Cisco N9300 Series Smart Switches NX-OS Feature Configuration Guide, Release 10.6(1s)

First Published: 2025-09-15

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

CHAPTER 1

New and Changed Information

Table 1: New and Changed Features

Feature	Description	Changed in Release	Where Documented
Support for interface features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.	Added support for these features. • Q-in-Q VLAN tunnels • UDLD and Port profiles • Virtual port channel (VPC) • Port channel and LACP • Single Hop BFD • Native VLAN traffic and Tagging, and Reserved VLAN • Layer 2 Access and trunk ports, Switch virtual interface (SVI) and VLAN logical interfaces	10.6(1s)	Interface and Layer 2 features, on page 5 Interface and Layer 2 feature guidelines, on page 6
Support for Layer 2 switching features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	Added support for these Layer 2 switching features. • VLAN • Spanning Tree Protocol (STP) • Rapid PVST+ • MST • MAC learning (global)	10.6(1s)	Layer 2 switching features, on page 9 Layer 2 switching feature guidelines, on page 10

Feature	Description	Changed in Release	Where Documented
Support for label switching features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	Added support for Segment Routing Layer 3 VPN feature with SR-MPLS underlay on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.	10.6(1s)	Label switching features, on page 11 Label switching feature guidelines, on page 11
Support for multicast routing features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	Added support for these multicast routing features. • IPv4 Layer 2 and Layer 3 multicast • Multicast consistency checker • Layer 3 Physical Interface, port channel, sub-interface, and SVI • Layer 2 port channel • Multicast flow counter	10.6(1s)	Multicast routing features, on page 13 Multicast routing feature guidelines, on page 14
Support for QoS features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	Added support for these QoS features. • QoS classification policies for system QoS • Queuing and Scheduling • QoS statistics	10.6(1s)	Quality of Service features, on page 15 QoS guidelines, on page 15
Support for security features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	Added support for these security features. • IP ACL • DHCP relay • Custom CoPP • MACsec	10.6(1s)	Security features, on page 17 Security feature guidelines, on page 18

Feature	Description	Changed in Release	Where Documented
Support for system management features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	Added support for these system management features. • Sampled Flow (sFlow) • Switched Port Analyzer (SPAN) • Encapsulated Remote SPAN (ERSPAN) • Link Layer Discovery Protocol (LLDP)	10.6(1s)	System management features, on page 21 System management feature guidelines, on page 22
Support for unicast routing features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	Added support for Unicast routing features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.	10.6(1s)	Unicast routing features, on page 25 Unicast routing feature guidelines, on page 28
Support for VXLAN features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	Added support for VXLAN features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.	10.6(1s)	VXLAN features, on page 31 VXLAN feature guidelines , on page 34
Support for common features on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches	Added support for these common features. • Power-On Auto Provisioning (POAP) • Smart Licensing Using Policy (SLP) • Telemetry	10.6(1s)	Commonly supported features, on page 39

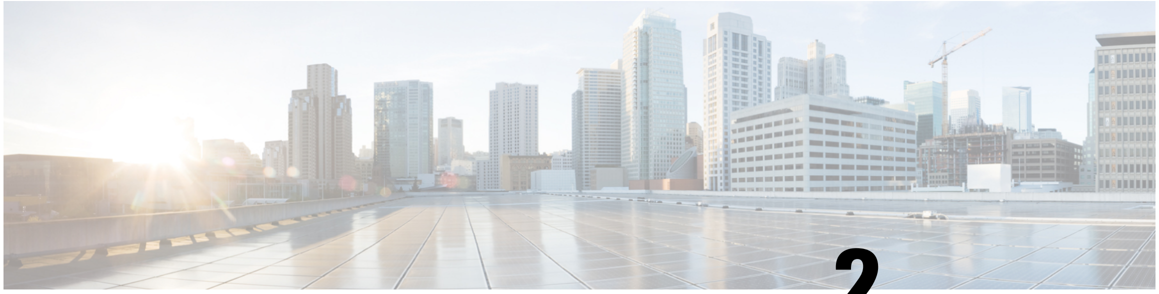

CHAPTER 2

Interface and Layer 2 features

Starting with Cisco NX-OS Release 10.6(1s), you can configure these interface and Layer 2 features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

Q-in-Q VLAN tunnels

A Q-in-Q VLAN tunnel enables a service provider to segregate the traffic of different customers in their infrastructure, while still giving the customer a full range of VLANs for their internal use by adding a second 802.1Q tag to an already tagged frame.

For information on configuring Q-in-Q VLAN tunnels, see Configuring Q-in-Q VLAN Tunnels.

Unidirectional Link Detection (UDLD)

Unidirectional Link Detection (UDLD) is a Cisco-proprietary protocol that enables devices connected by fiber-optic or copper Ethernet cables to detect and disable unidirectional links by periodically exchanging UDLD frames, thereby preventing network issues caused by misconnected or faulty cabling.

For information on UDLD, see Unidirectional Link Detection Parameter.

Port profiles

You can create a port profile that contains many interface commands and apply that port profile to a range of interfaces.

For information on port profiles, see Port Profiles

Layer 2 access and trunk ports

A Layer 2 port can be configured as an access port, which carries traffic for a single VLAN, or as a trunk port, which carries traffic for two or more VLANs simultaneously.

For information on access and trunk ports, see About Access and Trunk Interfaces.

Native VLAN ID and Tagging Native VLAN Traffic

The native VLAN ID on a trunk port is the VLAN that carries untagged traffic, meaning all untagged packets received on the trunk port are assigned to this VLAN, while the port can simultaneously carry both untagged and 802.1Q tagged packets.

Cisco software supports the IEEE 802.1Q standard on trunk ports, allowing untagged traffic to pass through a designated native VLAN, and provides an option to retain or strip 802.1Q tags on native VLAN packets globally on all trunk ports, ensuring flexible handling of tagged and untagged traffic.

For more information, see Configuring Layer 2 Interfaces

Virtual port channel (vPC)

A virtual port channel (vPC) allows links connected to two devices to function as a single logical port channel to a third device, enabling Layer 2 multipathing for redundancy, higher bandwidth, and load balancing.

