The bridge to possible

Catalyst C9500X and C9600X Deep Dive

Ninad Diwakar and Sai Zeya
C9K Technical Marketing Team
April 1, 2022

Digital trends shaping the future of business

Hybrid work

Work from home | Work from anywhere | Work from office

Industry 4.0

Wireless | Automation | Internet of Things | AI/ML

Hybrid cloud

Private cloud | Hybrid cloud | Public cloud

and the network is the core engine of hybrid work

Cisco Wi-Fi 6E and Catalyst 9000X

Enabling better business outcomes end-to-end with simplicity and choice

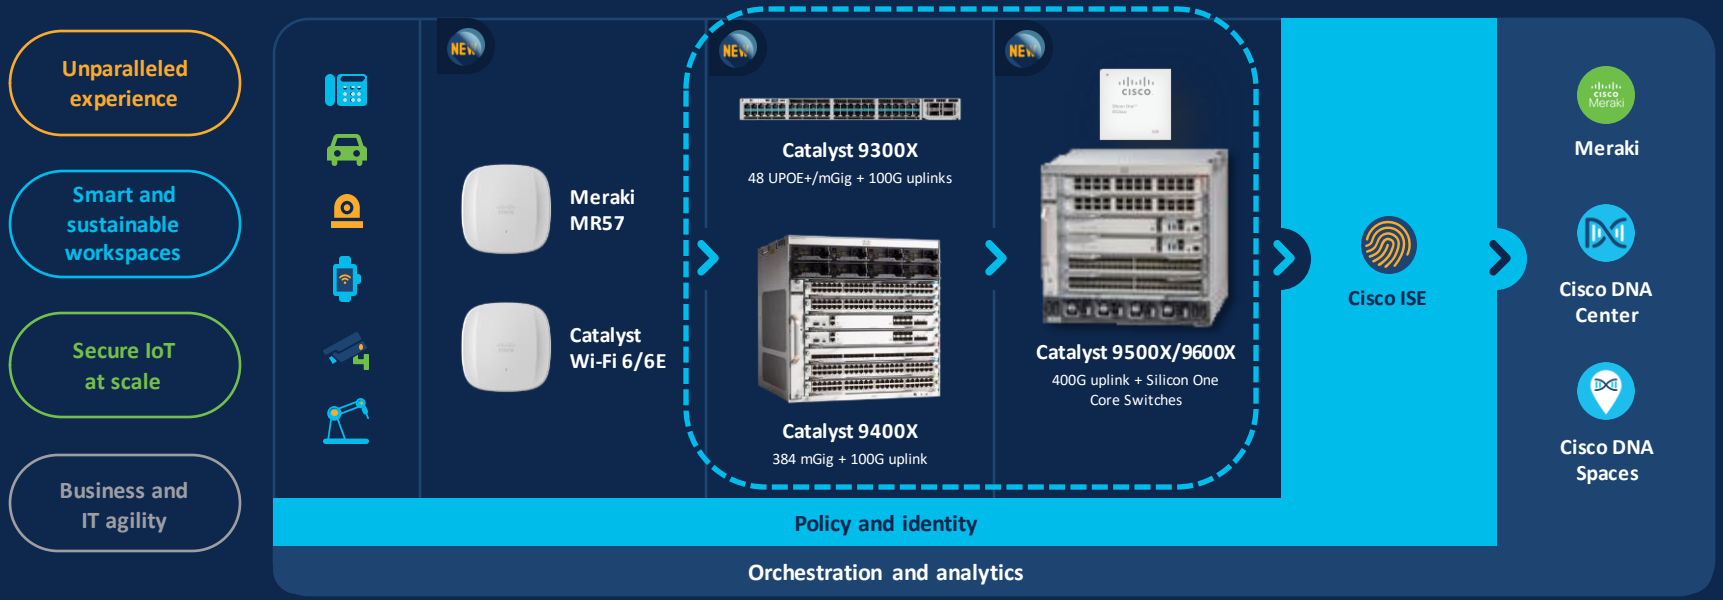

Cisco access networking stack

Catalyst 9000X – Expanding industry leadership

Adding the “X factor” to the industry’s leading switching family

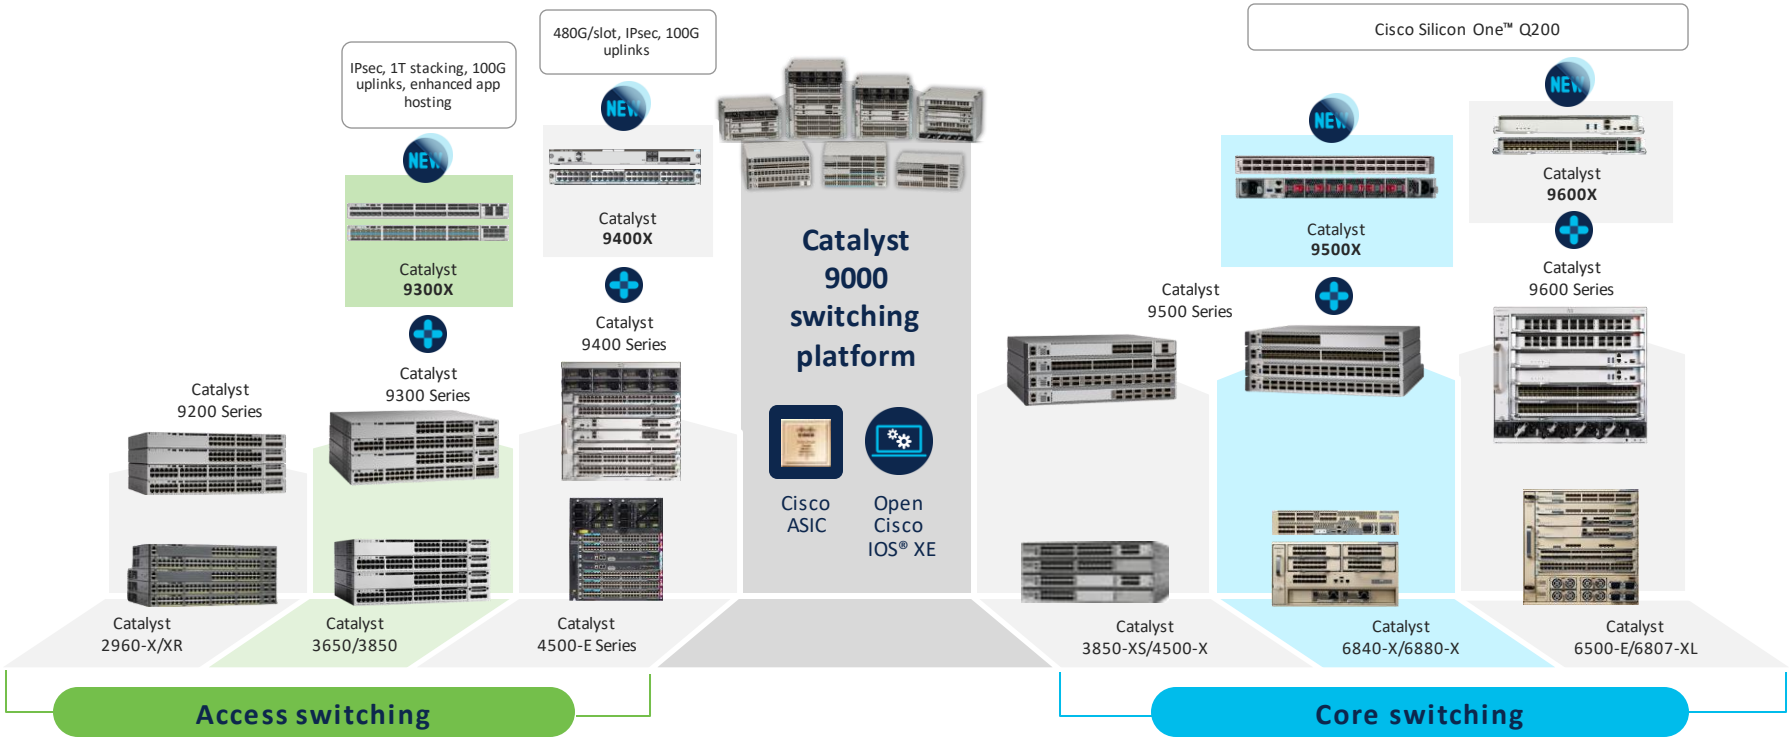

Introducing

Cisco Silicon One™

Cisco Silicon One™

Switching Silicon

- **High Throughput**
  extremely fast hardware-based L2-L4 forwarding and services (measured in Terabits per second)
- **Optimized Scale**
  optimized for Campus LAN environments with moderate IP & MAC scale (10s-100s of thousands)
- **Low Latency**
  extremely low hardware-based system latency (measured in Nanoseconds & Microseconds)
- **Streamlined Buffering**
  shallow buffering systems to reduce latency, with very high throughput

Routing Silicon

- **Flexible Features**
  complex, stateful L3-L7 forwarding and services (measured in Gigabits per second)
- **Massive Scale**
  optimized for WAN/SP environments with very high IP scale (100s of thousands - millions)
- **Mixed Interfaces**
  support for Ethernet, Serial, Cellular and other types and speeds in a single system
- **Deeper Buffering**
  deep buffers to accommodate different speeds, bursts and different flow patterns

Cisco Silicon One Bringing Switching and Routing convergence

Introducing Cisco Silicon One™

One architecture, multiple devices

www.cisco.com/c/en/us/solutions/service-provider/innovation/silicon-one.html

Q202

3.2 Tbps

Q201

6.4 Tbps

Q100

10.8 Tbps

Q200

12.8 Tbps

- First network silicon to break the 10 Tbps barrier
- Comprehensive routing, with switching efficiency
- Flexible P4 NPL Programmable packet processing

- Multiple functions: system-on-chip, line card or fabric
- Multiple form-factors: fixed or modular
- Multiple segments: enterprise and service provider

Cisco Silicon One™ Q200

Industry leading Switching and Routing Silicon

12.8T BW

8.1 Bpps

8G HBM for deep buffers

10M IPv4 or 5M IPv6 route scale

Fully P4 Programmable Pipeline

50G PAM4 Serdes

First 7nm ASIC providing lowest watts/GE power consumption

Fully P4 programmable enabling feature velocity

Multi slice architecture for flexibility and scale

Routing Capabilities with Switching Power and Performance

Catalyst 9000 Series – Common Building Blocks

Programmable x86 Multi-Core CPU

Application Hosting
Secure Containers

Open IOS XE® Polaris

Model-Driven APIs
Modular Patching

Cisco Silicon One™ Q200 ASIC

Programmable Pipeline
Flexible Tables

Same binary image for both UADP and Silicon One C9K platforms*

* C9200 uses IOS XE Lite

Extending Cisco Catalyst 9500 & 9600 Series

Powered by Cisco Silicon One™ Q200 ASIC

C9500 & **C9600-SUP1**
(w/ UADP 3.0)

Optimized for Features

✓ **Speed**
- 1/10 & **25G SFP**
- 40 & **100G QSFP**

✓ **Scale**
- Upto **128K MACs**
- Upto **256K Routes**
- **108MB Buffers** (3x 36MB)

✓ **Services**
- L2/L3 Routing, MPLS
- LAN MACsec, Netflow, NAT
- Custom ASIC Templates
- Campus Fabric (SDA & EVPN)

Extending Cisco Catalyst 9500 & 9600 Series

Powered by Cisco Silicon One™ Q200 ASIC

C9500 & C9600-SUP1
(w/ UADP 3.0)

Optimized for Features

✓ **Speed**
- 1/10 & **25G SFP**
- 40 & **100G QSFP**

✓ **Scale**
- Upto **128K MACs**
- Upto **256K Routes**
- **108MB Buffers** (3x 36MB)

✓ **Services**
- L2/L3 Routing, MPLS
- LAN MACsec, Netflow, NAT
- Custom ASIC Templates
- Campus Fabric (SDA & EVPN)

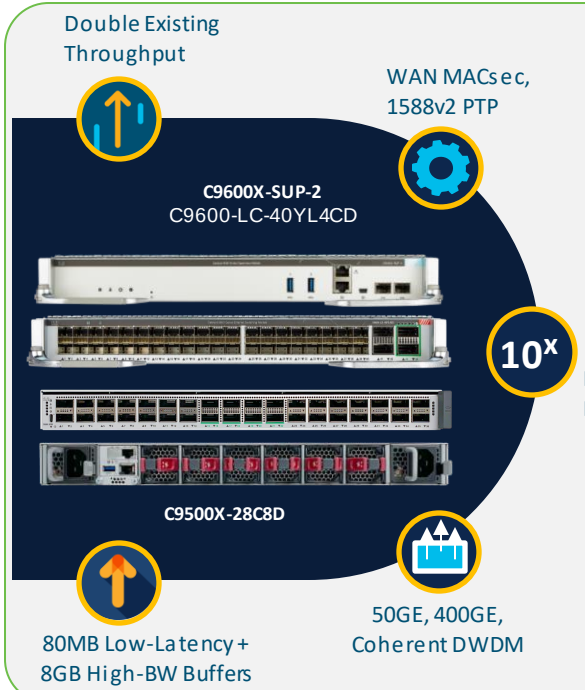

NEW

C9500X & C9600X-SUP2
(w/ S1 Q200)

Optimized for Scale

✓ **Speed**
- 10/25G & **50G**** SFP
- 40/100G & **200****/400G QSFP

✓ **Scale**
- Upto **256K MACs**
- Upto **2M Routes**
- **80MB + 8GB Buffers**

✓ **Services**
- L2/L3 Routing, MPLS/TE
- LAN & **WAN-MACSec***, Netflow, NAT**
- Custom ASIC Templates
- Campus Fabric (SDA & EVPN*)

