



The bridge to possible

Migration to Silicon One Q200- Based Platforms: C9500X/C9600X

Ninad Diwakar – Technical Marketing Engineer
BRKARC-2096

cisco *Live!*

#CiscoLive

Cisco Webex App


Questions?

Use Cisco Webex App to chat with the speaker after the session

How

- 1 Find this session in the Cisco Live Mobile App
- 2 Click “Join the Discussion”
- 3 Install the Webex App or go directly to the Webex space
- 4 Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until June 7, 2024.





Agenda

- Introduction
- Portfolio Overview
- Why migrate?
- Platform (HW) Migration
- Feature (SW) Migration
- Conclusion

The Catalyst 9000 Core Switch product line

Powering the cloud scale world

Catalyst 9300/9300X UADP 2.0



C9300X-12Y/24Y - 1G/10G/25G
optional 40G/100G uplinks

C9300-24S/48S - 1G SFP
optional 10G/25G/40G uplinks

Stackable with
C9300/C9300X
StackWise

Catalyst 9500/9500X UADP 3.0 and Q200




C9500X-60L4D - 10G/25G/50G SFP
C9500X-28C8D - 40G/100G/400G QSFP
Q200 ASIC

C9500-24Y4C/48Y4C - 10G/25G SFP
C9500-32QC/32C - 40G/100G QSFP
UADP 3.0 ASIC

StackWise Virtual

Catalyst C9600/9600X UADP 3.0 and Q200



C9600-SUP-1 - 1.2Tbps per line card
C9600X-SUP-2 - 3.2Tbps per line card

C9600-LC-24C - 40G/100G QSFP
C9600-LC-48YL - 1G/10G/25G/50G SFP
C9600-LC-40YL4CD - SFP and QSFP combo
C9600X-LC-32CD - 40G/100G/400G QSFP
C9600X-LC-56YL4C - 56x SFP with QSFP Uplinks

Quad Sup RPR*

Extending Core Leadership with Catalyst X Series

- 2.6x Bandwidth
- WAN MacSec Support
- 400G IPsec Support
- 200G/400G QSFP-DD Uplinks
- 2M IPv4 Route Scale
- 8GB HBM Deep Buffers

Cisco Catalyst 9000 Core Series leadership

- SD-Access
- MACsec-256 link encryption
- High Availability (SSO/ISSU)
- Application Hosting
- Model-driven programmability
- Patching and GIR
- Streaming telemetry

Powered by Cisco IOS XE

* Roadmap

Catalyst 9000 Core Switching

Campus switching for any type of network

Campus Core Distribution



Cisco Unified Access Data Plane
UADP 3.0



Catalyst 9500



Catalyst 9600

Catalyst 9500/9600 Models

L2 and L3 focused with MPLS, MVPN, NAT services, SDA LISP, EVPN, and StackWise virtualization.
Comprehensive features for most Campus core placement.

Campus Core Gateway



Cisco Silicon One
Q200 ASIC



Catalyst 9500X/9600X Models

L3 focused with 2M IPv4, 256K MAC with WAN MACsec, scaled MPLS, and deep QoS. EVPN Spine, LISP Border
Ideal for upgrading 6500/6800 XL deployments.

Cisco Campus innovations



SDA LISP, SDA EVPN
EVPN BGP, MPLS, GRE



Streaming Telemetry,
Programmability, PNP



MACsec 256
WAN MACSec*




On-box
App Hosting



HA, Hot Patching,
ISSU, GIR

* C9500X/C9600X only

Catalyst 9000 Core Deployments



New High Density Line Card for C9606



C9600X-LC-56YL4C

56 x 10G/25G/50G + 4x40G/100G

Combo line card with **dense 10G/25G/50G** and **100G** ports



256⁺ x SFP+ ports serving high port density along with LAN / WAN MACsec

Need even more density? Contact a Catalyst Switching PM today!



First line card with 3 rows of ports
Supported with SUP-2

- 3.2Tbps with SUP-2



Flexible deployment options

- Uplink and downlink
- 10G through 100G on the same slot
- Breakout and QSA support**




High-speed secure transport

- IPSEC*, MACsec, WAN-MACsec
- Hitless MUX (HMUX)


Generation of Platforms merged into IOS XE

Now is the time to refresh to Catalyst 9000 Core



Catalyst 9000X core series

Two upgrade paths based on use case and scale



No plans to EoS/EoL the first generation of Catalyst 9000 series of switches.

Don't hesitate to position a generation 1 Catalyst 9000 series core switch.

Thinking of moving to a Catalyst 9500X or 9600X?

