LET'S
BUILD
TOMORROW
TODAY

Nexus 9000 Series Switch Architecture

Lilian Quan

Principal Engineer

INSBU, Nexus 9000 and ACI

BRKARC-2222

Cisco live!

What is our Goal Today?

What is our Goal Today?

The Nexus 9000 series is the newest members family of Cisco's Data Centre switches designed for the current and evolving needs of the cloud and virtualized Data Centre. This family of switches are designed to be used in standard layer 2 and layer architectures, evolving VXLAN and dev-ops based architectures and will form the foundation of the Application Centric Infrastructure (ACI) that was developed by Cisco's Insieme acquisition.

This session will provide a technical discussion on the architecture of the Nexus 9500 modular and Nexus 9300 TOR switches including details on hardware architecture, packet forwarding, VXLAN overlay forwarding, QoS and buffering as well as newer capabilities such as the policy, security and service insertion features provided by the ACI fabric

Agenda – Nexus 9000 Architecture

- Nexus 9000 Series Switch Hardware
  - Nexus 9000 Hardware Overview
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9500 Buffer and Queuing
  - Nexus 9300 Hardware Architecture
  - Nexus 9300 Buffer and Queuing
- Nexus 9000 and 40GE
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 Open NX-OS
- ACI & Nexus 9000

Agenda – Nexus 9000 Architecture

- Nexus 9000 Series Switch Hardware
  - Nexus 9000 Hardware Overview
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9500 Buffer and Queuing
  - Nexus 9300 Hardware Architecture
  - Nexus 9300 Buffer and Queuing
- Nexus 9000 and 40GE
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 Open NX-OS
- ACI & Nexus 9000

Cisco Nexus 9000 Series Switches

High Performance 10/40/100 Gbps Switch Family

SCALABLE 1 GE/10 Gbps/40 Gbps/100 GE PERFORMANCE

Nexus® 9300

FCS Q12014 — 48 1/10G SFP+ & 12 QSFP+

FCS Q1 2014 — 96 1/10G-T & 8 QSFP+

FCS Q1 2014 — 12-port QSFP+ GEM

Nexus 9500

FCS Q4 2013 — Aggregation line card 36 40G QSFP+

FCS Q1 2014 — ACI Ready Leaf Line Card 48 1/10G-T & 4 QSFP+

FCS Q1 2014 — ACI-ready Leaf line card 48 1/10G SFP+ & 4 QSFP+

FCS Q4 2013 — C9500 8-Slot

FLEXIBLE FORM FACTORS CAN ENABLE VARIABLE DATA Centre DESIGN AND SCALING

PERFORMANCE | PORTS | PRICE | POWER | PROGRAMMABILITY

Merchant and Custom ASICs (Merchant+)

- ✧ Merchant+ Strategy

Merchant ASIC --- NFE (Broadcom T2)

Custom ASIC --- Cisco ALE (ACI Leaf Engine), ASE (ACI Spine Engi

- ✧ Best Performance and Functionalities
- ✧ Optimal Pricing

	NFE	ALE	ASE
ASIC Technology	40 nm	28 nm	28nm
40Gbps Ports	32 (24)	24 (24)	42(42)
Buffer (MB)	12 MB	40 MB	23 MB
L2/ L3	L2/ L3	L2/ L3	L3

Merchant +

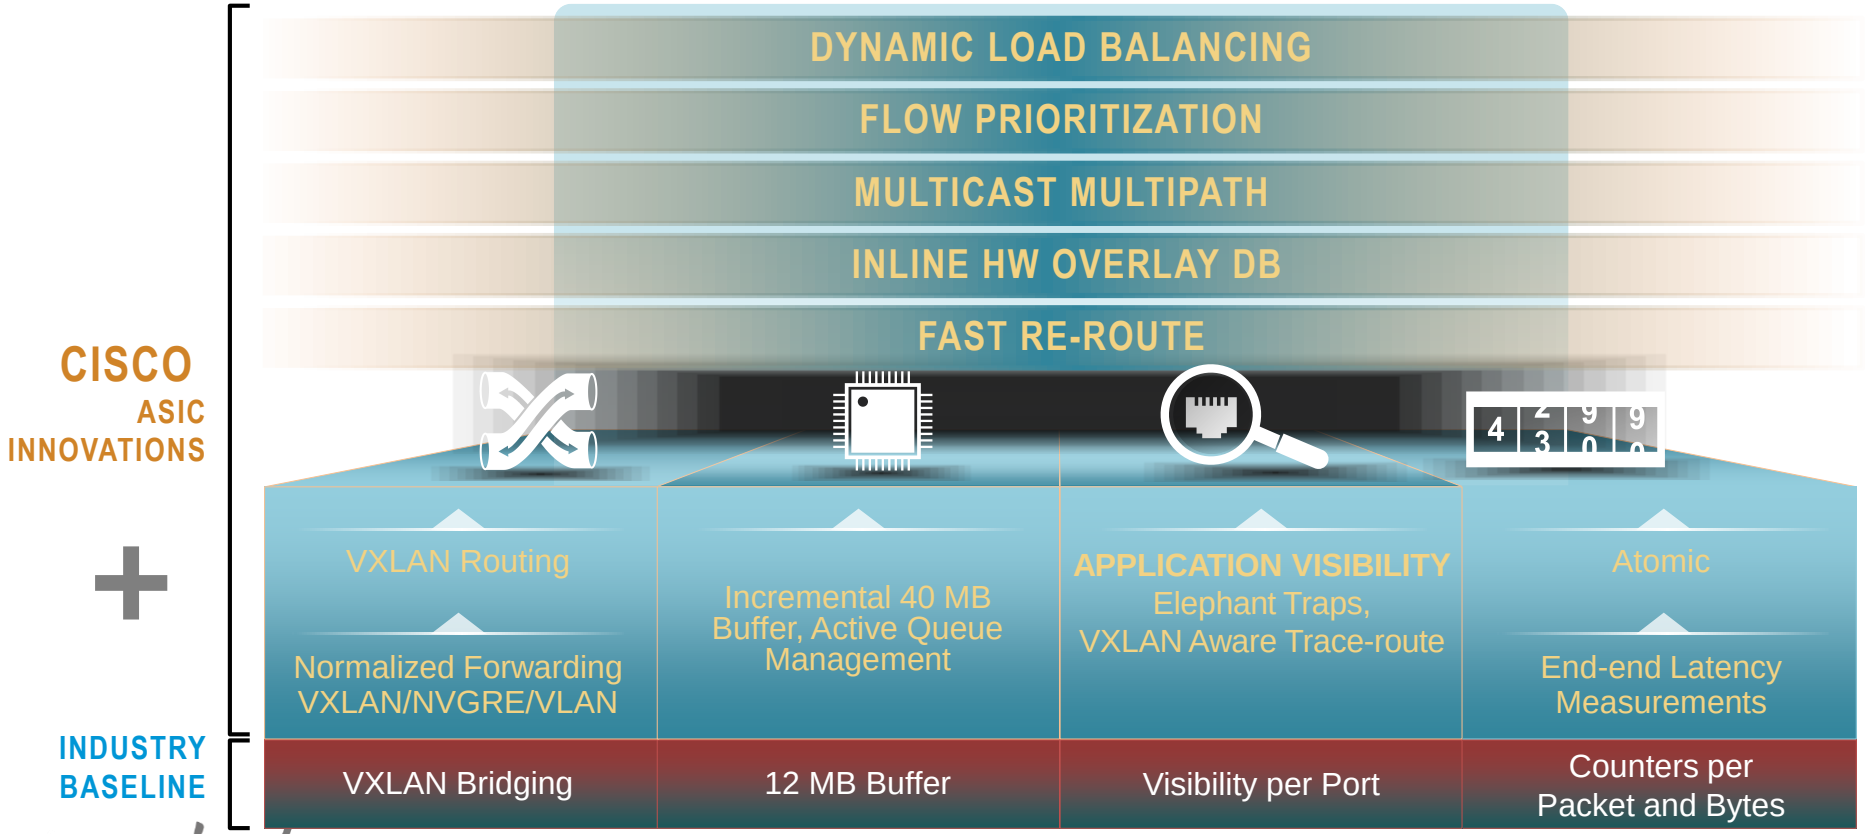

Agenda – Nexus 9000 Architecture

- Nexus 9000 Series Switch Hardware
  - Nexus 9000 Hardware Overview
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9500 Buffer and Queuing
  - Nexus 9300 Hardware Architecture
  - Nexus 9300 Buffer and Queuing
- Nexus 9000 and 40GE
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 Open NX-OS
- ACI & Nexus 9000

Nexus 9500 Platform Architecture

Nexus® 9508 Front View

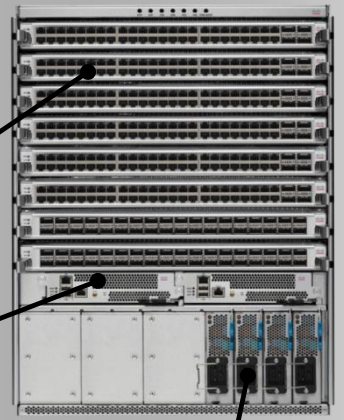

8 line card slots
Max 3.84 Tbps per slot
duplex

Redundant
supervisor engines

3000 W AC power supplies
2+0, 2+1, 2+2 redundancy
Supports up to 8 power supplies

Nexus 9508 Rear View

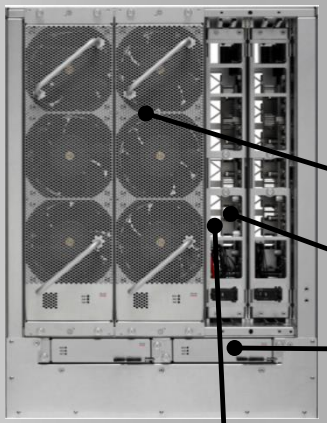

3 fan trays, front-to-back airflow

3 or 6 fabric modules
(behind fan trays)

Redundant system controller cards

No mid-plane for
LC-to-FM connectivity

Chassis Dimensions: 13 RU x 30 in. x 17.5 in (HxWxD)

Designed for Power and Cooling Efficiency
Designed for Reliability
Designed for Future Scale

Nexus 9500 Platform Architecture

Overview

- High port density -- 288 x 40 Gbps/Nexus 9508 -- 1152 x 10 Gbps/Nexus 9508
- Layer 2 and Layer 3 line-rate performance
- Low latency -- Up to 3.5 usec on the 36 x 40GE QSFP line card (N9K-X9636PQ)
- High power efficiency --3.5 W/10 Gbps port; 15 W/40 Gbps port
- First modular chassis without a mid-plane -- Unobstructed front-to-back airflow
- VXLAN bridging and routing
- Highly integrated switch and buffer functionality
  - Only 2 to 4 ASICs per line card
  - No buffer bloat
  - Mix of 28 nm Cisco® and 40 nm Broadcom ASICs

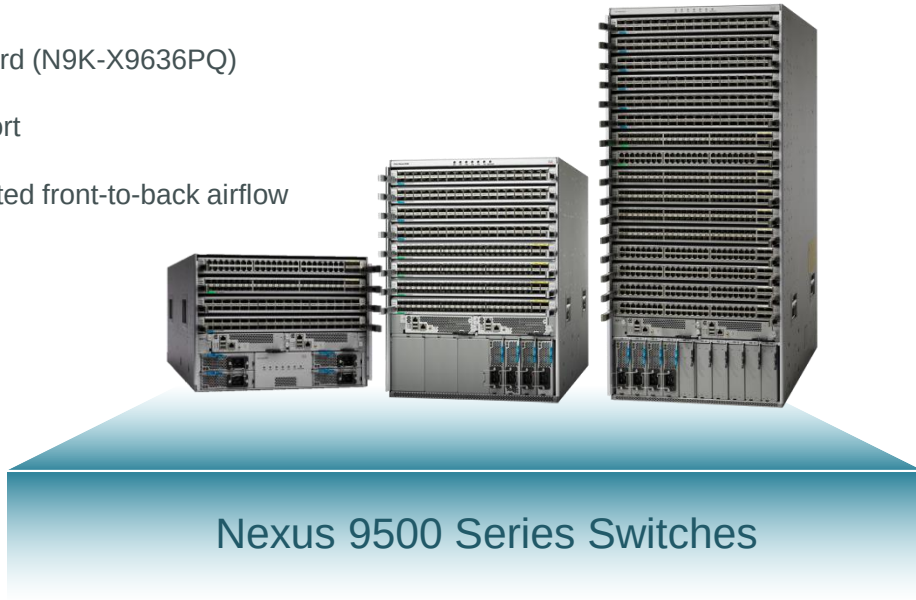

Nexus 9500 Platform Architecture

Chassis Design: No Mid-Plane

- Designed for:
  - Power & Cooling Efficiency
  - Designed for Reliability
  - Designed for Future Scale

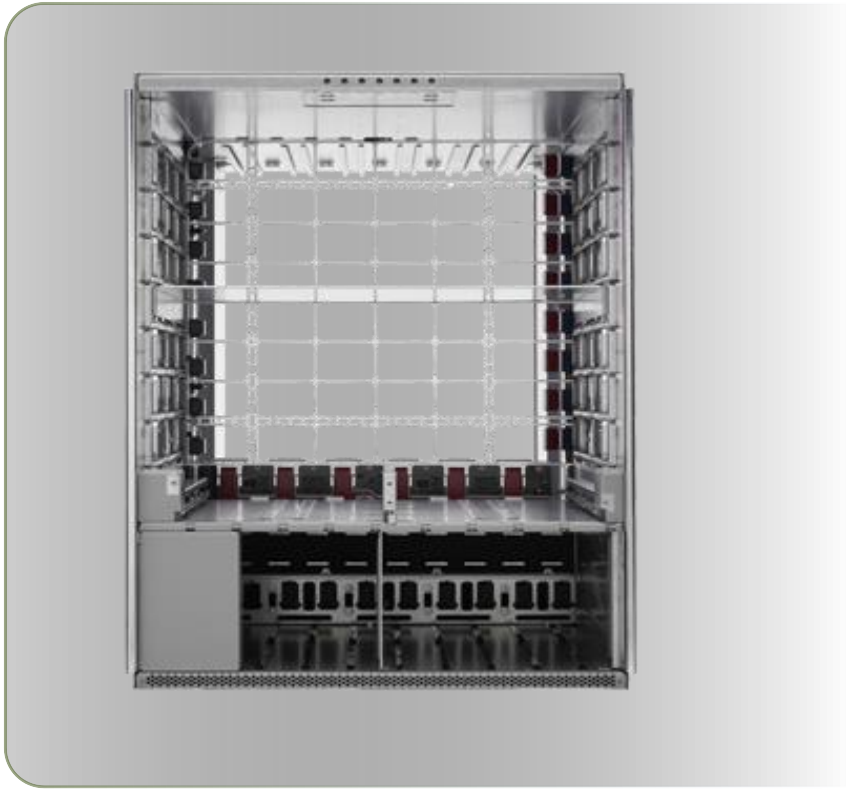

Nexus 9500 Platform Architecture

Supervisor Module Sup-A

- Redundant half-width supervisor engine
- Performance- and scale-focused
- Range of management interfaces
- External clock input (PPS)

Supervisor Module	
Processor	Romley, 1.8 GHz, 4 core
System Memory	16 GB, upgradable to 48 GB
RS-232 Serial Ports	One (RJ-45)
10/100/1000 Management Ports	One (RJ-45)
USB 2.0 Interface	Two
SSD Storage	64 GB

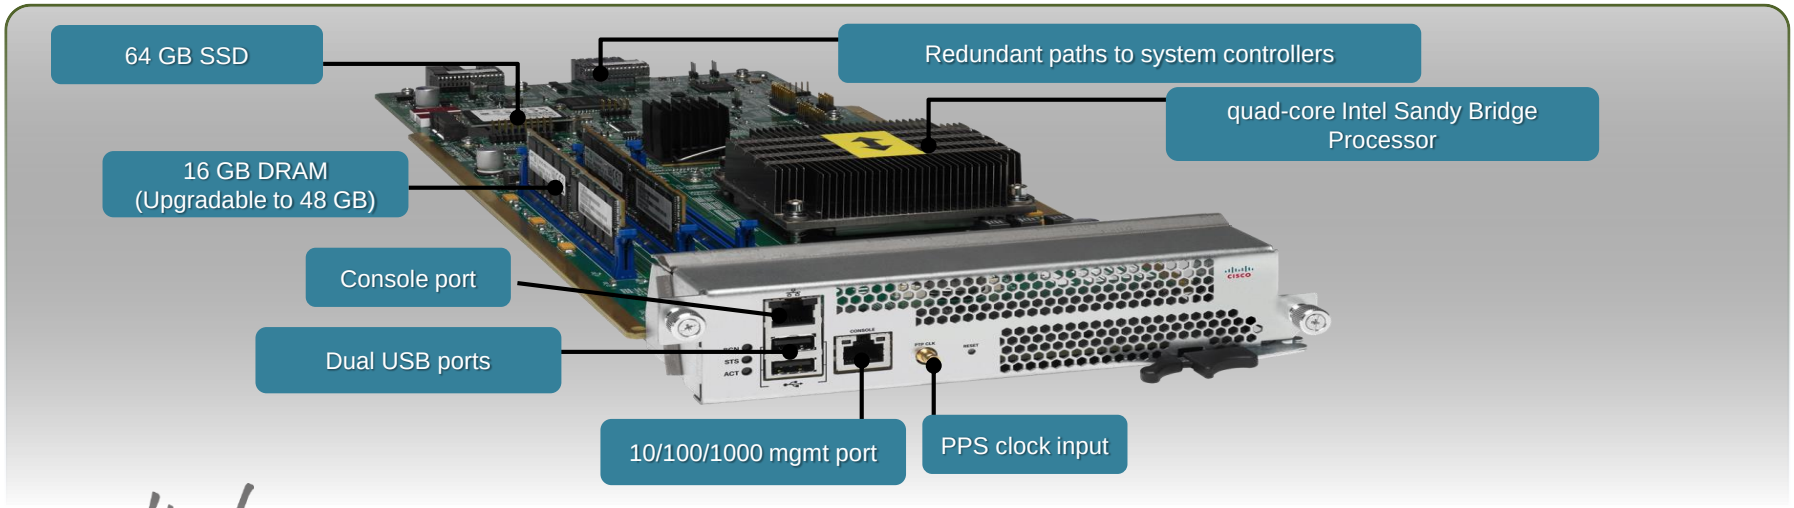

Nexus 9500 Platform Architecture

Supervisor Module Sup-B

- Redundant half-width supervisor engine
- Performance- and scale-focused
- Range of management interfaces
- External clock input (PPS)

Supervisor Module	
Processor	2.1 GHz, 6 cores 2.2GHz IVY Bridge
System Memory	24 GB, upgradable to 48 GB
RS-232 Serial Ports	One (RJ-45)
10/100/1000 Management Ports	One (RJ-45)
USB 2.0 Interface	Two
SSD Storage	256 GB

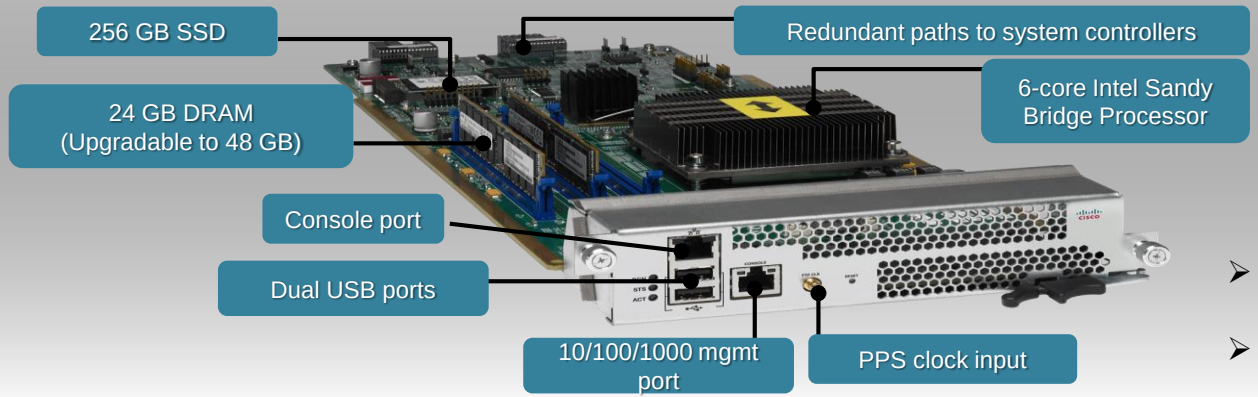

- 50% more CPU power
- 50% more memory space
- 300% more SSD storage

- Increase control protocols performance and convergence time.
- Ready for application intensive deployment

Nexus 9500 Platform Architecture

System Controller Module

- Redundant half-width system controller
- Offloads supervisor from device management tasks
  - Increased system resiliency
  - Increased scale
- Performance- and scale-focused
  - Dual core ARM processor, 1.3 GHz
- Central point-of-chassis control
- Ethernet Out of Band Channel (EOBC) switch:
  - 1 Gbps switch for intra-node control plane communication (device management)
- Ethernet Protocol Channel (EPC) switch:
  - 1 Gbps switch for intra-node data plane communication (protocol packets)
- Power supplies through system management bus (SMB)
- Fan trays

Nexus 9500 Platform Architecture

Power Supplies

3000W AC PSU

- Single 20A input – 220V
- Support for range of international cabling options
- 92%+ Efficiency
- Range of PS configurations
  - Minimum 1 PS
  - 2 PS for fully loaded chassis
  - N+1 redundancy
  - N+N grid redundancy
- Up to 2x head room for future port densities, bandwidth, and optics

** 80 Plus Platinum is equivalent to Climate Saver/ Green Grid Platinum rating*

Nexus 9500 Platform Architecture

Universal 3000W Power Supply

New!

Nexus 9500 Platform Architecture

Power Efficiency by Design

- 1st modular chassis w/o a mid-plane --- Unobstructed front-back airflow
- Platinum rated PS --- 90%-94% power efficiency across all work loads
- Highly integrated switch and buffer functionality --- Only 2 to 4 ASICs per line card

Traffic type	Power (watts)	Fan Speed
No traffic	3233	10%
100% line-rate with IMIX packets	4746	20%
100% line-rate with 64 byte packets	5470	25%

Test Results on a fully loaded Nexus 9508 switch with 288 40GE ports:

Nexus 9500 Platform Architecture

Fan Tray

- 3 fan trays
  - 3 dual fans per tray
  - Dynamic speed control driven by temperature sensors
  - Straight airflow across line cards and fabric modules
- N+1 redundancy per tray

Nexus 9500 Platform Architecture

Fan Tray (Cont'ed)

Fan trays are installed after the Fabric Module. To service a FM, the fan tray must be removed first.

1) If one fan tray is removed, the other two fan trays will speed up 100% to compensate for the loss of cooling power

1) Temperature Sensors in the chassis will shut down components once max temp is reached.

Nexus 9500 Platform Architecture

Fabric Modules

- Interconnect I/O module slots
- Installed at the rear of the chassis
- Uses Broadcom T2 as the network forwarding engine (NFE)
- Up to 3.84 Tbps duplex per line card slot
- All fabric cards are active and carry traffic

Chassis Type	Nexus 9504	Nexus 9508	Nexus 9516
NFEs per Fabric Module	1	2	4

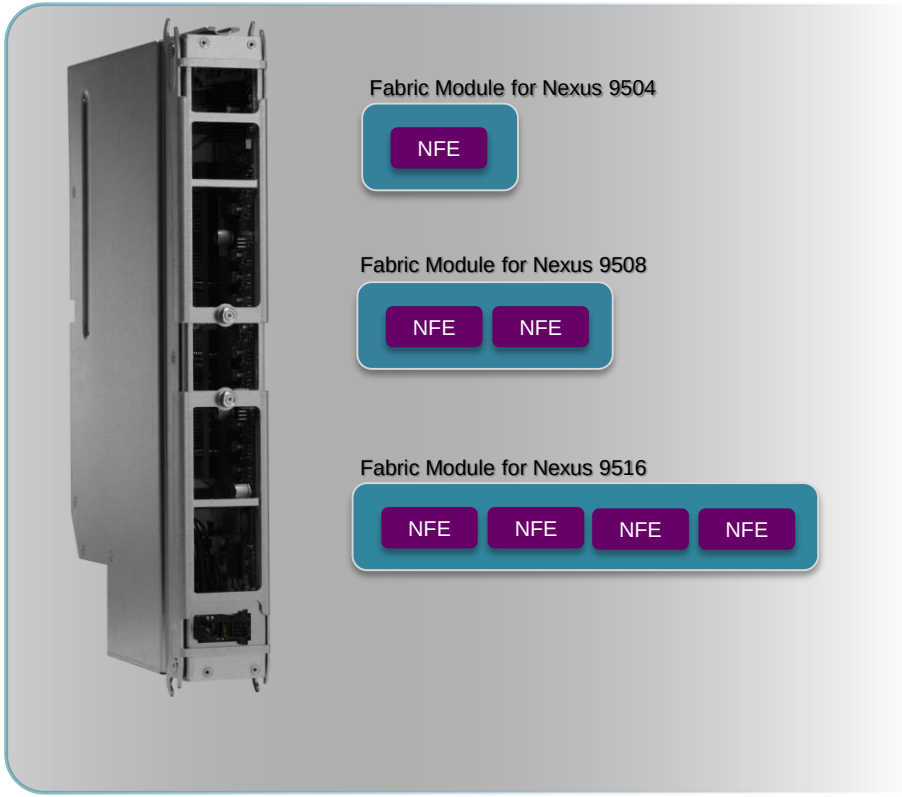

Nexus 9500 Platform Architecture

Fabric Modules

Fabric Module for Nexus 9516 Switches

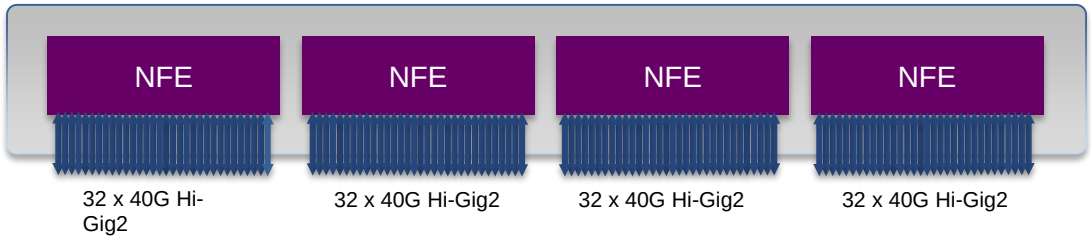

Fabric Module for Nexus 9508 Switches

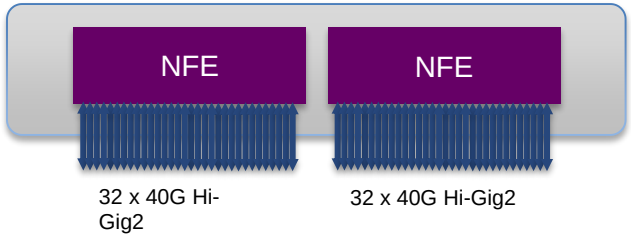

Fabric Module for Nexus 9504 Switches

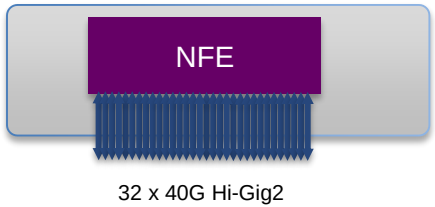

* The 40G Hi-Gig2 links can be clocked at 42G rate

Nexus 9500 Series Switch Architecture

Internal System Architecture: Bipartite Spine-Leaf network

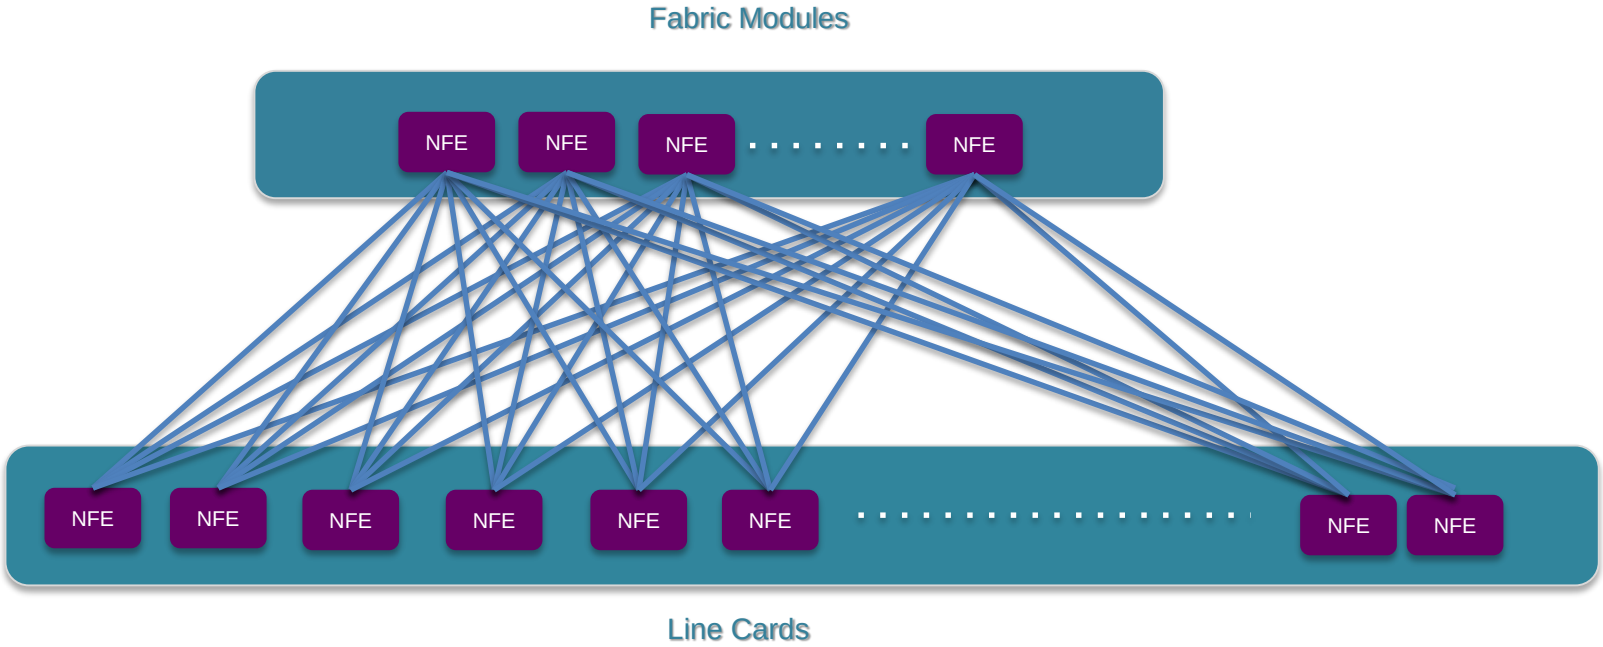

Nexus 9500 Series Switch Fabric Module

Data Plane Scaling (Using Nexus 9508 as an example)

- Each fabric module can provide up to 320 Gbps to each line card slot
- With 6 fabric modules, each lie card slot can have up to 1.92 Tbps forwarding bandwidth in each direction.

