You make **possible**

NCS5500

Deepdive in the Merchant Silicon High-end SP Routers

Nicolas FEVRIER
@CiscoIOSXR

BRKARC-3000

Cisco live!
June 9-13, 2019 · San Diego, CA

#CLUS

What We Hope To Achieve With This Session

- For a first approach
  - Getting familiar with the NCS5500 portfolio
  - Understand the implementation differences compared to traditional XR products (Buffering, Resource Management, …)
- For the experienced
  - Introducing the new platforms
  - Digging deeper in the architecture
- Some tips
- **For Reference**

Agenda

- Products Portfolio
- Fixed / Modular Platforms / Optics
- VOQ/FMQ and Life of a Packet
- Memory Structure
- Features: ACL / QoS
- Gotchas

Cisco Webex Teams

Questions?

Use Cisco Webex Teams to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

Webex Teams will be moderated by the speaker until June 16, 2019.

@CiscoIOSXR

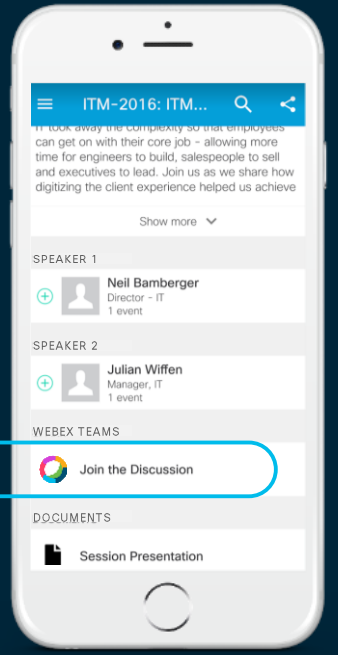

cs.co/ciscolivebot#BRKARC-3000

Introduction

You make networking **possible**

Cisco live!

IOS XR Routing Products

Cisco ASR9000, CRS & NCS6000

vRR/vPE
Universal Virtual Forwarder

Cisco IOS-XRv9000

Cisco NCS5500, NCS500
and NCS5000

Custom

Virtual

Merchant

Cisco XR Software

Network Convergence System

Vast Product Line

Platform / Series
NCS 520
NCS 540
NCS 560
NCS 1000
NCS 2000
NCS 4000
NCS 4200
NCS 5000
NCS 5500
NCS 6000

NCS...

At a Glance

Platform / Series	Role
NCS 520	Ethernet Access Device (IOS XE)
NCS 540	Access Router
NCS 560	Aggregation Router
NCS 1000	DCI / IP-DWDM
NCS 2000	Packet Optical DWDM / TDM to IP / CEM
NCS 4000	
NCS 4200	
NCS 5000	Top of Rack Router
NCS 5500	Core, Edge, Agg, Peering Router
NCS 6000	Core Router

NCS5500 and NCS5000

Two Very Different Platforms

- Both based on Merchant Silicon forwarding ASICs and running IOS XR 64-bit
- Still they are very different in nature and in position in networks
- NCS5500
  - High scale routing and features
  - Exists in Fixed and Modular form factors (Fabric Engine)
  - Hybrid Architecture with Deep Buffers
- NCS5000
  - Lower scale and small buffers
  - No Chassis with Fabric Engine
  - Cost optimized
  - Can be used as a nV Satellite for ASR9000 and NCS6000

NCS5500 and NCS500

Much Closer Platforms

- Both based on same Merchant Silicon ASIC family (DNX)
- A lot of commonalities in the architecture and feature support
- Some difference in scale and features related to specific additional hardware parts
- NCS540
  - based on Qumran-AX (lower scale)
- NCS560
  - Based on Qumran-MX with OP eTCAM (2nd Generation eTCAM)

But What is Merchant? Really...

You make networking **possible**

Components

Merchant/Commodity, Proprietary, Custom

- **Merchant**
  - Not designed by a system vendor
  - Available on the open market to any system vendor or network operator
- **Proprietary**
  - Designed or acquired by a router vendor
  - Not available to others
- **Custom**
  - Designed in concert with a specific router in mind
  - Usually proprietary but may be merchant with extensions

Custom and Merchant

Cisco Platforms Internal Components

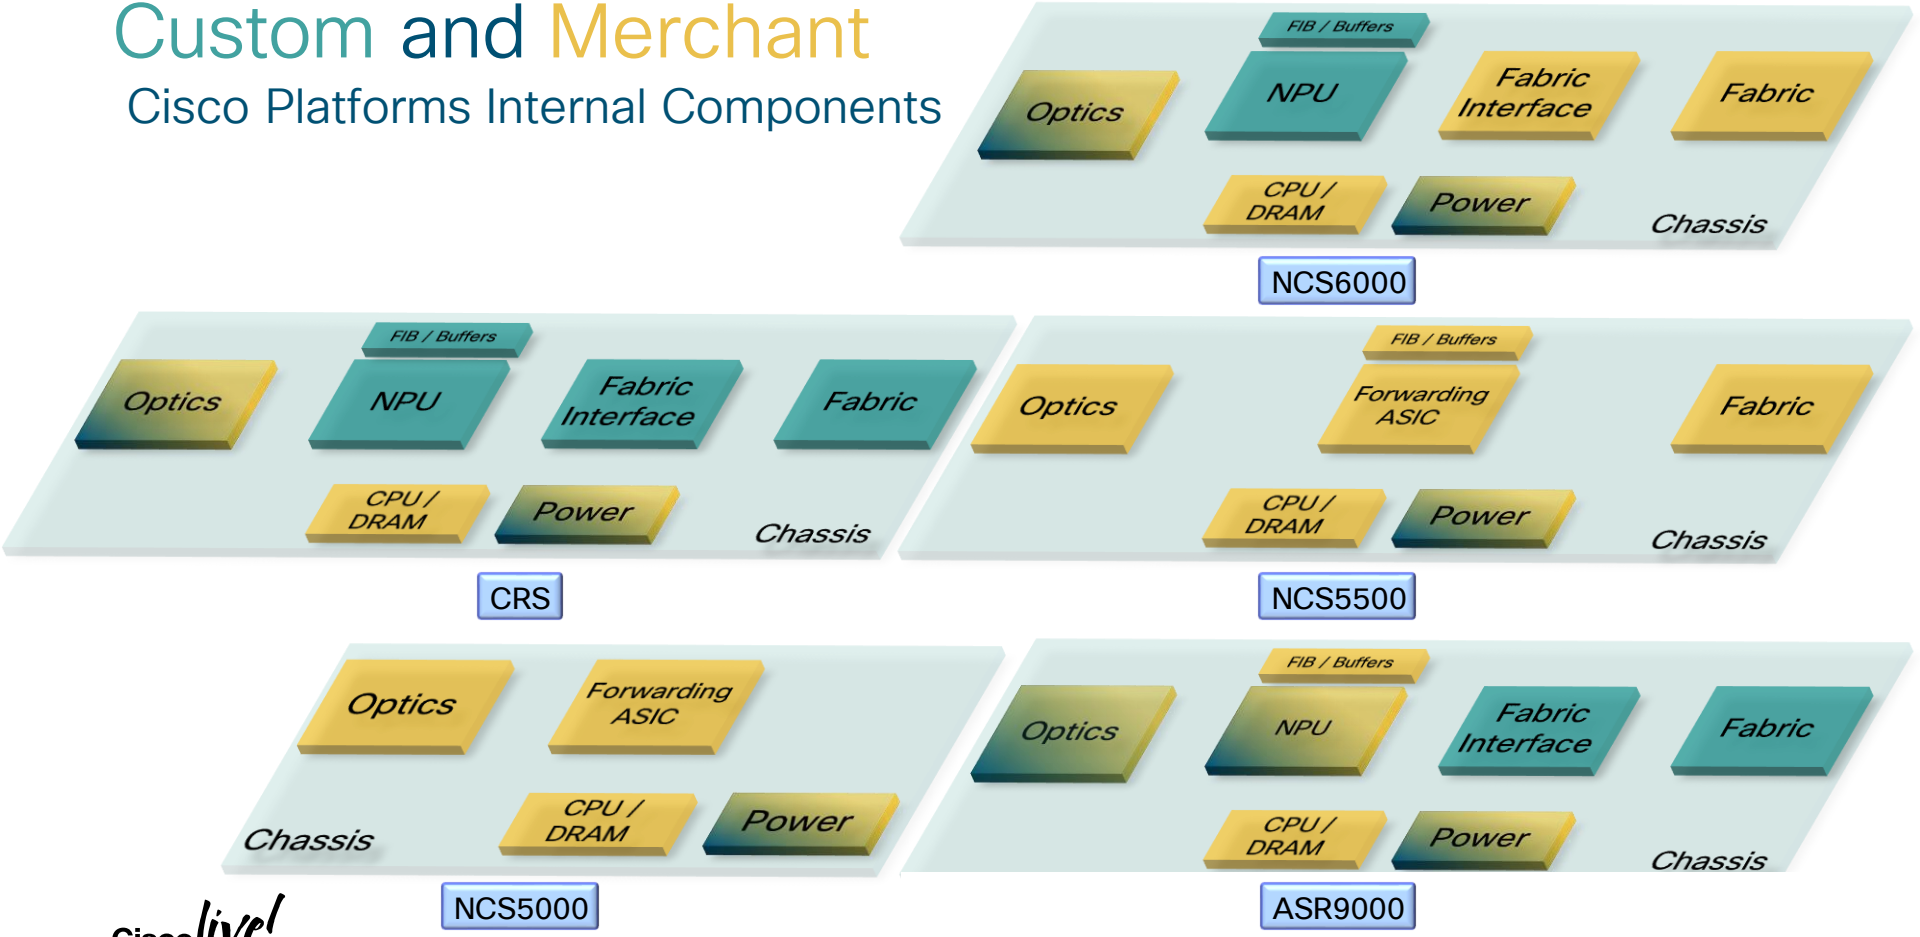

NCS5500 Portfolio

You make networking **possible**

Cisco live!

NCS5500 Products Family

- 13x Fixed Routers
  - NCS-5501(-SE)
  - NCS-5502(-SE)
  - NCS-55A1-24H
  - NCS-55A1-36H(-SE)-S
  - NCS-55A2-MOD(-SE)
  - NCS-55A2-MOD-HD(-SE)
  - NCS-55A1-48Q6H
  - NCS-55A1-24Q6H-S
- 3x Modular Routers
  - NCS-5504
  - NCS-5508
  - NCS-5516
- 11x Line Cards
  - NC55-36X100G
  - NC55-36X100G-S
  - NC55-24X100G-SE
  - NC55-18H18F
  - NC55-24H12F-SE
  - NC55-6x200-DWDM-S
  - NC55-36X100G-A-SE
  - NC55-MOD-A(-SE)-S
  - NC55-24D
  - NC55-18D12TH-SE

NCS5500 Products Family

Base and Scale Concept

- Both exist for modular and fixed systems
- **Base**
  - On-chip FIB and small TCAM for ACLs / QoS
- **Scale** (-SE) have increased FIB and ACL
  - off-chip TCAM
- External TCAM is a shared resource
  - IPv4 & IPv6 route scale
  - Ingress ACL / QoS matching scale

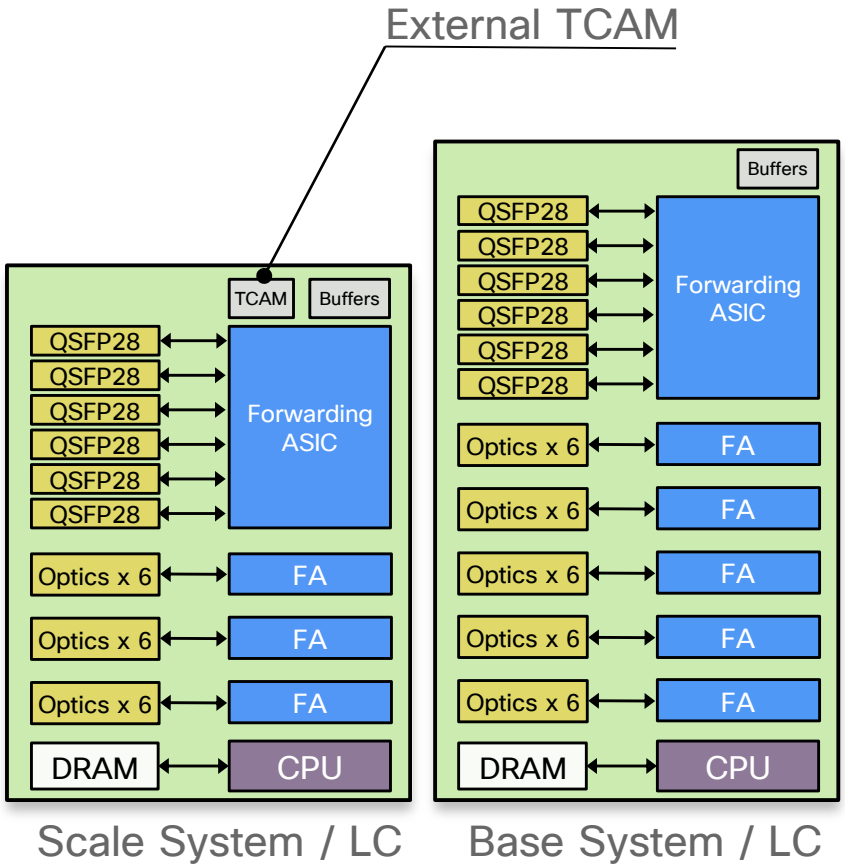

Base / Scale

So, it's like -TR/-SE on ASR9000?

- Yes and No
- On both platforms: -SE will support more features with higher scale
- But scale will be different
  - ASR9000: different QoS capability (because higher classifier scale)
  - NCS5500: different FIB scale (because TCAM is used to store routing information, not only classifiers)

NCS5500 Basics Concepts on NPU

You make networking **possible**

Cisco live!

Simplification is Key

Fewer Components: Cost Optimization and Lower Power Consumption

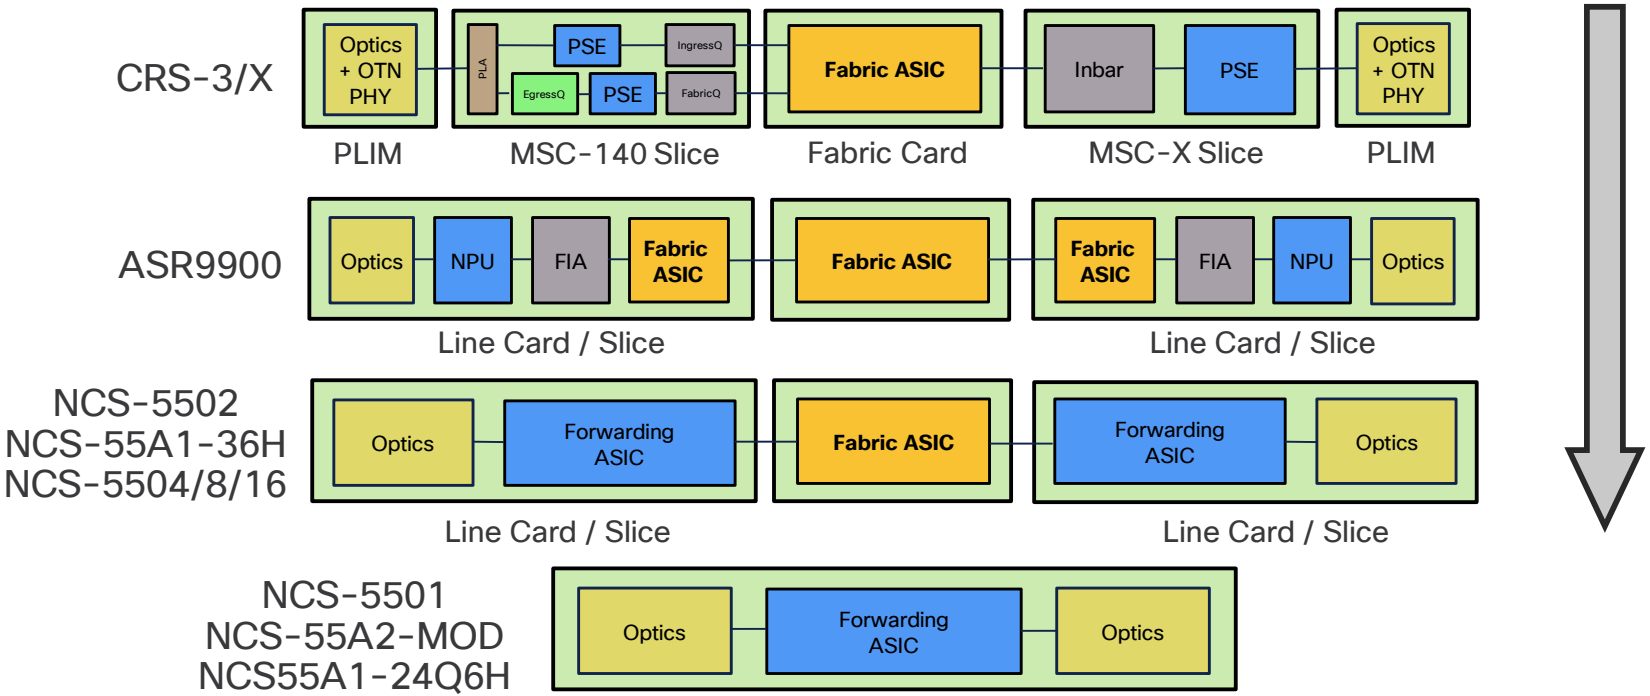

DNX Forwarding ASIC

- Broadcom StrataDNX Family
- From 2009 Dune Networks acquisition
- Standalone (SOC) or leaf-spine ASIC / Fabric Engine

DNX Forwarding ASIC in NCS5500

Standalone

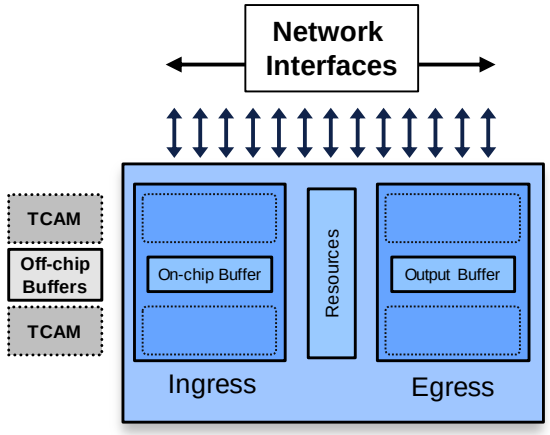

DNX Forwarding ASIC in NCS5500

Leaf-spine ASIC / Fabric Engine

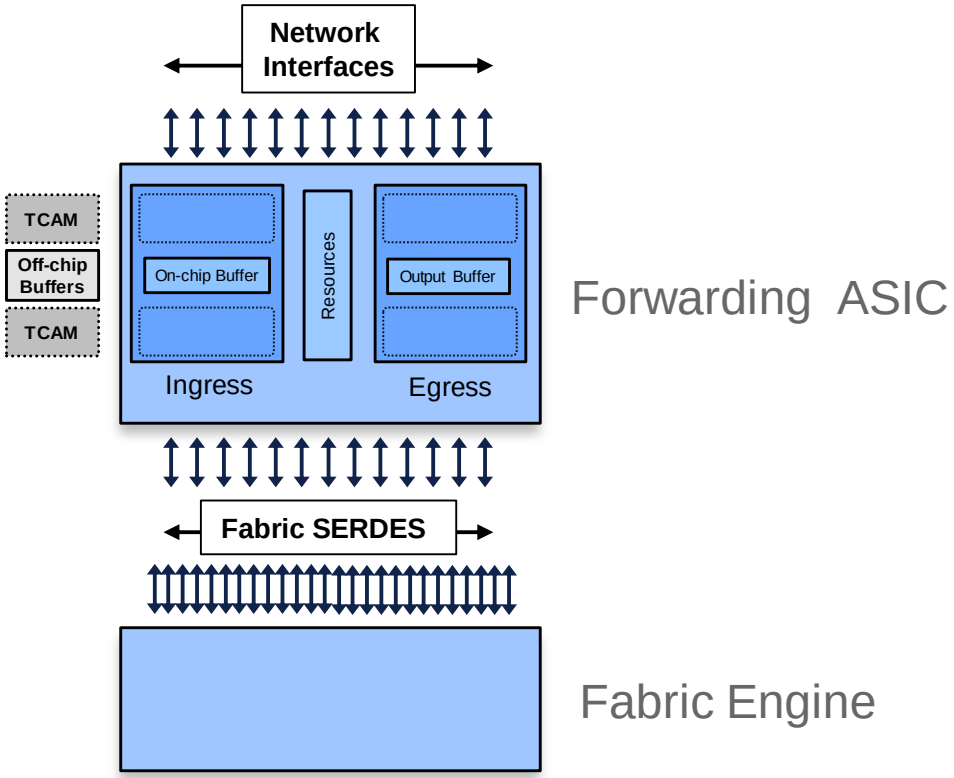

DNX Forwarding ASIC in NCS5500

Back to Back

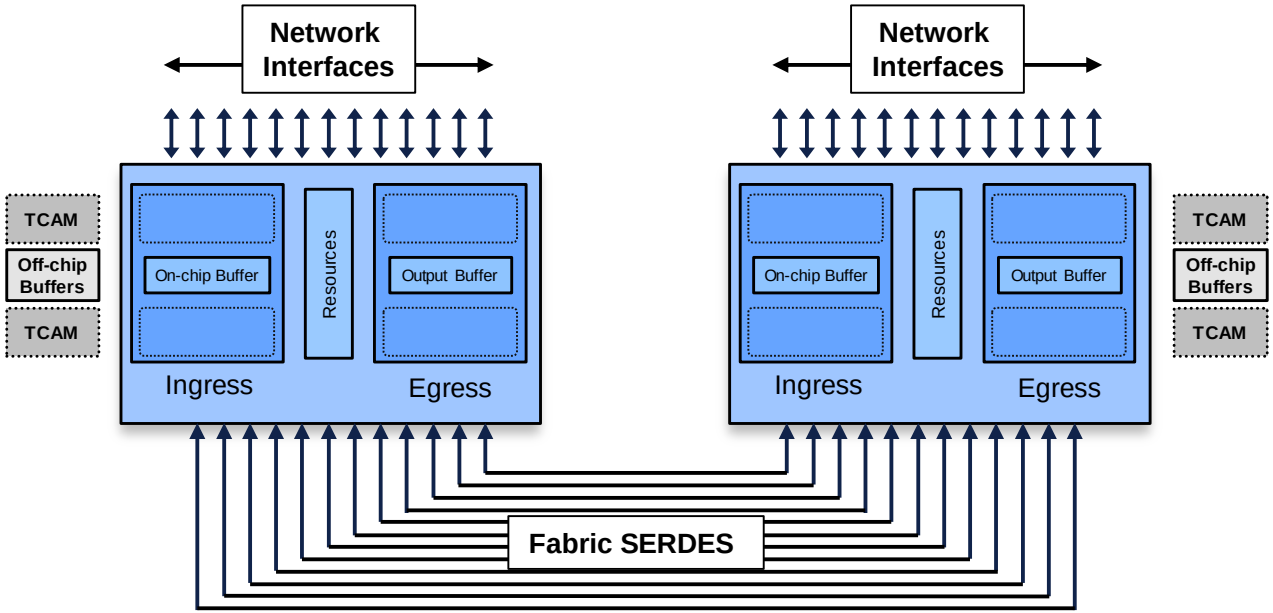

ASIC Architecture

RTC Scheduler or Pipeline?

- Run to Completion: many cores, each does everything for a packet
- Pipeline: many stages/block, each has a specialized role (NCS5500)

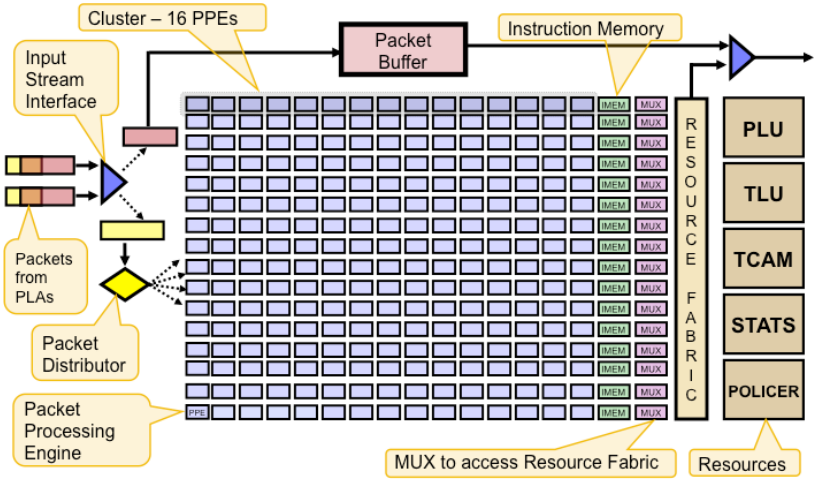

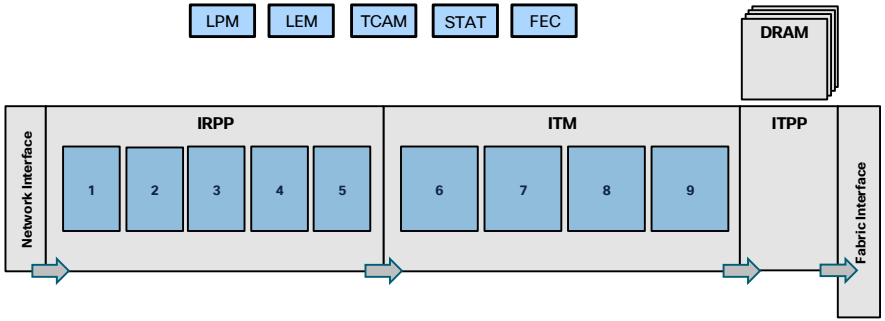

NCS5500 Series

NCS5500 Forwarding ASIC

Internal Components

- Integrated Forwarding and Fabric Interface
- 1 or 2 cores
- Separate ingress and egress Pipelines
- PP: Packet Processor
  - Lookup, features, …
- TM: Traffic Manager
  - QoS: WRED, hierarchical scheduling, shaping, policing

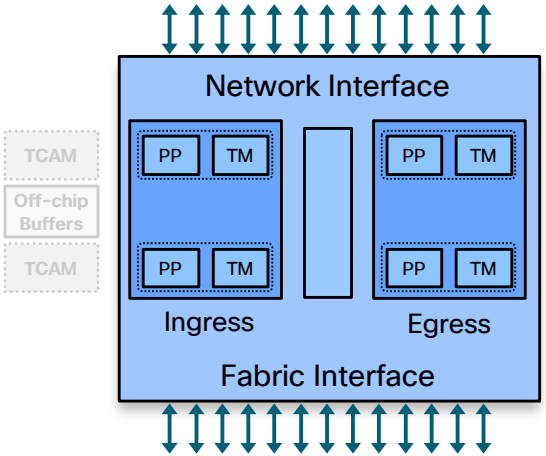

NCS5500 Forwarding ASIC

Internal Components

- Buffers
  - Used to store packets only
- On-chip resource
  - Small internal buffers
- Off-chip resource
  - Deep GDDR5 packet buffers external buffers
- Not optional

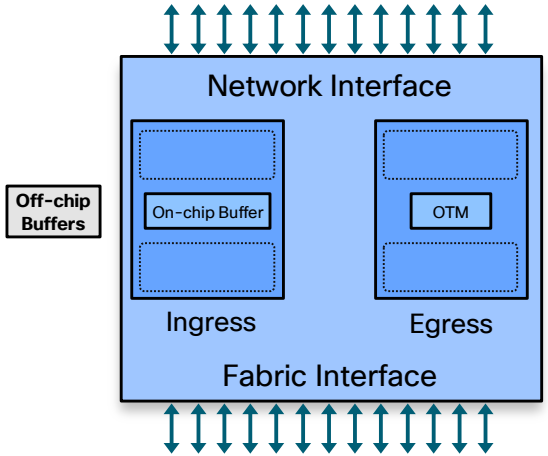

NCS5500 Forwarding ASIC

Internal Components

- "Information" memories
- On-chip databases used for
  - Route table: prefixes / nexthop / load-balancing
  - Classifiers / filters
  - Statistics
- Off-chip resources
  - Optional TCAMs for route/ACL scale

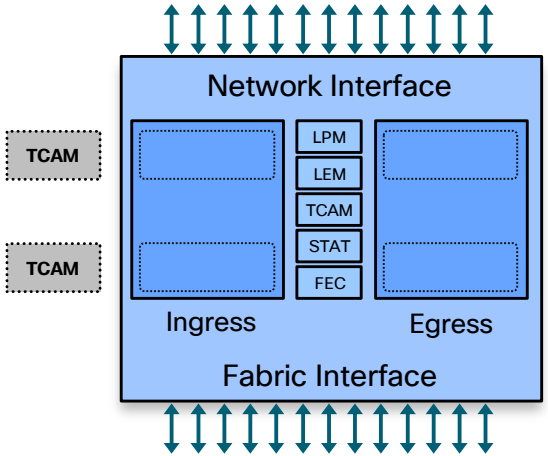

NCS5500 Forwarding ASICs

Jericho / Jericho+ / Jericho+ w/ Large LPM

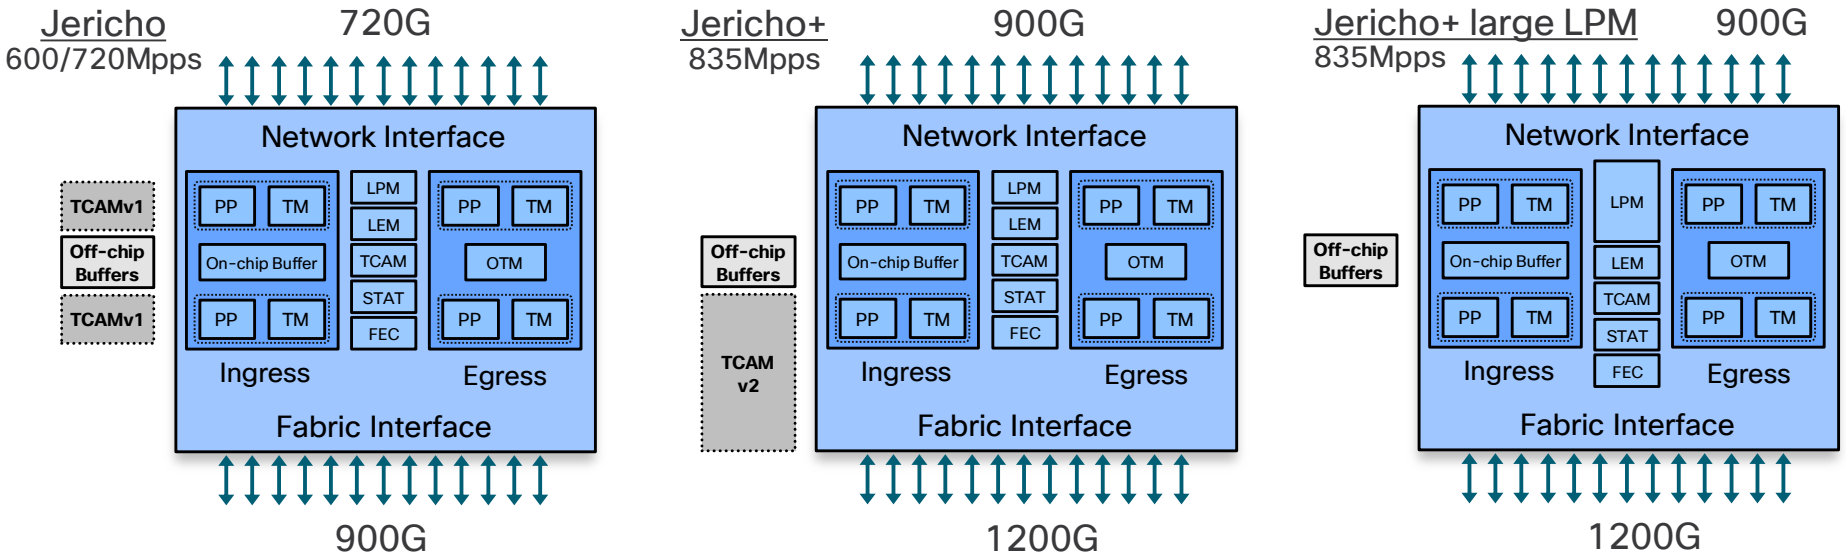

NCS5500 and NCS500 Forwarding ASICs

Qumran-MX / Qumran-AX

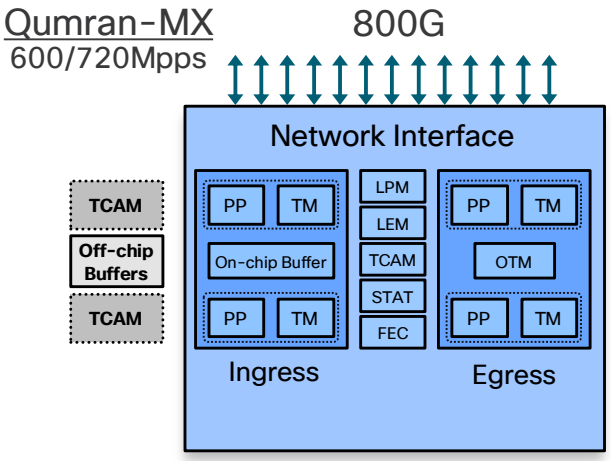

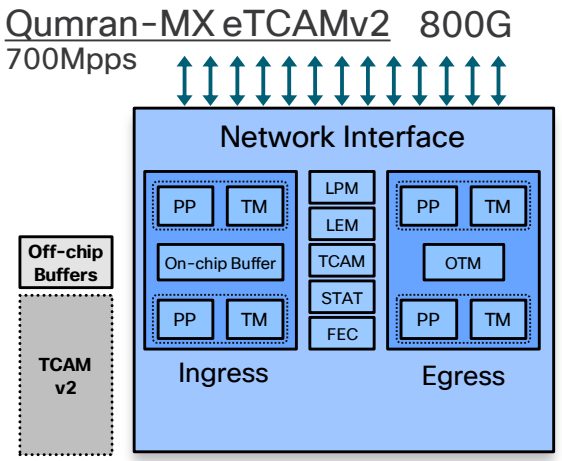

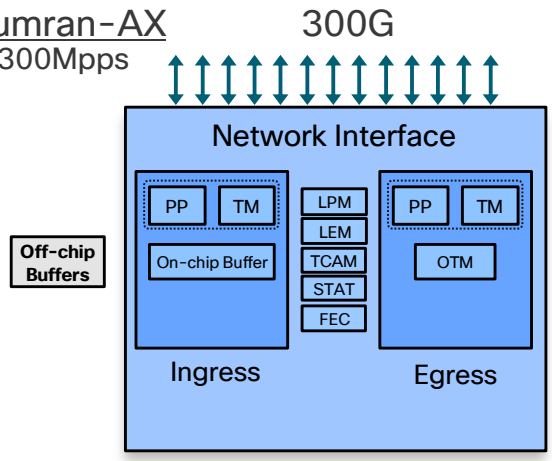

For Reference

In Summary

NCS5000
XGS ASICs

NCS5500
J/J+/Q-MX

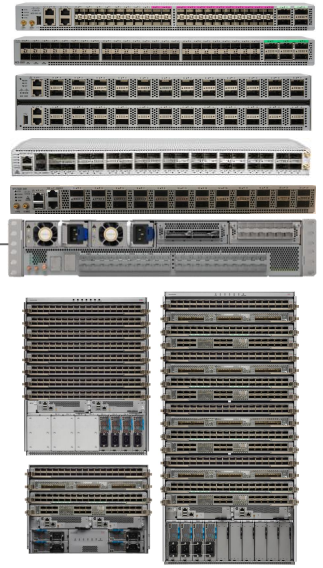

NCS540
Q-AX

NCS560
Q-MX

NCS5500 Forwarding ASIC

J/J+/Q-MX Pipeline Architecture

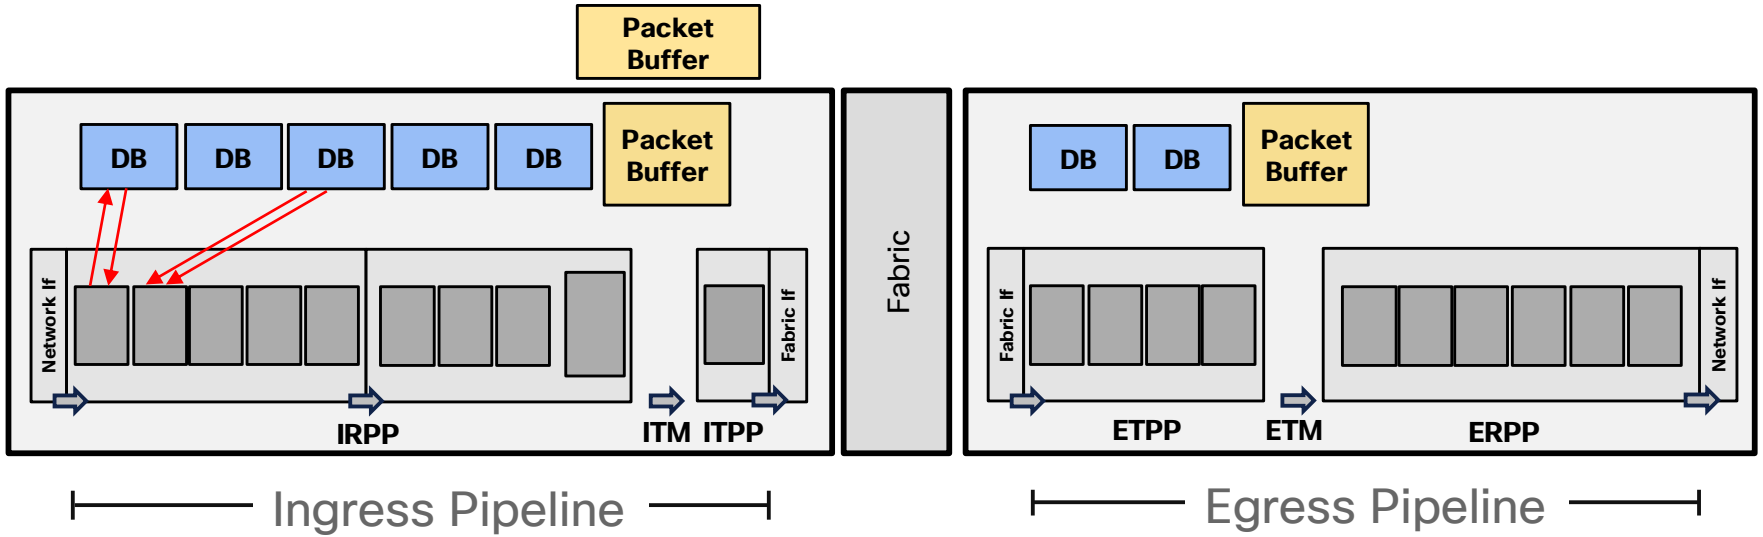

NCS5500 Forwarding ASIC

Pipeline Architecture

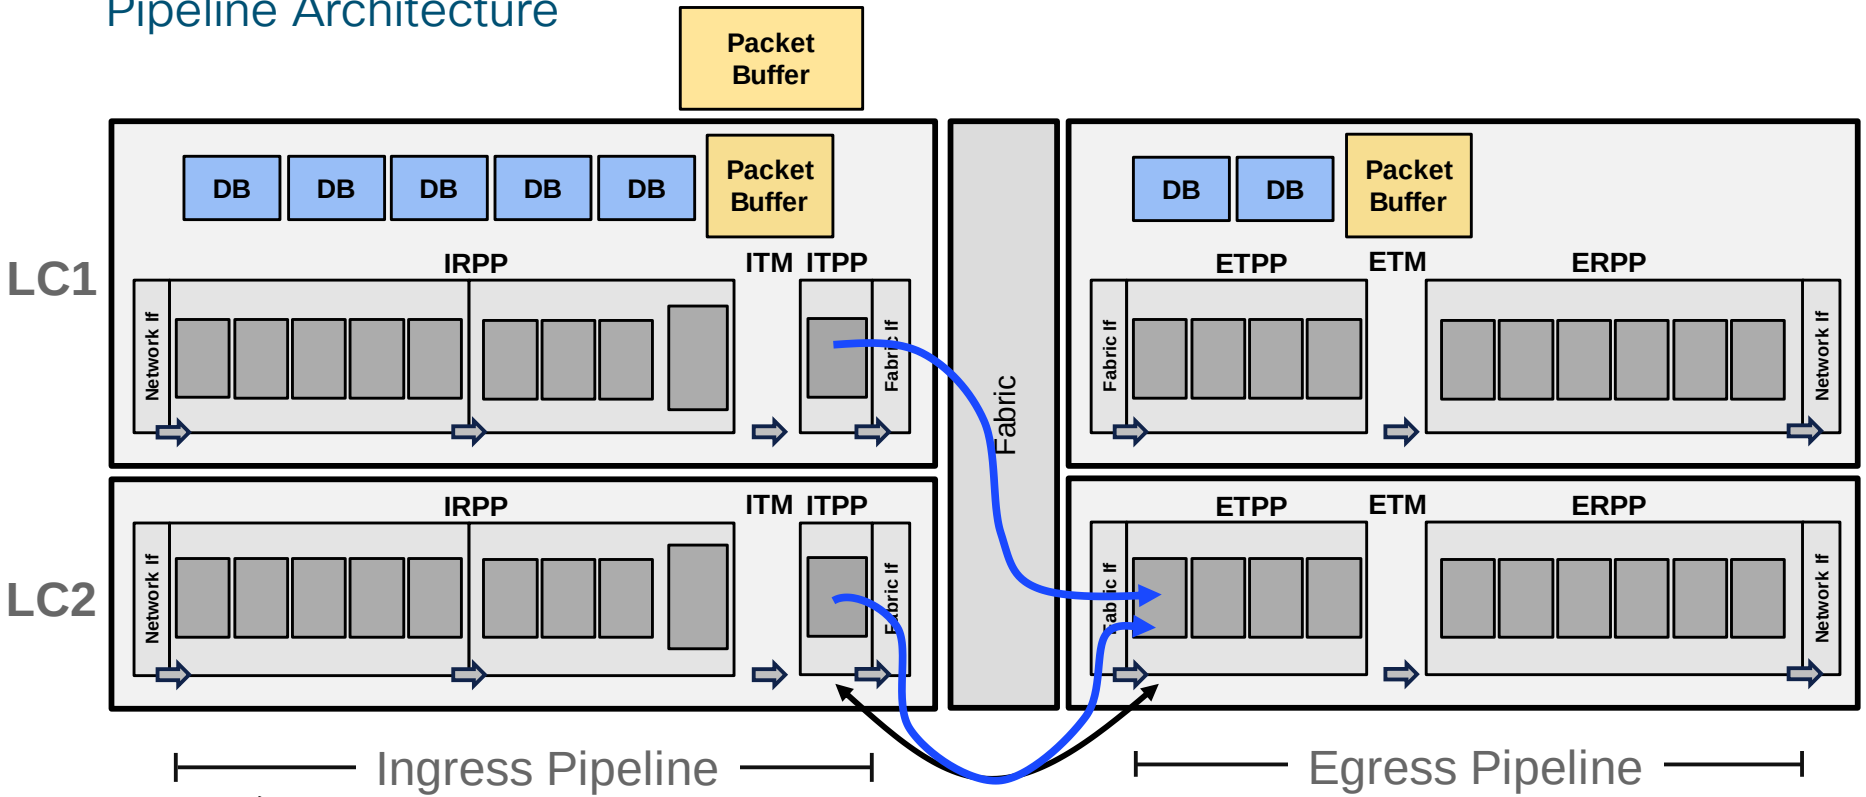

NCS5500 Forwarding ASIC

Pipeline Architecture

For Reference

```
RP/0/RP0/CPU0:5500#sh contr npu diag counters graphical instance 0 loc 0/1/CPU0

 Statistics Rack: 0, Slot: 1, Asic instance: 0

                                              |
                                              |           J E R I C H O   N E T W O R K   I N T E R F A C E
                                             \|/
+---------------------------------------------+--------------------------------------------+
|                                                                                         NBI
|  RX_TOTAL_BYTE_COUNTER                             = 0                                  |  TX_TOTAL_BYTE_COU
|  RX_TOTAL_PKT_COUNTER                              = 0                                  |  TX_TOTAL_PKT_COUN
|  RX_TOTAL_DROPPED_EOPS                             = 0                                  |
+---------------------------------------------+--------------------------------------------+
<SNIP>
+---------------------------------------------+--------------------------------------------+--------------------------------------------+
<SNIP>
+---------------------------------------------+--------------------------------------------+--------------------------------------------+
|                                                                                          |                                          FDA                                          |
|                                                                                          |  CELLS_IN_CNT_P1           = 0           |  CELLS_OUT_CNT_P1          = 0           |
|                                                                                          |  CELLS_IN_CNT_P2           = 22          |  CELLS_OUT_CNT_P2          = 20          |
+---------------------------------------------+--------------------------------------------|  CELLS_IN_CNT_P3           = 0           |  CELLS_OUT_CNT_P3          = 0           |
|                                            IPT                                           |  CELLS_IN_TDM_CNT          = 0           |  CELLS_OUT_TDM_CNT         = 0           |
|                                                                                          |  CELLS_IN_MESHMC_CNT       = 0           |  CELLS_OUT_MESHMC_CNT      = 0           |
|  EGQ_PKT_CNT                                       = 0                                 -->  CELLS_IN_IPT_CNT          = 0           |  CELLS_OUT_IPT_CNT         = 0           |
|  ENQ_PKT_CNT                                       = 0                                  |  EGQ_DROP_CNT                                      = 0                          |
|  FDT_PKT_CNT                                       = 0                                  |  EGQ_MESHMC_DROP_CNT                               = 0                          |
|  CRC_ERROR_CNT                                     = 0                                  |  EGQ_TDM_OVF_DROP_CNT                              = 0                          |
|  CFG_EVENT_CNT                                     = 0                                  |                                                                                   |
|  CFG_BYTE_CNT                                      = 0                                  |                                                                                   |
+---------------------------------------------+--------------------------------------------+--------------------------------------------+
|                                            FDT                                           |                                          FDR                                          |
|  IPT_DESC_CELL_COUNTER                             = 0                                  |  P1_CELL_IN_CNT                                    = 0                          |
|                                                                                          |  P3_CELL_IN_CNT                                    = 0                          |
|  TRANSMITTED_DATA_CELLS_COUNTER                    = 0                                  |  CELL_IN_CNT_TOTAL                                 = 22                         |
+---------------------------------------------+--------------------------------------------+--------------------------------------------+
                                              |                                                                    /|\
                                              |           J E R I C H O   F A B R I C   I N T E R F A C E          |
                                             \|/                                                                   |
```

Fixed Platforms

You make networking **possible**

Cisco live!

