cisco

The bridge to possible

Migrating from Cisco Catalyst 6500/6800 to 9600 Series Switches

ılıılıı
CISCO
The bridge to possible

Contents

Introduction

The Cisco® Catalyst® 9000 switching platform is the next generation in the legendary Cisco Catalyst switching family. It is designed for the new era of networking, with ASIC and software innovations to deliver an intent-based network. Within the Cisco Catalyst 9000 switching family, the Cisco Catalyst 9600 Series Switches are Cisco's leading modular enterprise switching core and distribution platform, built for intent-based architecture, security, Internet of Things (IoT), and cloud.

Purpose of this guide

This document is intended to help network planners and engineers who are familiar with the Cisco Catalyst 6500 and 6800 Series to deploy Cisco Catalyst 9600 Series Switches in the enterprise networking environment (Figure 1).

Figure 1.   Migration to the Cisco Catalyst 9600 Series

CISCO

The bridge to possible

Contents

Why migrate?

The Cisco Catalyst 9600 Series Switches are Cisco's leading modular core/distribution enterprise switching platform, built for security, IoT, and cloud. These switches form the foundational building block for Software-Defined Access (SD-Access), Cisco's leading intent-based architecture. The platform provides unparalleled investment protection with a chassis architecture that is capable of supporting up to 25.6 Tbps of system bandwidth. The 9600 Series expands the end-to-end Cisco Catalyst 9000 family, further aligning our comprehensive solution with our customer carrying business and network strategies. When you need the most secure, most resilient, and most intelligent network, you need Cisco Catalyst 9600 Series Switches, the industry's first purpose-built modular-configuration 1G, 2.5G, 5G, 10G, 25G, 40G,50G,100G and 400G line of switches for enterprise-class core and distribution layers.

The Cisco Catalyst 9600 Series offers an industry-leading supervisor engine built for secure networks, IoT applications, next-generation mobility, and cloud adoption. Supervisor Engine 1 (C9600-SUP-1) is built with the Cisco Unified Access® Data Plane 3.0 (UADP 3.0) Application-Specific Integrated Circuit (ASIC), and Supervisor Engine 2 (C9600X-SUP-2) is built with the Cisco Silicon One Q200 ASIC. Both Supervisors are ready for next-generation technologies with its programmable pipeline, microengine capabilities, and customizable allocations of Layer 2, Layer 3, forwarding, access control lists (ACLs), and quality-of-service (QoS) entries.

Migration overview

The Cisco Catalyst 9600 Series Switches use a centralized architecture. All forwarding, security, and queueing is done on the supervisor, while the line cards are considered transparent, with only PHYs and control logic. The simplicity of this centralized design allows easy upgrade of features as well as additional bandwidth by upgrading just the supervisor while keeping the existing line cards. The combination of the centralized architecture and transparent line cards also provides uninterrupted supervisor switchover, which is the foundation for In-Service Software Upgrades (ISSU). With the high-capacity bandwidth provided by the Supervisor Engine 1 and Supervisor Engine 2, all ports (including 400G) are line rate with the supported configuration. The latency between all ports are less than 5 microseconds, which is more than sufficient for any time-sensitive application such as real-time video conferencing and IP telephony.

This guide lists the different considerations when migrating from the Cisco Catalyst 6500/6800 Series to the 9600 Series.

Contents

Cisco Catalyst 9606R port density

The Cisco Catalyst 9606R is a 6-slot chassis. Two middle slots are dedicated for supervisors only, and they work in redundant mode. The top and bottom two slots are for line cards. The chassis is designed to provide up to 6.4 Tbps from each of the supervisor slots to each of the line card slots. This means the system will be able to provide 32 ports of 100G at line rate per line card slot. In total, the 9606R is capable of providing either:

- 32x line rate 400G ports or

- 128x line rate 100G/40G ports or

- 192x line rate 50G/25G/10G/5G/2.5G/1G ports.

The Cisco Catalyst 9600 Supervisor Engine 1 is powered with three UADP 3.0 ASICs. The aggregated bandwidth from the three ASICs is 9.6 Tbps. The port density with Supervisor Engine 1 is as follows:

- 48x line rate 100G/40G ports or

- 192x line rate 25G/10G/5G/2.5G/1G ports

The Cisco Catalyst 9600 Supervisor Engine 2 is powered with One Cisco Silicon One ASIC. The aggregated bandwidth from ASIC is 12.8 Tbps. The port density with Supervisor Engine 2 is as follows:

- 128x line rate 100G/40G ports or

- 192x line rate 50G/25G/10G ports or

- 8x line rate 400G + 96x line rate 100/40G ports or

- 8x line rate 400G + 8x line rate 100/40G + 160 x line rate 50/25/10G ports

The bridge to possible

Contents

Supervisor hardware

The Cisco Catalyst 9600 Supervisor Engine 1 is based on Cisco's UADP 3.0 ASIC architecture and an x86 CPU architecture. The Cisco Catalyst 9600 Supervisor Engine 2 is based on Cisco Silicon One ASIC architecture and an x86 CPU architecture. Both Supervisor Engines also provide options for additional external storage, which enables the device to host containers and run third-party applications and scripts natively within the switch. Table 1 compares the hardware of the Cisco Catalyst 6500/6800 and 9600 Series.

Table 1.   Hardware comparison

	Cisco Catalyst 6500/6800 Series (Supervisor Engine 6T)	Cisco Catalyst 9600 Series (Supervisor Engine 1)	Cisco Catalyst 9600 Series (Supervisor Engine 2)
CPU	Dual-core x86, 2.5 GHz	8-core x86, 2.0 GHz	8-core x86, 2.7 GHz
SDRAM	4 GB	16 GB	32 GB
Internal flash	4 GB	16 GB	16 GB
External storage	4 GB USB	480, or 960 GB*	480, or 960 GB*

* With Cisco certified SSD drives

The bridge to possible

Contents

Cisco Catalyst 9600 Series flexible and customized SDM templates

Unlike the supervisor on the Cisco Catalyst 6500/6800 Series, Supervisor Engines on the 9600 Series enables flexible Software Database Manager (SDM) templates for universal deployments by leveraging the ASIC's ability to create resources to optimize table sizes for different purposes in the network. Based on how the switch is used in the network, an appropriate flexible ASIC template may be selected to configure the switch for specific features.

Cisco Catalyst 9600 Supervisor Engine 1 supports the following templates:

· Distribution: Balances resources between Layer 2 and Layer 3

· Core: Maximizes system resources for unicast and multicast routing

· NAT: Maximizes system resources for Layer 3 and Network Address Translation (NAT)

· SD-Access: Maximizes system resources to support fabric deployment (till 17.3 1 release)

Cisco Catalyst 9600 Supervisor Engine 2 supports the following template:

· Core: Maximizes system resources for unicast and multicast routing

In the campus, the most commonly deployed design is the three-tiered design: core, distribution, and access. The core layer is based on Layer 3 IP routing and functions as a high-speed interconnection point to other network domains (data center, WAN, branch, etc.). The distribution layer traditionally consists of IP routing upstream to the core and Layer 2 switching downstream to the access layer. The access layer is purely Layer 2 and provides connectivity for the endpoints. With this design, the Core template would be the best fit for the core devices, and the Distribution template is most appropriate for the distribution devices.

