

Dubrovnik, Croatia, South East Europe
20-22 May, 2013

Nexus 5500 and 6000

Max Ardica
Technical Leader



Cisco's Data Center Portfolio

Leading with Innovation



Nexus 5K




Nexus 2K



Nexus 6K



Nexus 3K



Unified Fabric



Nexus 1K



Blade Offerings



MDS 9000



Nexus 7K

Financial



Low Latency

MSDC/Cloud



Scalability

Service Provider



Agile Service Delivery

Enterprise



Private Cloud
Automation

Cisco Nexus 5500 and 2000 Update



Cisco Nexus 5500

Offering Choice and Flexibility for Varied Deployments



NEXUS 5500 Platform

Nexus 5548UP 1RU Switch (Unified Ports)

Nexus 5596UP 2RU Switch (Unified Ports)



Nexus 5548 Layer 3 Daughter Card

New Version Now Available



Nexus 5596 Layer 3 Expansion Module

New Version Now Available




**Nexus 5500 Unified Ports
Expansion Module**

Cisco Nexus 5500

Nexus 5596T

Nexus 5596T Switch



Flexible

- 10GBASE-T Ecosystem (Intel, Panduit, Commscope)
- Support existing Fiber GEM and new Copper GEM

Scalable

- 96 Access Ports for Direct Server Attach
- Up to 68 1/10GBase-T Ports per chassis
- FP+ Ports supporting FEX Technology

Functionality

- Hardware (1/10GBASE-T Ports)
- Supports FCoE
- Software leverages NX-OS L2/L3 feature set

Industry First for FCoE over 10G BASE-T

Cisco Nexus 5500

4p QSFP+ GEM



4 ports QSFP+*
Each QSFP+ can support 4x 10G
ports



4 ports QSFP+ GEM

- ✓ Supported on all Nexus 5500 Series Chassis
- ✓ Only 4x10G mode supported


FEX Architecture

Simplified Operations and Lower cost

Scalability

Manageability

- Simultaneous support of EoR, MoR and ToR
- Reduction in the number of management points for ToR architectures
- Support for Rack and Blade server connectivity



Managing up to 1152 Server ports with 1RU switch

Cisco Nexus 2000 Series

Simplified Operations with Revolutionary Scale

Indicates Lead Product

2013
2012
2011
2010
2009

100M/1G FEX

1/10G FEX

Blade FEX



N2K-C2248TP-E-1GE
Expanded Memory FEX



N2K-C2248PQ
48 port 1/10G FEX SFP+
4xQSFP



N2K-C2224TP-1GE



N2K-C2232TM-E-10GE
RJ45 downlinks



N2K-C2232TM-10GE
RJ45 downlinks



N2K-C2248TP-1GE



N2K-C2232PP-10GE
SFP+ downlinks

100M/1G FEX



B22 HP FEX
1/10G FEX for
HP Blade servers




B22F FEX
10G FEX for
Fijitsu Blade servers



B22 Dell FEX
1/10G FEX for
Dell Blade servers

Nexus 2000 Series Fabric Extenders

Nexus 2248PQ-10GE



48 ports 10/1G Fabric Extender

- 48x 1/10GE SFP+ host interfaces
- 4x QSFP+ (16x 10GE) on network interfaces
- Front-to-back airflow and back-to-front airflow
- N5500 and N6000 parent switch at FCS
- Feature parity with Nexus 2232PP
- Additional uplink buffers (2x16MB)
- FCoE supported

Design Scenario

- High Density 10GE SFP+ ToR
- Connectivity Flexibility
- Virtualized Environments
- Storage consolidation
- Predictable low latency

Nexus 2000 Series Fabric Extenders

Nexus 2232TM-E



32 1G/10GT Downlinks

- ✓ Supports 1/10GBASE-T on all host ports
- ✓ Supports 10GBASE-T standard (cabling to 100m)
- ✓ Consistent FEX architecture benefits
- ✓ A newer and better 10GBASE-T PHY than 2232TM (40nm):
 - ✓ Better BER characteristics than 2232TM: Supports FCoE
 - ✓ Reduction in power consumption

Technology	BER	Distance	PHY Power (each side)	Transceiver Latency
SFP+ CU Copper	$\sim 10^{-18}$	10m	$\sim 0.1W$	$\sim 0.25\mu s$
SFP+ SR short reach	$\sim 10^{-18}$	300m	1W	$\sim 0.1\mu s$
SFP+ LR long reach	$\sim 10^{-18}$	10km	1W	$\sim 0.1\mu s$
10GBASE-T – 40nm	$\sim 10^{-15}$	100m	$\sim 3-5w$	$\sim 2.5-3\mu s$ $\sim 2.5-3\mu s$

Industry First for FCoE over 10G BASE-T

Cisco Nexus 6004/6001 Overview



Introducing Cisco Nexus 6004



High Performance

- Line rate L2 and L3 with all ports and all features and all frame sizes
- 1-microsecond port-to-port latency with all frame sizes
- 40-Gbps flow
- 40-Gbps FCoE
- Cut-through switching for 40 and 10 GE
- 25-MB buffer per three QSFP interfaces

High Scalability

- 96 x 40G in 4RU
- 384x10 GE in 4RU
- Up to 256,000 MAC (115k L2 / 64k L3 at FCS)
- Up to 128,000 ARP (64k at FCS)
- 32,000 LPM (24k at FCS)
- 16,000 bridge domains
- 31 Bidirectional SPAN sessions (16 at FCS)

Feature-Rich

- L2 and L3 features
- FEXlink
- vPC FabricPath TRILL
- FabricPath with segment ID
- Adapter-FEX/VM-FEX

Visibility and Analytics


- Line-rate SPAN
- Sampled NetFlow*
- Micro-burst and buffer monitoring*
- Latency monitoring*
- Conditional SPAN-SPAN on drop-SPAN on higher latency*

* Some features in this slide are not supported by software at FCS

Nexus 6004 Chassis

Port-Side View

- Chassis depth: 30 in.
- Chassis width: 17.5 in.



Nexus 6004 Chassis

Power Supply and Fans

- The chassis has six power supply slots; a minimum of three is required. They support both 3 + 1 and 3 + 3 redundancy.
- Each power supply is rated 1100W, 90—240 VAC.
- The chassis has four fan trays. A minimum of three is required.



Power Supply 3 + 3
Grid Redundancy
or 3 + 1 Redundancy

Console Mgmt0 USB

Fan Module 3+1
Redundancy

Nexus 6004 Chassis


Twelve QSFP Line-Card Expansion Modules

Provide 12 QSFP interfaces

Support 40, 10GE, FCoE

Offer same performance, features,
and scalability as fixed ports

Support OIR



Introducing Cisco Nexus 6001



High Performance

- Line rate L2 and L3 with all ports and all features and all frame sizes
- 1-microsecond port-to-port latency with all frame sizes
- 40-Gbps flow
- 40-Gbps FCoE
- Cut-through switching for 40 and 10 GE
- 25-MB buffer per three QSFP interfaces

High Scalability

- 48x10 GE + 4x40 GE in 1RU
- Up to 256,000 MAC (115k L2 / 64k L3 at FCS)
- Up to 128,000 ARP (64k at FCS)
- 32,000 LPM (24k at FCS)
- 16,000 bridge domains
- 31 Bidirectional SPAN sessions (16 at FCS)

Feature-Rich

- L2 and L3 features
- FEXlink
- vPC FabricPath TRILL
- FabricPath with segment ID
- Vinci leaf, spine, and border node
- Adapter-FEX/VM-FEX

Visibility and Analytics


- Line-rate SPAN
- Sampled NetFlow*
- Micro-burst and buffer monitoring*
- Latency monitoring*
- Conditional SPAN-SPAN on drop-SPAN on higher latency*