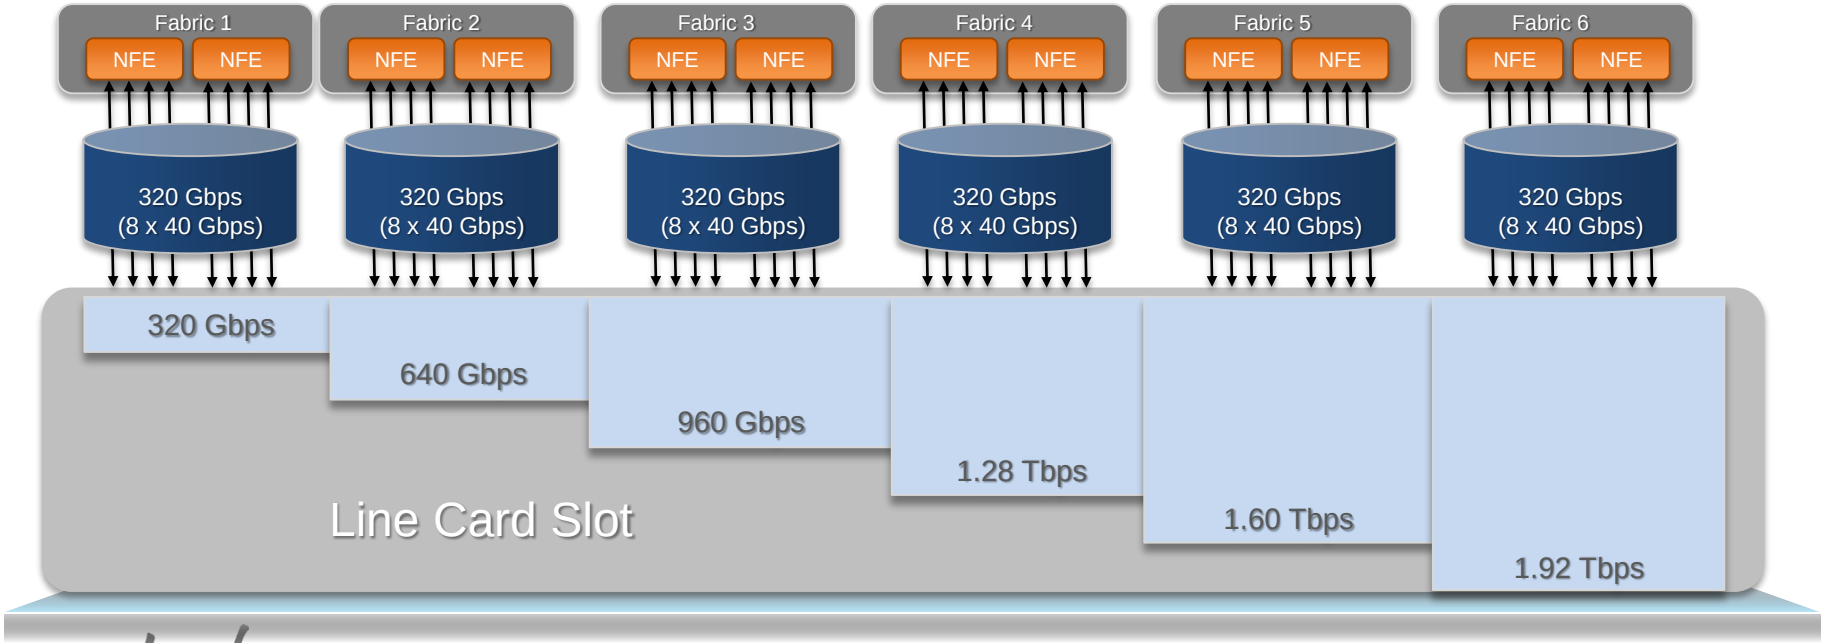

Agenda – Nexus 9000 Architecture

- Nexus 9000 Series Switch Hardware
  - Nexus 9000 Hardware Overview
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9500 Buffer and Queuing
  - Nexus 9300 Hardware Architecture
  - Nexus 9300 Buffer and Queuing
- Nexus 9000 and 40GE
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 Open NX-OS
- ACI & Nexus 9000

NFE Forwarding Capacity

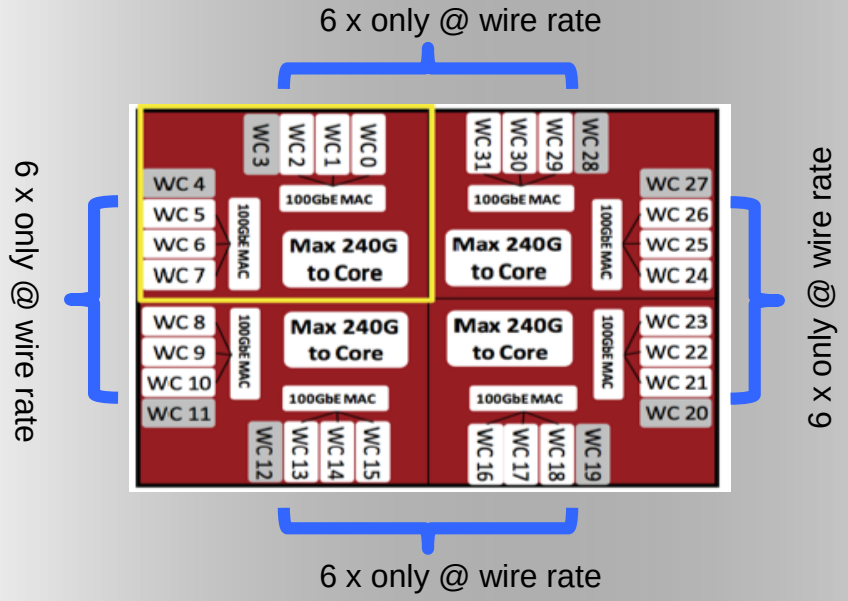

- 24 of 32 ports are full line rate for all packet sizes

Two forwarding Modes on NFE

- Full Late-Rate Mode (FLM)
- Over-subscribed Mode (OSM)

Full Late-Rate Mode (FLM):

- Only 24 40GE ports are used
- Every port is full line-rate for all packet sizes

Over-subscribed Mode (OSM)

- All 32 40GE ports are used
- Every ports is line-rate for packet sizes > 193 Bytes

Nexus 9500 Switch Line Cards

X9600 Series Line Cards (NFE only)
High Performance 40G Aggregation

- N9K-X9636PQ

Merchant only

NX-OS Mode Only

X9500 Series Line Cards (NFE + ALE)
Performance 10G/40G Access/Aggregation

- N9K-X9564PX
- N9K-X9564TX
- N9K-X9536PQ

Merchant +

NX-OS Mode & ACI Leaf Ready

X9400 Series Line Cards (NFE only)
Entry Level 10GE/40GE Access and Aggregation

- N9K-X9432PQ
- N9K-X9464PX
- N94-X9464TX

Merchant only

NX-OS Mode Only

X9700 Series Line Cards (ASE only)
High Performance 40G ACI Spine

- N9K-X9736PQ

Merchant +

ACI Spine Only

Nexus 9500 NX-OS Line Cards

Merchant+

9564 (48x10G+4x40G)

9564 (48x10G+4x40G)

9536 (36x40G)

- **Additional buffer**
- **VXLAN routing***
- **ACI Enabled***

Merchant

9636PQ (36x40G)

9464TX (48x10G+4x40G)

9464PX (48x10G+4x40G)

9432PQ (32x40G)

9408PC-CFP2 (8x 100G)

10G Copper **10G Fiber** **40G Fiber** **100G Fiber**

Nexus 9500 N9K-X9600 Series Line Cards

N9K-X9636PQ

N9K-X9636PQ line card needs 6 fabric modules to operate at full line rate on all 36 ports and for all packet sizes.

- 36x 40 Gbps QSFP ports
- 2.88 Tbps full-duplex fabric connectivity
- Layer 2 and 3 line-rate performance on all ports for all packet sizes
- Supports 4x 10 Gbps break-out mode
- Supported in 4 and 8-slot Nexus 9500 chassis

Nexus 9500 N9K-X9600 Series Line Cards

N9K-X9636PQ

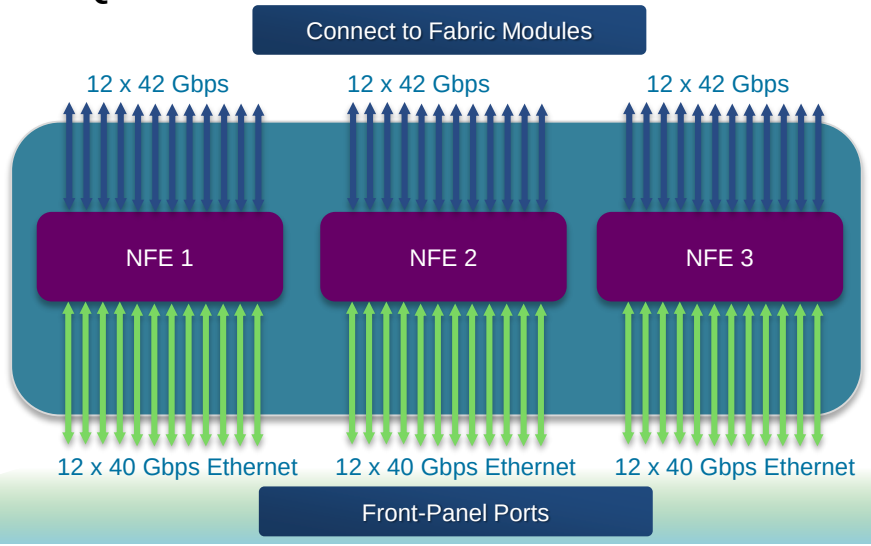

N9K-X9636PQ line card needs 6 fabric modules to operate at line rate on all 36 ports.

- 3 network forwarding engines (NFE)
- Each NFE runs in full-line-rate mode, providing 12 x 40 Gbps links to the front panel and 12 x 40 Gbps internal links to the fabric modules

Nexus 9500 N9K-X9600 Series Line Cards

N9K-X9636PQ Fabric Connectivity

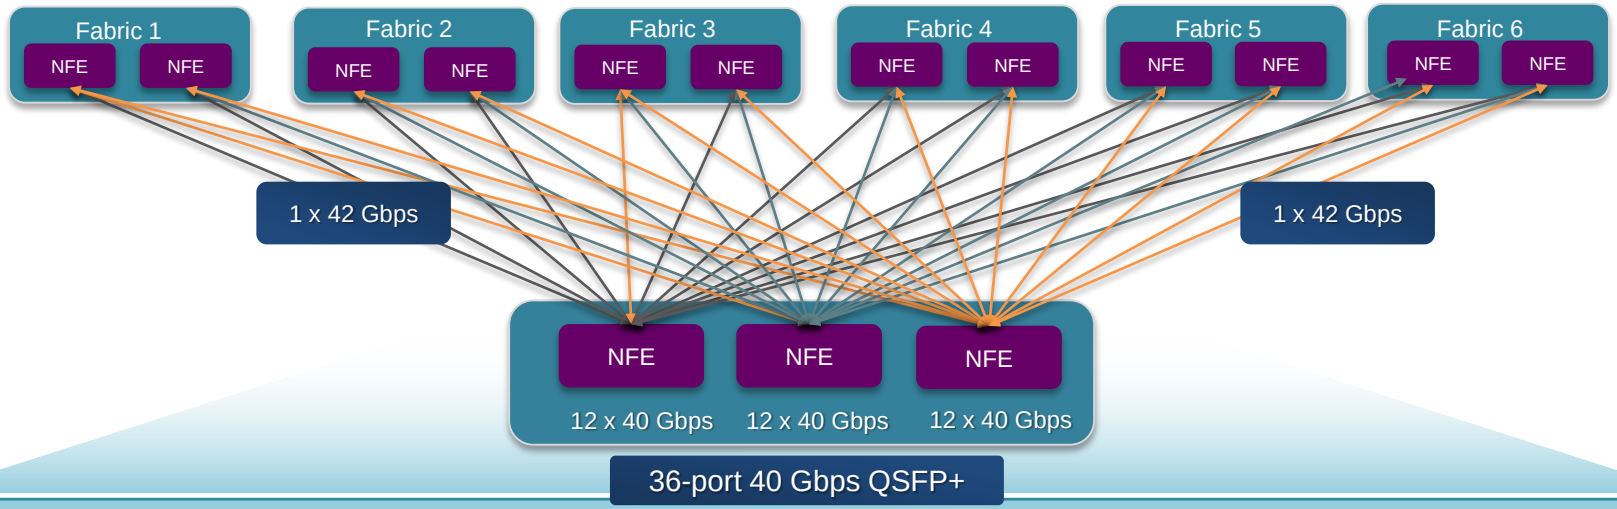

- All ports on the line card can operate at line rate for any packet sizes with 6 fabric modules
- Each NFE has 12 x 40 Gbps internal links to fabric modules - one to each Fabric NFE
- The Internal 40 Gbps links are running at 42 Gbps clock rate to compensate the internal overhead

Nexus 9500 N9K-X9600 Series Line Cards

Full Line Rate on All 40GE Ports

Unprecedented Full Line Rate Performance:

- Proved with RFC 2544/ RFC 2889/ RFC 3918 Throughput Test Results on a fully loaded Nexus 9508 switch with 288 40GE ports:
  - All ports are line rate at 100% unicast traffic load
  - All ports are line rate at 100% multicast traffic load
  - Full line rate for all packet sizes (64~9216 Bytes)

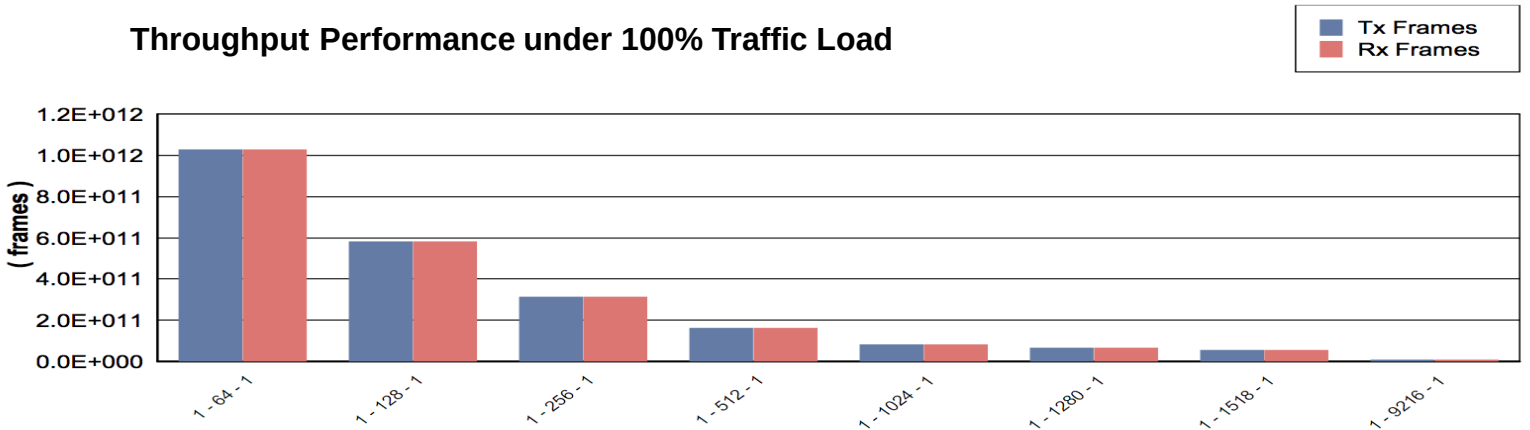

Nexus 9500 N9K-X9600 Series Line Cards

Low Latency Performance on N9K-X9636PQ

Low Latency (same for both Unicast and Multicast):

- Proved with RFC 2544/ RFC 2889/ RFC 3918 Throughput Test Results on a fully loaded Nexus 9508 switch with 288 40GE ports.
- Consistent unicast and multicast Latency at 100% traffic load:
  - 1.6 usec (64-Byte packets)
  - 3.5 usec (9216-Byte packets)

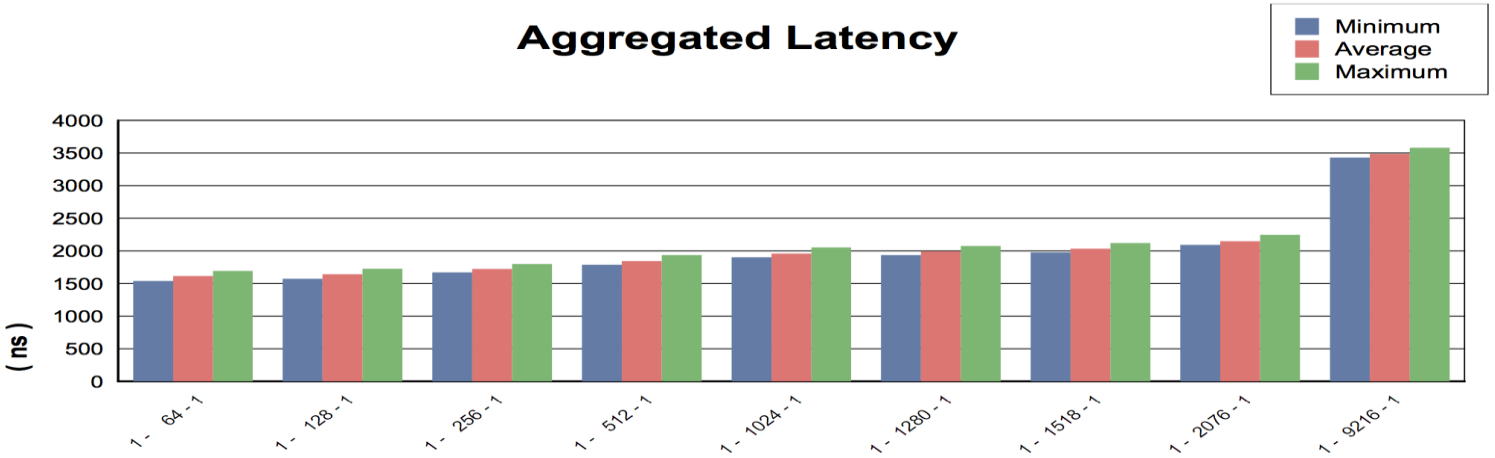

Nexus 9500 N9K-X9500 Series Line Cards

N9K-X9564PX & N9K-X9564TX

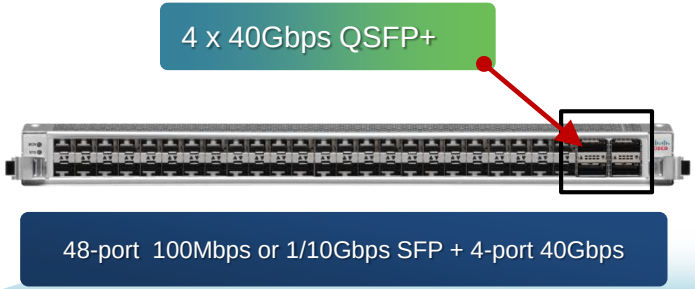

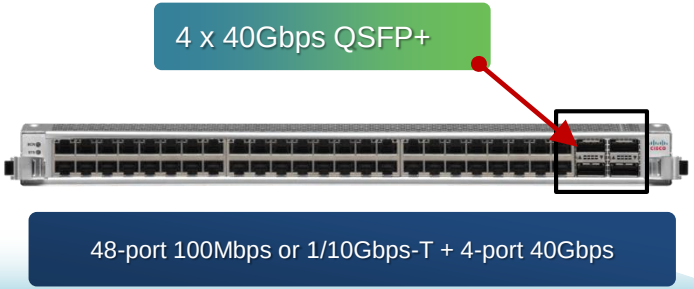

- N9K-X9564PX --- 48 100Mbps or 1/10Gbps SFP+ ports + 4 40Gbps QSFP+ ports
- N9K-X9564TX --- 48 100Mbps or 1/10Gbps Base-T ports + 4 40Gbps QSFP+ ports
- 1.92 Tbps duplex fabric connectivity
- Layer 2 and 3 line-rate performance on all ports for all packet sizes
- Cisco® NX-OS and Application Centric Infrastructure (ACI) mode
- Supported in 4, 8 and 16-slot Nexus 9500 chassis

Nexus 9500 N9K-X9500 Series Line Cards

N9K-X9564PX & N9K-X9564TX (Cont'ed)

2 network forwarding engines (NFEs)

2 application leaf engines (ALEs) for additional buffering and packet handling

Works in 4, 8 and 16 slot chassis
Line rate performance on all ports and all packet sizes with 3 or 6 fabric modules

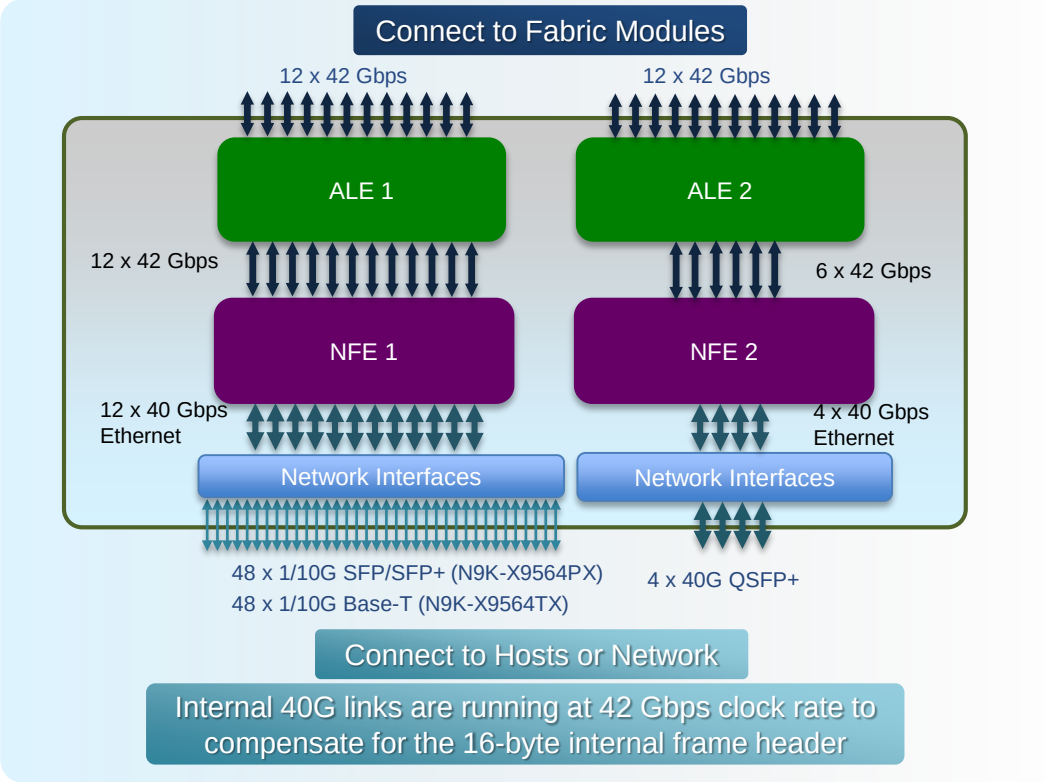

Nexus 9500 N9K-X9500 Series Line Cards

N9K-X9564PX & N9K-X9564TX (Cont'ed)

NFE ASICs act as main forwarding engines for standalone mode.

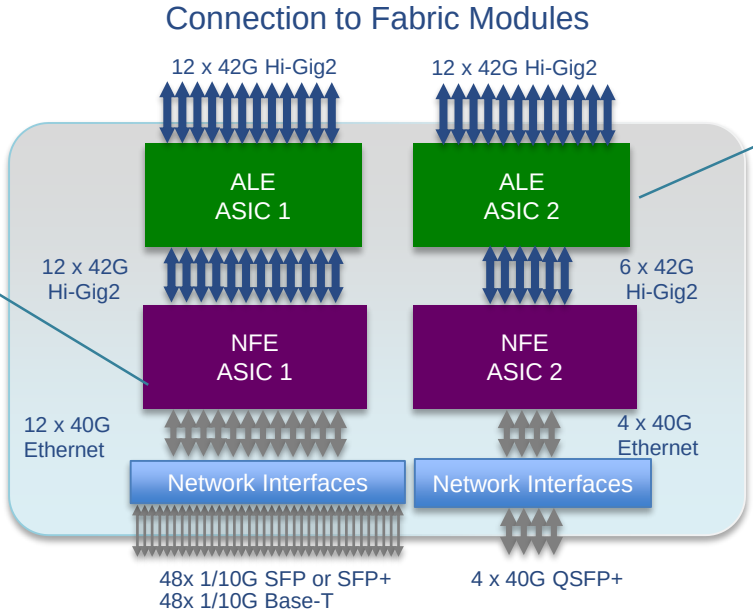

ALE ASICs perform additional packet processing and buffering for standalone mode.

Nexus 9500 N9K-X9500 Series Line Cards

N9K-X9564PX & N9K-X9564TX Fabric Connectivity

Fabric Connectivity with 3 Fabric Modules in Nexus 9508 Chassis

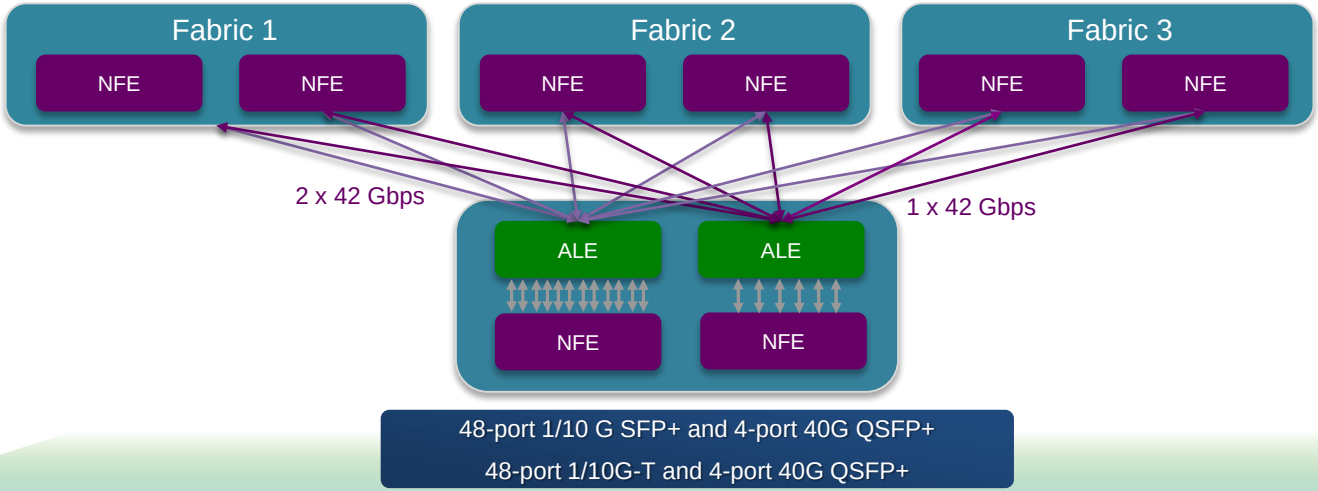

- Minimum of 3 fabric modules to get all line-rate ports
- When a chassis has only 3 fabric modules, they need to be positioned in fabric module slot 2, 4 and 6.

Nexus 9500 N9K-X9500 Series Line Cards

N9K-X9564PX & N9K-X9564TX Fabric Connectivity

Fabric Connectivity with 6 Fabric Modules in Nexus 9508 Chassis

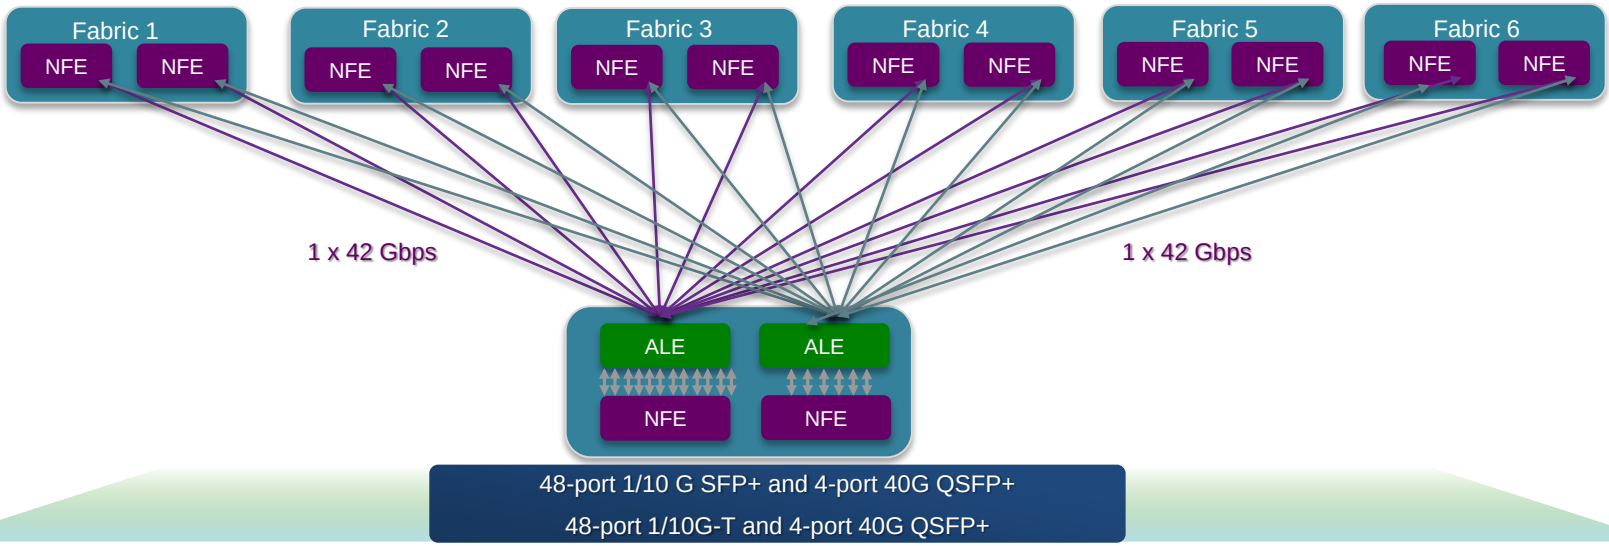

When a chassis has 6 fabric modules, N9K-X9564PX and N9K-X9564TX are connected to and utilize all 6 fabric modules

The above diagram shows the line card fabric connectivity in a Nexus 9508 chassis as an example

Nexus 9500 N9K-X9500 Series Line Cards

N9K-X9536PQ

N9K-X9536PQ is supported in all Nexus 9500 Chassis

- 36x 40 Gbps QSFP ports, 1.5:1 oversubscribed
- 1.92 Tbps duplex fabric connectivity
- Full Layer-2 & Layer-3 functions
- Supports 4x 10 Gbps break-out mode
- Support Cisco® NX-OS mode and ACI ready
- Sported in all Nexus 9500 chassis types, including 4, 8 and 16-slot

Nexus 9500 N9K-X9500 Series Line Cards

N9K-X9536PQ Line Card

- 2 network forwarding engines (NFEs)
- 2 application leaf engines (ALEs) for additional buffering and packet handling
- Need 3 fabric modules, can work with 6

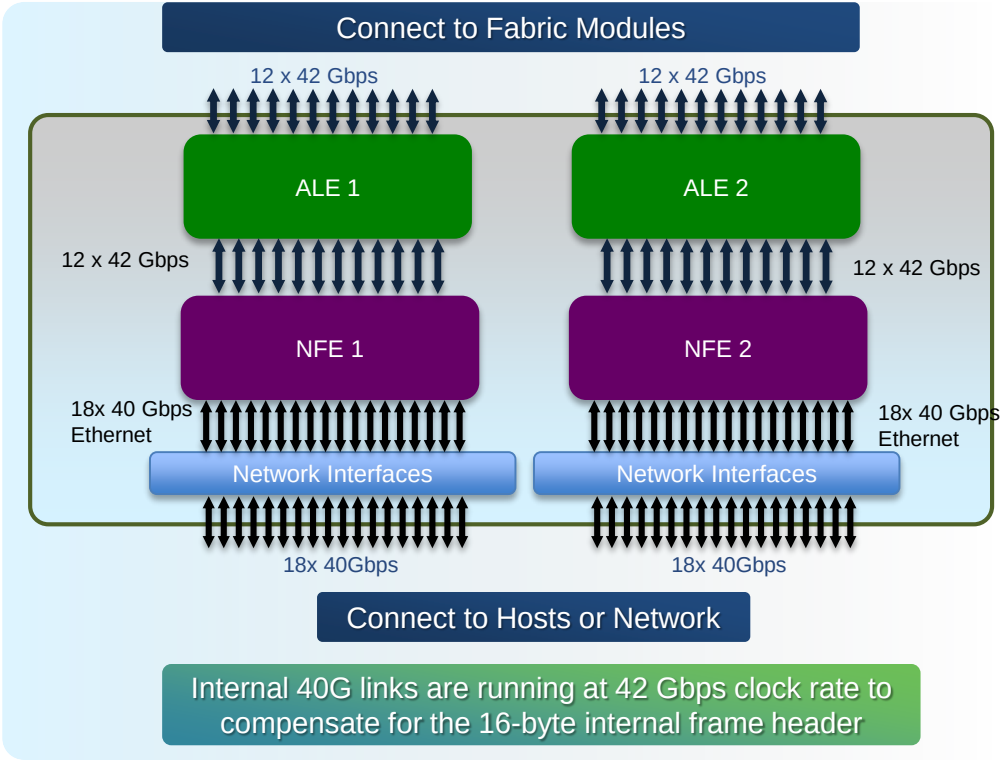

Nexus 9500 N9K-X9400 Series Line

N9K-X9432PQ Line Card

N9K-X9432PQ is supported in all Nexus 9500 Chassis

- 32x 40 Gbps QSFP ports
- 2.56 Tbps full-duplex fabric connectivity
- Full Layer-2 and Layer-3 feature sets
- Hardware supports 4x 10 Gbps break-out mode
- Cisco® NX-OS mode only
- Supported by all Nexus 9500 chassis, including Nexus 9504, 9508 and 9516
- Operates with 4 fabric modules for maximum performance (in fabric module slots 2,3,4 and 6)

Nexus 9500 N9K-X9400 Series Line Cards

N9K-X9432PQ Line Card (Cont'ed)

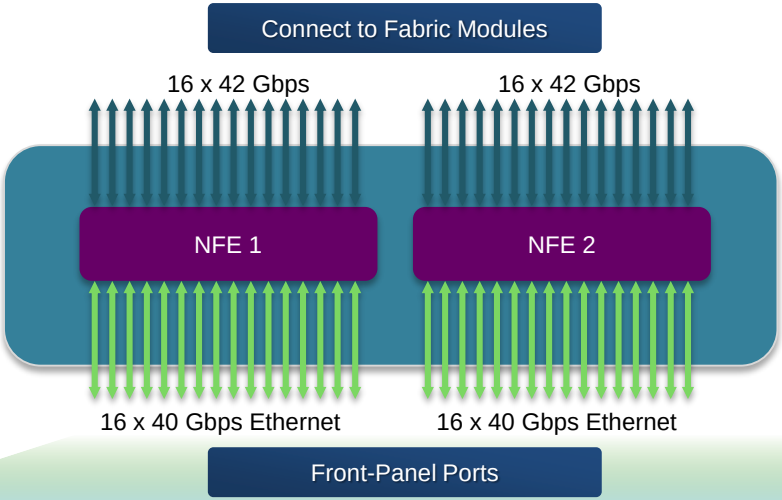

Internal 40G links are running at 42 Gbps clock rate to compensate for the 16-byte internal frame header

N9K-X9432PQ is supported in all Nexus 9500 chassis types.