Migration Checklist

Top Things to Consider

- 1 HW Scale
- 2 Feature Availability on Q200
- 3 QoS & Buffers
- 4 ACLs & Scale
- 5 NetFlow & Scale
- 6 MACsec & WAN-MACsec
- 7 1G support

Larger LPM

Larger MAC

Lower ACL

Check Software Support

VoQ Architecture with HBM

8K TCAM with OGACL & SGACL


Software-based Netflow

Encryption in CDR5M PHY

No Native 1G Support

Migration

Hardware Considerations



Cisco Silicon One™ Q200

Industry leading Switching and Routing Silicon



12.8 Tbps



8.1 Bpps



8GB HBM
for deep
buffers



2M IPv4
or 1M IPv6
routes



P4 NPL
Programmable
Pipeline



56G PAM4
256
Serdes

Cisco Silicon One Q200
Industry Leading
12.8T System on Chip



First 7nm ASIC

providing lowest Watts/GE
power consumption



Fully P4 programmable

enabling feature velocity




Multi-Slice architecture

for flexibility and scale

Routing Capabilities with Switching Power and Performance

Cisco Silicon One Q200

ASIC Architecture & Block Diagram




Cisco Silicon One Q200 – 9500X & 9600X


Combining multiple ASICs in one SOC



Catalyst 9500X-28C8D & 60L4D



Catalyst 9600X-SUP2




Catalyst 9500 & 9600 Series Core Positioning


Cisco Next Generation Core + Edge Switching



Feature Optimized **C9500 & C9600-SUP-1**



C9600-SUP-1



C9500


- ✓ Best-in-class Enterprise Distribution & Core Features
- ✓ Lower speeds (1G – 40G) and port density
- ✓ Comprehensive SDA, EVPN and MPLS, and MACsec
- ✓ Best for Campus Core, Collapsed-Core & Distribution

C6K non-XL deployment migration

Performance Optimized **C9500X & C9600X-SUP-2**



C9600X-SUP-2





C9500X

- ✓ Unmatched Forwarding Scale and Performance
- ✓ Higher speeds (10G – 400G) and port density
- ✓ Scaled SDA, EVPN and MPLS, and WAN-MACsec
- ✓ Best for Campus Core + Edge, or Centralized WLC

C6K XL deployment migration



Modular vs. Fixed Platforms



Modular

PROs

- More Flexible
- Longer Life-Cycle
- Higher Port Density
- More Power/Cooling
- Redundant Processors

CONs

- More Complex
- BW limit by Chassis
- Slow(er) Dev & Test
- Lower MTBF
- Higher COGs

Fixed

PROs

- Less Complex
- Swap Chassis for BW
- Faster Dev & Test
- Higher MTBF
- Lower COGs

CONs

- Less Flexible
- Shorter Life-Cycle
- Lower Port Density
- Less Power/Cooling
- Single Processor

Modular Platform Features & Benefits

Redundancy, Expansion, Efficiency & Flexibility




Highest Resiliency



- Redundant Supervisors
- StackWise® Virtual
- Easy Upgrades with ISSU & GIR
- Redundant Fans (Fan-Tray)
- Redundant PSUs (1:1, N+1)



Highest Flexibility



- SUP1 for Small Designs
- SUP2/XL for Large Designs
- Custom ASIC Scale Templates
- Traditional Multi-Layer Designs
- Fabric Overlay Designs



Highest Efficiency



- Lowest Watts per Port
- 3000W Power Supplies
- Titanium Rated (95%) PSUs
- AC and/or DC Power
- Configurable Power Priority



Longest Lifecycle



- Start w/ SUP1 & few Gen1 LCs
- Add Gen1 LCs as Access grows
- Replace SUP1 with SUP2
- Gen1 LCs get a 2X boost
- Add new Gen2 LCs as Core grows



Most Port Options

Mixes of RJ45, SFP & QSFP



C9600-LC-40YL4CD
40x 50G SFP + 2x 100G + 2x 400G QSFP



C9600X-LC-32CD
32x 100G or 24x + 8x 400G QSFP



C9400-LC-48XS
48x 1/10G SFP



C9400-LC-48HX
48 x 10G mGig + UPOE®

Catalyst 9000 Core

Three row units for high port density requirements

NEW

Catalyst 9600X – Supervisor Engine 2



C9600X-LC-56YL4C

IOS-XE: 17.13.1


56 x 10G/25G/50G SFP ports
4 x 40G/100G QSFP ports

Catalyst 9500X-60L4D



C9500X-60YL4D

60 x 10G/25G/50G SFP ports
4 x 40G/100/400G QSFP ports



Row 1

Row 2

Row 3

- Three rows of ports, allowing high port densities without sacrificing form factor.
- Combo SFP and QSFP ports.
- Secure high-density deployments – MACSec, WAN-MACSec and IPsec (hardware capable) supported.