* Some features in this slide are not supported by software at FCS

Nexus 6001 Chassis

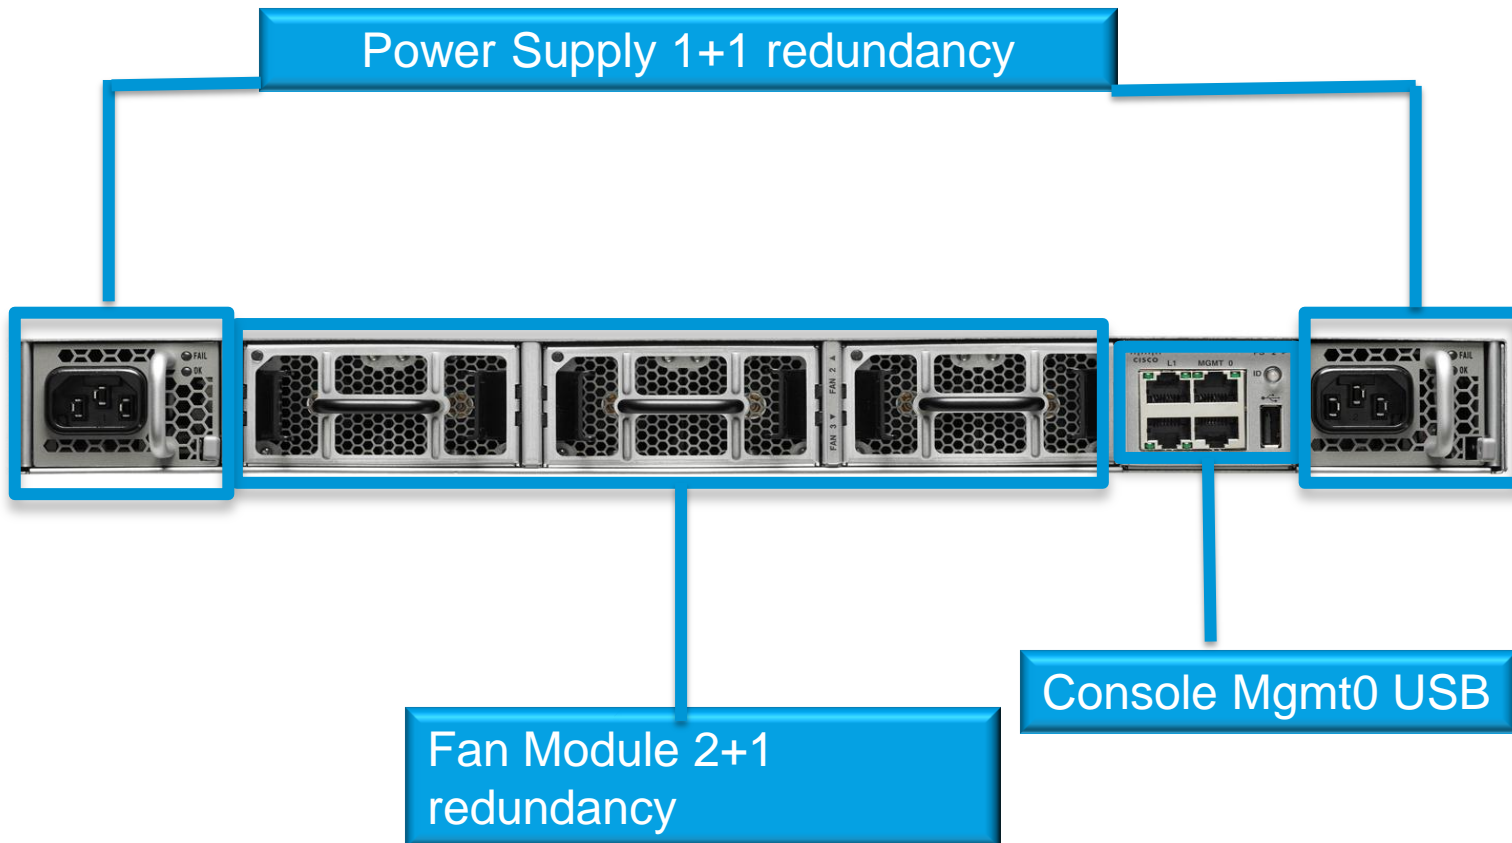
Port-Side View

- Chassis depth: 30 in.
- Chassis width: 17.3 in.



Nexus 6001 Chassis

Power Supply and Fans

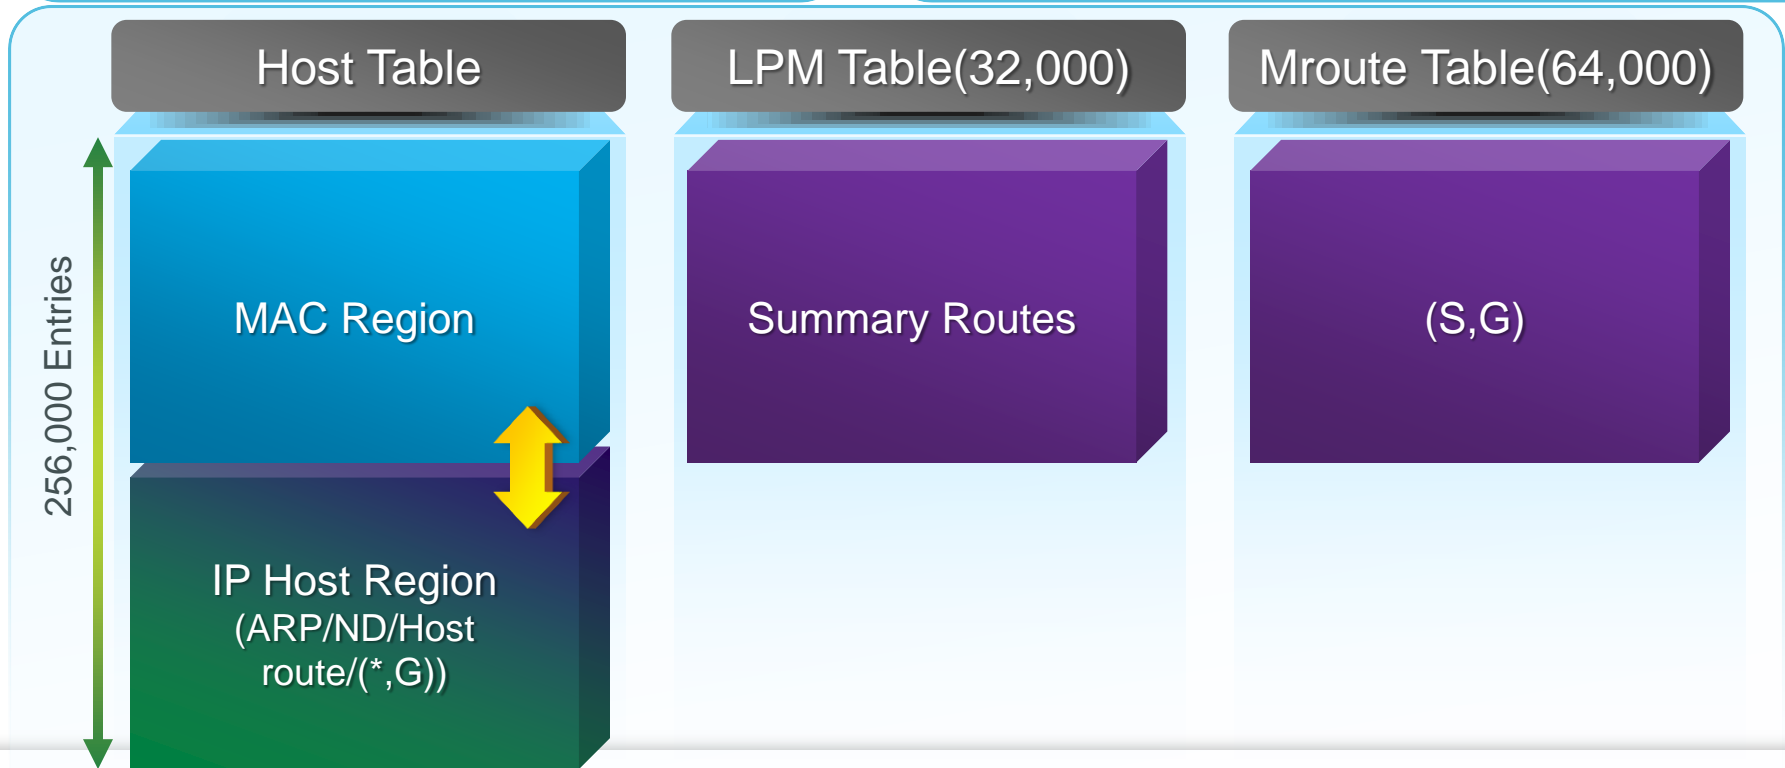


Nexus 6000

Key Forwarding Tables

- Host table: 256,000 -entry hashing table; actual capacity is slightly less than 256,000
- Host table: Shared between MAC, ARP, and ND and /32 host route