The routed access design, which also can be the three-tiered design, moves the Layer 2 and Layer 3 boundaries to the access layer. In this case, the distribution layer is purely Layer 3 and doesn't require the Layer 2 features. In this design, both the core and distribution layer devices should use the Core template.

Contents

Standard SDM templates can be used to configure system resources and optimize support for specific features. However, SDM templates are defined based on how the device is deployed in the network. A custom SDM template allows you to configure several features of the template based on your requirements and not the location of the device in the network. Both Supervisor Engine 1 and Supervisor Engine 2 support customizable SDM template for both FIB and TCAM allocations. Details information of customizable SDM template is available on C9600 datasheet.

System default behaviors

The system default behaviors on the Cisco Catalyst 9600 Series are similar to those of the Cisco Catalyst 6500/6800 Series. For example, IP routing is enabled, the management interface is in a dedicated Virtual Routing and Forwarding (VRF) instance, and so on. However, there are also some differences, as described in this section.

Table 2.   Interface default state

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Default port type	Layer 3	Layer 2
Default port state	Off (shutdown)	On (no shutdown)

Contents

Table 3.   Port-channel default load balance

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series

EtherChannel load balance	EtherChannel Load-Balancing Configuration: src-dst-ip enhanced
	
	mpls label-ip
	
	EtherChannel Load-Balancing Addresses Used Per-Protocol:
	
	Non-IP: Source XOR Destination MAC address
	
	   IPv4: Source XOR Destination IP address
	
	   IPv6: Source XOR Destination IP address
	
	   MPLS: Label or IP
		EtherChannel Load-Balancing Configuration: src-dst-mixed-ip-port
		
		EtherChannel Load- Balancing Addresses Used Per-Protocol:
		
		Non-IP: Source XOR Destination MAC address
		
		   IPv4: Source XOR Destination IP address and TCP/UDP (layer-4) port number
		
		   IPv6: Source XOR Destination IP address and TCP/UDP (layer-4) port number

Link-status logging: The logging for link-status changes is on by default in the Cisco Catalyst 9600 Series, and the behavior can be changed per interface in the configuration. In the 6500/6800 Series, the logging for link-status changes is off by default and can be changed globally. See Table 4.

Contents

Table 4.   Link-status logging comparison

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Default	Off for all ports	On for all ports
Configuration	Per system C6500(config)#logging event link-status global	On (no shutdown) C9600(config)#int tw 1/0/1 C9600(config-if)#no logging event link-status
	C6500(config)#no logging event link-status global	C9600(config-if)#logging event link-status

Power redundancy

The Cisco Catalyst 9600 Series provides four slots for the power supply, compared to two slots in the 6500/6800 Series. Power supplies can operate in the following two modes:

· Combined mode: This is the default mode. All power supply modules in the system are active and sharing power.

· N+1 redundant mode: The system enters this mode if the output from the N power supply is sufficient.

For more details on power redundancy, please see the Environmental Monitoring and Power Management chapter of the System Management Configuration Guide.

Contents

ROMmon and config-register

The Cisco Catalyst 9600 Series uses the x86 CPU architecture, which can be used in the future to enable hosting containers and third-party applications. With the CPU architecture change, there are also changes in the ROM monitor (ROMmon).

Prompts and file system

In ROMmon, the prompt on the Cisco Catalyst 9600 Series is "rommon >" and "bootflash:" is the memory partition for local storage. On the Cisco Catalyst 6500/6800 Series, the prompt is "rommon>" and the memory partition for the local storage is "bootdisk:". Table 5 shows outputs in ROMmon from the Cisco Catalyst 6500/6800 and 9600 Series.

Table 5.   ROMmon outputs

Cisco Catalyst 6500/6800 Series			
rommon 20 > dir bootdisk:			
File System: FAT32			
3	33554432	-rw-	sea_console.dat
8195	162870776	-rw-	s6t64-ipservicesk9-mz.SPA.154-1.SY2.bin
47959	33554432	-rw-	sea_log.dat
56151	169094712	-rw-	s6t64-adventerprisek9-mz.SPA.153-1.SY2.bin
97434	7766	-rw-	startup-config.converted_vs-20171005-221623

The bridge to possible

Contents

Cisco Catalyst 6500/6800 Series			
rommon 20 > dir bootdisk:			
File System: FAT32			
97436	13820	-rw-	C6807-D1-SUP6T-VSS-ECMP
97440	14171	-rw-	C6807-D1-SUP6T-VSS-VSS
97444	14191	-rw-	C6807-D1-S6T-IPV6-VSS-VSS
97448	11151	-rw-	startup-config.converted_vs-20180813-221352
97451	11051	-rw-	startup-config.converted_vs-20180813-225913
rommon 21 >			

Cisco Catalyst 9600 Series		
rommon 5 >	dir bootflash:	
Size	Attributes Name	
4096	-rw-	.installer
962	-rw-	bootloader_evt_handle.log

The bridge to possible

Contents

Cisco Catalyst 9600 Series		
rommon 5 >	dir bootflash:	
Size	Attributes Name	
4096	-rw-	.ssh
4096	-rw-	core
4096	-rw-	.prst_sync
4096	-rw-	.rollback_timer
4096	-rw-	gs_script
4096	-rw-	tech_support
4096	-rw-	dc_profile_dir
324	-rw-	boothelper.log
132095	-rw-	memleak.tcl
3	-rw-	.wlc_air_lic
4096	-rw-	onep

Contents

Cisco Catalyst 9600 Series		
rommon 5 >	dir bootflash:	
Size	Attributes Name	
0	-rw-	rdope_out.txt
91	-rw-	rdope.log
242	-rw-	smart_overall_health.log
35	-rw-	pnp-tech-time
71415	-rw-	pnp-tech-discovery-summary
60556	-rw-	vlan.dat
90523228	-rw-	sf-linux-1017.SSA
25124888	-rw-	issg_v232_1114.zb
4096	-rw-	tan
11359240	-rw-	cat9k-cc_srdriver.16.11.01.SPA.pkg
84354052	-rw-	cat9k-espbase.16.11.01.SPA.pkg

Contents

Cisco Catalyst 9600 Series		
rommon 5 >	dir bootflash:	
Size	Attributes Name	
1676292	-rw-	cat9k-guestshell.16.11.01.SPA.pkg
466576384	-rw-	cat9k-rpbase.16.11.01.SPA.pkg
38552418	-rw-	cat9k-rpboot.16.11.01.SPA.pkg
29877252	-rw-	cat9k-sipbase.16.11.01.SPA.pkg
57259008	-rw-	cat9k-sipspa.16.11.01.SPA.pkg
19936260	-rw-	cat9k-srdriver.16.11.01.SPA.pkg
12321792	-rw-	cat9k-webui.16.11.01.SPA.pkg
9216	-rw-	cat9k-wlc.16.11.01.SPA.pkg
7612	-rw-	packages.conf
4096	drw-	images
rommon 6 >		

Boot variables

The Cisco Catalyst 6500/6800 Series uses the traditional "config-register" command in both Cisco IOS® and ROMmon to control the booting behavior. The Cisco Catalyst 9600 Series uses a parallel set of commands in Cisco IOS XE Software, which creates the equivalent ROMmon variables. See Table 6 and 7.

