# Feature History for BGP EVPN VXLAN


• Feature History for BGP EVPN VXLAN, on page 1



# Feature History for BGP EVPN VXLAN


This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

| Release | Feature | Feature Information |
|---|---|---|
| Cisco IOS XE Fuji 16.9.1 | EVPN VXLAN Overlay Network for IPv4 Bridged Traffic | EVPN VXLAN overlay network for IPv4 bridged traffic is a Layer 2 overlay network that allows host devices within the same subnet to send IPv4 bridged traffic to each other using a Layer 2 virtual network instance (VNI). |
| | EVPN VXLAN Overlay Network for IPv4 Routed Traffic | EVPN VXLAN overlay network for IPv4 routed traffic is a Layer 3 overlay network that allows host devices in different Layer 2 networks to send IPv4 routed traffic to each other using a Layer 3 VNI and an IP VRF. |
| | Layer 2 Broadcast, Unknown Unicast, and Multicast (BUM) Traffic Forwarding using Underlay Multicast | Multi-destination Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic in an EVPN VXLAN network is replicated through a multicast group in the underlay network and forwarded to all the endpoints of the network. |
| | Leaf Functionality | A leaf switch sits on the edge of a BGP EVPN VXLAN fabric and is connected to the host or access devices. It functions as a virtual tunnel end point (VTEP) and performs encapsulation and decapsulation. |
| | EVPN VXLAN Integrated Routing and Bridging | EVPN VXLAN integrated and bridging (IRB) allows the VTEPs in a VXLAN network to forward both Layer 2 or bridged traffic and Layer 3 or routed traffic. It is implemented as symmetric and asymmetric IRB. |
| | EVPN VXLAN Distributed Anycast Gateway | EVPN VXLAN distributed anycast gateway is a default gateway addressing mechanism that enables the use of the same gateway IP address across all the leaf switches that are part of a VXLAN network.<br>Support was introduced with manual MAC address configuration on the Layer 2 VNI VLAN's switch virtual interface (SVI) on all VTEPs as the only method to enable the feature. |
| | DHCP Relay for IPv4 Traffic in BGP EVPN VXLAN Fabric | The VTEP in a BGP EVPN VXLAN fabric is configured as a DHCP relay agent to provide DHCP relay services for IPv4 traffic in a multi-tenant VXLAN environment. |

| Release | Feature | Feature Information |
|---|---|---|
| Cisco IOS XE Gibraltar 16.11.1 | EVPN VXLAN Overlay Network for IPv6 Bridged Traffic | EVPN VXLAN overlay network for IPv6 bridged traffic is a Layer 2 overlay network that allows host devices within the same subnet to send IPv6 bridged traffic to each other using a Layer 2 VNI. |
| | EVPN VXLAN Overlay Network for IPv6 Routed Traffic | EVPN VXLAN overlay network for IPv6 routed traffic is a Layer 3 overlay network that allows host devices in different Layer 2 networks to send IPv6 routed traffic to each other using a Layer 3 VNI and an IP VRF. |
| | Layer 2 Broadcast, Unknown Unicast, and Multicast (BUM) Traffic Forwarding using Ingress Replication | Ingress replication is a unicast approach to handle multi-destination Layer 2 BUM traffic in an EVPN VXLAN network. It involves an ingress device replicating every incoming BUM packet and sending them as a separate unicast to the remote egress devices. |

| Release | Feature | Feature Information |
|---|---|---|
| Cisco IOS XE Gibraltar 16.12.1 | MAC Aliasing for EVPN VXLAN Distributed Anycast Gateway | MAC aliasing allows the leaf switches in an EVPN VXLAN network to advertise the MAC addresses of their Layer 2 VLAN's SVI as the gateway MAC address to all the other leaf switches in the network.<br>MAC aliasing removes the need to explicitly configure the same MAC address on the Layer 2 VNI VLAN's SVI on all VTEPs in order to enable distributed anycast gateway. |
| | Border Leaf Functionality | A border leaf switch is a leaf switch in a BGP EVPN VXLAN fabric that enables external connectivity with other Layer 2 and Layer 3 networks by acting as the connecting node between the two networks. |
| | Autonomous System Number Rewrite | The **rewrite-evpn-rt-asn** command was introduced to enable the rewrite of the autonomous system number (ASN) of the EVPN route target that originates from the current autonomous system with the ASN of the target eBGP EVPN peer. |
| | VRF-Lite Border Leaf Handoff | VRF-Lite border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 3 external connectivity with a VRF-Lite network through a border leaf switch. |
| | MPLS Layer 3 VPN Border Leaf Handoff | MPLS Layer 3 VPN border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 3 external connectivity with an MPLS Layer 3 VPN network through a border leaf switch. |
| | IEEE 802.1Q Border Leaf Handoff | IEEE 802.1Q border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with an IEEE 802.1Q network through a border leaf switch. |
| | Access Border Leaf Handoff | Access border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with an Access network through a border leaf switch. |
| | VPLS over MPLS Border Leaf Handoff | VPLS over MPLS border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with a VPLS over MPLS network through a border leaf switch. |

