You make **possible**

Catalyst 9400 Architecture

Kenny Lei
Technical Marketing Engineer
BRKARC-3873

Session Goal

- To provide a thorough understanding of the Catalyst 9400 hardware architecture, packet flows and key forwarding components

- This session will also examine Catalyst 9400 software architecture, ACL and QoS

- No discussion on other Catalyst platforms

Agenda

- Overview

- Architecture

- Forwarding

- Features (ACL, QoS, Security, High Availability)

- Conclusion

Overview

You make networking **possible**

Cisco live!

Cisco Catalyst 9000 Family

Cisco Catalyst
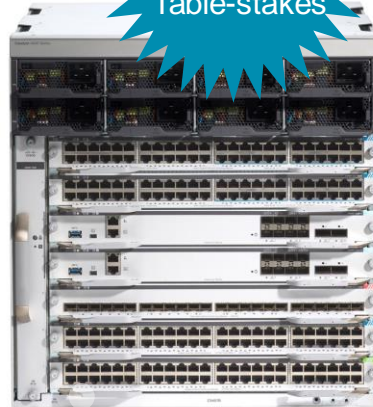9200 Series

Cisco Catalyst
9300 Series

Cisco Catalyst
9400 Series

Cisco Catalyst
9500 Series

Cisco Catalyst
9600 Series

Access switching

Core switching

IOS-XE
Common Software Architecture

UADP ASIC
Common Hardware Architecture

*C9300, C9400, C9500 and C9600 run the same binary IOS-XE image

Catalyst 9400

480G BW per slot

Redundancy is now Table-stakes

5KW PoE Per slot

4-Slot

7-Slot

10-Slot

Hardware

You make networking **possible**

Chassis
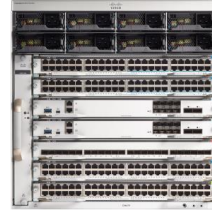
	4 – Slot	7 - Slot	10 - Slot
Supervisor	2 (Redundant)		
Line Cards	2	5	8
Ports	96x 10/100/1000 48 mGig; 56 SFP/SFP+ 2x QSFP+	240x 10/100/1000 120 mGig; 128 SFP/SFP+ 2x QSFP+	384x 10/100/1000 192 mGig; 200 SFP/SFP+ 2x QSFP+
Dimension	W:17.5"; D:16.25"; H:6RU	W:17.5"; D:16.25"; H:10RU	W: 17.5"; D:16.25"; H: 13RU
BW per LC Slot	480G	480G	480G
BW between Sup Slots	720G		
Power Supply	4 PS (N+1 and N+N)	8 PS (N+1 and N+N)	8 PS (N+1 and N+N)
PoE per slot	4,800W		
Cooling	Side to Side (Front-to-Back for PS)		

High Density 10G Ports, 100G Uplinks

Ready for future higher power PoE devices

Blue Beacon

- Chassis/Fan tray, Supervisor, Line cards, and PS have the beacon

- Identify the hardware during configuration and trouble shooting

- Turn on and off by using software control (No button except FanTray)
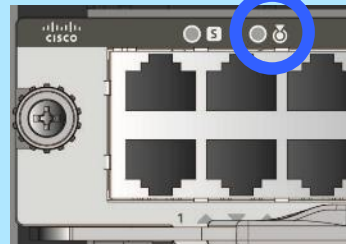
Chassis/FanTray	Supervisor	Line Card	PS

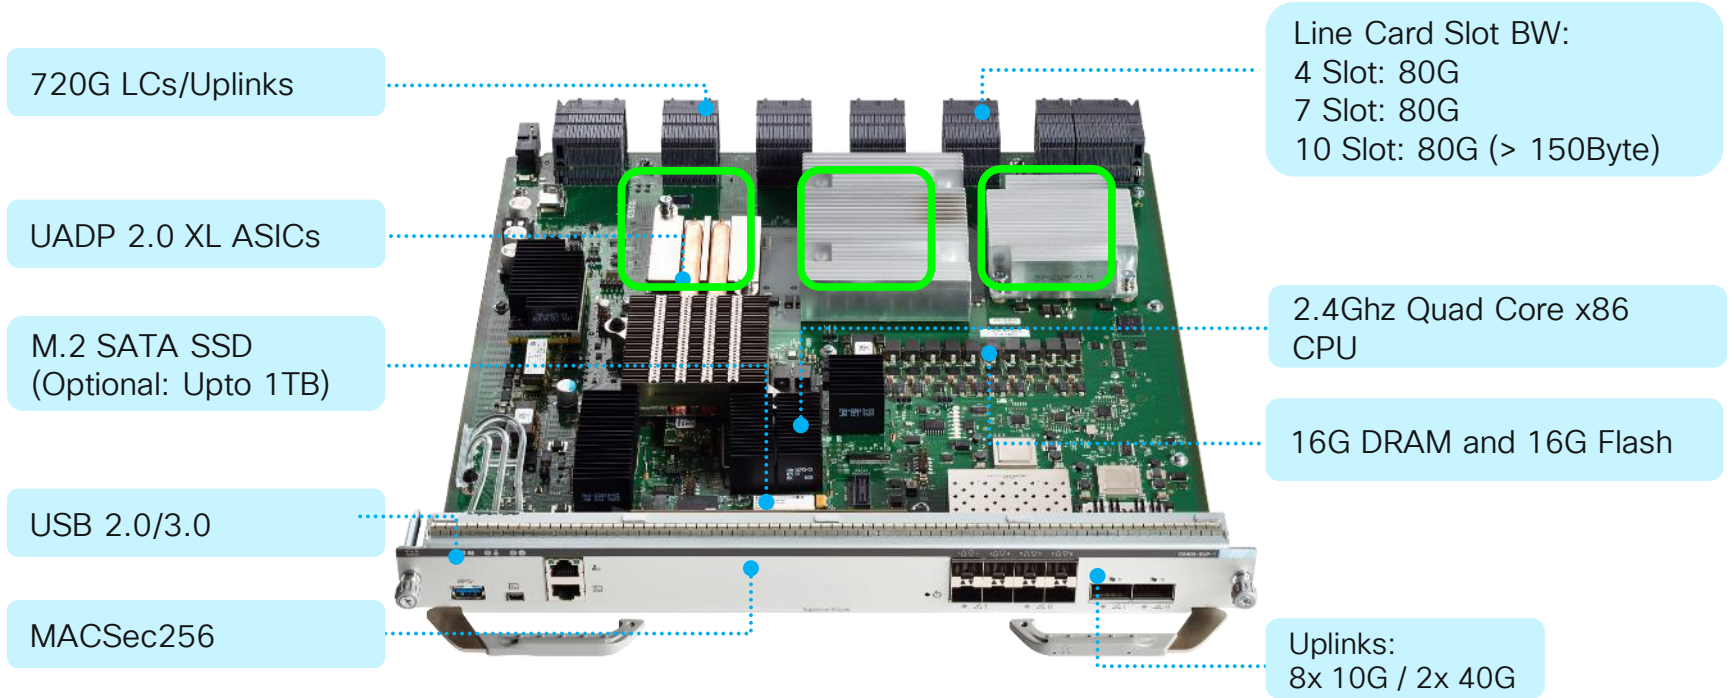# RFID

RFID on Every FRUable Components of Catalyst 9400

Inventory Management (Tracking) has never been Easier

C9400-SUP-1 - Overview

720G LCs/Uplinks

UADP 2.0 XL ASICs

M.2 SATA SSD
(Optional: Upto 1TB)

USB 2.0/3.0

MACSec256

Line Card Slot BW:
4 Slot: 80G
7 Slot: 80G
10 Slot: 80G (> 150Byte)

2.4Ghz Quad Core x86 CPU

16G DRAM and 16G Flash

Uplinks:
8x 10G / 2x 40G

C9400-SUP-1XL - Overview

720G LCs/Uplinks

UADP 2.0 XL ASICs

M.2 SATA SSD
(Optional: Upto 1TB)

USB 2.0/3.0

MACSec256

Line Card Slot BW:
4 Slot: **240G** (>150Byte)
7 Slot: **120G** (>150Byte)
10 Slot: 80G (> 150Byte)

Templates: Distribution,
Core, SD Boarder, NAT

16G DRAM and 16G Flash

2.4Ghz Quad Core x86
CPU

Uplinks:
8x 10G / 2x 40G

Optimized for Distribution/Core deployment

C9400-SUP-1XL-Y – Overview

720G LCs/Uplinks

UADP 2.0 XL ASICs

M.2 SATA SSD
(Optional: Upto 1TB)

USB 2.0/3.0

MACSec256

Line Card Slot BW:
4 Slot: **240G** (>150Byte)
7 Slot: **120G** (>150Byte)
10 Slot: 80G (> 150Byte)

Templates: Distribution,
Core, SD Boarder, NAT

16G DRAM and 16G Flash

2.4Ghz Quad Core x86
CPU

Uplinks:
2x 25G / 8x 10G / 2x 40G
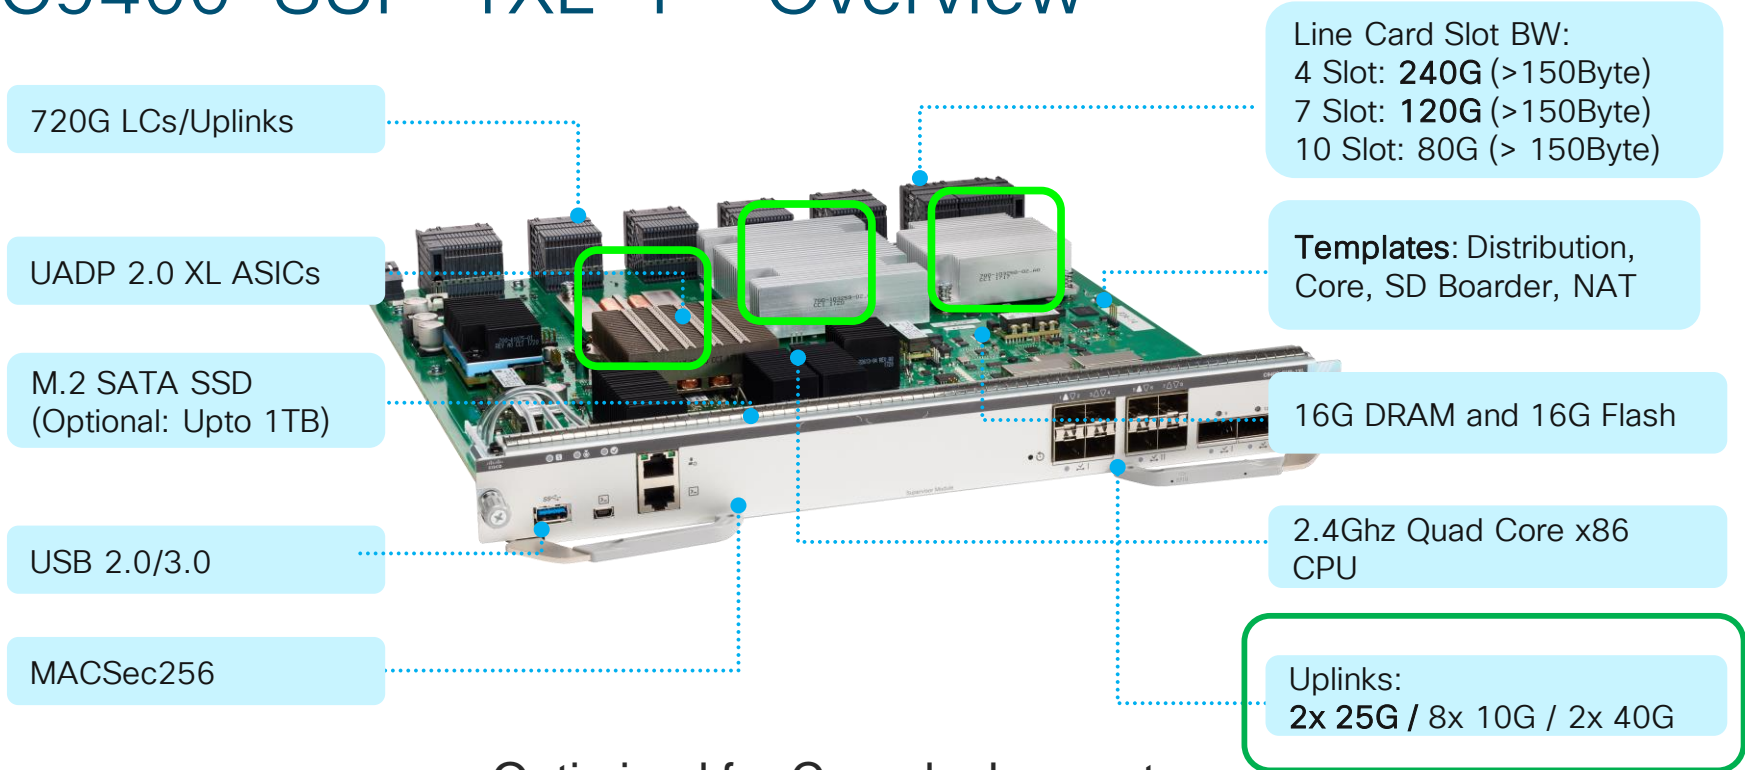
Optimized for Core deployment

Supervisors

Line Card Slot BW:	Template:	Uplinks:

C9400-SUP-1

4 Slot: 80G
7 Slot: 80G
10 Slot: 80G

Access

1G
10G
40G

C9400-SUP-1XL

4 Slot: **240G**
7 Slot: **120G**
10 Slot: **80G**

Access,
**Core,
SD Boarder,
NAT**

C9400-SUP-1XL-Y
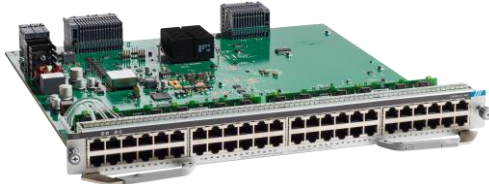
1G
10G
25G
40G
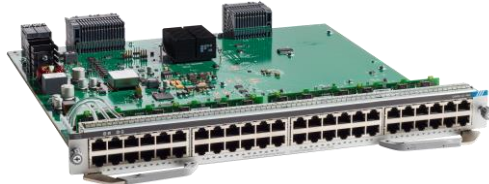
Cisco *live!*

Line Cards - Copper

C9400-LC-48T

RJ45 (Data)
48x 10/100/1000
TrustSec and MACSec(256)

C9400-LC-48P
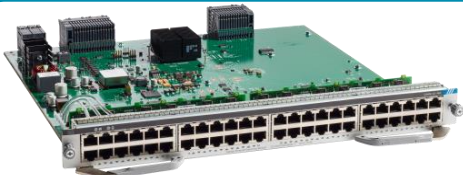

C9400-LC-48U

RJ45 (UPoE)
48x 10/100/1000
PoE/PoE+; PoE/PoE+/UPoE
TrustSec and MACSec(256)

C9400-LC-48UX

RJ45 (mGig)
24x 10/100/1000 + 24x
100/1G/2.5G/5G/10G
PoE/PoE+/UPoE
TrustSec and MACSec(256)

Line Cards – Fiber

C9400-LC-24S

C9400-LC-48S

SFP (1G)
24x or 48x 100/1000
TrustSec and MACSec(256)

C9400-LC-24XS

Fiber (1G/10G)
24x 1G/10G
TrustSec and MACsec(256)

Power Supplies

- Modular Design: 4 PS for 4 slot chassis; 8 PS for 7 and 10 slot chassis

- Shared: Power for both Data and Inline Power

- Platinum PS: 90%+ efficiency

- PS:
  - 3200W AC PS With 240V input. (1570W with 120V input. 16A input)
  - 2100W AC PS With 240V input. (940W with 120V input. 10.4A input)
  - 3200W DC PS With –40V to –72V input.

Fan Tray

- Redundant – N+1 fan

- Flexible Service – fan tray can be replaced from the portside or the back

- Efficient – Variable speed per fan depends on the load, temperature and altitudes (=>lower noise).

- Air flow – Side to side air flow

Architecture

You make networking **possible**

Cisco live!

Architecture

Centralized Architecture

Supervisor

Forwarding
Features
Embedded Controllers

Open
IOS-XE

Control Plane
Container
HA communication

Passive BackPlane
Upto 480G BW per slot

Line Card
STUB
ASIC

Line Card
STUB
ASIC

. . .

Line Card
STUB
ASIC
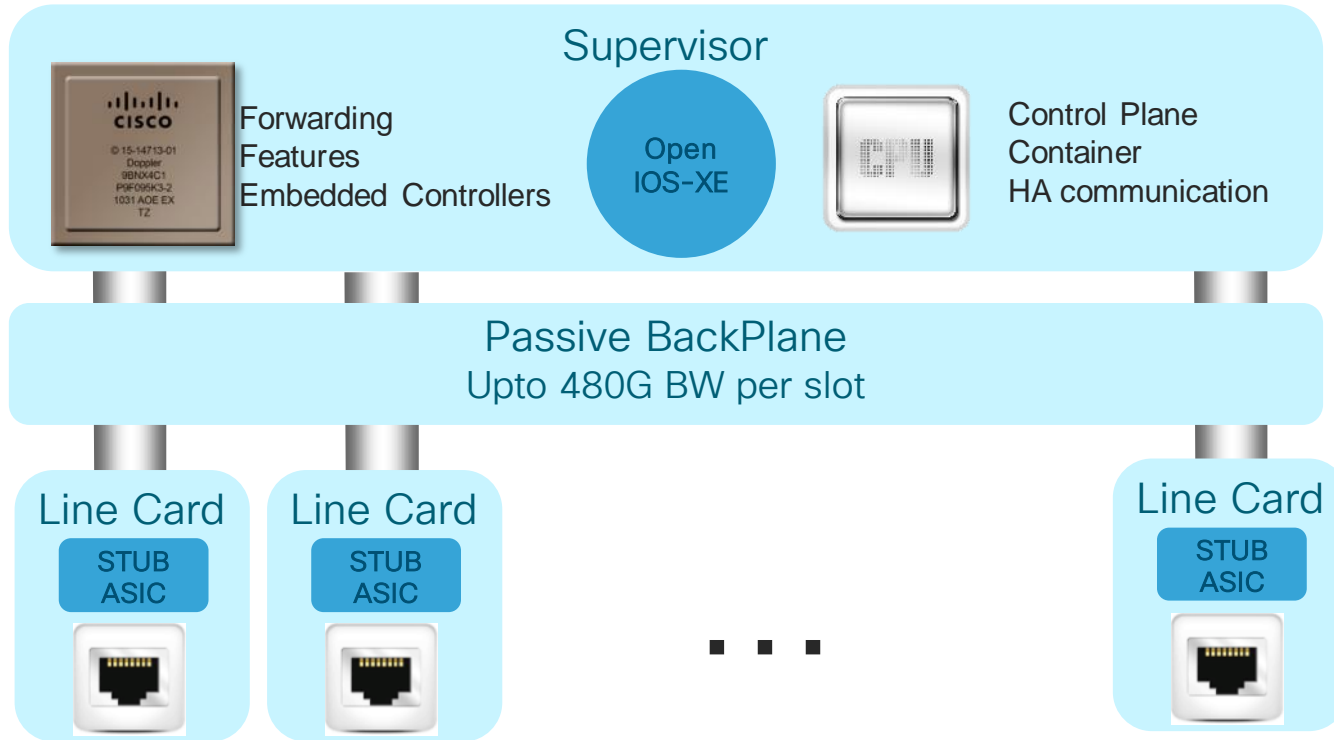
- Uninterrupted Sup switchover and software update

- Upgrade Supervisor unlock new capability

- High MTBF with passive backplane

© 15-14713-01
Doppler
9BNX4C1
P9F095K3-2
1031 AOE EX
TZ

CPU

Cisco live!