Naming Rules for Fixed Platforms

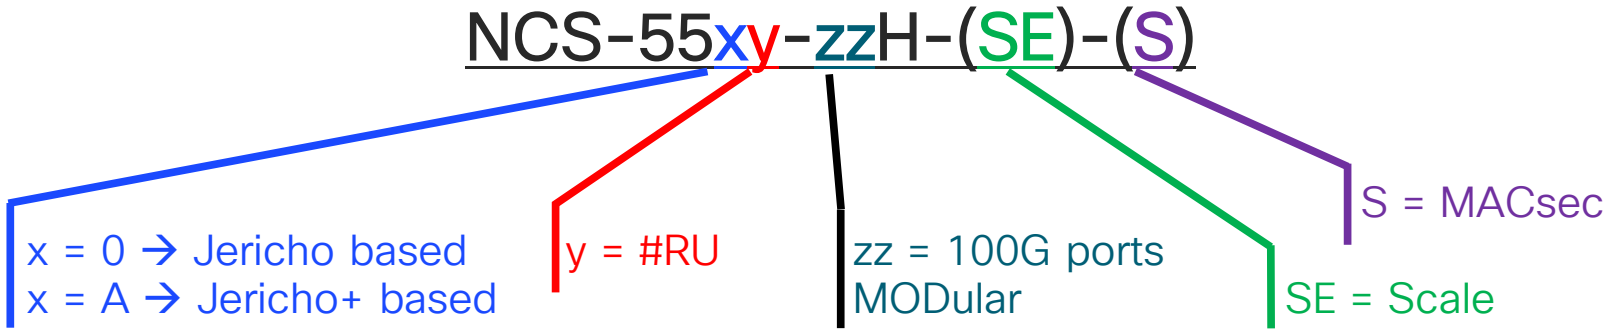

Jericho -SE → 2M extra IPv4 addresses

Jericho+ -SE → total 4M IPv4 addresses (more possible in future releases)

NCS5500 Fixed Platforms

NCS-5501-SE

- Single 800 Gbps FA, 4GB packet buffer
  - 600 Mpps
- No Oversubscription, total interfaces: 800G
  - 40x 1/10G SFP ports
  - 4x 40/100G QSFP ports
- Support of Timing and DWDM interfaces

16 regular ports (ports 0 to 15)

24 ports DWDM/ZR capable (ports 16 to 39)

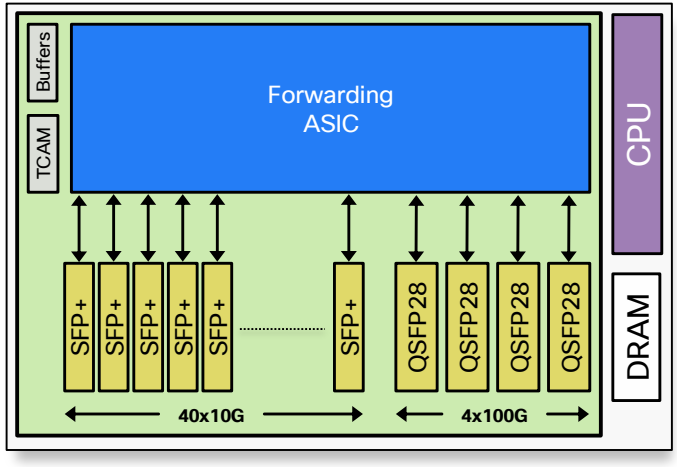

Product	LEM	LPM	eTCAM
NCS-5501-SE	786k	256k-350k	2M

NCS5500 Fixed Platforms

NCS-5501

- Single 800 Gbps forwarding ASIC, 4GB packet buffer
  - 720 Mpps
- Oversubscribed design, total bandwidth of 1.08 Tbps
  - 48x 1/10G SFP ports
  - 6x 40/100G QSFP ports
- No DWDM support
- No timing support

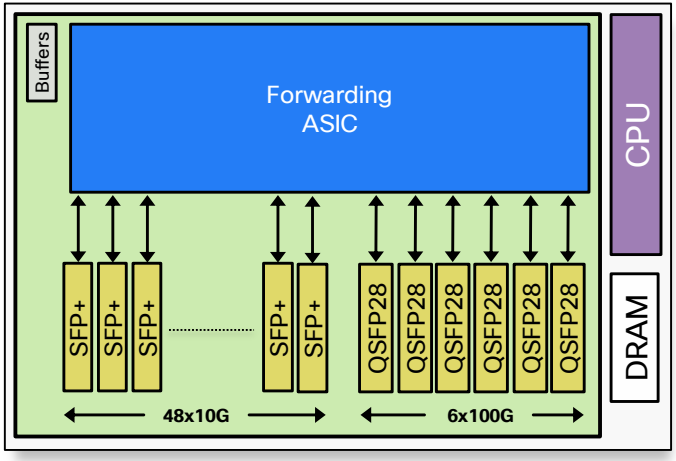

Product	LEM	LPM	eTCAM
NCS-5501	786k	256k-350k	-

NCS5500 Fixed Platforms

NCS-5502-SE

- 4.8 Tbps line-rate 100G < 1850W (Typical, SR optics)
- 48x 100G QSFP28 (or QSFP+)
- 8x 600 Gbps Forwarding ASICs (Common FA with modular chassis)
- 600 Mpps per FA

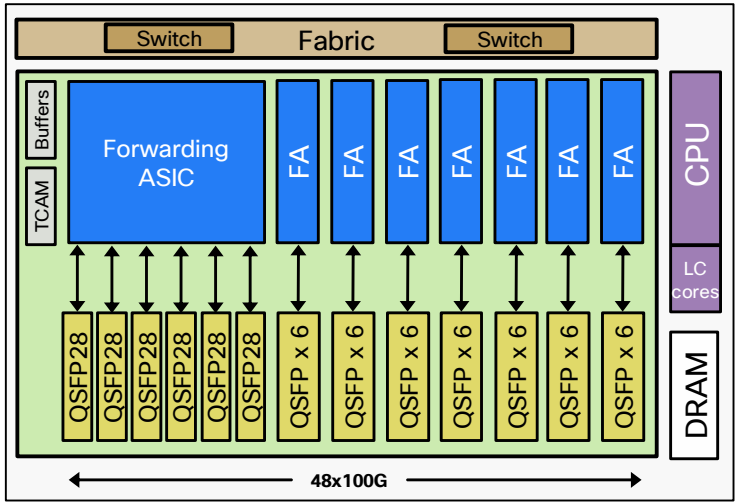

Product	LEM	LPM	eTCAM
NCS-5502-SE	786k	256k-350k	2M

NCS5500 Fixed Platforms

NCS-5502

- 4.8 Tbps line-rate 100G < 1450W (Typical, SR optics)
- 48x 100G QSFP28 (or QSFP+)
- De-pop'd version without external TCAM
- 8x 600 Gbps Forwarding ASICs
- 720 Mpps per FA

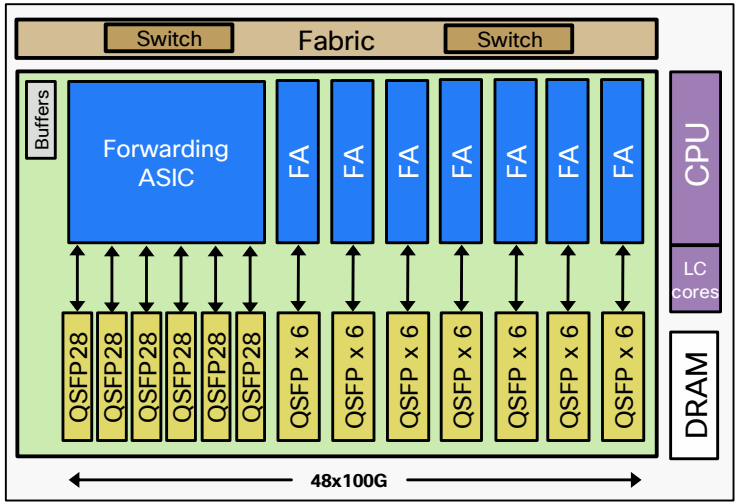

Product	LEM	LPM	eTCAM
NCS-5502	786k	256k-350k	-

NCS5500 Fixed Platforms

NCS-5502 Internal Architecture

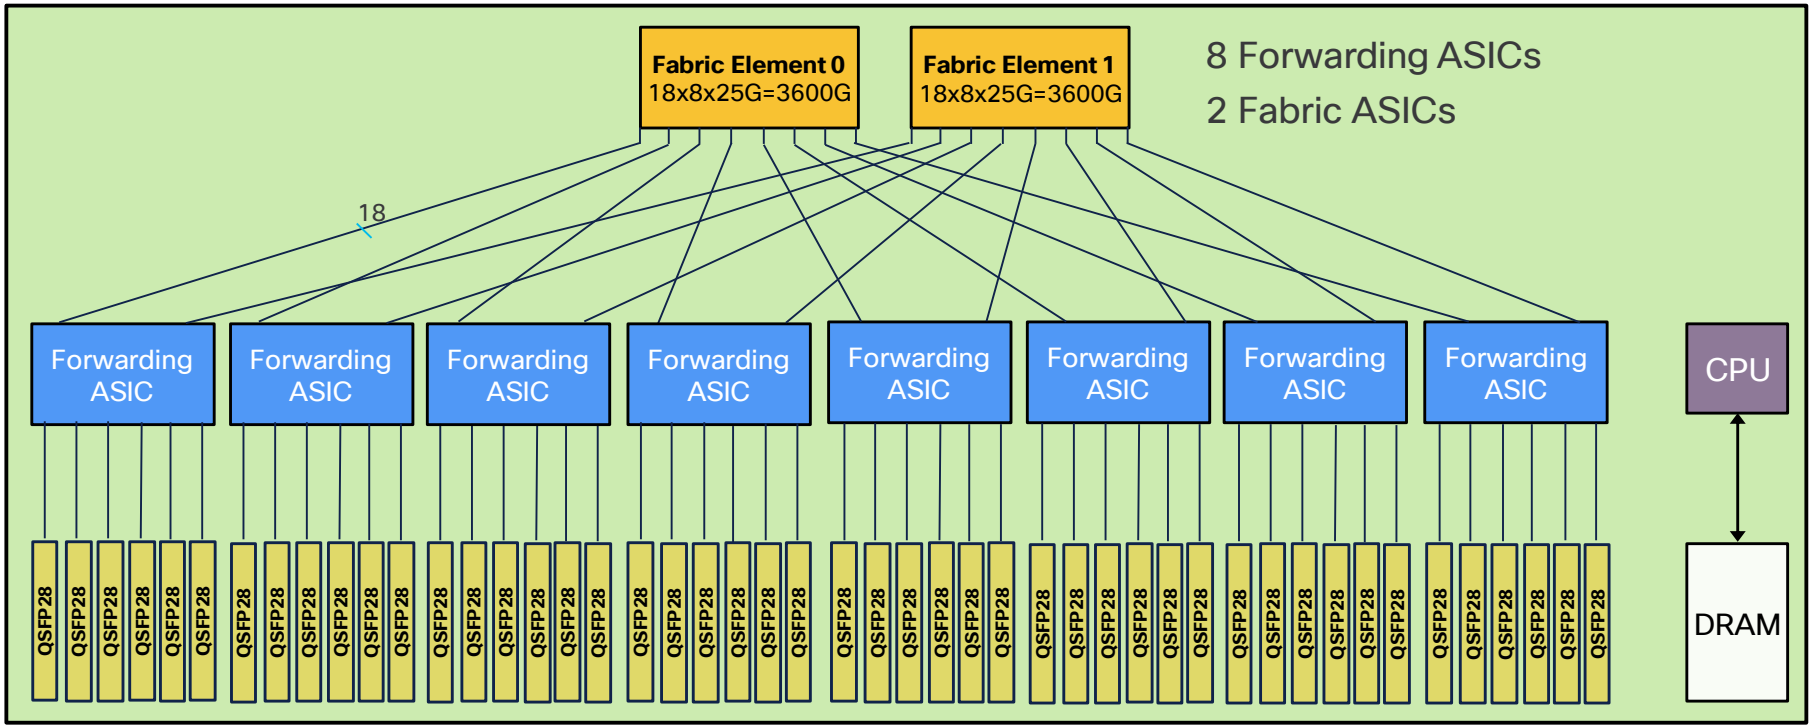

NCS5500 Fixed Platforms

NCS-5501 and NCS-5502 Back View

NCS5500 Fixed Platforms

NCS-55A1-36H-S / NCS-55A1-36H-SE-S

- 36x QSFP28 ports in 1 RU
- Single Intel Broadwell-DE D1577 CPU
  - 8-core @ 1.6GHz
  - 32GB RAM, 64GB SSD
- 2 Redundant Power Modules: 2kW AC or DC
  - Base system: Typical= 1100W / Max Power= 1450W
  - Scale system: Typical= 1300W / Max Power= 1700W
- 3 Redundant (N+1)
  - Front to Back Fan Modules
- Depth: 30 inches

Product	LEM	LPM	eTCAM
55A1-36H-S	786k	256k-350k	-
55A1-36H-SE-S	786k	256k-350k	4M+

NCS-55A1-36H-S / NCS-55A1-36H-SE-S

Internal Architecture

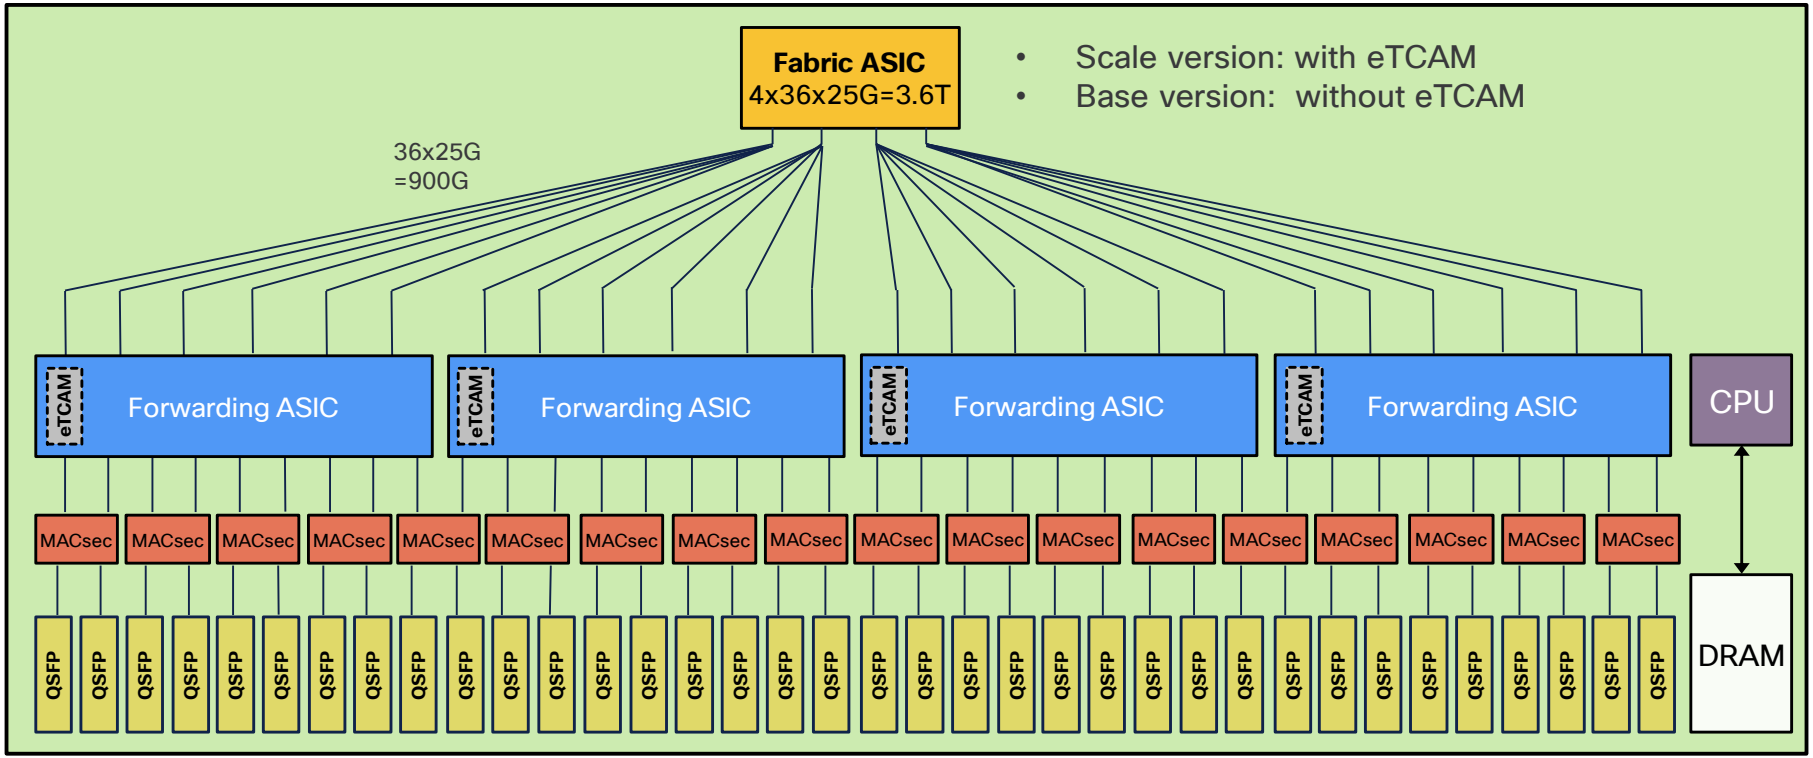

NCS5500 Fixed Platforms

NCS-55A1-24H

- 1 Rack Unit Fixed System: 24x QSFP28 ports
- Base version only and no MACSEC capability
- 1588 / Sync-E Capable
- 2x 900 Gbps Forwarding ASICs
- No Fabric ASIC, Forwarding ASICs are directly connected
- Dimension: 1RU / Depth: 21 inches

Product	LEM	LPM	eTCAM
NCS-55A1-24H	786k	1M-1.5M	-

NCS5500 Fixed Platforms

NCS-55A1-24H

- Single Intel Broadwell-DE D1577 CPU
  - 8-core @ 1.6GHz
  - 32GB RAM, 128GB SSD
- 2 Redundant Power Modules: AC or DC
  - Typical= 600W / Max Power= 800W
- 2 Redundant (N+1) Fan Modules: Front to Back (B2F planned)

NCS5500 Fixed Platforms

NCS-55A1-24H

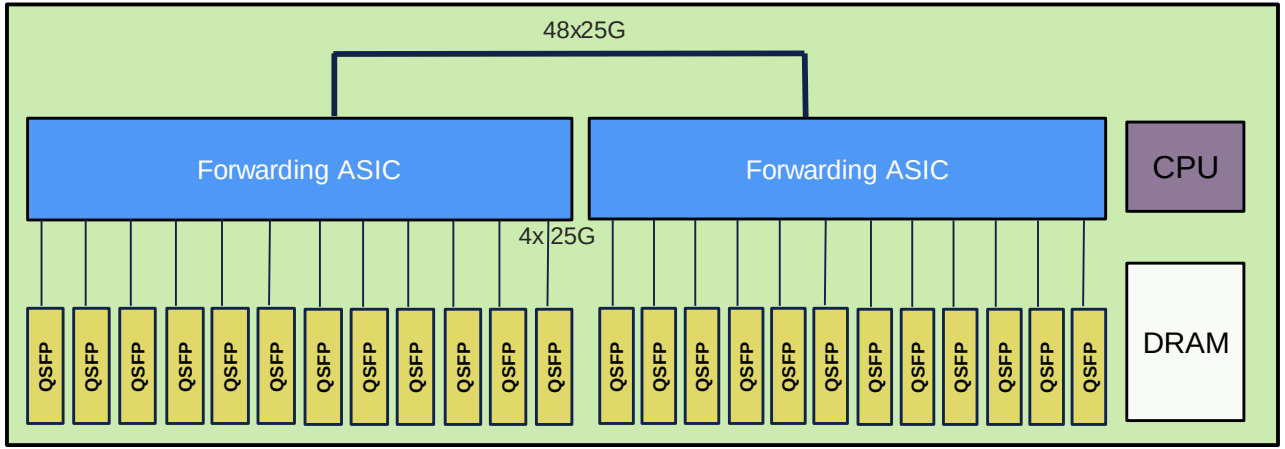

Oversubscription of 12x100G ports per 900G Forwarding ASIC

NCS5500 Fixed Platforms

NCS-55A2-MOD-S Series

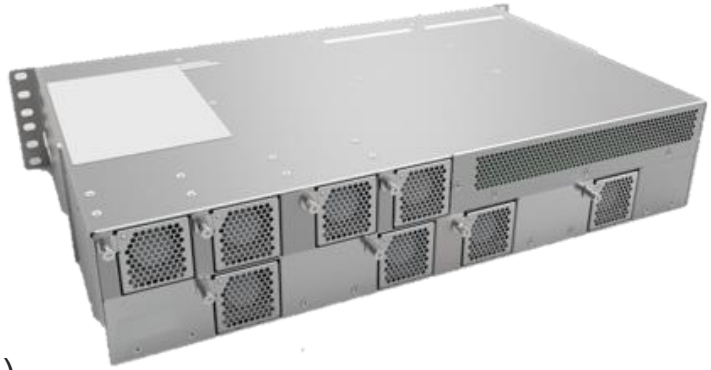

- 2RU, 11 inches deep (280mm)
- 1x Jericho+ Forwarding ASIC
  - 835 Mpps / 900 Gbps (160% max oversubscribed)
- Fixed 40x 1/10G SFP/SFP+ DWDM capable
  - 24x 1/10G
  - 16x 1/10/25G (MACsec at 10/25G)
- 2x 400G Modular Port Adaptor bays
- Timing 1588/SyncE and MACsec Capable
- 8x Fan Modules (F2B), 2x Power Supply AC/DC (Front)

NCS-55A2-MOD-S

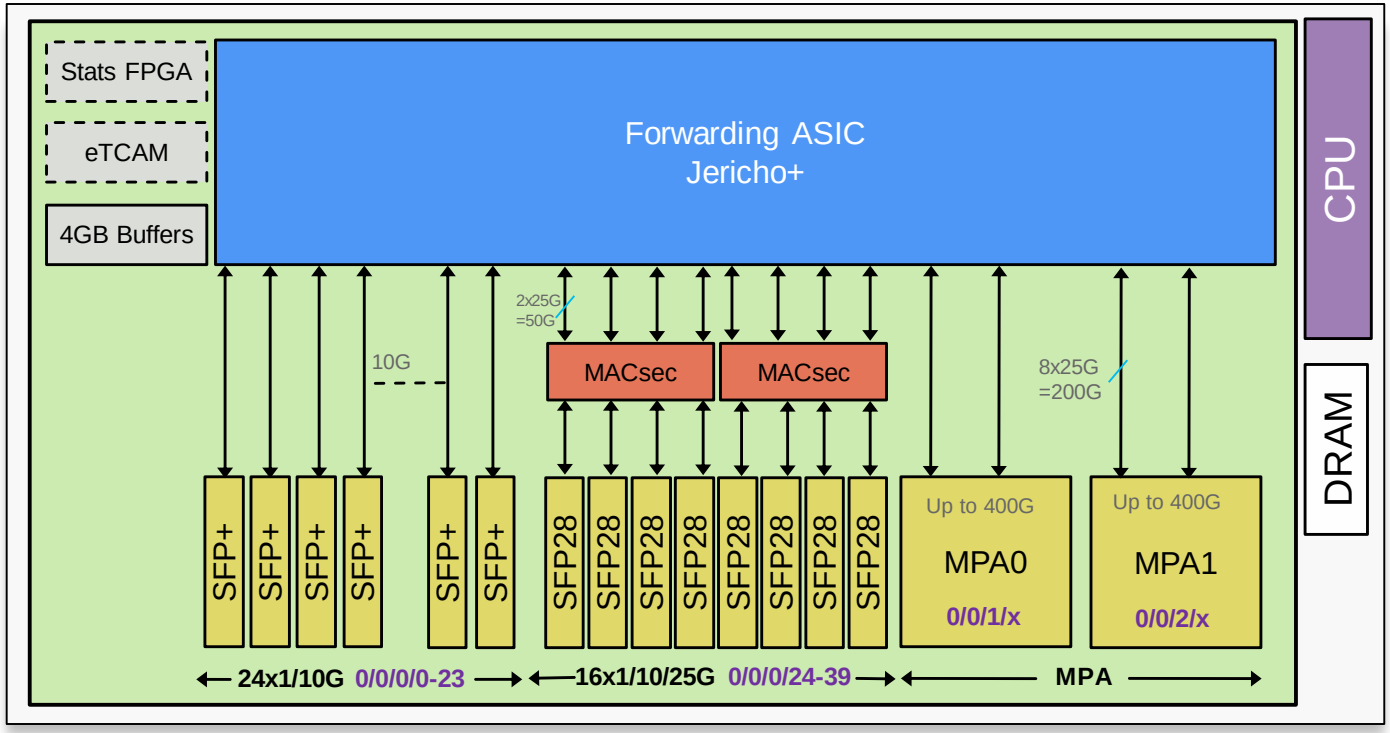

NCS-55A2-MOD Series

NCS-55A2-MOD-S

- Base version
- Single Intel Broadwell CPU (6 cores @ 2GHz), 32GB RAM, 128GB SSD

Product	LEM	LPM	eTCAM
NCS-55A2-MOD-S	786k	256k-350k	-

NCS-55A2-MOD Series

NCS-55A2-MOD-HD-S

- Base Hardened version
  - GR 3108 Class 2
  - Expected temperature range: around -40C to +70C
- Single Intel Broadwell CPU (6 cores @ 2GHz), 32GB RAM, 128GB SSD
- Single Temp Hardened MPA option
  - MPA 4x QSFP28 (4x10G / 40G / 100G)

Product	LEM	LPM	eTCAM
NCS-55A2-MOD-S	786k	256k-350k	-

NCS-55A2-MOD Series

NCS-55A2-MOD-SE-S

- Scale version
- Single Intel Broadwell CPU (8 cores @ 2GHz), 32GB RAM, 128GB SSD
- External TCAM and FPGA for statistics (future use)

Product	LEM	LPM	eTCAM
NCS-55A2-MOD-S	786k	256k-350k	4M+

Modular Port Adapters (MPA)

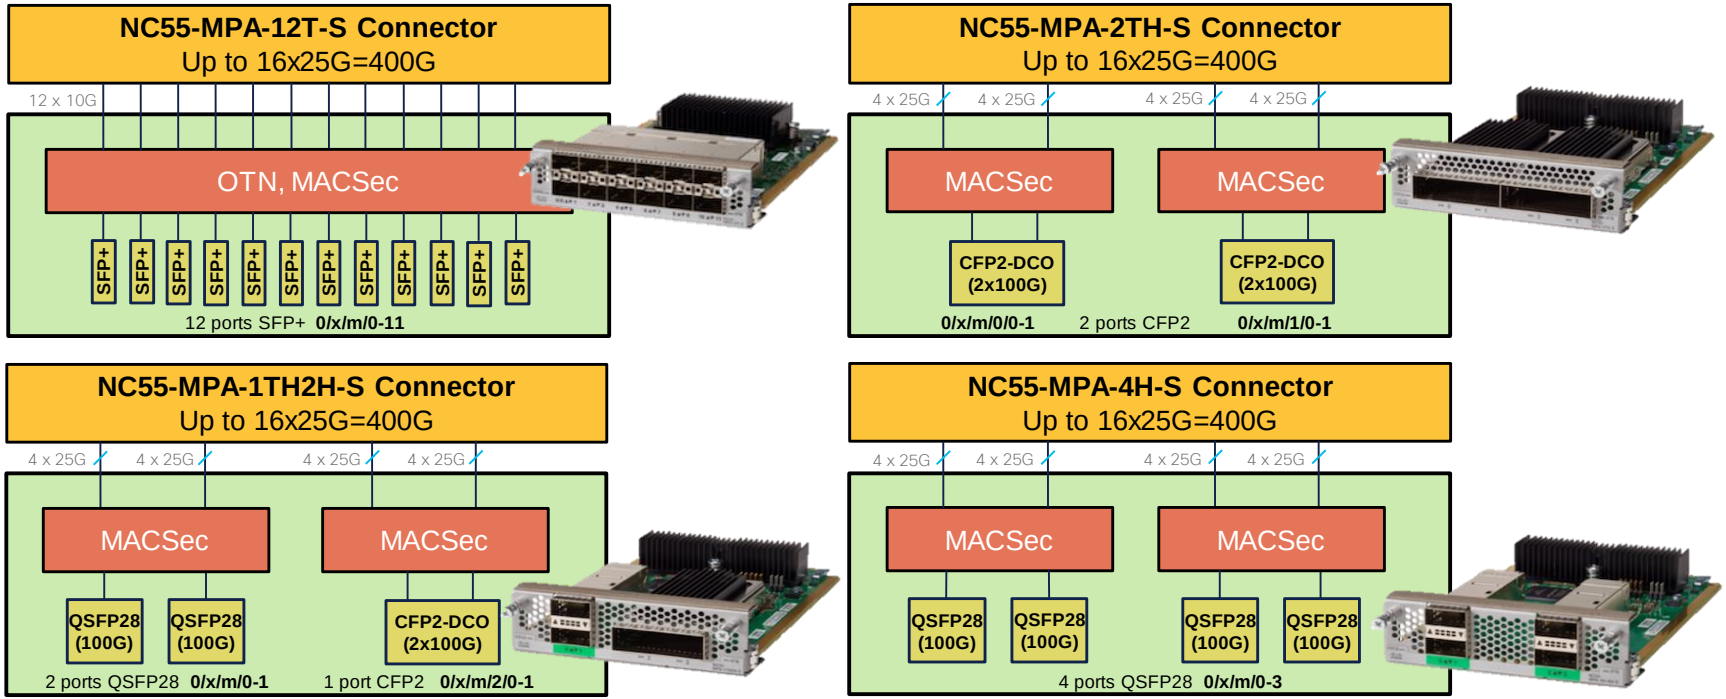

NCS-55A2-MOD Series

Timing Capabilities

- IEEE 1588-2008 PTP support
- External Satellite Inputs – 1PPS, 10MHz, TOD
- No BITS inputs
- Built-in GNSS/GPS Receiver (Trimble) Hardware
- ZL30363 IEEE 1588 and SyncE Packet Clock Network Synchronizer
  - with Stratum 3E OCXO Clock

NCS-55A2-MOD Series

MACsec Support

- Not on first 24xSFP+
- Capable on last 16xSFP28 fixed ports except 1GE mode
- MACsec on all MPA ports except 1GE mode
- MACsec support introduced in 6.6.1

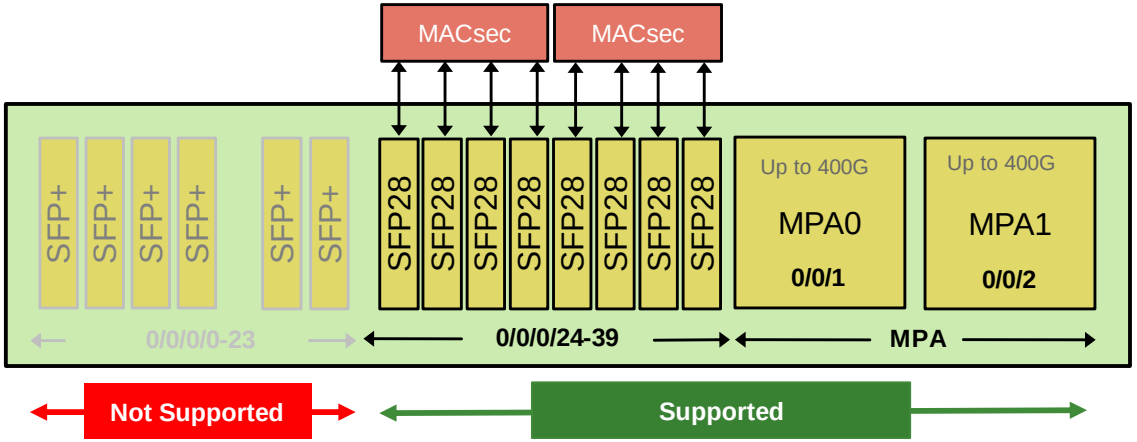

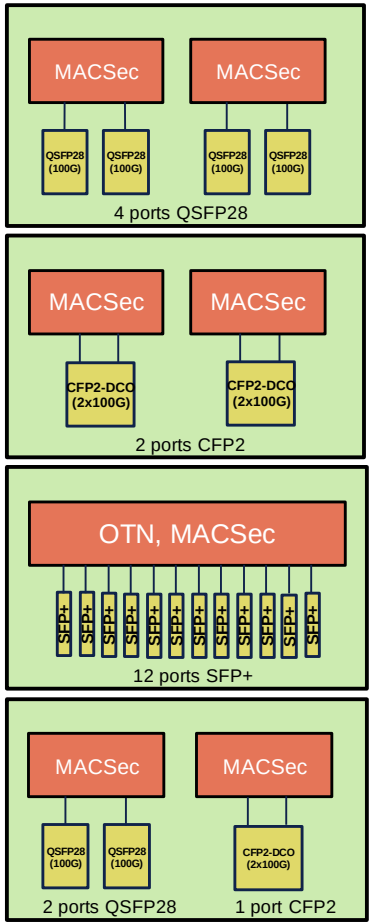

NCS5500 Fixed Platforms

NCS-55A1-24Q6H-S

New in 2019

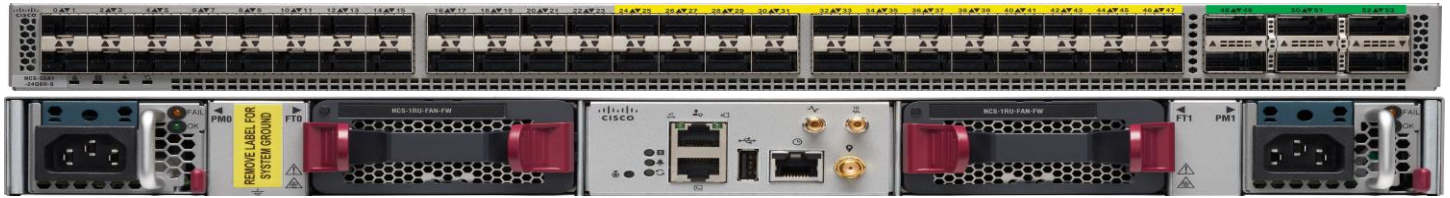

- 1RU: 48 ports SFP + 6 ports QSFP
  - 24x 1G/10G/25G + 24x 1G/10G + 6x 100G
  - Base version only
- 1x Jericho+ Forwarding ASIC (SoC)
  - Jericho Sclale
  - 835 Mpps / 900 Gbps
  - Oversubscription of 1.44Tbps ports
- Timing:
  - 1588 / Sync-E Capable (Class B)
- MACsec:
  - 100G ports
  - 16 out of the 24x SFP28

Product	LEM	LPM
NCS-55A1-24Q6H-S	786k	256k-350k

NCS5500 Fixed Platforms

NCS-55A1-24Q6H-S

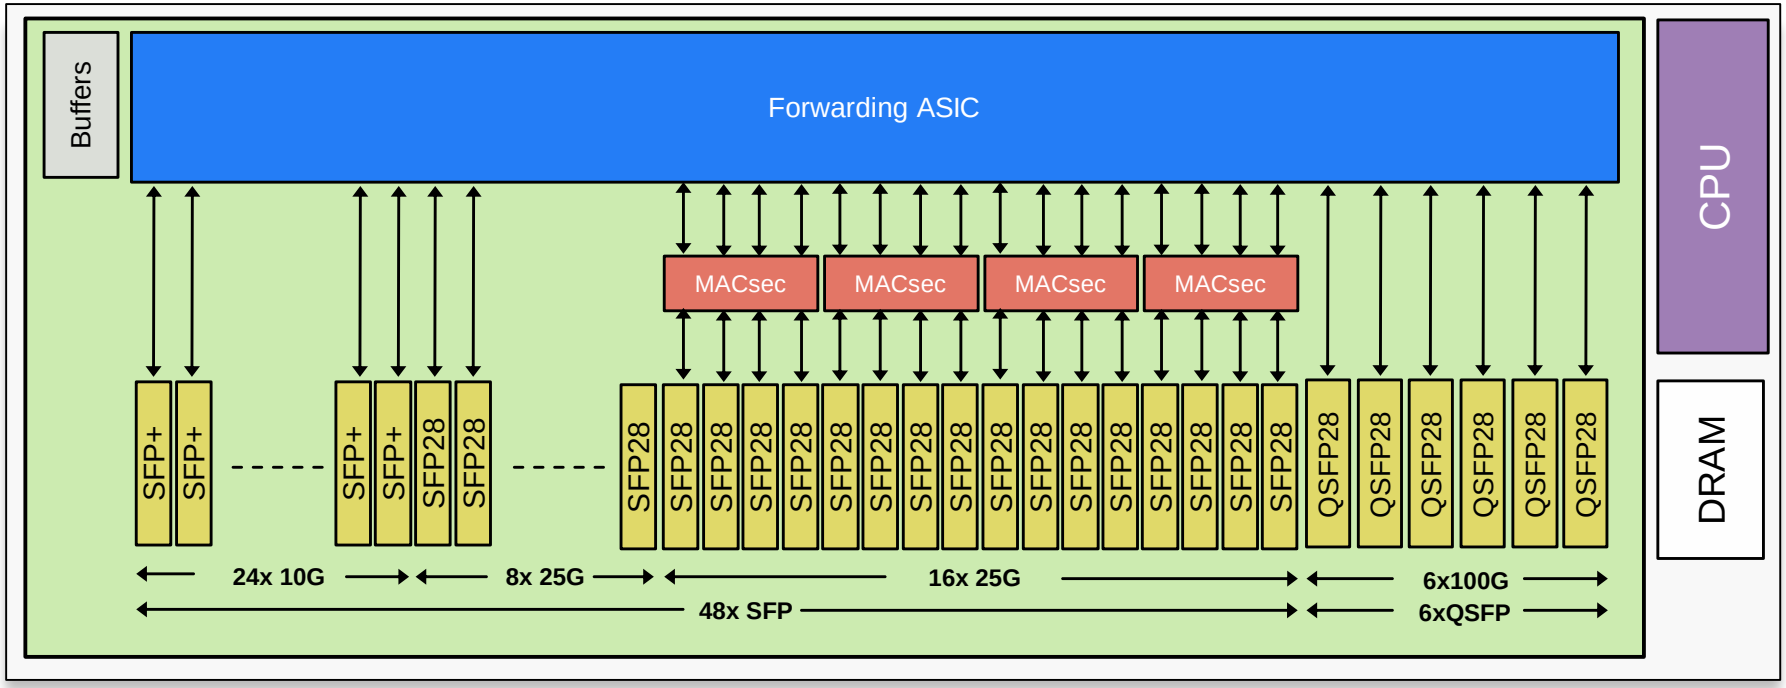

NCS5500 Fixed Platforms

NCS-55A1-48Q6H

NEW in 2019

- 1RU: 48 ports SFP + 6 ports QSFP
  - 48x 1G/10G/25G + 6x 100G
- Base version only
- 2x Jericho+
  - no fabric, back-to-back
  - 835 Mpps / 900 Gbps each
  - Large LPM
- Timing:
  - 1588 / Sync-E Capable (Class B)
- MACsec:
  - 100G ports only

Product	LEM	LPM
NCS-55A1-24Q6H-S	786k	1M-1.5M

NCS5500 Fixed Platforms

NCS-55A1-48Q6H

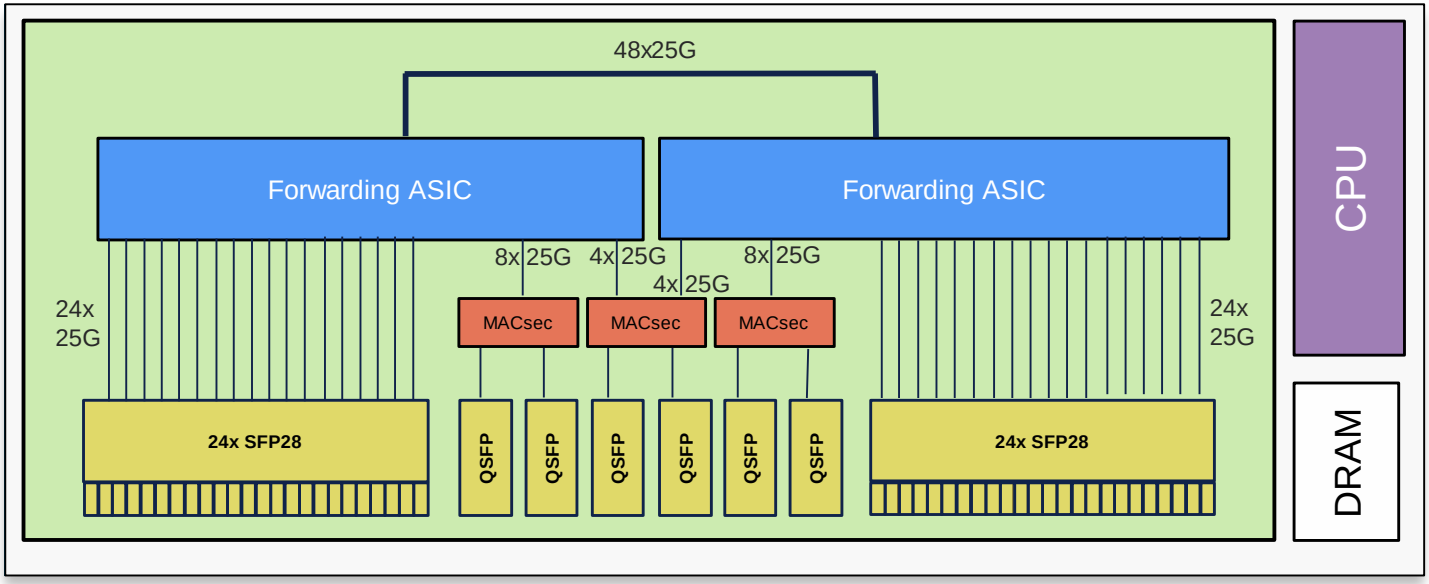

For Reference

NCS5500 Fixed Systems Comparison

	ASIC	QSFP	SFP	eTCAM	Capacity	Forwarding Capacity
NCS-5501	QMx	6	48	-	1.08 Tbps	800 Gbps
NCS-5501-SE	QMx	4	40	Yes	800 Gbps	800 Gbps
NCS-5502	8x J	48	-	-	4.8 Tbps	4.8 Tbps
NCS-5502-SE	8x J	48	-	Yes	4.8 Tbps	4.8 Tbps
NCS-55A1-24H	2x J+	24	-	-	2.4 Tbps	1.8 Tbps
NCS-55A1-36H-S	4x J+	36	-	-	3.6 Tbps	3.6 Tbps
NCS-55A1-36H-SE-S	4x J+	36	-	Yes	3.6Tbps	3.6 Tbps
NCS-55A2-MOD(-SE)-S	1x J+	Up to 8	40	Yes (-SE)	1.4Tbps	900 Gbps
NCS-55A1-24Q6H-S	1x J+	6	48	-	1.4 Tbps	900 Gbps
NCS-55A1-48Q6H	2x J+	6	48	-	1.8 Tbps	1.8 Tbps

NCS5500 Forwarding ASIC

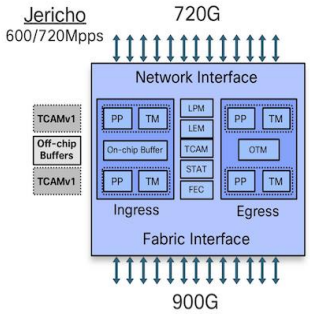

NCS5502
NCS5502-SE

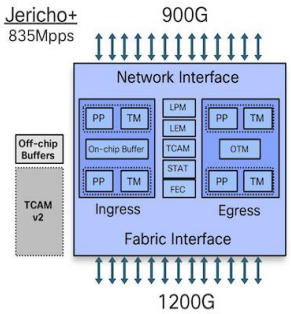

NCS55A2-MOD-S
NCS55A2-MOD-SE-S
NCS55A1-24Q6H-S
NCS55A1-36H-S
NCS55A1-36H-SE-S

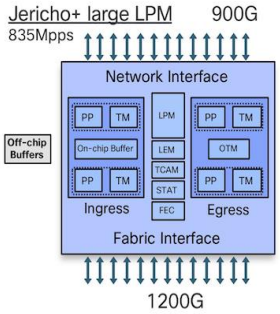

NCS55A1-24H
NCS55A1-48Q6H

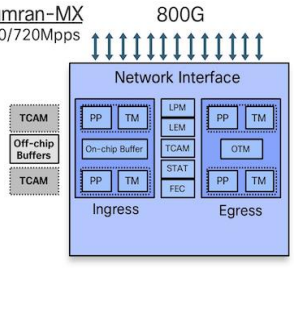

NCS5501
NCS5501-SE

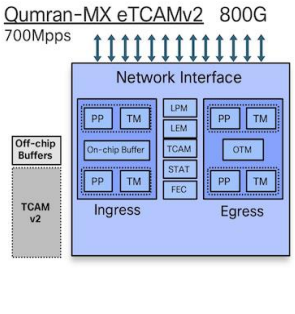

NCS560

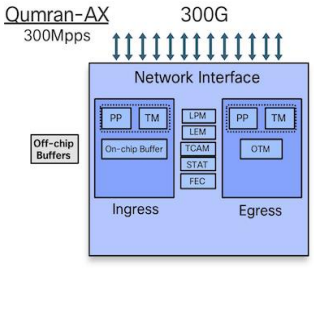

NCS540

NCS5500 Modular Chassis

You make networking **possible**

Cisco live!