Table 6.   Boot variables

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Cisco IOS Software	Confreg 0x???Y Autoboot if Y!=0	[no] boot manual
ROMmon	Confreg 0x???Y Autoboot if Y!=0	MANUAL_BOOT=[no | yes]

Baud rate

With the Cisco Catalyst 9600 Series, the user can set the baud rate in either the Cisco IOS XE command-line interface (CLI) or ROMmon. See Table 7.

Table 7.   Setting the baud rate

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Cisco IOS Software	Confreg 0x???? or Line con 0 Speed 9600	Line con 0 Speed 9600
ROMmon	Confreg Use the interactive prompt to set the baud rate	BAUD=9600

Contents

"Break" processing

At the beginning of the bootup process, the user can use Ctrl+C to break out of the booting process and drop the system back into ROMmon if the break sequence is enabled. See Table 8.

Table 8.   "Break" processing

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Cisco IOS Software	Confreg 0x????	[no] boot enable-break
ROMmon	Confreg Use the interactive prompt to set the baud rate	ENABLE_BREAK=[no | yes]

Ignoring the startup configuration

With the Cisco Catalyst 9600 Series, the user can ignore the startup configuration in either the Cisco IOS XE CLI or ROMmon. (See Table 9.)

Table 9.   Ignoring the startup configuration

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Cisco IOS Software	Confreg 0x8000 or 0x0040	C9600(config)#system ignore startupconfig C9600(config)#no system ignore startupconfig
ROMmon	Confreg Use the interactive prompt to enable/disable ignore startup configuration	SWITCH_IGNORE_ STARTUP_ CFG=1

Contents

Operations

Interface reference

The Cisco Catalyst 6500/6800 Series has two levels of interface numbering:

interface <Type><Slot#>/<Port#>.

The 9600 Series has three levels:

interface <Type><Slot#>/<Bay#>/<Port#>.

As of Release 16.12.1, the bay number is unused and is always 0. For example, FortyGigabit Ethernet port 1 on slot 1 is referenced as Fo1/1 in the 6500/6800 Series and as Fo1/0/1 in the 9600 Series. See Table 10.

Table 10.   Interface numbering

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
TenGigabit Ethernet	TenGigabitEthernet1/1	Te1/0/1
FortyGigabit Ethernet	FortyGigabitEthernet5/1	Fo5/0/1

Contents

Management interface

The management interface on the Cisco Catalyst 9600 Series can be a copper Gigabit Ethernet or fiber 10 Gigabit Ethernet interface. The Cisco Catalyst 6500/6800 Supervisor Engine 6T provides copper or fiber Gigabit Ethernet. The management port on both platforms has its own VRF for separation of management traffic from normal data traffic. However, the name of the VRF for the management port is different between the 9600 Series and 6500/6800 Series. Note also that the names of the VRFs are case sensitive. Table 11 lists the management port differences between the two platforms.

Note that the management interface is out of band and will not offer the same features as a full-fledged front-panel interface (for example, security like MACsec, NetFlow operations, etc.).

Table 11.   Management interfaces and VRF names

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Interface	Mgmt0	GigabitEthernet0/0 TenGigabitEther0/1
VRF	management	Mgmt-vrf

Contents

Software features

For details on the software features supported on the Cisco Catalyst 9600 Series, please use the feature navigator on Cisco.com. Some of the features behave differently on the 9600 Series compared to the 6500/6800 Series. Following are some of these differences.

System Maximum Transmission Unit (MTU)

System MTU is a Layer 2 MTU. On the Cisco Catalyst 9600 Series, the global command "system mtu <1500-9216>" changes the Layer 2 MTU on all the interfaces within the system. There is no support for an interface-level command to set the MTU for individual interfaces as of Cisco IOS XE Software Release 16.12.1. On the Cisco Catalyst 6500/6800 Series, the global command "system jumbo <1500-9216>" sets the global baby giant MTU for all interfaces. The default system jumbo MTU is 9216. The 6500/6800 Series also supports a per-interface MTU. The per-interface MTU command takes precedence. See Table 12.

Table 12. System MTU

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
"system jumbomtu <>"	Changes Layer 2 MTU on all interfaces	N/A
"system mtu"	N/A	Changes Layer 2 MTU on all interfaces
System jumbomtu/ MTU value	1500 to 9216	1500 to 9216
Interface-level MTU (layer 2)	• Range "1500 to 9216" • Takes precedence over system MTU	• Range "1500 to 9216" • Takes precedence over system MTU

StackWise Virtual

The Cisco Catalyst 6500/6800 Series supports a Virtual Switching System (VSS) that combines two physical switches into a single logical switch. The equivalent functionality on the 9600 Series is Cisco StackWise® Virtual. StackWise Virtual and VSS have the same functionalities but different implementations.
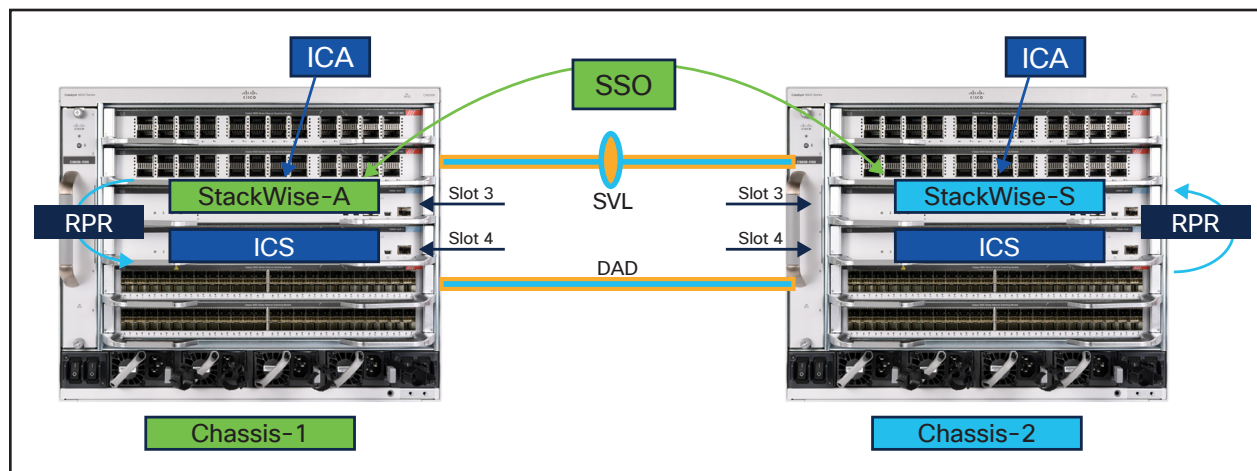

Figure 2.   StackWise Virtual

Cisco StackWise Virtual is a network system virtualization technology that pairs two Cisco Catalyst 9000 family switches into one virtual switch. Using Cisco StackWise Virtual to stack the switches in this way simplifies operational efficiency with a single control and management plane, scales system bandwidth with a distributed forwarding plane, and assists in building resilient networks using the recommended network design. Cisco StackWise Virtual allows two physical Cisco Catalyst 9600 Series Switches to operate as a single logical virtual switch using a 40G or 10G connection.

