Cisco Connect

Hong Kong
March 6-7, 2013

Cisco ASR 9000 Architecture

Vincent Ng
TME SPNG

Cisco ASR9000 – Next-Gen Edge Routing Platform

Key Design Goals & System Benefits

- Architectural Design for Longevity
- Product Portfolio with significant HW and SW commonality
- Highly integrated Network Processors for High Speed Scale and Feature Flexibility
- Cisco IOS XR based
  - Truly modular, full distributed OS
  - Enhanced for the Edge (L2 and L3)
- nV (Network Virtualization) for Operational Simplicity

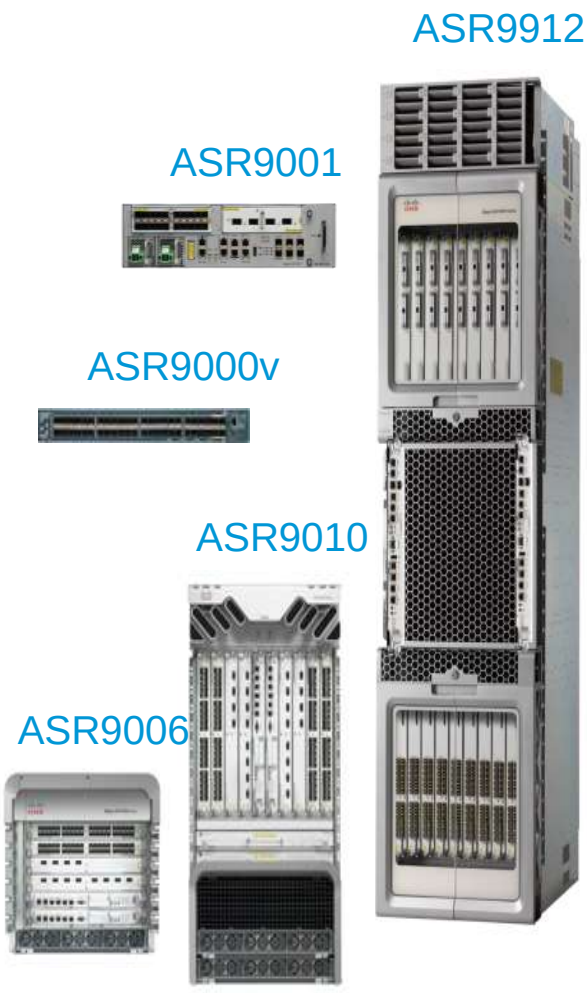

Agenda

- ASR9000 Hardware Overview
  - System Introduction and Chassis Overview
  - System Components
- Carrier Class, Scalable System Architecture
  - Fabric Architecture
  - Linecard Architecture
- Cisco nV – Network Virtualization
- Summary

ASR 9000 Hardware Overview

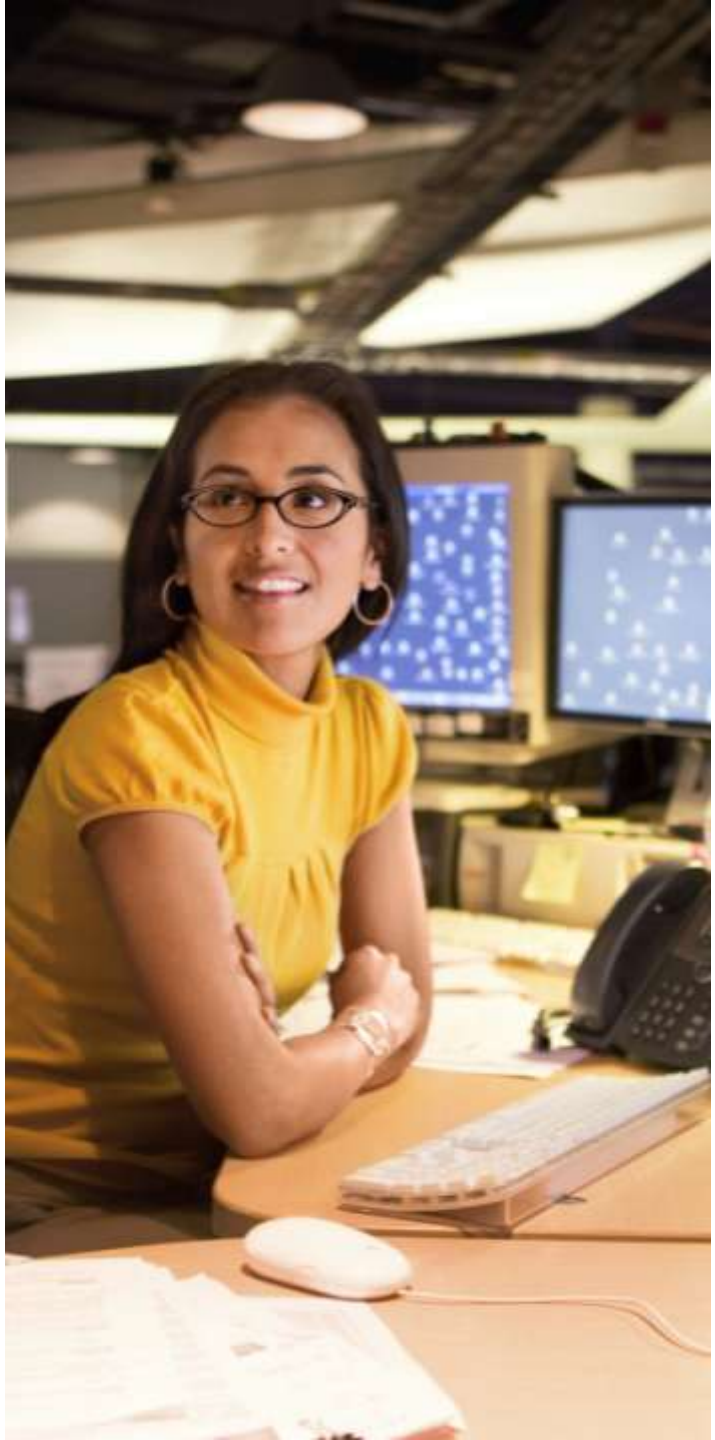

System Introduction and Chassis Overview

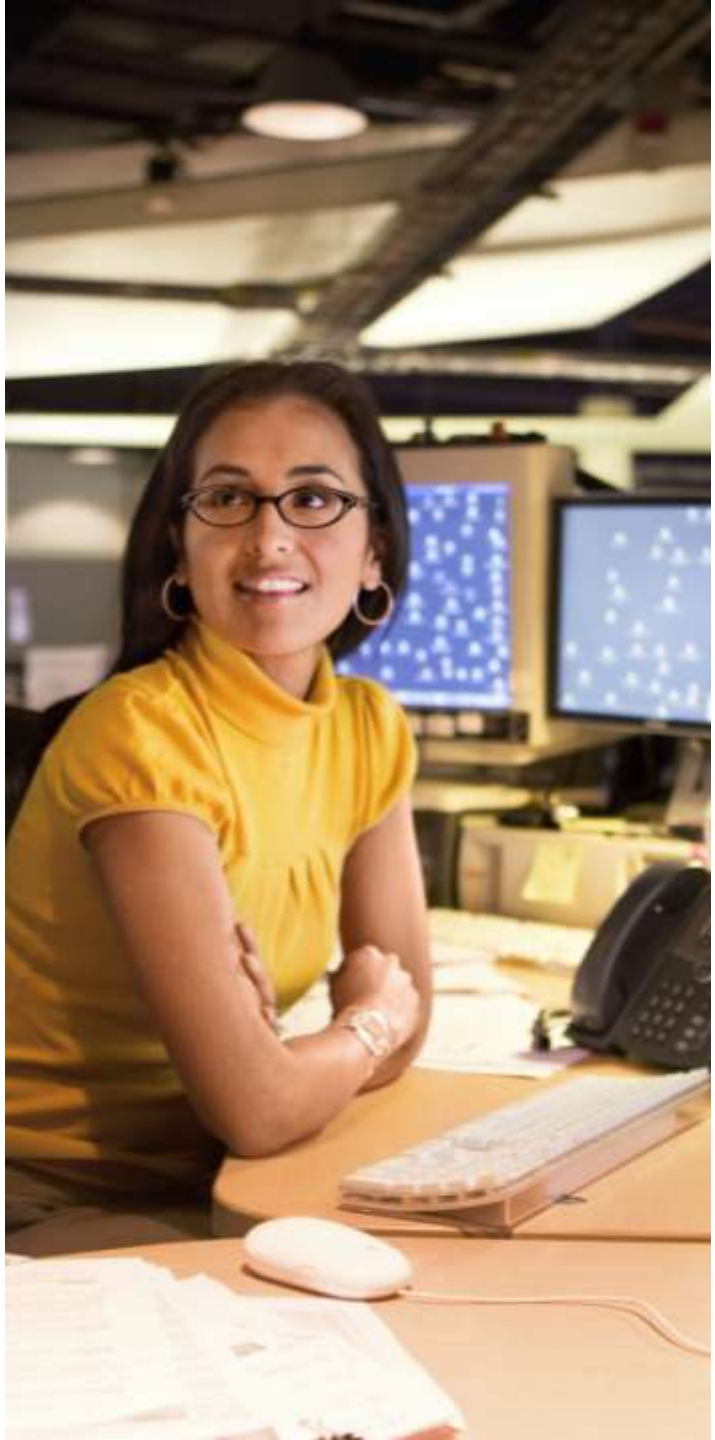

ASR 9000 Chassis Overview

Identical hardware, software, and forwarding
Across multiple physical form factors
Flexibility in data- to control-plane relationship

120 Gbps*

1.7 Tbps*

3.5 Tbps*

15Tbps*

	ASR 9001	ASR 9006	ASR 9010	ASR 9922
Switch fabric capacity	120Gbps/system 4x10 + 2 I/O bays	440G/slot 4 I/O slots	440G/slot 8 I/O slots	770G/slot 20 I/O slot
Size	2RU	10RU	21RU	44RU
Max 1/10/100GE	40/12*	160/144/8*	320/288/16*	800/720/40*
Hardware in development to double fabric capacity and port density over 18-24mo				

