Cisco Catalyst 9600 Series Architecture

Contents

Introduction

Enterprise campus networks are undergoing profound changes to support ever-increasing bandwidth demands on the access layer while moving toward supporting Wi-Fi 6 and the rapid growth of powerful endpoints. With access layer bandwidth moving from speeds of 1G to 2.5G, 5G, and 10G, higher bandwidths such as 25G, 40G, and 100G will become the de facto speeds in distribution and core layers to maintain a similar over subscription ratio.

The Cisco® Catalyst® 9600 Series Switches are the next generation of the industry-leading business-critical modular enterprise campus core and distribution platform. The 9606R chassis is hardware-ready to support a switching capacity of up to 25.6 Tbps. The 9600 Series switches support granular port densities that fit diverse campus needs, including nonblocking 40G and 100G (QSFP28); 1G, 10G, and 25G (SFP28); and 10M, 100M, 1G, 2.5G, 5G, and 10G (RJ-45). It is architected to support all the latest Cisco optics innovations, such as dual-rate 40G/100G and 10G/25G optics. The platform delivers high availability with field-replaceable dual supervisors, redundant power supplies, and fan tray. The platform is campus-optimized with an innovative dual-serviceable fan tray design and side-to-side airflow and is closet friendly with a depth of approximately16 inches.

This white paper provides an architectural overview of the Cisco Catalyst 9600 Series chassis, including system design, power, cooling, and storage options.

Platform overview

The Cisco Catalyst 9600 Series platform is a modular switch based on the Cisco Unified Access® Data Plane (UADP) 3.0 ASIC, which not only protects your investment but also allows a larger scale and higher throughput than the UADP 2.0 (Figure 1). The platform runs on the modern modular, open Cisco IOS® XE operating system, which supports model-driven programmability, has the capacity to host containers with support for up to 1 TB of solid-state (SSD) storage, and can run third-party applications and scripts natively within the switch (by virtue of the x86 CPU architecture, local storage, and a higher memory footprint).

The modern operating system offers enhanced high availability features such as In-Service Software Upgrade (ISSU), Software Maintenance Upgrades (SMU), Graceful Insertion and Removal (GIR) and Cisco StackWise® Virtual technology. Improved high availability is also added via Platinum-efficient, redundant power supplies as well as variable-speed, highly efficient redundant fans.

Figure 1.
Cisco Catalyst 9600 Series

Chassis: The Cisco Catalyst 9606R is a 6-slot chassis. Two middle slots (slots 3 and 4) 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 (3.2 Tbps for transmitting and 3.2 Tbps for receiving) from each of the supervisor slots to each of the line card slots. This means the system can provide 32 ports of 100G at line rate for each line card slot. The 9606R chassis can provide a maximum of 128x 100G/40G (QSFP) ports or 192x 25G/10G/5G/2.5G/1G/100M/10M ports. The backplane of the chassis is passive, which brings the following benefits:

- Lower power consumption
- Higher Mean Time Between Failures (MTBF)
- Fan tray, power supplies, and line cards are field replaceable and can be replaced nondisruptively

Supervisors: The Cisco Catalyst 9600 Series Supervisor Engine 1 (Sup-1) is powered with three UADP 3.0 Application-Specific Integrated Circuits (ASICs) (Figure 2). Each ASIC is capable of 3.2 Tbps (1.6 Tbps full duplex) switching capacity and 1 Bpps of forwarding performance. Together, the three UADP 3.0 ASICs with Sup-1 provide 9.6 Tbps of switching capacity and 3 Bpps of forwarding performance. There are no uplinks on the Sup-1 as the ASIC connections are dedicated to the line cards.

Figure 2.
Cisco Catalyst 9600 Series Sup-1

Line cards: Cisco Catalyst 9600 Series Switches offer the ability to mix and match a range of line cards to support different core and aggregation deployments (Figure 3).

Fiber line cards:

- C9600-LC-24C: 24-port 40G (QSFP+), 12-port 100G (QSFP28) line card
- C9600-LC-48YL: 48-port 25G/10G/1G SFP28/SFP+/SFP line card

Copper line card:

- C9600-LC-48TX: 48-port 10G/5G/2.5G/1G/100M/10M RJ-45 line card

Figure 3.
Line cards

Chassis overview

This section briefly describes the highlights of the Cisco Catalyst 9600 Series chassis.

Table 1 provides information about the capabilities of the chassis.

Table 1. Chassis specifications

Cisco Catalyst 9606R	
Supervisor slots	2 (slots 3 and 4)
Line card slots	4 (slots 1 and 2 and slots 5 and 6)
Port density	128x QSFP28 (100G), QSFP+ (40G) (48x/96x with Sup-1) 192x SFP28 (25G), SFP+ (10G), SFP (1G) 192x RJ-45 (10G/5G/2.5G/1G/100M/10M)
Dimensions (HxWxD)	13.95 in. x 17.4 in. x 16.1 in. (8RU)
Bandwidth per LC slot	6.4 Tbps
Bandwidth between supervisor slots	400G
Power supplies	4 (N+1 and Combined mode)
Cooling	Side to side

The power supplies are "Platinum efficient" (90% or higher efficiency).

An ACT2 chip for module authenticity is supported on all supervisors, line cards, and fan trays.

Fan tray with N+1 redundant fans and flexible options to service the fan from the front or back (Figure 4).

Figure 4.
Flexible fan-tray servicing

Chassis power

The Cisco Catalyst 9600 Series uses a modular design for power. The 9606R has four slots for power supplies. Each power supply is very compact but highly efficient. The system provides support for both Combined and N+1 redundant mode.

By default, the system operates in Combined mode. In this mode, all power supplies are active and sharing the load. In N+1 redundant mode, one of the power supplies is configured as the standby power supply.

To enable a diverse range of deployments, the Cisco Catalyst 9600 Series also supports combinations of AC and DC units. When combining AC and DC power supplies, both types of power supplies need to have the same power output level.

The Cisco power calculator (https://cpc.cloudapps.cisco.com/cpc/launch.jsp) can help you determine the power supplies required for a given configuration. The tool also provides heat dissipation information.

Power supply unit

The maximum output power per Power Supply Unit (PSU) for the Cisco Catalyst 9600 Series is listed below, and each PSU has a power holdup time of approximately 20 milliseconds at 100 percent load. Each PSU comes with front-to-back variable-speed cooling fans and has a push-release lock for simple and secure online insertion and removal (Figure 5).