- Two network forwarding engines (NFE)
- Each NFE supports 16x 40 Gbps front panel ports
- Oversubscribed for small packets (<193 Bytes)
- Line rate performance for larger packet sizes (> 193 Bytes)

Nexus 9500 N9K-X9400 Series Line

N9K-X9464PX and N9K-X9464TX Line Cards

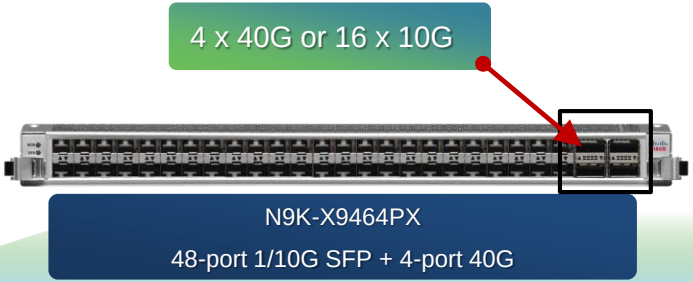

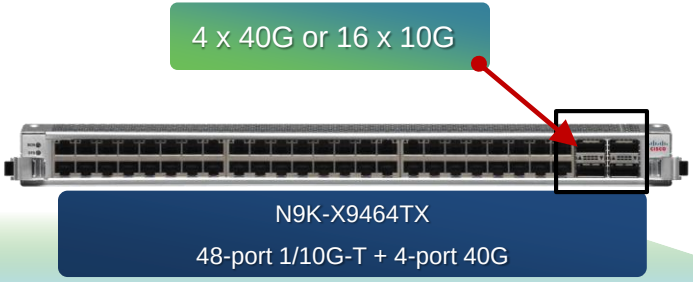

- N9K-X9464PX
  - 48x 1/10 SFP+ Gbps and 4x 40 Gbps QSFP ports
- N9K-X9464PX
  - 48x 1/10 SFP+ Gbps and 4x 40 Gbps QSFP ports

- 1.92 Tbps full-duplex fabric connectivity
- Full Layer-2 and Layer-3 feature sets
- Hardware supports 4x 10 Gbps break-out mode on 40 Gbps ports
- Cisco® NX-OS mode only
- Supported by all Nexus 9500 chassis, including Nexus 9504, 9508 and 9516
- Operate with 4 fabric modules for maximum performance (in fabric module slots 2, 3, 4 and 6)

Nexus 9500 N9K-X9400 Series Line Cards

N9K-X9464PX and N9K-X9464TX Line Cards (Cont'ed)

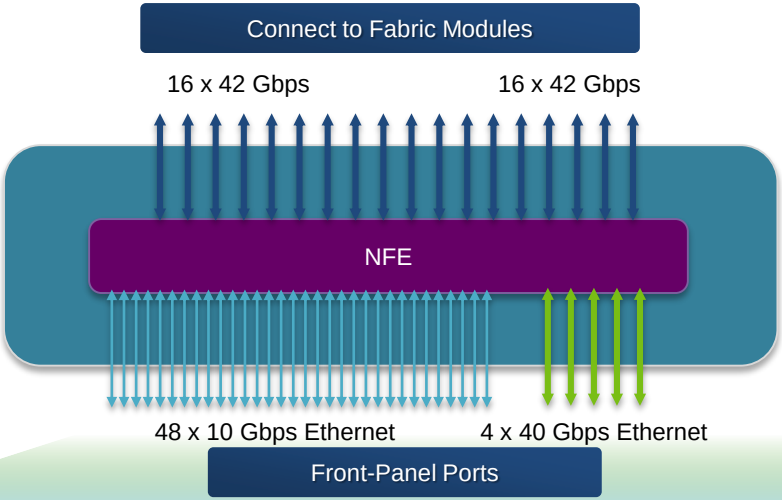

Internal 40G links are running at 42 Gbps clock rate to compensate for the 16-byte internal frame header

N9K-X9464PX/TX line cards are supported in all Nexus 9500 chassis types.

- One NFE supports all 48x 1/10 Gbps and 4x 40 Gbps front panel ports
- Oversubscribed for smaller packet sizes (<193 Bytes)
- Line rate performance for larger packet sizes (> 193 Bytes)

Nexus 9500 N9K-X9400 Series Line Cards

N9K-X9408PC-CFP2 Line Card

New!

N9K-X9408PC-CFP2 is supported in all Nexus 9500 Chassis

- 8-port 100 Gigabit Ethernet CFP2 line card
- Line rate > 193-byte packets
- Supported in 4,8 and 16slot chassis
- Compatible with the current generation of fabric modules (N9K-C9504-FM, N9K-C9508-FM, N9K-C9516-FM)
- Requires 4 fabric modules for maximum bandwidth (in fabric module slots 2, 3, 4 & 6)
- Supports NX-OS mode only
- Currently do not support break-out mode

Nexus 9500 N9K-X9400 Series Line Cards

N9K-X9408PC-CFP2 Line Card (Cont'ed)

New!

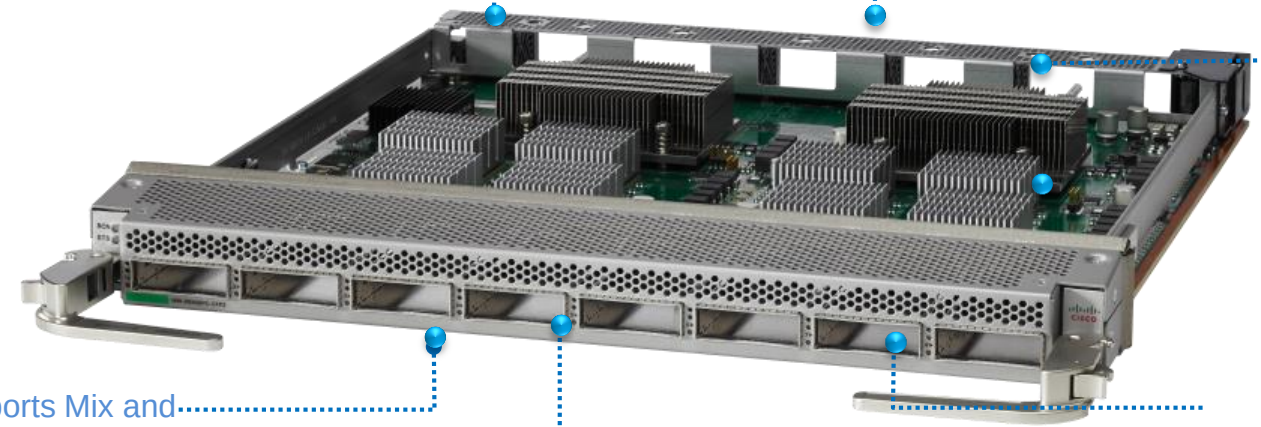

Nexus 9500 N9K-X9400 Series Line Cards

N9K-X9408PC-CFP2 (Cont'ed)

- **Supported by all Nexus 9500 chassis types, including Nexus 9504/9508/9516**

- **8 x 100G CFP2 ports**

Show module information:

```
N9500-1# show module
Mod  Ports  Module-Type                    Model              Status
---  -----  ----------------------------- -----------------  ------
1    8      8p 100G Ethernet Module       N9K-X9408PC-CFP2   ok
```

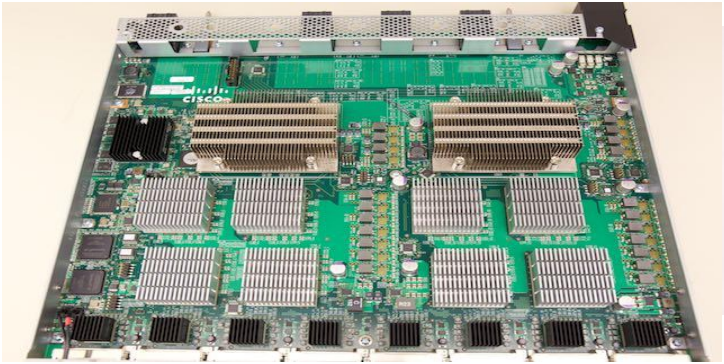

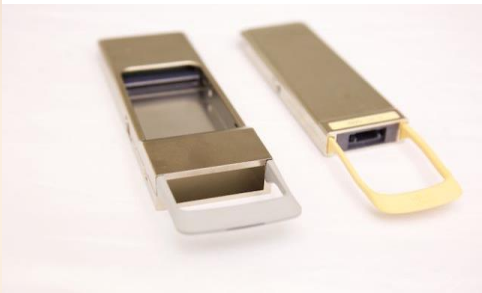

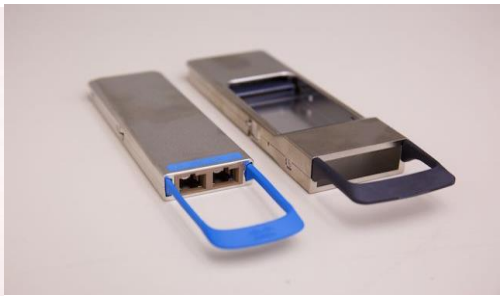

Nexus 9500 N9K-X9400 Series Line Cards

N9K-X9408PC-CFP2 Line Card (Cont'ed)

New!

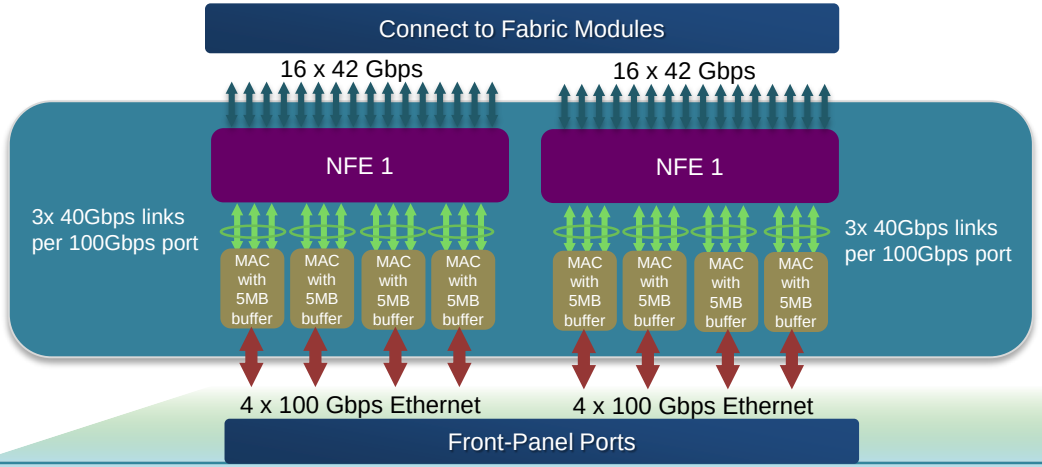

Internal links to fabric modules are running at 42 Gbps clock rate to compensate for the 16-byte internal frame header

N9K-X9408PC-CFP2 is supported in all Nexus 9500 chassis types.

- Two network forwarding engines (NFE)
- Each NFE supports 4x 100 Gbps front panel ports
- Oversubscribed for small packets (<193 Bytes)
- Line rate performance for larger packet sizes (> 193 Bytes)
- Each 100GE front panel port is essentially 3x 40GE ports on NFE
- Supports up to 40GE flows
- The 100GE MAC ASIC per front panel port has additional 5MB buffer

Agenda – Nexus 9000 Architecture

- Nexus 9000 Series Switch Hardware
  - Nexus 9000 Hardware Overview
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9500 Buffer and Queuing
  - Nexus 9300 Hardware Architecture
  - Nexus 9300 Buffer and Queuing
- Nexus 9000 and 40GE
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 Open NX-OS
- ACI & Nexus 9000

NFE Unified Forwarding Table

- NFE has a 16K traditional LPM TCAM table.
- Additionally NFE has the following Unified Forwarding Table for ALPM (Algorithm LPM) Mode

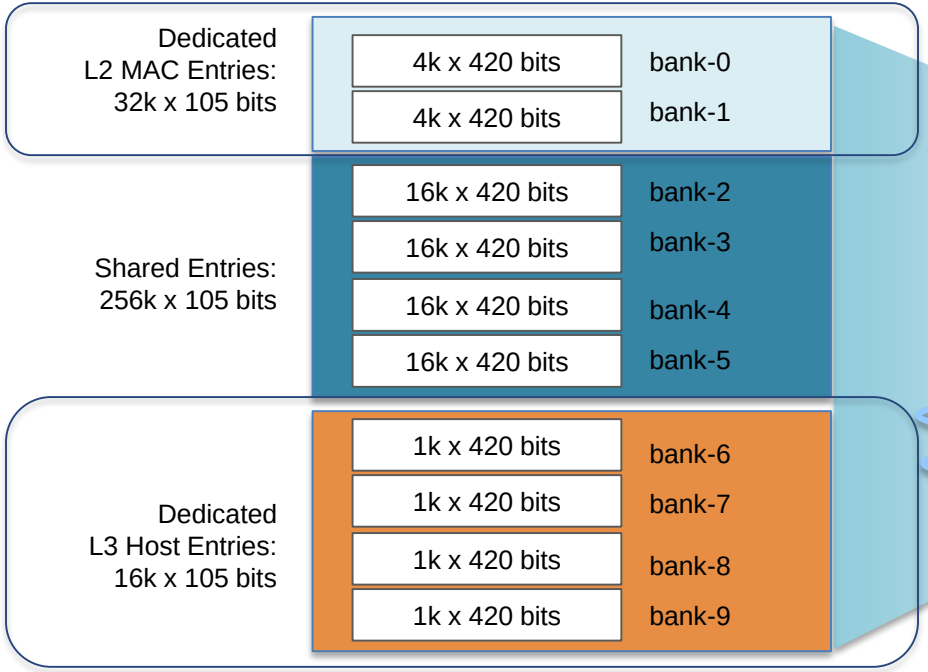

SUPPORTED COMBINATIONS

Mode	L2	L3 Hosts	LPM
0	288K	16K	0
1	224K	56K	0
2	160K	88K	0
3	96K	120K	0
4	32K	16K	128K

In default setting, N9500 line card NFE uses Mode 3, fabric module NFE uses Mode 4.

Nexus 9500 Forwarding Programming Mode

Hierarchical Routing Mode (Default)

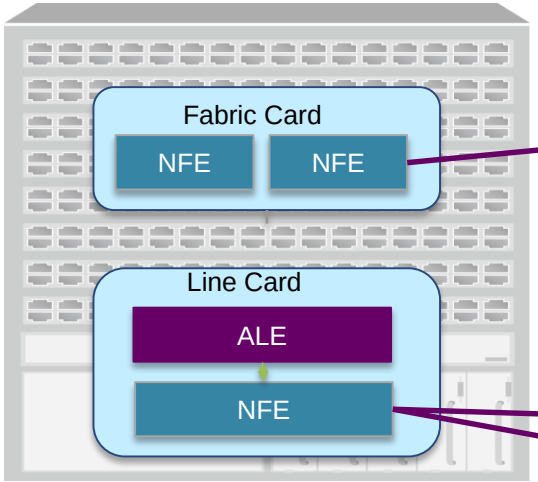

Feature	Scale	NFE Mode
IPv4/v6 LPM Routes	128K	4

Feature	Scale	NFE Mode
IPv4/v6 Host Routes*	120K*	3
MAC addresses	96K	

* Shares the same table with multicast routes

Nexus 9500 Forwarding Programming Mode

Nexus 9500 switches consist of two tiers of NFEs internally, the line card tier and the fabric module tier. on its line cards and fabric modules. Forwarding tables can be programmed on these two tiers differently based on the programming mode of the Nexus 9500 switch.

Forwarding tables:

- MAC Table
- IPv4/IPv6 Host Table
- IPv4/IPv6 LPM Table
- IPv4 Multicast table

Forwarding Programming Modes:

- Hierarchical mode ---- default system routing mode
- Hierarchical Max-host routing mode
- Non-Hierarchical routing mode
- Non-Hierarchical routing max-l3 mode
- 64-bit ALPM routing mode

Nexus 9500 Forwarding Programming Mode

	MAC Table		IPv4/IPv6 Host Table		IPv4/IPv6 LPM Route Table		Multicast Route Table	
	Location	NFE Mode	Location	NFE Mode	Location	NFE Mode	Location	NFE Mode
Hierarchical routing mode (default)	LC	3	LC	3	FM	4	LC+FM	3
Hierarchical 64-bit ALPM mode	LC	3	LC	3	FM	4	LC+FM	3
Hierarchical Max-host routing mode	LC	2	IPv4 on FM	3	IPv4 on FM	3	LC+FM	
			IPv6 on LC	2	IPv6 on LC	2		
Non-hierarchical routing mode	LC	3	LC	3	LC	3	LC	3
Non-hierarchical routing Max-L3 mode	LC	4	LC	4	LC	4	LC	4

Forwarding Programming Mode	Configuration Command
Default Hierarchical routing mode	Default
Hierarchical 64-bit ALPM mode	9508(config)# system routing hierarchical max-mode l3 64b-alpm
Hierarchical Max-host routing mode	9508(config)# system routing max-mode host
Non-hierarchical routing mode	9508(config)# system routing non-hierarchical
Non-hierarchical routing Max-L3 mode	9508(config)# system routing non-hierarchical max-mode l3

CLI to Show Forwarding Programming Mode

```
9508# sh system routing mode
Configured System Routing Mode: Non-Hierarchical (Default)
Applied System Routing Mode: Hierarchical (Default)
Configured SVI post-routed unknown-unicast hardware flood mode: enabled
US-DUR-LC01-9508#
```

```
9508# show forwarding route summary module 1

Module Type                    : Line-Card
Module Mode                    : Mode-3
Module Route Download-type     : Host only
(IPv4+IPv6) (1)

IPv4 routes for table default/base

'**' denotes routes NOT programmed in hardware
due to hierarchical routing

Cumulative route updates: 1005038
Cumulative route inserts: 1005005
Cumulative route deletes: 143
Total number of routes: 24
Total number of paths : 25

Number of routes per mask-length:
  /32 : 24
```

```
9508# show forwarding route summary module 26

Module Type                    : Fabric-Module
Module Mode                    : ALPM (Mode-4)
Module Route Download-type     : LPM only
(IPv4+IPv6) (2)

IPv4 routes for table default/base

'**' denotes routes NOT programmed in hardware due
to hierarchical routing

Cumulative route updates: 1005043
Cumulative route inserts: 1004930
Cumulative route deletes: 54
Total number of routes: 8
Total number of paths : 8

Number of routes per mask-length:
  /8  : 1          /30 : 5

US-DUR-LC01-9508#
```

CLI to Check Forwarding Table Sizes

	Software	Hardware (BCM-shell)
MAC Table	show mac address-table count	l2 show
IP Host Table	show ip route sum sh forwarding route summary mod <#>	l3 l3table show [on LC]
IP LPM Table	show ip route sum show forwarding route sum mod <#>	l3 defip show [on FM]
egress next-hop table		l3 egress show [on both LC and FM]

Commands to check hardware table size:
Leverage NX-OS “| count” to get the account of the hardware entries. Example”

```
TME-1-9508-1# bcm-shell mod 1 "l3 l3table show" | count
```

Nexus 9500 Forwarding Program Mode

Hierarchical Routing Mode (Default) (Cont'ed)

FYI

Default Forwarding Table Scalability in NX-OS 6.1.2(I2)1 and Later

FM: BCM T2 mode 4 LC: BCM T2 mode 3	Nexus® 9500
	Default
LPM Routes	128K (on Fabric Module)
IP Host Routes	120K (on Line Card) (208K for protocol learned host routes)
IP ARP Entries	45K (in default mode)
MAC Address Entries	90K (On Line Card)
Multicast Routes	32K* (Hardware capable of 72K)

* Shared with IP hosts

This table shows system maximum scalability. For validated scale, refer to the Nexus 9000 verified scalability white papers

Nexus 9500 Forwarding Program Mode

Hierarchical Routing Mode (Default) (Cont'ed)

FYI

To increase IPv6 LPM table scale, can enable 64-Bit ALPM routing mode

FM: BCM T2 mode 4 LC: BCM T2 mode 3	Nexus® 9500
	Default
LPM Routes	128K IPv4(on Fabric Module) 80K IPv6 (< /64)
IP Host Routes	120K (on Line Card) (208K for protocol learned host routes)
IP ARP Entries	45K (in default mode)
MAC Address Entries	90K (On Line Card)
Multicast Routes	32K* (Hardware capable of 72K)

* Shared with IP hosts

This table shows system maximum scalability. For validated scale, refer to the Nexus 9000 verified scalability white papers

```
9508(config)# system routing hierarchical max-mode l3 64b-alpm
```

Nexus 9500 Forwarding Program Mode

FYI

Hierarchical Routing Max-host Mode

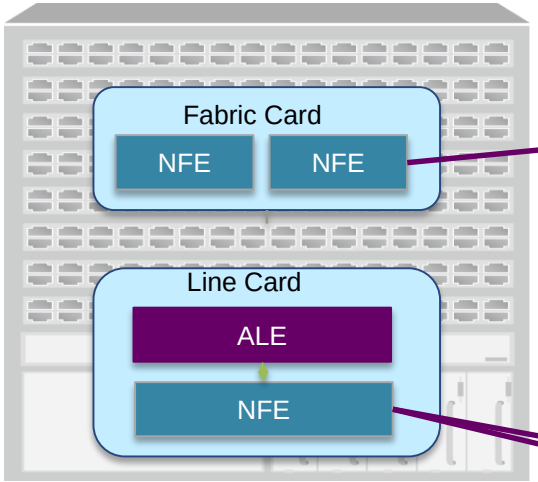

Feature	Scale	NFE Mode
IPv4 Host Routes	120K	3
IPv4 LPM Routes	16K	

Feature	Scale	NFE Mode
IPv6 Host Routes*	40K	2
IPv6 LPM Routes	8K (</64) 4K (> /64)	
MAC addresses	90K	

* Shared the same table with multicast routes

Nexus 9500 Forwarding Program Mode

FYI

Hierarchical Routing Max-host Mode (Cont'ed)

Max-Host mode is to support maximum IP host scalability, especially in IPv4/ IPv6 dual-stack mode.

```
9508(config)# system routing max-mode host
```

Forwarding Tables in Max Host Mode:

- IPv4 Host Table and IPv4 LPM Table are programmed on fabric modules (BCM T2 mode 3)
- IPv6 Host Table, IPv6 LPM Table and MAC Table are programmed on line cards (BCM T2 mode 2)
- IP multicast routes share the same table as IP Hosts. It's programmed on both fabric modules and line cards.

	Nexus® 9500	
	Max-Host Mode	
	IPv4	IPv6
LPM Routes	16K	8K (< /64) 4K (> /64)
IP Host Entries	120K (208K for protocol learned host routes) (60K IPv4 ARP Entries)	40K (can scale up in future release)
MAC Address Entries	90K	
Multicast Routes	32K* (programmed on both FM and LC)	

* Shared with IP hosts

This table shows system maximum scalability. For validated scale, refer to the Nexus 9000 verified scalability white papers

Nexus 9500 Forwarding Program Mode

FYI

Non-Hierarchical Routing (Max-Host) Mode

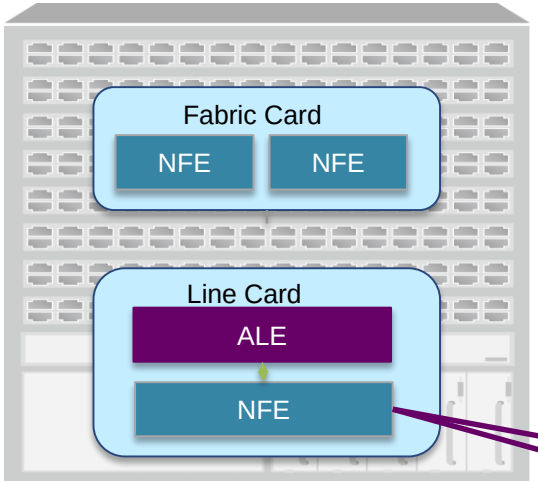

Feature	Scale	NFE Mode
IPv4/v6 Host Routes*	120K	3
IPv4/v6 LPM Routes	16K	
MAC addresses	96K	

* Shares the same table with multicast routes

Nexus 9500 Forwarding Program Mode

FYI

Non-Hierarchical Routing (Max-host) Mode (Cont'ed)

Non-hierarchical mode is created to move LPM table from fabric modules to line cards so that line card NFE can locally forward LPM routed traffic, avoiding the latency of sending traffic to fabric module for LPM lookup.

Forwarding Tables in Non-Hierarchical Mode

- No unicast forwarding tables are programmed on fabric modules
- Line cards are programmed symmetrically with IPv4/v6 LPM + Host + Mac tables
- MAX Host -- Line cards are programmed for maximum IP Host Scale (BCM T2 mode 3)

```
9508(config)# system routing non-hierarchical
```

	Nexus® 9500	
	Non-Hierarchical Mode	
	Max LPM Mode	Max IP Host Mode
LPM Routes	128K	16K
IP Host Entries	16K	120K (208K for protocol learn host routes) (60K directly attached hosts)
MAC Address Entries	32K	90K
Multicast Routes	8K*	32K*

* Shared with IP hosts
This table shows system maximum scalability. For validated scale, refer to the Nexus 9000 verified scalability white papers

FYI

Nexus 9500 Forwarding Program Mode

Non-Hierarchical Routing Max-L3 Mode

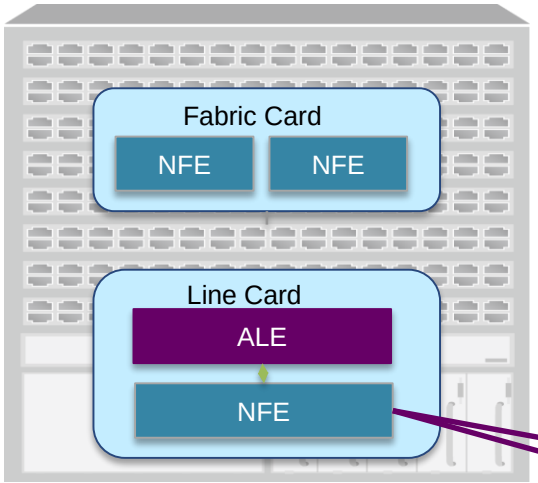

Feature	Scale	NFE Mode
IP Host Routes*	16K	4
IP LPM Routes	128K	
MAC addresses	32K	

* Shares the same table with multicast routes

Nexus 9500 Forwarding Program Mode

FYI

Non-Hierarchical Routing Mode Max-L3 Mode (Cont'ed)

Non-hierarchical mode is created to move LPM table from fabric modules to line cards so that line card NFE can locally forward LPM routed traffic, avoiding the latency of sending traffic to fabric module for LPM lookup.

Forwarding Tables in Non-Hierarchical Mode

- No unicast forwarding tables are programmed on fabric modules
- Line cards are programmed symmetrically with IPv4/v6 LPM + Host + Mac tables
- MAX LPM -- Line cards are programmed for maximum LPM scale (BCM T2 mode 4)

```
9508(config)# system routing non-hierarchical max-mode l3
```

	Nexus® 9500	
	Non-Hierarchical Mode	
	Max LPM Mode	Max IP Host Mode
LPM Routes	128K	16K
IP Host Entries	16K	120K (208K for protocol learn host routes) (60K directly attached hosts)
MAC Address Entries	32K	90K
Multicast Routes	8K*	32K*

* Shared with IP hosts

This table shows system maximum scalability. For validated scale, refer to the Nexus 9000 verified scalability white papers

FYI

Nexus 9300 Switch Forwarding Tables

	Nexus 9300	
	Default Mode	**Max-mode L3 7.0(3)I1(1) or Later**
LPM Routes	16K	128K
IP Host Entries	120K entries (208K for protocol learned hosts or 45K directly attached hosts)	16K
IP ARP Entries	45K	16K
MAC Address Entries	96K	32K
Multicast Routes	16K* (Hardware capable of 72K)	8K*

* Shared with IP hosts

This table shows system maximum scalability. For validated scale, refer to the Nexus 9000 verified scalability white papers

```
n9300(config)# system routing max-mode l3
```

Nexus 9500 Layer-2 Unicast Packet Walk

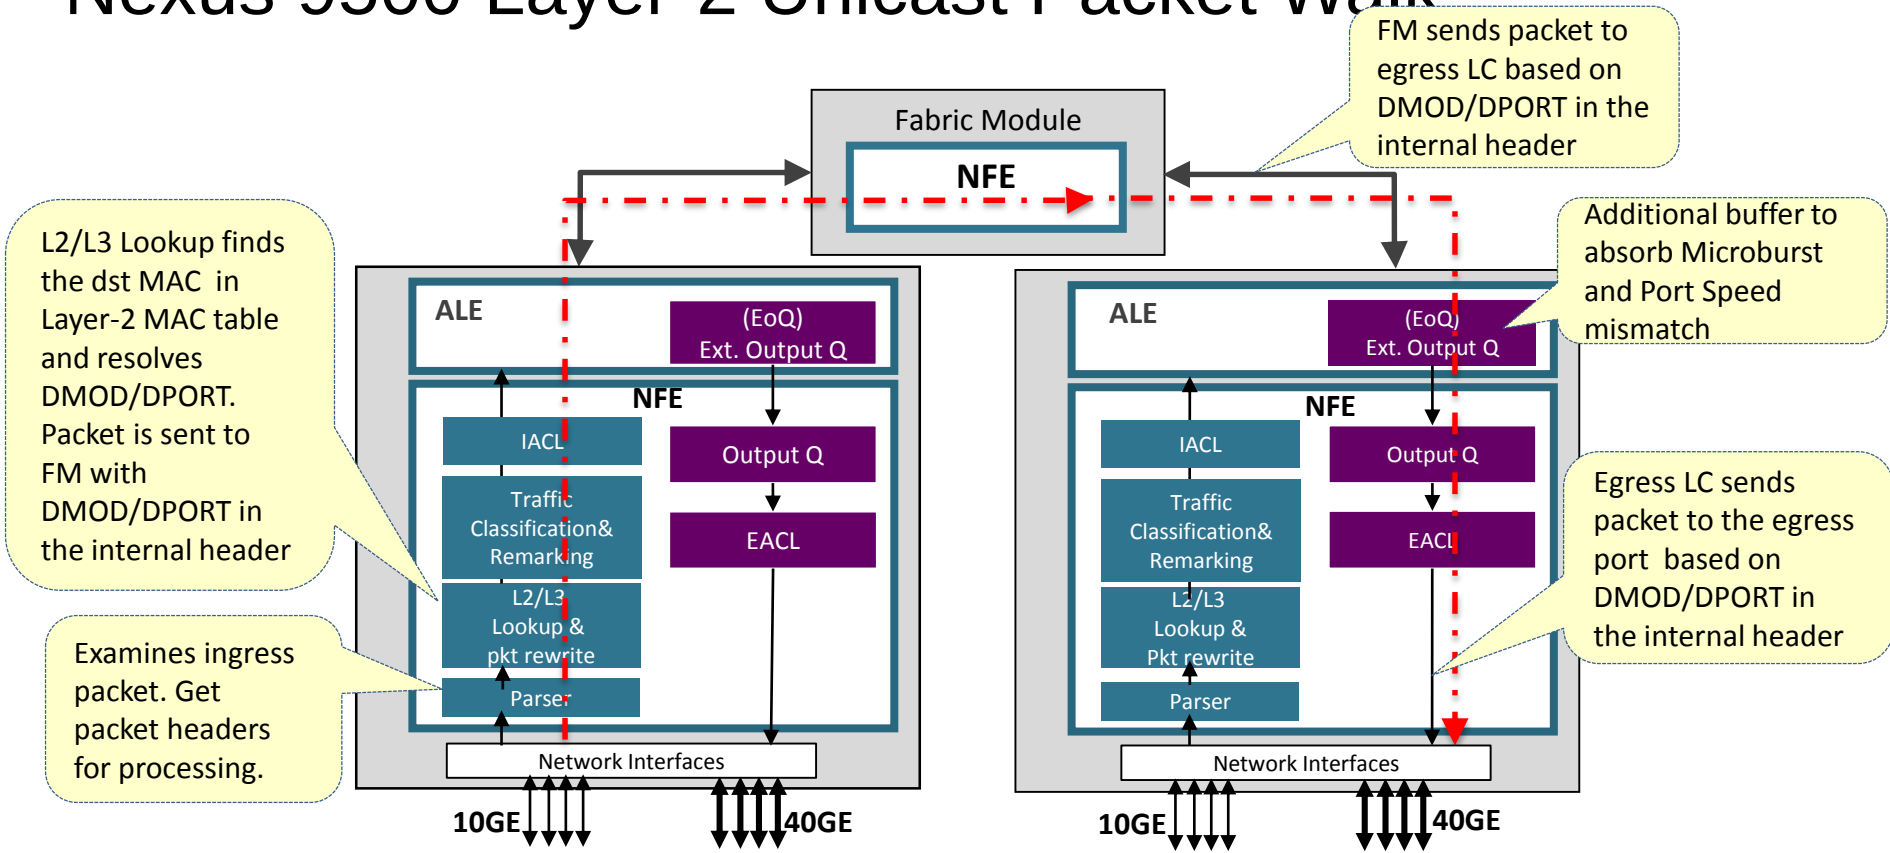

For Line Cards w/n ALE, EoQ provided by ALE does not apply.