For information on VPC, see Configuring vPCs.

LACP and Port channels

A port channel is an aggregation of multiple physical interfaces that creates a logical interface.

Link Aggregation Control Protocol (LACP) for Ethernet is defined in IEEE 802.1AX and IEEE 802.3ad. This protocol controls how physical ports are bundled together to form one logical channel.

For more information on LACP and Port channels, see Configuring Port Channels

Single Hop BFD

Single-hop Bidirectional Forwarding Detection (BFD) is a protocol that provides rapid detection of failures in the path between two directly connected network devices, enabling fast convergence and improved network reliability.

For more information, see Configuring Bidirectional Forwarding Detection.

- Interface and Layer 2 feature guidelines, on page 6

Interface and Layer 2 feature guidelines

Q-in-Q limitations

Beginning with Cisco NX-OS Release 10.6(1s), you can configure Q-in-Q on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

These limitations apply to Q-in-Q.

- You cannot configure a range of allowed VLANs by using **switchport trunk allowed vlan** *vlan_list* command.

  ```
  ..!
  interface Ethernet1/1  switchport mode trunk
  switchport vlan mapping all dot1q-tunnel 30
  switchport trunk allowed vlan 30-40
   ..!
  ```

  In the configuration example, trunk VLAN 30 is the provider VLAN. The VLANs 31 through 40 filter regular trunk traffic; these VLANs operate in sparse mode.

- You cannot use VLAN ACL with Q-in-Q.

- Multicast is not supported, and IGMP snooping is not supported.

- Custom EtherType is not supported.
- Variations of QinQ are not supported:
  - Q-in-VNI and Selective Q-in-VNI are not supported.
  - Selective Q-in-Q is not supported.
- You do not need the **system dot1q-tunnel transit** command for Q-in-Q tunneling when the switch acts as a transit device.

CHAPTER 3

Layer 2 switching features

Starting with Cisco NX-OS Release 10.6(1s), you can configure these Layer 2 switching features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

VLAN

You can use VLANs to divide the network into separate logical areas at the Layer 2 level. VLANs can also be considered as broadcast domains.

For more information, see Configuring VLANs.

Spanning Tree Protocol (STP)

STP is a Layer 2 link-management protocol that provides path redundancy while preventing loops in the network.

For more information, see STP.

Rapid PVST+

Rapid PVST+ is the default spanning tree mode for the software and is enabled by default on the default VLAN and all newly created VLANs.

For more information, see Rapid PVST+.

MST

MST is the IEEE 802.1 standard protocol that allows you to assign two or more VLANs to a spanning tree instance. MST maps multiple VLANs into a spanning tree instance, with each instance having a spanning tree topology independent of other spanning tree instances.

For more information, see Configuring MST Using Cisco NX-OS

MAC learning (global)

MAC learning is the process by which a device dynamically discovers and records the source MAC addresses of frames received on its interfaces, enabling efficient Layer 2 forwarding decisions.

For more information, see Configuring Layer 2 Switching.

Layer 2 switching feature guidelines

These features are not supported in Cisco NX-OS Release 10.6(1s) on the Cisco N9300 Series smart switches.

- Static MAC address
- Disable MAC learning at Layer 2 interfaces or per VLAN
- Flex links
- VTP
- Private VLAN
- Storm traffic control
- Reflective relay

CHAPTER 4

Label switching features

Label switching features are a suite of networking technologies, most notably Segment Routing Multiprotocol Label Switching (SR-MPLS), that forward traffic based on a short, pre-assigned label rather than performing a complex network-layer address lookup at every hop. This approach enables high-speed packet forwarding and supports advanced services such as Traffic Engineering, Quality of Service, and the creation of Virtual Private Networks (VPNs).

Starting with Cisco NX-OS Release 10.6(1s), you can configure these label switching features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

- Segment Routing Layer 3 VPN feature with SR-MPLS underlay

Segment Routing Layer 3 VPN feature with SR-MPLS underlay

Segment Routing Layer 3 VPN with an SR-MPLS underlay is a network service that creates isolated routing domains for multiple tenants by using BGP to distribute VPN routes while leveraging a simplified and scalable Segment Routing-based MPLS core for packet transport, eliminating the need for traditional protocols like LDP.

For more information, see Configuring Segment Routing.

Label switching feature guidelines

This section outlines feature support, guidelines, and limitations for label switching functionalities on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

Table 2: Supported features and releases

Features	Release
Segment Routing Layer 3 VPN feature with SR-MPLS underlay	10.6(1s)

Supported Segment Routing Layer 3 VPN feature with SR-MPLS underlay features:

- The chassis can be positioned as either a LEAF or a SPINE switch in the SR-MPLS fabric.
- SR-MPLS underlay is supported with BGP-LU, OSPF, and ISIS underlay protocols.

- L3VPN and L3 EVPN overlays over SR-MPLS underlay are supported using eBGP.
- The implementation supports Node-SID, Prefix-SID, and Adj-SID.
- SR-MPLS features are supported on L3 physical, L3 sub-interface, L3 Port-channel (PO), and L3 PO sub-interface types.
- Hierarchical ECMP (Level-1 and Level-2) is supported for SR-MPLS paths.
- Per-VRF VPN label encapsulation is supported.
- MPLS Decap Statistics are supported for VPN label termination.
- SVI (Switched Virtual Interface) is supported as an MPLS interface type.

Segment Routing Layer 3 VPN feature with SR-MPLS underlay limitations:

- MPLS TTL propagation operates in Uniform Mode.
- DSCP-EXP handling is Uniform during encapsulation and Pipe during decapsulation (tentative).
- Default load-sharing for SR-MPLS traffic is based on Label and IP (up to 5-tuple).

CHAPTER 5

Multicast routing features

Multicast routing features are a comprehensive suite of protocols and mechanisms designed to efficiently forward IP packets from a single source to a group of interested receivers. Unlike unicast (one-to-one) or broadcast (one-to-all), multicast (one-to-many) optimizes network resources by creating a targeted delivery path.