Three Chassis

- Common parts
  - RP
  - SC
  - Line Cards
  - Power Supply Modules
- Specific
  - Chassis
  - 3x Fan Tray Modules
  - 6x Fabric Line Cards

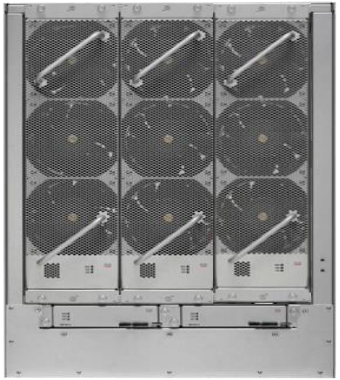

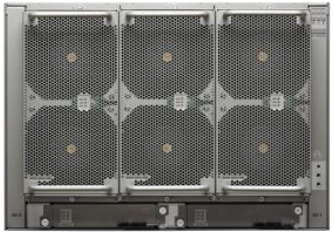

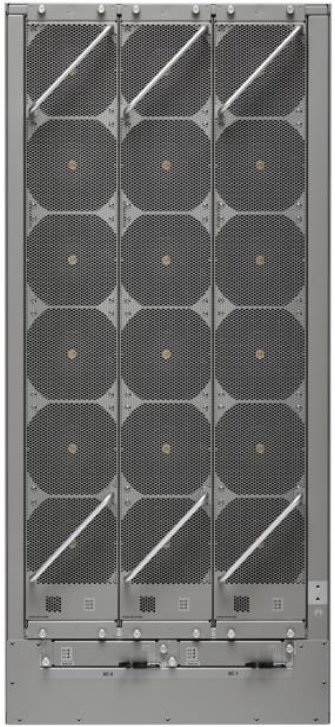

Orthogonal Design

- No backplane/midplane for data path
- Direct connection between LC to fabric cards at 90 degrees
- Air inlets above and between optics
- Air inlets on RP & power supplies

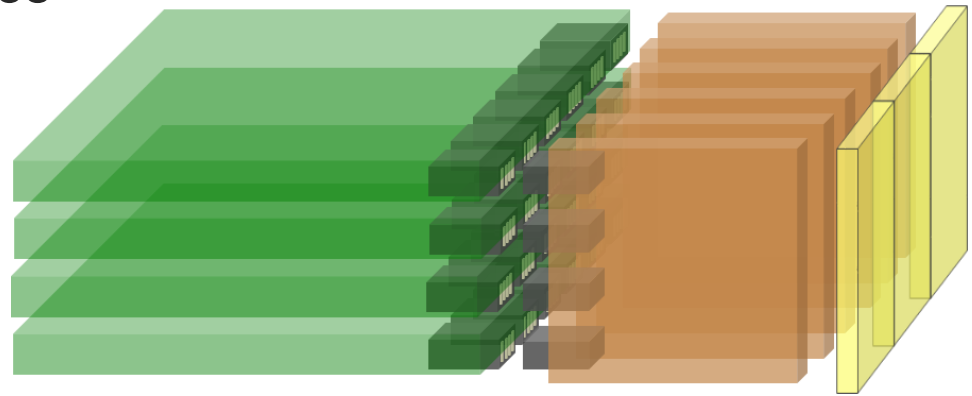

NCS5500 Modular Chassis

Mechanical Layout

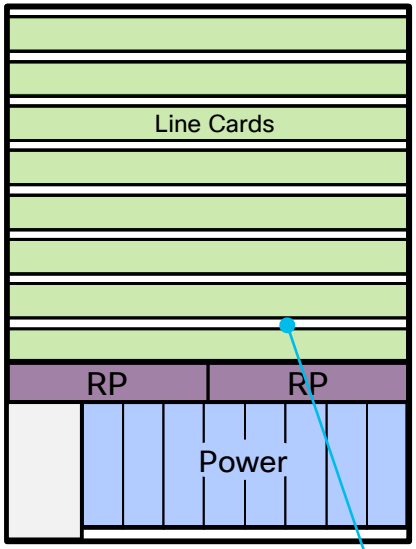

Front View

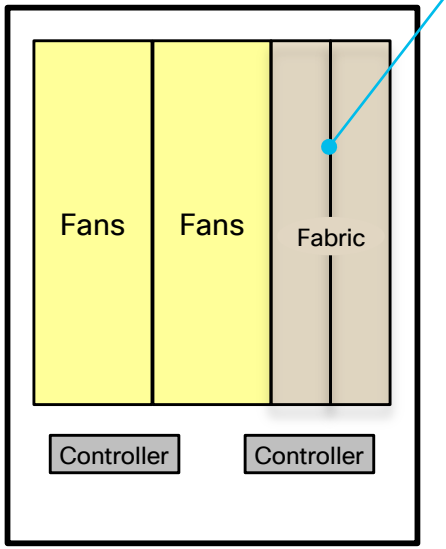

Rear View

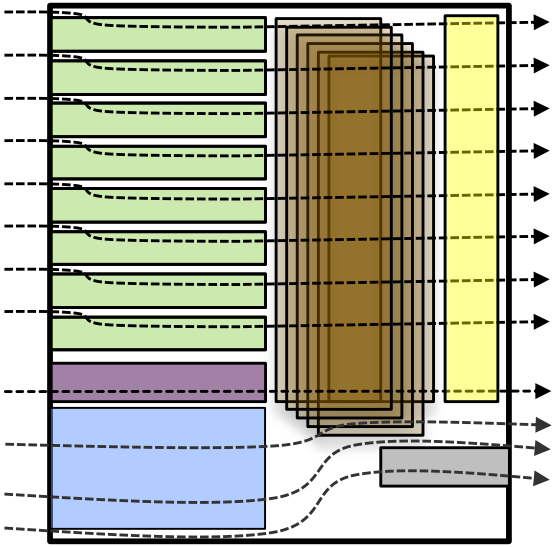

Side View w/ Airflow

Fabric Behind Fans

Air Intake

NCS-5504 Chassis

Up to 14.4Tbps

- Dimensions – 7RU
  - H x W x D: 12.25 x 17.5 x 31.7“
  - (31.1 x 44.50 x 84.20 cm)
- Power Supplies
  - 4 supplies
  - AC or DC

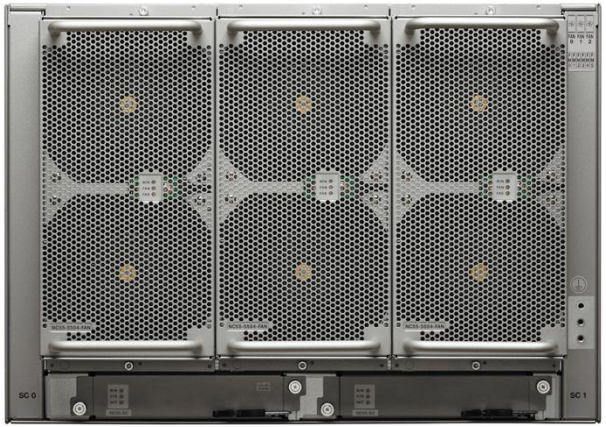

For Reference

NCS-5508 Chassis

Up to 28.8Tbps

- Dimensions – 13RU (1/3 rack)
  - H x W x D: 22.7 x 17.5 x 31.7”
  - 57.78 x 44.50 x 80.67 cm
  - Depth: 34.78 in / 88.34 cm
    (from linecard ejector to fantray handles)
- Power Supplies
  - 8 supplies
  - NEBS via air filter door and enclosure
- 28.8 Tbps @ 6920 W = 0.24 W/Gbps
  - 288 QSFP28 or QSFP+ ports

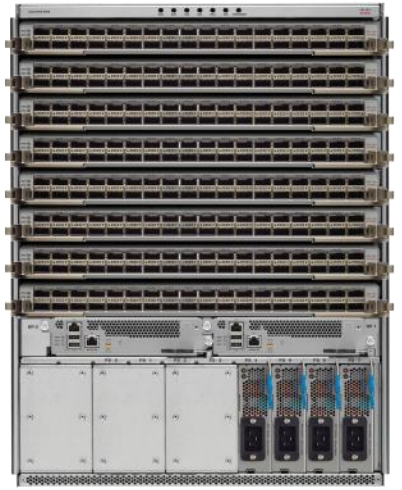

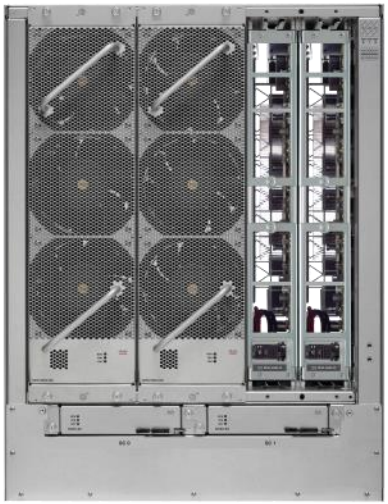

NCS-5516 Chassis

Up to 57.6Tbps

- Dimensions – 21 RU (1/2 rack)
  - H x W x D: 36.7 x 17.5 x 31.7”
  - 93.41 x 44.50 x 80.67 cm
  - Depth: 34.78 in / 88.34cm (from LC ejector to FT handles)
- Power supplies
  - 10 power supplies AC or DC
- 57.6 Tbps @ ~18000W = 0.31 W/Gbps
  - 576 QSFP28 or QSFP+ ports

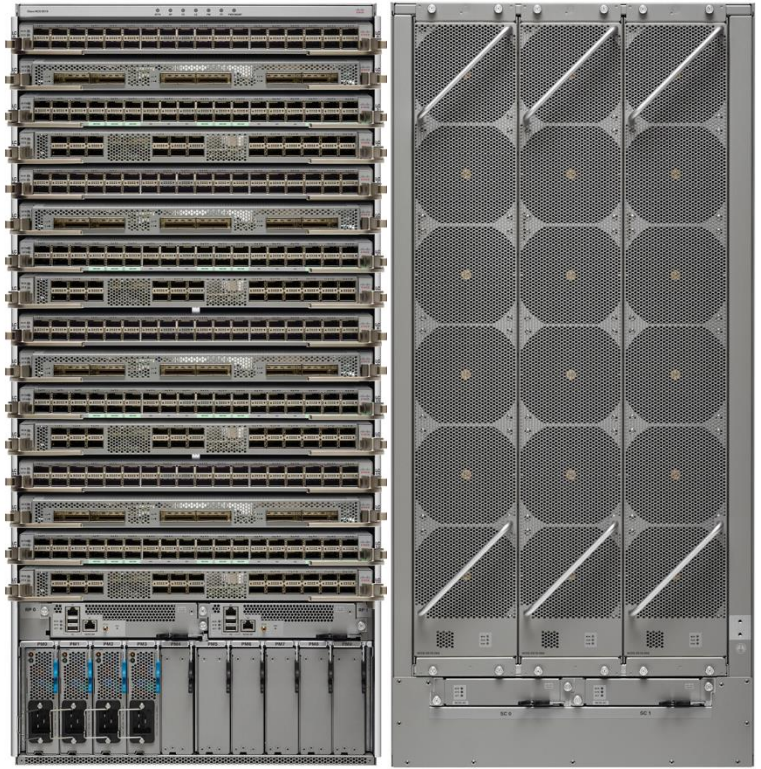

Switch Fabric Cards

- Cell-based fabric
  - FE3600 fabric ASIC
  - Next Gen “Ramon” ASIC
- 6 Fabric Cards per chassis
- FE3600
  - Same Switch Fabric Cards for both Jericho and Jericho+
- Ramon
  - Required for J2 Line Cards

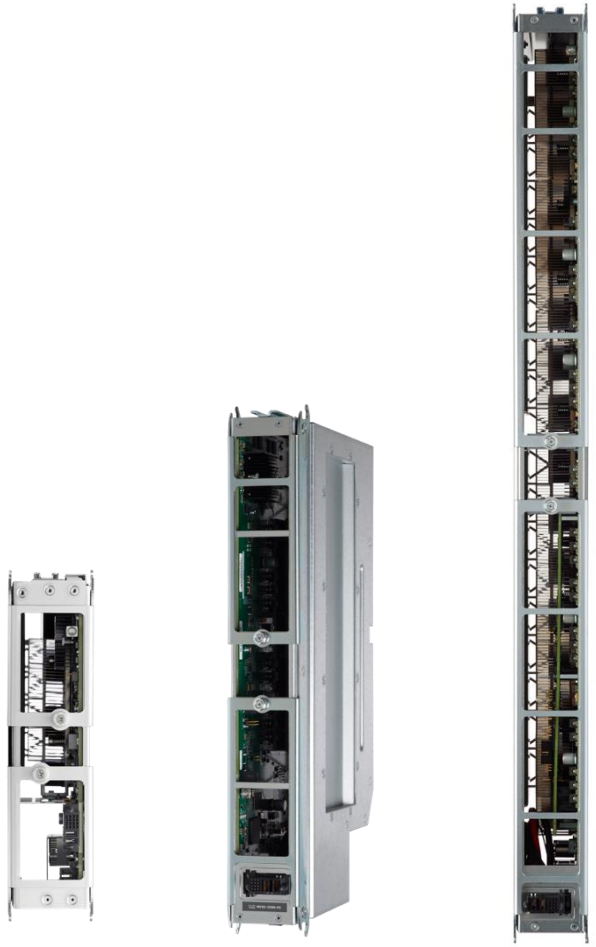

Switch Fabric Cards

Number of Fabric ASICs per Fabric Cards

- FE3600

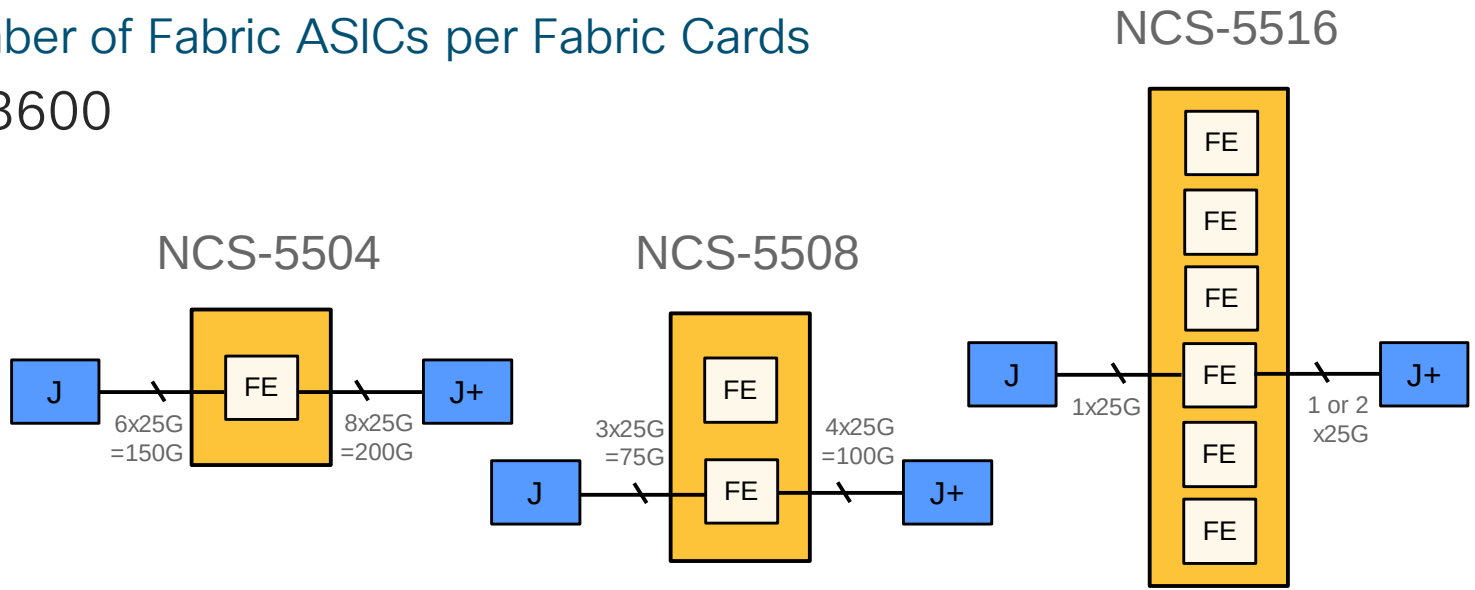

- Support: J/J+ cards, not J2
- PIDs: NC55-5504-FC / NC55-5508-FC / NC55-5516-FC

Switch Fabric Cards

Number of Fabric ASICs per Fabric Cards

- Ramon

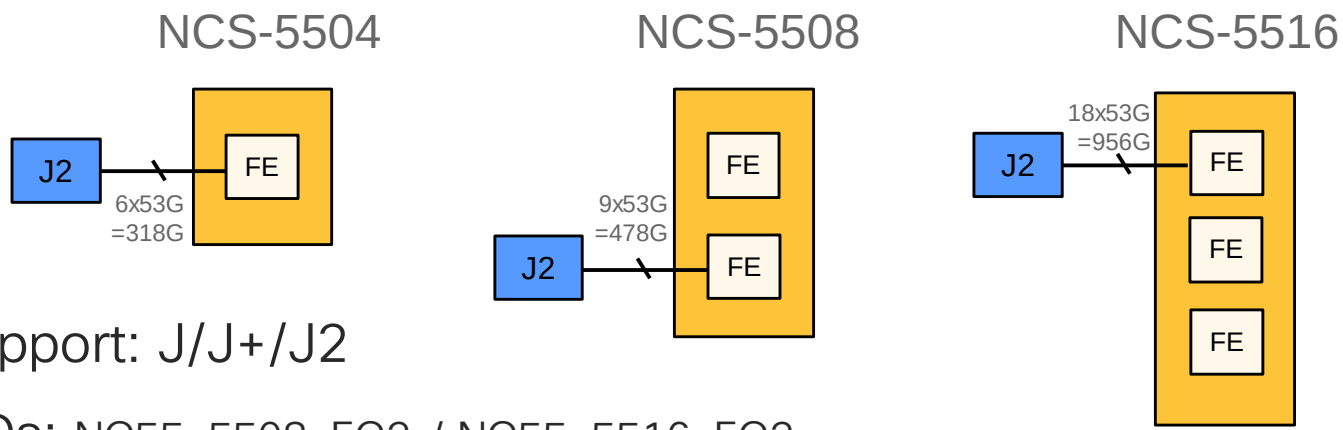

- Support: J/J+/J2
- PIDs: NC55-5508-FC2 / NC55-5516-FC2
  - NCS55-5504-FC2 in Roadmap

Switch Fabric Cards and Fan Trays

- v2 Fabric Cards requires v2 Fan Trays
  - NC55-5508-FC + NC55-5508-FAN
  - NC55-5508-FC2 + NC55-5508-FAN2
  - NC55-5516-FC + NC55-5516-FAN
  - NC55-5516-FC2 + NC55-5516-FAN2

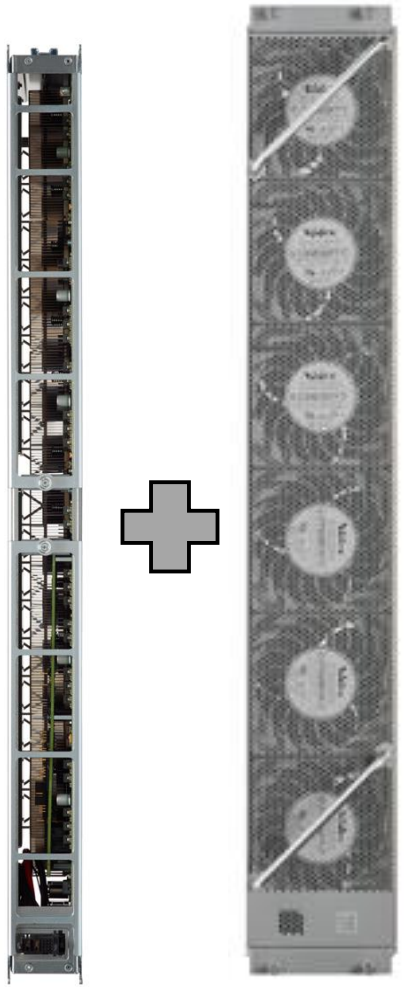

NCS5500 Modular Chassis

36x 100G Line Card Bandwidth Example

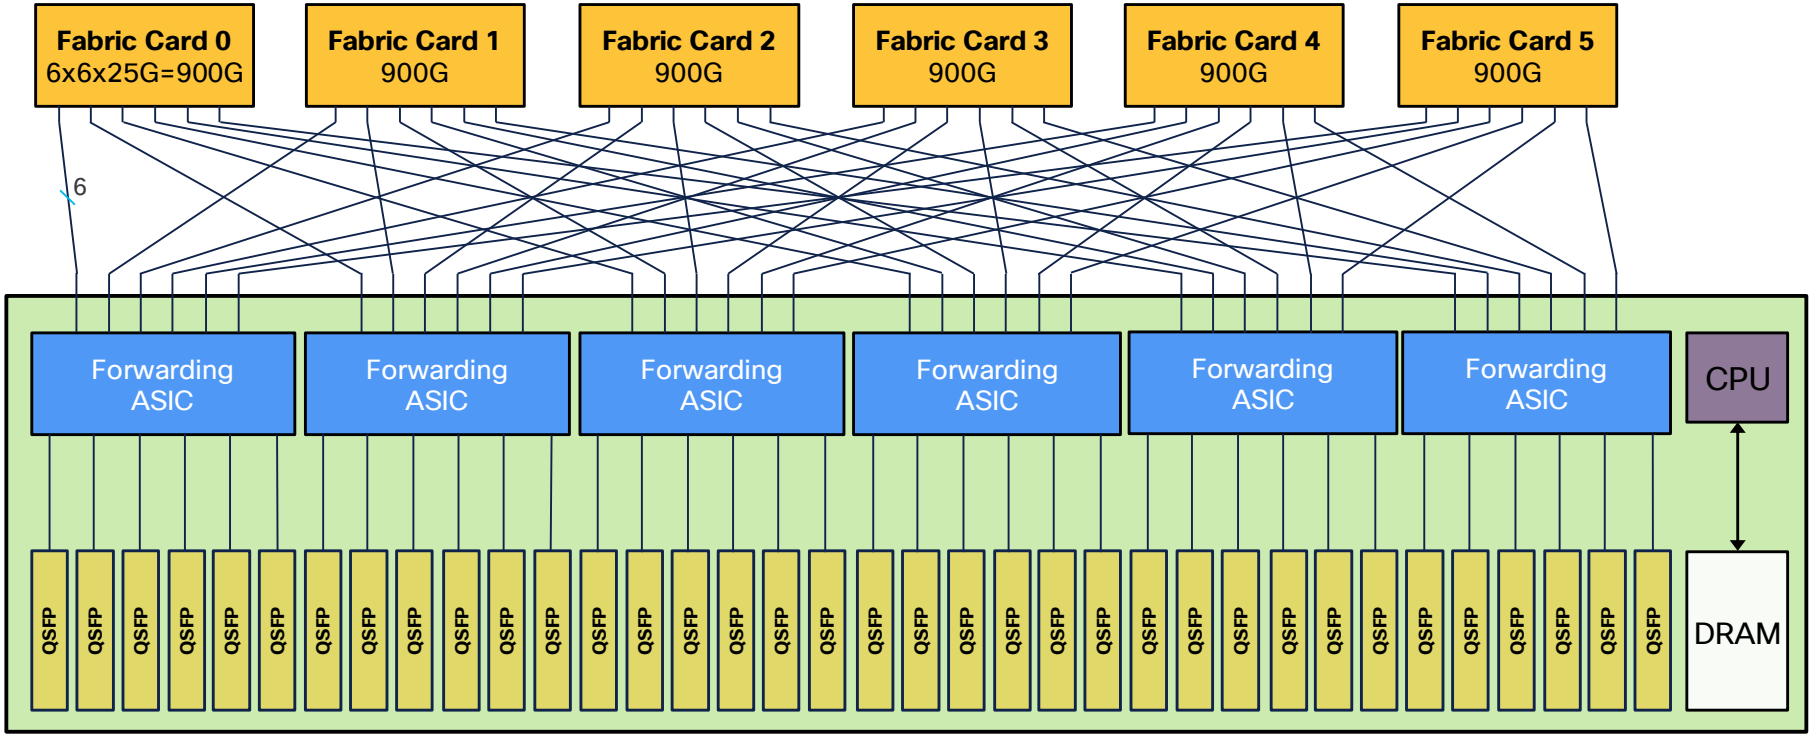

NCS5500 Modular Chassis

Common System Controller and Route Processor

- Route Processor
  - Ivy Bridge with 24GB RAM
  - Routing and management tasks
- System Controller
  - Chassis control and monitoring
    - Fan trays / Power supply
  - Ethernet Out-of-Band Channel (EOBC)
  - Ethernet Protocol Channel (EPC)

Modular Line Cards based on Jericho

36 Port 100GE no eTCAM (QSFP)
NC55-36X100G-BA

24 Port 100GE & 12 Port 40GE
External TCAM (QSFP) - NC55-24H12F-SB

24 Port 100GE External TCAM (QSFP)
NC55-24X100G-SB

36 Port 100GE with MACsec
No eTCAM (QSFP) - NC55-36X100G-BM

18 Port 100GE & 18 Port 40GE
No eTCAM (QSFP) - NC55-18H18F-BA

6 Port 100/150/200GE with MACsec
No eTCAM (CFP2) - NC55-6x200-DWDM-S

Modular Line Cards based on Jericho+

36 Port 100GE External TCAM (Scale)
NC55-36X100G-A-SB (-SE)

12X10, 2X40 & 2XMPA Line Card Base
NC55-MOD-A-S

12X10, 2X40 & 2XMPA Line Card Scale
NC55-MOD-A-SE-S

For Reference

NCS5500 Line Card Comparison

	ASIC	100G	40G	10G	eTCAM	MAC sec	Line Card Capacity	Forwarding Capacity
NC55-36X100G	6x J	36	-	-	-	-	3.6 Tbps	3.6 Tbps
NC55-36X100G-S	6x J	36	-	-	-	Yes	3.6 Tbps	3.6 Tbps
NC55-18H18F	3x J	18	18	-	-	-	2.52 Tbps	2.16 Tbps
NC55-24X100G-SE	4x J	24	-	-	Yes	-	2.4 Tbps	2.4 Tbps
NC55-24H12F-SS	4x J	24	12	-	Yes	-	2.88 Tbps	2.88 Tbps
NC55-6X2H-DWDM-S	2x J	6x 100/150/200	-	-	-	Yes	1.2 Tbps	1.2 Tbps
NC55-36X100G-A-SE	4x J+	36	-	-	Yes	-	3.6 Tbps	3.6 Tbps
NC55-MOD-A-S	1x J+	Up to 8	2	12	-	Yes	1 Tbps	900 Gbps
NC55-MOD-A-SE-S	1x J+	Up to 8	2	412	Yes	Yes	1 Tbps	900 Gbps

Coming Soon

You make networking **possible**

Cisco live!

Coming Soon: Modular LCs based on Jericho2

24 Port 400GE Base LC
NC55-24D

30 Ports (18x 400GE + 12x 200G) Scale LC
NC55-18D12TH-SE

J2 Line Cards

NC55-24D

- 9.6T capacity line card with 2 Jericho2 chipsets
- Requires
  - 2nd Gen 16/8/4 Fabric card (Ramon-based)
  - 2nd Gen Fan-trays
- Post-FCS: MACsec support on all ports

J2 Line Cards

NC55-18D12TH-SE

- 18 x 400G QSFPDD (or 30 x 200G/100G) TCAM Line Card
- OP2 TCAM (BCM16000 KBP) for Lookup and Stats
- 7.2T capacity line card with 2 Jericho2 chipsets (3.6T per JR2)
- Requires
  - 2nd Gen 16/8/4 Fabric card (Ramon-based)
  - 2nd Gen Fan-trays
- Post-FCS: MACsec support on all ports

J2 Line Cards

NC55-18D12TH-SE

- Port utilization

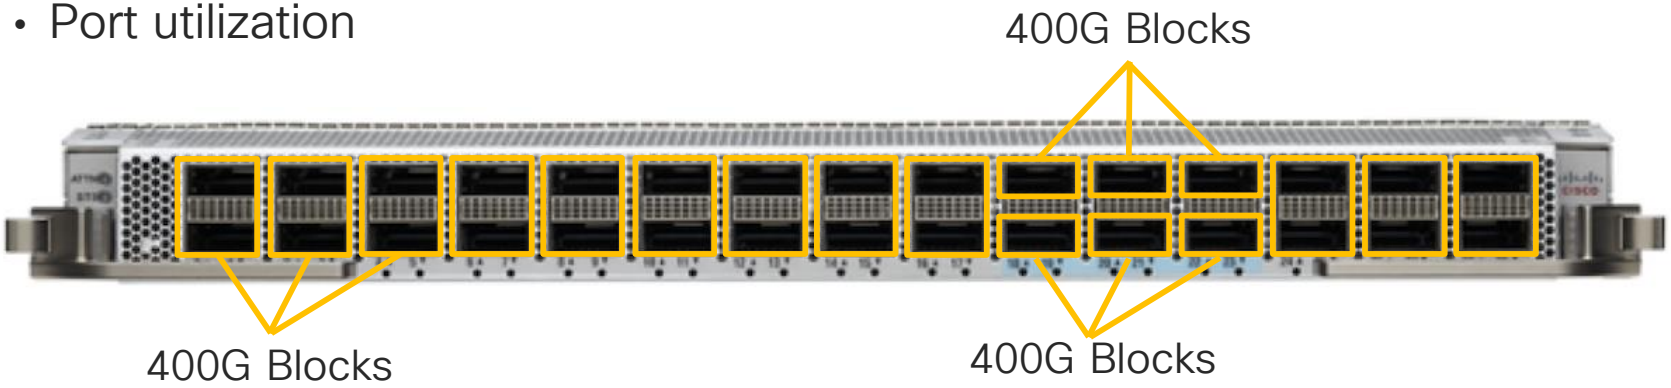

- Blocks of 400G
  - one port at 400G, other is disabled
  - 200G+200G or 100G+100G
- you can use all 30 ports in 100G or 200G mode
  or a mix of 100/200 or 400G up to a total of 7.2T to backplane

Mixing Different Generations of Line Cards

- Same chassis
  - Keep existing RP (or RP-E) and SC cards
  - Requires new fan trays and fabric cards
- At FCS
  - Capability to mix J, J+ and J2 cards
  - Scale numbers will be aligned on J+ for the J2 cards
  - “Compatibility mode”
- Future releases
  - J2-native mode with higher scale

NCS5500 Optics

You make networking **possible**

Cisco live!

NCS5500 Interfaces

Ethernet Only Platforms

- SFP optics slot: offering 1G or 10G (with SFP+) on the following platforms
  - NCS-5501 / NCS-55A2-MOD
  - NCS-55A1-24Q6H-S / NCS-55A1-48Q6H

- QSFP optics slot: offering 100G (with QSFP28), 40G (with QSFP+) and 4x 10G (QSFP+ with break-out cables) on the following platforms or LC
  - NCS-5502(-SE) / NCS-55A1-24H
  - NCS-55A1-36H(-SE)-S
  - Line Cards

NCS5500 Interfaces

Ethernet Only Platforms

- QSA: QSFP to SFP Adaptor
- 25GE only supported on J+ Platforms with 4x25G break-out
- CFP2 optics slot:
  - First on the 6 ports 100/150/200GE DWDM Line Cards
  - Now in MPA for MOD Line Cards (2x CFP2 or combo 1x CFP2 + 2 grey ports)
  - ACO vs DCO

NCS5500 Interfaces

Introducing 400G

- Based on QSFP-DD

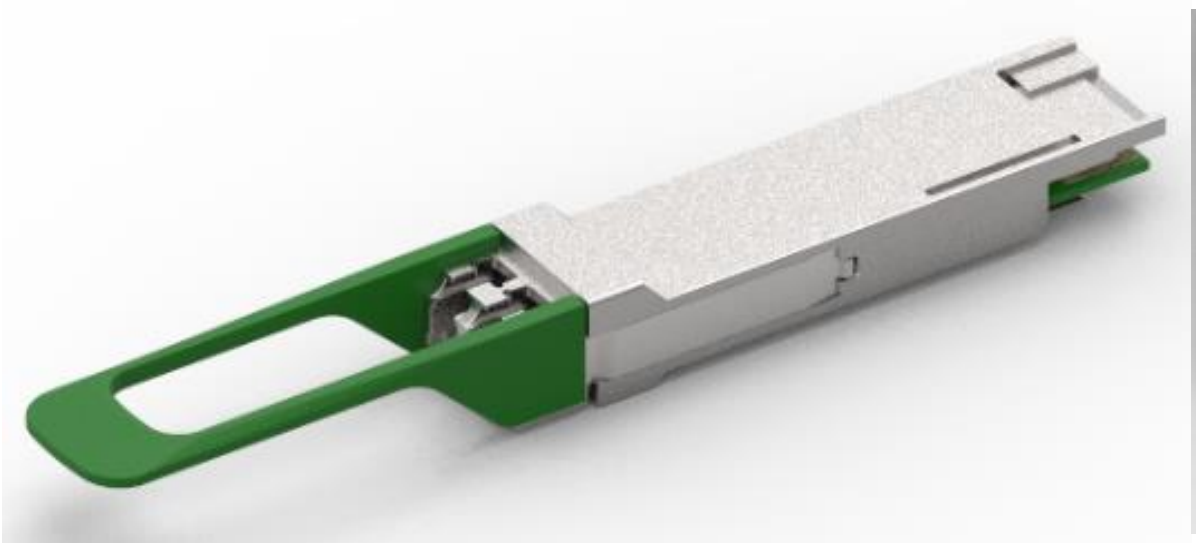

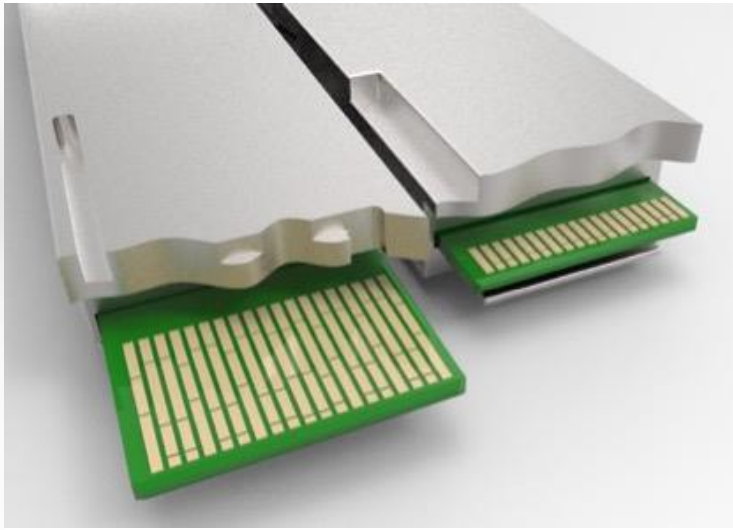

Modulation

NRZ

- On/Off Keying
- Non Return to Zero (NRZ)

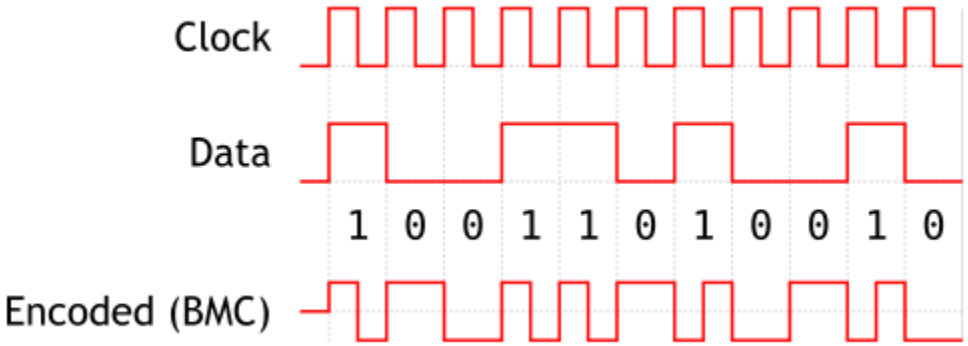

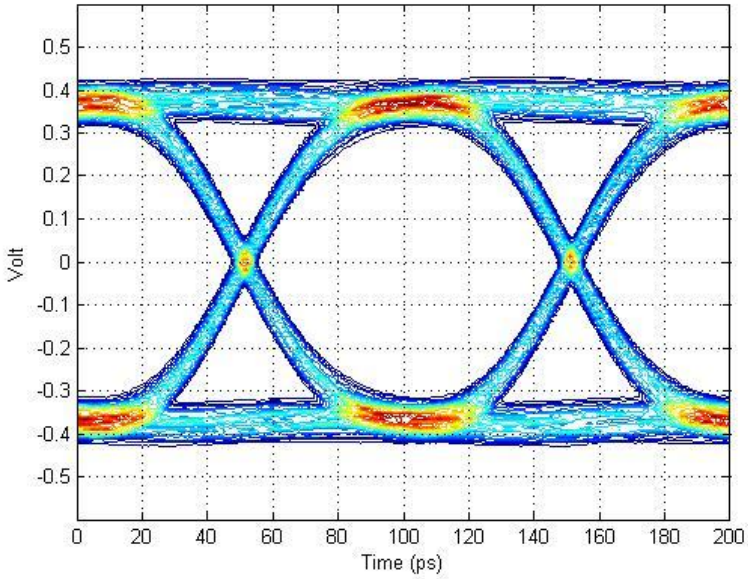

Modulation

PAM4

- PAM4 (Pulse Amplitude Modulation)
  - for 400G electrical Signals and DR4, FR4, and 100G FR optical

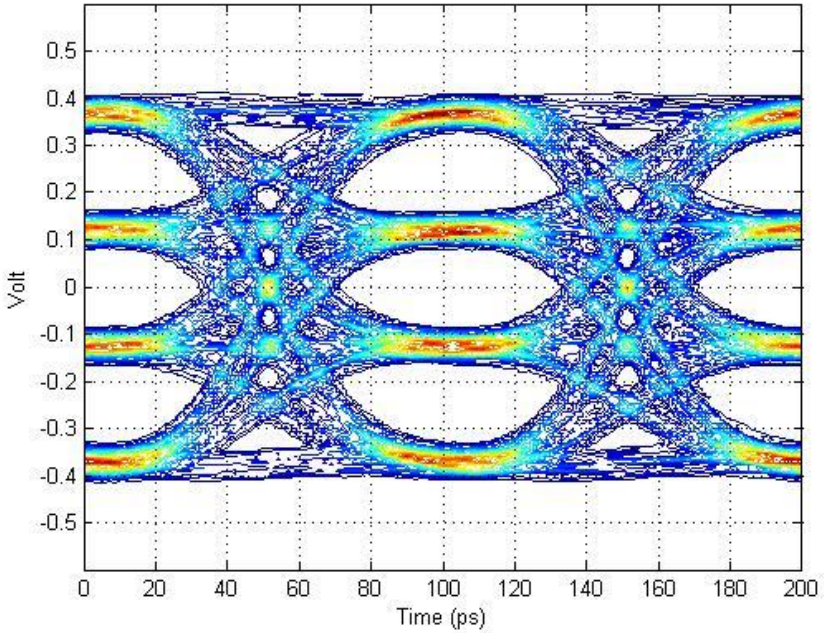

400G

QSFP56-DD & QSFP28-DD

- QSFP plus a second row of pins
- Same faceplate, slightly deeper
- backward compatible to QSFP+, QSFP28, QSFP56
- QSFP56-DD for 400G
  - 8 electrical lanes at 50G (56 w/ overhead)
- QSFP28-DD for 200G or 2x 100G
  - 8 electrical lanes at 25G (28 w/ overhead)
- Support breakout
- Cisco modules will be multi-sourced

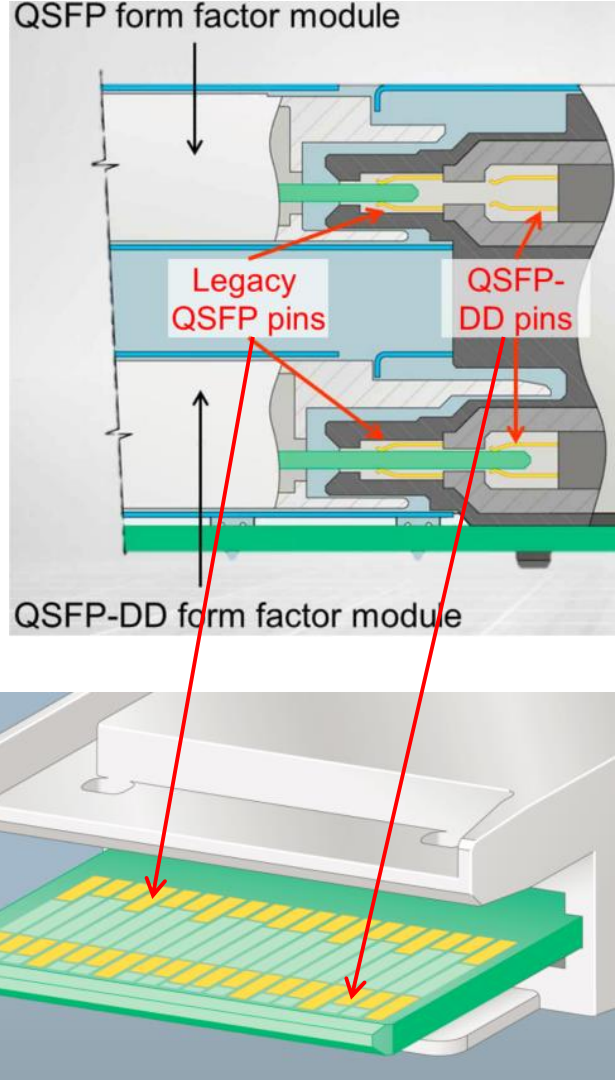

NCS5500 and 400G

All-in QSFP: QSFP56-DD

PMD	Reach	Media	Lasers	Modulation	λ
LR8	10km (6db)	Duplex SM	8	PAM4	1310nm
FR4	2km (5db)	Duplex SM	4	PAM4	1310nm
DR4	500m (4db)	PSM	4	PAM4	1310nm
ZR	40-80km	Duplex SM	1	DP 16QAM	1550nm
ZR+	Varies	Duplex SM	1	Varies	1550nm
DAC	3m	Copper	N/A	PAM4	N/A
AOC	100m	Fiber Cable	Black box	PAM4	1310nm

QSFP56-DD DR4 break-out to QSFP28/CPAK FR(1)

400G Breakout Options

- Today, 400G doesn't connect to existing 100G (25G based)
- 400GBASE-DR4 to 100GBASE-DR/FR Breakout (100G lambda)
- New 100G required: 1-lamba

Duplex LC SMF
Connector
1 lane of 1λ-100G-PAM4 signals

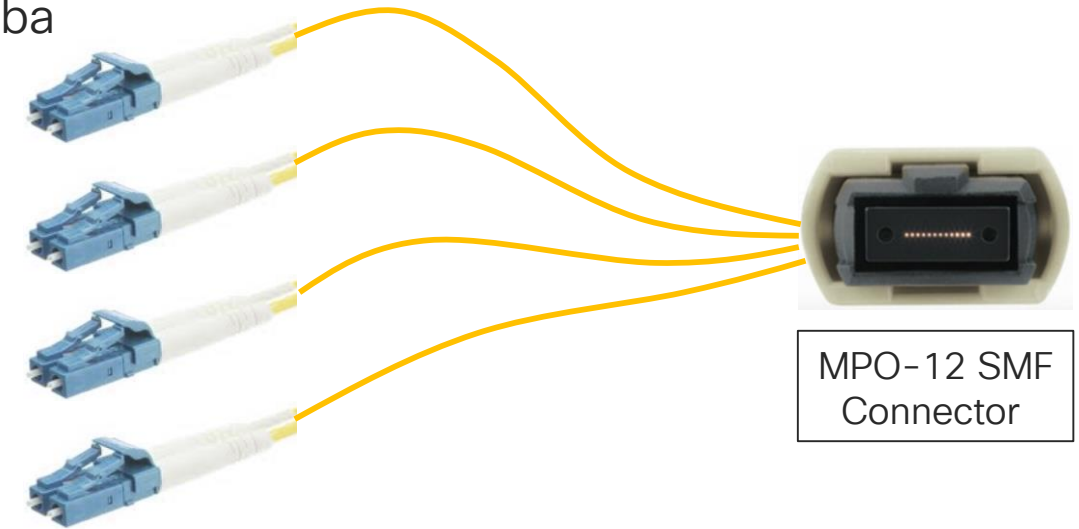

MPO-12 SMF
Connector

NCS5500 and 200G

All-in QSFP: QSFP28-DD

PMD	Reach	Media	Lasers	Modulation	λ
LR4	10km	Duplex SM	2x4	NRZ	1310nm
CWDM4	2km	Duplex SM	2x4	NRZ	1310nm
SR4	100m	Parallel MM	(2x4) MM	NRZ	850nm

QSFP28-DD optics are backward compatibility with current 100G optics generation (25Gbps-based)

200G Breakout Options

- Provides ability to connect legacy 100G modules

Module Type	Optical Connector
2x 100G-LR4	Dual Duplex CS Connector
2x 100G-CWDM4	Dual Duplex CS Connector
2x 100G-SR4	MMF MPO-24 Connector

NCS5500 Positioning

You make networking **possible**

Cisco live!

