Cisco C9000 Smart Switches Architecture

Kenny Lei
Technical Marketing Engineer

BRKARC-2668

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until June 13, 2025.

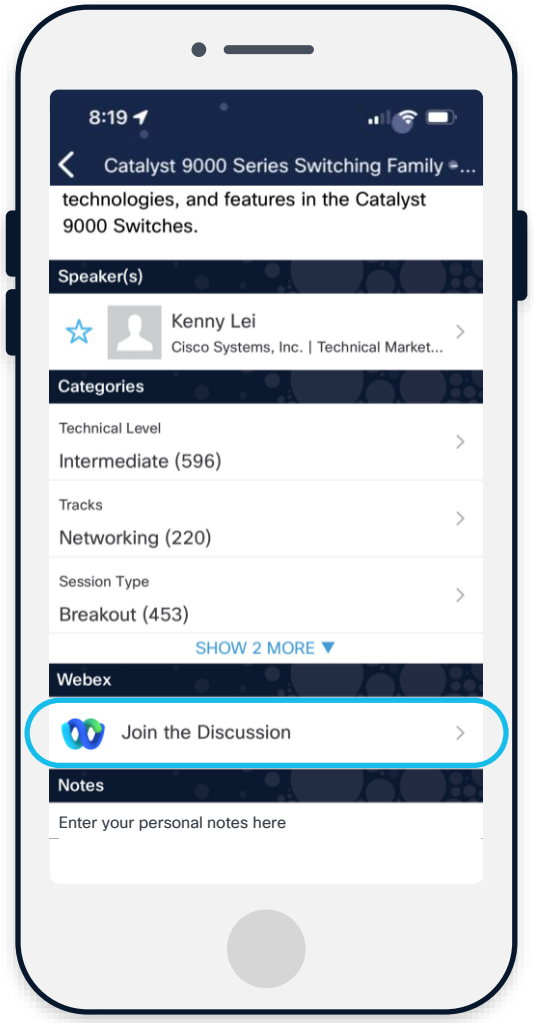

Agenda

01 Cisco Campus Smart Switch Overview

02 Cisco C9350 access smart switch

03 Cisco C9610 core/distribution smart switch

04 Cisco ASICs (A100, K100/E100)

Agenda

05 Features
- HCAM
- Security
- QoS
- High Availability
- Application hosting

06 IOSXE

07 Design consideration

08 Summary

Cisco C9000 Smart Switches

Cisco C9000 Smart Switch Family

Campus-Driven Silicon OneASICs

Cisco Silicon A100/L (C9350) and **E100/K100** (C9610-SUP3/3XL)

AI-Ready X86 & DDR5 CPU & DRAM

Intel Gen 12 Core Series (C9350) and Intel Gen10 Xeon Series (C9610-SUP3/3XL)

Feature-Rich IOS XE Software

UNX, Programmability, NG Stackwise, SD-Access, SmartPower

C9350 Series Smart Switches

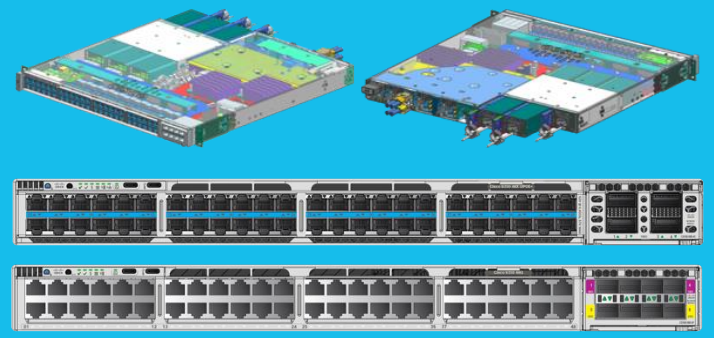

C9610 Series Smart Switches

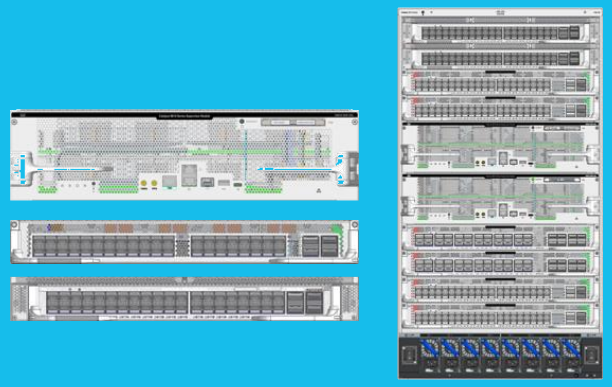

Next Generation of Cisco Campus Switching

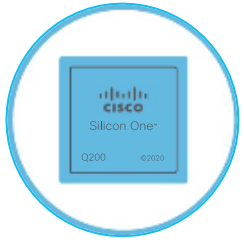

Next-Gen S1 ASICs

- NPU 2.0
- A100/L
- K100/E100

- Increases Performance & Scale in smaller ASIC footprint
- Rich Campus-focused Features built into ASIC hardware
- More Flexible ASIC architecture for efficiency & sustainability

Next-Gen x86 CPUs

- Multi-Core
- DDR4/5 DRAM
- AppGig

- Improves Control-Plane Software processing scale
- Enhances Software-assisted Services (AVC/XDR)
- Enables App-Hosting with AI-Hosting (Edge Compute)

Next-Gen IOS XE

- FED 3.0
- Unified onboarding
- NG StackWise

- Native Model-Driven Telemetry, Programmability & AI-Ops
- TDL API abstraction for Code Reusability and Stability
- New OS Infrastructure Virtualization & High-Availability

Cisco Silicon One Powers Campus Networking

Innovative silicon features that deliver unmatched control, security, efficiency, and scalability

- **Unified architecture**
- **Adaptable tables**
- Shared memory
- Flexible Serdes
- Silicon One SDK
- Programable networking
- Intelligent load balancing
- Enhanced troubleshooting
- **Hardware Root of Trust**
- **Native encryption and security**

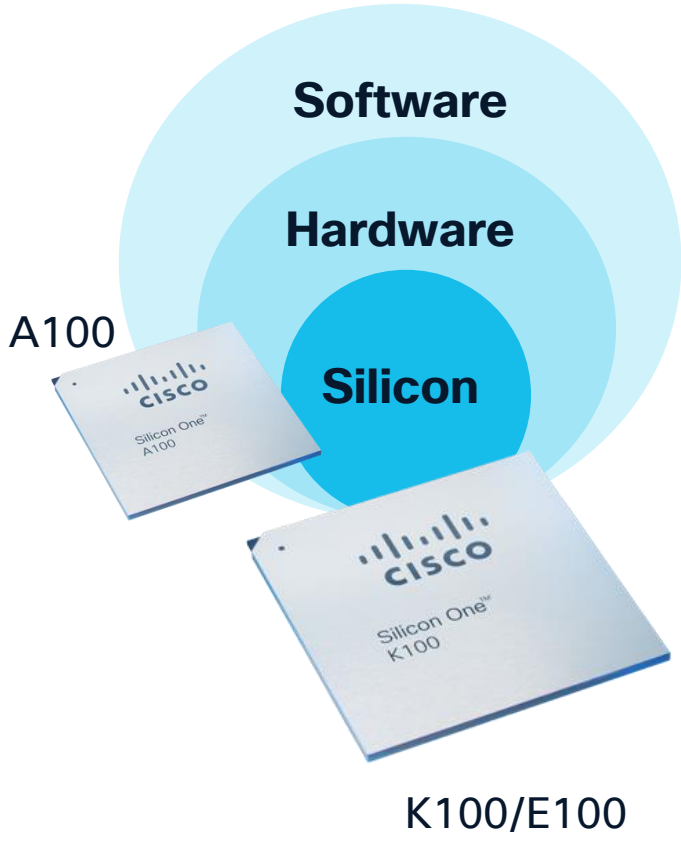

Enhanced Control-Plane & AI App-Hosting

Making Gen3 Cisco 9000 Series into Campus “Smart Switches”

x86

Cisco 9000 Series Switches have supported **Application Hosting** & **Edge Computing** for years

This expertise is being unleased with latest **X86 CPUs** & **DDR DRAM** to enable hosting distributed **Edge AI/ML, security apps**.

C9350

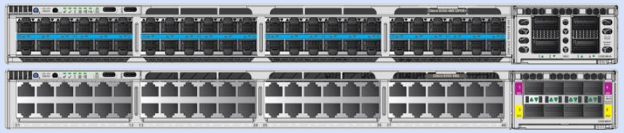

C9350 Standard SKUs

- **Class:** Low/Med CPU
- **Intel** RPL **i3-1316URE**
  - **4E+2P-Core** 2.3 GHz
  - **16GB DDR5**-5200
  - **4x10G NIC** (X710)

C9610-SUP-3/XL

C9610-SUP-3/XL SKUs

- **Class**: Med/High CPU
- **Intel** ICX **D-1733NT**
  - **8P-Core** 2.0GHz
  - **32GB DDR4**-2400
  - **4x25G-KR**

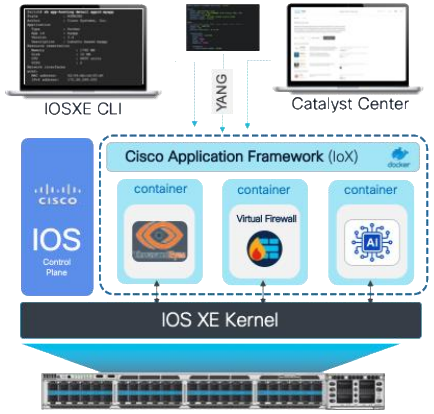

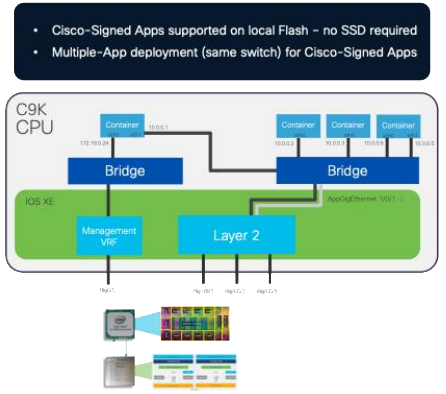

NG IOSXE

Open Cisco IOS XE built for Programmability, Virtualization and AI-Ready

Introducing a NG IOSXE architecture with 30 years of experience!

IOSXE builds on existing infrastructure to better abstract & modularize internal software processes, with a focus on 'model-driven' objects and databases, to support advanced **programmability**, **clustering** & **resiliency**.

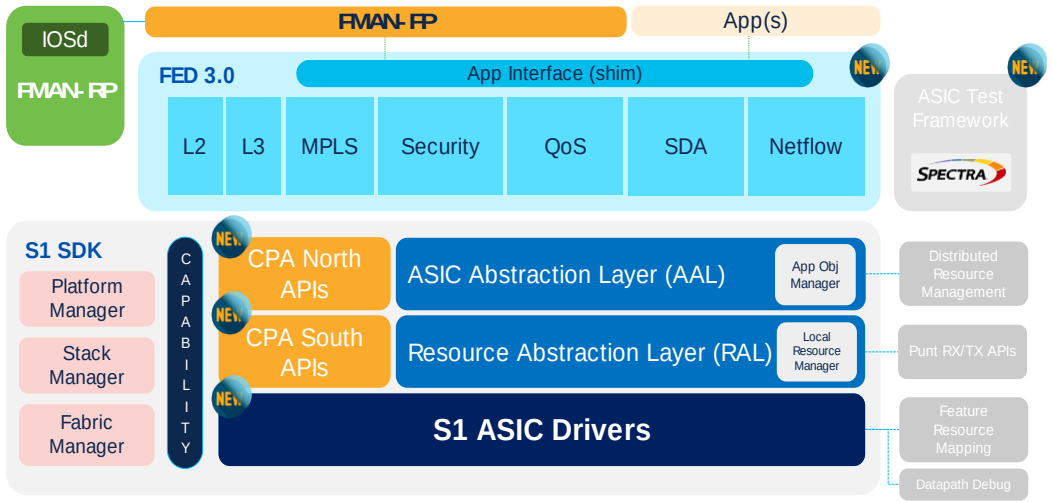

Common Platform Abstraction (CPA)

Model-driven HW to SW API disaggregation layer (CPA) hides platform hardware details from upper-layers of IOSXE features. Acts as a "single source of truth" for all hardware data.

Model-Driven Hardware & SDK APIs

FED programs hardware and manages the forwarding objects. FED 3.0 is model-driven, ASIC & SDK agnostic, and enhanced for various TDL databases with automatic code generation.

Unified Platform & Licensing (UNX)

NG C9K platforms will support future management controller models (cloud and/or on-premise) and will support a single Unified Licensing (UNX) model for all management options.

Cloud-Native Management & AI/ML

Automation, Assurance and AI-Ops models will support entirely cloud-based, entirely on-premise (e.g. air-gap) or a hybrid of both. All available through a single role-based UI/UX.

NG-Stacking & Cluster Infrastructure

Current StackWise and SVL (protocol & encapsulation) have been redesigned using standard VXLAN based protocol & encap, allowing >8 node ethernet-based (SPF) HA clusters.

Micro-Services Process Modularity

Modern multi-threaded IOS XE can start/stop individual SW processes on-demand, to enable new services and support real-time SMU and ISSU and enhanced infrastructure HA.

C9350 Smart Switches

Introducing C9350

Smart & Energy Efficient

EnergyStar Rated
Smart power
System and Endpoint Hibernation
EEE/Auto-off Components

High feature scale

Mac scale up to 128K*
IP route scale up to 256K*
ACL scale up to 36K*

Full throttle - Downlinks

All 48 ports simultaneous Full mgig/10G with UPOE+, Perpetual POE, Fast POE

25/50G & 100G uplinks

Dense uplinks: 2 or 4x 100G/40G, 50G at access: 8x 25G/10G or 4x 50G, 8x 10G/mGig

Silicon One A100

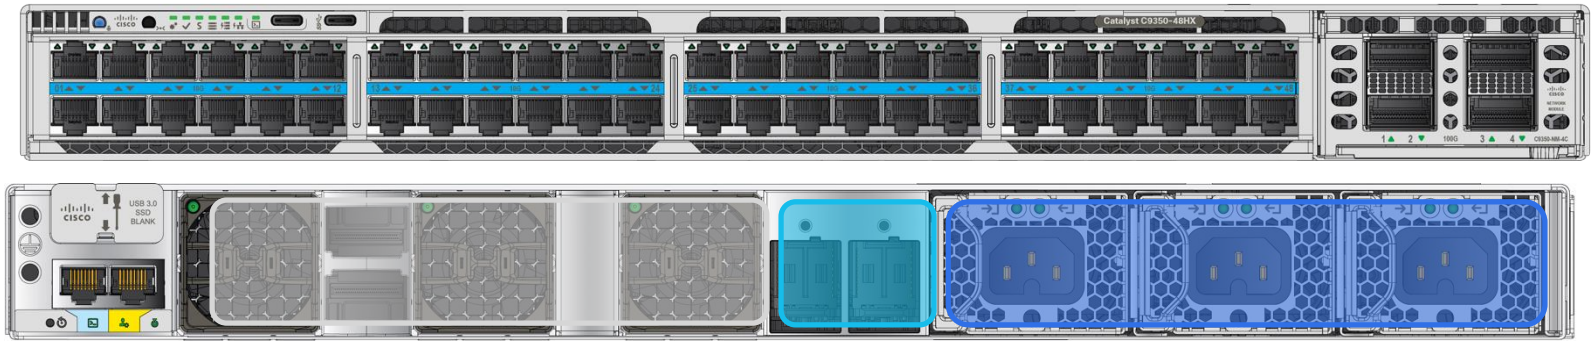

1.6T Stacking

Next Gen Stackwise Architecture
N+1 redundant fan module with reversible air flow options.

StackPower+

Enhanced power pooling with StackPower

3 Power Supplies

N+1 redundancy

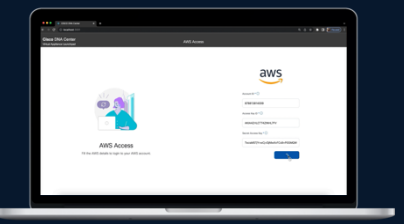

Single hardware, single license

On-prem / Cloud management

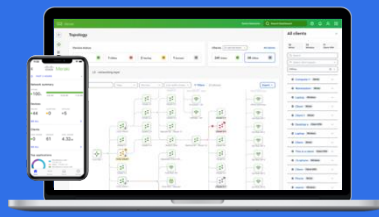

*HW Capability

Cisco C9350 Series

Gigabit SKUs

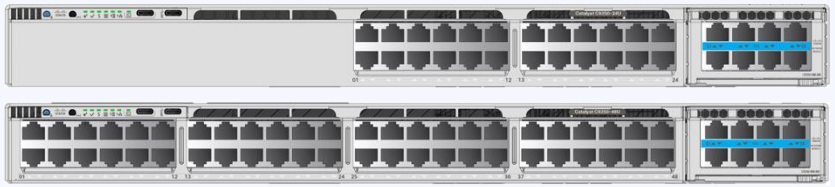

C9350-24P/48P
24/48x 1G/100M/10M copper
30W PoE+

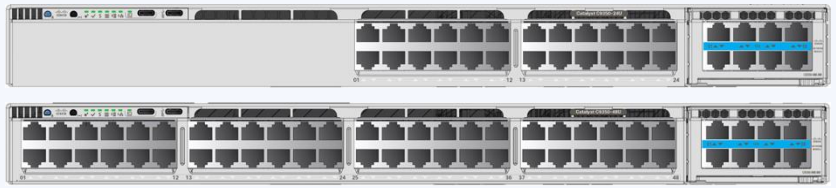

C9350-24U/48U
24/48x 1G/100M/10M copper
60W UPoE

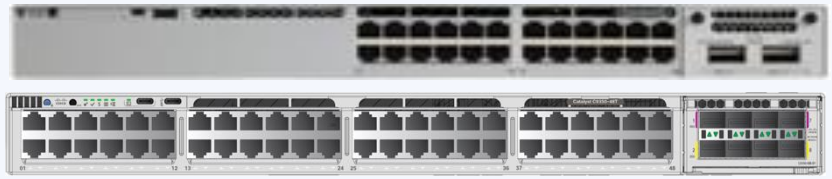

C9350-24T/48T
24/48x 1G/100M/10M copper
Data only

Modular Uplink 50G at access

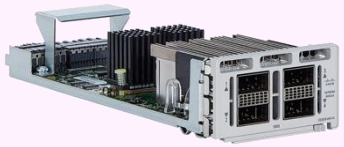

C9350-NM-4C
4x40/100G

C9350-NM-2C
2x40/100G

C9350-NM-8Y
8x10/25G or 4x50G

MultiGigabit SKUs

C9350-48HX
48x 10G mGig
90W UPOE+

C9350-48TX
24/48x 10G mGig
Data only

Note: Uplinks/PSU's/Stack Cables are not backward compatible with C9300(X) series

C9350 New Modular Fan Modules

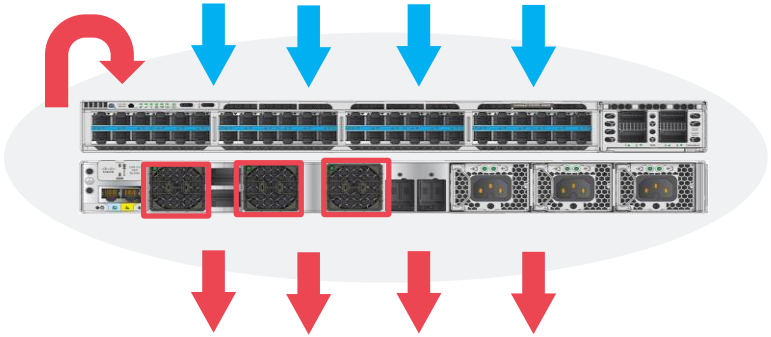

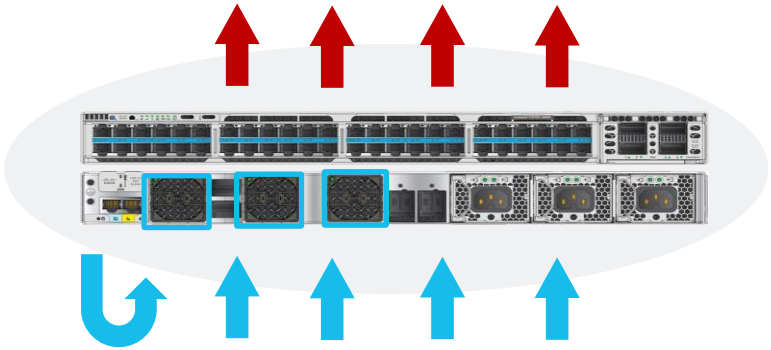

Front-to-Back Airflow -Port Side Intake (PSI)

- Ideal for enclosed racks in wiring closets, ensuring front intake cooling.
- **Red Handle(C9350-FAN-I)**
- Supported on **All SKU's**

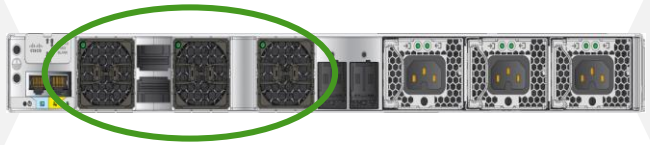

Back-to-Front Airflow -Port Side Exit (PSE)

- Suited for open-rack MDF/IDF rooms, directing airflow toward rear ventilation.
- **Blue Handle(C9350-FAN-E)**
- Supported on **Data/Fibre SKU's**

30% Boost in Airflow Capacity for Enhanced Thermal Efficiency

Cisco C9350-48HX/TX models

48 port 1/2.5/5/10G copper

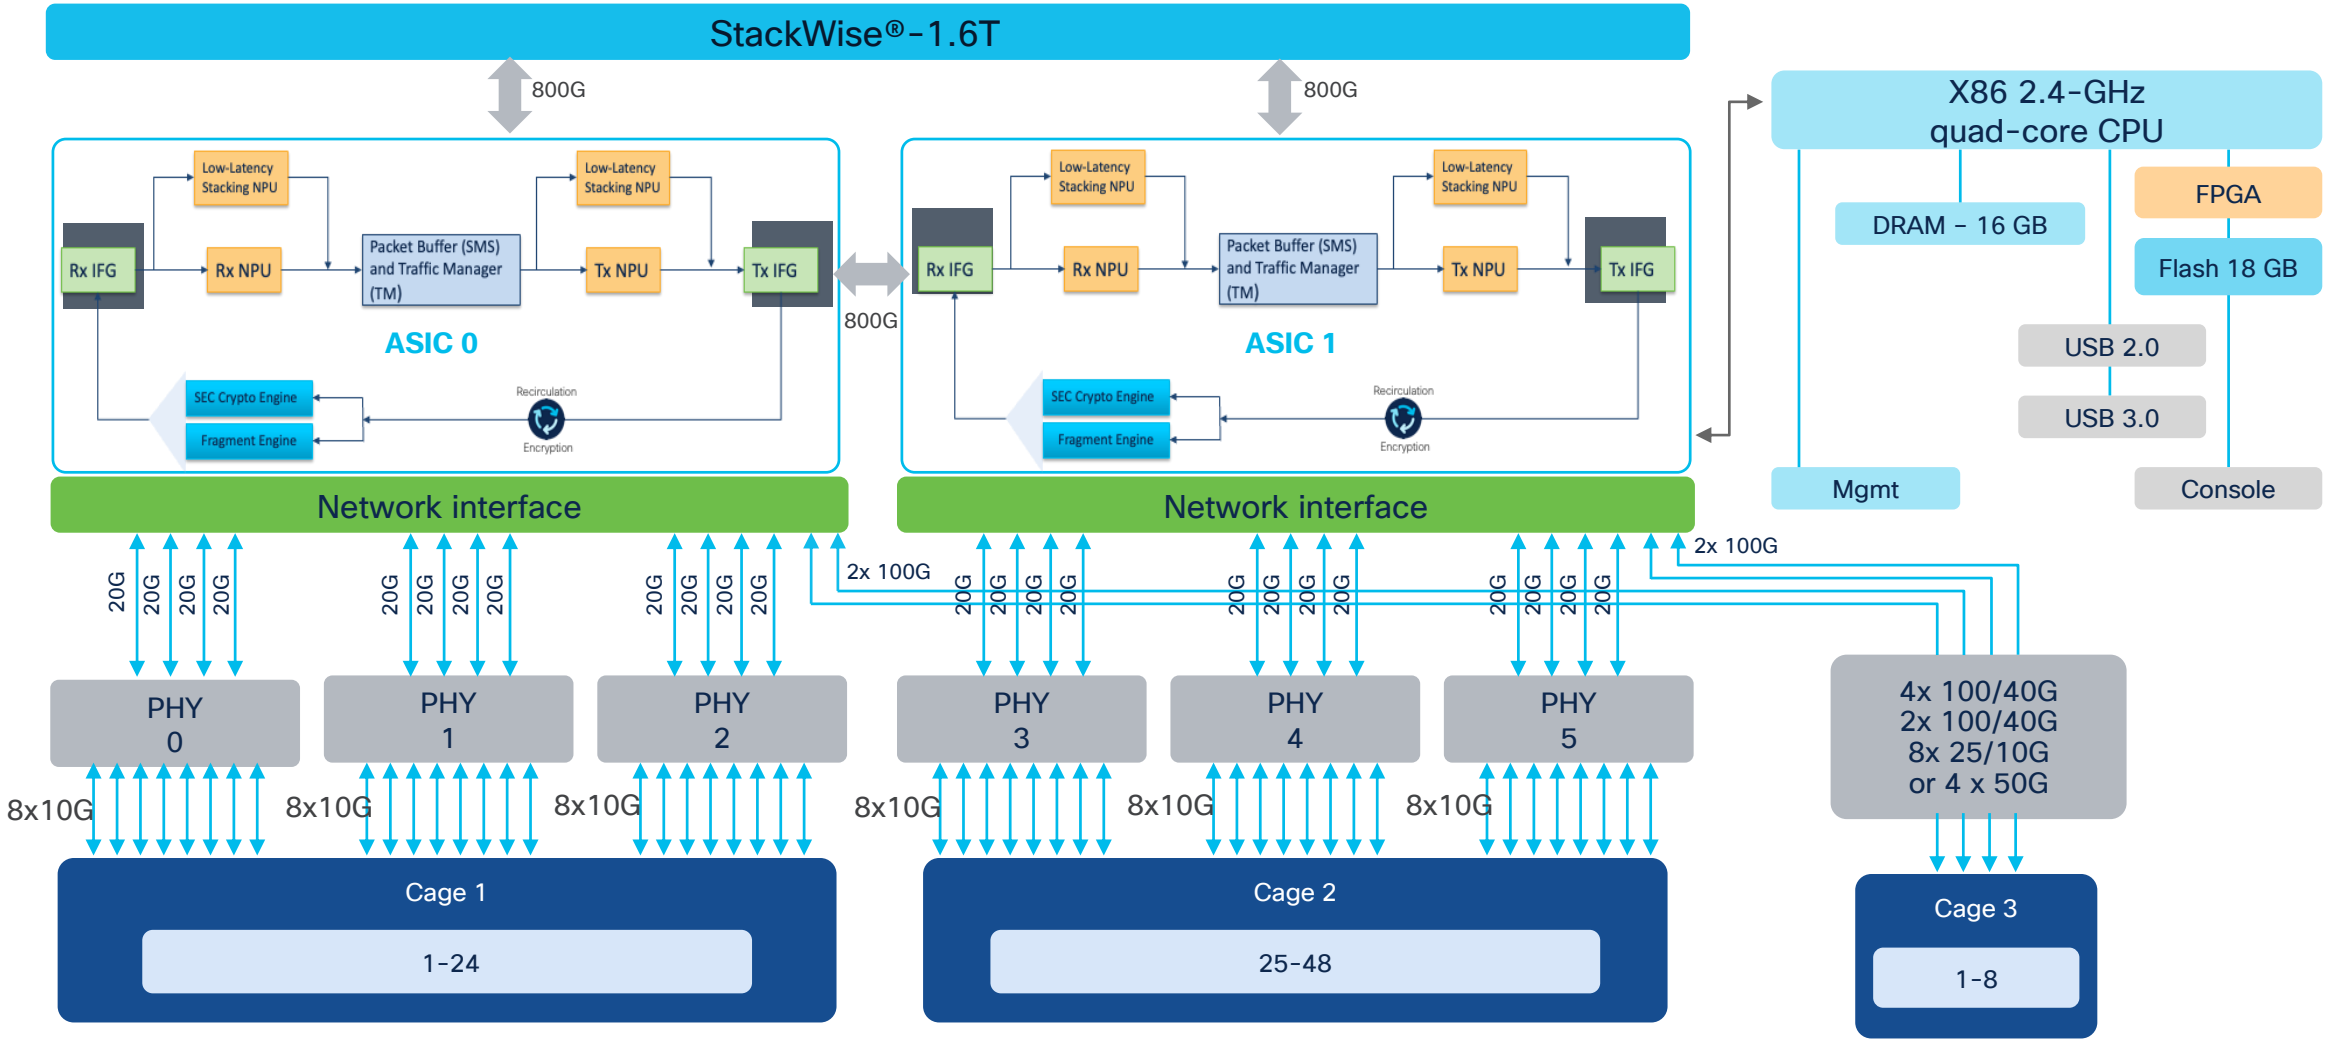

Cisco C9610 Series Core Smart Switches

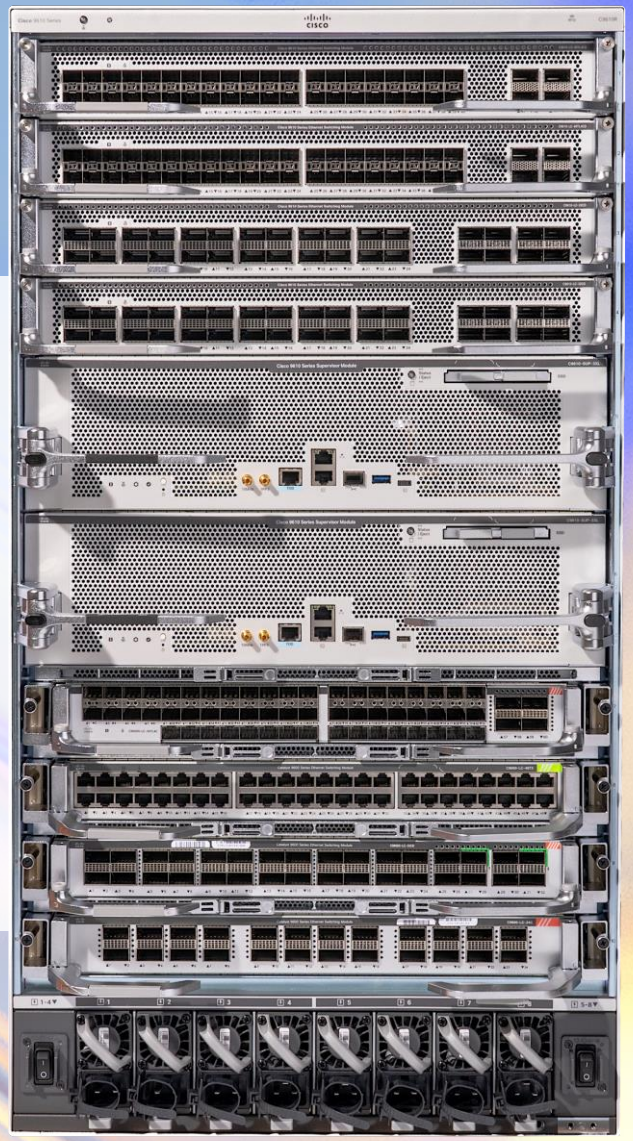

Introducing C9610

Generation 3 modular core – New Hardware launch

NOTE: Now with optional RFID for all components. (Choose on ordering)

