The bridge to possible

Cisco UADP & Silicon One ASIC Architecture & Innovations

Catalyst 9000 Series

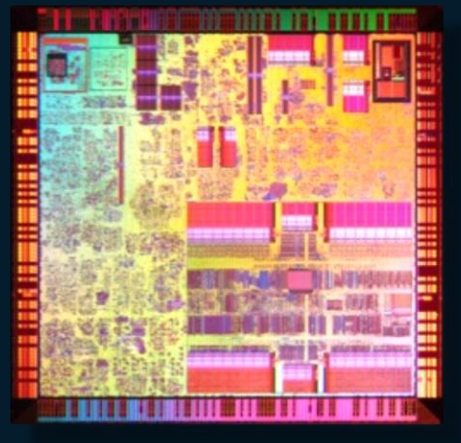

Kenny Lei – Technical Marketing Engineer

BRKARC-2091

Content by Shawn Wargo

cisco Live!

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

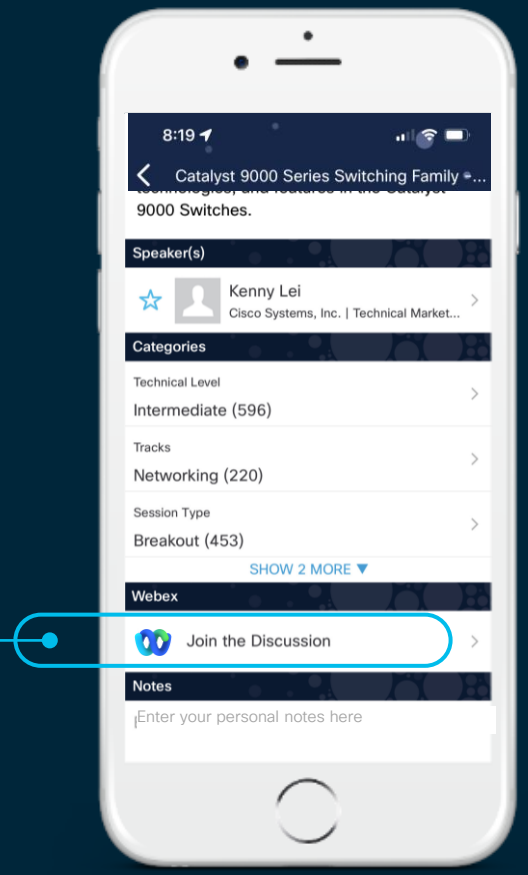

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKARC-2091

Webex spaces will be moderated by the speaker until June 17, 2022.

Agenda

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

Cisco Catalyst 9000 Switching Portfolio

One Family from Access to Core – Common Hardware & Software

February 2022

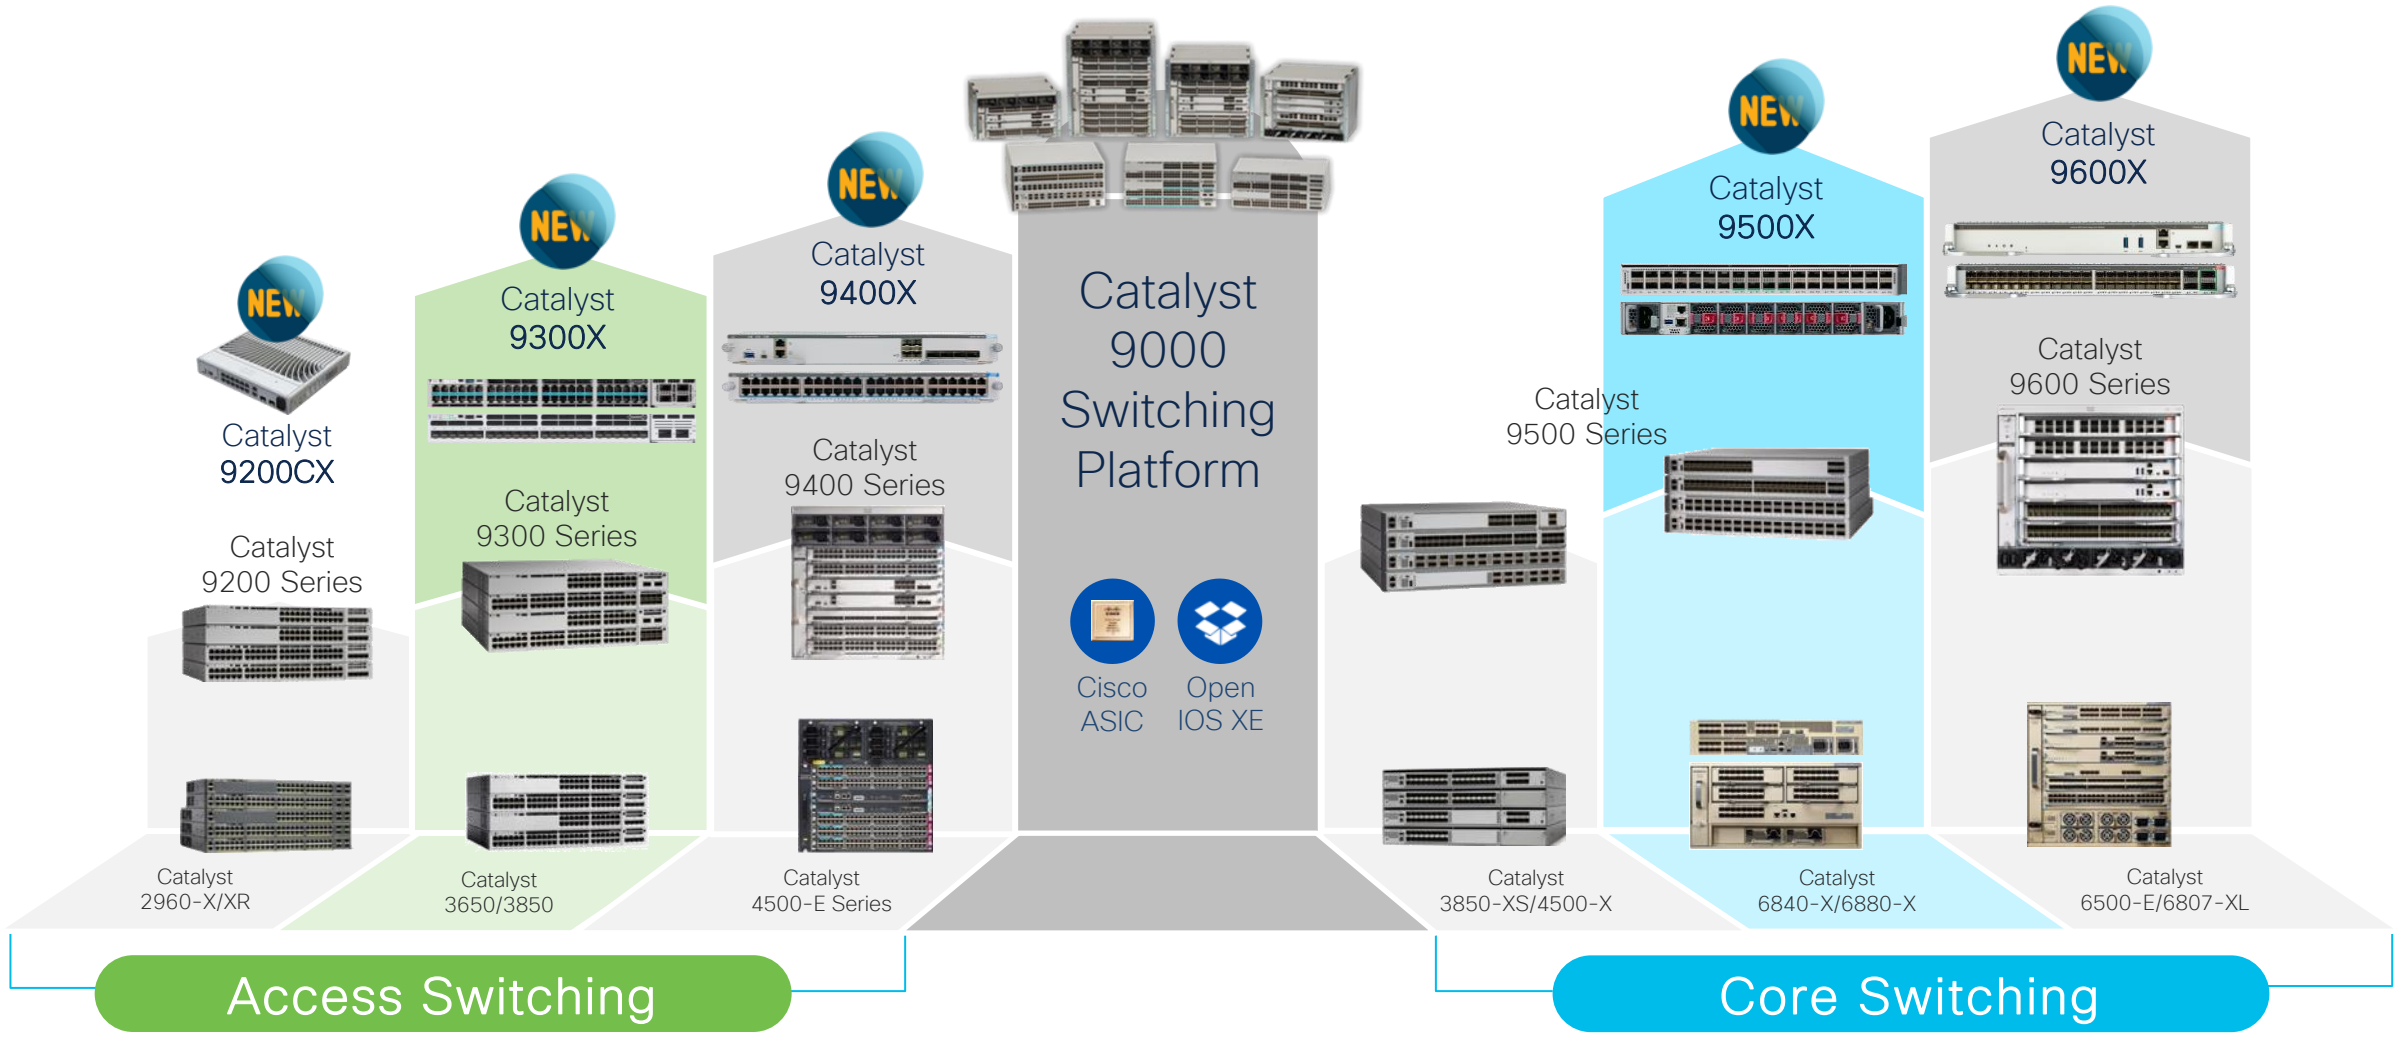

Catalyst 9000 Series – Common Building Blocks

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

Custom ASICs – Programmable Silicon

Cisco **Unified Access Data-Plane** (UADP®)

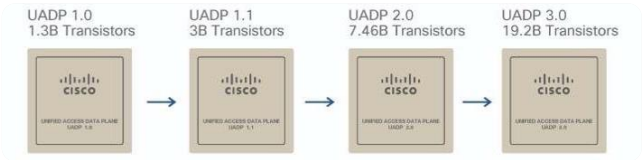

Cisco **Silicon One**™

Flexible Pipelines
Investment Protection

Adaptable Tables
Universal Deployment

Scalable Resources
Enhanced Scale and Buffering

Flexible & Programmable ASICs – Adapt to New Technologies

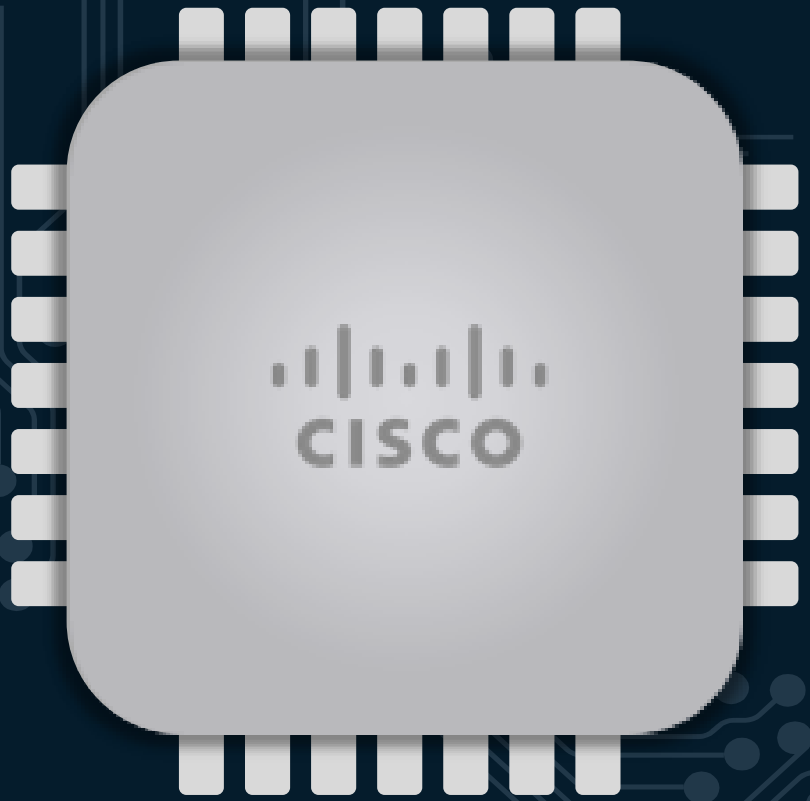

Why do we need ASICs?

What is an ASIC?

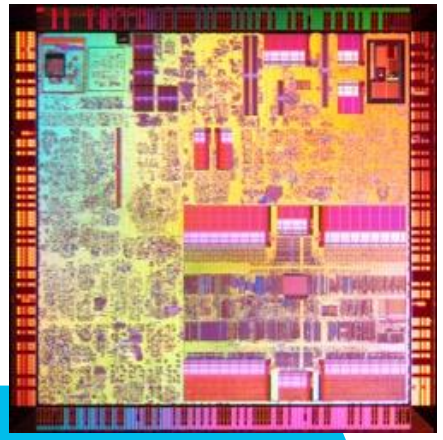

An Application Specific Integrated Circuit (ASIC) is a silicon microchip designed for a specific task …

… 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.

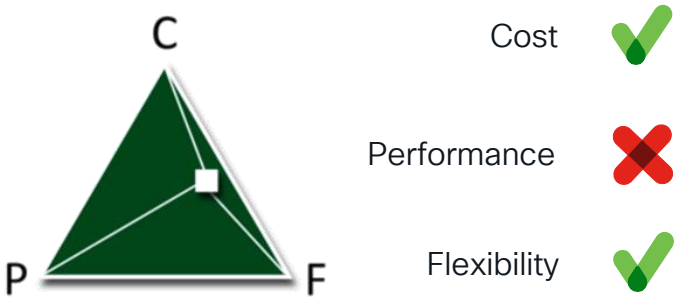

Remember, this is **Millions - Billions** of packets per second

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 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

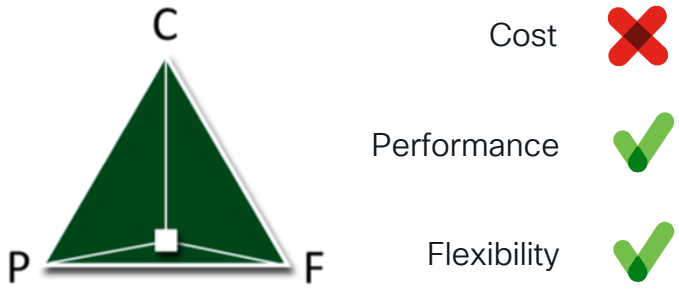

What does an ASIC do?

ASICs are fundamental to network devices

Network 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”

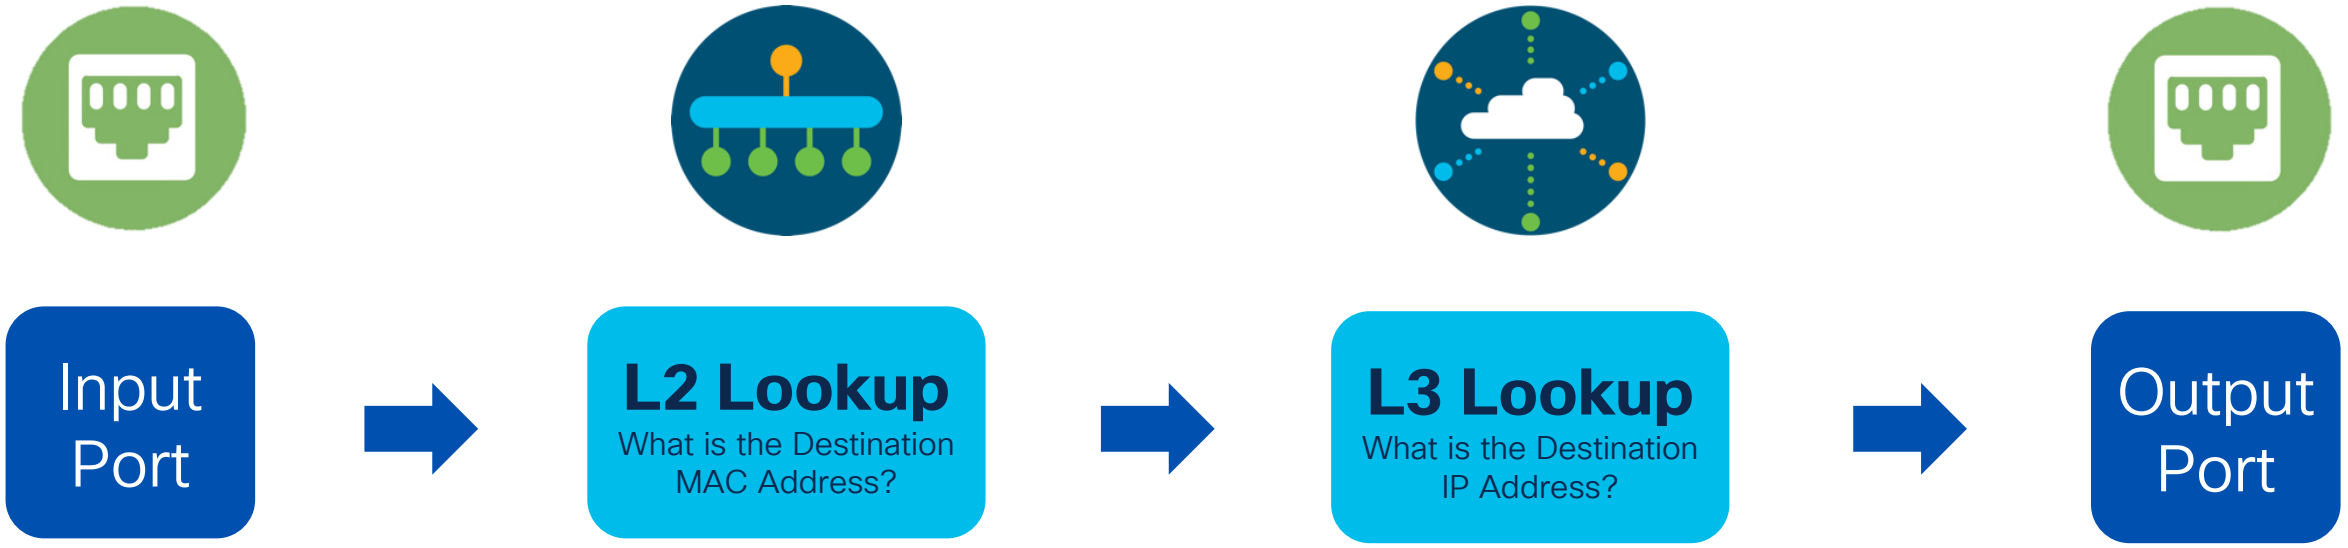

Modern speeds are **Gigabits per second** (Gbps) - trending to **Terabits** (Tbps)

What does an ASIC do?