- 2000W AC PS with 240V input (1050W with 120V input; 10.5A input)
- 2000W DC PS with 48V input (50A input)

Figure 5.
PSUs

Chassis cooling

The Cisco Catalyst 9600 Series Switches support a hot-swappable and field-replaceable fan tray that can be replaced from the front or back, which offers significant flexibility with different cable management options. The chassis supports side-to-side airflow. The fan unit is responsible for cooling the entire chassis and for interfacing with environmental monitors to trigger alarms when conditions exceed thresholds. The fan modules contain thermal sensors to detect ambient temperature and adjust the fan speed. The chassis supports a hardware failure of up to one individual fan, and if a fan fails the remaining fans will automatically increase their rpm to compensate and maintain sufficient cooling. If the switch fails to meet the minimum number of required fans, it shuts down automatically to prevent the system from overheating.

Cisco Catalyst 9600 Series chassis are equipped with onboard thermal sensors to monitor the ambient temperature at various points and report thermal events to the system so that it can adjust the fan speeds.

Chassis airflow

The Cisco Catalyst 9600 Series fan tray supports side-to-side airflow for the modules and front-to-back airflow for the power supplies (Figure 6). The system can also be made into front to back airflow with an optional airflow kits - sold separately.

Figure 6.
Chassis airflow

Architecture

The Cisco Catalyst 9600 Series Switches are based on a centralized architecture (Figure 7). All forwarding, security, and queueing are done on the supervisor, while the line cards are considered transparent, containing only PHYs and control logic. Each line card slot has up to a 6.4-Tbps connection to each of the supervisor slots. The simplicity of this centralized design allows easy upgrade of features as well as additional bandwidth by just upgrading 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 the In-Service Software Upgrade feature.

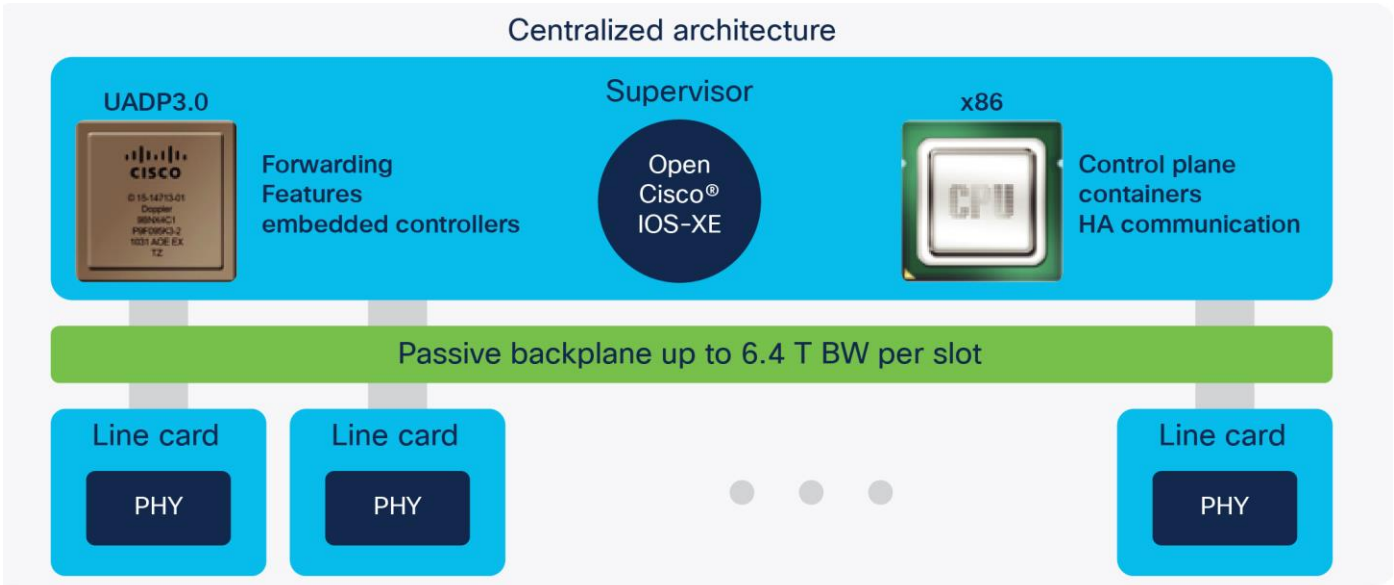

Figure 7.
Cisco Catalyst 9600 Series architecture

The majority of the Cisco Catalyst 9600 Series components (chassis, supervisors, line cards, and fan tray) come with a built-in passive RFID for inventory management. All those components also have Blue Beacon, which can be turned on and off by the software (Figure 8). This helps locate the components when they need to be serviced. The power supply doesn't have passive RFID or Blue Beacon.