Maximum Investment Protection | **400G Leadership in Campus** | **Unmatched Flexibility**

* Targeted 17.8.1
** Hardware capable. Roadmap (not committed)

Introducing

Catalyst 9500X

Catalyst 9500 Series

Extending the Catalyst 9500 High-Performance Fixed Core

Non-XL Scale

Catalyst 9500H

(UADP 3.0)

C9500-32C / C9500-32QC

C9500-48Y4C / C9500-24Y4C

Core + Edge

XL Scale

NEW

Catalyst 9500X

(S1 Q200)

C9500X-28C8D

Total Capacity

3.2 Tbps

32 x 100G or
48 x 25G + 4 x 100G

2x

Total Capacity

6.0 Tbps

36 x 100G or
28 x 100G + 8 x 400G

C9500X-28C8D

Gen2 Fixed 1RU QSFP Switch - 36x 100G / 28x 100G + 8x 400G

6.0 Tbps

Target FCS
Q1CY22
Software
IOS-XE 17.7.1

- 1x **Cisco SiliconOne Q200** ASIC
  - **6.0 Tbps** System Throughput
  - **28x QSFP28** ports - 40/100GE
  - **8x QSFPDD** ports - 100/200*/400GE
- 1x **8C 2.4GHz x86** CPU with 2x **16GB (32GB) DDR4** DRAM
- **16GB Flash;** Optional SSD (480G, 960G)
- 12x **CDR5M** PHYs
  - MACSec, WAN-MACSec, ClearTag v3.4
  - IEEE 1588 & PTPv2*
- Various SFP Breakout & QSA support*

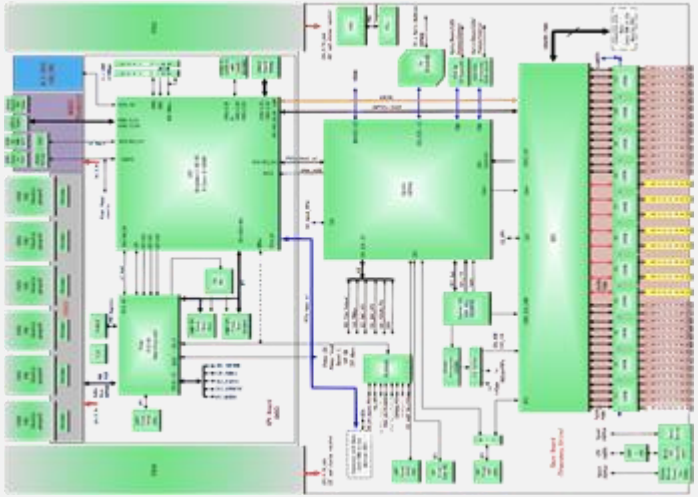

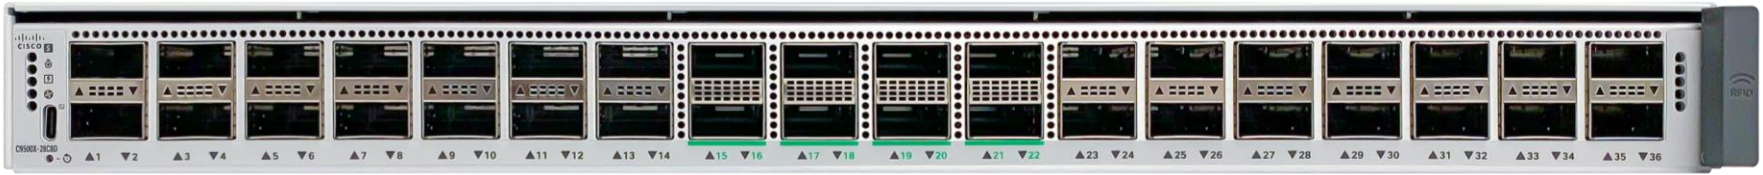

*Hardware capable. Roadmap (not committed)

C9500X-28C8D Block Diagram

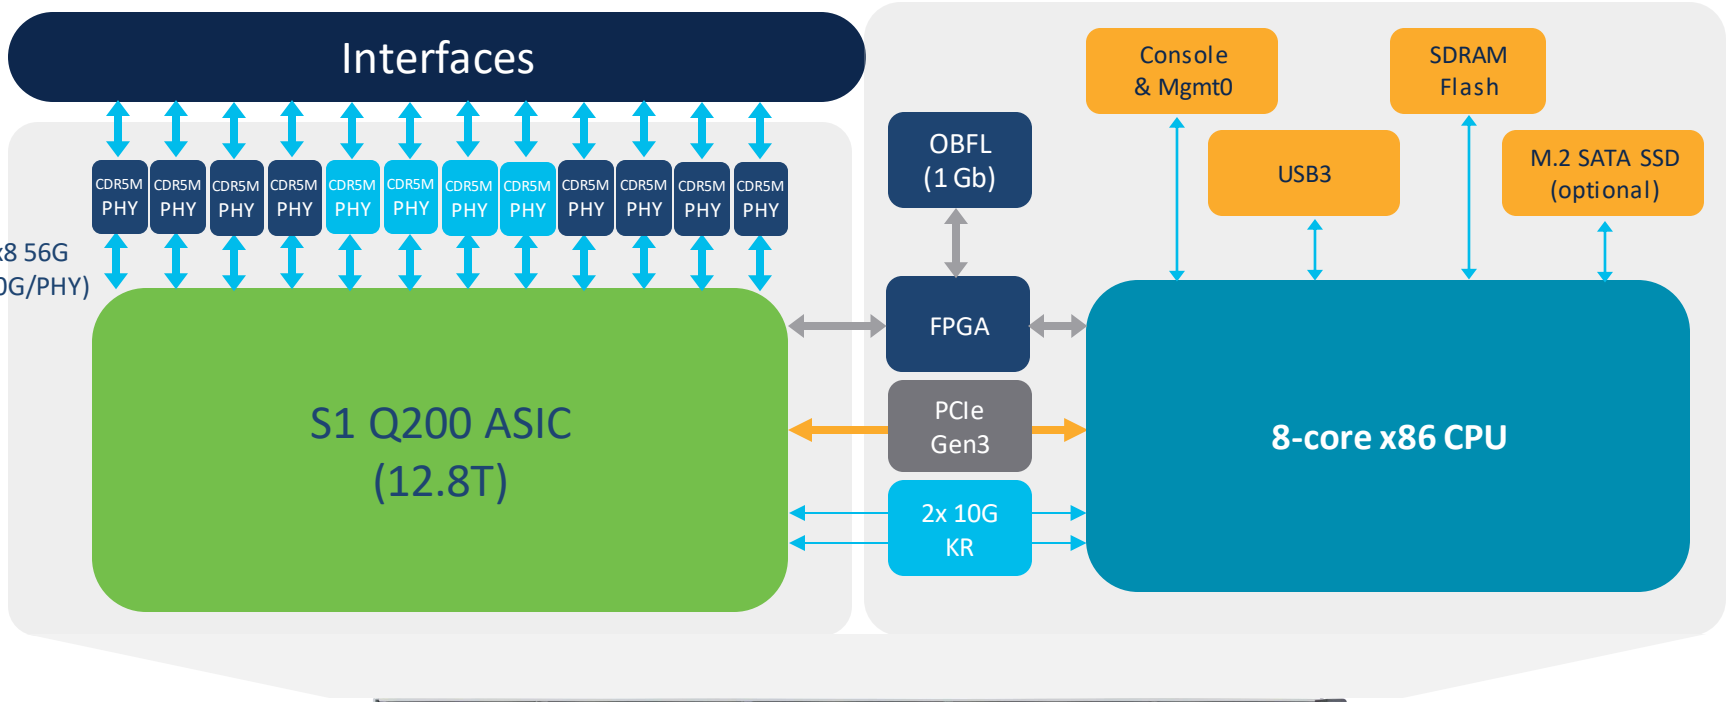

C9500X-28C8D

SFP Breakout & QSA* Support

100GE Ports

400GE Ports

100GE Ports

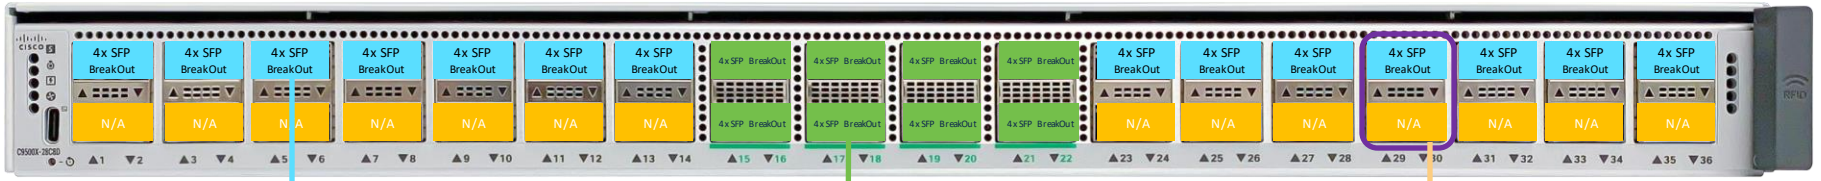

- Breakout is supported on **Top row (odd numbered) 100GE (QSFP28) ports**
- Breakout mode: 4x 10/25G
- **Maximum ports**: 14 x 4 = 56

- Breakout supported on **All 400GE (QSFPDD) ports**
- Breakout mode: 4x 10/25G; 8x 10/25G*; 4x 100G*
- **Maximum ports**: 8 x 4 = 32

- Breakout not supported on **Bottom row (even numbered) 100GE (QSFP28) ports**
- The bottom port is disabled when the corresponding top port is configured as break-out

Maximum ports with breakout at FCS: 88 (56 +32)

* Roadmap (not committed). System can support up to 120x 10G/25G

C9500X – Reversible Airflow

Back to Front
Port-side Exhaust

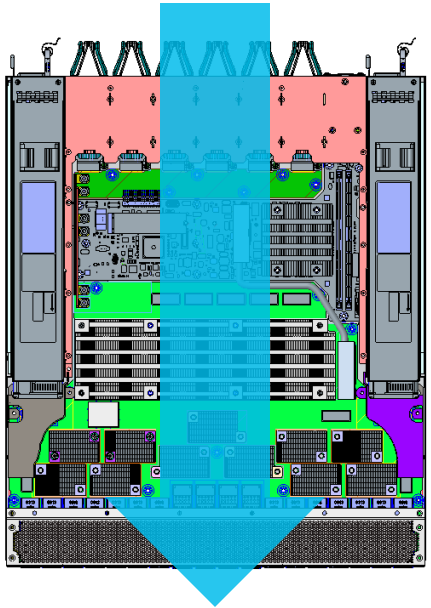

- Color of Fan Unit handle/latch represents direction of airflow
- Different Fan PIDs for different airflow directions
  - **Royal Blue – Back to Front**
  - **Burgundy – Front to Back**
- All Fans must be the same color (direction) to work correctly

Single **1500W AC/DC PSU** with **Cisco Grey** latch for both airflow directions

Front to Back
Port-side Intake

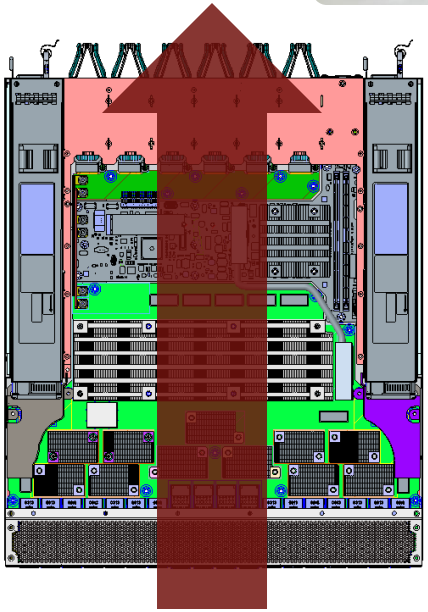

FYI

C9500X and C9500 - Physical

Cisco C9500X (S1 Q200)

Cisco C9500 (UADP 3)

	28C8D	24Y4C	48Y4C	32QC	32C
Height (1.75" RU)	1RU	1RU			
CPU (number)	**2.7GHz 8C Intel** (BDW-NS)	2.0GHz 8C Intel (BDW)			
DRAM (type)	**32GB** (DDR4)	16GB (DDR4)			
ASIC (number)	**Q200** (1x)	UADP3 (1x)			UADP3 (2x)
Capacity	**6.0T**	1.2T	1.6T	1.6T	3.2T
10G max	**120*^** (88 @ FCS)	24	48	16^	16^
25G max	**120*^**(88 @ FCS)	24	48	16^	16^
50G max	**120*^**	--	--	--	--
40G max	**28 + 32^***(36 @ FCS)	4	4	32	32
100G max	**28 + 32^*** (36 @ FCS)	4	4	16	32
400G max	**8**	--	--	--	--

* Roadmap (not committed)
^ Breakout Cable ports

FYI

C9500X and C9500 – Features and Scales

Cisco C9500X (S1 Q200)

Cisco C9500 (UADP 3)

	Default	Maximum (Custom)	Default	Maximum (Custom)
MAC Addresses	**128K**	256K	**80K**	128K
IP Host Routes	**128K**	256K	**80K**	128K
Multicast L2 groups	**16K**	64K*	**16K**	48K
Multicast L3 routes	**32K**	64K*	**32K**	48K
IP LPM Routes	**2M**	2M	**212K**	256K
MPLS Labels	**256K**	512K	**32K**	64K
SGT/OG Labels	**32K**	64K	**32K**	64K
NAT* Sessions	**16K***	128K*	**3K**	16K
Sec ACL Entries	**8K**	10K*	**12K**	27K
QoS ACL Entries	**8K**	10K*	**8K**	21K
PBR* ACL Entries	**8K***	10K*	**3K**	16K