Sustainability Impact – High port density in a smaller area-> Lower cooling requirements

Catalyst 9500X Fixed Core

Interface arrangement

Catalyst 9500X-28C8D



C9500X-28C8D

28 x 40/100G QSFP ports
48x 40G/100/400G QSFP-DD
ports

- Ports 1-14, 23-36 are 40/100G QSFP ports
- Ports 15-22 are 40/100/400G QSFP-DD ports

Catalyst 9500X-60L4D



C9500X-60YL4D

60 x 10G/25G/50G SFP ports
4 x 40G/100/400G QSFP-DD
ports

- Ports 1-30, 35-64 are 10/25/50G SFP ports
- Ports 31-34 are 40/100/400G QSFP-DD ports

High Speed uplinks are in center. Optimized arrangement for efficient cooling

Cisco Catalyst 9500 & 9600 with Q200

SDM templates and scale numbers

IOSXE 17.7.1

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

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

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

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




Silicon One ASIC supports 10G and above speed.




Need a few ports on 1G on C9500X/C9600X ?

1G Support on 9500X/9600X Platforms

For low port density requirements




¹ Contains PHY to convert 10G electrical interface to 1G optical interface



Only supported on native SFP 10/25/50G ports. Not supported on QSFP ports using QSA/breakout

Software

Feature Migration




Catalyst 9000 Series – Common Building Blocks



Programmable x86 Multi-Core CPU

Application Hosting
Secure Containers



Open IOS XE® Polaris

Model-Driven APIs
Modular Patching



Cisco UADP & Silicon One™

Programmable Pipeline
Flexible Tables

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

Most features are similar in syntax and implementation from a software standpoint

Top 12 Feature (Services) Categories

1

L2 Forwarding

MAC Learning, PVST/MST, VTP, DTP, 802.1Q, PVLAN, LLDP/CDP, UDLD, LACP/PAGP, L2 Portchannel, mDNS, BUM, L2 MTU, Storm Control

2

L3 Forwarding

IPv4/v6 FIB, L3 SVI, L3 Sub-Intf, ARP/NDP, IGP, MP-BGP, Static, ECMP, L3 Portchannel, HSRP/VRRP, DHCP & Relay, SISF, BFD, FRR, uRPF, VRF, PBR, NAT, IP MTU, IPSLA

3

Multicast & MVPN

IPv4/v6 MFIB, PIM-DM, ASM, SSM, BiDir, PIM BFD, RPF, MSDP, AutoRP/BSR, IGMP/MLD, Snooping, Boundary/Filtering, MoFRR, MSR, MVPN (MDT), mLDP, TRM/HER

4

Fabric (SDA/EVPN)

LISP, VXLAN-GPO, XTR/PXTR, EVPN, VXLAN, BGP, RR/ASBR, NVE, L2/L3 VNI, IPv4/v6, VRF/Extranet, Anycast GW, MAC/IP Binding, IRB, Host Mobility

5

Management (CLI/GUI/API)

CLI, WebUI, SNMP, HTTP, TLS, SSH/Telnet, Syslog, Debugs, EEM, YANG, NET/RESTConf, gNMI/gRPC, App Hosting, IOX, Docker, LSPAN, FPSAN, ERSPAN, EPC

6

HA & SVL

SSO, RF/CF, NSF/NSR, ISSU/xFSU, SMU, GIR, ECMP, MLAG/VPC Stackwise, SVL, NGSVL, L2/L3 MEC

7

QoS & Buffering

MQC, AutoQoS, HQoS, Policy, Class, Match, Queuing/VoQ, Marking, COS, TOS/IPP, DSCP, EXP, Policing/Shaping, Sub-Intf, Tunnels, Ingress/Egress, IPv4/v6, Multicast, L2/L3, L4/App

8

MPLS (L2/L3VPN)

MPLS, LDP, L2/L3 VPN, EoMPLS, VPLS/HVPLS, PWR, FAT, Inter-AS, MPLS-TE, TE-FRR, PE-CE, 6PE/6VPE,

9

Timing (NTP/PTP)

NTP, IPv4/v6, Server/Client, Stratum, PTP, gPTP, 1588, AES67, 802.1as, AVB, TSN, G8275, GM, OC/BC/TC, SyncE, 1PPS/10MHZ, TOD, GNSS

10

Security & Crypto

AAA, RADIUS, TACACS, IPSG/FHS, 802.1x/MAB, EAP/TLS, NEAT, TWS, ACL, PACL, VACL, RACL, OGACL, SGACL, FQDN, IPv4/v6, CTS/CMD, SXP MACsec/MKA, ClearTag, IPsec/IKE

