The bridge to possible

Cisco IOS XE & ASIC Architecture

Catalyst 9000 Series

Shawn Wargo, Principal TME
@Shawn_Wargo

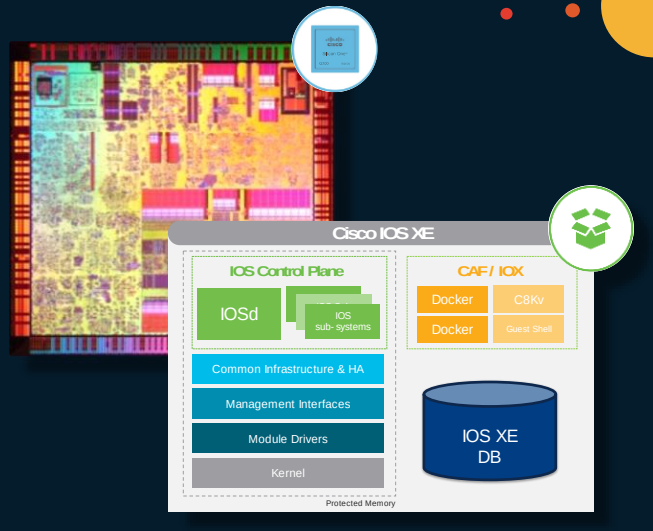

cisco Live!

BRKARC-2092

Who am I?

I'm a Principal Engineer of Technical Marketing (Principal TME) for Cisco Enterprise Networks (EN) and Intent-Based Network Group (IBNG) Product Marketing 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 Catalyst Switching, Wireless & Routing products, and solutions like Software-Defined Access (SDA) & Cisco DNA.

Shawn Wargo
Principal TME
swargo@cisco.com @shawn_wargo

Agenda

Agenda

The bridge to possible

Cisco UADP & Silicon One ASIC Architecture & Innovations

Catalyst 9000 Series

Shawn Wargo, Principal TME
@Shawn_Wargo

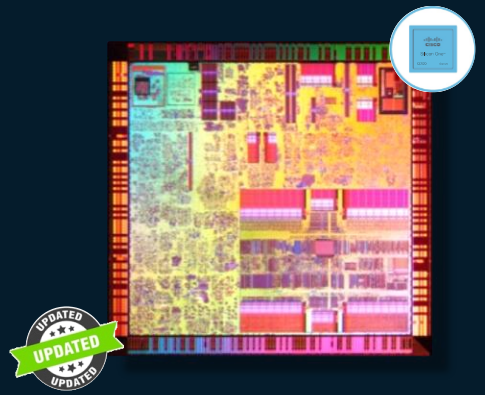

cisco Live!

BRKARC-2091

Session Abstract

Why do we even need ASICs?

- Can't we just do everything on multi-core CPUs? FPGAs?
- Application-Specific Integrated Circuits (ASICs) are purpose-built processors, capable of delivering Tbps of hardware-based network forwarding & services.
- ASICs are the heart of modern Campus, SP & Data Center networks.

If you know ASICs - then you know real networking!

- You may already know about the Cisco **UADP family of ASICs** on Catalyst 9000 series, but do you know about the new Cisco **SiliconOne family of ASICs**?
- Do you want to know how SiliconOne works in Catalyst 9000 series switches?

This session will cover the following areas:

- Why do we even need ASICs? - a brief history
- UADP and Programmable ASICs - the spirit of C9K
- Introducing Cisco SiliconOne ASICs - what's new/different?
- A glimpse into the future of Catalyst 9000 series - what's coming?

What this session is NOT

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**:

- Catalyst 9000 Series Switching Family – Access – BRKARC-2098
- Catalyst 9000 Series Switching Family – Core and Distribution – BRKARC-2099
- Cisco Programmable Silicon – Current and Future Use Cases – BRKENT-2000

Agenda

1. Why do we need ASICs?
2. Flexible ASICs & Cisco UADP
3. Cisco Silicon One ASICs NEW
4. Catalyst 9000 “X” Series NEW
5. A Glimpse into the Future
6. Summary & References

Cisco Catalyst 9000 Switching Portfolio

One Family from Access to Core – Common Hardware & Software

2022-2023 NEW

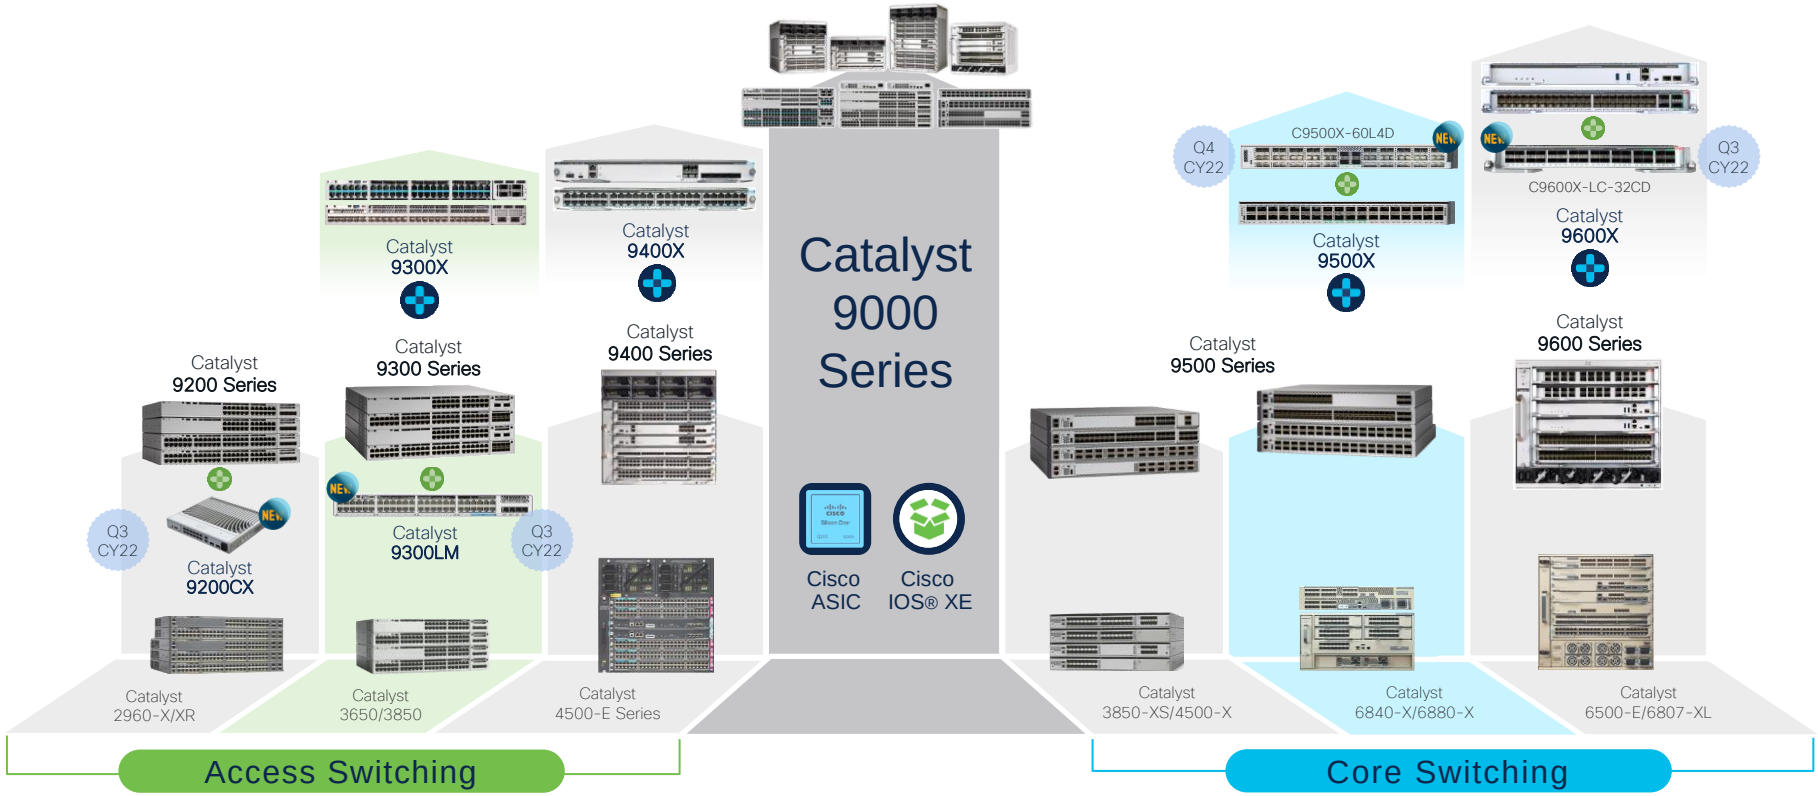

Catalyst 9000 Series – Common Building Blocks

Programmable x86 Multi-Core CPU

Application Hosting
Secure Containers

Open IOS XE® Polaris

Model-Driven APIs
Modular Patching

Cisco UADP & Silicon One™

Programmable Pipeline
Flexible Tables

Same IOS XE image for both **UADP*** and **Silicon One** C9K platforms

* Catalyst 9200 Series uses IOS-XE Lite

Software vs. Hardware

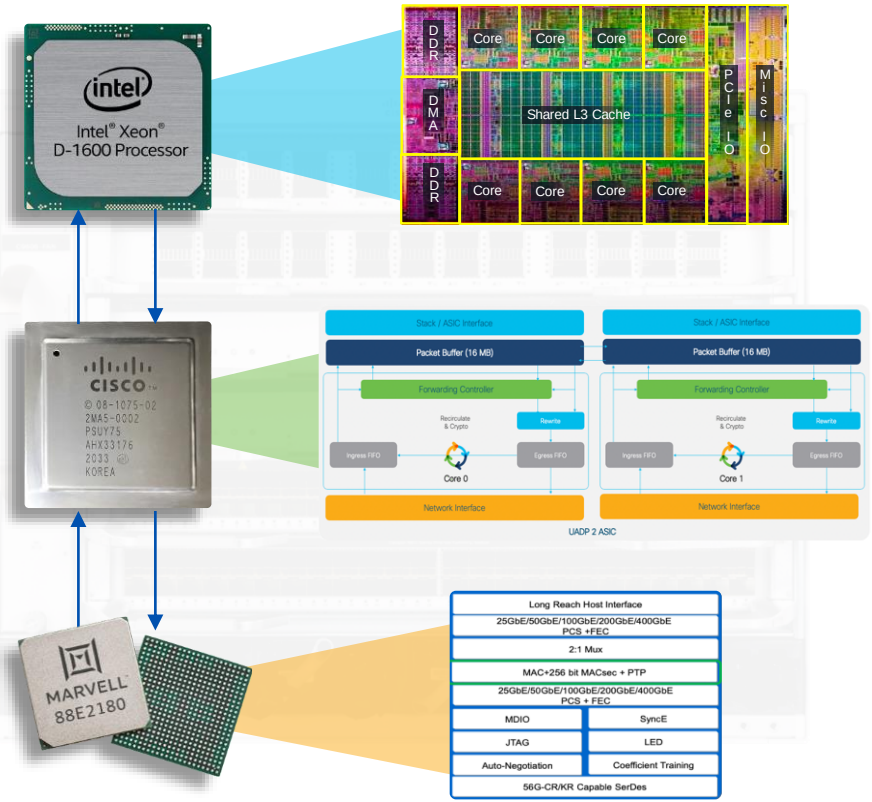

CPU/DRAM

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.

Stub/PHY(s)

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

Custom ASICs – Programmable Silicon

Access Data-Plane

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?

What is an ASIC?

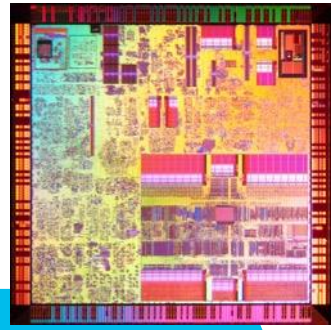

An Application Specific Integrated Circuit (ASIC) is a silicon microchip designed for a specific task …

… rather than ‘general-purpose’ processing in a CPU.

Why do we need ASICs?

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

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.

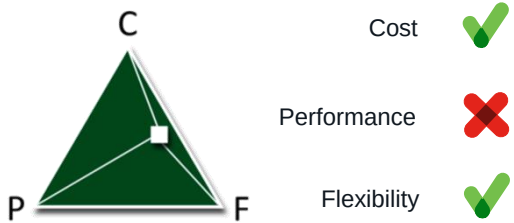

Remember, this is **Millions** – **Billions** of packets per second

Multi-Core ASIC Design

Combining multiple Processors in same ASIC

Multiple-Core processors have been used in personal computers since the 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 they can be cooled), by **load-sharing across multiple processors**.

However – there are **practical limits to how many cores yield improvements**, and cores must share all external components.

TIP

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

Intel Xeon BDW 8-Core

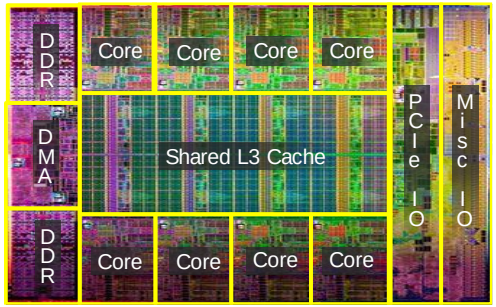

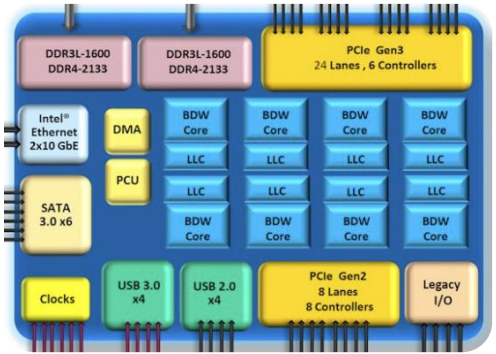

UADP (C9K)

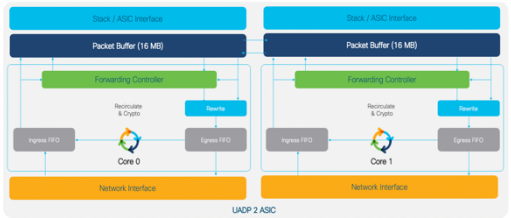

QFP (ASR/C8K)

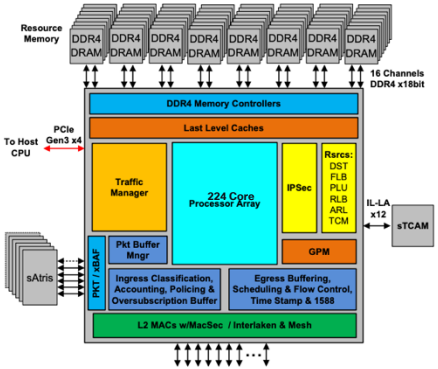

Hardware vs. Software Data-Plane

ASIC-based vs. CPU-based

UADP

QFP

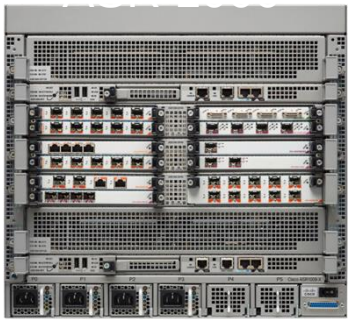

X86/ARM

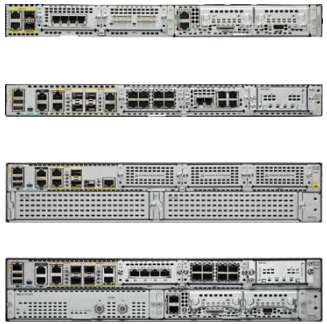

Hardware-Based (ASIC)

PROs	CONs
• High Throughput	• Less Flexible
• Gbps – Tbps	• PI + PD Development
• High Port Density	• Longer Dev & Test
• HW-based Services	• Scale limited by HW
• Lower COGs	• Limited Services

Software-Based (CPU)

PROs	CONs
• More Flexible	• Low Throughput
• PI Only Development	• Mbps – Gbps
• Faster Dev & Test	• Low Port Density
• Scale limited by CPU	• SW-based Services
• Lots of Services	• Higher COGs

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 **2X - 4X** more than an equal ASIC

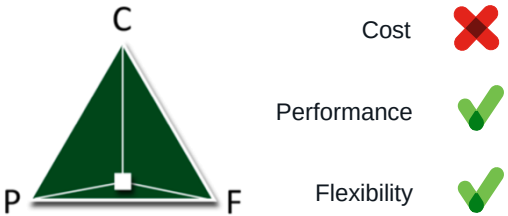

Field Programmable Gate Arrays

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 then 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 not 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, in order to add a new feature.

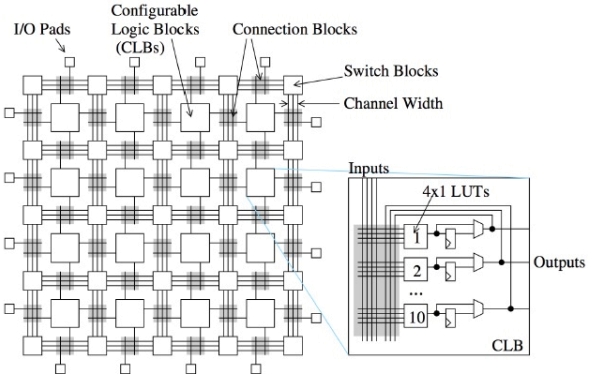

intel.com/content/dam/www/programmable/us/en/pdfs/literature/misc/fpgas_for_dummies_ebook.pdf

What does an ASIC do?

ASICs are fundamental to network devices

ASICs provide 2 basic packet processing functions

Forwarding

Receive, Process & Transmit

Services

Special Non-Forwarding Tasks

What does an ASIC do?

ASICs are fundamental to network devices

At a basic level – network devices forward data “as fast as possible”

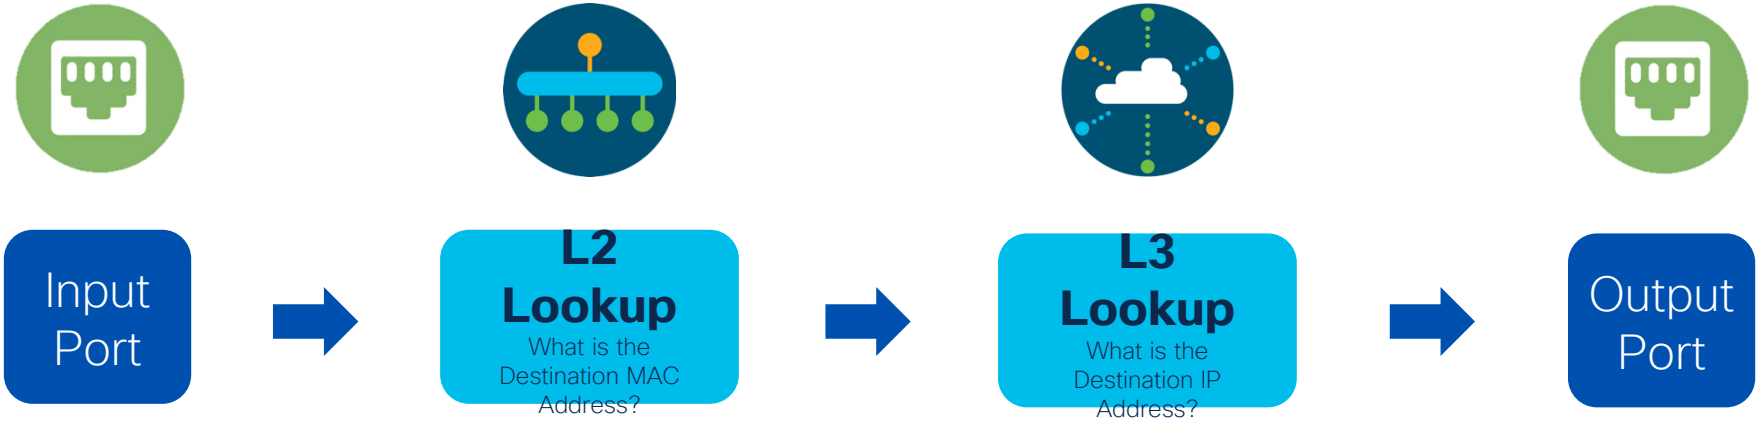

Modern processing speeds are **Terabits per second** (Tbps)

What does an ASIC do?

ASICs are fundamental to network devices

In addition – network devices can perform special processing tasks

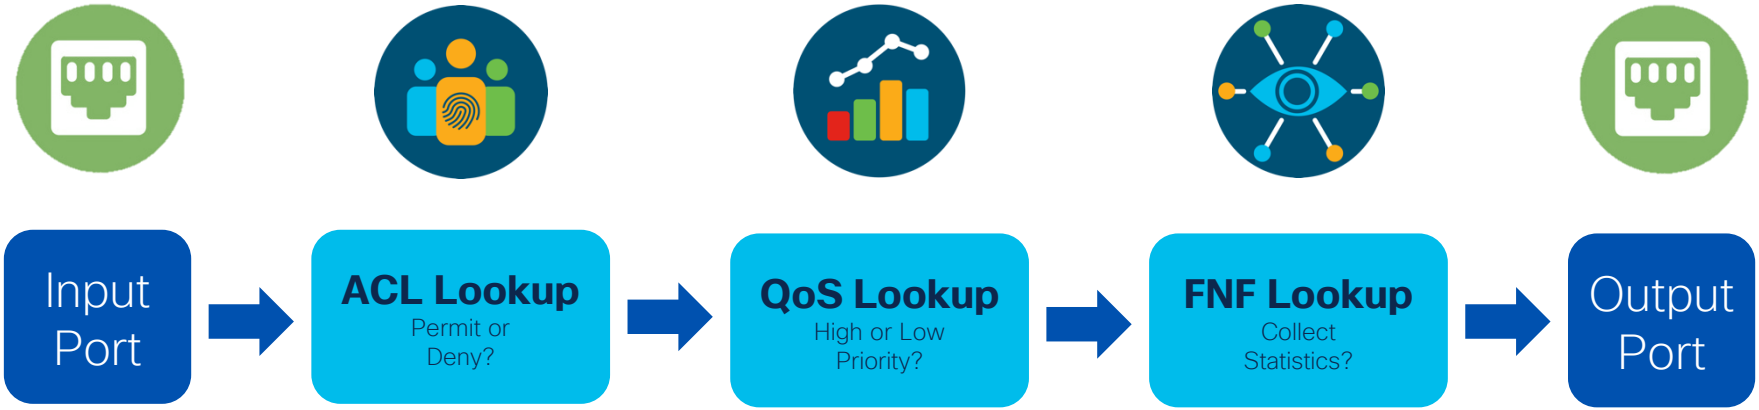

Common services are **Access Control**, **Quality of Service** and **Flow Analytics**

