



Cisco Support Community Expert Series Webcast:

Understanding Cisco ASR 9000 Series Aggregation Services Routers Platform Architecture and Packet Forwarding Troubleshooting

Xander Thuijs, CCIE #6775 ISP/VoIP
Principal Engineer, ASR9000 & IOS-XR

5/14/2013

Cisco Support Community – Expert Series Webcast

- Today's featured expert is Cisco Engineer “**Expert**”
- Ask him and the team questions now about the ASR9000 and IOS-XR



Xander Thuijs

CCIE #6775 in ISP and VoIP

Who are we

Members of the ASR9000 Escalation team present today:

- Aleksandar Vidakovic
Sr Technical Leader, Spain
- Sadanande Phadke
Technical Leader, Boston, MA
- Krishna Eranti
SW Engineer, RTP, NC

Thank You for Joining Us Today

Today's presentation will include audience polling questions

We encourage you to participate!



Thank You for Joining Us Today

If you would like a copy of the presentation slides, click the PDF link in the chat box on the right or go to

<https://supportforums.cisco.com/community/netpro/network-infrastructure/routing>

Or, <https://supportforums.cisco.com/docs/DOC-32967>



Thank You for Joining Us Today

Everyone who joins today's webcast will receive:

125 Cisco Preferred Access Points!



Polling Question 1

What is your level of experience with ASR9000?

- a) I am new to the ASR9000 and IOS-XR
- b) I have a good understanding of IOS-XR, but I am relatively new to ASR9000
- c) I have a good understanding of IOS-XR and I know how to work with the system
- d) I am using ASR9000 already and I am pretty good with it

Submit Your Questions Now!

Use the Q&A panel to submit your questions. Experts will start responding those





Cisco Support Community Expert Series Webcast:

ASR9000 Architecture

Xander Thuijs

Principal Engineer

Cisco's Product Security Incident Response Team (PSIRT)

Security Research and Operations

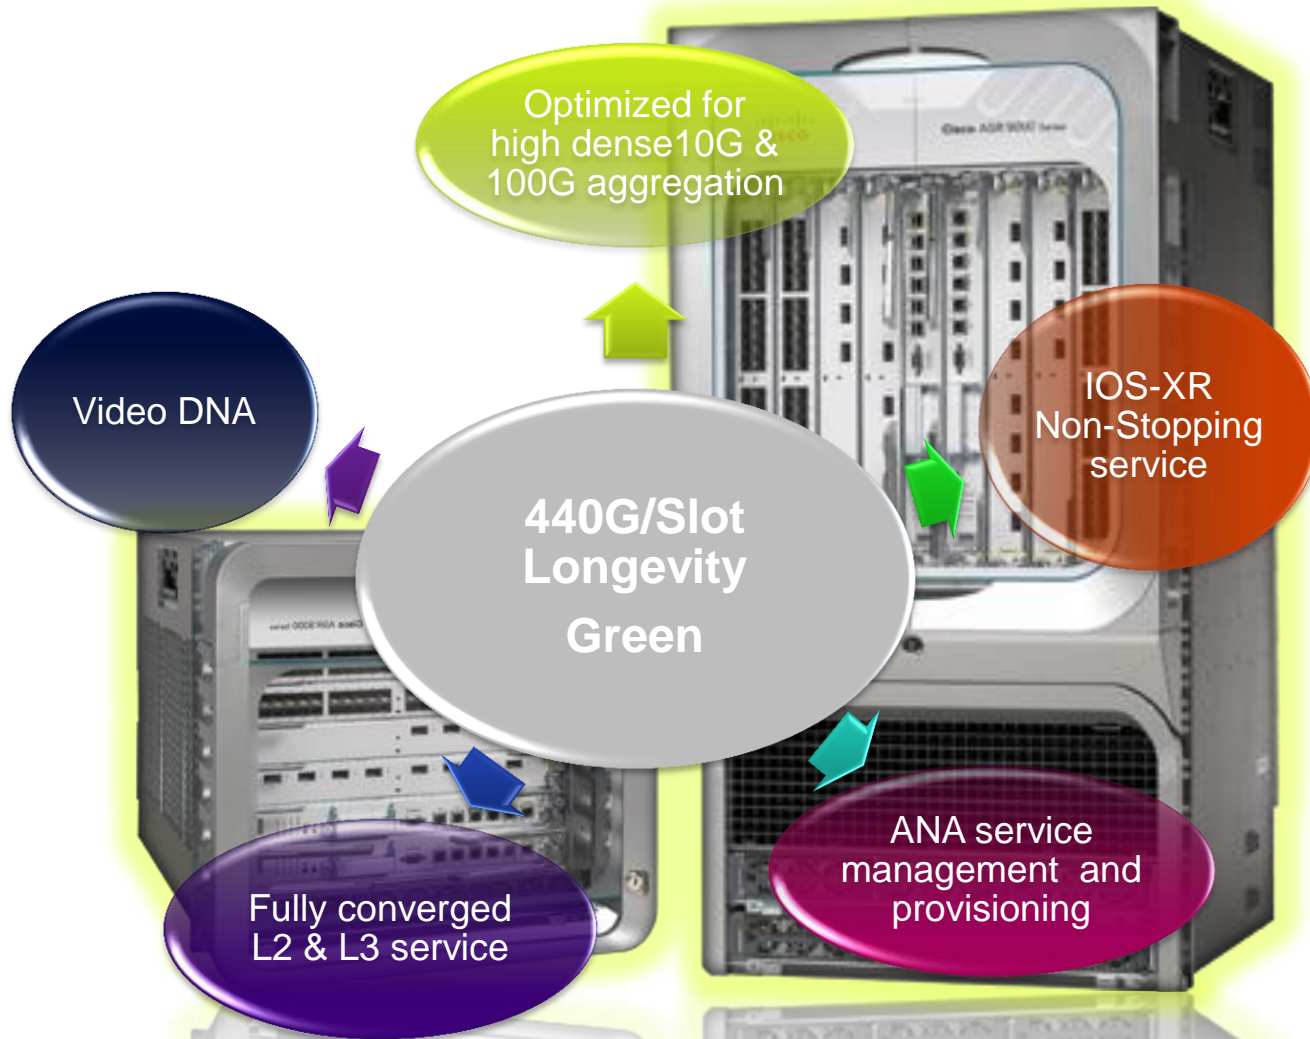
Agenda

- Platform architecture
- Linecard Architecture
- Packet flow through the Linecard and Fabric
- Feature order of operation
- NPU Architecture
- Scale
- QOS implementation
- FIB implementation

*We will not be discussing the in depths of IOS-XR or configuration today
(to be scheduled for a future session)*

ASR 9000 “At-a-Glance”

Next-Generation SP Edge & Aggregation




ASR 9000 System Architecture “At-a-Glance”

Fully Distributed Architecture for High Performance and High Multi-dimensional Control Plane Scale

- Data forwarding is fully distributed on the line cards
- Control plane split among RSP and LC CPU (same type of CPU as RSP)
- L2 protocols, BFD, CFM, Netflow runs on the LC CPU for high scale


True Modular OS for HA and Operational Simplicity

- Micro-kernel based, true modular OS
- High availability and System stability
- SW patch granularity for operational simplicity



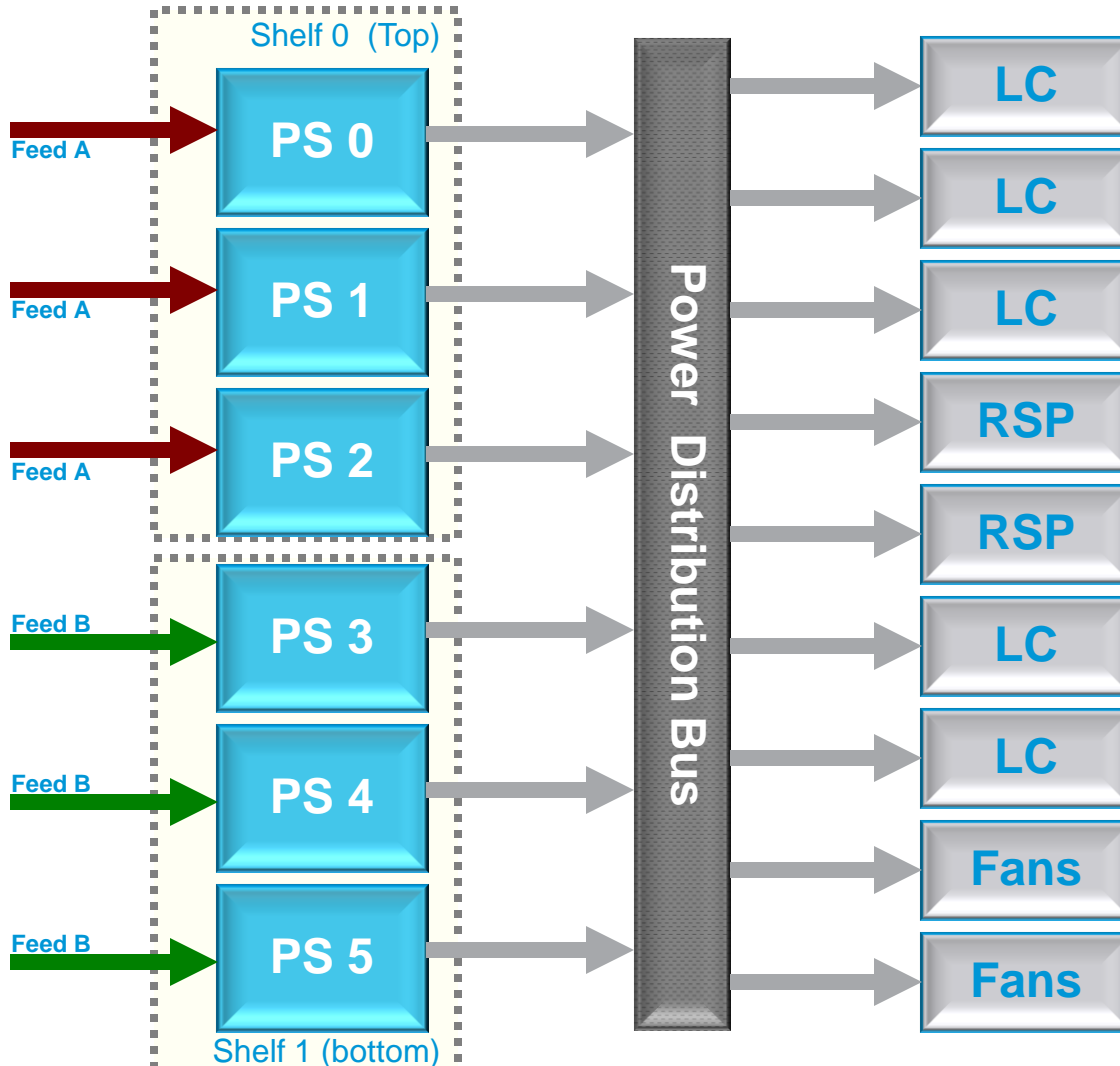
Active-Active Switch Fabric
Guarantee “0” packet loss during RSP failover

ASR 9010 and ASR 9006 Chassis



Power Distribution (AC 1:1 protection)

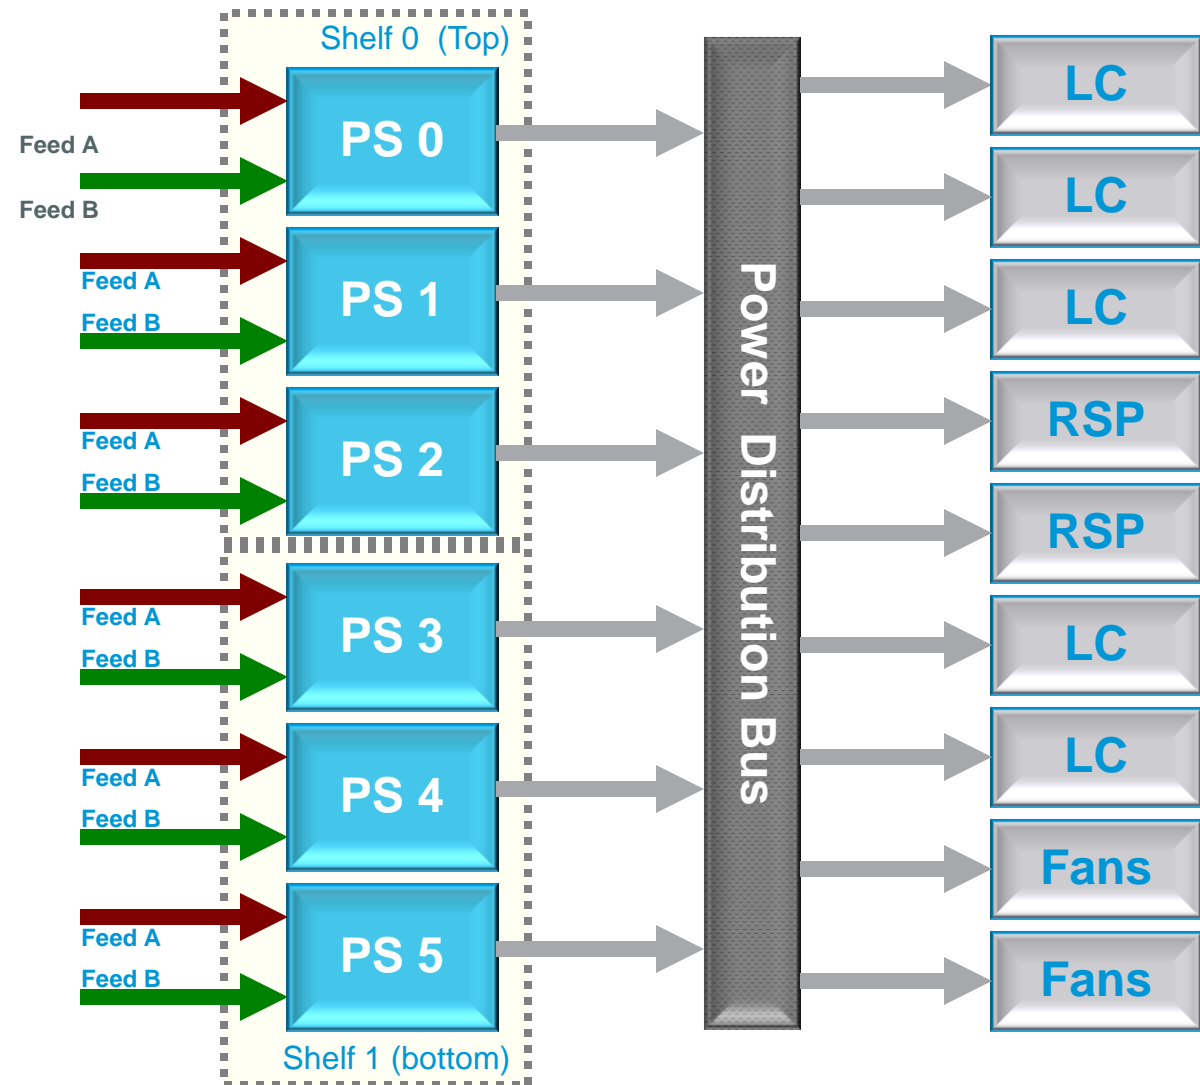
10 slots chassis (v1 power layout shown)



- Single power zone, one distribution bus
- All modules load share
- AC power supplies are rated @ 3KW each
- 'A' feed wired to top power shelf
- 'B' feed wired to bottom power shelf

Power Distribution (DC N:1 protection)

10 slots chassis (v1 power layout shown)



- Single power zone, one distribution bus
- All modules load share
- 2kW and 1.5kW supplies
- Each power supply is wired to both 'A' and 'B' feed
- Feed failure doubles draw on remaining feed
- supply failure increases draw on remaining supplies

ASR 9K RSP (Route/Switch Processors) Overview

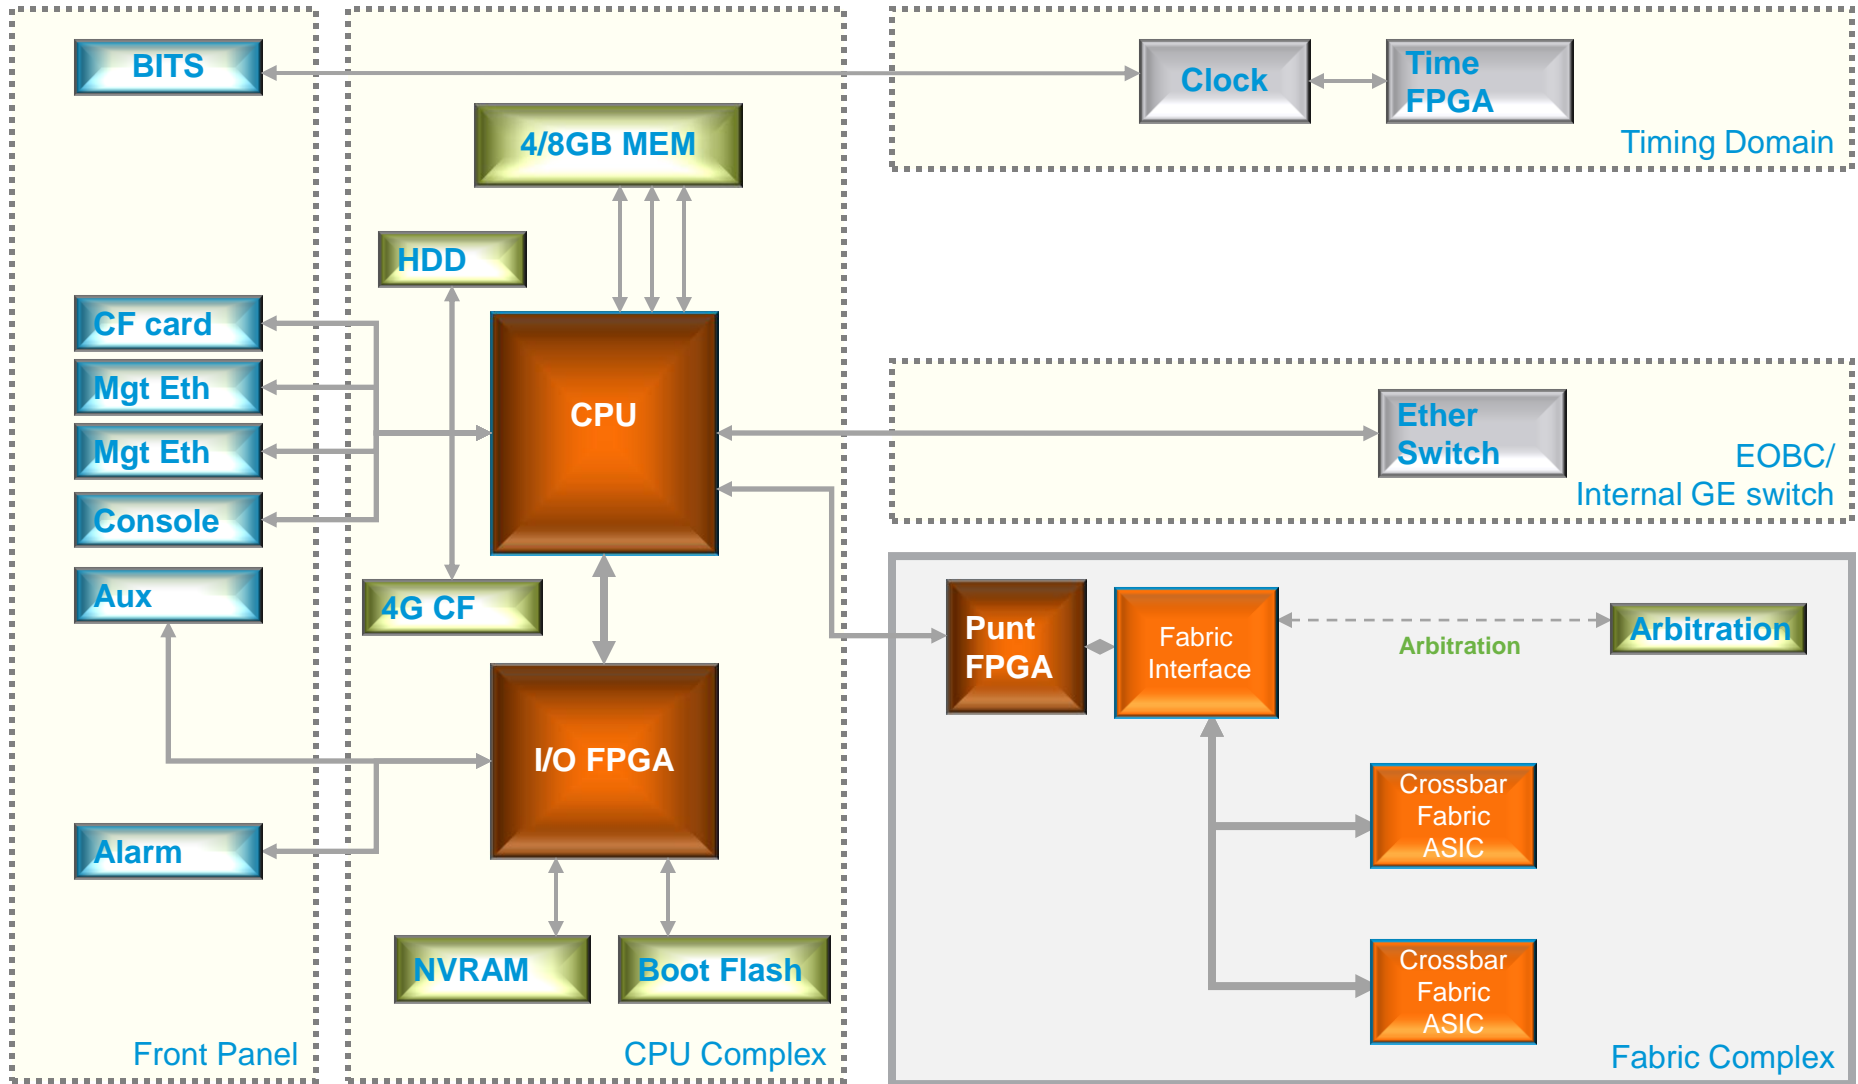
	Current RSP2	RSP440
Processors	2 x 1.5GHz Freescale 8641D CPU	Intel x86 Jasper Forest 4 Core 2.27 GHz
RAM (user expandable)	4GB @133MHz SDR 8GB	6GB (RSP440-TR) and 12GB (RSP440-SE) version @1066MHz DDR3
Cache	L1: 32KB L2: 1MB	L1: 32KB per Core L2: 8MB shared
Primary persistent storage	4GB	16GB - SDD
Secondary persistent storage (HD/SSD)	30GB - HDD	16GB - SDD
USB 2.0 port	No	Yes
Acceleration / Security	No	Yes
HW assisted CPU queues	No	Yes
nV Cluster – EOBC ports	No	Yes, 2 x 1G/10G SFP+
Switch fabric bandwidth	184G/slot (with dual RSP)	440G/slot (with dual RSP)



RSP440

Generic RSP Engine Architecture

(RSP2 shown)



