

# Cisco Silicon One and IOS XE Architecture and Innovation

**cisco** Live !

Cisco Catalyst 9000 Series

Shawn Wargo  
Principal TME

Session ID: BRKARC-2092

# Cisco Webex App

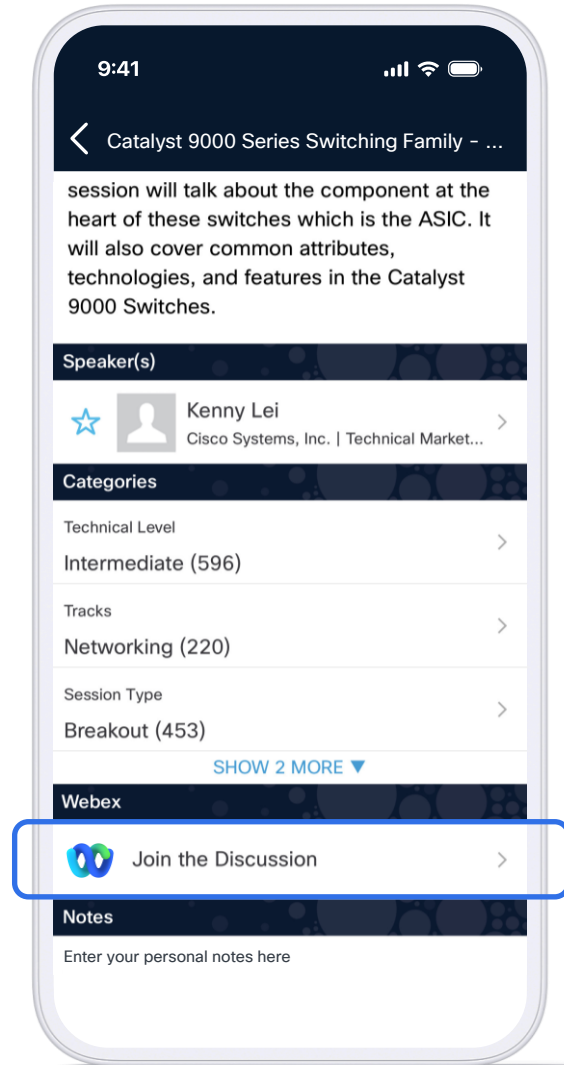
## Questions?

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 14 November 2025.**



<https://cislive.ciscoevents.com/cislivebot/#BRKARC-2092>

# Who am I?

I'm a **Principal Engineer of Technical Marketing** (Principal TME) for Cisco Enterprise 'Network Experience' (NX) Product Management team. I've been with Cisco **since 1999**.

I mainly focus on **Enterprise Switching & Routing** technology areas, with a special emphasis on 'next generation' **Hardware & Software** products and solutions.

As a Principal TME - I'm currently working on the next generation of **Cisco Switching, Wireless & Routing** products, and solutions like Software-Defined Access (SDA) & Cisco UNX.

**#HighBitRate**

**Shawn Wargo**  
Principal TME

swargo@cisco.com  
@shawn\_wargo



# Why BRKARC-2092?

Learn about ASICs & IOSXE in one session



The slide features the Cisco logo and tagline 'The bridge to possible' in the top left. The main title is 'Cisco IOS XE Software Architecture & Innovations' with the subtitle 'Catalyst 9000 Series'. The speaker information is 'Shawn Wargo - Principal TME @Shawn\_Wargo BRKARC-2090'. A diagram of the Cisco IOS XE architecture is shown, including components like IOS Control Plane, CAF / IOS, Common Infrastructure & HA, Management Interfaces, Module Drivers, Kernel, and IOS XE DB. The Cisco Live! logo and #CiscoLive hashtag are at the bottom.

BRKARC-2090



The slide features the Cisco logo and tagline 'The bridge to possible' in the top left. The main title is 'Cisco UADP & Silicon One ASIC Architecture & Innovations' with the subtitle 'Catalyst 9000 Series'. The speaker information is 'Shawn Wargo - Principal TME @Shawn\_Wargo BRKARC-2091'. An image of a silicon chip is shown in the bottom right. The Cisco Live! logo and #CiscoLive hashtag are at the bottom.

BRKARC-2091

# Cisco ASICs - What this session is NOT about

This session is **NOT** a detailed ASIC “deep-dive”!

- **Level 2** - Intermediate
- **Limited Time** (only 45 minutes)

This session also does not go into detail about the (many) IOS XE features

Goal is to *introduce & familiarize* you – so you want to *learn more* 😊

## Other Related Sessions:

- **BRKARC-1011** – Cisco Silicon One: Innovation at Speed and Scale
- **BRKARC-2093** – Cisco Innovations in Silicon & Software
- **BRKARC-2098** – Catalyst 9000 Series Switching Family - Access
- **BRKARC-2099** – Catalyst 9000 Series Switching Family - Core and Distribution
- **BRKARC-2668** – Campus Switching Architecture for Future Proofed Workspaces

# Cisco IOS XE - What this session is NOT about

This session is **NOT** a detailed IOS XE “deep-dive”!

- **Level 2** - Intermediate
- **Limited Time** (only 45 minutes)

This session also does not go into detail about the (many) IOS XE features

Goal is to *introduce & familiarize* you – so you want to *learn more* 😊

## Other Related Sessions:

- **BRKENS-2004** - Catalyst 9000 IOS XE Innovations
- **BRKARC-2098** - Catalyst 9000 Series Switching Family - Access
- **BRKARC-2099** - Catalyst 9000 Series Switching Family - Core and Distribution
- **DEVNET-1110** - Modern approaches for IOS XE network device management on Cat9k

# Cisco 9000 Series – Common Building Blocks



## Multi-Core x86 CPU

Application Hosting  
Secure Containers



## Cisco IOS XE<sup>®</sup>

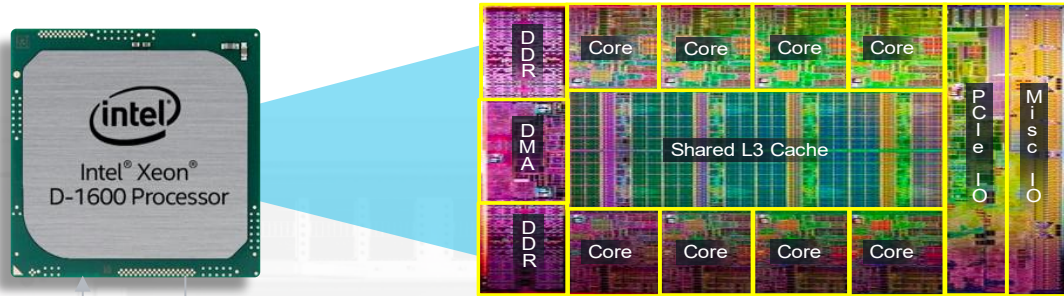
Model-Driven APIs  
Modular Patching



## Cisco UADP & Silicon One™

Programmable Pipeline  
Flexible Tables

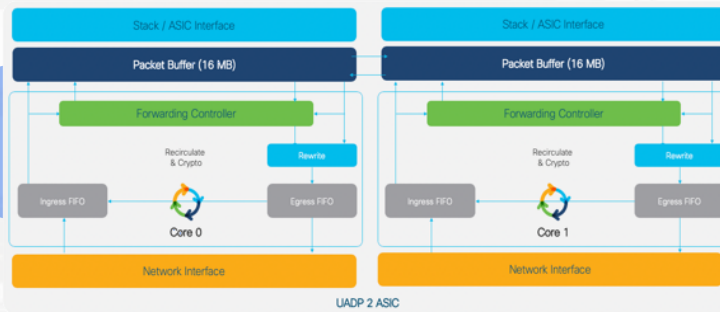
# Software vs. Hardware



## CPU/DRAM

Where the OS “software” runs. Includes control-plane, data-plane and system-management functions.


- **OS layer** – IOSXE (IOSd) and Features, etc.
- **System layer** – FMAN, CMAN, IOMD, FED, etc.



## ASIC(s)

Where the “hardware” processing of traffic & services runs. Uses forwarding and state tables programmed by the software.

- **Forwarding** – L2, L3, ECMP, Encap, etc.
- **Services** – ACLs, QoS, Analytics, Encryption, etc.



Long Reach Host Interface	
25GbE/50GbE/100GbE/200GbE/400GbE	
PCS + FEC	
2:1 Mux	
MAC+256 bit MACsec + PTP	
25GbE/50GbE/100GbE/200GbE/400GbE	
PCS + FEC	
MDIO	SyncE
JTAG	LED
Auto-Negotiation	Coefficient Training
56G-CR/KR Capable SerDes	

## Stub/PHY(s)

Transforms electrical and optical signals, splits or combines signals, and other various “physical” layer functions, such as encryption and timestamping.

# What's New?

CiscoLive 2025 - San Diego



## New Cisco C9000 Smart Switch Series



Campus-Driven **S1 NPU 2.0** ASICs

A100 (C9350) and K100, E100 (C9610-SUP3)



AI-Ready **X86** & **DDR5** CPU & DRAM


Intel RPL I3 (C9350) and Intel ICX-D (C9610-SUP3)




Cloud-Native **IOS XE 3.0** Software

UNIX, Multi-Mode, Programmability, HyperStack, SD-Access, Smart Power

### C9350 Series Smart Switches

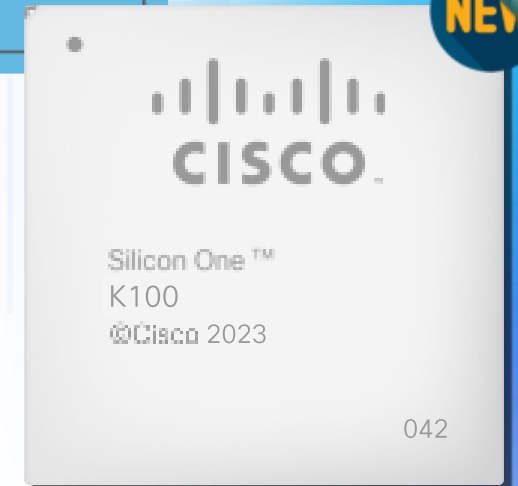



### C9610 Series Smart Switches



# Cisco UADP & Silicon One

ASIC Architecture & Innovations  
Cisco C9000 Series



# Agenda

- 01 Why do we need ASICs
- 02 Flexible ASICs & Cisco UADP
- 03 Cisco Silicon One ASICs
- 04 Cisco 9000 Smart Switches
- 05 A Glimpse into the Future

# Catalyst 9000 Series – Common Building Blocks



## Multi-Core x86

Application Hosting  
Secure Containers



## Cisco IOS XE<sup>®</sup>

Model-Driven APIs  
Modular Patching




## Cisco UADP & Silicon One<sup>™</sup>

Programmable Pipeline  
Flexible Tables

# Custom ASICs – Programmable Silicon



## Cisco Unified Access Data-Plane (UADP)



## Cisco Silicon One™



**Flexible Pipelines**  
Investment Protection



**Adaptable Tables**  
Universal Deployment




**Scalable Resources**  
Enhanced Scale and Buffering


Flexible & Programmable ASICs – Adapt to New Technologies

# Why do we need ASICs

- 1 Why ASICs?
- 2 Cisco UADP
- 3 Cisco S1
- 4 Cisco 9000 Series
- 5 ASIC Futures



# What is an ASIC?



Application Specific Integrated Circuit (ASIC) is a silicon microchip designed for a specific task ...

... rather than '*general-purpose*' processing in a **CPU**.

# Why not use CPUs?


A 'general-purpose' CPU may be fast at running random-access applications, on a laptop or server, but **processing and forwarding 'network traffic'** is a different matter.

Network traffic requires **constant searching of large memory** tables (e.g. L2 tables for MAC addresses, L3 tables for IP routes, L4 ACLs for Security and QoS, etc.)

In a CPU - there are **limited data paths** and tables are held in **off-chip memories** (e.g. DRAM) that can incur significant performance penalties for frequent access.

Remember, this is **Millions - Billions** of packets per second

**CPUs** are **Flexible** but **Slow**



# Multi-Core ASIC Design

Combining multiple Processors in same ASIC

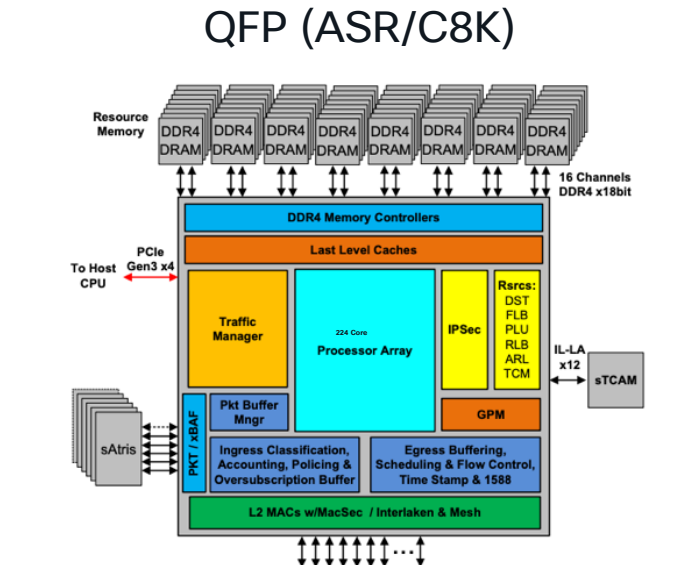
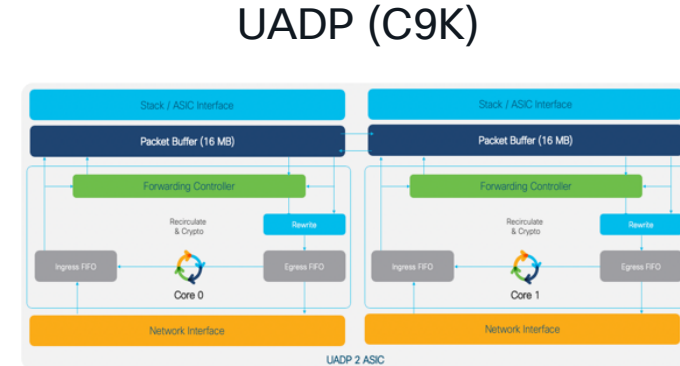
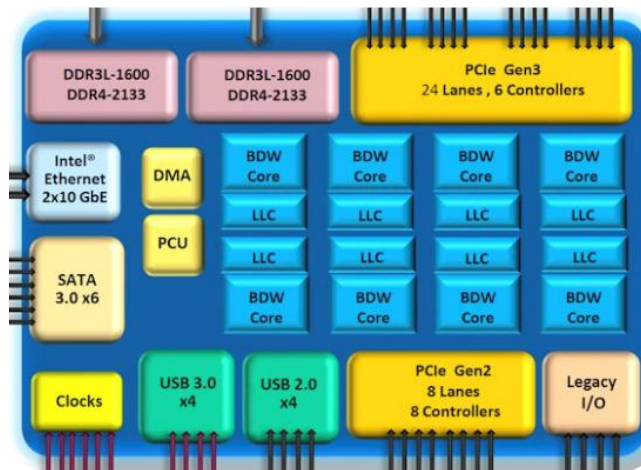

Multiple-Core processors have been used in personal computers since the mid 1990s.

So, while the concept is not new – it's good to briefly revisit.

Multi-Core design addresses the physical limitations of per-processor clock speed (e.g. how effectively can they be cooled), by load-sharing across multiple processors.

However – there are practical (physical) limits to how many cores yield improvements, and cores must share all external

components. Cisco UADP & QFP are examples of using multiple Processor Cores to boost overall ASIC performance



# Why not use FPGAs?


**FPGAs** are **Flexible** but **Expensive**

FPGAs do provide a lot of design flexibility, but they can be **very expensive to develop and support**. They are not built for any specific task and must be reprogrammed for each new task.

FPGAs also **have little or no onboard memory**, requiring other components to provide memory access.

These limits generally relegate FPGAs to a “**special-purpose**” role in most network devices. FPGAs are most often used to augment other ASICs, for the “**one extra feature**” the primary processor does not have.

FPGAs typically cost **far more (10X)** more than an equal ASIC



# Field Programmable Gate Arrays

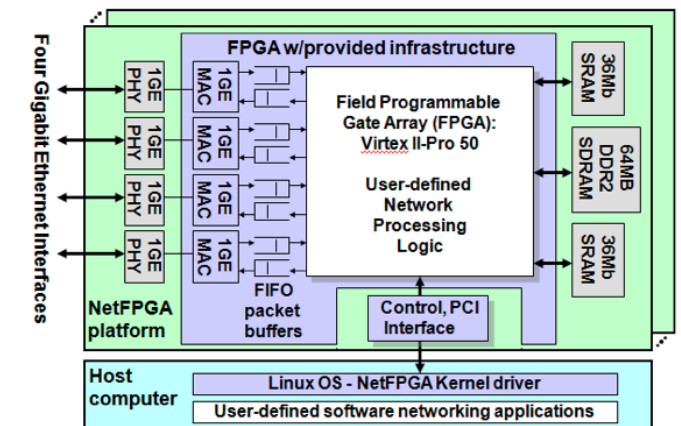

FPGAs can be re-programmed after Manufacturing

As the name suggests, **an FPGA is an integrated circuit** (like an ASIC) with **sets of logic gates and memories** (a gate array), which a **user can re-program** (in the field), after silicon manufacturing, to perform one or more logical operations.

With a fixed ASIC, the microchip is fully baked. It can't be reprogrammed; you get what you get. The software can be deleted or replaced, but the hardware is unchanged.

With an FPGA, there are **no predefined hardware circuits**. The user programs the circuits. The programming can be a single, simple logic gate (e.g. just an AND or OR function), or it can involve multiple complex functions that, together, act as a complete multi-core processor.

You might use an FPGA if you may need to make changes at the chip-level during development (e.g. prototypes), or even to augment another ASIC, to add a new feature.






[intel.com/content/dam/www/programmable/us/en/pdfs/literature/misc/fpgas\\_for\\_dummies\\_ebook.pdf](http://intel.com/content/dam/www/programmable/us/en/pdfs/literature/misc/fpgas_for_dummies_ebook.pdf)

# CPU vs. GPU vs. DPU (& SOC)

## Normalizing Terminology – Data Processors

### CPU vs. GPU vs. DPU

CPU	GPU	DPU
Several cores	Many cores	Dozens of cores
Low latency	High throughput	Higher degree of versatility
Ideal for serial processing	Ideal for parallel processing	Ideal for big data processing
Handles a handful of operations at once	Handles thousands of operations at once	Handles thousands of operations at scale



## Central Processing Unit

A **central processing unit (CPU)** is commonly referred to as the ‘brain’ of a computer. A CPU is essential to all modern computing systems and may use one or multiple processing cores.

A CPU **performs all general-purpose computations centrally**, including low-level Operating System (OS) infrastructure, as well as random Application software programs.

Thus, the system performance is the sum of all software programs running on the CPU.

Major Vendors: Intel & AMD

## Graphics Processing Unit

A **graphics processing unit (GPU)** was originally designed to offload the computer CPU for **video graphics processing**.

A GPU has many smaller, but specialized cores. These cores deliver massive performance by working together and dividing processing tasks across many cores simultaneously (in parallel).

A GPU **excels at highly parallel tasks** like rendering visuals, manipulating data during content creation, and computing results in intensive AI workloads.

Major Vendors: NVIDIA & AMD

## Data Processing Unit

A **data processing unit (DPU)** was originally designed to offload the computer CPU for **network processing (NIC)**. This gave it the earlier term: SmartNIC.

A DPU is focused on moving and processing data within a local network (e.g. data center).

DPU is a new class that combines elements of both CPU & GPU while adding **acceleration engines for unique functionalities** like network path processing, network security, load-balancing and storage networking.

Major Vendors: Marvell, Mellanox

## System on Chip

A system on chip is a combination of all three (**CPU+GPU+DPU**) in a single integrated system.

- 1. Multi-Core CPU:** A standard, high-performance, programmable, multi-core CPU (often based on ARM architecture).
- 2. Programmable Micro-Engines:** To boost performance for apps like AI and Machine Learning, security, physical media, crypto and storage.
- 3. High-Performance NIC:** This ensures efficient buffering and data transfer to GPUs and CPUs at the network’s speed.




# Hardware vs. Software Data-Plane

ASIC-based vs. CPU-based Packet Forwarding


S1/UADP



QFP



X86/ARM



## Hardware-Based (ASIC)

### PROs

- **High Throughput**
- Gbps - Tbps
- High Port Density
- HW-based Services
- Lower COGs

### CONs

- **Less Flexible**
- PI + PD Development
- Longer Dev & Test
- Scale limited by HW
- Limited Services

## Software-Based (CPU)

### PROs

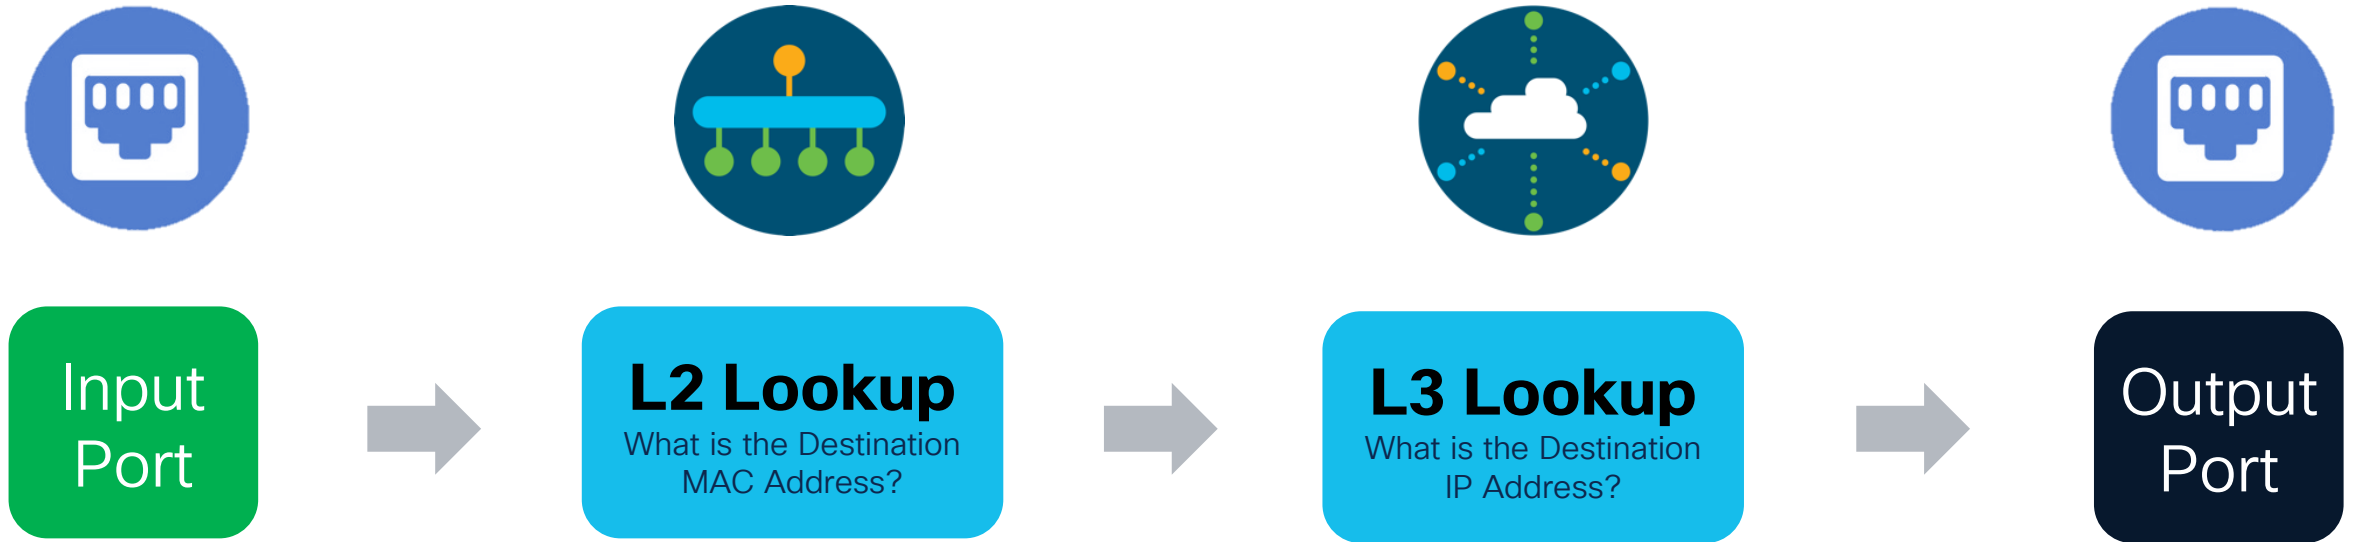
- **More Flexible**
- PI Only Development
- Faster Dev & Test
- Scale limited by CPU
- Lots of Services

### CONs

- **Low Throughput**
- Mbps - Gbps
- Low Port Density
- SW-based Services
- Higher COGs

# What does an ASIC do?

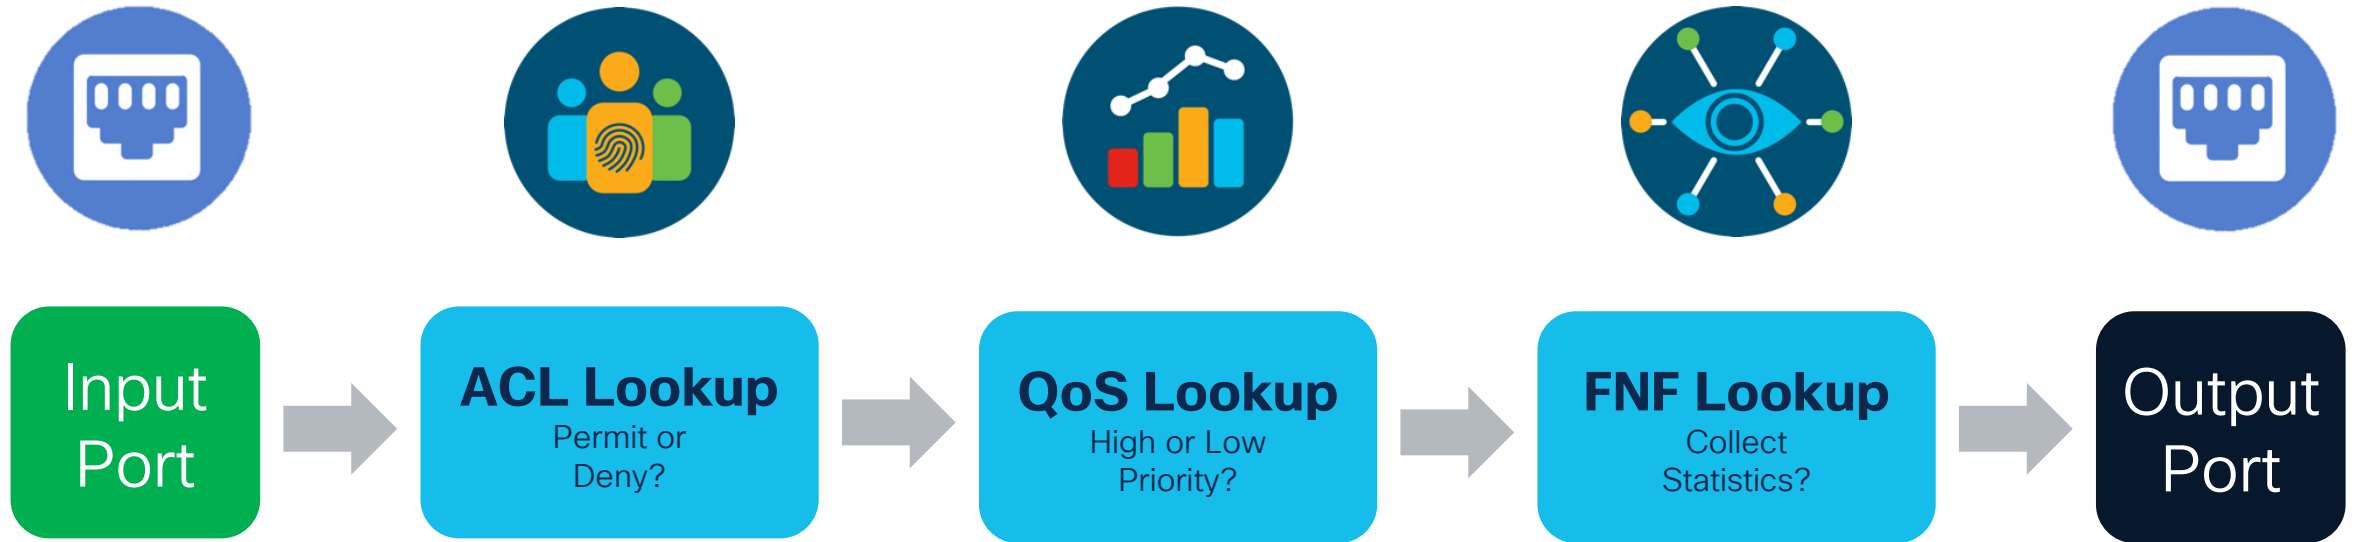
ASICs are fundamental to network devices



Modern processing speeds are **Terabits per second (Tbps)**

# What does an ASIC do?

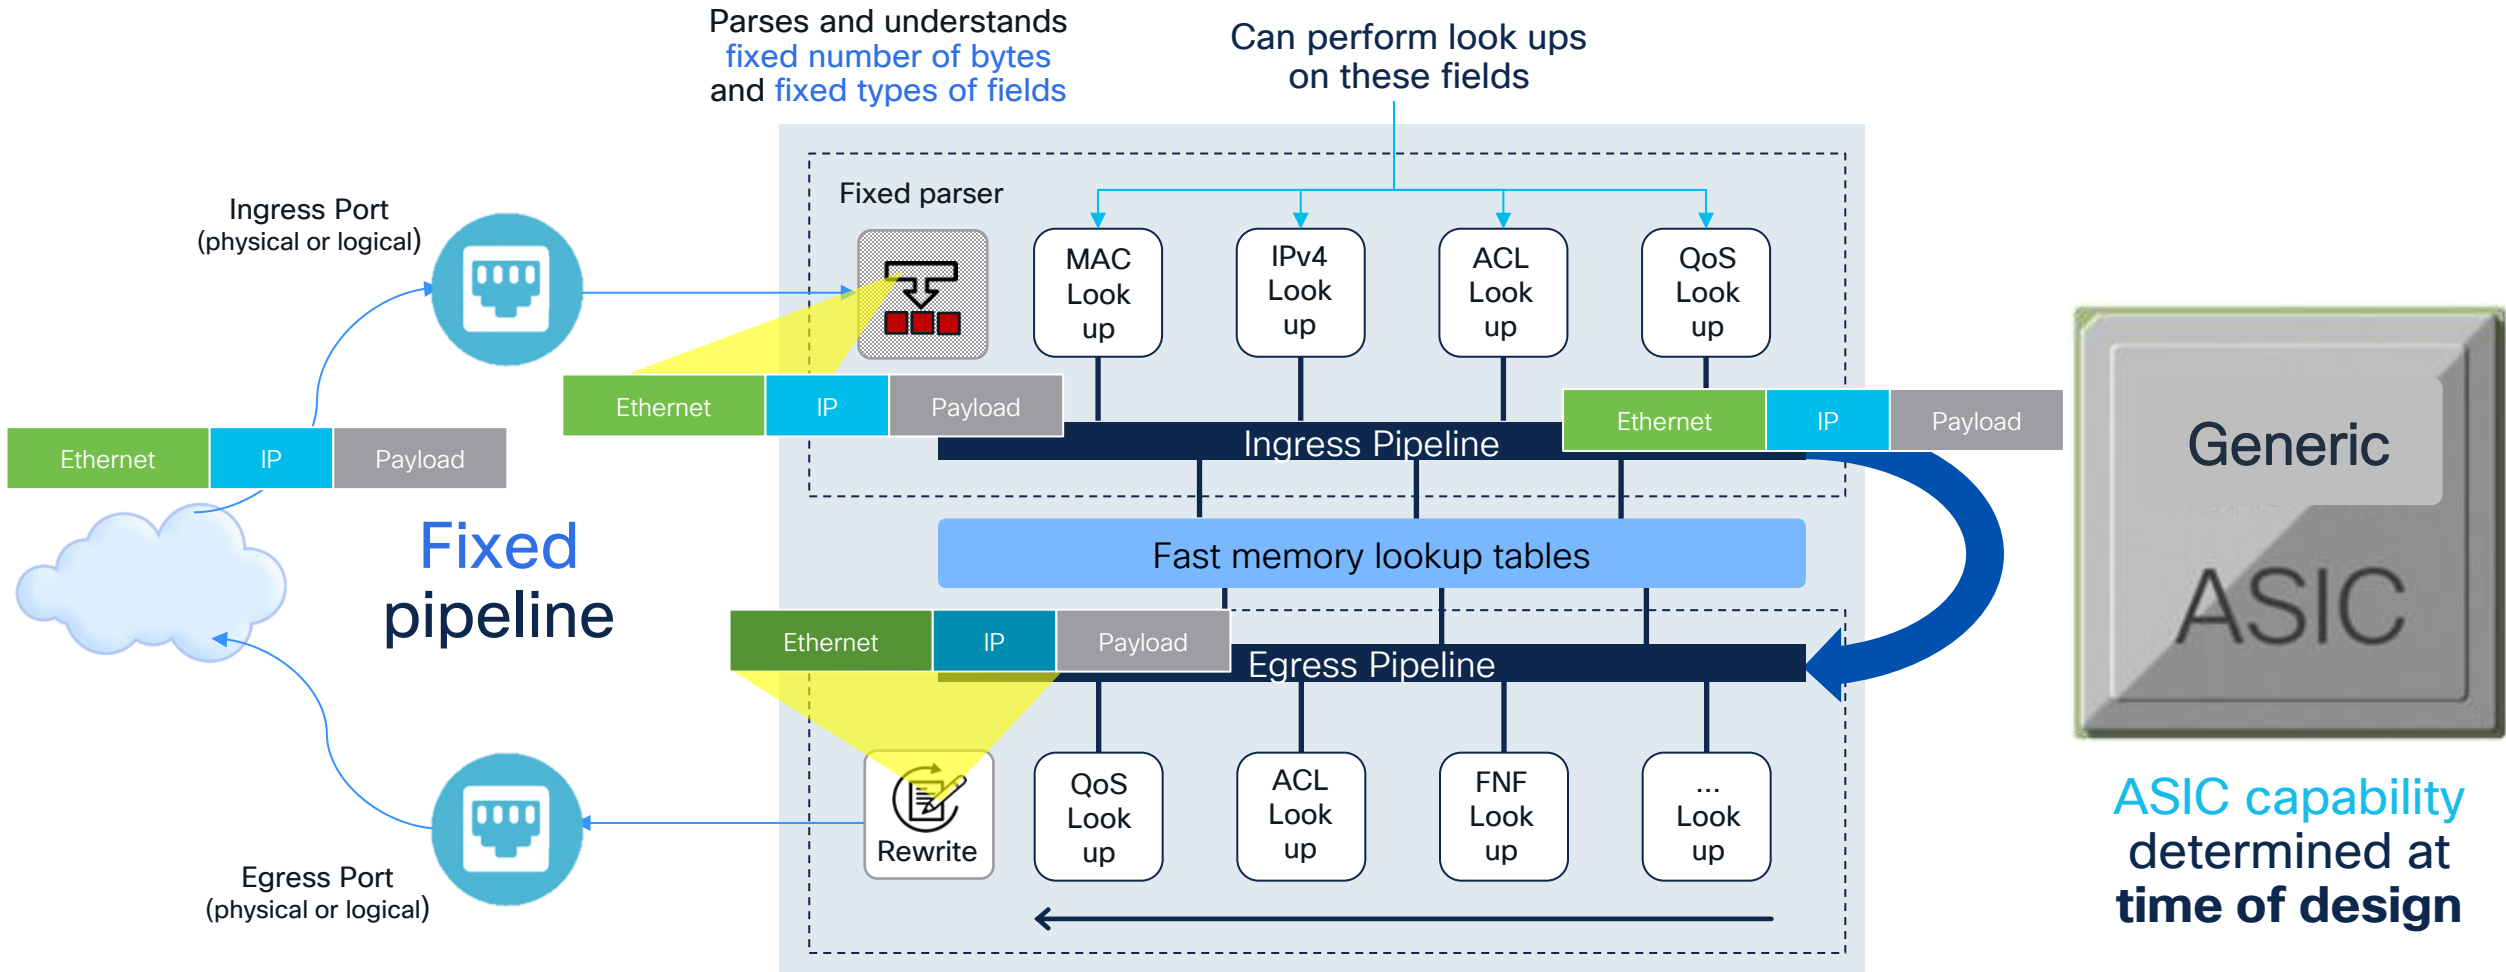
ASICs are fundamental to network devices



Common services are **Access Control, Quality of Service and Flow Analytics**


# The Past - Fixed Network ASICs

## Fixed Processing Pipeline



# Flexible ASICs for Enterprise Switching

- 1 Why ASICs?
- 2 Cisco UADP
- 3 Cisco S1
- 4 Cisco 9000 Series
- 5 ASIC Futures



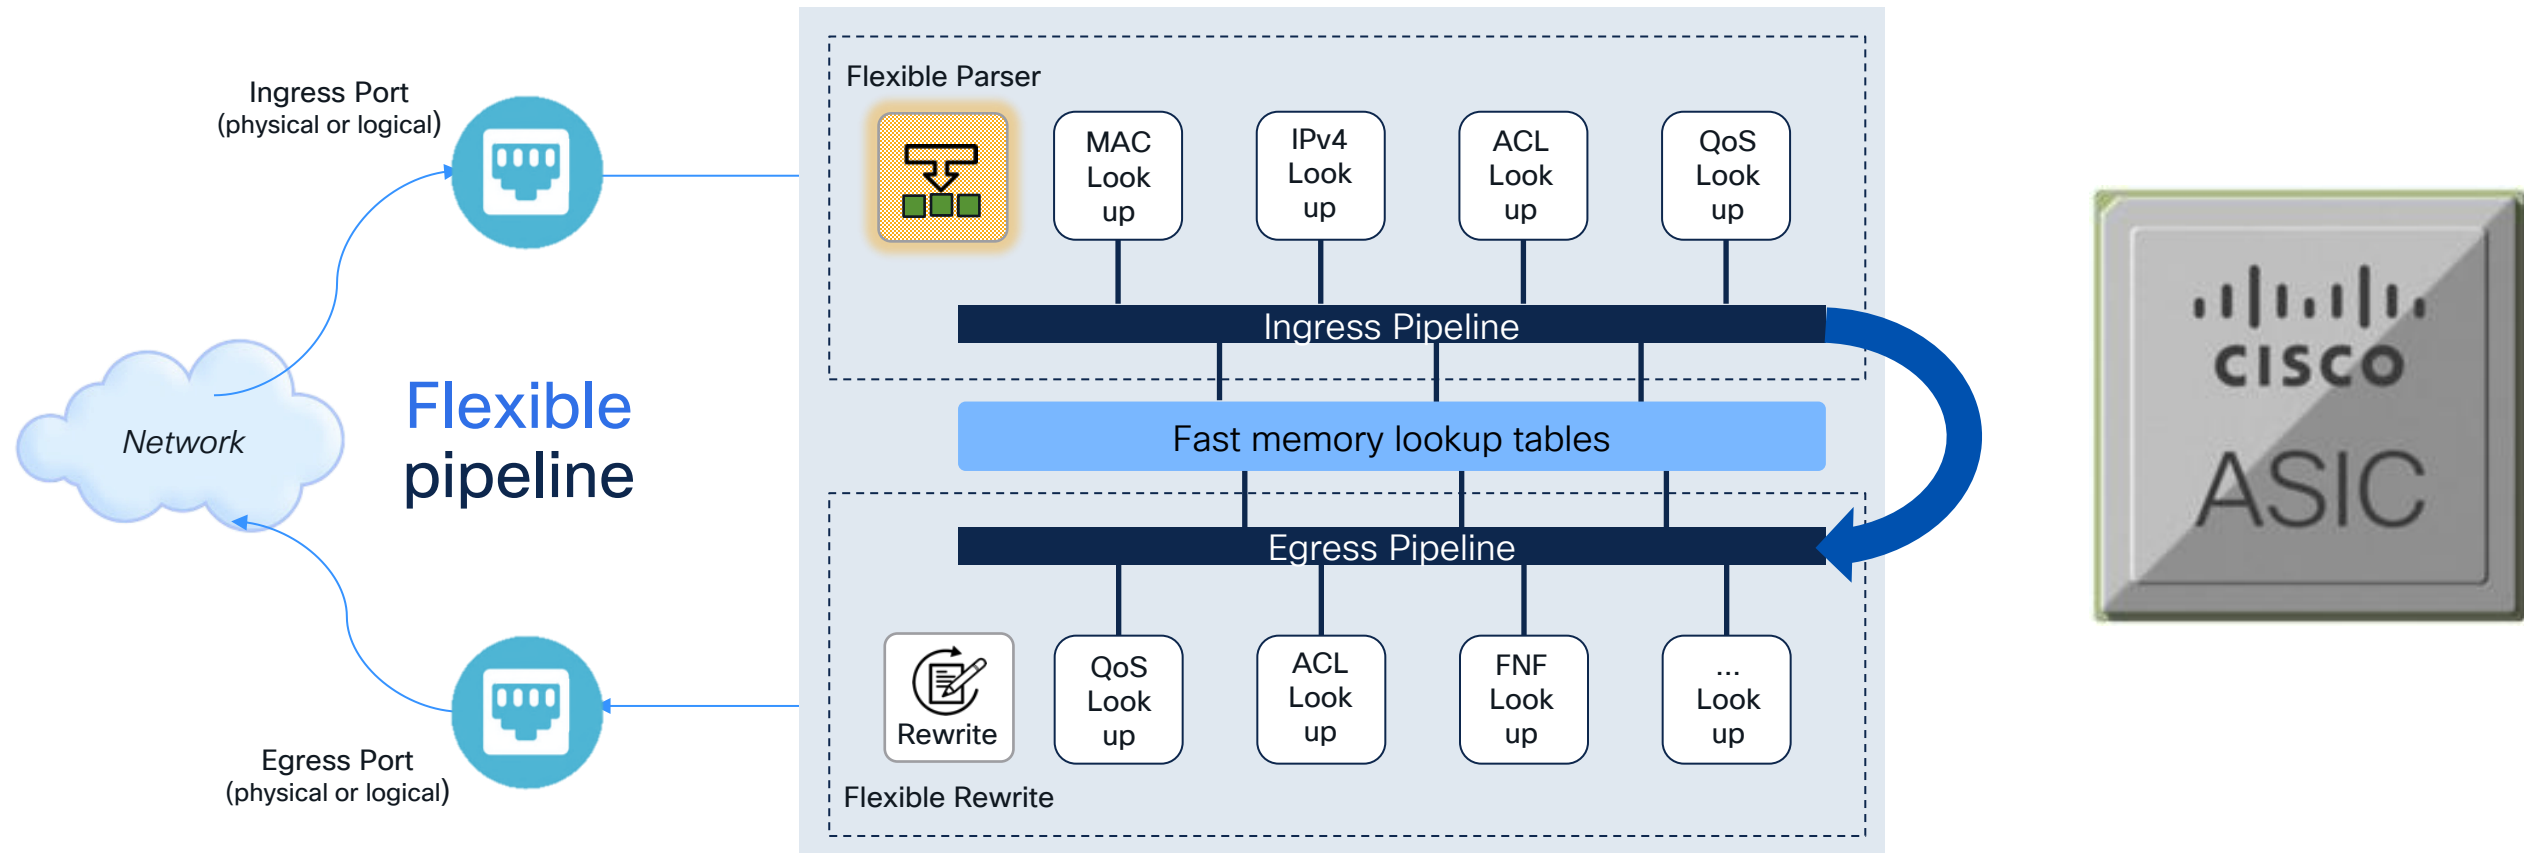
# Programmable Network ASICs

## Flexible Parsing

## Flexible Parsing

Look deep into the packet header, with **programmable field parsing**

Parses and understands  
multiple programmable headers  
with flexible field definitions

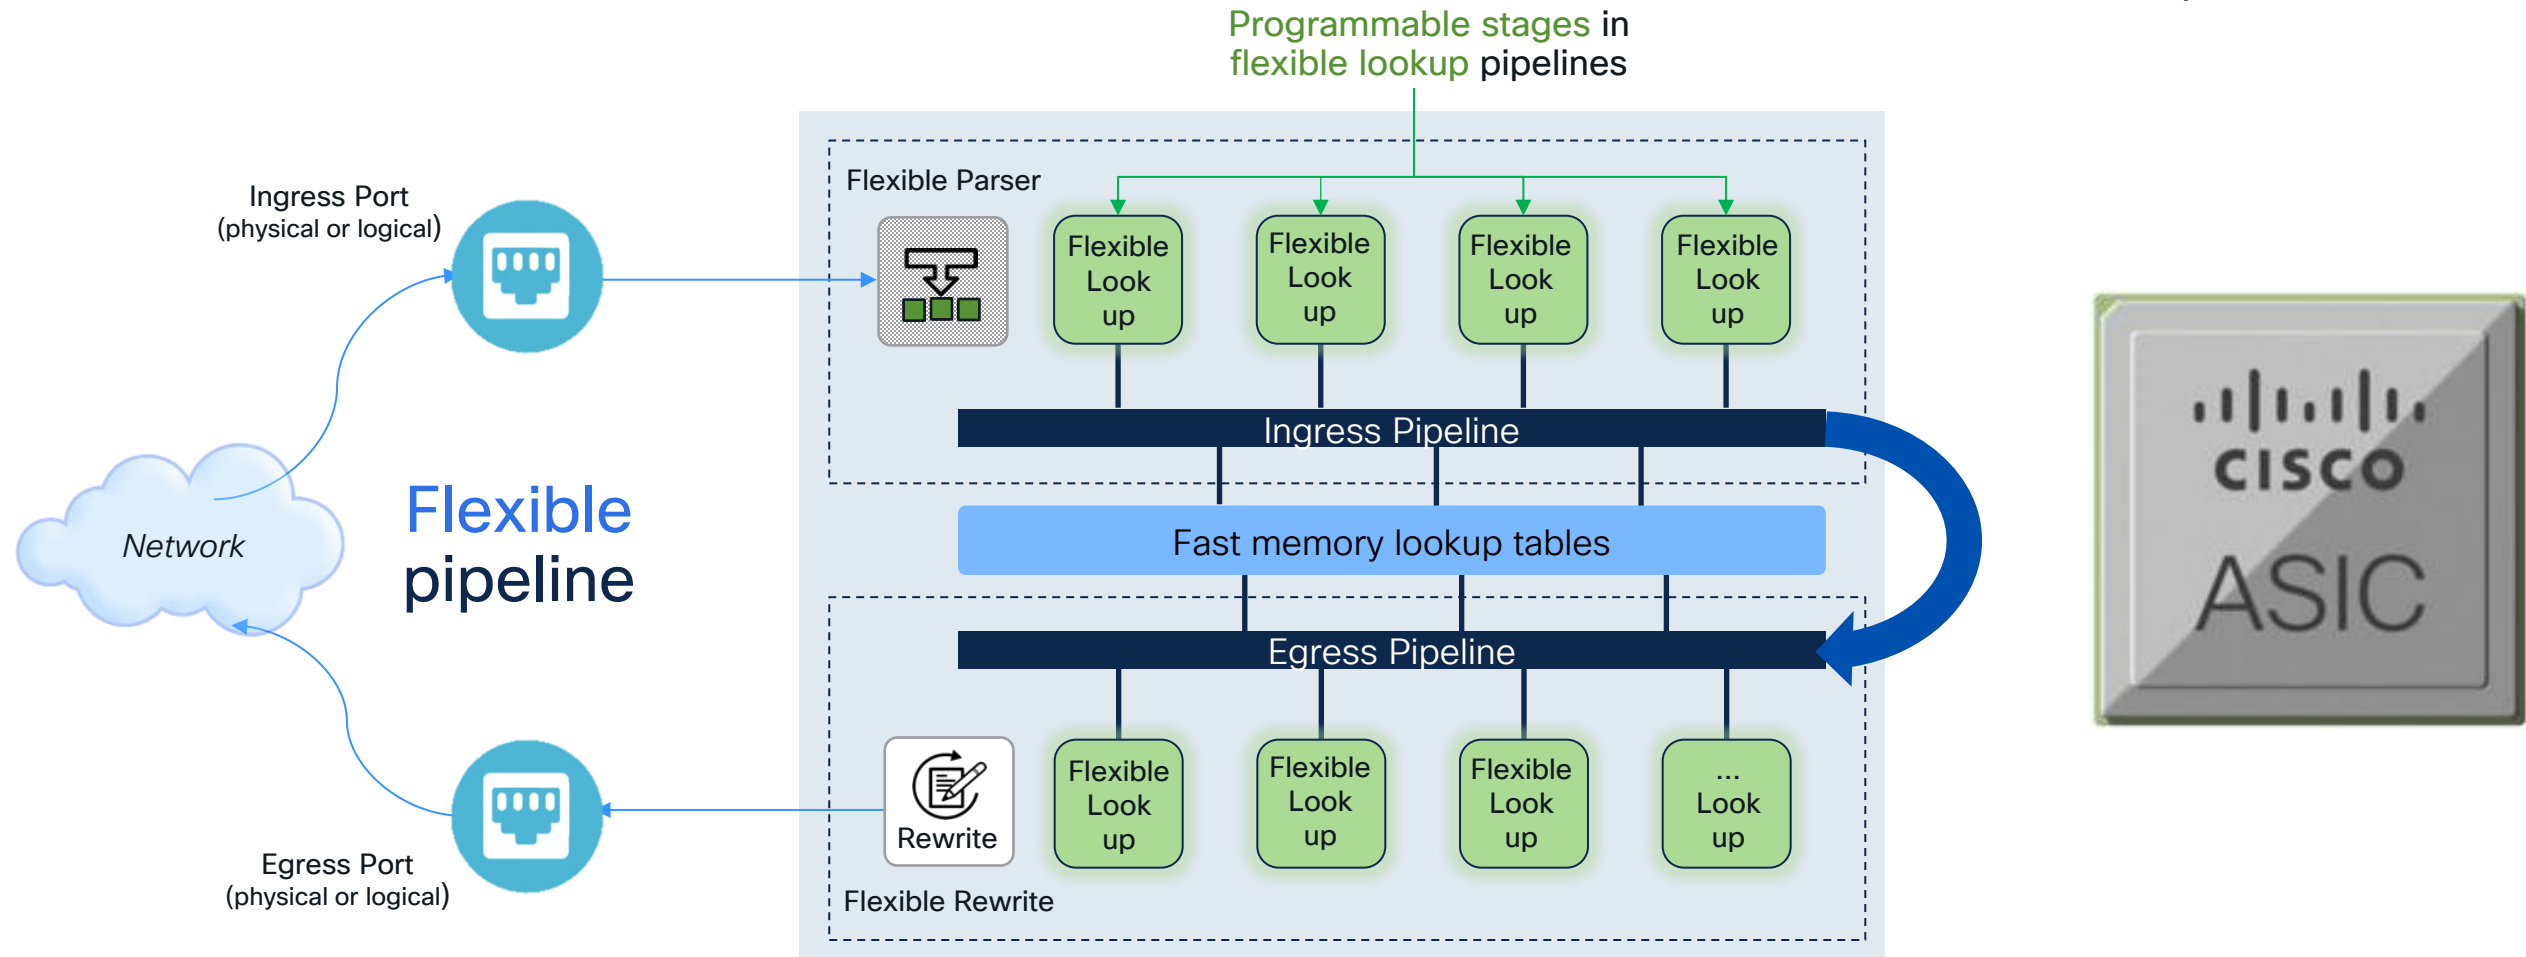


# Programmable Network ASICs

## Flexible Lookups

## Flexible Lookups

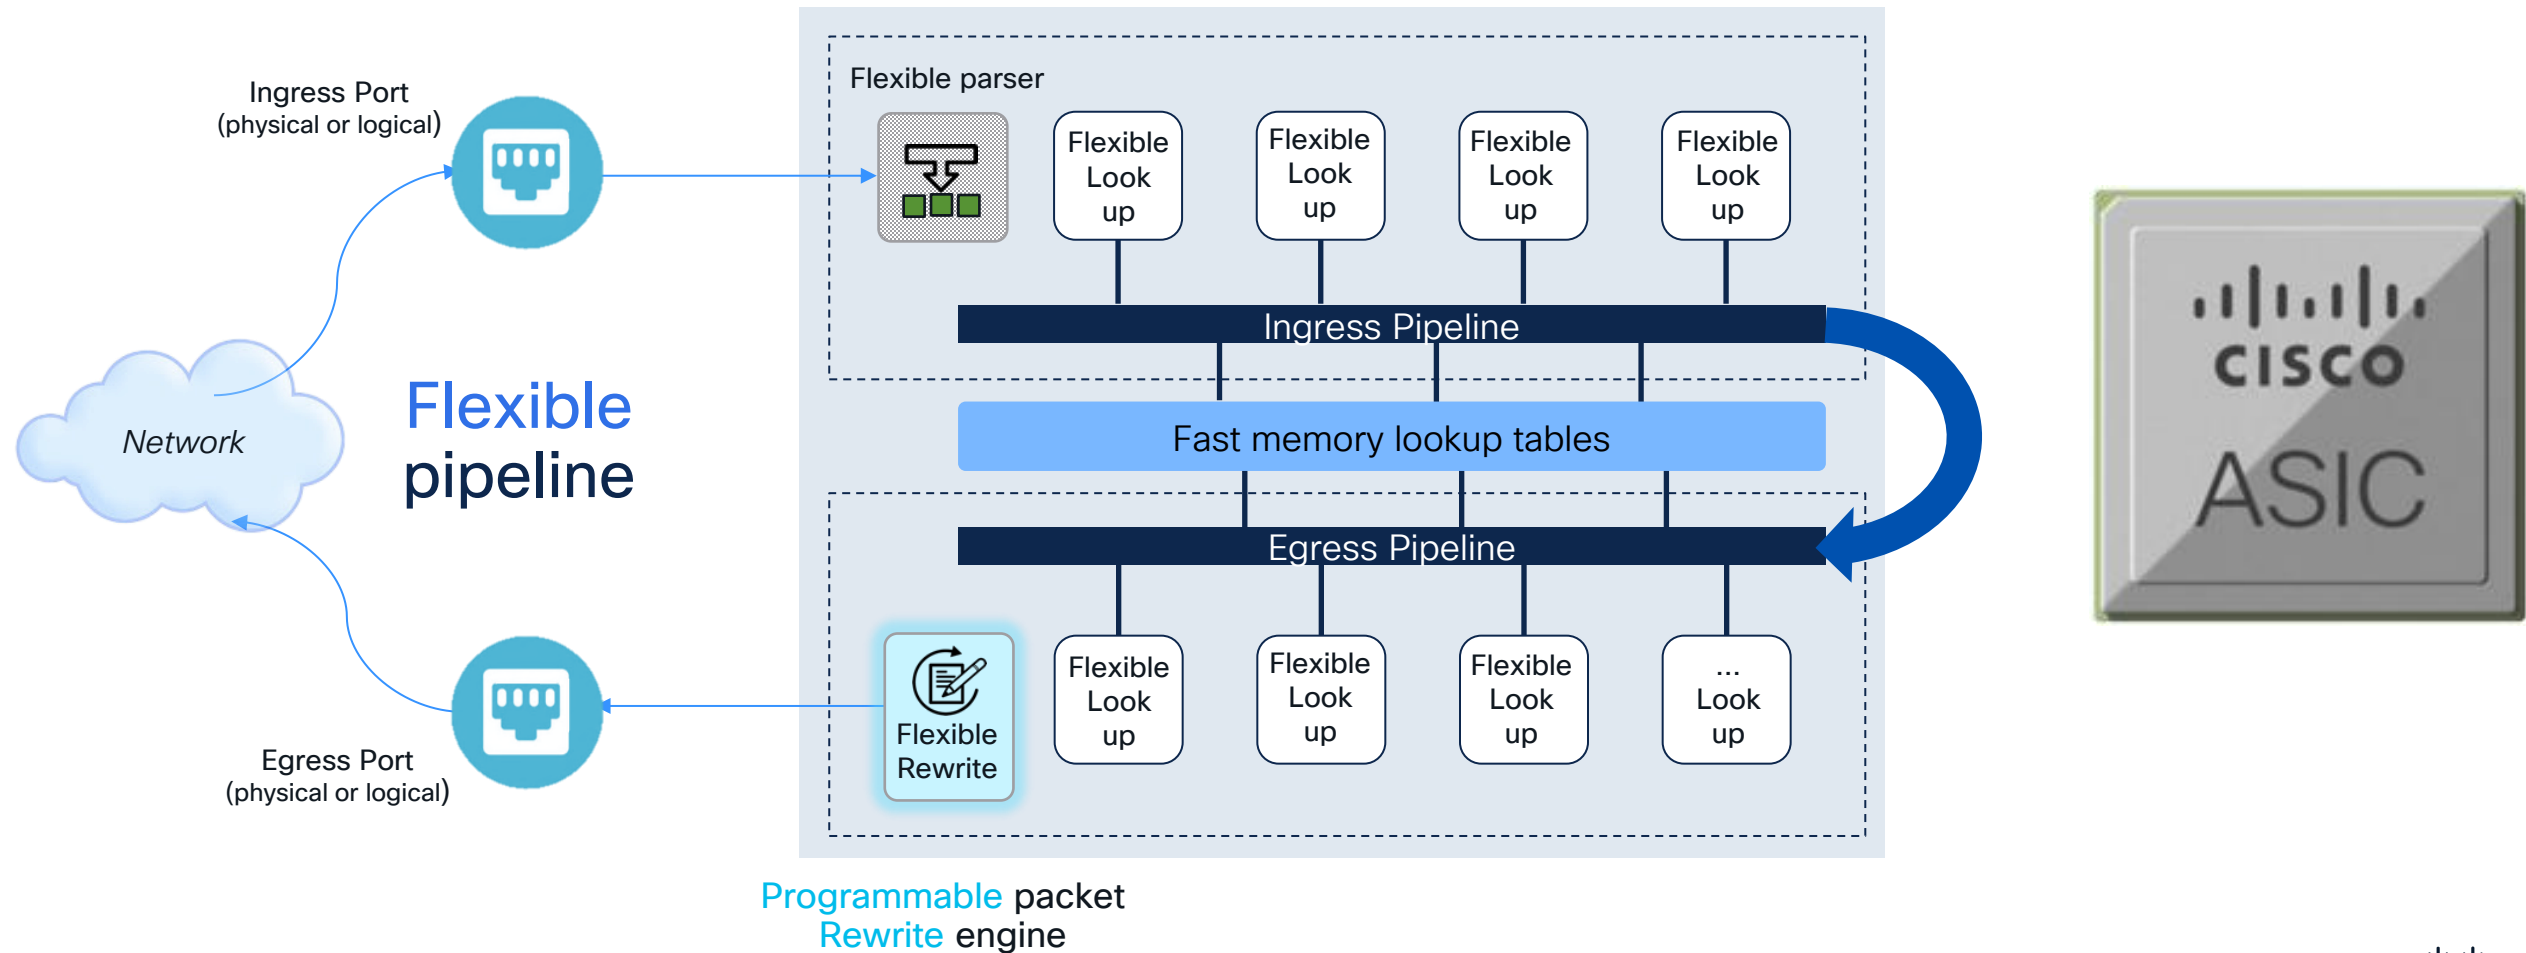
Multi-stage packet handling, with **flexible packet lookups** at every step