Traditional Network ASICs

Fixed Pipelines

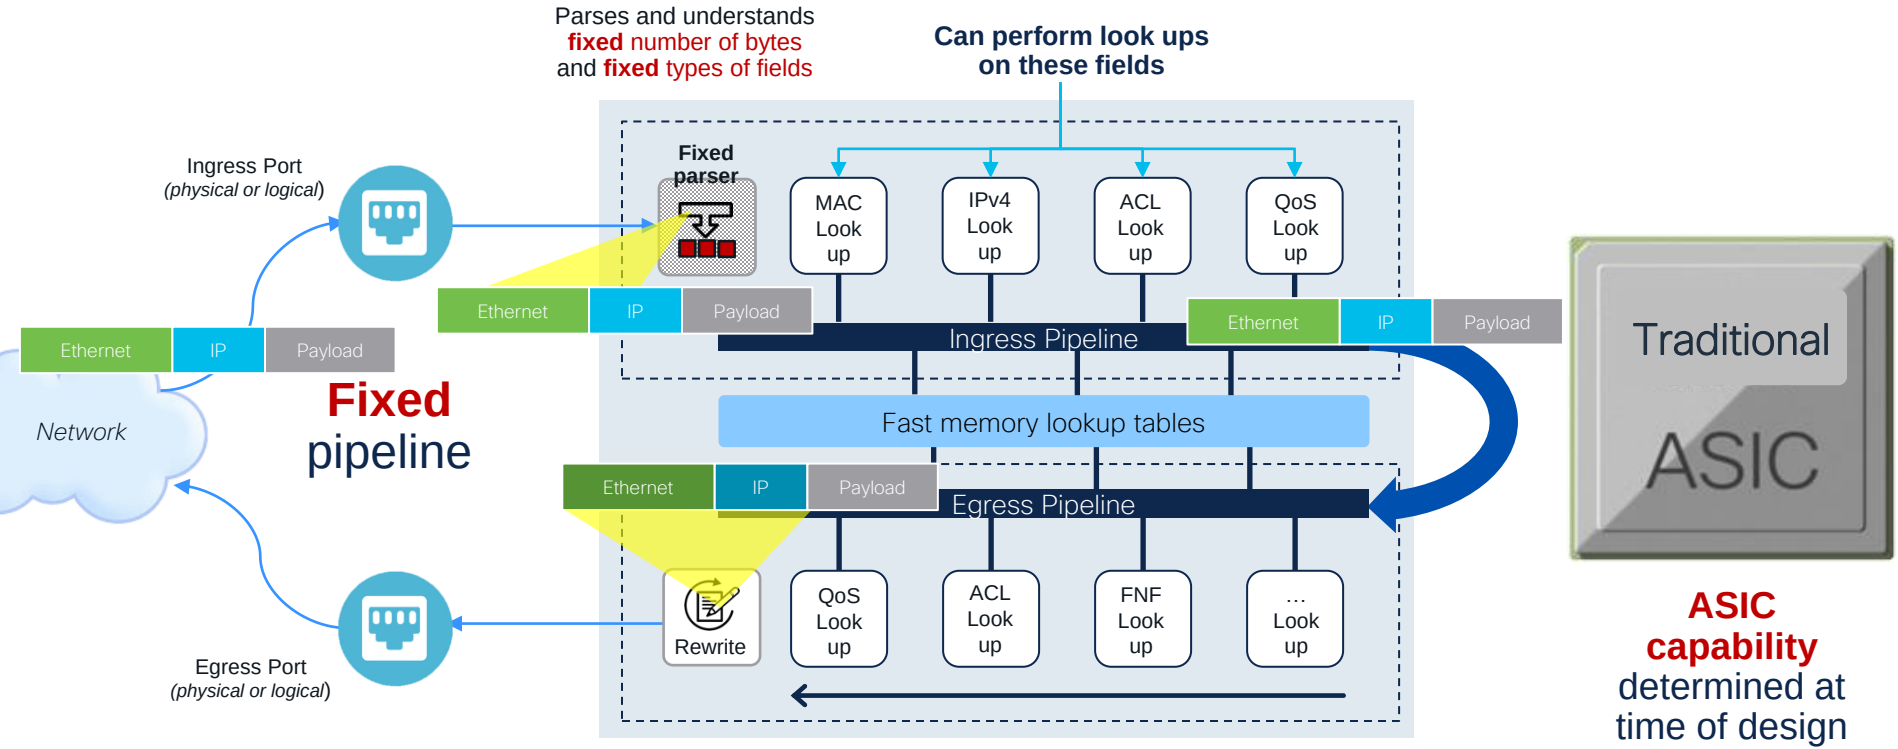

Flexible ASICs

for Enterprise Switching

Programmable Network ASICs

Flexible Parsing

Flexible Parsing

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

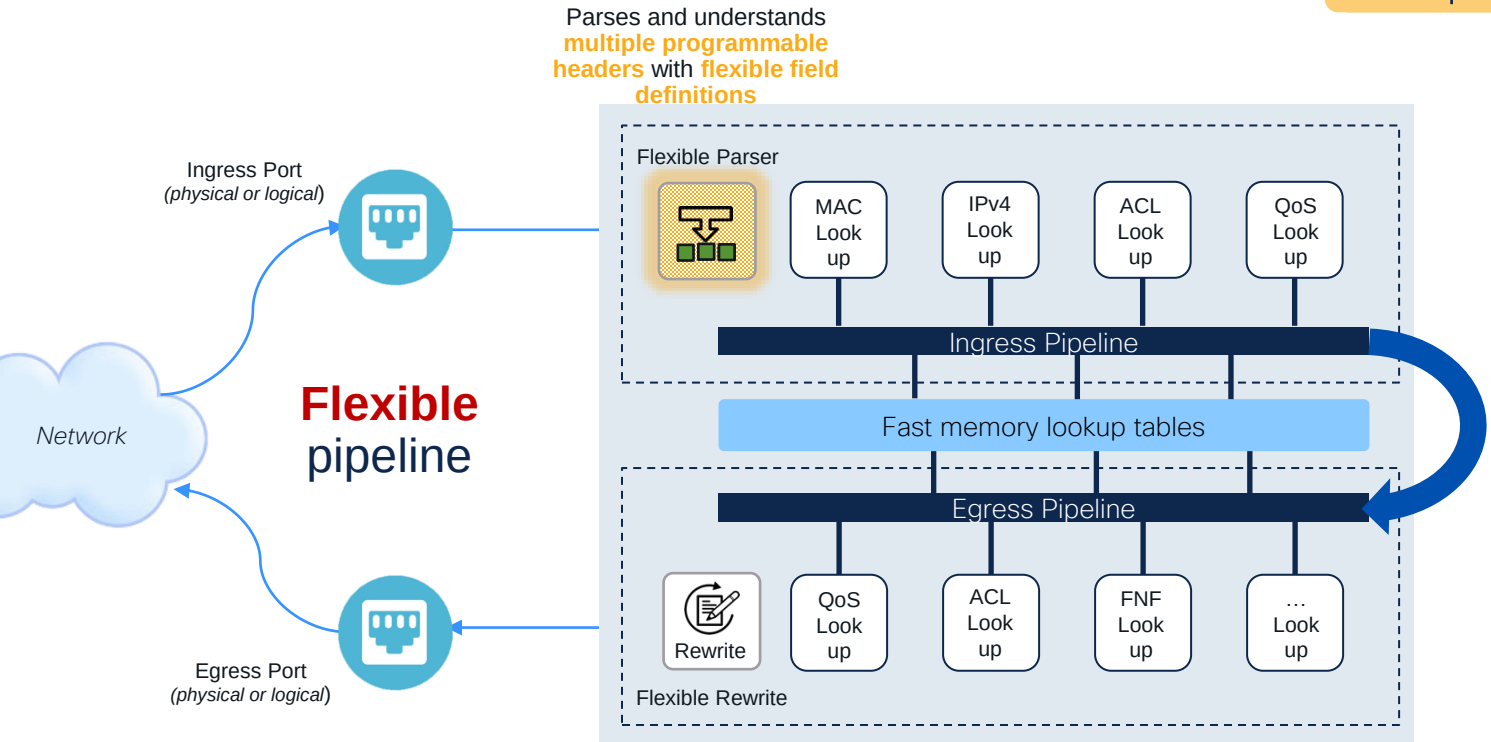

Programmable Network ASICs

Flexible Lookups

Flexible Lookups

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

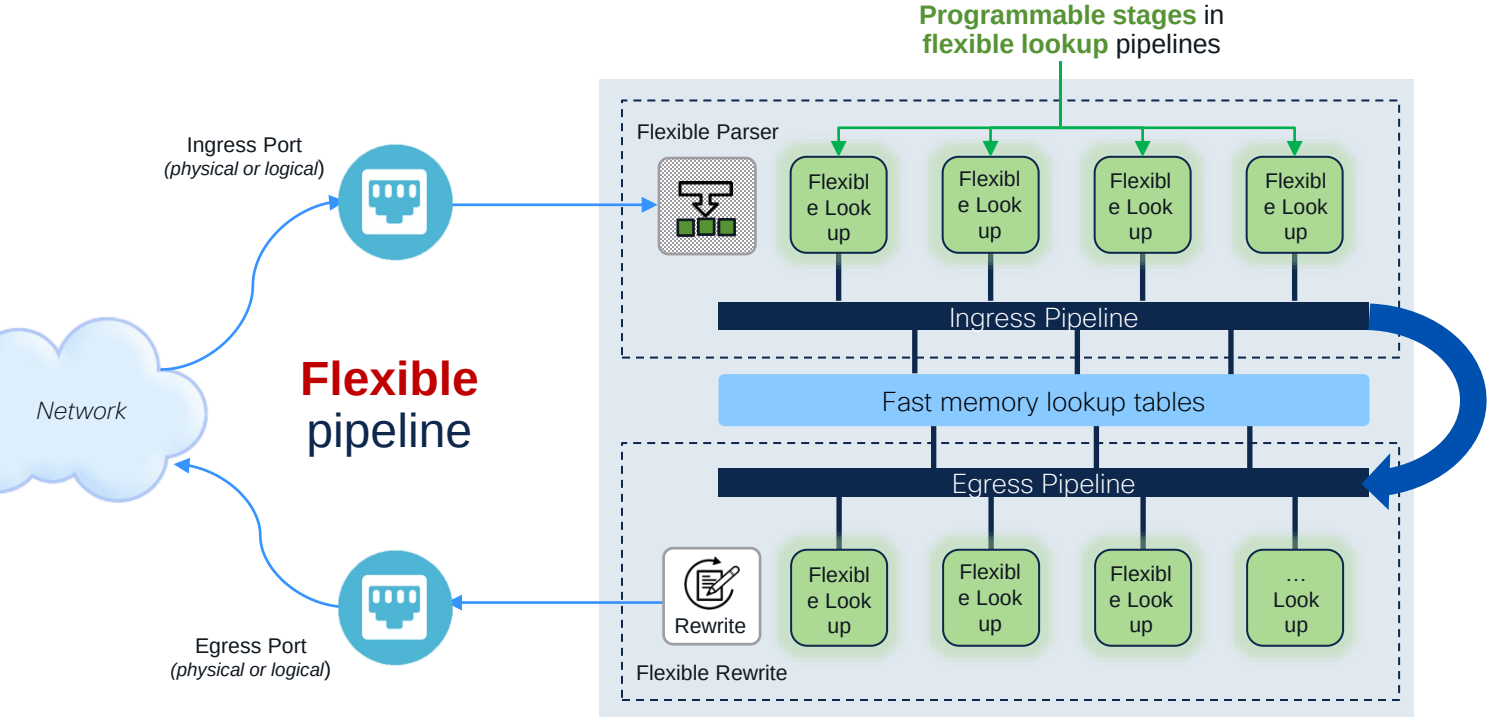

Programmable Network ASICs

Flexible Rewrites

Flexible Rewrites

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

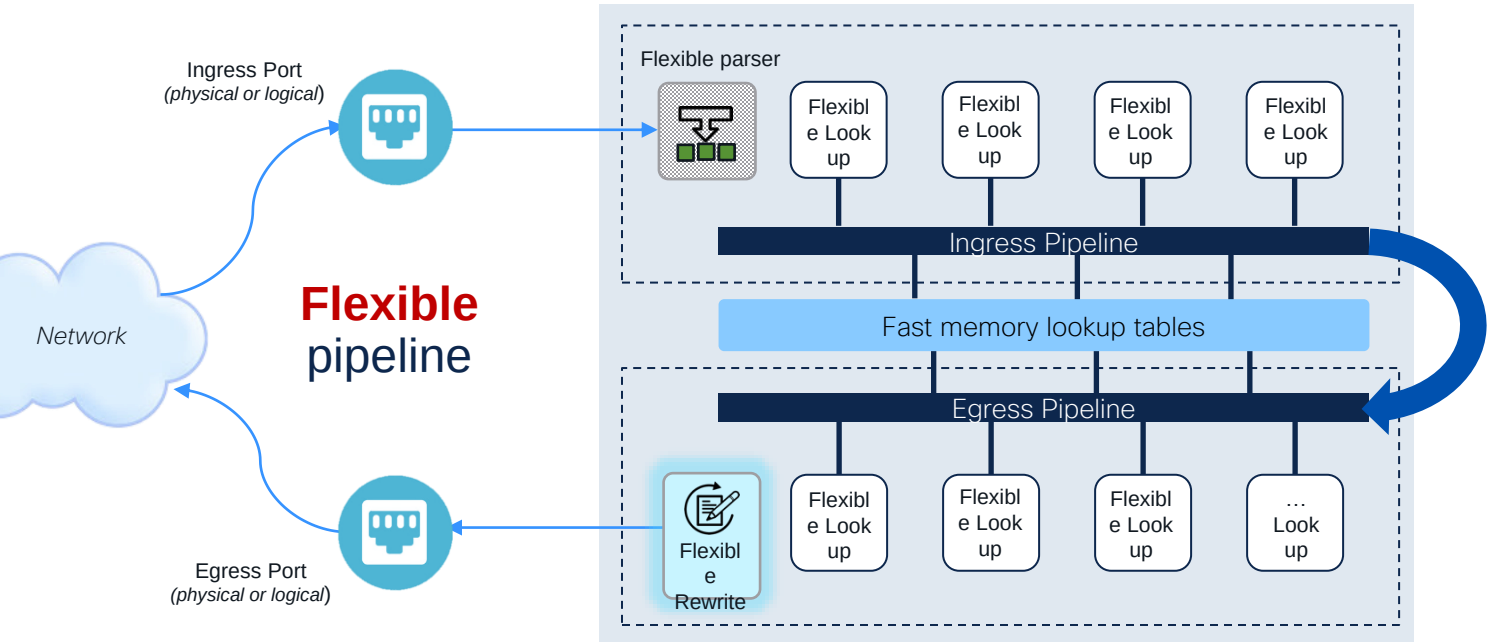

Programmable Network ASICs

Optimized Recirculation

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

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

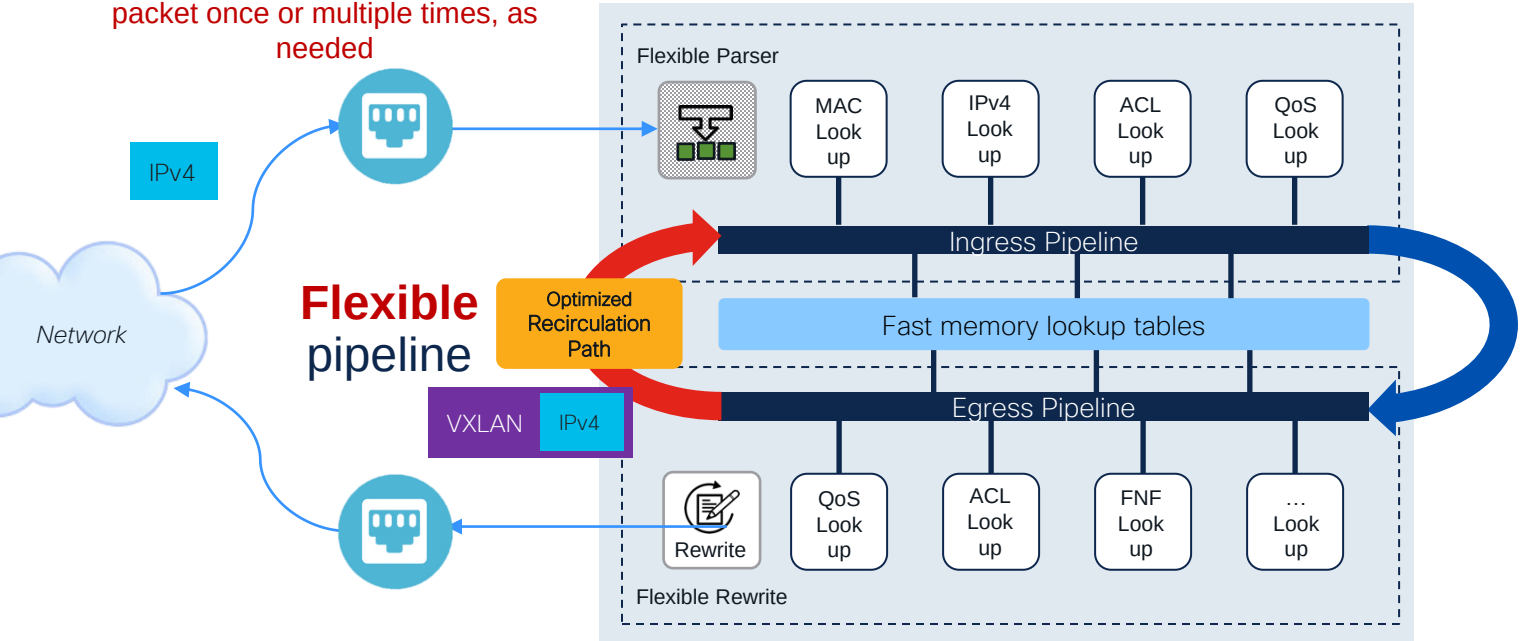

Programmable Network ASICs

Balancing Cost, Performance & Flexibility

Programmable ASIC

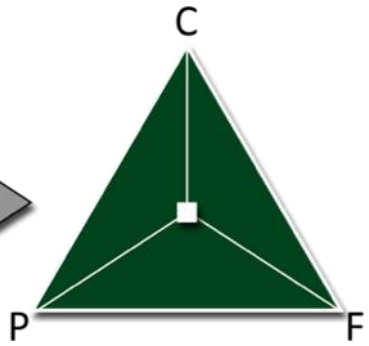

Cost ✓

Performance ✓

Flexibility ✓

ASIC Engineer

Microcode Update

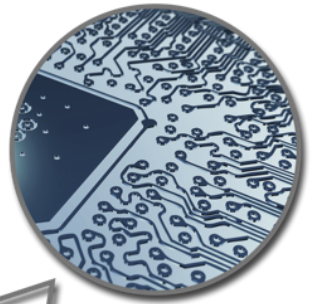

New ASIC Functionality

Creating Custom ASICs

From Definition to Deployment

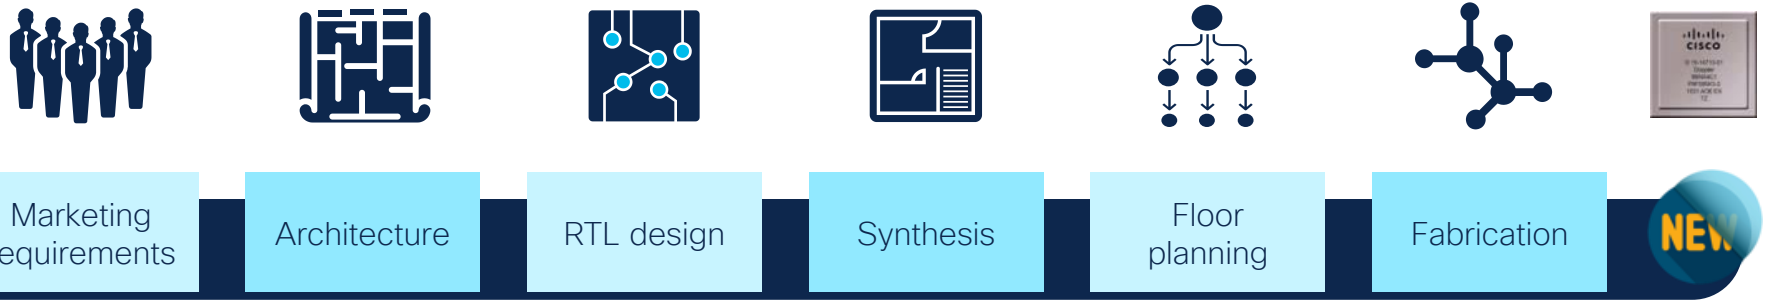

3 to 5 years

Building a **new ASIC** takes a lot of **time** and **money**

Cisco UADP

for Enterprise Switching

Cisco Unified Access Data-Plane (UADP®)

Common ASIC Architecture for Switching Access, Distribution & Core

UADP 2.0m	UADP 2.0/XL	UADP 2.0sec	UADP 3.0	UADP 3.0sec
120 Gbps	240 Gbps	480 Gbps	1.6 Tbps	1.6 Tbps
16nm FinFET	28nm FinFET	16nm FinFET	16nm FinFET	16nm FinFET
1.3B Transistors	7.6B Transistors	7.6B Transistors	192.B Transistors	192.B Transistors
1 Core + ARM CPU	2 Core	1 Core2 + SEC	2 Core	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

community.cisco.com/t5/networking-blogs/uadp-the-powerhouse-of-catalyst-9000-family/ba-p/3764605

Cisco UADP Platform Evolution

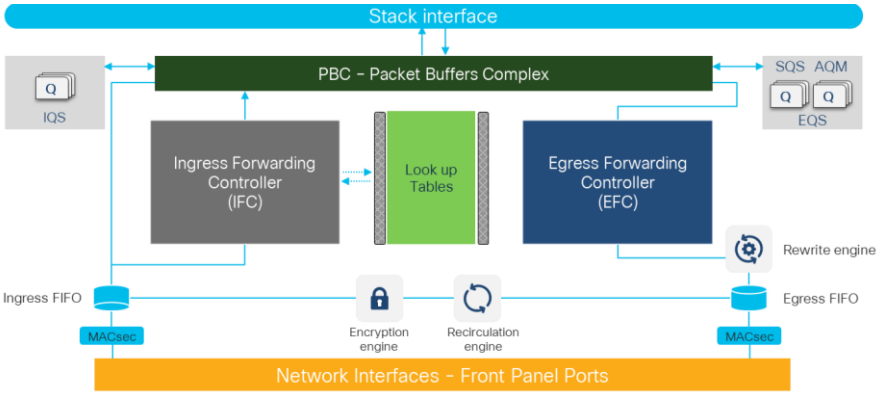

Catalyst 3850 – 2013
UADP 1.0 – 1.3B transistors

Catalyst 3850 mGig – 2015
UADP 1.1 – 3.0B transistors

Catalyst 9300 – 2017
UADP 2.0 – 7.5B transistors

Catalyst 9400/9500 – 2018
UADP 2.0 – 7.5B transistors

Catalyst 9500/9600 – 2019
UADP 3.0 – 19.2B transistors

Catalyst 9200 – 2019
UADP 2.0 mini – 3.2B transistors

Cisco UADP ASICs

ASIC Architecture & Block Diagram

up to 1 BILLION times per second!

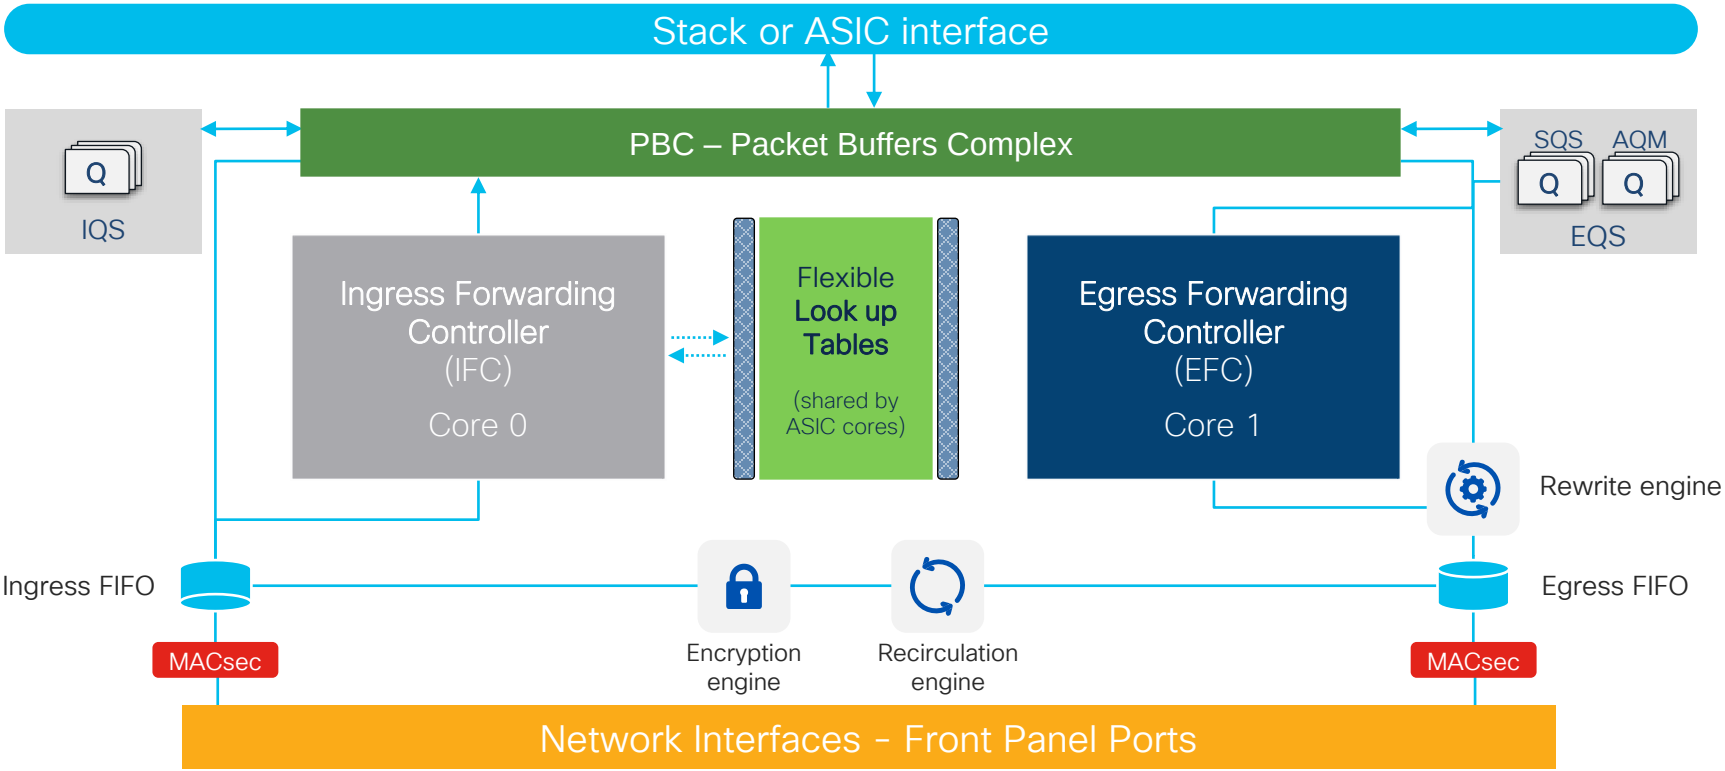

Cisco UADP – Dual Cores

Multiple ASIC Cores & Unified Packet Buffers

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

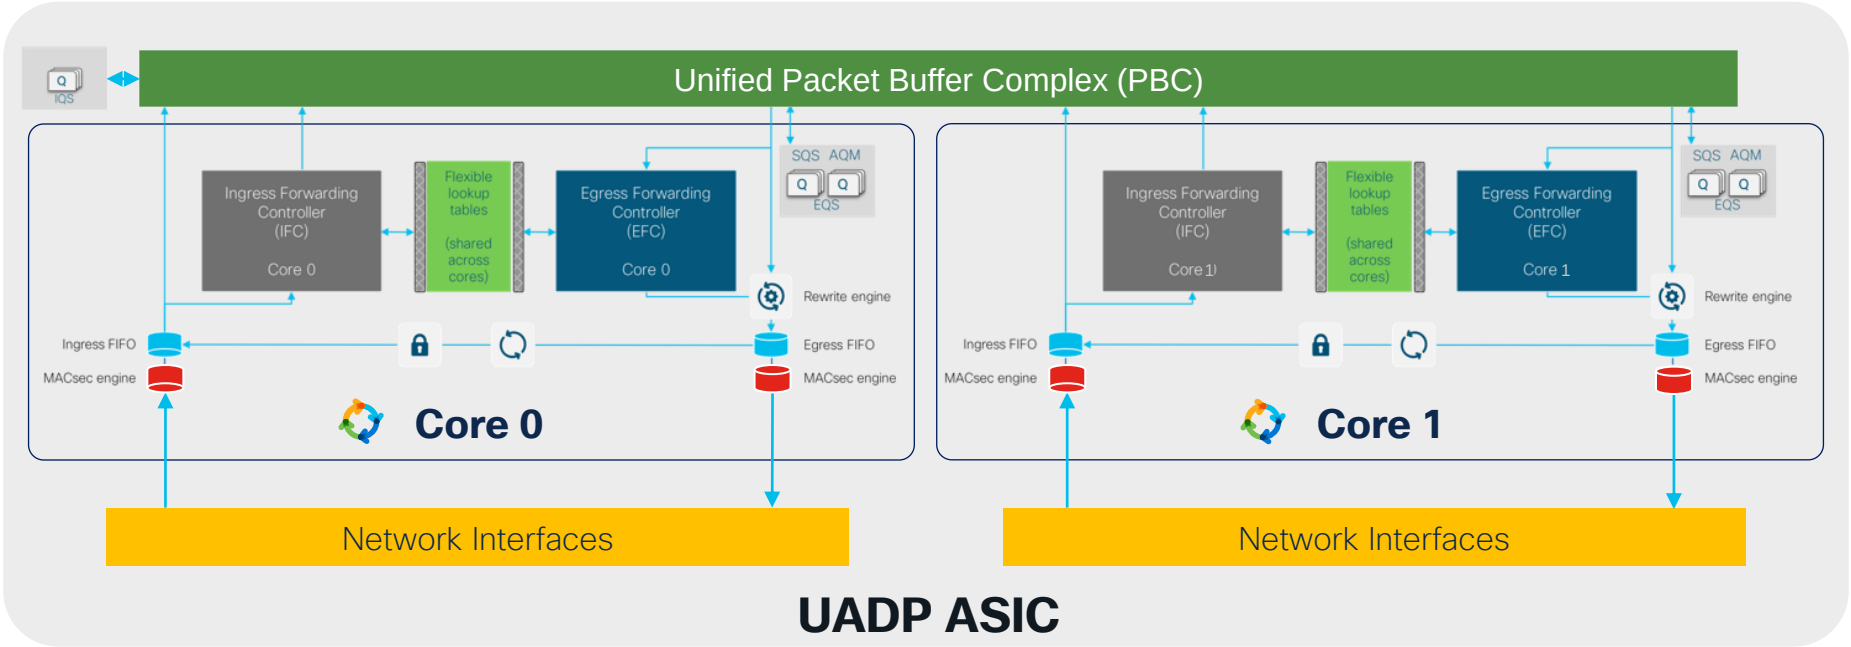

Cisco UADP – Flexible Lookups

Programmable Ingress and Egress Processing Stages

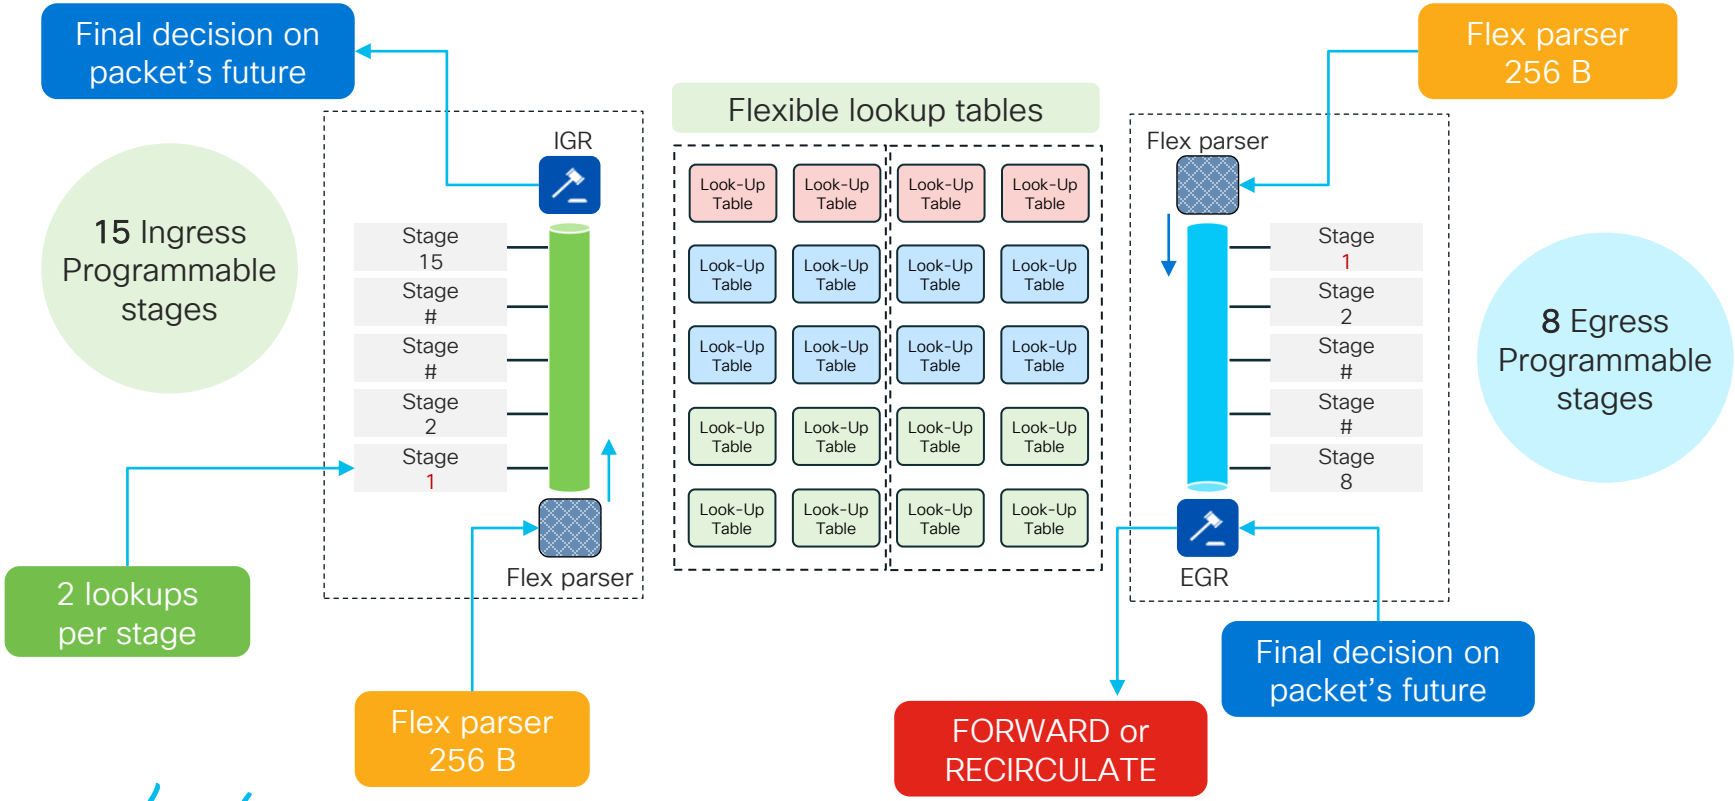

Cisco UADP – Microcode

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

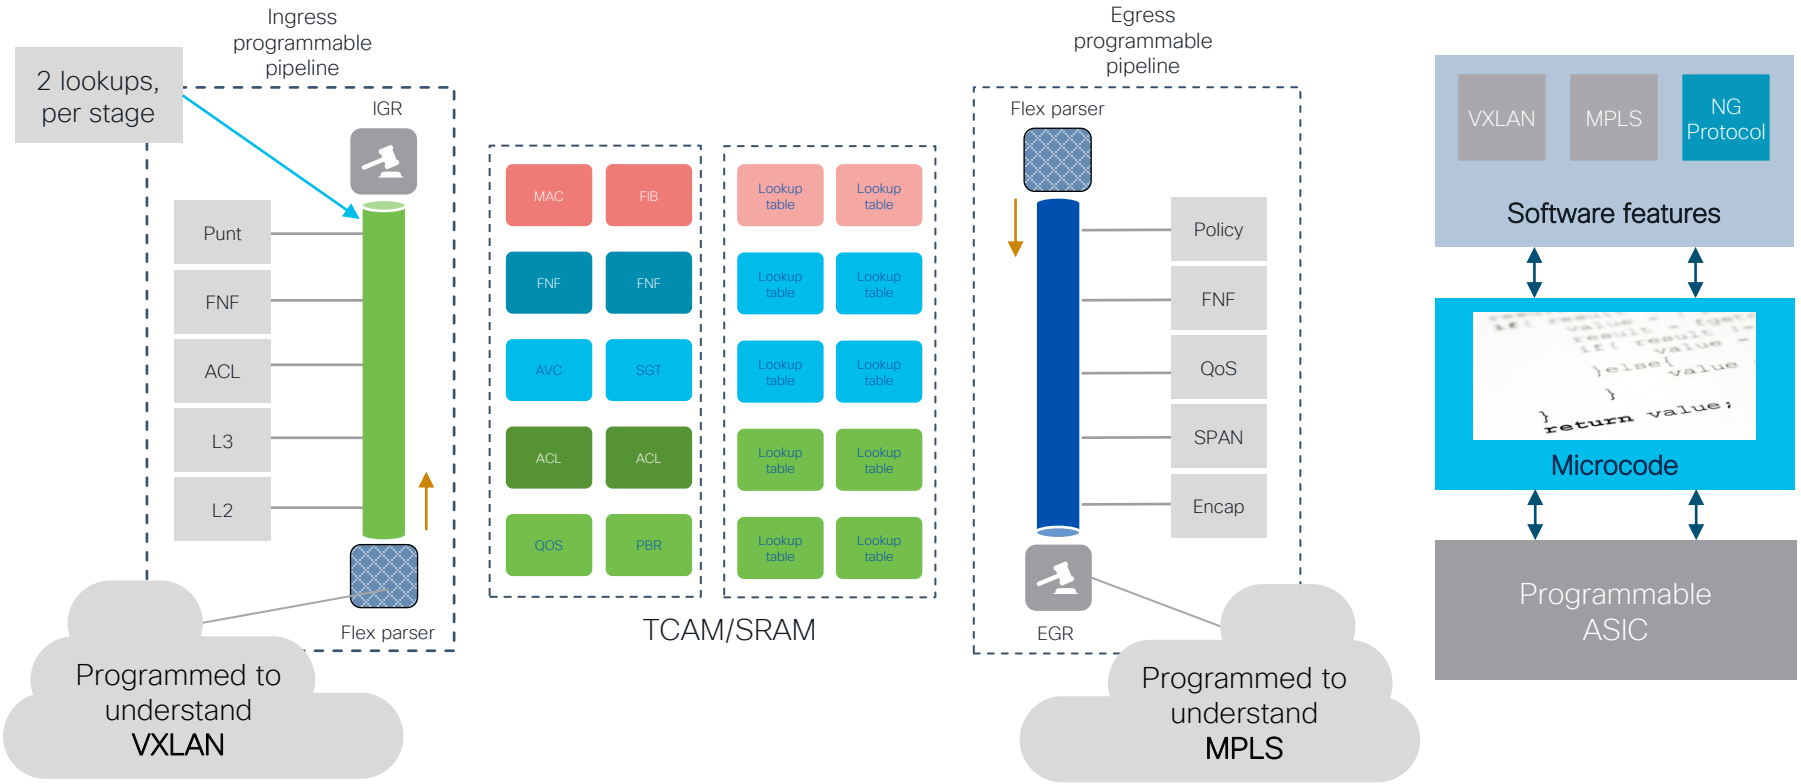

Cisco UADP – Flexible Tables

Customizable ASIC tables for universal deployment flexibility

SRAM/TCAM

- 64K
- 256K
- 32K
- 48K
- 16K
- 8K
- 4K

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

Flexible ASIC templates

Access template	Core template	Distribution template	NAT template
Layer 2 and Client scale	Layer 3 and Policy scale	Mix of Layer 2 and Layer 3 capabilities	Layer 3 and NAT scale

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

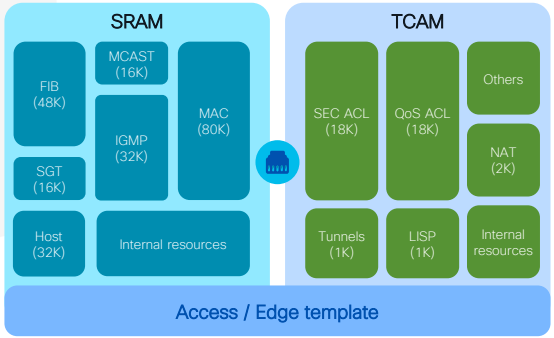

Table sizes can be tailored to support multiple templates

Cisco Catalyst 9500 & 9600 with UADP3

SDM templates and scale numbers

IOSXE 17.4.1

Feature		Distribution	Core (default)	SDA	NAT
Routes (IPv4/IPv6)		114K/114K	212K/212K	212K/212K	212K/212K
Multicast routes (IPv4/IPv6)		16K/16K	32K/32K	32K/32K	32K/32K
MAC address table		82K	32K	32K	32K
IGMP/MLD snooping		2K	2K	2K	2K
Flexible NetFlow (Ingress)		49K/ASIC	32K/ASIC	32K/ASIC	32K/ASIC
Flexible NetFlow (Egress)		49K/ASIC	32K/ASIC	32K/ASIC	32K/ASIC
SGT label		32K	32K	32K	32K
Security ACL	Ingress	12K		8K	12K
	Egress	15K		19K	8K
QOS ACL	Ingress	8K		8K	4K
	Egress	8K		8K	4K
NetFlow ACL	Ingress	1K		1K	1K
	Egress	1K		1K	1K
SPAN	Ingress	0.5K		0.5K	0.5K
	Egress	0.5K		0.5K	0.5K
PBR/NAT		3K		2K	15.5K
CPP		1K		1K	1K
Tunnel termination and MACsec		3K		3K	2K
LISP		1K		2K	1K