A hard requirement for the StackWise Virtual link, in cases of modular systems, is that both switches in the Cisco StackWise Virtual pair must be of the same switch model and version.

The bridge to possible

Contents

On the Cisco Catalyst 9600 Series Supervisor Engine 1 module, the stack can be formed over a 40G or 100G link to help ensure that the distribution or aggregation switches can be deployed over a long distance. Additionally, on the Cisco Catalyst 9600 Series Supervisor Engine 2 module (C9600X-SUP-2), the stack can be formed over a 400G link.

StackWise Virtual is supported on C9600X-SUP-2 starting with Cisco IOS XE Release 17.10.1 and onward.

Use the steps below to configure StackWise Virtual.

Step 1: Configure the StackWise Virtual domain.

SW-1	SW-2
9500-Dist-1(config)# **stackwise-virtual**	9500-Dist-2(config)# **stackwise-virtual**

Step 2: Configure the StackWise Virtual link.

SW-1	SW-2
9500-Dist-1(config)# **interface range FortyG x/y/z**	9500-Dist-2(config)# **interface range FortyG x/y/z**
9500-Dist-1(config-if)# **stackwise-virtual link <1 | 255>**	9500-Dist-2(config-if)# **stackwise-virtual link <1 | 255>**

Contents

Step 3: Specify dual-active detection (DAD).

SW-1	SW-2
9500-Dist-1(config)# **interface range TenG x/y/z**	9500-Dist-2(config)# **interface range TenG x/y/z**
9500-Dist-1(config-if)# **stackwise-virtual dual-active-detection**	9500-Dist-2(config-if)# **stackwise-virtual dual-active-detection**

Step 4: Save and reload to convert.

SW-1	SW-2
9500-Dist-1# **copy run start**	9500-Dist-2# **copy run start**
9500-Dist-1# **reload**	9500-Dist-2# **reload**

Note: On the Cisco Silicon One™ Q200-based Catalyst 9500X switches, the StackWise Virtual and DAD links can be dynamically edited without requiring a system reload.

The initial conversion from a standalone system to a Stackwise Virtual system still requires a reload.

For more information, refer to the StackWise Virtual configuration guide for the Cisco Catalyst 9600 Series:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst9600/software/release/17-14/configuration_guide/ha/b_1714_ha_9600_cg/configuring_cisco_stackwise_virtual.html.

Contents

Port-channel load balancing

The Cisco Catalyst 9600 Series provides additional combinations of different header fields as input for port-channel load balancing. Table 13 lists the supported load-balancing methods for the 6500/6800 and 9600 Series.

Table 13.   Port-channel load balancing

Cisco Catalyst 6500/6800 Series		Cisco Catalyst 9600 Series	
C6500E:		**C9600:**	
dst-ip	Dst IP Addr	dst-ip	Dst IP Addr
dst-mac	Dst Mac Addr	dst-mac	Dst Mac Addr
dst-mixed-ip-port Port	Dst IP Addr and TCP/UDP	dst-mixed-ip-port Port	Dst IP Addr and TCP/UDP
dst-port	Dst TCP/UDP Port	dst-port	Dst TCP/UDP Port
mpls packets	Load Balancing for MPLS		
		extended	Load Balance Methods
src-dst-ip	Src XOR Dst IP Addr	src-dst-ip	Src XOR Dst IP Addr

Contents

Cisco Catalyst 6500/6800 Series		Cisco Catalyst 9600 Series	
src-dst-mac	Src XOR Dst Mac Addr	src-dst-mac	Src XOR Dst Mac Addr
src-dst-mixed-ip-port TCP/UDP Port	Src XOR Dst IP Addr and	src-dst-mixed-ip-port TCP/UDP Port	Src XOR Dst IP Addr and
src-dst-port	Src XOR Dst TCP/UDP Port	src-dst-port	Src XOR Dst TCP/UDP Port
src-ip	Src IP Addr	src-ip	Src IP Addr
src-mac	Src Mac Addr	src-mac	Src Mac Addr
src-mixed-ip-port Port	Src IP Addr and TCP/UDP	src-mixed-ip-port Port	Src IP Addr and TCP/UDP
src-port	Src TCP/UDP Port	src-port	Src TCP/UDP Port
vlan-dst-ip	Vlan, Dst IP Addr	vlan-dst-ip	Vlan, Dst IP Addr
vlan-dst-mixed-ip-port TCP/UDP Port	Vlan, Dst IP Addr and	vlan-dst-mixed-ip-port TCP/UDP Port	Vlan, Dst IP Addr and
vlan-src-dst-ip	Vlan, Src XOR Dst IP Addr	vlan-src-dst-ip	Vlan, Src XOR Dst IP Addr

Contents

Cisco Catalyst 6500/6800 Series		Cisco Catalyst 9600 Series	
vlan-src-dst-mixedip-port and TCP/UDP Port	Vlan, Src XOR Dst IP Addr	vlan-src-dst-mixedip-port and TCP/UDP Port	Vlan, Src XOR Dst IP Addr
vlan-src-ip	Vlan, Src IP Addr	vlan-src-ip	Vlan, Src IP Addr
vlan-src-mixed-ip-port TCP/UDP Port	Vlan, Src IP Addr and	vlan-src-mixed-ip-port TCP/UDP Port	Vlan, Src IP Addr and

The "extended" option on the Cisco Catalyst 9600 Series provides additional combinations of different fields:

C96000(config)#port-channel load-balance extended ?

dst-ip	Dest IP
dst-mac	Dest MAC
dst-port	Dest Port
ipv6-label	IPV6 Flow Label
l3-proto	L3 Protocol
src-ip	Src IP
src-mac	Src MAC
src-port	Src Port
<cr>	<cr>

C96000(config)#

Contents

Port-channel services

Table 14. Port-channel services

Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
QoS policies are configured on the port-channel interfaces	QoS policies are configured on the individual port-channel member ports
Port ACLs (PACLs) are configured on the port-channel interfaces	PACLs are configured on the individual port-channel member ports

Virtualization: VNET

The Cisco Catalyst 6500 Series supports the Easy Virtual Network (EVN) feature, which uses VNET. This feature is no longer available on the Cisco Catalyst 9000 switch family. The alternative is to configure VRF-Lite with subinterfaces, which are supported on the 9600 Series switches.

Host tracking feature

The Cisco Catalyst 6500/6800 Series supports IP Device Tracking (IPDT) for keeping track of connected hosts (association of MAC and IP addresses). In the Cisco Catalyst 9600 Series with the latest Cisco IOS XE release, the new Switch Integrated Security Features (SISF)-based IPDT feature acts as a container policy that enables snooping and device-tracking features available with First Hop Security (FHS) in both IPv4 and IPv6, using IP-agnostic CLI)commands.