* Roadmap for C9500X (not committed)

Introducing

Catalyst 9600X

Catalyst 9600 Series

Extending Modular Core with a Performance-Optimized Supervisor 2

Supervisor 1

Total Capacity

4.8 Tbps

Slot B/W

1.2 Tbps

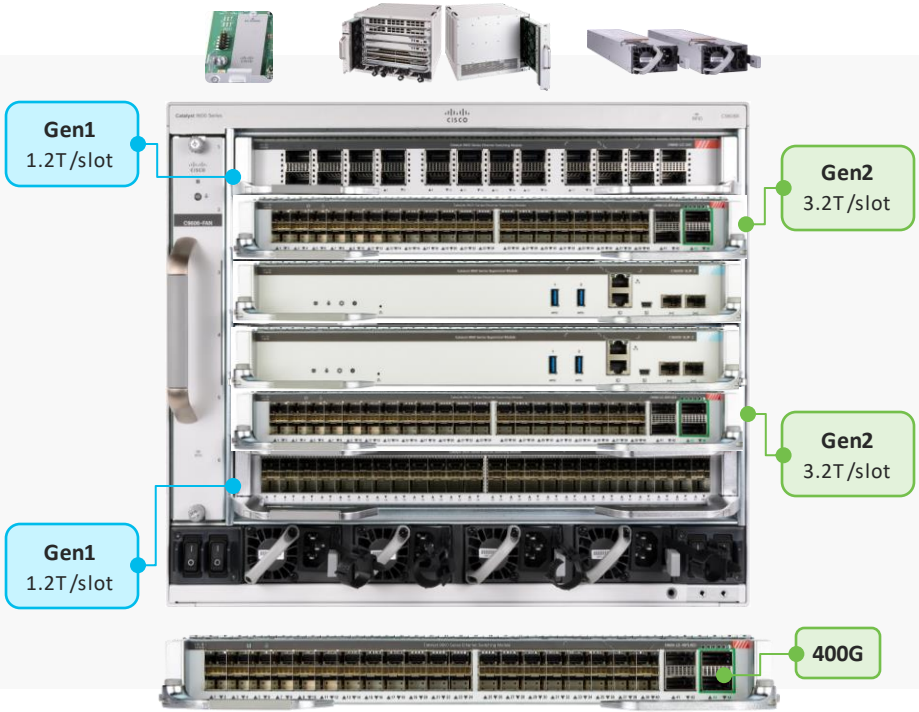

NEW

Supervisor 2

Total Capacity

12.8 Tbps

Slot B/W

3.2 Tbps

3x

C9600X-SUP-2

Gen2 Supervisor Module with Silicon One™ Q200

12.8 Tbps

Target FCS
Q1CY22
Software
IOS-XE 17.7.1

- **1x Cisco SiliconOne Q200 ASIC** (12.8Tbps)
  - 3.2Tbps per Slot
  - Optimized for 10G to 400G
- **1x x86 2.7GHz CPU with 2x 16GB (32GB) DDR4 DRAM**
- **Management**
  - 2x 10G SFP+ Mgmt ports to CPU (App Hosting)
  - 1x 10/100/1000M RJ45 Mgmt0 port
  - 1x RJ45 Console port, USB Type-B-Mini port
- **Storage**
  - 2x USB3.0 Type-A SSD ports
  - 480-960GB M.2 SATA Drive (optional)

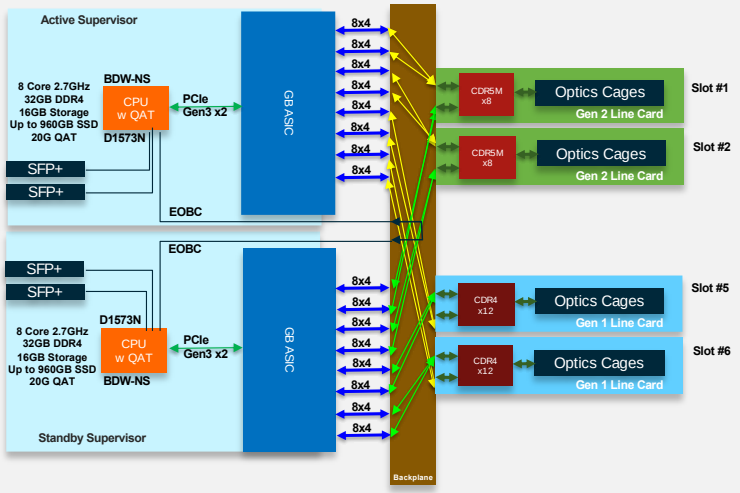

C9600X-SUP-2 Block Diagram

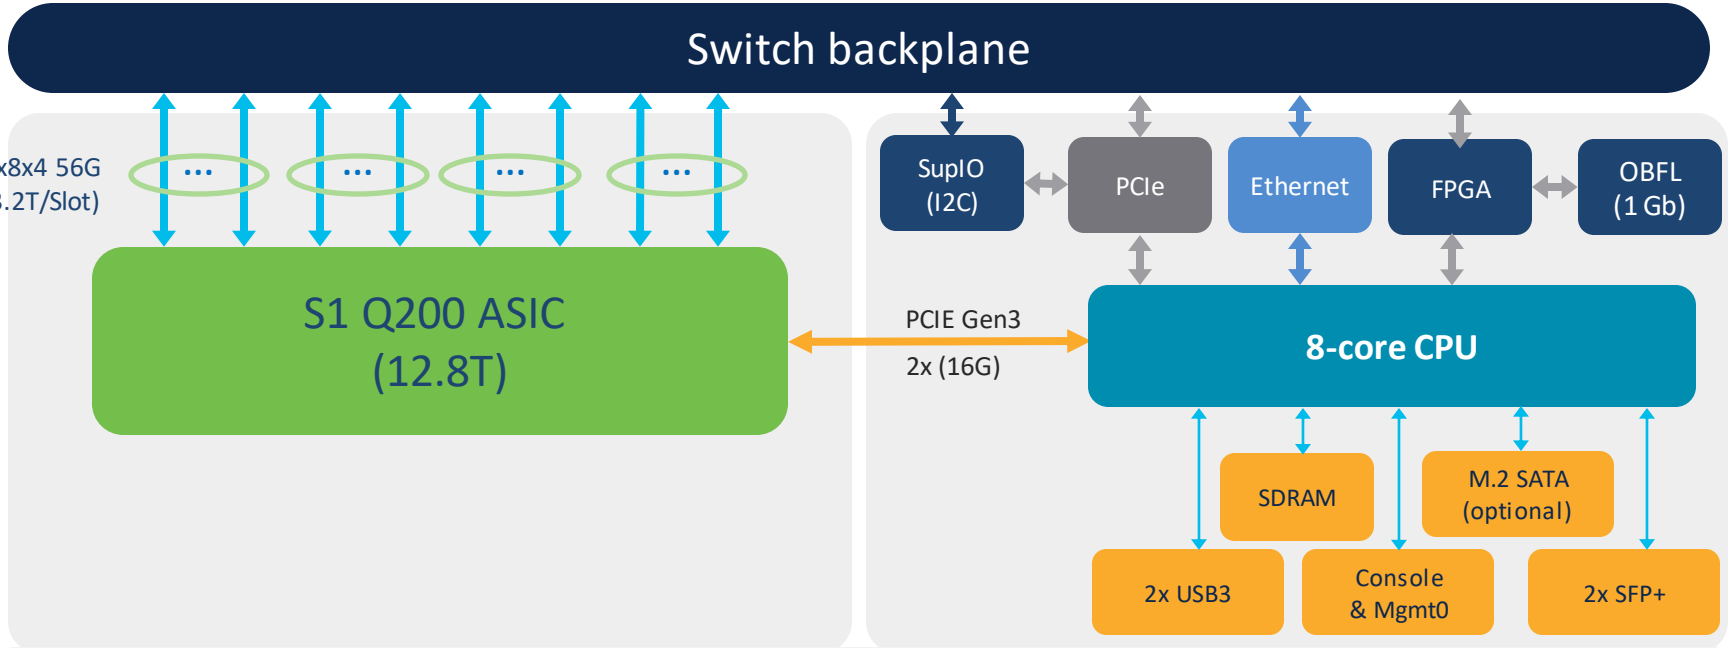

C9600X: Introducing 1st **400G** Line-Card in Campus

Shipping

Catalyst 9606R

C9600-LC-24C

C9600-LC-48YL

C9600-LC-48S

C9600-LC-48TX
(Multigigabit)

NEW

Combo LC

C9600-LC-40YL4CD

40x SFP56 +
2x QSFP56 + 2x QSFPDD

3.2 Tbps with SUP2

1.2 Tbps with SUP1

Provides for **Uplink requirement** without using additional slot

Dual personality Line-Card

Flexible speed support

Lower total BoM cost, save on additional LC

C9600-LC-40YL4CD

Gen2 Combo Line-Card – 40xSFP56 + 2xQSFP56 + 2xQSFPDD

3.2 Tbps

Target FCS
Q1CY22
Software
IOS-XE 17.7.1

- **3.2 Tbps** With Gen2 Sup
  - 40x SFP56 ports - 10/25/50*GE
  - 2x QSFP56 ports - 40/100/200*GE
  - 2x QSFPDD ports - 40/100/200*/400GE
- **1.2 Tbps With Gen1 Sup**
  - 40x SFP28 ports - 1/10/25GE
  - 2x QSFP28 ports - 40/100GE
- 8x **CDR5M** PHYs
  - MACSec, WAN-MACSec, ClearTag v3.4
  - IEEE 1588 & PTPv2*
  - Hitless MUX (HMUX)
- Various Breakout & QSA support (QSFP ports)*

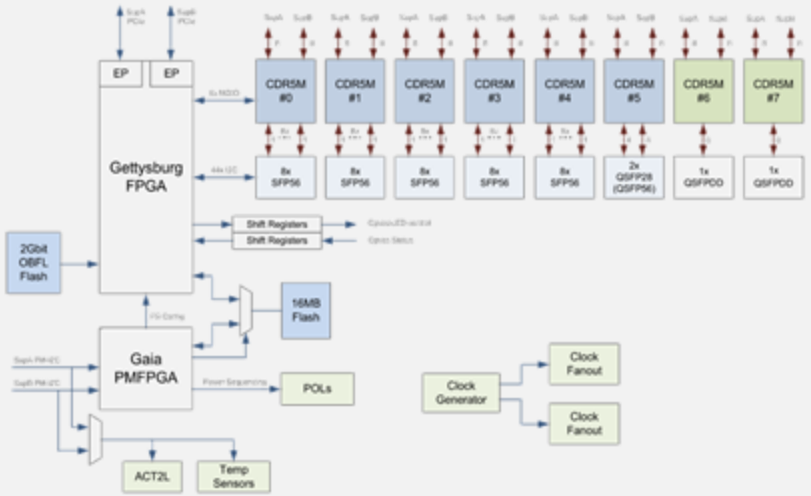

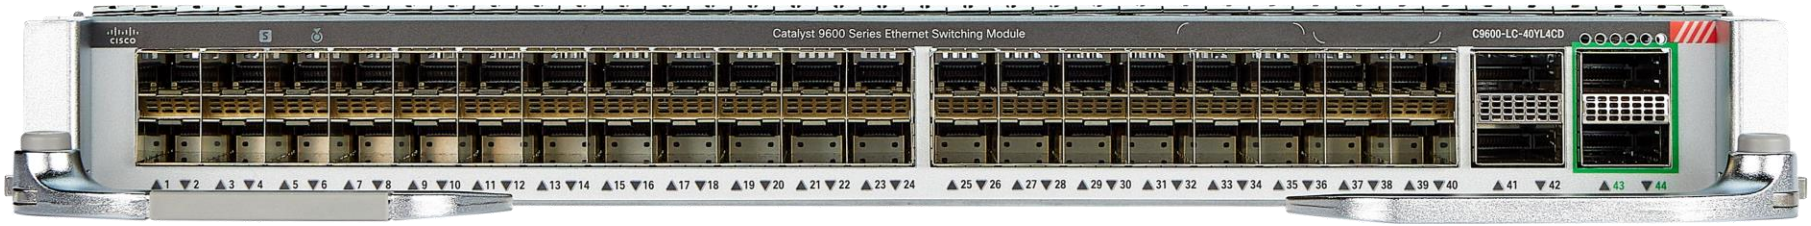

* Hardware capable. Roadmap (not committed)