Higher scale

Customizable resources

Cisco UADP – Packet Walks

Generic Unicast Packet Walk – Same ASIC

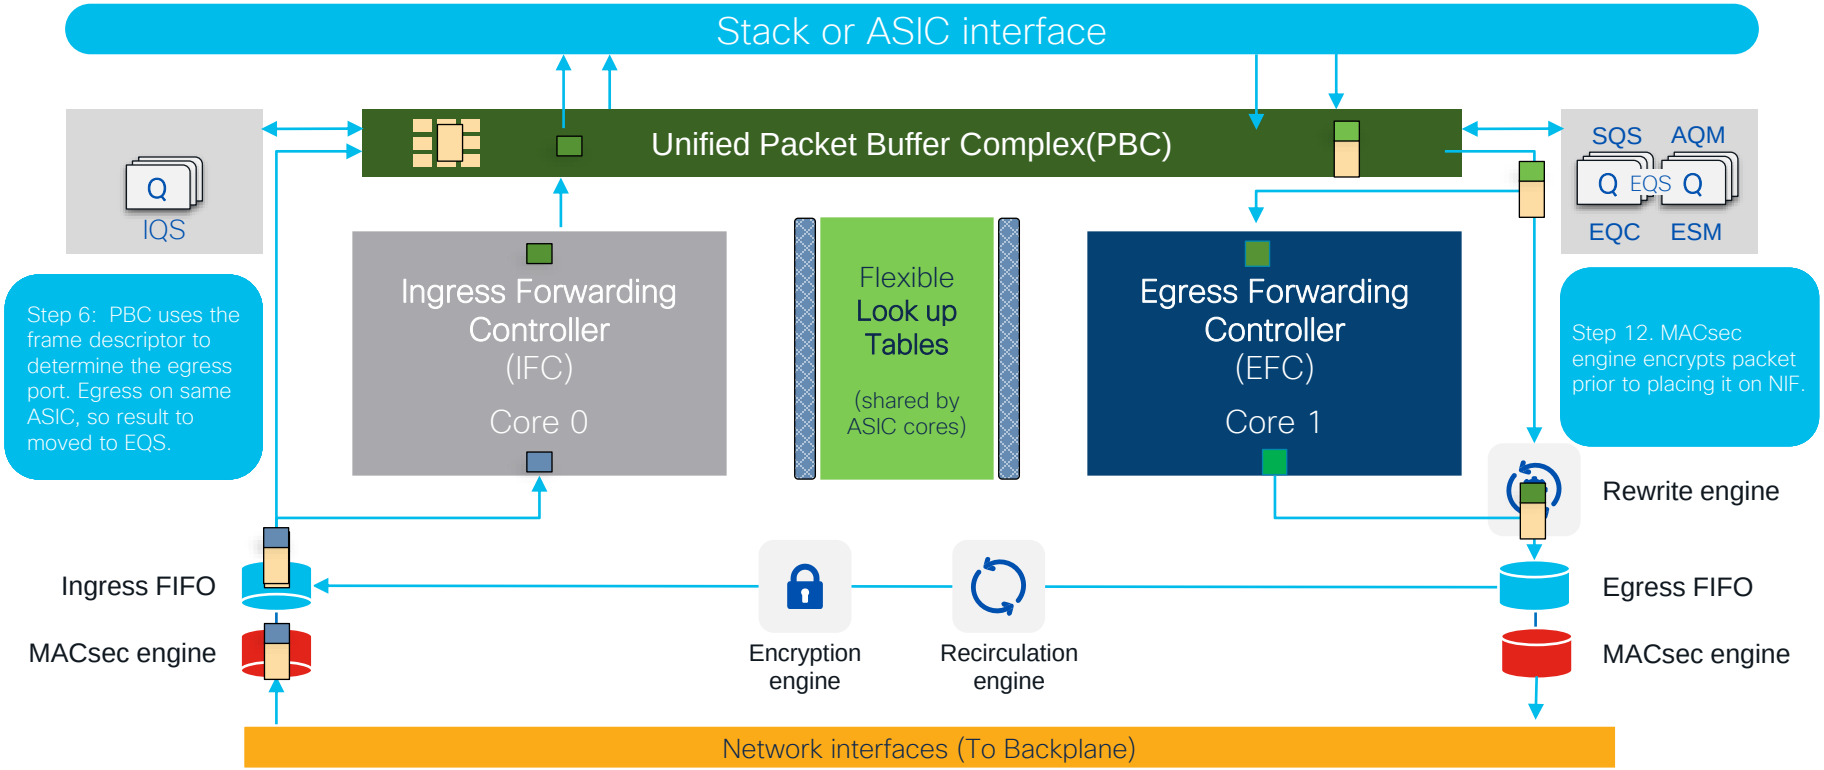

Cisco UADP – Multiple ASICs

Interconnecting Multiple ASIC Cores (Stacking or DPP)

Catalyst 9300X-48HX

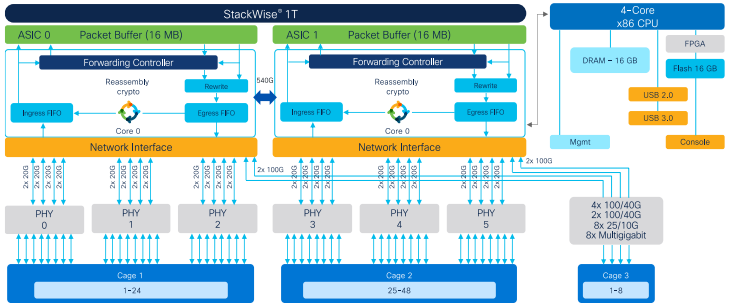

Catalyst 9400X-SUP2/XL

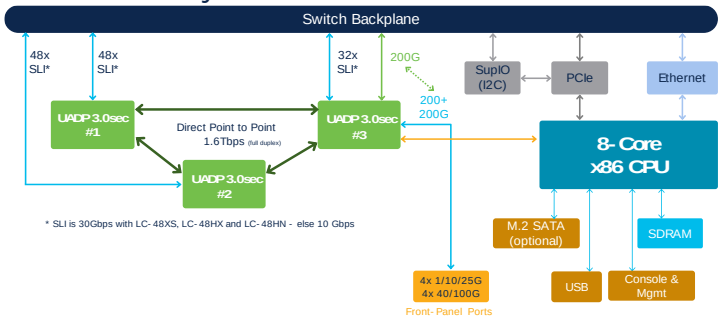

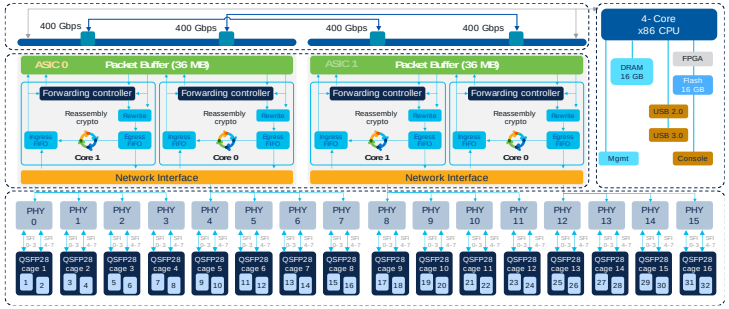

Catalyst 9500-32C

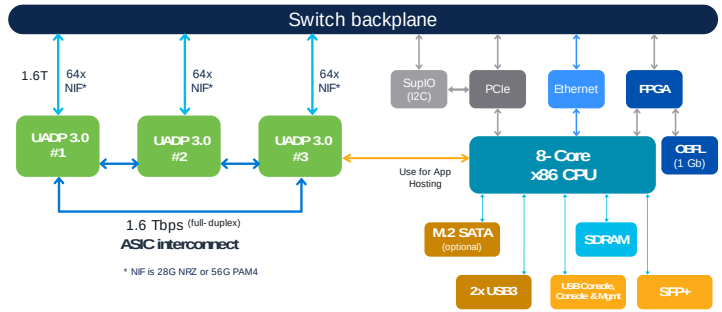

Catalyst 9600-SUP1

Cisco UADP – Packet Walks

Generic Unicast Packet Walk – Between ASICs

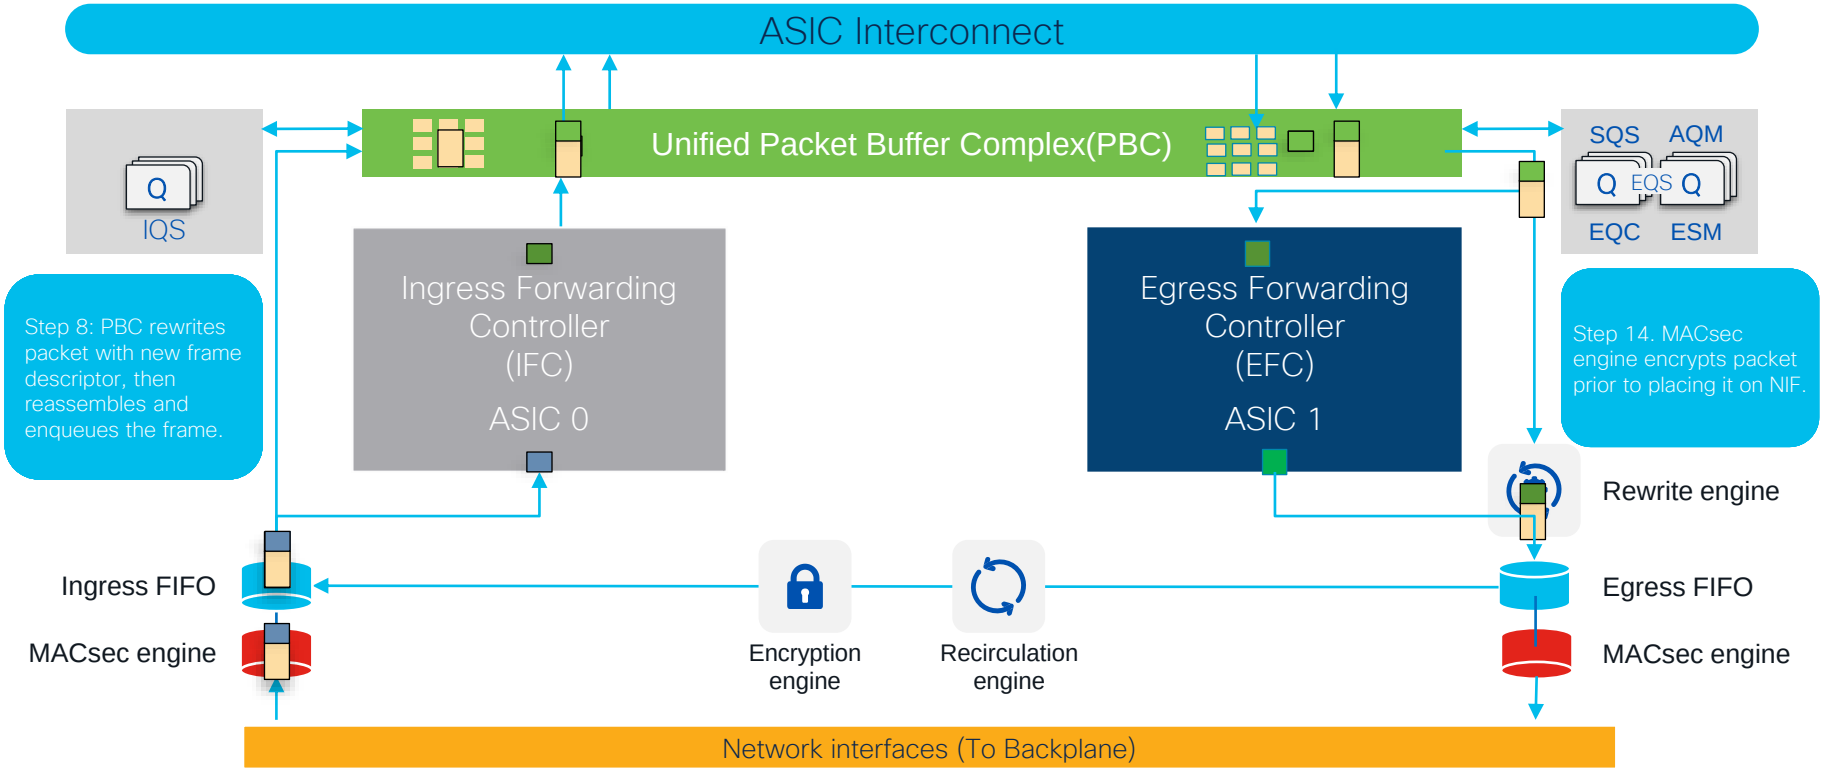

Cisco Silicon One

for Enterprise Switching

Introducing Cisco Silicon One™

One Architecture - Multiple Devices

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

www.cisco.com/c/dam/en/us/solutions/collateral/silicon-one/white-paper-sp-product-family.pdf

Introducing Cisco Silicon One™

One Architecture – Multiple Devices

Q202

3.2 Tbps

7nm FinFET

1 Slice SOC

Q201

6.4 Tbps

7nm FinFET

3 Slice SOC

Q100

10.8 Tbps

16nm FinFET

6 Slice SOC

Q200

12.8 Tbps

7nm FinFET

6 Slice SOC

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

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

www.cisco.com/c/dam/en/us/solutions/collateral/silicon-one/white-paper-sp-product-family.pdf

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 are connected via an integrated crossbar "fabric", using an integrated Virtual Output Queue (VoQ) buffer architecture, to manage traffic between Slices.

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

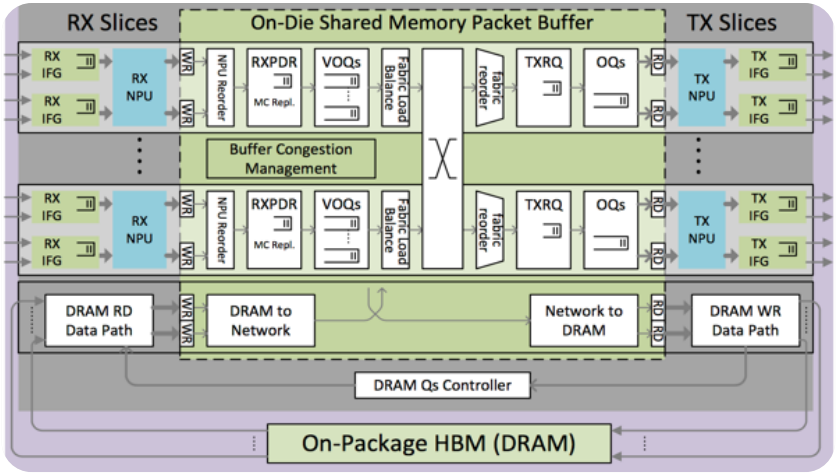

TIP

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

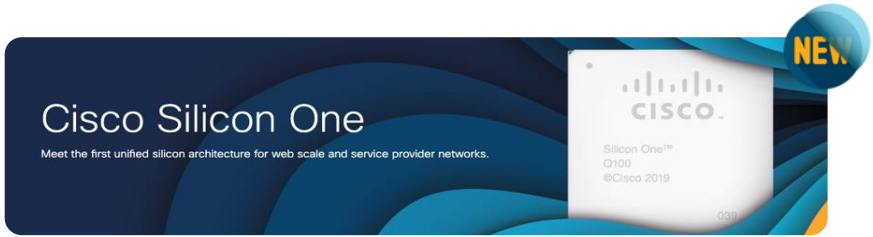

www.cisco.com/c/en/us/solutions/service-provider/innovation/silicon-one.html

Cisco Silicon One™ Q200

Industry leading Switching and Routing Silicon

Switching Silicon

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

Routing Silicon

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

Cisco Silicon One Bringing Switching and Routing convergence

Cisco Silicon One™ Q200

Industry leading Switching and Routing Silicon

12.8 Tbps

8.1 Bpps

8GB HBM for deep buffers

2M IPv4 or 1M IPv6 routes

P4 NPL Programmable Pipeline

56G PAM4 Serdes

First 7nm ASIC
providing lowest Watts/GE power consumption

Fully P4 programmable
enabling feature velocity

Multi-Slice architecture
for flexibility and scale

Routing Capabilities with Switching Power and Performance

Cisco Silicon One™ Q200

ASIC Architecture & Block Diagram

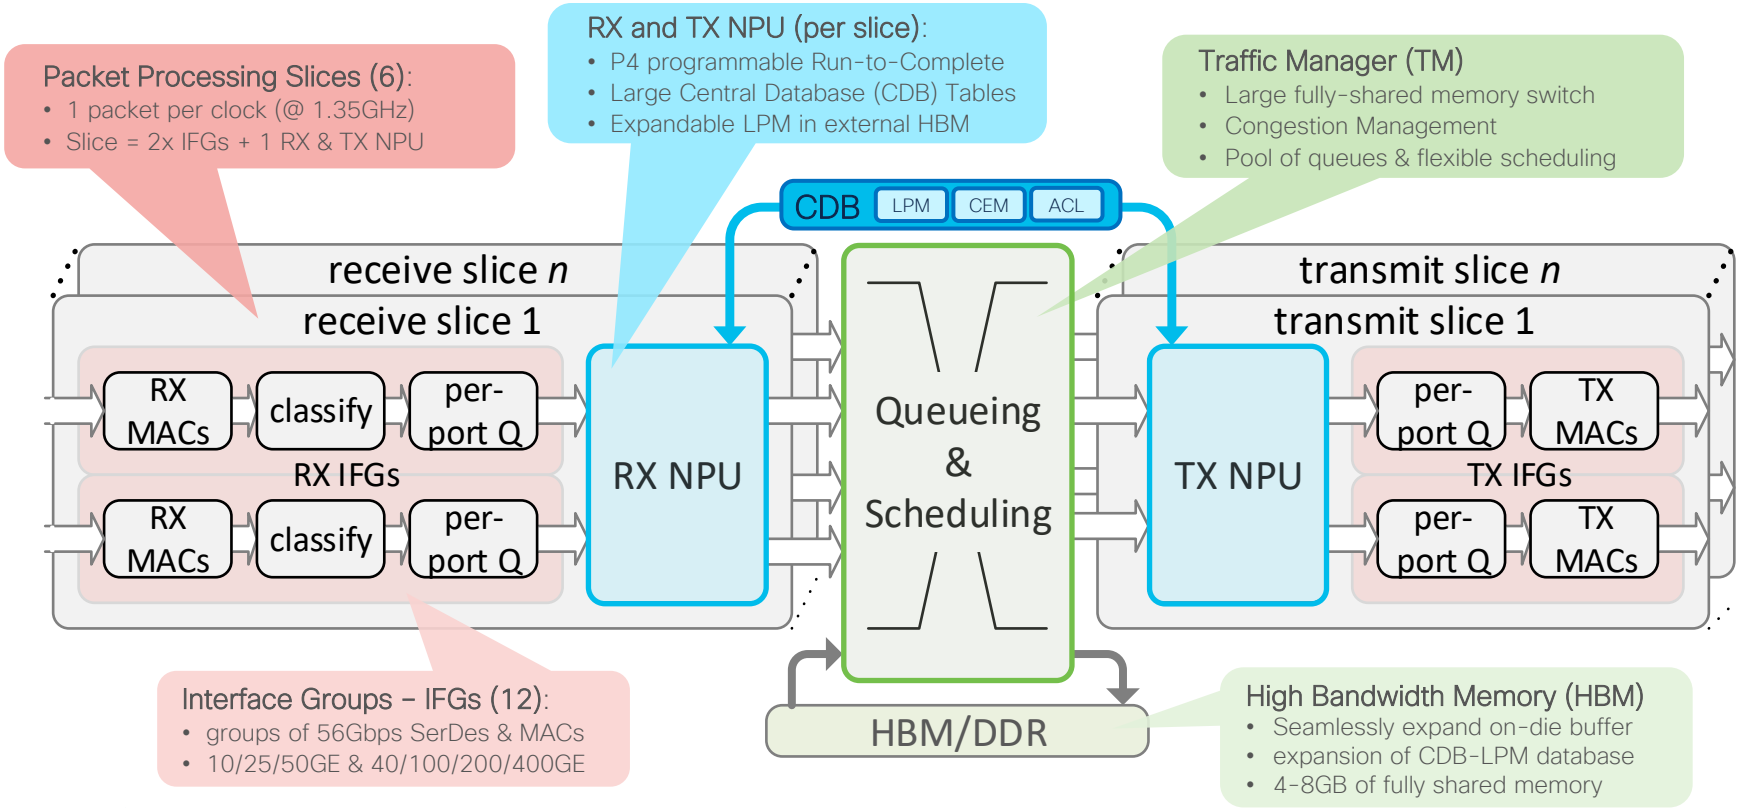

Cisco Silicon One™ Q200 – Buffer Queuing

Local SMS + On-Package HBM

Shared Memory Subsystem (SMS)

- Fully shared across slices
- **108MB** (80MB for data) internal buffers
- **multicast** buffering – **write once, read many**

~12.8T bps = ~8.1B packets-per-second

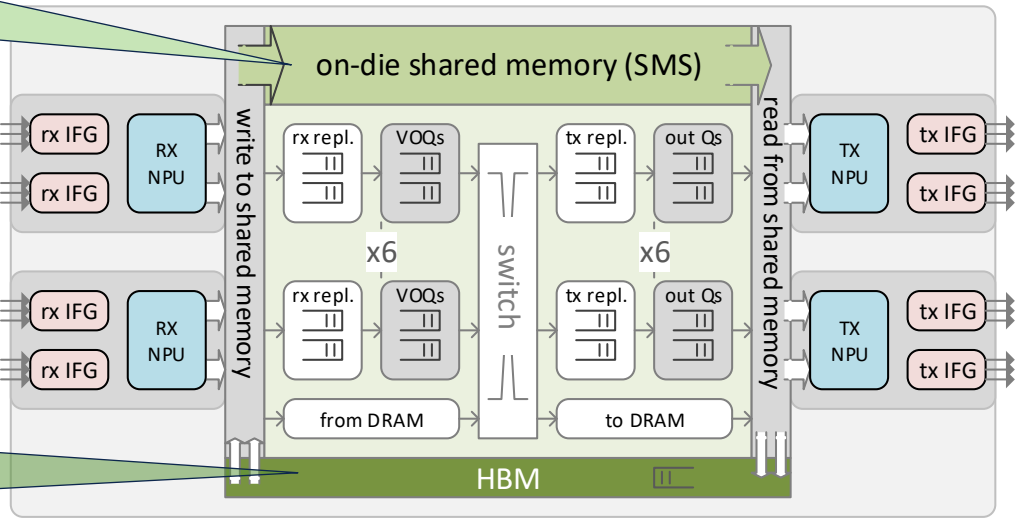

High Bandwidth Memory (HBM)

- Fully shared across slices
- **2.4 Tbps full-duplex** (in and out)
- **8GB memory** external buffers

blogs.cisco.com/sp/optimize-real-world-throughput-with-cisco-silicon-one

Cisco Silicon One – HBM

High Bandwidth Memory

- Augments local on-die memory
  - use local (SMS) buffers until full
  - use HBM for bursts or congestion
- Deep buffering + FIB expansion on-package, at high-speed
- 2 x stacks of 2.5D memory with wide-bus interposer = ~2.4 Tbps full duplex
- Interposer connects ASIC die to on-package HBM memory

en.wikipedia.org/wiki/High_Bandwidth_Memory

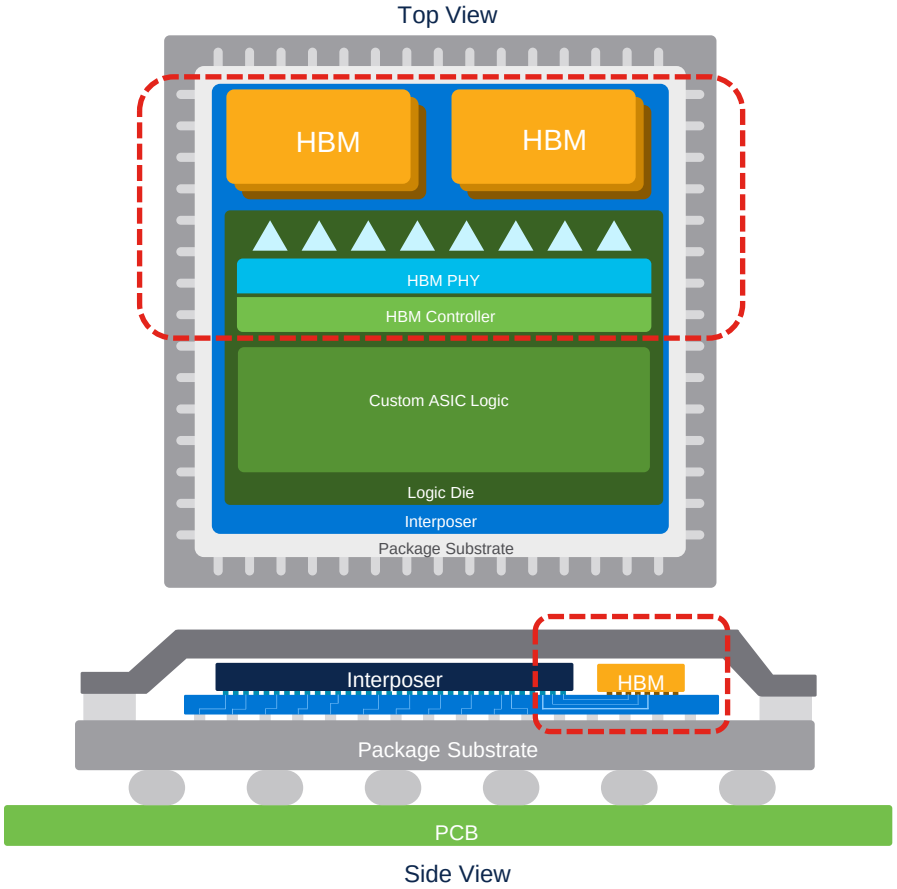

Cisco Silicon One™ Q200 – Central Databases

Onboard LPM, CEM & ACL memory

Q200 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
- Up to* 2M IPv4 route entries, or 1M IPv6 route entries
- 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)
- Up to 608K IPv4 entries, or 304K IPv6 entries
- 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
- Up to 8K IPv4 ACL entries, or 4K IPv6 ACL entries
- OG/SGACLs use CEM, with only action ACEs in TCAM

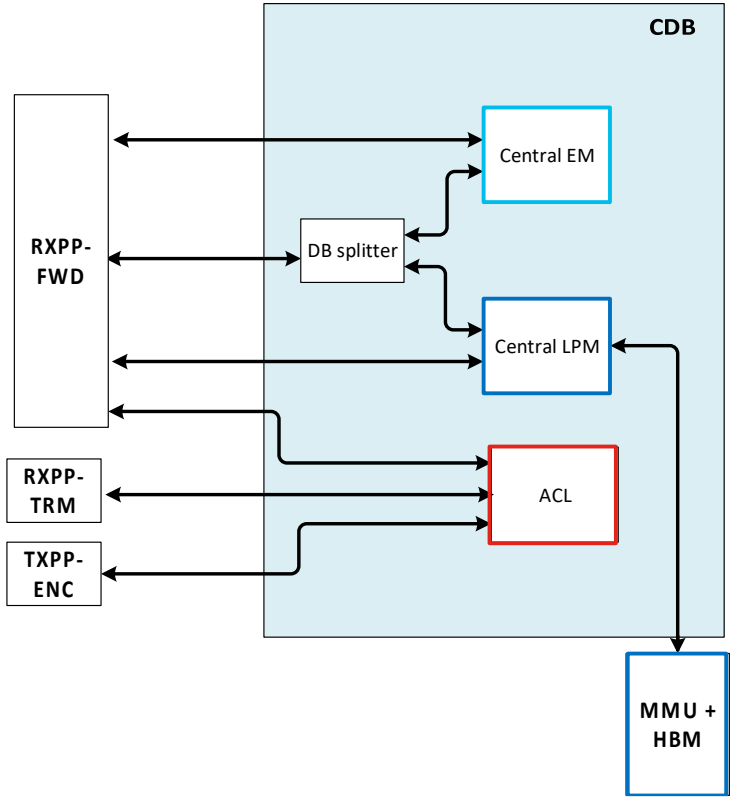

* Exact scale depends on IP/mask distribution (contiguous vs. random) and hash efficiency. Sample tests with IPv4 GRT is ~1.85M

Cisco Catalyst 9500 & 9600 with Q200

IOSXE 17.7.1

SDM templates and scale numbers

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

L3 Routes (LPM)	FCS+	2 million / 1 million	1 million to 2 million	@ 1 million step

