
cisco
You make possible

Cisco ASR 9000 System Architecture

Yongzhong Peng
Manager, Technical Marketing
SP Routing Infrastructure

BRKARC-2003

cisco Live!
Barcelona | January 27-31, 2020

Cisco Webex Teams

Questions?

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

How

1. Find this session in the Cisco Events 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

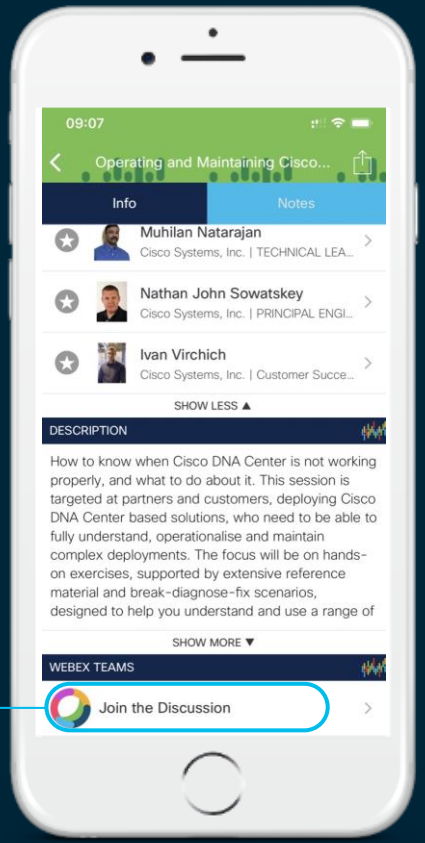

Agenda

- ASR9000 Products Introduction
- ASR9000 System Hardware Architecture
- ASR9000 Distributed Control Plane
- ASR9000 Data Packet Processing
- ASR9000 QoS Architecture & TCAM Usage
- Conclusion

ASR 9000 Products Introduction

Cisco ASR 9000 System Portfolio

Compact & Powerful Access/Aggregation	Flexible Service Edge	High Density Service Edge and Core
• Small footprint with full IOS-XR for distributed environments	• Optimized for MSE with high M-D scale for medium to large sites	• Scalable, ultra high density service routers for large, high-growth sites

One Platform, One OS, One Family

nV Satellites
ASR 9000v2, NCS5000

ASR 9901

456Gbps

ASR 9001

Fixed 2RU

240 Gbps

ASR 9904

2 LC/6RU

16 Tbps

ASR 9006

4 LC/10RU

16 Tbps

ASR 9906

4 LC/14RU

32 Tbps

ASR 9010

8 LC/21RU

32 Tbps

ASR 9910

8 LC/21RU

64 Tbps

ASR 9912

10 LC/30RU

80 Tbps

ASR 9922

20 LC/44RU

160 Tbps

MSE | DCI | Peering | P/PE | CE | Mobility | Broadband

ASR 9000 Modular Systems

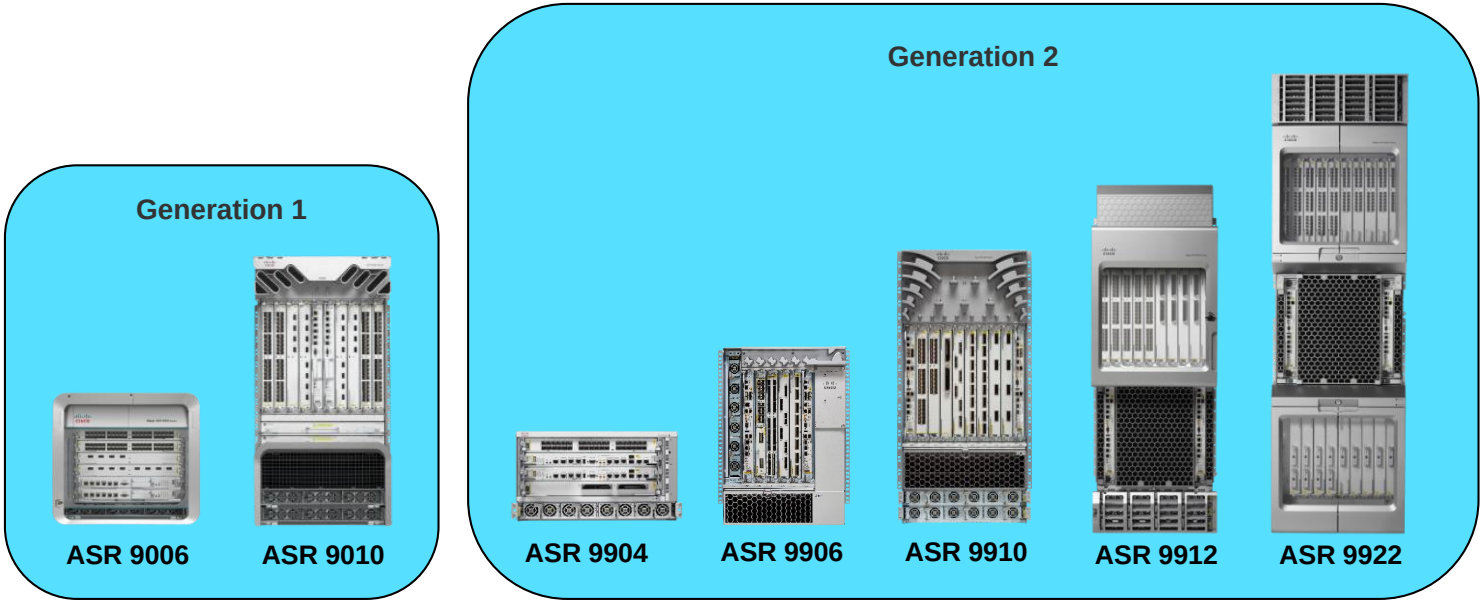

- Up to 2T/slot
- A9K 5-Fabric Card Support

- Up to 4T/slot
- A9K 5-Fabric & A99 7-Fabric Card Support

ASR 9000 Hardware Components – Control Module

RSP: Route Switch Processor

ASR 9910

ASR 9010

ASR 9006

ASR 9904

ASR 9906

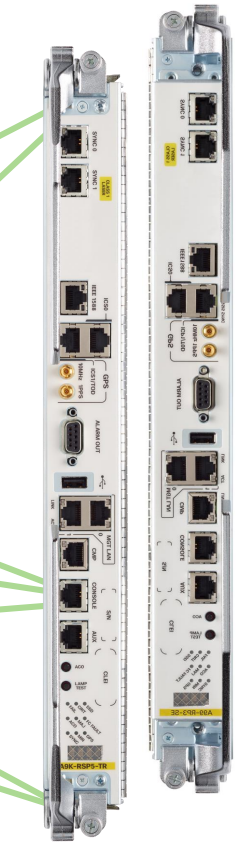

ASR 9922

ASR 9912

RP: Route Processor

RSP5/RP3 Front Panel Configuration

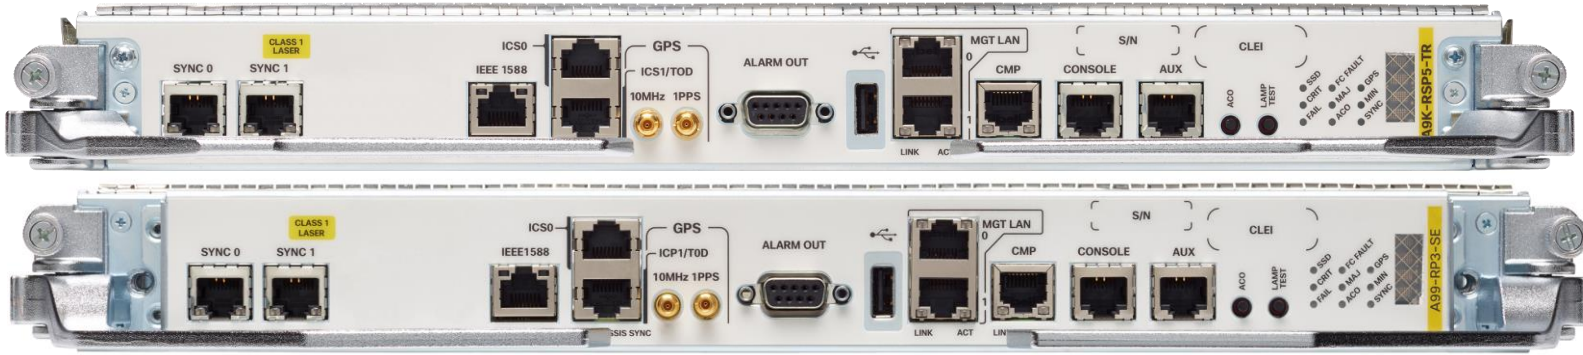

- 2x BITS ports on RJ-45
- 100Mbps, 1588 port – RJ-45
- TOD – RJ-45
- 10Mhz on SMA
- 1x CMP
- 1PPS on SMA

- Alarm output serial port
- USB
- 2x Management ports on RJ-45
- AUX & Console on RJ-45 connectors
- LED's for major/critical and normal oper alarms or states

ASR 9000 Hardware Components – Linecards

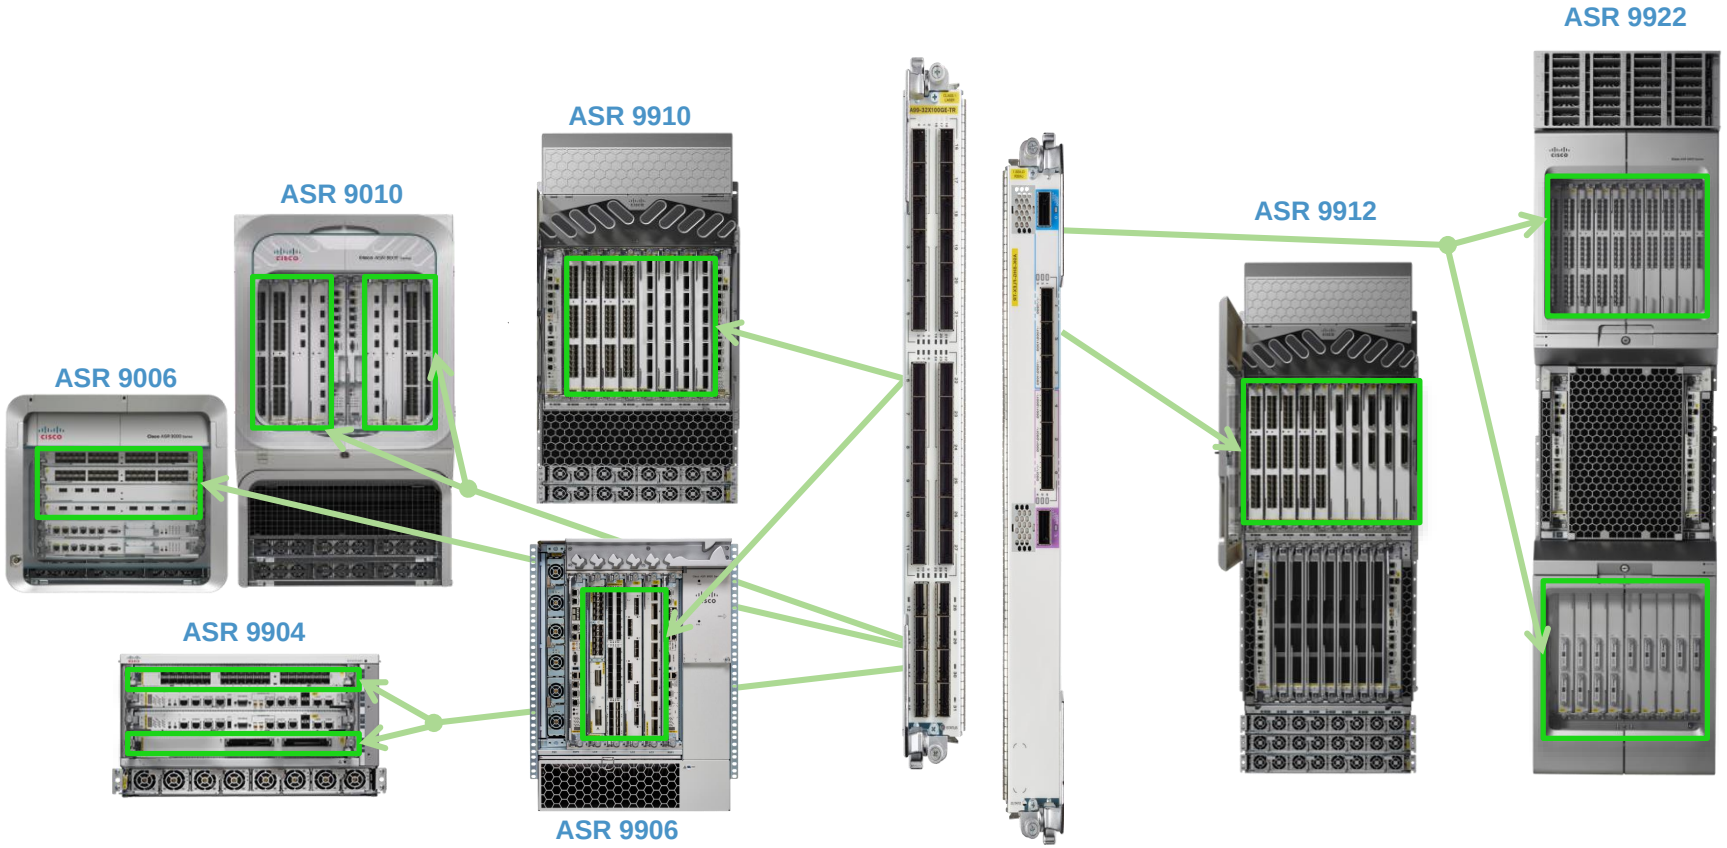

ASR 9900 – Switch Fabric Cards

ASR 9906

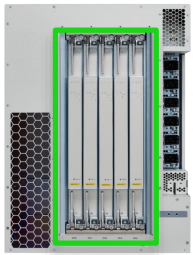

ASR 9910

ASR 9912

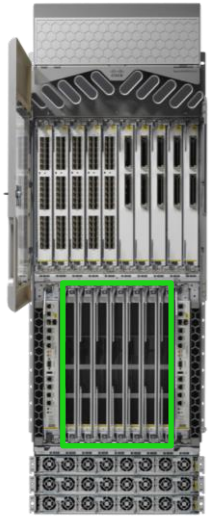

ASR 9922

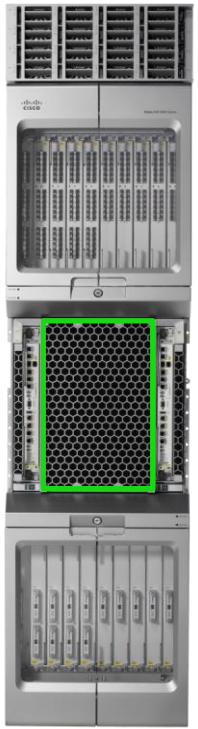

- ✓ Common for ASR 9912 and ASR 9922
- ✓ SFC-S for ASR 9910
- ✓ 7 Fabric Card Slots
- ✓ 5 Fabric Card Slots for ASR 9910 (2 on RSPs => total 7)
- ✓ Decoupled, multi-stage switch fabric hardware
- ✓ True HW separation between control and data plane
- ✓ Add bandwidth per slot easily & independently

	SFC2/SFC-S	SFC3
Fabric Capacity per SFC	230G	600G
Fabric Capacity Per Line Card Slot	1.38T N+1 1.61T N+0	3.6T N+1 4.2T N+0
Fabric Redundancy	N+1	N+1
LC Support	Typhoon Tomahawk	Tomahawk LightSpeed LightSpeed Plus

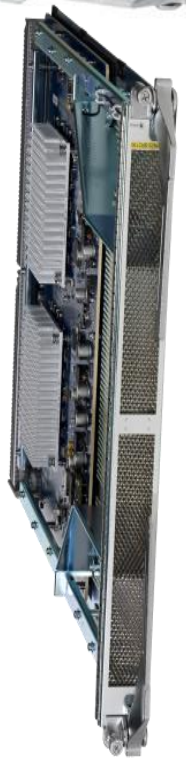

cisco Live!

ASR 9000 Systems Switch Fabric Overview

Integrated Fabric on RSP

ASR 9904 ASR 9006 ASR 9010

Hybrid Systems

ASR 9906 ASR 9910

Dedicated Fabric Cards

ASR 9912 ASR 9922

Integrated Fabric/RP/LC

ASR 9001 ASR 9901

4th Generation Hardware Commons PIDs

- *ASR-9922*
  - A99-RP3-SE/TR
  - A99-SFC3
  - ASR-9922-FAN-V3
- *ASR-9912*
  - A99-RP3-SE/TR
  - A99-SFC3
- *ASR-9910*
  - A9K-RSP5-SE/TR
  - A99-SFC3-S
- *ASR-9906*
  - A9K-RSP5-SE/TR
  - A99-SFC3-T
- *ASR-9904 / ASR-9010 / ASR-9006*
  - A9K-RSP5-SE/TR

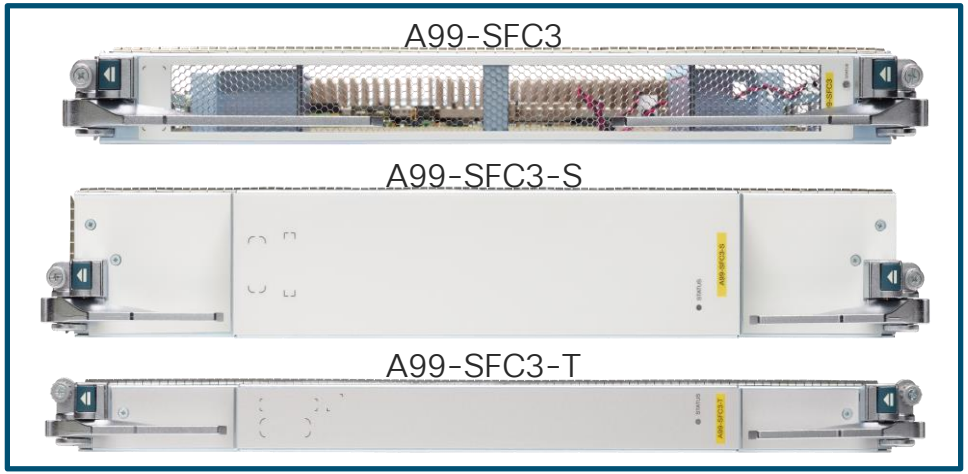

Cisco ASR 9906 Overview

Feature	Description
Total Capacity	>32T
Capacity per Slot	>4T
Slots	6 slots - 4 Line Cards and 2 RSPs
Rack size	14RU, front-to-back airflow.
Power	1 Power Tray 4.4 KW DC supplies (max 4) 6.0 KW AC supplies (max 3)
Fan	Front-to-Back Airflow 2 Fan Trays, FRU
RSPs	Integrated Fabric, 1+1 Redundancy
Fabric Cards	5 Fabric Cards + 2 RSP Integrated Fabric Total 7 Fabric for 6+1 Redundancy 230G per FC2 or 600G per FC3
Line cards	Tomahawk LightSpeed LightSpeed Plus Service Card: VSM with 32-bit XR