***actual, shipping, year 2012, engineering-math, not doubled, usable densities**

ASR 9010 and ASR 9006 Chassis

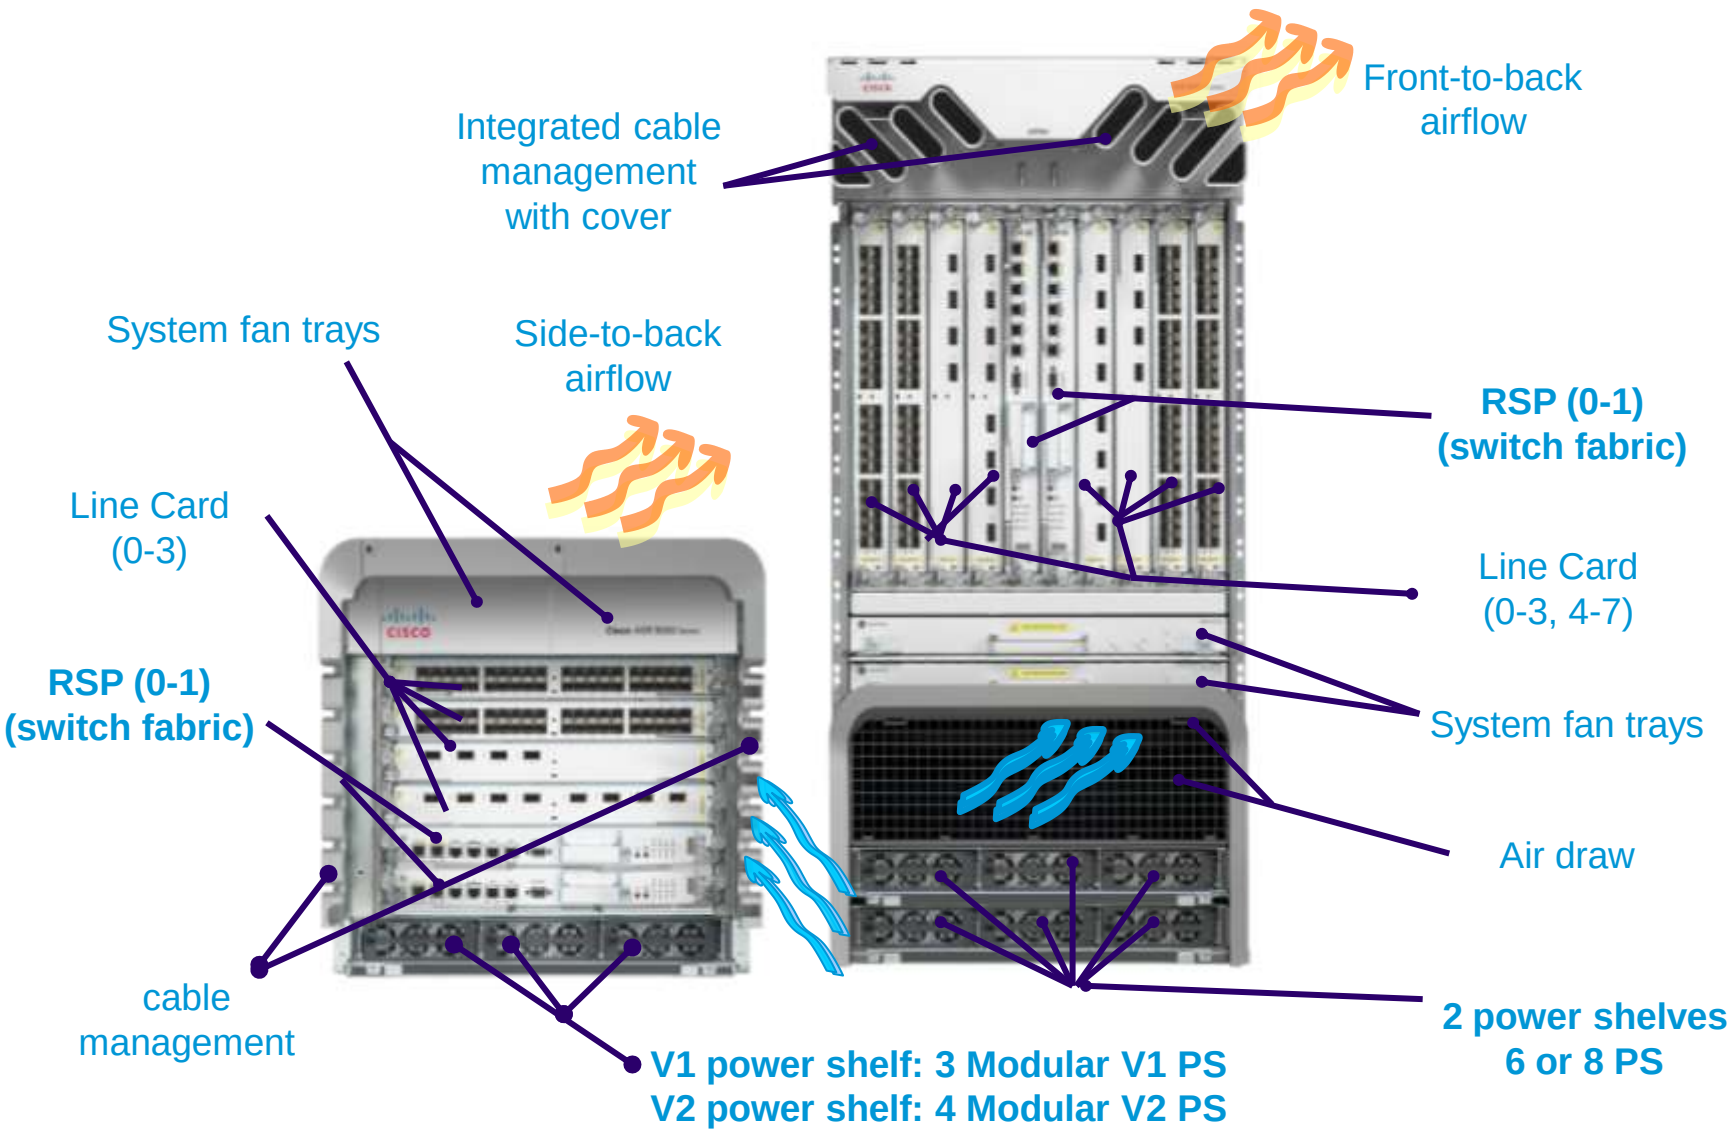

ASR 9001 Compact Chassis

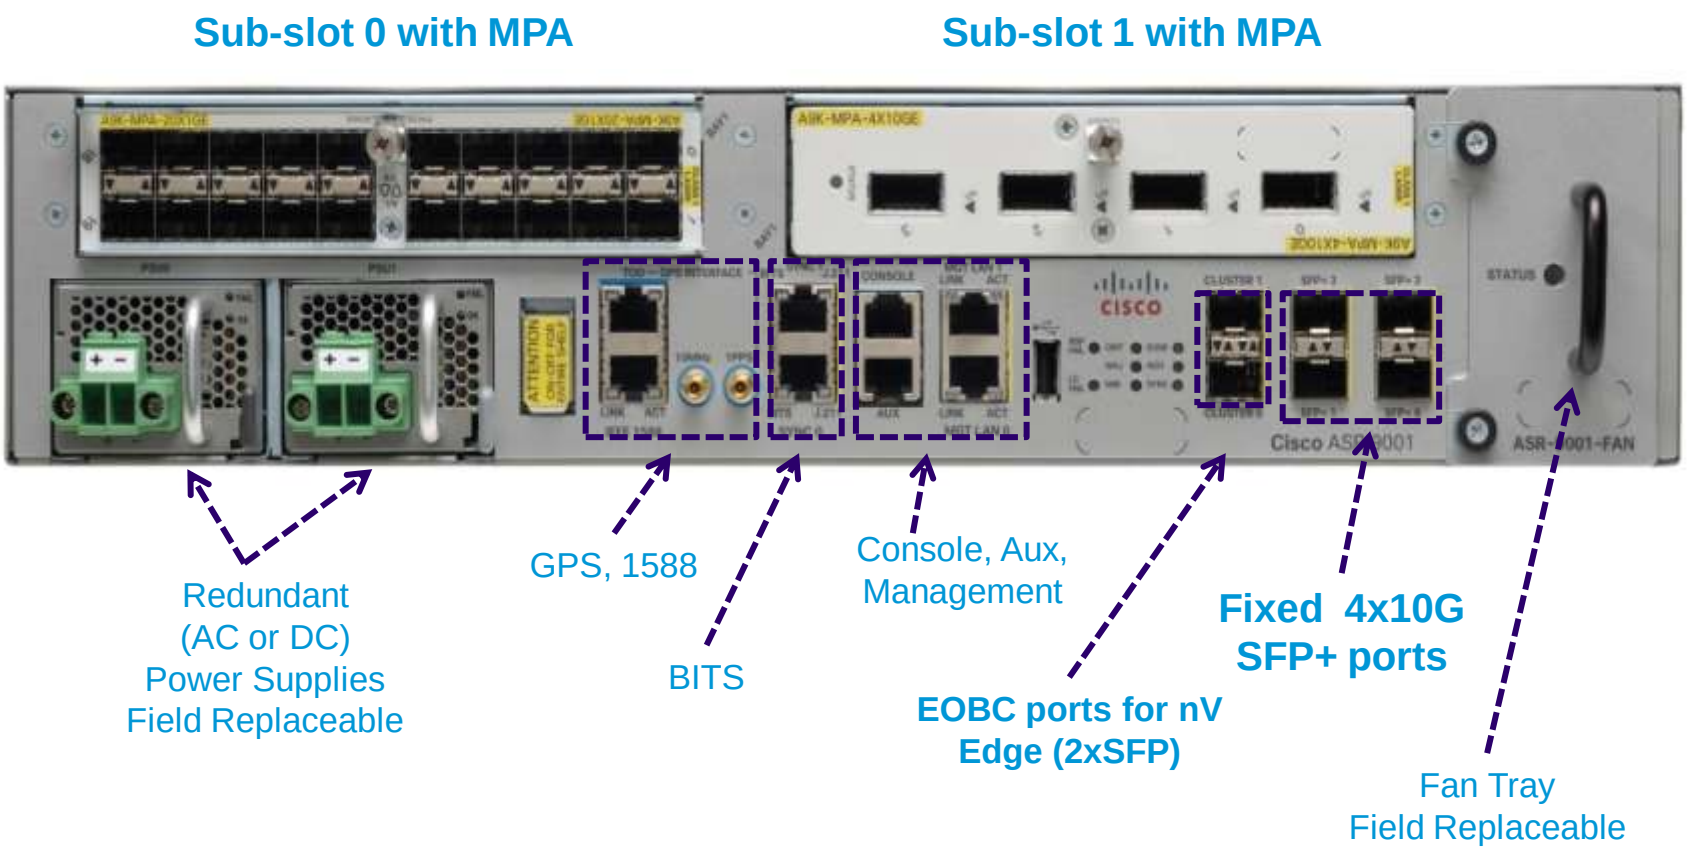

ASR9922 Large Scale Chassis

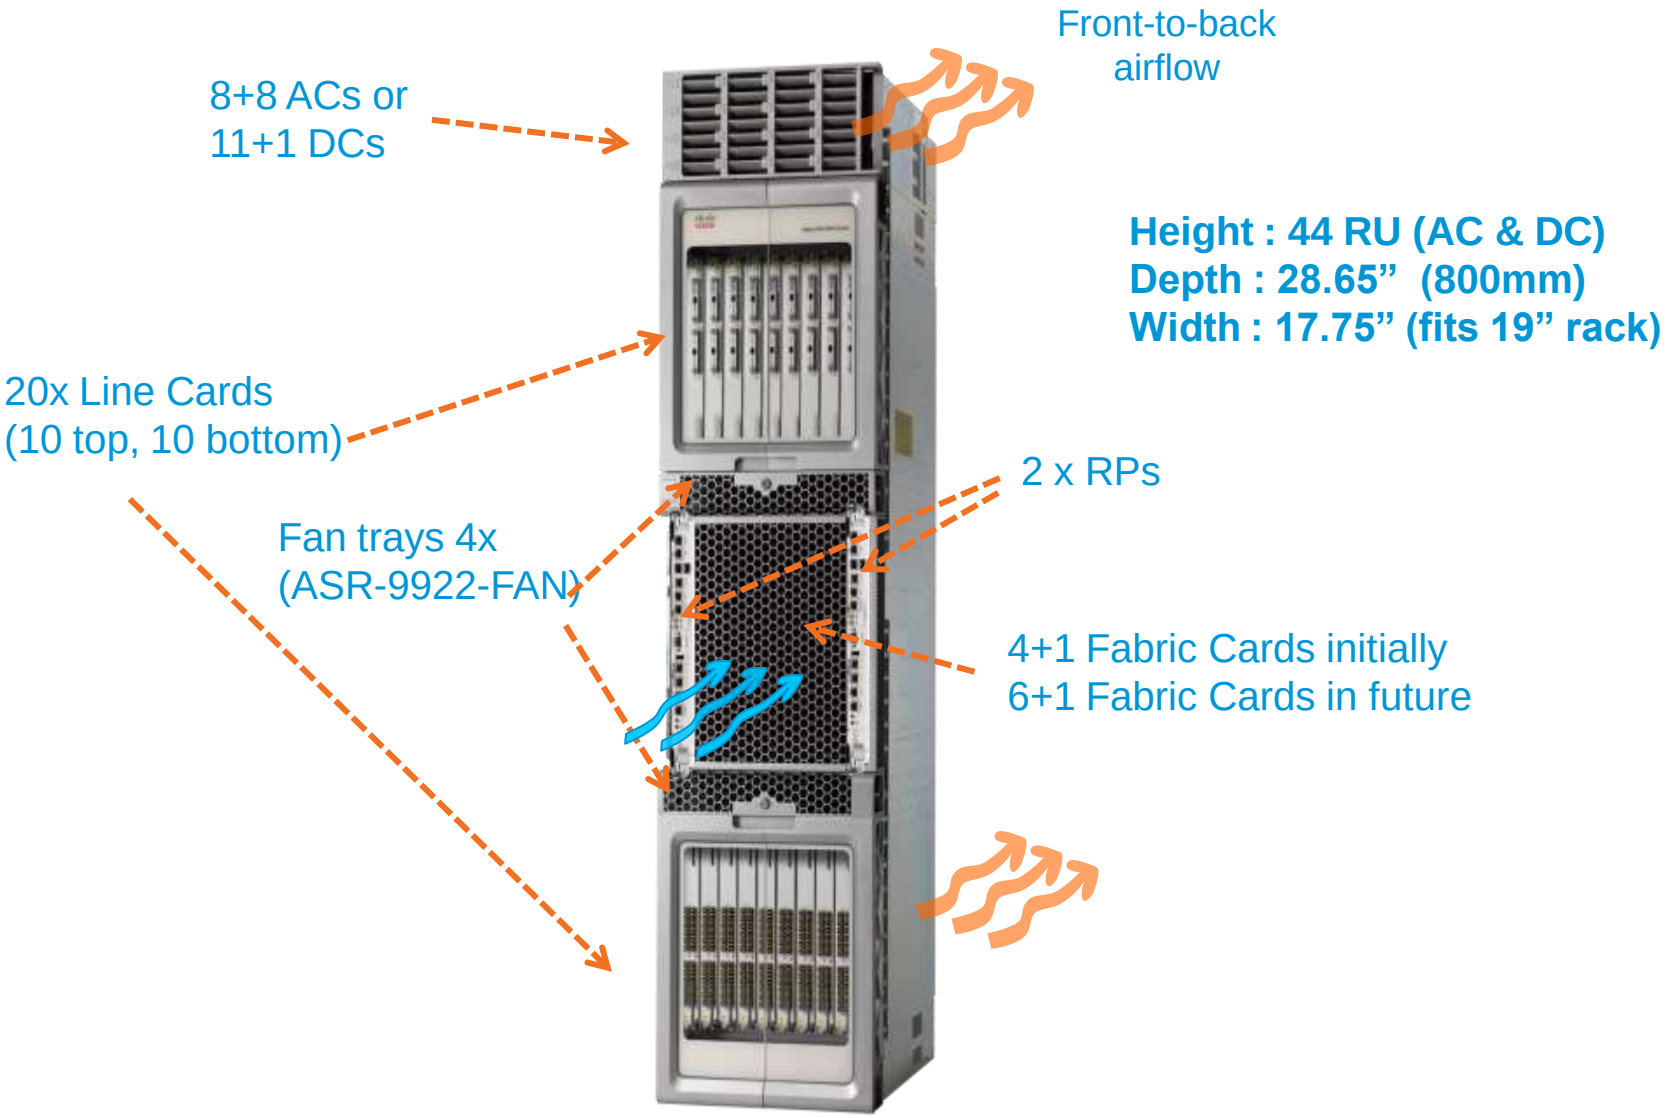

System Components

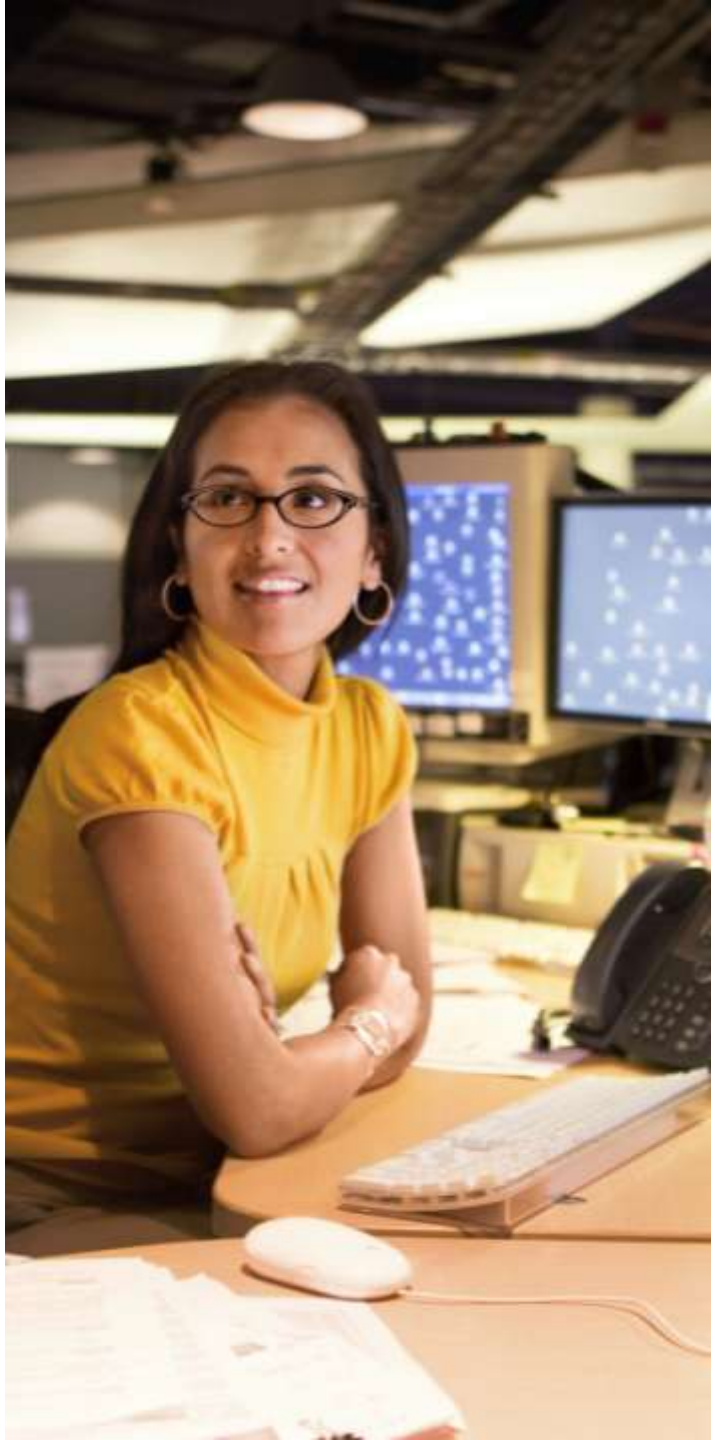

Power and Cooling

ASR-9010-FAN

ASR-9006-FAN

- Fans unique to chassis
- Variable speed for ambient temperature variation
- Redundant fan-tray
- Low noise, NEBS and OSHA compliant

Power Supply

DC Supplies

A
B
1.5 kW [1)]

A
B
2.1 kW

AC Supplies

A
3 kW

B
3 kW

- Single power zone
- All power supplies run in active mode
- Power draw shared evenly
- 50 Amp DC Input or 16 Amp AC for Easy CO Install

1) Version 1 only

Version 1 Power vs Version 2 Power System

Power Switch:
V1 → in the back
V2 → in the front

PEM Insertion from the Front

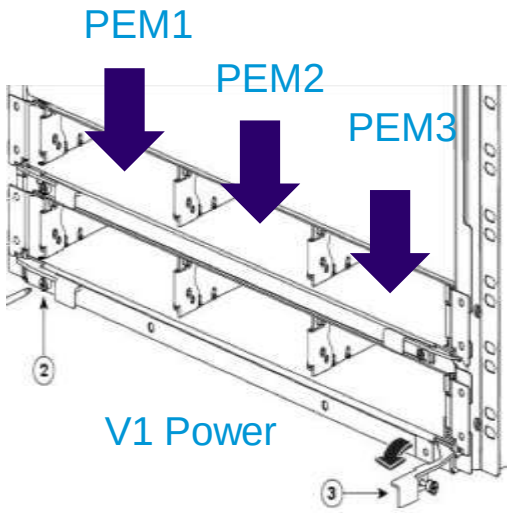

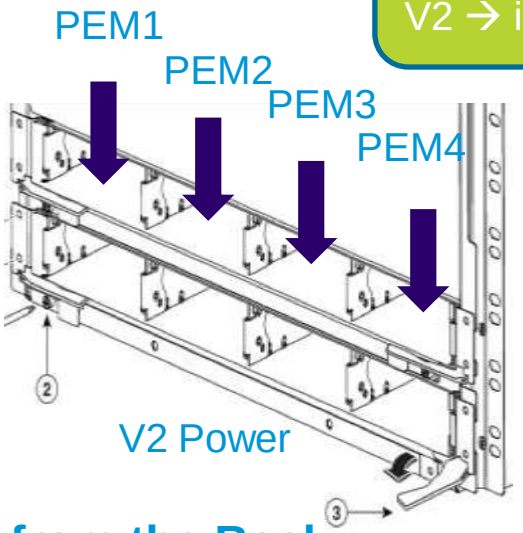

Power Feed Cabling from the Back

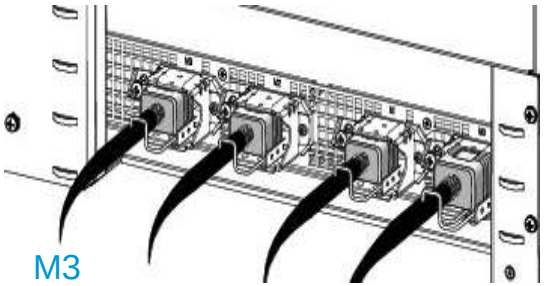

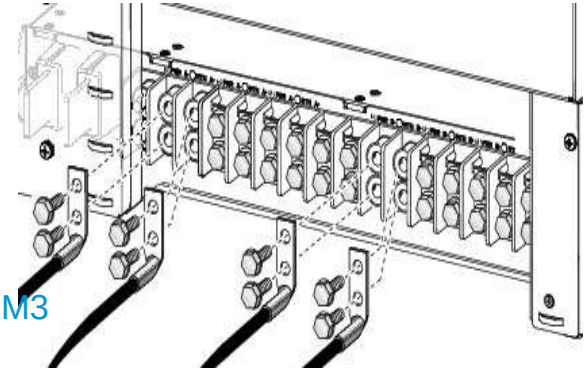

Route Switch Processors (RSPs) and Route Processors (RPs)

RSP used in ASR9006/ASR9010, RP used in ASR9922

	RSP	RSP440	9922-RP
		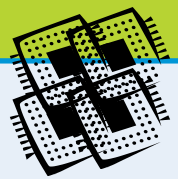	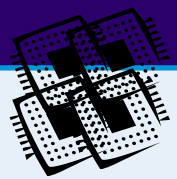
Processors	PPC/Freescale 2 Core 1.5GHz	Intel x86 4 Core 2.27 GHz	Intel x86 4 Core 2.27 GHz
RAM	RSP-4G: 4GB RSP-8G: 8GB	RSP440-TR: 6GB RSP440-SE: 12GB	-TR: 6GB -SE: 12GB
nV EOBC ports	No	Yes, 2 x 1G/10G SFP+	Yes, 2 x 1G/10G SFP+
Switch fabric bandwidth	92G + 92G (with dual RSP)	220+220G (with dual RSP)	660+110 (7-fabric model)

RSP440 – Faceplate and Interfaces

Same Front Panel on RP for ASR9922

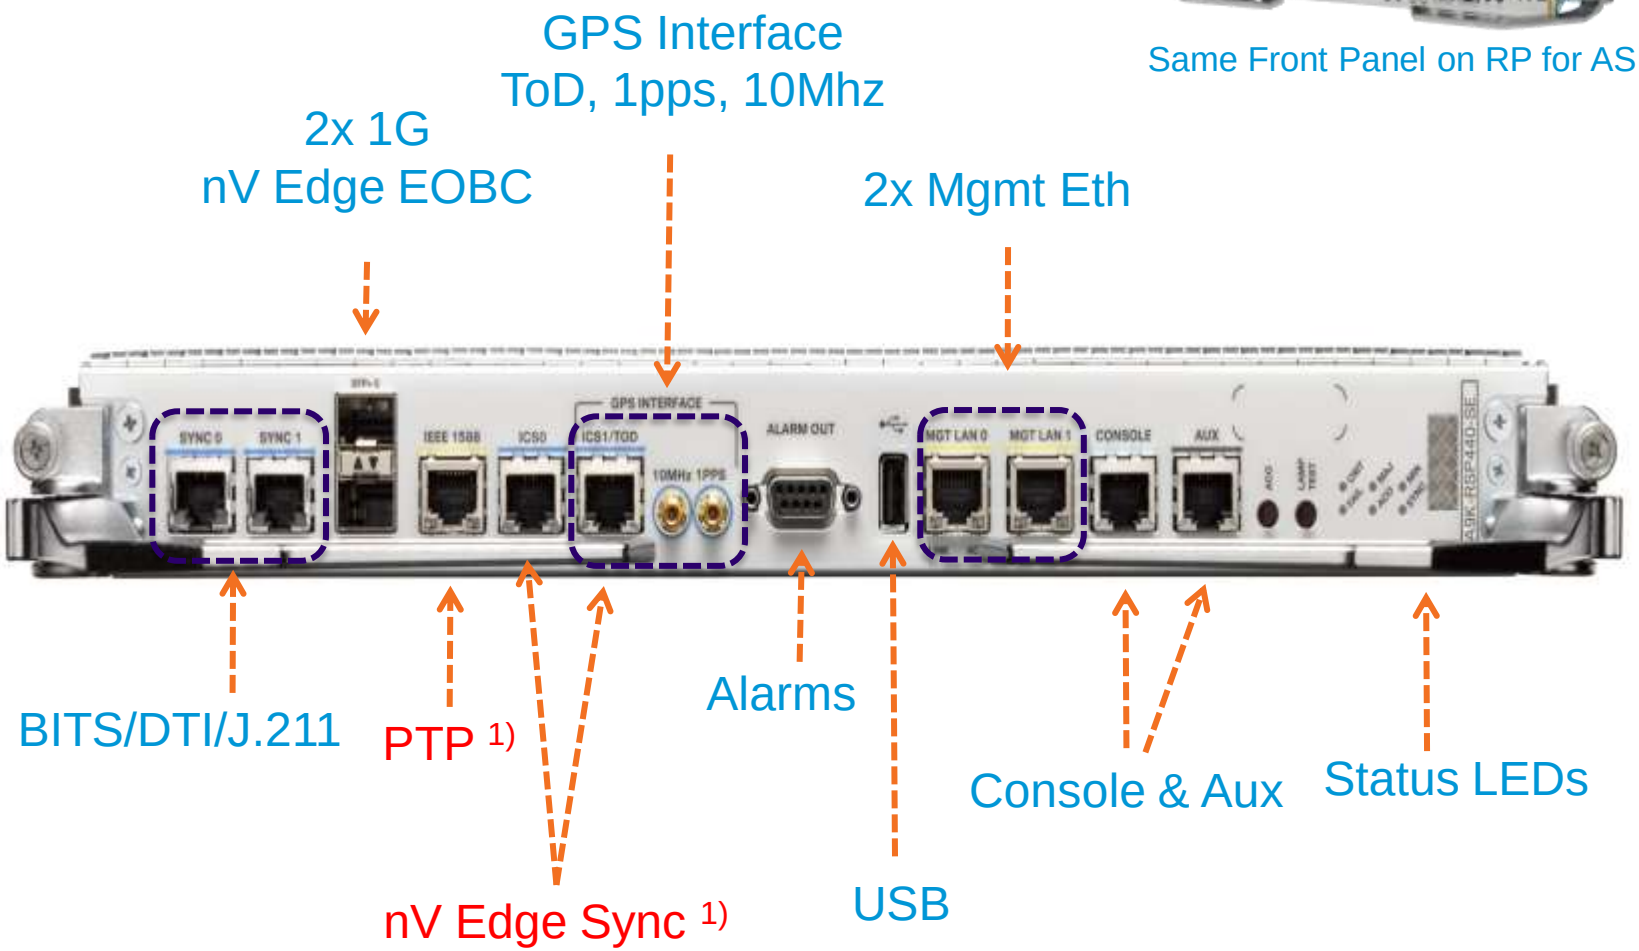

1) Future SW support

RSP Engine Architecture

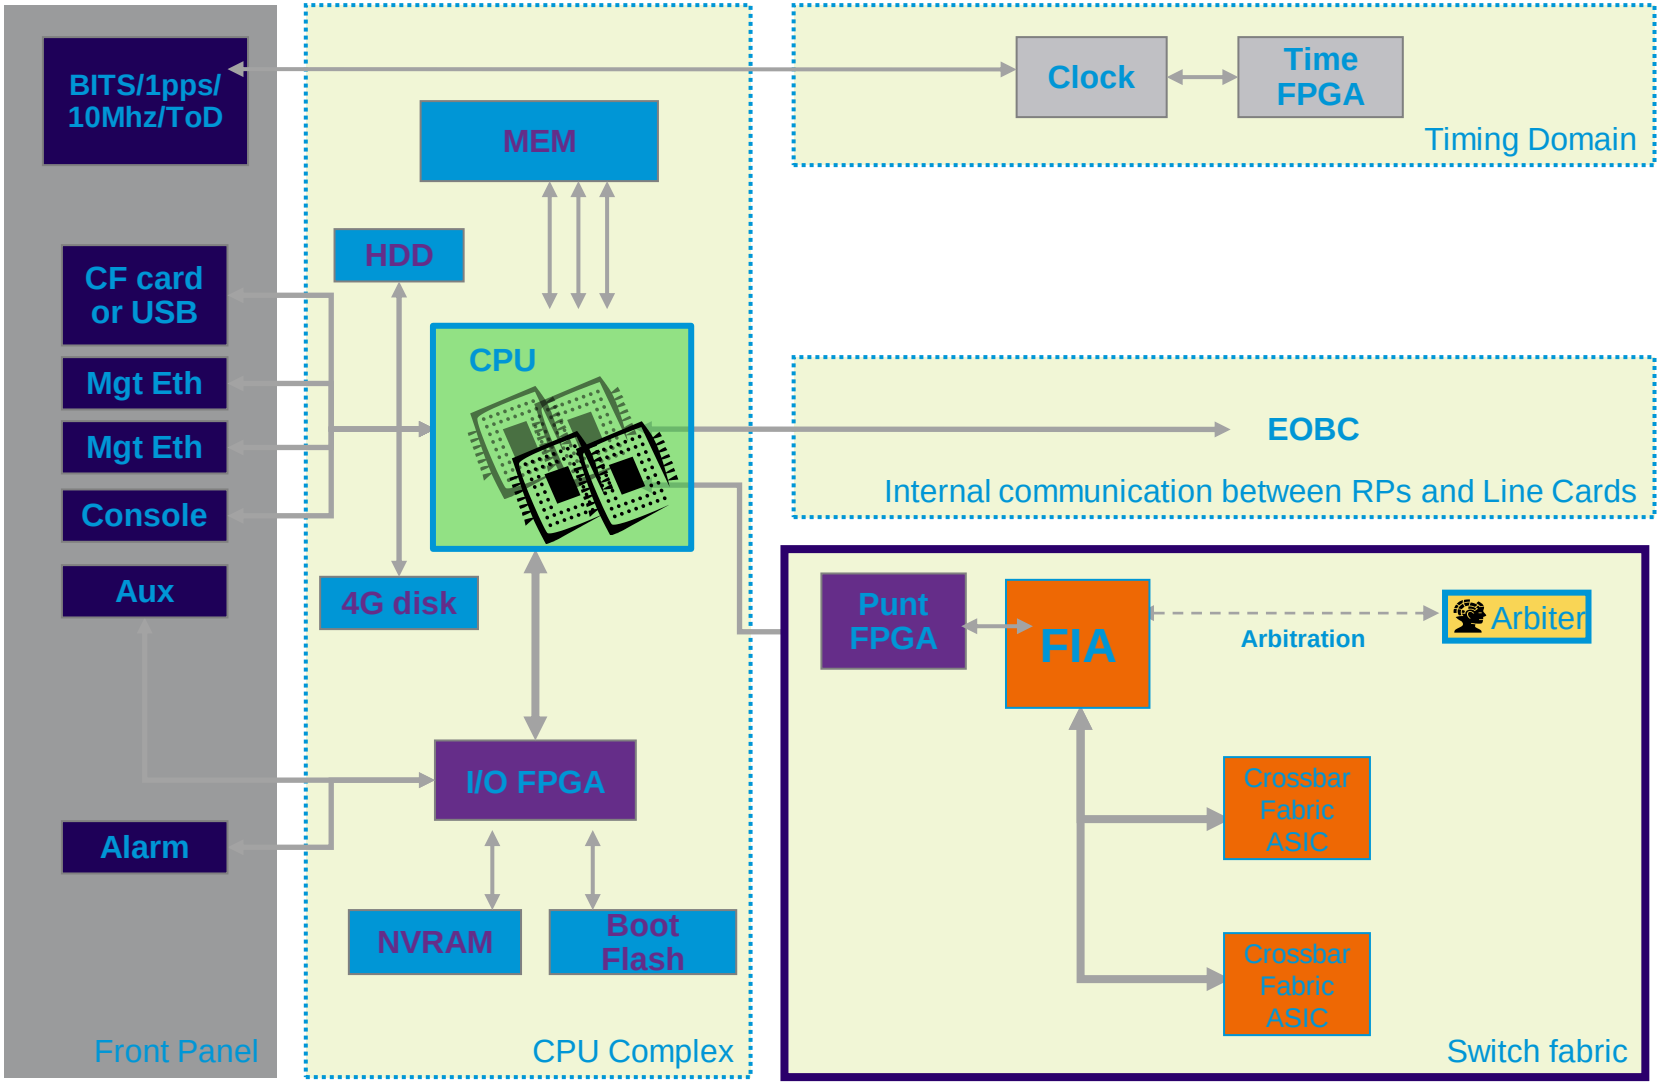

ASR 9000 Ethernet Line Card Overview

-L, -B, -E

First-generation LC (Trident)

A9K-40G A9K-4T A9K-8T/4 A9K-2T20G A9K-8T A9K-16T/8

-TR, -SE

Second-gen LC (Typhoon)

A9K-24x10GE A9K-2x100GE

A9K-MOD160

A9K-MOD80

A9K-36x10GE

MPAs
20x1GE
2x10GE
4x10GE
1x40GE
2x40GE

Modular SPA Linecard

20Gbps, feature ritch, high scale, low speed Interfaces

Quality of Service

- 128k Queues
- 128k Policers
- H-QoS
- Color Policing

Scalability

- Distributed Control and Data Plane
- 20Gbits, 4 SPA Bays
- L3 i/f, route, session protocol – scaled for MSE needs

High Availability

- IC-Stateful Switch Over Capability
- MR-APS
- IOS-XR base for high scale and Reliability

Powerful & Flexible QFP Processor

- Flexible uCode Architectue for Feature Richness
- L2 + L3 ServicesL FR, PPP, HDLC, MLPPP, LFI
- L3VPN, MPLS, Netflow, 6PE/6VPE