# Programmable Network ASICs

Flexible Rewrites

**Flexible Rewrites**  
Flexible packet handling and forwarding,  
with a programmable packet rewrite



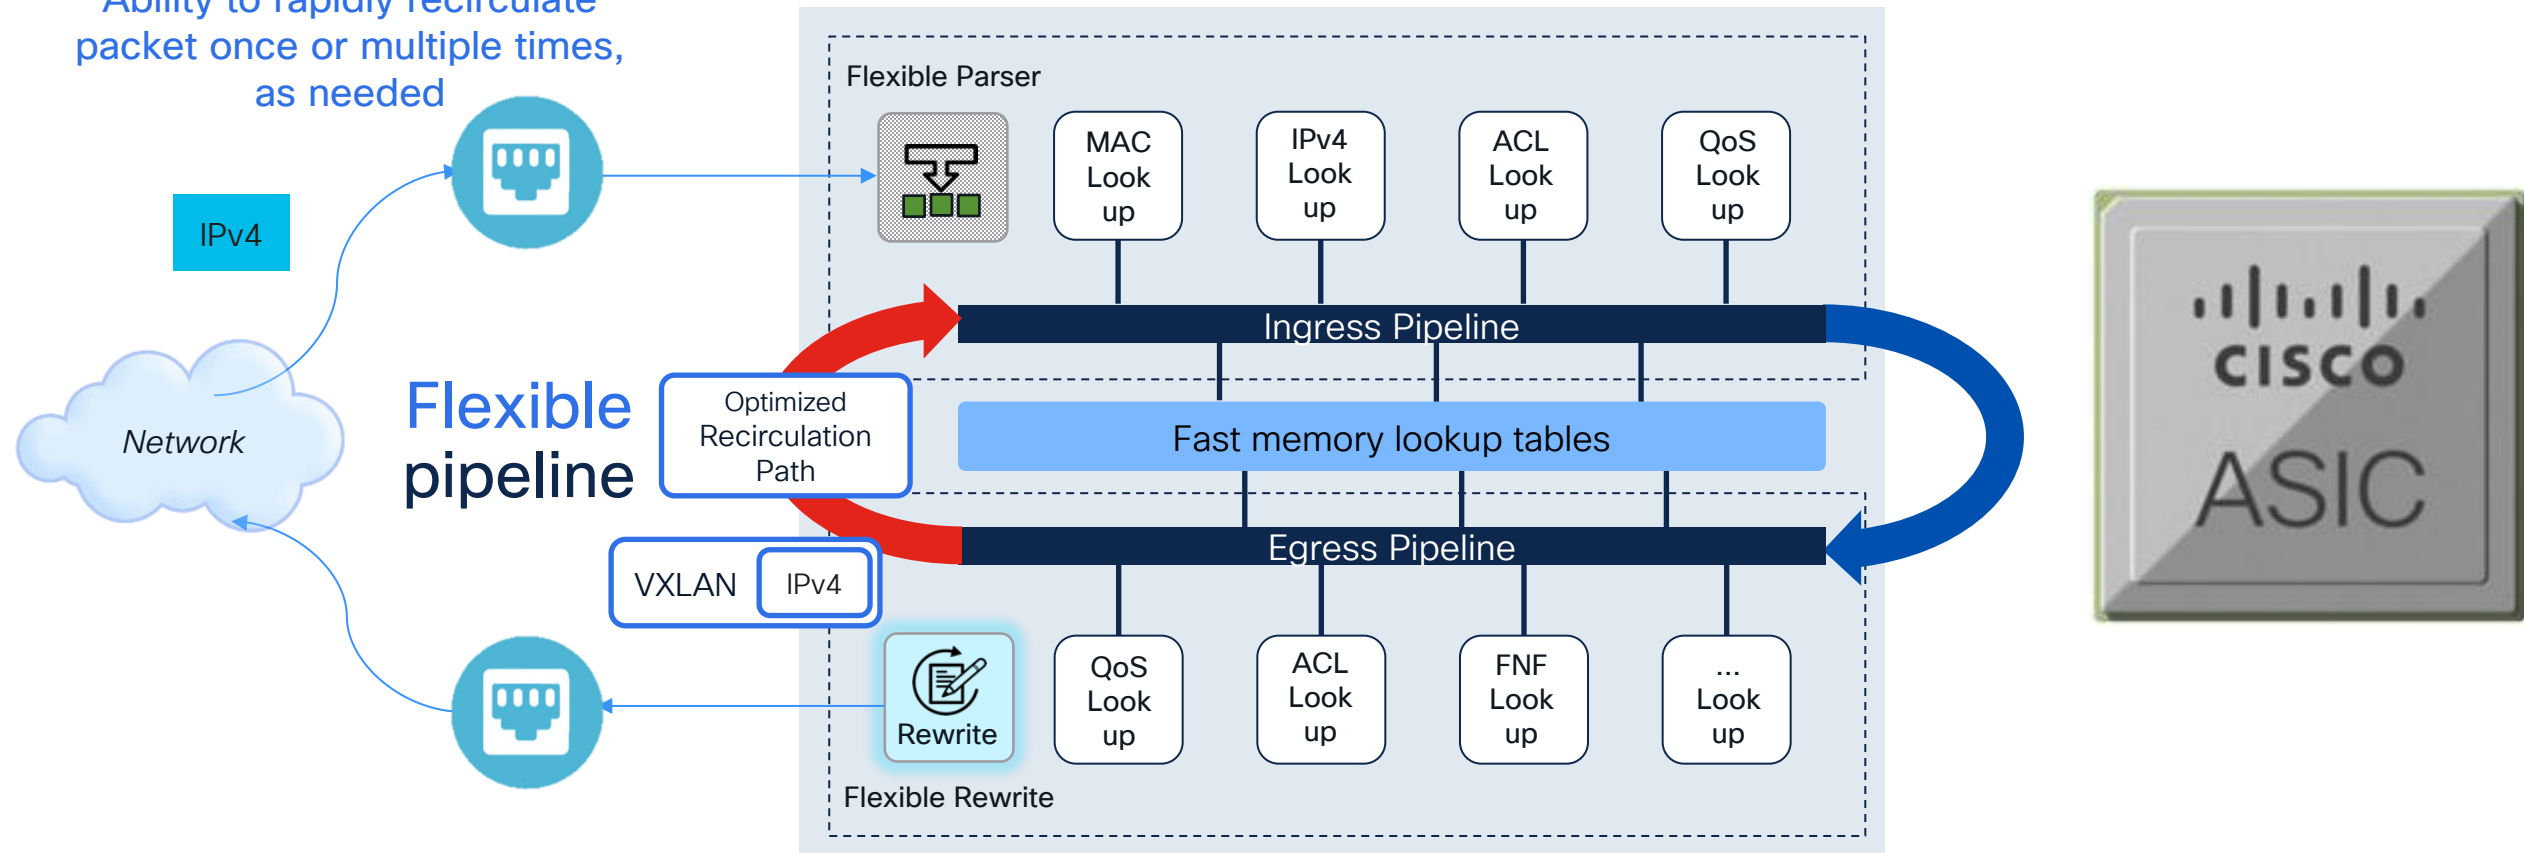
# Programmable Network ASICs

Optimized Recirculation

## Optimized Recirculation

Highly **optimized recirc path** for packet header addition / removal / forwarding


Ability to rapidly recirculate packet once or multiple times, as needed



# Programmable Network ASICs

Balancing Cost, Performance & Flexibility


Programmable  
ASIC



Cost ✓

Performance ✓

Flexibility ✓



New ASIC  
Functionality





ASIC  
Engineer



Microcode  
Update

# Creating Custom ASICs

From Definition to Deployment




3 to 5 years

**Building a new Custom ASIC takes a lot of Time and Money**

Yes... we are already building the 'Next Generation' 🤪

# Cisco UADP for Enterprise Switching

- 1 Why ASICs?
- 2 Cisco UADP
- 3 Cisco S1
- 4 Cisco 9000 Series
- 5 ASIC Futures



# Cisco Unified Access Data-Plane (UADP®)

Common ASIC Architecture for Campus Switching Access, Distribution & Core



[community.cisco.com/t5/networking-blogs/uadp-the-powerhouse-of-catalyst-9000-family/ba-p/3764605](https://community.cisco.com/t5/networking-blogs/uadp-the-powerhouse-of-catalyst-9000-family/ba-p/3764605)



**UADP 2.0m**

**120 Gbps**

16nm FinFET  
1 Core + ARM CPU



**UADP 2.0/XL**

**240 Gbps**

28nm FinFET  
2 Core



**UADP 2.0sec**

**480 Gbps**

16nm FinFET  
1 Core2 + SEC



**UADP 3.0**

**1.6 Tbps**

16nm FinFET  
2 Core



**UADP 3.0sec**

**1.6 Tbps**

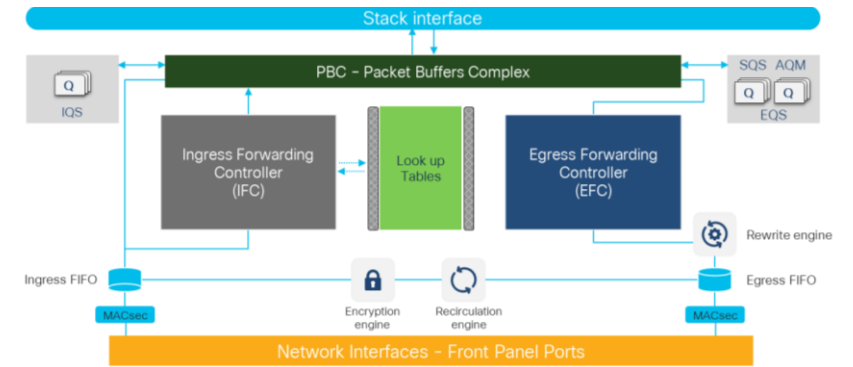
16nm FinFET  
2 Core + SEC

- Multiple generations and formats, same architecture
- Rich flexible forwarding & services memories
- First fully programmable microcode network silicon

- Multiple functions: system-on-chip or line-card
- Multiple form factors: fixed or modular
- Multiple places: Access, Distribution and Core

# Catalyst 9000 with UADP

The industry's first Programmable Campus ASIC Family



**Catalyst 9200CX**  
UADP 2.0 mini



**Catalyst 9300X/LM**  
UADP 2.0sec



**Catalyst 9400X-SUP2**  
UADP 3.0sec



**Catalyst 9200/L**  
UADP 2.0 mini



**Catalyst 9300/L**  
UADP 2.0



**Catalyst 9400-SUP1 & 9500**  
UADP 2.0XL



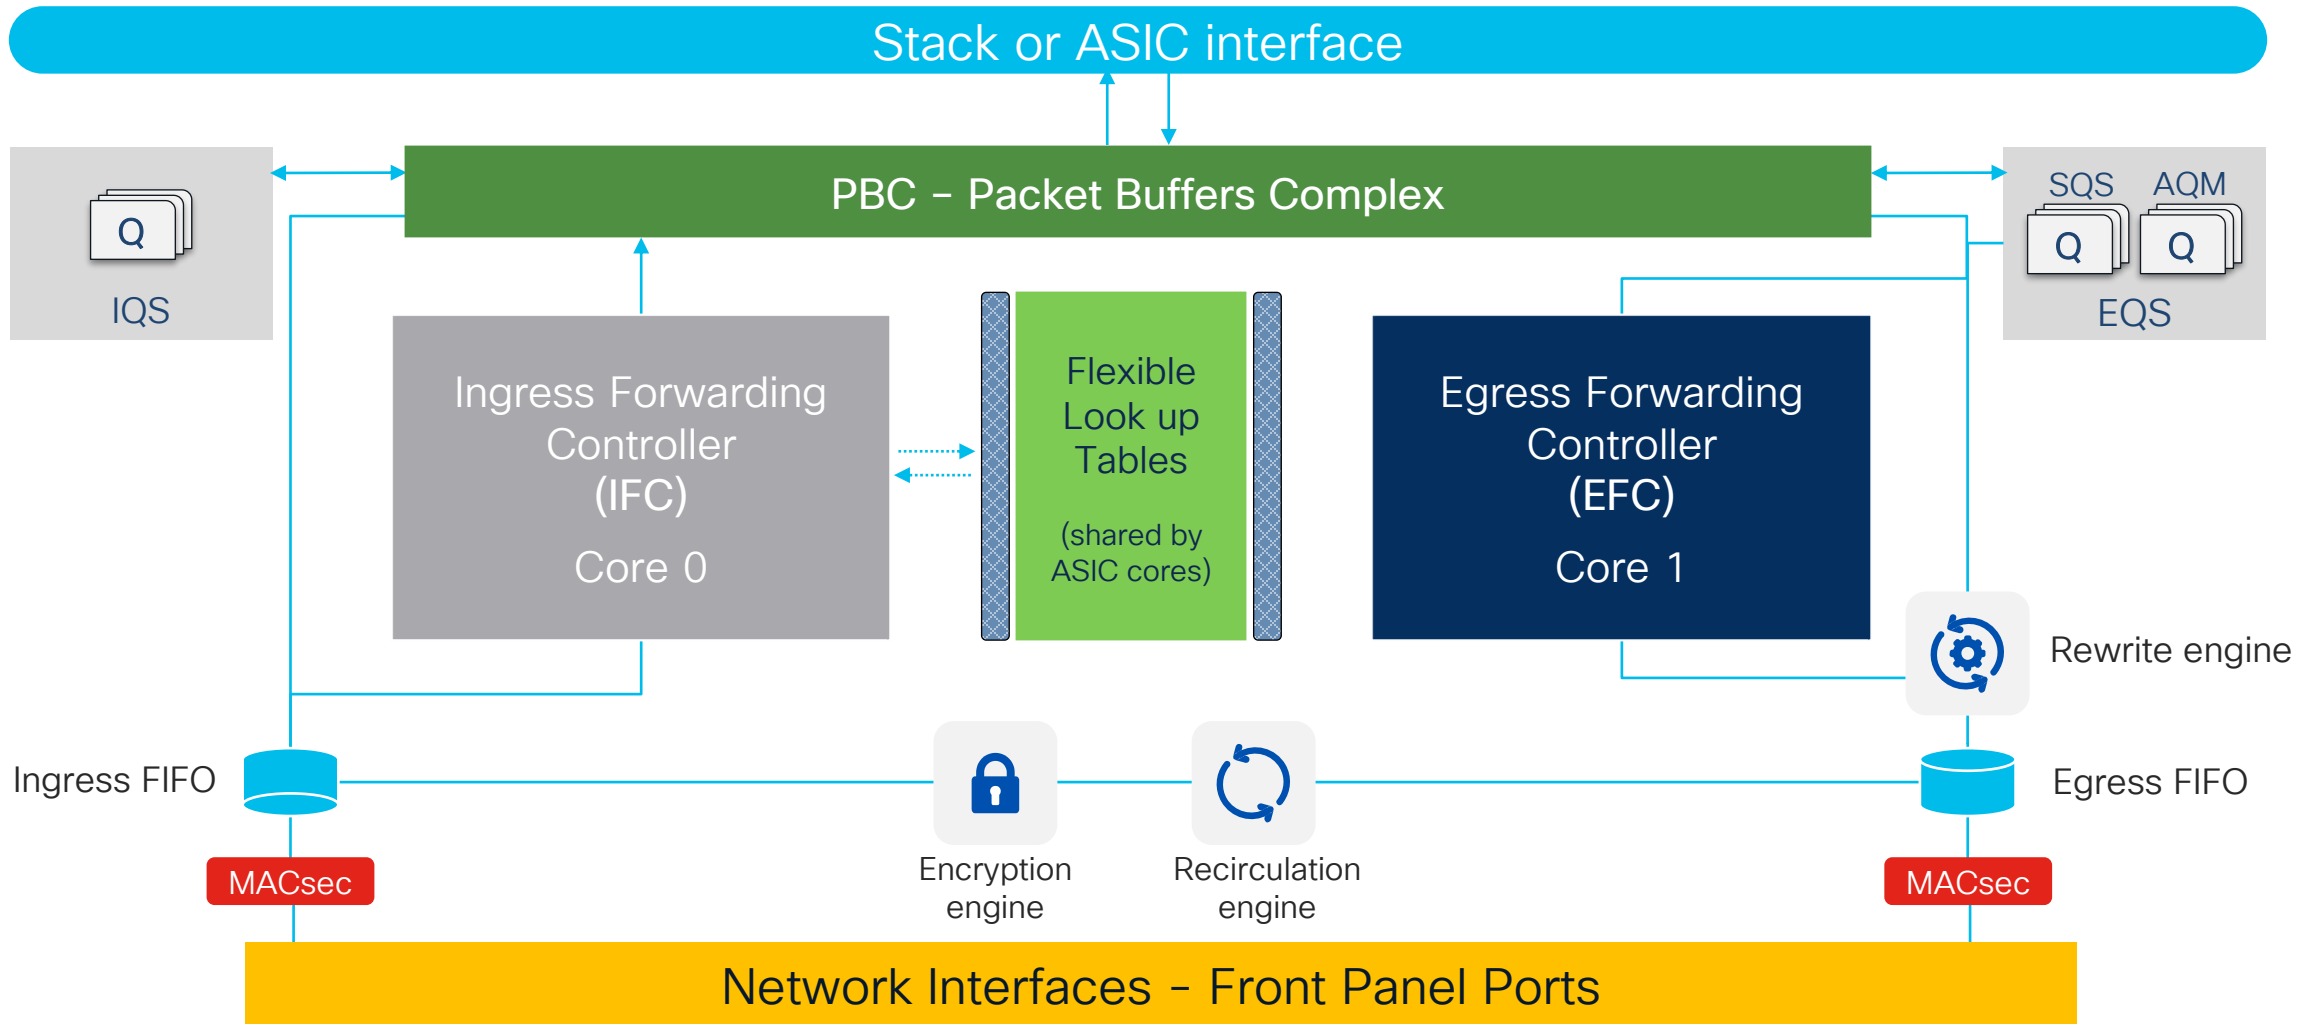
**Catalyst 9500(H) & 9600-SUP1**  
UADP 3.0



# Cisco UADP ASICs

## ASIC Architecture & Block Diagram

up to 1 BILLION times per second!

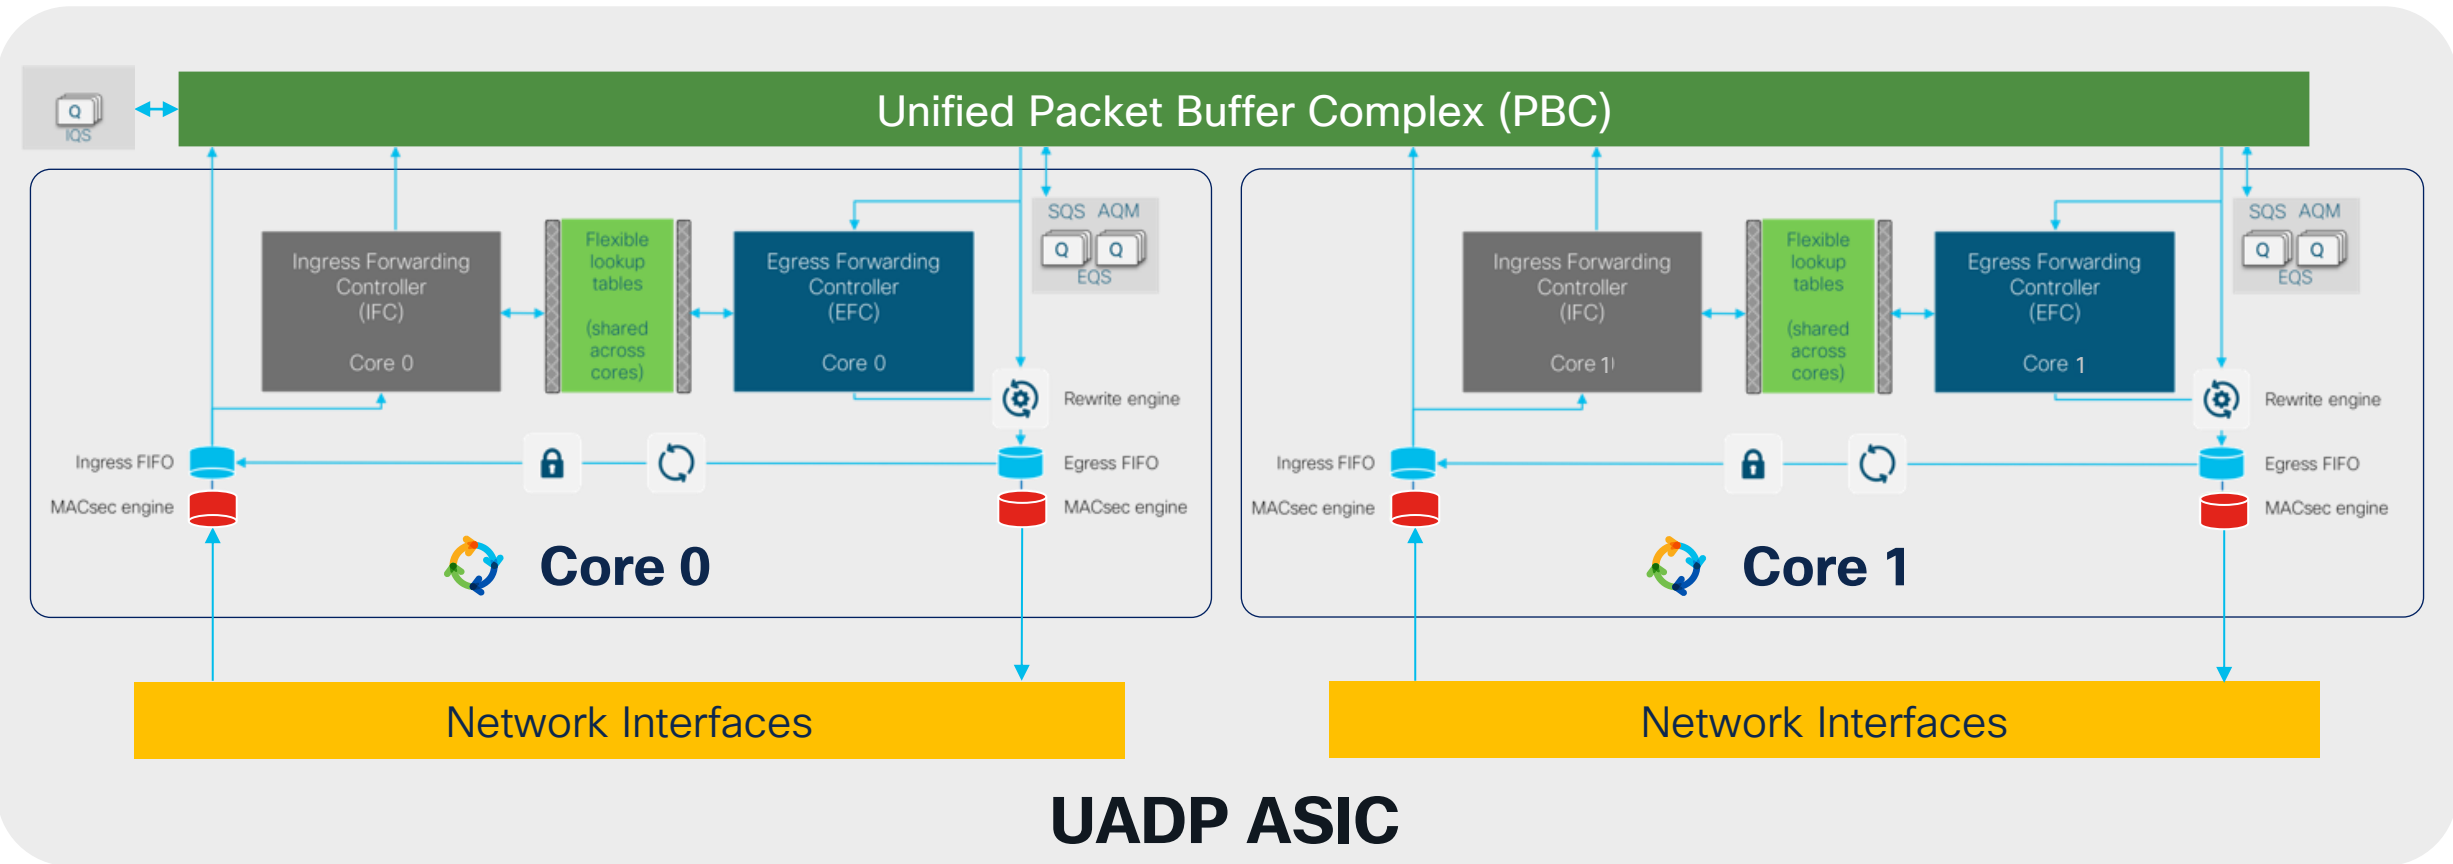


# Cisco UADP – Dual Cores

Multiple ASIC Cores & Unified Packet Buffers

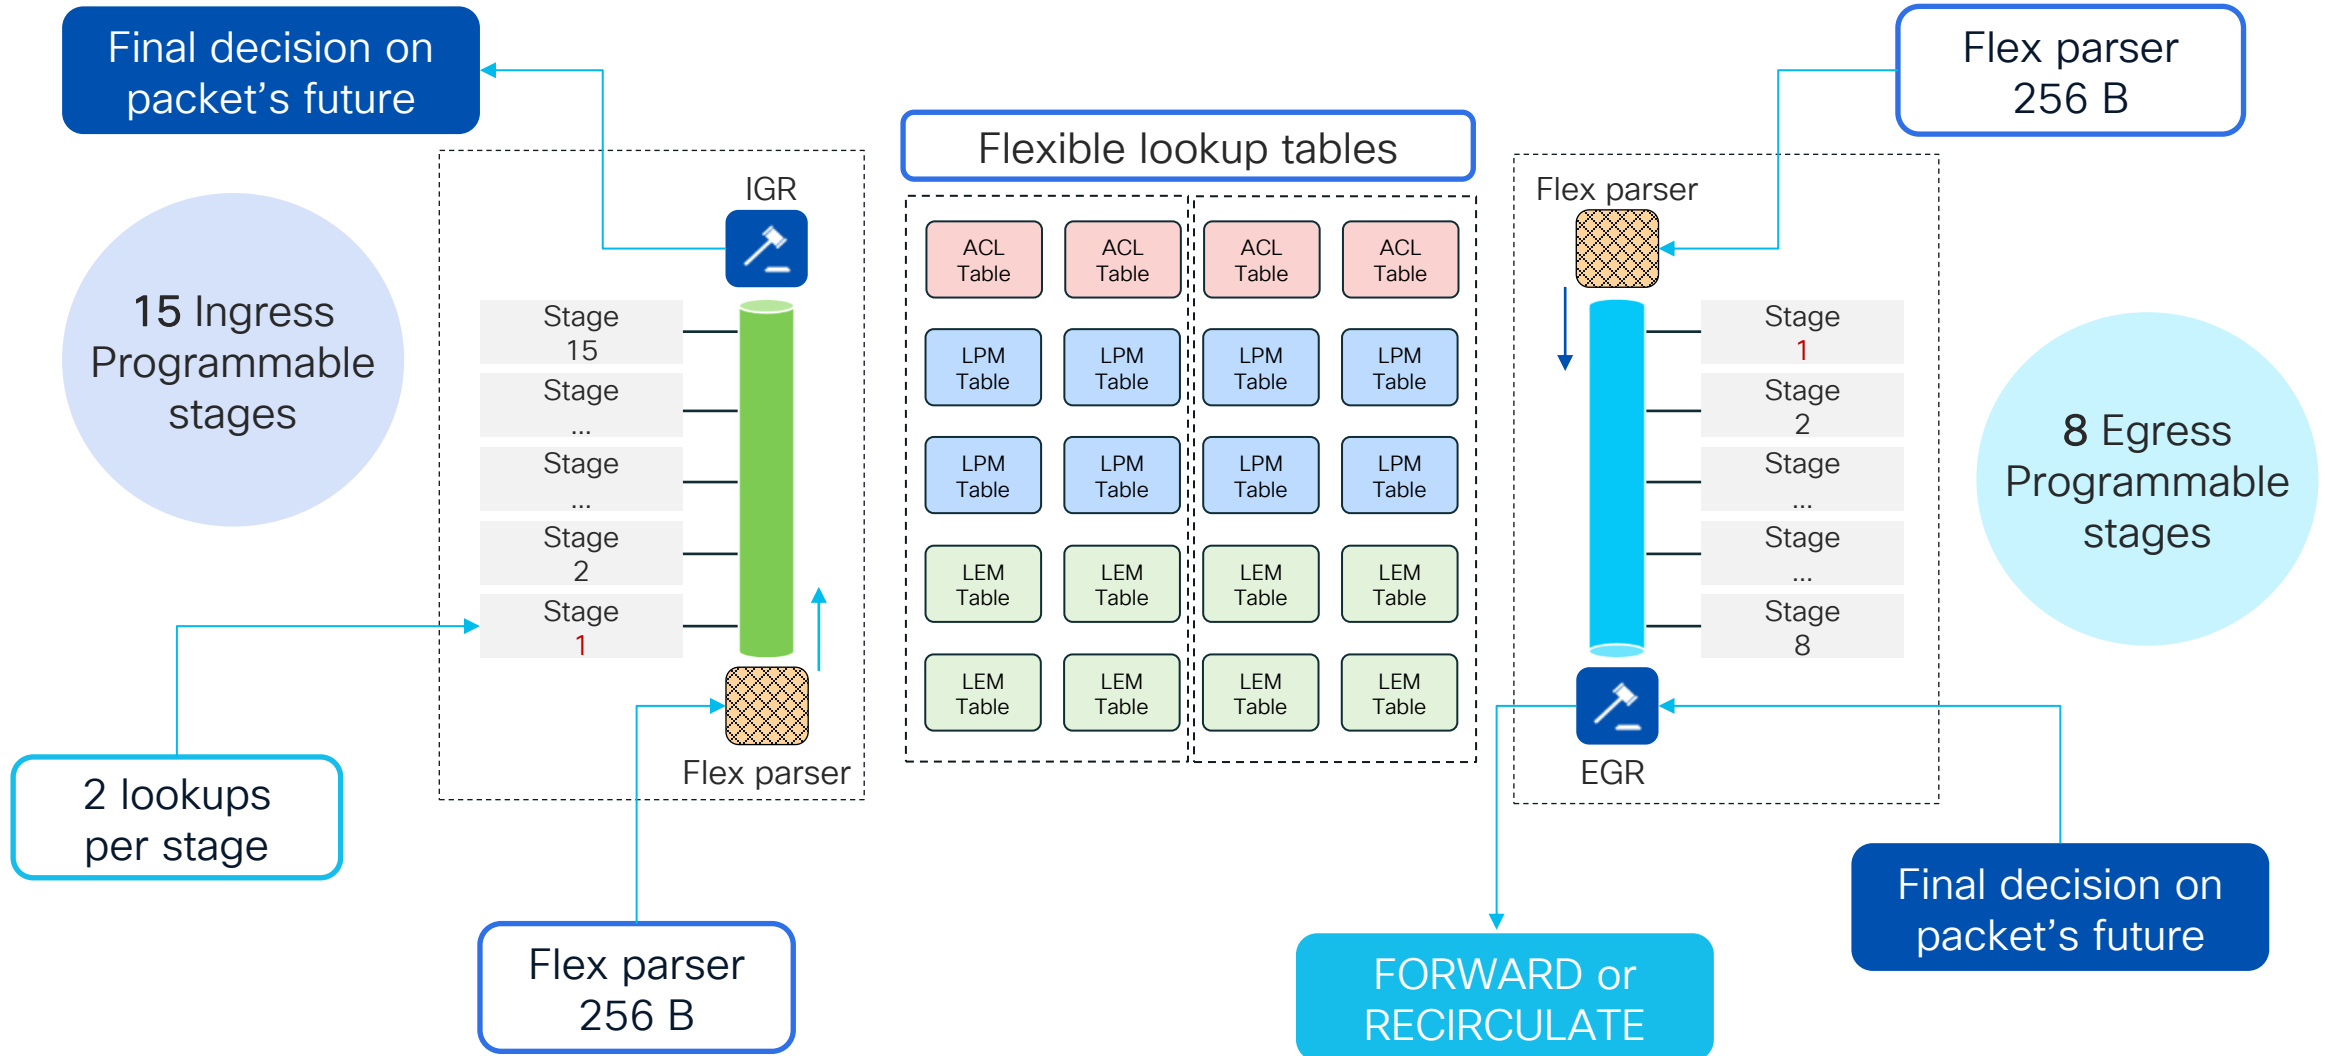


Each ASIC “Core” services a certain number of front-panel ports



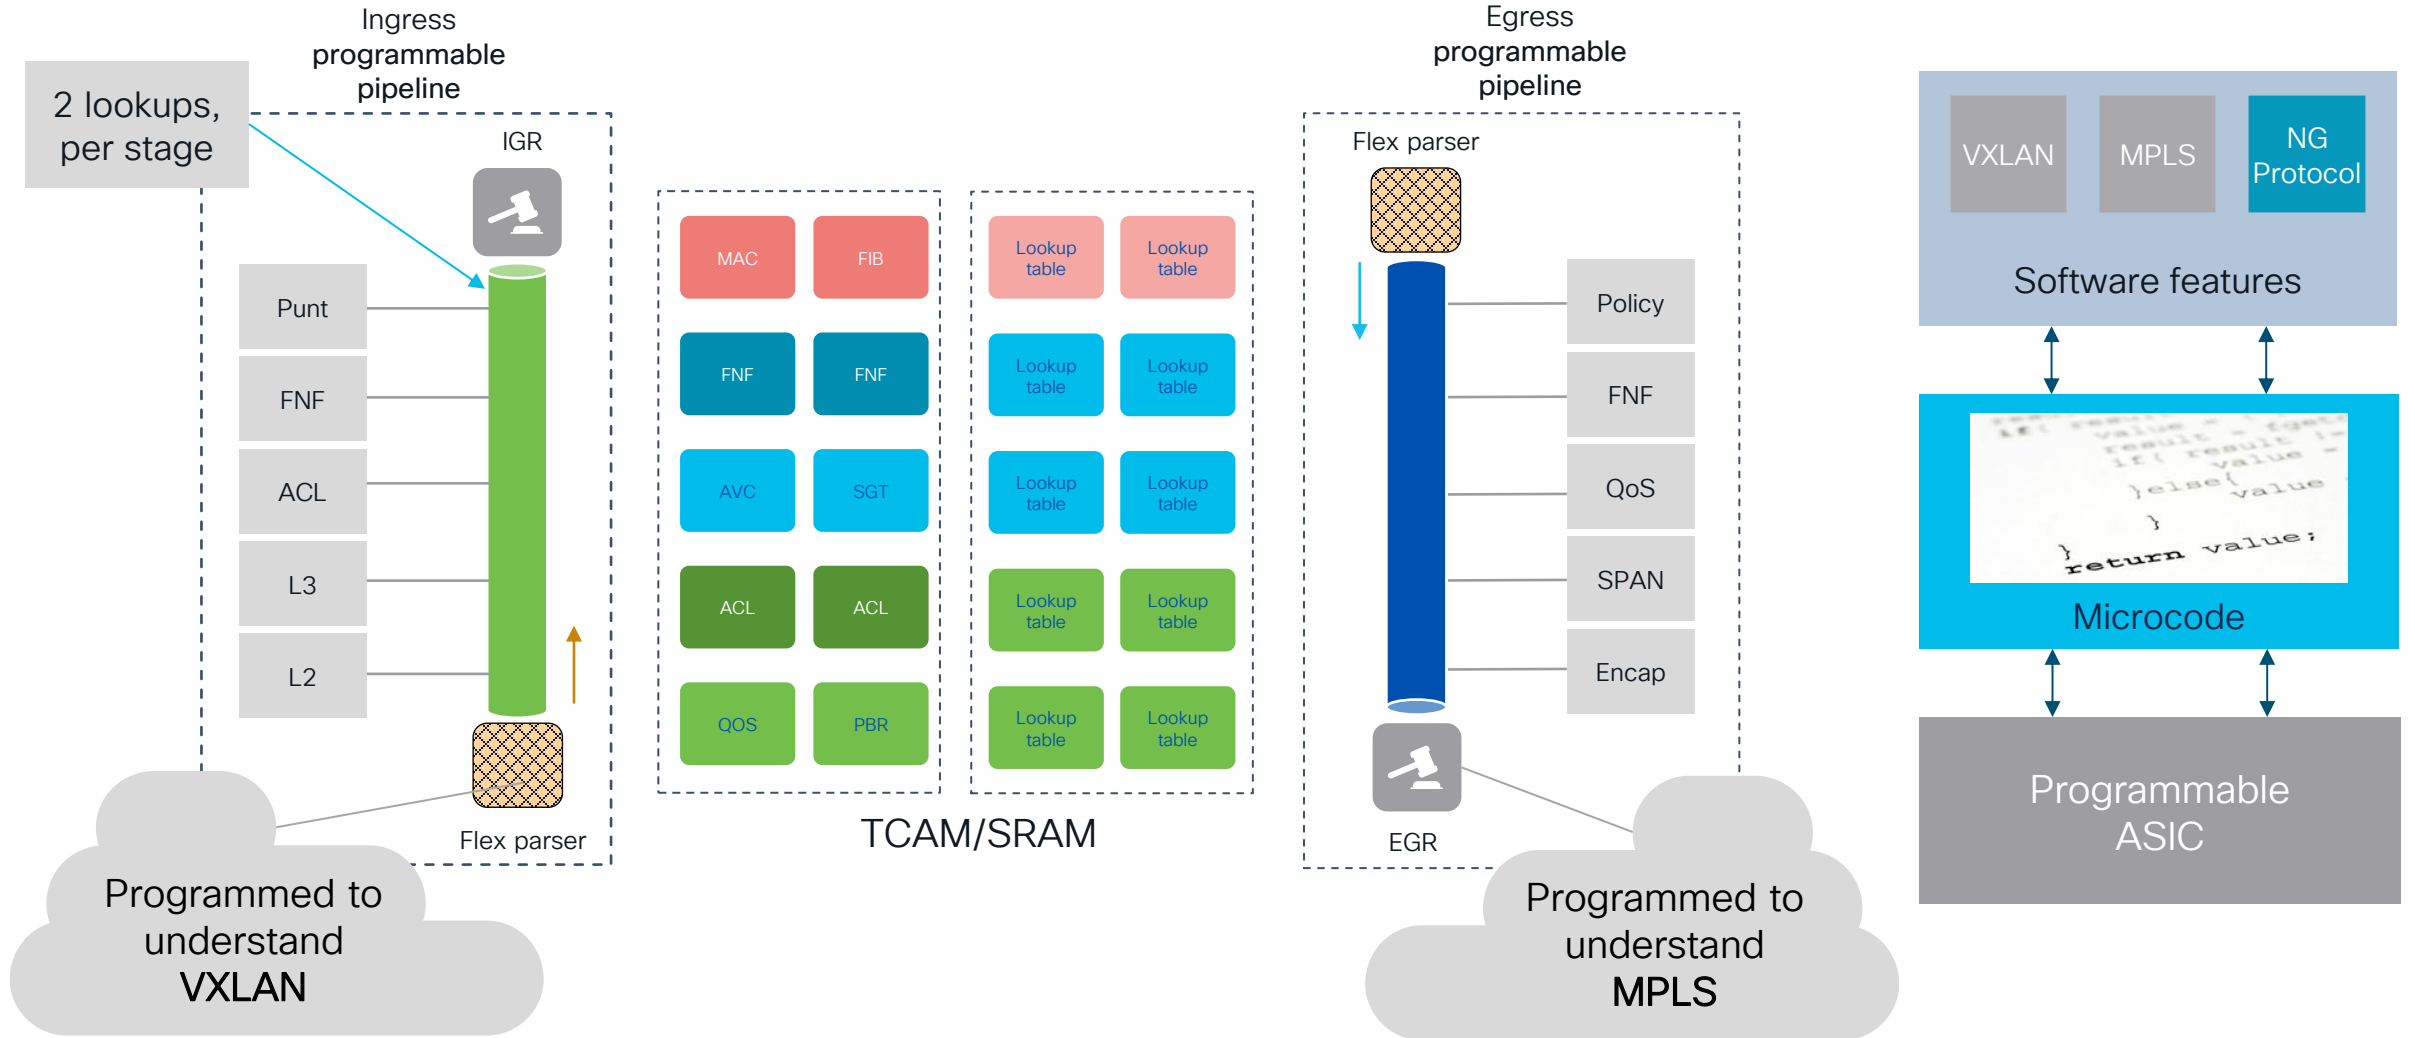
# Cisco UADP - Flexible Lookups

Programmable Ingress and Egress Processing Stages




# Cisco UADP – Microcode

ASIC Microcode (NPL/SDK) can be upgraded to add new features



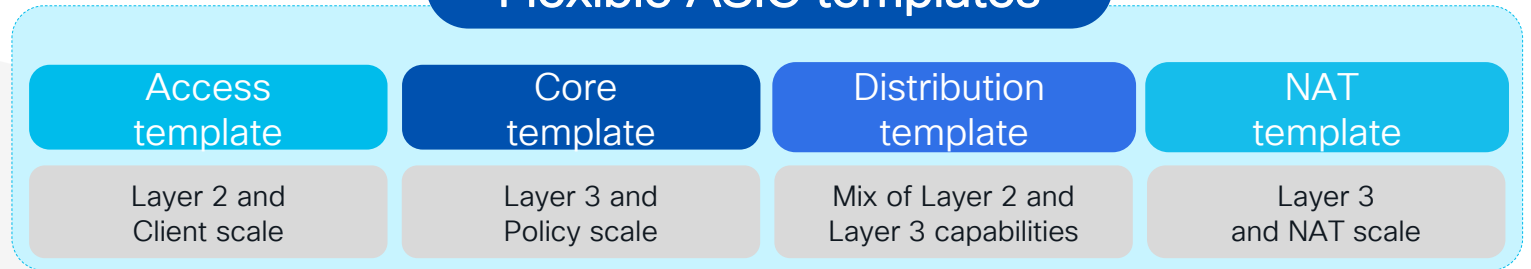
# Cisco UADP - Flexible Tables

Customizable ASIC tables for universal deployment flexibility



- MAC
- IPv4/v6
- Unicast
- Multicast
- NetFlow
- ACL
- SGACL
- QoS
- NAT
- SPAN

## Flexible ASIC templates



Customize table sizes for each function, based on the Place in Network

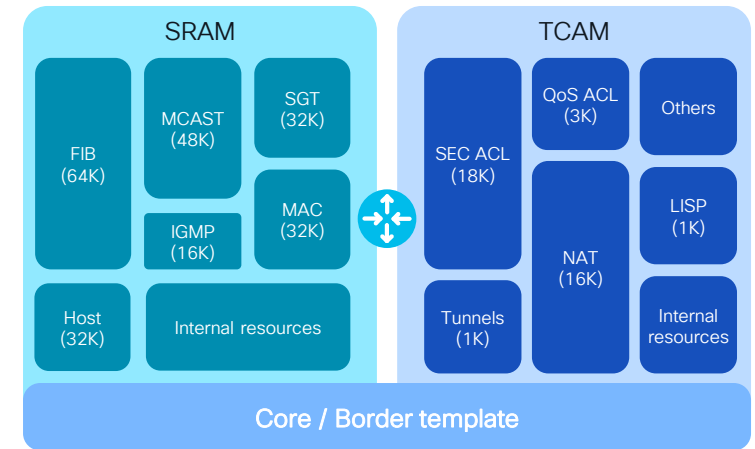
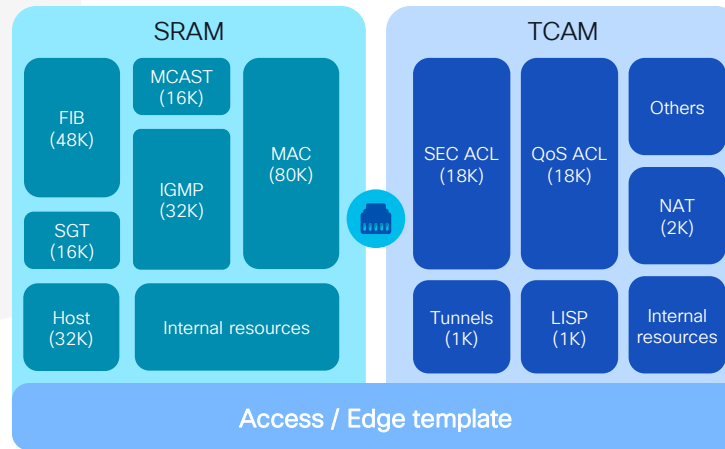
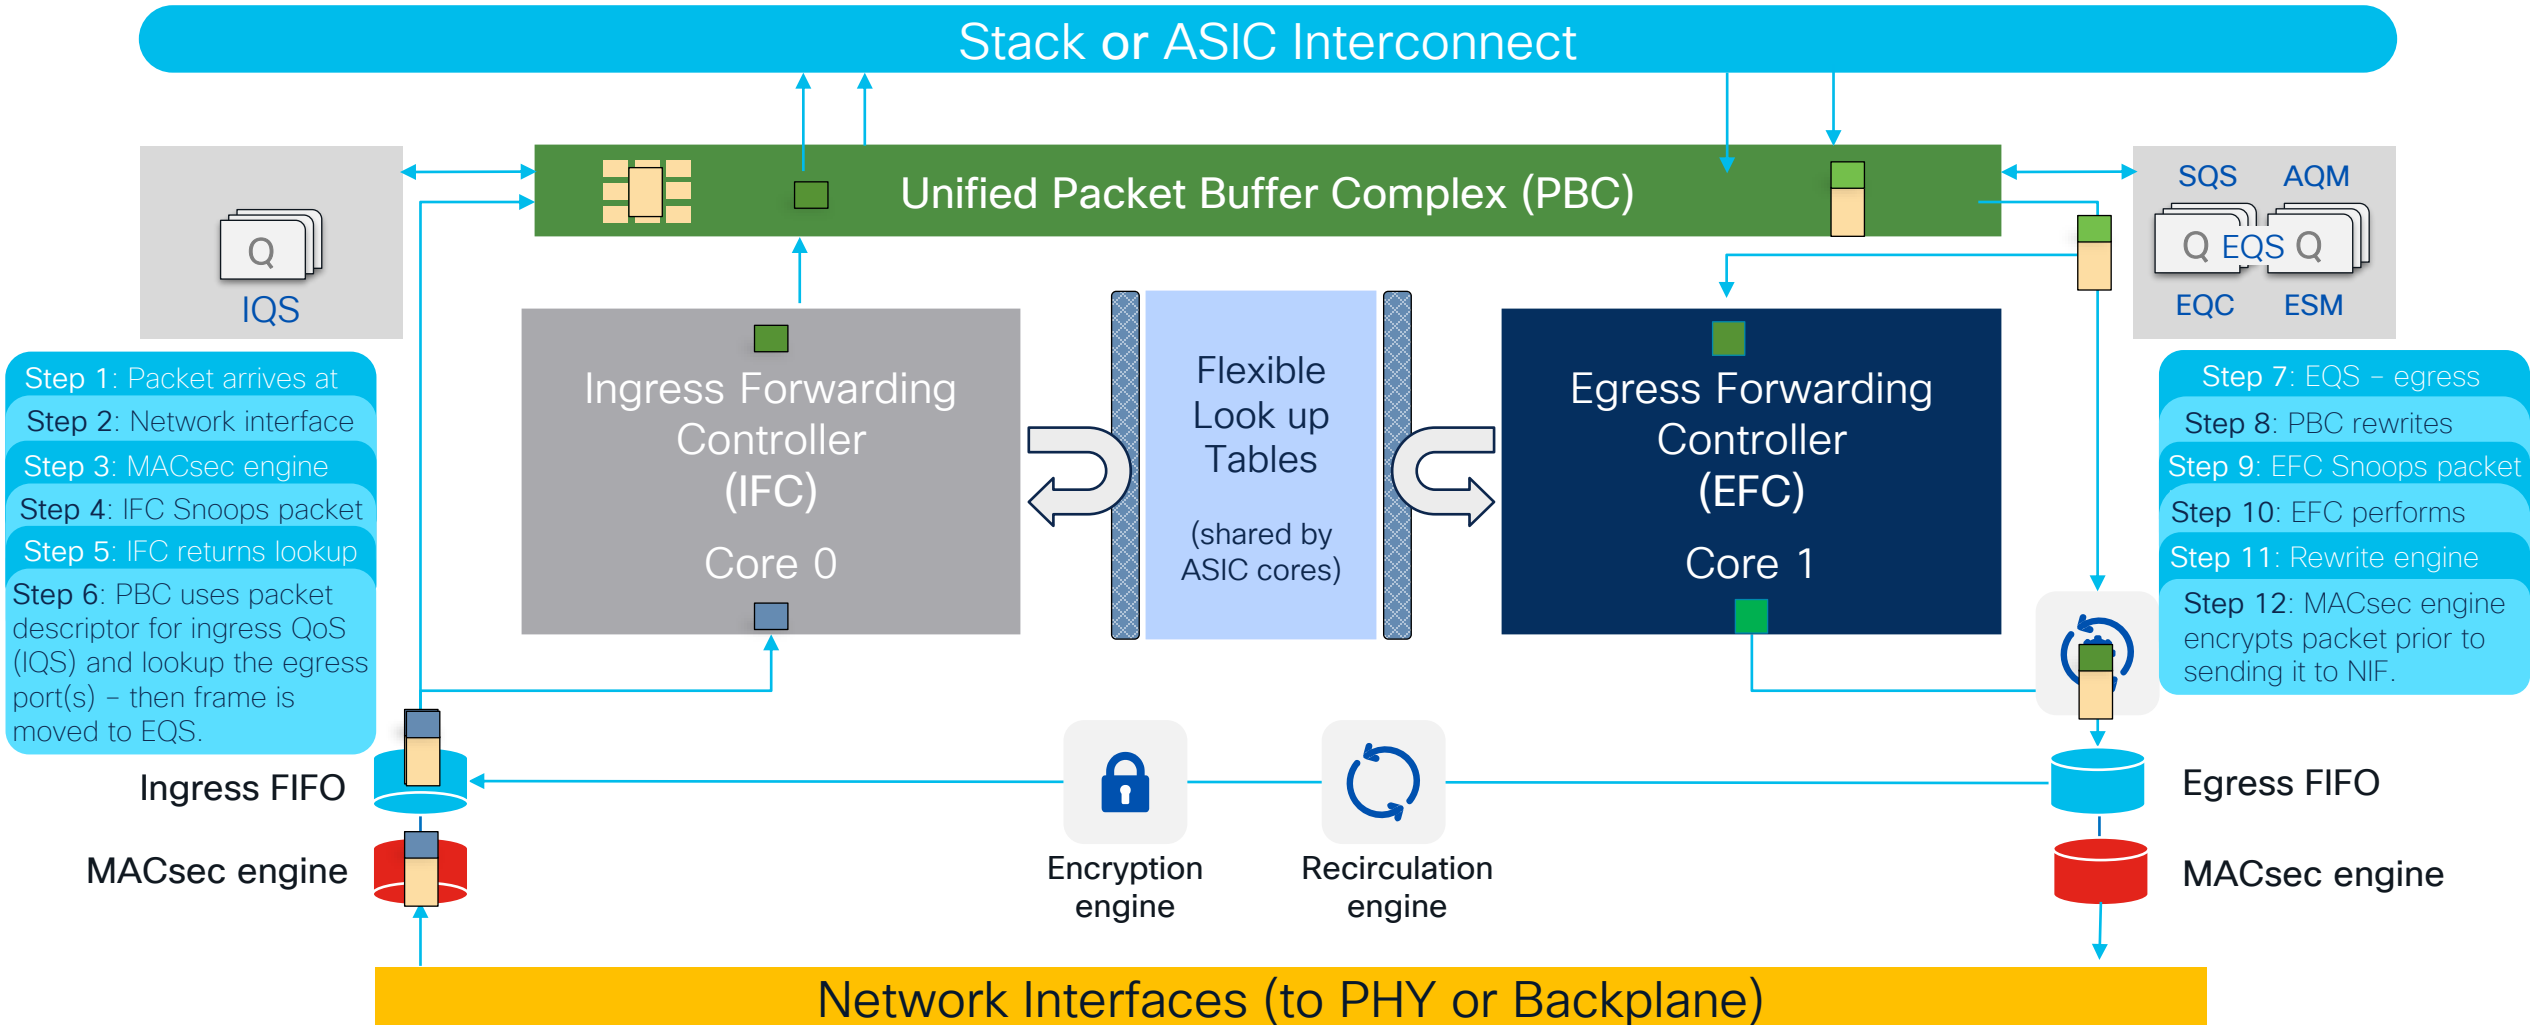


Table sizes can be tailored to support multiple templates

# Cisco UADP – Packet Walks

Generic Packet Walk – Unicast, Same ASIC

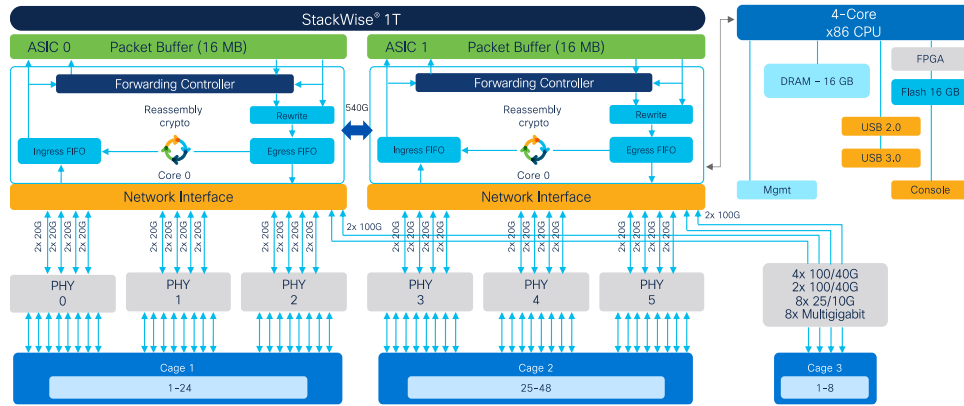


# Cisco UADP on Cisco 9000 Switches

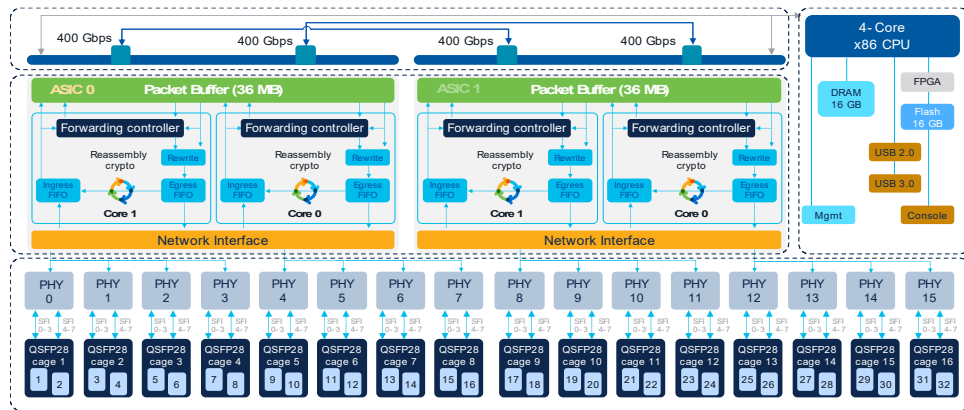
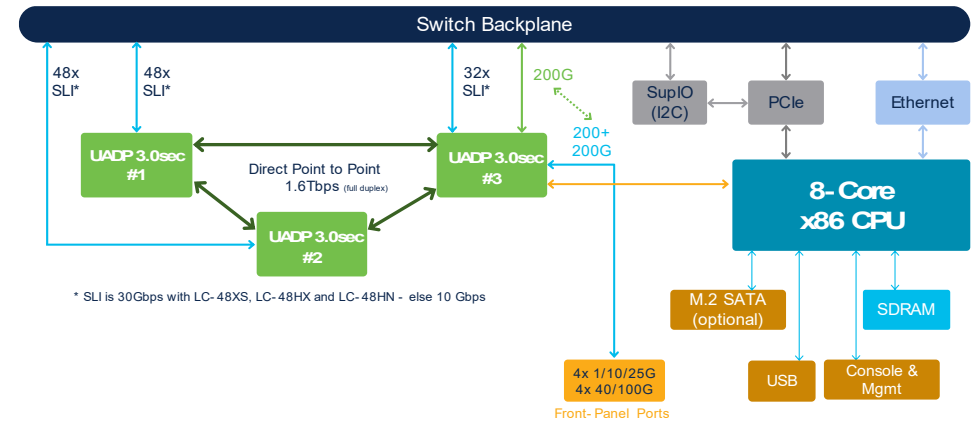
## Common ASIC Architecture