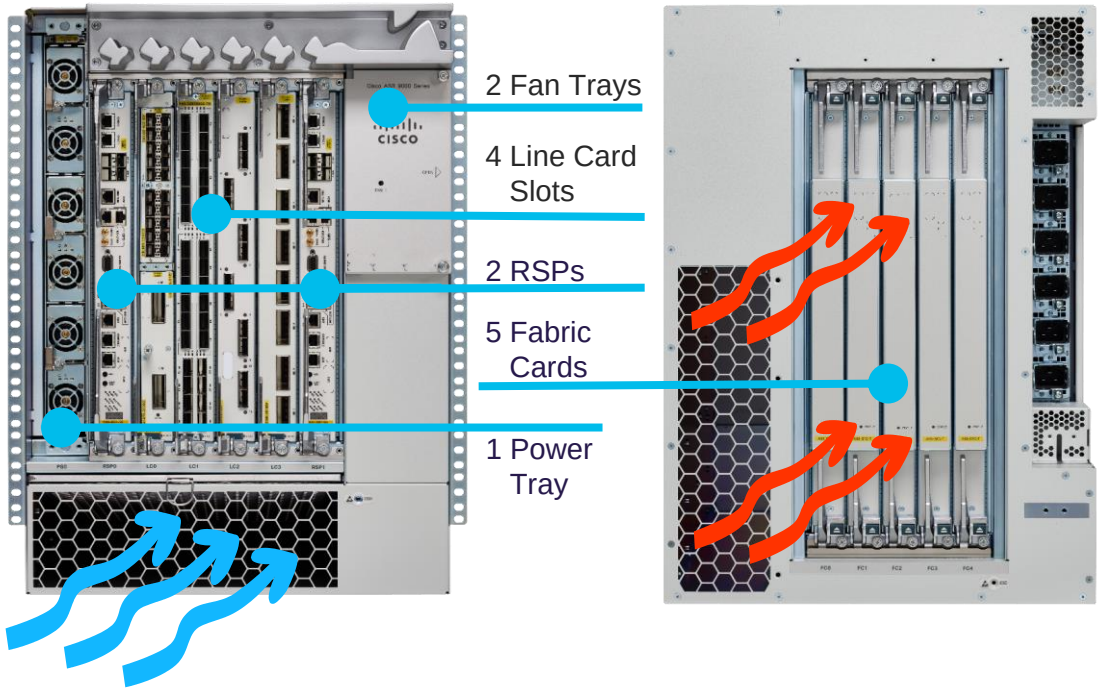

ASR 9000 Silicon Evolution

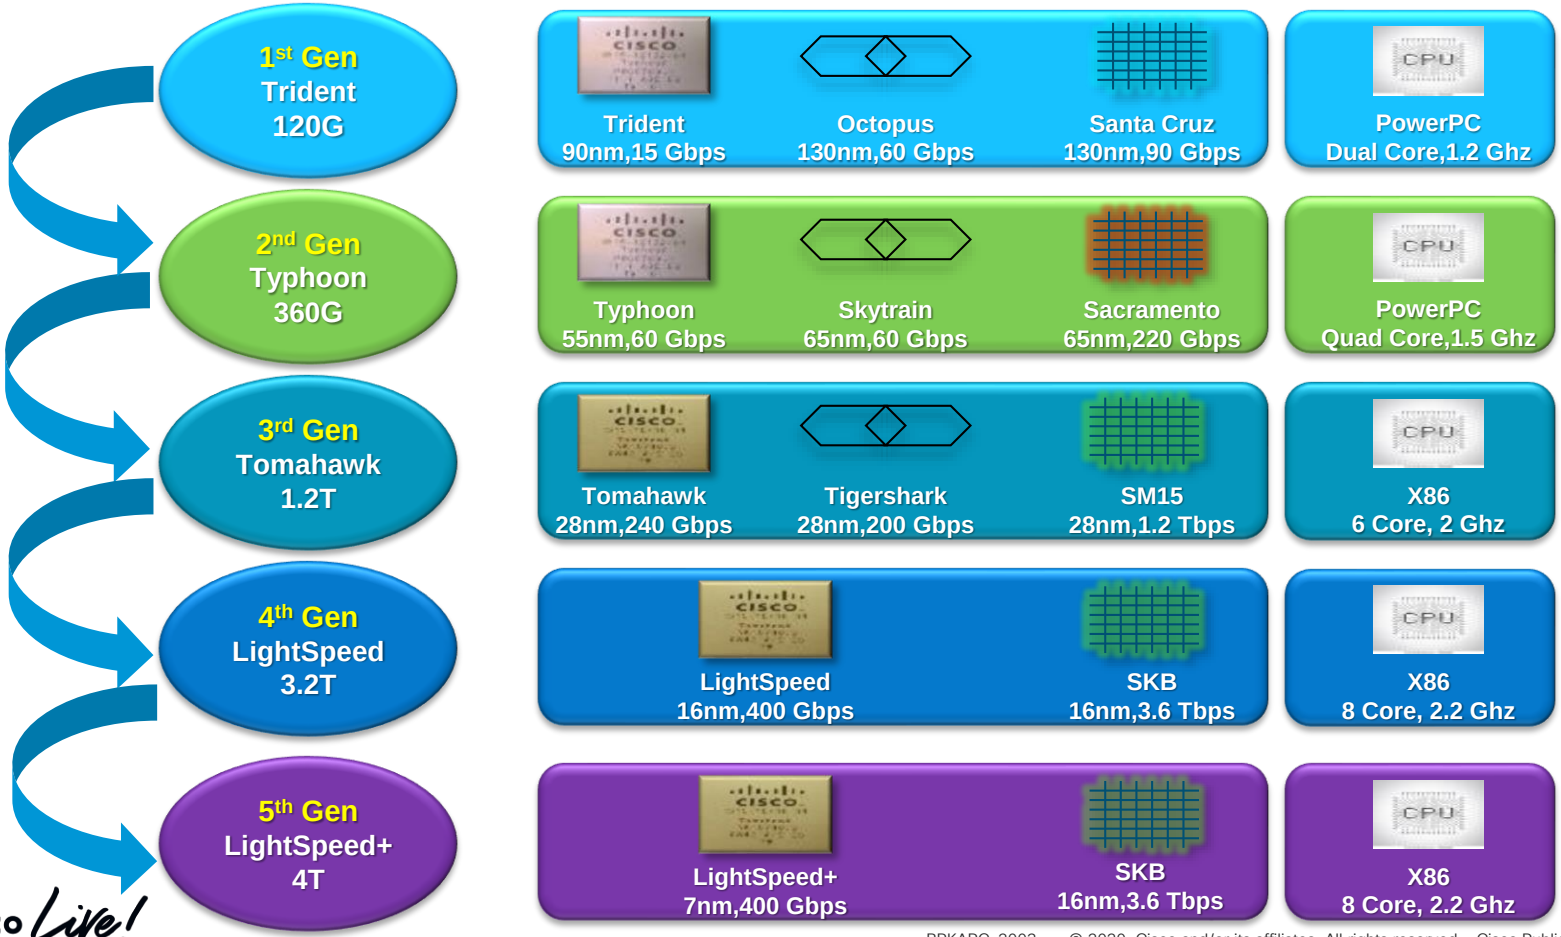

ASR 9000 3rd Generation Line Card Overview

MPAs
20x1GE
2x10GE
4x10GE
1x40GE
2x40GE

MPAs
1x100GE
2x100GE
20x10GE

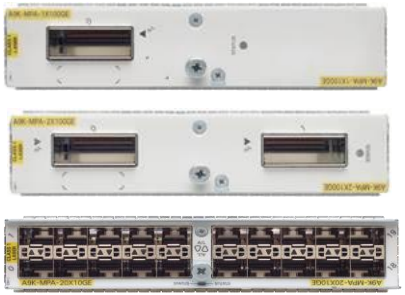

MPA
32x1GE

3rd gen LC Tomahawk
NPU: 240Gbps,
~150Mpps

4x100GE & 8x100G CPAK OTN Line Card

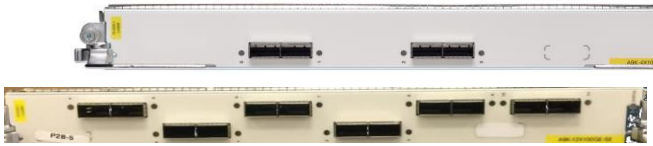

4x100GE & 12x100G QSFP28 LAN Line Card

-TR, -SE

A9K-24X10GE-1G

A9K-48X10GE-1G

MOD400/MOD200

A9K-400G-DWDM

ASR 9000 4th Generation Line cards

1.6T-3.2T per slot

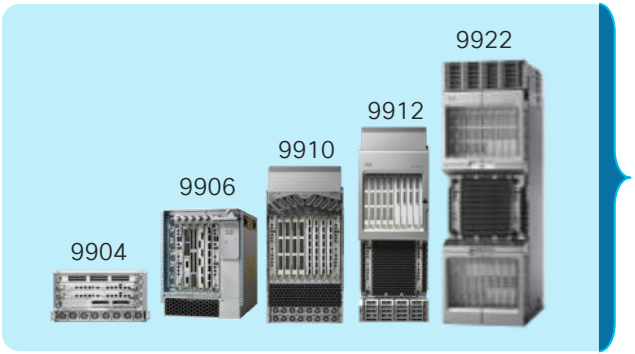

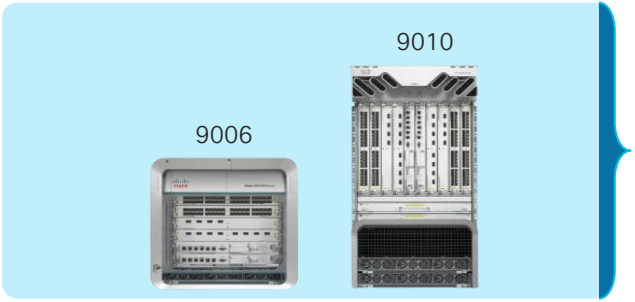

32 x 100GE TR line card

16 x 100GE TR line card

16 x 100GE SE line card

16 x 100GE TR line card

ASR 9000 5th Generation Line cards

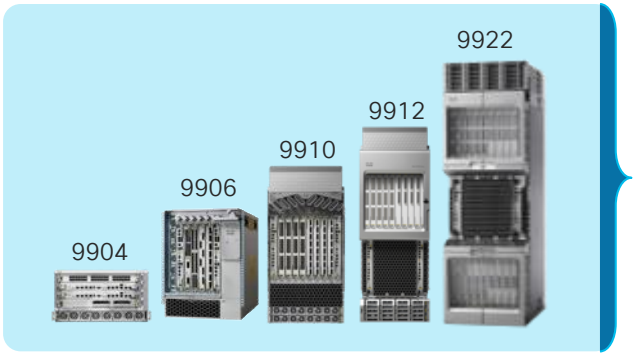

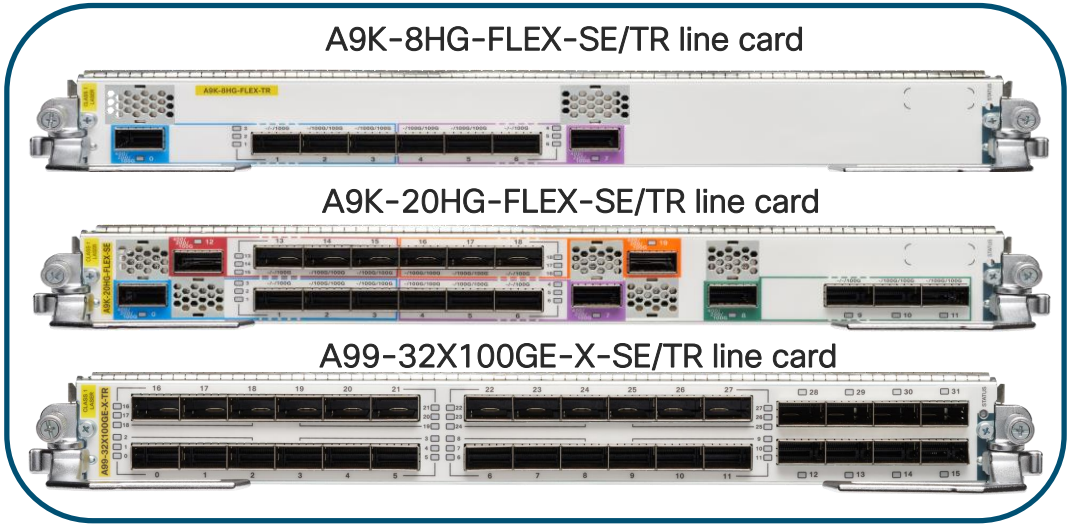

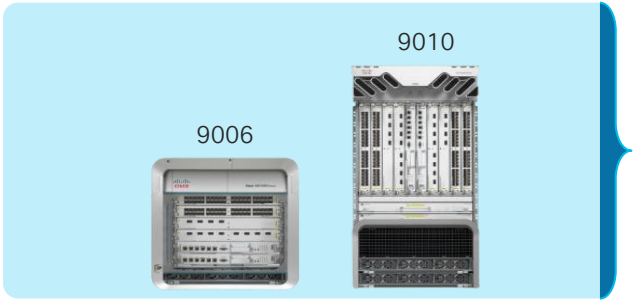

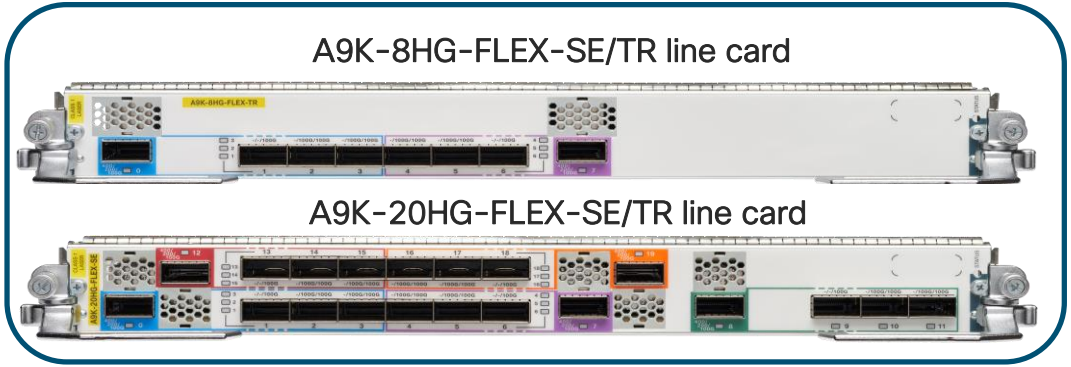

ASR 9000 5th Generation 3.2T Card

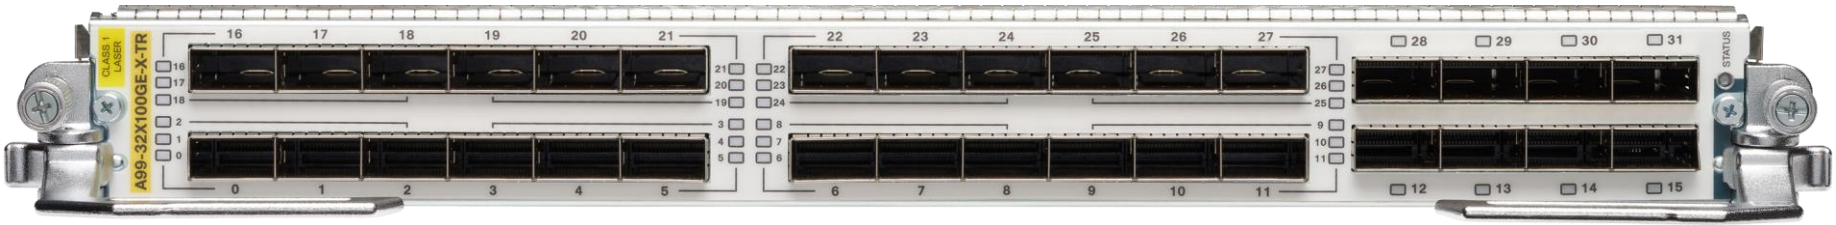

A99-32X100GE-X-SE/TR line card

- 32 Ports 100G support
- Full Feature Parity with Tomahawk Except for MACsec
- IPv6 Forwarding Optimized

ASR 9000 5th Generation 2T Combo Card

400G/200G/100G
Multi-rate Ports:
0/7/8/12/19

5 Slices:
1x 400G/200G/100G
multi-rate port
+ 3x100G ports per Slice

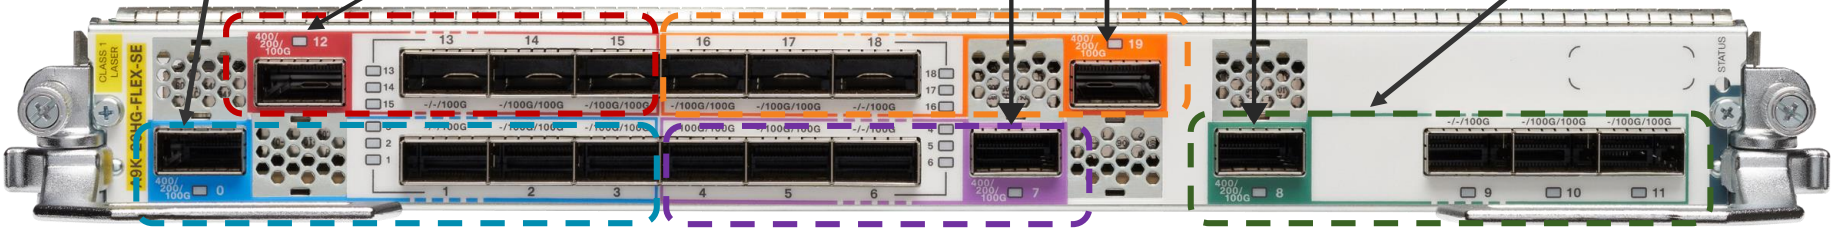

A9K-20HG-FLEX-SE/TR line card

400G Ready
- 10G/25G/40G/100G/200G/400G Support

Each Slice Independently Configured as:
- 1x400G
- 1x200G + 2x100G
- 4x100G

Each 100G Breakout into 4x25G or 4x10G

ASR 9000 5th Generation 800G Combo Card

400G/200G/100G
Multi-rate Ports:
0/7

2 Slices:
1x 400G/200G/100G
multi-rate port
+ 3x100G ports per Slice

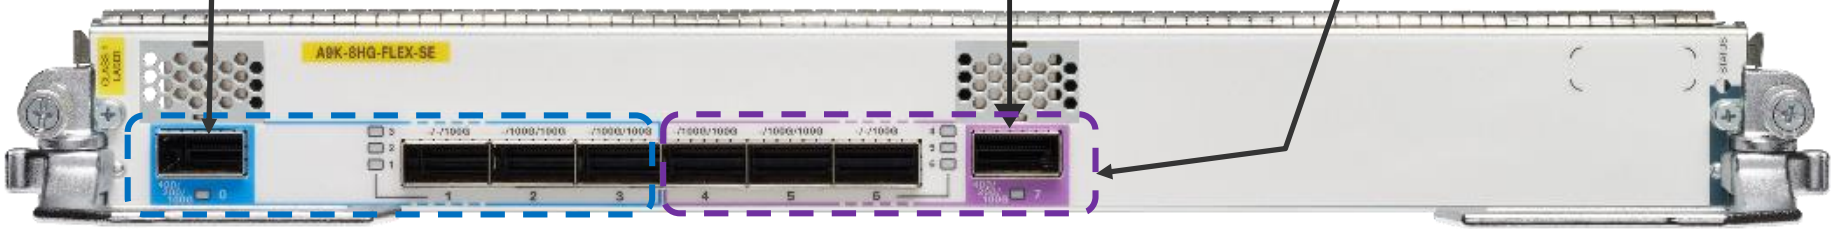

A9K-8HG-FLEX-SE/TR line card

400G Ready
- 10G/25G/40G/100G/200G/400G Support

Each Slice Independently Configured as:
- 1x400G
- 1x200G + 2x100G
- 4x100G

Each 100G Breakout into 4x25G or 4x10G

Pop Quiz ????

What different interface rates can ASR9000 LightSpeed Plus 2T Combo Linecard Support?

- 400G &100G
- 400G, 200G & 100G
- 400G, 100G & 40G
- 400G, 100G, 25G & 10G
- 400G, 200G, 100G, 40G, 25G & 10G

ASR 9000 Hardware Architecture

ASR 9000 System Architecture "At-a-Glance"

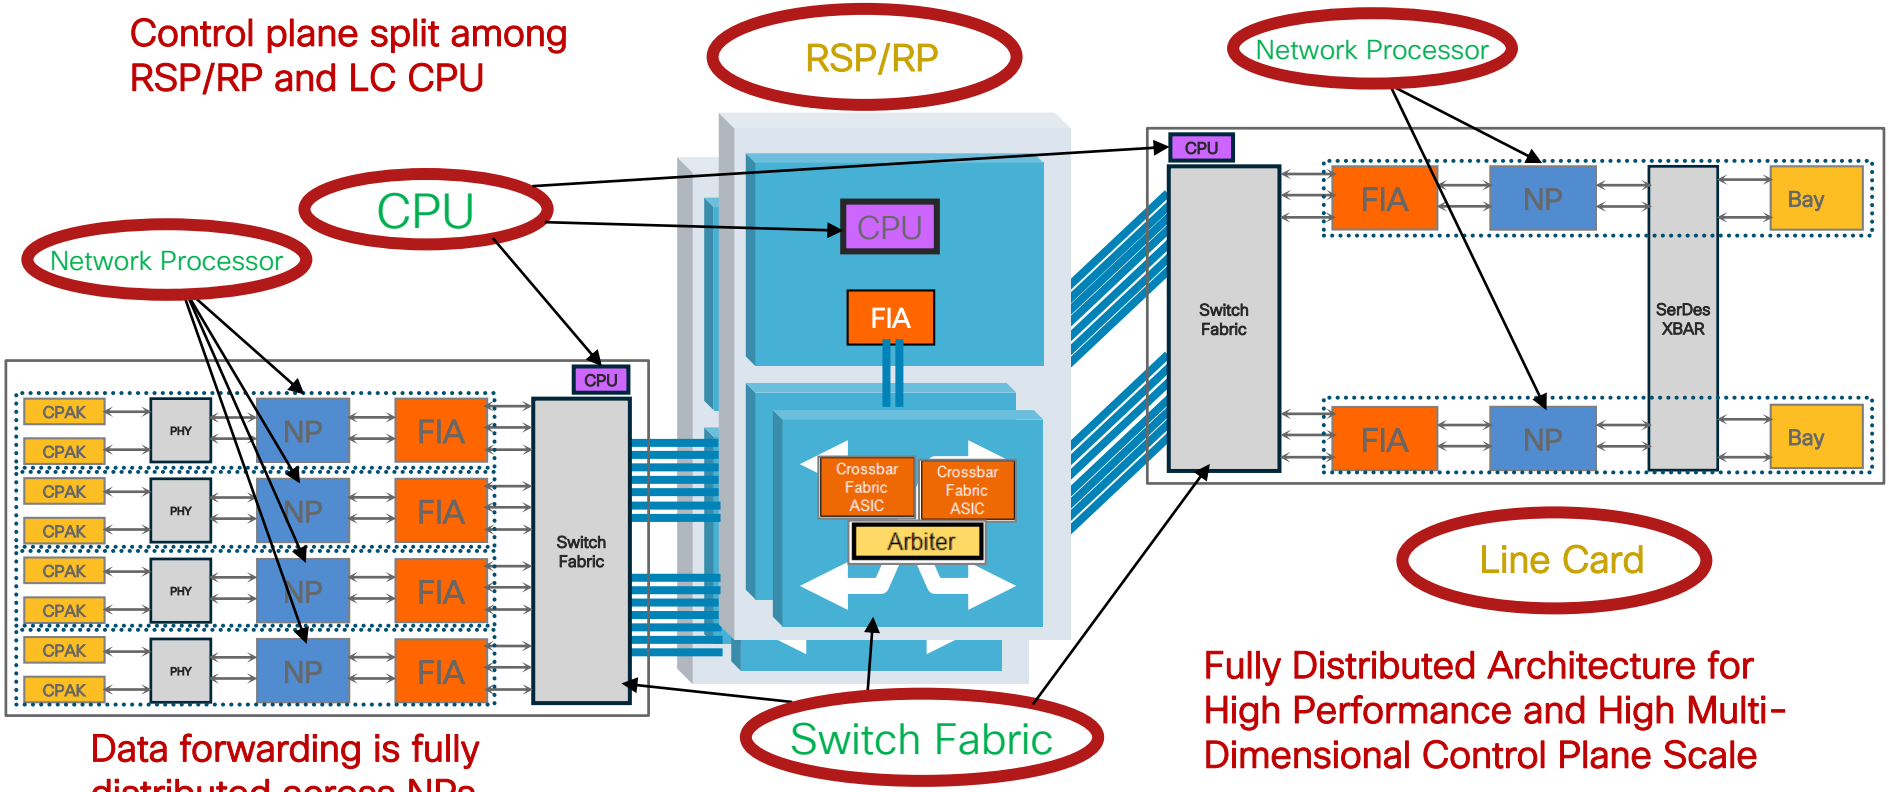

A99-32X100GE-X-SE/TR (7-fabric) LC Architecture (in 9922, 9912, 9910 & 9906)