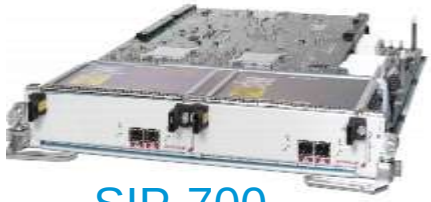

SIP-700

SPAs

SPA Support

- ChOC-3/12/48 (STM1/4/16)
- POS: OC3/STM1, OC12/STM4, OC-48/STM16, OC192/STM64
- ChT1/E1, ChT3/E3, CEoPs, ATM

ASR 9000 ISM (Integrated Service Module)

CDS Streaming:
TV and internet streaming
Error repair

CGN (carrier grade NAT):
NAT44, DS-Lite
NAT64

Feature	ASR 9000 ISM Capabilities
Applications	Ultra-Dense VoD, TV, Internet Streaming, Error Repair, CGv6
Bandwidth	30-40 Gbps streaming capacity ~3 Gbps cache fill rate
Compatibility	Works with all CDS appliances
Concurrent Streams	Up to 8,000 SD equivalent
Content Cache	3.2 TBytes at FCS - Modular Design
Video Formats	MPEG2 & AVC/H.264
Transport	MPEG over UDP / RTP
Session Protocols	RTSP / SDP
Environmental	NEBS / ETSI compliant

CDS: Manage 8,000 streams up to 40G per second
CGv6: 20M translations, 1M translations/sec., ~15Gbps throughput / ISM

ASR 9000 Optical Interface Support

All Linecards use Transceivers

Based on Density and Interface Type the Transceiver is different

- 1GE (SFP) T, SX, LX, ZX, CWDM/DWDM
- 10GE (XFP & SFP+): SR, LR, ZR, ER, DWDM
- 40GE (QSFP): SR4, LR4
- 100GE (CFP): SR10, LR4, DWDM [1)]

XFP

SFP, SFP+

QSFP

CFP

All 10G and 40G Ports do support G.709/OTN/FEC

1) Using Optical Shelf (ONS15454 M2/M6)

For latest Transceiver Support Information

http://www.cisco.com/en/US/prod/collateral/routers/ps9853/data_sheet_c78-624747.html

Carrier Class, Scalable System Architecture

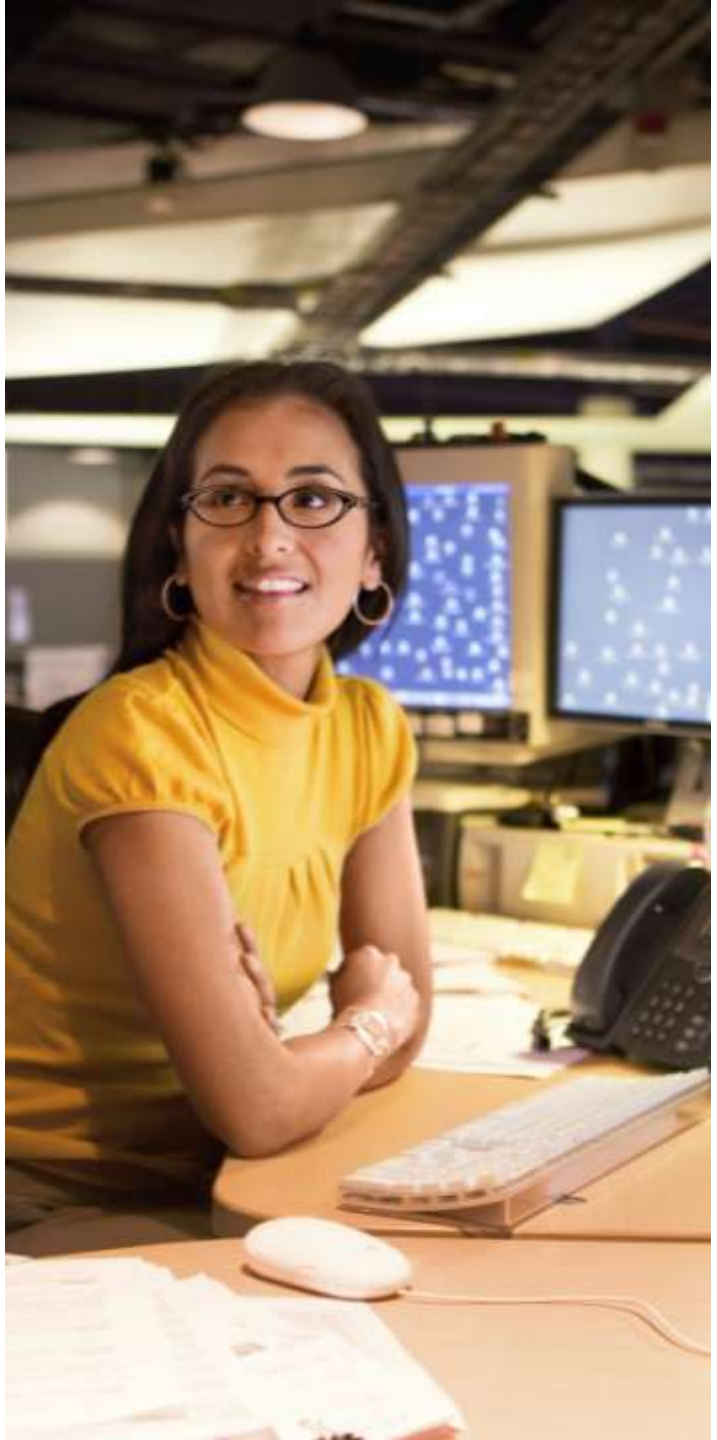

Fabric Architecture

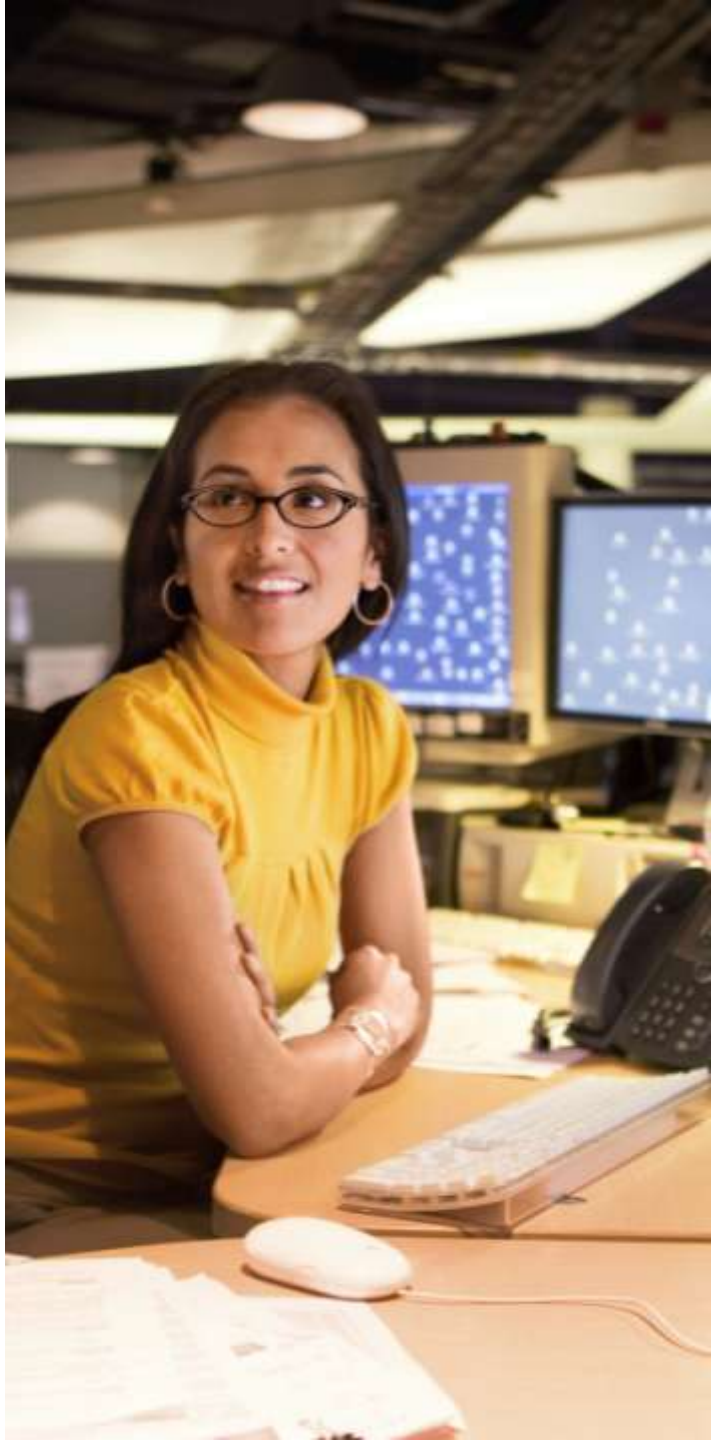

Cisco ASR9000 System Architecture

"At a Glance"

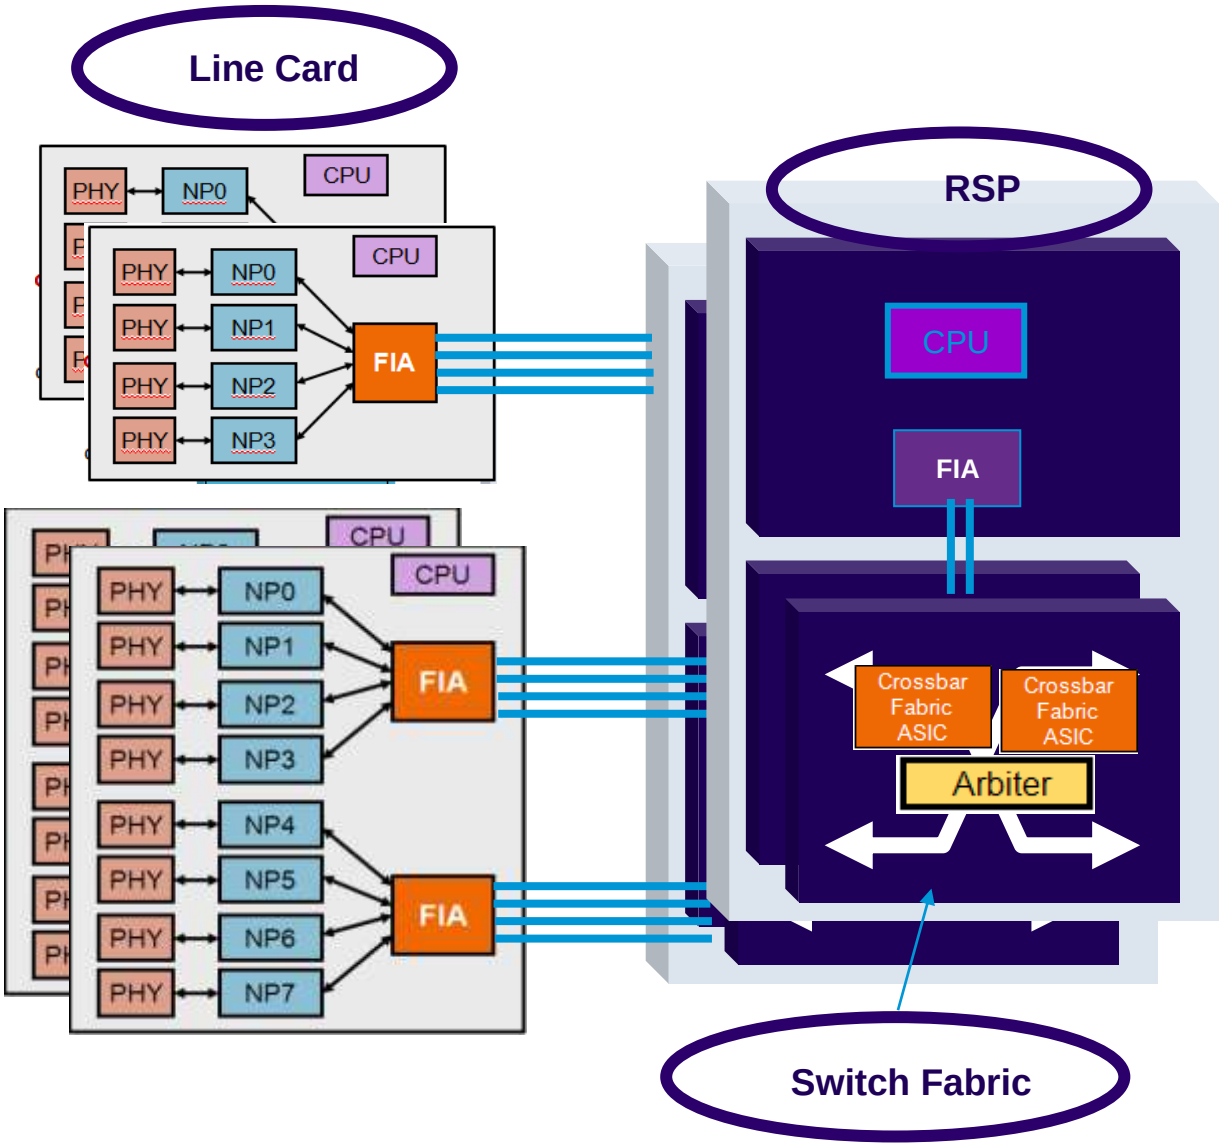

ASR9000 Switch Fabric Overview

3-Stage Fabric

Fabric frame format:
Super-frame
Fabric load balancing:
Unicast is per-packet
Multicast is per-flow

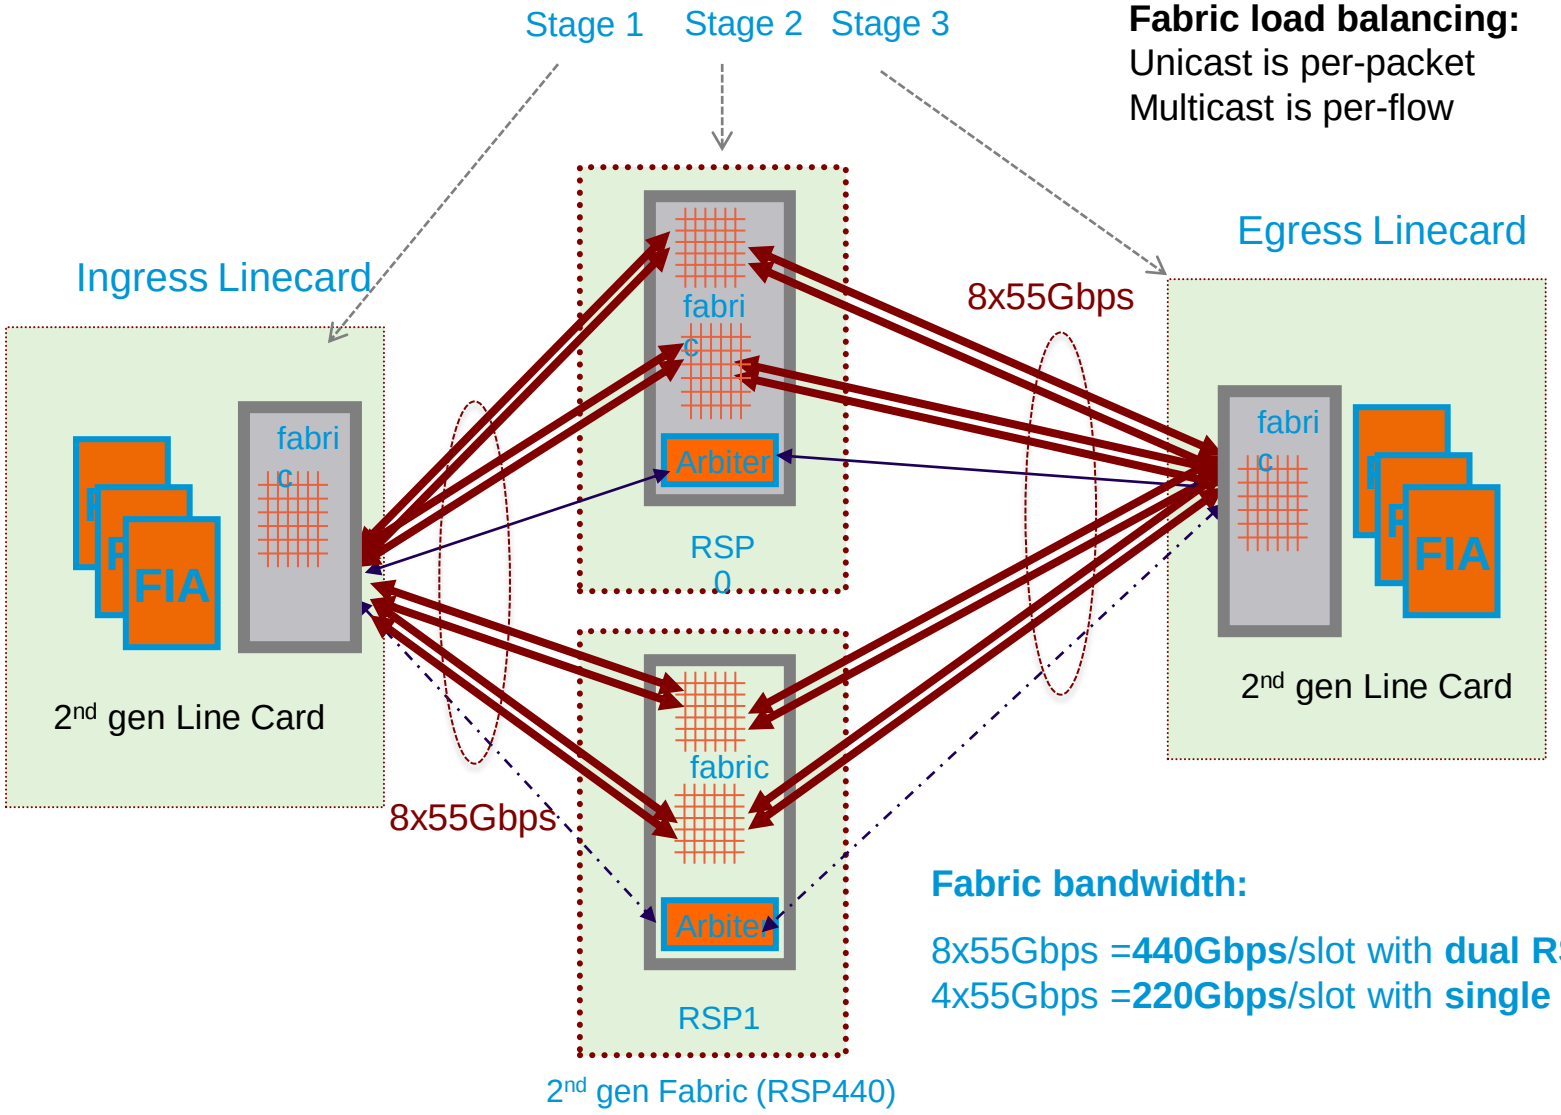

Fabric bandwidth:

8x55Gbps =**440Gbps**/slot with **dual RSP**
4x55Gbps =**220Gbps**/slot with **single RSP**