Open IOS-XE

IOS	IOS XE 3.7.x	Open IOS XE 16.x.x

IOS

Features Components

Common Infrastructure / HA

Management Interface

Module Drivers

Kernel

IOSd

Features Components

Hosted Apps

WCM

Wireshark

IOSd

IOSd Blob

IOS Sub Systems

Hosted Apps

Container

Container

Common Infrastructure / HA

Management Interface

Module Drivers

Kernel

IOS-XE DB

Same Look & Feel, Enhanced & Modern Architecture

Open IOS XE – IOS Sub Systems

IOSd Blob

IOS Sub Systems

BGP

OSPF

MPLS, etc

Failure of a Sub Systems is isolated
=>
Rest of the System intact
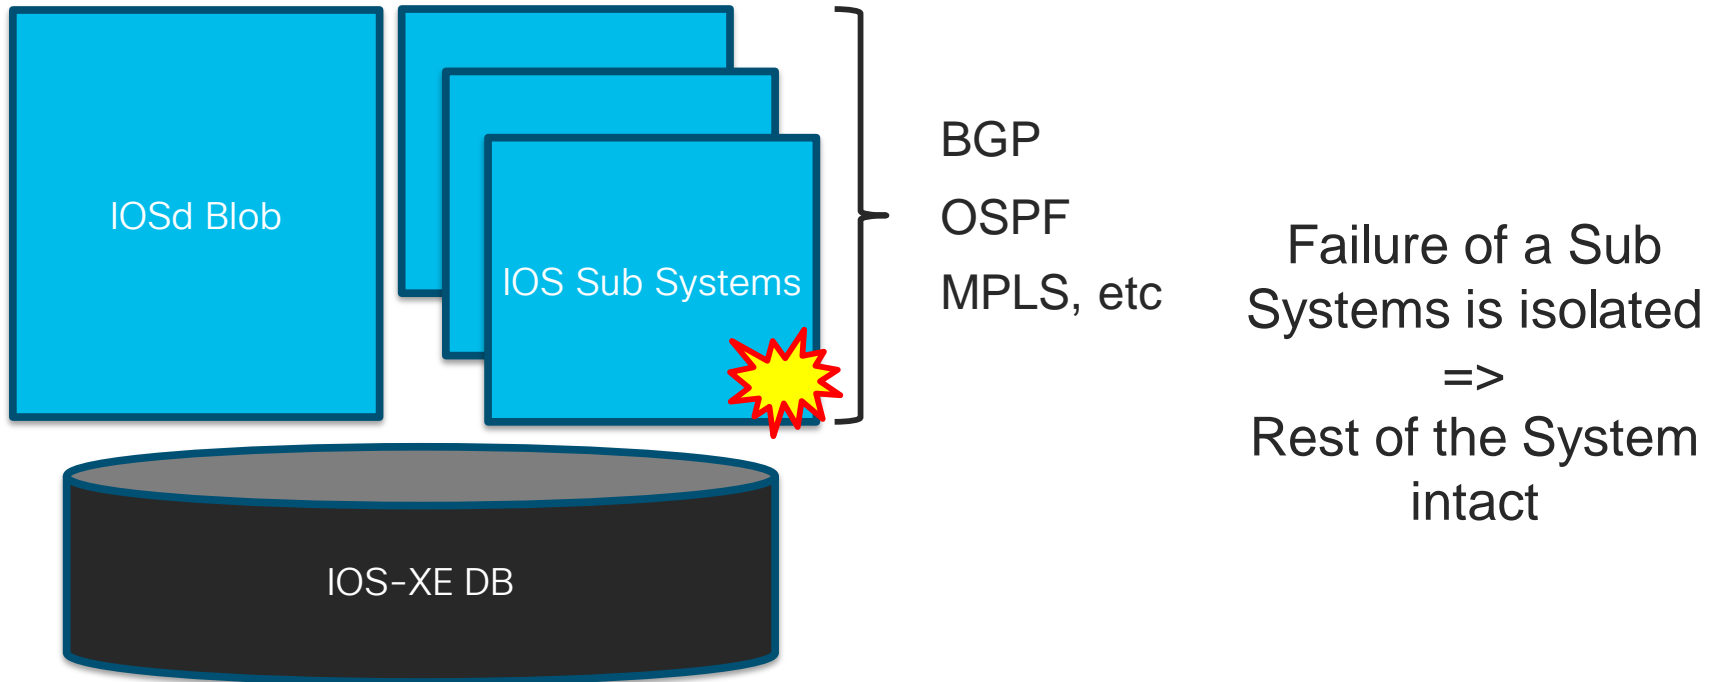
IOS-XE DB

IOS Sub Systems Enhances IOS Resiliency

Open IOS XE – DB

IOSd Sub Systems

IOS-XE DB

✔ Config & Operational States
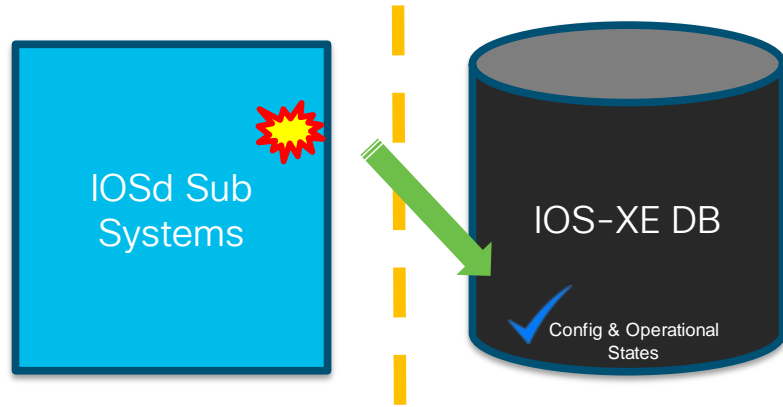
Decoupling Code & Data protects the Operational & Configurational States

Higher Application UP Time

Quicker Recovery

Better Convergence

Application Hosting in Catalyst 9K
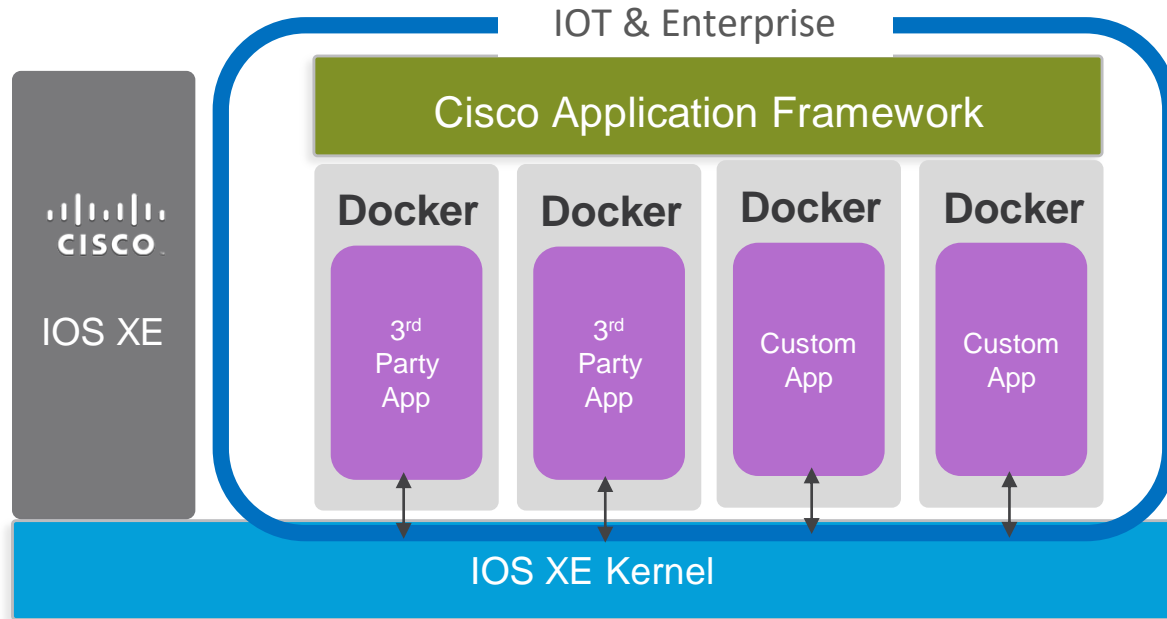
UADP 2.0XL – Next Generation of ASIC Innovation

Investment Protection	
Flexible Pipeline	

Universal Deployments	
Adaptable Tables	

Enhanced Scale/Buffering	
Multicore resource share	
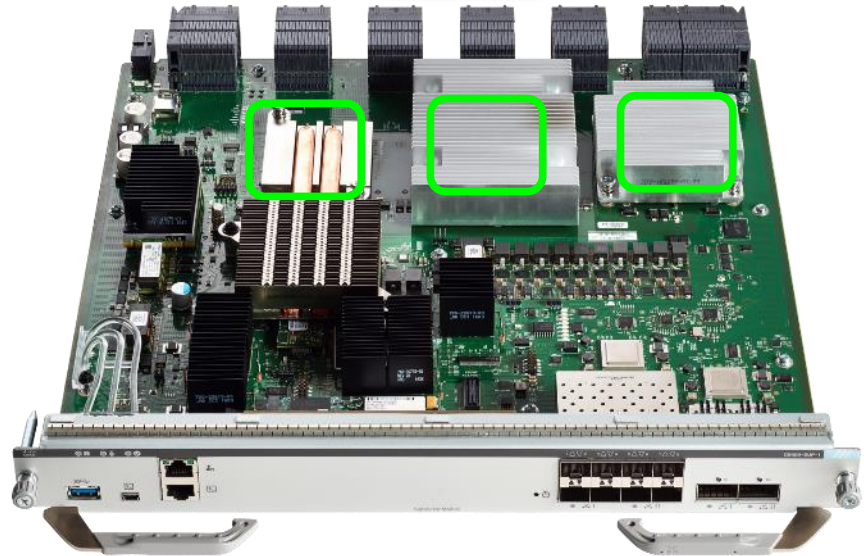
384K Flex Counters

Shared Lookup

Up to 240GE Bandwidth

Up to 2X to 4X forwarding + TCAM

Embedded Microcontrollers

Up to 32MB Packet Buffer

Up to 64K x2 Netflow Records

UADP 2.0 – 2 Cores

Shared Lookup Tables

Lookup Tables shared across the cores

Scales better at High Performance
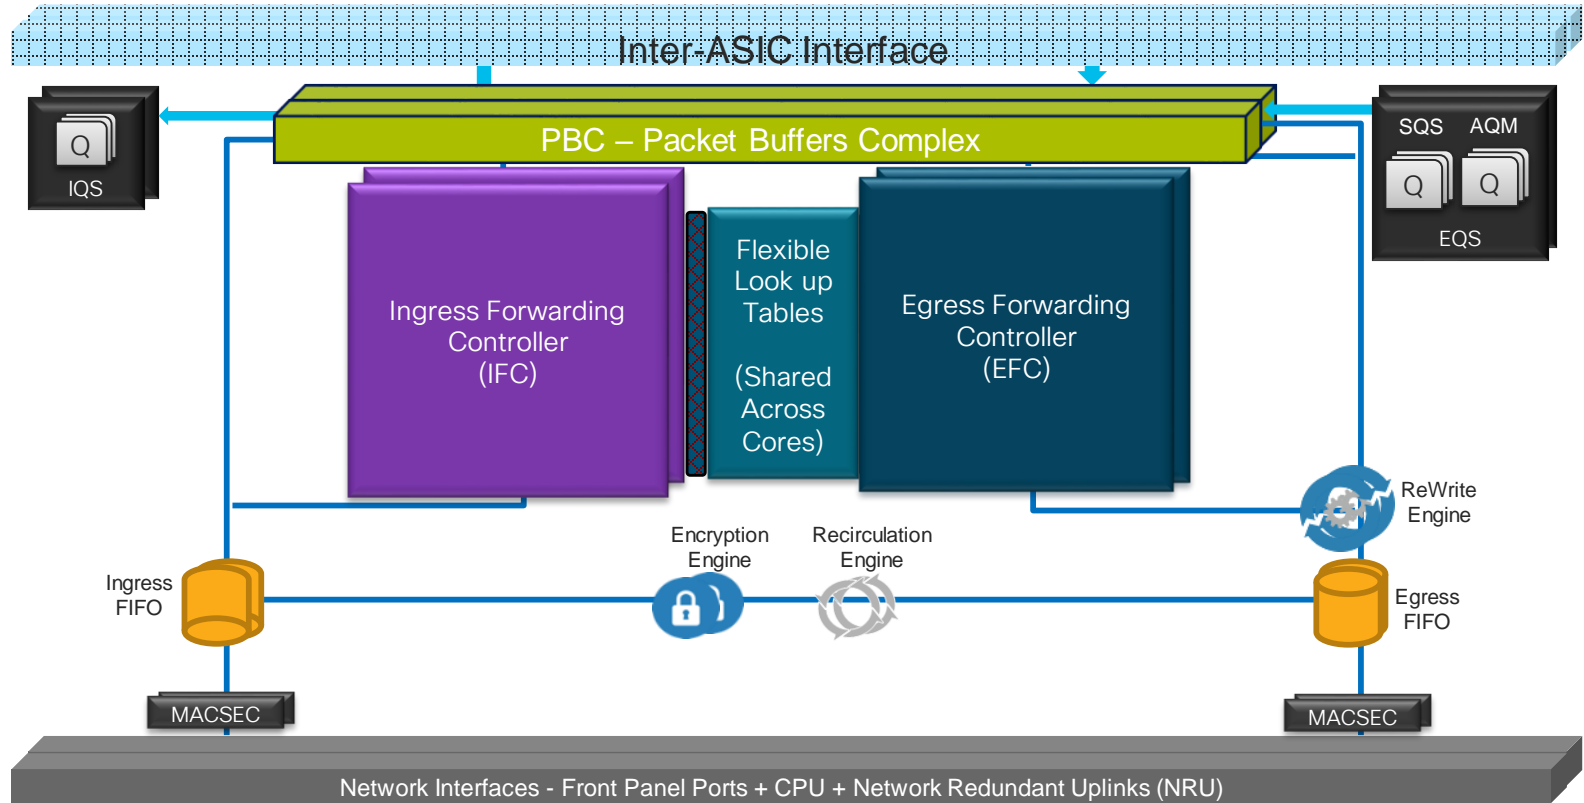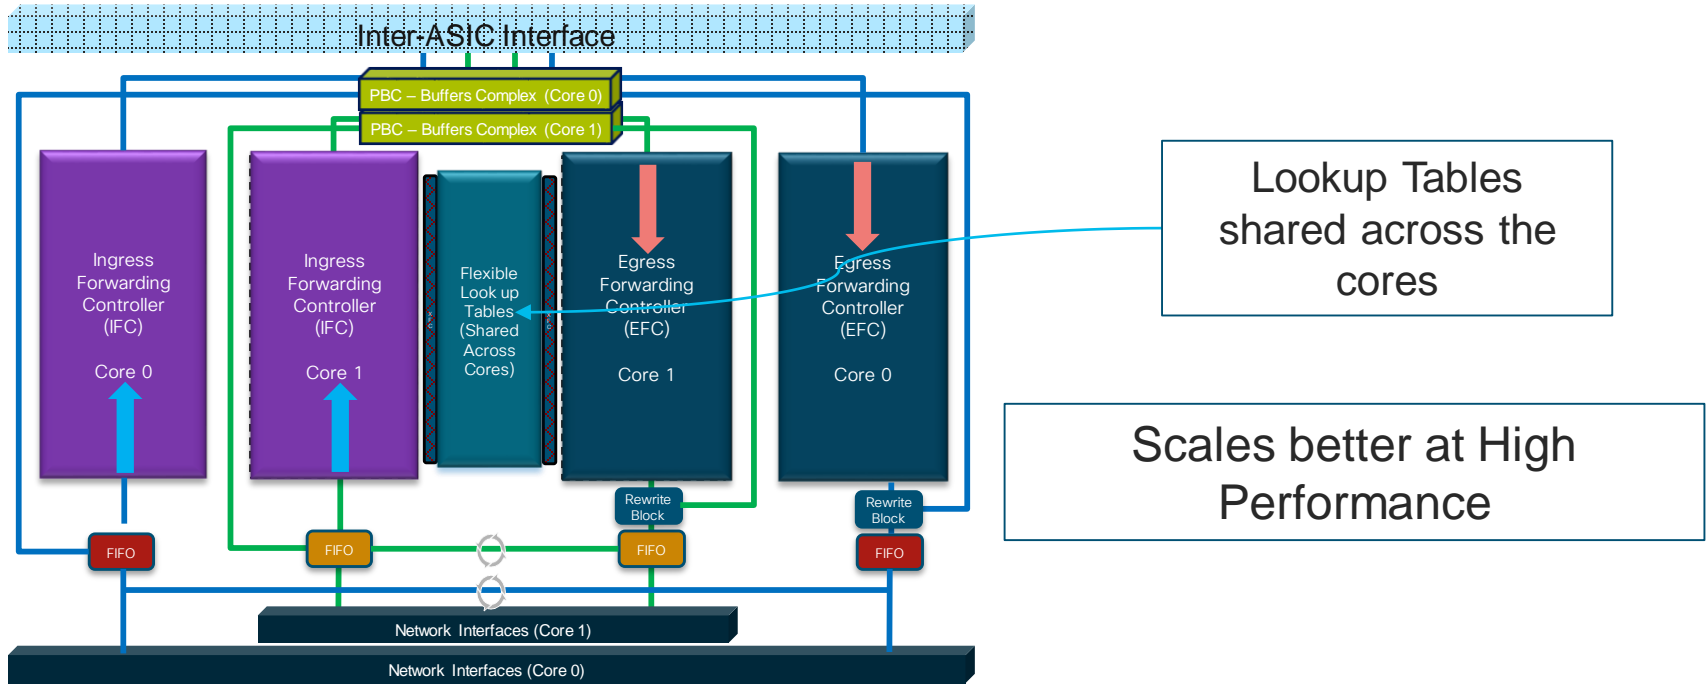
UADP 2.0 – Forwarding Pipelines
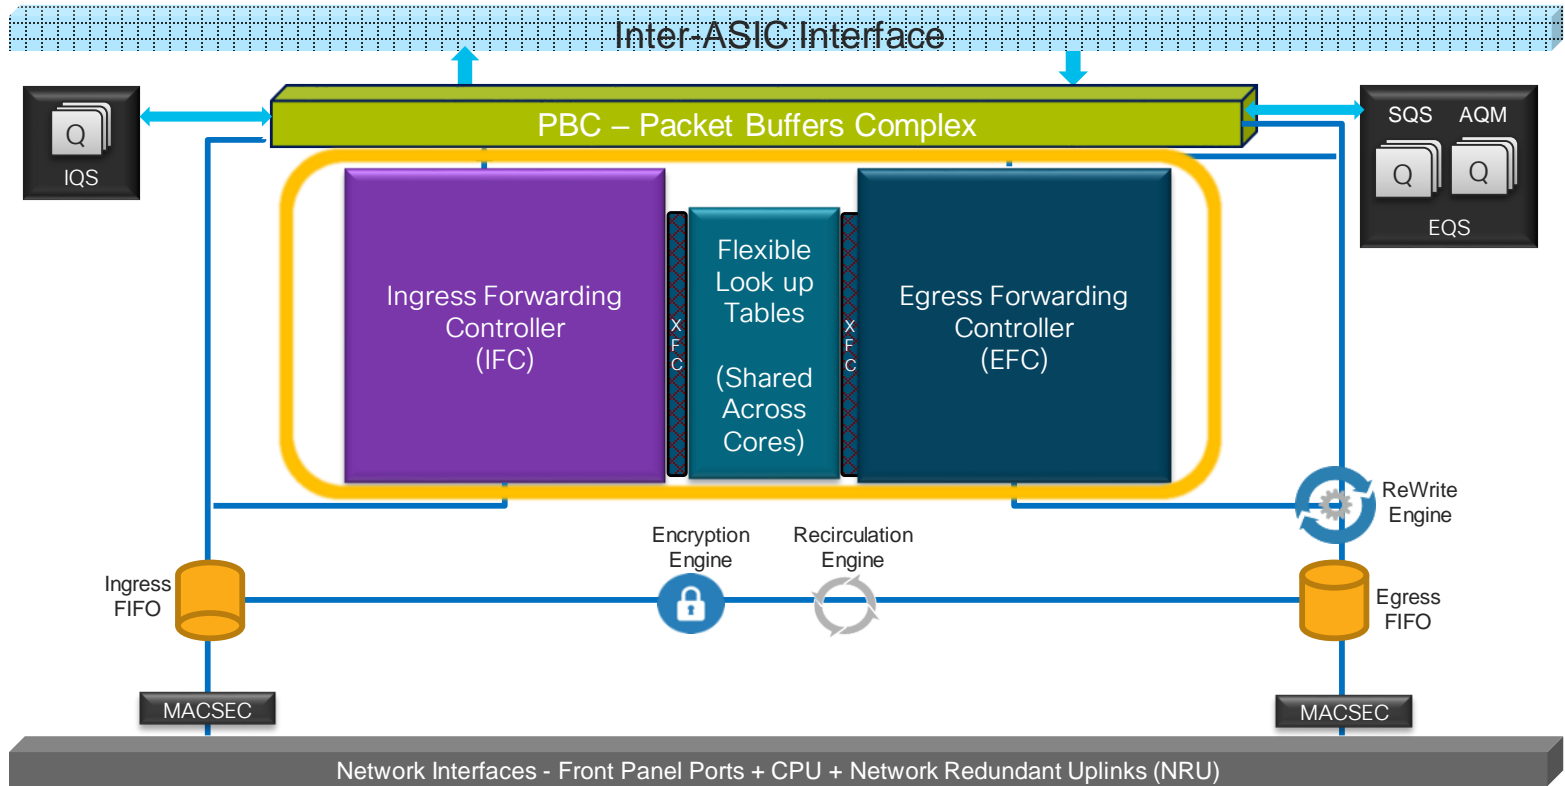
Forwarding - Under the covers

Final Decision on Packet's Future

IGR

Flex Parser

Flex Parser 256 B

15 Ingress Programmable Stages

Stage #15
Stage #..
Stage #..
Stage #2
Stage #1

Flexible Look up Tables

(Shared Across Cores)

Stage #1
Stage #2
Stage #..
Stage #..
Stage #8

8 Egress Programmable Stages

Flex Parser

Flex Parser 256 B

EGR

Final Decision on Packet's Future

Cisco live!
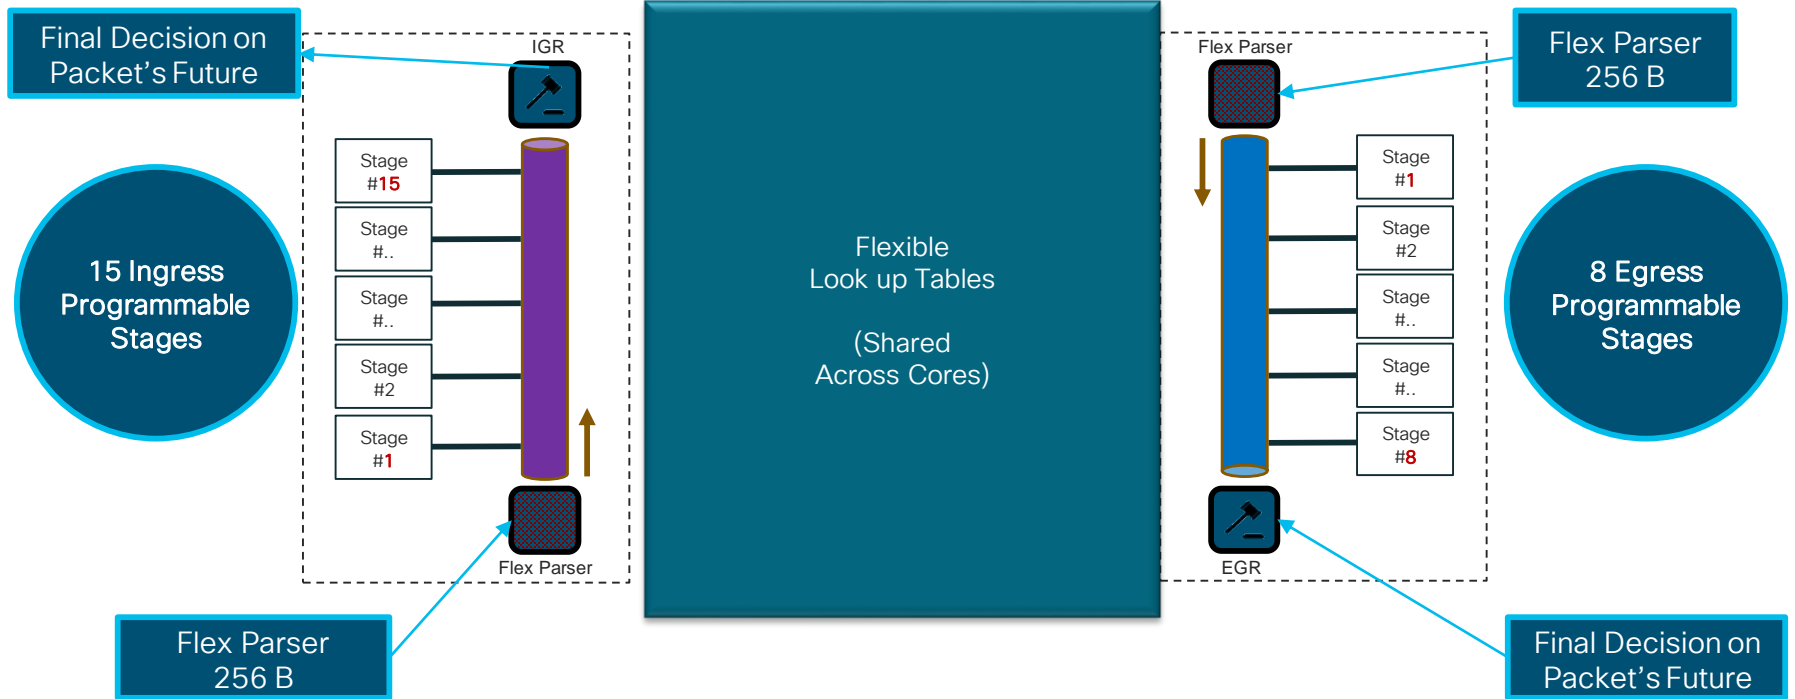
Microcode programs UADP 2.0

Ingress Programmable Pipeline

IGR

Punt
NF
SPAN
MC
L3/L2

Flex Parser

Programmed to understand VXLAN

Lookup Table (×20)