C9600-LC-40YL4CD Ports and Speeds Support

with **Sup2**: **3.2Tbps**

40x 10/25/50*GE + **2**x 40/100/200*GE + **2**x 40/100/200*/400GE

IOS-XE 17.7.1

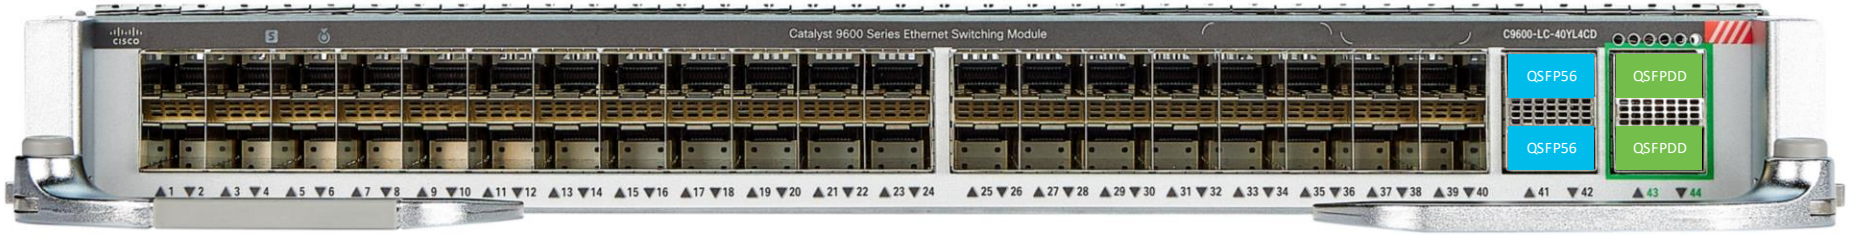

with **Sup1**: **1.2Tbps**

40x 1/10/25GE + **2**x 40/100GE

IOS-XE 17.8.1

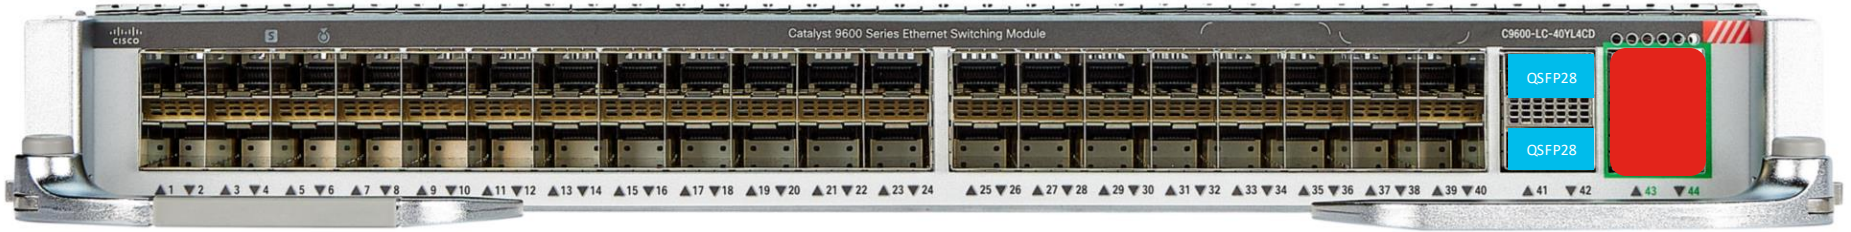

* Roadmap (not committed).

Gen1 Line-Cards Support with SUP2

Centralized Architecture

- **Gen1 Line-Cards supported***
- Only PHYs on the Line-Cards
- All forwarding on the Supervisor (ASIC)

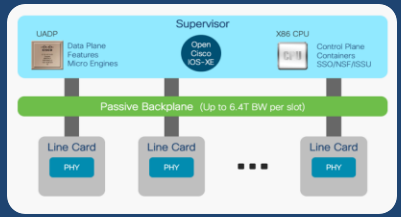

Additional Bandwidth ☺

- C9606 backplane traces support up to 56G PAM4
- Gen1 Line Cards now support up to 2.4T per Slot
  - **24 x 100G** QSFP on LC-24C
  - **48 x 50G**** SFP on LC-48YL

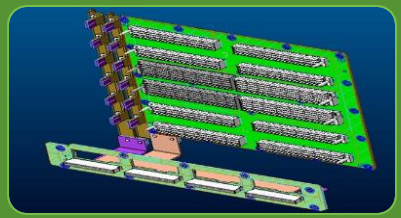

No MACsec support ☹

- Q200 does not have onboard Crypto engine
  - Gen2 LCs use newer CDR5M PHY for MACsec
- UADP has onboard MACsec engine
  - Gen1 LCs use older CDR4 PHY (not MACsec capable)

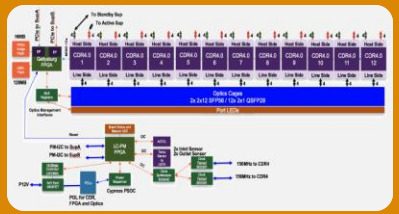

* C9600-LC-48S is not supported with Sup2

** Roadmap (not committed)

C9600 Line Card – Supervisor Support Matrix

	SUP 1	SUP 2
C9600-LC-24C	✓ 24x 40G or 12x 100G	✓ 24x 40G & **100G** (No MACsec)
C9600-LC-48YL	✓ 48x 1/10 & 25G	✓ 48x 10/25 & **50G*** (No MACsec, No 1G)
C9600-LC-48TX	✓ 48x 1/2.5/5 & 10G (mGig)	✓ 48x 10G (No MACsec, No 1/2.5/5G)
C9600-LC-48S	✓ **48x 1G SFP**	✗
C9600-LC-40YL4CD	✓ 40x 1/10 & 25G + 2x 40 & 100G	✓ 40x 10/25 & **50G*** + 2x 40/100 & **200G*** + 2x 40/100/**200*** & **400G** MACsec & **WAN-MACsec** (No 1G)

TIP

SUP2 does not support 1GE or below speeds

If 1GE downlinks are required, position SUP1

* Roadmap (not committed)

FYI

C9600X and C9600 - Physical

Cisco C9600X (S1 Q200)

Cisco C9600 (UADP 3)

	Sup2	Sup1
CPU (number)	**2.7GHz 8C Intel** (BDW-NS)	2.4GHz 8C Intel (BDW)
DRAM (type)	**32 GB** (DDR4)	16 GB (DDR4)
ASIC (number)	**Q200** (1x)	UADP3 (3x)
Capacity (chassis)	**12.8 Tbps** (full-duplex)	4.8 Tbps (full-duplex)
Capacity (per slot)	**3.2 Tbps** (full-duplex)	1.2 Tbps (full-duplex)
10G max	**256** (40x4 + 24x4^*)	192 (48x4)
25G max	**256** (40x4 + 24x4^*)	192 (48x4)
50G max	**256** (40x4 + 24x4^*)	--
40G max	**96** (Sup2 can support **128**)	96 (24x4)
100G max	**96** (Sup2 can support **128**)	48 (12x4)
400G max	**8** (4x2)	--

* Roadmap (not committed)
^ Breakout Cable ports

FYI

C9600X and C9600 – Features and Scales

Cisco C9600X (S1 Q200)

Cisco C9600 (UADP 3)

	Default	Maximum (Custom)	Default	Maximum (Custom)
MAC Addresses	**128K**	256K	**80K**	128K
IP Host Routes	**128K**	256K	**80K**	128K
Multicast L2 groups	**16K**	64K*	**16K**	48K
Multicast L3 routes	**32K**	64K*	**32K**	48K
IP LPM Routes	**2M**	2M	**212K**	256K
MPLS Labels	**256K**	512K	**32K**	64K
SGT/OG Labels	**32K**	64K	**32K**	64K
NAT* Sessions	**16K***	128K*	**3K**	16K
Sec ACL Entries	**8K**	10K*	**12K**	27K
QoS ACL Entries	**8K**	10K*	**8K**	21K
PBR* ACL Entries	**8K***	10K*	**3K**	16K

* Roadmap for C9600X (not committed)

C9600 and C6800 Scale

Features	C9600– Sup2	C6K-Sup6T-XL	C9600– Sup1	C6K-Sup6T
Switching capacity	12.8T	6T	9.6T	6T
Forwarding rate (IPv4/IPv6)	8 Bpps 8 Bpps	780Mpps 390Mpps	3 Bpps 3 Bpps	780Mpps 390Mpps
MAC addresses	256,000*	128,000	82,000*	128,000
LPM/host routes (IPv4/IPv6)	2,000,000* 1,000,000*	1,000,000 500,000	212,000* 212,000*	256,000 128,000
Multicast routes	64,000*	128,000	32,000*	128,000
Security ACLs	Shared with QoS & PBR 10,000*	Shared with QoS 256,000	27,000*	Shared with QoS 64,000
QOS ACLs	Shared with Security & PBR 10,000*	Shared with Security 256,000	16,000*	Shared with Security 64,000
Flexible NetFlow (per ASIC)	2,000,000 (Sampled)	1,000,000/ASIC	96,000*	512,000
VLAN ID	4,000	4,000	4,000	4,000
Spanning Tree instances	4,000	4,000	4,000*	4,000

* Depends on SDM template

Customizable SDM Template

Switch Database Management (SDM) template

User-customizable template
Allows customizable
ACL TCAM resources

Cisco® Catalyst® 9600X and 9500X Series

Default template
Maximizes system resources for Layer 3 unicast and multicast **routes**

Custom template
Provide flexibility for customizing TCAM space for specific requirements

Silicon One Q200 SDM template – 17.7.1

Features	Customizable	DEFAULT (core + edge)	Custom (min to max @ step)	
MAC addresses		**128,000**	32,000 to **256,000**	@ 1000 step
Host routes (ARP/NDP)		**128,000/64,000**	32,000 to **256,000**	@ 1000 step
Layer 2 multicast entries (IGMP/MLD)	FCS+	**16,000/8000**	0 to **64,000**	@ 1000 step
Layer 3 multicast routes (IPv4/IPv6)	FCS+	**32,000/16,000**	0 to **64,000**	@ 1000 step
ACL compression (SGT, DGT, OGID/v6)		**32,000/16,000**	0 to **64,000**	@ 1000 step
MPLS labels		**256,000**	0 to **512,000**	@ 1000 step
Reserved (PBR/NAT)	FCS+	**16,000/8,000**	0 to **256,000**	@ 1000 step
	CEM	**608,000** (288,000 for LPM)		
Layer 3 unicast routes (IPv4/IPv6)	FCS+	**2 million/ 1 million**	1 million to 2 million	@ 1 million step

Features	Customizable	DEFAULT (core + edge)	Custom (min to max @ step)	
Security ACL (IPv4/IPv6)	FCS+	**8000/4000** shared*	0 to 11,000/5000	@ 1 step
Quality of service (IPv4/IPv6)	FCS+	**8000/4000** shared*	0 to 11,000/5000	@ 1 step
Policy-based routing (IPv4/IPv6)	FCS+	**8000/4000** shared*	0 to 11,000/5000	@ 1 step
Lawful intercept (IPv4/IPv6)	FCS+	**1000** (2x 512) reserved	1000 to 5000/2500	@ 1 tap (2 ACE)
LPTS, EPC, FSPAN, NFL (IPv4/IPv6)	FCS+	**1000** (2x 512) reserved	1000	@ 1 step
	TCAM	**10,000** (2000 for LPM)		

* Shared is an unreserved space, first come, first served per feature.

FIB Allocation Examples

Feature	Customer 1 (L2 focus)	Customer 2 (L3 focus)
MAC addresses	**256K**	32K
Host routes (ARP/NDP)	32K	**256K**
Layer 2 multicast entries (IGMP/MLD)	64K	16K
Layer 3 multicast routes (IPv4/IPv6)	16K	64K
ACL compression (SGT, DGT, OGID/v6)	64K	64K
MPLS labels	32K	128K
Reserved (PBR/NAT)	15K	48K
Total Resources	**608K**	