Nexus 9500 Layer-3 Host Unicast Packet Walk

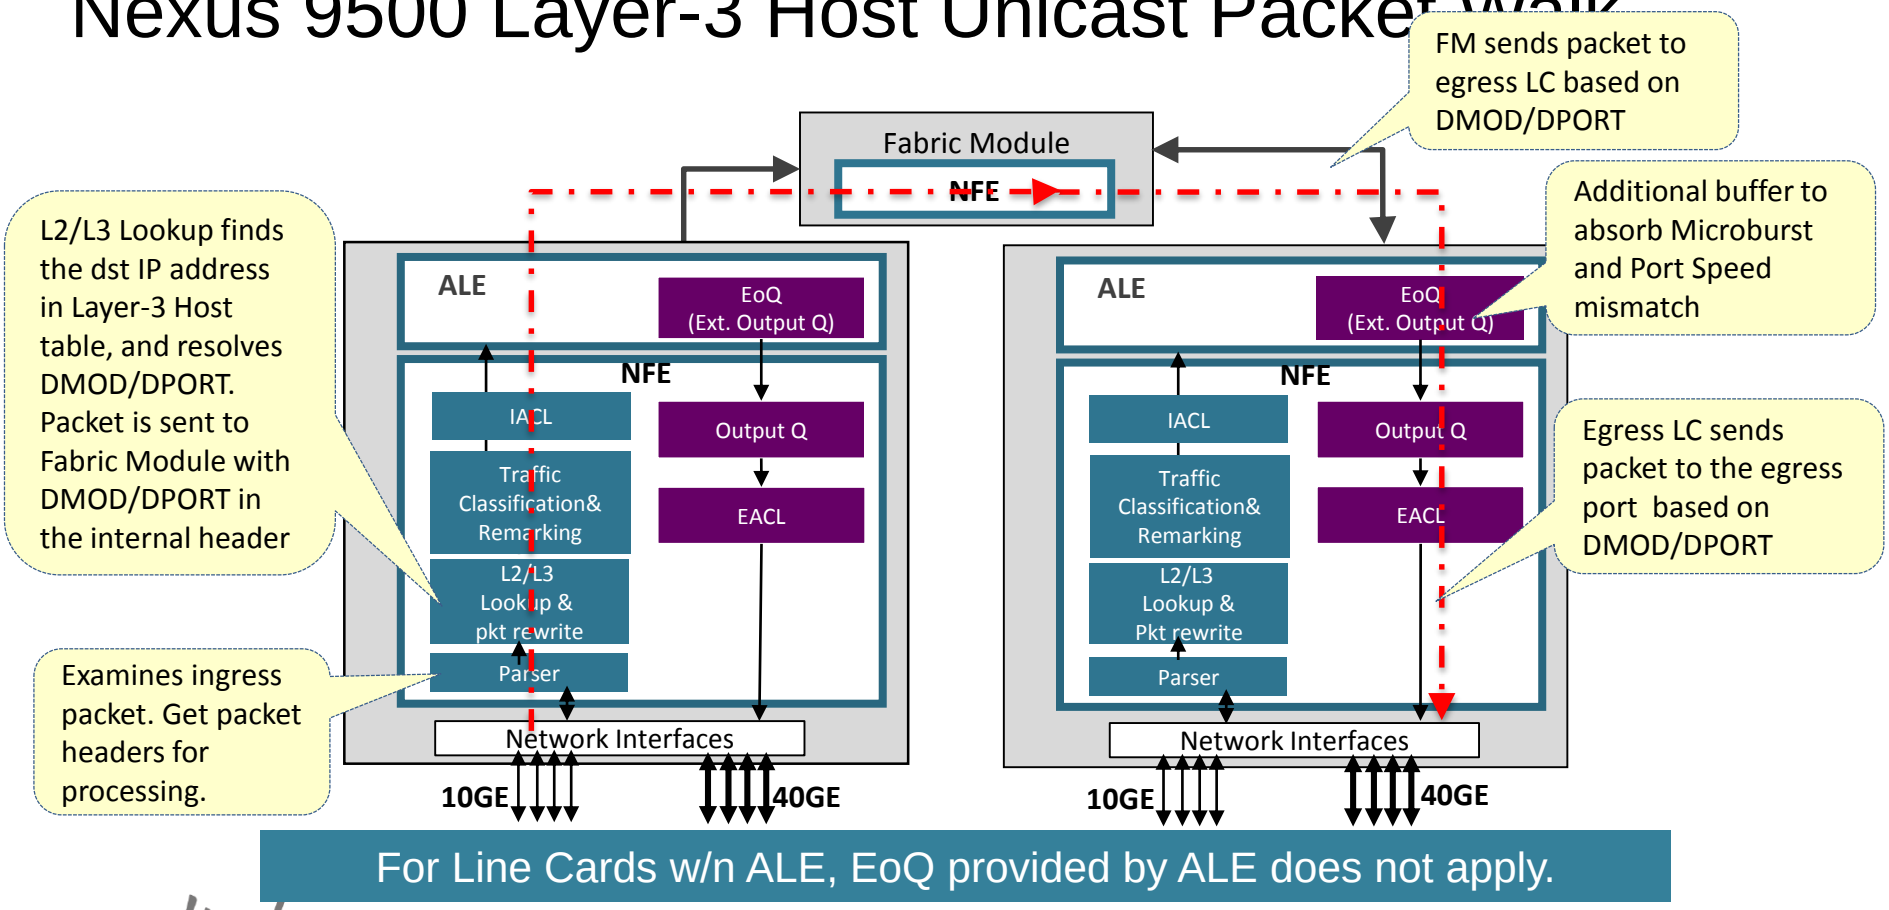

Nexus 9500 Layer-3 LPM Unicast Packet Walk

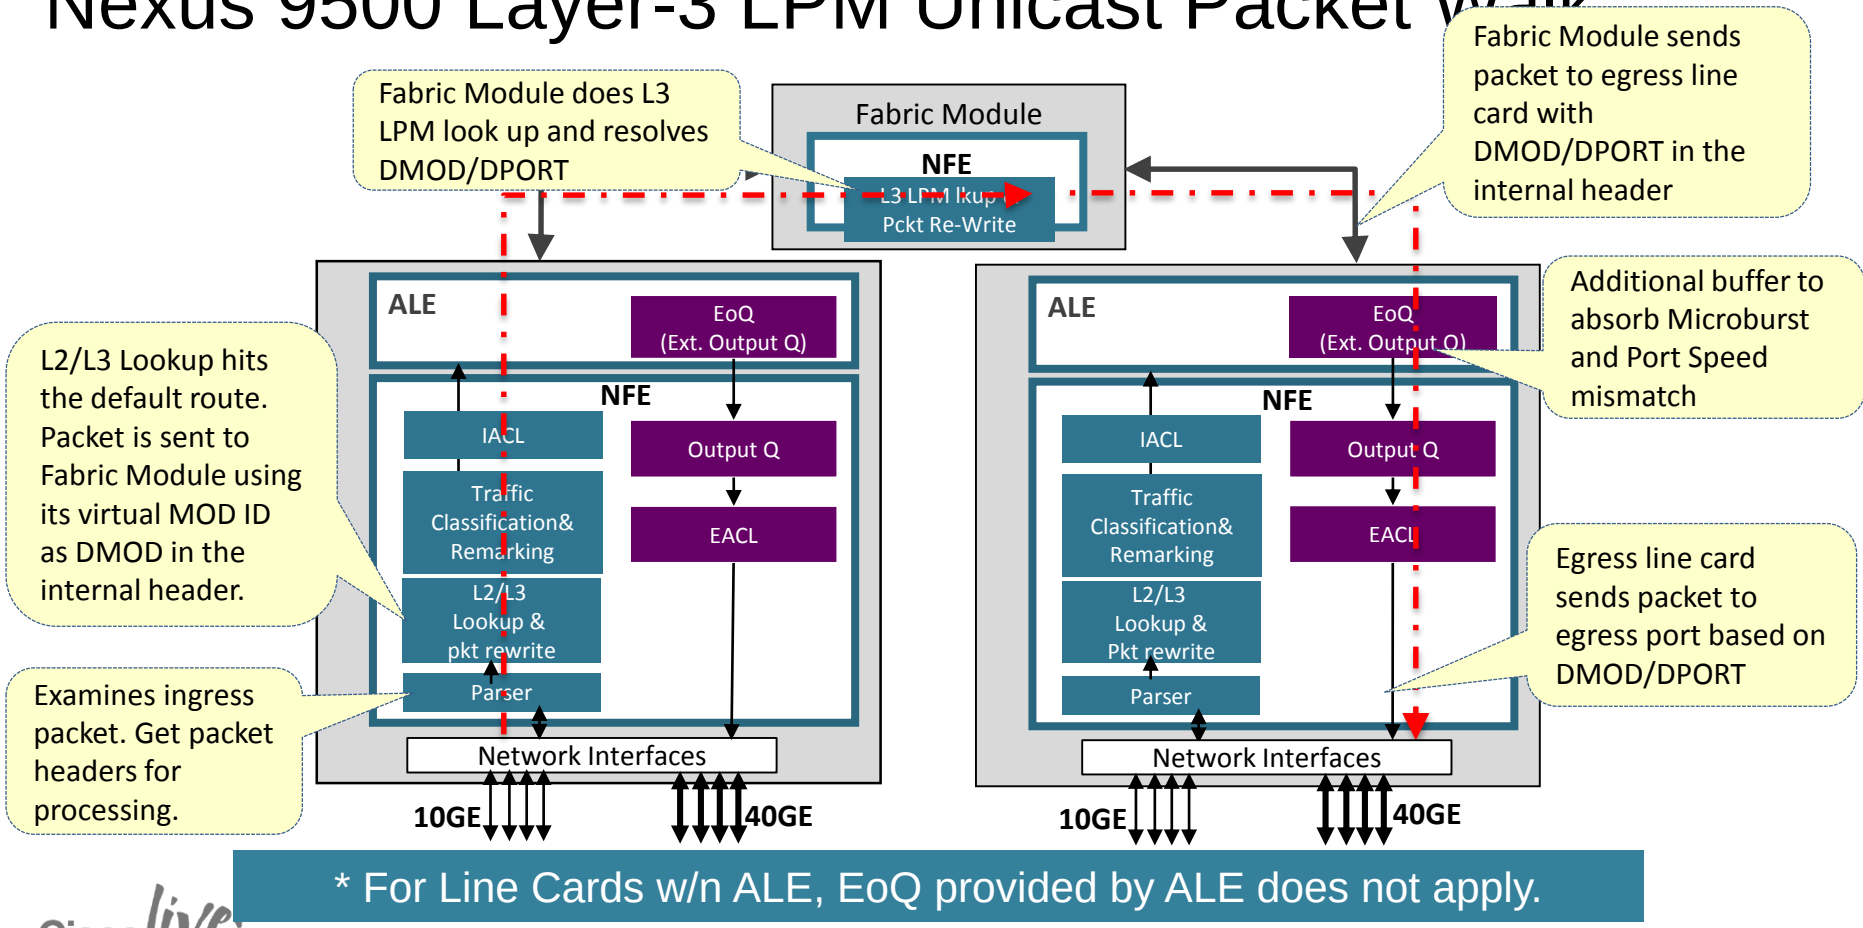

* For Line Cards w/n ALE, EoQ provided by ALE does not apply.

Nexus 9500 L2/L3 Multicast Packet Walk

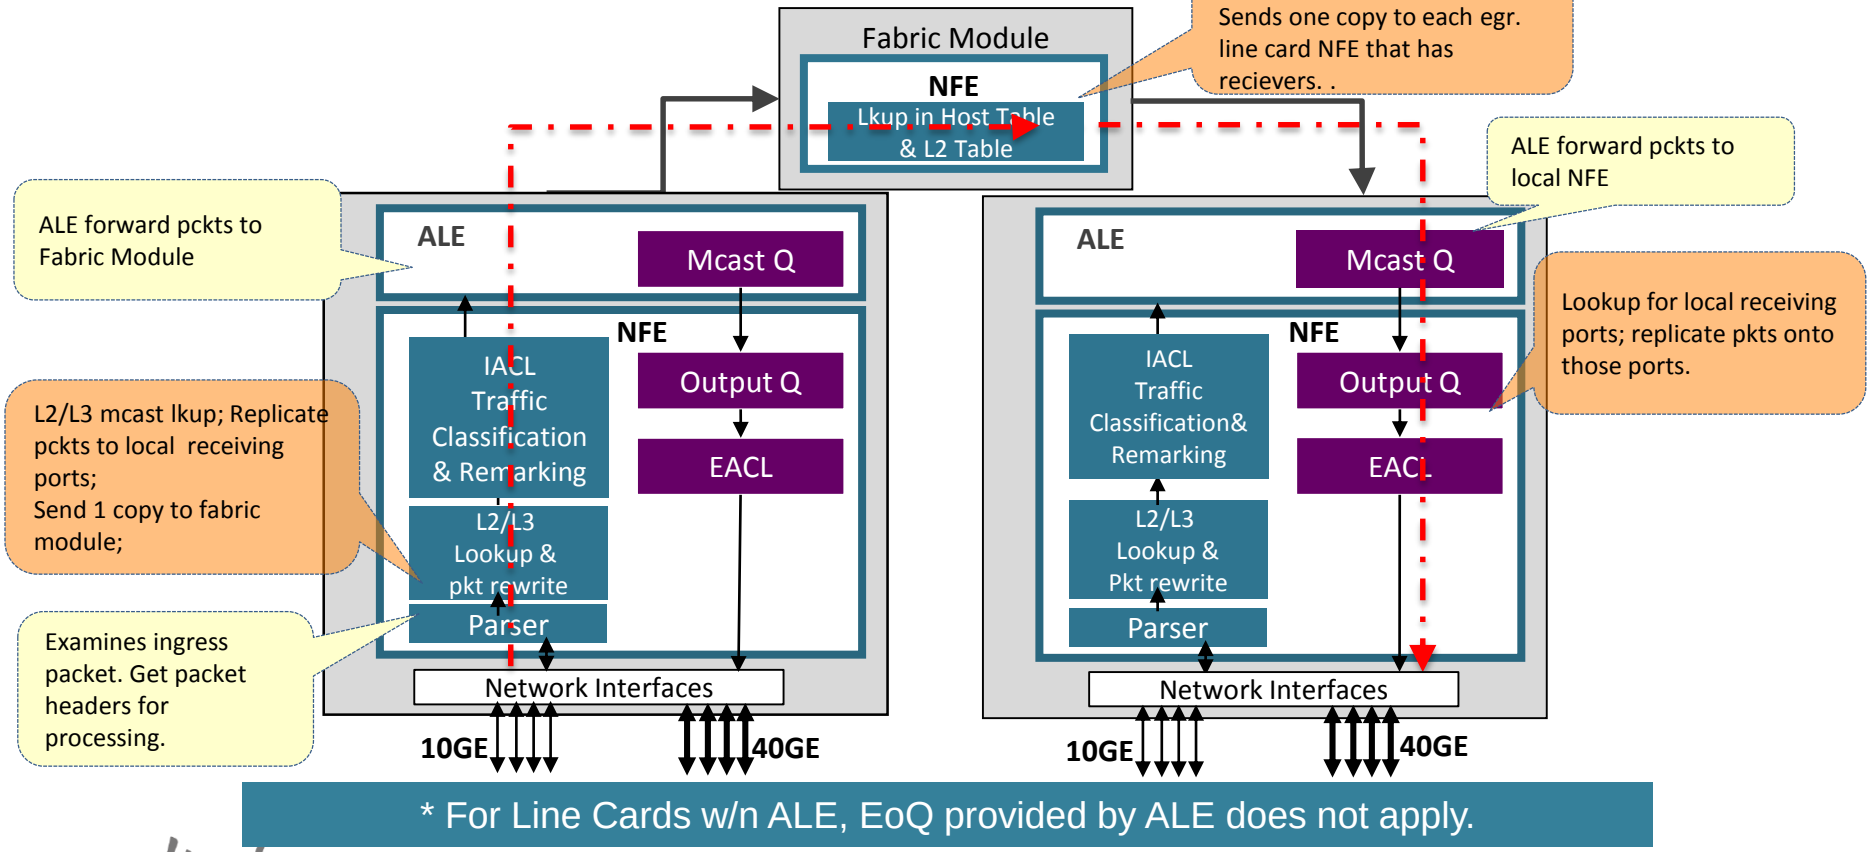

* For Line Cards w/n ALE, EoQ provided by ALE does not apply.

Agenda – Nexus 9000 Architecture

- Nexus 9000 Series Switch Hardware
  - Nexus 9000 Hardware Overview
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9500 Buffer and Queuing
  - Nexus 9300 Hardware Architecture
  - Nexus 9300 Buffer and Queuing
- Nexus 9000 and 40GE
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 Open NX-OS
- ACI & Nexus 9000

Buffer Resources on Nexus 9500 Switch Line Cards

Buffer On NFE (12MB)

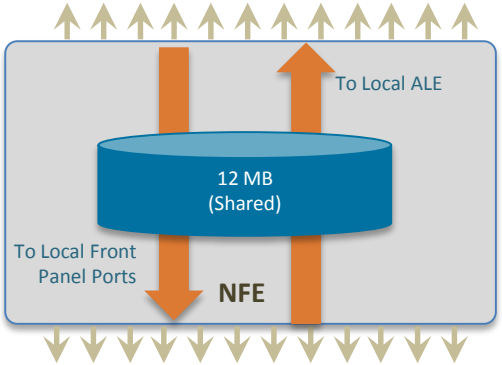

Buffer On ALE (40MB)

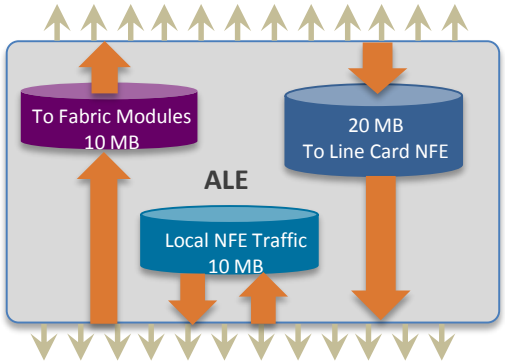

- Shared 10MB ingress buffer for fabric bound traffic;
- Shared 20MB egress buffer for traffic coming from fabric and going to front panel port.
- Shared 10MB egress buffer for local NFE traffic

Module Type	Buffer Size
N9K-X9636PQ (36x 40GE)	36 MB
N9K-X9432PQ (32x 40GE)	24 MB
N9K-X9464PX/TX (48x 1/10GE + 4x 40GE	12 MB
N9K-X9408PC-CFP2 (8x 100GE)	24 MB + 5MB/port

Module Type	Buffer Size
N9K-X9564TX/PX (48x 1/10G T + 4x 40G)	104 MB
N9K-X9536PX (36x 40G)	104 MB

Module Type	Buffer Size
Fabric Module	12 MB (9504 FM) 24 MB (9508 FM) 48 MB (9516 FM)

Buffer Boost on N9K-X9500 Series Line Cards

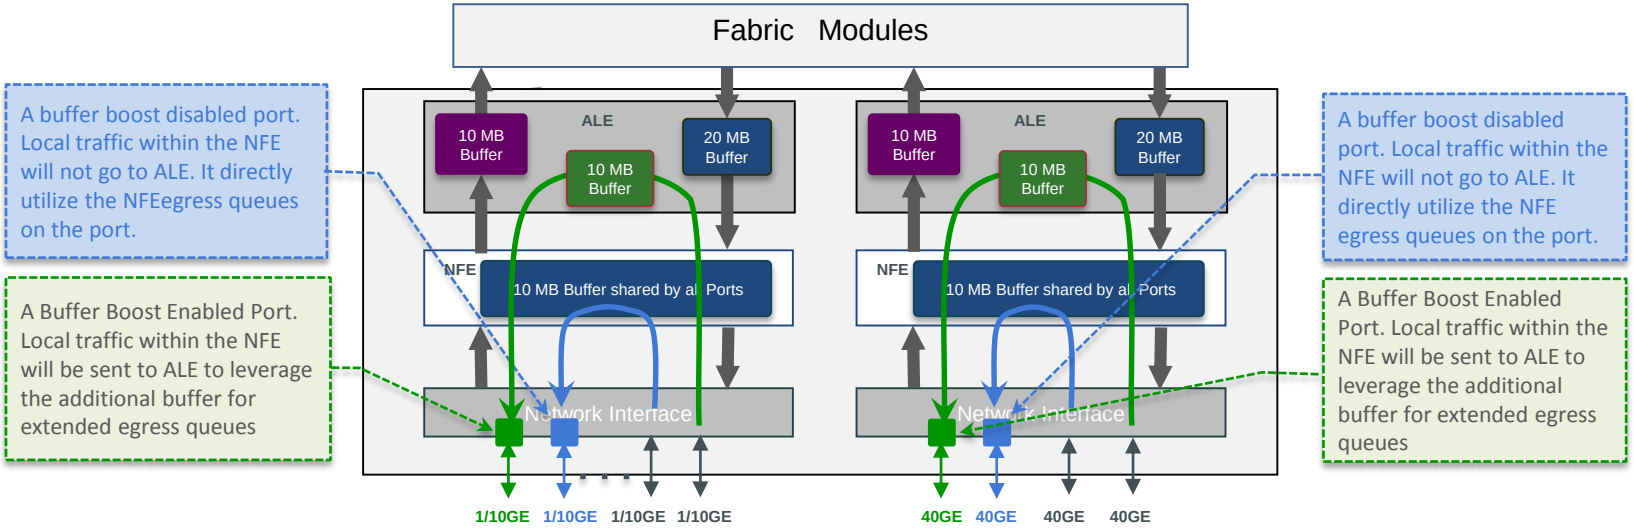

- Buffer boost function is introduced to allow line cards with ALEs to leverage the additional buffer on ALE for NFE local traffic.
- When buffer boost is enabled on a port, NFE local traffic to this port is sent to ALE for additional buffer space.
- When buffer boots is disabled on a port, NFE local traffic to this port remains local on this NFE.
- Buffer boots is enabled on ports of line cards with ALEs by default. It is a port property and can be disabled on a per-port basis.

Queuing & Scheduling Nexus 9500 Switch Line Cards

Egress Queuing on N9K-X9600/ N9K-X9400 Series Line Cards

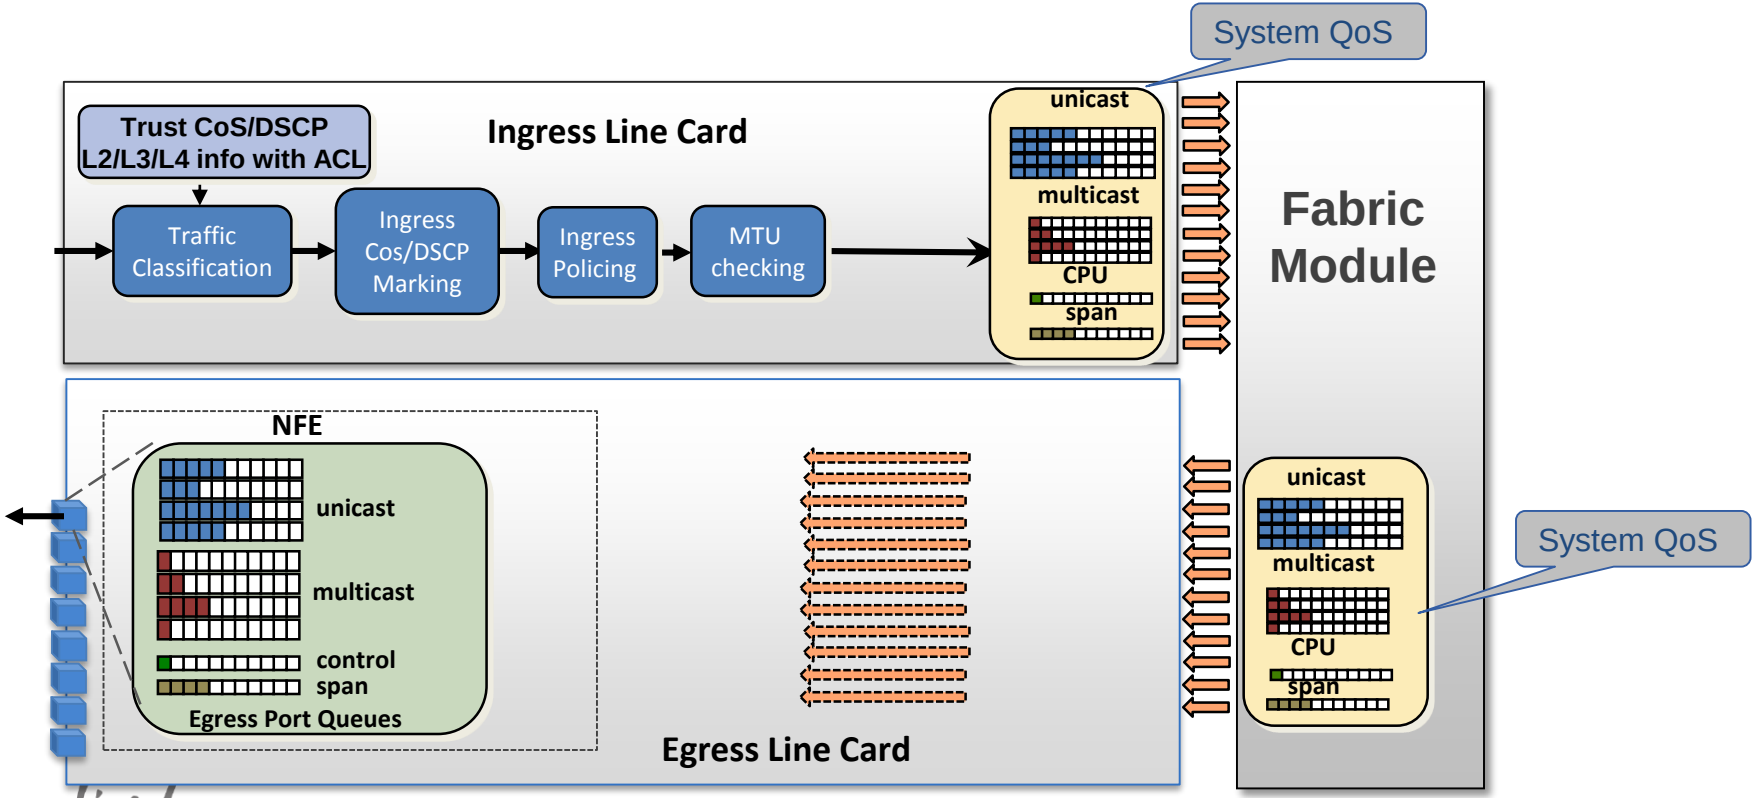

Queuing & Scheduling on Nexus 9500 Switch Line Cards

Extended Output Queues (EoQ) on N9K-9500 Series Line Cards (Cont'ed)

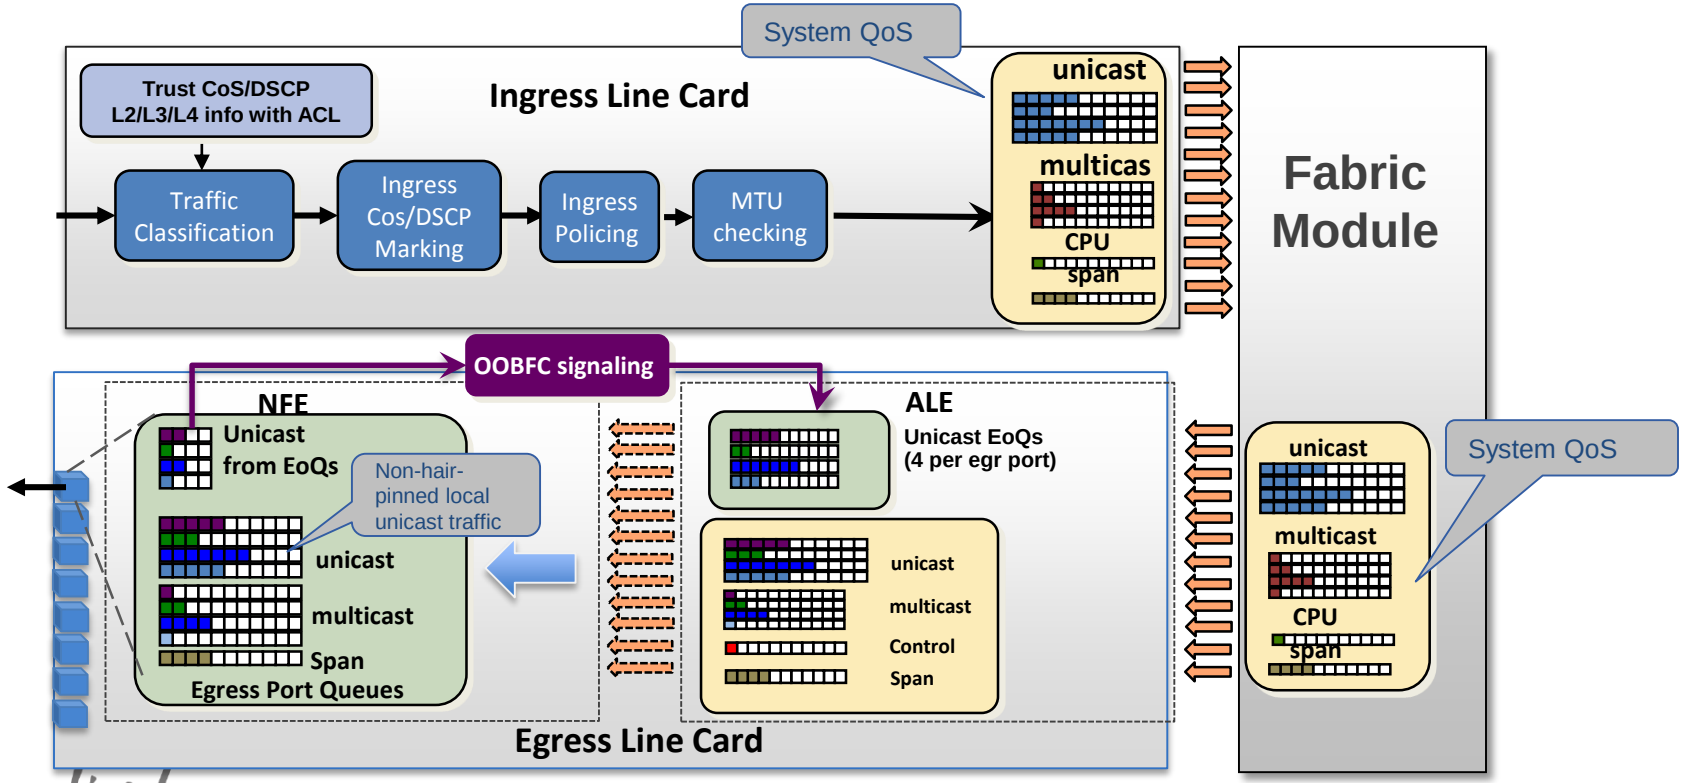

Queuing & Scheduling on Nexus 9500 Switch Line Cards

Extended Output Queues (EoQ) on N9K-9500 Series Line Cards (Cont'ed)

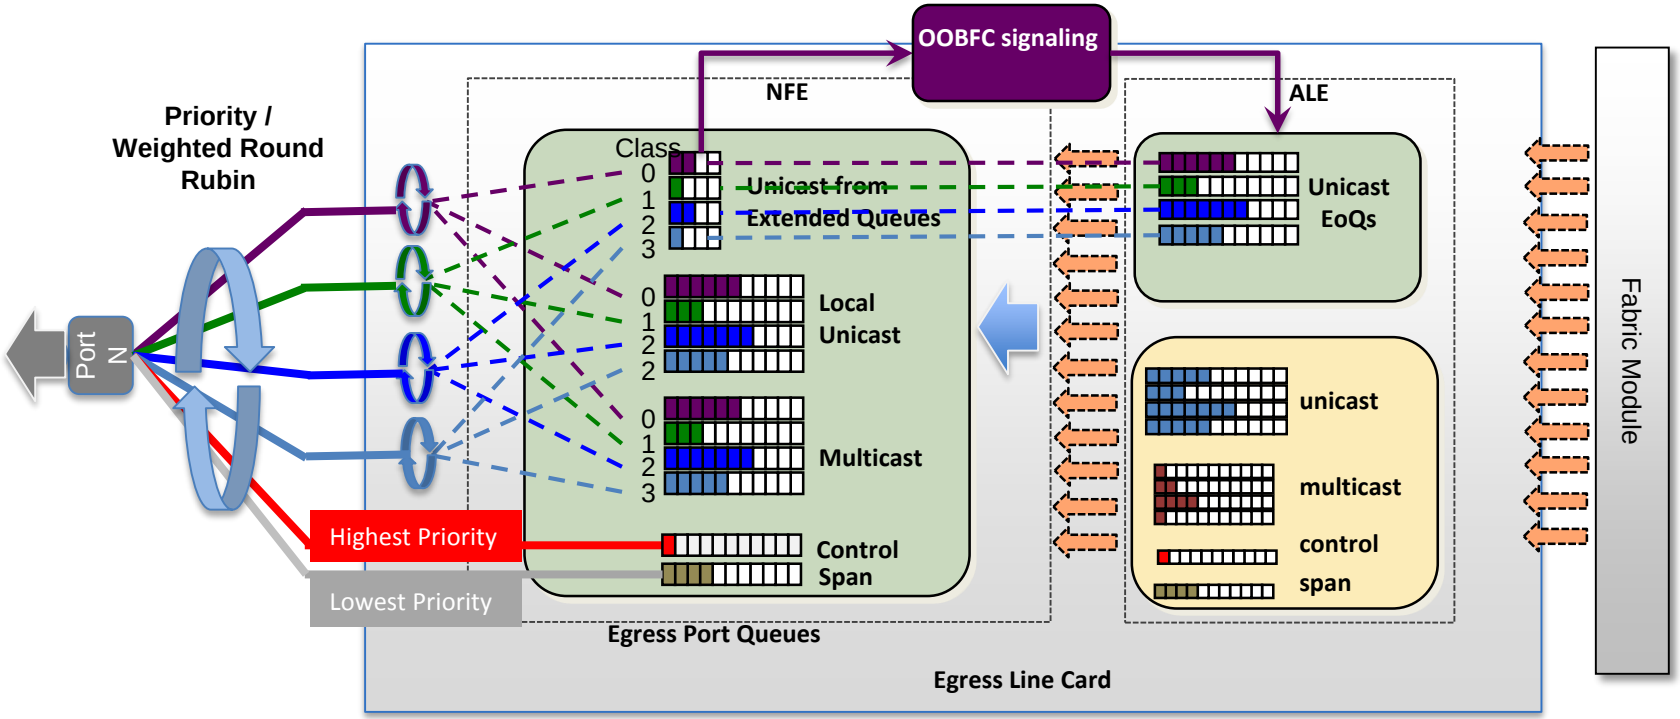

Cisco Nexus 9500 Series Switch Ingress Classification and Egress Queuing Operation in 4-Queue Mode

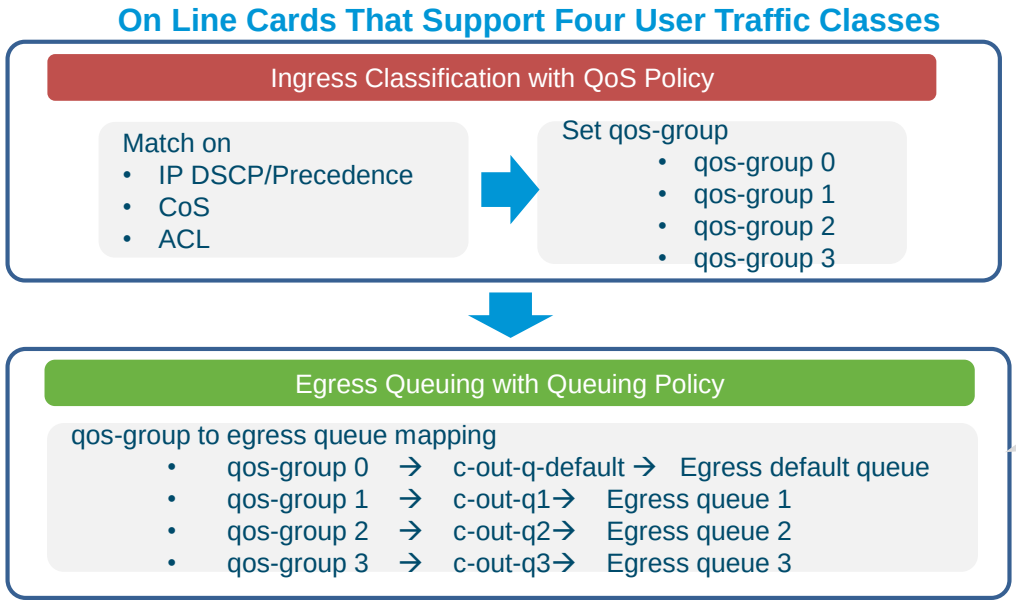

Sample Queueing Policy Configuration:

```
policy-map type queuing egr-queuing-policy
    class type queuing c-out-q3
      priority level 1
    class type queuing c-out-q2
      priority level 2
    class type queuing c-out-q1
      bandwidth remaining percent 60
    class type queuing c-out-q-default
      bandwidth remaining percent 0
```

