

What You Make Possible











Cisco ASR9000 Architecture and Network Virtualisation BRKARC-2003











Session Goals

• Understand the Cisco ASR 9000 system architecture

Key components RSP, switch fabric, Line card

Packet flows, control plane, data plane

nV (network virtualisation)

IOS XR software high level overview

Not a dedicated IOS-XR troubleshooting session

Please refer to BRKSPG-3612 (Troubleshooting Cisco IOS XR) for IOS-XR related technical information

Not a dedicated platform configuration / troubleshooting session





Hardware Overview

Carrier Class, Scalable System Architecture

- Fabric architecture
- Line card architecture
- Packet flow, control plane and data plane
- Multicast packet flow and replication
- Internal system Qos
- nV Technology Overview



ASR 9000 System Introduction and Chassis Overview









Cisco ASR 9000 – Next-Gen Edge Routing Platform High level system design goals

- System longevity : 10-15 years in the network
- Full HW portfolio / form factors
- Software modularity, portability, flexibility
- High power control plane / data plane scale
- 10/100/??? Gigabit ethernet aggregation
- Complex feature paths : Video, Mobile, BNG
- Data centre internet and cloud gateway







Cisco ASR 9000 – Next-Gen Edge Routing Platform Achieving design goals

- Significant hardware and software commonality
- Built on a repeatable/extensible "slice" arch.
- Edge-enhanced IOS XR software
- Shift to x86 CPU for improved power and virtualisation capabilites
- Highly integrated network processors for high data plane scale and feature flexibility
- Complete set of L2/L3 protocols and forwarding microcode





System Capacity : key components

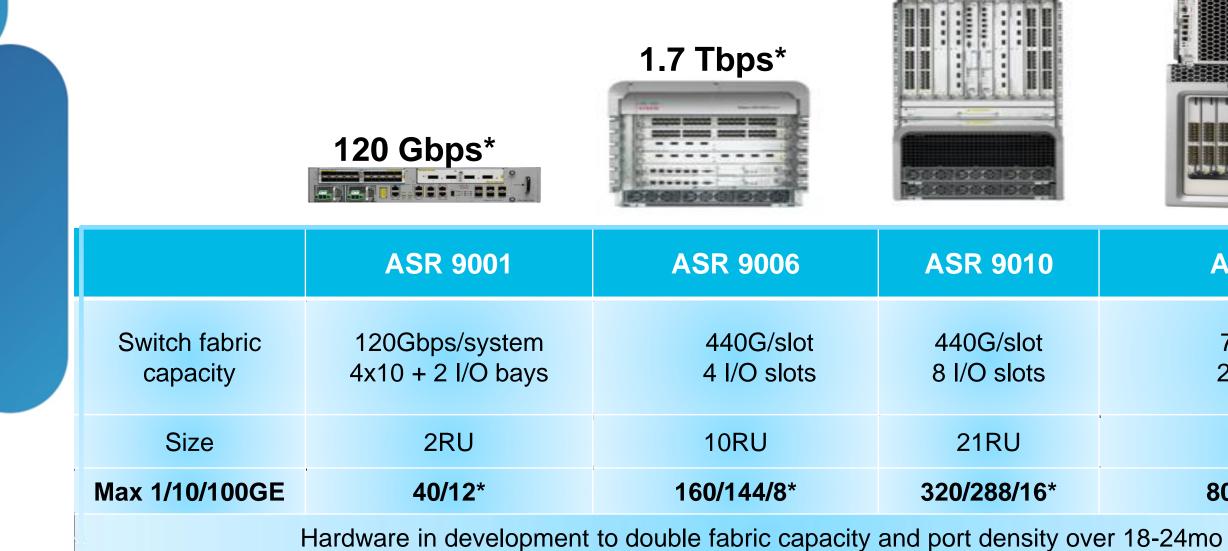
- Modular systems have multiple components with independent capabilities and interactions
- Important to understand how these fit together and relative costs
- Backplane
- System interconnect (generally a switch fabric)
- Linecards (or choose whatever name you like)
- ASR 9000 series has 5 current form factors, each with multiple current and future interconnect/fabric options, and multiple linecard options





ASR 9000 Chassis Overview

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



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

BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved.

3.5 Tbps*



ASR 9922

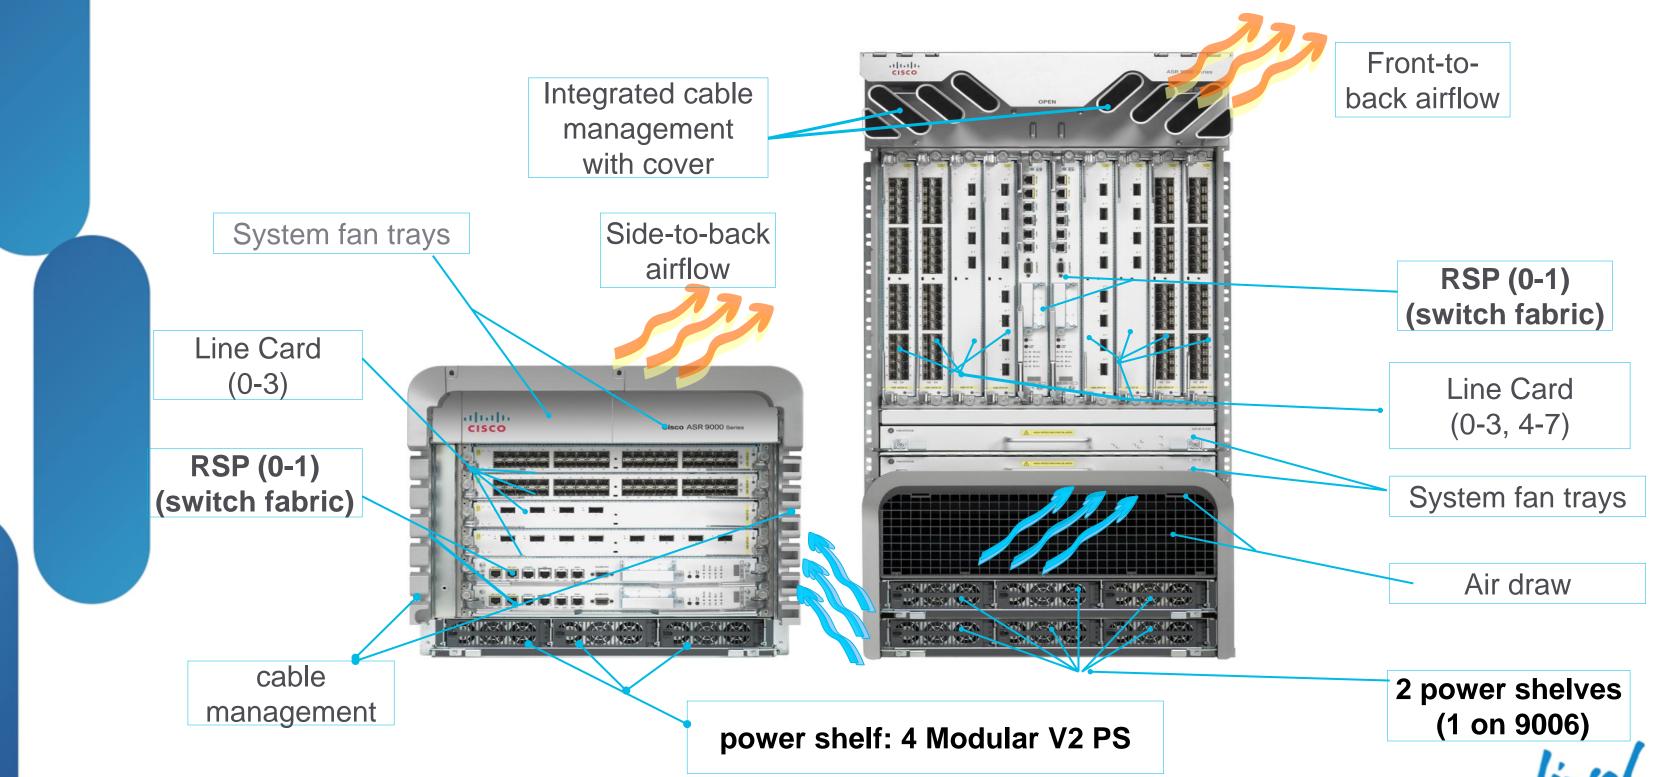
770G/slot 20 I/O slot

44RU

800/720/40*



ASR 9010 and ASR 9006 Chassis

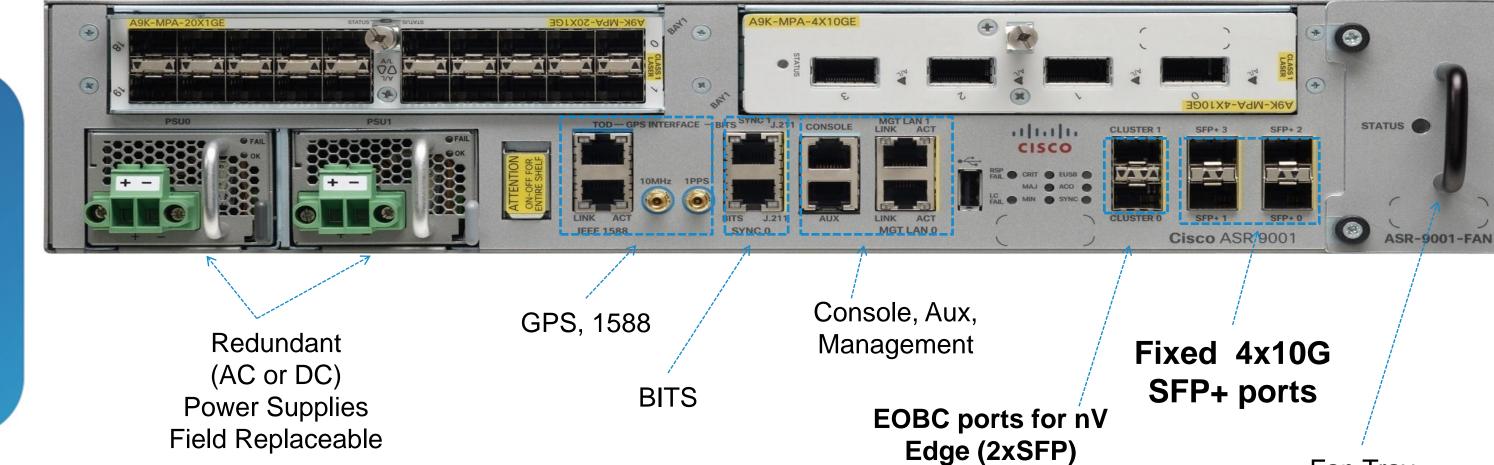


Cisco

ASR 9001 "Iron Man" Chassis

Sub-slot 0 with 20xGEMPA

Sub-slot 1 with 4x10GE MPA



Supported MPA (Media Port Adapter): 20xGE, 2/4x10GE, 1x40GE

BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved.

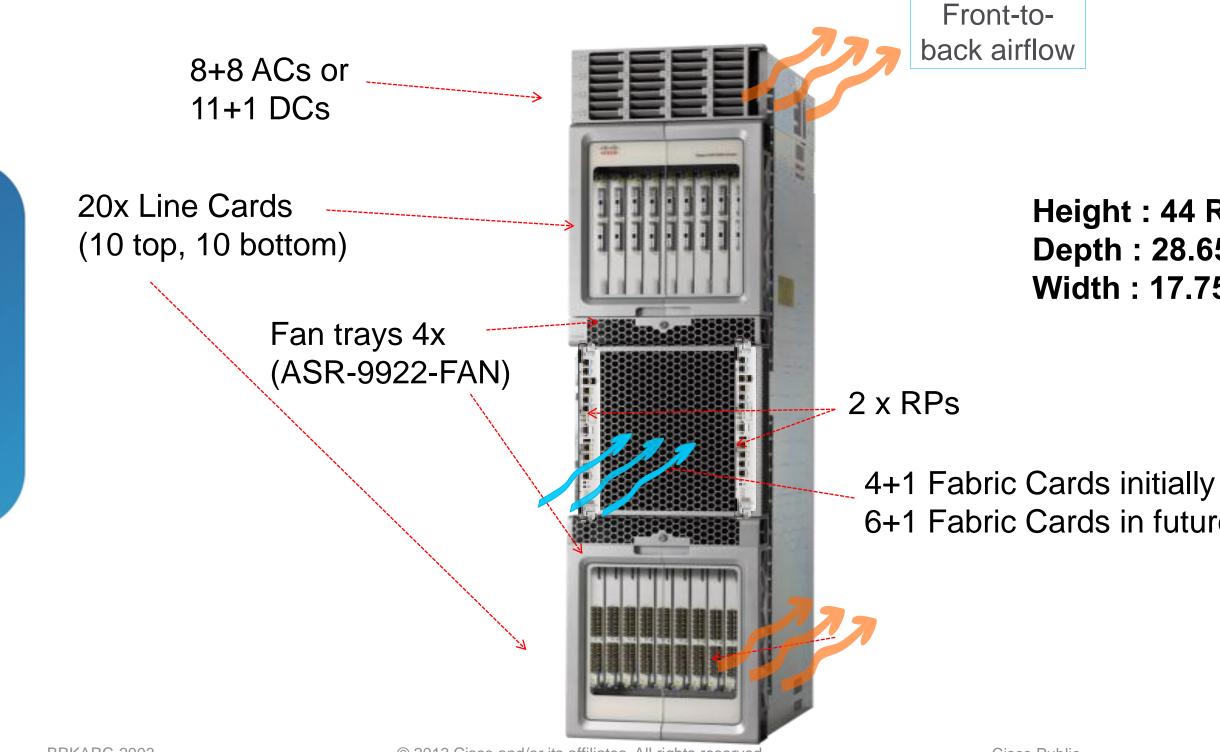




Fan Tray Field Replaceable

Cisco

ASR 9922 "Megatron" Chassis bandwidth system



BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved.



Height : 44 RU (AC & DC) Depth : 28.65" (800mm) Width : 17.75" (fits 19" rack)

6+1 Fabric Cards in future



Cisco Public

ASR 9000 System Components









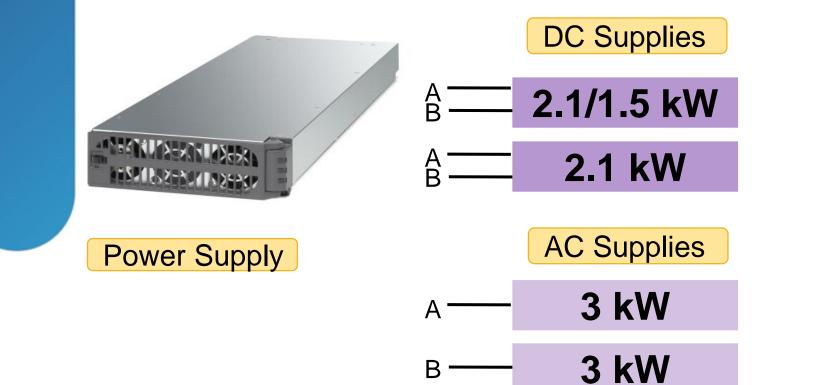
Power and Cooling



ASR-9010-FAN



- Fans unique to chassis
- Variable speed for
- Redundant fan-tray



Single power zone Power draw shared evenly for Easy CO Install

ambient temperature variation Low noise, NEBS and OSHA compliant

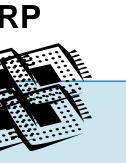
All power supplies run in active mode 50 Amp DC Input or 16 Amp AC



ASR 9006/9010/9922 RSP (Route/Switch Processors)

	RSP	RSP440	9922-R
Processors	PPC/Freescale	Intel x86	Intel x8
	2 Core 1.5GHz	4 Core 2.27 GHz	4 Core
RAM	RSP-4G: 4GB RSP-8G: 8GB	RSP440-TR: 6GB RSP440-SE: 12GB	-TR: 60 -SE: 12
nV EOBC ports	No	Yes, 2 x 1G/10G SFP+	Yes, 2 SFP+
Switch fabric bandwidth	92G + 92G (with dual RSP)	220+220G (with dual RSP)	660+11 (7-fabri