Layer 3 unicast routes (IPv4/IPv6)	**1M/500K**	2M/1M

16K – 16384
32K – 32768

Application Hosting

Enhanced app-hosting infrastructure on Catalyst 9500

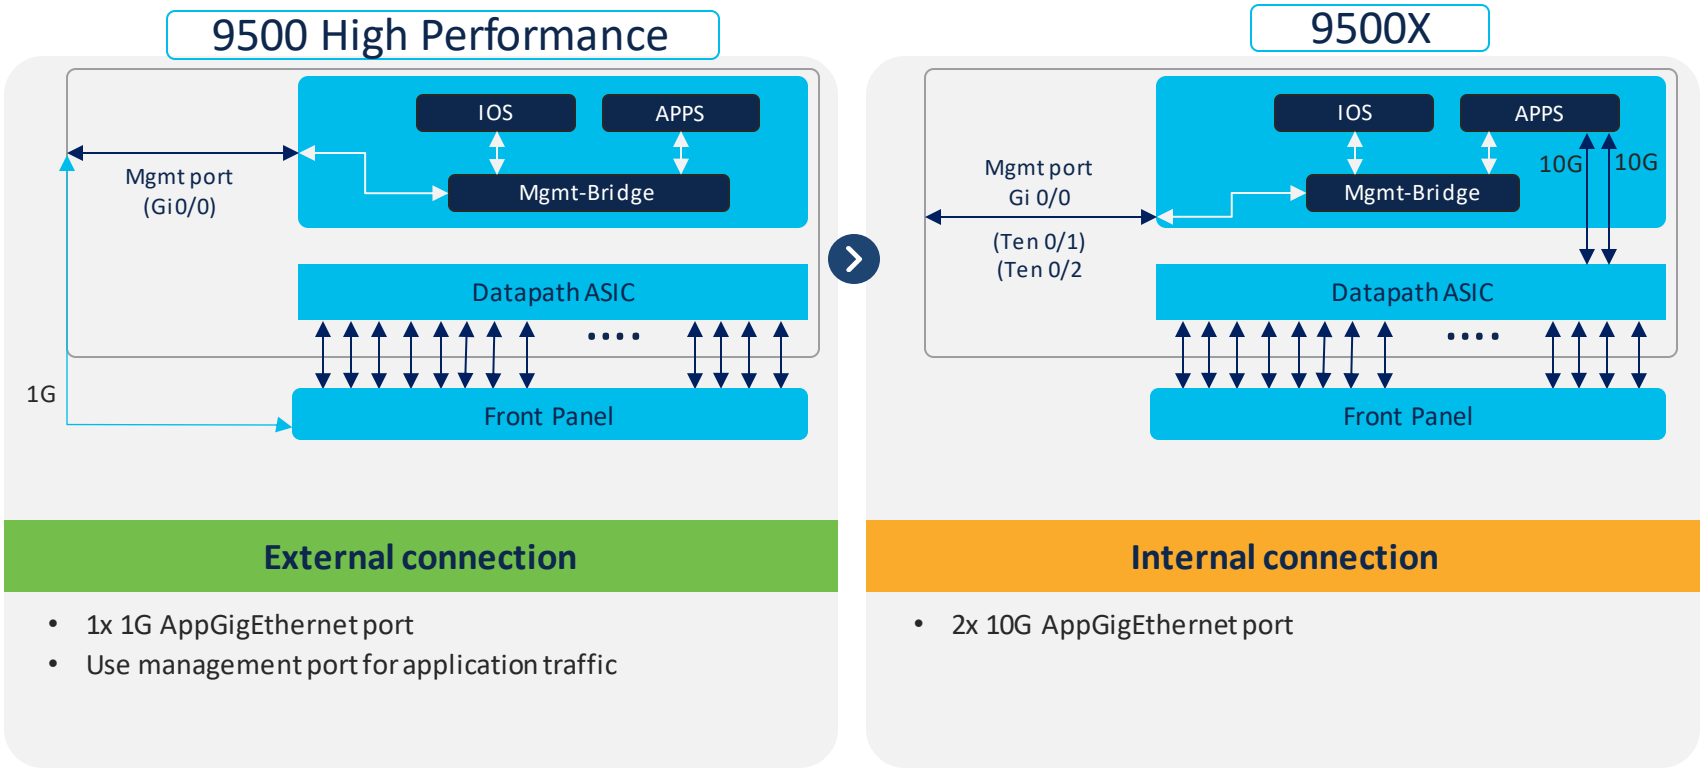

External connection

- 1x 1G AppGigEthernet port
- Use management port for application traffic

Internal connection

- 2x 10G AppGigEthernet port

Enhanced app-hosting infrastructure on Catalyst 9600

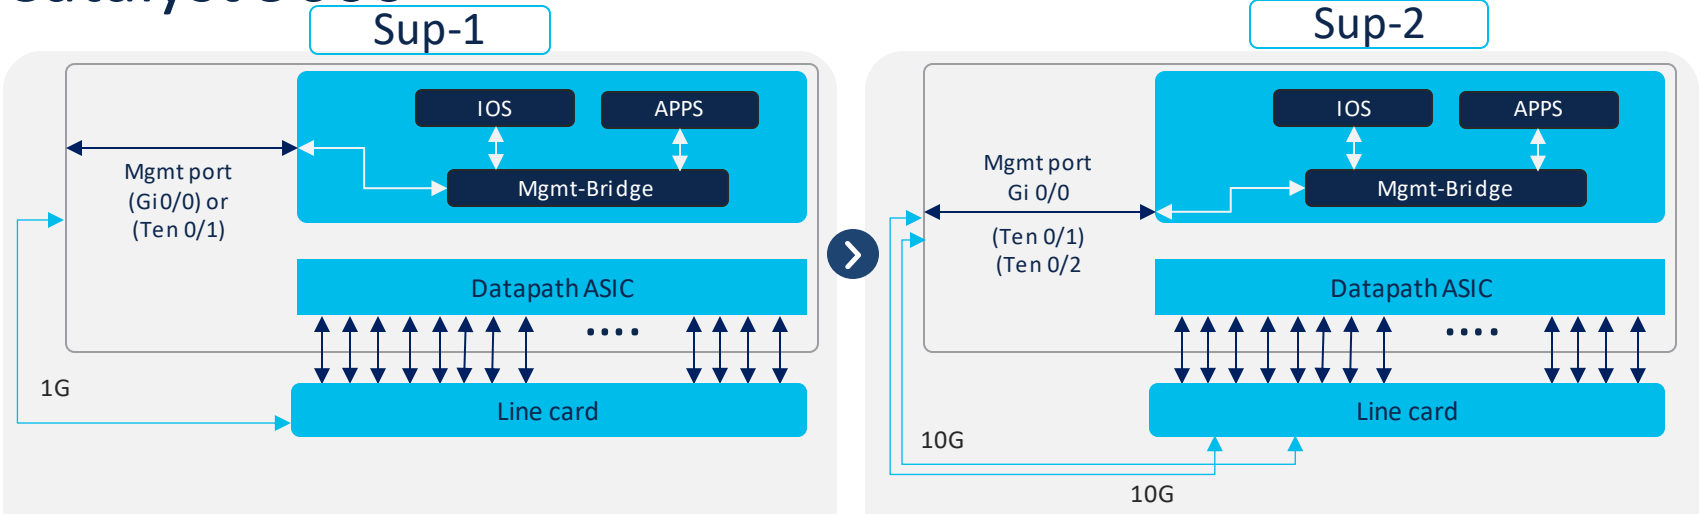

External connection

- 1x 1G AppGigEthernet port
- Use one of the management ports for application traffic

External connection

- 2x 10G AppGigEthernet port
- Dedicated management port (Gi 0/0)
- Additional management ports (Te 0/1 and Te 0/2) can be used for application traffic

High Availability

High availability
Protect business continuity

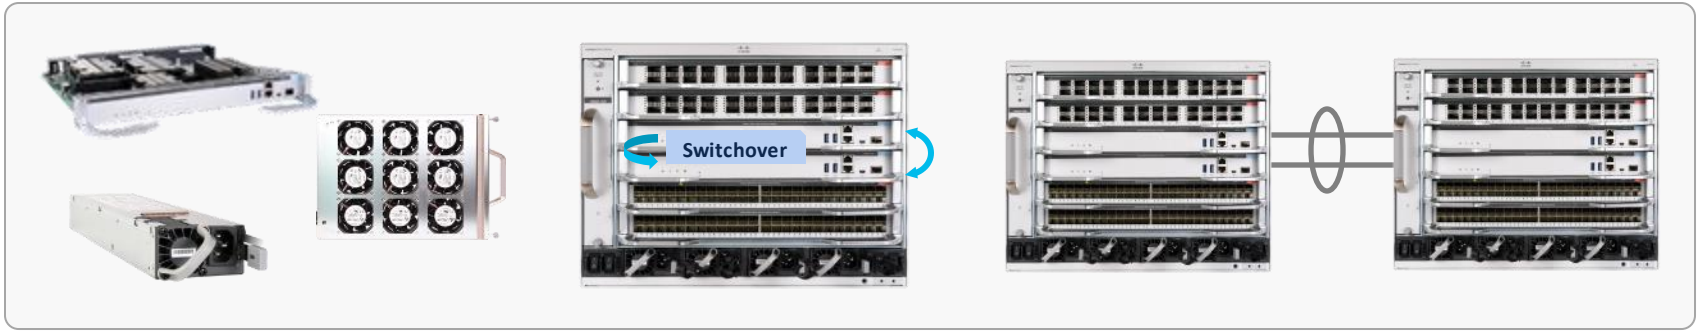

Physical redundancy	Stateful Switchover (SSO)	Non-Stop Forwarding (NSF)	In-Service Software Upgrade (ISSU)	StackWise®-Virtual*
Redundant hardware • Redundant power supplies • Redundant fan in the fan tray • Redundant supervisors	**Sub-second failover** • Between supervisors within chassis (<5ms) • Between chassis with StackWise-Virtual*	**Resilient L3 topologies** • NSF support for OSPF, EIGRP, ISIS, BGP	**Minimize upgrade downtime** • SMU • ISSU* • GIR*	**Infrastructure resilience** • Multi-chassis EtherChannel (MEC) provides hardware-based failover

* Roadmap for Sup 2

Security & ACL

Security Access Control List (ACL)

Hardware Pattern Matching

RECAP

RACL

IP-Based Router ACL

PACL

Port-Based ACL

Flexible **resource sharing** for **different ACL types** in TCAM

VACL

VLAN-Based ACL

GACL

Group-Based ACL

RECAP

Access Control List Terminology

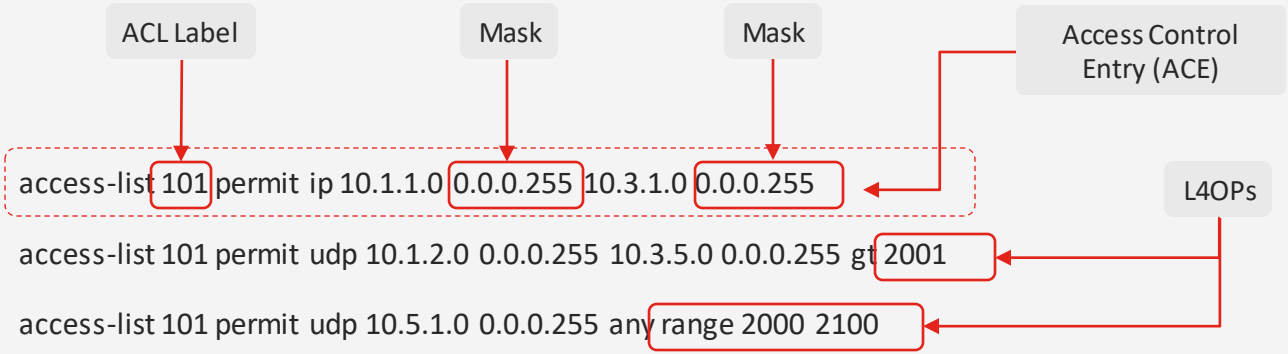

What is an L4OP?	**GT** (greater than)	**LT** (less than)	**NE** (not equal to)	**Range** (from - to)
What is **NOT** an L4OP?	**EQ** (equal to)			
VCU (Value Comparison Unit)	**Range** -> 2 VCUs | Source & Destination L4OPs -> Separate VCUs			

Traditional ACL vs Group-Based ACL

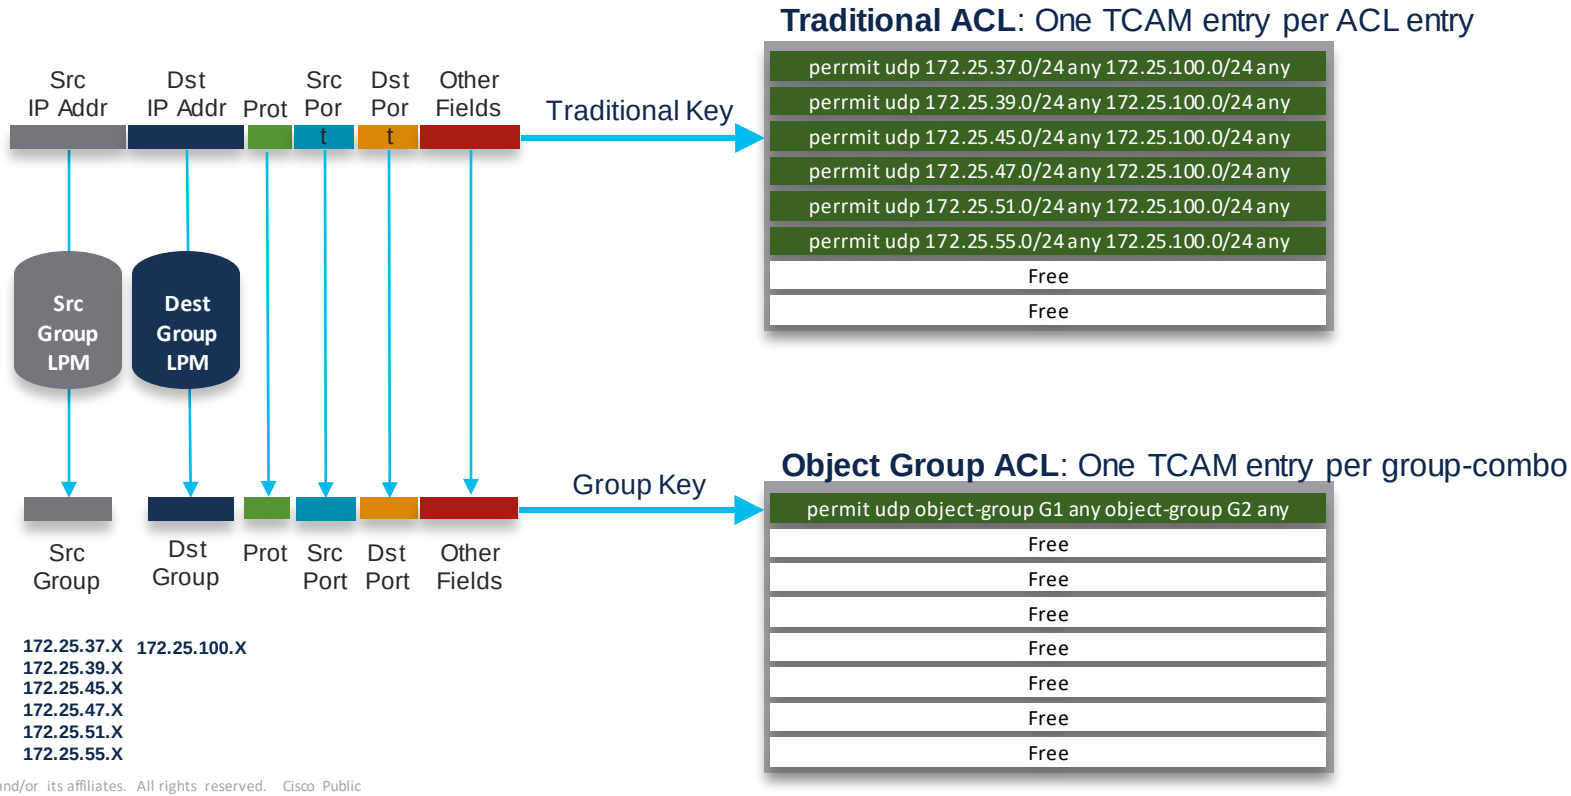

Why OGACL/SGACL in Campus Core?

