You make **possible**

Catalyst 9600 Architecture

Kenny Lei
Technical Marketing Engineer
BRKARC-3010

Cisco live!
June 9-13, 2019 • San Diego, CA
#CLUS

Agenda

- Overview
- Architecture
- Forwarding
- Features (ACL, QoS, Security, High Availability)
- Catalyst 9600 Design
- Closing

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

Overview

You make networking **possible**

Cisco Catalyst 9600 Series Switches

Modular platform for Campus Core and Distribution

Cisco Catalyst 9600 Series
Chassis

Dual-serviceable fan tray

Built-in RFID

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

Modular power supplies

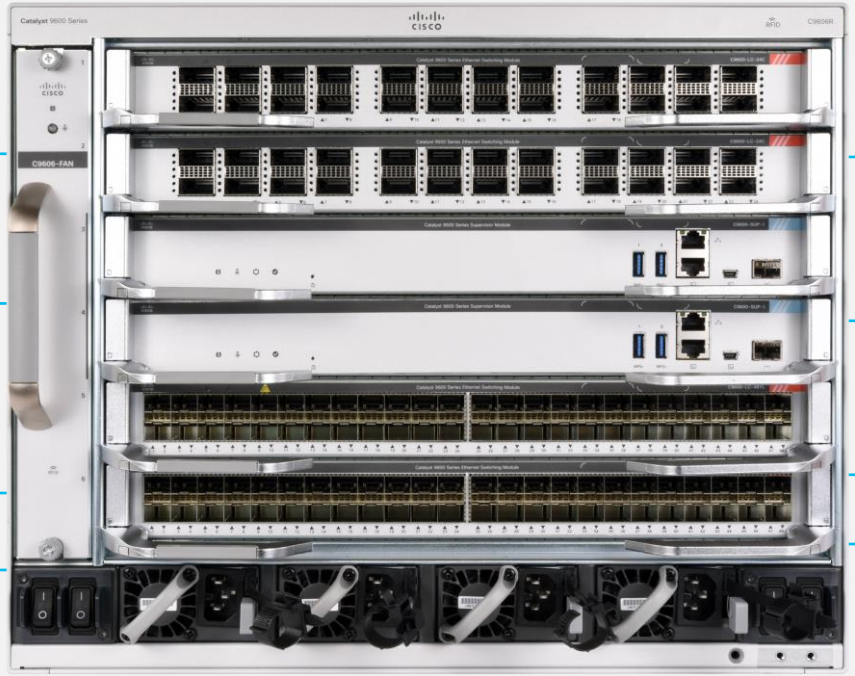

4 line card slots

2 supervisor slots (dedicated)

6.4 Tbps per slot from each supervisor slot

Dimensions (HxWxD inches) 13.95 x 17.4 x 16.1 (8RU)

Cisco Catalyst 9600 Series
C9606R chassis port density

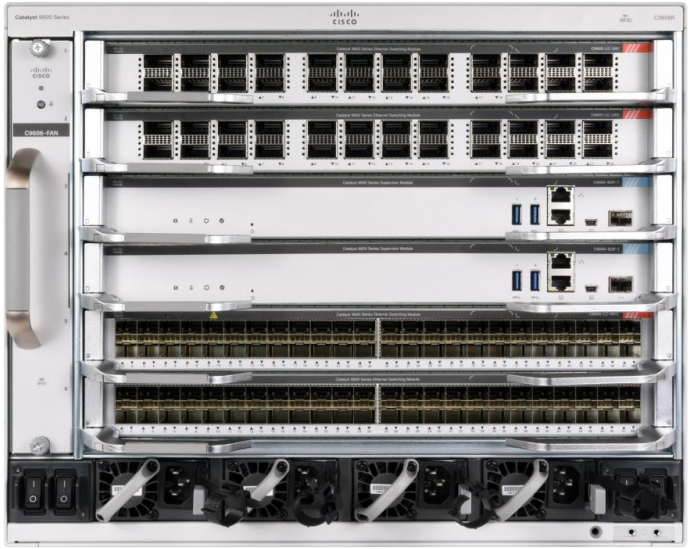

Port speed	Density with supervisor 1	Maximum chassis density
100G	48	128
40G	96	128
25G	192	192
10G	192	192
1G*	192	192

Line Rate non-blocking

*Roadmap

Cisco Catalyst 9600 Series
Supervisor 1

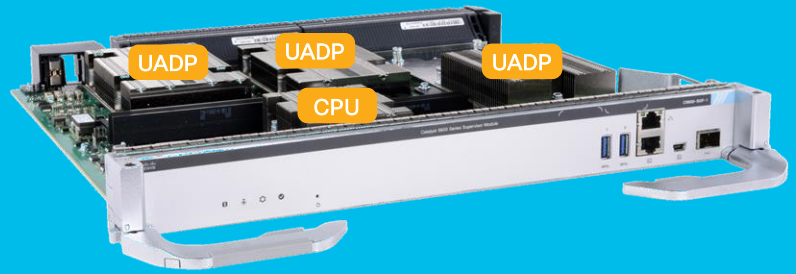

9.6 Tbps **3** Bpps	2.4 Tbps per slot
3x UADP 3.0 ASIC	**8 core** X86 CPU @2.0 Ghz
M.2 SATA SSD (optional: up to 1 TB)	16G DDR4 memory
Built-in RFID	Mgmt ports: copper and **fiber**
Blue Beacon	2x USB3 1x mini-B USB console

Cisco Catalyst 9600 Series
Line cards

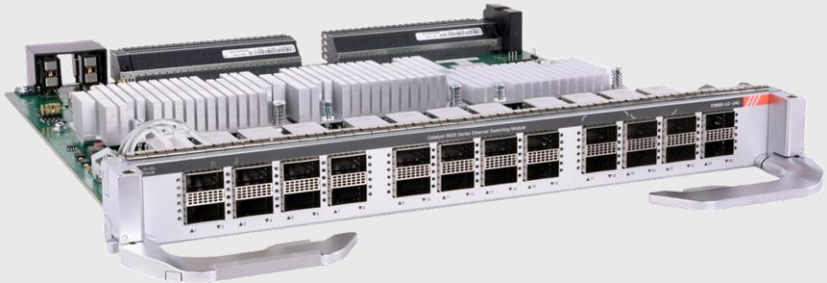

C9600-LC-24C – 100G/40G (fiber)

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

C9600-LC-48YL – 25G/10G/1G* (fiber)

- 48 ports
- SFP28/SFP+/SFP
- Supports 25G, 10G, and 1G

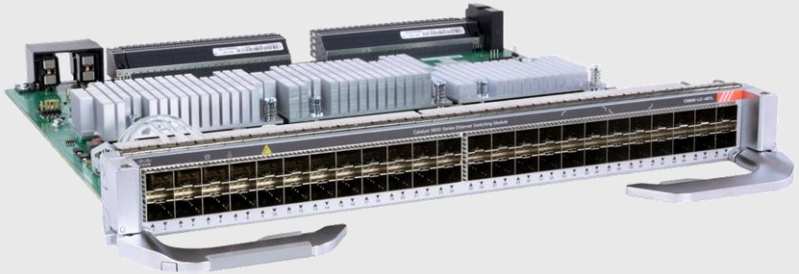

*Roadmap

The Y in the product ID (PID) indicates the hardware capability

Cisco Catalyst 9600 Series
100G/40G Line card – C9600-LC-24C

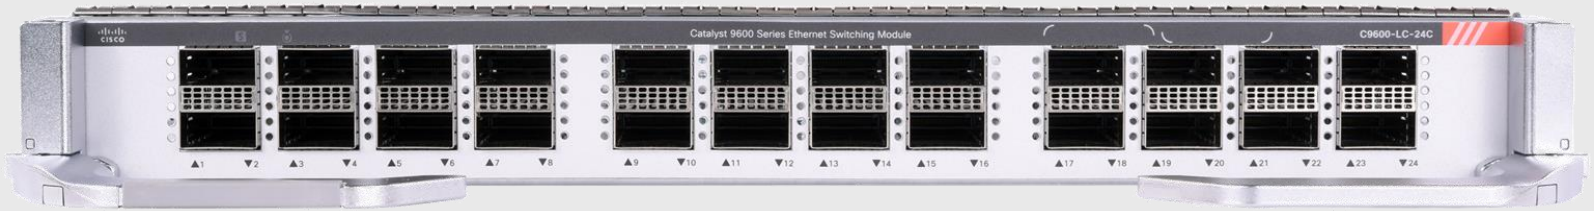

- All 24 ports are capable of 100G (QSFP28)/40G (QSFP+)
- Hardware-ready with QSA (for 1G/10G)
- With Supervisor Engine 1
  - 100G: Every 2 ports in a port-group. The odd number of ports can be 100G and the next even number port is disabled. (Maximum of 12x 100G, line rate with 187 byte or higher)
  - 40G – 24x 40G (line rate with 148 byte or higher)

C9600-LC-24C – Port Numbering with Supervisor Engine 1

- 40G numbering from 1 to 24
- 100G number from 25 to 48

40G port numbering

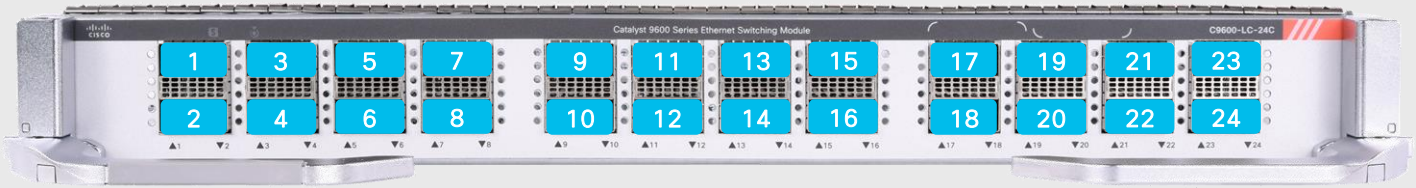

100G port numbering

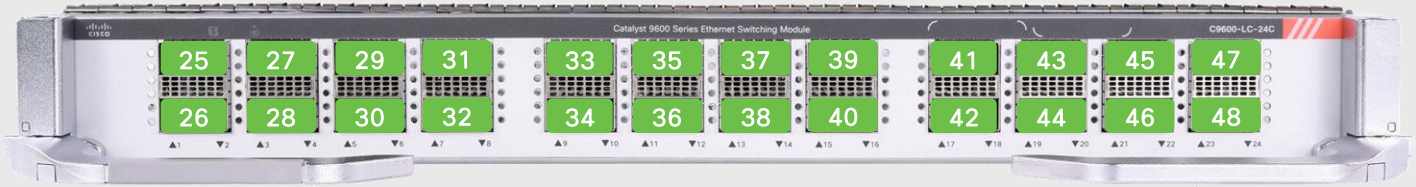

C9600-LC-24C with supervisor engine 1

- This line card appears in 40G mode by default
- Future supervisors can support 100G speed on all ports at the same time

Default mode
(all ports 40G)

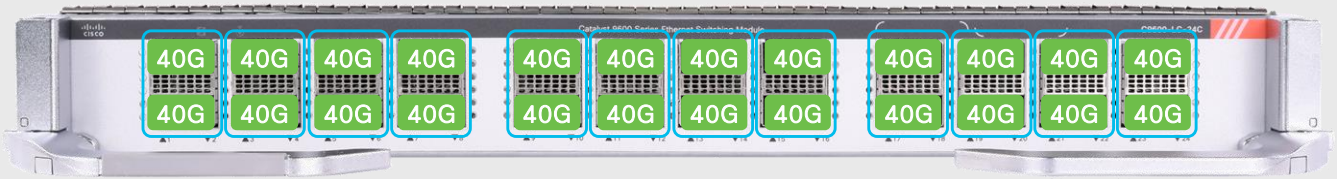

100G
configuration

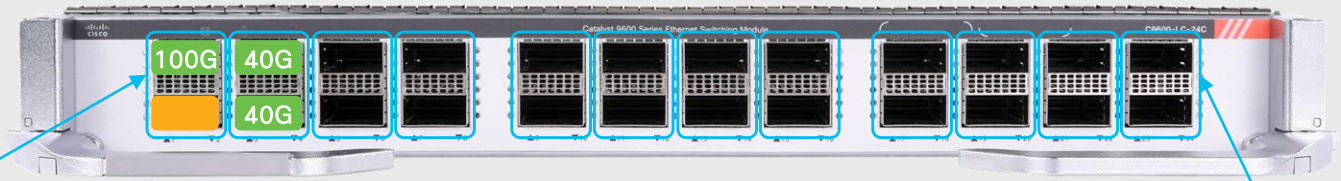

Fo<slot#>/0/1
Hu <slot#>/0/25

interface HundredGigE1/0/25 enable

- Enable Hu1/0/25 as 100G
- Disabled Fo1/0/1 and 1/0/2

Fo<slot#>/0/23
Hu <slot#>/0/47

Cisco Catalyst 9600 Series
QSA adapter CVR-QSFP-SFP10G support

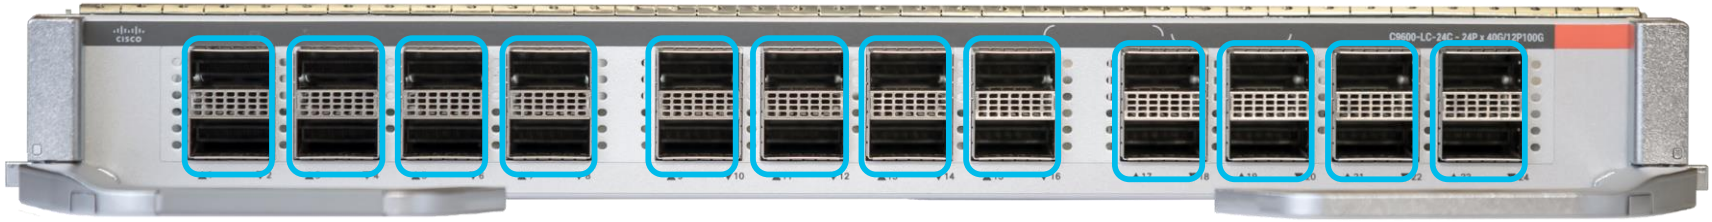

QSA on both top bottom 	QSFP on both top and bottom 	QSA on top and QSFP on bottom 	QSFP on top and QSA on bottom 
✓	✓	✓	✗

Cisco Catalyst 9600 Series
25G/10G/1G Line card – C9600-LC-48YL

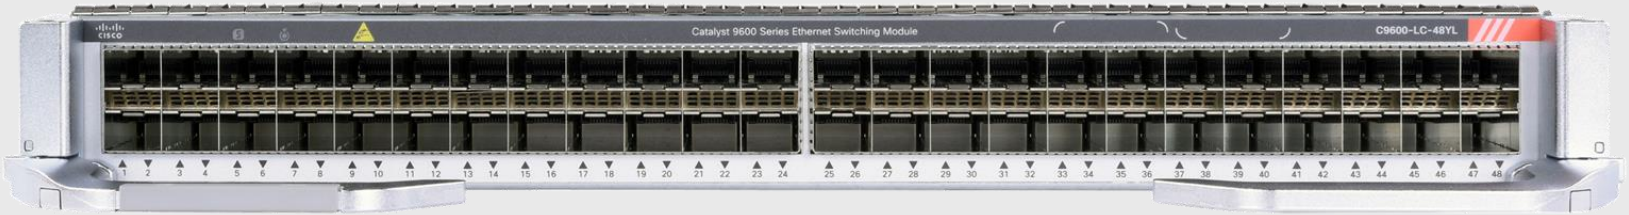

- All 48 ports support 25G/10G/1G
- Hardware capable of 10/100M
- Line rate with 25G/10G/1G (at 187 bytes for 25G; any packet size with 10G/1G)
- Any port, any supported speed
- Port reference is always ”TwentyFive<slot#>/0/<port#>” and port speed is auto-detected based on the inserted transceiver

Cisco Catalyst 9600 Series
Fan tray

- N+1 (8+1) fan redundancy
- 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)
- Airflow – side-to-side airflow

Fan tray hot-swappable needs to be done within 120 seconds

Cisco Catalyst 9600 Series
Power supplies

- Chassis has 4 slots for power supply
- Individual on/off switch for each power supply
- Supports a mix of AC (@220V) and DC power supplies

AC

- Supports both 110V and 220V input
- 2 KW output with 220V (1050W with 110V)
- Platinum rate power supply
- Redundant mode: Combined and N+1

DC

- Supports input range of -40V to -72V
- 2 KW output
- Platinum rate power supply
- Redundant mode: Combined and N+1

Cisco Catalyst 9600 Series
Power supply redundancy

	Normal operation	Power supply failure	AC and DC power supplies
			
Requirement	• AC: No mixing of 110V input and 220V input		• Mix DC and AC: AC input needs to be 220V
Operation	• Equal load sharing and all active	• Equal load sharing among the remaining power supplies	• Equal load sharing
Power budgeting	• Combined mode: Use all available power supplies for system budgeting • N+1 mode: Use N power supplies for system budgeting	• Combined mode: Line card can shut down if there isn't enough power • N+1 mode: Always enough power with single power supply outage	• Combined mode: Line card can shut down if there isn't enough power • N+1 mode: Always enough power with single power supply outage

Power priority

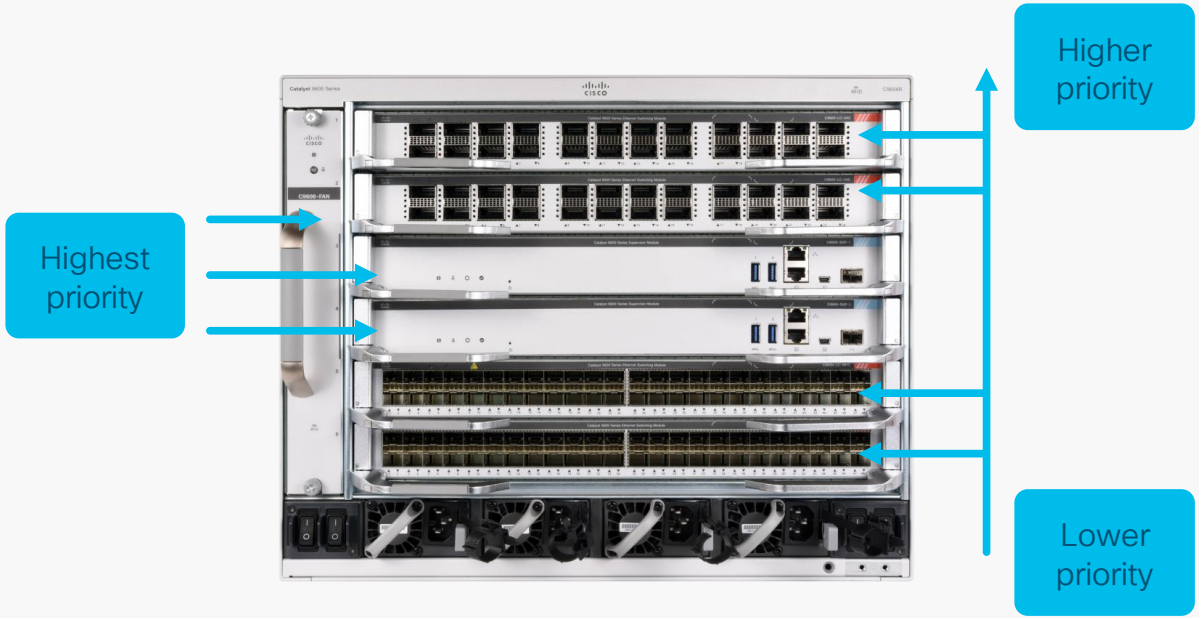

- All components in the system are assigned with a power priority level
- Supervisors and the fan tray have the same highest priority level
- Line cards with lower slot numbers have the higher power priority level by default
- User-configurable power priority for line card slots is on the roadmap

Architecture

You make networking **possible**

Cisco live!

