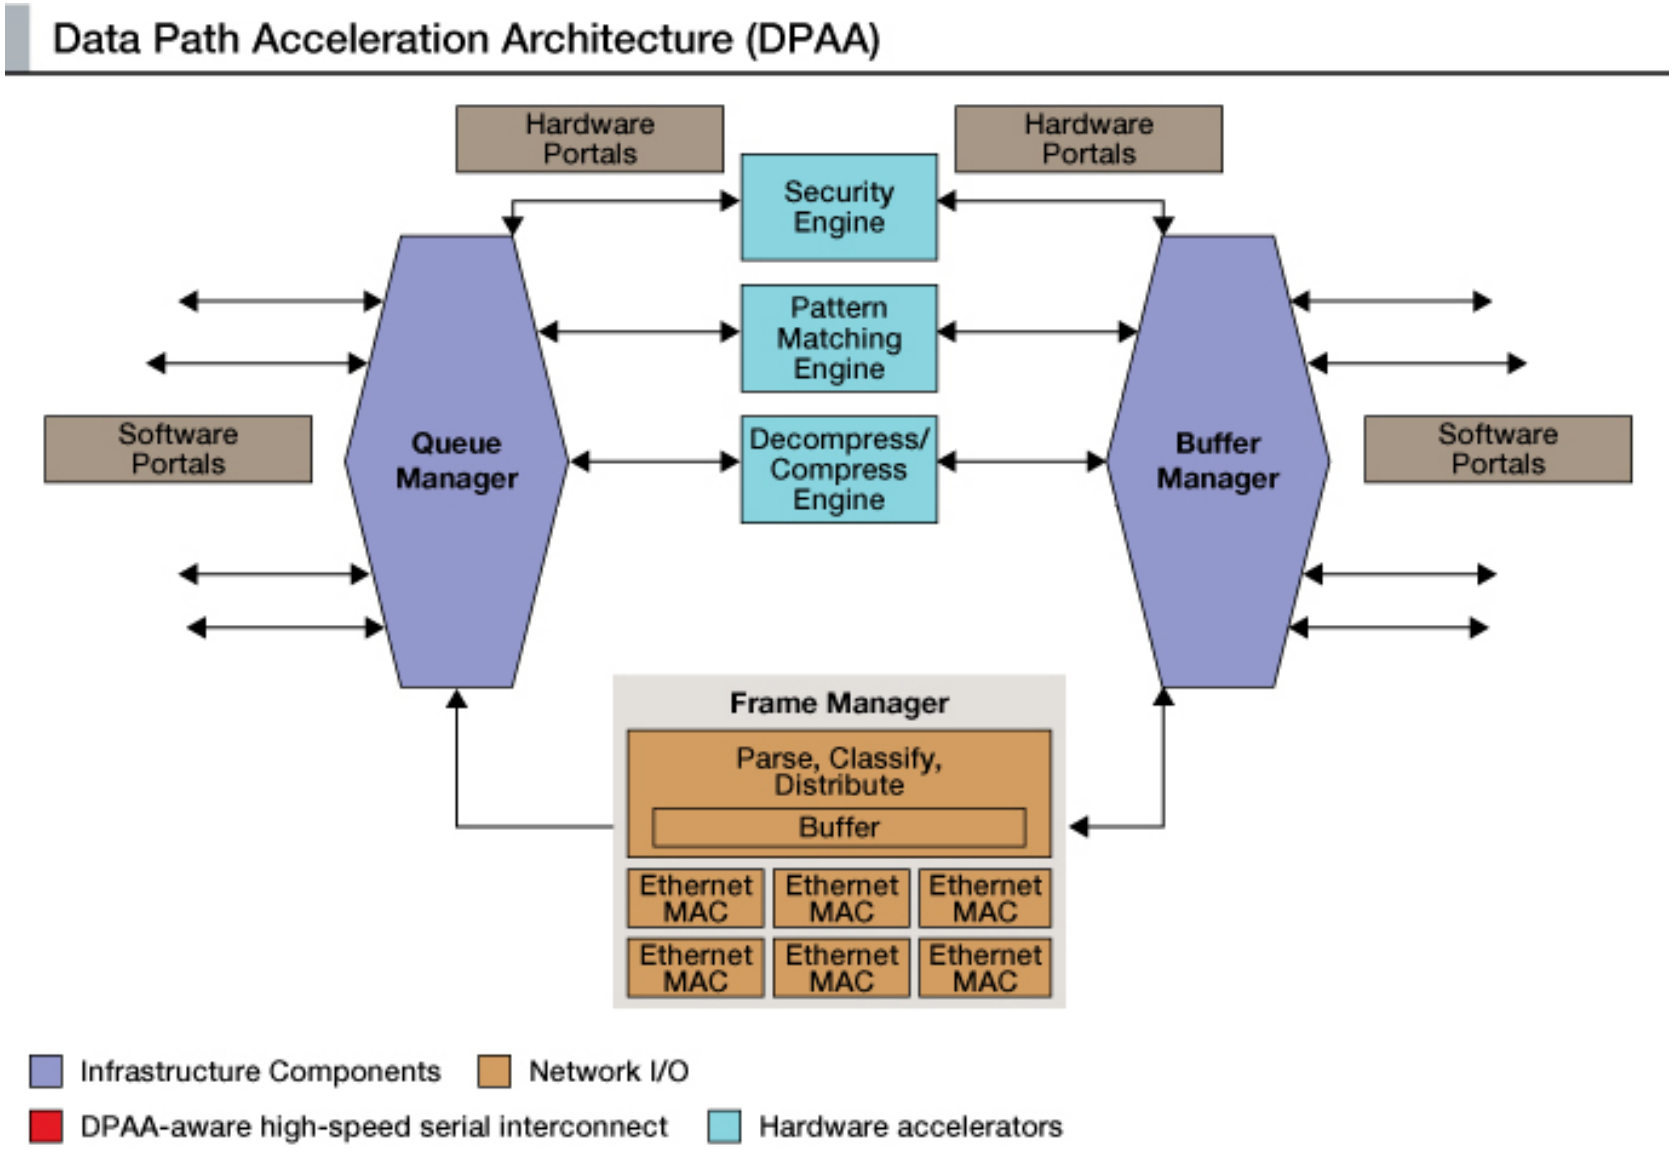
Data Path Acceleration Architecture (DPAA)
Hardware Portals
Hardware Portals
Security Engine
Pattern Matching Engine
Decompress/ Compress Engine
Queue Manager
Buffer Manager
Software Portals
Software Portals
Frame Manager
Parse, Classify, Distribute
Buffer
Ethernet MAC
Ethernet MAC
Ethernet MAC
Ethernet MAC
Ethernet MAC
Ethernet MAC
Infrastructure Components
Network I/O
DPAA-aware high-speed serial interconnect
Hardware accelerators

