Revised: March 10, 2026

VXLAN

VXLAN

This article outlines the set of Virtual Extensible LAN (VXLAN) features for the Cisco N9324C-SE1U and N9348Y2C6D-SE1U switches, enabling the creation of highly scalable, agile, and multitenant data center networks. VXLAN extends Layer 2 domains over Layer 3 infrastructure using MAC-in-UDP encapsulation, which supports seamless tenant isolation, workload mobility, and efficient resource utilization across modern data center fabrics.

VXLAN Overlay Control Plane and Segmentation

The VXLAN Overlay Control Plane and Segmentation features in Cisco 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 BGW
- VXLAN BGP EVPN Border Spine
- SR-MPLS to VXLAN 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 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 and 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.

Guidelines

- VLAN based L3VNI configuration is deprecated. Only **vni** *vnid* **l3** command is supported.

  For information on L3VNI, see Configure New L3VNI Mode.

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

SR MPLS to VXLAN Handoff feature

The SR MPLS to VXLAN Handoff feature enables seamless routing and forwarding between MPLS Segment Routing (SR) domains and VXLAN overlays, 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 an MPLS SR core and a VXLAN-based overlay network.

For information on SR MPLS to VXLAN Handoff, see SR MPLS to VXLAN Handoff feature.

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 Cisco 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** 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.
- **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 and Configure VXLAN.

Guidelines

- Distributed Anycast Gateway feature supports the **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

VXLAN Overlay and Underlay Traffic Handling

The VXLAN Overlay and Underlay Traffic Handling features in Cisco 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, 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, see Configure the Underlay and Configure VXLAN.

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 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, see 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.

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

Guidelines

- BUM traffic supports these underlay on Fabric and DCI:
  - IPv4 unicast underlay (IR)
  - IPv4 multicast underlay with PIM ASM
- VXLAN uplinks is supported on only Ethernet and Port-channel routed interfaces.
- VXLAN uplinks do not support SVI or L3-subinterfaces.

Multicast and Broadcast Optimization

The Multicast and Broadcast Optimization features in Cisco 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 further optimizes broadcast domains by intercepting and responding to ARP requests locally, preventing the need to flood these requests across the entire 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 Multicast Underlay
- BUD Node
- IGMP Snooping
- ARP Suppression
- 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.

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 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.

For more information on IPv4 multicast underlay with PIM ASM, see 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 Multicast underlay BUD node, see 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.

For more information on IGMP snooping, see 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, 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.

Guidelines for Cisco N9324C-SE1U and Cisco N9348Y2C6D-SE1U smart switches

- IGMP snooping do not support Anycast BGW
- Do not support TRMv6 and Data MDT.
- TRMv4 and TRMv4 L3 supports these functionalities on vPC leaf, vPC fabric peering leaf, Anycast BGW, and standalone leaf.
  - 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.

Advanced services and tools

The advanced services and tools in Cisco NX-OS empower administrators with enhanced operational flexibility, monitoring, and automation for VXLAN deployments. Features such as DHCP Relay in VXLAN enable seamless IP address assignment for hosts across different subnets by forwarding DHCP requests between VXLAN overlays and external servers. VXLAN Counters provide detailed traffic statistics for encapsulated and decapsulated packets, aiding in real-time performance analysis and troubleshooting.

VXLAN NGOAM (Overlay Operations, Administration, and Maintenance) delivers proactive monitoring and diagnostic tools for overlay network paths, ensuring high reliability and rapid issue resolution. DSVNI (Downstream VNI and Route Leak) supports complex multitenant environments by enabling asymmetric VNI communication and controlled route leaking between isolated VXLAN segments. Additionally, BGW Advertisement features—such as advertise-pip and advertise using PIP towards fabric or Data Center Interconnect (DCI)—enhance multi-site connectivity and optimize route propagation across distributed data center fabrics.

Key features include:

- DHCP Relay in VXLAN
- VXLAN Counters

- VXLAN NGOAM (Overlay OAM/Monitoring)
- DSVNI (Downstream VNI and Route Leak)
- BGW Advertisement (advertise-pip, advertise using PIP towards fabric/DCI)

Together, these advanced services and tools streamline the management, visibility, and scalability of VXLAN-based networks.

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.

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.

For more information on VXLAN counters, see 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.

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

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.

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 and Configure VXLAN sections.

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 and Configure VXLAN sections.

Guidelines

- VXLAN counters supports VXLAN peer-based total packet/byte counters and VNI based total packet and byte counters.
- VXLAN counters do not support peer counters or per-peer-per-vni counters.
- VXLAN NGOAM supports VTEP and host reachability
- BGW advertise supports **advertise-pip** and **fabric-advertise-pip l3** commands

Unsupported VXLAN features

These VXLAN functionalities are not supported on Cisco N9324C-SE1U and 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: IPv6 IR underlay (DCI), EVPN multi-site storm control, and CloudSec.
- 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