ACL   QOS

TCAM/ SRAM

At each stage, **2** simultaneous lookups

Flex Parser

Policy
NF
Sec
Stage #..
SPAN

EGR

Egress Programmable Pipeline

Programmed to understand MPLS

VXLAN   MPLS   NG Protocol

Software Features

Micro Code
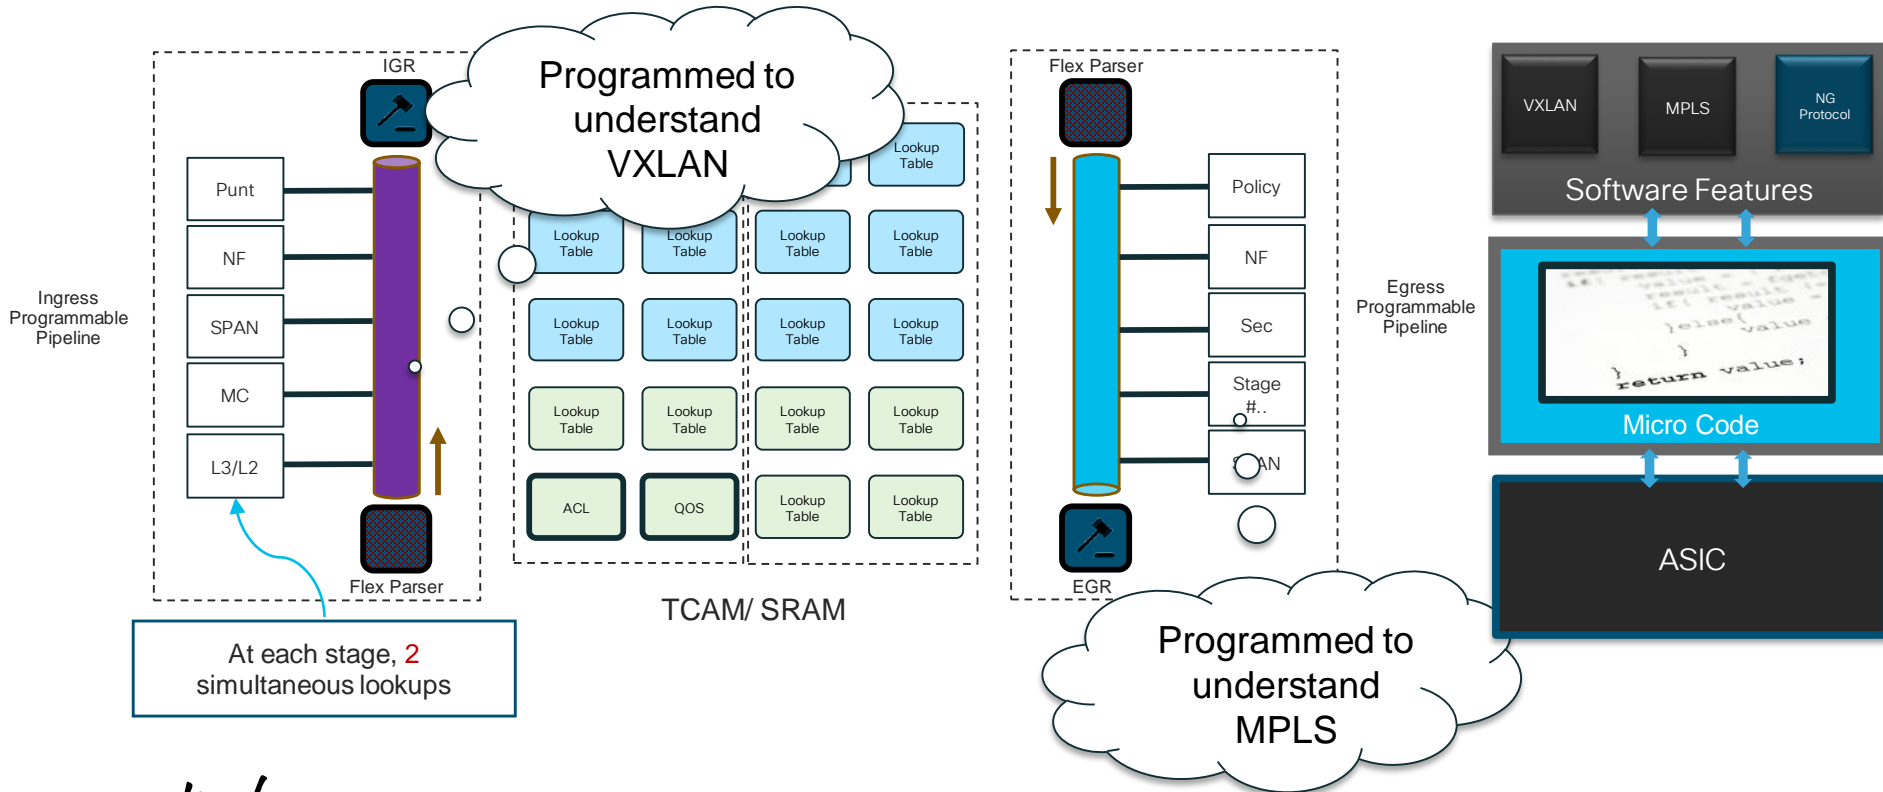
ASIC

UADP 2.0 – Look up Tables
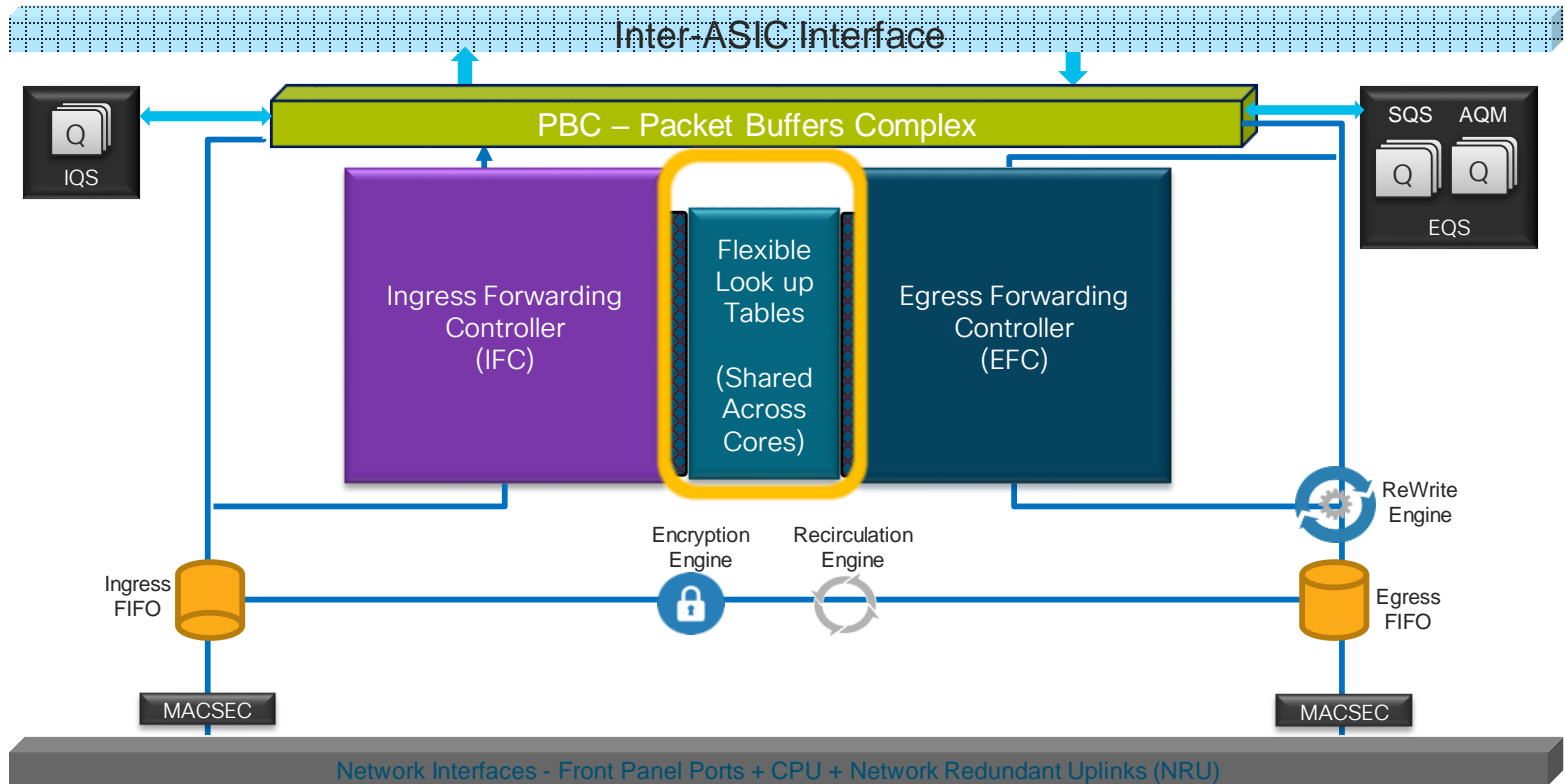
Flex Tables

Custom ASIC Templates for Universal Deployment

SRAM/TCAM

64K	
16K	
32K	
32K	
48K	
8K	
64K	
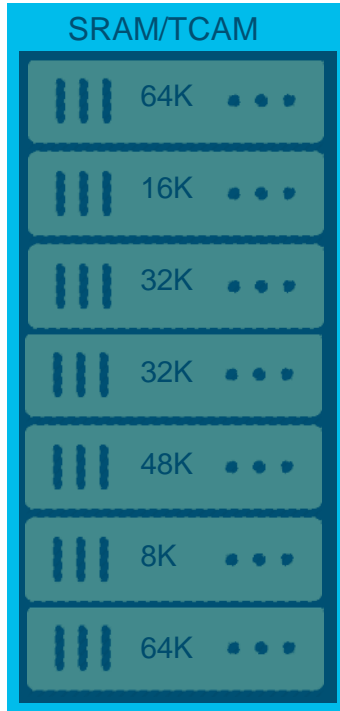
- MAC
- IPv4/IPv6
- VACL
- PACL
- RACL
- SGACL
- QoS
- NAT
- SPAN
- CoPP

Flexible ASIC Templates

Access-Edge Template	Core-Border Template	Aggregation Template	Collapsed Core-WAN Template
Client Scale	L3 & Cross Domain Policy	Mix of L2/L3 Capabilities	L3 & NAT
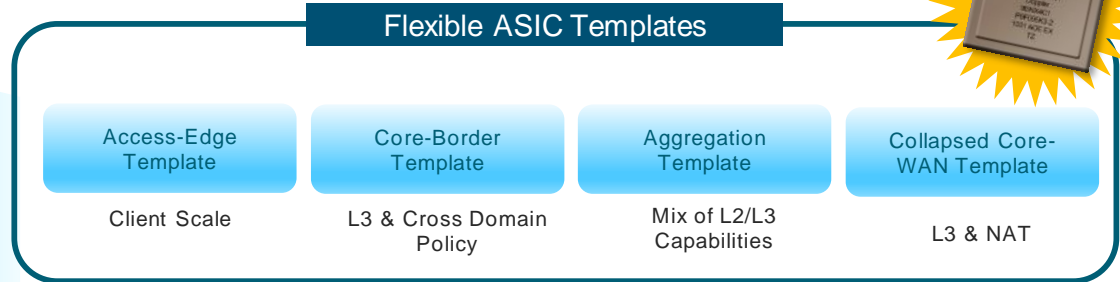
Customize table size for each function based on the place in the network

Access-Edge Template

Core-Border Template
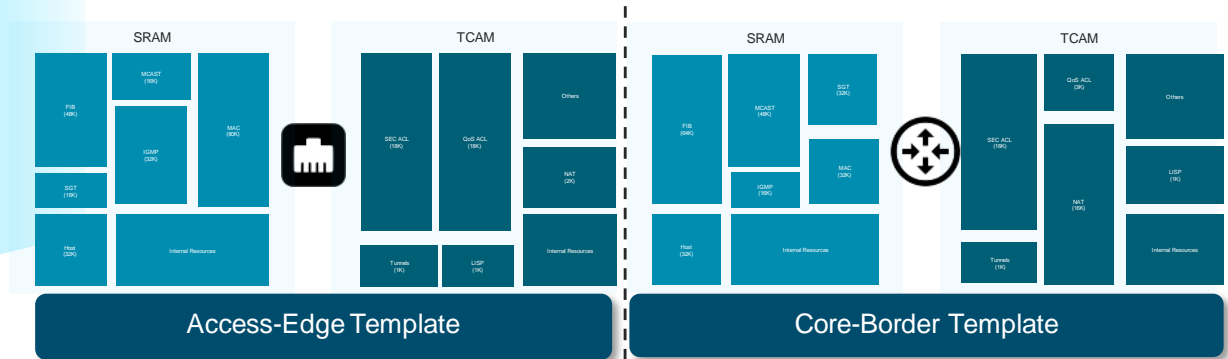
Cisco Catalyst 9400 Series
SDM templates and scale numbers Sup-1XL

Feature		Access/Distribution template (default)	Core template	SDA template	NAT template
Indirect/LPM Routes (IPv4/IPv6)		64K / 32K	64K / 32K	64K / 32K	64K / 32K
Direct/Host Routes (IPv4/IPv6)		48K / 24K	32K / 16K	80K / 16K	48K / 24K
Multicast routes (IPv4/IPv6)		16K / 8K	16K / 8K	16K / 8K	16K / 8K
MAC address table		65K	16K	16K	16K
Flexible netflow		128K/ASIC	128K/ASIC	128K/ASIC	128K/ASIC
SGT label		8K	8K	8K	8K
Security ACL	Ingress	18K			18K
	Egress				
QOS ACL	Ingress	18K			3K
	Egress				
Netflow ACL	Ingress	1K			1K
	Egress	2K			2K
SPAN	Ingress	1K			1K
	Egress				
PBR/NAT		2K			16K
CPP		1K			1K
Tunnel termination and MACSEC		1K			1K
LISP		1K			1K

Sup-1/Sup-1XL Block Diagram

ASIC #	4-Slot	7-Slot	10-Slot
UADP #1	Slot 1	Slot 2 and 7	Slot1, 9, and 10
UADP #2	Slot 4	Slot 1 and 5	Slot 2, 3 and 4
UADP #3	Uplinks	Slot 6 and Uplinks	Slot 7, 8, and Uplinks
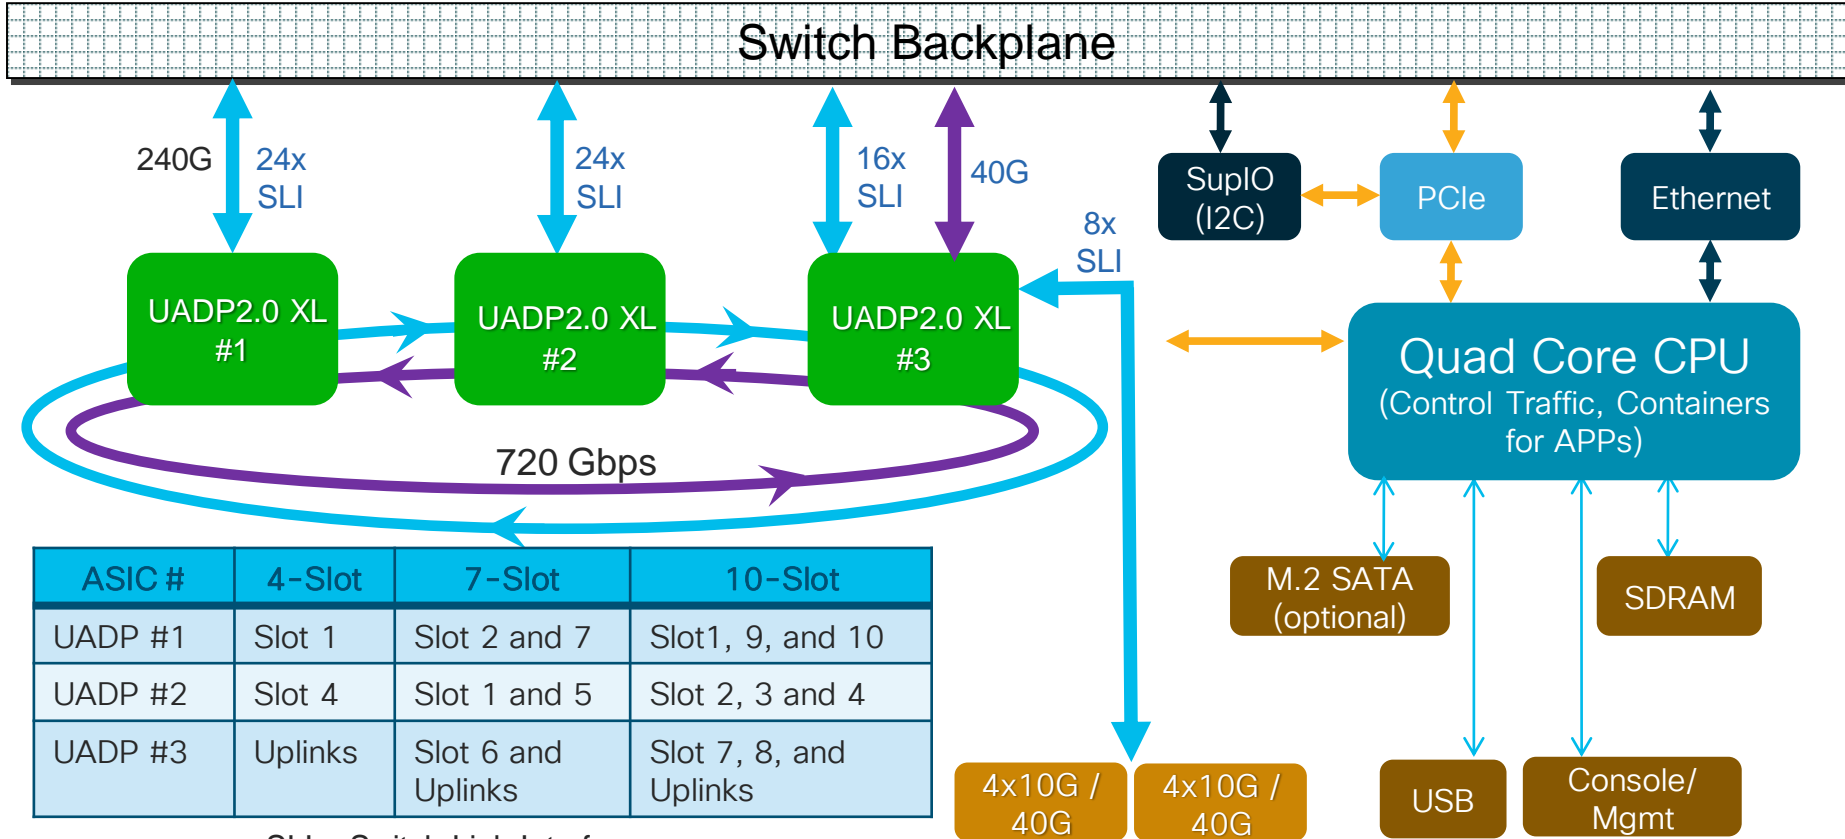
SLI – Switch Link Interfaces

Port to ASIC Mapping

```
switch#show platform software fed active ifm mappings
Interface               IF_ID   Inst  Asic  Core  Port  SubPort  Mac  Cntx  LPN  GPN  Type  Active
GigabitEthernet1/0/1    0x7     2     1     0     0     0        4    4     1    101  NIF   Y
GigabitEthernet1/0/2    0x8     2     1     0     1     1        4    4     2    102  NIF   Y
GigabitEthernet1/0/3    0xf5    2     1     0     2     2        4    4     3    103  NIF   Y
GigabitEthernet1/0/4    0xa     2     1     0     3     3        4    4     4    104  NIF   Y
GigabitEthernet1/0/5    0xb     2     1     0     4     4        4    4     5    105  NIF   Y
GigabitEthernet1/0/6    0xc     2     1     0     5     5        4    4     6    106  NIF   Y
GigabitEthernet1/0/7    0xd     2     1     0     6     6        4    4     7    107  NIF   Y
GigabitEthernet1/0/8    0xe     2     1     0     7     7        4    4     8    108  NIF   Y
GigabitEthernet1/0/9    0xf     2     1     0     8     0        0    0     9    109  NIF   Y
GigabitEthernet1/0/10   0x10    2     1     0     9     1        0    0     10   110  NIF   Y
GigabitEthernet1/0/11   0x11    2     1     0     10    2        0    0     11   111  NIF   Y
GigabitEthernet1/0/12   0x12    2     1     0     11    3        0    0     12   112  NIF   Y
GigabitEthernet1/0/13   0x13    2     1     0     12    4        0    0     13   113  NIF   Y
GigabitEthernet1/0/14   0x14    2     1     0     13    5        0    0     14   114  NIF   Y
GigabitEthernet1/0/15   0x15    2     1     0     14    6        0    0     15   115  NIF   Y
GigabitEthernet1/0/16   0x16    2     1     0     15    7        0    0     16   116  NIF   Y
GigabitEthernet1/0/17   0x17    2     1     0     16    0        5    5     17   117  NIF   Y
GigabitEthernet1/0/18   0x18    2     1     0     17    1        5    5     18   118  NIF   Y
GigabitEthernet1/0/19   0x19    2     1     0     18    2        5    5     19   119  NIF   Y
GigabitEthernet1/0/20   0x1a    2     1     0     19    3        5    5     20   120  NIF   Y
GigabitEthernet1/0/21   0x1b    2     1     0     20    4        5    5     21   121  NIF   Y
GigabitEthernet1/0/22   0x1c    2     1     0     21    5        5    5     22   122  NIF   Y
GigabitEthernet1/0/23   0x1d    2     1     0     22    6        5    5     23   123  NIF   Y
GigabitEthernet1/0/24   0x1e    2     1     0     23    7        5    5     24   124  NIF   Y
GigabitEthernet1/0/25   0x1f    2     1     0     24    0        1    1     25   125  NIF   Y
<SNIP>

switch#
```