ASR 9K Ethernet Line Card Overview

First-generation
LC (Trident NP)
-L, -B, -E

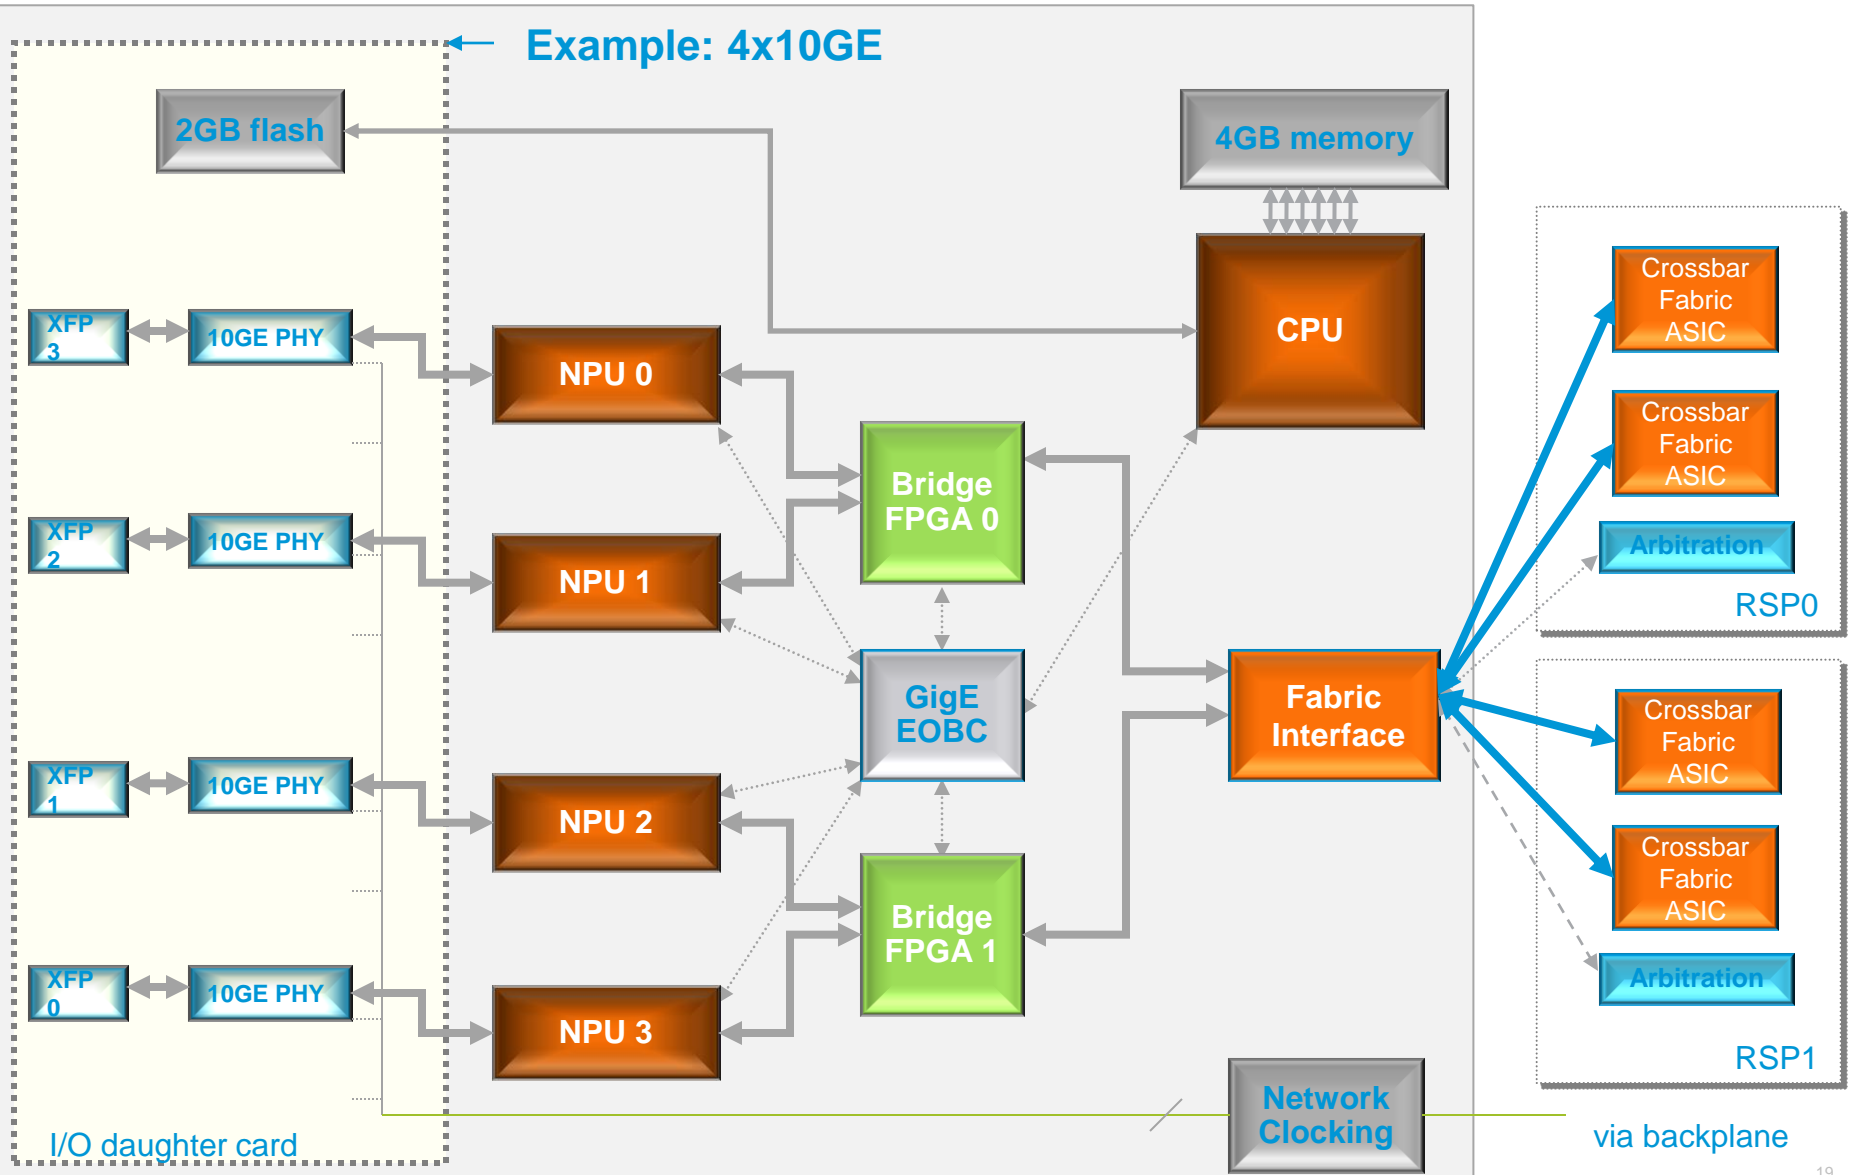


Second-generation
LC (Typhoon NP)
-TR, -SE



MPAs
20x1GE
2x10GE
4x10GE
1x40GE
2x40GE

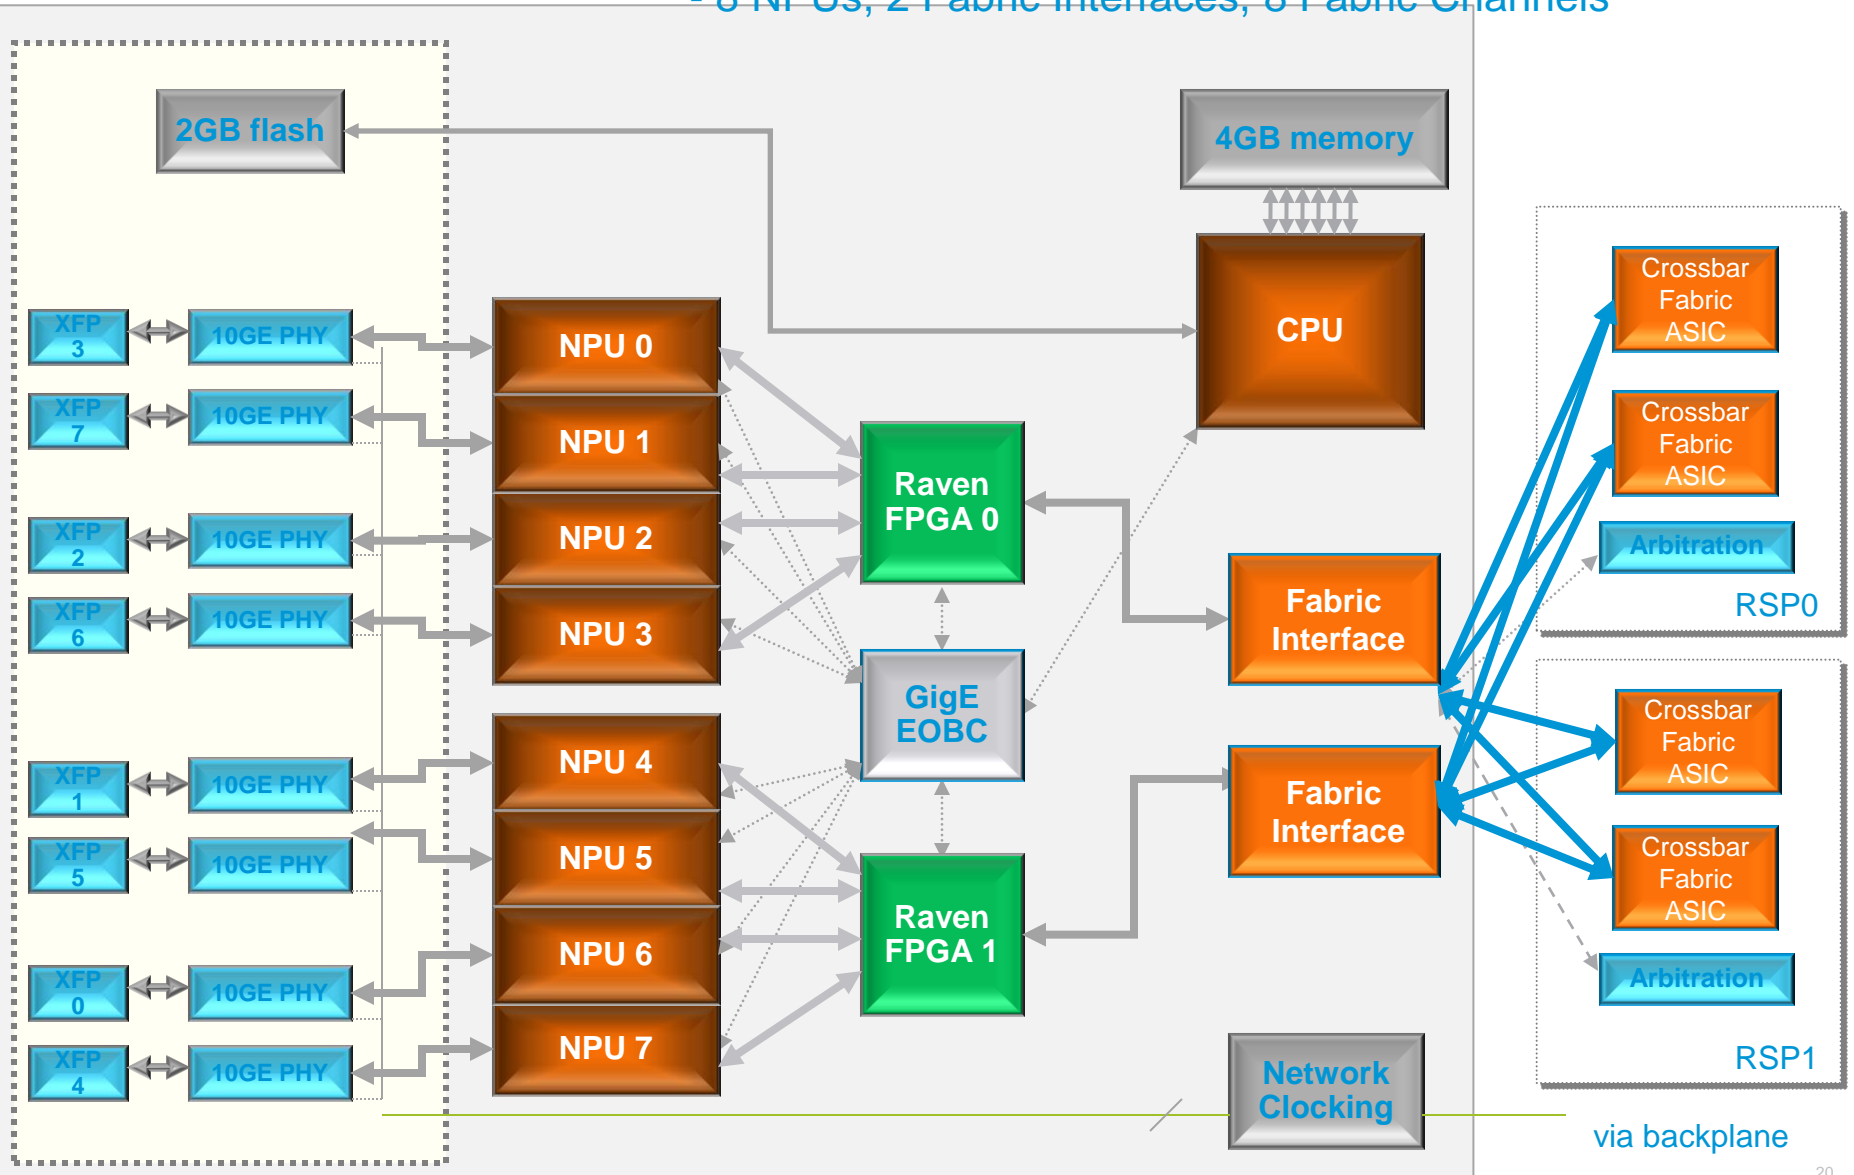
40G Line Card Hardware Architecture



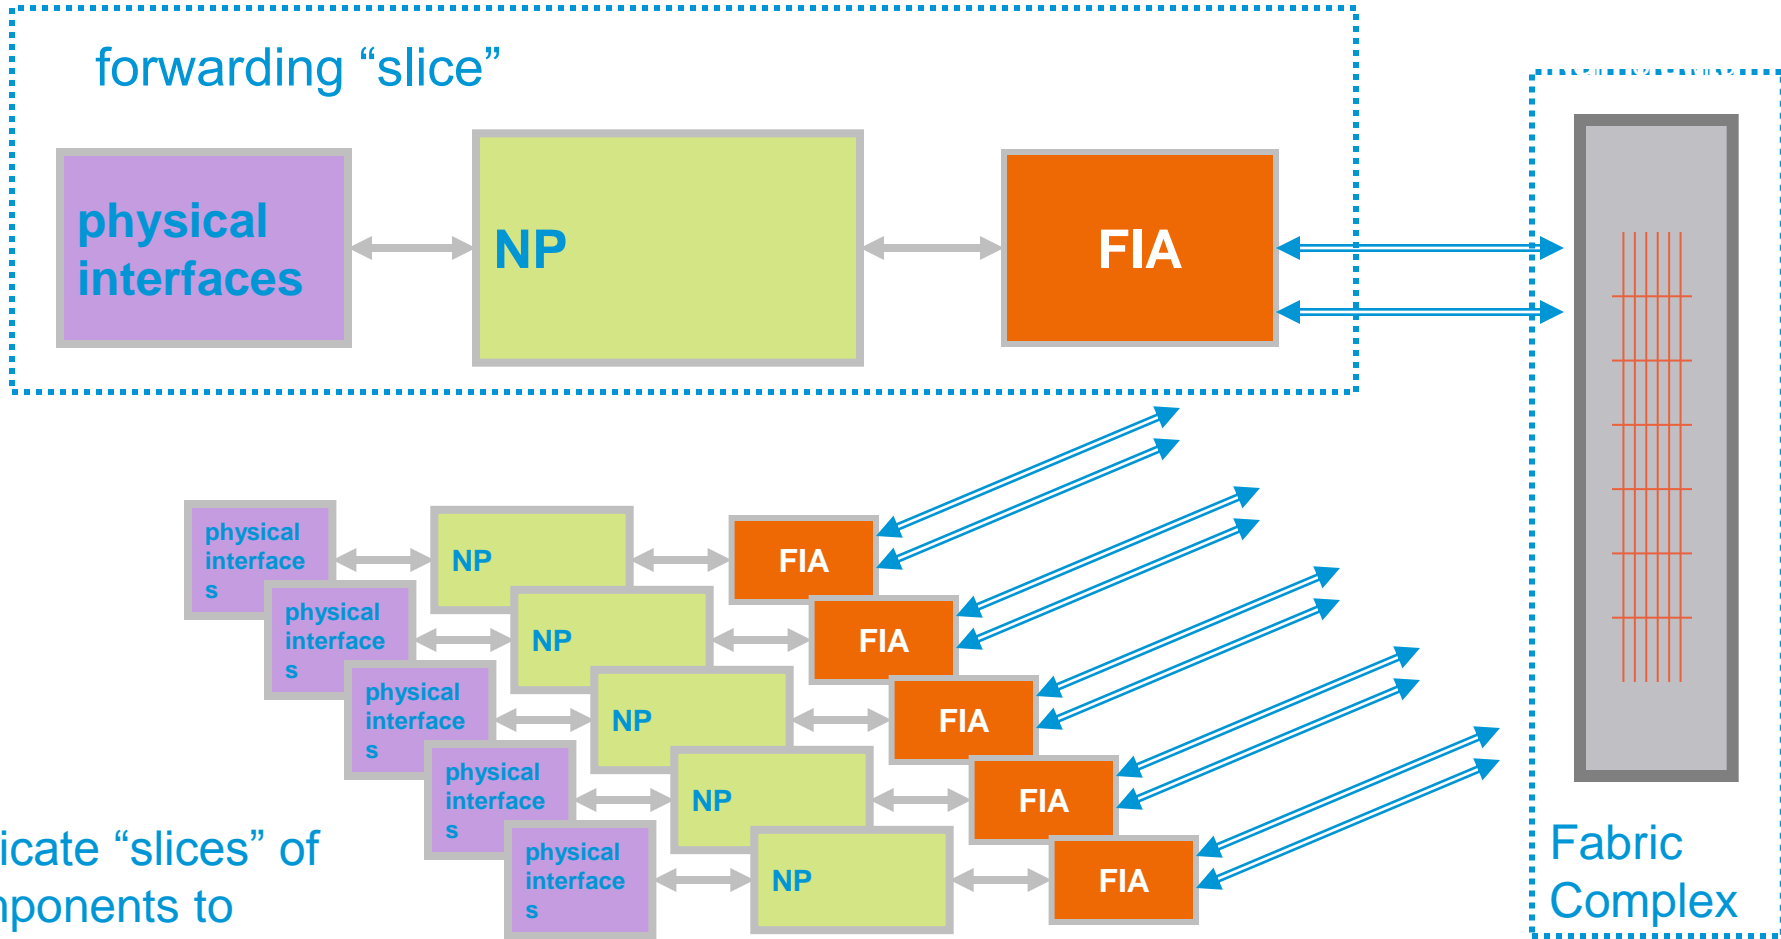
80G Line rate Line Card Hardware Architecture

Number of HW elements on LC doubles:

- 8 NPUs, 2 Fabric Interfaces, 8 Fabric Channels

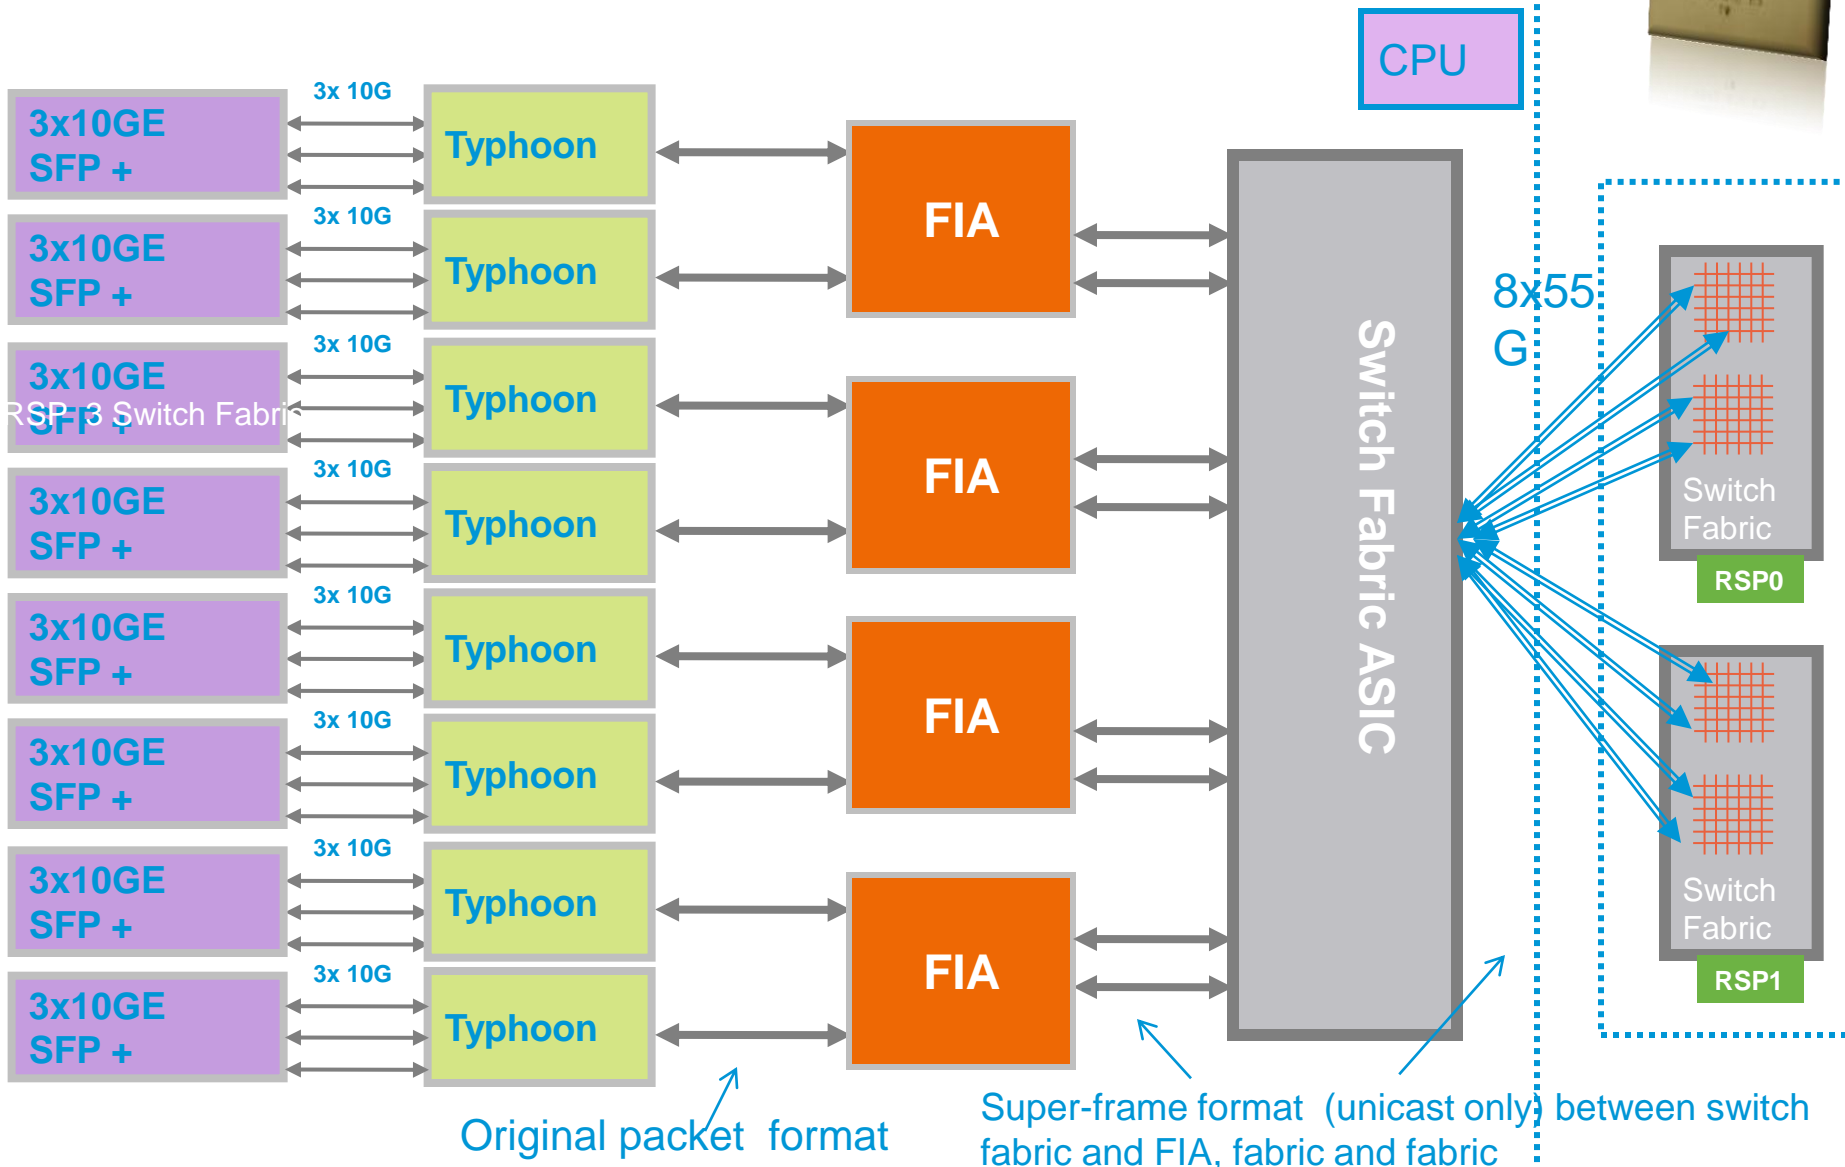


Generic LC Architecture (2) – Forwarding Slice

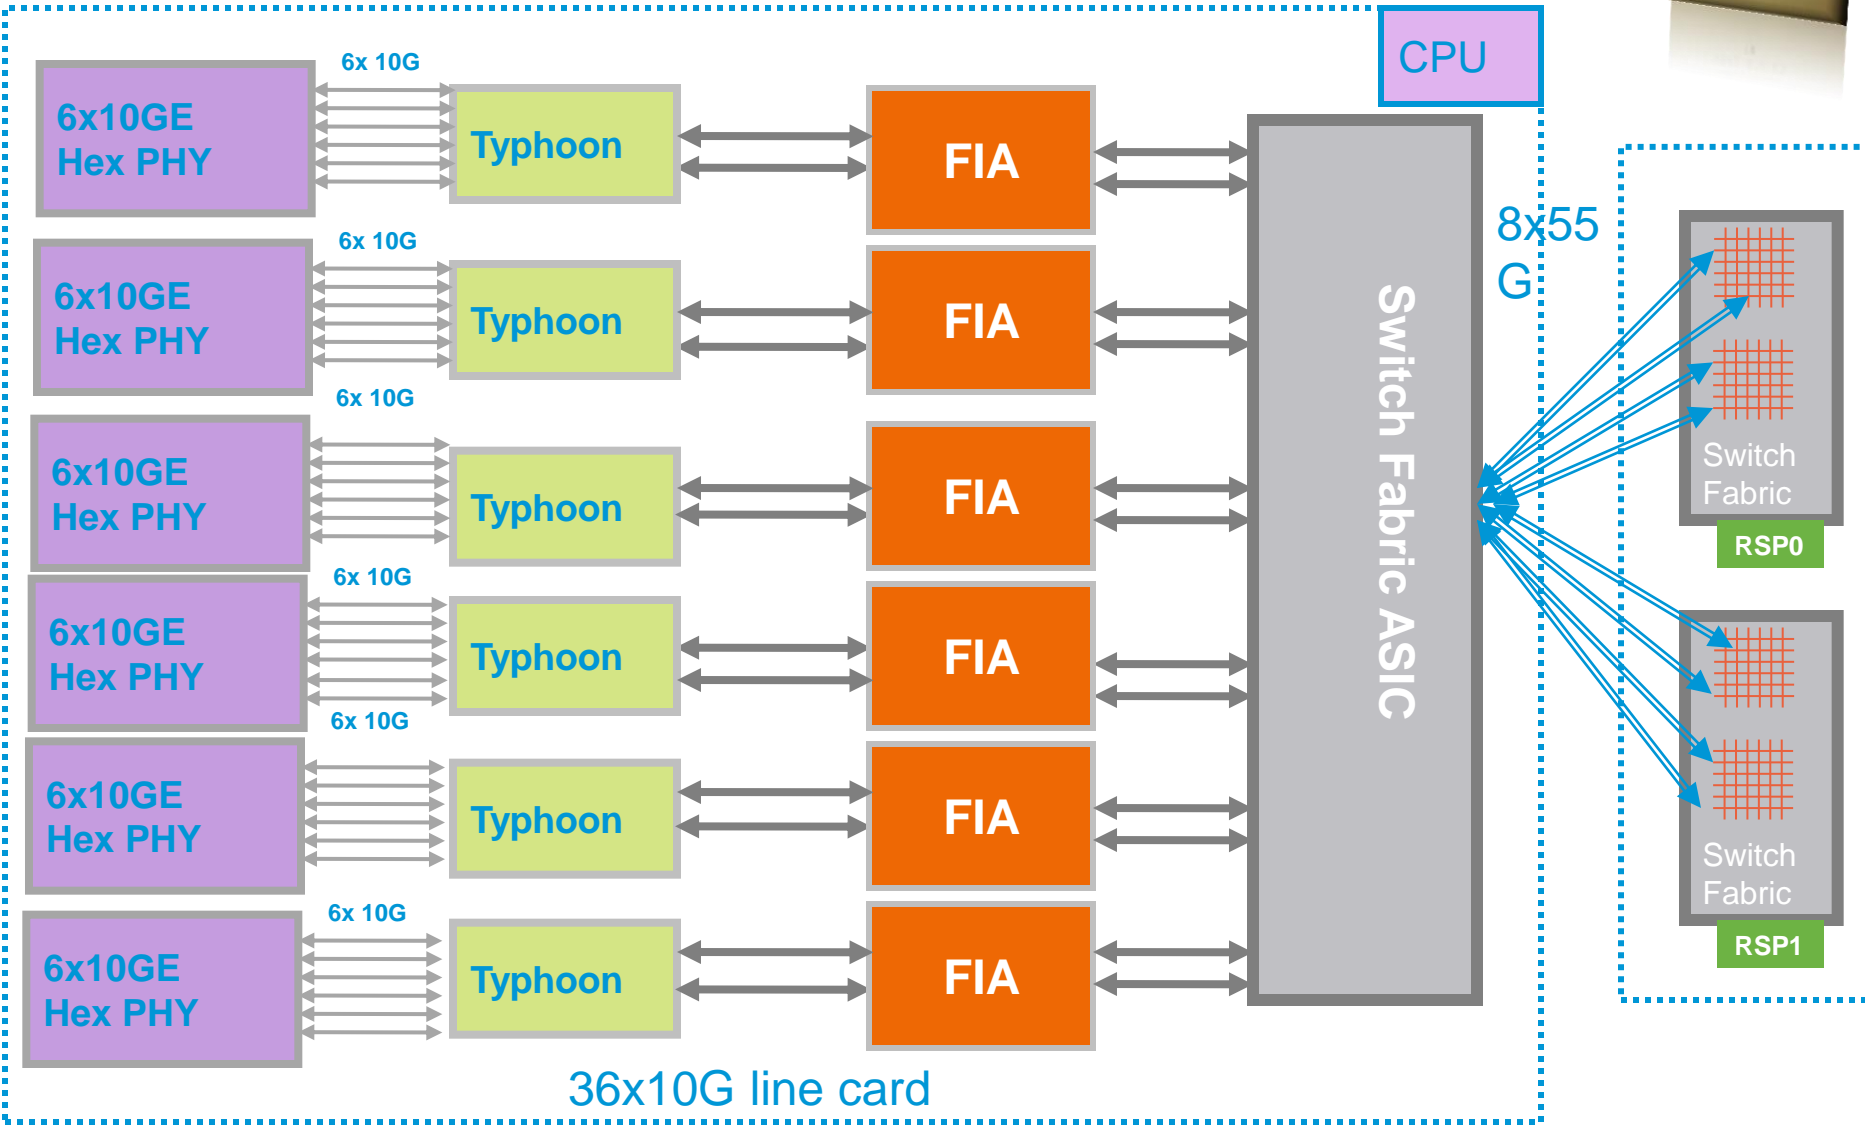


Typhoon LC CPU: Freescale Quad core P4040

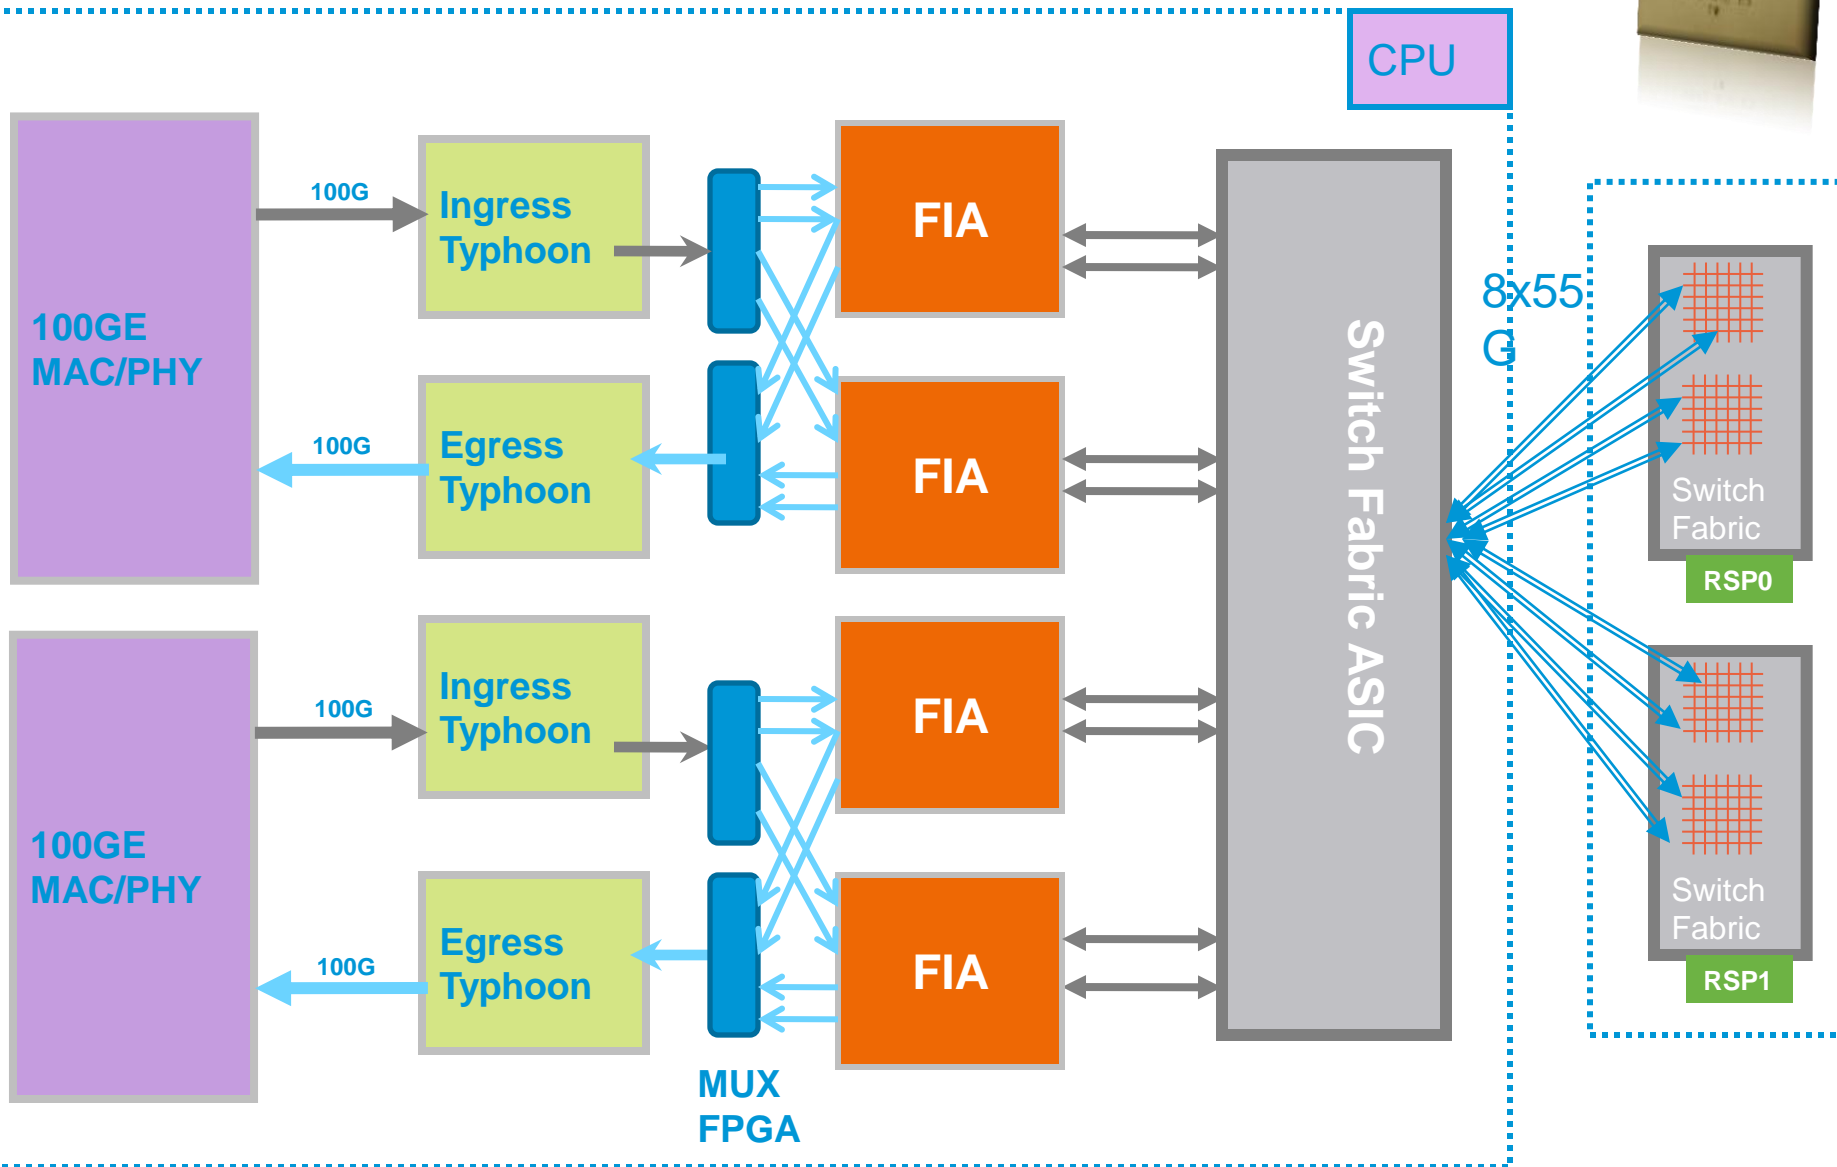
LC Architecture – 24x10G



LC Architecture – 36x10G

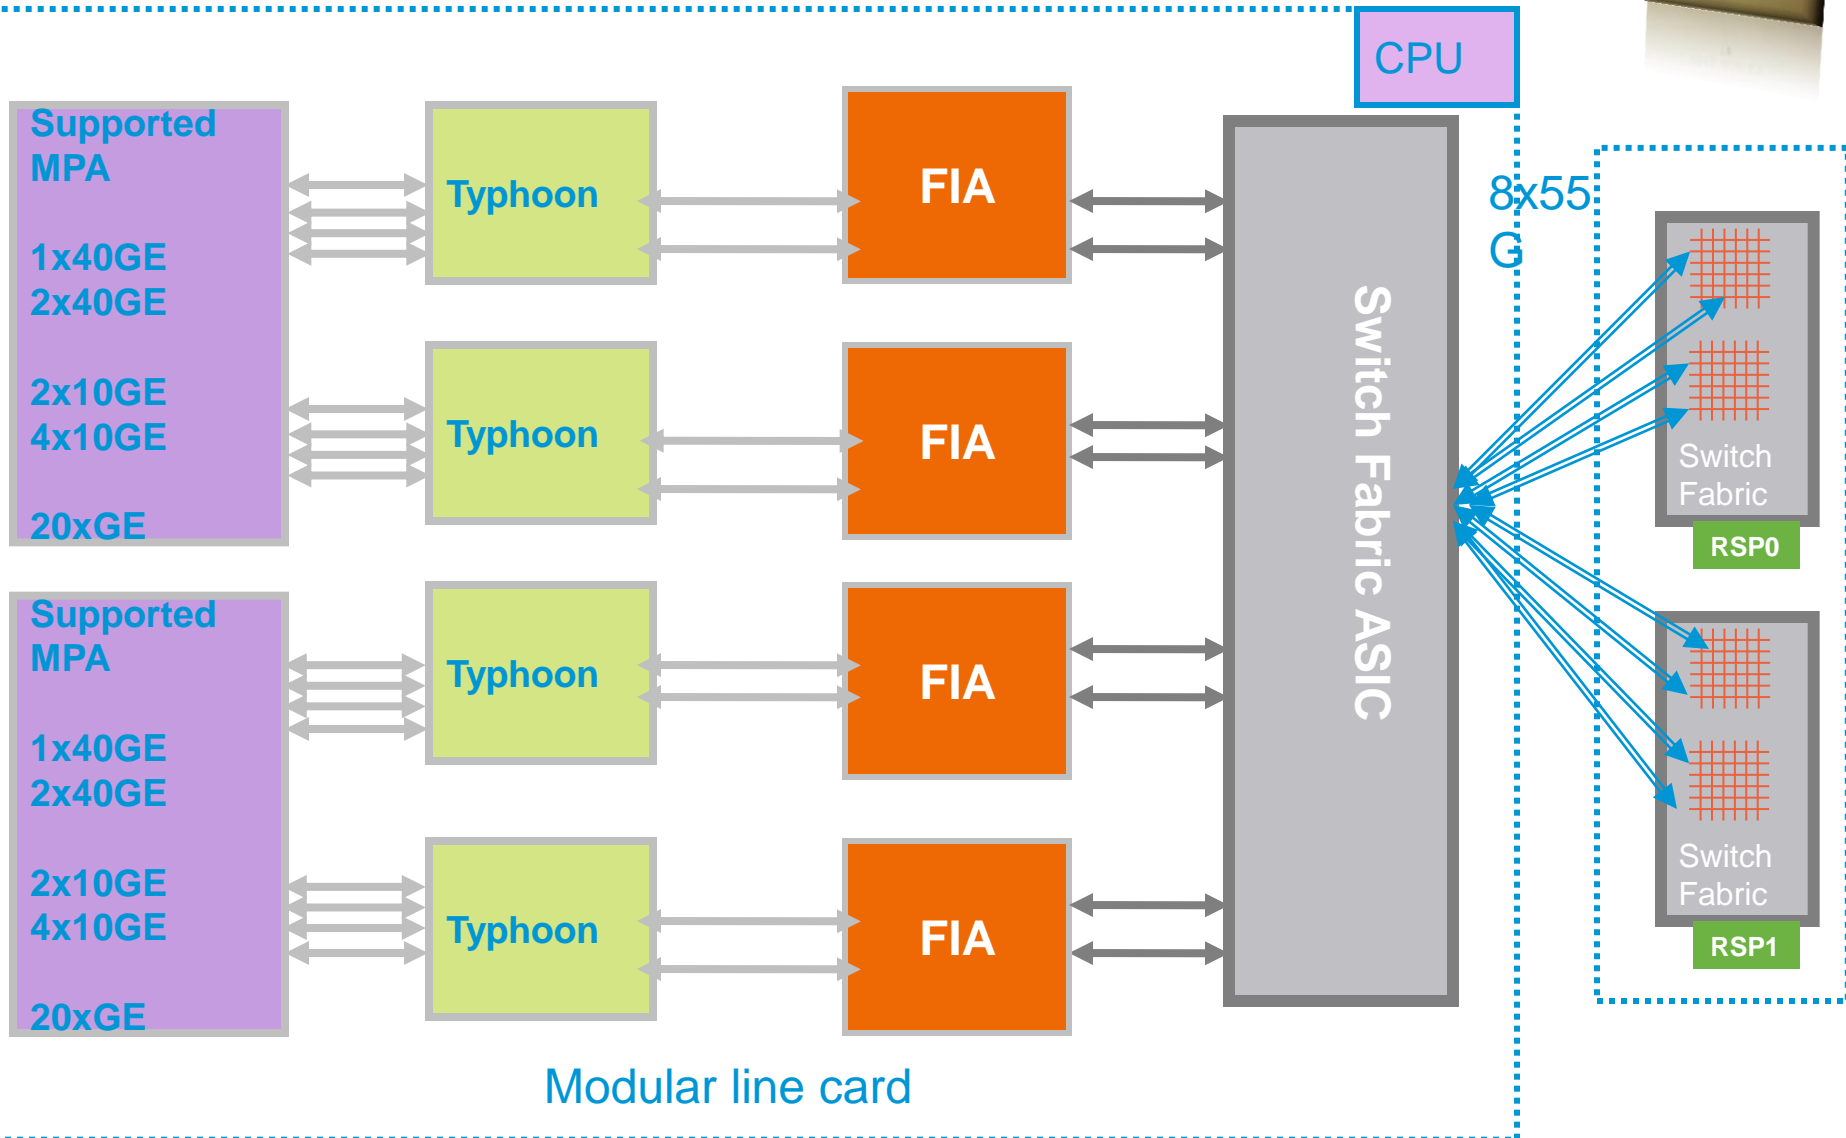


LC Architecture – 2x100G



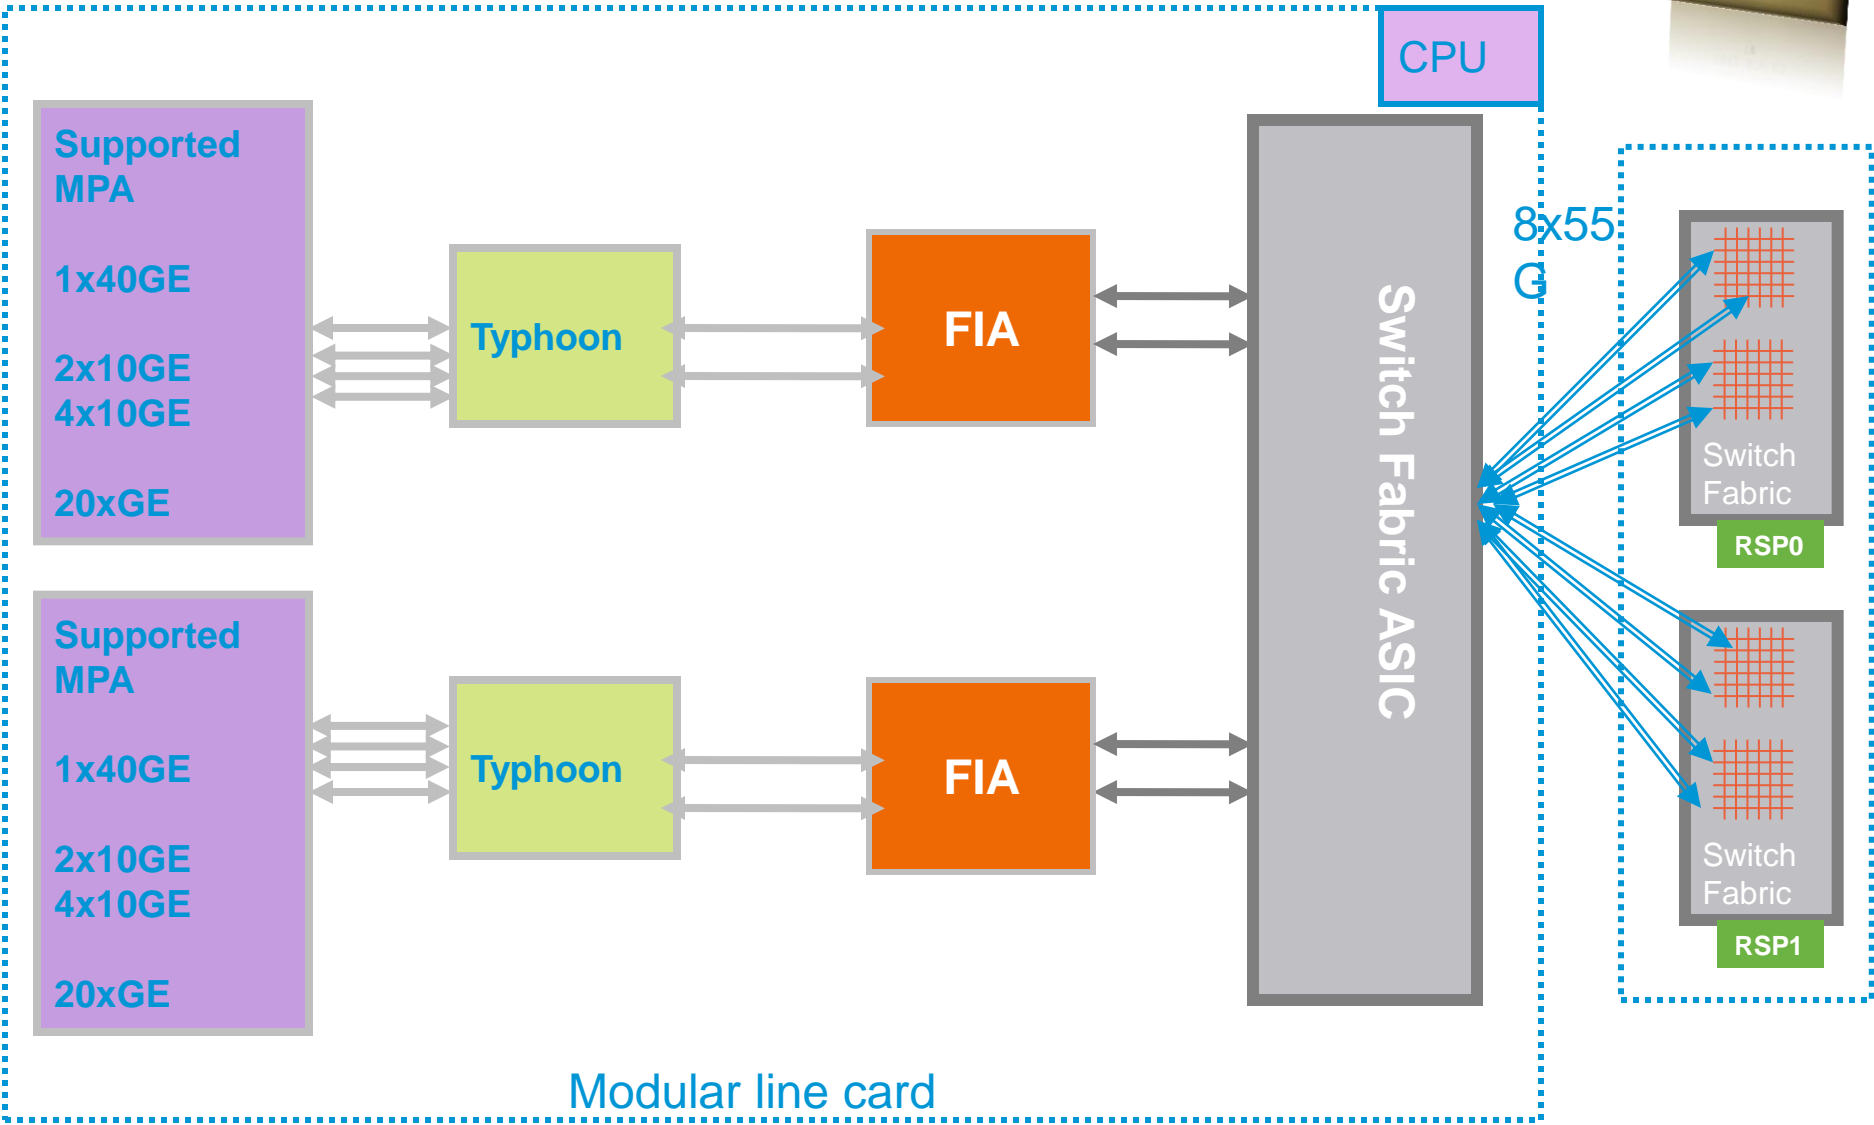


LC Architecture – Modular Ethernet MOD160



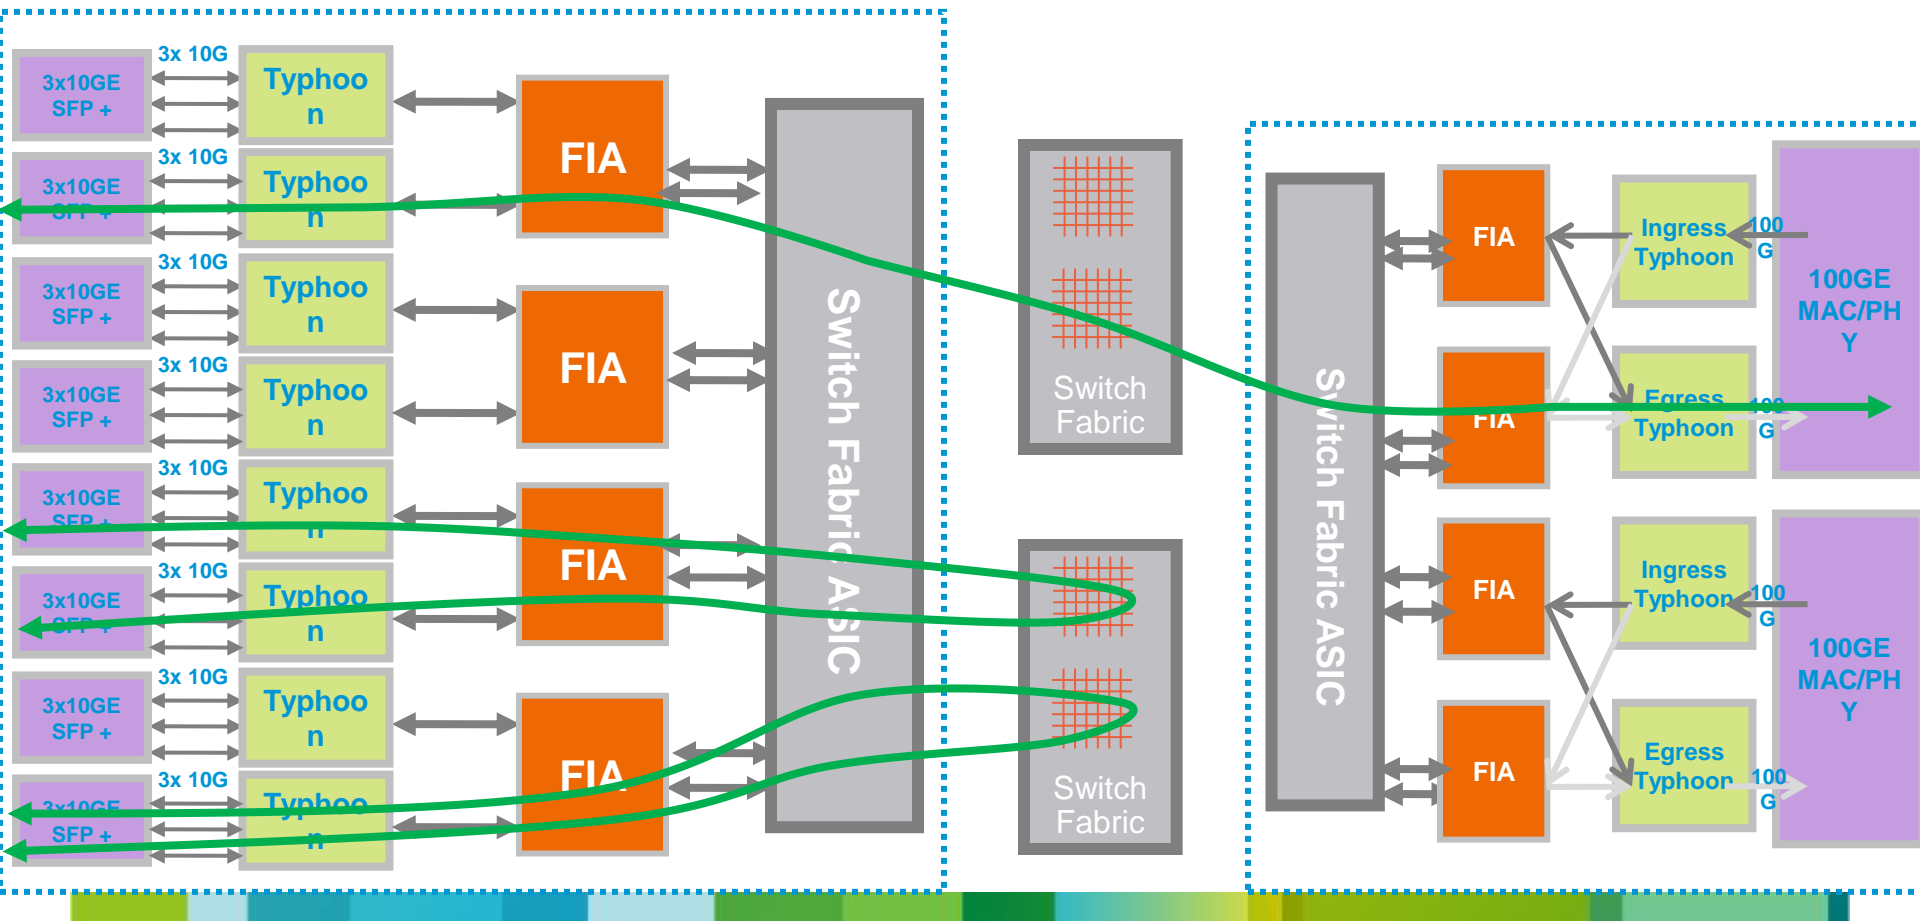


LC Architecture – Modular Ethernet MOD80

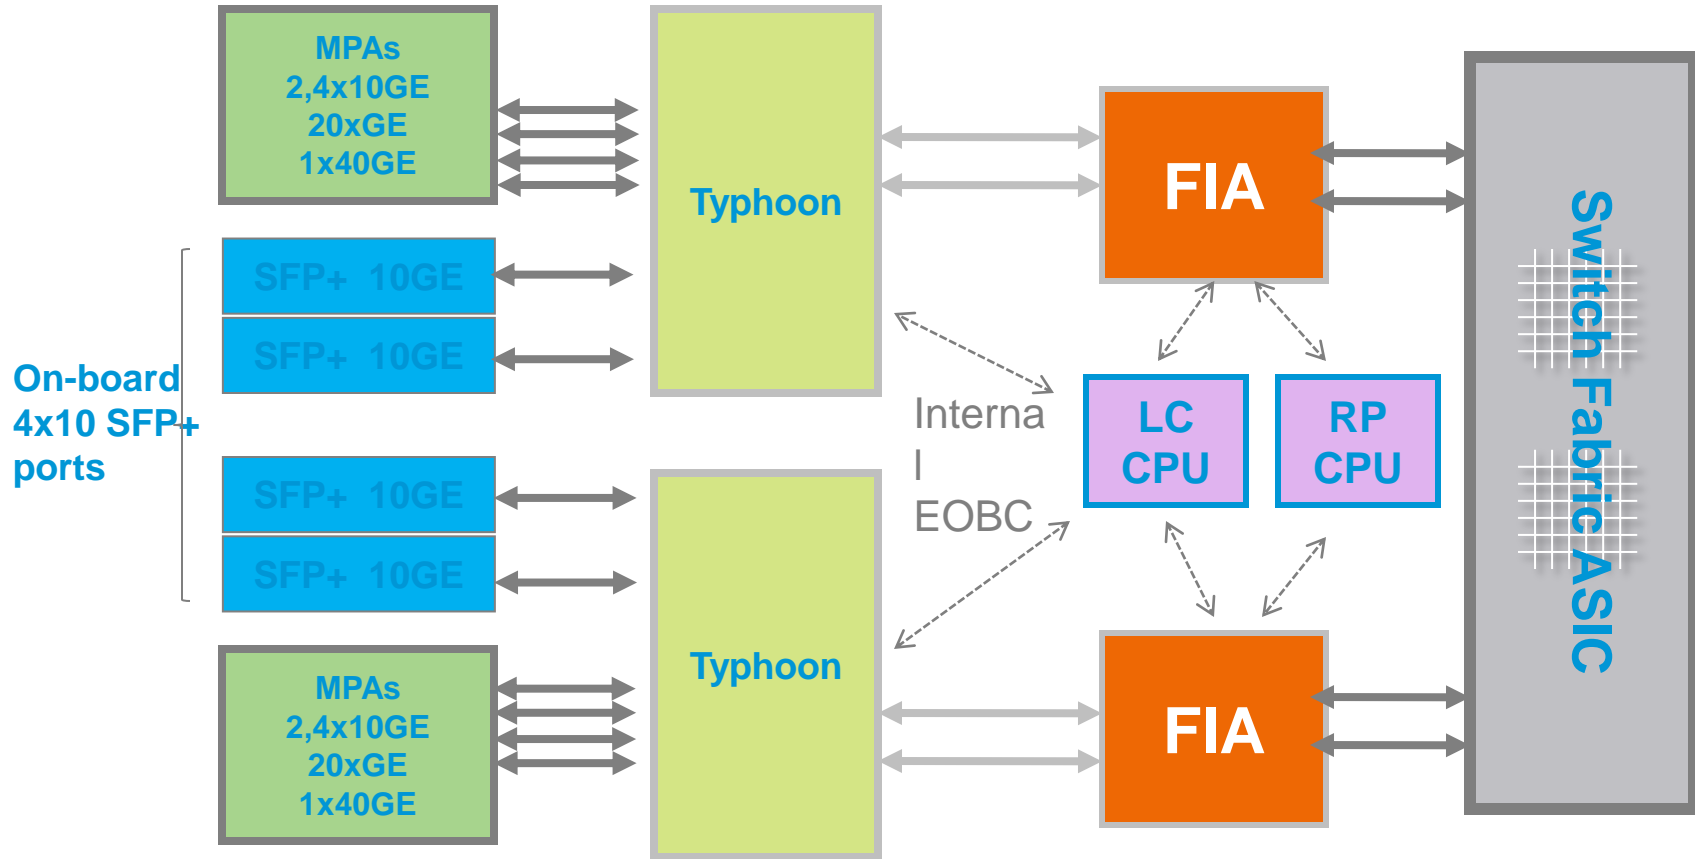


Packet Flow Overview

Same as existing system: Two-stage IOS-XR packet forwarding
Uniform packet flow: All packet go through central fabric on the RP



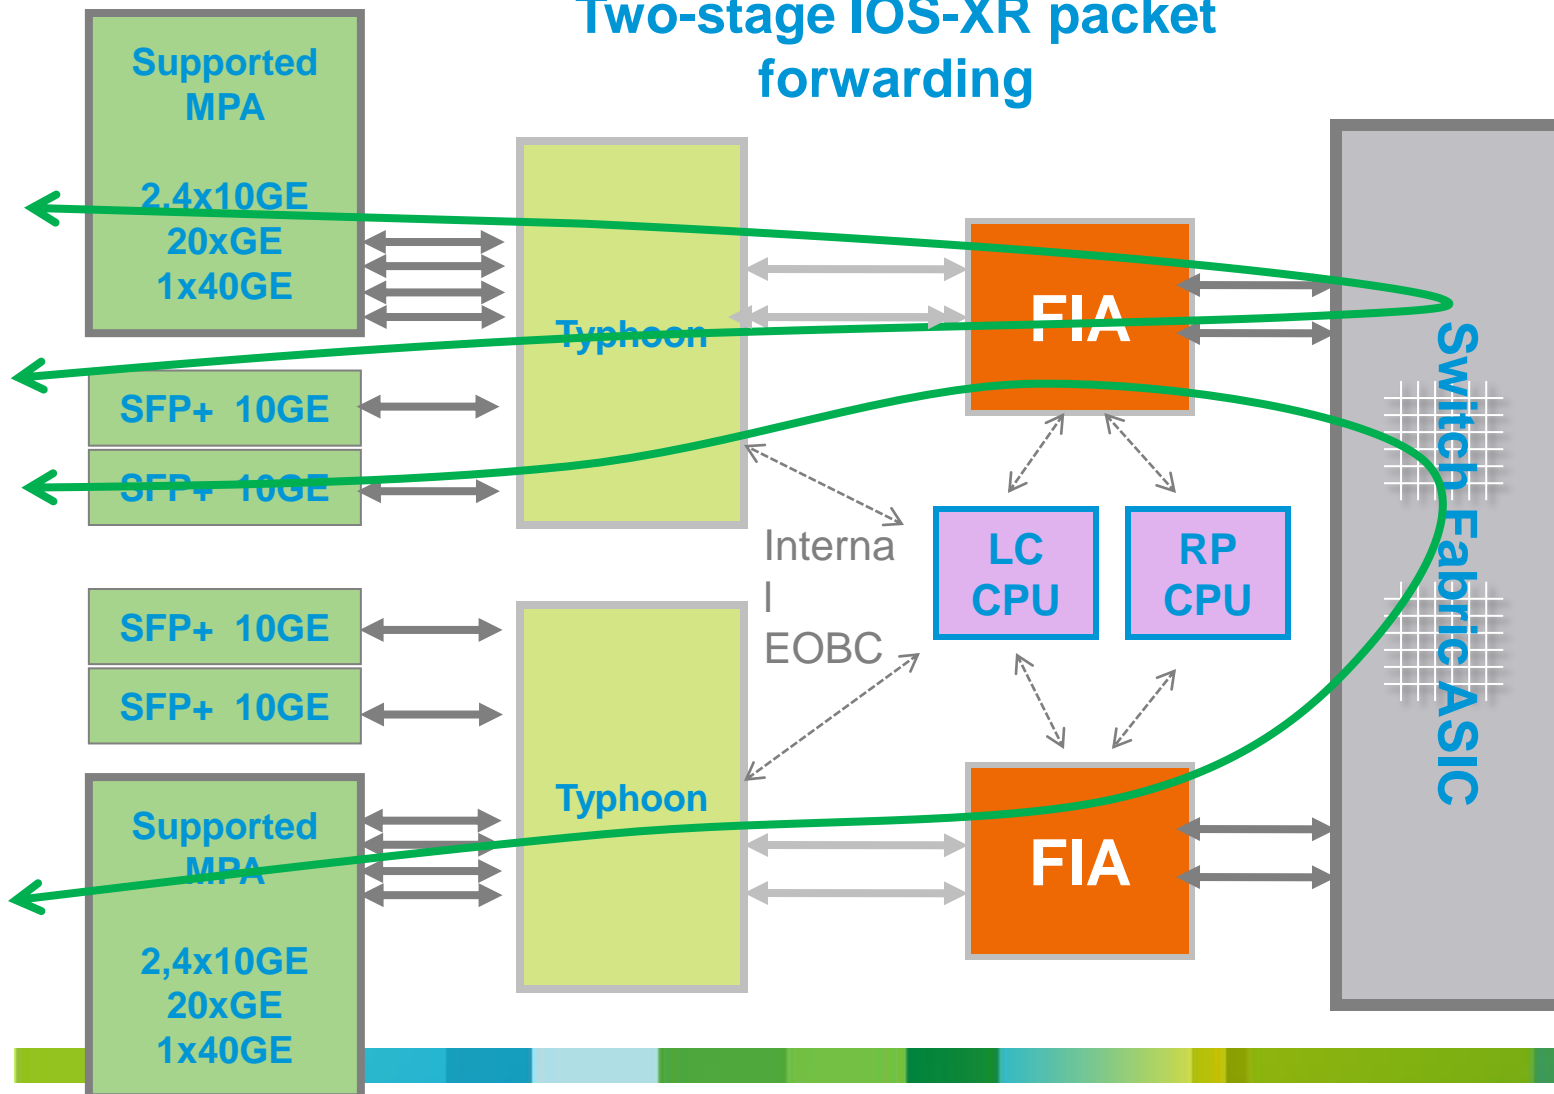
ASR 9001 System Architecture Overview



It has both central RP and LC CPU like big chassis
But it only have central switch fabric, no LC fabric
Maximum 120Gbps bi-directional system

ASR 9001 Packet Flow Overview

Same as big chassis system:
Two-stage IOS-XR packet forwarding




Line Card Memory Options

Queue Scale

- 3 memory options for each line card:
 - Extended (or high queue)
 - Base (medium queue)
 - Low (low queue)*
- Different memory options have different:
 - QoS queue scale
 - L2 sub-interface scale
- All other system wide scale parameters remain same:
 - FIB
 - MAC address
 - Bridge-domain
 - L3 sub-interface
 - VRF, etc
- All line cards have the same HW → Identical features
- Mixed LC types supported on same chassis

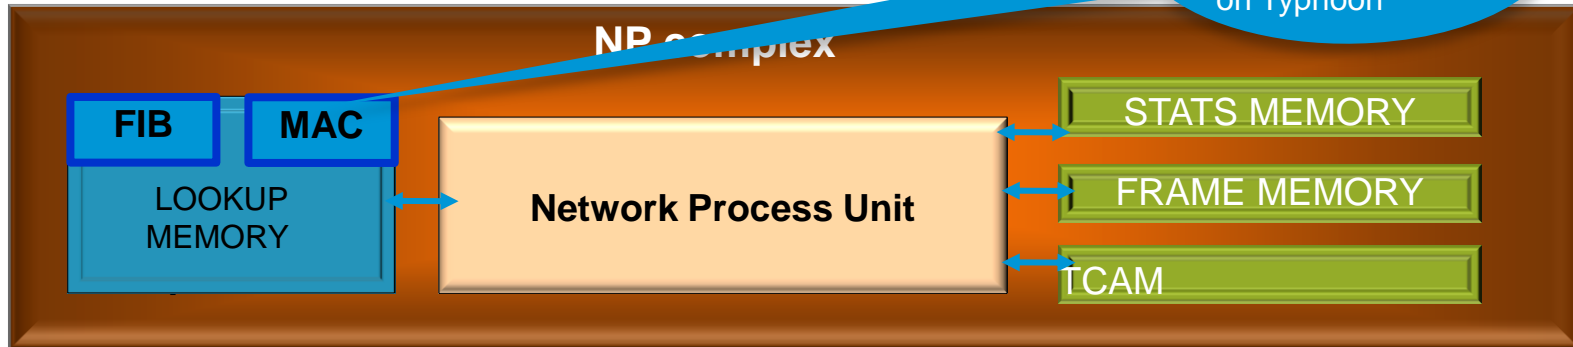
* A9K-16T/8 only have “B” option. It doesn’t have “E” or “L” option as of the 3.9.1 release



L/B/E Line Cards

What's the Difference?

Shared between L2/L3 on Trident.