Gen 3 Supervisors

C9610-SUP3XL/3

Enterprise Focus
K100/E100

- **2M IPv4 / 1M IPv6**
- 3.2 Tbps per slot
- **512K hosts or clients***
- 32G DDR4 memory
- **256K ACL entries***
- **8 core** X86 CPU @2.0 Ghz

Gen 3 Chassis

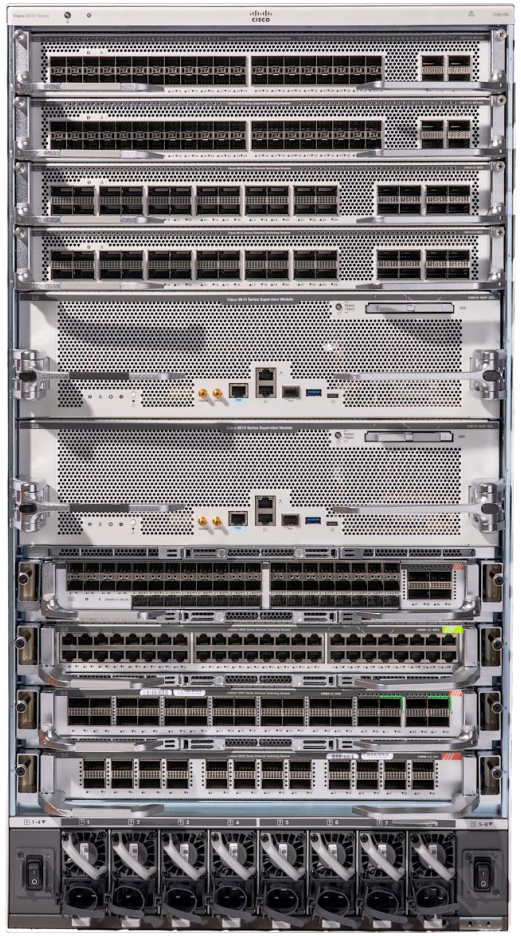

C9610R

- 51.2 Tbps System Bandwidth
- 8-line card slots (1.25 RU)
- 2 supervisor slots (2.5 RU)
- 4 serviceable fan trays in rear side
- Blue Beacons (system/fan tray, sup, line cards)
- 8 Modular power supplies

Gen 3 Line cards

C9610-LC-40YL4CD

C9610-LC-32CD

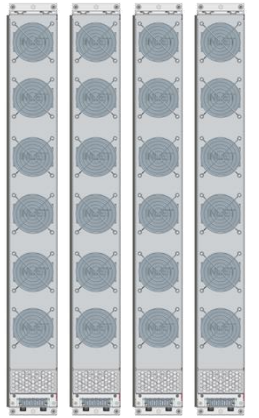

4x FAN Trays
with Front to Back Airflow

Cisco C9610R Chassis

Four serviceable fan trays in rear side

Built-in RFID (optional)

Blue Beacons (system, fan tray, sup, line cards)

Modular power supplies

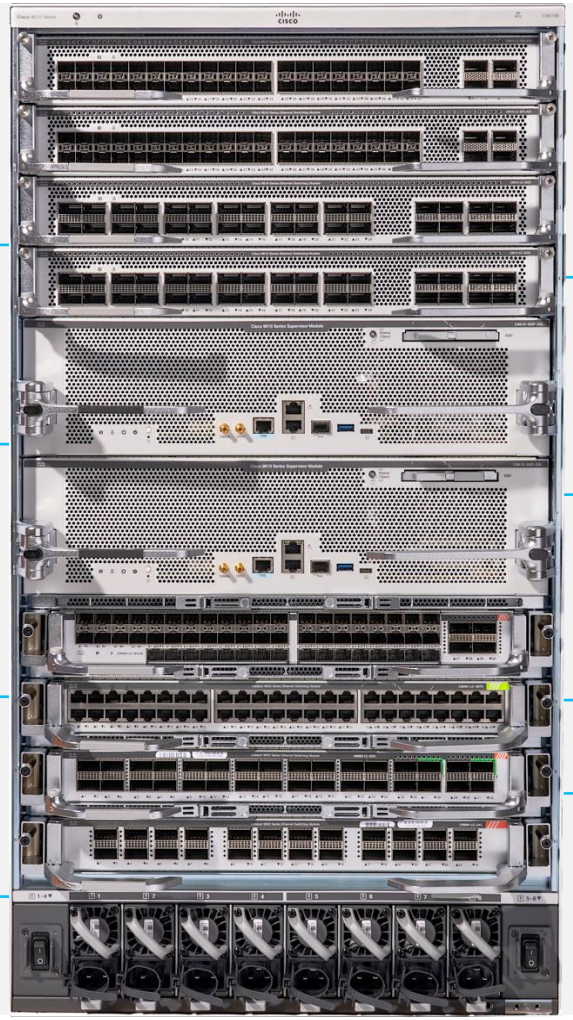

8 line-card slots (1.25 RU)

2 supervisor slots (dedicated) (2.5 RU)

51.2 Tbps
System Bandwidth

Dimensions
(H x W x D inches)
31.47 x 17.4 x 26.1
(18RU)

Industry’s First Centralized Cable Backplane

Back View

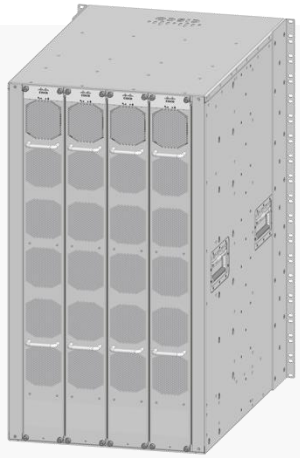

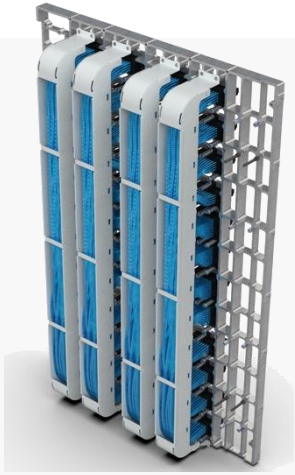

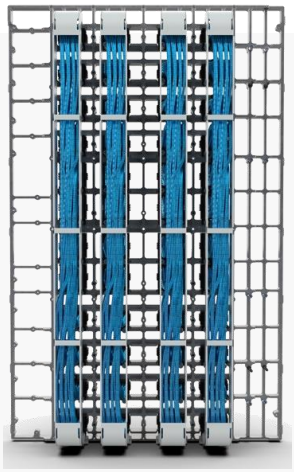

Front View

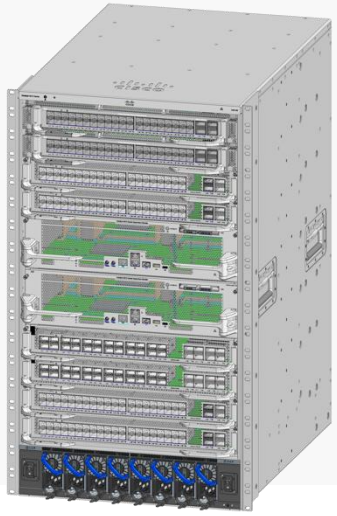

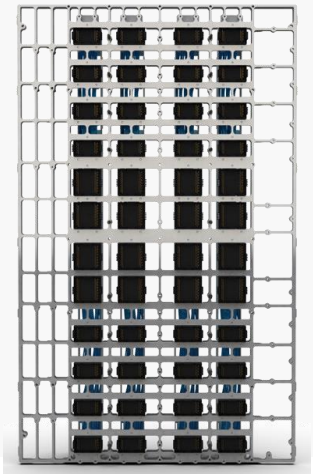

- Passive cable backplane
- High performance cable 100G SerDes
- Modular design
- Front to back airflow

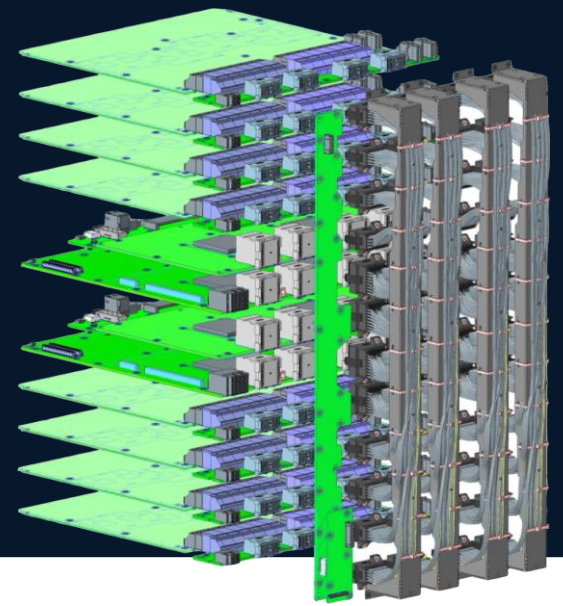

Cisco C9610R Series Supervisors

Supervisor 3 and Supervisor 3 XL

Powered by Silicon One E100

C9610-SUP3 - E100

- 1M IPv4 Prefixes
- 768K Exact Match
- 256K Policies

25.6 Tbps (full duplex)	**4x K100/E100** + 1x Fabric ASIC
3.2 Tbps per slot	**8x Line Card** slots + **2x Supervisor** slots
8 core X86 CPU @ 2.5 GHz	**32GB** DDR4 DRAM
Built-in RFID (optional)	Copper and Fiber Mgmt ports
SATA SSD (Field Replacable)	1 x USB 3.0 storage 1x USB C console

Powered by Silicon One K100

C9610-SUP3XL - K100

- 2M IPv4 Prefixes
- 1.2M Exact Match
- 512K Policies

Cisco 9600 Architecture

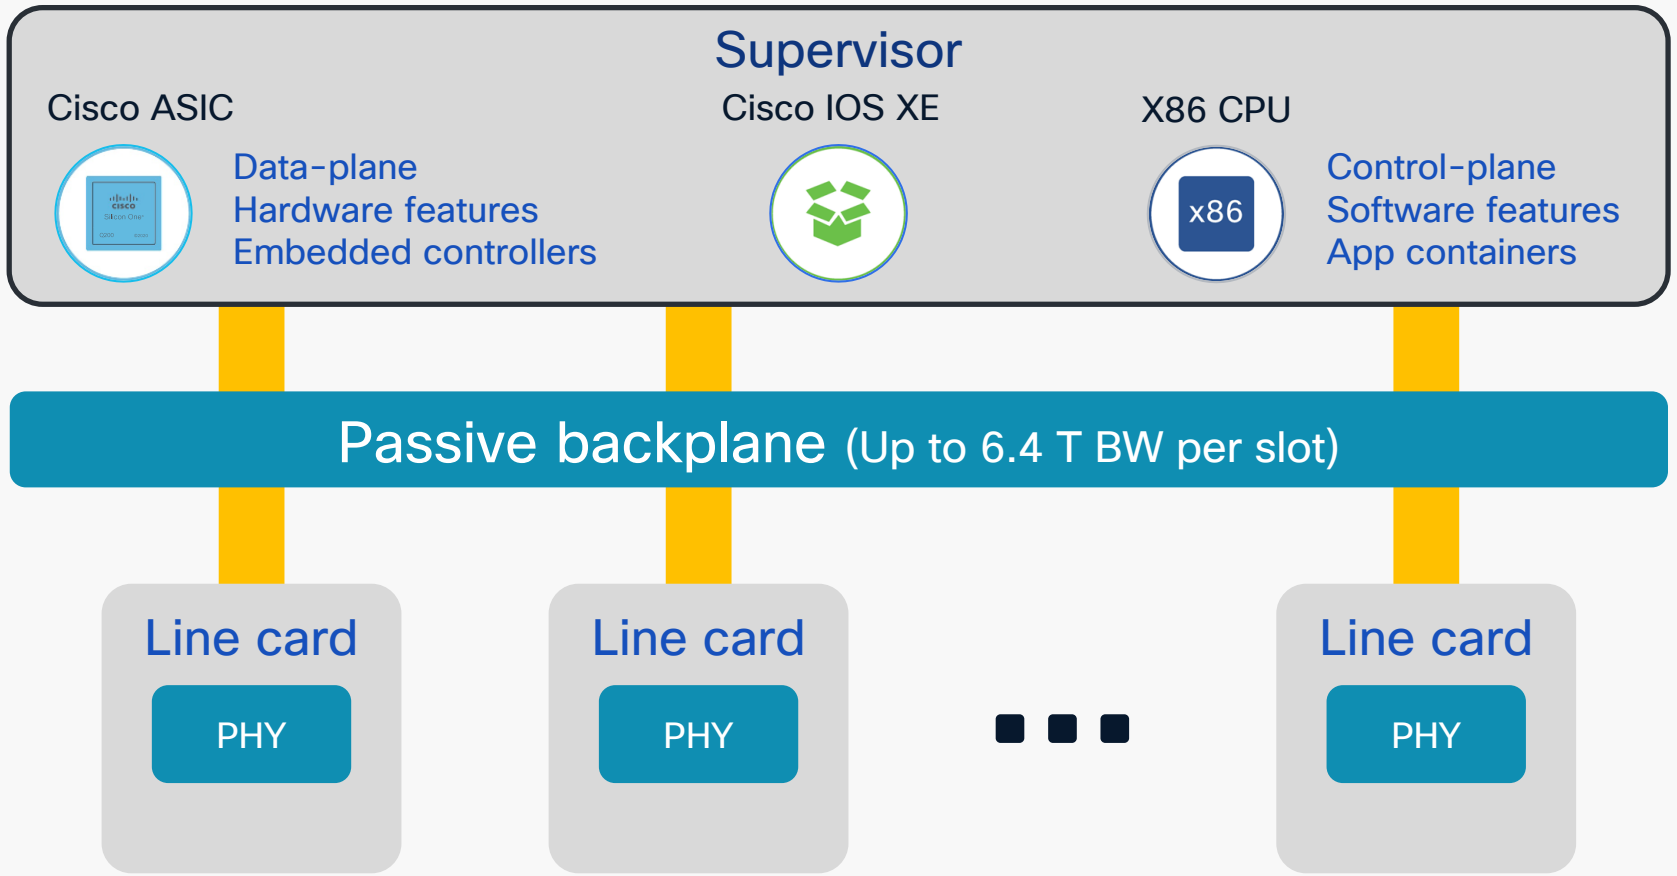

- Centralized architecture:
  - Uninterrupted supervisor switchover
  - Forwarding, queuing, and security are done on the supervisor
  - Unlock new capabilities with a supervisor upgrade
- Transparent line cards: Compatible with new sup
- Passive backplane: Higher MTBF
- X86 CPU + storage: App hosting

C9610-SUP-3/3XL

Supervisor Block diagram

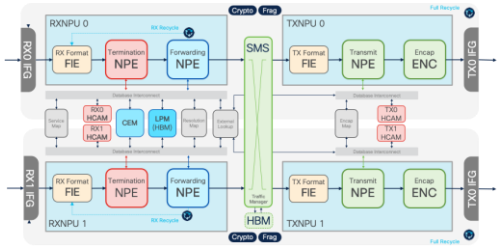

Switch backplane

128x50G (6.4T)

128x50G

128x50G

128x50G

K100/E100 #1

K100/E100 #2

K100/E100 #3

K100/E100 #4

64x50G (3.2T)

64x50G

64x50G

64x50G

Fabric ASIC

For App Hosting

SupIO (I2C)

PCIe

Ethernet

FPGA

8-core CPU

OBFL (1 Gb)

Ext. SATA (optional)

SDRAM

2x USB3

USB Console, Console, Mgmt

SFP+

- 4x K100/E100 + Fabric ASIC
- 25.6Tbps (512 x 50G)
- 2:1 ASIC interconnect
- 64 SerDes/Slot

Port Mapping

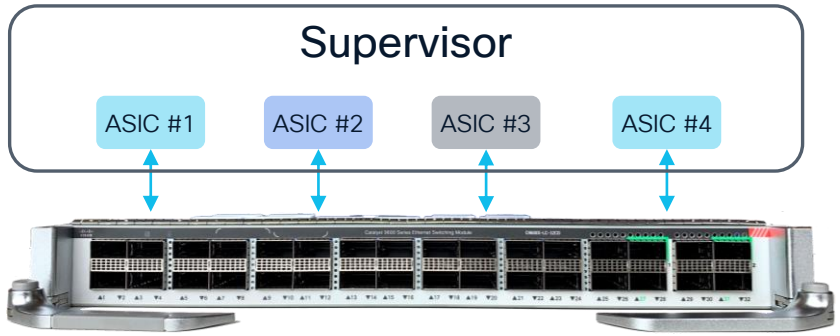

- Each slot has ¼ ports to each K100 ASICs
- Each slot has access to all 4 K100 ASICs resources
- Simplify deployment

```
EPCOT-P2D-TME#sho platform software fed active ifm mappings
Interface                IF_ID  Inst Asic Core IFG_ID Port SubPort Mac
HundredGigE3/0/1         0x4b6     3    3    1      0    0       0   129
HundredGigE3/0/2         0x4b7     3    3    1      0    0       0   130
HundredGigE3/0/3         0x40c     3    3    1      0    0       0   131
HundredGigE3/0/4         0x40d     3    3    1      0    0       0   132
HundredGigE3/0/5         0x40e     1    1    1      0    0       0   133
HundredGigE3/0/6         0x40f     1    1    1      0    0       0   134
HundredGigE3/0/7         0x410     1    1    1      0    0       0   135
HundredGigE3/0/8         0x411     1    1    1      0    0       0   136
HundredGigE3/0/9         0x412     3    3    0      0    0       0   137
HundredGigE3/0/10        0x413     3    3    0      0    0       0   138
HundredGigE3/0/11        0x414     3    3    0      0    0       0   139
HundredGigE3/0/12        0x415     3    3    0      0    0       0   140
HundredGigE3/0/13        0x416     1    1    0      0    0       0   141
HundredGigE3/0/14        0x417     1    1    0      0    0       0   142
HundredGigE3/0/15        0x418     1    1    0      0    0       0   143
HundredGigE3/0/16        0x419     1    1    0      0    0       0   144
HundredGigE3/0/17        0x41a     2    2    1      0    0       0   145
HundredGigE3/0/18        0x41b     2    2    1      0    0       0   146
HundredGigE3/0/19        0x41c     2    2    1      0    0       0   147
HundredGigE3/0/20        0x41d     2    2    1      0    0       0   148
HundredGigE3/0/21        0x41e     0    0    1      0    0       0   149
HundredGigE3/0/22        0x41f     0    0    1      0    0       0   150
HundredGigE3/0/23        0x420     0    0    1      0    0       0   151
HundredGigE3/0/24        0x421     0    0    1      0    0       0   152
HundredGigE3/0/25        0x422     2    2    0      0    0       0   153
HundredGigE3/0/26        0x423     2    2    0      0    0       0   154
FourHundredGigE3/0/27    0x424     2    2    0      0    0       0   155
HundredGigE3/0/28        0x425     2    2    0      0    0       0   156
HundredGigE3/0/29        0x426     0    0    0      0    0       0   157
HundredGigE3/0/30        0x427     0    0    0      0    0       0   158
FourHundredGigE3/0/31    0x428     0    0    0      0    0       0   159
HundredGigE3/0/32        0x429     0    0    0      0    0       0   160
```

ASIC# Slice#

Cisco C9610R Support C9606R Line Cards

C9600-LC-24C - 100G/40G(fiber)

- 24 ports
- QSFP28/QSFP+
- 100G and 40G

17.18.2

C9600-LC-40YL4CD – 400G/200G*/100G/50G/40G/ 25G/10G (fiber)

- 40+2+2 ports
- SFP56/QSFP56+/QSFPDD
- 400G,200G**,100G,50G,40G,25G, 10G and 1G*

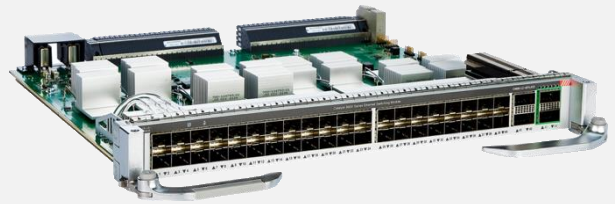

C9600-LC-48YL – 50G/ 25G/10G/1G (fiber)

- 48 ports
- SFP56/SFP28/SFP+
- 50G,25G,10G,and 1G*

17.18.2

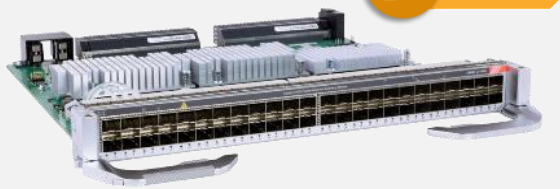

C9600X-LC-32CD – 400G/200G*/100G/40G(fiber)

- 30+2 ports
- QSFP28/QSFPDD
- Supports 400G,200G*,100G, and 40G

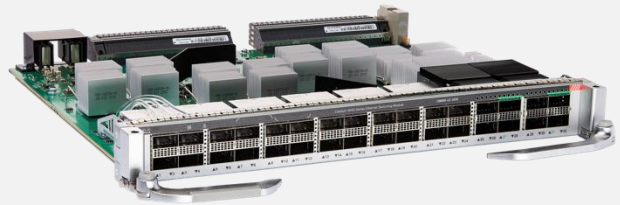

C9600-LC-48TX - mGig (copper)

- 48 ports
- Copper 10G (NBASE-T/10BASE-T)
- 10G and 1G*

C9600X-LC-56YL4C - 100G/50G/40G/25G/10G (fiber)

- 56 50G+4 100G ports
- SFP56/QSFP28
- 100G, 50G, 40G, 25G, 10G and 1G*

* Native 1G 17.18.2

** 200G hardware capable

C9610-LC-ADPT

Adapter for C9606 linecards

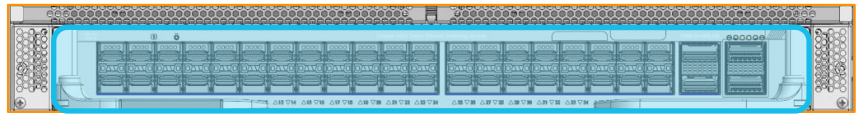

C9606 line-cards were 1RU form factor

1.25RU Adapter house the C9606 line-cards

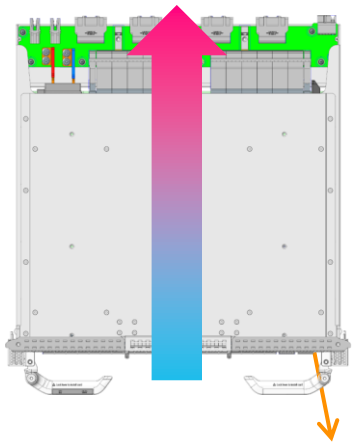

C9606 line-cards were designed for left-to-right air flow

Adapter internally redirects airflow from front-to-back

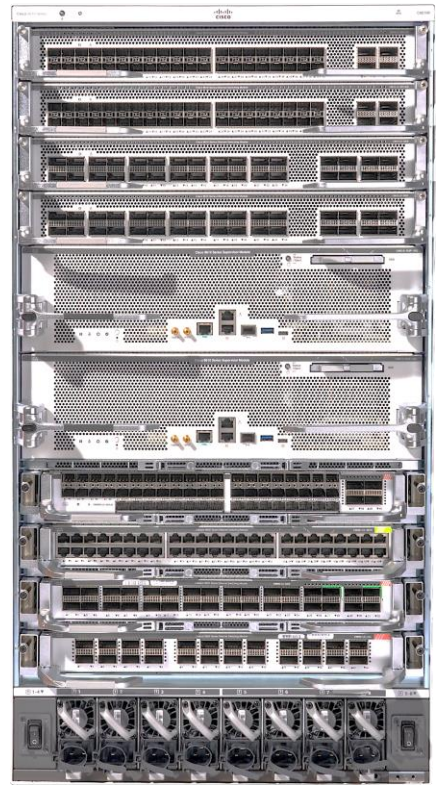

- C9606 and C9610 have the same chassis width.
- No changes in **rack width** required

Backward Compatibility with C9606 LCs

C9610X-LC-32CD Line Card block diagram

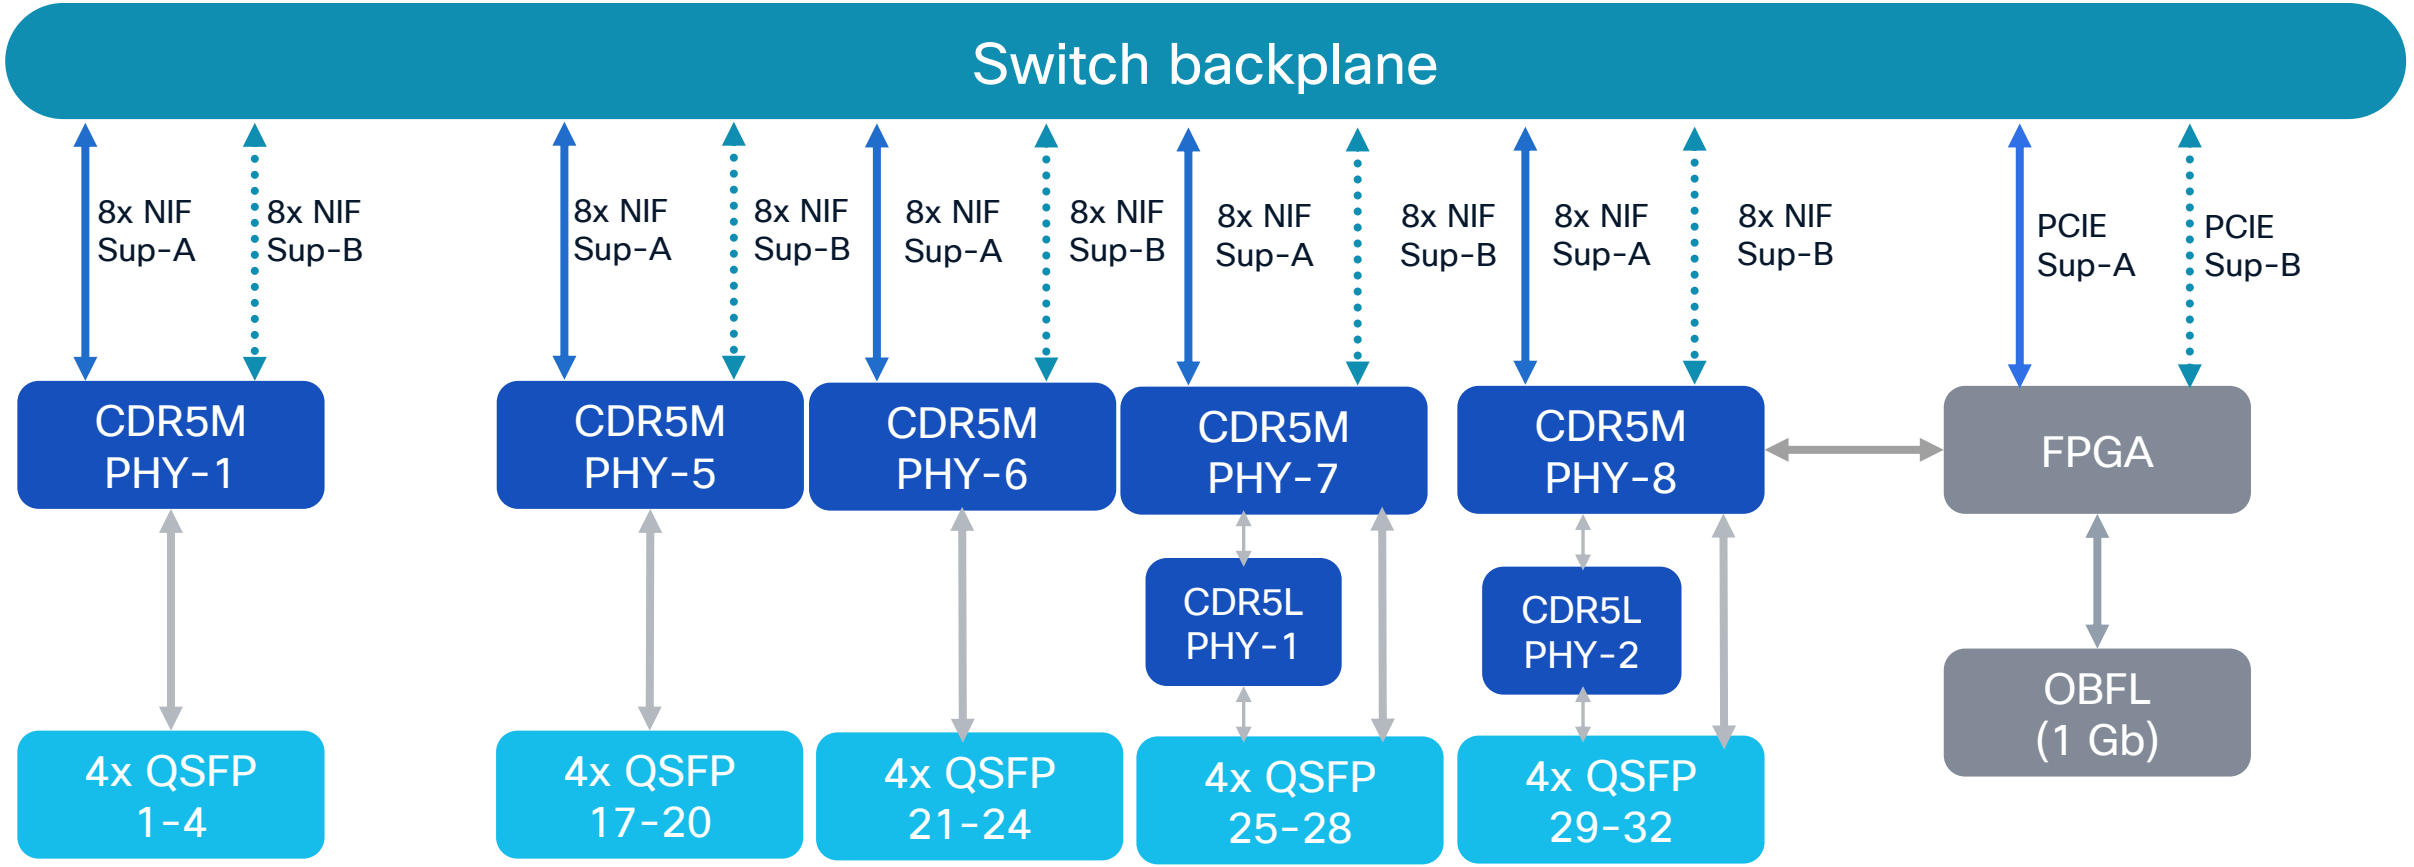

C9610R Power Supplies and FanTrays

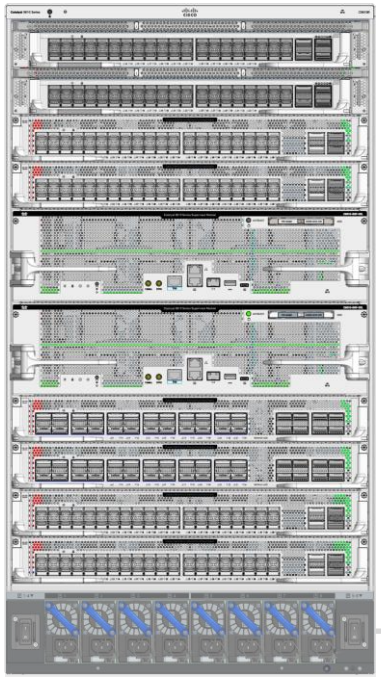

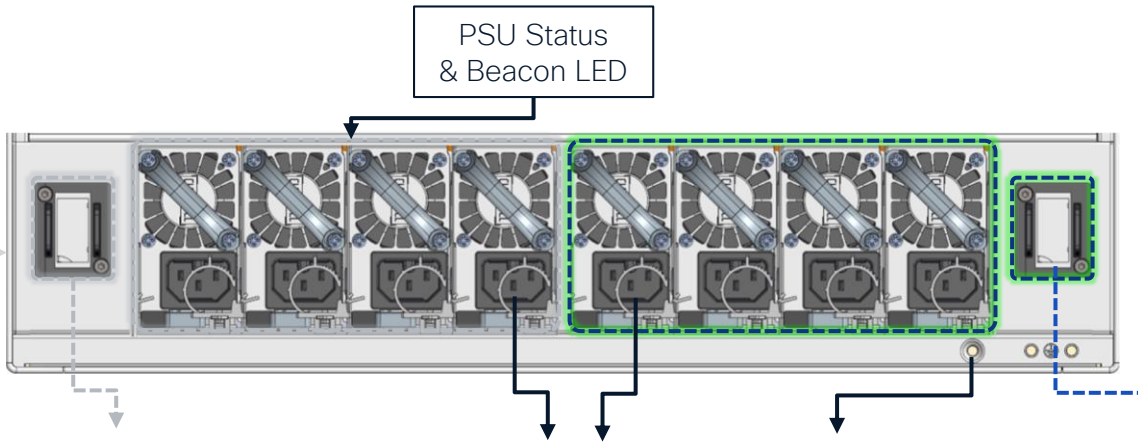

3KW – Titanium Rated
2KW – Platinum Rated

Same PSUs as C9606R (2KW AC/DC or 3KW AC)