ASICs are fundamental to network devices

In addition – network devices can perform special processing tasks

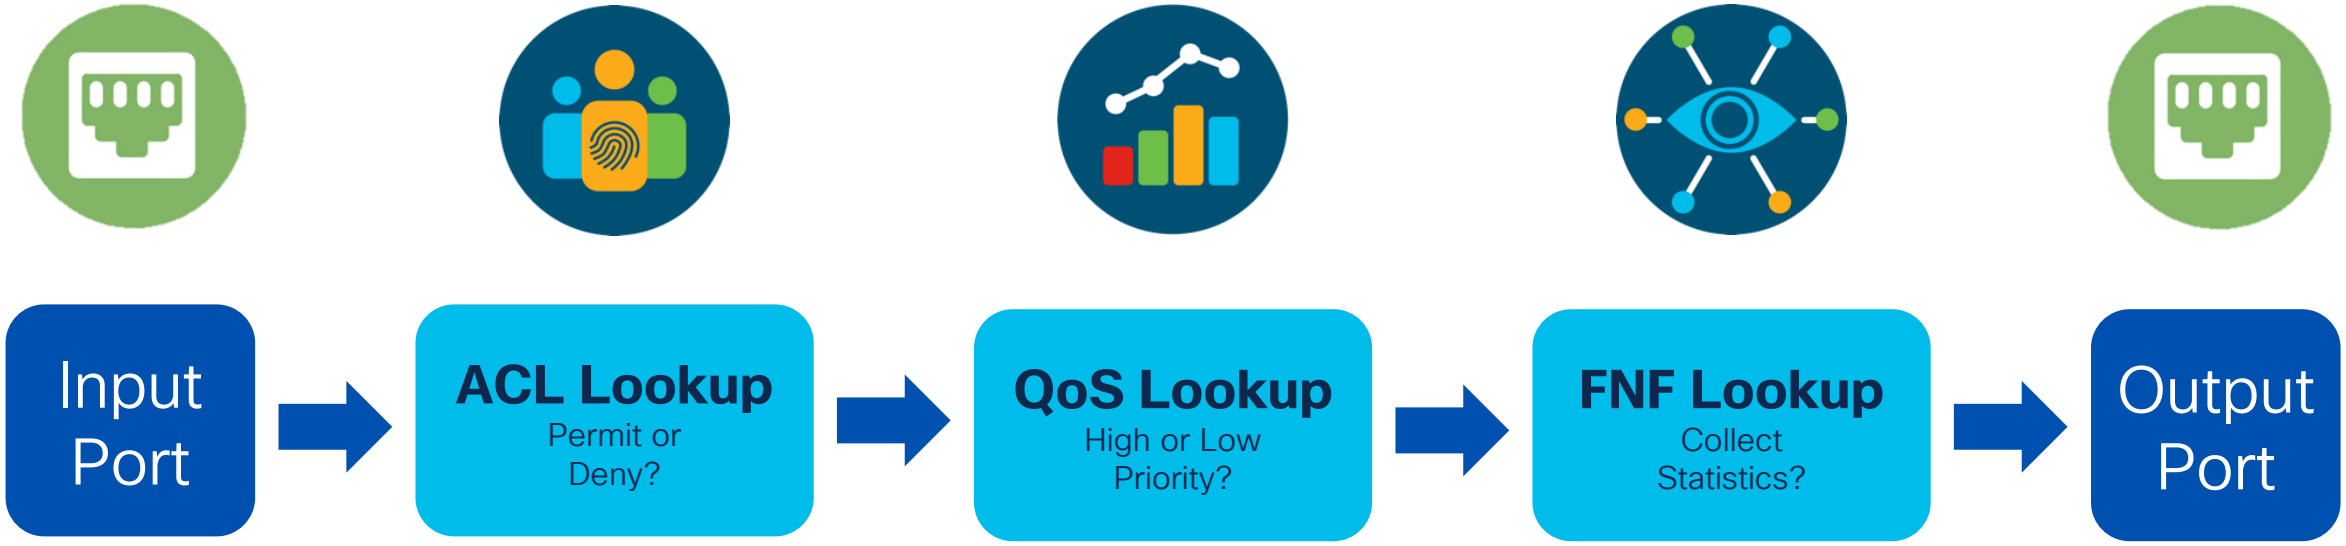

Common services are **Access Control**, **Quality of Service** and **Flow Analytics**

Traditional Network ASICs

Fixed Pipelines

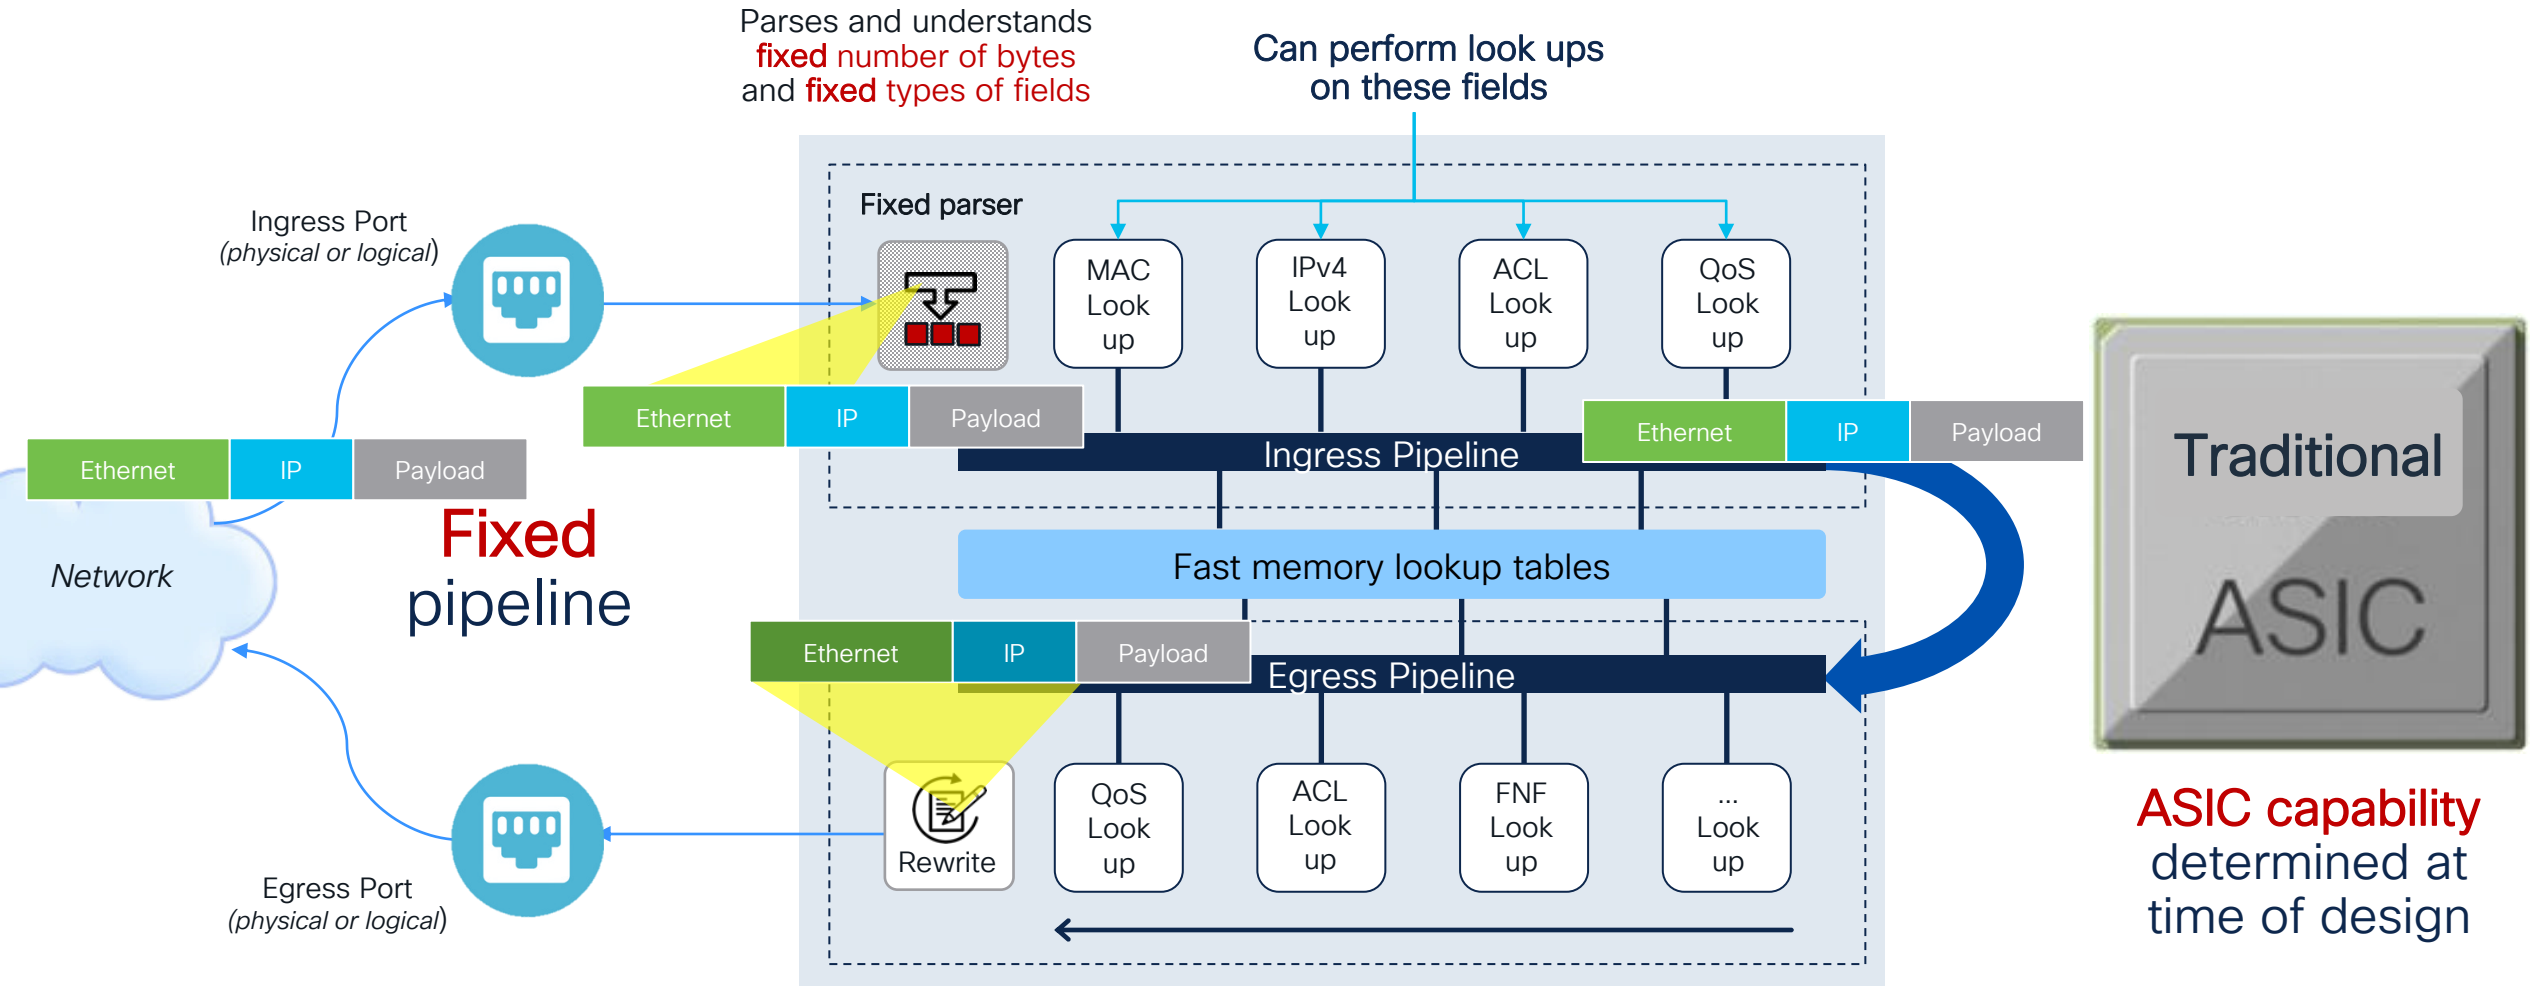

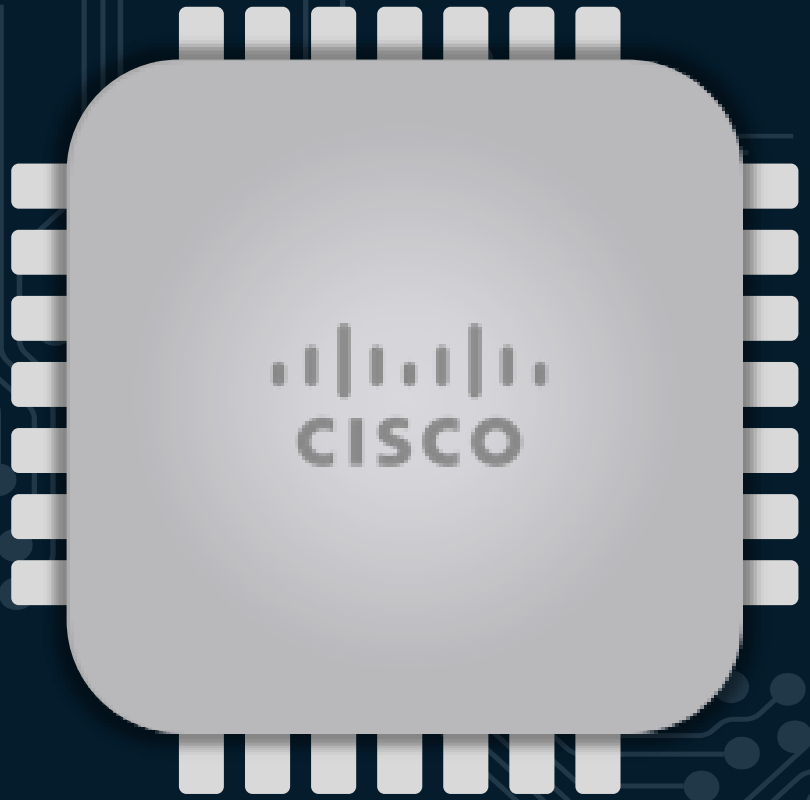

Flexible ASICs

for Enterprise Switching

Programmable Network ASICs

Flexible Parsing

Flexible Parsing

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

Parses and understands multiple programmable headers with flexible field definitions