| Release | Feature | Feature Information |
|---|---|---|
| Cisco IOS XE Amsterdam 17.1.1 | Spine Functionality | A spine switch acts as the connecting node between all the leaf switches in a BGP EVPN VXLAN fabric, forwards the traffic between the leaf switches and provides redundancy to the network. |
| | Border Spine Functionality | A border spine switch in a BGP EVPN VXLAN fabric enables external connectivity with other Layer 2 and Layer 3 networks by acting as the connecting node between the two networks. |
| | Layer 3 Tenant Routed Multicast for IPv4 Traffic | Layer 3 tenant routed multicast (TRM) for IPv4 traffic enables multicast forwarding for IPv4 traffic in a BGP EVPN VXLAN fabric. It provides multi-tenancy-aware multicast forwarding between senders and receivers within the same subnet or different subnets, locally or across VTEPs.<br>You can configure TRM with PIM sparse mode (PIM-SM) using anycast RP mode and TRM with PIM source specific mode (PIM-SSM). |
| | VRF-Lite Border Spine Handoff | VRF-Lite border spine handoff in a BGP EVPN VXLAN fabric allows Layer 3 external connectivity with a VRF-Lite network through a border spine switch. |
| | MPLS Layer 3 VPN Border Spine Handoff | MPLS Layer 3 VPN border spine handoff in a BGP EVPN VXLAN fabric allows Layer 3 external connectivity with an MPLS Layer 3 VPN network through a border spine switch. |
| | IEEE 802.1Q Border Spine Handoff | IEEE 802.1Q border spine handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with an IEEE 802.1Q network through a border spine switch. |
| | Access Network Border Spine Handoff | Access border spine handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with an Access network through a border spine switch. |
| | VPLS over MPLS Border Spine Handoff | VPLS over MPLS border spine handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with a VPLS over MPLS network through a border spine switch. |
| | VXLAN-Aware Flexible Netflow | VXLAN-aware Flexible Netflow captures the VXLAN flow information for both bridged and routed traffic. |
| | BGP EVPN VXLAN MIB support | Support was introduced for the MIB. |

| Release | Feature | Feature Information |
|---|---|---|
| Cisco IOS XE Amsterdam 17.2.1 | ARP and IPv6 Neighbor Discovery Flooding Suppression | Flooding suppression avoids the flooding of ARP and IPv6 neighbor discovery packets over the VXLAN network to the local and remote host or access devices. |
| | EVPN VXLAN Centralized Default Gateway | EVPN VXLAN Centralized Default Gateway allows a single VTEP in the EVPN VXLAN network to act as the Layer 3 gateway for all the Layer 2 VNIs in the network. |
| Cisco IOS XE Amsterdam 17.3.1 | Layer 3 Tenant Routed Multicast for IPv6 Traffic | Layer 3 TRM for IPv6 traffic enables multicast forwarding for IPv6 traffic in a BGP EVPN VXLAN fabric. It provides multi-tenancy-aware multicast forwarding between senders and receivers within the same subnet or different subnets, locally or across VTEPs.<br>Support was introduced to configure TRM with PIM-SSM and TRM with PIM-SM. |
| | Enhanced RP Functionality for Layer 3 TRM for IPv4 and IPv6 traffic | Enhancement was introduced for the configuration of an RP in the overlay network for TRM with PIM-SM. The enhancement allows you to configure an RP on a single or multiple VTEPs inside the BGP EVPN VXLAN fabric or on a device outside the fabric. |
| | Interworking of Layer 3 TRM with MVPN Networks for IPv4 Traffic | Interworking of Layer 3 TRM with MVPN Networks allows you to forward IPv4 Layer 3 multicast traffic between sources and receivers of an EVPN VXLAN network and an MVPN network. |
| | Broadcast, Unknown Unicast, and Multicast Traffic Rate Limiting | BUM Traffic Rate Limiting allows you to use a policer and set the flood rate limit of the BUM traffic in the network to a predefined value. |
| Cisco IOS XE Bengaluru 17.6.1 | Private VLANs (PVLANs) in a BGP EVPN VXLAN Fabric | BGP EVPN VXLAN allows you to extend a PVLAN across the VTEPs in an EVPN fabric overlay. The extension uses isolated and community VLANs to provide features like subnet management and sub segmentation of a broadcast domain in the fabric. |
| Cisco IOS XE Bengaluru 17.6.1 | EVPN VXLAN Multi-homing in Single-active Redundancy Mode | Multi-homing provides redundancy in the connection between a customer edge (CE) device and a VTEP by connecting the customer network with multiple VTEPs in an EVPN VXLAN network.<br>In single-active redundancy mode, only one VTEP, among a group of VTEPs that are attached to the particular ethernet segment, is allowed to forward traffic to and from that ethernet segment.<br>Multi-homing in single-active redundancy mode was introduced only in the form of dual-homing, allowing a CE device to be connected to two VTEPs. |