Sup-1/Sup-1XL Uplink - Single Sup

Default Mode (8x 10GE)

2x 40G
`interface FortyGigabitEthernet<slot>/0/[9-10]`
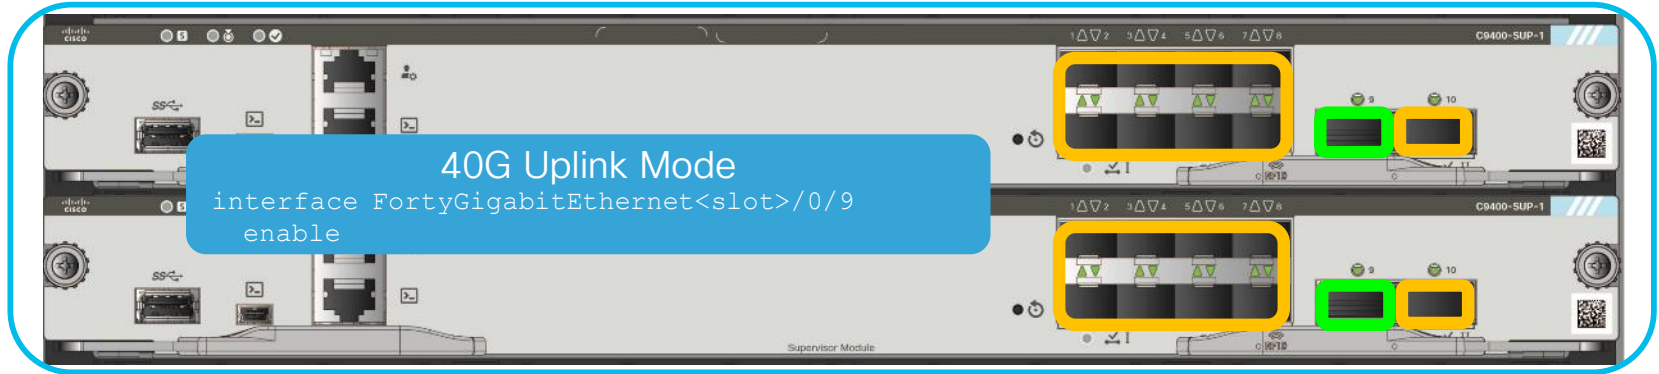`enable`

Mix Mode (4x 10GE + 1x 40G)
`interface FortyGigabitEthernet<slot>/0/10`
`enable`

Mix Mode (4x 10GE + 1x 40G)
`interface FortyGigabitEthernet<slot>/0/9`
`enable`

Active

Disabled

Sup-1/Sup-1XL Dual Sups - Uplink Redundancy

Default Mode (8x 10GE)

40G Uplink Mode
```
interface FortyGigabitEthernet<slot>/0/9
    enable
```

Active

Disabled

Sup-1/Sup-1XL Dual Sups - Mix Uplink Mode
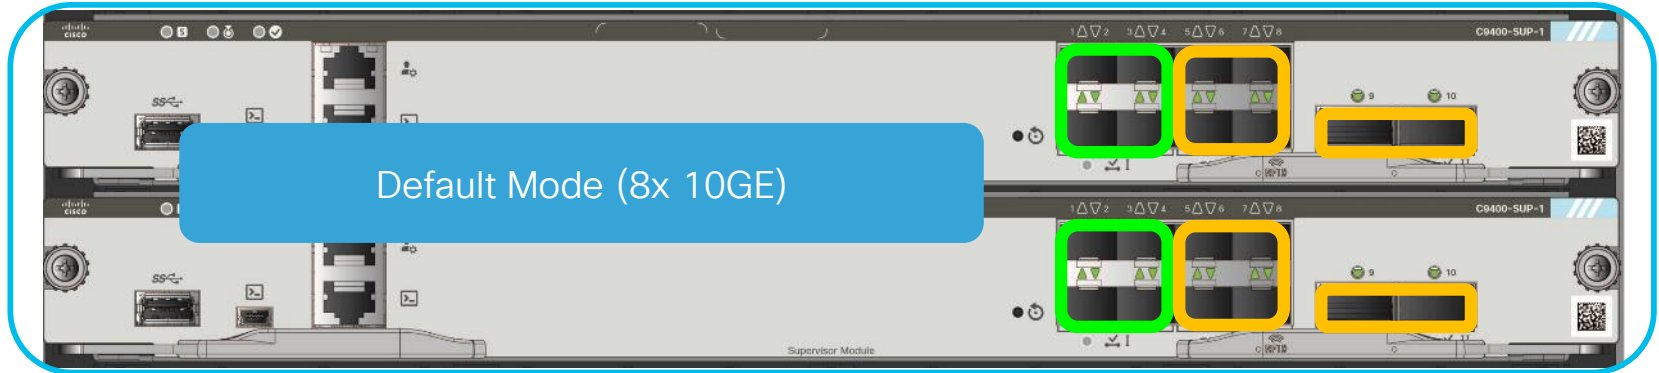

Mix Mode (4x 10GE + 1x 40GE)
```
interface FortyGigabitEthernet<slot>/0/9
   enable
```

Mix Mode (1x 40GE + 4x 10GE)
```
interface FortyGigabitEthernet<slot>/0/9
   enable
```
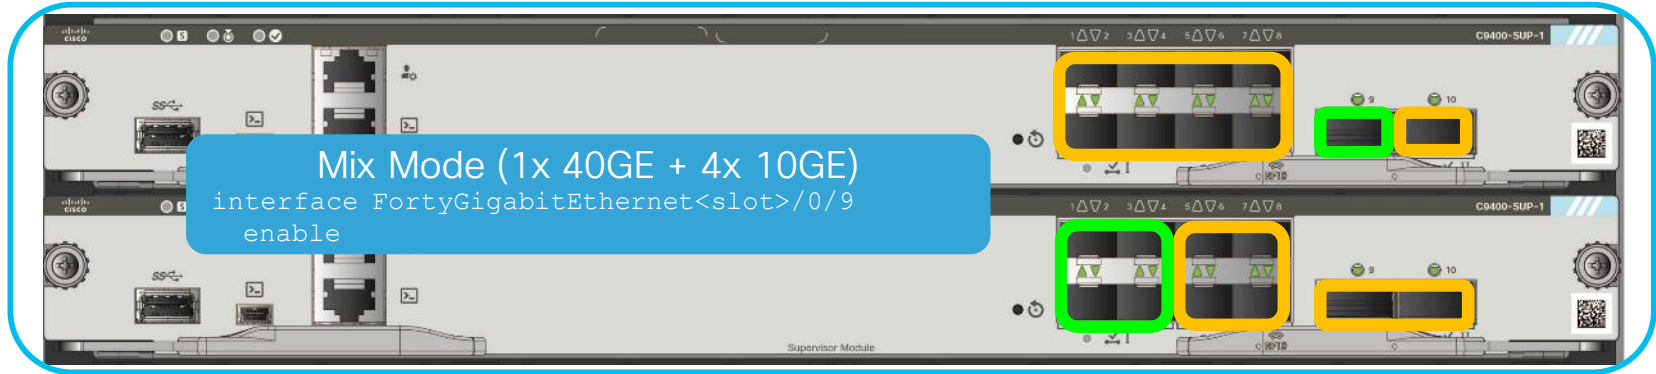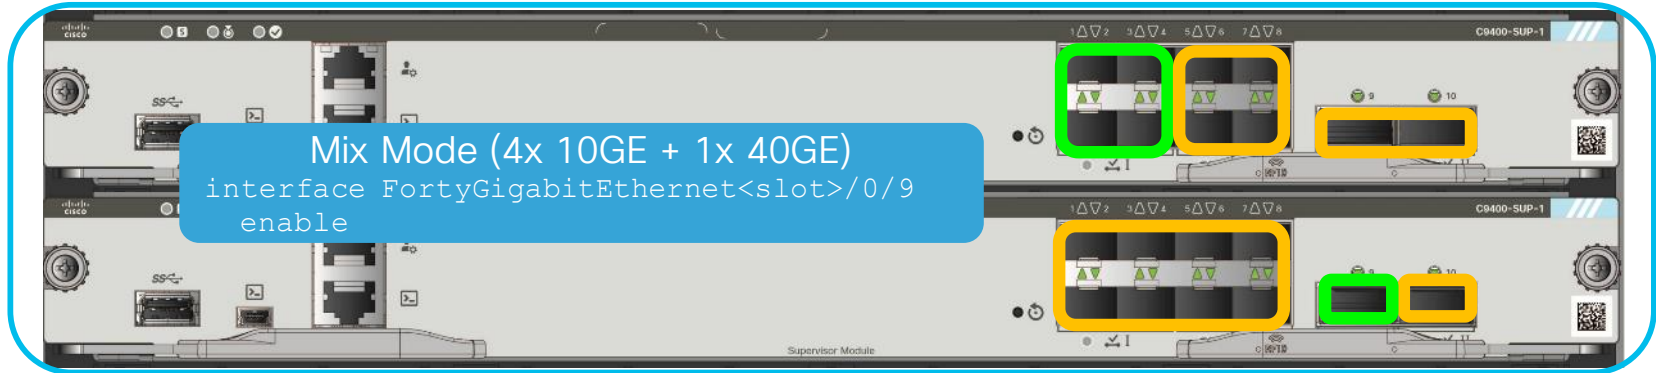
Active

Disabled

C9400-Sup-1XL-Y (25G Uplinks)

Tw <slot#>/0/1

Tw <slot#>/0/5

Ten <slot#>/0/1
Ten <slot#>/0/2
Ten <slot#>/0/3
Ten <slot#>/0/4

Ten <slot#>/0/5
Ten <slot#>/0/6
Ten <slot#>/0/7
Ten <slot#>/0/8

Four <slot#>/0/9
Four <slot#>/0/10
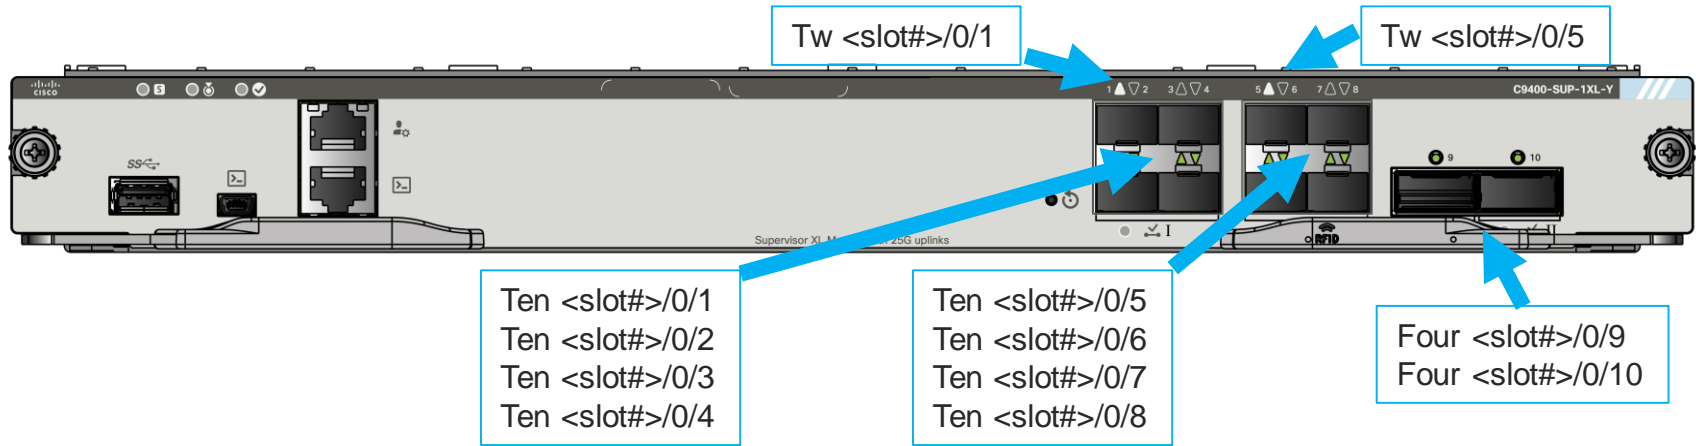
If port 9 (and/or 10) is enabled, port 1-4 (and/or 5-8) are disabled.
    => 40G ports are enabled.

If port 9 (and/or 10) is disabled (which is default)
    If Tw <slot#>/0/1(and/or 5) is enabled, Ten <slot#>/0/1 – 4 (and/or 5-8) are disabled.
    If Tw <slot#>/0/1(and/or 5) is disabled, Ten <slot#>/0/1 – 4 (and/or 5-8) are enabled. (which is default)

Note:
1.   TW ports are not auto-sense with speed for 10G/1G, manual speed configuration is needed. (same for dual-rate SFP28)

48x1G RJ45 Line Card (PoE+/UPoE)

FYI

Switch Backplane

6x
SLI

Stub ASIC

Management
Sub-System

IEEE
1588

10G
(USGMII)

Octal PHY	Octal PHY	Octal PHY	Octal PHY	Octal PHY	Octal PHY
RJ45 1-8	RJ45 9-16	RJ45 17-24	RJ45 25-32	RJ45 33-40	RJ45 41-48

Power
Sub-system

PoE
Sub-System

Speed10/100/1000
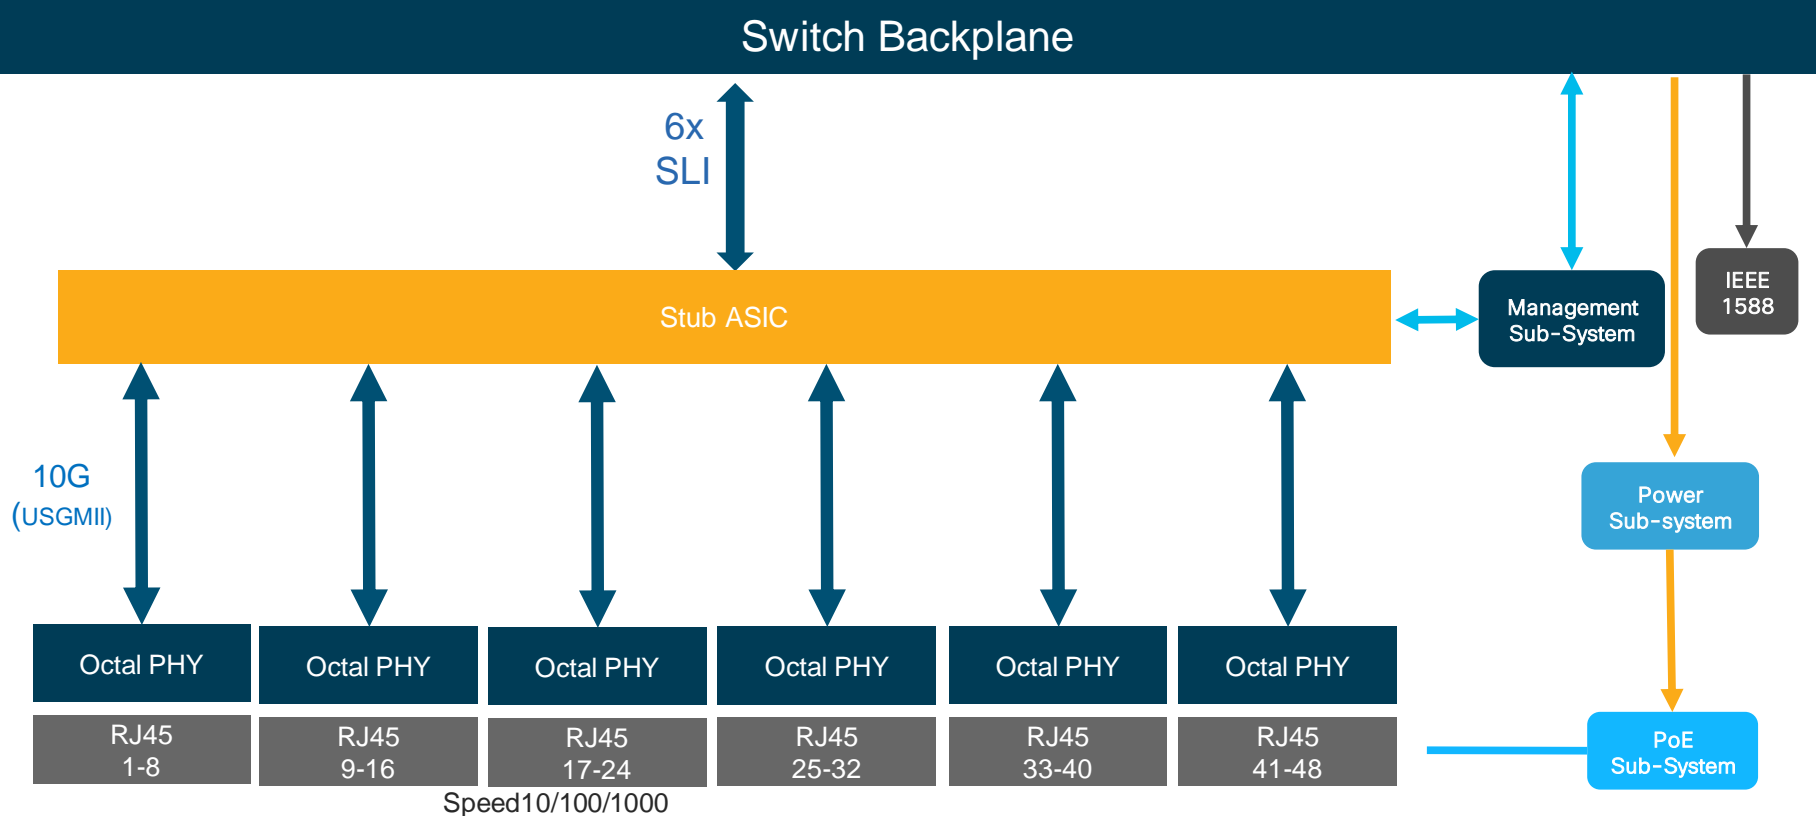
Line Rate on all Ports; UPoE on all Ports

48x1G Line Card (RJ45 Data or SFP)
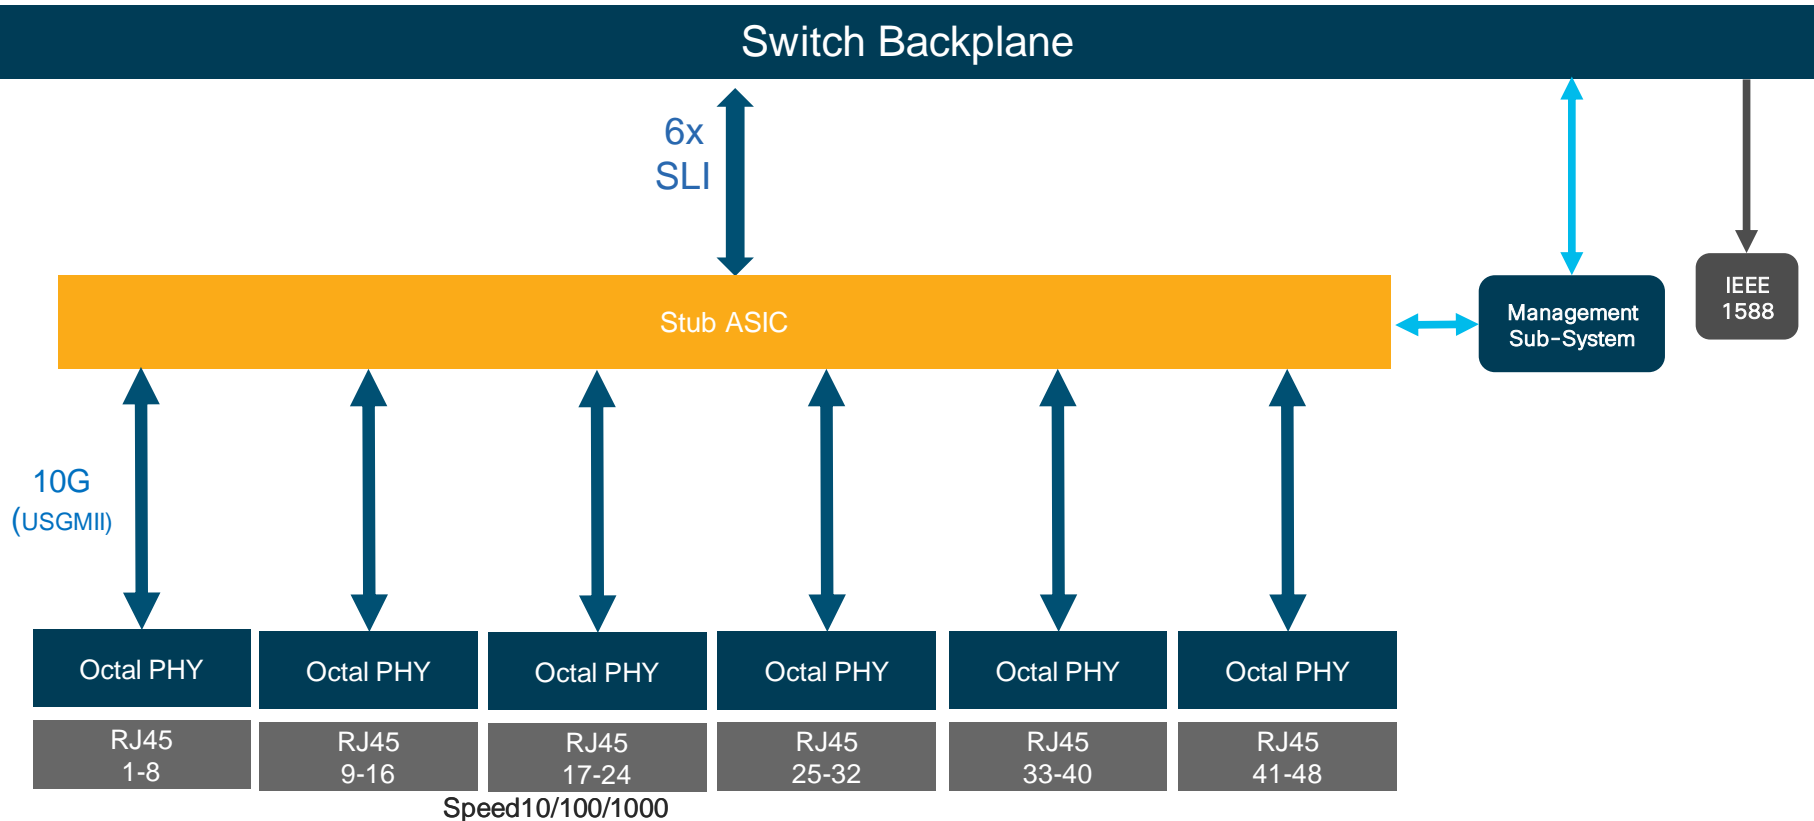
Switch Backplane