600G/fab x 7 fab
= 4.2T

8x Slices

QSFP28 0	Retimer Retimer	PHY	LightSpeedPlus NP0 FIA
QSFP28 4	Retimer Retimer	PHY	LightSpeedPlus NP1 FIA
QSFP28 8	Retimer Retimer	PHY	LightSpeedPlus NP2 FIA
QSFP28 12	Retimer Retimer	PHY	LightSpeedPlus NP3 FIA
QSFP28 16	Retimer Retimer	PHY	LightSpeedPlus NP4 FIA
QSFP28 20	Retimer Retimer	PHY	LightSpeedPlus NP5 FIA
QSFP28 24	Retimer Retimer	PHY	LightSpeedPlus NP6 FIA
QSFP28 28	Retimer Retimer	PHY	LightSpeedPlus NP7 FIA

CPU

Switch Fabric 0 (SKB)

Switch Fabric 1 (SKB)

x8

A99-32X100GE-X-SE/TR LC Architecture (in 9904)

1.8T/RSP x 2RSP
= 3.6T

8x Slices

QSFP28 0 — Retimer Retimer — PHY LightSpeedPlus NP0 FIA
QSFP28 4 — Retimer Retimer — PHY LightSpeedPlus NP1 FIA
QSFP28 8 — Retimer Retimer — PHY LightSpeedPlus NP2 FIA
QSFP28 12 — Retimer Retimer — PHY LightSpeedPlus NP3 FIA
QSFP28 16 — Retimer Retimer — PHY LightSpeedPlus NP4 FIA
QSFP28 20 — Retimer Retimer — PHY LightSpeedPlus NP5 FIA
QSFP28 24 — Retimer Retimer — PHY LightSpeedPlus NP6 FIA
QSFP28 28 — Retimer Retimer — PHY LightSpeedPlus NP7 FIA

CPU

Switch Fabric 0 (SKB) — x8 x8 x8 x8 x8 x8 — RSP0

Switch Fabric 1 (SKB) — x8 x8 x8 x8 x8 x8 — RSP1

A9K-20HG-FLEX-SE/TR (5-fabric) LC Architecture (in 9922, 9912, 9910 & 9906)

5x Slices

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP0 FIA

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP1 FIA

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP2 FIA

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP3 FIA

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP4 FIA

CPU

Switch Fabric 0 (SKB)

600G/fab x 5 fab
= 3.0T (actual)

2x8
2x8
2x8
2x8
2x8

A9K-20HG-FLEX-SE/TR (5-fabric) LC Architecture (in 9010 & 9006)

5x Slices

QSFP-DD
QSFP28
QSFP28
QSFP28

PHY

LightSpeedPlus NP0 FIA

LightSpeedPlus NP1 FIA

LightSpeedPlus NP2 FIA

LightSpeedPlus NP3 FIA

LightSpeedPlus NP4 FIA

CPU

1.1T/RSP x 2RSP
= 2.2T (actual)

Switch Fabric 0 (SKB)

x8
x8
x8
x8
x8
x8
x8
x8

RSP0

RSP1

A9K-20HG-FLEX-SE/TR (5-fabric) LC Architecture (in 9904)

5x Slices

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP0 FIA

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP1 FIA

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP2 FIA

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP3 FIA

QSFP-DD
QSFP28
PHY
LightSpeedPlus NP4 FIA

CPU

Switch Fabric 0 (SKB)

1.4T/RSP x 2RSP
= 2.8T (actual)

x8
x8
x8
x8
x8

x8
x8
x8
x8
x8

RSP0

RSP1

A9K-8HG-FLEX-SE/TR (5-fabric) LC Architecture (in 9922, 9912, 9910 & 9906)

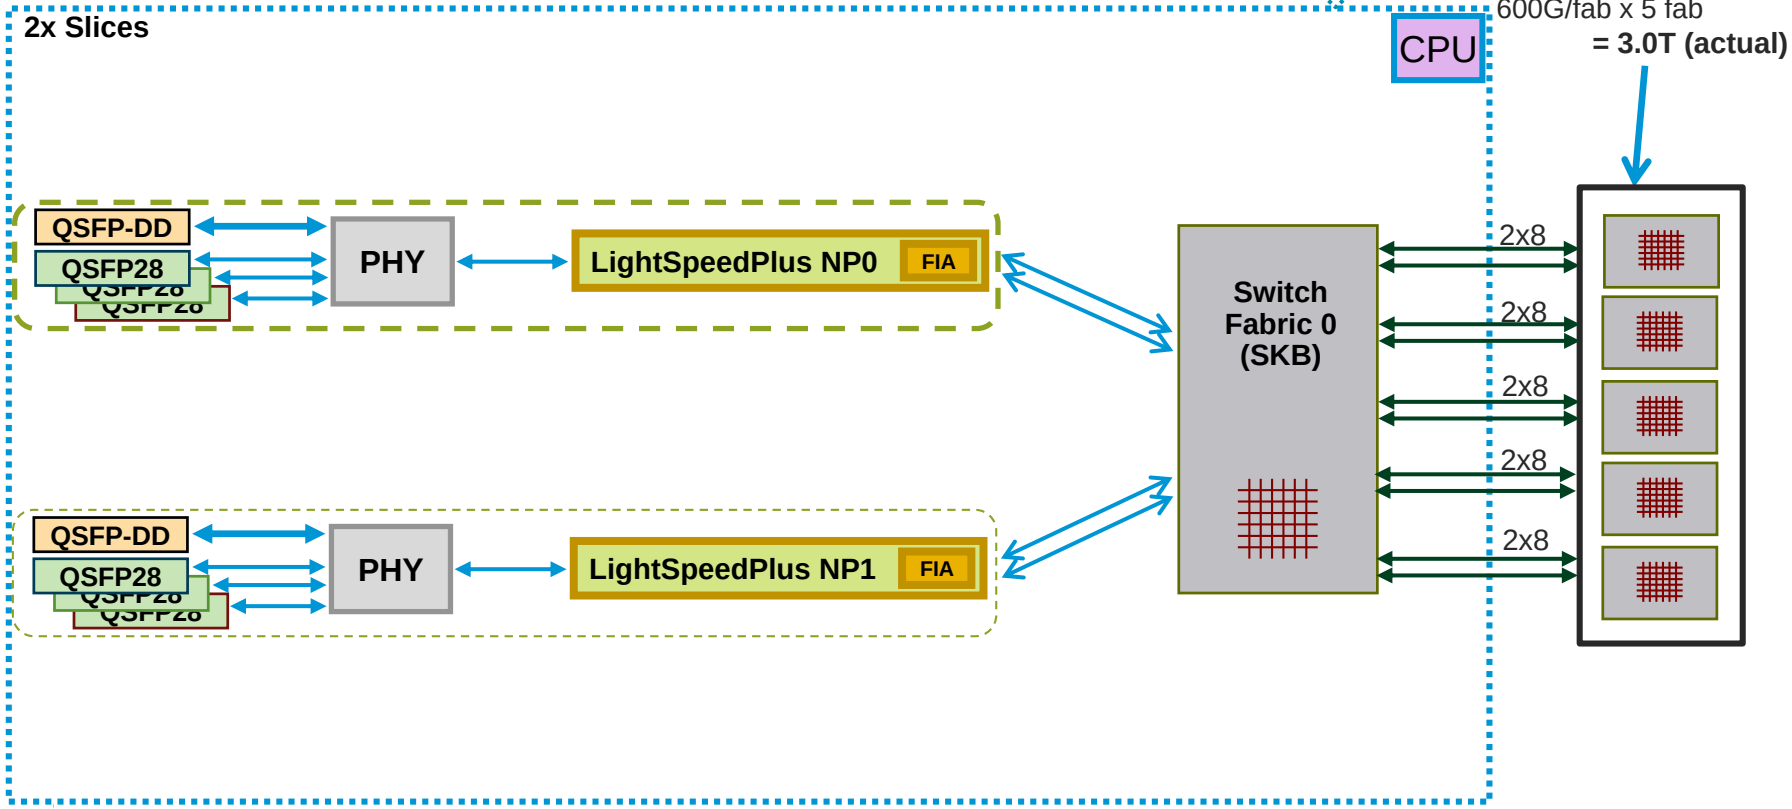

A9K-8HG-FLEX-SE/TR (5-fabric) LC Architecture (in 9010 & 9006)

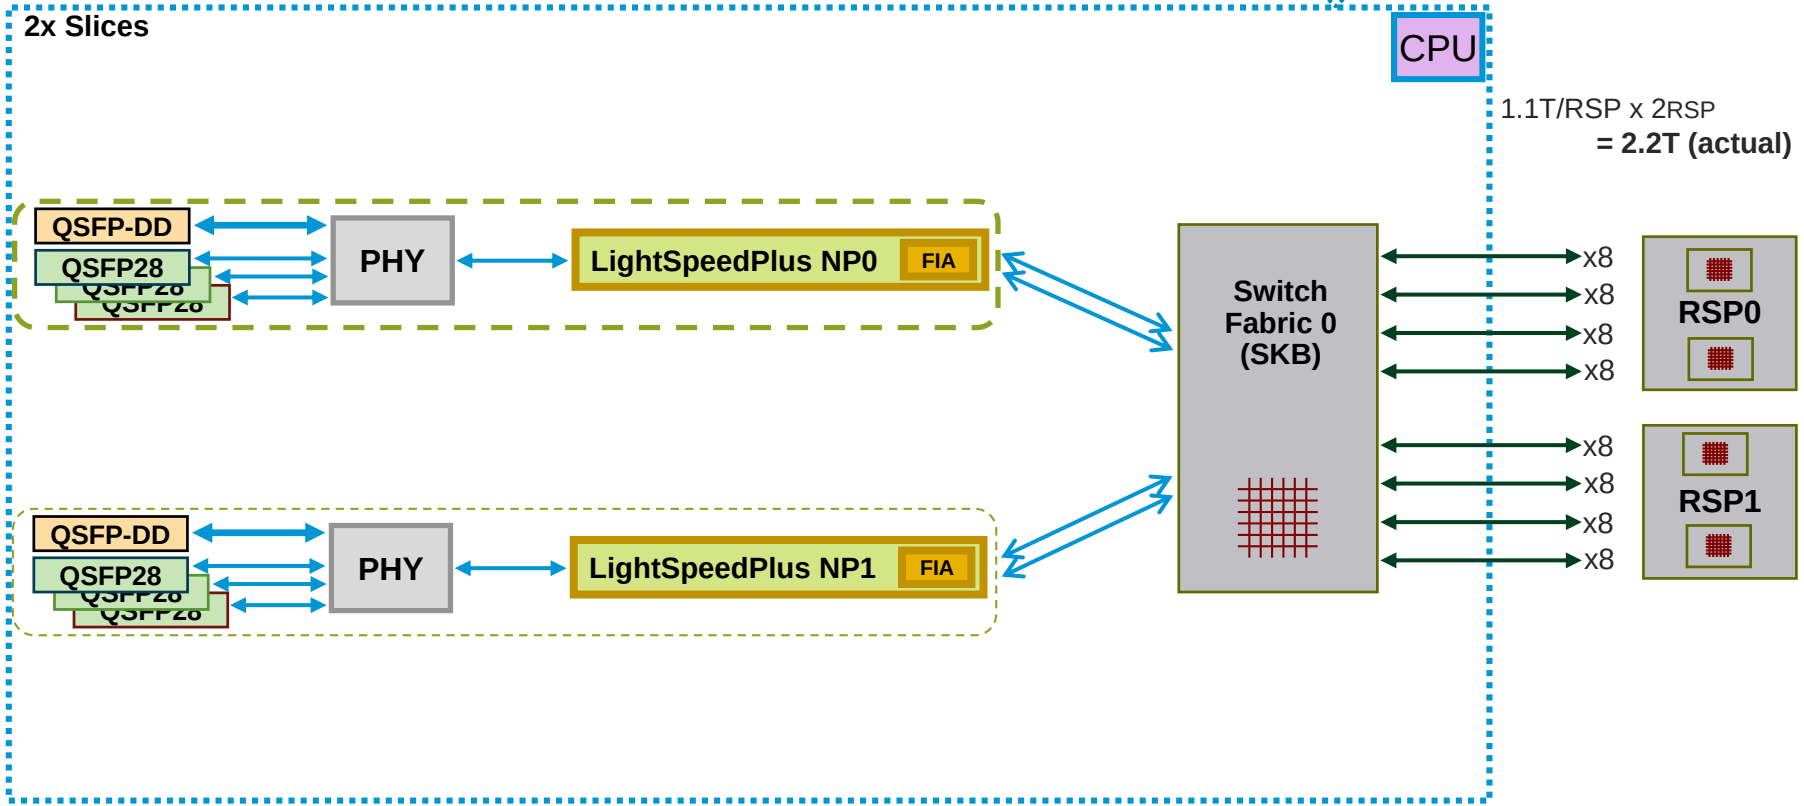

A9K-8HG-FLEX-SE/TR (5-fabric) LC Architecture (in 9904)

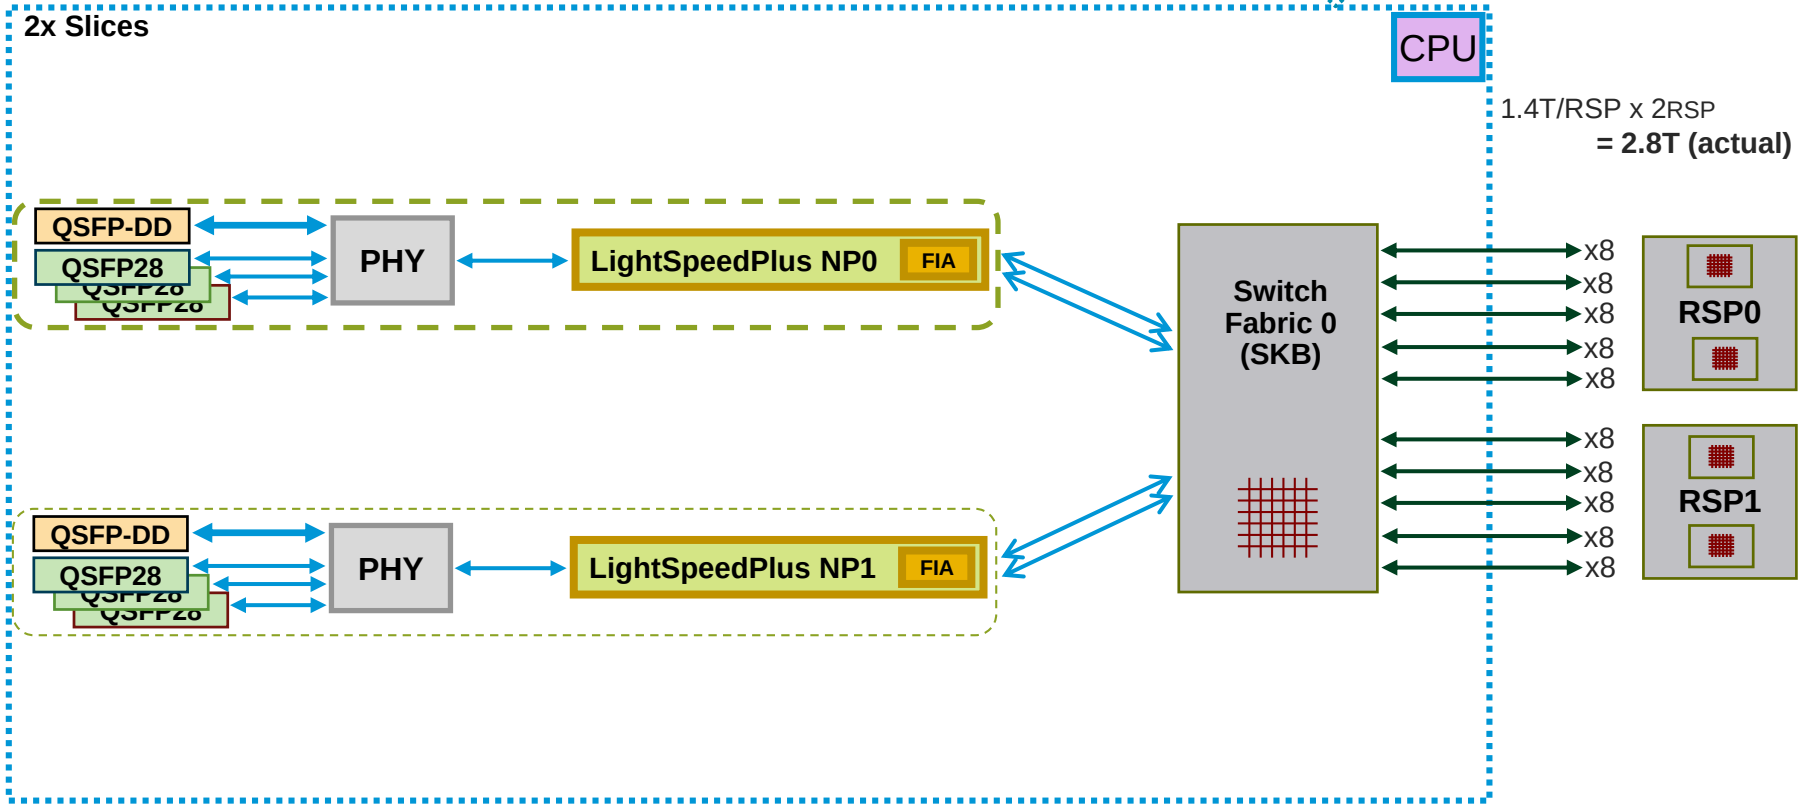

ASR90xx – RSP5 and Mixed LC

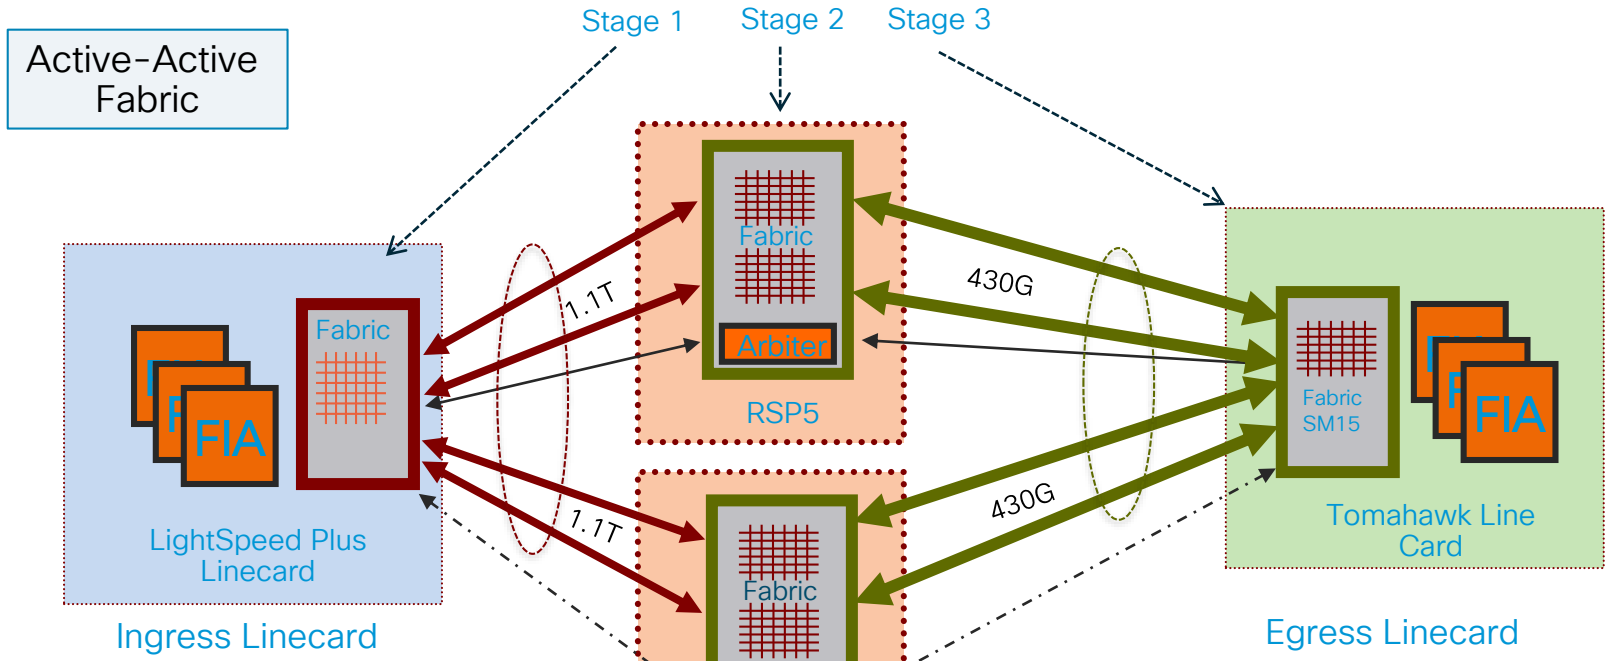

LightSpeed Plus LC Fabric bandwidth:

dual RSP: 2.2T/slot

single RSP: 1.1T/slot

Tomahawk LC Fabric bandwidth:

Dual RSP: 860Gbps/slot

Single RSP: 430Gbps/slot

ASR9922/12 System – LightSpeed Plus & Tomahawk

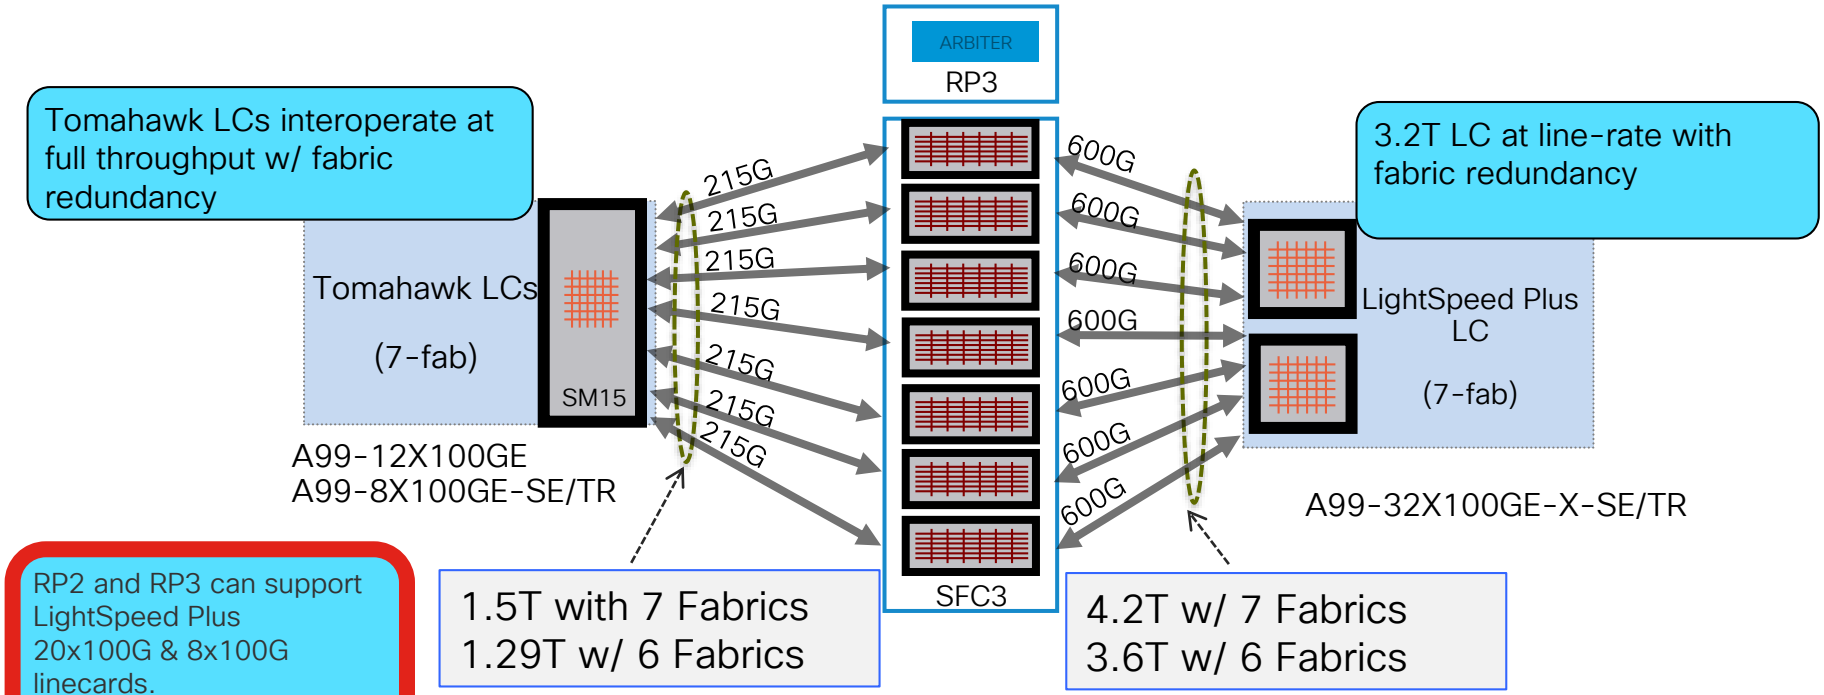

ASR9910/06 System – LightSpeed Plus & Tomahawk

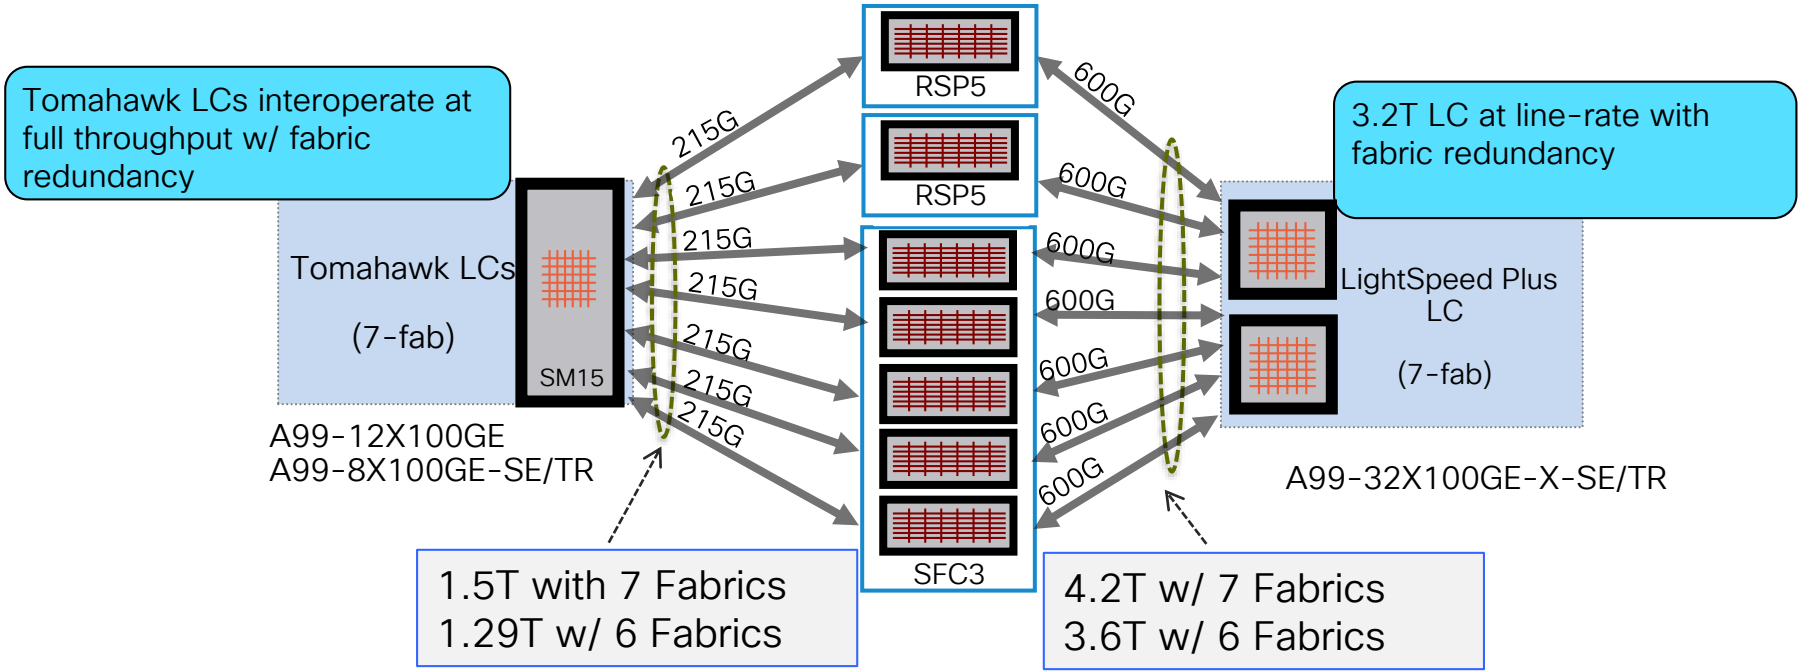

ASR99xx Mix Mode

All Tomahawk LCs interoperate at full throughput w/ fabric redundancy; except for 12x100GE LC, which will be at 1.075T

1.075T with 5 fabrics
860G w/ 4 fabrics

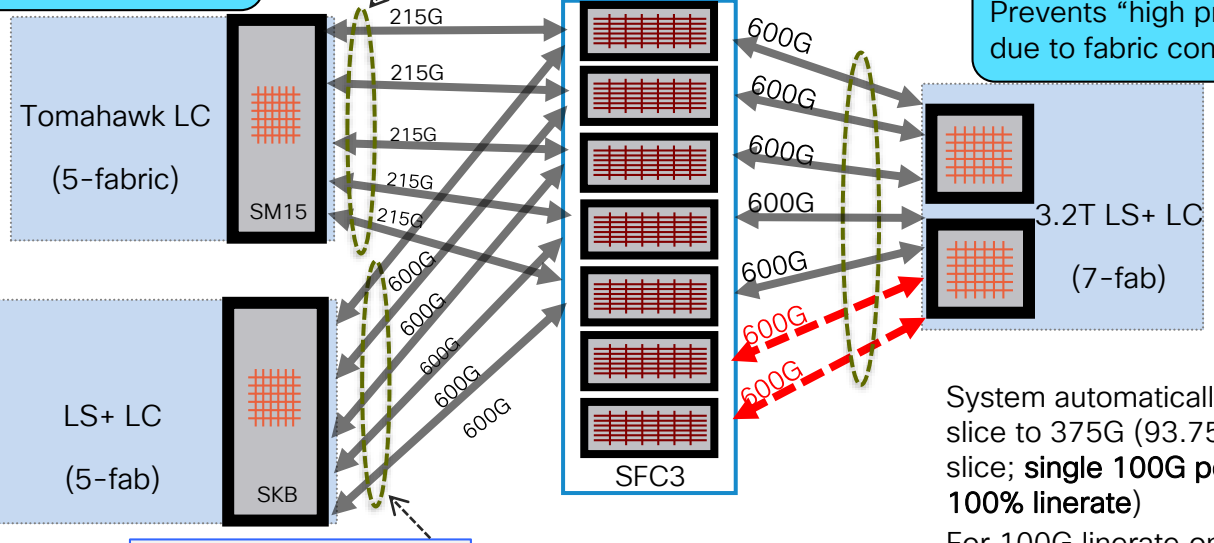

3T with 5-fabrics
Prevents “high priority” packet drop due to fabric congestion

2.0T LC at line-rate w/ fabric redundancy
800G LC at line-rate w/ fabric redundancy

3T w/ 5 fabrics
2.4T w/ 4 fabrics

System automatically rate-limits each slice to 375G (93.75% linerate per slice; **single 100G port can reach 100% linerate**)

For 100G linerate on all ports:

- Shut down one slice

cisco Live!

Egress FIA Rate-Limiter

- Triggered by Switch Fabric Degradation: Not enough switch fabric bandwidth to support linerate
- Rate-Limiter is applied in the direction NP -> FIA

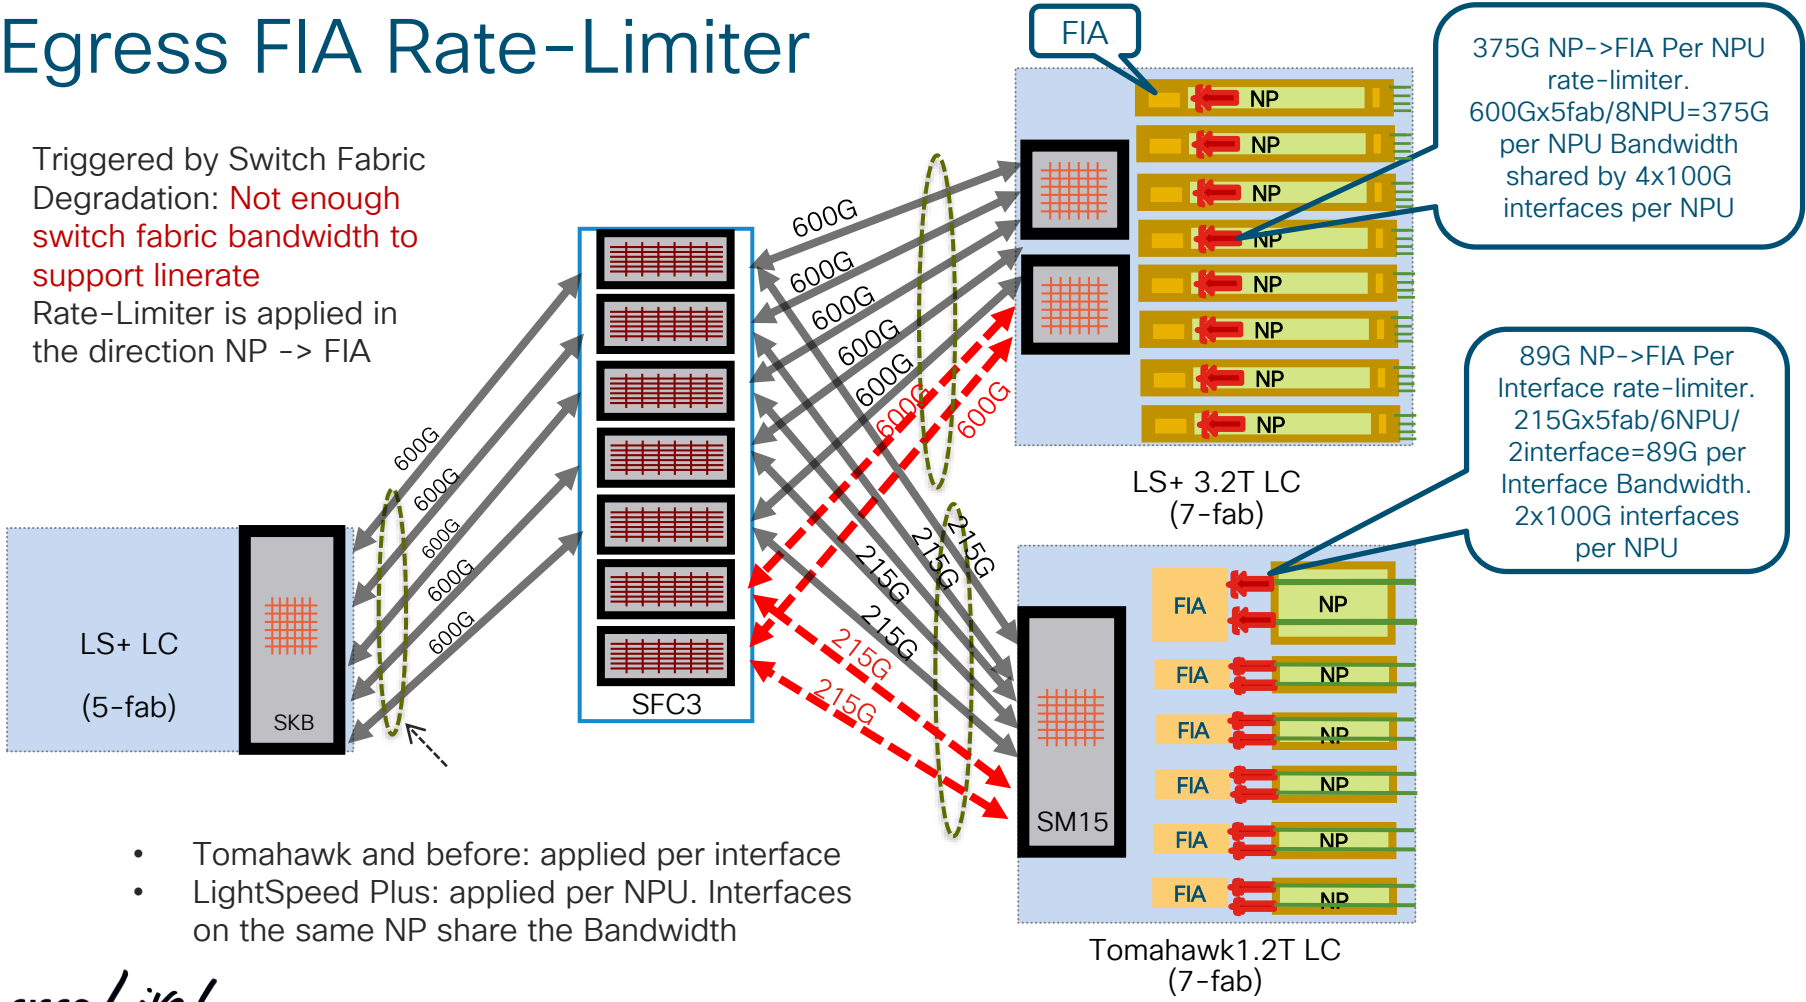

- Tomahawk and before: applied per interface
- LightSpeed Plus: applied per NPU. Interfaces on the same NP share the Bandwidth

How do you know Rate-Limiter Triggered?

Console Message:

```
LC/0/3/CPU0: pfm_node_lc[261]: %FABRIC-FIA-1-RATE_LIMITER_ON : Set|fialc[4795]|0x108a000|Insufficient fabric capacity for card types in use - FIA egress rate limiter applied
LC/0/5/CPU0:pfm_node_lc[207]: %FABRIC-FIA-1-RATE_LIMITER_ON : Set|fialc[4798]|0x108a000|Insufficient fabric capacity for card types in use - FIA egress rate limiter applied
```

Command to check if Rate-limiter has been applied

"show pfm location all"

```
RP/0/RSP0/CPU0:PE-3#show pfm loc all
Thu Aug 15 16:29:00.169 PHT

node: node0_0_CPU0
----------------------
CURRENT TIME: Aug 15 16:29:00 2019
PFM TOTAL: 1   EMERGENCY/ALERT(E/A): 1   CRITICAL(CR): 0   ERROR(ER): 0
------------------------------------------------------------------------------------------
Raised Time         |S#|Fault Name                        |Sev|Proc_ID|Dev/Path Name |Handle
--------------------+--+----------------------------------+---+-------+--------------+----------
Jul 31 16:24:34 2019|2 |RATE_LIMITER_ON                   |E/A|4812   |Fabric Interfa|0x108a000
```

Slice Level Management

- CLI to shut down slice
  - hw-module location ? slice ? power-savings
- Benefits:
  - Ports on remaining slices are not affected
  - Power Saving
  - Switch Fabric Bandwidth Re-allocation
    - Example: 32x100G LC mix with 5-fab cards
      - Total switch fabric bandwidth = 5x600G =3T
      - Rate-limiter triggered: 3000G/8 = 375G per NPU
      - One Slice Shut: 3000G/7 > 400G per NPU, Linerate Guaranteed
      - Two Slice Shut: (3000G-600G)/6 = 400G per NPU. N+1 Switch Fabric Redundancy

Line Card Components

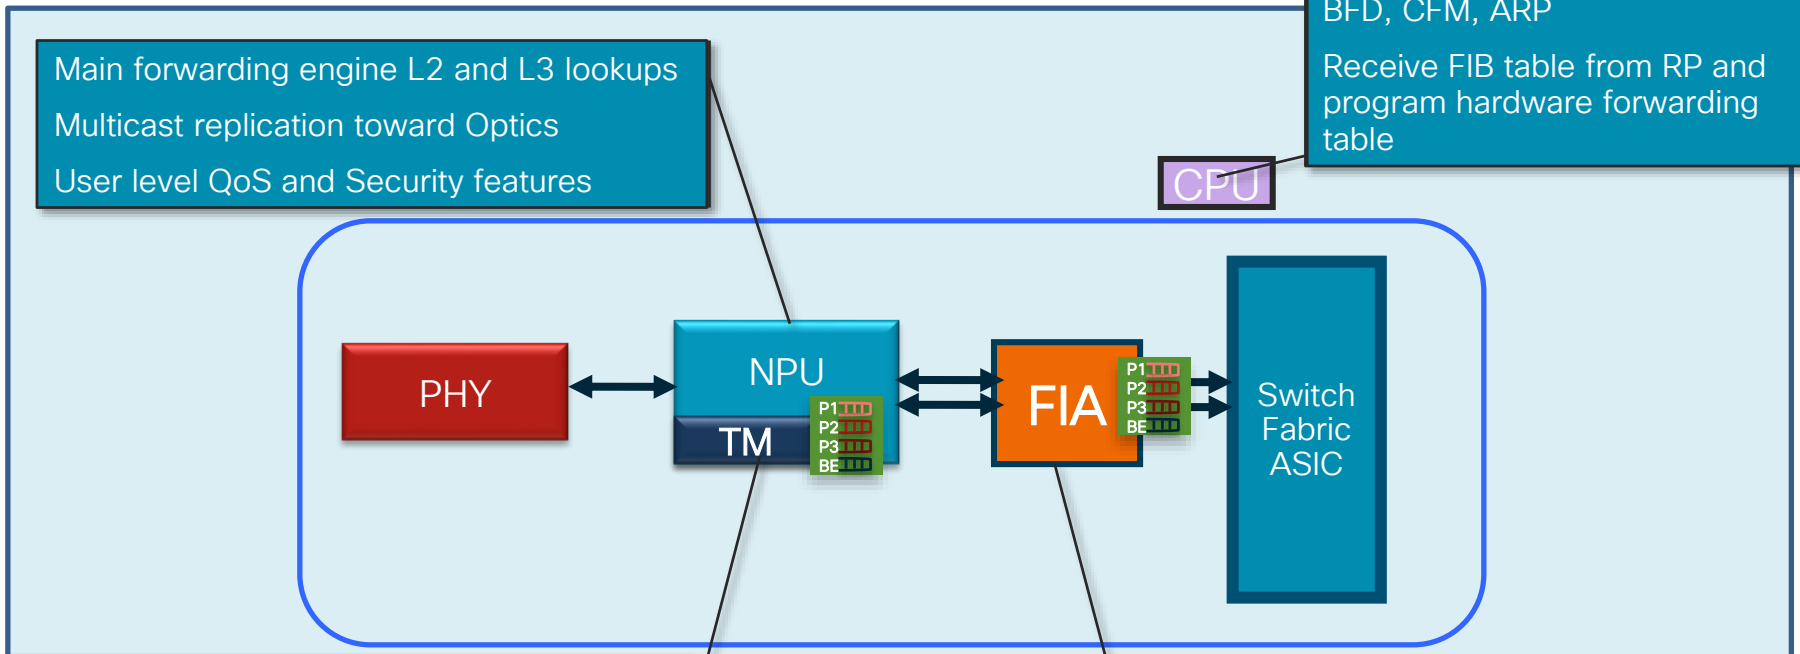

Network Processor Architecture Details

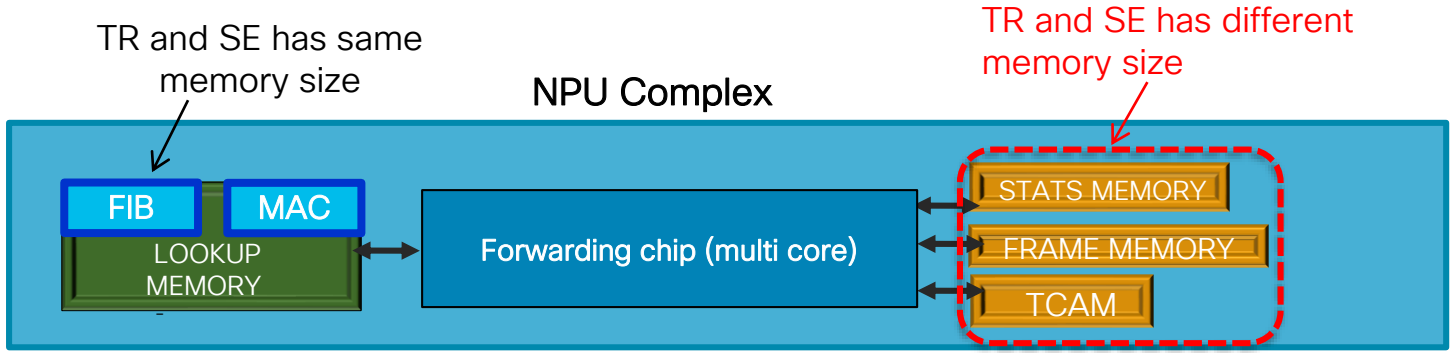

- TCAM: VLAN tag, QoS and ACL classification
- Stats memory: interface statistics, forwarding statistics etc
- Frame memory: buffer, Queues
- Lookup Memory: forwarding tables, FIB, MAC, ADJ
- TR/SE
  - Different TCAM/frame/stats memory size for different per-LC QoS, ACL, logical interface scale
  - Same lookup memory for same system wide scale mixing different variation of LCs doesn't impact system wide scale

-TR: transport optimized, -SE: Service edge optimized

Pop Quiz ????