- Host table FCS carving: 128,000 MAC, 128,000 IP host
- LPM table: 32,000 entries. Also known as summary routes
- Mroute table: 64,000 entries



* Hardware table size. Please check configuration limit for software scaling

Nexus 6000


FEX Scalability

- 24 FEX per Cisco Nexus® 6004 in L2 and L3 at FCS (32 planned at Harbord+)
- Support for all different types of FEX including HP FEX and Fujitsu FEX for blade servers (except first model Cisco Nexus 2148)
- Dell FEX support with Harbord Maintenance release (Q2CY13)

Cisco Nexus® 6000



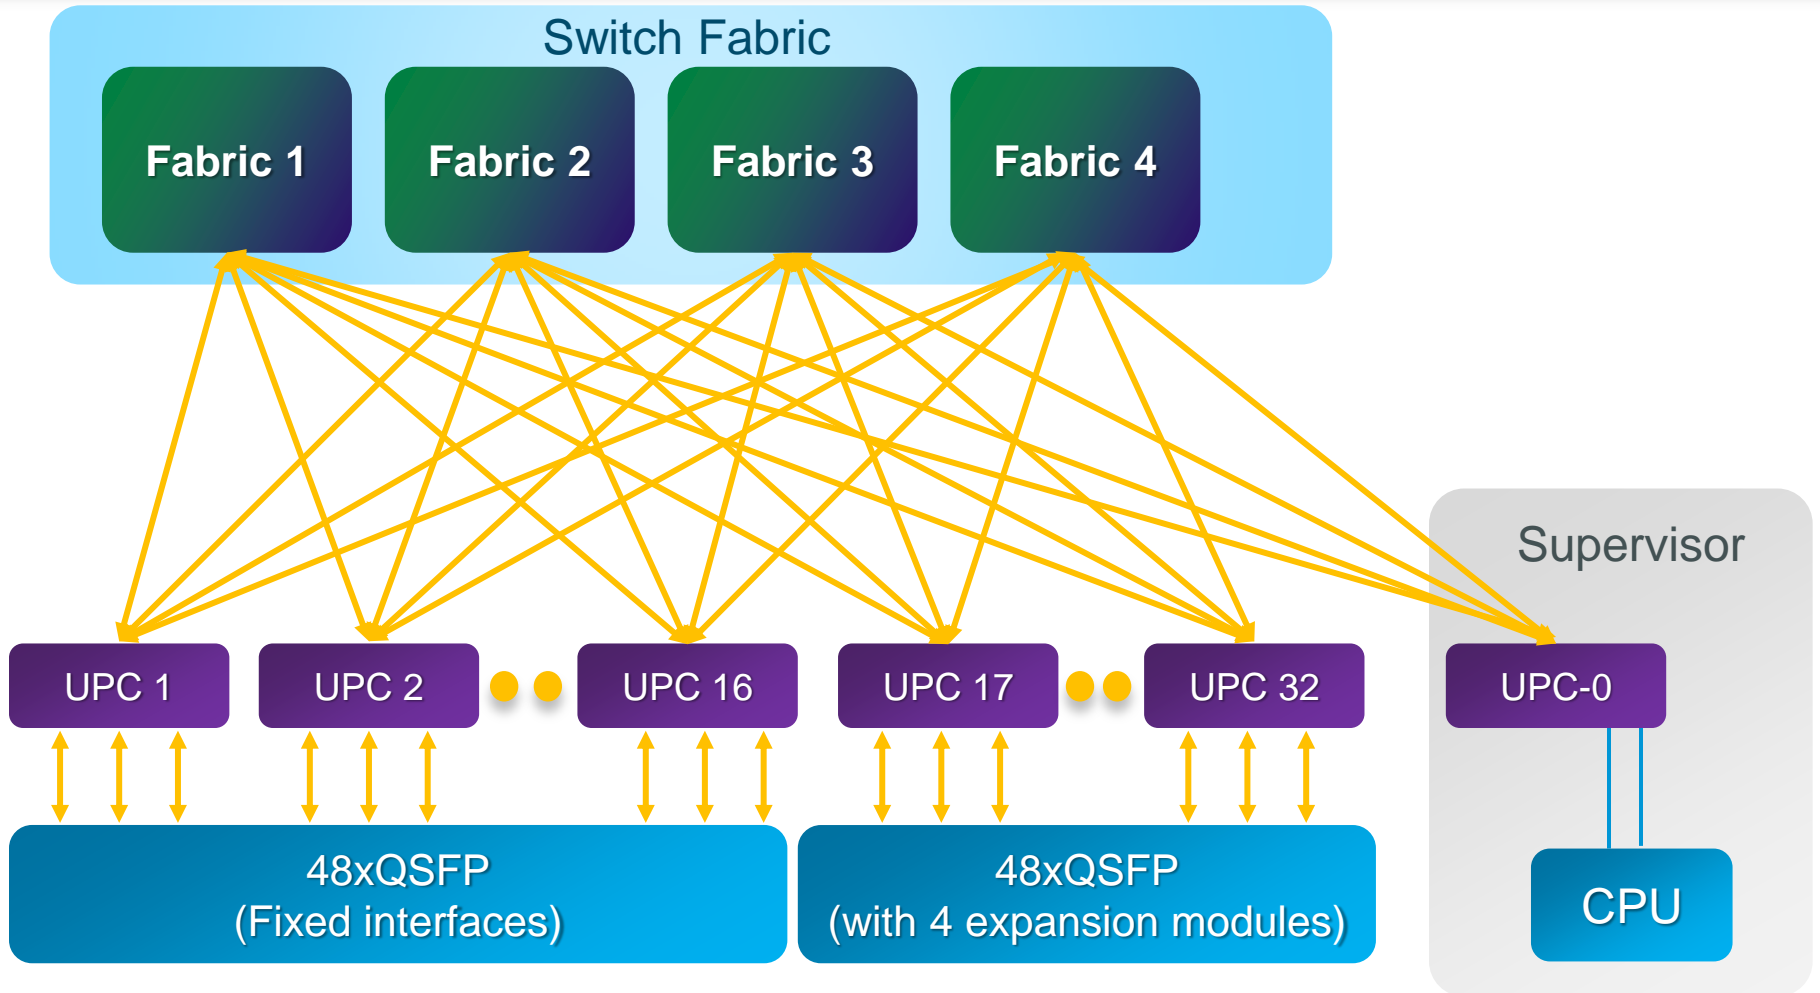
Cisco Nexus® 2000 FEX



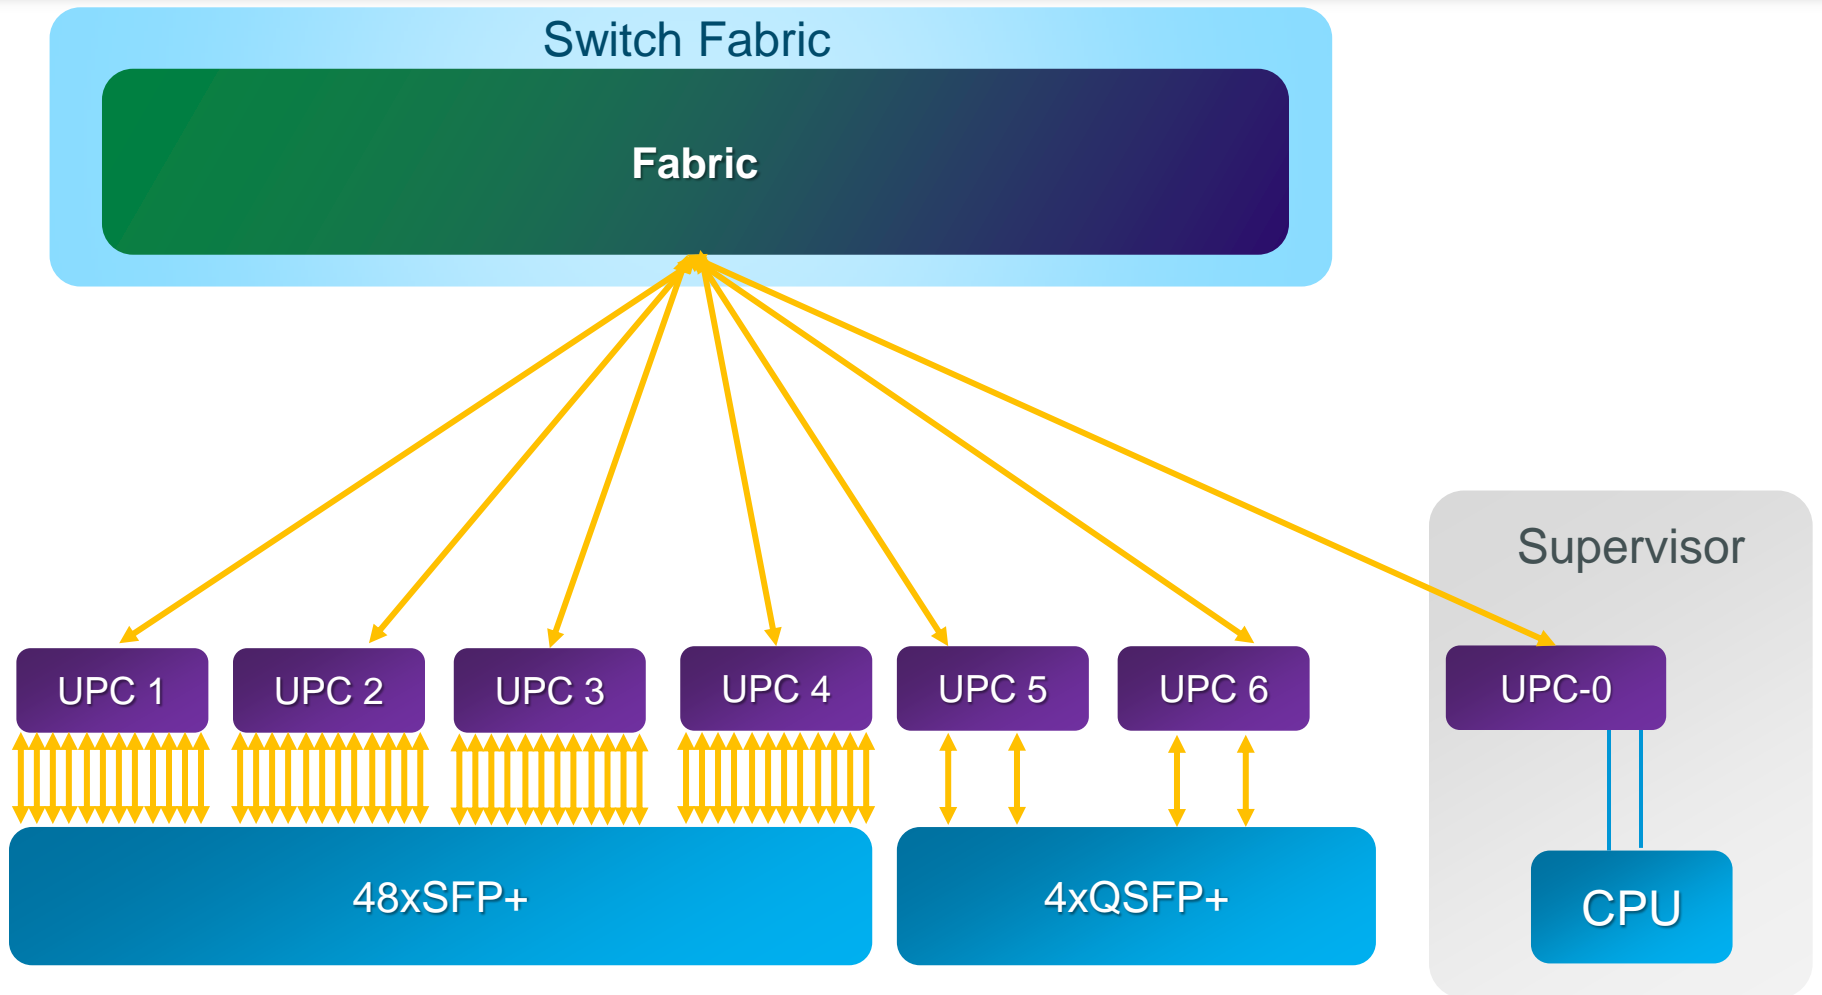
Cisco Nexus 6000 Internal Architecture



Nexus 6004 Internal Architecture




Nexus 6001 Internal Architecture



Unified Port Controller (UPC)

UPC ASIC

- Multimode MAC; built-in PHY for 1, 10, and 40 GE
- Packet parsing and rewriting
- Lookup engine and access control: L2, L3, FabricPath, TRILL, ACL, FCoE, and policing
- Buffering and queuing: Buffer management, PFC for lossless traffic, queuing (Strict Priority Queuing and DWRR), and packet replication (SPAN and multicast)
- Extra fabric bandwidth for SPAN and multidestination traffic




Cisco Nexus 6000 Transceivers and Cables



Interface Speed Mode

- By default port is in 40 GE mode.
- Port speed can be changed at a group of three QSFP ports.
- The group of 12 QSFP ports needs to be reset after port mode change.
- The fixed 48 x QSFP is named in the same way as the ports on the expansion module.



Unified Port Controller (UPC)

Converting 40GE Interface to Four 10 GE Interfaces

Convert the QSFP Interfaces to Four 10 GE Interfaces When Connecting to a SFP+ Port or to the QSFP Interfaces That Operate in Four 10 GE Mode, such as the QSFP Uplink of N2248PQ-10G or the Cisco Nexus® 5500 QSFP GEM.

1 Apply global CLI to change interface types to 10 GE.

- Every three contiguous QSFP interfaces resides on one UPC ASIC.
- The port range specified in the CLI has to include all ports on the ASIC.

2 Power off the affected modules.

- Every group of 12 QSFP interfaces is managed as one module, even for the fixed interfaces.