Combined or Redundant mode (N+1 or N+N)

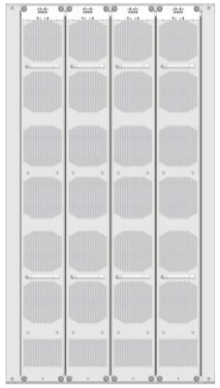

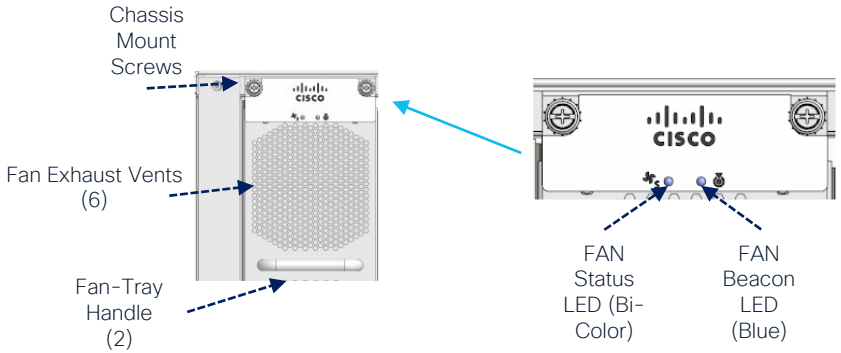

Front to back Airflow

N+ 1 Redundancy

Cisco C9610R Chassis

Port Density (Without breakout)

Reference

Port speed	Density with C9606 Sup 2	Density with C9610 Sup3/3XL
400G	8	**16**
200G	8**	**16****
100G	128	**256**
50G	224	**448**
40G	128	**256**
10/25G fiber	224	**448**
10G Copper	192	**384**
1G	8 (192 with Sup1)	**384***

2X

* Native 1G 17.18.2

** 200G Hardware capable

Silicon One ASICs

What's new in A100 & K100/E100?

Cisco Silicon One NPU 2.0 ASICs - Built for Enterprise Features

Shared innovative **S1 Architecture**

- Multi-Slice Run-To-Complete NPU
- Low Power with High-Performance
- P4-based NPL and Common SDK
- VOQ-based SMS and HBM/DDR
- Standalone, Module or Fabric mode

Newer **Longest Prefix Match (LPM)**

- High-scale IP/mask with HBM extension
- Improved efficiency with embedded Hash

Newer & larger **Exact Match (CEM)**

- High-scale MAC, ARP, Multicast, SGT, NAT, etc.
- More flexible for various 'exact match' features

New & larger **Policy Match (HCAM)**

- High-scale Algorithmic TCAM + Hash SRAM
- Robust support for ACL, QoS & FNF features
- Better Ingress and Egress policy support

NEW

Hardware **Flexible NetFlow**

- FNF records use HCAM (wide-entries)
- Support for AVC/NBAR and ETA/XDR

Hardware **Data Encryption**

- Embedded AES-GCM 256-Bit Crypto Engines
- Line-rate (LAN) MACsec
- Hardware WAN MACsec and IPsec

Hardware **Precision Time**

- Hardware PTP 1588, AVB, AES67 & G8275
- Support for IT/OT Frame Preemption

Advanced **QoS & Buffering**

- More VoQs, Policers and Shapers
- Enhanced WRED and HQoS

Flexible **Telemetry Counters**

- Millions of Programmable counters
- Configurable Feature allocation

Lower & Higher-Speed **Ports**

- 10/100/1000M and 2.5/5/10G mGig RJ45
- 1/10/25G/50G SFP and 40/100/400G QSFP

... and much more

K100 ASIC – Overview

Reference

- Generation 2.2, 7nm ASIC
- ~9.6 Tbps Bandwidth @ ~1.3GHz
  - ~6.4 Tbps Network + ~3.2 Tbps Fabric
- 3 Slices for ~3 Bpps Throughput
  - 2x 1.3B Slices Network + 1x Slice Fabric
- 192 Serdes* @ 56G PAM4
  - 128x Serdes Network + 64x Fabric
- ≤64MB shared Buffer (up to 8GB w/ HBM)
- ≤2.0M internal LPM (up to 8M w/ HBM)
- ≤1.25M shared CEM
- ≤512K shared HCAM (8K TCAM + 128K Tiles x4)

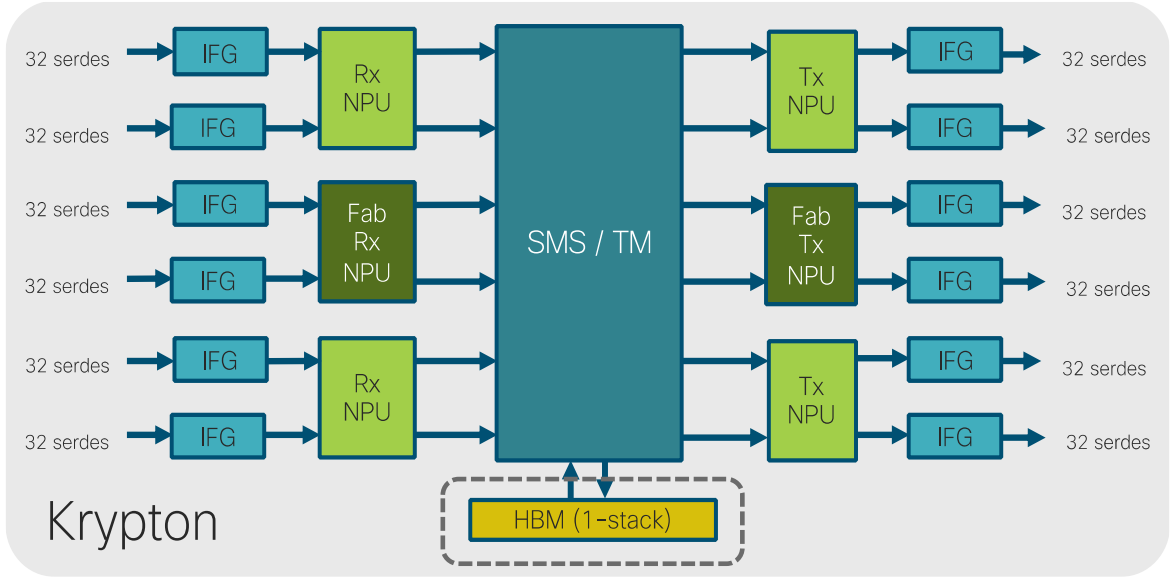

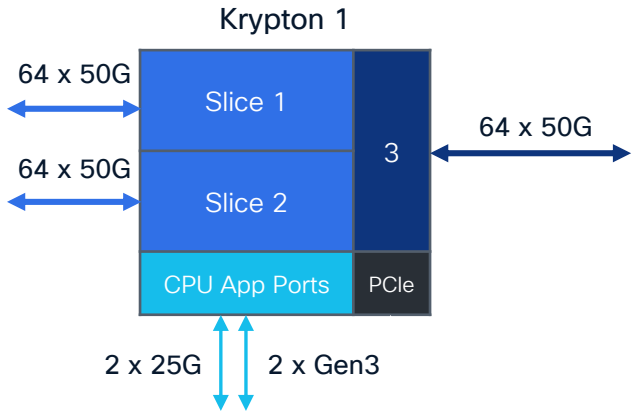

* K100/E100 has an additional 2x Serdes just for CPU (KR) ports

A100 ASIC Overview

Reference

- Generation 2.1, 16nm ASIC
- ~1.3 Tbps Bandwidth @ ~1.0GHz
  - ~500G Network + ~800G Stack
- 1 Slice for ~1.5 Bpps Throughput
  - ~500M RXNPU + ~1B Stack
- 48 SerDes @ 56G PAM4
- ≤18MB shared Buffer (up to 8GB DDR4/5)
- ≤256K internal LPM (no external)
- ≤256K shared CEM
- ≤136K shared HCAM (8K TCAM + 128K Tiles)

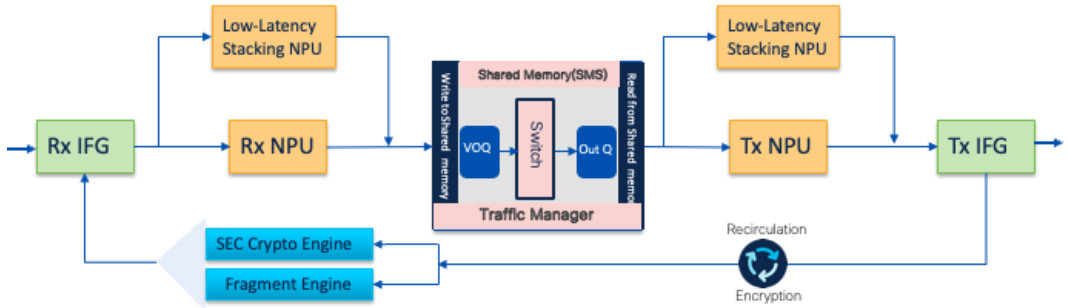

Cisco Silicon One™

ASIC Architecture & Block Diagram

Packet Processing Slices:
- 1 packet per clock
- Slice = 2x IFGs + 1 RX & TX NPU

RX & TX NPU (per slice):
- P4 programmable Run-to-Complete
- Large Central Database (CDB) Tables
- Expandable LPM in external HBM

Traffic Manager (TM)
- Large fully-shared memory switch
- Congestion Management
- Pool of queues & flexible scheduling

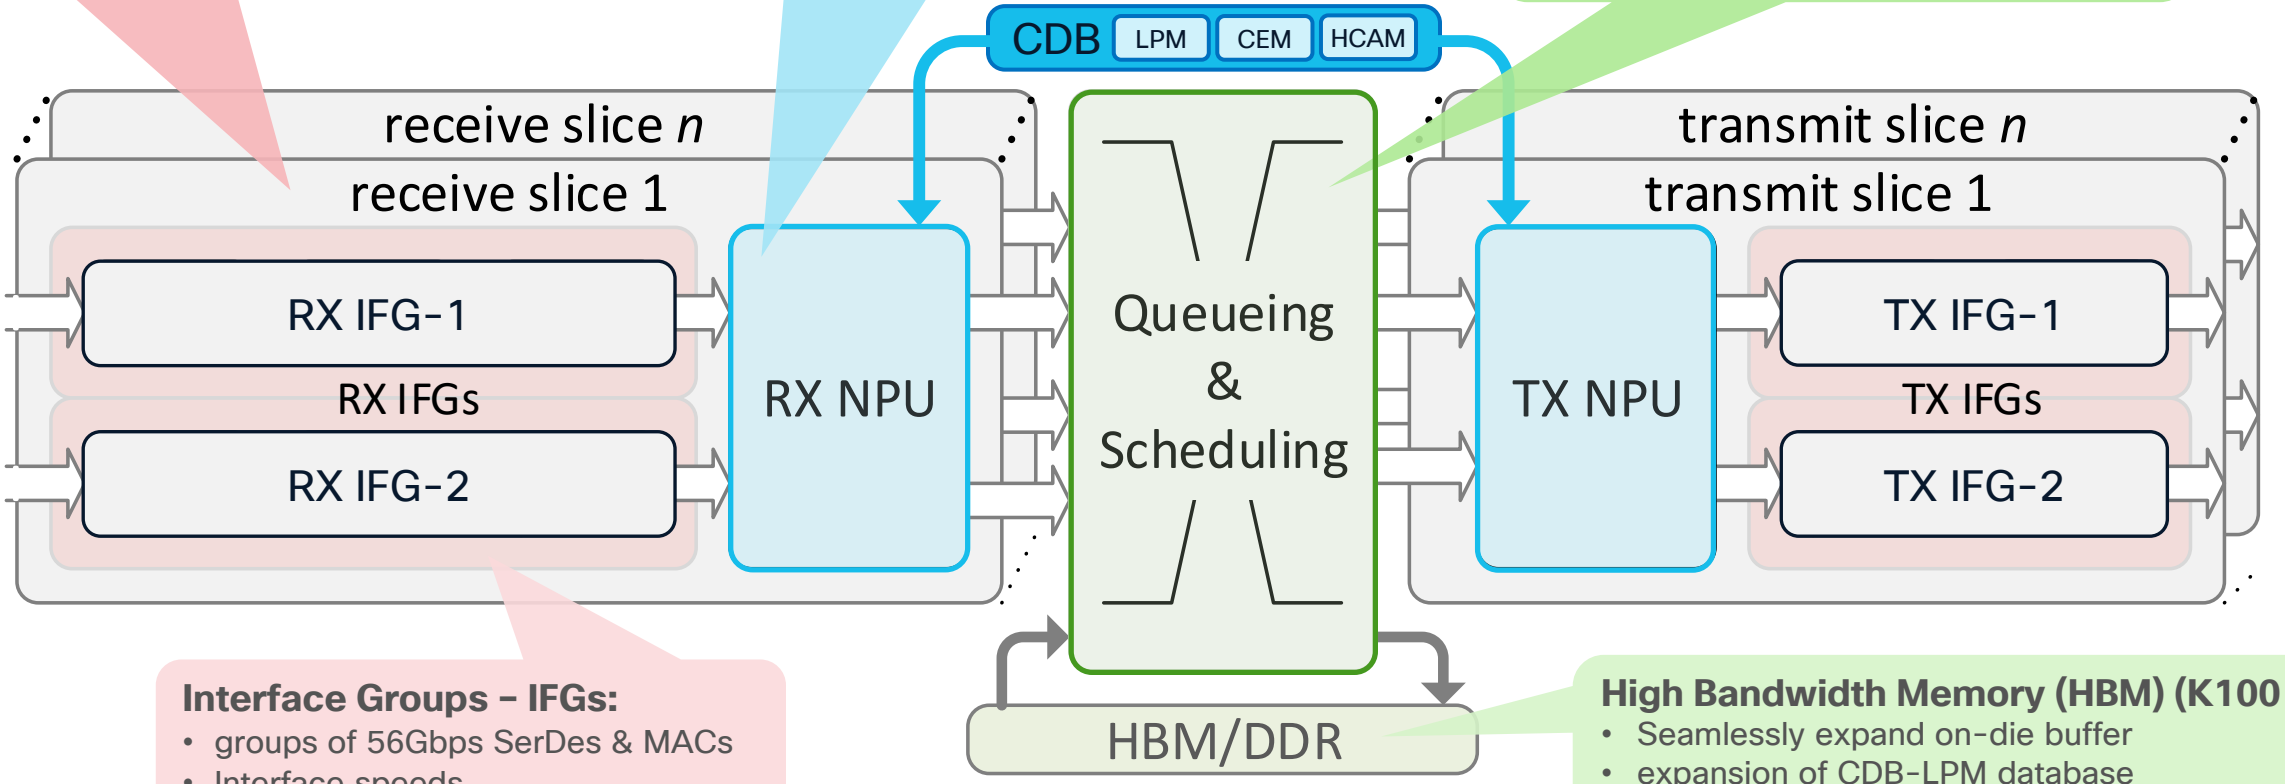

Interface Groups – IFGs:
- groups of 56Gbps SerDes & MACs
- Interface speeds

High Bandwidth Memory (HBM) (K100 only)
- Seamlessly expand on-die buffer
- expansion of CDB-LPM database
- 8GB of fully shared memory

Cisco S1 – Standalone SOC Architecture

Generic Unicast Packet Walk

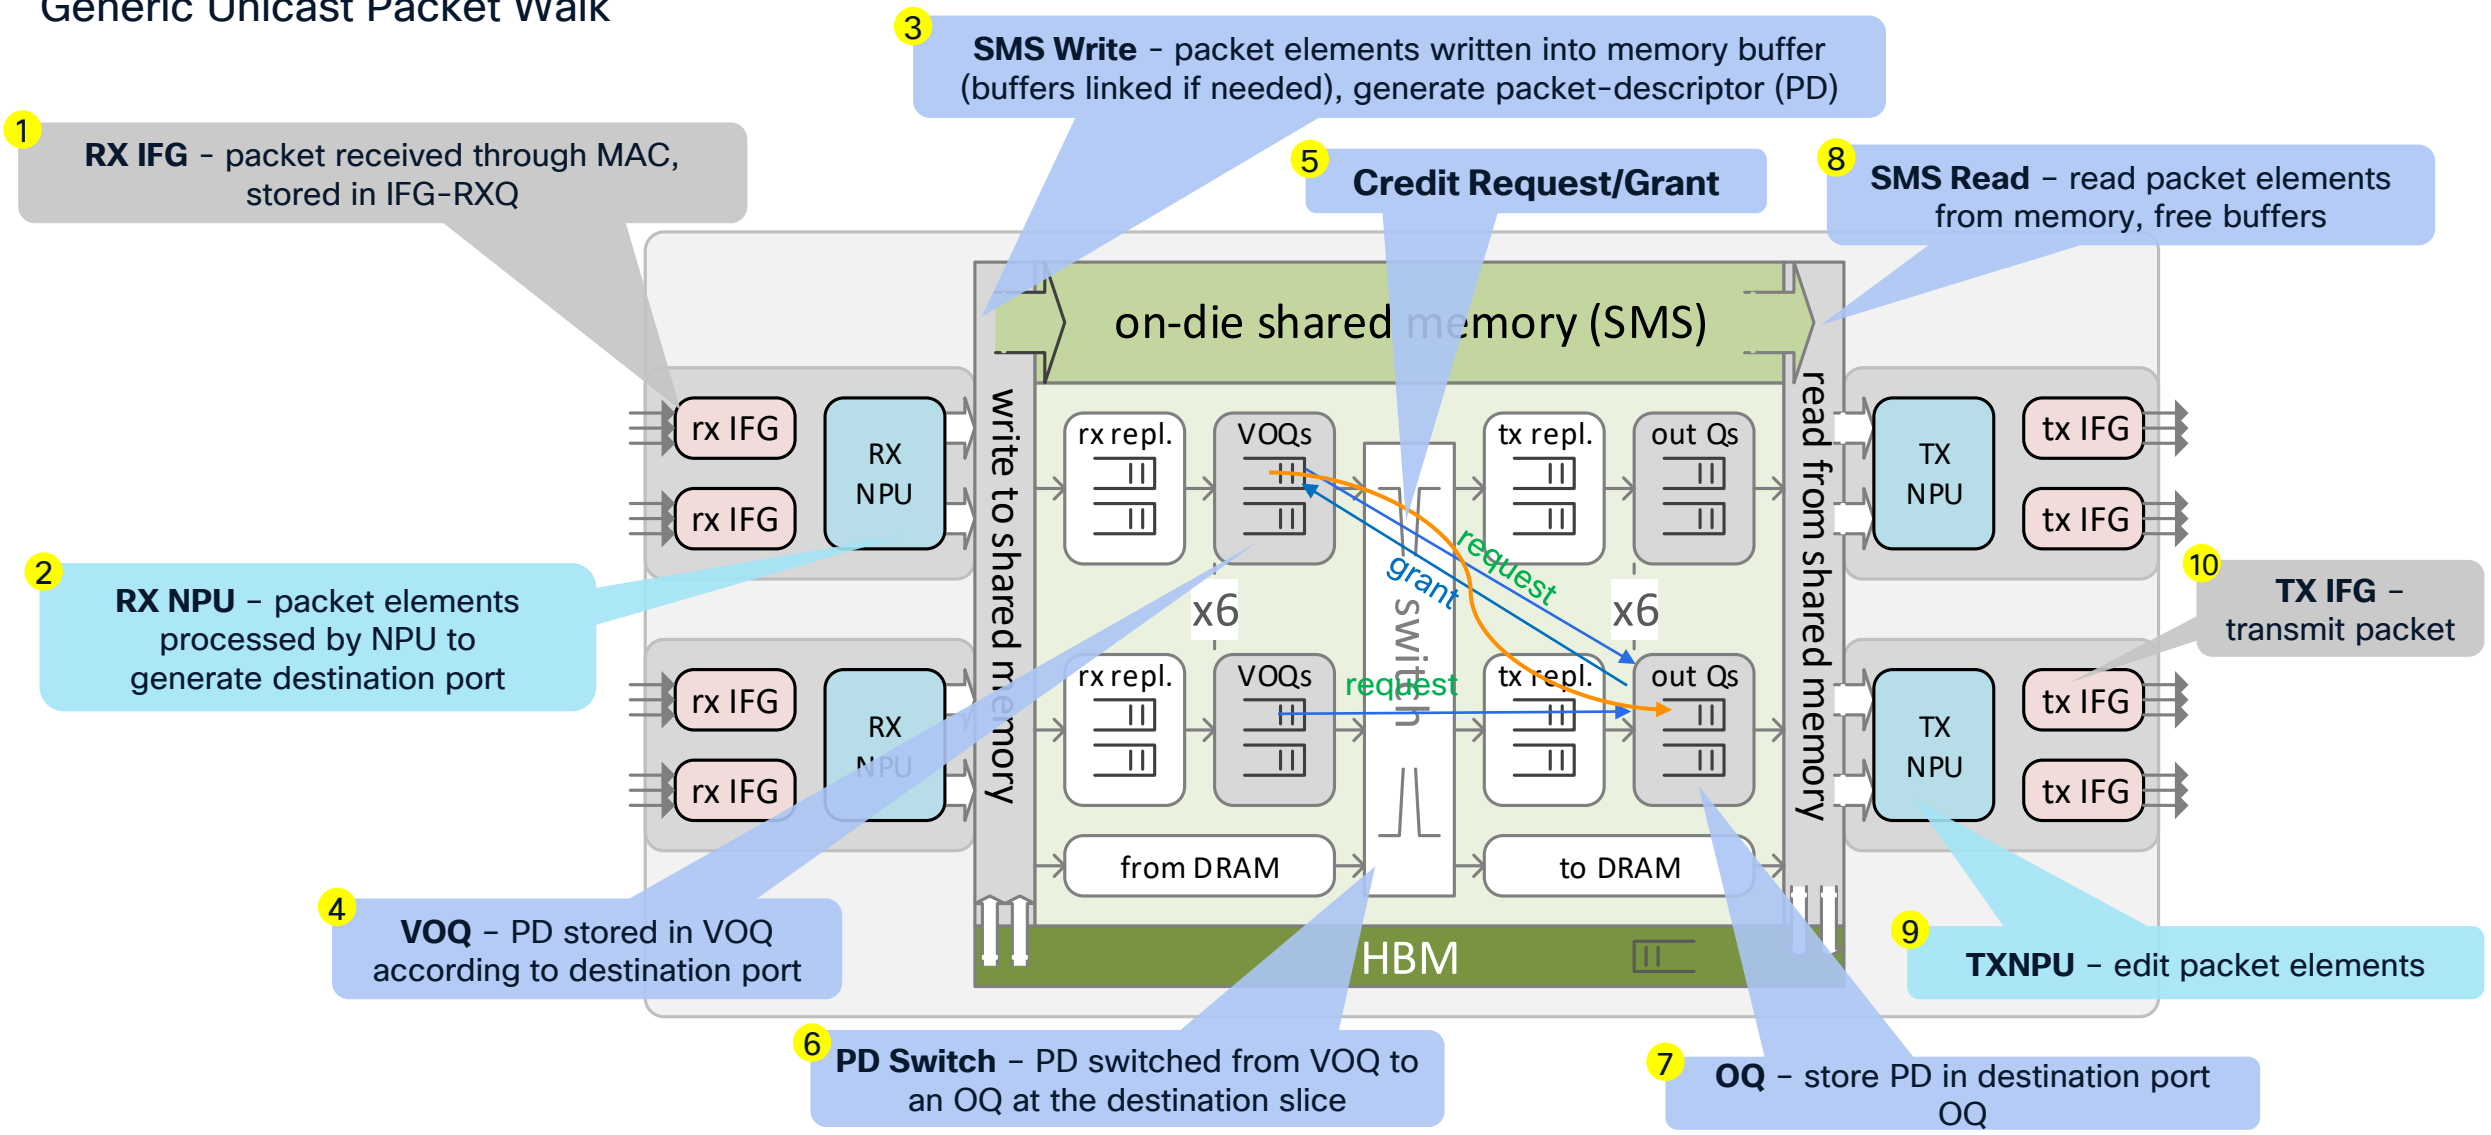

SDM Template

Cisco Silicon One™ – Central Databases

Onboard LPM, CEM & HCAM memory

CDB includes the **Central L2/L3 Forwarding** and **ACL** databases:

LPM – SRAM database for IP/mask routing implemented by **L**ongest **P**refix **M**atch algorithm

- Primarily used by IPv4 and IPv6 unicast routing
- IPv4 use single entry and IPv6 use two entries
- LPM can be extended (from CDB) to HBM

CEM – SRAM database for MAC & Host (/48, /32 or /128), Multicast & Labels implemented by **E**xact **M**atch algorithm

- For features using an exact match (every bit, no mask)
- IPv4 use single entry and IPv6 use two entries
- CEM can be flexibly reallocated for different tables

HCAM – TCAM plus SRAM for Security, QoS and Services **A**ccess **C**ontrol **L**ist entries, Netflow cache...

- For features that use (match criteria + action) policies
- Contains both TCAM (single width) and SRAM (double width)
- Flexibility with different allocation of resources

Resources	A100	K100
LPM (v4/v6)	256K/128K	2M/1M
CEM (v4/v6)	256K/128K	1.25M/620K
HCAM (TCAM/Tiles)	8K (TCAM)	8K (TCAM)
	128K (Tiles)	512K (Tiles)

* Exact scale depends on IP/mask distribution (contiguous vs. random) and hash efficiency. Sample tests with IPv4 GRT is ~1.85M

C9350 default SDM Template

```
C9350# show sdm prefer
Showing SDM Template Info
This is the Access template.
Feature-Name                                        Reserved-Scale
Unicast MAC addresses                               65536
FIB Host Route                                      65536
OG                                                  2048
SGACL Hosts                                         24576
SGACL Cells                                         24576
L3 Multicast entries                                8192
L2 Multicast entries                                8192 (**)
Number of VLANs                                     4094 (**)
Overflow Unicast MAC addresses                      512 (**)
Ipv4/Ipv6 shared unicast routes                     98304 (**)
Overflow shared unicast routes                      196608 (**)
STP Instances                                       1024 (**)
Tunnels                                             1024 (**)
VRF                                                 256 (**)
Max L3 adjacency                                    65536 (**)
Max L3 Interface                                    8192 (**)
Max MPLS TE TUNNEL                                  4096 (**)
MAX NAT ENTRIES                                     63488
Security ACL OVERALL                                36864
QOS ACL OVERALL                                     5120
PBRNAT ACL OVERALL                                  16384
MIRROR ACL OVERALL                                  1024
FNF ACL OVERALL                                     2048
FNF FLOW OVERALL                                    65536

(**) - SDM library is referred to only obtain scale

Resource scale information
EM                                                  262144
HCAM                                                131072
OTHER                                               262144
```

C9610 Default SDM Template

Reference

This is the Core template.