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

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

Cisco Silicon One Q200

Generic Unicast Packet Walk

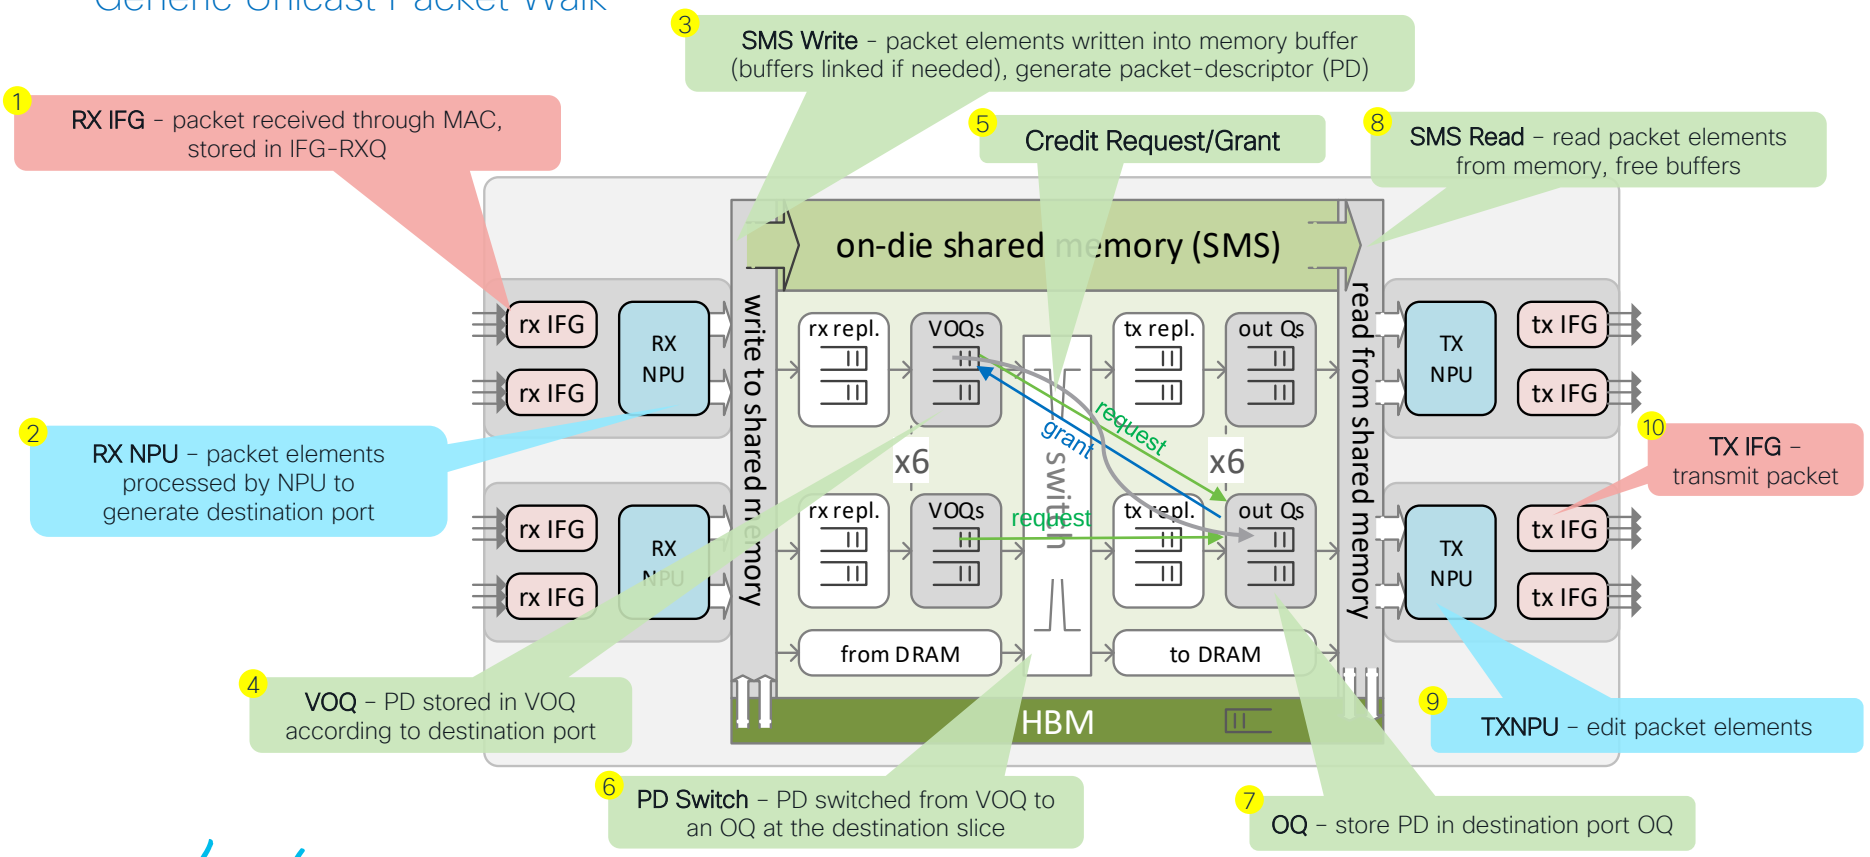

Cisco Silicon One – Q200

Combining multiple ASICs in One SOC

Catalyst 9500X-28C8D &

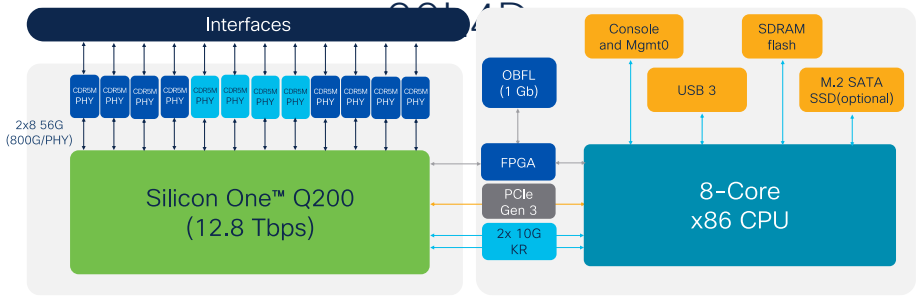

Catalyst 9600X-SUP2

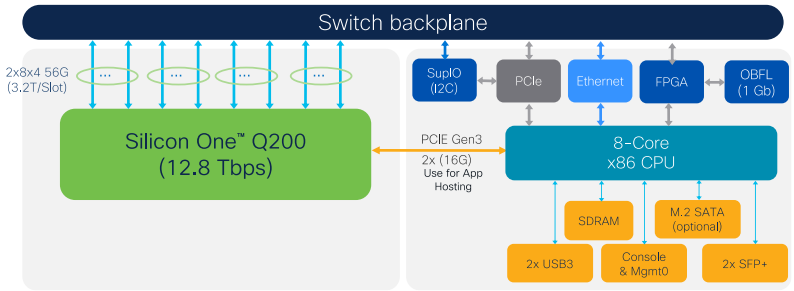

Cisco Catalyst 9000 “X” Series

Extending Enterprise Switching

New additions to Catalyst 9300X Series

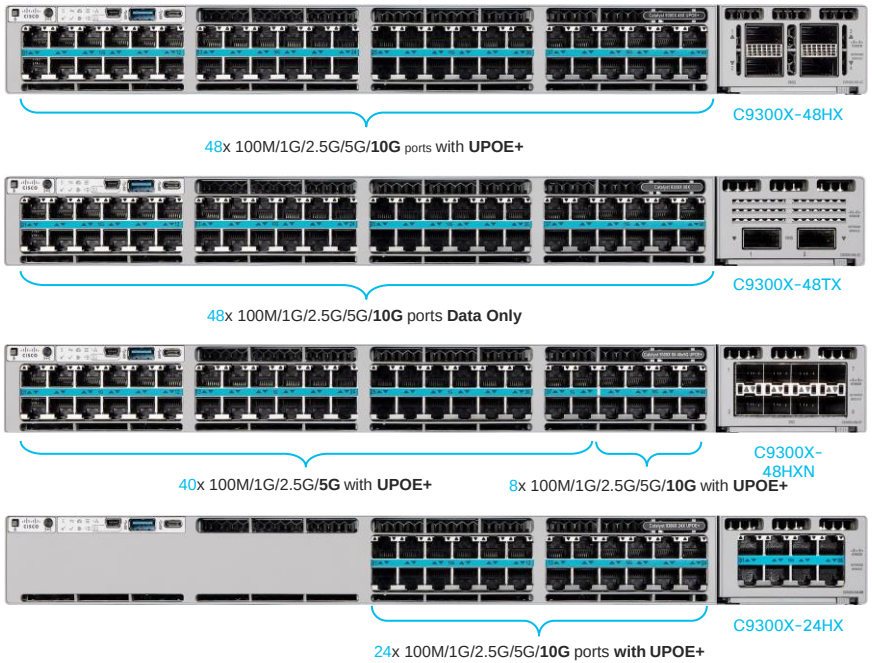

100/40G Modular Uplinks

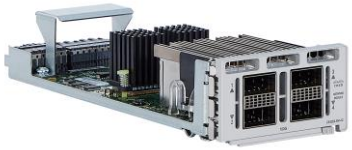

C9300X-NM-4C
4 x 100/40G QSFP

C9300X-NM-2C
2 x 100/40G QSFP

Multigigabit Uplinks

C9300X-NM-8M
8 x 10G-mGig

10/25 G Modular Uplinks

C9300X-NM-8Y
8 x 25/10/1G

Introducing Catalyst 9400X Gen2 family

Performance meets flexibility and investment protection

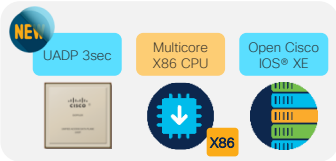

6
No Oversubscription
480 Gbps per Slot

100G support

Compatible with Gen1 Line-Cards

Cisco UPOE® across all ports

Catalyst 9400 SUP-2/SUP-2XL

C9400-LC-48HX
(48x 10G Multigigabit UPOE+)

480G /Slot

C9400-LC-48XS
(48x 10G SFP+)

100G Layer 3 encryption

108 MB unified buffer

Flexible ASIC Scale

Densest MultiGigabit + Fiber

Cloud-native agility

- **6x increase in throughput**
- 5x increase in uplink density with 100G
- Line rate across all chassis and line cards
- Up to 384 ports of Multigigabit and UPOE
- 108 MB unified packet buffer

Industry first

Secure, connected experience

- 100G Layer 3 encryption
- Cloud-scale multilayer security
- 256-bit LAN/WAN MACsec
- 2x increase in app hosting with QAT

Industry first

Business-first resiliency

- High Availability with ISSU & sub-second SSO
- Built-in Compute (App Hosting) resources
- End-to-end Analytics and Assurance

Catalyst 9500 & 9500X Series

Extending High-Performance Fixed Core with a Performance-Optimized Edge Switch

Non-XL Scale

Catalyst 9500 (UADP 3.0)

C9500-32C / 32QC

C9500-48Y4C / 24Y4C

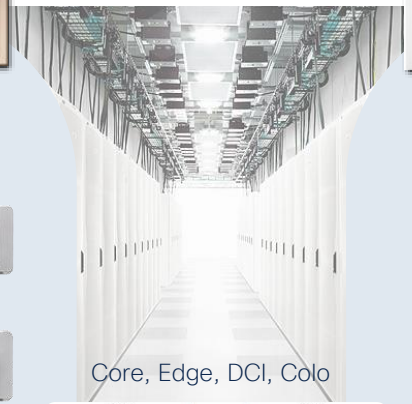

Core, Edge, DCI, Colo

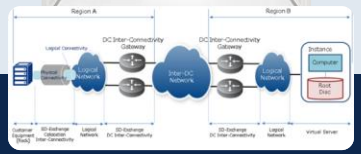

Total Capacity

3.2 Tbps

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

XL Scale

NEW

Catalyst 9500X (S1 Q200)

C9500X-28C8D — 17.7.1

C9500X-60L4D — 17.10.1

2x

Total Capacity

6.0 Tbps

32 x 100G + 8 x 400G
or 60 x 50G + 4 x 400G

Catalyst 9600 & 9600X Series

Extending Modular Core with a Performance-Optimized Supervisor & Cards

Supervisor 1

Total Capacity

4.8 Tbps

Slot B/W

1.2 Tbps


SS
D


Fan-Tray


4x
PSU

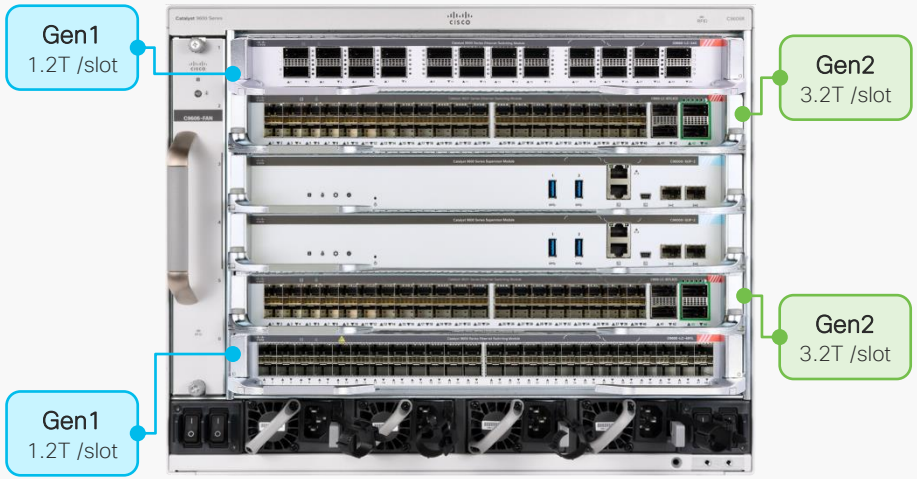

NEW

Supervisor 2

Total Capacity

12.8 Tbps

Slot B/W

3.2 Tbps

3 X

40x 1/10/25/50G + 2x 200G + 2x 400G NEW

32x 40/100G or 24x 40/100G + 2x 400G NEW

Glimpse into the Future

for Enterprise Switching

Where are things going?

Speeds and Feeds

More details in PPT Notes

25/50G → 100G SFP

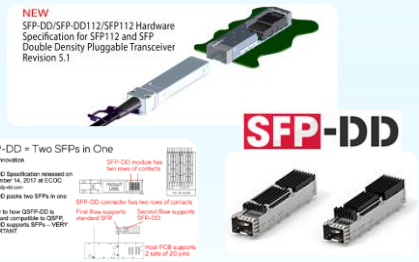

SFP-DD(112): 100Gbit/s using PAM4

100/400G → 800G QSFP

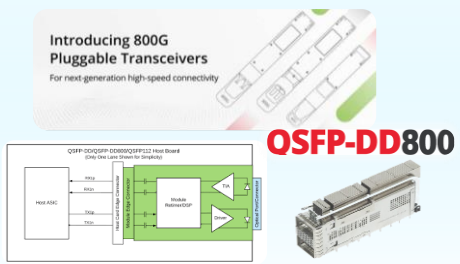

QSFP-DD(112): 800Gbit/s using PAM4

Coherent Optics

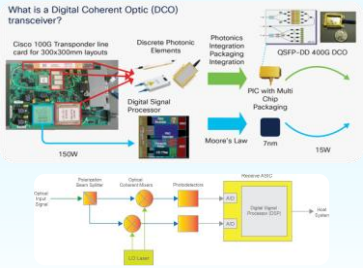

Uses phase of optical signal vs. power level

Direct-Drive (PHYLESS)

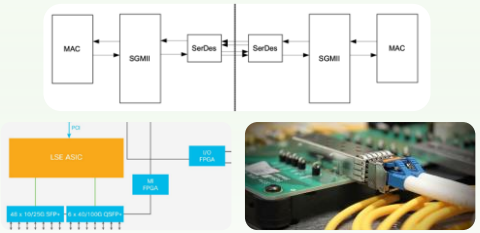

Transceiver uses SFI/SGMII to connect directly to the ASIC SerDes (no PHY)

4x100G & 8x50G Breakout

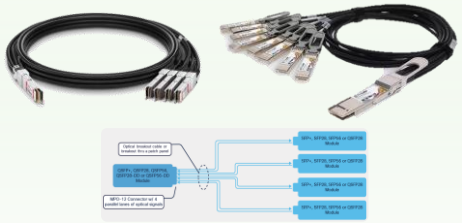

QSFP-DD (400G) can be split into 4 x 100G QSFP28 or 8 x 50G SFP56

Single Lambda

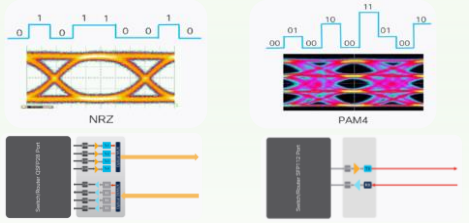

Uses a single 112G PAM4 signal vs. 4 x 28G NRZ signals

The bridge to possible

Cisco IOS XE Software Architecture & Innovations

Catalyst 9000 Series

Shawn Wargo, Principal TME
@Shawn_Wargo

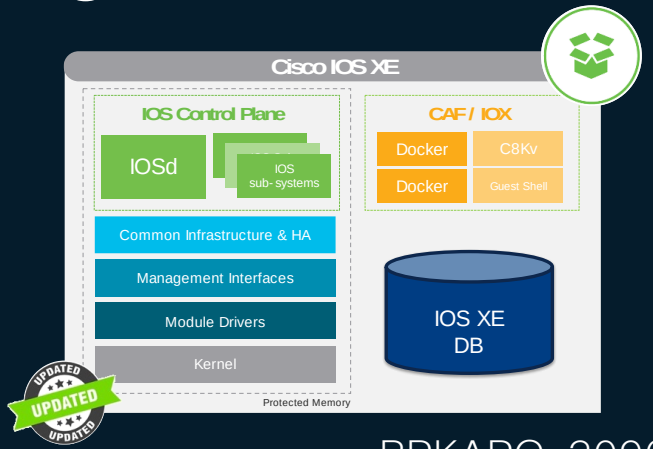

Session Abstract

You've used the CLI commands. You've tried the Programmable API's.
You've seen the GUI screens.... but what's **REALLY** happening inside **IOS XE**?

- How is Cisco IOS XE similar or different from other Cisco OS architectures?
- Do you wish you could look –inside- the different software layers, and understand how they all work together (or when they don't, for troubleshooting)?

This session will focus on the **software components** (processes) within IOS XE on the Catalyst 9000 series:

1. Recap of the basic principles and history of IOS
2. Summary of basic components of Catalyst 9000 IOS XE
3. Summary of key technologies enabled by IOS XE
4. Catalyst 9000 IOS XE architecture up to 17.6.1 release
5. Catalyst 9000 IOS XE architecture after 17.7.1 release

What this session is NOT

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**:

- Catalyst 9000 Series Switching Family – Access – BRKARC-2098
- Catalyst 9000 Series Switching Family – Core and Distribution – BRKARC-2099
- Catalyst 9000 IOS XE Innovations – BRKENS-2004

Agenda

1. Brief History of IOS XE
2. Basic IOS XE Components
3. IOS XE Technologies
4. C9K IOS XE upto 17.6.1
5. C9K IOS XE after 17.7.1 NEW
6. Summary & References

Catalyst 9000 Series – Common Building Blocks

Programmable x86 Multi-Core CPU

Application Hosting
Secure Containers

Open IOS XE® Polaris

Model-Driven APIs
Modular Patching

Cisco UADP & Silicon One™

Programmable Pipeline
Flexible Tables

Same IOS XE image for both **UADP*** and **Silicon One** C9K platforms

* Catalyst 9200 Series uses IOS XE Lite

Multi-Core CPU – Built for App Hosting

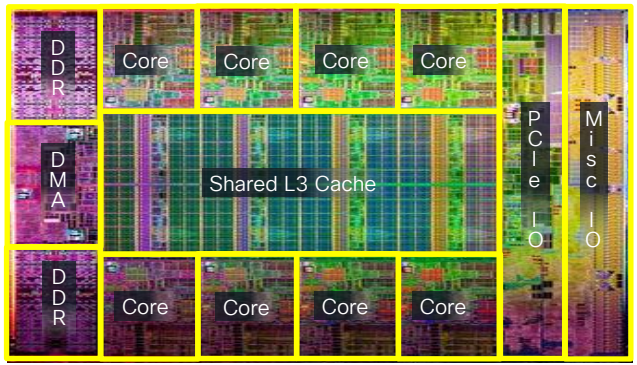

x86 CPU

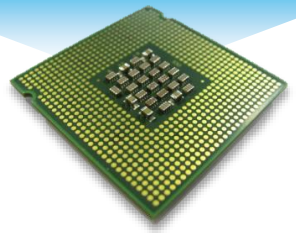

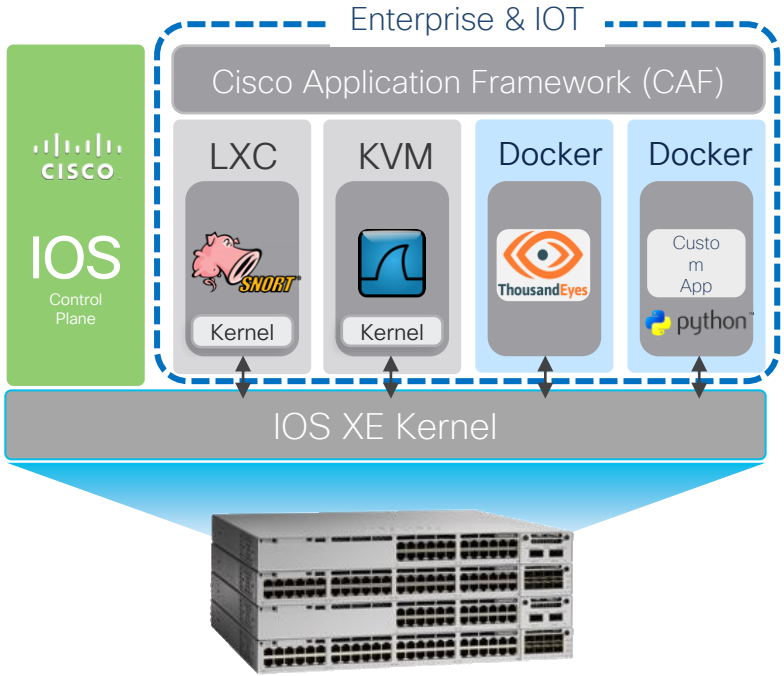

x86 CPU enables hosting NFV devices, Containers and 3rd-party Apps

cisco Live!

* Catalyst 9200 Series uses an embedded ARMv8 CPU

Cisco IOS XE – A Modern Operating System

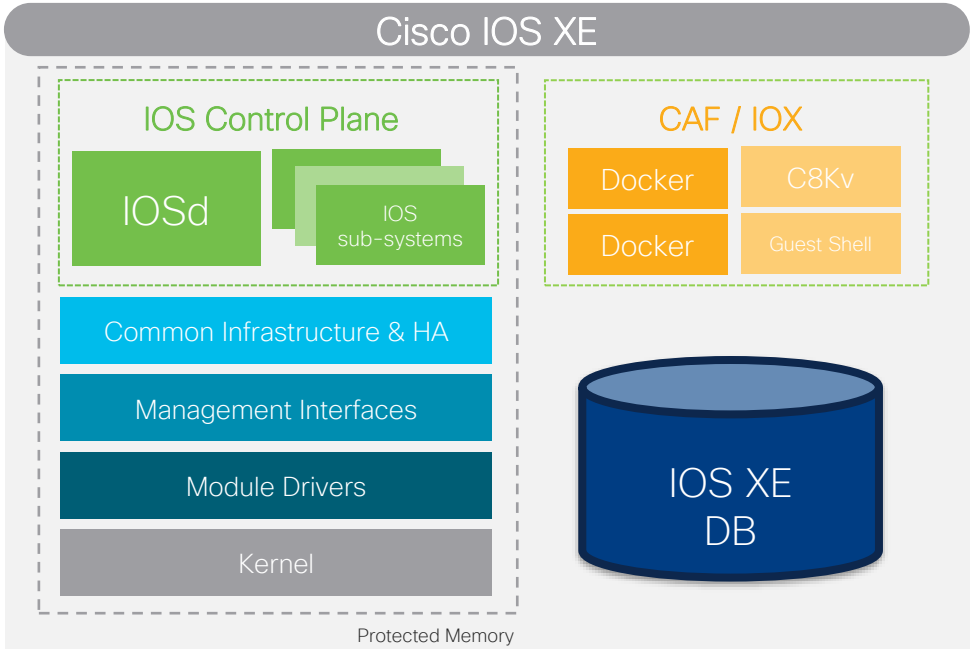

Cisco IOS subsystems
Resiliency and High Availability

Cisco IOS XE database
Programmability and Open models

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

Open, Model Driven & Secure Operating System

Cisco Catalyst 9000 Series IOS XE Software

History of IOS XE

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

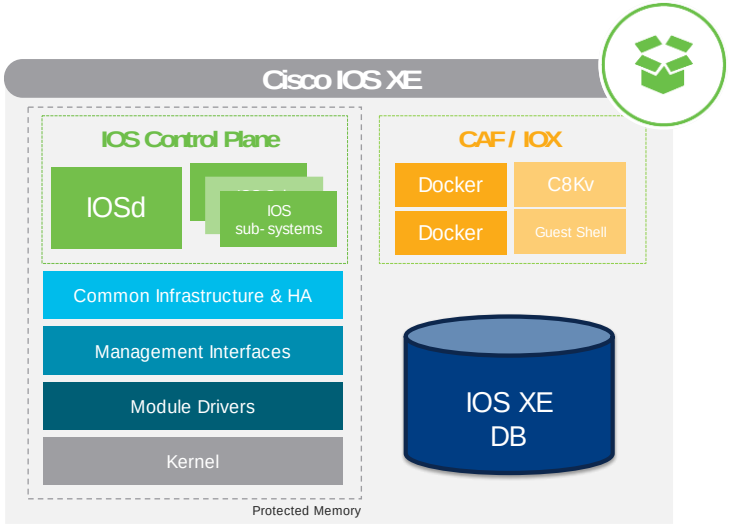

Brief History of Cisco IOS

NOTE: Timeline is just an approximation

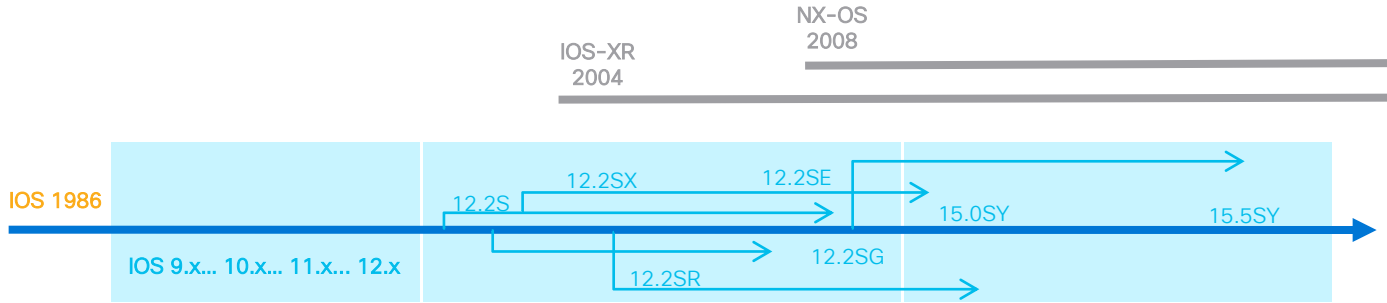

1984	1986	1993	2000	2007	2009	2010	2015	2017
Cisco was born	Operating System 6.0	IOS 9.x – 11.x	IOS 12.x S Release	IOS XE BinOS 3.x	IOS 15.x M&T Release	IOS XE NOVA 3.x	IOS XE Polaris 16.1.1	Open IOS XE Polaris 16.5.1
	Cisco Ships First AGS Router	LAN Switching, WAN Switching	Cat6500, 7600 Series	ASR1000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


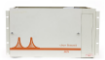






Brief History of IOS XE

Across Cisco Enterprise Platforms

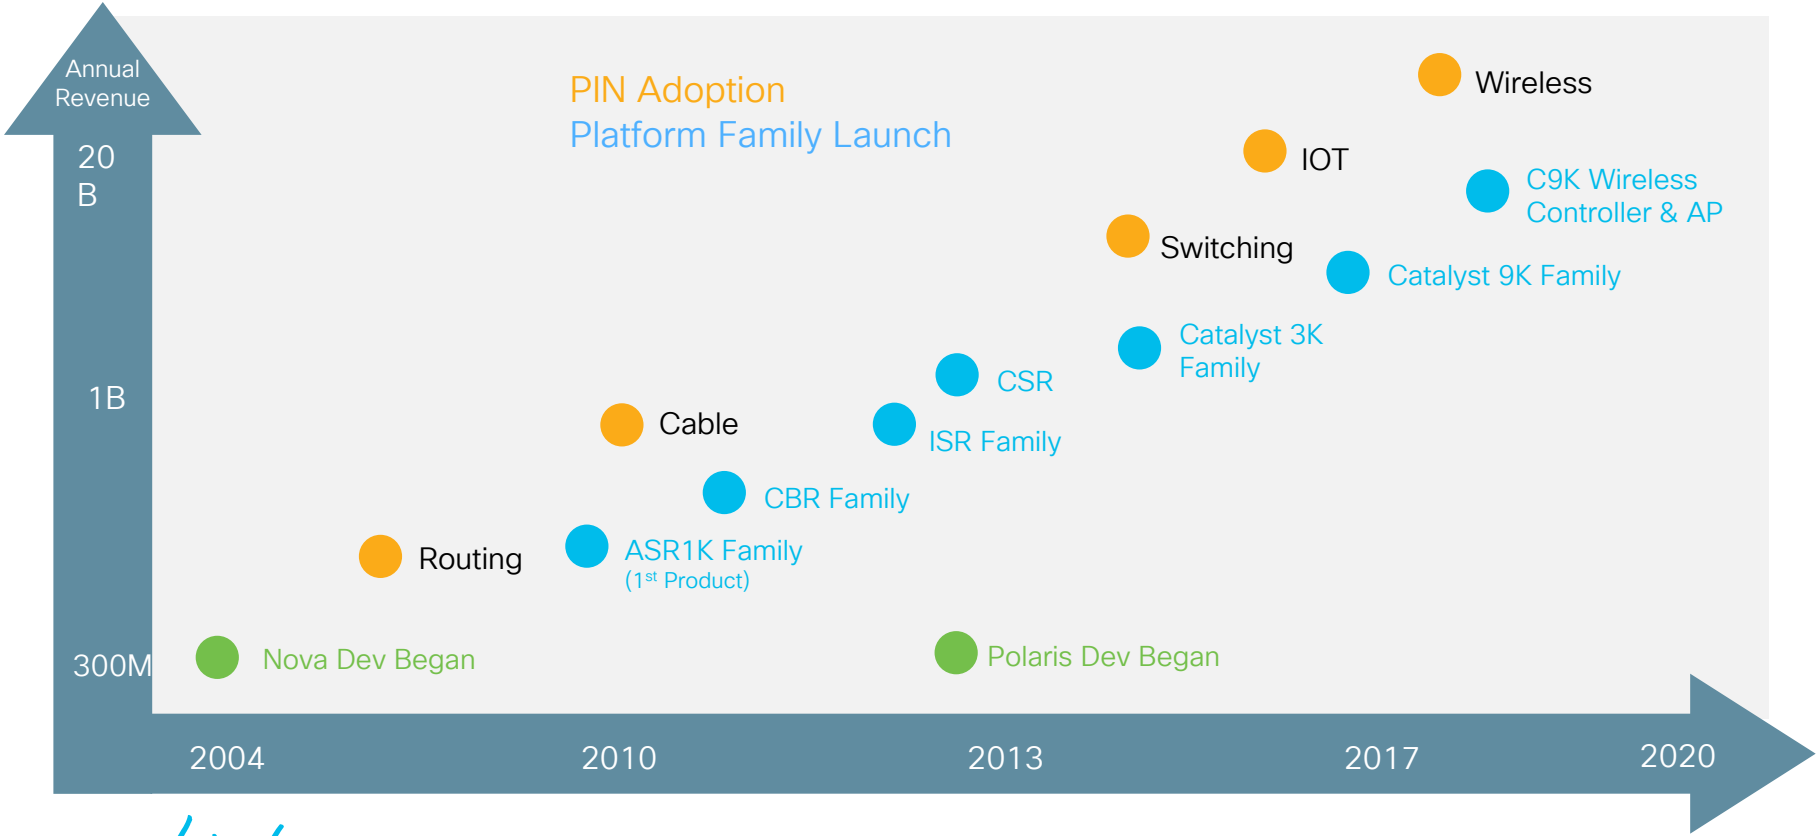

Cisco IOS XE – Architecture Evolution

Same look and feel – more powerful architecture

Cisco IOS	Cisco IOS XE 3.7.x(SE)	Cisco IOS XE 16.x
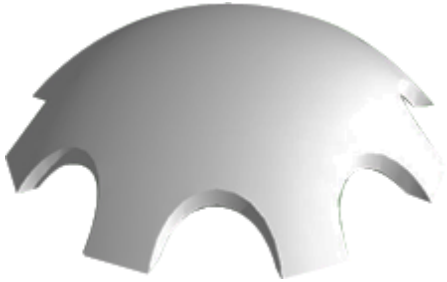		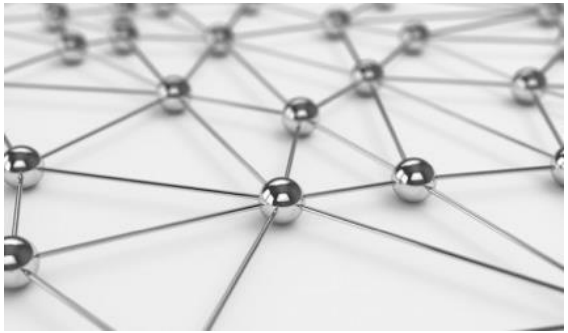
• **Monolithic IOS** • Compact, Streamlined • High performance	• **Monolithic IOSd**: Control-plane • Sub-packages for data plane • Linux daemons hosting capability • Message parsing capability	• **IOSd**: Component assemblies • **Modularized features**: Sub-packages • **Distributed Operating System** • **IOS XE (Crimson) Database** • Radioactive tracing and events

What is “Open” Cisco IOS XE?