3 Power on the affected modules.

```
N6004(config)# interface breakout slot  
1 port 1-6 map 10g-4x
```

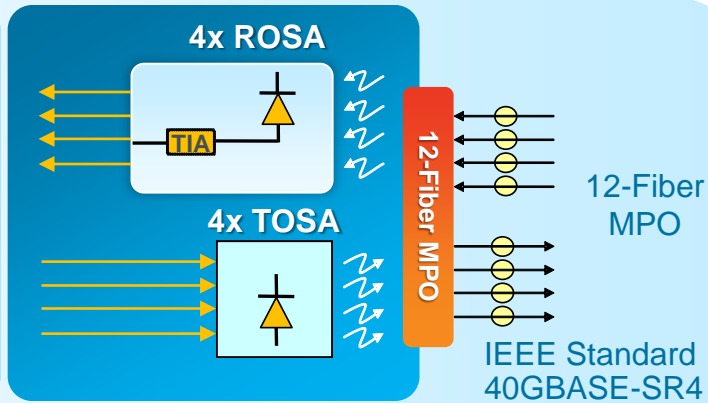
```
N6004(config)# poweroff module ?  
<1-8> Please enter module number
```

```
N6004(config)# poweroff module 1
```

```
N6004(config)# no poweroff module 1  
N6004(config)#
```

QSFP-40G-SR4

Host Card



Key Specs

40GBASE-SR4

Electrical interface

Four 10 GE

Media type

Multimode ribbon fiber

Optical technology

850-nm VCSEL technology/array

Reach

100m with OM3 MMF
150m with OM4 MMF

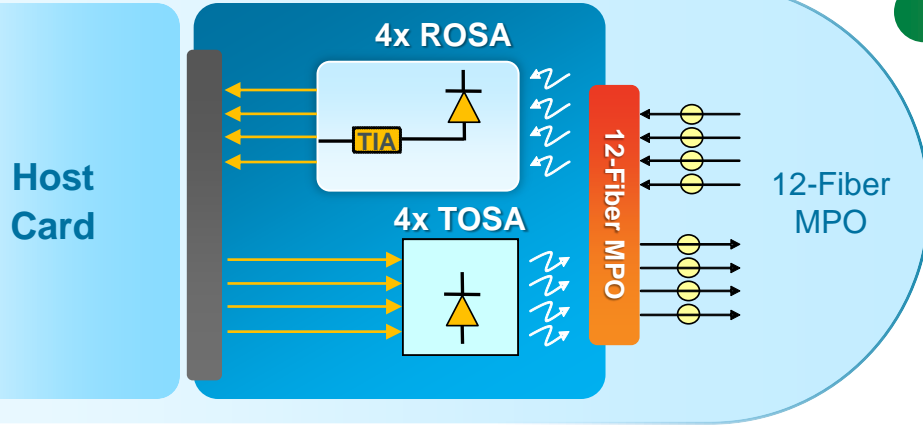
Optical interface

12 fiber MPO/MTP

Power consumption

<1.5W

QSFP-40G-CSR4



FEATURES

- IEEE 10 GE-SR compliance
- Configurable to run in 40 GE mode
 - (Cisco to Cisco connectivity)
- Supported Connectivity
 - Between QSFP-40G-CSR4
 - Between QSFP-40G-CSR4 and 10GE-SR
- Up to 300m reach on MMF OM3
 - Up to 400 m reach on MMF OM4

