
TECHNICAL DECISION MAKER PRESENTATION

Cisco Catalyst 9600 Series

Prepared as a convenient PDF reference for network planning, platform evaluation, and procurement discussions.

Source presentation: Cisco. Layer23-Switch added this cover and closing resource page. Cisco and related marks belong to Cisco Systems, Inc. This material may include historical roadmap content; verify current specifications, lifecycle status, licensing, and availability before purchase or deployment.



Cisco Catalyst 9600 Series

Technical Decision Maker Presentation

August 26, 2026


Agenda

- Hardware
- Architecture
- Platform design
- SDM templates
- Cisco® UADP 3.0 ASIC
- Cisco® Silicon One ASIC
- Platforms Comparison
- C9600 App Hosting
- Competitive Info
- High Availability
- Security
- QoS
- ACL
- Multicast
- Catalyst 9000 Design
- Intent-based Networking
- Closing

Hardware

Cisco Catalyst 9000 Switching Portfolio

Adding the “X factor” to the industry’s leading switching family



Catalyst 9000 Modular Value Proposition

Key Modular Differentiation



Speed Transition

- 10/25/40G → 100/400G



Mission Critical HA

- Platform, OS, Architecture resiliency



Max Investment Protection

- 15-20 Year Lifecycle



Unmatched Flexibility

- Flexible templates, table sizes



Catalyst 9600

Cisco Catalyst 9600 Series Switches

Ready for Catalyst 6500/6800 migration with Industry-leading feature set

20 years of innovation
\$50B of install base



Catalyst 6500/6800

Resiliency

Security

Breadth of
services

Catalyst 6500/6800 –
Gold standard for campus core

Catalyst® 6500/6800 parity



L2, L3, Multicast, ACL



Resiliency: ISSU
and SVL



MPLS



SD-Access



IBN-ready



ETA*, Trustworthy Systems
MACsec-256 encryption,



Programmability/
telemetry



NBAR2 AVC*,
Assurance



Hot patching, GIR



EVPN BGP VxLAN

Catalyst 9600 Series — Raising the bar on campus core


Powered by
UADP 3.0/Silicon One
and open
Cisco IOS XE



Catalyst 9600
Series

* Roadmap

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)

3.2 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



Line Rate non-blocking


Port speed	Density with C9600-Sup-1	Density with C9600X-Sup-2	Maximum chassis density
400G	NA	8	32
200G	NA	16*	64
100G	48	96	128
50G	NA	224	224
40G	96	96	128
25G	192	224	224
10G	192	224	224
5G	192	NA	192
2.5G	192	NA	192
1G	192	24**	192

* Roadmap

** w. SFP-1G-SX/LH optic. In SVL setup, limit is only 24. Not 48.

Cisco Catalyst 9600 Series

Supervisor 1



4.8 Tbps (full duplex)

1.2 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

Addressing the Campus Core + Edge Market

More than just “Speeds & Feeds”



System Capacity

- **Higher Port Speeds**
enabling newer Ethernet speeds and types, such as 50G SFP and 400G QSFP
- **More Access Bandwidth**
higher density of 802.11ac moving to 11ax, HD Video (1080p) moving to UHD (4K)
- **Non-Blocking**
Core needs non-blocking (1:1) performance



System Scale

- **High-Scale Routing**
support for full IPv4 and IPv6 Internet + LAN + VRF routing
- **Flexible Transport Options**
support for MPLS-VPN, SD-Access and BGP-EVPN with 4K VRFs
- **Centralized Wireless**
support for high MAC and ARP/NDP scale



Port Density

- **More Uplinks**
support for multiple QSFP uplinks, with LR/ZR optics for redundant edge connections
- **More Downlinks**
support for hundreds of SFP ports/breakouts
- **Port Options**
mix of SFP & QSFP ports, at various interface speeds



Flexibility

- **Programmable ASIC**
Enable new features and encapsulations
- **Model-driven Microcode**
hardware functions use model-driven APIs
- **App Hosting Infrastructure**
support Docker applications
- **Architecture-level HA**
L2/L3 redundancy with Stackwise Virtual and GIR



Core + Edge



Core + Distro



Cisco Silicon One™



Switching Silicon

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



Routing Silicon

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

Cisco Silicon One Bringing Switching and Routing convergence

Cisco Silicon One™ Q200

Industry leading Switching and Routing Silicon



12.8T BW



8.1 Bpps



8G HBM for
deep buffers



10M IPv4
or 5M IPv6
route scale



Fully P4
Programmable
Pipeline



50G PAM4
Serdes

Cisco Silicon ONE Q200

**Industry Leading
12.8T System on Chip**



**First 7nm ASIC providing lowest
watts/GE power consumption**



**Fully P4 programmable enabling
feature velocity**




**Multi slice architecture for
flexibility and scale**

Routing Capabilities with Switching Power and Performance

Catalyst 9600 Series

Extending modular core with a performance-optimized Supervisor 2

UADP 3.0



Supervisor 1


Total capacity

4.8 Tbps

Slot bandwidth

1.2 Tbps

Gen 1
1.2T per slot



Gen 2
3.2T per slot

Gen 2
3.2T per slot

Gen 1
1.2T per slot

NEW

Silicon
One™
Q200



Supervisor 2

Total capacity

12.8 Tbps

Slot bandwidth

3.2 Tbps

>2.5X

Cisco Catalyst 9600 Series Supervisor 2



12.8 Tbps (full duplex)

3.2 Tbps per slot

1x Silicon One™ Q200 ASIC

8 core X86 CPU
@2.7 Ghz

M.2 SATA SSD
(optional: up to 1 TB)

32G DDR4 memory

Built-in RFID

Mgmt ports: copper
and 2x10G **fiber**

Blue Beacon

2x USB3
1x mini-B USB console


Cisco Catalyst 9600 Series

Line cards

Gen1 line cards


C9600-LC-24C - 100G/40G (fiber)

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



C9600-LC-48YL - 50G/ 25G/10G/1G (fiber)

- 48 ports
- SFP56/SFP28/SFP+
- Supports 50G,25G,10G,and 1G



C9600-LC-48TX - mGig (copper)

- 48 ports
- Copper 10G (NBASE-T/10BASE-T)
- Supports 10G,5G,2.5G,1G,100M and 10M



C9600-LC-48S- 1G (fiber)


- 48 ports
- SFP
- Supports 1G



Gen2 line cards

C9600-LC-40YL4CD - 400G/200G³/100G/50G/40G/ 25G/10G (fiber)


- 40+2+2 ports
- SFP56/QSFP56+/QSFPDD
- Supports 400G,200G³,100G,50G,40G,25G and 10G



17.7.1

C9600X-LC-32CD - 400G/200G³/100G/40G(fiber)

- 30+2 ports
- QSFP28/QSFPDD
- Supports 400G,200G³,100G, and 40G



17.9.1

C9600X-LC-56YL4C - 100G/50G/40G/25G/10G (fiber)

- 56 50G+4 100G ports
- SFP56/QSFP28
- Supports 100G, 50G, 40G, 25G, 10G




17.13.1

³ HW Capable

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)
- With Supervisor Engine 2
 - All ports support 40G or 100G (line rate with 136 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 Port Speeds 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

C9600-LC-24C – Port Numbering with Supervisor Engine 2

- 100G/40G numbering from 1 to 24


**Port
Numbering**



interface HundredGigE1/0/1-24

Cisco Catalyst 9600 Series

50G/25G/10G/1G Line card - C9600-LC-48YL




- All 48 ports support 25G/10G/1G (Sup1)
- All 48 ports support with 50G,25G,10G (Sup2)
- Hardware capable of 10/100M
- Line rate with 25G/10G/1G (at 187 bytes for 25G; any packet size with 10G/1G)
- Port reference is always "TwentyFive<slot#>/0/<port#>" and port speed is auto-detected based on the inserted transceiver (Sup 1)
- Port reference is always "FiftyGigE<slot#>/0/<port#>" and port speed is auto-detected based on the inserted transceiver (Sup 2)

Cisco Catalyst 9600 Series

mGig Line Card – C9600-LC-48TX

For collapsed core deployment and connectivity to application servers



- All 48 ports can support 10G/5G/2.5G/1G/100M/10M (Sup1)
- All 48 ports can support 10G (Sup2)
- Line rate on all ports. Any port, any supported speed.
- No PoE
- Port reference is "Ten<slot#>/0/<port#>" and port speed is auto (default).

Cisco Catalyst 9600 Series

1G Line card - C9600-LC-48S



- All 48 ports support 1G (Sup1 Only)
- Hardware capable of 10/100M*
- Line rate with 1G (any packet size with 1G)
- Port reference is always "Gig<slot#>/0/<port#>".