Can LightSpeed Plus A99-32X100GE-X-SE/TR be used in ASR-9010 or ASR9006 Chassis?

- Yes
- No

ASR 9000 Distributed Control Plane

IOS-XR 6.X: A New Software Infrastructure

- 64-bit OpenEmbedded Linux support.
  - Processes containerization.
  - Standard Linux toolchain.
  - Third-Party applications.
- NCS5500, NCS5000 and NCS1002 support 64-bit Linux.
- ASR 9000 supports 64-bit Linux starting with 6.1.1.
  - Will still have 32-bit QNX support.

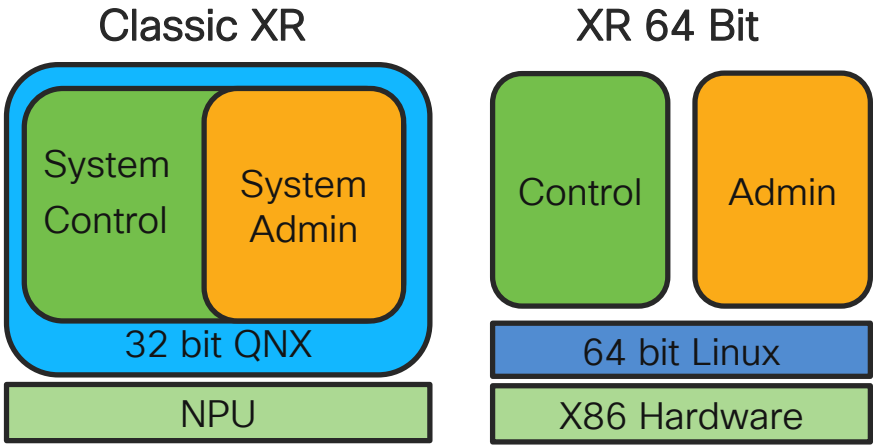

IOS XR 32-Bit and 64-Bit Differences

- IOS XR Exists in two flavors
  - 32-bit in XR12k, CRS, ASR9000
    - QNX-based
    - No Virtualization
    - No ISSU
  - 64-bit in ASR9000, NCS 5500, NCS 5000, NCS 1000 and NCS 6000
    - Linux based
    - Larger addressable memory
    - Separation Networking OS and Admin Plane
    - Virtualization: VM or Container
    - ASR9000 Running with VMs

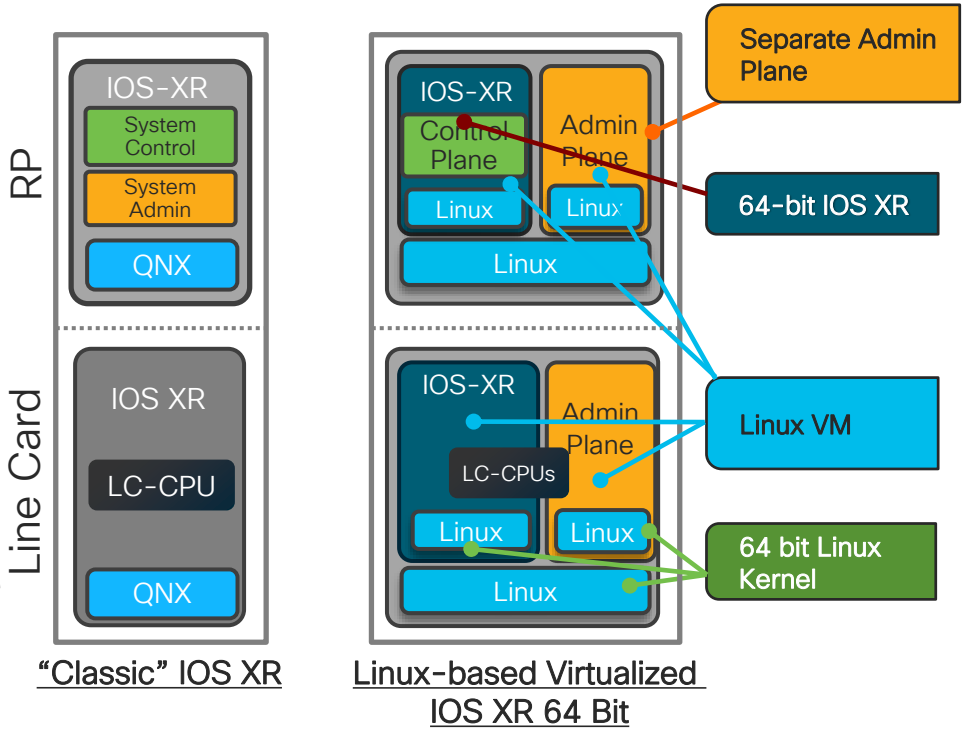

ASR9000 Fully Distributed Control Plane

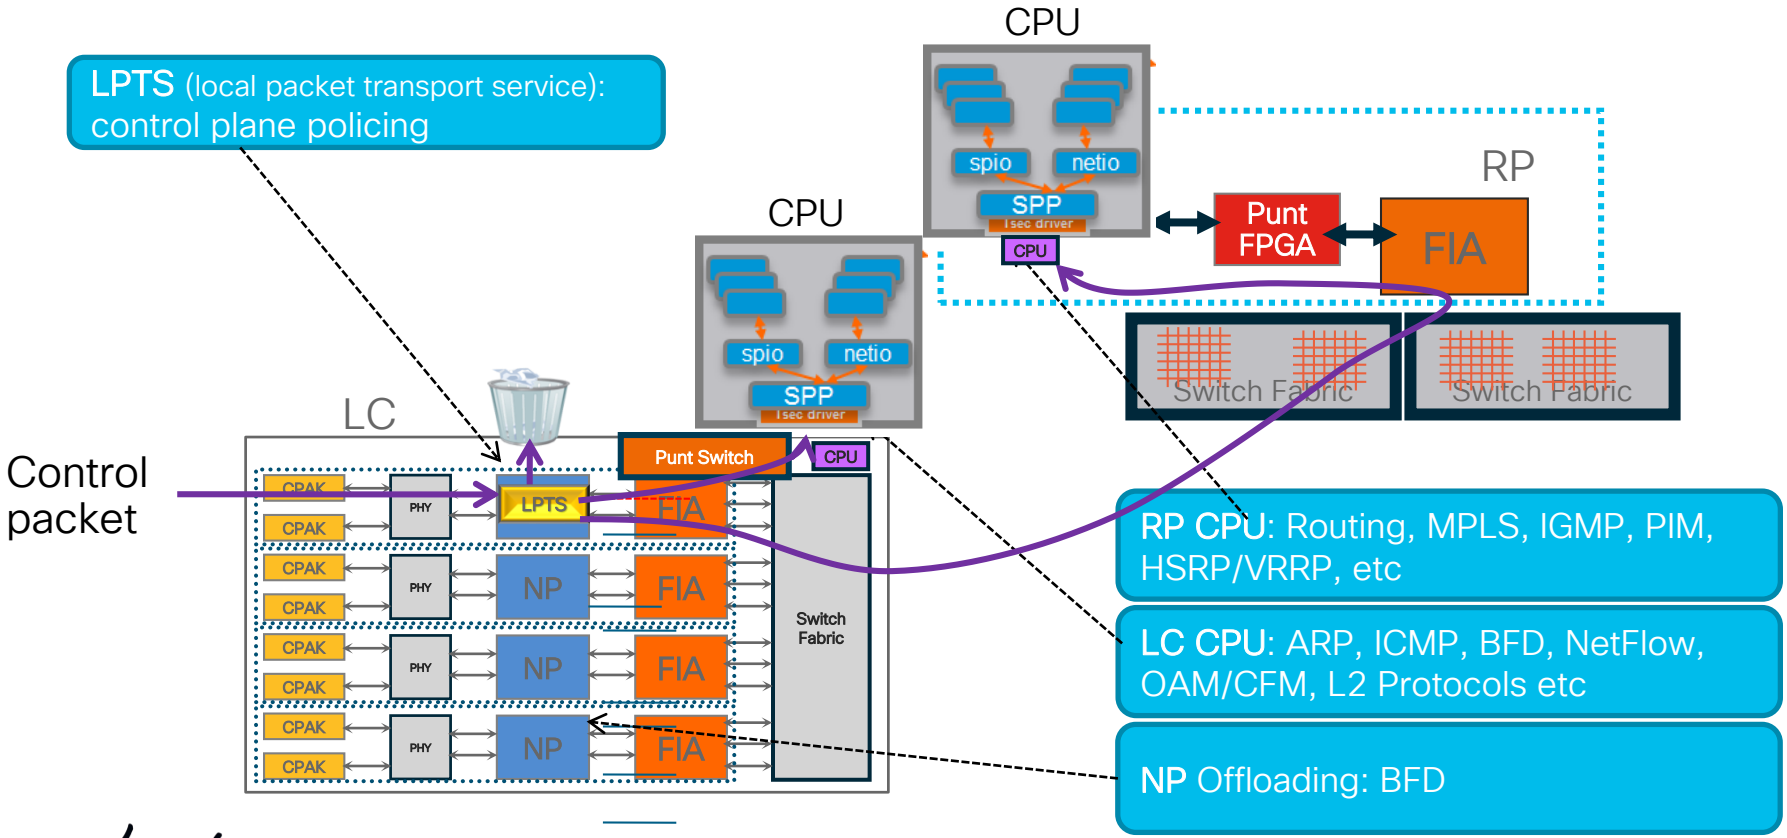

L3 Control Plane Architecture

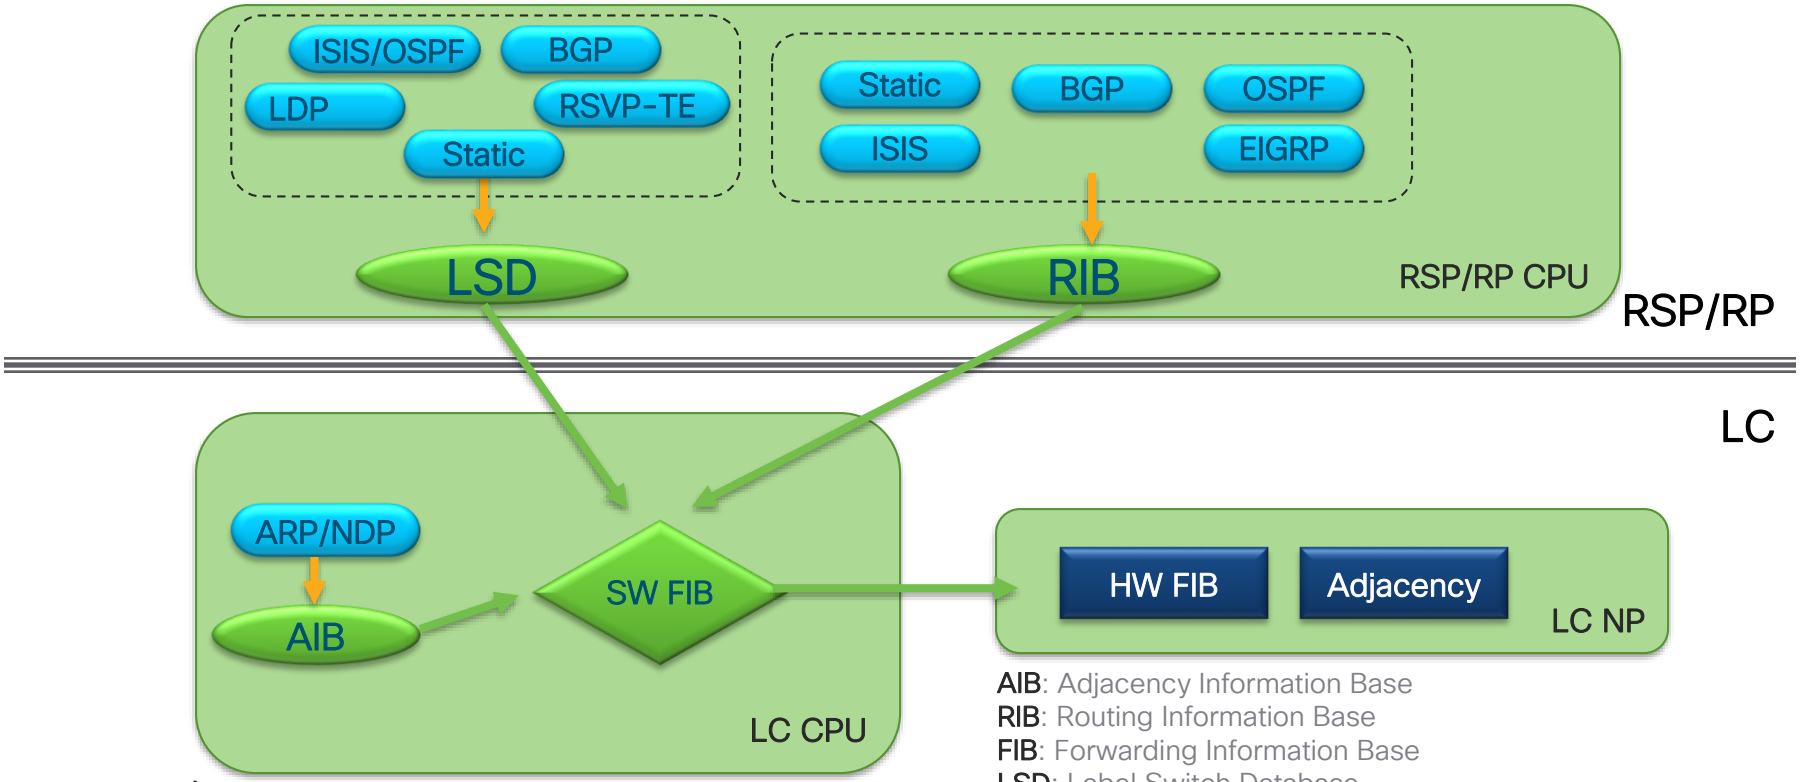

AIB: Adjacency Information Base
RIB: Routing Information Base
FIB: Forwarding Information Base
LSD: Label Switch Database

Distributed ARP Processing

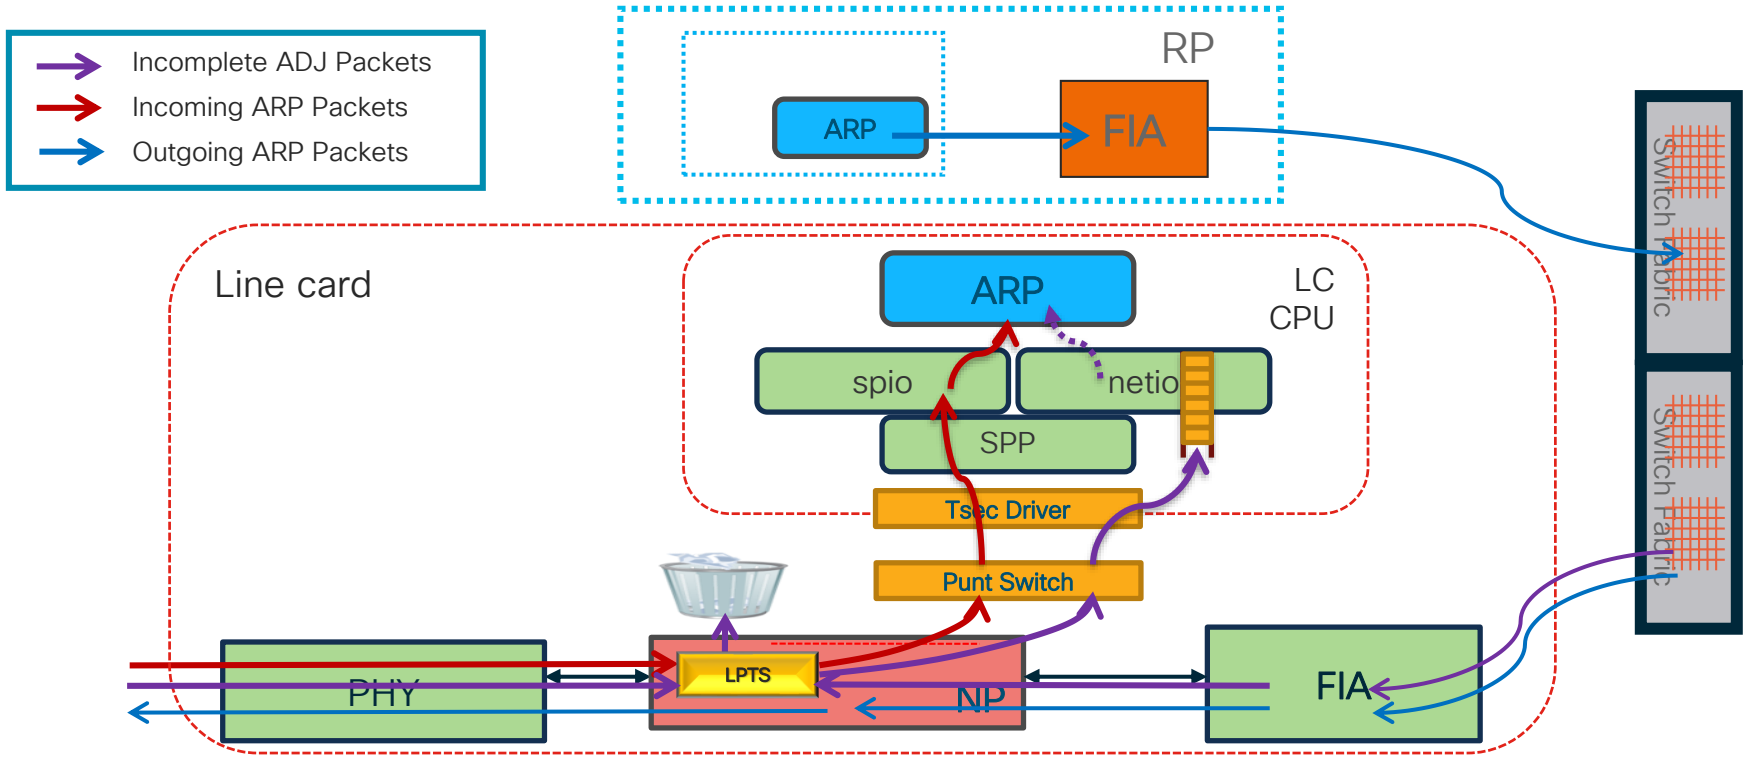

MAC Learning and Sync

Hardware based MAC learning: ~4Mpps/NP

1. NP learn MAC address in hardware (around 4M pps)

2. NP flood MAC notification (data plane) message to all other NPs in the system to sync up the MAC address system-wide. MAC notification and MAC sync are all done in hardware