6x SLI

Stub ASIC

Management Sub-System

IEEE 1588

10G (USGMII)

Octal PHY	Octal PHY	Octal PHY	Octal PHY	Octal PHY	Octal PHY
RJ45 1-8	RJ45 9-16	RJ45 17-24	RJ45 25-32	RJ45 33-40	RJ45 41-48

Speed10/100/1000

Line Rate on all Ports

24x 1/10G SFP/SFP+ Line Card
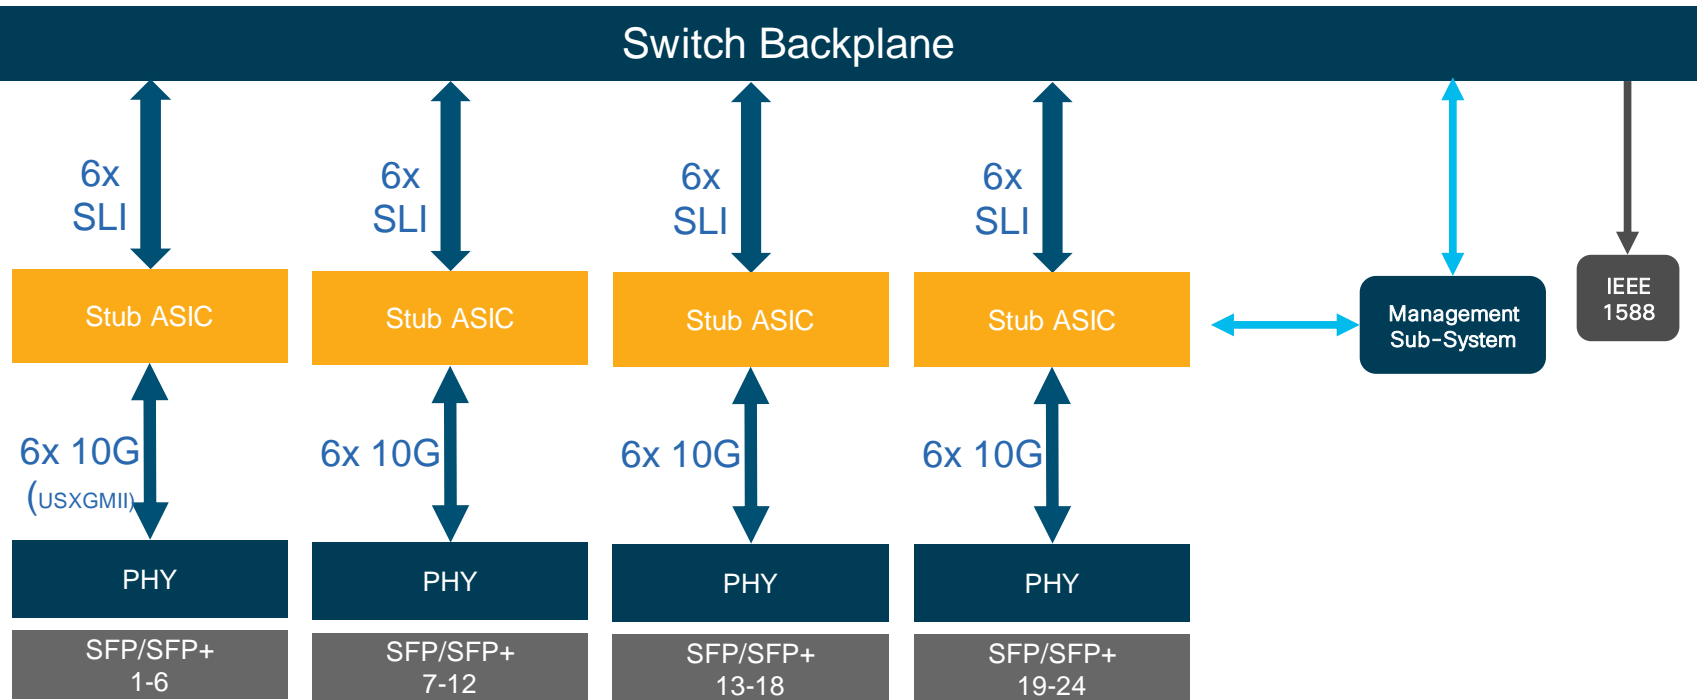
Switch Backplane

6x SLI — 6x SLI — 6x SLI — 6x SLI

Stub ASIC — Stub ASIC — Stub ASIC — Stub ASIC

Management Sub-System

IEEE 1588

6x 10G (USXGMII) — 6x 10G — 6x 10G — 6x 10G

PHY — PHY — PHY — PHY

SFP/SFP+ 1-6 — SFP/SFP+ 7-12 — SFP/SFP+ 13-18 — SFP/SFP+ 19-24

Sup-1: 80G with C9410R, C9407R and C9404R.
Sup-1XL: 80G with C9410R; 120G with C9407R; 240G with C9404R.

C9400-LC-24XS Port-Group
With XL Supervisors

7 Slot Chassis: 12 Port-Group

2 Ports in a port-group

10 Slot Chassis: 8 Port-Group

3 Ports in a port-group

C9400-LC-24XS Port-Group – 7 Slot Chassis

- Bandwidth shared within port-group

- 12 port-group in the 7 slot chassis

- Modes: dynamic, performance and static

- For 10G line rate performance:
  - Configure: "`hw-module subslot <slot#/0> mode performance`"
  - 8 Port @ line-rate, other ports are disabled

```
R4-C94-2041#show platform hardware iomd 5/0 portgroups
Port    Interface                       Status      Interface   Group
Max
Group                                               Bandwith    Bandwidth

1       TenGigabitEthernet5/0/1         up          10G
1       TenGigabitEthernet5/0/2         down        10G         10G

2       TenGigabitEthernet5/0/3         up          10G
2       TenGigabitEthernet5/0/4         down        10G         10G

3       TenGigabitEthernet5/0/5         up          10G
3       TenGigabitEthernet5/0/6         down        10G         10G

4       TenGigabitEthernet5/0/7         up          10G
4       TenGigabitEthernet5/0/8         down        10G         10G
<SNIP>
11      TenGigabitEthernet5/0/21        up          10G
11      TenGigabitEthernet5/0/22        down        10G         10G


12      TenGigabitEthernet5/0/23        up          10G
12      TenGigabitEthernet5/0/24        down        10G         10G

R4-C94-2041#show
```
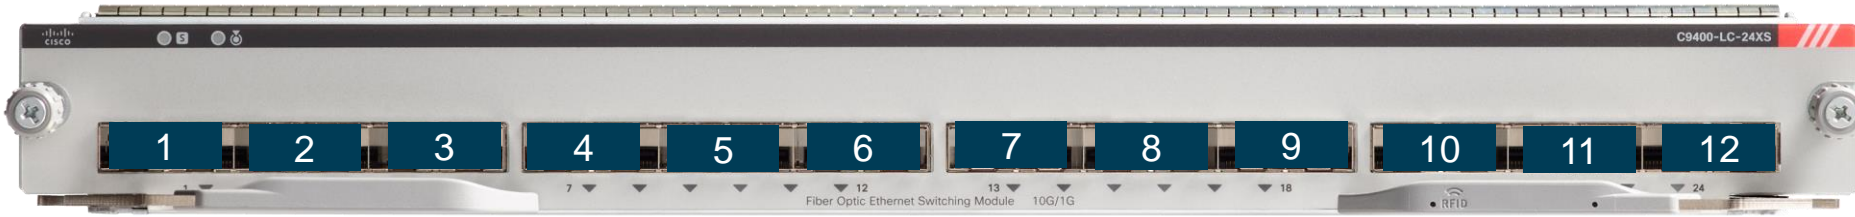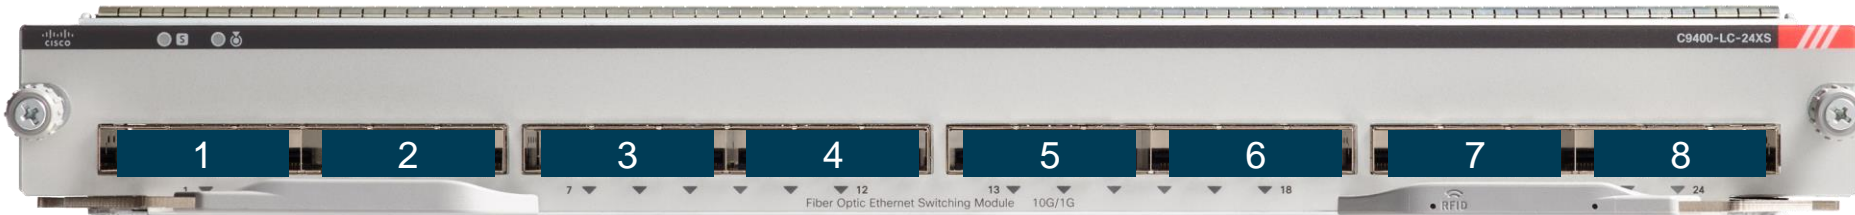
C9400-LC-24XS Port-Group – 10 Slot Chassis

- Bandwidth shared within port-group

- 8 port-group in the 10 slot chassis

- Modes: dynamic, performance and static

- For 10Gline rate performance:
  - Configure: "`hw-module subslot <slot#/0> mode performance`"
  - 5 Port @ line-rate, other ports are disabled

```
mac1#show platform hardware iomd 10/0 portgroups
Port    Interface                       Status     Interface   Group Max
Group                                              Bandwith    Bandwith

1       TenGigabitEthernet10/0/1        admindown  10G
1       TenGigabitEthernet10/0/2        admindown  10G
1       TenGigabitEthernet10/0/3        admindown  10G         10G

2       TenGigabitEthernet10/0/4        admindown  10G
2       TenGigabitEthernet10/0/5        admindown  10G
2       TenGigabitEthernet10/0/6        admindown  10G         10G

3       TenGigabitEthernet10/0/7        admindown  10G
3       TenGigabitEthernet10/0/8        admindown  10G
3       TenGigabitEthernet10/0/9        admindown  10G         10G
<SNIP>
7       TenGigabitEthernet10/0/19       admindown  10G
7       TenGigabitEthernet10/0/20       admindown  10G
7       TenGigabitEthernet10/0/21       admindown  10G         10G

8       TenGigabitEthernet10/0/22       admindown  10G
8       TenGigabitEthernet10/0/23       down       10G
8       TenGigabitEthernet10/0/24       admindown  10G         10G


mac1#
```

mGig RJ45 Line Card

Switch Backplane

6x SLI

Stub ASIC — 2x / 6x — 4x Octal PHY (use 6 out of 8) — RJ45 1-24 — Speed10/100/1000

Stub ASIC — 2x / 6x — 2x Quad mGig PHY — RJ45 25-32

Stub ASIC — 2x / 6x — 2x Quad mGig PHY — RJ45 33-40

Stub ASIC — 2x / 6x — 2x Quad mGig PHY — RJ45 40-48

Speed100/1G/2.5G/5G/10G

Management Sub-System

IEEE 1588

Power Sub-system

PoE Sub-System

Sup-1: 80G with C9410R, C9407R and C9404R.
Sup-1XL: 80G with C9410R; 120G with C9407R; 240G with C9404R.

C9400-LC-48UX Port-Group
With XL Supervisors

7 Slot Chassis: 12 Port-Group

4 Ports in a port-group

10 Slot Chassis: 8 Port-Group

6 Ports in a port-group

C9400-LC-48UX Port-Group – 7 Slot Chassis

FYI

- Bandwidth shared within port-group

- 12 port-group in the 7 slot chassis

- Modes: dynamic, performance and static

- For 10G line rate performance:
  - Configure: "`hw-module subslot <slot#/0> mode performance`"
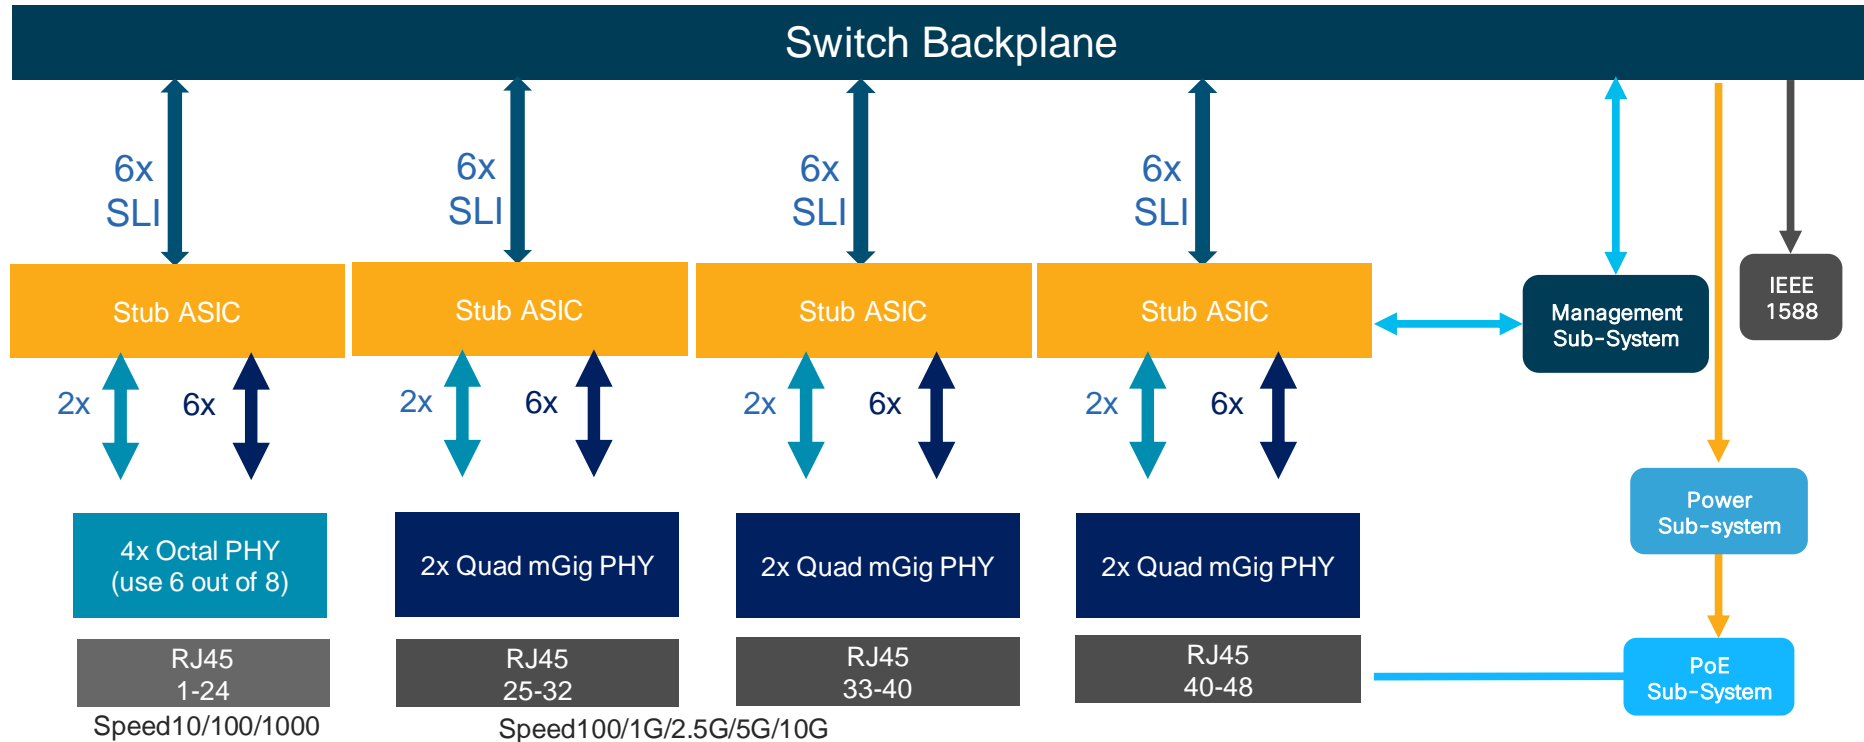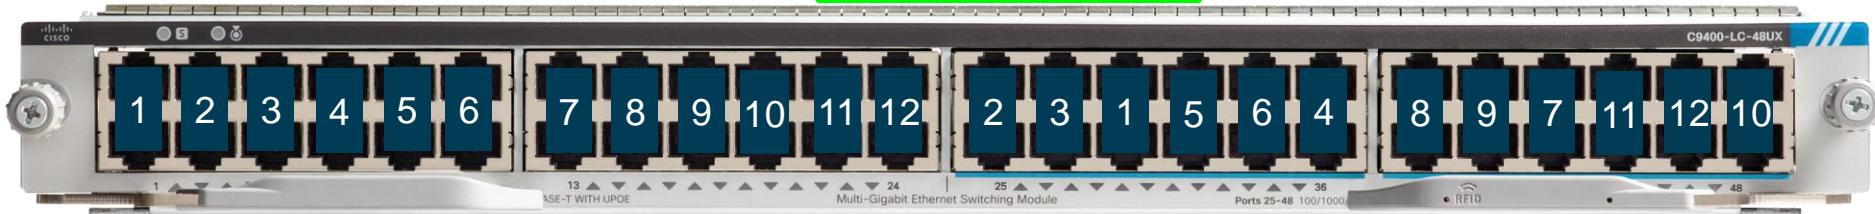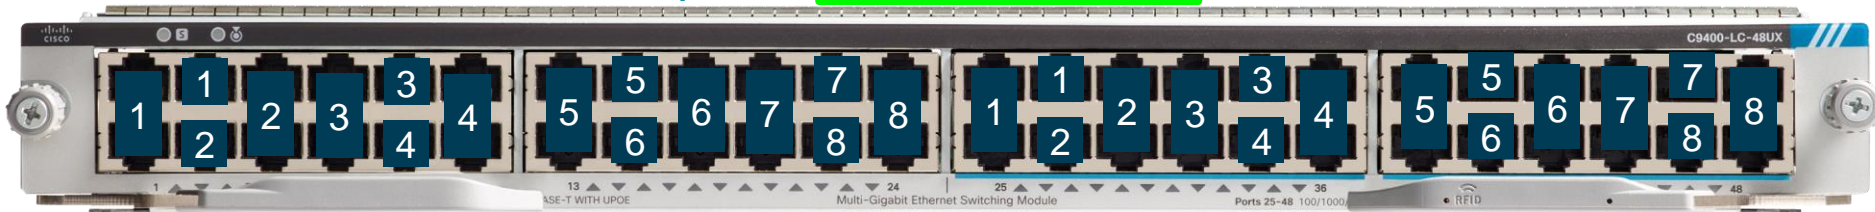 - 8 Port @ line-rate, other ports are disabled

```
R4-C94-2041#show platform hardware iomd 6/0 portgroups
Port   Interface                      Status    Interface   Group
Max
Group                                           Bandwith    Bandwidth

1      GigabitEthernet6/0/1           up        1G
1      GigabitEthernet6/0/2           up        1G
1      TenGigabitEthernet6/0/29       up        10G
1      TenGigabitEthernet6/0/30       up        10G          10G

2      GigabitEthernet6/0/3           up        1G
2      GigabitEthernet6/0/4           up        1G
2      TenGigabitEthernet6/0/25       up        10G
2      TenGigabitEthernet6/0/26       up        10G          10G
<SNIP>
11     GigabitEthernet6/0/21          down      1G
11     GigabitEthernet6/0/22          down      1G
11     TenGigabitEthernet6/0/43       up        10G
11     TenGigabitEthernet6/0/44       up        10G          10G

12     GigabitEthernet6/0/23          down      1G
12     GigabitEthernet6/0/24          down      1G
12     TenGigabitEthernet6/0/45       up        10G
12     TenGigabitEthernet6/0/46       up        10G          10G

R4-C94-2041#
```

C9400-LC-48UX Port-Group – 10 Slot Chassis

- Bandwidth shared within port-group

- 8 port-group in the 10 slot chassis

- Modes: dynamic, performance and static

- For 10Gline rate performance:
  - Configure: "`hw-module subslot <slot#/0> mode performance`"
  - 5 Port @ line-rate, other ports are disabled

```
mac1#show platform hardware iomd 9/0 portgroups
Port    Interface                       Status     Interface  Group Max
Group                                              Bandwith   Bandwith

1       GigabitEthernet9/0/1            admindown  1G
1       GigabitEthernet9/0/2            admindown  1G
1       GigabitEthernet9/0/3            admindown  1G
1       TenGigabitEthernet9/0/25        admindown  10G
1       TenGigabitEthernet9/0/26        admindown  10G
1       TenGigabitEthernet9/0/27        admindown  10G        10G

2       GigabitEthernet9/0/4            admindown  1G
2       GigabitEthernet9/0/5            admindown  1G
2       GigabitEthernet9/0/6            admindown  1G
2       TenGigabitEthernet9/0/28        admindown  10G
2       TenGigabitEthernet9/0/29        admindown  10G
2       TenGigabitEthernet9/0/30        admindown  10G        10G
<SNIP>
8       GigabitEthernet9/0/22           admindown  1G
8       GigabitEthernet9/0/23           admindown  1G
8       GigabitEthernet9/0/24           admindown  1G
8       TenGigabitEthernet9/0/46        admindown  10G
8       TenGigabitEthernet9/0/47        admindown  10G
8       TenGigabitEthernet9/0/48        admindown  10G        10G

mac1#
```

Power

	Normal	PS failure
Combined (Default)	Load sharing on all PSs	Load sharing on functional PSs
Redundant	Load sharing on active PSs Standby PS in output disabled	Standby PS becomes active System enters alarm state
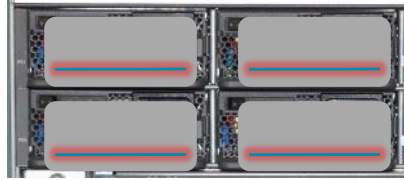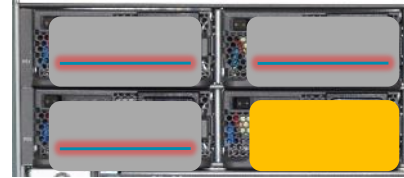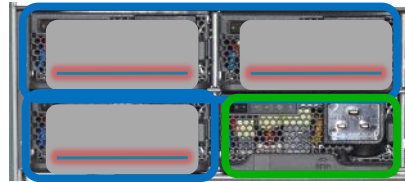
Failed PS   ACTIVE   STANDBY