Architecture

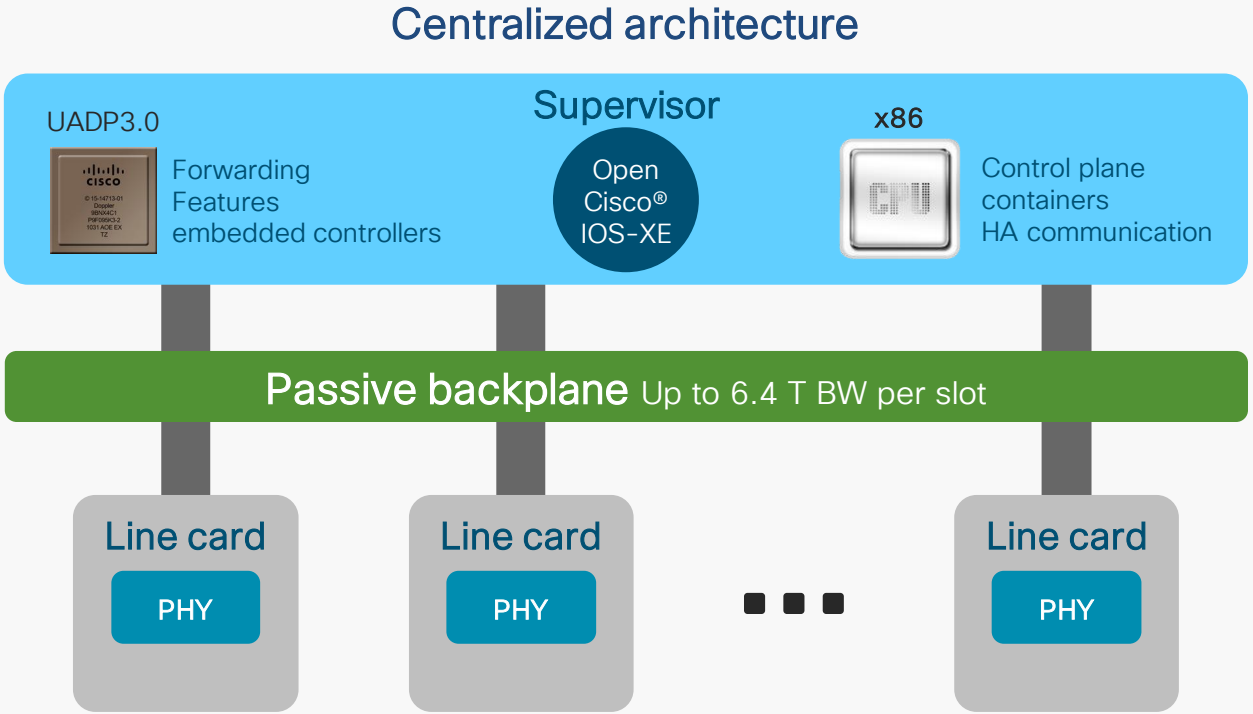

- Centralized architecture => **Uninterrupted supervisor switchover**
- Centralized architecture (Forwarding, queuing, and security are done on the supervisor) => **Unlock new capability** with a supervisor upgrade
- Transparent line cards => **Compatible** with new sup
- Passive backplane => **High MTBF**
- X86 CPU + storage => **App hosting**

Open IOS-XE

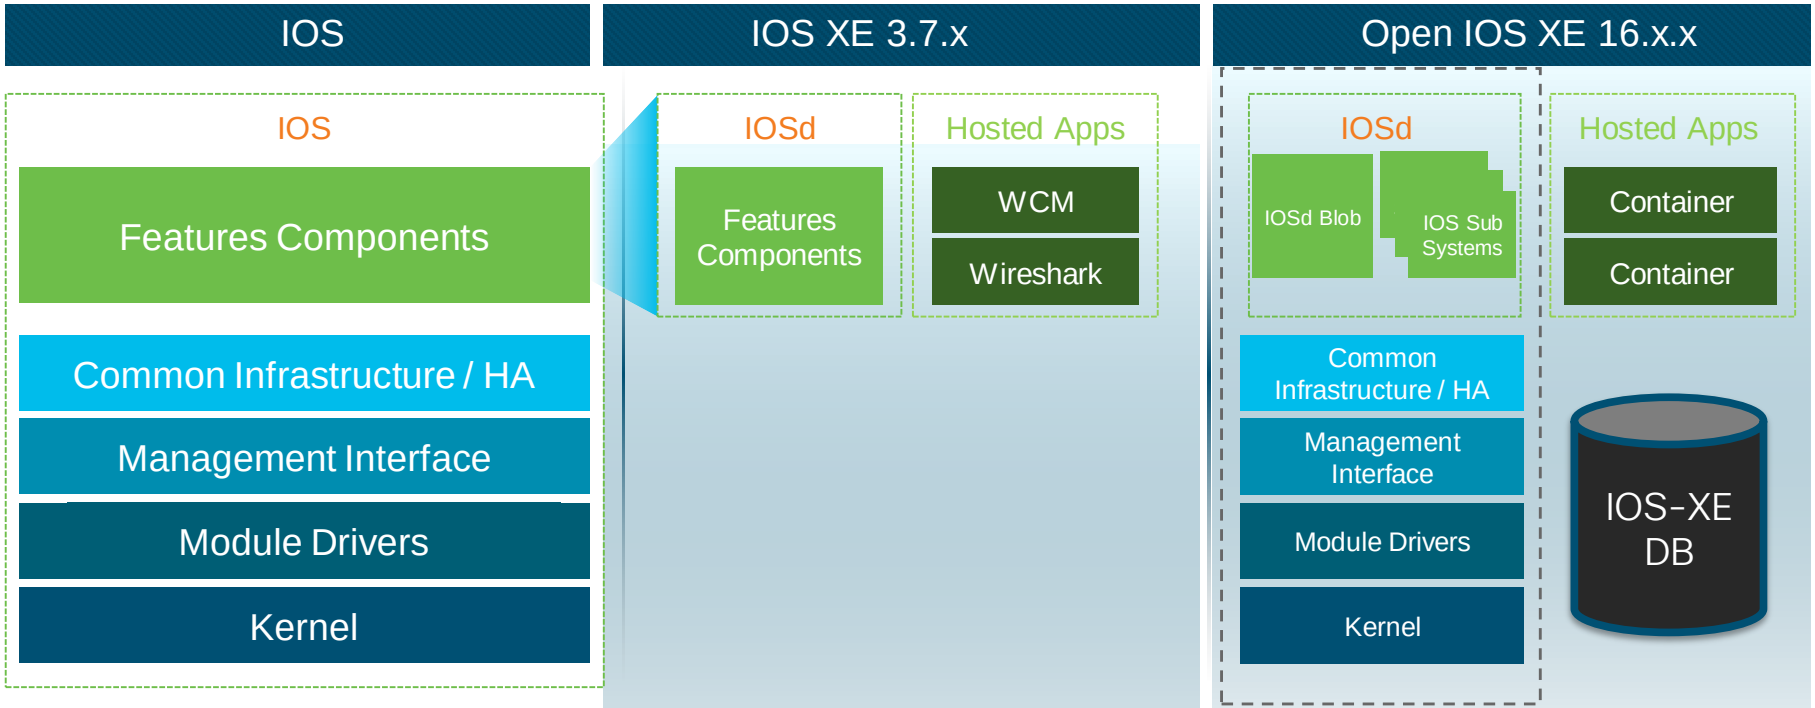

Same Look & Feel, Enhanced & Modern Architecture

Open IOS XE – IOS Sub Systems

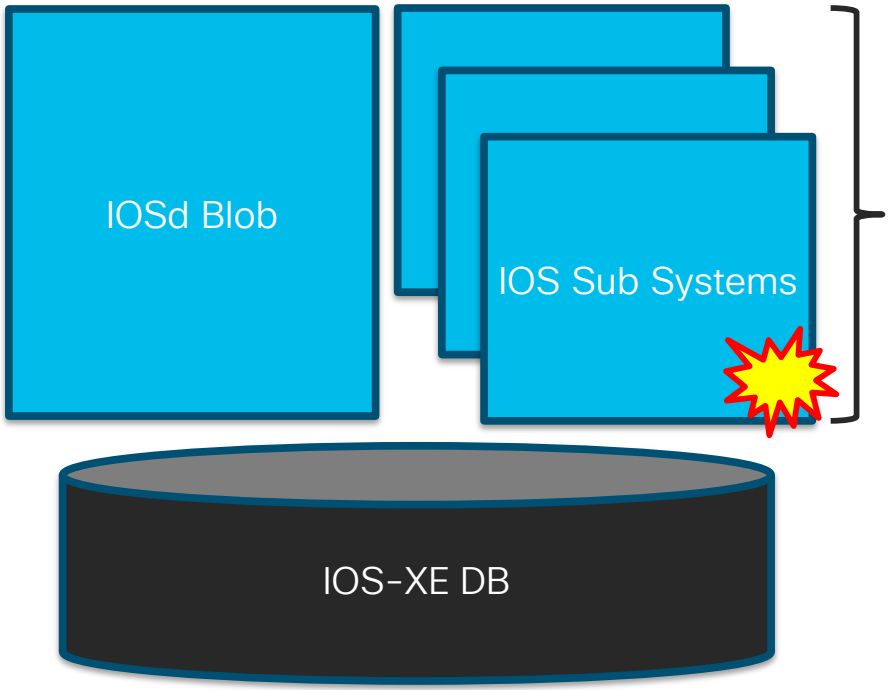

BGP

OSPF

MPLS, etc

Failure of a Sub Systems is isolated
=>
Rest of the System intact

IOS Sub Systems Enhances IOS Resiliency

Open IOS XE – DB

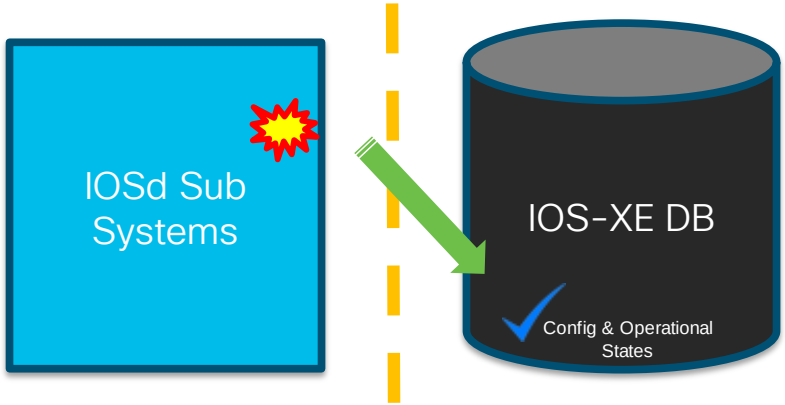

Decoupling Code & Data protects the Operational & Configurational States

Higher Application UP Time

Quicker Recovery

Better Convergence

Model-driven telemetry

- Support for any YANG subtree
- Structured data
- XML encoding
- Periodic or on change
- Reduced CPU load

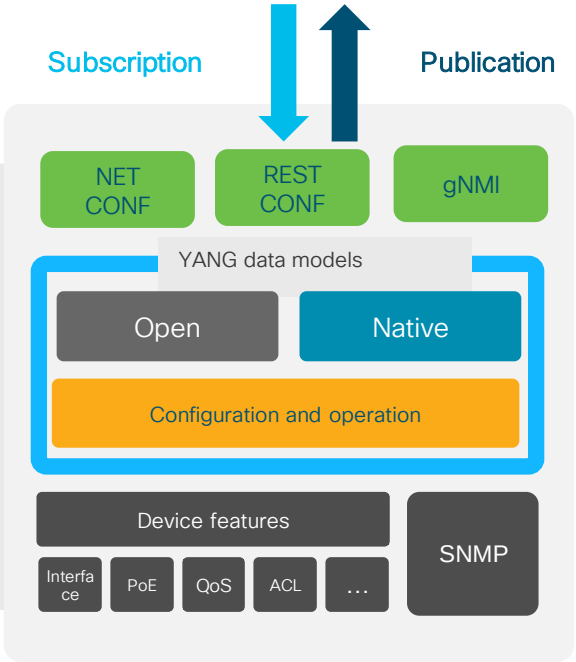

Export enriched, consistent, and concise data with context from devices for a better user and operator experience