Feature-Name	Reserved-Scale
Unicast MAC addresses	**131072**
FIB Host Route	**131072**
OGACL	2048
SGACL Hosts	24576
SGACL Cells	24576
Max MPLS Label	65536
L3 Multicast entries	32768
L2 Multicast entries	16384
Number of VLANs	4094 (**)
Overflow Unicast MAC addresses	512 (**)
Overflow L2 Multicast entries	512 (**)
Overflow L3 Multicast entries	512 (**)
Ipv4/Ipv6 shared unicast routes	**131072 (**)**
Overflow shared unicast routes	**1000000 (**)**
STP Instances	4094 (**)
Tunnels	1024 (**)
VRF	3839 (**)
Max MPLS VPN Routes Per-Vrf label mode	1000000 (**)
Max MPLS VPN Routes Per-Prefix label mode	65536 (**)
Max L3 adjacency	131072 (**)
Max L3 Interface	12288 (**)
Max MPLS TE TUNNEL	4096 (**)
Security ACL IN	**21504**
Security ACL OUT	**21504**
QOS ACL IN	5120
QOS ACL OUT	5120
PBRNAT ACL IN	8192
PBRNAT ACL OUT	8192
MIRROR ACL IN	1024
MIRROR ACL OUT	1024
FNF ACL IN	2048
FNF ACL OUT	2048
FNF FLOW IN	**32768**
FNF FLOW OUT	**32768**

(**) - SDM library is referred to only obtain scale

Resource scale information

EM	1280000
HCAM	131072
OTHER	0

CLI: *show sdm prefer*

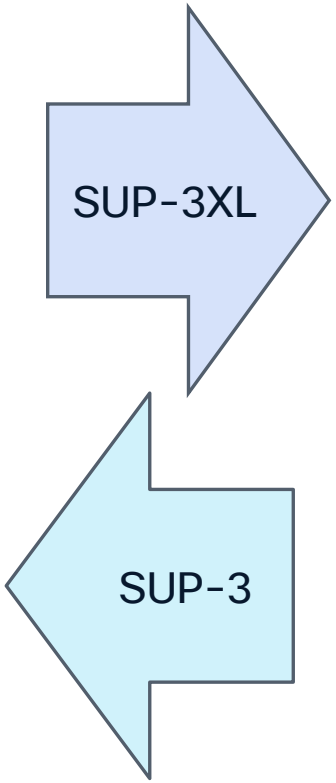

This is the Core template.

Feature-Name	Reserved-Scale
Unicast MAC addresses	**131072**
FIB Host Route	**131072**
OGACL	2048
SGACL Hosts	24576
SGACL Cells	24576
Max MPLS Label	131072
L3 Multicast entries	32768
L2 Multicast entries	16384
Number of VLANs	4094 (**)
Overflow Unicast MAC addresses	512 (**)
Overflow L2 Multicast entries	512 (**)
Overflow L3 Multicast entries	512 (**)
Ipv4/Ipv6 shared unicast routes	**131072 (**)**
Overflow shared unicast routes	**2000000 (**)**
STP Instances	4094 (**)
Tunnels	1024 (**)
VRF	3839 (**)
Max MPLS VPN Routes Per-Vrf label mode	1000000 (**)
Max MPLS VPN Routes Per-Prefix label mode	65536 (**)
Max L3 adjacency	131072 (**)
Max L3 Interface	12288 (**)
Max MPLS TE TUNNEL	4096 (**)
Security ACL IN	**38912**
Security ACL OUT	**38912**
QOS ACL IN	8192
QOS ACL OUT	8192
PBRNAT ACL IN	16384
PBRNAT ACL OUT	16384
MIRROR ACL IN	1024
MIRROR ACL OUT	1024
FNF ACL IN	2048
FNF ACL OUT	2048
FNF FLOW IN	**65536**
FNF FLOW OUT	**65536**

(**) - SDM library is referred to only obtain scale

Resource scale information

EM	1280000
HCAM	131072
OTHER	0

HCAM

HCAM (Algorithmic Hash based TCAM)

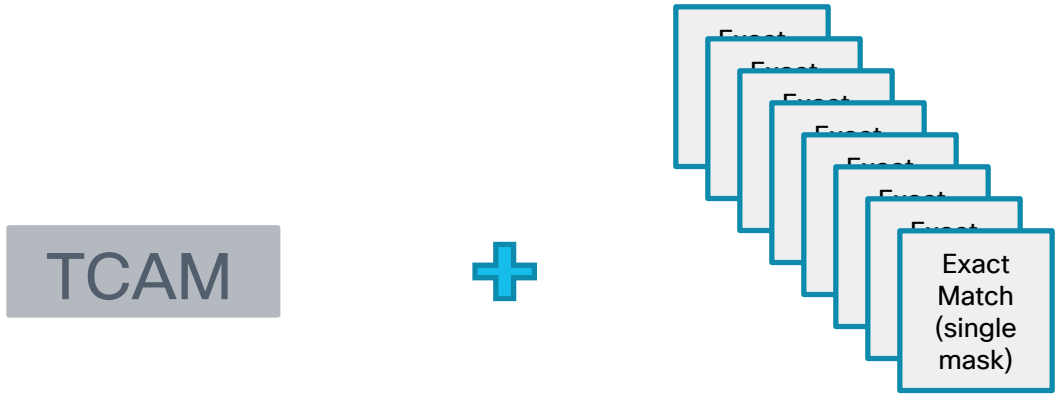

TCAM (8K)	EM = Tile (8x 16K)
Key based – variable mask	**Hash based – full mask**
Supports flexible matching logic	Simple, specific searches
Complex & versatile searches	Rigid exact matching logic
Less-Dense (fewer bits in same area)	More-Dense (more bits in same area)
Faster lookup speed	Slightly slower speed
Power & Heat intensive	Lower Power & Heat
Expensive to scale	**Less costly to scale**

HCAM combines the **speed** and **capability** of TCAM and the **flexibility** of hash table

- ACEs which share the same popular masks can be searched with a relatively small number of exact-match tiles
- Remaining ACEs with unpopular mask combinations can be installed in the TCAM

ACLs implementation with HCAM (TCAM+ Tile)

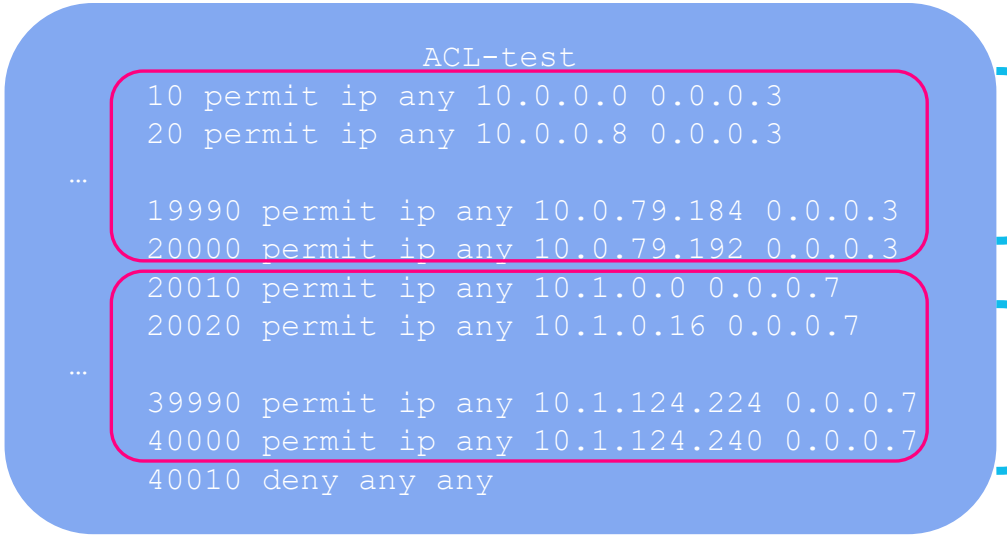

Tile 0

Index	ACL label	Rule
1	ACL-test	Permit ip any 10.0.0.0 0.0.0.3
...	ACL-test	...
2000	ACL-test	Permit ip any 10.0.79.192 0.0.0.3

Tile 1

Index	ACL label	Rule
1	ACL-test	Permit ip any 10.1.0.0 0.0.0.7
...	ACL-test	...
2000	ACL-test	Permit ip any 10.1.124.240 0.0.0.7

TCAM

Index		
1	ACL-test	
2		
3		
4		
...		

Flexible Scale with HCAM

ACL and FNF entry space sharing

Algorithmic TCAM – the Best of Both Worlds

HCAM = TCAM + SRAM

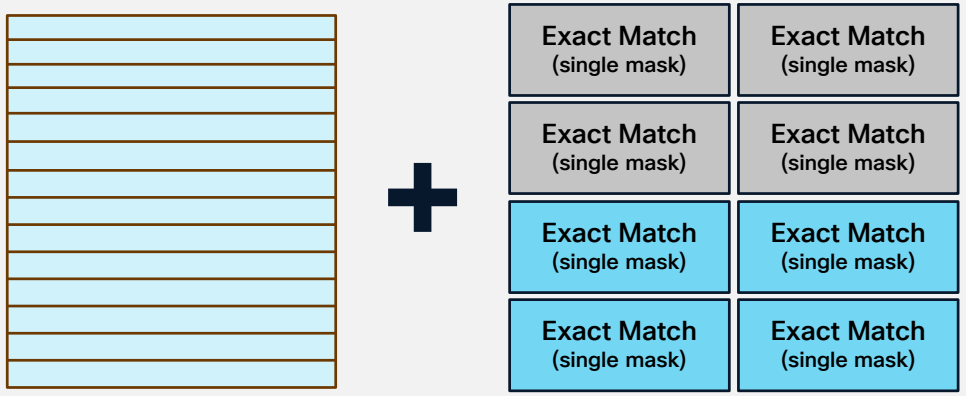

C9350	C9610
System (total) • **16K** Ingress ACL Entries • **16K** Egress ACL Entries	System (total) • **32K** Ingress ACL Entries • **32K** Egress ACL Entries
Per ASIC (x ASICs) • **32K** Ingress NetFlow • **32K** Egress NetFlow	Per ASIC (x ASICs) • **64K** Ingress NetFlow • **64K** Egress NetFlow

***Default SDM: Half of each tile shared**

Half reserved for ACL entries

Half reserved for NetFlow cache

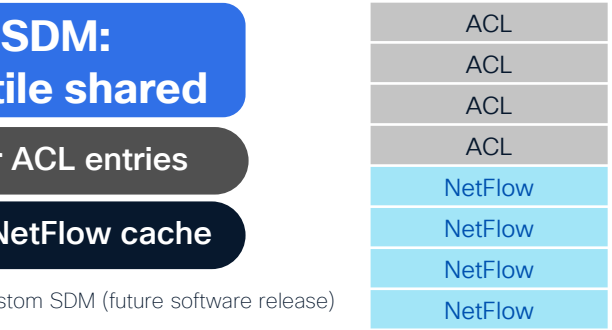

A100	K100
1x HCAM Shared by RX & TX	4x HCAM 2x RX 2x TX

* Based on Default SDM. ASIC capable of higher scale with Custom SDM (future software release)

Security

What Pervasive Security means to us

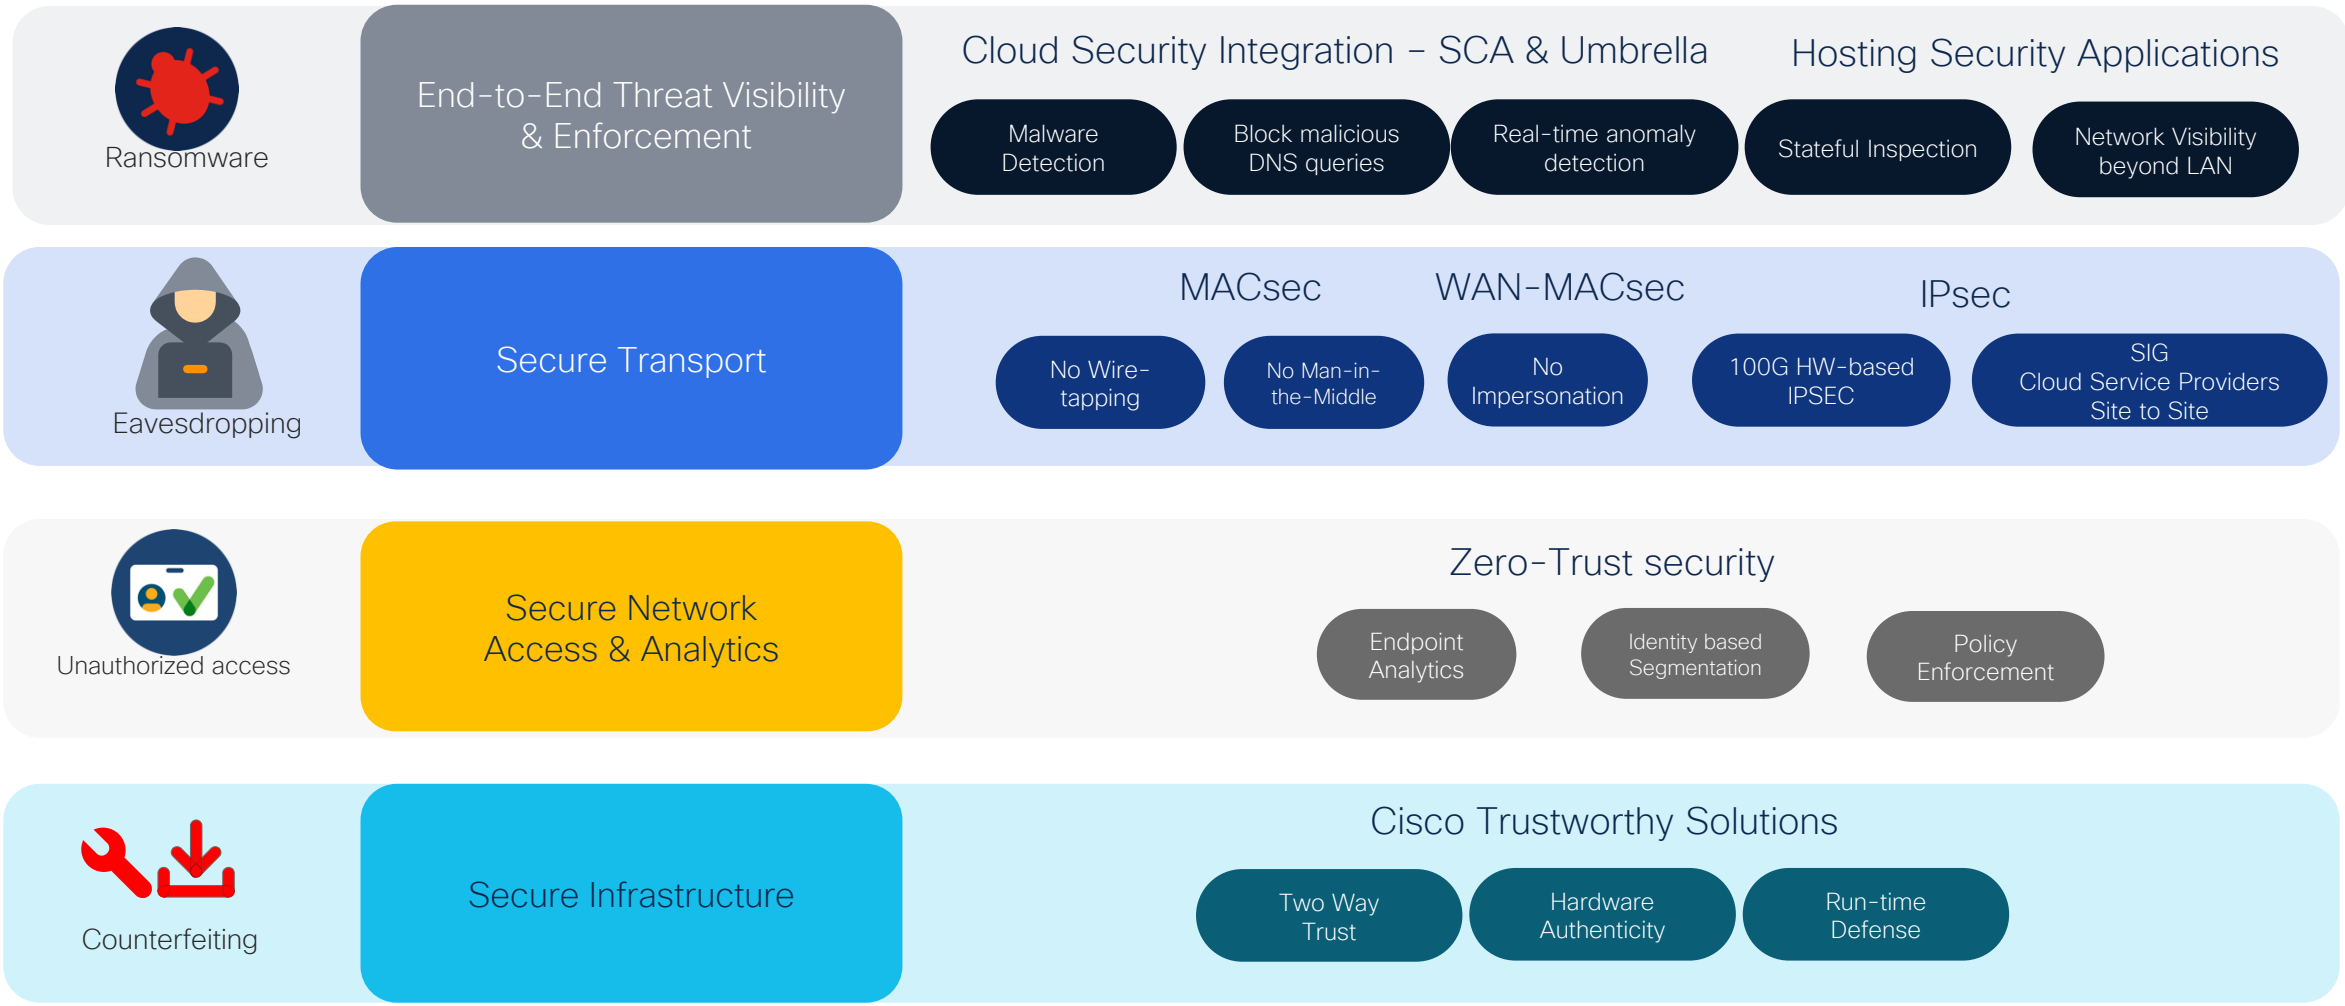

Trustworthy Solutions on Cisco C9000 Smart Switches

What we are doing with Trust 2.0

Why

- Cyber attacks targeting network infrastructure are a growing concern.
- Digitization causes an increased need for protection and integrity of network infrastructure

What

Cisco Trustworthy Technologies provide a foundational layer of security

- ✓ Reduces vulnerabilities
- ✓ Enhances security
- ✓ protects against counterfeit and software modification
- ✓ verifies that Cisco networking products are operating as intended.

Trust 1.0 — Embedded & Implicit

- Tell us our hardware is safe and authentic
- But with attacks always evolving
- Customer requirements and awareness change as well

Trust 2.0 — Transparent & Customer Managed

- How do customer verify authenticity for themselves?
- Allow customers to manage their own trust

PQC initiatives within Trustworthy Systems

PQC Hardware Ready

Attack	Protect with Cisco Trustworthy Solutions	PQC Initiatives for Solutions
Counterfeit Hardware	**Secure Device ID & Trust Anchor** Built-in HW ID Root-of-trust, HW is Cisco, Secure PnP	• **SUDI**: PQC-signed certs (e.g., ML-DSA-87)
Compromised Software	**Image Signing** The Open IOS XE is authentic & unmodified	• ML-DSA image signing key storage to enable verification of OS
Compromised Boot Code	**Secure Boot** The boot-code is authentic & unmodified	• LMS signing and verification capability on bootloader

Note: * Hw Capable –Sw Support on Roadmap

MAC Security (LAN MACsec)

- Hop-by-Hop vs End-to-End “Bump-in-the-wire” model
  - -Packets are decrypted on ingress port
  - -Packets are encrypted on egress port
- Allows the network to continue to perform all the packet inspection features currently used.

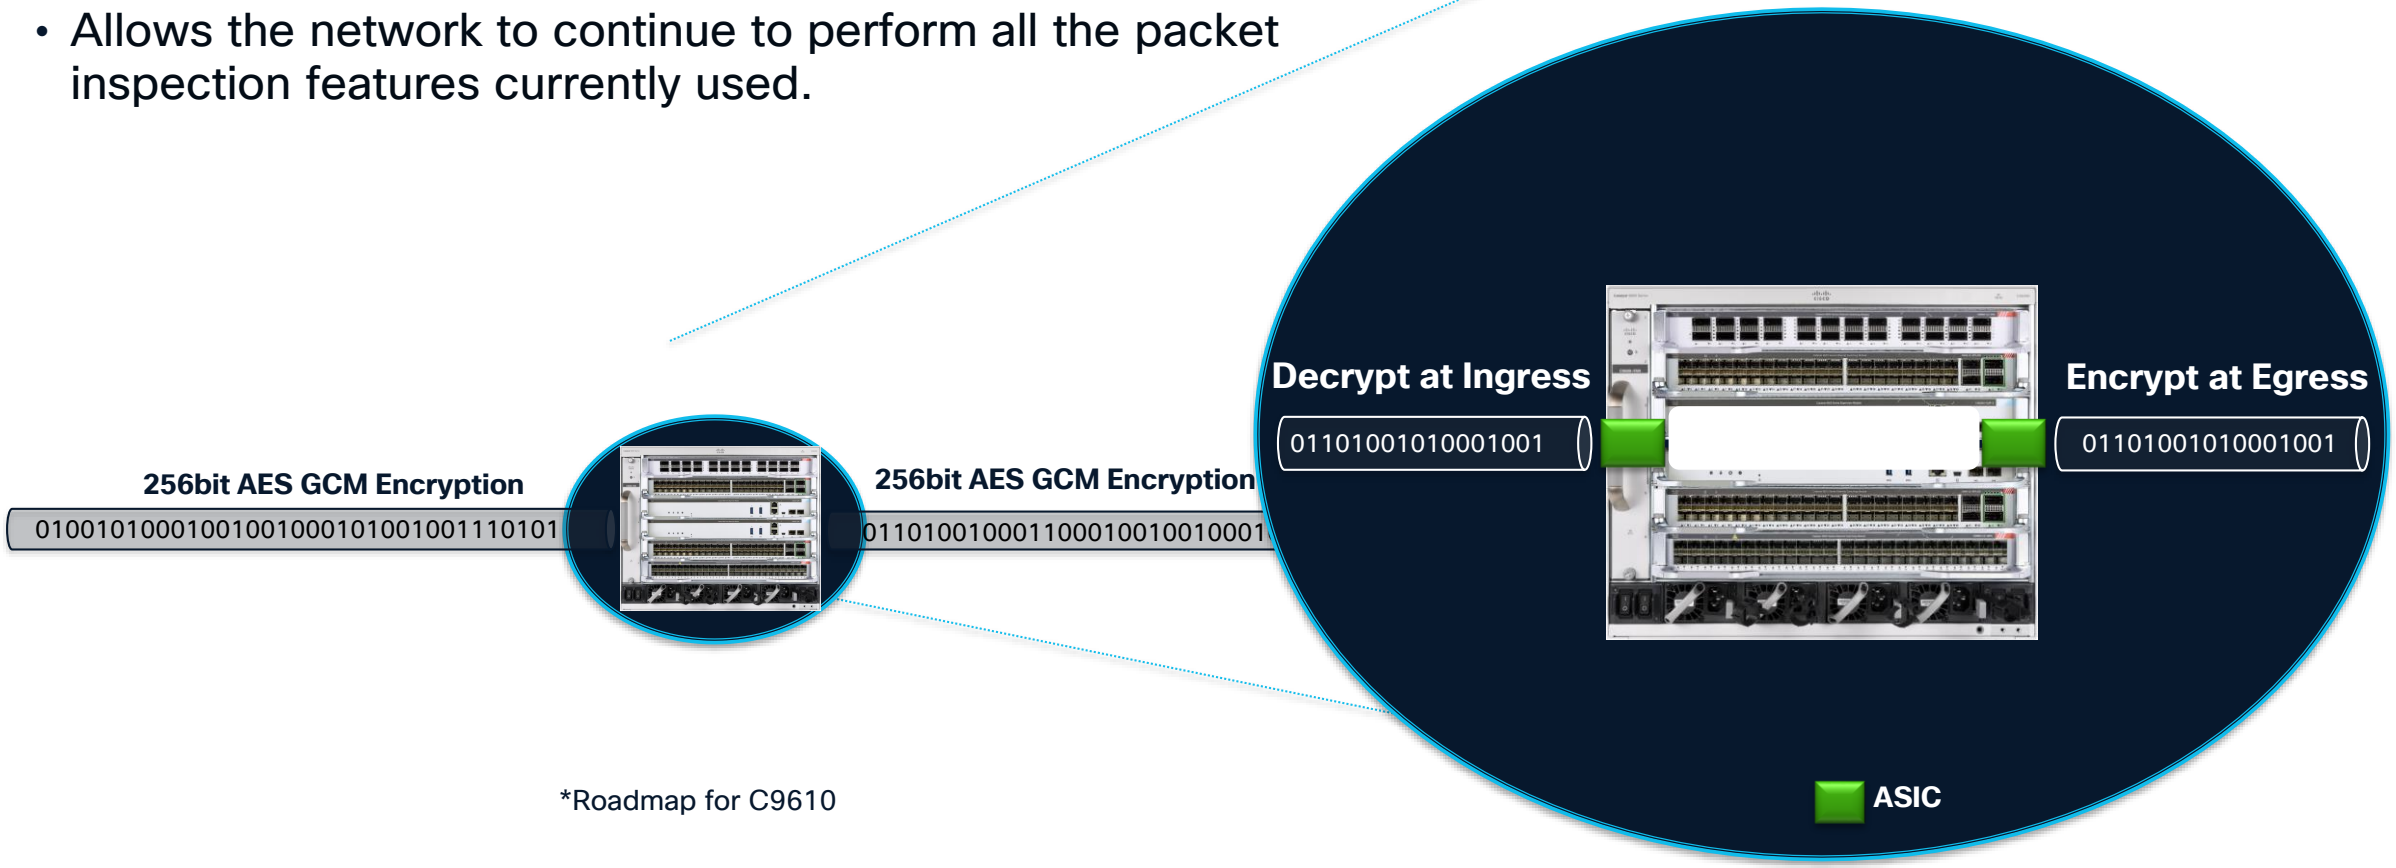

*Roadmap for C9610

Object Group Access List (OGACL)

- OGACLs allow user to classify IP addresses, protocol & ports into groups and use these object-groups in Access List Entries.
- In large networks, the number of ACLs can be large (hundreds of lines) and difficult to configure and manage, especially if the ACLs frequently change
- This reduces the size of ACL and improves the manageability significantly for network admins.

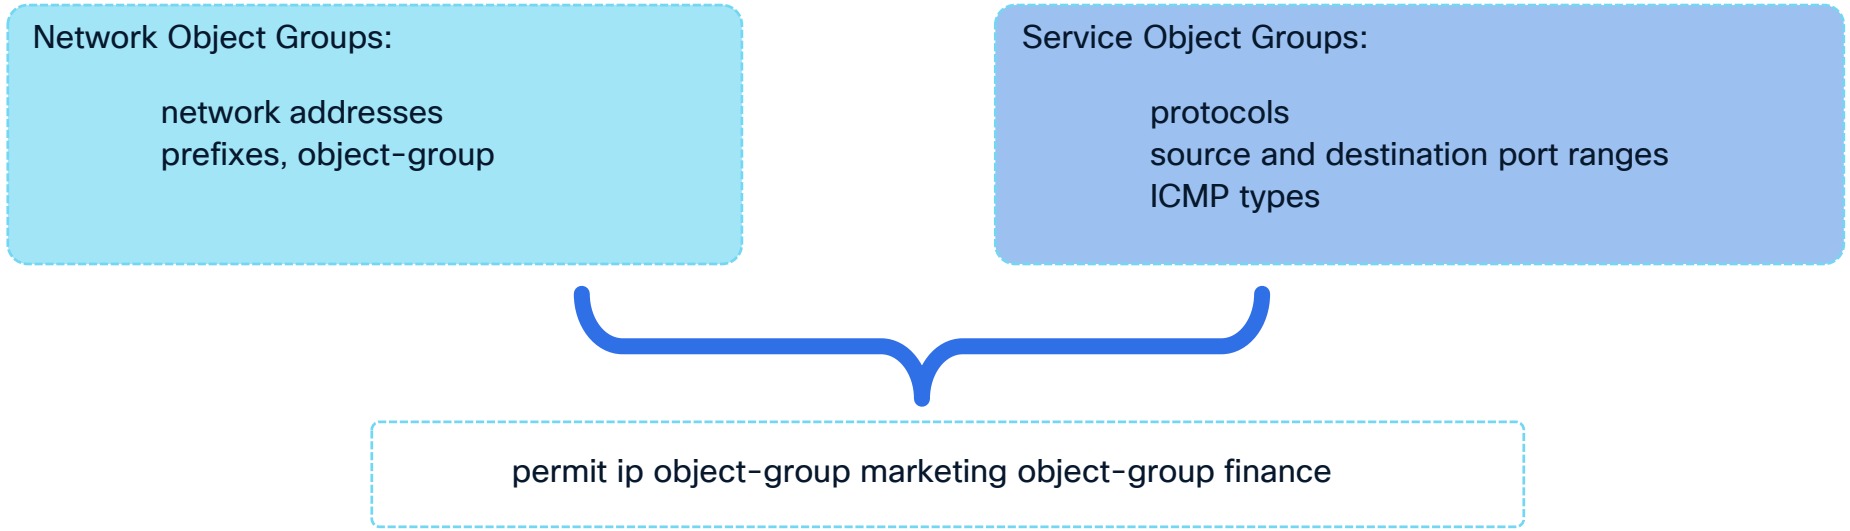

Traditional ACL vs OGACL

Traditional security ACL

```
ip access-list extended hospital_security_ACL
  permit ip host 11.96.172.238 host 10.10.80.182
  permit ip host 11.96.172.238 host 10.10.90.190
  permit ip host 11.96.172.238 11.96.172.0 255.255.255.0
  permit ip host 11.96.172.238 host 11.96.172.254
  permit ip host 11.96.172.240 host 10.10.80.182
  permit ip host 11.96.172.240 host 10.10.90.190
  permit ip host 11.96.172.240 11.96.172.0 255.255.255.0
  permit ip host 11.96.172.240 host 11.96.172.254
```

Object Group ACL

```
object-group network OPD_DEPT
     host 11.96.172.238
     host 11.96.172.240

object-group network ACCOUNTING_DEPT
    host 10.10.80.182
    host 10.10.90.190
    11.96.172.0 255.255.255.0
    host 11.96.172.254

ip access-list extended hospital_ACL
     permit tcp object-group OPD_DEPT
object-group ACCOUNTING_DEPT
```

ACL - Optimizing the hardware resources

Optimized resource utilization using OG-ACL on Silicon One

```
object-group network TestGroup
 host 10.5.64.140
 host 10.5.64.86
 host 10.5.75.124
 10.10.170.32 255.255.255.224
 10.179.252.0 255.255.252.0
 10.5.200.240 255.255.255.240
 10.53.199.0 255.255.255.0
 10.56.100.224 255.255.255.240
... Truncated (1,000 entries total)
 host 10.56.223.56
```

```
ip access-list extended SampleACL
 permit ip object-group TestGroup any
 deny ip any any (implicit)
!
```

```
interface Vlan100
 ip access-group SampleACL in
```

Contents of the Object-groups are expanded to the CEM space.