11

NetFlow & AVC

NetFlow v5/v9, IPFIX, NDE, L2 FNF, L3 FNF, IPv4/v6, L4 FNF, BGP AS/NH, VRF/VPN, Multicast, Tunnels, Sampling, Ingress/Egress, AVC/NBAR, NaaS/ETA

12

Tunnels (GRE/QinQ)

GRE, mGRE, IPv4/v6, VRF-Aware, Multicast, L2oMGRE, MPLSoGRE, QinQ, VLAN Translation, Auto 6to4/4to6 Tunnels, ISATAP, NAT Traversal

Top Features NOT supported on Q200

up to **17.12.x**

1

L2 Forwarding

- Private VLANs
- REP
- Flexlinks
- Bonjour / SDG
- CDP Bypass

2

L3 Forwarding

- NAT / PAT
- IP FRR / LFA
- Adv. PBR (some features)
- WCCP
- GLBP

3

Multicast & MVPN

- PIM Bidir
- PIM Snooping
- MVPN (Rosen & mLDP)
- Multicast mGRE
- MSR
- MVR

4

Fabric (SDA/EVPN)

- LISP Extranet
- L2 EVPN TRM

5

Management (CLI/GUI/API)

- RSPAN
- ERSPAN + IPsec Encap
- Adv. SPAN (some features)
- Adv. EPC (some features)
- Open Flow 1.3

6

HA & SVL

- L3 NSR
- GIR - L3 EIGRP
- Quad Sup SVL (9600X)

7

QoS & Buffering

- Egress QoS ACL TCAM
- Egress Aggregate Policing
- Egress WTD, WRED Thresholds
- QoS on Portchannel interface
- AVC/NBAR based QoS
- Netflow based Microflow Policing

8

MPLS (L2/L3VPN)

- VPLS & H-VPLS
- EoMPLS over LACP & PaGP
- EoMPLS over GRE & mGRE
- MPLS over GRE & mGRE
- TE FRR (some features)

9

Timing (NTP/PTP)

- PTP
- AVB

10

Security & Crypto

- PACL & VACL
- MAC ACL
- FQDN & Reflexive ACL
- L2 CMD
- ACL Hitless Update
- 802.1x (some features)
- RADSEC

11

NetFlow & AVC

- AVC
- Egress (output) FNF
- BGP NH & AS records
- sFlow

12

Tunnels (GRE/QinQ)

- L2oMGRE
- 802.3ad full QinQ
- L3VPN over GRE
- IPv6 Auto Tunnels

Catalyst 9000 Core

Silicon One Q200 positioning

Can you run your network without these features for the foreseeable future?

Do you need these features and are fine with lower speeds and scale of the UADP 3.0 ASIC?

Do you need both the features and the increased scale Q200 based platforms provide?


Position the Silicon One Q200 powered Catalyst core.

Position the UADP 3.0 powered Catalyst core.

Talk to us.

ACL

Access control lists



Key difference on Silicon One Q200

Q200 has minimal (space-optimized) TCAM memory: 8K IPv4/4K IPv6 ACEs.



Object-Group & Security-Group ACLs use CEM to map IP-to-Group label, TCAM only uses L4 ACEs.

OG/SG ACL design is optimal for Core.

Scale


126 Unique ACLs in Ingress
16 Unique ACLs in Egress

Access-Lists on Silicon One Q200

Traditional IP ACL

```
ip access-list extended TestACL
 permit ip host 10.5.64.140 any
 permit ip host 10.5.64.86 any
 permit any host 10.5.75.124
 permit ip 10.10.170.32 255.255.255.224 any
 permit ip 10.179.252.0 255.255.252.0 any
 permit 10.5.200.240 255.255.255.240 any
 permit ip any 10.53.199.0 255.255.255.0
 permit ip any 10.56.100.224 255.255.255.240
... Truncated (total 10,000 entries)
 permit host 10.56.223.56
```

```
interface Vlan100
 ip access-group TestACL in
```




permit/deny operations are
programmed into the ACL space

10,000 entries programmed in an 8,000 entry TCAM space


ACL - Optimizing the hardware resources

Optimized space utilization using OG-ACL/SG-ACLs*



Entries in object-groups use LPM space instead of TCAM space

Traditional ACL vs Object Group ACL



ACL – Benefits

OG-ACL/SG-ACL vs Traditional ACL

Traditional IP ACL

- Easy to get started.
- Granular control but higher TCAM utilization.
- High compatibility.