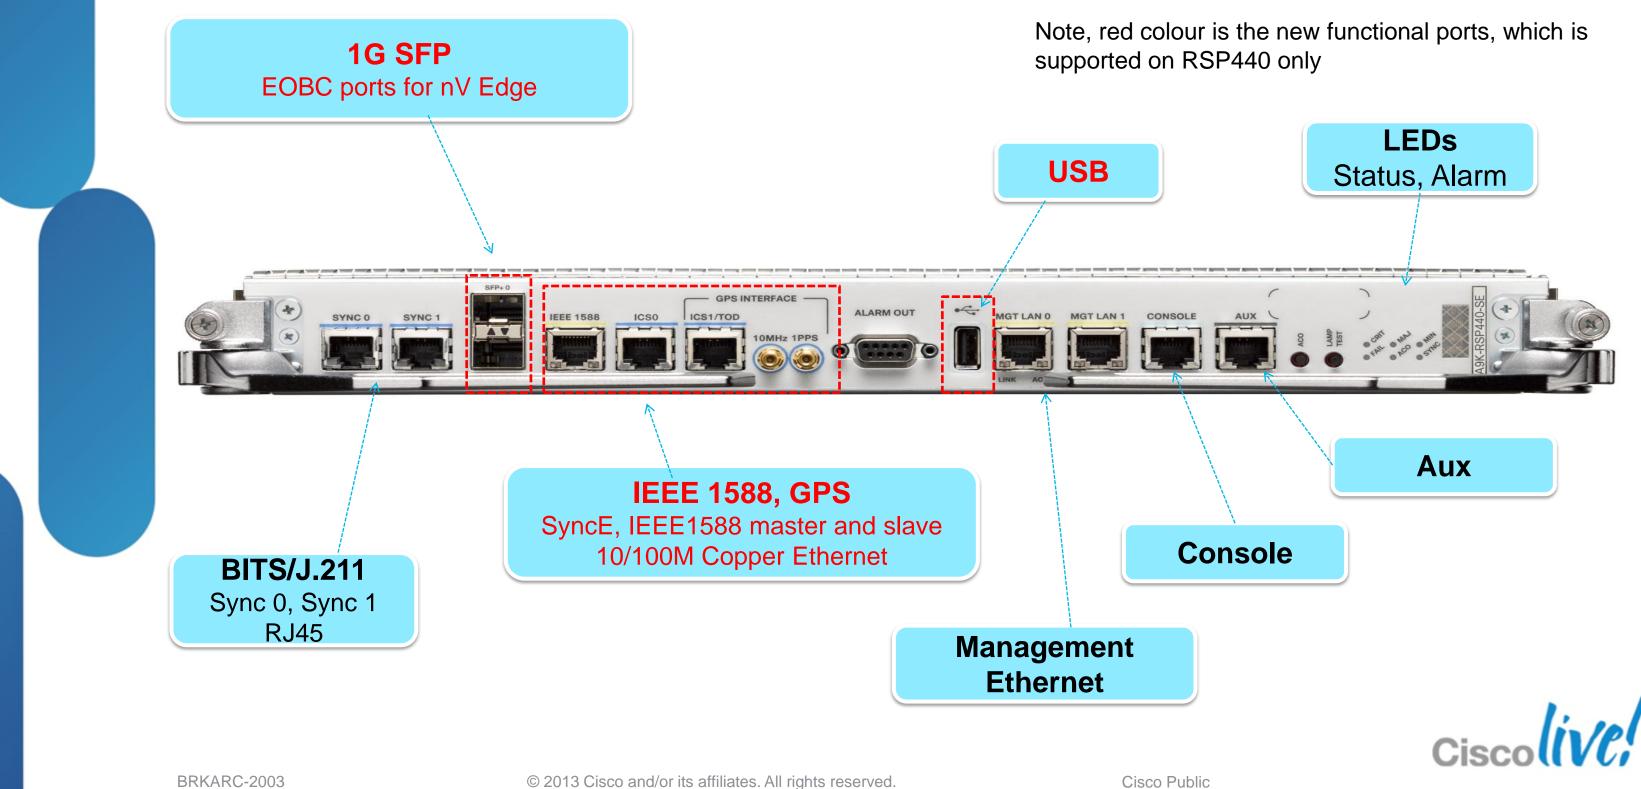




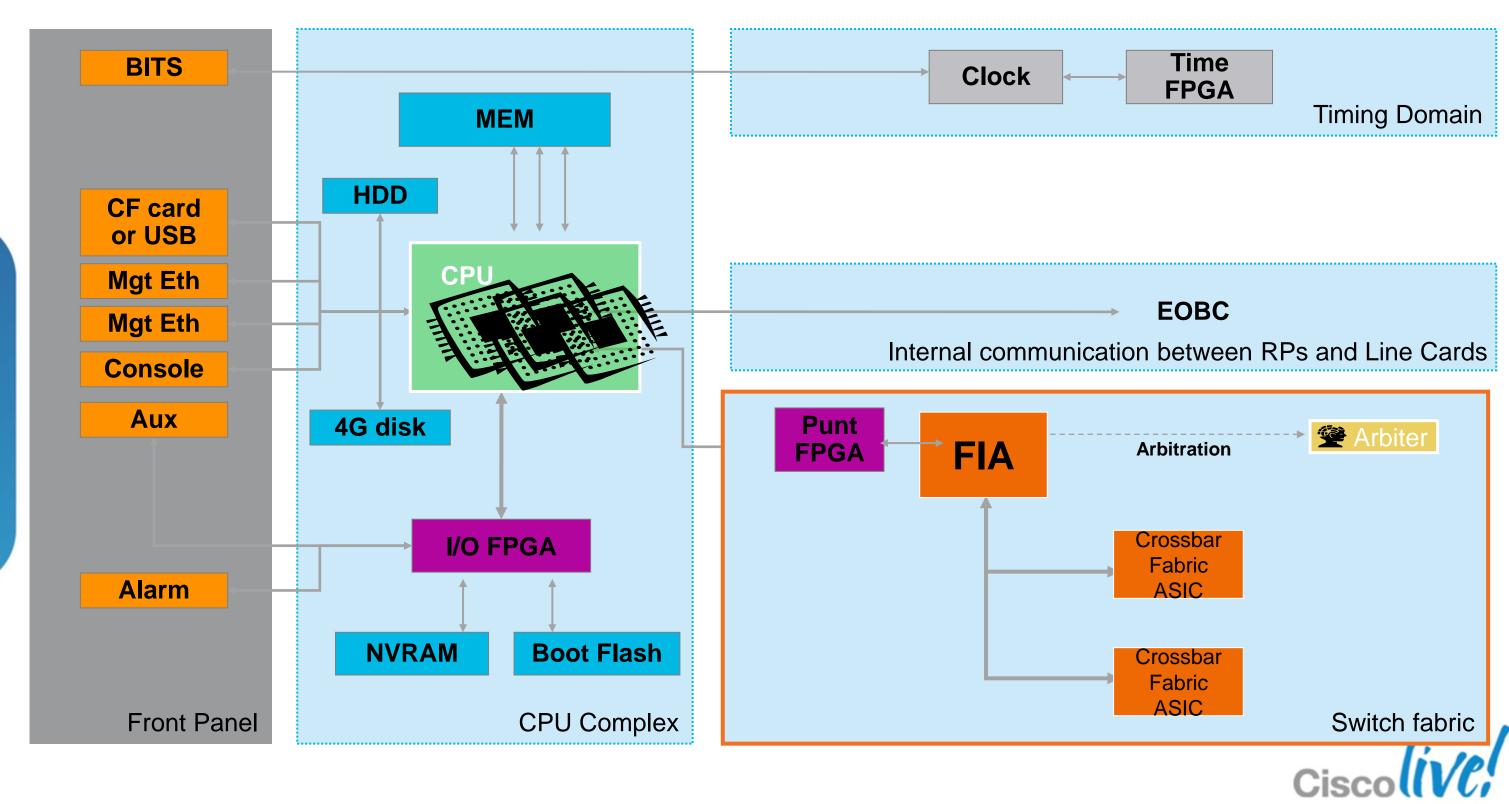
- 86
- e 2.27 GHz
- GB
- 2GB
- x 1G/10G
- 10 ric model)



RSP440 – Front Ports

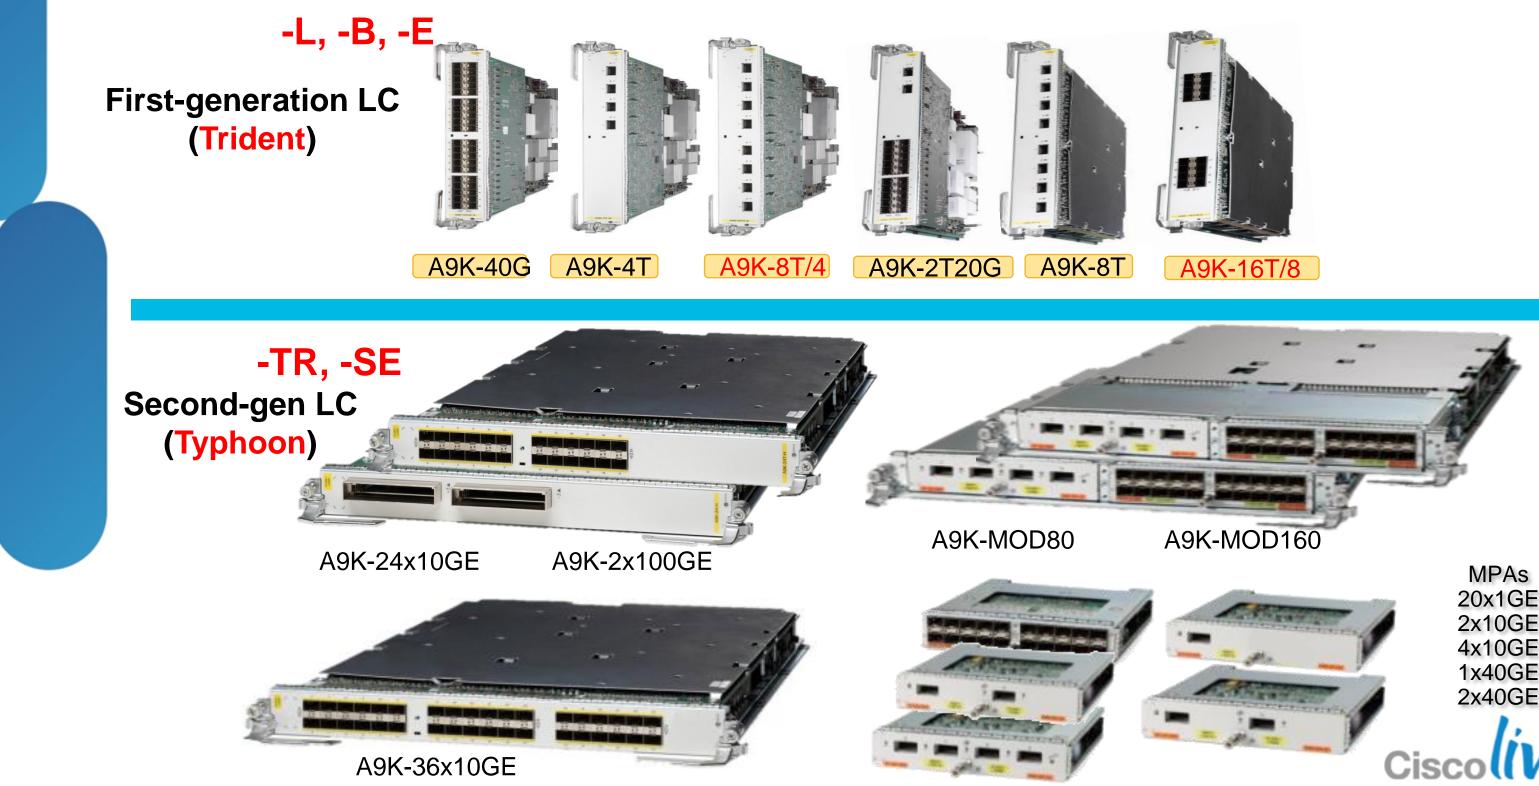


RSP Engine Architecture



BRKARC-2003

ASR 9000 Ethernet Line Card Overview



BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved.

Cisco Public

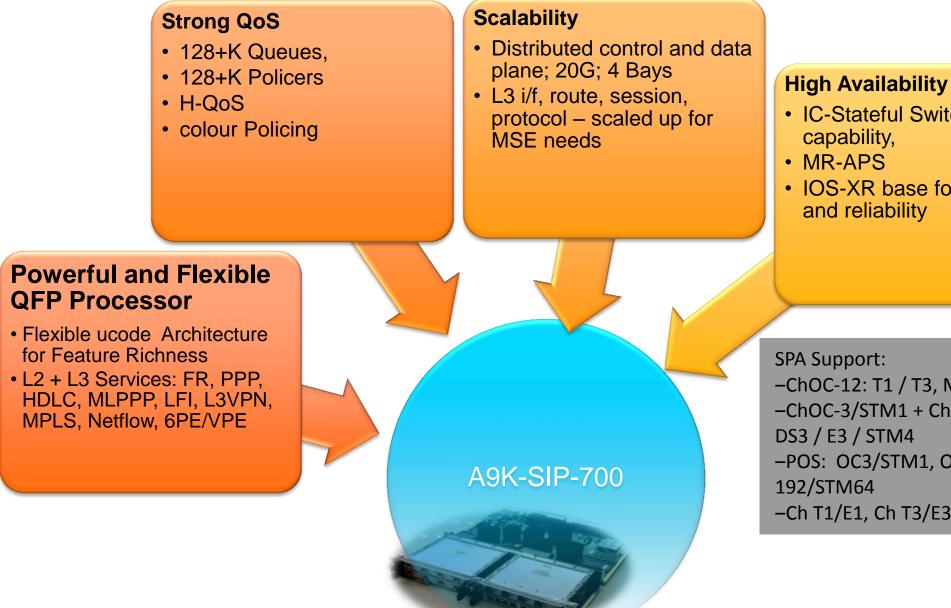


20x1GE 2x10GE 4x10GE 1x40GE 2x40GE

ASR 9000 SIP-700 and SPA **20Gbps, feature rich, high scale, low speed interfaces**

SPA support list:

http://www.cisco.com/en/US/partner/prod/collateral/routers/ps9853/data_sheet_c78-573452.html



 IC-Stateful Switch Over IOS-XR base for high scale

-ChOC-12: T1 / T3, MLPPP, LFI, IC-SSO, MR-APS -ChOC-3/STM1 + ChOC-12: DS0 / STM1 + ChOC-48:

-POS: OC3/STM1, OC-12/STM4, OC-48/STM16, OC-

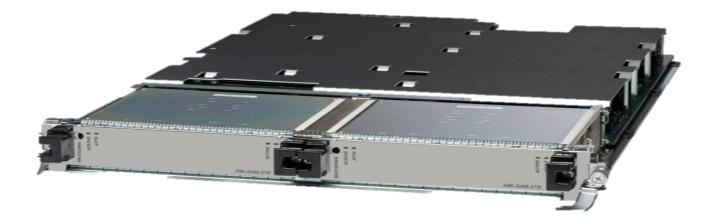
-Ch T1/E1, Ch T3/E3, CEoPs, ATM



ASR 9000 ISM (Integrated Service Module)

CDS Streaming: TV and internet streaming Error repair

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



Feature
Applications
Bandwidth
Compatibility
Concurrent Streams
Content Cache
Video Formats
Transport
Session Protocols
Environmental

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



ASR 9000 ISM Capabilities

Ultra-Dense VoD. TV. Internet Streaming, Error Repair, CGv6

30-40 Gbps streaming capacity ~3 Gbps cache fill rate

Works with all CDS appliances

Up to 8,000 SD equivalent

3.2 TBytes at FCS - Modular Design

MPEG2 & AVC/H.264

MPEG over UDP / RTP

RTSP / SDP

NEBS / ETSI compliant

Optical Networking Overview









Key concepts for high end optical routers/switches

- Optics are now a very significant part of system design.
- Mechanical design issues (space, power, heat) are becoming or in some cases more significant as packet forwarding
- The market for modular optics is exceptionally competitive and wildly variable in cost, quality, specifications
- This generates complexity and pain (we're sorry!) but ...
- lower costs and improves densities (you're welcome!)





ASR 9000 Optical Interface Support

- Modular optics used everywhere
- GE: SFP interfaces (T, SX, LX, ZX)
- IOGE: SFP+ and XFP modules
 - XFP (for line card with less than 8x10GE) ports) : SR, LR, ZR, DWDM, CWDM, Tunable
 - SFP+ (for line card with more than 8X10GE) ports): SR, LR, ZR

SFP, SFP+







ASR 9000 40/100GE Optics

- ASR9k 40GE: QSFP form factor
 - SR4
 - LR4
- ASR9k 100GE: CFP form factor
 - LR4
 - SR10
- **DWDM** support:
 - 40GE has g.709/FEC built-in at FCS needs Q-SFP DWDM optics (future)
 - 100GE is "gray" only and no FEC/g.709 use external transponder

QSFP (40GE)



CFP (100GE)



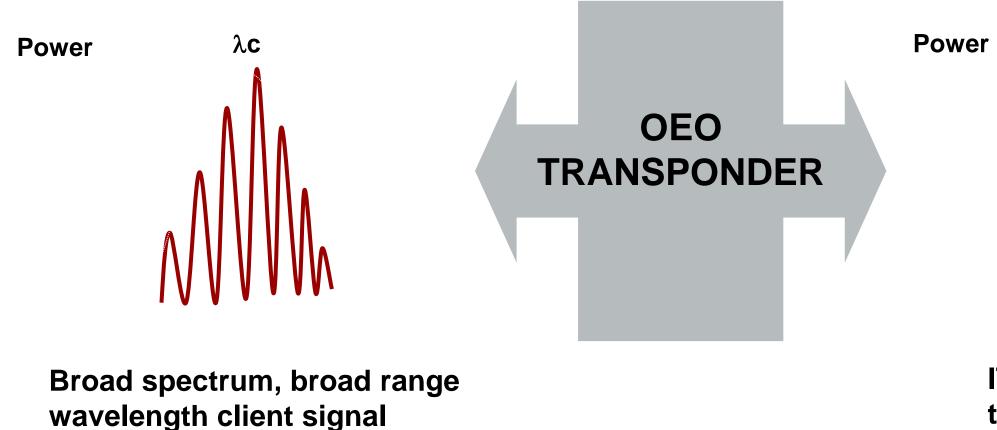
IPoDWDM overview

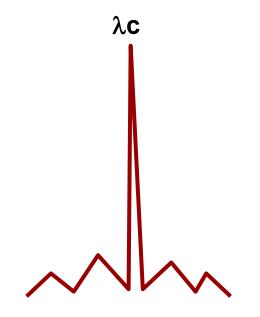
- "IP over DWDM" is slightly mis-named....
- implies two things which are really orthogonal
 - Long-reach, high-precision, coloured optics
 - Enhanced error-correction algorithms to extend signal integrity ("transponder")
- (normally) requires some type of DWDM mux device (almost always dedicated HW)



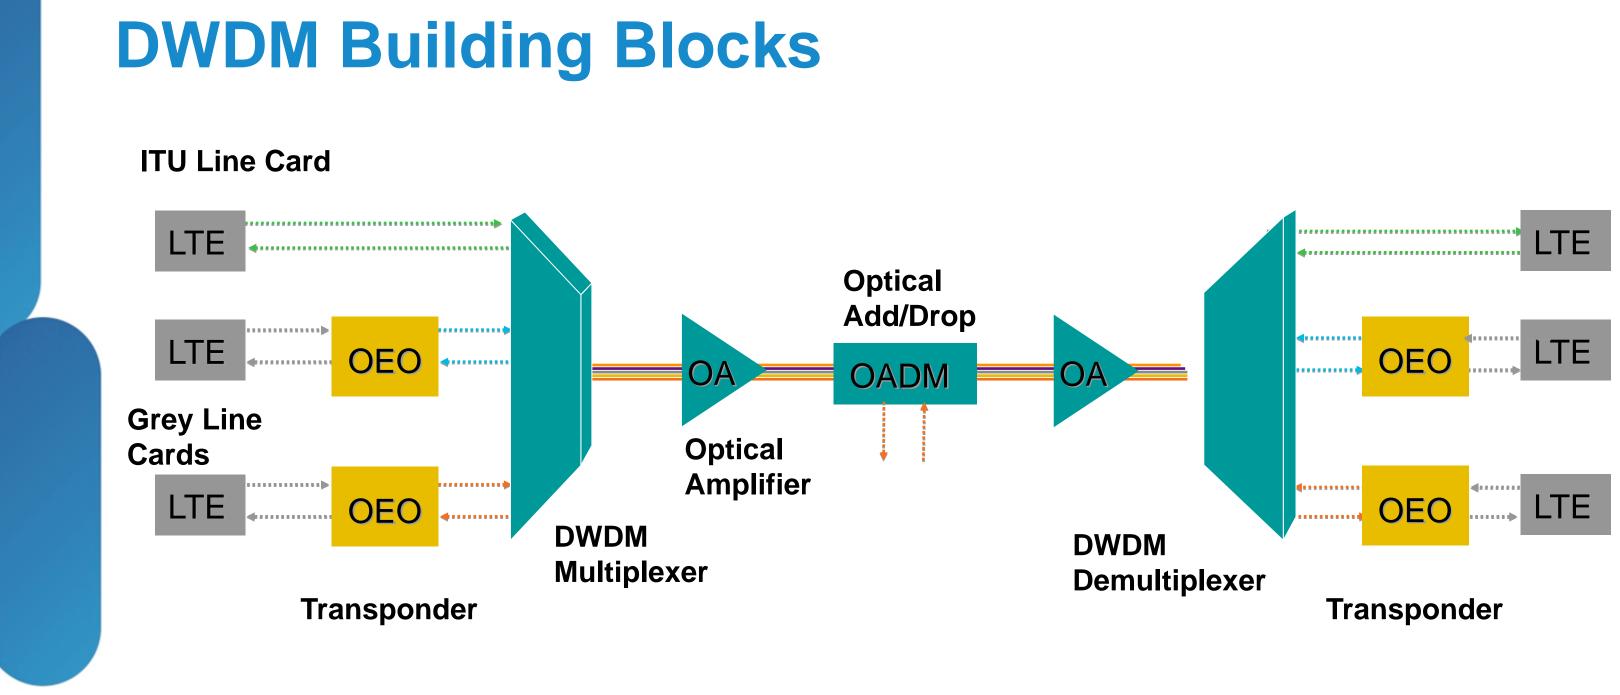
Transponder OEO Converter

- Converts the wide spectrum client laser input into a DWDM ITU compliant channel
- Converts the DWDM channel into a client compliant output signal
- Down side is that they add significant cost