Their primary goal is to conserve network bandwidth and reduce host processing overhead by building an optimal **distribution tree** that delivers traffic only to the network segments where there are active listeners, rather than flooding it everywhere.

Starting with Cisco NX-OS Release 10.6(1s), you can configure these multicast routing features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

- IPv4 Layer 2 and Layer 3 multicast
- Multicast consistency checker
- Layer 3 physical interface, port channel, sub-interface, and SVI
- Layer 2 port channel
- Multicast flow counter

IPv4 Layer 2 and Layer 3 multicast

IPv4 Layer 2 and Layer 3 multicast is a comprehensive multicast solution where **Layer 2 multicast** (using IGMP Snooping) provides efficient traffic delivery within a single broadcast domain (VLAN), and **Layer 3 multicast** (using routing protocols like PIM) provides end-to-end routing of that traffic between different subnets across the network.

For more information on IPv4 Layer 2 and Layer 3 multicast, see Overview.

Multicast consistency checker

Multicast consistency checker is a diagnostic tool that periodically verifies the consistency of multicast configurations and states (such as PIM RPs or IGMP queriers) between routers on a shared network segment, helping to identify and troubleshoot misconfigurations.

For more information on multicast consistency checker, see Overview.

Layer 3 physical interface, port channel, sub-interface, and SVI

Layer 3 physical interface, port channel, sub-interface, and SVI are the various types of Layer 3 interfaces on which multicast routing protocols, such as PIM, can be enabled. This allows the device to participate in multicast routing, build distribution trees, and forward multicast packets on physical ports, aggregated links, logical sub-interfaces, or VLAN interfaces (SVIs).

For more information on Layer 3 physical interface, port channel, sub-interface, and SVI, see Overview.

Layer 2 port channel

Layer 2 port channel An aggregated link operating at Layer 2 that can carry multicast traffic for one or more VLANs. IGMP snooping can be enabled on the port channel to learn which multicast groups are needed by downstream devices, ensuring that multicast frames are forwarded efficiently across the aggregated link.

For more information on Layer 2 port channel, see Overview.

Multicast flow counter

Multicast flow counter feature is a hardware-assisted capability that enables a network device to track traffic statistics for individual multicast flows, identified by their unique Source, Group (S,G) pair. This provides granular visibility into the packet and byte volume of specific streams, allowing for precise monitoring and troubleshooting of multicast applications.

For more information on multicast flow counter, see Multicast Counters.

- Multicast routing feature guidelines, on page 14

Multicast routing feature guidelines

This section outlines feature support, guidelines, and limitations for multicast routing functionalities on Cisco N9300 Series smart switches.

Guidelines and limitations for multicast routing Features

Table 3: Supported features and releases

Features	Release
IPv4 Layer 2 and Layer 3 multicast	10.6(1s)
Multicast consistency checker	10.6(1s)
Layer 3 physical interface, port channel, sub-interface, and SVI	10.6(1s)
Layer 2 port channel	10.6(1s)
Multicast flow counter for IPv4	10.6(1s)

CHAPTER 6

Quality of Service features

Quality of Service (QoS) is a set of networking tools and mechanisms designed to manage network traffic and ensure predictable performance by classifying, marking, policing, and prioritizing specific data flows to avoid congestion.

Starting with Cisco NX-OS Release 10.6(1s), you can configure these QoS features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

- Queuing and Scheduling
- QoS statistics

QoS Queuing and Scheduling

QoS Queuing and Scheduling is a two-part process where **Queuing** first orders packets into different traffic classes, and **Scheduling** then methodically controls the output of those packets based on priority to ensure a consistent and efficient flow of network traffic.

For information on Queuing and Scheduling, see Configure Queuing and Scheduling.

QoS statistics

QoS statistics are a collection of performance counters, enabled by default on a network device, used to monitor and display the behavior and effectiveness of applied Quality of Service policies.

For information on QoS statistics, see Monitoring the Statistics.

QoS guidelines

This section outlines feature support, guidelines, and limitations for QoS functionalities on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

Guidelines and Limitations for Queuing and Scheduling

Table 4: Supported features and releases

Features	Release
Queuing and Scheduling	10.6(1s)
QoS statistics	10.6(1s)

Supported or unsupported Queuing and Scheduling features

Supported features:

- Eight queues - SPAN and CPU Queues that are overloaded with eight user queues are supported.
- SP and DWRR are supported. However, the shaper and DWRR accuracy will have a 5% variance.
- QoS statistics are supported.
- Maximum shaper and static limit are supported.

Unsupported features:

- Micro-Burst monitoring is not supported.
- Link Level Flow Control (LLFC) is not supported.
- Dynamic queue-limit is not supported.
- Multicast queuing statistics is not supported.

QoS statistics limitation

On the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches, the **show queuing** command operates independently of the **qos statistics** configuration.

For information on QoS statistics, see Monitoring the Statistics.

CHAPTER 7

Security features

Security features are a collection of integrated tools, protocols, and policies designed to safeguard network systems and data by mitigating threats, preventing unauthorized access, and ensuring data integrity.

Starting with Cisco NX-OS Release 10.6(1s), you can configure these security features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

- IP Access Control List (IP ACL)
- Dynamic Host Configuration Protocol (DHCP) relay
- Control Plane Policing (CoPP)
- Media Access Control Security (MACsec)

IP ACL

An IP ACL is an ordered set of rules that filters network traffic by specifying conditions based on IP packet information, such as source and destination IP addresses and protocol type, to determine whether a packet is permitted or denied.

For information on IP ACLs, see Configure IP ACLs.

DHCP relay

DHCP relay is a feature that enhances network control by centralizing IP address allocation, which enables consistent policy enforcement and aids in the prevention of rogue DHCP servers by forwarding client requests exclusively to trusted, remote servers.

For information on DHCP relay, see Configuring DHCP.

Control Plane Policing

CoPP is a security feature that protects a network device's control plane from Denial of Service (DoS) attacks by applying a policy map to limit the rate of traffic destined for the CPU, ensuring network stability and reachability.