The command "device-tracking upgrade-cli" allows you to migrate the existing IPDT configuration to the new SISF-based device-tracking CLI commands.

See Appendix A for more information on migrating from the IPDT CLI configuration to the new SISF-based device-tracking CLI configuration.

Contents

Access control lists

Object group ACLs: Both the Cisco Catalyst 6500/6800 Series and 9600 Series support object group ACLs. There are, however, syntax differences. Table 15 shows some examples of object group ACLs with source and destination port groups.

Table 15.   Object group ACLs

Object group ACL features	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Object group ACL with source port group	object-group ip address g1 host 10.20.20.1 host 10.20.21.1 object-group ip port p1 gt 100 lt 200 ip access-list extended test1 permit tcp host 1.1.1.1 port-group p1 adrgroup g1	object-group network g1 host 10.20.20.1 host 10.20.21.1 object-group service p1 tcp source gt 100 tcp source lt 200 ip access-list extended test1 permit object-group p1 host 1.1.1.1 object-group g1

Contents

Object group ACL features	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Object group ACL with destination port group	object-group ip address g2 host 10.30.20.1 host 10.30.21.1 object-group ip port p2 gt 300 lt 400 ip access-list extended test2 permit tcp host 1.1.1.1 addrgroup g2 port-group p2	object-group network g2 host 10.30.20.1 host 10.30.21.1 object-group service p2 tcp gt 300 tcp lt 400 ip access-list extended test2 permit object-group p2 host 1.1.1.1 object-group g2

ılıılı
CISCO
The bridge to possible

Contents

Object group ACL features	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Object group ACL with source and destination port groups	object-group ip address g1 host 10.20.20.1 host 10.20.21.1 object-group ip port p1 gt 100 lt 200 object-group ip address g2 host 10.30.20.1 host 10.30.21.1 object-group ip port p2 gt 300 lt 400 ip access-list extended test3 permit tcp addrgroupt g1 portgroup p1 addrgroup g2 portgroup p2	object-group network g1 host 10.20.20.1 host 10.20.21.1 object-group service p3 tcp source gt 100 gt 300 tcp source gt 100 lt 400 tcp source lt 200 gt 300 tcp source lt 200 lt 400 object-group network g2 host 10.30.20.1 host 10.30.21.1 ip access-list extended test3 permit object-group p3 object-group g1 object-group g2

Object group ACL features	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Object group ACL with "established" keyword	object-group ip address g4 10.22.33.0 255.255.255.0 10.33.44.0 255.255.255.0 object-group ip port p4 eq 500 eq 600 ip access-list extended test4 permit tcp addrgroup g4 portgroup p4 10.30.40.0 0.0.0.255 established	object-group network g4 10.22.33.0 255.255.255.0 10.33.44.0 255.255.255.0 ip access-list extended test4 permit tcp object-group g4 eq 500 10.30.40.0 0.0.0.255 established permit tcp object-group g4 eq 600 10.30.40.0 0.0.0.255 established

Note: To view the expanded ACL for object group ACLs, enable "service internal" and use the command "show ip access-list <list_name> expand."

Access group mode: When PACLs are applied to the physical port, VLAN ACLs (VACLs) are applied to the VLAN, and Router ACLs (RACLs) are applied to the Switch Virtual Interface (SVI), the Cisco Catalyst 6000 Series offers options to merge all three ACLs with the option "merge mode" or to ignore VACLs and RACLs with the option "prefer port mode." With the 9600 Series, the ACLs will always be applied in the following order: PACL, VACL, and then RACL.

TCAM exhaustion

In the Cisco Catalyst 6000 Series, once the Access Control Entries (ACE) exceed the maximum free available Ternary Content Addressable Memory (TCAM) spaces, ACL reduction will occur, and any traffic hitting the overflowed ACL will be software switched.

In the Cisco Catalyst 9000 family, the overflowed ACL will not be programmed and will have a default of "deny all" for the overflowed ACL.

The bridge to possible

Contents

Flexible NetFlow

Both the Cisco Catalyst 9600 Series and 6500/6800 Series support Flexible NetFlow. Beside the scalability differences, there are a few configuration differences. They are listed in Table 16.

Table 16.   Flexible NetFlow differences

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Timestamp	Use system uptime	Use absolute time [0 is at time 00:00:00 January 1, 1970]
NetFlow on port channel	Configuration under port channel	Configuration under L3 port channel and member of L2/L3 port channel
Bridged traffic	Apply the flow monitor to the Layer 2 interface with keyword "layer2-switched"	Apply the flow monitor to a VLAN
NetFlow on tunnel	Supported	Not supported
NetFlow collect options: collect routing next-hop address ipv4 collect ipv4 source prefix collect ipv4 source mask collect ipv4 destination mask collect flow sampler	Supported	Supported

The NetFlow implementation on the Silicon One Q200-based Catalyst 9600X-SUP-2 switches is a software-based implementation. Details can be found in the dedicated NetFlow section later in this document.

Contents

Embedded Logic Analyzer Module (ELAM)

The Cisco Catalyst 6500/6800 Series supports the ELAM feature, which captures packets in real time on the switch without disruptions to performance. The more comprehensive feature on the Cisco Catalyst 9600 Series is Wireshark.

The following are the differences between the 6500/6800 Series ELAM and the 9600 Series Wireshark:

- The ELAM feature configures and displays commands through show commands only. In Wireshark capture, configuration is through the monitor-level CLI and display is though show commands.

- ELAM is capable of packet capture at the ASIC level, whereas packet capture occurs at the interface level with Wireshark.

- ELAM captures only the first packet that hits the switch that matches the configuration. Wireshark is capable of capturing packets over a duration.

- Wireshark can capture both data plane and control plane packets.

Switched Port Analyzer (SPAN) filter

Both the Cisco Catalyst 6500/6800 and 9600 Series support SPAN filters. The 6500/6800 Series supports a filter option of "good/bad," which isn't supported on the 9600 Series.

Quality of Service (QoS)

The ASICs that power the Cisco Catalyst 65006800 and 9600 Series are different, so there are some differences in QoS behaviors, as described below.

Hardware rate limiters

On the Cisco Catalyst 6500/6800 Series, the "platform rate-limit" command enables rate limiting in hardware. All rate limiting or policing is done in hardware with the 9600 Series, so this command is not needed.

The bridge to possible

Contents

Control plane policing

Control Plane Policing (CoPP) is enabled on the Cisco Catalyst 9600 Series, with default policing rates for different classes of traffic. These policing rates are optimized for a typical campus environment.

The policing rates can be changed or disabled to meet the requirements of different application environments. On the Cisco Catalyst 6500/6800 Series with Supervisor Engine 2T or 6T, CoPP is also enabled by default and can be disabled. The 6500/6800 Series also allows class maps under CoPP to be added, modified, or removed. Table 17 lists the differences between the two platforms.