* Roadmap

Catalyst 9600/X: New Combo Line Cards

Catalyst® 9606R



C9600-LC-24C



C9600-LC-48YL



C9600-LC-48S



C9600-LC-48TX
(Multigigabit)



Combo line cards

C9600-LC-40YL4CD



IOS XE 17.7.1

3.2 Tbps with SUP-2
1.2 Tbps with SUP-1

C9600X-LC-32CD



IOS XE 17.9.1

With SUP-2 only
(3.2Tbps)

C9600X-LC-56YL4C



IOS XE 17.13.1

Provides for **uplink requirement**
without using additional slot



Dual-personality
line card



Flexible speed
support




Lower total BoM
cost, save on
additional line
cards



Cisco Catalyst 9600 Series

400G/200G/100G/50G/25G/10G Line card - C9600-LC-40YL4CD



- 40x 50G/25G/10G + 2x200G*/100G/40G + 2x400G/200G*/100G/40G
- 1.2 Tbps with Sup 1 (IOS-XE 17.8.1)
 - 40x SFP28 ports - 1/10/25GE
 - 2x QSFP28 ports - 40/100GE
- 3.2 Tbps with Sup 2 (IOS-XE 17.7.1)
 - 40x SFP56 ports - 10/25/50*GE
 - 2x QSFP56 ports - 40/100/200*GE
 - 2x QSPDD ports - 40/100/200*/400GE
- **CDR5M** PHYs
 - MACsec, WAN-MACsec, ClearTag v3.4
 - IEEE 1588 & PTPv2*
 - Hitless MUX (HMUX)


* Roadmap

C9600-LC-40YL4CD ports and speeds support

with SUP-2: 3.2 Tbps

40x 10/25/50G + 2x 40/100/200G* / 2x 40/100/200*/400G


Cisco IOS® XE 17.7.1



with SUP-1: 1.2 Tbps

40x 1/10/25G + 2x 40/100G

Cisco IOS XE 17.8.1




* Roadmap

Port-group Command-Line Interface (CLI)

Q200 SerDes has a limit that requires adjacent SFP ports to be the same speed

- Requires a new configuration for a “port group” of two ports to set the same speed
- Default selection is the per-module native port speed (for example, 50G for 40YL4CD)



[no] hw-module slot <1-6> port-group <1-20|range> select <10G|25G|50G>

```
C9600X-SUP2(config)# hw-mod slot 6 ?
logging      Logging onboard
port-group   configure feature on Port-group
shutdown     Shutdown a fru

C9600X-SUP2(config)# hw-mod slot 6 port-group ?
<1-20>      port-group to apply feature
range       Range of port-group

C9600X-SUP2(config)# hw-mod slot 6 port-group 1 ?
select      configure selected speed

C9600X-SUP2(config)# hw-mod slot 6 port-group 1 select ?
10G         Speed of 10Gbps
25G         Speed of 25Gbps
50G         Speed of 50Gbps

C9600X-SUP2(config)# hw-mod slot 6 port-group 1 select 10g ?
<cr>       <cr>

C9600X-SUP2(config)# hw-mod slot 6 port-group 1 select 10g
C9600X-SUP2(config)#end
C9600X-SUP2#
```

```
C9600X-SUP2# sho hw-mod slot 6 port-group
-----
Port-group Ports                               Speed
-----
1             Fif6/0/1, Fif6/0/2               10G
2             Fif6/0/3, Fif6/0/4               NONE
3             Fif6/0/5, Fif6/0/6               NONE
4             Fif6/0/7, Fif6/0/8               NONE
5             Fif6/0/9, Fif6/0/10              NONE
6             Fif6/0/11, Fif6/0/12              NONE
7             Fif6/0/13, Fif6/0/14              NONE
8             Fif6/0/15, Fif6/0/16              NONE
9             Fif6/0/17, Fif6/0/18              NONE
10            Fif6/0/19, Fif6/0/20              NONE
11            Fif6/0/21, Fif6/0/22              NONE
12            Fif6/0/23, Fif6/0/24              NONE
13            Fif6/0/25, Fif6/0/26              NONE
14            Fif6/0/27, Fif6/0/28              NONE
15            Fif6/0/29, Fif6/0/30              NONE
16            Fif6/0/31, Fif6/0/32              NONE
17            Fif6/0/33, Fif6/0/34              NONE
18            Fif6/0/35, Fif6/0/36              NONE
19            Fif6/0/37, Fif6/0/38              NONE
20            Fif6/0/39, Fif6/0/40              NONE

C9600X-SUP2#
```

C9600X-LC-32CD

Gen2 Module – 30 x QSFP28 + 2 x QSFP-DD

3.2
Tbps


- **3.2Tbps with Supervisor 2 Only**
 - 30 x QSFP28 ports, 40/100GE
 - 2 x QSFP-DD ports, 40/100/200*/400GE
- **NOT supported with Supervisor 1**
- 8 x **CDR5M** + 2 x **CDR5L** PHYs
 - MACSec, WAN-MACSec, ClearTag v3.4
 - IEEE 1588 & PTPv2 (Precision Time Protocol)
 - Hitless MUX (HMUX)



Supervisor 2



Supervisor 1




* SW Roadmap (FCS+)


C9600X-LC-32CD

100G / 400G Port Modes

100G (QSFP28) Mode (Default):



Two 400G (QSFP-DD) Mode:



New Port-Group CLI

C9600X-LC-32CD



Enabling 1 x 400G ports – will disable 3 x 100G ports in that port group.

- Default mode will be 100G – use port-group command to enable 400G mode
- Use the 'no' form to disable 400G mode and re-enable 4 x 100G mode

[no] hw-module slot <1-6> port-group <1-2|range> select <400G>

```
C9600X-SUP2(config)# hw-mod slot 5 port-group ?
```

```
<1-2> port-group to apply feature
range Range of port-group
```

```
C9600X-SUP2(config)# hw-mod slot 5 port-group 1 ?
```

```
mode configure selected speed
```

```
C9600X-SUP2(config)# hw-mod slot 5 port-group 1 mode ?
```

```
400G Enable 400G port & disable 100G ports
```

```
C9600X-SUP2(config)# hw-mod slot 5 port-group 1 mode 400g ?
```

```
<cr> <cr>
```

```
C9600X-SUP2(config)# hw-mod slot 5 port-group 1 mode 400g
```

```
< OUTPUT TO BE DETERMINED >
```

```
C9600X-SUP2(config)#end
```

```
C9600X-SUP2#
```

```
C9600X-SUP2# show running | include hw-module slot 5 port-group
```

```
hw-module slot 5 port-group 1 mode 400G
```

```
C9600X-SUP2#
```

```
C9600X-SUP2# show interface status module 5 | inc Port|5/0/
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Hu5/0/25		inactive	1	full	100G	QSFP 100G AOC5M
Hu5/0/26		inactive	1	full	100G	QSFP 100G AOC3M
Fou5/0/27		connected	1	full	100G	QSFP-DD 400G CU2_5M
Hu5/0/28		inactive	1	auto	a-100G	unknown
Hu5/0/29		notconnect	routed	full	100G	QSFP 100G AOC3M
Hu5/0/30		notconnect	routed	full	40G	QSFP 40G AOC3M
Fou5/0/31		notconnect	routed	full	100G	QSFP 100G CU1M
Hu5/0/32		notconnect	routed	full	100G	QSFP 100G CU3M

```
C9600X-SUP2#
```

C9600X-LC-56YL4C

Gen2 Module – 56 x SFP56 + 4 x QSFP

3.2
Tbps


- **3.2Tbps with Supervisor 2 Only**
 - 56 x SFP56 ports, 10/25/50GE
 - 4 x QSFP ports, 40/100GE
- **NOT supported with Supervisor 1**
- **8 x CDR5M 2.0 PHYs**
 - IPsec*, MACSec, WAN-MACSec, ClearTag v3.4
 - IEEE 1588 & PTPv2 (Precision Time Protocol)**
 - Hitless MUX (HMUX)



Supervisor 2




Supervisor 1



* 400G IPSEC HW Capable

** Post-FCS Features

C9600X-LC-56YL4C Transceiver Power



Port 1-8 : Up to 2.5W

Port 9-56 : Up to 1.5W

Syslog Message up on insertion of high-power optic (>1.5W)


Advice to leave adjacent port empty to avoid thermal issue if
high-power optic is used

Port 57-60 : No restrictions on transceiver power

Port-group Command-Line Interface (CLI)

Q200 SerDes has a limit that requires adjacent SFP ports to be the same speed

- Requires a new configuration for a “port group” of two ports to set the same speed
- Default selection is the per-module native port speed (for example, 50G for 56YL4C)



```
[no] hw-module slot <1-6> port-group <1-20|range> select <10G|25G|50G>
```

```
C9600X-SUP-2(config)# hw-mod slot 6 ?
logging      Logging onboard
port-group   configure feature on Port-group
shutdown     Shutdown a fru

C9600X-SUP-2(config)# hw-mod slot 6 port-group ?
<1-20>      port-group to apply feature
range       Range of port-group

C9600X-SUP-2(config)# hw-mod slot 6 port-group 1 ?
select      configure selected speed

C9600X-SUP-2(config)# hw-mod slot 6 port-group 1 select ?
10G         Speed of 10Gbps
25G         Speed of 25Gbps
50G         Speed of 50Gbps

C9600X-SUP-2(config)# hw-mod slot 6 port-group 1 select 10g ?
<cr>       <cr>


C9600X-SUP-2(config)# hw-mod slot 6 port-group 1 select 10g
C9600X-SUP-2(config)#end
C9600X-SUP-2#
```