NCS5500 Position in Network

Multi-dimensional Equation

- The position decision of a platform should be based on:
  - Ports types / density requirement for X years
  - Scale requirements
  - Buffering capability
  - Supported features
  - Power consumption
  - Network OS preference (IOS XR)
- No simple rule of thumb

NCS5500 Position in Network

Think about…

- QoS
- ECMP-FEC
- Multi-Dimensional scale
- Counters
- Hw-profiles

For Reference

NCS5500 Platform Comparison

	NCS-5501	NCS-5501-SE	NCS-5502/-SE	NCS-5504	NCS-5508	NCS-5516
10G	48+6x4	40+4x4	48x4	4x36x4	8x36x4	16x36x4
25G	-	-	-	4x36x4	8x36x4	16x36x4
40G	6	4	48	4x36	8x36	16x36
100G	6	4	48	4x36	8x36	16x36
BW Gbps	800	800	4,800	14,400	28,800	57,600
Total Mpps	720	600	5,760	17,280	34,560	69,120
Power W	240	260	1,850	3,990	7,980	17,100
Pfx scale	1.1M+	2.75M	2.75M	Depends on LC (J/J+ w/ w/o eTCAM)		
100G	6	4	48	4x 36	8x 36	16x 36
Queues	96k					
Buffer	4GB per Forwarding ASIC					

NCS5500 Platform Comparison

	NCS-55A1-36H-S	NCS-55A1-36H-SE-S	NCS-55A2-MOD-S	NCS-55A2-MOD-SE-S
10G	36x4	36x4	40 (+2x 12)	40 (+2x 12)
25G	36x4	36x4	16	16+ 2x4
40G	36	36	8	8
100G	36	36	8	8
BW Gbps	3,600	3,600	1,440	1,440
Total Mpps	3,340	3,340	835	835
Power W	1,100	1,300	270 + 2x (50-75)	320 + 2x MPA
Pfx scale	1.1M+	4M	1.1M+	4M
100G	36	36	2x 4	2x 4
Queues	96k			
Buffer	4GB per Forwarding ASIC			

NCS5500 Platform Comparison

	NCS-55A1-24H	NCS-55A1-24Q6H-S	NCS-55A1-48Q6H
10G	24x4	48 + 6x4	48 + 6x4
25G	24x4	24 + 6x4	48 + 6x4
40G	24	6	6
100G	24	6	6
BW Gbps	2,400	1,440	1,440
Total Mpps	1,670	835	1x670
Power W	600	360	550
Pfx scale	2.2M+	1.1M+	2.2M+
Queues	96k		
Buffer	4GB per Forwarding ASIC		

VOQ and Life of a Unicast Packet

You make networking **possible**

Cisco live!

NCS5500 Architecture

Local Routing

- Local traffic on NCS5500 series can be routed by the FA without going through the fabric: lower latency

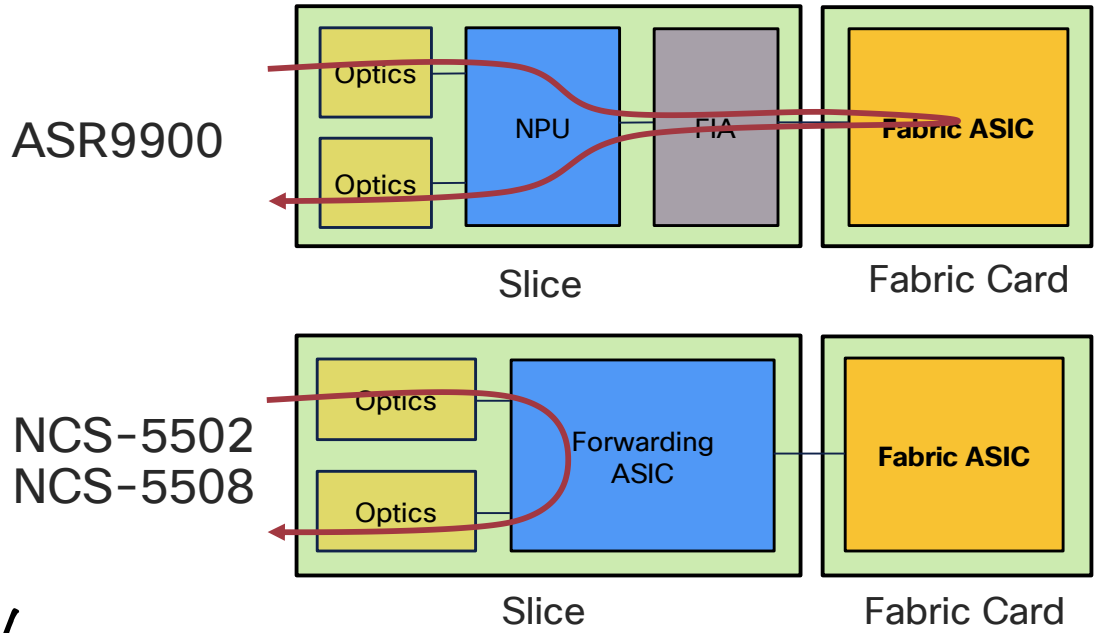

NCS5500 Architecture

Comparison with Traditional XR Platforms

- Two-lookup architecture on traditional XR platforms

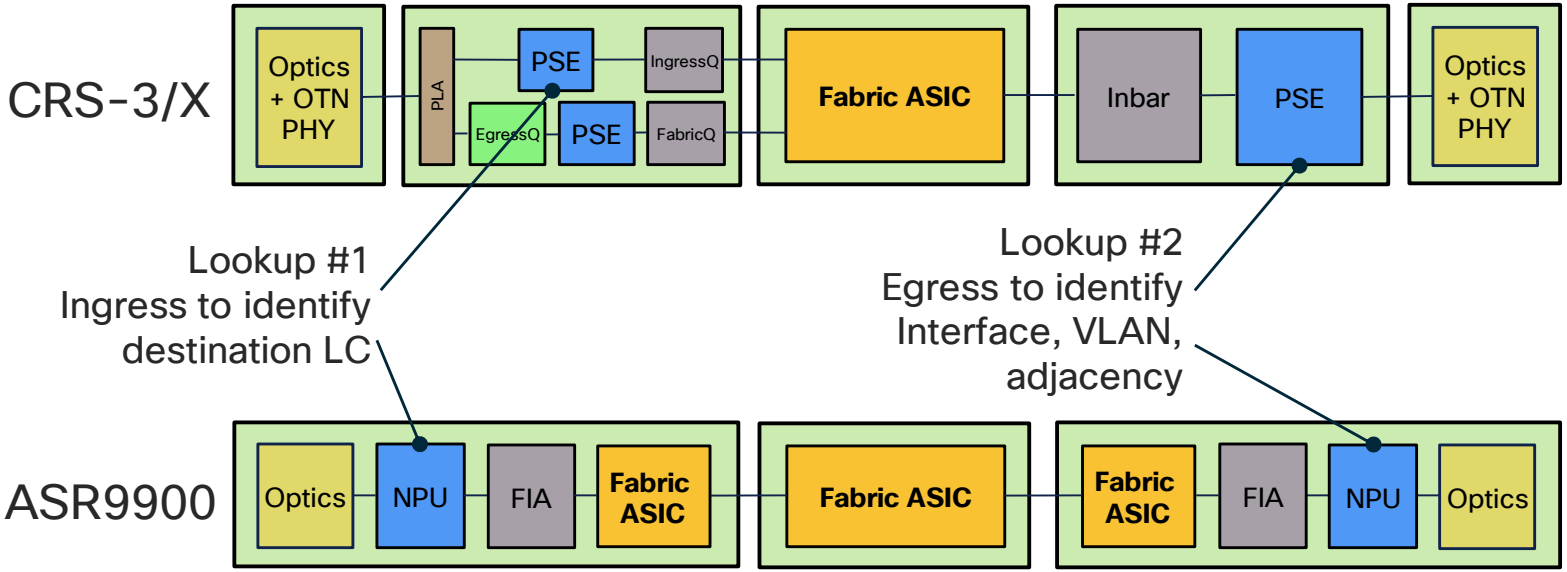

NCS5500 Architecture

Comparison with Traditional XR Platforms

- ASR9K: Buffering in two places but mostly in egress

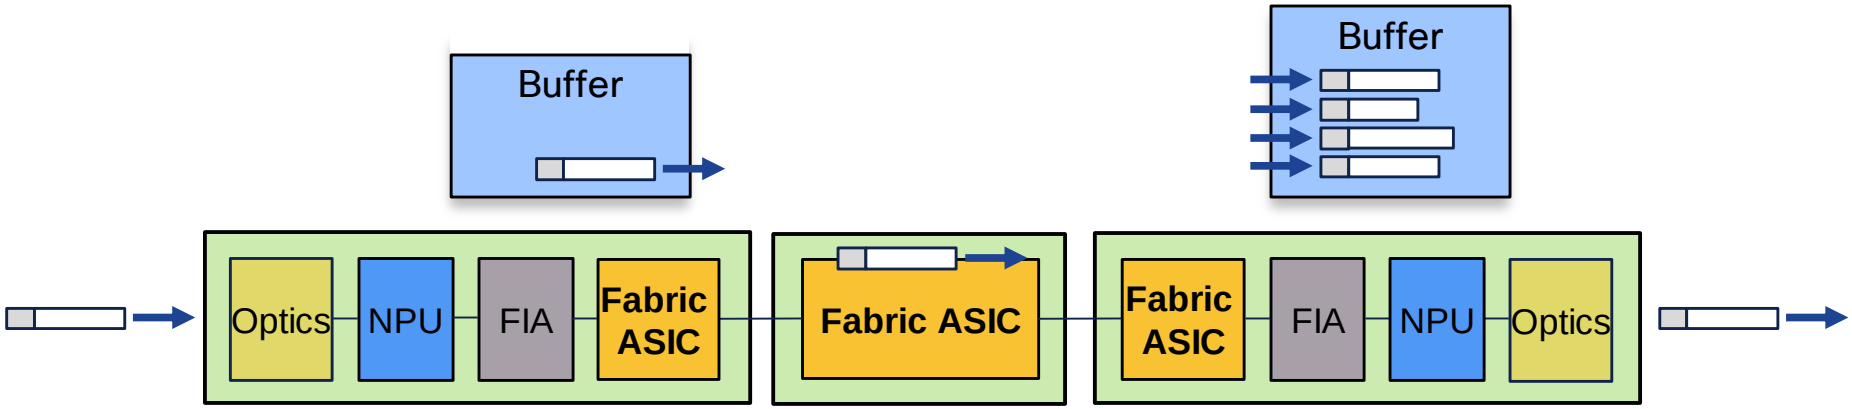

NCS5500 Architecture

Comparison with Traditional XR Platforms

- Single-lookup architecture at ingress on NCS5500
  - VOQ-only Model

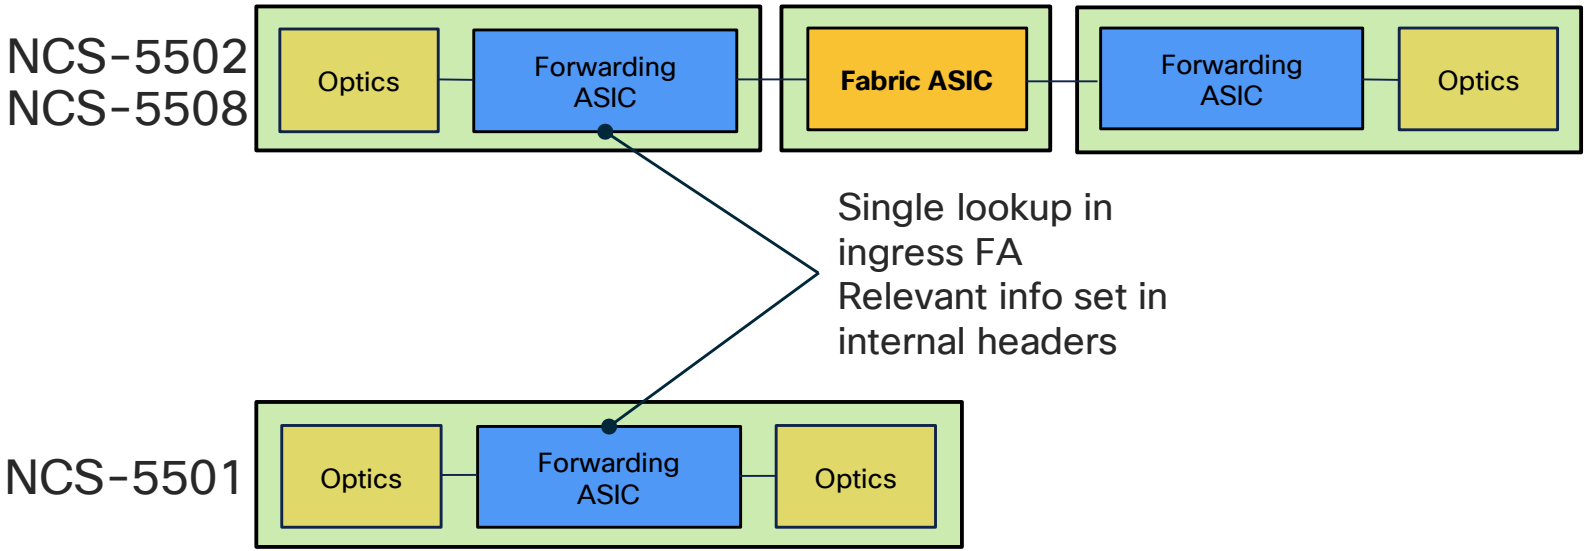

NCS5500 Architecture

Comparison with Traditional XR Platforms

- NCS5500 is using ingress buffering

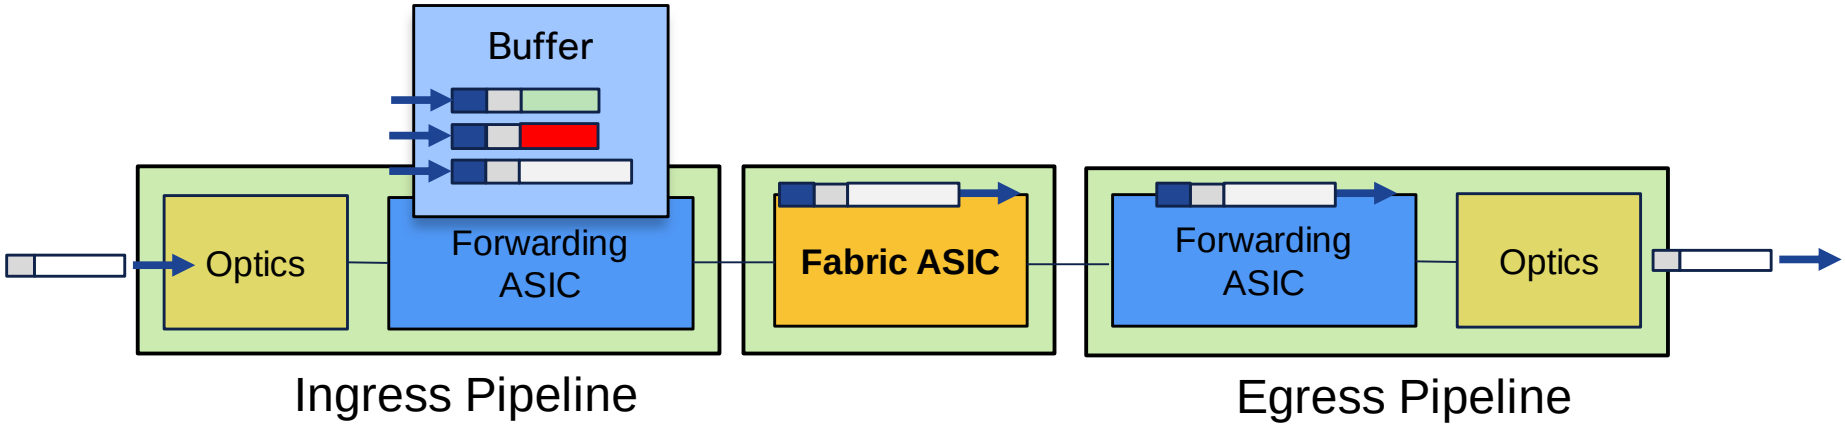

NCS5500 System Architecture

Three Packet Buffers / Hybrid Model

- Ingress On-chip Buffer: 16MB
- Ingress Off-chip Buffer: 4GB
- Egress On-chip port Buffer: 6MB (3MB per Core)

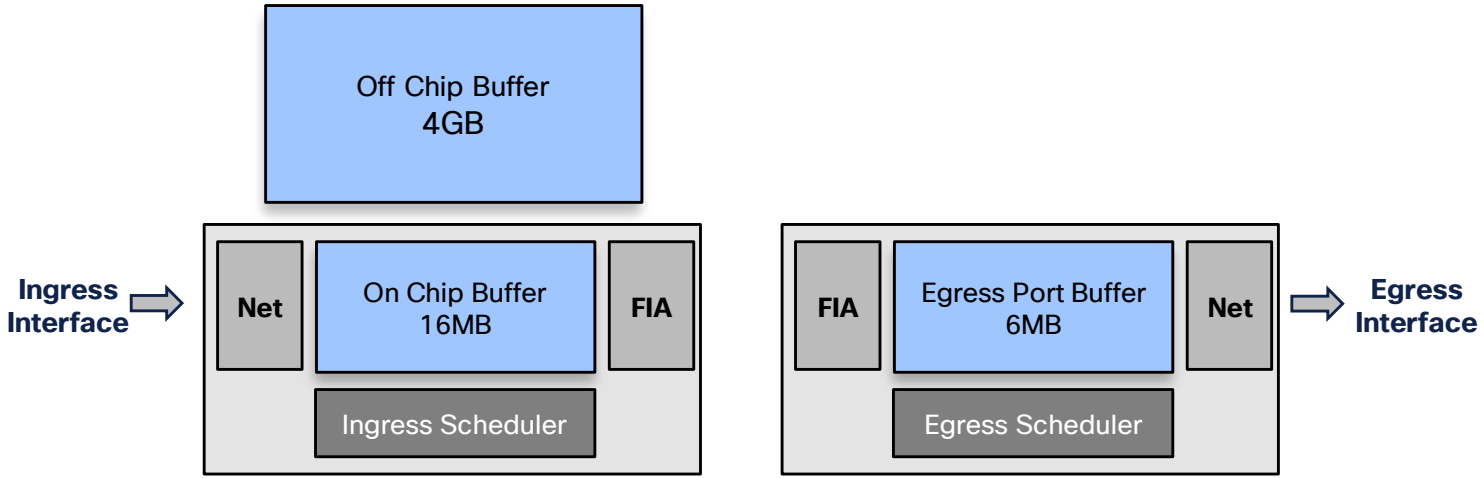

NCS5500 System Architecture

Three Packet Buffers / Hybrid Model

- Normal traffic condition (no congestion)
  - Packets stored in on-chip buffers only
  - That's the 99.85% of the packets

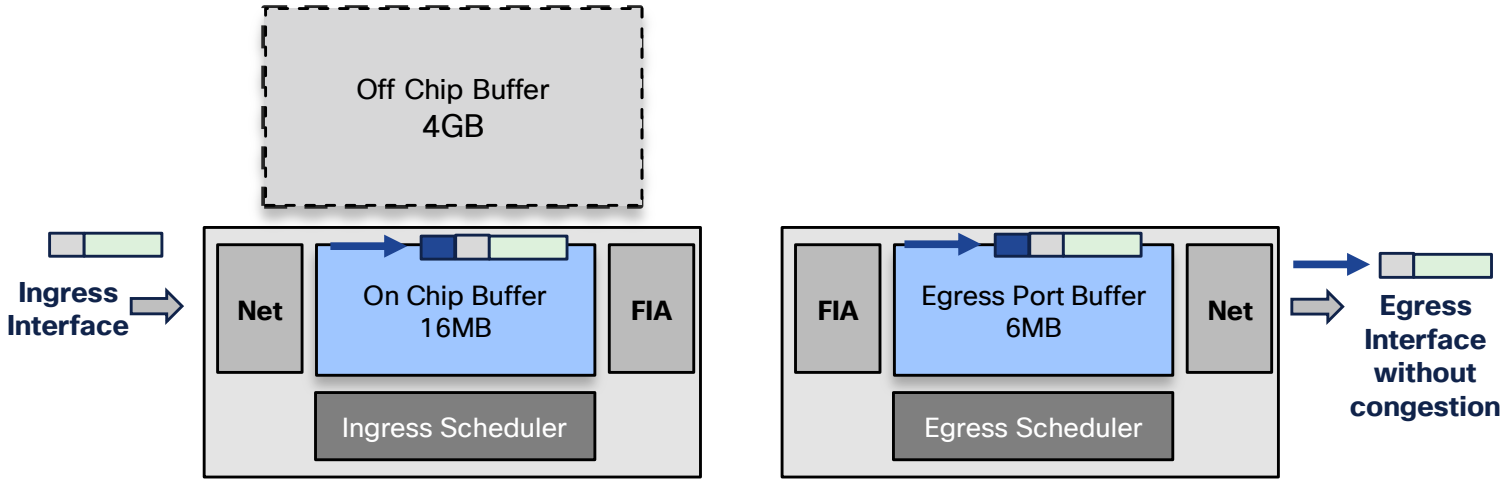

NCS5500 System Architecture

Three Packet Buffers / Hybrid Model

- In case of egress queue congestion
  - Packets stored in ingress off-chip buffers until they receive permission

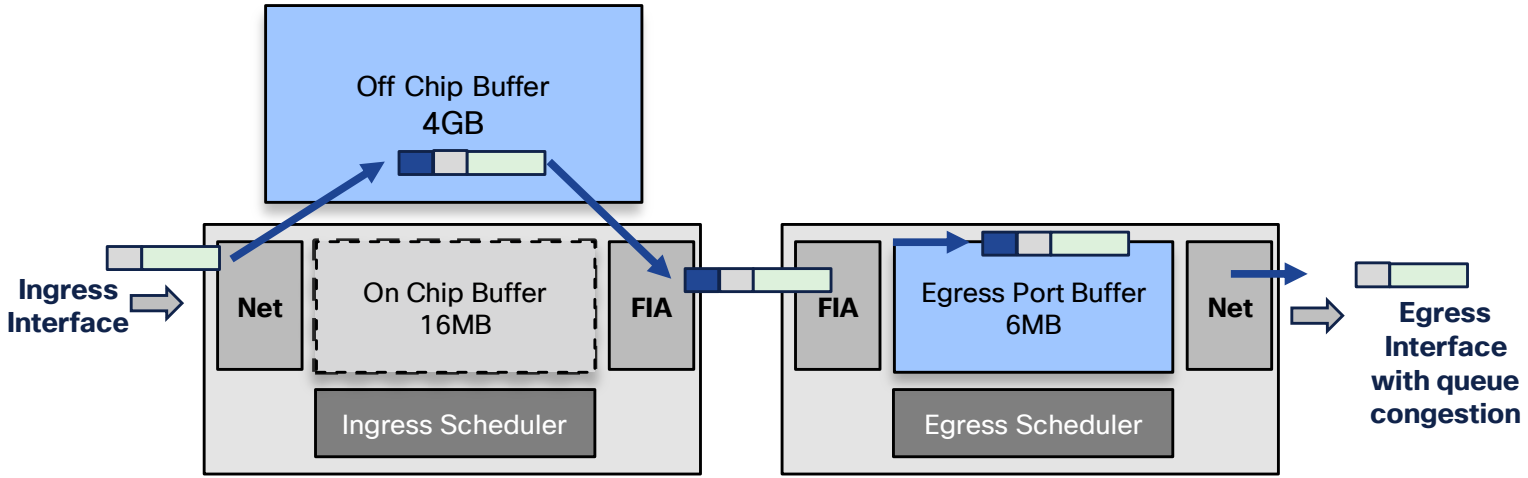

NCS5500 System Architecture

Three Packet Buffers / Hybrid Model

- Eviction to DRAM
  - Per virtual output queue

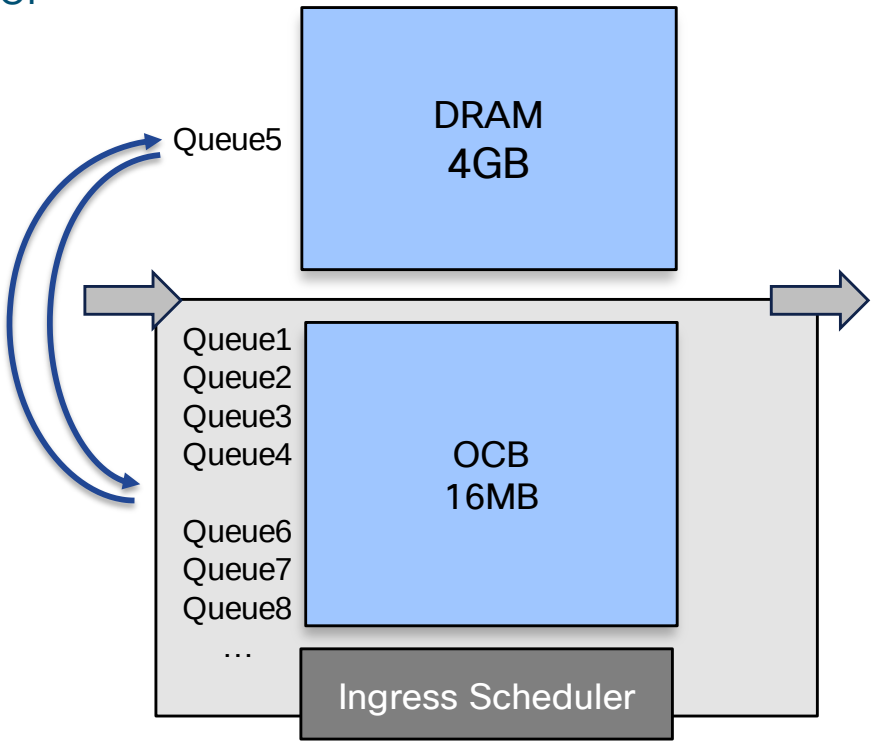

NCS5500 System Architecture

Three Packet Buffers / Hybrid Model

- Contrary to traditional XR platforms: very short egress buffering
  - 4 priorities on the egress port buffer
    - High Unicast
    - High Multicast
    - Low Unicast
    - Low Multicast
  - High >> Low
  - In case of tie-break
    - 80% Unicast
    - 20% Multicast

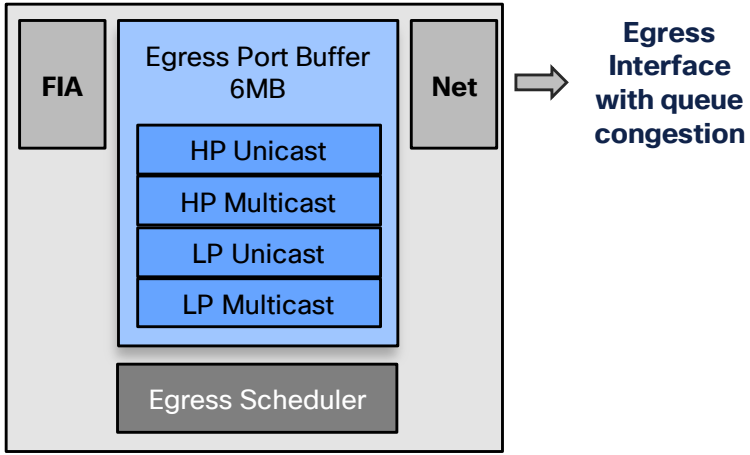

NCS5500 System Architecture

VOQ-Only Architecture (Virtual Output Queues)

- We have 8 queues per attachment point
- Attachment points are
  - L2/L3 interfaces (physicals, bundles, BVI, ...)
  - Sub-interfaces (L2/L3)
- Example here:
  - Hu0/7/0/0.2 dot1q interface

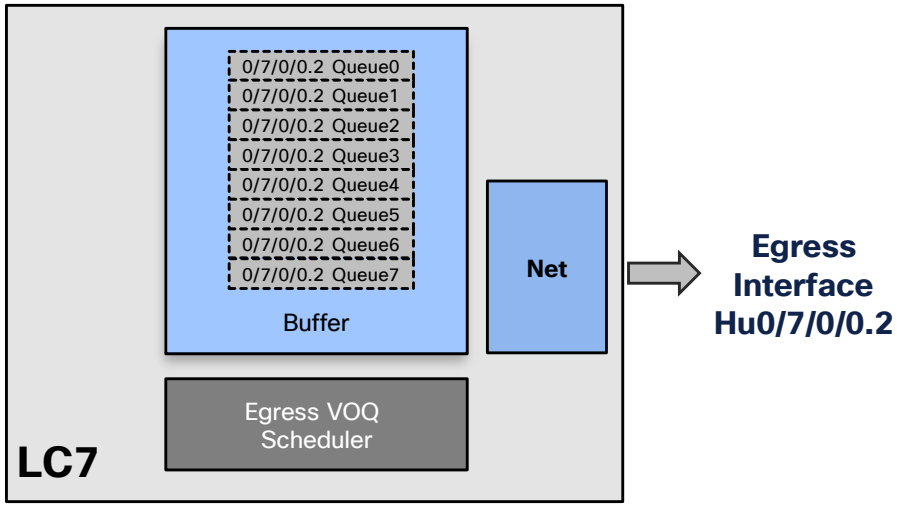

NCS5500 System Architecture

VOQ-Only Architecture (Virtual Output Queues)

- Every NPU will have a logical (virtual) representation of these egress queue locally where packets are actually stored in congestion situation → VOQ

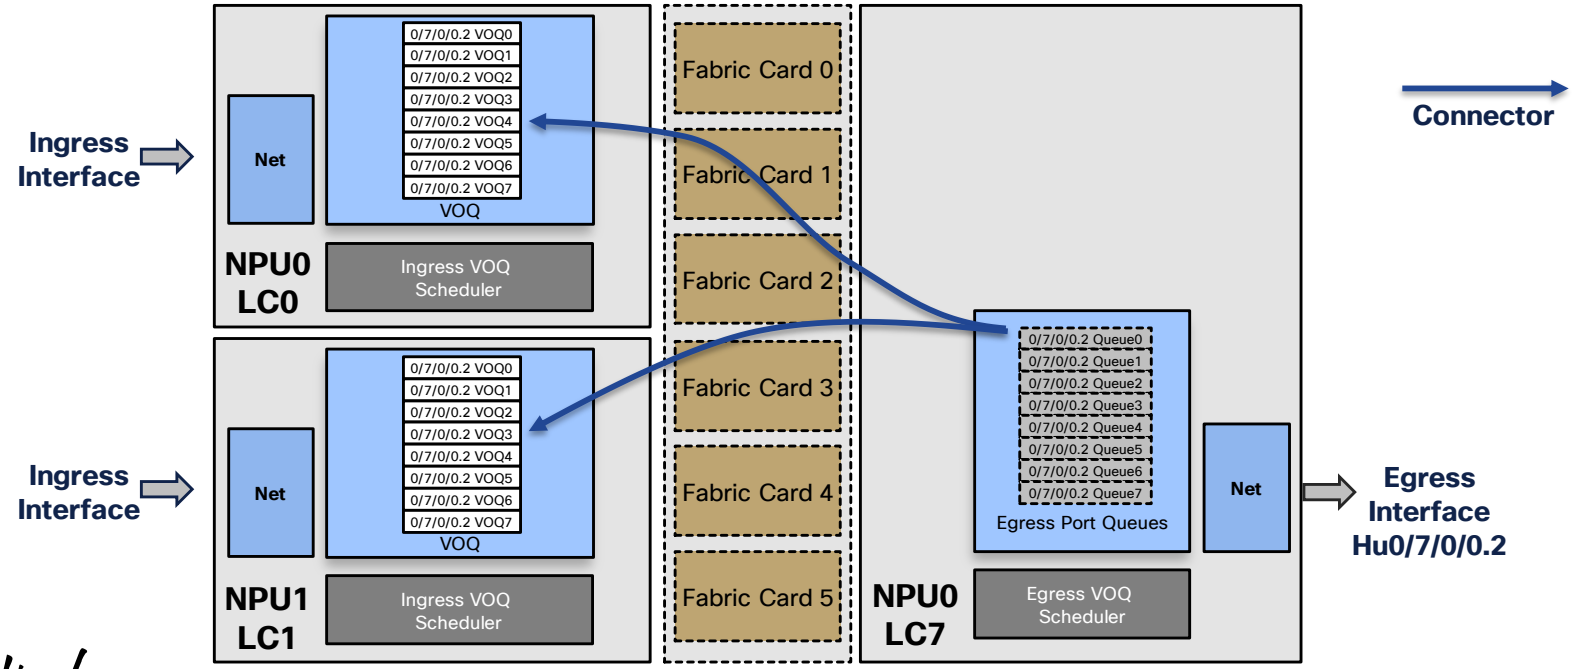

NCS5500 System Architecture

VOQ-Only Architecture (Virtual Output Queues)

- Even for the same NPU 0 on the same LC7, the ingress pipeline uses this virtual representation (Local VOQ)

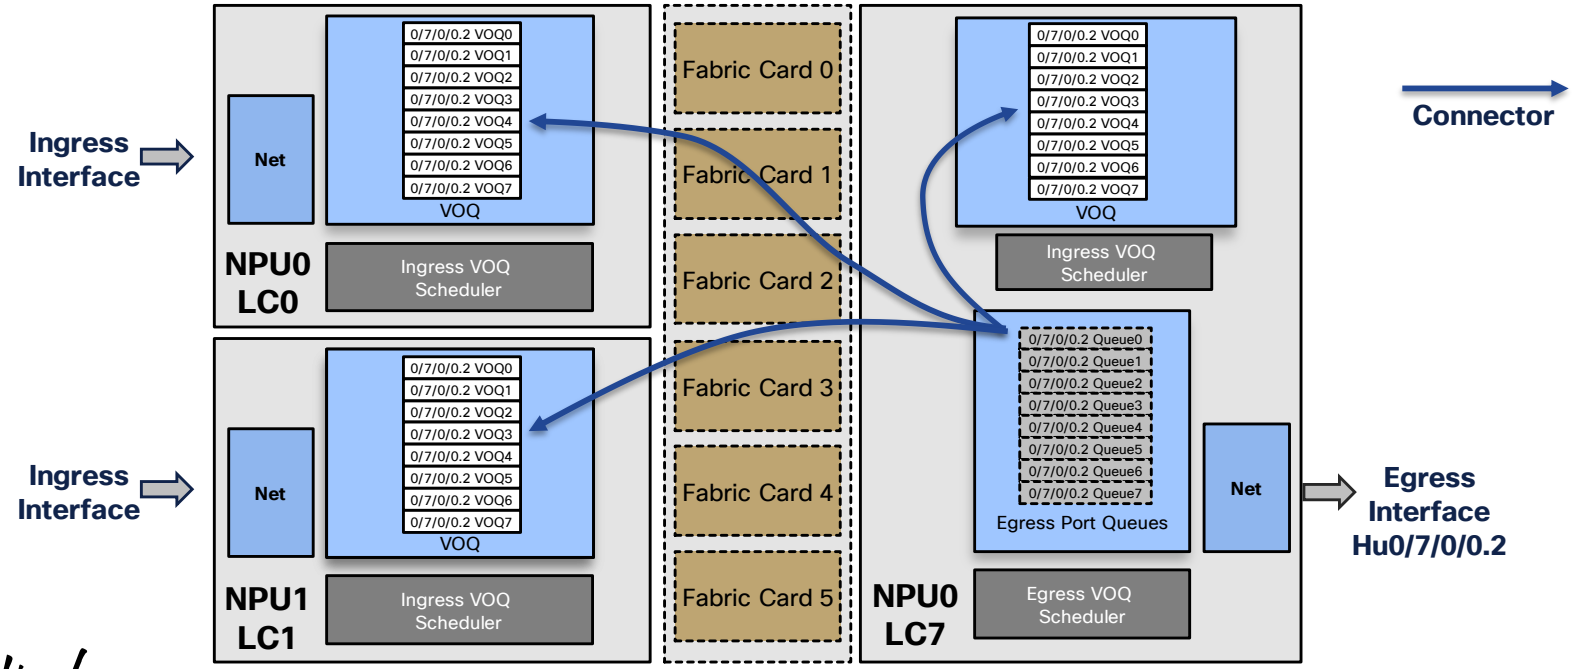

NCS5500 System Architecture

VOQ-Only Architecture (Virtual Output Queues)

- CLI illustration: Local and Remote visibility of the Output Queues

```
RP/0/RP0/CPU0:NCS5508-1_PE1#sh contr npu voq-usage interface all instance 0 location 0/0/CPU0

--------------------------------------------------------------------
Node ID: 0/0/CPU0
Intf          Intf      NPU NPU   PP    Sys    VOQ    Flow    VOQ     Port
name          handle     #  core Port  Port   base   base    port    speed
              (hex)                                      type    (Gbps)
--------------------------------------------------------------------------
Hu0/3/0/5     1800100    0   0     1   1537    1072   10280 remote  100
Hu0/0/0/26    200        4   1    17    273    1424    4136 local   100
Hu0/3/0/6     1800108    1   1    21   1621    1080    1064 remote  100
Hu0/0/0/27    208        4   0     9    265    1432    5416 local   100
Hu0/3/0/7     1800110    1   1    13   1613    1088    2344 remote  100
Hu0/0/0/28    210        4   0     5    261    1440    7208 local   100
Hu0/3/0/8     1800118    1   1    17   1617    1096    4136 remote  100
Hu0/0/0/29    218        4   0     1    257    1448    8488 local   100
Hu0/3/0/9     1800120    1   0     9   1609    1104    5416 remote  100
Hu0/0/0/30    220        5   1    21    341    1456    2344 local   100
```

NCS5500 Forwarding ASIC Detail

Deep Buffer

- Expansion via off-chip resources
  - Deep GDDR5 packet buffers external packet buffers
- In normal conditions
  - Packets are stored in On-Chip Buffers only
- In case of egress congestion
  - Packets are moved to the Off-Chip Buffer in Virtual Output Queues
  - Packets are identified by packet descriptors
  - Each ASIC can manage 3M of these descriptors
- A single queue can take up to 25% of the 1.5M descriptors of a core
- Decision to move packets from on-chip to off-chip buffer is made (today)
  - When a queue exceeds 200kB
  - When a queue exceeds 6000 packets

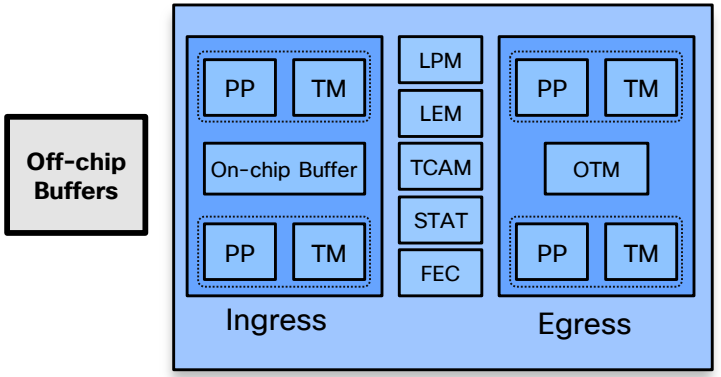

NCS5500 VOQ-Only Architecture

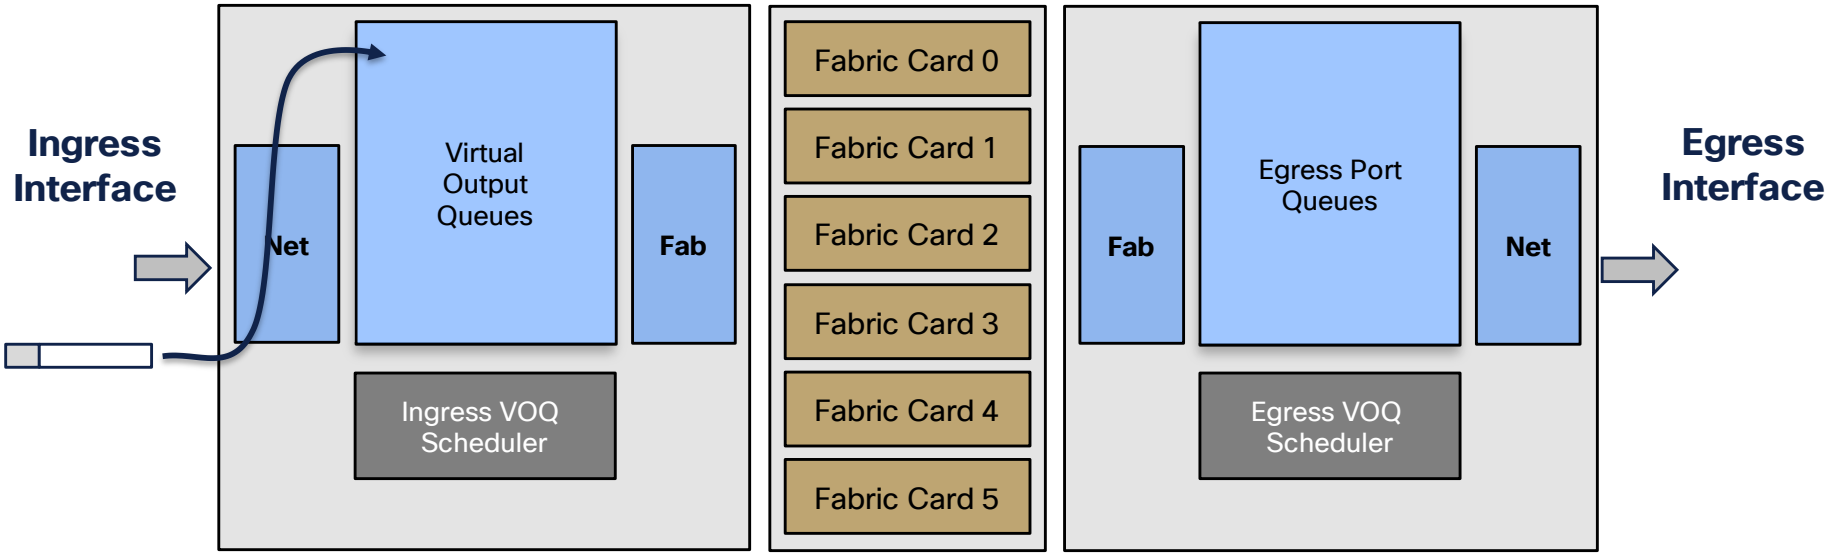

- Packet is received on ingress interface, classified, and stored in an internal buffer
- Single lookup
- Queuing is based on credit request and grant scheme
- Actual buffering happens on ingress devices

NCS5500 VOQ-Only Architecture

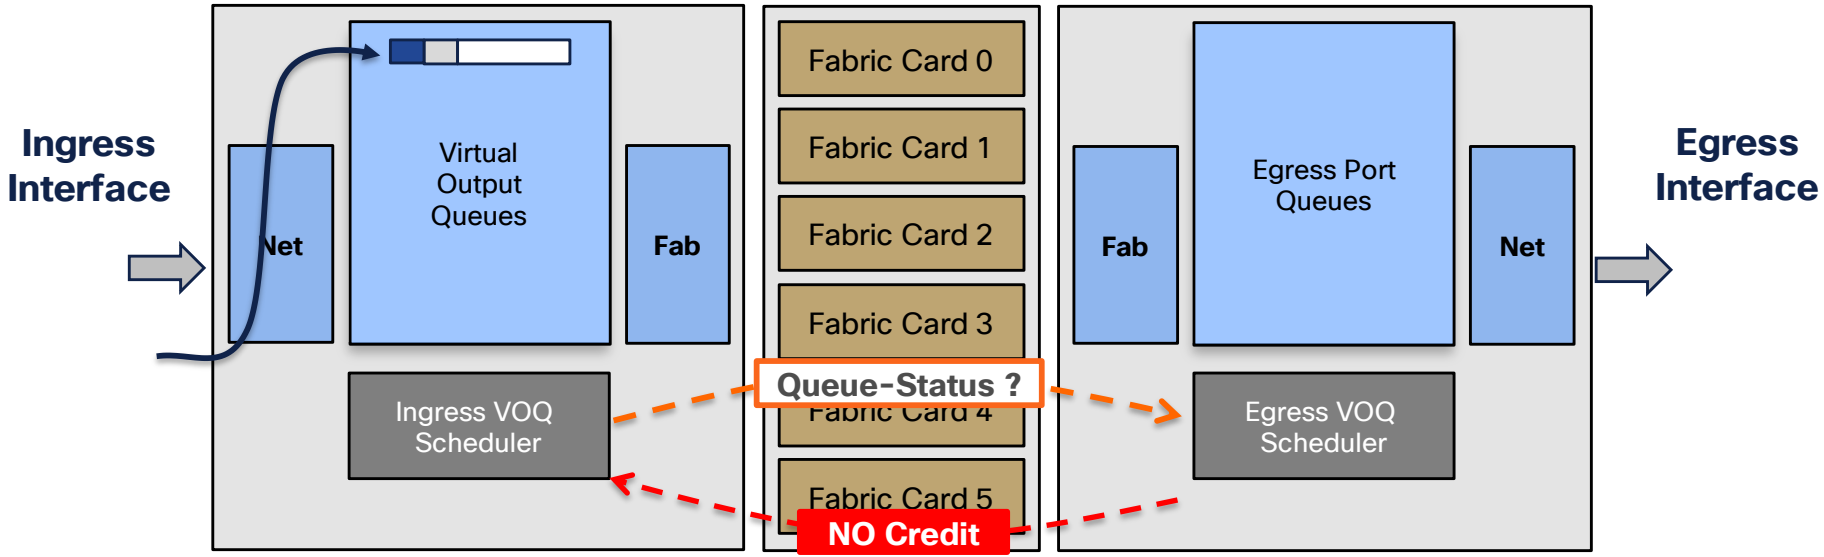

- Ingress VOQ scheduler polls Egress scheduler (maintaining a local VOQ DB)
- Egress answers with a credit-message (or not, in our example)
- Egress device decides how much traffic can be sent by granting credits to any ingress requesting Forwarding ASIC