A custom CoPP policy is an administrator-defined ruleset that tailors the protection of a device's control plane by classifying specific traffic types and applying unique rate limits to meet a network's particular security and operational requirements.

For information on CoPP, see Configure Control Plane Policing.

MACsec

MACsec is an IEEE 802.1AE security standard that provides confidentiality and data integrity for Ethernet links by encrypting traffic at Layer 2 on a hop-by-hop basis.

For information on MACsec, see Configure MACsec.

- Security feature guidelines, on page 18

Security feature guidelines

This section outlines feature support, guidelines, and limitations for security functionalities on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

Table 5: Supported features and releases

Features	Release
PACL	10.6(1s)
RACL on L3 interfaces, L3 Port-channel interfaces, subinterfaces, and SVI interfaces	10.6(1s)
PBR ACL	10.6(1s)
DHCP relay	10.6(1s)
Custom CoPP	10.6(1s)
MACsec	10.6(1s)

ACL limitations

- The **mac packet-classify** command is not supported.
- Each TCAM slice supports 7136 entries for RACL or PACL. The Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches have two slices.

DHCP relay limitations

When a DHCP relay is configured on a switch and a RACL is also applied, the RACL will drop unicast DHCP relay traffic. This occurs because RACLs and other security policies have a higher priority in the processing order than control plane traffic, such as DHCP relay.

CoPP limitations

- Policer rates are specified in packets per second (PPS).
- Destination IP-based access-list matching is not supported in Custom CoPP.
- User-defined MAC access lists are not supported; only built-in MAC ACLs can be used.
- A maximum of 28 class maps can be configured under a Custom CoPP policy.

- Custom CoPP supports maximum 360 IPv4 (uni-dimensional) and 180 IPv6 (uni-dimensional) TCAM entries.
- The CoPP Consistency Checker is not supported for Custom CoPP.
- Zero Committed Information Rate (CIR) is not supported. If zero CIR is configured, the hardware sets it to the minimum possible value.

MACsec limitations

- On Cisco N9324C-SE1U switch, MACsec is supported on all 24 QSFP28 ports.
- On Cisco N9348Y2C6D-SE1U switch, MACsec is supported on all 48 SFP28, 6 QSFP-DD, and 2 QSP28 ports.
- The MACsec SecY statistics do not support any statistical information related to egress Secure Association (SA) counters. Egress statistical data always displays as zero.
- For EAPOL configuration, ensure the following requirements are met:
    - The EAPOL MAC address must be in the range 0180.C200.0000 to 0180.C200.00FF (the last byte can be from 0x00 to 0xFF) and the EtherType can be any value between 0x600 and 0xFFFF.
    - Any MAC address is allowed when the EtherType is set to the default value 0x888E.
    - The broadcast MAC address can be used with any EtherType value in the range 0x600 to 0xFFFF.

CHAPTER 8

System management features

System management features in Cisco devices encompass a variety of capabilities designed to facilitate efficient configuration, monitoring, and maintenance of network systems.

Starting with Cisco NX-OS Release 10.6(1s), you can configure these security features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

- Sampled Flow (sFlow)
- Switched Port Analyzer (SPAN)
- Encapsulated Remote SPAN (ERSPAN)
- Link Layer Discovery Protocol (LLDP)

sFlow

sFlow (Sampled Flow) is a standards-based network traffic monitoring technology that enables real-time monitoring of traffic in data networks containing switches and routers.

For information on sFlow, see Configuring sFlow.

SPAN

SPAN is a network monitoring feature on Cisco devices that allows duplication of traffic from specified source ports to a designated destination port. This enables analysis of network traffic by sending a copy of the packets to an external analyzer connected to the destination port.

For information on SPAN, see Switched Port Analyzer.

ERSPAN

ERSPAN is used to transport mirrored traffic over an IP network, enabling remote monitoring of multiple switches across a network by sending traffic from source ports or VLANs on one device to destination ports or analyzers on another device.

For information on ERSPAN, see Configuring ERSPAN.

LLDP

LLDP is a vendor-neutral, one-way device discovery protocol defined by the IEEE 802.1AB standard. It enables network devices to advertise information about themselves to other devices on the same local network segment, facilitating device discovery and network topology mapping.

For information on LLDP, see Configuring ERSPAN.

- System management feature guidelines, on page 22

System management feature guidelines

This section outlines the supported system management features, corresponding switches, software releases, and known limitations for Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

Table 6: Supported System management features and releases

Features	Release
SPAN	10.6(1s)
ERSPAN	10.6(1s)
sFlow	10.6(1s)
LLDP	10.6(1s)

PTP is not supported on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

SPAN limitations

- **Sessions**—A maximum of 10 active monitor (SPAN) sessions are supported at a time.
- **Packet mirroring**—Sharing of the same source port or interface across multiple sessions is not supported. SPAN mirrored packets use the default egress queue and do not have a dedicated SPAN egress queue.
- **SPAN to CPU**—Monitor statistics are not displayed for SPAN to CPU. Both Rx and Tx mirroring are supported for SPAN to CPU.
- **Port-channel interface**—When a port-channel interface with multiple member ports is configured as a SPAN destination, only one member interface is used for mirrored traffic. Member selection is handled in software, which results in packet loss when membership changes.
- **MTU truncation**—MTU truncation is supported only for 144 bytes in Rx mirroring and 80 bytes in Tx mirroring, excluding FCS.
- **Unsupported features**—The features that are not supported include:
  - SPAN on subinterfaces,
  - sharing of the same source port or interface across sessions,
  - tunnel ports,
  - VLAN source,
  - UDF, and
  - ACL filter.