```
C9600X-SUP-2# sho hw-mod slot 6 port-group
```

Port-group	Ports	Speed
1	Fif6/0/1, Fif6/0/2	10G
2	Fif6/0/3, Fif6/0/4	NONE
3	Fif6/0/5, Fif6/0/6	NONE
4	Fif6/0/7, Fif6/0/8	NONE
5	Fif6/0/9, Fif6/0/10	NONE
6	Fif6/0/11, Fif6/0/12	NONE
7	Fif6/0/13, Fif6/0/14	NONE
8	Fif6/0/15, Fif6/0/16	NONE
9	Fif6/0/17, Fif6/0/18	NONE
.		
.		
.		
21	Fif6/0/41, Fif6/0/42	NONE
22	Fif6/0/43, Fif6/0/44	NONE
23	Fif6/0/45, Fif6/0/46	NONE
24	Fif6/0/47, Fif6/0/48	NONE
25	Fif6/0/49, Fif6/0/50	NONE
26	Fif6/0/51, Fif6/0/52	NONE
27	Fif6/0/53, Fif6/0/54	NONE
28	Fif6/0/55, Fif6/0/56	NONE


```
C9600X-SUP-2#
```

Gen 1 line cards support with SUP-2




Centralized architecture

- Gen 1 line cards supported*
- Only PHYs on the line cards
- All forwarding on the supervisor (ASIC)



Additional bandwidth

- Catalyst® 9606 backplane traces support up to 56G PAM4
- Gen 1 line cards now support up to 2.4T per slot
 - 24x 100G QSFP on LC-24C
 - 48x 50G SFP on LC-48YL



No MACsec support

- Q200 does not have an onboard crypto engine
 - Gen 2 line cards use newer CDR5M PHY for MACsec
- UADP has an onboard MACsec engine
 - Gen 1 line cards use older CDR4 PHY (not MACsec capable)

* C9600-LC-48S is not supported with SUP-2.

C9600-LC-24C

Gen 1 modules with SUP-2

2.4
Tbps

2.4 Tbps when working with Gen 2 supervisor

- 24x QSFP ports, 40/100G



12x Inphi Vega CDR4 PHYs

- IEEE 1588 and PTP v2
- HMUX
- No support for MACsec

FPGA

- Gettysburg: I/O FPGA (to supervisor)
- LC-PM: IDPROM and power control

Various breakout and QSA support*



C9600-LC-48YL

Gen 1 modules with SUP-2

2.4
Tbps

2.4 Tbps (capable) when working with Gen 2 supervisor



- 48x SFP ports, 10/25/50*G
- No support for 1G

12x Inphi Vega CDR4 PHYs

- IEEE 1588 and PTP v2
- HMUX
- No support for MACsec

FPGA

- Gettysburg: I/O FPGA (to supervisor)
- LC-PM: IDPRM and power control



* Roadmap

C9600-LC-48TX

Gen 1 modules with SUP-2

480 Gbps when working with Gen 2 supervisor



- 48x RJ-45 ports, 10G (Multigigabit)
- No support for 1/2/5G

12x Inphi Vega CDR4 PHYs

- IEEE 1588 and PTP v2
- HMUX
- No support for MACsec

FPGA

- Gettysburg: I/O FPGA (to supervisor)
- LC-PM: IDPROM and power control



Catalyst 9600 Line Card Support Matrix



	C9606 Chassis	
	C9600-SUP-1	C9600X-SUP-2
C9600-LC-24C	24x 40G or 12x 100G ✓	24x 40G/ 100G (No MACsec) ✓
C9600-LC-48YL	48x 1/10G and 25G ✓	48x 10/25G and 50G (No MACsec, no native 1G) ✓
C9600-LC-48TX	48x 1/2.5/5G and 10G (Multigigabit) ✓	48x 10G (No MACsec, no 1/2.5/5G) ✓
C9600-LC-48S	48x 1G SFP ✓	✗
C9600-LC-40YL4CD	40x 1/10G and 25G + 2x 40G and 100G ✓	40x 10/25G and 50G + 2x 40/100G and 200G* + 2x 40/100/ 200G* and 400G MACsec and WAN MACsec (no native 1G) ✓
C9600X-LC-32CD	✗	30x 40/100G + 2x 40/100/ 200G* and 400G MACsec and WAN MACsec (no native 1G) ✓
C9600X-LC-56YL4C	✗	56x 10/25/50G + 4x 40/100G MACsec and WAN MACsec (no native 1G) IPSec* ✓



SUP2 hardware does not support 1GE or below speeds natively.

Can support 1GE speed up to 24 ports per system (standalone or SVL) with SFP-1G-SX/LH

If 1GE downlinks are required, use SUP1

17.7.1

17.9.1

17.13.1

* Hardware capable

Introducing 50GE - first time in Campus

Provides seamless migration path from 10/25GE

Designation	Speed
D	400GE
C	100GE
L	50GE
Y	25GE
X	10GE

Supported on



C9600X-SUP2 & LC-48YL*



C9600X-SUP2 & LC-40YL4CD
LC-56YL4C



C9500X-60L4D

IOS-XE 17.12.1

Compatible 50G SFPs

SFP-50G-CUxM(1/1.5/2/2.5/3/4/5M)

Only on C9500X-60L4D (Roadmap for 9600X)

SFP56-25/50G-SL

Only on C9500X-60L4D (Roadmap for 9600X)

SFP-50G-SR-S

Only on C9500X-60L4D (Roadmap for 9600X)

SFP-50G-LR-S

C9500X60L4D and C9600 Sup 2 (with 48YL/40YL4CD)

Use Case

- Help decrease the oversubscription ratio between the distribution and the 400G core links
- Increase SVL bandwidth

Reduced Capex

Reuse of existing cable infrastructure

Single Lane Optics

Same port densities as 10/25G offerings

Dual Rate Optics

Staggered/Gradual upgrades and future proofing

Reduced Opex

Optimized power and cooling requirements

*No MACsec when using LC-48YL with Supervisor 2

Catalyst 9600 Series

Reversible fan tray



Fan tray is fully hot-swappable within 120 seconds (2 minutes)

* Optional side-to-back air baffle

- N+1 (8+1) fan redundancy
- Flexible service: Fan tray can be replaced from the front or back
- Efficient: Variable speed per fan depends on load, temperature, and altitude (=> lower noise)
- Airflow is side to side*


C9600 - 3KW Power Supply




	Min	Normal	Max
Vin for 1500W	90 Vac	115 Vac	140Vac
Vin for 3000W	180 Vac	230Vac	264Vac
Max current @ 110Vac	-	16A	17.6A
Max current @ 220Vac	-	16A	17.6A
MTBF (hours)	300,000		

Highest Power Efficiency – Titanium rated Power Supply


Higher Power Budget and Resiliency with 3KW PSU



Fully loaded Gen 1 LC



Mix of Gen1& Gen2 LCs



Fully Loaded Gen2 LCs

220 V	2K	2K	2K		2K	2K	2K		2K	2K	2K	
110 V	2K	2K	2K	2K	2K	2K	2K	2K	2K	2K	2K	2K
110 V	3K	3K	3K		3K	3K	3K	3K		3K	3K	3K



In Use



Standby

Catalyst 9600 Series

Power supplies



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

3kW AC



- Supports both 110V and 220V input
- 3 kW output with 220V (1500W with 110V)
- Titanium-rated power supply
- Redundant mode: Combined and N+1


NEW

2kW AC



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

2kW DC



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

Cisco Catalyst 9600 Series

Power budget


		Sup 1 Reserved power (Watt)	Sup 2 Reserved power (Watt)	Typical power (Watt)	System always budget the power with two supervisors (by default)
Required	 Fan tray	450	450	308	
	 Supervisor Slot3 Slot4	775 775	660 660	191 191	
	 400G/100G/25G line card	420	420	336	
Optional	 100G/40G line card	200	280	159	
	 25G/10G/1G line card	230	300	182	
	 mGig line card	315	315	320	
	 1G line card	160	NA	94	

Reserved power: Software allocated at bootup
Typical power: ATIS-TEER at 100% traffic rate

Architecture


Architecture

Centralized 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**

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#\$

Supervisor engine 2 Block Diagram



Silicon One Q200 – Sup2 IFG Map (CLI)