Object-Group ACLs for IP

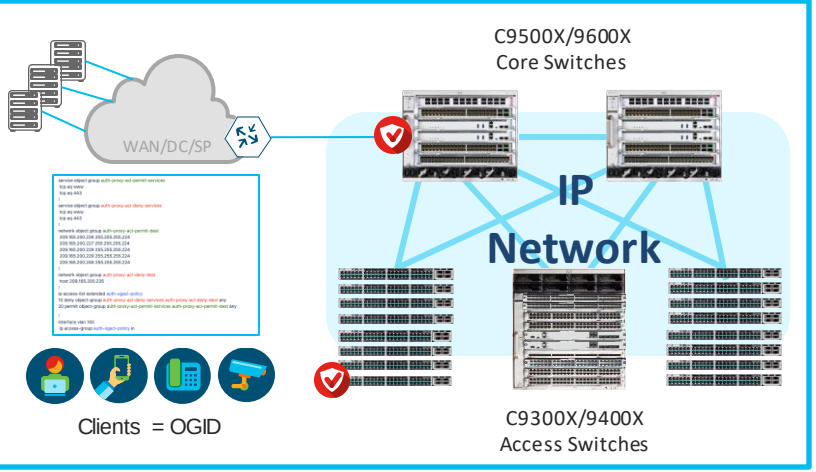

Object-Groups map IP/mask to Labels in CEM

- User defines IP/masks to simple OG name
- OGID labels are stored in Exact Match table

OGACL ACEs take minimal space in ACL TCAM

- Only the Permit/Deny ACEs in TCAM
- OGACLs with same ACEs can reuse entries

Scalable-Groups for SDA

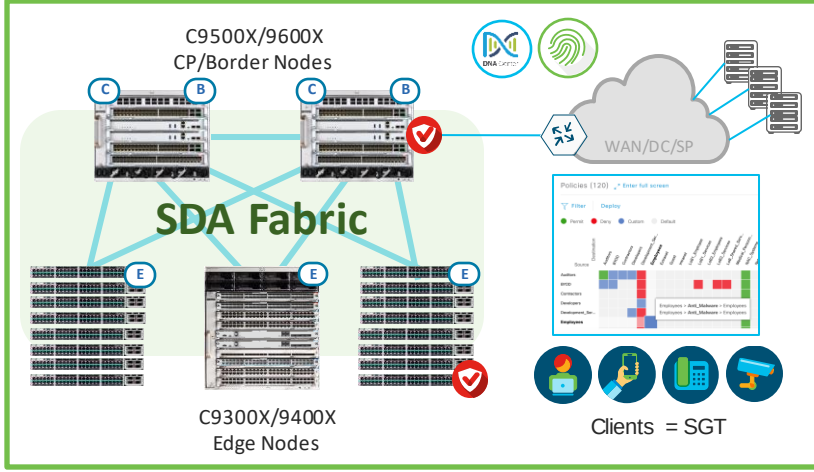

Scalable-Groups map IP/mask to Labels in CEM

- ISE/DNAC defines IP/masks to simple SG name
- SGT labels are stored in Exact Match table

SGACL ACEs take minimal space in ACL TCAM

- Only the Permit/Deny ACEs in TCAM
- SGACLs with same ACEs can reuse entries

WAN MACsec overview and use cases

Enabled in hardware on Catalyst 9000 Switches

MACsec

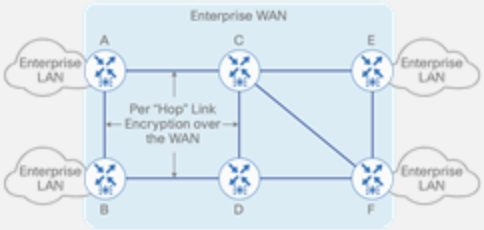

- Hop-by-hop encryption
- Directly connected Layer 2 links **only**
- Secure any direct link in campus/ WAN, including DWDM
- 256/128-bit AES encryption

Supported on **all Catalyst® 9000** models
All ports and speeds simultaneously

WAN MACsec

- End-to-end encryption across Layer 2 Ethernet WAN service
  - EoMPLS/VPLS, Q-in-Q
  - Point-to-point or point-to-multipoint
- Applications include secure site interconnect, DCI, storage replication
- 256/128-bit AES encryption

Catalyst 9600X

Industry's first

Catalyst 9500X

Supported on **Silicon One™** platforms*
All ports/speeds (up to 400G) simultaneously

Catalyst 9400X

Catalyst 9300X

Roadmap support access **Catalyst 9300X and 9400X**
Aggregate speed up to 100G

* Requires Release 17.8.1 and HSEC key ($1 PID).

QoS

Catalyst 9000 Series

QoS & Buffering Technologies

QoS Features

- Trust / Conditional Trust
- Classify Traffic
- Police Traffic
- Mark / Conditional Remark

Queuing Features

- Prioritize strict traffic (SPQ)
- Schedule traffic based on weight (WRR)
- Shape the traffic rate (SRR)
- Manage congestion (WRED/WTD)
- Extra buffering for traffic bursts

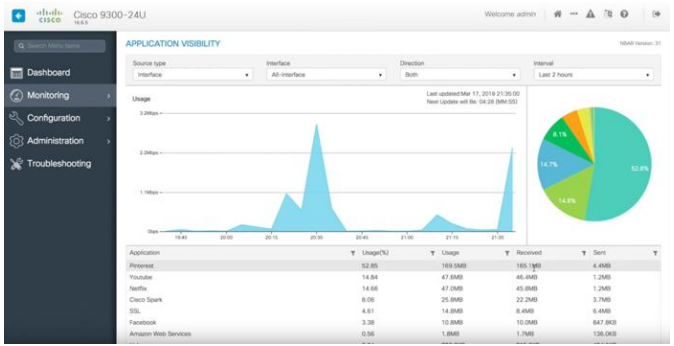

Head of line blocking

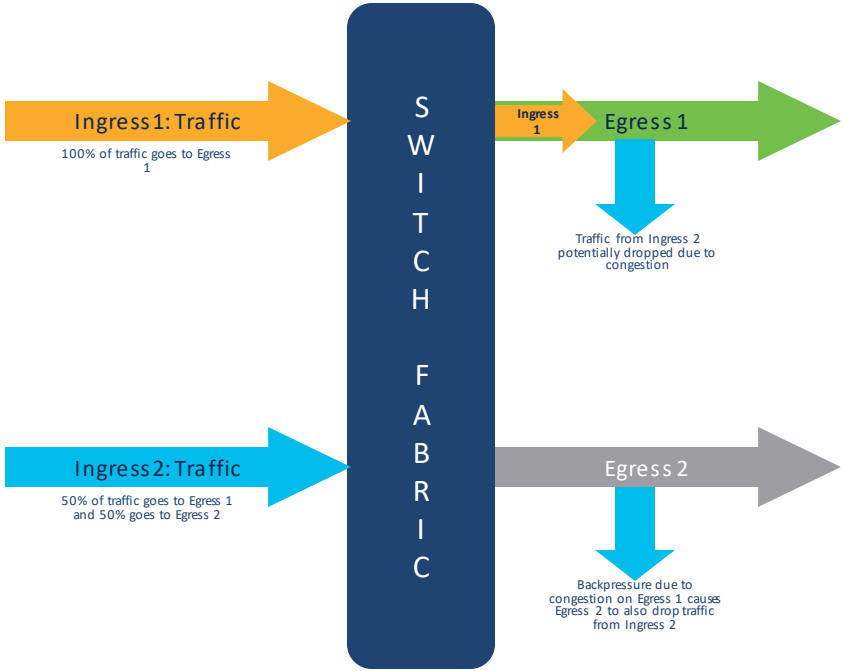

VoQ Architecture

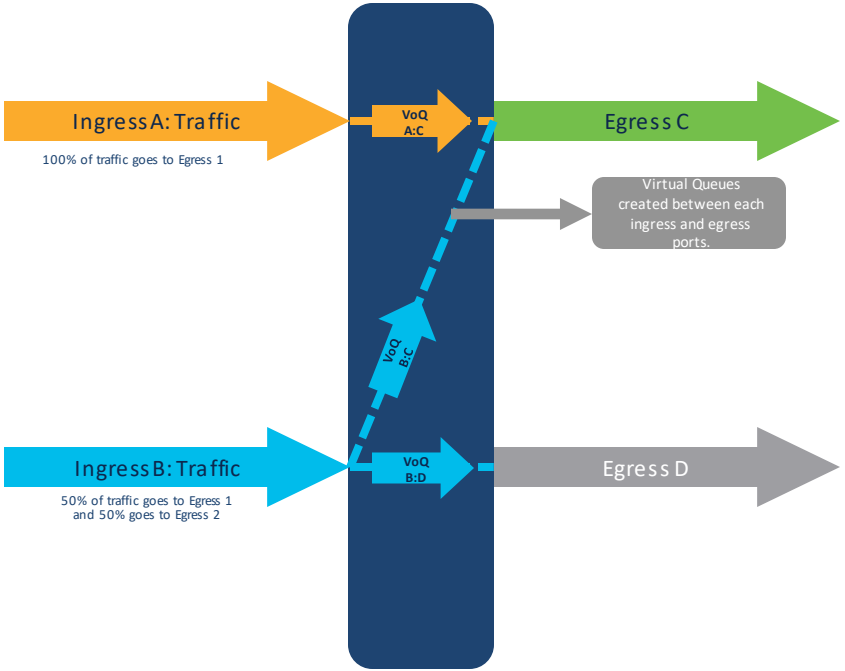

Why VoQ QoS in Campus Core?

No Head-of-Line Blocking

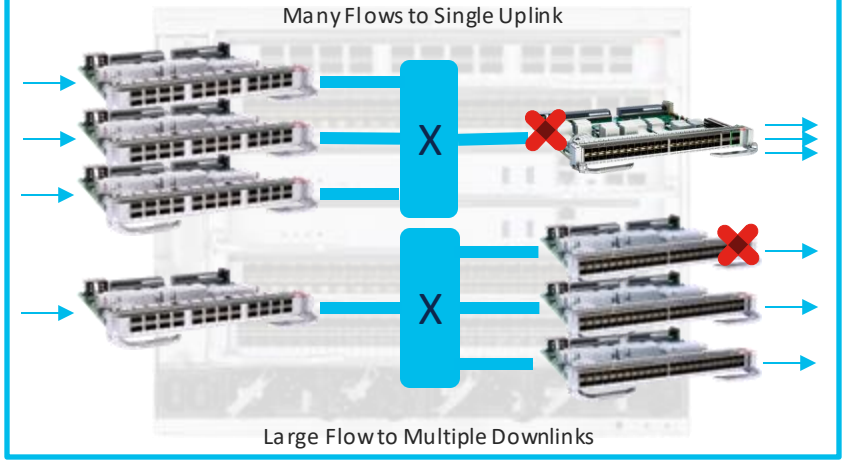

Many Flows to a Single Uplink

- Common on (expensive) WAN/Edge uplinks
- Even if bandwidth available, buffers can fill up

Large Flow to Multiple Downlinks

- Common for Multicast & Broadcast traffic
- One slow receiver can penalize other ports

Local vs. HBM Buffers

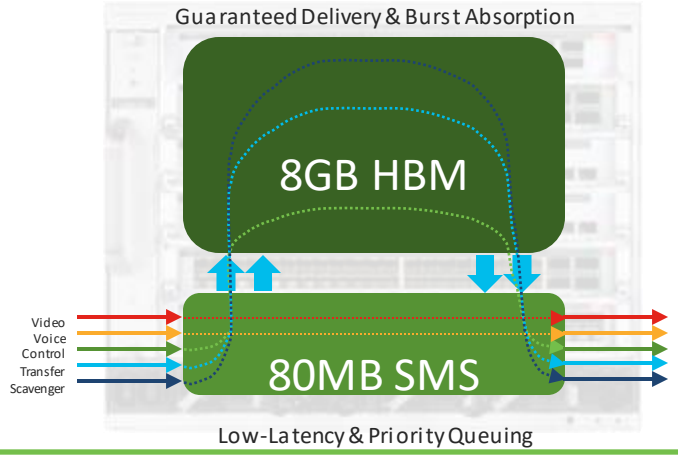

Low-Latency Local Shared Memory Buffers

- Voice & Video are very latency-sensitive
- Multiple levels of Strict Priority Queuing

Deep High-Bandwidth Memory Buffers

- Guarantee delivery of session-oriented flows
- Reserve buffers to absorb occasional bursts

S1 Q200 - Standalone SOC Architecture

Generic Unicast Packet Walk

1 **RX IFG** - packet received through MAC, stored in IFG-RXQ

2 **RX NPU** – packet elements processed by NPU to generate destination port

3 **SMS Write** - packet elements written into memory buffer (buffers linked if needed), generate packet-descriptor (PD)