Cisco Nexus 9500 Series Switch Ingress Classification and Egress Queuing Operation in 8-Queue Mode

On Line Cards That Support Eight User Traffic Classes

Ingress Classification with QoS Policy

Match on
- IP DSCP/Precedence
- CoS
- ACL

Set qos-group
- qos-group 0
- qos-group 1
- qos-group 2
- qos-group 3
- qos-group 4
- qos-group 5
- qos-group 6
- qos-group 7

Egress Queuing with Queuing Policy

qos-group to Egress Queue mapping

• qos-group 0	→	c-out-8q-q-default	→ Egress default queue
• qos-group 1	→	c-out-8q-q1	→ Egress queue 1
• qos-group 2	→	c-out-8q-q2	→ Egress queue 2
• qos-group 3	→	c-out-8q-q3	→ Egress queue 3
• qos-group 4	→	c-out-8q-q4	→ Egress queue 4
• qos-group 5	→	c-out-8q-q5	→ Egress queue 5
• qos-group 6	→	c-out-8q-q6	→ Egress queue 6
• qos-group 7	→	c-out-8q-q7	→ Egress queue 7

```
policy-map type queuing default-8q-
out-policy
  class type queuing c-out-8q-q7
    priority level 1
  class type queuing c-out-8q-q6
    priority level 2
  class type queuing c-out-8q-q5
    priority level 3
  class type queuing c-out-8q-q4
    bandwidth remaining percent 20
  class type queuing c-out-8q-q3
    bandwidth remaining percent 20
  class type queuing c-out-8q-q2
    bandwidth remaining percent 0
  class type queuing c-out-8q-q1
    bandwidth remaining percent 0
  class type queuing c-out-8q-q-
default
    bandwidth remaining percent 0
```

```
9508-1(config)# system qos
9508-1(config-sys-qos)# service-policy type queuing
output default-8q-out-policy
```

QoS on Nexus 9500 Series Switches

- **Ingress QoS (Classification and Remarking)**
  - Policy-map type qos
  - Match on CoS/ IP Precedence/ DSCP /ACL
  - Set qos-group
  - Remark CoS/ IP Precedence/ DSCP
  - Ingress policing

- **Network-QoS**
  - Policy-map type network-qos
  - Match on qos-group
  - Enable PFC

- **Egress Qos (Queuing and Shaping)**
  - Policy-map type queuing
  - Enable WRED or ECN (default is tail drop)
  - 4 or 8 user-defined classes based on qos-group
  - 1 control class for CPU and 1 class for SPAN traffic
  - Up to 3 PFC non-drop queues
  - Up to 3 Strict priority queues
  - Egress policing* and shaping (* only on NFE ports)

Agenda – Nexus 9000 Architecture

- Nexus 9000 Series Switch Hardware
  - Nexus 9000 Hardware Overview
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9500 Buffer and Queuing
  - Nexus 9300 Hardware Architecture
  - Nexus 9300 Buffer and Queuing
- Nexus 9000 and 40GE
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 Open NX-OS
- ACI & Nexus 9000

Nexus 9300 Series Switches Portfolio

Nexus® 9372PX/ 9372TX
- 1 RU w/n GEM module slot
- 720Gbps
- 6-port 40 Gb QSFP+
- 48-port 1/10 Gb SFP+ on Nexus 9372PX
- 48-port 1/10 G-T on Nexus 9372TX

Nexus 9332PQ
- 1 RU w/n GEM module slot
- 1,280Gbps
- 32-port 40 Gb QSFP+

Nexus 93120TX
- 2 RU w/n GEM module slot
- 1200Gbps
- 6-port 40 Gb QSFP+
- 96-port 1/10 G-T

Nexus® 9396PX/ 9396TX
- 2 RU with 1 GEM module slot
- 960Gbps
- 48-port 1/10 Gb SFP+ on Nexus 9396PX
- 48-port 1/10 G-T on Nexus 9396TX
- 6 ports 40 Gb QSFP+ on N9K-M6PQ GEM module
- 12 ports 40 Gb QSFP+ on N9K-M12PQ GEM module
- 4 ports 100 Gb CFP2 on N9K-M4PC-CFP2 GEM module

Nexus 93128TX/ 93128PX
- 3 RU with 1 GEM module slot
- 1,280Gbps
- 96-port 1/10 G-T on Nexus 93128TX
- 96-port 1/10 SFP+ on Nexus 93128P
- 6 ports 40 Gb QSFP+ on N9K-M6PQ GEM module
- 8 ports 40 Gb QSFP+ on N9K-M12PQ GEM module
- 2 ports 100 Gb CFP2 on N9K-M4PC-CFP2 GEM module

Nexus 9300 Series Switch GEM Modules

Generic Extension Module (GEM)

N9K-M12PQ

- 12-port 40 Gb QSFP+
- 1 ALE
- Additional 40 MB buffer on ALE
- Full VXLAN gateway/bridging and routing capability
- Supported by all Nexus 9300 series switches with GEM module slots
- ACI capable

N9K-M6PQ

- 6-port 40 Gb QSFP+
- 1 ALE-2
- Additional 25 MB buffer on ALE-2
- Full VXLAN gateway/ bridging and routing capability
- Supported by all Nexus 9300 series switches with GEM module slots
- ACI capable

N9K-M4PC-CFP2

- 4-port 100 Gb CFP2
- 4 100GE-MAC
- Additional 5 MB buffer/ 100 Gb port
- Full VXLAN gateway/ bridging capability
- Supported by all Nexus 9300 series switches with GEM module slots
- Not ACI capable
- When used with N9K-C93128TX, only the right two ports are active

Nexus 9300 Series Switch GEM Modules(Cont'ed)

Generic Extension Module (GEM)

N9K-M12PQ

12x 40GE
Front Panel Ports

12x 42GE
Internal Links to NFE

- Supports both NX-OS and ACI
- Supports both VXLAN bridging and routing
- 40GE front panel ports don't support break-out mode

N9K-M6PQ

6x 40GE
Front Panel Ports

6x 42GE
Internal Links to NFE

- Supports both NX-OS and ACI
- Supports both VXLAN bridging and routing
- 40GE front panel ports don't support break-out mode

N9K-M4PC-CFP2

New!

4x 100GE (CFP2)
Front Panel Ports

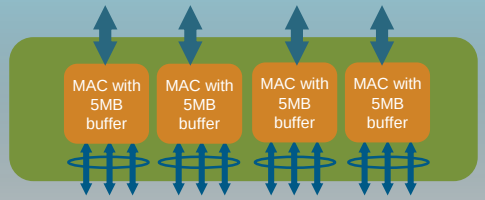

12x 40GE
Internal Ethernet Links to NFE

- Supports only NX-OS
- Supports VXLAN bridging, but not VXLAN routing
- 100GE front panel ports don't support break-out mode

Nexus 9300 Platform Architecture

Cisco Nexus® 9396PX / 9396TX

- 2 RU height
- 48x 1Gb SFP / 10 Gb SFP+ ports on Nexus 9396PX
- 48x 1/10 Gb Base-T ports on Nexus 9396TX
- 12x 40 Gb QSFP+ ports on N9K-M12PQ GEM module
- 6x 40 Gb QSFP+ ports on N9K-M6PQ GEM module
- 4x 100 Gb CFP2 ports on N9K-M4PC-CFP2 GEM module
- 1 100/1000baseT management port
- 1 RS232 console port
- 2 USB 2.0 ports
- Front-to-back and back-to-front airflow options
- 1+1 redundant power supply options
- 2+1 redundant fans
- No-blocking architecture with full line-rate performance on all ports for all packet sizes
- VXLAN bridging & routing with N9K-M12PQ or N9K-M6PQ GEM module
- VXLAN bridging only with N9K-M4PC-CFP2 GEM module

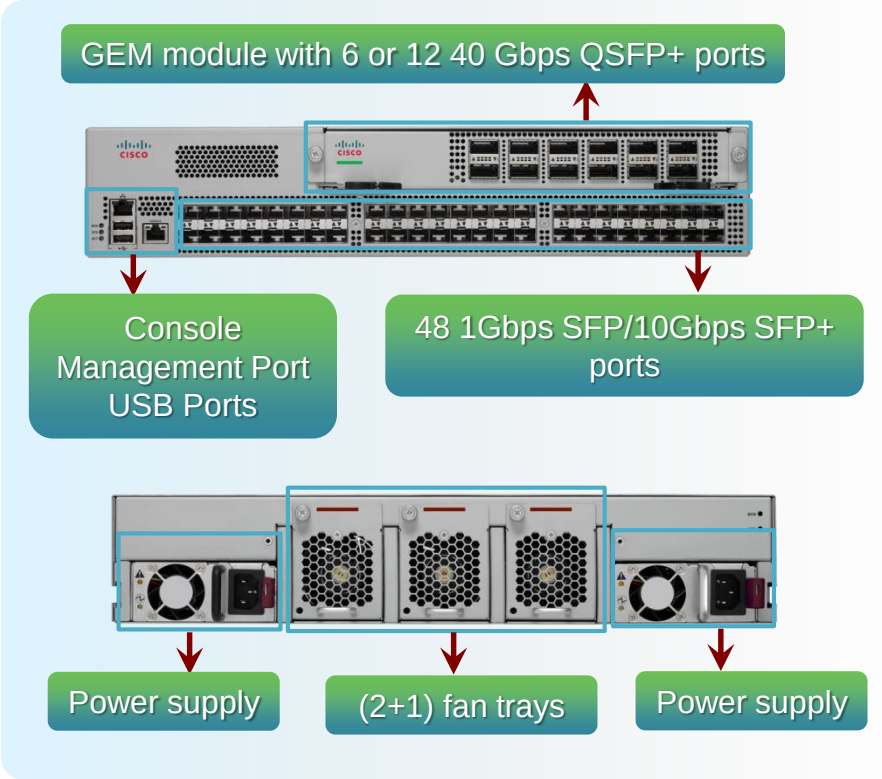

Nexus 9300 Series Switch Architecture

Nexus 9396PX/TX Block Diagram with N9K-M12PQ or N9K-M6PQ GEM Module

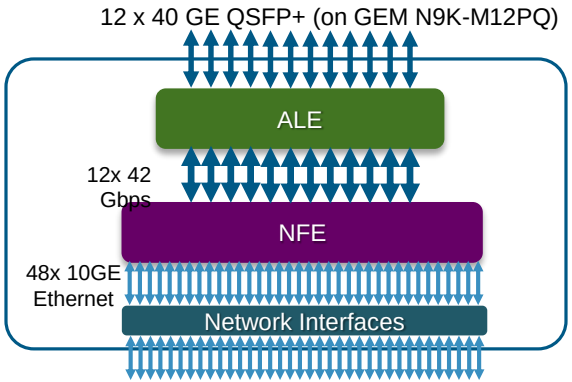

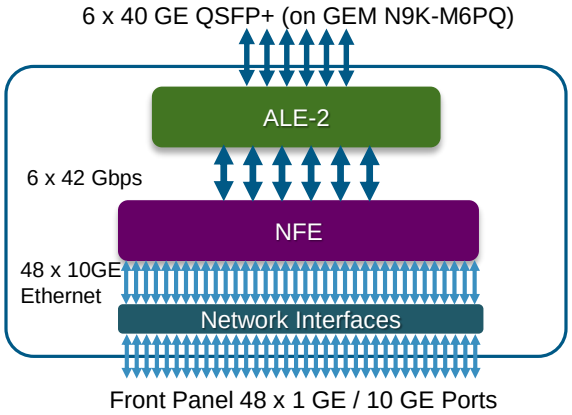

Nexus® 9396PX/TX with N9K-M12PQ GEM Module

Nexus® 9396PX/TX with N9K-M6PQ GEM Module

- Hardware is capable of VXLAN bridging and routing
- Hardware is capable of supporting both NX-OS and ACI
- Line rate performance for packet sizes > 200-Bytes

Nexus 9300 Series Switch Architecture

Nexus 9396PX/TX Block Diagram with N9K-M4PC-CFP2 GEM Module

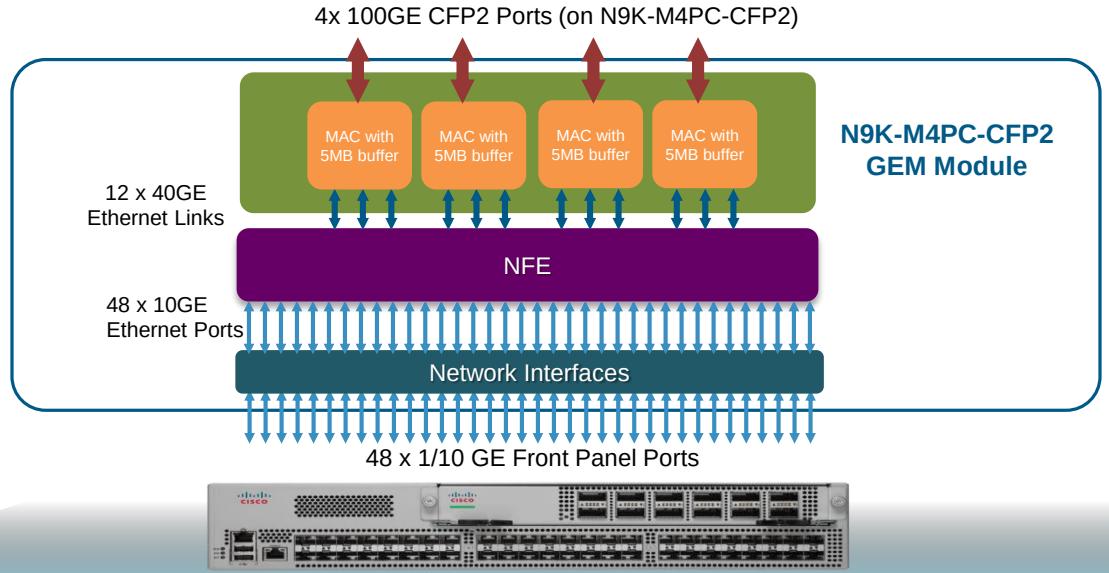

Nexus® 9396PX/TX with N9K-M4PC-CFP2 GEM Module

- Hardware is capable of VXLAN bridging only
- Hardware is capable of supporting NX-OS only
- Line rate performance for packet sizes > 200-Bytes

Nexus 9300 Series Switch Architecture

Cisco Nexus® 93128TX

- 3 RU height
- 96 1/10 Gbase-T ports
- 6 active 40GE QSFP+ ports on N9K-M6PQ GEM module
- 8 active 40GE QSFP+ ports on N9K-M12PQ GEM module
- 2 active 100GE CFP2 ports on N9K-M4PC-CFP2 GEM module
- 1 100/1000baseT management port
- 1 RS232 console port
- 2 USB 2.0 ports
- Front-to-back and back-to-front airflow options
- 1+1 redundant power supply options
- 2+1 redundant fans
- Line-rate performance for packet sizes > 200-Bytes
- VXLAN bridging & routing with N9K-M12PQ or N9K-M6PQ GEM module
- VXLAN bridging only with N9K-M4PC-CFP2 GEM module

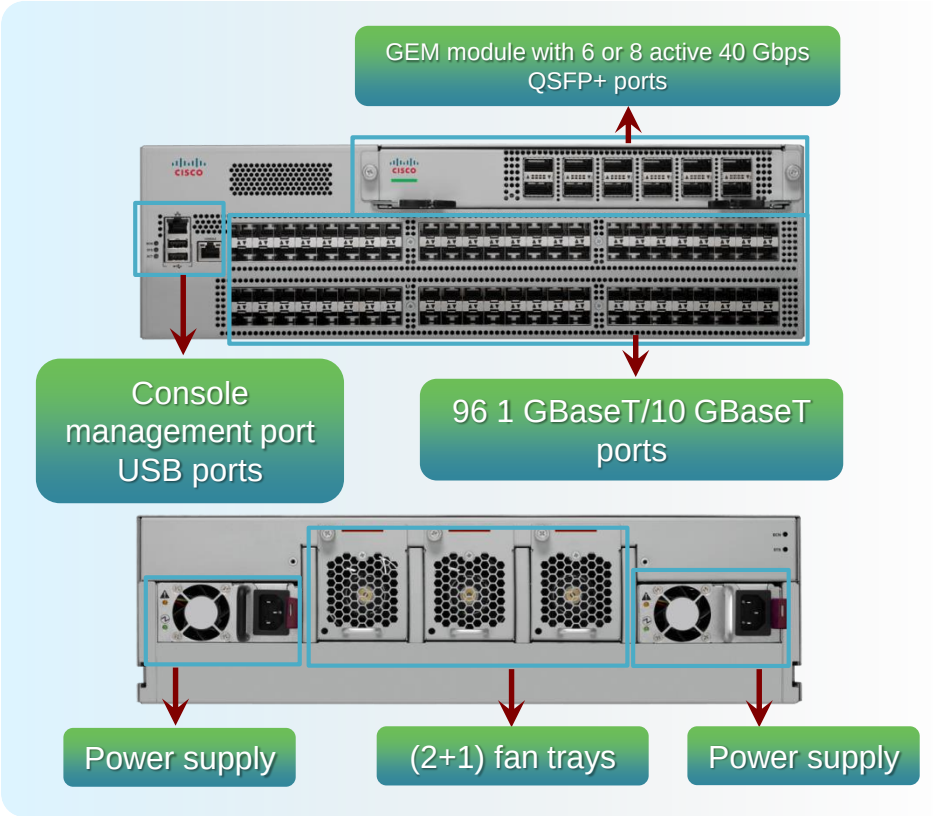

Nexus 9300 Series Switch Architecture

Nexus 93128TX Block Diagram with N9K-M12PQ or N9K-M6PQ GEM Module

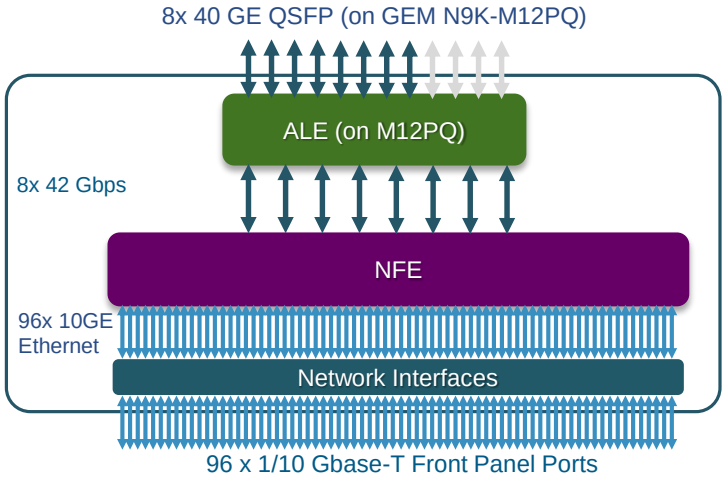

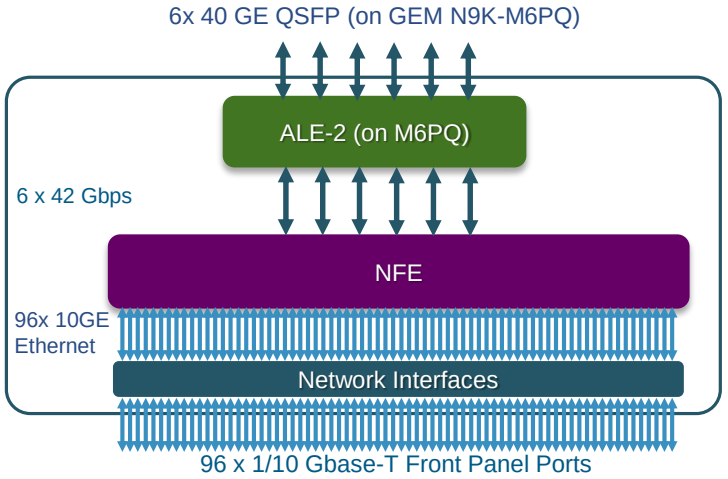

Nexus® 93128TX switch with N9K-M12PQ GEM Module

Nexus® 93128TX switch with N9K-M6PQ GEM Module

- Hardware is capable of VXLAN bridging and routing
- Hardware is capable of supporting both NX-OS and ACI
- Line rate performance for packet sizes > 200-Bytes

Nexus 9300 Series Switch Architecture

Nexus 93128TX Block Diagram with N9K-M4PC-CFP2 GEM Module

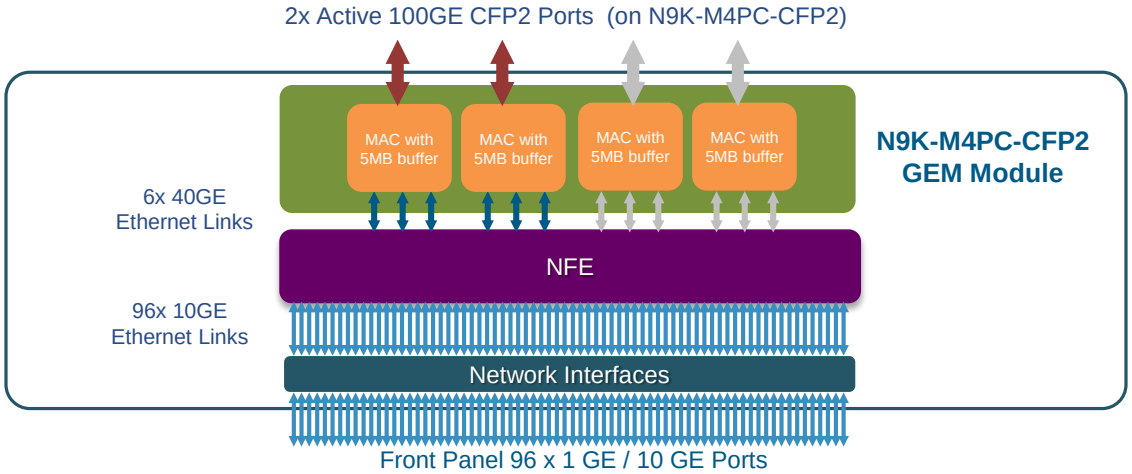

Nexus® 93128TX switch with M9K-M4PC-CFP2 GEM Module

- Hardware is capable of VXLAN bridging only
- Hardware is capable of supporting NX-OS only
- Line rate performance for packet sizes > 200-Bytes

Nexus 9300 Series Switch Architecture

Cisco Nexus® 9372PX / 9372TX

- 1 RU height
- No GEM module
- 48x 1Gb SFP / 10 Gb SFP+ ports on Nexus 9372PX
- 48x 1/10 Gb Base-T ports on Nexus 9372TX
- 6x 40 Gb QSFP+ ports
- 1 100/1000baseT management port
- 1 RS232 console port
- 2 USB 2.0 ports
- Front-to-back and back-to-front airflow options
- 1+1 redundant power supply options
- 2+1 redundant fans
- Full line rate performance for all packet sizes
- VXLAN bridging and routing
- Capable of supporting both NX-OS and ACI modes

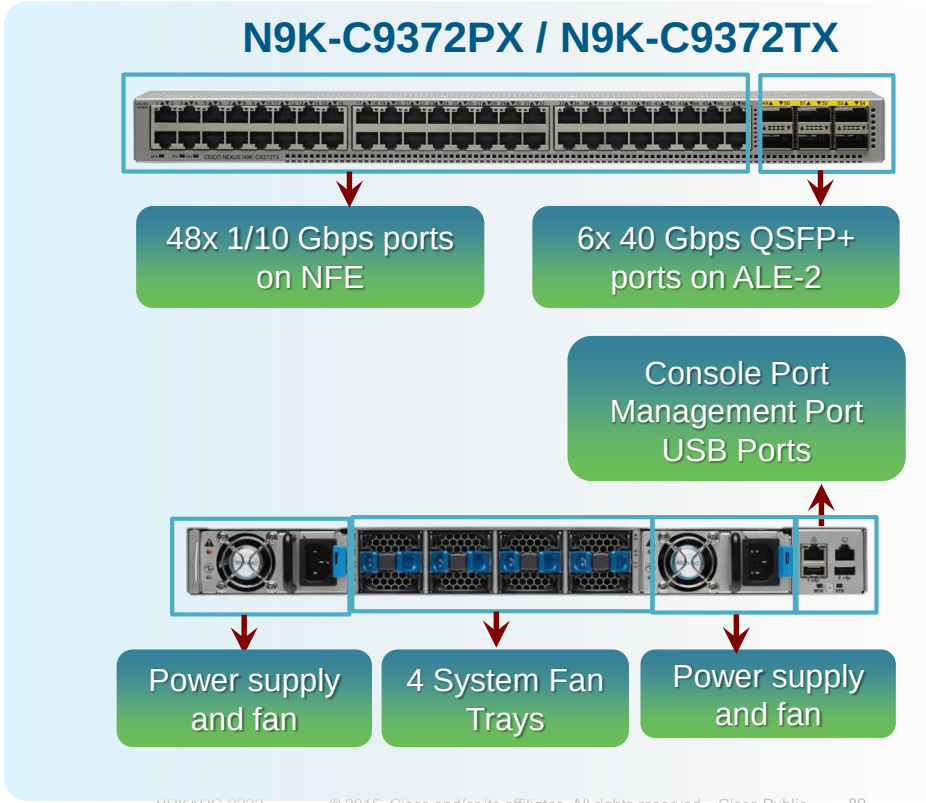

Nexus 9300 Series Switch Architecture

Nexus 9372PX/ Nexus 9372TX Block Diagram

1 application 1leaf engines (ALE-2)) for additional buffering and packet handling

1 network forwarding engine (NFE)

1 RU with redundant power supplies and fan.
6 QSFP+ 40GE ports and 40 SFP+ 10GE ports

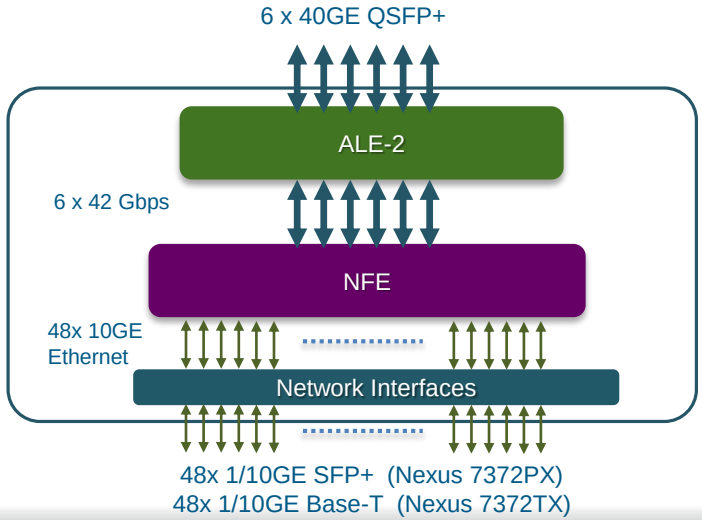

Nexus® 9372PX, Nexus 9372TX

- The 6 40GE links between NFE and ALE-2 run at 42Gbps clock rate to accommodate the internal packet header.
- Hardware is capable of VXLAN bridging and routing
- Hardware is capable of supporting both NX-OS and ACI modes
- Full line rate performance for all packet sizes

Nexus 9300 Series Switch Architecture

Cisco Nexus® 9332PQ

- 1 RU height
- No GEM module
- 32x 40 Gb QSFP=
- 1 100/1000baseT management port
- 1 RS232 console port
- 2 USB 2.0 ports
- Front-to-back and back-to-front airflow options
- 1+1 redundant power supply options
- 2+1 redundant fans
- Line-rate performance for packet sizes > 200-Bytes
- VXLAN bridging & routing
- Capable of supporting both NX-OS and ACI modes

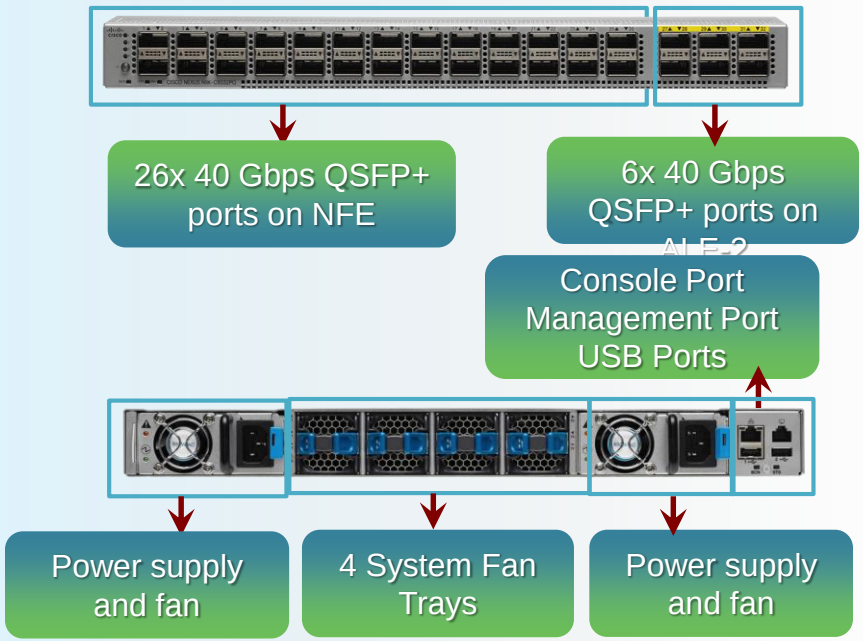

Nexus 9300 Series Switch Architecture

Nexus 9332PQ Block Diagram

1 application 1leaf engines (ALE-2)) for additional buffering and packet handling

1 network forwarding engine (NFE)

1 RU with redundant power supplies and fan trays.
32 QSFP+ 40GE ports. 6 are on ALE-2 and 26 are on NFE.

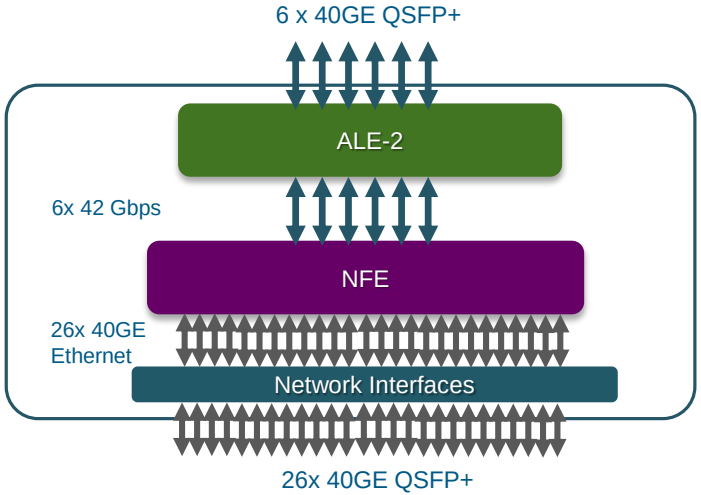

Nexus® 9332PQ

- The 6 40GE links between NFE and ALE-2 run at 42Gbps clock rate to accommodate the internal packet header.
- Hardware is capable of VXLAN bridging and routing
- Hardware is capable of supporting both NX-OS and ACI modes
- Line rate performance for packet sizes > 200-Bytes

Nexus 9300 Platform Architecture

New!

Cisco Nexus® 93120TX

- 2 RU height
- No GEM module
- 96x 1/10 GBase-T ports
- 6x 40GE QSFP ports
- 1 100/1000baseT management port
- 1 RS232 console port
- 2 USB 2.0 ports
- Front-to-back and back-to-front airflow options
- 1+1 redundant power supply options
- 2+1 redundant fans
- Line-rate performance for packet sizes > 200-Bytes
- VXLAN bridging & routing
- Capable of supporting both NX-OS and ACI modes

N9K-C93120TX

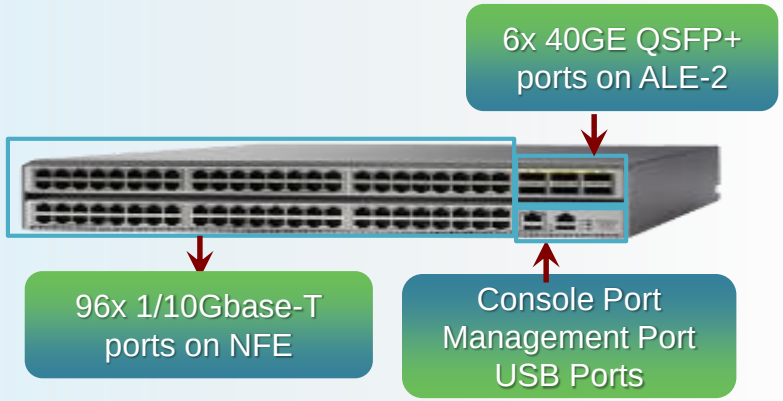

Nexus 9300 Series Switch Architecture

Nexus 93120TX Block Diagram

New!

1 application 1leaf engines (ALE-2)) for additional buffering and packet handling

1 network forwarding engine (NFE)

1 RU with redundant power supplies and fan trays.
32 QSFP+ 40GE ports. 6 are on ALE-2 and 26 are on NFE.

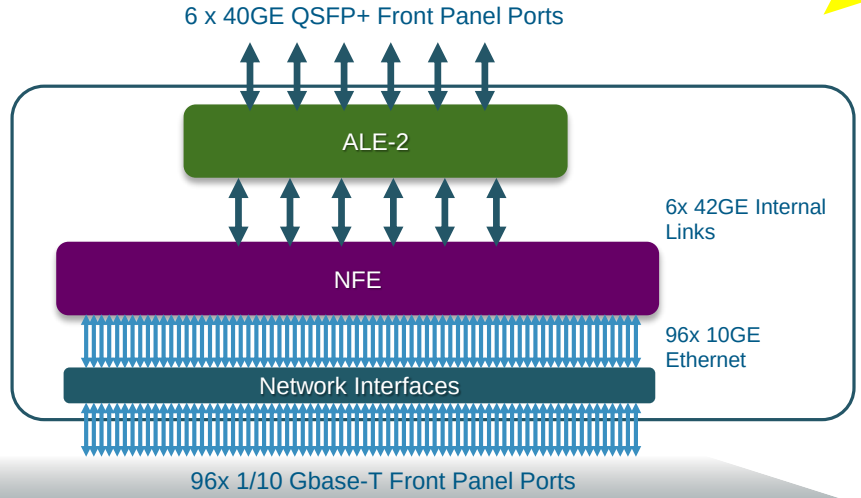

- The 6 40GE links between NFE and ALE-2 run at 42Gbps clock rate to accommodate the internal packet header.
- Hardware is capable of VXLAN bridging and routing
- Hardware is capable of supporting both NX-OS and ACI
- Full line rate performance for all packet sizes

Nexus 9300 Series Switch Forwarding Tables

- In Nexus 9300 system there is no separate NFE ASIC that would distribute LPM Route learning from the rest of the system.
- As a result of this, the forwarding tables on a single NFE ASIC is completely responsible for LPM Routes, Host Routes, and MAC Address learning.
- However, it is possible to adjust the allocation of table space based on defined templates.