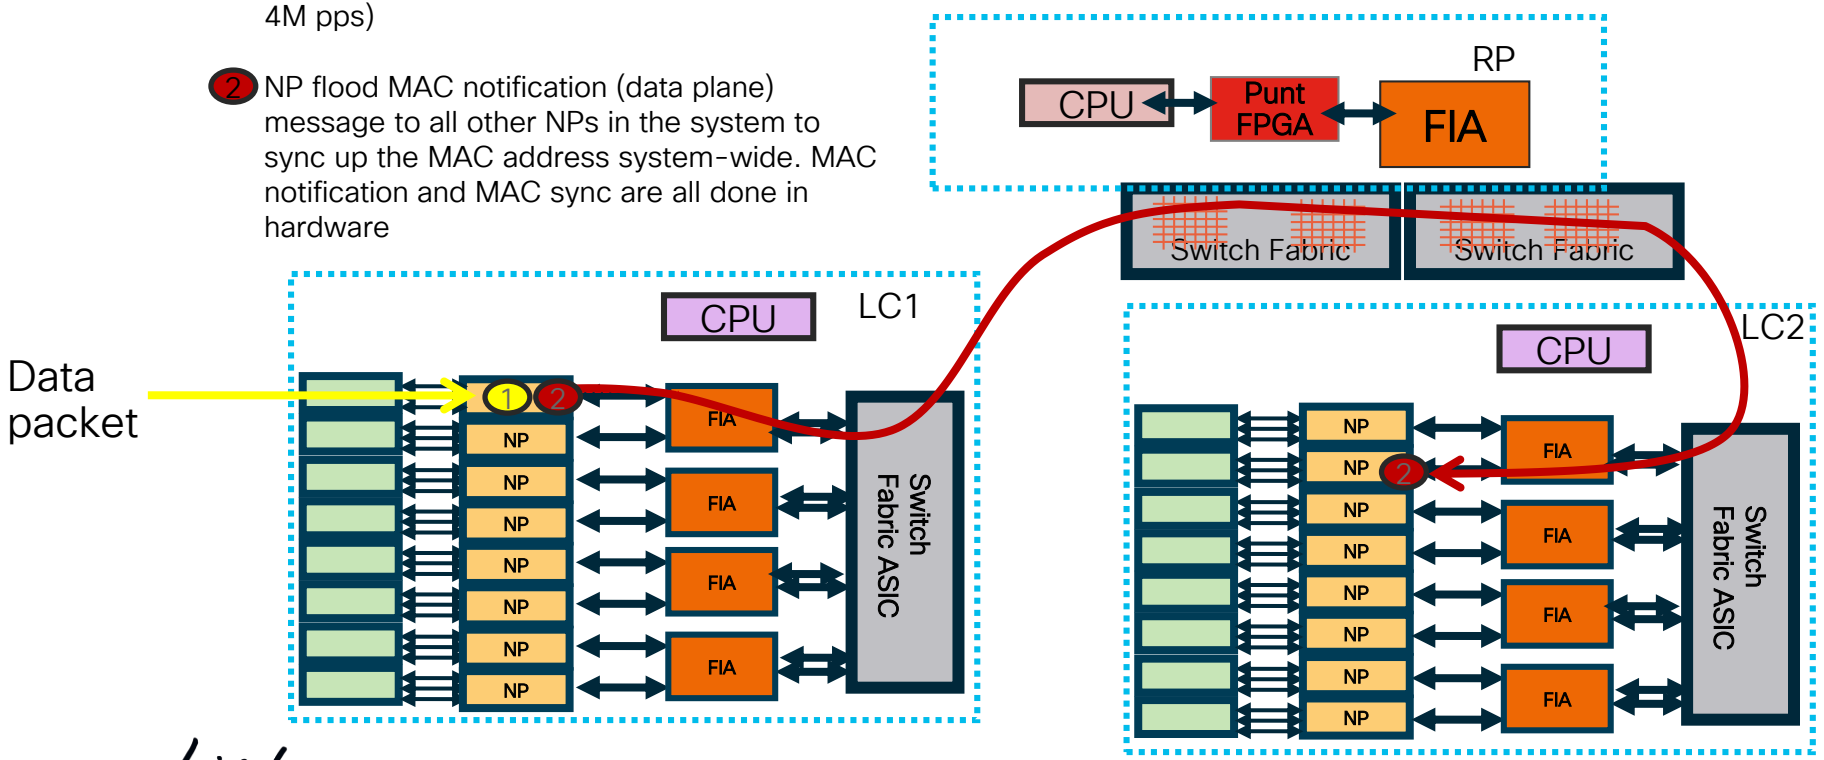

Distributed Netflow Architecture

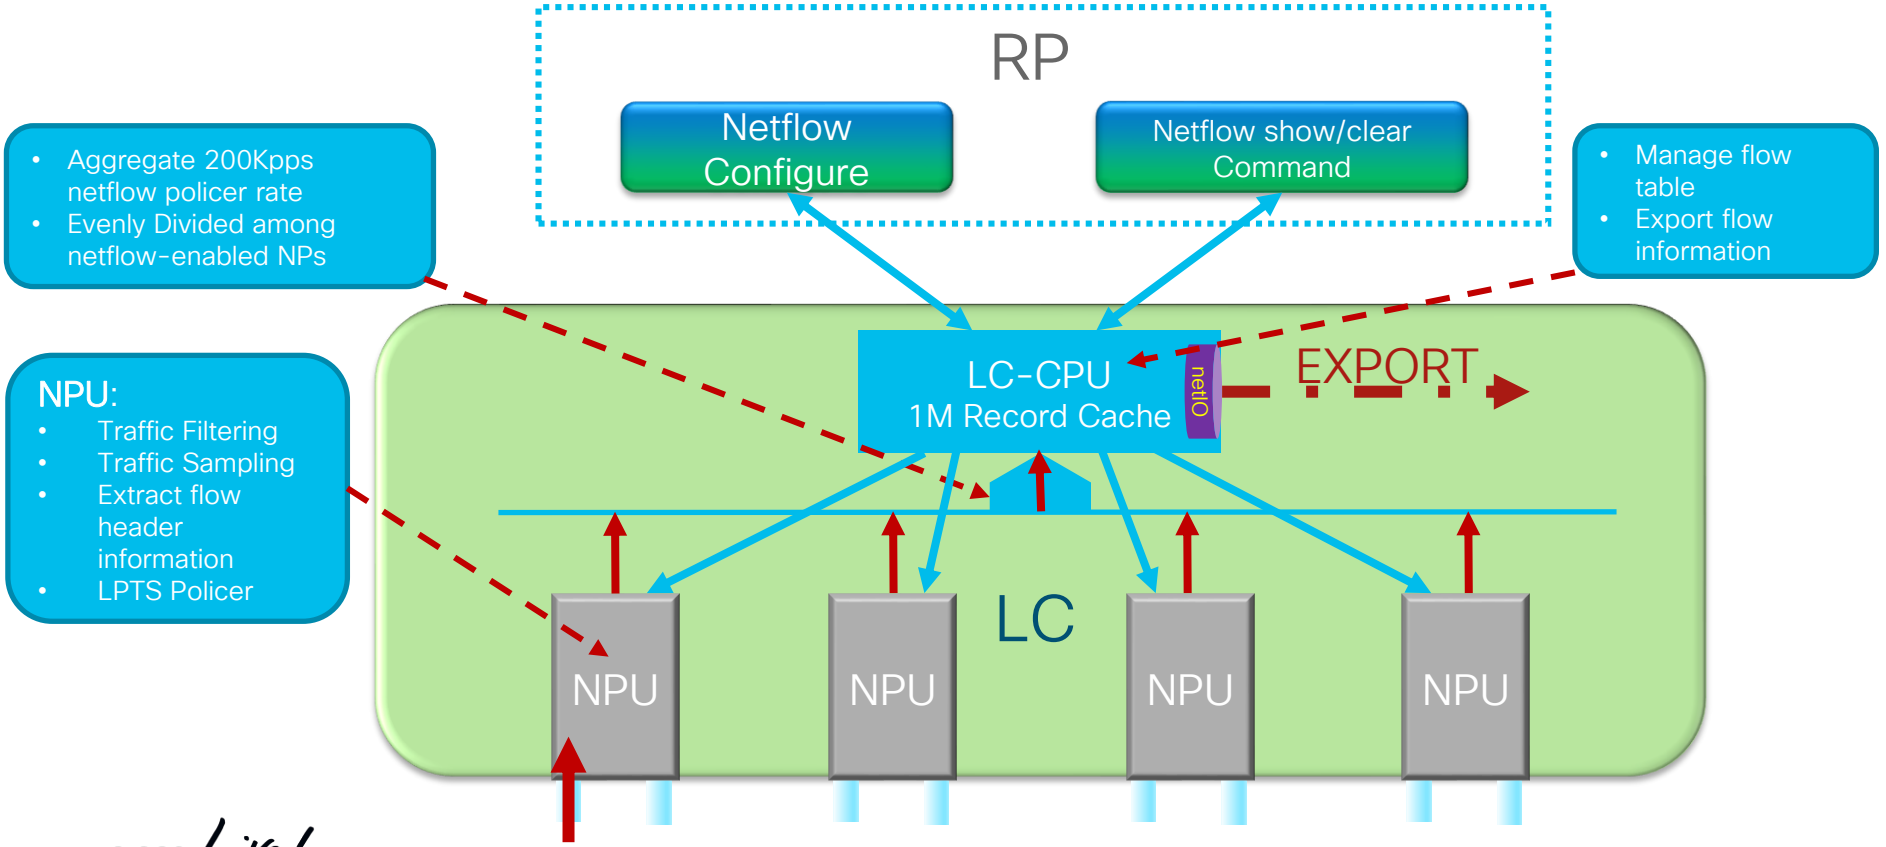

Pop Quiz ????

When **A9K-8HG-FLEX-SE/TR** is used in the following chassis with RSP5, which one(s) provide full fabric redundancy?

- ASR 9904 (1.4T/RSPx2=2.8T)
- ASR 9006 (1.1T/RSPx2=2.2T)
- ASR 9010 (1.1T/RSPx2=2.2T)
- All of above

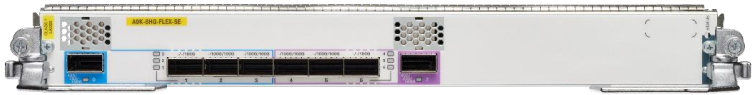

ASR 9000 Data Packet Processing

Distributed Two-Stage Packet Processing

- Ingress lookup yields packet egress port and applies ingress features
- Egress lookup performs packet-rewrite and applies egress features

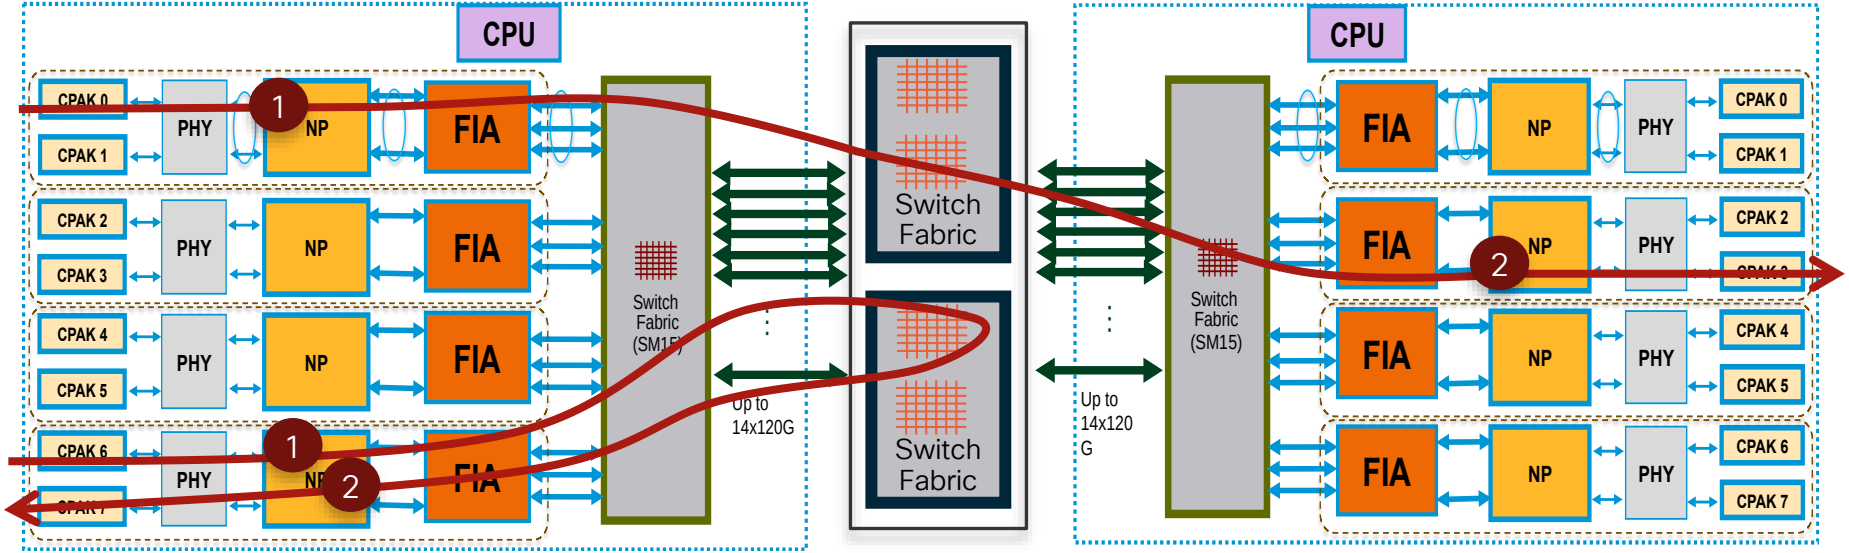

Uniform packet flow for simplicity and predictable performance

ASR9000 Life of a Packet

2
- Ingress L2/L3 FIB lookup, ACL/QoS lookup
- Ingress PBR/ABR, ACL, uRPF
- Ingress QoS: classification, marking, policing
- Packet Punting
- Ingress ECMP/LAG hashing

4
- Buffering packet from NP
- Requesting fabric credit
- Manage superframe and load-balancing packet across fabric
- Manage system VoQ

6
- Egress L2/L3 FIB lookup, ACL/QoS lookup
- Egress PBR/ABR, ACL, uRPF
- Egress QoS: classification, marking, policing, shaping
- Incomplete Adj Packet Punting
- Egress ECMP/LAG hashing

8
- MACSEC Emcryption
- G.709/OTN/WAN-PHY/LAN-PHY
- Line Clocking

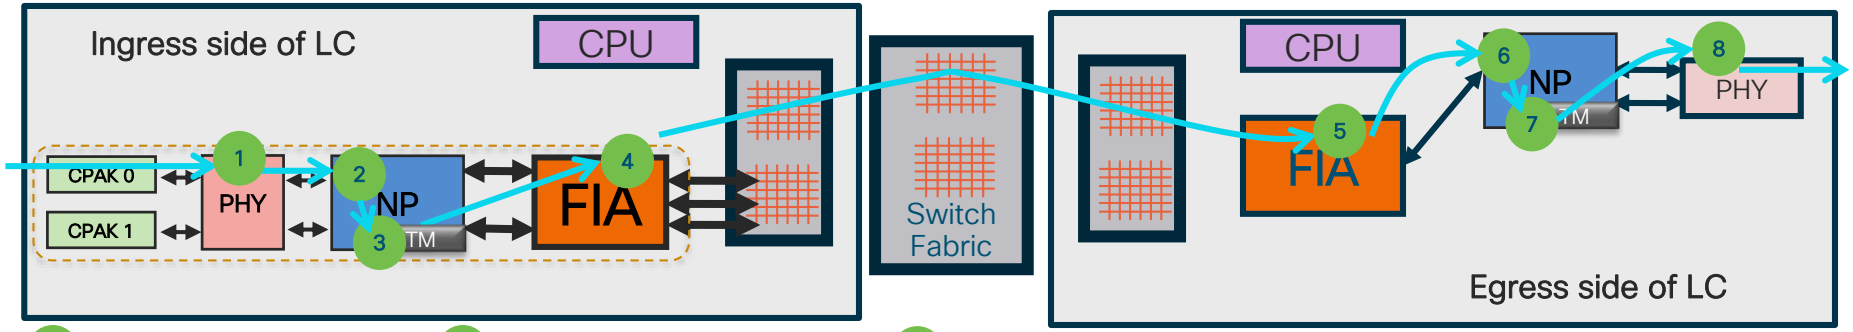

1
- MACSEC Decryption
- G.709/OTN/WAN-PHY/LAN-PHY
- Line Clocking

3
- Ingress Queuing Processing
- Bypassed in case no ingress queuing support

5
- Re-assembling packets from superframe
- Send packet to corresponding NP
- Release buffer and fabric credit

7
- Egress Queuing Processing

Switch Fabric Arbitration

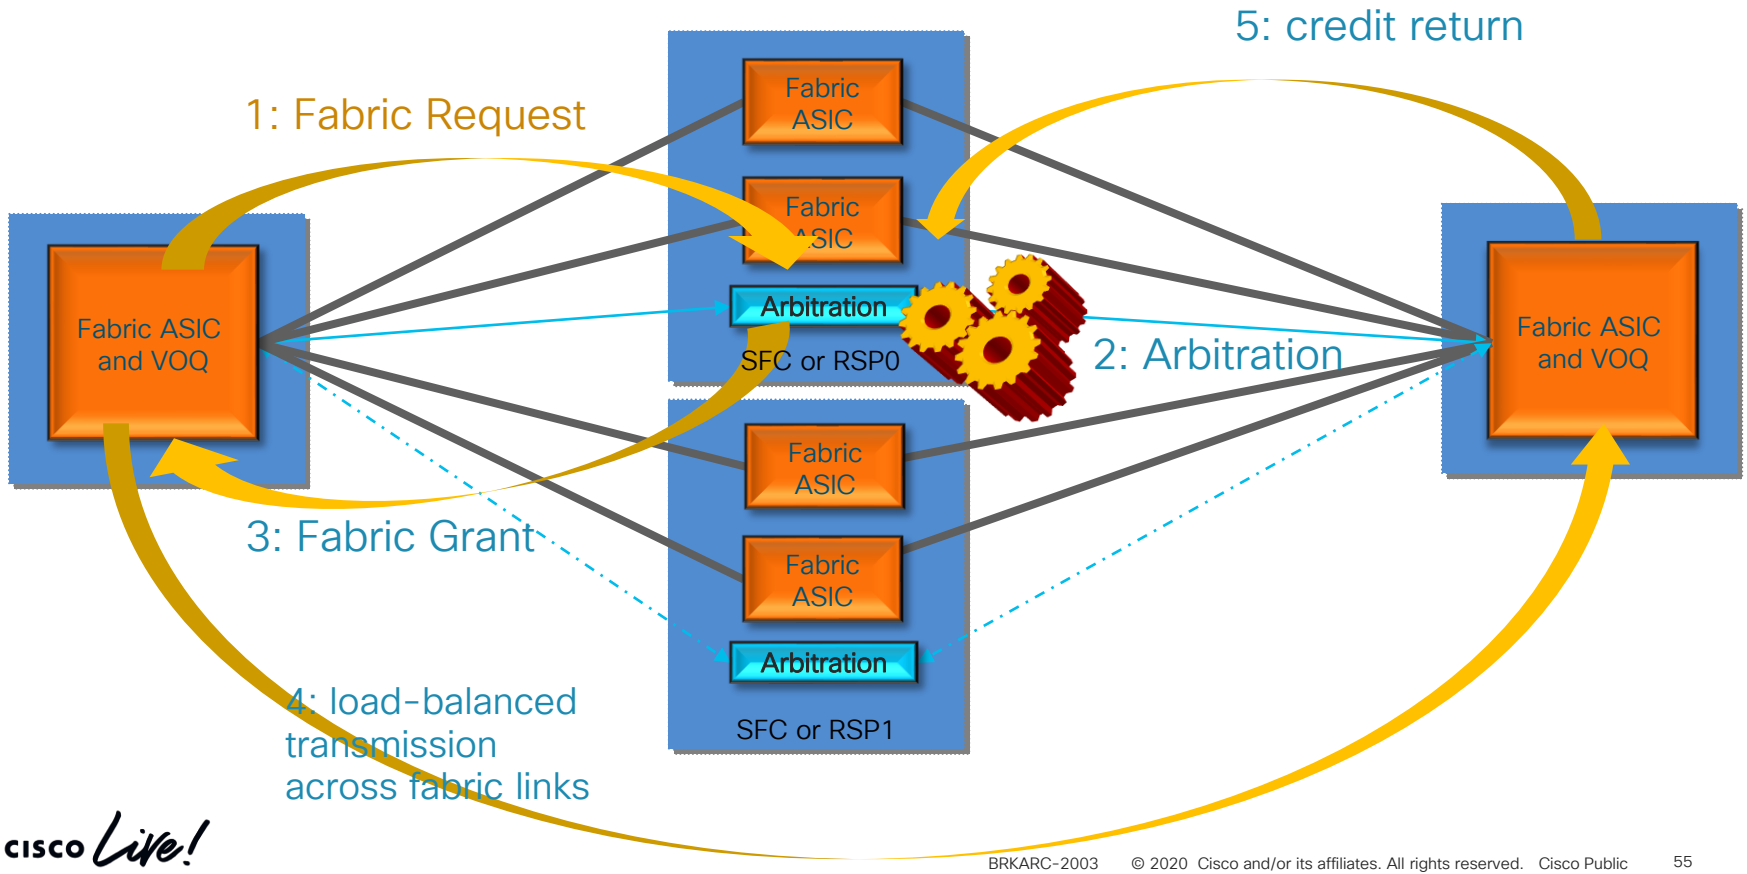

Fabric Load Balancing – Unicast

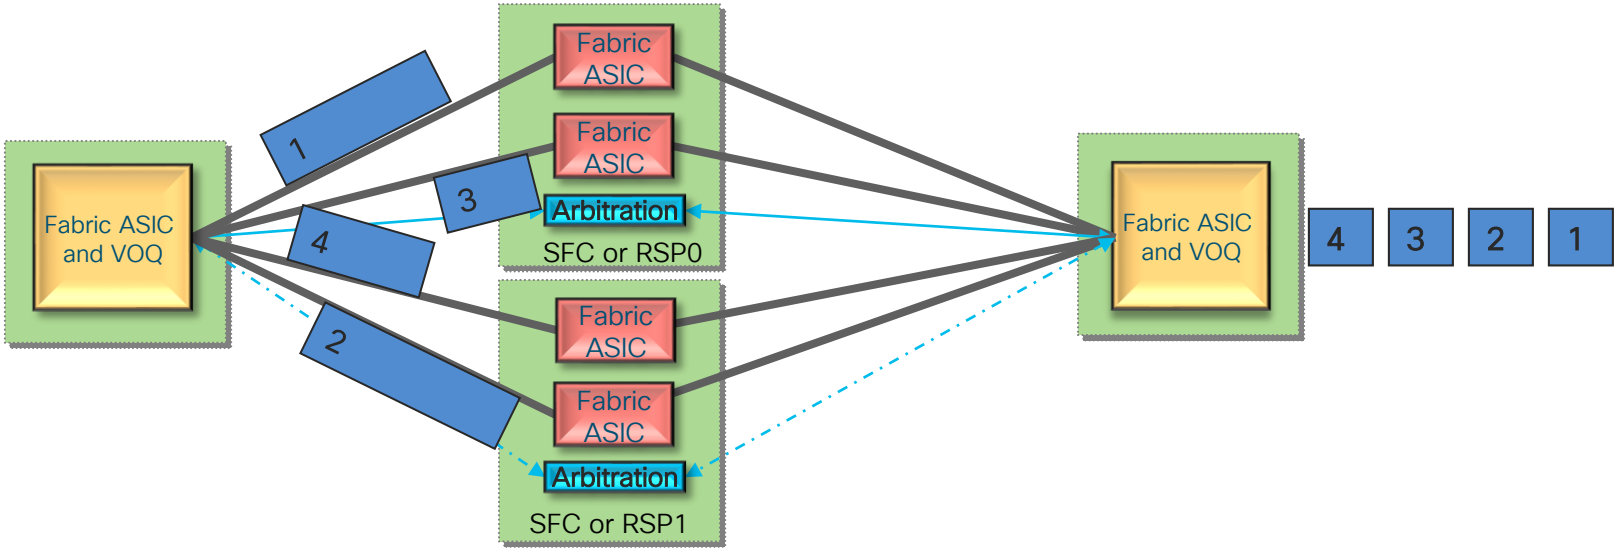

- Unicast traffic sent across first available fabric link to destination (maximizes efficiency)
- Each frame (or super frame) contains sequencing information
- All destination fabric ASIC have re-sequencing logic
- Additional re-sequencing latency is measured in nanoseconds