UADP 3.0

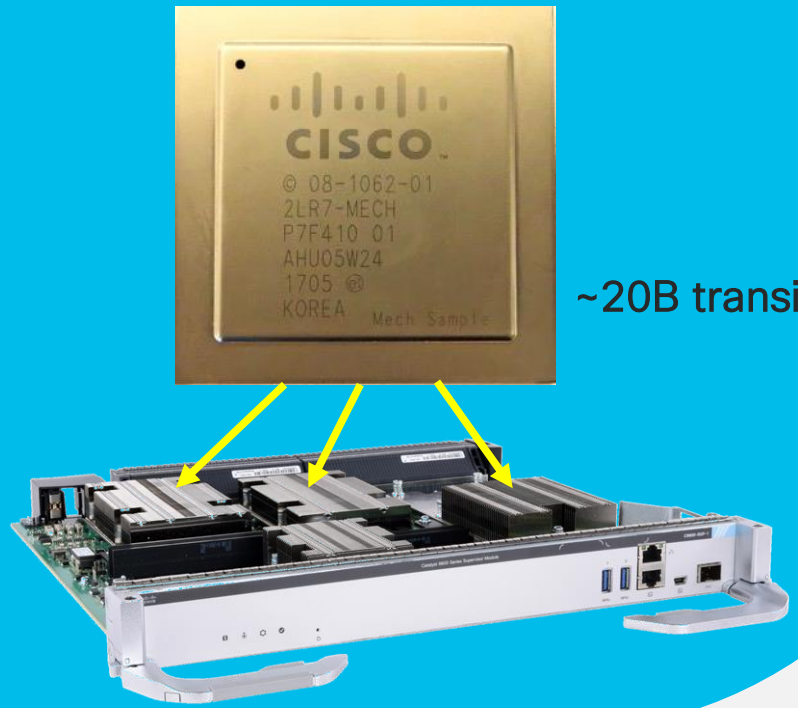

~20B transistors

Customizable ASIC templates

36-MB unified buffer

IP v4/v6

Double-width tables

3x more FIB scale

Up to 1.6 TB bandwidth

1G, 10G, 25G, 40G, and 100G speeds

Up to 1000Mpps throughput

Up to 1.6Tbps ASIC Interconnect

Common Capabilities of UADP ASIC

Flex Parser & Programmable Pipelines

Recirculation Capability

Micro Engines

Adaptable Tables

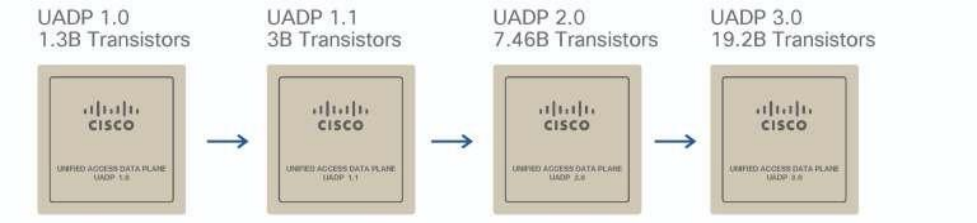

UADP 3.0 – Under the covers
showing both cores

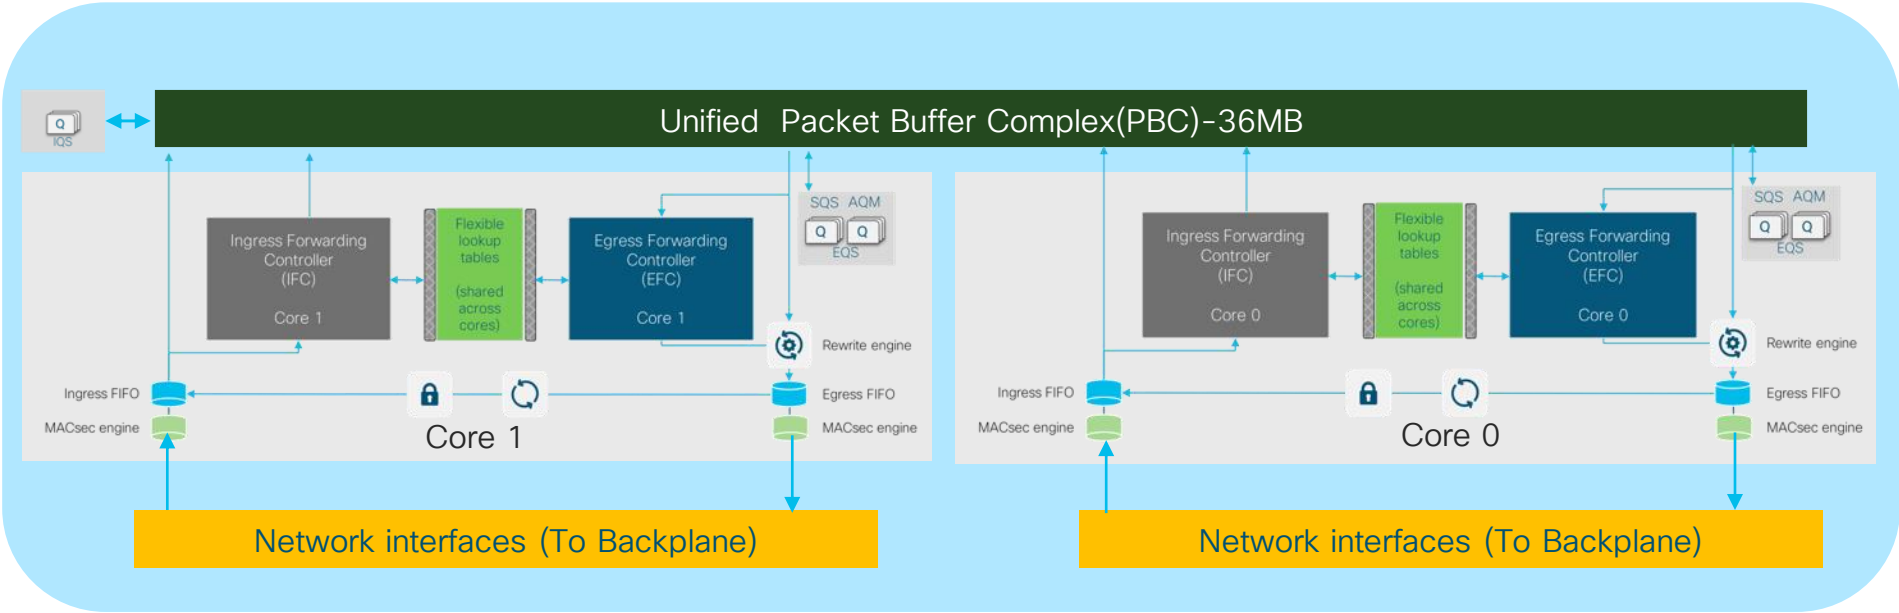

Cisco Catalyst 9600 Series – Sup1

UADP 3.0 – Under the covers showing one of the two cores

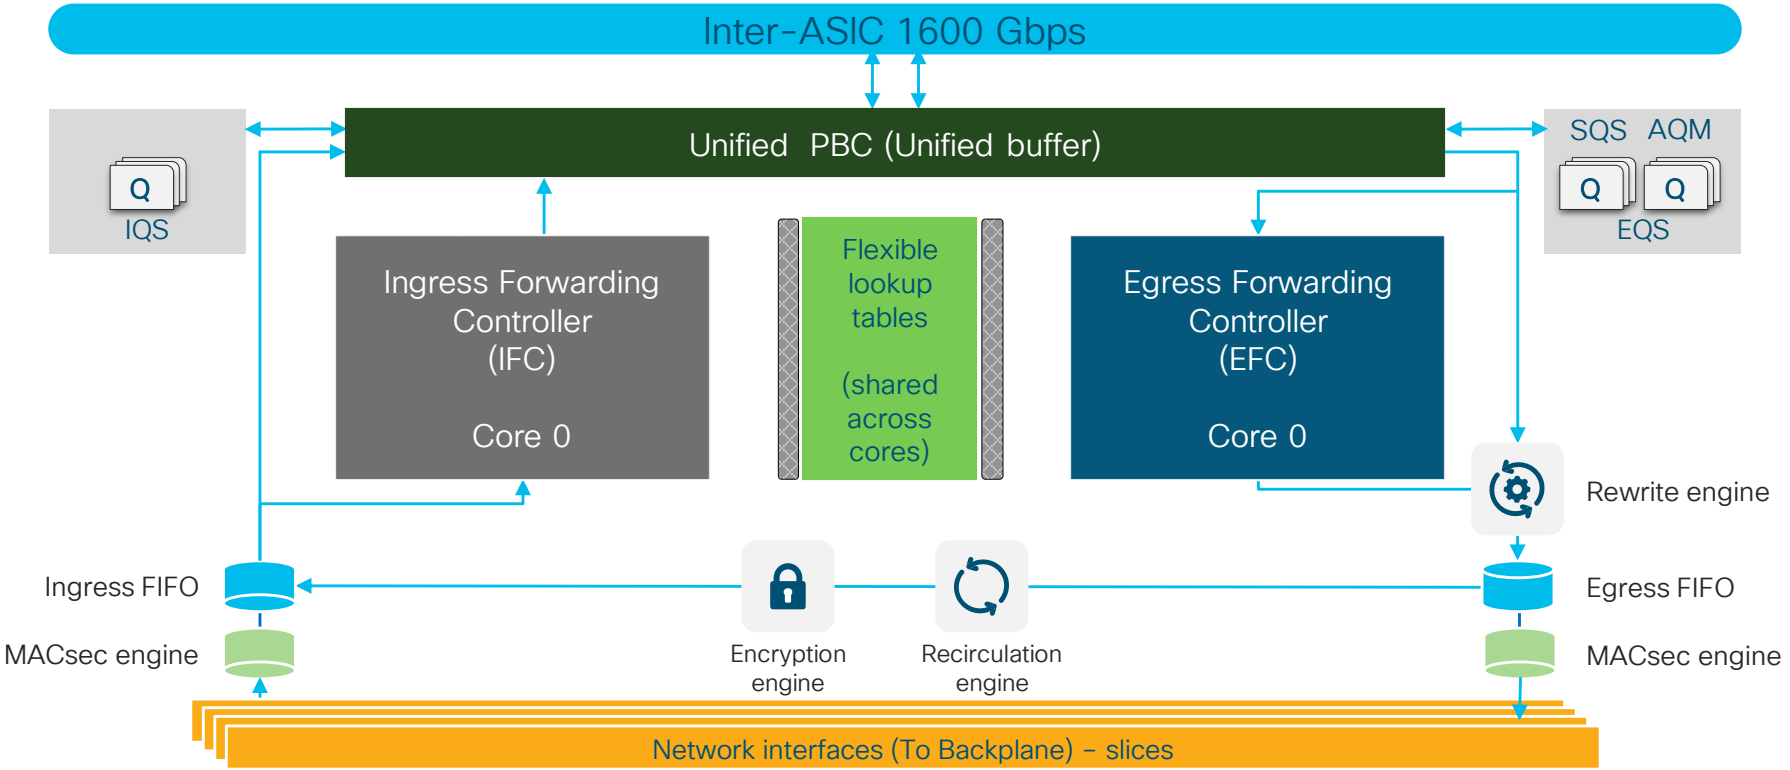

Supervisor engine 1 – Block diagram

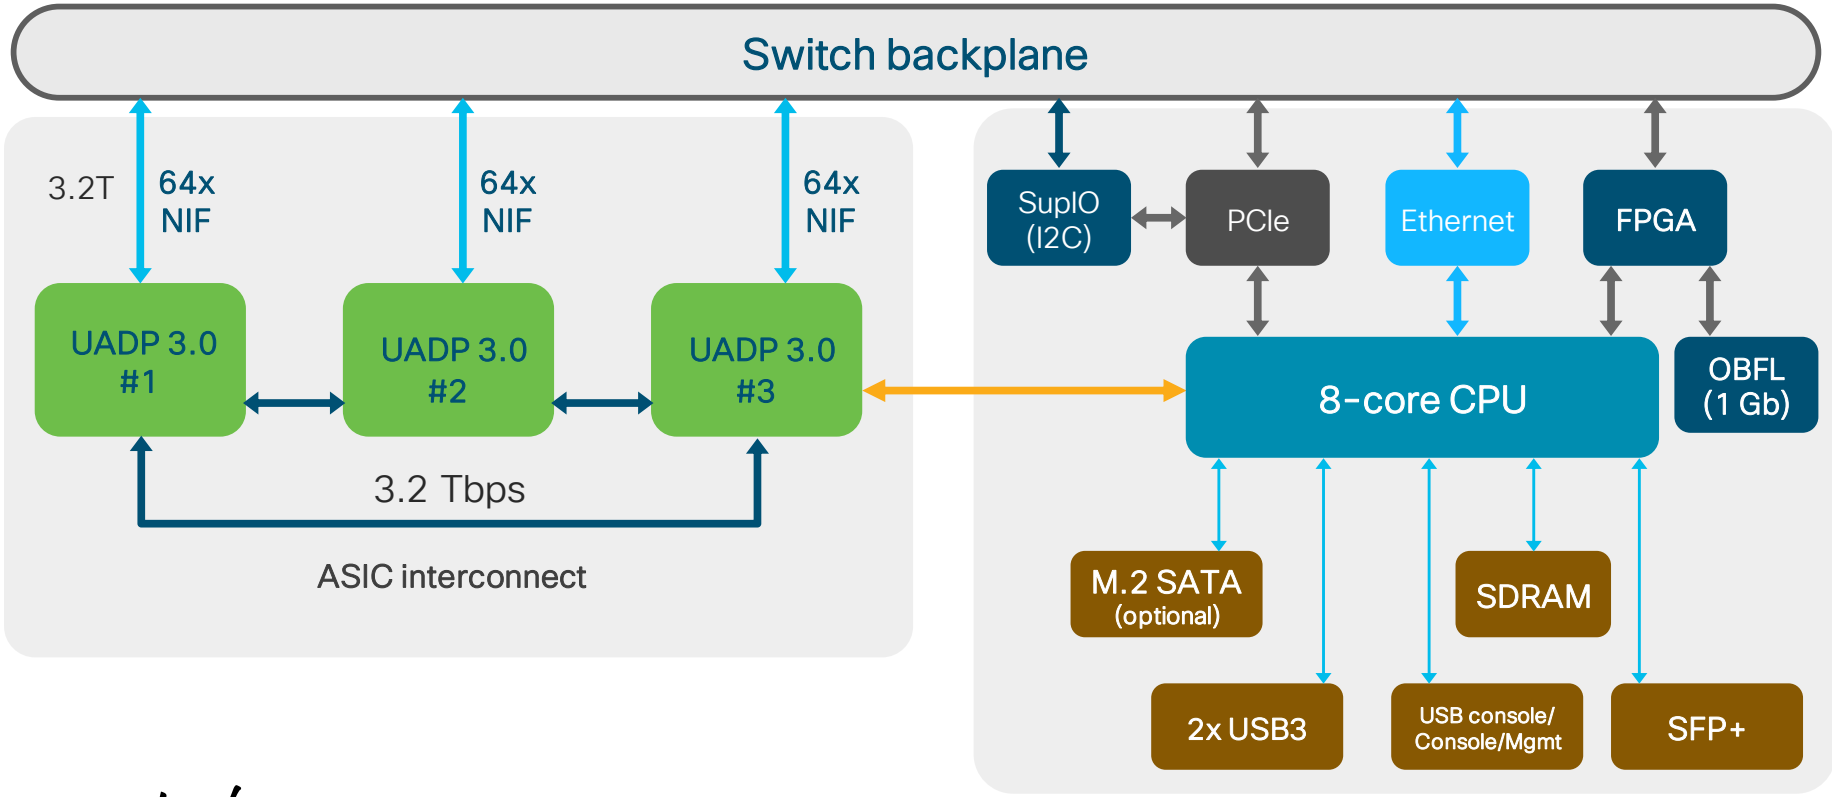

Supervisor engine 1 – ASICs to LC mapping

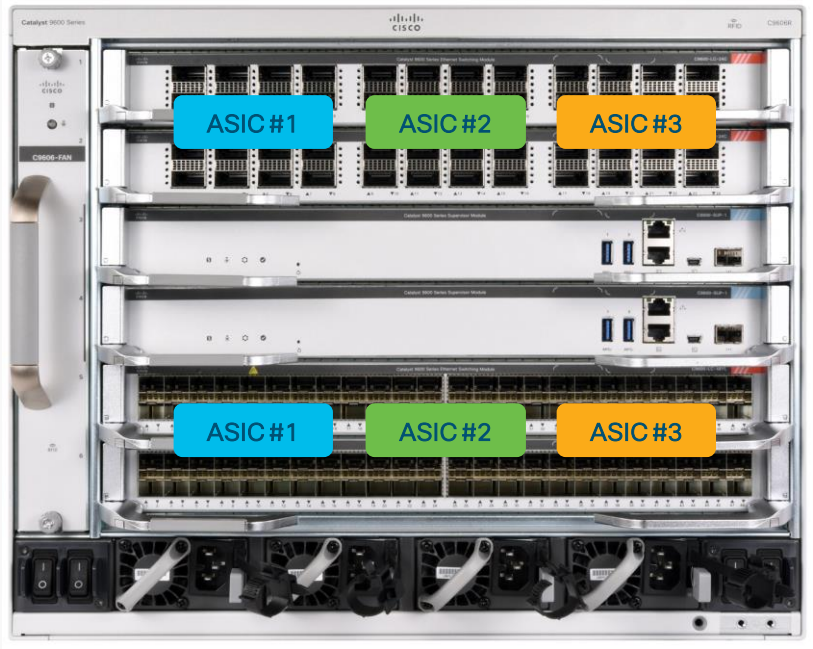

- ASIC #1: First third of the ports
  - 48-port module: 1-16
  - 24-port module: 1-8
- ASIC #2: Middle third of the ports
  - 48-port module: 17-32
  - 24-port module: 9-16
- ASIC #3: Last third of the ports
  - 48-port module: 33-48
  - 24-port module: 17-24

Cisco Catalyst 9600 – Supervisor 1
Port-to-ASIC mapping

Command to verify the port-to-ASIC mapping:
show platform software fed active ifm mappings

```
C9600-Bottom#show platform software fed active ifm mappings
Interface                IF_ID    Inst Asic Core Port SubPort Mac  Cntx LPN  GPN  Type Active
FortyGigabitEthernet1/0/1 0x7       0    0    0    0      0      0    0    1    101  NIF  N
FortyGigabitEthernet1/0/2 0x8       0    0    0    8      0      2    1    2    102  NIF  N
FortyGigabitEthernet1/0/3 0x9       0    0    0    16     0      16   0    3    103  NIF  N
FortyGigabitEthernet1/0/4 0xa       0    0    0    24     0      18   1    4    104  NIF  N
FortyGigabitEthernet1/0/5 0xb       1    0    1    8      0      14   1    5    105  NIF  Y
FortyGigabitEthernet1/0/6 0xc       1    0    1    0      0      12   0    6    106  NIF  Y
FortyGigabitEthernet1/0/7 0xd       1    0    1    24     0      30   1    7    107  NIF  Y
FortyGigabitEthernet1/0/8 0xe       1    0    1    16     0      28   0    8    108  NIF  Y
FortyGigabitEthernet1/0/9 0xf       2    1    0    0      0      0    0    9    109  NIF  Y
<SNIP>
FortyGigabitEthernet1/0/16 0x16      3    1    1    16     0      28   0    16   116  NIF  Y
FortyGigabitEthernet1/0/17 0x17      4    2    0    0      0      0    0    17   117  NIF  Y
<SNIP>
FortyGigabitEthernet1/0/24 0x1e      5    2    1    16     0      28   0    24   124  NIF  N
HundredGigE1/0/25        0x1f      0    0    0    0      0      0    0    25   125  NIF  Y
<SNIP>

C9600-Bottom#$
```

100G/40G line card block diagram

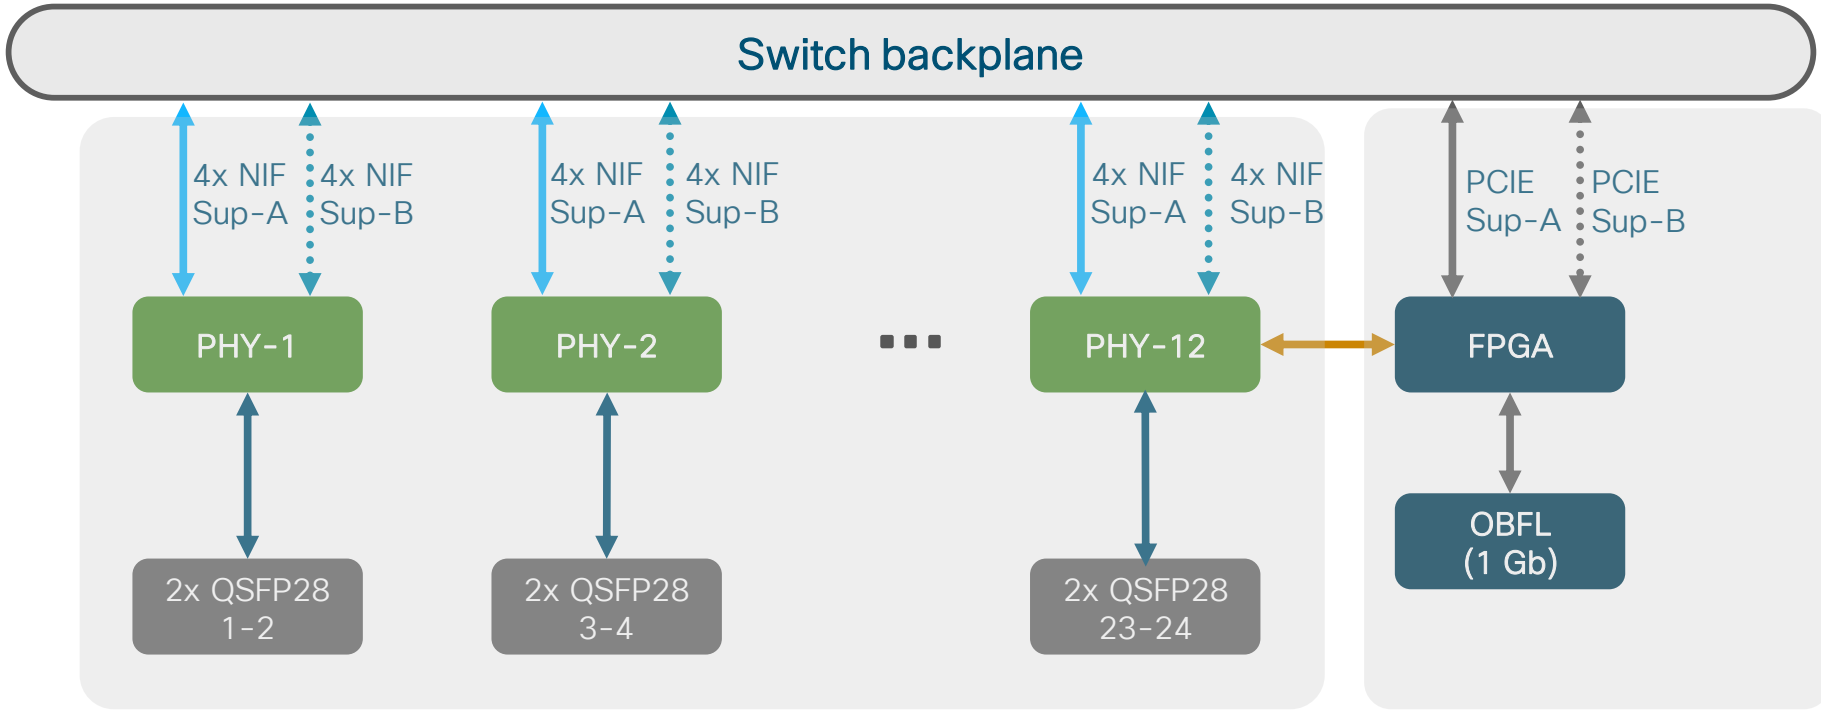

25G/10G/1G line card block diagram

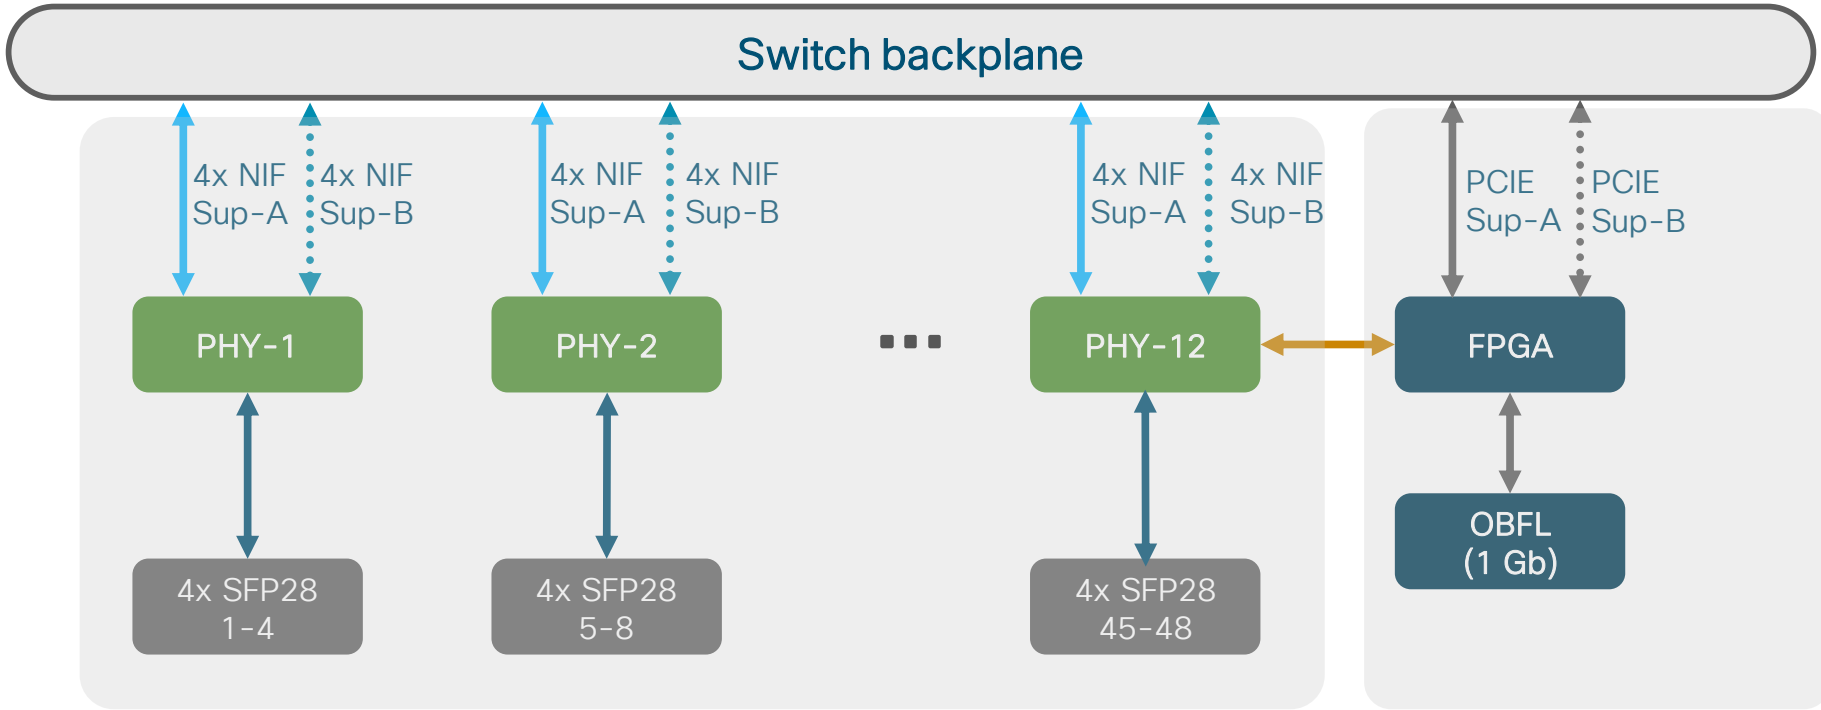

SDM Templates

You make networking **possible**

Cisco live!

Cisco Catalyst 9600 Series – Supervisor engine 1
Switch Database Management (SDM) template

Core template
Maximizes system resources for Layer 3 unicast and multicast **routes (default)**

SD-Access template
Maximizes system resources for **policy** to support **fabric** deployment

User-customizable template
Allows customizable ACL TCAM resources

Cisco® Catalyst® 9600 Series

Distribution template
Balances system resources between Layer 3 **routes** and Layer 2 **MAC** and **Netflow**

NAT template
Maximizes the **NAT** configurations on the switch

Cisco Catalyst 9600 Series
SDM templates and scale numbers

Feature		Distribution template	Core template (default)	SDA template	NAT template
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
Flexible NetFlow		**98K/ASIC**	64K/ASIC	64K/ASIC	64K/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

Customizable ACL TCAM resources

Cisco Catalyst 9600 Series
SDM template – Customizable TCAM section

C9600-Bottom#sho sdm prefer
Showing SDM Template Info

This is the Core template.