### C9300X-48HX

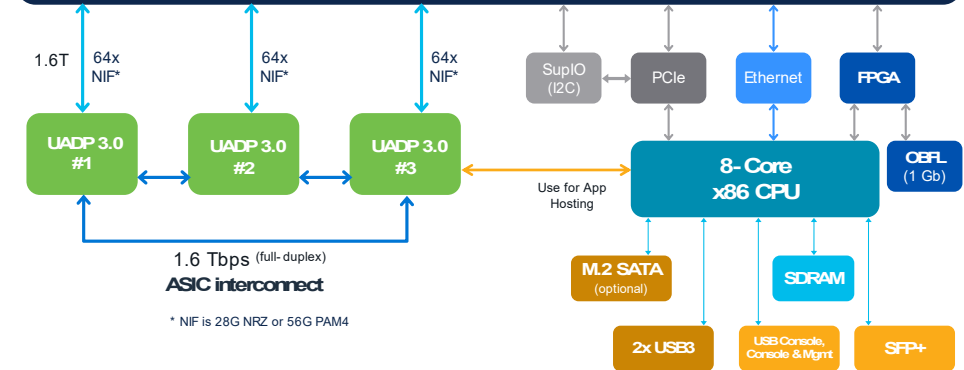


### C9400X-SUP-2/XL



### C9500-32C


### Switch backplane



### C9600-SUP-1

# Cisco Silicon One™ for Enterprise Switching

- 1 Why ASICs?
- 2 Cisco UADP
- 3 Cisco S1
- 4 Cisco 9000 Series
- 5 ASIC Futures

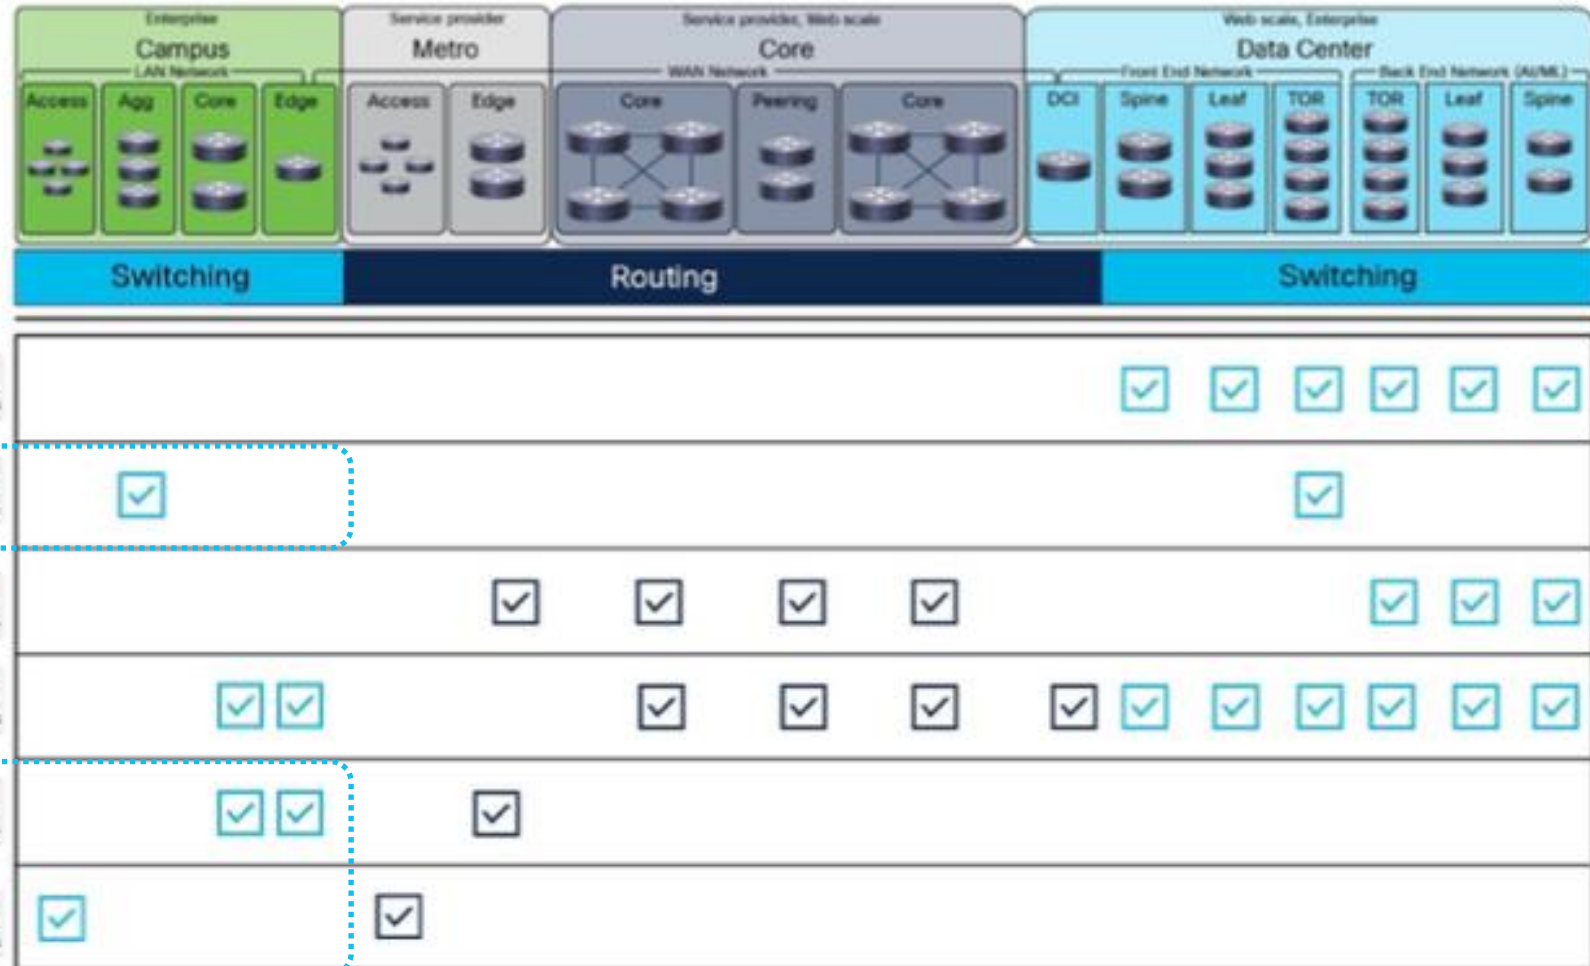


# Introducing Cisco Silicon One™

One Architecture - Multiple Devices



[www.cisco.com/c/en/us/solutions/silicon-one.html](http://www.cisco.com/c/en/us/solutions/silicon-one.html)



NEW

NEW

One Architecture, One SDK, One Experience

# Multi-Slice ASIC Design

Combining multiple Pipelines in same ASIC



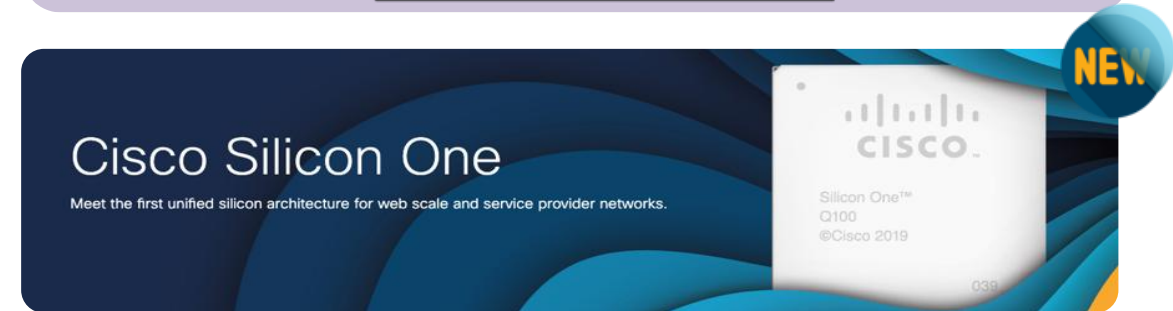
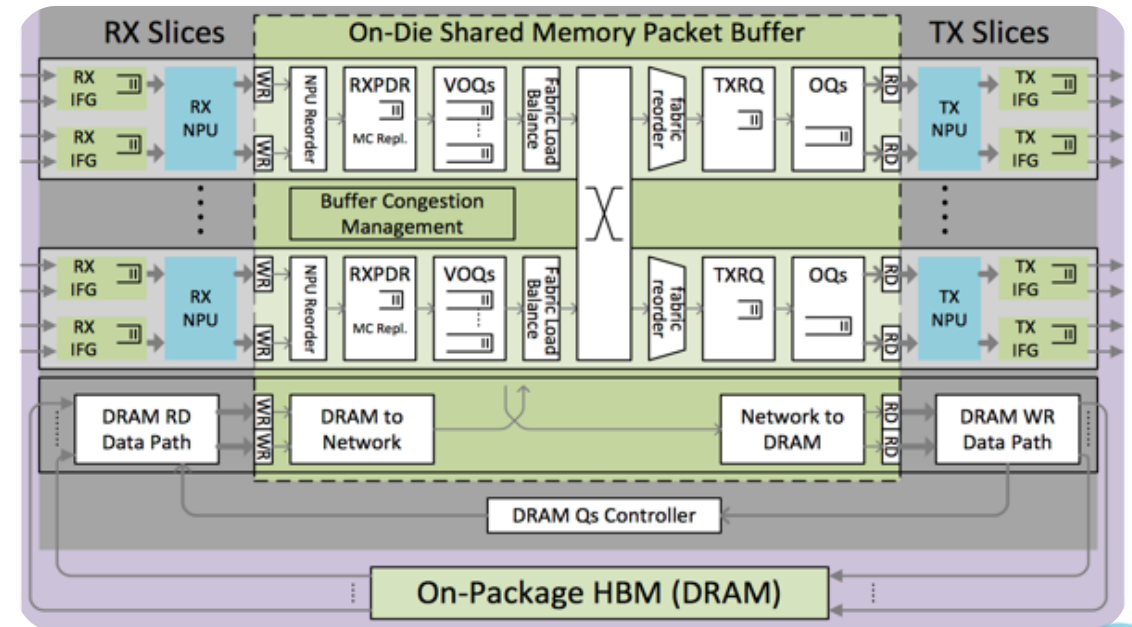
A new approach (like Multi-Core & Multi-ASIC) is to use **multiple full Network Processing Unit (NPU) pipelines** on a single ASIC package.

Each NPU pipeline (or ‘Slice’) **operates independently and connected via an integrated crossbar “fabric”**, using an integrated Virtual Output Queue (VoQ) buffer architecture, to manage traffic between Slices.

Unlike Multi-Core (only multiplies “processing”) **each Slice has dedicated RX/TX resources** for parsing, QoS, replication and other ASIC “forwarding” components.



The new Cisco Silicon One Q100 & Q200 are examples of a Multi-Slice ASIC design



[www.cisco.com/c/en/us/solutions/service-provider/innovation/silicon-one.html](http://www.cisco.com/c/en/us/solutions/service-provider/innovation/silicon-one.html)

# Cisco Silicon One™

One Architecture - Multiple Devices



[www.cisco.com/c/en/us/solutions/collateral/silicon-one/silicon-one-wp.html](http://www.cisco.com/c/en/us/solutions/collateral/silicon-one/silicon-one-wp.html)



## Q100

### 10.8 Tbps


16nm FinFET  
6 Slice SOC  
SMS + HBM\*



## Q200

### 12.8 Tbps


7nm FinFET  
6 Slice SOC  
SMS + HBM2\*



## A100

### 1.3 Tbps


16nm FinFET  
1 Slice SOC  
SMS + DDR4\*



## E100

### 6.4 Tbps

7nm FinFET  
2 Slice SOC  
SMS



## K100

### 6.4 Tbps

7nm FinFET  
2 Slice SOC  
SMS + HBM2E

- **Comprehensive routing, with switching efficiency**
- Onboard and expandable memories
- Flexible P4 NPL programmable packet processing

- **Multiple functions: system-on-chip, line-card or fabric**
- Multiple form factors: fixed or modular
- Multiple networks: Enterprise, Data Center and SP

# 10 Ways Cisco Silicon One Powers Campus Networking

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



Unified architecture



Programmable networking



Adaptable tables



Intelligent load balancing



Shared memory



Enhanced troubleshooting



Flexible Serdes




Hardware Root of Trust



Silicon One SDK

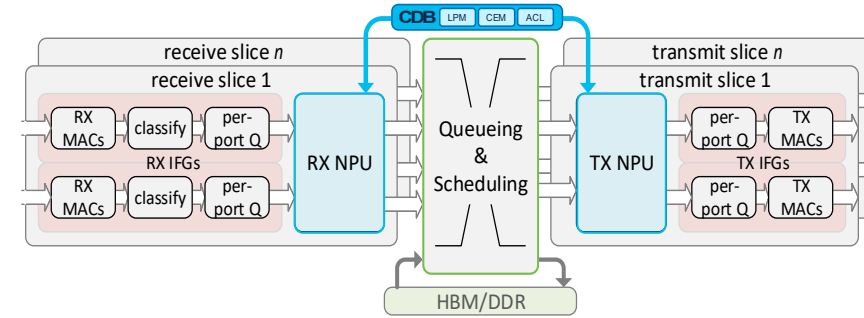


Native encryption and security



# Catalyst 9000 with S1


Introducing the next-generation of Campus ASICs



Cisco C9350 Series  
Powered by Cisco Silicon One A100



Cisco C9610-SUP-3 & 3XL  
Powered by Cisco Silicon One E100 & K100



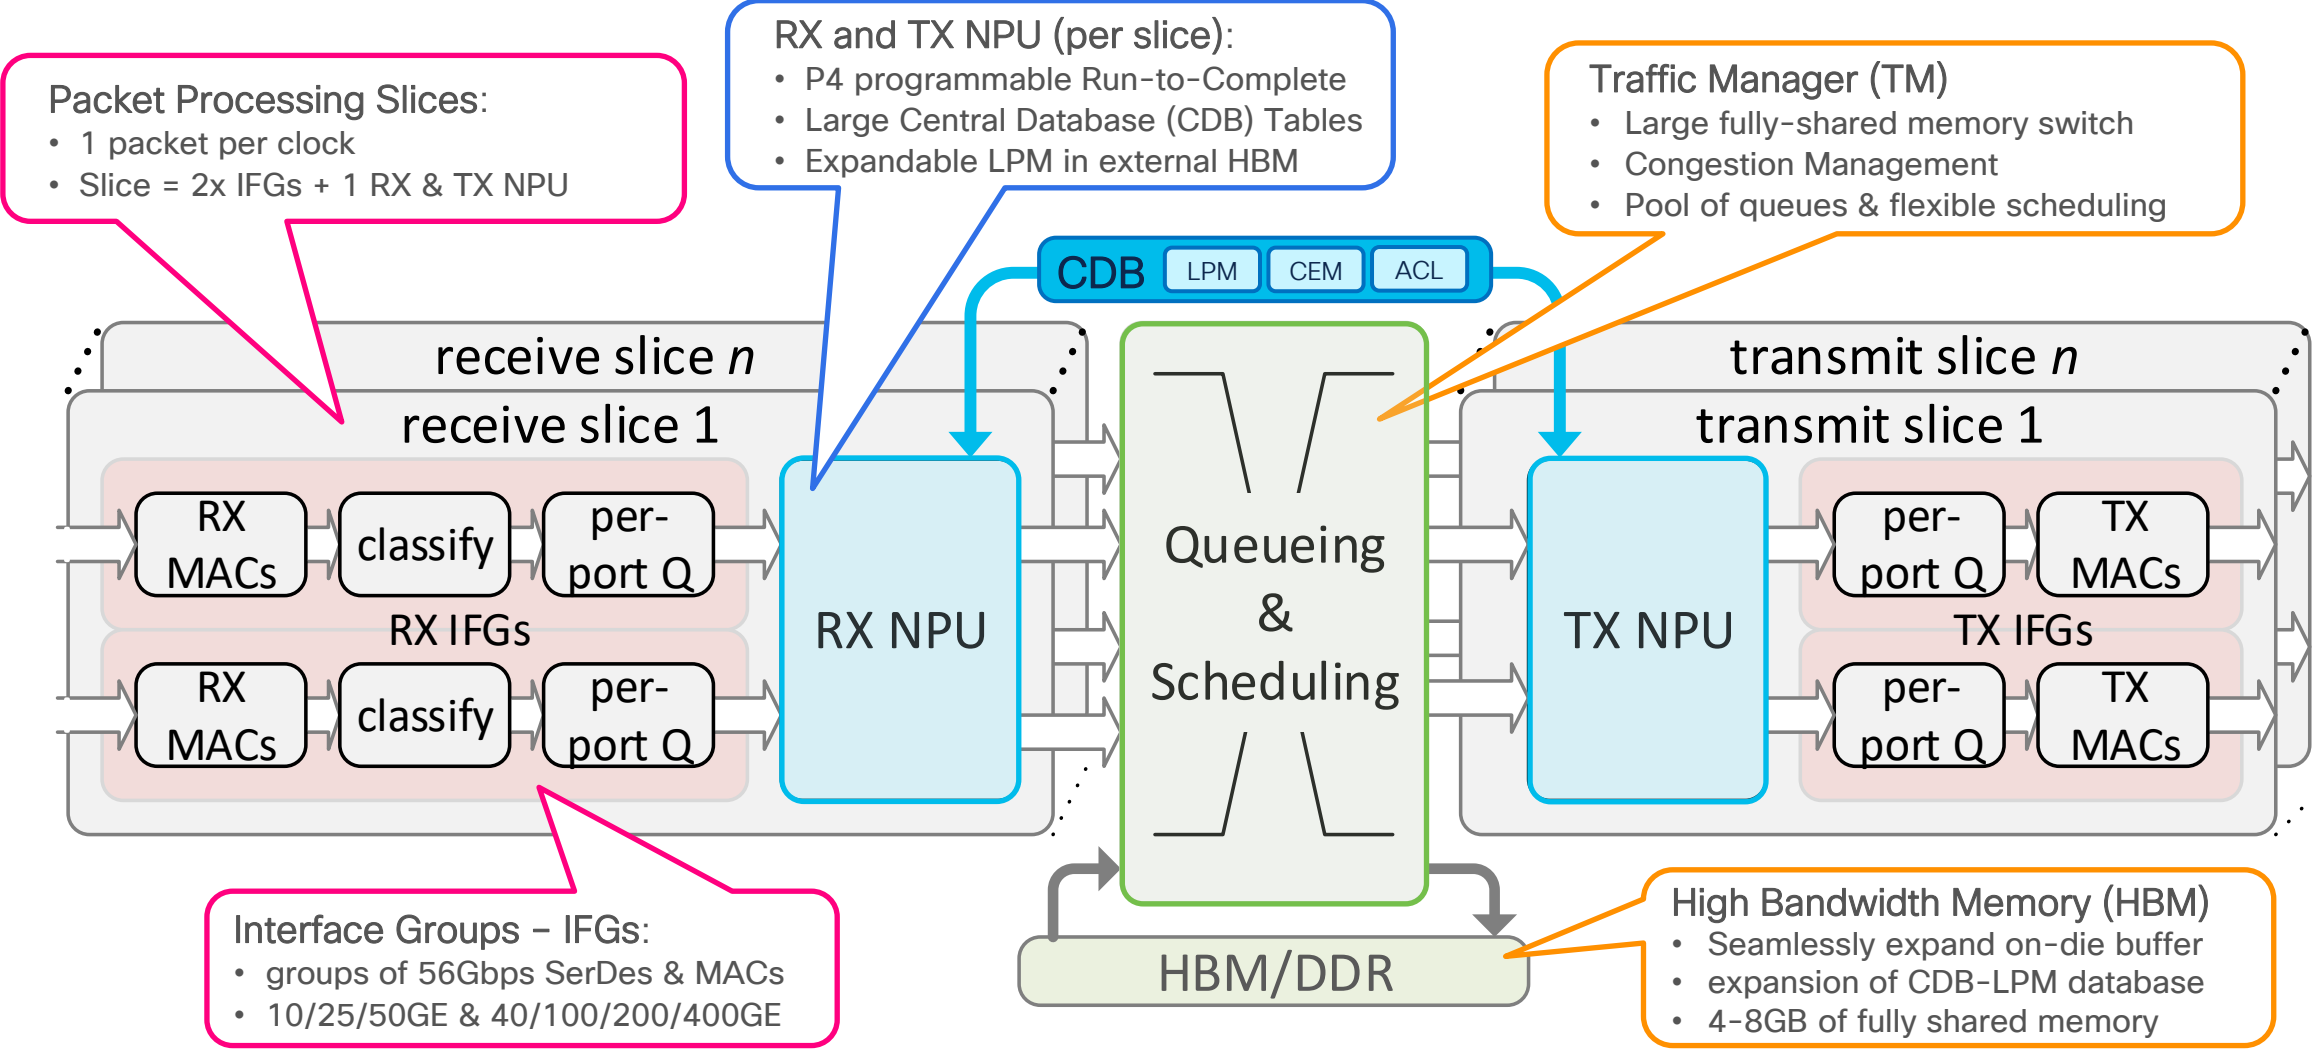
Catalyst 9500X & 9600X-SUP-2  
Powered by Cisco Silicon One Q200



# Cisco Silicon One™

## ASIC Architecture & Block Diagram

up to 4 BILLION times per second!



**Packet Processing Slices:**

- 1 packet per clock
- Slice = 2x IFGs + 1 RX & TX NPU

**RX and 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
- 10/25/50GE & 40/100/200/400GE

**High Bandwidth Memory (HBM)**

- Seamlessly expand on-die buffer
- expansion of CDB-LPM database
- 4-8GB of fully shared memory

# Cisco Silicon One™ - Central Databases

Onboard LPM, CEM & ACL memory



## S1 CDB includes the Central L2/L3 Forwarding and ACL databases:

**LPM** - SRAM database for IP/mask routing implemented by **Longest Prefix Match** algorithm

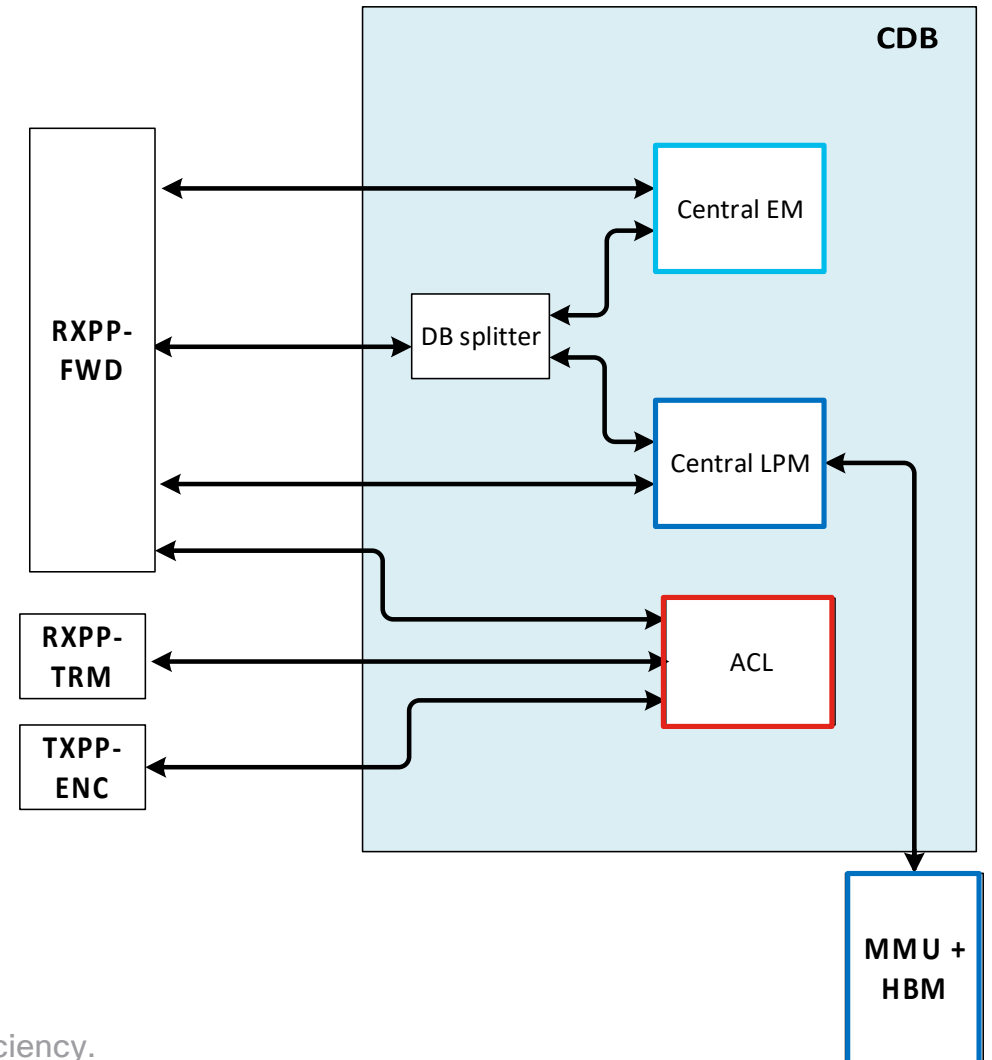
- Primarily used by IPv4 and IPv6 unicast routing
- LPM can be extended (from CDB) to HBM

**CEM** - SRAM database for MAC & Host (/48, /32 or /128), Multicast & Labels implemented by **Exact Match** algorithm

- For features using an exact match (every bit, no mask)
- CEM can be flexibly reallocated for different tables

**ACL** - TCAM classification database, contains Security, QoS and Services **Access Control List** entries

- For features that use (match criteria + action) policies
- Includes Security ACL (permit or deny)



\* Exact scale depends on IP/mask distribution (contiguous vs. random) and hash efficiency.

# Cisco S1 - Packet Processing Flow

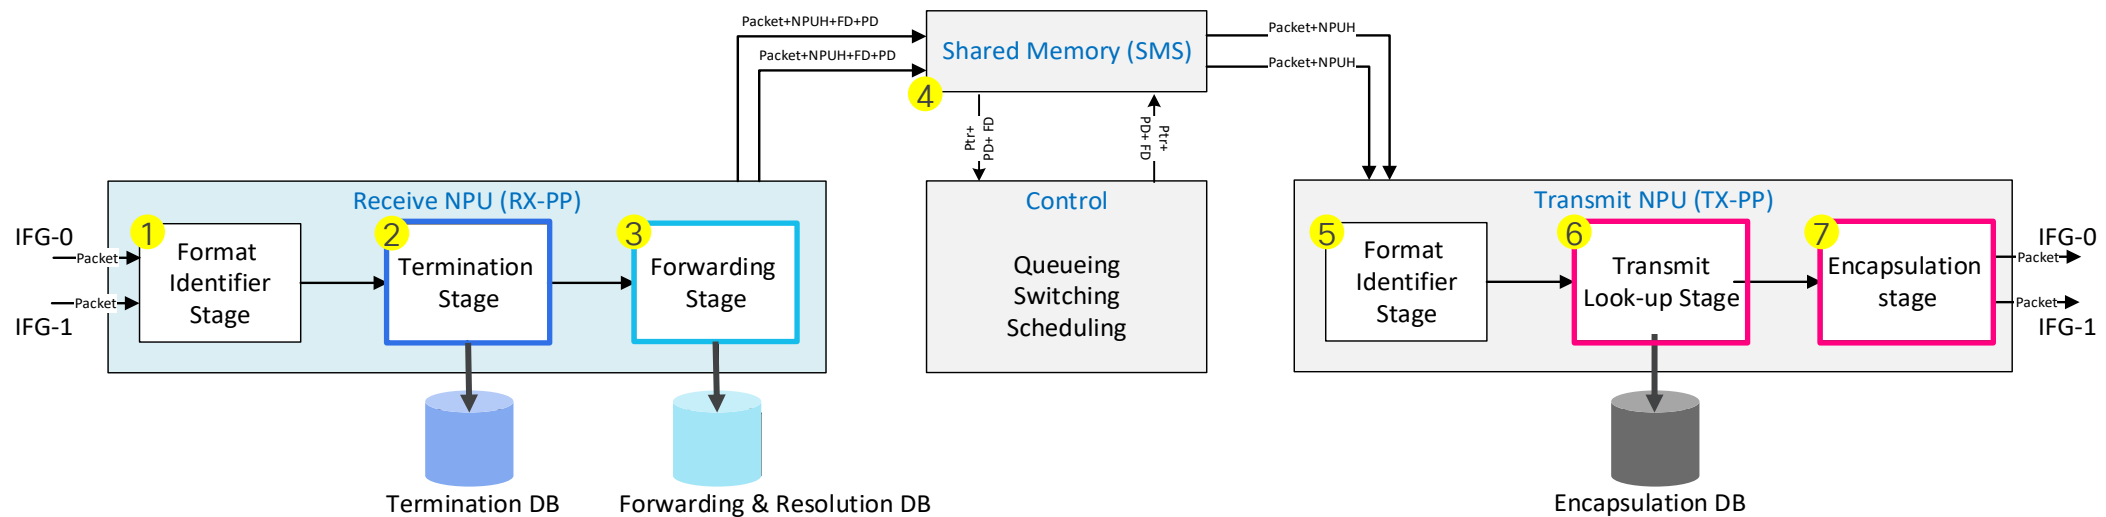
## System-wide Processing Context



S1 NPU packet processing (pipeline) has 3 main stages:

RX: 1. Termination and 2. Forwarding & Resolution and TX: 3. Transmit + Encapsulation

- Termination** mainly used for **MAC & Tunnel termination** - to obtain the **L2 Logical Port (LP)** and **L2 Relay-ID** attributes, and/or (if MAC is not terminated) to obtain the **L3 LP** and **L3 Relay-ID** attributes.
- Forwarding** **forwarding** (1<sup>st</sup> step) determines the **Destination Logical Port (DLP)** - then passes to **resolution** (2<sup>nd</sup> step) to resolve the destination logical port to the **Destination System Port (DSP)** & encap data.
- Encapsulation** determines the final **transmit encapsulation** (based on data from the forwarding & resolution stage).



# Cisco Silicon One – HBM

## High Bandwidth Memory

- **HBM achieves higher bandwidth\*** using less power, in a substantially smaller form factor:
  - **2x stacks of HBM2E memory** = ~4.8 Tbps full duplex

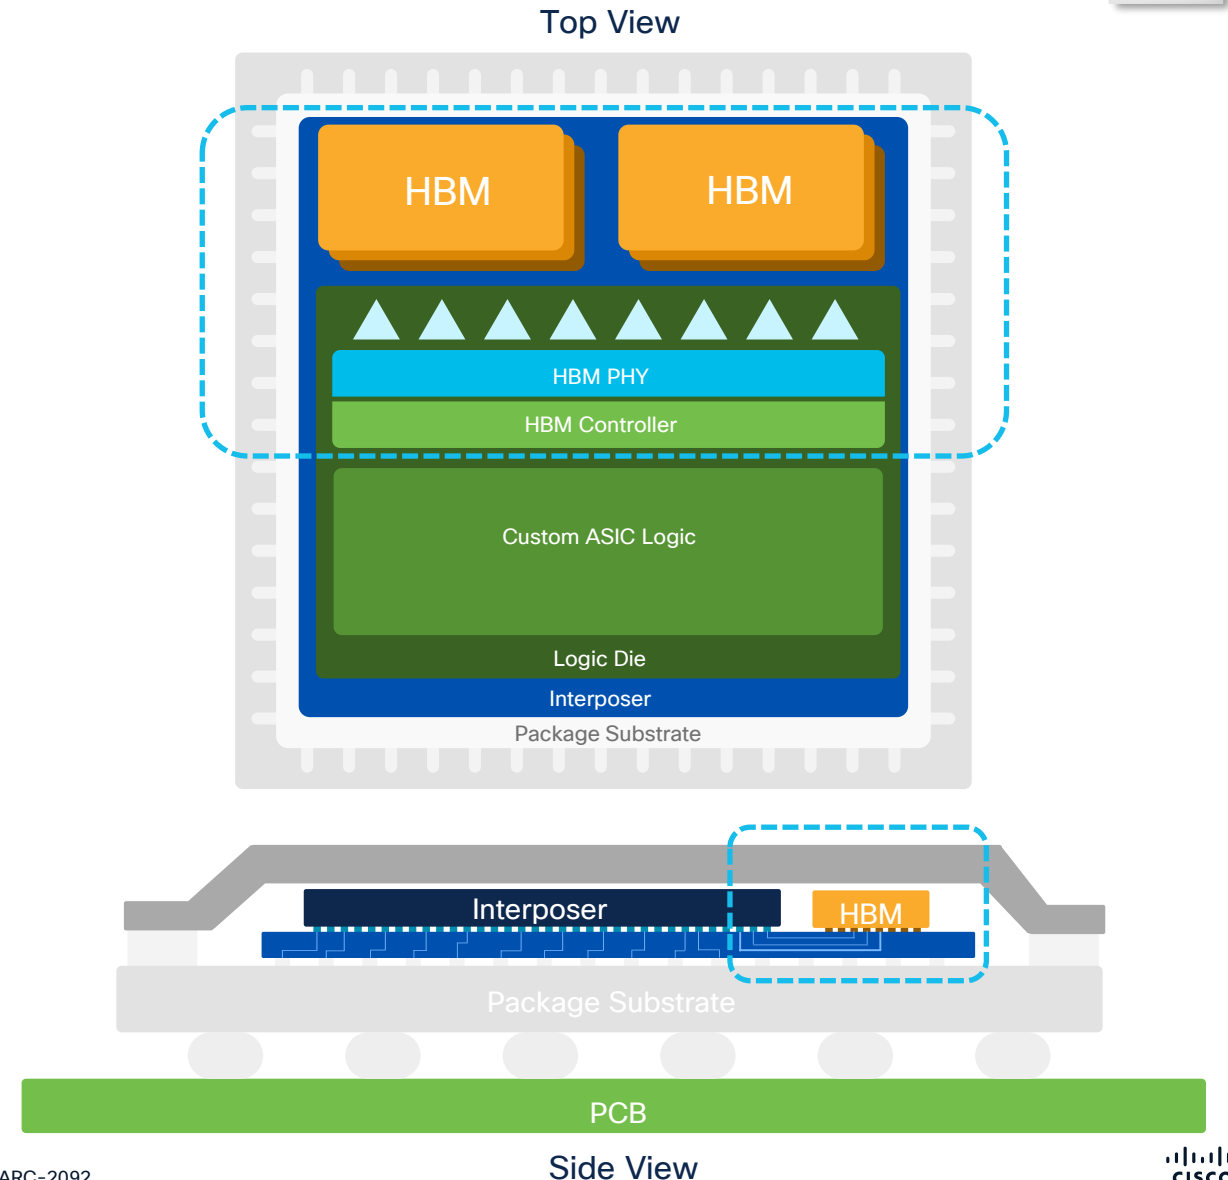
- **Augments local on-die memory**

- buffer uses local (SMS) buffers until full
- buffer uses HBM for bursts & congestion

- **Deep buffering + FIB expansion**

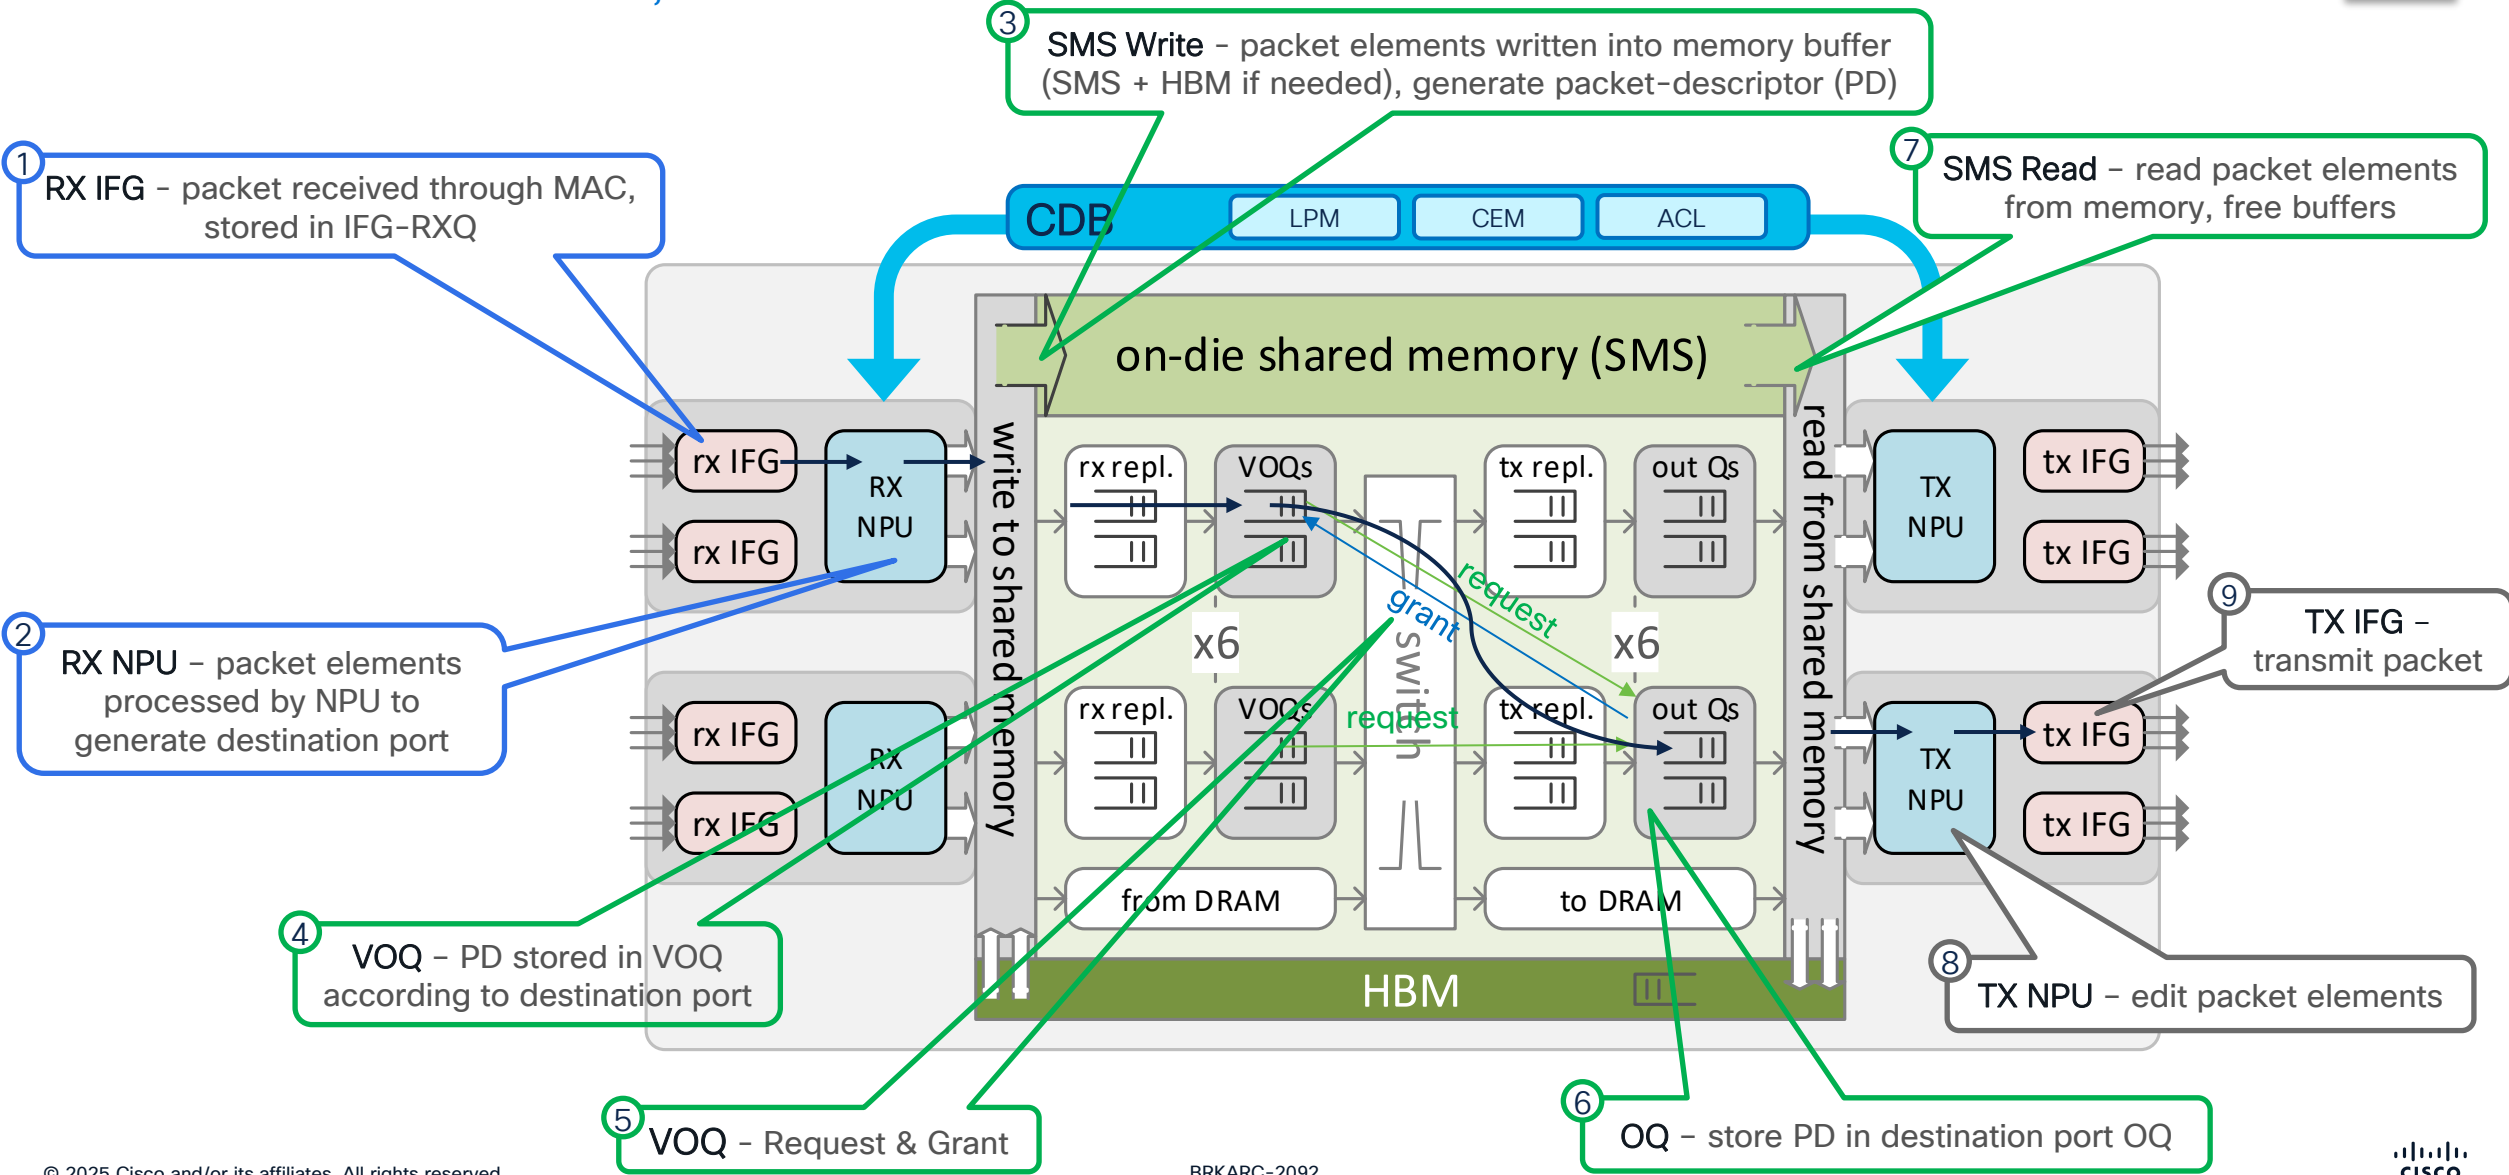
- Interposer connects ASIC die to on-package HBM memory
- local access at high-speed

[en.wikipedia.org/wiki/High\\_Bandwidth\\_Memory](https://en.wikipedia.org/wiki/High_Bandwidth_Memory)



# Cisco Silicon One - Packet Walks

Generic Packet Walk - Unicast, Same ASIC



# What's new in A100 & K100/E100?

Cisco Silicon One NPU 2.0 ASICs - Built for Enterprise Features



\* Compared to S1 Q200 (NPU 1.0)



## 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

# What is Cisco S1 HCAM?

Algorithmic Hash-based TCAM for High-Scale ACL & NetFlow



## 8K TCAM


- **Key based - variable mask**
- Supports flexible matching logic
- Complex & versatile searches
- Less-Dense (fewer bits in same area)
- Faster lookup speed
- Power & Heat intensive
- **Expensive to scale**

### TCAM



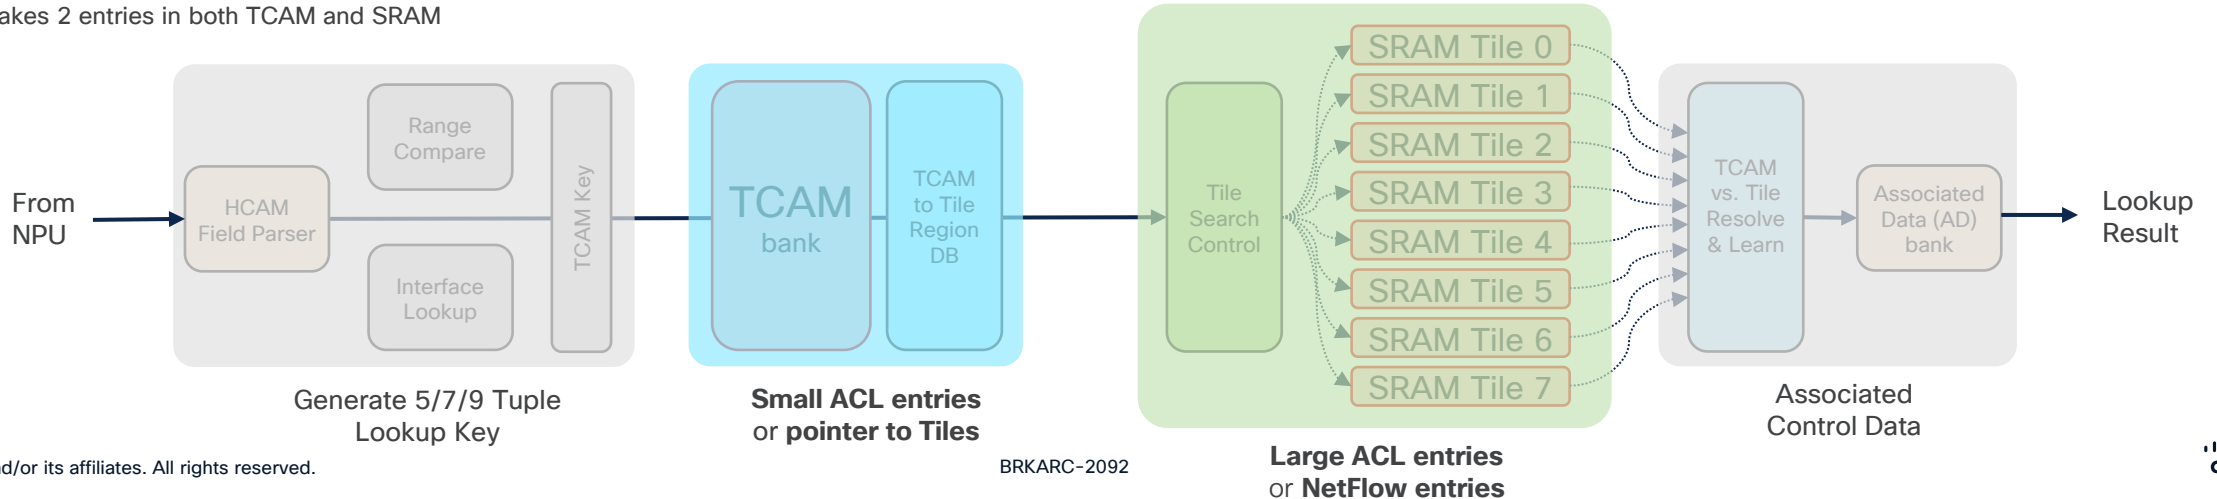
## 128K SRAM (8x Tiles of 16K)

### SRAM (Tiles)



- **Hash based - full mask**
- Simple, specific searches
- Rigid exact matching logic
- More-Dense (more bits in same area)
- Slightly slower speed
- Lower Power & Heat
- **Less costly to scale**

NOTE: IPv6 takes 2 entries in both TCAM and SRAM

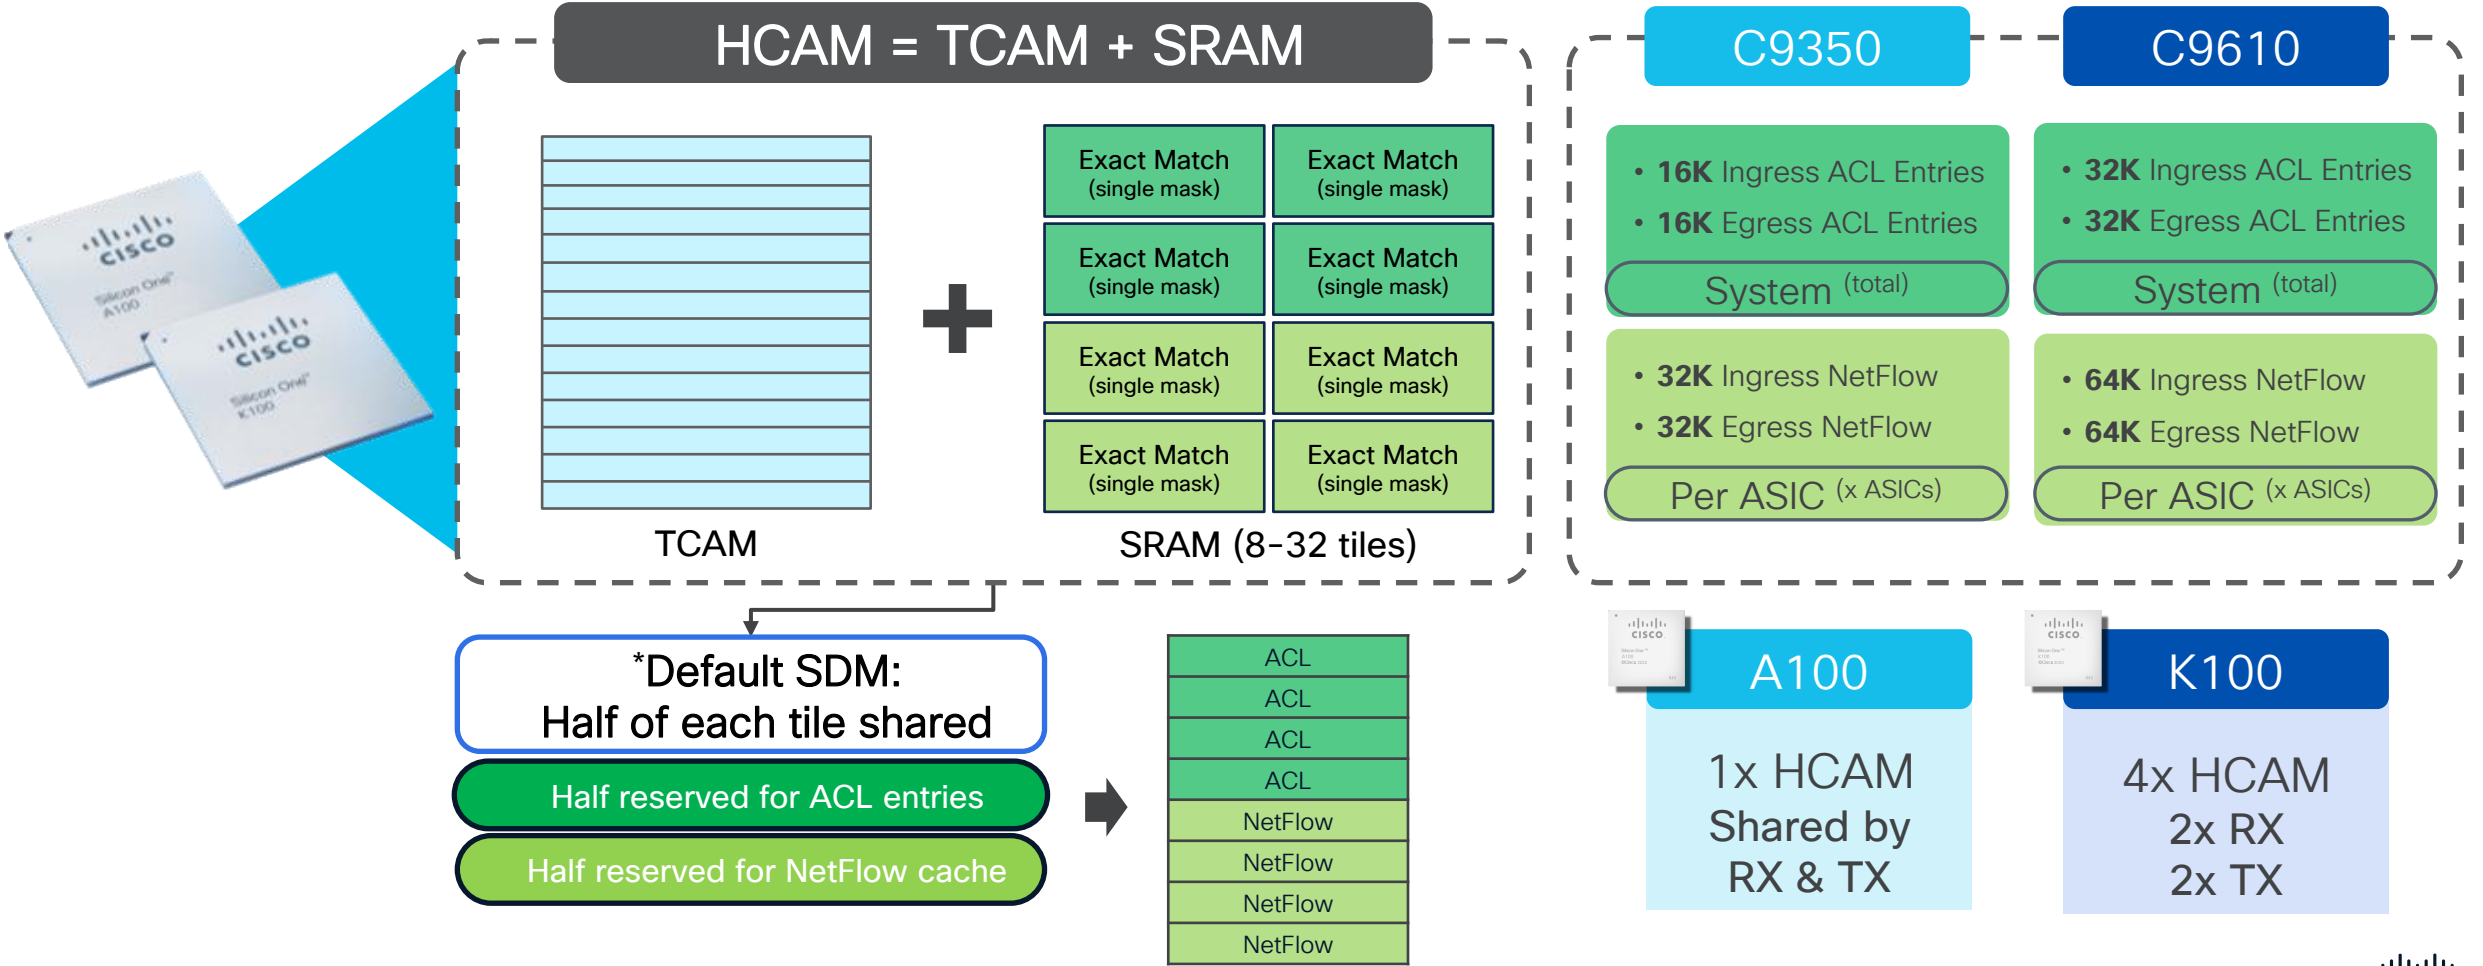


# Flexible Scale with HCAM

ACL and FNF entry space sharing



## Algorithmic TCAM - the Best of Both Worlds



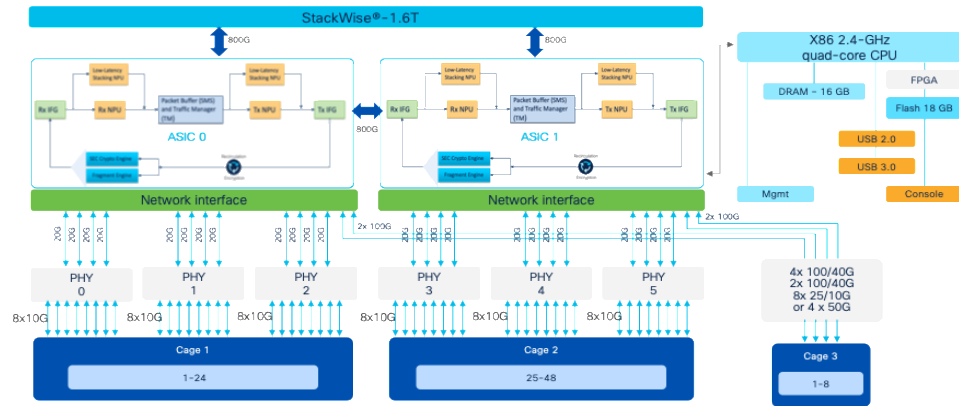
\* Based on Default SDM. ASIC capable of higher scale with Custom SDM (future software release)