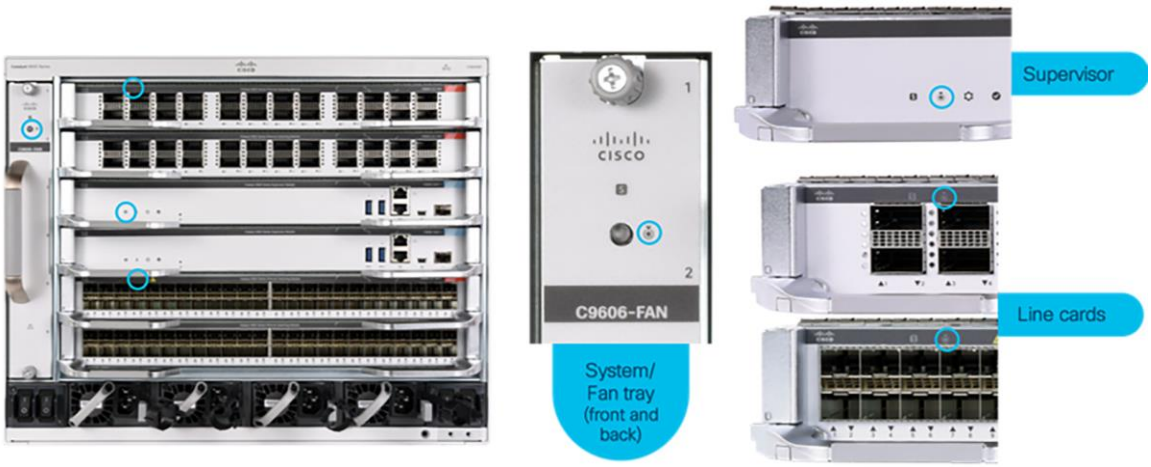

Figure 8.
Blue Beacon location on Cisco Catalyst 9600 Series components

Supervisor

The Cisco Catalyst 9600 Series Supervisor Engine 1 is powered by three UADP 3.0 ASICs and one x86 CPU processor (Figure 9). The three ASICs are interconnected with a 3.2-Tbps ASIC interconnect on each ASIC. Sup-1 provides 9.6 Tbps (4.8 Tbps full duplex). With the Cisco Catalyst 9606R chassis, each slot has 2.4 Tbps with Sup-1. Due to high-performance line card requirements, the Sup-1 module doesn't have any dedicated uplink ports (any port on any line card can be used as an uplink).

The supervisor architecture consists of the following main components

- UADP ASIC
- X86 CPU complex
- ASIC interconnect
- Memory

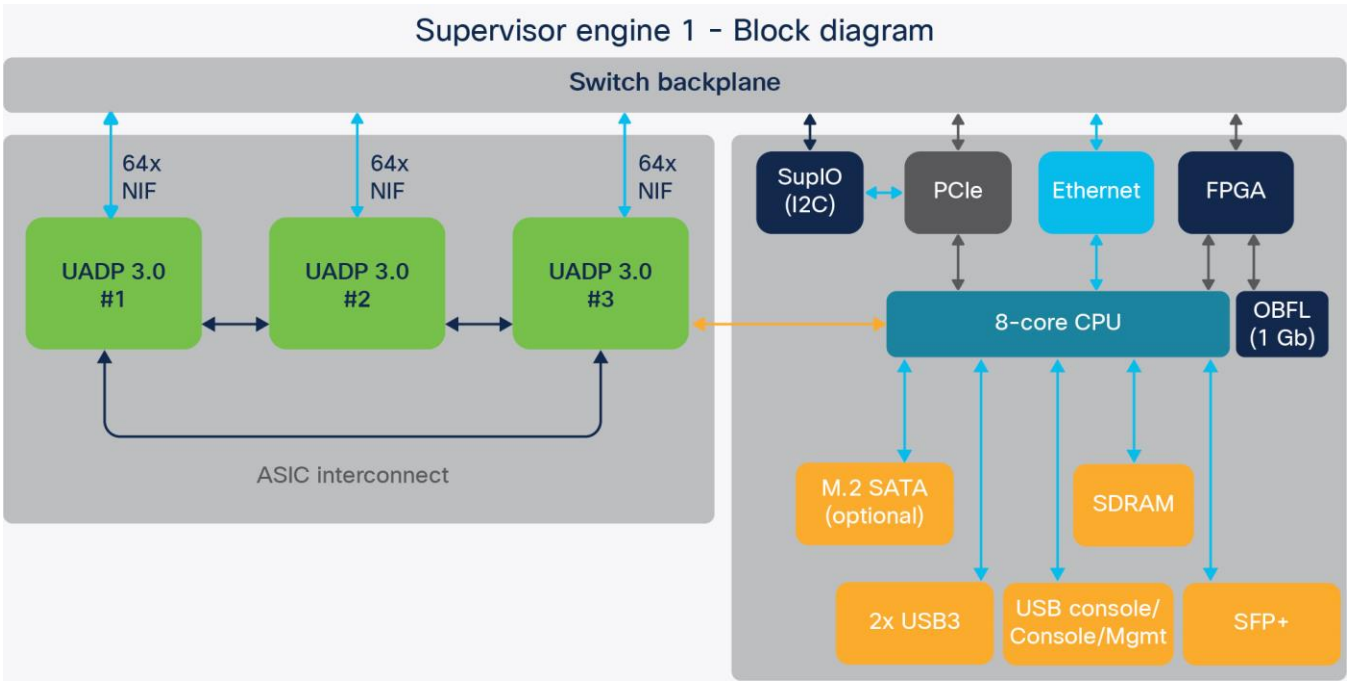

Figure 9.
Sup-1 diagram

UADP ASIC

The supervisors are built on the UADP 3.0 ASIC, which is based on a system-on-chip architecture (SOC) (Figure 10).

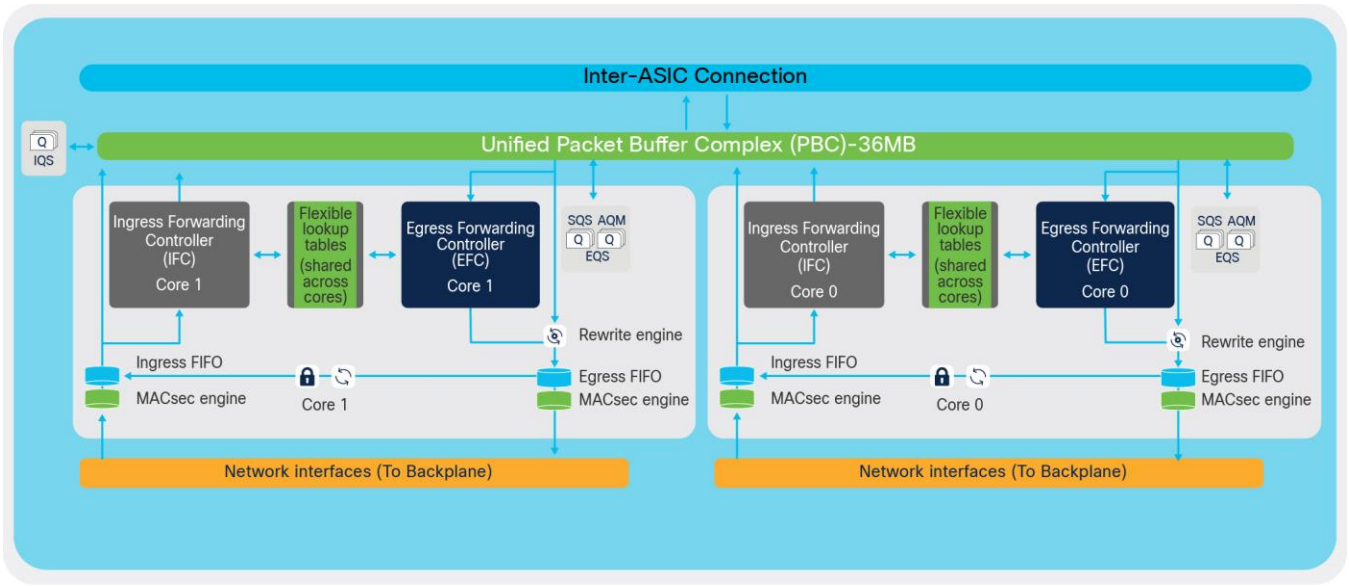

Figure 10.
UADP 3.0 ASIC diagram

The UADP 3.0 ASIC is the latest generation in the UADP family. It is built with 16-nanometer technology that offers significantly larger tables and bandwidth compared to the other UADP ASICs. The UADP 3.0 continues to offer programmable pipelines and flexible allocation of hardware resources for different needs in different places in the network.