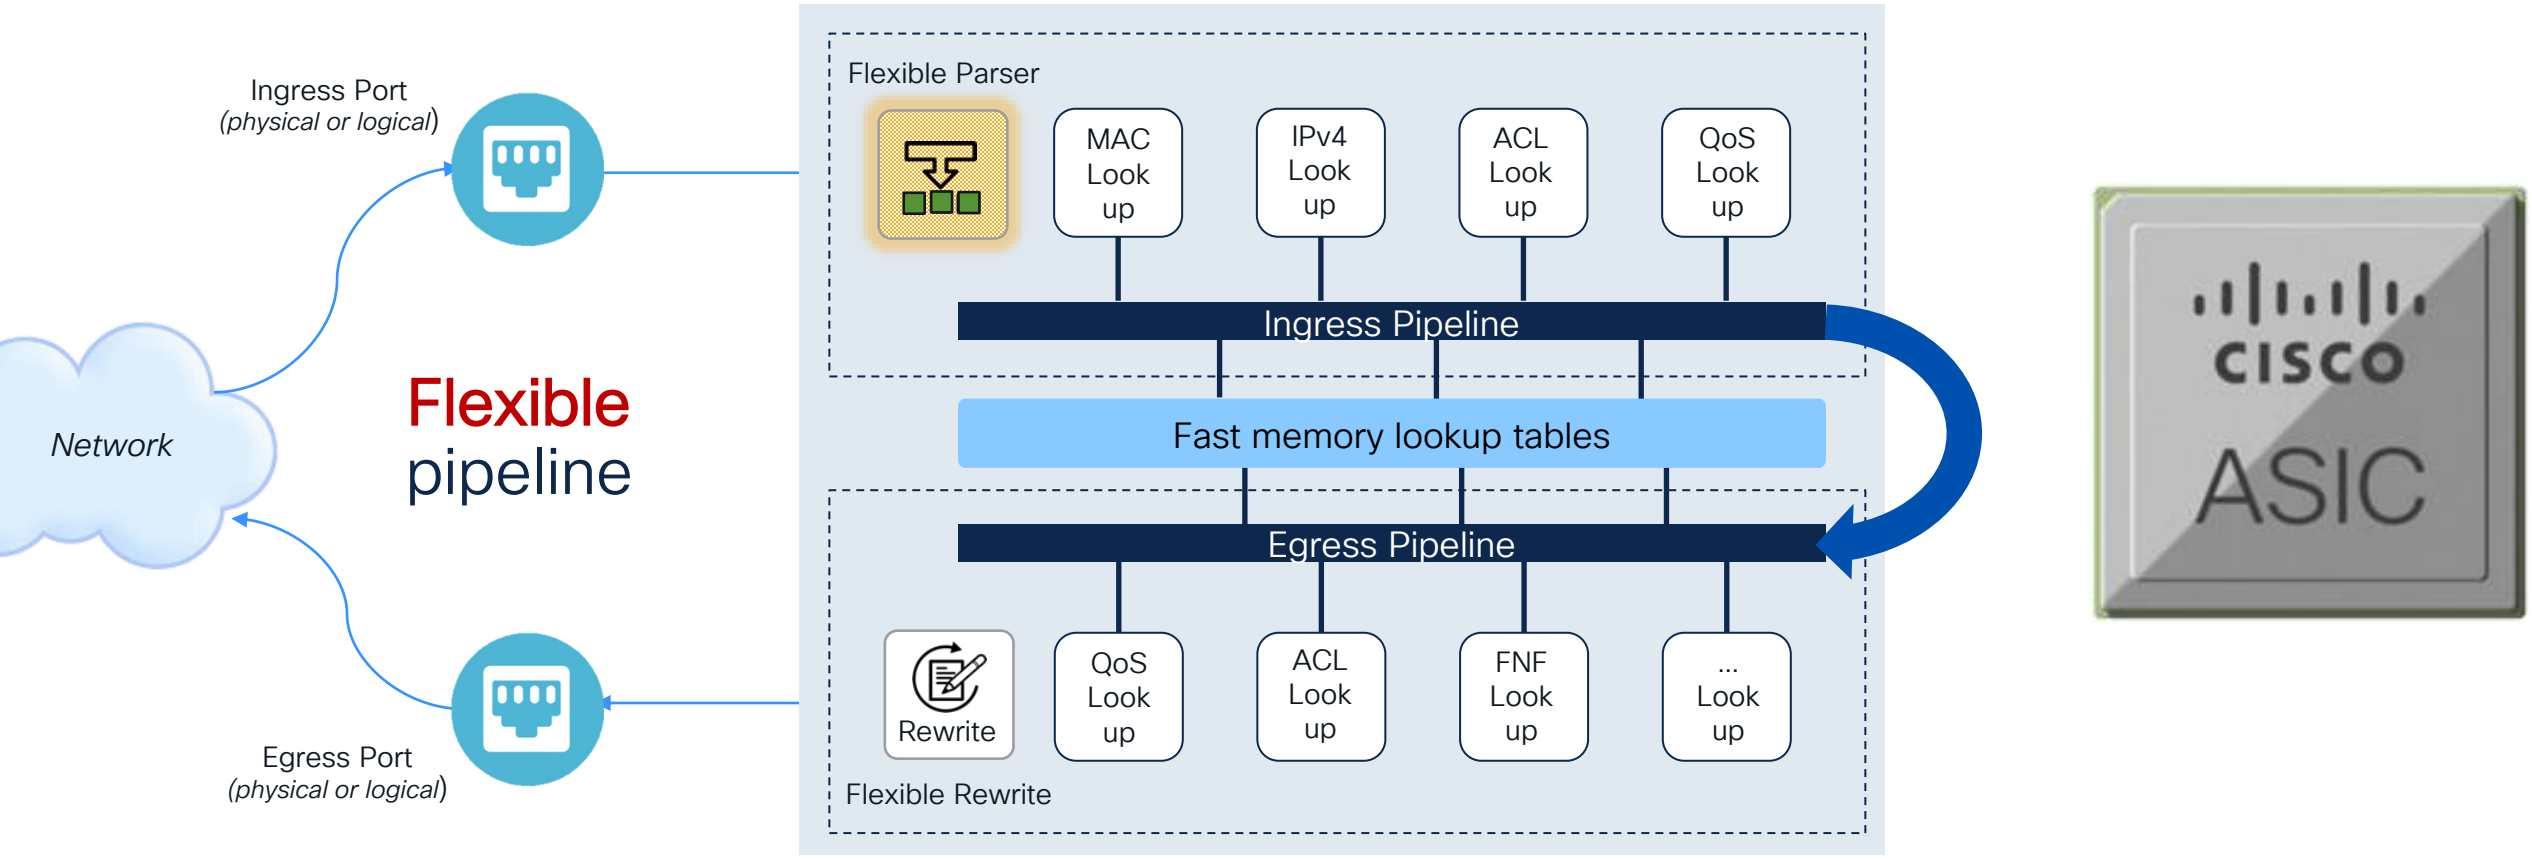

Programmable Network ASICs

Flexible Lookups

Flexible Lookups

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

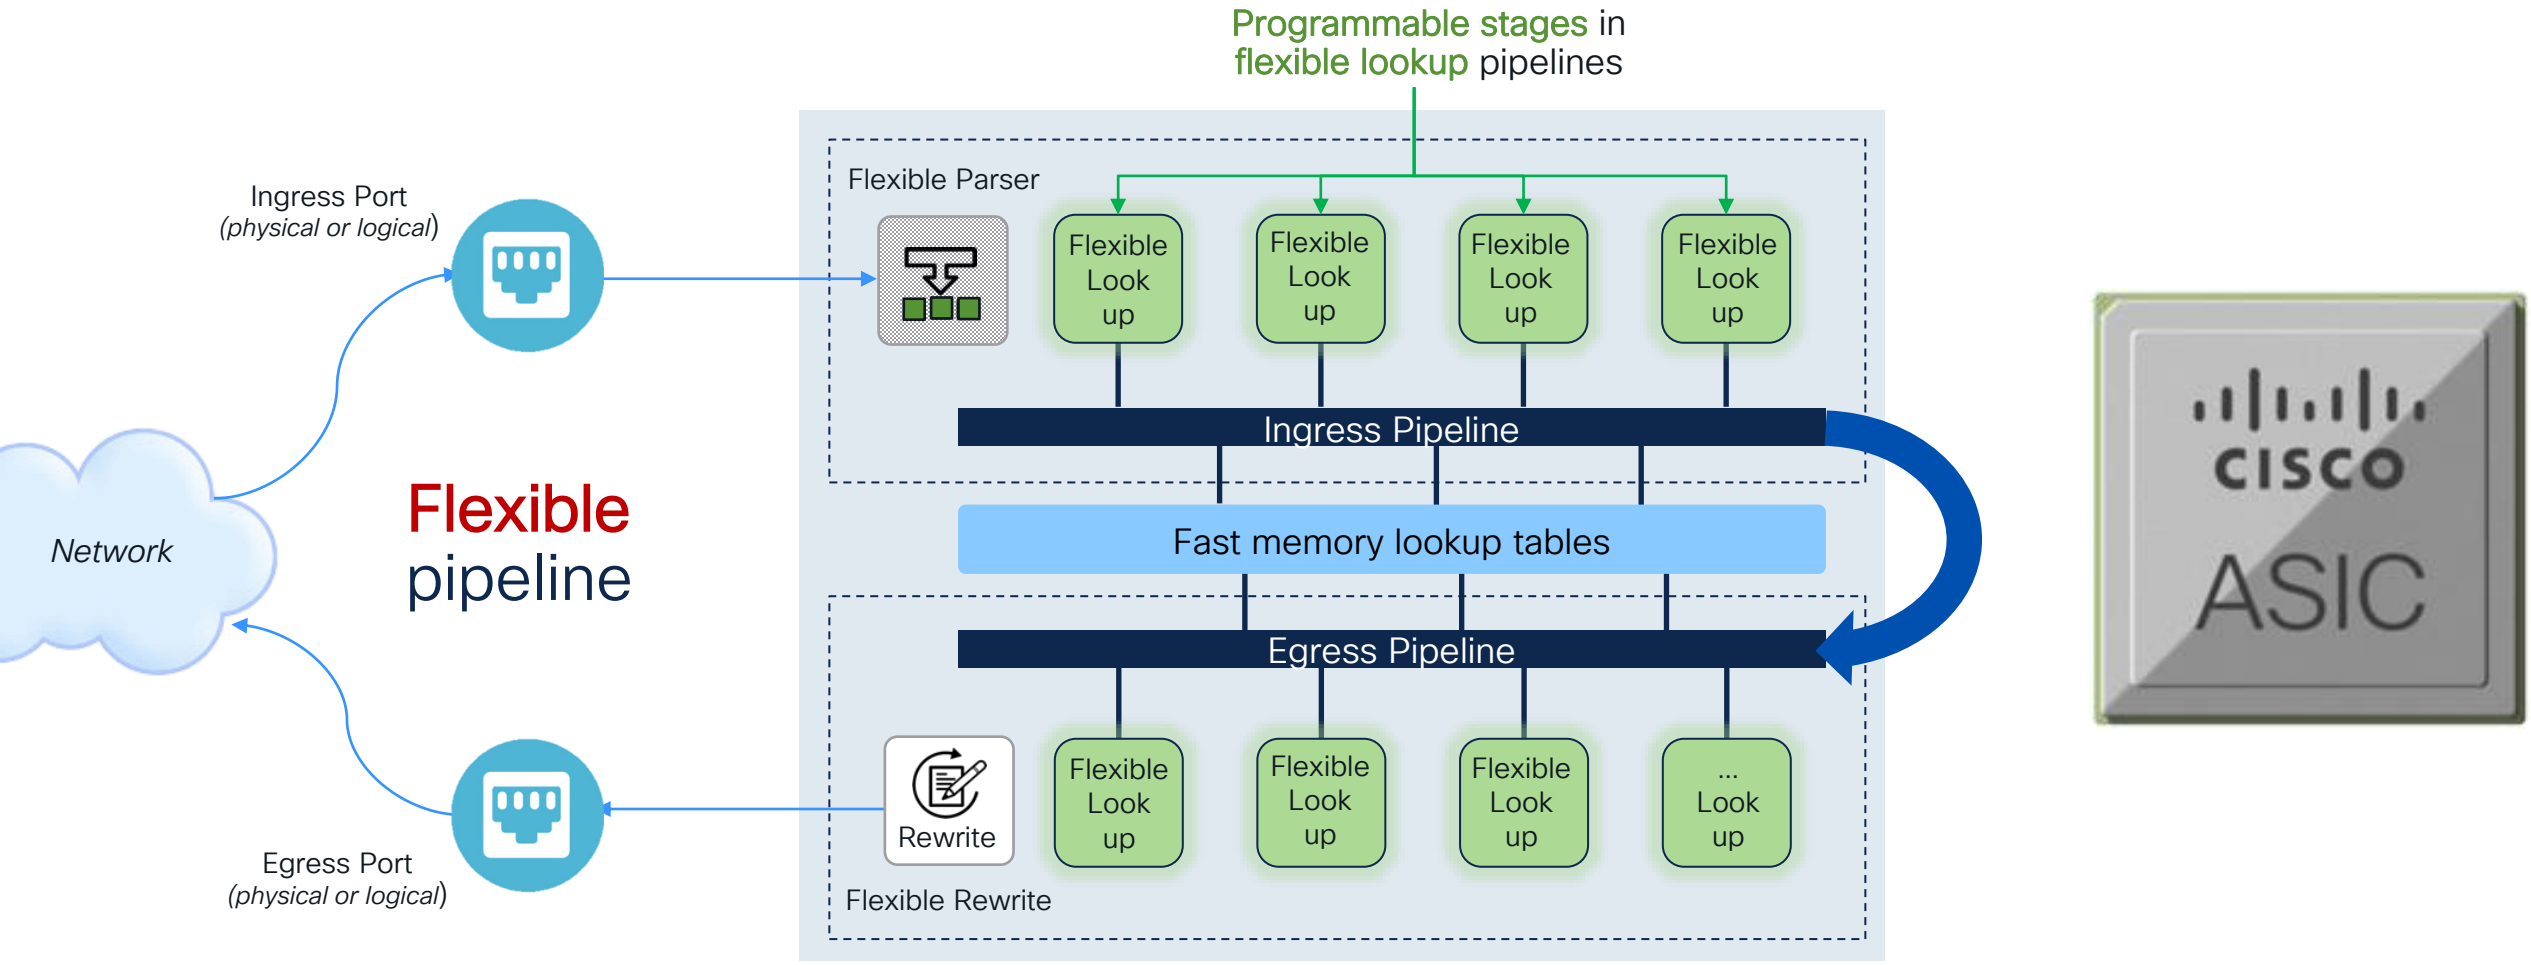

Programmable Network ASICs

Flexible Rewrites

Flexible Rewrites

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

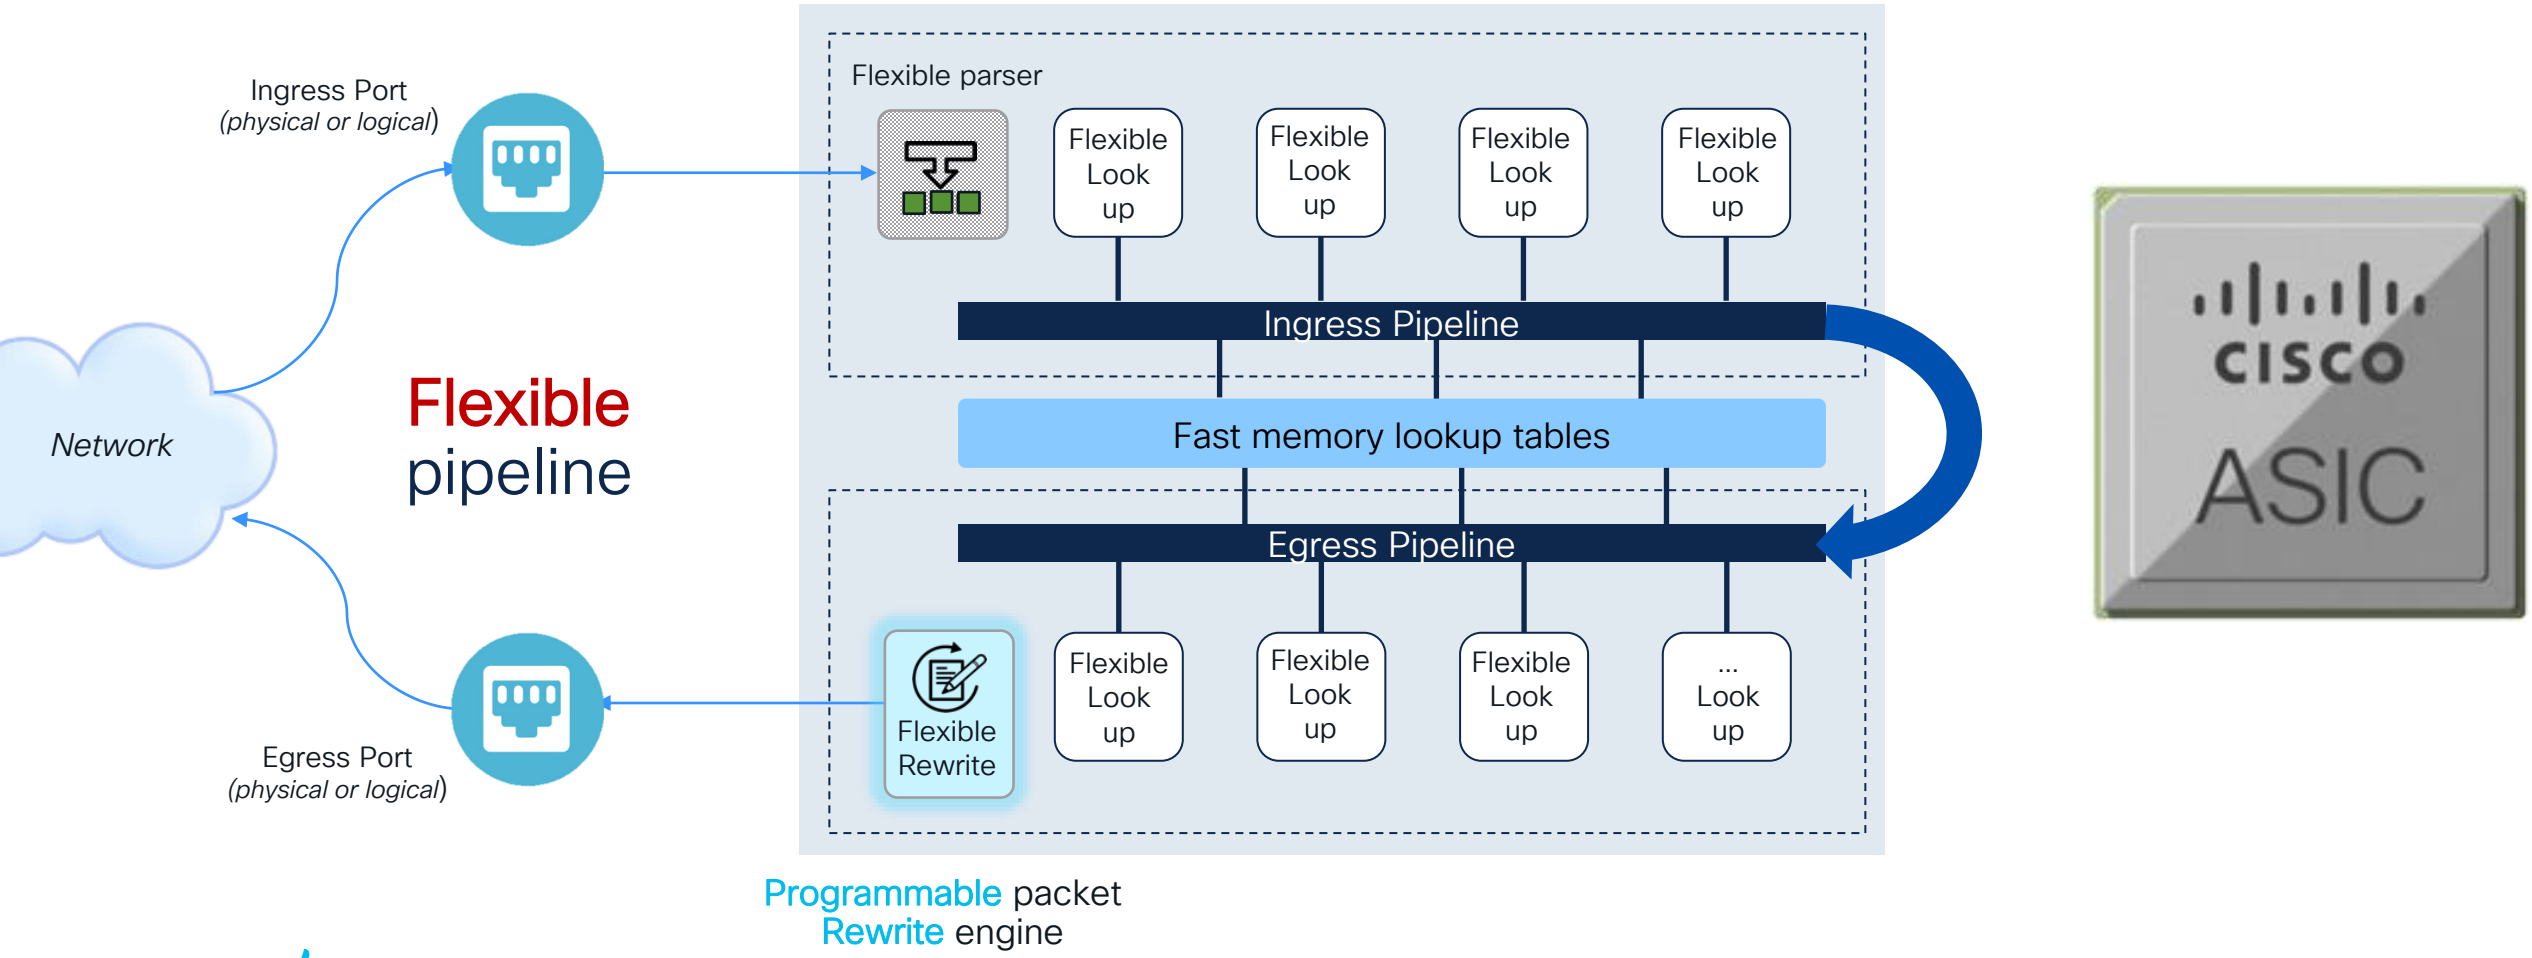

Programmable Network ASICs

Optimized Recirculation

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

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

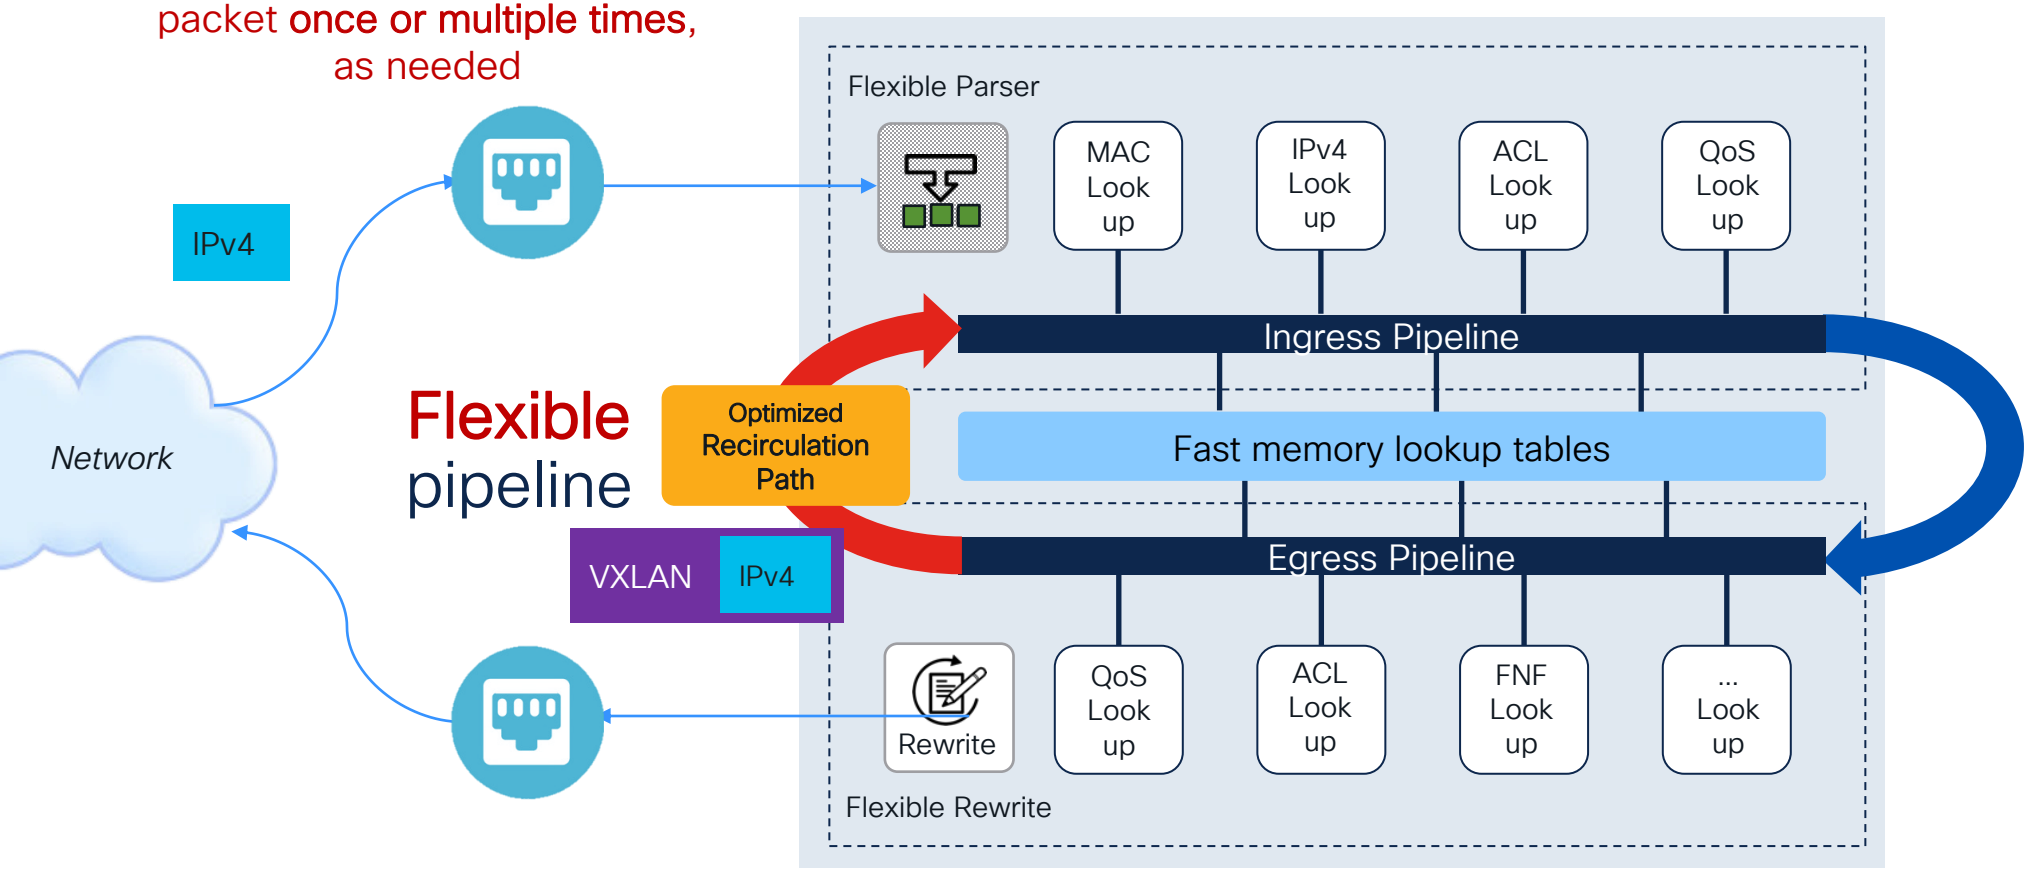

Programmable Network ASICs

Balancing Cost, Performance & Flexibility

Programmable ASIC

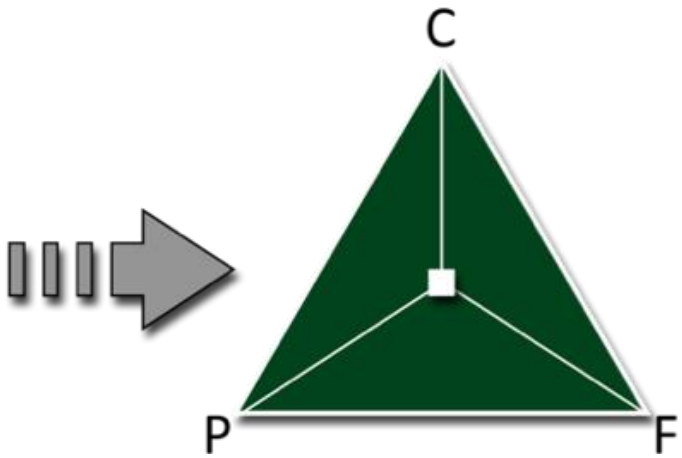

Cost ✓

Performance ✓

Flexibility ✓

ASIC Engineer

Microcode Update

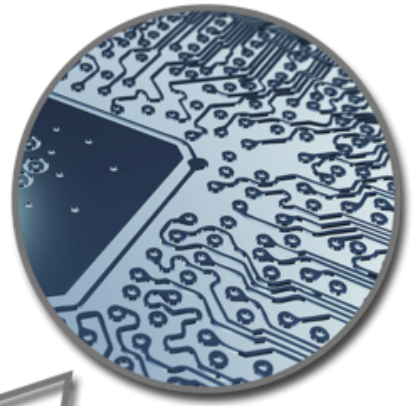

New ASIC Functionality

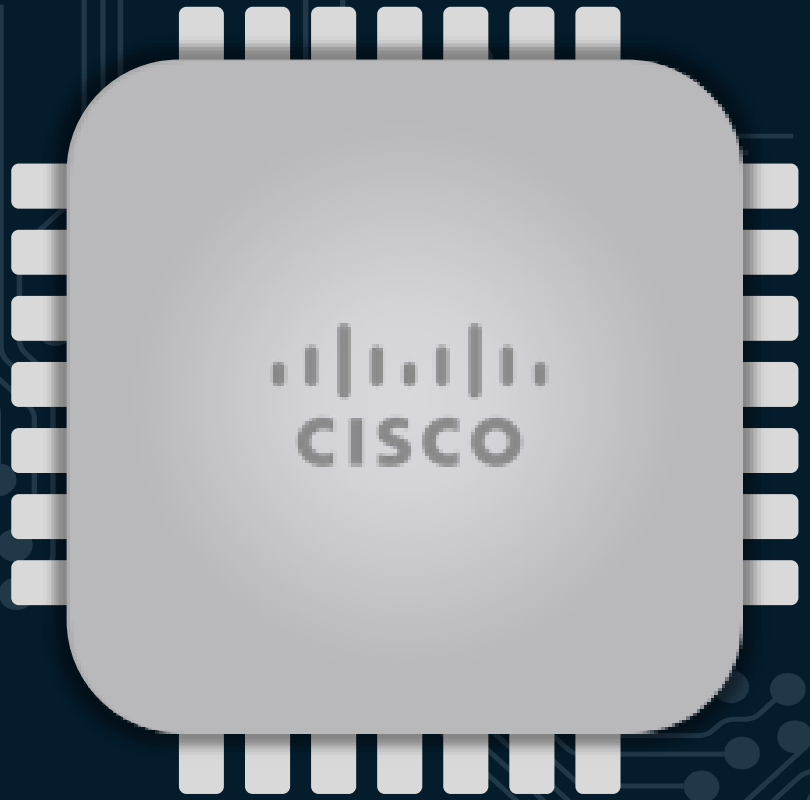

Cisco UADP

for Enterprise Switching

Cisco Unified Access Data-Plane (UADP®)