ERSPAN limitations

- **Sessions**—A maximum of 10 active monitor (ERSPAN) sessions are supported at a time.
- **Packet mirroring**—Sharing of the same source port or interface across multiple sessions is not supported. ERSPAN mirrored packets use the default egress queue and do not have a dedicated ERSPAN egress queue.
- **MTU truncation**— MTU truncation is supported only for 144 bytes in Rx mirroring and 80 bytes in Tx mirroring, excluding FCS.
- **Port-channel interface**—When port-channel interface with multiple member ports is configured as an ERSPAN destination, only one member interface is used for mirrored traffic. Member selection is handled in software, which results in packet loss when membership changes.
- **Unsupported features**—The features that are not supported include:
  - ERSPAN on subinterfaces
  - sharing of the same source port or interface across sessions,
  - tunnel ports,
  - VLAN as source,
  - UDF, and
  - ACL filter.

sFLOW limitations

- For egress sampled packet, re-written information is not available in sFlow record.
- Egress Layer 2 source interface is not supported.
- sFlow is not supported on the sub-interface traffic.

CHAPTER 9

Unicast routing features

Unicast routing features are a comprehensive suite of protocols, mechanisms, and services responsible for determining the optimal path and forwarding IP packets from a single source to a single, specific destination across one or more interconnected networks. Their primary function is to build and maintain a routing table that enables intelligent, hop-by-hop packet delivery.

Starting with Cisco NX-OS Release 10.6(1s), you can configure these unicast routing features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

Core IP services and addressing

These are the fundamental protocols and concepts required for basic IP communication on a local network segment.

- **IPv4/IPv6 unicast routing**is the fundamental process of forwarding IP packets from a single source device to a single, specific destination device using their unique IPv4 or IPv6 addresses. Layer 3 uses either the IPv4 or IPv6 protocol. IPv6 increases the number of network address bits from 32 bits (in IPv4) to 128 bits.

  For more information on IPv4 Unicast routing, see IPv4 addresses.

  For more information on IPv6 Unicast routing, see Configuring IPv6.

- **Host route** is a specific type of route in a routing table that directs traffic to a single host IP address using a /32 (for IPv4) or /128 (for IPv6) prefix length.

  For more information on host route, see IPv4 addresses and Configuring IPv6.

- **Address Resolution Protocol (ARP)** is a protocol used to map a known IPv4 address to its corresponding Layer 2 MAC address within a local network segment.

  For more information on host route, see IPv4 addresses and Configuring IPv6.

- **Neighbor Discovery (ND)** is an IPv6 protocol suite that replaces ARP and provides additional functions such as router discovery, address autoconfiguration, and neighbor reachability tracking.

  For more information on host route, see IPv4 addresses and Configuring IPv6.

Routing protocols

These are the dynamic protocols used to build and maintain routing tables by exchanging reachability information with other routers.

- **Border Gateway Protocol (BGP)**is an inter-autonomous system routing protocol. A BGP router advertises network reachability information to other BGP routers using Transmission Control Protocol (TCP) as its reliable transport mechanism. The network reachability information includes the destination network prefix, a list of autonomous systems that needs to be traversed to reach the destination, and the next-hop router. Reachability information contains additional path attributes such as preference to a route, origin of the route, community and others.

  For more information on BGP, see Configuring BGP.

- **Open Shortest Path First (OSPF)** protocol is a link-state routing protocol used to exchange network reachability information within an autonomous system. Each OSPF router advertises information about its active links to its neighbor routers. Link information consists of the link type, the link metric, and the neighbor router that is connected to the link. The advertisements that contain this link information are called link-state advertisements. For more information on OSPF, see Configuring OSPFv2.

- **Intermediate System-to-Intermediate System (IS-IS)** protocol is an intradomain Open System Interconnection (OSI) dynamic routing protocol specified in the International Organization for Standardization (ISO) 10589. The IS-IS routing protocol is a link-state protocol. IS-IS features are as follows:

  - Hierarchical routing
  - Classless behavior
  - Rapid flooding of new information
  - Fast Convergence
  - Very scalable

  For more information on ISIS, see Configuring IS-IS.

Traffic forwarding and path control

These features influence how a router forwards packets based on the information in its routing table.

- **Equal-Cost Multi-Path (ECMP)** is a routing mechanism that allows a device to load-balance traffic across multiple paths when they all have the same routing cost to a given destination.

  For more information on ECMP, see Overview.

- **IP Directed Broadcast** is a feature that enables a packet, addressed to the broadcast IP of a remote subnet, to be routed as a unicast packet across the network until the final router translates it into a local Layer 2 broadcast on the destination segment.

  For more information on IP directed broadcast, see Configuring IP Directed Broadcasts.

- **Policy-Based Routing (PBR)** is a feature that enables a network device to override its standard destination-based routing table for specific traffic flows. By using policies, typically defined by an access list, to classify unicast packets based on criteria like source/destination IP, protocol, or port number, PBR can redirect that traffic to a designated next-hop IP address or egress interface, allowing for customized traffic paths for purposes like service chaining or traffic engineering.

  For more information on IP directed broadcast, see Configuring Policy-Based Routing.

Security features

These features are designed to protect the network from common Layer 3 attacks.

- **Unicast Reverse Path Forwarding (uRPF)** is a security feature that helps prevent IP address spoofing by verifying that the source IP address of an incoming packet is reachable through the same interface on which the packet was received.

  For more information on IP directed broadcast, see Configuring HSRP.

Interface configuration features

This category includes specialized ways to configure IP on an interface.

- **IP Unnumbered (non-SVI) interface** is a feature that enables a point-to-point physical interface (non-SVI) to process IP traffic without having its own assigned IP address, typically by "borrowing" the IP address from another interface.

  For more information on non SVI IP unnumbered interface, see IPv4 addresses and Configuring IPv6.

High Availability and redundancy

This category includes features designed to prevent single points of failure, provide resilient network paths, and ensure continuous operation.

- **Hot Standby Router Protocol (HSRP)** is a first-hop redundancy protocol (FHRP) that provides transparent, high-availability gateway services for IP hosts. It achieves this by creating a virtual router, with a shared virtual IP and virtual MAC address, from a group of physical routers. Within the group, one router is elected as the **active** router to forward traffic, while another becomes the **standby** router, ready to take over instantly if the active router fails.

- **Virtual Router Redundancy Protocol (VRRP)** allows for a transparent failover at the first-hop IP router by configuring a group of routers to share a virtual IP address. VRRP selects an allowed router in that group to handle all packets for the virtual IP address. The remaining routers are in standby and take over if the allowed router fails.