1st/2nd Generation switch fabric compatibility

System With 2nd Generation Fabric

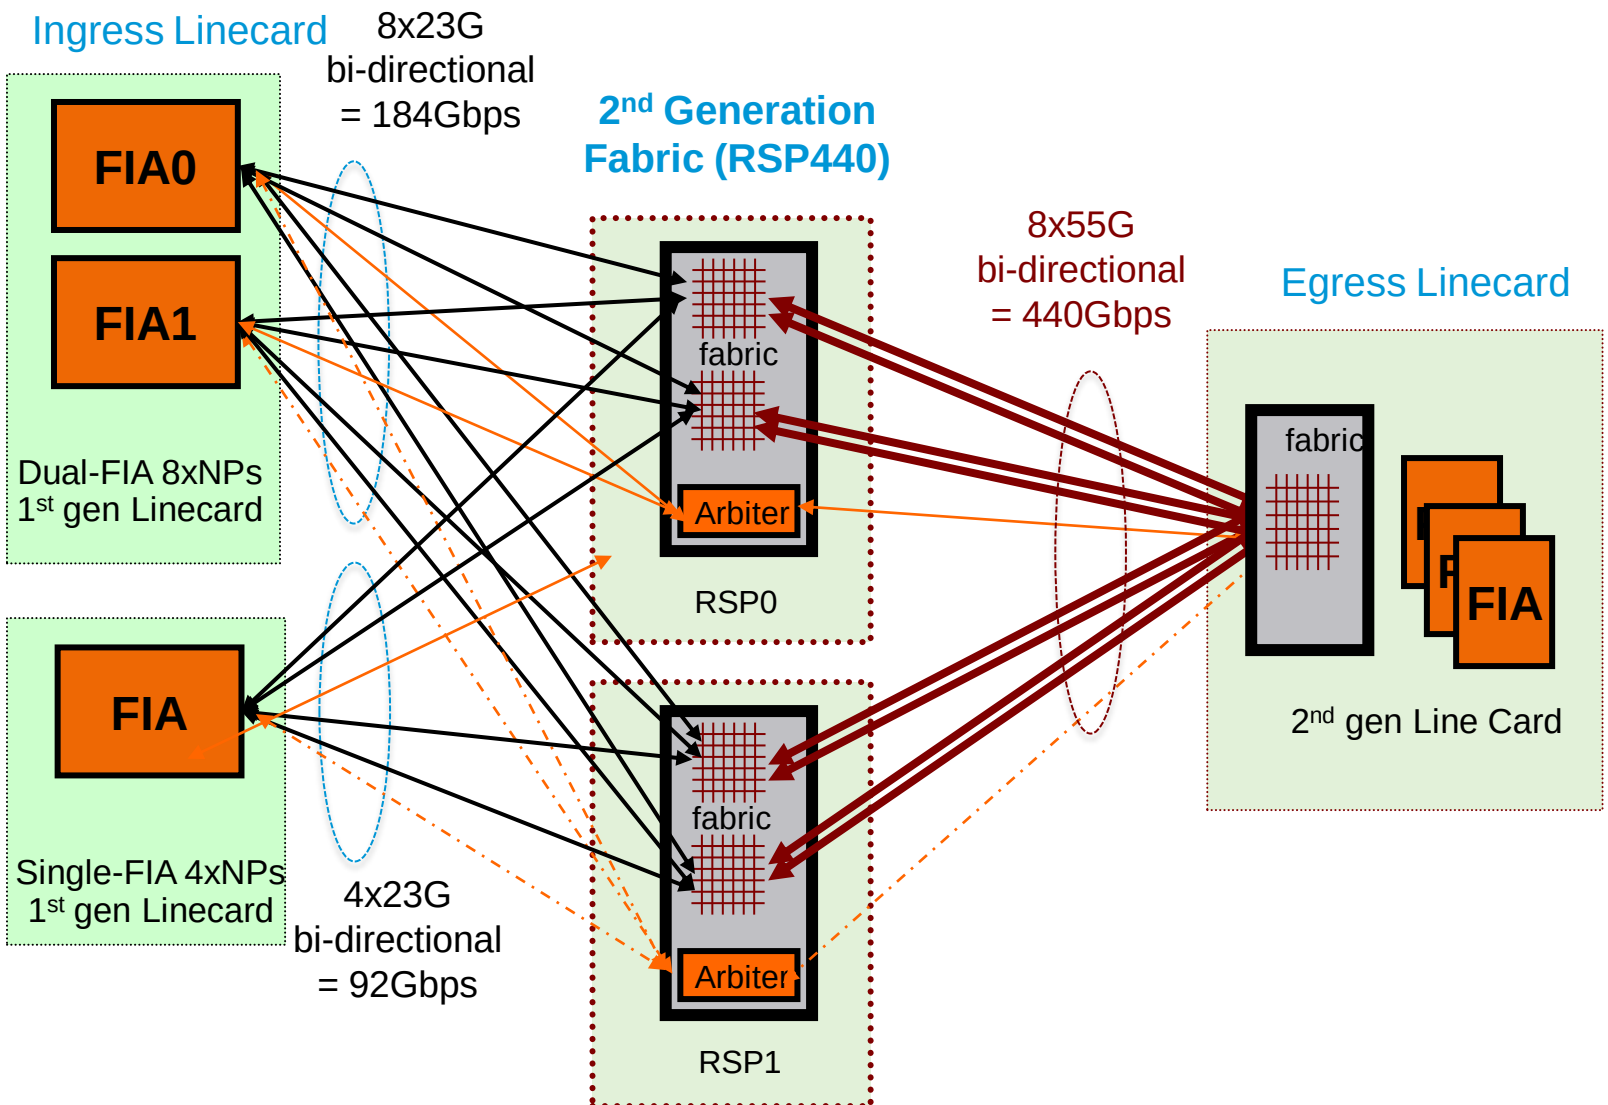

1st/2nd Generation switch fabric compatibility

System with 1st Generation Fabric

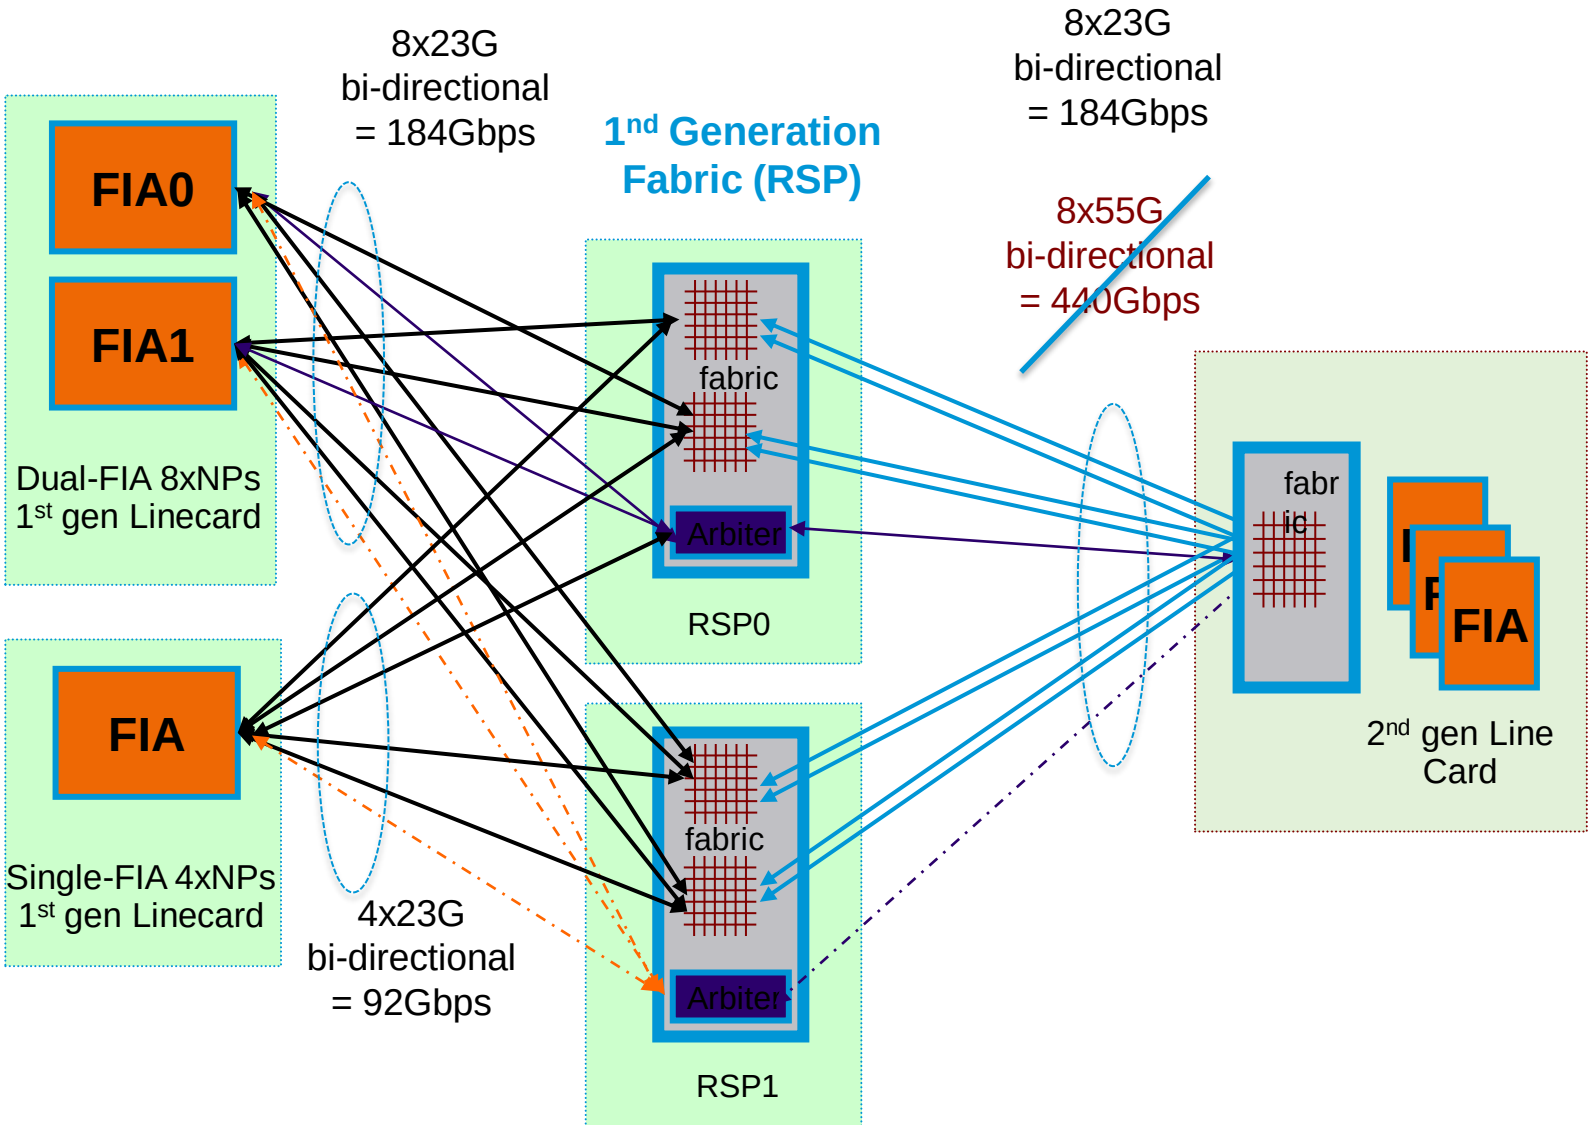

ASR 9922 Fabric Architecture : 5-plane System

Supported Today

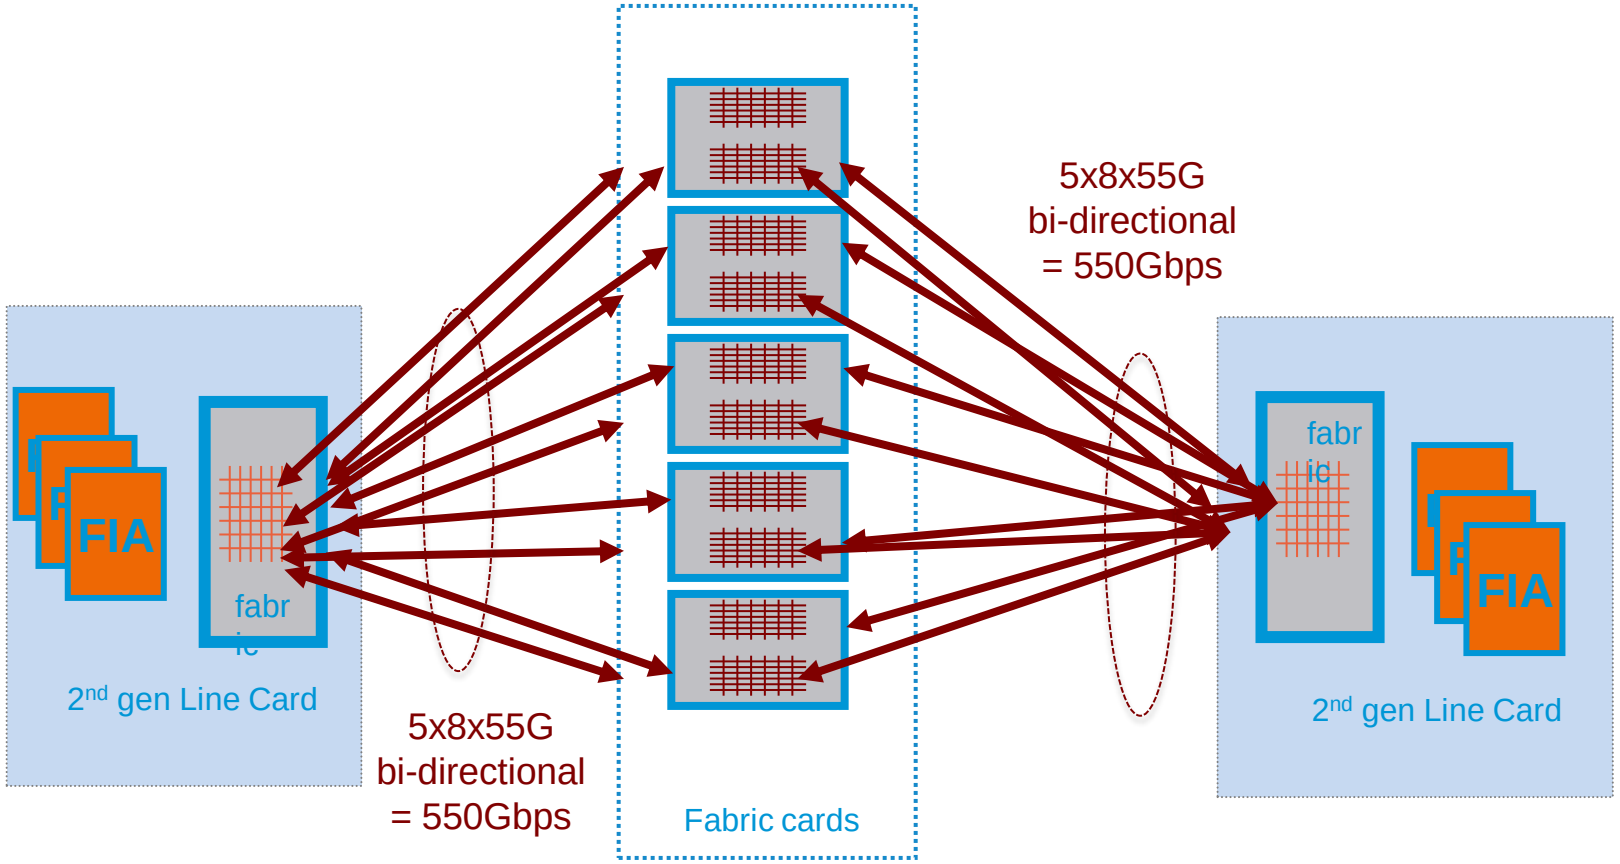

550Gbps/LC or 440Gbps/LC with fabric redundancy

ASR 9922 Fabric Architecture : 7-plane System

Not supported today !

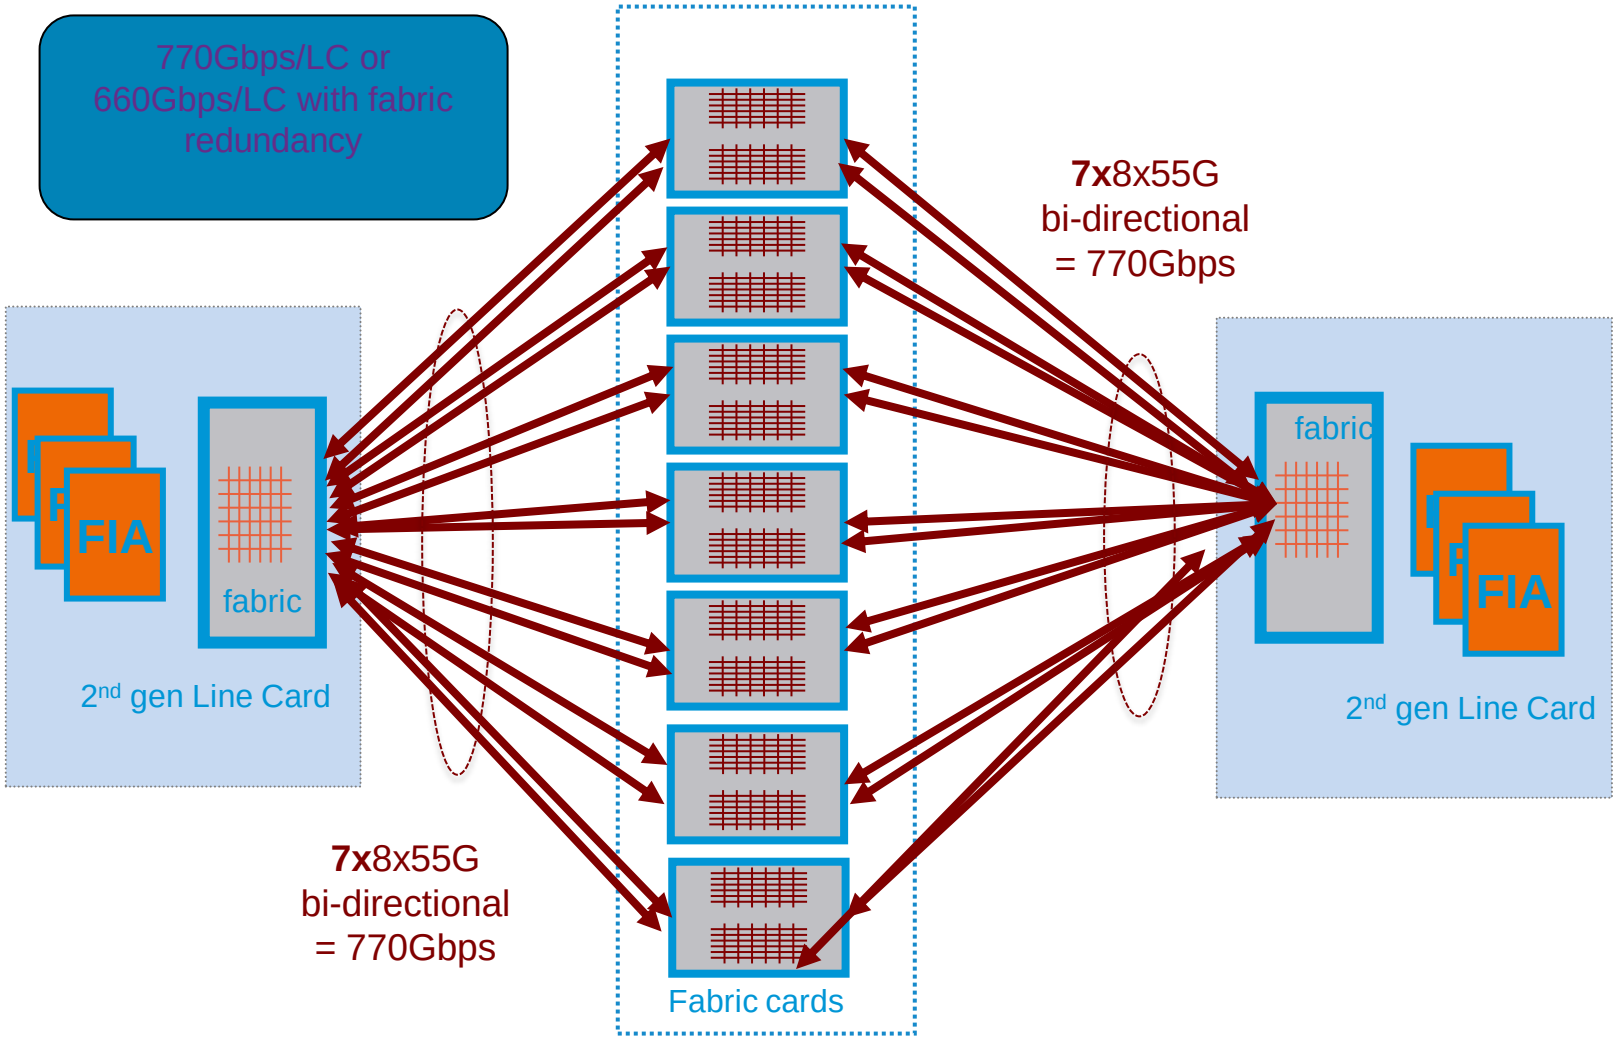

Linecard Architecture

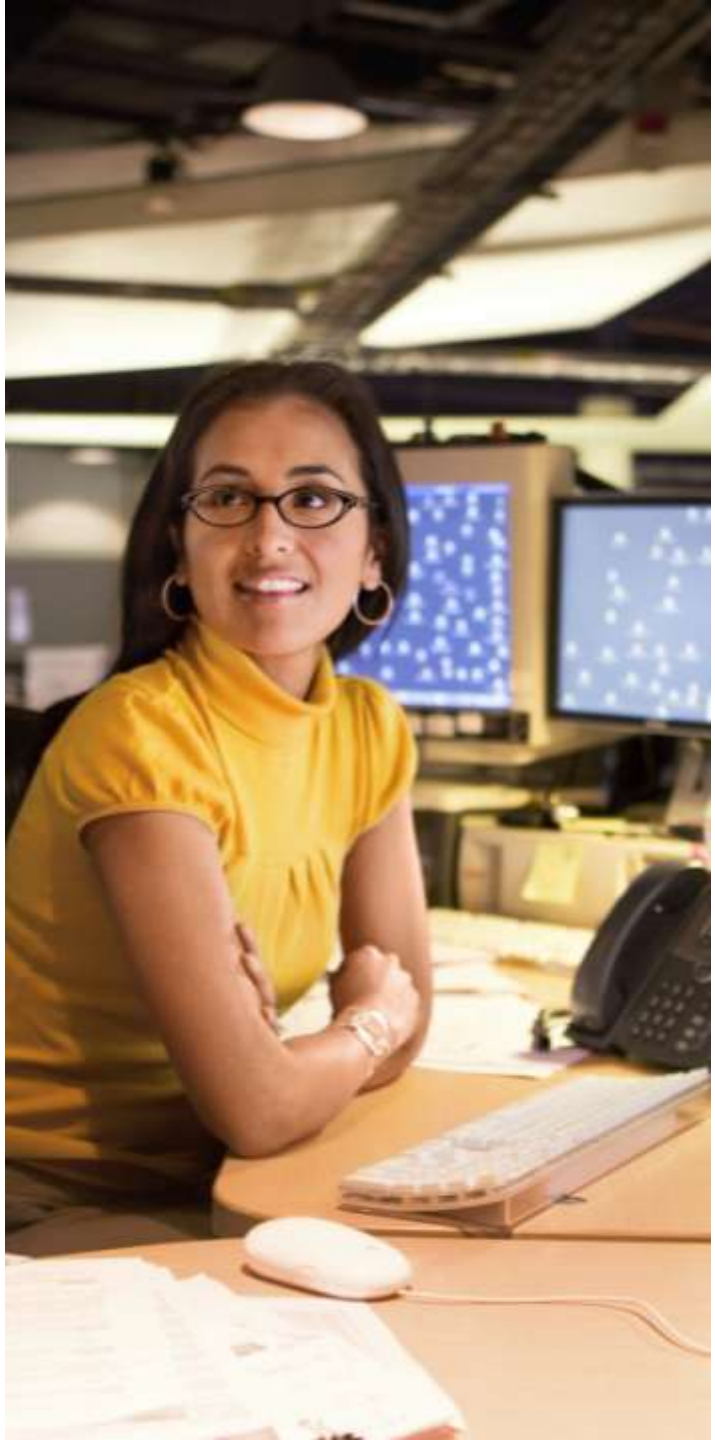

Generic Linecard Architecture – Components

Pluggable physical interfaces

- speeds: GE, 10GE, 40GE, 100GE
- form factors: SFP, SFP+, XFP, QSFP, CFP
- media/reach: T, SR, LR, ZR, LR4, SR10
- colors: gray, CWDM, DWDM, Tunable

CPU

CPU

- Distributed Control planes
- SW switched packets
- Inline Netflow
- Program HW forwarding tables

Network Processor

NP

- forwarding and feature engine for the LC
- scales bandwidth via multiple NPs
  - up to 8 NPs/LC for performance vs. density options
- highly integrated silicon as opposed to multiple discrete components
  - shorter connections, faster communication channels
  - higher performance, density with lower power draw
  - simplified software development model

Fabric Interface ASIC

FIA

- interface between forwarding processor and system switch fabric
- arbitration, framing, accounting in HW
- provides buffering and virtual output queuing for the switch fabric
- QoS awareness for Hi/Lo and ucast/mcast
  - total flexibility regarding relative priority of unicast vs. multicast