ITU Wavelength specific DWDM trunk channel



LTE: line terminating equipment **OEO: optical-electrical-optical converter**

BRKARC-2003



ASR 9000 System Interconnect and Switch Fabric Architecture

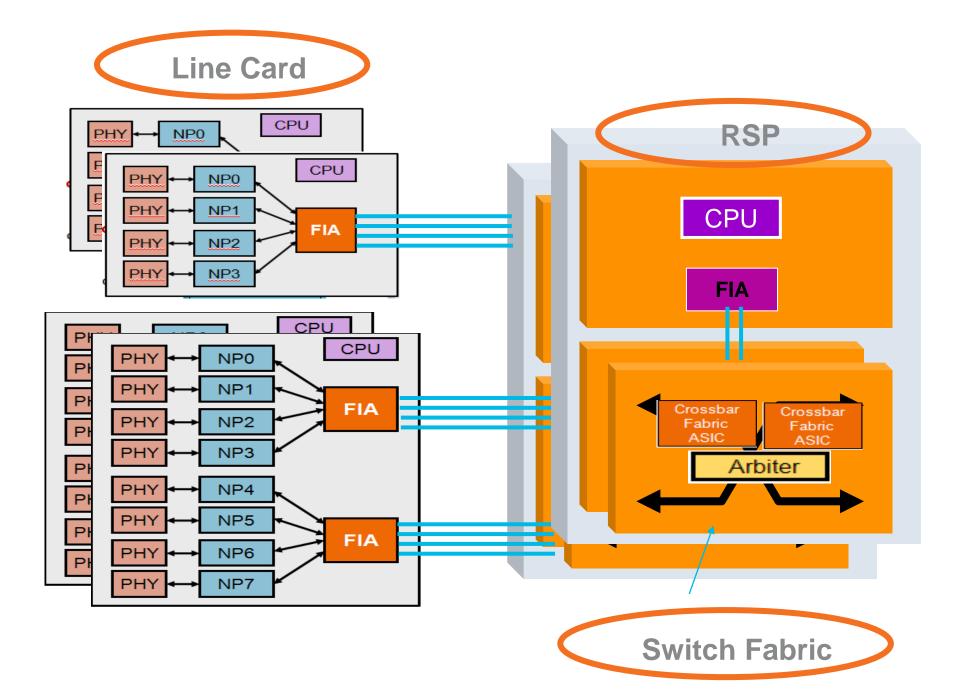




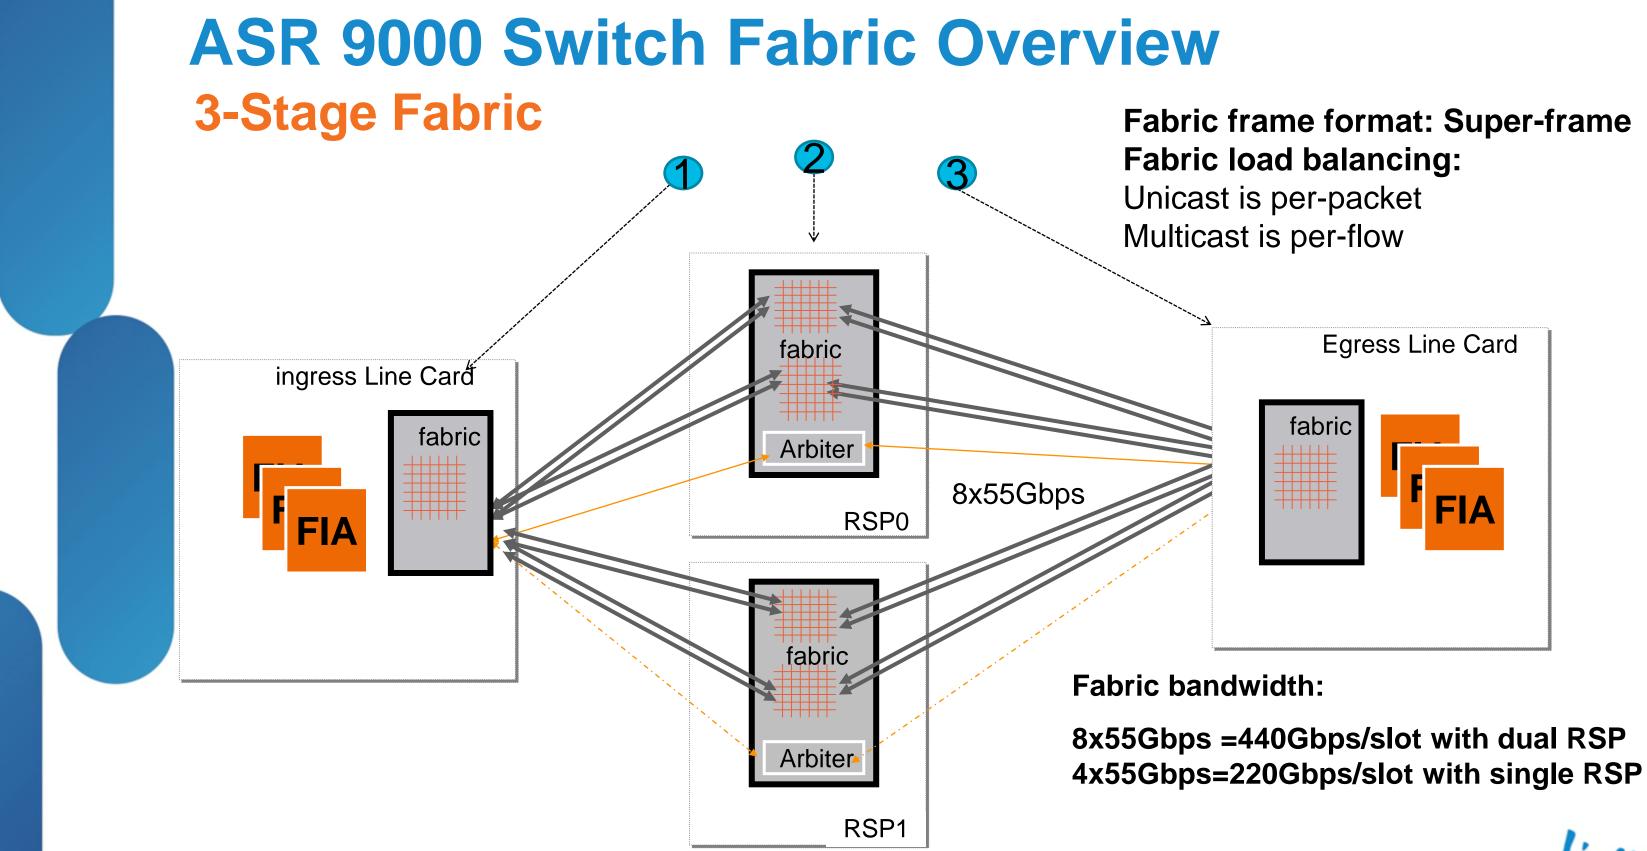




ASR 9000 System Architecture "At-a-Glance"



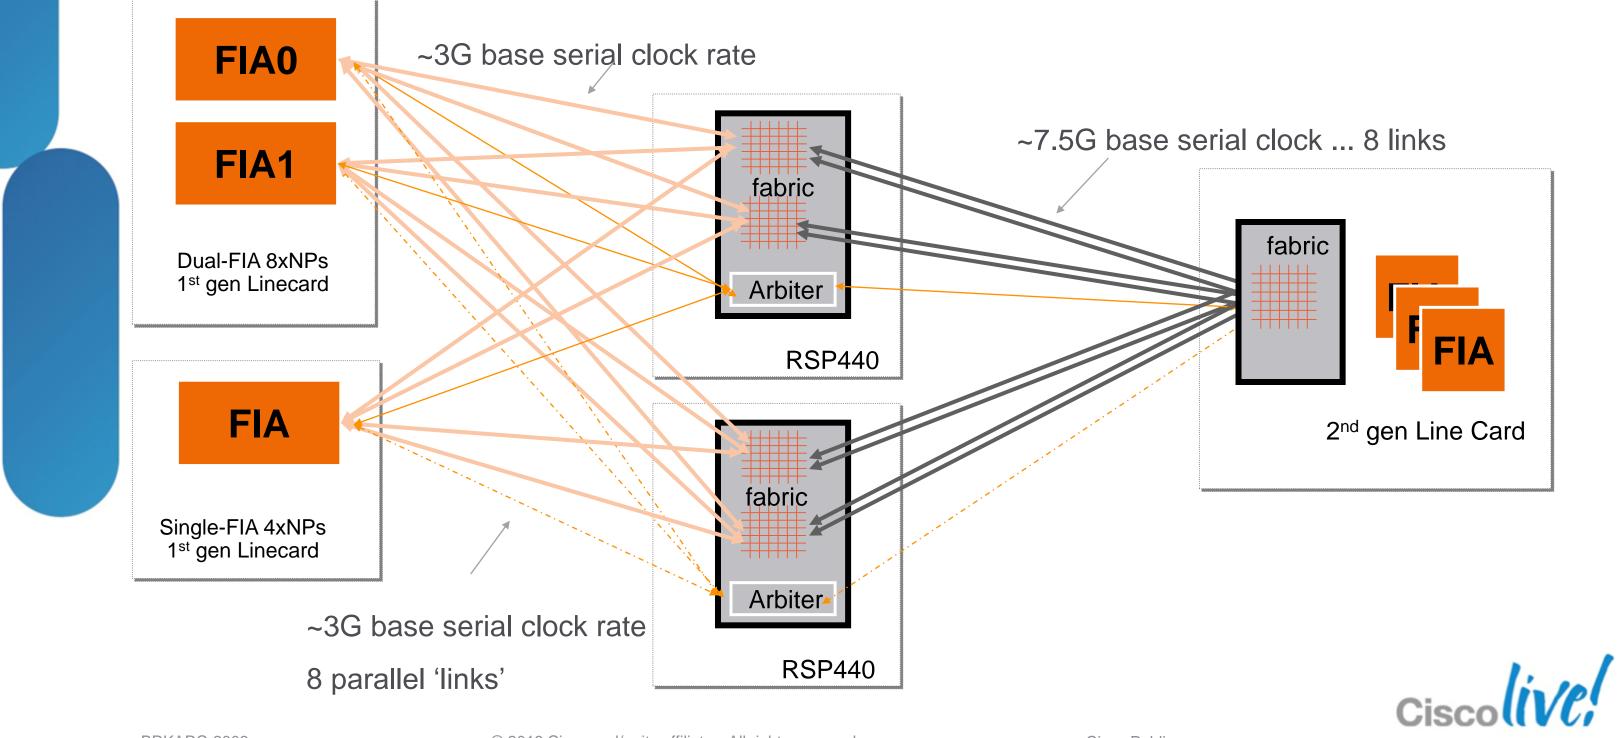




Fabric frame format: Super-frame



1st/2nd Generation Switch Fabric Compatibility





Fabric Forwards/Backwards Compatibility

- 1st generation linecards based on 3.25Ghz serial links
- 2nd gen linecards/fabric based on 7.5Ghz
- In 2nd generation, both the linecard fabric interface ASICs and the fabric crossbars can negotiate to higher or slower clock rate
- Provides forward/backwards compatibility in fabric

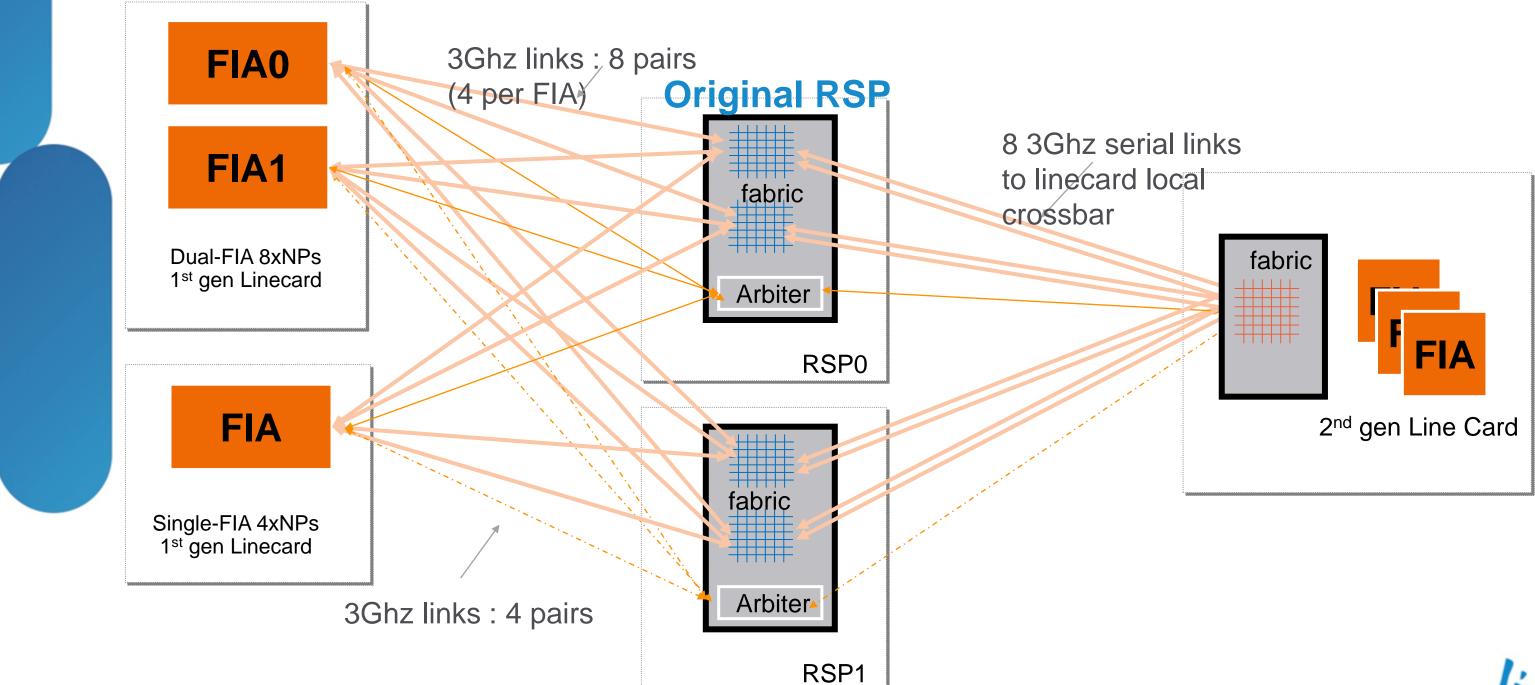


Switch Fabric Scalability in 3 Dimensions

- Base clock rate of the discrete serial components
 - 3.125Ghz, 6.25Ghz, 7.5Ghz, 15Ghz... ???
- Number of parallel serial links that make up a "fabric port" (usually either 4x or 8x on the asr9k)
- Number of parallel fabric "planes" in the system...
- Fabric planes provide redundancy (N+1) for switch fabric complex
- Example: 4 planes * 8 links * 3.125Ghz
 - Yields $\rightarrow \sim 92$ Gb (80g linecards @ 15% fab speedup)

omponents

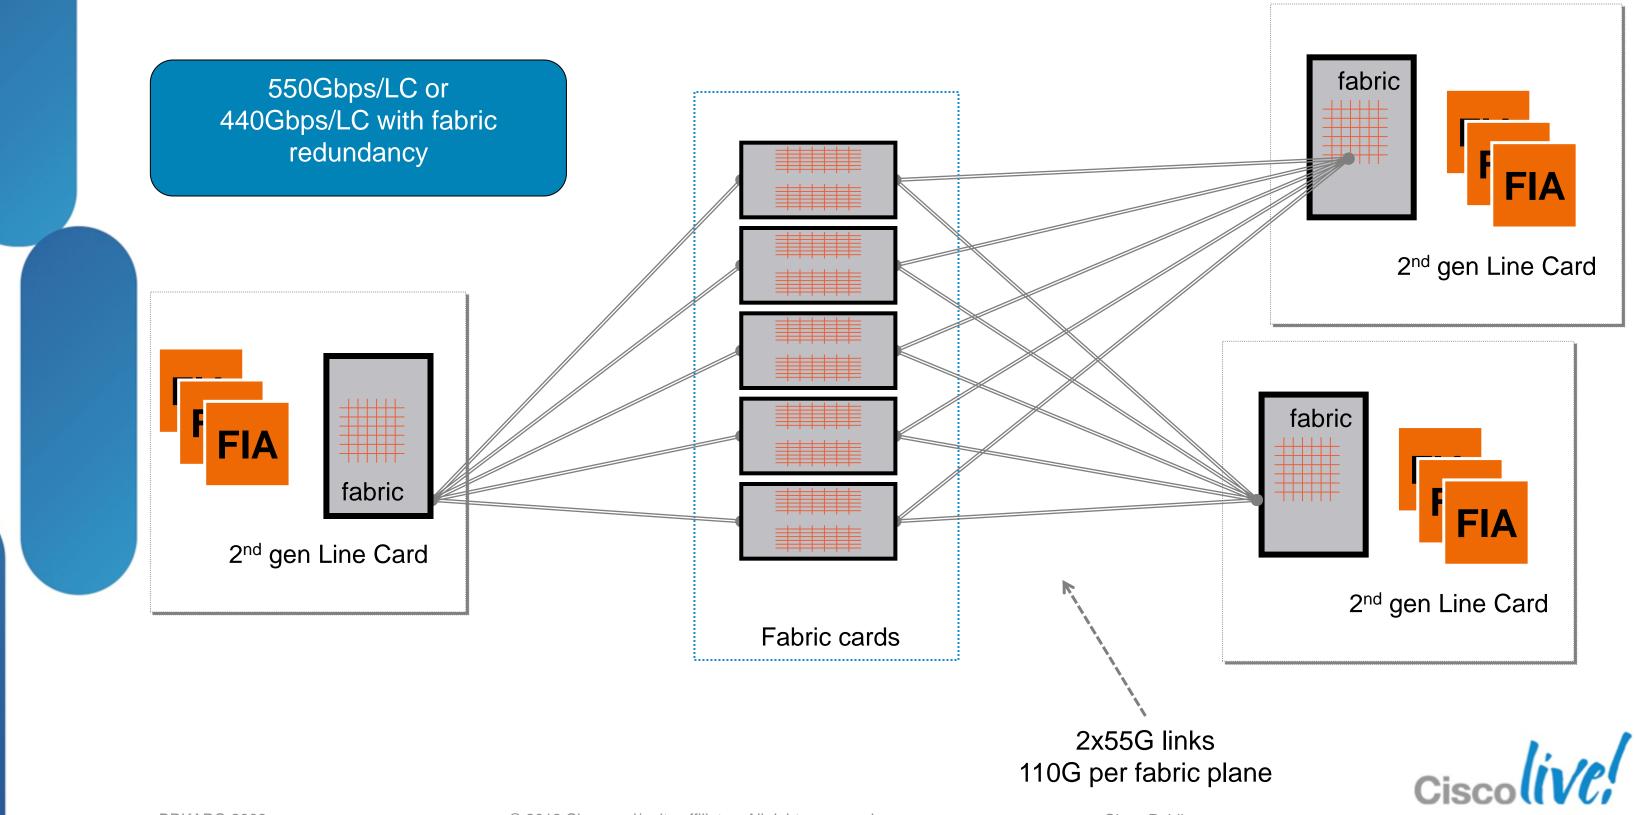
1st/2nd Generation Switch Fabric Compatibility



