Revised: September 15, 2026

Routing and Switching for Network Mode

Supported routing and switching features

This article outlines the supported routing and switching features, their functions, and associated limitations for the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches. The features cover Layer 2 switching, VXLAN, unicast, multicast, and label switching. Use the related articles in this guide for configuration requirements and feature-specific limitations.

Layer 2 switching

This article describes the core Layer 2 switching features supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches and explains how Layer 2 switching provides segmentation, loop prevention, and frame forwarding.

Layer 2 switching concepts

Layer 2 switching uses MAC addresses and VLAN membership to forward Ethernet frames within a broadcast domain. Spanning-tree features control redundant paths so segmentation and forwarding remain stable.

Table 1: Layer 2 switching terms

Term	Definition	How it fits
Global MAC learning	Recording a source MAC address against the interface where the frame arrived.	Builds the forwarding information used to send known unicast frames.
VLAN	A logical broadcast domain within the switching network.	Defines the Layer 2 scope in which MAC learning and flooding operate.
Spanning tree	A loop-prevention mechanism for redundant bridged paths.	STP, Rapid PVST+, and MST determine which redundant paths forward.

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.

Beginning with NX-OS Release 10.6(2)F, VLANs are supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information, refer to Configuring VLANs.

Global 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.

Beginning with NX-OS Release 10.6(2)F, global MAC learning is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information, refer to Configuring Layer 2 Switching.

STP

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

Beginning with NX-OS Release 10.6(2)F, STP is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information, refer to 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.

Beginning with NX-OS Release 10.6(2)F, Rapid PVST+ is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information, refer to 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.

Beginning with NX-OS Release 10.6(2)F, MST is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information, refer to Configuring MST Using NX-OS

Unicast routing

This article explains how unicast routing uses dynamic routing protocols, static routes, and VRFs to direct traffic across scalable Layer 3 networks.

Unicast routing concepts

Unicast routing forwards traffic from one source to one destination by using IP addresses and a routing table. Interfaces provide adjacency, routing protocols or static routes provide reachability, and policy or redundancy features influence the selected path.

Table 2: Unicast routing terminology

Term	Definition	Relationship
Routing table	A collection of routes that describes reachable destinations.	Forwarding uses the best available route to select the next hop or interface.
Static routing	A routing method that uses manually configured routes to reach destinations.	Static routes provide reachability without dynamic route exchange.
Virtual Routing and Forwarding (VRF)	A virtualization technology that maintains separate routing tables on the same device.	Each VRF provides an isolated routing context for its associated interfaces and routes.
Routing protocol	A protocol that exchanges reachability information.	BGP, EIGRP, OSPF, RIP, and IS-IS can populate the routing table dynamically.

Term	Definition	Relationship
Enhanced Interior Gateway Routing Protocol (EIGRP)	A dynamic routing protocol that exchanges reachability information between neighboring routers.	EIGRP can populate the routing table with learned routes.
Routing Information Protocol (RIP)	A routing protocol that exchanges reachability information between neighboring routers.	RIP can populate the routing table with learned routes.
ECMP	Equal-Cost Multi-Path forwarding across routes with the same cost.	Uses multiple eligible paths for load sharing and resilience. The smart switches support 8192 ECMP groups and 512-way ECMP.
First-hop redundancy	A shared gateway design that provides failover for hosts.	HSRP, VRRP, and vPC-related features provide different redundancy roles.

IPv4 Layer 3 unicast features

IPv4 Layer 3 unicast features provide the capabilities required to route IPv4 traffic across enterprise and data center networks. The feature relationships are organized into routing and addressing, neighbor and address resolution, and directed broadcast.

Beginning with NX-OS Release 10.6(2)F, IPv4 Layer 3 unicast features are supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

- **Routing and addressing**:

  Includes core mechanisms such as IPv4 unicast routes for packet forwarding, host routes for specifying individual device destinations, and non-SVI IP unnumbered interfaces that optimize IP address usage on point-to-point links.

  - **IPv4 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 addresses.

    For more information on IPv4 unicast routing, refer to IPv4 addresses.

  - **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) prefix length.

    For more information on host route, refer to IPv4 addresses.

  - **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 interfaces, refer to IPv4 addresses.

- **Neighbor and address resolution**:

  Features like IPv4 Neighbor Discovery (ND) and Address Resolution Protocol (ARP) enable devices to dynamically map IP addresses to MAC addresses, ensuring seamless communication across the local network segment.