Power Redundancy: N+N and N+1

- Default active is PS1-4 and standby is PS5-8

- Standby power slots are configurable

- Default active is PS1-7 and standby is PS8

- Standby power slot is configurable

Support mix of 3200W AC @ 220V input and 3200W DC

```
SW(config)#power redundancy-mode redundant ?
  N+N   Redundant N+N (N is active, N is standby)
  N+1   Redundant N+N (N is active, 1 is standby)
SW(config)#power redundancy-mode redundant N+1 ?
  <1-8>  standby slot in N+N mode
SWR(config)#
```

```
SW(config)#power redundancy-mode redundant ?
  N+N   Redundant N+N (N is active, N is standby)
  N+1   Redundant N+N (N is active, 1 is standby)
SW(config)#power redundancy-mode redundant N+1 ?
  <1-8>  standby slot in N+1 mode
SWR(config)#
```

| ACTIVE | STANDBY |

Power Priority

- All components in the system are assigned with power priority level

- Supervisors and Fan Tray has the same highest priority level

- Lower slot# has the higher power priority level by default if "`power supply autoLC shutdown`" is configured

- Configurable power priority for line card slots

```
C94(config)#power supply autoLC priority ?
  <1-7>   Physical slot number
  <cr>
C94(config)
```


Highest Priority

Configurable Priority
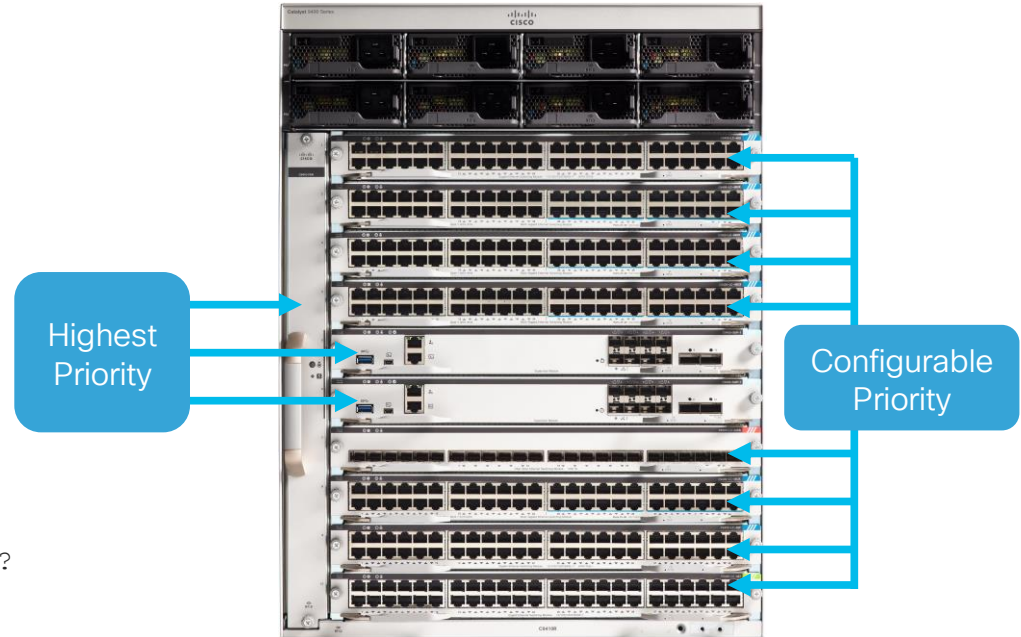
Forwarding – Unicast

You make networking **possible**

Flex Tables

Forwarding Resources

- MAC: 64K
- Host Route: 48K – 112K
- IGMP Groups: 16K
- LPM Route: 64K
- Multicast Route: 16K
- SGT: 16K

Feature Resources

- Security ACL: 18K
- QoS ACL: 18K
- Service ACL: 18K
  - PBR/NAT
  - Netflow ACL
  - SPAN
  - MACsec
  - CoPP
  - Tunnel
  - LISP

Netflow

Netflow Entries: 128K per ASIC

Unicast – within ASIC

Inter-ASIC Interface

PBC – Packet Buffers Complex

④

IQS

SQS AQM

EQS ⑤

Ingress Forwarding Controller (IFC) ③

X F C

Flexible Look up Tables (Shared Across Cores)

X F C

Egress Forwarding Controller (EFC)

②

⑥

1. Received, processed by MACSec and into FIFO

2. A copy to buffer and a copy to IFC

3. Goes through IFC, result descriptor send to PBC

4. Descriptor has local destination, PBC sends the info to EQS

5. EQS schedule PBC to send a copy to EFC and a copy to ReWrite (includes descriptor)

6. EFC sends results to ReWrite

7. Rewrite the packet and send though the egress FIFO

Encryption Engine

Recirculation Engine

ReWrite Engine ⑦

Ingress FIFO

Egress FIFO

① MACSEC

MACSEC

Network Interfaces - Front Panel Ports + CPU + Network Redundant Uplinks (NRU)

Unicast – Across ASICs on Input

Inter-ASIC Interface

PBC – Packet Buffers Complex

5 IQS (Q)

4

SQS AQM (Q) (Q)

EQS

3

Ingress Forwarding Controller (IFC)

X F C

Flexible Look up Tables

(Shared Across Cores)

X F C

Egress Forwarding Controller (EFC)

1. Received, processed by MACSec and into FIFO

2. A copy to buffer and a copy to IFC

3. Goes through IFC, result descriptor send to PBC

4. Descriptor has remote destination, PBC sends the info to IQS

5. IQS schedule PBC to send the packet with descriptor to Stack Interface

2

Encryption Engine

Recirculation Engine

ReWrite Engine

Ingress FIFO

Egress FIFO

1 MACSEC

MACSEC

Network Interfaces - Front Panel Ports + CPU + Network Redundant Uplinks (NRU)

Unicast – Across ASICs on Output

Inter-ASIC Interface 6

PBC – Packet Buffers Complex

IQS

SQS AQM

Q Q

EQS 7

Ingress Forwarding Controller (IFC)

Flexible Look up Tables (Shared Across Cores)

Egress Forwarding Controller (EFC)

X F C

X F C

8

6. PBC received the frame and sends the info to EQS

7. EQS schedule PBC to send a copy to EFC and a copy to ReWrite (includes descriptor)

8. EFC sends results to ReWrite

9. Rewrite the packet and send out though the egress FIFO

Encryption Engine

Recirculation Engine

Ingress FIFO

Egress FIFO

9

MACSEC

MACSEC

Network Interfaces - Front Panel Ports + CPU + Network Redundant Uplinks (NRU)

Multicast

You make networking **possible**

Multicast Features

Features	Catalyst 9K
IGMP (Internet Group Management Protocol)v1, v2, v3	Yes
IGMP snooping (v1, v2, v3)	Yes
MLD (Multicast Listener Discovery) v1, v2	Yes
MLD snooping (v1, v2)	Yes
PIM (Protocol Independent Multicast) SM (Sparse Mode)	Yes
PIM Dense Mode	Yes
PIM SSM (Source Specific Mode)	Yes
PIM Bi-Dir	Roadmap

Multicast Key Components

Inter-ASIC Interface

PBC – Packet Buffers Complex

SQS AQM

Q Q

EQS

IQS

- Schedule the packet for Stack Interface (Remote).
- Notify PBC when it no longer needs the packet

- Receive result descriptor from IFC.
- Notify IQS and/or EQS
- Remove the packet after notification from IQS and/or EQS

- Generate a list of egress ports
- Schedule the packet for each egress port
- Notify PBC when it no longer needs the packet.

Encryption Engine

Recirculation Engine

Ingress FIFO

Egress FIFO

MACSEC

MACSEC

Network Interfaces - Front Panel Ports + CPU + Network Redundant Uplinks (NRU)

Multicast - Egress Local

Inter-ASIC Interface

Only a single copy of packet in buffer memory during replication

PBC – Packet Buffers Complex 4

SQS AQM
Q Q
EQS 5

IQS

Ingress Forwarding Controller (IFC) 3

X F C

Flexible Look up Tables
(Shared Across Cores)

X F C

Egress Forwarding Controller (EFC)

1. Received, processed by MACSec and into FIFO

2. A copy to buffer and a copy to IFC

3. Goes through IFC, result descriptor send to PBC

4. Descriptor has local destination, PBC sends the info to EQS

5. AQM within EQS generate the list of egress port based on descriptor, schedule for each egress port

6. For each egress port, frame goes though the EFC, ReWrite and Egress FIFO

2

6

ReWrite Engine

Encryption Engine

Recirculation Engine

Ingress FIFO

Egress FIFO

1 MACSEC

MACSEC

Network Interfaces - Front Panel Ports + CPU + Network Redundant Uplinks (NRU)

Multicast – Egress Remote on Input

Inter-ASIC Interface

PBC – Packet Buffers Complex

IQS

SQS   AQM

EQS

1. Received, processed by MACSec and into FIFO

2. A copy to buffer and a copy to IFC

3. Goes through IFC, result descriptor send to PBC

4. Descriptor has remote destination, PBC sends the info to IQS

5. IQS schedule PBC to send the packet with descriptor to Stack Interface

Ingress Forwarding Controller (IFC)

Flexible Look up Tables

(Shared Across Cores)

Egress Forwarding Controller (EFC)

Encryption Engine

Recirculation Engine

ReWrite Engine

Ingress FIFO

Egress FIFO

MACSEC

MACSEC

Network Interfaces - Front Panel Ports + CPU + Network Redundant Uplinks (NRU)

Descriptor can contains both local and remote destinations

Multicast – Egress Remote Output

Inter-ASIC Interface

Replication done on egress => Efficient use of BW

PBC – Packet Buffers Complex

IQS

SQS   AQM

Q   Q

EQS

⑥

⑦

Ingress Forwarding Controller (IFC)

Flexible Look up Tables

(Shared Across Cores)

XFC

XFC

Egress Forwarding Controller (EFC)

⑧

6. PBC received the frame and sends the info to EQS

7. AQM within EQS generate the list of egress ports based on descriptor, schedule for each egress port

8. For each egress port, frame goes though the EFC, ReWriet and Egress FIFO

Encryption Engine

Recirculation Engine

Ingress FIFO

Egress FIFO

MACSEC

MACSEC

Network Interfaces - Front Panel Ports + CPU + Network Redundant Uplinks (NRU)

ACL

You make networking **possible**

Cisco live!

ACL Types and Features

- Security ACL (MAC, IPv4, and IPv6)
  - PACL: ACL enabled under L2 interface
  - VACL: ACL enabled for L2 VLAN traffic
  - RACL: ACL enabled for routed traffic
  - GACL: ACL for Clients group
  - SGACL: ACL for CTS/SGT
  - WCCP-Egress
- ACL for QoS classification and Policing (including CoPP)
- Service ACLPolicy
  - PBR/NAT/WCCP-Ingress
  - Netflow ACL
  - SPAN
  - MACSec
  - User ACL to redirect traffic (CoPP)
  - Tunnel
  - LISP

Security ACLs

```
interface gigabitethernet1/1
  ip access-group PACL-1 in
  ip access-group PACL-2 out
  switchport access vlan 100
vlan access-map VACL-map
  match ip address VACL-1
  action forward
vlan filter VACL-map vlan-list 100
interface Vlan100
  no shutdown
  ip access-group RACL-1 in
  ip access-group RACL-2 out
  ip address 100.1.1.1/24
```

PACL: Direction indicated in CLI

VACL: VLAN ACL for Both Ingress and Egress; Apply to Pre- and Postrouted Traffic

RACL: Routed ACL; Apply Only to Traffic That Requires Routing; Direction Indicated in the CLI

ACL Type	Attach Point	Direction
PACL	L2 interface	IN OUT
VACL	VLAN	Always IN & OUT
RACL	L3 interface, L3 PortChannel, sub-interface SVI	IN OUT

Security ACL Processing Order and Priority

FYI

- The following is a conceptual illustration. In the ASIC the lookup for different types of ACL takes place concurrently.

- A packet is dropped if it hits the deny rule in any of these types of ACLs.

- RACL is applied only to traffic that is L3 forwarded.

Ingress Processing

GACL In → PACL In → VACL In → RACL In →

Within ASIC Or Cross ASICs

Egress Process

RACL Out → VACL Out → PACL Out → GACL Out → SGACL Out
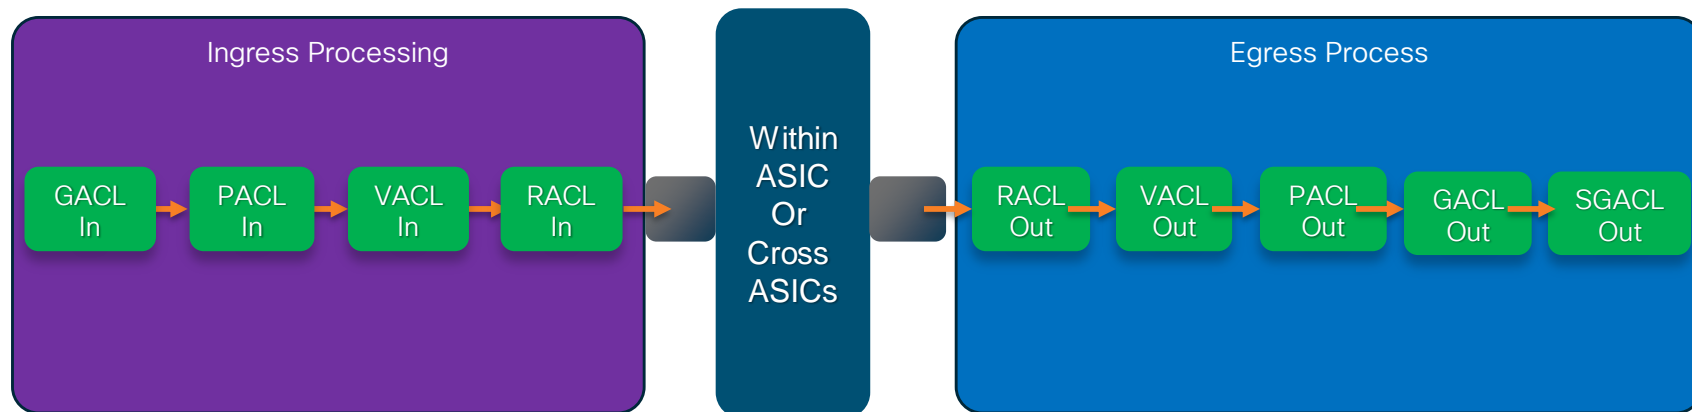
Sharing ACE

- Each ACL policy is reference by a label.
- Same ACL policy (Security ACL like GACL, PACL, VACL, RACL and QoS ACL) is applied to multiple interfaces or VLAN, it uses the same label.
- Label sharing is within the core
- Ingress and egress use different label

```
ip access-list extended ip-list-1
  deny   ip 100.1.1.0 0.0.0.255 200.1.1.0 0.0.0.255
  deny   ip 100.1.2.0 0.0.0.255 200.1.2.0 0.0.0.255
  deny   ip 100.1.3.0 0.0.0.255 200.1.3.0 0.0.0.255
  permit ip any any

interface GigabitEthernet1/0/1
  ip access-group ip-lists-1 in
interface GigabitEthernet1/0/2
  ip access-group ip-lists-1 in
interface GigabitEthernet1/0/3
  ip access-group ip-lists-1 in
```

Label	
101	Gi1/0/1
101	Gi1/0/2
101	Gi1/0/3
101	IPv4 ACL ip-list-1 in

Resource Utilization

```
C9400#show platform hardware fed active fwd-asic resource tcam utilization
CAM Utilization for ASIC Instance [0] Table          Max Values          Used Values
--------------------------------------------------------------------------------
 Unicast MAC addresses                                65536/512           17/21
 IGMP and Multicast groups                            8192/512            0/0
 L2 Multicast groups                                  8192/512            0/0
 Directly or indirectly connected routes              16384/65024          0/14
 NAT/PAT SA address and Port                              0                0
 QoS Access Control Entries                           18432                0
 Security Access Control Entries                      18432               114
 Ingress Netflow ACEs                                 1024                 8
 Policy Based Routing ACEs                            2048                 0
 Egress Netflow ACEs                                  2048                 0
 Input Microflow policer ACEs                          512                 0
 Output Microflow policer ACEs                        1024                 7
 Flow SPAN ACEs                                       1024                 0
 Control Plane Entries                                1024                200
 Tunnels                                              1024                17
 Lisp Instance Mapping Entries                        1024                 3
 Input Security Associations                           512                 0
 Output Security Associations and Policies             512                 5
 SGT_DGT                                              8192/512            0/0
 CLIENT_LE                                            4096/256            0/0
 INPUT_GROUP_LE                                       1024                 0
 OUTPUT_GROUP_LE                                      1024                 0
 Macsec SPD                                            256                 2
```
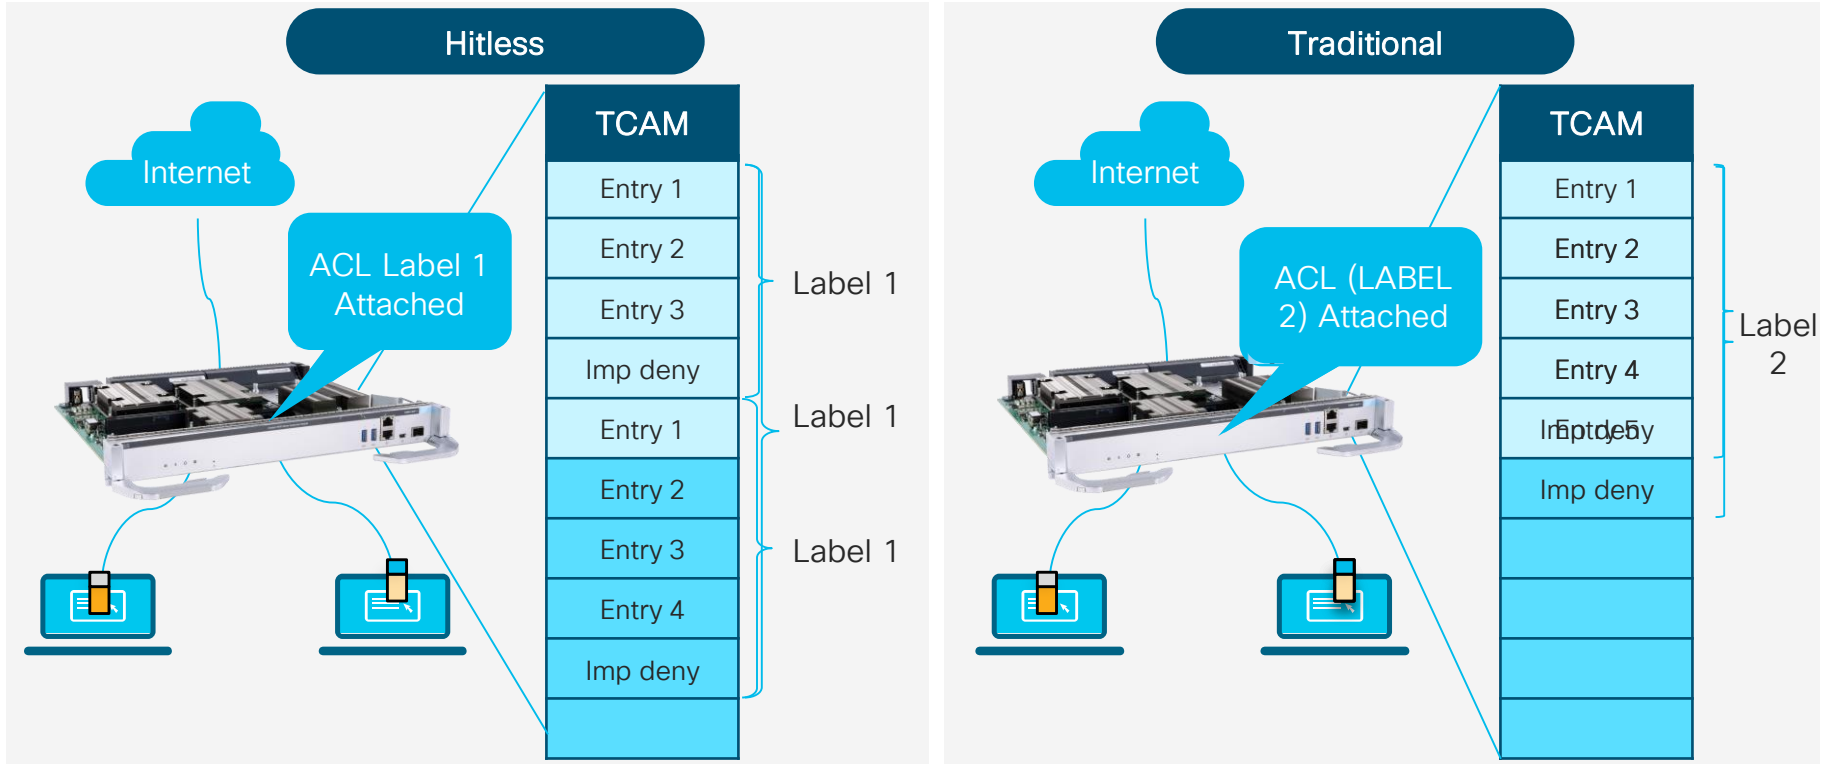
ACL Programming Scope

- RACL, VACL, SGACL: on all cores and ASICs.

- VACL: programmed for both ingress and egress (IPv6 VACL takes 4 entries – 2 for ingress and 2 for egress).