4 **VOQ** – PD stored in VOQ according to destination port

5 **Credit Request/Grant**

6 **PD Switch** – PD switched from VOQ to an OQ at the destination slice

7 **OQ** – store PD in destination port OQ

8 **SMS Read** – read packet elements from memory, free buffers

9 **TXNPU** – edit packet elements

10 **TX IFG** – transmit packet

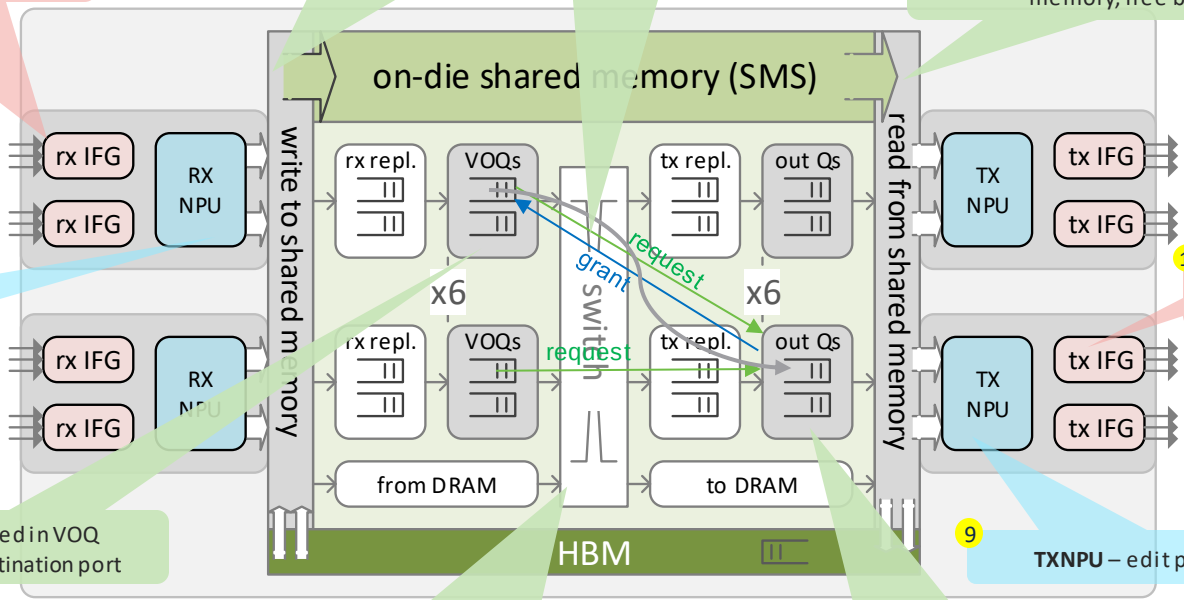

S1 Q200 - ASIC mapping of QoS features

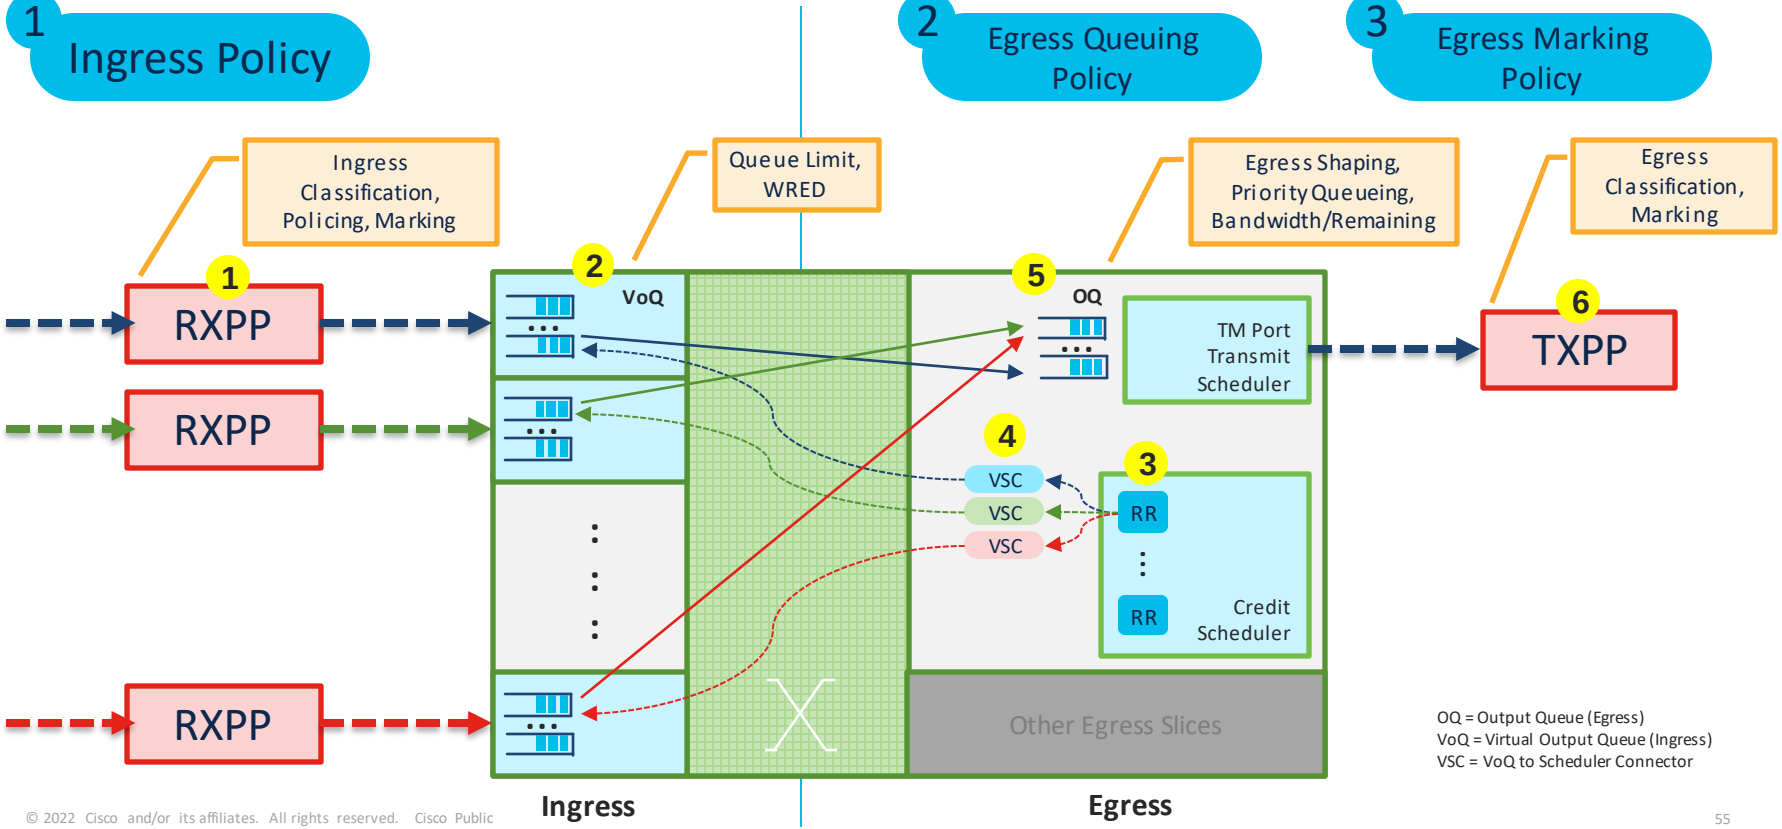

OQ = Output Queue (Egress)
VoQ = Virtual Output Queue (Ingress)
VSC = VoQ to Scheduler Connector

Netflow

Silicon One vs UADP

Capabilities

- Filtering of traffic to be sampled
- Random sampling and selection of one out of N filtered packets
- Mirroring of the selected packets along with their NPU context to CP CPU.
- 1 dedicated CPU core for Netflow.

Limitations vs UADP

- UADP builds and updates flow records in H/W cache.
- Entries moved from H/W to S/W cache for aggregation.
- Aggregated entries exports to collector.

S1 Netflow

- Sampled Netflow.
- ASIC samples packets on configured interfaces.
- Selected packet passed to CP CPU for parsing.
- Parsed data populated in flow cache and exported in required format.

S1 Q200 - Sampled FNF Collection

Flexible NetFlow data collection happens at configurable "**sampling rate**"

- Sampling rate:
  - 1 out of (2-16384). **Default is 1:1024**
  - Balance of Accuracy vs. CPU Load
- Sampling mode:
  - Deterministic – a fixed packet within sample (e.g., always 10th)
  - Random – a random packet within sample (e.g., 4, 13, 67, etc.)

Each flow **checked during packet parsing**

- Is it a new flow? If not new - within sample rate?

S1 Q200 sends a copy of the packet to a **dedicated X86 CPU Core**

- IOSXE (CPP) software builds a FNF cache entry

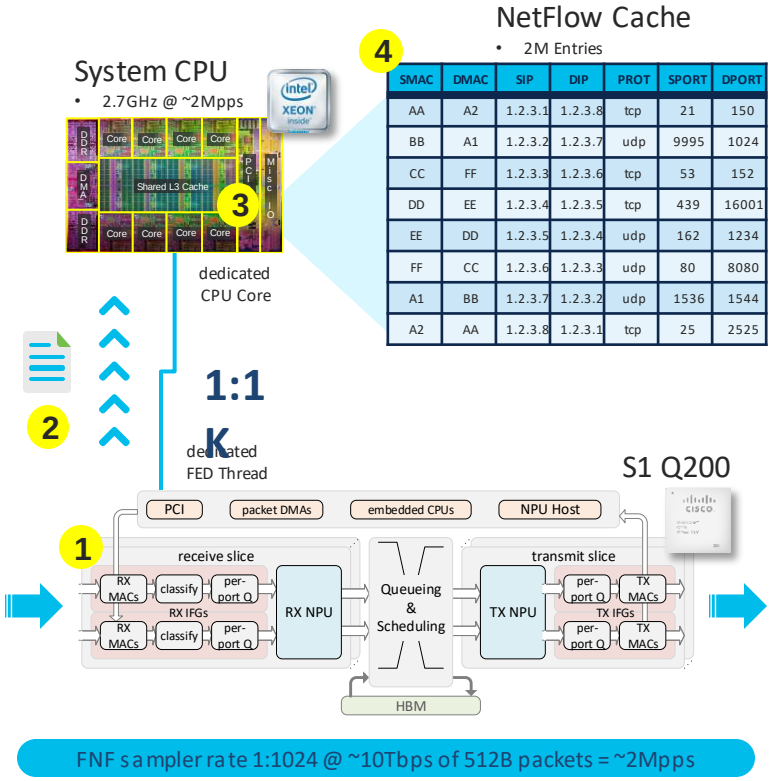

SMAC	DMAC	SIP	DIP	PROT	SPORT	DPORT
AA	A2	1.2.3.1	1.2.3.8	tcp	21	150
BB	A1	1.2.3.2	1.2.3.7	udp	9995	1024
CC	FF	1.2.3.3	1.2.3.6	tcp	53	152
DD	EE	1.2.3.4	1.2.3.5	tcp	439	16001
EE	DD	1.2.3.5	1.2.3.4	udp	162	1234
FF	CC	1.2.3.6	1.2.3.3	udp	80	8080
A1	BB	1.2.3.7	1.2.3.2	udp	1536	1544
A2	AA	1.2.3.8	1.2.3.1	tcp	25	2525

FNF sampler rate 1:1024 @ ~10Tbps of 512B packets = ~2Mpps

S1 Q200 - Sampled FNF Export

NetFlow data export is based on aging timers

Dedicated X86 CPU Core builds UDP packets for the FNF records that are aged-out

- Active timer
- Inactive timer

Dedicated X86 CPU Core sends the UDP packets to the configured Export IP address

- Up to 4 FNF flow exporters
- Support for NFv9 or IPFIX format
- Support for IPv4 or IPv6 exports

TIP Same export process/method on all platforms!

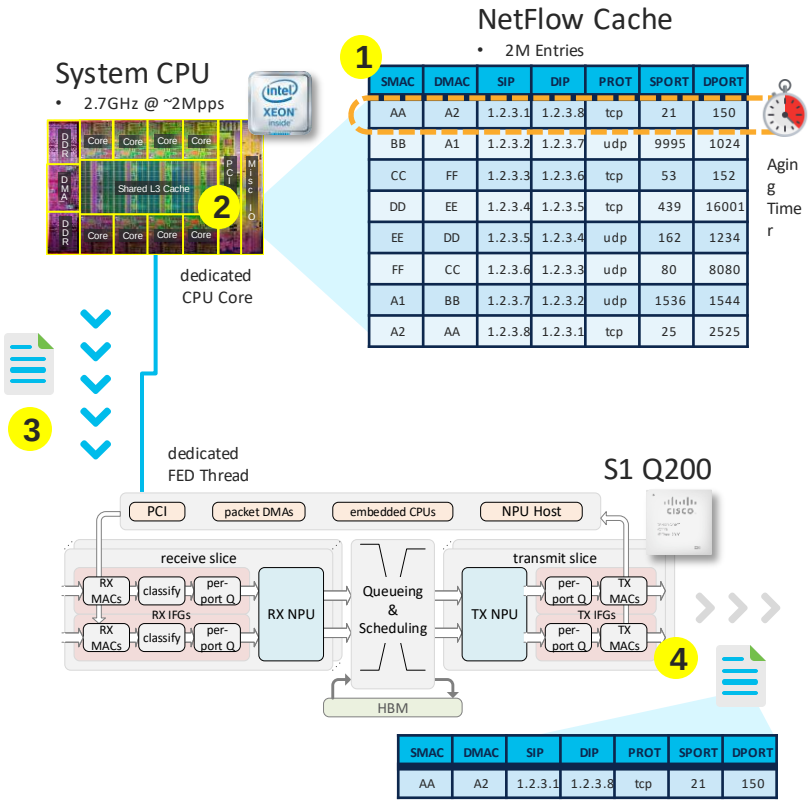

Silicon One sFlow

Overview

- Added capability on Catalyst 9500X and 9600X in addition to Sampled Netflow.
- Like Sampled Netflow, sFlow randomly samples one packet out of N packets on target interface.
- No cache associated with sFlow.
- Packet header for sampled packet along with packet metadata encoded into a datagram and sent collectors.
- Supported with DNA Advantage License

Netflow Advantages

- Per flow export.
- Consumes less data – packet data parsed and sent.
- More options for configuration.
- Less load on collectors.

sFlow disadvantages

- Per packet export.
- Consumes more data – full packet headers sent.
- Limited configurable options.
- More load on collectors.