- **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, refer to IPv4 addresses.

- **Neighbor Discovery (ND)** is an 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, refer to IPv4 addresses.

- **IP directed broadcast** 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, refer to Configuring IP Directed Broadcasts.

Together, these IPv4 Layer 3 unicast features support robust, scalable, and flexible IP routing, address management, and network interoperability in modern NX-OS environments.

For more information on IPv4 Layer 3 unicast features, refer to IPv4 addresses.

IPv6 Layer 3 unicast features

The IPv6 Layer 3 unicast features in NX-OS provide the capabilities required to route IPv6 traffic across enterprise and data center networks. The feature relationships are organized into routing and addressing and neighbor and address resolution.

Beginning with NX-OS Release 10.6(2)F, IPv6 Layer 3 unicast features are supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

- **Routing and addressing**:

  IPv6 unicast routes enable the forwarding of traffic to specific destinations using globally unique IPv6 addresses. IPv6 host routes provide precise control by directing traffic to individual devices, utilizing the /128 prefix length for single-host reachability.

- **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 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 IPv6 Unicast routing, refer to 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 /128 (for IPv6) prefix length.

  For more information on host route, refer to Configuring IPv6.

- **Neighbor and address resolution**:

  IPv6 ND replaces traditional ARP and facilitates critical functions such as address resolution, router discovery, and neighbor reachability, enabling seamless communication and network self-configuration. Additionally, ARP provides backward compatibility for environments where IPv4 and IPv6 coexist.

- **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, refer to 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, refer to Configuring IPv6.

Together, these IPv6 Layer 3 unicast features ensure robust connectivity, efficient address management, and dynamic device interaction for next-generation, dual-stack, or IPv6-exclusive NX-OS deployments.

For more information on IPv6 Layer 3 unicast features, refer to Configuring IPv6.

Static routing

Static routing uses manually configured routes to direct traffic to specific destinations. Beginning with NX-OS Release 10.6(2)F, static routing is supported on N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

VRF

Virtual Routing and Forwarding (VRF) maintains separate routing tables on the same device. Each VRF provides an isolated routing context for its associated interfaces and routes. Beginning with NX-OS Release 10.6(2)F, VRF is supported on N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Unicast routing protocols

Unicast routing protocols exchange reachability information and build dynamic routing tables. The related topics cover dynamic routing protocols, equal-cost path selection, and policy-based forwarding.

The table shows the first NX-OS release that supports each feature on the Smart switches. It is a navigation aid; review the related feature topic for protocol-specific capabilities, restrictions, and configuration guidance.

Table 3: Smart-switch unicast feature support by release

Feature	First supported release	Scope
BGP	10.6(2)F	N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.
OSPF	10.6(2)F	N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.
EIGRP	10.6(2)F	N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.
RIP	10.6(2)F	N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.
IS-IS	10.6(2)F	N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.
ECMP	10.6(2)F	Includes 8192 ECMP groups and 512-way ECMP on the Smart switches.
Policy-based routing	10.6(2)F	Supported PBR capabilities and limitations are listed in the feature topic.

BGP

Border Gateway Protocol (BGP) is an inter-autonomous system routing protocol. BGP routers exchange network prefixes, autonomous system paths, next hops, and other route attributes over Transmission Control Protocol (TCP).

Use BGP to exchange routes between autonomous systems or to provide a policy-driven control plane for a data center fabric. Route attributes and policy determine which advertised path is selected and installed.

Beginning with NX-OS Release 10.6(2)F, BGP is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Guidelines and limitations

Plan the local autonomous system, neighbor relationships, address families, route policy, and next-hop reachability before configuring BGP. For BGP configuration guidelines and limitations, refer to Configuring BGP.

OSPF

Open Shortest Path First (OSPF) is a link-state routing protocol. OSPF routers advertise link information to their neighbors, and the resulting link-state database supports path calculation within an autonomous system.

Use OSPFv2 to exchange IPv4 reachability information and OSPFv3 to exchange IPv6 reachability information. OSPF forms adjacencies with neighboring routers, synchronizes link-state information, and calculates paths within an area hierarchy.

Beginning with NX-OS Release 10.6(2)F, OSPF is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Guidelines and limitations

Plan the area structure, router IDs, interface participation, and adjacency requirements before configuring OSPF. For OSPFv2 and OSPFv3 configuration guidelines and limitations, refer to Configuring OSPFv2 and Configuring OSPFv3.