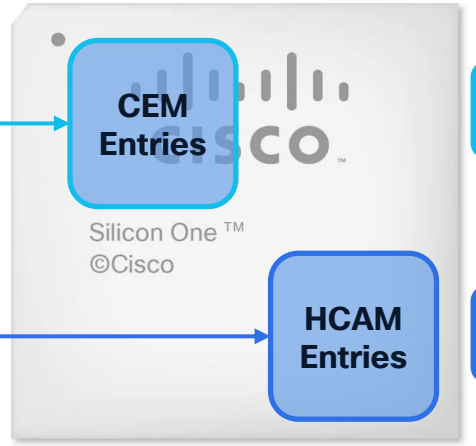

1K CEM entries consumed

3 HCAM entries consumed

Entries in object-groups use CEM space instead of TCAM space

QoS

Buffer types – Silicon One vs UADP

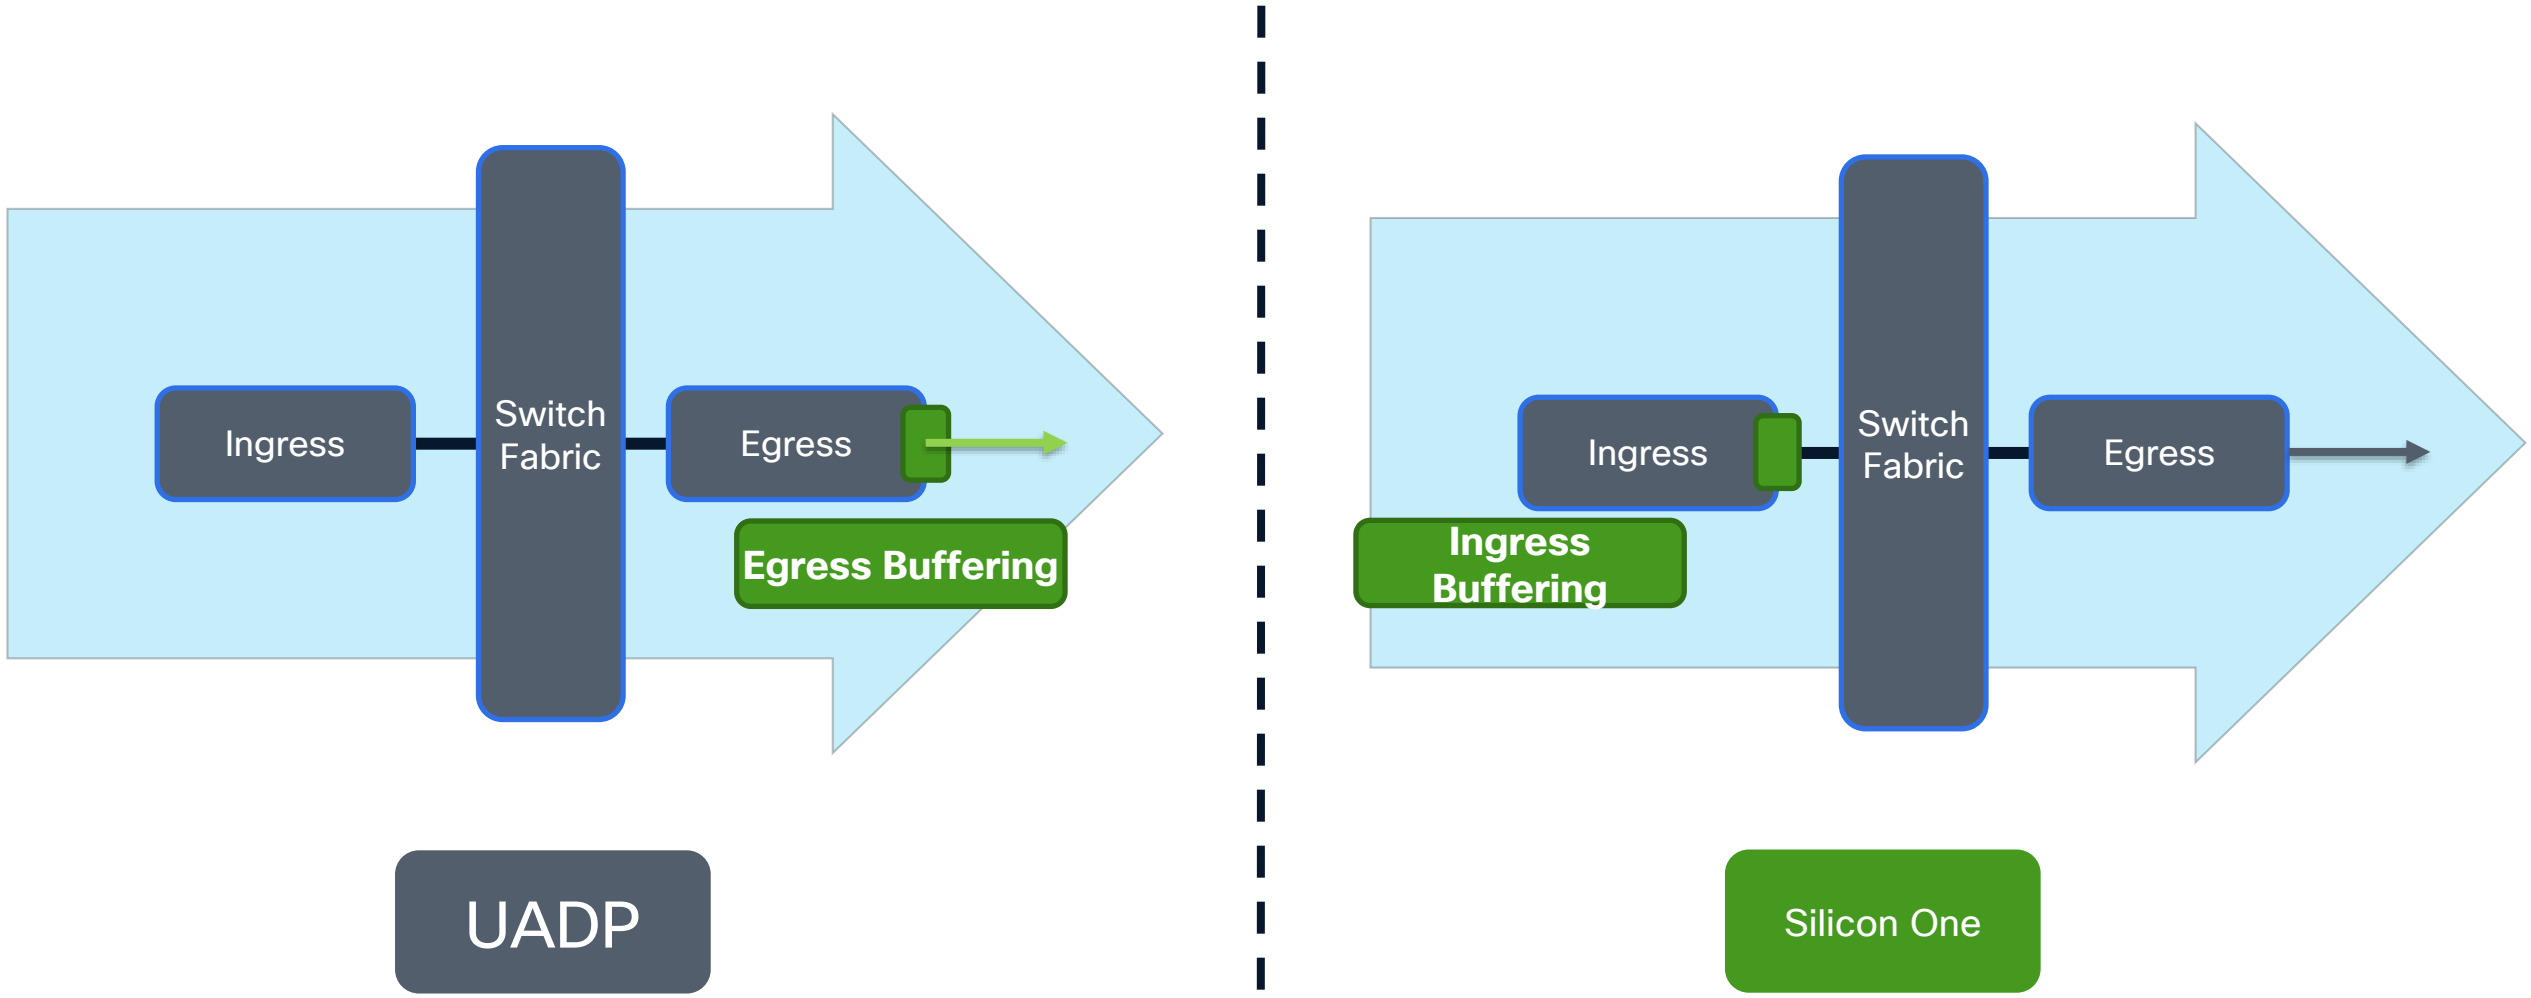

Ingress buffering – Head of Line Blocking

What is the Problem?

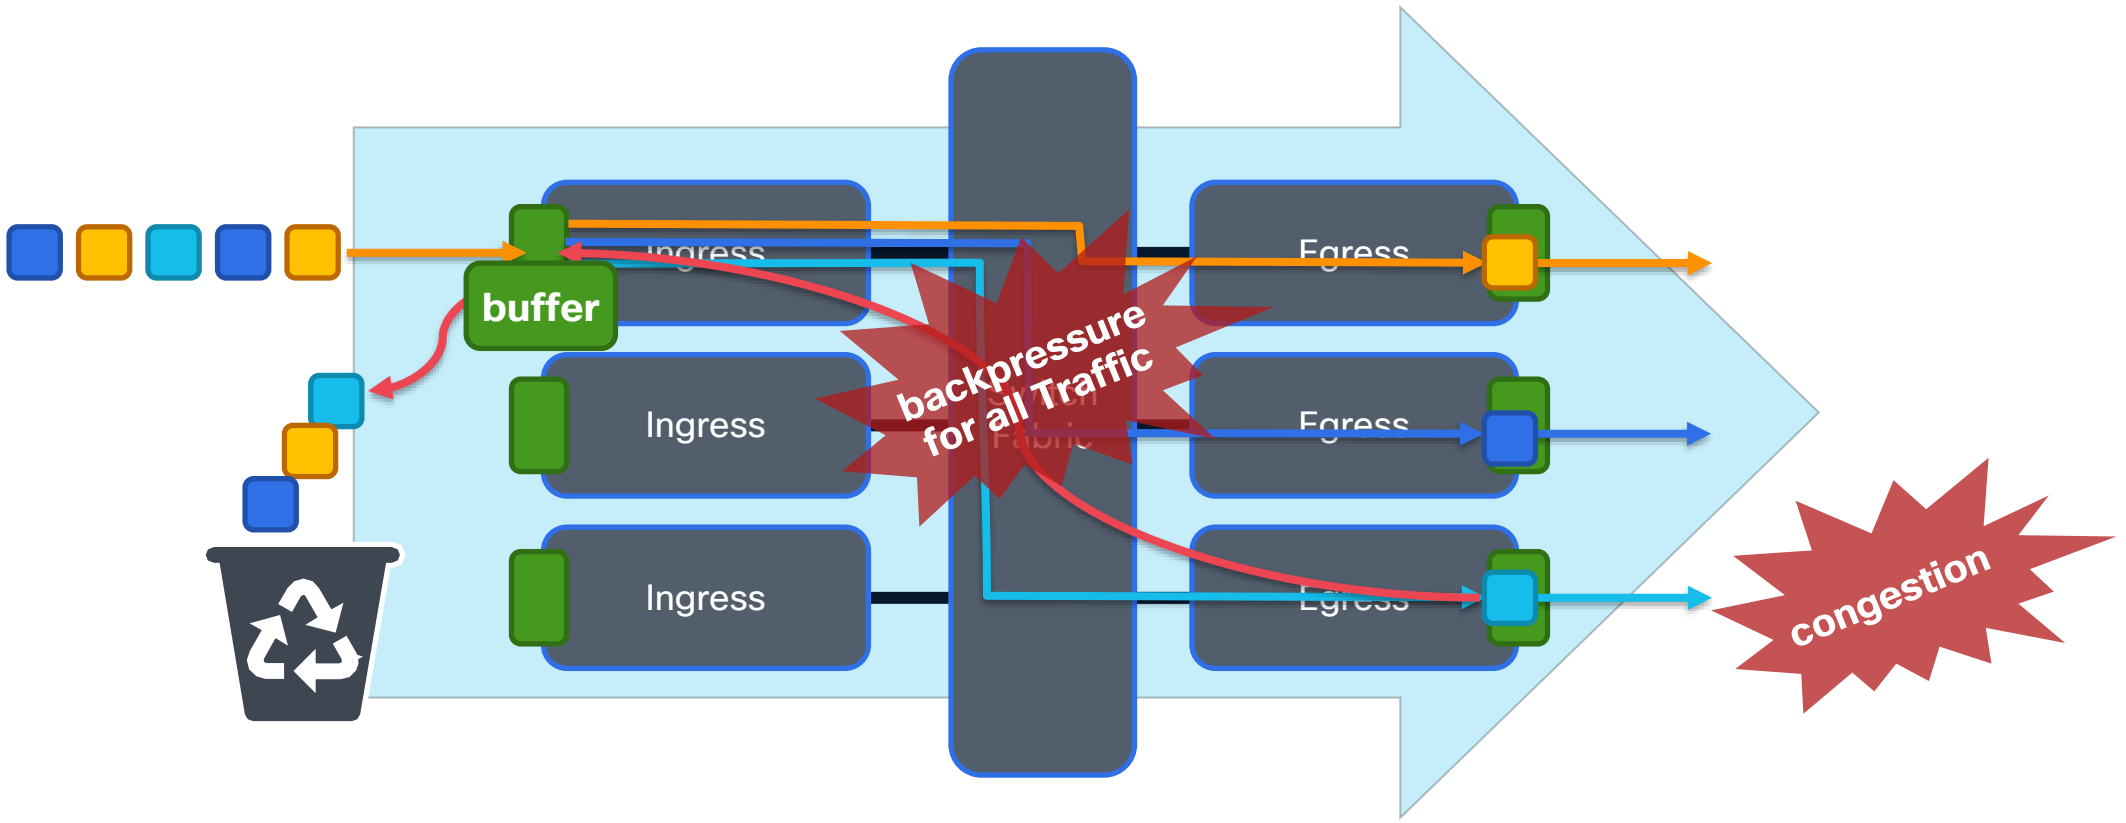

Virtual Output Queuing – Silicon One ASIC

Address Head of Line block

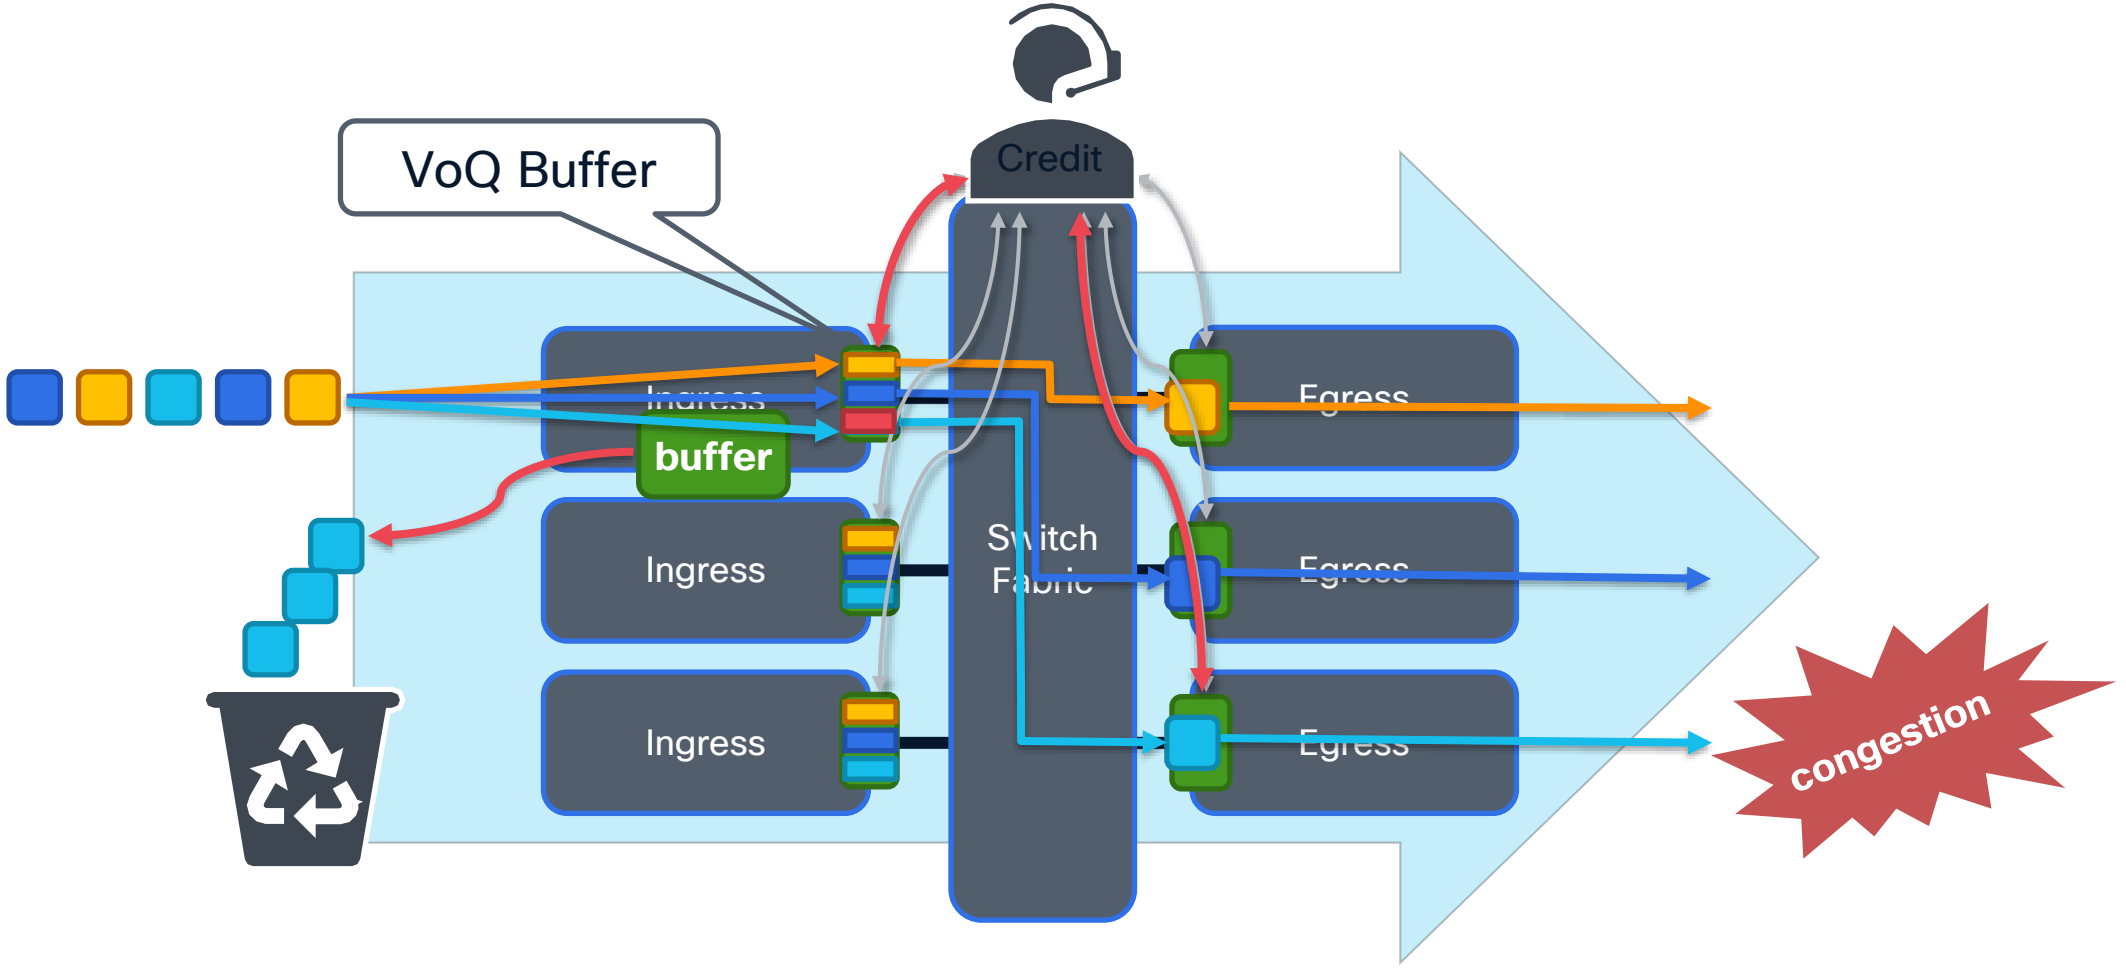

Reference

Silicon One QoS

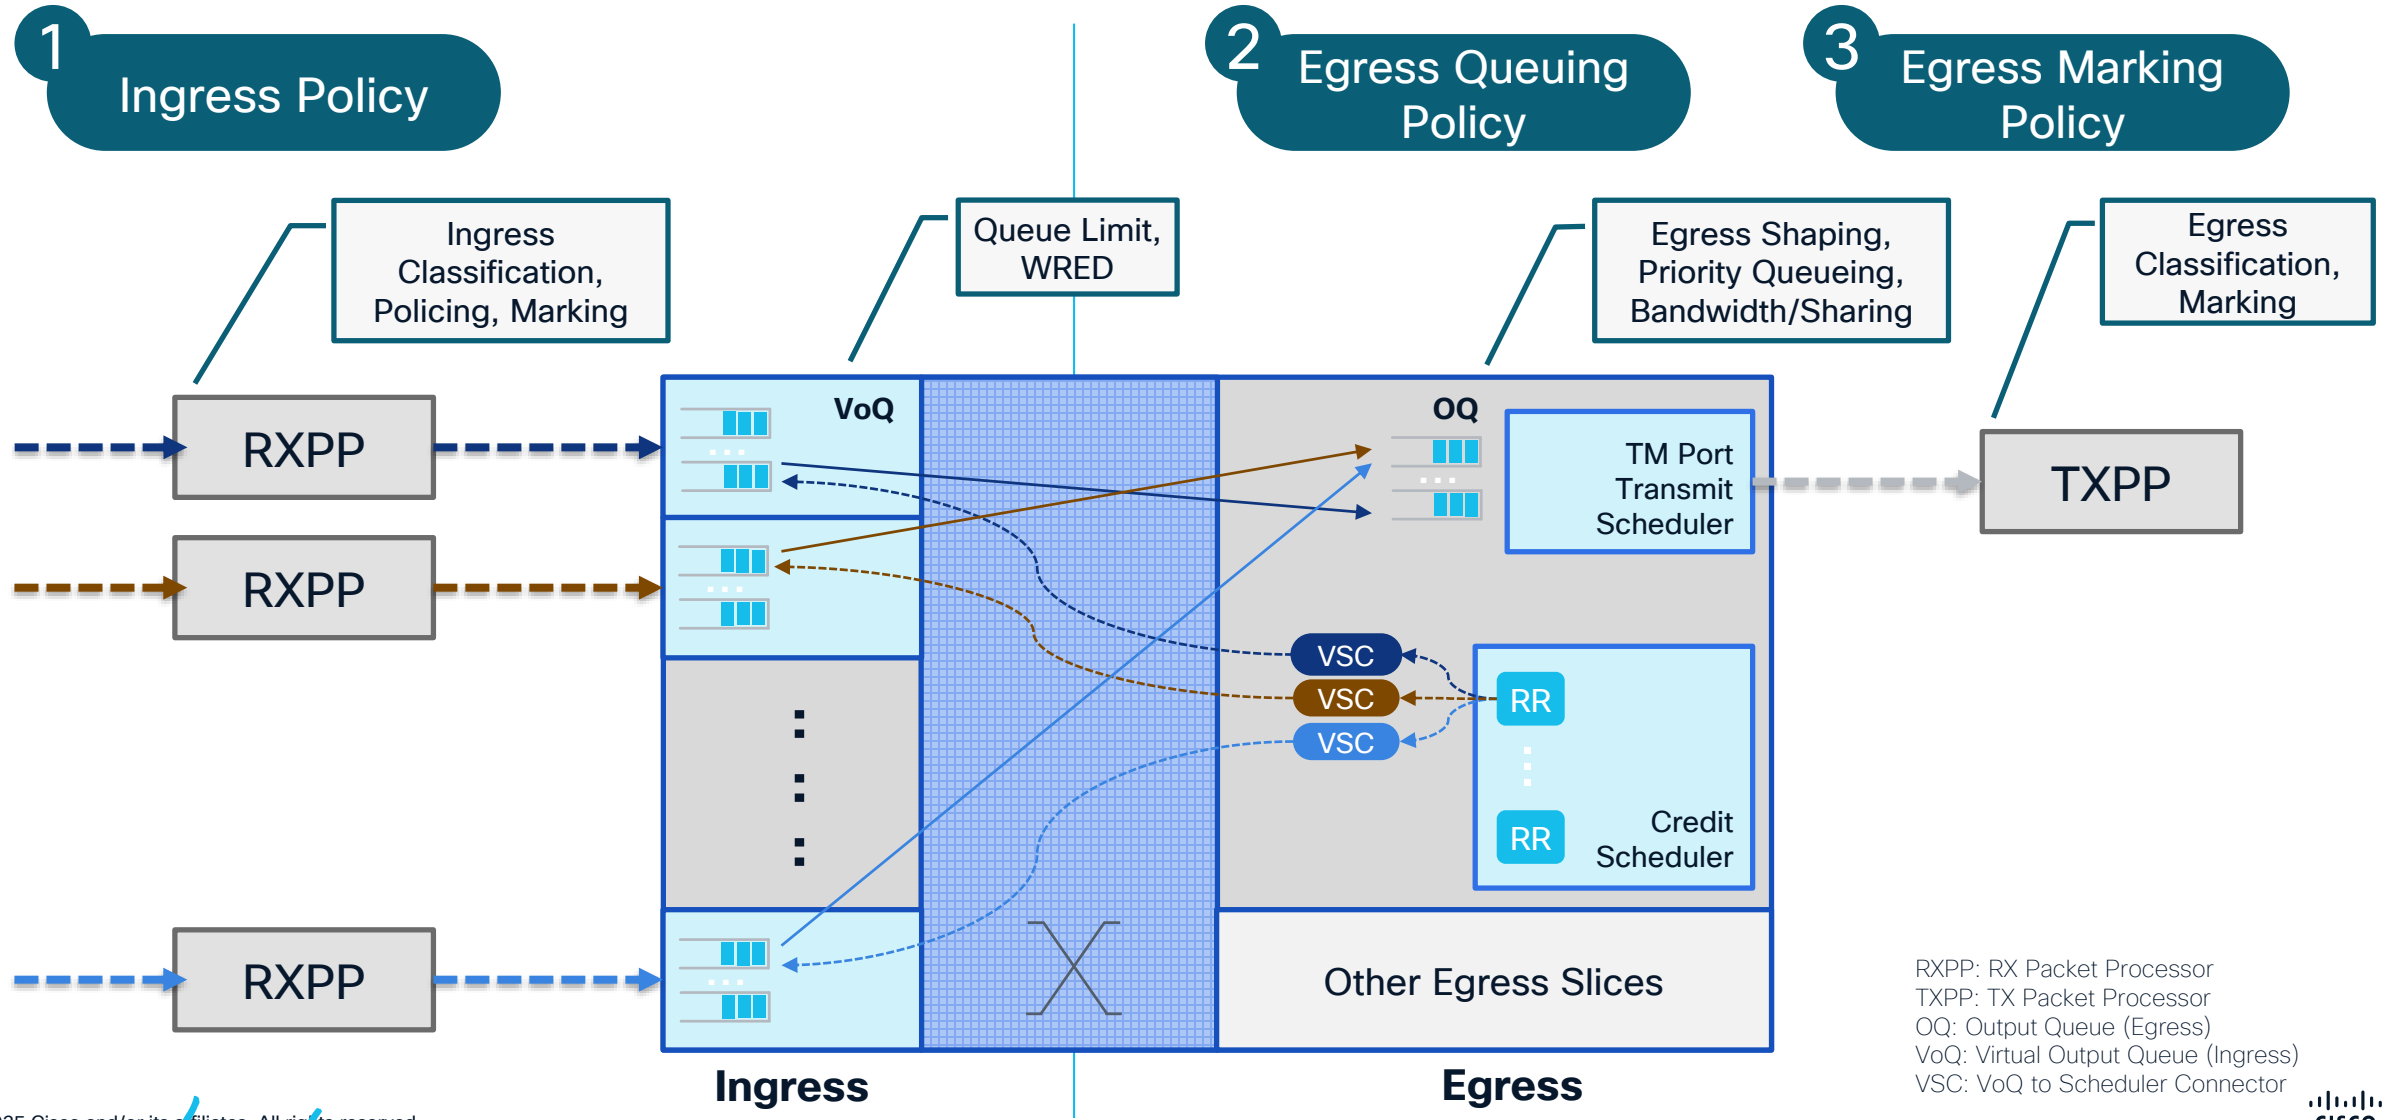

RXPP: RX Packet Processor
TXPP: TX Packet Processor
OQ: Output Queue (Egress)
VoQ: Virtual Output Queue (Ingress)
VSC: VoQ to Scheduler Connector

Queues and Traffic class

- 8 Queue and 7 can be priority queues
- Uses traffic classes to map traffic to different queues. ”**traffic-class**” is local significant to the switch only
- Traffic-class 0 - lowest priority (maps to class-default); traffic-class 7 - highest priority (traffic-class 1 to 6 can be non-priority)
- **Ingress policies** classify packets to specific traffic classes
- Class-maps in **egress queuing policy** can **ONLY** match traffic-class

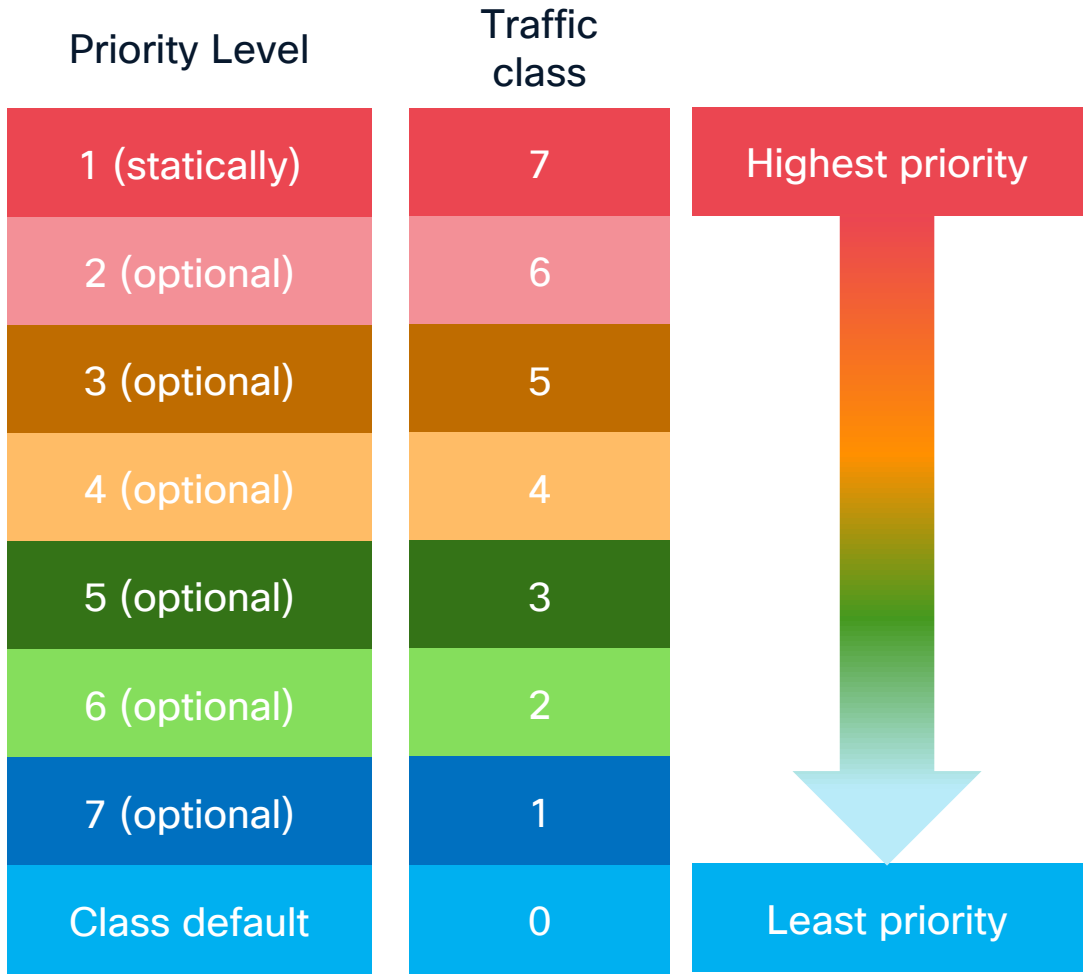

K100 Buffers

- Two different buffers to address two different requirements.
  - 64 MB of Shared Memory Sub-system (SMS) buffers:
    - Low latency packet queueing (video/voice packets)
    - Shallow specialized pool of buffers for quick queueing.
  - 8 GB of High Bandwidth Memory (HBM) buffers:
    - Deep pool of on-demand buffers for guaranteed delivery.
    - Reserve to absorb occasional bursts or address speed over-subscription between ingress and egress.