- Intended for smaller configurations.
- Increase in complexity for large configurations.
- Requires granular source and destination configurations for full security control

OG-ACL

- Simplified configuration by clubbing hosts/IPs into logical groups.
- Improved readability and management.
- Flexible design allows easy modification of entries.
- Group-based policy-enforcement.

- Full benefits do not apply for smaller configurations.

SG-ACL

- Ties ACL security to SGTs -> Identity of the end device
- Dynamic control of security based on end device.
- Simplified policy management
- Integrates port-based security and user authentication.

- Requires an ISE integration to really get all the benefits.
- Manual SGT tagging without ISE quickly ramps up complexity.

Access-Lists on Silicon One Q200

Configuration Considerations

Configuration considerations

Use OG-ACL/SG-ACLs to simplify operations

Keep track of ACL-ID limitations


Small ACLs for debug/troubleshooting can be configured as-is

Scale

126 Unique ACLs in Ingress
16 Unique ACLs in Egress

QoS

Quality of Service



QoS on Silicon One Q200

Considerations and benefits

Feature support – As similar as possible to UADP based platforms

Benefits

Ingress queueing (Virtual output Queueing)

Drops packets before using the switching backplane.
Control what uses the fabric.

Higher Buffers

Better burst absorption. 80MB/8GB vs 36MB (UADP ASIC)

Sub-Interface Queueing

Apply policies on sub-interface rather than on the parent interface.

Upto 7 priority queues

Granular control over latency priority of packets leaving the system.

Considerations

Different QoS model


Requires change in configuration from UADP as well as older Cat6K, 4K etc

Ingress buffering

Buffering is performed on ingress moving the drop point for the traffic at ingress

Upgrading to S1 from Cat6K/UADP

Traffic class and traffic color




- S1 Q200 appends two separate tags to the packet header.
- These tags are significant only within the switch.

Traffic-class and color must be tagged MANUALLY

VoQ QoS


The fundamental difference vs UADP

UADP



Can forward and define priority using [dscp, cos or ip prec](#)

Silicon One Q200




- Additional field appended to the packet header while processing within the switch.
- Field added when packet enters the switch
- Field removed when packet exits the switch
- Neighbouring switches have **no idea about traffic-class**

Can **ONLY** forward and define priority using [Traffic-class](#)


Head of Line Blocking

Issue with Ingress queueing



Virtual Output Queueing (VoQ)

Addressing Head of line Blocking




VoQ Architecture first for a campus switch. Used in our routing and data centre products

Traffic class


Define packet priority

Priority Level

Traffic class



Traffic color – discard-class



Traffic must be colored at ingress to use WTD/WRED

Putting it all together

Silicon One Q200 QoS config

Step 1

When the traffic comes in, define the dscp tags which will use traffic-class. Optionally color the traffic.

Example:

- Traffic with dscp EF goes to *traffic-class 7*
- Traffic with dscp af11 goes to *traffic-class 6* and is colored GREEN
- Traffic with dscp af12 goes to *traffic-class 6* and is colored YELLOW
- Traffic with dscp 31 goes to *traffic-class 5*...

Step 2

Define the priority and queueing for the traffic.

Example:

- *Traffic-class 7* is priority level 1
- *Traffic-class 6* is normal queue with 35% weighted bandwidth ...
- YELLOW traffic in Traffic-class 6 has lower priority than GREEN traffic

Putting it all together

Silicon One Q200 config considerations

INGRESS POLICY

Marking

Map traffic header/ACL/VLANs to traffic-class. Optionally color traffic.

EGRESS POLICY


Queueing

Perform queueing based on traffic-class. Perform WRED/WTD based on packet color.

Putting it all together – Contd.

Silicon One Q200 config considerations

Map dscp to traffic-class and then use traffic-class at egress




Putting it all together

Building a Silicon One Q200 configuration - 2P2Q2T Example

```
class-map match-any VOICE-PQ1
  match dscp ef
class-map match-any VIDEO-PQ2
  match dscp 24
class-map match-any USER-TRAFFIC
  match dscp cs3 af31 af32
Class-map match-any USER-TRAFFIC-YELLOW
  match dscp af33
```

```
policy-map CAMPUS_INGRESS_POLICY
  class VOICE-PQ1
    set traffic-class 7
  class VIDEO-PQ2
    set traffic class 6
  class USER-TRAFFIC
    set traffic class 5
  Class USER-TRAFFIC-YELLOW
    set traffic class 5
    set discard-class 1
```

```
Interface gig1/0/1
  service-policy input CAMPUS_INGRESS_POLICY
```



```
class-map match-any tc7
  match traffic-class 7
class-map match-any tc6
  match traffic-class 6
class-map match-any tc5
  match traffic-class 5

policy-map CAMPUS_EGRESS_POLICY
  class tc7
    priority level 1
  class tc6
    priority level 2
  class tc5
    bandwidth remaining ratio 65
    random-detect discard-class-based
    random-detect discard-class 0 percent 80 100
    random-detect discard-class 1 percent 40 100
  class class-default
    bandwidth remaining ratio 35
```

```
Interface gig1/0/1
  service-policy output CAMPUS_EGRESS_POLICY
```