EIGRP

Enhanced Interior Gateway Routing Protocol (EIGRP) combines distance-vector and link-state protocol capabilities to exchange reachability information between neighboring routers.

Use EIGRP to exchange routes with directly connected neighbors and select paths using the protocol metric. Neighbor relationships and the advertised network prefixes determine which routes become available for forwarding.

Beginning with NX-OS Release 10.6(2)F, EIGRP is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Guidelines and limitations

Plan the autonomous system, neighbor interfaces, advertised networks, and route redistribution before configuring EIGRP. For EIGRP configuration guidelines and limitations, refer to Configuring EIGRP.

RIP

Routing Information Protocol (RIP) exchanges reachability information between neighboring routers. RIPv2 uses User Datagram Protocol (UDP) to exchange IPv4 routing information.

Use RIP to exchange IPv4 routes in networks that use a simple distance-vector routing design. RIP selects routes using hop count and periodically exchanges routing updates with neighboring routers.

Beginning with NX-OS Release 10.6(2)F, RIP is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Guidelines and limitations

Plan the RIP version, participating interfaces, advertised networks, and redistribution boundaries before configuring RIP. For RIP configuration guidelines and limitations, refer to Configuring RIP.

IS-IS

Intermediate System-to-Intermediate System (IS-IS) is an intradomain link-state routing protocol specified in ISO 10589. IS-IS supports hierarchical and classless routing, rapid flooding, fast convergence, and scalable path calculation.

Use IS-IS to distribute link-state information across an intradomain routing topology. Level 1 and Level 2 areas provide hierarchy, while the link-state database supports shortest-path calculation.

Beginning with NX-OS Release 10.6(2)F, IS-IS is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Guidelines and limitations

Plan the area hierarchy, NET addresses, participating interfaces, and level behavior before configuring IS-IS. For IS-IS configuration guidelines and limitations, refer to Configuring IS-IS.

ECMP

Equal-Cost Multi-Path (ECMP) forwarding load-balances traffic across multiple paths when the paths have the same routing cost to a destination.

ECMP uses multiple next hops with the same routing cost to distribute traffic and provide path redundancy. The routing protocol supplies the equal-cost paths, and the forwarding plane selects among them for traffic distribution.

Beginning with NX-OS Release 10.6(2)F, ECMP, 8192 ECMP groups, and 512-way ECMP are supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Guidelines and limitations

Ensure that the underlay or routed topology presents equal-cost next hops and account for the supported ECMP group and path-width limits. For ECMP configuration and path-selection guidelines, refer to Overview and Configuring BGP.

Policy-based routing

Policy-Based Routing (PBR) classifies unicast packets by criteria such as source or destination IP address, protocol, or port. PBR then redirects selected traffic to a designated next hop or egress interface.

Use PBR when traffic must follow a forwarding policy based on packet attributes instead of only the destination routing table. A route map matches the selected traffic and applies the supported forwarding action.

Beginning with NX-OS Release 10.6(2)F, PBR is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

PBR guidelines and limitations

The Smart switches support these PBR capabilities:

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

The Smart switches do not support these PBR capabilities:

- Default next hop.
- Set VRF.
- Fast convergence.
- MPLS.
- SR-TE.
- VXLAN.
- Tunnel support.

For more information on PBR, refer to Configuring Policy-Based Routing.

Security and path validation

The security and path validation features in NX-OS are essential for safeguarding network infrastructure against common Layer 3 threats, such as IP address spoofing and unauthorized traffic. A key capability in this area is **Unicast Reverse Path Forwarding (uRPF)**, which enhances security by validating that the source IP address of each incoming packet is reachable via the interface on which it arrives. This mechanism helps prevent malicious actors from injecting spoofed traffic into the network, thereby strengthening the integrity and reliability of IP communications.

Beginning with NX-OS Release 10.6(2)F, uRPF and the related Layer 3 path-validation features are supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information on IP directed broadcast, refer to Configuring HSRP.

First-hop redundancy and multichassis connectivity

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

Beginning with NX-OS Release 10.6(2)F, HSRP, VRRP, vPC, and vPC peer gateway support is available on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

- **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, refer to 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, refer to Configuring HSRP.

Multicast features

This article describes the supported multicast features on N9324C-SE1U and N9348Y2C6D-SE1U Smart switches. It explains how multicast forwards traffic from one source to multiple receivers while reducing unnecessary unicast replication and broadcast flooding.