Security Ingress IPv4 Access Control Entries*	: 6656 (current) - 6656 (proposed)
Security Ingress Non-IPv4 Access Control Entries*	: 5632 (current) - 5632 (proposed)
Security Egress IPv4 Access Control Entries*	: 6656 (current) - 6656 (proposed)
Security Egress Non-IPv4 Access Control Entries*	: 8704 (current) - 8704 (proposed)
QoS Ingress IPv4 Access Control Entries*	: 4608 (current) - 4608 (proposed)
QoS Ingress Non-IPv4 Access Control Entries*	: 3584 (current) - 3584 (proposed)
QoS Egress IPv4 Access Control Entries*	: 4608 (current) - 4608 (proposed)
QoS Egress Non-IPv4 Access Control Entries*	: 3584 (current) - 3584 (proposed)
Netflow Input Access Control Entries*	: 1024 (current) –1024 (proposed)
Netflow Output Access Control Entries*	: 1024 (current) – 1024 (proposed)
Flow SPAN Input Access Control Entries*	: 512 (current) – 512 (proposed)
Flow SPAN Output Access Control Entries*	: 512 (current) – 512 (proposed)

Cisco Catalyst 9600 Series
SDM customizable template – CLI

Customizable range: 10% - 90%

- Between input and output
- Between IPv4 and non-IPv4

Security-ACL allocation	Default	
27K	12K (input)	7K (v4)
		5K (non-v4)
	15K (output)	7K (v4)
		8K (non-v4)

Example 1

Security-ACL allocation	Input =10% Input V4 – 75% Output v4 – 75%	
27K	3K (input)	2K (v4)
		1K (non-v4)
	24K (output)	18K (v4)
		6K (non-v4)

Example 2

Security-ACL allocation	Input = 50% Input V4 – 75% Output v4 – 75%	
27K	13K (input)	9.5K (v4)
		3.5K (non-v4)
	14K (output)	10.5K (v4)
		3.5K (non-v4)

Example 3

Security-ACL allocation	Input = 90% Input V4 – 75% Output v4 – 75%	
27K	24K (input)	18K (v4)
		6K (v4)
	3K (output)	2K (v4)
		1K (non-v4)

Cisco Catalyst 9600 Series – Supervisor Engine 1
SDM customizable template – CLI

Command to modify ACL TCAM allocation

```
C9600(config)#sdm prefer template-modification?
  default        Default prefered template
  fspan          Filter Span
  nfl            NFL ACLs
  qos            QOS
  security-acl   Security ACLs
```

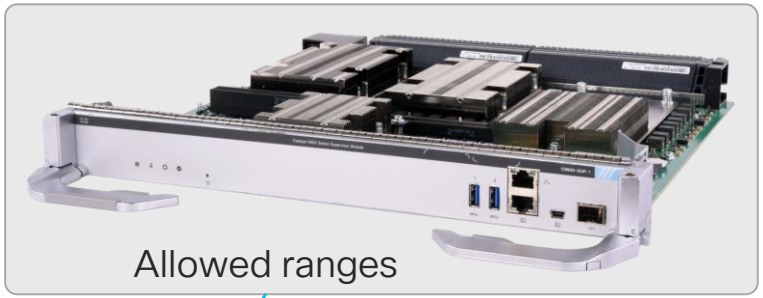


```
C9600 (config)# sdm prefer template-modification security-acl input allowed-range
Total_size: 27648 Suggested split percentage for input: 11 18 22 25 33 37 40 48 49 52 60 63 67 75 78 82 89

C9600(config)#sdm prefer template-modification security-acl input 15 input-ipv4 15 output-ipv4 85
Allocated Security Acl Input (IPv4:1024, Non-IPv4:4096) entries, Output (IPv4:18432, Non-IPv4:4096) entries
input=18.52 input_ipv4=20.00, output_ipv4=81.82
Modifications to the preferred template have been stored, but cannot take effect until the next reload. Allocations will be an
approximation of user-specified percentages. Use 'show sdm prefer' to see proposed values.

C9600(config)#
```

Forwarding

You make networking **possible**

Cisco live!

Cisco Catalyst 9600 Series
Unicast forwarding within ASIC (ingress and egress)

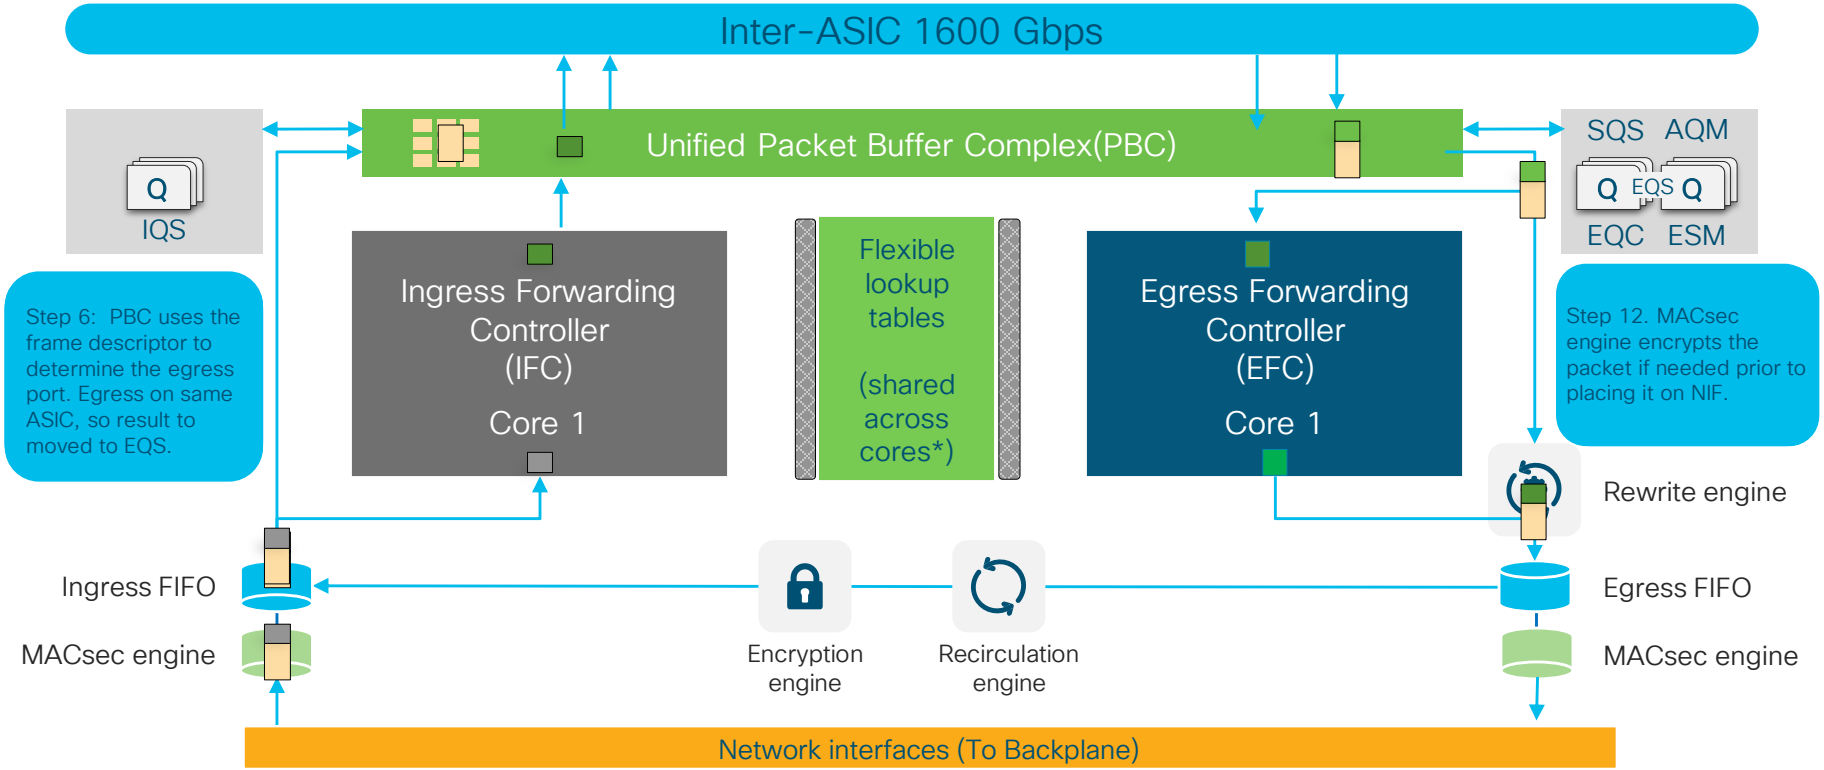

Cisco Catalyst 9600 Series
Unicast forwarding within ASIC (ingress and egress)

Reference

Step 1: Packet arrives at ingress port, PHY converts the signal and serializes the bits, and then it sends to network interface ports.

Step 2: Network interface passes packet to ingress MACsec engine.

Step 3: MACsec engine decrypts the packet if needed and passes unencrypted packet to ingress FIFO.

Step 4: IFC snoops packet between FIFO and PBC.

Step 5: IFC returns lookup result (frame descriptor) to PBC.

Step 6: PBC uses the frame descriptor to determine the egress port. Egress on same ASIC, so result to moved to EQS.

Step 7. EQS – replication, scheduling, and queue management.

Step 8: PBC sends packet with new frame descriptor and enqueues the frame.

Step 9. EFC snoops packet between PBC and rewrite engine.

Step 10. EFC performs egress lookup functions to learn SRC MAC, egress SPAN, etc. and sends results to rewrite engine.

Step 11. Rewrite engine rewrites packets and sends through the egress FIFO.

Step 12. MACsec engine encrypts packet prior to placing it on NIF.

Cisco Catalyst 9600 Series
Unicast forwarding across ASIC (ingress and egress)

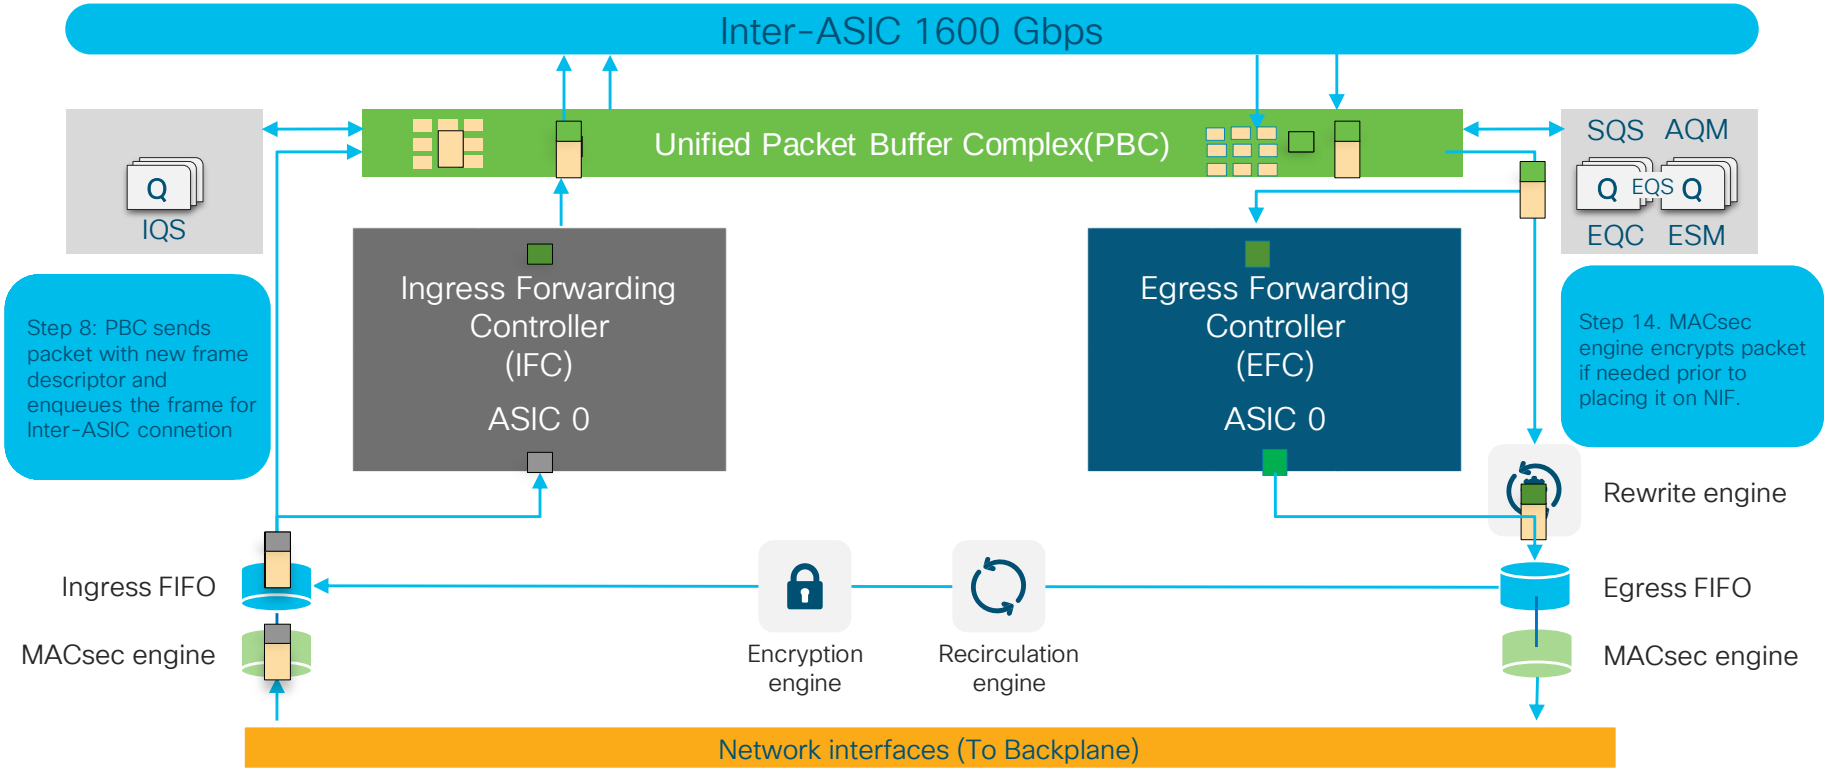

Cisco Catalyst 9600 Series
Unicast forwarding across ASIC (ingress and egress)

Reference

Step 1: Packet arrives at ingress port, PHY converts the signal and serializes the bits, and then it sends to network interface ports

Step 2: Network interface passes packet to ingress MACsec engine.

Step 3: MACsec engine decrypts the packet if needed and passes unencrypted packet to ingress FIFO.

Step 4: IFC snoops packet between FIFO and PBC.

Step 5: IFC returns lookup result (frame descriptor) to PBC.

Step 6: PBC uses the frame descriptor to determine the egress port. Egress port across ASIC enqueues result to IQS.

Step 7: IQS provides queuing and scheduling functions for packet to be enqueued to Inter-ASIC connection.

Step 8: PBC sends packet with new frame descriptor and enqueues the frame to Inter-ASIC connection.

Step 9. Packet arrives from inter-ASIC connection, PBC parses header and sends to EQS.

Step 10. EQS performs replication, scheduling, and queue management and sends header to PBC.

Step 11. EFC snoops packet between PBC and rewrite engine.

Step 12. EFC performs egress lookup functions to learn SRC MAC, egress SPAN, etc. and sends results to rewrite engine.

Step 13. Rewrite engine rewrites packets and sends through the egress FIFO.

Step 14. MACsec engine encrypts packet prior to placing it on NIF.

Cisco Catalyst 9600 Series
Multicast forwarding

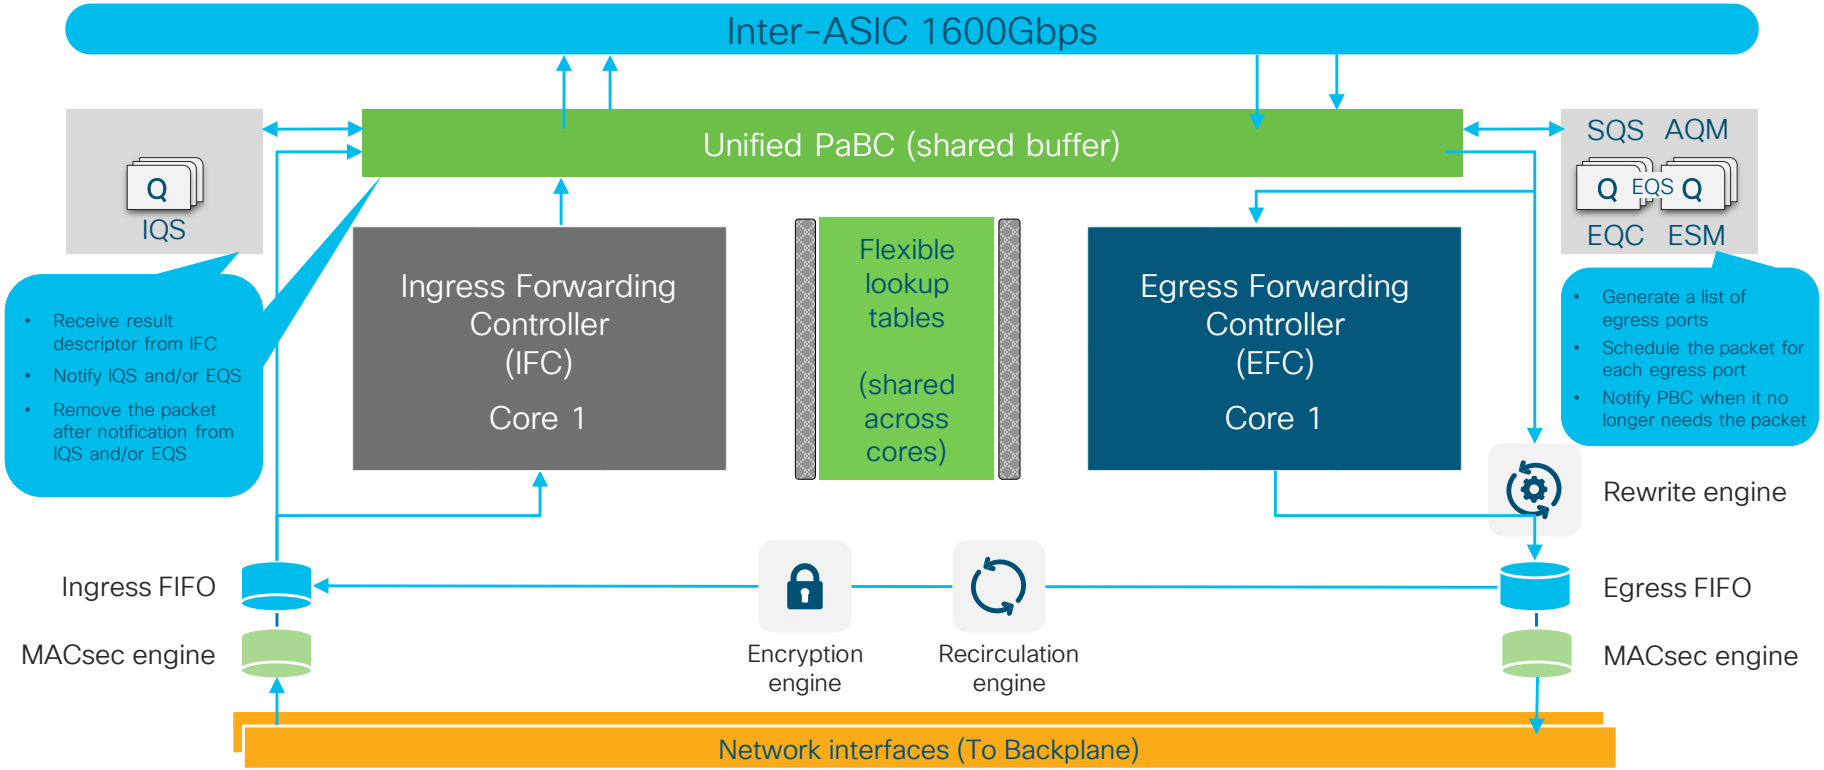

Cisco Catalyst 9600 Series
Multicast forwarding within ASIC (ingress and egress)

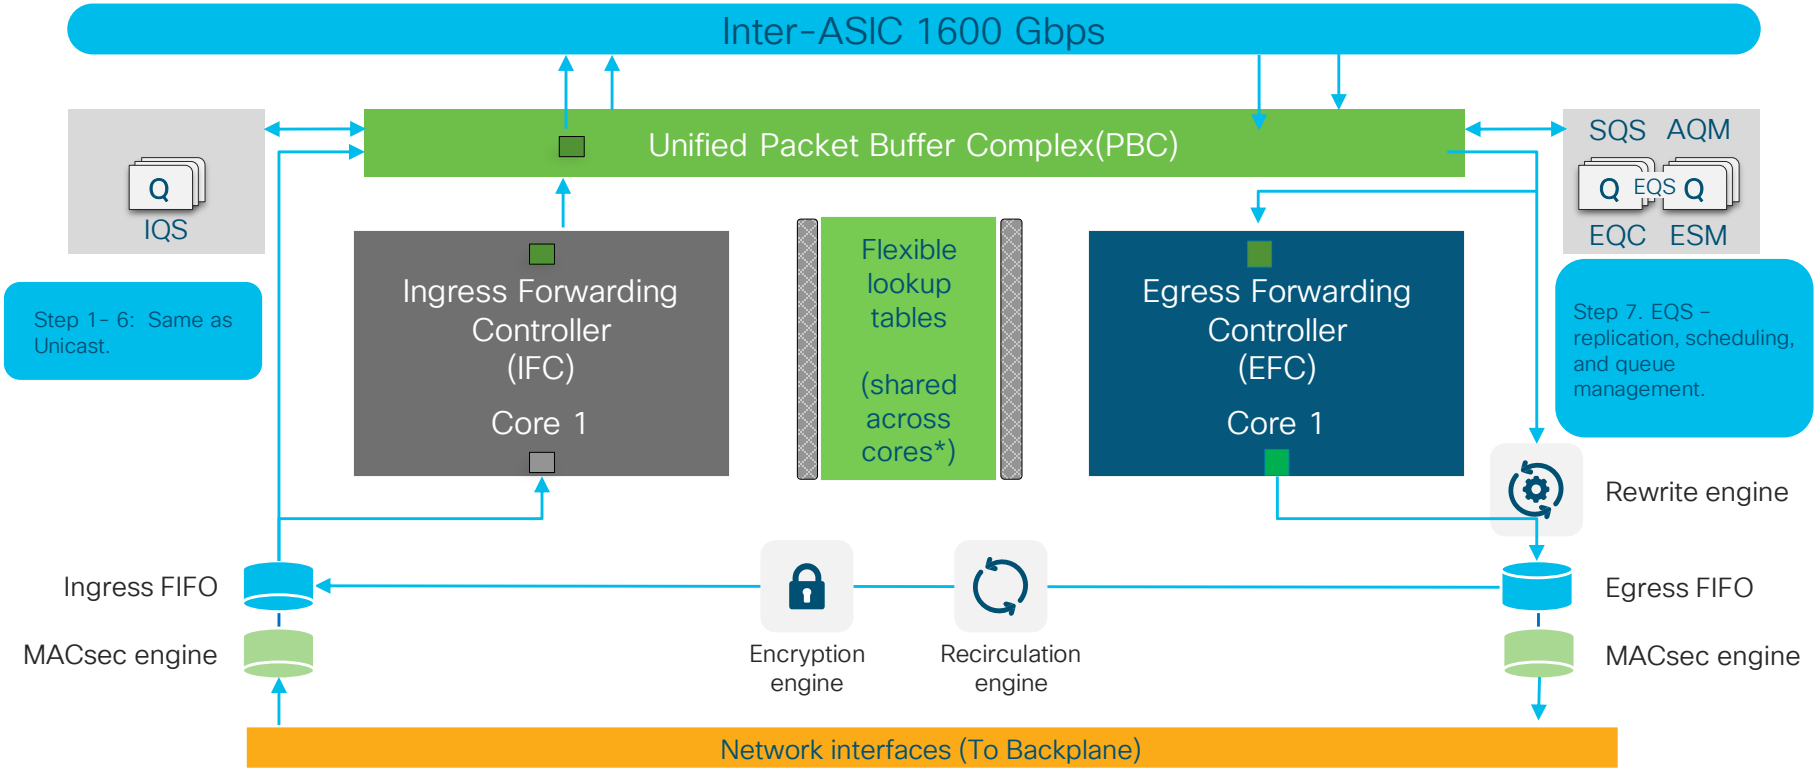

Access Control Lists

You make networking **possible**

Cisco live!