DPAA 1.x zoom in: Frame Manager

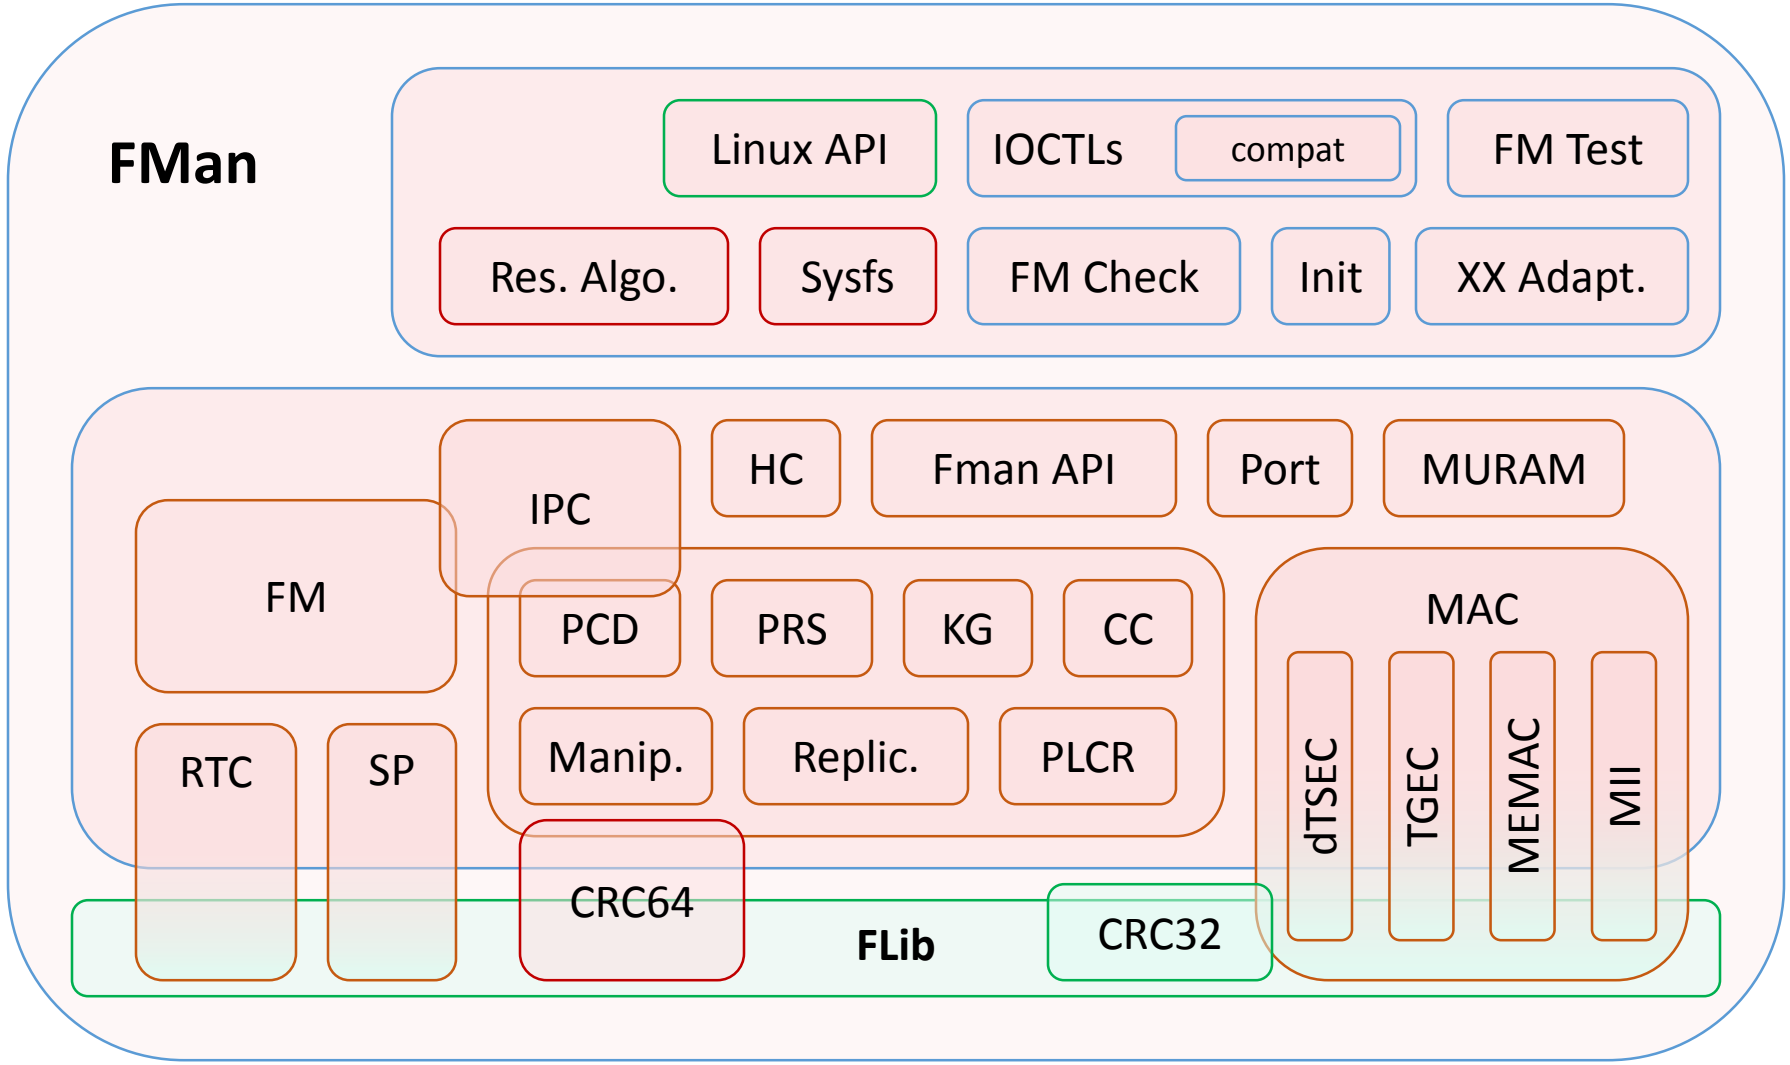

DPAA QoS Support

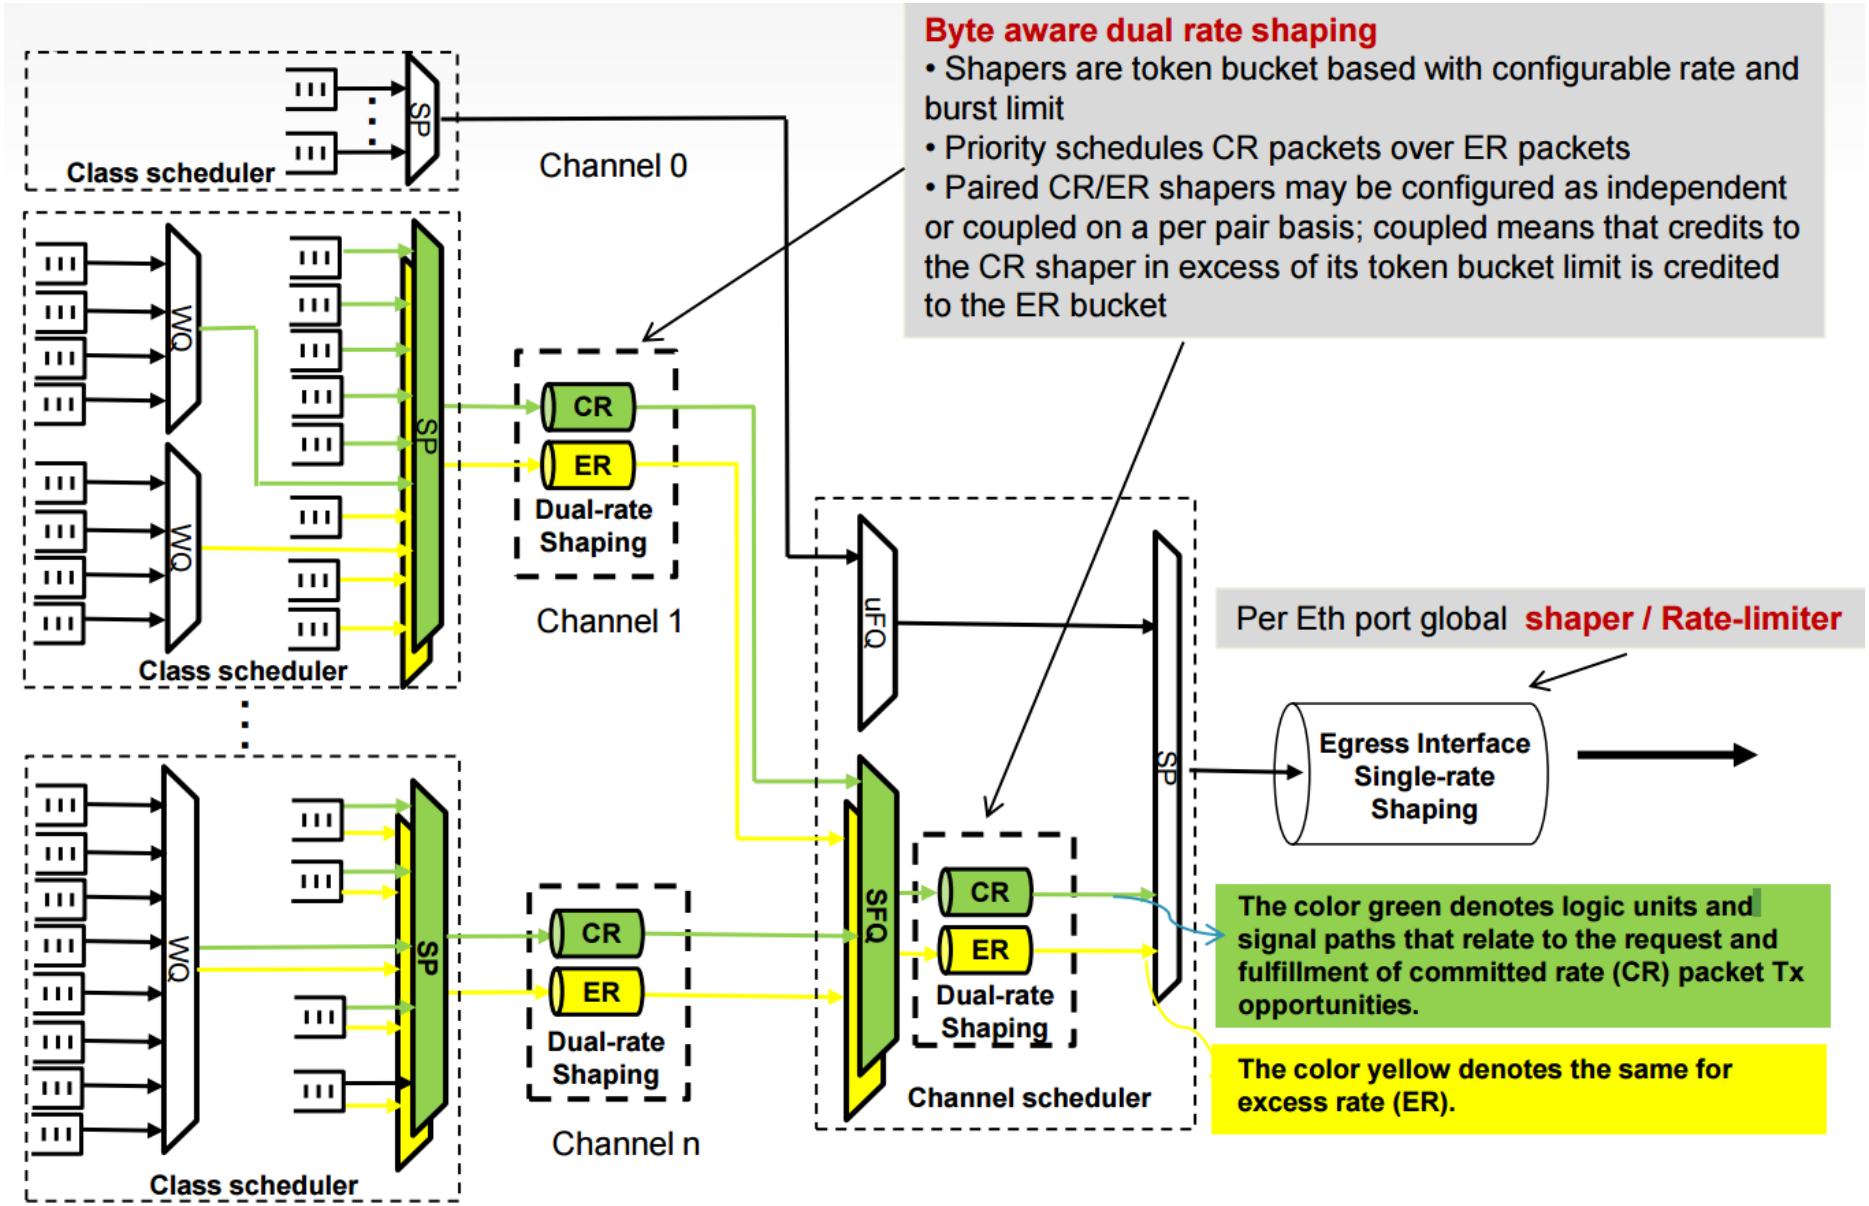

DPAA - Future

- Offload TC filters in HW
  - Flower
    - Maps well to HW abilities
    - Likely first priority
  - U32
    - Supported by HW as well
- Rx hash for CPU flow affinity
- Debug capabilities

- Offload shaping in HW
  - TBF
    - HW supports shaping per interface
  - HTB
    - HW also supports hierarchical shaping
- Expand coverage of HW offerings
  - Within a traffic class, HW supports:
    - Strict Priority
    - Weighted Fair Queuing
    - RED, WRED

Challenges and future research areas

- HW features and limitations
  - How to expose them to the user?
  - Configuring HW features for a specific scenario
    - DCB?
    - Devlink?
- HW to SW TC objects mapping
  - Adapt/extend the existing ones
  - Add new Qdiscs that mirror the custom HW
- Mixing SW and HW offloads

Questions?

madalin.bucur@nxp.com