NCS5500 VOQ-Only Architecture

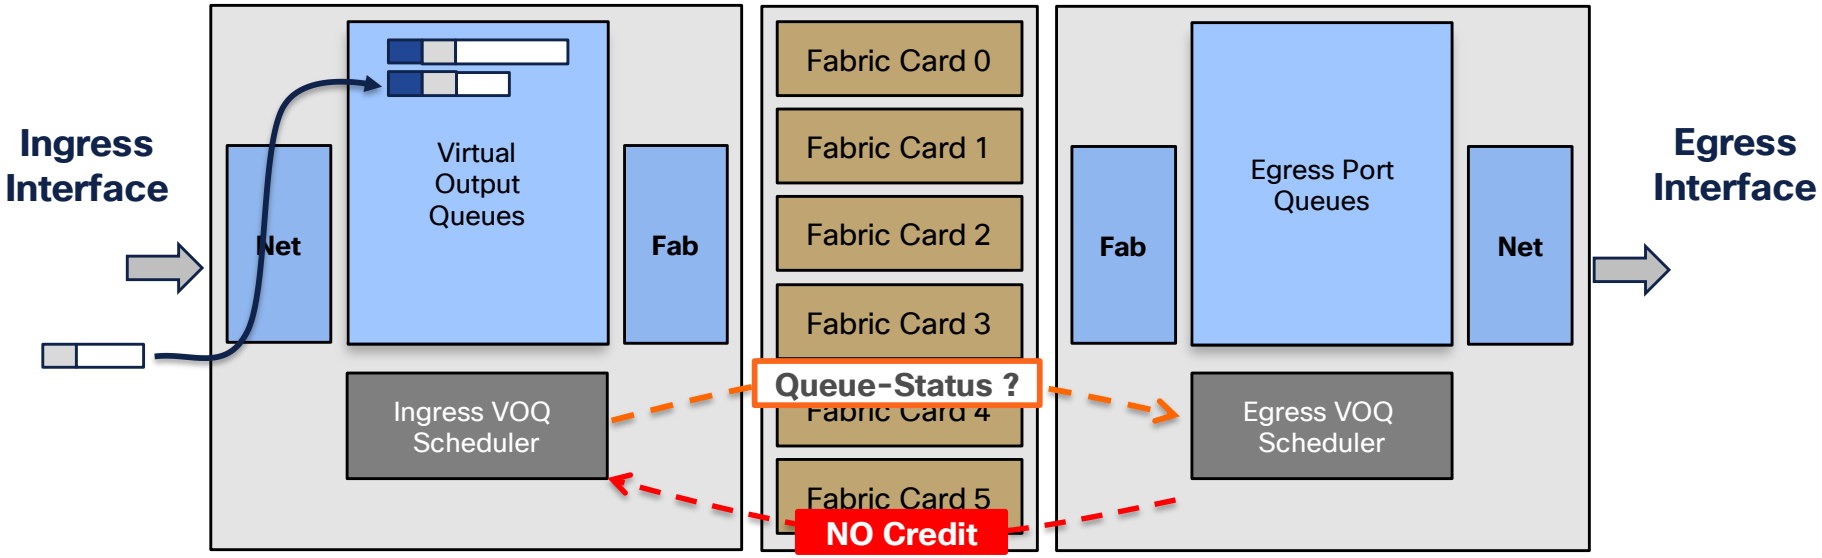

- Packets are piling up in the ingress buffer

NCS5500 VOQ-Only Architecture

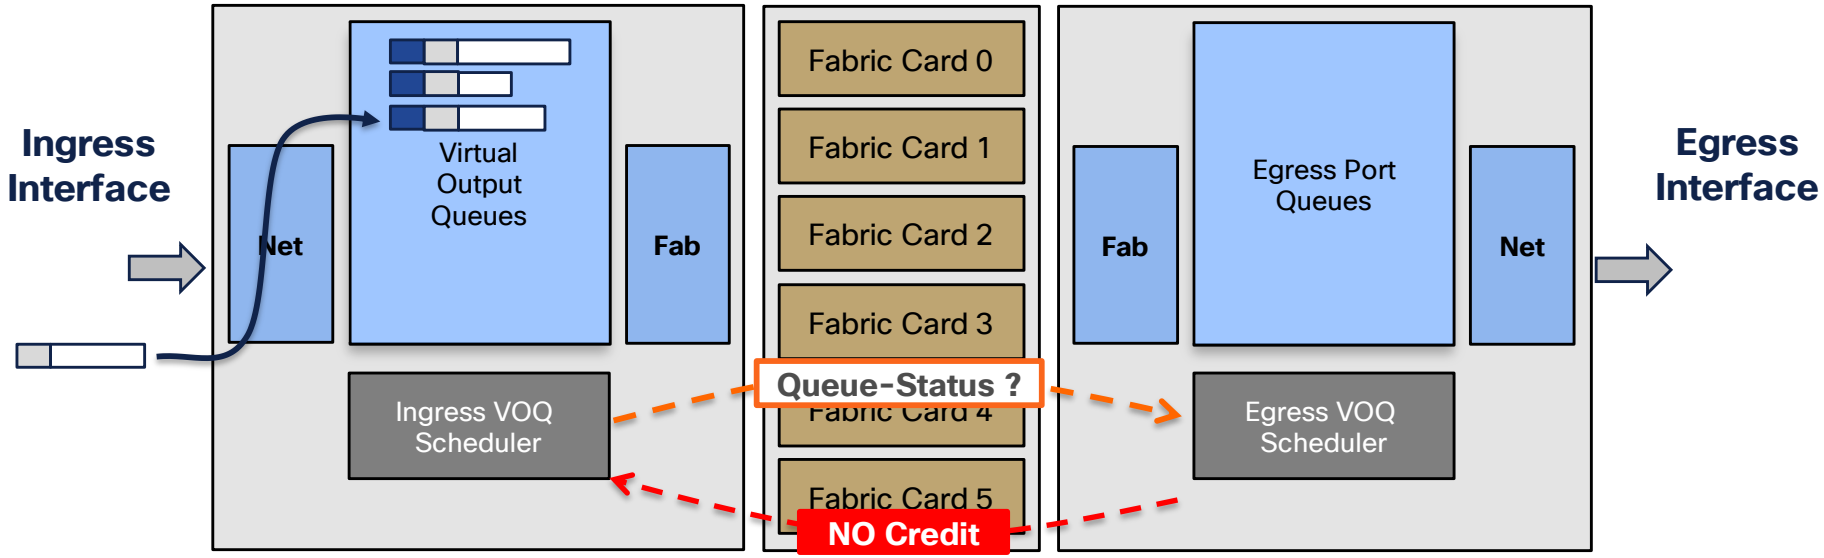

- Packets are piling up in the ingress buffer
- If a given queue size is exceeded, new packets are tail-dropped

NCS5500 VOQ-Only Architecture

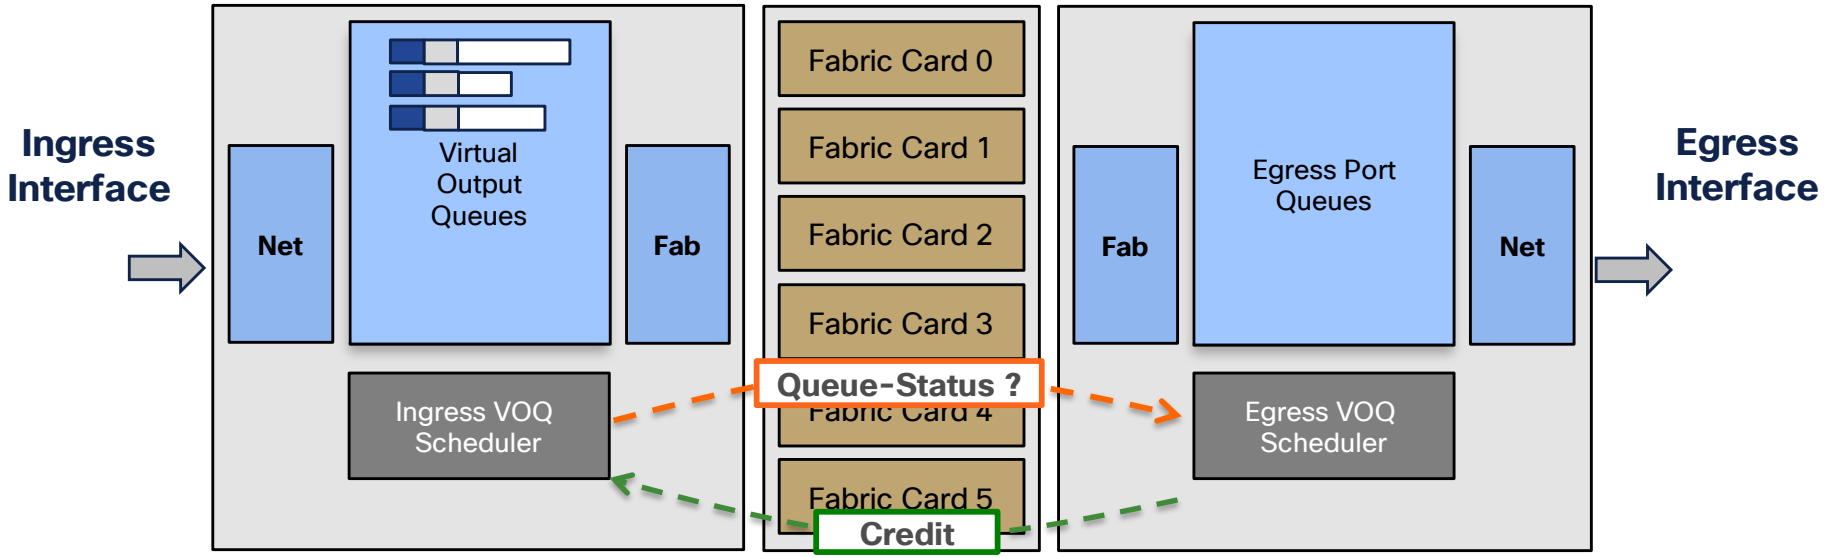

- Finally, the egress scheduler grants the credit for packet transmission

NCS5500 VOQ-Only Architecture

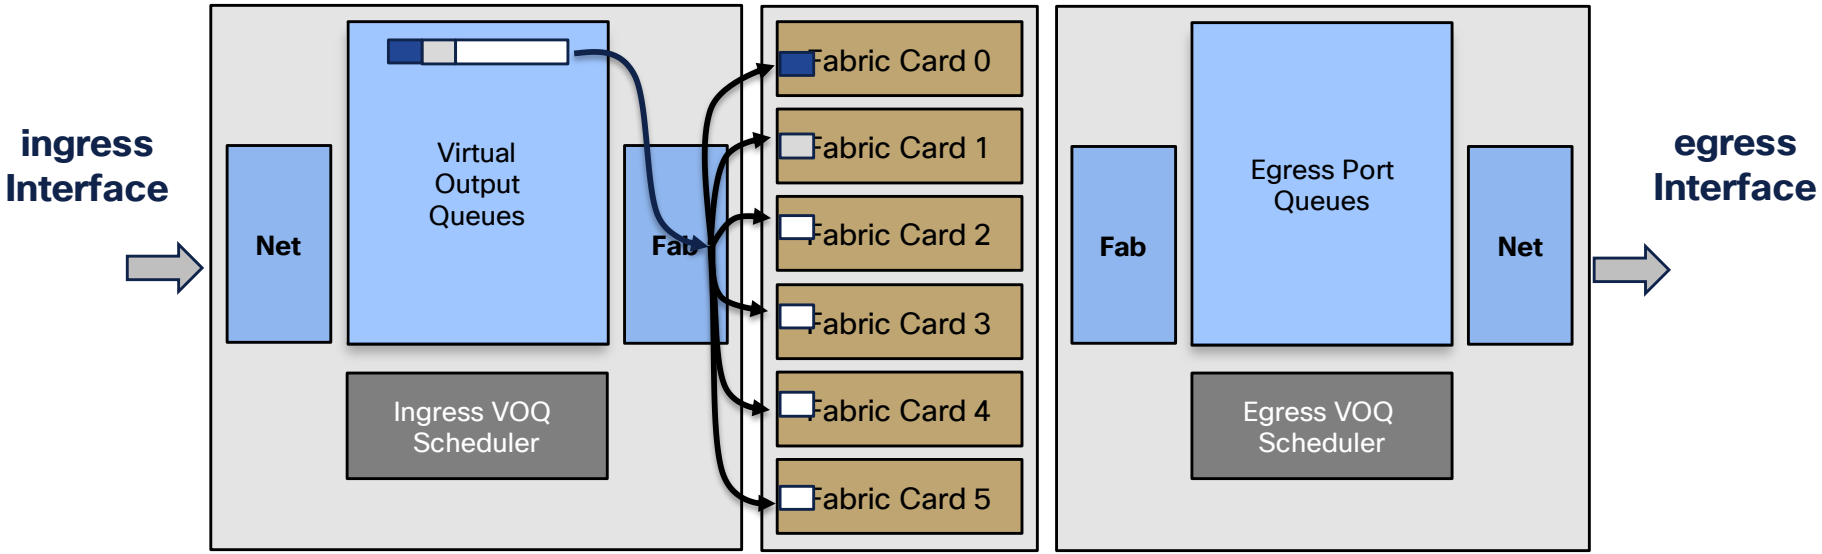

- Packet is split in cells and load balanced among the fabric cards
- Cells are transported to the egress line card

NCS5500 VOQ-Only Architecture

- Let's take the example of a 1400B packet

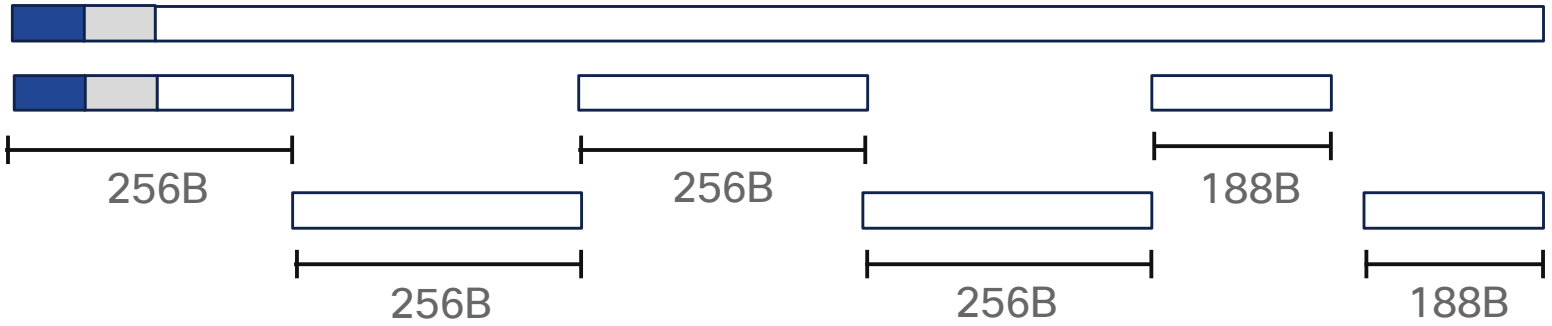

- If the last part is between 256B and 512B, we divide by 2

  1400 – 4 x 256 = 376 = 2 x 188

NCS5500 VOQ-Only Architecture

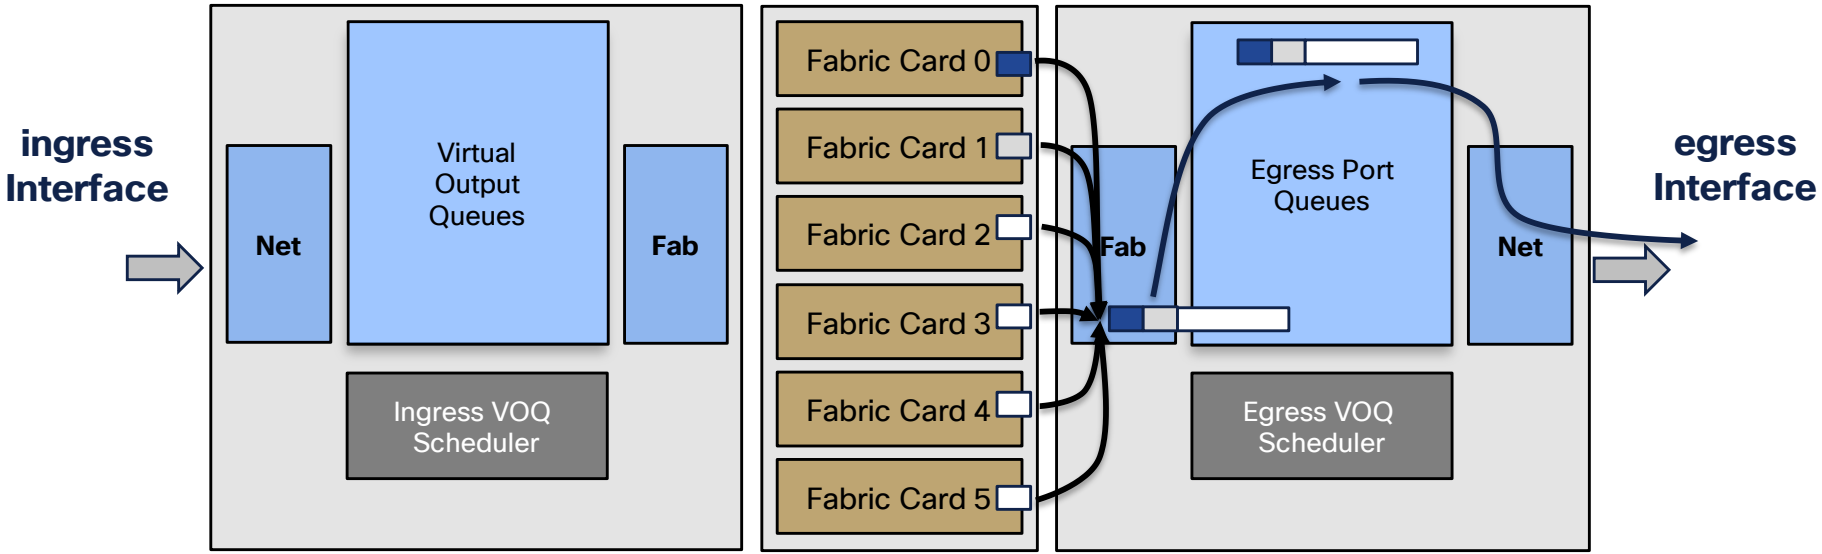

- Cells are collected and packet re-assembled
- Packet is stored in the port queue
- Finally packet is transmitted through the egress interface

NCS5500 VOQ-Only Architecture in SoCs

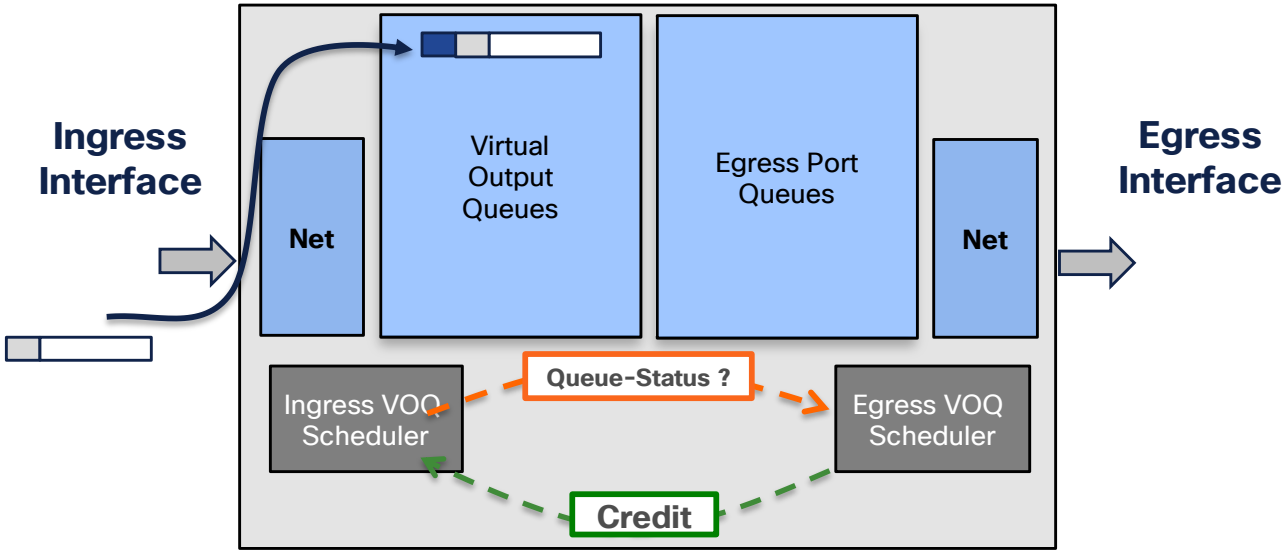

- Packet is received on ingress interface, classified, and stored in internal buffer
- Ingress VOQ scheduler polls Egress scheduler (maintaining a local VOQ DB)
- Egress answers with a credit-message

NCS5500 VOQ-Only Architecture in SoCs

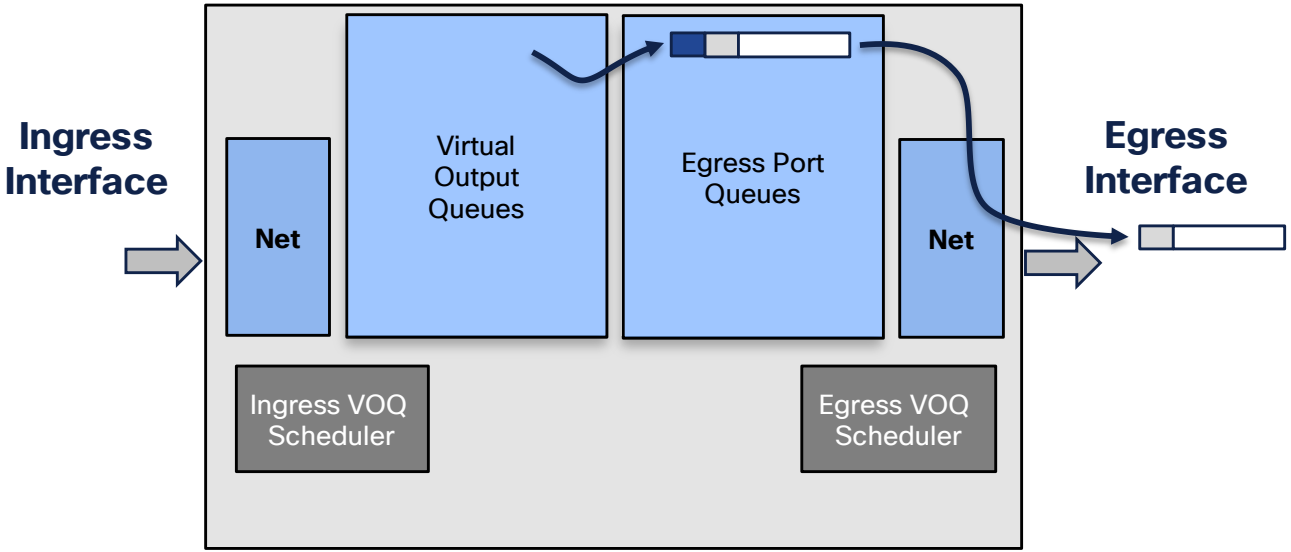

- Packet is stored in the port queue
- Finally packet is transmitted through the egress interface

FMQ and Life of a Multicast Packet

You make networking **possible**

Cisco live!

Multicast in NCS5500

- (S,G) information stored in LPM and takes one entry each
  - IPv4 key (VRF, S, G)
  - IPv6 key (VRF, G)
- MCID / FGID
- Replication performed at two levels
  - Fabric level
  - Egress Forwarding ASIC level

NCS5500 System Architecture

Control Plane

- IGMP and PIM joins are punted to RP CPU process (igmp/pim)
  - Packets use EPC internal network to reach the process executed on RP LXC

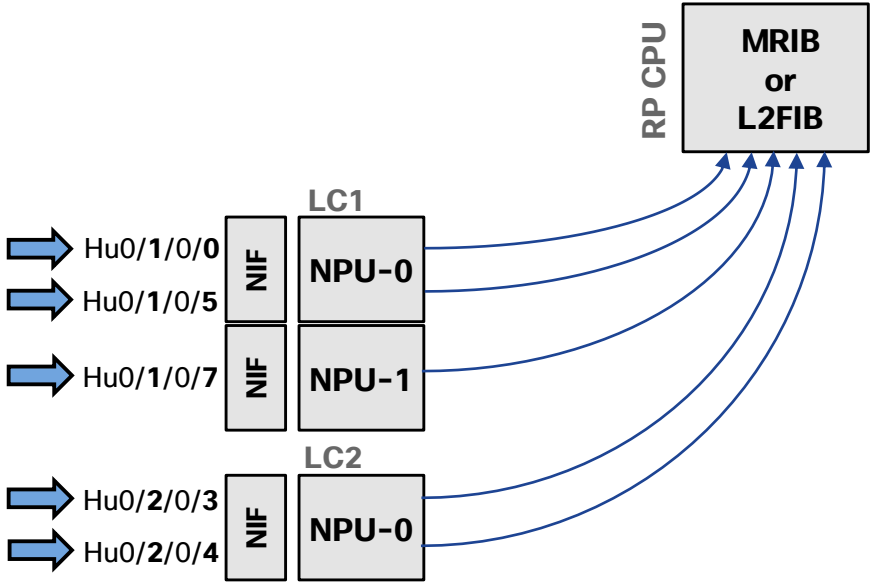

NCS5500 System Architecture

Control Plane

- If it's a new group, the process (MRIB or L2FIB) will allocate a Multicast ID (MCID)
- If a MCID is already allocated, information will be updated based on join/leave

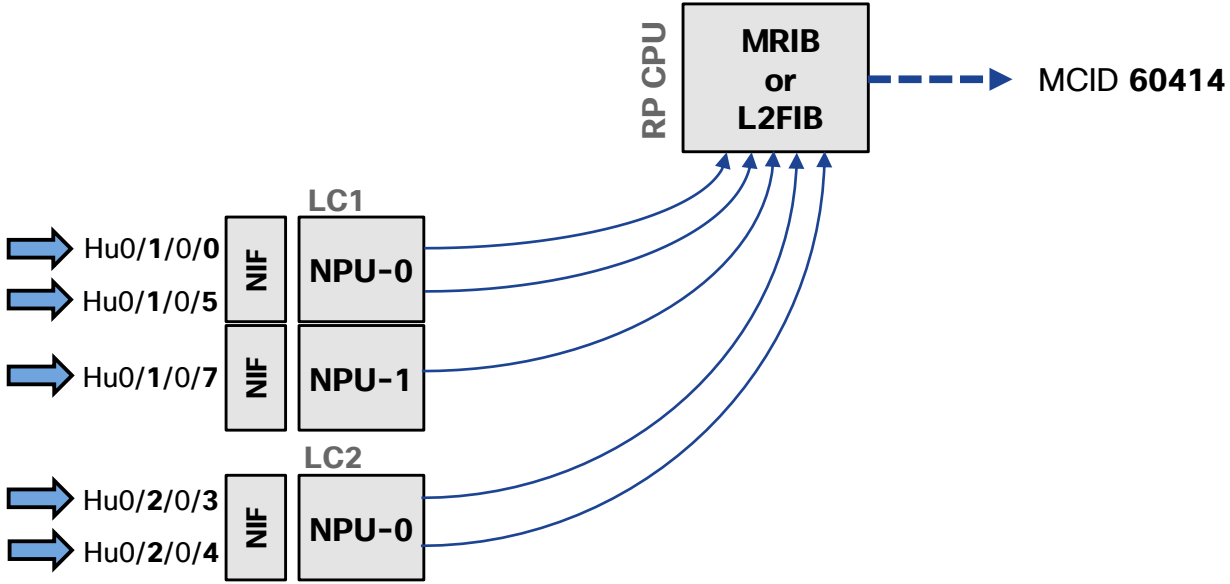

NCS5500 System Architecture

Control Plane: Identifying MCID

- MCID is often times referred as FGID internally
  - You can find the MCID associated to a (*,G) or (S,G) pair with the following CLI:

```
RP/0/RP0/CPU0:Router#sh mrib route 50.41.13.11 232.31.0.12 detail

IP Multicast Routing Information Base
<SNIP>

(50.41.13.11,232.31.0.12) Ver: 0xef18 RPF nbr: 16.2.4.1 Flags: RPF, FGID: 9155
  Up: 04:20:11
  Incoming Interface List
    Bundle-Ether162.4 Flags: A, Up: 04:20:11
  Outgoing Interface List
    Bundle-Ether361.6 Flags: F NS, Up: 04:20:11
RP/0/RP0/CPU0:Router#
```

NCS5500 System Architecture

Control Plane

- The process running on RP CPU will dynamically compute two tables for each MCID
- MCID mapping is a 128 bitmap mask where Ones represent NPUs who received a join and who expect a copy of the packet from the fabric
- MCID-DB associates ports where a replication is expected

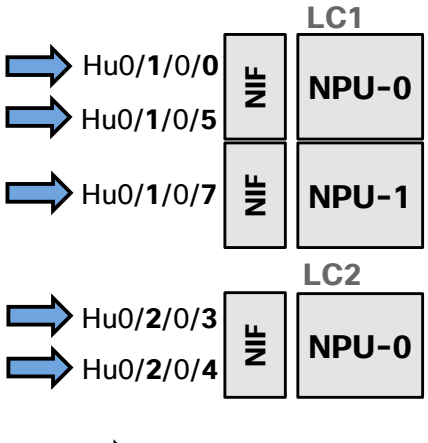

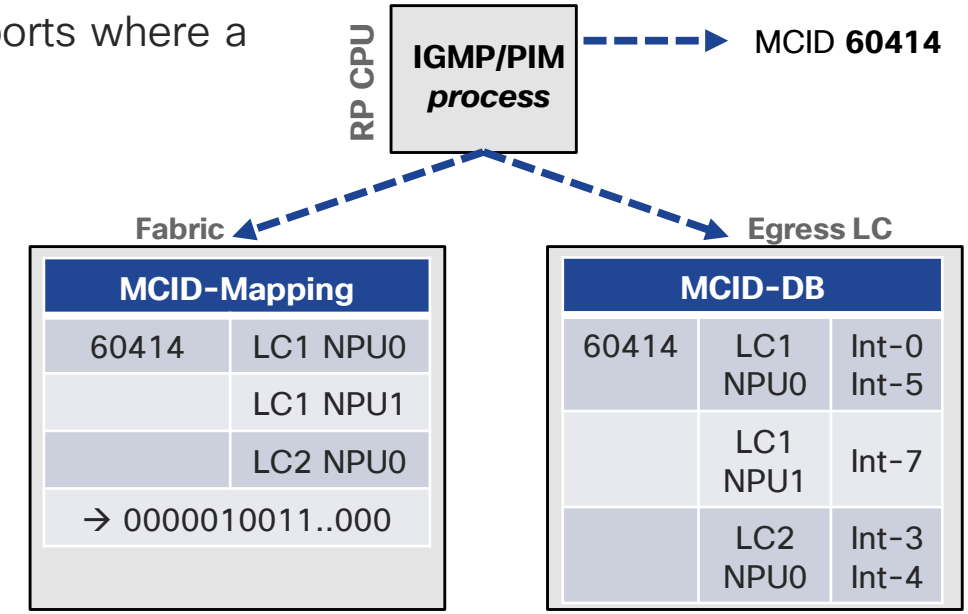

NCS5500 System Architecture

Show Commands

```
RP/0/RP0/CPU0:ios#show mrib route detail
<SNIP>
(25.1.1.2,232.1.1.4) Ver: 0x6632 RPF nbr: 25.1.1.2 Flags: RPF, FGID: 3177
  Up: 2w4d
    Incoming Interface List
      BVI1 Flags: A, Up: 2w4d
    Outgoing Interface List
      TenGigE0/3/0/3/0.100 Flags: F NS LI, Up: 2w4d
RP/0/RP0/CPU0:ios#
RP/0/RP0/CPU0:ios#show mfib route 232.1.1.4 location 0/3/CPU0

(25.1.1.2,232.1.1.4),  Flags:
  Up: 2w4d
  Last Used: never
  SW Forwarding Counts: 0/0/0
  SW Replication Counts: 0/0/0
  SW Failure Counts: 0/0/0/0/0
  TenGigE0/3/0/1/0.100 Flags:  A, Up:2w4d
  TenGigE0/3/0/2/0.200 Flags:  NS EG, Up:2w4d
```

NCS5500 System Architecture

MCID Bitmap

```
RP/0/RP0/CPU0:ios#show mrib fgid info 3177

FGID information
----------------
FGID (type)      : 3177 (Primary)
Context          : IP (0xe0000000, 25.1.1.2, 232.1.1.4/32)
Members[ref]     :  0/3/0[1]
LineCard Slot    :  3 ::  Npu Instance 0
FGID bitmap
0x0000000000040000  0x0000000000000000  0x0000000000000000  0x0000000000000000
0x0000000000000000  0x0000000000000000  0x0000000000000000  0x0000000000000000
0x0000000000000000  0x0000000000000000  0x0000000000000000  0x0000000000000000
0x0000000000000000  0x0000000000000000  0x0000000000000000  0x0000000000000000
FGID chkpt context valid : TRUE
FGID chkpt context :
table_id 0xe0000000 group 0xe8010104/32 source 0x19010102
FGID chkpt info : 0x23000000
Fgid in batch          : NO
Secondary node count   : 0
RP/0/RP0/CPU0:ios#
```

NCS5500 System Architecture

MCID Bitmap

```
sysadmin-vm:0_RP0# show controller fabric fgid information id 10927 detail

Displaying FGID: 10927
FGID Information:

FGID number:            10927

FGID Hex bitmap:
        0x00001fffc0000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000
        0x0000000000000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000
        0x0000000000000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000
        0x0000000000000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000
        0x0000000000000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000 0000000000000000-0000000000000000

FGID Binary bitmap:
        0000000000000000000111111111111111000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
        0000000000000000000000000000000000000000000000000000000000000000
```

NCS5500 System Architecture

Data Plane

- Multicast Packet is received on ingress interface
- Lookup provides a FEC-ID itself pointing to MCID
  - In LPM for L3 packets (we will use it as an example)
  - In iTCAM for L2 packets (future plans to move them to LPM too)

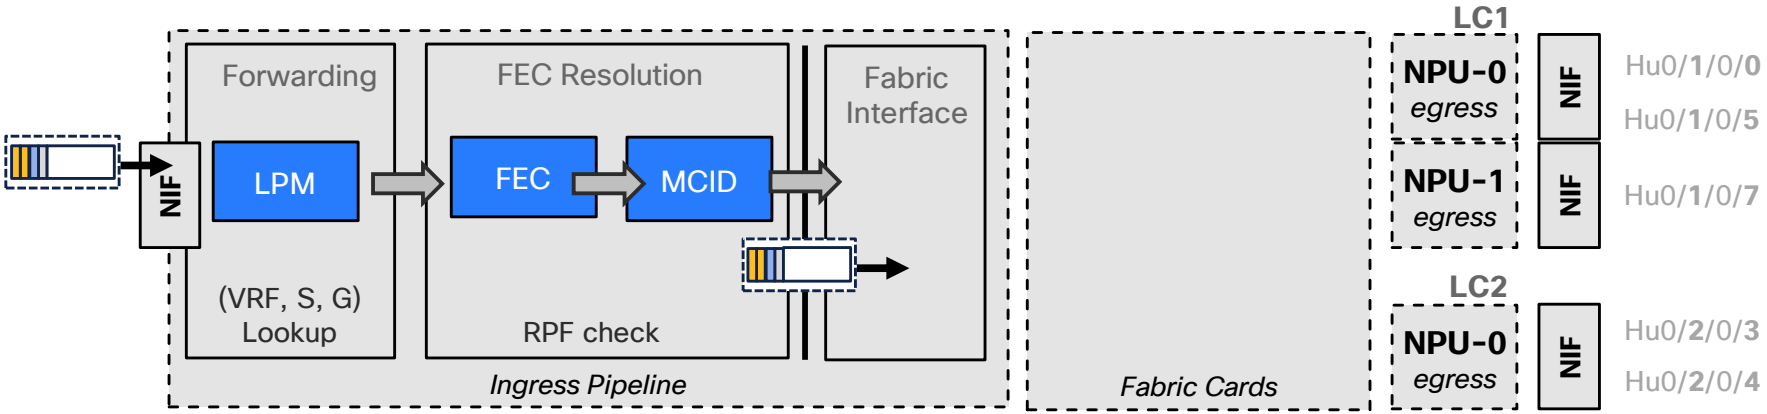

NCS5500 System Architecture

Data Plane

- Internal Header has been marked with MCID
- Packet is passed to the fabric interface and split in cells
- Based on MCID-Mapping bitmap, the cells are replicated in the fabric to the NPUs where they are re-assembled by fabric interfaces

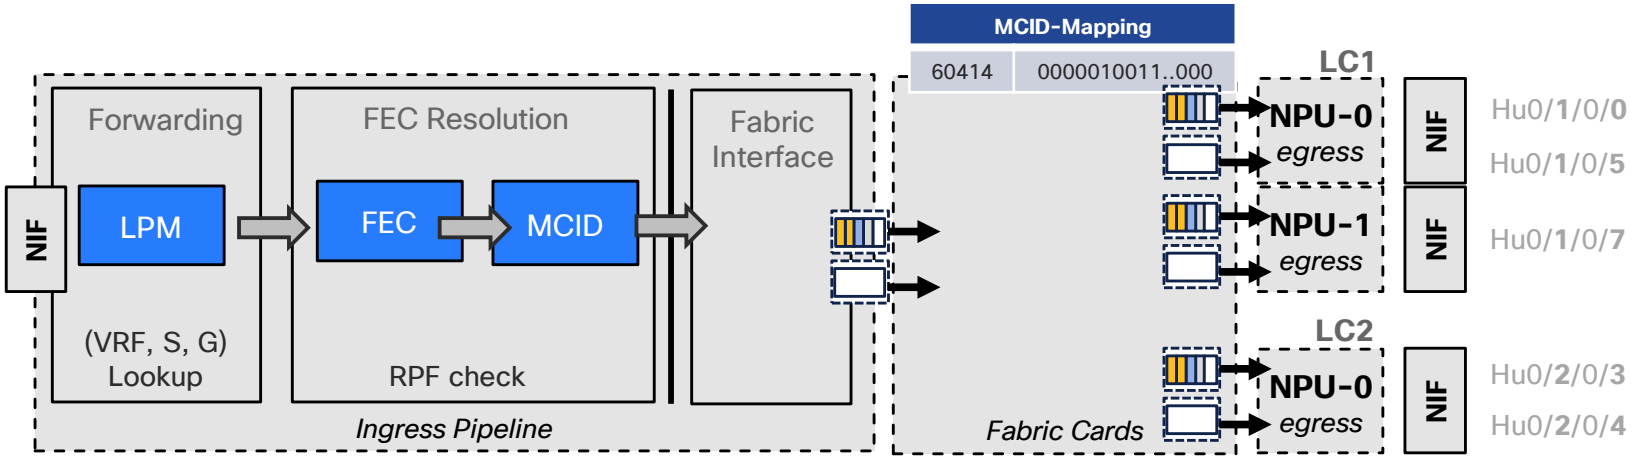

NCS5500 System Architecture

Data Plane

- Re-assembled packets will be replicated on egress NPU based on MCID-DB information
- It's the second level of replication

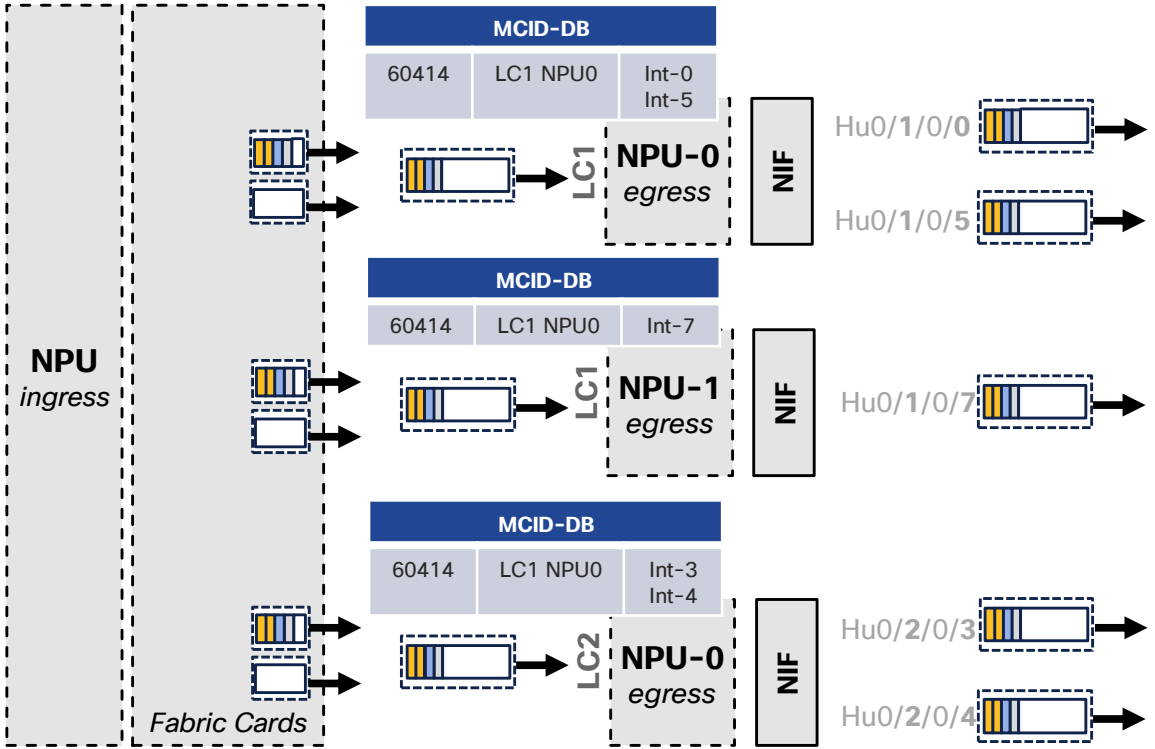

Multicast Packet Queueing in NCS5500

- Based on Fabric Multicast Queues
  - Pairs of Traffic Class mapped into FMQ
    - TC 0 and 1 to FMQ 0
    - TC 2 and 3 to FMQ 1
    - TC 4 and 5 to FMQ 2
    - TC 6 and 7 to FMQ 3
- Not scheduled / Not handled by QoS scheduling configuration (but classification and remarking is supported)
  - Back pressure mechanism needed
  - Tie-break rule in case of egress congestion