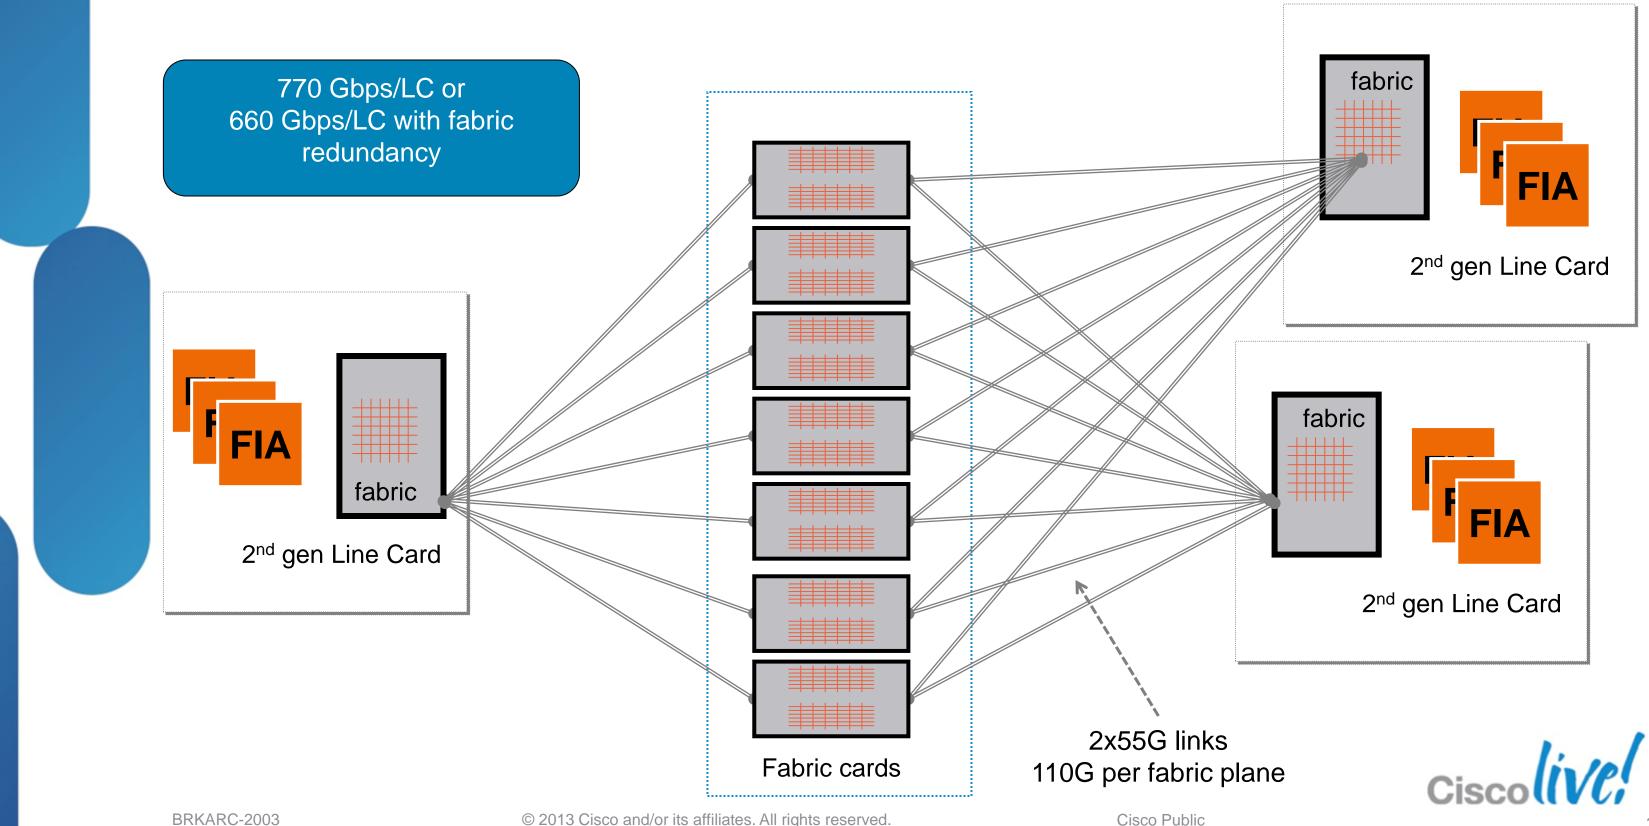




ASR 9922 Fabric Architecture : 5-plane System



ASR 9922 Fabric Architecture : 7-plane System



© 2013 Cisco and/or its affiliates. All rights reserved.

Generic LC Architecture (1) – Components

Pluggable physical interfaces

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

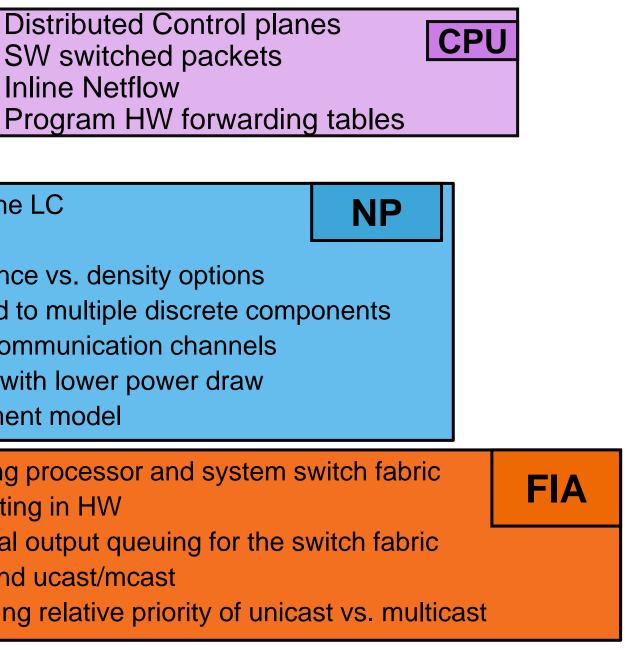
SW switched packets **Inline Netflow**

- forwarding and feature engine for the LC
- scales bandwidth via multiple NPs

PHY

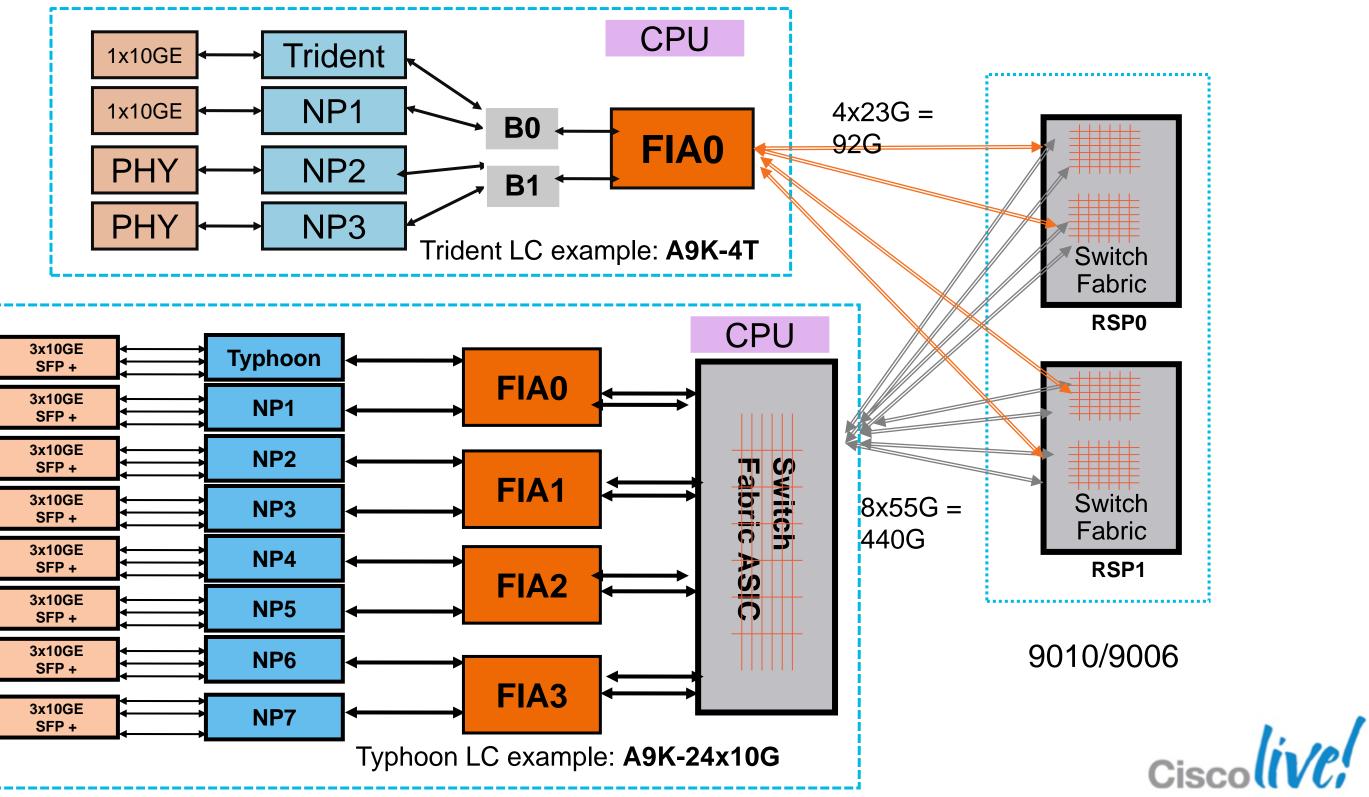
- up to 8 NPs/LC for performance vs. density options
- highly integrated silicon as opposed to multiple discrete components
 - shorter connections, faster communication channels
 - higher performance, density with lower power draw
 - simplified software development model
 - interface between forwarding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





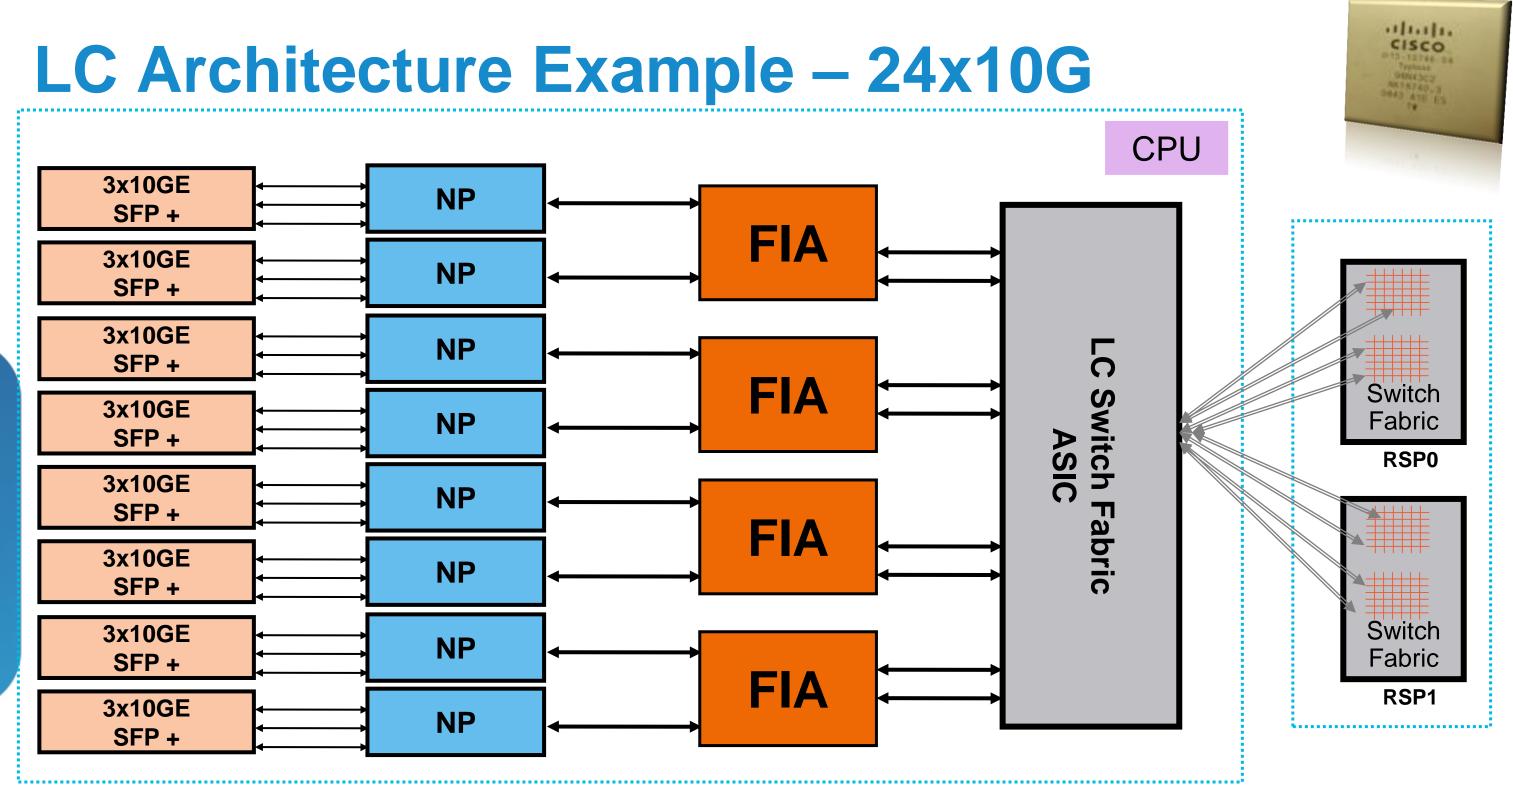
ASR 9000 Line Card Architecture Overview



BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved.





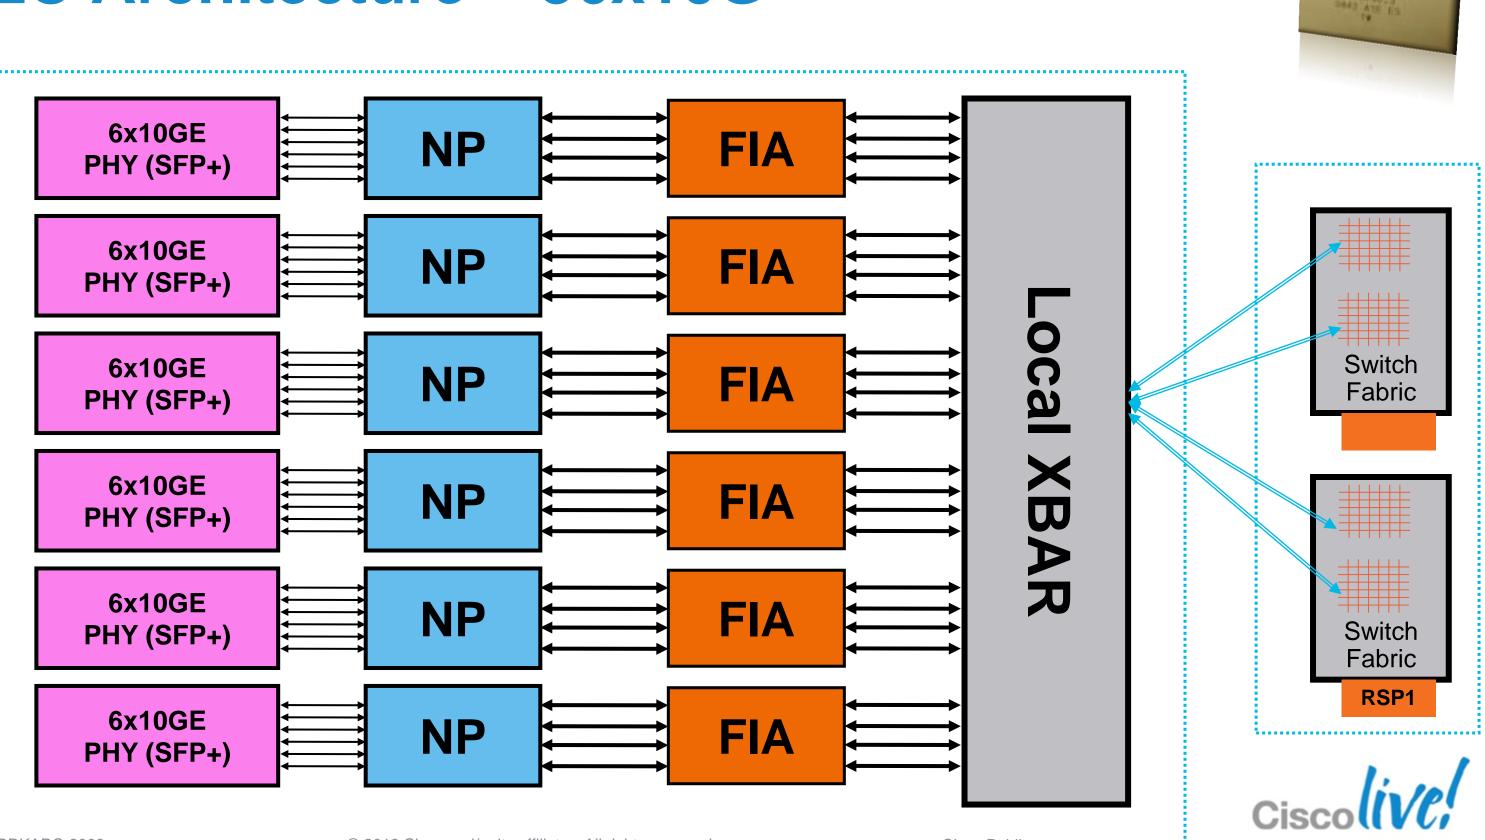
Each NP: 60Gbps bi-directional 120Gbps uni-directional 2x45Mpps

Each FIA: 60Gbps bi-directional

© 2013 Cisco and/or its affiliates. All rights reserved.