Catalyst 9500X and 9600X will be the first IOS-XE based platforms to support sFlow

Why Sampled FNF in Campus Core?

ID @ Access – Monitor @ Core

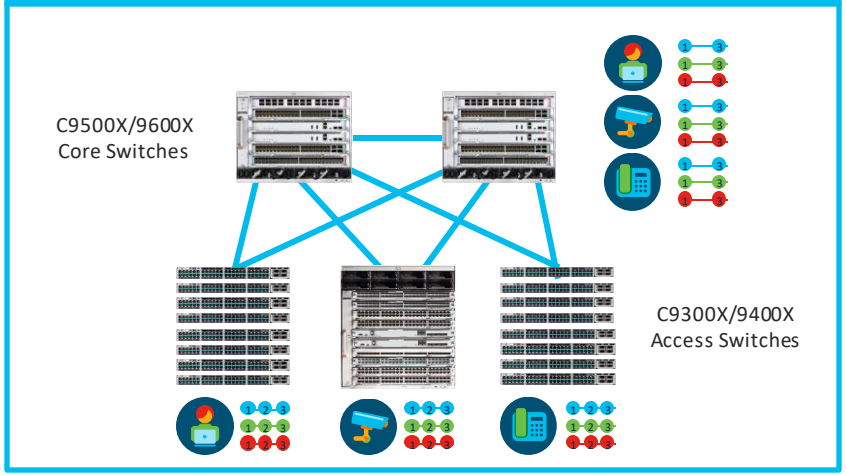

Detailed (1:1) flow identification at the Access

- Better to ID flows as they enter the network
- Full accounting of every client/flow (ETA/AVC)

Aggregate (1:1K) monitoring of flows at the Core

- Just need to monitor the overall network usage
- Adjust sample rates to balance scale & load

Campus-wide Scale

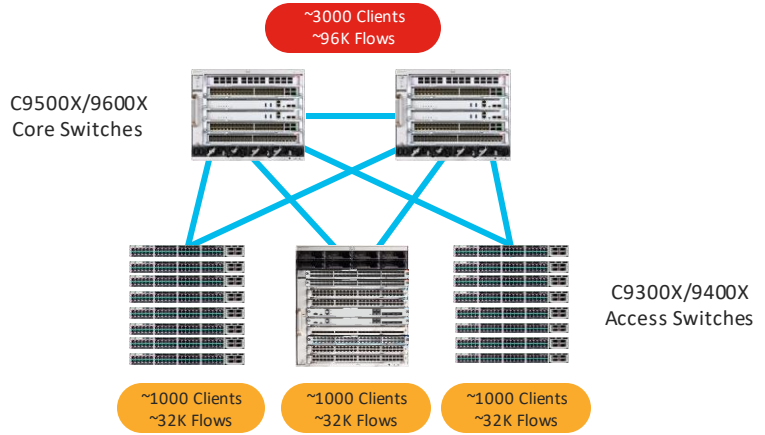

Low-Moderate scale at the Access

- Fewer number of connected clients/flows
- Average ~1K clients x ~32K flows per Access

Medium-High scale at the Core

- Need to aggregate all clients/flows (# Access x 32K)
- Adjust cache aging to increase overall scale

Catalyst 9500X & 9600X

Summary

Catalyst 9500X & 9600X – Things to know

Next Generation Core + Edge Switching with Silicon One™ Q200

C9600X-SUP-2 + LC-40YL4CD

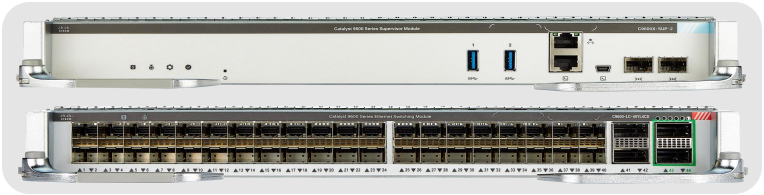

C9500X-28C8D

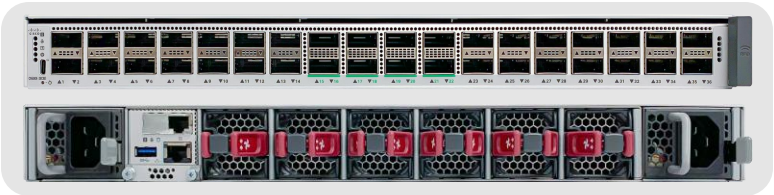

Technology	Brief Description	Diffs from UADP 3.0
Large LPM Table	• **Up to 2M IPv4 or 1M IPv6 (hash efficiency is about 80%)** • **Dedicated Memory for LPM**	• Up to 256K IPv4/IPv6 with Custom SDM template • LPM and other features share 416K ASIC memory
Large MAC Table	• **Up to 256K MAC entries Custom SDM template**. • Shared with other features	• Up to 128K MAC entries with Custom SDM template • MAC shared with LPM, etc. in same 416K ASIC memory
VoQ QoS + HBM	• **Q200 has 80MB local (low-latency) + 8GB HBM (High Bandwidth Memory) buffer memory.** • **Q200 uses a Virtual Output Queue (VoQ) architecture.** All queuing and policing policies applied on Ingress.	• Max 36MB unified buffer memory per ASIC • Supports both Ingress/Egress queuing & policing
OGACL & SGACL	• **8K IPv4, 4K IPv6 ACL TCAM entries.** • **Object-Group & Security-Group ACLs use CEM to map IP-to-Group label, TCAM only uses L4 ACEs**. (OG/SG ACL design is optimal for layer 3 environment).	• 64K ACL TCAM entries per ASIC • Object-group expand into the TCAM space
LAN & WAN-MACsec*	• **Q200 does not have built-in crypto engine.** • **C9500X & C9600X-LC uses new CDR5M PHY (400Gbps Full-Duplex)**. CDR5M provides line-rate (8x 400G = 3.2T) 802.1ae (LAN) MACsec and WAN-MACsec.	• UADP3 has built-in MACsec crypto (speed of ASIC) • UADP3 only supports LAN MACsec (no WAN-MACsec)
Flexible NetFlow*	• **Q200 does not have built-in Flow Cache memory (no hardware-based Netflow).** • **C9500X & C9600X uses new Software-based FNF (≤ 2M entries), with a dedicated CPU core (~2Mpps)**. FNF sampler rate 1:1000, ~10Tbps of 512-Byte packets = ~2Mpps.	• UADP3 has built-in (HW) FNF, max 64K entries per ASIC • FNF shared with LPM, etc. in same 416K ASIC memory

* Roadmap. Target 17.8.1

Catalyst 9500 & 9600 Series Core Positioning

Next Generation Core + Edge Switching with Silicon One™ Q200

Feature Optimized
C9500 & C9600-SUP-1

C9600-SUP-1

C9500

- ✓ Best-in-class Enterprise Core feature set
- ✓ Low speeds (1G – 40G) and port density
- ✓ Comprehensive MPLS, EVPN and SDA
- ✓ Ideal for Campus Core, Collapsed Core + Agg

Ideal for **C6K non-XL** deployment migration

Performance Optimized
C9500X & C9600X-SUP-2

C9600X-SUP-2

C9500X

- ✓ Unmatched forwarding scale and performance
- ✓ High speeds (10G – 400G) and port density
- ✓ Scaled MPLS and SDA, WAN-MACsec
- ✓ Ideal for Campus Core + Edge, or Centralized WLC

Ideal for **C6K XL** deployment migration

C9600/X C9500/X - Place In Network (PIN)

BGP IGP STP

1 Layer 3 Core

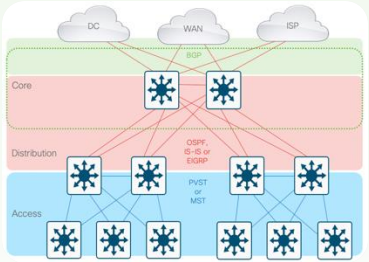

- **Base L3 Routing**
- High-BW, Port Density
- Simple ACL
- Simple QoS

2 Layer 3 Core + Edge

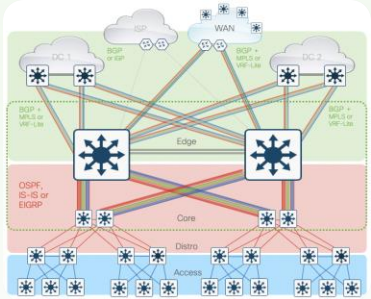

- **L3 Core + Edge Services**
- DCI, WAN, Internet
- Edge Security, VPN & OGACL
- Complex H-QoS

3 Fabric (SDA & EVPN)

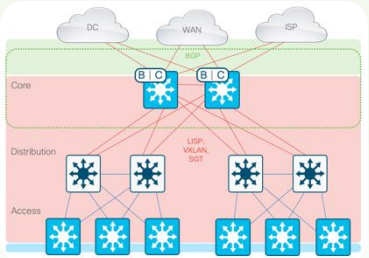

- **L3 Core + VXLAN Fabric***
- Border, Edge, CP/RR
- L2 & L3 VNI & SGACL
- App-based QoS

4 Collapsed Core

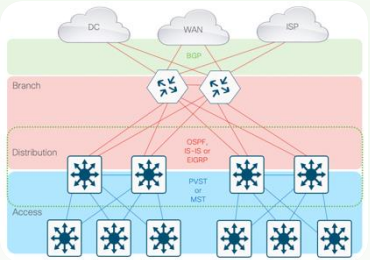

- **L3 Core + Distribution**
- L2 Services to Access
- First Hop Security
- Access QoS & AVC

Edge + Core

Catalyst 9600X with SUP-2

Catalyst 9500X models

Catalyst 9600 SUP-1

Catalyst 9500 models

Core + Distribution

* Roadmap for C9600X/C9500X

Campus + Edge Design

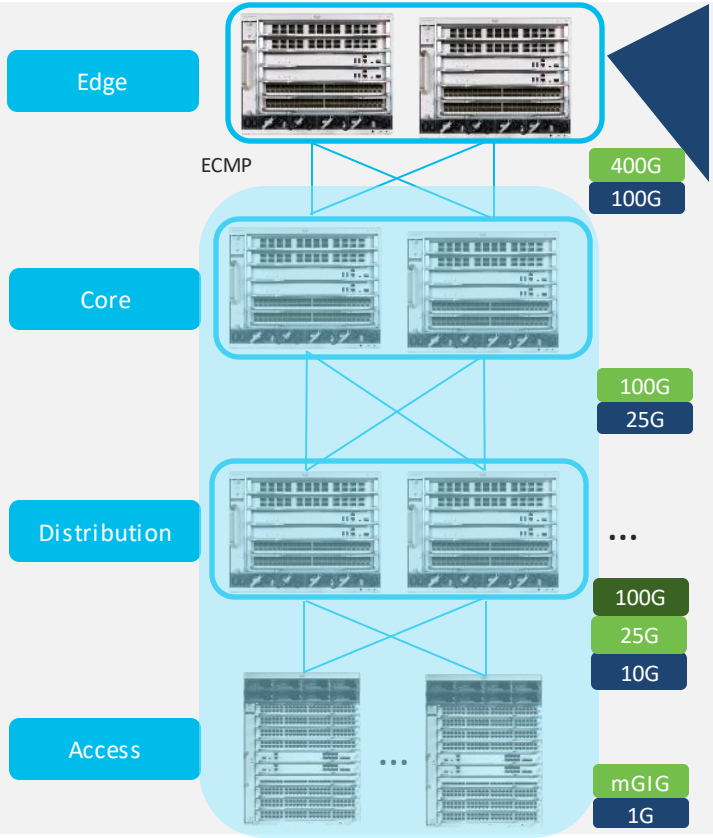

Edge

Catalyst® C9600X and 9500X for campus Edge.

High-Scale Routing:
Support for full IPv4 and IPv6 Internet + LAN + VRF routing

Port speeds:
Support port speed up to 400G

Flexible Transport Options :
Support for MPLS-VPN, SD-Access and BGP-EVPN* with 4K VRFs

Low-Latency:
Support for local optimized low-latency shared memory

Deep buffer:
Support large buffers for micro-burst & congestion

* Roadmap for C9600X/C9500X

Closing

Catalyst 9000 Switching – New and Upcoming Features

Reference

IOS XE 17.7.1
Available on CCO

SECURE

Enhanced Security Controls
- Single policy approach for Dynamic PVLAN
- IPSEC PBR Support
- AAA cache for dot1x for Catalyst 9000

CONNECT

Platform/ Infra
- GIR on 9500H/9600
- xFSU Support for 9300X
- Increased virtual port scale (upto 30K) on 9400
- AES67 timing profile
- PTPv2/ gPTP on 9600 w/o SSO
- Interface template cli Loop Detection Guard

Flexibile Network Segmentation
- Data MDT Support for L3 TRM- IPv4
- EVPN to Global IP Route Leaking
- Interface template based secure onboarding of devices from extended node

AUTOMATE

Programmability
- gNOI reset
- Leaf level filtering for telemetry

IOS XE 17.8.1
Mar '22

Enhanced Security Controls
- Transparently pass MKA BPDUs on non-MACSEC interfaces
- VRF Aware Centralized Web Authentication (CWA)*

SECURE

Platform/ Infra
- G8275.x Timing profile
- Catalyst 9300/9400: Dynamic Power consumption reporting

Flexible Network Segmentation
- EVPN-VxLAN over IPSEC
- IPSEC - Multicast over SVTI
- Layer 2 IPv4/IPV6 TRM support with External RP

CONNECT

Programmability
- ZTP config through YANG
- Native CLI to XPATH conversion

AUTOMATE

* Limited Availability

The bridge to possible