  For more information on VRRP, see Configuring VRRP.

- **Virtual Port-Channel (vPC)** is a virtualization technology that allows two physical network switches to appear as a single logical switch to a connected device, enabling all-active link aggregation for increased bandwidth, providing high availability, and creating a loop-free topology.

- **vPC Peer Gateway** is a feature that allows a vPC switch to act as the active gateway for packets addressed to the shared router MAC address, enabling local forwarding of traffic without sending it across the vPC peer-link, which optimizes routing and improves compatibility with network-attached storage (NAS) devices.

For more information on HSRP, vPC, and vPC peer gateway, see Configuring HSRP.

Layer 2 addressing and interface configuration

This category includes features related to the configuration and management of Layer 2 addresses and interface properties.

- **User-defined MAC address** is a manually configured Layer 2 address that overrides the default hardware address for a logical interface, such as a Switched Virtual Interface (SVI) or a vPC system, to ensure a stable and predictable identity in the network.

  For more information on user-defined MAC address, see Configuring HSRP.

- Unicast routing feature guidelines, on page 28

Unicast routing feature guidelines

This section outlines feature support, guidelines, and limitations for unicast routing functionalities on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

Table 7: Supported features and releases

Features	Release
IPv4 Layer-3 features • **IPv4 unicast routes** • **IPv4 host routes** • **IPv4 Neighbor Discovery** • **ARP** • **IP directed broadcast** • **Non SVI IP unnumbered interface**	10.6(1s)
IPv6 Layer-3 features • **IPv6 unicast routes** • **IPv6 host routes** • **IPv6 Neighbor Discovery** • **ARP**	10.6(1s)
BGP	10.6(1s)
ECMP	10.6(1s)
OSPFv2 and OSPFv3	10.6(1s)
OSPFv2	10.6(1s)
IS-IS	10.6(1s)
Policy-based routing	10.6(1s)
uRPF	10.6(1s)
HSRP	10.6(1s)

Features	Release
VRRP	10.6(1s)
vPC and vPC peer gateway	10.6(1s)
User-defined MAC address	10.6(1s)

Supported PBR features:

- Base redirection
- 32 way ECMP
- NULL routes
- IPv4 and IPv6 PBR policies

Supported PBR features:

- Default next hop
- Set VRF
- Fast convergence
- MPLS
- SRTE
- VXLAN
- Tunnel support

User-defined MAC address limitations:

- For Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches, of the total Router MAC scale, the number of unique 32 bit MSBs allowed for user-defined MAC addresses is 11.

CHAPTER 10

VXLAN features

Virtual Extensible LAN (VXLAN) is a network virtualization technology that extends Layer 2 segments over a Layer 3 infrastructure using MAC-in-UDP encapsulation, enabling the creation of highly scalable and flexible multitenant data center fabrics.

Starting with Cisco NX-OS Release 10.6(1s), you can configure these VXLAN features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

VXLAN BGP EVPN

VXLAN BGP EVPN is a network virtualization overlay solution that uses BGP with EVPN address family as the control plane to distribute Layer 2 and Layer 3 reachability information across a VXLAN data plane.

For information on VXLAN BGP EVPN, see Configure VXLAN BGP EVPN.

VXLAN L2VNI and L3VNI

VXLAN Layer 2 Virtual Network Identifier (L2VNI) is a VXLAN segment identifier used to create a Layer 2 broadcast domain for bridging traffic.

For information on L2VNI, see Configure VXLAN BGP EVPN.

Layer 3 Virtual Network Identifier (L3VNI) is associated with a VRF to provide Layer 3 routing services between different L2VNIs.

For information on L3VNI, see Configure New L3VNI Mode.

VXLAN BGP EVPN multi-site Anycast BGW

VXLAN BGP EVPN multi-site Anycast BGW is a high-availability architecture where multiple Border Gateway (BGW) nodes, connecting separate VXLAN EVPN sites, share a common anycast IP address to provide resilient and optimized data traffic forwarding and control plane signaling between sites.

For information on VXLAN BGP EVPN multi-site Anycast BGW, see Configure VXLAN EVPN Multi-Site.

VXLAN BGP EVPN border spine

VXLAN BGP EVPN border spine is a network device that functions as both a spine switch within a VXLAN fabric and a BGW, responsible for connecting the fabric to external networks or other data center sites.

For information on VXLAN BGP EVPN border spine, see Configure VXLAN BGP EVPN.

VXLAN standalone or vPC VTEP, vPC, and vPC fabric peering

Standalone, vPC, and vPC fabric peering are methods for establishing network connectivity and redundancy for VXLAN Tunnel Endpoints (VTEPs).

- **VXLAN standalone** refers to a single-node VTEP.

  For information on standalone or vPC VTEP, see Configure VXLAN.

- **vPC** allows two switches to act as a single logical VTEP for dual-homed hosts. For more information on vPC, see vPC Considerations for VXLAN Deployment, Configure vPC Multi-Homing.

- **vPC fabric peering** establishes underlay routing between vPC pairs. For more information on vPC fabric peering, see Configure vPC Fabric Peering.

Distributed Anycast Gateway

Distributed Anycast Gateway is a VXLAN EVPN feature where the default gateway IP and MAC addresses for a subnet are identically configured on all VTEPs within that Layer 2 segment, enabling optimal east-west traffic routing.

For more information on distributed Anycast Gateway, see Distributed Anycast Gateway.

DHCP relay

DHCP relay is a feature that forwards DHCP broadcast requests from clients within a VXLAN overlay network to a DHCP server located in a different subnet, often outside the fabric, by encapsulating the requests and sending them across the Layer 3 underlay.

For more information on DHCP relay, see DHCP Relay in VXLAN BGP EVPN.

IPv4 and IPv6 unicast overlay traffic

IPv4 and IPv6 unicast overlay traffic is a transport standard of IPv4 and IPv6 unicast packets, encapsulated within a VXLAN header, across the overlay network between source and destination endpoints.