Multicast Packet Queueing in NCS5500

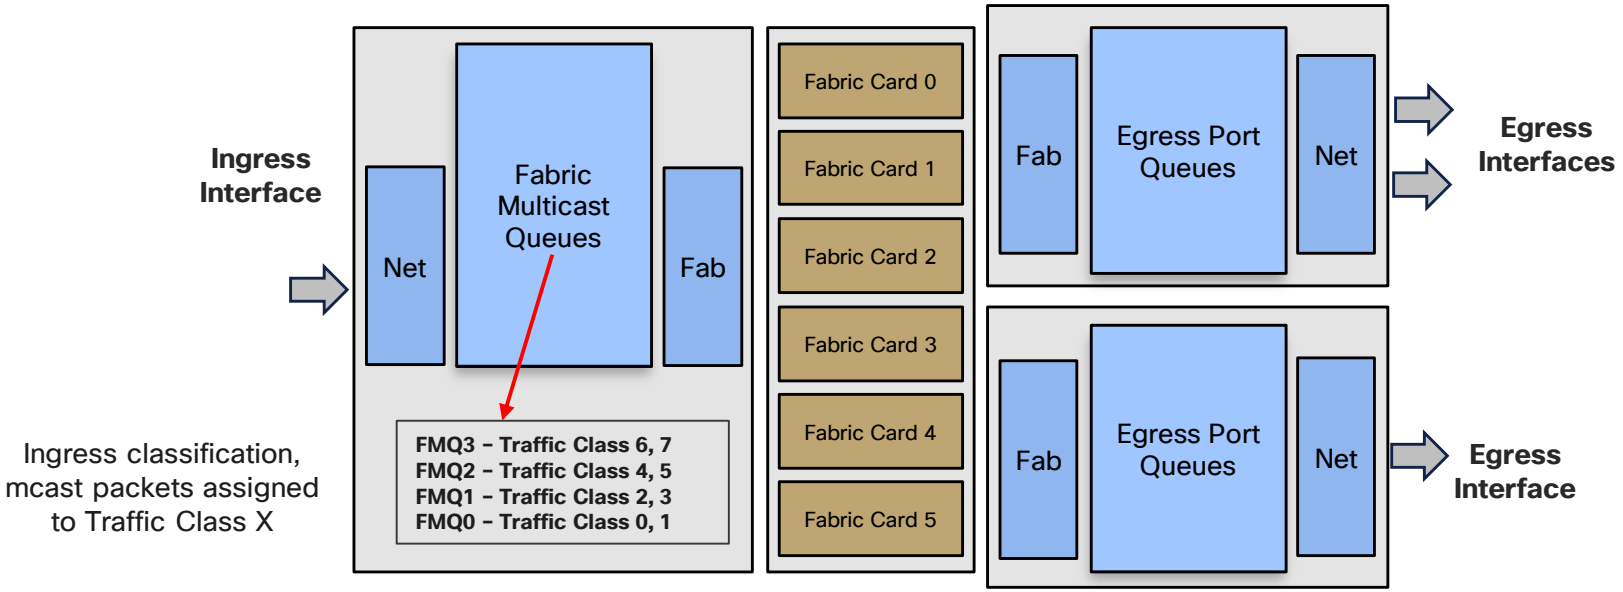

- Input policy-map sets traffic class
- Traffic Class mapped in one of the 4 FMQs, by default: goes to FMQ0

Multicast Packet Queueing in NCS5500

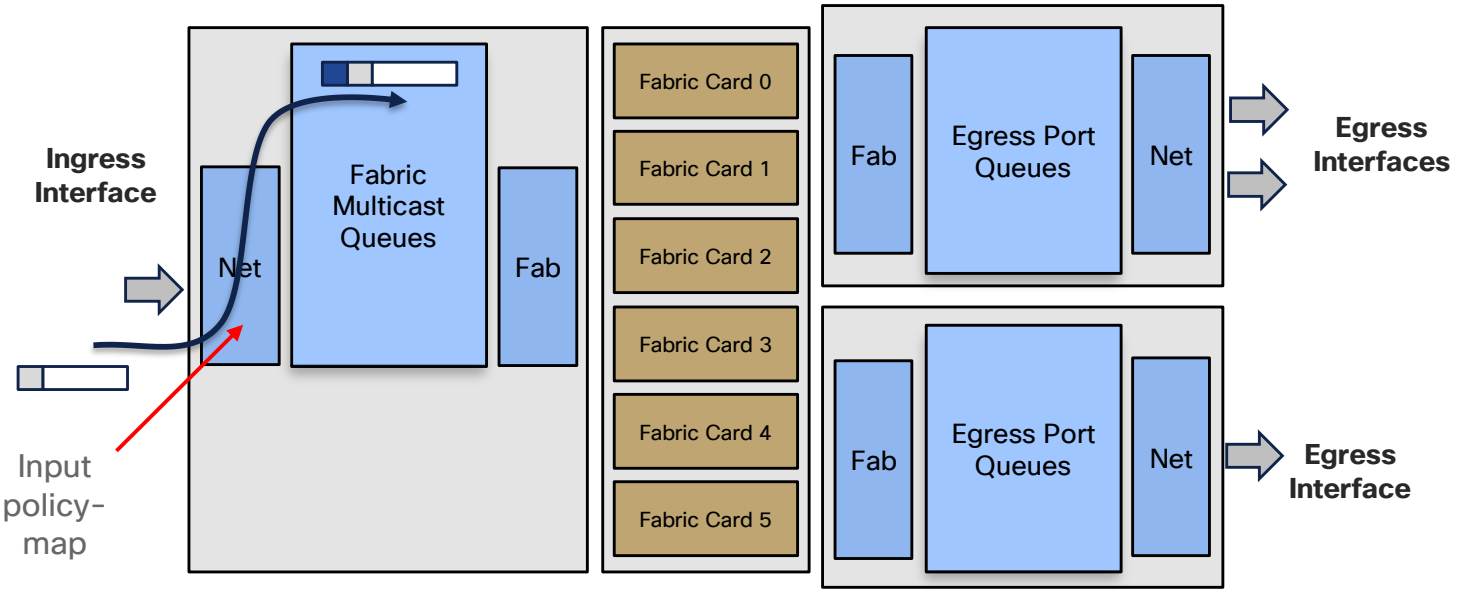

- Ingress Interface receives packet, applies input policy-map
- Then it makes forwarding decision and selects FMQ based on traffic class value

Multicast Packet Queueing in NCS5500

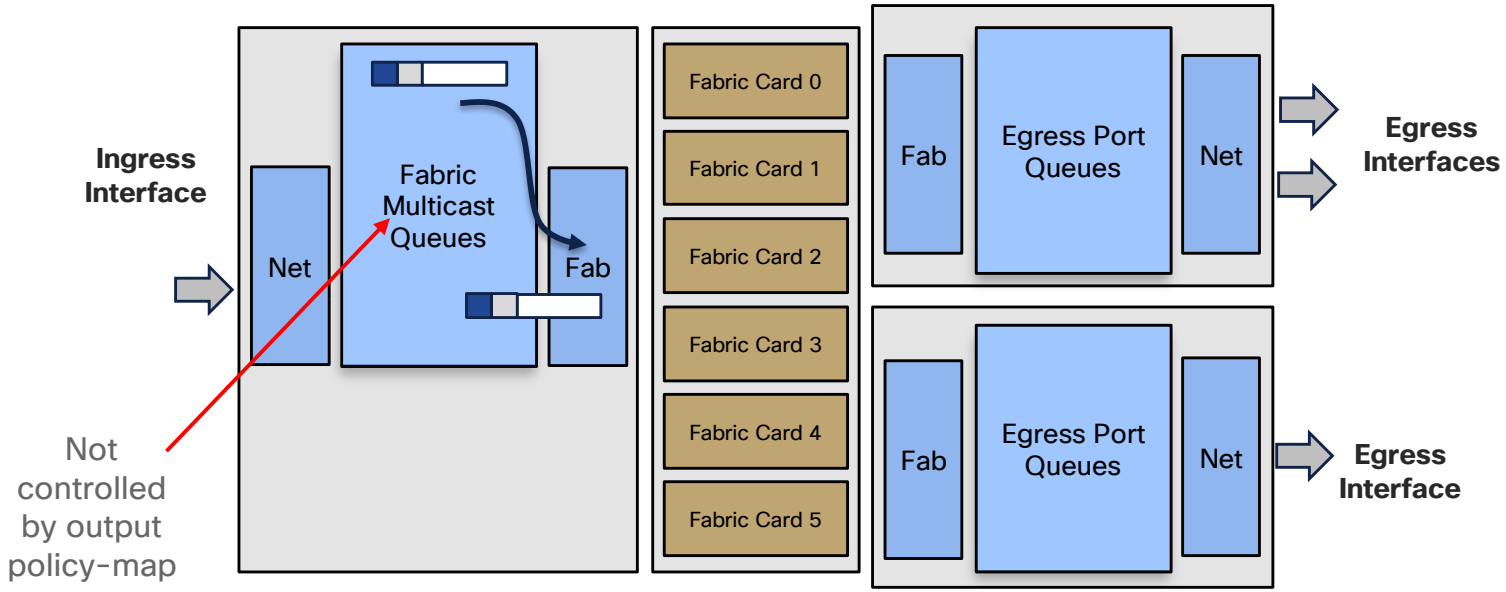

- Ingress Traffic Manager selects packet from an FMQ and gives it to Ingress Fab

Multicast Packet Queueing in NCS5500

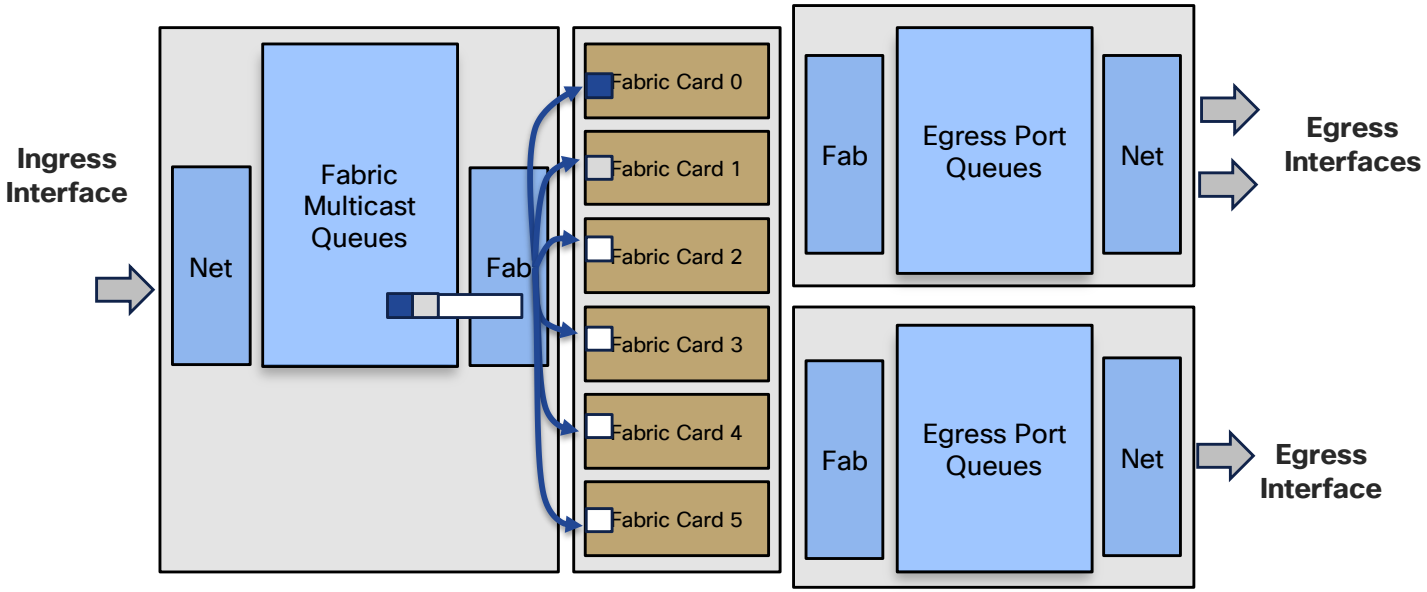

- Ingress Fab splits packet into cells and load balances them across the fabric cards

Multicast Packet Queueing in NCS5500

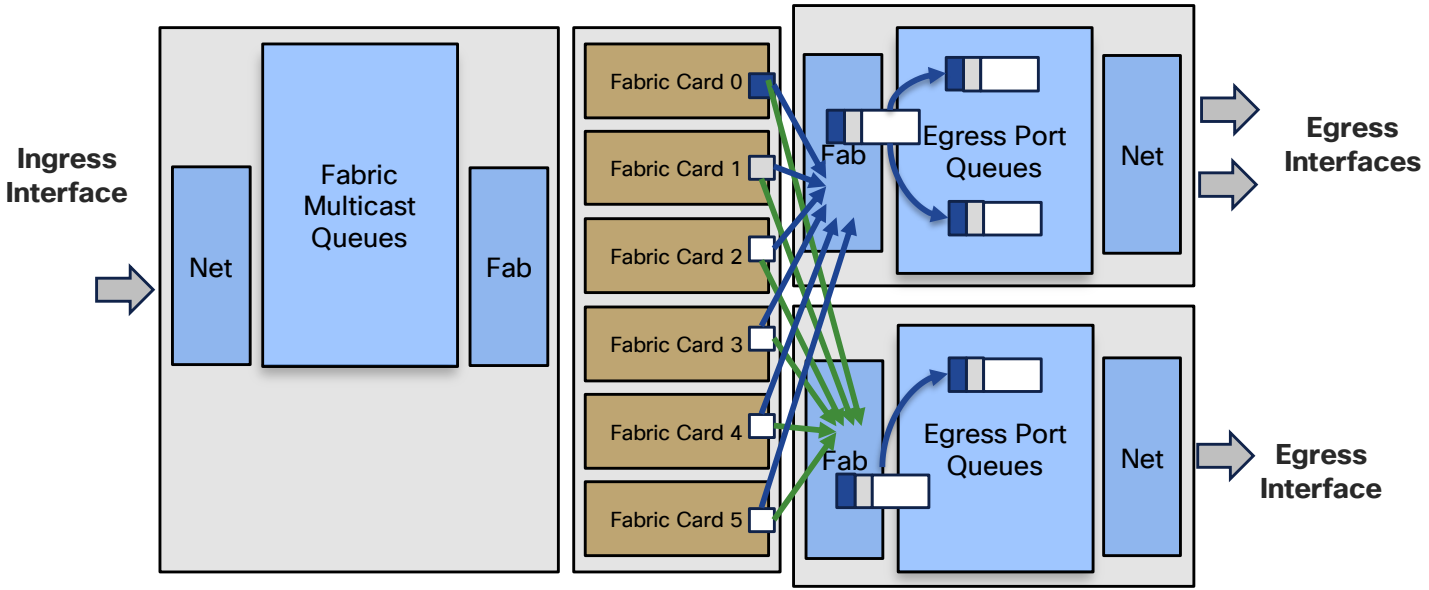

- Fabric cards replicate cells to each egress card
- Egress Fab reassembles and replicates to each interface's egress queues

Multicast Packet Queueing in NCS5500

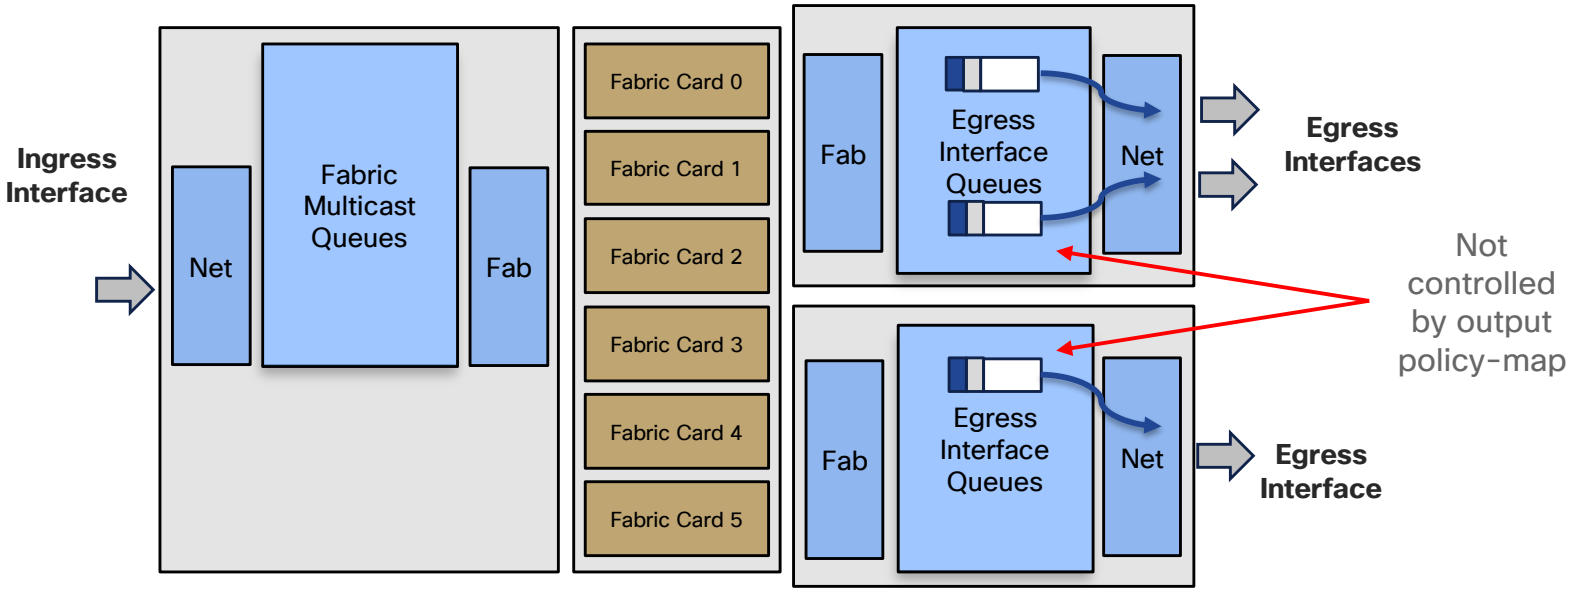

- Egress Traffic Manager selects packets from egress interface queues
- Egress Net transmits packets
- No ingress replication (one at the fabric, one at the egress NPU level)

Multicast Packet Queueing in NCS5500

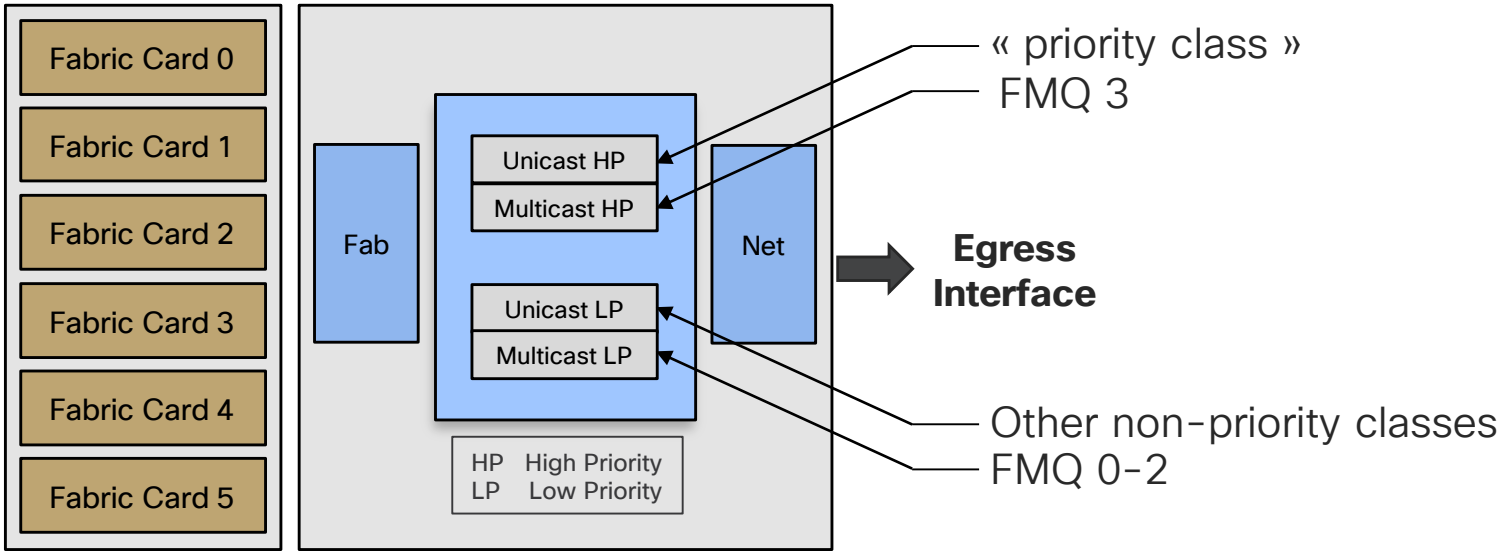

NCS5500
Memory Structure

You make networking **possible**

Cisco live!

Route Scale per Platform

Hardware		Scale
NCS-5501		1.1M pfx
NCS-5501-SE		2.75M pfx
NCS-5502		1.1M pfx
NCS-5502-SE		2.75M pfx
NCS-55A1-36H-S		1.1M pfx
NCS-55A1-36H-SE-S		4M pfx
NCS-55A1-24H		2M+ pfx
NCS-55A2-MOD-S		1.1M pfx
NCS-55A2-MOD-HD-S		1.1M pfx
NCS-55A2-MOD-SE-S		4M pfx
NCS-55A1-48Q6H		2M+ pfx
NCS-55A1-24Q6H-S		1.1M pfx

Route Scale per Platform

Hardware		Scale
NC55-36X100G		1.1M pfx
NC55-24X100G-SE		2.75M pfx
NC55-18H18F		1.1M pfx
NC55-24H12F-SE		2.75M pfx
NC55-36X100G-S		1.1M pfx
NC55-6x200-DWDM-S		1.1M pfx
NC55-36X100G-A-SE		4M pfx
NC55-MOD-A-S		1.1M pfx
NC55-MOD-A-SE-S		4M pfx

NCS5500 Forwarding ASIC Details

Memory / Databases

- Longest Prefix Match Database (LPM or KAPS)
  - Used to store IPv4 and IPv6 prefixes
  - Algorithmic memory: 256k-350k / 1M-1.5M entries (IPv6 uses 2)
- Large Exact Match Database (LEM)
  - Used to store MAC addresses, MPLS labels and IPv4 host prefix (but also /24, /23, /20… Database size: 786k entries)
- Internal TCAM (iTCAM)
  - Packet classification (ACL, QoS, VLAN ranges, tunnels. Database size: 48k entries)
- External TCAM (eTCAM, not on all line cards / systems)
  - Used for unicast route scale up to 2M or 4M+ IPv4 Routes
  - Used to extend ACL and classification

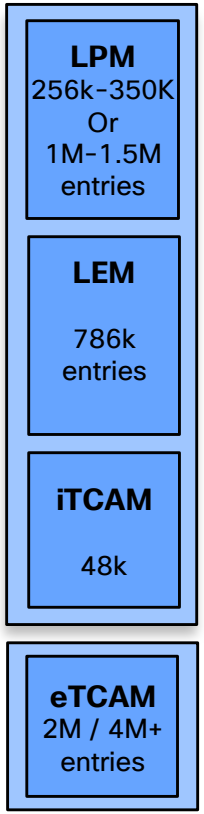

NCS5500 Forwarding ASIC Details

Algorithmic Database

- LPM memory is qualified for 256k IPv4 or 128k IPv6 addresses worst case
- Algorithmic memory scaling higher: around 350k with Internet v4 distribution and 160k with Internet v6 distribution

```
RP/0/RP0/CPU0:Router#show contr fia diagshell 0 "kbp kaps_db_stats" location 0/0/CPU0

Node ID: 0/0/CPU0

Table Configuration

Table-ID      Table-Name                Size         Table Width   AD Width      Entry Count   ~Capacity
8 - Public    FLP IPv4 UC KAPS          256000       50            20            308390        342530
8 - Private   FLP IPv4 UC KAPS          256000       50            20            308390        342530

<SNIP>

53 - Public   FLP IPv4 UC SCALE SHORT KAPS 256000     42            20            308390        342530
53 - Private  FLP IPv4 UC SCALE SHORT KAPS 256000     42            20            308390        342530
54 - Public   FLP IPv4 UC SCALE LONG KAPS 256000      50            20            308390        342530
54 - Private  FLP IPv4 UC SCALE LONG KAPS 256000      50            20            308390        342530
RP/0/RP0/CPU0:Router#
```

NCS5500 Forwarding ASIC Details

Algorithmic Database – Specific Cases

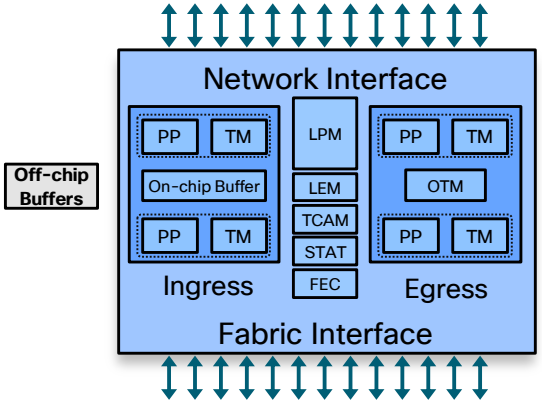

- Platforms with large LPM J+
  - NCS-55A1-24H
  - NCS-55A1-48Q6H
- LPM is algorithmic memory too and is qualified for minimum of 1M IPv4 prefixes and could scale up to 1.5M+

```
HW Resource Information
    Name                                : lpm

OOR Information
    NPU-0
        Estimated Max Entries           : 1686996
        Red Threshold                   : 95
        Yellow Threshold                : 80
        OOR State                       : Green
```

```
Current Usage
    NPU-0
        Total In-Use                    : 287936    (17 %)
        iproute                         : 287904    (17 %)
        ip6route                        : 11        (0 %)
        ipmcroute                       : 1         (0 %)
        ip6mcroute                      : 0         (0 %)
```

NCS5500 Forwarding ASIC Details

Memory / Databases

- FEC
  - Used for NextHop and ECMP (128k entries)
  - Contains the FEC ECMP (4k entries)
- Egress Encapsulation DB (EEDB)
  - Used for egress rewrites (96k entries)
    - Link Local – ARP, ND
    - Tunnel – MPLS label, GRE, etc
- Ingress/Egress Small Exact Match (ISEM/ESEM)
  - Used for tunnel termination and egress VLAN translation
- Statistics
  - Used to store all counters (256k entries)

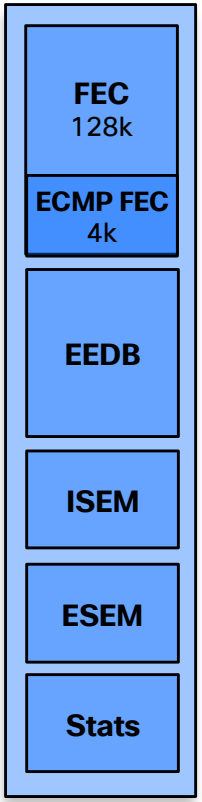

NCS5500 Databases

For Packet Lookup

- Prefix lookup points to FEC Entry
- FEC Entry contains VOQ / Egress Interface and EEDB (encapsulation entry)
- EEDB indicates the encapsulation for the packet (ARP, ND or GRE, MPLS, …)

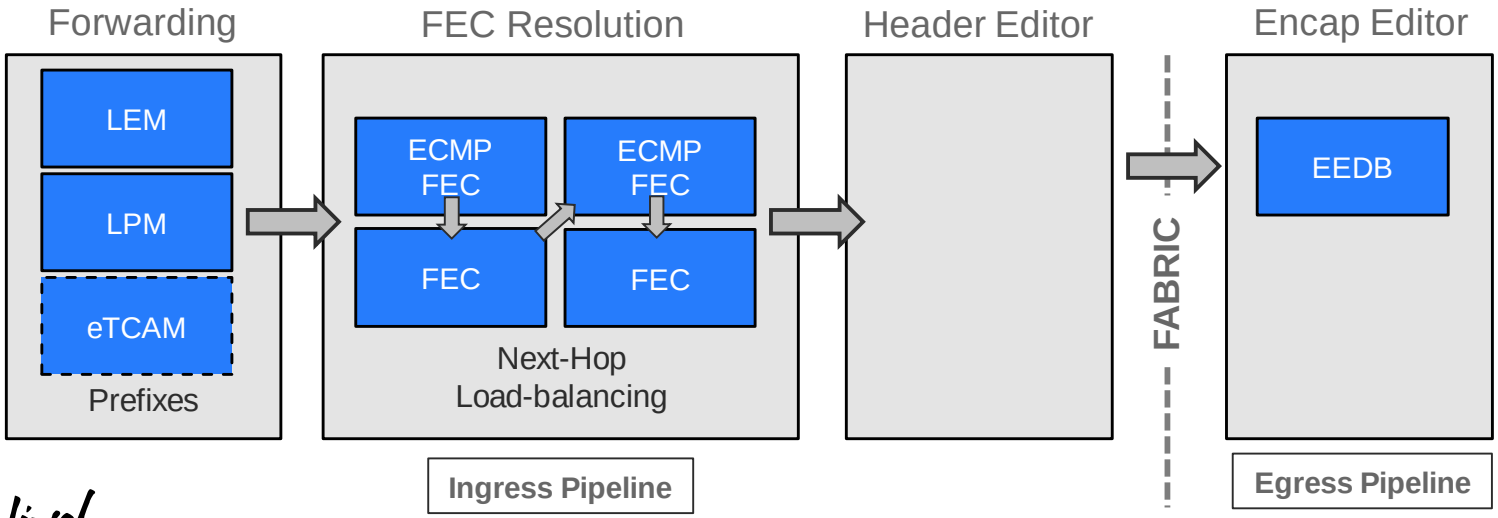

Memory Structure for Non-eTCAM Systems / LC

Host Optimized Mode (Default)

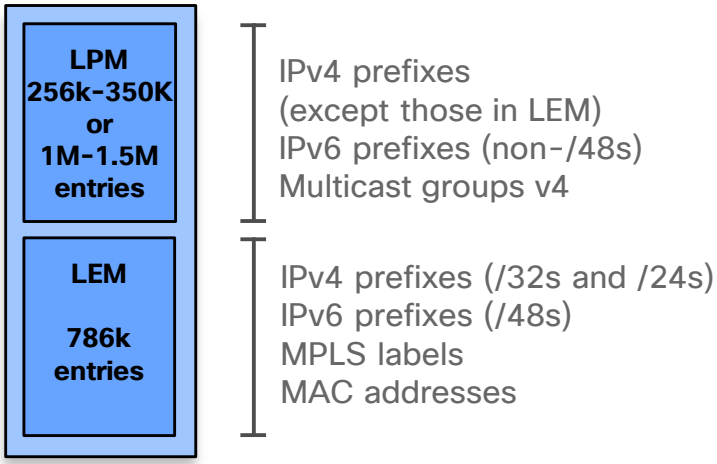

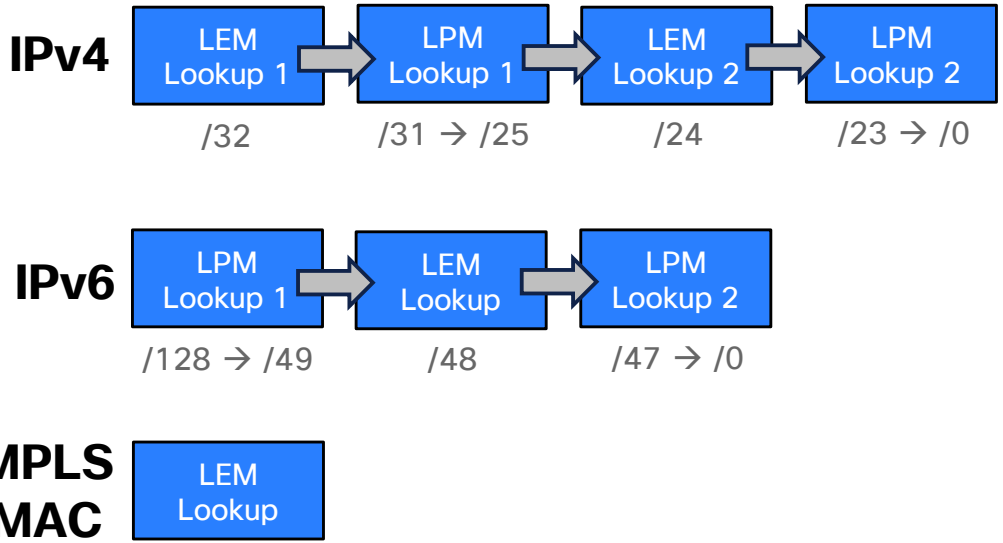

Qumran-MX no eTCAM / Jericho no eTCAM / Jericho+ no eTCAM

Non-eTCAM Systems / LC Host Optimized Mode

Illustration with 2018 Internet View: 655815 v4 and 58966 v6 real routes

- Jericho / Qumran-MX / Jericho+ with “normal” LPM

```
HW Resource Information
    Name                              : lem

Current Usage
    NPU-0
        Total In-Use                  : 386610    (49 %)
        iproute                       : 367385    (47 %)
        ip6route                      : 19222     (2 %)
        mplslabel                     : 5         (0 %)

HW Resource Information
    Name                              : lpm

Current Usage
    NPU-0
        Total In-Use                  : 328236    (83 %)
        iproute                       : 288456    (73 %)
        ip6route                      : 39767     (10 %)
        ipmcroute                     : 0         (0 %)
```

v4/32 and v4/24

v6/48

Other v4 routes

Other v6 routes

Non-eTCAM Systems / LC Host Optimized Mode

Illustration with 2019 Internet View: 751665 v4 and 42856 v6 real routes

- Jericho+ with “large” LPM

```
HW Resource Information
    Name                                : lem

Current Usage
    NPU-0
        Total In-Use                    : 396636    (50 %)
        iproute                         : 376997    (48 %)
        ip6route                        : 19650     (2 %)
        mplslabel                       : 0         (0 %)

HW Resource Information
    Name                                : lpm

Current Usage
    NPU-0
        Total In-Use                    : 397915    (24 %)
        iproute                         : 374680    (23 %)
        ip6route                        : 23214     (1 %)
        ipmcroute                       : 1         (0 %)
        ip6mcroute                      : 0         (0 %)
```

v4/32 and v4/24

v6/48

Other v4 routes

Other v6 routes

Memory Structure for Jericho non-eTCAM

Internet Optimized Mode

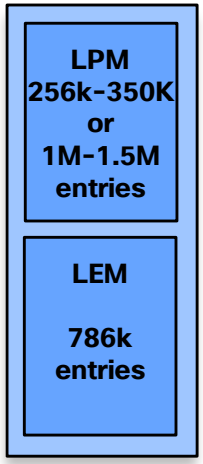

IPv4 prefixes
(except those in LEM)
IPv6 prefixes (non-/48s)
Multicast groups v4

IPv4 prefixes (/20s, /23s - /24s)
IPv6 prefixes (/48s)
MPLS labels
MAC addresses

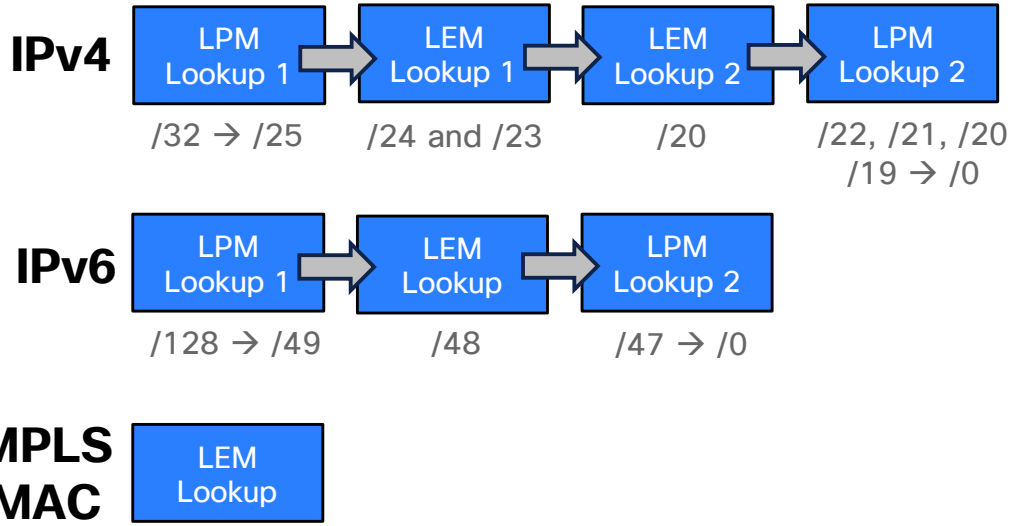


```
RP/0/RP0/CPU0:NCS55A1-24H-6.5.1(config)# hw-module fib ipv4 scale ?
  host-optimized-disable  Configure Host optimization by default
  internet-optimized      Configure Intetrnet optimized
RP/0/RP0/CPU0:NCS55A1-24H-6.5.1(config)#
```

Non-eTCAM Systems / LC Internet Optimized Mode

Illustration with Public Internet View: 655815 v4 and 58966 v6 real routes

- Jericho / Qumran-MX / Jericho+ with “normal” LPM

```
HW Resource Information
    Name                                : lem

Current Usage
    NPU-0
        Total In-Use                    : 530670    (67 %)
        iproute                         : 518495    (66 %)
        ip6route                        : 19222     (2 %)
        mplslabel                       : 5         (0 %)

HW Resource Information
    Name                                : lpm

Current Usage
    NPU-0
        Total In-Use                    : 231172    (51 %)
        iproute                         : 194021    (43 %)
        ip6route                        : 39768     (9 %)
        ipmcroute                       : 0         (0 %)
```

v4/24, v4/23 expanded v4/20

v6/48

Other v4 routes v4/20 with overlaps

Other v6 routes

Non-eTCAM Systems / LC Internet Optimized

Illustration with 2019 Internet View: 751665 v4 and 42856 v6 real routes

- Jericho+ with “large” LPM: not recommended

```
HW Resource Information
    Name                                : lem

Current Usage
    NPU-0
        Total In-Use                    : 546064    (69 %)
        iproute                         : 526417    (67 %)
        ip6route                        : 19650     (2 %)
        mplslabel                       : 0         (0 %)

HW Resource Information
    Name                                : lpm

Current Usage
    NPU-0
        Total In-Use                    : 297077    (18 %)
        iproute                         : 273842    (17 %)
        ip6route                        : 23214     (1 %)
        ipmcroute                       : 1         (0 %)
        ip6mcroute                      : 0         (0 %)
```

v4/24, v4/23 and v4/20

v6/48

Other v4 routes

Other v6 routes

Profile Recommendation For Base Systems

Hardware	NPU	Profile
NCS-5501	Qumran-MX	Internet-optimized
NCS-5502	Jericho	Internet-optimized
NCS-55A1-36H-S	Jericho+	Internet-optimized
NCS-55A1-24H	Jericho+ Large LPM	Host-optimized
NCS-55A2-MOD-S	Jericho+	Internet-optimized
NCS-55A1-48Q6H	Jericho+ Large LPM	Host-optimized
NCS-55A1-24Q6H-S	Jericho+	Internet-optimized

Memory Structure for J w/ eTCAM Systems / LC

Default Distribution

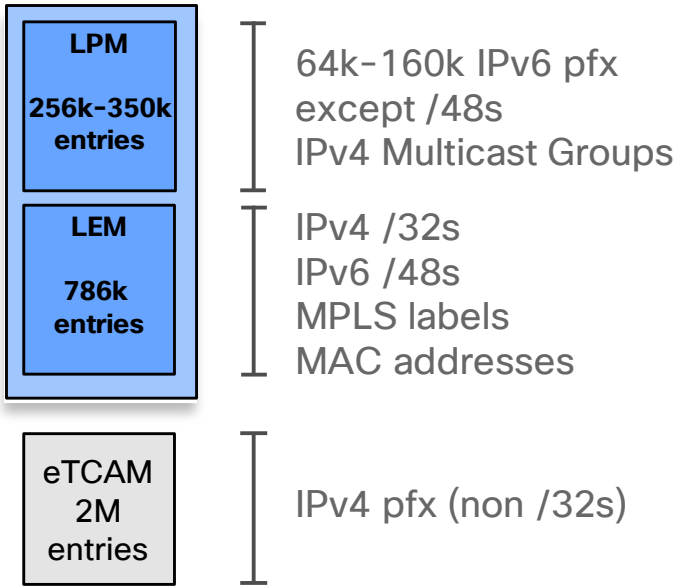

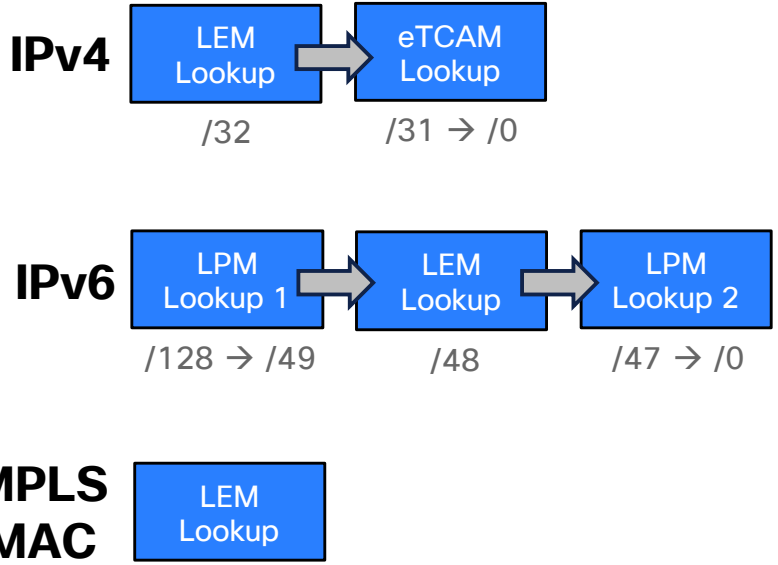

Memory Structure for J w/ eTCAM

Illustration with Public Internet View: 655815 v4 and 58966 v6 real routes

```
HW Resource Information
    Name                          : lem

Current Usage
    NPU-0
        Total In-Use              : 20132     (3 %)
        iproute                   : 904       (0 %)
        ip6route                  : 19222     (2 %)
        mplslabel                 : 5         (0 %)

HW Resource Information
    Name                          : lpm

Current Usage
    NPU-0
        Total In-Use              : 39786     (10 %)
        iproute                   : 0         (0 %)
        ip6route                  : 39767     (10 %)
        ipmcroute                 : 0         (0 %)

HW Resource Information
    Name                          : ext_tcam_ipv4

Current Usage
    NPU-0
        Total In-Use              : 654937    (40 %)
        iproute                   : 654937    (40 %)
        ipmcroute                 : 0         (0 %)
```

v4/32

v6/48

No v4 routes in LPM

Other v6 routes

All v4 routes except v4/32

Memory Structure for J+ w/ eTCAM Systems / LC

IOS XR 6.3.2 Onwards

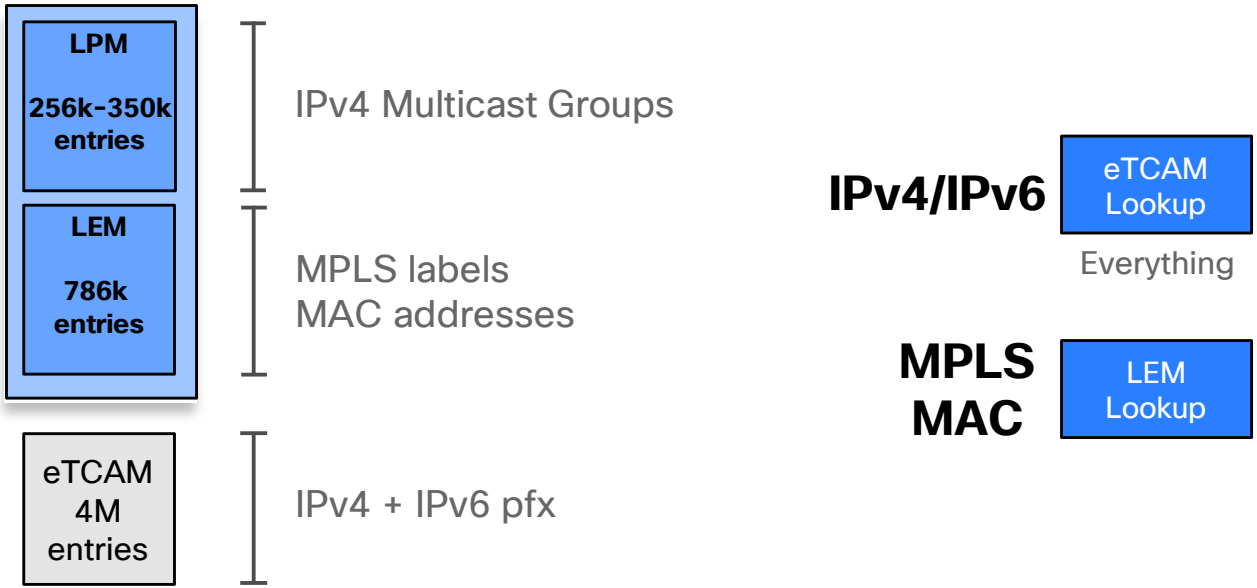

Demos

http://iosxr.io/ncs5500/

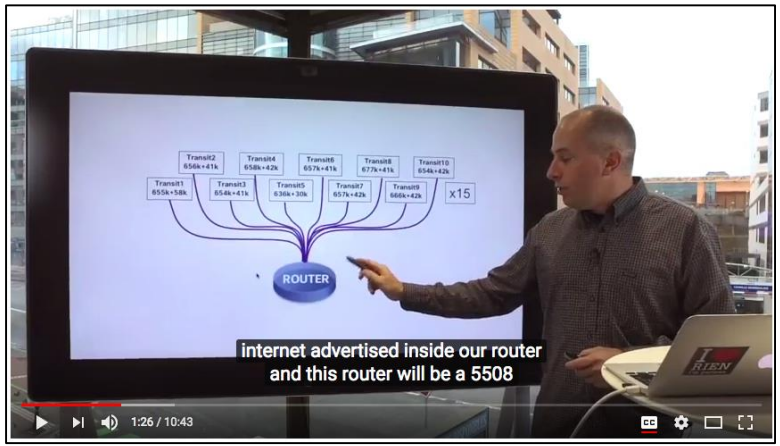

http://bit.ly/ncs5500-base

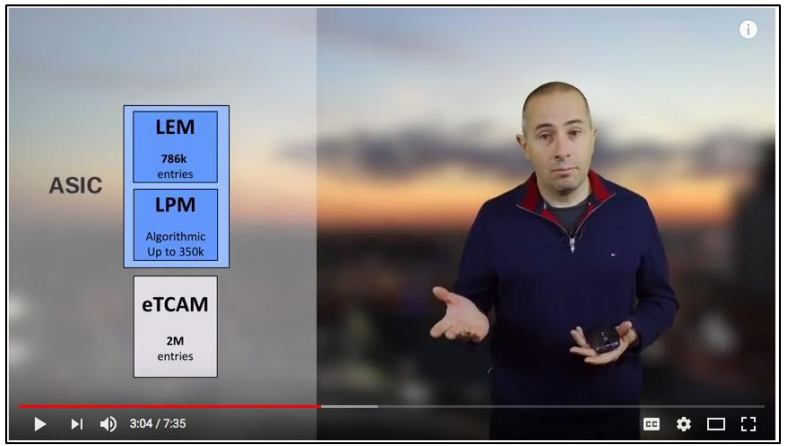

http://bit.ly/ncs5500-scale

NCS5500 Resource Monitoring

You make networking **possible**

Cisco live!