Cisco Catalyst 9600 Series access control lists
Four forms of security ACLs

The Cisco Catalyst 9600 Series supports four forms of security ACL: RACL, VACL, PACL, Group ACL

Router ACL (RACL)	VLAN ACL (VACL)	Port ACL (PACL)	Group ACL (GACL)
Used to permit or deny the movement of traffic between Layer 3 subnets	Used to permit or deny the movement of traffic between Layer 3 subnets and VLANs or within a VLAN	Used to permit or deny the movement of traffic between Layer 3 subnets and VLANs or within a VLAN	Used to permit or deny the movement of traffic based on the groups that are assigned

• Direction: **In, Out**
• Attach Point:
  • **Layer 3 interface**
  • **SVI,**
  • **Layer 3 EtherChannel interface**
	• Direction: **Inherently both In and Out**
	• Attach Point: **VLAN**
		• Direction: **In, Out**
		• Attach Point:
		  • **Layer 2 switch port interface**
		  • **Layer 2 EtherChannel interface**
			• Direction: **In**
			• Attach Point:
			  • **Layer 3 interface**
			  • **SVI**
			  • **Layer 3 EtherChannel Interface**

Standard/extended ACLs	Standard/extended ACLs	Standard/extended/ MAC ACLs	Standard/extended

Cisco Catalyst 9600 Series access control lists
Order of processing

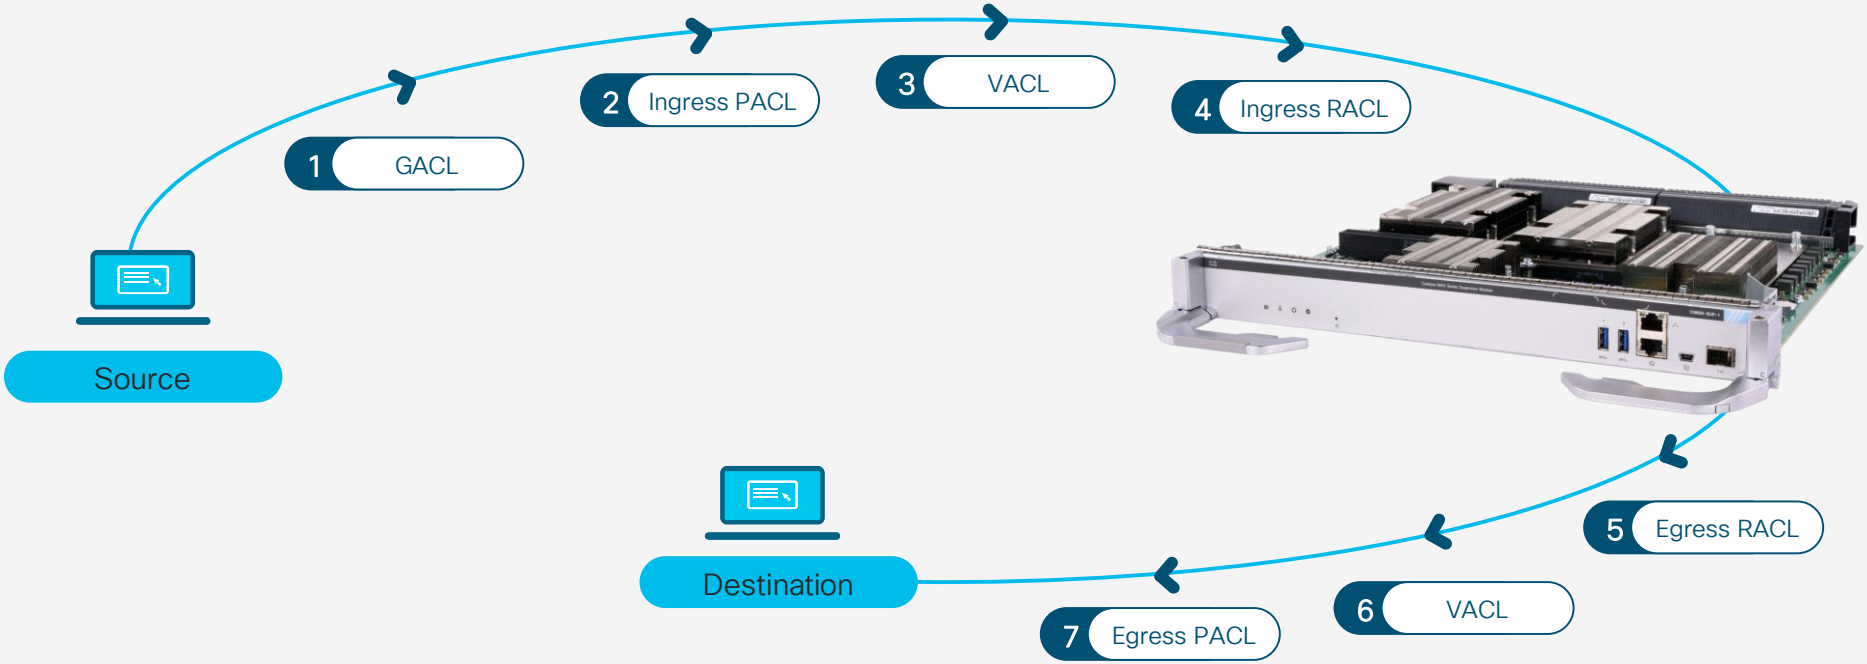

Cisco Catalyst 9600 Series – access control lists
Hardware support

Create an ACL or classification policy, using the CLI or Network Management system (NMS)

```
ip access-list extended Internet
 permit ip any host 10.2.2.4
 permit ip any host 10.5.2.33
 permit ip any host 10.11.0.0
 permit ip any host 10.4.0.0
```

1

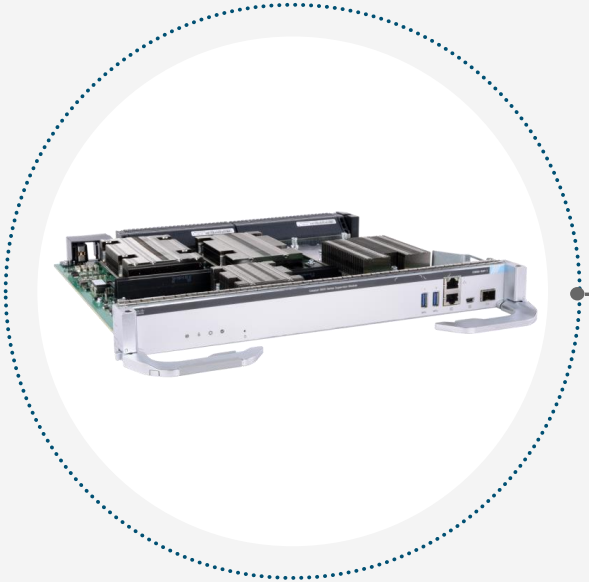

No hardware resources
Drop packet* 2

Full hardware support
- Router ACLs
- VLAN ACLs
- Port-based ACLs
- Group ACLs

3

Hardware-assisted Service ACL features
- NetFlow
- NAT and PAT
- PBR
- WCCP
- Cisco TrustSec®

4

* If TCAM exhaustion happens, packets are dropped.

Resource Utilization

Reference

```
C9600-Top#$rm hardware fed active fwd-asic resource tcam utilization
CAM Utilization for ASIC  [0]
 Table                                              Max Values          Used Values
 ------------------------------------------------------------------------------------
 Unicast MAC addresses                              32768/768           24/21
 L3 Multicast entries                               32768/768           0/5
 L2 Multicast entries                               2304                6
 Directly or indirectly connected routes            212992/1536         10/12
 Input Ipv4 QoS Access Control Entries              5632                5
 Input Non Ipv4 QoS Access Control Entries          2560                15
 Output Ipv4 QoS Access Control Entries             6144                5
 Output Non Ipv4 QoS Access Control Entries         2048                15
 Input Ipv4 Security Access Control Entries         7168                12
 Input Non Ipv4 Security Access Control Entries     5120                76
 Output Ipv4 Security Access Control Entries        7168                10
 Output Non Ipv4 Security Access Control Entries    8192                27
 Ingress Netflow ACEs                               1024                8
 Policy Based Routing ACEs                          3072                20
 Egress Netflow ACEs                                1024                9
 <SNIP>
```

Cisco Catalyst 9600 Series – hitless TCAM update

- **Allows updates to an ACL without interrupting traffic**
- Multiple features updated at once
  - IPv4, IPv6, MAC
  - PACL, RACL, VACL, and SG ACL
- Hitless update is enabled by default; can't be disabled
- Hitless update feature requires free ACL TCAM space for reprogramming but doesn't consume any additional TCAM resources
- If not enough space in TCAM, falls back to old ACL method (deny all while updating)

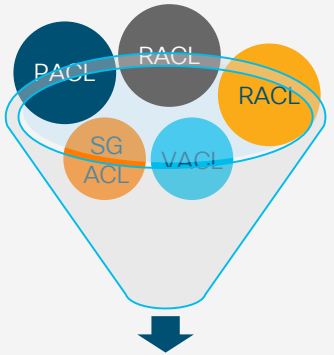

ACL updates

Cisco Catalyst 9600 Series– hitless TCAM update

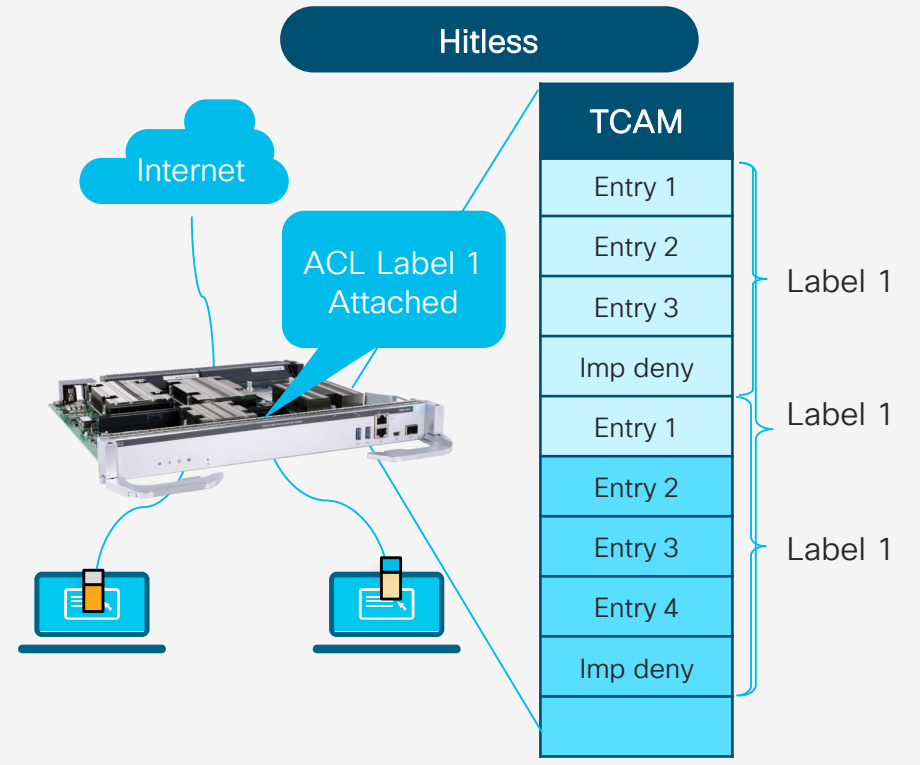

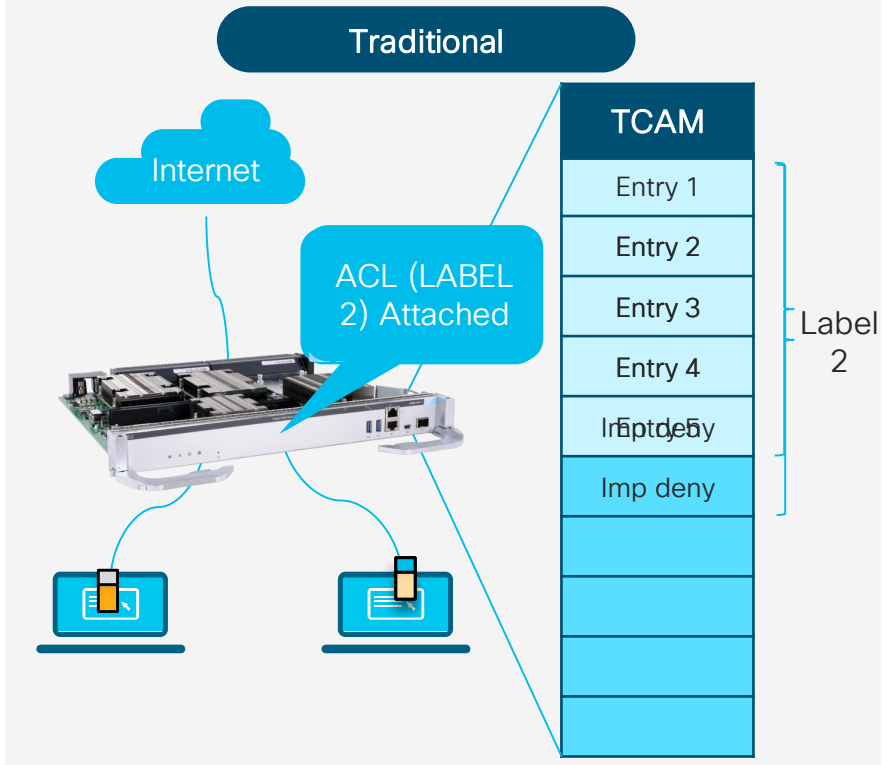

Quality of service

You make networking **possible**

Cisco live!

Cisco Catalyst 9600 Sup1 – Quality of service

- QoS is enabled by default
- All ports are trusted at Layer 2 and Layer 3 by default

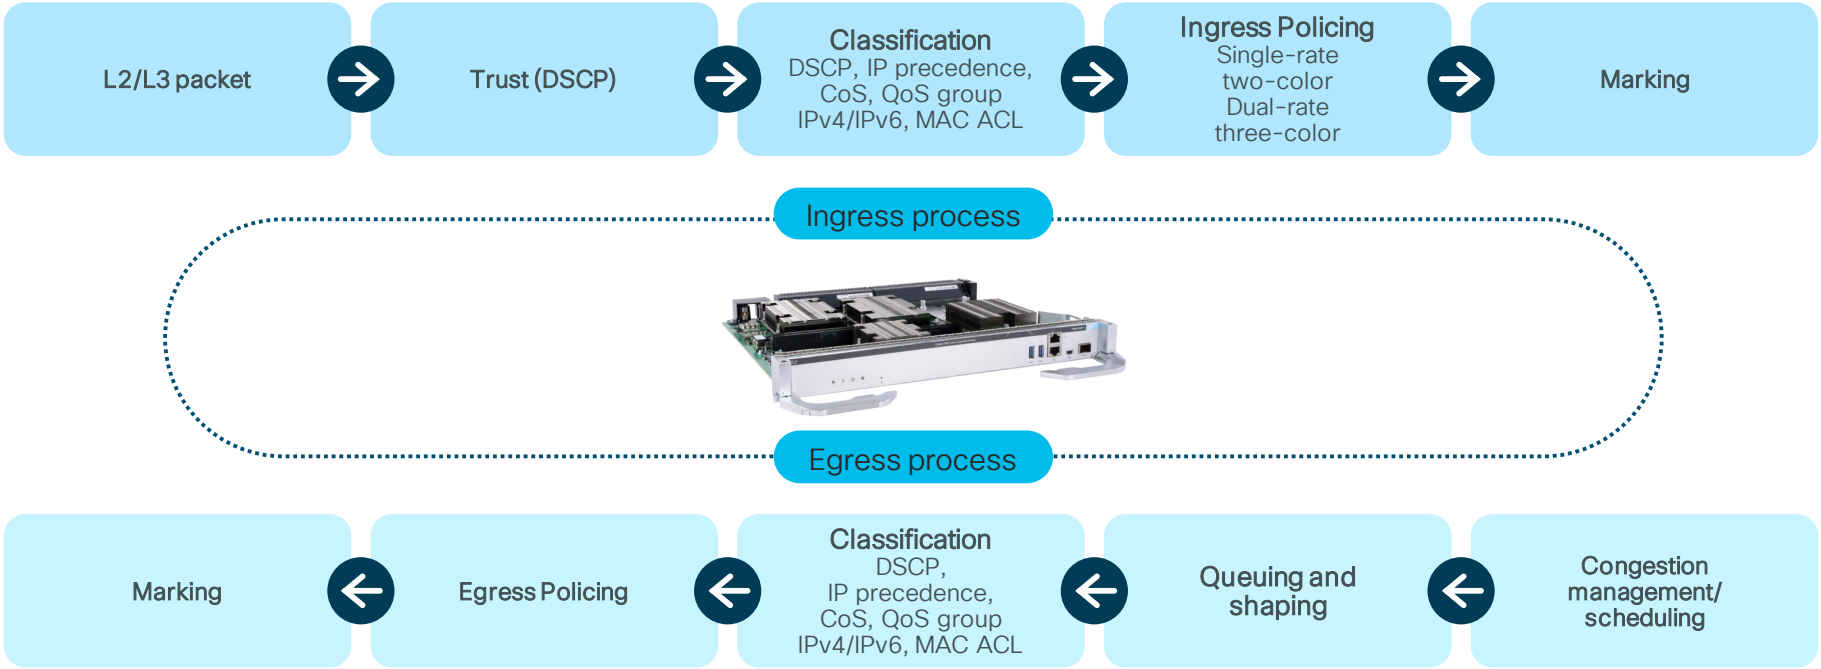

Cisco Catalyst 9600 Sup1

QoS forwarding (ingress and egress)