Common ASIC Architecture for Switching Access, Distribution & Core

UADP 2.0m

120 Gbps
16nm FinFET
1.3B Transistors
1 Core + ARM CPU

UADP 2.0/XL

240 Gbps
28nm FinFET
7.6B Transistors
2 Core

UADP 2.0sec

480 Gbps
16nm FinFET
7.6B Transistors
1 Core[2] + SEC

UADP 3.0

1.6 Tbps
16nm FinFET
19.2B Transistors
2 Core

UADP 3.0sec

1.6 Tbps
16nm FinFET
19.2B Transistors
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

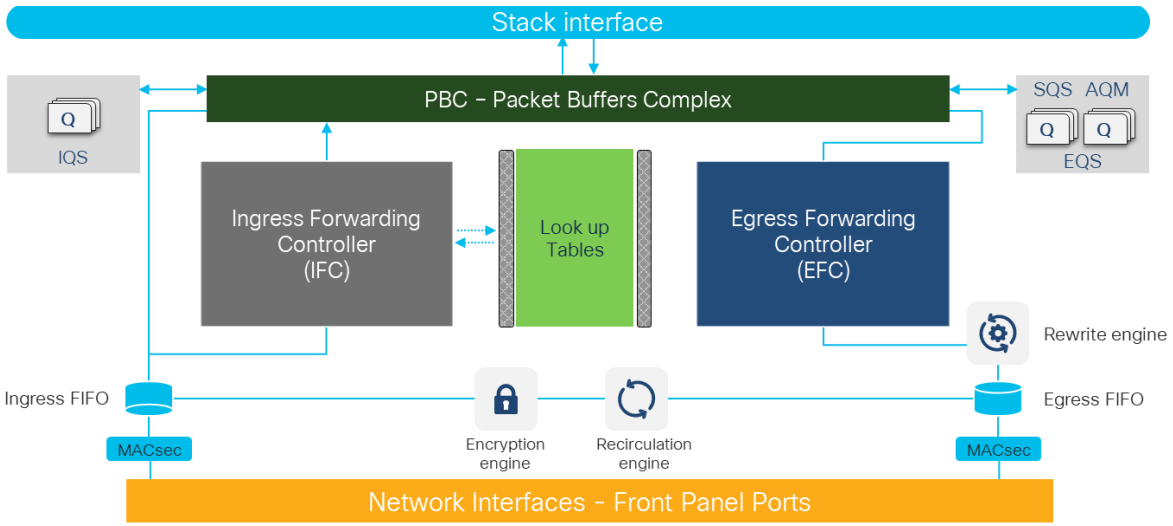

Catalyst 9500/9600 – 2019
UADP 3.0 – 19.2B transistors

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

Catalyst 9300 – 2017
UADP 2.0 – 7.5B transistors

Catalyst 3850 – 2013
UADP 1.0 – 1.3B transistors

Catalyst 3850 mGig – 2015
UADP 1.1 – 3.0B transistors

Catalyst 9200 – 2019
UADP 2.0 mini – 3.2B transistors

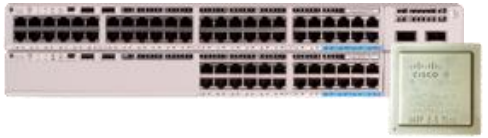

Cisco UADP ASICs

ASIC Architecture & Block Diagram

up to **1 BILLION** times per second!