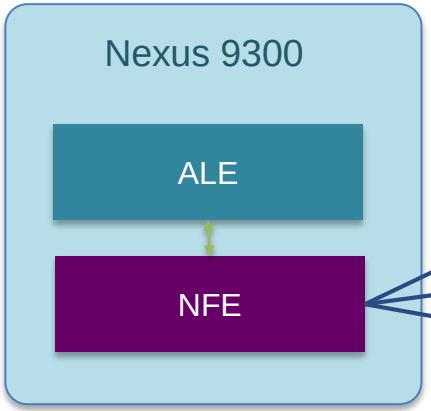

Feature	
LPM Routes	16K
Host Routes	120K
MAC Addresses	96K

Nexus 9300 Forwarding Scalabilities

	Nexus 9300	
	Default	Maximum Layer-3 Mode
LPM Routes	16K	128K
IP Host Entries	120K (208K protocol learned IPv4 host routes)	16K
MAC Address Entries	96K	32K
Multicast Routes	32K* (hardware capable of 72K)	8K*
Multicast Fan Outs	8K (no vPC)	8K (no vPC)
IGMP Snooping Groups	32K* (hardware capable of 72K)	8K*
ACL TCAM	Hardware: 4K ingress, 1K egress Available to user: 3K ingress, 768 egress	Hardware: 4K ingress, 1K egress Available to user: 3K ingress, 768 egress
VRF	1000	1000
Max Links in Port Channel	32	32
Max ECMP Paths	64	64
Max vPC Port Channels	528	528
Max Active SPAN Sessions	4	4
Max RPVST Instances	507	507
Max HSRP Groups	490	490
MAX VLANs	4K	4K
SPAN/ERSPAN	4 active sessions	4 active sessions

* Shared with IP hosts

Agenda – Nexus 9000 Architecture

- Nexus 9000 Series Switch Hardware
  - Nexus 9000 Hardware Overview
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9500 Buffer and Queuing
  - Nexus 9300 Hardware Architecture
  - Nexus 9300 Buffer and Queuing
- Nexus 9000 and 40GE
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 Open NX-OS
- ACI & Nexus 9000

Why Speedup in Network Design

Higher speed links improve ECMP efficiency

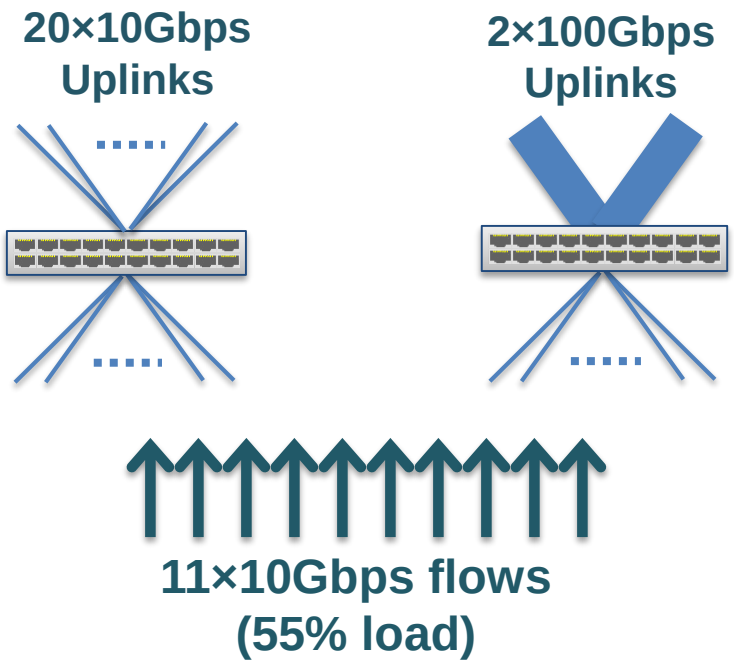

Prob of 100% throughput = 3.27%

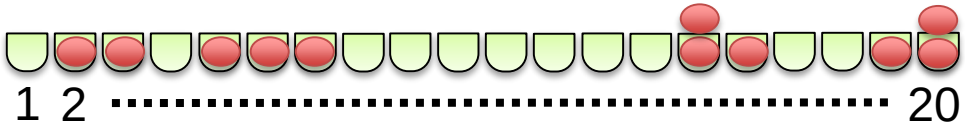

Prob of 100% throughput = 99.95%

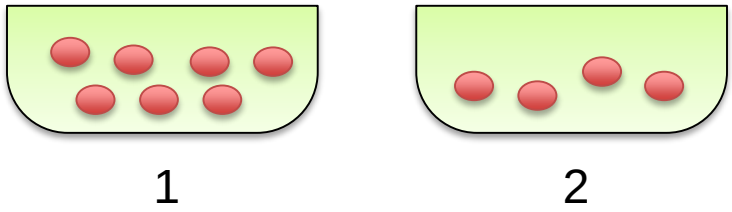

Impact of Buffering

Where are large buffers more effective for Incast?

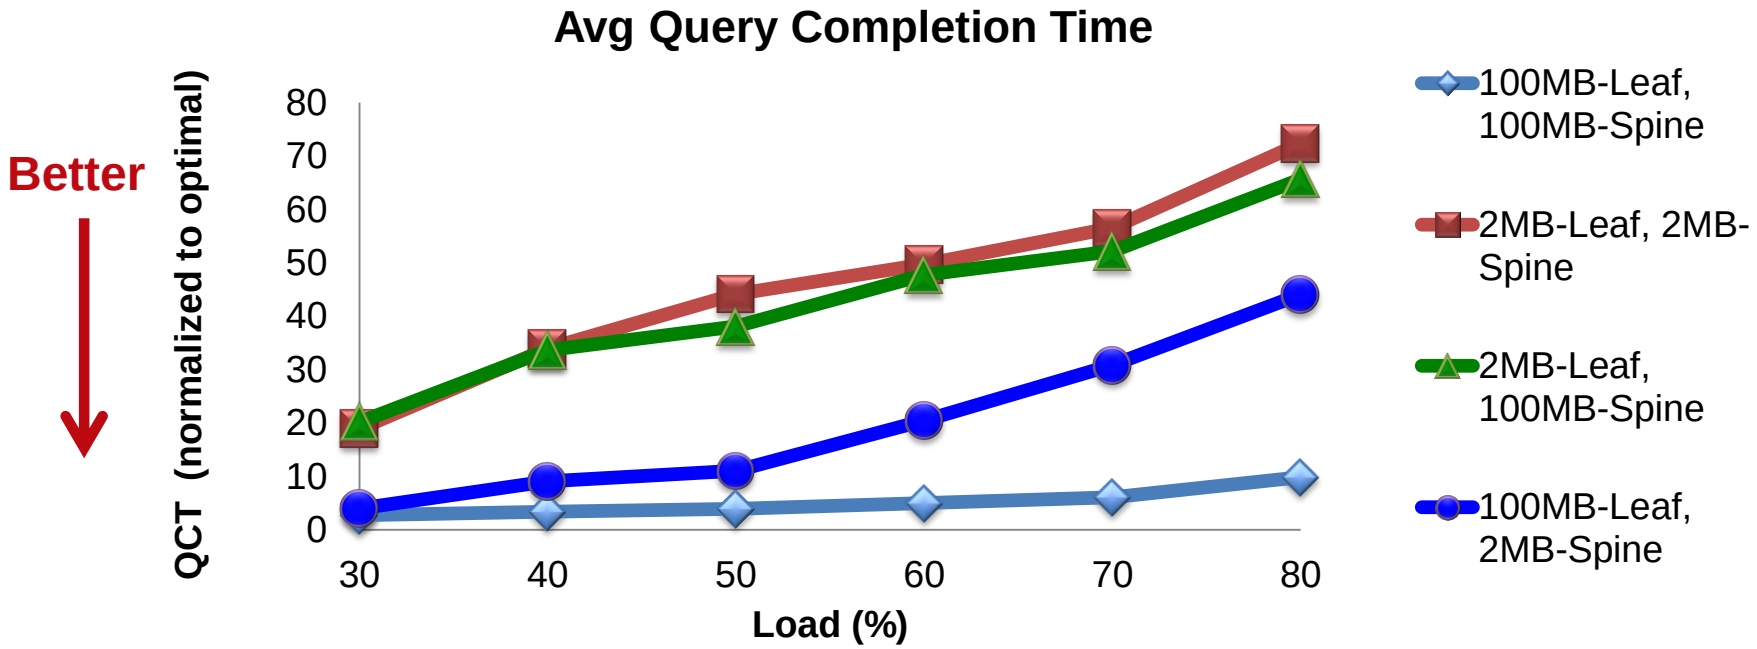

Nexus 9300 Series Switch Buffer Architecture

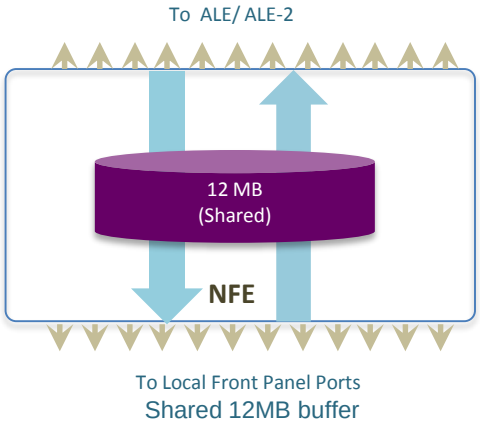

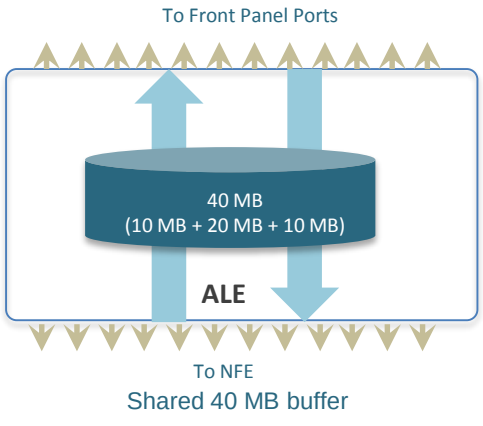

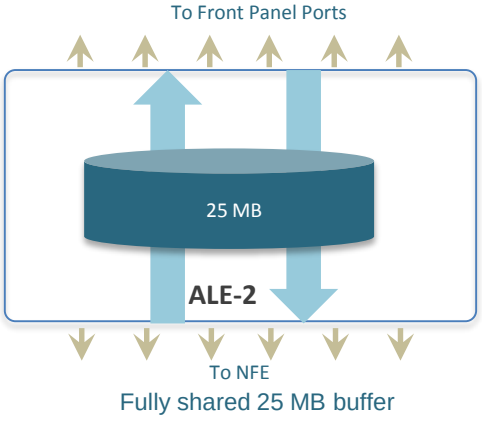

Buffer Location	Buffer Size
NFE (Network Forwarding Engine)	12 MB
ALE (Application Leaf Engine)	40 MB
ALE-2 (Application Leaf Engine-2)	25 MB

Nexus 9300 Series Switch Buffer Architecture (Cont'ed)

NFE + ALE Buffer Structure

Buffer on NFE:
- Fully shared 12 MB buffer for ports on NFE

Buffer on ALE
- Shared 10 MB egress buffer for traffic going to all ALE front-panel ports
- Shared 20 MB egress buffer for traffic coming from ALE front-panel ports and going to NFE front panel port.
- Shared 10MB egress buffer for local NFE traffic

ALE can be found on the N9K-M12PQ GEM module. Therefore, the NFE+ ALE buffer structure can be seen in the following Nexus 9300 switch platforms when they're equipped with N9K-M12PQ GEM module:
- Nexus 9396PX
- Nexus 9396TX
- Nexus 93128TX

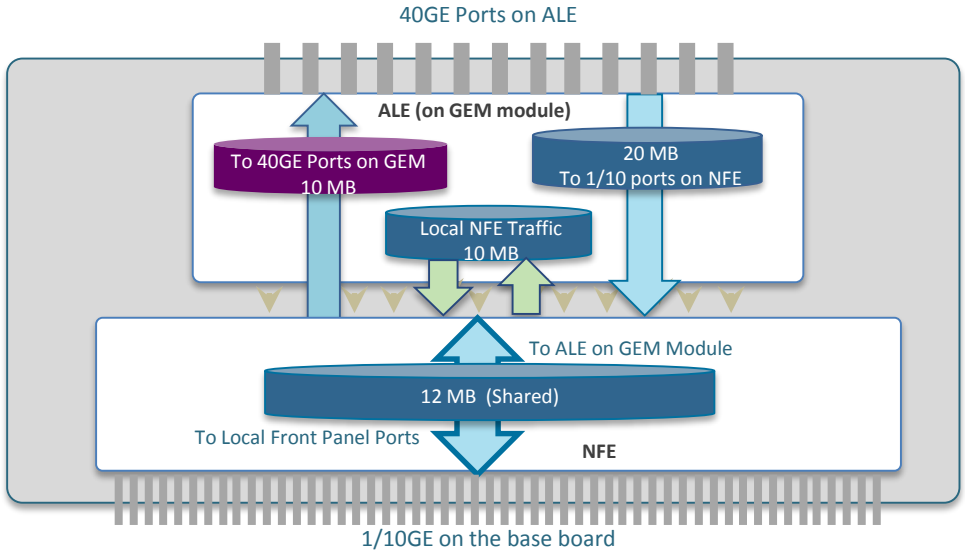

Nexus 9396PX/TX with N9K-M12PQ GEM Module
Nexus 93128PTX with N9K-M12PQ GEM Module

Nexus 9300 Series Switch Buffer Architecture (Cont'ed)

NFE + ALE-2 Buffer Structure

Buffer on NFE:
- Fully shared 12 MB buffer for all ports on NFE

Buffer on ALE-2
- Fully shared 25 MB buffer for all ports on ALE-2

ALE-2 can be found on the N9K-M6PQ GEM module or selected Nexus 9300 switch base-board. The NFE+ ALE-2 buffer structure can be seen in the following Nexus 9300 switch platforms:

- Nexus 9396PX with N9K-M6PQ GEM Module
- Nexus 9396TX with N9K-M6PQ GEM Module
- Nexus 93128TX with N9K-M6PQ GEM Module
- Nexus 9372PX
- Nexus 9372TX
- Nexus 9332PQ

The 25MB buffer on ALE-2 is shared by traffic in all directions through ALE-2. This fully dynamically buffer sharing allows ALE-2 to provide more per-port buffer in each direction with less total buffer size in comparison to ALE, achieving better performance and optimized buffer resource utilization.

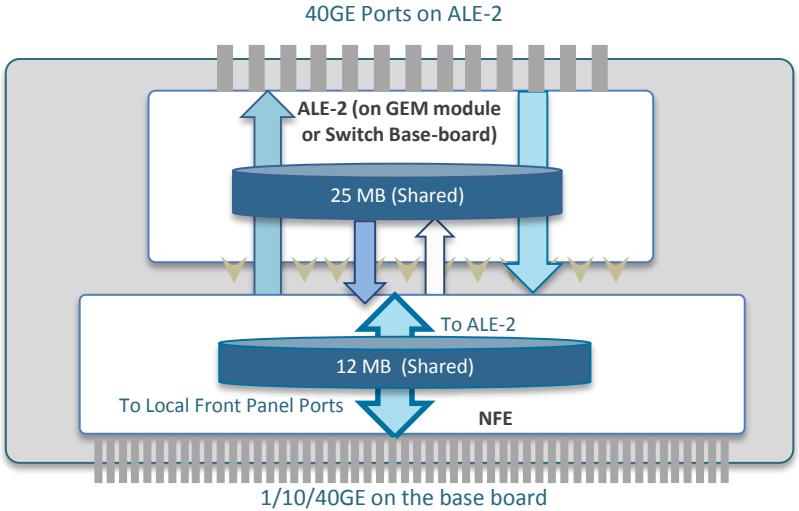

Nexus 9396PX/TX with N9K-M6PQ GEM Module
Nexus 9396PX/TX with N9K-M6PQ GEM Module
Nexus 93128TX with N9K-M6PQ GEM Module
Nexus 9372PX/TX,
Nexus 9332PQ

Buffer Boost on Nexus 9300 Serials Switches

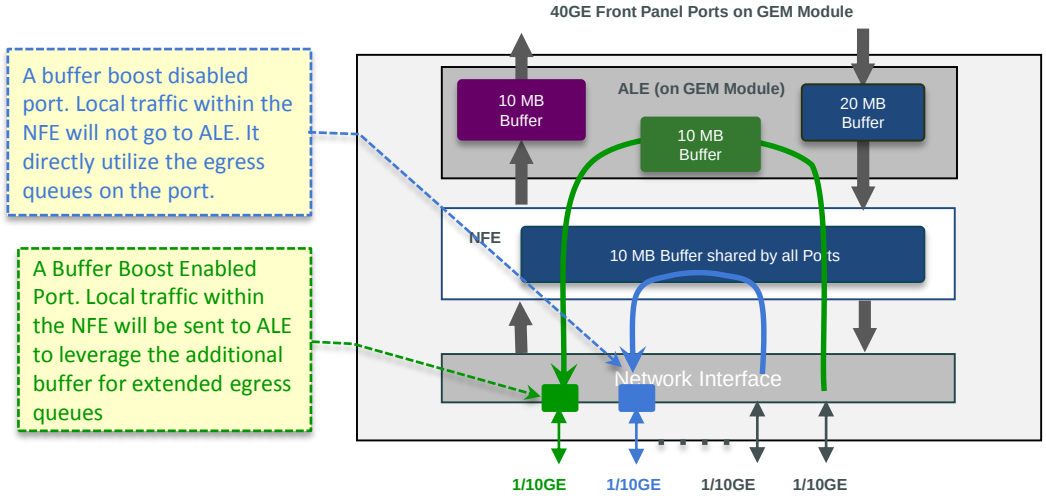

- Buffer boost function is introduced on 1/10GE ports to leverage the additional buffer on ALE for local traffic.
- When buffer boost is enabled on a port, local traffic from other 1/10GE ports to this port is sent to ALE for additional buffer space.
- When buffer boots is disabled on a port, local traffic from other 1/10GE ports to this port remains local on NFE.
- Buffer boots is enabled on all 1/10GE ports by default. It is a port property and can be disabled on a per-port basis.

Queuing & Scheduling on Nexus 9300
Extended Output Queues

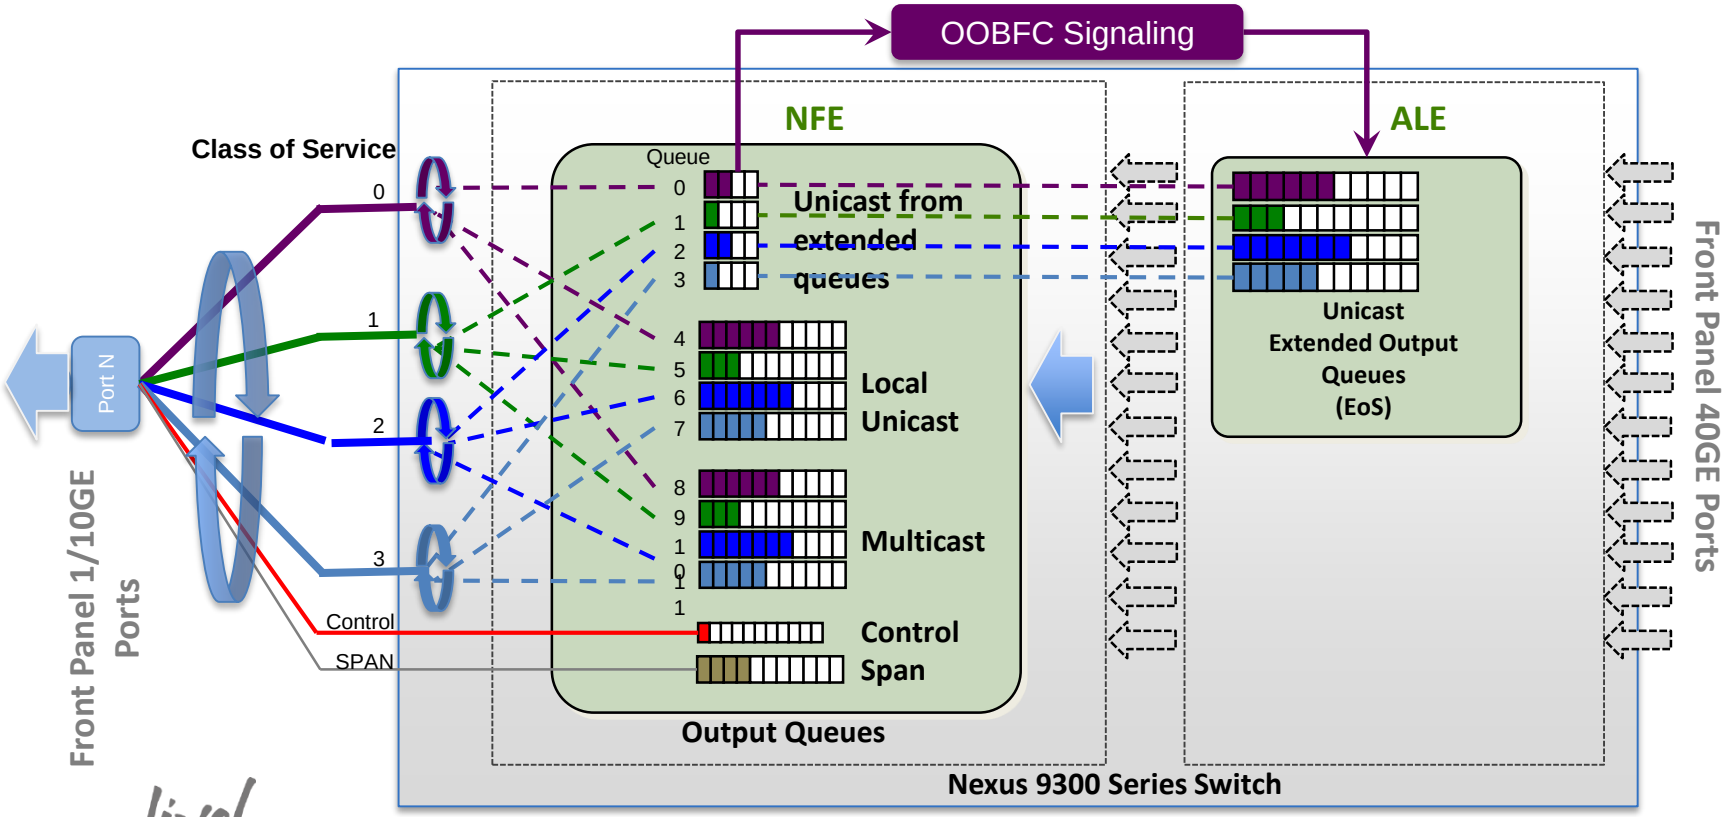

QoS on Nexus 9300 Series Switches

- **Ingress QoS (Classification and Remarking)**
  - Policy-map type qos
  - Match on CoS/ IP Precedence/ DSCP /ACL
  - Set qos-group
  - Remark CoS/ IP Precedence/ DSCP
  - Ingress policing

- **Network-QoS**
  - Policy-map type network-qos
  - Match on qos-group
  - Enable PFC

- **Egress Qos (Queuing and Shaping)**
  - Policy-map type queuing
  - Enable WRED or ECN (default is tail drop)
  - 4 user-defined classes based on qos-group
  - 1 control class for CPU and 1 class for SPAN traffic
  - Up to 3 PFC non-drop queues
  - Up to 3 Strict priority queues

Agenda – Nexus 9000 Architecture

- Nexus 9000
  - Nexus 9000 Hardware
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9300
- Nexus 9000 and 40G
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 & Dev-Ops
- ACI & Nexus 9000

Optical Innovation --- Removing 40 Gb Barriers

Problem

- 40 Gb optics are a significant portion of capital expenditures (CAPEX)
- 40 Gb optics require new cabling

Solution

- Re-use existing 10 Gb MMF cabling infrastructure
- Re-use patch cables (same LC connector)

Cisco® 40 Gb SR-BiDi QSFP

- QSFP, MSA-compliant
- Dual LC connector
- Support for 100 m on OM3 and upto 150m on OM4
- TX/RX on two wavelengths at 20 Gb each

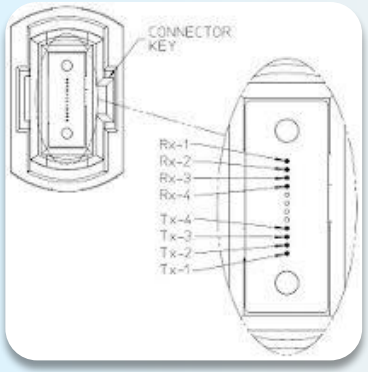

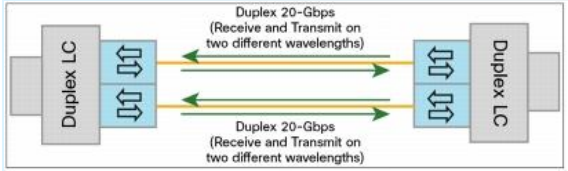

Available end of CY13 and supported across all Cisco QSFP ports

QSFP-BIDI vs. QSFP-40G-SR4

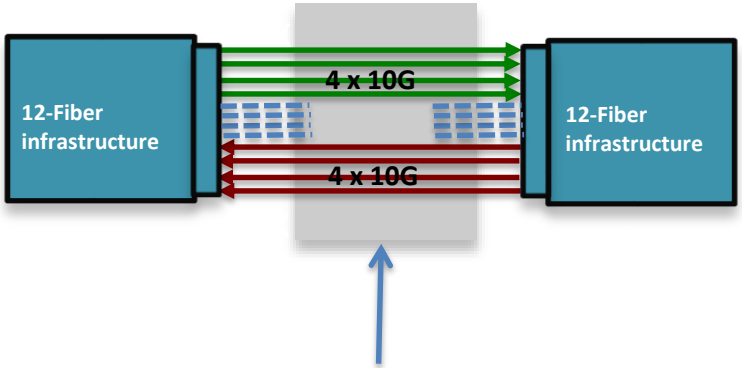

12-Fiber ribbon cable with MPO connectors at both ends

Higher cost to upgrade from 10G to 40G due to 12-Fiber infrastructure

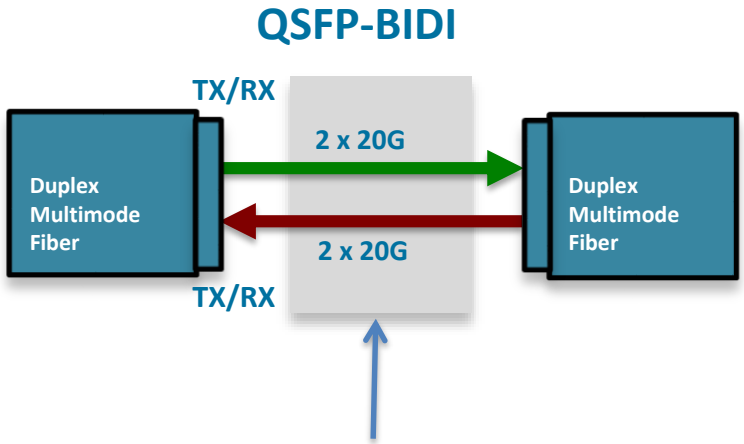

Duplex multimode fiber with Duplex LC connectors at both ends

Use of duplex multimode fiber lowers cost of upgrading from 10G to 40G by leveraging existing 10G multimode infrastructure

Significant Cabling & Transceiver Savings

10G Optical Link

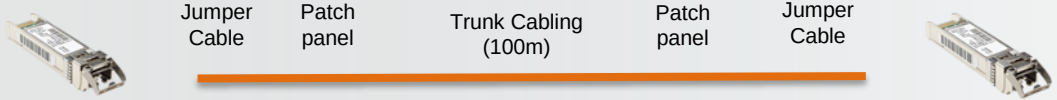

+$6,259*

Traditional 40G Optical Link—Complete Replacement

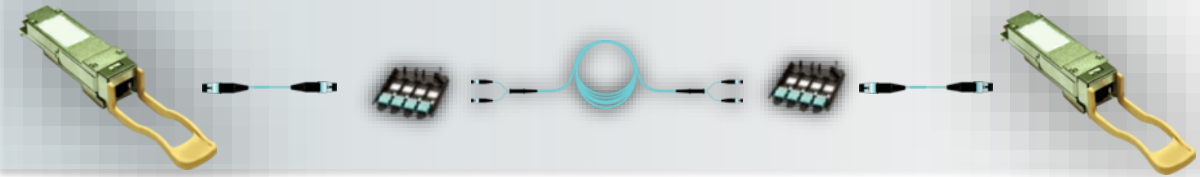

40G BiDi Optical Link—Reuse all 10G Cabling/Patch Panels

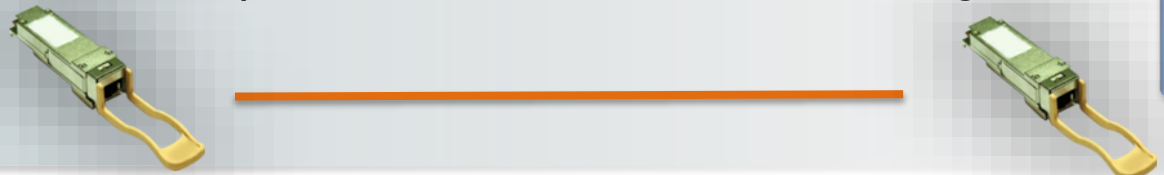

+$2,200*

Cost Avoidance Brownfield

Cost Savings Greenfield

$4,059
SAVINGS (LIST) PER 40G LINK

Source: Corning OM3 Cable & Patch Panel list prices, Cisco 40G BiDi list price, Competitors 40G SR4 list price

Agenda – Nexus 9000 Architecture

- Nexus 9000
  - Nexus 9000 Hardware
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9300
- Nexus 9000 and 40G
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 & Dev-Ops
- ACI & Nexus 9000

Nexus 9500/9300 FEX Support

- Optimized Model for Long Term TCO during Evolution

Cisco Nexus® 5500

Cisco Nexus® 2000 FEX

Cisco Nexus® 9300

Cisco Nexus® 2000 FEX

- Migration of Nexus 5500 to Nexus 9300 provides
- Increased scalability
  - 160K MAC
  - 16K Routes
  - 32K MRoutes
  - 160K IGMP Groups
- Addition of 40G uplinks for lower oversubscription
- Addition of VXLAN Bridging, Gateway and Routing capabilities
- Line Rate Layer 2 and Layer 3
- Reduction of Latency

Cisco live!