For more information on IPv4 and IPv6 unicast overlay traffic, see Configure the Underlay.

BUM traffic

Broadcast/Unknown unicast/Multicast (BUM) traffic refers to Broadcast, Unknown Unicast, and Multicast traffic within a VXLAN segment, which is typically handled by replicating the traffic and forwarding it to all relevant VTEPs, either through multicast replication in the underlay or ingress replication (head-end replication).

For more information on BUM traffic, see Configure the Underlay.

IPv4 unicast underlay (IR)

IPv4 unicast underlay (IR) is a standard IPv4 unicast routing protocol in the physical network (underlay) to provide reachability between VTEPs, where BUM traffic is handled via Ingress Replication (IR), meaning the source VTEP unicasts a copy of the packet to every other relevant VTEP.

For more information on IPv4 unicast underlay, see Configure the Underlay.

IPv4 multicast underlay with PIM ASM

IPv4 multicast underlay with Protocol Independent Multicast - Any-Source Multicast (PIM ASM) is an IPv4 multicast-enabled physical network (underlay), typically running PIM ASM, to efficiently handle BUM traffic by forwarding a single copy of a BUM packet to a multicast group that all relevant VTEPs have joined.

For more information on IPv4 multicast underlay with PIM ASM, see Multicast Routing in the VXLAN Underlay.

VXLAN uplinks

VXLAN uplinks are physical or port-channel Layer 3 interfaces on a VTEP (typically a leaf switch) that connect to the underlay network (typically spine switches) and carry the encapsulated VXLAN traffic.

VXLAN counters

VXLAN counters is a statistical counters maintained by a network device to track the volume of VXLAN traffic, including encapsulated and decapsulated packets and bytes, used for monitoring, performance analysis, and troubleshooting.

Underlay ECMP and Overlay ECMP (L3)

Underlay ECMP refers to the use of multiple equal-cost paths in the physical network to load-balance traffic between VTEPs, while Overlay ECMP is a BGP EVPN feature that enables load-balancing of traffic across multiple remote VTEPs that are advertising reachability to the same destination prefix.

For more information on Underlay ECMP and Overlay ECMP, see Configure the Underlay.

VXLAN NGOAM

VXLAN NGOAM refers to a suite of tools and protocols designed for proactive monitoring and troubleshooting of VXLAN overlay networks, such as traceroute and ping for overlay paths.

For more information on VXLAN NGOAM, see VXLAN OAM or VXLAN NGOAM.

Multicast underlay BUD node

Multicast underlay Bridge and Drop (BUD) node is a device in a multicast underlay that is not a VTEP for a given VNI but is on the multicast tree path, which forwards the VXLAN multicast traffic without decapsulating it.

For more information on Multicast underlay BUD node, see Configure Bud Node.

DSVNI

VXLAN EVPN with downstream VNI provides the following solutions:

- Enables asymmetric VNI communication across nodes in a VXLAN EVPN network
- Provides customers access to a common shared service outside of their domain (tenant VRF)
- Supports communication between isolated VXLAN EVPN sites that have different sets of VNIs

For more information on DSVNI and route leak, see Configure VXLAN BGP EVPN.

IGMP snooping

Internet Group Management Protocol (IGMP) snooping is a feature applied within a VXLAN overlay that allows a VTEP to monitor IGMP messages from hosts, learn which hosts are interested in specific multicast groups, and prune multicast traffic to only forward it to VTEPs with interested receivers.

For more information on IGMP snooping, see Optimized Layer 2 Overlay Multicast.

ARP suppression

Address Resolution Protocol (ARP) suppression is an efficiency feature in VXLAN EVPN where a VTEP intercepts ARP requests and, if it already knows the MAC-to-IP binding from the BGP control plane, responds directly to the host, thereby suppressing (preventing) the ARP request from being flooded across the entire VXLAN segment.

For more information on ARP suppression, see Configure VXLAN BGP EVPN.

TRMv4 L3 mode

Tenant Routed Multicast for IPv4 (TRMv4) L3 mode is a mode of TRM for IPv4 that enables efficient and scalable multicast forwarding across different subnets (VNIs) within a VXLAN EVPN fabric using a Layer 3 overlay.

For more information on TRMv4 L3 mode, see Configure Tenant Routed Multicast.

BGW advertisement

The **advertise-pip** is a BGP EVPN command used on Border Gateways (BGWs) in a multi-site deployment to advertise the BGW's Primary IP (PIP) as the next-hop for routes learned from other sites, ensuring symmetric traffic flows for stateful services.

For more information on **advertise-pip**, see Configure vPC Multi-Homing.

BGW advertise using PIP towards fabric and DCI

BGW advertise using PIP towards fabric and Data Center Interconnect (DCI) is a specific BGW behavior in a multi-site architecture where the PIP is advertised as the next-hop for external routes both internally towards the local fabric and externally towards the DCI.

For more information on **fabric-advertise-pip l3**, see the Advertise Using PIP Towards Fabric section

- VXLAN feature guidelines , on page 34

VXLAN feature guidelines

This section outlines feature support, guidelines, and limitations for VXLAN functionalities on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

Table 8: Supported features and releases with limitations

Features	Note	Release
VXLAN BGP EVPN	—	10.6(1s)

Features	Note	Release
VXLAN L2VNI and L3VNI	VLAN based L3VNI configuration is deprecated. Only **vni** *vnid* **l3** command is supported. For information on L3VNI, see Configure New L3VNI Mode.	10.6(1s)
VXLAN BGP EVPN Multi-Site Anycast Border Gateway	—	10.6(1s)
VXLAN BGP EVPN Border Spine	—	10.6(1s)
Standalone, vPC, and vPC fabric peering	Supported as leaf or border leaf	10.6(1s)
Distributed Anycast Gateway	This feature is supported with **fabric forwarding anycast-mode** command. This configuration is supported with the following combination of configuration • On VXLAN-VLAN only • With SVI configured or up on such VXLAN-VLAN • With global Fabric Anycast MAC configuration in system	10.6(1s)
DHCP Relay	—	10.6(1s)
IPv4 and IPv6 unicast overlay traffic	—	10.6(1s)
BUM traffic • **IPv4 unicast underlay (IR)** • **IPv4 multicast underlay with PIM ASM**	Fabric and DCI	10.6(1s)
VXLAN uplinks	• Only Ethernet and Port-channel routed interfaces are supported as uplinks • SVI or L3-subinterfaces as uplinks are not supported	10.6(1s)