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.

The following table identifies the first release that supports the feature on the Smart switches. The release applies to the Smart-switch platform scope; use the feature-specific topic for configuration requirements and limitations.

Table 4: Smart-switch multicast feature support by release

Feature	First supported release	Scope or restriction
IPv4 multicast routing	10.6(2)F	N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.
IGMP snooping	10.6(2)F	Supported on the Smart switches; VXLAN Anycast BGW is not supported.
Multicast route-aliveness	10.6(2)F	Supported on the Smart switches.
Multicast hitbit and route statistics	10.6(2)F	Supported on the Smart switches.
Network Load Balancing (NLB)	10.6(2)F	Supported on the Smart switches.
IPv6 multicast routing	10.6(3)F	Support begins in this release on the Smart switches.
MLD snooping	10.6(3)F	Support begins in this release; use **ipv6 mld snooping vxlan** for VXLAN VLANs.

Multicast concepts and relationships

Multicast sends one source stream to a receiver group. Layer 2 membership control limits delivery within a VLAN, while Layer 3 routing builds paths between subnets.

Table 5: Multicast terminology

Term	Definition	Relationship
Source	The device or application that sends the multicast stream.	Provides traffic that the network replicates toward interested receivers.
Receiver	A host or application that joins a multicast group.	Membership information guides Layer 2 forwarding and Layer 3 routing.
IGMP snooping	A Layer 2 function that learns receiver membership from IGMP messages.	Works within a VLAN; PIM or another multicast routing function connects routed subnets.
Distribution tree	The forwarding path from a source toward multicast receivers.	Multicast protocols build the tree so traffic avoids unrelated segments.

Supported multicast features

The supported multicast feature set for N9324C-SE1U and N9348Y2C6D-SE1U Smart switches is summarized in the support table. The related topics provide feature-specific descriptions and configuration guidance.

- Multicast route-aliveness.
- Hitbit and route statistics for packets and bytes.
- Network load balancing (NLB).
- IPv6 multicast routing and Multicast Listener Discovery (MLD) snooping beginning with NX-OS Release 10.6(3)F.

Unsupported multicast features

The following multicast features are not supported on N9324C-SE1U and N9348Y2C6D-SE1U Smart switches in NX-OS Release 10.6(2)F. A feature can be unsupported in the baseline release and supported in a later release; refer to the support table before treating a feature as permanently unsupported.

- PIM Bidir Support.
- L2 mcast.
- Multicast over GRE.
- Multicast Ingress/Egress NAT.
- MVPN.
- FPV for IPv6 multicast.
- MU/UM NAT.
- FEX.
- ND ISSU.

Multicast Layer 2 and Layer 3 interfaces

Multicast features apply to specific Layer 2 and Layer 3 interface types on Cisco N9300 Series smart switches. Identify the interface role first so that you configure multicast routing on routed interfaces or IGMP snooping on Layer 2 segments where receivers connect.

Beginning with NX-OS Release 10.6(2)F, IPv4 multicast features on the supported Layer 2 and Layer 3 interface types are available on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches. IPv6 multicast routing and Multicast Listener Discovery (MLD) snooping are supported beginning with NX-OS Release 10.6(3)F.

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, refer to 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, refer to Overview.

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.

Beginning with NX-OS Release 10.6(2)F, IPv4 Layer 2 and Layer 3 multicast is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

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

IGMP snooping

IGMP snooping software examines Layer 2 IP multicast traffic within a VLAN to discover the ports where interested receivers reside. Using the port information, IGMP snooping can reduce bandwidth consumption in a multi-access LAN environment to avoid flooding the entire VLAN. IGMP snooping tracks which ports are attached to multicast-capable routers to help the routers forward IGMP membership reports. The IGMP snooping software responds to topology change notifications. By default, IGMP snooping is enabled on the device.

Beginning with NX-OS Release 10.6(2)F, IGMP snooping is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information on IGMP snooping, refer to Configuring IGMP Snooping

MLD snooping

Multicast Listener Discovery (MLD) snooping examines MLD messages in a Layer 2 VLAN and builds a forwarding database for IPv6 multicast listeners. The switch uses that membership information to forward IPv6 multicast traffic only to ports that have interested receivers.