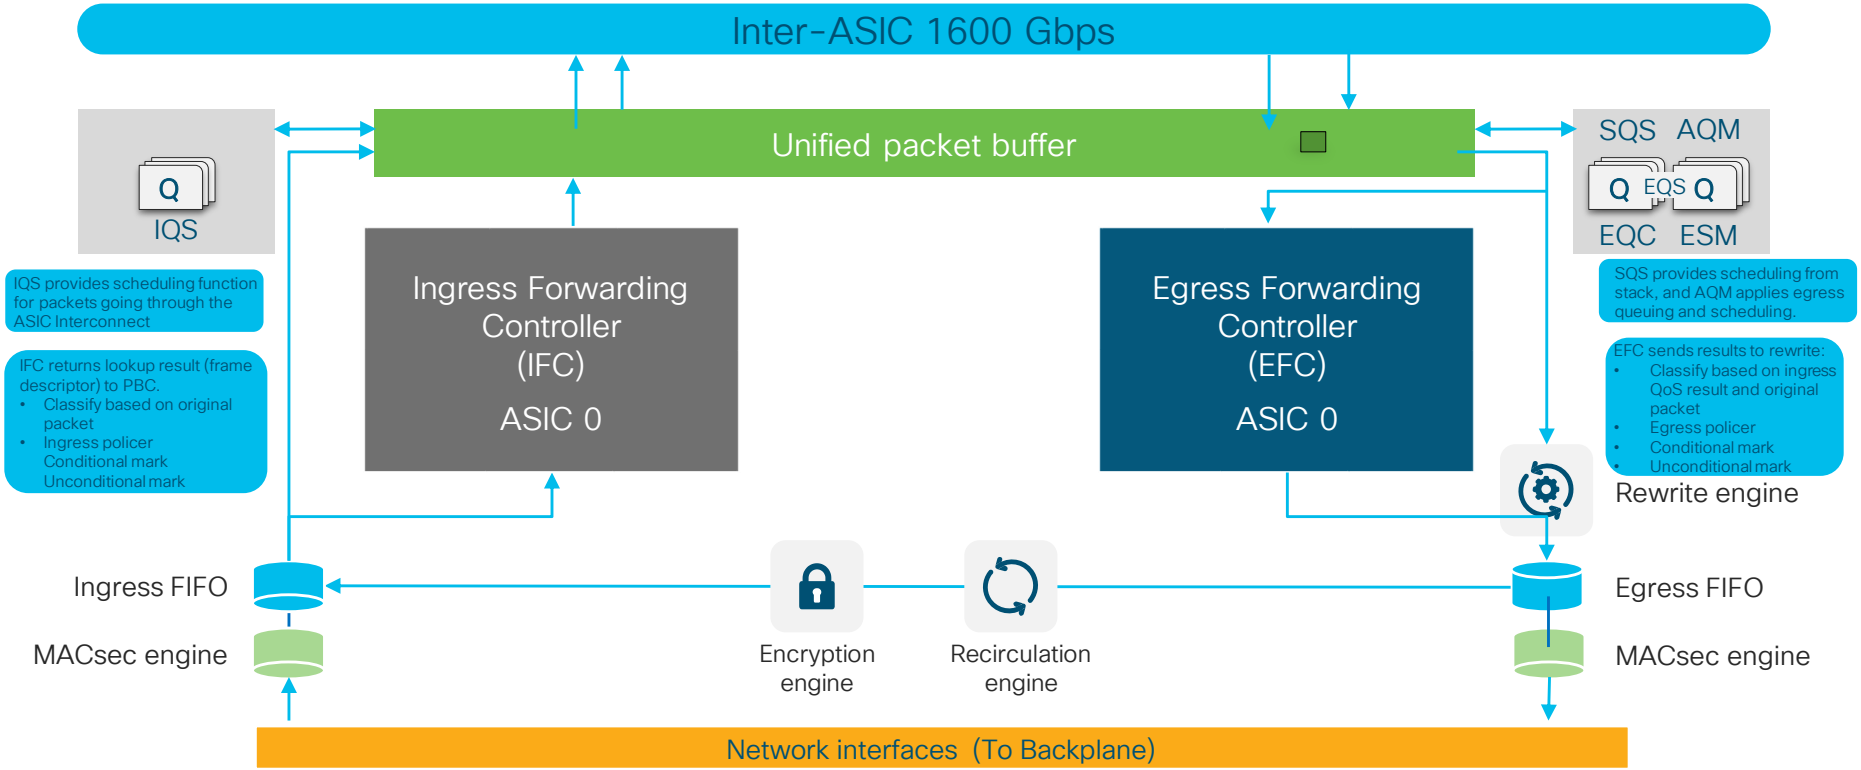

Cisco Catalyst 9600 Series Sup1

Hardware queues mapping – default

Default (2Q3T)

1 priority queue
1 nonpriority queue

Number of Queues

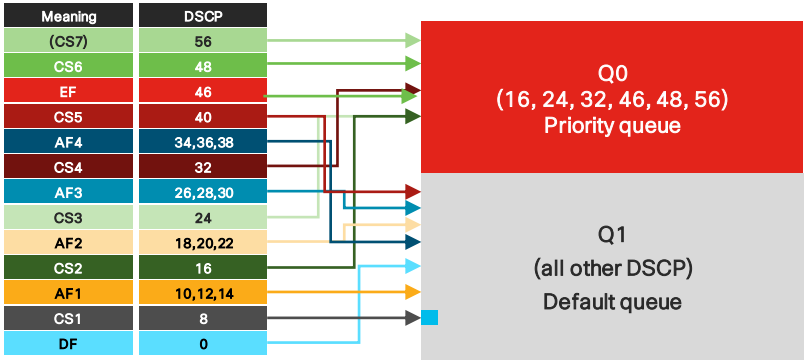

Mapping

Cisco Catalyst 9600 Series Sup1
Hardware queues mapping – Configurable with 2P6Q3T

Configured Example (8Q3T/2P6Q3T)

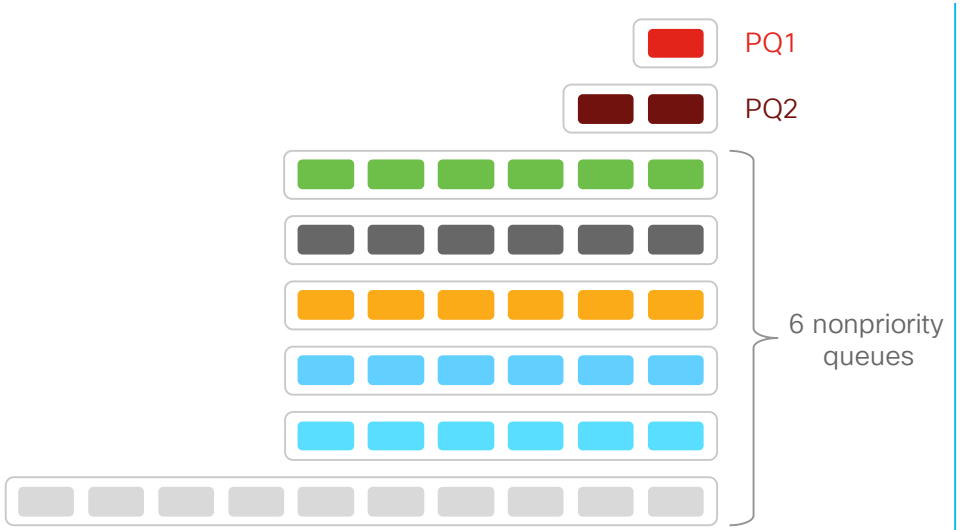

Number of Queues

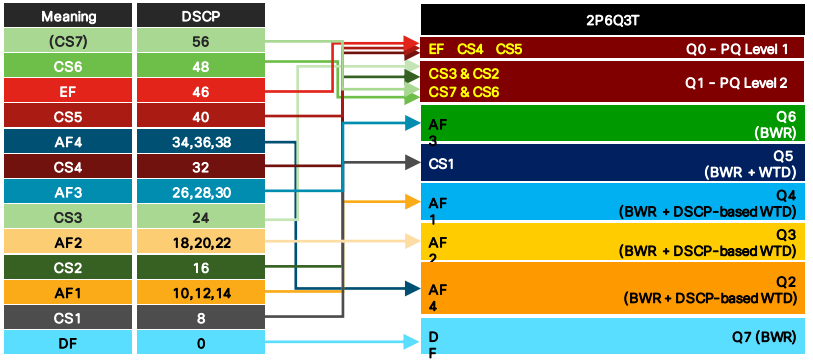

Mapping

Cisco Catalyst 9600 Sup1 – buffer complex

- Total of 108MB buffer on Sup1
- 36-MB unified packet buffer per ASIC is shared by ingress and egress data paths and between both cores
- Resources consuming packet buffer
  - Ingress buffers (IQS)
  - Egress stack buffers (SQS)
  - Egress port buffers (AQM)
  - Temporary buffers
  - Common buffers

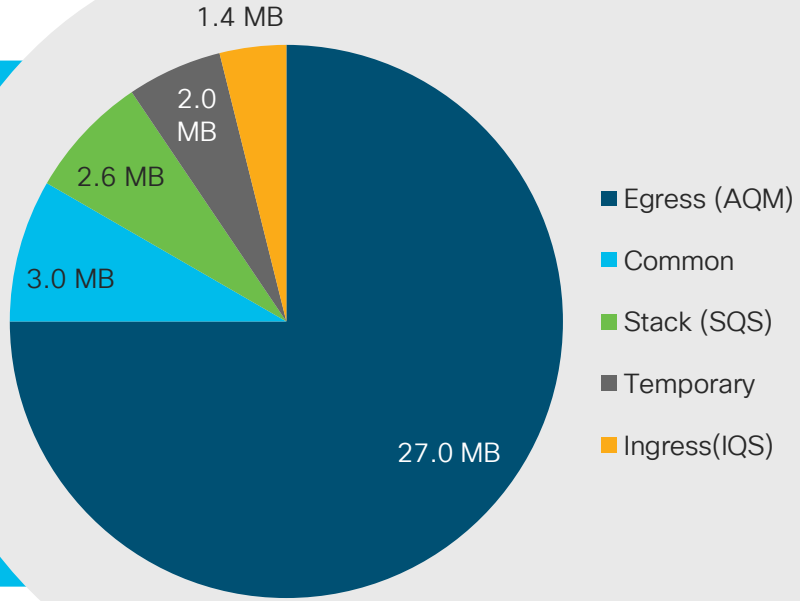

Software support for unified buffer sharing is on the roadmap

Buffer – Dedicated

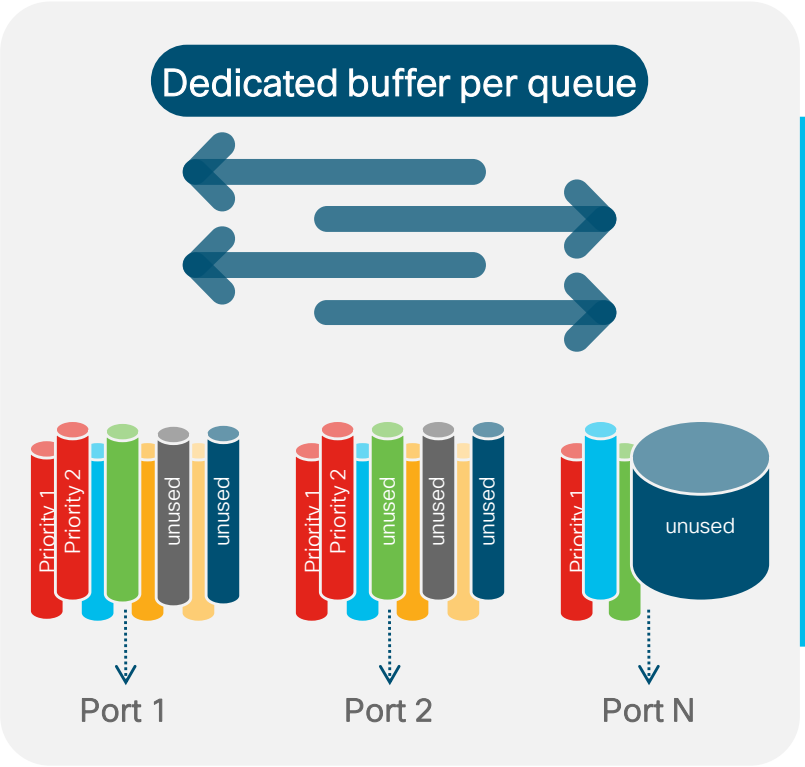

- Dedicated buffers are statically allocated for each configured queue
- Reminding buffer are allocated to the shared pool
- System use dedicated buffer first. Once a queue exhausted the dedicated buffer, then it use the shared buffer

Buffer – Shared
DTS – Dynamic Threshold and Scaling

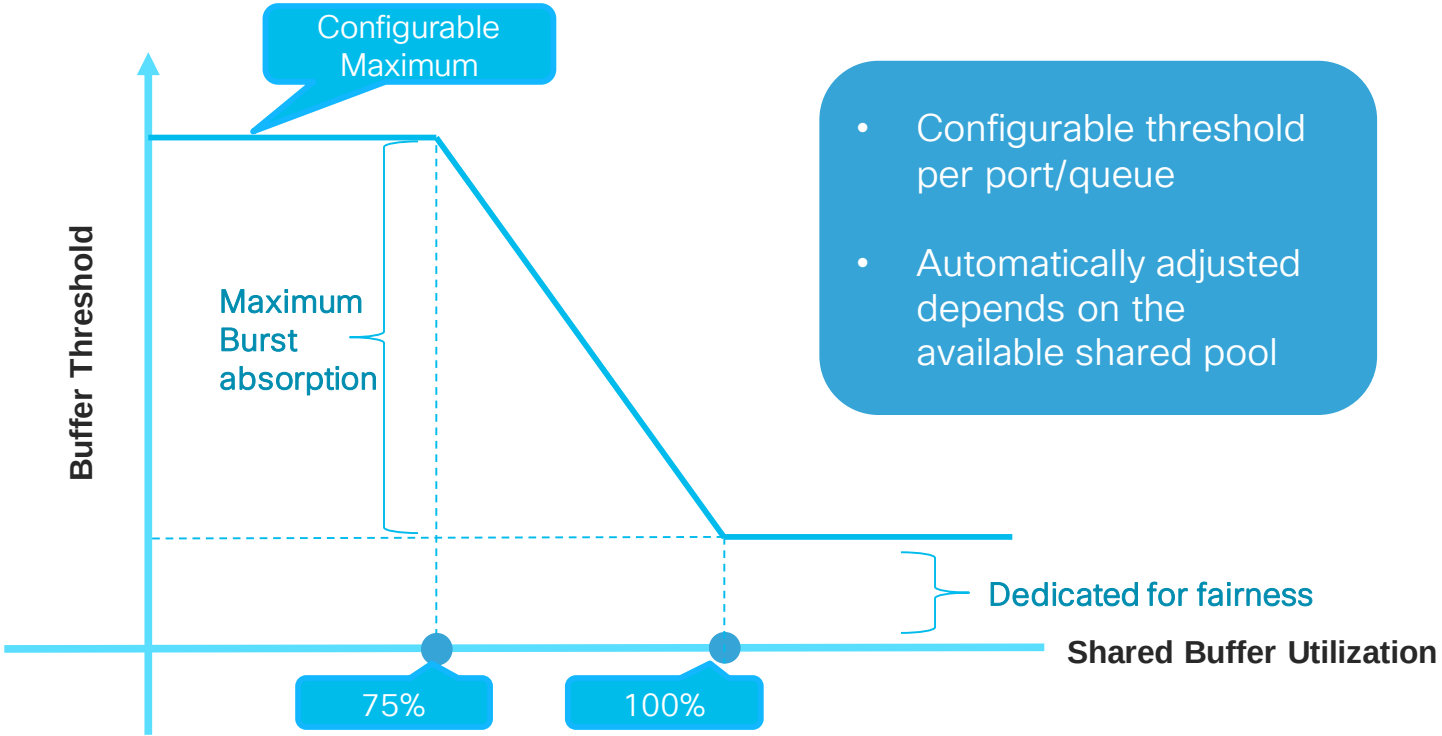

- Configurable threshold per port/queue
- Automatically adjusted depends on the available shared pool

Reference

Default buffer allocation per port speed

Platform	Port speed	100 Mbps, 1, 2.5, 5 Gbps		10 Gbps		25 Gbps		40 Gbps		100 Gbps	
	Queue	**Hard max**	**Soft max**	**Hard max**	**Soft max**	**Hard max**	**Soft max**	**Hard max**	**Soft max**	**Hard max**	**Soft max**
Cisco® Catalyst® 9300 Series	Q0	100	400	600	2400	–	–	2400	9600	–	–
Cisco Catalyst 9400 Series	Q0	176	700	176	700	–	–	176	700	–	–
Cisco Catalyst 9500 Series	Q0	200	800	1200	4800	–	–	4800	19,200	–	–
Cisco Catalyst 9500 High End	Q0	–	–	240	960	480	1920	720	2880	1920	7680
Cisco Catalyst 9600 – Sup1	Q0	–	–	240	960	480	1920	720	2880	1920	7680
		Soft min	**Soft max**	**Soft min**	**Soft max**	**Soft min**	**Soft max**	**Soft min**	**Soft max**	**Soft min**	**Soft max**
Cisco Catalyst 9300 Series	Q1	150	600	300	1200	–	–	3600	14,400	–	–
Cisco Catalyst 9400 Series	Q1	225	3600	264	1056	–	–	337	10,800	–	–
Cisco Catalyst 9500 Series	Q1	800	3600	1800	7200	–	–	7200	28,800		–
Cisco Catalyst 9500 High End	Q1	–	–	360	1440	720	2880	1080	4320	2880	11,520
Cisco Catalyst 9600 – Sup1	Q1	–	–	360	1440	720	2880	1080	4320	2880	11,520

Notes:
All allocation in units (each unit is 256-byte storage)
Hard = Dedicated; Soft=Shared
Q0: Soft max = 4x hard max
Q1: Soft max = 4x soft min
Hard max (hard buffer allocation): Do not participate in DTS/priority queue only

Port speed	Buffer (KB)	Number of buffers
100G	1200	4800
40G	450	1800
25G	300	1200
10G	150	600
1GE	70	280

Cisco Catalyst 9600 Series – Traffic policing

1 rate 2 color

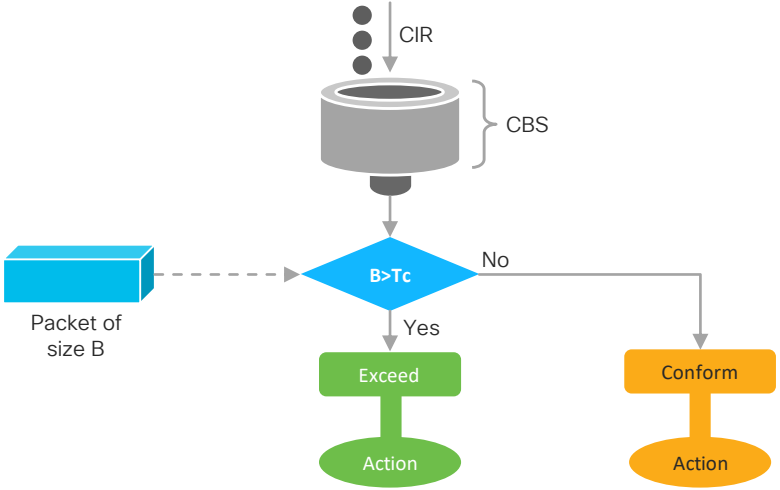


```
police cir 100000000 bc 3125000
  conform-action set-dscp-transmit af41
  exceed-action drop
```

2 rate 3 color

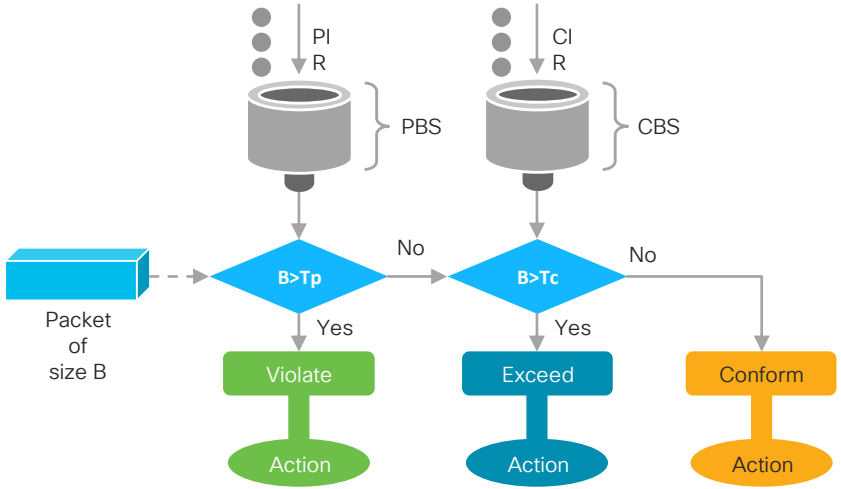


```
police cir percent 10 pir percent 50
  conform-action transmit
  exceed-action set-dscp-transmit dscp table MARKDOWN
  violate-action drop
```

CIR – Committed Information Rate
PIR – Peak Information Rate

PBS- Peak Burst Size
CBS – Committed Burst Size

Cisco Catalyst 9600 Series – Hierarchical QoS (HQoS)

HQoS (two-level hierarchy) allows you to perform the following functions:

- Classification
- Policing
- Shaping

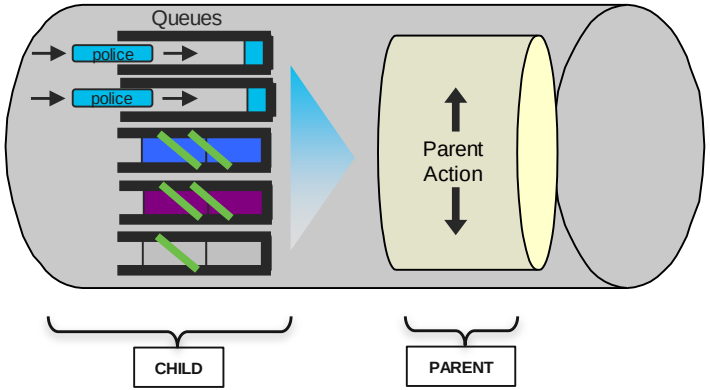

Child Action	Parent Action
Classification + Policing	Shaping
	Marking
Classification + Marking	Policing
	Shaping

Security

You make networking **possible**

Cisco live!