Table 17.   CoPP differences

	Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
Default	Enabled (can be disabled)	Enabled (can't be disabled, but policing rates can be modified)
CoPP class map	Can be added, modified, or removed	System predefined
Policing rate	Can be modified	Can be modified

Table 18 lists the buffer differences between the Cisco Catalyst 9600 Supervisor Engine 1, Supervisor Engine 1 and the Cisco Catalyst 6500 Supervisor Engine 6T.

The bridge to possible

Contents

Table 18. Buffers

	Cisco Catalyst 6500 Supervisor Engine 2T	Cisco Catalyst 9600 Supervisor Engine 1	Cisco Catalyst 9600 Supervisor Engine 2
Buffer	Varies depending on the line card	3x36 MB	80 MB (Shared Memory System) + 8 GB (High Bandwidth Memory)
Buffer Sharing	Buffers are dedicated per port	Buffer sharing is within the ASIC; there are 3 ASICs in Supervisor Engine 1	Buffer sharing is within the ASIC.

Cisco Catalyst 6000 Series platform-specific commands

Table 19 lists commands that are specific to the Cisco Catalyst 6000 Series and are not available on the 9600 Series.

Table 19. Cisco Catalyst 6500/6800 Series platform-specific commands

Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
mls <...>	Not applicable; the 9600 Series provides hardware-enabled feature by default
Auto qos default	Auto qos global compact
diagnostic fpga soft-error recover conservative	Not applicable

The bridge to possible

Contents

Cisco Catalyst 6500/6800 Series	Cisco Catalyst 9600 Series
ntp update-calendar	clock calendar-valid
ip device tracking	Please see Appendix A
Platform ip cef load-sharing full	Not applicable
Flow hardware usage notify <...>	Not applicable
Flow hardware usage notify <...>	Not applicable
Vlan internal allocation policy asending	Not applicable
Vlan access-log ratelimit <...>	Not applicable
Ip domain-name	Ip domain name
Ip domain-lookup	Ip domain lookup

Contents

Appendix A. IPDT/SISF

If your device has no legacy IP device tracking or IPv6 snooping configurations, you can use only the new SISF-based device-tracking commands for all your future configurations. The legacy IPDT commands and IPv6 snooping commands are not available. For details on SISF configuration, please refer to the configuration guide.

IPDT, IPv6 snooping, and device-tracking CLI compatibility

Table 20 displays the new SISF-based device-tracking commands and the corresponding IPDT and IPv6 snooping commands.

Table 20.   Device-tracking and corresponding IPDT and IPv6 snooping commands

IP Device Tracking (IPDT)	IPv6 snooping	SISF-based device tracking
IP device tracking probe count	Not supported	Not supported
IP device tracking probe delay	IPv6 neighbor binding reachable-lifetime	Device-tracking policy reachable-lifetime
IP device tracking probe interval	IPv6 snooping tracking retry-interval	Device-tracking policy retry-interval
IP device tracking probe use-svi	Accepted and interpreted as IP device tracking probe auto-source override	Accepted and interpreted as IP device tracking auto-source override
IP device tracking probe auto-source fallback	Not supported	Not supported

The bridge to possible

Contents

IP Device Tracking (IPDT)	IPv6 snooping	SISF-based device tracking
IP device tracking probe auto-source override	Not supported	Not supported
IP device tracking tracebuffer	Not supported	Not supported
IP device tracking maximum	IPv6 snooping policy <name> limit	Device-tracking snooping policy <name> limit
IP device tracking probe count	Not supported	Not supported
IP device tracking probe interval	Not supported	Not supported
Clear ip device tracking all	Not supported	Not supported

Appendix B. Unsupported features

The Cisco Silicon One Q200 ASIC brings a host of new capabilities and enhancements to enterprise networks. However, from a feature support perspective, certain features are yet to be ported over from the UADP (and older) platforms.

Note: While most of these features are expected to be supported eventually, there are no defined timelines for adding this support.

If any of these features are critical to your network operation, it might be prudent to consider whether the UADP-based Catalyst core 9000 switches can satisfy your network requirements. The UADP-based Catalyst 9500 high-performance series and 9600 Supervisor Engine 1 will be supported for the near future, and there are no plans to announce EoS/EoL of these devices soon.

ı|ı.ı|ı.
CISCO
The bridge to possible

Contents

The following are key features not supported as of IOS XE Release 17.12.1:

· **Multicast:** PIM Dense mode, PIM Bidir, PIM Snooping, and MVPN

· **MPLS:** VPLS, EoMPLS over tunnels/port channels, and Traffic Engineering FRR

· **Layer 3 forwarding:** NAT/PAT, many advanced PBR features, and WCCP

· **Layer 2 forwarding:** REP, private VLANs, FlexLinks, and Bonjour support

· **Security:** PACL and VACL, FQDN and reflexive ACLs, and RadSec

Note that this is not an exhaustive list. Please validate with the **Cisco Feature Navigator** for support information on specific features. When in doubt, please check with your accounts team to determine whether your existing switch configuration (Catalyst 6000 or UADP-based Catalyst 9000) can be ported to a Silicon One Q200-based switch without any issues with feature compatibility and scale limits.

Best practices and caveats while migrating

Transitioning to the Cisco Silicon One Q200 represents a significant leap forward in network infrastructure capabilities. The following sections delve into the best practices and key "gotchas" to be aware of during the upgrade process. The objective is to focus on the key features and configurations that are known to require some input during the update process.

We will go over each feature in a separate section below.

· Quality of service (QoS)

· Access control lists (ACLs)

· Flexible NetFlow (FNF)

Contents

Quality of Service (QoS)

In a first for campus switches, the Cisco Silicon One Q200-based switching platforms use a virtual output queueing (VOQ) model. The new ASIC comes with a new forwarding model, with QoS closely tied into the capabilities defined by the ASIC, requiring a redesign of the old configuration.

The major difference between the QoS models lies in how the traffic is differentiated and prioritized internally within the switch. The UADP-based Catalyst 9000 switches, as well as older-generation Catalyst 2000, 3000, 4000, and 6000 level platforms, all differentiate traffic on the basis of predefined fields available in the IP or the Layer 2 packet header (for example, using a differentiated services code point [DSCP] tag in the Layer 3 packet header). If we cannot or don't want to use those details in the packet header, we leverage other tools like ACLs or VLANs to first "tag" the packet at ingress using a qos-group and then calling the same qos-group in the egress.