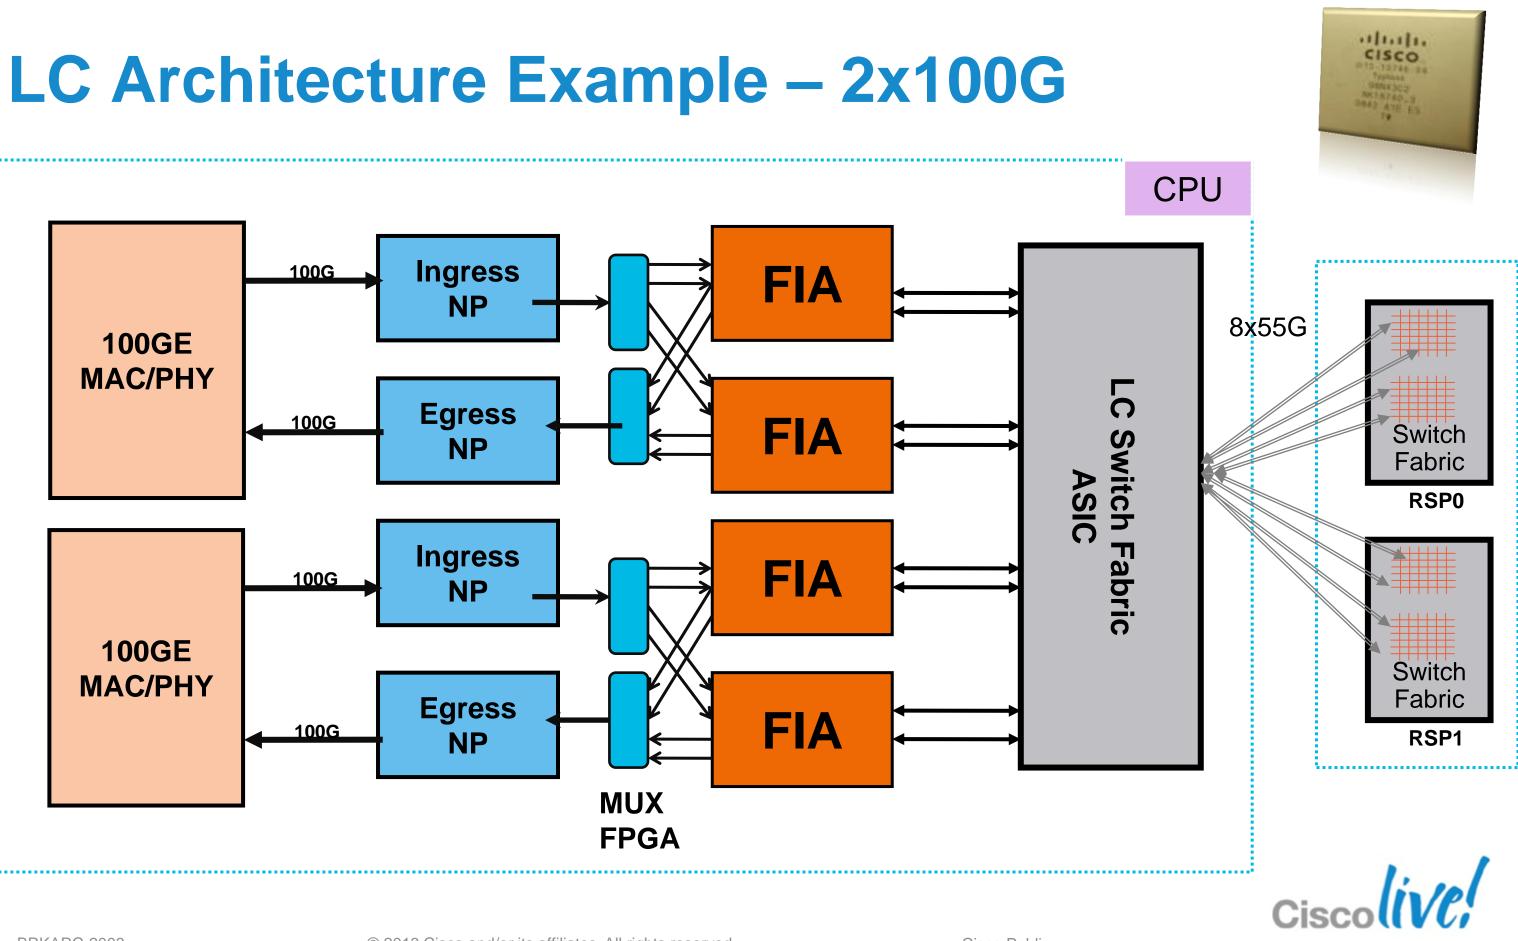
LC Architecture – 36x10G



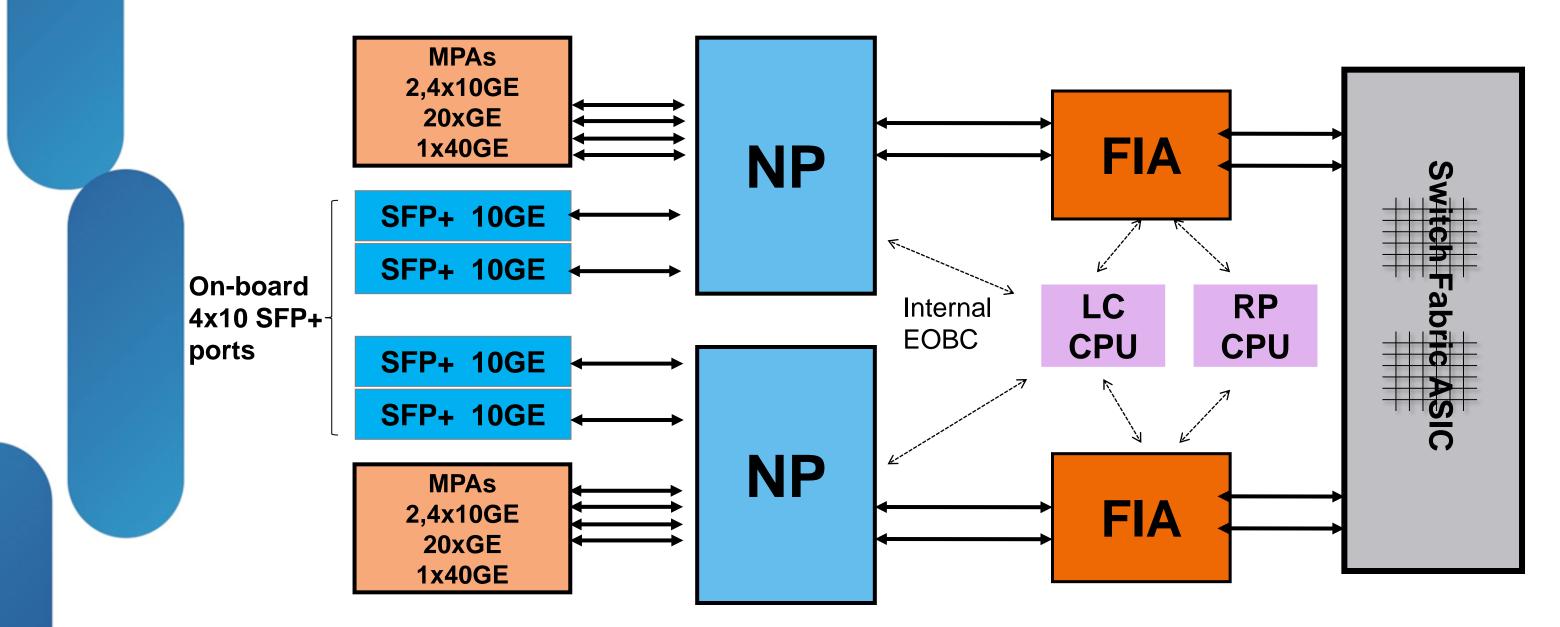
© 2013 Cisco and/or its affiliates. All rights reserved. Cisco Public

Almala. CISCO

KTETED. 7

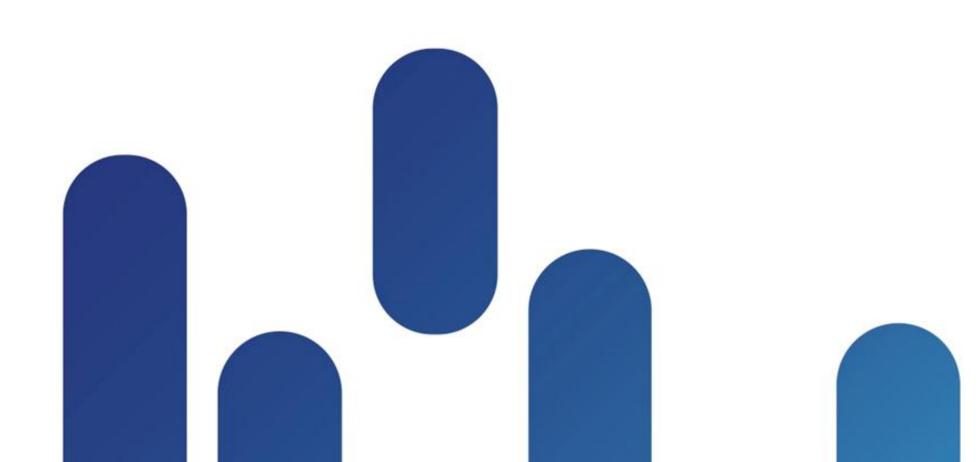


ASR 9001 Architecture Overview Identical components and architecture





ASR 9000 Packet Forwarding Infrastructure









Architectural Choices: Forwarding Paradigms

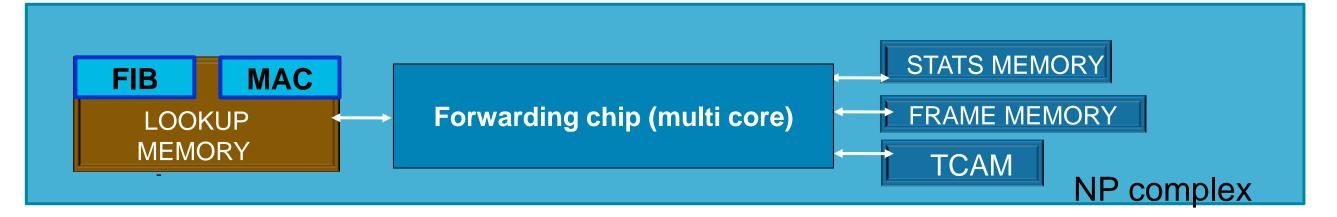
- Three high level options
- Fully integrated ASIC
 - Optimised for : highest performance, unit cost/power
 - Sacrifices: TTM, flexibility, extensibility
- **Network processor / customized microcode**
 - Optimised for : good flexibility with high performance / BW
 - Sacrifices: moderate cost, programming complexity
- General purpose CPU / stack programming model
 - Optimised: development, most flexible, TTM, complex code
 - Sacrifices: performance, bandwidth, unit cost/power

BRKARC-2003

Cisco and/or its affiliates. All rights reserved



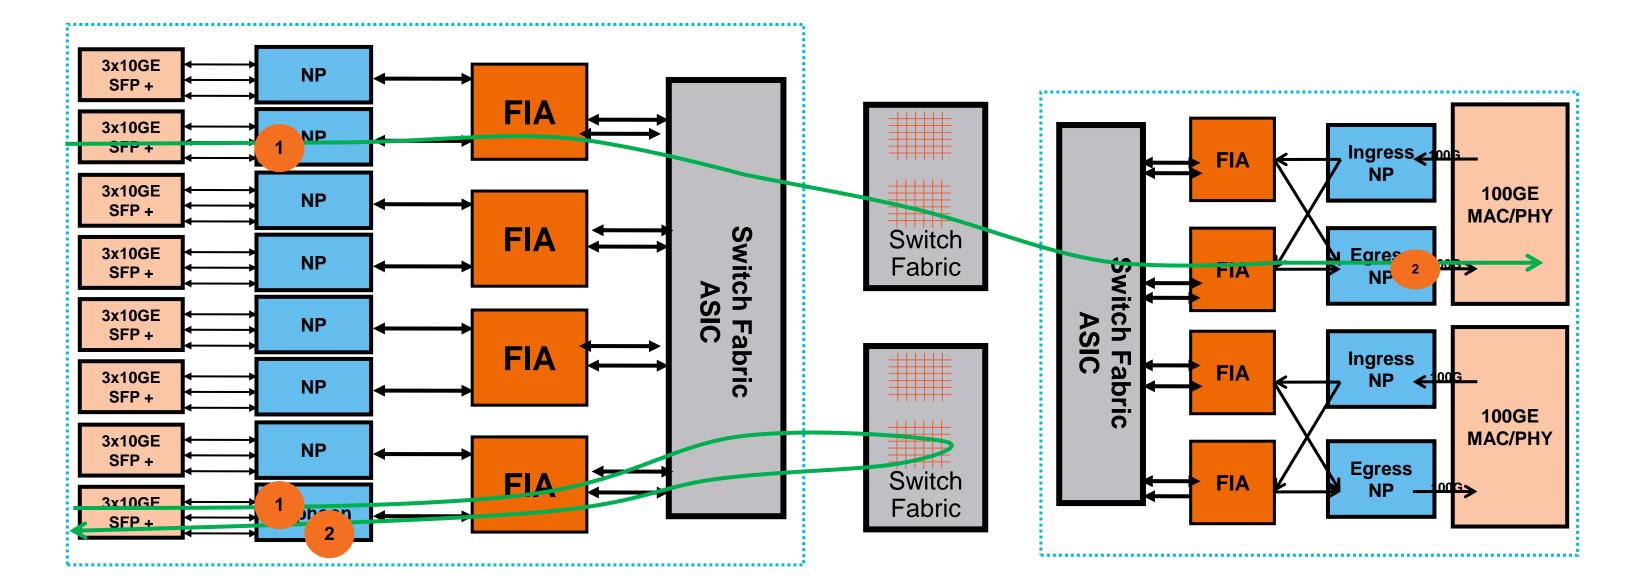
NP Architecture Details



- 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 \bullet interface scale
 - Same lookup memory for same system wide scale \rightarrow mixing different variation of LCs doesn't impact system wide scale



IOS-XR 2-Stage Forwarding Overview For high system scale



Uniform packet flow for simplicity and predictable performance

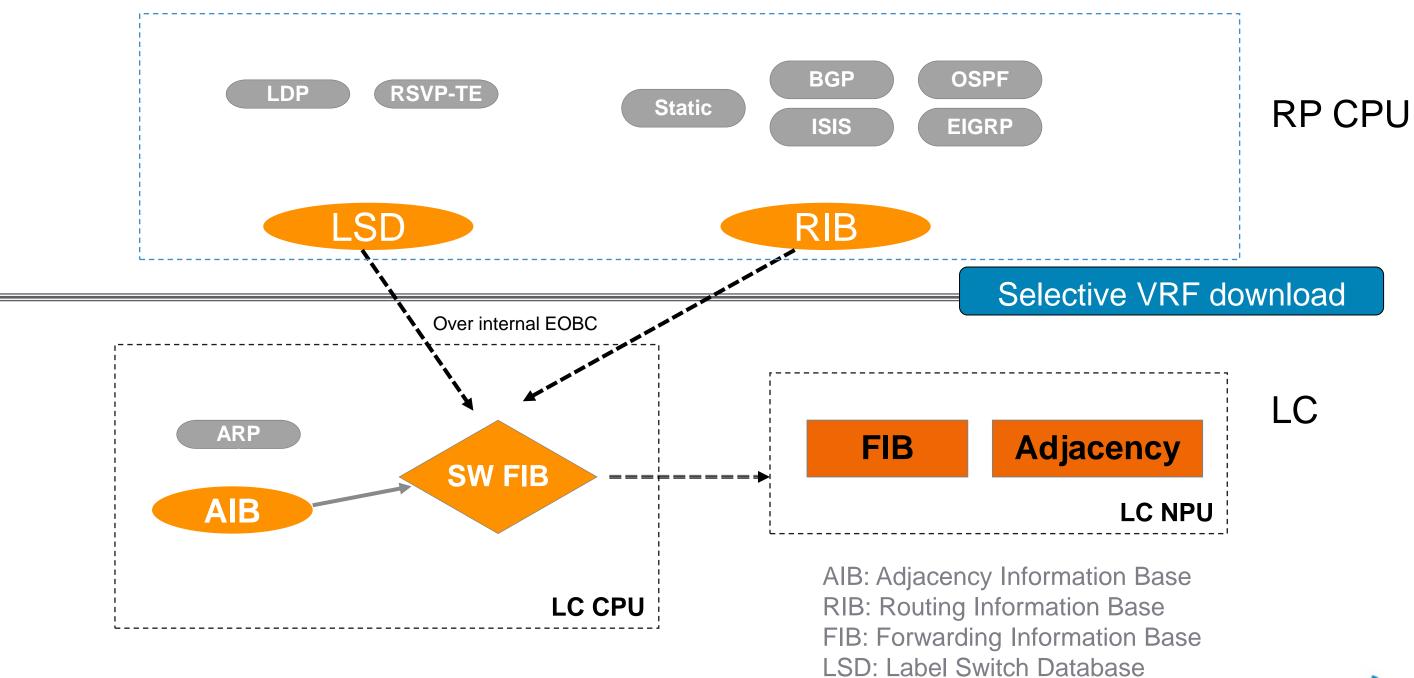
BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved.