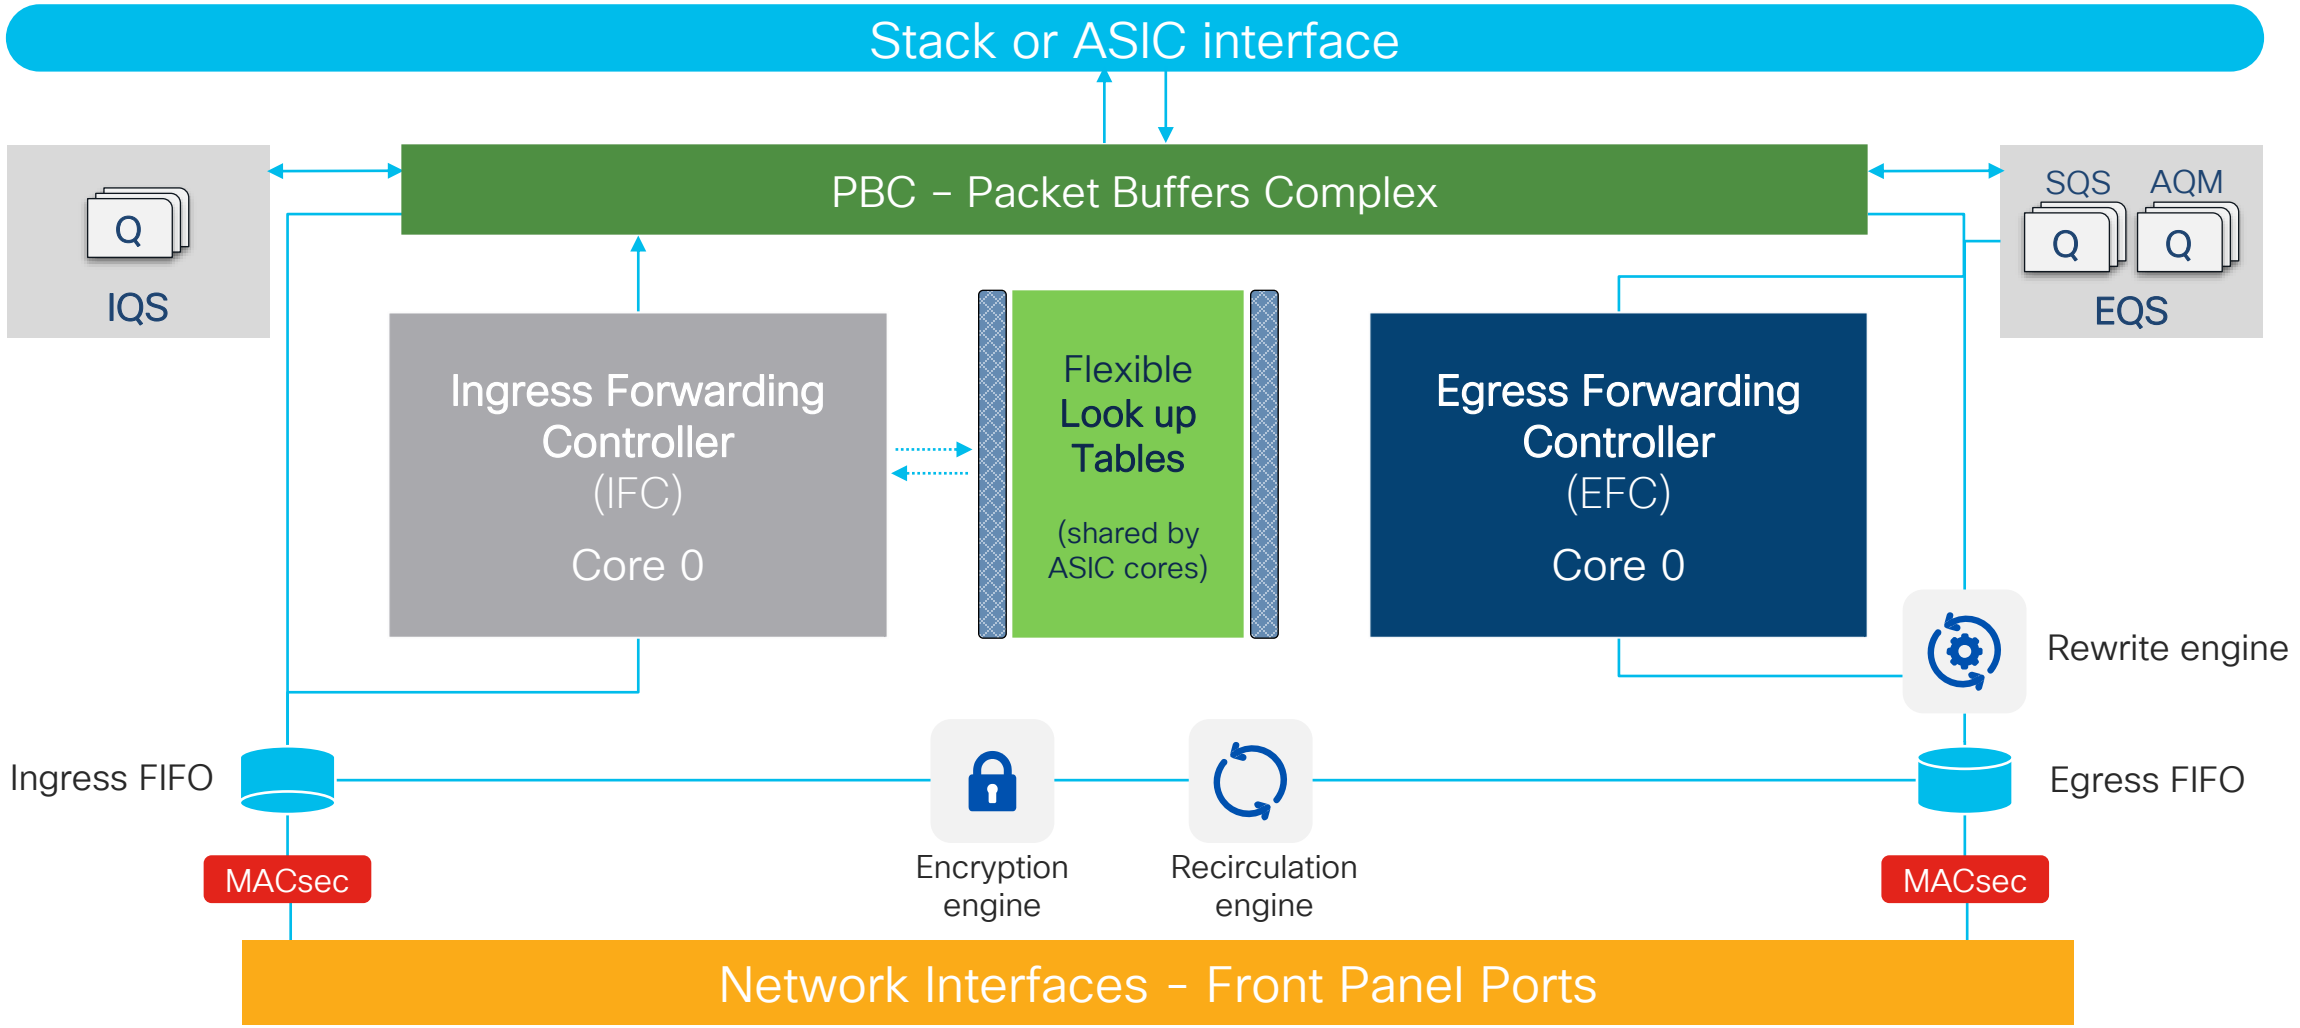

Cisco UADP – Flexible Lookups

Programmable Ingress and Egress Processing Stages

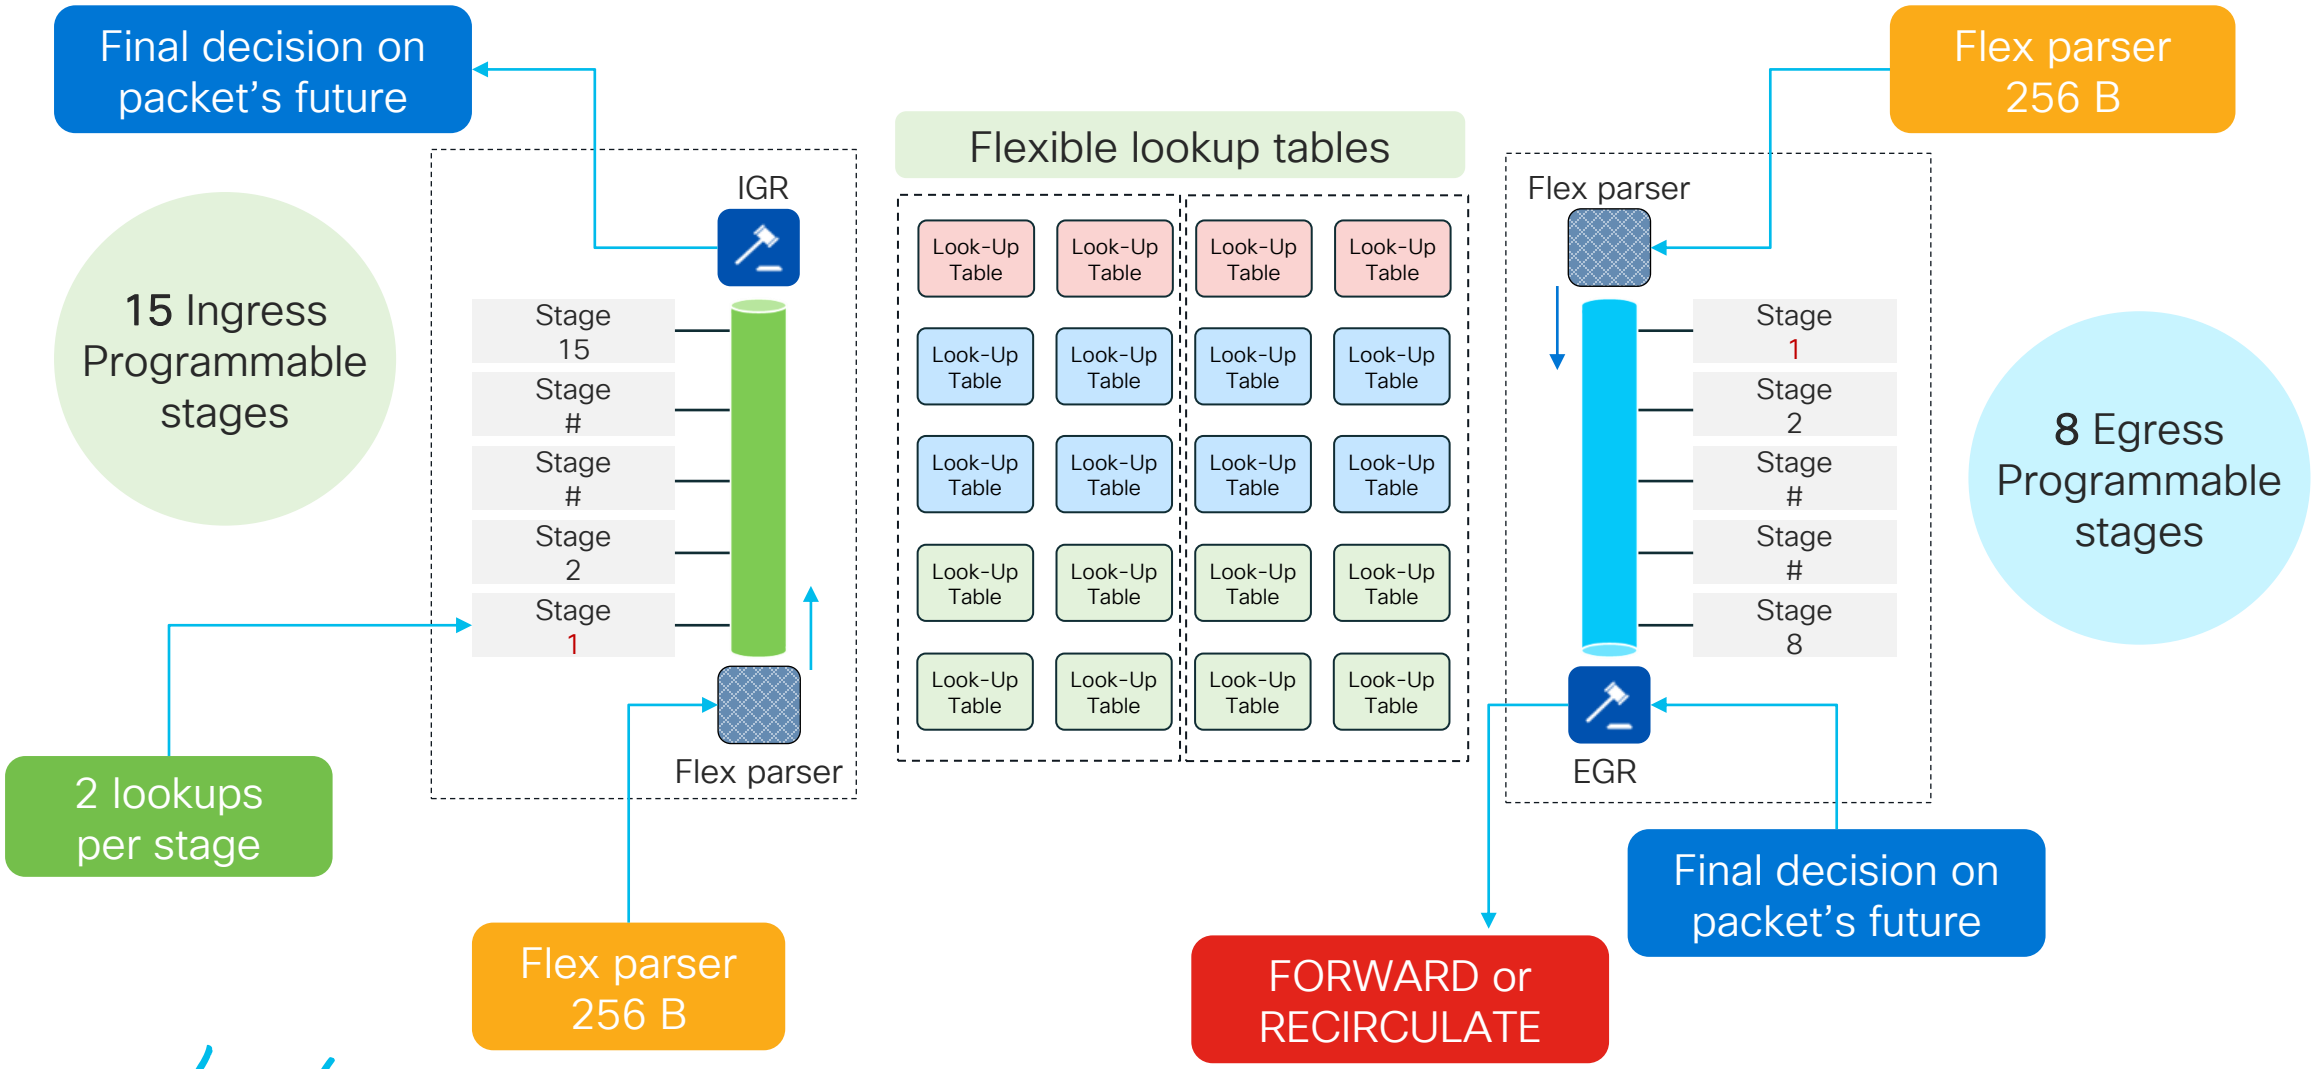

Cisco UADP – Flexible Tables

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 SDM 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

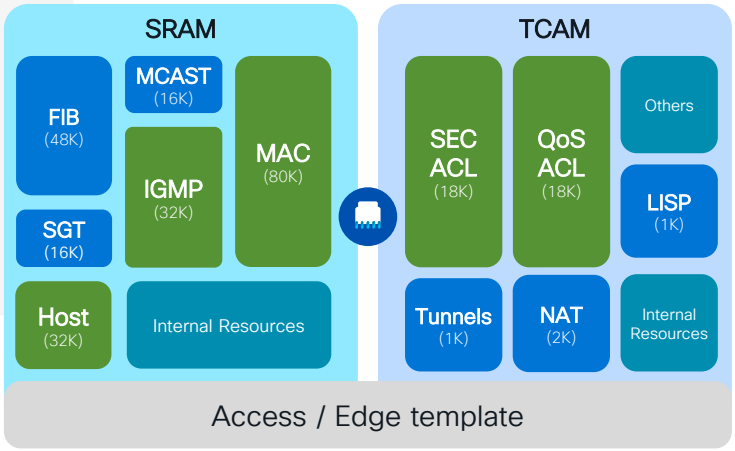

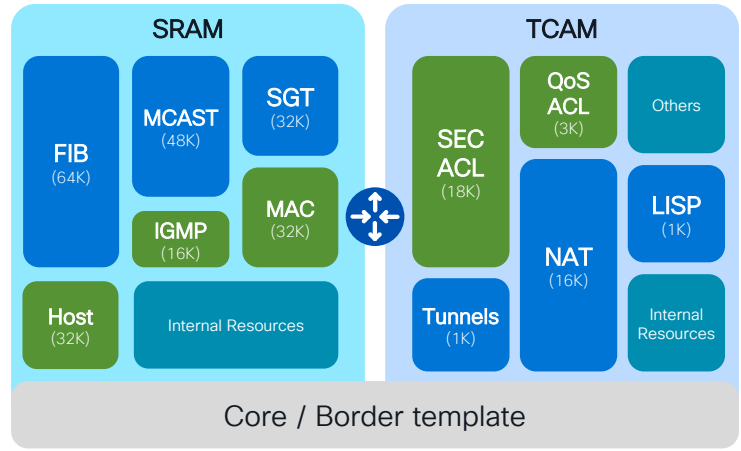

Table sizes can be tailored to support multiple templates

Cisco UADP – Packet Walks

Generic Unicast Packet Walk – Same ASIC

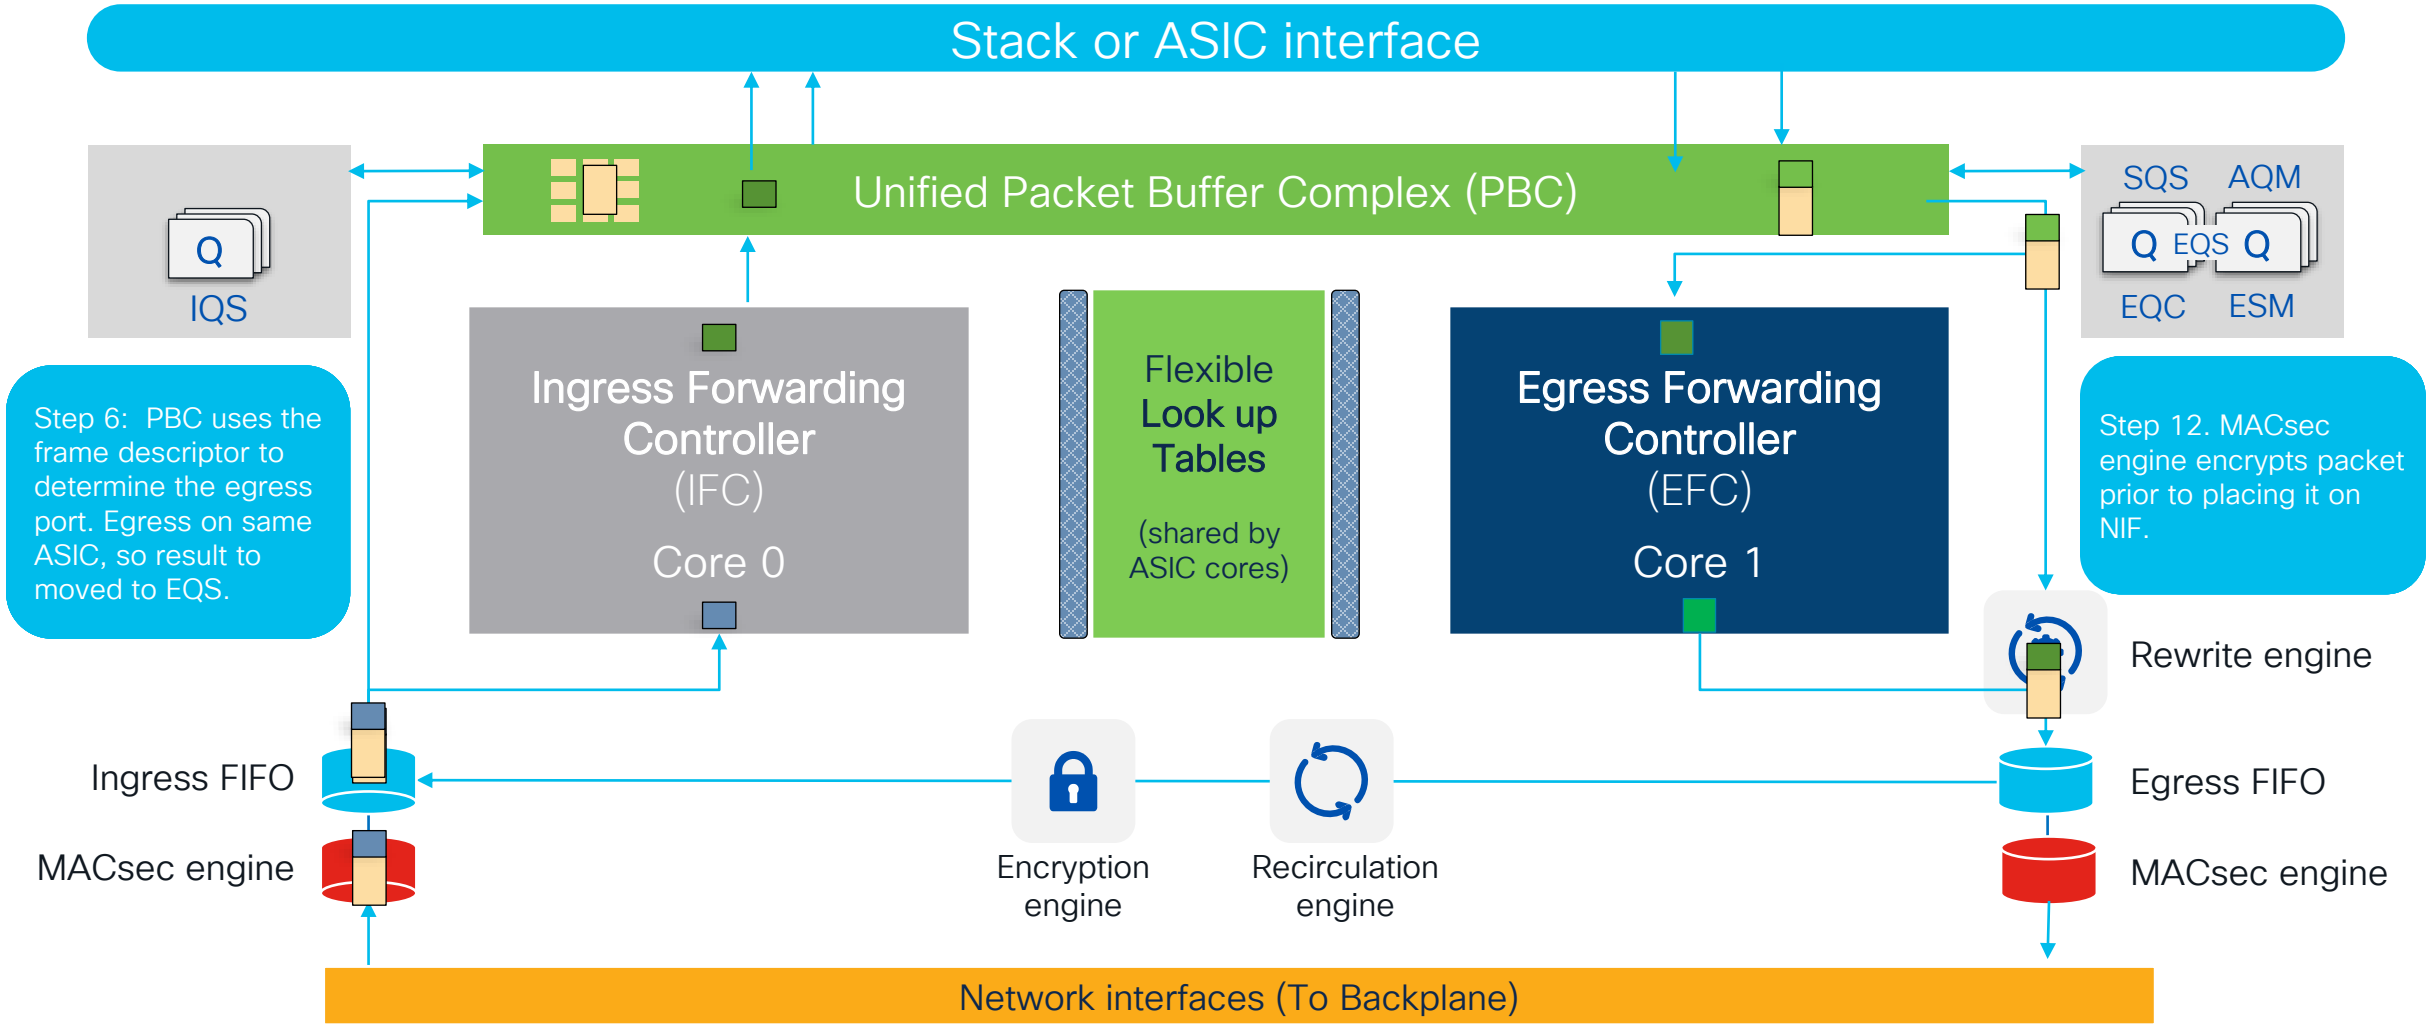

Cisco UADP – Multiple ASICs

Interconnecting Multiple ASIC Cores (Stacking or DPP)