Cisco Catalyst 9000 Platform Trustworthy Solutions

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

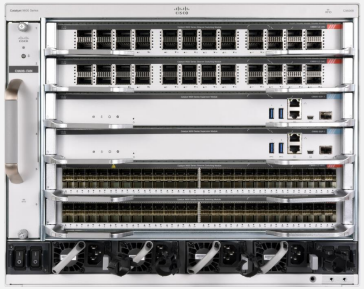

Secure boot
Boot sequence check

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

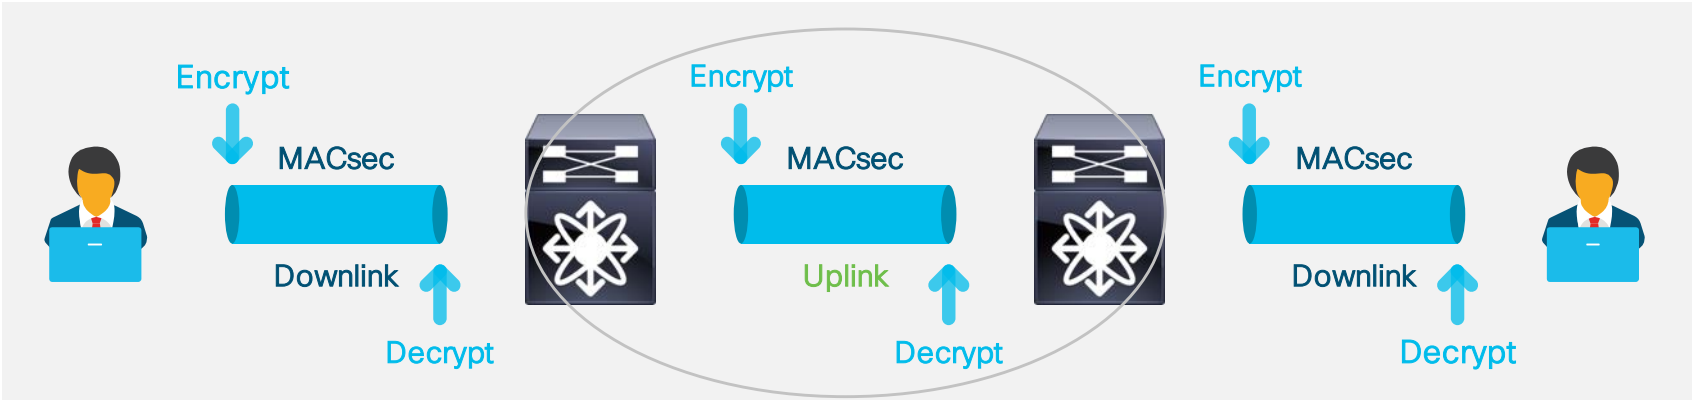

- Packets are encrypted on egress; decrypted on ingress
- Offers line-rate encryption on all ports and speeds (1G, 10G, 25G, 40G, and 100G)
- Transparent to all upper-layer protocols
- Supports switch-to-switch and switch-to-host MACsec
- 256-bit MACsec-capable between switch to switch
- Manual or 802.1X modes supported

Switch and Switch:
- 128 bit: MKA, SAP
- 256 bit: MKA
- Passthrough / ClearTag

Switch and Host:
- 128 bit: MKA
- 256 bit: MKA

ETA – Finding Malicious Activity in Encrypted Traffic

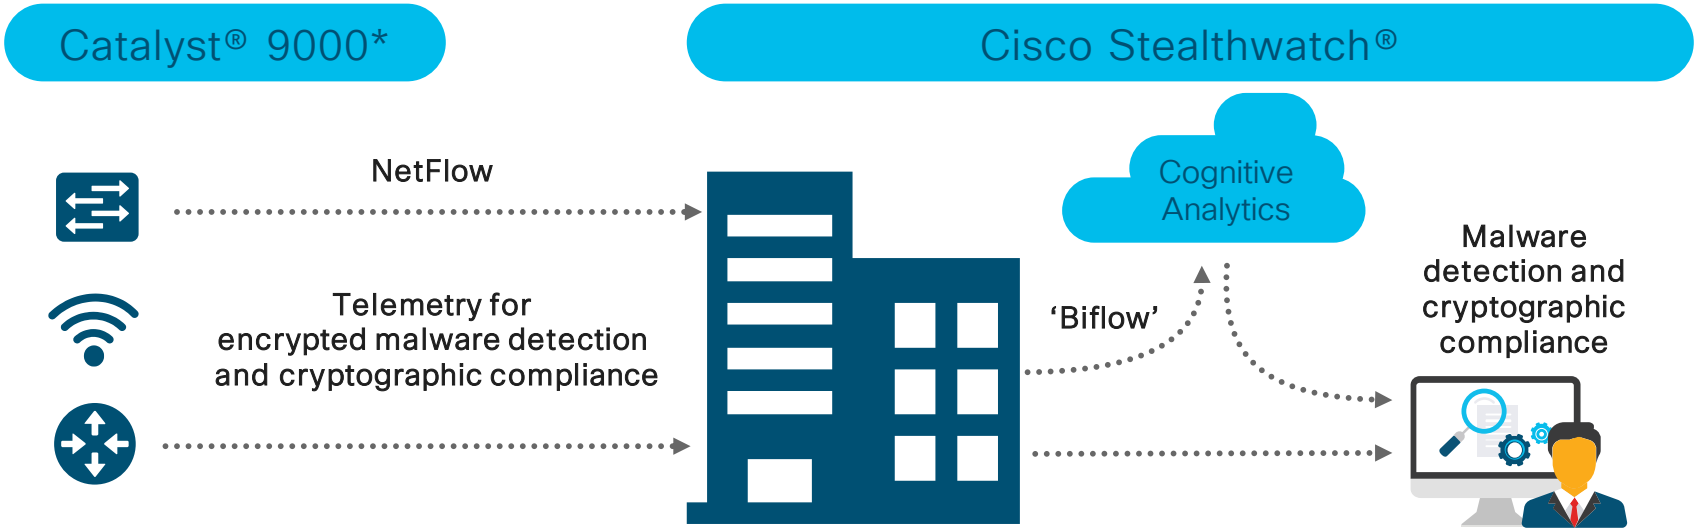

Leveraged network	Faster investigation	Higher precision	Stronger protection
Enhanced NetFlow from Cisco's newest switches and routers	Enhanced analytics and machine learning	Global-to-local knowledge correlation	Continuous enterprise-wide compliance

*ETA support for the Catalyst 9600 is on the roadmap

High Availability

You make networking **possible**

Cisco live!

High availability
Protect business continuity

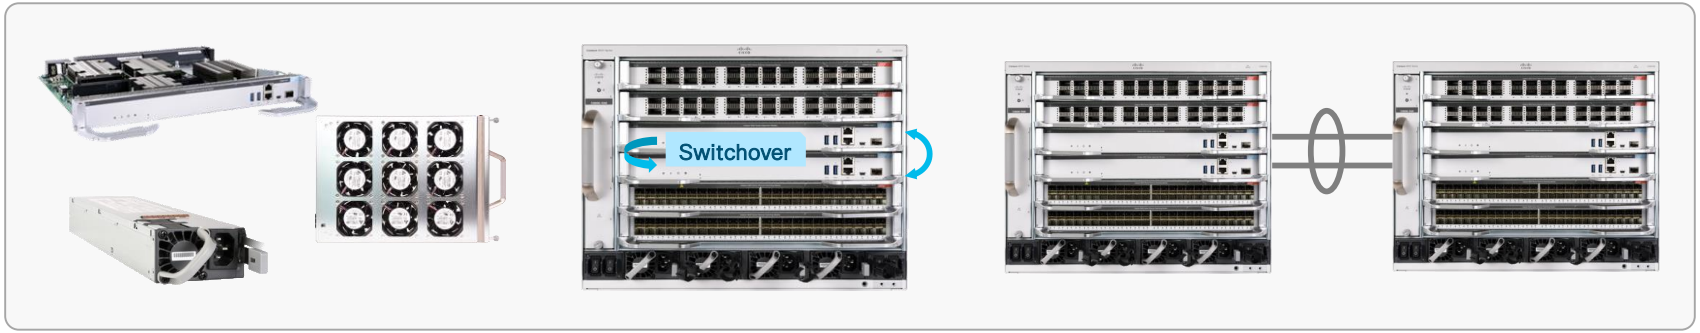

Physical redundancy	Stateful Switchover (SSO)	Non-Stop Forwarding (NSF)	In-Service Software Upgrade (ISSU)	StackWise®-Virtual*
Redundant hardware • Redundant power supplies • Redundant fan in the fan tray • Redundant supervisors	**Sub-second failover** • Between supervisors within chassis (<5ms) • Between chassis with StackWise-Virtual *	**Resilient L3 topologies** • NSF support for OSPF, EIGRP, ISIS, BGP	**Minimize upgrade downtime** • SMU • ISSU • GIR *	**Infrastructure resilience** • Multi-chassis EtherChannel (MEC) provides hardware-based failover

* Roadmap

SSO – Stateful Switchover

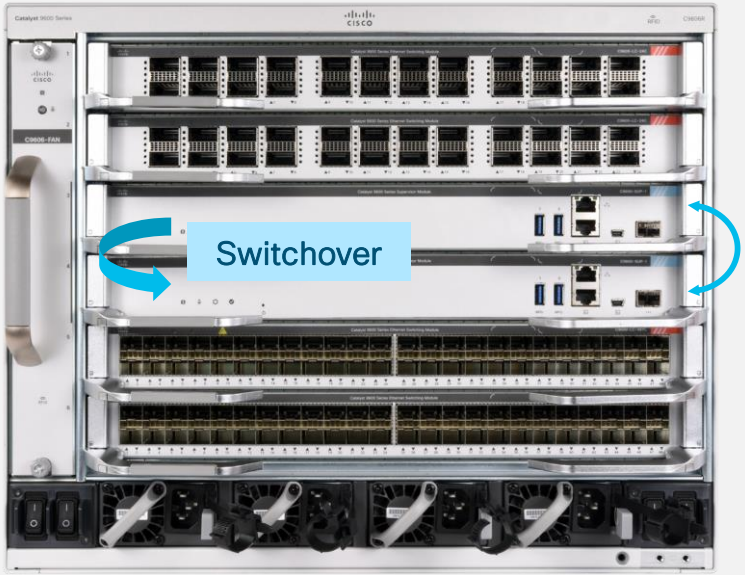

SSO is the default and only redundancy mode with two supervisors in the system

1. The active supervisor is responsible for all control plane processing
2. The active supervisor is responsible for hardware programming on both the active and standby supervisors

Supervisors and line cards: Data path

Receiving

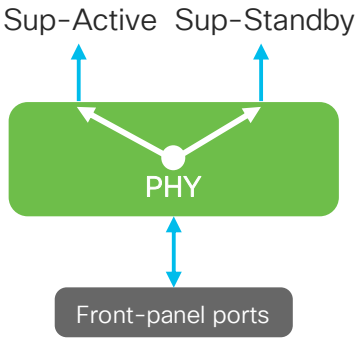

Transmitting

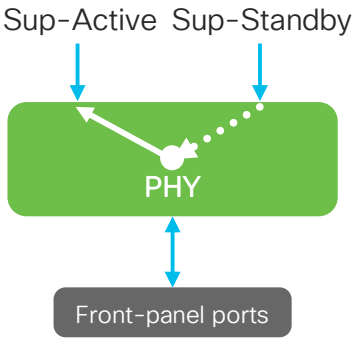

- Both active and standby supervisors receive data from line cards
- Line cards select the transmitting data from the active supervisor

Hitless supervisor switchover

StackWise Virtual

Topology Comparisons

Traditional

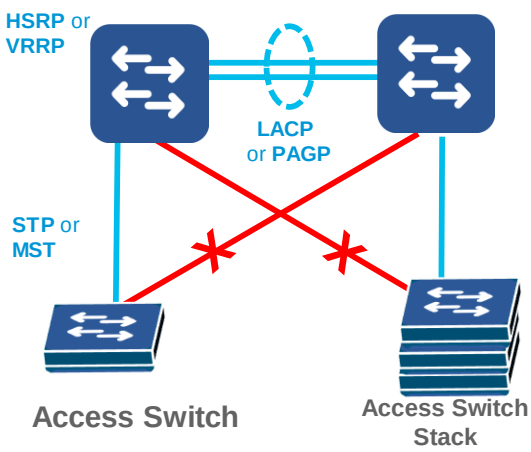

SV - Physical

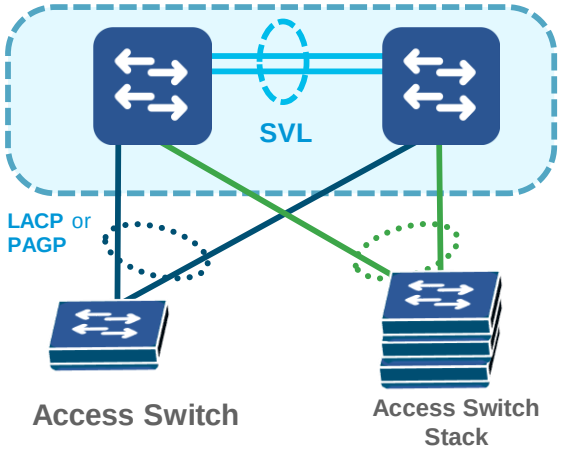

SV- Logical

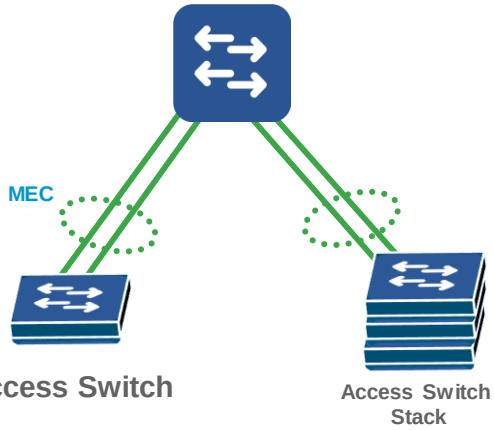

Benefits of StackWise Virtual

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 – C9600 *

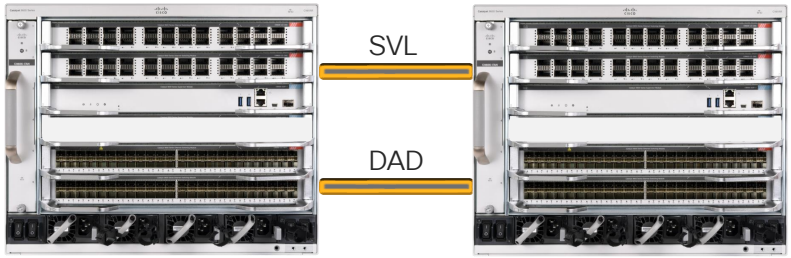

* StackWise Virtual is currently targeted for IOSXE 16.12.1

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

Cisco Catalyst 9000 Series ISSU workflow

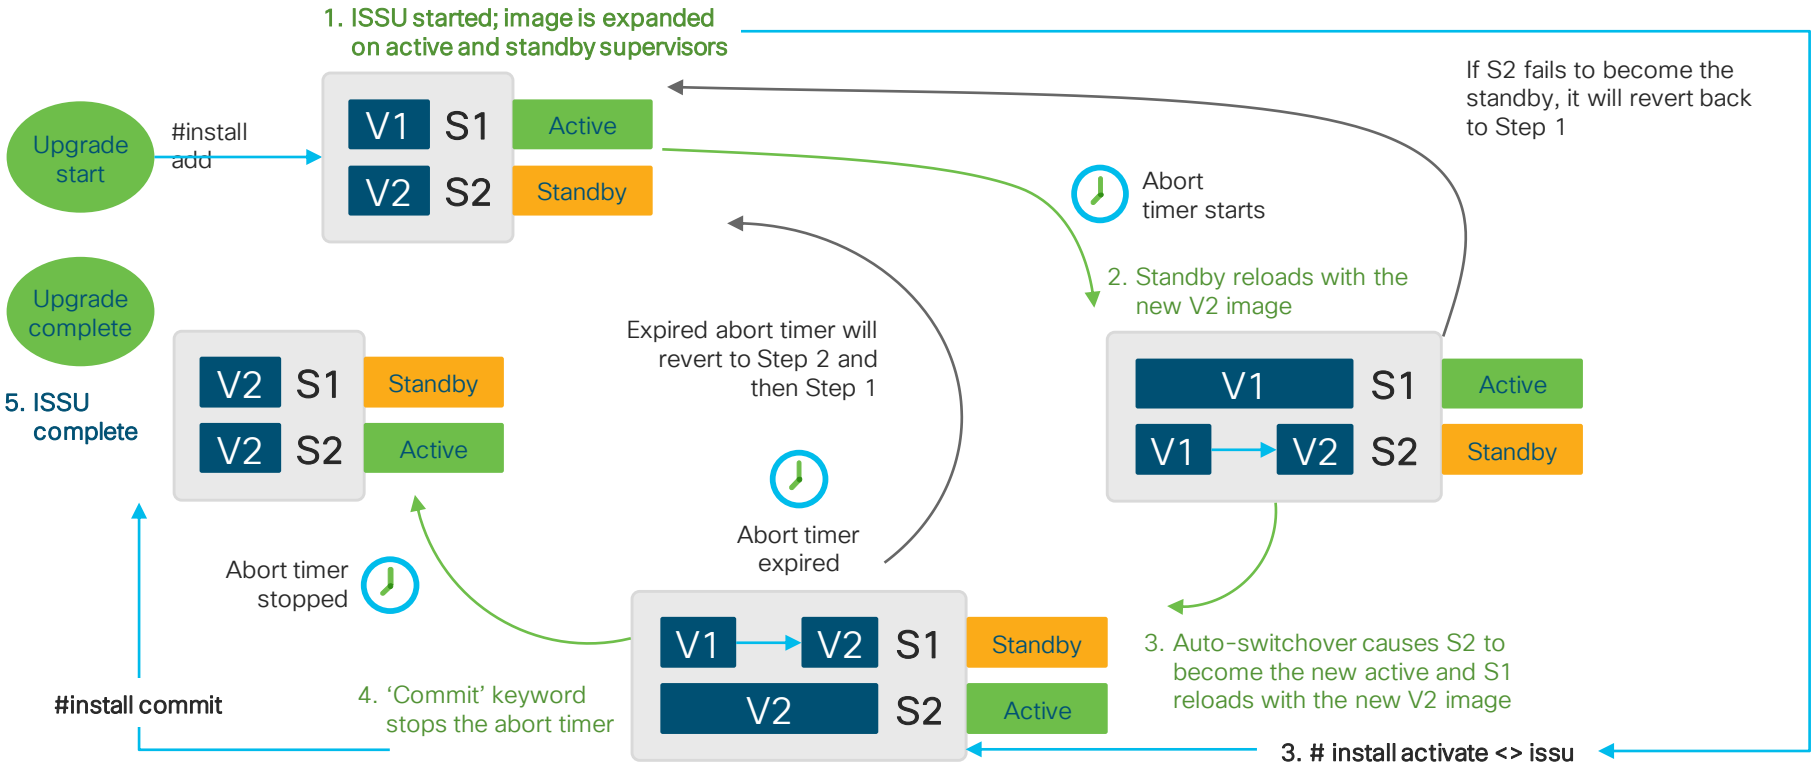

* ISSU is currently targeted for IOSXE 16.12.1

ISSU Upgrade steps

Three-step process:

- Install add file <tftp/ftp/flash/disk:*.bin>
- Install activate ISSU
- Install commit

Granular control on the upgrade process with the ability to roll back

One-step process:

- Install add file <tftp/ftp/flash/disk:*.bin>activate ISSU commit

Single command to perform a complete ISSU

Install Command-Line Interface (CLI) commands

Step-by-step workflow:

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

install activate issu

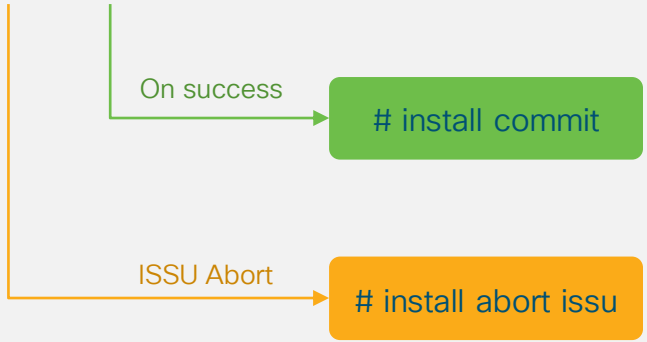

Workflow steps details:

- The **"Install Add"** command will perform the image download from the cisco.com posted location
- The **"Install activate"** command will upgrade the chassis with a new software version
- The **"Install commit"** command makes the changes permanent and deletes the older version of software from the chassis
- **"install abort issu"**: The customer can issue the abort command to revert the software back to the original state

Catalyst 9600 Design Consideration

You make networking **possible**

Cisco live!