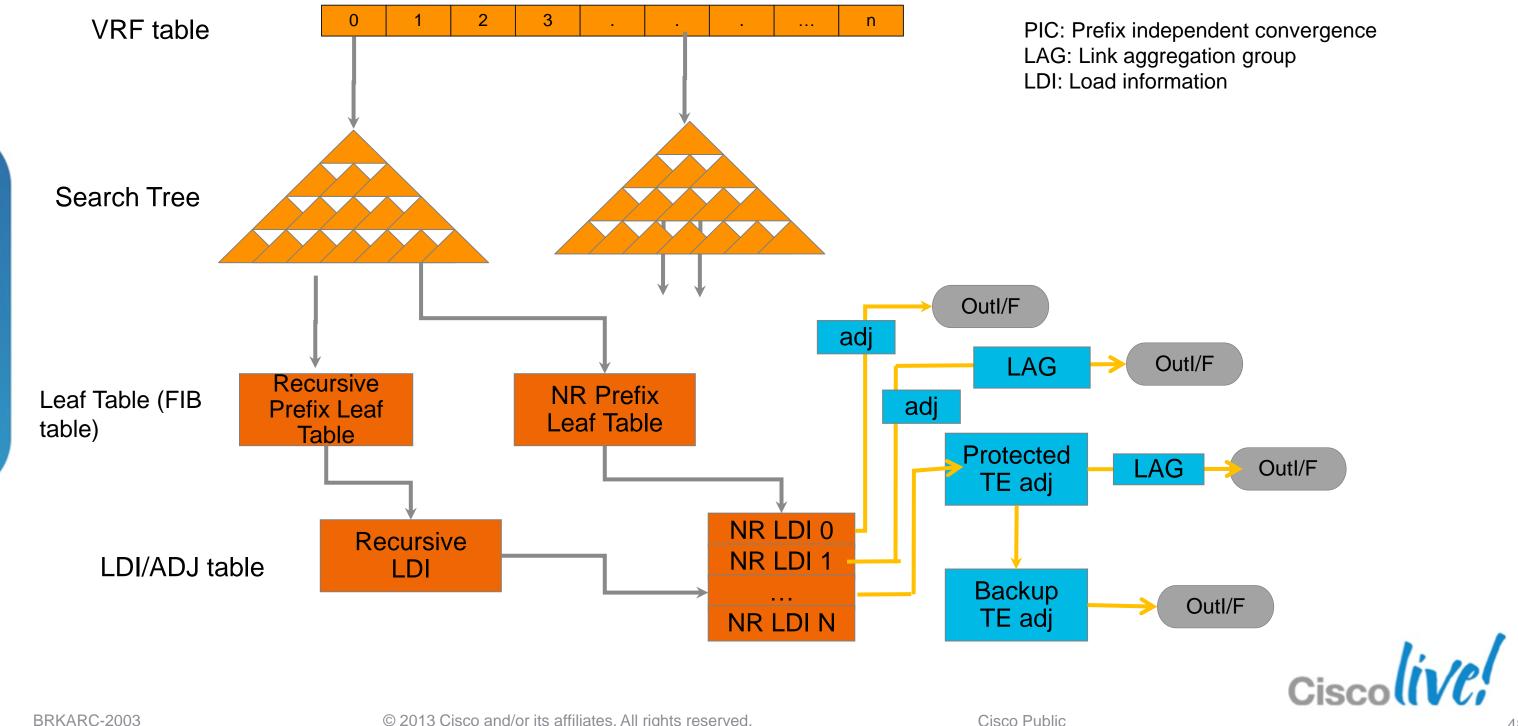


L3 Control Plane Overview





L3 Forwarding – Hierarchical Data Structure for prefix independent fast convergence: BGP, TE/FRR, LAG



MAC Learning and Sync

CPU

FIA

FIA

FIA

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

NP

NP

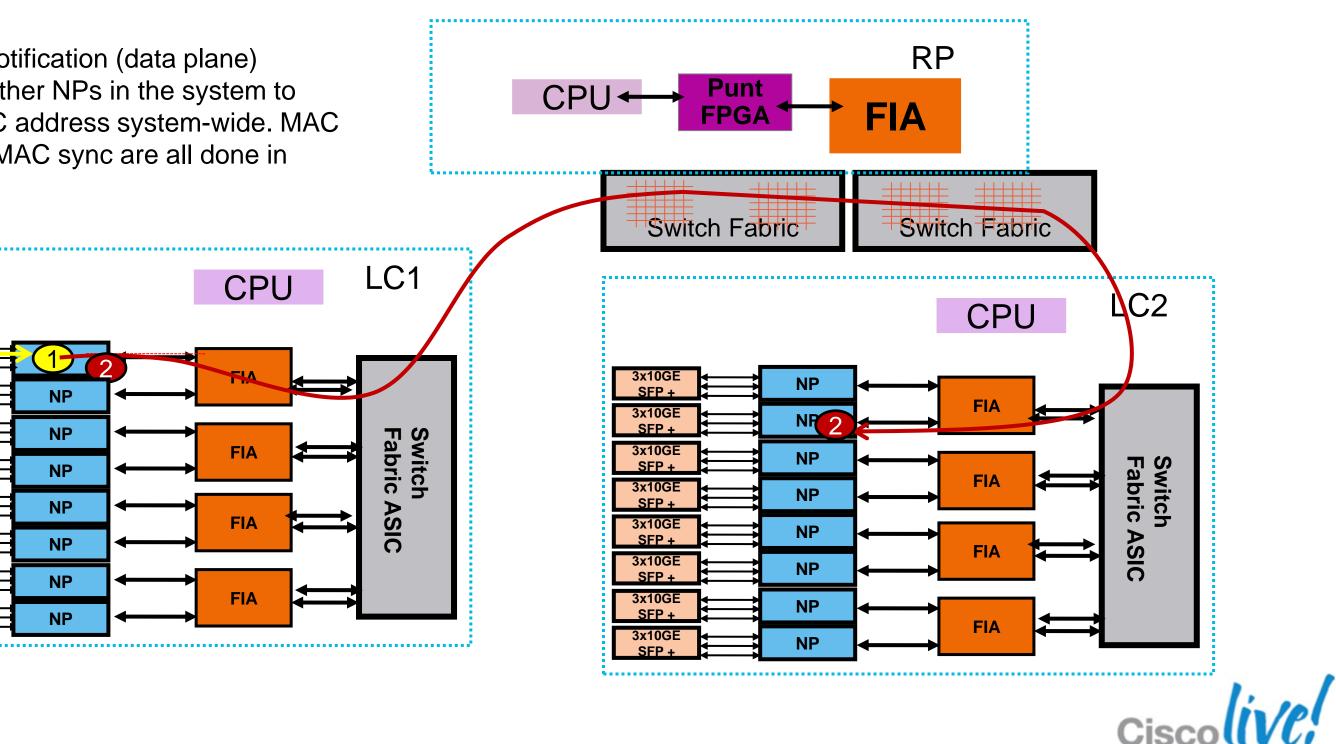
NP

NP

NP

NP

NP



3x10G

3x10GE

SFP +

3x10GE

3x10GE

3x10GE

SFP +

3x10GE

SFP +

3x10GE

SFP +

3x10GE

SFP -

SFP +

SFP +

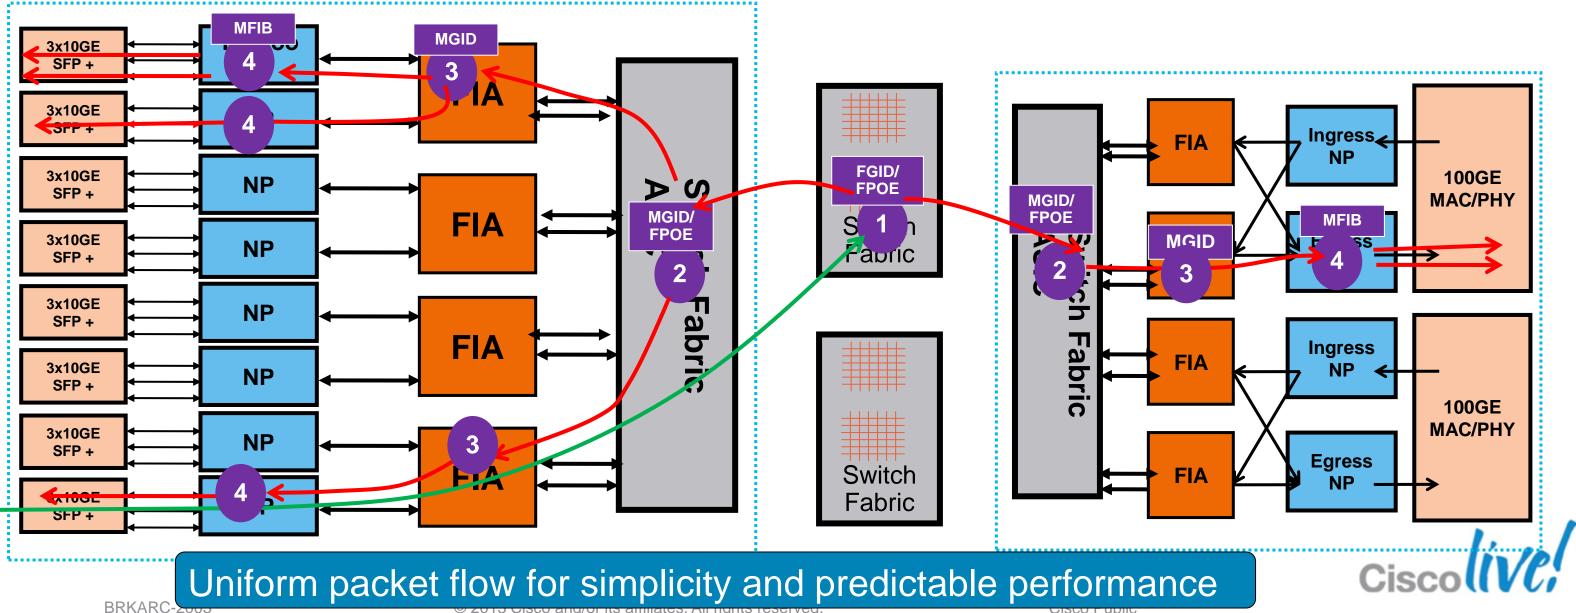
Hardware based MAC learning: ~4Mpps/NP

Internal Multicast Replication Overview

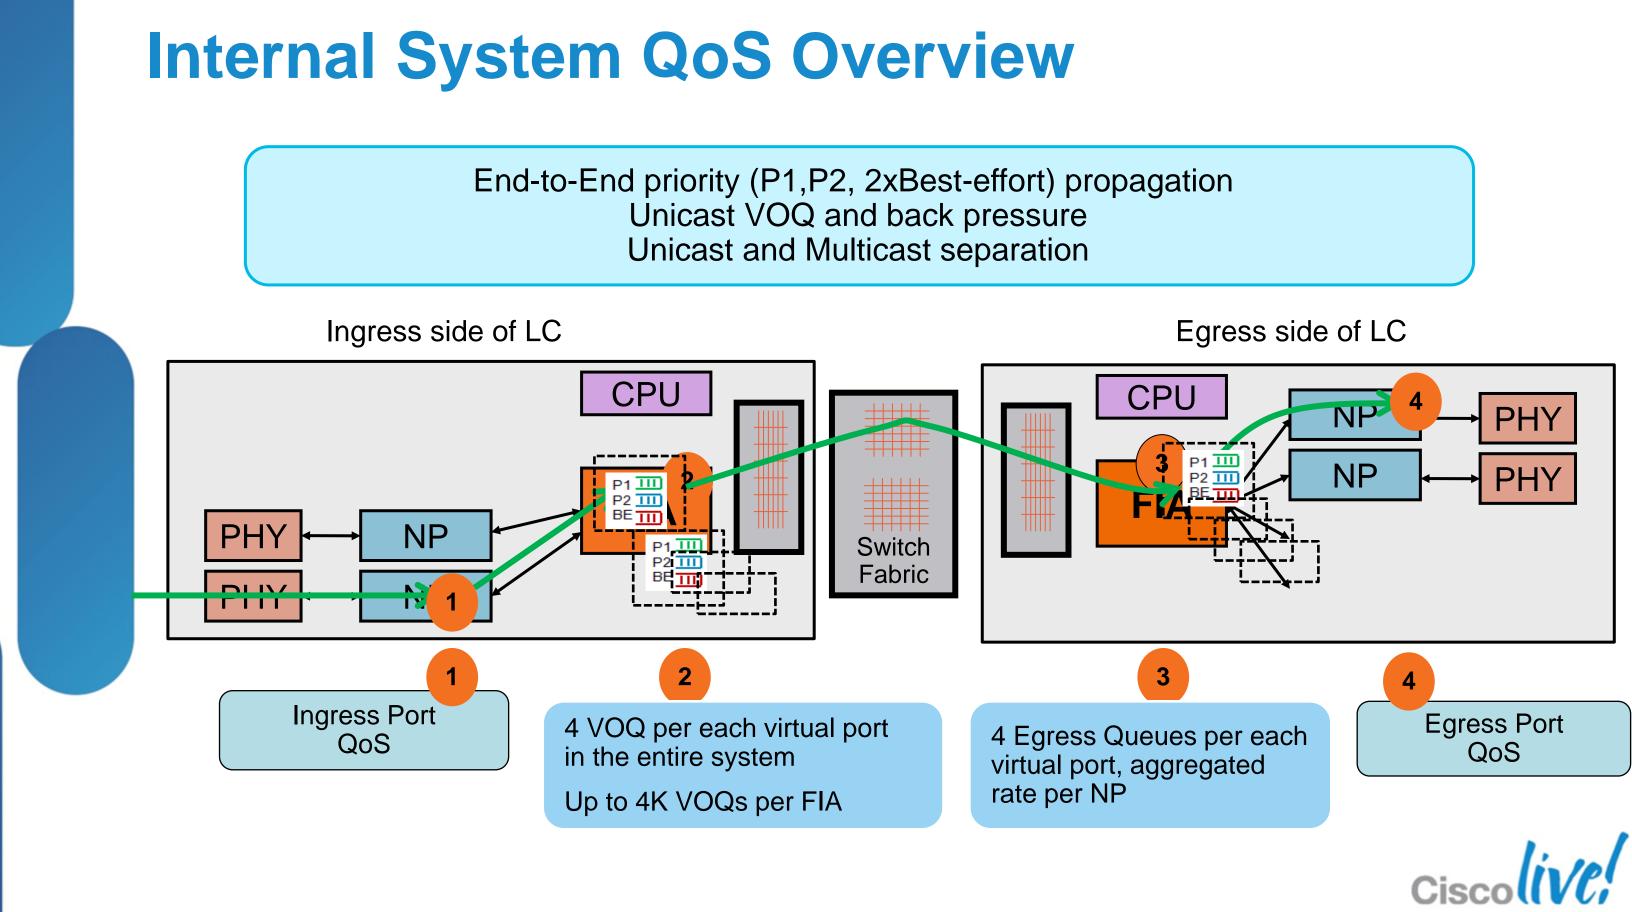
- 2
 - Fabric to LC Replication
 - LC fabric to FIA Replication
- **FIA to NP Replication** 3
- **NP to egress port Replication** 4

FGID – Fabric Group ID MGID – Multicast Group ID MFIB – Multicast Forwarding Information Base

Efficiency: replicate if required Line rate: for fully loaded chassis Simple and clean architecture **Predictable performance**







ASR 9000 nV: network Virtualisation

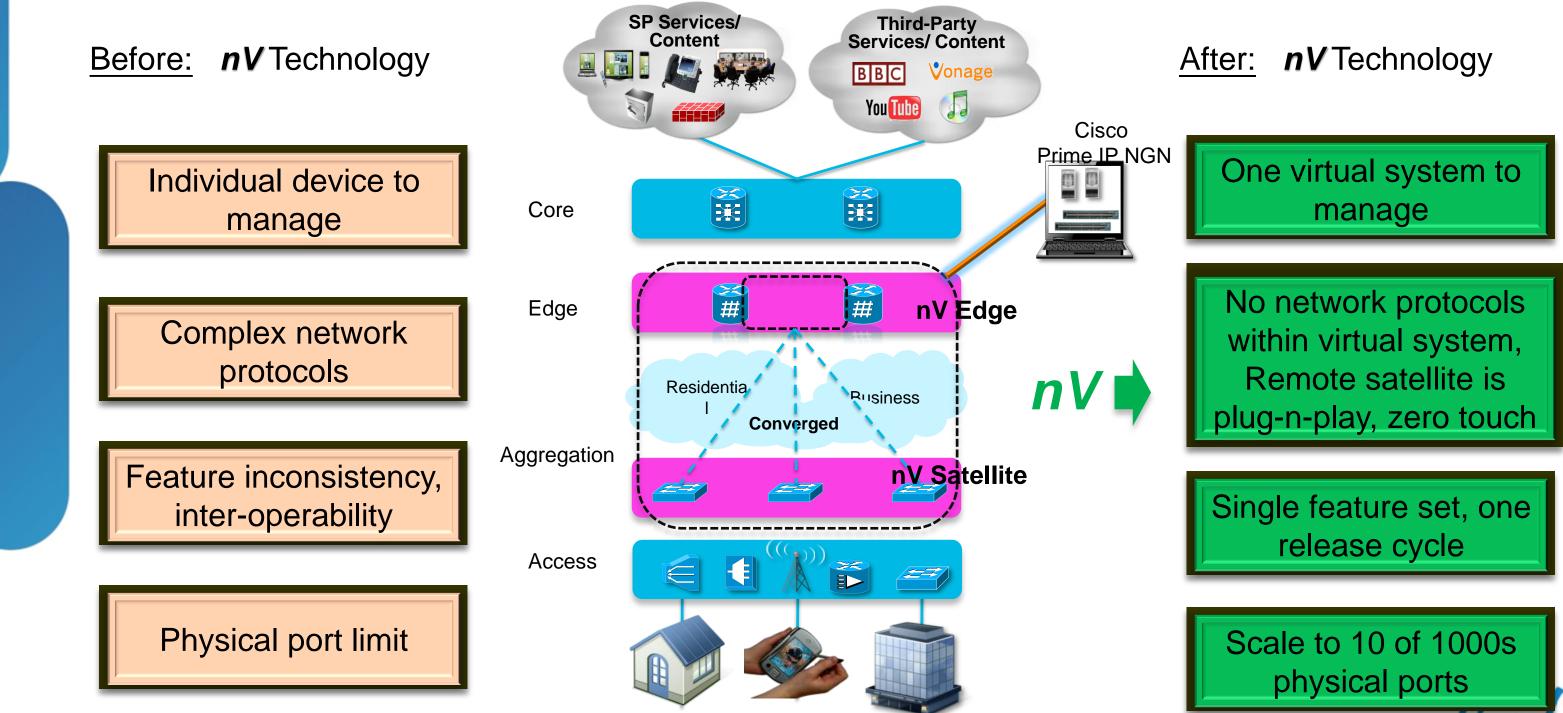








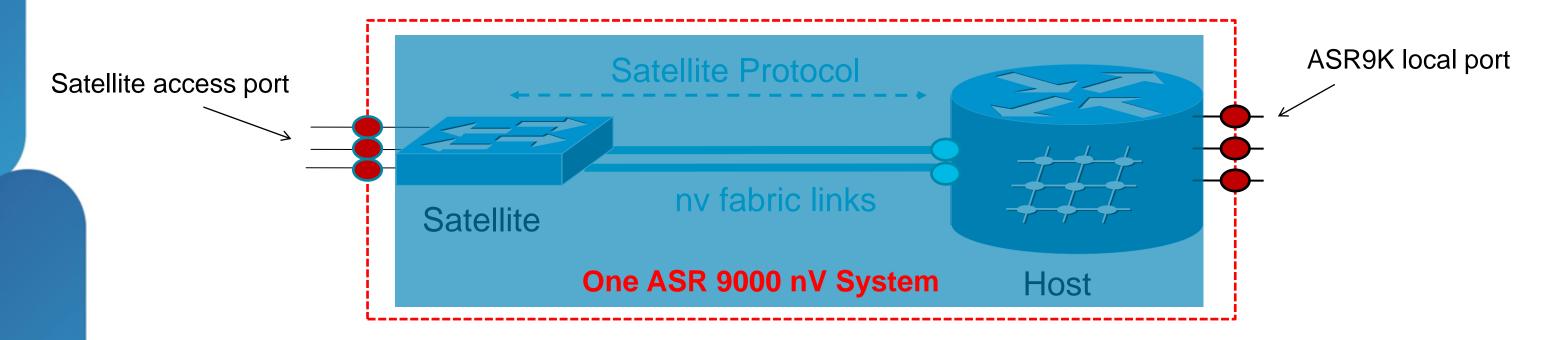
ASR 9000 nV Technology Overview





Cisco

ASR 9000 nV Satellite Overview Plug-n-play, zero touch satellite access device



- Satellite and ASR 9000 Host run satellite protocol for auto-discovery, provisioning and management
- Satellite and Host could co-locate or in different location. There is no distance limit 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)