Catalyst 9300X-48HX

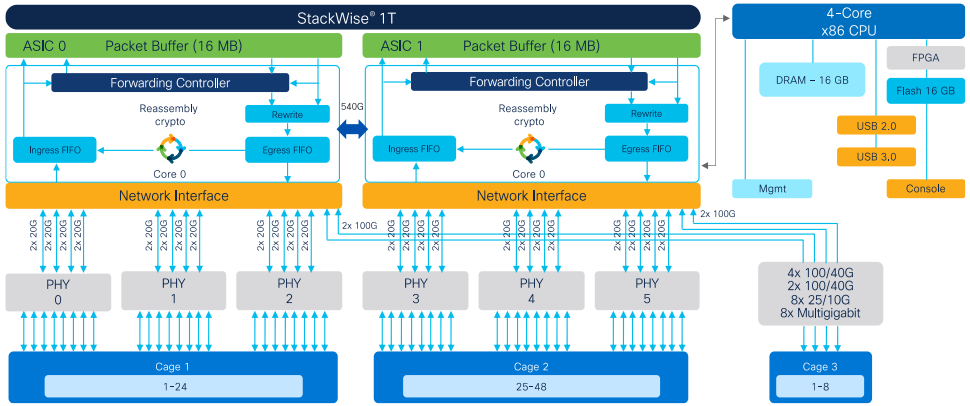

Catalyst 9400X-SUP2/XL

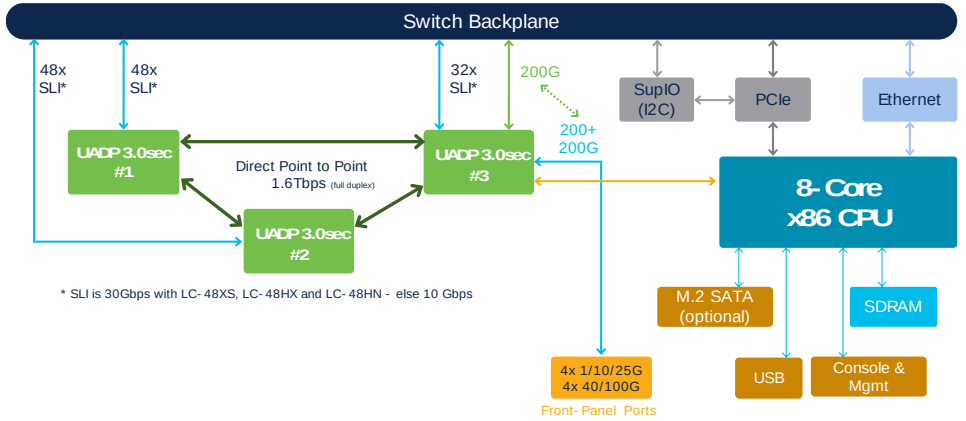

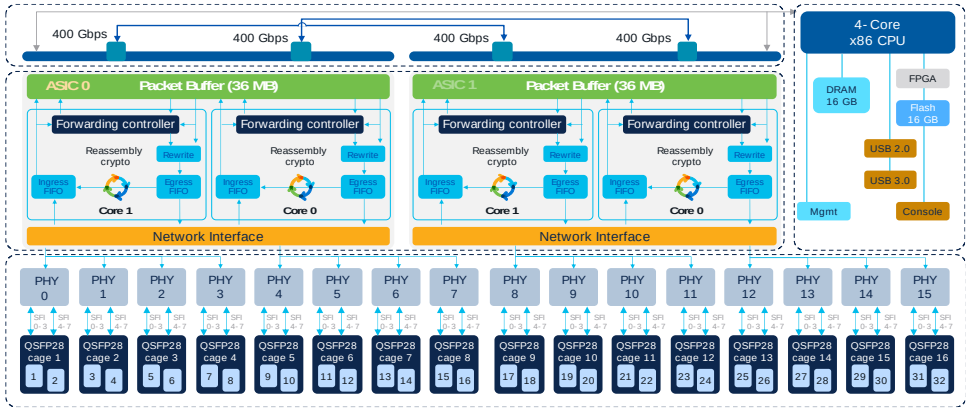

Catalyst 9500-32C

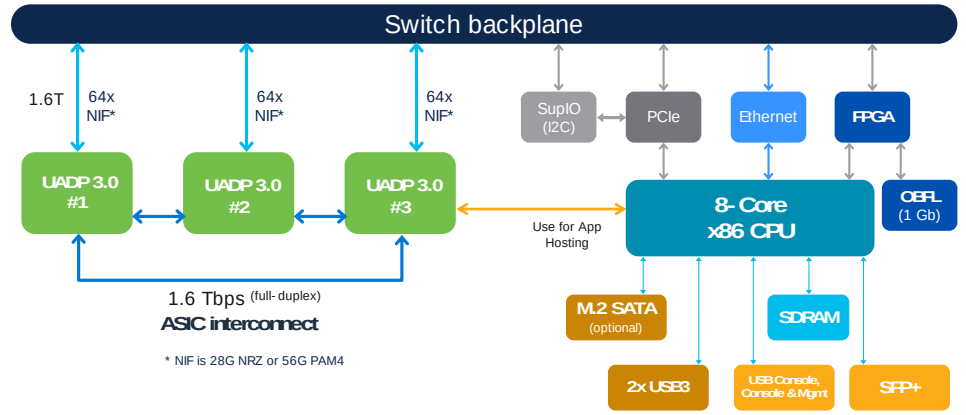

Catalyst 9600-SUP1

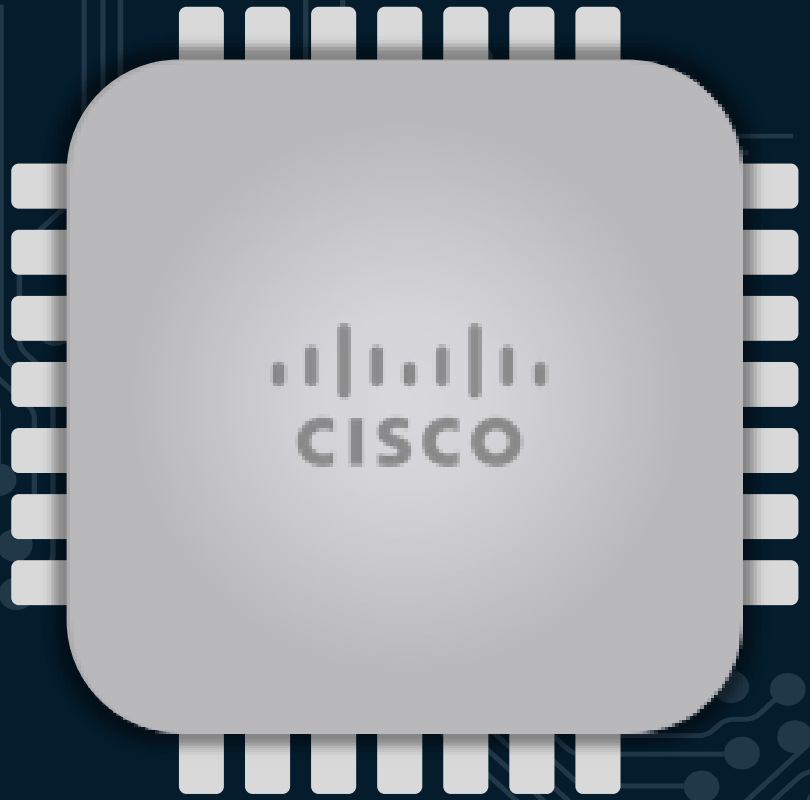

Cisco Silicon One

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

NEW

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

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

50G PAM4 Serdes

Cisco Silicon ONE Q200

Industry Leading

12.8T System on Chip

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

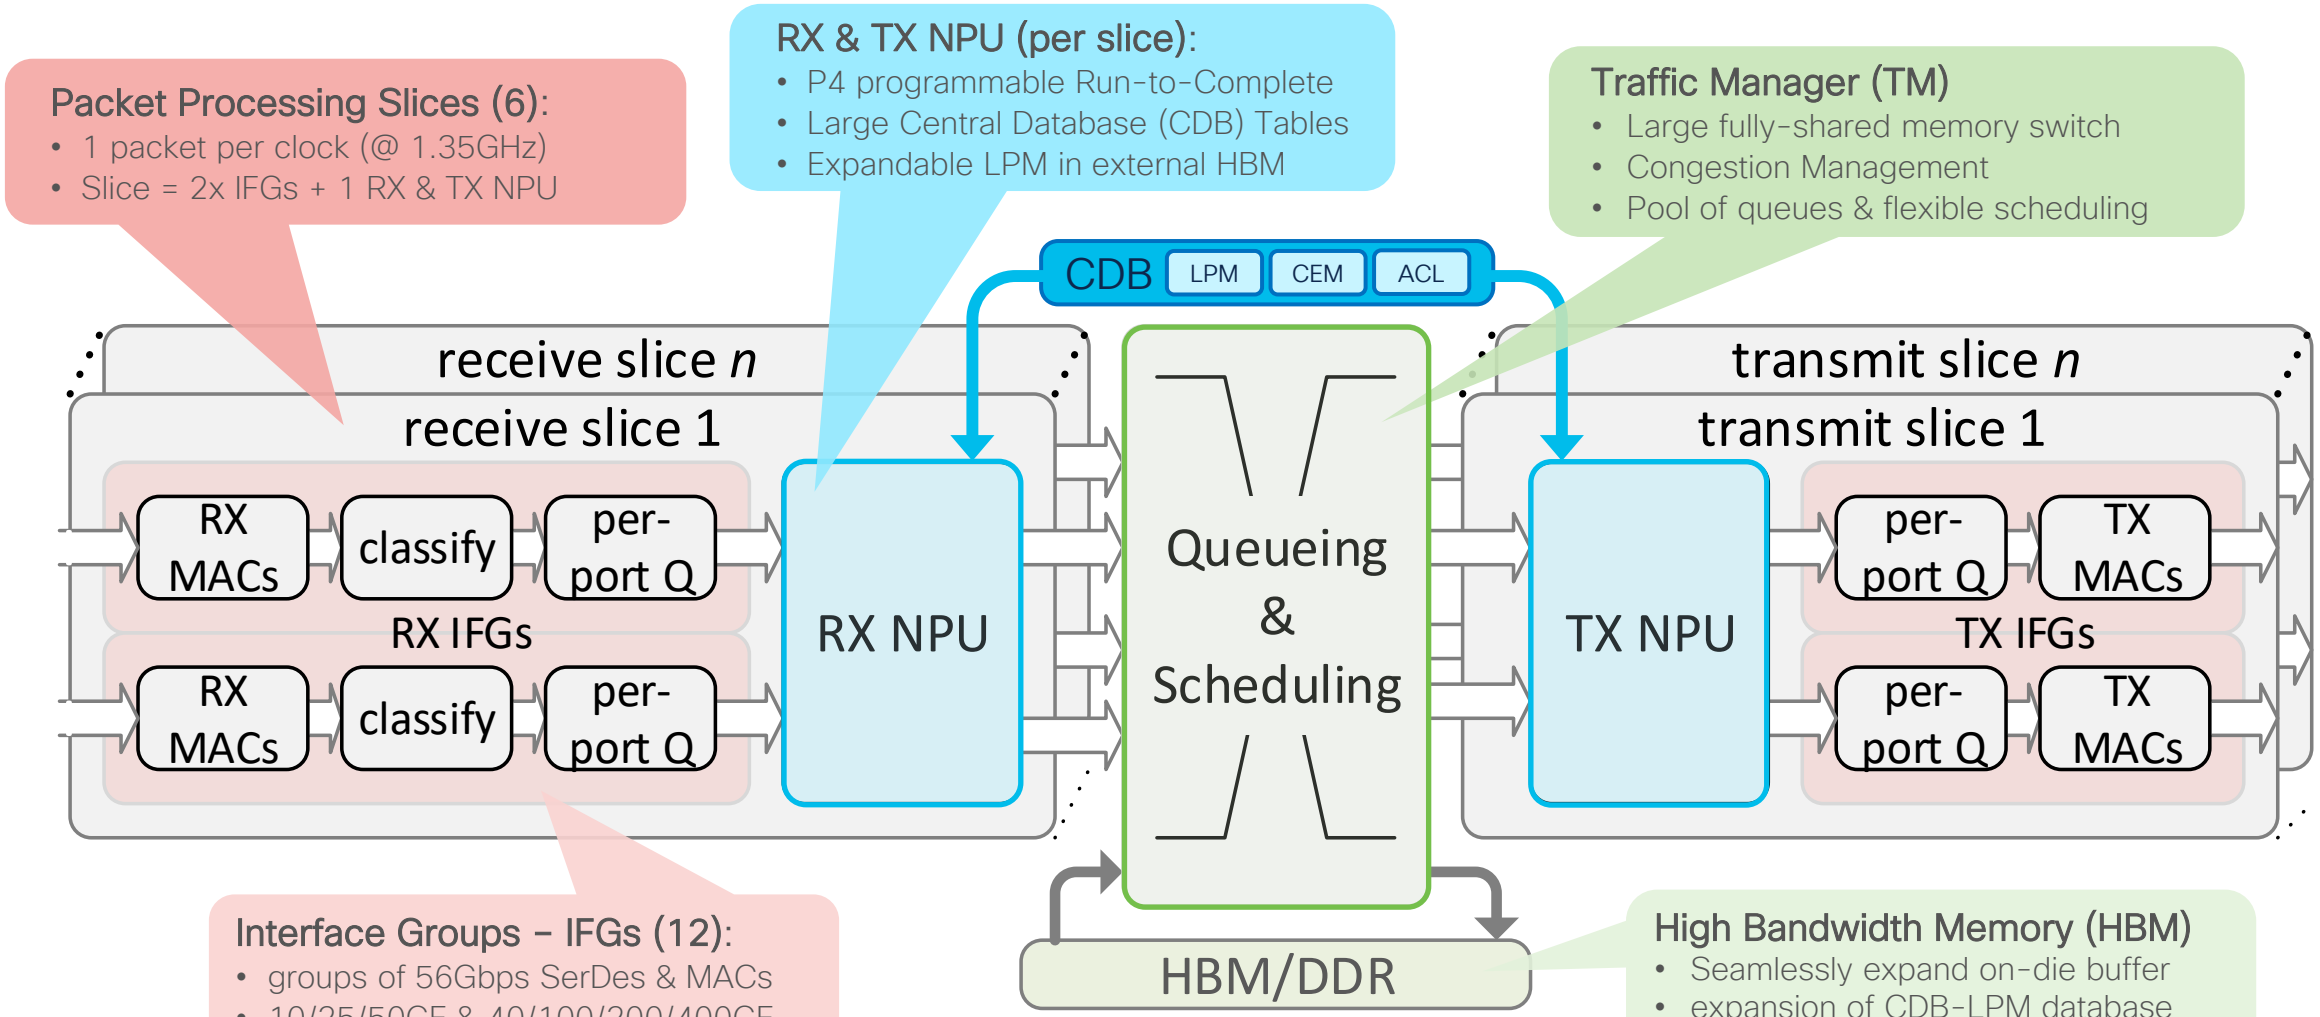

Cisco Silicon One™ Q200 – HBM

On-Package High Bandwidth Memory

- **Augments local on-die memory**
  - use local (SMS) buffer until full
  - use HBM for bursting or congestion
- **Deep buffering** (+LPM expansion) on-package at high-speed
- **2 x 4GB stacks of 2.5D** memory with wide-bus interposer = ~2.4 Tbps full duplex
- Interposer connects ASIC die to on-package HBM memory