Beginning with NX-OS Release 10.6(3)F, MLD snooping is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Guidelines and limitations

For VXLAN VLANs, use the **ipv6 mld snooping vxlan** command to enable MLD snooping for IPv6 multicast traffic in the overlay. Review the VLAN, interface, and multicast-routing requirements before enabling the feature.

For more information on MLD snooping, refer to Configuring MLD.

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.

Beginning with NX-OS Release 10.6(2)F, the multicast consistency checker is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information on multicast consistency checker, refer to 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.

Beginning with NX-OS Release 10.6(2)F, IPv4 multicast flow counters are supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information on multicast flow counter, refer to Multicast Counters.

VXLAN features

This article describes Virtual Extensible LAN (VXLAN) concepts and feature relationships. VXLAN extends Layer 2 domains across Layer 3 infrastructure using MAC-in-UDP encapsulation to support scalable, multitenant data center networks.

VXLAN concepts and relationships

VXLAN carries tenant traffic across a Layer 3 underlay by encapsulating Layer 2 frames in UDP. The control plane distributes endpoint reachability, VNIs identify overlay segments, and VTEPs connect the overlay to the underlay.

Table 6: VXLAN terminology

Term	Definition	Relationship
VXLAN	An overlay technology that extends Layer 2 and Layer 3 services across a Layer 3 network.	Uses encapsulation to separate tenant traffic from the physical underlay.
VTEP	A VXLAN Tunnel Endpoint that encapsulates or decapsulates overlay traffic.	Connects host-facing interfaces and the VXLAN overlay to the underlay.
VNI	A VXLAN Network Identifier for an overlay segment.	L2VNI identifies Layer 2 segments; L3VNI supports routed tenant traffic.
BGP EVPN	A control plane that distributes endpoint and tenant reachability.	Provides the information used by VTEPs to forward overlay traffic.

VXLAN overlay and underlay traffic handling

The VXLAN overlay and underlay traffic handling features in NX-OS are designed to ensure efficient, scalable, and reliable data transport within VXLAN-based networks. By supporting both IPv4 and IPv6 unicast overlay traffic, the solution enables seamless communication between endpoints across the virtualized fabric. The system effectively manages BUM (broadcast, unknown unicast, and multicast) traffic, utilizing methods such as head-end replication or multicast underlay to optimize traffic delivery and bandwidth utilization.

Load balancing and redundancy are further enhanced through the use of Equal-Cost Multi-Path (ECMP) routing in both the underlay (physical network) and overlay (VXLAN layer), distributing traffic across multiple paths for improved performance and fault tolerance. VXLAN uplinks, implemented via physical or port-channel Layer 3 interfaces, connect VTEPs to the underlay network, providing robust pathways for encapsulated VXLAN packets. Together, these features deliver the high throughput, resilience, and scalability required for modern data center fabrics

Key features include:

- IPv4 and IPv6 unicast overlay traffic
- BUM (broadcast, unknown unicast, and multicast) traffic
- Underlay ECMP and overlay ECMP (L3)
- VXLAN uplinks

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, refer to Configure the Underlay and Configure VXLAN.

BUM traffic

Broadcast, unknown unicast, and multicast (BUM) traffic refers to traffic within a VXLAN segment. The traffic is typically handled by replicating packets and forwarding them to all relevant VTEPs through multicast replication in the underlay or ingress replication (head-end replication).

BUM traffic supports these underlays in the fabric and DCI:

- IPv4 unicast underlay (IR)
- IPv4 multicast underlay with PIM ASM

For more information on BUM traffic, refer to Configure the Underlay and Configure VXLAN.

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, refer to Configure the Underlay and Configure VXLAN.

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 uplinks are supported only on Ethernet and port-channel routed interfaces. SVI and Layer 3 subinterfaces are not supported as VXLAN uplinks.

For more information on VXLAN uplinks, refer to Configure VXLAN and Configure VXLAN BGP EVPN.

VXLAN overlay control plane and segmentation

The VXLAN overlay control plane and segmentation features in NX-OS provide a scalable, flexible foundation for building modern, multitenant data center networks. Leveraging BGP EVPN as the control plane, these features enable dynamic distribution of Layer 2 and Layer 3 reachability information across the VXLAN fabric, simplifying network provisioning and mobility.