Nexus 9500/9300 FEX Support (Cont'ed)

Supported Topology

Supported FEX Topology:
- Single-homed FEX
- vPC port channel to hosts

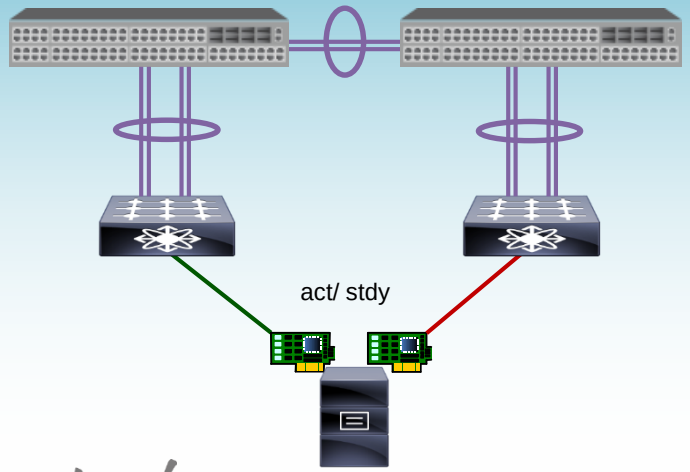

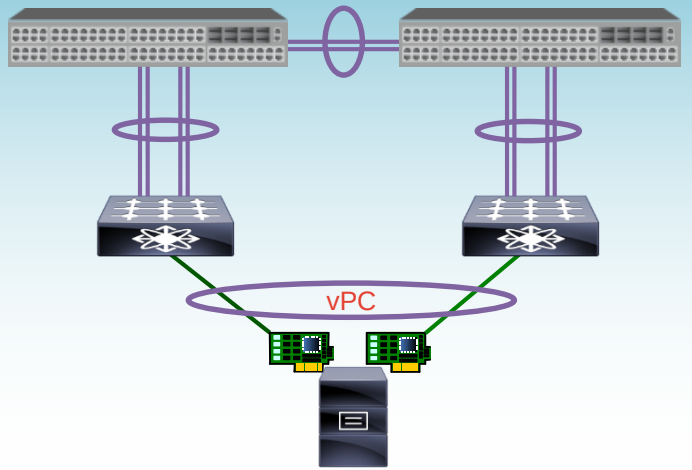

Migration and Interop with Existing Nexus Platforms

- Pod Design Migration with vPC
  - Nexus 9000 is fully compatible with all existing Nexus vPC & FEX designs
  - Nexus 9500 provides a migration path to high performance and high density 40Gbps aggregation

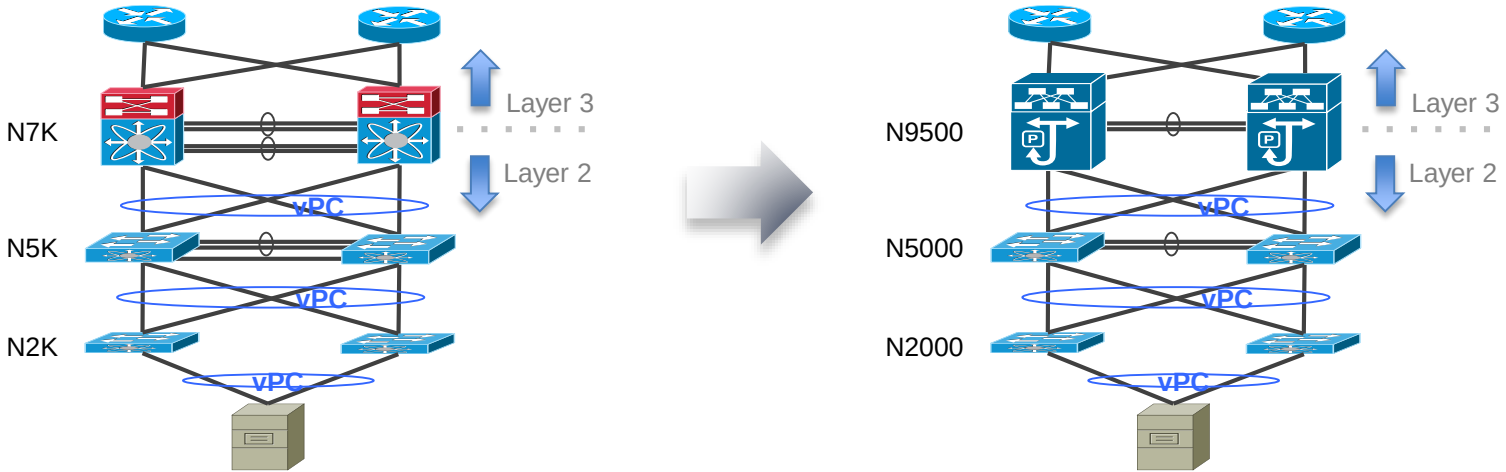

Collapsed Aggregation and Access

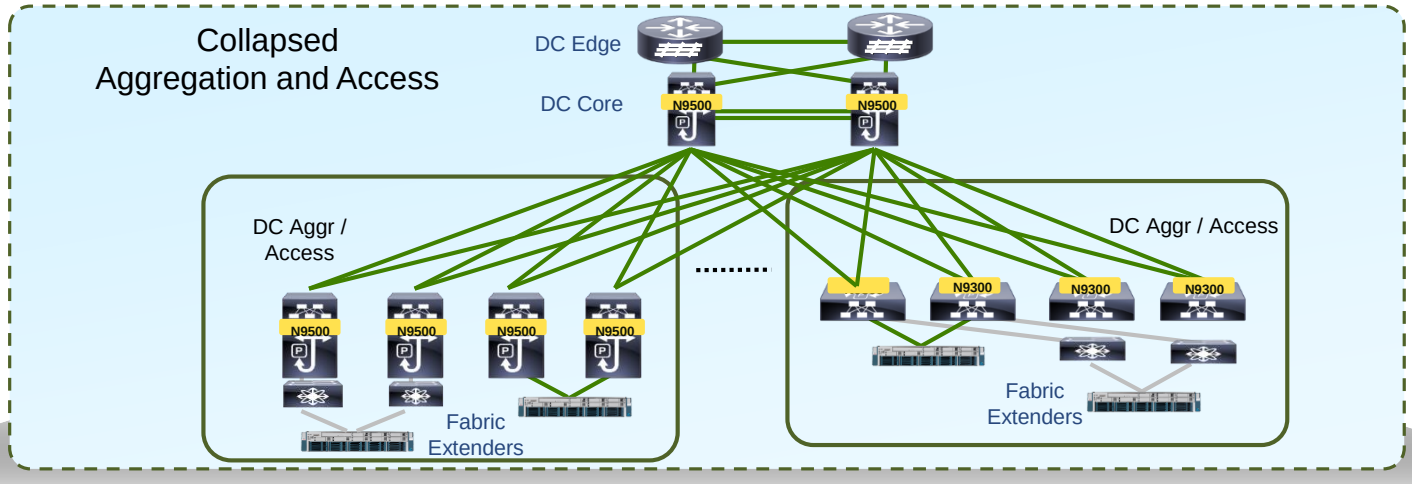

Configuration Options

Option 1:

- Layer 2 and 3 boundary is on the core
- vPC between aggregation and core
- Nexus 9500 can be EoR with FEX in the racks
- Nexus 9300 can be MoR and EoR with FEX in the racks

Option 2:

- Layer 3 all the way to access
- ECMP between access and aggregation
- Nexus 9500 can be EoR with FEX in the racks
- Nexus 9300 can be MoR and EoR with FEX in the racks

Traditional 3-Tier Network Design

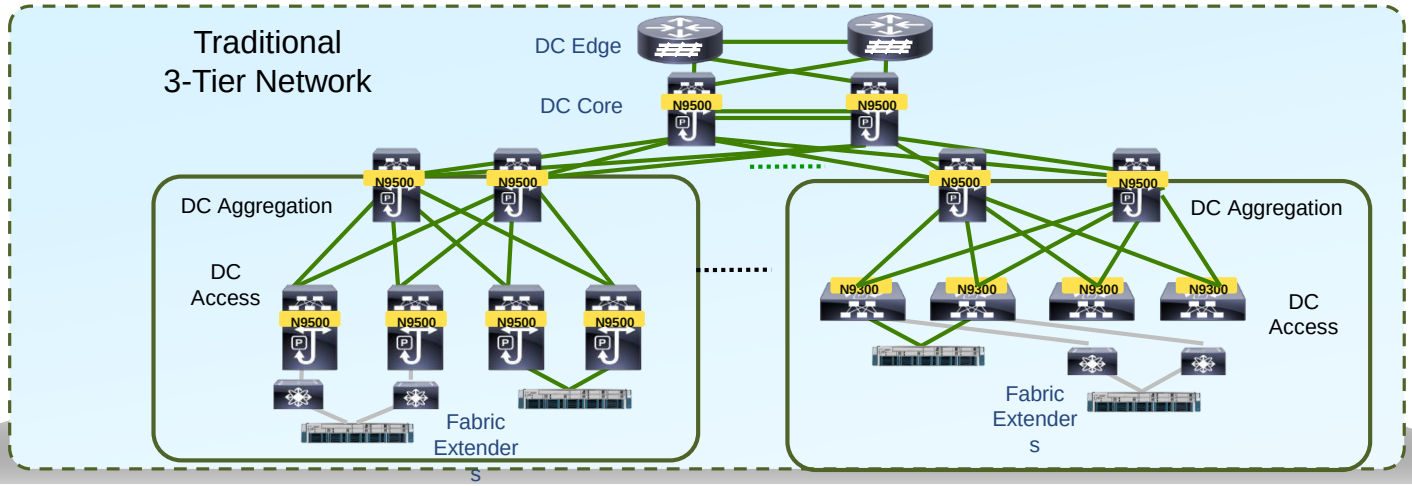

Configuration Options

Option 1:

- Layer 2 or 3 boundary is on aggregation
- vPC between access and aggregation
- Nexus 9500 can be EoR
- Nexus 9300 can be ToR or MoR

Option 2:

- Layer 2 or 3 boundary is on aggregation
- vPC between access and aggregation
- Nexus 9500 can be EoR
- Nexus 9300 can be ToR or MoR

Spine-Leaf Fabric Design

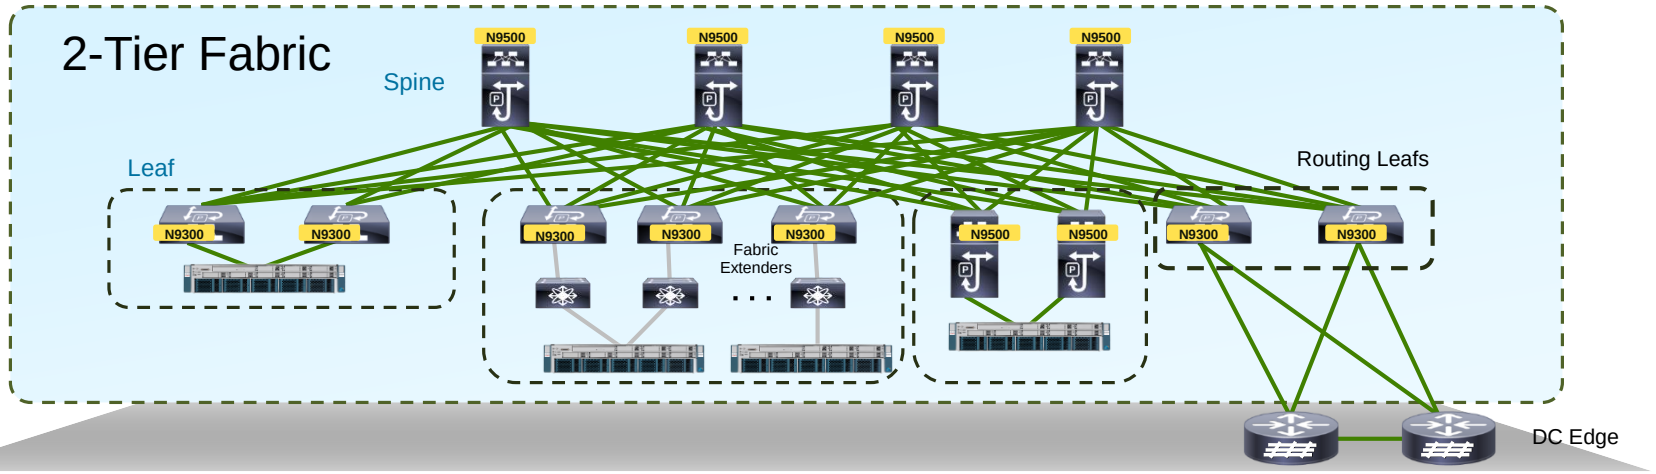

Configuration Options:

- Layer 3 IP fabric
- Layer 3 all the way to access
- Layer 2 extensibility can be provided by VXLAN
- Nexus® 9500 can be EoR
- Nexus 9300 can be ToR and MoR

VXLAN EVPN Enabled Fabric Design for Host-based Forwarding

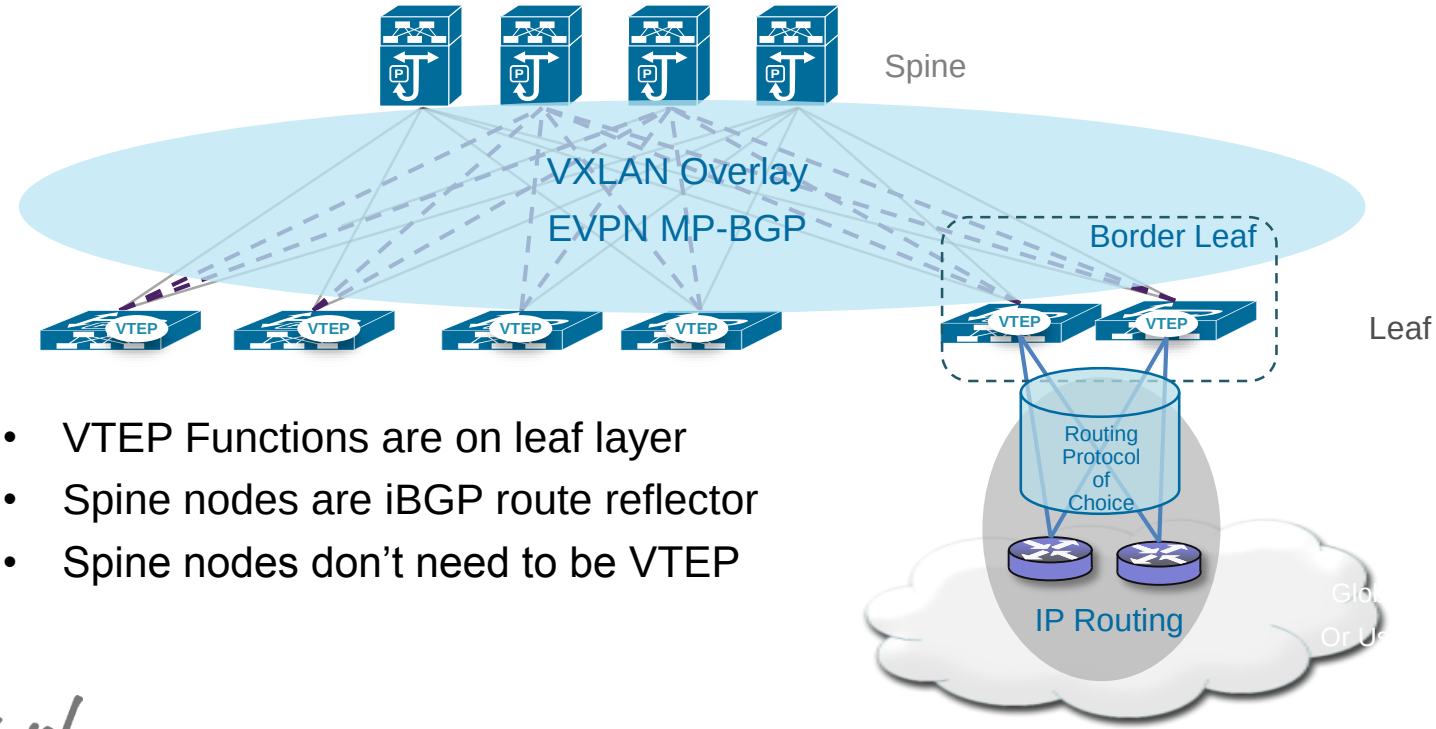

- VTEP Functions are on leaf layer
- Spine nodes are iBGP route reflector
- Spine nodes don’t need to be VTEP

DCI with VXLAN EVPN (Option A)

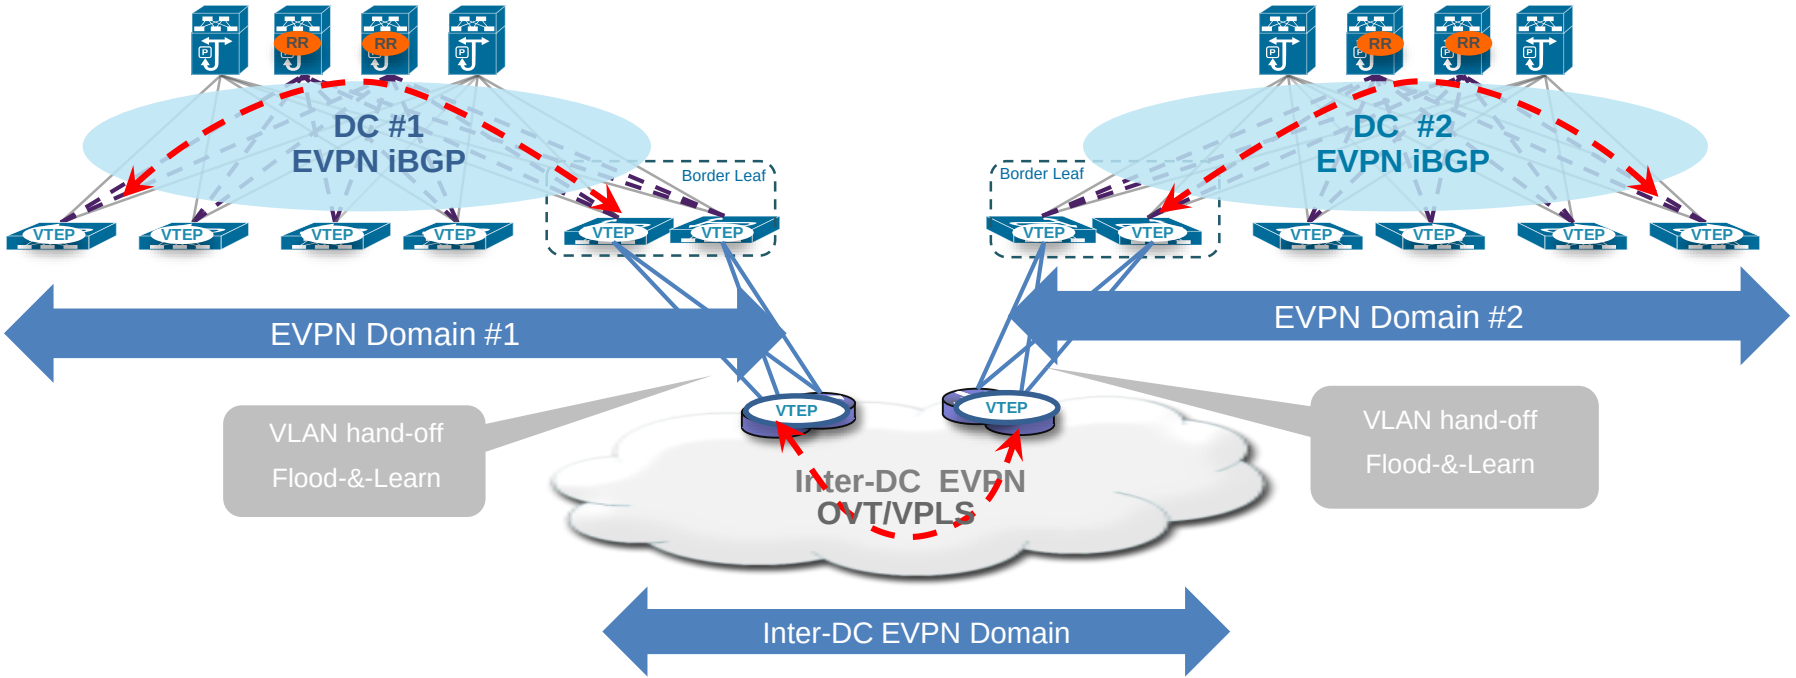

DCI with VXLAN EVPN (Option B)

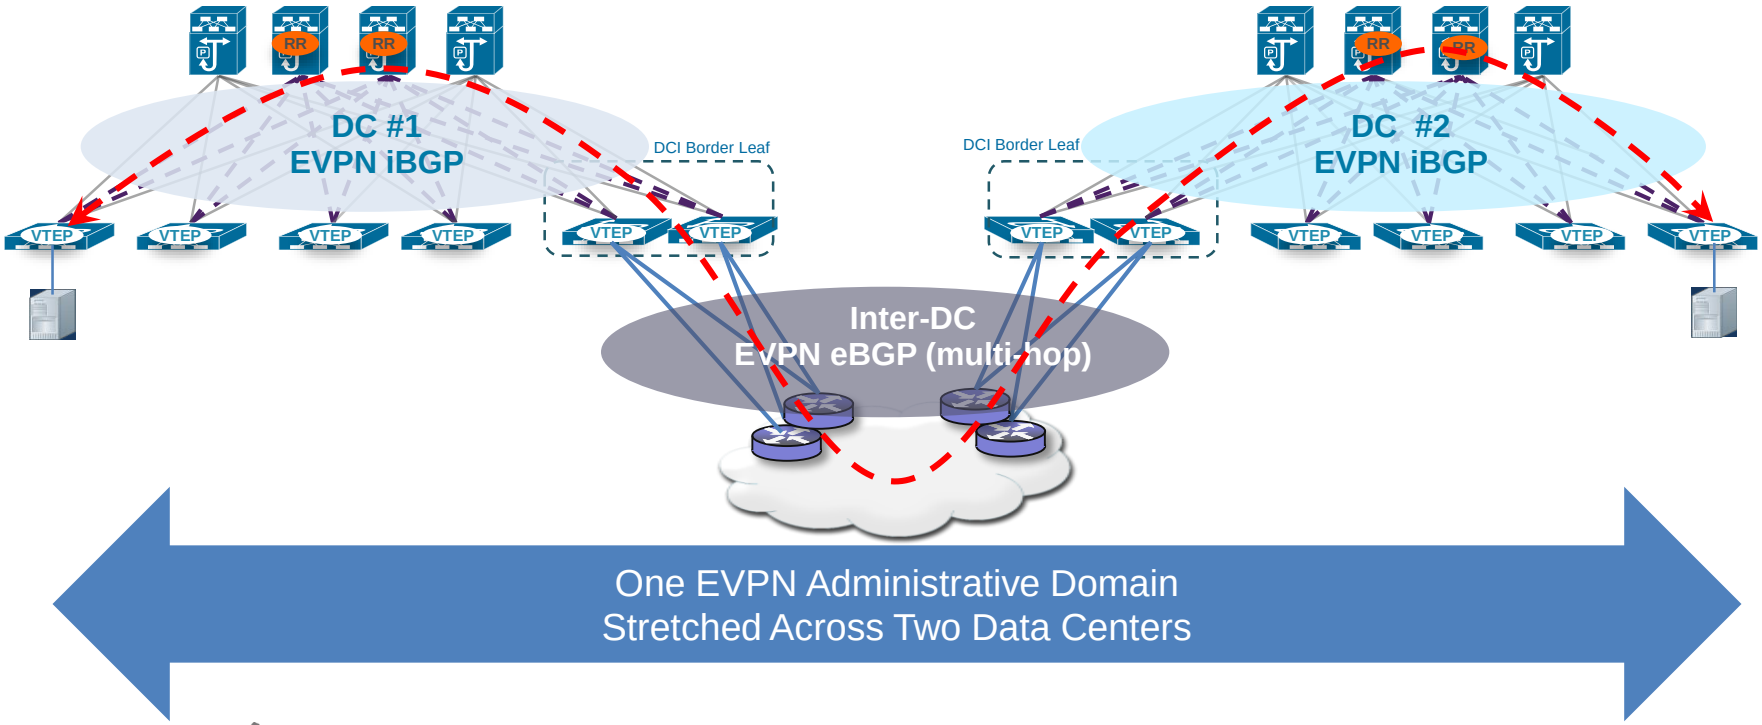

Agenda – Nexus 9000 Architecture

- Nexus 9000
  - Nexus 9000 Hardware
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9300
- Nexus 9000 and 40G
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 & Dev-Ops
- ACI & Nexus 9000

OPEN NX-OS - Extensible, Open, Programmable

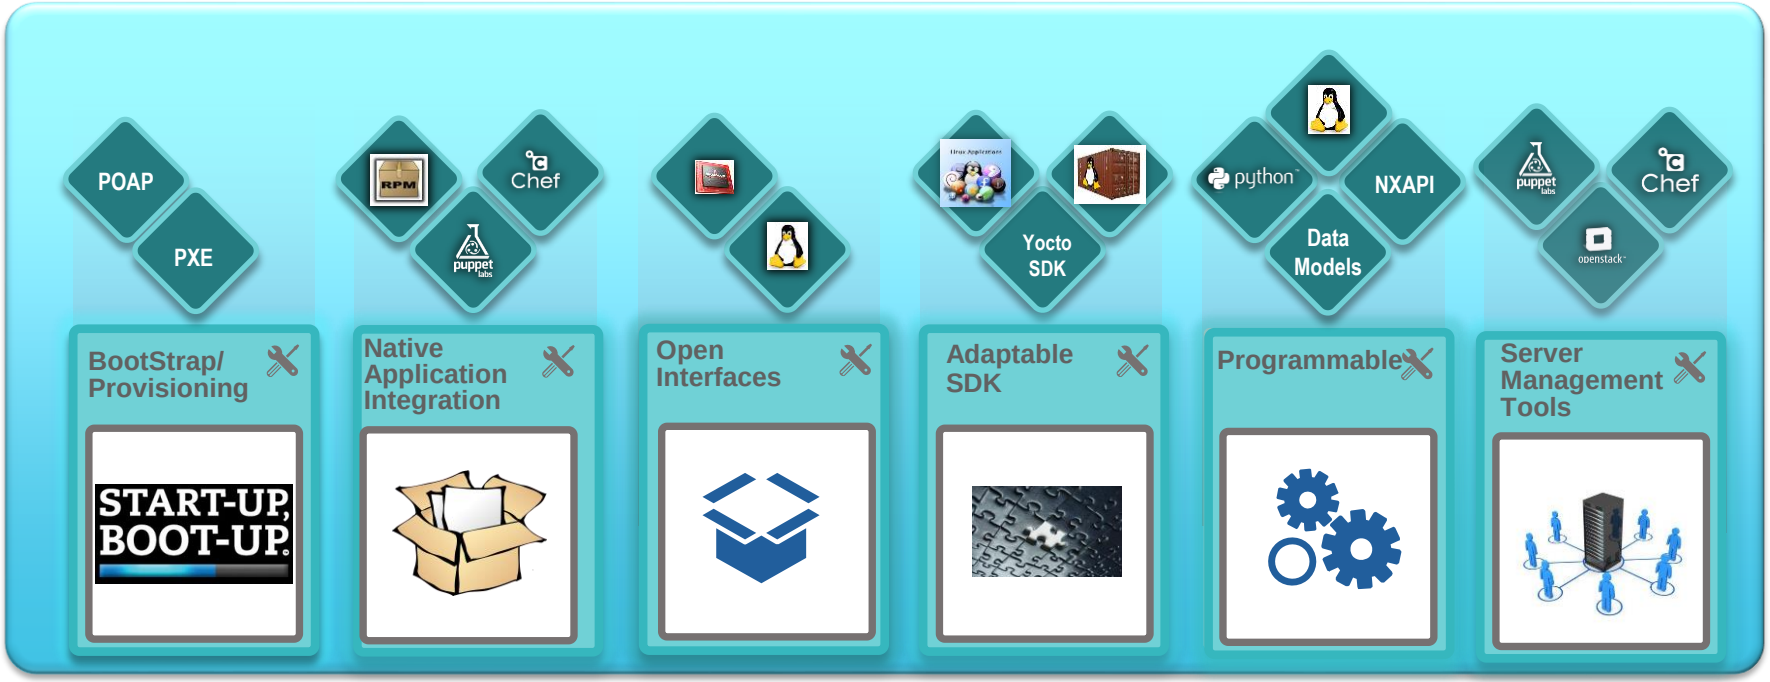

Auto Deployment Options | **Extensibility** | **Open** | **Application Integration** | **Programmability Tool Choice** | **DevOps Enabling**

Open NX-OS: Infrastructure Layer Enhancements

OPEN BOOTLOADERS & PROVISIONING

OPEN PACKAGE/APPLICATION INTEGRATION

OPEN INTERFACES

OPEN OBJECT BASED API's (NX-API, Model Driven)

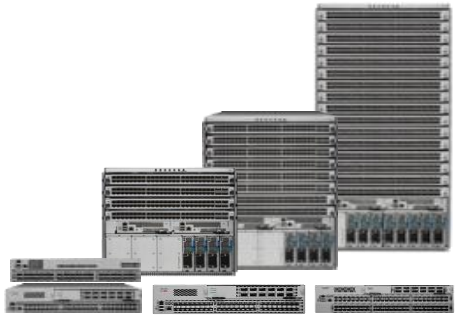

Open NX-OS consistent across both ToR and Modular

Enhanced NX-OS NX-API

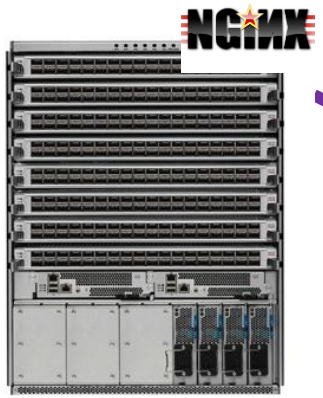

N9K


```
{
  "ins_api": {
    "type": "cli_show",
    "version": "0.1",
    "sid": "session1",
    "outputs": {
      "output": {
        "TABLE_interface": {
          "ROW_interface": [
            {
              "interface": "mgmt0",
              "state": "up",
              "ip_addr": "172.21.128.227",
              "speed": "1000",
              "mtu": "1500"
            },
            {
              "interface": "loopback0",
              "state": "up"
            }
          ]
        }
      }
    }
  }
}
```


CLI Input

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<ins_api>
  <type>cli_conf</type>
  <version>0.1</version>
  <sid>1</sid>
  <input>show interface brief </input>
  <output_format>xml</output_format>
</ins_api>
```

HTTP

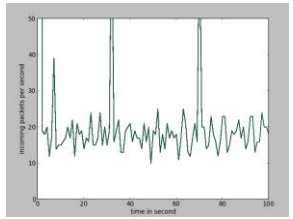

Programmability

Open NX-OS Programmability
Model Driven Automation

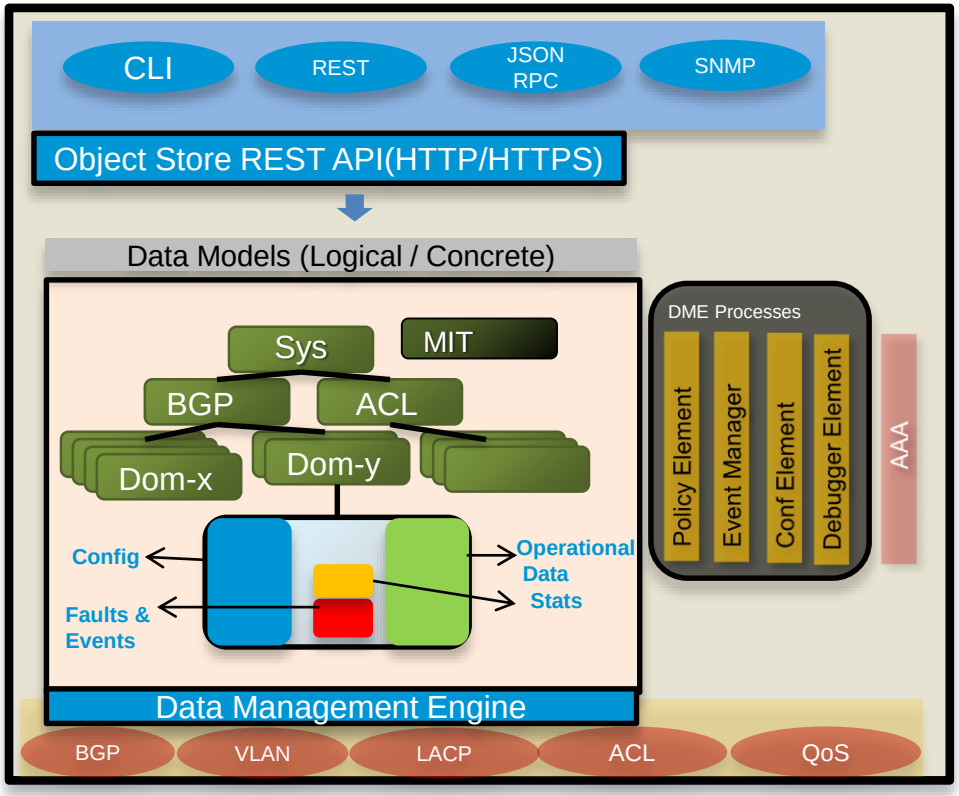

- Automate at scale using REST API access to Nexus object store
- Automate operations leveraging object create/update/delete operations
- Benefits of Model Based automation
  - Software productivity
  - Software maintainability
  - Software quality and code reuse
- Asynchronous model driven programmability:
  - BGP, VLAN, LACP, ACL, QoS, UDLD, CDP, MAC, DHCP, DNS, RBAC, AAA,SVI, Logging, NTP, VRRP
- Object Model Specification available at http://developer.cisco.com
- Push based model for event reporting, leveraging websockets interface.

Object Based Programmability
BGP Configuration

CLI	POST Request (Pre Camden)	POST Request BGP Object (Camden)

router bgp 11
router-id 1.1.1.1
	POST http://Switch-IP/ins
	{'content-type':'application/json-rpc'}.json()
	
	{ "jsonrpc": "2.0",
	"method": "cli",
	"params": {
	"cmd": "config t",
	"version": 1 }, "id": 1},
	{ "jsonrpc": "2.0",
	"method": "**cli**",
	"params": {
	"cmd": "router bgp 11",
	"version": 1 },"id": 1},
	{ "jsonrpc": "2.0",
	"method": "**cli**",
	"params": {
	"cmd": "router-id 1.1.1.1",
	"version": 1}, "id": 2}]
		POST http://Switch-IP/api/mo/sys/bgp/inst.json
		
		{ "bgpInst" : {
		"children" : [{
		"bgpDom" : { 11
		"attributes" : {
		"name":"default",
		"rtrId" : "1.1.1.1"
		}
		}
		}
		]
		}
		}

Open NX-OS Linux Interfaces Bash Access

- Leverage Linux command toolkit for switch monitoring configuration and troubleshooting
  - **# tcpdump -w file.pcap -i eth1-1**
- Use ethtool to display detailed interface statistics:
  - **#ethtool –S eth2-1**
- Use ifconfig to change mtu for an interface to jumbo MTU:
  - **#ifconfig eth2-1 mtu 9000**
- Use ip route to add a static route for a given interface:
  - **#ip route add 203.0.113.0/24 via 198.51.100.2 dev eth2-1**
- Leverage bash for NX-OS scripting automation
  - **vsh –c “show interface brief” | grep up | awk/sed**

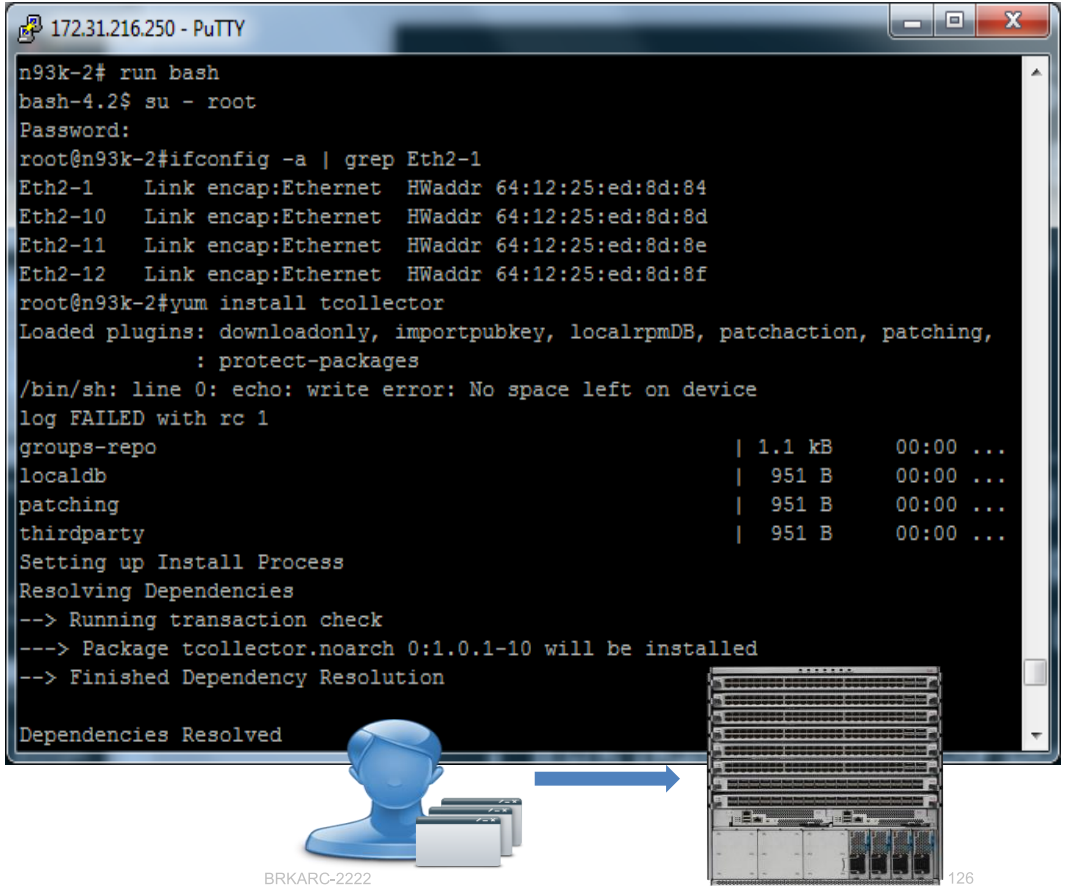

Open NX-OS Linux Interfaces Using Linux Utilities

- Capability to configure native switch linux devices with standard off the shelf linux tools.
- **Lots of tools, lots of options**, some examples below:
- For troubleshooting use tcpdump to capture all packets on a given port, and dump to a file:
  - **# tcpdump -w file.pcap -i eth1**-1
  - tcpdump: listening on eth1, link-type EN10MB (Ethernet), capture size 65535 bytes 4 packets captured 4 packets received by filter 0 packets
- Use ethtool to display detailed interface statistics:
  - **#ethtool –S eth2-1**
- Use ifconfig to change mtu for an interface to jumbo MTU:
  - **#ifconfig eth2-1 mtu 9000**
- Use ip route to add a static route for a given interface:
  - **#ip route add 203.0.113.0/24 via 198.51.100.2 dev eth2-1**
- Leverage bash for NX-OS scripting automation
  - **vsh –c “show interface brief” | grep up | awk/sed**

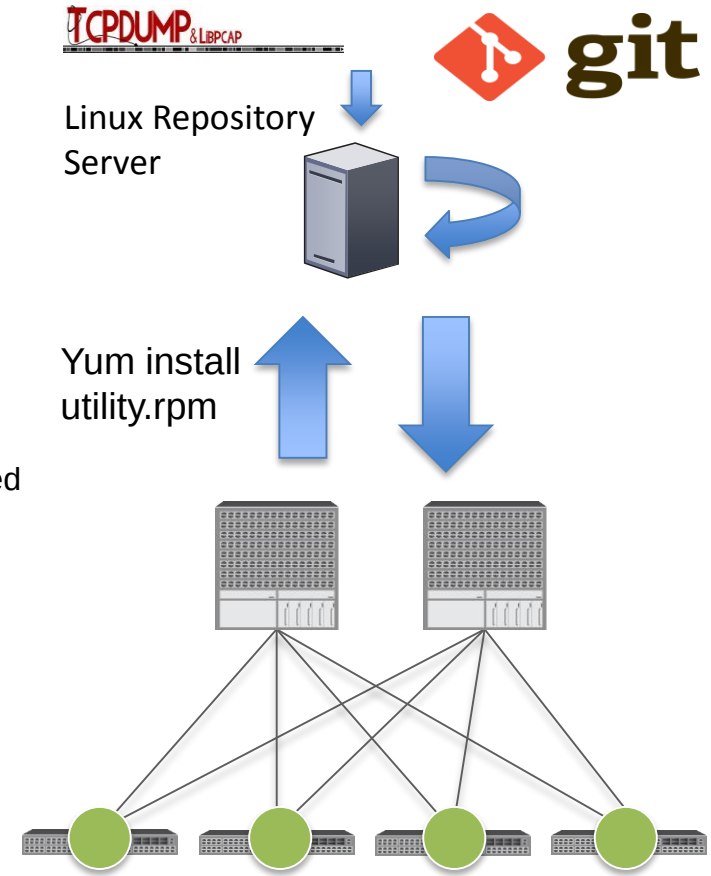

BCM Shell Access

- Issue a CLI to get shell access to underlying BCM chips
- Direct read/write access to hardware tables
- Can Peek/Poke underlying registers
- Python wrapper to get BCM Shell output