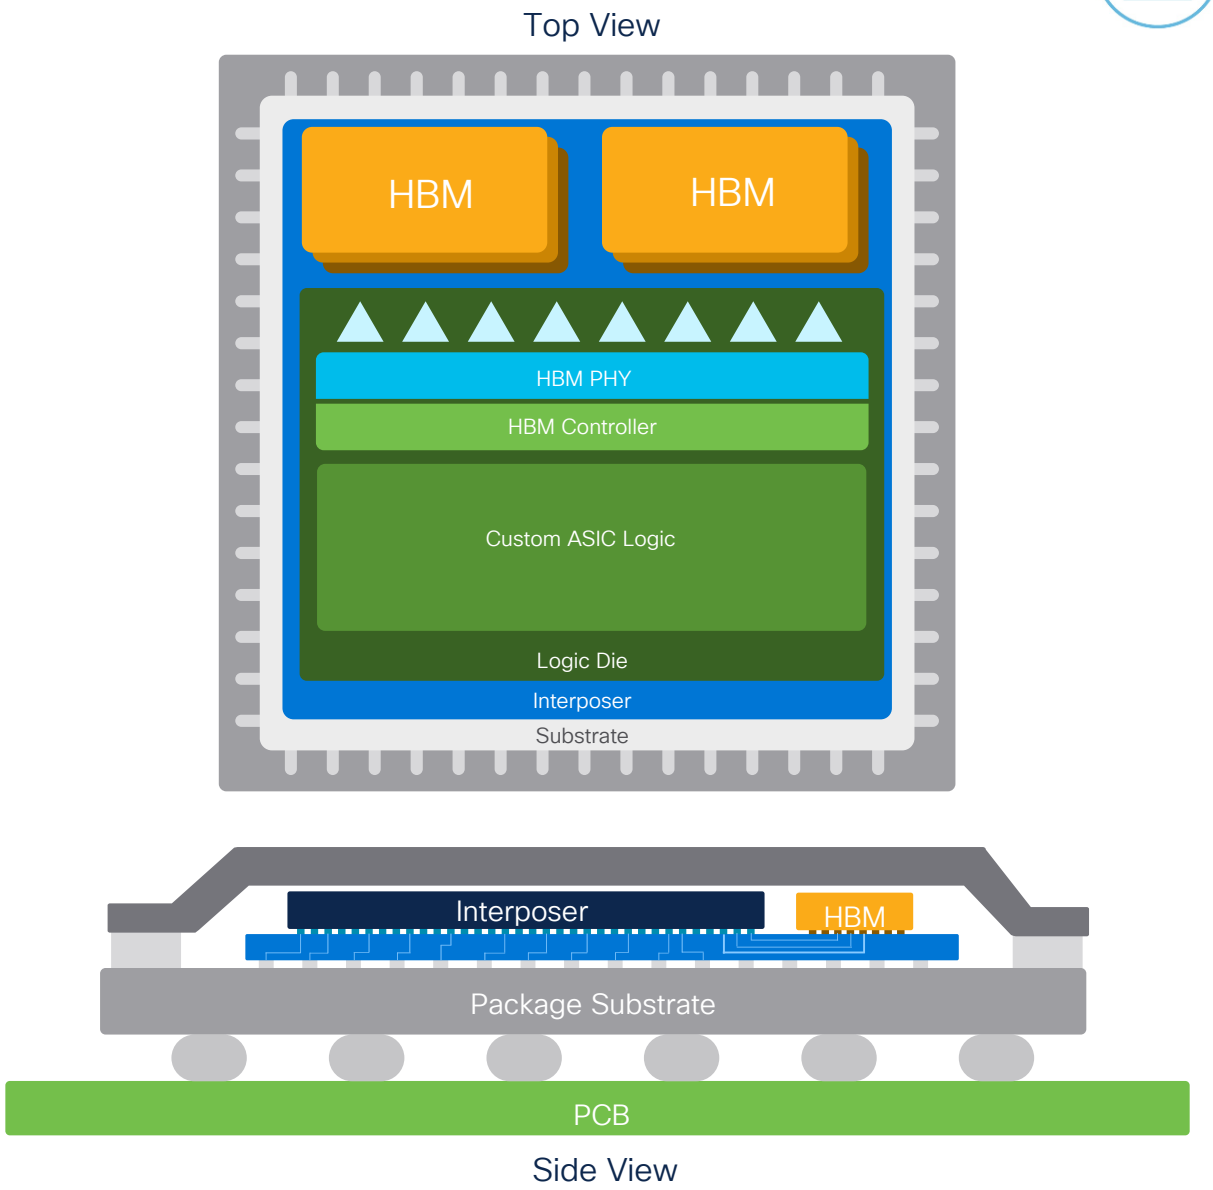

Cisco Silicon One™ Q200 – Central Databases

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

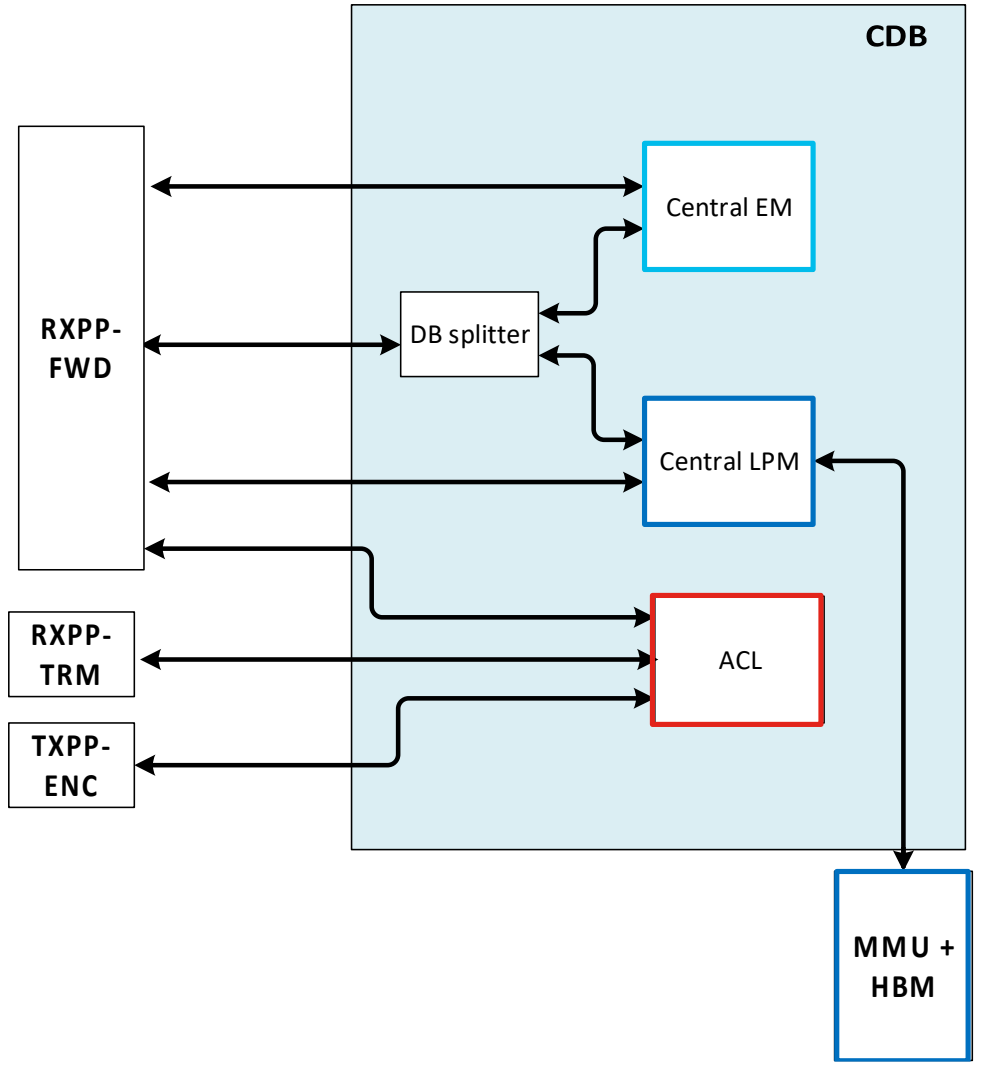

* Exact scale depends on IP/mask distribution (contiguous vs. random) and hash efficiency. Sample tests with IPv4 GRT is ~1.85M