The following are the key UADP 3.0 capabilities.

- Packet bandwidth/switching throughput (full duplex): 1.6 Tbps (800 Gbps per core)
- Forwarding performance: 1 Bpps (500 Mpps per core)
- ASIC interconnects: 3.2 Tbps (1.6 Tbps full duplex)
- Forwarding Information Base (FIB) table: 416,000 (double-width entries optimized for IPv6 deployments)
- Unified packet buffer: 36 MB
- NetFlow: IPv4 and IPv6 parity with double-width shared tables
- ACL TCAM: 54,000 entries

x86 CPU complex

As with other products in the Cisco Catalyst 9000 family, the 9600 Series uses the x86 CPU. The CPU complex has the following highlights:

- 2.0 GHz x86 8-core CPU
- 16-GB DDR4 RAM
- 16-GB internal enhanced USB flash
- M.2 SATA internal storage (up to 960 GB)
- Console supports mini USB and RJ-45 connectivity
- Supports two USB 3.0 ports
- Management port supports RJ-45 (1G) and SFP/SFP+ (1G and 10G) (only one can be active)
- System reset switch for manually resetting the supervisor

ASIC interconnect

The Cisco Catalyst 9600 Series supervisors are built with three UADP 3.0 ASICs (Figure 11). Communication within a core or between cores is locally switched within the ASIC, meaning that packets destined to local ports within the ASIC do not use the ASIC interconnect link. The purpose of the ASIC interconnect is to move data between multiple UADP ASICs.

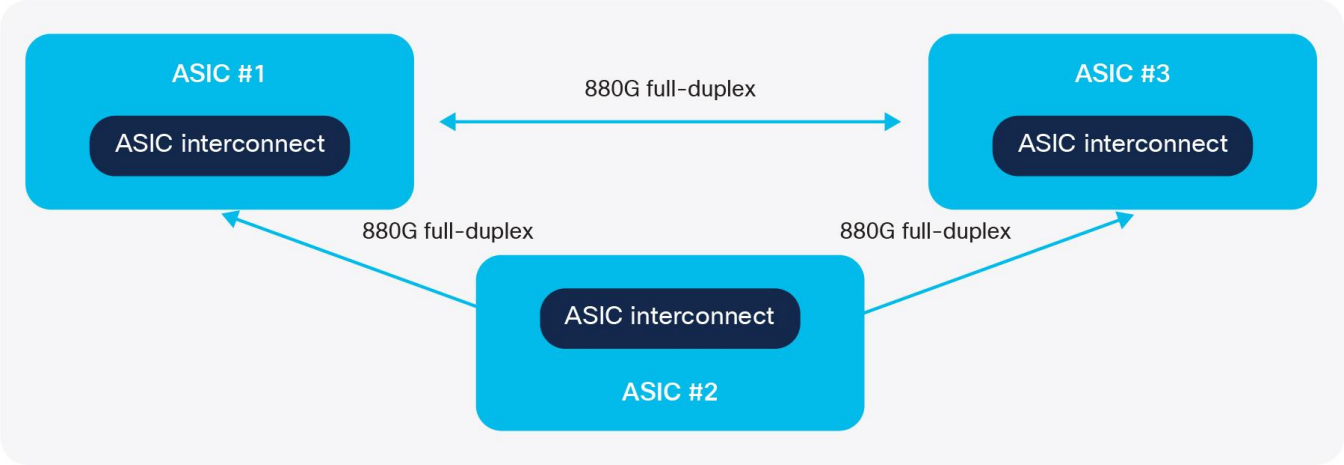

Figure 11.
ASIC interconnect diagram

External storage

Cisco Catalyst 9600 Series Switches provide two types of external storage:

- USB 3.0 SSD on the front-panel of the supervisor
- M2 SATA that can be plugged into the removable supervisor (up to 1 TB)

This external storage can be used for general-purpose storage for packet capture, operation system trace logs, and graceful insertion and removal (GIR) snapshots. Most importantly, the M2 SATA SSD can be used for application hosting. An application hosted on a network device can serve a variety of purposes, ranging from automation, configuration management monitoring, and integration with existing tool chains.

Internal flash storage cannot be used to store third-party applications, as it is not supposed to be formatted as an EXT2 or EXT4 file system. But the M2 SATA SSD can support an EXT2 or EXT4 (default) file system and application hosting. It also has the ability to monitor the health of the SSD storage through S.M.A.R.T.

Supervisor and line card connections

Cisco Catalyst 9600 Series chassis line card slots have dedicated connections to both supervisor slots. Once the line cards are up and running, all traffic entering the line cards is sent to both the active and hot standby supervisors. The hot standby supervisor processes those packets just like the active supervisor does, and the resulting packets are sent to the egress line cards. The egress line cards select the packets from the active supervisor and send them out of the front panel ports (Figure 12). When there is a switchover between the supervisors, the PHYs in the line cards just need to switch the connection to the new active supervisor. As a result, the outage during this event is very minimal. This capability, together with the centralized architecture, enables the Cisco Catalyst 9600 Series to provide uninterrupted In-Service Software Upgrades.