Dedicated mem on Typhoon



- Each NPU has Four Main memories:
 - **Lookup/Search Memory (RLDRAM):** stores MAC, FIB, and Adjacencies Tables
 - **TCAM:** classification (Vlan Tag (EVCs), QoS and Security ACL
 - **Stats QDR memory:** interface and forwarding statistics, policers data, etc
 - **Frame memory:** buffer memory for Queues
- 3 LC versions – low, base and extended - differ for size of memories
 - TCAM, QDR and Frame memory sizes depend on LC version
 - Affects number of QoS queues and L2 sub-interfaces supported
 - Search Memory is same
 - System level scale (unicast, multicast, MPLS label) adjacency and MAC address) not affected by a mix of LCs

ASR 9000 Ethernet Linecards

Capability Comparison

Metric	Low Queue	Medium Queue	High Queue
MAC Addresses	512K	512K	512K
IPv4 Routes (total/per VRF)	1M/128K	1M/128K	1M/128K
ARP entries	32K	32K	32K
IPv6 Routes (total/per VRF)	512K/128K	512K/128K	512K/128K
VRFs	4k	4k	4k
L3 Subif/Port	4k	4k	4k
Bridge Domains	8k	8k	8k
MPLS Labels	128k	128k	128k
EFPs (L2 sub-int)	4k	16k	32k
Queues	8/port	64k/32k	256k/128k
Policers	8k	128k	256k

Common

Different

What's the Difference Between "-SE" and "-TR"?

Feature	-TR	-SE	Comments	
FIB (V4+V6)		4M	V4 and V6 share the same table V6 uses two FIB entries Support per-VRF FIB table download per LC (SVD)	System wide scale
Multicast FIB		128K		
MAC		2M	Support per-LC MAC learning in the future	
L3 VRF		4K	8K in 4.2.1	
BD/VFI		64K		
PW		128K		
L3 interface	8K/LC	20K/LC		
L2 interface	16K/LC	64K/LC		
QoS	8 queues/port (I and O) 8K policers/NP 1G frame memory/NP	256K queues (I+O) / NP 256K policers/NP 2G frame memory/NP		
ACL	24k ACE	96k ACE	Note XR431 has ACL compression (hybrid)	

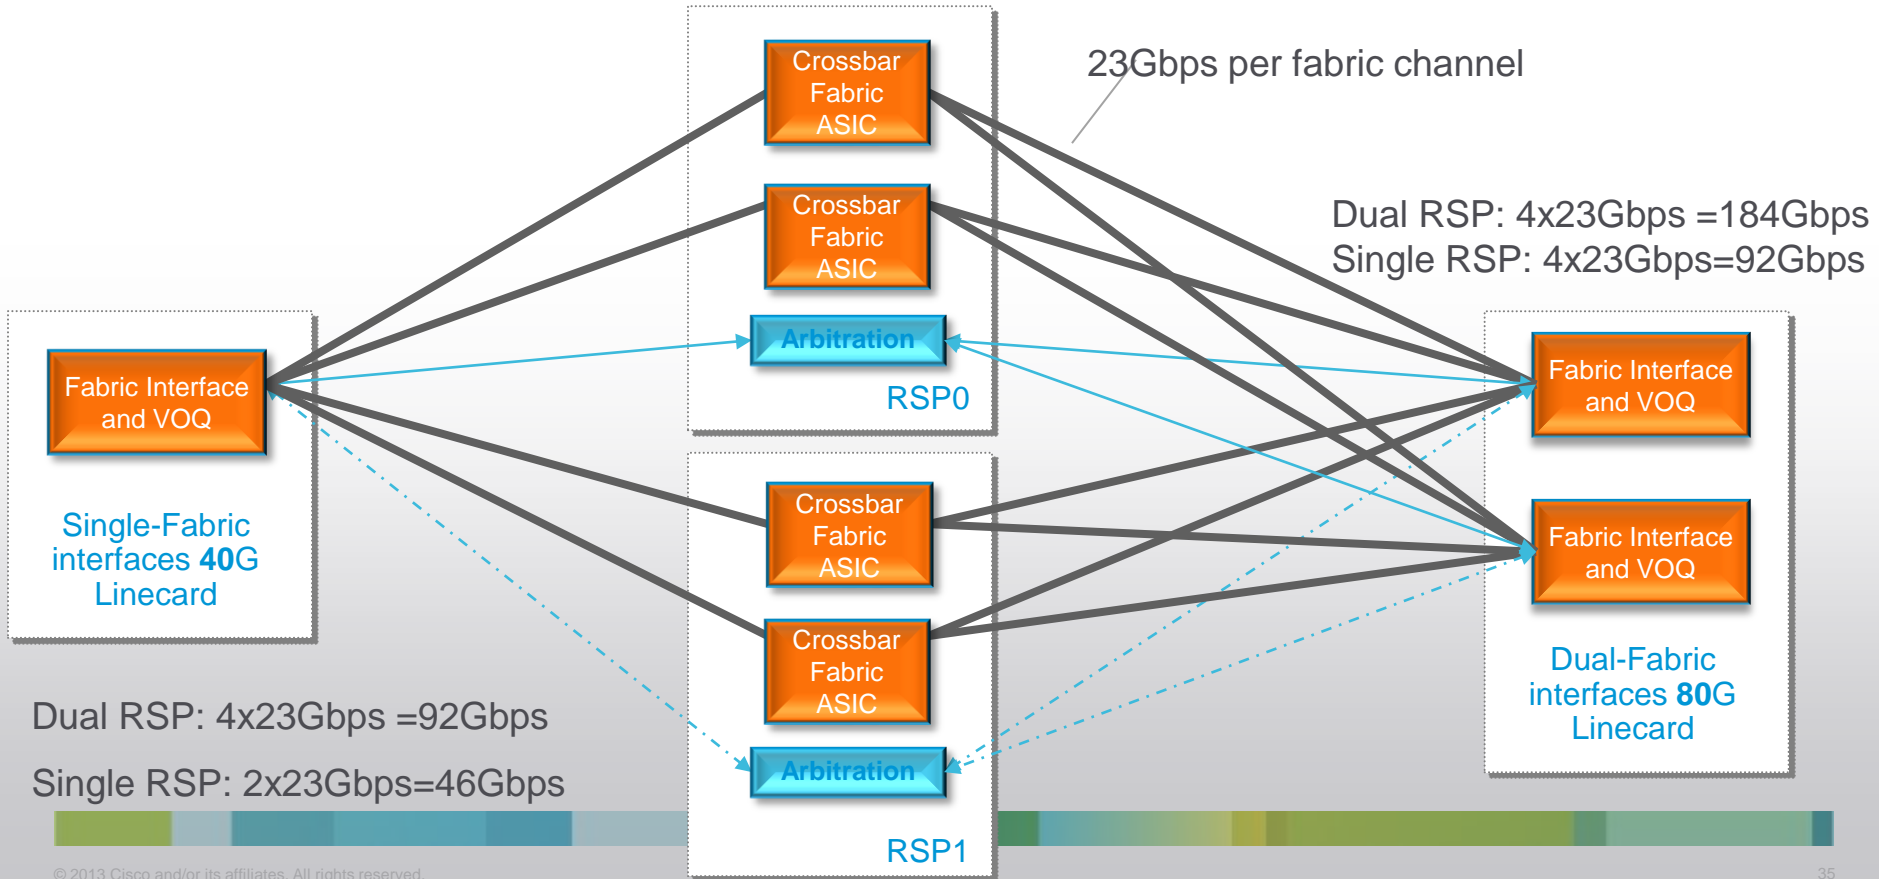
Polling Question 2

How are you using the ASR9000?