*Text is colored corresponding to queue colors used in the graphic

Buffering and use case

Silicon One Q200


8 GB HBM – Secondary Buffers
Addresses speed imbalances and traffic bursts

80MB SMS – Primary Buffers
Takes care of low latency traffic requirements

Buffering at ingress – Helps optimize fabric utilization

Historically, all our switches (Catalyst 2K, 3K, 4K, 6K and UADP 9K) buffer in the EGRESS

Unique hybrid buffer complex enables Core/Campus Gateway positioning of Silicon One Q200 Devices



Config Migration steps

1 Define the problem/behavior addressed with QoS.

Simply copy-pasting existing configs between platform families will always throw errors due to differences in syntax and supported actions between platforms.

2 How many times do you want to split your traffic – Upto 8 queues possible with our ASIC

Its often not as much as you think you need. Broad generalized splits often are more efficient than granular splits

3 Do you want multiple strict priority classes?

Know what strict means. All traffic coming into it will be serviced at the expense of other classes.

4 Define traffic shaping/policing or sharing between queues.

Police/shape priority queues. Use weights to control bandwidth sharing with remaining queues

5 Do you want differential priority for packets within same class? – Use WTD or WRED

Advanced configuration options, not required for most use cases.

Migration from Catalyst 6K to Silicon One Q200

Catalyst 6K Configuration

```
class-map type lan-queuing match-all REALTIME
  match dscp ef
class-map type lan-queuing match-all NETWORK_CONTROL
  match dscp cs6 cs7
class-map type lan-queuing match-all VIDEO
  match dscp af31 af32 af33
...
policy-map type lan-queuing CAMPUS_EGRESS_6800_POLICY
  class type lan-queuing REALTIME
    priority level 1
  class type lan-queuing NETWORK_CONTROL
    bandwidth remaining percent 10
  class type lan-queuing VIDEO
    bandwidth remaining percent 20
  class class-default
    random-detect dscp-based
    random-detect dscp af11 percent 80 100
```

1. Classified Based on DSCP value
2. 4 classes (3 defined + 1 class-default)

3. 4 total queues
4. 1 priority queue and 3 normal queues.
5. Scheduling is WRR with “bandwidth remaining”
6. WRED used on the class-default

Note: C6K WRED default min is 40, max is 100.

Migration from Catalyst 6K to Silicon One Q200

Understanding the WRED configurations

Catalyst 6K Configuration

```
...  
...  
policy-map type lan-queuing CAMPUS_EGRESS_6800_POLICY  
...  
class class-default  
  random-detect dscp-based  
  random-detect dscp af11 percent 80 100
```

Note: C6K WRED default min is 40, max is 100.

- Dscp af11 uses threshold of 80 to 100.
- All other dscp values hitting class-default uses threshold of 40 to 100

Color af11 green

Color everything else yellow

Config Migration from Catalyst 6K to Silicon One Q200

Catalyst 6K Configuration

```
class-map type lan-queuing match-all REALTIME
  match dscp ef
class-map type lan-queuing match-all NETWORK_CONTROL
  match dscp cs6 cs7
class-map type lan-queuing match-all VIDEO
  match dscp af31 af32 af33
...
policy-map type lan-queuing CAMPUS_EGRESS_6800_POLICY
  class type lan-queuing REALTIME
    priority level 1
  class type lan-queuing NETWORK_CONTROL
    bandwidth remaining percent 10
  class type lan-queuing VIDEO
    bandwidth remaining percent 20
  class class-default
    random-detect dscp-based
    random-detect dscp af11 percent 80 100
```

1. Classified Based on DSCP value
2. 4 classes (3 defined + default)
3. 4 queues (traffic-class), traffic-7 is priority level 1



Apply policy on the ingress interface

Catalyst 9K(Q200) Configuration

```
class-map match-all REALTIME
  match dscp ef
class-map match-all NETWORK_CONTROL
  match dscp cs6 cs7
class-map match-all VIDEO
  match dscp af31 af32 af33
class-map match-all default-green
  match dscp af11

policy-map INGRESS
  class REALTIME
    set traffic-class 7
  class NETWORK_CONTROL
    set traffic-class 6
  class VIDEO
    set traffic-class 5
  class default-green
    set traffic-class 0
  class class-default
    set discard-class 1
```