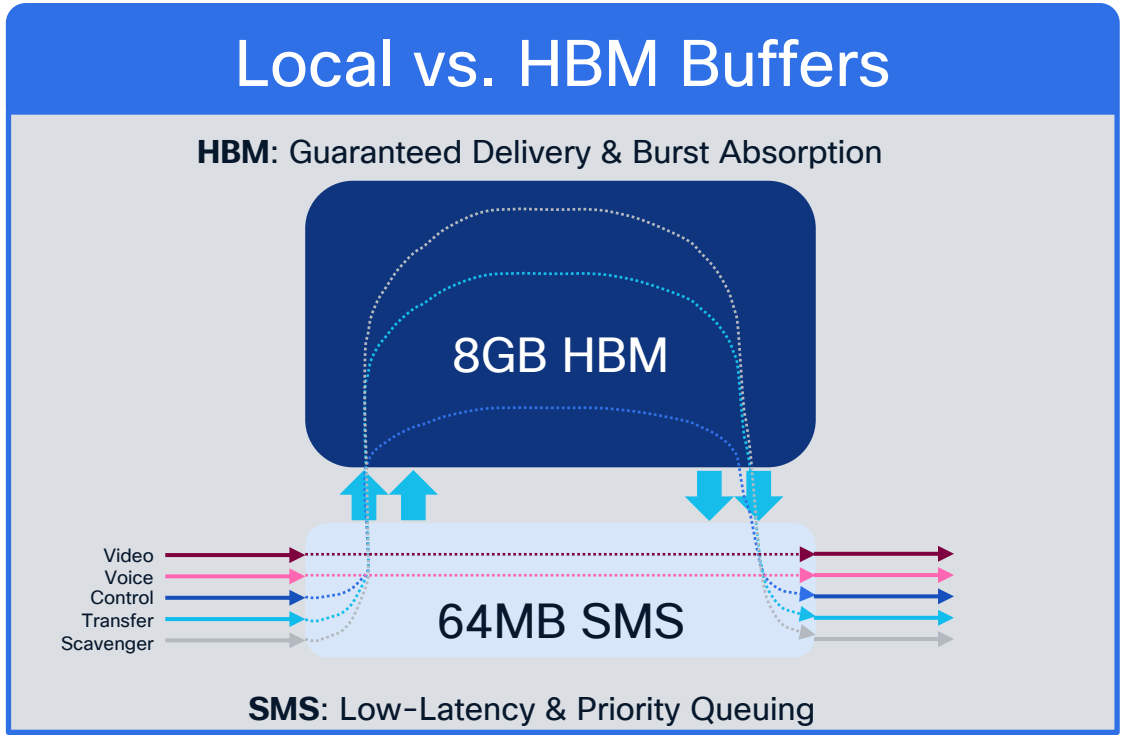

- Packet will always hit the SMS buffers first.
- SMS send the packet to HBM if additional buffers are needed.
- HBM **CANNOT** send the packet to the output queue, it has to be sent to the SMS again to be sent to the egress.

High Availability

High availability

Protect business continuity

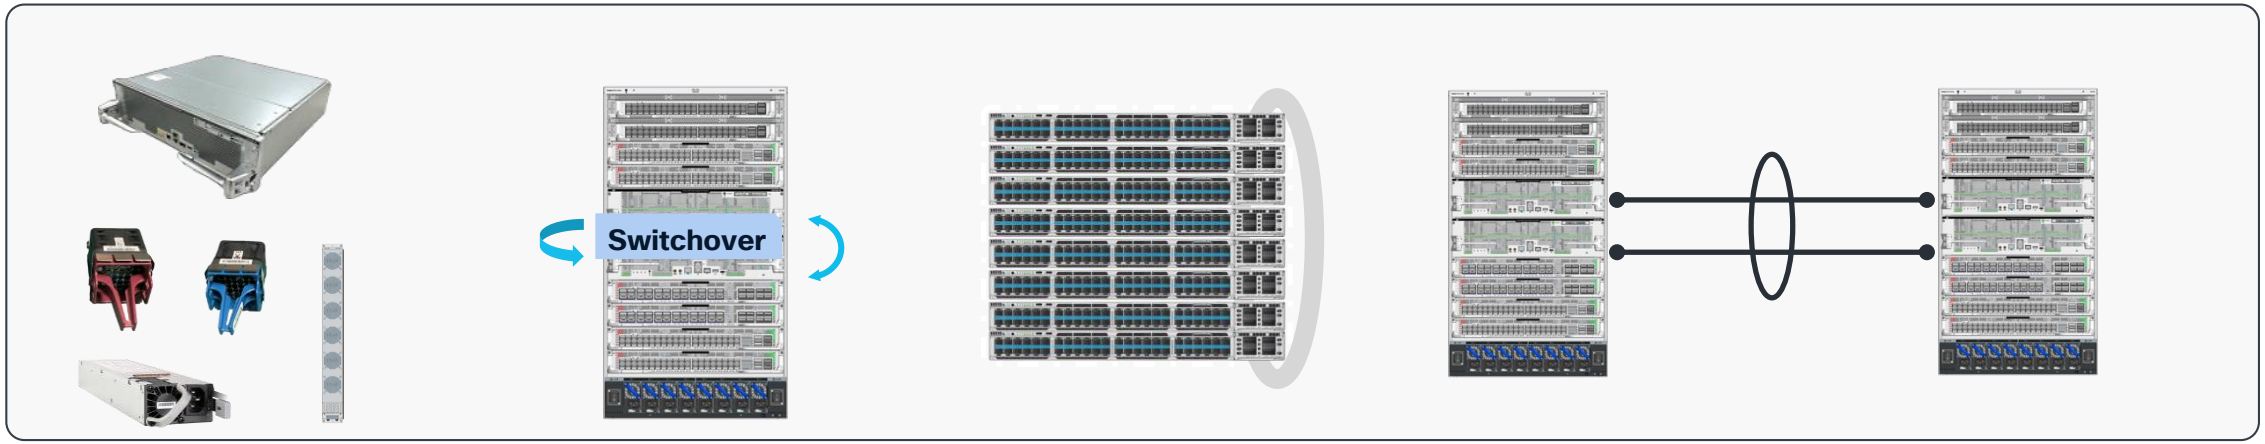

Physical Resiliency	Design Resiliency	Operational Resiliency		
Redundant hardware • Redundant power supplies • Redundant fan / fan tray • Redundant supervisors	**Infrastructure resilience with Next Generation StackWise / Virtual *** • Multi-chassis EtherChannel (MEC) provides hardware-based failover	**Resilient L3 topologies with Non-Stop Forwarding (NSF)** • NSF support for OSPF, EIGRP, ISIS, BGP	**Sub-second failover with Stateful Switch Over (SSO)** • Between supervisors within chassis (<5ms) • Between chassis with NG StackWise Virtual*	**Minimize upgrade downtime** • SMU • ISSU ** / XFSU ** • GIR

* Limited availability with C9610R at GA

** Future Software Release

NG StackWise – New architecture

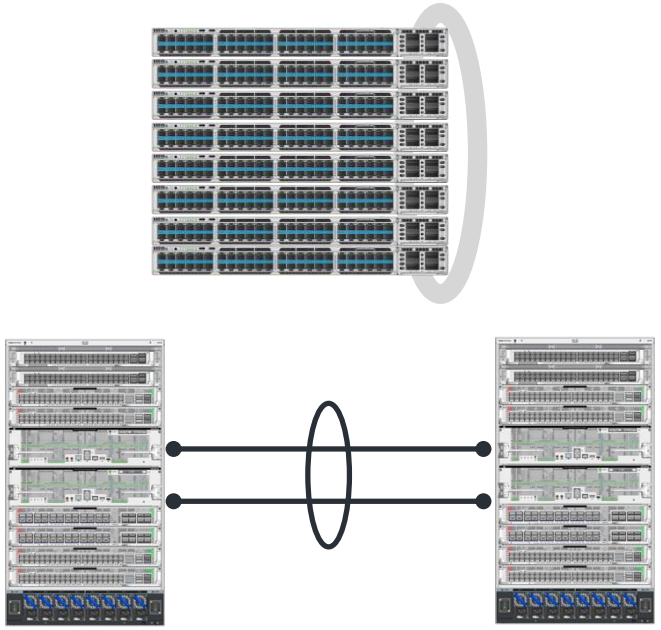

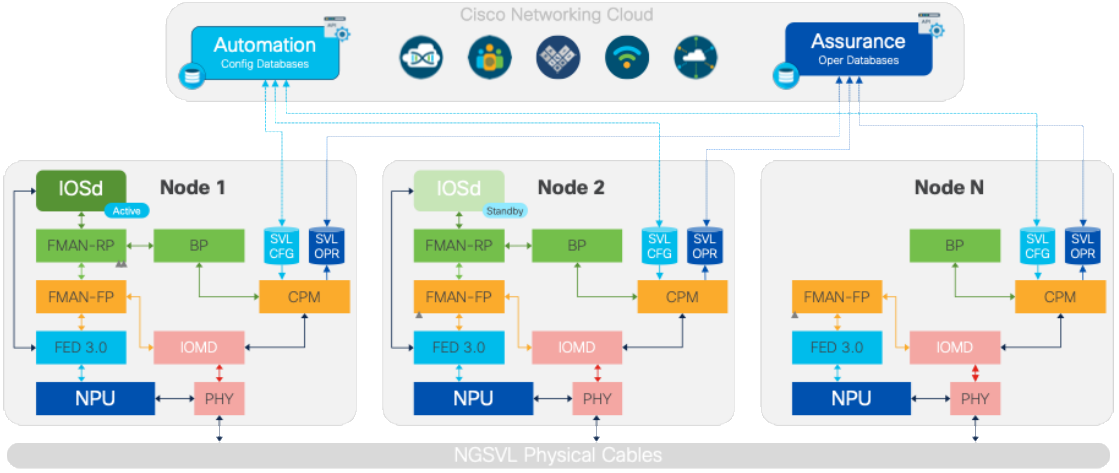

- Standard Ethernet Based (VXLAN)
- Use shortest-path-first (SPF) (ISIS)
- Support stacking with both front panel and back-side connections
- Capable of supporting more than 8 members

Cisco C9350 Series Switches

New stacking architecture for the new generation of campus switching

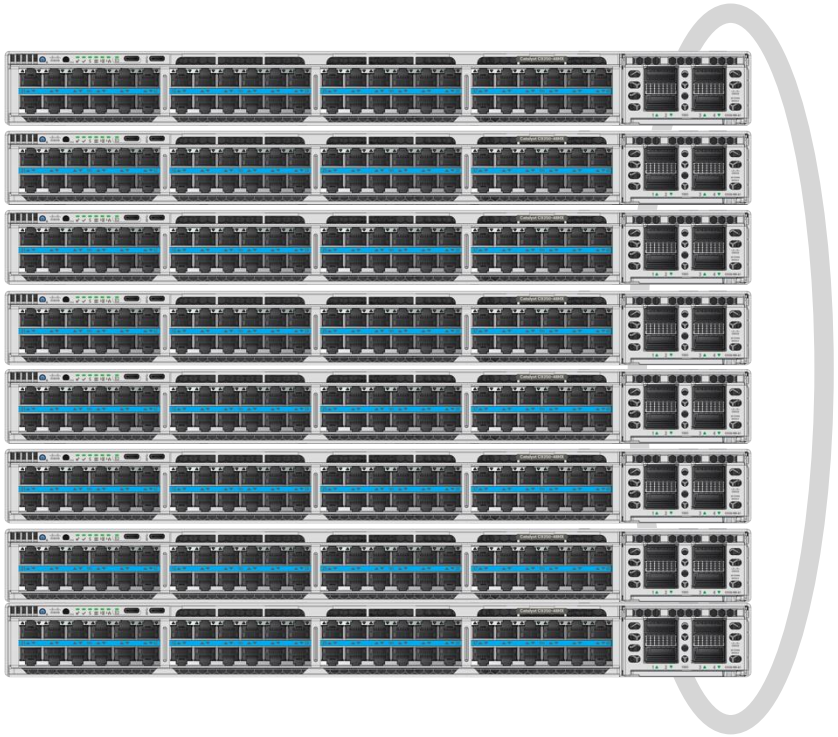

1.6 Tbps stacking bandwidth

User friendly stacking design
Eliminate mandatory long stack cable requirement

Ethernet based stacking

Isolated bandwidth loss in event of failure

Enhanced stackpower
More power draw per cable

Cisco C9350 StackWise Configurations options

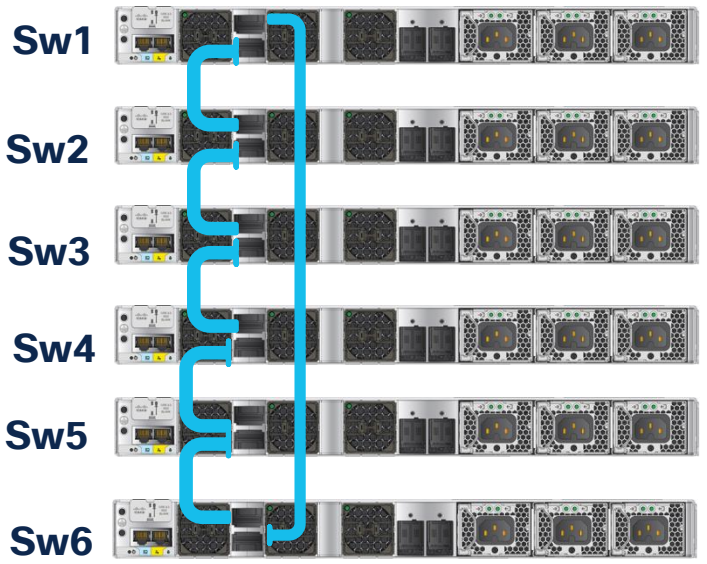

Typical Stack Cables Configurations

Sw1 Port1- Sw2 Port0
Sw2 Port1- Sw3 Port0
Sw3 Port1- Sw4 Port0
Sw4 Port1- Sw5 Port0
Sw5 Port1- Sw6 Port0
Sw6 Port1- Sw1 Port0

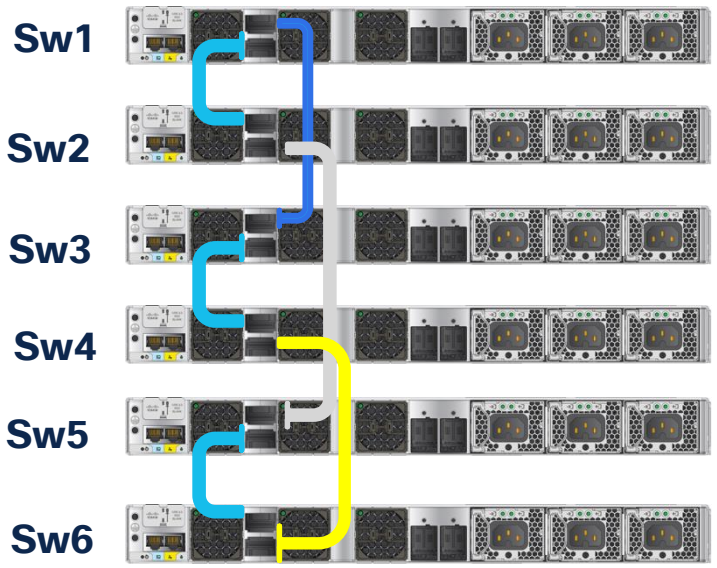

Sw1 Port0- Sw3 Port0
Sw1 Port1- Sw2 Port0
Sw2 Port1- Sw5 Port0
Sw3 Port1- Sw4 Port0
Sw4 Port1- Sw6 Port1
Sw5 Port1- Sw6 Port0

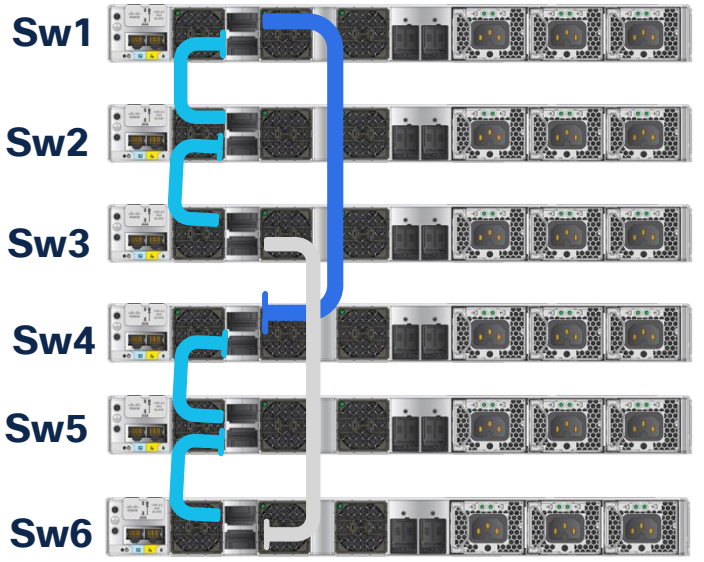

Sw1 Port0- Sw4 Port0
Sw1 Port1- Sw2 Port0
Sw2 Port1- Sw3 Port0
Sw3 Port1- Sw6 Port1
Sw4 Port1- Sw5 Port0
Sw5 Port1- Sw6 Port0

As stack grows, it provides greater flexibility in connecting cables in different configurations with shorter stack cables. No restrictions on how or where you connect to form a full ring.

Considering 6 x C9350 Switches in Stack

Enhanced stacking architecture

Failure handling

Current StackWise architecture

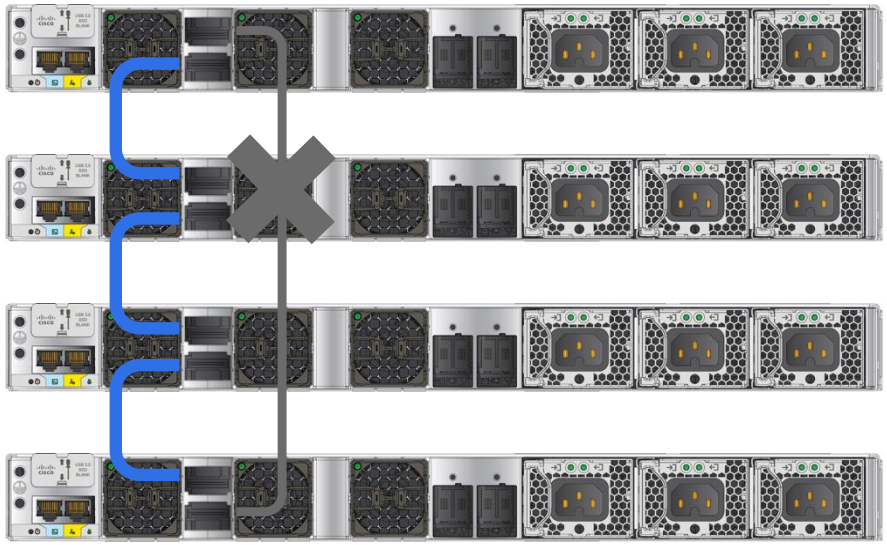

Since packets are fragmented and load-balanced between stack ports, a single failure will effectively halve the bandwidth for the full stack.

Next Generation StackWise architecture

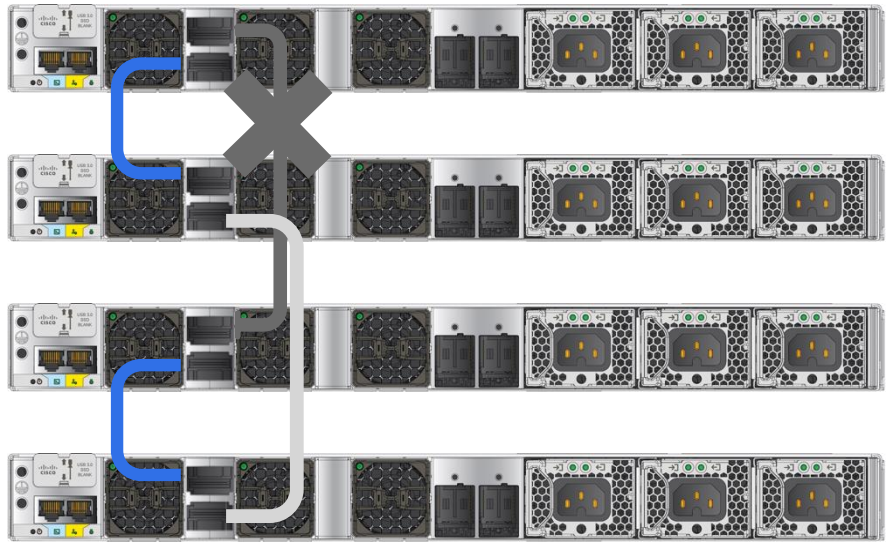

Switch 1 to 3 bandwidth is halved. Other switches operate at full bandwidth.

```
Switch#sh switch stack-bandwidth
                  Stack      Current
Switch#   Role    Bandwidth  State
-------------------------------------
*1        Active  800G       Ready
 2        Standby 1600G      Ready
 3        Member  800G       Ready
 4        Member  1600G      Ready
```

A link failure only affects bandwidth for directly connected switches, all other switches operate at full bandwidth

Cisco C9350 Series – Enhanced Stackpower+

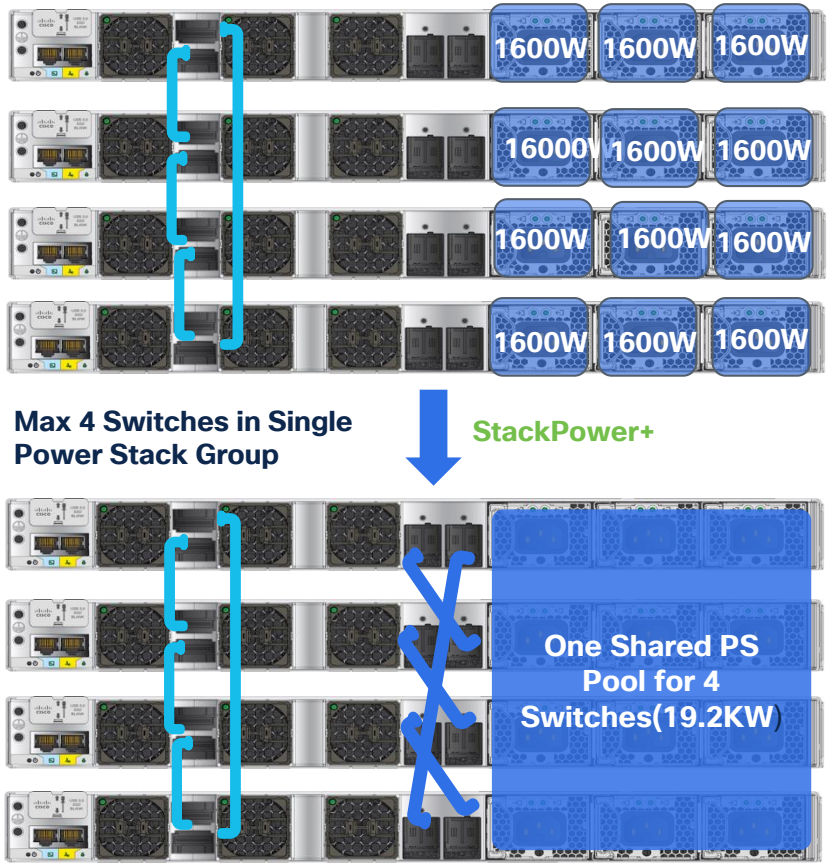

- New Stack Cables Capable of Handling Higher Current (55A)
- Enhanced Stackport Circuitry for reliable power distribution
- Support for up to 4 Switches in a single StackPower Group

Zero-Footprint RPS – Built-in redundancy without extra hardware

Power Pooling – Shared power across switches for efficient use

1+N Redundancy – Inline backup for higher reliability and uptime

Priority-Based Power Allocation – with unused units on standby

StackPower+ delivers over 30% more power via StackPower ports

Hyper Stack – StackWise Virtual

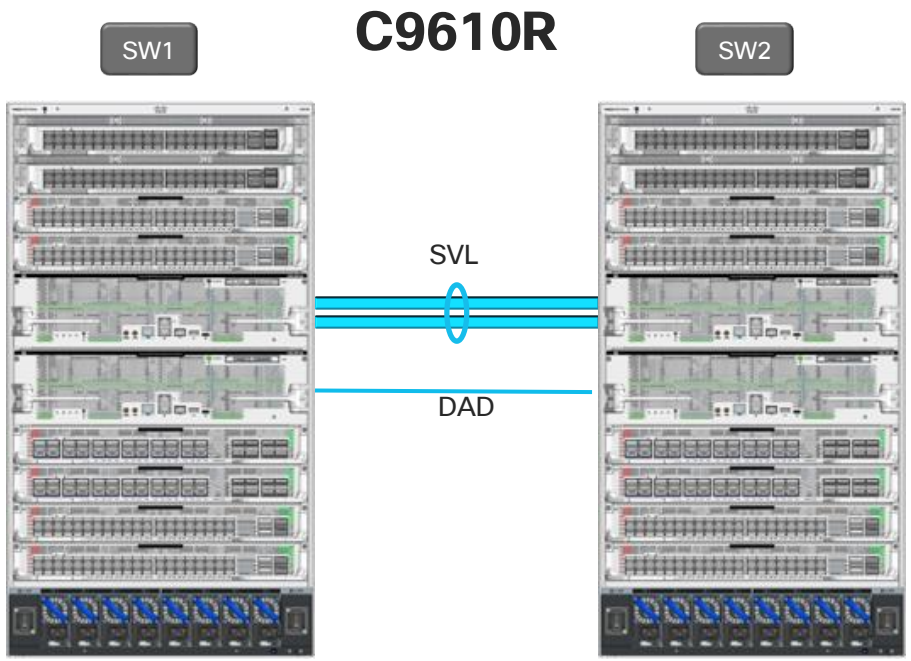

Standard Ethernet Based (VXLAN)

- Use shortest-path-first (SFP – ISIS)
- Dynamically adding and removing
  - SVL members
  - DAD links Dynamic adding and removing

- Simplify Operations by Eliminating STP, FHRP and Multiple Touch-Points
- Double Bandwidth & Reduce Latency with Active-Active Multi-chassis EtherChannel (MEC)
- Minimizes Convergence with Sub-second Stateful and Graceful Recovery (SSO/NSF)