# Cisco Silicon One on Cisco 9000 Switches

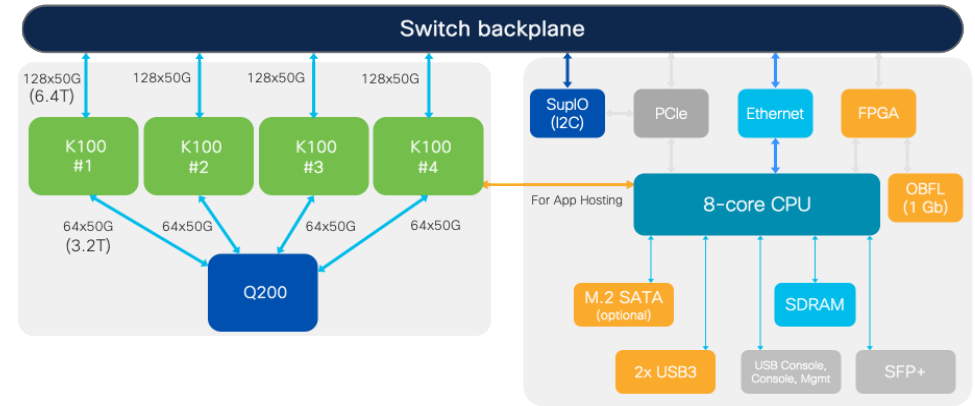
## Common ASIC Architecture



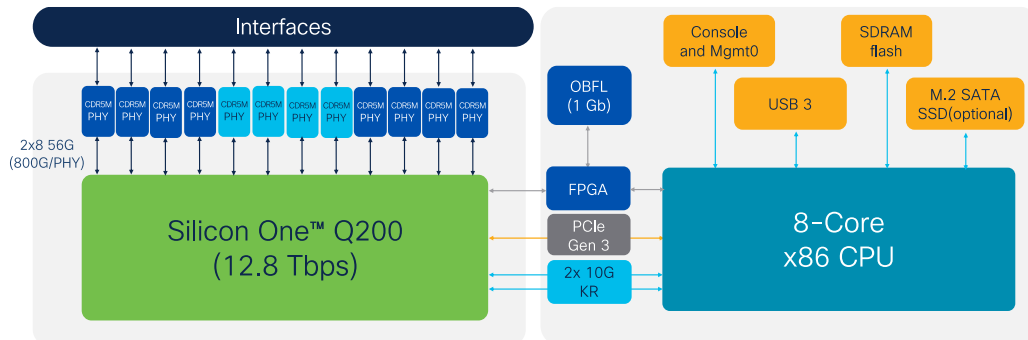
### C9350-48HX & TX



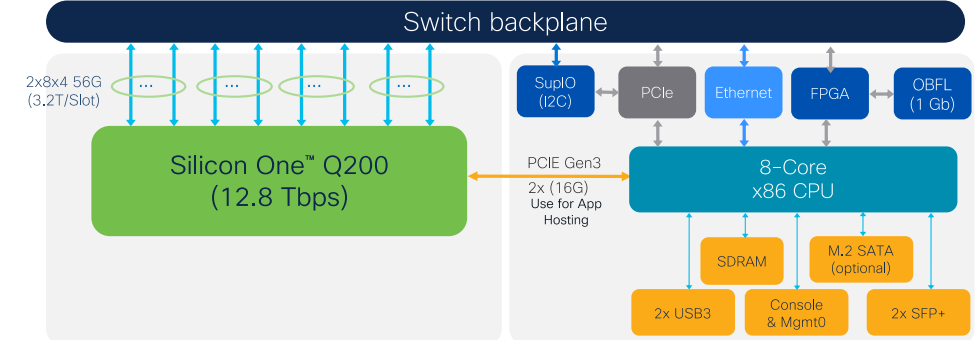
### C9610-SUP-3/XL



### C9500X-28C8D & 60L4D



### C9600X-SUP-2



# Cisco Silicon – ASIC Comparison

Contact PM/TME for latest status

 FCS+ HW & SW Committed

**NOTICE:** All scale numbers assume the ‘Default’ ASIC template

	S2 A100/L		UADP 2.0sec		UADP 2.0		UADP 1.1	
<b>IP LPM Routes</b>	256K/128K		15K/7.5K		8K/4K		8K/4K	
<b>IP Host Routes</b>	96K/48K		32K/16K		32K/16K		32K/16K	
<b>MAC Address</b>	64K		32K		32K		32K	
<b>Multicast</b> <small>L2 groups</small>	8K		8K		8K		4K	
<b>Multicast</b> <small>L3 routes</small>	8K/4K		8K/4K		8K/4K		4K/2K	
<b>MPLS</b> <small>VPN Labels</small>	32K		15K		7K		7K	
<b>NAT/PAT</b>	13K		5K		5K		--	
<b>PBR</b>	3K		3K		3K		1K	
<b>SGT/OG</b> <small>ACL Labels</small>	12K		5K		5K		3K	
<b>Security ACL</b> <small>Entries</small>	16K	16K	4K	4K	2.5K	2.5K	1.5K	1.5K
<b>QoS ACL</b> <small>Entries</small>	1K	1K	2K	2K	2K	2K	1K	1K
<b>NetFlow</b> <small>Entries</small>	32K	32K	32K	32K	32K	32K	24K	24K
<b>Policers</b>	4K		4K		4K		1K	
<b>PVST</b> <small>Instances</small>	1K		300		300		128	
<b>VLANs</b>	4K		1K		1K		1K	
<b>VRFs/VPNs</b>	256		256		256		64	

# Cisco Silicon – ASIC Comparison

Contact PM/TME for latest status

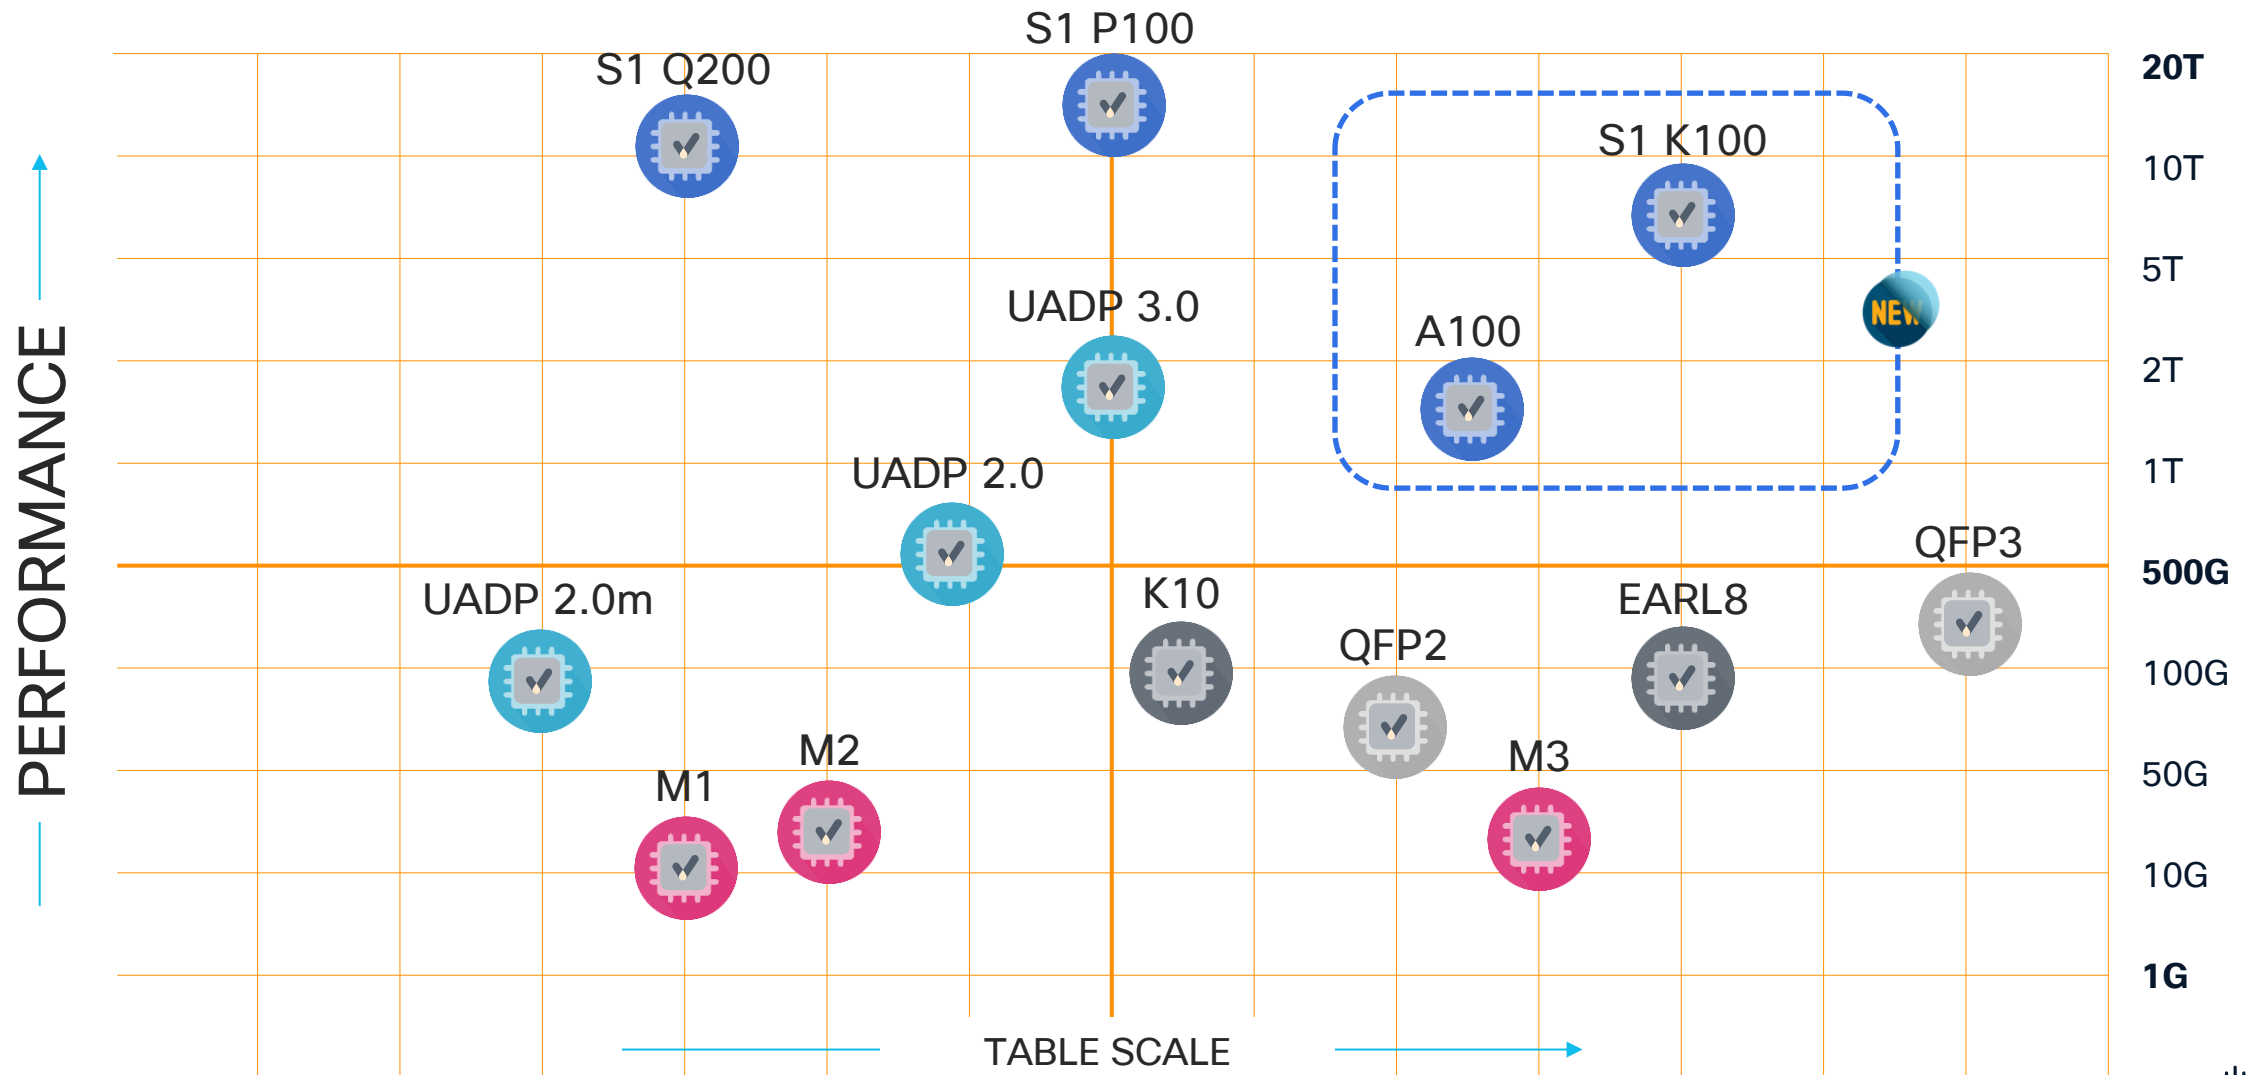
 FCS+ HW & SW Committed

**NOTICE:** All scale numbers assume the ‘Default’ ASIC template

	S1 E100		S1 K100		S1 Q200		UADP 3	
IP LPM Routes	1M / 512K		2M / 1M		1.5M / 600K		212K	
IP Host Routes	128K		128K		128K		32K	
MAC Address	128K		128K		128K		32K	
Multicast L2 groups	8K		16K		16K		16K	
Multicast L3 routes	32K		48K		32K		32K	
MPLS VPN Labels	64K		64K		64K		32K	
NAT/PAT	128K		256K		---		15K	
PBR	4K		16K		16K		3K	
SGT/OG ACL Labels	24K		24K		24K		24K	
Security ACL Entries	21K	21K	38K	38K	8K	1K	12K	15K
QoS ACL Entries	1K	512	1K	512	1K	512	8K	8K
Flexible NetFlow Entries	64K (64K x2)	64K (64K x2)	128K (64K x2)	128K (64K x2)	-- 2M (SW @ 10Mpps)		96K (32K x3)	96K (32K x3)
Policers	8K		16K		8K		4K	
PVST Instances	4K		4K		4K		1K	
VLANs	4K		4K		4K		4K	
VRFs/VPNs	4K		4K		4K		4K	

# Cisco Forwarding ASICs


Processor Performance vs. Features Landscape



# Cisco 9000 Series Smart Switches

Extending Enterprise Switching

- 1 Why ASICs?
- 2 Cisco UADP
- 3 Cisco S1
- 4 Cisco 9000 Series
- 5 ASIC Futures



# Catalyst 9200/CX, 9300/X & 9400/X



## BRKARC-2098

### Cisco 9000 Series Switching Family – Access

Manish Sharma – Technical Marketing, Cisco

This session will cover the platform overview of Cisco Catalyst 9000 Series access switches.

It will share the details of the Catalyst 9000 product portfolio, which will include new additions in **fixed and modular access series** – **Catalyst 9200, Catalyst 9300X and Catalyst 9400X**.

The session will talk about the component at the heart of these switches, which is the ASIC. It will also cover common attributes, technologies, and features in the Catalyst 9000 Series switches.

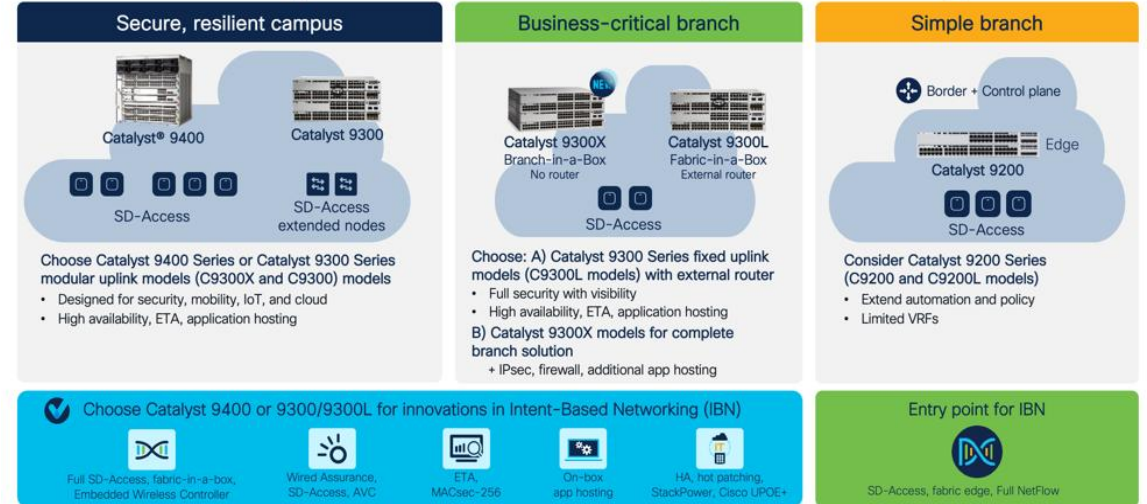


CiscoLive US 2025

Introducing **Cisco C9350 Series Smart Switches**

[www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2025/pdf/BRKARC-2098.pdf](http://www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2025/pdf/BRKARC-2098.pdf)

### Cisco Catalyst Access Switching Positioning



# Introducing C9350 Series

3rd Generation Fixed Access/Aggregation – Stacking up Industry Firsts



## MultiGigabit Downlinks



C9350-48HX



C9350-48TX

## 10/100/1000M Downlinks



C9350-48U



C9350-24U



C9350-48P



C9350-24P



C9350-48T



C9350-24T

## Gen 3 Chassis

**1.5 Tbps**  
System Capacity


**48x MGig and**  
**48x UPOE+**

**16GB DDR4**  
memory

**4 core X86 CPU**  
@2.0 Ghz

**1.6 Tbps**  
NG Stacking

**8-12\* Member**  
NG Stacking



**3x Modular**  
Fan Units

**New Stacking**  
& StackPower+

**3x Modular**  
Power Supplies



## Powered by Silicon One



Enterprise Focused  
A100/L

256K IPv4 / 128K IPv6

64K MAC

64K ARP

32K ACL

32K FNF

## Modular Uplinks



NM-4C



NM-2C



NM-8Y

# Cisco 9500/X, 9600/X



## BRKARC-2099

### Cisco 9000 Series Switching Family - Core & Distribution

Kenny Lei - Leader Technical Marketing, Cisco

This session will cover the platform overview of Catalyst 9000 Series core and distribution switches.

It will share the details of the Catalyst 9000 Series product portfolio, which will include new additions in **fixed and modular core and distribution switching series - Catalyst 9500/X and Catalyst 9600/X**.

The session will discuss the component at the heart of these switches, which is the ASIC, and it will also cover common attributes, technologies, and features in Catalyst 9000 switches.



CiscoLive US 2025

Introducing **Cisco C9610 Series Smart Switches**

[www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2025/pdf/BRKARC-2099.pdf](http://www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2025/pdf/BRKARC-2099.pdf)

### Catalyst 9000 Series Core Portfolio

	UADP 3.0	Silicon One Q200
<b>Core + Distribution</b>	<b>Catalyst 9600</b> C9600-SUP-1 C9600-LC-24C C9600-LC-48YL	<b>Catalyst 9500</b> C9500-32C / C9500-32QC C9500-48Y4C / C9500-24Y4C
<b>Core + Campus Edge</b>	<b>Catalyst 9600X</b> C9600X-SUP-2 C9600X-LC-32CD C9600X-LC-56YL4C	<b>Catalyst 9500X</b> C9500X-28C8D C9500X-60L4D
	<b>Total capacity 4.8 Tbps</b> Slot bandwidth 1.2 Tbps	<b>Total capacity 12.8 Tbps</b> Slot bandwidth 3.2 Tbps
	<b>Highest capacity 3.2 Tbps</b>	<b>Highest capacity 6 Tbps</b>

Also check out **BRKARC-2668**

# Introducing C9610 Series

3<sup>rd</sup> Generation Modular Core – with maximum port density



## Gen 3 Supervisors



C9610-SUP-3/3XL



Enterprise Focus  
K100/E100



2M IPv4 / 1M IPv6

3.2 Tbps per slot


512K MAC & ARP\*

32GB DDR4  
memory

256K ACL & FNF\*

8 core X86 CPU  
@2.0 Ghz

\*Hardware Capable



C9610R

## Gen 3 Chassis

51.2 Tbps  
System Bandwidth

8-line card slots  
(1.25 RU)

2 supervisor slots  
(2.5 RU)

4 fan trays  
serviceable in rear

Blue Beacons  
(Chassis, Fans, Sups,  
& Line-Cards)

8 power supplies  
AC/DC


## Gen 3 Line cards



C9610-LC-40YL4CD




C9610-LC-32CD




4x FAN Trays  
with Front to Back  
Airflow

# Glimpse into the Future

ASIC Innovations for Enterprise Switching



- 1 Why ASICs?
- 2 Cisco UADP
- 3 Cisco S1
- 4 Cisco 9000 Series
- 5 ASIC Futures



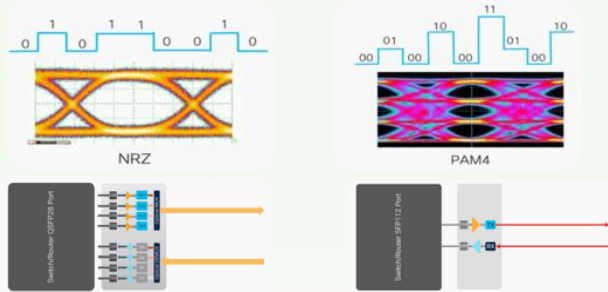
# Where are things going?

## Speeds and Feeds

More details in PPT Notes



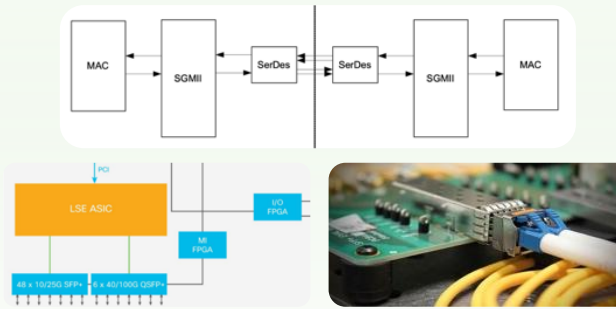
### Single Lambda



Uses a **single 112G PAM4** signal vs. 4 x 28G NRZ signals



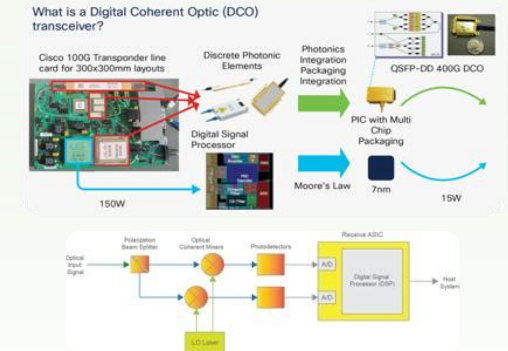
### Direct-Drive (PHYLESS)



Transceiver uses **SFI/SGMII** to connect directly to the **ASIC SerDes** (no PHY)



### Coherent Optics



Uses **phase of optical signal** vs. power level



### 25/50G → 100G SFP


**NEW**  
SFP-DD/SFP-DD112/SFP112 Hardware Specification for SFP112 and SFP Double Density Pluggable Transceiver Revision 5.1



**SFP-DD**

SFP-DD = Two SFPs in One

- SFP-DD Specification released on September 14, 2017 at ECOC
- SFP-DD packs two SFPs in one cage
- Similar to how QSFP-DD is backward compatible to QSFP, SFP-DD supports SFPs - VERY IMPORTANT




**SFP-DD(112): 100Gbit/s using PAM4**



### 100/400G → 800G QSFP

Introducing 800G Pluggable Transceivers

For next-generation high-speed connectivity




**OSFP-DD800**



**QSFP-DD(112): 800Gbit/s using PAM4**



### 4x100G & 8x50G Breakout

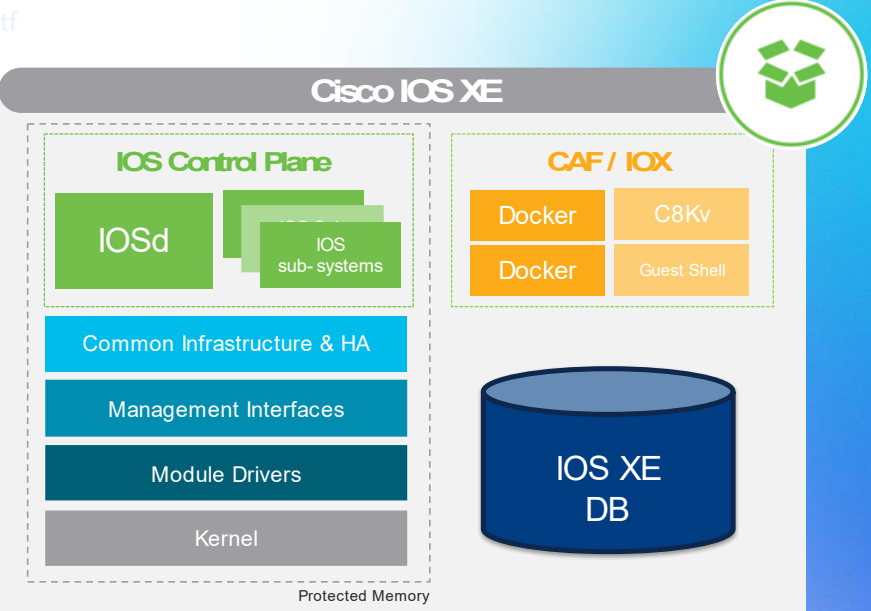


**QSFP-DD (400G) can be split into 4 x 100G QSFP28 or 8 x 50G SFP56**

```
extract_num-
ber_and_incr (destination, source) int
*destination; unsigned char **source; { extract_num-
ber (destination, *source); *source += 2; } #ifndef EXTRACT_MAC-
ROS #undef EXTRACT_NUMBER_AND_INCR #define EXTRACT_NUM-
BER_AND_INCR(dest, src) \ extract_number_and_incr (&dest, &src) #endif /*
not EXTRACT_MACROS */ #endif /* DEBUG */ /* If DEBUG is defined, Regex prints
many voluminous messages about what it is doing (if the variable `debug' is nonzero). If
linked with the main program in `iregex.c'; you can enter patterns and strings interactively.
And if linked with the main program in `main.c' and the other test files, you can run the al-
ready-written tests. */ #ifdef DEBUG /* We use standard I/O for debugging. */ #include <stdio.h>
/* It is useful to test things that `must' be true when debugging. */ #include <assert.h> static int
debug = 0; #define DEBUG_STATEMENT(e) e #define DEBUG_PRINT1(x) if (debug) printf (x) #define
DEBUG_PRINT2(x1, x2) if (debug) printf (x1, x2) #define DEBUG_PRINT3(x1, x2, x3) if (debug) printf
(x1, x2, x3) #define DEBUG_PRINT4(x1, x2, x3, x4) if (debug) printf (x1, x2, x3, x4) #define DE-
BUG_PRINT_COMPILED_PATTERN(p, s, e) \ if (debug) print_partial_compiled_pattern (s, e) #defin-
BUG_PRINT_DOUBLE_STRING(w, s1, sz1, s2, sz2) \ if (debug) print_double_string (w, s1, sz1, s2, sz2)
extern void printchar(); /* Print the fastmap in human-readable form. */ void print_fastmap (fastmap
print_fastmap; (fastmap & was_a_range) { if (fastmap[i] & was_a_range = 1; i++;) if
(was_a_range) { printchar (i - 1); } putchar ('\n'); } /* Print a compiled pattern string in hu-
man-readable form, starting at the START pointer into it and ending just before the pointer END. */ void
print_compiled_pattern_string (s, end) const char *s; const char *end; { int mcnt; un-
signed char *p; while (p < end) { switch ((re_opcode_t) *p++) { case no_op: printf ("/no_op");
break; case charset: case charset_not: { register int c; printf ("/charset%s", (re_opcode_t) *(p -
1) == charset_not ? "_not" : ""); assert (p + *p < end); for (c = 0; c < *p; c++) { unsigned bit;
unsigned char map_byte = p[1 + c]; putchar ('/'); for (bit = 0; bit < BYTEWIDTH; bit++) if
(map_byte & (1 << bit)) putchar (c * BYTEWIDTH + bit); } p += 1 + *p; break; } case beg-
line: printf ("/begline"); break; case endline: printf ("/endline"); break; case on_failure_-
jump: extract_number_and_incr (&mcnt, &p); printf ("/on_failure_jump/0/%d", mcnt);
break; case on_failure_keep_string_jump: extract_number_and_incr (&mcnt, &p); printf
("/on_failure_keep_string_jump/0/%d", mcnt); break; case dummy_failure_jump: ex-
tract_number_and_incr (&mcnt, &p); printf ("/dummy_failure_jump/0/%d", mcnt); break;
case push_dummy_failure: printf ("/push_dummy_failure"); break; case may-
be_pop_jump: extract_number_and_incr (&mcnt, &p); printf
("/maybe_pop_jump/0/%d", mcnt); break; case pop_failure_-
jump: extract_number_and_incr (&mcnt, &p); printf ("/pop_-
failure_jump/0/%d", mcnt); break; case jump_past_alt:
extract_number_and_incr (&mcnt, &p); printf ("/-
```

# Cisco IOS XE Software Architecture & Innovations

## Cisco C9000 Series



# Agenda

## 01 Brief History of IOS XE

## 02 Basic IOS XE Components

## 03 IOS XE Technologies

## 04 C9K IOS XE upto 17.15.1

## 05 C9K IOS XE after 17.18.1

## 06 Summary & Closing

# Catalyst 9000 Series – Common Building Blocks



## Multi-Core x86 CPU

Application Hosting  
Secure Containers



## Cisco IOS XE<sup>®</sup>


Model-Driven APIs  
Modular Patching



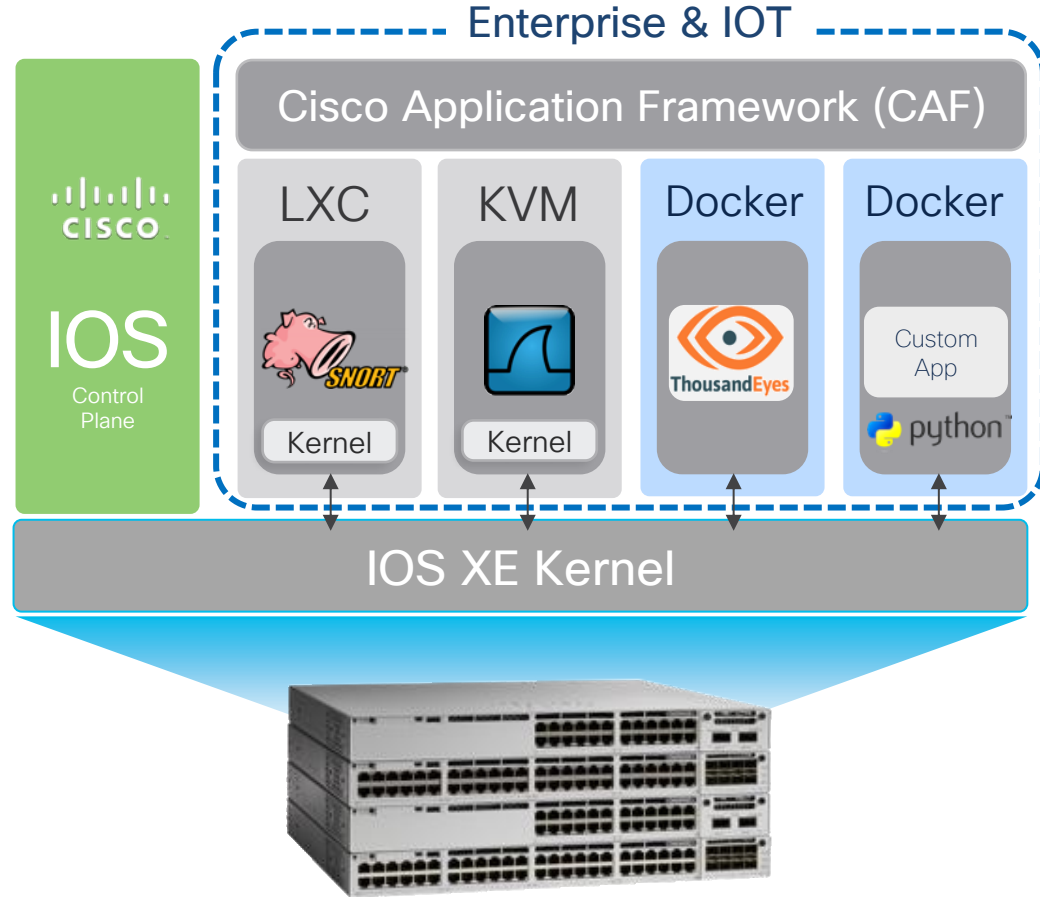

## Cisco UADP & Silicon One<sup>™</sup>

Programmable Pipeline  
Flexible Tables

# Multi-Core CPU – Built for App Hosting

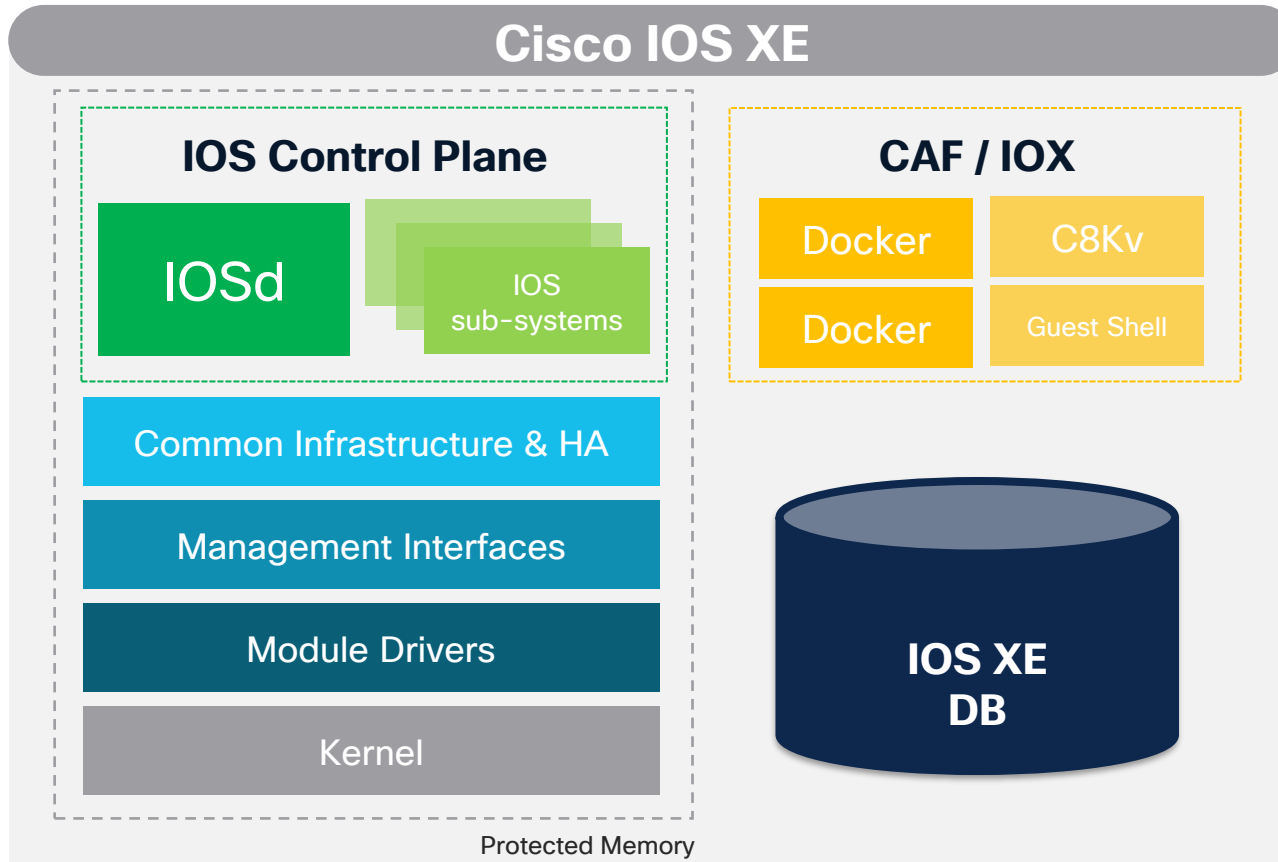


**x86**  
**CPU**



x86 CPU enables hosting NFV devices, Containers and 3<sup>rd</sup>-party Apps

# Cisco IOS XE – A Modern Operating System



 **Cisco IOS subsystems**  
Resiliency and High Availability

 **Cisco IOS XE database**  
Programmability and Open models


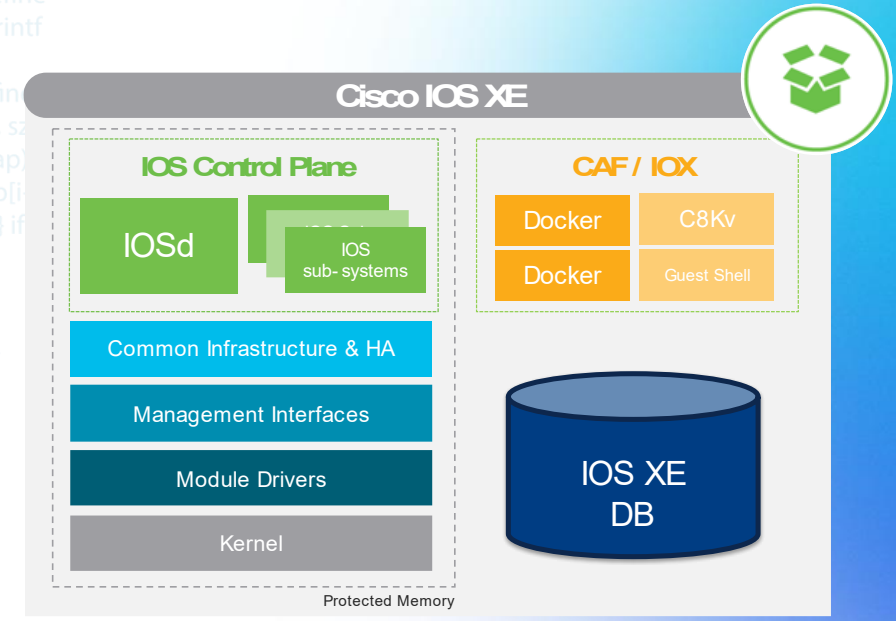
 **IOX + Docker containers**  
Cisco and 3rd-party App hosting

Open, Model Driven & Secure Operating System – Over 3000 Features

```
extract_number_and_incr (destination, source) int
*destination; unsigned char **source; { extract_number
(destination, *source); *source += 2; } #ifndef EXTRACT_MAC
ROS #undef EXTRACT_NUMBER_AND_INCR #define EXTRACT_NUM
BER_AND_INCR(dest, src) \ extract_number_and_incr (&dest, &src) #endif /*
not EXTRACT_MACROS */ #endif /* DEBUG */ /* If DEBUG is defined, Regex prints
many voluminous messages about what it is doing (if the variable `debug' is nonzero). If
linked with the main program in `iregex.c'; you can enter patterns and strings interactively.
And if linked with the main program in `main.c' and the other test files, you can run the al-
ready-written tests. */ #ifdef DEBUG /* We use standard I/O for debugging. */ #include <stdio.h>
/* It is useful to test things that `must' be true when debugging. */ #include <assert.h> static int
debug = 0; #define DEBUG_STATEMENT(e) e #define DEBUG_PRINT1(x) if (debug) printf (x) #define
DEBUG_PRINT2(x1, x2) if (debug) printf (x1, x2) #define DEBUG_PRINT3(x1, x2, x3) if (debug) printf
(x1, x2, x3) #define DEBUG_PRINT4(x1, x2, x3, x4) if (debug) printf (x1, x2, x3, x4) #define DE
BUG_PRINT5(x1, x2, x3, x4, x5) if (debug) print_partial_compiled_pattern (s, e) #defin
DEBUG_PRINT6(x1, x2, x3, x4, x5, x6) \ if (debug) print_double_string (w, s1, sz1, s2, sz2)
extern void printchar(); /* Print the fastmap in human-readable form. */ void print_fastmap (fastmap)
char *fastmap; { unsigned was_a_range = 0; unsigned i = 0; while (i < (1 << BYTEWIDTH)) { if (fastmap[i]
was_a_range = 1; i++; } if (fastmap[i]
a compiled pattern string in hu-
man-readable form, starting at the START pointer into it and ending just before the pointer END. */ void
print_partial_compiled_pattern (start, end) unsigned char *start; unsigned char *end; { int mcnt, mcnt2; un-
signed char *p; while (p < end) { if (start == p) printf ("(null)\n"); return; } /* Loop over
opcodes. */ while (p < end) { case no_op: printf ("/no_op");
break; case exactn: mcnt = *p++; printf ("/exactn/%d", mcnt); do { putchar ('/'); printf (*p++);
} while (--mcnt); break; case start_memory: mcnt = *p++; printf ("/start_memory/%d/%d", mcnt,
*p++); break; case anychar: printf ("/anychar");
break; case charset: case charset_not: { register int c; printf ("/charset%s", (re_opcode_t) *(p -
1) == charset_not ? "_not" : ""); assert (p + *p < pend); for (c = 0; c < *p; c++) { unsigned bit;
unsigned char map_byte = p[1 + c]; putchar ('/'); for (bit = 0; bit < BYTEWIDTH; bit++) if
(map_byte & (1 << bit)) printf (c * BYTEWIDTH + bit); } p += 1 + *p; break; } case beg-
line: printf ("/begline"); break; case endline: printf ("/endline"); break; case on_failure_-
jump: extract_number_and_incr (&mcnt, &p); printf ("/on_failure_jump/0/%d", mcnt);
break; case on_failure_string: extract_number_and_incr (&mcnt, &p); printf
("/on_failure_string/0/%d", mcnt); break; case dummy_failure_jump: ex-
tract_number_and_incr (&mcnt, &p); printf ("/dummy_failure_jump/0/%d", mcnt); break; case may-
be_pop_jump: extract_number_and_incr (&mcnt, &p); printf
("/maybe_pop_jump/0/%d", mcnt); break; case pop_failure_-
jump: extract_number_and_incr (&mcnt, &p); printf ("/pop_-
failure_jump/0/%d", mcnt); break; case jump_past_alt:
extract_number_and_incr (&mcnt, &p); printf ("/-
```

1	2	3	4	5	6
Why IOS XE?	IOS XE Design	Up to 17.15.x	After 17.18.x	IOS XE Features	Summary


- History of Cisco IOS®
- Cisco IOS evolved into IOS XE
- Nova IOS XE (Catalyst 3K)
- Polaris IOS XE (Catalyst 9K)



# Brief History of Cisco IOS



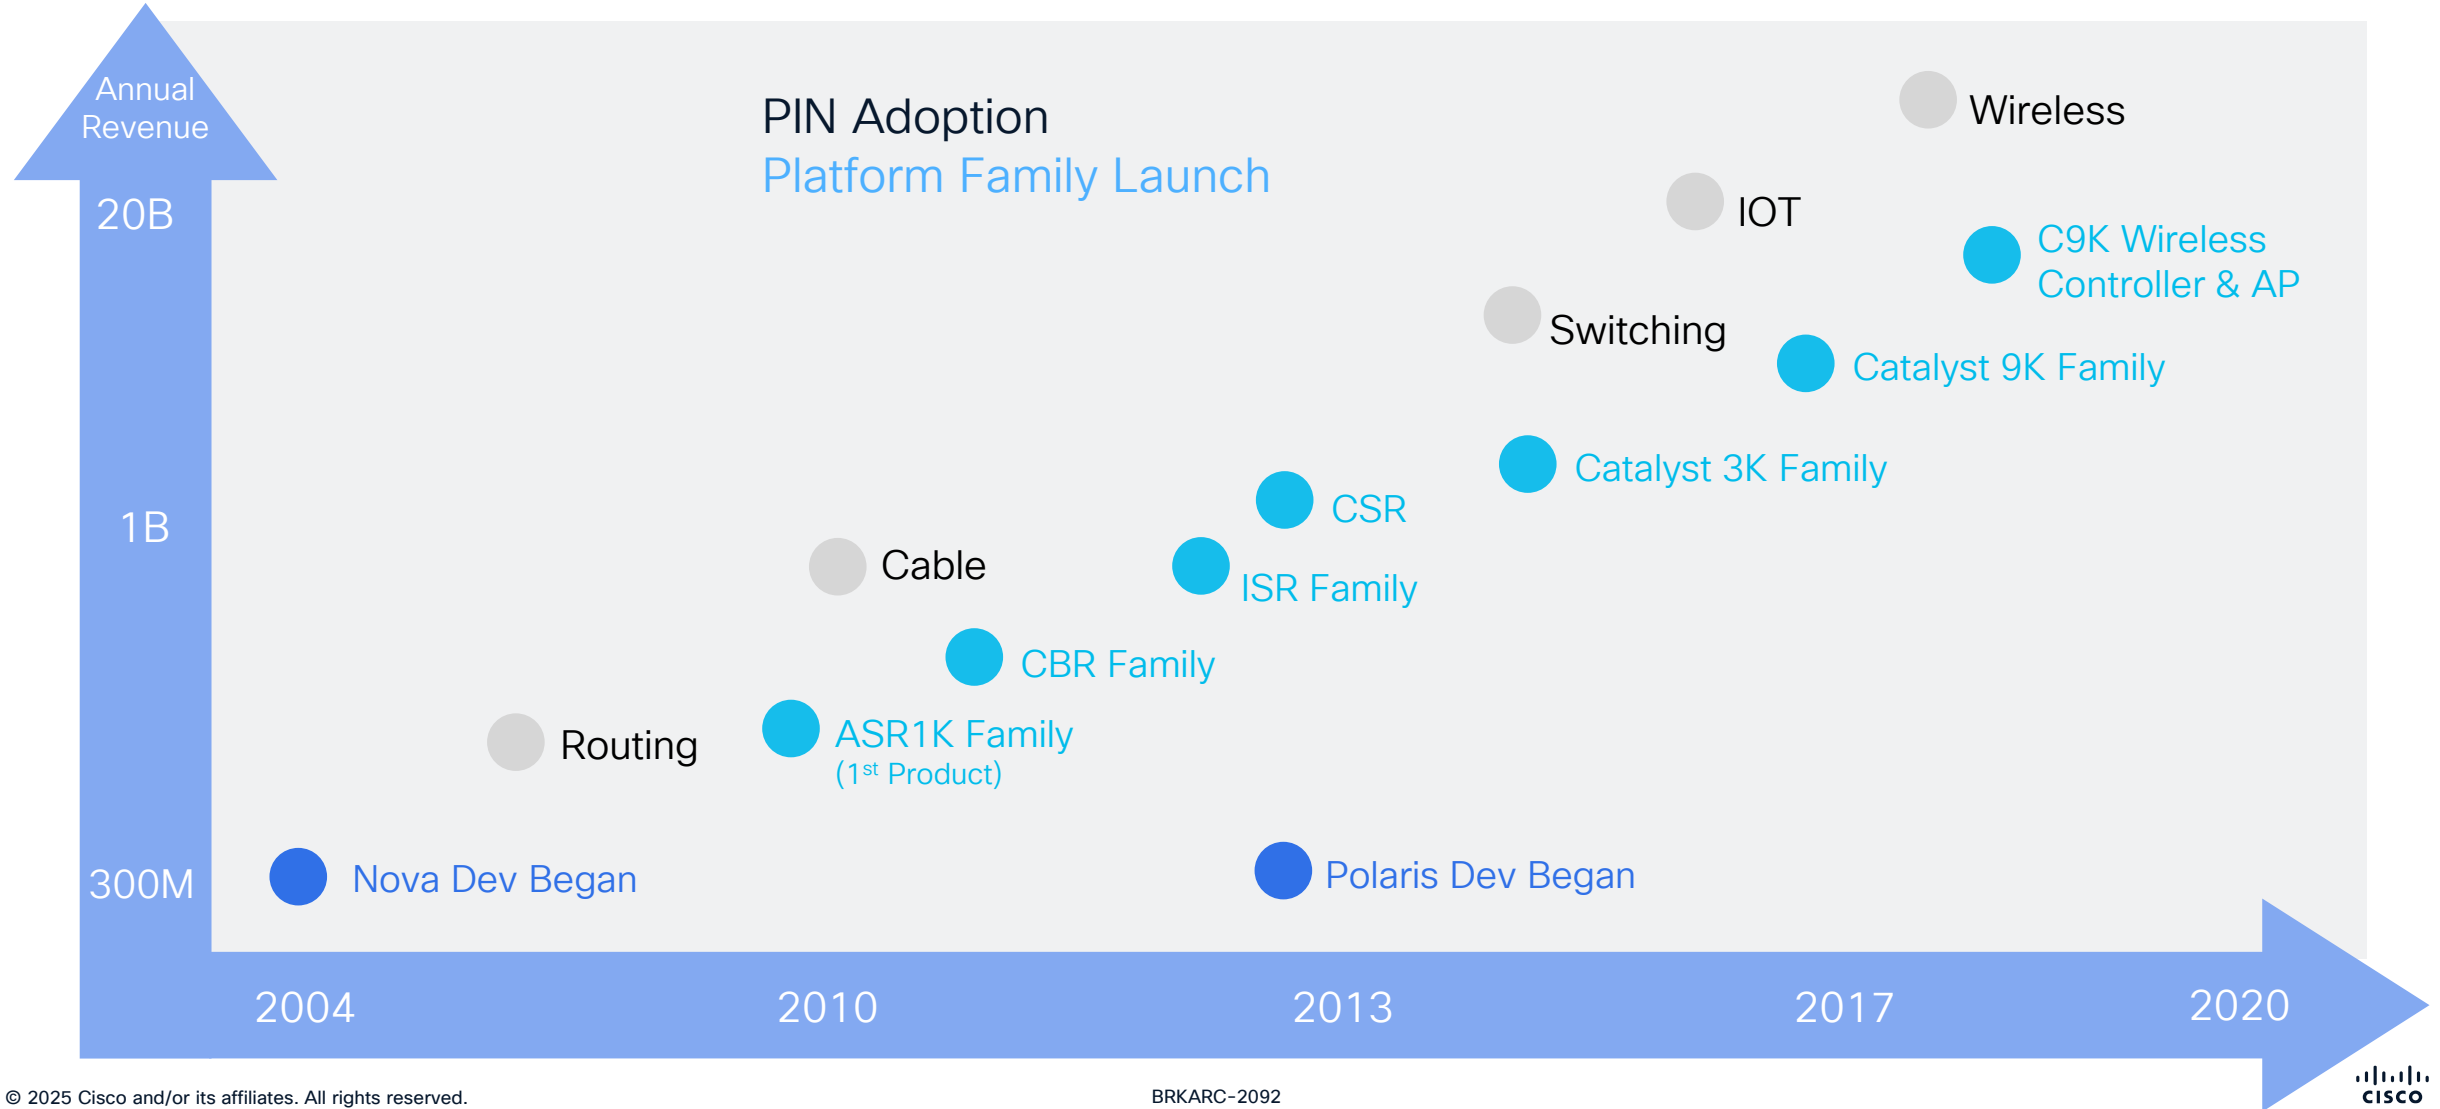
NOTE: Timeline is just an approximation



Year	Operating System	Key Features / Milestones	Product Series
1984	Cisco was born		
1986	Operating System 6.0	Cisco Ships First AGS Router	
1993	IOS 9.x - 11.x	LAN Switching, WAN Switching	
2000	IOS 12.x S Release	Cat6500, 7600 Series	
2007	IOS XE BinOS 3.x	ASR1000 Series	
2009	IOS 15.x M&T Release	Cat6800 Series	
2010	IOS XE NOVA 3.x	Cat4500 Series, Cat3850	
2015	IOS XE Polaris 16.1.1	Cat3850-XS, ASR-X, ISR	
2017	Open IOS XE Polaris 16.5.1	Catalyst 9000 Series	

# Brief History of IOS XE

Across Cisco Enterprise Platforms



# Cisco IOS XE - Architecture Evolution

Same look and feel - more powerful architecture



## Cisco IOS




- **Monolithic IOS**
- Compact, Streamlined
- High performance

## Cisco IOS XE 3.7.x (SE)



- **Monolithic IOSd:** Control-plane
- Sub-packages for data plane
- Linux daemons hosting capability
- Message parsing capability

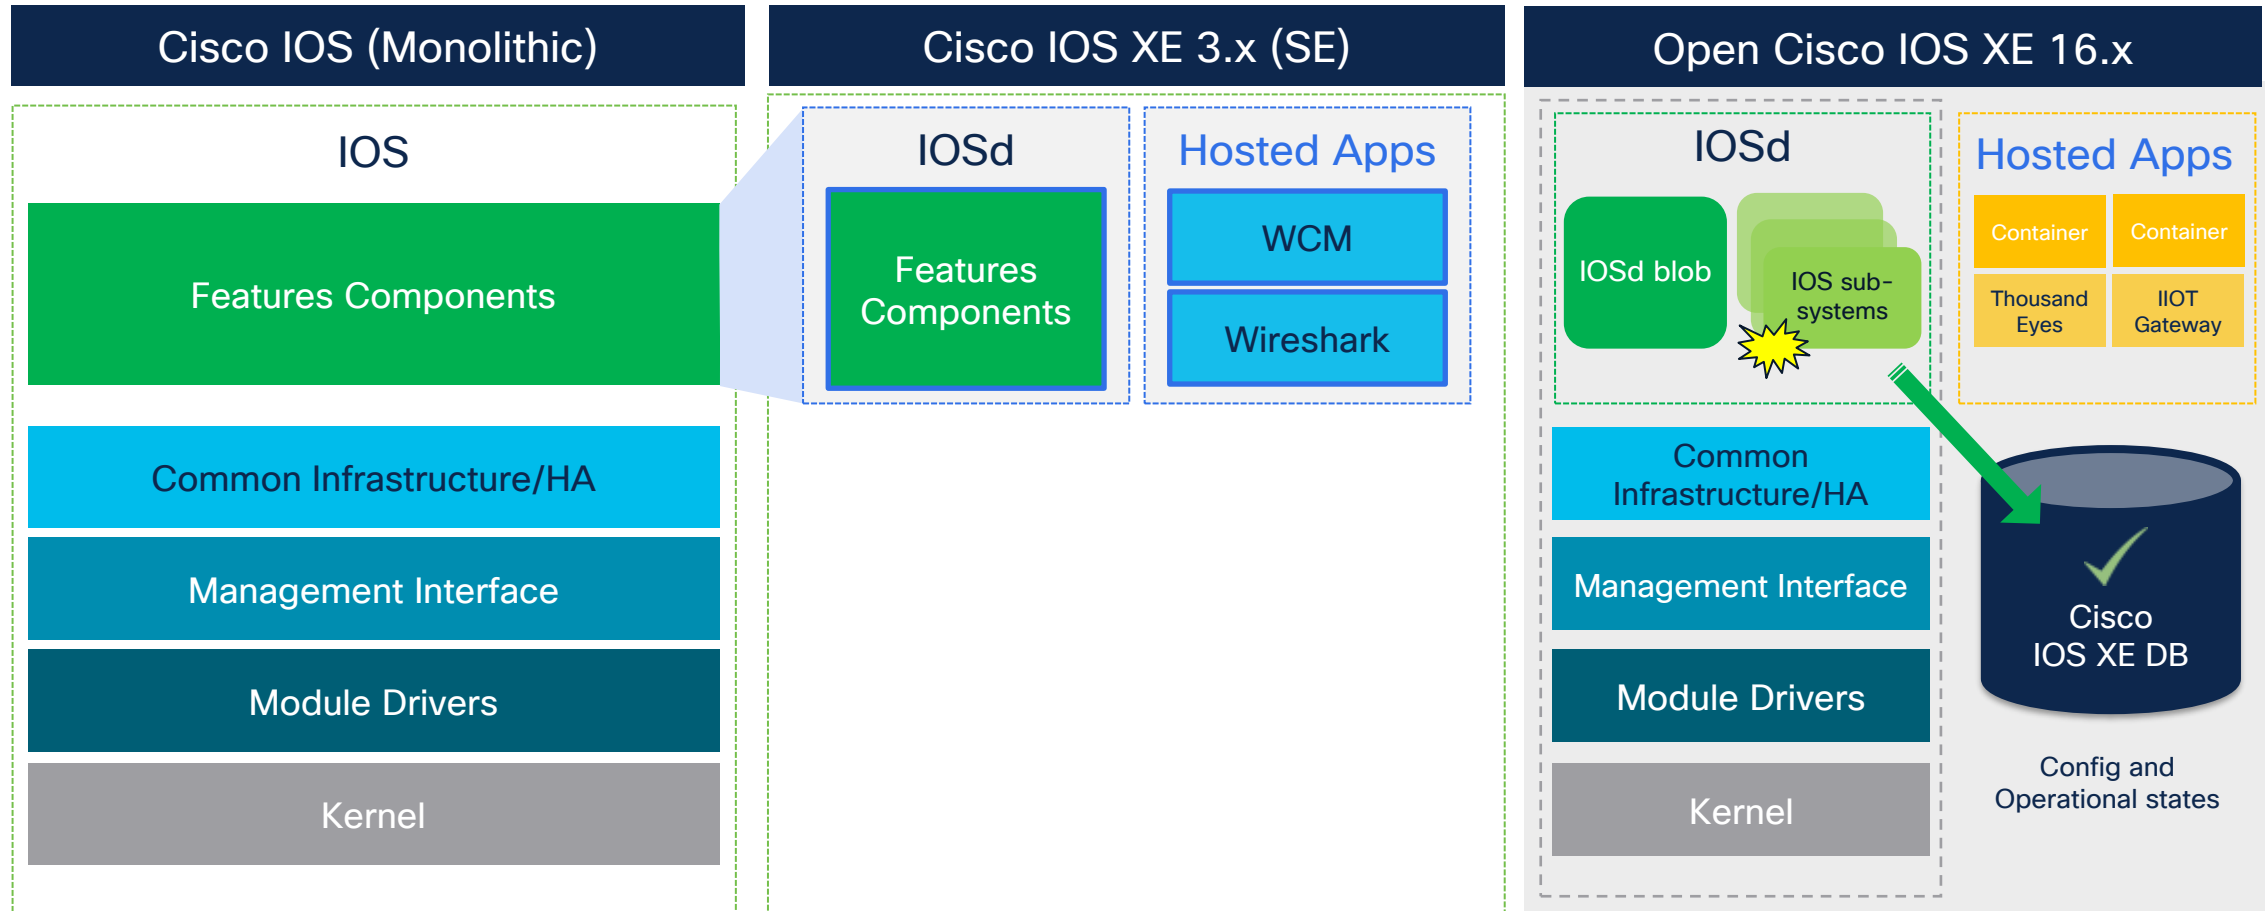
## Cisco IOS XE 16.x



- **IOSd:** Component assemblies
- **Modularized features:** Sub-packages
- **Distributed Operating System**
- **IOS XE (Crimson) Database**
- Radioactive tracing and events

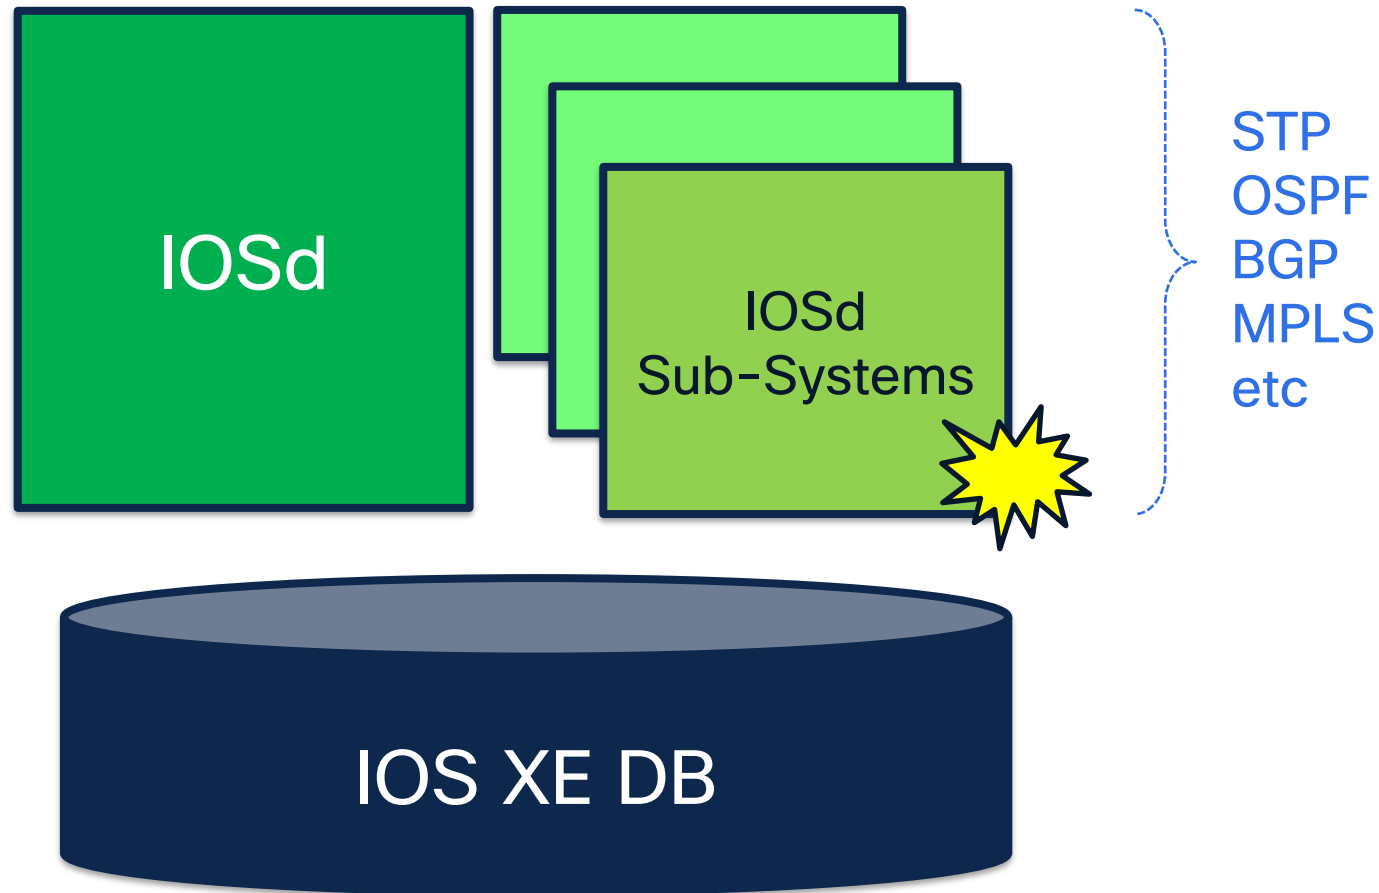
# What is Cisco IOS XE?