```
C9600X-SUP2# sho platf soft fed active if map | inc Int|1/0
Interface          IF_ID  Inst Asic Core IFG ID Port SubPort Mac First Serdes Last Serdes Cntx LPN GPN Type
Active
HundredGigE1/0/1   0x408  0 0 0 0 0 0 1 0 1 0 1 1 NIF Y
HundredGigE1/0/2   0x409  0 0 0 0 0 0 2 2 3 0 2 2 NIF Y
HundredGigE1/0/3   0x40a  0 0 0 0 0 0 3 4 5 0 3 3 NIF Y
HundredGigE1/0/4   0x40b  0 0 0 0 0 0 4 6 7 0 4 4 NIF Y
HundredGigE1/0/5   0x40c  0 0 0 1 0 0 5 8 8 0 5 5 NIF Y
HundredGigE1/0/6   0x40d  0 0 0 1 0 0 6 10 11 0 6 6 NIF Y
HundredGigE1/0/7   0x40e  0 0 0 1 0 0 7 12 13 0 7 7 NIF Y
HundredGigE1/0/8   0x40f  0 0 0 1 0 0 8 14 15 0 8 8 NIF Y
HundredGigE1/0/9   0x410  0 0 1 0 0 0 9 16 17 0 9 9 NIF Y
HundredGigE1/0/10  0x411  0 0 1 0 0 0 10 18 19 0 10 10 NIF Y
HundredGigE1/0/11  0x412  0 0 1 0 0 0 11 20 21 0 11 11 NIF Y
HundredGigE1/0/12  0x413  0 0 1 0 0 0 12 22 23 0 12 12 NIF Y
HundredGigE1/0/13  0x414  0 0 3 1 0 0 13 2 3 0 13 13 NIF Y
HundredGigE1/0/14  0x415  0 0 3 1 0 0 14 0 1 0 14 14 NIF Y
HundredGigE1/0/15  0x416  0 0 3 1 0 0 15 6 7 0 15 15 NIF Y
HundredGigE1/0/16  0x417  0 0 3 1 0 0 16 4 5 0 16 16 NIF Y
HundredGigE1/0/17  0x418  0 0 4 1 0 0 17 10 11 0 17 17 NIF Y
HundredGigE1/0/18  0x419  0 0 4 1 0 0 18 8 9 0 18 18 NIF Y
HundredGigE1/0/19  0x41a  0 0 4 1 0 0 19 14 15 0 19 19 NIF Y
HundredGigE1/0/20  0x41b  0 0 4 1 0 0 20 12 13 0 20 20 NIF Y
HundredGigE1/0/21  0x41c  0 0 5 0 0 0 21 6 7 0 21 21 NIF Y
HundredGigE1/0/22  0x41d  0 0 5 0 0 0 22 4 5 0 22 22 NIF Y
HundredGigE1/0/23  0x41e  0 0 5 0 0 0 23 2 3 0 23 23 NIF Y
HundredGigE1/0/24  0x41f  0 0 5 0 0 0 24 0 1 0 24 24 NIF Y
C9600-R4R1-top#
```

IFG 3,4,7,8 are IFG16
Other IFGs are IFG24

Single ASIC & Instance = 0

Core # = Slice #

IFG_ID = 1/0 (2 per Slice)




C9600-R4R1-top#sho mod 1 | inc LC
1 24 24-Port 40GE / 100GE
C9600-R4R1-top#

C9600-LC-24C CAT2313L2W2

Silicon One Q200 Slice/IFG Map

IFG 3,4,7,8 are IFG16
Other IFGs are IFG24


Faceplate Left Edge




	48-Port Line-Cards							
	Ports 0-7	Ports 8-15	Ports 16-23	Ports 24-31	Ports 32-39	Ports 40-47	Ports 48-55	Ports 56-63
Slot 1	IFG0 S0	IFG1 S0	IFG2 S1	IFG7 S3	IFG9 S4	IFG10 S5	IFG5 S2	IFG6 S3
Slot 2	IFG0 S0	IFG1 S0	IFG3 S1	IFG8 S4	IFG9 S4	IFG11 S5	IFG5 S2	IFG7 S3
Slot 3	N/A							
Slot 4	N/A							
Slot 5	IFG0 S0	IFG2 S1	IFG3 S1	IFG8 S4	IFG10 S5	IFG11 S5	IFG5 S2	IFG6 S3
Slot 6	IFG1 S0	IFG2 S1	IFG4 S2	IFG9 S4	IFG10 S5	IFG11 S5	IFG4 S2	IFG6 S3
	64-Port Line-Cards							

CISCO CONFIDENTIAL (NDA REQUIRED)


C9600-LC-24C line card block diagram




C9600-LC-48YL line card block diagram




C9600-LC-48TX card block diagram




C9600-LC-48S line card block diagram




C9600-LC-40YL4CD line card block diagram



C9600X-LC-32CD Line Card block diagram




C9600X-LC-56YL4C line card block diagram



User Centric Platform Design

Cisco Catalyst 9600 Series

Blue beacon



System/
Fan tray (front
and back)



Supervisor



Line cards



Cisco Catalyst 9600 Series

Blue beacon

C9600-Bottom#hw-module beacon ?

RP Route-Processor
fan-tray Fan Tray
slot Slot information

C9600-Bottom#hw-module beacon rp active ?

off Turn off
on Turn on

C9600-Bottom#hw-module beacon rp active on

C9600-Bottom#

*Dec 5 22:01:03.156: %CMAN_LED-6-
BEACON_LED_TURNED: R0/O: cmand: Slot 3 Beacon
LED turned ON

C9600-Bottom#show beacon all
Switch# Beacon Status

C9600-Bottom#

C9600-Bottom#conf t

C9600-Bottom(config)#service internal

C9600-Bottom(config)#end

C9600-Bottom#hw-module beacon rp active on

C9600-Bottom#hw-module beacon rp active status

BLUE

C9600-Bottom#

**System/
Fan tray**
(front and back)



**Line
cards**



Supervisor



- Blue beacon LED allows easy identification of the switch and components
- Color/state – solid blue
- Blue beacon can be turned on or off using exec-mode CLI

*Blue Beacon DDTs CSCvo57488

Cisco Catalyst 9600 Series

Supervisor LEDs

1 Status LED

- Green** All diagnostic tests passed
- Amber** Major environmental warning
- Red** Major fault

2 Beacon LED

- Solid Blue** The operator has indicated that the system needs attention

3 System LED

- Green** No environmental monitors issue
- Amber** PS or fan failure (minor)
- Red** Major fault

4 Active LED

- Green** Operational and active
- Amber** ROMMON prompt/Standby mode
- Blinking amber** Major fault

5 Reset

Reset button

Reset button: complete reboot of the system



Cisco Catalyst 9600 Series

Line card, PSU, and Fan LEDs




Fan status LED

- Green** Fans operating normally
- Amber** One fan failed
- Red** More than one fan failed

Line card status LED

- Green** Module is up and online
- Amber** Minor environmental alarms
- Red** Major environmental alarms



PSU LED


- Green** Active and normal operation
- Blink green** Standby and normal operation
- Amber** PS failure



New Catalyst 9600X Gen 2 line card ejectors

Based on feedback and usability studies, the original Catalyst® 9600 Gen 1 line card ejectors were **redesigned for Gen 2**:

- **Enhanced locking mechanism:** Prevents inadvertent ejection, without using thumb screws.
- **Greater extension arm depth:** Reduces ejector interference with optics and cables.
- **Extension arm swivel mechanism:** Allows ejector to swing completely free, for easier optic and cable insertion and removal.
- **Easy thumb-release mechanism:** Enables faster release or lock action.
- **RFID tag is embedded** on the front of the vertical limb of the ejector.



Cisco Catalyst 9600 Series

SSD storage

Capabilities	
Form factor	M2 SATA
Capacity	480 GB, or 960 GB
Performance	300 MB/s read, 290 MB/s write
Power write	4.5W
Security	Hardware-based AES-256-bit

Product ID	Release
C9K-F1-SSD-480G	Cisco IOS® XE 16.11.1
C9K-F1-SSD-960G	



Use cases

- General-purpose storage
- App hosting*
- Wireshark on-board
- Logging local NetFlow record storage

* Roadmap

Cisco Catalyst 9600 Series

USB 3.0 storage



USB 3.0 120G SSD



Capabilities	
Form factor	USB 3.0
Capacity	120 GB
Performance	400 MB/s read, 140 MB/s write
Power write	4.5W maximum
Security	Hardware-based AES-256-bit

Product ID	Release
SSD-120G/ SSD-120G=	Cisco IOS® XE 16.11.1

Use cases

- General-purpose storage
- Wireshark on-boarding
- Logging local NetFlow record storage

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
- Combined mode: Line card can shut down if there isn't enough power
- N+1 mode: Use N power supplies for system budgeting
- 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

SDM Templates

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)

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

SD-Access template

Maximizes system resources
for **policy** to support **fabric**
deployment

NAT template

Maximizes the **NAT**
configurations on the switch

Cisco Catalyst 9600 with Sup 1

SDM templates and scale numbers