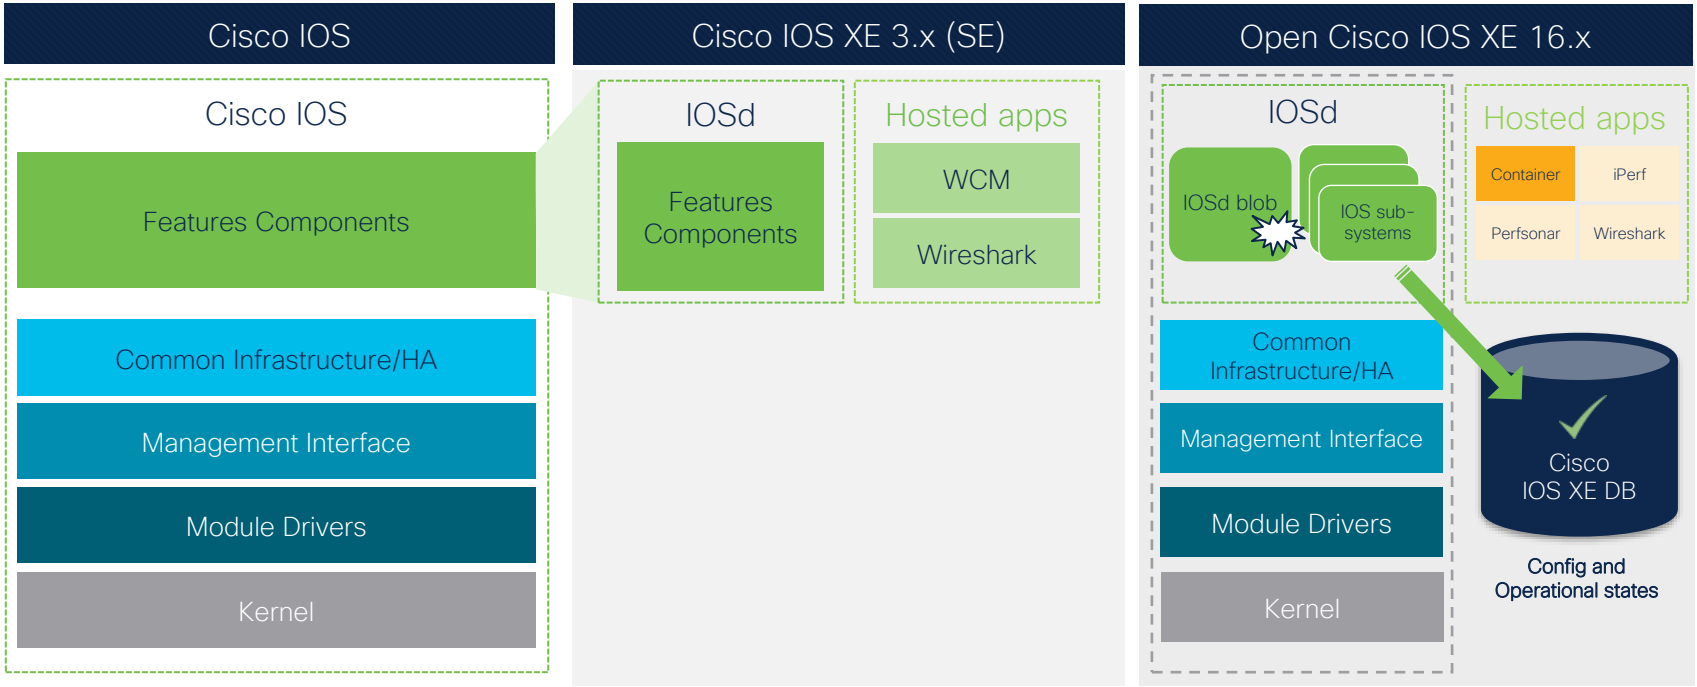

Modern Software Architecture - with the same look and feel

Open IOS XE – IOS Sub Systems

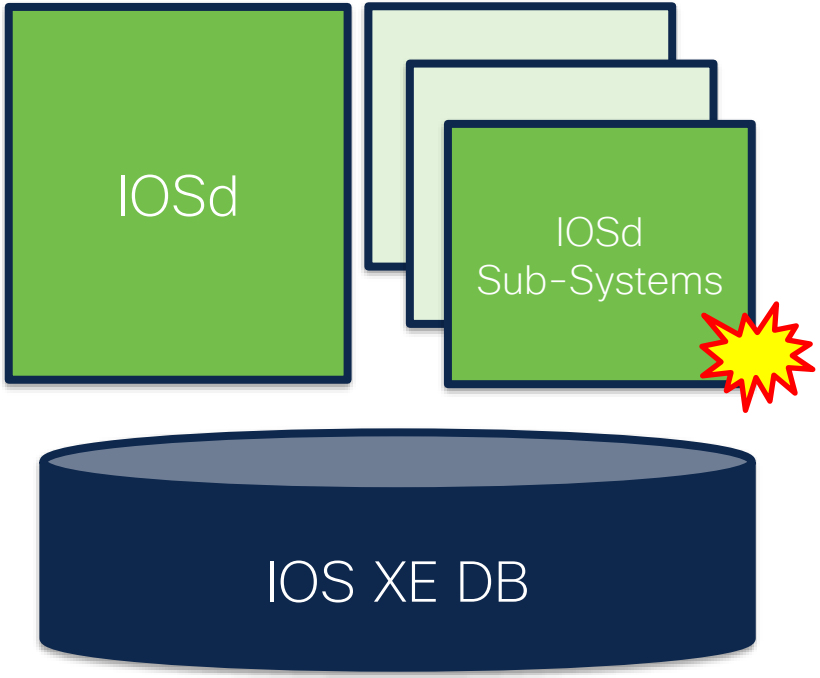

Failure of one IOS XE Sub-System(s) - keeps rest of the system intact

IOSd Sub-Systems enhance IOS Resiliency

Open IOS XE – Hardware DB

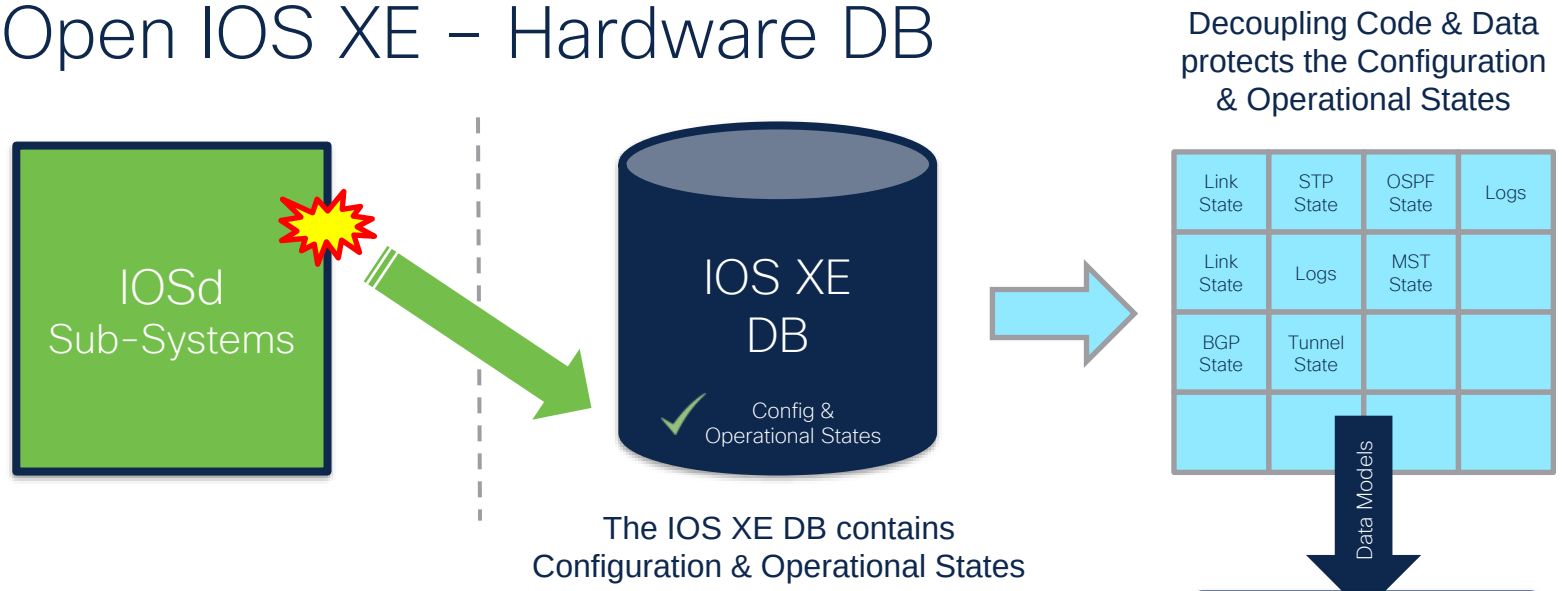

Decoupling Code & Data protects the Configuration & Operational States

The IOS XE DB contains Configuration & Operational States

- Higher Application UP Time
- Quicker Recovery
- Better Convergence

Basic Components

- **IOS XE Architecture**
  - Control Plane
  - Data Plane
  - System Plane
  - Management Plane
- **IOS XE on Catalyst 9K**

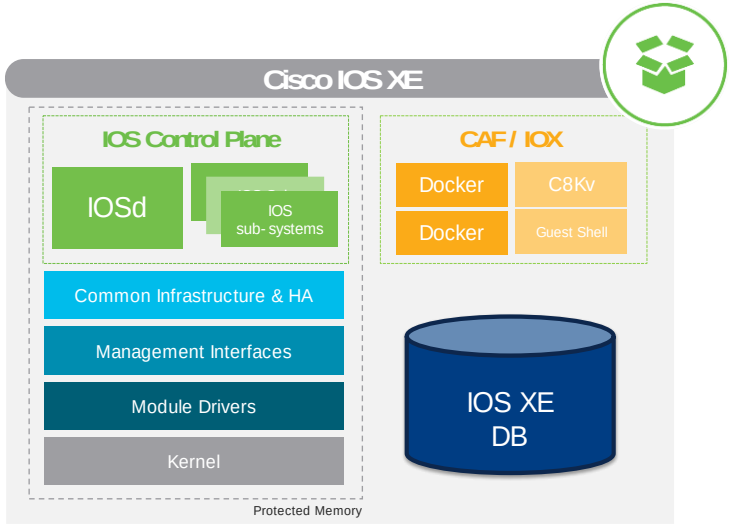

Cisco IOS XE Architecture

Modularized Components for Software Abstraction

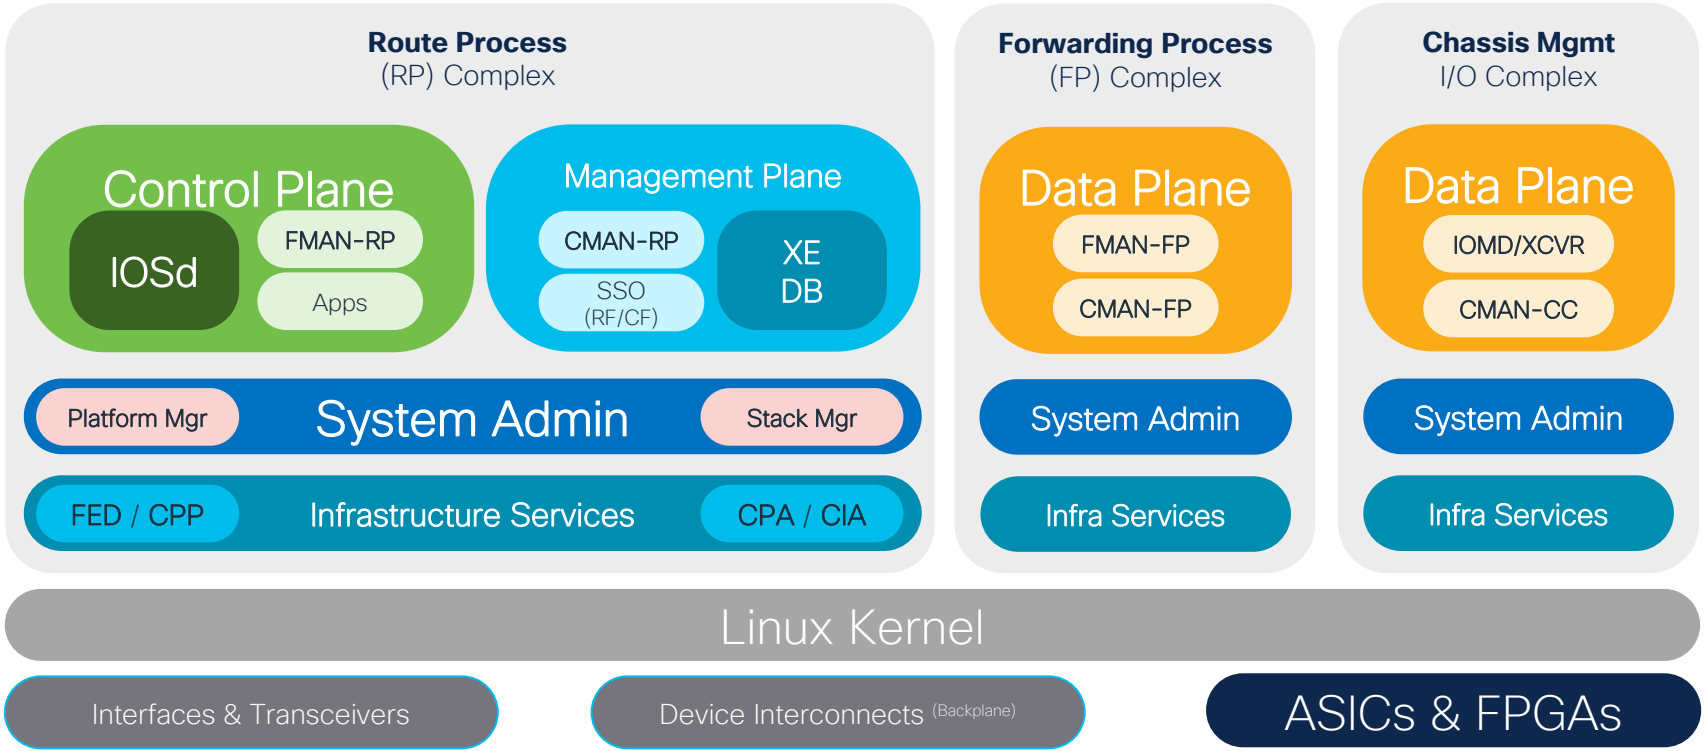

Cisco IOS XE Software

PI vs. PD Software Components

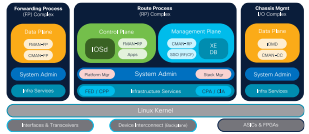

Platform Independent (PI)

- IOS – Internetwork Operating System
- FMAN – Forwarding Manager
- RP – Routing Process
- FP – Forwarding Process
- CGM – Classification Group Manager
- WCM – Wireless Controller Module
- IFM – Interface Manager
- PDS – Packet Distribution Service
- LSMPI – Linux Shared Memory

Platform Dependent (PD)

- CPA – Common Platform Abstraction
- 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
- FED – Forwarding Engine Driver

Cisco IOS XE Architecture

High Level Overview

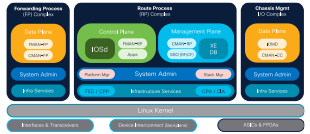

1. Control Plane
2. Data Plane
3. System Plane
4. Management Plane
5. Infrastructure Services
6. Linux Kernel & Platform Drivers

Cisco IOS XE – Control Plane

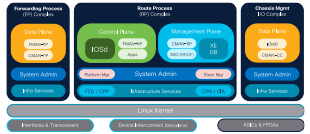

1 Control Plane

Data Plane

System Plane

Management Plane

Infrastructure Services

Linux Kernel & Platform Drivers

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

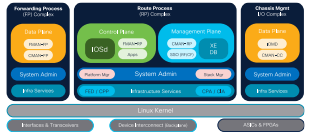

Control Plane

2

Data Plane

System Plane

Management Plane

Infrastructure Services

Linux Kernel & Platform Drivers

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 special hardware (e.g. UADP & QFP)

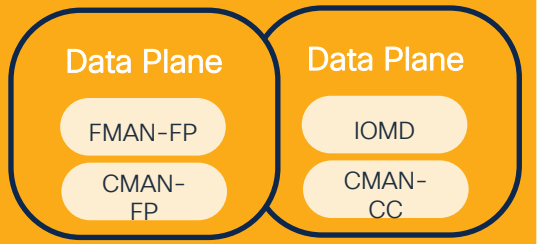

Cisco IOS XE - System Plane

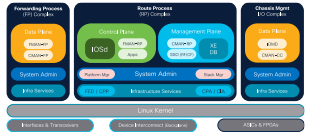

Control Plane

Data Plane

3

System Plane

Management Plane

Infrastructure Services

Linux Kernel & Platform Drivers

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

System Admin

Platform Mgr

Stack Mgr

Cisco IOS XE – Management & Infra

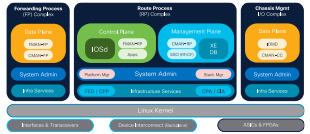

Control Plane

Data Plane

System Plane

4 Management Plane

5 Infrastructure Services

Linux Kernel & Platform Drivers

Device-Specific Services

- XE Database
- TDL DSL
- HW Drivers
- HA & ISSU
- Messaging
- Licensing
- Compiler
- Btrace
- Buildtime tools

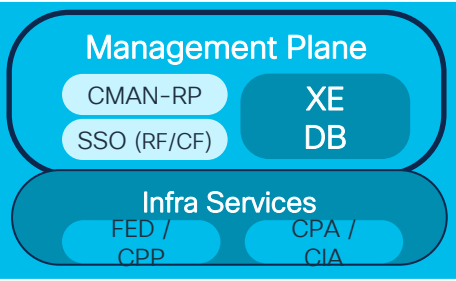

Cisco IOS XE on Catalyst 9000 Series

Hardware Forwarding Architecture

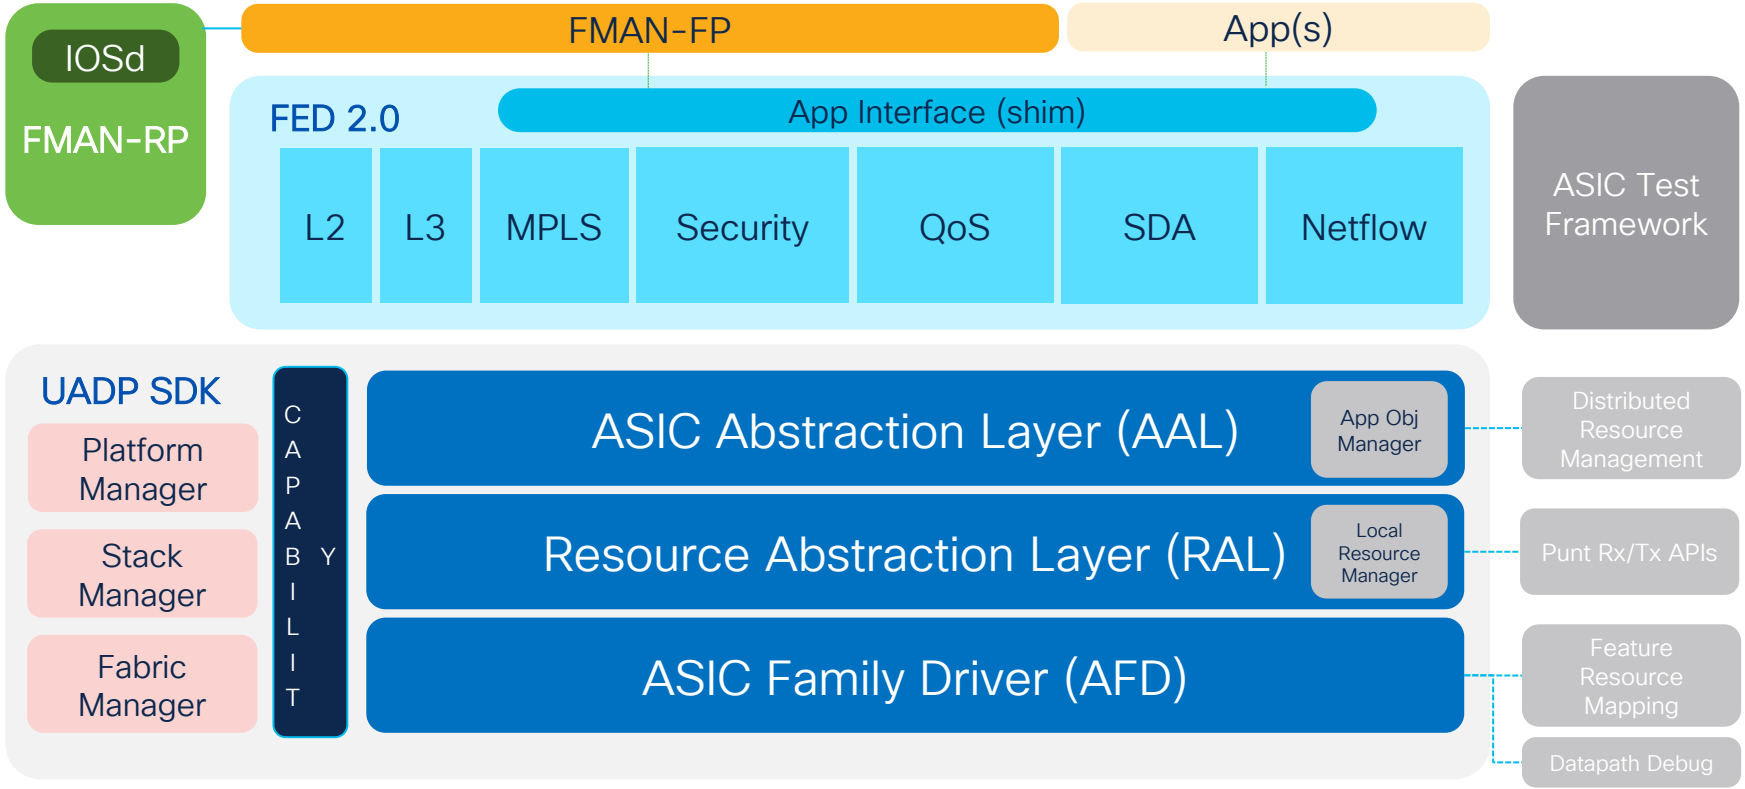

Cisco IOS XE – PD Abstraction Layer

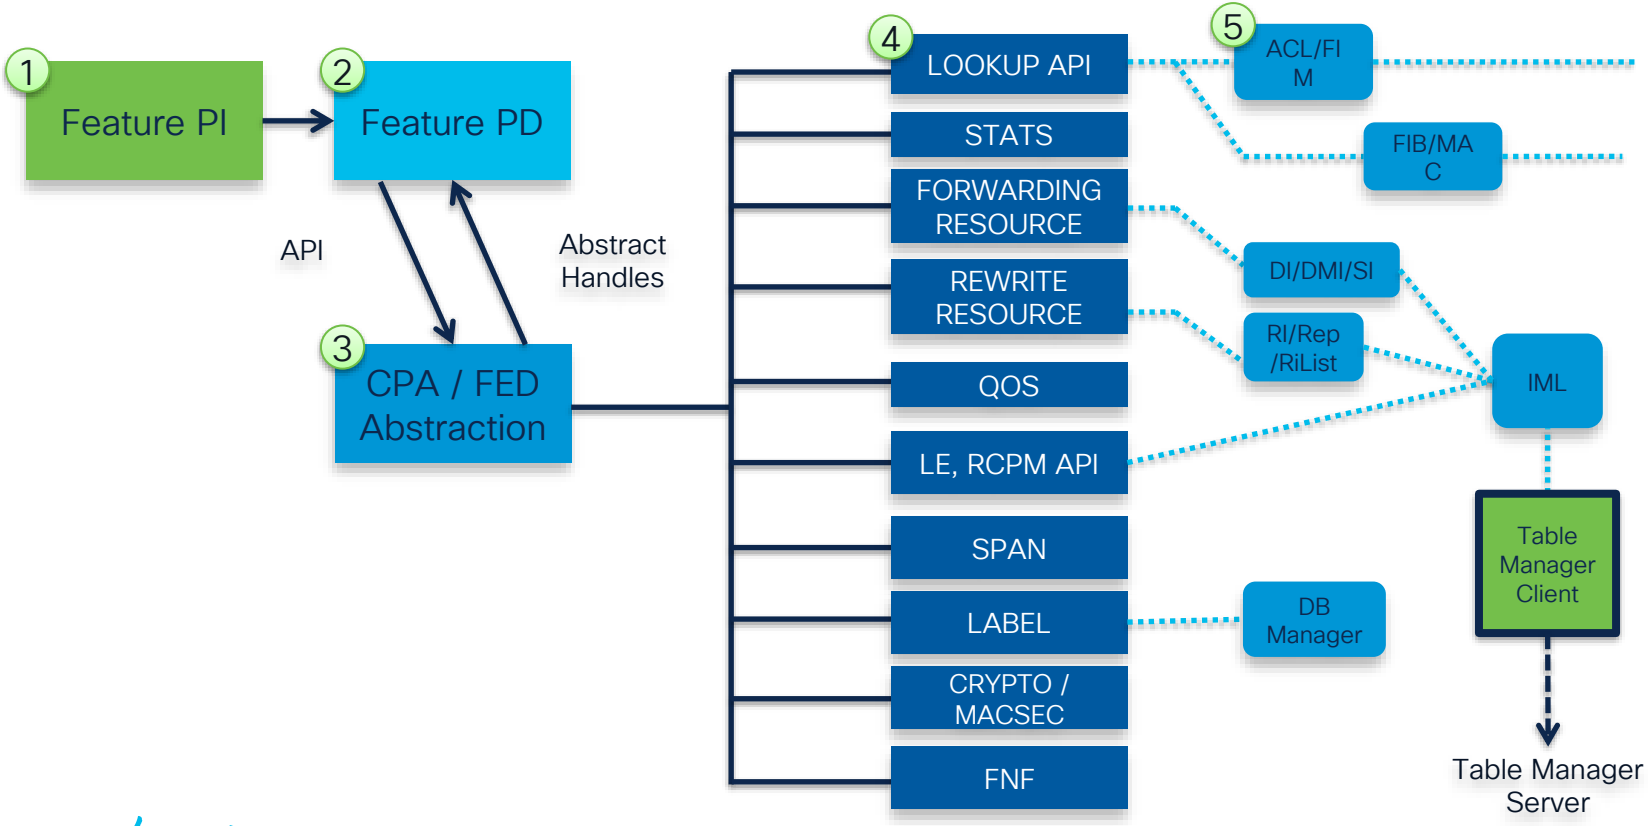

Cisco IOS XE “Lite” Architecture

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

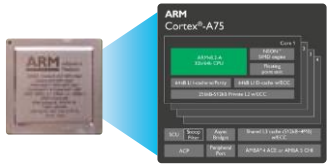

Route Process
(RP) Complex

Control Plane

IOSd

FMAN-RP

Apps

Management Plane

CMAN-RP

SSO
(RF/CF)

XE
DB

Stack Mgr

Infrastructure Services

Forwarding Process
(FP) Complex

Data Plane

FMAN-FP

FED

Infra Services

Chassis Mgmt
I/O Complex

Data Plane

IOMD-Lite

Infra Services

Linux Kernel

Interfaces & Transceivers

Device Interconnects (Backplane)

ASICs & FPGAs

Fixed vs. Modular vs. Stacking

Reusing Common Elements of OS Architecture

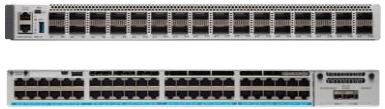

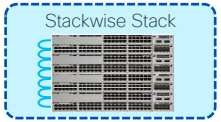

Fixed Platform

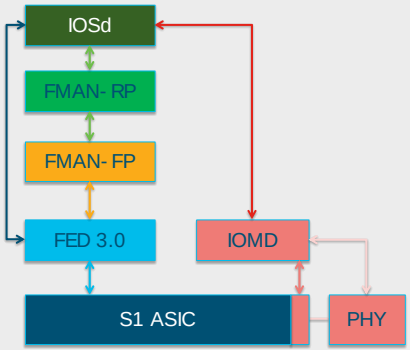

Modular Platform

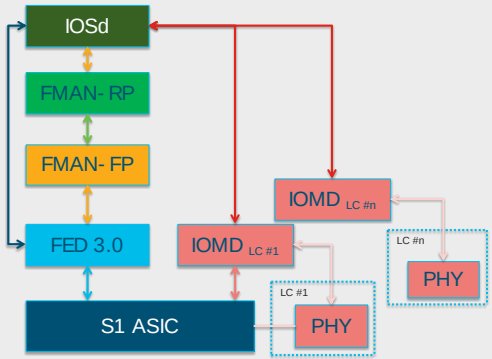

Stacked Platform(s)

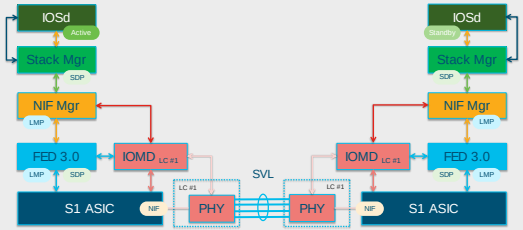

Cisco IOS XE

16.x - 17.x

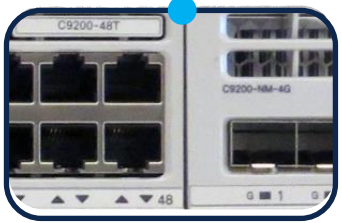

Catalyst 9200/L/CX

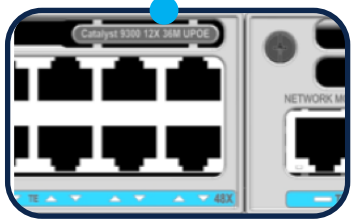

Catalyst 9300/L/X

Catalyst 9400/X

Catalyst 9500/X

Catalyst 9600/X

IOS XE Lite
Binary Image

IOS XE
Binary Image

Catalyst 9000 runs the same IOS XE Operating System

Cisco Catalyst 9000 Series IOS XE Software

Features of IOS XE

- **High Availability** (SSO & StackWise)
- **Install Mode** (SMU & ISSU/xFSU)
- **Model-Driven Telemetry**
- **Application Hosting**

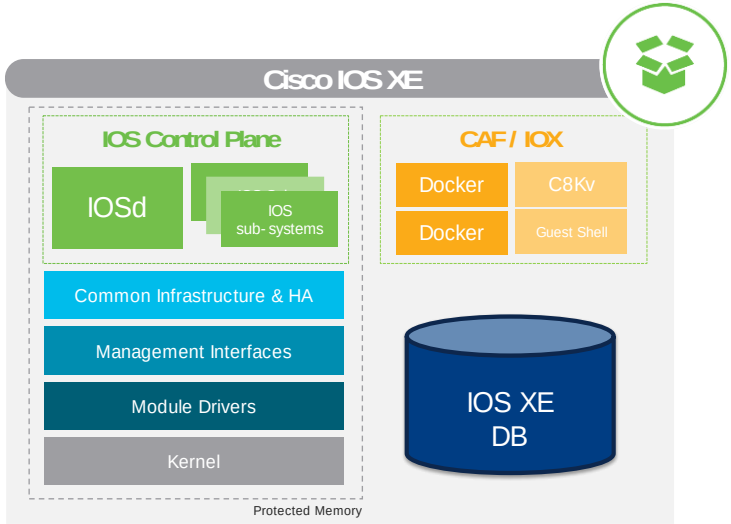

Mission-Critical Resiliency

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)

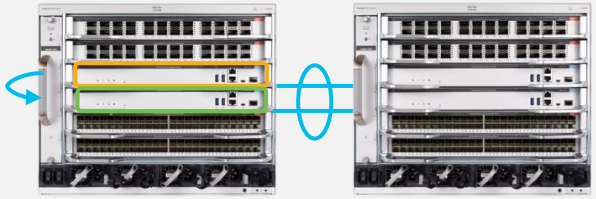

Catalyst 9400 Series

Catalyst 9500 Series

Architecture

StackWise® and StackWise Virtual

- Virtualized redundant systems for simplified configuration & protocols

Graceful Insertion/Removal (GIR)

- No downtime when device in maintenance mode

Operating System

Hot Patching (SMU)

- Minimal or no downtime for critical fixes

In-Service Software Upgrade (ISSU)