Note: class-default is always assigned with traffic-class 0

Config Migration from Catalyst 6K to Silicon One Q200

Map the traffic-class marking defined on the ingress

Catalyst 6K Configuration

```
policy-map type lan-queuing CAMPUS_EGRESS_6800_POLICY
  class type lan-queuing REALTIME
    priority level 1
  class type lan-queuing NETWORK_CONTROL
    bandwidth remaining percent 10
  class type lan-queuing VIDEO
    bandwidth remaining percent 20
  class class-default
    random-detect dscp-based
    random-detect dscp af11 percent 80 100
```



Apply policy on the egress interface

Catalyst 9K(Q200) Configuration


```
class-map tc7
  match traffic-class 7
class-map tc6
  match traffic-class 6
class-map tc5
  match traffic-class 5

policy-map type queuing EGRESS
  class tc7
    priority level 1
  class tc6
    bandwidth remaining ratio 1
  class tc5
    bandwidth remaining ratio 2
  class class-default
    random-detect discard-class-based
    random-detect discard-class 0 percent 80 100
    random-detect discard-class 1 percent 40 100
```

3. 1 priority queue
4. Scheduling is WRR with “bandwidth remaining”
5. Congestion management is WRED with the default class

Note: C6K WRED default min is 40, max is 100.

NetFlow



Key difference on Silicon One Q200





Q200 does not have built-in Flow Cache memory.

C9500X & C9600X uses new Software-based FNF ($\leq 2\text{M}$ entries), with a dedicated CPU core ($\sim 2\text{Mpps}$).

Upgrading to S1 from Cat6K/UADP

NetFlow - It is in the name



Upgrading to S1 from Cat6K/UADP

Key difference

Software based NetFlow

Number of flows, rate of traffic has an impact on the CPU

Sampled NetFlow

Inspecting every packet (full flexible NetFlow) is not possible with 100s of GB of traffic.

Ingress NetFlow

Only supported in ingress direction. Egress NetFlow unsupported.

Random Sampling

One packet at random out of sample size picked for parsing to the CPU

Silicon One Q200 NetFlow

True Random Sampling

True Random Sampling

Packet 320 in the first 1000 packets (sample size 1000) is selected and parsed. In the next 1000 packets, packet 20 is selected, then packet 456 and so on)

Deterministic Sampling (flow-based sampling) is NOT supported


NetFlow requires the Network-Advantage license set

***Higher cache size - over a period, most flows will be cached**

Silicon One Q200 NetFlow

Choosing the right sample size

Software NetFlow -> Traffic hits the CPU -> Traffic rate can cause CPU to spike up to 100%



- Default and Recommended sample size of 1:1000
- More aggressive sampling can cause dedicated CPU core to spike to 100%
- Sampling can go as low as 1:2

Silicon One Q200 NetFlow

Building a sample configuration

```
flow record v4_core
  description Used for basic IPv4 traffic analysis
  match ipv4 tos
  match ipv4 protocol
  <truncated>
```

```
flow exporter ISO-SWOLF
  destination 192.1.1.1
  dscp 23
  ttl 15
  transport udp 2055
  template data timeout 1200
```

```
flow monitor V4_FLOWS
  exporter ISO-SWOLF
  cache timeout inactive 20
  cache timeout active 180
  cache entries 2000000
  record v4_core
```

```
sampler 1in1000
  mode random 1 out-of 1000
```

```
interface FiftyGigE1/0/9
  ip flow monitor V4_FLOWS_109 sampler 1in1000 input
end
```

Same syntax and commands as UADP


Mandatory Sampler configuration

Apply on Interface in INGRESS direction.
NOT supported on EGRESS.

Silicon One Q200 Sampled NetFlow

Role in the network

Access layer – Identify the traffic




Granular (1:1) flow identification at the Access

- Full Flexible NetFlow for granular identification of traffic
- Full accounting of every client/flow (ETA/AVC)

Low-Moderate scale at the Access

- Fewer number of connected clients/flows
- Average ~1K clients x ~32K flows per Access

Core/Gateway Layer – Monitor



Aggregate (1:1K) monitoring of flows at the Core


- Just need to monitor the overall network usage
- Adjust sample rates to balance scale & load

Medium-High scale at the Core

- Need to aggregate all clients/flows (# Access x 32K)
- Adjust cache aging to increase overall scale

Secure Transport

MACSec and WAN-MACSec



Key difference on Silicon One Q200

- Q200 does not have built-in crypto engine.
- This causes some Gen-1 line cards on 9600X to not support MACSec when used with Supervisor engine 2.
- C9500X & C9600X-LC uses new CDR5M PHY (400Gbps Full-Duplex).
- CDR5M provides line-rate (8x 400G = 3.2T) 802.1ae (LAN) MACsec and WAN-MACsec.
- C9500X-60L4D and C9600X-LC-56YL4C* support hardware-based IPsec.