- a) Mainly in an L2 PE environment
- b) As a replacement for a 7600
- c) Mainly as a L3 PE
- d) A core router
- e) Anything it can do

Fabric Overview

- Centralized fabric resides on RSP
- Logically separated from LC and RSP
 - All fabric ASICs run in active mode regardless of RSP Redundancy status
 - Extra fabric bandwidth and instant fabric switch over
- Each 40G LC/RSP has one fabric interface ASIC
- 80G line rate LCs have 2 fabric interface ASICs



Fabric Arbitration and Redundancy

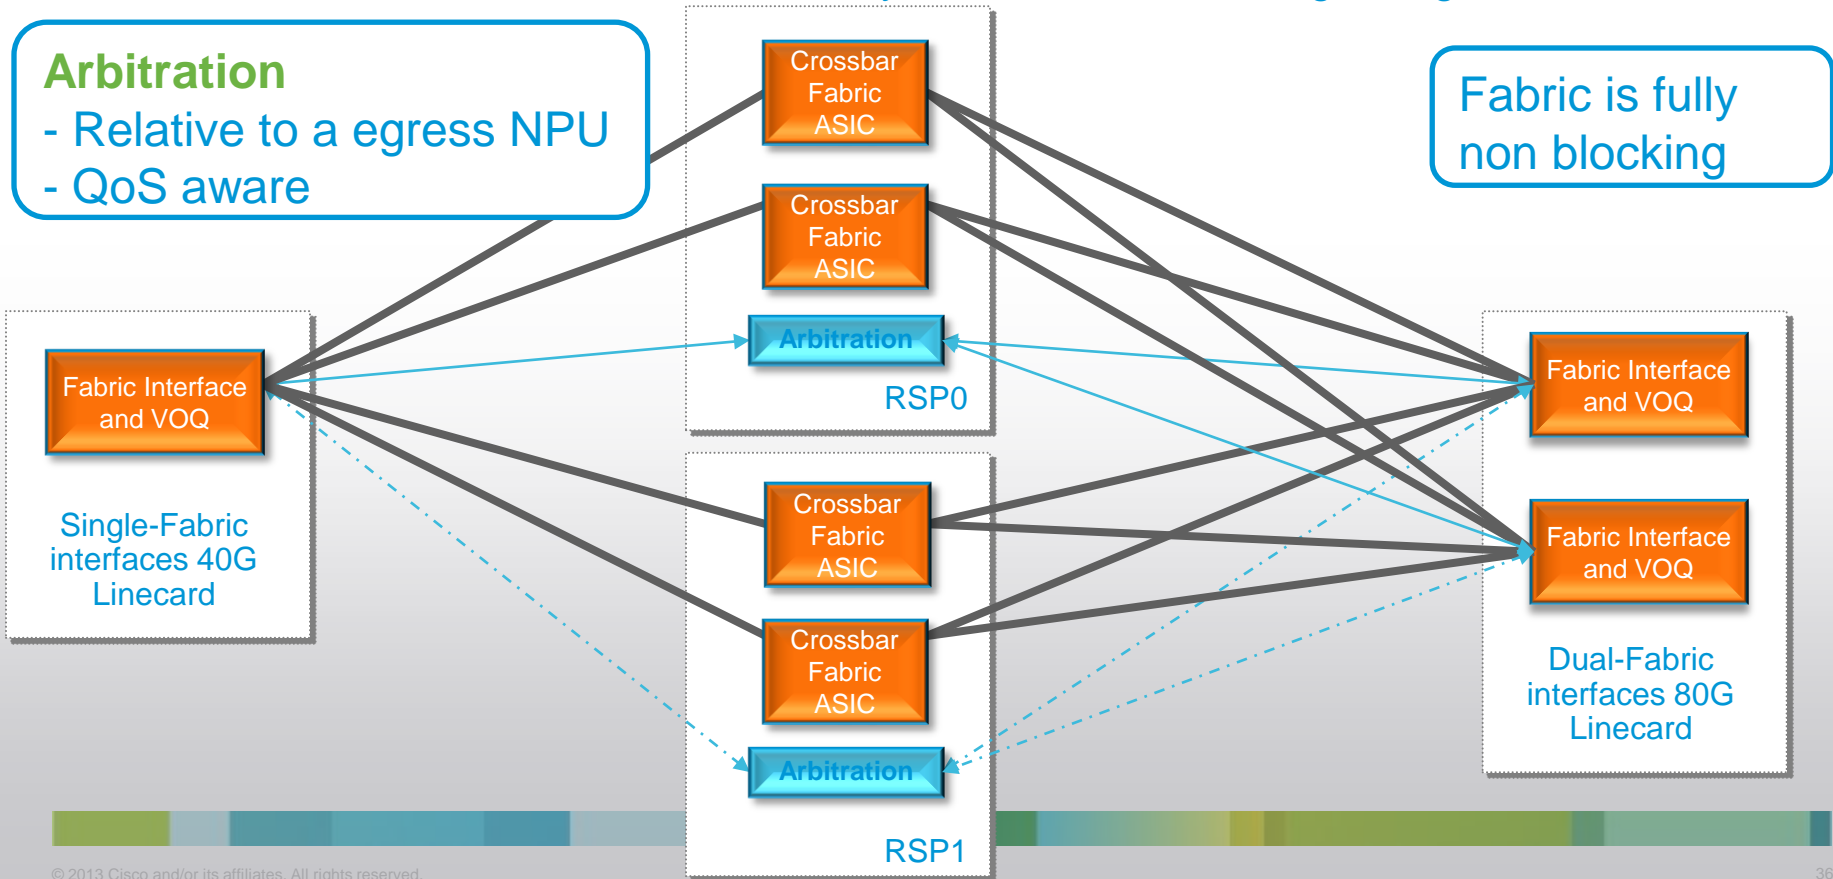
“0” packet loss guarantee during RSP failover and OIR

- Access to fabric controlled using **central arbitration**.
 - One Arbitration ASIC (Arbiter) per RSP
 - Both Arbiters work in parallel – both answer to requests to transmit
 - FIAs follow active Arbiter, and switch to backup if needed
 - Arbiter switchover controlled by low level hardware signalling

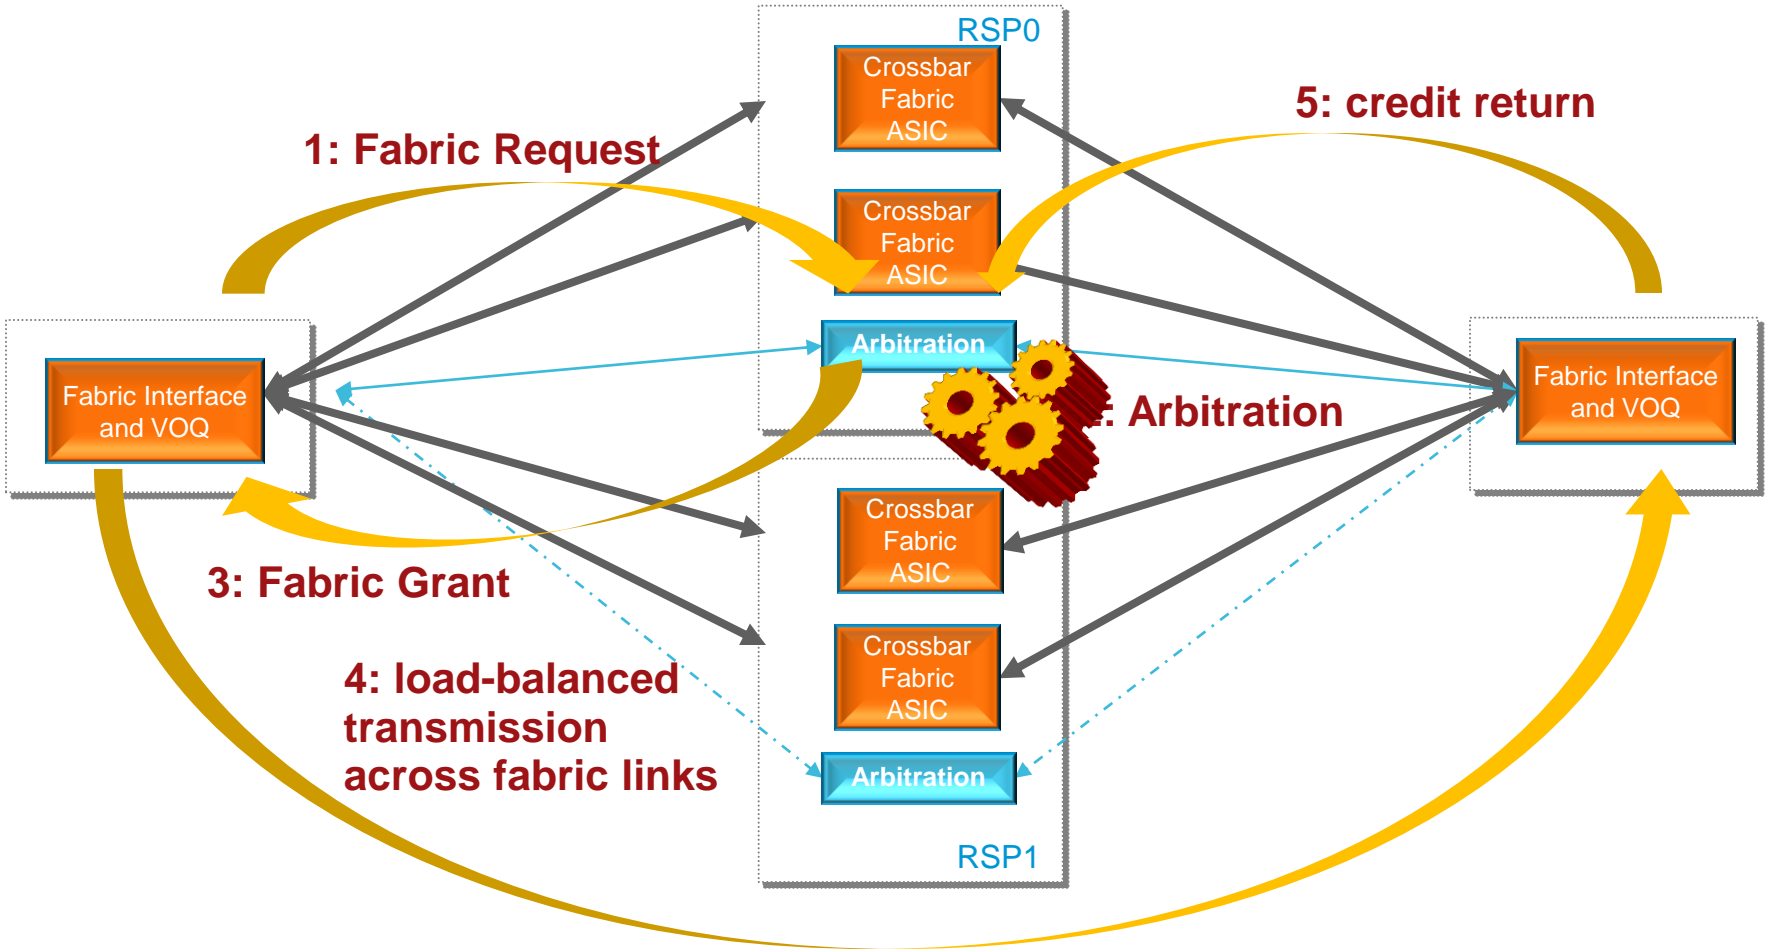
Arbitration

- Relative to a egress NPU
- QoS aware

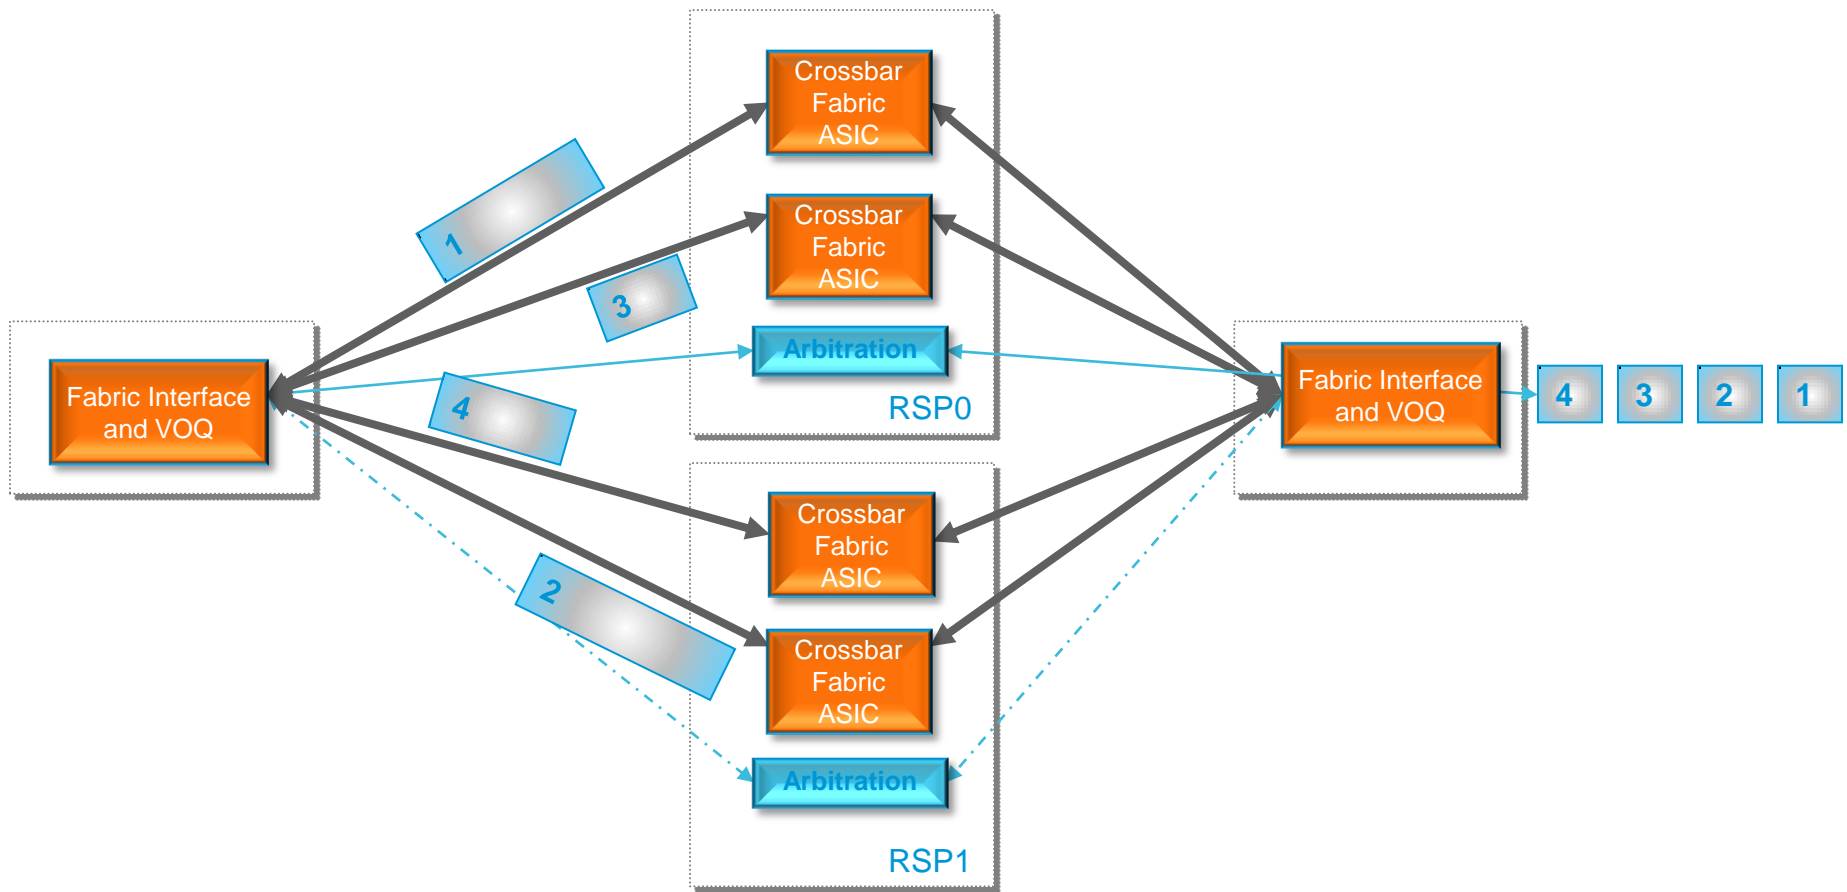
Fabric is fully non blocking



Fabric Arbitration

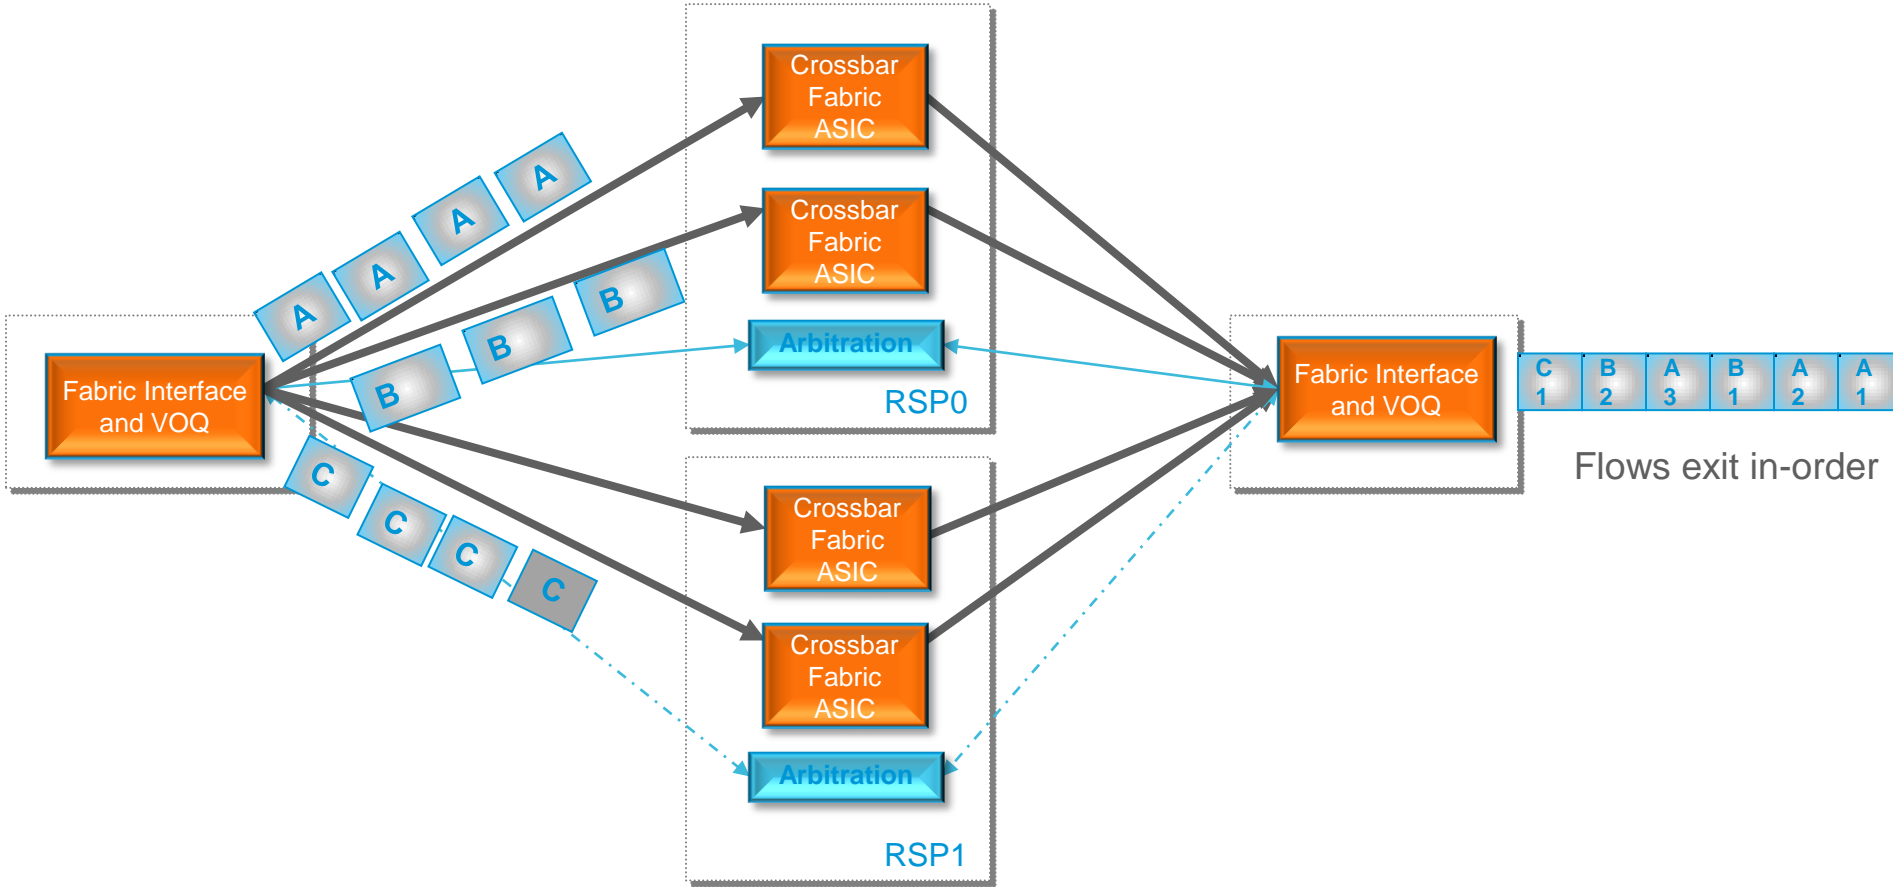


Fabric Load Sharing – Unicast



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

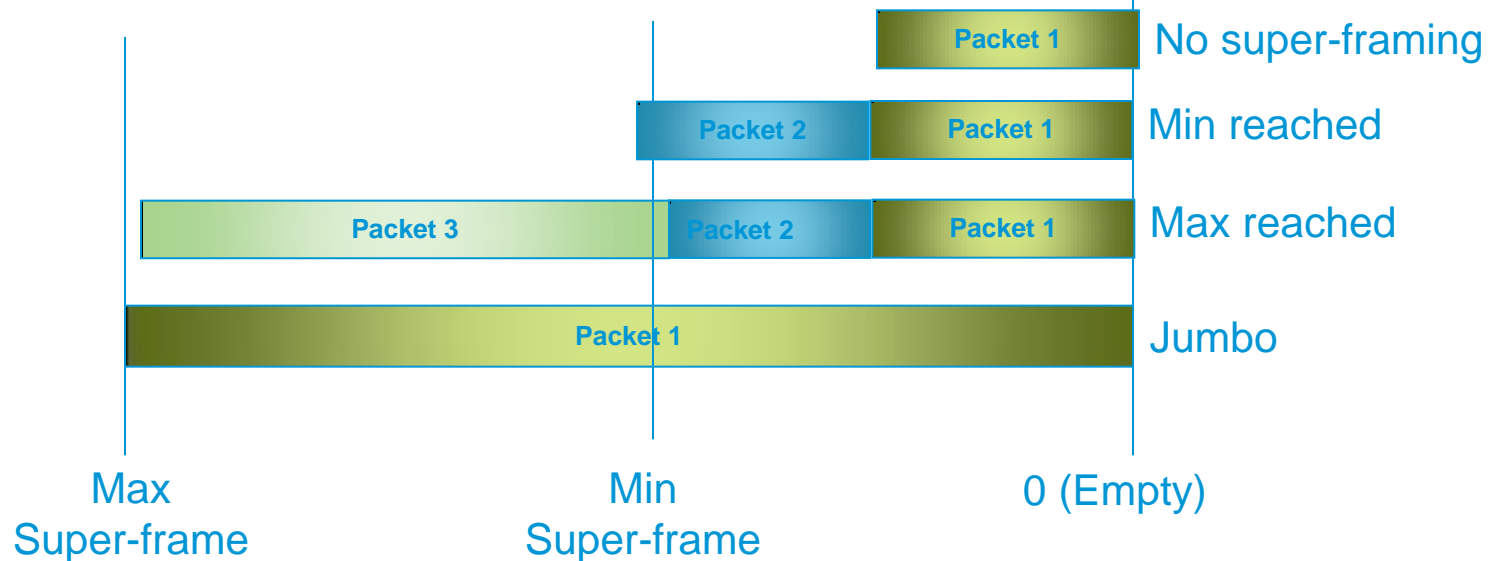
Fabric Load Sharing – Multicast



- 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

Fabric Super-framing Mechanism

- Multiple unicast frames from/to same destinations aggregated into one super frame
- Super frame is created if there are frames waiting in the queue, up to 32 frames or when min threshold met, can be aggregated into one super frame
- Super frame only apply to unicast, not multicast
- Super-framing significantly improves total fabric throughput

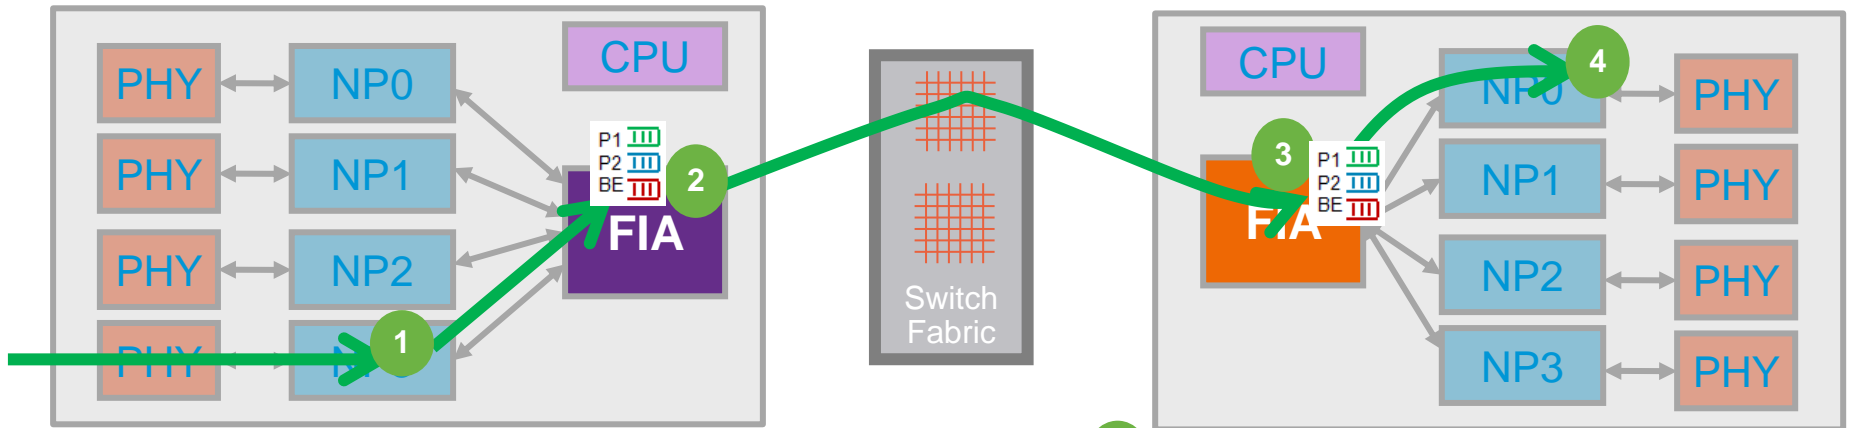


End-to-End System QoS Queuing

End-to-End priority (P1,P2, Best-effort) propagation →
 Guarantee bandwidth, low latency for high priority traffic
 at any congestion point
 3 strict priority level across all internal HW components