Key Specs

QSFP-40G-CSR4

Standard

Four 10GBASE-SR

Connector

MPO-12

Reach

OM3: ≤ 300m
OM4: ≤ 400m

Fiber

Multimode ribbon

Power consumption

1.5W

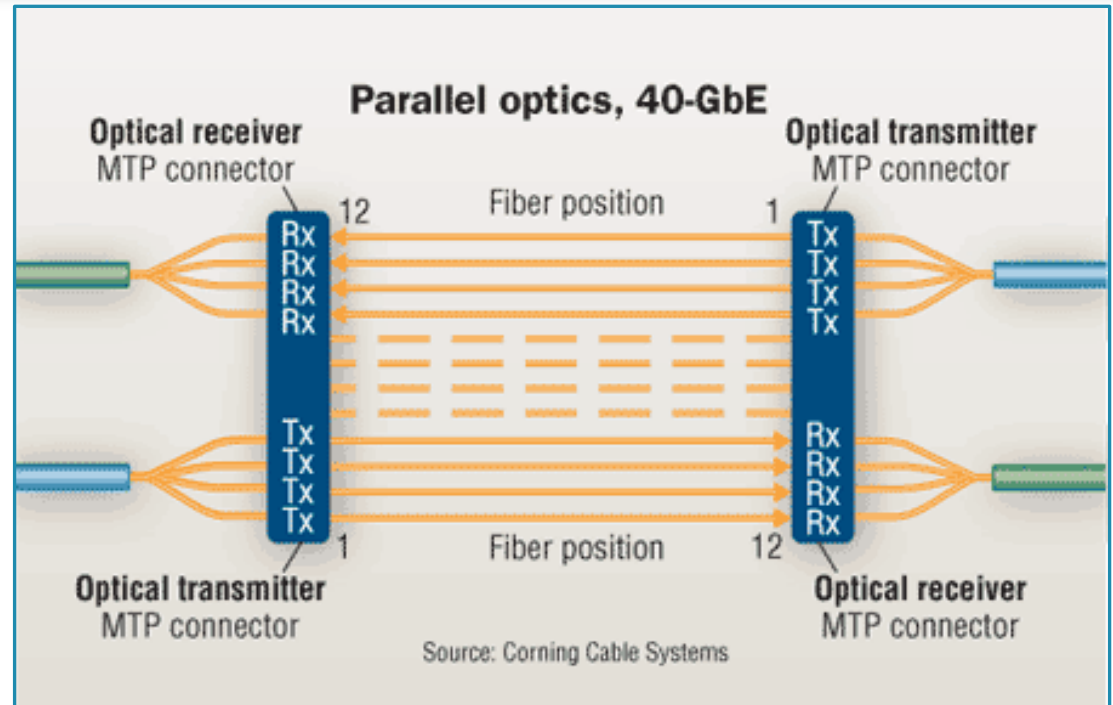
Multimode Ribbon Fiber

40 GE

For QSFP-40G-SR4 and
QSFP-40G-CSR4

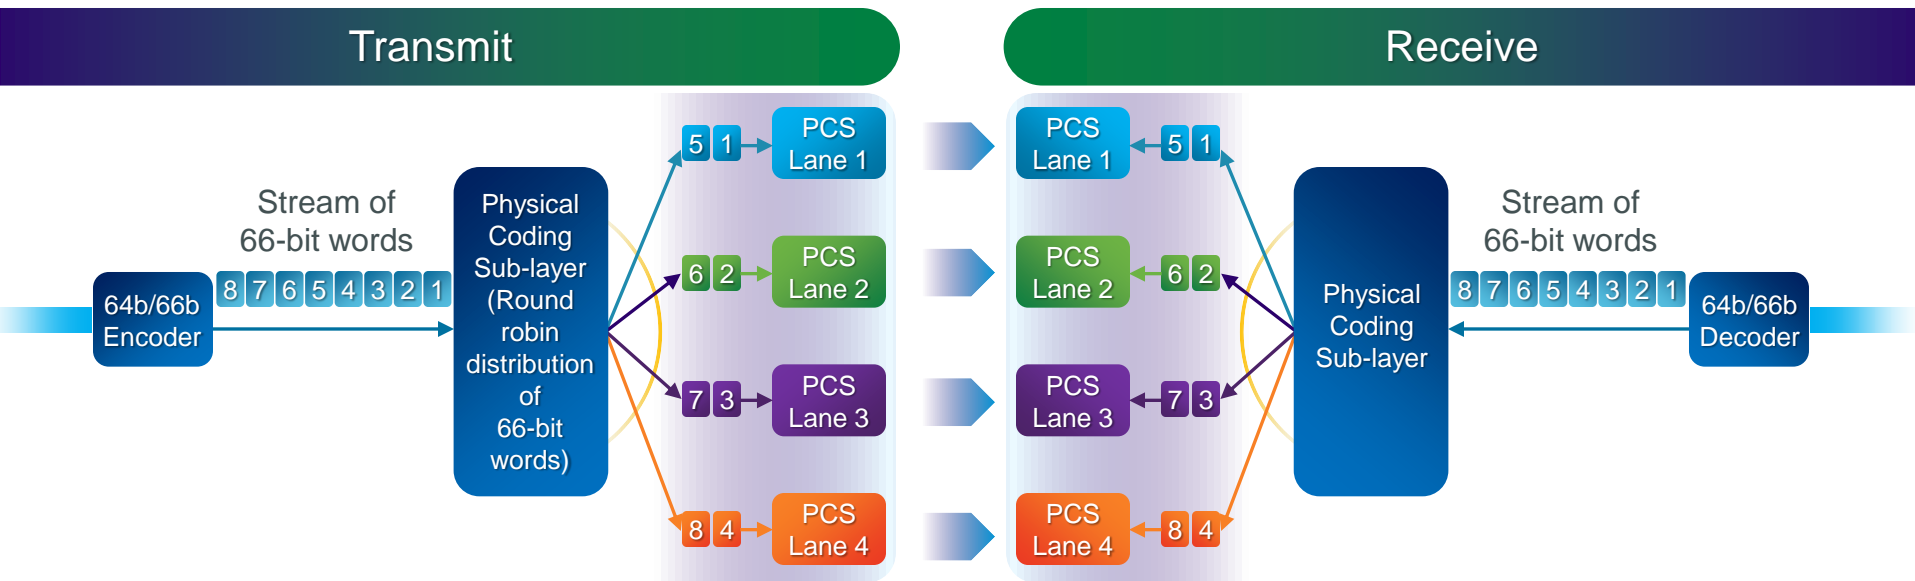
12-Fiber MPO Connector

- MPO/MTP connector with 12 pins
- Use 4 fiber pairs
 - 4 TX and 4 RX allows for 40 GE
 - 4 unused fibers in the center



Bit Spray Over QSFP-40G

- Cisco Nexus 6004 and all its components are designed to carry 40-Gbps flow.
- Each packet is sent over four fibers. No packet hash algorithm is involved.
- Each packet is stripped to 64-bit blocks first; then it encodes to a 66-bit block.
- All the 66-bit blocks are sent over the four fibers in the round-robin fashion.
- The bit spray depicted here applies to QSFP interfaces running at 40 GE mode. When the QSFP interface is configured as four 10 GE interfaces, each fiber and copper cable will carry the packet for one 10 GE interface.



Physical Connection Choices

- FCS support: QSFP-SR4 (100m over OM3); copper 40 GE–40 GE cable and copper breakout cable
- FCS stretch: QSFP-CSR4 (300m over OM3); post FCS: QSFP-LR4 (10 KM over SMF)

40G QSFP+ Interfaces



QSFP-SR4



QSFP-SR4

40G QSFP+ Interfaces

40G QSFP+ Interfaces

1M, 3M, 5M Passive
7M, 10M Active




40G QSFP+ Interfaces

40G QSFP+ Interfaces



QSFP-SR4



10G SFP+ Interfaces

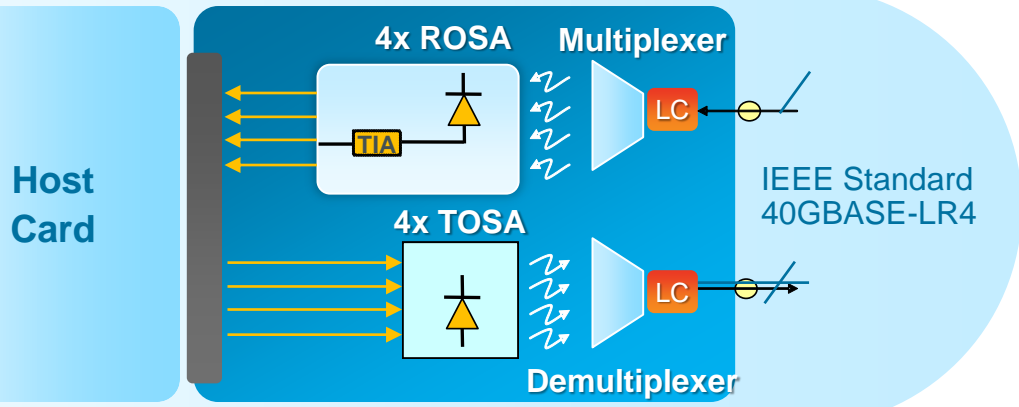
40G QSFP+ Interfaces

QSFP-4x10G-7M
QSFP-4x10G-10M



10G SFP+ Interfaces

QSFP-40G-LR4



Key Specs

40GBASE-LR4

Cisco Will Also Support OTN Data Rate

Media type

Standard single-mode fiber

Optical technology

1300-nm built-in CWDM laser technology
(1271-, 1291-, 1311-, and 1331-nm grid)
Integrated optical multiplexer/de-multiplexer

Reach

10-km single mode

Optical interface

LC duplex connector

Power consumption

3.5W


FET-40G

Planned for Q3CY13

- Low-cost QSFP optical transceiver connecting FEX to Cisco Nexus® 6004
- Supported on Cisco Nexus 6004 and Nexus 2248PQ-10G
- Interoperable with FET-10G
- Support for 100m distance with OM3