Monitoring Memory Resources

Thresholds Yellow / Red

- For both base and scale systems
- Hardware programming is done through an abstraction layer: DPA

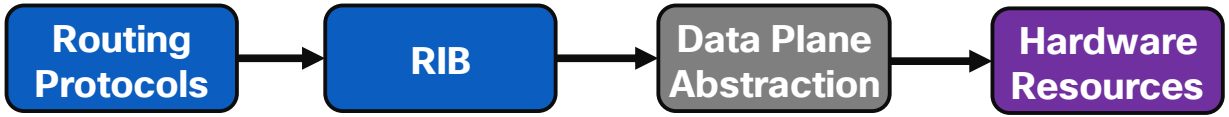

- Each database is using two thresholds: yellow at 80% and red at 95%

```
LC/0/0/CPU0:Jan 18 23:41:56.750 : fia_driver[279]: %PLATFORM-DPA-1-OOR_RED : NPU 0, Table iproute
LC/0/0/CPU0:Jan 18 23:41:56.750 : fia_driver[279]: %PLATFORM-DPA-4-OOR_YELLOW : NPU 0, Table iproute
LC/0/0/CPU0:Jan 18 23:41:56.750 : fia_driver[279]: %PLATFORM-DPA-1-OOR_RED : NPU 0, Table iproute
LC/0/0/CPU0:Jan 18 23:42:00.336 : fia_driver[279]: %PLATFORM-DPA-1-OOR_RED : NPU 2, Table iproute
LC/0/0/CPU0:Jan 18 23:42:00.418 : fia_driver[279]: %PLATFORM-DPA-1-OOR_RED : NPU 4, Table iproute
LC/0/0/CPU0:Jan 18 23:42:00.438 : fia_driver[279]: %PLATFORM-DPA-4-OOR_YELLOW : NPU 4, Table iproute
LC/0/0/CPU0:Jan 18 23:42:00.439 : fia_driver[279]: %PLATFORM-DPA-1-OOR_RED : NPU 4, Table iproute
```

Monitoring Memory Resources

Exceeding a Database Capacity

- DPA will not program new prefixes and “Hw failures” counter will increment
- Example: advertising 800k IPv4 /24s (in LEM database):
  - 784k prefixes are actually programmed and 16k are generating failures

```
RP/0/RP0/CPU0:NCS5508#sh dpa resources iproute location 0/0/CPU0

<SNIP>
                          NPU ID: NPU-0          NPU-1          NPU-2          NPU-3          NPU-4          NPU-5
<SNIP>
                          Errors
                     HW Failures: 16131          16131          16131          16132          16131          16131
                Resolve Failures: 0              0              0              0              0              0
                 No memory in DB: 0              0              0              0              0              0
                 Not found in DB: 0              0              0              0              0              0
                    Exists in DB: 0              0              0              0              0              0
RP/0/RP0/CPU0:NCS5508#
RP/0/RP0/CPU0:NCS5508#sh contr npu resources lem location 0/0/CPU0
<SNIP>
Current Usage
    NPU-0
        Total In-Use                : 783898   (100 %)
        iproute                     : 783898   (100 %)  (Prefix Count: 783898)
        mplslabel                   : 0        (0 %)  (Prefix Count: 0)
<SNIP>
```

Monitoring Memory Resources

CLI to Check LEM Database Usage

```
RP/0/RP0/CPU0:5508-6.3.2#sh contr npu resources all loc 0/1/CPU0

HW Resource Information
    Name                                : lem

OOR Information
    NPU-0
        Estimated Max Entries           : 786432
        Red Threshold                   : 95
        Yellow Threshold                : 80
        OOR State                       : Green
    NPU-1
        Estimated Max Entries           : 786432
        Red Threshold                   : 95
        Yellow Threshold                : 80
        OOR State                       : Green
<...>
    NPU-3
        Estimated Max Entries           : 786432
        Red Threshold                   : 95
        Yellow Threshold                : 80
        OOR State                       : Green
<...>
```

```
<...>

Current Usage
    NPU-0
        Total In-Use                    : 434785    (55 %)
        iproute                         : 434784    (55 %)
        ip6route                        : 0         (0 %)
        mplslabel                       : 0         (0 %)
    NPU-1
        Total In-Use                    : 434785    (55 %)
        iproute                         : 434784    (55 %)
        ip6route                        : 0         (0 %)
        mplslabel                       : 0         (0 %)

<...>

    NPU-3
        Total In-Use                    : 434785    (55 %)
        iproute                         : 434784    (55 %)
        ip6route                        : 0         (0 %)
        mplslabel                       : 0         (0 %)

<...>
```

Monitoring Memory Resources

CLI to Check LPM Database Usage

```
HW Resource Information
    Name                            : lpm

OOR Information
    NPU-0
        Estimated Max Entries       : 338879
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green
    NPU-1
        Estimated Max Entries       : 338879
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<SNIP>

    NPU-3
        Estimated Max Entries       : 338879
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<...>
```

```
Current Usage
    NPU-0
        Total In-Use                : 26        (0 %)
        iproute                     : 0         (0 %)
        ip6route                    : 0         (0 %)
        ipmcroute                   : 1         (0 %)
    NPU-1
        Total In-Use                : 26        (0 %)
        iproute                     : 0         (0 %)
        ip6route                    : 0         (0 %)
        ipmcroute                   : 1         (0 %)

<SNIP>

    NPU-3
        Total In-Use                : 26        (0 %)
        iproute                     : 0         (0 %)
        ip6route                    : 0         (0 %)
        ipmcroute                   : 1         (0 %)

<...>
```


Monitoring Memory Resources

CLI to Check EEDB/Encap Database Usage

```
HW Resource Information
    Name                            : encap

OOR Information
    NPU-0
        Estimated Max Entries       : 80000
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green
    NPU-1
        Estimated Max Entries       : 80000
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<SNIP>

    NPU-3
        Estimated Max Entries       : 80000
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<...>
```

```
Current Usage
    NPU-0
        Total In-Use                : 2         (0 %)
        ipnh                        : 0         (0 %)
        ip6nh                       : 0         (0 %)
        mplsnh                      : 2         (0 %)
    NPU-1
        Total In-Use                : 2         (0 %)
        ipnh                        : 0         (0 %)
        ip6nh                       : 0         (0 %)
        mplsnh                      : 2         (0 %)

<SNIP>

    NPU-3
        Total In-Use                : 2         (0 %)
        ipnh                        : 0         (0 %)
        ip6nh                       : 0         (0 %)
        mplsnh                      : 2         (0 %)
```


Monitoring Memory Resources

CLI to Check eTCAM Usage

```
HW Resource Information
    Name                            : ext_tcam_ipv4

OOR Information
    NPU-0
        Estimated Max Entries       : 4000000
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green
    NPU-1
        Estimated Max Entries       : 4000000
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<SNIP>

    NPU-3
        Estimated Max Entries       : 4000000
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<...>
```

```
<...>

Current Usage
    NPU-0
        Total In-Use                : 1186457   (30 %)
        iproute                     : 1186472   (30 %)
    NPU-1
        Total In-Use                : 1186457   (30 %)
        iproute                     : 1186472   (30 %)
    NPU-2
        Total In-Use                : 1186457   (30 %)
        iproute                     : 1186472   (30 %)
    NPU-3
        Total In-Use                : 1186457   (30 %)
        iproute                     : 1186472   (30 %)

<...>
```


Monitoring Memory Resources

CLI to Check FEC Database Usage

```
HW Resource Information
    Name                            : fec

OOR Information
    NPU-0
        Estimated Max Entries       : 126976
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green
    NPU-1
        Estimated Max Entries       : 126976
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<SNIP>

    NPU-3
        Estimated Max Entries       : 126976
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<...>
```

```
<...>

Current Usage
    NPU-0
        Total In-Use                : 68        (0 %)
        ipnhgroup                   : 55        (0 %)
        ip6nhgroup                  : 13        (0 %)
    NPU-1
        Total In-Use                : 68        (0 %)
        ipnhgroup                   : 55        (0 %)
        ip6nhgroup                  : 13        (0 %)
    NPU-2
        Total In-Use                : 68        (0 %)
        ipnhgroup                   : 55        (0 %)
        ip6nhgroup                  : 13        (0 %)
    NPU-3
        Total In-Use                : 68        (0 %)
        ipnhgroup                   : 55        (0 %)
        ip6nhgroup                  : 13        (0 %)


<...>
```


Monitoring Memory Resources

CLI to Check ECMP FEC Database Usage

```
HW Resource Information
    Name                            : ecmp_fec

OOR Information
    NPU-0
        Estimated Max Entries       : 4096
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green
    NPU-1
        Estimated Max Entries       : 4096
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<SNIP>

    NPU-3
        Estimated Max Entries       : 4096
        Red Threshold               : 95
        Yellow Threshold            : 80
        OOR State                   : Green

<...>
```

```
<...>

Current Usage
    NPU-0
        Total In-Use                : 0         (0 %)
        ipnhgroup                   : 0         (0 %)
        ip6nhgroup                  : 0         (0 %)
    NPU-1
        Total In-Use                : 0         (0 %)
        ipnhgroup                   : 0         (0 %)
        ip6nhgroup                  : 0         (0 %)
    NPU-2
        Total In-Use                : 0         (0 %)
        ipnhgroup                   : 0         (0 %)
        ip6nhgroup                  : 0         (0 %)
    NPU-3
        Total In-Use                : 0         (0 %)
        ipnhgroup                   : 0         (0 %)
        ip6nhgroup                  : 0         (0 %)


<...>
```


Monitoring Memory Resources

CLI to Check ECMP FEC Database Usage before 6.3.15

```
RP/0/RP0/CPU0:ios#show contr npu diag alloc all instance 0 location 0/7/CPU0

Node ID: 0/7/CPU0
<SNIP>
 Pool FECs for global use                    Total number of entries: 126976  Used entries 14     Lowest entry ID is: 4096(0x1000)
 Pool VLAN translation ingress usage          is unavalible.
 Pool VLAN translation egress usage           is unavalible.
 Pool VSIs for TB VLANS                      Total number of entries: 4096    Used entries 0      Lowest entry ID is: 1(0x1)
 Pool VSIs for MSTP                          Total number of entries: 28672   Used entries 1      Lowest entry ID is: 4096(0x1000)
 Pool FEC Failover id (Jericho)              Total number of entries: 65533   Used entries 1      Lowest entry ID is: 1(0x1)
 Pool Ingress Failover id (Jericho)          Total number of entries: 32767   Used entries 0      Lowest entry ID is: 1(0x1)
 Pool Egress Failover id (Jericho)           Total number of entries: 32767   Used entries 0      Lowest entry ID is: 1(0x1)
 Pool Failover id (Arad+ and below)           is unavalible.
 Pool QOS INGRESS LABEL MAP ID               Total number of entries: 1       Used entries 0      Lowest entry ID is: 0(0x0)
 Pool QOS INGRESS LIF/COS IDs                Total number of entries: 63      Used entries 0      Lowest entry ID is: 1(0x1)
 Pool QOS INGRESS PCP PROFILE IDs            Total number of entries: 15      Used entries 0      Lowest entry ID is: 1(0x1)
 Pool QOS INGRESS COS OPCODE IDs             Total number of entries: 7       Used entries 0      Lowest entry ID is: 0(0x0)
 Pool QOS EGRESS REMARK QOS IDs              Total number of entries: 15      Used entries 0      Lowest entry ID is: 1(0x1)
 Pool QOS EGRESS MPLS PHP QOS IDs            Total number of entries: 3       Used entries 0      Lowest entry ID is: 1(0x1)
 Pool number of meters in processor A        Total number of entries: 65536   Used entries 442    Lowest entry ID is: 0(0x0)
 Pool number of meters in processor B        Total number of entries: 65536   Used entries 12     Lowest entry ID is: 0(0x0)
 Pool SW handles of policer                  Total number of entries: 7       Used entries 0      Lowest entry ID is: 1(0x1)
 Pool ECMP id                                Total number of entries: 4095    Used entries 0      Lowest entry ID is: 1(0x1)
 Pool QOS EGRESS L2 I TAG PROFILE IDs        Total number of entries: 1       Used entries 0      Lowest entry ID is: 0(0x0)
 Pool QOS EGRESS DSCP/EXP MARKING PROFILE ID,s Total number of entries: 4     Used entries 0      Lowest entry ID is: 0(0x0)
<SNIP>
```

Monitoring Memory Resources

Alternative CLI to Check eTCAM Database Usage

```
RP/0/RP0/CPU0:NCS5508-1-631#show controllers npu diag kbp dbstats instance 0 location 0/1/CPU0
...
Table Configuration

Tbl-ID    Tbl-Name             Size      Width     AD Width  Num ent.  ~Capacity Shuffles
-------------------------------------------------------------------------------
 0        IPv4 UC              1024000   80        64        37        75591    0
 1        IPv4 RPF             1024000   80        32        0         0        0
 18       IPV4 UC DUMMY        0         80        32        0         0        0
...
RP/0/RP0/CPU0:NCS5508-1-631#show controllers npu diag kbp dbstats instance 0 location 0/6/CPU0
...
Table Configuration

Tbl-ID    Tbl-Name             Size      Width     AD Width  Num ent.  ~Capacity Shuffles
-------------------------------------------------------------------------------
 15       IPV4 DC              2048000   80        24        8         2048000  0
 20       IPV4 DC DUMMY        0         80        32        0         0        0
...
```

Monitoring Memory Resources

CLI to Check Statistics Database Usage in 6.3.x

```
RP/0/RP0/CPU0:NCS5508-1-631#sh contr npu resources stats instance 0 loc 0/7/CPU0

System information for NPU 0:
  Counter processor configuration profile: Default
  Next available counter processor:        4

Counter processor: 0                          | Counter processor: 1
  State: In use                               |   State: In use
                                              |
  Application:              In use   Total    |   Application:              In use   Total
    Trap                       97     300     |     Trap                       97     300
    Policer (QoS)               0    6976     |     Policer (QoS)               0    6976
    ACL RX, LPTS              171     915     |     ACL RX, LPTS              171     915
                                              |
                                              |
Counter processor: 2                          | Counter processor: 3
  State: In use                               |   State: In use
                                              |
  Application:              In use   Total    |   Application:              In use   Total
    VOQ                       104    8191     |     VOQ                       104    8191
                                              |
                                              |
Counter processor: 4                          | Counter processor: 5
  State: Free                                 |   State: Free
                                              |
                                              |
Counter processor: 6                          | Counter processor: 7
  State: Free                                 |   State: Free
```


For Reference

Monitoring Memory Resources

CLI to Check Statistics Database Usage in 6.3.x

```
Counter processor: 8                          | Counter processor: 9
  State: Free                                 |   State: Free
                                              |
                                              |
Counter processor: 10                         | Counter processor: 11
  State: In use                               |   State: In use
                                              |
  Application:              In use   Total    |   Application:              In use   Total
    L3 RX                        0    8191   |     L3 RX                        7    8191
    L2 RX                        0    8192   |     L2 RX                        0    8192
                                              |
                                              |
Counter processor: 12                         | Counter processor: 13
  State: In use                               |   State: In use
                                              |
  Application:              In use   Total    |   Application:              In use   Total
    Interface TX                 0   16383   |     Interface TX                14   16383
                                              |
                                              |
Counter processor: 14                         | Counter processor: 15
  State: In use                               |   State: In use
                                              |
  Application:              In use   Total    |   Application:              In use   Total
    Interface TX                 0   16384   |     Interface TX                 0   16384
                                              |
                                              |
RP/0/RP0/CPU0:NCS5508-1-631#
```

Monitoring Memory Resources via YANG

```
<?xml version="1.0"?>
<rpc-reply message-id="urn:uuid:4883a370-4115-4779-ac18-636371bb7bef"
xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<data>
  <dpa xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-fretta-bcm-dpa-hw-resources-oper">
   <stats>
    <nodes>
     <node>
      <node-name>0/0/CPU0</node-name>
      <hw-resources-datas>
       <hw-resources-data>
        <resource>lem</resource>
        <resource-id>0</resource-id>
        <name>lem</name>
        <num-npus>6</num-npus>
        <npu-hwr>
         <max-allowed>0</max-allowed>
         <npu-id>0</npu-id>
         <max-entries>750000</max-entries>
         <red-oor-threshold>712500</red-oor-threshold>
         <red-oor-threshold-percent>0</red-oor-threshold-percent>
         <yellow-oor-threshold>600000</yellow-oor-threshold>
         <yellow-oor-threshold-percent>0</yellow-oor-threshold-percent>
         <inuse-objects>13</inuse-objects>
         <num-lt>2</num-lt>
         <oor-change-count>0</oor-change-count>
         <oor-state-change-time1>N/A</oor-state-change-time1>
         <oor-state-change-time2>N/A</oor-state-change-time2>
         <oor-state>Green</oor-state>
...
```

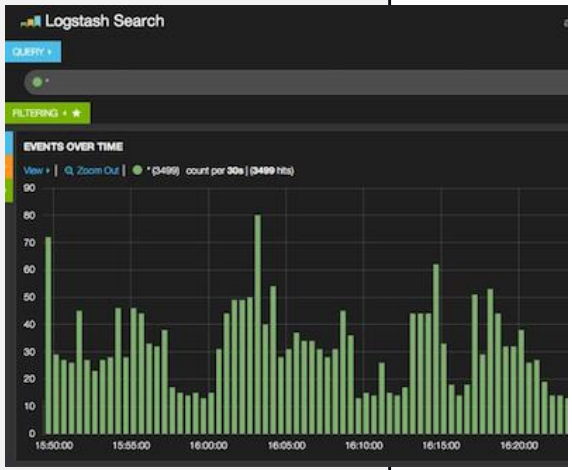

Memory Resources

Need More Info ?

- XRDOCS: https://xrdocs.io/ncs5500/tutorials/

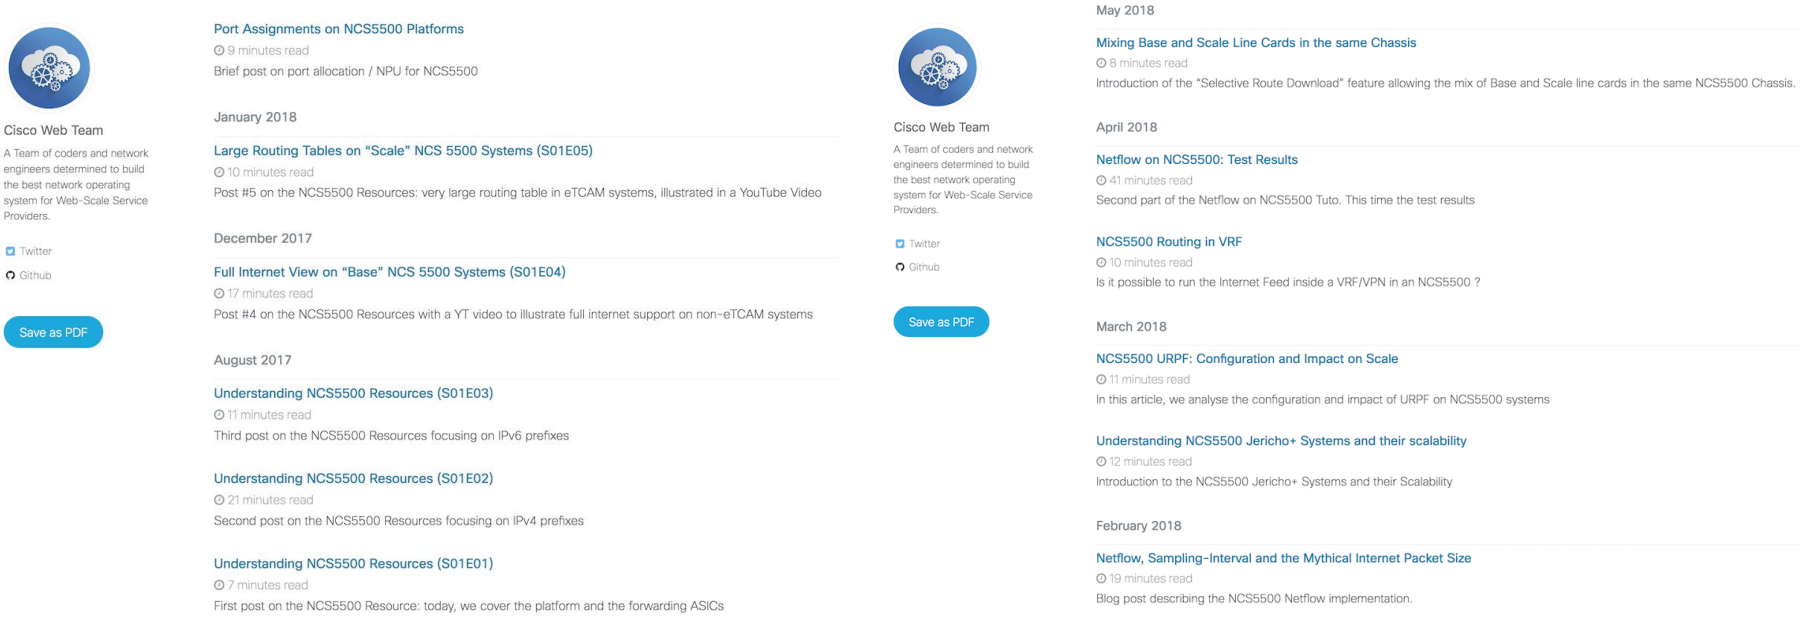

NCS5500 Access-Lists

You make networking **possible**

Cisco live!

Using Access-Lists

With Jericho and Jericho+ LC / Systems

- Traditional ACLs
  - Supported on systems with or without eTCAM
  - ACEs are stored in iTCAM only
- Hybrid / Scale ACLs
  - Supported on scale systems only (with eTCAM)
  - Part of the ACE will be stored and compress on eTCAM
  - Other part of the ACE will be in iTCAM (2-step look-up mechanism)
  - Ingress ACL only

Traditional ACLs

Using Only Internal TCAM (iTCAM)

- 12 large banks (0-11): 2k entries each
- 4 small banks (12-15): 128 entries each
  - Shared between ingress and egress features configured. First come, first served
- Same ACL used on several ingress interfaces are counted once
- Same ACL used on X egress interfaces are counted X times
- Support of 32 ingress and 32/255 egress ACLs per NPU
  - More with recent version of IOS XR
- Support 4000 IPv4 or 2000 IPv6 ACEs per NPU
  - Smaller than potential 12k entries (bundles spread among multiple NPUs)

Traditional ACLs

Checking Internal TCAM (iTCAM) in 6.2.2 Onwards

```
RP/0/RP0/CPU0:NCS5508-2-622#sh contr npu internaltcam location 0/7/CPU0
Internal TCAM Resource Information
NPU   Bank   Entry   Owner        Free      Per-DB   DB    DB
      Id     Size                 Entries   Entry    ID    Name
=============================================================
0     0\1    320b    pmf-0        2006      36       7     INGRESS_LPTS_IPV4
0     0\1    320b    pmf-0        2006      2        12    INGRESS_RX_ISIS
0     0\1    320b    pmf-0        2006      2        32    INGRESS_QOS_IPV6
0     0\1    320b    pmf-0        2006      2        34    INGRESS_QOS_L2
0     2      160b    pmf-0        2044      2        31    INGRESS_QOS_IPV4
0     2      160b    pmf-0        2044      1        33    INGRESS_QOS_MPLS
0     2      160b    pmf-0        2044      1        42    INGRESS_ACL_L2
0     3      160b    egress_acl   2022      10       3     EGRESS_RECEIVE
0     3      160b    egress_acl   2022      16       4     EGRESS_QOS_MAP
0     4\5    320b    pmf-0        2024      24       8     INGRESS_LPTS_IPV6
0     6      160b    Free         2048      0        0
0     7      160b    Free         2048      0        0
0     8      160b    Free         2048      0        0
0     9      160b    Free         2048      0        0
0     10     160b    Free         2048      0        0
0     11     160b    Free         2048      0        0
0     12     160b    pmf-1        90        37       11    INGRESS_RX_L2
0     12     160b    pmf-1        90        1        13    INGRESS_MCAST_IPV4_ASM
0     13     160b    pmf-0        112       2        10    INGRESS_DHCP
0     13     160b    pmf-0        112       13       26    INGRESS_MPLS
0     13     160b    pmf-0        112       1        41    INGRESS_EVPN_AA_ESI_TO_FBN_DB
0     14     160b    Free         128       0        0
0     15     160b    Free         128       0        0
```

Free Space
No ACL configured

Traditional ACLs

Checking Internal TCAM (iTCAM) in 6.2.2 Onwards

```
RP/0/RP0/CPU0:NCS5508-2-622#sh contr npu internaltcam location 0/7/CPU0
Internal TCAM Resource Information
NPU   Bank   Entry   Owner       Free      Per-DB   DB   DB
      Id     Size                Entries   Entry    ID   Name
=============================================================
0     0\1    320b    pmf-0       2006      36       7    INGRESS_LPTS_IPV4
0     0\1    320b    pmf-0       2006      2        12   INGRESS_RX_ISIS
0     0\1    320b    pmf-0       2006      2        32   INGRESS_QOS_IPV6
0     0\1    320b    pmf-0       2006      2        34   INGRESS_QOS_L2
0     2      160b    pmf-0       2044      2        31   INGRESS_QOS_IPV4
0     2      160b    pmf-0       2044      1        33   INGRESS_QOS_MPLS
0     2      160b    pmf-0       2044      1        42   INGRESS_ACL_L2
0     3      160b    egress_acl  2022      10       3    EGRESS_RECEIVE
0     3      160b    egress_acl  2022      16       4    EGRESS_QOS_MAP
0     4\5    320b    pmf-0       2024      24       8    INGRESS_LPTS_IPV6
0     6      160b    pmf-0       997       1051     16   INGRESS_ACL_L3_IPV4
0     7      160b    Free        2048      0        0
0     8      160b    Free        2048      0        0
0     9      160b    Free        2048      0        0
0     10     160b    Free        2048      0        0
0     11     160b    Free        2048      0        0
0     12     160b    pmf-1       90        37       11   INGRESS_RX_L2
0     12     160b    pmf-1       90        1        13   INGRESS_MCAST_IPV4_ASM
0     13     160b    pmf-0       112       2        10   INGRESS_DHCP
0     13     160b    pmf-0       112       13       26   INGRESS_MPLS
0     13     160b    pmf-0       112       1        41   INGRESS_EVPN_AA_ESI_TO_FBN_DB
0     14     160b    Free        128       0        0
0     15     160b    Free        128       0        0
```

1000 ACEs configured

Traditional ACLs

Counters

- Limitations with packets targeted to the router
- For-us packets matching deny ACE
  - Counted and dropped
- For-us packets matching permit ACE
  - Punted and not counted

Traditional ACLs

Counting with permit ACEs

- By default only deny ACEs are allocated counters
- Permit entries can be allocated counters via configuration

```
RP/0/RP0/CPU0:NCS5508-1-631(config)#hw-module profile stats acl-permit
RP/0/RP0/CPU0:NCS5508-1-631(config)#commit
```

- Requires a reload of the line card to be activated

```
RP/0/RP0/CPU0:NCS5508-1-631#sh access-lists ipv4 PERMIT-TEST hardware ingress location 0/7/CPU0

ipv4 access-list PERMIT-TEST
10 permit icmp any host 1.1.1.1
15 permit icmp any host 1.1.1.3
16 permit tcp any any eq telnet (2 matches) <-<-<-
17 permit tcp any eq telnet any
20 permit udp any any
30 permit tcp any any
40 deny ipv4 any any (1169 matches)
RP/0/RP0/CPU0:NCS5508-1-631#
```

Hybrid ACLs

Only on eTCAM Systems

- In 6.3.2, requires a carving
- IPv4 and IPv6
- Ingress only
- Two-step look-up
  - First in eTCAM
  - Second in iTCAM

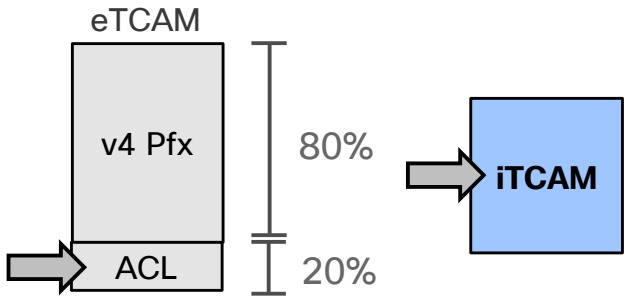

Hybrid ACLs

Example

- CLI to display an expanded version of the access-list

```
object-group network ipv4 netobj1
 10.2.1.0/24
 host 1.3.5.7
 host 1.11.111.1
!
object-group port portobj1
 eq telnet
 eq bgp
 range 100 200
!
```

```
RP/0/RP0/CPU0:R1#sh access-lists ipv4 network-object-acl
ipv4 access-list network-object-acl
 10 deny tcp net-group netobj1 port-group portobj1 any
 20 permit ipv4 net-group netobj1 any

RP/0/RP0/CPU0:R1#sh access-lists ipv4 network-object-acl expanded
ipv4 access-list network-object-acl
 10 deny tcp 10.2.1.0 0.0.0.255 eq telnet any
 10 deny tcp 10.2.1.0 0.0.0.255 eq bgp any
 10 deny tcp 10.2.1.0 0.0.0.255 range 100 200 any
 10 deny tcp host 1.11.111.1 eq telnet any
 10 deny tcp host 1.11.111.1 eq bgp any
 10 deny tcp host 1.11.111.1 range 100 200 any
 10 deny tcp host 1.3.5.7 eq telnet any
 10 deny tcp host 1.3.5.7 eq bgp any
 10 deny tcp host 1.3.5.7 range 100 200 any
 20 permit ipv4 10.2.1.0 0.0.0.255 any
 20 permit ipv4 host 1.11.111.1 any
 20 permit ipv4 host 1.3.5.7 any
RP/0/RP0/CPU0:R1#
```

Hybrid ACLs

Monitoring Resource: 1- On eTCAM

```
RP/0/RP0/CPU0:NCS5508-1-631#sh contr npu externaltcam loc 0/7/CPU0

External TCAM Resource Information
=================================================================
NPU   Bank    Entry   Owner        Free      Per-DB   DB    DB
      Id      Size                 Entries   Entry    ID    Name
=================================================================
0     0       80b     FLP          983784    654616   15    IPV4 DC
0     1       80b     FLP          28634     38       81    INGRESS_IPV4_SRC_IP_EXT
0     2       80b     FLP          28671     1        82    INGRESS_IPV4_DST_IP_EXT
0     3       160b    FLP          26624     0        83    INGRESS_IPV6_SRC_IP_EXT
0     4       160b    FLP          26624     0        84    INGRESS_IPV6_DST_IP_EXT
0     5       80b     FLP          28664     8        85    INGRESS_IP_SRC_PORT_EXT
0     6       80b     FLP          28672     0        86    INGRESS_IPV6_SRC_PORT_EXT
...
```

Hybrid ACLs

Monitoring Resource: 2- On iTCAM

```
RP/0/RP0/CPU0:NCS5508-1-631#sh contr npu internaltcam loc 0/7/CPU0
Internal TCAM Resource Information
============================================================
NPU   Bank    Entry   Owner        Free      Per-DB   DB    DB
      Id      Size                 Entries   Entry    ID    Name
============================================================
0     0\1     320b    pmf-0        1963      49       7     INGRESS_LPTS_IPV4
0     0\1     320b    pmf-0        1963      2        12    INGRESS_RX_ISIS
0     0\1     320b    pmf-0        1963      11       32    INGRESS_QOS_IPV6
0     0\1     320b    pmf-0        1963      23       34    INGRESS_QOS_L2
0     2       160b    pmf-0        2030      11       31    INGRESS_QOS_IPV4
0     2       160b    pmf-0        2030      6        33    INGRESS_QOS_MPLS
0     2       160b    pmf-0        2030      1        42    INGRESS_ACL_L2
0     3       160b    egress_acl   2032      16       4     EGRESS_QOS_MAP
0     4\5     320b    pmf-0        2021      27       8     INGRESS_LPTS_IPV6
0     6\7     320b    pmf-1        2045      3        49    INGRESS_HYBRID_ACL
0     8       160b    Free         2048      0        0
0     9       160b    Free         2048      0        0
0     10      160b    Free         2048      0        0
0     11      160b    Free         2048      0        0
0     12      160b    pmf-1        88        40       11    INGRESS_RX_L2
0     13      160b    pmf-0        84        3        10    INGRESS_DHCP
0     13      160b    pmf-0        84        1        13    INGRESS_MCAST_IPV4_ASM
0     13      160b    pmf-0        84        13       26    INGRESS_MPLS
0     13      160b    pmf-0        84        1        41    INGRESS_EVPN_AA_ESI_TO_FBN_DB
0     13      160b    pmf-0        84        26       79    INGRESS_BFD_IPV4_NO_DESC_TCAM_T
0     14      160b    Free         128       0        0
0     15      160b    Free         128       0        0
```

NCS5500 Introduction to QoS

You make networking **possible**

Cisco live!

Quality of Service on NCS5500

- Ingress direction supports classification and remarking
- Egress direction supports the same with less flexibility
- Policing only in ingress
- Shaping only in egress

Quality of Service

Internal Markers

- We use internal markers at ingress to take egress actions

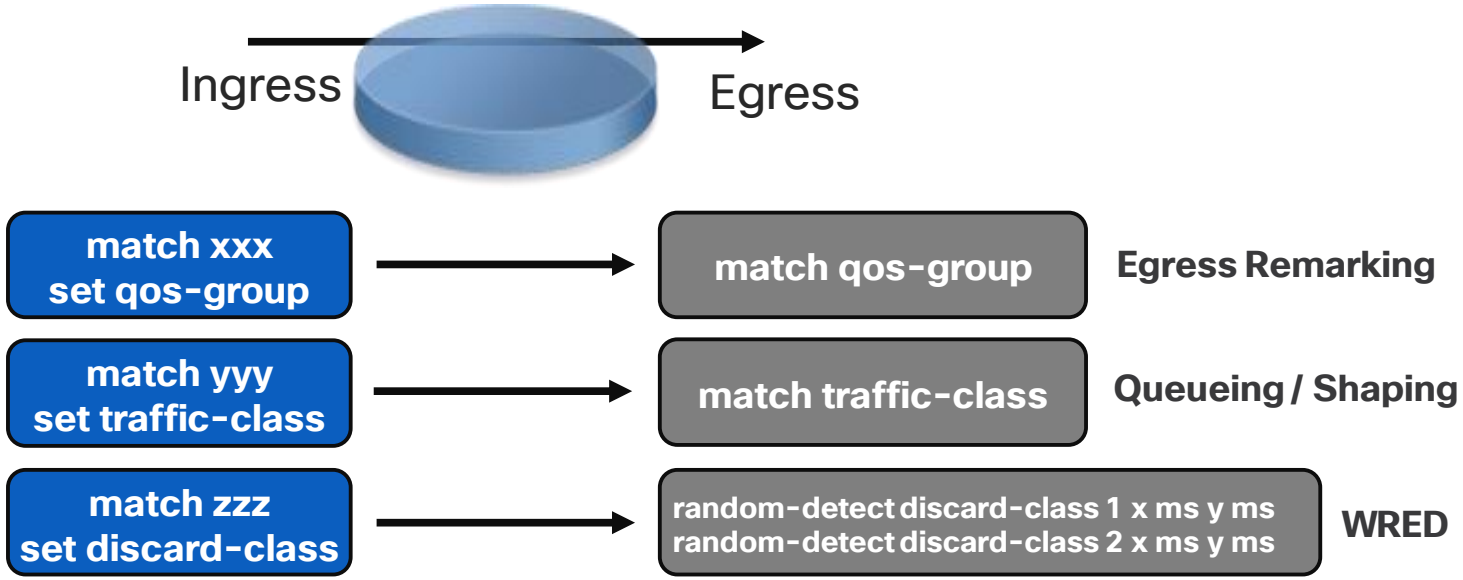

Configuring Quality of Service

Policer Configuration


```
class-map classify1
  match precedence 1

policy-map Pol1
  class classify1
    set qos-group 1
    set dscp ef
    police rate percent 10

interface hu 0/0/0/0
  service-policy input Pol1
```

Class-Map
Match criterias

set qos-group (optional)

set dscp/... (optional)

Policer

Configuring Quality of Service

Policer Configuration

```
class-map classify1
     match precedence 1
class-map classify2
     match precedence 2
class-map classify3
     match precedence 3
policy-map  ingress-policy
    class classify1
         set qos-group 1
         police rate percent 10 peak-rate percent 20
     class classify2
         set qos-group 2
     class classify3
         set qos-group 3
interface hu 0/0/0/0
    service-policy input ingress-policy
```

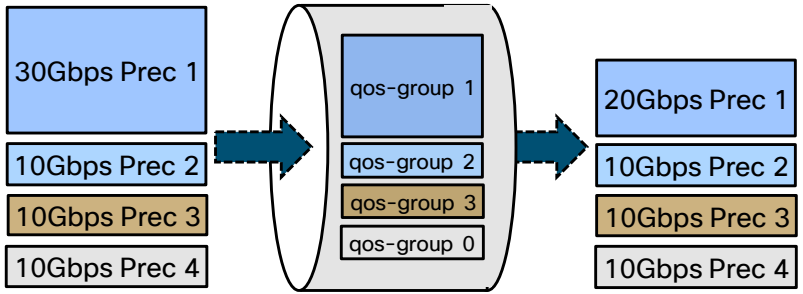

Configuring Quality of Service

Shaper Configuration

```
class-map match-any classify1
 match precedence 1
 end-class-map
!
class-map match-any classify2
 match precedence 2
 end-class-map
!
class-map match-any classify3
 match precedence 3
 end-class-map
!
policy-map Pol1
 class classify1
  set traffic-class 1
 !
 class classify2
  set traffic-class 2
 !
 class classify3
  set traffic-class 3
 !
 class class-default
  set traffic-class 7
 !
 end-policy-map
interface bundle-ether 1
  service-policy input Pol1
```

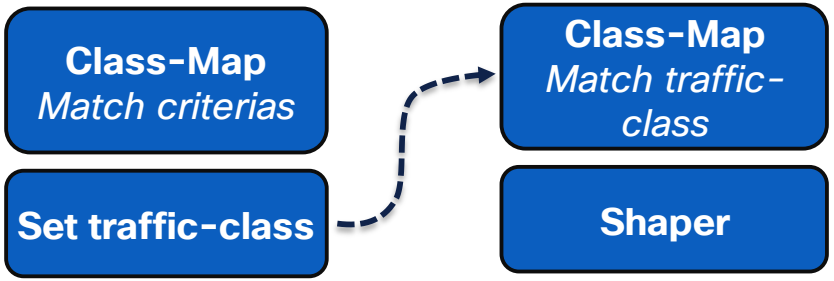

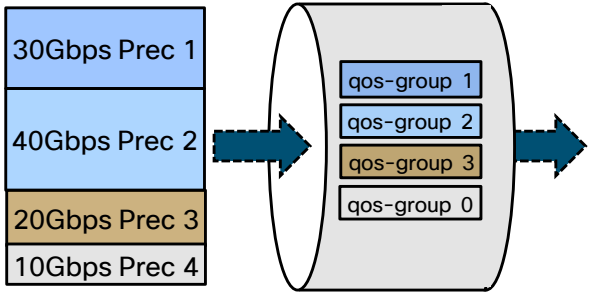


```
class-map match-any tc1
 match traffic-class 1
 end-class-map
!
class-map match-any tc2
 match traffic-class 2
 end-class-map
!
class-map match-any tc3
 match traffic-class 3
 end-class-map
!
policy-map Pol1
 class tc1
  priority level 1
  shape average percent 20
 !
 class tc2
  shape average percent 50
 !
 class tc3
  shape average percent 30
 !
 class class-default
 !
 end-policy-map
!
interface hu 0/0/0/0
  service-policy output Pol1
```

Configuring Quality of Service

Egress Dual-Policy Example

```
class-map match-any cos1
   match cos 1
end-class-map
!
class-map match-any cos2
   match cos 2
end-class-map
!
policy-map ingress-classify
   class cos1
      set qos-group 1
      set traffic-class 3
  !
  class cos2
      set qos-group 2
      set traffic-class 5
  !
class class-default
!
```

```
class-map match-any qos1
   match qos-group 1
end-class-map
!
class-map match-any qos2
    match qos-group 2
end-class-map
!
policy-map egress-marking
   class qos1
       set cos 1
   !
   class qos2
       set cos 2
       set dei 1
   !
class class-default
       set cos 7
   !
end-policy-map
```

```
class-map match-any tc3
    match traffic-class 3
end-class-map
!
class-map match-any tc5
    match traffic-class 5
end-class-map
!
policy-map egress-queuing
   class tc3
       priority level 1
       shape average 10 mbps
   !
   class tc5
       bandwidth remaining <>
   !
   class class-default
   !
end-policy-map
!

interface TenGigE0/0/1/0/0
 service-policy input ingress-classif
 service-policy output egress-marking
 service-policy output egress-queuing
!
```

Configuring Quality of Service

Shaper Configuration on Bundles

- All QoS rules applied to a bundle are applied to all members

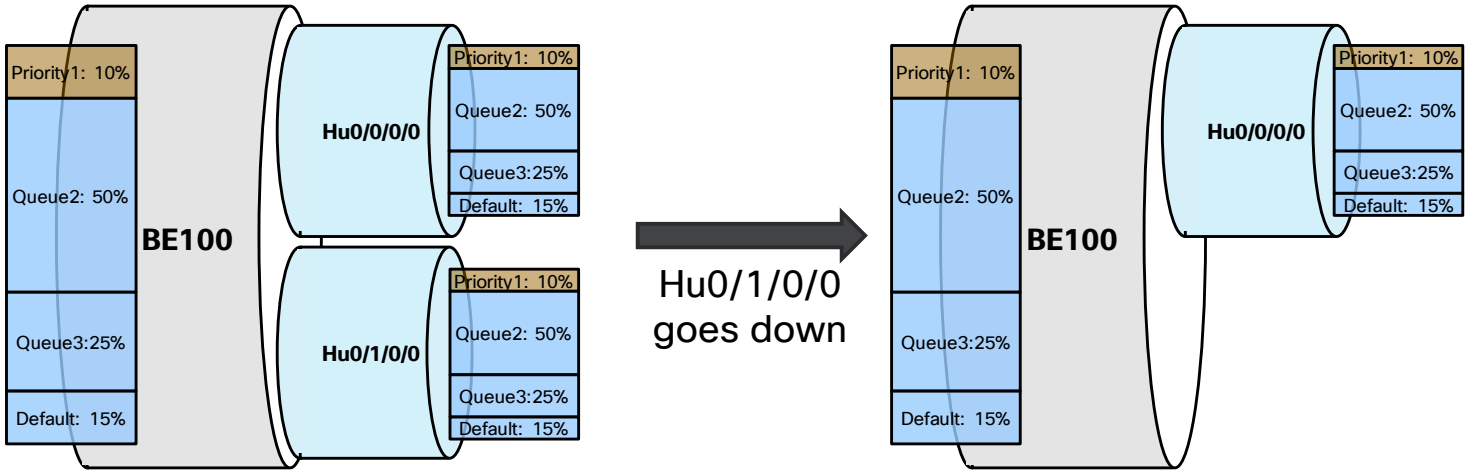

Configuring Quality of Service

Shaper Configuration on Bundles

- If we use absolute values, they are applied to each member too

→ Use percent

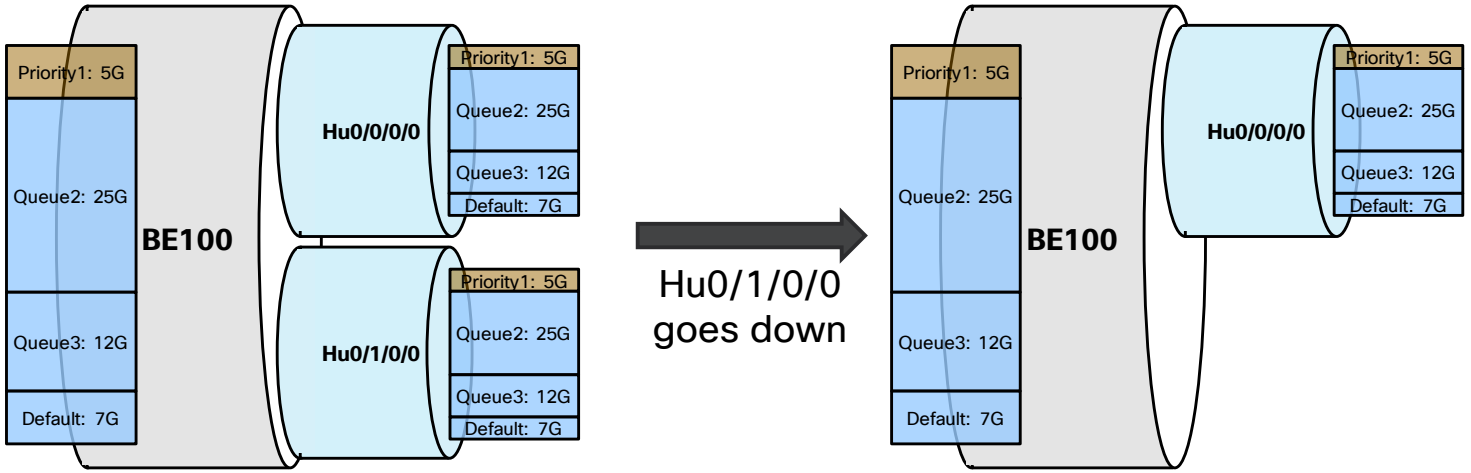

Key Differences with Traditional XR Platforms

Unicast is Scheduled but Multicast Traffic doesn't Follow VOQ-only Model

- In case of egress interface congestion
  - If unicast or multicast is high priority, it will take full precedence over the other
  - If same priority (HP/HP or LP/LP), then the forwarding will be 80% ucast / 20% mcast

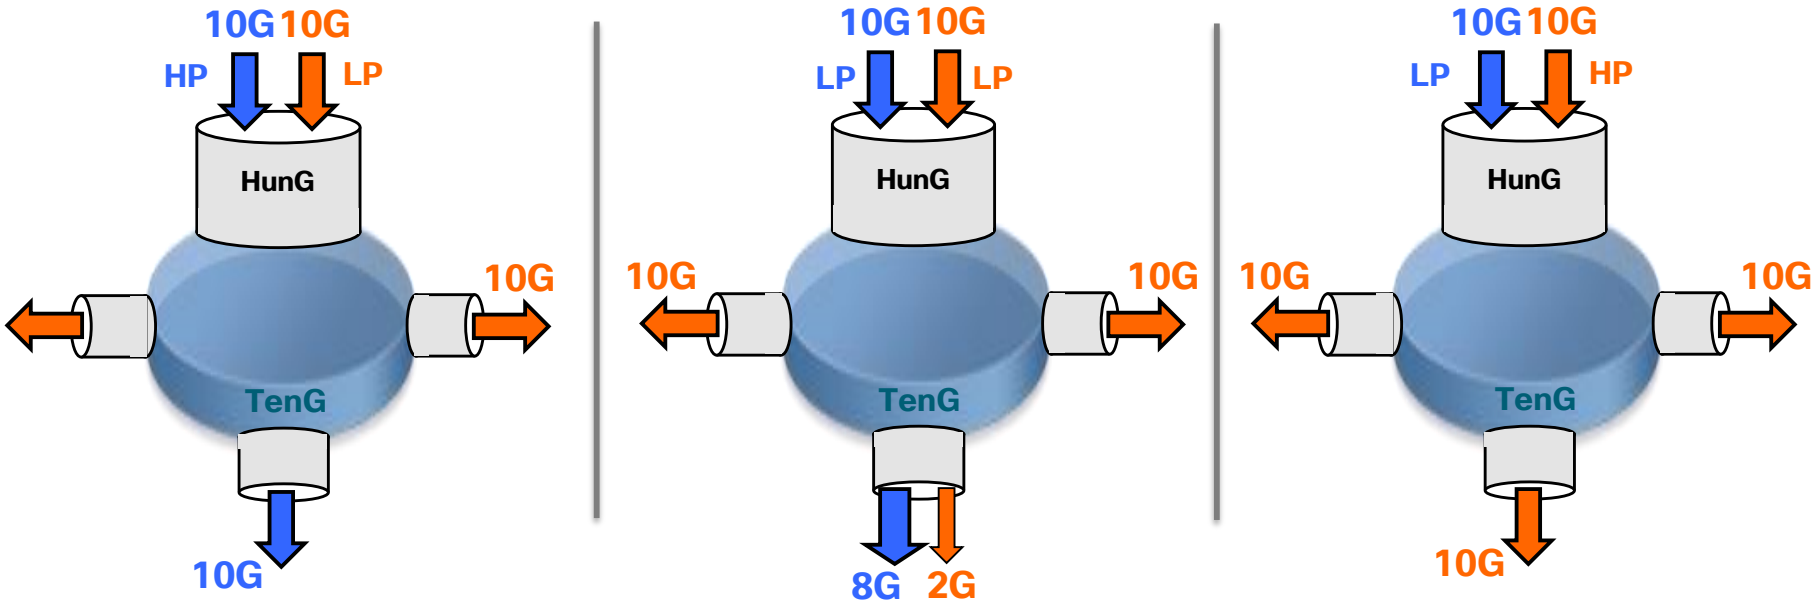

NCS5500 Gotchas / Good-to-know

You make networking **possible**

Cisco live!