Ingress side of LC

Egress side of LC



1
 Ingress (sub-)interface
 QoS Queues

2
 Virtual
 Output
 Queues

3
 Egress FIA
 Queues

4
 Egress (sub-)interface
 QoS Queues

Implicit Configuration
 Two strict high priority +
 Normal priority

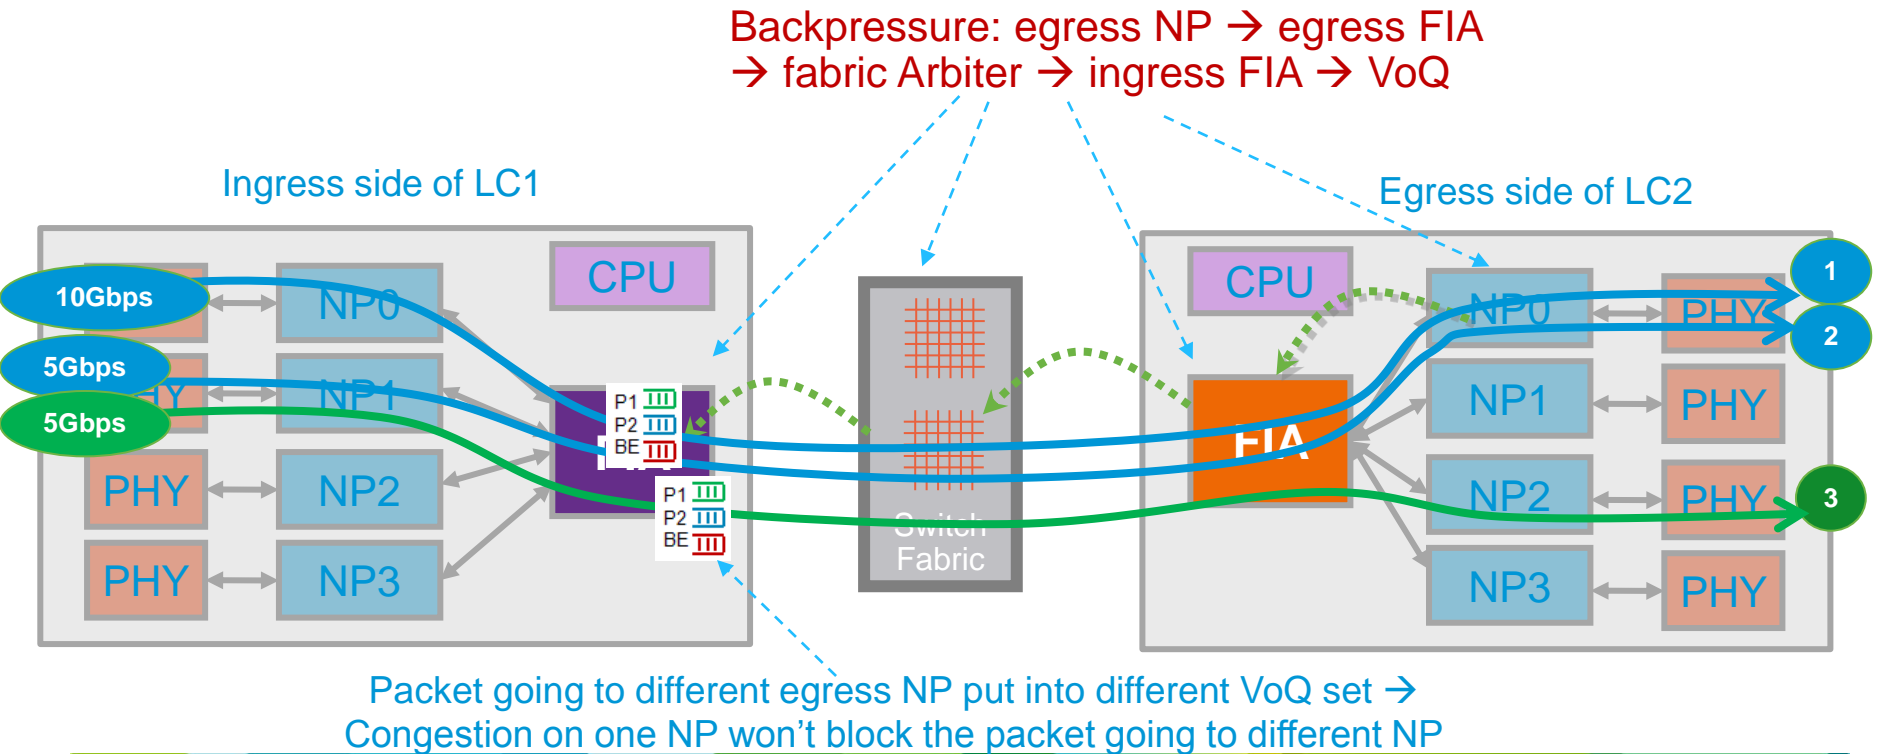
Configure with
 Ingress MQC
 4-layer hierarchy
 Two strict high priority +
 Normal priority

Configure with
 Egress MQC
 4-layer hierarchy
 Two strict high priority +
 Normal priority

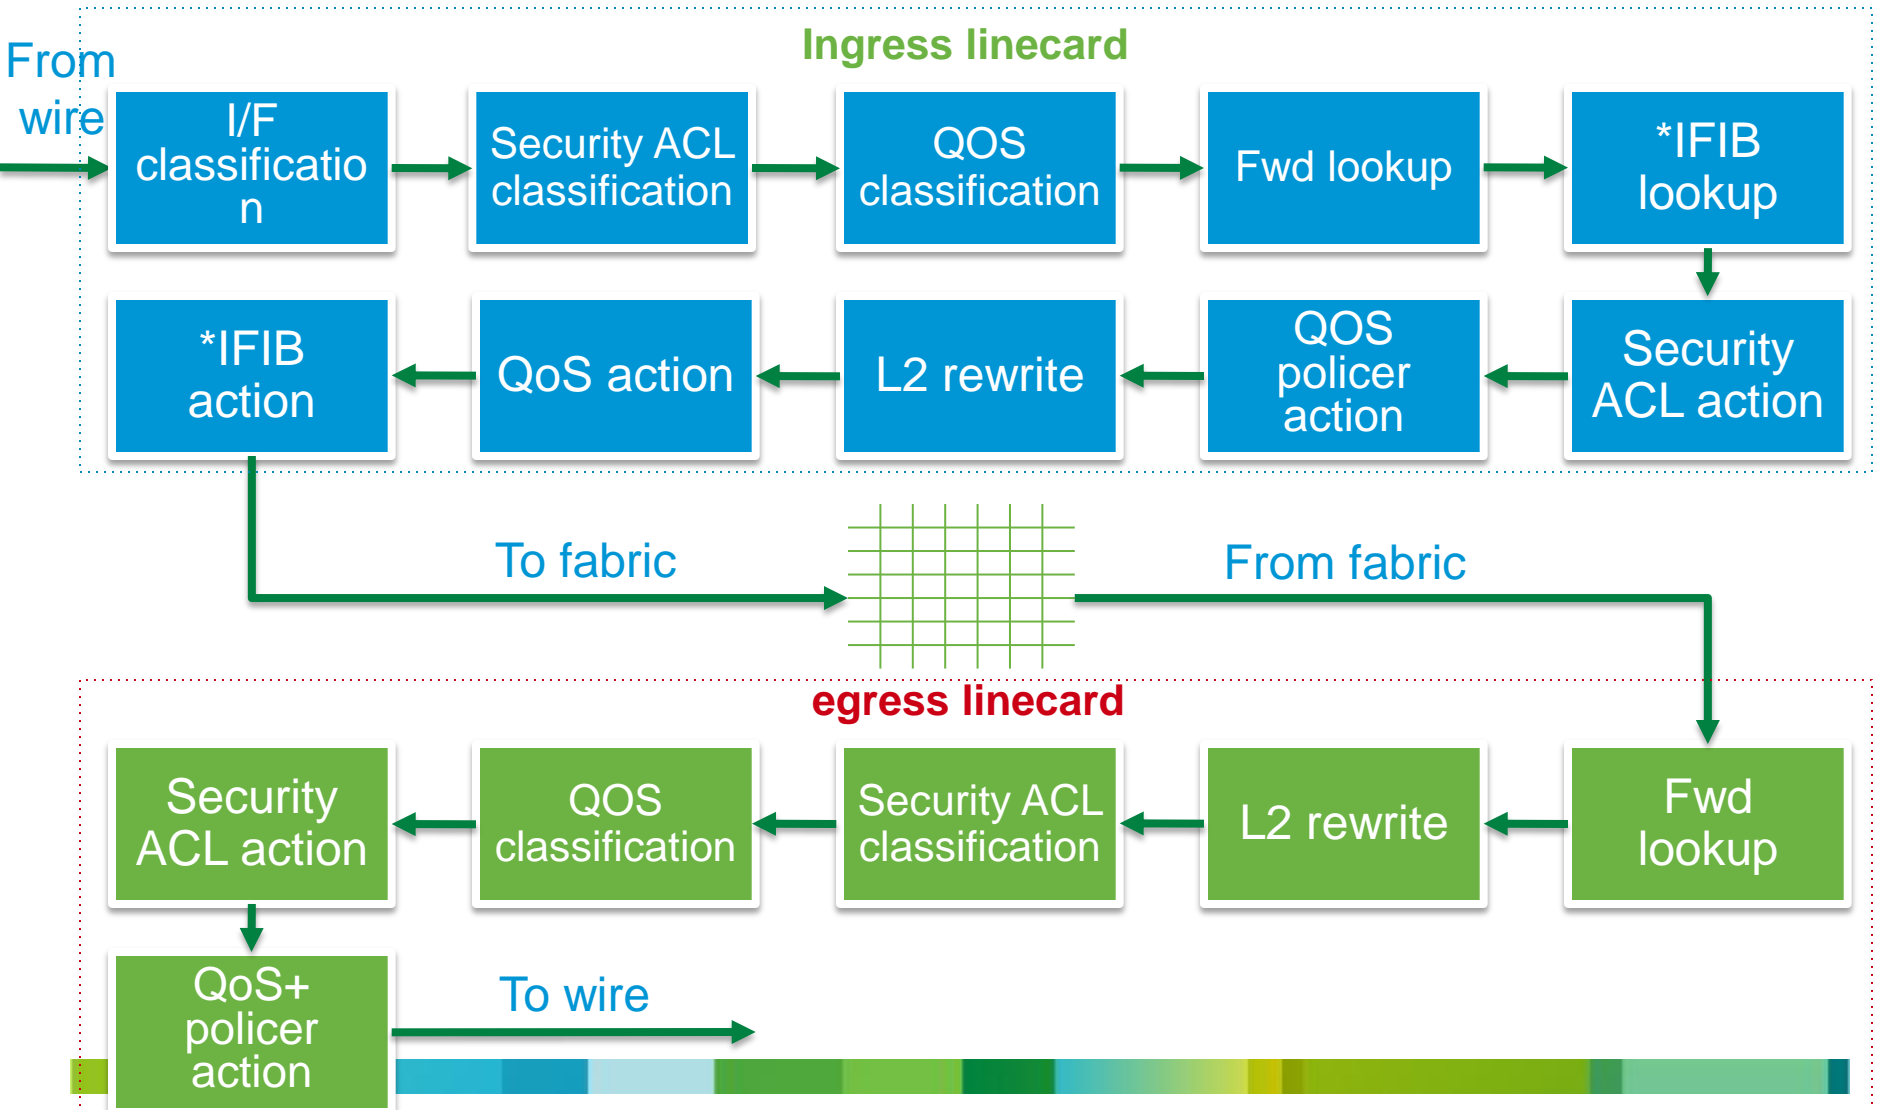
Backpressure and VoQ Mechanism

VoQ Scale: Each FIA has P1/P2/BE queue set for every NP and RSPs in the entire system

Egress NP congestion → → backpressure to ingress FIA →
Packet is en-queued in the dedicated VoQ →
No impact of the packet going to different egress NP →
No head-of-line-block issue

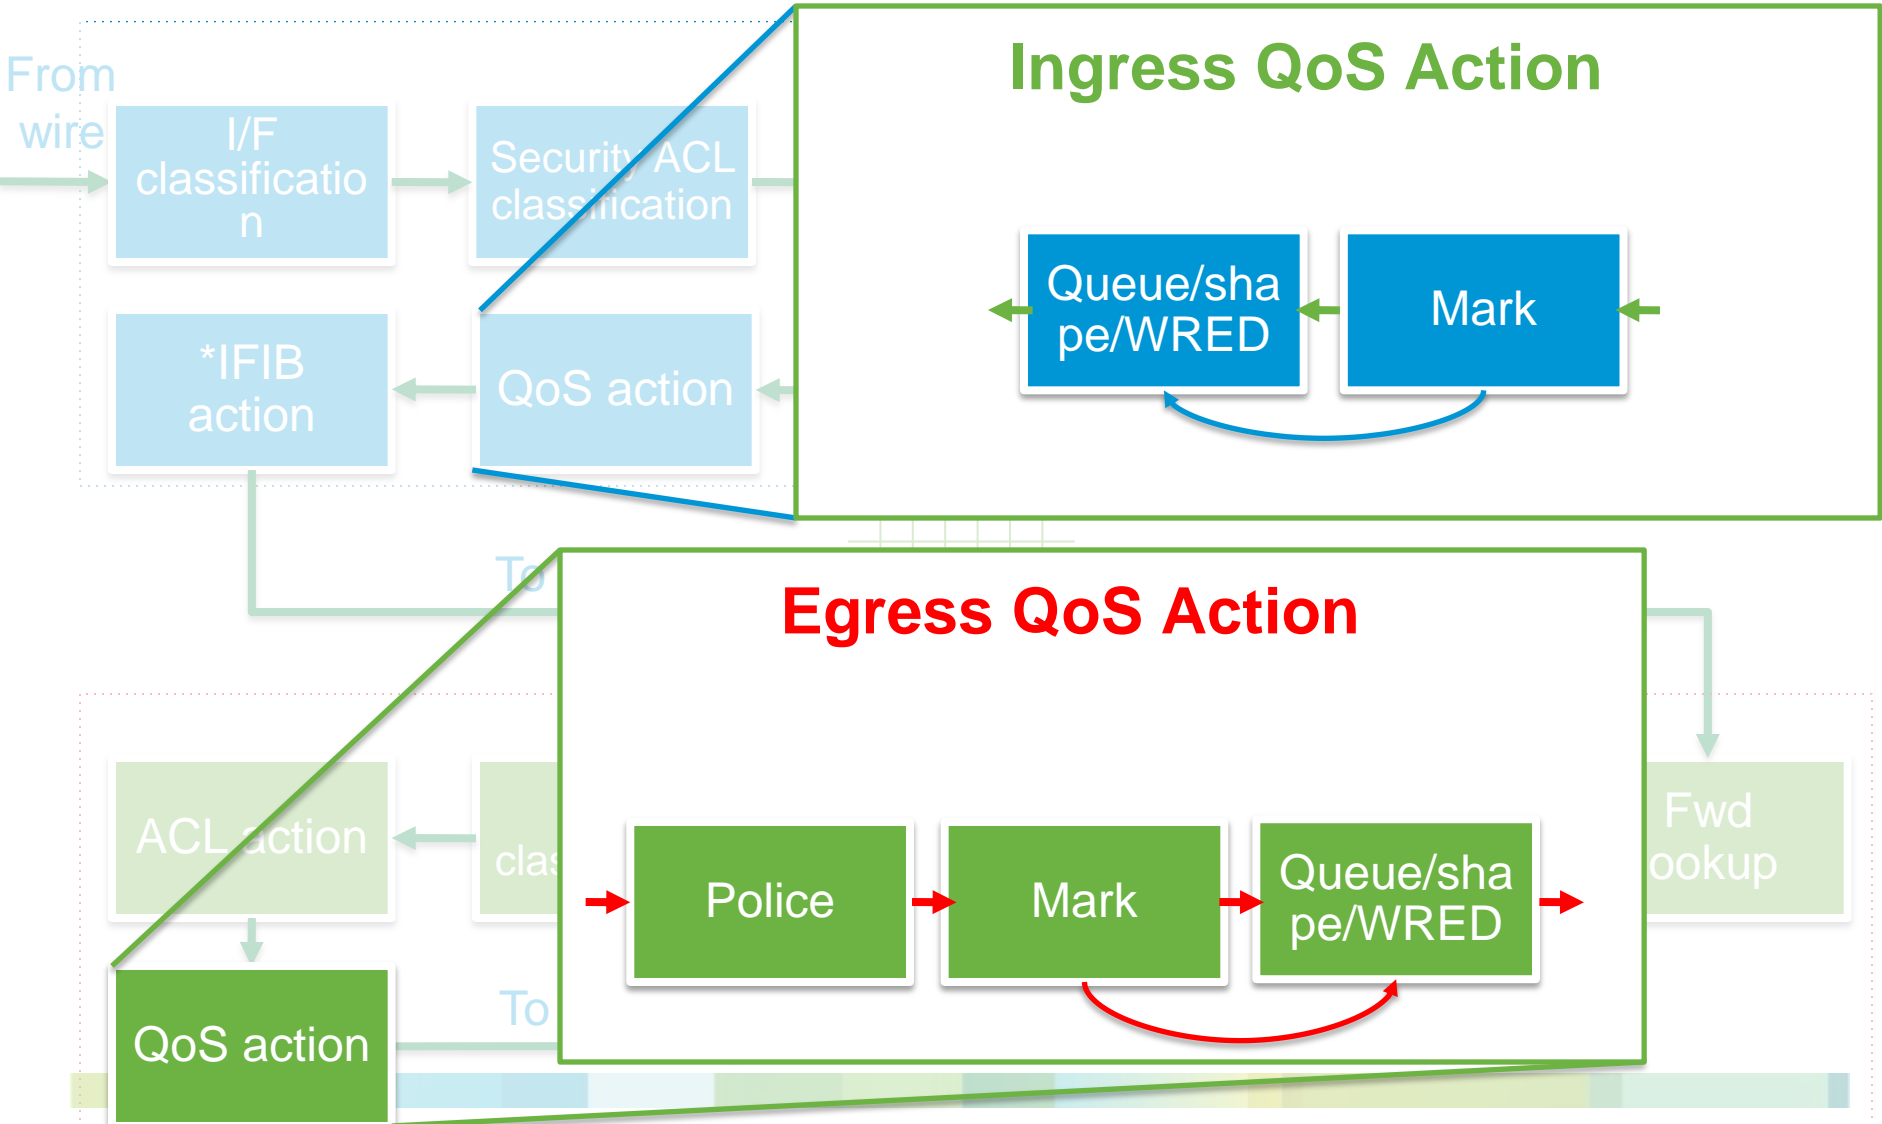


Feature order on ASR 9000 NPU (simplified)



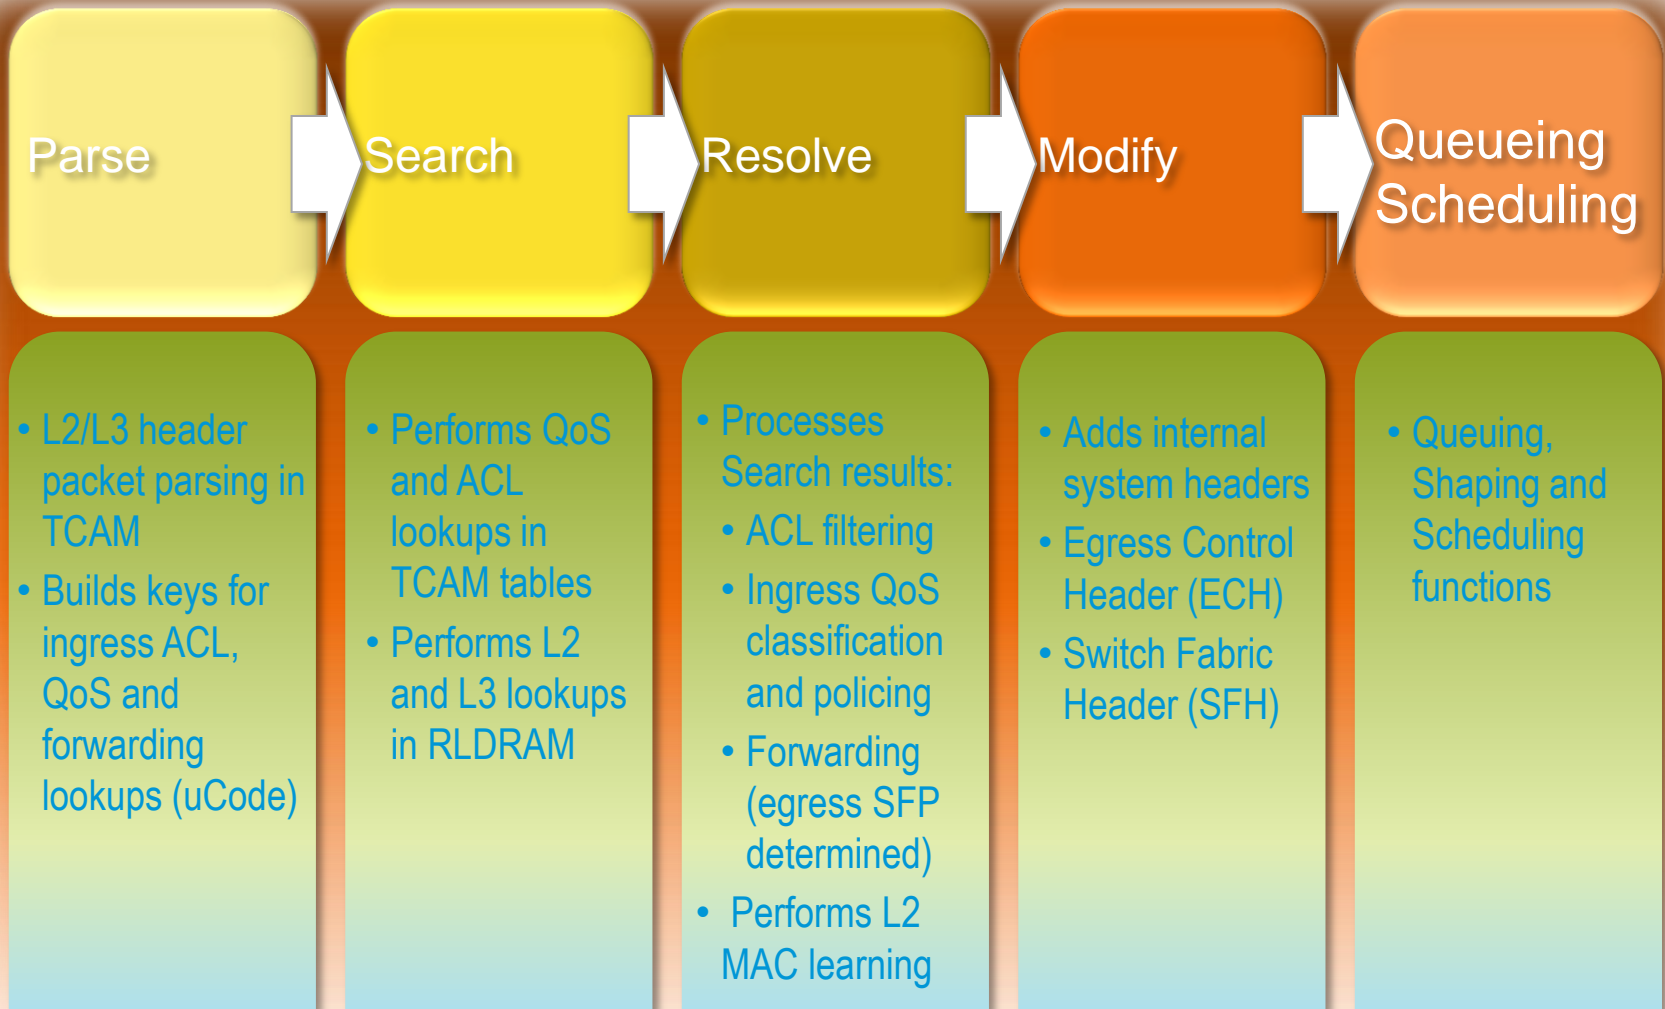
Feature order on ASR 9000 NPU

QoS Action Order



NPU Packet Processing - Ingress

5 Stages:



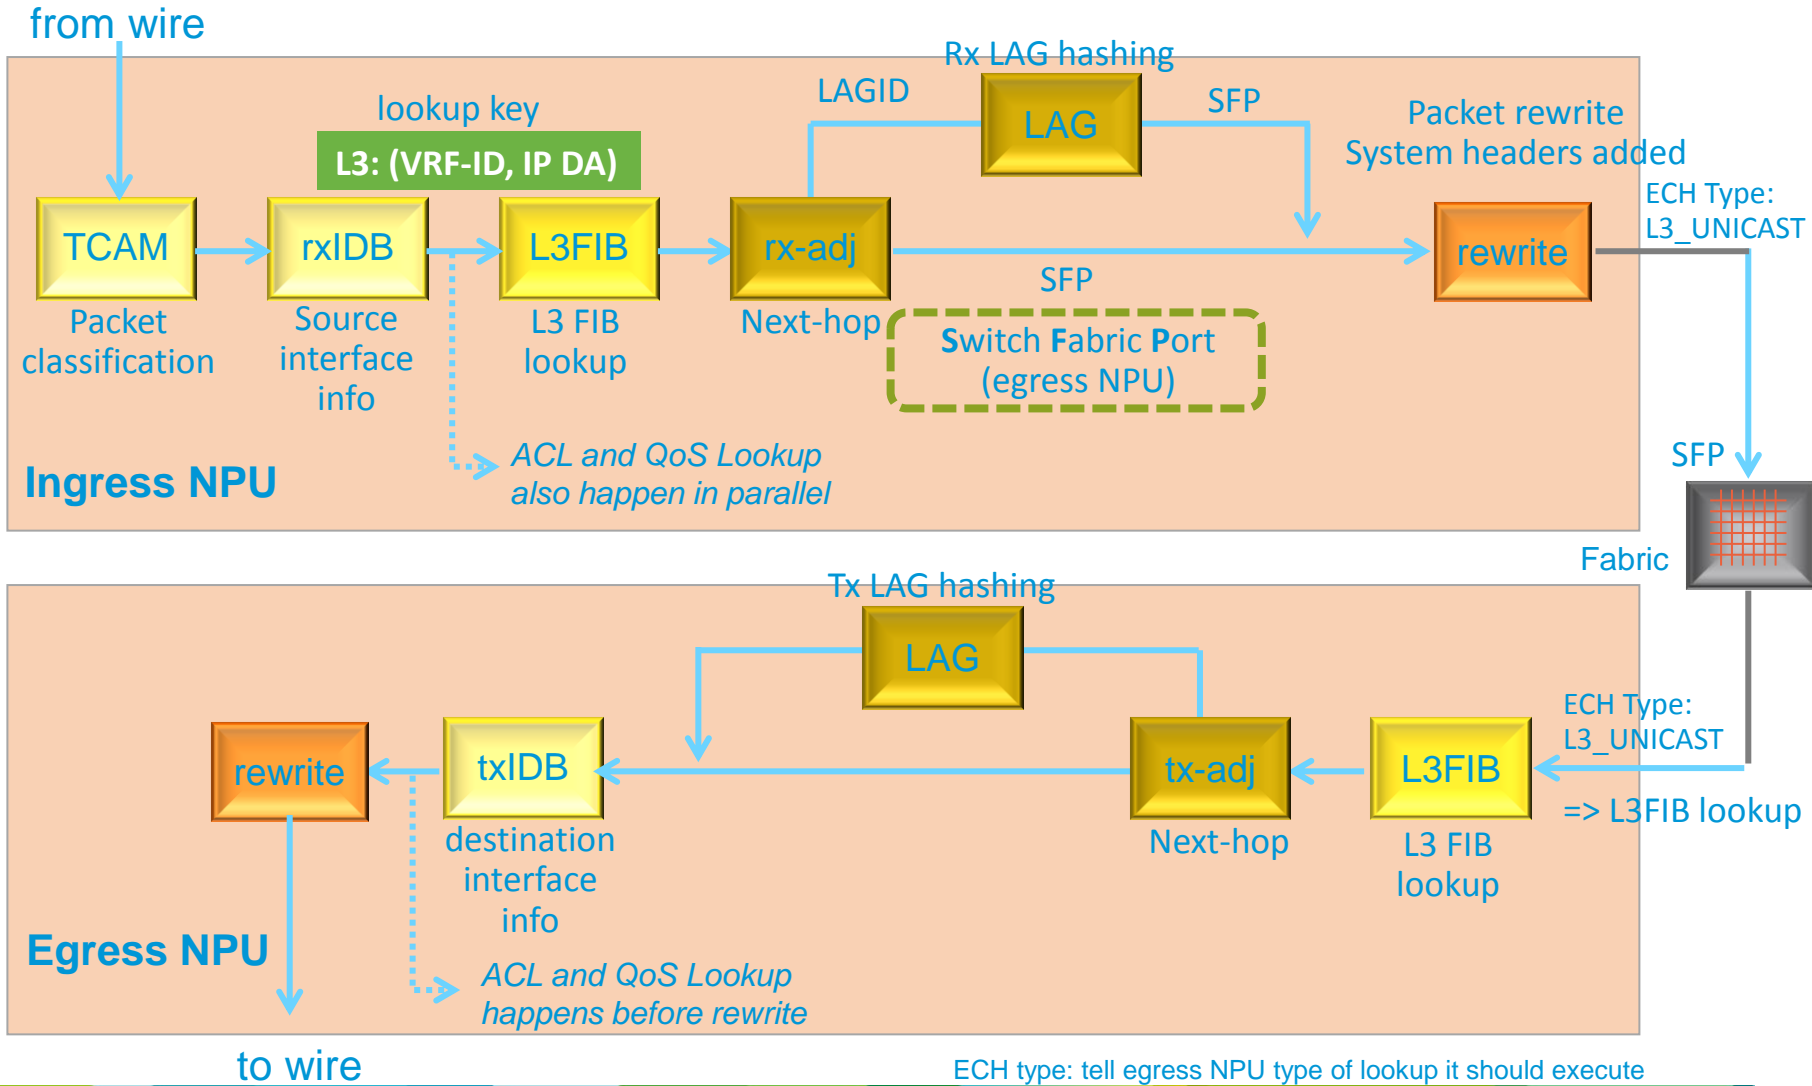
NPU Packet Processing - Egress

5 Stages:



L3 Unicast Forwarding

Packet Flow (Simplified)