Same look and feel - more powerful architecture



Modern Software Architecture - with the same look and feel

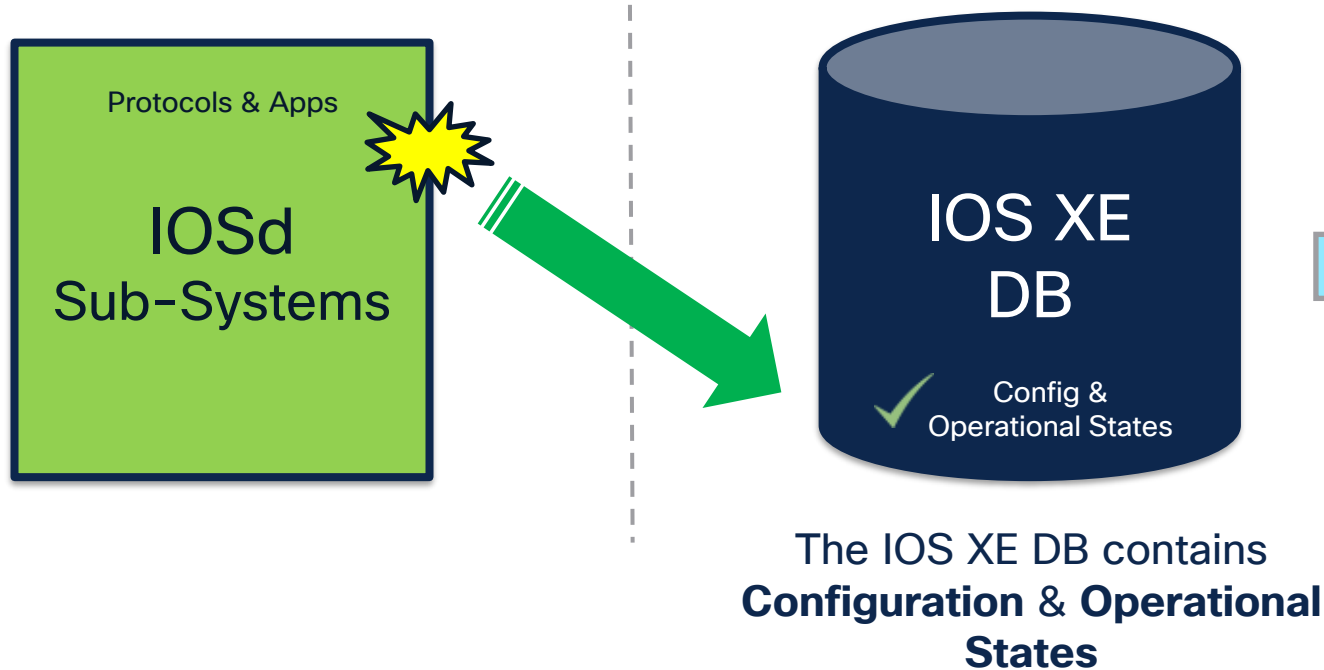
# Cisco IOS XE – IOS Sub Systems



Failure of one  
IOS XE  
Sub-System(s) -  
keeps rest of the  
system intact

IOSd Sub-Systems enhance IOS Resiliency

# Cisco IOS XE – Hardware DB



Decoupling Code & Data protects the Configuration & Operational States

Link State	STP State	OSPF State	Logs
Link State	Logs	MST State	
BGP State	Tunnel State		

Data Models

```

<nodes xmlns="urn:opendaylight:inventory">
  <node>
    <id>controller-config</id>
  </node>
  <node>
    <id>openflow1</id>
    <table xmlns="urn:opendaylight:flow:inventory">
      <id>0</id>
      <flow>
        <id>561183150</id>
        <match>
          <ethernet-match>
            <ethernet-destination>
              <address>00:00:00:00:00:02</address>
            </ethernet-destination>
            <ethernet-source>
              <address>00:00:00:00:00:01</address>
            </ethernet-source>
          </ethernet-match>
        </match>
        <flow-name>mac2mac</flow-name>
        <priority>512</priority>
        <instructions>
          <instruction>
            <order>0</order>
            <apply-actions>
              <action>
                <order>0</order>
                <output-action>
                  <output-node>connector>openflow:11</output-node>
                  <max-length>65535</max-length>
                </output-action>
              </action>
            </apply-actions>
          </instruction>
        </instructions>
        <idle-timeout>0</idle-timeout>
        <hard-timeout>0</hard-timeout>
        <table-id>0</table-id>
        <cookie>38264189495927313</cookie>
        <buffer-id>0</buffer-id>
      </flow>
    </table>
  </node>
</nodes>

```


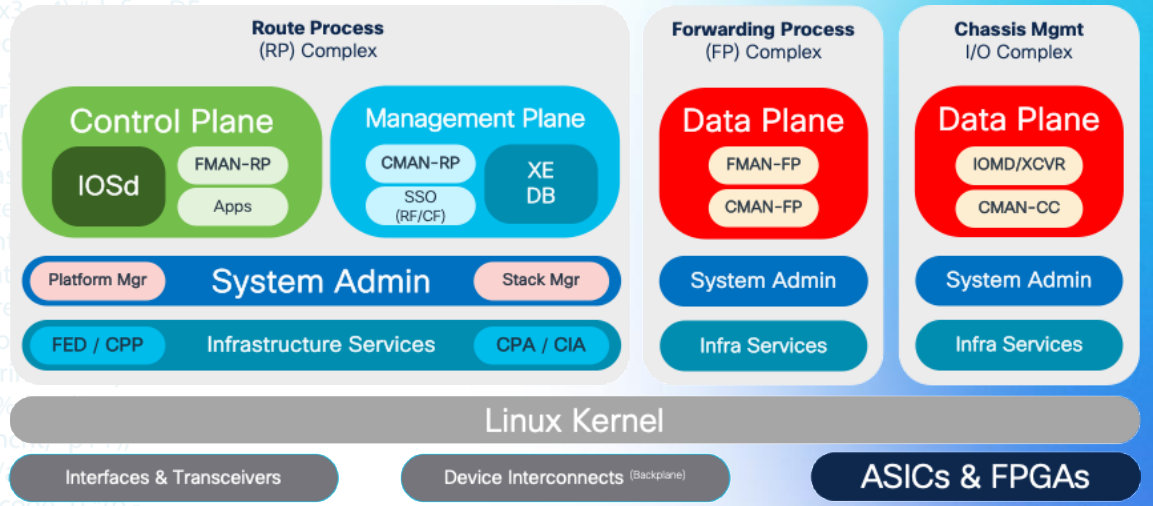
- Higher Application UP Time
- Quicker Recovery
- Better Convergence

# IOS XE Architecture

- Control Plane
- Data Plane
- System Plane
- Management Plane

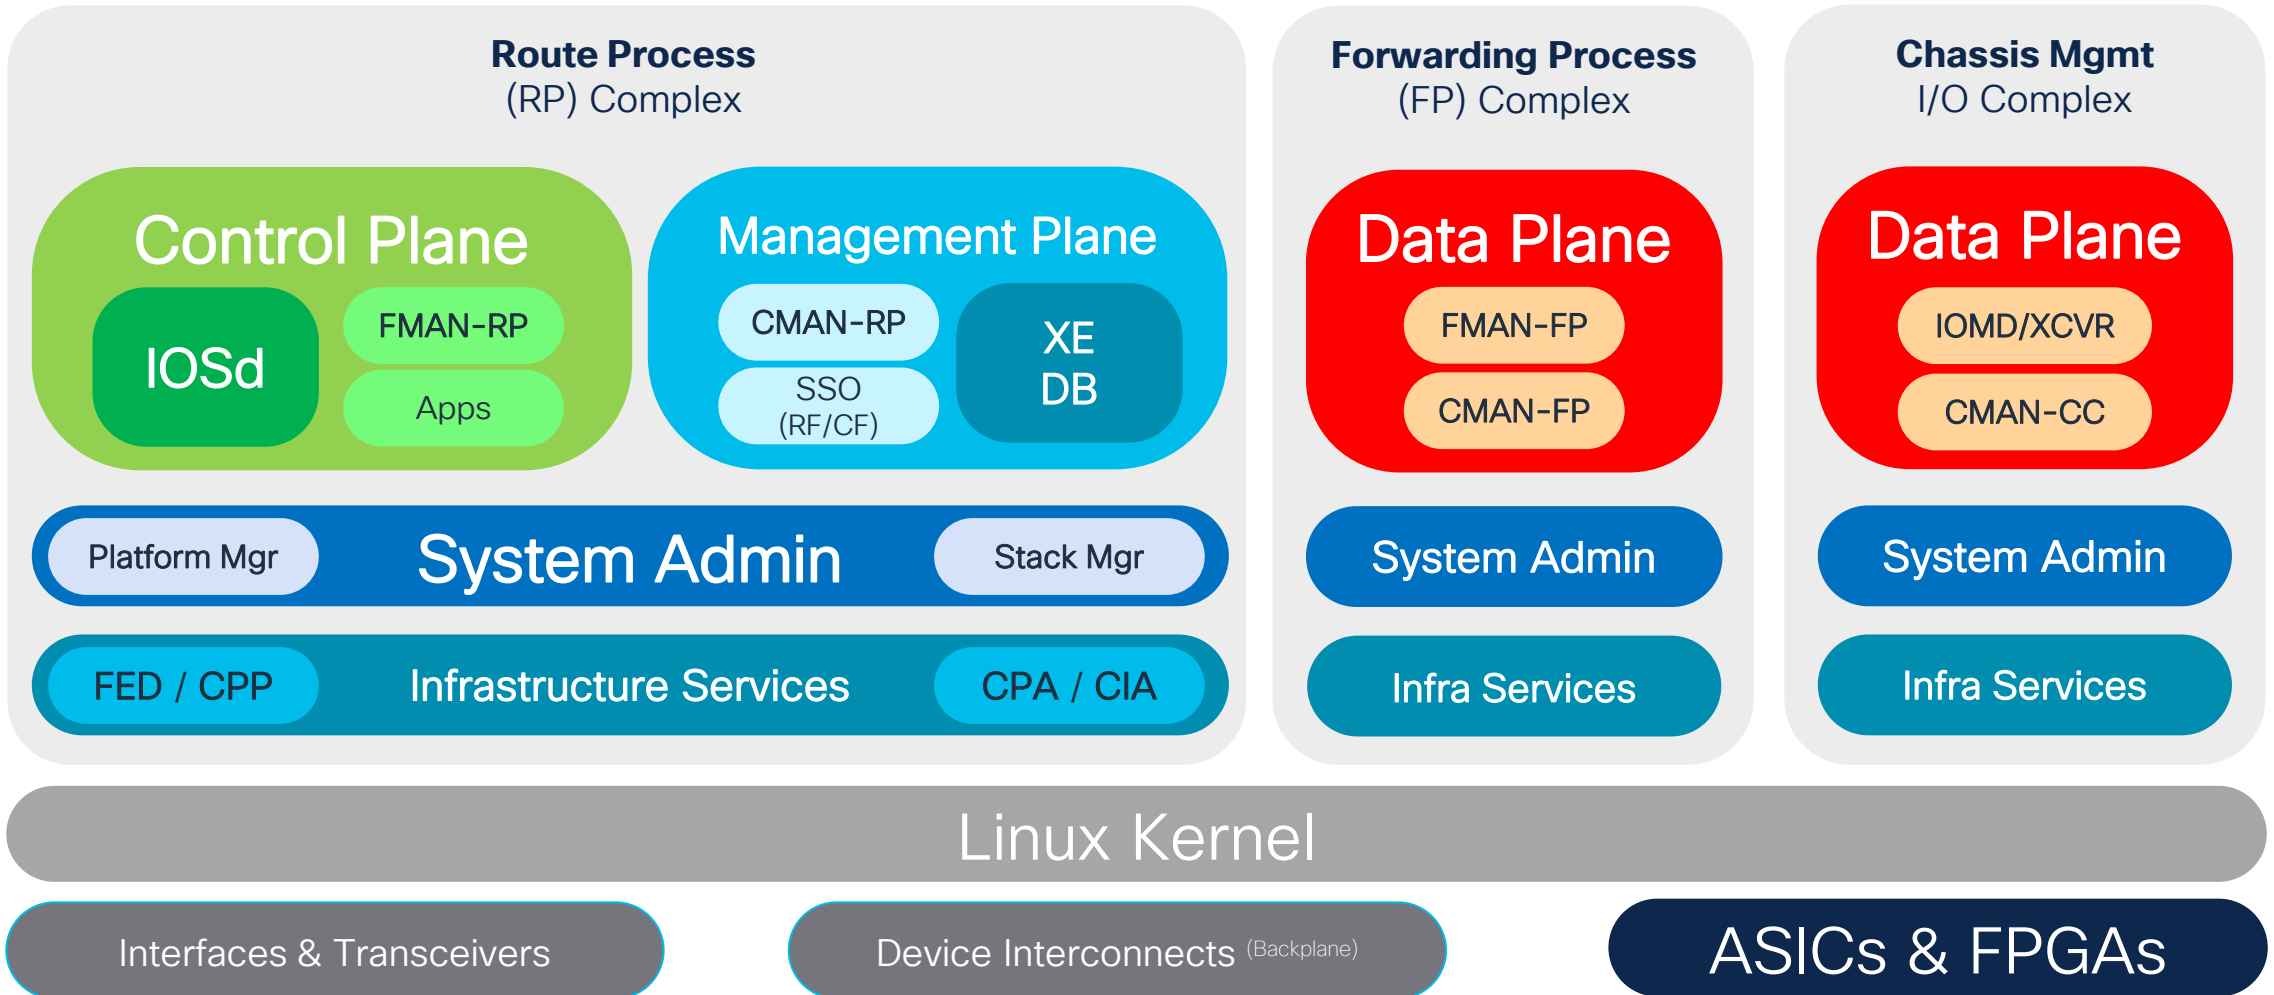
# IOS XE on Cisco C9K

## IOS XE Lite




# Cisco IOS XE Architecture

Modularized Components for Software Abstraction



# Cisco IOS XE Software

## PI vs. PD Software Components



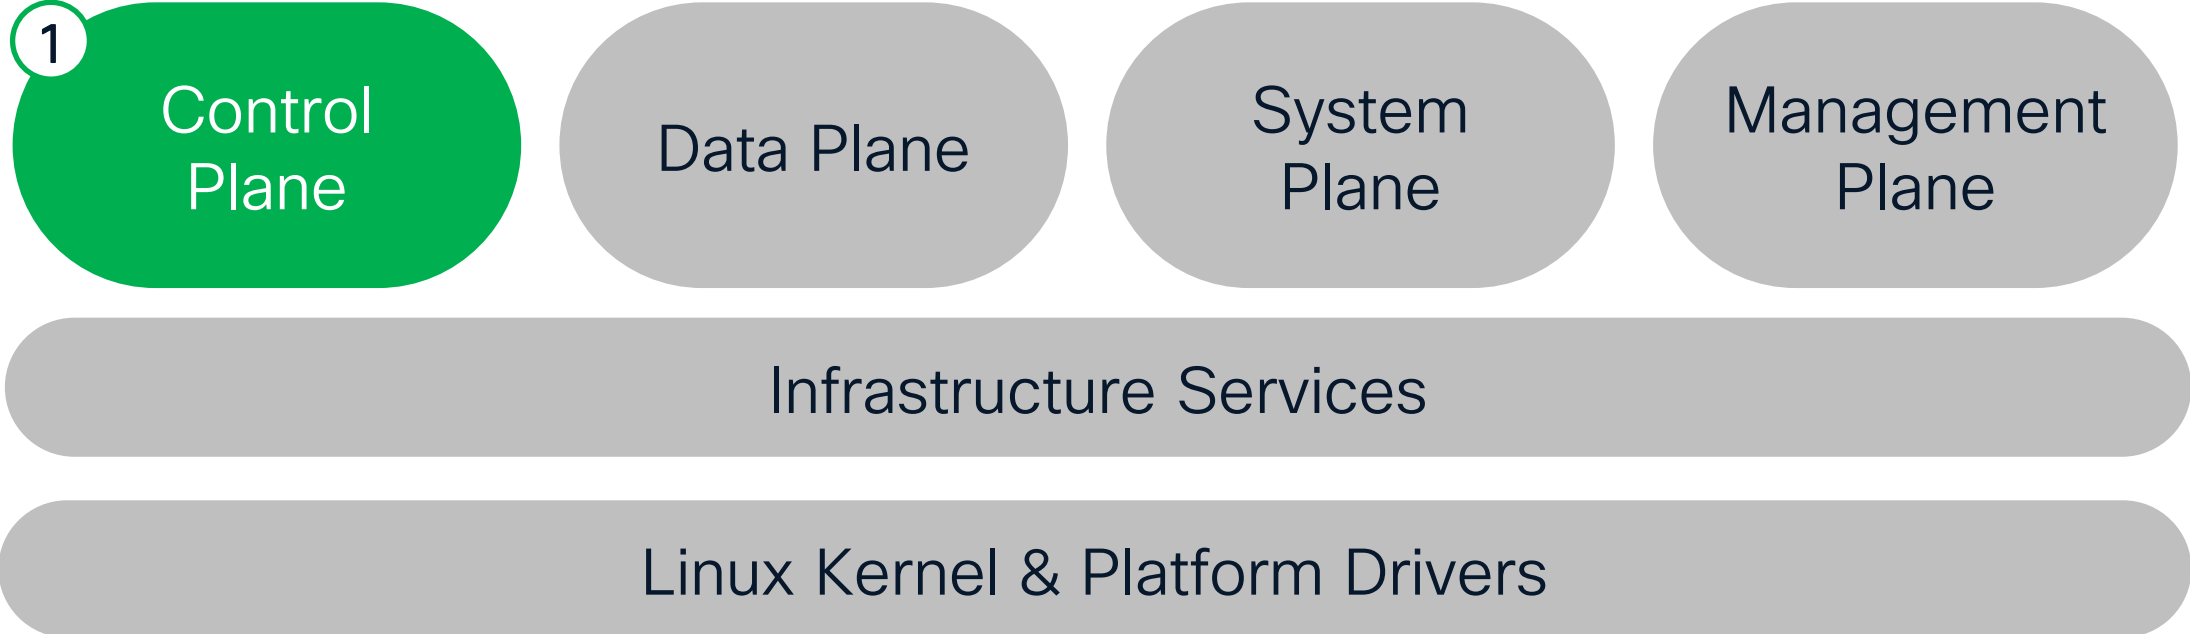

### Platform Independent (PI)

- IOS - Internetwork Operating System
- FMAN - Forwarding Manager
- RP - Routing Process
- FP - Forwarding Process
- CGM - Classification Group Manager
- IFM - Interface Manager
- AOM - Abstract Object Manager
- PDS - Packet Distribution Service
- LSMPI - Linux Shared Memory

### Platform Dependent (PD)

- CPA - Common Platform Abstraction
- FED - Forwarding Engine Driver
- IOMD - I/O Manager
- CMAN - Chassis Manager
- PMAN - Platform Manager
- SMAN - Stack Manager
- XCVR = Transceiver/Optics
- Table Manager - Client & Server
- Punject - Punt+Inject (CPU) interface

# Cisco IOS XE - Control Plane

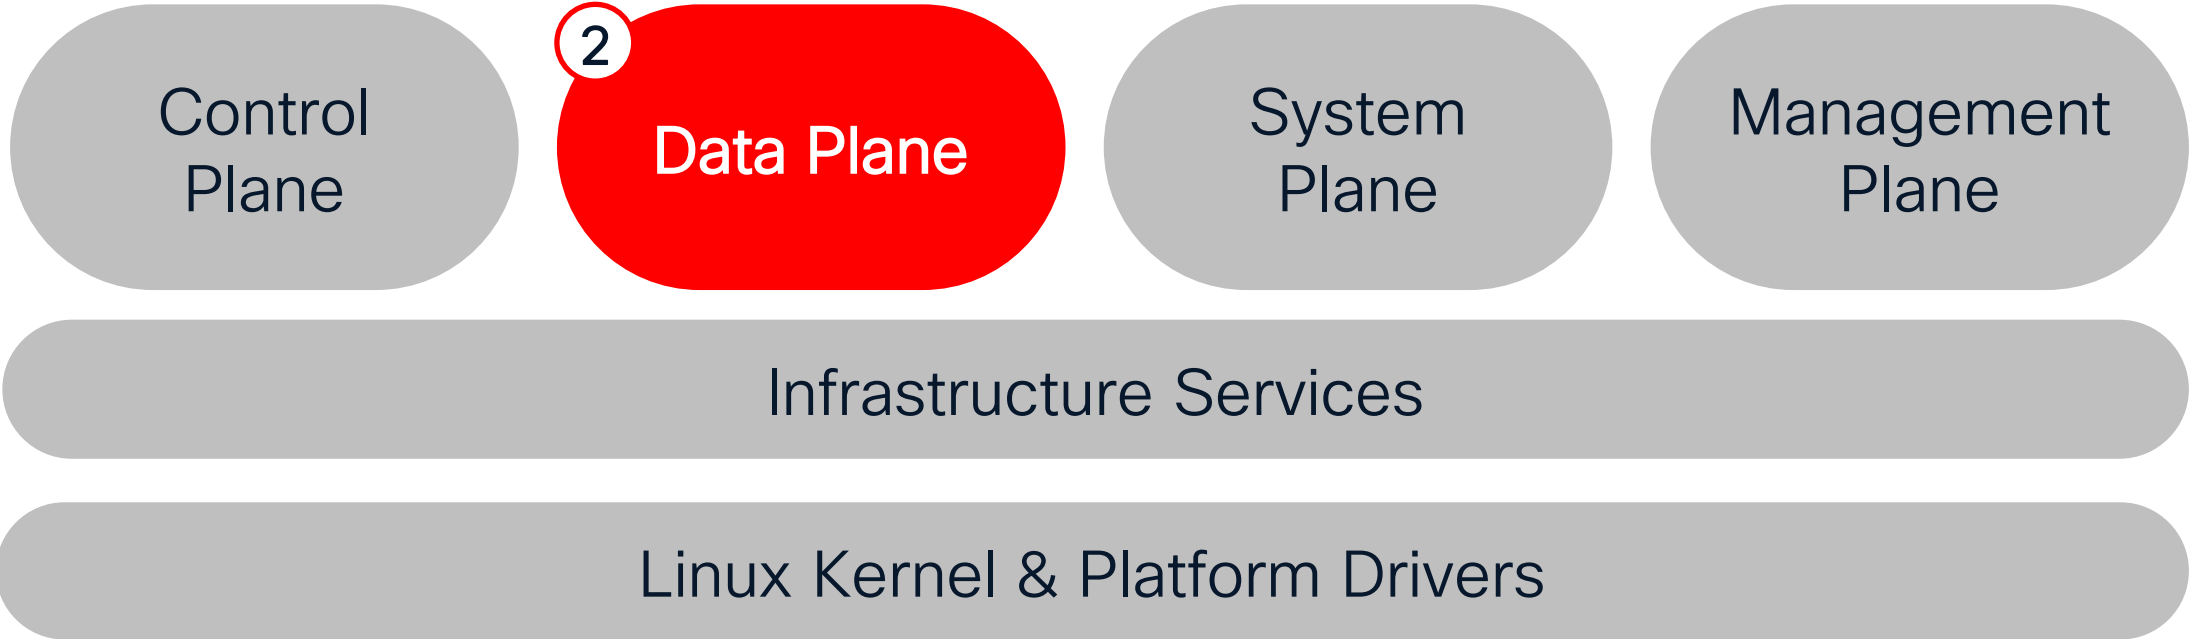



## This is the 'brain' of the network stack

- Most control-plane logic runs within IOSd
- Home to routing & bridging protocols (network learning)
- Richest networking features in industry (~5000 features)
- Distributes protocol (RP) forwarding states to data-plane (FP)




# Cisco IOS XE - Data Plane

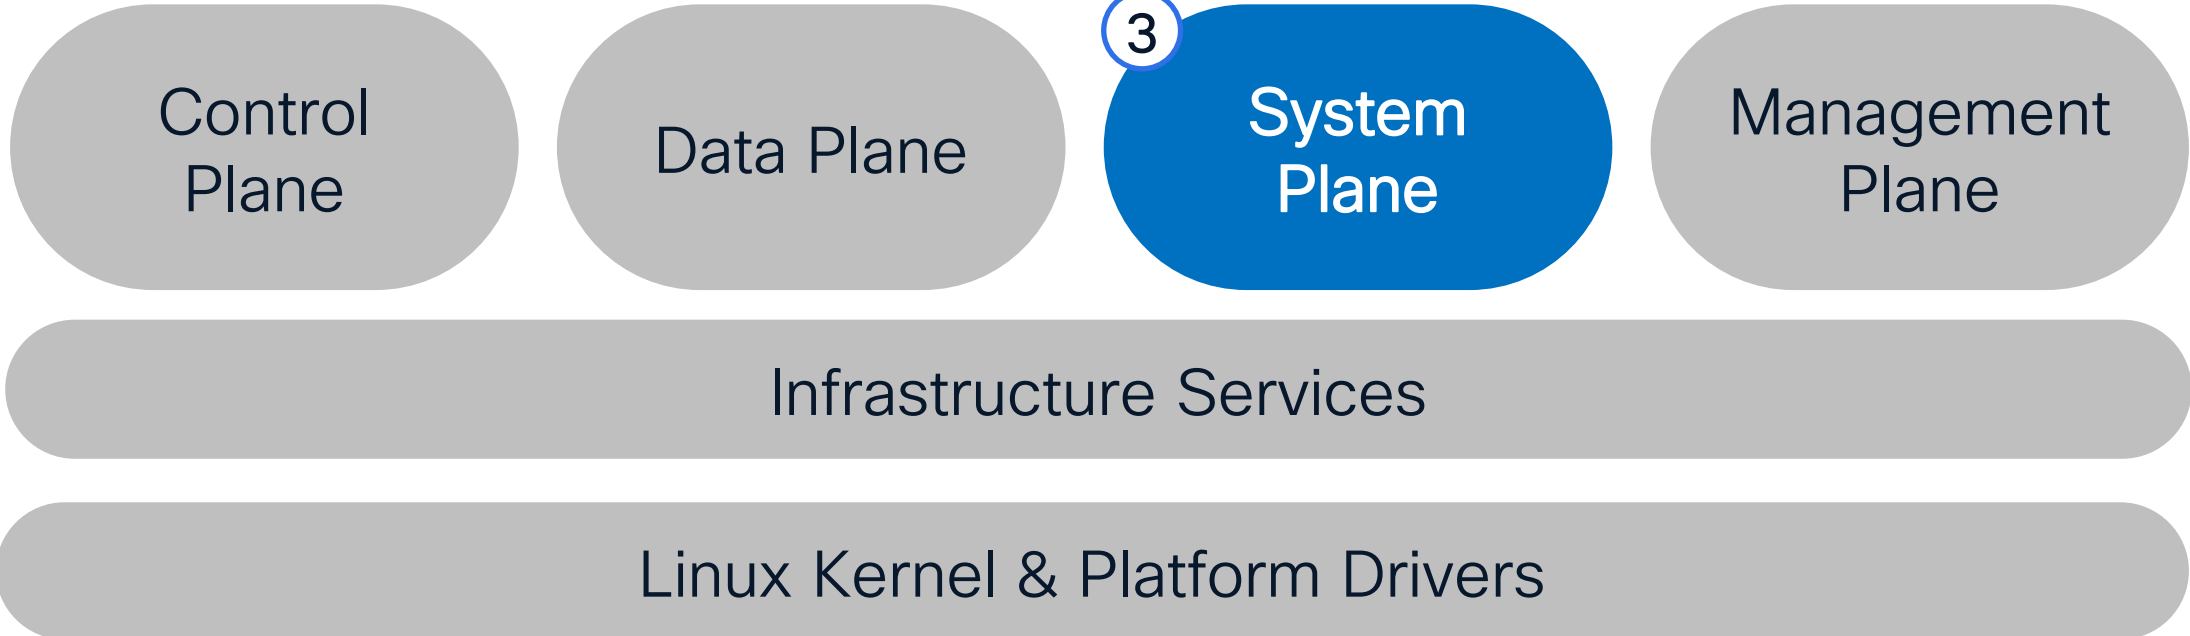



## Handles high-speed Packet Forwarding

- Touches every packet! High-throughput and low-latency forwarding
- Programming from control-plane abstracted by well defined APIs
- Supports multiple forwarding architectures: standalone & modular
- Forwarding is generally handled in custom hardware (e.g. UADP & S1)



# Cisco IOS XE - System Plane



## General Administration & functions of the System

- Manages the Chassis, Modules, I/O, Power, Fans
- Manages Stacking & Virtual Chassis processing
- Also manages software image management & patching

CMAN-FP

CMAN-CC

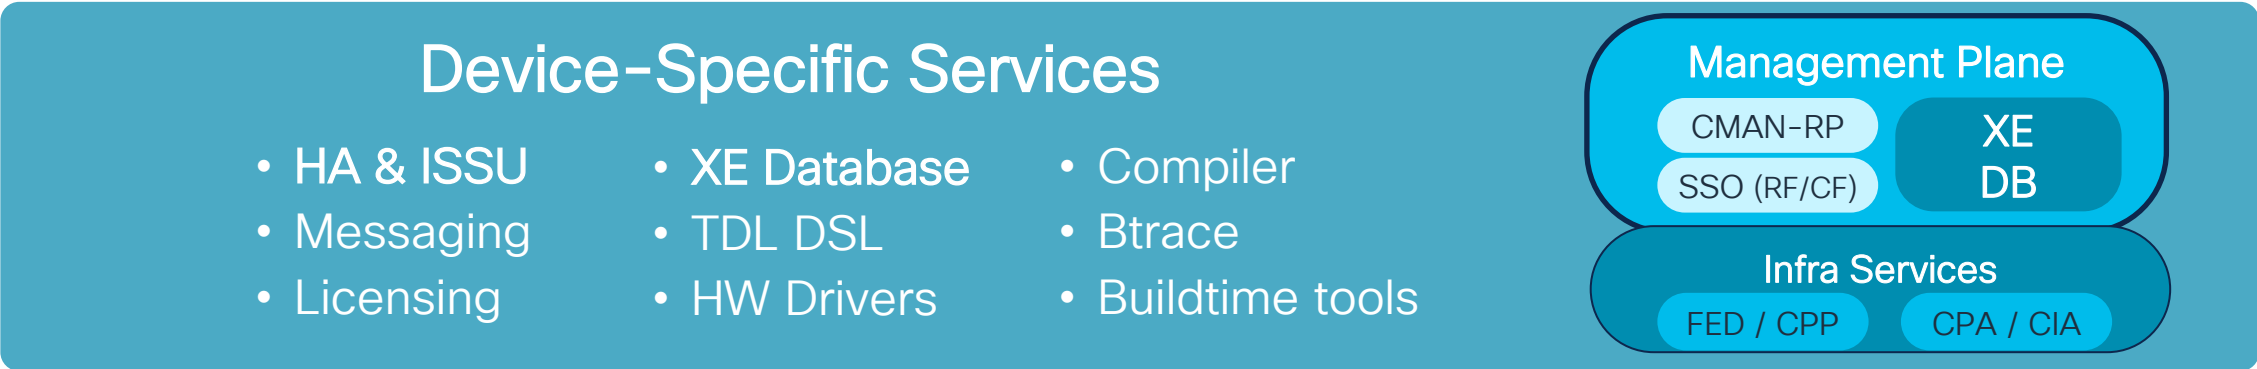

IOMD

### System Admin

Platform Mgr

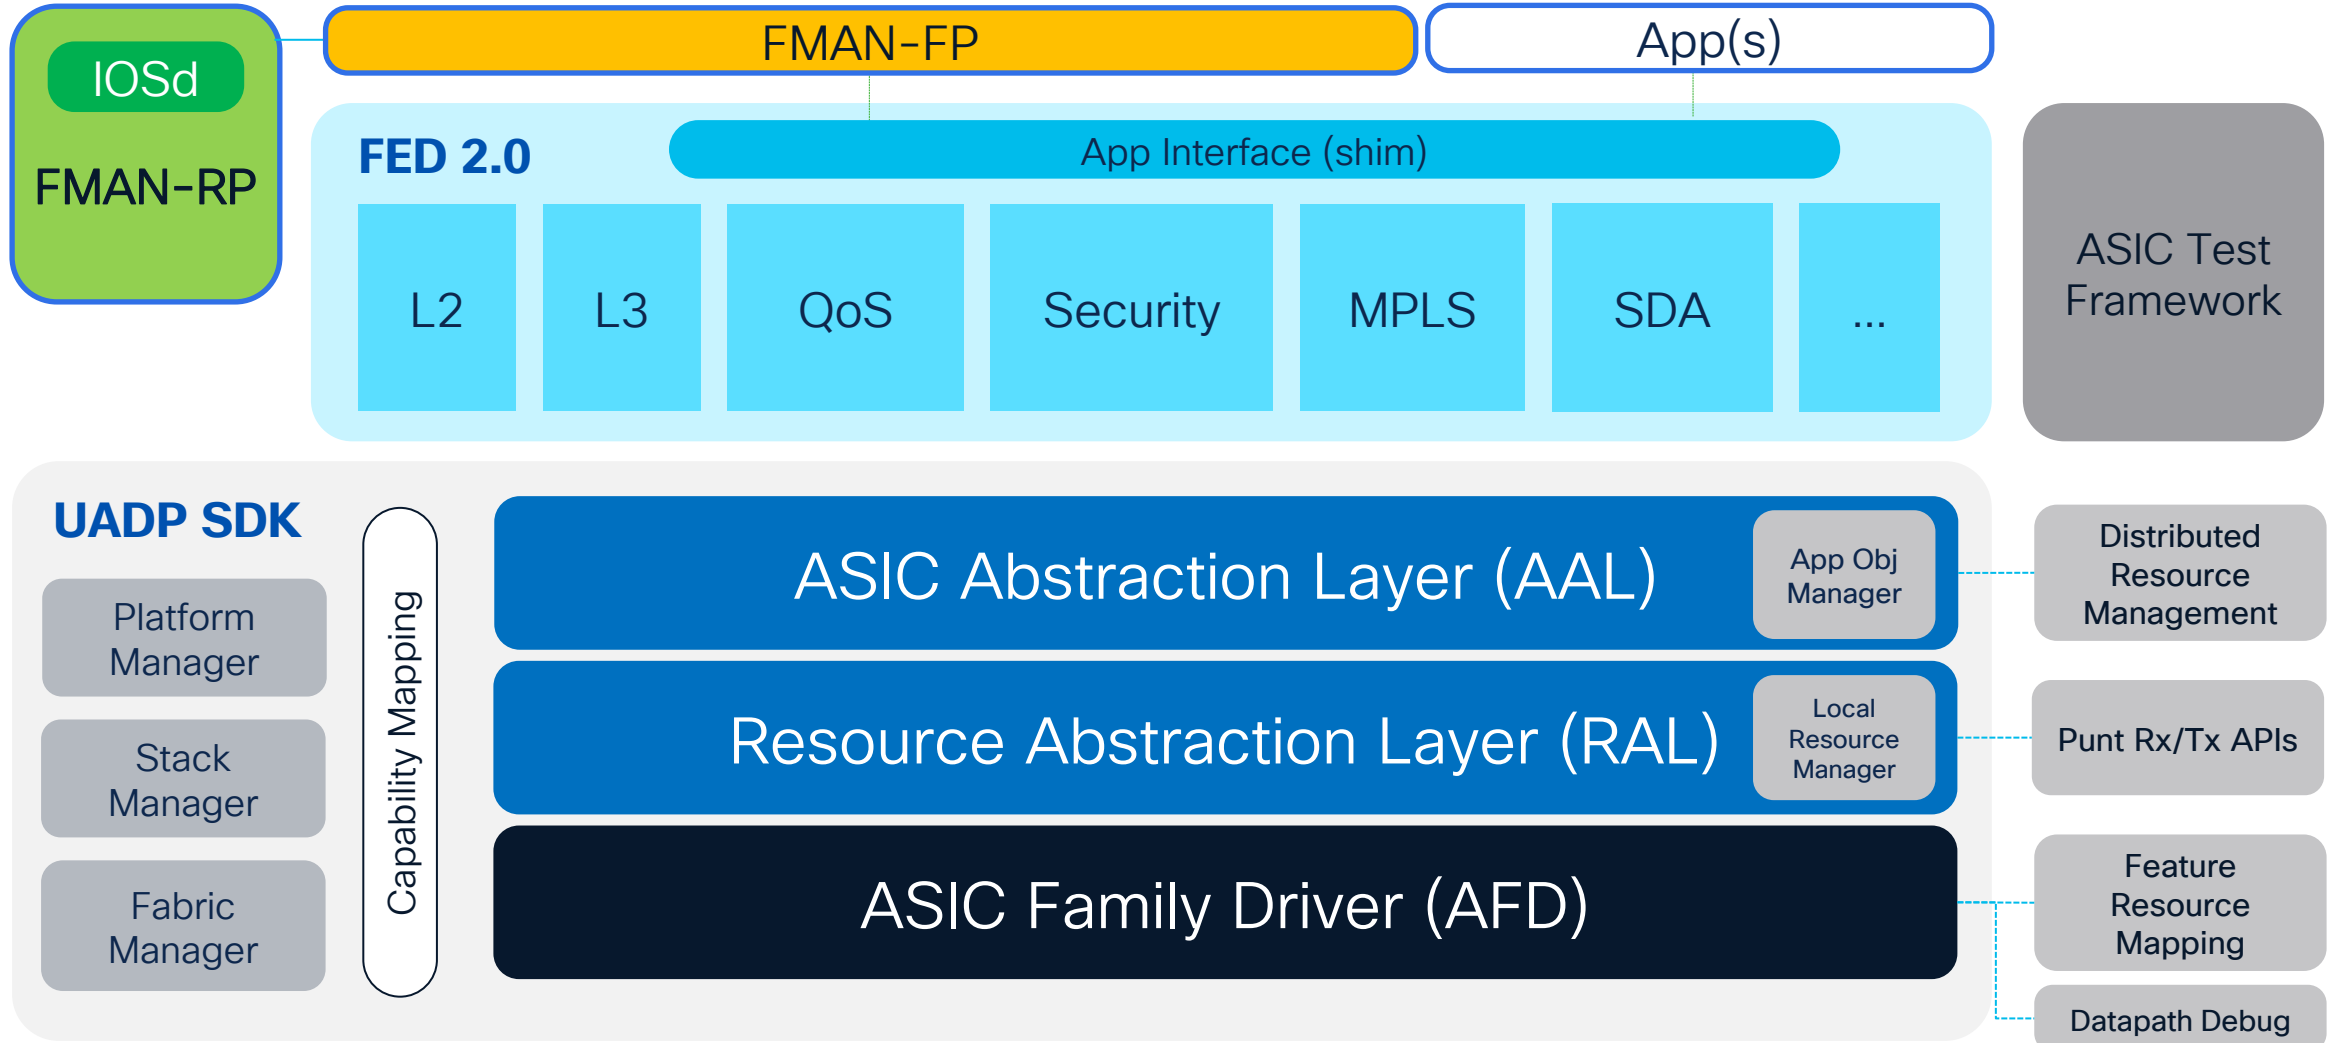
Stack Mgr

# Cisco IOS XE – Management & Infra

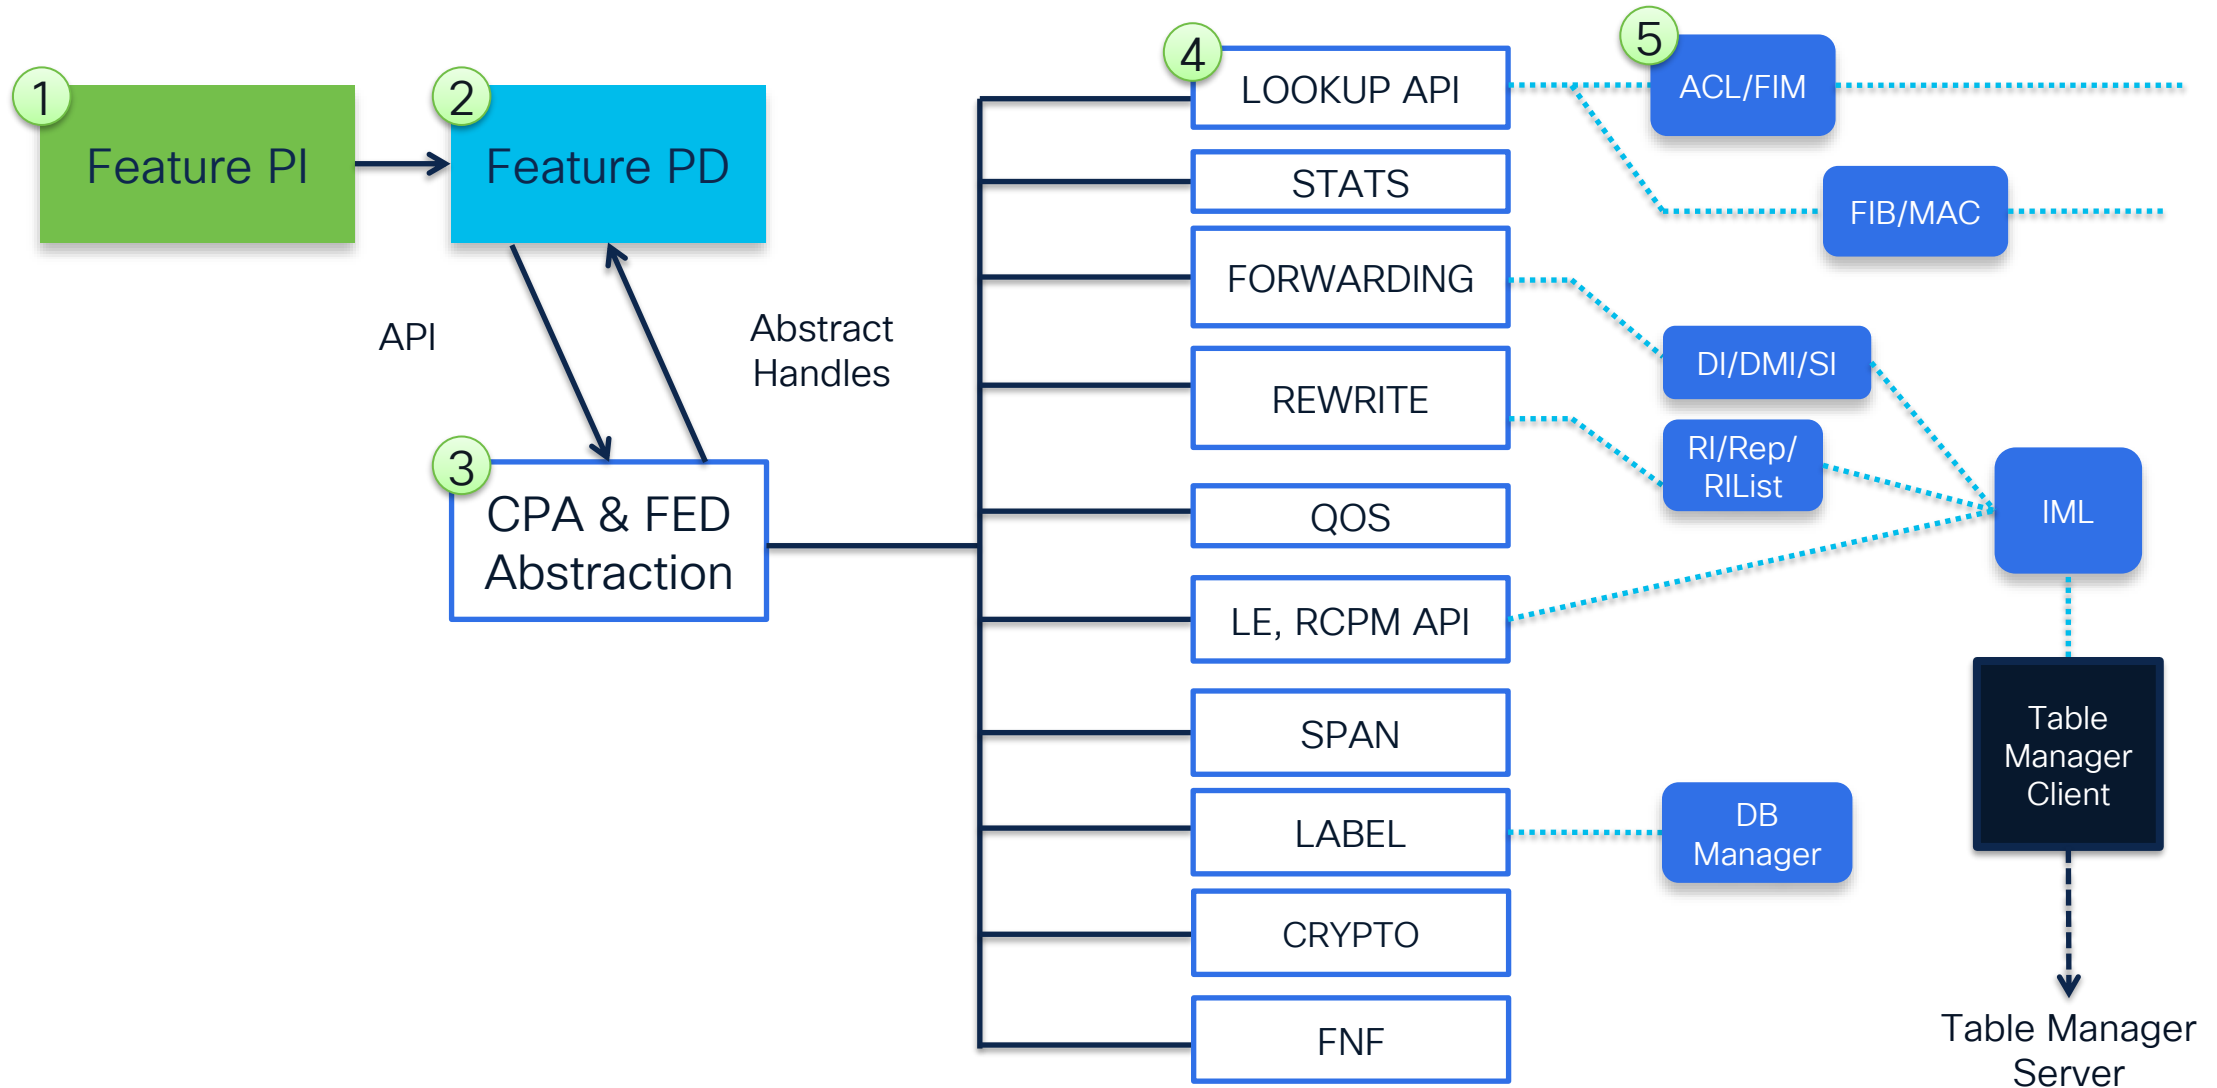


# Cisco IOS XE on Catalyst 9000 Series

## Hardware Forwarding Architecture

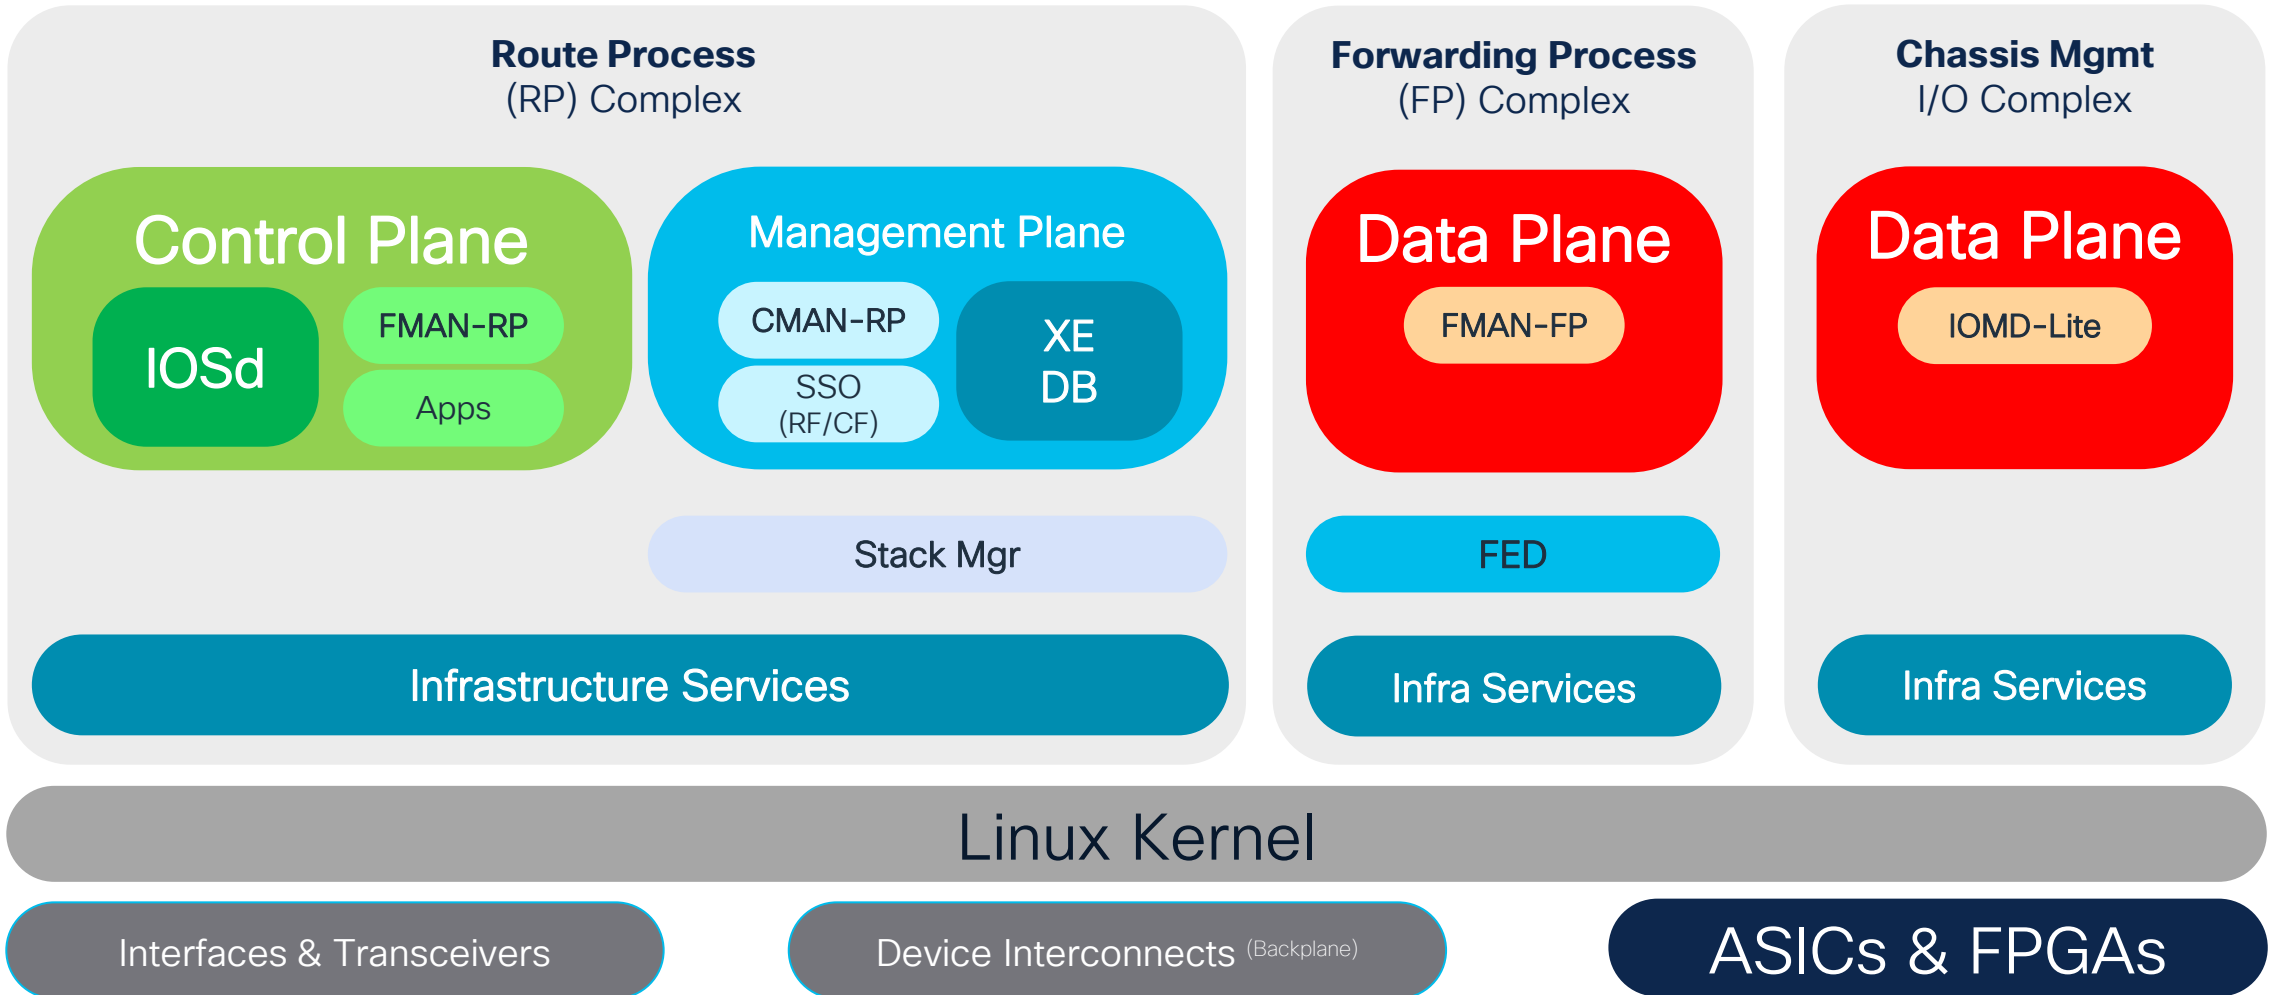



# Cisco IOS XE - PD Abstraction Layer






# Cisco IOS XE "Lite" Architecture

Same code base as Cisco IOS XE – Optimized for ARM CPU & Memory

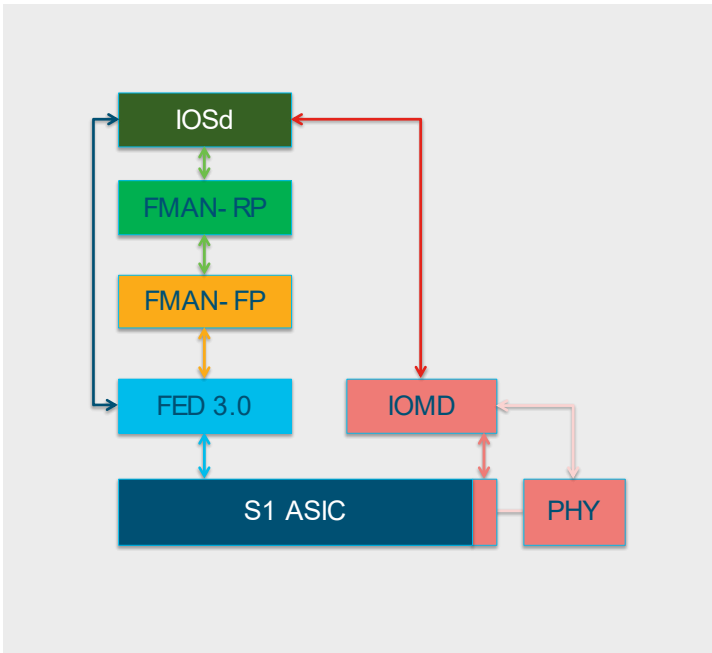


# Fixed vs. Modular vs. Stacking

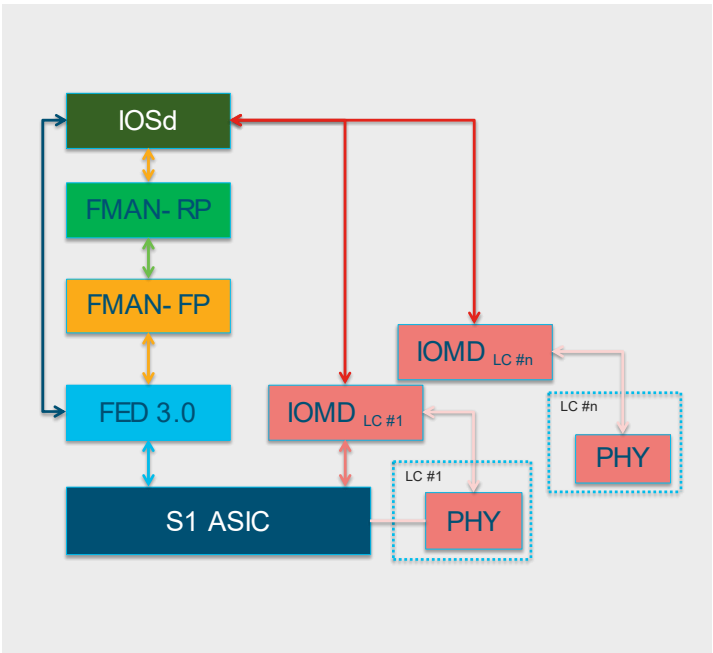
Reusing Common Elements of OS Architecture