* Limited Availability at GA with single sup per chassis

Application hosting

Cisco C9000 Switches – App hosting

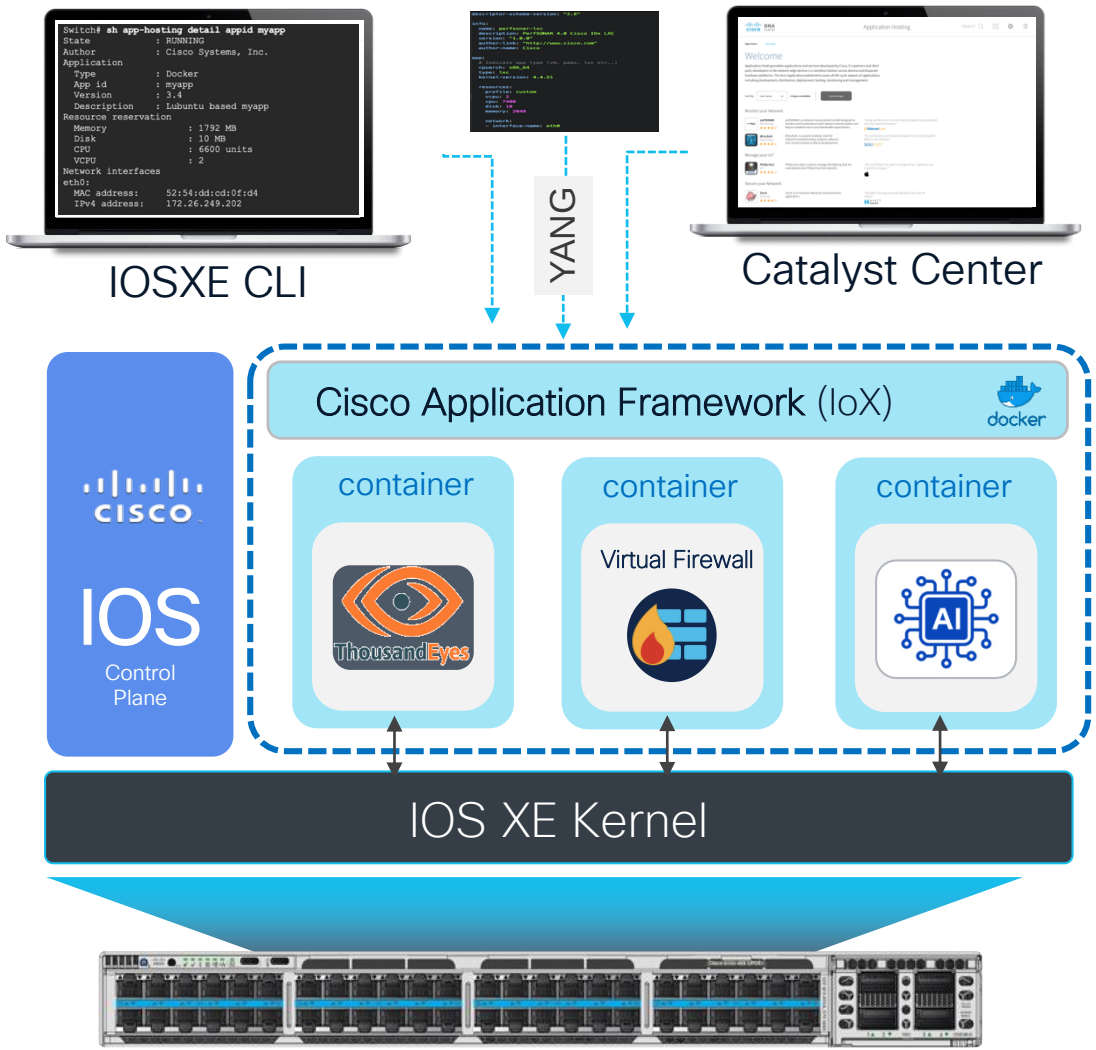

- Cisco-Signed Apps supported on local Flash – no SSD required
- Multiple-App deployment (same switch) for Cisco-Signed Apps

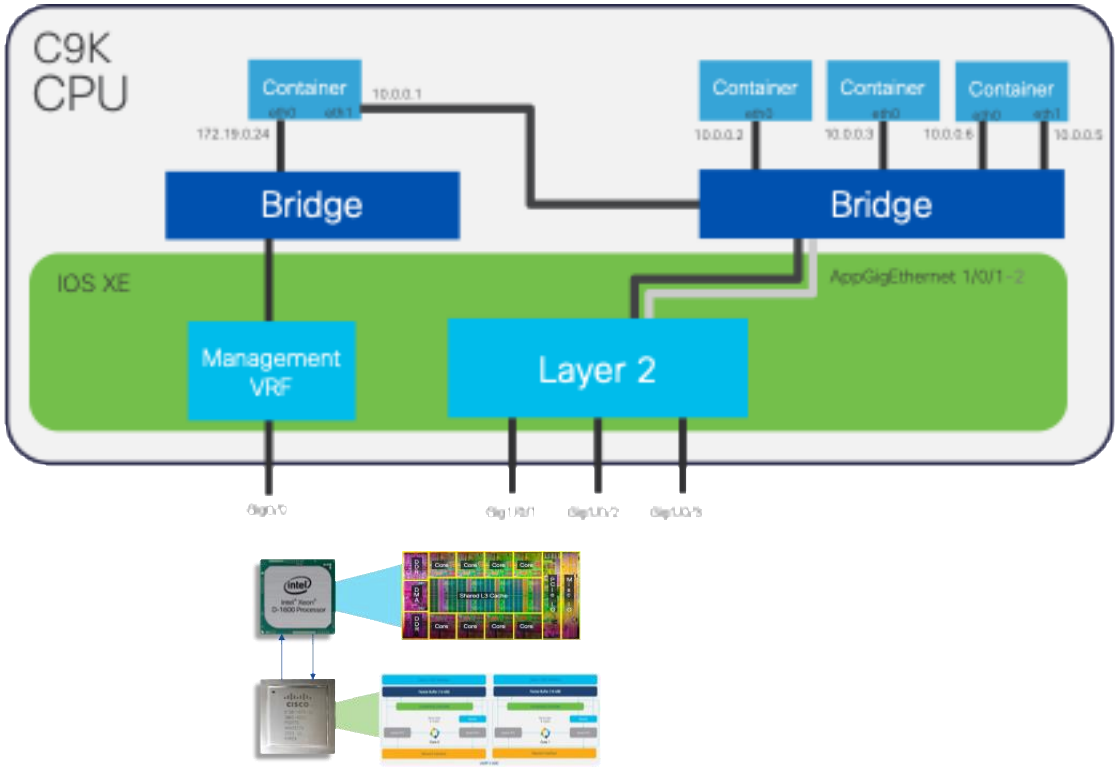

C9350 Enhanced Application-Hosting

Unlocking emerging use cases

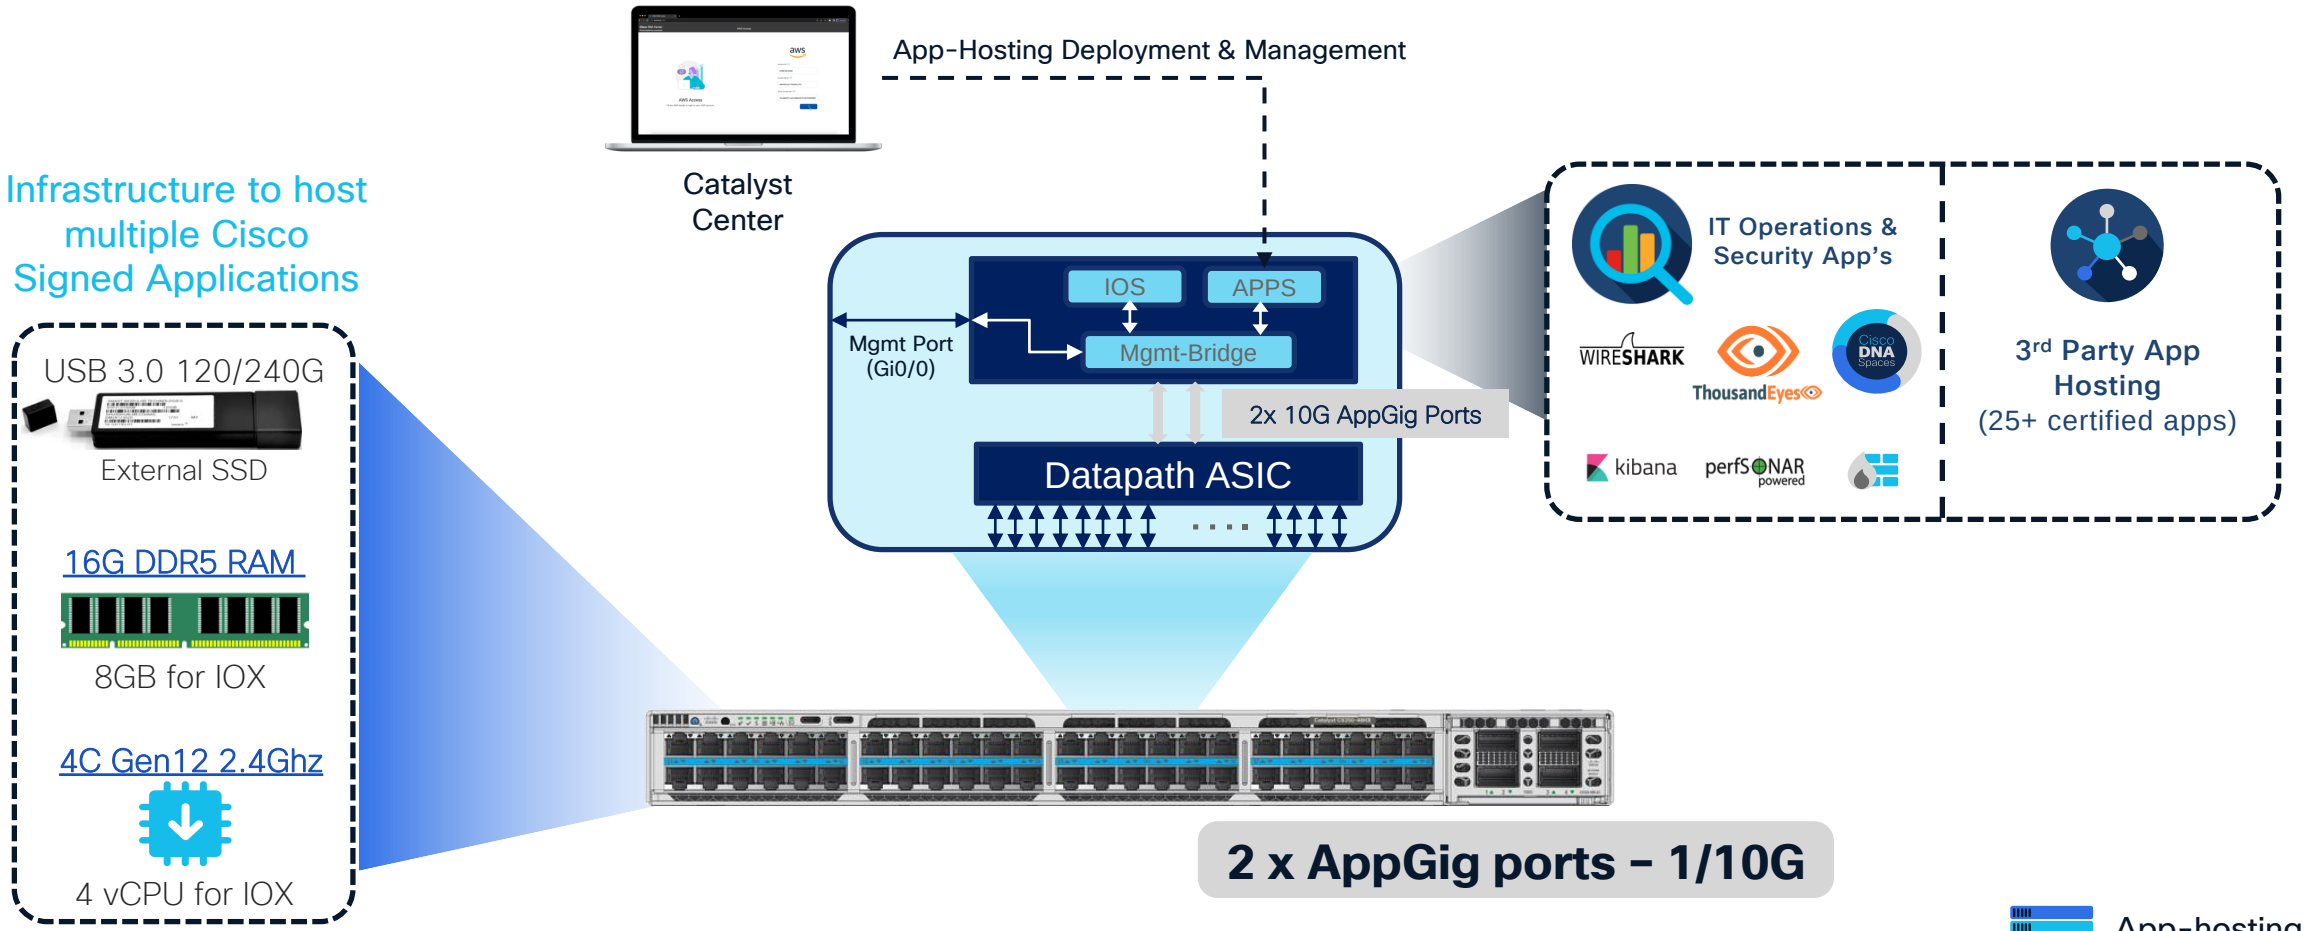

2x App Memory | ≥13% faster[1] DDR5 | ≥35% faster[2] Gen12

Expanding capabilities for the next generation

17.18.2

Enhanced app-hosting infrastructure on Cisco 9610-Sup 3/XL

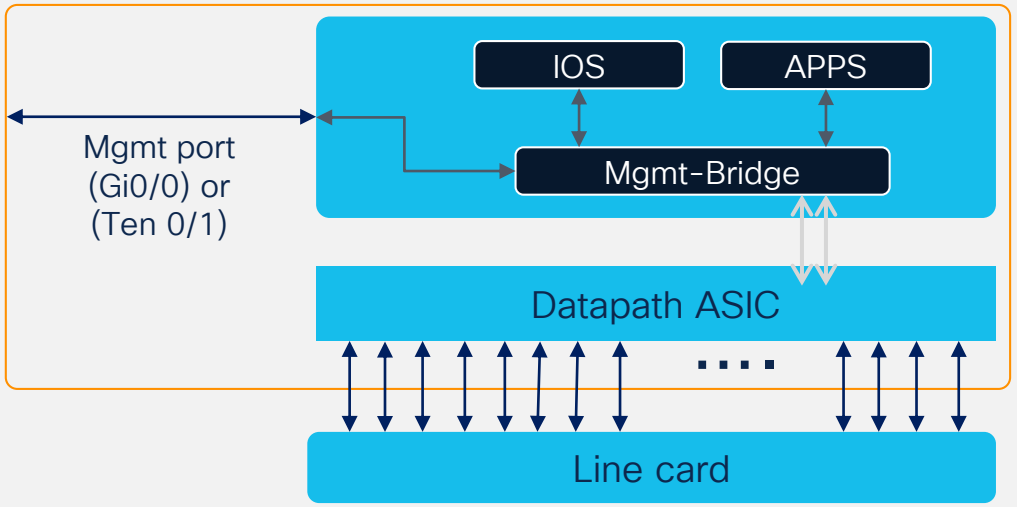

Internal connection
2x 10G AppGigEthernet Ports

NEW

External connection
Management port (gi0/0 or Ten 0/1)