- Upgrade with minimal or no traffic loss

xFSU on C9300/L Stack New

- < 30 sec downtime - Stack upgrade

Platform

Redundant Supervisors

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

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

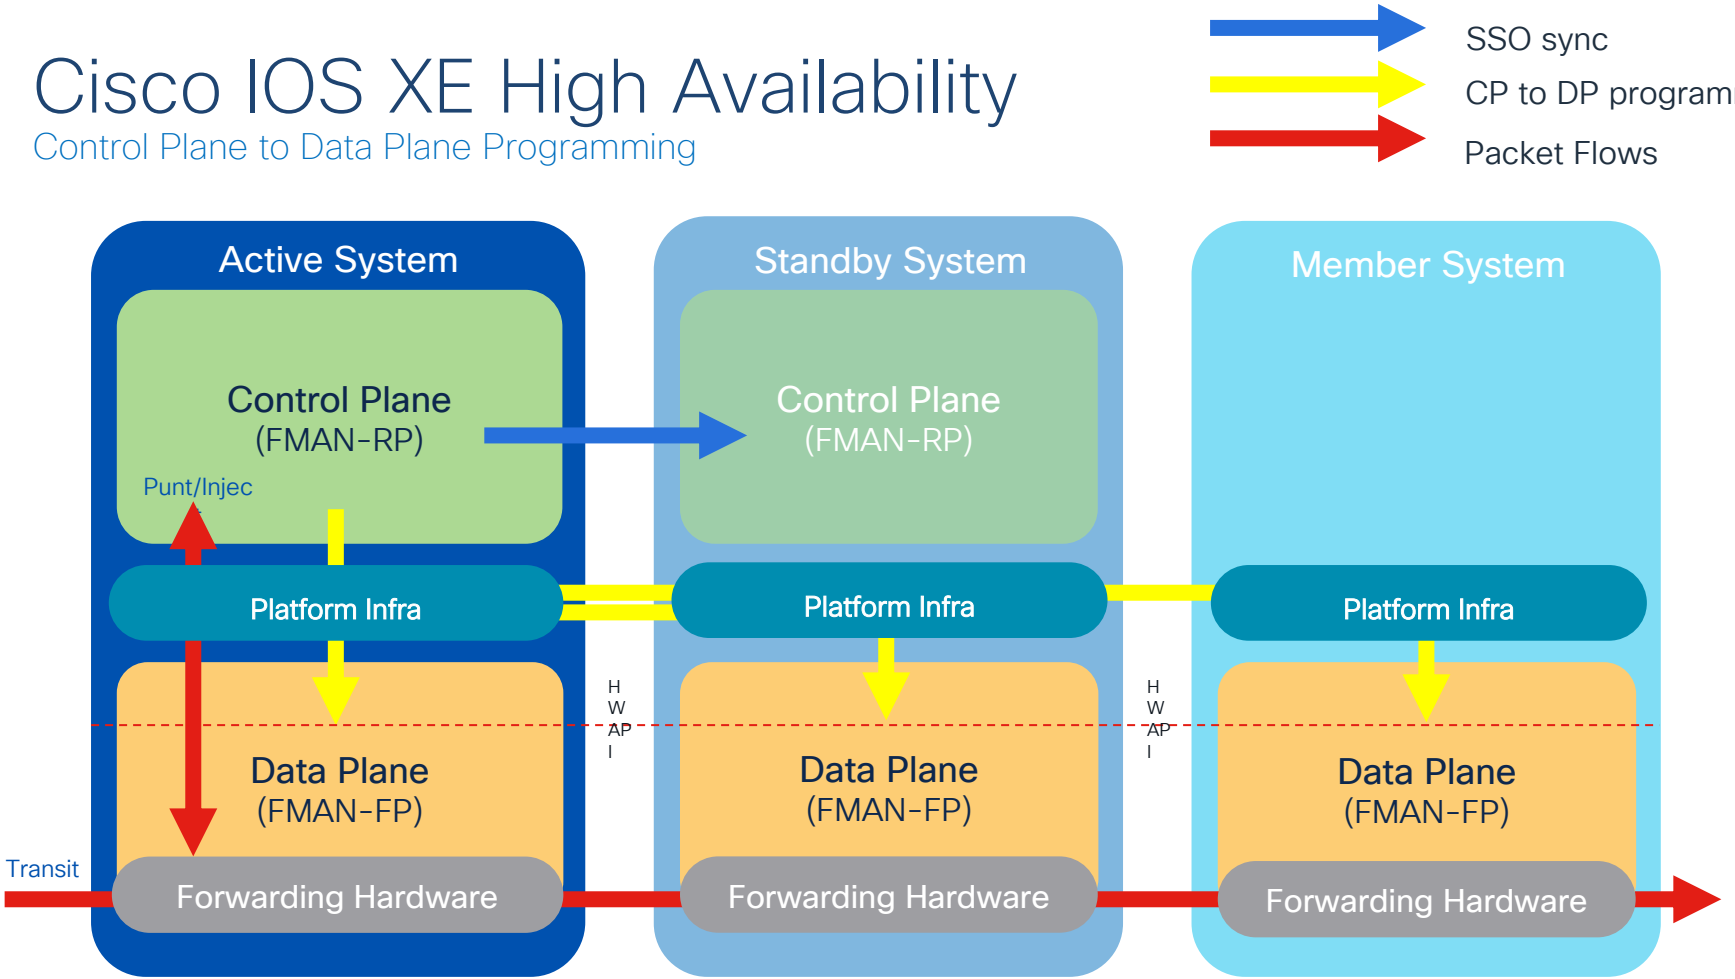

Cisco IOS XE - Install Mode

Single CLI set for Software Install, Patch & Upgrade

Workflow Steps:

- **Install Add** command will **perform the image download** from Cisco CCO Posted location
- **Install Activate** command **upgrade the control plane** with new software version
- **Install Commit** command **makes the changes permanent** (and deletes the older version)
- **Install Abort** you can issue an abort command to **revert the software** back to the original state

```
# install add <tftp://cisco.com/image.bin>
# install activate
```

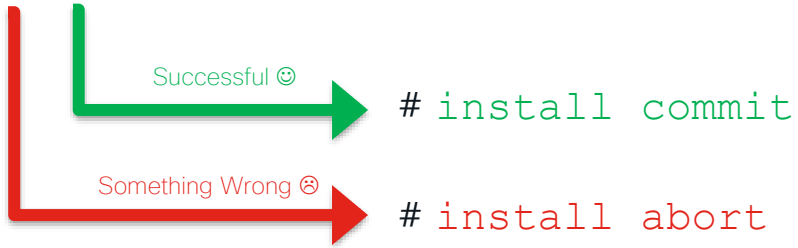

IOS XE Install – SMU patches

Ready for software patching

A **Software Maintenance Update (SMU)** is an emergency point fix positioned for expedited delivery to a customer in case of a network down or revenue-affecting scenario.

Cold Patching:
Install of an SMU requires a system reload in the first release. It is traffic impacting.

Hot Patching:
Install of an SMU does not require a reload. No traffic impact.

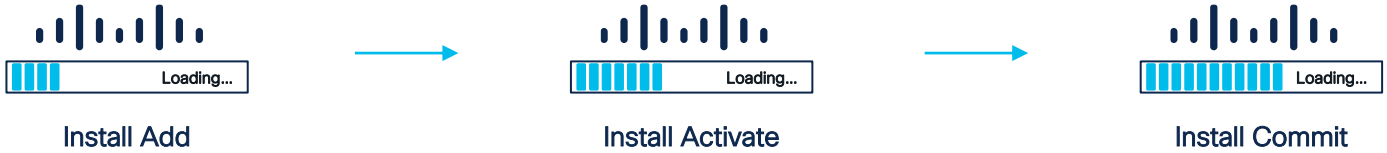

Install Add → Install Activate → Install Commit

In-Service Software Upgrade (ISSU)

Leverages SSO between IOS XE versions for seamless upgrade

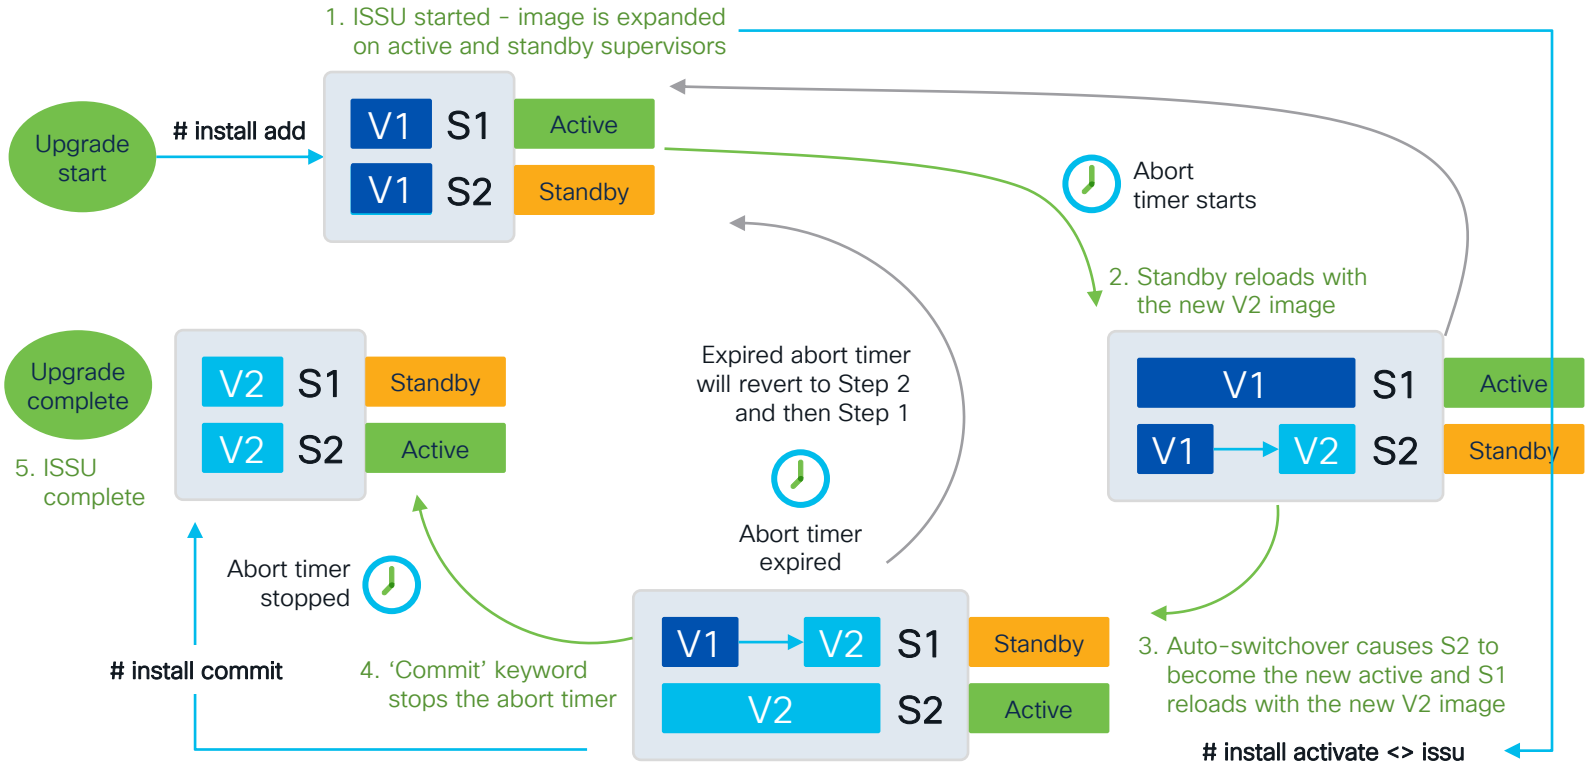

C9300/L- 17.3.2

C9300X- 17.7.1

Extended Fast Software Upgrade (xFSU)

Catalyst® 9300/9300L/9300X standalone

#install add file image activate reloadfast commit

Control plane

Data plane

Catalyst 9300/9300L/9300X stack

#install add file image activate reloadfast commit

Active Control plane

Data plane

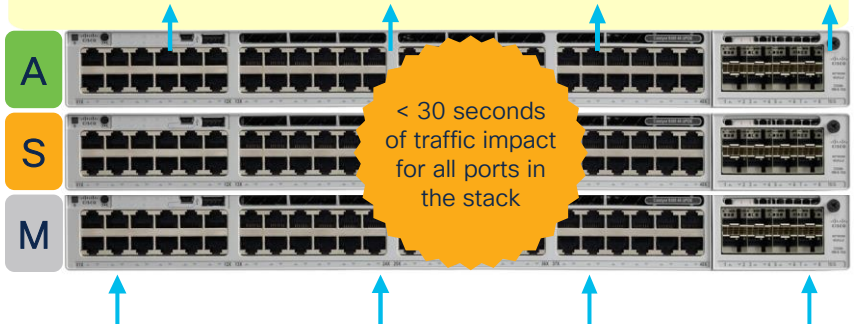

Cisco IOS XE Programmability

Telemetry "Stack"

DEVNET

developer.cisco.com/site/IOS XE

CLI

WebUI

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

YANG data models define the what's available for configuration and streaming telemetry

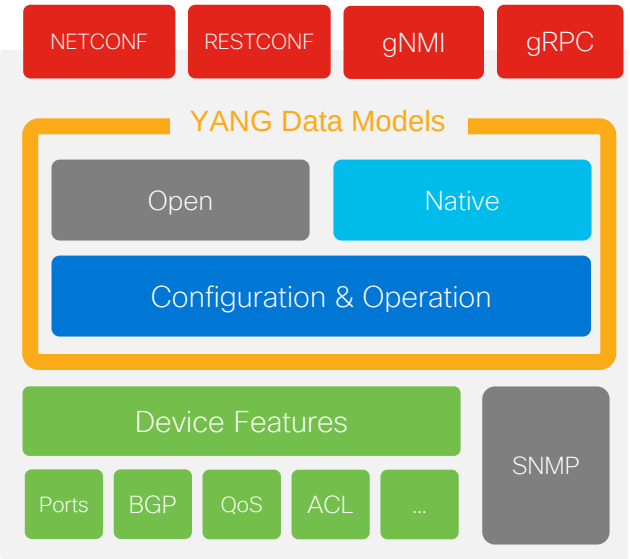

IETF

OpenConfig

IEEE

CISCO

Cisco IOS XE - Management

Management Plane – High-Level Overview

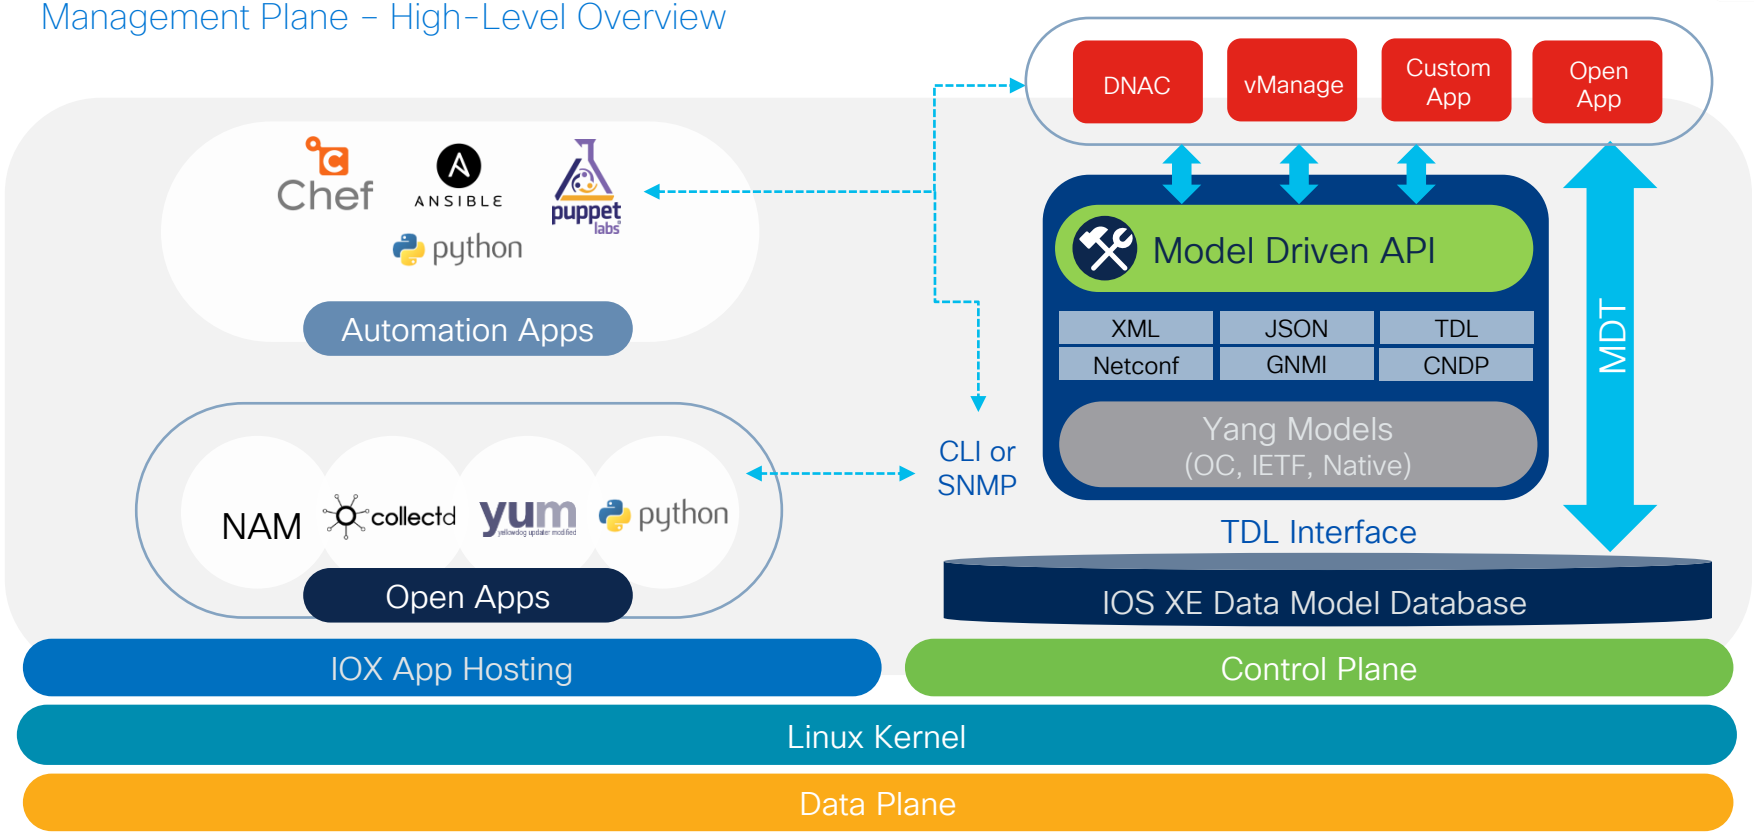

Cisco IOS XE - Model-Driven Telemetry

MDT Architecture

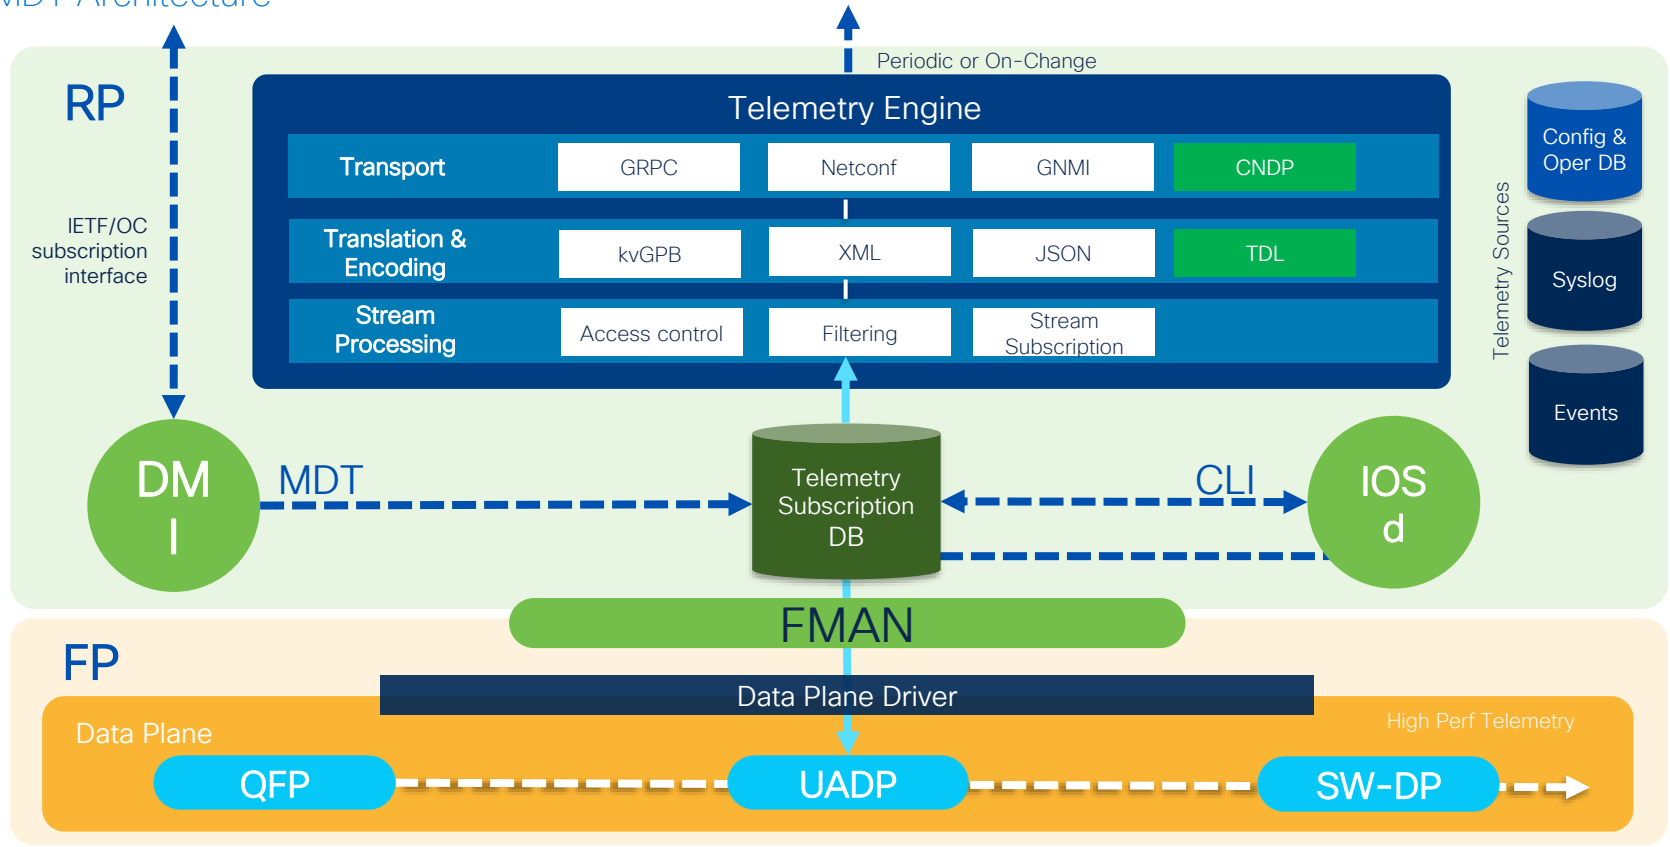

Cisco IOS XE

Application Hosting

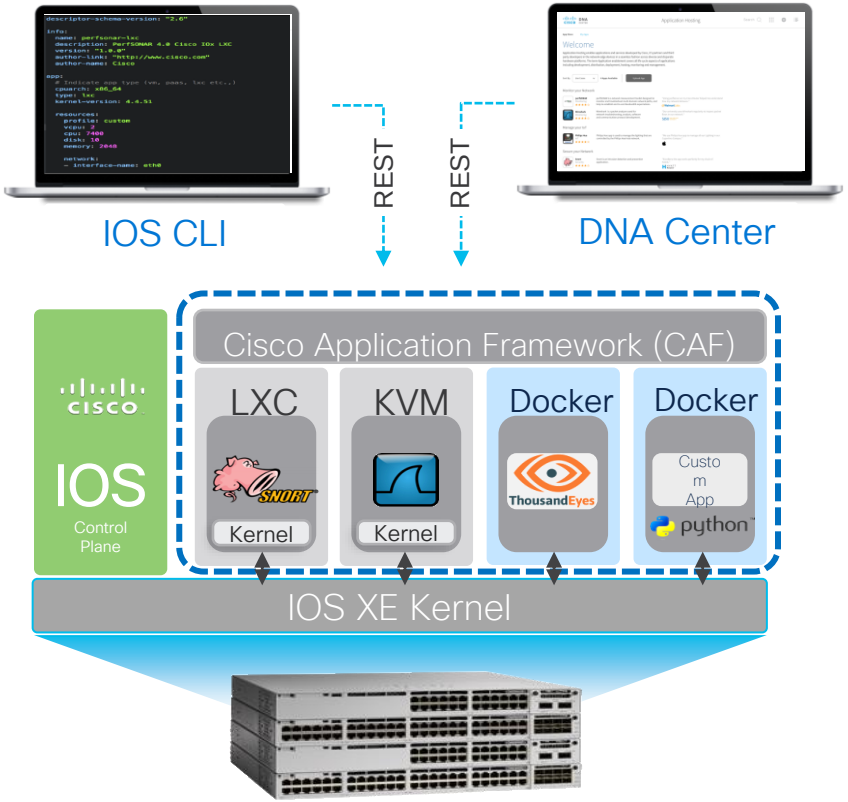

Catalyst 9000

Application Ecosystem

TRex Realistic Traffic Generator
Attivo Networks
Hawkeye
TRAFip
ThousandEyes
NetBeez
perfSONAR
Snort
AWS Greengrass
ISC DHCP Server
BIND 9
iPerf
Nagios
Alleantia
More...

- 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

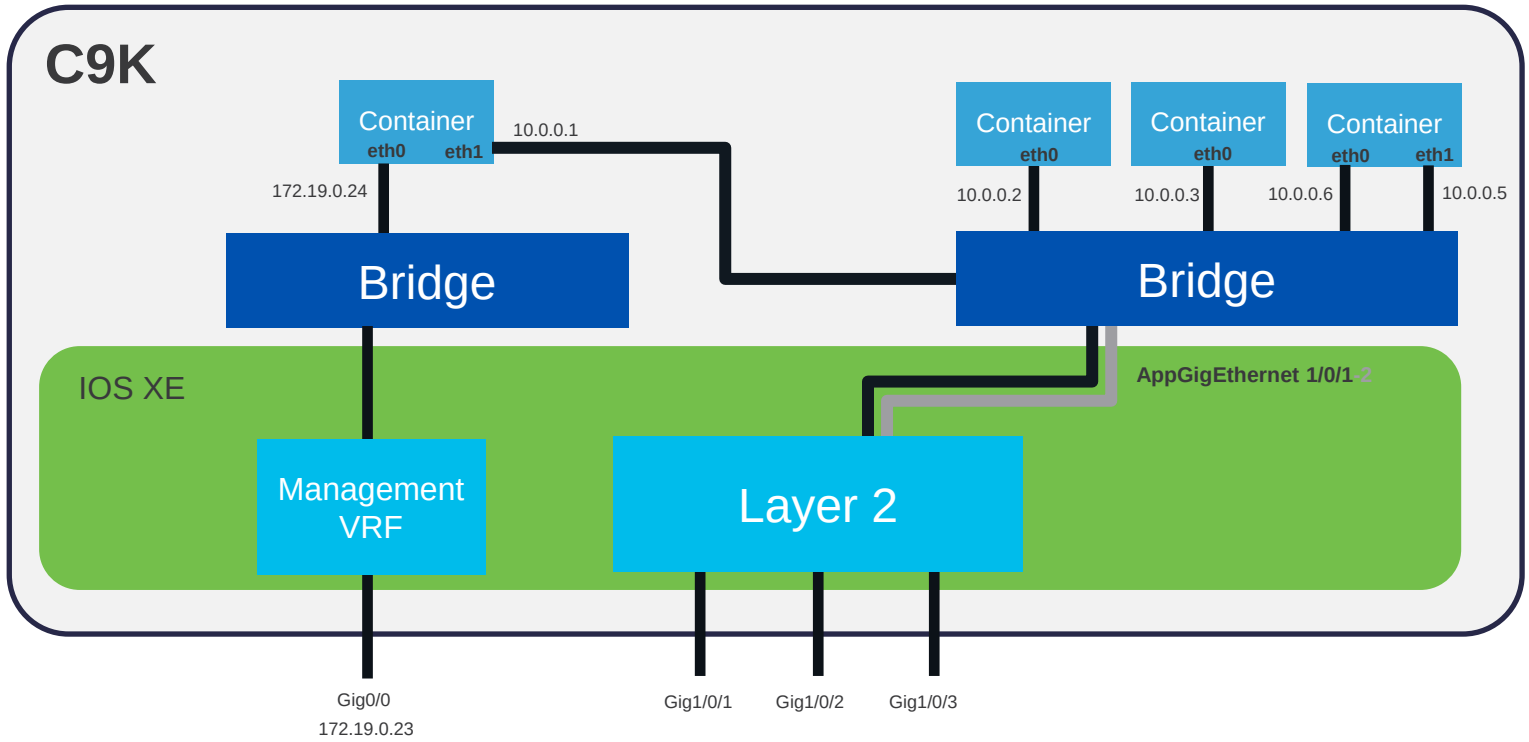

Media Access Control Security (MACsec 256)

Available on all Catalyst 9000 Series Switches

Enhanced security with AES-GCM 256-bit MACsec encrypted data plane

1. Extended packet numbering
2. Standard MKA key exchange (802.1ae)
3. "AES-256" keys Stronger security

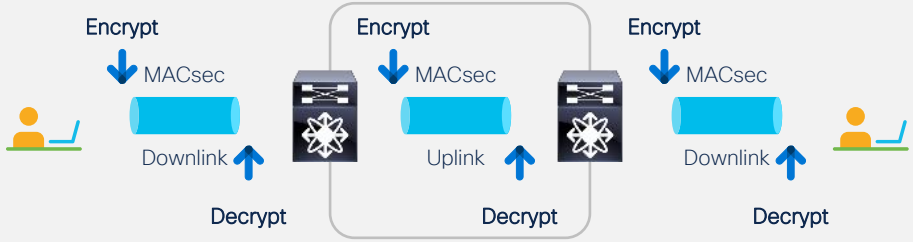

	Cisco Catalyst 9200 Series	Cisco Catalyst 9300 Series	Cisco Catalyst 9400 Series	Cisco Catalyst 9500 Series	Cisco Catalyst 9600 Series
Switch to switch	Supported (128-bit)	Supported (128-bit, 256-bit)	Supported (128-bit, 256-bit)	Supported (128-bit, 256-bit)	Supported (128-bit, 256-bit)
Switch to host	Supported (128-bit)	Supported (128-bit, 256-bit)	Supported (128-bit, 256-bit)	Supported (128-bit, 256-bit)	Supported (128-bit, 256-bit)

Benefits

Complete access security

Complete cross-platform alignment with uplink/downlink support

Protection against "Inside threats"