## Fixed Platform




## Modular Platform



## Stacked Platform(s)




# Cisco IOS XE 16.x - 17.x



Catalyst 9200/L/CX

IOS XE Lite  
Binary Image



Catalyst 9300/L/X



Catalyst 9400/X



Catalyst 9500/X



Catalyst 9600/X


IOS XE  
Binary Image

## Catalyst 9000 runs the same IOS XE Operating System


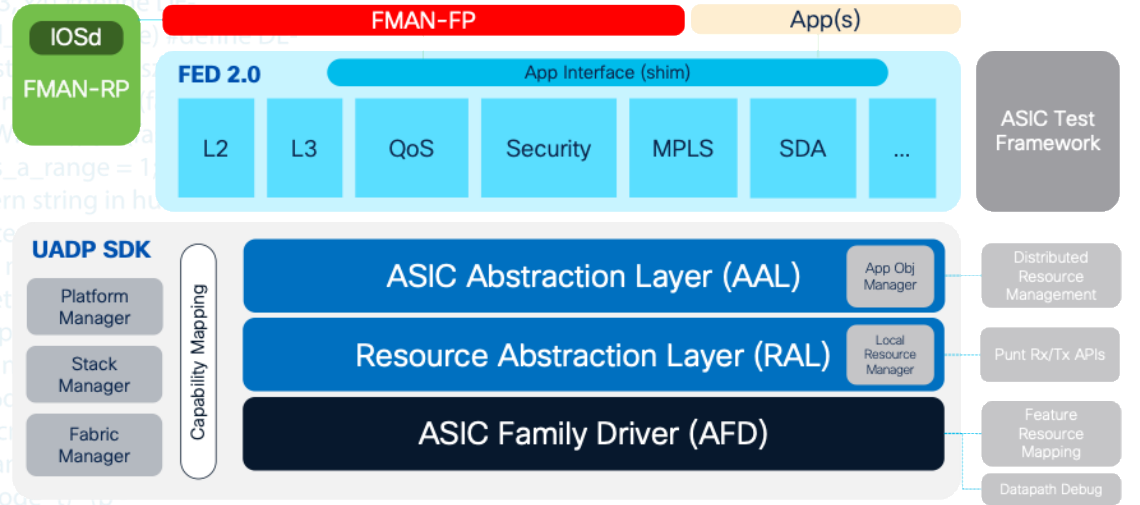
```

extract_num-
ber_and_incr (destination, source) int
*destination; unsigned char **source; { extract_num-
ber (destination, *source); *source += 2; } #ifndef EXTRACT_MAC-
ROS #undef EXTRACT_NUMBER_AND_INCR #define EXTRACT_NUM-
BER_AND_INCR(dest, src) \ extract_number_and_incr (&dest, &src) #endif /*
not EXTRACT_MACROS */ #endif /* DEBUG */ /* If DEBUG is defined, Regex prints
many voluminous messages about what it is doing (if the variable `debug' is nonzero). If
linked with the main program in `iregex.c'; you can enter patterns and strings interactively.
And if linked with the main program in `main.c' and the other test files, you can run the al-
ready-written tests. */ #ifdef DEBUG /* We use standard I/O for debugging. */ #include <stdio.h>
/* It is useful to test things that `must' be true when debugging. */ #include <assert.h> static int
debug = 0; #define DEBUG_STATEMENT(e) e #define DEBUG_PRINT1(x) if (debug) printf (x) #define
DEBUG_PRINT2(x1, x2) if (debug) printf (x1, x2) #define DEBUG_PRINT3(x1, x2, x3) if (debug) printf
(x1, x2, x3) #define DEBUG_PRINT4(x1, x2, x3, x4) if (debug) printf (x1, x2, x3, x4)
extern void printchar(); /* Print the fastmap in human-readable form. */ void pri
while (i < (1 << BYTEWIDTH)) { unsigned was_a_range = unsigned j = 1; while (i < (1 << BYTEWIDTH)) {
(unsigned char) (was_a_range) { printf ("-"); printchar (i); } putchar (i); } /* Print a compiled pattern string in hu-
man-readable form, starting at the START pointer into it and ending just before the pointer
unsigned char *end; { int
pattern commands. */ while (p < pend) { switch ((re_opcode_t) *p++) { case no_op
break; case exactn: mcnt = *p++; printf ("/exactn/%d", mcnt); do { putchar ('/'); prin
break; case exactn: mcnt = *p++; printf ("/exactn/%d", mcnt); do { putchar ('/'); prin
break; case duplicate: printf ("/duplicate/%d", *p++); break; case anychar: printf ("/ar
break; case charset: case charset_not: { register int c; printf ("/charset%is", (re_opcode_t) *p
1) == charset_not ? "_not" : ""); assert (p + *p < pend); for (c = 0; c < *p; c++) { unsigned bit;
unsigned char map_byte = p[1 + c]; putchar ('/'); for (bit = 0; bit < BYTEWIDTH; bit++) if
(map_byte & (1 << bit)) printchar (c * BYTEWIDTH + bit); } p += 1 + *p; break; } case beg-
line: printf ("/begline"); break; case endline: printf ("/endline"); break; case on_failure_-
jump: extract_number_and_incr (&mcnt, &p); printf ("/on_failure_jump/0/%d", mcnt);
break; case failure_keep_stack_jump: extract_number_and_incr (&mcnt, &p); printf
("/on_failure_keep_stack_jump/0/%d", mcnt); break; case dummy_failure_jump: ex-
tract_number_and_incr (&mcnt, &p); printf ("/dummy_failure_jump/0/%d", mcnt); break;
case push_dummy_failure: printf ("/push_dummy_failure"); break; case may-
be_pop_jump: extract_number_and_incr (&mcnt, &p); printf
("/maybe_pop_jump/0/%d", mcnt); break; case pop_failure_-
jump: extract_number_and_incr (&mcnt, &p); printf ("/pop_-
failure_jump/0/%d", mcnt); break; case jump_past_alt:
extract_number_and_incr (&mcnt, &p); printf ("/-

```



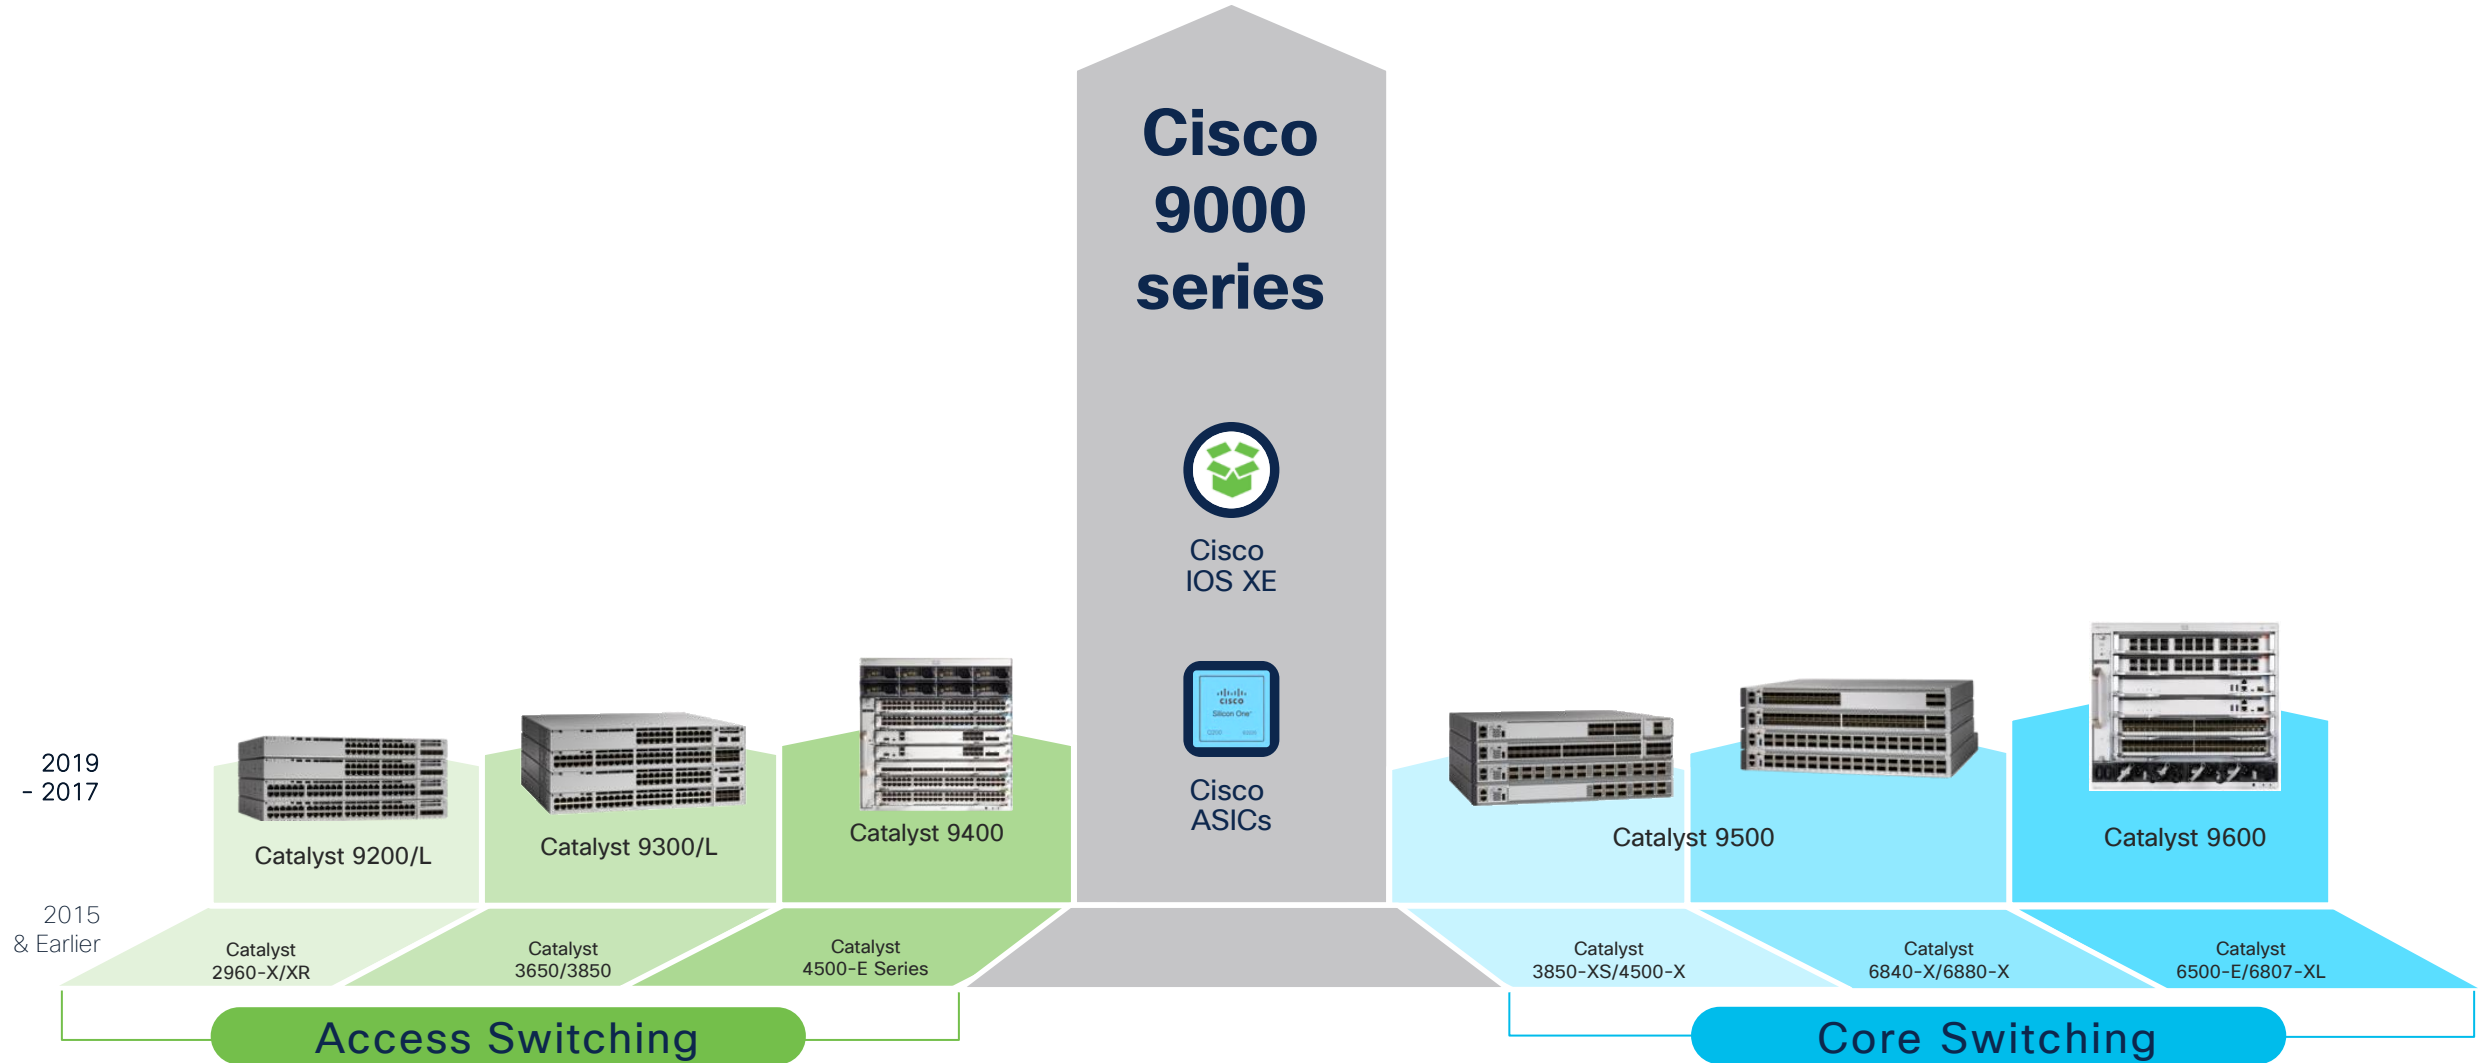
- ✓ Existing C9K Portfolio
- ✓ IOS XE Release Model
- ✓ IOS XE 17.1.1 - 17.6.1
- ✓ IOS XE 17.7.1 - 17.15.1



# Cisco Catalyst 9000 Switching Portfolio

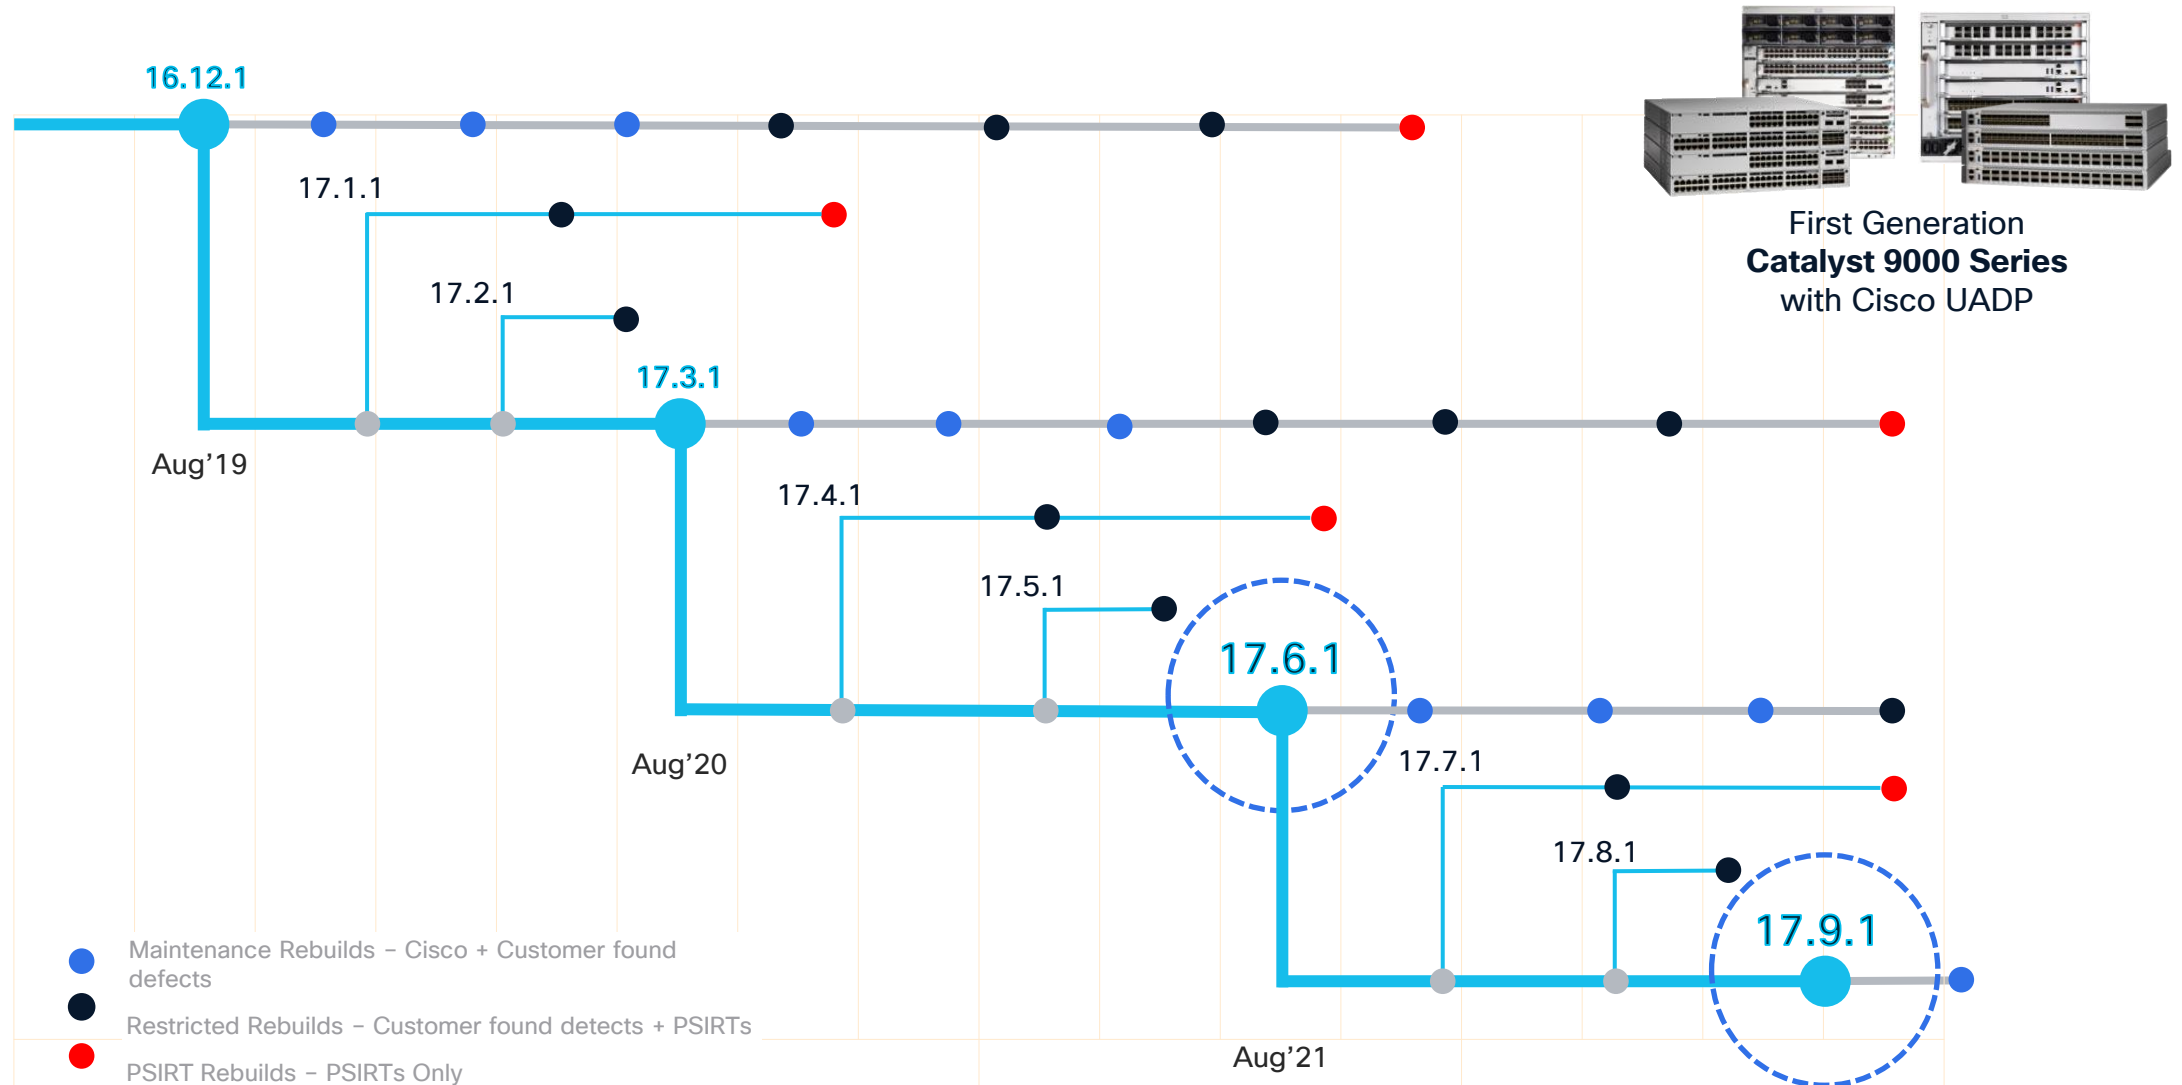
One Family from Access to Core – Common Hardware & Software

16.6<sup>·x</sup> – 17.6<sup>·x</sup>



# Cisco IOS XE - Release Schedule

Graphical Overview - 17.1.x to 17.9.x



# Catalyst 9000 Switching – Key Features

\* Limited Availability (LA) only

<b>IOS XE 17.1.1</b> (Dec'19) <span style="float: right;">SMR</span>	<b>IOS XE 17.2.1</b> (Apr'20) <span style="float: right;">SMR</span>	<b>IOS XE 17.3.1</b> (Aug'20) <span style="float: right;">EMR</span>
<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ 9200/9300 - Umbrella Integration</li> <li>❖ MACSEC over EoMPLS</li> <li>❖ ERSPAN to v6 Destination</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ TWS - Secure Swipe Clean DoD 5220.22-M standard</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ Enhanced ACL Logging</li> <li>❖ Wired Client Sensor*</li> <li>❖ 9200/9300 - Umbrella Switch Connector with AD Integration</li> </ul>
<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ Inter-AS Option A (VRF-Lite)</li> <li>❖ VPLS Flow Aware Transport (FAT) PseudoWire</li> <li>❖ Extranet mVPN</li> <li>❖ VXLAN aware Flexible Netflow</li> <li>❖ EVPN to VRF-Lite handoff for Border Spine</li> <li>❖ EVPN to MPLS handoff for Border Spine</li> <li>❖ EVPN Tenant Routed Multicast (TRM)</li> </ul>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ EVPN - VxLAN ARP/ND flooding suppression.</li> <li>❖ EVPN to MPLS hand off on Cat9K in Border spine role (single box)</li> <li>❖ Hierarchical VPLS</li> <li>❖ VPLS Multiple VCs per Spoke</li> </ul>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ mLDP: Multicast LDP*</li> <li>❖ VPLS Routed PseudoWire (IRB): IPv6 Unicast</li> <li>❖ MVPNv6 (Multicast 6VPE)</li> <li>❖ MPLS VPN - Inter-AS Option AB</li> <li>❖ BGP EVPN w VxLAN BUM rate-limiting support</li> <li>❖ BGP-EVPN w VXLAN MAC/IP learning on Access</li> <li>❖ Wide Area Bonjour with BGP-EVPN over VXLAN</li> </ul>
<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ 9600 - VRF aware PBR</li> <li>❖ 9400 - NAT Profile</li> </ul>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ NAT - VRF aware NAT (VRF to Global)</li> </ul>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ 9500H/9600 : Customized SDM Template Ph1 (FIB)</li> <li>❖ IP-FRRv4: LFA EIGRP and OSPFv2 per prefix</li> <li>❖ Non-Stop Routing (NSR): L3 Forwarding Redundancy</li> <li>❖ LACP 1:1 redundancy and dampening</li> </ul>
<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ 9600 - Quad Sup SVL Support (RPR)</li> <li>❖ 9300 - xFSU Standalone</li> </ul>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ 9300 - xFSU reload with backside stacking*</li> <li>❖ 9300 - xFSU support with dot1x, MAB, Webauth*</li> <li>❖ 9300 - xFSU : LACP Protocol support*</li> </ul>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ Flexlink+ with VLAN Load Balancing</li> <li>❖ 9300 - xFSU Reload: Stacked and Standalone (17.3.2)</li> </ul>
<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ Per port MTU support</li> <li>❖ 9500H/9600 - Unified Port Buffer</li> <li>❖ 9400 - Native Docker for App Hosting</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ 9400 - 9216 bytes MTU</li> <li>❖ 9600 - Breakout Support</li> <li>❖ gPTP/ PTPv2 support on Port Channels</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ PVLAN on Trunks and Port Channels</li> <li>❖ PTPv2 and gPTP support on 9400*</li> <li>❖ ETA and AVC Interoperability on same port</li> <li>❖ SHA-512 secure image-bootup integrity check</li> <li>❖ gRPC Model Driven Telemetry (MDT) with TLS</li> </ul>
<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ C9300L mGig SKUs</li> <li>❖ C9600-LC-48TX mGig Linecard</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ 9500/9600 - AOC/DAC, QSFP-4SFP10G</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ None</li> </ul>

# Catalyst 9000 Switching – Key Features

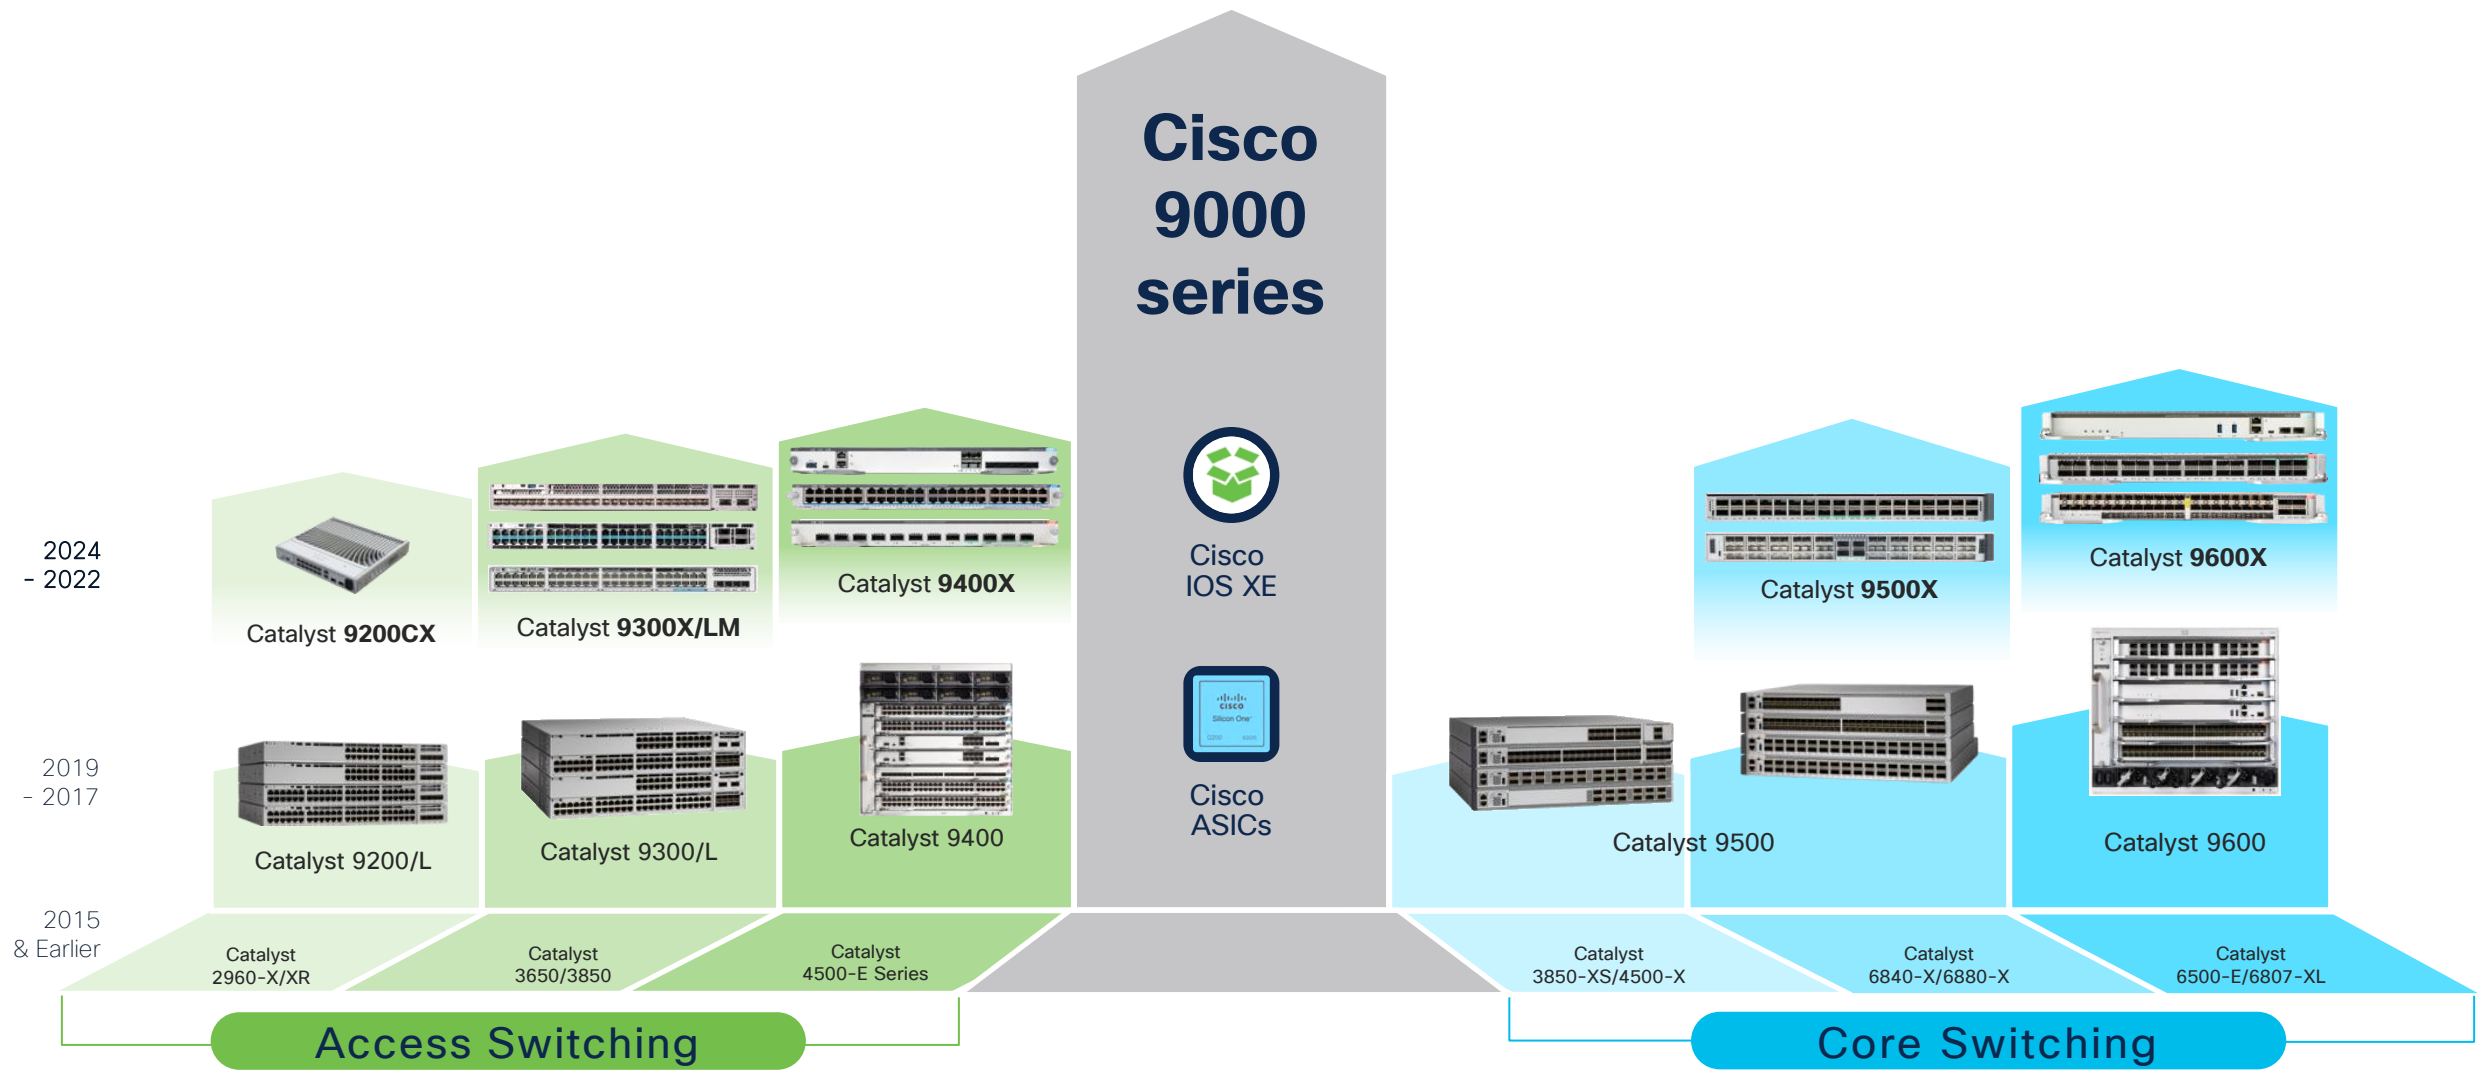
\* Limited Availability (LA) only

<b>IOS XE 17.4.1</b> (Nov'20) <span style="float: right;">SMR</span>	<b>IOS XE 17.5.1</b> (Apr'21) <span style="float: right;">SMR</span>	<b>IOS XE 17.6.1</b> (Aug'21) <span style="float: right;">EMR</span>
<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ FQDN ACLs</li> <li>❖ RADSEC - Radius over TLS and DTLS</li> <li>❖ Stealthwatch Cloud Integration*</li> <li>❖ Wired Client Sensor with Flash*</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ FQDN Redirect ACL</li> <li>❖ Wired Dynamic VLAN</li> <li>❖ Secure Network Analytics Connector</li> <li>❖ DSCP Marking for RADIUS Packets</li> <li>❖ Session timers AV Pair</li> <li>❖ Interface Templates</li> <li>❖ Trustworthy Systems</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ 9300X - IPsec Phase1 - SVTI, IKEv2</li> <li>❖ IPv6 FQDN Redirect ACL</li> <li>❖ RADSEC CoA Enhancement</li> </ul>
<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ PVLAN with BGP EVPN over VxLAN</li> </ul>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ Selective Q-in-Q</li> <li>❖ BGP EVPN L2/L3 VNI scale</li> </ul>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ MPLS Traffic Engineering (TE) - Phase1</li> <li>❖ LACP/PAGP over EoMPLS</li> <li>❖ MLD snooping over VPLS</li> </ul>
<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ 9500H/9600 - Customized SDM Template Ph2 (ACL)</li> </ul>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ 9500H/9600 - Customized SDM Template Ph3 (4K VLAN)</li> <li>❖ Enhanced NAT scale</li> <li>❖ BGP Monitoring Protocol</li> <li>❖ WCCP Over GRE</li> </ul>	<b>Platform Features</b> <ul style="list-style-type: none"> <li>❖ VRF Aware WCCP</li> <li>❖ Enhanced NAT Session Monitoring</li> <li>❖ NAT Precedence</li> <li>❖ Bonjour mDNS SSO, FHRP Service Peer Support</li> </ul>
<b>High Availability</b>	<b>High Availability</b>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ VRRPv3 SSO</li> </ul>
<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ YANG model updates</li> <li>❖ Smart Licensing using Policy</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ gPTP over L3 Unicast</li> <li>❖ Disable USB SSD</li> <li>❖ App Hosting Updates</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ App Hosting on 9300X</li> <li>❖ Thousand Eyes - 4.0 Version Agent</li> <li>❖ Perpetual PoE/UPOE with StackPower</li> <li>❖ PTP on StackWise*, PTP over SDA</li> <li>❖ Programmability &amp; Automation updates</li> </ul>
<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>9500H/9600</b> - SFP-10G-TX, QSFP-40/100-SR4,</li> <li>❖ <b>9600</b> - 4 x 25G Breakout, GLC-GE-100FX and GLC-TE-100M</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>C9400-LC-48HN</b> - 5G MGIG line card with 90W PoE</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>C9300X-12Y / 24Y</b> - 10/25G Fiber Switch with Cisco UADP2.0sec</li> <li>❖ <b>C9300X-NM</b> - 2x 40/100G, 8x 10/25G, 8x mGiG uplinks</li> </ul>

# Cisco Catalyst 9000 Switching Portfolio

One Family from Access to Core – Common Hardware & Software

17.7<sup>·x</sup> – 17.15<sup>·x</sup>




# Cisco IOS XE - Release Schedule

Graphical Overview - 17.7.x to 17.15.x



Second Generation Catalyst 9000 X Series with Cisco UADP & S1



# C9K IOS XE 3.0 - Highlights



**Secure**

- Secure Boot, Image Signing
- SELinux, X.509

**Zero Copy Punt**


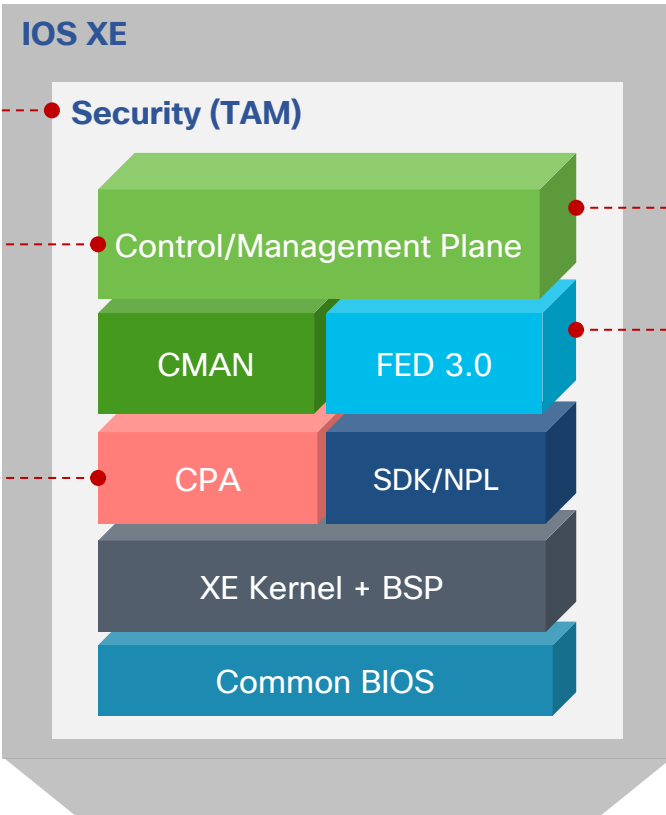
- High Speed zero-copy punt to enable Software-based Apps (e.g. NetFlow)

**CPA - Single Source of Truth, HW Abstraction**

- CPA architecture for sharing common software across multiple platforms
- Single Source of Truth - for various devices and interconnects

**Custom Cisco ASIC + SDK**

- High capacity, programmable ASIC
- Common SDK as an integration layer




**Programmable**

- Managed Through Models
- Programmable through YANG

**FED3.0 - Model Driven Forwarding**

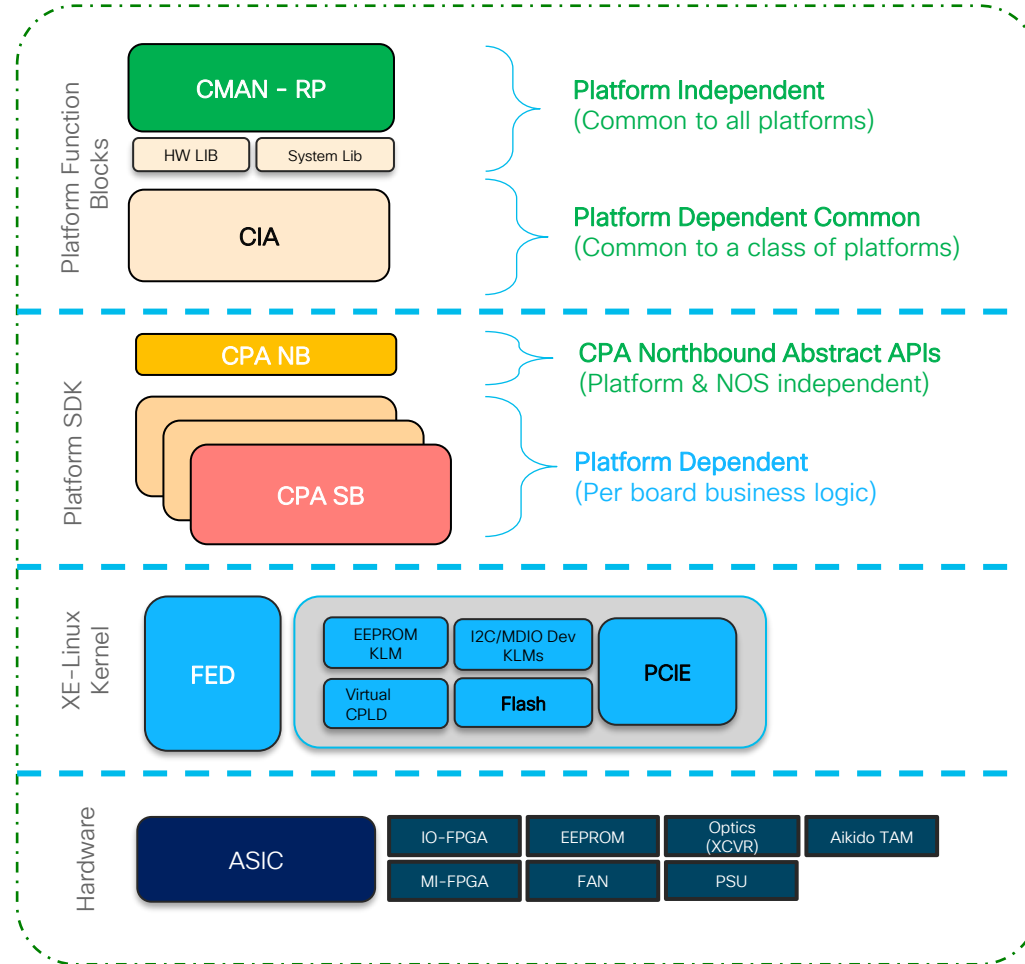
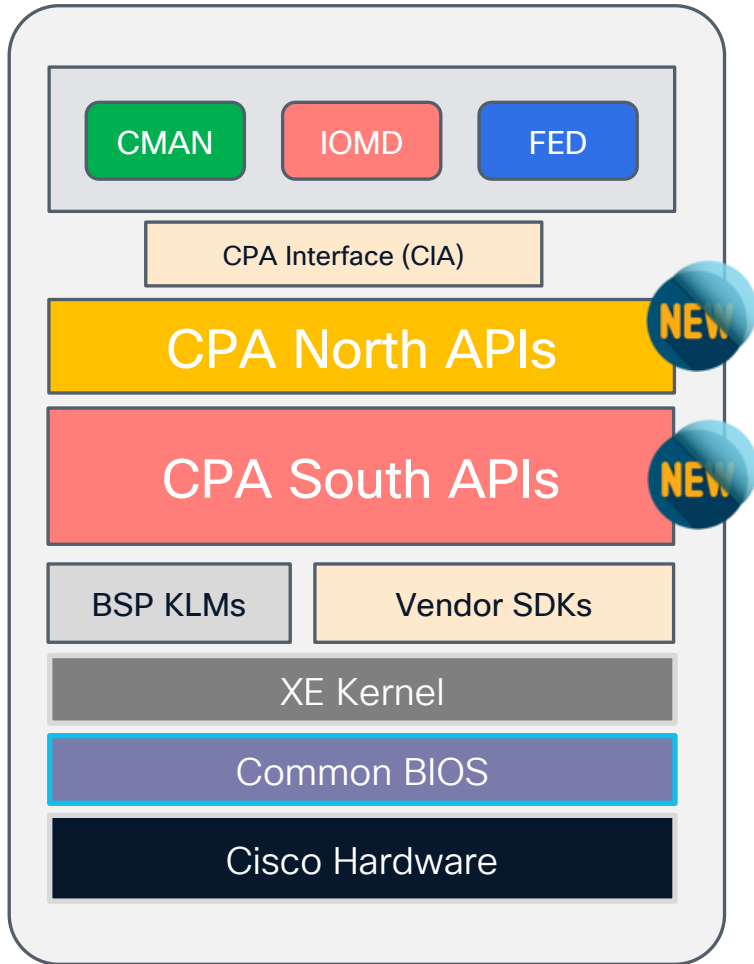
- Bring Polaris infrastructure to FED
- Bring FED closer for stateful restart

**SPECTRA** • Integrated ASIC SDK/NPL Testing



# C9K IOS XE 3.0 - Platform Infra

Architecture Evolution - Common Platform Abstraction (CPA)



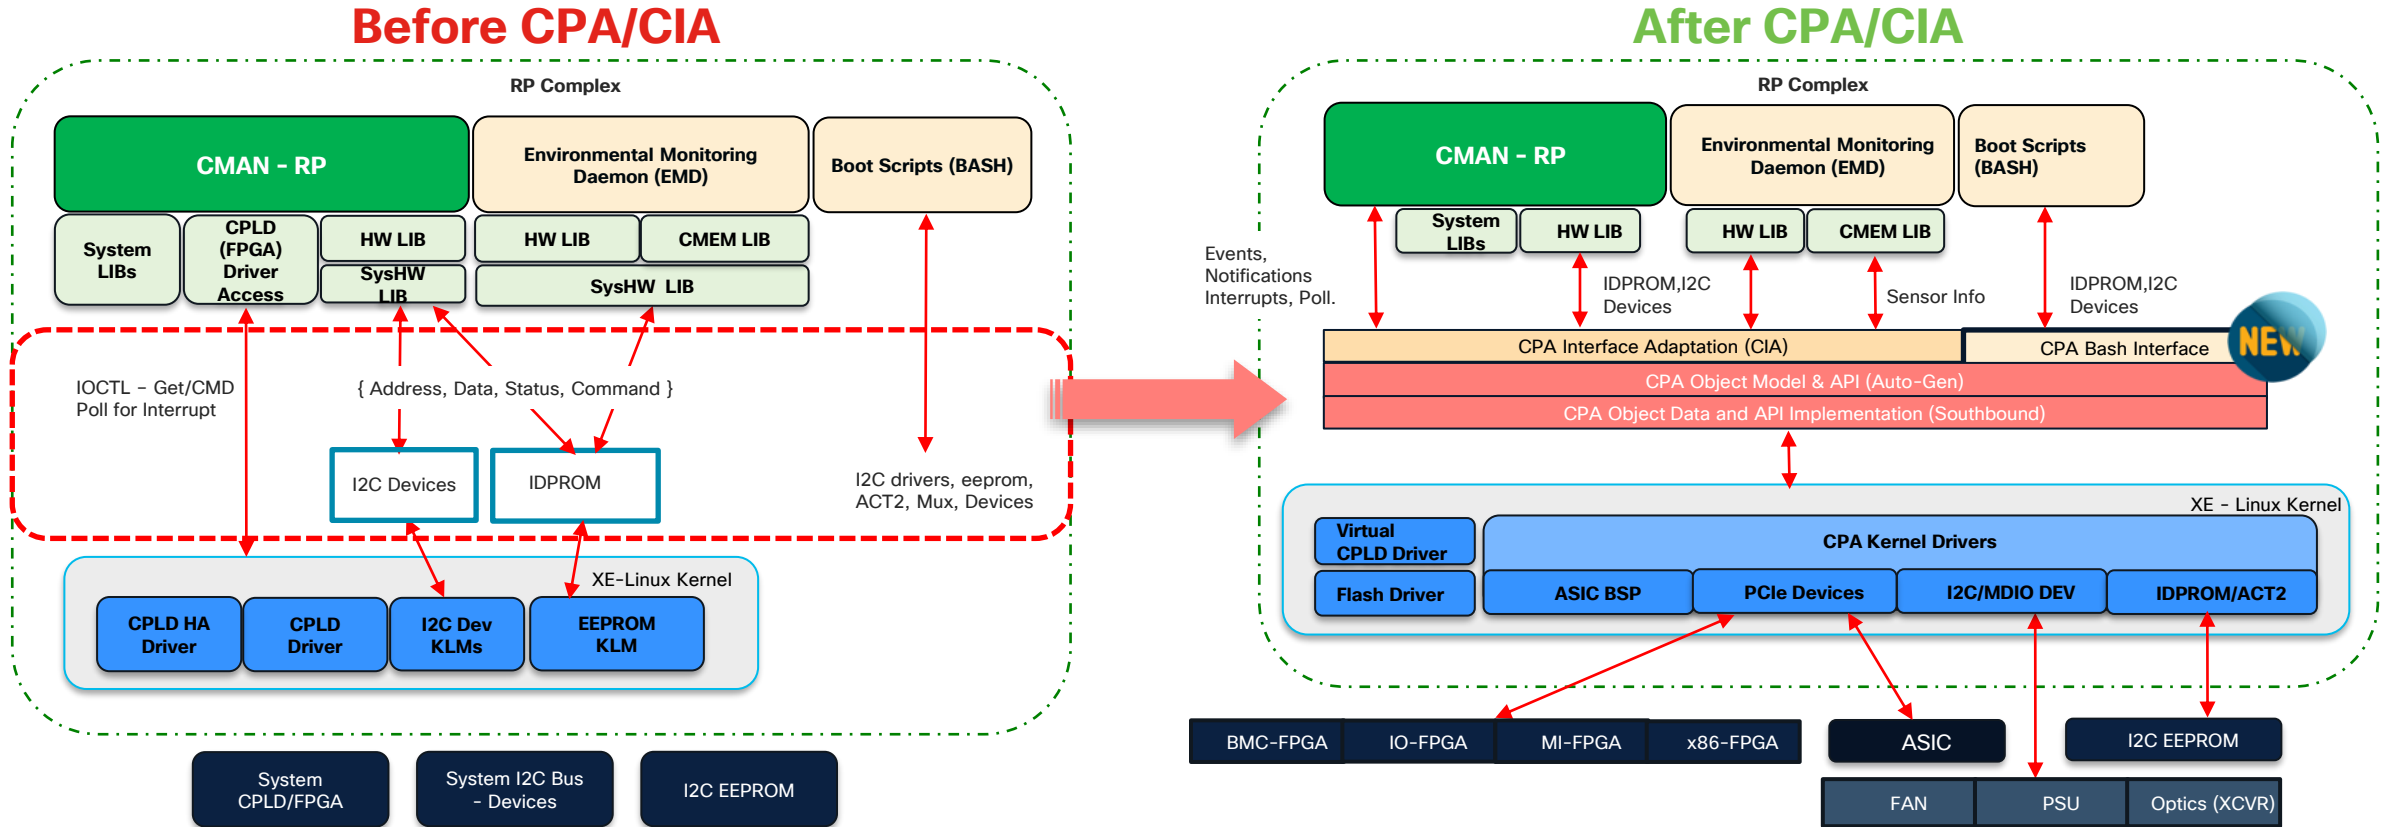
- Leveraged by NG Platforms
- PD code

# C9K NG IOS XE- CPA Overview

## Architecture Evolution - before & after



Cisco Platform Abstraction (CPA) is a model-driven HW to SW disaggregation layer for platform devices. It hides board and device connection details from upper-layers of platform features - acts as a “single source of truth” for hardware data.