- PACL and GACL: only programmed on the core that has the port. (If L4OP is part of the PACL/GACL, then it will be programmed in both cores)

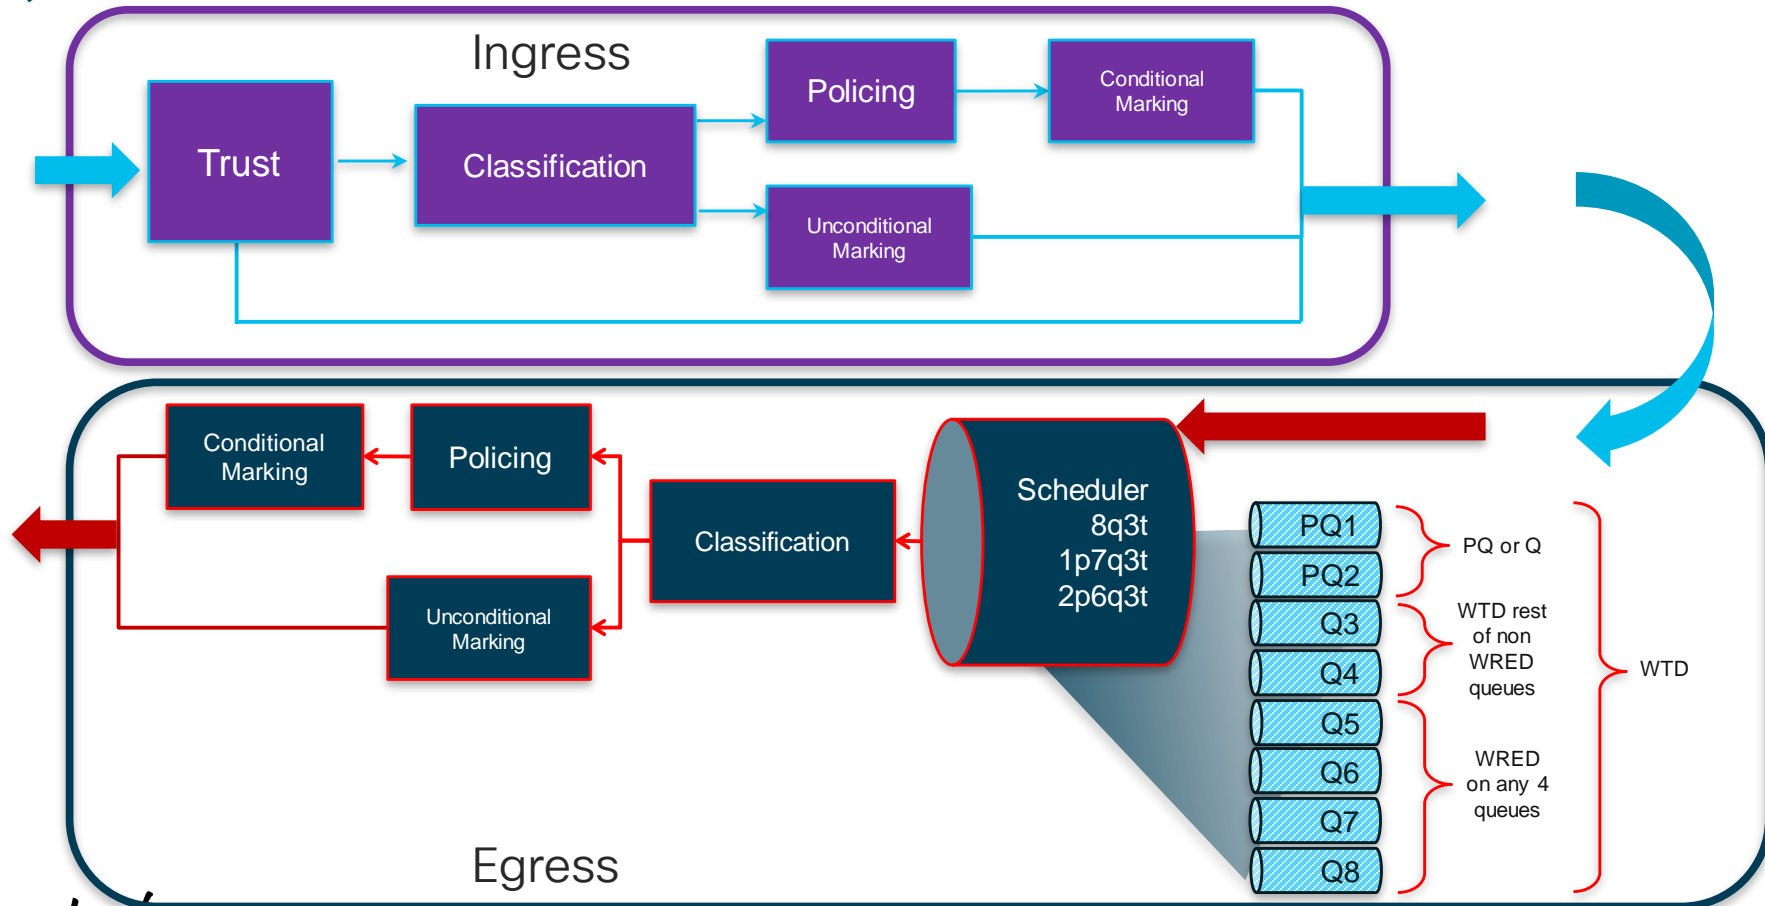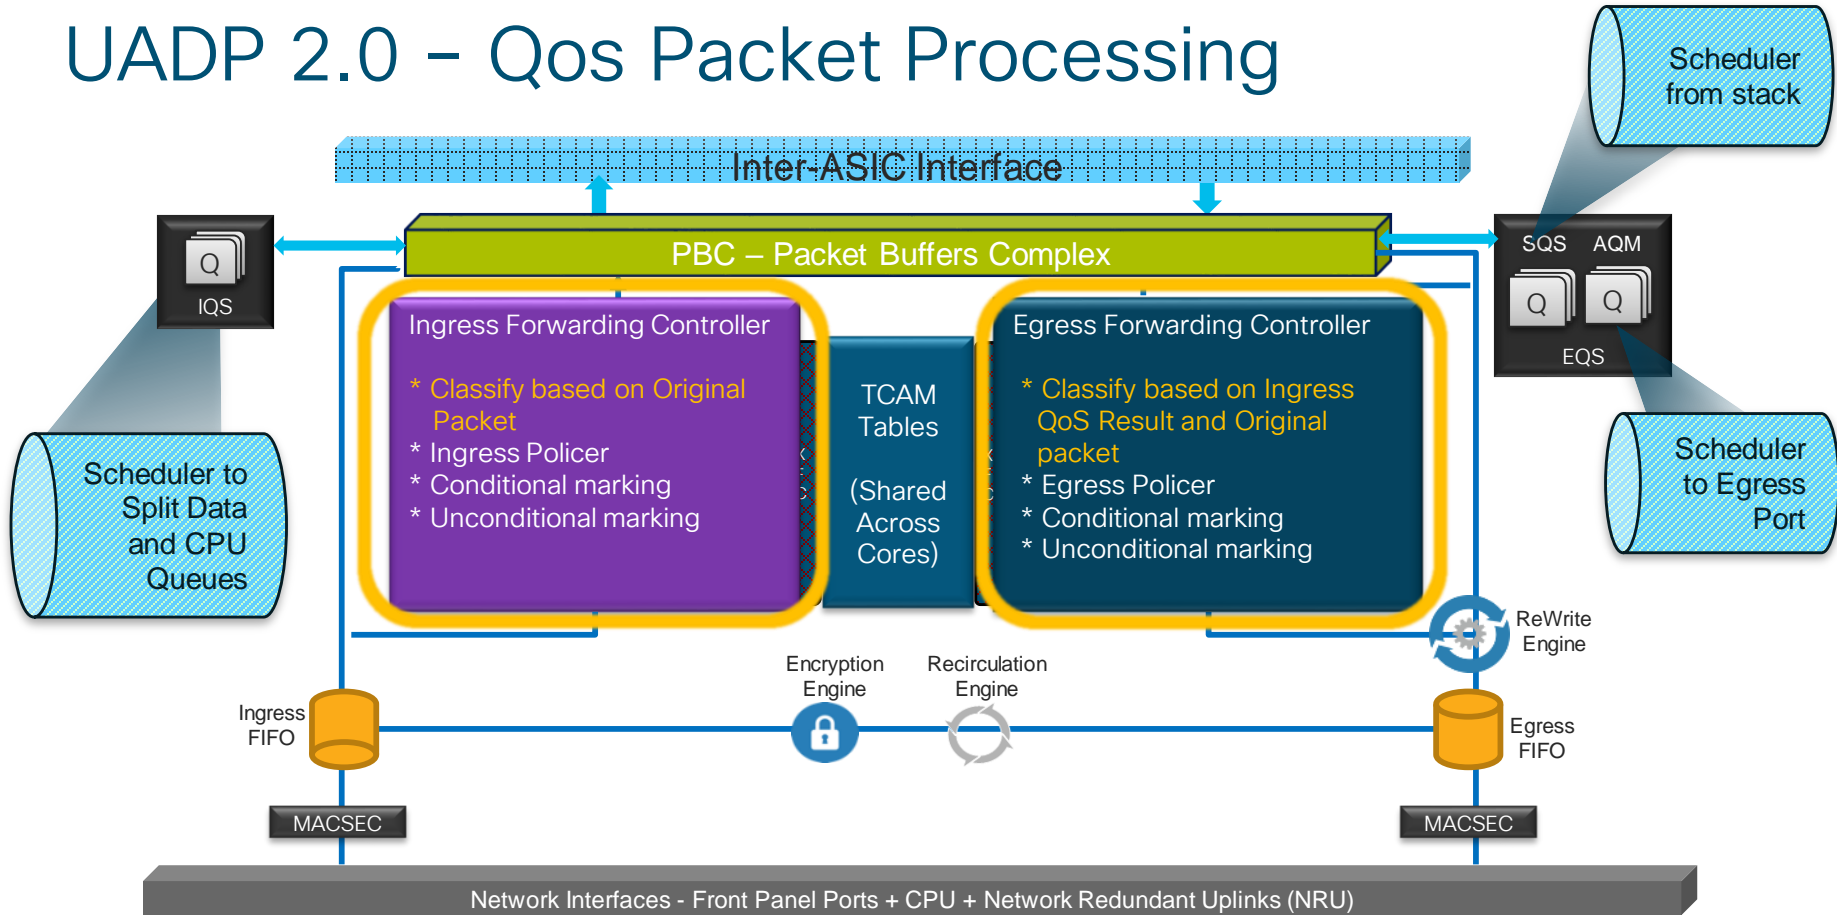# Cisco Catalyst 9400 Series- hitless TCAM update
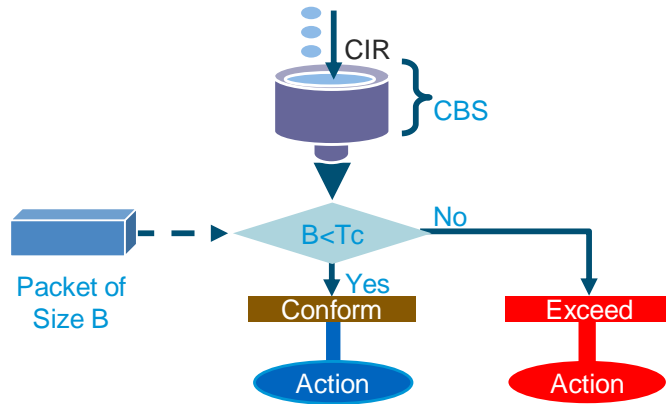
QoS

You make networking **possible**

Cisco live!

QoS Fundamental

UADP 2.0 – Qos Packet Processing

Inter-ASIC Interface

PBC – Packet Buffers Complex

Scheduler from stack

IQS

SQS AQM

Q Q

EQS

Ingress Forwarding Controller

* Classify based on Original Packet
* Ingress Policer
* Conditional marking
* Unconditional marking

TCAM Tables

(Shared Across Cores)

Egress Forwarding Controller

* Classify based on Ingress QoS Result and Original packet
* Egress Policer
* Conditional marking
* Unconditional marking

Scheduler to Split Data and CPU Queues

Scheduler to Egress Port

ReWrite Engine

Encryption Engine

Recirculation Engine

Ingress FIFO

Egress FIFO

MACSEC

MACSEC

Network Interfaces - Front Panel Ports + CPU + Network Redundant Uplinks (NRU)

Policing

1 Rate 2 Color	2 Rate 3 Color


```
police cir 100000000 bc 3125000 conform-
action set-dscp-transmit af41 exceed-action
drop
```
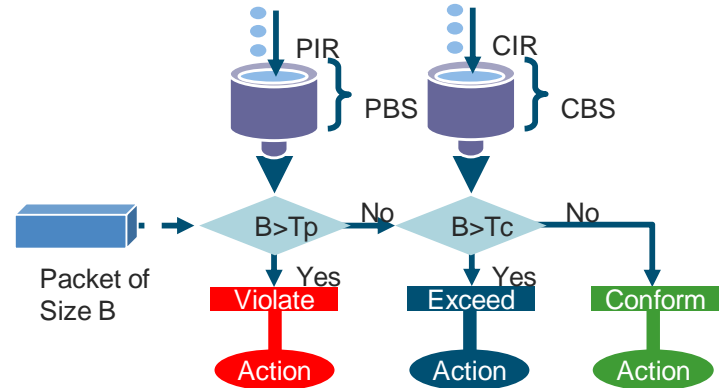
```
police cir percent 10 pir percent 50
conform-action transmit exceed-action set-
dscp-transmit af11 violate-action drop
```

Packet Buffer

Types:

- Ingress Buffers: For ingress queues and scheduler (IQS) (dedicated and shared)

- Egress Stack Buffers: For Egress Stack Queues and Scheduler (SQS) (dedicated and shared)

- Egress Port Buffers: For Egress port queues (AQM)

- Temporary Buffers: For packet from the Stack or CPU.

Allocation:

- Use dedicated first then shared.

10 MB Egress

1.5 MB FIFO

0.4MB– 1.5 MB Ingress

1.5MB – 3.5 MB Stack
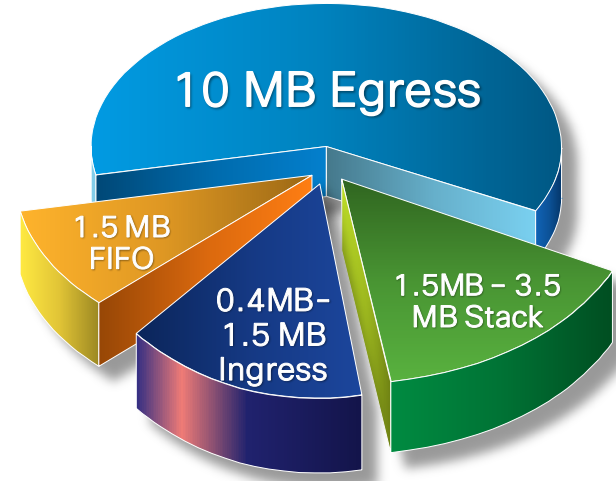
16MB per Core
32MB per ASIC

2-Level HQoS
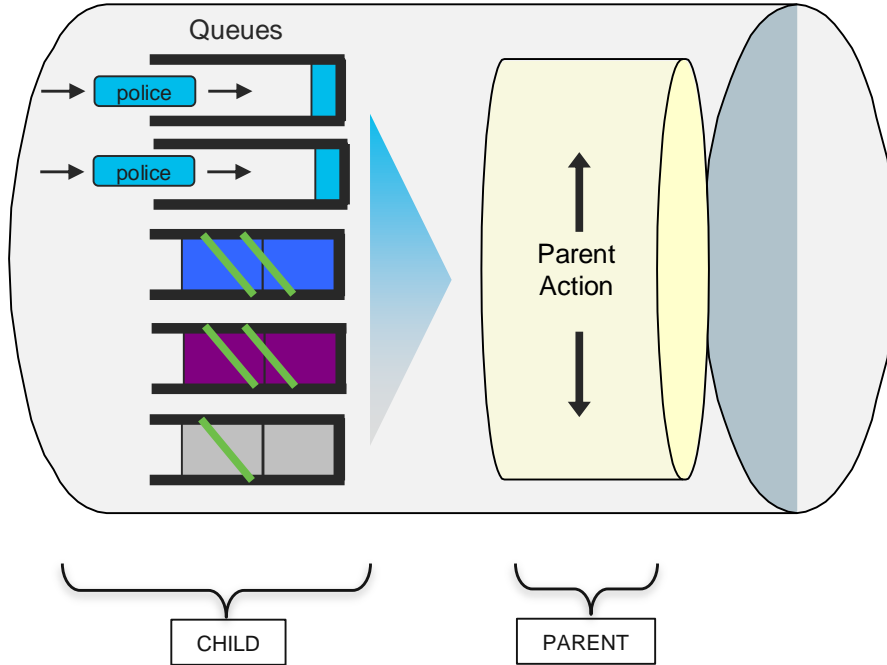
Child	Parent
BW, Policing	Shape
Marking	Policing
Policing	Marking

Security

You make networking **possible**

Cisco Catalyst 9000 Platform Trustworthy Systems

Design/Develop	Plan/Order	Source	Make	Quality	Delivery	Service/End of Life (EOL)

Physical security practices + security technology innovations + logical security processes

PnP SUDI support
Two-way trust

Image signing
Authentic OS

Hardware authenticity
Genuine hardware

Secure boot
Boot sequence check

Integrity verification
Malware protection

Runtime defenses
64-bit ASLR

Cisco® trustworthy systems use industry best practices to help ensure full development lifecycle integrity and end-to-end security

MACsec
Hop-by-hop encryption via 802.1AE
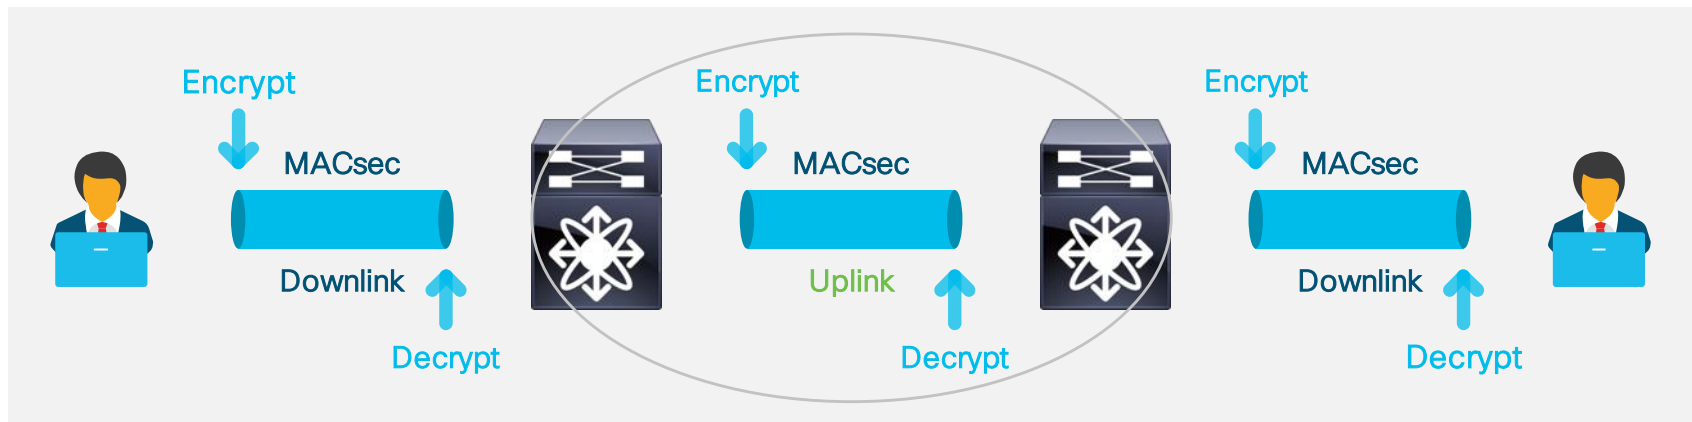
- Packets are encrypted on egress; decrypted on ingress
- Offers line-rate encryption on all ports and speeds (1G, 10G, and 40G) *
- Transparent to all upper-layer protocols
- Supports switch-to-switch and switch-to-host MACsec
- 256-bit MACsec-capable between switch to switch
- Manual or 802.1X modes supported

* C9400-Sup1-XL-Y doesn't support MACSec on the uplink ports

ETA - Finding Malicious Activity in Encrypted Traffic

Catalyst® 9400

Cisco Stealthwatch®

NetFlow

Cognitive Analytics

Telemetry for encrypted malware detection and cryptographic compliance

'Biflow'

Malware detection and cryptographic compliance

Leveraged network	Faster investigation	Higher precision	Stronger protection
Enhanced NetFlow from Cisco's newest switches and routers	Enhanced analytics and machine learning	Global-to-local knowledge correlation	Continuous enterprise-wide compliance

High Availability

You make networking **possible**

Cisco *live!*

SSO – Stateful Switch Over

AP PC/Laptop IP Phone

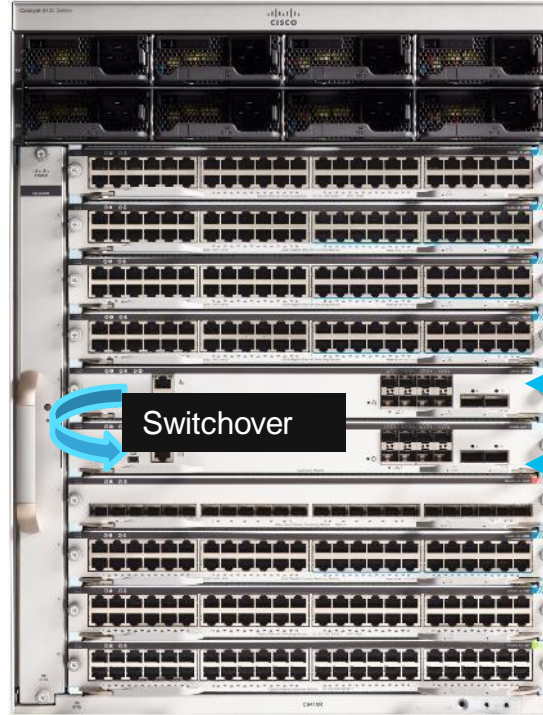**Links Stay up**
Continue Communication
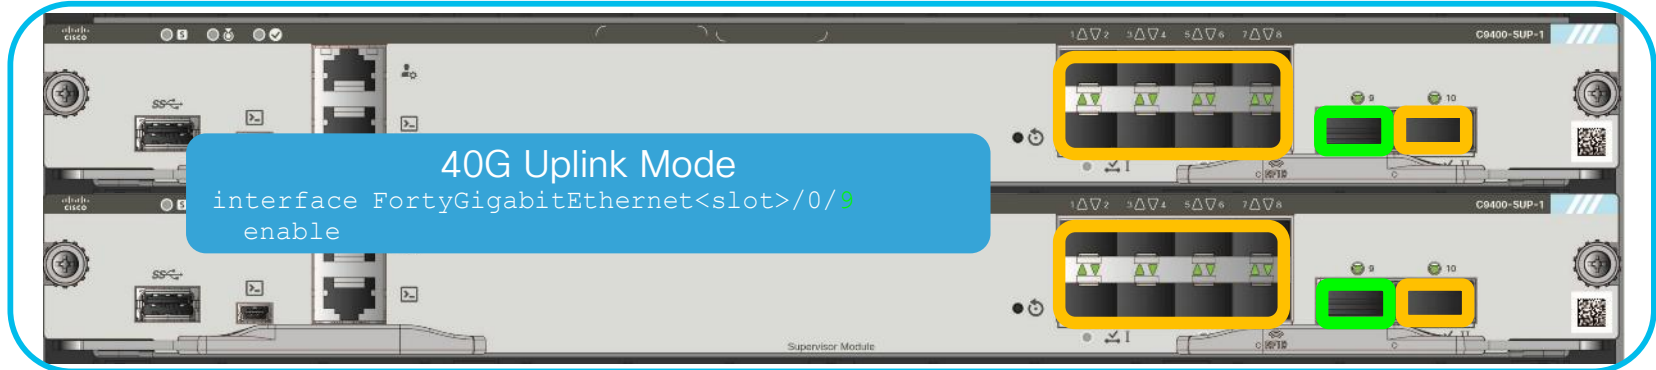
LED Lights
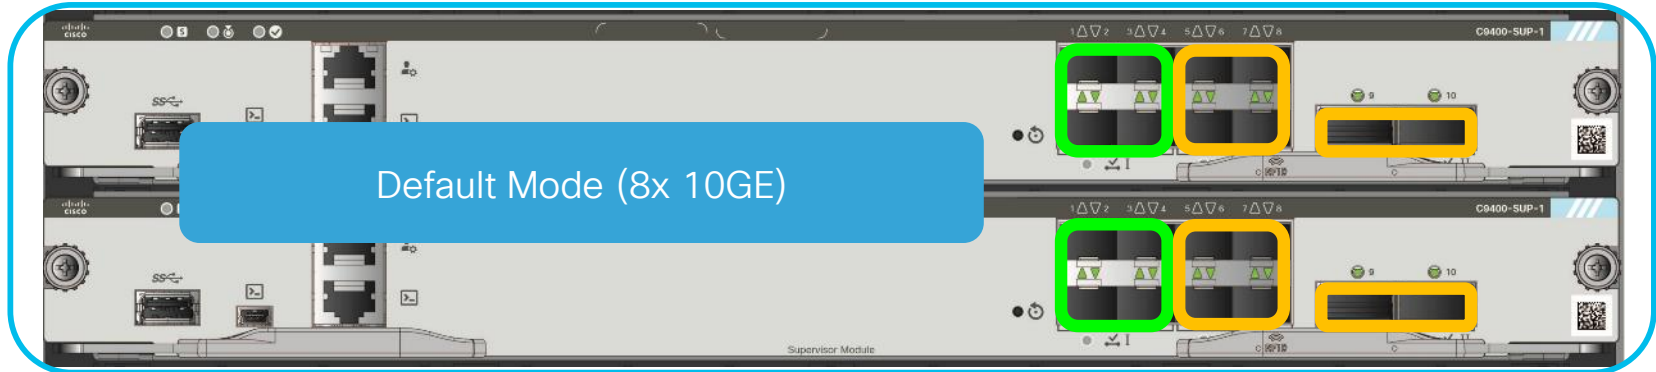
Lights Stay on