Cisco Nexus 6000

FET-40G



Cisco Nexus 2248PQ-10G

FET-40G

Cisco Nexus 6000

FET-40G



FET-10G



Cisco Nexus 2232PP and 2232TM-E 2232TM
Cisco Nexus 2248TP-E 2248TP

Cisco Nexus 6000 Quality of Service



Cisco Nexus 6000 QoS Features

Eight classes of service; 2 reserved for control traffic, 6 for data traffic

Traffic classification

DSCP, CoS, and ACL

Strict Priority Queuing and DWRR

DCBX 802.1Qaz

Packet marking


DSCP, CoS, and ECN

Ingress and egress policing

4096 policers per ASIC

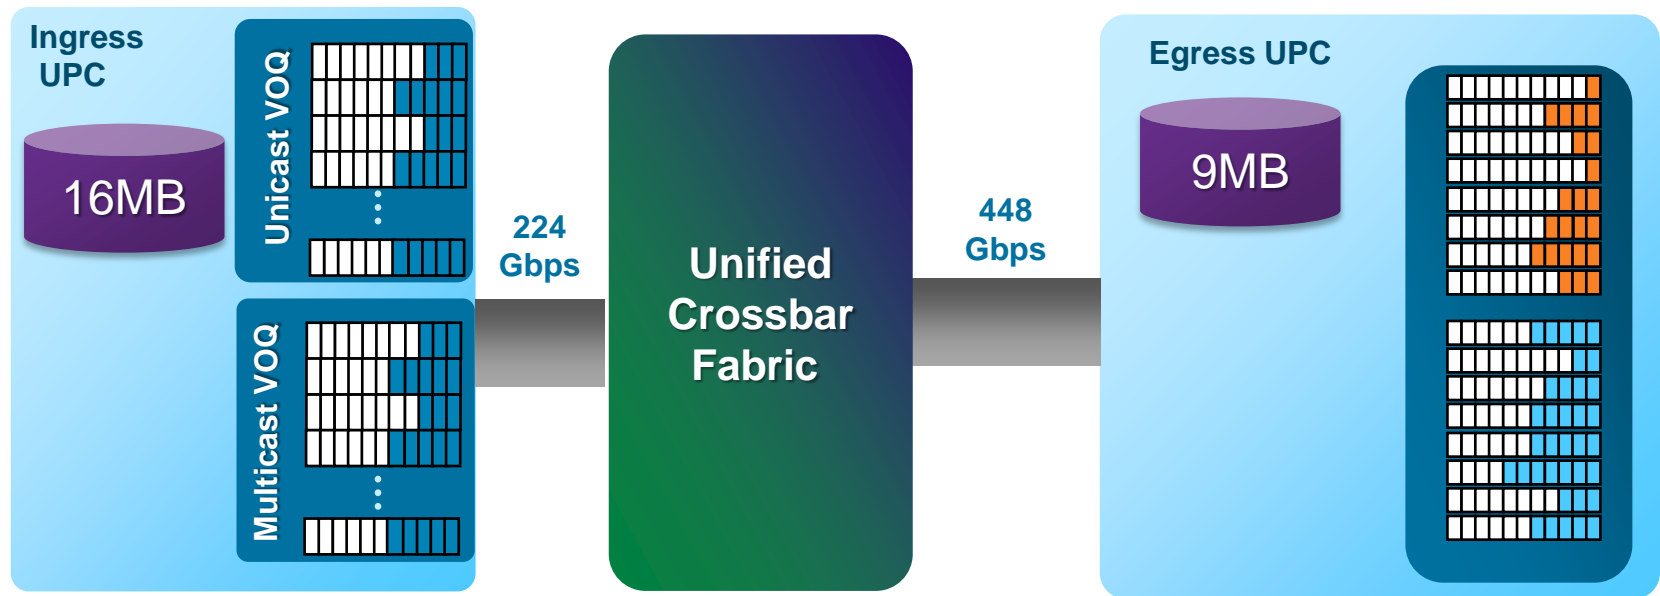
No drop system class

Flexible buffer management



Increased Packet Buffer

- 25-MB packet buffer is shared by every three 40 GE ports or twelve 10 GE ports.
- Buffer is 16 MB at ingress and 9 MB at egress.
- Unicast packet can be buffered at both ingress and egress.
- Multicast is buffered at egress.



Flexible Buffer Management

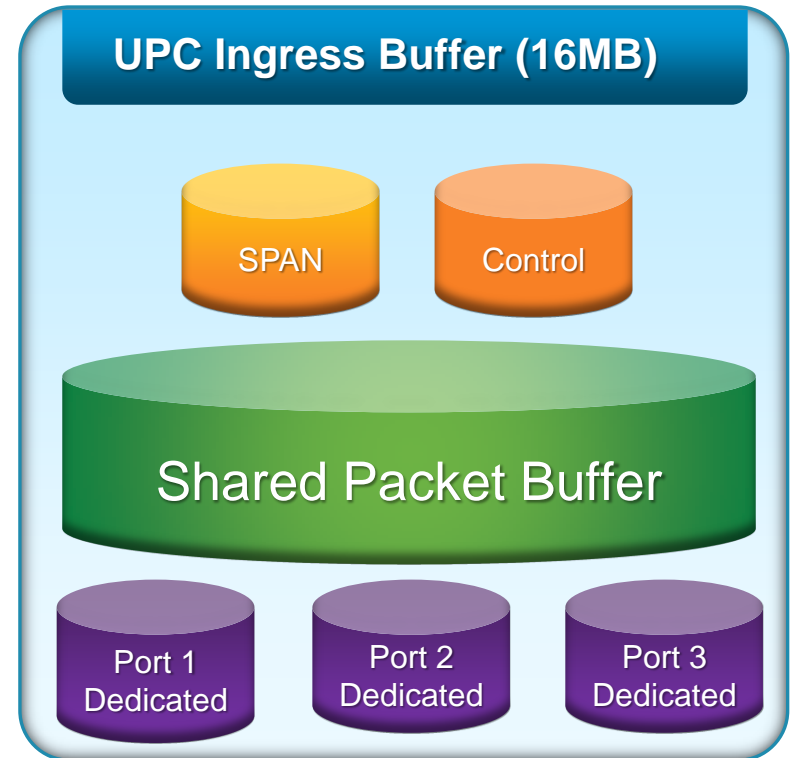
Ingress Buffer

Shared buffer is good for burst absorption.

Dedicated buffer is good for predictable performance for each port.

Buffer management is flexible:
Dedicated plus shared.

Long-distance FCoE, video editing (i.e., AVID), Big Data, and distributed storage



Flexible Buffer Management

Egress Buffer

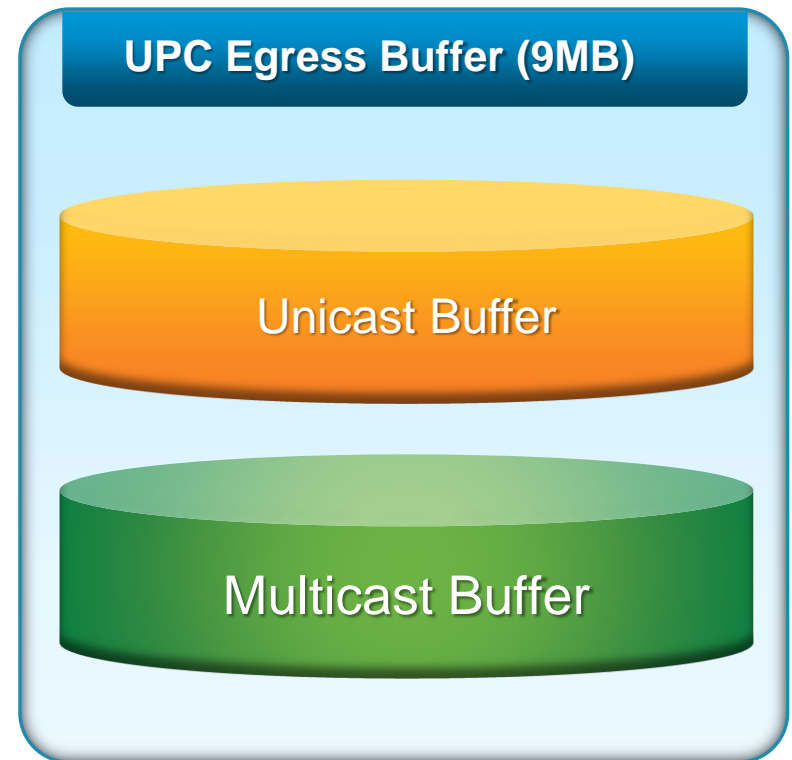
9-MB packet buffer at egress
UPC is shared among three 40
GE or twelve 10 GE.

By default majority of egress buffer is
allocated for multicast traffic

CLI is provided to allocate buffer
between unicast and multicast
(future).

Unicast traffic can be buffered at
egress and ingress.

Multicast is buffered at egress in
case of interface
oversubscription.