*Hardware capable.

Cisco Silicon One Q200 – 9500X & 9600X


Combining multiple ASICs in one SOC



Catalyst 9500X-28C8D & 60L4D



Catalyst 9600X-SUP2



Catalyst 9600

Secure Transport support




	C9606 Chassis	
	C9600-SUP- 1	C9600X-SUP- 2
C9600-LC-24C	✓ MACSec Supported	✗ MACSec/WAN MACSec not Supported
C9600-LC-48YL	✓ MACSec Supported	✗ MACSec/WAN MACSec not Supported
C9600-LC-48TX	✓ MACSec Supported	✗ MACSec/WAN MACSec not Supported
C9600-LC-48S	✓ MACSec Supported	LINE CARD NOT SUPPORTED
C9600-LC-40YL4CD	✓ MACSec Supported	✓ MACSec/WAN MACSec Supported
C9600X-LC-32CD	LINE CARD NOT SUPPORTED	✓ MACSec/WAN MACSec Supported
C9600X-LC-56YL4C	LINE CARD NOT SUPPORTED	✓ MACSec/WAN MACSec/IPSec* Supported

* Hardware capable


WAN MACsec and LAN MACsec Comparison

	WAN MACsec	LAN MACsec
Connection Type	Over L2 MPLS, VPLS, EoMPLS, QinQ, Multiple Point to Point	Only Directly Connected Devices
AES 128	✓	✓
AES 256	✓	✓
IPv4 and IPv6 Independent	✓	✓
Performance	Line rate on all Ports (C9600X, C9500X)	Line rate on all Ports
Operation and Configuration	Simple	Simple
Overhead	32 Bytes	32 Bytes

End-to-end security with MACSec and WAN-MACSec



Conclusion



Silicon One Q200 vs UADP

Benefits and conclusion



Feature optimized UADP ASIC

- Best-in-class enterprise core feature set
- Medium speeds (1G to 100G) and port density
- Comprehensive MPLS, EVPN, and SD-Access
- Campus Core, **Collapsed-Core or Distribution**



Performance optimized S1 Q200 ASIC


- Unmatched forwarding scale and performance
- High speeds (10G to 400G) and port density
- Scaled MPLS, EVPN and SD-Access, WAN MACsec, IPsec*
- **Campus Core & Campus Gateway**, or centralized WLC

Complement NOT Compete

Benefits of Upgrading to Silicon One Q200

Core/Campus Gateway Positioning

- Learn full routing table (2M IPv4 routes)
- Core focussed feature table sizes (256K Mac scale, 256K host table)
- Secure transport with WAN-MACSec and IPsec*



Port Scale

- 8 x 400G ports – Both 9500X and 9600X
- 128 x 100G ports (without breakout) – On 9600X
- 224 x 10/25/50G ports (without breakout) – On 9600X



C9500X-60L4D




C9600X-LC-56YL4C

Q1
CY24

Buffering

- Ingress buffering for better control over switch fabric utilization
- Hybrid dual buffering complex.
- Buffers catering to both deep burst absorption and quick low latency switching



*IPsec is only supported on C9500X-60L4D switch model

NOTE: C9600 with C9600X-LC-56YL4D line card is hardware capable of IPsec

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

73

cisco *Live!*

#CiscoLive

BRKARC-2096

Thinking of moving to a Catalyst 9500X or 9600X?

Migration Checklist

Top Things to Consider

- 1 HW Scale
- 2 Feature Availability on Q200
- 3 QoS & Buffers
- 4 ACLs & Scale
- 5 NetFlow & Scale
- 6 MACsec & WAN-MACsec
- 7 1G support

Larger LPM

Larger MAC

Lower ACL

Check Software Support

VoQ Architecture with HBM

8K TCAM with OGACL & SGACL

Software-based Netflow

Encryption in CDR5M PHY

No Native 1G Support

Miercom Report

Catalyst 9500X and 9600X performance reports



Independently verified and tested performance and latency for Cisco 9500X and 9600X series switches

Miercom.com/cisco

Complete Your Session Evaluations



Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to **win 1 of 5 full conference passes** to Cisco Live 2025.



Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.



Level up and earn **exclusive prizes!**



Complete your surveys in the **Cisco Live mobile app**.

Continue your education

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

Contact me at: ndiwakar@cisco.com



The bridge to possible

Thank you

CISCO *Live!*

#CiscoLive