Distributed Multi-Stage Multicast Replication

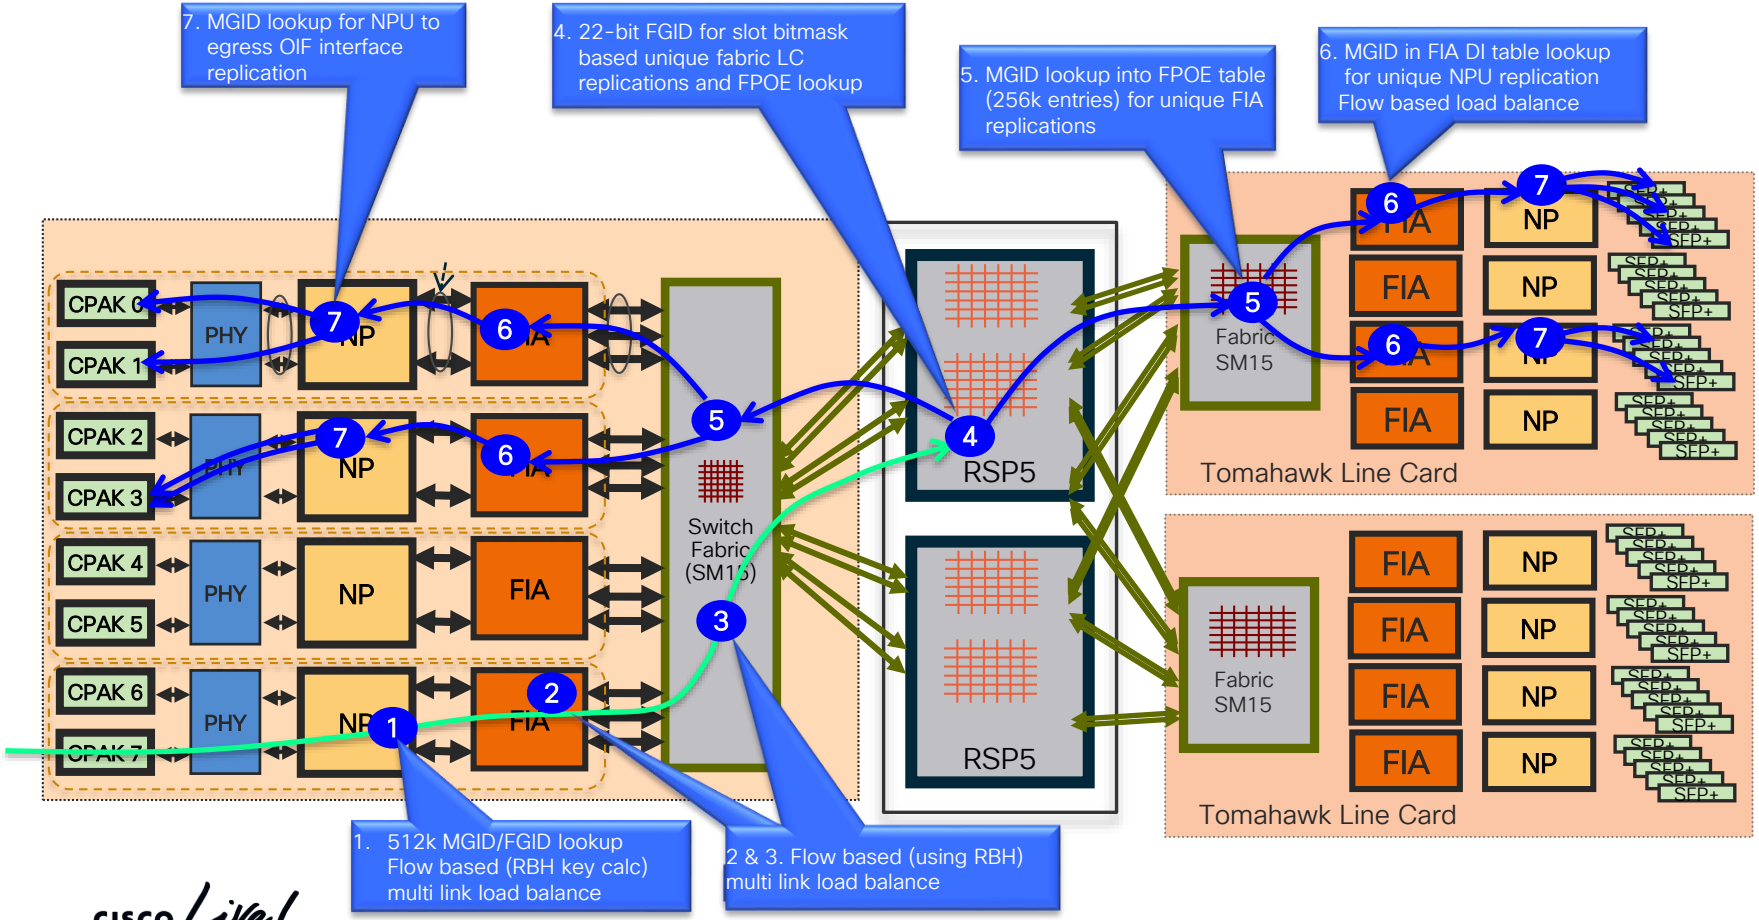

Fabric Load Balancing – Multicast

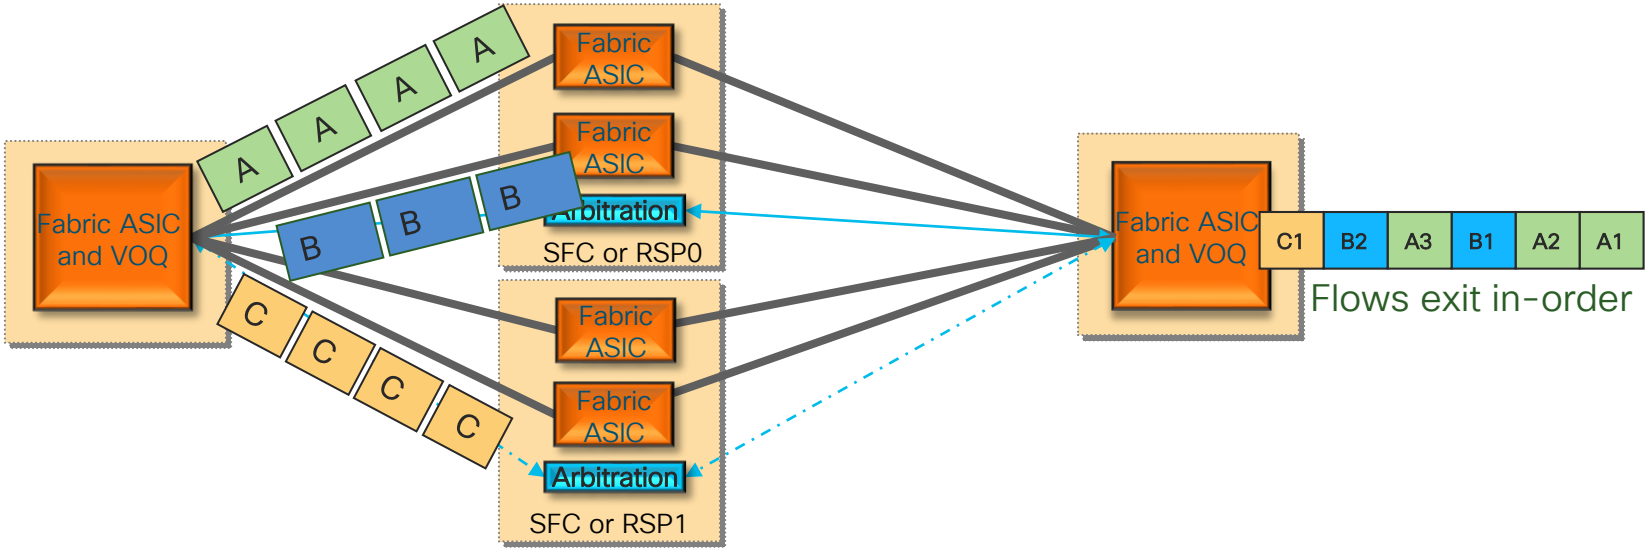

- Multicast traffic hashed based on (S,G) info to maintain flow integrity
  - Very large set of multicast destinations preclude re-sequencing
- Multicast traffic is non arbitrated – sent across a different fabric plane

ASR 9000 QoS Architecture & TCAM usage

ASR9000 Priority-Based QoS Architecture

- Dedicated Traffic Manager(TM) for Traffic Queuing
- User Configurable QoS Policy on Ingress/Egress NP
- End-to-End priority propagation → Guarantee bandwidth, low latency for high priority traffic
- Unicast VOQ and back pressure

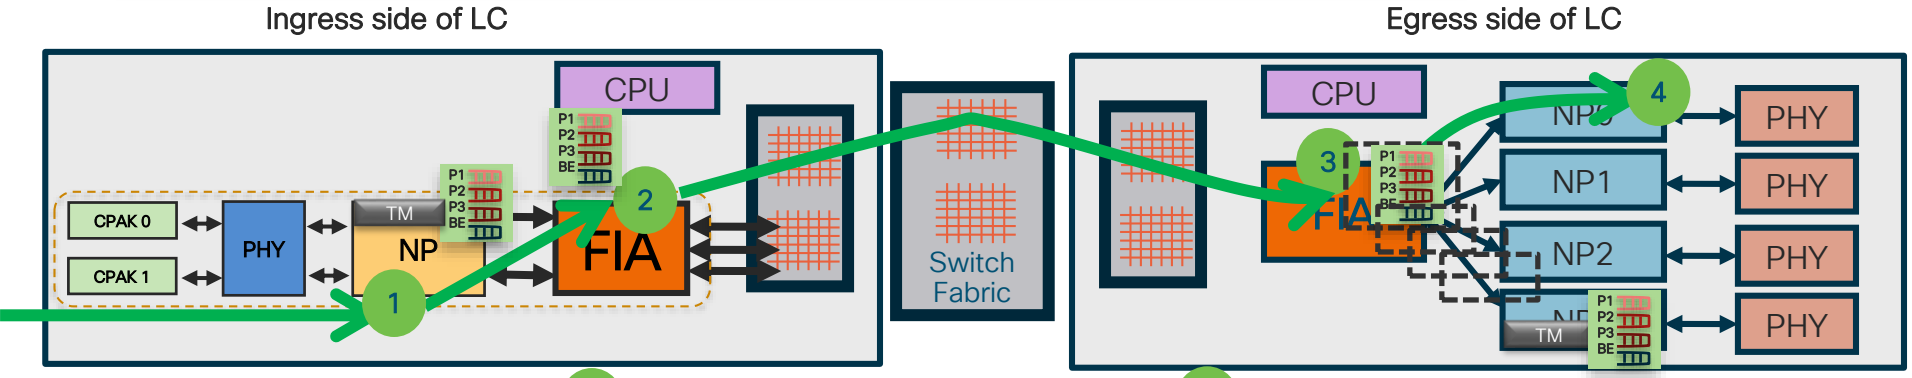

1
- Ingress (sub-)interface QoS Queues
- User Configurable Ingress QoS Policy

User-configuration with Ingress MQC

2
- 4xVOQ per VQI
- Up to 8K VOQs per TSK FIA (vs 4k per SKT FIA)

3
4x Egress Destination Qs per VQI, aggregated at egress port rate

4
Egress (sub-)interface QoS Queues

User-configuration with Egress MQC

Implicit Configuration
Not User-controllable

Default Implicit Trust Model

L2 IF: trust outer Cos
L3 IF: trust DSCP
L3 MPLS: trust outer EXP

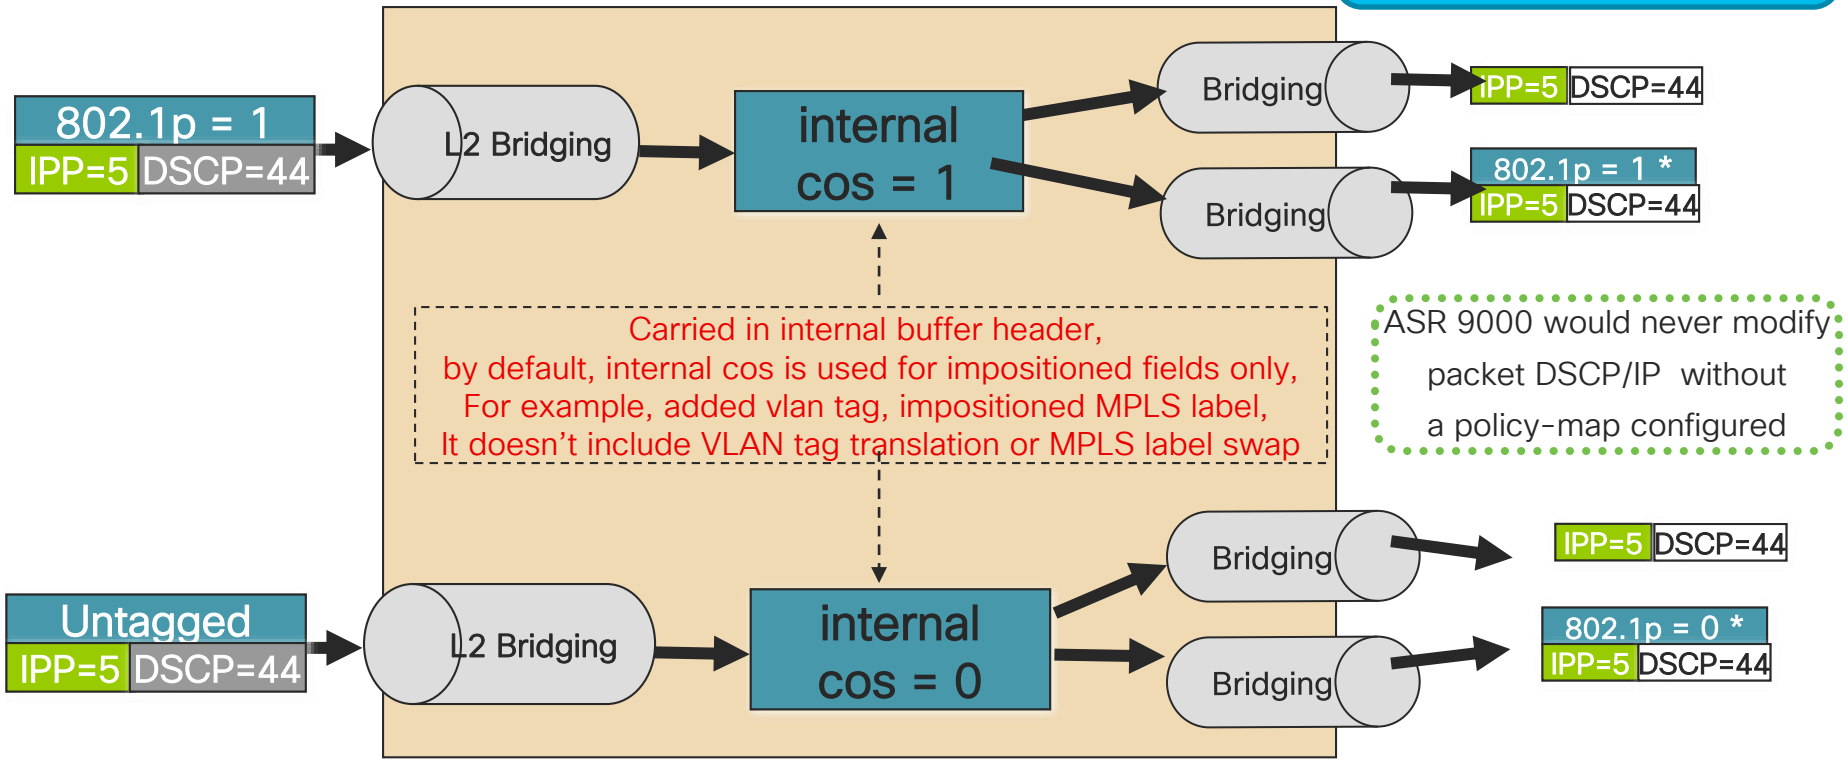

Ingress line card

Egress line card

3-Layer Hierarchical QoS (H-QoS)

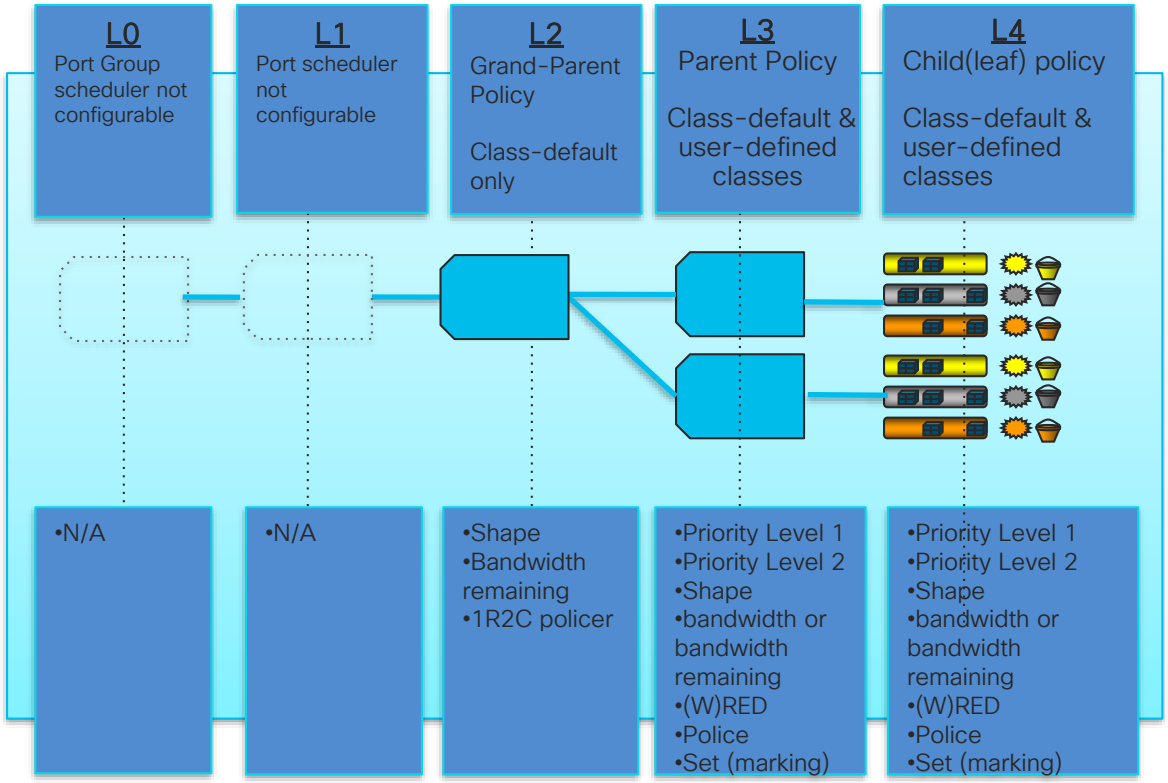


```
policy-map child
 class Pr1
    police rate 64 kbps
    priority level 1
 class Pr2
    police rate 10 mbps
    priority level 2
 class Cl3
    bandwidth 3 mbps
 class Cl4
    set precedence routine
    bandwidth 1 mbps
!
policy-map parent
 class parent1
    shape average 100 mbps
    service-policy child
 class parent2
    shape average 25 mbps
    service-policy child
 class class-default
!
policy-map grand-parent
 class class-default
     shape average 500 mbps
     service-policy parent
```

H-QoS – Supported Classification/Policy

Policy-map hierarchy level	Classification support	Policy Support
Grand-parent	Only class-default	• Shape Average • Bandwidth remaining • 1R2C policer with only drop/transimit action(no set/mark)
Parent	User defined classes with restrictions based on format/interface types.	• Priority/WRED Queue and Queue-limit on Leaf only • Policer/Shaper/Marking/non-Priority Queue/Bandwidth/Bandwidth Remaining
Child	User defined classes with restrictions based on format/interface types.	• Priority/WRED Queue and Queue-limit on Leaf only • Policer/Shaper/Marking/non-Priority Queue/Bandwidth/Bandwidth Remaining

Default Interface Queues

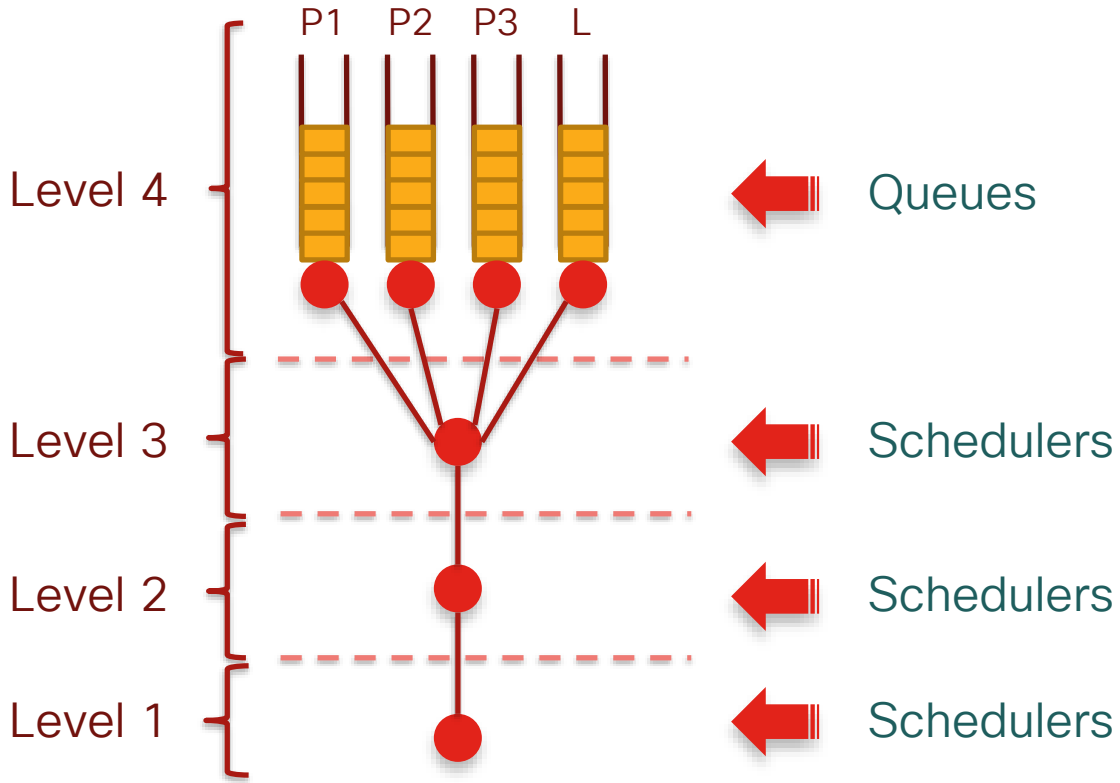

MQC Hierarchy in Queuing ASIC

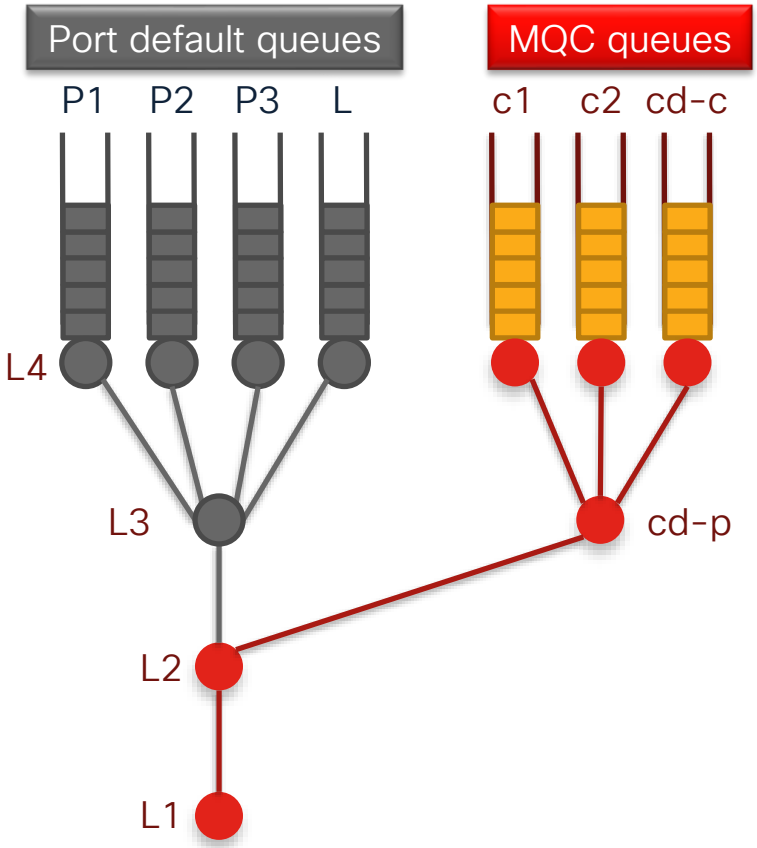


```
policy-map child
 class c1
  priority level 1
  police rate 640 kbps
 class c2
  bandwidth 20 mbps
 class class-default                cd-c
  bandwidth 1 mbps
!
policy-map parent
 class class-default                cd-p
  shape average 35 mbps
  service-policy child
!
interface GigabitEthernet0/0/0/0
 service-policy output parent
```