NCS-5501-SE

100Mbps / 1Gbps Limitations

- NCS-5501-SE ports 0/8 to 0/15
  - Don't support 100Mbps copper SFP modules (GLC-T)
  - Don't support auto-neg for 1G optical SFP

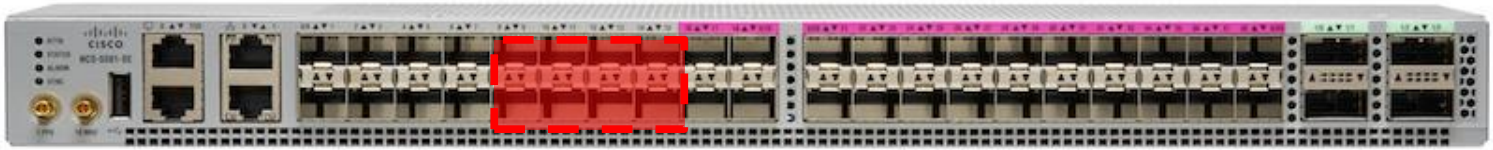

- NCS-5501-SE other SFP ports
  - Support 1G and 100M speeds
  - Support 1G Auto Neg (Clause 37)
- No limitation on the 48 ports of NCS-5501

NCS-55A2-MOD 10G/25G QUADs

- Ports 0/0/0/24 to 0/0/0/39 supports 1/10G or 25G
- Config per block of 4 ports aka « Quad »
  - Quad 0 : ports 0/0/0/24 to 0/0/0/27
  - Quad 1 : ports 0/0/0/28 to 0/0/0/31
  - Quad 2 : ports 0/0/0/32 to 0/0/0/35
  - Quad 3 : ports 0/0/0/36 to 0/0/0/39

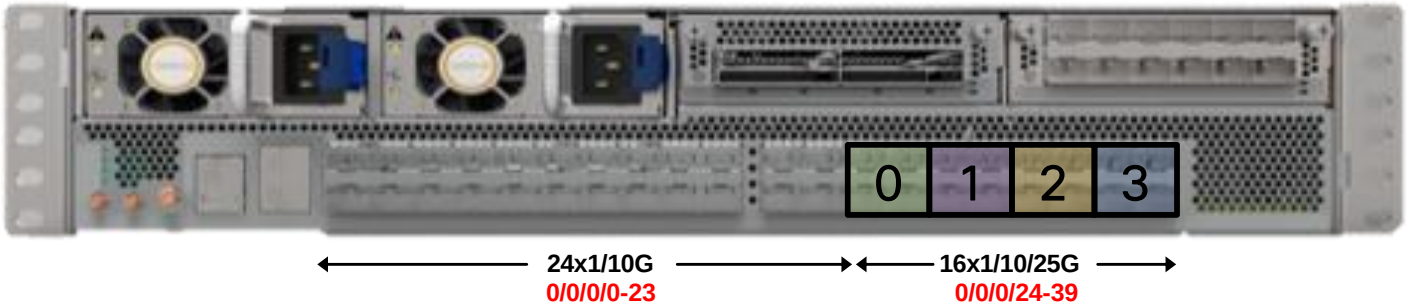

NCS-55A2-MOD 10G/25G QUADs

- Default on these ports is 25G « TF0/0/0/x »
- 1G/10G optics can NOT be mixed with 25G optics in the same quad
- 1G and 10G optics CAN co-exist in the same quad
- Configuration does not require reboot

NCS-55A2-MOD 10G/25G QUADs

```
RP/0/RP0/CPU0:55A2-MOD-SE(config)#do sh int brief | i 0/0/0/2
<SNIP>
         Te0/0/0/23  admin-down  admin-down              ARPA  1514   10000000
         TF0/0/0/24  admin-down  admin-down              ARPA  1514   25000000
         TF0/0/0/25  admin-down  admin-down              ARPA  1514   25000000
         TF0/0/0/26  admin-down  admin-down              ARPA  1514   25000000
         TF0/0/0/27  admin-down  admin-down              ARPA  1514   25000000
         TF0/0/0/28  admin-down  admin-down              ARPA  1514   25000000
         TF0/0/0/29  admin-down  admin-down              ARPA  1514   25000000
RP/0/RP0/CPU0:55A2-MOD-SE(config)#hw-module quad 0 location 0/0/CPU0 mode ?
  WORD  10g or 25g, (10g mode also operates 1g transceivers)
RP/0/RP0/CPU0:55A2-MOD-SE(config)#hw-module quad 0 location 0/0/CPU0 mode 10g
RP/0/RP0/CPU0:55A2-MOD-SE(config)#commit
RP/0/RP0/CPU0:55A2-MOD-SE(config)#do sh int brief | i 0/0/0/2
<SNIP>
         Te0/0/0/23  admin-down  admin-down              ARPA  1514   10000000
         Te0/0/0/24  admin-down  admin-down              ARPA  1514   10000000
         Te0/0/0/25  admin-down  admin-down              ARPA  1514   10000000
         Te0/0/0/26  admin-down  admin-down              ARPA  1514   10000000
         Te0/0/0/27  admin-down  admin-down              ARPA  1514   10000000
         TF0/0/0/28  admin-down  admin-down              ARPA  1514   25000000
         TF0/0/0/29  admin-down  admin-down              ARPA  1514   25000000
RP/0/RP0/CPU0:55A2-MOD-SE(config)#
```

Modular Port Adaptors

100Mbps / 1Gbps Limitations

- 1G supported on 8 ports out of 12 ports on MPA-12T-S
  - Ports 0-3 and 8-11

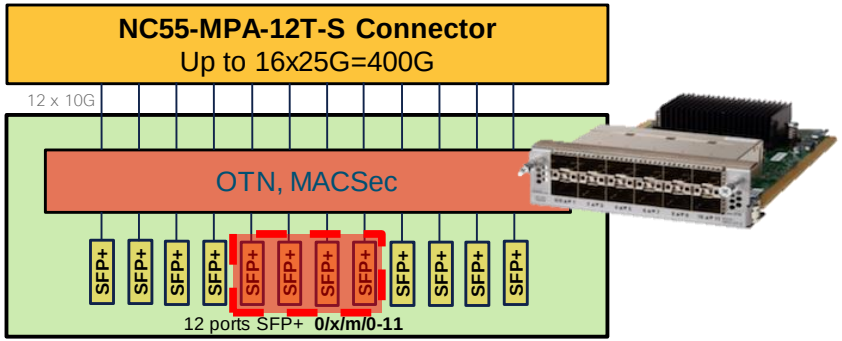

25G Support

- No support for 4x25G breakout on Jericho
  - Only supported on J+ systems and line cards
- Native SFP28 ports in
  - NCS-55A2-MOD*
  - NCS-55A1-48Q6H
  - NCS-55A1-24Q6H

Breakout

- First, interface name depends on optics inserted
  - QSFP+: Fo0/x/y/z
  - QSFP28: Hu0/x/y/z
- Breakout requires
  - Appropriate optics
  - Configuration

```
RP/0/RP0/CPU0:NCS5500(config)#controller optics 0/0/0/2
RP/0/RP0/CPU0:NCS5500(config-Optics)# breakout 4x10
```

  - Interface name is changed to 25G (TF) or 10G (Te) with a 5th tuple
    - Fo0/x/y/z becomes **Te0/x/y/z/b**

100G ER4L Configuration

- ER4L uses Forward Error Correction (FEC) to reach 40km
- No configuration required with similar systems back to back

```
RP/0/RP0/CPU0:router#show controllers HundredGigE0/0/0/13 all | in Forward
Forward error correction: Standard (Reed-Solomon)
```

- If remote system does not support RS-FEC, reach is 25km

```
RP/0/RP0/CPU0:router(config)#interface HundredGigE <0/2/0/8>
RP/0/RP0/CPU0:router(config-if)#fec ?
base-r    Enable BASE-R FEC
none      Disable any FEC enabled on the interface
standard Enable the standard (Reed-Solomon) FEC
RP/0/RP0/CPU0:router(config-if)#fec none
```

Timing

- RP-E + J+ line cards needed on chassis
- Not supported
  - On Jericho-based systems, except NCS-5501-SE
  - On 1G mode on SFP28 and QSFP28 interfaces
  - On 1G SFP copper interfaces
- Roadmap
  - Support on logical interfaces: bundle, BVI and loopback
  - Support on MPLS interfaces

QoS on Sub-Interface

- By default egress QoS can be applied on main interfaces only
- To enable it on L2/L3 sub-if you need to configure HQoS mode:

```
RP/0/RP0/CPU0:Router(config-subif)#show config failed
!! SEMANTIC ERRORS: This configuration was rejected by
!! the system due to semantic errors. The individual
!! errors with each failed configuration command can be
!! found below.

interface TenGigE0/0/0/0.1
service-policy output CORE-OUTPUT-QOS
!!% 'DNX_QOSEA' detected the 'warning' condition 'QoS is supported on sub-interface(s) only in
Hierarchical QoS Mode.'
!
end

RP/0/RP0/CPU0:Router(config-subif)#exit
RP/0/RP0/CPU0:Peyto-SE(config)#hw-module profile qos hqos-enable
In order to activate this new qos profile, you must manually reload the chassis/all line cards
RP/0/RP0/CPU0:Peyto-SE(config)#
```

HealthCheck

- Some useful show commands to track the router health

```
show hw-module fpd
show media
show watchdog memory-state location all
show health gsp
show health cfgmgr
show health sysdb
show asic-errors all summary location <LC>
show dpa resource iproute location <LC>
show dpa resource ip6route location <LC>
show contr npu resource all loc <LC>
admin show controllers fabric health
admin show environment temperatures
admin show environment fan
admin show environment power
admin show vm
```

For Reference

Hardware Profiles

```
RP/0/RP0/CPU0:Peyto-SE(config)#hw-module ?
  fib                Forwarding table to configure
  oversubscription   Configure oversubscription
  profile            Configure profile.
  quad               Configure quad.
  route-stats        Configure multicast per-route statistics
  service            Configure service role.
  subslot            Configure subslot h/w module
  tcam               Configure profile for TCAM LC cards
  vrrpscale          to scale VRRP sessions
RP/0/RP0/CPU0:Peyto-SE(config)#hw-module profile ?
  acl               Configure acl profile
  bundle-scale      Max number of bundles supported
  bw-threshold      Asic Fabric Link Bandwidth Availability Threshold
  flowspec          Configure support for v6 flowspec
  l2                Configure l2 profile
  load-balance      Configure load balance parameters
  netflow           Configure Netflow profile.
  offload           Offload profile in NCS5501-SE
  qos               Configure qos profile
  segment-routing   Segment routing options
  sr-policy         SR Policy options
  stats             Configure stats profile.
  tcam              Configure profile for TCAM LC cards
RP/0/RP0/CPU0:Peyto-SE(config)#
```

Conclusion

You make networking **possible**

Conclusion

- Merchant silicon is not something new in SP portfolio
- Many form factors
- NCS5500 can be used in multiple roles in Networks such as
  - Core, Peering, SP DC, Aggregation and Edge: You decide.
- Architecture based on VOQ-only for unicast and FMQ for multicast
- Compared to traditional IOS XR platforms
  - Resources needs to be monitored differently
  - Features can have a different implementation

Complete your online session evaluation

- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on Thursday) to receive your Cisco Live water bottle.
- All surveys can be taken in the Cisco Live Mobile App or by logging in to the Session Catalog on ciscolive.cisco.com/us.

Cisco Live sessions will be available for viewing on demand after the event at ciscolive.cisco.com.

Continue your education

Demos in the Cisco campus

Walk-in labs

Meet the engineer 1:1 meetings

Related sessions

NDA Roadmap Sessions at Cisco Live
Customer Connection Member Exclusive

Join Cisco's online user group to ...

Connect online with 29,000 peer and Cisco experts in private community forums

Learn from experts and stay informed about product roadmaps

- Roadmap sessions at Cisco Live
- Monthly NDA briefings

Give feedback to Cisco product teams

- Product enhancement ideas
- Early adopter trials
- User experience insights

NETWORKING ROADMAPS	SESSION ID	DAY / TIME
Roadmap: SD-WAN and Routing	CCP-1200	Mon 8:30 – 10:00
Roadmap: Machine Learning and Artificial Intelligence	CCP-1201	Tues 3:30 – 5:00
Roadmap: Wireless and Mobility	CCP-1202	Thurs 10:30 – 12:00

Join at the Customer Connection Booth
(in the Cisco Showcase)

Member Perks at Cisco Live
- Attend NDA Roadmap Sessions
- Customer Connection Jacket
- Member Lounge

Join online: www.cisco.com/go/ccp

Thank you

You make **possible**

NCS5500 TCAM Carving

You make networking **possible**

Cisco live!

Default eTCAM Carving

Jericho w/ eTCAM

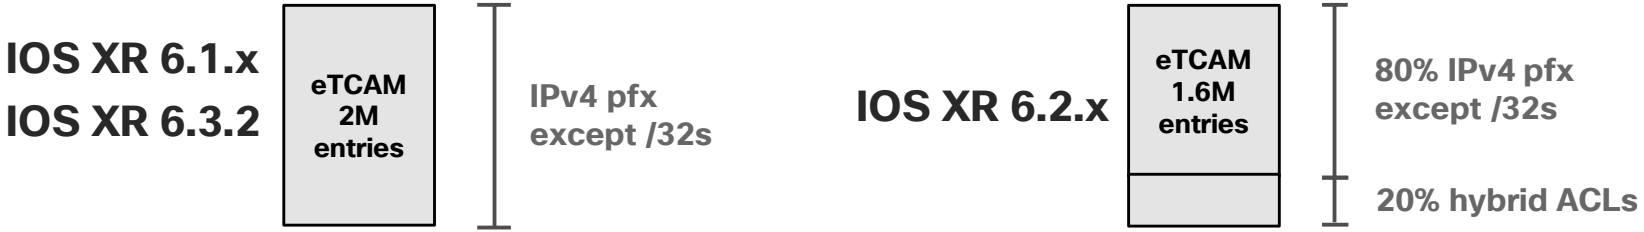


```
RP/0/RP0/CPU0:NCS5508-6.3.2#sh contr npu ext loc 0/6/CPU0

External TCAM Resource Information
=============================================================
NPU   Bank    Entry   Owner        Free       Per-DB   DB    DB
      Id      Size                 Entries    Entry    ID    Name
=============================================================
0     0       80b     FLP          2047993    7        15    IPV4 DC
1     0       80b     FLP          2047993    7        15    IPV4 DC
2     0       80b     FLP          2047993    7        15    IPV4 DC
3     0       80b     FLP          2047993    7        15    IPV4 DC

RP/0/RP0/CPU0:NCS5508-6.3.2#
```

```
RP/0/RP0/CPU0:TME-5508-6.2.3#sh contr npu externaltcam loc 0/6/CPU0

External TCAM Resource Information
=============================================================
NPU   Bank    Entry   Owner        Free       Per-DB   DB    DB
      Id      Size                 Entries    Entry    ID    Name
=============================================================
0     0       80b     FLP          498950     1139450  15    IPV4 DC
0     1       80b     FLP          28672      0        76    INGRESS_IPV4_SRC_IP_EXT
0     2       80b     FLP          28672      0        77    INGRESS_IPV4_DST_IP_EXT
0     3       160b    FLP          26624      0        78    INGRESS_IPV6_SRC_IP_EXT
0     4       160b    FLP          26624      0        79    INGRESS_IPV6_DST_IP_EXT
0     5       80b     FLP          28672      0        80    INGRESS_IP_SRC_PORT_EXT
0     6       80b     FLP          28672      0        81    INGRESS_IPV6_SRC_PORT_EXT
...
```

Default eTCAM Carving

Jericho w/ eTCAM with URPF Loose

- Activating URPF requires to disable the eTCAM dual capacity mode

```
RP/0/RP0/CPU0:NCS5508(config)#hw-module tcam fib ipv4 scaledisable
RP/0/RP0/CPU0:NCS5508(config)#commit
```

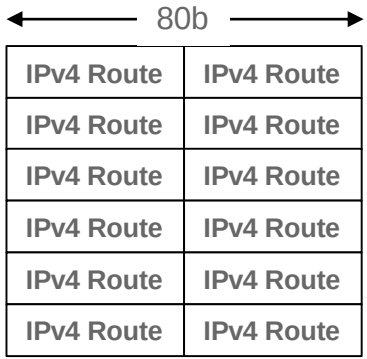

Double capacity mode

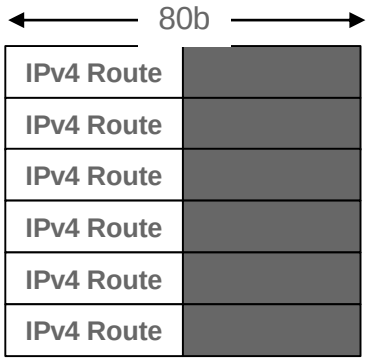

Double capacity mode disabled

Default eTCAM Carving

Jericho w/ eTCAM with URPF Loose

- It effectively reduces the eTCAM size by half

IOS XR 6.1.x
IOS XR 6.3.2

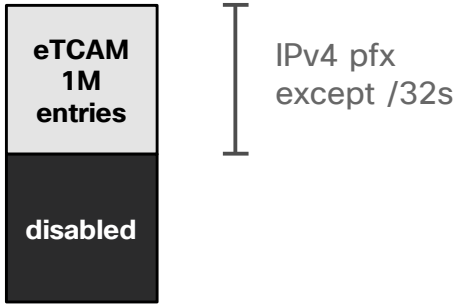

IOS XR 6.2.x

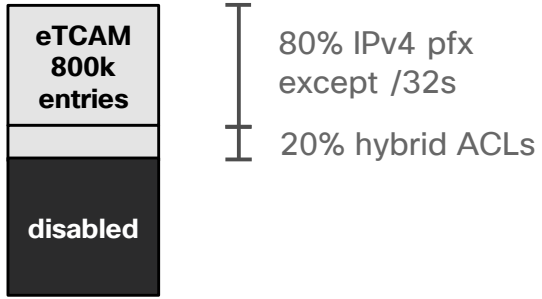

Default eTCAM Carving

Jericho+ w/ eTCAM

- In 6.3.2, the system is validated for 4M v4 routes (with or without uRPF)
- Hybrid ACL objects are stored in a different zone and don't impact the scale

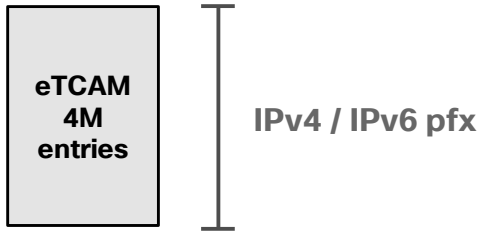

Modifying eTCAM Carving

Jericho w/ eTCAM

- It's advised to configure a total of 100% for predictable results

```
RP/0/RP0/CPU0:R(config)#hw-module profile tcam fib ipv4 unicast percent 50
RP/0/RP0/CPU0:R(config)#hw-module profile tcam fib ipv6 unicast percent 50
RP/0/RP0/CPU0:R(config)#commit
```

- After reload of the line cards

```
RP/0/RP0/CPU0:R#show controllers npu diag kbp dbstats instance 0 location 0/7/CPU0

 Statistics Rack: 0, Slot: 7, Asic instance: 0

Table Configuration

Tbl-ID    Tbl-Name       Size      Width     AD Width  Num ent.  ~Capacity Shuffles
-----------------------------------------------------------
 3        IPv6 UC        256000    160       64        7         51200     0
 4        IPv6 RPF       256000    160       32        0         51200     0
 15       IPV4 DC        1024000   80        48        5         1024000   0
```

Modifying eTCAM Carving

Jericho w/ eTCAM

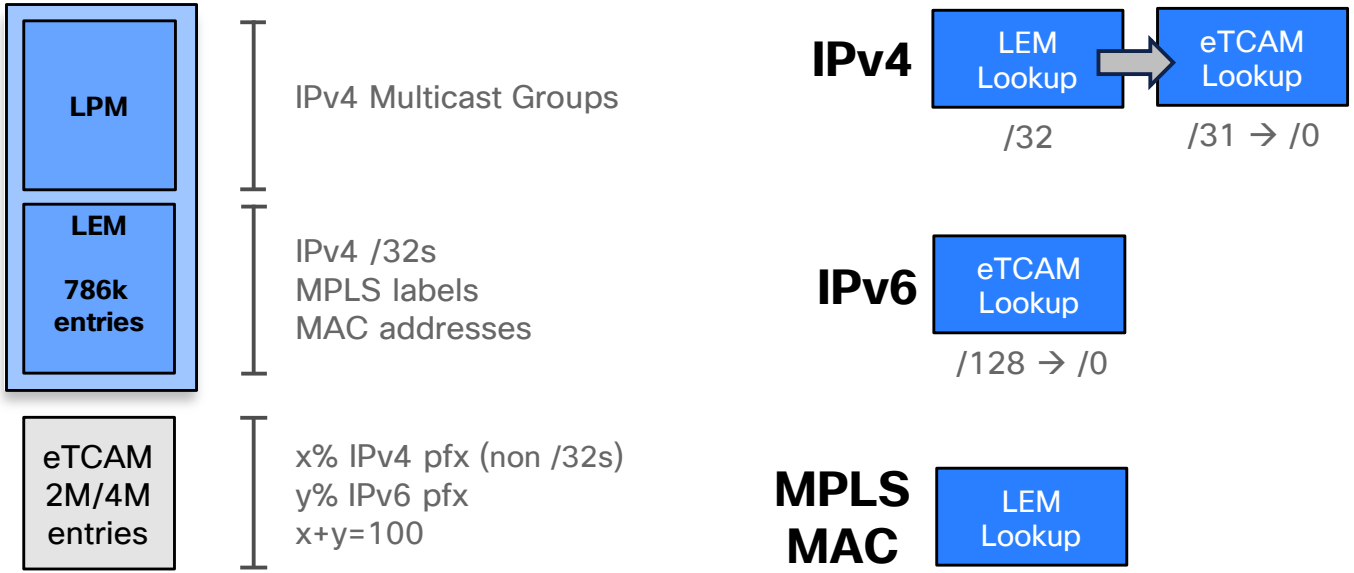

Only v4/32s are programmed in LEM
All other v4/v6 routes go to eTCAM except
if x=100 / y=0, IPv6 will be moved to LEM/LPM

Configuring 100% IPv6 in eTCAM is not possible,
but 1% / 99% is accepted

Mixing Scale and Base LineCards

You make networking **possible**

Cisco live!

Selective Route Download Feature

- eTCAM and non-eTCAM can co-exist in the same chassis
  - It's possible to select routes that will be programmed in scale line cards only
  - In BGP configuration
    - using a table-policy and a specific path-color "external-reach"
- With this feature
  - IGP routes will be programmed in both LC types
  - BGP routes with path-color external-reach will be programmed in Scale LC only
  - Other BGP routes will programmed in both LC types

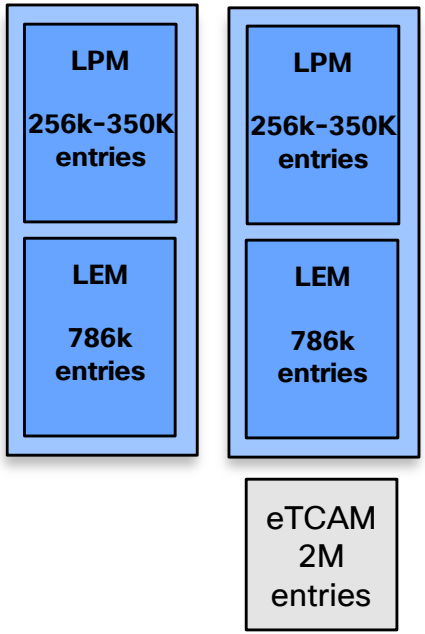

Selective Route Download Configuration

```
route-policy PEER-EXT
  set community PEER-EXT-comm
end-policy
!
route-policy HILO-FIB
  if community matches-any PEER-EXT-comm then
    set path-color external-reach
    pass
  else
    pass
  endif
end-policy!

router bgp 100
address-family ipv4 unicast
  table-policy HILO-FIB
!
!
neighbor 192.168.100.151
  address-family ipv4 unicast
   route-policy PEER-EXT in
   maximum-prefix 8000000 75
   route-policy PERMIT-ANY out
!
```

Selective Route Download Verification

- Check a route

```
RP/0/RP0/CPU0:NCS5508-1-631#sh route 1.0.144.0/20

Routing entry for 1.0.144.0/20
  Known via "bgp 100", distance 200, metric 0, external-reach-lc-only
  Tag 2914, type internal
  Installed Nov 27 22:48:56.925 for 00:00:45
  Routing Descriptor Blocks
    192.168.100.151, from 192.168.100.151
      Route metric is 0
  No advertising protos.
RP/0/RP0/CPU0:NCS5508-1-631#
```

Selective Route Download Verification

```
RP/0/RP0/CPU0:NCS5508-1-631#sh cef 1.0.144.0/20 detail

1.0.144.0/20, version 25081094, external-reach-lc-only, internal 0x5000001 0x0 (ptr 0x8f485390) [1], 0x0
(0x0), 0x0 (0x0)
 Updated Nov 27 22:48:56.929
 local adjacency 192.168.100.151
 Prefix Len 20, traffic index 0, precedence n/a, priority 4
  gateway array (0x8e0e9250) reference count 655801, flags 0x2010, source rib (7), 0 backups
                [1 type 3 flags 0x48501 (0x8e18f758) ext 0x0 (0x0)]
  LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
  gateway array update type-time 1 Nov 27 22:48:56.929
 LDI Update time Nov 27 22:48:56.929
   via 192.168.100.151/32, 2 dependencies, recursive [flags 0x6000]
    path-idx 0 NHID 0x0 [0x8e0bf1b0 0x0]
    next hop 192.168.100.151/32 via 192.168.100.151/32

    Load distribution: 0 (refcount 1)

    Hash  OK  Interface                 Address
    0     Y   MgmtEth0/RP0/CPU0/0       192.168.100.151

RP/0/RP0/CPU0:NCS5508-1-631#
```

Selective Route Download Use-Case

- Lookup executed in ingress only
- Position of the Base and Scale line card is opposite than ASR9k or CRS
- Internet-facing interface could be DWDM card or MACsec card

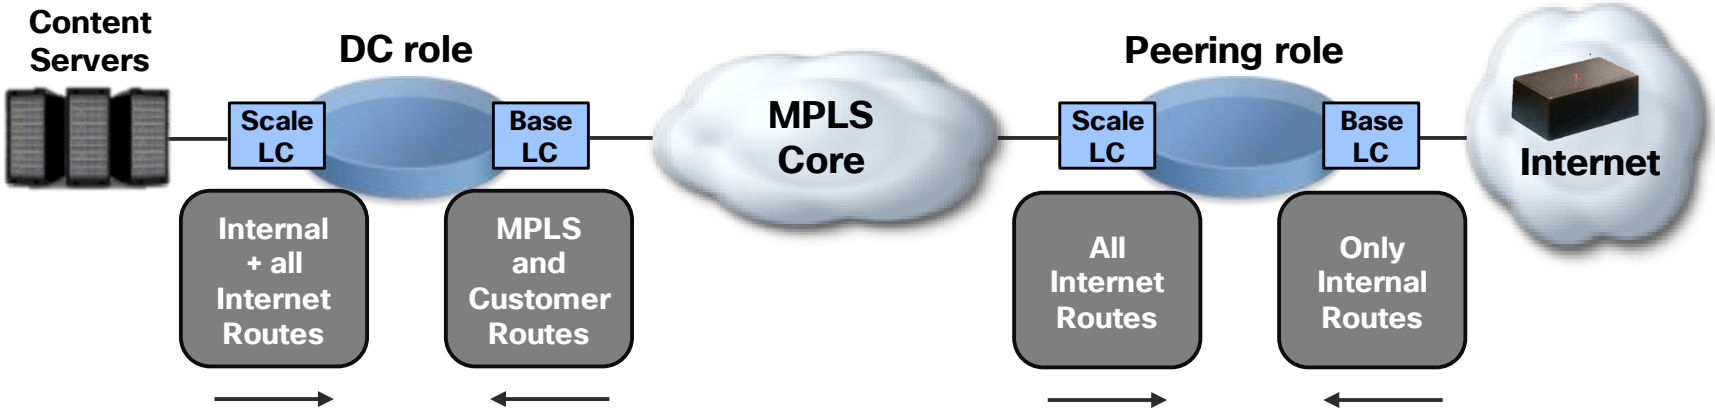

Demo

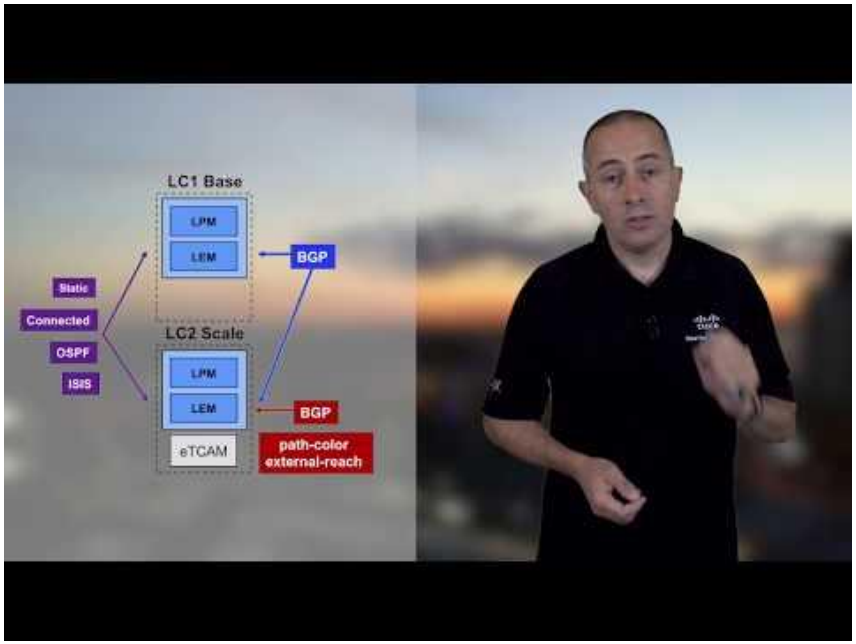

http://bit.ly/ncs5500-mix

NCS5500 Internals

You make networking **possible**

Cisco live!

NCS5500 System Architecture

Intra-Chassis Communication

- EOBC and EPC: two isolated networks
- EOBC network: Ethernet Out-of-Band Channel
  - Used for inter-process communication (IPC)
- EPC network: Ethernet Protocol Channel
  - Used for packet punt (all “for-us packets”)
- EMON
  - Kernel process running on all cards and managing the path
  - Replaces spanning tree to offer loop free topology
  - HeartBeat (HB) every 40ms, 5 misses → failure
- System Controller
  - All these messages are going through the SC cards in NCS-5508 chassis

NCS5500 Internals

EOBC in Modular Chassis

- Ethernet Out-of-Band Channel
  - Intra-system management communication
  - EOBC channel is provided via a switch chipset on the System Controllers that inter-connects all modules together, including RPs, Fabric Cards and Line Cards

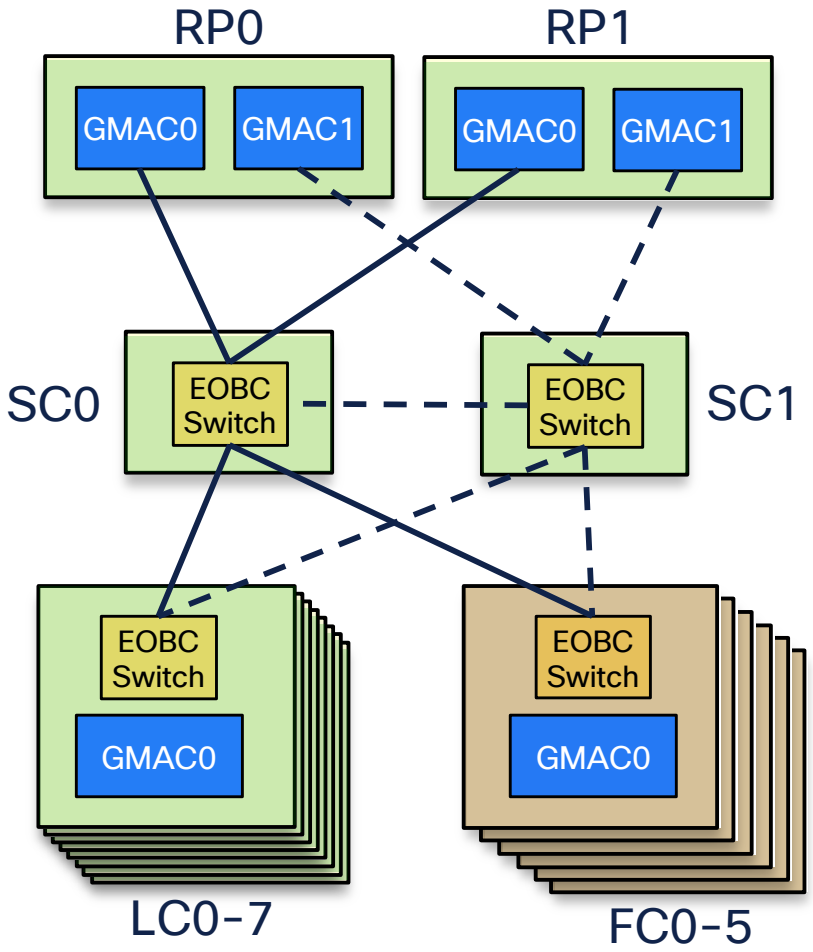

NCS5500 Internals

EPC in Modular Chassis

- Ethernet Protocol Channel
  - Intra-system data plane protocol communication
  - EPC switch only connects Fabric Cards to RPs
  - If protocol packets need to be sent to RP, line cards utilize the internal data path to transfer packets to Fabric Cards first, Fabric Cards then redirect them via the EPC channel to supervisor engines
  - Uses different VLAN for different traffic types (one VLAN per NPU for Netflow sampled packets)

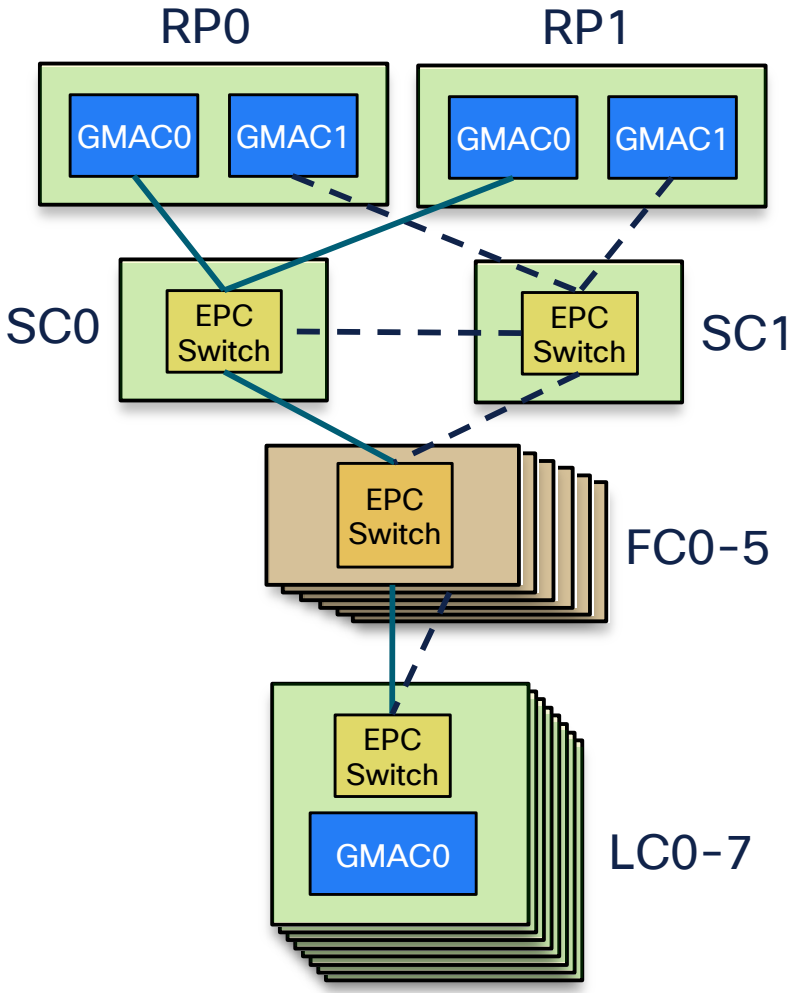

NCS5500 Internals

Internal Switches in Modular Chassis

```
sysadmin-vm:0_RP0# show controller switch reachable

Rack  Card  Switch
---------------------
0     SC0   SC-SW
0     SC0   EPC-SW
0     SC0   EOBC-SW
0     SC1   SC-SW
0     SC1   EPC-SW
0     SC1   EOBC-SW
0     LC0   LC-SW
0     LC1   LC-SW
0     LC3   LC-SW
0     FC0   FC-SW
0     FC1   FC-SW
0     FC2   FC-SW
0     FC3   FC-SW
0     FC4   FC-SW
0     FC5   FC-SW

sysadmin-vm:0_RP0#
```

EPC switch
EOBC swith
Both EOBC and EPC

NCS5500 Internals

EPC/EOBC Switches

- In Line Cards, switches are shared for EPC/EOBC
- Different bandwidth depending on the LC type (1G, 2.5G)
- Only one Fabric Card link is forwarding

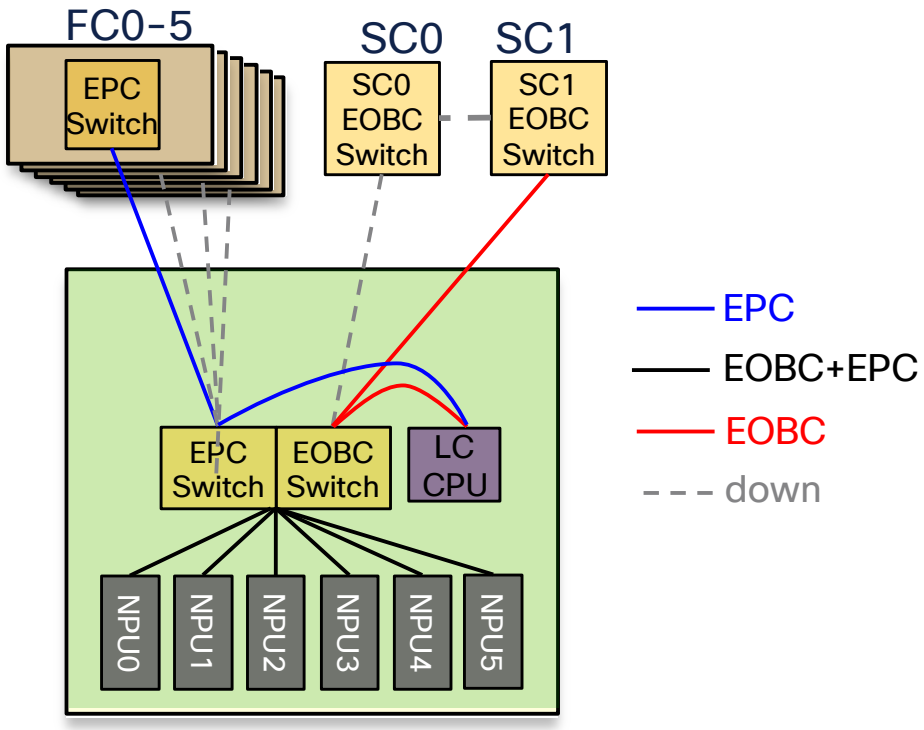

Example: EPC/EOBC in 24x100G Line Cards

```
sysadmin-vm:0_RP0# show controller switch summary location 0/LC7/LC-SW
Rack  Card  Switch  Rack Serial Number
-------------------------------------
0     LC7   LC-SW   FGE194714QQ

      Phys   Admin  Port       Protocol  Forward
Port  State  State  Speed      State     State       Connects To
-------------------------------------------------------------------------
4     Up     Up     2.5-Gbps   -         Forwarding  LC CPU (EPC 0)
5     Up     Up     2.5-Gbps   -         Forwarding  LC CPU (EPC 1)
6     Up     Up     2.5-Gbps   -         Forwarding  LC CPU (EPC 2)
7     Up     Up     2.5-Gbps   -         Forwarding  LC CPU (EOBC)
8     Up     Up     2.5-Gbps   -         Forwarding  NPU2
9     Up     Up     2.5-Gbps   -         Forwarding  NPU1
10    Up     Up     2.5-Gbps   -         Forwarding  NPU0
11    Up     Up     2.5-Gbps   -         Forwarding  NPU3
12    Up     Up     1-Gbps     -         Forwarding  FC0
13    Down   Down   1-Gbps     -         -           FC1
14    Down   Down   1-Gbps     -         -           FC2
15    Down   Down   1-Gbps     -         -           FC3
16    Down   Down   1-Gbps     -         -           FC4
17    Down   Down   1-Gbps     -         -           FC5
18    Up     Up     1-Gbps     -         Forwarding  SC0 EOBC-SW
19    Down   Down   1-Gbps     -         -           SC1 EOBC-SW
sysadmin-vm:0_RP0#
```