IOSXE 17.2.x or earlier

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



Customizable resources

Cisco Catalyst 9600 with Sup 1

SDM templates and scale numbers

IOSXE 17.3.1



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

Higher scale



Customizable resources

FIB Scale Customization

* Total Multicast (L2+L3) =/ < 48K

Feature	Scale Values (Min – Max)	Default Value	Step Units
L2 MAC	32K – 128K	32K	16K
L3 Unicast	64K – 256K	64K	16K
L3 Multicast *	0 - 32K	16K	16k
L2 Multicast *	0 - 32K	16K	16K
SGT & MPLS IP-to-Tag labels	0 - 64K	32K	32K
Netflow - Input	0 - 64K	32K	32K
Netflow - Output	0 - 64K	0	32K
Total Resources	416K		

Scale Values: minimum and maximum possible values
Step Units: granularity (block size) of scale per type
Default Value: default value, if it's not specified by CLI

16K – 16384
 32K – 32768

FIB Allocation Examples

* Netflow requires 2X for flow records

Feature	Customer 1 (L2 focus)	Customer 2 (L3 focus)	Customer 3 (FNF focus) *
L2 MAC	128K	32K	32K
L3 Unicast	144K	256K	64K
L3 Multicast	32K	16K	16K
L2 Multicast	16K	32K	16K
SGT & MPLS IP-to-Tag labels	64K	64K	32K
Netflow - Input	32K	32K	64K *
Netflow - Output	0	0	64K *
Total Resources	416K		

16K – 16384

32K – 32768

Cisco Catalyst 9600 with Sup 1

SDM Templates, Scale & Customization

IOSXE 17.4.1

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


Higher scale



Customizable resources

ACL Scale Considerations

IOSXE 17.4.1



ACL Features

Feature	Scale Values (Min – Max)	Default Value	Step Size	Customizable
Ingress ACL	4K – 26K, 27K	4K	2K	Y
Egress ACL	4K – 26K, 27K	4K	2K	Y
Ingress QoS	1K, 2K - 16K	1K	2K	Y
Egress QoS	1K, 2K - 16K	1K	2K	Y
PBR/NAT	1K, 2K - 16K	1K	2K	Y
LISP	1K - 2K	1K	1K	Y
Tunnels	1K - 3K	1K	1K	Y
NFL ACL	1K – 2K	1K	1K	Y
Ingress SPAN	NA	0.5K	NA	N
Egress SPAN	NA	0.5K	NA	N
CPP	NA	0.5K	NA	N
Sec. Assoc	NA	0.5K	NA	N
Macsec SPD	NA	0.25K	NA	N



ACL Allocation Examples

* 2K entries are reserved for System

Feature *	Customer 1 (Security focus)	Customer 2 (QoS focus)	Customer 3 (NAT focus)
Security - Input	27K	6K	12K
Security - Output	18K	6K	8K
QoS - Input	1K	16K	6K
QoS - Output	1K	16K	4K
PBR/NAT	2K	3K	16K
LISP	1K	2K	2K
Tunnels	1K	2K	1K
Netflow ACL	1K	1K	1K
Total Resources *	54K		

1K - 1024
16K - 16384

Cisco Catalyst 9600 Series – Supervisor engine 2

Switch Database Management (SDM) template

Default template

Maximizes system resources
for Layer 3 unicast and
multicast **routes**

User-customizable template

Allows customizable
ACL TCAM resources



**Cisco® Catalyst®
9600 Series**

Custom template

Provide flexibility for
customizing TCAM space for
specific requirements



Silicon One Q200 SDM template – 17.7.1

Features	Customizable	DEFAULT (core + edge)	Custom (min to max @ step)	
MAC addresses	✓	128,000	32,000 to 256,000	@ 1000 step
Host routes (ARP/NDP)	✓	128,000/64,000	32,000 to 256,000	@ 1000 step
Layer 2 multicast entries (IGMP/MLD)	FCS+	16,000/8000	0 to 64,000	@ 1000 step
Layer 3 multicast routes (IPv4/IPv6)	FCS+	32,000/16,000	0 to 64,000	@ 1000 step
ACL compression (SGT, DGT, OGID/v6)	✓	32,000/16,000	0 to 64,000	@ 1000 step
MPLS labels	✓	256,000	0 to 512,000	@ 1000 step
Reserved (PBR/NAT)	FCS+	16,000/8,000	0 to 256,000	@ 1000 step
	CEM	608,000 (288,000 for LPM)		
Layer 3 unicast routes (IPv4/IPv6)	FCS+	2 million/ 1 million	1 million to 2 million	@ 1 million step
Features	Customizable	DEFAULT (core + edge)	Custom (min to max @ step)	
Security ACL (IPv4/IPv6)	FCS+	8000/4000 shared*	0 to 11,000/5000	@ 1 step
Quality of service (IPv4/IPv6)	FCS+	8000/4000 shared*	0 to 11,000/5000	@ 1 step
Policy-based routing (IPv4/IPv6)	FCS+	8000/4000 shared*	0 to 11,000/5000	@ 1 step
Lawful intercept (IPv4/IPv6)	FCS+	1000 (2x 512) reserved	1000 to 5000/2500	@ 1 tap (2 ACE)
LPTS, EPC, FSPAN, NFL (IPv4/IPv6)	FCS+	1000 (2x 512) reserved	1000	@ 1 step
	TCAM	10,000 (2000 for LPM)		

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

FIB Allocation Examples

Feature	Customer 1 (L2 focus)	Customer 2 (L3 focus)
MAC addresses	256K	32K
Host routes (ARP/NDP)	32K	256K
Layer 2 multicast entries (IGMP/MLD)	64K	16K
Layer 3 multicast routes (IPv4/IPv6)	16K	64K
ACL compression (SGT, DGT, OGID/v6)	64K	64K
MPLS labels	32K	128K
Reserved (PBR/NAT)	15K	48K
Total Resources	608K	

Layer 3 unicast routes (IPv4/IPv6)	1M/500K	2M/1M
------------------------------------	---------	-------

16K – 16384
32K – 32768

Cisco UADP 3.0 Architecture

Common Capabilities of UADP ASIC



Flex Parser
&
Programmable
Pipelines




Recirculation
Capability



Micro Engines



Adaptable Tables



UADP 3.0



~20B transistors

16-nm technology



Customizable ASIC
templates



Double-width tables



Up to 1.6 TB
bandwidth



Up to 1000Mpps
throughput



36-MB
unified buffer



3x more FIB scale




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



Up to 1.6Tbps ASIC
Interconnect


Cisco Catalyst 9600 Series – Sup1

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




UADP 3.0 – Under the covers

showing both cores




Cisco Catalyst 9600 Series

Unicast forwarding within ASIC (ingress and egress)



Cisco Catalyst 9600 Series

Unicast forwarding across ASIC (ingress and egress)




Cisco Catalyst 9600 – Supervisor 1

ASIC interconnects


ASIC interconnects highlights:

- 1600G packet bandwidth point-to-point links
- No packet size limitations
- No packet fragmentation
- No buffering on interconnects
- Provides a credit-based mechanism to slow Packet Buffer Capture (PBC) and align to the data rate
- No prioritization between the packet types



Cisco Catalyst 9600 – Supervisor 1

ASIC interconnects



Cisco Silicon One Q200 Architecture



Introducing Cisco Silicon One

One architecture, multiple devices



Q202

3.2 Tbps



Q201

6.4 Tbps



Q100

10.8 Tbps



Q200

12.8 Tbps

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

- Multiple functions: System-on-chip, line card, or fabric
- Multiple form factors: Fixed or modular
- **Multiple segments: Enterprise and service provider**

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

Cisco Silicon One Q200

Industry-leading switching and routing silicon



12.8-Tbps
bandwidth



8.1 Bpps



8-GB HBM for
deep buffers



10 million IPv4
or 5 million
IPv6 route scale



Fully P4
programmable
pipeline



50G PAM4
serializer/
deserializer (SerDes)

Cisco Silicon
One™ Q200
Industry-leading
12.8-Tbps system on chip



First 7-nm ASIC providing lowest
watts/G power consumption



Fully P4 programmable, enabling
feature velocity




Multislice architecture for flexibility
and scale

Routing capabilities with switching power and performance

Cisco Silicon One Q200


Programmable next-generation ASIC

- **Generation 1.1, 7-nm ASIC**
- **6 slices** for **≤8.1 Bpps** throughput