Cisco Silicon One Q200

Generic Unicast Packet Walk

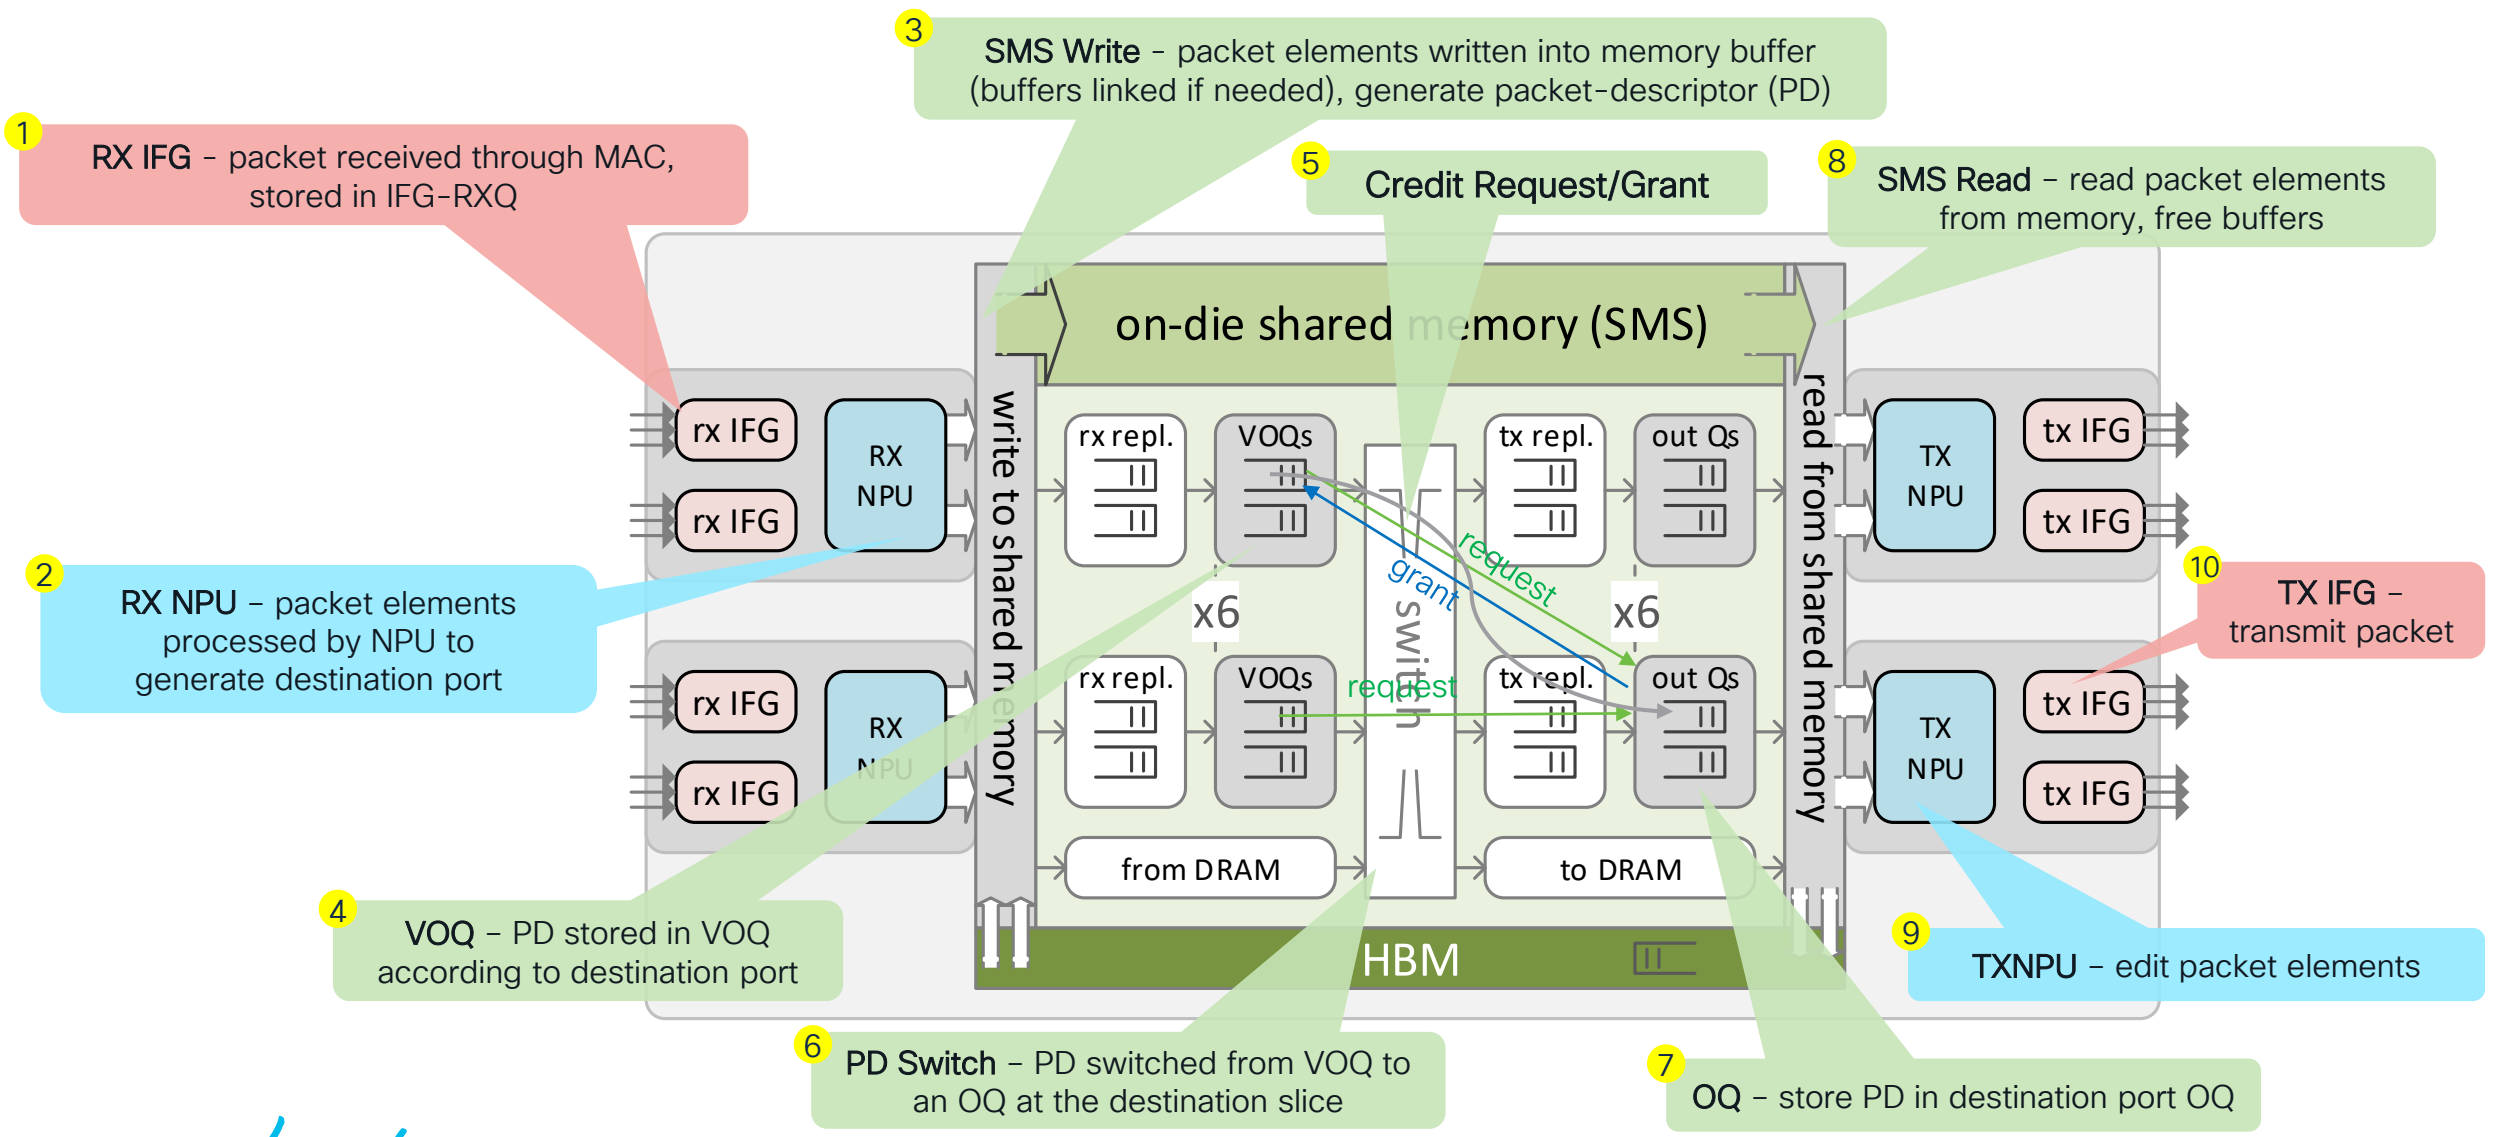

Cisco Silicon One – Q200

Combining multiple ASICs in One SOC

Catalyst 9500X-28C8D

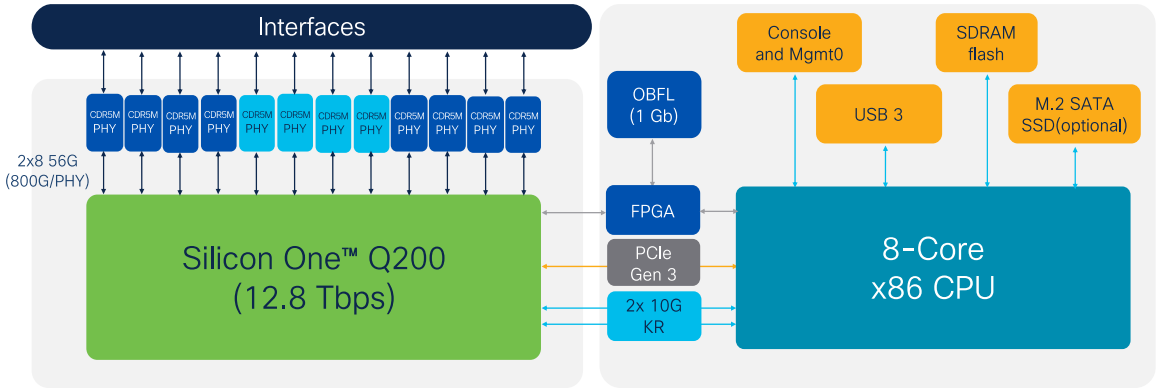

Catalyst 9600X-SUP2

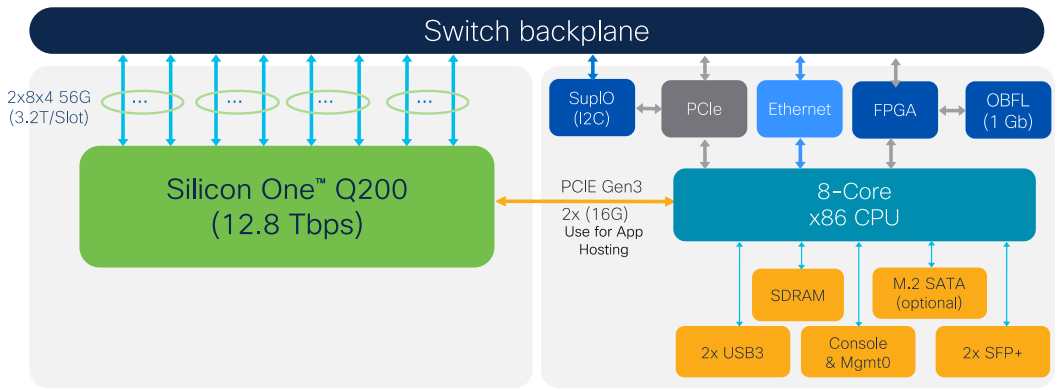

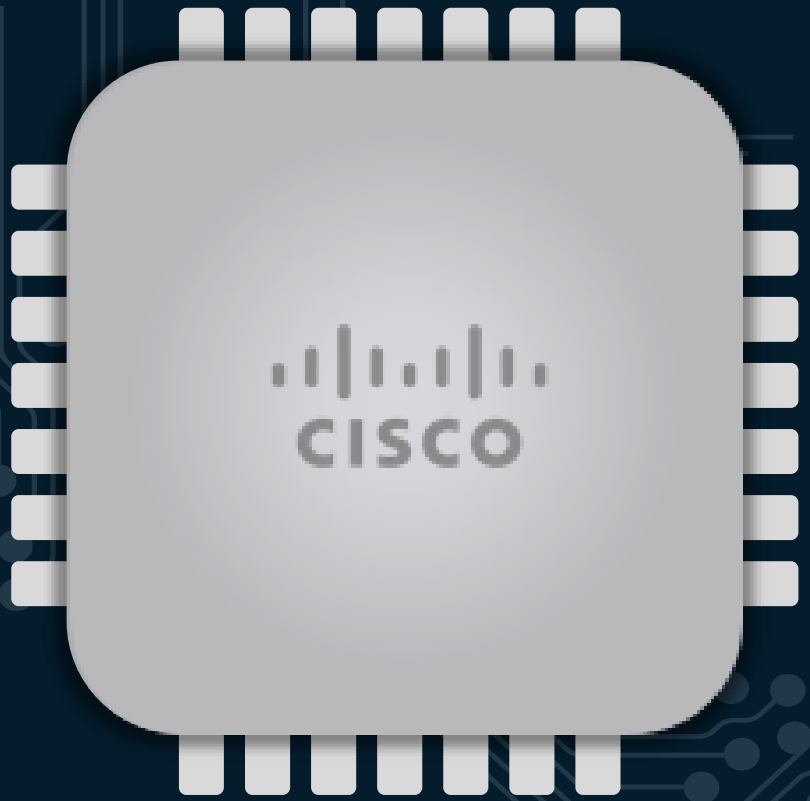

Cisco Catalyst 9000 “X” Series

Extending Enterprise Switching

Adding the “X factor” to the Catalyst 9000 switch family

Higher bandwidth & scale – powered by Cisco UADP and Silicon One ASICs

NEW

Catalyst 9300X
Powered by UADP 2.0sec

High-density MultiGigabit models with 90W UPOE+

ORDERABLE

C9300X-48HX

ORDERABLE

C9300X-48TX

ORDERABLE

C9300X-24Y

C9300X-12Y

C9300X-48HXN

C9300X-24HX

1T Stacking, 100G Uplinks, IPsec, Enhanced App Hosting

NEW

Catalyst 9400X

Supervisor-2/XL
Powered by UADP 3.0sec

ORDERABLE

480 Gbps per Slot

4.8T, 100G uplinks
Enhanced Scale
IPsec Capable

MultiGigabit line card

C9400-LC-48HX
48x 10G MultiGigabit, 90W UPOE+

Fiber line card

C9400-LC-48XS
48x10G SFP+

NEW

Catalyst 9500X
Powered by Silicon One Q200

High-density QSFP model with 400G

ORDERABLE

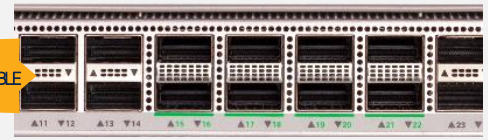

100/400G SKU

6T, 400G uplinks, Deep Buffers, Internet Routing Scale

NEW

Catalyst 9600X

Supervisor 2
Powered by Silicon One Q200

ORDERABLE

3.2 Tbps per Slot

12.8T, 400G, Deep Buffers, Internet Routing Scale

400G combo line card

ORDERABLE

C9600-LC-40YL4CD
Combo SFP + QSFP line card
compatible with SUP1 and SUP2

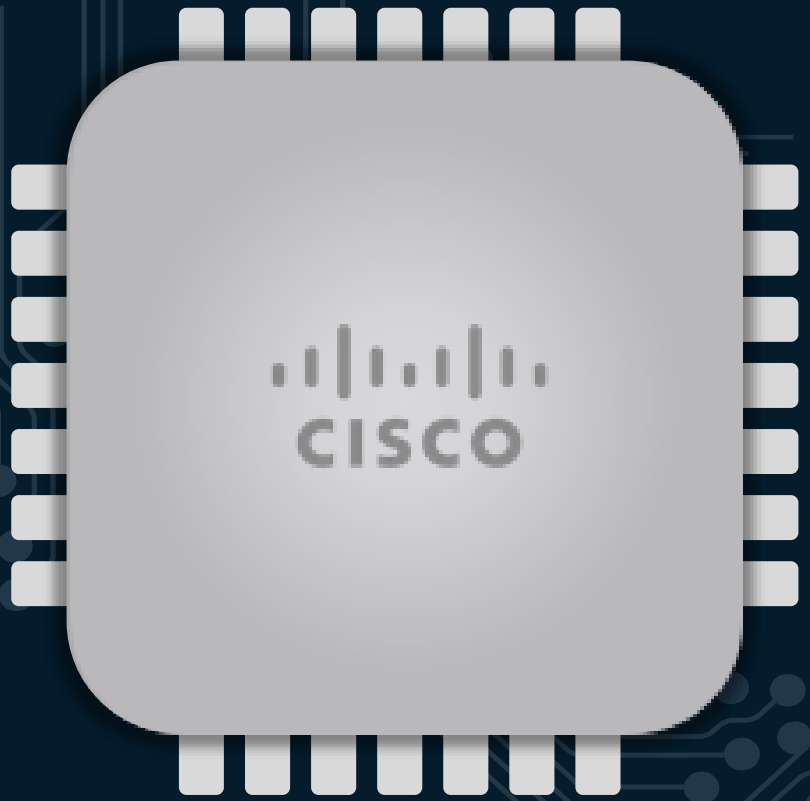

Glimpse into the Future

for Enterprise Switching

Where are things going?

Speeds and Feeds

More details in PPT Notes

25/50G → 100G SFP

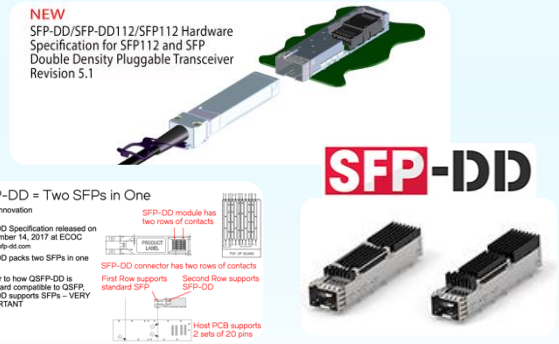

SFP-DD(112): 100Gbit/s using PAM4

100/400G → 800G QSFP

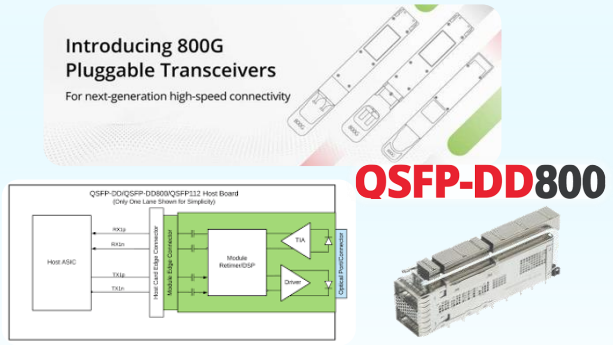

QSFP-DD(112): 800Gbit/s using PAM4

Coherent Optics

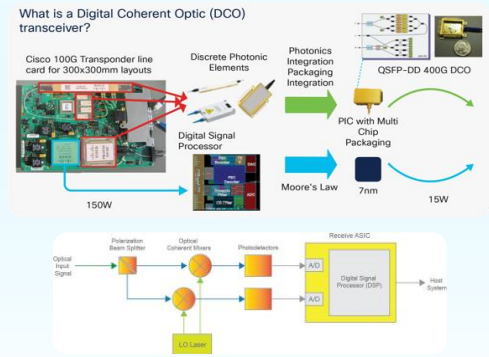

Uses phase of optical signal vs. power level

Direct-Drive (PHYLESS)

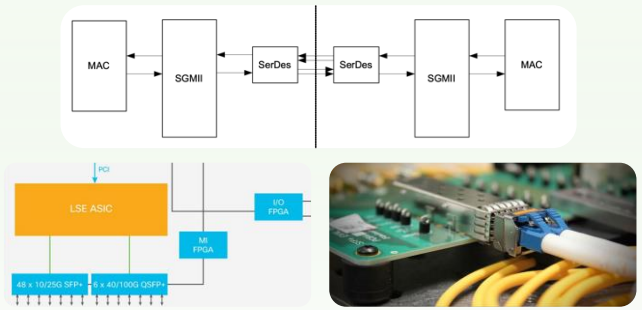

Transceiver uses SFI/SGMII to connect directly to the ASIC SerDes (no PHY)

4x100G & 8x50G Breakout

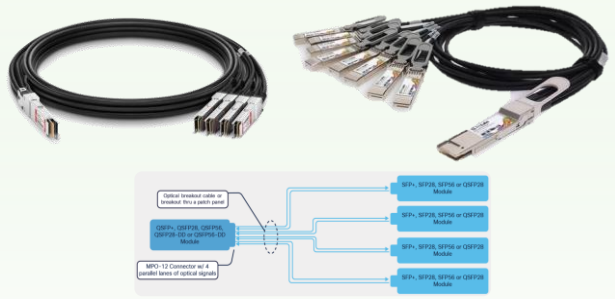

QSFP-DD (400G) can be split into 4 x 100G QSFP28 or 8 x 50G SFP56

Single Lambda

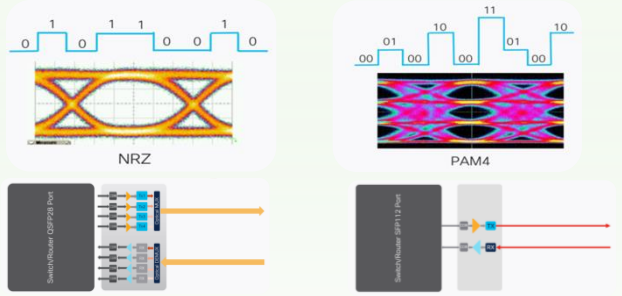

Uses a single 112G PAM4 signal vs. 4 x 28G NRZ signals

Where are things going?

Features & Scale

More details in PPT Notes

MAC & IPv4/v6 Scale

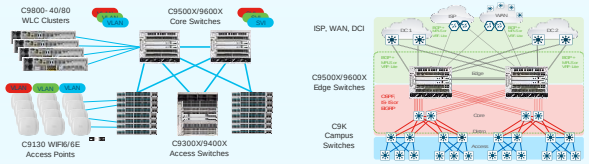

Increasing MAC scale for Wireless & IOT

- 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

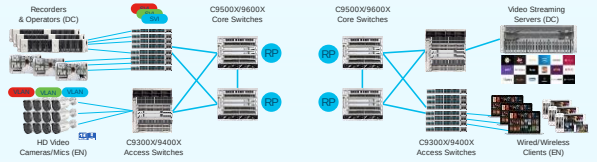

Higher L2 & L3 Multicast scale in Hardware

- 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

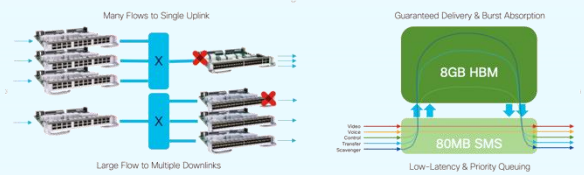

Virtual Output Queues (VoQ) for high throughput

- 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

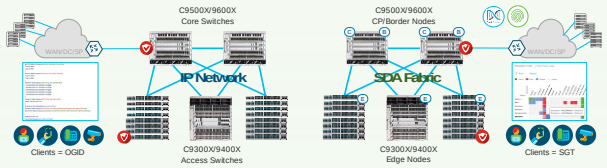

Object-Groups map IP/mask to Labels in CEM

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

OGACL ACEs take minimal space in ACL TCAM

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

WAN MACsec & IPsec

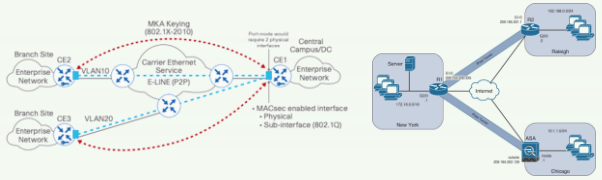

256-bit Hardware Encryption over L2

- 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

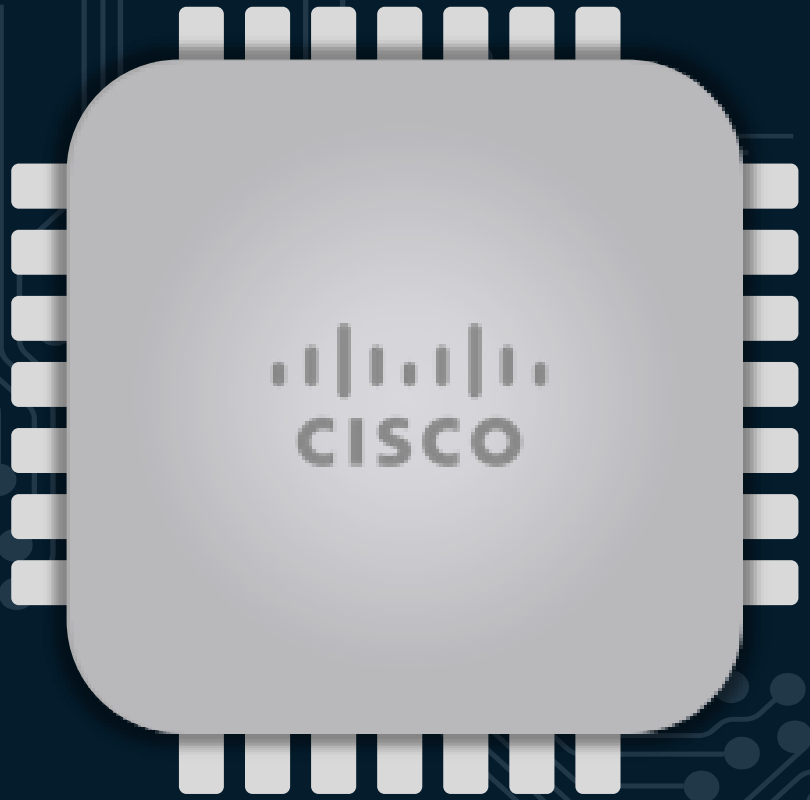

Summary

Wrap-Up

Cisco **Programmable Hardware** equals

FLEXIBILITY & ADOPTABILITY

Enabling **Network Evolution** on the journey to **Intent-Based Networking**

How Did We Do?

Inventing the Future
of Networking

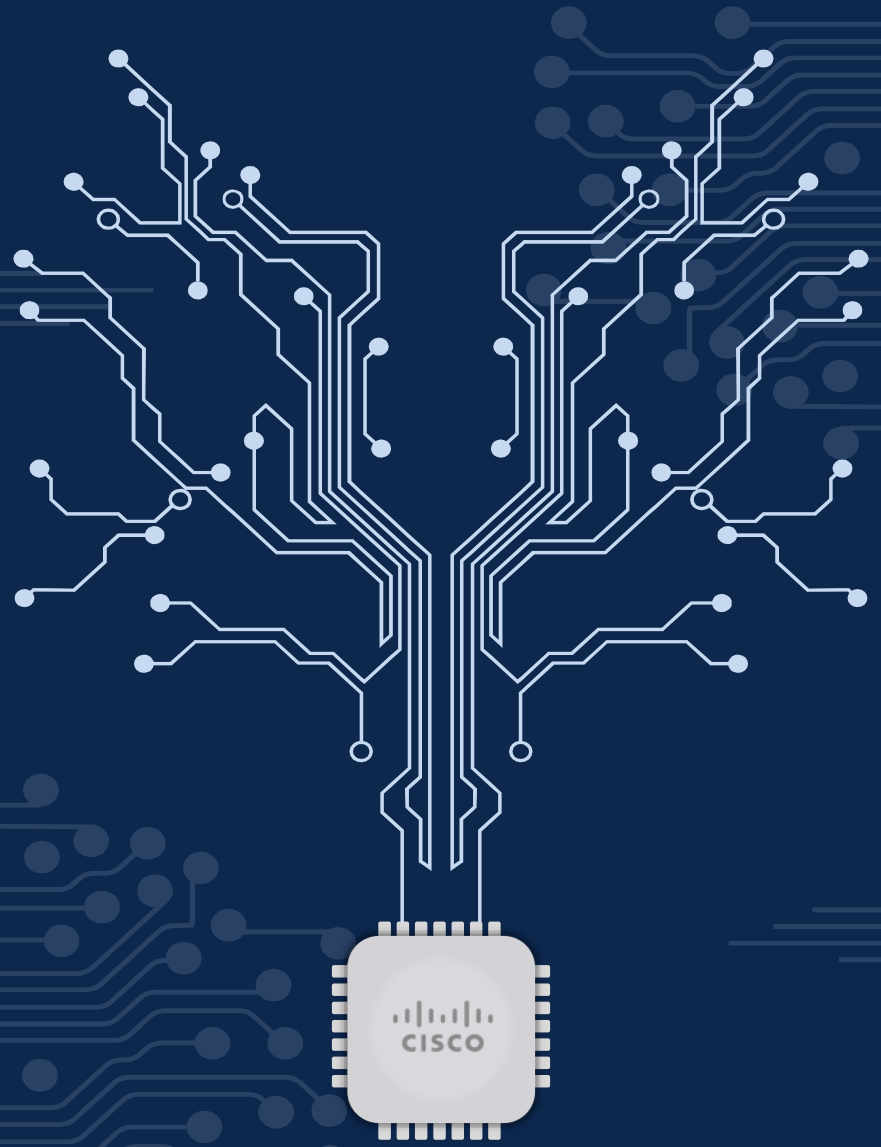

Do You Have a Better Understanding ...

... of why we need Programmable ASICs in Networking?

... of why ASIC and Hardware innovation are important?

... and how you can leverage these ASIC innovations in your own network?

Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.

Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. www.cisco.com/go/certs

Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.

Learn

Cisco U.
IT learning hub that guides teams and learners toward their goals

Cisco Digital Learning
Subscription-based product, technology, and certification training

Cisco Modeling Labs
Network simulation platform for design, testing, and troubleshooting

Cisco Learning Network
Resource community portal for certifications and learning

Train

Cisco Training Bootcamps
Intensive team & individual automation and technology training programs

Cisco Learning Partner Program
Authorized training partners supporting Cisco technology and career certifications

Cisco Instructor-led and Virtual Instructor-led training
Accelerated curriculum of product, technology, and certification courses

Certify

Cisco Certifications and Specialist Certifications
Award-winning certification program empowers students and IT Professionals to advance their technical careers

Cisco Guided Study Groups
180-day certification prep program with learning and support

Cisco Continuing Education Program
Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

cisco Live!

Thank you