Switch Fabric Port (SFP)

Getting SFP for a port

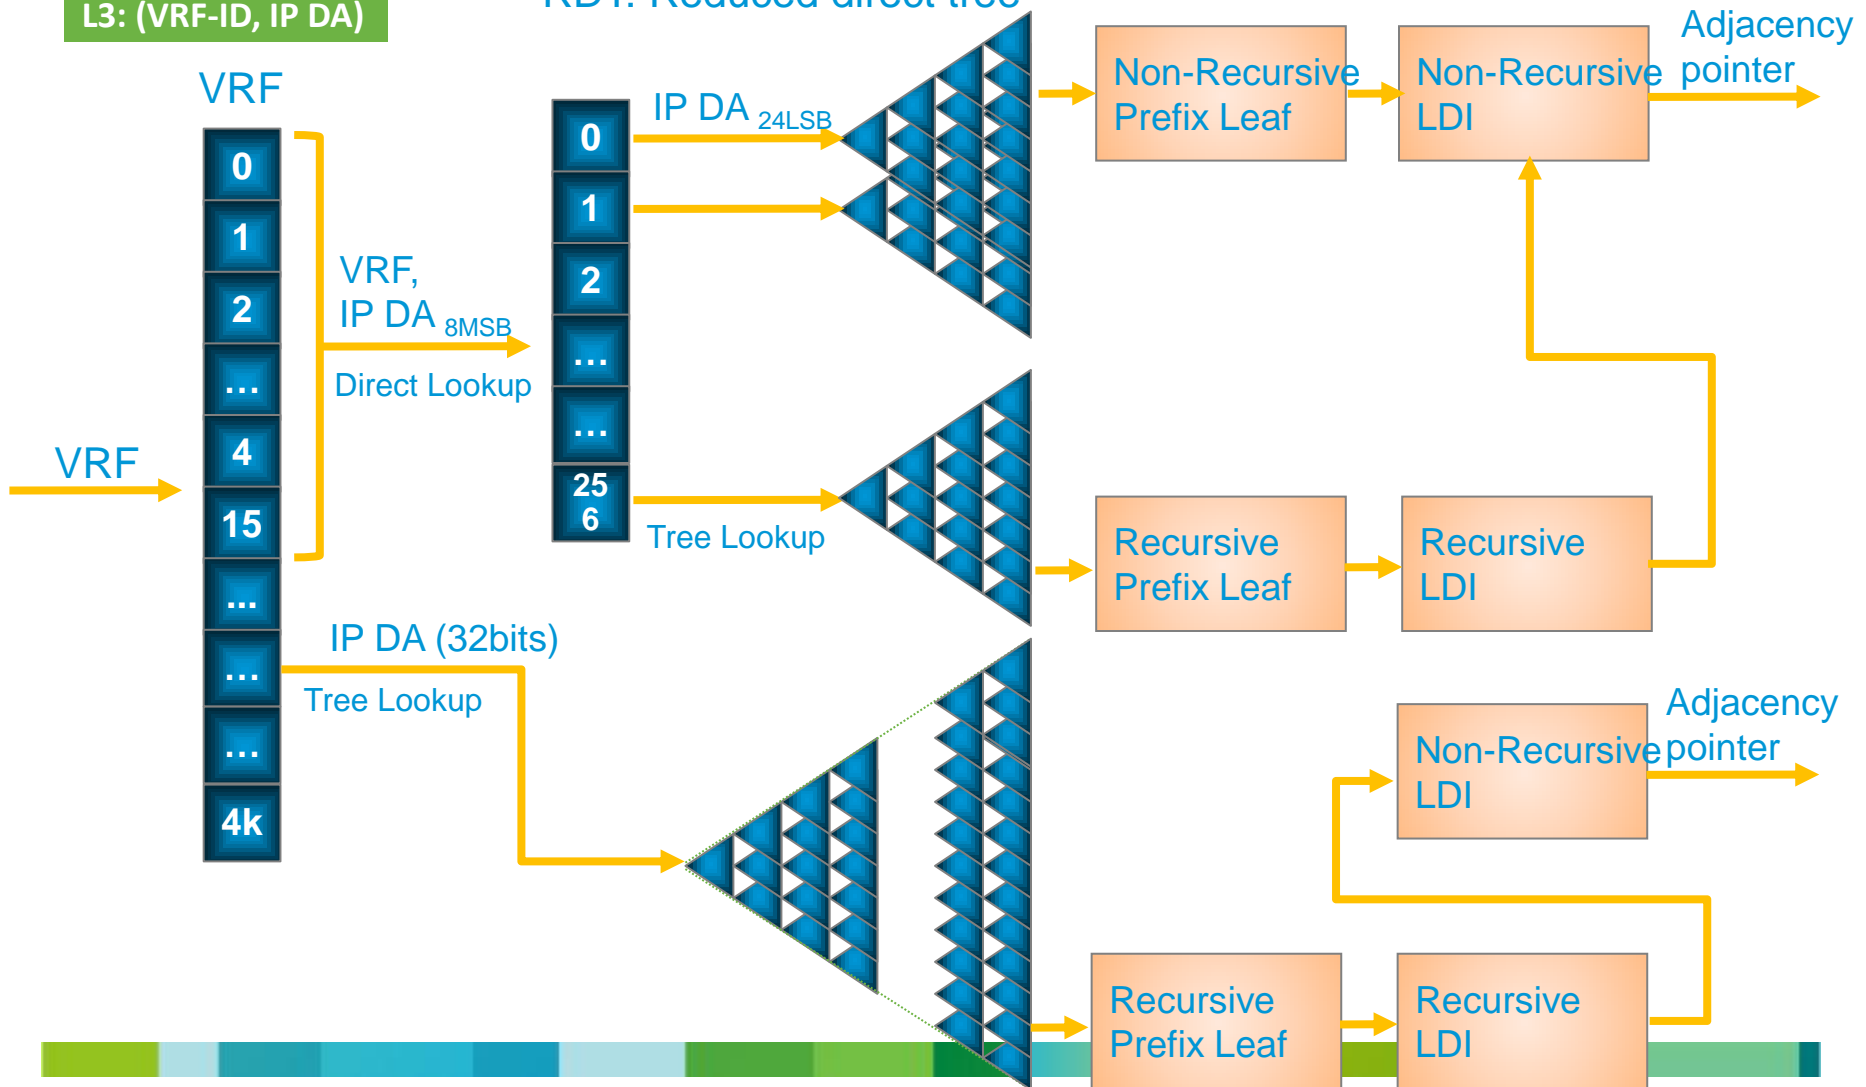
Shows uIDB information

```
RP/0/RSP1/CPU0:asr#sh controllers pm interface gig 0/0/0/1 loc 0/0/CPU0
Tue Aug  3 13:20:19.853 UTC
Ifname(1): GigabitEthernet0_0_0_1, ifh: 0x40000c0 :
iftype                0xf
egress_uidb_index     0x3
ingress_uidb_index    0x3
port_num              0x1
phy_port_num          0x1
channel_id            0x3
lag_id                0x0
virtual_port_id       0x0
switch_fabric_port 0x3
in_tm_qid_fid0        0x10000202
in_tm_qid_fid1        0xffffffff
in_qos_drop_base     0x10600
out_tm_qid_fid0       0x82
out_tm_qid_fid1      0xffffffff
out_qos_drop_base    0x38480
bandwidth             1000000 kbps
ing_stats_ptrs        0x0, 0x0, 0x0, 0xd8414
egr_stats_ptrs        0x0, 0x0, 0x0, 0xd8423
l2_transport          0x0
ac_count              0x0
parent_ifh            0x0
parent_bundle_ifh    0x0
```

Ports connected to the same NPU share the same SFP value

L3 NP3 IPv4 FIB Architecture

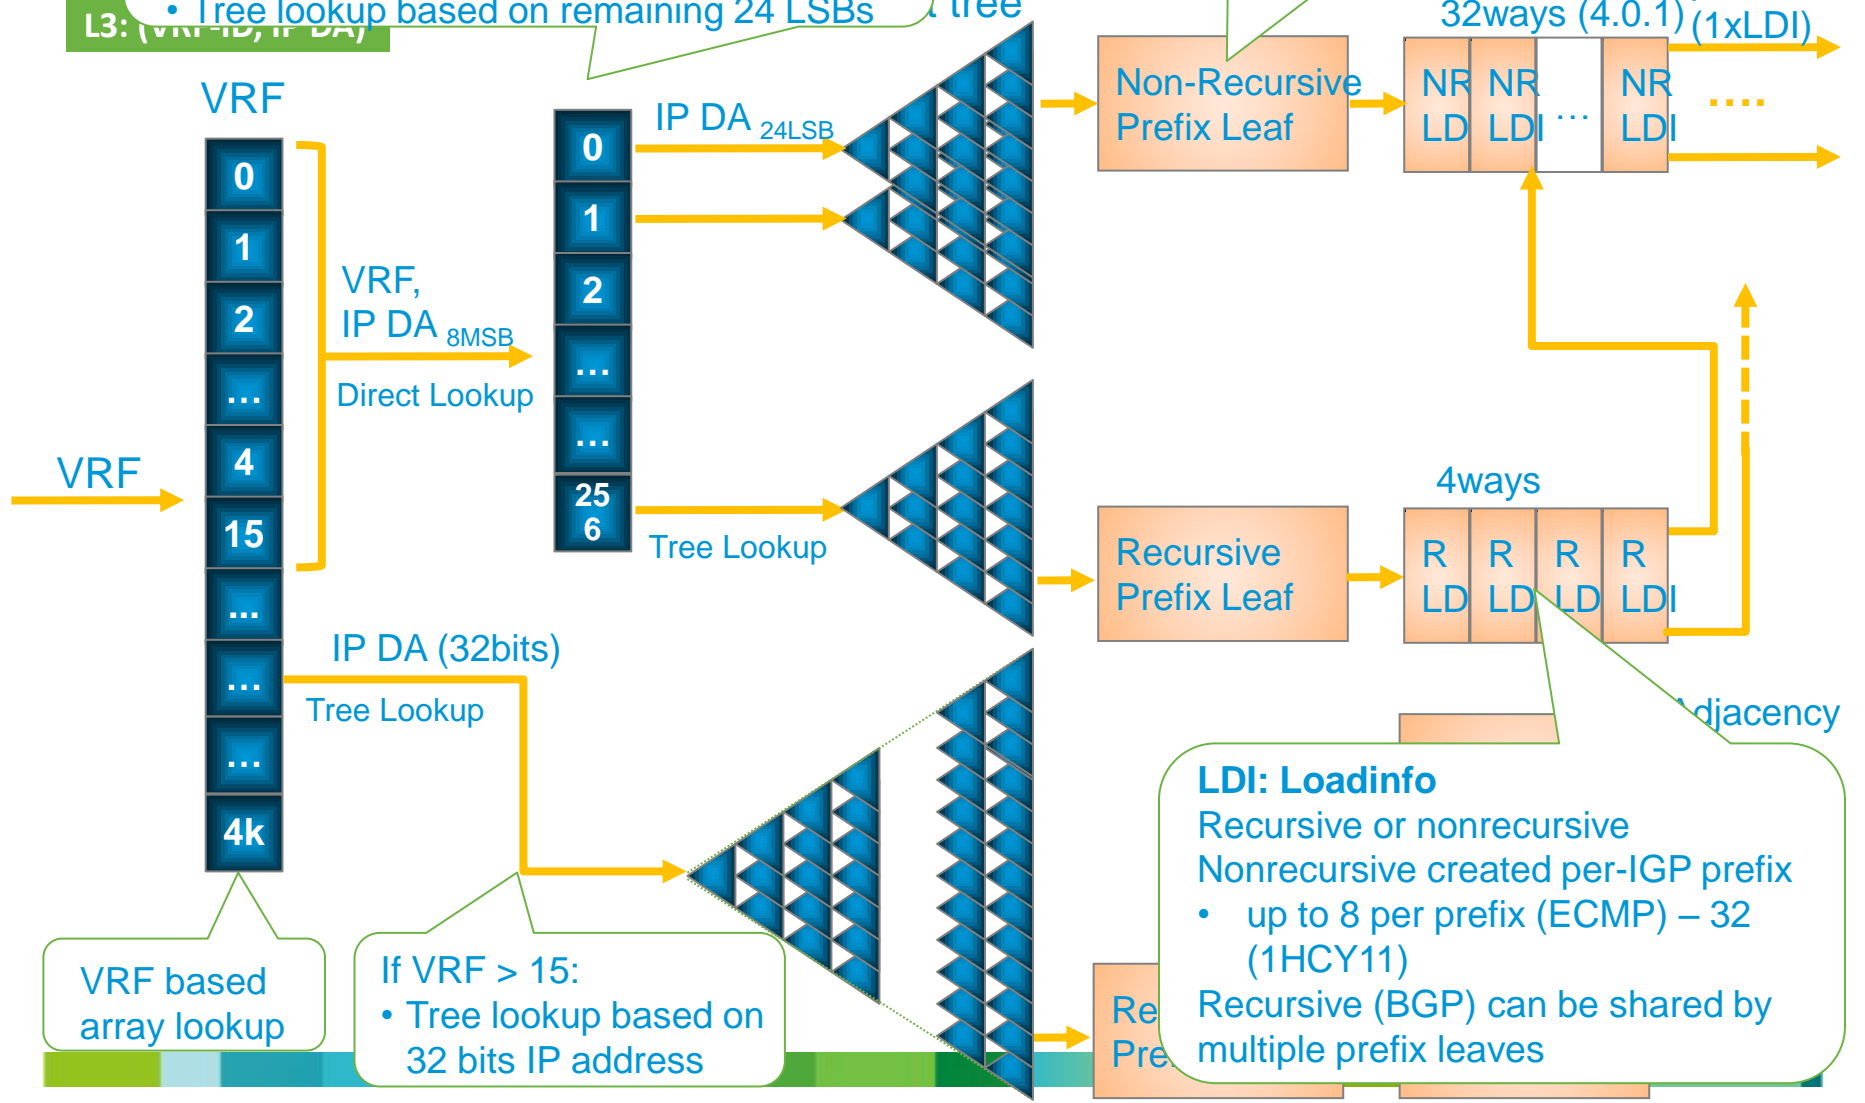
lookup key
L3: (VRF-ID, IP DA)



If VRF < 15, 2 steps route lookup for faster search:

- direct look up based on 8 MSBs of IP DA
- Tree lookup based on remaining 24 LSBs

Leaf: 1 per IPv4 prefix
Endpoint of IPv4 lookup
Points to LDI



VRF based array lookup

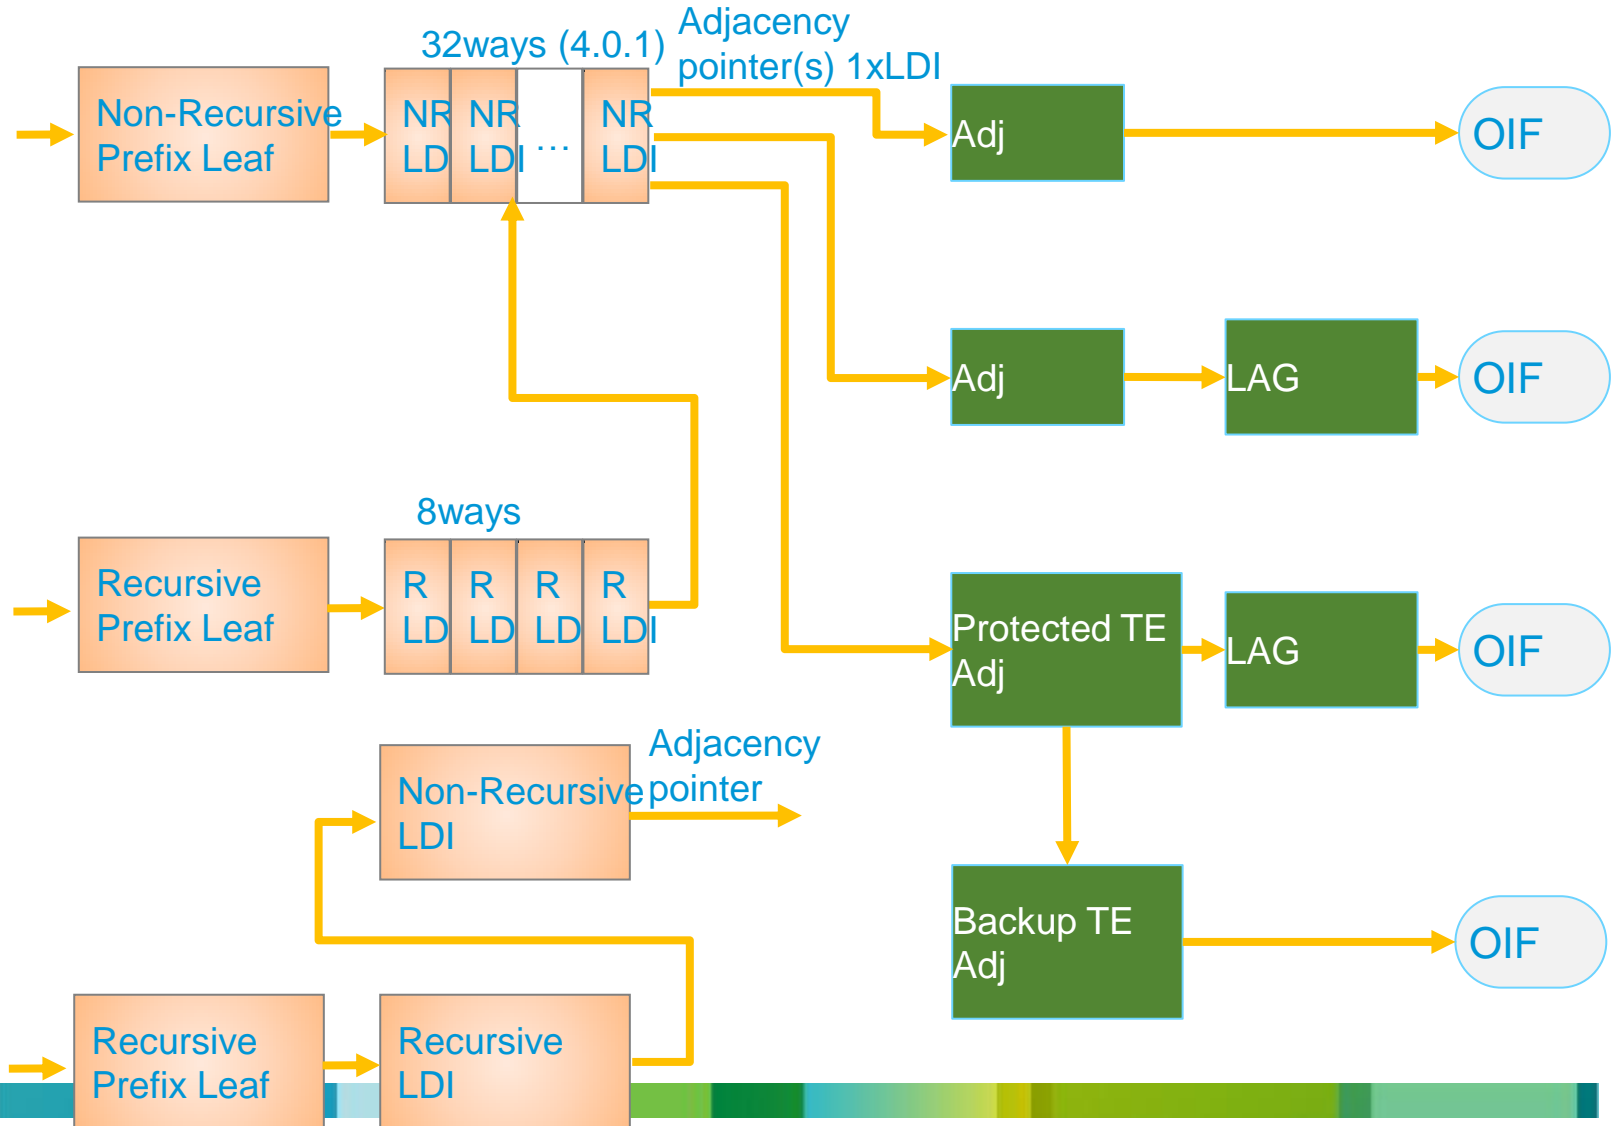
If VRF > 15:
• Tree lookup based on 32 bits IP address

LDI: Loadinfo
Recursive or nonrecursive
Nonrecursive created per-IGP prefix

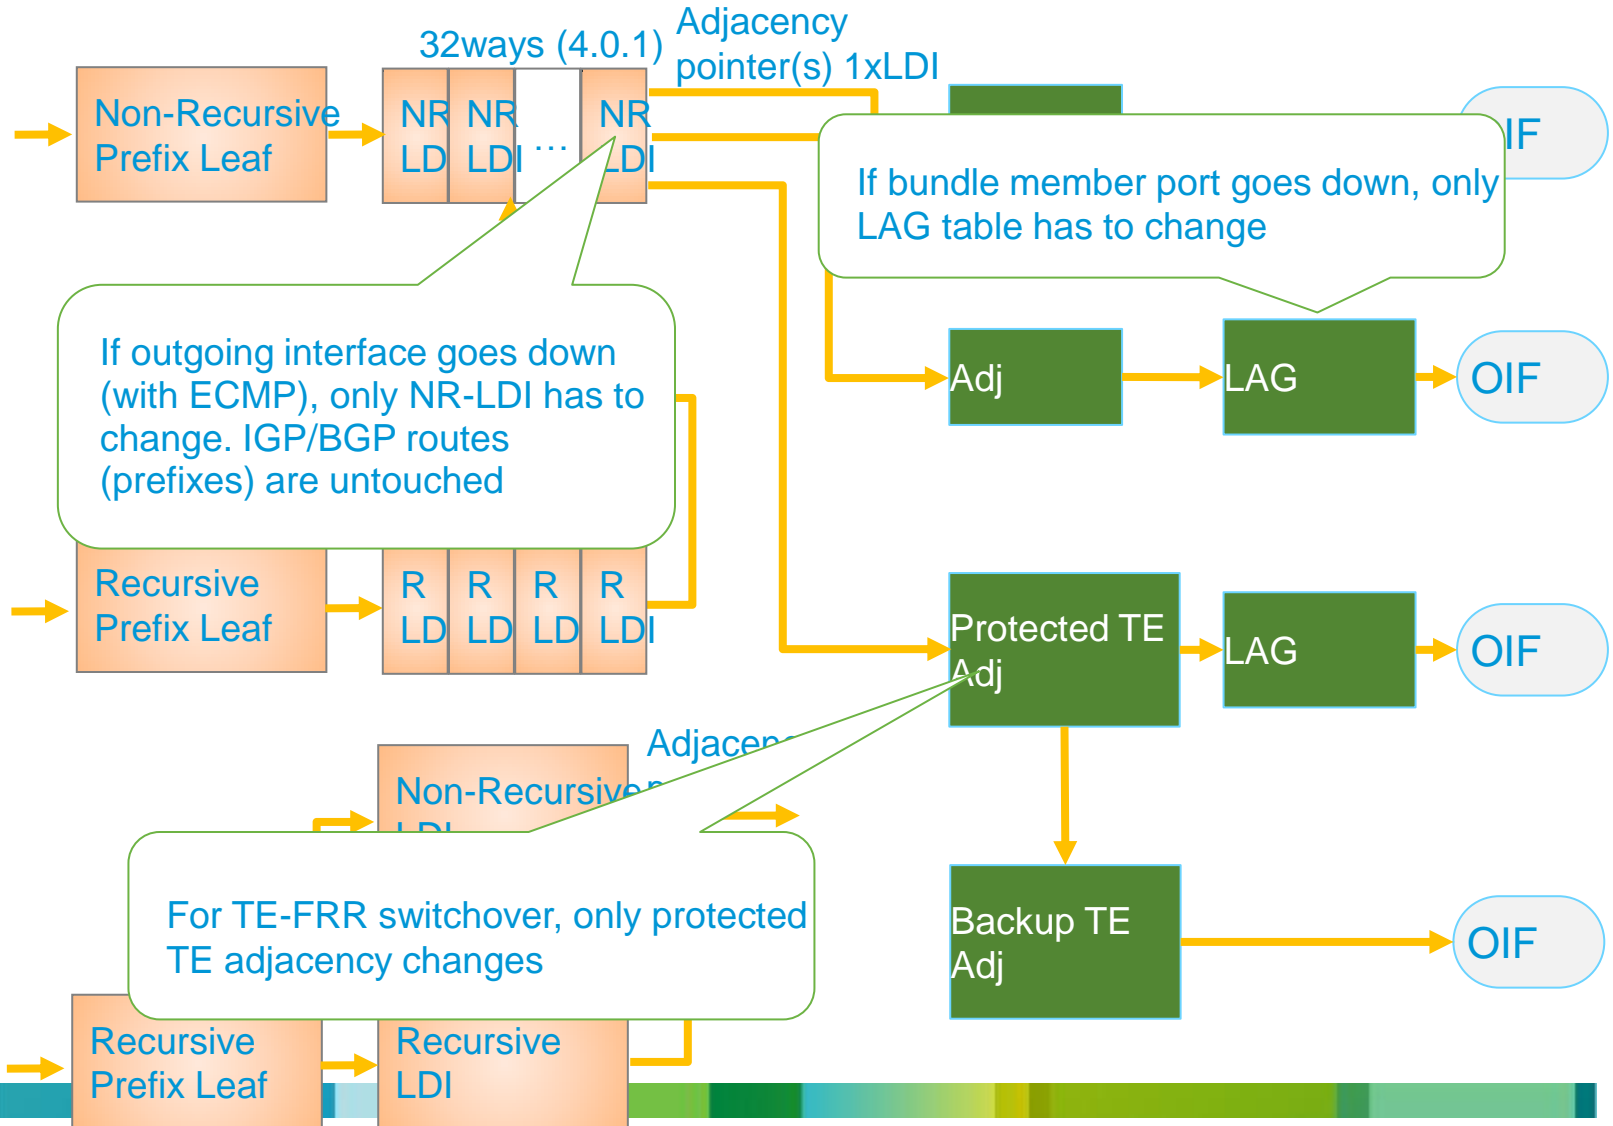
- up to 8 per prefix (ECMP) – 32 (1HCY11)