Securing campus infrastructure

Hop-by-hop Ethernet encryption

Line-rate performance on all ports

256-bit MACsec – Network Advantage

128-bit MACsec – Network Essentials

Cisco Catalyst 9000 Series IOS XE Software

IOS XE up to 17.6.x

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

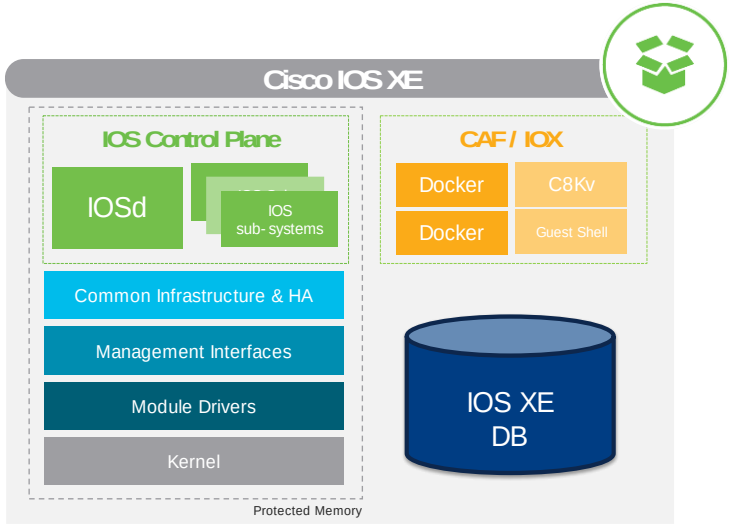

Cisco Catalyst 9000 Switching Portfolio

One Family from Access to Core – Common Hardware & Software

IOS XE
17.1.x - 17.6.x

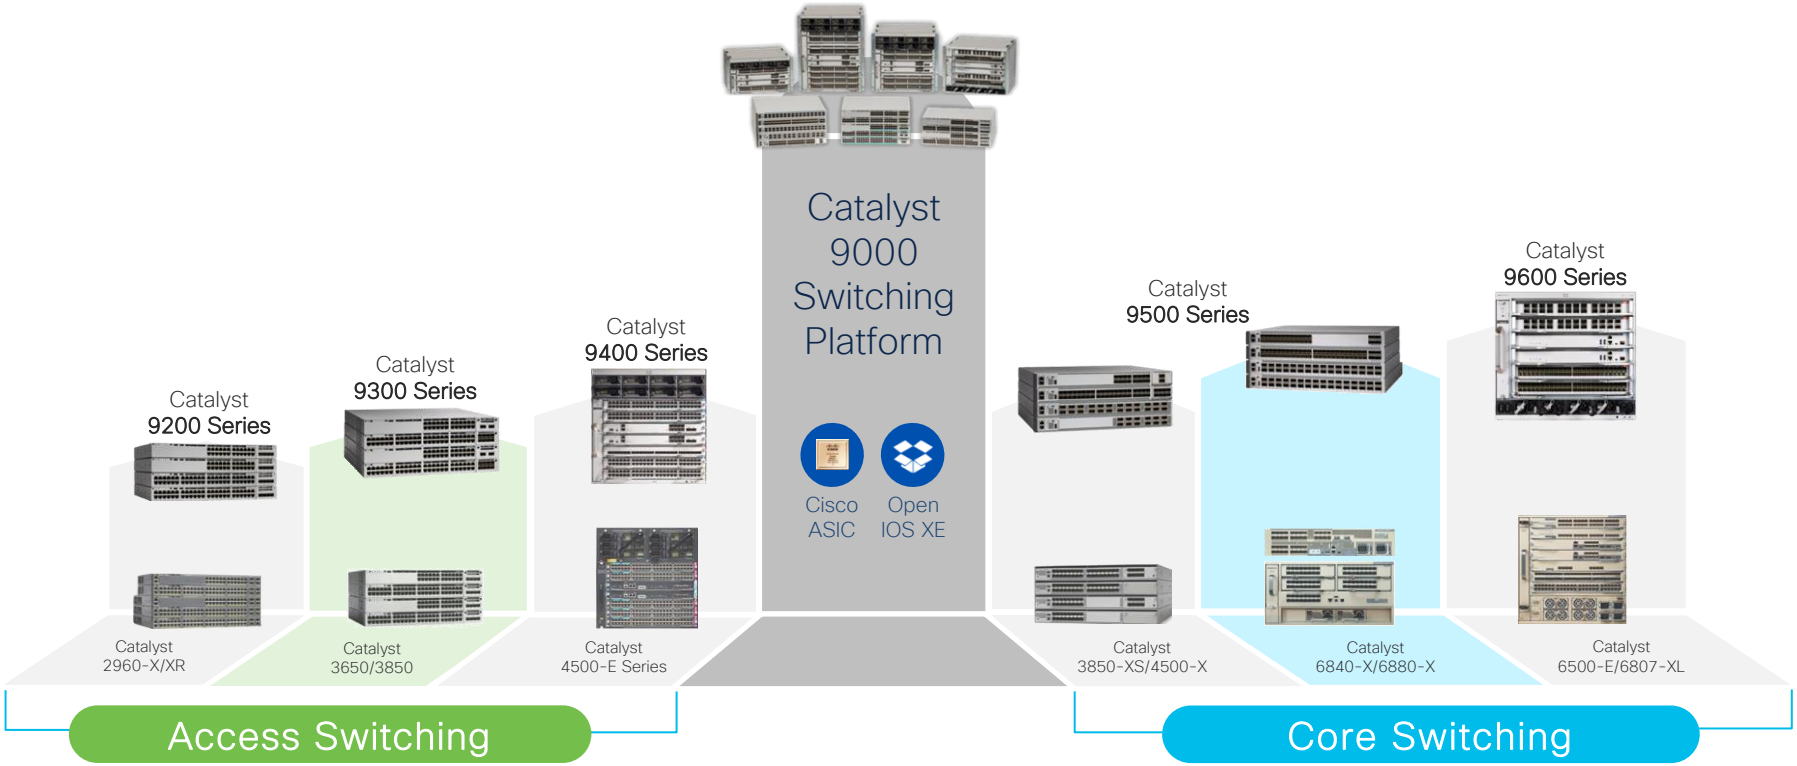

Cisco IOS XE – Release Schedule

3 Releases Annually (approx. every 4 months)

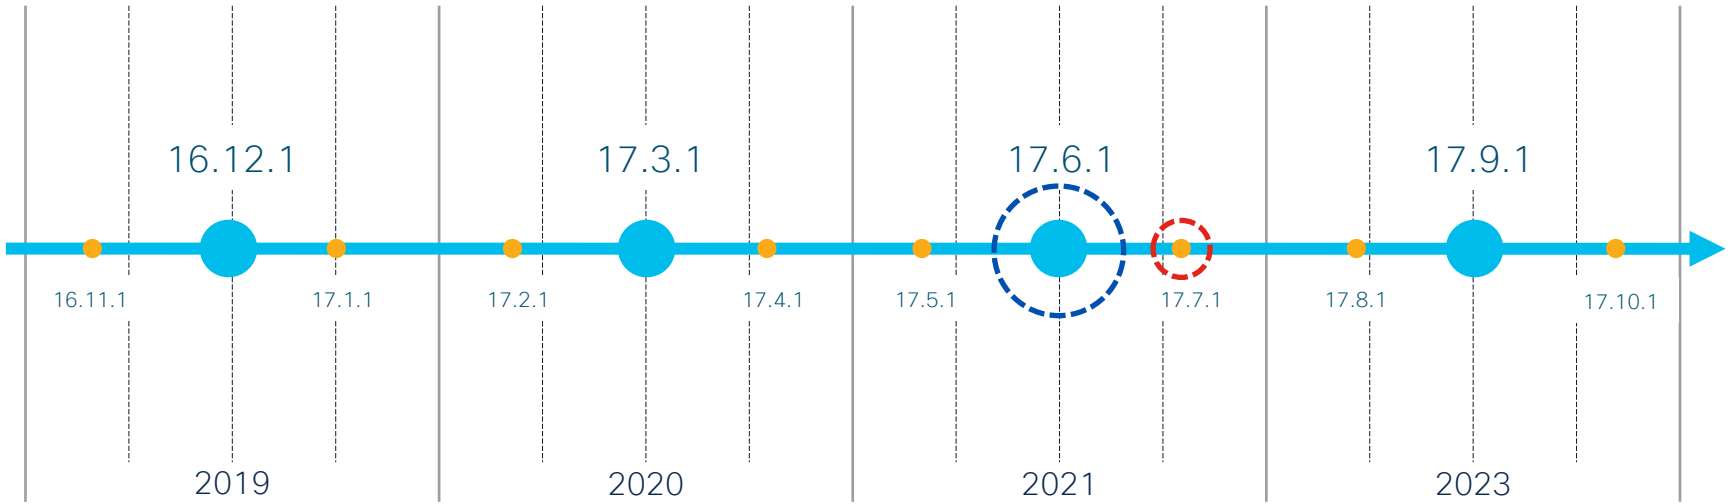

Extended Maintenance Release (EMR) – 36 months support

Recommended for wide-scale production deployments – Supports patches (SMU) and ISSU

Standard Maintenance Release (SMR) – 12 months support

Reference

Catalyst 9000 Switching – Key Features

* Limited Availability (LA) only

IOS XE 17.1.1 (Nov'19) SMR	IOS XE 17.2.1 (Mar'20) SMR	IOS XE 17.3.1 (July'20) EMR
Enhanced Security ❖ 9200/9300 – Umbrella Integration ❖ MACSEC over EoMPLS ❖ ERSPAN to v6 Destination	**Enhanced Security** ❖ TWS – Secure Swipe Clean DoD 5220.22-M standard	**Enhanced Security** ❖ Enhanced ACL Logging ❖ Wired Client Sensor* ❖ 9200/9300 – Umbrella Switch Connector with AD Integration
Overlays & Segmentation ❖ Inter-AS Option A (VRF-Lite) ❖ VPLS Flow Aware Transport (FAT) PseudoWire ❖ Extranet mVPN ❖ VXLAN aware Flexible Netflow ❖ EVPN to VRF-Lite handoff for Border Spine ❖ EVPN to MPLS handoff for Border Spine ❖ EVPN Tenant Routed Multicast (TRM)	**Overlays & Segmentation** ❖ EVPN – VxLAN ARP/ND flooding suppression. ❖ EVPN to MPLS hand off on Cat9K in Border spine role (single box) ❖ Hierarchical VPLS ❖ VPLS Multiple VCs per Spoke	**Overlays & Segmentation** ❖ mLDP: Multicast LDP* ❖ VPLS Routed PseudoWire (IRB): IPv6 Unicast ❖ MVPNv6 (Multicast 6VPE) ❖ MPLS VPN – Inter-AS Option AB ❖ BGP EVPN w VxLAN BUM rate-limiting support ❖ BGP-EVPN w VXLAN MAC/IP learning on Access ❖ Wide Area Bonjour with BGP-EVPN over VXLAN
Forwarding & Features ❖ 9600 – VRF aware PBR ❖ 9400 – NAT Profile	**Forwarding & Features** ❖ NAT – VRF aware NAT (VRF to Global)	**Forwarding & Features** ❖ 9500H/9600 : Customized SDM Template Ph1 (FIB) ❖ IP-FRRv4: LFA EIGRP and OSPFv2 per prefix ❖ Non-Stop Routing (NSR): L3 Forwarding Redundancy ❖ LACP 1:1 redundancy and dampening
High Availability ❖ 9600 – Quad Sup SVL Support (RPR) ❖ 9300 – xFSU Standalone	**High Availability** ❖ 9300 – xFSU reload with backside stacking* ❖ 9300 – xFSU support with dot1x, MAB, Webauth* ❖ 9300 – xFSU : LACP Protocol support*	**High Availability** ❖ Flexlink+ with VLAN Load Balancing ❖ 9300 – xFSU Reload: Stacked and Standalone (17.3.2)
Platform & Programmability ❖ Per port MTU support ❖ 9500H/9600 – Unified Port Buffer ❖ 9400 – Native Docker for App Hosting	**Platform & Programmability** ❖ 9400 – 9216 bytes MTU ❖ 9600 – Breakout Support ❖ gPTP/ PTPv2 support on Port Channels	**Platform & Programmability** ❖ PVLAN on Trunks and Port Channels ❖ PTPv2 and gPTP support on 9400* ❖ ETA and AVC Interoperability on same port ❖ SHA-512 secure image-bootup integrity check ❖ gRPC Model Driven Telemetry (MDT) with TLS
Hardware & Optics ❖ **C9300L** mGig SKUs ❖ **C9600-LC-48TX** mGig Linecard	**Hardware & Optics** ❖ **9500/9600** – AOC/DAC, QSFP-4SFP10G	**Hardware & Optics** ❖

Reference

Catalyst 9000 Switching – Key Features

* Limited Availability (LA) only

IOS XE 17.4.1 (Nov'20) SMR	IOS XE 17.5.1 (Apr'21) SMR	IOS XE 17.6.1 (Aug'21) EMR
Enhanced Security ❖ FQDN ACLs ❖ RADSEC – Radius over TLS and DTLS ❖ Stealthwatch Cloud Integration* ❖ Wired Client Sensor with Flash*	**Enhanced Security** ❖ FQDN Redirect ACL ❖ Wired Dynamic VLAN ❖ Secure Network Analytics Connector ❖ DSCP Marking for RADIUS Packets ❖ Session timers AV Pair ❖ Interface Templates ❖ Trustworthy Systems	**Enhanced Security** ❖ 9300X – IPsec Phase1 – SVTI, IKEv2 ❖ IPv6 FQDN Redirect ACL ❖ RADSEC CoA Enhancement
Overlays & Segmentation ❖ PVLAN with BGP EVPN over VxLAN	**Overlays & Segmentation** ❖ Selective Q-in-Q ❖ BGP EVPN L2/L3 VNI scale	**Overlays & Segmentation** ❖ MPLS Traffic Engineering (TE) – Phase1 ❖ LACP/PAGP over EoMPLS ❖ MLD snooping over VPLS
Forwarding & Features ❖ 9500H/9600 – Customized SDM Template Ph2 (ACL)	**Forwarding & Features** ❖ 9500H/9600 – Customized SDM Template Ph3 (4K VLAN) ❖ Enhanced NAT scale ❖ BGP Monitoring Protocol ❖ WCCP Over GRE	**Platform Features** ❖ VRF Aware WCCP ❖ Enhanced NAT Session Monitoring ❖ NAT Precedence ❖ Bonjour mDNS SSO, FHRP Service Peer Support
High Availability	**High Availability**	**High Availability** ❖ VRRPv3 SSO
Platform & Programmability ❖ YANG model updates ❖ Smart Licensing using Policy	**Platform & Programmability** ❖ gPTP over L3 Unicast ❖ Disable USB SSD ❖ App Hosting Updates	**Platform & Programmability** ❖ App Hosting on 9300X ❖ Thousand Eyes – 4.0 Version Agent ❖ Perpetual PoE/UPOE with StackPower ❖ PTP on StackWise*, PTP over SDA ❖ Programmability & Automation updates
Hardware & Optics ❖ **9500H/9600** – SFP-10G-TX, QSFP-40/100-SR4, ❖ **9600** – 4 x 25G Breakout, GLC-GE-100FX and GLC-TE-100M	**Hardware & Optics** ❖ **C9400-LC-48HN** – 5G MGIG line card with 90W PoE	**Hardware & Optics** ❖ **C9300X-12Y / 24Y** – 10/25G Fiber Switch with Cisco UADP2.0sec ❖ **C9300X-NM** – 2x 40/100G, 8x 10/25G, 8x mGiG uplinks

Cisco Catalyst 9000 Series IOS XE Software

IOS XE after 17.7.x

- **IOS XE Continues to Evolve**
- **New C9K Platforms**
- **Common Platform Abstraction**
- **Forwarding Engine Driver**

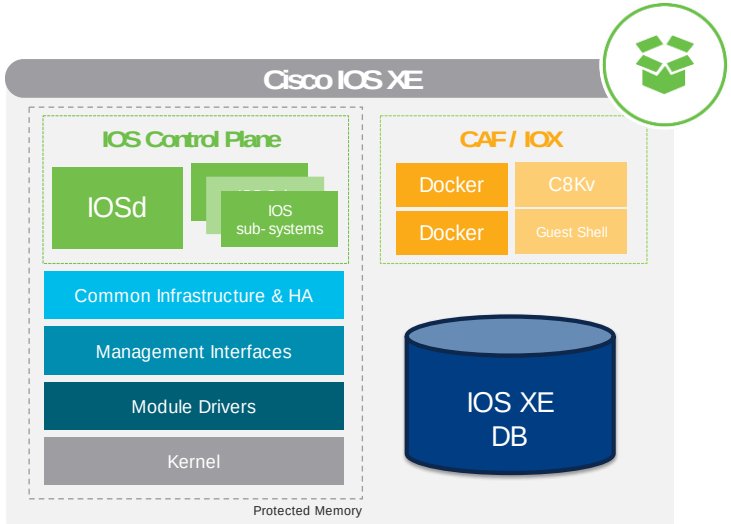

Cisco Catalyst 9000 Switching Portfolio

One Family from Access to Core – Common Hardware & Software

IOS XE 17.7.x - 17.9.x

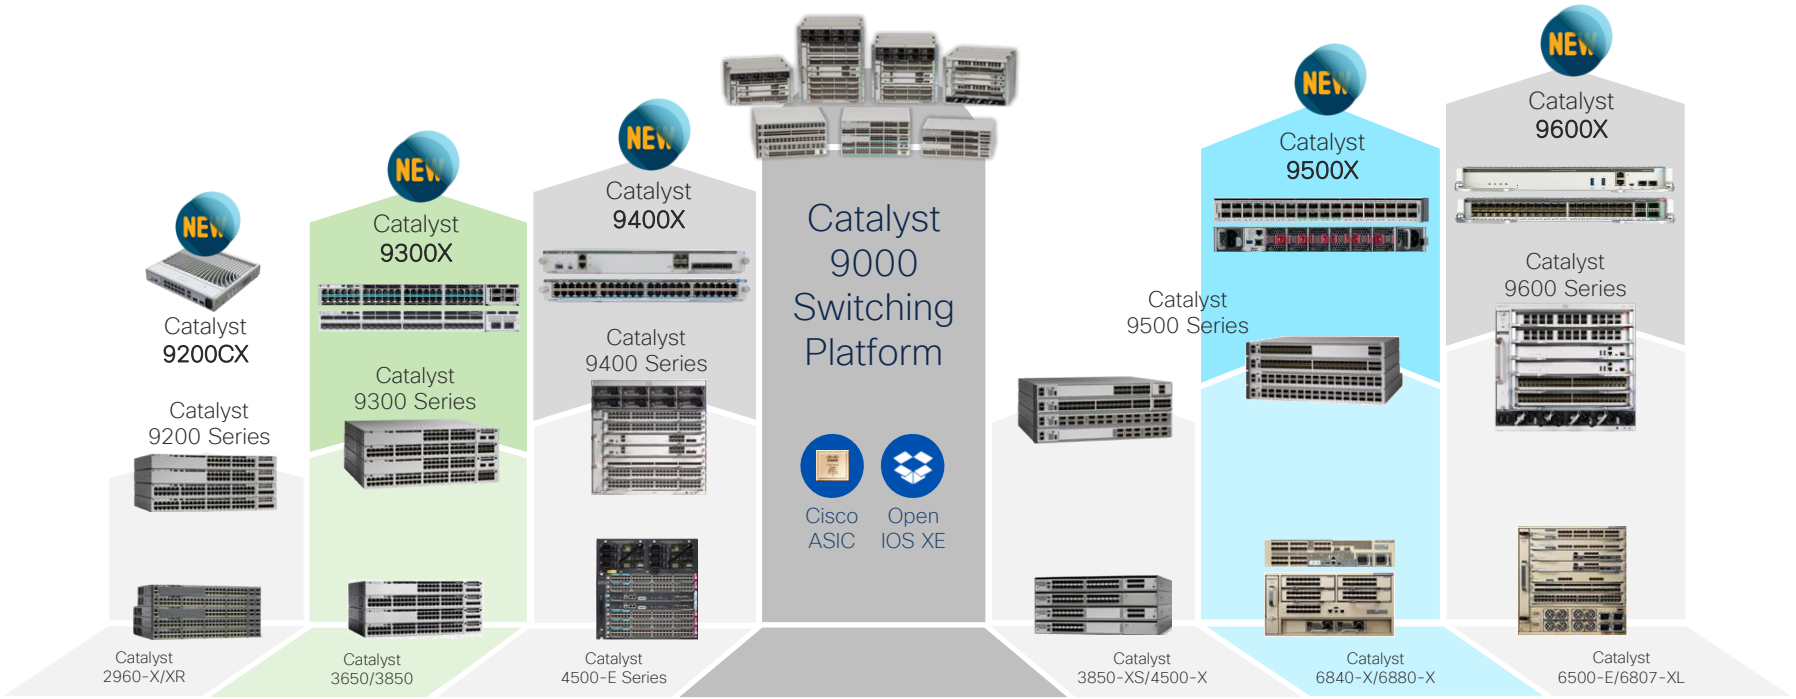

Cisco IOS XE – Release Schedule

3 Releases Annually (approx. every 4 months)

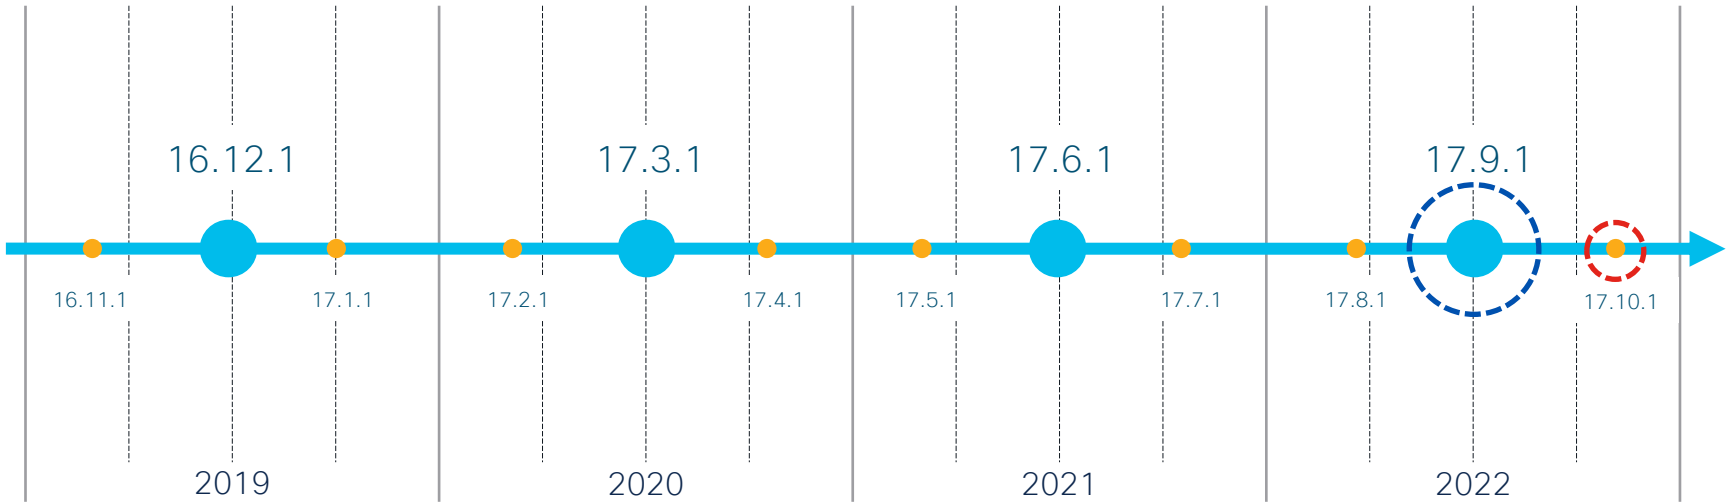

Extended Maintenance Release (EMR) – 36 months support

Recommended for wide-scale production deployments – Supports patches (SMU) and ISSU

Standard Maintenance Release (SMR) – 12 months support

Cisco IOS XE - Release Schedule

Graphical Overview

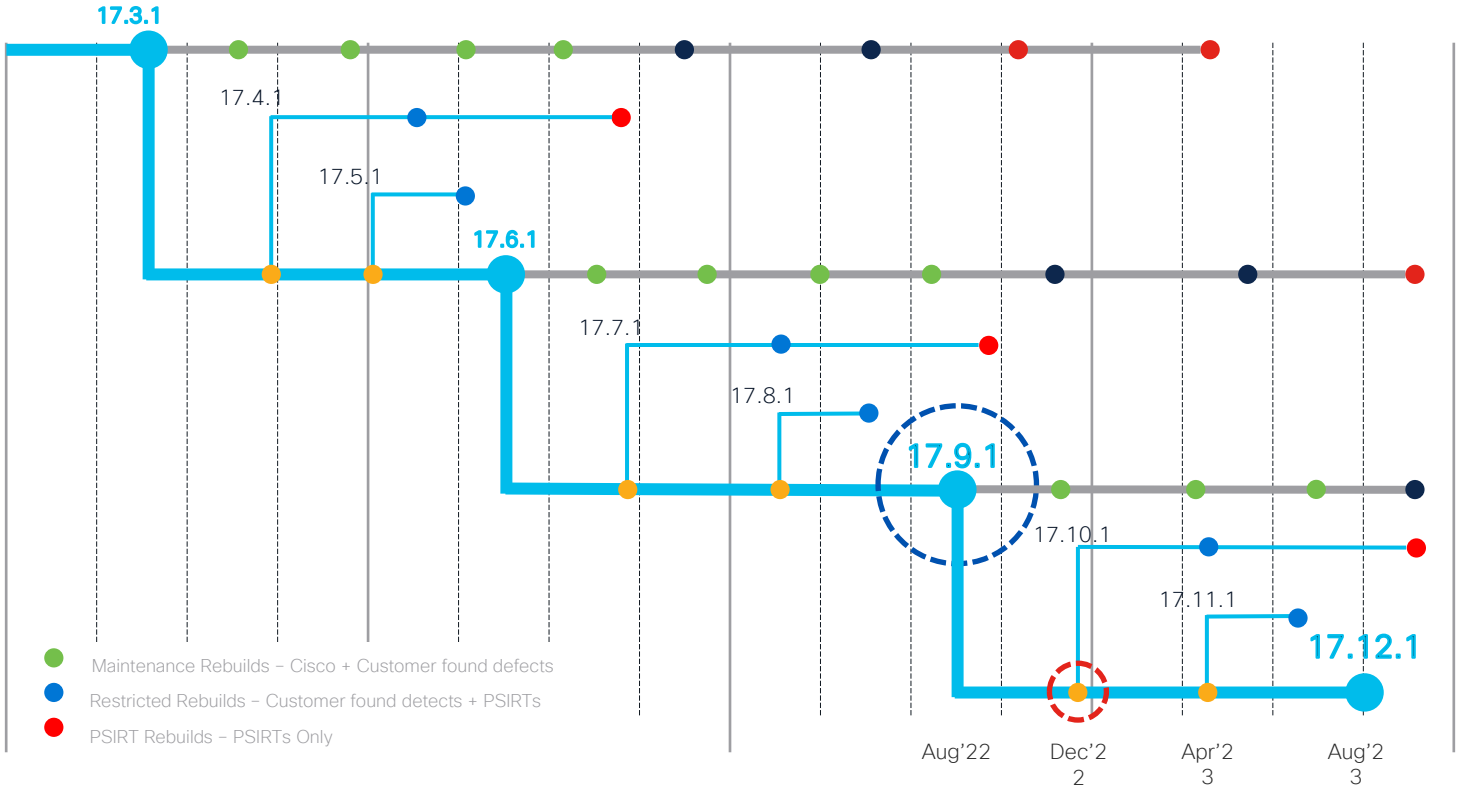

Reference

Catalyst 9000 Switching – Key Features

* Limited Availability (LA) only

IOS XE 17.7.1 (Dec'21) SMR	IOS XE 17.8.1 (Apr'22) SMR	IOS XE 17.9.1 (Aug'22) EMR
Enhanced Security ❖ 9500X/9600X – MACsec ❖ 9200/9300 – API Registration for Umbrella Switch connector	**Enhanced Security** ❖ 9300X – IPsec Phase2 – Multicast (SVTI) ❖ 9500X/9600X – WAN-MACsec, with HSEC license ❖ SW SUDI 2099 Enablement	**Enhanced Security** ❖ 9300X – IPsec Phase3 – NAT Traversal ❖ 9300X – VRF-aware IPsec ❖ Reflexive ACL
Overlays & Segmentation ❖ 9500X/9600X – MPLS and TE Phase1 ❖ 9500X/9600X – EoMPLS ❖ EVPN L3 TRM with MDT Data	**Overlays & Segmentation** ❖ EVPN L2 TRM	**Overlays & Segmentation** ❖ SDA LISP Graceful Restart for MAC cache ❖ SDA VN Extranet across SDA Transit
Forwarding & Features ❖ 9500X/9600X – L3 Routing (IGP, BGP) feature set ❖ Low priority Control packet mapping to Non-LLQ ❖ Bonjour – Micro-Location services	**Forwarding & Features** ❖ 9500X/9600X – Sampled Flexible NetFlow	**Forwarding & Features** ❖ Destination IP NAT scale enhancement ❖ PAT support for Enhanced NAT scale ❖ Conditional Static NAT using Route-map
High Availability ❖ 9400X/9600X – SSO & ISSU ❖ 9300X – xFSU	**High Availability** ❖ 9500H/9600 – Graceful Insertion & Removal (GIR)	**High Availability** ❖ 9400X – StackWise Virtual (Dual-Sup)
Platform & Programmability ❖ PTP on 9300 StackWise ❖ PTP on 9600 ❖ PTP AES67 compliance ❖ AVNU Certification – 9300 & 9500 ❖ gNOI reset.proto – tooling	**Platform & Programmability** ❖ PTP – G8275.1 ITU Telcom Profile on 9300/9300X ❖ 9500X/9600X – L3 Sub-Interface Queuing ❖ C9300 System Power-Consumption Reporting ❖ gNMI Native Configuration Yang Model ❖ Guest Shell HA – Guest-Share Folder Sync	**Platform & Programmability** ❖ 9500H – AVB support ❖ 9400X – Perpetual PoE support ❖ 9400X – Support for hosting multiple applications ❖ 9400X – 432 Port-Channels ❖ 9400X – 4K VLANs support
Hardware & Optics ❖ **C9300X-48TX / 48HX** – 48x mGig Switch (Cisco UADP2.0sec) ❖ **C9400X-SUP-2/XL** – Supervisor 2 (Cisco UADP3.0sec) ❖ **C9500X-28C8D** – 100/400G Fiber switch (Cisco S1 Q200) ❖ **C9600X-SUP-2** – Supervisor 2 (Cisco S1 Q200) ❖ **C9600-LC-40YL4CD** – 40x SFP + 4x QSFP Combo Linecard	**Hardware & Optics** ❖ **C9400-LC-48HX** – 48x mGig Linecard ❖ **C9400-LC-48XS** – 48x 1/10G Linecard ❖ **C9600-PWR-3KW** – 3000W AC PSU	**Hardware & Optics** ❖ **C9200CX** – 8-12x Compact switch (Cisco UADP2.0mini) ❖ **C9300LM** – 24-48x + Fixed Uplinks, ≤14 in. Deep (Cisco UADP2.0sec) ❖ **C9600-LC-32CD** – 32x 100G or 24x 100G + 2 x 400G QSFP Linecard

Catalyst 9000 Switching – Key Features

* Limited Availability (LA) only

IOS XE 17.10.1 (Dec'22) SMR	IOS XE 17.11.1 (Apr'23) * SMR	IOS XE 17.12.1 (Aug'23) EMR
Enhanced Security ❖ DHCP Snooping Glean ❖ Reflexive ACL for IPv4 ❖ MACsec Transparent Pass-Through ❖ MACsec Fallback-Key with HA support ❖ Secure Data Wipe (NIST 3-pass) ❖ 9400X – IPsec support	**Enhanced Security** ❖ 9300X - GRE over IPsec support ❖ 9400X - NAT-Traversal support for IPSEC ❖ 9500X/9600X - IPv6 SGACL support	
Overlays & Segmentation ❖ BGP EVPN over IPv6 Underlay* ❖ BGP EVPN Dynamic Peering ❖ BGP EVPN over IPsec - TRM support	**Overlays & Segmentation** ❖ BGP EVPN Micro-segmentation (VLXAN-GPO) ❖ BGP EVPN Mobility-Convergence enhancements* ❖ 9500X/9600X - BGP EVPN with SVL support ❖ 9500X/9600X - MPLS VPN - Inter-AS Option A ❖ 9500X/9600X - L2VPN (EoMPLS) Pseudowire Redundancy	
Forwarding & Features ❖ Stateful NAT64 ❖ NAT on L3 EtherChannel ❖ LACP Standalone Mode on L3 EtherChannels ❖ PTPv2 on StackWise Virtual ❖ 9500X/9600X - Bonjour & mDNS Routing	**Forwarding & Features** ❖ 9500X/9600X - Policy Based Routing (PBR) ❖ 9200/9200CX – Bonjour Service Peer	
High Availability ❖ 9500X/9600X - StackWise Virtual (Dual Sup SSO)	**High Availability** ❖ StackWise Virtual (Dual Sup SSO) on 9500X-60L4D	
Platform & Programmability ❖ ERSPAN on AppGig interface ❖ YANG 1.1 Support ❖ YANG Config Model support for native gNMI ❖ Guest Shell HA Guest-Share Folder Sync	**Platform & Programmability** ❖ 9400X - Customized SDM Template ❖ 9500X/9600X – ERSPAN ❖ YANG Oper model support for PTPv2 and AVB ❖ Switching Telemetry for Cisco DNA Center*	
Hardware & Optics ❖ **C9500X-60L4CD** - 50G + 400G Fiber switch (Cisco S1 Q200) ❖ **9600X** - 50G SFP Linecard support	**Hardware & Optics** ❖ **9500X/9600X** – Embedded-PHY 1G SFP support ❖ Support for QSFP-100G-FR-S ❖ Support for SFP-10/25BXD-1 and BXU-1 ❖ Support for QSFP-4x10G on 100G PSM4	

* Software Roadmap - Subject to Change!