- **~12.8-Tbps** bandwidth at **1.35 GHz**
- **256 SerDes at 56G PAM4**
- **~80-MB internal buffer** (up to 8-GB HBM)
- **~2 million internal Longest Prefix Match (LPM) routes** (up to 4 million in HBM)
- **~608,000 shared CEM** (896,000, with 288,000 for LPM)
- **~10,000 shared TCAM entries** (96,000/6, with 16,000 for LPM)





Silicon One ASIC architecture overview






S1- Interface Groups (IFG) Overview

Interface Groups (IFGs):

- pool of 56Gbps SerDes & MACs
- 10/25/50GE or 40/100/200/400GE
- Q200 (Gibraltar) – 256 total SerDes

Per port buffer:


- admission test based on Traffic Class
- absorb/smooth temporary burst oversubscription
- flow-control generation for lossless to core
- 4Mb – partitioned between ports



Q200 – Intrepid (Stargazer) Slice/IFG Map

IFG 3,4,7,8 are IFG16
Other IFGs are IFG24

Faceplate Left Edge



	48-Port Line-Cards							
	Ports 0-7	Ports 8-15	Ports 16-23	Ports 24-31	Ports 32-39	Ports 40-47	Ports 48-55	Ports 56-63
Slot 1	IFG0 S0	IFG1 S0	IFG2 S1	IFG7 S3	IFG9 S4	IFG10 S5	IFG5 S2	IFG6 S3
Slot 2	IFG0 S0	IFG1 S0	IFG3 S1	IFG8 S4	IFG9 S4	IFG11 S5	IFG5 S2	IFG7 S3
Slot 3	N/A							
Slot 4	N/A							
Slot 5	IFG0 S0	IFG2 S1	IFG3 S1	IFG8 S4	IFG10 S5	IFG11 S5	IFG5 S2	IFG6 S3
Slot 6	IFG1 S0	IFG2 S1	IFG4 S2	IFG9 S4	IFG10 S5	IFG11 S5	IFG4 S2	IFG6 S3
	64-Port Line-Cards							

CISCO CONFIDENTIAL (NDA REQUIRED)




S1 – Buffer Queuing Overview

Onboard SMS vs. HBM

Shared Memory Subsystem (SMS)

- Fully shared across slices:
- 14 x 128B write + 14 x 128B read buses
- **108MB** managed as **288K** x 384B buffers
- multicast buffering – write once, read many

~12.8T bps = ~8.1B packets-per-second (6 x 1.35GHz)



High Bandwidth Memory (HBM)

- fully shared across slices
- 2 x 128B write + 2 x 128B read buses
- ~2Tbps full-duplex (in and out)
- 8GB managed as 1M x 8KB buffers
- DRAM packets – combine up to 32 small packets




S1 Queuing – Ingress Queues

IFG RXQ – Interface-Group Receive-Queue

- Buffer data (FIFO) received from the various Ethernet ports, before injection into the core

RXRQ – Receive Replication-Queues

- Store PDs of multicast packets that require replication, before VOQs



PDSQ – PD Switch-Queues

- queue per crosspoint (mesh) of the switch

VOQs – Virtual Output Queues

- Pool of 64K/slice of Queues
- User assigned to a OQ or TXRQ
- Seamlessly extendable in HBM



Platform Comparison

Catalyst 9600X and 9600 – physical

	 Catalyst® 9600X (Silicon One™ Q200)	 Catalyst 9600 (UADP 3)
	SUP-2	SUP-1
CPU (BDW-NS, BDW)	2.7 GHz 8C Intel® (BDW-NS)	2.4 GHz 8C Intel (BDW)
DRAM (type)	32 GB (DDR4)	16 GB (DDR4)
ASIC (number)	Q200 (1x)	UADP3 (3x)
Capacity (chassis)	12.8 Tbps (full duplex)	4.8 Tbps (full duplex)
Capacity (per slot)	3.2 Tbps (full duplex)	1.2 Tbps (full duplex)
10G max	256 (40x4 + 24x4^*)	192 (48x4)
25G max	256 (40x4 + 24x4^*)	192 (48x4)
50G max	256 (40x4 + 24x4^*)	—
40G max	96 (SUP-2 can support 128)	96 (24x4)
100G max	96 (SUP-2 can support 128)	48 (12x4)
400G max	8 (4x2)	—

* Roadmap (not committed). ^ Breakout cable ports.

Catalyst 9600X and 9600 – features and scales

	 Catalyst® 9600X (Silicon One™ Q200)		 Catalyst 9600 (UADP 3)	
	Default	Maximum (Custom)	Default	Maximum (custom)
MAC addresses	128,000	256,000	80,000	128,000
IP host routes	128,000	256,000	80,000	128,000
Multicast L2 groups	16,000	64,000*	16,000	48,000
Multicast L3 routes	32,000	64,000*	32,000	48,000
IP LPM routes	2 millions	2 millions	212,000	256,000
MPLS labels	256,000	512,000	32,000	64,000
SGT/OG labels	32,000	64,000	32,000	64,000
NAT* sessions	16,000*	128,000*	3000	16,000
Security ACL entries	8000	10,000*	12,000	27,000
QoS ACL entries	8000	10,000*	8000	21,000
PBR* ACL entries	8000*	10,000*	3000	16,000

* Roadmap for Catalyst 9600X

C9600 Sup1 and C6K Sup2T/6T Scale

Features	C9600– Sup1	C6K-Sup720 3C	C6K-Sup2T	C6K-Sup6T
Switching capacity	9.6T	720G	2T	6T
Forwarding rate (IPv4/IPv6)	3 Bpps 3 Bpps	450Mpps 225Mpps	720Mpps 390Mpps	780Mpps 390Mpps
MAC addresses	82,000*	96,000	128,000	128,000
LPM/host routes (IPv4/IPv6)	212,000* 212,000*	256,000 128,000	256,000 128,000	256,000 128,000
Multicast routes	32,000*	32,000	128,000	128,000
Security ACLs	27,000*	Dedicated 32,000	Shared with QoS 64,000	Shared with QoS 64,000
QOS ACLs	16,000*	Dedicated 32,000	Shared with Security 64,000	Shared with Security 64,000
Flexible NetFlow (per ASIC)	96,000*	128,000	512,000	512,000
VLAN ID	4,000	4,000	4,000	4,000
Spanning Tree instances	4,000*	4,000	4,000	4,000

* Depends on SDM template


C9600 Sup2 and C6K Sup2T/6T XL Scale

Model	C9600– Sup2	C6K-Sup720 3C	C6K-Sup2T	C6K-Sup6T
Switching capacity	12.8T	720G	2T	6T
Forwarding rate (IPv4/IPv6)	8 Bpps 8 Bpps	450Mpps 225Mpps	720Mpps 390Mpps	780Mpps 390Mpps
MAC addresses	256,000*	96,000	128,000	128,000
LPM/host routes (IPv4/IPv6)	2,000,000* 1,000,000*	1,000,000 500,000	1,000,000 500,000	1,000,000 500,000
Multicast routes	64,000*	32,000	128,000	128,000
Security ACLs	Shared with QoS & PBR 10,000*	Dedicated 32,000	Shared with QoS 256,000	Shared with QoS 256,000
QOS ACLs	Shared with Security & PBR 10,000*	Dedicated 32,000	Shared with Security 256,000	Shared with Security 256,000
Flexible NetFlow	2,000,000 (Sampled)	256,000/ASIC	1,000,000/ASIC	1,000,000/ASIC
VLAN ID	4,000	4,000	4,000	4,000
Spanning Tree instances	4,000	4,000	4,000	4,000

* Depends on SDM customization

C9600 App Hosting

Enhanced app-hosting infrastructure on Catalyst 9600



Competitive Info

Catalyst 9600 Series for continued market leadership

Competitive challenges

ARISTA



Campus:
7300X3

Data center:
7500R3

aruba

a Hewlett Packard
Enterprise company



CX 8400

Threats

- 7300X3 in campus with higher 40/100G density
- Positioning data center switch: 7500R3 in campus core with higher scale and speeds/feeds

Threats

- Higher port densities
- Higher route, MAC scales

Leadership with Catalyst® 9600/9600X

	Cisco Catalyst 9600X	Arista 7500R3	Cisco Catalyst 9600	Arista 7300X3	Aruba CX8400
ASIC	Q200	Jericho 2	UADP	Trident 3	Jerich o
Scale, performance, hardware differentiation					
ACL/route scale	Best	Best	Partial	Partial	Partial
MAC scale	Best	Partial	Partial	Partial	Partial
Port buffers	Best	Best	Partial	Partial	Partial
Uplinks on LC/SUP: 400G	Best	Best	Bad	Bad	Bad
MACsec, WAN MACsec	Best	Bad	Partial	Bad	Partial
Enterprise-grade campus differentiation					
Line-rate encryption MACsec, WAN MACsec, NetFlow	Best	Bad	Partial	Bad	Bad
Maximum Investment protection Common chassis, fan tray, and power supply, all Gen 1 line cards supported, 3x performance with less than 25% investment	Best	Bad	Partial	Bad	Bad
400G leadership in campus Support for 10G to 400G speeds	Best	Best	Bad	Bad	Bad

Legend

Best

Partial

Bad

Lead with investment protection: Gen 1 backward compatibility.