Key features include:

- VXLAN Layer 2 Virtual Network Identifier (L2VNI)
- VXLAN Layer 3 Virtual Network Identifier (L3VNI)
- VXLAN BGP EVPN Multi-Site Anycast Border Gateway (BGW) and vPC BGW
- VXLAN BGP EVPN Border Spine
- VXLAN to SR-MPLS handoff

Together, these control plane and segmentation features facilitate robust tenant isolation, efficient traffic engineering, and support for large-scale, distributed data center deployments.

VXLAN BGP EVPN

VXLAN BGP EVPN provides the control plane that distributes Layer 2 and Layer 3 endpoint reachability across the VXLAN fabric.

Beginning with NX-OS Release 10.6(2)F, VXLAN BGP EVPN is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information on VXLAN BGP EVPN, refer to 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, refer to Configure VXLAN BGP EVPN.

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

VLAN-based L3VNI configuration is deprecated. Use only the **vni** *vnid* **l3** command.

For information on L3VNI, refer to Configure New L3VNI Mode.

VXLAN BGP EVPN Multi-Site Anycast BGW and vPC 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. The vPC BGW option extends this design to a vPC pair.

Beginning with NX-OS Release 10.6(2)F, VXLAN BGP EVPN Multi-Site Anycast BGW and vPC BGW are supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

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

Multicast fast convergence for vPC BGW failures

Multicast fast convergence helps restore multicast forwarding quickly when a vPC Border Gateway fails.

Beginning with NX-OS Release 10.6(2)F, multicast fast convergence for vPC Border Gateway failures is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

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, refer to Configure VXLAN.

VXLAN to SR-MPLS handoff

The VXLAN to SR-MPLS handoff feature enables seamless routing and forwarding between VXLAN overlays and MPLS Segment Routing (SR) domains, typically used in data center and WAN edge architectures. This feature acts as a gateway capability on border leaf or spine devices, facilitating the transition of routed traffic between a VXLAN-based overlay network and an MPLS SR core.

Beginning with NX-OS Release 10.6(2)F, VXLAN to SR-MPLS handoff is supported on N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For information on VXLAN to SR-MPLS handoff, refer to VXLAN to SR-MPLS handoff.

Supported features and limitations

- Coexistence of VXLAN EVPN and MPLS L3VPN (SR-MPLS) is required on the same device for DCI handoff.
- VPN label assignment supports only per-VRF label allocation. Per-prefix label allocation is not supported.
- Only EVPN Type 5 (IP prefix routes) are supported for handoff to the SR-MPLS core. Subnet (Type 2) handoff and Layer 2 extension are not supported.
- For VXLAN to SR traffic, TTL is handled in pipe mode during VXLAN decapsulation and in uniform mode during MPLS encapsulation.

  For SR to VXLAN traffic, TTL is handled in uniform mode during MPLS decapsulation and in pipe mode during VXLAN encapsulation.

  By default, QoS operates in uniform mode across all traffic flows.
- The following scale limits are supported:
  - Maximum 256 VXLAN peers
  - Maximum 1000 VRFs
- **Unsupported features**
  - SR EVPN to VXLAN handoff is not supported.
  - LDP-based MPLS handoffs are not supported.
  - TRM multicast is not supported.
  - vPC, VMCT, and pMCT configurations are not supported.
  - VPN label decap statistics are available; label switch router (LSR) and adjacency statistics are not supported.
  - Priority Flow Control (PFC) is not supported in DCI handoff mode.
  - Route leaking or VRF import/export between VXLAN and MPLS domains is not supported; only same-VRF handoff is allowed.
  - During a failure or shutdown of the NVE (VXLAN) interface, next-hop resolution falls back to the MPLS underlay. This behavior is expected for resiliency.

VXLAN endpoint connectivity and high availability

The VXLAN endpoint connectivity and high-availability features in NX-OS ensure robust and resilient network access for devices within a VXLAN fabric.

Key features include:

- VXLAN standalone
- vPC and vPC Fabric Peering
- Distributed Anycast Gateway

VXLAN standalone, vPC, and vPC fabric peering

Standalone, vPC, and vPC fabric peering are methods for establishing network connectivity and redundancy for VTEPs.

VXLAN standalone, vPC, and vPC fabric peering are supported only as leaf or border leaf.

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

  For information on standalone or vPC VTEP, refer to Configure VXLAN.

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

- **vPC fabric peering** establishes underlay routing between vPC pairs. For more information on vPC fabric peering, refer to 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.

Distributed Anycast Gateway supports the **fabric forwarding anycast-mode** command. Use this configuration only on a VXLAN VLAN with an SVI configured or operational and with a global Fabric Anycast MAC configured in the system.