Sup Uplinks stay up

Less than 200ms traffic interruption

Switchover

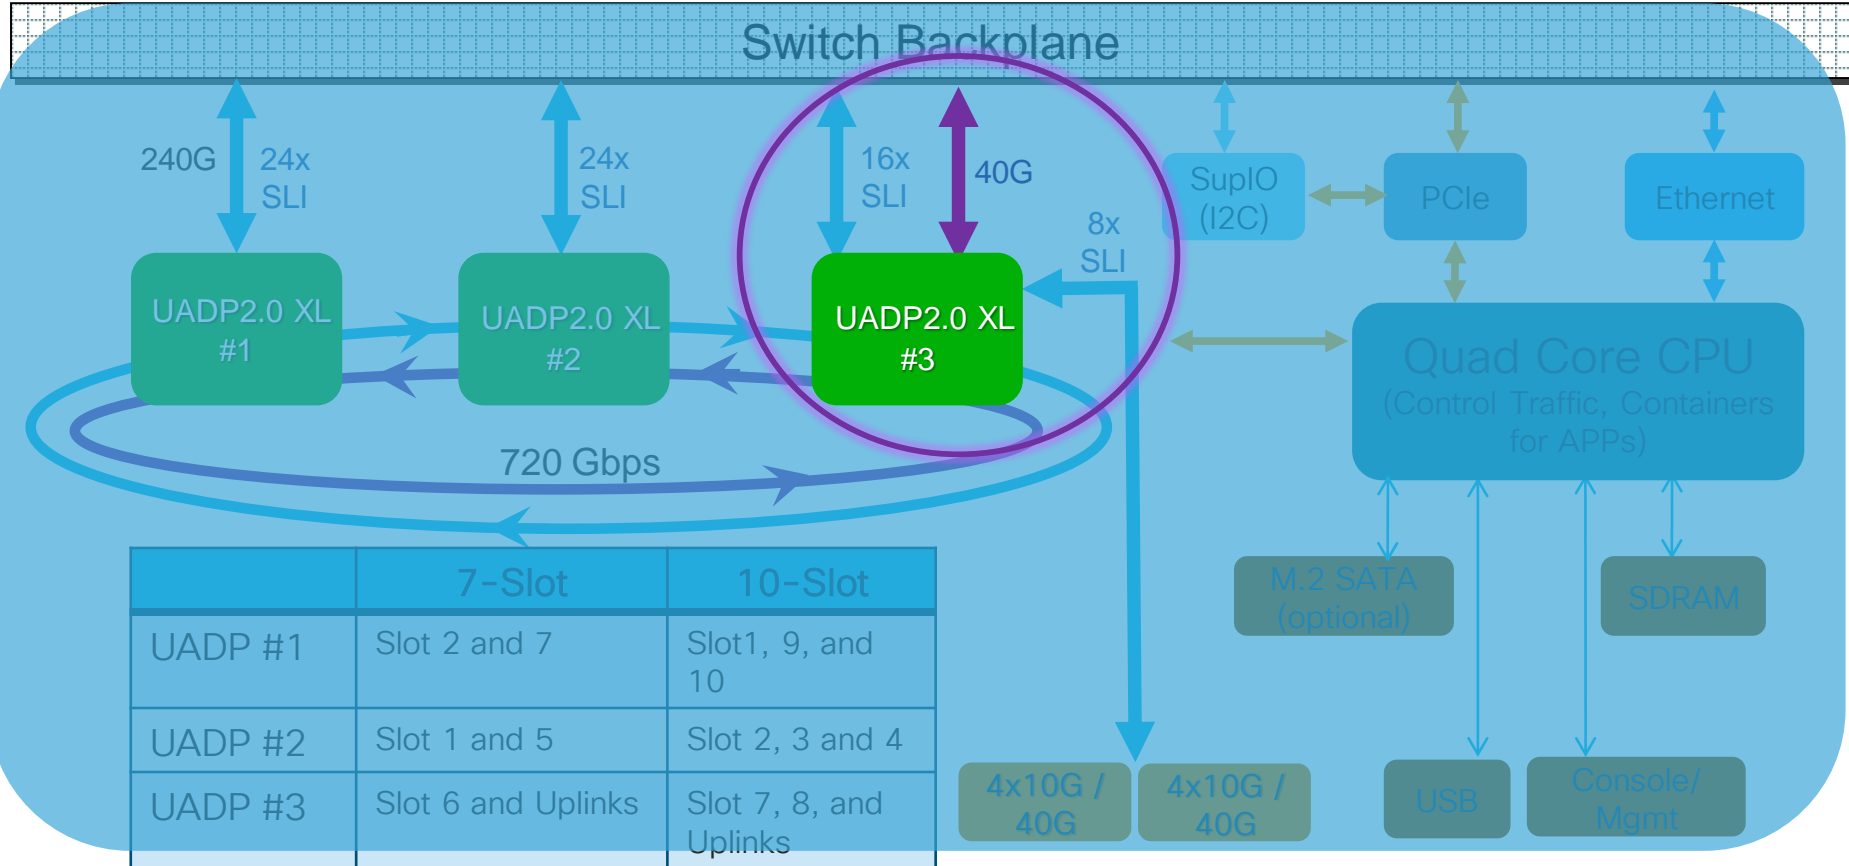Default Redundancy Mode

Continue Synchronization from Active to Standby

Sup-1 Uplink Redundancy – Dual Sups

Default Mode (8x 10GE)

40G Uplink Mode
```
interface FortyGigabitEthernet<slot>/0/9
  enable
```

Redundant Uplinks

	7-Slot	10-Slot
UADP #1	Slot 2 and 7	Slot1, 9, and 10
UADP #2	Slot 1 and 5	Slot 2, 3 and 4
UADP #3	Slot 6 and Uplinks	Slot 7, 8, and Uplinks

Redundant Uplinks

Backplane

Active Sup

UADP #3

Cores

Network Redundant Unlink

Uplinks

Standby Sup

UADP #3

Cores

Network Redundant Unlink

Uplinks

These components stay operational when Sup reboots

Standby Uplinks stays up when standby Sup reboots
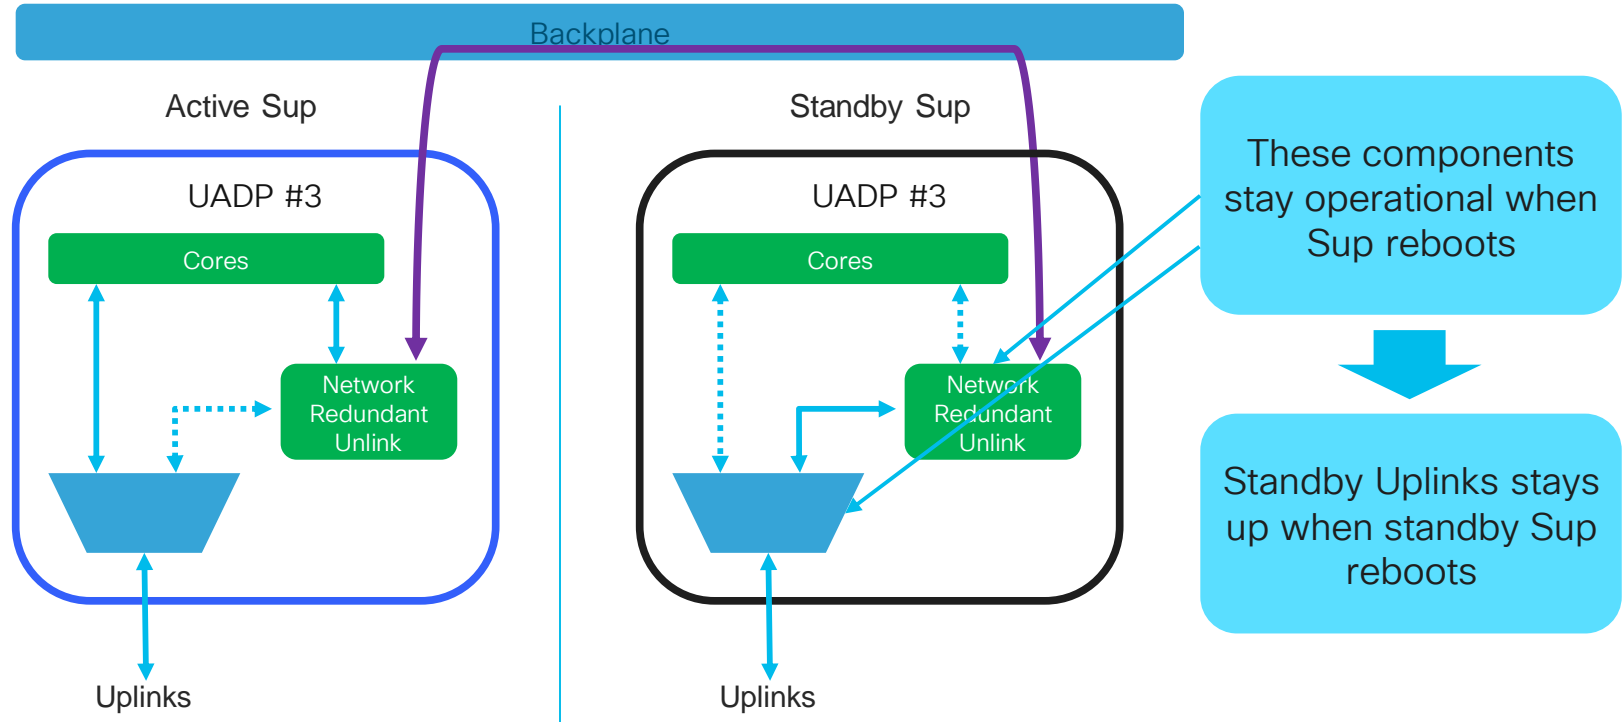
StackWise Virtual
Topology Comparisons

Traditional

HSRP or VRRP

LACP or PAGP

STP or MST

Access Switch

Access Switch Stack

SV - Physical

SVL

LACP or PAGP

Access Switch

Access Switch Stack

SV- Logical

MEC

Access Switch

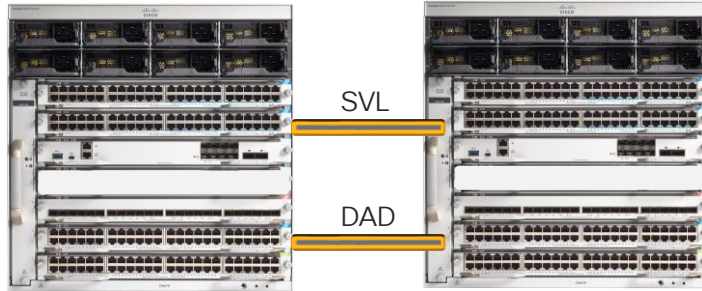**Access Switch Stack**
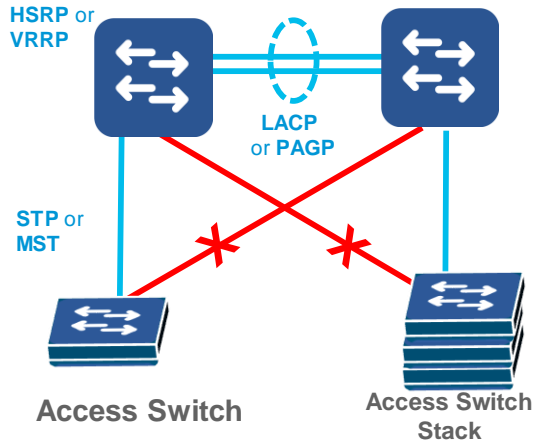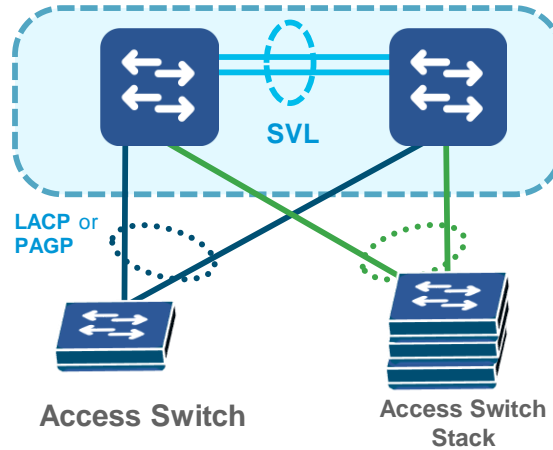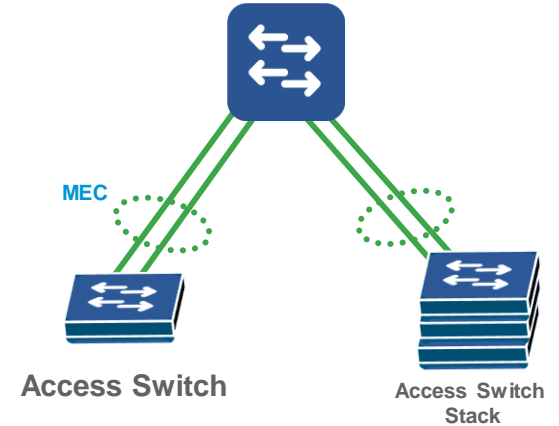
Benefits of StackWise Virtual

Simplify Operations by Eliminating STP, FHRP and Multiple Touch–Points

Double Bandwidth & Reduce Latency with Active-Active Multi-chassis EtherChannel (MEC)

Minimizes Convergence with Sub-second Stateful and Graceful Recovery (SSO/NSF)
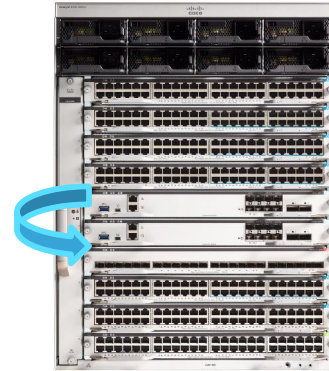
StackWise Virtual – C9400

SVL

DAD

StackWise Virtual is supported on all C9400 chassis and all Supervisors starting IOSXE 16.11.1

- SVL: StackWise Virtual Link
  - Same speed ports (10G or higher)
  - Up to 8 ports
- DAD: Dual Active Detection:
  - Fast Hello
    - Directly connected
    - Up to 4 links
  - Enhanced PAgP
    - EtherChannel with PAgP
    - Up to 4 port-channels

- In SVL mode, 2nd Supervisor is not supported in the chassis and will be powered off if inserted.

- A Distribution layer technology allowing stacking of 2 switches

- Same Chassis model and same supervisor model

- Supports flexible distances with support of all supported cables and optics

In Service Software Upgrade (ISSU)

- ISSU provides a mechanism

- to perform software upgrades and downgrades without taking the switch out of service

- Leverages the capabilities of NSF and SSO to allow the switch to forward traffic during Supervisor IOS upgrade (or downgrade)

- Key technology is the ISSU Infrastructure
  - Allows SSO between different software versions

Redundant Supervisors

SVL

40G/10G

DAD
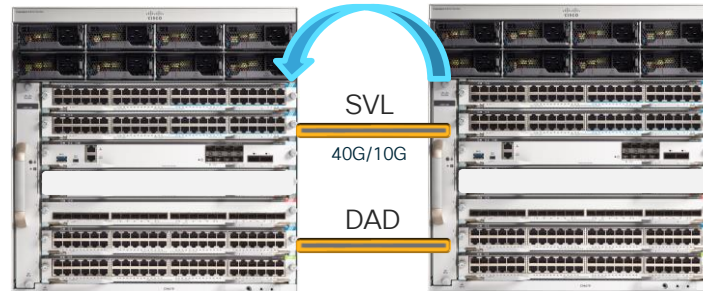
StackWise Virtual

Cisco Catalyst 9000 Series ISSU workflow

1. ISSU started; image is expanded on active and standby supervisors

Upgrade start

#install add

| V1 | S1 | Active |
| V2 | S2 | Standby |

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

🕐 Abort timer starts

2. Standby reloads with the new V2 image

| V1 | | S1 | Active |
| V1 → V2 | | S2 | Standby |

Upgrade complete

Expired abort timer will revert to Step 2 and then Step 1

🕐 Abort timer expired

5. ISSU complete

| V2 | S1 | Standby |
| V2 | S2 | Active |

🕐 Abort timer stopped

#install commit

4. 'Commit' keyword stops the abort timer

| V1 → V2 | S1 | Standby |
| V2 | S2 | Active |

3. # install activate <> issu

3. Auto-switchover causes S2 to become the new active and S1 reloads with the new V2 image

Cisco live!

Conclusion

You make networking **possible**

Cisco live!

Catalyst 9400

5KW PoE
Per slot

Redundancy
is now
Table-stakes

480G BW
per slot

4-Slot

7-Slot

10-Slot

Supervisor	Access Linecards	Core Linecards	Power Supply
Sup-1: 80G/Slot Access Optimized	24xmGig + 24xUPOE	24x 10G SFP+	3200W AC
Sup-1XL: 240G&120G/Slot Core Optimized	48xUPoE	48x1G SFP	2100W AC
Sup-1XL-Y: 240G&120G/Slot Core Optimized	48xPoE+	24x1G SFP	3200W DC
	48xData		

Catalyst 9K enables the New Era of Networking

X86 Based CPU

Larger Storage

Catalyst 9K Family

Usability & Operational Efficiency

UADP 2.0

IOS XE 16.x

Foundational Components

Encrypted Traffic Analytics (ETA)	Resiliency with Patching & GIR	Programmability & Automation
Standards Leadership .3bt, .3bz, etc	Real time Streaming Telemetry	SD Access
Trustworthy Systems	3rd Party App Hosting	Fabric Enabled Wireless

Catalyst 9K – Built to see you through Next Decade

Cisco Webex Teams

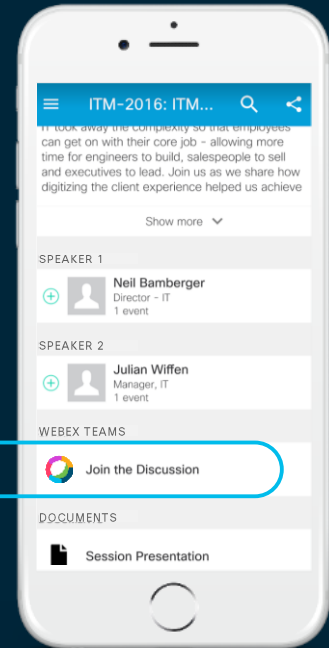## Questions?
Use Cisco Webex Teams to chat
with the speaker after the session

How

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

Webex Teams will be moderated
by the speaker until June 16, 2019.

cs.co/ciscolivebot#BRKARC-3873

Complete your online session evaluation

- Please complete your session survey after each session. Your feedback is very important.

- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on Thursday) to receive your Cisco Live water bottle.

- All surveys can be taken in the Cisco Live Mobile App or by logging in to the Session Catalog on ciscolive.cisco.com/us.

  Cisco Live sessions will be available for viewing on demand after the event at ciscolive.cisco.com.

Continue your education

Demos in the Cisco campus

Walk-in labs

Meet the engineer 1:1 meetings

Related sessions

NDA Roadmap Sessions at Cisco Live
Customer Connection Member Exclusive

Join Cisco's online user group to ...

Connect online with 29,000 peer and Cisco experts in private community forums

Learn from experts and stay informed about product roadmaps
- ➤ Roadmap sessions at Cisco Live
- ➤ Monthly NDA briefings

Give feedback to Cisco product teams
- ➤ Product enhancement ideas
- ➤ Early adopter trials
- ➤ User experience insights

Join online: www.cisco.com/go/ccp

NETWORKING ROADMAPS	SESSION ID	DAY / TIME
Roadmap: SD-WAN and Routing	CCP-1200	Mon 8:30 – 10:00
Roadmap: Machine Learning and Artificial Intelligence	CCP-1201	Tues 3:30 – 5:00
Roadmap: Wireless and Mobility	CCP-1202	Thurs 10:30 – 12:00

Join at the Customer Connection Booth
(in the Cisco Showcase)

Member Perks at Cisco Live
- Attend NDA Roadmap Sessions
- Customer Connection Jacket
- Member Lounge

Thank you

You make **possible**