ASR 9000 Line Card Architecture Overview

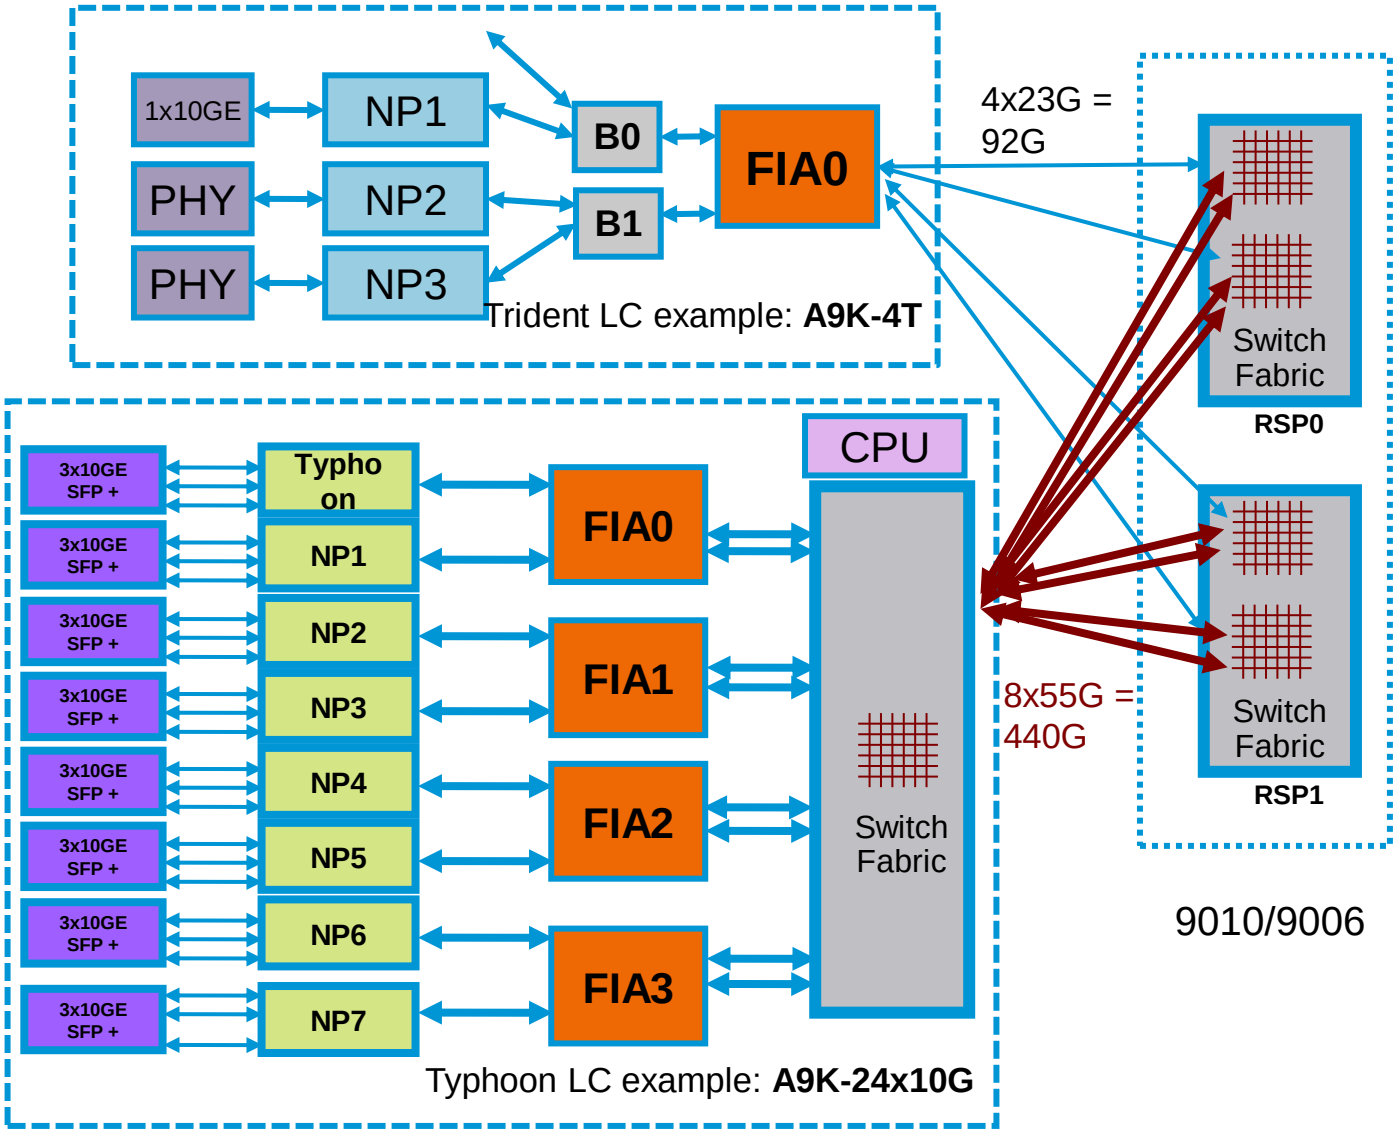

24port 10GE Linecard Architecture

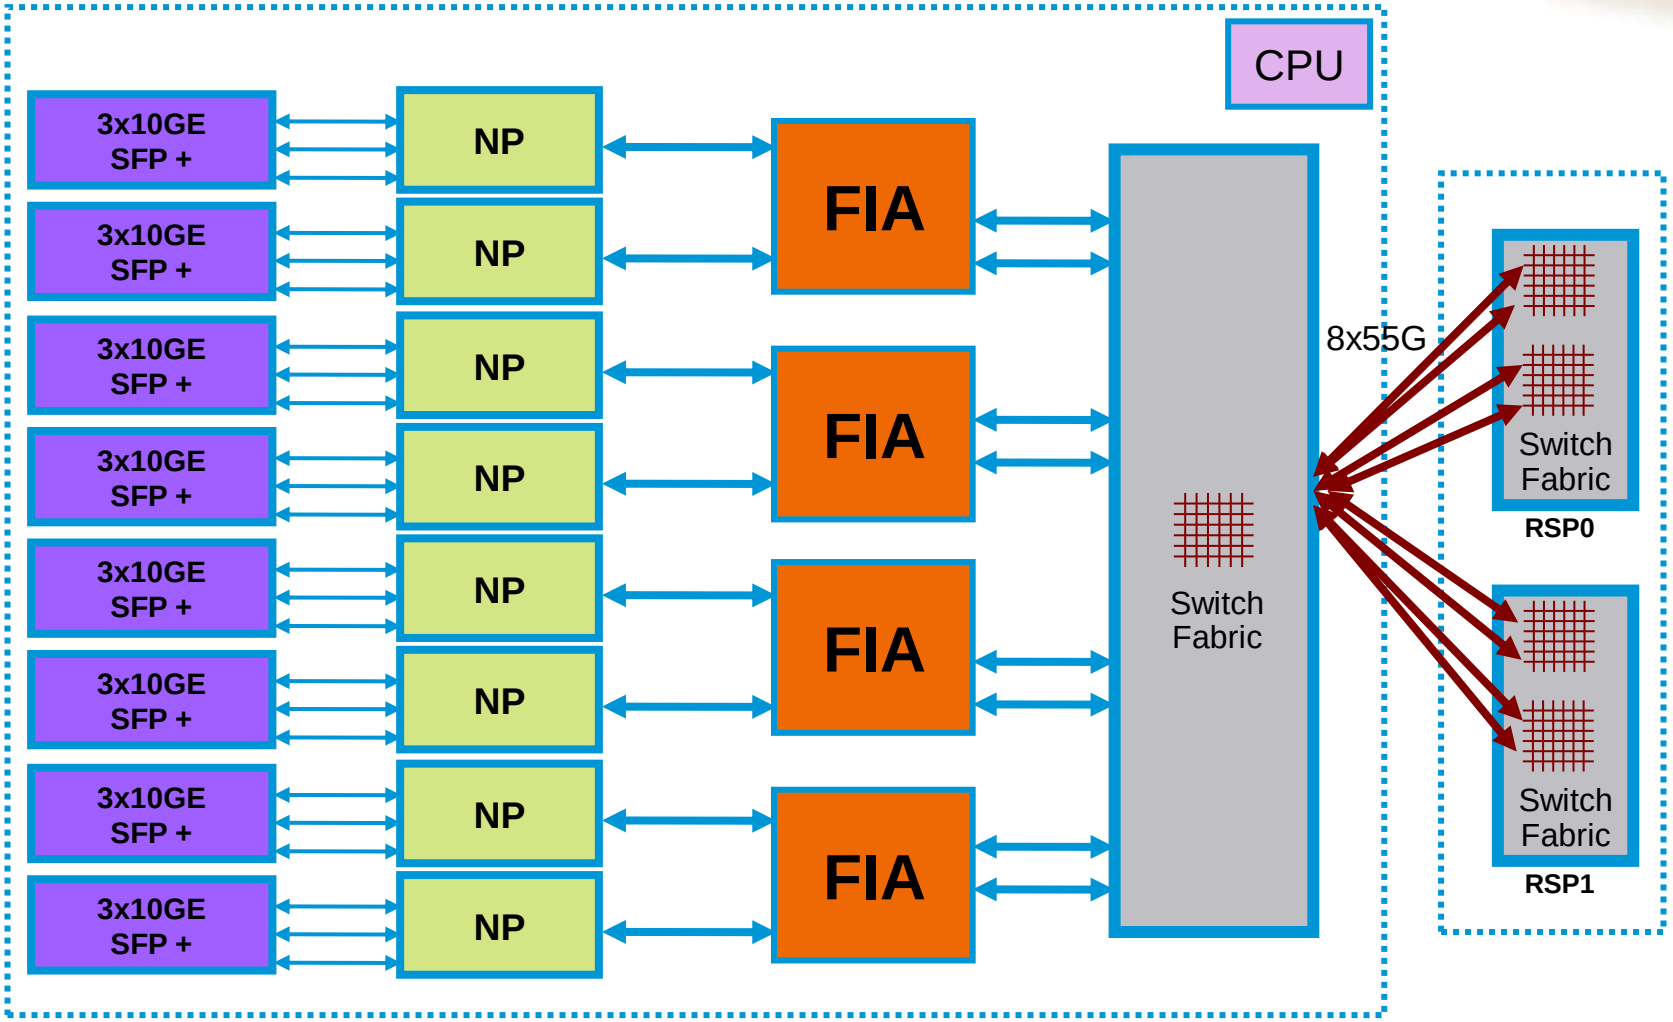

Each NP: 60Gbps bi-directional
120Gbps uni-directional
2x45Mpps

Each FIA: 60Gbps bi-directional

36port 10GE Linecard Architecture

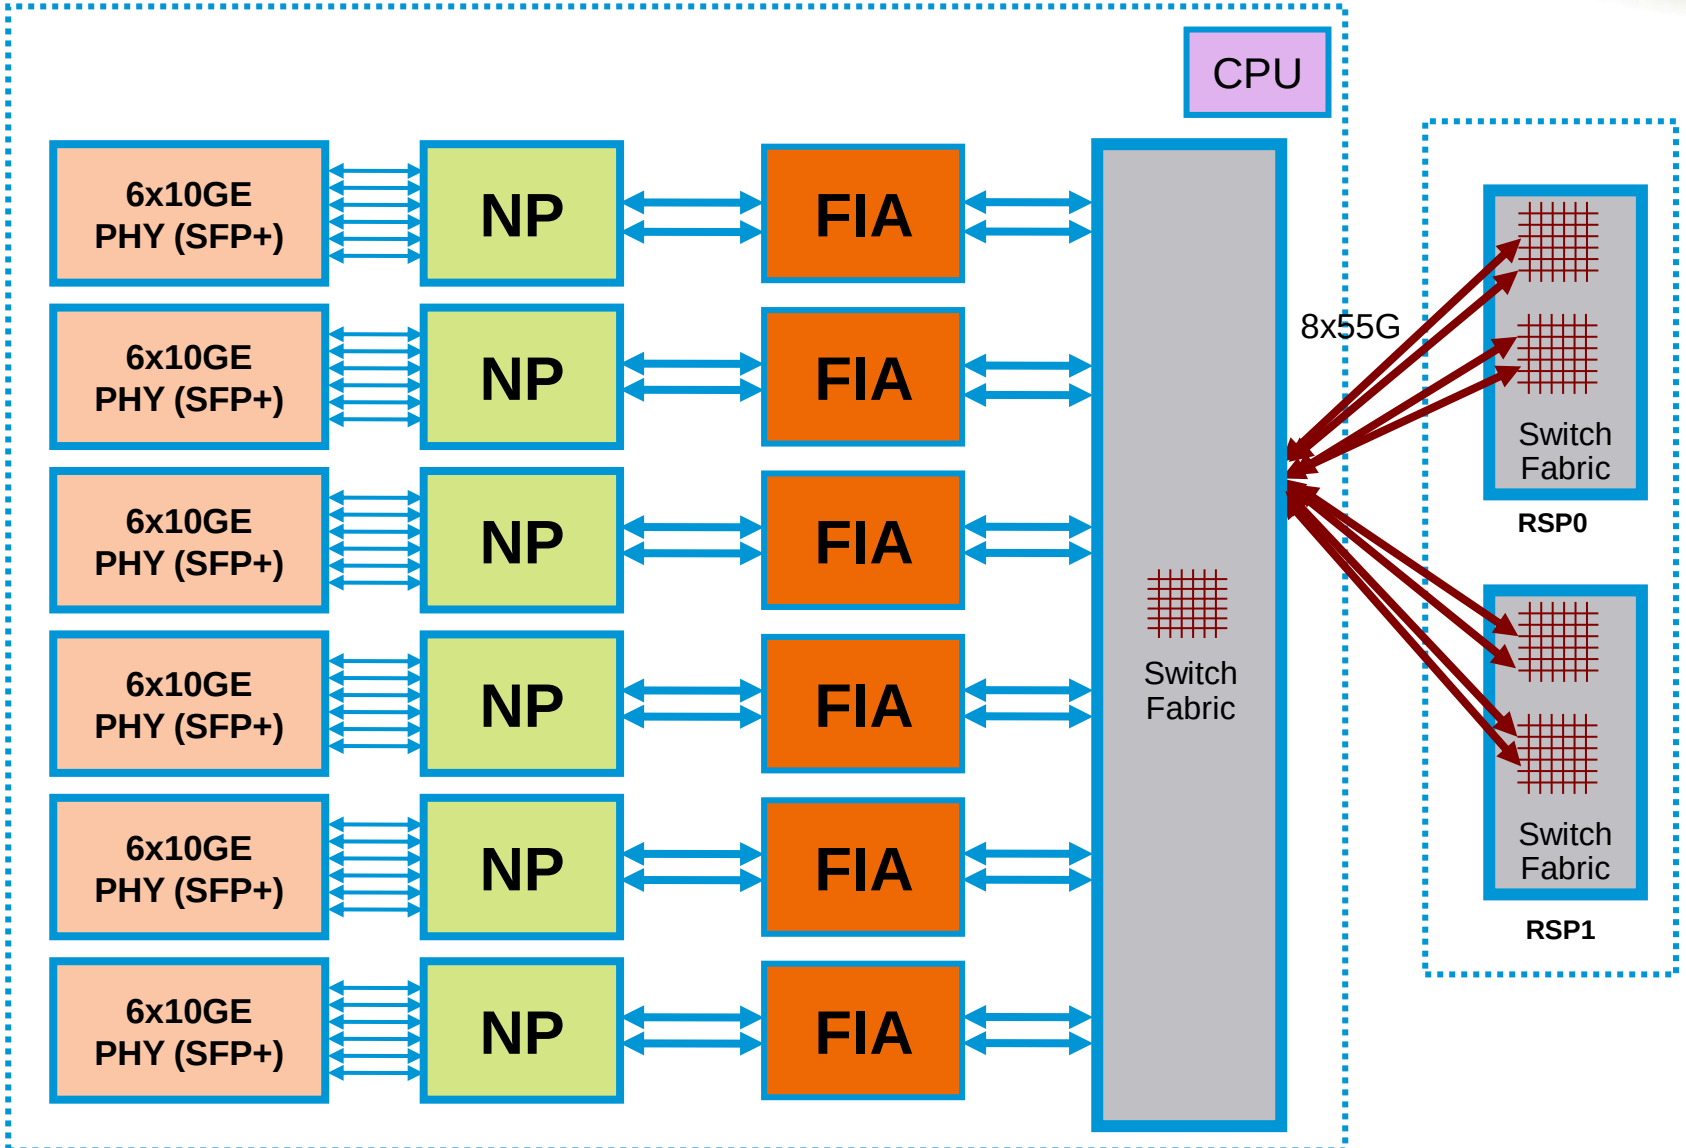

2port 100GE Linecard Architecture

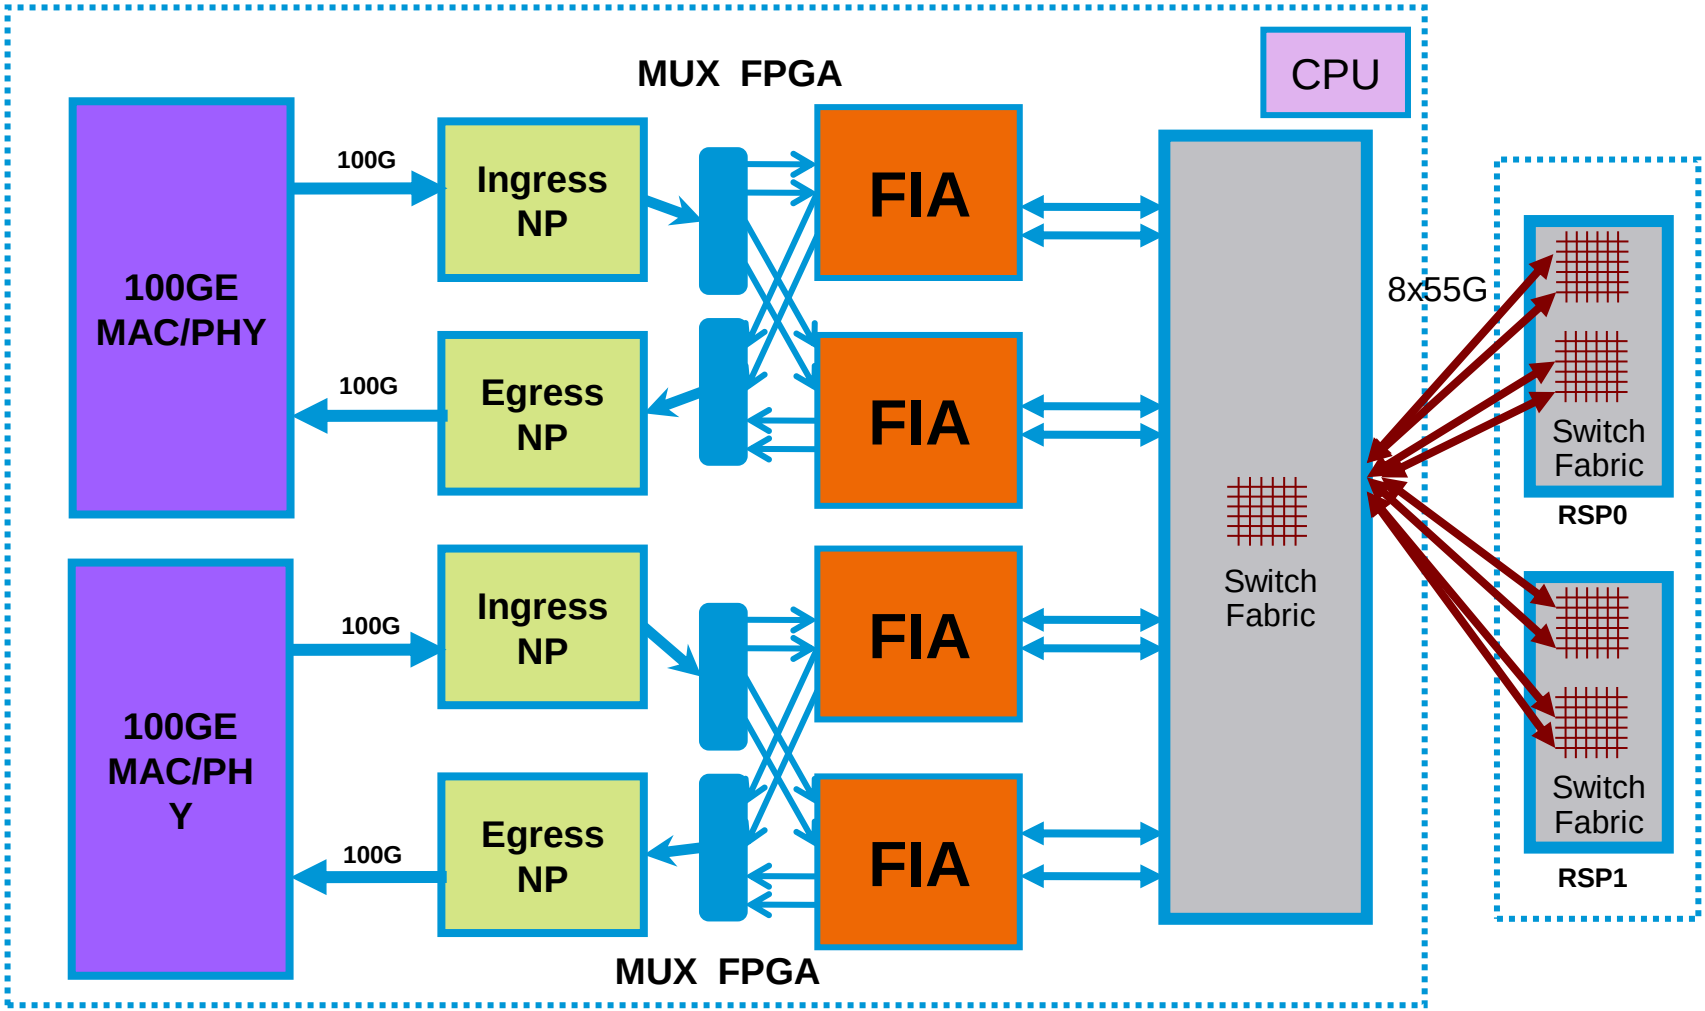

Module Cards – MOD80

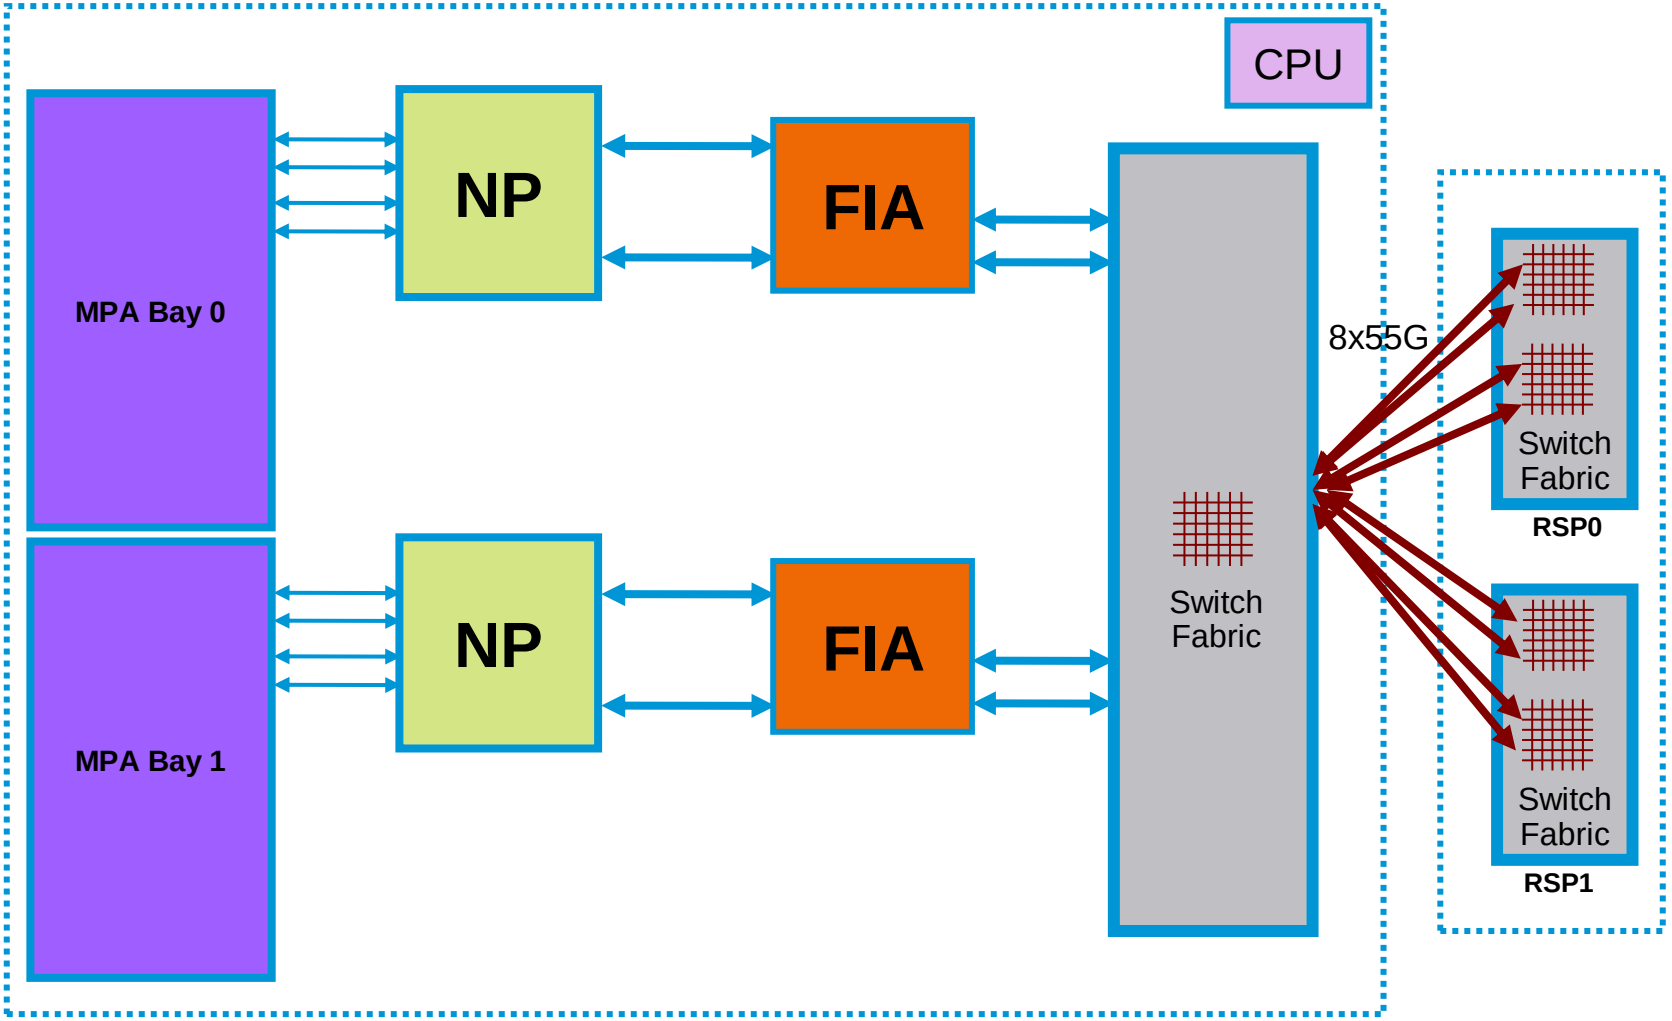

Module Cards – MOD160

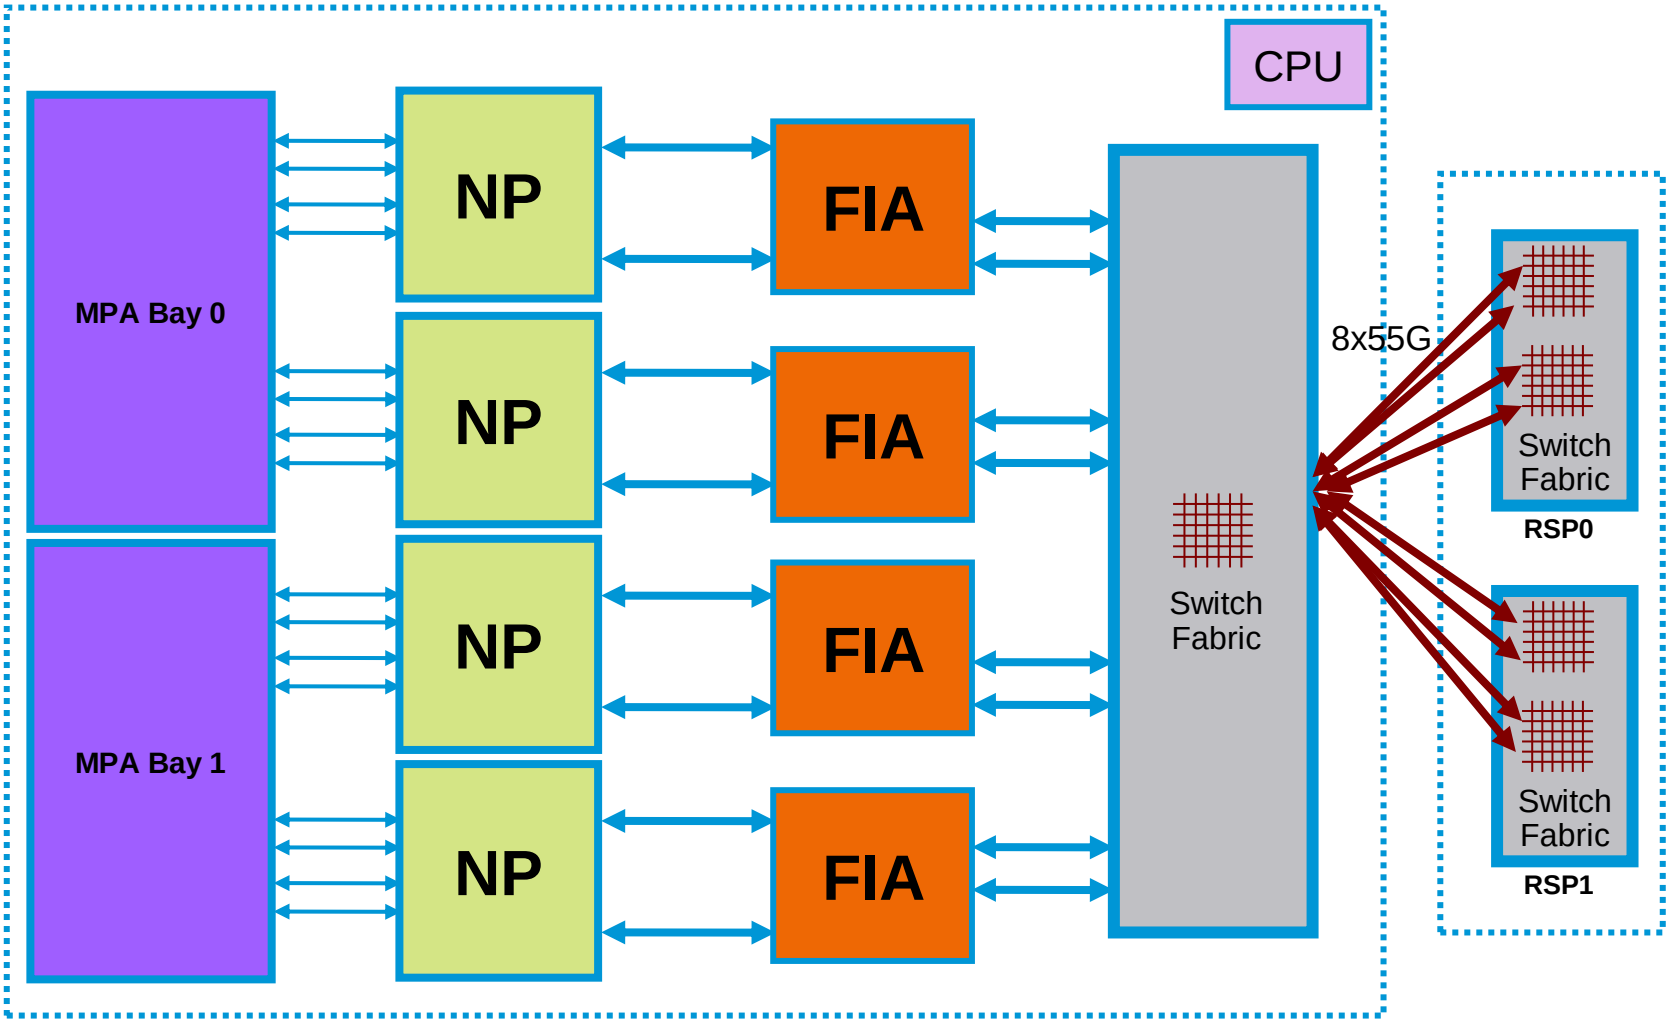

MPA Port Mapping Examples for 10GE Ports

MOD80

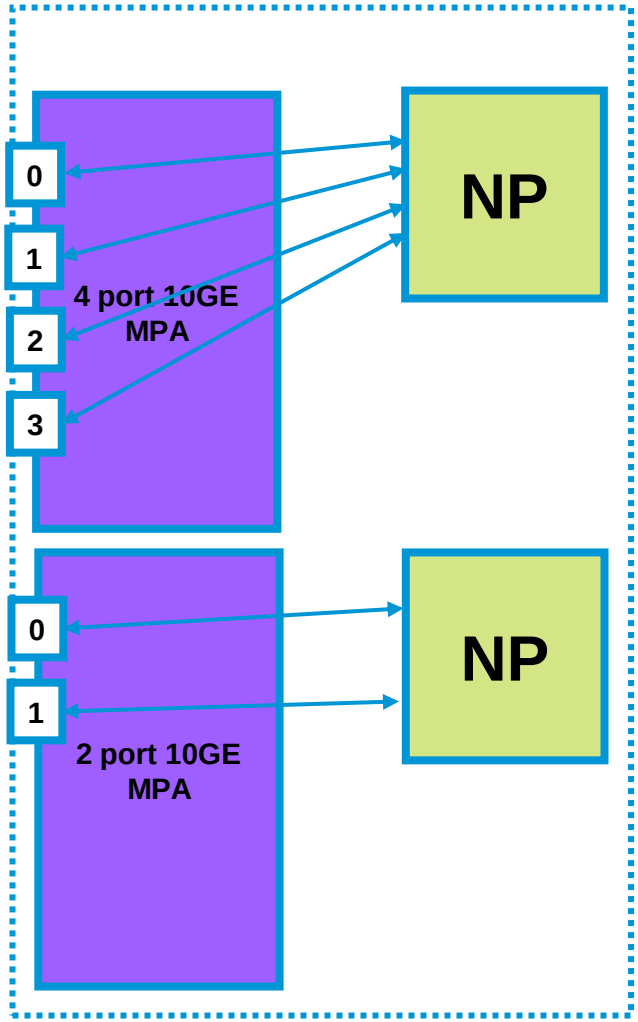

MOD160

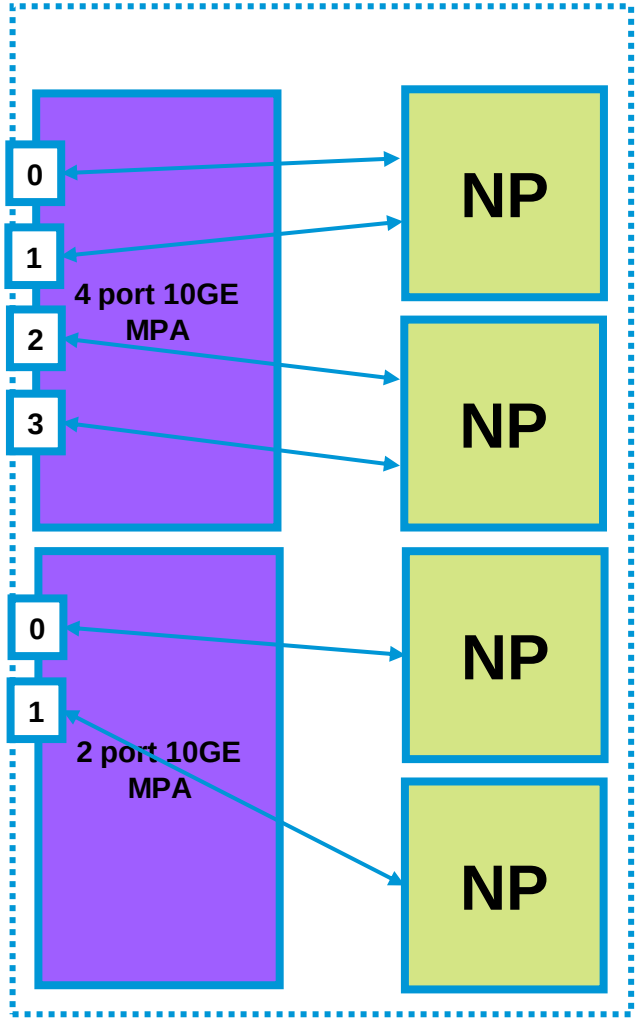

Network Processor Architecture Details

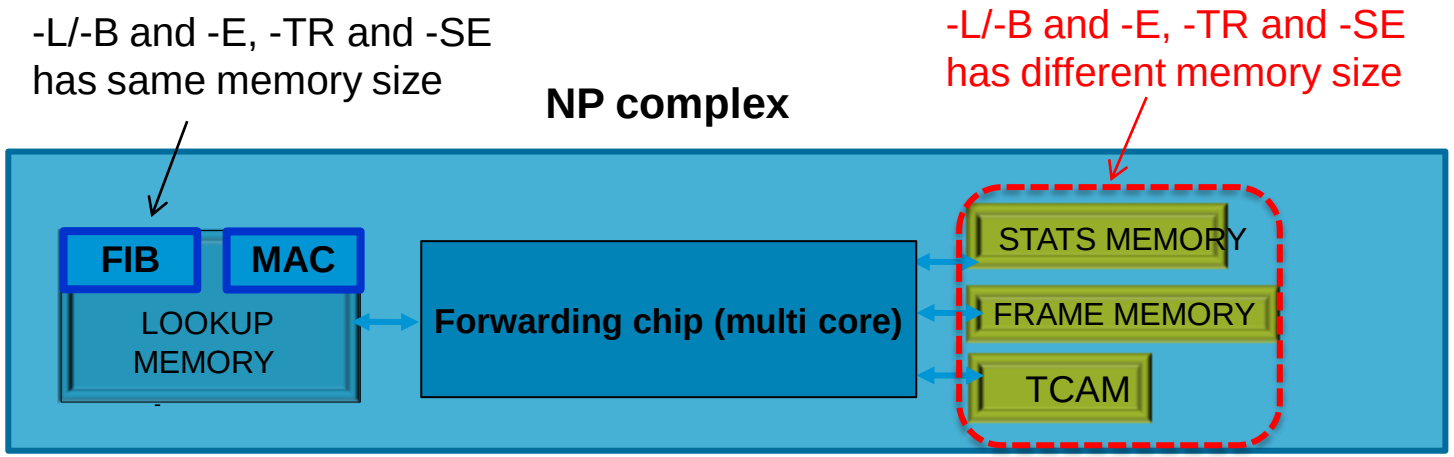

TCAM: VLAN tag, QoS and ACL classification