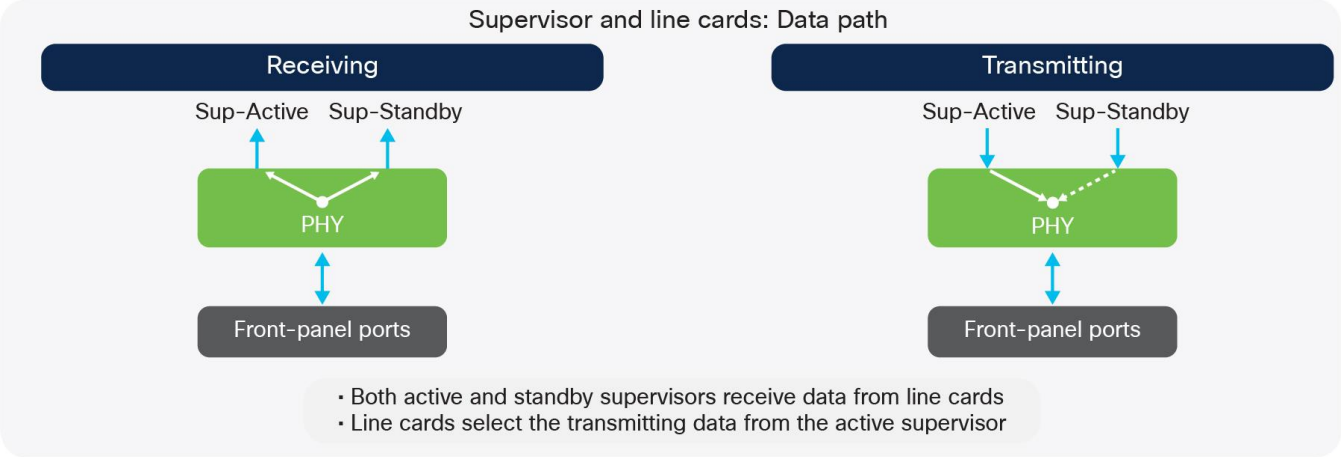

Figure 12.
Supervisors and line card connections

Line cards

The Ethernet PHY (physical layer) connects a link layer device (often a MAC) to a physical medium such as a transceiver. The PHY on the Cisco Catalyst 9600 Series Switches is a fully integrated Ethernet transceiver supporting steering and mapping of lanes back to the ASIC to enable multiple speeds depending on the optics inserted on the front panel ports or on whether copper ports are present.

C9600-LC-24C

Figure 13 shows the architecture of the C9600-LC-24C line card (40G/100G).

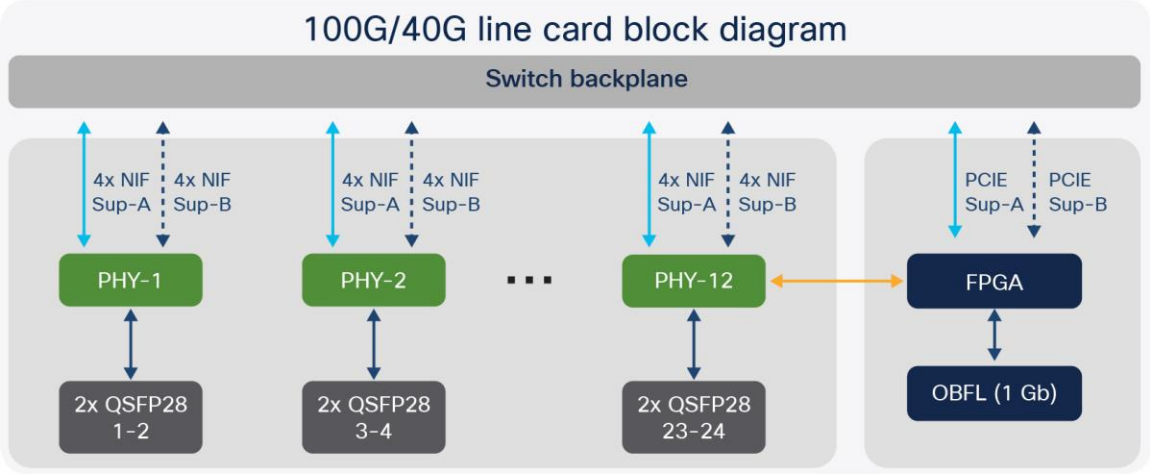

Figure 13.
Diagram for the C9600-LC-24C line card

- Up to 24 ports of 40G nonblocking
  - This is the default mode.
- Up to 12 ports of 100G nonblocking (upper ports)
  - The upper ports can be enabled for 100G. When a 100G port is enabled, the subsequent port will be disabled. (Example: If port 1 is enabled as 100G, port 2 will be disabled.)
- The line card will be capable of supporting all 24 ports at 100G nonblocking with a future supervisor

C9600-LC-48YL

Figure 14 shows the architecture of the C9600-LC-48YL line card.

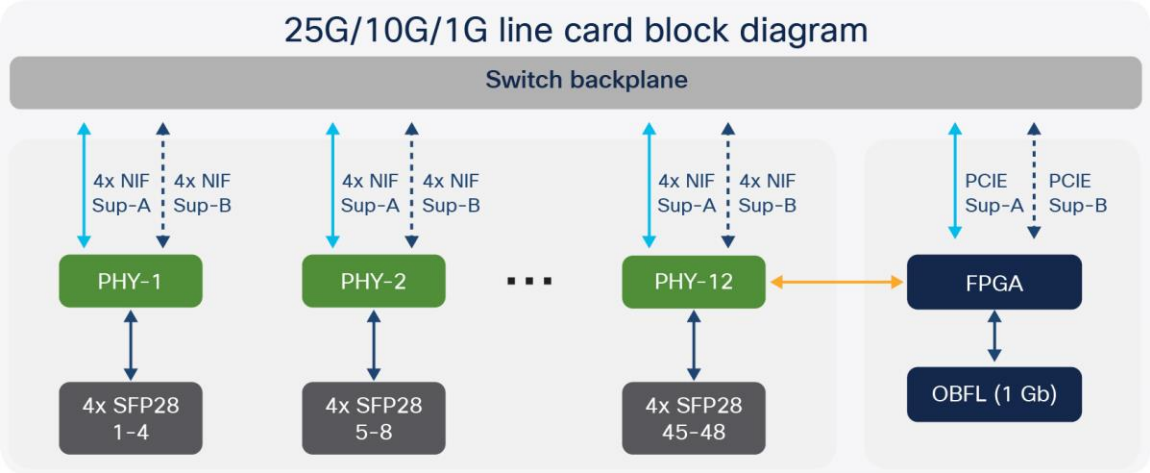

Figure 14.
Diagram for C9600-LC-48YL line card

- Up to 48 ports of 1G, 10G, or 25G nonblocking
- Speed is auto-negotiated depending on the inserted optics
- Supports 10G/25G dual-rate optics, optimized for campus distances (SFP-10G/25G-CSR-S supports 300 or 400 meters with OM3/4 cables at both10G and 25G)
- The line card will be capable of supporting all 48 ports at 50G nonblocking with a future supervisor

C9600-LC-48TX

Figure 15 shows the architecture of the C9600-LC-48TX line card.