Inactive entity

Active entity

MQC Hierarchy in Queuing ASIC

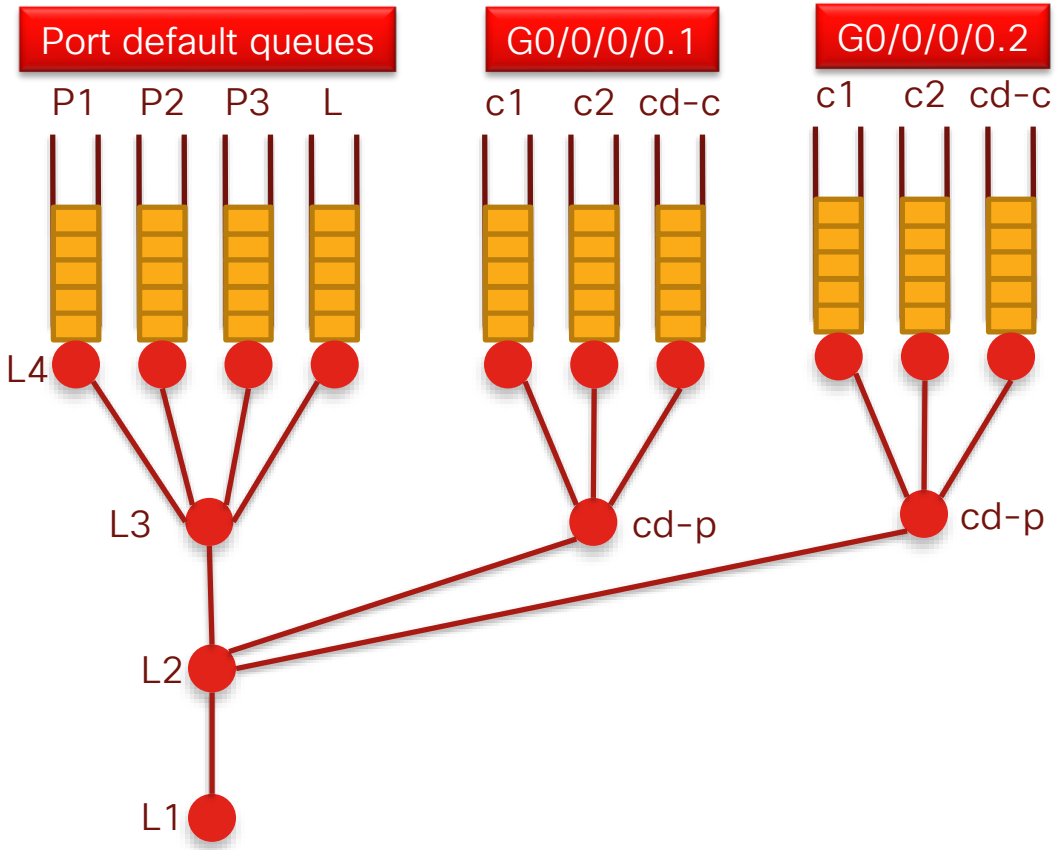


```
policy-map child
 class c1
  priority level 1
  police rate 640 kbps
 class c2
  bandwidth 20 mbps
 class class-default
  bandwidth 1 mbps
!
policy-map parent
 class class-default
  shape average 35 mbps
  service-policy child
!
interface GigabitEthernet0/0/0/0.1
 service-policy output parent
!
interface GigabitEthernet0/0/0/0.2
 service-policy output parent
```

cd-c

cd-p

Inactive entity

Active entity

cisco Live!

TCAM Used for Traffic Classification

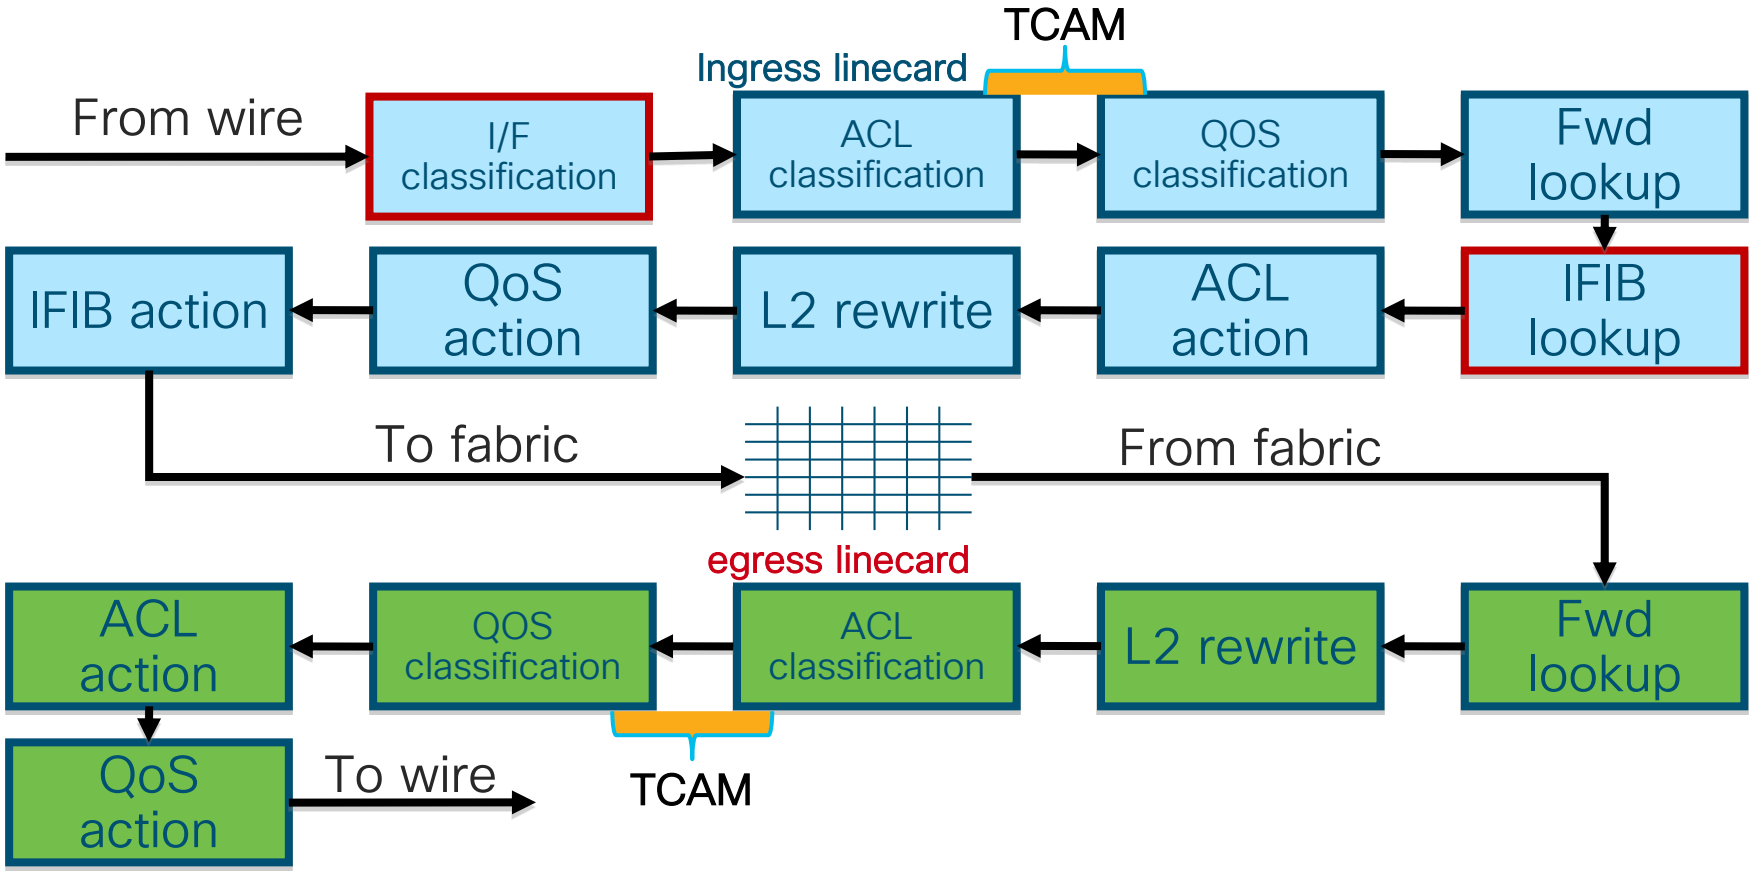

4th/5th Generation LC TCAM Bank Sets

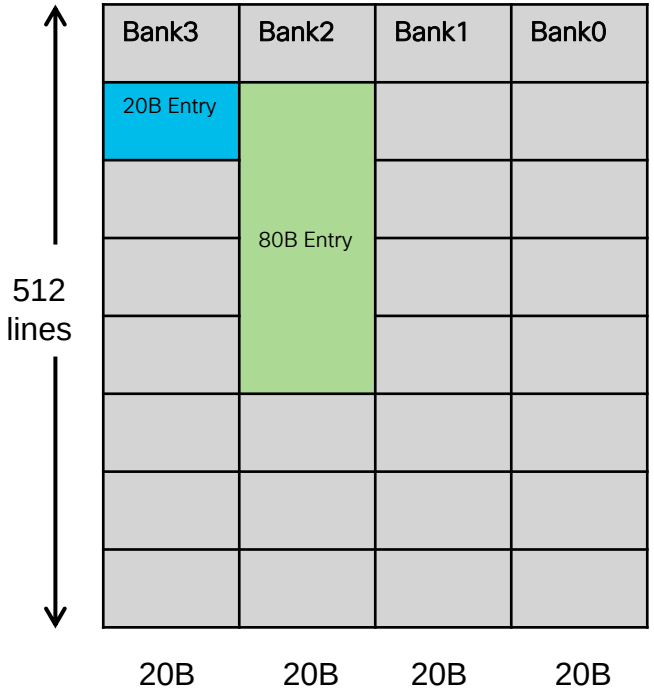

- TCAM organized into bank sets, with 4 arrays in each bank set.
- Each array composed of 512x160b entries or 128x640b entries
- Entries can be arranged as 20B/80B wide and can span multiple lines/blocks
- LPTS on 4th and 5th Generation LC uses hash table instead of iTCAM which saves iTCAM usage

3rd and 4th Generation LC TCAM Partitions

TCAM Regions		3rd Generation LC	4th Generation LC
L2 Partition	Physical Ports, Bundles	Reserved Partition for each type	N/A, L2 interface classification does not require TCAM resource
	Encap Default		
	Encap Untagged, Encap Any		
	Single VLAN, PWHE		
	Double VLAN, BVI		
ODS2 (160 bit entries)	ODS2 iFIB	Reserved Partition for iFIB (IPv4 LPTS)	N/A, LPTS does not require TCAM resource
	Common	Reserved Partition for all ODS2 features	Reserved Partition for all ODS2 features
ODS8(640 bit entries)	ODS8 iFIB	Reserved Partition for iFIB (IPv6 LPTS)	N/A, LPTS does not require TCAM resource
	Common	Reserved Partition for all ODS8 features	Reserved Partition for all ODS8 features

4th Generation LC TCAM Feature Lookup Region

TCAM Region	Features sharing resources	Search mode
160-ING	L2-ACL IPV4-ACL IPv4-QOS PBR-IPV4 PBR-MPLS PBR-L2 IPV4-LI	160 bits
160-EGR	L2-ACL IPV4-ACL IPv4-QOS	160 bits
640-ING	IPV6-ACL IPv6-QOS PBR-IPV6 IPV6-LI EDPL BGP Flowspec	640 bits
640-EGR	IPV6-ACL IPv6-QOS	640 bits

TCAM Partition Example –3rd Generation LC

```
show prm server tcam summary all all np0 location 0/0/CPU0
                    Node: 0/0/CPU0:
----------------------------------------------------------------------

TCAM summary for NP0:

  TCAM Logical Table: TCAM_LT_L2 (1)
     Partition ID: 0, priority: 2, valid entries: 2, free entries: 2046
     Partition ID: 1, priority: 2, valid entries: 0, free entries: 2048
     Partition ID: 2, priority: 0, valid entries: 0, free entries: 2048
     Partition ID: 3, priority: 0, valid entries: 8, free entries: 24568
     Partition ID: 4, priority: 0, valid entries: 5, free entries: 67579
  TCAM Logical Table: TCAM_LT_ODS2 (2), free entries: 89710, resvd 128
     ACL Common Region: 448 entries allocated. 448 entries free
     Application ID: NP_APP_ID_IFIB (0)
       Total: 1 vmr_ids, 8005 active entries, 8005 allocated entries.
     Application ID: NP_APP_ID_QOS (1)
       Total: 5 vmr_ids, 13 active entries, 13 allocated entries.
     Application ID: NP_APP_ID_ACL (2)
       Total: 0 vmr_ids, 0 active entries, 0 allocated entries.
     Application ID: NP_APP_ID_AFMON (3)
       Total: 0 vmr_ids, 0 active entries, 0 allocated entries.
     Application ID: NP_APP_ID_LI (4)
       Total: 0 vmr_ids, 0 active entries, 0 allocated entries.
     Application ID: NP_APP_ID_PBR (5)
       Total: 0 vmr_ids, 0 active entries, 0 allocated entries.
  TCAM Logical Table: TCAM_LT_ODS8 (3), free entries: 15200, resvd 128
     ACL Common Region: 448 entries allocated. 448 entries free
     Application ID: NP_APP_ID_IFIB (0)
       Total: 1 vmr_ids, 603 active entries, 603 allocated entries.
     Application ID: NP_APP_ID_QOS (1)
       Total: 0 vmr_ids, 0 active entries, 0 allocated entries.
     Application ID: NP_APP_ID_ACL (2)
       Total: 1 vmr_ids, 5 active entries, 5 allocated entries.
     Application ID: NP_APP_ID_PBR (5)
       Total: 0 vmr_ids, 0 active entries, 0 allocated entries.
     Application ID: NP_APP_ID_EDPL (6)
       Total: 0 vmr_ids, 0 active entries, 0 allocated entries.
```

Note: “show controllers rm tcam summary all all np X location X/X/X” for 4th Generation LC

TCAM Partition Example – 4th Generation LC

160-ING Region

```
TCAM summary for NP0:

 TCAM Region: 160-ING (0)
 max entries: 16384, num free: 16367
   Feature ID: IPV4-ACL (0)
     VMR ID:    1, used entries:   14, allocated entries:   14
     Total vmr_ids per feature id: 1, Total used entries per feature id: 14 Total allocated entries: 14
   Feature ID: L2-ACL (2)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: IPV4-LI (4)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: QOS-FMT-0 (8)
     VMR ID:    2, used entries:    3, allocated entries:    3
     Total vmr_ids per feature id: 1, Total used entries per feature id: 3 Total allocated entries: 3
   Feature ID: QOS-FMT-1 (9)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: QOS-FMT-2 (10)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: QOS-FMT-3 (11)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: QOS-FMT-15 (13)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: PBR-IPV4 (14)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: PBR-MPLS (15)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: PBR-L2 (17)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: PBR-FMT-0 (18)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
   Feature ID: IPV4-META-ACL (20)
     Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
```

TCAM Partition Example – 4th Generation LC

160-ING Region

```
TCAM summary for NP0:

 TCAM Region: 160-EGR (1)
 max entries: 8192, num free: 8186
    Feature ID: IPV4-ACL (0)
      VMR ID:    1, used entries:    6, allocated entries:    6
      Total vmr_ids per feature id: 1, Total used entries per feature id: 6 Total allocated entries: 6
    Feature ID: L2-ACL (2)
      Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
    Feature ID: QOS-FMT-0 (8)
      Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
    Feature ID: QOS-FMT-1 (9)
      Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
    Feature ID: QOS-FMT-2 (10)
      Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
    Feature ID: QOS-FMT-3 (11)
      Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
    Feature ID: QOS-FMT-15 (13)
      Total vmr_ids per feature id: 0, Total used entries per feature id: 0 Total allocated entries: 0
```

5th Gen LC Default Dynamic TCAM Carving

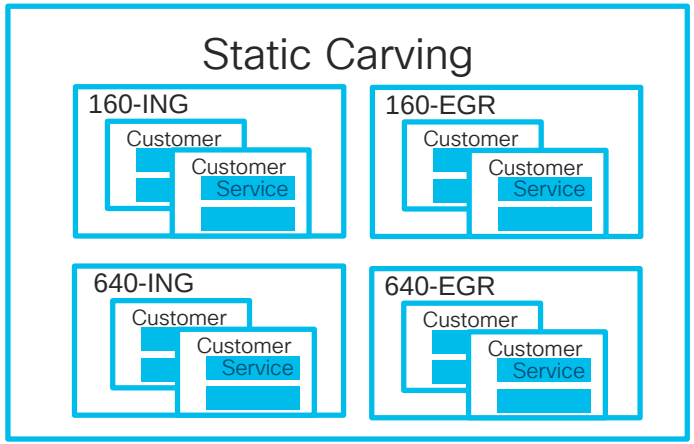

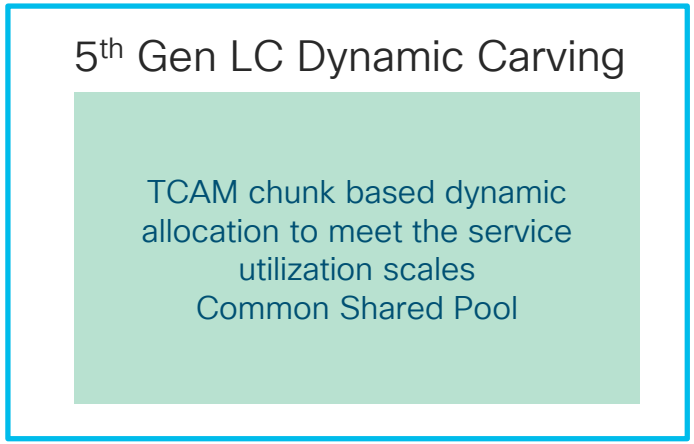

- LSP -SE by default with dynamic 64K TCAM allocation
  - TCAM usage on-demand, TCAM allocation at minimum block size 4K
  - TCAM has no ingress/egress, or IPv4/IPv6 boundary at LC bootup
  - LC reload required when switching to static carving
- LSP -TR by default with dynamic 40K TCAM allocation
  - LSP -TR has 40K TCAM space at LC bootup with block size 4K
  - LSP -TR re-carving TCAM into 16 blocks at minimum block size 2.5K
  - LC reload required when switching to static carving
- Per NPU Control

TCAM Usage Summary

- **5th generation LC** Supports Dynamic TCAM Carving/Allocation for Flexible MD-Scales
- **No TCAM resource** required for L2 Interface Classification and LPTS on 4th and 5th Generation LCs.
- **IPv4-ACL/QoS search mode:** 160bits(20Bs) Partition
- **IPv6-ACL/QoS search mode:** 640bits(80Bs) Partition
- **BGP FlowSpec (both v4 and v6):** 640bits(80Bs) Partition. Ingress Direction Only
- **PBR/Lawful Interception(LI):** Ingress Direction Only

Pop Quiz ????

For an incoming untagged layer 2 frame, what is the priority value used when the frame is processed inside ASR9000?

- 2
- COS bit
- 802.1p Value
- 0

Pop Quiz ????

As we discussed, most of IPv4 related features use ODS2 partition and most of IPv6 related features use ODS8 partition. When BGP flowspec is configured for IPv4 traffic and rules has been pushed, in which TCAM region are the rules programed?

- L2 Partition
- ODS8 Common
- ODS2 Common
- ODS2 iFIB

Conclusion

- ASR9000 – Truly Carrier-Class Edge Router Provides:
  - Rich Features, Flexible Service Capability
  - Variety of Hardware to Meet Different Capacity Requirements
- Fully Distributed Architecture for High Performance and Massive Scalability
- Uniform, Open and Modularized Software Architecture

Complete your online session survey

- 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 t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Content Catalog on ciscolive.com/emea.

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

cisco Live!

Continue your education

Demos in the Cisco campus

Walk-in labs

Meet the engineer 1:1 meetings

Related sessions

cisco Live!

Thank you

cisco Live!

You make **possible**