Highlight 400G leadership in campus.

Win with Catalyst 9600 differentiation against the competition.

Catalyst 9600 Series for continued market leadership

Competitive challenges

Arista



Campus:
7300X3

Data center:
7500R3

Threats

- 7300X3 in campus with higher 40/100G density
- Positioning data center switch: 7500R3 in campus core with higher scale and speeds/feeds

Aruba



CX 8400

Threats

- Higher port densities
- Higher route, MAC scales

Leadership with Catalyst® 9600/9600X

	Cisco Catalyst 9600X	Arista 7500R3	Cisco Catalyst 9600	Arista 7300X3	Aruba CX8400
ASIC	Q200	Jericho 2	UADP	Trident 3	Jerich o
Scale, performance, hardware differentiation					
ACL/route scale	Best	Best	Partial	Partial	Partial
MAC scale	Best	Partial	Partial	Partial	Partial
Port buffers	Best	Best	Partial	Partial	Partial
Uplinks on LC/SUP: 400G	Best	Best	Bad	Bad	Bad
MACsec, WAN MACsec	Best	Bad	Partial	Bad	Partial
Enterprise-grade campus differentiation					
Line-rate encryption MACsec, WAN MACsec, NetFlow	Best	Bad	Partial	Bad	Bad
Maximum Investment protection Common chassis, fan tray, and power supply, all Gen 1 line cards supported, 3x performance with less than 25% investment	Best	Bad	Partial	Bad	Bad
400G leadership in campus Support for 10G to 400G speeds	Best	Best	Bad	Bad	Bad

Legend

Best

Partial

Bad

Lead with investment protection: Gen 1 backward compatibility.


Highlight 400G leadership in campus.

Win with Catalyst 9600 differentiation against the competition.

High Availability

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 for Sup 2

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

Cisco Catalyst 9000 Series ISSU workflow



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

Software Maintenance Updates/ Patching

SMU is an emergency point fix positioned for expedited delivery to a customer in case of a network down or revenue affecting scenario.

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

Hot Patching: Install of a SMU does not require a reload.




Graceful Insertion and Removal

Wish to Isolate a node with no or minimal impact

Hardware
Replacement

Software
Upgrades

Configuration
Changes



Graceful Insertion and Removal



Upgrades with no or Minimal Traffic Loss



Comprehensive Node Isolation Framework



Easy Execution with a single command




Highly Customizable workflow

**Simple
Customizable
Non-Traffic Impacting**


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




Sup 1 StackWise Virtual is supported with IOSXE 16.12.1 or later
Sup 2 StackWise Virtual is in roadmap

- 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
- SVL and DAD are supported on any ports with 10G or high speed, including QSA.

StackWise Virtual - Quad Sup RPR – 9600




Role	Description	Control Plane	Data Plane
StackWise-A	StackWise Virtual Active In Chassis Active	Active	Active
StackWise-S	StackWise Virtual Standby In Chassis Active	Hot Standby	Active
ICS	In-Chassis Warm Standby	Warm Standby	Warm Standby

Sup 2 StackWise Virtual Quad Sup is in roadmap

StackWise Virtual Enhancements for C9500X/C9600X


IOS-XE 17.10.1

Adding/Deleting Additional SVL or DAD Link – Requires a Reload of the entire SVL Pair



Adding/Deleting Additional SVL or DAD Link – Reload not required

IOS-XE 17.10.1 **NEW**



Security

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

Segmentation Agility with Security


Secure onboarding of users and devices with flexible authentication and segmentation



SD-Access segmentation simplified through policy analytics




End-to-End Security for Campus




WAN MACsec overview and use cases

Enabled in hardware on Catalyst 9000 Switches

MACsec




- Hop-by-hop encryption
- Directly connected Layer 2 links **only**
- Secure any direct link in campus/ WAN, including DWDM
- 256/128-bit AES encryption



Supported on **all Catalyst® 9000** models
All ports and speeds simultaneously

WAN MACsec



- End-to-end encryption across Layer 2 Ethernet WAN service
- EoMPLS/VPLS, Q-in-Q
- Point-to-point or point-to-multipoint
- Applications include secure site interconnect, DCI, storage replication
- 256/128-bit AES encryption



Catalyst 9600X



Catalyst 9500X

Supported on **Silicon One™** platforms*
All ports/speeds (up to 400G) simultaneously



Catalyst 9400X




Catalyst 9300X

Roadmap support access **Catalyst 9300X and 9400X**
Aggregate speed up to 100G

* Requires Release 17.8.1 and HSEC key (\$1 PID).


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, 2.5G, 5G, 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

WAN MACsec – L2 Encryption



L2 MPLS / VPLS / EoMPLS / QinQ Service Provider




Quality of service

UADP 3.0 QoS


Cisco Catalyst 9600 Series – quality of service

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



Cisco Catalyst 9600 Series

QoS forwarding (ingress and egress)






Cisco Catalyst 9600 Series

Hardware queues


Default (2q3t)



Configurable (8Q3T/2P6Q3T)



Cisco Catalyst 9600 Series – unified packet buffer complex





- 36-MB unified packet buffer 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

Silicon One QoS

VoQ Architecture and Head of Line Blocking


Virtual Output Queueing – Egress Optimized



Cisco 9600X – Why VoQ QoS?



No Head-of-Line Blocking




Many Flows to a Single Uplink

- Common on (expensive) WAN/Edge uplinks
- Even if bandwidth available, buffers can fill up

Large Flow to Multiple Downlinks

- Common for Multicast & Broadcast traffic
- One slow receiver can penalize other ports

Local vs. HBM Buffers




Low-Latency Local Shared Memory Buffers

- Voice & Video are very latency-sensitive
- Multiple levels of Strict Priority Queuing


Deep High-Bandwidth Memory Buffers

- Guarantee delivery of session-oriented flows
- Reserve buffers to absorb occasional bursts


S1 Q200 – QoS overview



QoS Toolset



VoQ Mapping



Access control lists

Cisco Catalyst 9600 Series access control lists


Four forms of security ACLs

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

Router ACL (RACL)	VLAN ACL (VACL)	Port ACL (PACL)	Security group ACL (SG ACL)
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 SGTs that are assigned
Applied as an input or output policy to a Layer 3 interface , Switch Virtual Interface (SVI) , or Layer 3 EtherChannel interface	Applied as a policy to a VLAN; is inherently applied to both inbound and outbound traffic	Applied as a policy to a Layer 2 switch port interface or EtherChannel interface ; is applied to inbound traffic only	Applied as a policy to a Layer 2 switch port interface; is applied to inbound traffic only
Standard/extended ACLs	Standard/extended ACLs	Standard/extended/MAC ACLs	Standard/extended/MAC ACLs

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)

1