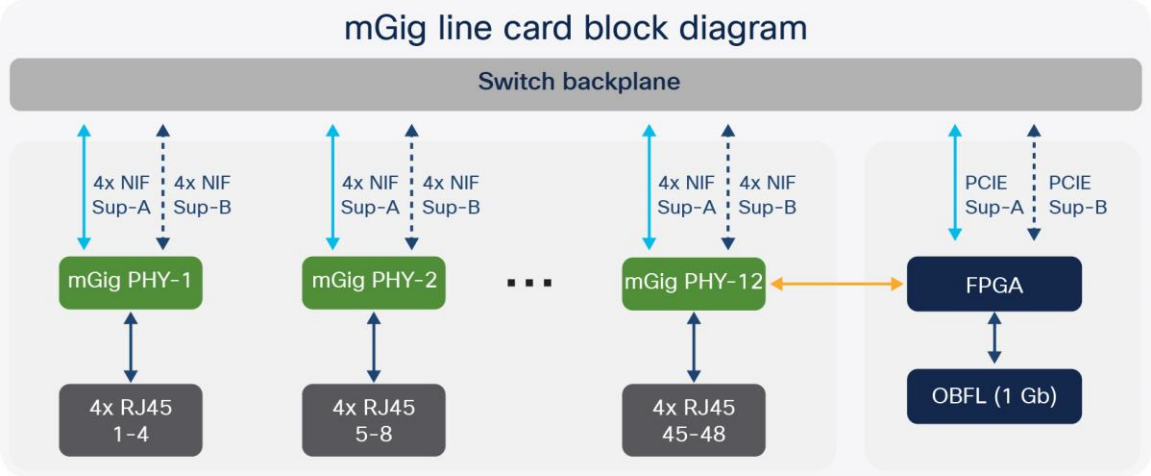

Figure 15.
Diagram for C9600-LC-48TX line card

- Up to 48 ports of 10M, 100M, 1G, 2.5G, 5G, or 10G nonblocking
- All ports are Multigigabit / 802.3bz
- No Power over Ethernet (PoE) on these ports

Mapping of front panel interfaces to ASICs

Figure 16 shows the mapping of the front panel ports to the ASICs. Spreading the ports within a port channel will maximize the utilization of hardware resources (buffer, forwarding, etc.) from each ASIC.

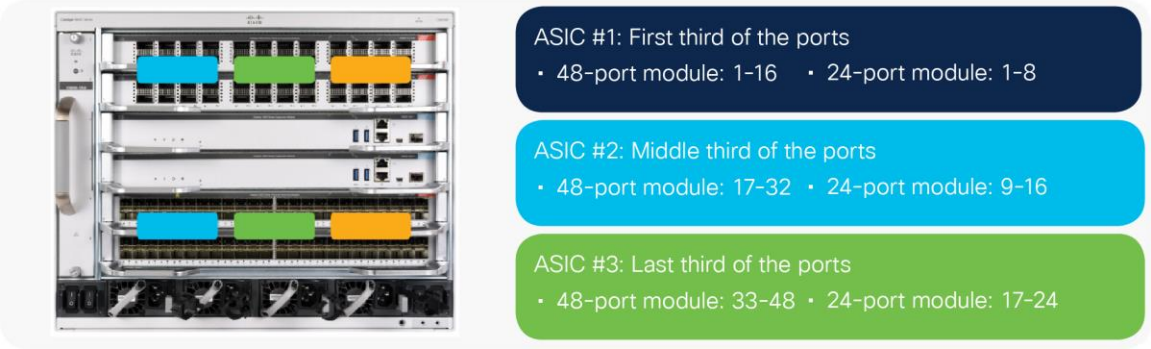

Figure 16.
Diagram for front panel interface to ASICs mapping

Packet walks

This section provides a high-level overview of how packet forwarding is performed on the Cisco Catalyst 9600 Series Switches.

Unicast forwarding within the ASIC

Figure 17 shows the basic sequence of events when packets enter the Cisco Catalyst 9600 Series front panel ports for unicast packet forwarding within the ASIC