For more information on distributed Anycast Gateway, refer to Distributed Anycast Gateway and Configure VXLAN.

Multicast and broadcast optimization

The multicast and broadcast optimization features in NX-OS enhance the efficiency and scalability of VXLAN networks by intelligently managing the delivery of broadcast, multicast, and unknown unicast (BUM) traffic. Utilizing options such as IPv4 unicast underlay with ingress replication (IR) and IPv4 multicast underlay with PIM ASM, the network can optimize how multicast and broadcast packets are distributed, minimizing unnecessary traffic and improving overall performance.

Specialized capabilities like the multicast underlay BUD node allow selective forwarding of multicast streams, while IGMP snooping ensures that multicast traffic is only delivered to interested receivers, reducing bandwidth consumption on the network. ARP suppression and ND suppression further optimize broadcast domains by intercepting and responding to address-resolution requests locally, preventing unnecessary flooding across the VXLAN segment. Additionally, TRMv4 L3 mode (Tenant Routed Multicast) enables efficient and scalable multicast routing across different tenant segments within the VXLAN fabric.

Key features include:

- IPv4 unicast underlay (IR)
- IPv4 multicast underlay with PIM ASM
- BUD node
- IGMP snooping
- ARP suppression
- ND suppression beginning with NX-OS Release 10.6(3)F
- TRMv4 L3 mode (Tenant Routed Multicast)

Together, these features provide robust mechanisms for traffic optimization, bandwidth conservation, and enhanced multicast scalability in large-scale, multitenant data center environments.

Table 7: Smart-switch VXLAN multicast and broadcast feature support by release

Feature	First supported release	Scope or restriction
IPv4 unicast underlay with ingress replication	10.6(2)F	Supported on the Smart switches; ingress replication between DCI peers is supported.
IPv4 multicast underlay with PIM ASM	10.6(2)F	Multicast underlay for fabric peers is supported.
Multicast underlay BUD node	10.6(2)F	Supported on the Smart switches.
IGMP snooping	10.6(2)F	Not supported with Anycast BGW.
ARP suppression	10.6(2)F	Use with the VXLAN BGP EVPN control plane.
ND suppression	10.6(3)F	Plain BGP EVPN only; not supported with Multi-Site, vMCT, IRB, Centralized Gateway, Firewall Clustering, or vPC.
TRMv4 L3 mode	10.6(2)F	Supported on vPC leaf, vPC fabric peering leaf, Anycast BGW, and standalone leaf devices. TRMv6 and Data MDT are not supported.

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.

Ingress Replication between DCI peers across the core is supported.

For more information on IPv4 unicast underlay, refer to Configure the Underlay and Configure VXLAN.

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.

Multicast underlay for fabric peers is supported.

For more information on IPv4 multicast underlay with PIM ASM, refer to Multicast Routing in the VXLAN Underlay and Configure VXLAN.

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 the multicast underlay BUD node, refer to Configure Bud Node and Configure VXLAN.

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.

IGMP snooping does not support Anycast BGW.

For more information on IGMP snooping, refer to Configure VXLAN.

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, refer to Configure VXLAN BGP EVPN.

ND suppression

Neighbor Discovery (ND) suppression reduces IPv6 multicast flooding across the VXLAN core by learning IPv6 address and MAC bindings and responding to eligible Neighbor Solicitation requests locally.

Beginning with NX-OS Release 10.6(3)F, ND suppression is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Use ND suppression only with plain BGP EVPN. Do not use it with Multi-Site, vMCT, IRB, Centralized Gateway, Firewall Clustering, or vPC. The associated VNI must have ARP suppression enabled, an operational SVI, and both IPv4 and IPv6 addresses configured.

For more information on ND suppression, refer to 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.

TRMv4 and TRMv4 L3 are supported on vPC leaf, vPC fabric peering leaf, Anycast BGW, and standalone leaf devices. TRMv6 and Data MDT are not supported.

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

VXLAN services and operational tools

Use the VXLAN services and operational tools in this topic after you establish the overlay. The sections follow the deployment journey from service configuration and connectivity extension to monitoring and troubleshooting.

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.