Stats memory: interface statistics, forwarding statistics etc

Frame memory: buffer, Queues

Lookup Memory: forwarding tables, FIB, MAC, ADJ

-TR/-SE, -L/-B/-E

- 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

ASR9001 Architecture

Identical HW Components as the Modular Systems

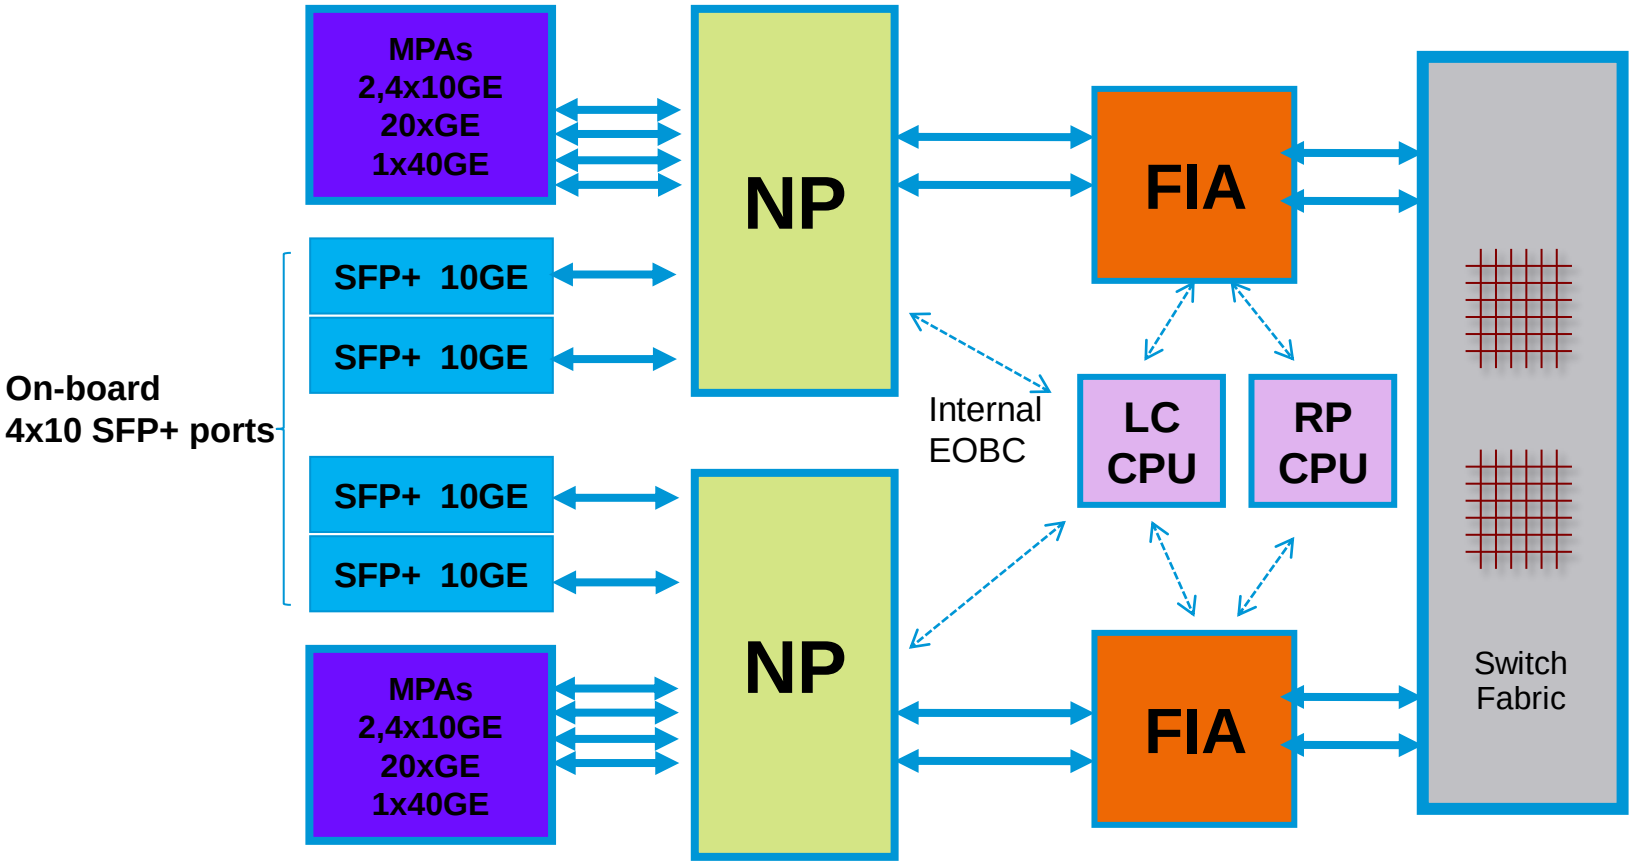

ASR 9000 nV Network Virtualization

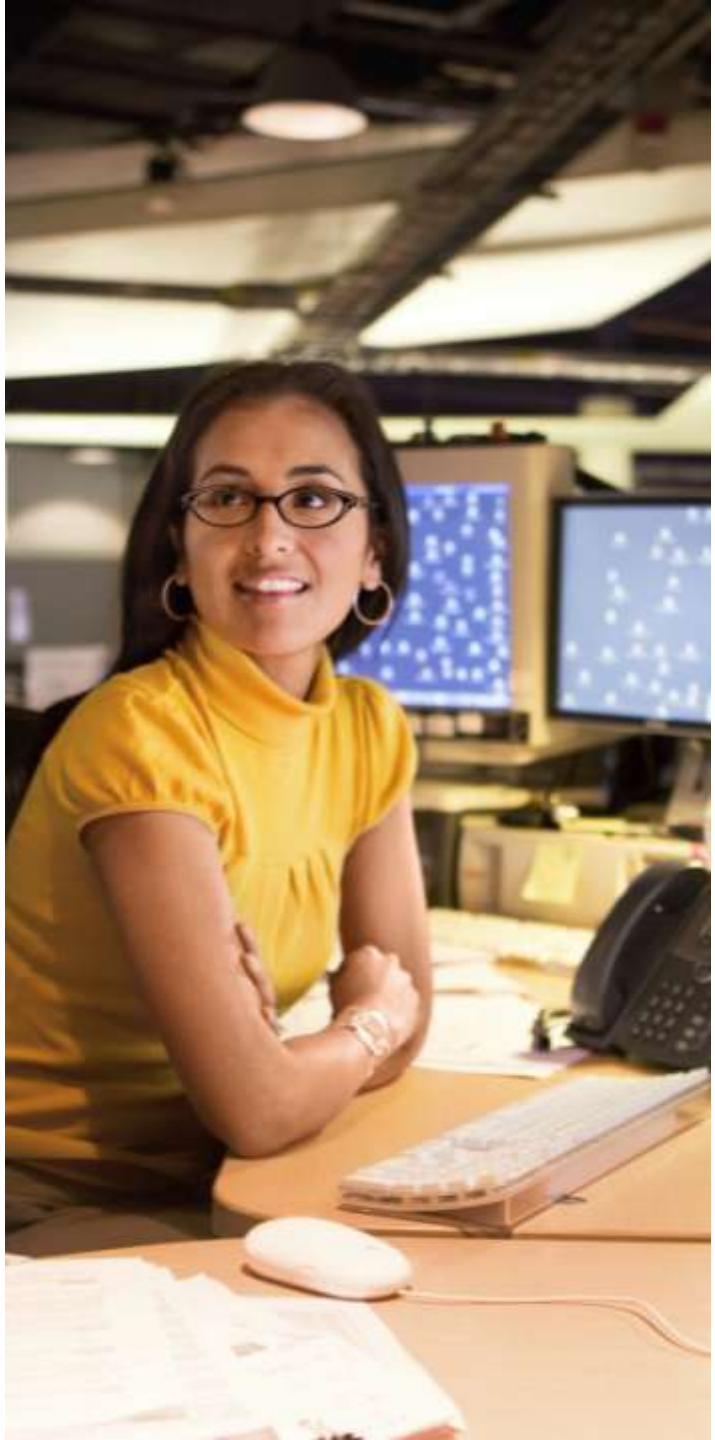

ASR 9000 nV Technology Overview

Before: ***nV*** Technology

Individual device to manage

Complex network protocols

Feature inconsistency, inter-operability

Physical port limit

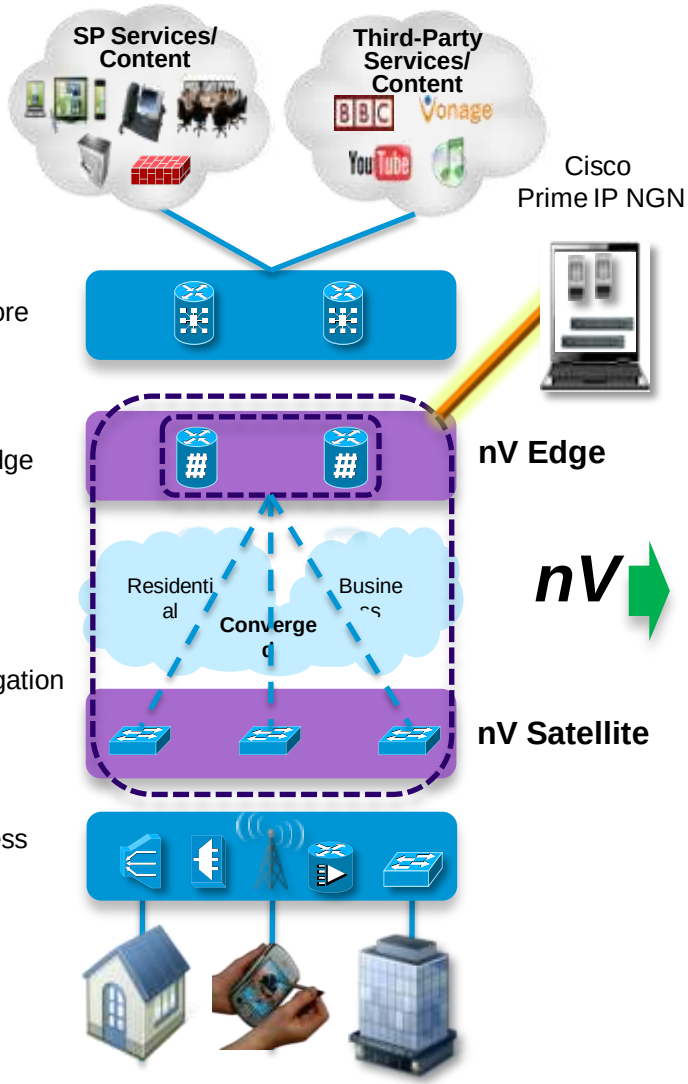

After: ***nV*** Technology

One virtual system to manage

No network protocols within virtual system, Remote satellite is plug-n-play, zero touch