The Silicon One Q200 QoS implementation is different in that we do not use the predefined field already present in the header, and instead we append two tags to the packet when it enters the switch and then remove them when the packet exits the switch. The two tags in question are:
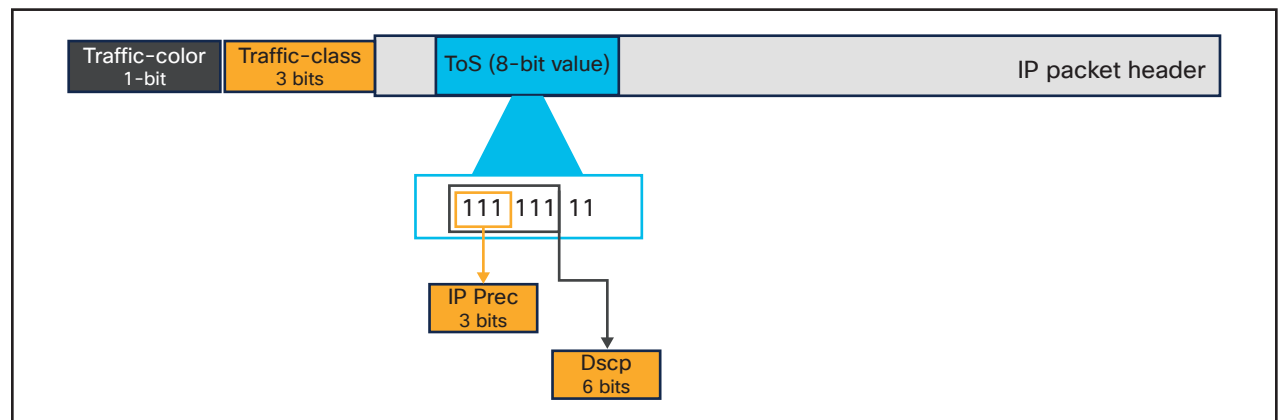
· Traffic-class

· Traffic-color

Figure 3.   Tags appended to the header on ingress

The bridge to possible

Contents

Traffic-class

Traffic-class is a 3-bit field used by the switch to determine the traffic priority. The values range from 0 to 7. The higher the traffic-class, the higher the priority of the traffic. There are two hard predefined settings within the switch that cannot be changed or disabled:

· Traffic-class 7 is a strict priority queue with priority level 1 (the highest priority).

· Traffic-class 0 is the class default, and any untagged traffic automatically gets sent here. This queue is always a normal queue and cannot be configured as a strict priority queue.

At ingress, the incoming traffic must be classified (using DSCP, class of service [CoS], IP precedence, ACLs, or VLANs) into the corresponding traffic classes. Any traffic not matching these classifications will be assigned to traffic- class 0.

At egress, we perform queueing and shaping operations based on the traffic-class.

Note: Directly configuring DSCP, CoS, IP precedence, ACLs, or VLANs at the egress for queueing will not work, as the default DSCP/CoS to traffic-class mapping will be used instead.

Traffic-color

Traffic-color is a 1-bit field (with a value of either 0 or 1) that is used by congestion management algorithms (weighted tail drop [WTD] and weighted random early detection [WRED]) to prioritize traffic within a queue in the event of congestion. All traffic, without exception, is colored green (value of 0) by default. Any traffic-color change to yellow (value of 1) must be explicitly performed via configuration at ingress. The traffic-color can be set using the command below.

```
Set discard-class <0-1>
```

Note: Trying to perform any congestion management on the egress policy will not work unless the traffic is assigned a color at egress. Congestion management based on packet tags like DSCP or CoS is not supported.

The bridge to possible

Contents

Configuration migration considerations

Let's examine a 2P2Q configuration using the UADP ASIC as a base. An example configuration would
look something like the following:

```
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 af31  af32  af33


policy-map CAMPUS _ EGRESS _ POLICY

 class VOICE-PQ1

    priority level 1 percent 10

 class VIDEO-PQ2

    priority level 2 percent 20

class USER-TRAFFIC

    bandwidth remaining percent 65

    random-detect dscp-based

    random-detect dscp af31 70 100

    random-detect dscp af32 60 80

    random-detect dscp af33 50 70
```

The bridge to possible

Contents

```
class class-default

    bandwidth remaining percent 35


Interface gig1/0/1

 service-policy output CAMPUS _ EGRESS _ POLICY
```

Let's ignore the WRED configuration for now and just focus on the queue configurations, translating those to an equivalent Silicon One configuration. There are some considerations here:

- Traffic-class 7 is a strict priority queue set to level 1. This cannot be changed or disabled.

- The UADP performs the queueing based on the DSCP tags here. On the Silicon One, we must map the DSCP tags to the corresponding traffic-class at the ingress and map the traffic-class to the queueing operations at the egress.

This requires two separate policies, one at the ingress and one at the egress. Again, ignoring the WRED configurations for now, the "translated" configurations would look like this.

Starting with the INGRESS policy-map (mapping the DSCP to the traffic-class):

```
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 af31 af32 af33


policy-map CAMPUS _ INGRESS _ POLICY
```

Contents

```
  class VOICE-PQ1

     set traffic-class 7

  class VIDEO-PQ2

     set traffic class 6

class USER-TRAFFIC

     set traffic class 5


Interface gig1/0/1

  service-policy input CAMPUS _ INGRESS _ POLICY
```

Next the EGRESS policy-map (queueing based on the traffic-class):

```
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
```

Contents

```
class tc6

    priority level 2

class tc5

    bandwidth remaining percent 65

class class-default

    bandwidth remaining percent 35
```

Interface gig1/0/3

 service-policy output CAMPUS _ EGRESS _ POLICY

The next step would be to migrate the WRED configurations as well. This requires some additional configurations, as the Silicon One Q200 does WRED based on the traffic-color and not based on DSCP. Again, there are some considerations here:

· Coloring of the packet must be done at INGRESS.

· Traffic-color can either be 0 (green) or 1 (yellow). So we can have a total of two thresholds vs. the three thresholds (one each for af31, af32, and af33) we configured on the UADP switch.

Editing the configuration with these considerations, we get the following "translated" configuration:

Starting with the INGRESS policy-map (mapping the DSCP to the traffic-class and coloring the traffic):

```
class-map match-any VOICE-PQ1

 match dscp ef

class-map match-any VIDEO-PQ2
```

The bridge to possible

Contents

```
 match dscp 24
class-map match-any USER-TRAFFIC-GREEN
 match dscp af31
class-map match-any USER-TRAFFIC-YELLOW
 match dscp af32 af33


policy-map CAMPUS _ INGRESS _ POLICY
 class VOICE-PQ1
    set traffic-class 7
 class VIDEO-PQ2
    set traffic class 6
class USER-TRAFFIC-GREEN
    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
```

Guide Cisco public

CISCO
The bridge to possible

Contents

Next the EGRESS policy-map (queueing based on the traffic-class and performing WRED based on traffic-color):

```
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 percent 65
    random-detect discard-class-based
    random-detect discard-class 0 percent 80 100
    random-detect discard-class 1 percent 50 70
```

Contents

```
class class-default

    bandwidth remaining percent 35


Interface gig1/0/1

 service-policy output CAMPUS _ EGRESS _ POLICY
```

For more details on QoS architecture and configuration, please refer to the QoS white paper for the Silicon One Q200 platforms. You can find the link in the references section later in this document.

Access control lists (ACLs)

From a security perspective, ACLs provide a granular method for network administrators to control traffic going into and out of specific nodes in the network. The ACLs are stored in a specialized type of memory called the TCAM (ternary content-addressable memory). This allows for rapid searches for specific fields. However, the space available on the TCAM is finite, and storing a large number of entries requires optimization to make the best use of the space available.

The Silicon One Q200 has a TCAM size of 8000 entries on the default SDM template. As of Release 17.12.1, there is no support for increasing this space using customizable SDM templates. In comparison, the UADP has a total space of 54,000 entries using customizable SDM templates.