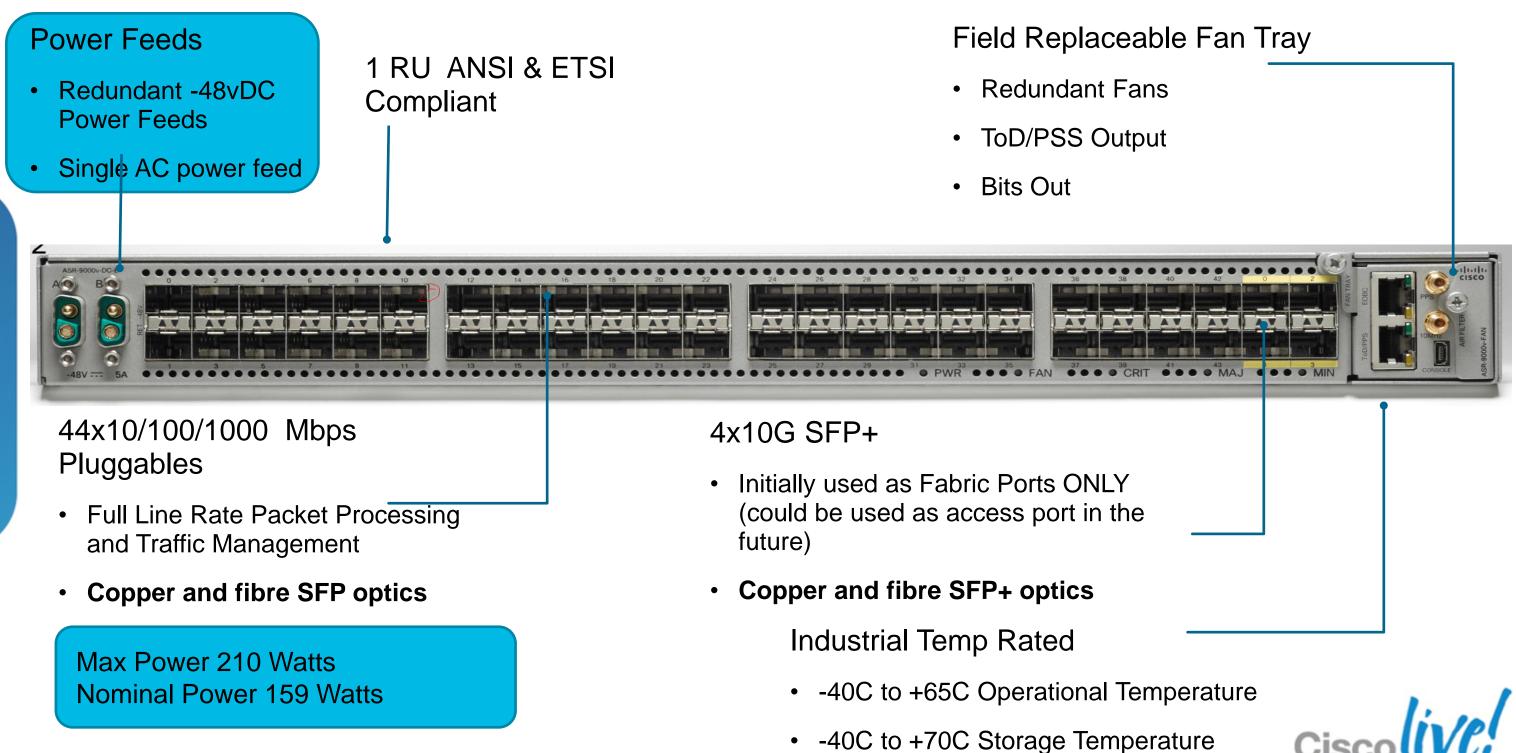
Satellite access port has the same feature parity as ASR9K local port, it works/feels just as local port

BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved.

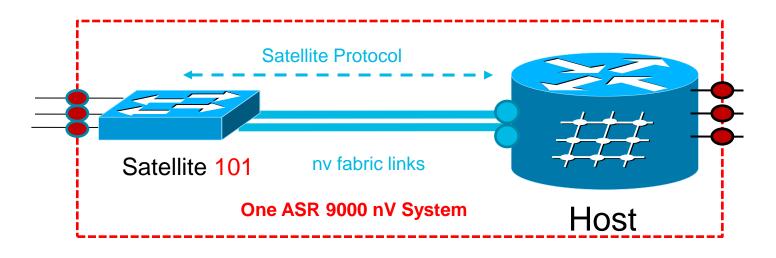


Satellite Hardware: ASR 9000v Overview





Satellite Operation – Initial Configuration



nv

satellite 101 \leftarrow 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



CLI example – Configure Satellite Access Port



Satellite access port configuration examples

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

Local port configuration examples

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

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

interface Bundle-ethernet 100 ipv4 address 1.1.1.1 255.255.255.0

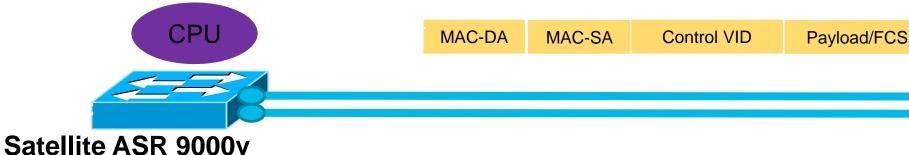
bundle-id 100



interface GigabitEthernet 0/0/0/2



Satellite – Host Control Plane Satellite discovery and control protocol



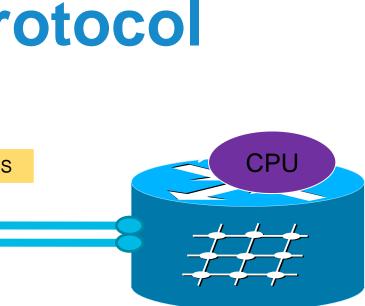
Discovery Phase

• 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 for the time being. It could move to standard in the future
- Get/ Set style messages to provision the satellites and also to retrieve notifications from the satellite

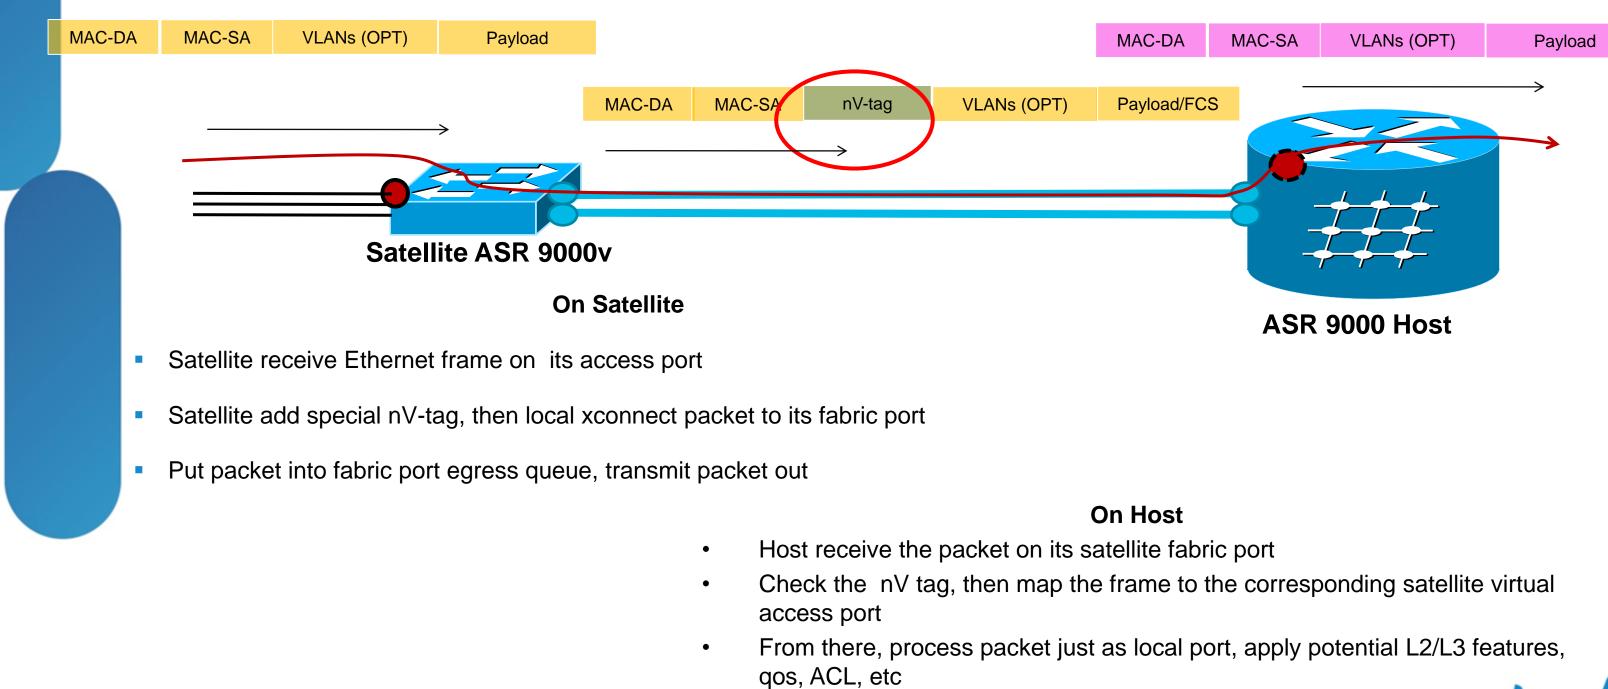
Standardisation is considered for future



ASR 9000 Host



Satellite – Host Data Plane Encapsulation



• different satellite

© 2013 Cisco and/or its affiliates. All rights reserved.

Packet is forwarded out of local port, or satellite fabric port to same or

Satellite Operation – Management

- Normal operation, like show CLIs are done on the Host directly, for example
 - Satellite inventory reporting, environmental monitoring
 - Interface counts, stats
 - **SNMP MIB**
 - NMS support, including ACT, ANA/ PRIME
- Low level debug could still be done directly on the satellite device

User can telnet into satellite via out-of-band management console, or in-band from Host, and run regular show/debug CLIs





Satellite Operation – Image Upgrade

RP/0/RSP0/CPU0:R1#sh install active Node 0/RSP0/CPU0 [RP] [SDR: Owner] Boot Device: disk0: Boot Image: /disk0/asr9k-os-mbi-4.2.1.22K.CSCtz10483-0.0.4.i/0x100305/mbiasr9k-rsp3.vm Active Packages: disk0:asr9k-px-4.2.1.22K.CSCtz10483-0.0.4.i disk0:asr9k-satellite-px-4.2.1.22K ← satellite image PIE disk0:asr9k-mini-px-4.2.1.22K disk0:asr9k-mpls-px-4.2.1.22K disk0:asr9k-mcast-px-4.2.1.22K disk0:asr9k-fpd-px-4.2.1.22K

RP/0/RSP0/CPU0:R1#install nv satellite?

<100-65534> Satellite ID All active satellites all

Golden satellite image is always there in satellite flash card for image fall back

RP/0/RSP0/CPU0:R1#install nv satellite 100?

activate Install a new image on the satellite, transferring first if necessary transfer Transfer a new image to the satellite, do not install yet

RP/0/RSP0/CPU0:R1#install nv satellite 100 active

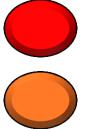




Satellite Operation – Rack Plug and Go

Initial satellite configuration (on ASR9K Host



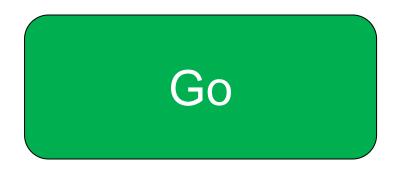


Critical Error LED ON \rightarrow bad hardware, RMA Major Error LED ON \rightarrow Unable to connect to ASR9K host

- Missing the initial satellite configuration? lacksquare
- L1 issue, at least one of the uplink port light green?
- Security check (optional), is the satellite SN# correct?

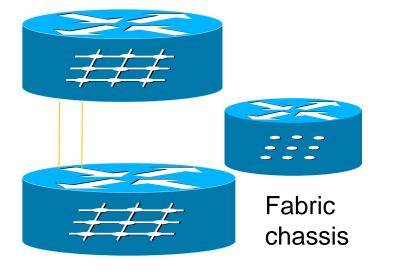
Status light green \rightarrow ready to go, satellite is fully managed by Host







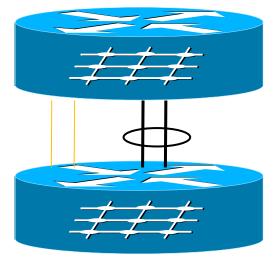
ASR 9000 nV Edge Overview Simple Network Dual-homing Resiliency & Extensiable Node Capacity



CRS Multi-Chassis

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

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

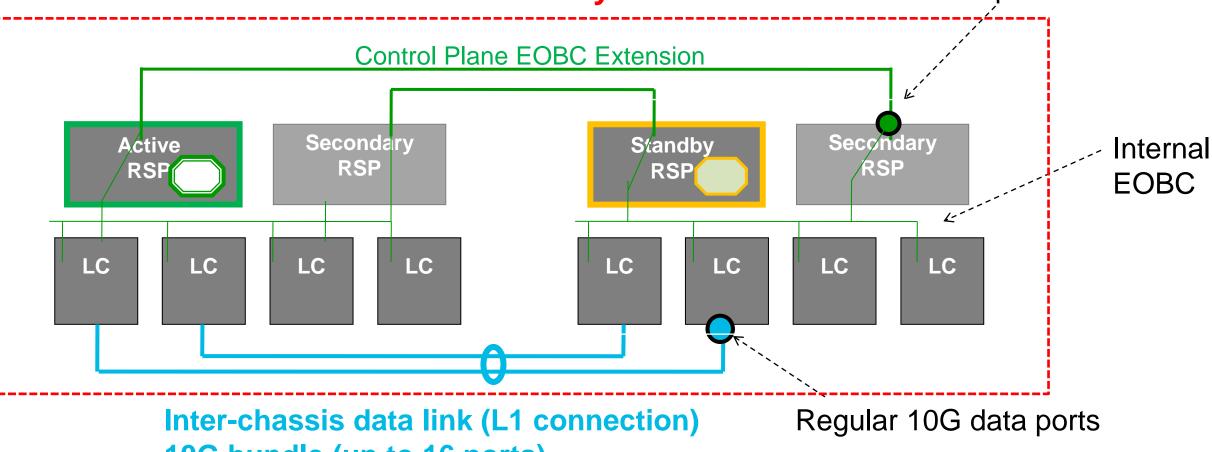


ASR 9000 nV Edge



nV Edge Architecture Overview

One Virtual ASR 9000 nV System



10G bundle (up to 16 ports)

- 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)

RSP440 nV EOBC ports for control plane connection



nV Edge Configuration



Configure nv Edge globally

nv

```
edge-system
serial FOX1439G63M rack 0
```

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 ③

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 !!!

BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved.





IOS XR Software Overview











Cisco IOS-XR High Level Themes

- Modular—Runtime SW upgrade/downgrade support
- Distributed—scalable with multi chassis support
- Platform Independent—POSIX compliant
- Management Interface—Unified Data Model (XML)
- High Availability—Hot Standby and Process Restart
- Security—Control, Data and Management Plane

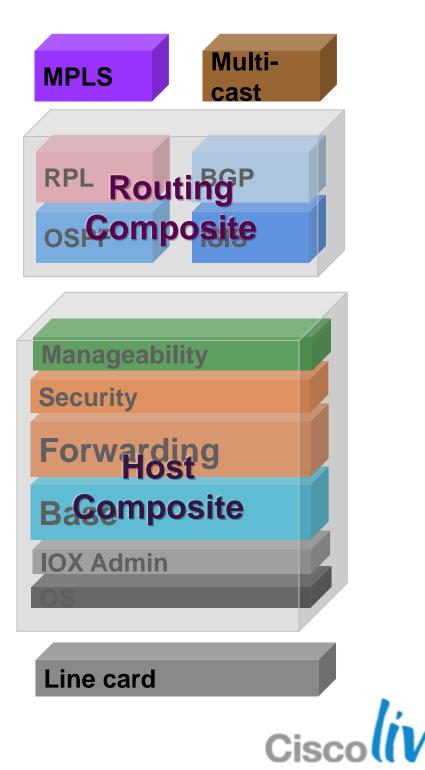




Cisco IOS-XR Software Modularity

- Ability to upgrade independently MPLS, Multicast, Routing protocols and Line Cards
- Ability to release software packages async
- Ability to have composites into one manageable unit if desired
- Notion of optional packages if technology not desired on device (Multicast, MPLS)



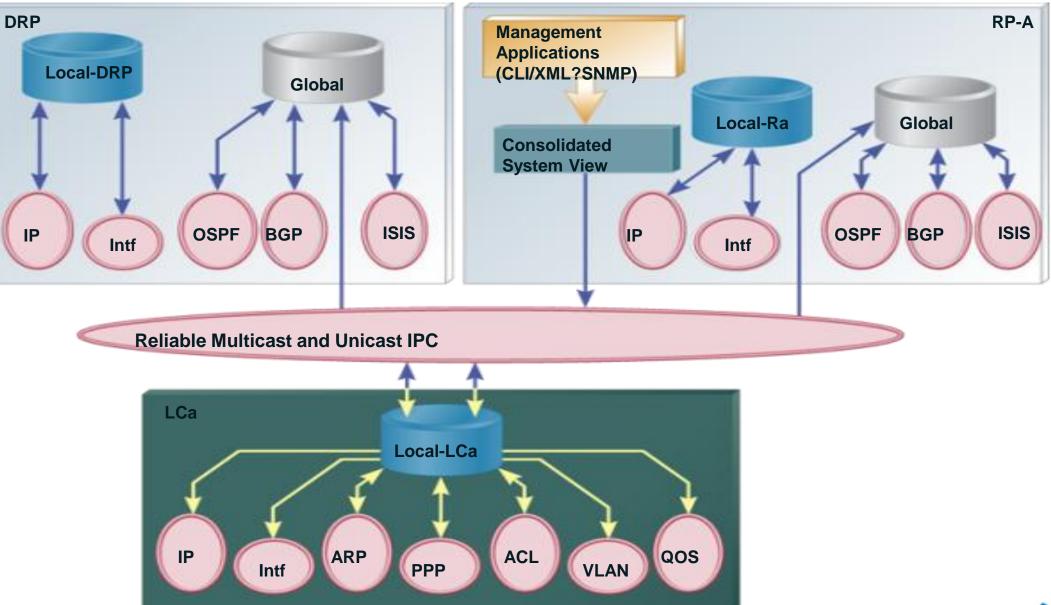


Distributed In-Memory Database (IMDB)