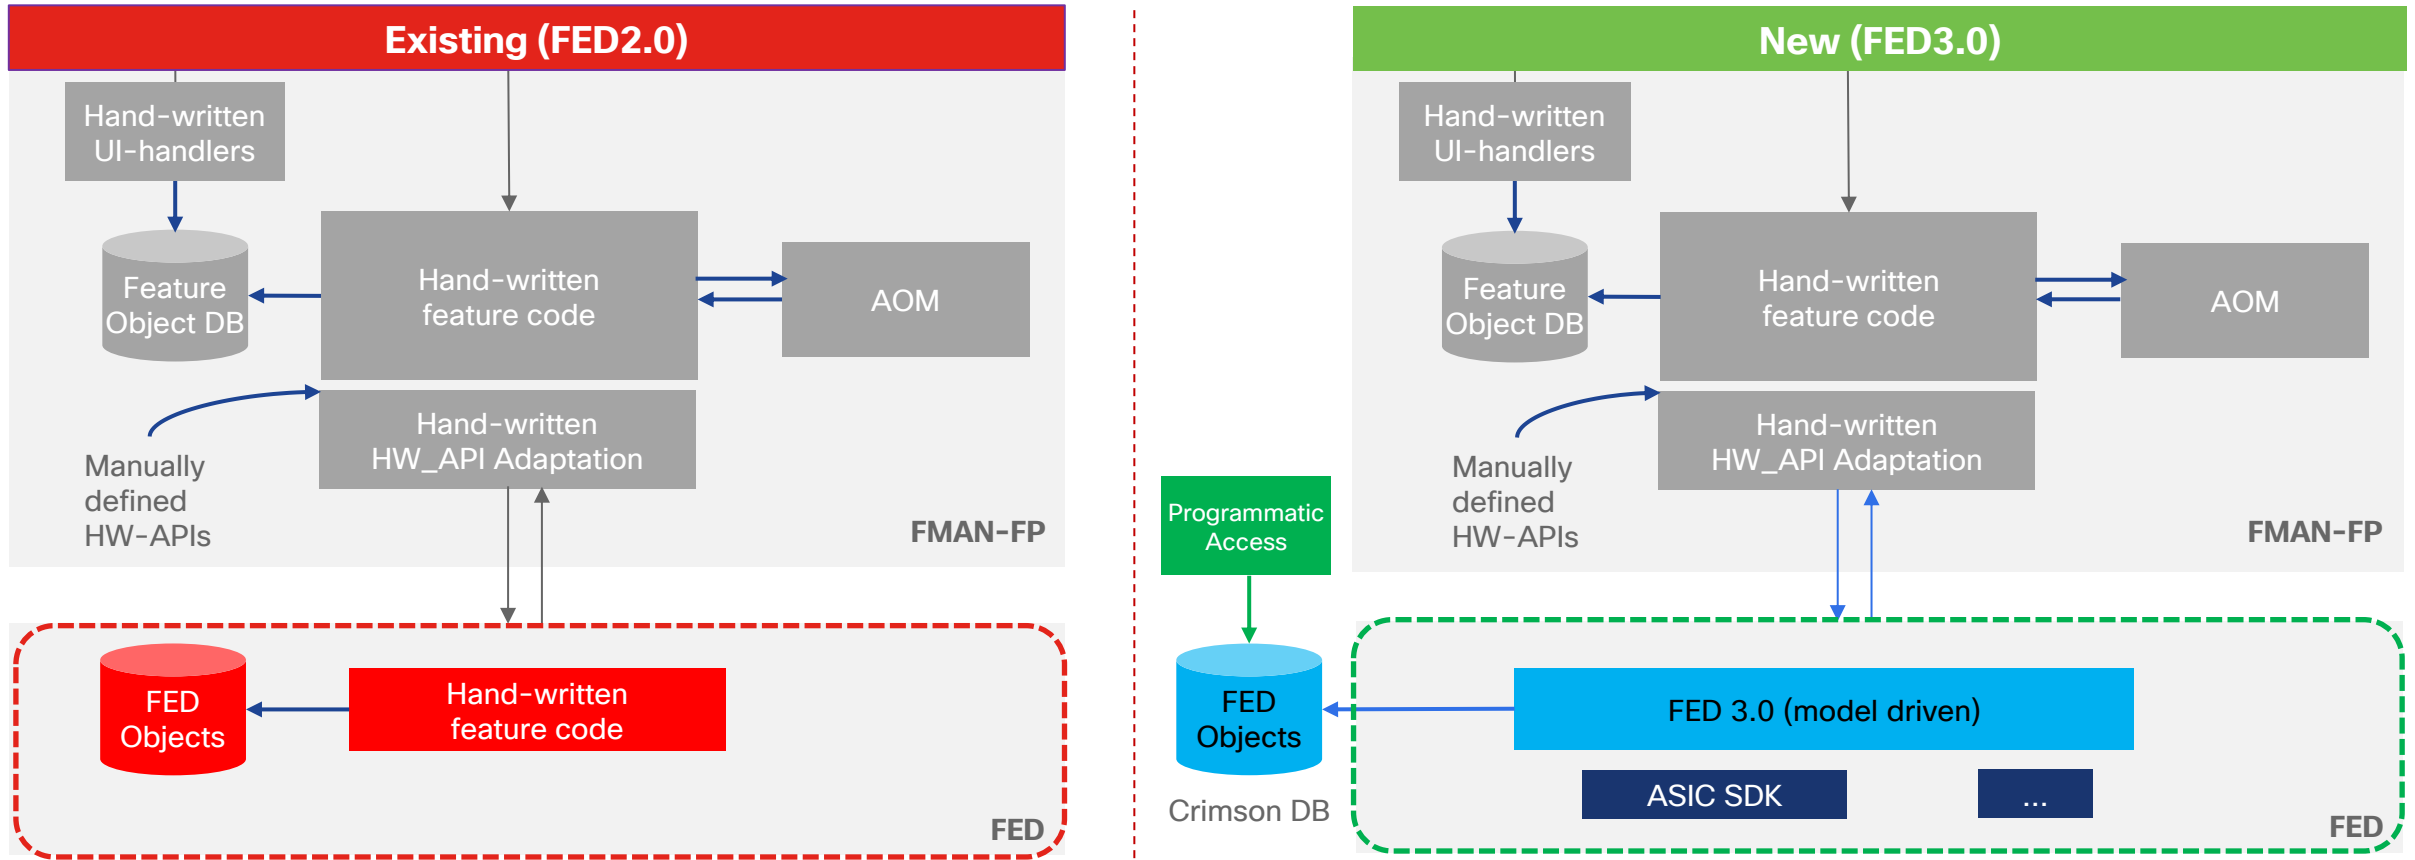
CPA provides common & consistent inventory and device-access helper functions (APIs) for admin-plane features

# C9K NG IOS XE – FED 3.0 Overview



## FED 3.0 – Before & After

IOS XE and FED have been enhanced to support various types of databases and to support automatic code-generation

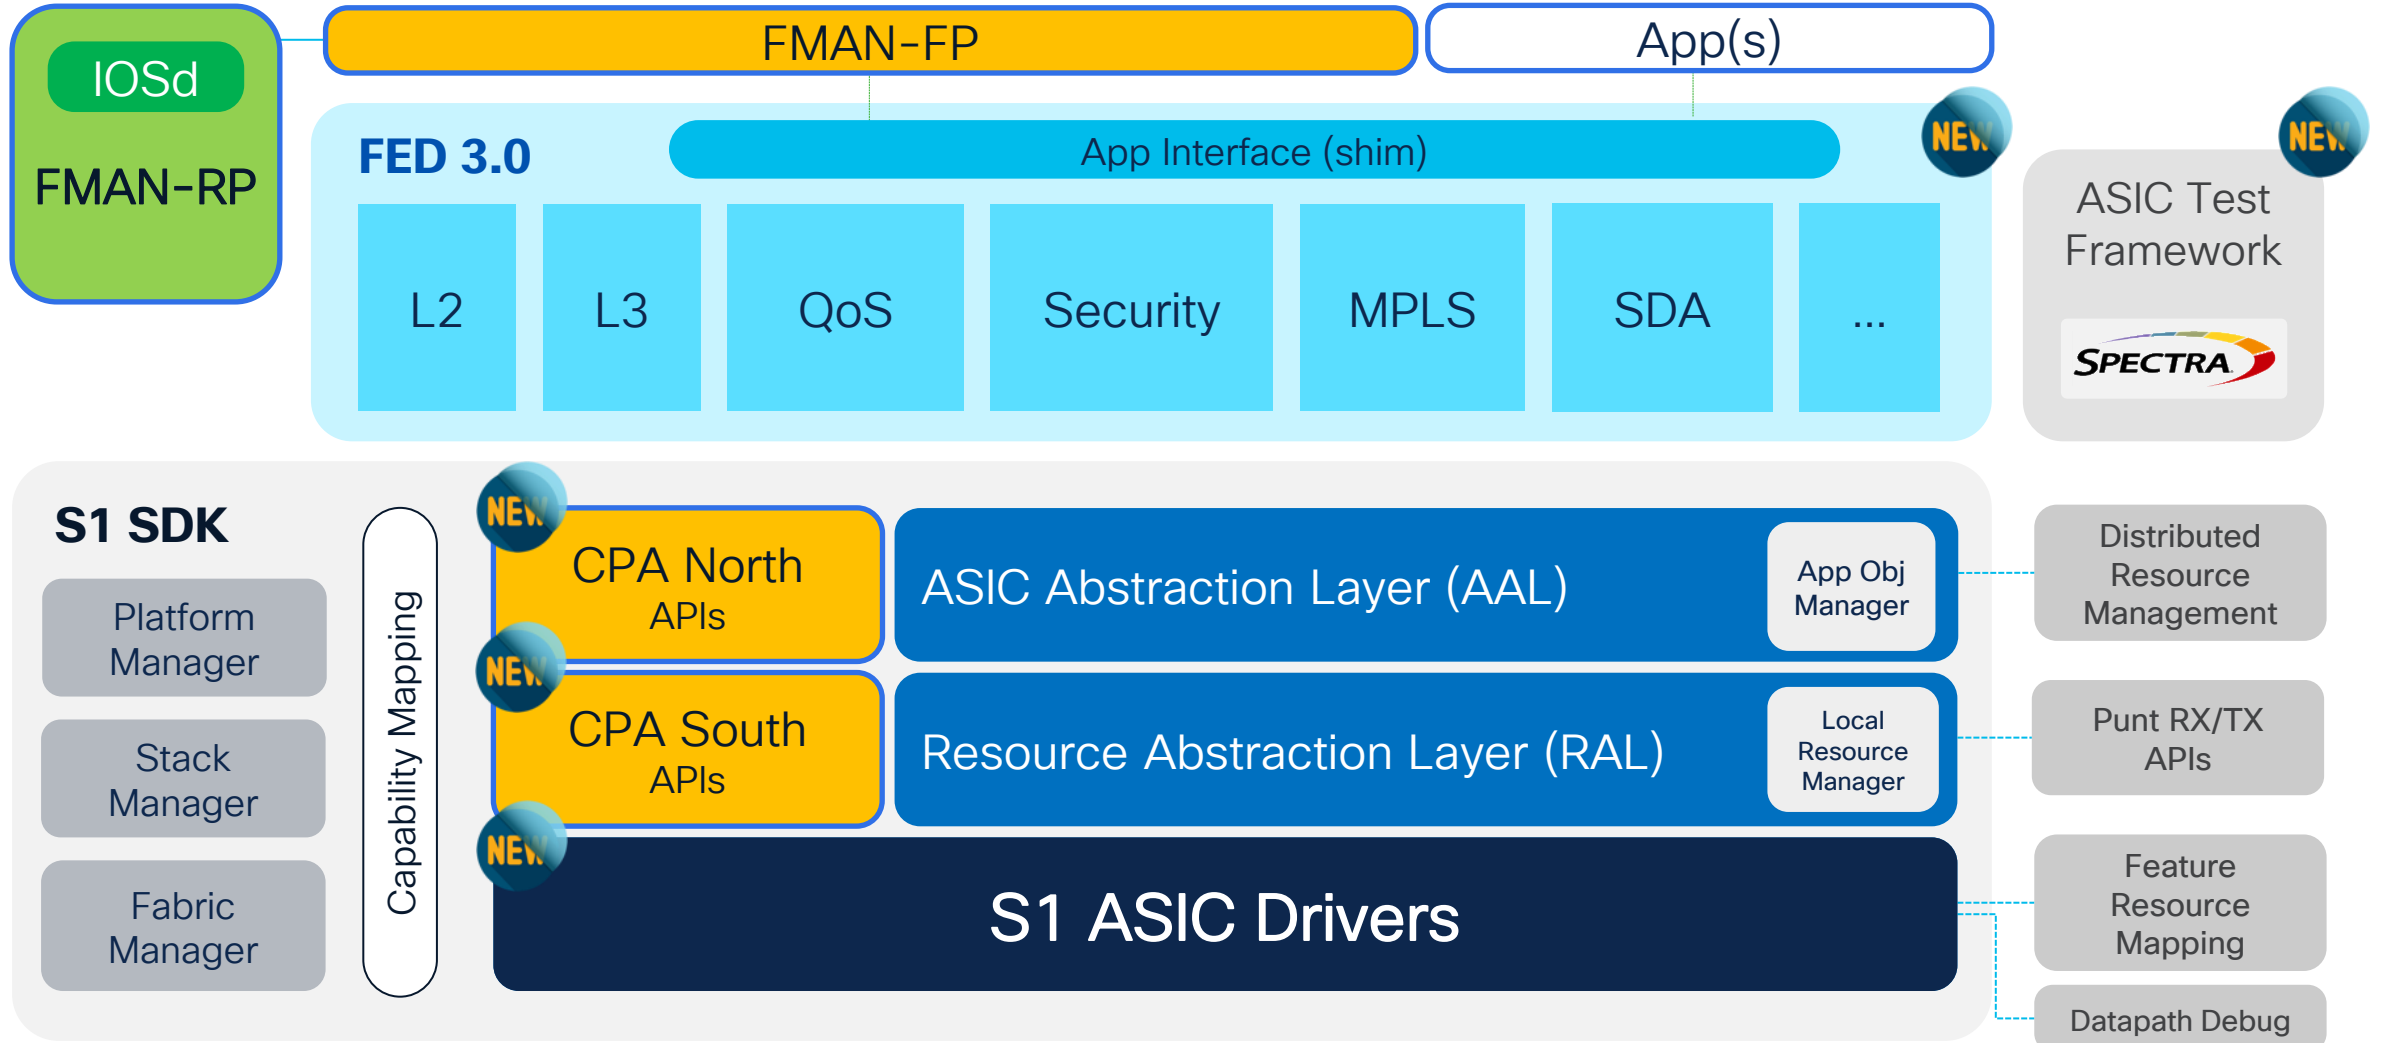


- Controller visibility of HW forwarding objects
- Time series for HW forwarding objects
- Stateful restart HW forwarding objects
- Automatic code-generation for object life-cycle

- Manual
- Auto-generated
- New capability

# C9K NG IOS XE - Overview

## Platform Forwarding Architecture



# Catalyst 9000 Switching – Key Features

\* Limited Availability (LA) only

<b>IOS XE 17.7.1</b> (Dec'21) <span style="float: right;">SMR</span>	<b>IOS XE 17.8.1</b> (Apr'22) <span style="float: right;">SMR</span>	<b>IOS XE 17.9.1</b> (Aug'22) <span style="float: right; color: white;">EMR</span>
<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ 9500X/9600X – MACsec</li> <li>❖ 9200/9300 – API Registration for Umbrella Switch connector</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ 9300X – IPsec Phase2 – Multicast (SVTI)</li> <li>❖ 9500X/9600X – WAN-MACsec, with HSEC license</li> <li>❖ SW SUDI 2099 Enablement</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ 9300X – IPsec Phase3 – NAT Traversal</li> <li>❖ 9300X – VRF-aware IPsec</li> <li>❖ Reflexive ACL</li> </ul>
<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ 9500X/9600X – MPLS and TE Phase1</li> <li>❖ 9500X/9600X – EoMPLS</li> <li>❖ EVPN L3 TRM with MDT Data</li> </ul>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ EVPN L2 TRM</li> </ul>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ SDA LISP Graceful Restart for MAC cache</li> <li>❖ SDA VN Extranet across SDA Transit</li> </ul>
<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ 9500X/9600X – L3 Routing (IGP, BGP) feature set</li> <li>❖ Low priority Control packet mapping to Non-LLQ</li> <li>❖ Bonjour – Micro-Location services</li> </ul>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ 9500X/9600X – Sampled Flexible NetFlow</li> </ul>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ Destination IP NAT scale enhancement</li> <li>❖ PAT support for Enhanced NAT scale</li> <li>❖ Conditional Static NAT using Route-map</li> </ul>
<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ 9400X/9600X – SSO &amp; ISSU</li> <li>❖ 9300X – xFSU</li> </ul>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ 9500H/9600 – Graceful Insertion &amp; Removal (GIR)</li> </ul>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ 9400X – StackWise Virtual (Dual-Sup)</li> </ul>
<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ PTP on 9300 StackWise</li> <li>❖ PTP on 9600</li> <li>❖ PTP AES67 compliance</li> <li>❖ AVNU Certification – 9300 &amp; 9500</li> <li>❖ gNOI reset.proto – tooling</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ PTP – G8275.1 ITU Telecom Profile on 9300/9300X</li> <li>❖ 9500X/9600X – L3 Sub-Interface Queuing</li> <li>❖ C9300 System Power-Consumption Reporting</li> <li>❖ gNMI Native Configuration Yang Model</li> <li>❖ Guest Shell HA – Guest-Share Folder Sync</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ 9500H – AVB support</li> <li>❖ 9400X – Perpetual PoE support</li> <li>❖ 9400X – Support for hosting multiple applications</li> <li>❖ 9400X – 432 Port-Channels</li> <li>❖ 9400X – 4K VLANs support</li> </ul>
<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>C9300X-48TX / 48HX</b> – 48x mGig Switch (Cisco UADP2.0sec)</li> <li>❖ <b>C9400X-SUP-2/XL</b> – Supervisor 2 (Cisco UADP3.0sec)</li> <li>❖ <b>C9500X-28C8D</b> – 100/400G Fiber switch (Cisco S1 Q200)</li> <li>❖ <b>C9600X-SUP-2</b> – Supervisor 2 (Cisco S1 Q200)</li> <li>❖ <b>C9600-LC-40YL4CD</b> – 40x SFP + 4x QSFP Combo Linecard</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>C9400-LC-48HX</b> – 48x mGig Linecard</li> <li>❖ <b>C9400-LC-48XS</b> – 48x 1/10G Linecard</li> <li>❖ <b>C9600-PWR-3KW</b> – 3000W AC PSU</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>C9200CX</b> – 12x Compact switch (Cisco UADP2.0mini)</li> <li>❖ <b>C9300LM</b> – 24/48x + Fixed Uplinks, ≤14 in. (Cisco UADP2.0sec)</li> <li>❖ <b>C9600-LC-32CD</b> – 32x 100G / 24x 100G + 2x 400G QSFP Linecard</li> </ul>

Products and Services Solutions Support Learn

---

Products & Services / Cisco IOS and NX-OS Software / Cisco IOS / End-of-Life and End-of-Sale Notices /

**End-of-Sale and End-of-Life  
Announcement for the Cisco IOS XE  
17.6.x**

# Catalyst 9000 Switching – Key Features

\* Limited Availability (LA) only


<b>IOS XE 17.10.1</b> (Dec'22) <span style="float: right;">SMR</span>	<b>IOS XE 17.11.1</b> (Apr'23) <span style="float: right;">SMR</span>	<b>IOS XE 17.12.1</b> (Aug'23) <sup>^</sup> <span style="float: right;">EMR</span>
<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ DHCP Snooping Glean</li> <li>❖ Reflexive ACL for IPv4</li> <li>❖ MACsec Transparent Pass-Through</li> <li>❖ Secure Data Wipe (NIST 3-pass)</li> <li>❖ 9400X – IPsec support</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ 9300X - GRE over IPsec support</li> <li>❖ 9400X - NAT-Traversal support for IPSEC</li> <li>❖ 9500X/9600X - IPv6 SGACL support</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ IC2M updates for FIPS 140-3</li> <li>❖ IPSEC OSFPv3 Auth support</li> <li>❖ Service Redirect (Firewall PBR)</li> <li>❖ 9500X-60L4D - IPsec support (P1*)</li> </ul>
<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ BGP EVPN over IPv6 Underlay*</li> <li>❖ BGP EVPN Dynamic Peering</li> <li>❖ BGP EVPN over IPsec - TRM support</li> </ul>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ BGP EVPN Micro-segmentation (VLXAN-GPO)</li> <li>❖ BGP EVPN Mobility-Convergence enhancements*</li> <li>❖ 9500X/9600X - BGP EVPN with SVL support</li> <li>❖ 9500X/9600X - MPLS VPN - Inter-AS Option A</li> <li>❖ 9500X/9600X - L2VPN (EoMPLS) Pseudowire Redundancy</li> </ul>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ BGP EVPN (VRF Aware) with Dynamic NAT64</li> <li>❖ BGP EVPN Config and Operation CLI simplicity</li> <li>❖ Increased VNI Scale for EVPN</li> <li>❖ 9500X/9600X - EVPN L2 TRM support</li> </ul>
<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ Stateful NAT64</li> <li>❖ NAT on L3 EtherChannel</li> <li>❖ LACP Standalone Mode on L3 EtherChannels</li> <li>❖ PTPv2 on StackWise Virtual</li> <li>❖ 9500X/9600X - Bonjour &amp; mDNS Routing</li> </ul>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ 9500X/9600X - Policy Based Routing (PBR)</li> <li>❖ 9200/9200CX - Bonjour Service Peer</li> </ul>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ PTPv2 on 9400X</li> </ul>
<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ 9500X/9600X - StackWise Virtual (Dual Sup SSO)</li> </ul>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ StackWise Virtual (Dual Sup SSO) on 9500X-60L4D</li> </ul>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ Standby Port Delay for SVL Bringup</li> </ul>
<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ ERSPAN on AppGig interface</li> <li>❖ YANG 1.1 Support</li> <li>❖ YANG Config Model support for native gNMI</li> <li>❖ Guest Shell HA Guest-Share Folder Sync</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ 9400X - Customized SDM Template</li> <li>❖ 9500X/9600X - ERSPAN</li> <li>❖ YANG Oper model support for PTPv2 and AVB</li> <li>❖ Switching Telemetry for Cisco DNA Center*</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ EPC Support of AppGigabitEthernet</li> <li>❖ Upgrade firmware (CPLD/FPGA/Rommon) with IOSXE in single reload</li> <li>❖ Bootup time enhancements</li> </ul>
<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>C9500X-60L4CD</b> - 60x 50G + 4x 400G switch (Cisco S1 Q200)</li> <li>❖ 9600X - 50G SFP Linecard support</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>9500X/9600X</b> - Embedded-PHY 1G SFP support *</li> <li>❖ Support for QSFP-100G-FR-S</li> <li>❖ Support for SFP-10/25BXD-1 and BXU-1</li> <li>❖ Support for QSFP-4x10G on 100G PSM4</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>C9400-LC-24XY</b> - 24x 10/25G Linecard</li> <li>❖ <b>C9400-LC-12QC</b> - 12x 40/100G Linecard</li> <li>❖ <b>C9200CX-12PD-2X2G &amp; 8PD-2G</b></li> <li>❖ 9500X/9600X - 50G SR-S/SL/Multirate support</li> </ul>

# Catalyst 9000 Switching – Key Features

\* Limited Availability (LA) only

Reference

<b>IOS XE 17.13.1</b> (Dec'23) <span style="float: right;">SMR</span>	<b>IOS XE 17.14.1</b> (Apr'24) <span style="float: right;">SMR</span>	<b>IOS XE 17.15.1</b> (Aug'24) <span style="float: right;">EMR</span>
<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ C9400X - GRE over IPSEC with VRF support</li> <li>❖ SGACL NetFlow drop action record</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ C9300X/C9400X - IPSEC Scale enhancements</li> <li>❖ Secure encryption for IP sftp password *</li> </ul>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ 9500X/9600X - MACsec/WAN-MACsec support on Port-Channel Sub-interfaces</li> <li>❖ AAA PAC-less Authentication</li> <li>❖ IOSXE SELinux Enforced Mode</li> </ul>
<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ VXLAN with IPv6 Underlay for Unicast &amp; Multicast</li> <li>❖ 9500X/9600X - Cisco TrustSec SGT inline tagging</li> </ul>	<b>Overlays &amp; Segmentation</b>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ BGP EVPN CLI Simplification</li> <li>❖ Multi-Cluster Fabric - Router MAC Rewrite with NextHop-Self BGP Attribute</li> </ul>
<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ BGP feature support on C9200CX</li> <li>❖ Inter-AS option B MVPN support</li> <li>❖ IPv6 Proxy Neighbor Discovery</li> </ul>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ PTP Support for C9200CX</li> <li>❖ IP SLA Probe Config modification via config replace</li> </ul>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ Link Debounce Timer</li> <li>❖ BGP Router-ID compliance with RFC6286</li> </ul>
<b>High Availability</b>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ NAT SSO with StackWise Virtual *</li> </ul>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ eXtended Fast Sw Upgrade (xFSU) ≤ 5 seconds - 17.15.2</li> </ul>
<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ Map SNMP mibs to YANG xpaths</li> <li>❖ AAA server MIB Support</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ Energy Consumption Visibility for 9200, 9300 and 9400</li> <li>❖ gNMI stream subscriptions with ON_CHANGE mode</li> <li>❖ AWS S3 and CloudWatch Integration</li> <li>❖ YANG Model Support for L2NAT</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ Atomic Config Replace (ACR) Phase 1* - 17.15.2</li> <li>❖ Enhanced Energy Monitoring</li> <li>❖ Shutdown Fiber Ports when no SFP Inserted</li> <li>❖ Shutdown Power Supplies in C9300 StackPower during Low Load</li> </ul>
<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>C9600X-56YL4C</b> - 56x 50G + 4x 100G Linecard</li> <li>❖ <b>C9400-LC-48TX</b> - 48x 10G Copper (Data-Only) Linecard</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ 1G Support in C9500X/C9600X</li> <li>❖ New optics support</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>C9400X-SUP-2/2XL</b> - 1G support with transceivers</li> </ul>



[www.cisco.com/c/en/us/support/docs/switches/.../214814-recommended-releases-for-catalyst-9200-9.html](http://www.cisco.com/c/en/us/support/docs/switches/.../214814-recommended-releases-for-catalyst-9200-9.html)


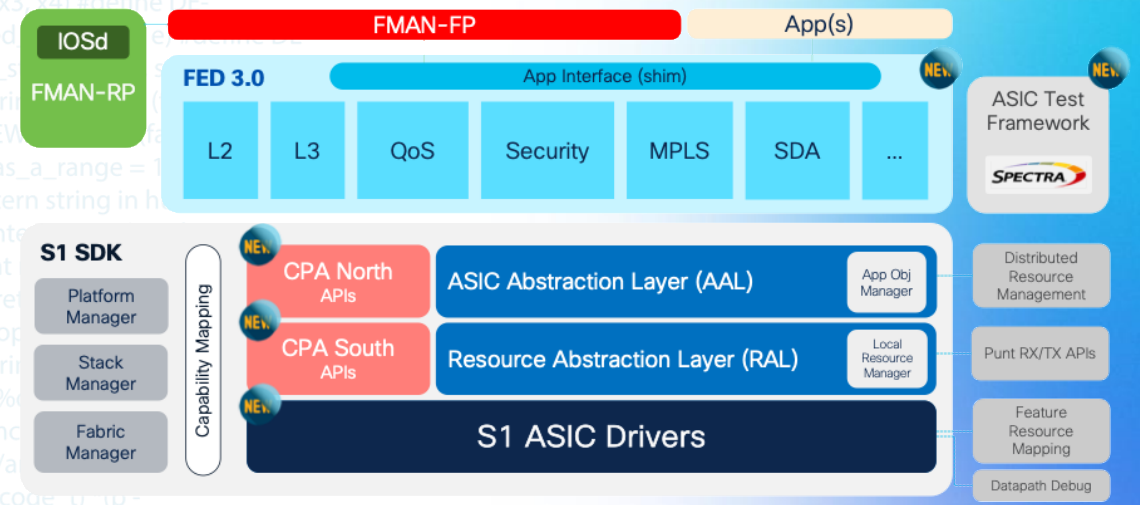
```

extract_number_and_incr (destination, source) int
*destination; unsigned char **source; { extract_number
number (destination, *source); *source += 2; } #ifndef EXTRACT_MAC
ROS #undef EXTRACT_NUMBER_AND_INCR #define EXTRACT_NUM
BER_AND_INCR(dest, src) \ extract_number_and_incr (&dest, &src) #endif /*
not EXTRACT_MACROS */ #endif /* DEBUG */ /* If DEBUG is defined, Regex prints
many voluminous messages about what it is doing (if the variable `debug' is nonzero). If
linked with the main program in `iregex.c'; you can enter patterns and strings interactively.
And if linked with the main program in `main.c' and the other test files, you can run the al-
ready-written tests. */ #ifdef DEBUG /* We use standard I/O for debugging. */ #include <stdio.h>
/* It is useful to test things that `must' be true when debugging. */ #include <assert.h> static int
debug = 0; #define DEBUG_STATEMENT(e) e #define DEBUG_PRINT1(x) if (debug) printf (x) #define
DEBUG_PRINT2(x1, x2) if (debug) printf (x1, x2) #define DEBUG_PRINT3(x1, x2, x3) if (debug) printf
(x1, x2, x3) #define DEBUG_PRINT4(x1, x2, x3, x4) if (debug) printf (x1, x2, x3, x4) #define DF
BUG_PRINT_COMPILED_PATTERN(p, s, e) \ if (debug) print_partial_compiled
BUG_PRINT_DOUBLE_STRING(w, s1, sz1, s2, sz2) \ if (debug) print_double_s
extern void print_compiled_pattern(const char *start, e) { unsigned char *end; { int
signed char *p = start; unsigned char *pend = end; if (start == NULL) { printf ("(null)\n"); re
pattern commands. */ while (p < pend) { switch ((re_opcode_t) *p++) { case no_op
while (--mcnt); break; case start_memory: mcnt = *p++; printf ("/start_memory/%
*p++; break; case stop_memory: mcnt = *p++; printf ("/stop_memory/%d/%d", mc
break; case duplicate: printf ("/duplicate/%d", *p++); break; case anychar: printf ("/a
break; case charset: case charset_not: { register int c; printf ("/charset%s", (re_opcode_t) (p
1) == charset_not ? "_not" : ""); assert (p + *p < pend); for (c = 0; c < *p; c++) { unsigned bit;
unsigned char map_byte = p[1 + c]; putchar ('/'); for (bit = 0; bit < BYTEWIDTH; bit++) if
(map_byte & (1 << bit)) putchar (c * BYTEWIDTH + bit); } p += 1 + *p; break; } case beg-
line: printf ("/begline"); break; case endline: printf ("/endline"); break; case on_failure_
jump: extract_number_and_incr (&mcnt, &p); printf ("/on_failure_jump/0/%d", mcnt);
break; case on_failure_jump_past_alt: extract_number_and_incr (&mcnt, &p); printf
("/on_failure_jump_past_alt/0/%d", mcnt); break; case dummy_failure_jump: ex-
tract_number_and_incr (&mcnt, &p); printf ("/dummy_failure_jump/0/%d", mcnt); break; case may-
be_pop_jump: extract_number_and_incr (&mcnt, &p); printf ("/maybe_pop_jump/0/%d", mcnt); break; case pop_failure_
jump: extract_number_and_incr (&mcnt, &p); printf ("/pop_
failure_jump/0/%d", mcnt); break; case jump_past_alt:
extract_number_and_incr (&mcnt, &p); printf ("/-

```

- 1  
Why  
IOS XE?
- 2  
IOS XE  
Design
- 3  
Up to  
17.15.x
- 4  
After  
17.18.x
- 5  
IOS XE  
Features
- 6  
Summary

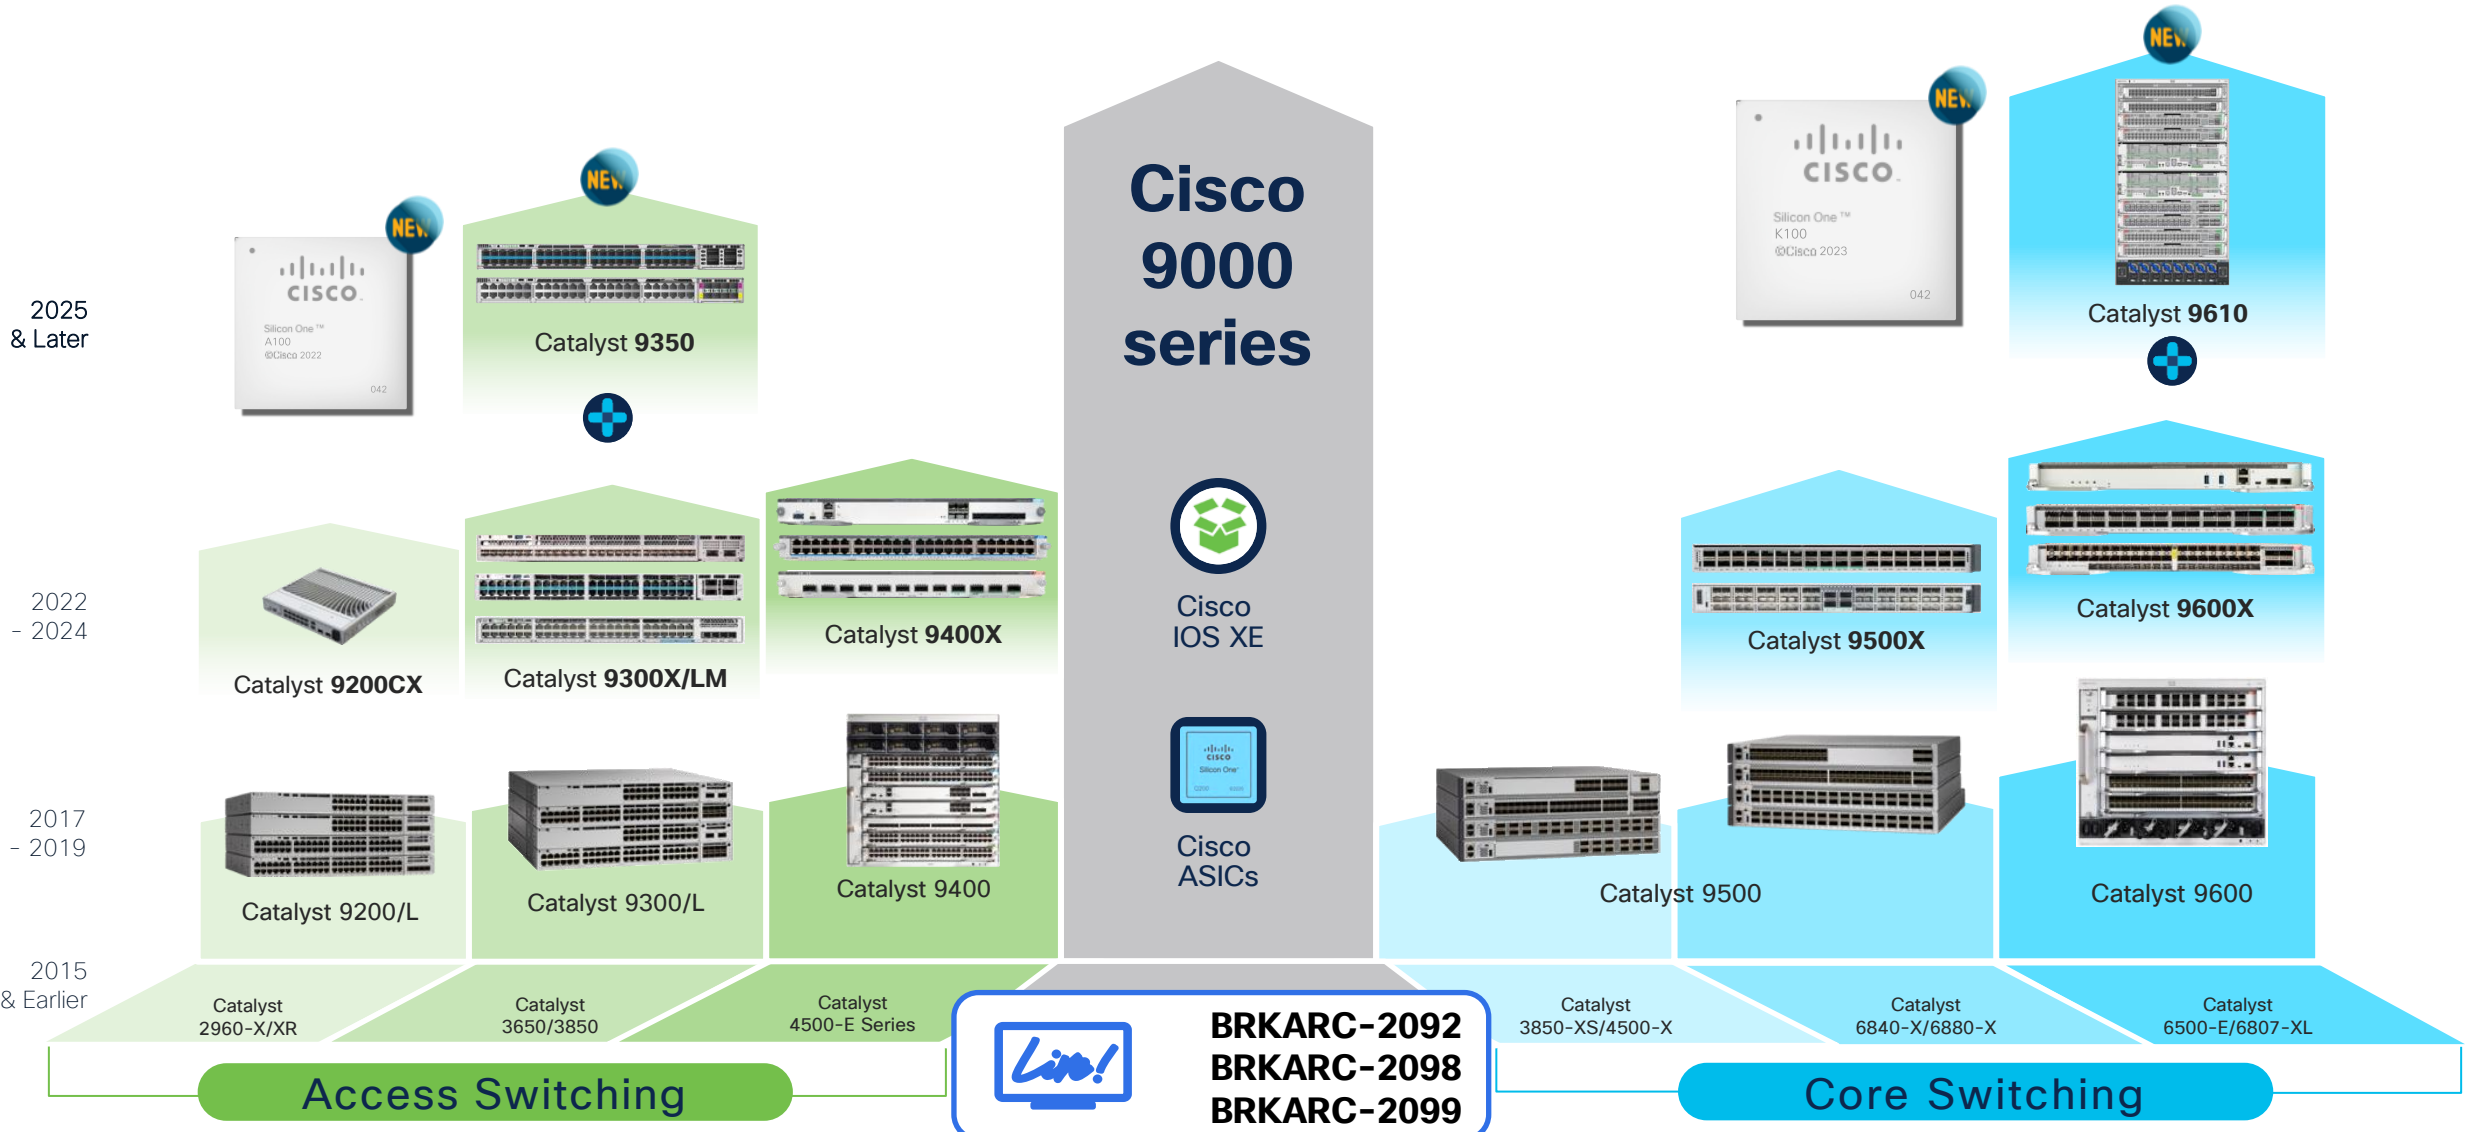
# New Cisco 9000 Series IOS XE 17.16.1 - 17.18.1 Cloud-Native IOS XE



# Cisco Catalyst 9000 Switching Portfolio

One Family from Access to Core – Common Hardware & Software

2025+ **NEW**




# Cisco IOS XE - Release Schedule

Graphical Overview - 17.16.x to 17.18.x



Third Generation  
Cisco C9000 Series  
with Cisco S1




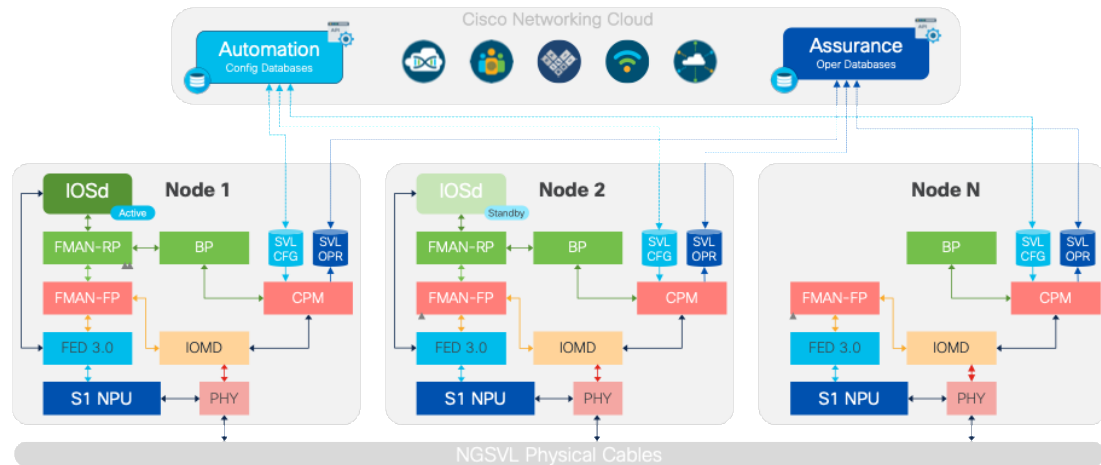
# Introducing Cloud-Native IOSXE

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




## Building a revolutionary NOS architecture with 30 years of experience!

**Cloud-Native IOSXE** builds on our existing infrastructure abstracting & modularizing internal software processes, with a focus on 'model-driven' objects and standard databases, to enable advanced **programmability, clustering & resiliency**.




## Common Platform Abstraction

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 Software & Licensing (UNX)

Cisco C9K platforms support Unified Management controller models (cloud and/or on-premises) and support a single, common Unified Licensing model.



## Cloud-Native Management & AI/ML

Automation, Assurance & AI-Ops models will support entirely cloud-based, entirely on-premise (e.g. air-gap) or a hybrid of both. Available through a Unified Onboarding.



## NG-Stacking & Cluster Infrastructure

Current StackWise and SVL (protocol & encapsulation) have been redesigned using standard VXLAN based protocols & 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.



# Catalyst 9000 Switching – Key Features

Shipping



\* Limited Availability (LA) only

<b>IOS XE 17.16.1</b> (Dec'24) SMR	<b>IOS XE 17.17.1</b> (Apr'25) SMR	 <b>IOS XE 17.18.1</b> (Aug'25) EMR
<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ CTS: Multiple Static SGTs on interface</li> </ul>	<b>Enhanced Security</b>	<b>Enhanced Security</b> <ul style="list-style-type: none"> <li>❖ TWS enhancements (PQC TAM)</li> <li>❖ SELinux with eBPF</li> <li>❖ TLS 1.3 for HTTPS, LDAP, RADsec, Syslog, etc.</li> </ul>
<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ Multi-cluster BGP EVPN VXLAN Fabric : Router MAC rewrite with Next-Hop Self BGP attribute</li> </ul>	<b>Overlays &amp; Segmentation</b>	<b>Overlays &amp; Segmentation</b> <ul style="list-style-type: none"> <li>❖ EVPN Multi-Homing (fabric mode)</li> <li>❖ EVPN Policy-Based Routing (PBR) with Next-Hop Self</li> <li>❖ LISP EID-level Pub-Sub</li> <li>❖ Security Service Insertion (SSI)</li> <li>❖ Multicast VPN (mVPN) on C9500X, C9600X-SUP-2</li> </ul>
<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ Multicast Flow-Aware SG-timer</li> <li>❖ ECMP support with NAT Scale</li> <li>❖ STP bridge assurance per-interface</li> </ul>	<b>Forwarding &amp; Features</b>	<b>Forwarding &amp; Features</b> <ul style="list-style-type: none"> <li>❖ L2, L3, Multicast for C9350 &amp; C9610</li> <li>❖ ACLs, QoS &amp; FNF for C9350 &amp; C9610*</li> <li>❖ Inter-AS Option B on C9500X, C9600X-SUP-2</li> </ul>
<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ Quad-Sup StackWise Virtual on 9600X-SUP-2</li> <li>❖ Enhanced XFSU for 9300/9300X (≤ 5 sec)</li> </ul>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ Uplink Sequencing with C9400</li> </ul>	<b>High Availability</b> <ul style="list-style-type: none"> <li>❖ SSO for C9350 Stacking, C9610-SUP-3</li> <li>❖ NG-Stackwise for C9350 &amp; C9610*</li> </ul>
<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ C9200CX ZTP: External Connectivity</li> <li>❖ FQDN Enhancements</li> <li>❖ Auto-off switchport LED C9200, C9300</li> <li>❖ Auto-off C9200, C9300 Fiber Ports when no SFP Inserted</li> <li>❖ Auto-off C9300 Power Supplies in StackPower during Low Load</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ LLDP support on C9300L management port for Zero-Touch Provisioning (ZTP)</li> <li>❖ C9400: 2100AC Titanium PSU Support</li> </ul>	<b>Platform &amp; Programmability</b> <ul style="list-style-type: none"> <li>❖ Meraki Dashboard for C9200/L, C9500</li> <li>❖ Unified Management for C9350 &amp; C9610*</li> <li>❖ Unified Licensing for C9350 &amp; C9610*</li> <li>❖ Unified Onboarding (Native TLS)</li> <li>❖ Enhanced Energy Monitoring – Smart Power</li> <li>❖ Atomic Config Replace (ACR) Phase 1 – GA</li> <li>❖ OpenConfig support for L2/L3, Multicast &amp; EVPN</li> </ul>
<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ QSFP-100G-FR-S support on C9400X-SUP-2/XL</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ QSFP-100G-LR4-S support on C9400X-SUP-2/XL</li> <li>❖ QSFP-100G-LR4-S support on C9600-LC-40YL4CD &amp; LC-32CD</li> </ul>	<b>Hardware &amp; Optics</b> <ul style="list-style-type: none"> <li>❖ <b>Cisco C9350 Series</b> – Phase 1</li> <li>❖ <b>Cisco C9610 Series</b> – SUP-3/XL, Native LCs, LC-ADPT</li> </ul>

```

extract_number_and_incr (destination, source) int
*destination; unsigned char **source; { extract_number
(destination, *source); *source += 2; } #ifndef EXTRACT_MAC
ROS #undef EXTRACT_NUMBER_AND_INCR #define EXTRACT_NUM
BER_AND_INCR(dest, src) \extract_number_and_incr (&dest, &src) #endif /*
not EXTRACT_MACROS */ #endif /* DEBUG */ #if defined, Regex prints
many voluminous messages about what it is doing (if the variable `debug' is nonzero). If
linked with the main program in `iregex.c'; you can enter patterns and strings interactively.
And if linked with the main program in `main.c' and the other test files, you can run the al-
ready-written tests. */ #ifdef DEBUG /* We use standard I/O for debugging. */ #include <stdio.h>
/* It is useful to test things that `must' be true when debugging. */ #include <assert.h> static int
debug = 0; #define DEBUG_STATEMENT(e) e #define DEBUG_PRINT1(x) if (debug) printf (x) #define
DEBUG_PRINT2(x1, x2) if (debug) printf (x1, x2) #define DEBUG_PRINT3(x1, x2, x3) if (debug) printf
(x1, x2, x3) #define DEBUG_PRINT4(x1, x2, x3, x4) if (debug) printf (x1, x2, x3, x4) #define DE
BUG_PRINT_COMPILED_PATTERN(p, s, e) \if (debug) print_partial_compiled_pattern (s, e)
DEBUG_PRINT_DOUBLE_STRING(w, s1, sz1, s2, sz2) \if (debug) print_double_string (w, s1, sz
1, s2, sz2) #endif #define DEBUG_PRINT_FASTMAP(f, s) \if (debug) print_fastmap (f, s) #endif
void printchar(); /* Print the fastmap in human-readable form. */ void print_fastmap (fa
stmap) { unsigned was_a_range = 0; unsigned i = 0; while (i < (1 << BYTEWIDTH)) { if (fastmap[i]
was_a_range) { printf ("-"); printchar (i - 1); } } putchar ('\n'); } /* Print a compiled patter
n-readable form, starting at the START pointer into it and ending just before the point
pattern commands. */ while (p < pend) { switch ((re_opcode_t) *p++) { case no_op:
while (--mcnt); break; case start_memory: mcnt = *p++; printf ("/start_memory/%d", mcnt);
*p++; break; case stop_memory: mcnt = *p++; printf ("/stop_memory/%d/%d", mcnt, *p);
*p++; break; case duplicate: printf ("/duplicate/%d", *p); *p++; break; case anychar: printf ("/
register int c; printf ("/charset%s", (re_opcode_t) *(p-
1)) == charset_not ? "_not " : ""; assert (p + *p < pend); for (c = 0; c < *p; c++) { unsigned bit;
unsigned char map_byte = p[1 + c]; putchar ('/'); for (bit = 0; bit < BYTEWIDTH; bit++) if
(map_byte & (1 << bit)) printchar (c * BYTEWIDTH + bit); } p += 1 + *p; break; } case beg-
line: printf ("/begline"); break; case endline: printf ("/endline"); break; case on_failure_
jump: extract_number_and_incr (&mcnt, &p); printf ("/on_failure_jump/0/%d", mcnt);
break; case on_failure_jump_past_alt: extract_number_and_incr (&mcnt, &p); printf
("/on_failure_jump_past_alt/0/%d", mcnt); break; case dummy_failure_jump: ex-
tract_number_and_incr (&mcnt, &p); printf ("/dummy_failure_jump/0/%d", mcnt); break;
case push_dummy_failure: printf ("/push_dummy_failure/0/%d", mcnt); break; case may-
be_pop_jump: extract_number_and_incr (&mcnt, &p); printf
("/maybe_pop_jump/0/%d", mcnt); break; case pop_failure_
jump: extract_number_and_incr (&mcnt, &p); printf ("/pop_
failure_jump/0/%d", mcnt); break; case jump_past_alt:
extract_number_and_incr (&mcnt, &p); printf ("/-

```

- 1  
Why  
IOS XE?
- 2  
IOS XE  
Design
- 3  
Up to  
17.15.x
- 4  
After  
17.18.x
- 5  
IOS XE  
Features
- 6  
Summary

# High Availability (SSO)


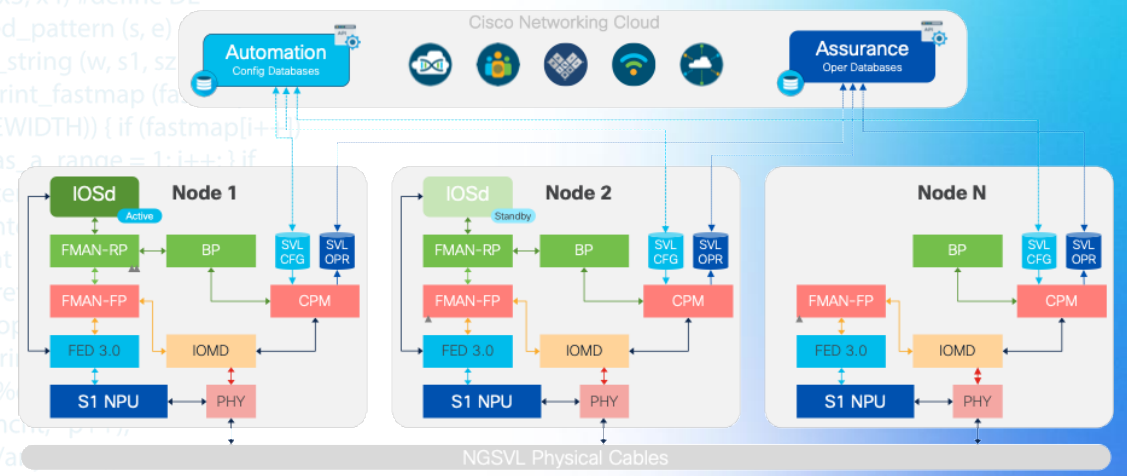
ISSU & xFSU

NG Stackwise

# Programmability

Model-Driven

# Application Hosting



# Campus Switching Innovations for Future Proofed Workspaces



## BRKENS-2609

### Campus Switching Innovations for Future Proofed Workspaces

**Minhaj Uddin** - Leader Technical Marketing, Cisco

Modern workspaces demand secure, scalable, and high-performing networks to support hybrid work, IoT, next-gen technologies, and sustainability goals.

This session introduces **Cisco's new switches** alongside **key innovations**, such as pervasive security, AI-driven insights, software-defined access, and cloud management.

Discover how these solutions simplify management, enhance scalability, boost energy efficiency, and enable smart, resilient, and sustainable environments ready to meet tomorrow's challenges.

CiscoLive Melbourne 2025  
Room 217, 1:00-2:30 PM, Wednesday 12th

[www.ciscolive.com/apjc/learn/session-catalog.html?search=BRKENS-2609](http://www.ciscolive.com/apjc/learn/session-catalog.html?search=BRKENS-2609)

#### Considerations for a future-proofed campus network

##### Wi-Fi 7 and Smart Spaces

- 10Gbps port throughput
- UPoE+ Power requirements
- Increased uplink and stack bandwidth
- Scaling MAC address density

##### Universal Zero Trust

- Seamless zero trust network access
- Least-privilege MFA & identity services
- Fabric architecture and scaled policy
- Rich telemetry and distributed enforcement

##### Campus Edge AI

- Move AI and ML workloads closer to client and data
- Minimize latency
- Defend against data export.
- Network-wide visibility and control

##### Future-proofed Scale

- Modern apps and clients are driving massive client growth
- 6Ghz Wireless, 4K/8K video, IoT, Sensors and autonomous robots
- Higher scale requirements for L2, L3, ACL, FNF, across Access and Core

© 2025 Cisco and/or its affiliates. All rights reserved.

BRKENS-2609

6

# Mission-Critical Resiliency


Your business stops if the network is down



Cost of only **one hour** of downtime to an average enterprise > **\$300,000**\*\*

\*\* Based on industry reports from Gartner and ITIC

**Catalyst 9600 Series**  
(Dual chassis w/ StackWise Virtual)



**Catalyst 9400 Series**



**Catalyst 9500 Series**

## Architecture

### StackWise® + StackWise Virtual

- Virtualized redundant systems for simplified configuration & protocols

### Graceful Insertion/Removal (GIR)

- No downtime when device in maintenance mode

## Operating System

### Software Maintenance Upgrade (SMU)

- Minimal or no downtime patches

### In-Service Software Upgrade (ISSU)

- Upgrade with minimal or no traffic loss

### Extended Fast Software Upgrade (xFSU)

- < 5 sec downtime - Stack upgrade

## Platform

### Redundant Supervisors

- Modular with SSO/NSF
- SVL Quad-SUP RPR

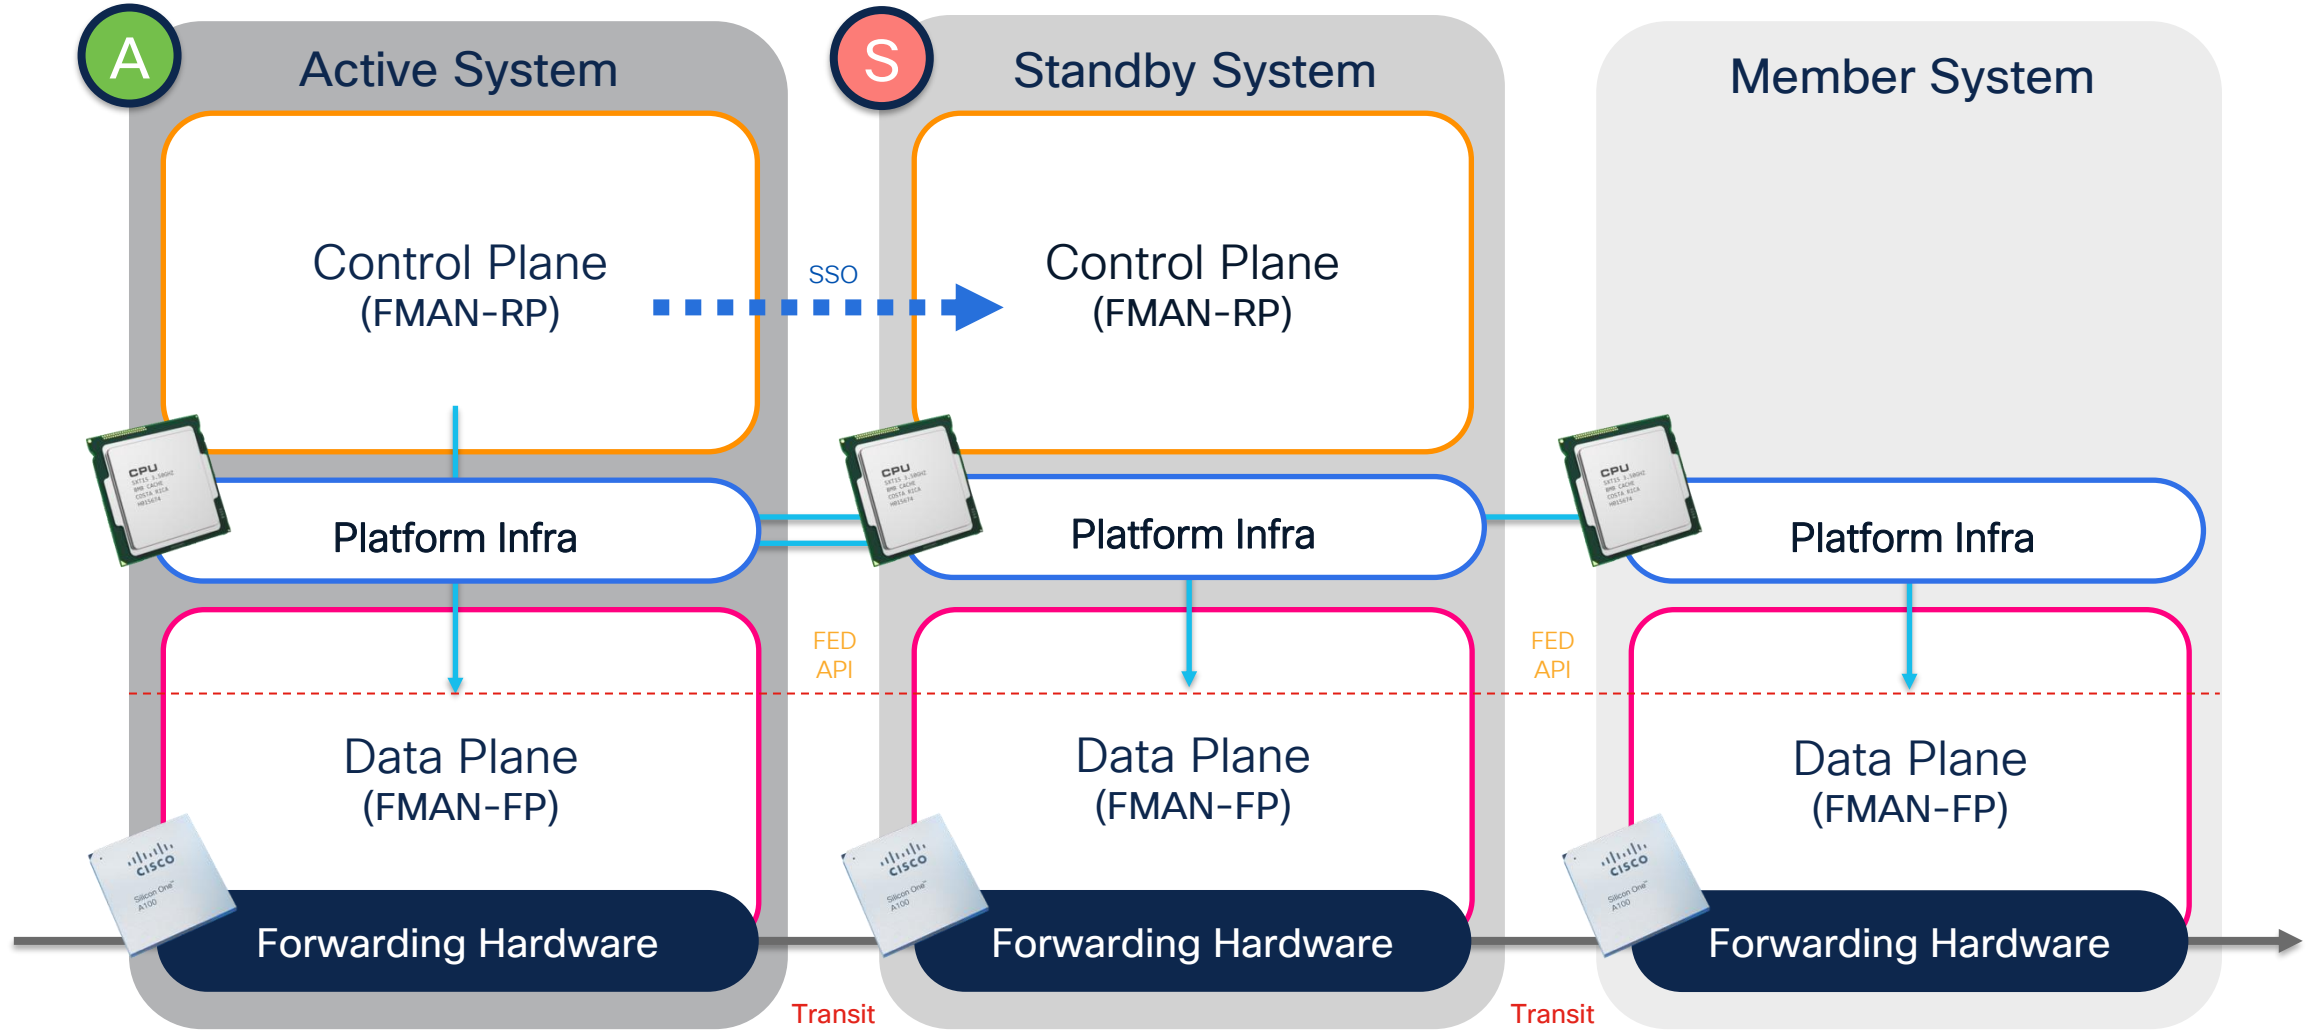

### Redundant Power & Fans

- In case of any hardware failure

Eliminate downtime with **High Availability** designed at every level

# Cisco IOS XE High Availability

## Control Plane to Data Plane Programming



# IOS XE Install – SMU patches

Patch a previously 'certified' image, without full upgrade



DNAC Version: Guardian 2.3.3  
License Level : DNA Advantage

**Software Maintenance Update (SMU)** is a small software package that can be installed for a **fix** or **security resolution** on a released version.