Another consideration is the total number of unique named ACLs (regardless of if it's an IP ACL, object group ACL [OG-ACL], or security group ACL [SG-ACL]) in egress and ingress. We can have a total of:

· 126 unique named ACLs (labels) applied in the ingress direction

· 16 unique named ACLs (labels) applied in the egress direction

Contents

Object group ACLs (OG-ACLs)

For traditional ACLs, we can leverage the 8000 entries in the TCAM space. However, for larger entries, migrating from a traditional IP ACL to an OG-ACL or SG-ACL would be more beneficial.

OG-ACLs leverage object groups (a collection of similar "objects" like network addresses, port numbers, and protocols) and perform permit or deny operations on them. The advantage of using OG-ACLs and SG-ACLs is that the content in the object group or security group is expanded to a different field (the CEM, or Central Exact Match), and only the permit and deny statements are programmed into the TCAM space. The CEM space is leveraged to learn the IPv4 and IPv6 routes, and in terms of scale, the Silicon One Q200 supports a sizable number of routes, up to 2 million IPv4 or 1 million IPv6 routes.

Let's look at an example to understand this better. Below is an example of an ACL calling on an object group.

```
ip access-list extended SampleACL

 permit ip object-group TestGroup any

 deny ip any any (implicit)

!
```

Expanding the object group shows up that it is composed of 10,000 unique entries.

```
object-group network TestGroup

 host 10.5.64.140

 host 10.5.64.86

 host 10.5.75.124

 10.10.170.32  255.255.255.224

 10.179.252.0  255.255.252.0
```

Contents

```
10.5.200.240 255.255.255.240

10.53.199.0 255.255.255.0

10.56.100.224 255.255.255.240

... Truncated (10,000 entries)

host 10.56.223.56
```

When we apply this OG-ACL to an interface on a UADP-based Catalyst 9000 switch, the contents of the OG-ACL are expanded and programmed to the TCAM space (up to 54,000 entries).

However, when we use the same ACL on a Catalyst 9000 switch based on the Silicon One Q200, the contents of the object group are expanded to the CEM table, and only the permit and deny (implicit) is programmed in the TCAM space. So, in terms of utilization,

- 10,000 entries are consumed from the 2 million CEM field and

- 2 entries are consumed from the 8000 TCAM field

In this way, we can program large ACLs into the Silicon One Q200-based core Catalyst 9000 switch by leveraging OG-ACLs or SG-ACLs.

Configuration migration considerations

If migrating from a Catalyst 6000 or 9000 switch, we need to consider the following:

- The total number of access control entries (ACEs) in IP ACLs. If fewer than 8000, we can program it as is. But if more than 8000, conversion to OG-ACL or SG-ACL is mandated.

- The number of ACLs applied in the ingress and egress directions. While ingress has a good number of labels (126 labels), it's easy to exhaust the egress labels. So if the source configuration has many egress ACLs, it is necessary to redesign the ACL to be applied in the ingress direction.

The bridge to possible

Contents

Flexible NetFlow (FNF)

The Silicon One Q200 supports a software-based Flexible NetFlow that uses flows to provide statistics for accounting, network monitoring, and network planning.

A flow is a unidirectional stream of packets that arrives on a source interface and has the same values for the keys. A key is an identified value for a field within the packet. You create a flow using a flow record to define the unique keys for your flow.

NetFlow on the Silicon One Q200

The following are the key differences in the NetFlow implementation on the Silicon One Q200-based Catalyst 9000 switches compared to the UADP-based Catalyst 9000 switches as well as the Catalyst 6000 Series switches:

- The NetFlow is software based – in other words, the packets are sent to the CPU for parsing and there is no dedicated hardware in the ASIC.

- The NetFlow is a sampled implementation – in other words, not every packet across the system is inspected; inspection occurs for only one packet for every N packets, where N is the sample size configured.

The Catalyst 9500X and Catalyst 9600 Supervisor Engine 2 have an 8-core CPU powering the system. One of these eight cores is explicitly reserved for NetFlow operations, along with a separate IOS XE processing thread. Having a dedicated CPU gives us the following benefits.

- Dedicated CPU resources for NetFlow operations: There is no CoPP or other process that can hinder and block NetFlow operations or packets.

- Isolated CPU for NetFlow: Even if NetFlow causes the CPU to spike up, normal management operations on the switch will continue as expected, as only the NetFlow core and operations will be impacted.

Example: For N = 1000, one packet is selected at random and sent to the CPU for parsing. Information from the header of the packet is inspected and keys are updated and exported out.

Contents

The supported sampler method is true random sampling. This means that packets are selected at random based on the sample size regardless of flow.

Note: This means that, due to the nature of randomness, all the flows are not captured immediately on enabling NetFlow, as packets belonging to that flow might simply not have been sent to the CPU. Over a period, due to the nature of randomness, the exporter will get details on the flow.

Sample size

The recommended sampler size is 1000. This recommendation is based on the typical load one can expect from a switch which is handling hundreds of gigabytes of traffic at any given time. If more aggressive samplers are required, the load on the CPU will increase correspondingly as more packets are sent to the CPU for parsing.

Aggressive samplers can be used when the rate of traffic is lower than the line rate supported on our devices. We recommend keeping an eye on the CPU load while tweaking the sampler, as configuring very aggressive values can overwhelm the dedicated CPU, causing NetFlow operations to break.

We support a sampler of up to 1:2, as in 1 in every 2 packets are parsed to the CPU for processing.

Transport security (encryption)

In a hybrid network, transport security is a critical piece to ensure end-to-end security. The Catalyst 9500X switches support WAN-MACsec in addition to MACsec. These help ensure that traffic is encrypted regardless of the destination of the packet. Within the same network, MACsec helps with the encryption. As the packet leaves the network, we have both WAN-MACsec and IPsec for encryption across the WAN boundary. WAN-MACsec will operate at Layer 2 and at line speed, making it a more appealing solution.

Contents

In terms of support:

IPsec is supported only on the Catalyst 9500X-60L4D. None of the other products, including the Catalyst 9500X-28C8D, support IPsec. At the time of writing this document, on the Catalyst 9600 with Supervisor Engine 2, the C9600X-LC-56YL4C has the capability to support IPsec, but the support will be added in a later IOS XE release.

WAN-MACsec is supported on all Silicon One Q200-based devices. (On the Catalyst 9600 Supervisor Engine 2, we require generation 2 line cards.)

Note that for both WAN-MACsec and IPsec, an HSEC k9 license is required to be installed on the system.

Links and references

- [**UADP QoS white paper**](#)

- [**Silicon One Q200 white paper**](#)

- [**Catalyst 9600 Series data sheet**](#)

- [**Catalyst 9600 Architecture white paper**](#)

Conclusion

The Cisco Catalyst 9600 Series Switches are Cisco's leading modular enterprise switching core and distribution platforms. They are the new generation of the core and distribution platform and provide many additional capabilities, making them well suited for enterprises looking to migrate from their existing Cisco Catalyst 6500/6800 Series deployment.