Reliable Multicast IPC improves scale and performance

Distributed data management model improves performance and Scale

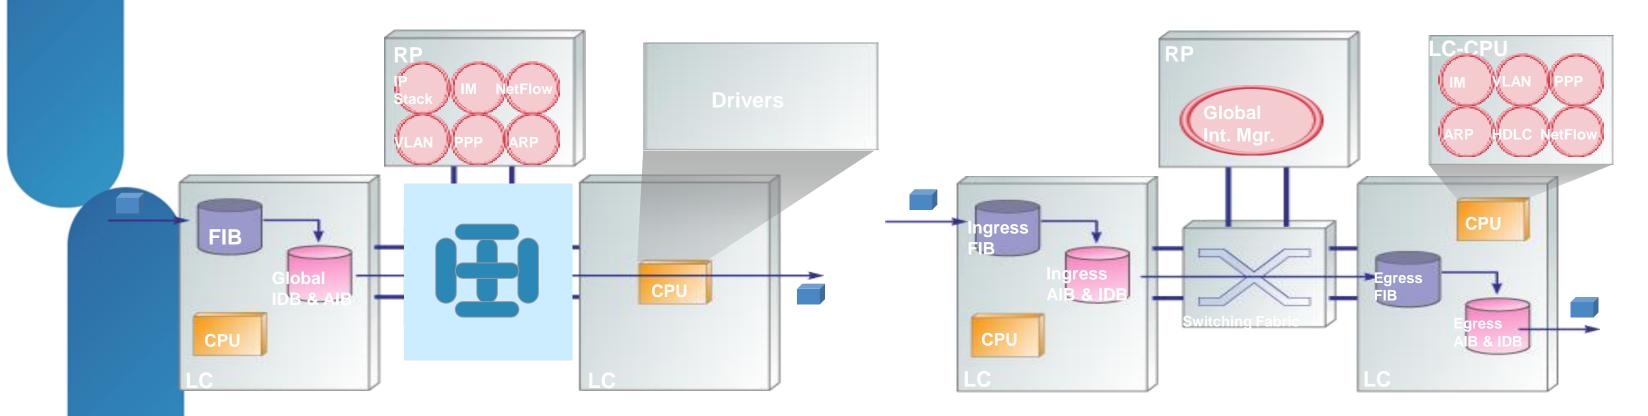
Single Consolidated view of the system eases maintenance







IOS XR Distributed Forwarding Infrastructure



Single Stage Forwarding

- Single global Adjacency Information Base (AIB) distributed to all line cards
- Single global Interface Management DB distributed to all line cards
- Only Ingress FIB forces forwarding features to be run in RP

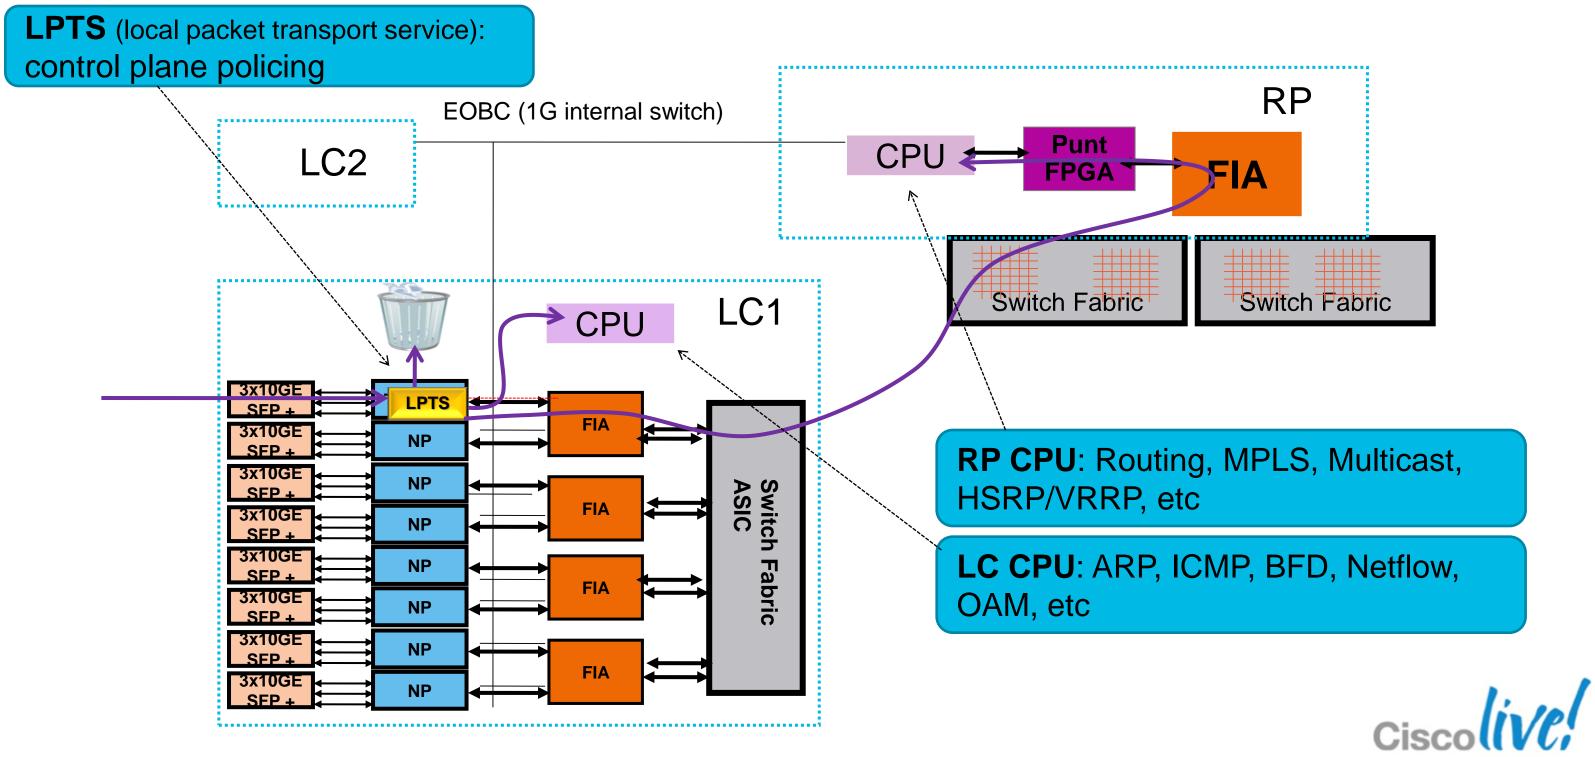
Two Stage Forwarding

- Each line card has independent AIB only for local interfaces
- Each line card has independent Interface DB for local interfaces
- LCs **Cisco** Public

BRKARC-2003

Both Ingress and Egress FIB – allows forwarding features to be independently rul

ASR 9000 Fully Distributed Control Plane



BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved.



IOS XR HA Software Design Principles

NSR (GR)

Non-Stop Forwarding

Separate Control and Data Planes

RP/DRP Redundancy Active/Standby Failover

Process Restartability: Active State Check pointing

All subsystems: Separate Address Spaces memory faults affect only 1 process, recovery = restart process

BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved



Process Restart

- Microkernel includes minimal functionality
- Non-kernel processes can be restarted
- Critical to HA, ISSU, and MDR functions
- Restarting many processes can be tricky Dependent processes may also need to restart





Software Maintenance Updates (SMU)

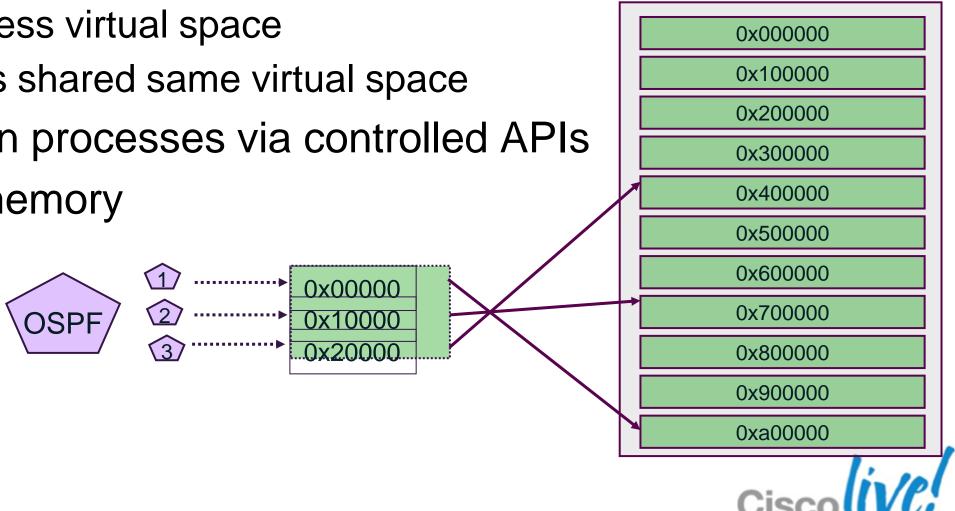
- Allows for software package installation/removal leveraging on Modularity and Process restart
- Redundant processors are not mandatory (unlike ISSU) and in many cases is non service impacting and may not require reload.
- Mechanism for
 - delivery of software features (e.g. Multicast, MPLS) delivery of critical bug fixes without the need to wait for next maintenance release





Protected Process Memory Space

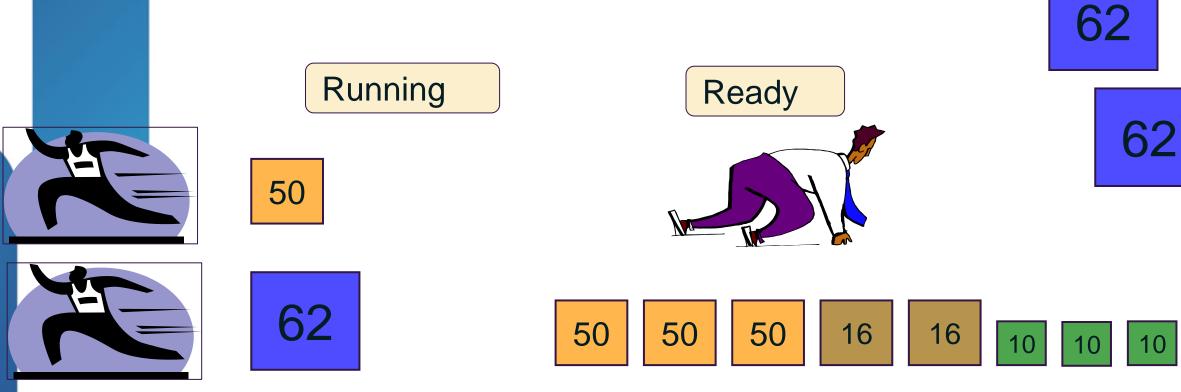
- Each process has a virtual memory space Kernel/MMU maps virtual address to physical address (at page level) Threads share the memory space
- One process cannot corrupt another' s memory Process can only access virtual space In IOS – all processes shared same virtual space
- Communication between processes via controlled APIs
- Limited use of shared memory

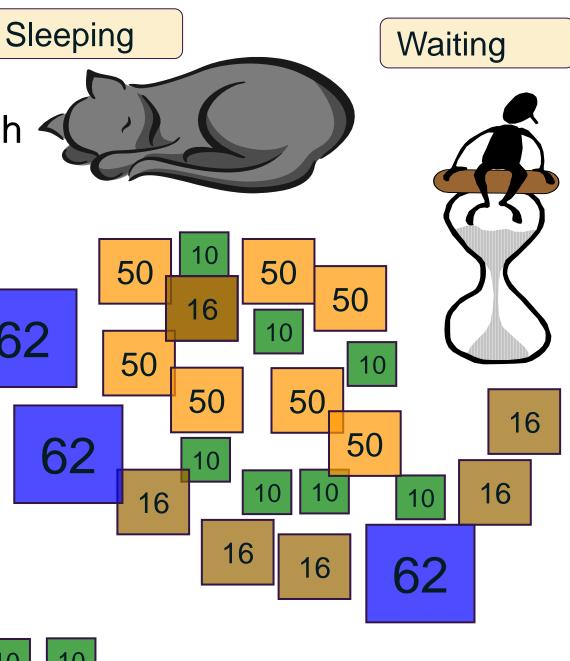




Preemptive Multitasking

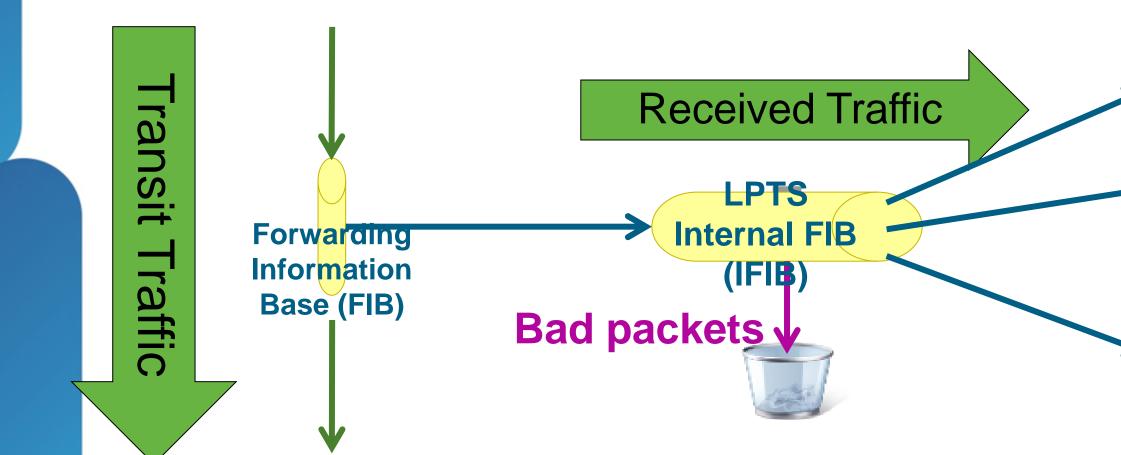
- Default priority is 10
- Higher priority processes can interrupt In IOS, must wait for running process to finish
- FIFO within same priority
- Threads run while parent process is running







Local Packet Transport Services (LPTS): **Protection of Control Plane**



LPTS enables applications to reside on any or all RPs, DRPs, or LCs Active/Standby, Distributed Applications, Local processing IFIB forwarding is based on matching control plane flows

Built in dynamic "firewall" for control plane traffic

LPTS is transparent and automatic





Local Stack on LC



Local Packet Transport Service

- There is no longer a single RP
- IOS XR is a fully distributed operating system with applications running in multiple physical locations
- Filters and polices (in hardware) local 'receive' packets and sends them only to the nodes that need them
- Packet rate correlates with trust
- Defaults for 90% of customers, configurability for the rest
- Handles fragments, also checks TTL/hop count
- High Availability for NSR (Non-Stop Routing)



Local Packet Transport Service

- LPTS has Hardware policers on line cards to limit traffic sent to local or remote nodes
- LPTS entries in TCAM classifies packets to select a policer to apply
- The policer value can be tuned to 0 (to drop all packet matching) classification criteria)
- Polices on protocol (BGP, OSPF, SSH) and flow state (BGP) established, BGP configured, and BGP listen)
- Policing done on the LC Hardware ASIC before packets hit **RP/LC**CPU
- All filters are automatically and dynamically installed by the IOS XR infrastructure

BRKARC-2003

© 2013 Cisco and/or its affiliates. All rights reserved

Summary **Cisco ASR 9000: Next-gen SP+DC Edge Routing**

Carrier Class

IOS-XR true modular OS

Superior Multicast replication

Advanced internal system QoS

And more

Full HW portfolio: chassis, LC, port

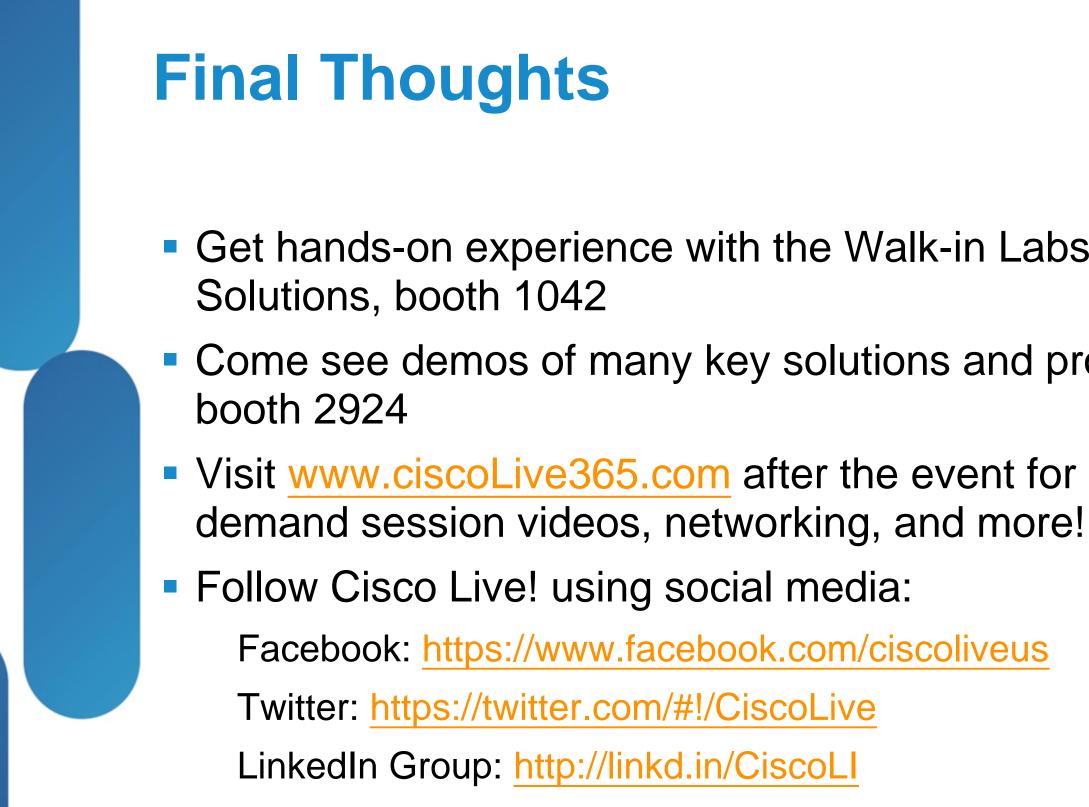
nV – expand the system capacity and physical ports

Fully distributed IOS-XR control plane

HW MAC learning, prefix independent fast convergence, and more ...

Scalable System Architecture

IOS-XR 2-stage forwarding



Get hands-on experience with the Walk-in Labs located in World of

Come see demos of many key solutions and products in the main Cisco

Visit www.ciscoLive365.com after the event for updated PDFs, on-



Q & A











Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm





communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button. www.ciscoliveaustralia.com/portal/login.ww



Don't forget to activate your Cisco Live 365 account for access to all session material,



CISCO