Single feature set, one release cycle

Scale to 10 of 1000s physical ports

nV Satellite

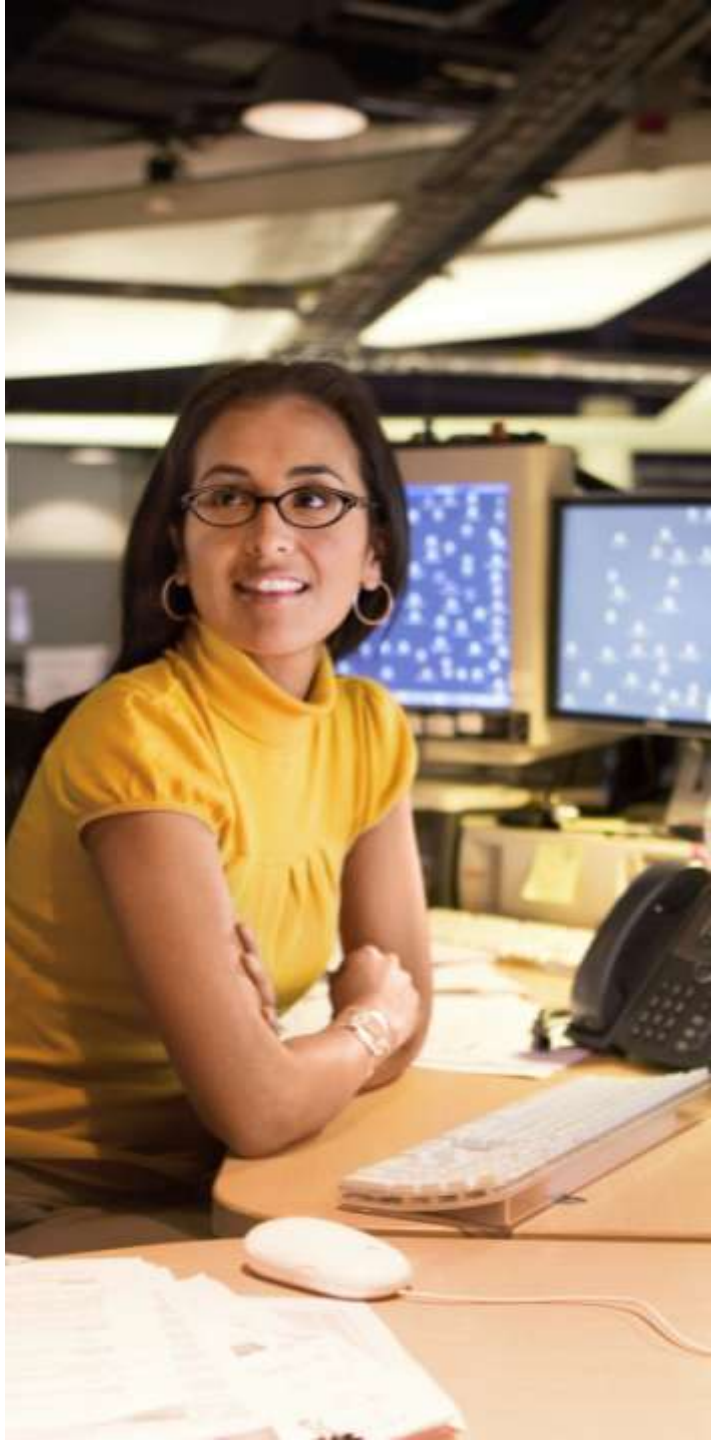

ASR 9000 nV Satellite Overview

Plug and Play, zero Touch Satellite Access Device

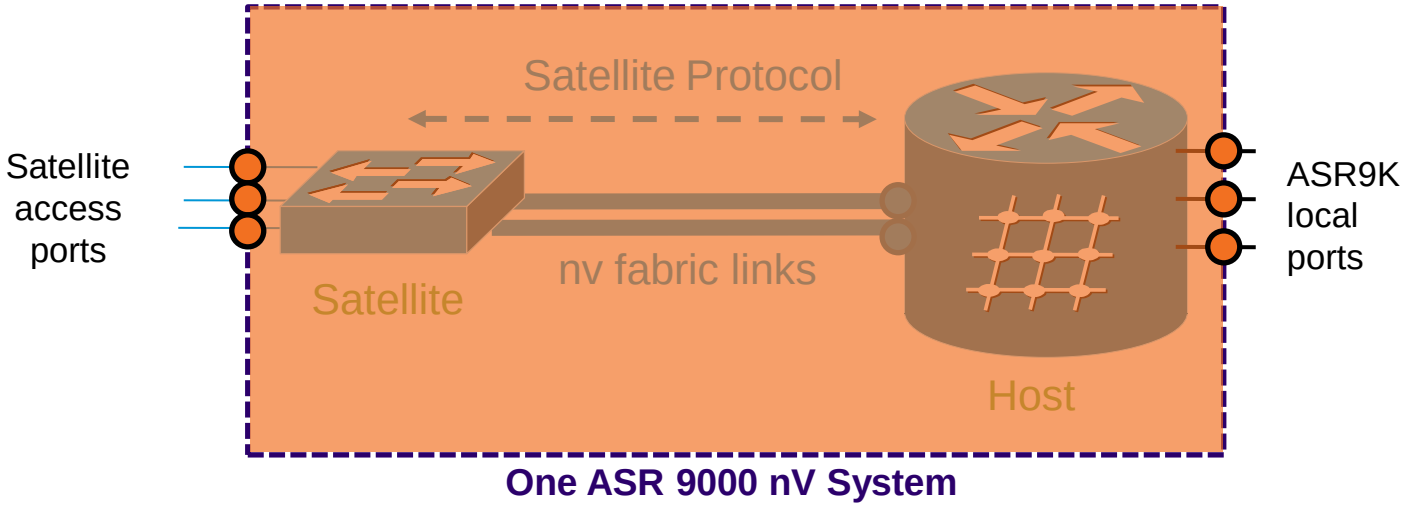

Satellite and ASR 9000 Host run satellite protocol for auto-discovery, provisioning and management

Satellite and Host could be co-located or in different location. There is no distance limitation between satellite and Host

The connection between satellite and host is called “nv fabric link”, which could be L1 or over L2 virtual circuit (future)

Satellite access port have feature parity with ASR9K local ports
→ it works/feels just as local port

Satellite Hardware – ASR 9000v Overview

Power Feeds

- Redundant -48vDC Power Feeds
- Single AC power feed
- Max Power 210W
- Nominal Power 159 W

1 RU ANSI & ETSI Compliant

Field Replaceable Fan Tray

- Redundant Fans
- ToD/PSS Output
- Bits Out

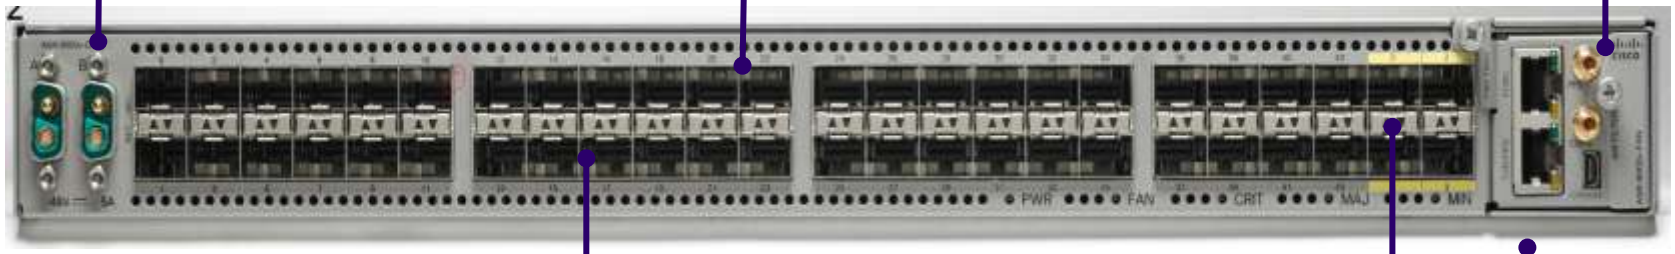

44x10/100/1000 Mbps Pluggables

- Full Line Rate Packet Processing and Traffic Management
- **Copper and fiber SFP optics**
- **Speed/duplex auto negotiation**

4x10G SFP+

- Initially used as Fabric Ports ONLY (could be used as access port in the future)
- **Copper and fiber SFP+ optics**

Industrial Temp Rated

- -40C to +65C Operational Temperature
- -40C to +70C Storage Temperature

SFP Support Information
http://www.cisco.com/en/US/docs/optical/cpt/pluggables/guide/b_cpt_pluggables.html#reference_A3D61C5FFC0A471BB39C3635CDC05E95

Satellite Hardware – ASR901 Overview

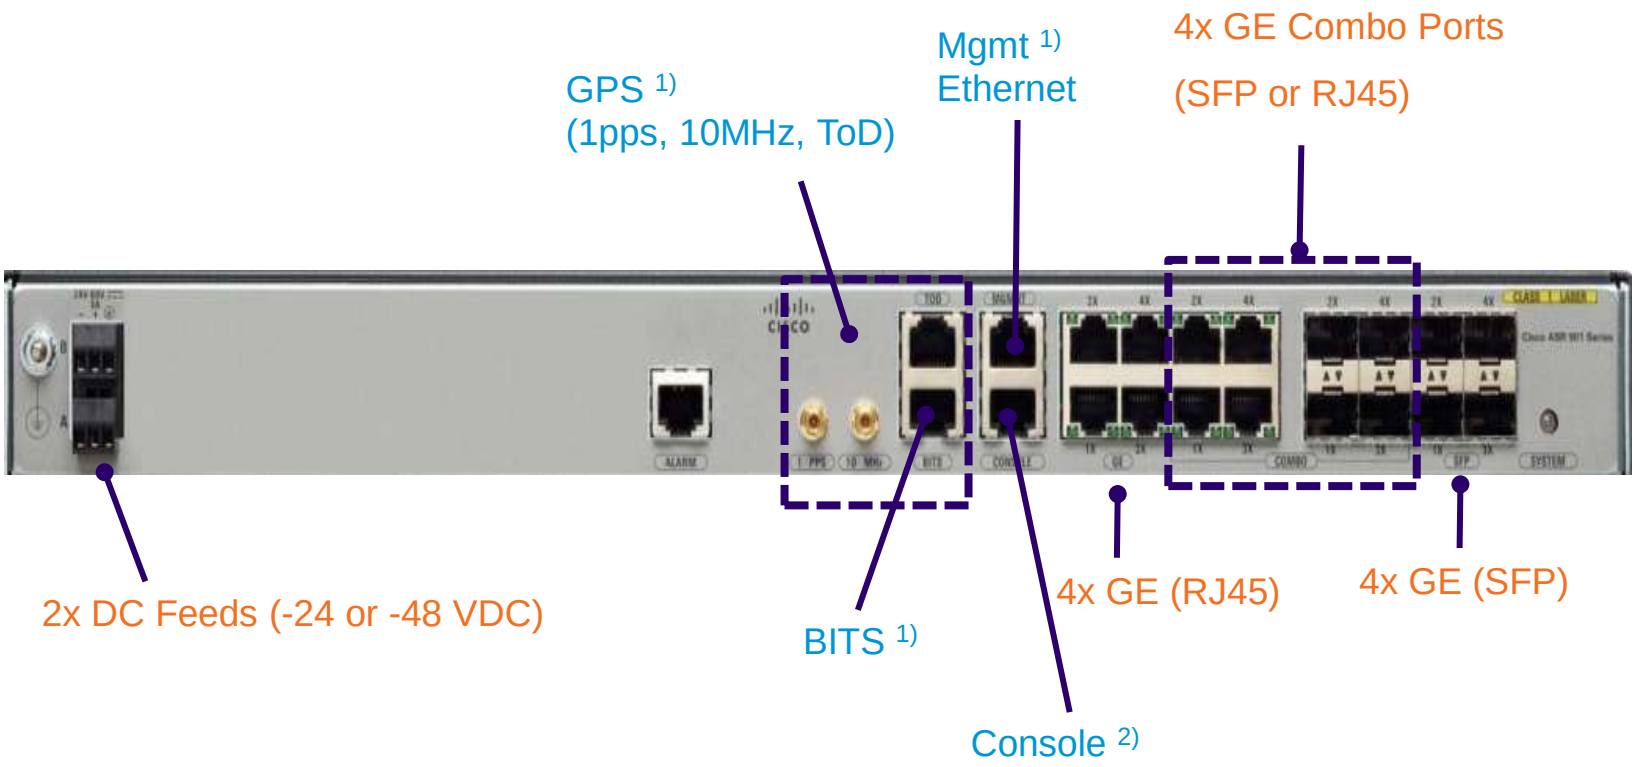

1) Not supported/used when operating in nV Satellite Mode
2) Used for low level debugging only

Satellite Hardware – ASR903 Overview

Router Switch Processor

- Currently only 1x RSP supported

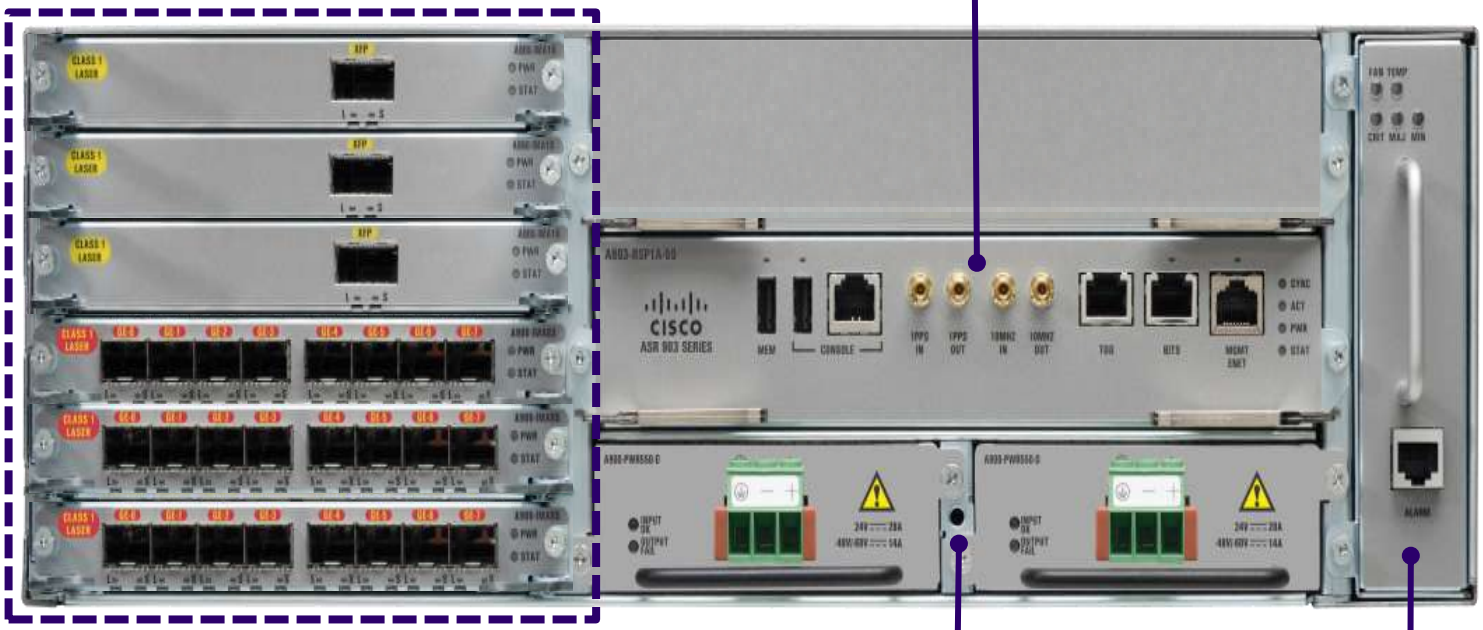

Six I/O Modules

- 1 port 10GE Module (XFP) – nV fabric links only
- 8 port 1GE Module (SFP) – access ports only
- 8 port 1GE Module (RJ45) – access ports only

2x Power Modules

- DC PEM, 1x -24 or -48 VDC
- AC PEM, 1x 115..230 VAC

Fan Module

Satellite – Host Control Plane

Satellite discovery and control protocol

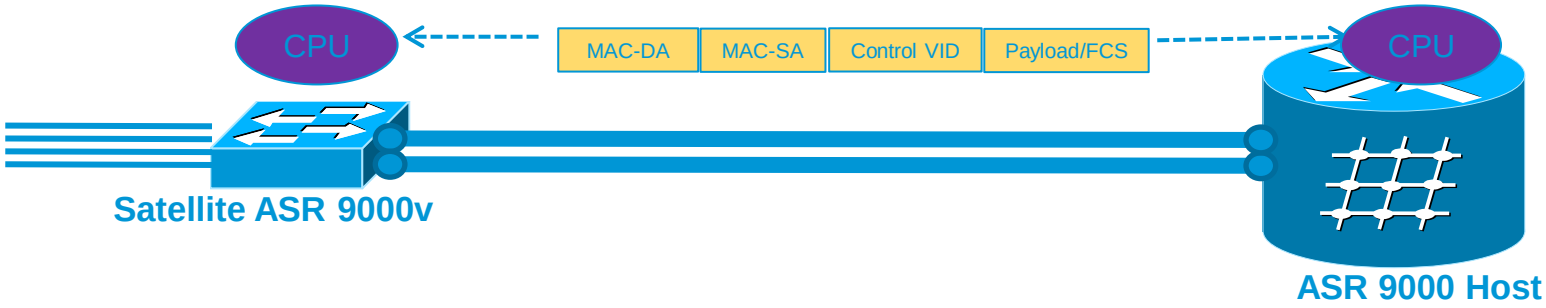

Discovery Phase

A **CDP-like** link-level protocol that discovers satellites and maintains a periodic heartbeat

Heartbeat sent once every second, used to detect satellite or fabric link failures. CFM based fast failure detection plan for future release

Control Phase

Used for **Inter-Process Communication** between Host and Satellite

Cisco proprietary protocol over TCP socket, it could get standardized in the future

Get/Set style messages to provision the satellites and also to retrieve notifications from the satellite

Satellite – Host Data Plane Encapsulation

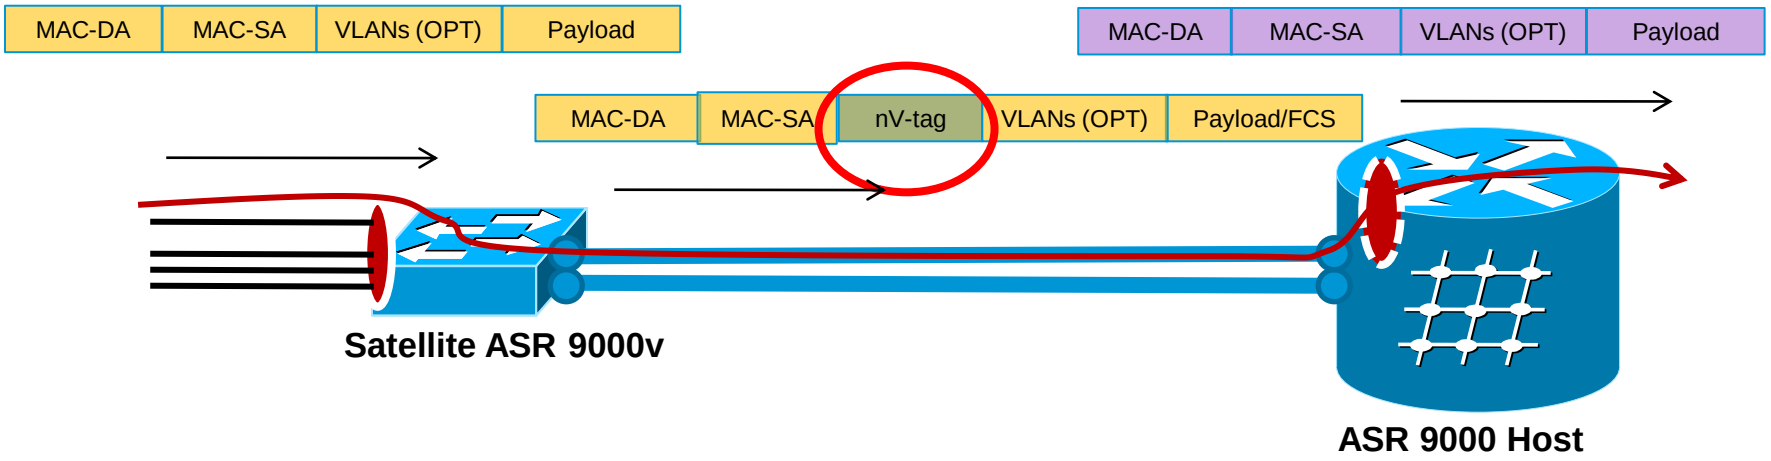

On the Satellite

Satellite receives Ethernet frame on its access port

Special nV-tag is added

Local xconnect between access and fabric port (no MAC learning !)