Features	Note	Release
VXLAN counters	- VXLAN peer-based total packet/byte counters are supported - VNI based total packet and byte counters are supported - Peer counters or per-peer-per-vni counters are not supported	10.6(1s)
Underlay ECMP and Overlay ECMP (L3)	—	10.6(1s)
VXLAN NGOAM	VTEP and Host reachability are supported.	10.6(1s)
Multicast underlay BUD node	-	10.6(1s)
DSVNI	-	10.6(1s)
IGMP snooping	Not supported for Anycast BGW	10.6(1s)
ARP Suppression	—	10.6(1s)
TRMv4 L3 Mode - **IPv4 unicast underlay (IR)** - **IPv4 multicast underlay with PIM ASM**	- **Fabric**: Only multicast underlay is supported - **DCI**: IR and multicast underlay are supported	10.6(1s)
advertise-pip	For more information on **advertise-pip** command, see Configure vPC Multi-Homing.	10.6(1s)
BGW advertise using PIP towards fabric and DCI	**fabric-advertise-pip l3** command. For more information on advertise PIP, see the Advertise Using PIP Towards Fabric section **dci-advertise-pip**	10.6(1s)

VXLAN unsupported features

These VXLAN functionalities are not supported on Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

- Core fabric and underlay design features: VXLAN flood and learn, VXLAN static tunnels, RFC 5549 underlay, VXLAN IPv6 underlay, IPv6 unicast or IR underlay (fabric), and IPv6 Multicast underlay and Multicast underlay with PIM BIDIR (fabric).
- Overlay routing and multi-homing features: ESI-RX or VXLAN ESI multi-homing, Proportional ECMP (Mixed path), and VXLAN Traffic Engineering (TE)

- Multi-Site and DCI features: vPC multi-site BGW, IPv6 IR underlay (DCI), EVPN multi-site storm control, CloudSec, and VXLAN to SR/MPLS handoff.
- Advanced multicast handling features: Multicast Listener Discovery (MLD) snooping, Neighbor Discovery (ND) suppression, and TRMv6, TRM L2 mode and TRM mixed mode, and TRM data MDT.
- Security features: First-Hop Security (FHS), ACL on VXLAN, Security Group ACL (SGACL), and Null route or Static remote MAC
- Overlay services & integrations features: VNF (gateway IP), VXLAN Policy-Based Routing (PBR), VXLAN Quality of Service (QoS) policy, and VXLAN distributed NAT support.
- Access and host connectivity features: VXLAN access features: A general category for features applied at the host-facing edge of the fabric such as Private VLAN (PVLAN), 802.1x, Multitag, Cross Connect, Port security, Port VLAN translation, QinVNI, Selective QinVNI, and Layer 2 Protocol Tunneling (L2PT), and Fabric Extender (FEX)
- Operations, Administration, & Maintenance (OAM) features: Separate counters for broadcast, multicast, and unicast traffic and Southbound loop detection

Support and unsupported features of TRMv4

- Beginning with Cisco NX-OS Release 10.6(1s), Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches supports TRMv4 on vPC leaf, vPC fabric peering leaf, Anycast BGW, and standalone leaf.
- Beginning with Cisco NX-OS Release 10.6(1s), Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches supports TRMv4 L3 on vPC leaf, vPC fabric peering leaf, Anycast BGW, and standalone leaf.

  Supported features are:

  - TRMv4,
  - Ingress Replication between DCI peers across the core,
  - Multicast underlay for fabric peers, and
  - VLAN based L3VNI configuration is deprecated. Only **vni vnid l3** command is supported.

  Unsupported features are:

  - TRMv6 and
  - Data MDT.

CHAPTER 11

Commonly supported features

Starting with Cisco NX-OS Release 10.6(1s), you can configure these common features on the Cisco N9324C-SE1U, Cisco N9348Y2C6D-SE1U switches.

- Power-On Auto Provisioning (POAP)
- Smart Licensing Using Policy (SLP)
- Telemetry
- Packet Tracer

POAP

POAP is a Cisco automation feature designed to simplify the initial deployment of Cisco Nexus switches and other supported devices. It automates the process of upgrading software images and installing configuration files on devices that are being deployed in the network for the first time.

For information on POAP, see Using PowerOn Auto Provisioning.

SLP

SLP is an enhanced version of Smart Licensing, the objective of which is to provide a licensing solution that does not interrupt the operations of your network and to enable a compliance relationship to account for the hardware and software licenses you purchase and use.

For information on SLP, see Smart Licensing Using Policy.

Telemetry

Telemetry is a modern network monitoring approach that enables continuous, automated data collection for analyzing and troubleshooting network health. Traditional mechanisms like SNMP, CLI, and Syslog use a pull model, where data requests originate from clients. This pull model has limitations in scalability and automation, especially when multiple network management stations (NMS) are involved, as it requires continual manual intervention and only sends data upon request.

Telemetry overcomes these limitations by using a push model that continuously streams data from network devices to clients, providing near-real-time access to monitoring data. This push model enhances automation, scalability, and efficiency in network monitoring.

For information on telemetry, see Telemetry.

Packet Tracer

The Packet Tracer is a troubleshooting tool that allows a packet to be captured from the Network Processor. Similar to the ELAM tool available on Cisco Nexus 9000 Cloud Scale switches, this tool provides information to understand how the ASIC forwarded the captured packet. This information is useful to troubleshoot packet flow.

While various tools like SPAN, ERSPAN, Ethanalyzer exist to debug packet flow issues, Packet Tracer allows troubleshooting within the forwarding pipeline of an ASIC without any performance penalty or disruption to the environment.

For information on packet tracer, see Troubleshooting Packet Flow Issues.