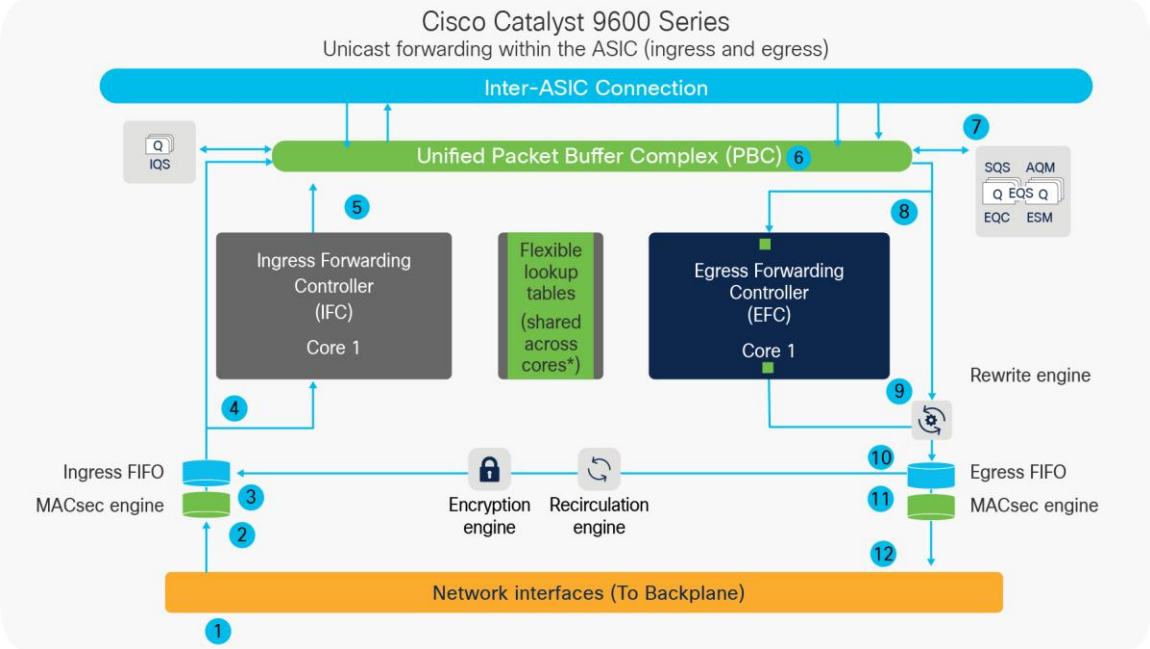

Figure 17.
Unicast packet walk within a single ASIC core

1. Packet arrives at the line card's ingress port; PHY converts the signal and serializes the bits, and then sends the packet to the Network Interface (NIF) that goes to the backplane.
2. The packet travels through the backplane and enters the NIF of one of the ASICs.
3. The NIF passes the packet to the ingress MACsec engine. The MACsec engine will decrypt the packet if needed. The decryption is done at line rate. The packet now enters the Ingress First In First Out (FIFO).
4. The Ingress FIFO sends the packet to both the Ingress Forwarding Controller (IFC) and the Packet Buffer Complex (PBC) in parallel.
5. The IFC performs Layer 2, Layer 3, Access Control List (ACL), and Quality-of-Service (QoS) lookups and more, then returns the forwarding result (frame descriptor header) to the PBC.
6. The PBC uses the frame descriptor to determine the egress port. As the egress port is on the same ASIC, the result is sent to the Egress Queueing System (EQS) on the same ASIC.
7. The EQS receives the notification from the PBC and schedules the packet to be sent for egress processing.

8. The EQS signals the PBC to send the packet and descriptor out to both the Egress Forwarding Controller (EFC) and the Rewrite Engine (RWE).
9. The EFC completes egress functions and sends the final rewrite descriptor to the RWE.
10. The RWE performs packet rewrite with the final descriptor and sends the packet to the Egress FIFO.
11. The Egress FIFO sends the packet to the Egress MACsec.
12. The Egress MACsec performs a wire-rate encryption if required and then passes the frame on to the NIF. The packet then goes through the backplane and is sent out from one of the line card ports.

Unicast forwarding across ASICs

Figure 18 shows the basic sequence of events when unicast packets enter the Cisco Catalyst 9600 Series front panel ports and are sent across the ASIC interconnect link.

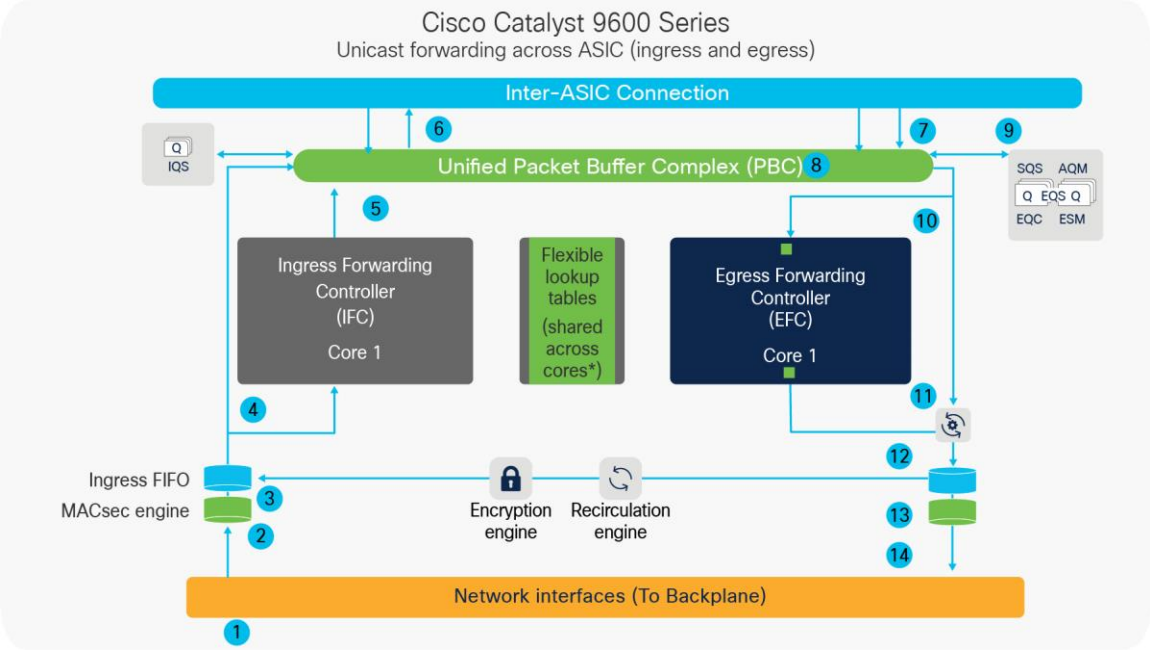

Figure 18.
Unicast packet walk across ASICs

1. Packet arrives at the line card's ingress port; PHY converts the signal and serializes the bits and then sends the packet to the network interface (NIF) that goes to the backplane.
2. The packet travels through the backplane and enters the NIF of one of the ASICs.
3. The NIF passes the packet to the ingress MACsec engine. The MACsec engine will decrypt the packet if needed. The decryption is done at line rate. The packet now enters the Ingress FIFO.
4. The Ingress FIFO sends the packet to both the Ingress Forwarding Controller (IFC) and the Packet Buffer Complex (PBC) in parallel.
5. The IFC performs Layer 2, Layer 3, ACL, and QoS lookups and more to return the forwarding result (frame descriptor header) to the PBC.