Packet is put into fabric port egress queue and transmitted out toward host

On the Host

- Host receives the packet on its satellite fabric port
- **Checks the nV tag**, then maps the frame to the corresponding satellite virtual access port
- Packet Processing identical to local ports (L2/L3 features, qos, ACL, etc all done in the NPU)
- Packet is forwarded out of a local, or satellite fabric port to same or different satellite

Initial Satellite Configuration

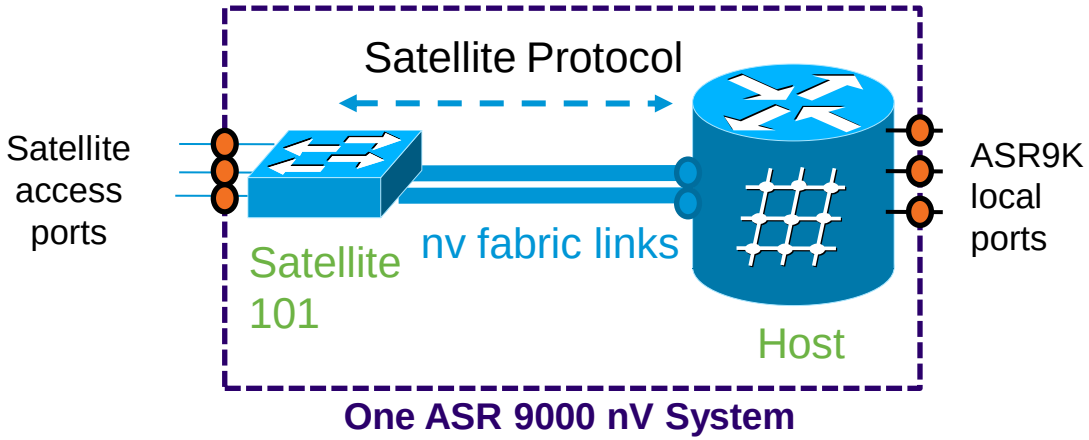


```
nv
 satellite 101 ← define satellite
  type asr9000v
  ipv4 address 10.0.0.101


interface TenGigE 0/2/0/2 ← configure satellite
fabric port
  ipv4 point-to-point
  ipv4 unnumbered Loopback100
  nv
    satellite-fabric-link satellite 101
     remote-ports ← satellite to fabric port mapping
       GigabitEthernet 0/0/0-9
```

Satellite Port Configuration

Comparison to Local Port Configuration

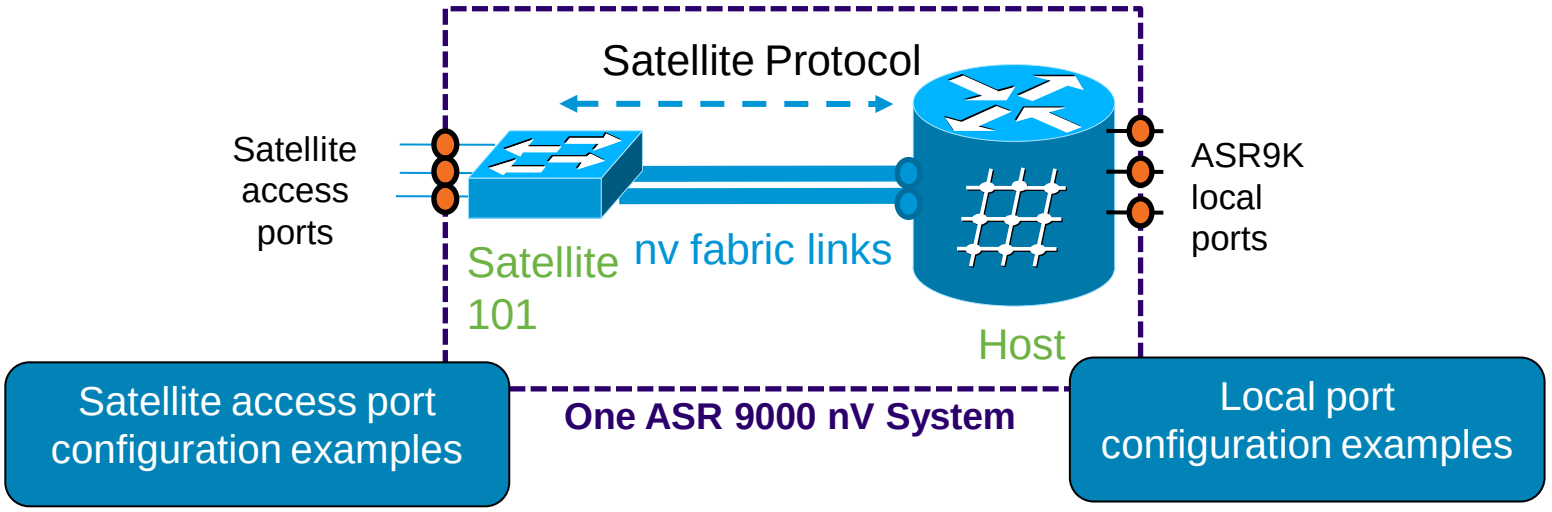


```
interface GigabitEthernet 101/0/0/1
  ipv4 address 1.2.2.2 255.255.255.0

interface TenGig 101/0/0/1.1
  encapsulation dot1q 101
  rewrite ingress tag pop 1 sym

interface Bundle-ethernet 200
  ipv4 address 1.1.1.1 255.255.255.0

interface GigabitEthernet 101/0/0/2
  bundle-id 200
```

```
interface GigabitEthernet 0/0/0/1
  ipv4 address 2.2.2.2 255.255.255.0

interface TenGig 0/0/0/1.1
  encapsulation dot1q 101
  rewrite ingress tag pop 1 sym

interface Bundle-ethernet 100
  ipv4 address 1.1.1.1 255.255.255.0

interface GigabitEthernet 0/0/0/2
  bundle-id 100
```

Satellite Deployment Models

ASR9000v Example

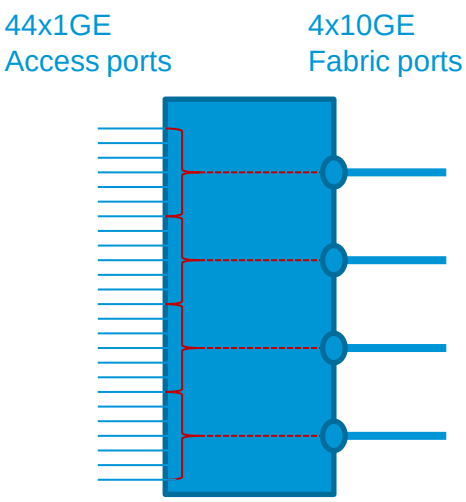

Mode 1: Static pinning
No fabric port redundancy

It can mix model 1 and 2 on the same satellite

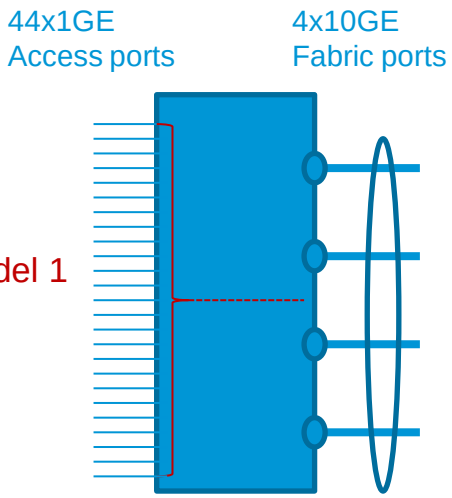

Mode 2: Fabric bundle
Fabric port redundancy

Access ports are mapped to a single Fabric Link

Fabric Link failure does bring Access Port down

Fabric links are forming a Link-Bundle

Access port traffic is “hashed” across Bundle Members

Fabric link failure keeps all Access Ports up, re-hashing of Traffic

Satellite Plug and Play

Configure, Install and Ready-to-Go

- Critical Error LED ON → bad hardware, RMA
- Major Error LED ON → Unable to connect to ASR9K host
  - Missing the initial satellite configuration?
  - L1 issue, at least one of the uplink port light green?
  - Security check (optional), is the satellite SN# correct?
- Status light green → ready to go, satellite is fully managed by Host

nV Edge

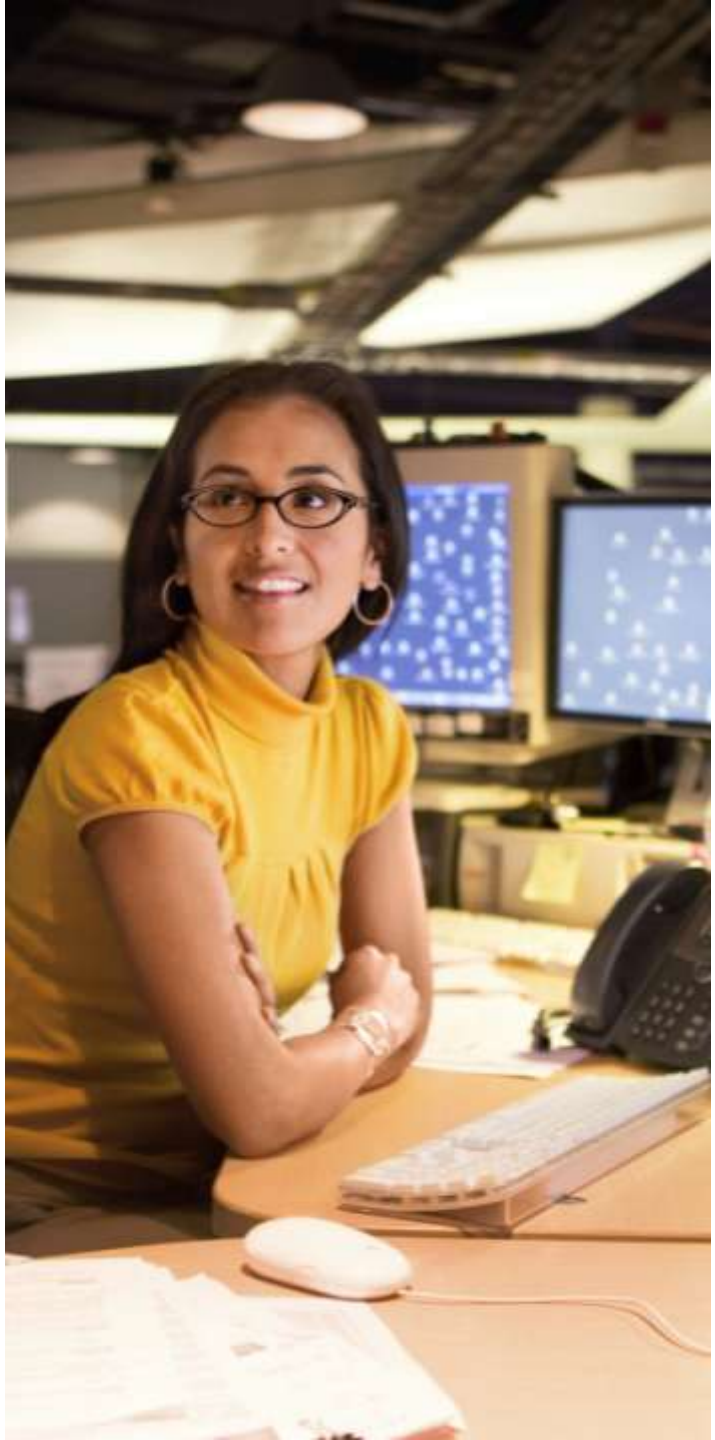

ASR9000 nV Edge Overview

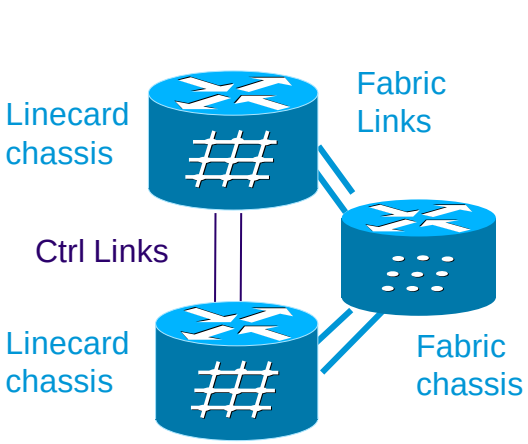

CRS Multi-Chassis

Leverage existing IOS-XR CRS multi-chassis SW infrastructure Simplified/Enhanced for ASR 9000 nV Edge

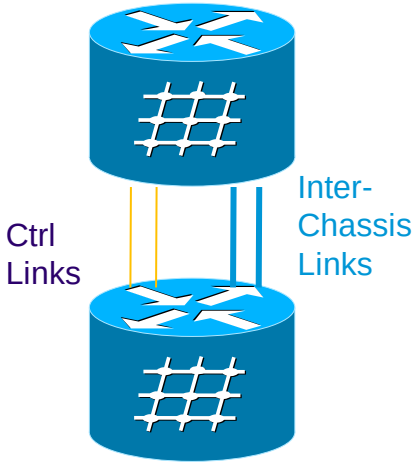

ASR 9000 nV Edge

Single control plane, single management plane, fully distributed data plane across two physical chassis → one virtual nV system

nV Edge Architecture Details

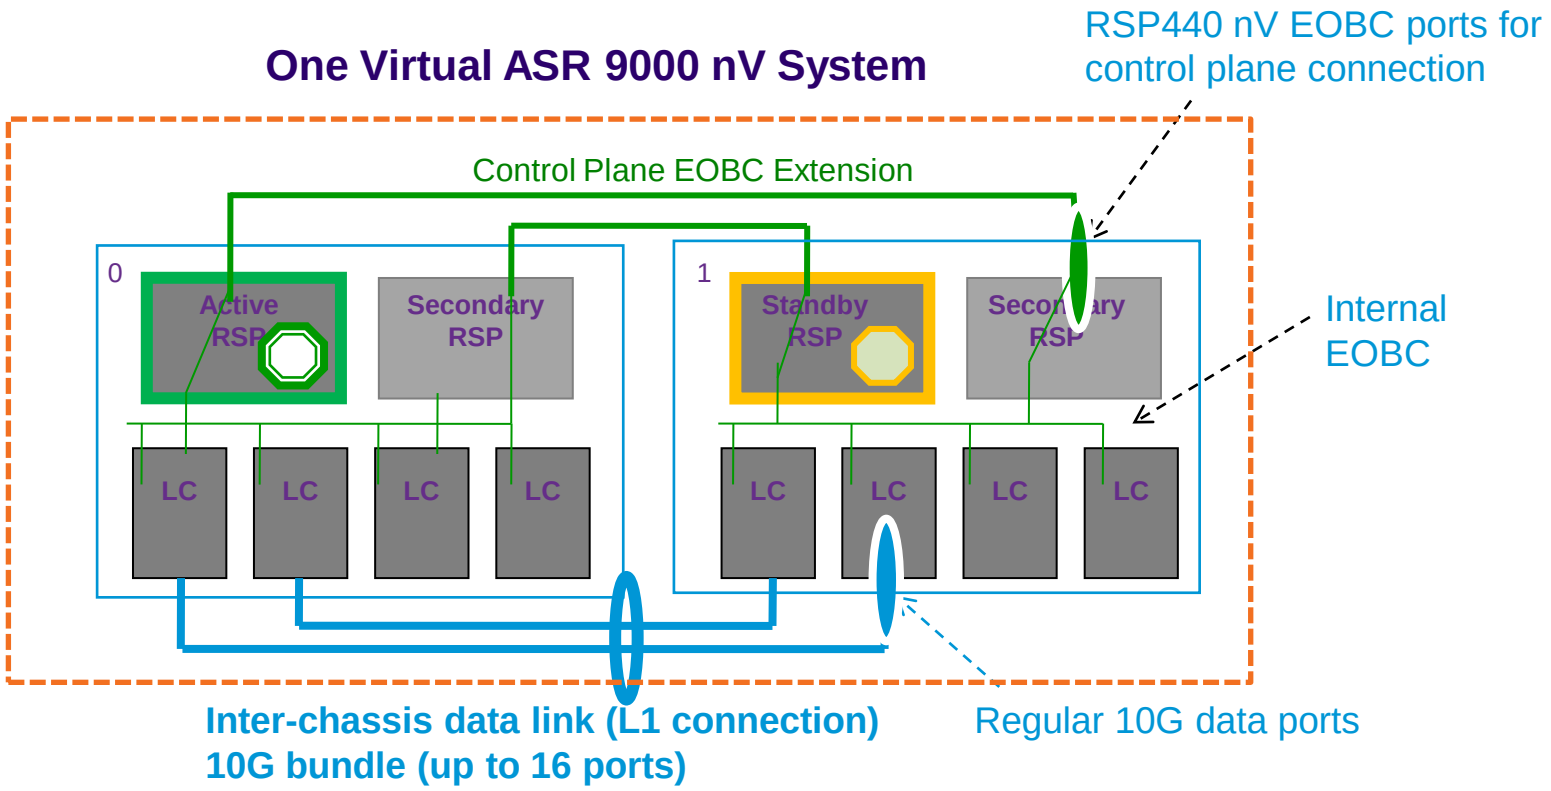

Control plane connection: Active RSP and standby RSP are on the different chassis, they communicate via external EOBC links

Data plane connection: bundle regular data links into special “nV fabric link” to simulate switch fabric function between two physical chassis for data packet

Flexible co-located or different location deployment (upto 10msec latency)

nV Edge Configuration

Configure nV Edge globally

```
nv
 edge-system
  serial FOX1437GC1R rack 1
  serial FOX1439G63M rack 0
```

← static mapping of chassis serial# and rack#

Configure the inter-chassis fabric(data plane) links

```
interface TenGigE1/2/0/0
   nv edge interface
interface TenGigE0/2/0/0
   nv edge interface
```

NO need to configure the inter-chassis control plane EOBC ports. It's plug-and-play ☺

After this configuration, rack 1 will reload and then join cluster after it boot up
Now you successfully convert two standalone ASR 9000 into one ASR 9000 nV Edge
As simple as this !!!

nV Edge Interface Numbering

Interfaces on 1st Chassis (Rack 0)

```
GigabitEthernet0/1/1/0          unassigned      Up                    Up
GigabitEthernet0/1/1/1.1        unassigned      Shutdown              Down
...
```

Interface on 2nd Chassis (Rack 1)

```
GigabitEthernet1/1/1/0          unassigned      Up                    Up
GigabitEthernet1/1/1/1.22       unassigned      Shutdown              Down
...
```

Interfaces on a Satellite connected to the nV Edge Virtual System

```
GigabitEthernet100/1/1/0        unassigned      Up                    Up
GigabitEthernet100/1/1/1.123    unassigned      Up                    Up
...
```

nV Edge System Monitoring

```
RP/0/RSP0/CPU0:ASR4-Rack0(admin)#show dsc
Thu Apr 12 03:01:12.225 UTC
-------------------------------------------------------------
          Node  (    Seq#)     Role      Serial# State
-------------------------------------------------------------
    0/RSP0/CPU0  (       0)   ACTIVE  FOX1545GRM1 PRIMARY-DSC
    0/RSP1/CPU0  (   31785)  STANDBY  FOX1545GRM1 NON-DSC
    1/RSP0/CPU0  (   31763)  STANDBY  FOX1325G77H NON-DSC
    1/RSP1/CPU0  (   32001)   ACTIVE  FOX1325G77H BACKUP-DSC

RP/0/RSP0/CPU0:ASR4-Rack0#show platform
Thu Apr 12 03:00:32.799 UTC
Node            Type                        State            Config State
-------------------------------------------------------------------------------
0/RSP0/CPU0     A9K-RSP440-SE(Active)       IOS XR RUN       PWR,NSHUT,NMON
0/RSP1/CPU0     A9K-RSP440-SE(Standby)      IOS XR RUN       PWR,NSHUT,NMON
0/0/CPU0        A9K-2x100GE-TR              IOS XR RUN       PWR,NSHUT,MON
0/1/CPU0        A9K-MOD160-TR               IOS XR RUN       PWR,NSHUT,NMON
0/1/0           A9K-MPA-2X40GE              DISABLED         PWR,SHUT,MON
0/1/1           A9K-MPA-20X1GE              OK               PWR,NSHUT,MON
0/3/CPU0        A9K-SIP-700                 IOS XR RUN       PWR,NSHUT,MON
0/3/0           SPA-8XOC12-POS              OK               PWR,NSHUT,MON
0/3/1           SPA-2XCHOC12/DS0            OK               PWR,NSHUT,MON
0/3/2           SPA-2XOC48POS/RPR           OK               PWR,NSHUT,MON
1/RSP0/CPU0     A9K-RSP440-SE(Standby)      IOS XR RUN       PWR,NSHUT,MON
1/RSP1/CPU0     A9K-RSP440-SE(Active)       IOS XR RUN       PWR,NSHUT,MON
1/3/CPU0        A9K-24x10GE-TR              IOS XR RUN       PWR,NSHUT,MON
1/4/CPU0        A9K-24x10GE-SE              IOS XR RUN       PWR,NSHUT,MON
```

nV Topologies

Single-homed, static pinning

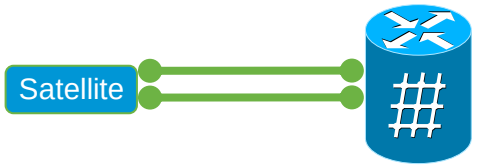

Dual-homed to nV Edge, static pinning

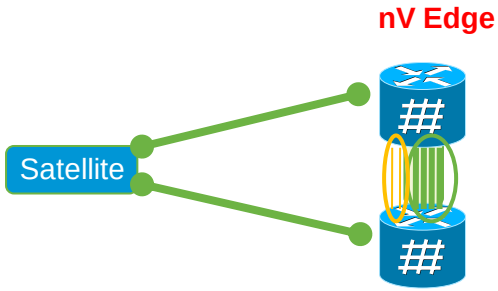

Single-homed, fabric bundle

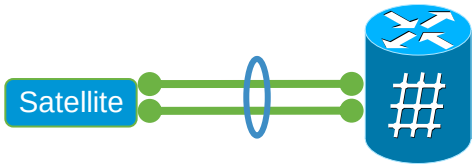

Dual-homed to nV Edge, fabric bundle

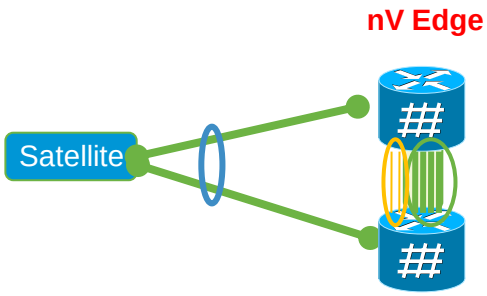

Cisco ASR9000 – Next-Gen Edge Routing Platform

Key Design Goals & System Benefits

- Architectural Design for Longevity
- Product Portfolio with significant HW and SW commonality
- Highly integrated Network Processors for High Speed Scale and Feature Flexibility
- Cisco IOS XR based
  - Truly modular, full distributed OS
  - Enhanced for the Edge (L2 and L3)
- nV (Network Virtualization) for Operational Simplicity

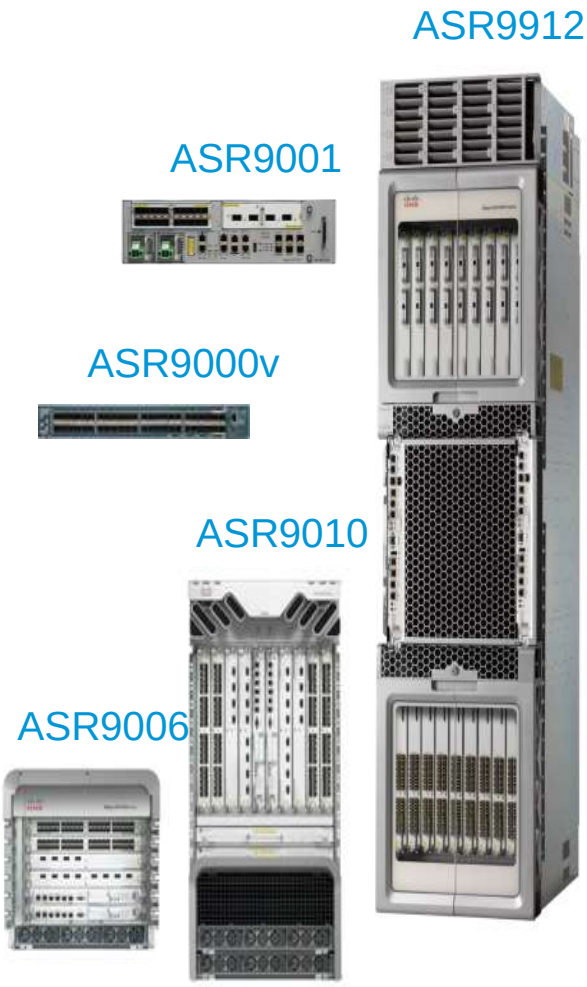

Thank you.