8x Core Intel Gen10 Xeon Series CPU

32G DDR4 RAM

App hosting
4 vCPU + 8G

IOS-XE/Software

Cisco IOS XE – A Modern Operating System

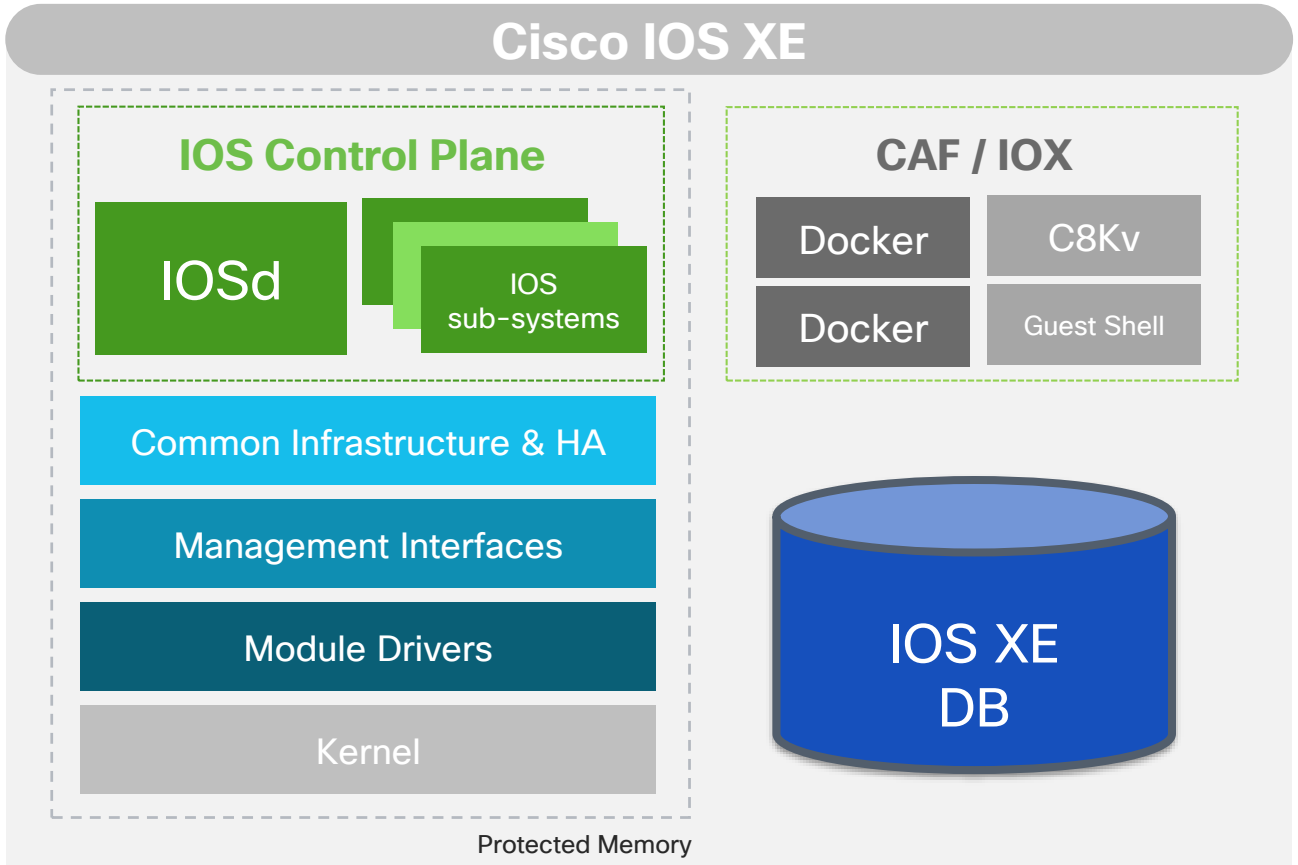

Cisco IOS subsystems
Resiliency and High Availability

Cisco IOS XE database
Programmability and Open models

IOX + Docker containers
Cisco and 3rd-party App hosting

Using a Single, Open, Model Driven & Secure Operating System

Cisco IOS XE Architecture

Modularized Components for Software Abstraction

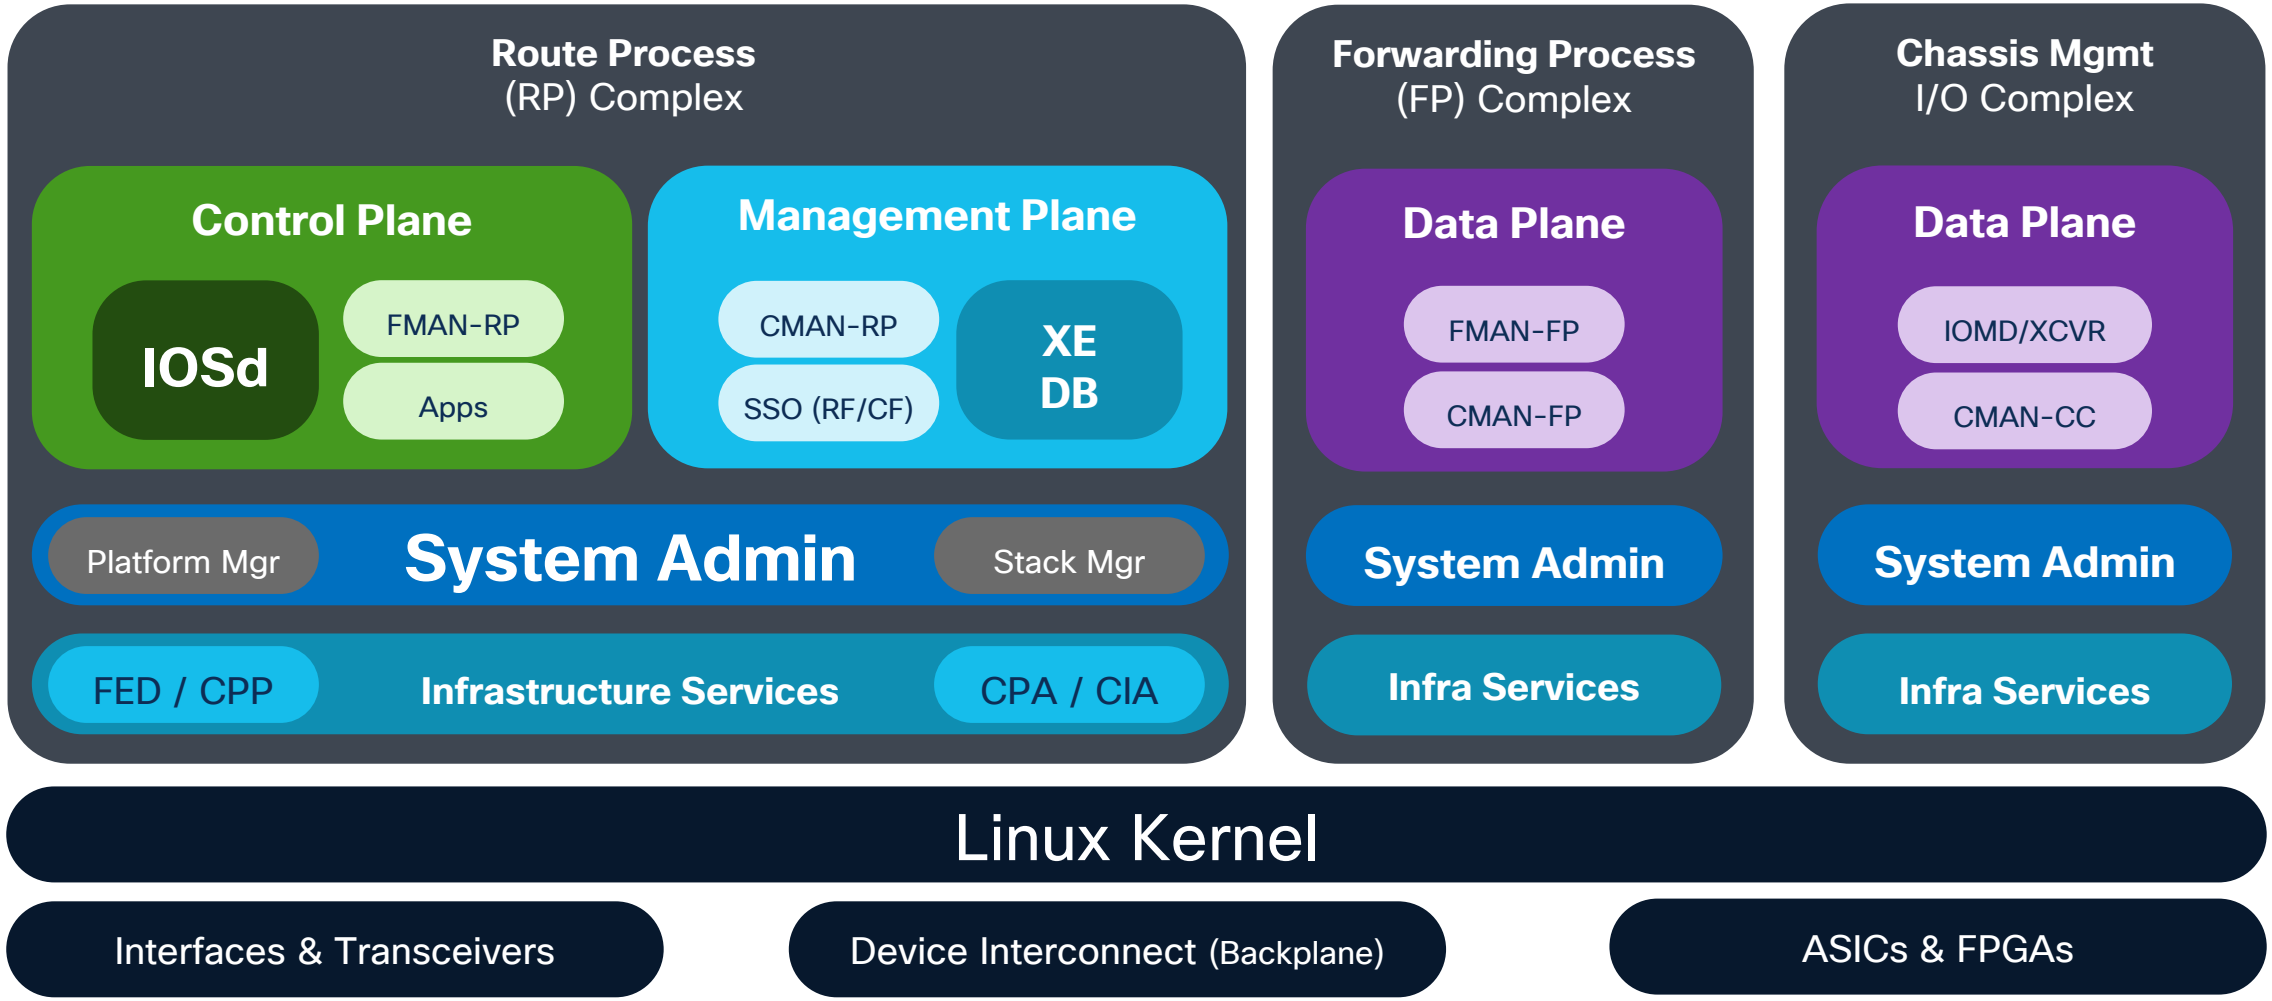

Cisco IOS XE High Availability

Control Plane to Data Plane Programming

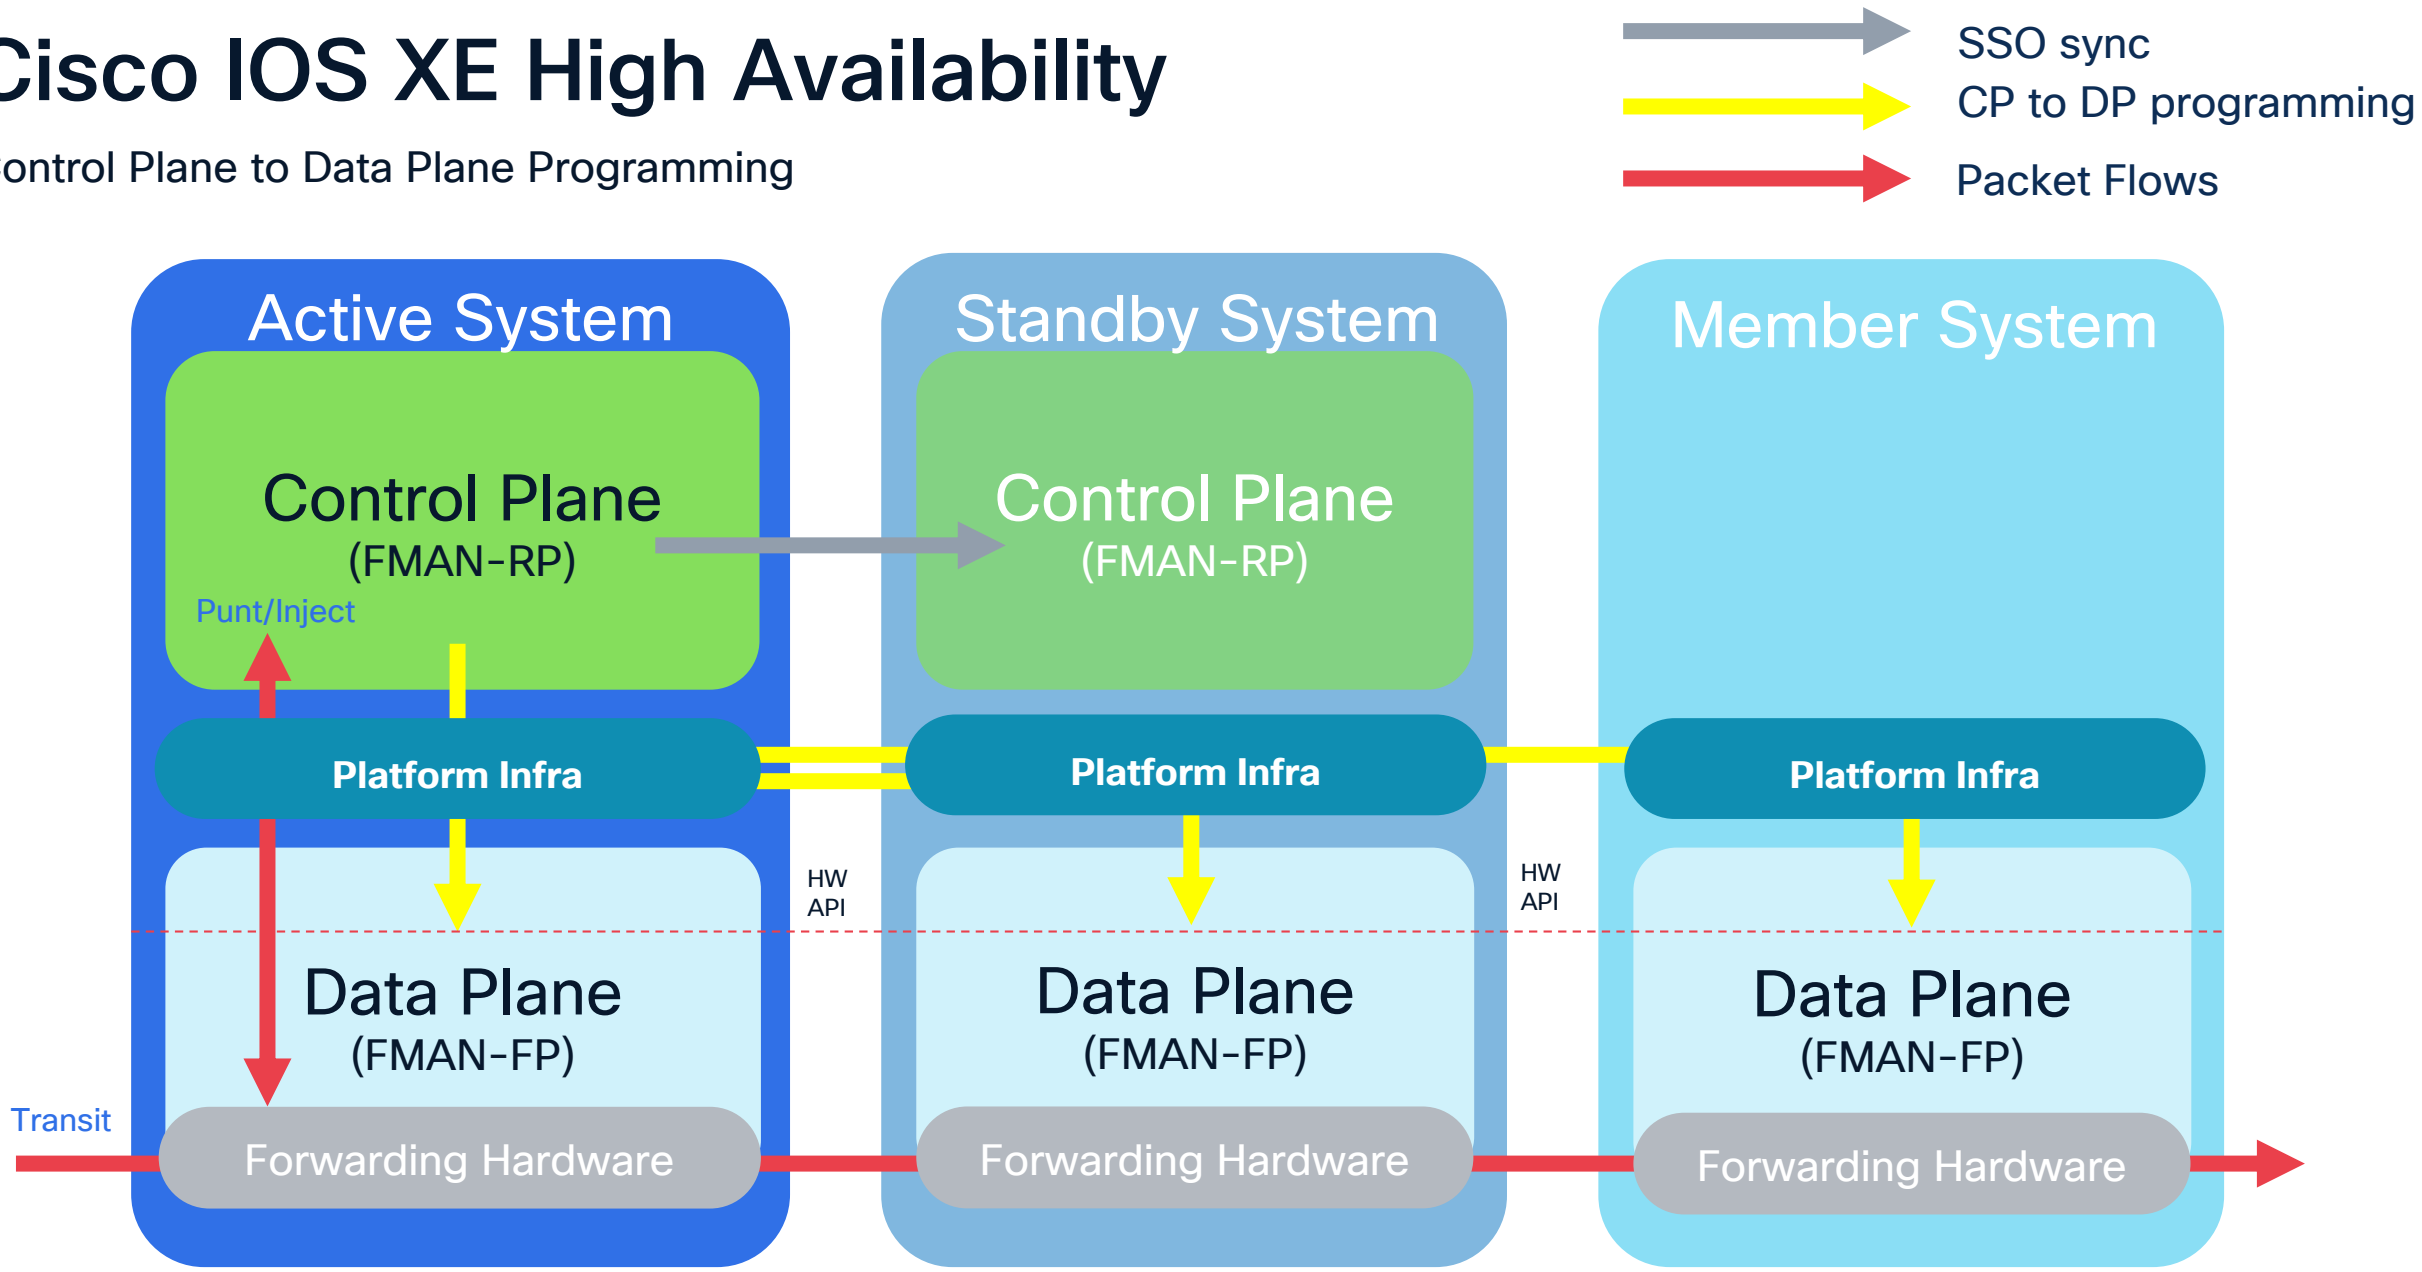

Design Consideration

Three-Tiers Campus Design

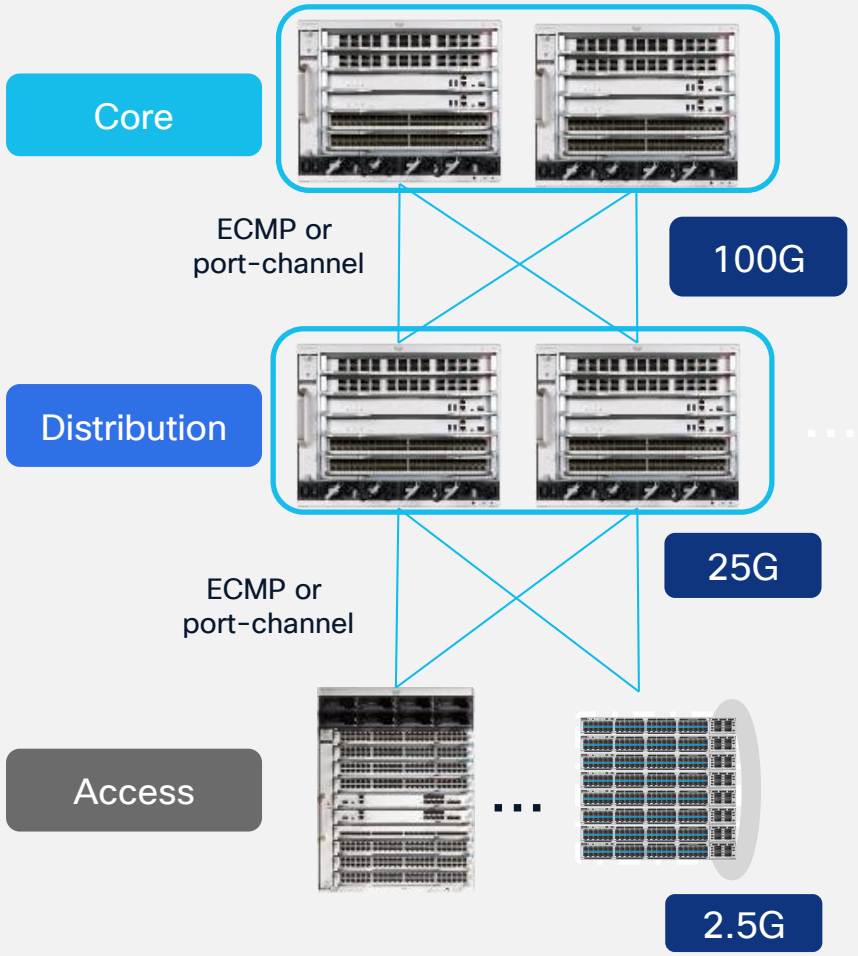

- Catalyst® 9606R Switch with 4x C9600-LC-32C for connection to the distribution layer, WAN, and between core devices
- Soft oversubscription ratio of 4:1 from distribution to core

Core

- Catalyst 9606R with 4x C9606-LC-56YL4C: 100G as the uplink
  connection to the core and SVL links; 10G/25G as the connections to the access layer
- Soft oversubscription ratio of 20:1 from access to distribution

Distribution

- Catalyst 9410 with 1G copper and/or mGig modules with Sup2/2XL
- 8x Catalyst 9300 StackWise-480 with 10G/25G uplinks

Access

384 ports x **112** Access Pods x **6** Distribution Blocks = **258K** ports of 2.5G

Three-Tiers Campus Design – C9610R and C9350

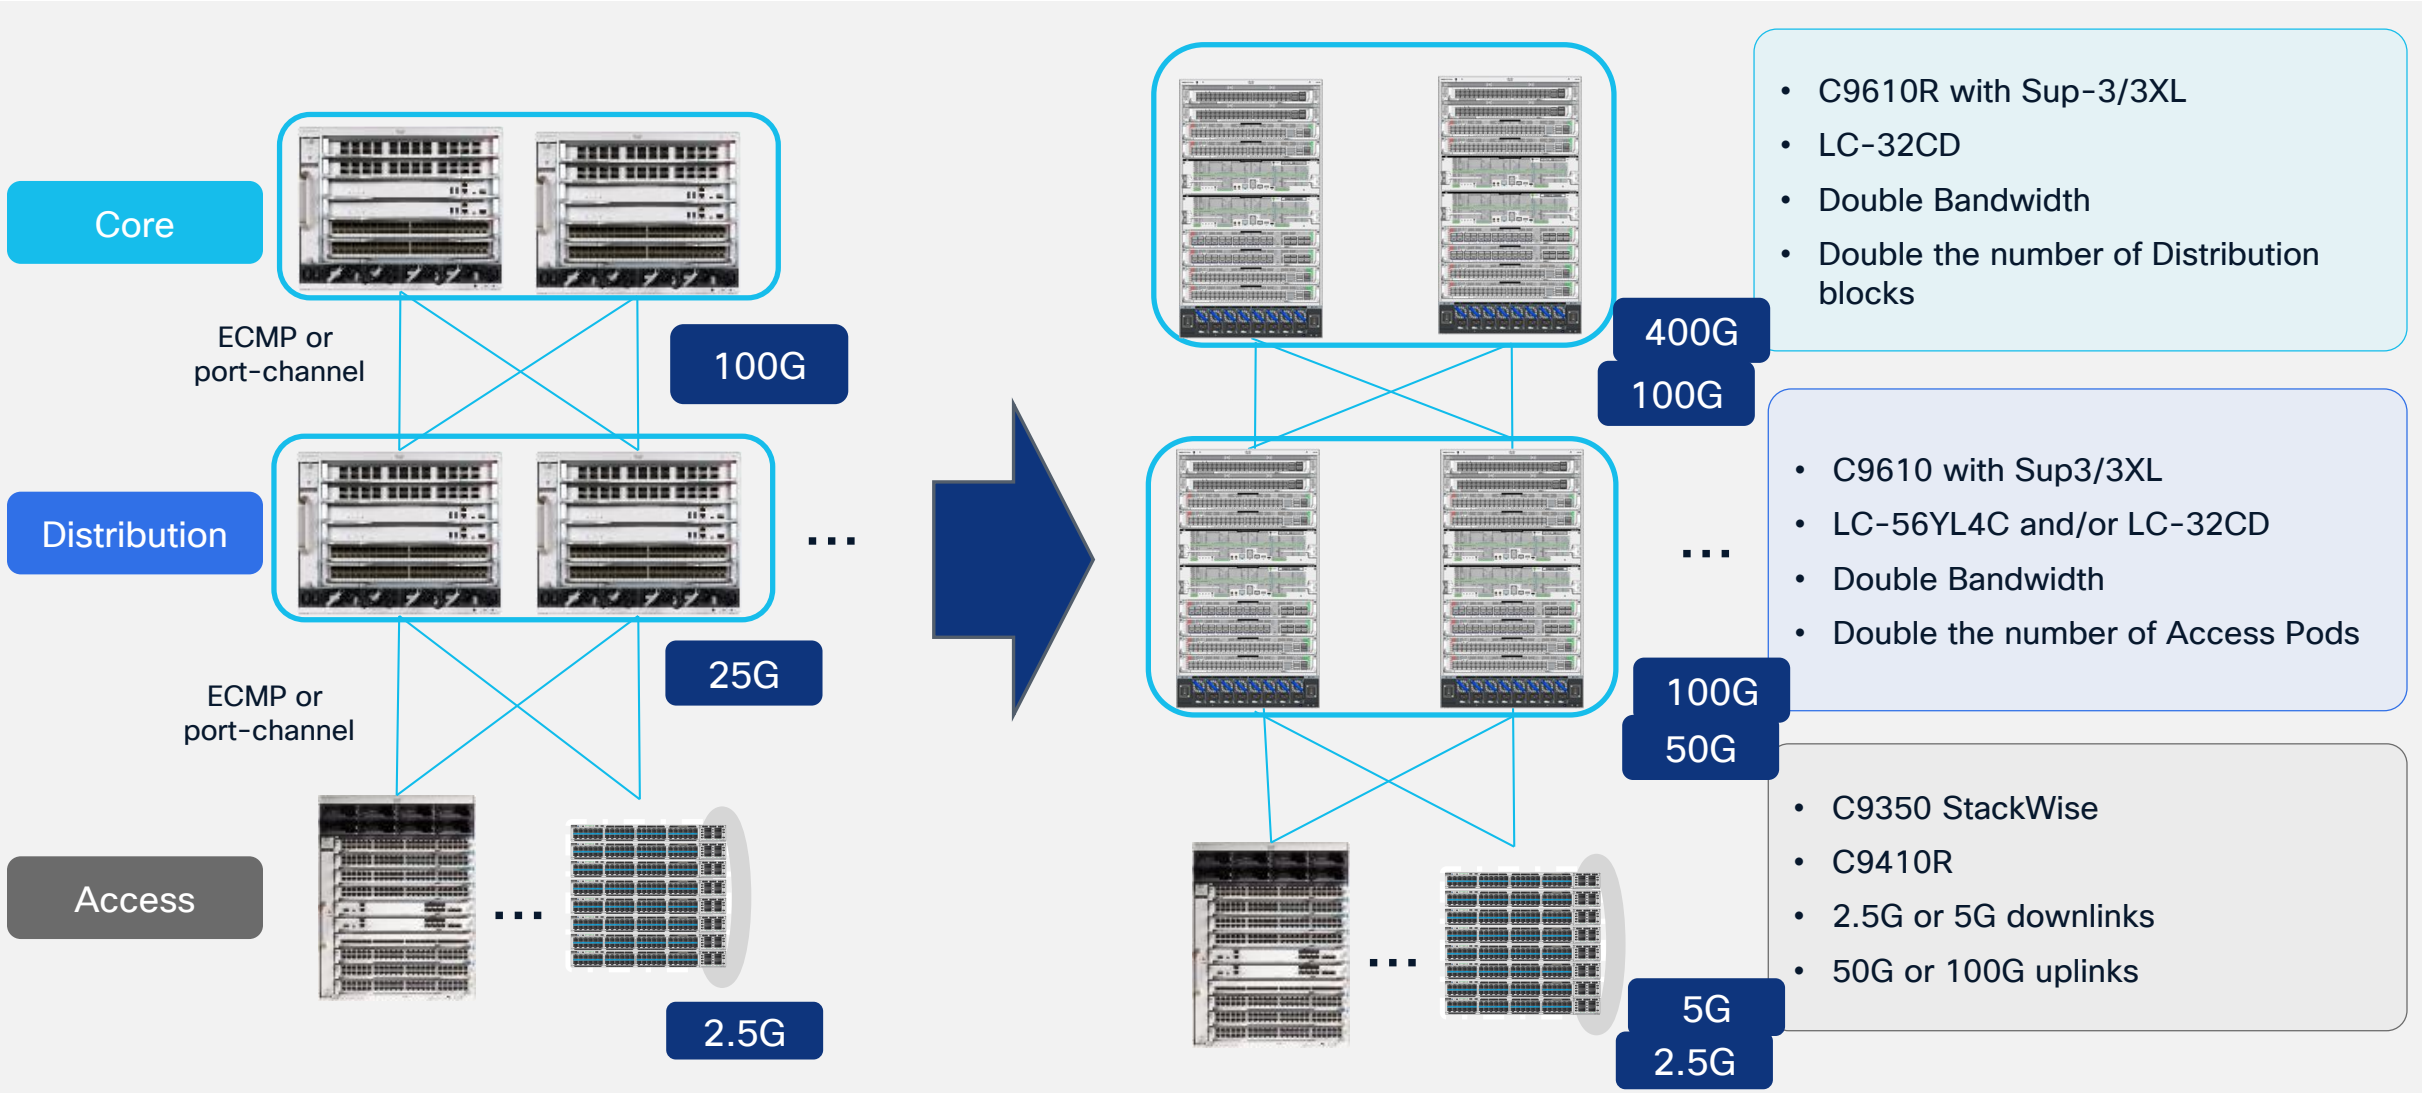

Unified Experience

Unified Experience: The New out-of-box Experience

Effortless Onboarding: Catalyst or Meraki mode!

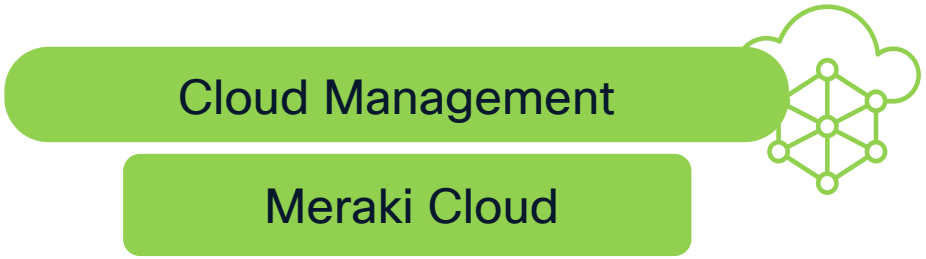

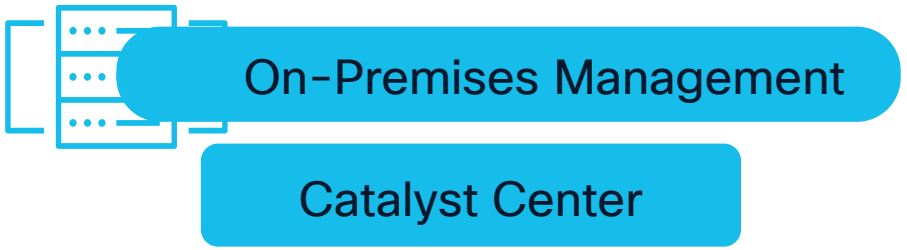

Single Unified Product

One box, your choice
No CLI conversion or -M SKUs needed!

Migration

Unified Onboarding Experience

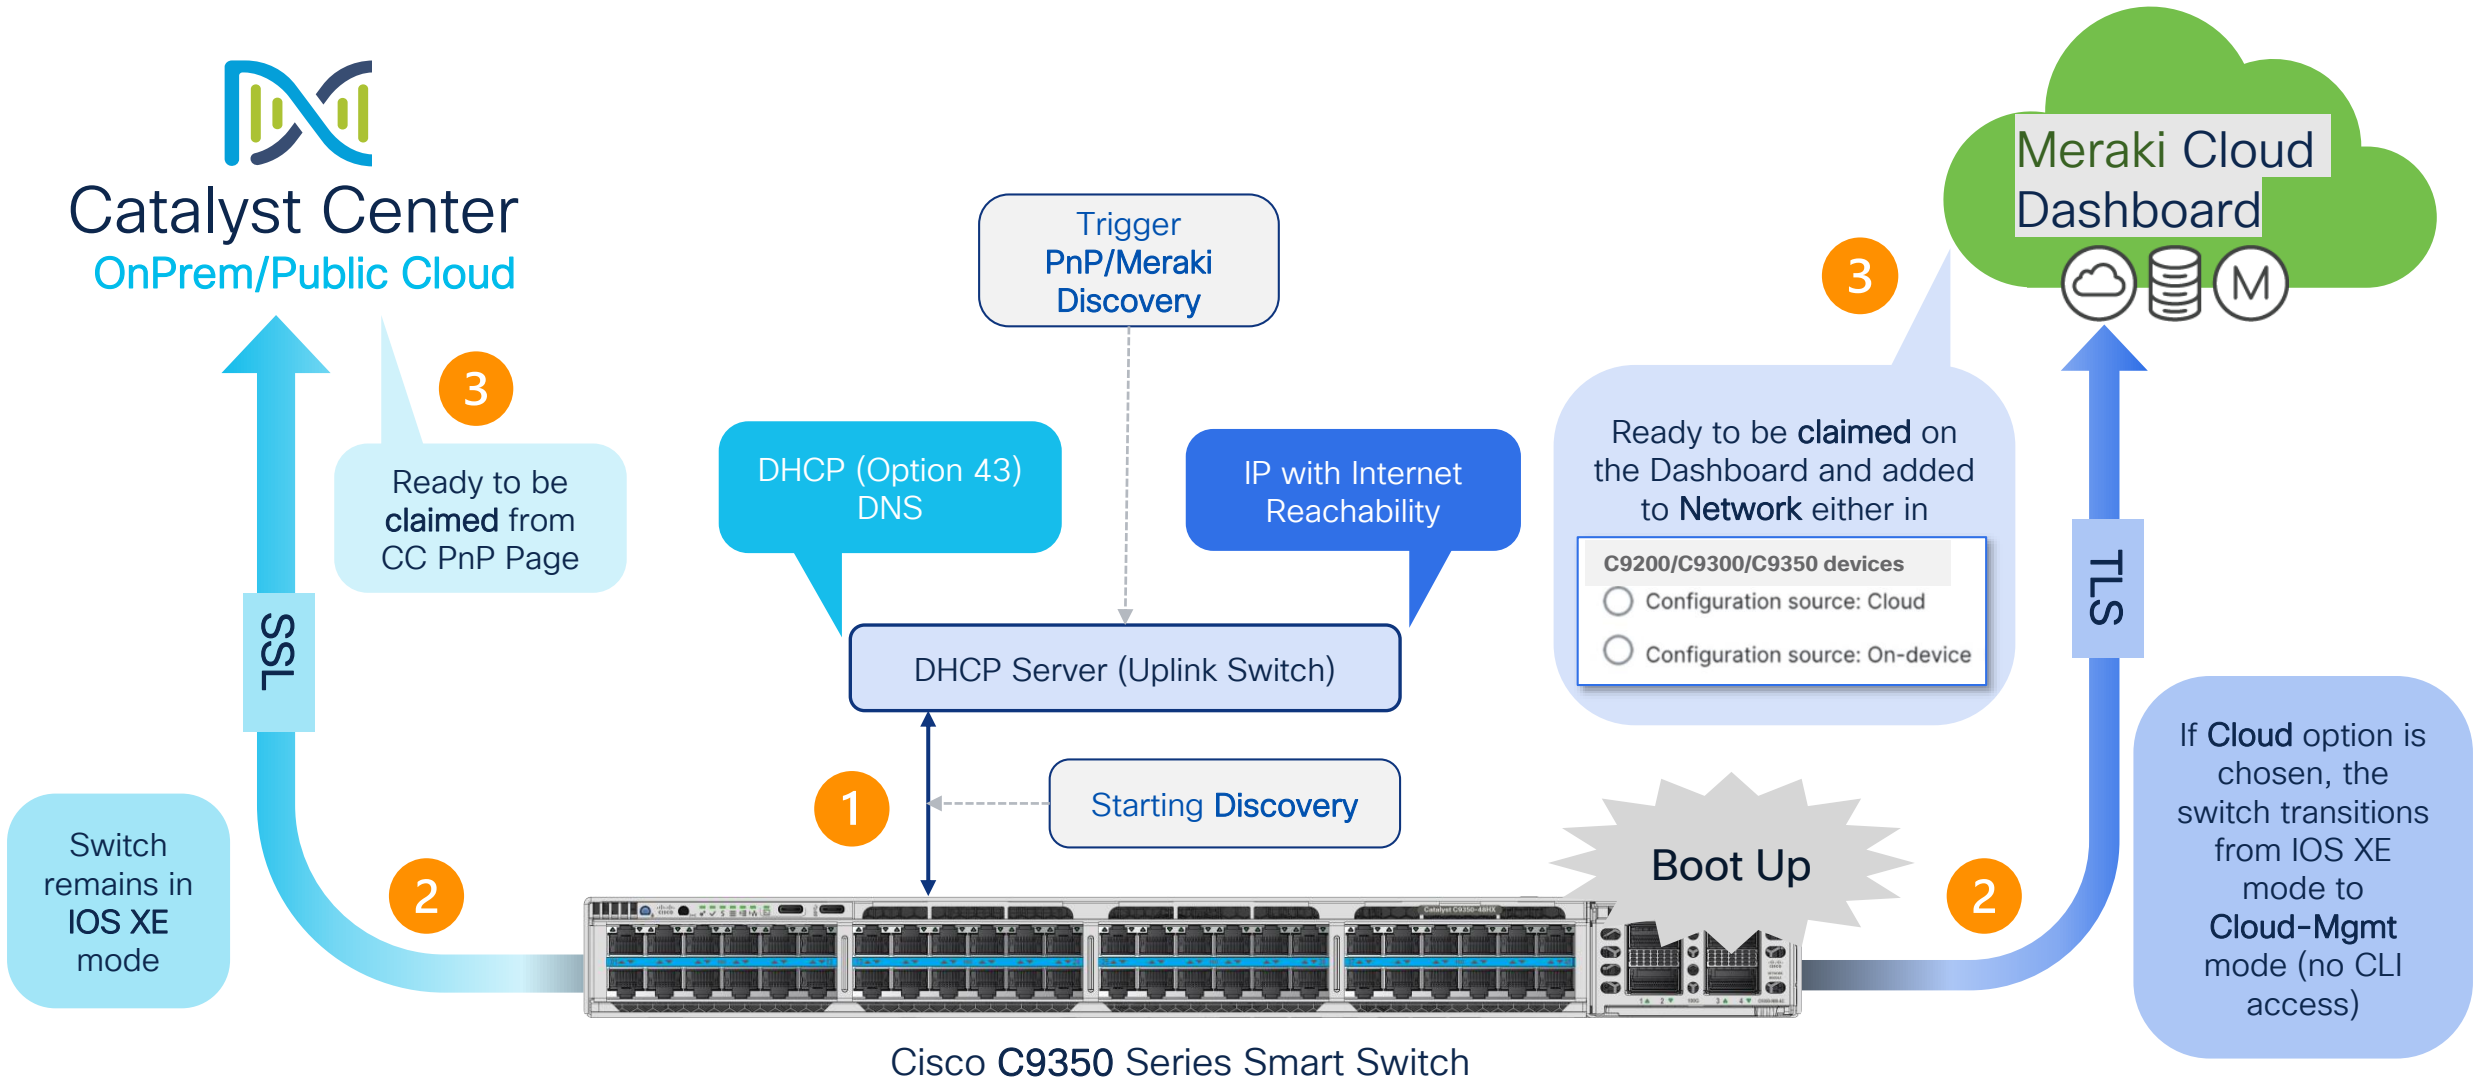

Summary

Choose the right Access Switch for legacy Migration

Low to medium scale

Cost focus – 9300/L (Gen1)

- 8K IPv4 or 4K IPv6 routes
- Best performance for price
- 32,000 MACs
- 32,000 /32 host routes
- 4,000 ACE entries
- Mature & Feature optimized

Medium to high scale

Feature focus – 9300X (Gen2)

- 16K IPv4 or 8K IPv6 routes
- Balances scale & performance
- 32,000 MACs
- 32,000 /32 host routes
- 8,000 ACE entries
- Mature & Feature optimized

Very high scale

Scale focus – 9350 (Gen3)

- 256K IPv4 or 128K IPv6 routes
- Max 10G down + 90W UPOE
- 64,000 MACs
- 96,000 /32 host routes
- 32,000 ACE entries
- HW-ready for all C3K features

Choose the right Core Switch for legacy Migration

Low to medium scale

Feature focus - C9600-SUP1

C9600-SUP-1
& GEN1 LCs

- 256K IPv4/IPv6 routes
- Links speeds up to 100Gs
- 128,000 MACs
- 128,000 /32 host routes
- 54,000 ACE entries
- Mature & Feature optimized

Medium to high scale

ACL/FNF focus - C9610-SUP3

- 1M IPv4 or 512K IPv6 routes
- Links speeds up to 400Gs
- 128,000 MACs
- 128,000 /32 host routes
- 64,000 ACE entries

IP route focus - C9600-SUP2

- 2M IPv4 or 1M IPv6 routes
- Links speeds up to 400Gs
- 128,000 MACs
- 128000 /32 host routes
- 8,000 ACE entries

Very high scale

Scale focus - C9610-SUP3XL

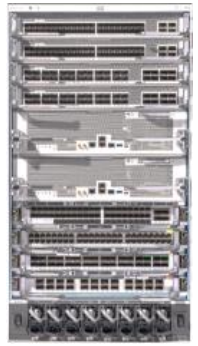

9610-SUP3–3XL
& GEN2 LCs

- 2M IPv4 or 1M IPv6 routes
- Links speeds up to 400Gs
- 256,000 MACs
- 256,000 /32 host routes
- 256,000 ACE entries
- HW-ready for all C6K features

Complete your session evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to win 1 of 5 full conference passes to Cisco Live 2026.

Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.

Level up and earn exclusive prizes!

Complete your surveys in the Cisco Live mobile app.

Continue your education

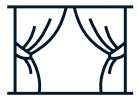

Visit the Cisco Showcase for related demos

Book your one-on-one Meet the Engineer meeting

Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs

Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

Thank you

cisco Live !


CISCO