Beginning with NX-OS Release 10.6(2)F, DHCP relay in VXLAN BGP EVPN is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

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

DSVNI

VXLAN EVPN with downstream VNI provides these 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

Beginning with NX-OS Release 10.6(2)F, DSVNI is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

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

BGW advertisement

The **advertise-pip** command advertises the Border Gateway (BGW) primary IP (PIP) as the next hop for routes learned from other sites in a multi-site deployment. This behavior helps maintain symmetric traffic flows for stateful services.

BGW advertisement using PIP toward the fabric and Data Center Interconnect (DCI) advertises the PIP as the next hop for external routes toward the local fabric and the DCI.

BGW advertisement supports the **advertise-pip** and **fabric-advertise-pip l3** commands.

Beginning with NX-OS Release 10.6(2)F, BGW advertisement is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information, refer to Configure vPC Multi-Homing, Advertise Using PIP Towards Fabric, and Configure VXLAN.

VXLAN counters

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

VXLAN counters support VXLAN peer-based total packet and byte counters and VNI-based total packet and byte counters. Peer counters and per-peer-per-VNI counters are not supported.

Beginning with NX-OS Release 10.6(2)F, VXLAN counters are supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information on VXLAN counters, refer to Configure VXLAN.

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.

VXLAN NGOAM supports VTEP and host reachability checks.

Beginning with NX-OS Release 10.6(2)F, VXLAN NGOAM is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information on VXLAN NGOAM, refer to VXLAN OAM or VXLAN NGOAM.

Proportional multipath for VNF

Proportional multipath for Virtual Network Functions (VNFs) enables a VXLAN EVPN fabric to advertise multiple next hops for a service network. The forwarding plane can then use the available paths to distribute traffic across the links that connect the fabric to the VNF.

Beginning with NX-OS Release 10.6(2)F, EVPN ECMP for VNF and Proportional VNF (PVNF) is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

Supported capability

The Smart switches support 64-way ECMP for EVPN VNF/PVNF.

Guidelines and limitations

Use the VNF-specific support information when evaluating this feature. The feature has separate requirements and limitations for BGP route advertisement, mixed multipath, VNF reachability, vPC, and VXLAN IPv6 underlays.

For the complete configuration procedures, support matrix, and guidelines, refer to Configuring Proportional Multipath for VNF.

Unsupported VXLAN features

These VXLAN functionalities are not supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart 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: IPv6 IR underlay (DCI), EVPN multi-site storm control, and CloudSec.
- Advanced multicast handling features: Multicast Listener Discovery (MLD) snooping, Neighbor Discovery (ND) suppression (supported starting from NX-OS Release 10.6(3)F , TRMv6, TRM Layer 2 mode, 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

Segment Routing Multiprotocol Label Switching

This article describes the supported Segment Routing Multiprotocol Label Switching (SR-MPLS) features and limitations for the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches. SR-MPLS forwards traffic using short, preassigned labels and supports services such as Traffic Engineering, Quality of Service, and Virtual Private Networks (VPNs).

Beginning with NX-OS Release 10.6(2)F, the Segment Routing Layer 3 VPN feature with an SR-MPLS underlay is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

SR-MPLS concepts and relationships

SR-MPLS forwards packets through an MPLS network using a sequence of segments or labels. The underlay supplies reachability, while Segment Routing policies and VPN services use that reachability for traffic steering or tenant isolation.

Table 8: SR-MPLS terminology

Term	Definition	Relationship
Segment	An instruction that identifies a destination or forwarding behavior.	Segments combine to express the path that traffic should follow.

Term	Definition	Relationship
MPLS label	A short identifier used to forward a packet through an MPLS domain.	Represents the forwarding instruction carried in the packet.
SR-MPLS underlay	The routed MPLS network that transports labeled traffic.	Provides the transport used by services such as SR Layer 3 VPN.

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.

Beginning with NX-OS Release 10.6(2)F, this feature is supported on the N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.

For more information, refer to Configuring Segment Routing.

SR-MPLS fabric and VPN features

The supported fabric and VPN capabilities are grouped by their role in the SR-MPLS design.

- 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.

Interfaces and forwarding features

The supported interface and forwarding capabilities are:

- 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.

MPLS encapsulation statistics

Beginning with NX-OS Release 10.6(3)F, MPLS encapsulation statistics are supported on Cisco N9324C-SE1U Smart switches.

MPLS encapsulation statistics are not available for MPLS packets that egress on SVI interfaces.

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).

- Static MPLS is not supported on N9324C-SE1U and N9348Y2C6D-SE1U Smart switches.