Catalyst 9600 for Multidomain Campus Core

L3/Collapsed Core

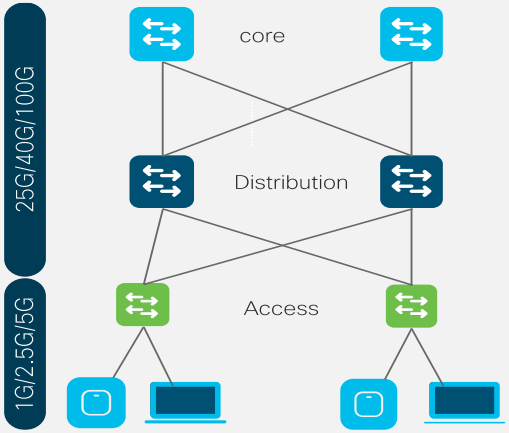

- Reduced Complexity, Resiliency & Scale
- IPv4/IPv6, Unicast & Multicast, QoS & ACL Scale

L3 Core + MPLS PE

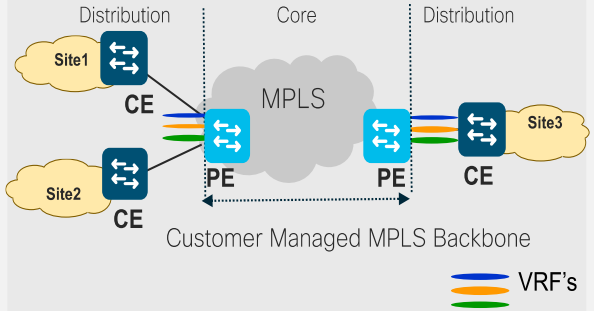

- Segmentation, Scale, LAN/WAN Consistency
- MPLS VPNs (L2 & L3), MPLS over GRE

Fabric Border

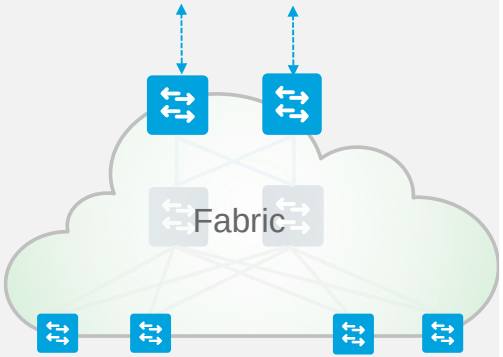

- Segmentation & Automation/Programmability
- Turnkey Solution: SD-Access
- DIY : BGP EVPN VxLAN

One Platform. Any Place. Any Speed (1G to 100G)

Example of a Traditional Three-Tiers Campus Design

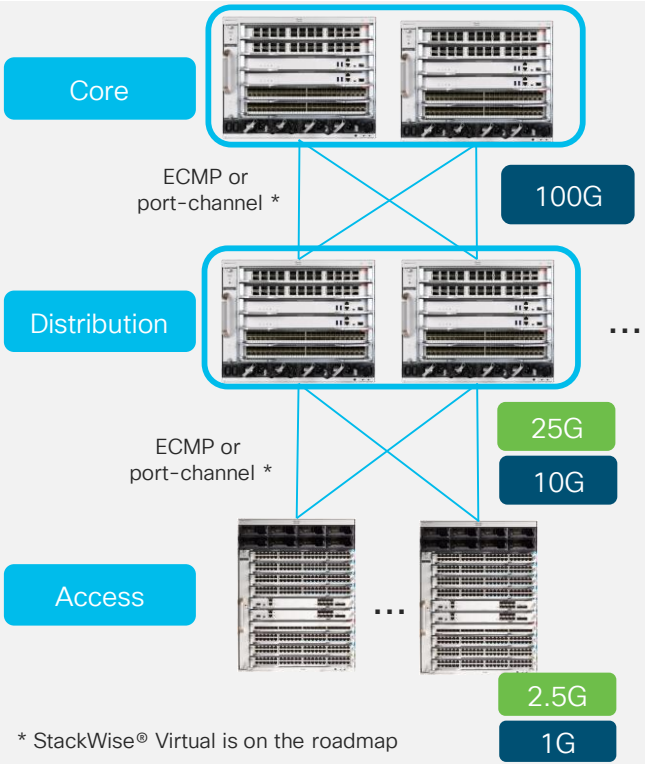

- Catalyst® 9606R Switch with 4x C9600-LC-24C for connection to the distribution layer, WAN, and between core devices
- Oversubscription ratio of 4:1 from distribution to core

Core

- Catalyst 9606R with 1x C9606-LC-24C as the uplink connection to the core and 3x C9606-LC-48YL as the downlink connection to the access layer
- Oversubscription ratio of 20:1 from access to distribution

Distribution

- Catalyst 9410 with 1G copper and/or mGig modules with Sup1-XL-Y
- 8x Catalyst 9300 StackWise-480 with 10G/25G uplinks

Access

Access Layer – POD

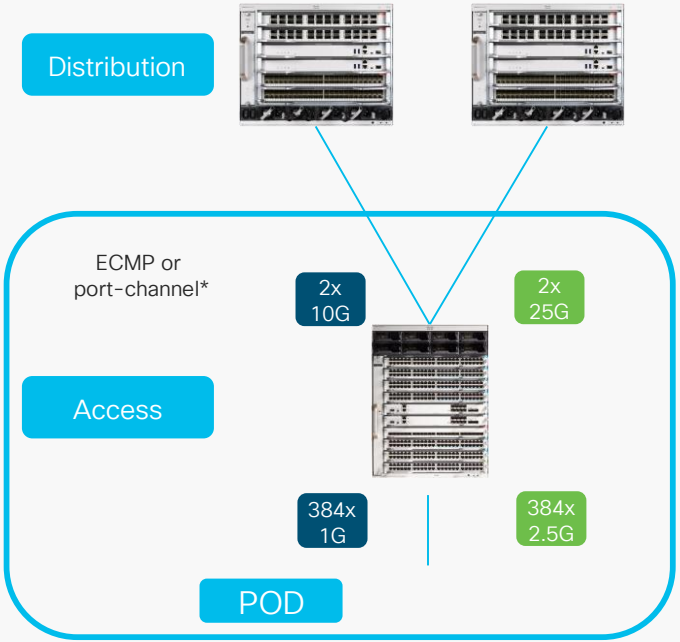

* StackWise® Virtual is on the roadmap

Cisco® Catalyst® 9400

- A Catalyst 9410 switch provides a total of 384 ports of 1G
- Catalyst 9410 can also provide 192x1G + 192x mGig ports (up to 10G)

Catalyst 9300 StackWise®-480

- Stack of 8 can provide a total of 384 ports of 1G or 2.5G (mGig)
- Stack of 8 can also provide 384 ports of 1G and mGig combination

Aggregated downlink BW:

- 384G with 384x 1G
- 960G with 384x 2.5G

Uplinks BW needed for 20:1 oversubscription from access to distribution

- 2x 10G for 384x 1G
- 2x 25G for 384x 2.5G

Distribution Layer – Block

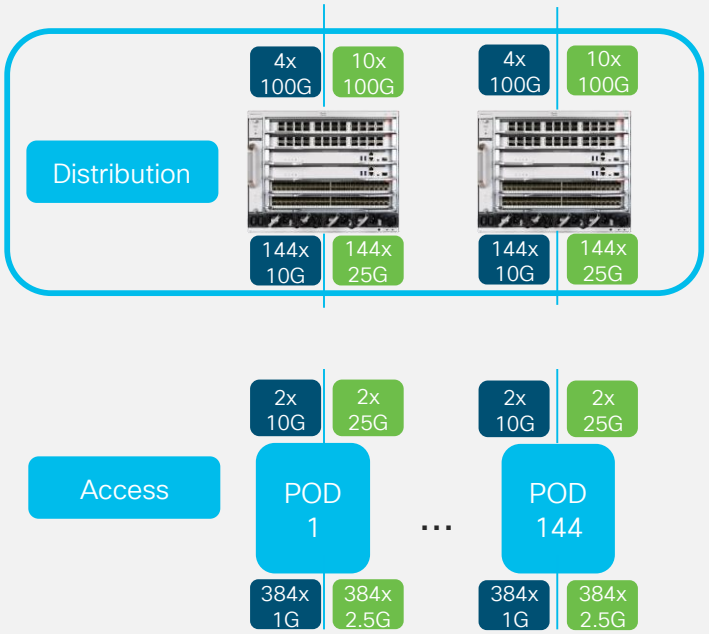

* StackWise® Virtual is on the roadmap

Cisco® Catalyst® 9606R Switch Downlinks:

- 3x C9600-LC-48YL per Catalyst 9606R
- A total of 144 x 10G/25G ports per chassis
- Aggregate downlink BW per Catalyst 9606 Switch
  1. With 10G uplinks: 144x 10G = 1.44T
  2. With 25G uplinks: 144x 25G = 3.6T

Uplinks

- 1x C9606-LC-24C per Catalyst 9606R
- To maintain 4:1 oversubscription between distribution and core layers
  1. With 10G uplinks: BW = 1.44T/4 = 360G => 4x 100G ports
  2. With 25G uplinks: BW = 3.6T/4 = 900G => 10x 100G ports

(The remaining 100G/40G ports can be used for ECMP or StackWise Virtual when it is available.)

Each distribution block can aggregate 144 access PODs.
That's 144 x 384 = **55,296** of 1G, or 2.5G ports

Core Layer with 1G in the Access Layer

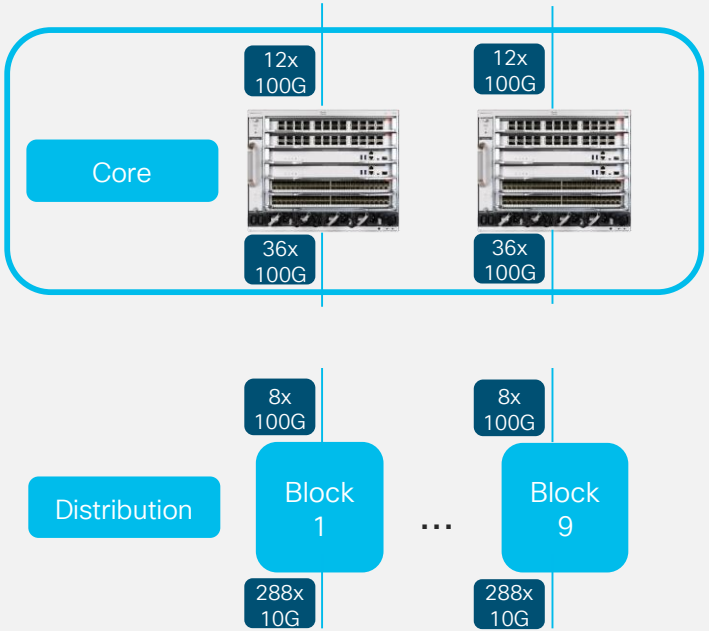

Cisco® Catalyst® 9606R Switch

- 4x C9606-LC-24C
  - 75% of ports (36x 100G) to distribution
  - 25% of ports (12x 100G) for connections between the two cores and the WAN
- Two of the core devices will provide 72x 100G for the distribution layer

- 1G aggregation
  - With 8x 100G per distribution block, two Catalyst 9606R Switches with the above configuration can aggregate 72/8, or 9 distribution blocks

The total number of 1G ports:
9x 144 x 384 = **497,664** of 1G ports

Core Layer with 2.5G in the Access Layer

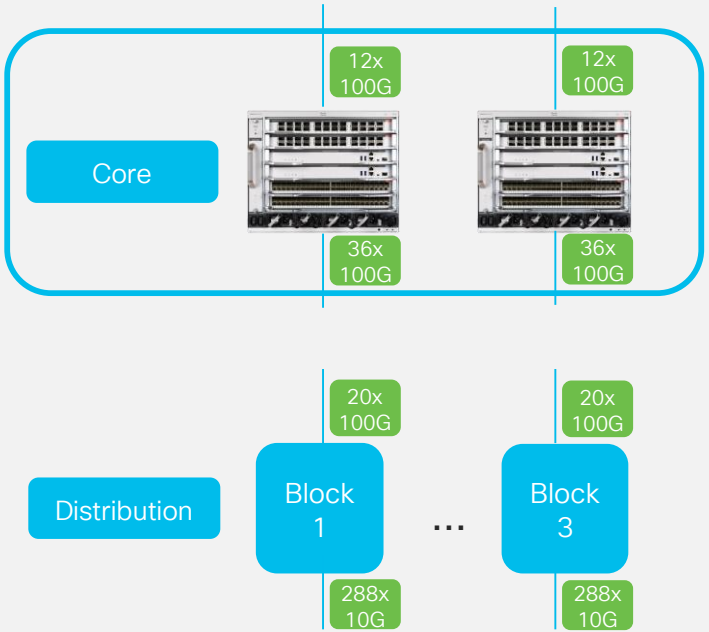

Cisco® Catalyst® 9606R Switch

- 4x C9606-LC-24C
  - 75% of ports (36x 100G) to distribution
  - 25% of ports (12x 100G) for connections between the two cores and the WAN
- Two of the core devices will provide 72x 100G for the distribution layer

- 2.5G aggregation
  - With 20x 100G per distribution block, two of Catalyst 9606R Switches with the above configuration can aggregate 72/20, or 3 distribution blocks

The total number of 2.5G ports:
3x 144 x 384 = **165,888** of 2.5G ports

Core Layer with 2.5G in the Access Layer

With 4x Catalyst 9606 in the core

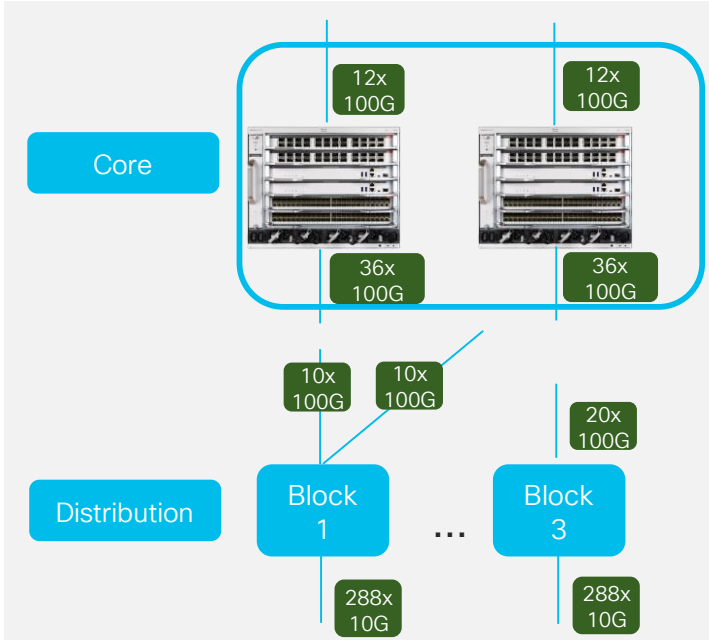

The total number of 2.5G ports:
3x 144 x 384 = **165,888** of 2.5G ports

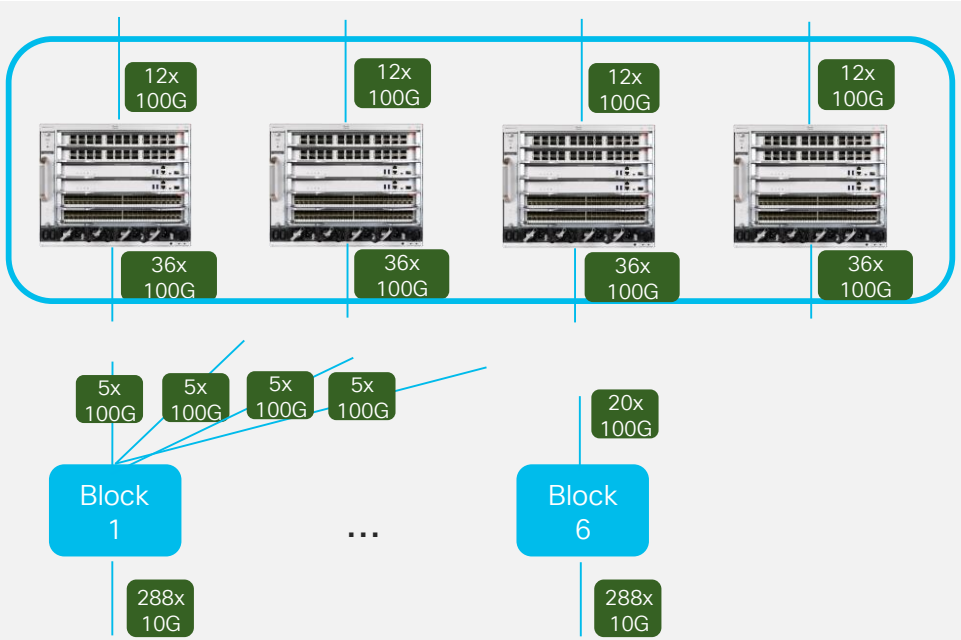

The total number of 2.5G ports:
6x 144 x 384 = **331,776** of 2.5G ports

Summary

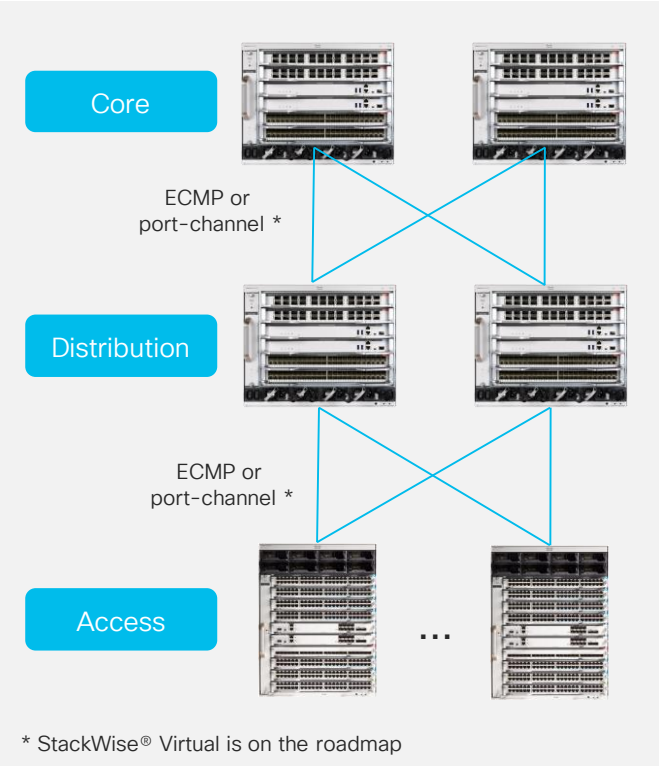

Two Cisco Catalyst 9606R Switches in the core can provide:

1. 497K of 1G ports, or
2. 165K of 2.5G ports

Oversubscription = 4:1		
Uplinks (40/100G module):	4x 100G	10x 100G
Downlinks (10/25G modules):	144x 10G	144x 25G

Oversubscription: 20:1		
Uplinks (Supervisor or uplink module):	2x 10G	2x 25G
Downlinks (1G/mGIG module):	384x 1G	384x 2.5G

Closing

You make networking **possible**

Cisco Catalyst 9600 Series Summary

Offering a comprehensive, high-density portfolio on campus with 100G, 40G, 25G, 10G

Architectural flexibility
- Broad support for 1G/10G, 25G, 40G, 100G from aggregation to core

Infrastructure investment protection
- Non-disruptive migration from 10G to 25G and beyond

Cost-effective optics
- Innovation in standards to support high-density, multilane optics

Platform benefits

Up to 1TB SSD storage

Customizable ASIC templates

Same Cisco IOS® XE image

Dual-serviceable fan tray

N:1 power supply redundancy

You make **possible**

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.

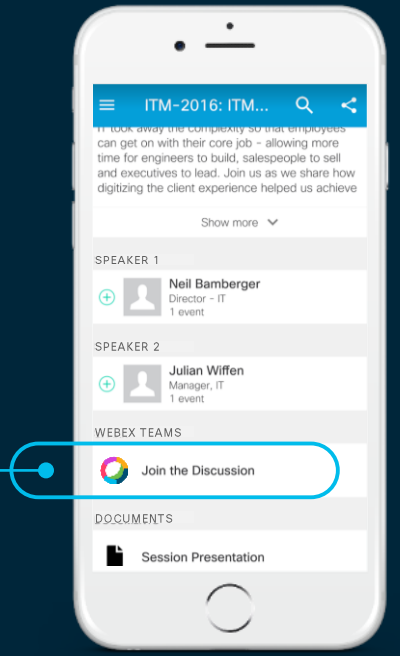

cs.co/ciscolivebot# BRKARC-3010

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

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

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

Thank you

You make **possible**