```
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
```



No hardware resources
Drop packet*

2

Full hardware support

- Router ACLs
- VLAN ACLs
- Port-based ACLs
- Role-based ACLs

3


Hardware-assisted ACL features

- NetFlow
- NAT and PAT
- PBR
- WCCP
- Cisco TrustSec®


4

* Packets must be forwarded by software for any reason (for example, not enough hardware resources).
The switch forwards the packets only after applying ACLs on the packets.

Cisco Catalyst 9600 Series security TCAM lookup



Access control list terminology



What is an L4OP?

GT
(greater than)

LT
(less than)

NE
(not equal to)

RANGE
(from - to)


What is **NOT** an L4OP?

EQ
(equal to)

Cisco Catalyst 9600 Series – Hitless TCAM update

Starting
16.8.1

- **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 does consume any additional TCAM resources
- If not enough space in TCAM, falls back to old ACL method




ACL updates




Cisco Catalyst 9600 Series– Hitless TCAM update

Example 1



Example 2






S1 Q200 - Central Database (CDB)

Onboard LPM, CEM & ACL memory


- **CDB** includes the **Central L2/L3 Forwarding** and **ACL** databases:
 - **LPM** – IP/mask routing database implemented by Longest Prefix Match algorithm (extensible to HBM, for scale)
 - **CEM** – MAC & IP Host (/48 or /32 - /128), Multicast and Label database implemented by an Exact Match algorithm
 - **ACL** – Access Control List database, contains service Access Control Entries in TCAM
- **LPM & CEM** are accessed by **RX Forwarding** stage via 1 of 3 interfaces:
 - **LPM** direct interface, or **CEM** direct interface
 - **Splitter** – if accessing both, splitter merges response
- **ACL** is accessed by **RX Termination**, **RX Forwarding** and **TX Encapsulation** stages





S1 Q200 – ACL Run-Time Flexibility (RTF)


- RTF allows **Security & QOS** (and others), **IPv4 & IPv6**, to share ACL key-fields within the same TCAM space
- Support for RTF for **Ingress ACL only** (API for egress, as well)
- **Multiple ACL Stages** - **input-port** (outer MAC\IP), **pre-forwarding** (on forwarding header), **pre-resolution** and **rx-end**





Catalyst 9600X – ACL TCAM (CDB) Architecture


1 TCAM Ring per Slice



- Central TCAM is 8 TCAM rings
 - **24 TCAM banks** per ring
 - $8 \times 24 \times 512$ entries = 96K
- **5 inputs** per TCAM lookup process
 - $1 \text{ TRM} + 2 \text{ FWD}^{(RX)} + 2 \text{ ENC}^{(TX)}$
- **FWD^(RX) & ENC^(TX)** given to **6 slices**
 - Single ring per slice ⁽¹⁻⁶⁾
 - each bank accepts 1 input
 - $6 \times 24 \times 512 = 72K$
- **TRM** uses the remaining **2 rings** ⁽⁷⁻⁸⁾
- **LPM** takes 2-4 TCAM banks (per ring)
 - 20-22 banks remaining
- 512 x 160b per TCAM bank
 - $512 \times 20 = \sim 10K$ (L2/IPv4)
- 512 x 320b per TCAM bank-pair
 - $512 \times 20/2 = \sim 5K$ (IPv6)



S1 Q200 - TCAM NPL Access




TCAM Pool	Classification Feature(s)
TERM	L2-LPTS
FW0	IPv4 L3-LPTS (35) IPv4 ACL IPv4/6 OGACL IPv4/6 SGACL
FW1	IPv4 QoS
FW0+1	IPv6 L3-LPTS (29) IPv6 ACL IPv6 QoS
TX0	IPv4 ACL
TX0+1	IPv6 ACL

- Key-Size 160b or 320b
- Forward (0+1) stage is large (18 banks)
- Transmit (0+1) stage is small (4 banks)
- Run To Completion – No hardware merge
- Multiple TCAM access impacts PPS
- Limited key-size requires per-feature access
- SGACL is ingress feature

Why OGACL/SGACL in Campus Core?

Object-Group ACLs for IP




Object-Groups map IP/mask to Labels in CEM

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

OGACL ACEs take minimal space in ACL TCAM

- Only the Permit/Deny ACEs in TCAM
- OGACLs with same ACEs can reuse entries

Scalable-Groups for SDA



Scalable-Groups map IP/mask to Labels in CEM

- ISE/DNAC defines IP/masks to simple SG name
- SGT labels are stored in Exact Match table


SGACL ACEs take minimal space in ACL TCAM

- Only the Permit/Deny ACEs in TCAM
- SGACLs with same ACEs can reuse entries

Multicast

Cisco Catalyst 9600 Series


Multicast forwarding within ASIC (ingress and egress)





S1 Q200 - Standalone SOC Architecture

Generic Multicast Packet Walk



Cisco Catalyst 9600 Series – IPv4 and IPv6 multicast

IPv4 multicast

- Multicast Source Discovery Protocol (MSDP)
- IPv4 Auto-RP, BSR RP
- PIM-SM, SSM, and Bidir-PIM
- IPv4 multicast HA
- Multicast VRF-Lite
- Multicast load splitting


Core

- Dual-stack IPv4/ Pv6
- PIM-SM, SSM, and Bidir-PIM
- IGMP v2,v3 and PIM snooping
- IGMP throttling
- Stub IP multicast routing
- PIM BFD
- IPv4 multicast HA

Distribution

- IGMP v1, v2, v3 and PIM snooping
- IPv4 multicast QoS and ACL
- IGMP v1, v2 filtering

Access



IPv6 multicast

- IPv6 BSR RP
- PIM-SM, SSM, and Bidir-PIM
- IPv6 embedded RP
- IPv6 multicast HA

Core

- Dual-stack IPv4/IPv6
- PIM-SM, SSM, and Bidir-PIM
- Multicast Listener Discovery (MLD) v1, v2 and PIM snooping
- Hardware register and RPF
- HSRP-aware PIM
- IPv6 multicast HA

Distribution

- MLD v1, v2 and PIM snooping
- IPv6 multicast QoS and ACL

Access

Catalyst 9000 Design Consideration

C9600 - Place In Network (PIN)


BGP

IGP

STP

1


Layer 3 Core



- Base L3 Routing
- High-BW, Port Density
- Simple ACL
- Simple QoS

2


Layer 3 Core + Edge



- L3 Core + Edge Services
- DCI, WAN, Internet
- Edge Security, VPN & OGACL
- Complex H-QoS

3


Fabric (SDA & EVPN)



- L3 Core + VXLAN Fabric
- Border, Edge, CP/RR
- L2 & L3 VNI & SGACL
- App-based QoS

4

Collapsed Core



- L3 Core + Distribution
- L2 Services to Access
- First Hop Security
- Access QoS & AVC



Edge +
Core



Catalyst 9600X
with SUP-2



Catalyst 9500X
models



Catalyst 9600
SUP-1




Catalyst 9500
models


Core +
Distribution



Example of a Traditional Three-Tiers Campus Design



Scale of Traditional Three-Tiers Campus Design



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

Example of a Four-Tiers - Campus + Edge Design



Edge

Catalyst® 9606R Switch with C9600X-Sup2 for connection to WAN Edge and between core devices.

High-Scale Routing:

Support for full IPv4 and IPv6 Internet + LAN + VRF routing

Flexible Transport Options :

Support for MPLS-VPN, SD-Access and BGP-EVPN with 4K VRFs

Low-Latency Buffers:

Support for local optimized low-latency shared memory

High-Bandwidth Memory:

Support large buffers for micro-burst & congestion

Virtual Output Queuing :

eliminate head-of-line blocking, with dedicated output queues for every egress port

Intent-based networking

Cisco Catalyst 9600

Primed for resilient, high-scale SD-Access deployments

Built for
SD-Access
border and
control plane



- Fabric border and control plane
- Fabric edge*
- Wireless controller*

Performance

- Higher bandwidth and port density - built for large and medium campuses
- Better scale with faster control plane – suitable for multi-site deployments

Resiliency

- Faster convergence at SD-Access border
- Higher border resiliency with SSO
- HA software upgrades with ISSU

Orchestration

- Simple, intuitive multi-site workflows with Cisco® DNA Center
- Extend services like AVC*, ETA*, app hosting



Platform Benefits

Extending Catalyst® 6500/6800 performance

- Faster CPU and larger ACL table
- Higher bandwidth (100G/40G/25G) and port density

Open, extensible Cisco IOS® XE


- Enhanced patching and ISSU
- Customizable ASIC templates
- Model-driven programmability

Campus fabric with SD-Access

- VRF and SGT scale
- Advanced assurance
- eWLC support*

* Roadmap

Software Defined – Open, Programmable, Flexible Scaled operations




Catalyst 9600 – Simplify network operations open-standards APIs



Day 0


Device onboarding



Automatic device bring-up using industry-standard interface

Day 1


Model-driven configuration mgmt



Open-source tool support for configuration and provisioning

Day N

Monitoring, analytics, SW management




Model-driven telemetry for monitoring and analytics; Docker support for third-party app hosting, On-Box python

Plug in to operational in minutes


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

Using Python with Cisco IOS XE devices



Guestshell

- Fault and resource isolation
- Secure Linux shell environment
- Python interpreter
- Integrated with ZTP and EEM
- Disabled by default*
 - `switch# guestshell enable`



Note: Additional config needed from 16.8.1 app-hosting appid guestshell vnic management guest-interface 0


Application hosting on the catalyst 9000

new strategic capabilities for Cisco devices

IT operations and monitoring tools



Consolidate physical infrastructure



IT operations and monitoring tools




Enhance visibility and security enforcement



Cloud gateways with server-less edge compute



Reduce app latency and optimize app traffic




Customer Specific Applications



Derive new insights and respond



Application hosting in the enterprise and IoT



Closing



Catalyst 9600X – things to know

Next-generation core + edge switching with Silicon One Q200

C9600X-SUP-2 + LC-40YL4CD



Technology	Brief description	Differences from UADP 3.0
✓ Large LPM table	- Up to 2 million IPv4 or 1 million IPv6 (hash efficiency is about 80%). - Dedicated memory for LPM.	- Up to 256,000 IPv4/IPv6 with custom SDM template. - LPM and other features share 416K ASIC memory.
✓ Large MAC table	- Up to 256,000 MAC entries, custom SDM template. - Shared with other features.	- Up to 128,000 MAC entries with custom SDM template. - MAC shared with LPM, etc. in same 416K ASIC memory.
✓ VOQ QoS + HBM	- Q200 has 80 MB local (low latency) + 8 GB HBM (high-bandwidth memory) buffer memory. - Q200 uses a virtual output queue (VOQ) architecture. All queuing and policing policies applied on ingress.	- Max 36 MB unified buffer memory per ASIC. - Supports both ingress and egress queuing and policing.
✓ OGACL and SGACL	8000 IPv4, 4000 IPv6 ACL TCAM entries. - Object group and security group ACLs use CEM to map IP to group label, TCAM uses only Layer 4 ACEs. (OG/SG ACL design is optimal for Layer 3 environment.)	64,000 ACL TCAM entries per ASIC. - Object group ACLs expand into the TCAM space.
✓ LAN and WAN MACsec	- Q200 does not have built-in crypto engine - Catalyst® 9600X-LC use new CDR5M PHY (400 Gbps full duplex). CDR5M provides line-rate (8x 400G = 3.2T) 802.1ae (LAN) MACsec and WAN MACsec.	- UADP3 has built-in MACsec crypto (speed of ASIC). - UADP3 supports only LAN MACsec (no WAN MACsec).
✓ Flexible NetFlow	- Q200 does not have built-in flow cache memory (no hardware-based NetFlow). - Catalyst 9600X use new software-based Flexible NetFlow (FNF) (≤ 2 million entries), with a dedicated CPU core (~2 Mpps). FNF sampler rate 1:1000, ~10 Tbps of 512-byte packets = ~2 Mpps.	- UADP3 has built-in (hardware) FNF, max 64,000 entries per ASIC. - FNF is shared with LPM, etc., in same 416K ASIC memory.

Catalyst 9600 Series core positioning

Next-generation core + edge switching with Silicon One Q200



Feature optimized
Catalyst® C9600-SUP-1



C9600-SUP-1

- Best-in-class enterprise core feature set
- Low speeds (1G to 40G) and port density
- Comprehensive MPLS, EVPN, and SD-Access
- Ideal for campus core, collapsed core + aggregation

**Ideal for Catalyst 6500/6800 Series
non-XL deployment migration**



Performance optimized
Catalyst C9600X-SUP-2



C9600X-SUP-2

- Unmatched forwarding scale and performance
- High speeds (10G to 400G) and port density
- Scaled MPLS and SD-Access, WAN MACsec
- Ideal for campus core + edge, or centralized WLC

**Ideal for Catalyst 6500/6800-XL
deployment migration**





Turn this technical overview into a purchase-ready BOM

Send the exact model or BOM and your deployment requirements. Layer23-Switch can review compatibility, availability, and quote options.

- ✓ Exact model or Cisco PID and quantity
- ✓ Port count, media type, PoE budget, and uplink speed
- ✓ Redundancy, licensing, support, and accessory needs
- ✓ Ship-to country, target date, and lead-time constraints

REQUEST A BOM REVIEW

www.layer23-switch.com



SCAN TO REQUEST A QUOTE

Layer23-Switch is an independent reseller and is not affiliated with or endorsed by Cisco Systems, Inc. Cisco and related marks belong to Cisco Systems, Inc. This reference may include historical information; verify the exact SKU, current specifications, lifecycle status, licensing, and availability before purchase or deployment.