6. The PBC uses the frame descriptor to determine the egress port. As the egress port is on a different ASIC, the Ingress Queueing System (IQS) schedules the packet to be sent to the destination ASIC using an inter-ASIC connection.
7. The PBC on the destination ASIC receives the packet from the source ASIC via the inter-ASIC connection.
8. The PBC sends the frame descriptor to the EQS.
9. The EQS receives the notification from the PBC and schedules the packet to be sent for egress processing.
10. The EQS signals the PBC to send the packet and descriptor out to both the Egress Forwarding Controller (EFC) and the Rewrite Engine (RWE).
11. The EFC completes egress functions and sends the final rewrite descriptor to the RWE.
12. The RWE performs packet rewrite with the final descriptor and sends the packet to the Egress FIFO.
13. The Egress FIFO sends the packet to the Egress MACsec.
14. The Egress MACsec performs a wire-rate encryption if required and then passes the frame on to the NIF. The packet then goes through the backplane and is sent out from one of the line card ports.

Multicast forwarding

Figure 19 shows the basic sequence of events when packets enter the Cisco Catalyst 9600 Series front panel ports for multicast packet forwarding within the ASIC.

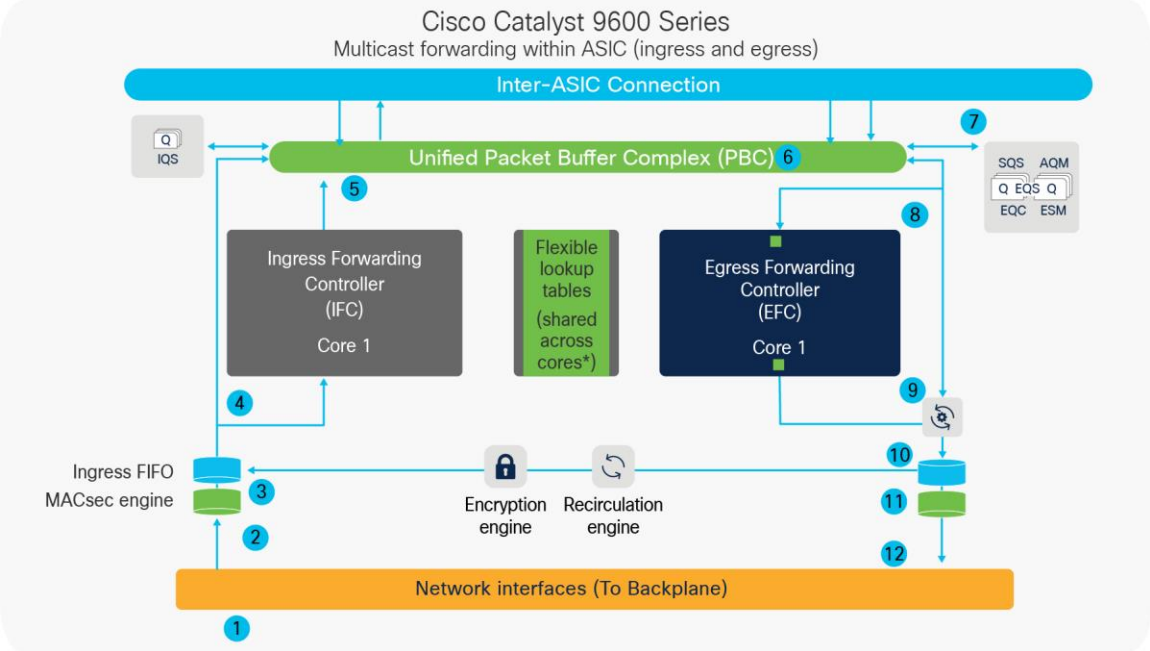

Figure 19.
Unicast packet walk

1. Packet arrives at the line card's ingress port; PHY converts the signal and serializes the bits and then sends the packet to the Network Interface (NIF) that goes to the backplane.
2. The packet travels through the backplane and enters the NIF of one of the ASICs.
3. The NIF passes the packet to the ingress MACsec engine. The MACsec engine will decrypt the packet if needed. The decryption is done at line rate. The packet now enters the Ingress FIFO.
4. The Ingress FIFO sends the packet to both the Ingress Forwarding Controller (IFC) and the Packet Buffer Complex (PBC) in parallel.
5. The IFC performs Layer 2, Layer 3, ACL, and QoS lookups and more, then returns the forwarding result (frame descriptor header) to the PBC. The frame descriptor in this case is a pointer to the replication table.
6. The PBC uses the frame descriptor to determine the egress port. (If there are receivers on other ASICs, the IQS will schedule the packet for the destination ASICs via the inter-ASIC connection.) For the local receivers, the result is sent to the Egress Queueing System (EQS).
7. The EQS receives the notification from the PBC. Based on the result, Active Queue Management (AQM) generates a list of egress ports and schedules the packet for each of those egress ports. The following steps are repeated for each of the egress ports in that list.
8. The EQS signals the PBC to send the packet and descriptor out to both the Egress Forwarding Controller (EFC) and the Rewrite Engine (RWE).
9. The EFC completes the egress functions and sends the final rewrite descriptor to the RWE.
10. The RWE performs packet rewrite with the final descriptor and sends the packet to the Egress FIFO.
11. The Egress FIFO sends the packet to the Egress MACsec.
12. The Egress MACsec performs a wire-rate encryption if required and then passes the frame on to the NIF. The packet then goes through the backplane and is sent out from one of the line card ports.

Conclusion

Cisco Catalyst 9600 Series Switches are enterprise-class core and distribution switches in the Cisco Catalyst 9000 family, offering a comprehensive portfolio and architectural flexibility with interface speeds from 10M to 100G. This platform is based on Cisco's next-generation programmable UADP ASIC for increased bandwidth, scale, security, and telemetry. The platform also supports infrastructure investment protection with nondisruptive migration from 10G to 25G and beyond. The Cisco Catalyst 9600 Series is built on a modular system architecture designed to provide high performance to meet the evolving needs of highly scalable and growing enterprise networks.

References

The following websites offer more details on the Cisco Catalyst 9600 Series and its capabilities.

Cisco Catalyst 9600 Series Switches Data Sheet

Cisco Catalyst 9600 Supervisor Engine-1 Data Sheet

Cisco Catalyst 9600 Series Line Cards Data Sheet

Cisco Catalyst 9600 Series Switches Hardware Installation Guide

25GE and 100GE – Enabling Higher Speeds in Enterprise with Investment Protection White Paper

Cisco Catalyst 9000 Switching Platforms: QoS and Queuing White Paper

Cisco Catalyst 9000 Switching Platforms: StackWise Virtual White Paper

Cisco Catalyst 9000 - Switching for a New Era of Intent-based Networking

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA C11-743272-00 02/20