C9K Next-Gen IOS XE
Platform Infra Overview

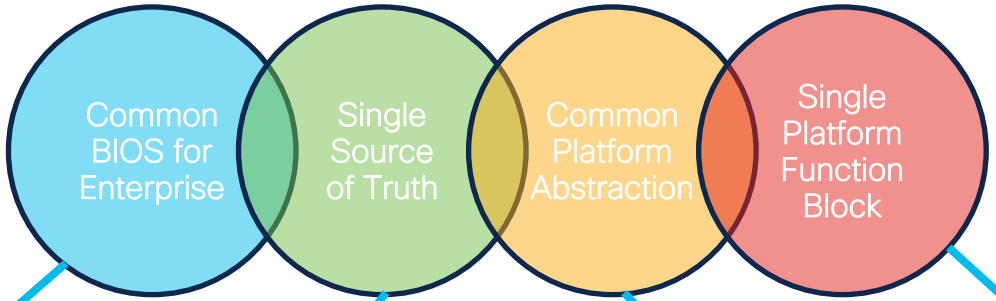

Common BIOS

- Common BIOS repository for all Enterprise platforms
- Standardized customizations

Single Source of Truth

- Partially-automated way to transfer hardware attributes to software
- Single place for all engineers to go to

Common Platform Abstraction

- Model-driven Device APIs consistent across Cisco
- Easier to adopt for subsequent platform (limit changes to certain files)

Platform Function Block

- Common CMAN, CMCC, EMD, IOMD
- For Enterprise platforms running on IOS XE (Polaris)

C9K NG IOS XE – 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, Abstraction

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

Differentiated ASIC + SDK

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

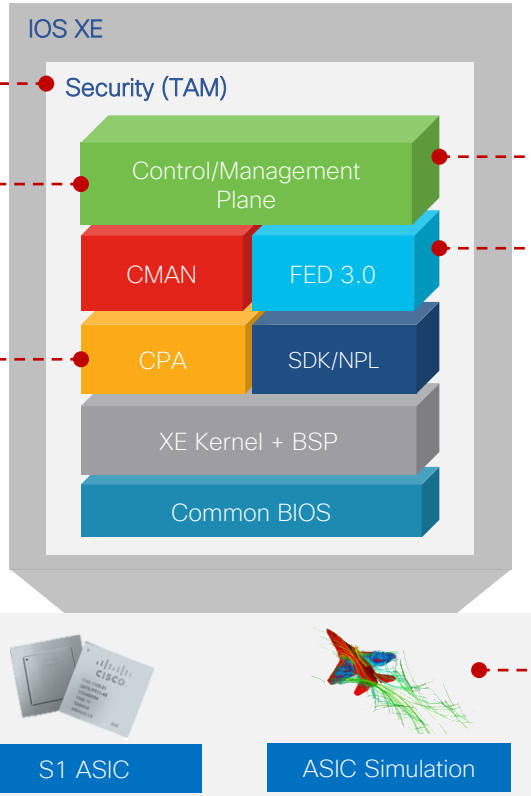

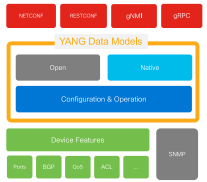

Programmable

- 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 NG IOS XE – Platform Infra

Architecture Evolution – Common Platform Architecture (CPA)

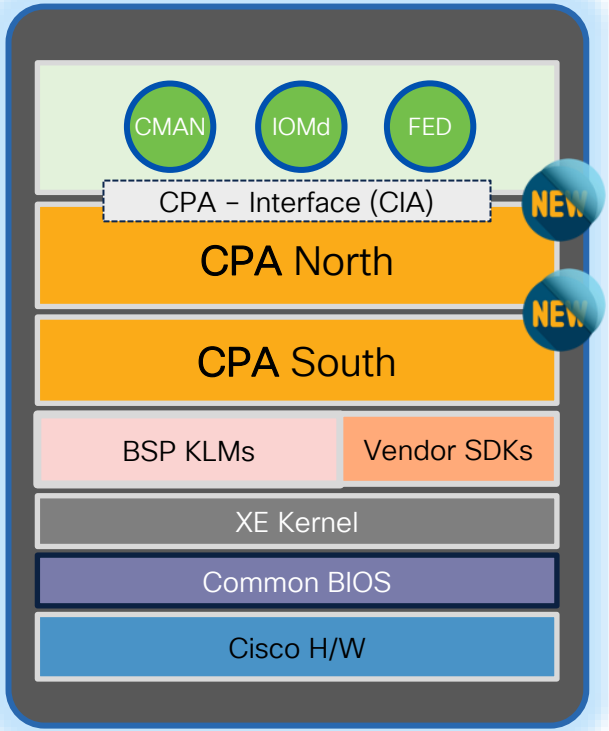

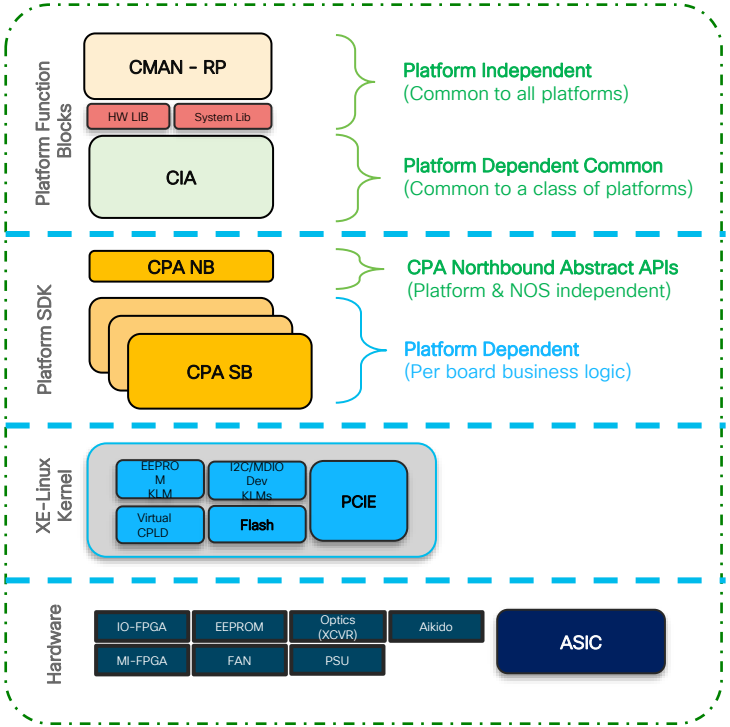

Leveraged by NG Platforms

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 - and **acts as a "single source of truth" for hardware data.**

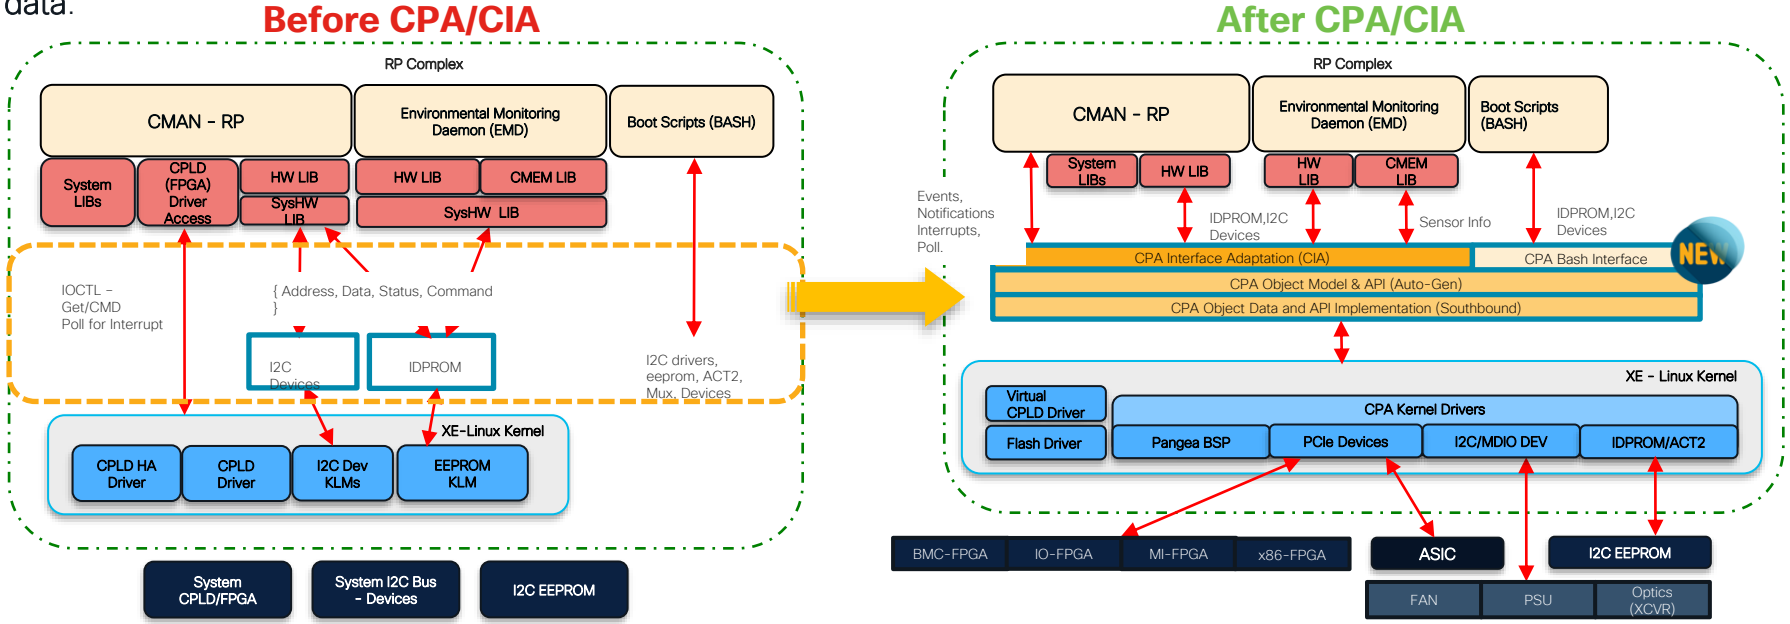

CPA provides common & consistent inventory and device-access helper functions (APIs) for admin-plane features

C9K NG IOS XE – What is CPA?

Common Platform Architecture (CPA)

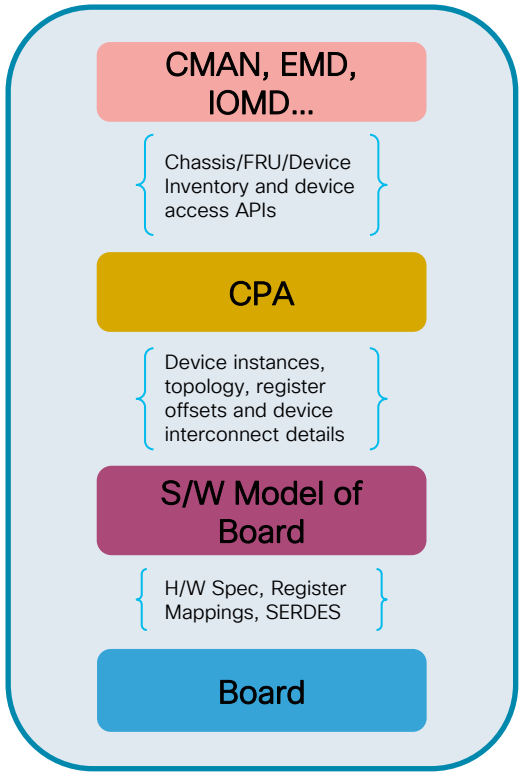

- **Model driven H/W – S/W Disaggregation layer**
- **Single Source of Truth (SST) for platform data**
- **Has 3 layers** : OS-Adoption, North Bound and South Bound
- Provides common/consistent inventory and device access helper functions/APIs for admin plane features (Platform SDK)
- Hides board/device interconnection details from NOS and provides abstract board/device access APIs
- Code generation to generate platform abstraction APIs & data objects
- CPA inventory scope has a complete system view
  - A chassis has slots, which has boards/trays, which has devices…
- Abstraction APIs to set/get attributes and events from platform devices
- SB Implementation (manual) to perform platform/device specific operation using the underlying driver interface layer
  - Kernel Loadable Modules/User Mode Drivers
- CPA Model and APIs are compatible to ONLP
- Platform with BSP and CPA can be independently verified

C9K NG IOS XE- What is CPA?

Common Platform Architecture (CPA)

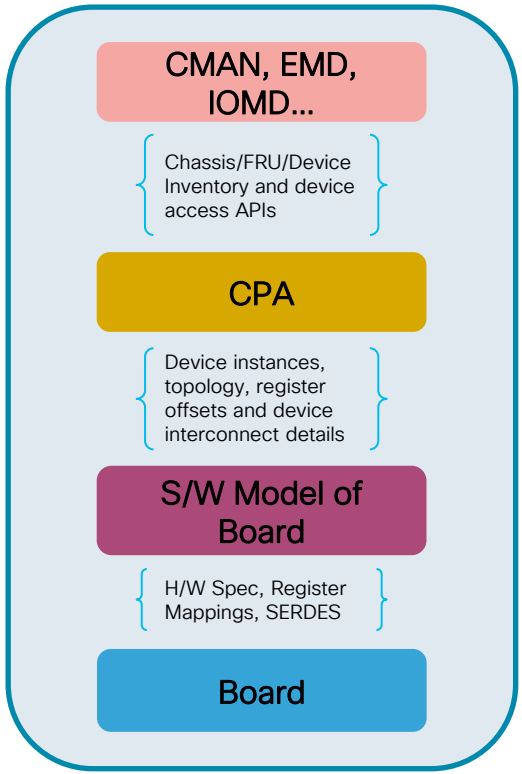

Platform Functional Blocks

- CPA NB/SB APIs provide a consistent way of accessing inventory and device helper functions. This allow a single source and single instance of Platform Functional Blocks for all EN platforms

Model Driven

- Platform devices are modeled into objects and attributes, along with the inter-relationships using YANG
- Per-board instantiation and inter-relationship is described in JSON files.
- A chassis is a collection of board JSONs, dynamically instantiated
- YANG model is generic and reusable across platforms
- Device models use access method to describe connectivity details

Board Specification

- Capture all attributes in a single JSON file as a single source of truth per board; JSON file is filled in by hardware and platform engineers
- Auto converted into YAML file to be consumed by CPA layer

C9K NG IOS XE – Modular CPA Layers

Common Platform Architecture (CPA)

CPA framework consists of different layers as shown in this diagram:

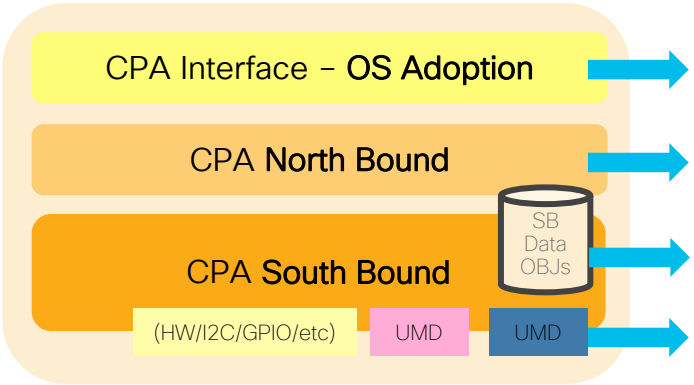

Translation layer to maintain the existing Polaris platform functions/VFTs, data structures, trace/log mechanism and error codes (NOS Specific)

Consists of device YANG models, codegen plugins, CPA infra source files, public header files and Test code (NOS-independent Platform APIs across different platforms)

Consists of per FRU yaml files, common implementation and platform specific implementation logic and driver interface code (Board specific)

Common/Re-usable Cisco and Vendor-supplied Userspace drivers and KLMs (Device specific)

CPA NB APIs – SB Objects and corresponding implementations are organized into 3 groups:

- **Core** : Consists of CMAN, EMD functions (inventory, chassis, sprom, Fan, PSU, LED, watchdog, FPD etc.)
- **Optics** : Consists of Port controller, Port, PHY and Transceiver functions
- **NPU** : Consists of ASIC discovery/initialization/reset/serdes tuning functions

C9K NG IOS XE – FED 3.0 Overview

FED Before & After

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

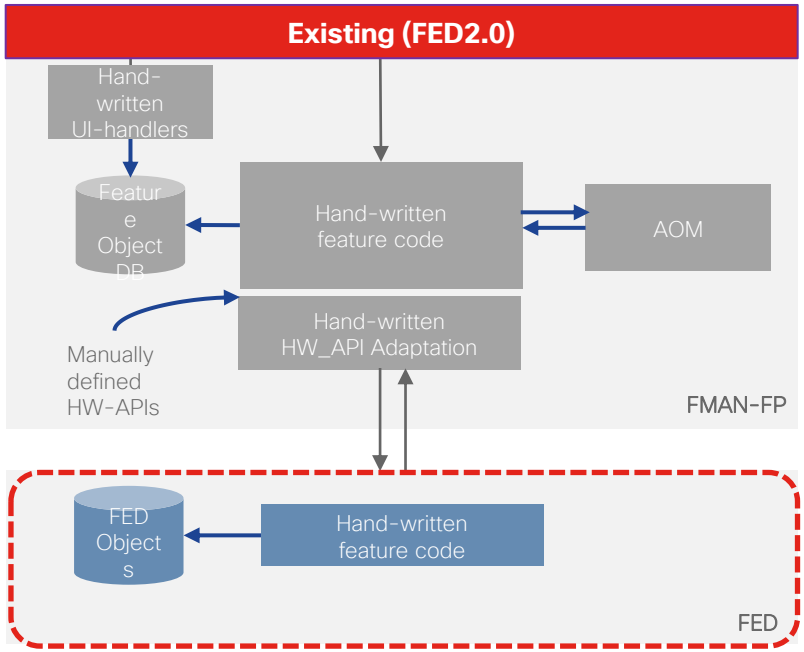

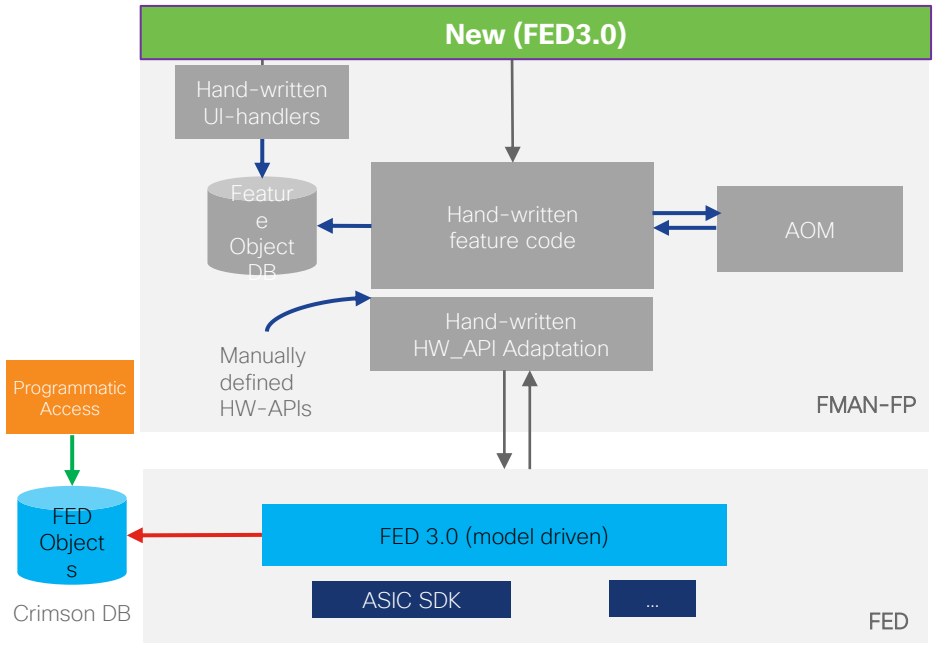

- Controller visibility of HW forwarding objects
- Time series for HW forwarding objects
- Stateful restart HW forwarding objects
- Automatic code-generation for object life-cycle

Manual
Auto-generated
New capability

C9K NG IOS XE - Overview

Hardware Forwarding Architecture

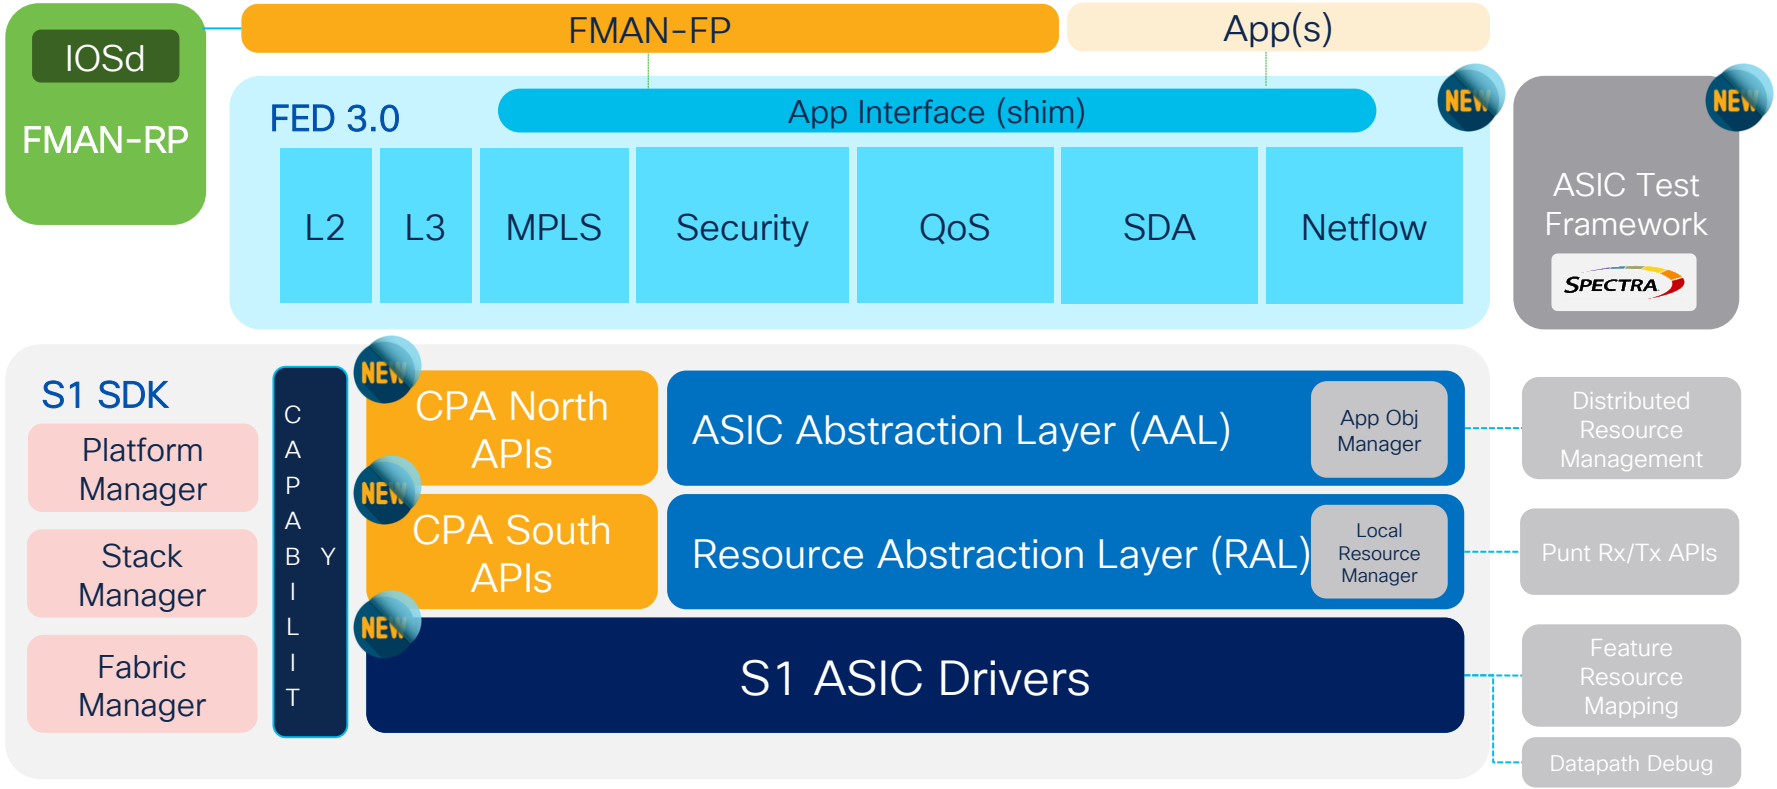

Glimpse into the Future

for Enterprise Switching

Where are things going?

Features & Scale

More details in PPT Notes

MAC & IPv4/v6 Scale

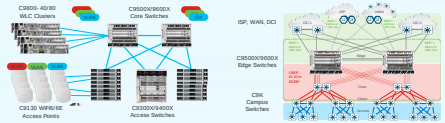

Increasing MAC scale for Wireless & IOT

- WLC to Core SVI scale growing ≥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 ≥850K and IPv6 GRT is ≥50K

Multicast & MVPN

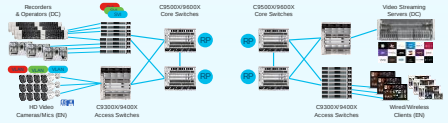

Higher L2 & L3 Multicast scale in Hardware

- Flexible L2/L3 multicast group allocation in CEM
- Hardware Replication and specialized QoS

Multicast VPN with GRE or mLDP in Hardware

- MVPN MDT with Profile 0, 1, 12 & 13
- NG-MVPN with Profiles 14 & 15

VoQ & HBM

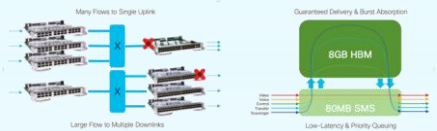

Virtual Output Queues (VoQ) for high throughput

- 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
- HBM buffers for burst absorption & guaranteed delivery

OGACL & SGACL

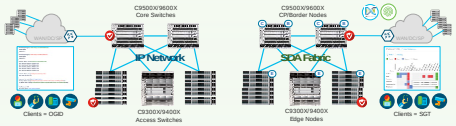

Object-Groups map IP/mask to Labels in CEM

- 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

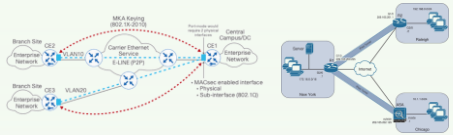

256-bit Hardware Encryption over L2

- 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
- P2MP DVTI IPsec with IKEv2, ESP

SDA, AVC & ZTN

SD-Access Fabric with Group-Based Policy

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

App & Endpoint Identification & Analytics

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

Summary

Wrap-Up

Catalyst 9000 with Programmable ASICs

Benefits for your network

FLEXIBILITY and ADOPTABILITY

Enabling Network Evolution
on your journey to
Intent-Based Networking

Catalyst 9000 with Cisco IOS XE

Benefits for your network

One Release Train

Operational Efficiency,
Consistency in Behavior,

Run Any Feature Anywhere

Feature Velocity
across Platforms

Software Update & Patching

Sub-package upgrades,
Cold and Hot Patching

Comprehensive Programmability

Object based model,
Netconf/REST Interfaces

Trustworthy & Secure Platform

64-bit ASLR, Secure Boot,
Hardware TAM

Would You Like to Know More?

Catalyst 9000 X Series Collateral

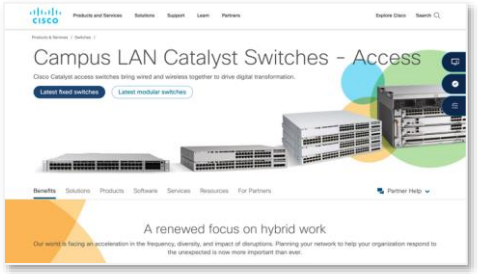

- Cisco.com - Enterprise Networks - Switching
- Catalyst 9000 Switches At-a-Glance
- Catalyst 9000 Frequently Asked Questions

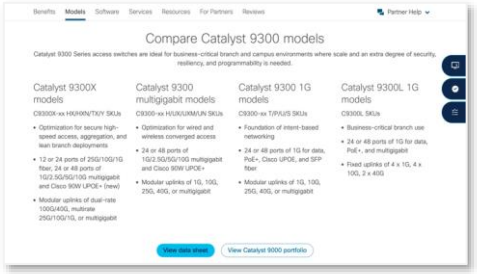

- Catalyst 9300 Series Switches data sheet
- Catalyst 9400 Series Switches data sheet
- Catalyst 9500 Series Switches data sheet
- Catalyst 9600 Series Switches data sheet

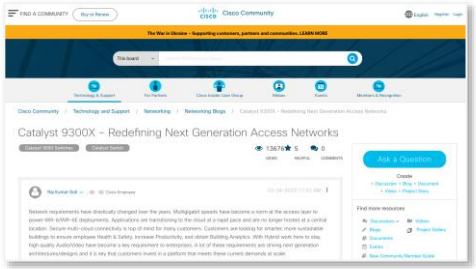

- Catalyst 9300X Technical Blog (Community)
- Catalyst 9400X Technical Blog (Community)
- Catalyst 9500X Technical Blog (Community)
- Catalyst 9600X Technical Blog (Community)

- Under the Hood of the Catalyst 9000X (TFD)
- Network Insiders Podcast
- Meet the C9300X
- Meet the C9400X
- Meet the C9500X and C9600X

Would You Like to Know More?

Cisco IOS-XE & Cisco DNA Resources

cisco.com/go/dna

cisco.com/go/iosxe

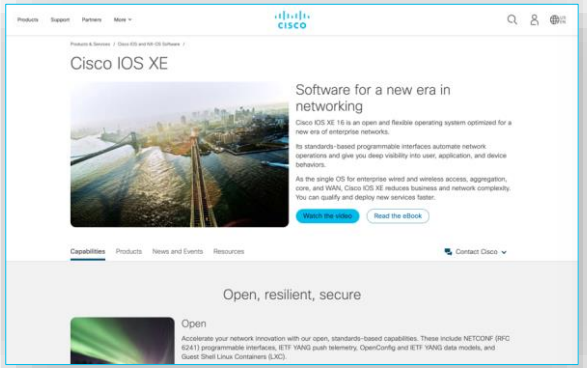

cisco.com/go/smartlicensing

cisco.com/go/dnacenter

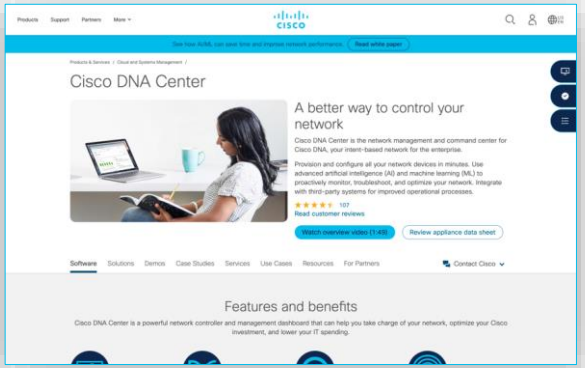

Would You Like to Know More?

Catalyst 9000 Series Enterprise Switches

- **cisco.com/go/cat9k**
- Cisco Catalyst 9000 At-a-Glance
- Cisco Catalyst 9000 Family FAQ
- Catalyst 9000 Series - Cisco Community
- Catalyst 9000 Series – CiscoLive Library

3rd Edition - 2022

cs.co/cat9kbook

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 until February 24, 2023.

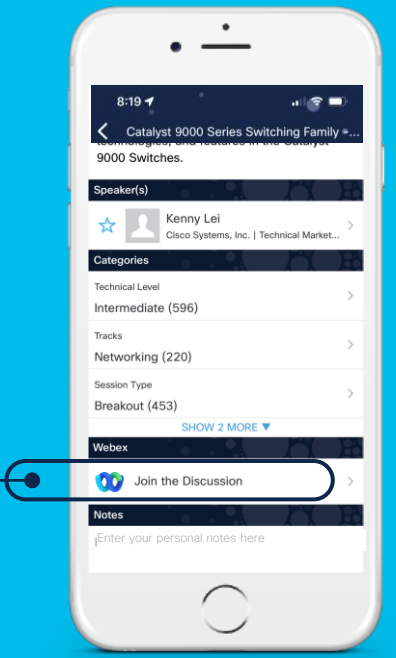

Complete your Session Survey

- Please complete your session survey after each session. Your feedback is important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (open from Thursday) to receive your Cisco Live t-shirt.

- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Session Catalog and clicking the "Attendee Dashboard" at https://www.ciscolive.com/emea/learn/sessions/session-catalog.html

Continue Your Education

Visit the Cisco Showcase for related demos.

Book your one-on-one Meet the Engineer meeting.

Attend any of the related sessions at the DevNet, Capture the Flag, and Walk-in Labs zones.

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


cisco
The bridge to possible
Thank you
cisco Live!


CISCO Live!
ALL IN