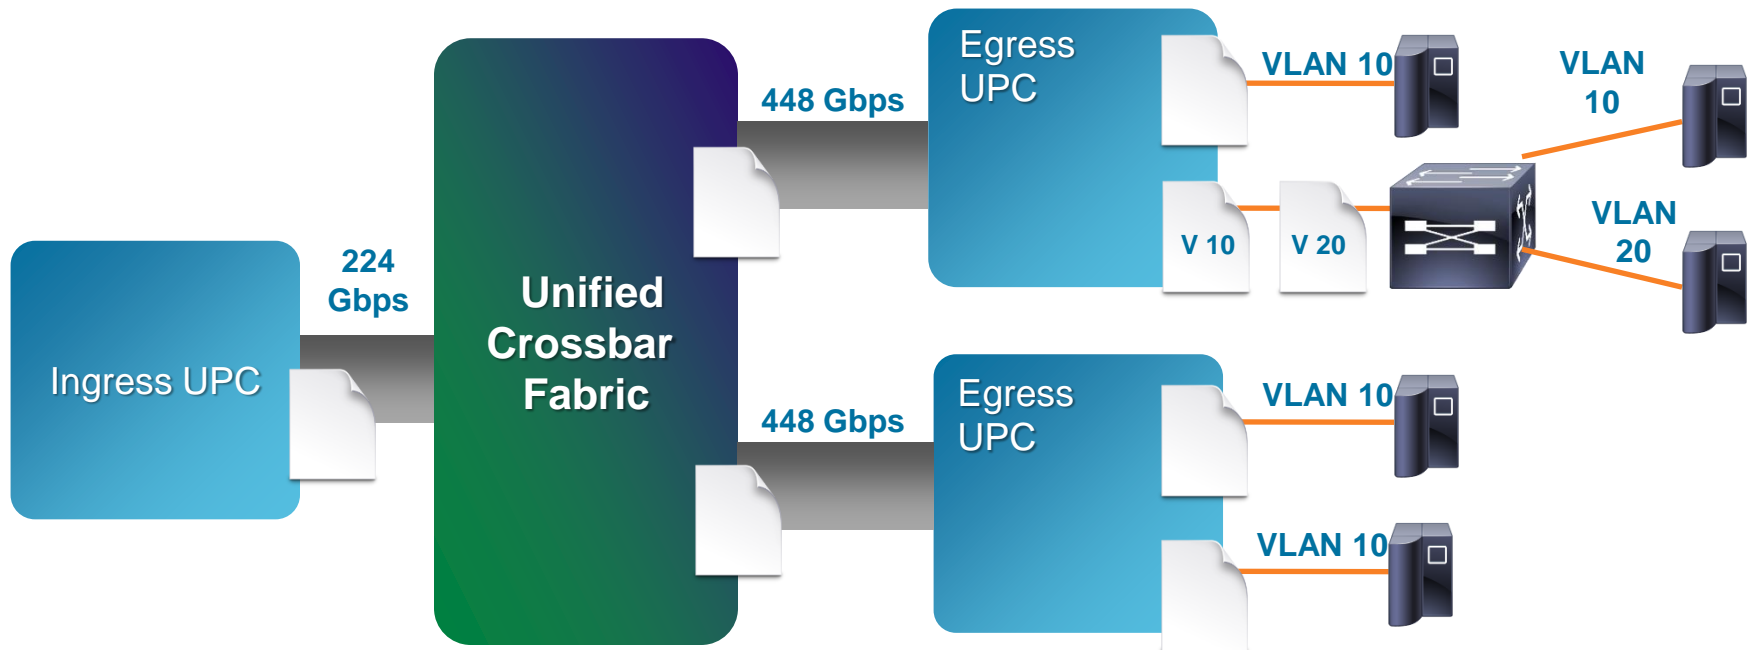


Cisco Nexus 6000 Multicast



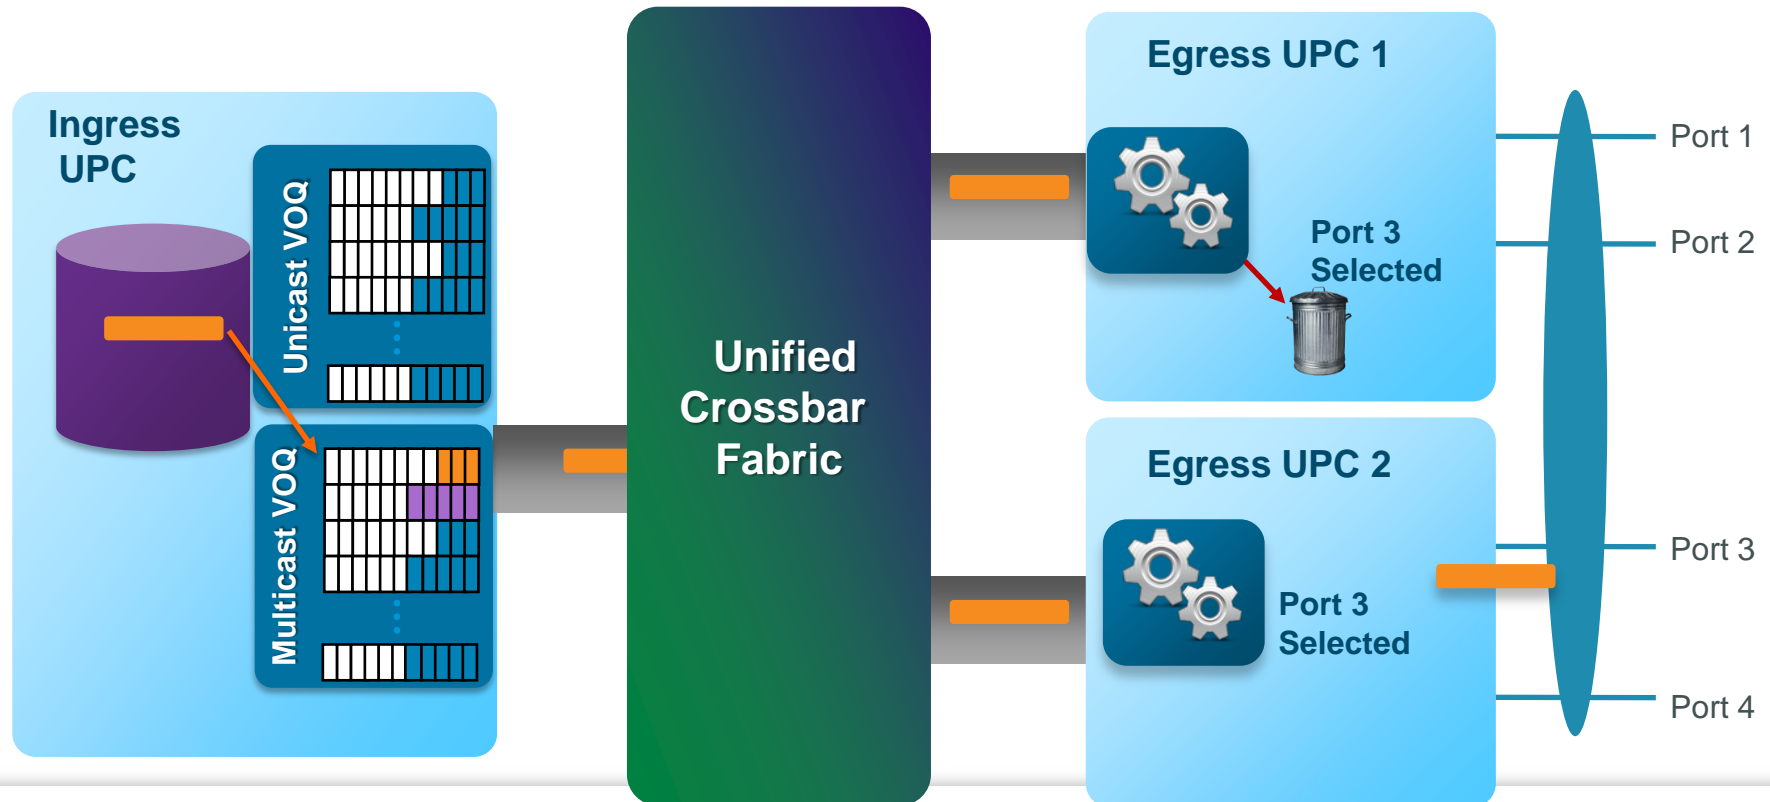
Efficient Multicast Replication

- Optimized multicast replication throughout the system
- Fabric replication and egress replication; one copy is replicated to egress UPC, where there is a receiver—minimizing the traffic load on the switch fabric and eliminating the switch fabric congestion
- Line-rate multicast replication at fabric and egress UPC for all frame sizes



Multicast Hashing over PortChannel

- N6004 implements flow based hashing for multi-destination traffic and it support multicast load sharing over PortChannel with 5-tuple packet header.
- Traffic is replicated to all egress UPC where PortChannel member resides
- Egress UPC run hash calculation and one egress port is chosen to send out multicast packets. The UPC ASIC that is not supposed to send out packet will drop packet (Egress UPC 1 in the example)



Cisco Nexus 6000 SPAN



Nexus 6000 SPAN Differentiators

Large number of active SPAN sessions

- 31 bidirectional active SPAN sessions supported by hardware
- 16 bidirectional active SPAN sessions supported at FCS

Line-rate SPAN

- Line-rate SPAN for multiple 40 GE ports
- Built-in extra capacity in fabric to accommodate SPAN traffic

Intelligent SPAN

- Prioritize data over SPAN in case of fabric link oversubscription due to SPAN traffic
- Conditional SPAN



For More Information

For more information about Nexus
6001/6004 collateral please visit

<http://www.cisco.com/go/nexus6000>



Thank you.