Recursive (BGP) can be shared by multiple prefix leaves

L3 NPU IPv4 FIB Architecture



L3 NPU IPv4 FIB Architecture



ECMP Load balancing

IPv6 uses first 64 bits in 4.0 releases, full 128 in 4.2 releases

1: IPv4 Unicast or IPv4 to MPLS

- No or unknown Layer 4 protocol: IP SA, DA and Router ID
- UDP or TCP: IP SA, DA, Src Port, Dst Port and Router ID

2: IPv4 Multicast

- For (S,G): Source IP, Group IP, next-hop of RPF
- For (*,G): RP address, Group IP address, next-hop of RPF

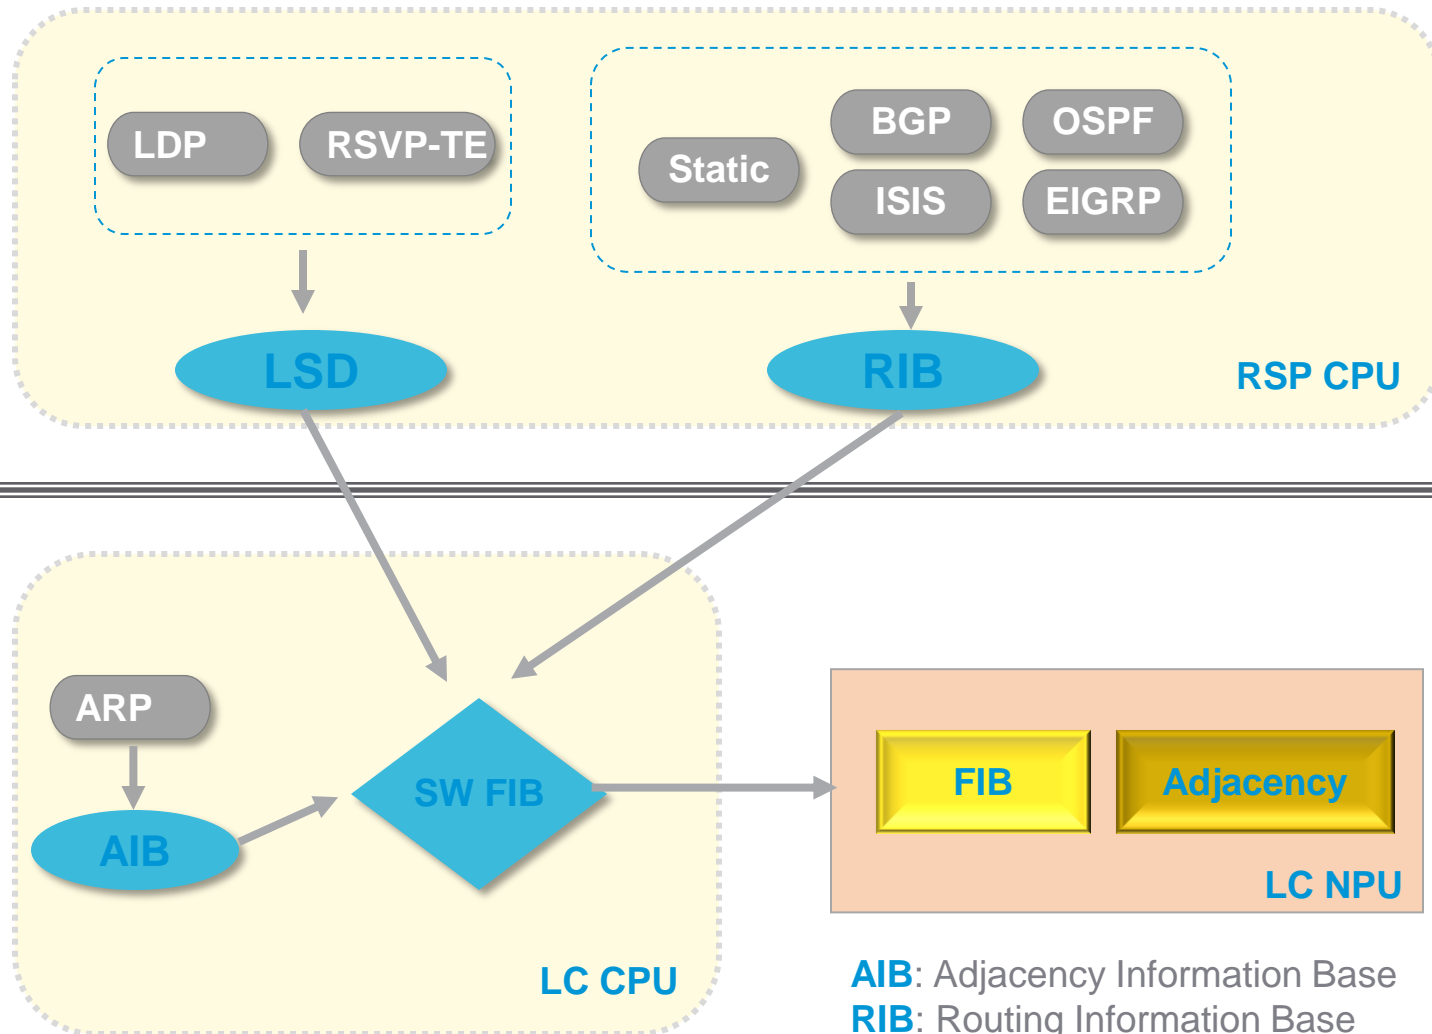
3: MPLS to MPLS or MPLS to IPv4

- # of labels ≤ 4 : same as IPv4 unicast (if inner is IP based, EoMPLS, etherheader will follow: 4th label+RID)
- # of labels > 4 : 4th label and Router ID

Bundle Load balancing

- L3 bundle uses 5 tuple as “1” (eg IP enabled routed bundle interface)
- MPLS enabled bundle follows “3”
- L2 access bundle uses access S/D-MAC + RID, OR L3 if configured (under I2vpn)
- L2 access AC to PW over mpls enabled core facing bundle uses PW label (not FAT-PW label even if configured)
 - FAT PW label only useful for P/core routers

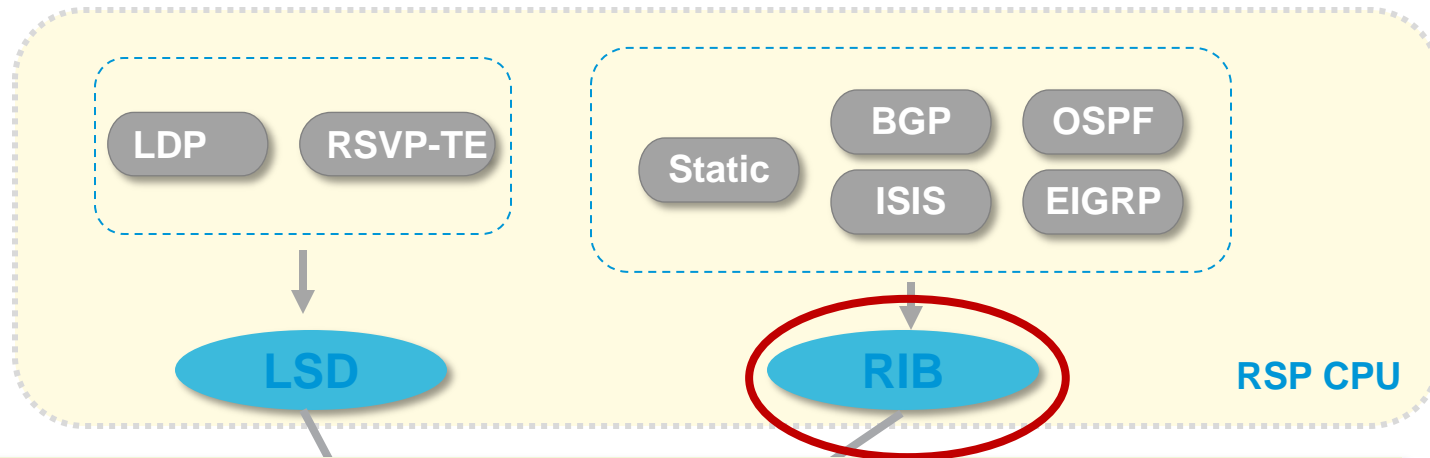
L3 IPv4 Control Plane Architecture



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

L3 IPv4 Control Plane Architecture

Show commands



```
RP/0/RSP0/CPU0:asr#sh route 222.0.0.6/31
```

```
Routing entry for 222.0.0.6/31
```

```
Known via "isis isis1", distance 115, metric 20, type level-1
```

```
Installed Mar  2 17:58:12.251 for 00:00:47
```

```
Routing Descriptor Blocks
```

```
  222.0.0.2, from 222.2.2.1, via TenGigE0/1/0/3
```

```
    Route metric is 20
```

```
No advertising protos.
```

FIB: Forwarding Information Base

LSD: Label Switch Database

L3 IPv4 Control Plane Architecture

Show commands

```
RP/0/RSP0/CPU0:asr#show adjacency summary location 0/1/CPU0
```

```
Adjacency table (version 26) has 19 adjacencies:
```

```
 11 complete adjacencies
```

```
  8 incomplete adjacencies
```

```
  0 deleted adjacencies in quarantine list
```

```
  8 adjacencies of type IPv4
```

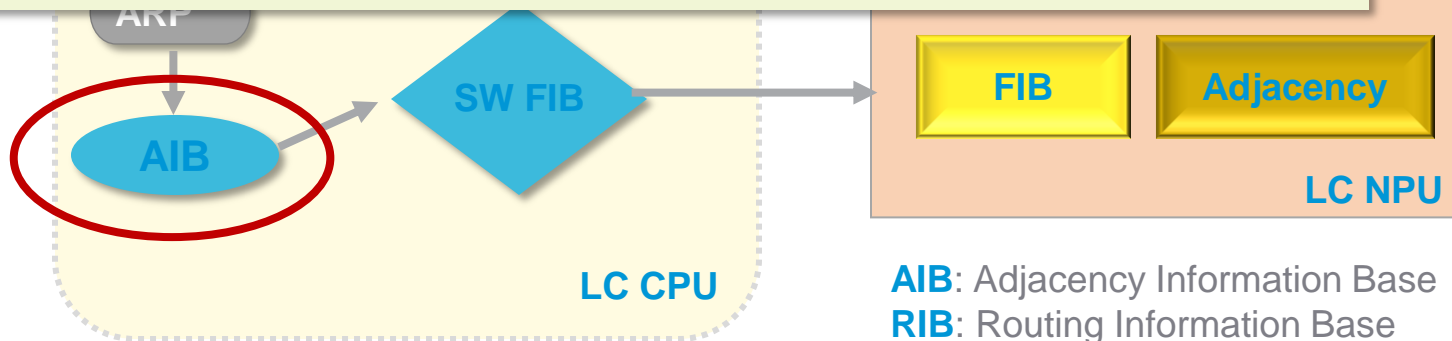
```
    8 complete adjacencies of type IPv4
```

```
    0 incomplete adjacencies of type IPv4
```

```
    0 deleted adjacencies of type IPv4 in quarantine list
```

```
    0 interface adjacencies of type IPv4
```

```
    4 multicast adjacencies of type IPv4
```



AIB: Adjacency Information Base

RIB: Routing Information Base

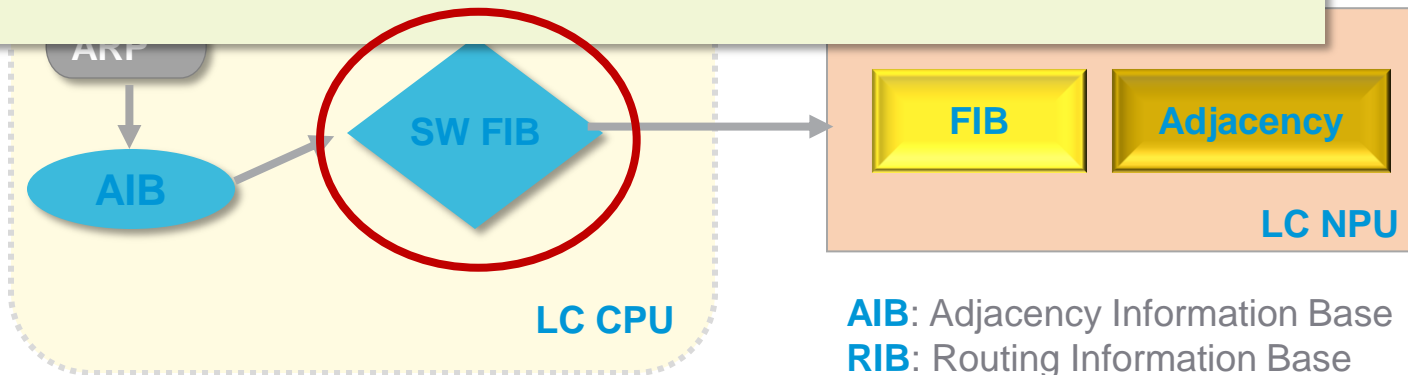
FIB: Forwarding Information Base

LSD: Label Switch Database

L3 IPv4 Control Plane Architecture

Show commands

```
RP/0/RSP0/CPU0:vikings-1#sh cef 222.0.0.6 location 0/1/CPU0
222.0.0.6/31, version 1, internal 0x40000001
Updated Mar  2 17:58:11.987
  local adjacency 222.0.0.2
  Prefix Len 31, traffic index 0, precedence routine (0)
    via 222.0.0.2, TenGigE0/1/0/3, 5 dependencies, weight 0,
class 0
  next hop 222.0.0.2
  local adjacency
```



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

L3 IPv4 Control Plane Architecture

Show commands

```
RP/0/RSP0/CPU0:asr#sh cef 222.0.0.6 hardware ingress lo 0/1/CPU0
222.0.0.6/31, version 1, internal 0x40000001 (0xb1d66c6c) [1], 0x0 (0xb1b4f758), 0x0 (0x0)
Updated Mar  2 17:58:11.987
local adjacency 222.0.0.2
Prefix Len 31, traffic index 0, precedence routine (0)
  via 222.0.0.2, TenGigE0/1/0/3, 5 dependencies, weight 0, class 0
  next hop 222.0.0.2
  local adjacency
```

EZ:0 Leaf

=====

Search ctrl-byte0: 0x3 ctrl-byte1: 0x8 ctrl-byte2:0x5

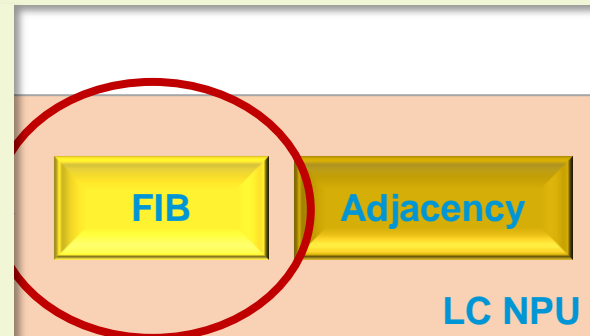
Leaf Action : FORWARD

Search Control-Flag31:

```
match      : 1      valid: 1
done       : 0      ifib_lookup: 0
ext_lsp_array : 0    match_all_bit: 0
recursive  : 0      nonrecursive : 1
default_action: 1
```

Non Recursive Leaf:

```
ldi ptr    : 10936 (0x2ab8)    igp statsptr:0
rpf ptr    : 0x0000
BCP policy a/c : 0          AS number : 0
```



AIB: Adjacency Information Base

RIB: Routing Information Base

FIB: Forwarding Information Base

LSD: Label Switch Database

References

- [Cisco Support Community XR OR and Platforms](https://supportforums.cisco.com/community/netpro/service-providers/ios-xr)
<https://supportforums.cisco.com/community/netpro/service-providers/ios-xr>
- [ASR9000/XR Feature Order of operation](#)
- [ASR9000/XR Frequency Synchronization](#)
- [ASR9000/XR: Understanding SNMP and troubleshooting](#)
- [Cisco BGP Dynamic Route Leaking feature Interaction with Juniper](#)
- [ASR9000/XR: Cluster nV-Edge guide](#)
- [Using COA, Change of Authorization for Access and BNG platforms](#)
- [ASR9000/XR: Local Packet Transport Services \(LPTS\) CoPP](#)
- [ASR9000/XR: How to capture dropped or lost packets](#)
- [ASR9000/XR Understanding Turboboot and initial System bring up](#)
- [ASR9000/XR: The concept of a SMU and managing them](#)
- [ASR9000/XR Using MST-AG \(MST Access Gateway\), MST and VPLS](#)
- [ASR9000/XR: Loadbalancing architecture and characteristics](#)
- [ASR9000/XR Netflow Architecture and overview](#)
- [ASR9000 Understanding the BNG configuration \(a walkthrough\)](#)
- [ASR9000/XR NP counters explained for up to XR4.2.1](#)
- [ASR9000/XR Understanding Route scale](#)
- [ASR9000/XR Understanding DHCP relay and forwarding broadcasts](#)
- [ASR9000/XR: BNG deployment guide](#)

References

- [ASR9000/XR: Understanding and using RPL \(Route Policy Language\)](#)
- [ASR9000/XR What is the difference between the -p- and -px- files ?](#)
- [ASR9000/XR: Migrating from IOS to IOS-XR a starting guide](#)
- [ASR9000 Monitoring Power Supply Information via SNMP](#)
- [ASR9000 BNG Training guide setting up PPPoE and IPoE sessions](#)
- [ASR9000 BNG debugging PPPoE sessions](#)
- [ASR9000/XR : Drops for unrecognized upper-level protocol error](#)
- [ASR9000/XR : Understanding ethernet filter strict](#)
- [ASR9000/XR Flexible VLAN matching, EVC, VLAN-Tag rewriting, IRB/BVI and defining L2 services](#)
- [ASR9000/XR: How to use Port Spanning or Port Mirroring](#)
- [ASR9000/XR Using Task groups and understanding Priv levels and authorization](#)
- [ASR9000/XR: How to reset a lost password \(password recovery on IOS-XR\)](#)
- [ASR9000/XR: How is CDP handled in L2 and L3 scenarios](#)
- [ASR9000/XR : Understanding SSRP Session State Redundancy Protocol for IC-SSO](#)
- [ASR9000/XR: Understanding MTU calculations](#)
- [ASR9000/XR: Troubleshooting packet drops and understanding NP drop counters](#)
- [Using Embedded Event Manager \(EEM\) in IOS-XR for the ASR9000 to simulate ECMP "min-links"](#)
- [XR: ASR9000 MST interop with IOS/7600: VLAN pruning](#)

Submit Your Questions Now!

Use the Q&A panel to submit your questions. Experts will start responding those



Q & A

Expert responding some of your questions verbally. Use the Q&A panel to continue asking your questions



We Appreciate Your Feedback!

Those who fill out the Evaluation Survey will enter a raffle to win:

\$50 Amazon Gift Card

To complete the evaluation, please click on link provided in the chat or in the pop-up once the event is closed.

Trivia Question (select the correct answer)

What do Cisco Router Technology, McDonalds and Disney have in common?

In this same year

Cisco earned its first patent for its method and apparatus for routing communications among computer networks. The first McDonalds in Beijing China opened and Euro Disney opened in France.

What year was it?

- a) 1992
- b) 1994
- c) 1995



Ask The Experts Event (with Xander Thuijs)



If you have additional questions, you can ask them to Xander. He will be answering from May 14 to May 24, 2013

<https://supportforums.cisco.com/thread/2216914>

You can watch the video or read the Q&A 5 business days after the event at


<https://supportforums.cisco.com/community/netpro/ask-the-expert/webcasts>



Next Webcast in Russian

Configuration and Features of Border Gateway Protocol (BGP) Prefix-Independent Convergence

Tuesday May 21st at



10:00 a.m. Brussels
12:00 p.m. Moscow

With Cisco expert: **Nataliya Omelyanyuk**

During this live event session explains the advantages and features of Border Gateway Protocol (BGP) prefix-independent convergence technology in service provider networks. She will also provide configuration examples for Cisco IOS Software and Cisco IOS XR Software.

Join the discussion for these Ask The Expert webcasts at:

<https://supportforums.cisco.com/community/netpro/expert-corner#view=webcasts>

Next Webcast in Japanese

Cisco Catalyst High CPU Troubleshooting Training

Tuesday May 28th

10:00 a.m. Japan Standard Time

Tuesday May 27th

6:00 p.m. PDT (San Francisco)

With the Cisco expert: **Yasuhiro Nakajima**



This session discusses several troubleshooting examples for High CPU Utilization on Cisco Catalyst Series. Nakajima will note major checkpoints and specific ways to narrow down the problems on this issue.

Join the discussion for these Ask The Expert webcasts at:

<https://supportforums.cisco.com/community/netpro/expert-corner#view=webcasts>

Next Webcast in English

Configuration, Design, and Troubleshooting of Cisco Nexus 1000



Tuesday, June 4 at

7:00 a.m. PDT (San Francisco)
10:00 a.m. EDT (New York)
3:00 p.m. Paris

With Cisco expert: **Louis Watta**

During this live event, Cisco expert Louis Watta will go over the design, configuration, and troubleshooting of Cisco Nexus 1000V Series Switches operating inside VMware ESXi and Hyper-V..

Join the discussion for these Ask The Expert webcasts at:

http://tools.cisco.com/gems/cust/customerQA.do?METHOD=E&LANGUAGE_ID=E&SEMINAR_CODE=S18361&PRIORITY_CODE=

Ask the Expert Events – English



Topic: Fiber Channel over Ethernet (FCoE)

With Cisco expert Ozden Karakok

Learn how to design, plan, configure, implement, and troubleshoot Fibre Channel over Ethernet

Ends May 17, 2013



Topic: SSO with CWMS, IRP and ELM Solutions

With Cisco expert Arun Kumar

Learn and ask questions about Cisco WebEx Meetings Server (CWMS)

Ends May 17, 2013



Topic: Understanding, configuring and troubleshooting IP Multicast and MVPN

With Cisco expert Pulikkal Sekharan Raju

Learn and ask questions on how to deploy, configure and troubleshoot Single Sign On (SSO), Internet Reverse Proxy (IRP), Enterprise License Manager (ELM)

Starts May 20, 2013



Topic: Deploying Cisco FabricPath in Data Center NetworkFabricPath

With Cisco Anees Mohamed/Viral Bhutta

Learn and ask questions about how to plan, design, and implement Cisco Overlay Transport Virtualization (OTV) in your Data Center Network.

Starts May 20, 2013

Join the discussion for these Ask The Expert Events at:

<https://supportforums.cisco.com/community/netpro/expert-corner#view=ask-the-experts>

Upcoming Events on Cisco Live 365

June 25-27

Cisco Live 2013 Orlando – Virtual Event

Full agenda including live WebEx sessions will be published at the end of May.

Agenda & Details: <https://www.ciscolive365.com/connect/agenda.www>



We invite you to actively collaborate in the Cisco Support Community and social media

<https://supportforums.cisco.com>



<http://www.facebook.com/CiscoSupportCommunity>



http://twitter.com/#!/cisco_support



<http://www.youtube.com/user/ciscosupportchannel>



<https://plus.google.com/110418616513822966153?prsrc=3#110418616513822966153/posts>



<http://itunes.apple.com/us/app/cisco-technical-support/id398104252?mt=8>



https://play.google.com/store/apps/details?id=com.cisco.swtg_android



<http://www.linkedin.com/groups/CSC-Cisco-Support-Community-3210019>



Newsletter Subscription:

https://tools.cisco.com/gdrp/coiga/showsurvey.do?surveyCode=589&keyCode=146298_2&PHYSICAL%20FULFILLMENT%20Y/N=NO&SUBSCRIPTION%20CENTER=YES

We have communities in other languages

If you speak **Spanish, Portuguese, Japanese, Polish or Russian**, we invite you to ask your questions and collaborate in your language:

- Spanish → <https://supportforums.cisco.com/community/spanish>
- Portuguese → <https://supportforums.cisco.com/community/portuguese>
- Japanese → <https://supportforums.cisco.com/community/csc-japan>
- Polish → <https://supportforums.cisco.com/community/etc/netpro-polska>
- Russian → <https://supportforums.cisco.com/community/russian>

Trivia Answer

What do Cisco Router Technology, McDonalds and Disney have in common?

In this same year

Cisco earned its first patent for its method and apparatus for routing communications among computer networks. The first McDonalds in Beijing China opened and Euro Disney opened in France.

What year was it?

- a) 1992
- b) 1994
- c) 1995



Thank You for
Your Time

Please Take a Moment to Complete the Evaluation



Thank you.