Patching can also be performed using Cisco Catalyst Center

- **Hot Patching** – does not require explicit reload after installation to get activated. (not traffic-affecting)
- **Cold Patching** – requires a system reload after installation to get activated. (traffic-affecting). The 9200 family supports only this.

- **Bundle SMU** – One single SMU file containing and applying multiple SMU fixes at once.
- **Independent SMU** – One single SMU file tailored to address a specific bug or vulnerability.




C9300-Access.cisco.com (172.26.192.131) Image Update

Date: Mar 16, 2023 12:15 PM      Duration: 4 minutes 20 seconds      Status: ● Successfully Activated cat9k\_iosxe.17.03.04.SPA.bin  
cat9k\_iosxe.17.03.04.CSCwa53947.SPA.smu.bin

---

Operations      Checks

- > ● SMU Distribution 1 minute 51 seconds
- ∨ ● SMU Activation 2 minutes 27 seconds
  - SMU Activation of image: cat9k\_iosxe.17.03.04.CSCwa53947.SPA.smu.bin on device: 172.26.192.131 completed successfully
  - > ● Pre Activation Script Execution 1 second
  - > ● Activation 2 minutes 26 seconds



Reduces both time and scope for [re]certification



Self-sufficient and reliant bug fix solution for PSIRT, CFD & IFD



Hitless Zero-Downtime with Hot Patching



EMR Releases with 371 SMU files available

# LAN High Availability

## In-Service Software Upgrade (ISSU)



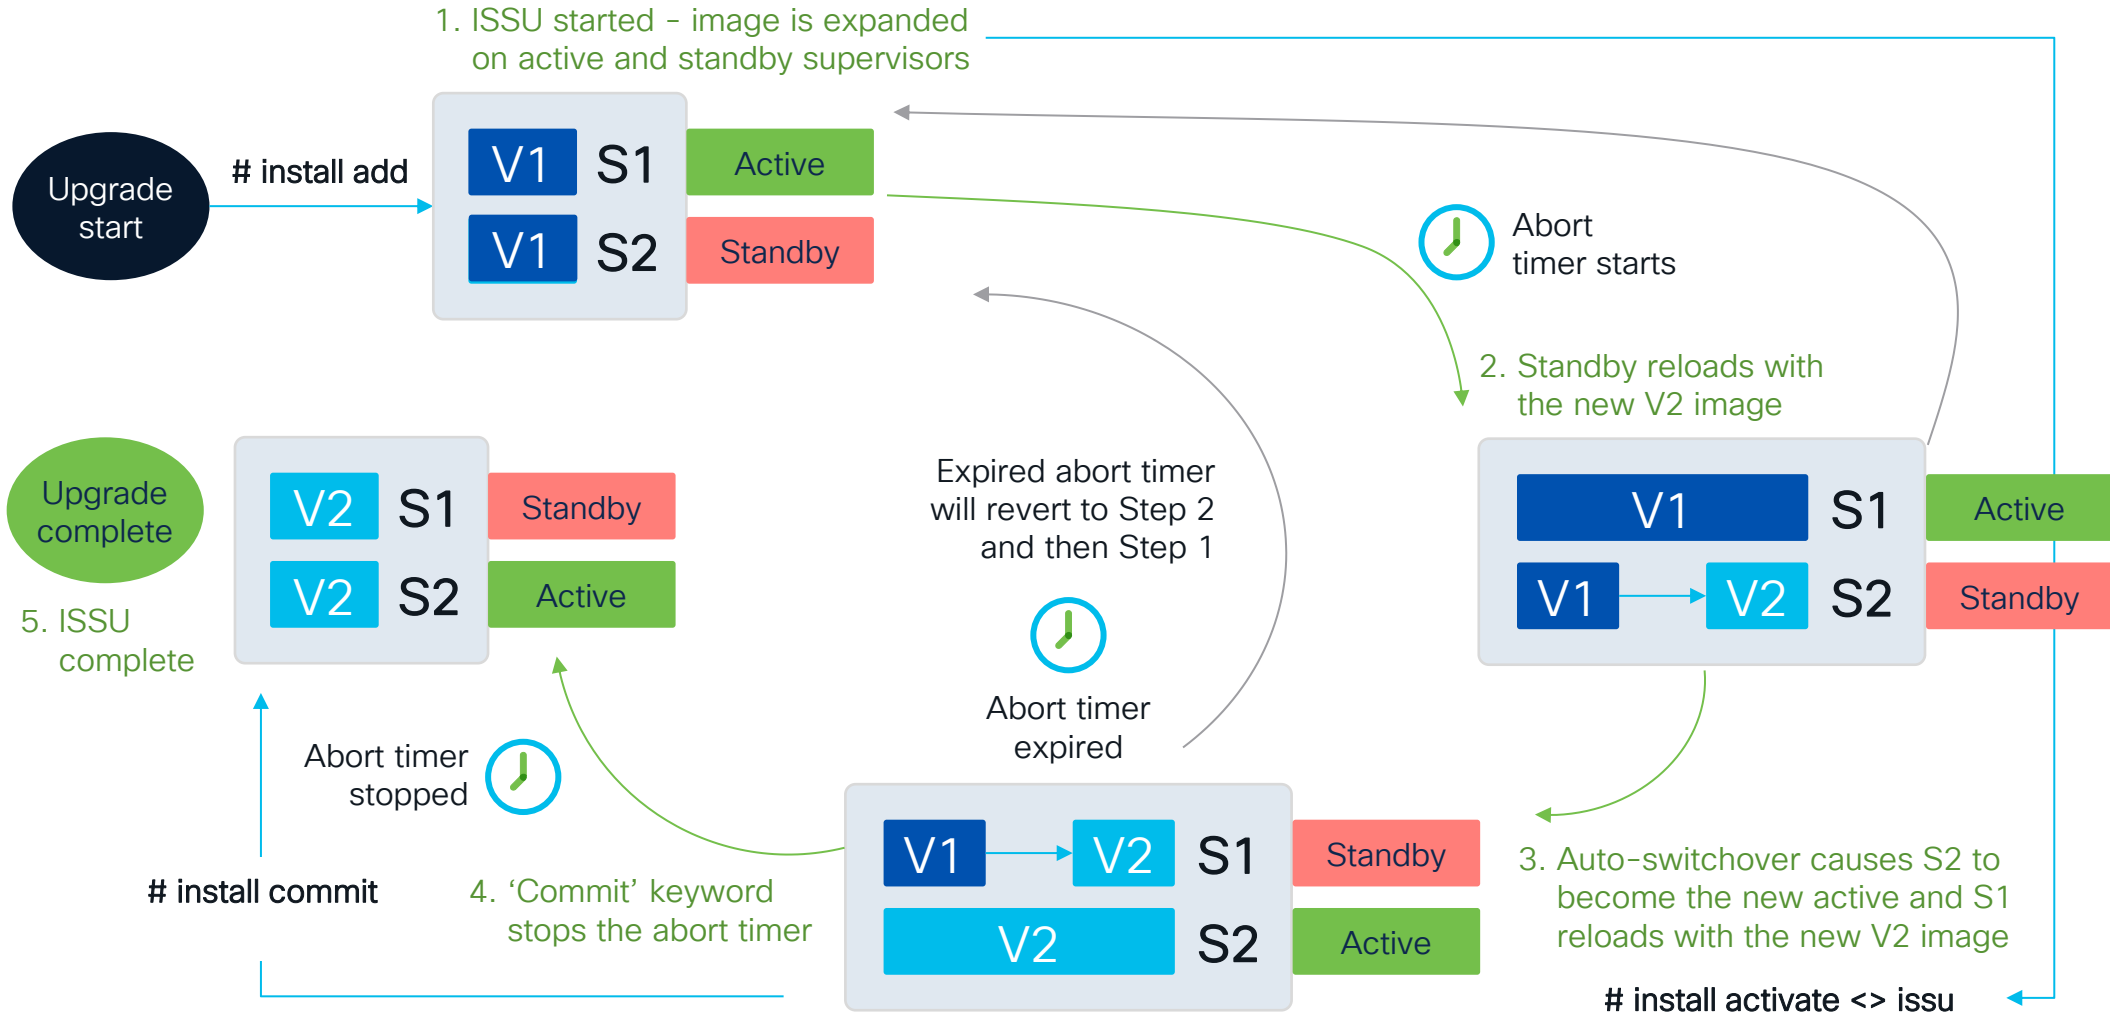
Catalyst 9600/X



Catalyst 9500/X



Catalyst 9400/X



If S2 fails to become the active, it will revert back to Step 1

# LAN High Availability

Extended Fast Software Upgrade (xFSU)




**NEW** ≤ 5 Seconds - 17.15.2

C9300X (≤ 30s) - 17.7.1

C9300/L (≤ 30s) - 17.3.2

```
>_
C9300# install add file <image> activate xfsu commit
```

Control Plane Upgrade  
V1 → V2



Data Plane upgrade  
V1 → V2

C9300 | C9300L | C9300X


Cisco xFSU minimizes downtime to **less than 5 seconds** (standalone or stack)

# Cisco IOS XE Programmability

Full Telemetry “Stack”

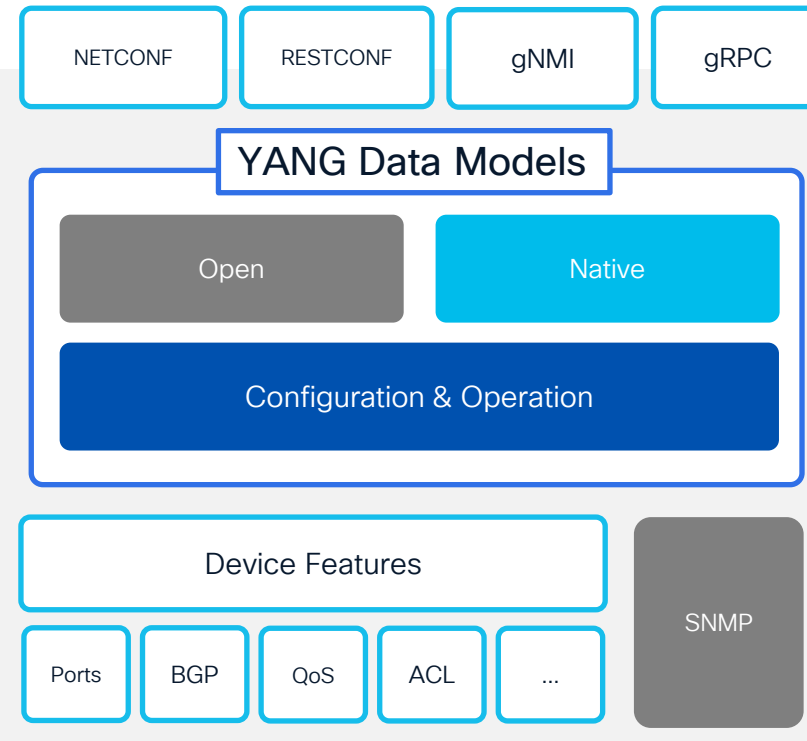
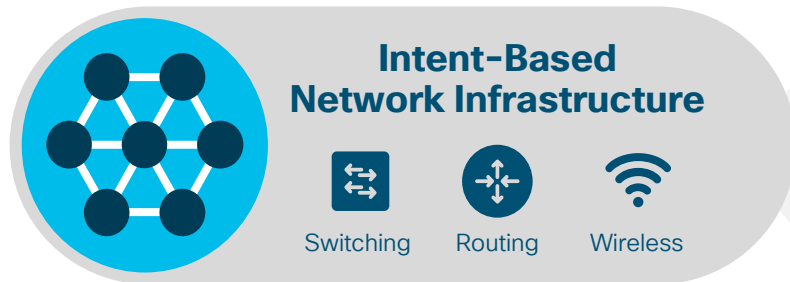


[developer.cisco.com/site/IOS\\_XE](https://developer.cisco.com/site/IOS_XE)



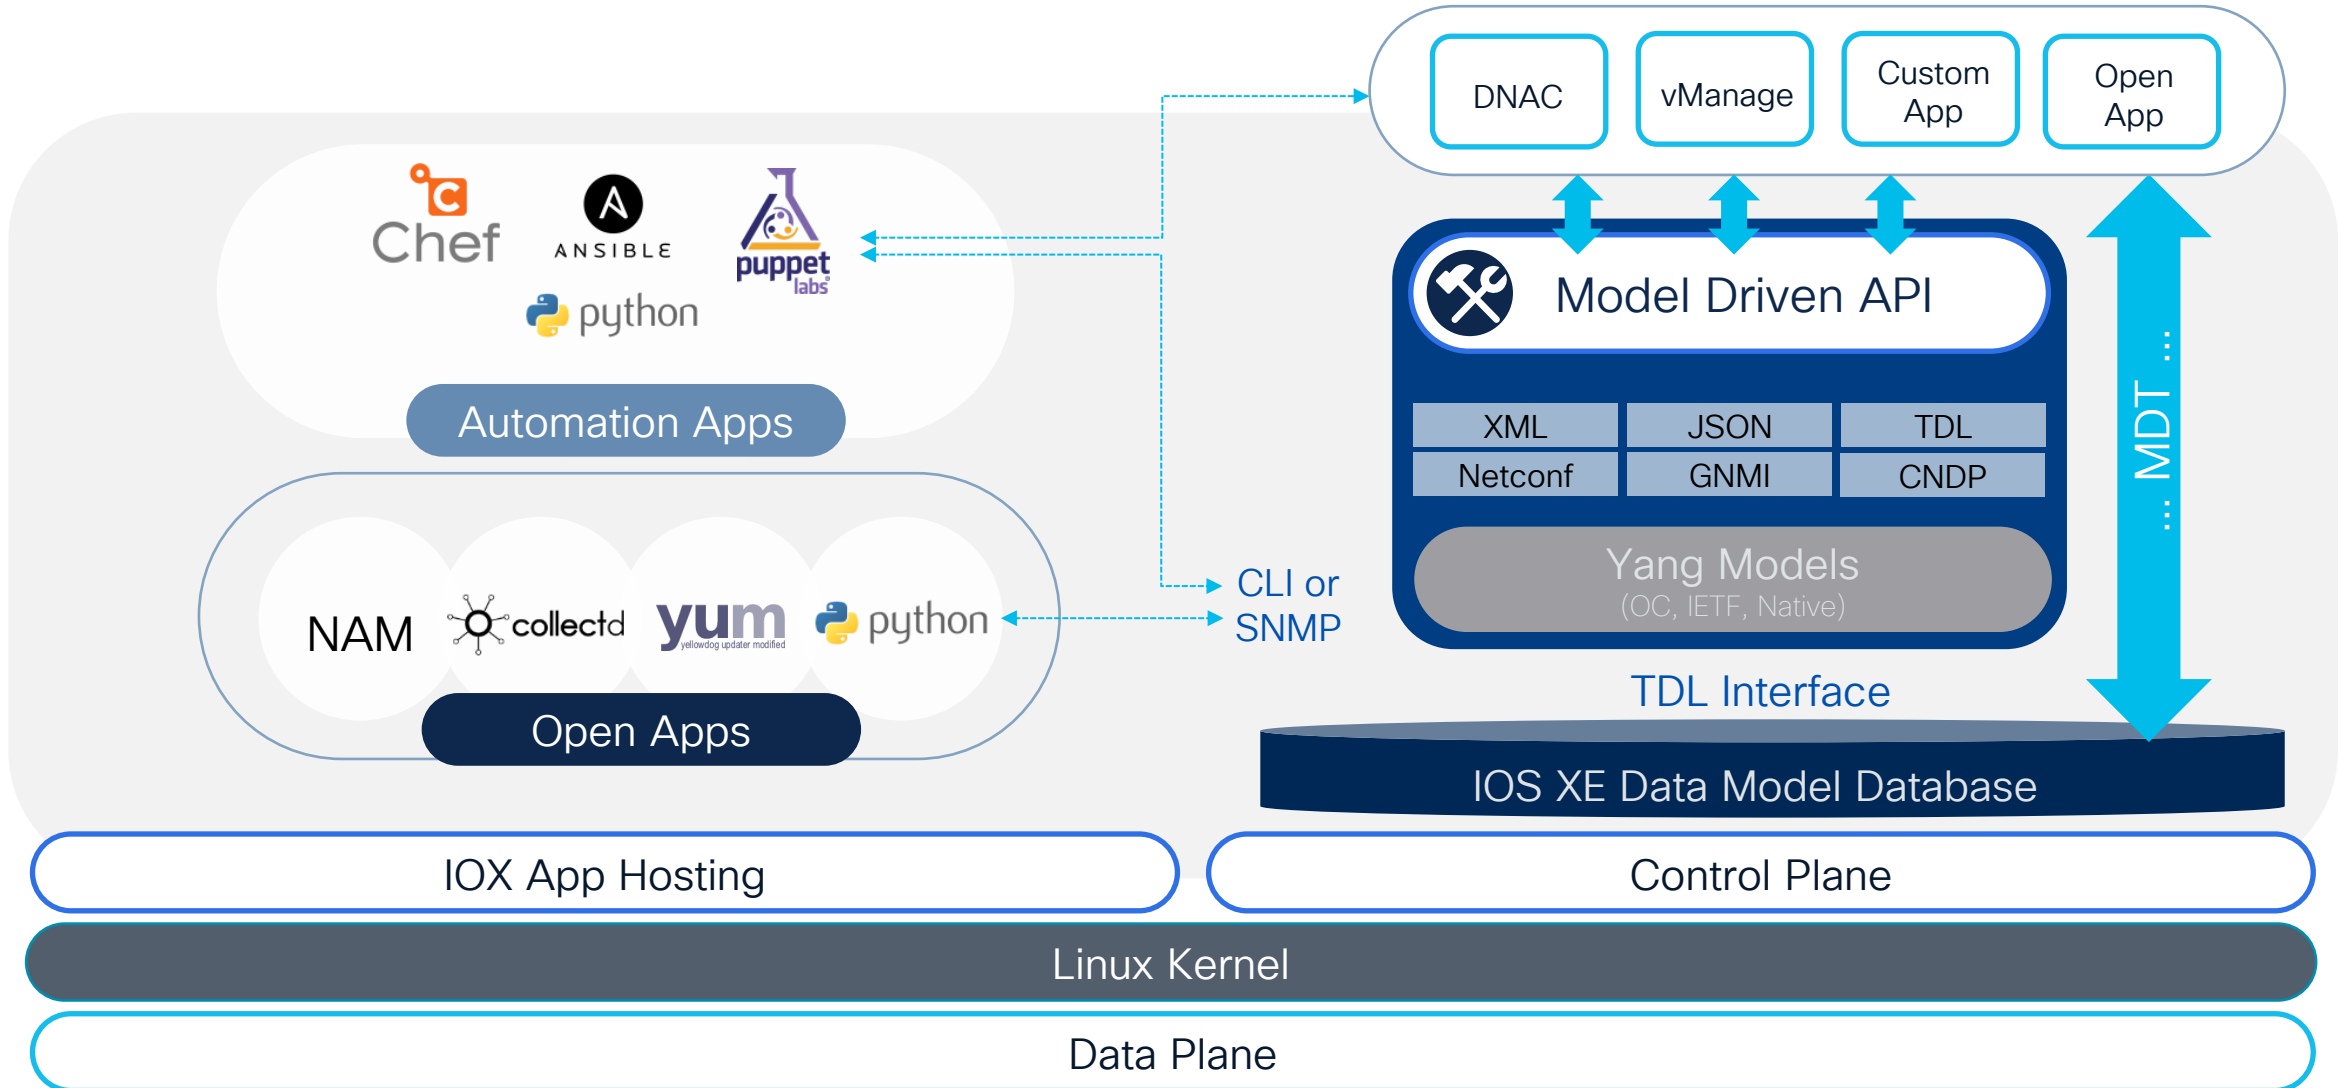
NETCONF, RESTCONF, gNMI & gRPC are programmatic interfaces that provide additional methods for interfacing with an IOS XE device

YANG data models define what is available for configuration and operational telemetry



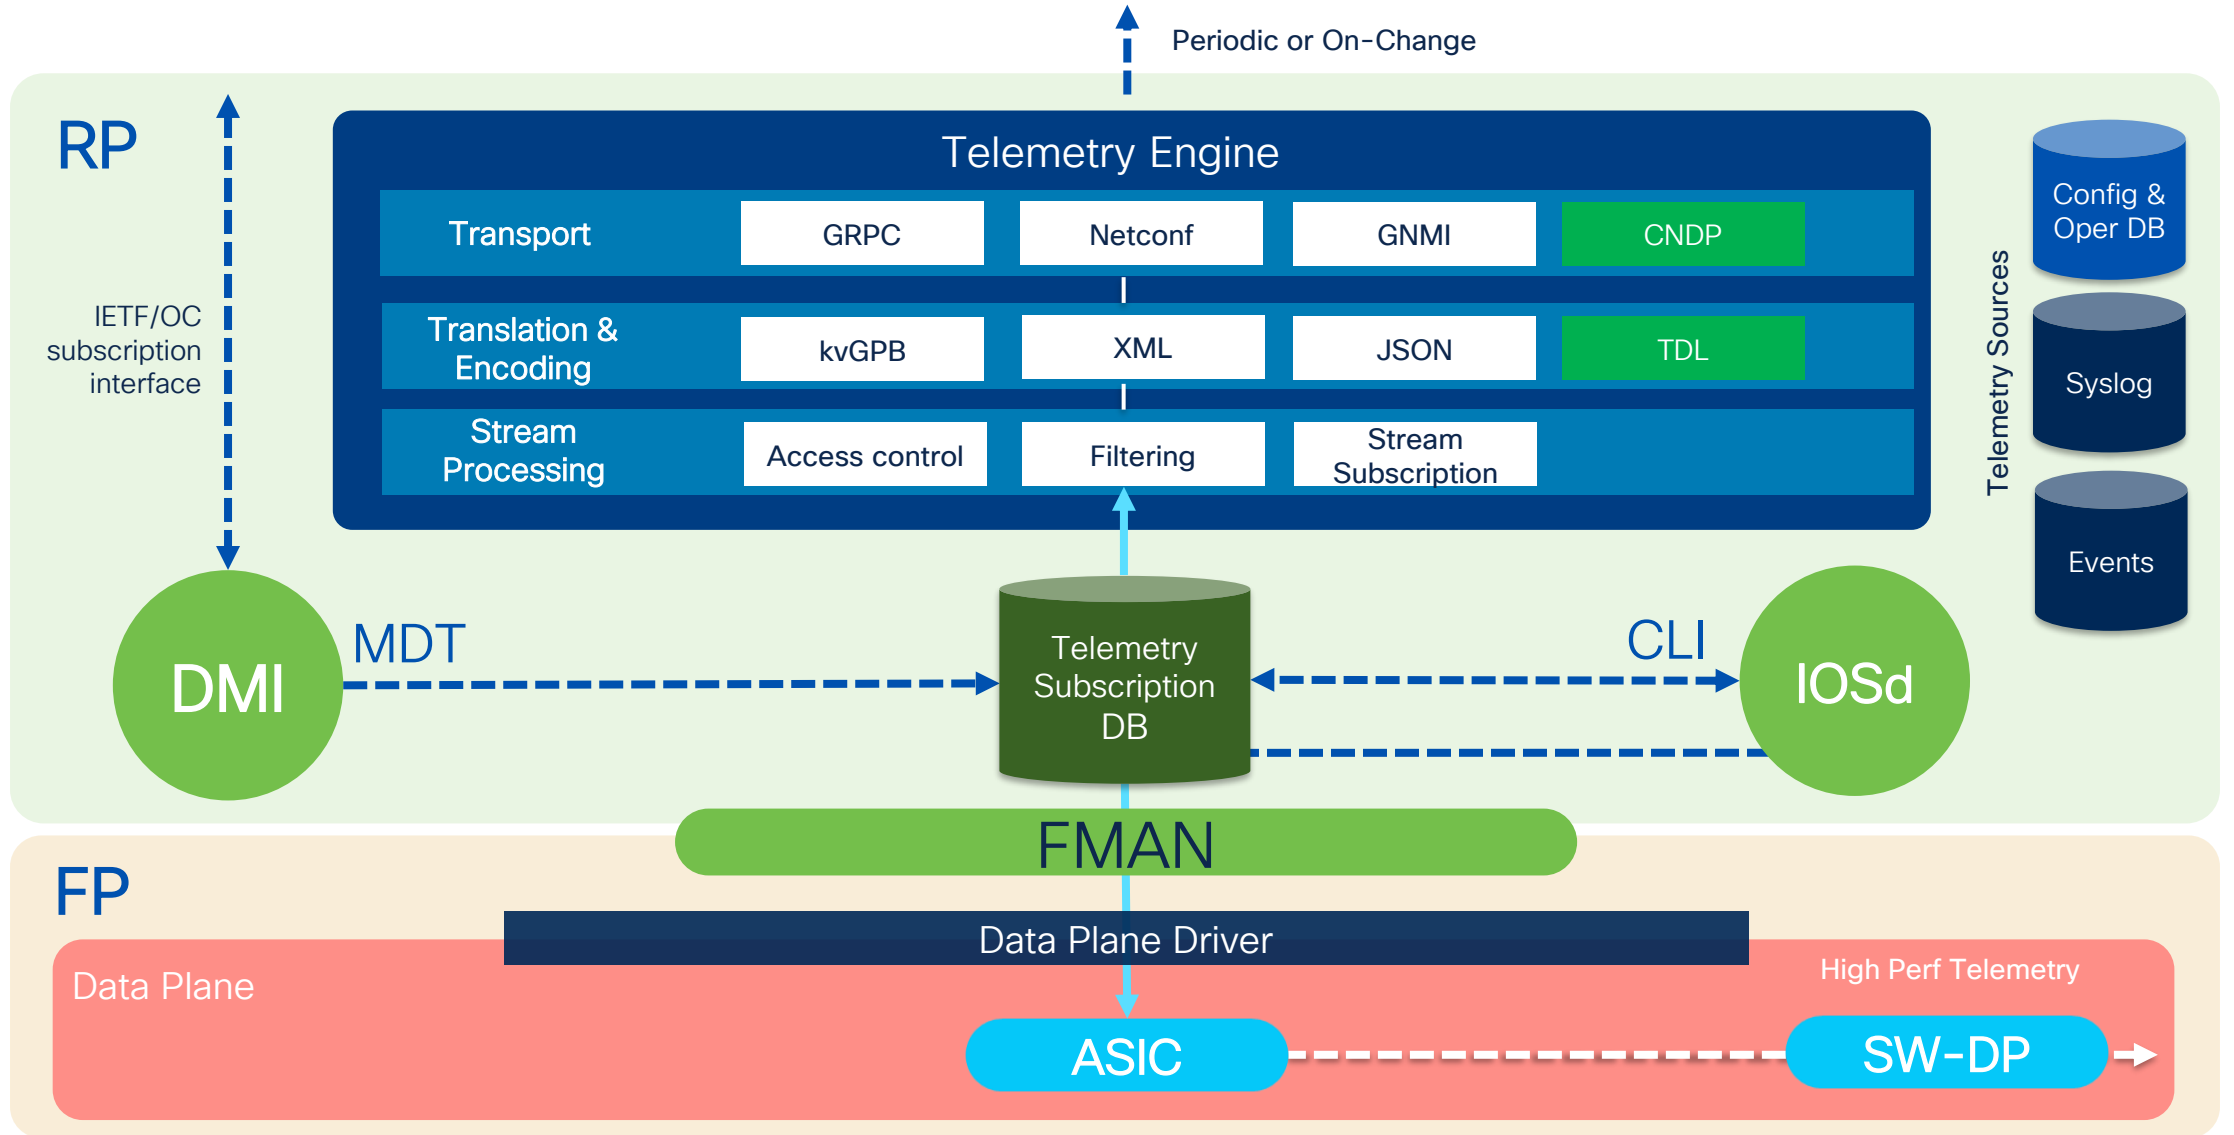
# Cisco IOS XE – Management

## Management Plane – High-Level Overview



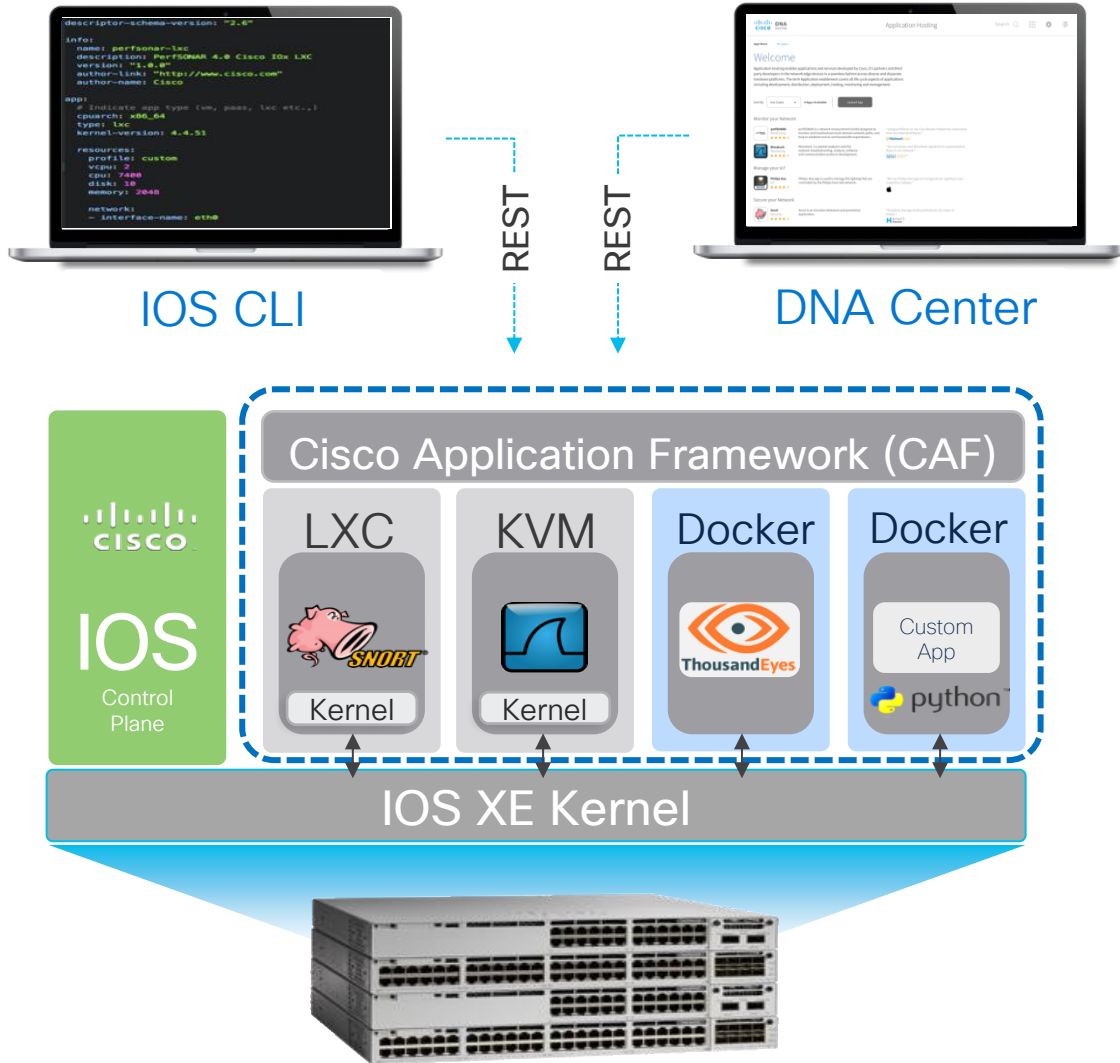
# Cisco IOS XE - Model-Driven Telemetry

## MDT Architecture



# Cisco IOS XE

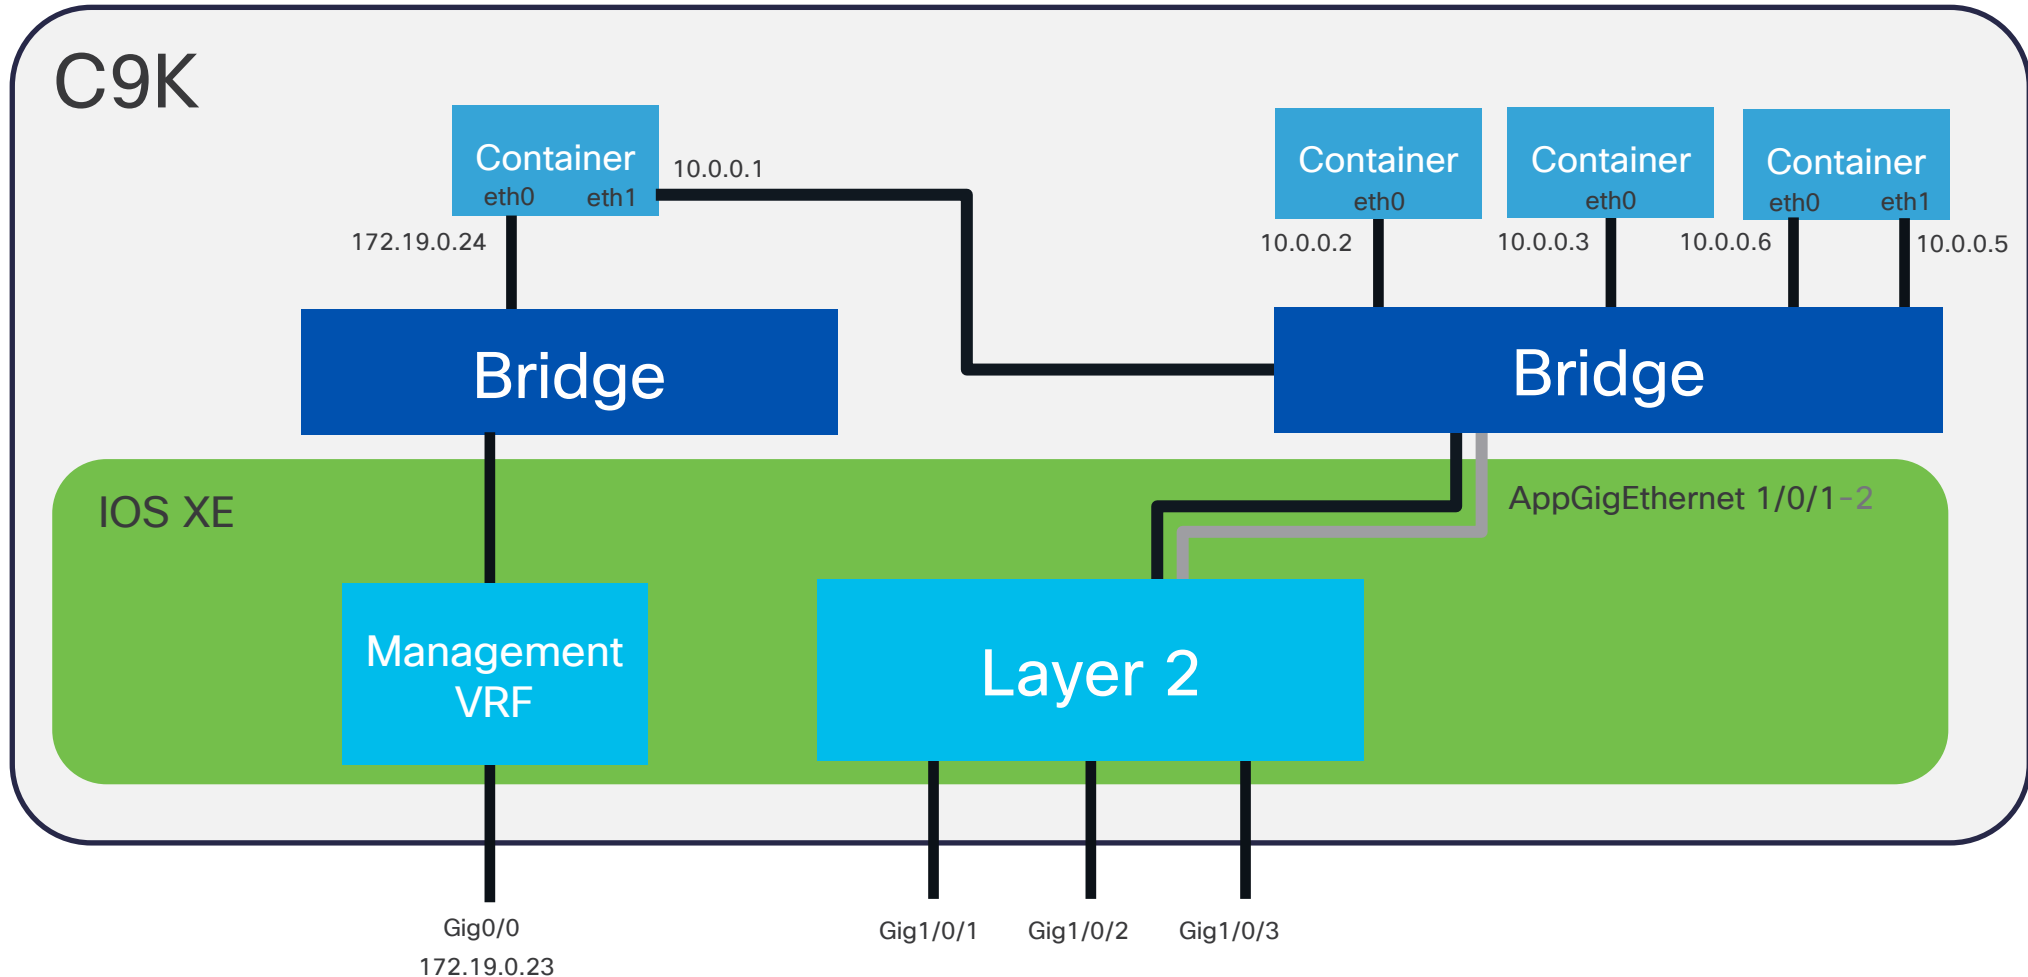
## Application Hosting



- Cisco will not support third-party apps or open-source apps, unless specifically called out
- Such apps, however, will be validated for compatibility on Catalyst 9000 switches
- DevNet ecosystem will indicate the partners who have worked on Catalyst 9000 switches

# Catalyst 9000 Series - App Hosting

## Container Networking





# Introducing Next Generation Stackwise

A new era of micro-services device clustering



## A new **Stacking Architecture** based on standard Ethernet protocols!

**Cloud-Native IOSXE** includes a **new Stack architecture** based on standard Ethernet protocols and encapsulation. Using either back-side or front-side ethernet links - for **greater scale, stability, better SSO convergence, and ISSU/XFSU & SMU** for hitless software upgrades.



### How is the Architecture new?

- A new Bootstrap Process (BP) & Cluster Process Manager (CPM), as a separate Linux thread.
- Runs IP SPF + VXLAN-GPO over standard Ethernet

### Is it front panel or back panel?

- Both! Dedicated cables or Transceivers + cables.
- Stack Interface (SIF) is treated as any ethernet port

### Can I still manage it the same?

- Familiar (same) config and show commands
- Uses same 'stackwise' & 'stackwise-virtual' CLI

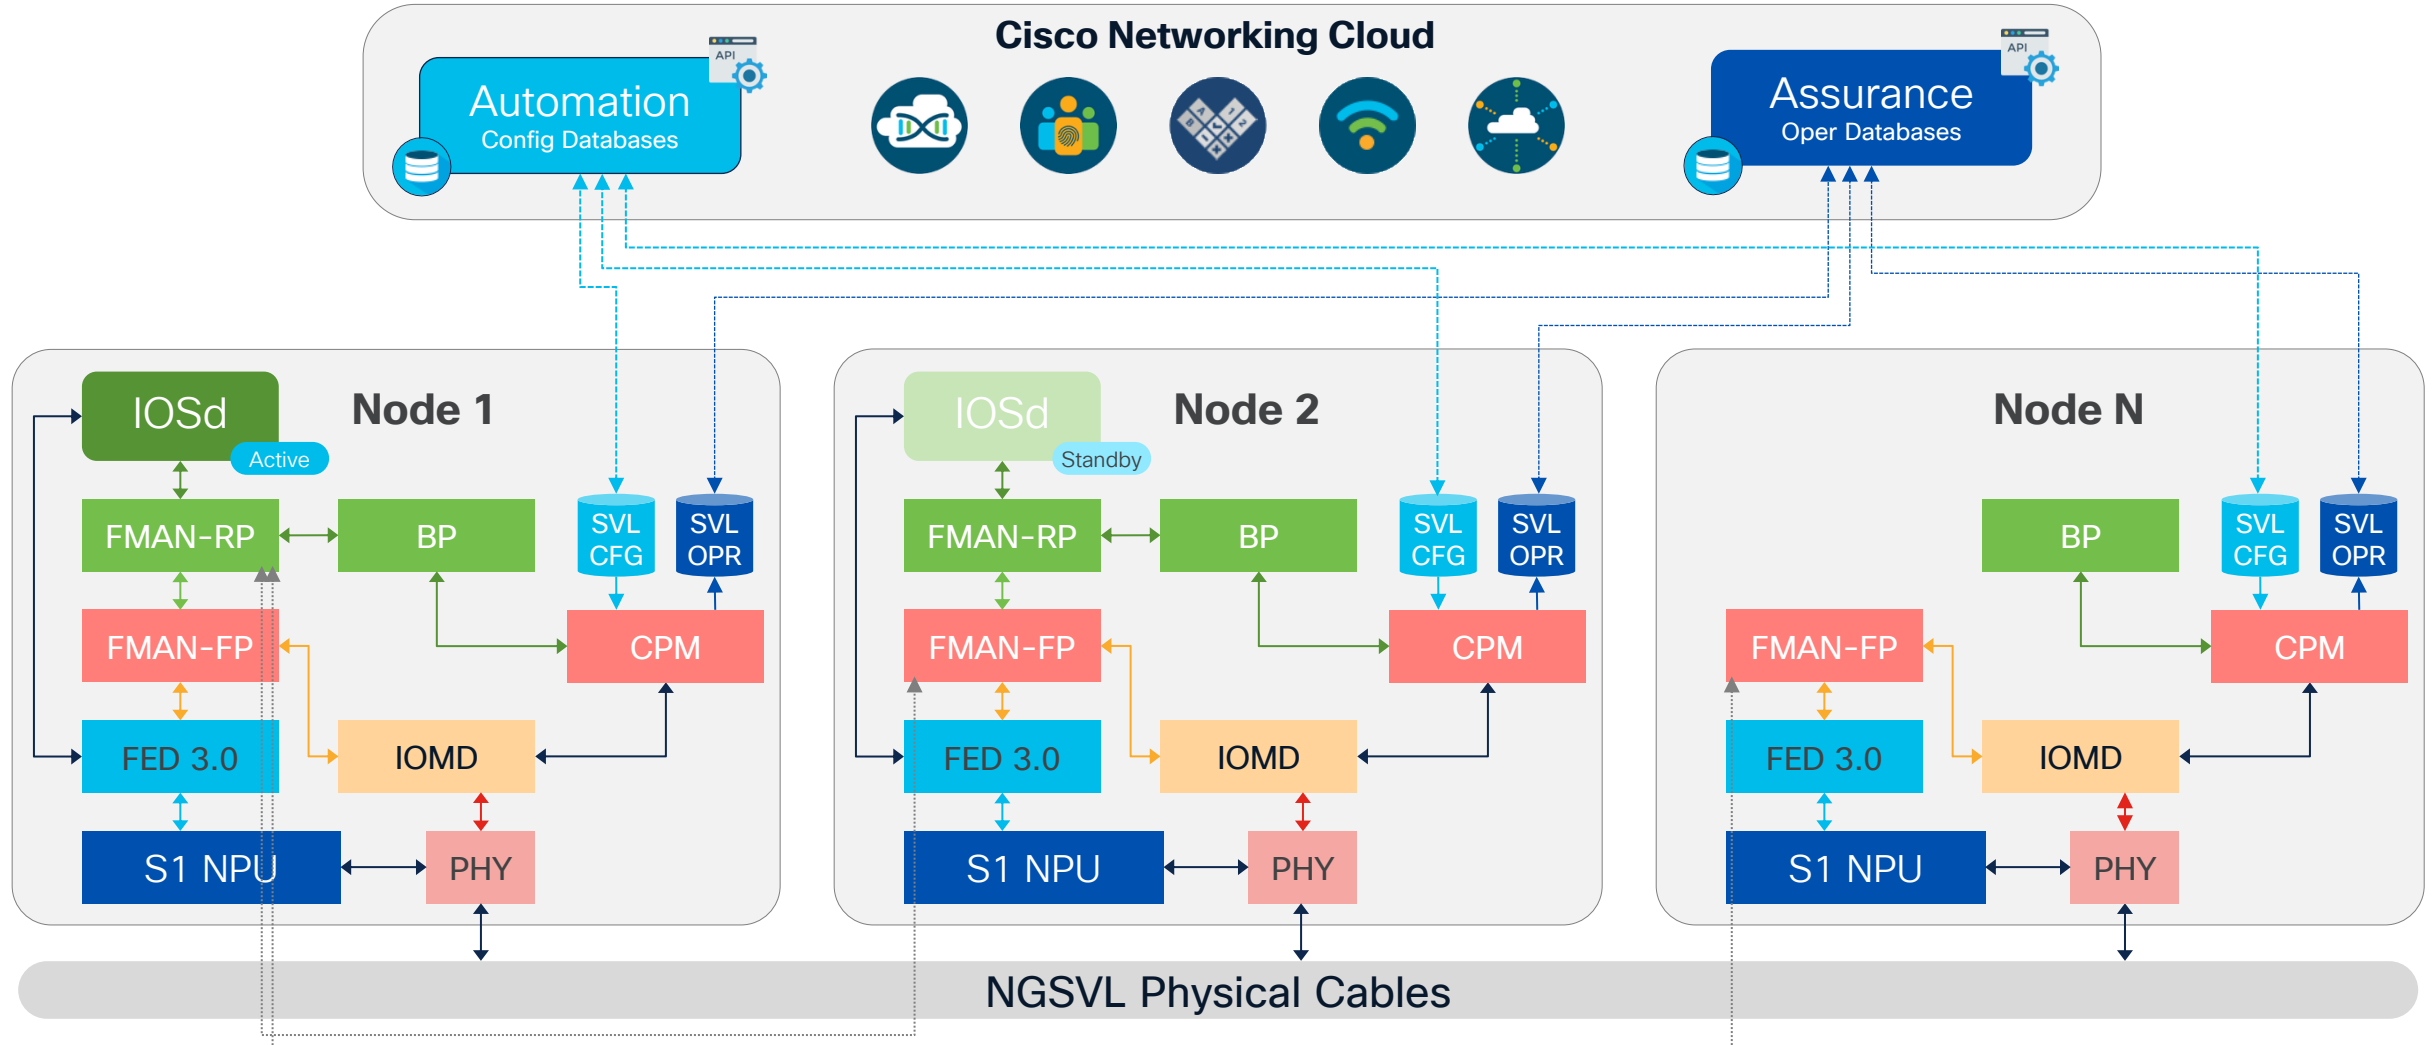
### How does it benefit me?

- Common to both C9350 & C9610
- Simplified cabling & dynamic link add/change
- Improved SSO and ISSU/SMU convergence

... and much more, coming soon


# What is Next Gen Stackwise?

Introducing an entirely new Stacking Architecture

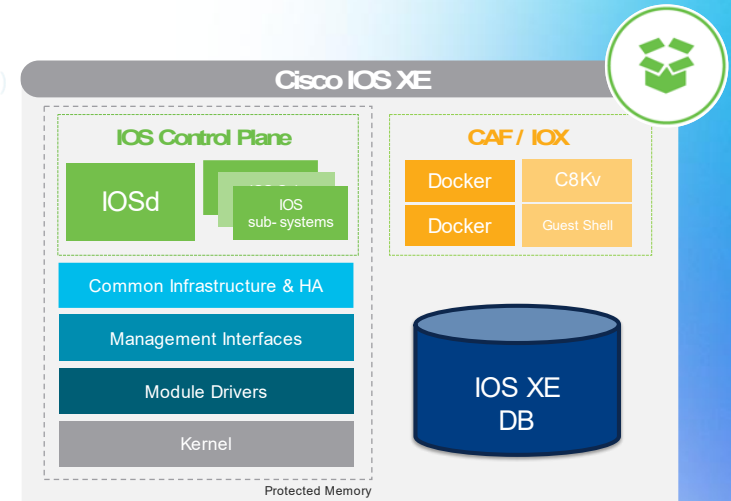


# Glimpse into the Future

## Software Innovations for Enterprise Switching



- 1 Why IOS XE?
- 2 IOS XE Design
- 3 IOS XE Features
- 4 Up to 17.15.x
- 5 After 17.18.x
- 6 Summary



# Making Cisco C9000 Series into Campus “Smart Switches”

Enhanced Control-Plane and AI-ready App-Hosting

x86

**Cisco C9000 Series Switches** have supported **Application Hosting & Edge Computing** for years!

We are bringing these capabilities to a new level with more Apps, Multi-App hosting and hardware-support for **Edge AI/ML agents** or **Hypershield agents**

## C9350



### C9350 Standard SKUs

- **Class:** Low/Med CPU
- **Intel Gen12 RPL Core i3**
  - **4E+2P-Core** (2.5 GHz)
  - **16GB DDR5** (4800 MHz)
  - **4x 10G NIC\*** (X710)

- **USB3.0 SSD** (External)
- \* Hardware capability. 2x AppGig for FCS

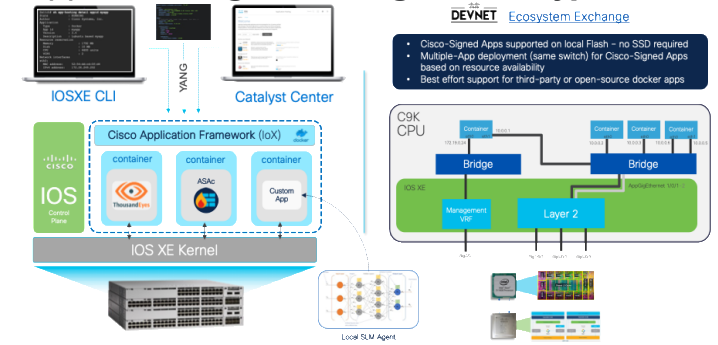


## C9610-SUP-3/XL



### C9610-SUP-3/XL SKUs

- **Class:** Med/High CPU
- **Intel Gen10 ICX Xeon D**
  - **8P-Core** (2.0GHz)
  - **32GB DDR4** (2400 MHz)
  - **4x 25G KR\***
- **M2 SATA SSD** (External)




# Where are things going?

## Features & Scale

More details in PPT Notes



## SMU & XFSU/ISSU



### Patch an existing (certified) IOS XE version


- Bug fixes and new (rebuild) features
- Critical software vulnerabilities (PSIRT)

### Sub-second In-Service Software Upgrades

- Redundant Supervisors and SVL
- Standalone (xFSU) and Stackwise



## MAC & IPv4/v6 Scale



### Increasing MAC scale for Wireless & IOT


- WLC to Core SVI scale growing  $\geq 256K$  MACs
- WiFi6E, 5G and IOT Devices & Sensors

### Increasing IPv4/v6 scale for Internet & VPN

- Collapsed LAN Core + SP/WAN Edge designs
- IPv4 GRT is  $\geq 850K$  and IPv6 GRT is  $\geq 50K$



## VoQ & HBM



### Virtual Output Queues (VoQ) for high throughput


- Eliminate Head-of-Line Blocking at egress
- Support for logical (e.g. sub) interfaces with HQoS

### Local and Expandable HBM for optimal buffering

- Local buffers for low-latency strict-priority queuing



## OGACL & SGACL



### Object-Groups map IP/mask to Labels in CEM


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

### OG/SGACL ACEs take minimal space in TCAM

- Only the Permit/Deny ACEs stored in TCAM
- OG/SGACL with same ACEs can reuse entries



## WAN MACsec & IPsec



### 256-bit Hardware Encryption over L2

- P2P LAN MACsec with 802.1ae
- P2MP WAN MACsec with 802.1q ClearTag

### 256-bit Hardware Encryption over L3

- P2P SVTI IPsec with IKEv2, ESP
- Supports Site-to-Cloud (AWS, Azure, Zscaler)



## SDA & ZTNA



### SD-Access Fabric with Group-Based Policy


- L2/L3 virtual overlay, with macro & micro segments
- Full automation, assurance & policy with Cisco DNA

### Application & Endpoint Identification & Analytics

- Hardware Flexible NetFlow and IPFIX
- AVC/NBAR2 and SDAVC/CBAR to ID clients & apps

# Summary

- 1
- 2
- 3
- 4
- 5
- 6



# Catalyst 9000 with Programmable ASICs

Benefits for your network



- ✓ **FLEXIBILITY** and
- ✓ **ADOPTABILITY**



Enabling **Network Evolution**  
on your journey to  
**AI-Ready Secure Networking**

# Catalyst 9000 with Cisco IOS XE

Benefits for your network



## One Release Train

Operational Efficiency,  
Consistency in Behavior,



## Common Features

Feature Velocity  
across Platforms



## Updates & Patching

Sub-package upgrades,  
Cold and Hot Patching



## Native Programmability

Object based model,  
Netconf/REST Interfaces



## Trustworthy & Secure

64-bit ASLR, Secure Boot,  
Hardware TAM

# Catalyst Leadership in Enterprise Networks

- New Feature
- Enhanced


## A Platform based Approach

### Catalyst Center and Meraki Dashboard


**32M** Network Devices Managed

↑ 19M APs | 6M Switches | 2.5M Routers | 830M Clients


**17M**  
Devices on  
Catalyst Center



**15.3M**  
Devices on  
Meraki Dashboard




### Cisco C9000 Family



**100,000+** Customers, **Millions of** Switches

Catalyst 9K continues to be the fastest ramping product in the company's history

Secure Networking	Digital Experience	Operational Simplicity
● Common Policy	● Campus Automation	● Cloud Managed Catalyst
● Secure Equipment Access	● AI Endpoint Analytics	● Infrastructure as a Code
● SD-Access (LISP & EVPN)	● ThousandEyes Digital Experience	● S3 & CloudWatch Integration
● High-speed Encryption	● AI Ops & Assurance	● Visibility, Control & Rollback



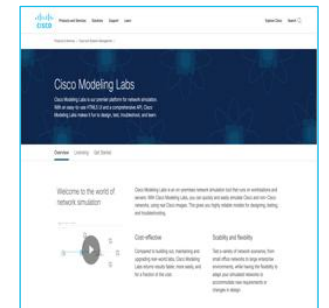
Cisco Validated Profiles (CVP)  
BRKARC-2092



Industry Validated Reports



Industry Certifications



Cisco Modeling Labs



# Catalyst 9000 Switching - Useful Resources

**Cisco Catalyst TV**  
@CiscoCatalystTV 1.52K subscribers 71 videos  
Subscribe

Welcome to Cisco Catalyst TV! This channel is all about Cisco Catalyst Pla... >

HOME VIDEOS PLAYLISTS COMMUNITY CHANNELS ABOUT


Latest Popular

**Introducing the Catalyst 9200CX Series Switches**  
111 views · 5 days ago

**IOS XE 17.11.1 Catalyst Switching Highlights**  
136 views · 7 days ago

**IOS-XE Power Telemetry and Automation**  
143 views · 1 month ago

**Trustworthy Systems on Catalyst 9000 Platforms(Part-II)**  
56 views · 1 month ago



**SUBSCRIBE**



[www.youtube.com/c/CiscoCatalystTV](https://www.youtube.com/c/CiscoCatalystTV)



[cs.co/cat9kbook](https://cs.co/cat9kbook)

[Cisco Catalyst 9000 switches - Your Network, Your Way \(3rd Edition\)](#)  
[Cisco SD-Access for Industry Verticals - From Design to Migration](#)

# Complete your session evaluations



**Complete** a minimum of 4 session surveys and the Overall Event Survey to claim a Cisco Live T-Shirt.



**Earn** up to 800 points by completing all surveys and climb the Cisco Live Challenge leaderboard.



**Level up** and earn exclusive prizes!



**Complete your surveys** in the Cisco Live Events app.

# Continue your education



**Visit** the Cisco Showcase for related demos



**Book** your one-on-one Meet the Expert meeting



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



**Visit** the On-Demand Library for more sessions at [www.CiscoLive.com/on-demand](https://www.CiscoLive.com/on-demand)

**Thank you**

**CISCO** Live !