```
TME-1-9508-1# bcm-shell module 1
Warning: BCM shell access should be used with caution
Entering bcm shell on module 1
Available Unit Numbers: 0 1 2
bcm-shell.0> ^[[A^[[A^[[B

bcm-shell.0> l3 l3table show
l3 l3table show
Unit 0, free L3 table entries: 212960
Entry VRF IP address       Mac Address           INTF MOD PORT    CLASS HIT
147488 1    192.168.1.2       00:00:00:00:00:00  100006    0    0          0 y
149300 1    30.1.1.255        00:00:00:00:00:00  149150    0    0          0 n
150696 1    30.1.1.1          00:00:00:00:00:00  100012    0    0          0 n
152696 1    10.1.1.3          00:00:00:00:00:00  100007    0    0          0 y
154860 1    192.168.1.15      00:00:00:00:00:00  149150    0    0          0 n
156336 1    192.168.1.0       00:00:00:00:00:00  100000    0    0          0 n      (LOCAL ROUTE)
163452 1    192.168.1.13      00:00:00:00:00:00  149151    0    0          0 y      (LOCAL ROUTE)
165120 1    192.168.1.3       00:00:00:00:00:00  149150    0    0          0 n
166280 1    30.1.1.0          00:00:00:00:00:00  100000    0    0          0 n      (LOCAL ROUTE)
168280 1    10.1.1.2          00:00:00:00:00:00  100006    0    0          0 y
170444 1    192.168.1.14      00:00:00:00:00:00  100010    0    0          0 y
173968 1    192.168.1.1       00:00:00:00:00:00  149151    0    0          0 y      (LOCAL ROUTE)
174872 1    30.1.1.2          00:00:00:00:00:00  149151    0    0          0 n      (LOCAL ROUTE)
179036 1    192.168.1.12      00:00:00:00:00:00  100000    0    0          0 n      (LOCAL ROUTE)
183716 1    192.168.1.11      00:00:00:00:00:00  149150    0    0          0 n
184680 1    192.168.1.6       00:00:00:00:00:00  100007    0    0          0 y
186876 1    10.1.1.10         00:00:00:00:00:00  149151    0    0          0 n      (LOCAL ROUTE)
192308 1    192.168.1.9       00:00:00:00:00:00  149151    0    0          0 n      (LOCAL ROUTE)
193012 1    10.10.10.10       00:00:00:00:00:00  149151    0    0          0 y      (LOCAL ROUTE)
193528 1    192.168.1.4       00:00:00:00:00:00  100000    0    0          0 n      (LOCAL ROUTE)
201800 1    192.168.1.7       00:00:00:00:00:00  149150    0    0          0 n
```

Linux Containers

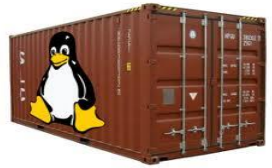

- Enable Nexus switches to host customer applications using LXC virtualization technologies
- Customers can use ‘Application Development Toolkit’ to compile & package their applications into Cisco OVAs for deployment on Nexus switches

LXC

This is an operating system virtualization technology that shares the host kernel with the guest but provides isolation through namespace extensions to the Linux kernel.
http://linuxcontainers.org/

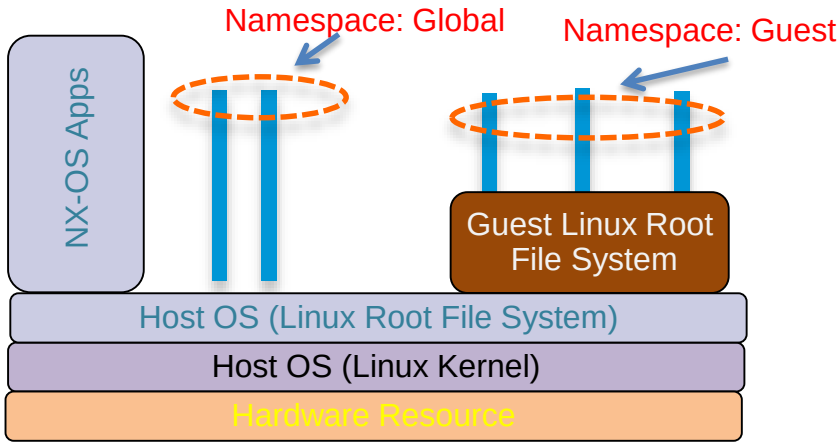

FYI

LinuX Containers (LXC) on Nexus 9000/3000

- Provides a secure and segregated operating environment for applications
- Can run either Cisco or Open Source applications
- Can use standard Linux distros
- OS Level Virtualization
- Shared Kernel
- Shared physical resources
- Isolation through name spaces

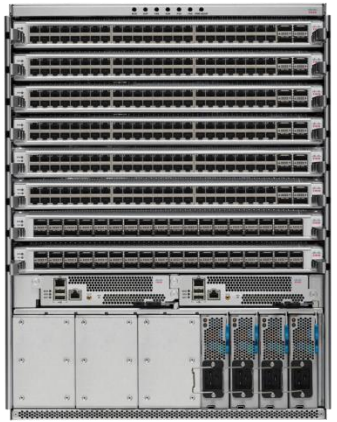

Open NX-OS Puppet/Chef

- Cisco Puppet Agent RPM/software package posted to Puppetforge and Open Sourced to Github
- Install Cisco Puppet Module on Puppet Master
- Yum install Puppet Agent rpm on switches
- Switch Agent periodically will poll Puppet/Chef Master for updated catalog/cookbooks and attempt to converge switch to desired state

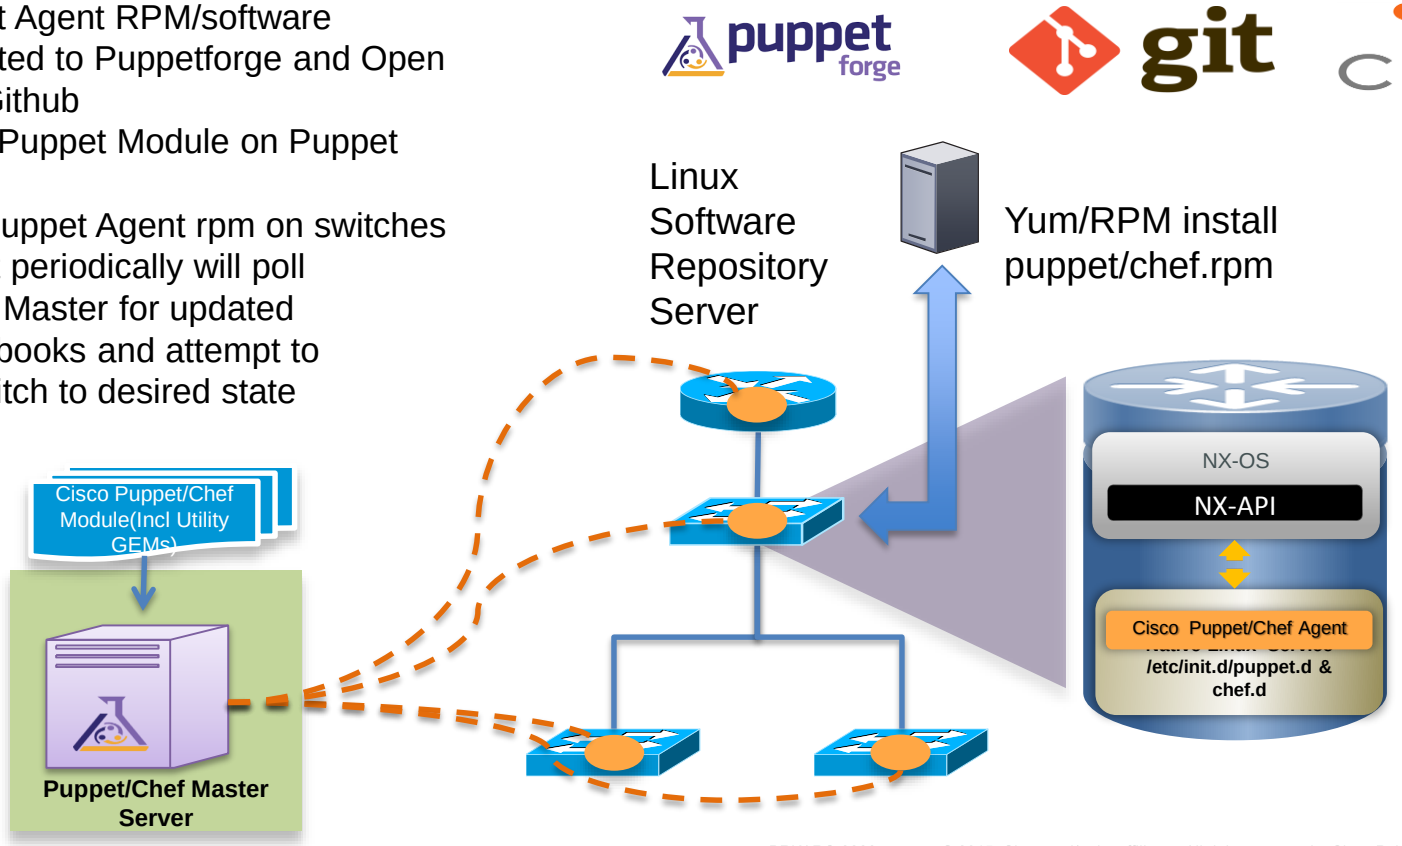

Cisco Puppet Resource Type Coverage: 1

Feature	Resource Name	Description
Cisco Device Access	cisco_device	Allows credentials for user access control & accounting
Base L2/L3 interface	cisco_interface	General interface & L2/L3 base settings
VLAN	cisco_vlan	Create/destroy of VLANs and general settings
Interface-vlan (SVI)	cisco_interface_vlan	Create/destroy of SVIs and SVI specific interface settings
VLAN Trunking Proto (VTP)	cisco_vtp	VTP global settings
SNMP	cisco_snmp_server cisco_snmp_community cisco_snmp_group cisco_snmp_user	SNMP monitoring settings. Notification receiver settings not covered as of now.
OSPF	cisco_ospf cisco_ospf_vrf cisco_interface_ospf	OSPF instance create/destroy, per-VRF settings, and interface settings (area, cost, msg digest, etc)

OpenStack Network (Neutron) Plugin

- Enables fully automated compute, storage and network resource orchestration
- Support for Grizzly OpenStack release
- Enable VLAN-based tenant separation
- Enhance efficient resource usage
- Leverages NX-OS NetConf-XML programmatic interface

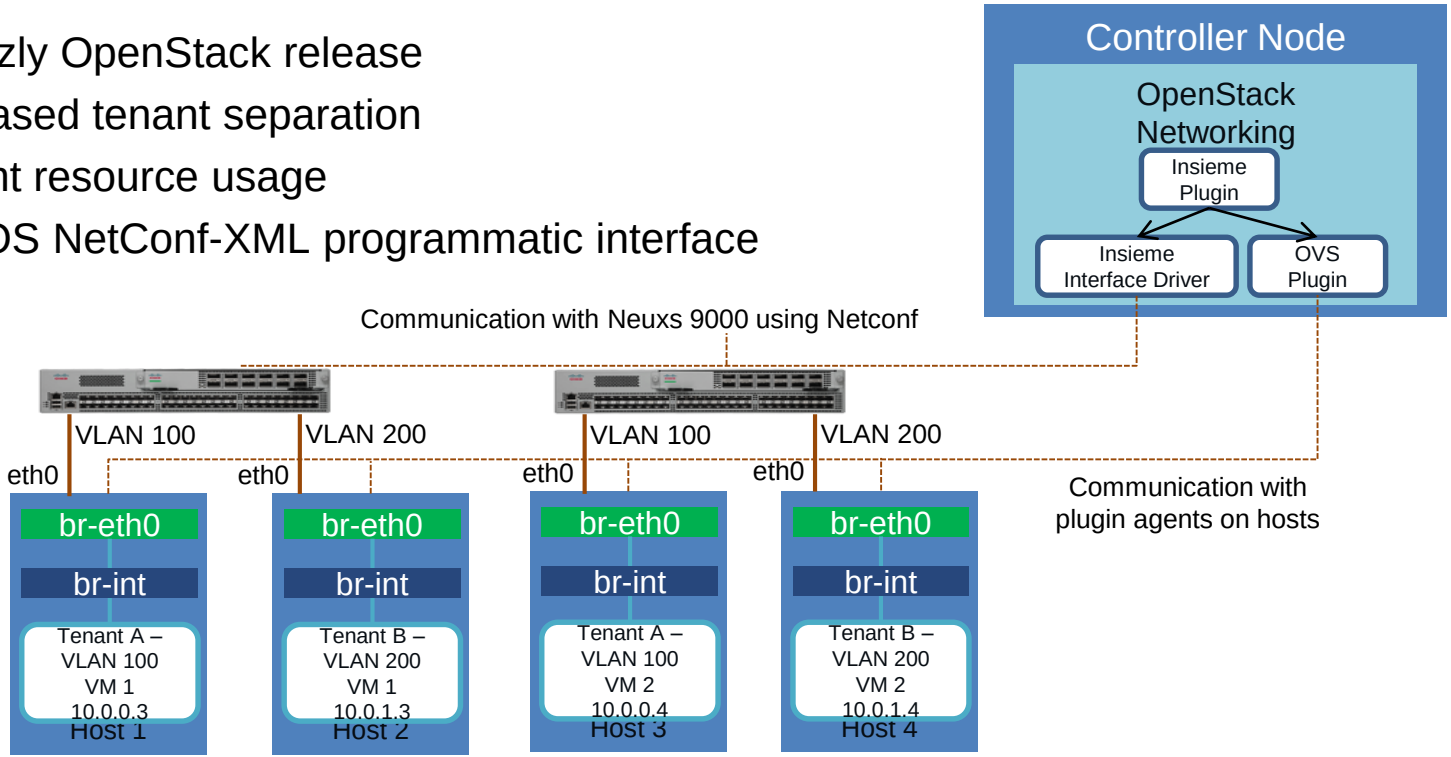

Agenda – Nexus 9000 Architecture

- Nexus 9000
  - Nexus 9000 Hardware
  - Nexus 9500 Chassis
  - Nexus 9500 Line Cards
  - Nexus 9500 Packet Forwarding
  - Nexus 9300
- Nexus 9000 and 40G
- Nexus 9000 Designs: FEX, vPC & VXLAN
- Nexus 9000 & Dev-Ops
- ACI & Nexus 9000

Common Platform: Two Modes of Operation

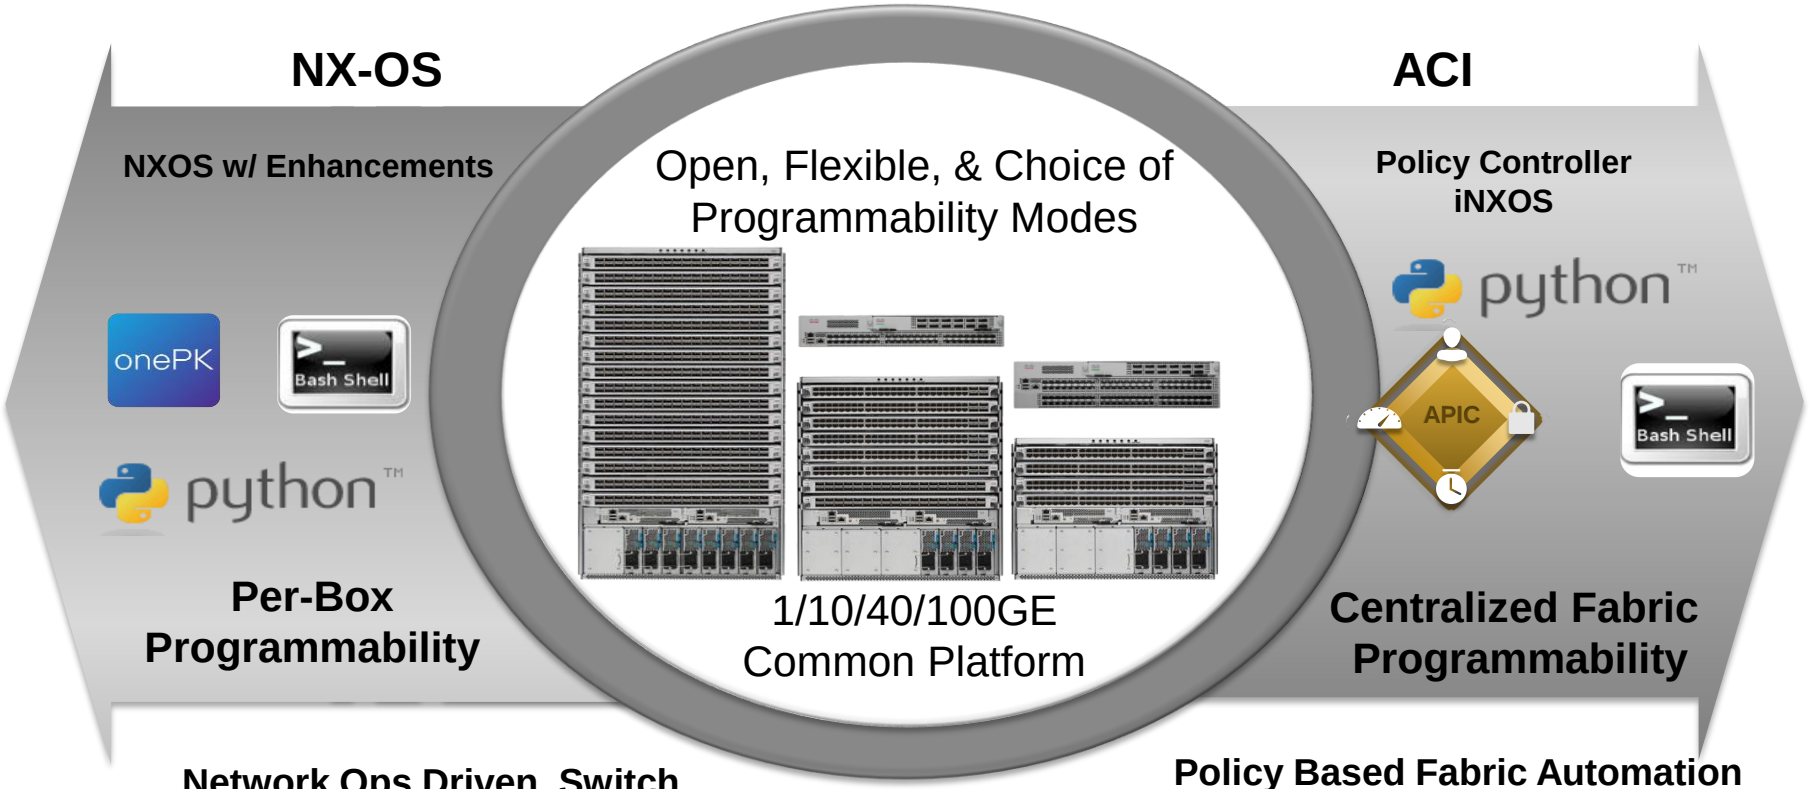

Cisco ACI

Logical Network Provisioning of Stateless Hardware

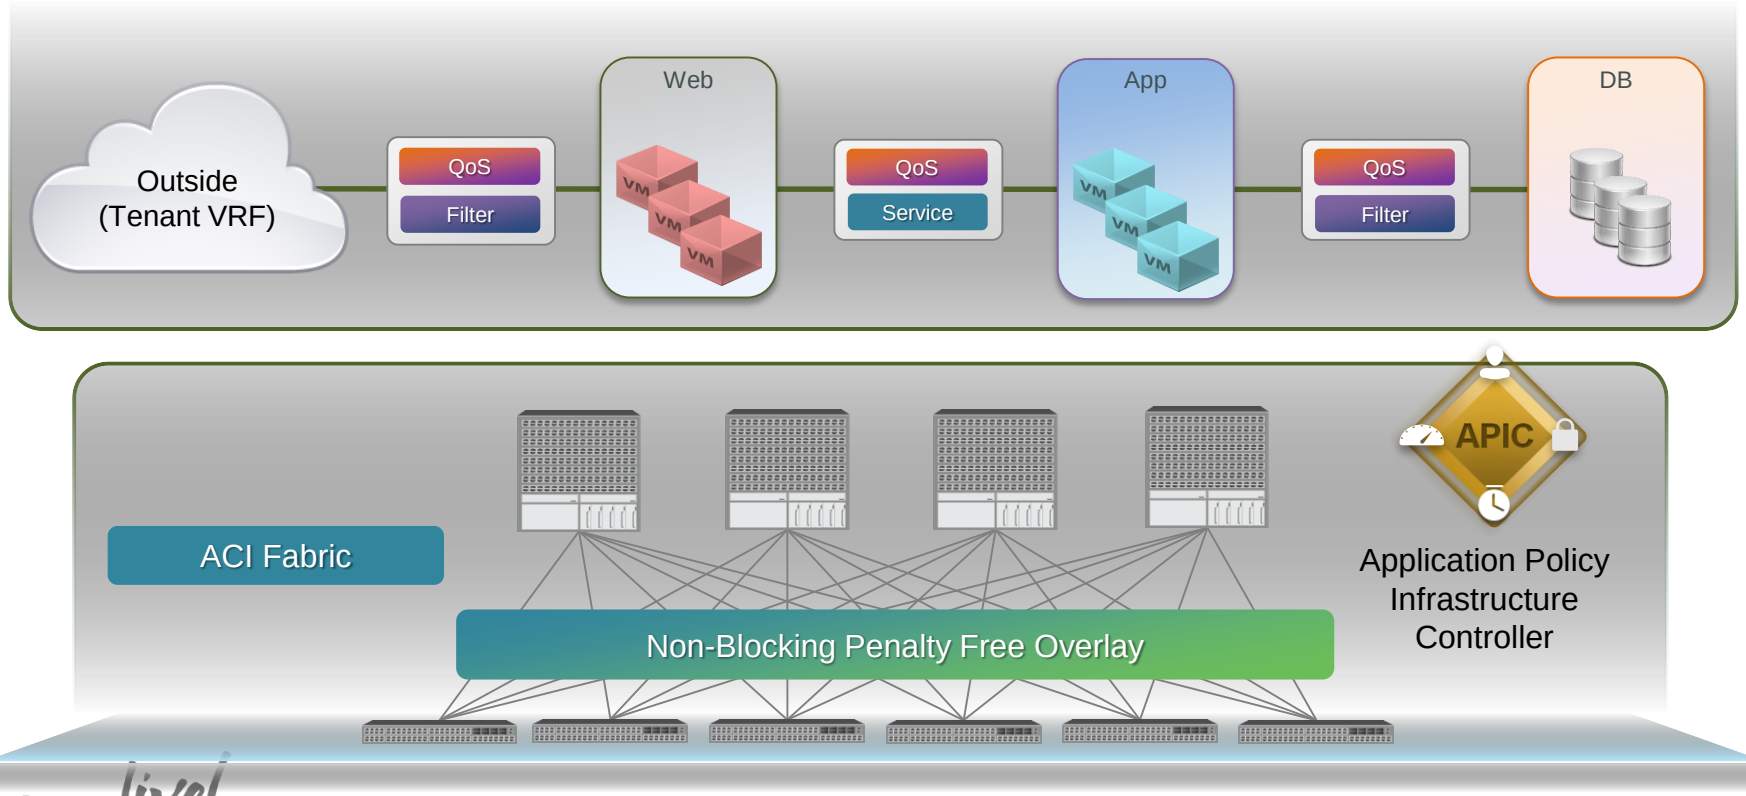

Foundational Switching Platforms for the Next Decade

Nexus 9000

1/10/40/100G

Standalone / ACI Ready

Industry Leading Price/Performance, Port Density: **Fastest 10G/40G /100G Platform with Merchant+**

15% Better Power & Cooling–**2.8X** Better Reliability

Programmability/ Open APIs: Linux Containers, Python, Power Shell, Puppet, Chef… **Ideal for DevOps!!**

Innovations: Object Model, No Backplane, No Midplane, Health scores

$ Multi-million Savings 40/100G on Existing Cables using BiDi Optics. **Non disruptive migration to 40GE**

Where You Can Learn More?

Nexus Platforms and NX-OS Related Sessions

Session ID	Session Title	Speaker	Day / Time
BRKDCT-2458	Nexus 9000/7000/6000/5000 Operations and Maintenance Best Practices	Arvind Durai Nicholas Garner	Tuesday, Jun 9, 8:00 AM - 10:00 AM
BRKDCT-3101	Nexus 9000 (Standalone) Architecture Brief and Troubleshooting	Shridhar Dhodapkar Sukumar Subburayan	Monday, Jun 8, 1:00 PM - 3:00 PM
CCSACI-2552	The Journey to Nexus 9k and ACI: NetApp Global Engineering Cloud	Michael McKee	Thursday, Jun 11, 8:30 AM - 9:30 AM
BRKAPP-2027	Big Data Architecture and Deployment	Sean McKeown	Tuesday, Jun 9, 3:30 PM - 5:00 PM
BRKCOC-1000	Inside Cisco IT: Secure and Simplified Cloud Services with ACI	Benny Van De Voorde Erich Latchford	Wednesday, Jun 10, 1:00 PM - 3:00 PM
PSODCT-1013	Building flexible 40/100G Fabrics with Nexus Switching Portfolio	Anant Shah	Thursday, Jun 11, 11:30 AM - 12:30 PM

Where You Can Learn More?

ACI Related Sessions

Session ID	Session Title	Speaker	Day / Time
BRKACI-1024	Dev-Ops and the Application Centric Infrastructure - Open Standards and Open API's	Mike COHEN	Thursday, Jun 11, 1:00 PM - 2:30 PM
BRKACI-1102	L4-L7 Services Insertion Design and Deployment	Azeem Suleman	Tuesday, Jun 9, 3:30 PM - 5:00 PM
BRKACI-1502	Simplify Operations with ACI	Azeem Suleman Igino Cafiero	Wednesday, Jun 10, 3:30 PM - 5:00 PM
BRKACI-1789	How to Perform Common Tasks in ACI	Jody Davis	Monday, Jun 8, 8:00 AM - 9:30 AM
BRKACI-2001	Integration and Interoperation of Existing Nexus Networks into an ACI Architecture	Mike Herbert	Tuesday, Jun 9, 3:30 PM - 5:00 PM
BRKACI-2006	Integration of Hypervisors and L4-7 Services into an ACI Fabric	Maurizio Portolani Bradley Wong	Thursday, Jun 11, 8:00 AM - 9:30 AM
BRKACI-2101	Understanding, Verifying, and Troubleshooting ACI Configuration Policies	Daniel Pita	Tuesday, Jun 9, 1:00 PM - 3:00 PM
BRKACI-2102	ACI Troubleshooting	Mioljub Jovanovic	Thursday, Jun 11, 1:00 PM - 2:30 PM
BRKACI-2410	Management and Automation of Application Centric Infrastructure (ACI) with Cisco UCS Director	Shankar Varanasy Michael Zimmerman	Thursday, Jun 11, 8:00 AM - 9:30 AM
BRKACI-2501	Operationalize ACI	Jerry Ye	Thursday, Jun 11, 1:00 PM - 2:30 PM
BRKACI-3204	Automating Operational Tasks in Cisco ACI	Paul Lesiak	Monday, Jun 8, 1:00 PM - 3:00 PM
BRKACI-3456	Mastering OpenStack and ACI	Lucien Avramov	Monday, Jun 8, 10:00 AM - 12:00 PM
BRKACI-3503	Extending ACI to Multiple Sites - Dual Site Deployment Deep Dive	Santiago Freitas	Wednesday, Jun 10, 3:30 PM - 5:00 PM
LTRACI-1210	Evolve your Datacenter with ACI	Erum Frahim, Tim Garner Liyuan Quan	Tuesday, Jun 9, 1:00 PM - 5:00 PM
LTRACI-2202	Deploying and Implementing ACI	Luis Flores Cesar Obediente	Wednesday, Jun 10, 8:00 AM - 12:00 PM
LTRACI-2225	OpenStack Installation, Integration and Deployment with Cisco ACI using Group Based Policy	Luis Flores, Rafael Muller Cesar Obediente	Monday, Jun 8, 1:00 PM - 5:00 PM
TECACI-2009	Application Centric Infrastructure (ACI) - The Policy Driven Data Center	Maurizio Portolani Carlos Pereira Max Ardica, Rolf Schaerer	Sunday, Jun 7, 8:00 AM - 5:00 PM
TECACI-2010	ACI Day 2 Operations & Management	Daniel Pita, Rafael Muller Mike Brown,Rob Burns	Thursday, Jun 11, 8:00 AM - 12:00 PM
TECACI-2101	ACI Concepts, Profiles and Constructs for the Next Generation Data Center	Rafael Muller, Daniel Pita Mike Brown	Sunday, Jun 7, 1:00 PM - 5:00 PM

Call to Action

- Visit the World of Solutions for:
  - Cisco Campus – Building Overlay Fabric with VXLAN /Programmability /ACI demos
  - Walk in Labs – LABWISP-1000/2000/3000/4000 (Monday~ Thursday, 8:00am~4:45pm)
  - Technical Solution Clinics
- Meet the Engineer – Wednesday & Thursday @ 10:00 – 11:00
- Lunch time Table Topics – Monday, Tuesday & Wednesday, 11:30~12:3

Participate in the “My Favorite Speaker” Contest

- Promote Your Favorite Speaker and You Could be a Winner
- Promote your favorite speaker through Twitter and you could win $200 of Cisco Press products (@CiscoPress)
- Send a tweet and include
  - Your favorite speaker’s Twitter handle <@lilianquan>
  - Two hashtags: #CLUS #MyFavoriteSpeaker
- You can submit an entry for more than one of your “favorite” speakers
- Don’t forget to follow @CiscoLive and @CiscoPress
- View the official rules at http://bit.ly/CLUSwin

Thank you

TOMORROW starts here.