| Release | Feature | Feature Information |
|---|---|---|
| Cisco IOS XE Bengaluru 17.6.1 | Route Leaking between EVPN and Global Routing Table | The route leaking feature provides the import and export of routes between an EVPN-enabled VRF and the Global Routing Table (GRT). By enabling the import and export of the EVPN learned routes into the GRT, the border node provides access to services such as DNS or DHCP to the hosts in the fabric. |
| Cisco IOS XE Bengaluru 17.6.2 | Data MDT Support for L3 TRM | Data MDTs are purpose built underlay MDTs to provide optimized forwarding in the MVPN and EVPN core. |
| Cisco IOS XE Cupertino 17.8.1 | Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 traffic | Optimized Layer 2 Overlay Multicast forwards multicast traffic within the Layer 2 Virtual Network Instance (L2VNI). IPv4 and IPv6 Layer 2 overlay multicast traffic is optimized both at the access level and within the VXLAN fabric. |
| Cisco IOS XE Cupertino 17.9.1 | BGP EVPN VXLAN over IPsec | An IPsec-based underlay network securely transports the VXLAN-encapsulated packets between the source and destination VTEPs. Securing the BGP EVPN VXLAN data traffic using IPsec tunnel encrypts the data and maintains data integrity.<br>Support for this feature was introduced only on the Cisco Catalyst 9300X Series Switches. |
| Cisco IOS XE Dublin 17.10.1 | BGP EVPN VXLAN with IPv6 in the Underlay (VXLANv6) | Introduced support for IPv6 addressing in the underlay of a BGP EVPN VXLAN fabric. In a new deployment, you can build your BGP EVPN VXLAN fabric with IPv6 underlay. For an existing BGP EVPN VXLAN fabric with IPv4 underlay, you can seamlessly migrate to an IPv6 or dual stack underlay. |

| Release | Feature | Feature Information |
|---|---|---|
| Cisco IOS XE Dublin 17.11.1 | EVPN Microsegmentation | BGP EVPN VXLAN integrates Cisco TrustSec to provide microsegmentation and end-to-end access control with the propagation of the security group tag (SGT). Using security group-based access control lists (SGACLs), you can control the operations that a user can perform, based on the security group assignments and destination resources in a VXLAN campus fabric. |
| | EVPN VXLAN Multi-Homing | BGP EVPN is enhanced to restrict the ethernet segment operations to the EVPN-controlled VLANs on the trunk port. This allows traditional Layer 2 domains to co-exist with the Layer 2 VNI-enabled VLANs at access layer. It also allows selective VLAN migration to overlay VXLAN segmentation. |
| | Dynamic BGP Peering for EVPN | Support for BGP dynamic neighbor sessions is extended to the L2VPN EVPN address family. |
| | EVPN Route Map Support | The Leaf, Spine and Border nodes of a BGP EVPN fabric now support route map for the L2VPN address family.<br>With route map support, the BGP attributes and their values can be modified to customize the routing policy based on the requirement. The routing policy can be applied for both inbound and outbound EVPN routes. |
| Cisco IOS XE Dublin 17.11.99sw | TRM over BGP EVPN VXLANv6 | TRM over BGP EVPN VXLANv6 follows the industry standard IETF RFCs 6513 and 6514 to support rich multicast applications over the BGP EVPN VXLAN fabric network.<br>The Cisco Catalyst 9000 series switches provide flexibility to transport IPv4 multicast applications in overlay networks while the underlay network can be built on single-stack IPv6 only. |
| Cisco IOS XE Dublin 17.12.1 | BGP EVPN VRF Auto RD and Auto RT | BGP EVPN Layer 3 overlay VRF configuration is simplified with the introduction of new CLIs to auto generate the route distinguisher (RD) and route target (RT) for a VRF. |
| | Dynamic ARP inspection (DAI) and DHCP Rogue Server Protection in BGP EVPN VXLAN fabric | BGP EVPN VXLAN fabric now supports ARP inspection and DHCP Rogue Server Protection. To configure these features, enable ARP inspection and DHCP Snooping on the VTEPs in the corresponding VLANs of the EVPN VXLAN fabric. |

| Release | Feature | Feature Information |
|---|---|---|
| Cisco IOS XE Dublin 17.12.2 | TRM over BGP EVPN VXLANv6 | TRM over BGP EVPN VXLANv6 follows the industry standard IETF RFCs 6513 and 6514 to support rich multicast applications over the BGP EVPN VXLAN fabric network.<br>The Cisco Catalyst 9000 series switches provide flexibility to transport IPv4 multicast applications in overlay networks while the underlay network can be built on single-stack IPv6 only. |
| | DHCP Relay on VTEPs in Layer 2 Overlay | In a BGP EVPN VXLAN Layer 2 overlay, DHCP Relay is configured on the centralized gateway (CGW). The CGW is configured either on an external router or on the Layer 2 border VTEP. |

Use Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to Cisco Feature Navigator.

Feature History for BGP EVPN VXLAN