



BGP EVPN over MultiProtocol Label Switching

The BGP EVPN over MultiProtocol Label Switching (MPLS) feature utilizes the functionality defined in RFC 7432 (BGP MPLS-based Ethernet VPN), to achieve EVPN over MPLS functionality between Provider Edge (PE) and Customer Edge (CE) devices.

- [Feature Information for BGP EVPN Over MPLS, on page 1](#)
- [Information about BGP EVPN over MultiProtocol Label Switching, on page 1](#)
- [How to Configure BGP EVPN over MultiProtocol Label Switching, on page 6](#)
- [Configuration Examples for BGP EVPN over MPLS, on page 9](#)
- [Additional References for BGP EVPN over MultiProtocol Label Switching, on page 9](#)

Feature Information for BGP EVPN Over MPLS

Feature Name	Releases	Feature Information
BGP EVPN Over MPLS	Cisco IOS XE Fuji 16.9.x	The BGP EVPN over MPLS feature utilizes the functionality defined in RFC 7432, to achieve EVPN over MPLS functionality between Provider Edge (PE) and Customer Edge (CE) devices.

Information about BGP EVPN over MultiProtocol Label Switching

BGP MPLS based Ethernet VPN (EVPN) Overview

EVPN (RFC 7432) addresses the following requirements:

- PE node redundancy with load-balancing based on L2/L3/L4 flows from CE to PE.
- Flow-based multi-pathing of traffic from local PE to remote PEs across core and vice-versa
- Geo-redundant PE nodes with optimum unicast forwarding.

- Flexible redundancy grouping, where a PE can be a member of multiple redundancy groups, each containing a different set of CEs.

EVPN Building Blocks

There are three fundamental building blocks for EVPN technology, EVPN Instance (EVI), Ethernet Segment (ES), EVPN BGP routes and extended communities:

- EVI is a VPN connection on a PE router. It is the equivalent of IP VPN Routing and Forwarding (VRF) in Layer 3 VPN. It is also known as MAC-VRF.
- ES is a connection with a customer site (device or network) and is associated with access-facing interfaces. Access-facing interfaces are assigned unique IDs that are referred to as Ethernet Segment Identifiers (ESI). A site can be connected to one or more PEs. The ES connection has the same ESI in each PE connected to the site.
- RFC 7432 defines routes and extended communities to enable EVPN support. In Cisco IOS XE Fuji 16.8.x Software Release, Route Type 2 and Route Type 3 are supported.

In BGP MPLS-based EVPN, an EVI is configured for every PE device for each customer associated with the PE device. In this case, a customer is any customer edge device that is attached to the PE device. The CE device can be a host, a switch or a router. Each EVI has a unique Route Distinguisher (RD) and one or more Route Targets (RT).

For EVPN Single-Homing feature, a CE device is attached to a single PE device and has an Ethernet Segment with ESI=0.


Service Interfaces

The following are types of EVPN VLAN service interfaces:

VLAN-based Service Interface

In VLAN-based service interface, each VLAN is associated to one bridge domain and one EVI.

Figure 1: VLAN-Based Service Interface



For VLAN-based Service Interface, Type 1 Route Distinguisher, a unique number used to distinguish identical routes in different VRFs, is used for EVIs as recommended by the RFC 7432. The Route Distinguishers and Router Targets, which are used to share routes between different VRFs, are autogenerated to ensure unique Route Distinguisher numbers across EVIs.

VLAN Bundle Service Interface

In VLAN Bundle Service Interface, multiple VLANs share the same bridge table.

Figure 2: VLAN Bundle Service Interface




Each EVPN instance corresponds to multiple broadcast domains maintained in a single bridge table per MAC-VRF. For VLAN Bundle Service Interface service to work, MAC addresses must be unique across all VLANs for an EVI.

VLAN-Aware Bundle Service Interface

For VLAN-aware Bundle Service Interface, each VLAN is associated with one bridge domain, but there can be multiple bridge domains associated with one EVI.

Figure 3: VLAN-Aware Bundle Service Interface



An EVPN instance consists of multiple broadcast domains where each VLAN has one bridge table. Multiple bridge tables (one per VLAN) are maintained by a single MAC-VRF that corresponds to the EVPN instance.

BGP EVPN over MPLS Inter-AS and Prefix SID

BGP EVPN over MPLS Route Target 5 Inter-Autonomous Systems (Inter-AS) Option C is used to advertise the site prefixes across the network that forms the overlay. BGP Segment Routing Prefix segment identifiers (SID) is used to advertise the Node SID of the PE device sourcing the prefix across the network that forms the underlay.

Inter-AS Option C for EVPN

The BGP EVPN over MPLS feature supports InterAS option C for EVPN session spanning across AS domains. When the session spans across multiple AS domains, the next hop of the BGP EVPN update that is set by the egress PE remains unchanged across the AS domains, and is be the next hop for the underlay for ingress (headend) PE. Underlay could be either established through BGP Labeled Unicast (BGP-LU) across AS domains or by a centralised controller.

Route Type 5 – IP Prefix

BGP EVPN over MPLS feature implements Route Type 5, as defined by RFC 7432.

The BGP EVPN Route Type 5 update carries IP prefixes and includes a tunnel encapsulation attribute that indicates the VPN session encapsulation. Absence of tunnel encapsulation attribute in BGP EVPN Route Target 5 update indicates MPLS VPN encapsulation for the BGP EVPN Route Type 5 VPN prefix.

Route Type 5 is used to advertise prefixes independently of the MAC advertisement routes, and possible prefix next-hops in the overlay topology, namely -ESI, IRB IP address, Floating IP address.

Route type 5 implementation has the following features:

- Prefix advertisement is not linked to MAC mobility.
- Network Virtual Interfaces (NVE) that have EVPN configured but do not support the optional prefix-advertisement route, can easily identify the route and ignore it without processing the route.
- When selecting routes, MAC information is not compared by BGP.
- Flexible overlay next-hop (IRB, floating IP, ESI) can be configure to address different end-point types.

The below table shows the format of the prefix-advertisement:

Table 1: Prefix Advertisement Route Format

Field	Value	LEngth (Octets)
Route Type	0x05	1
Length	25	1
EVI RD	Type 1 (IPv4Addr) RD unique across all EVIs on the PE.	8
ESI	Ethernet Segment Identifier	10
Ethernet Tag	0 or valid Ethernet Tag	4
IP Address	Gateway IP Address	4 or 16
Label	Valid MPLS label allocated per [EVI, ESI, EtherTag] tuple	3
EVI RT	Type 0 (2byteAS) route target	8

Importing IP Routes to EVPN

To import locally sourced (redistributed) and provider edge to customer edge BGP IPv4 and IPv6 routes into EVPN, add these routes to VPNv4 table using the following configuration:

```
router bgp
address-family {ipv4 | ipv6} vrf <vrfname>
advertise l2vpn evpn
```

To import VPNv4 or VPNv6 routes learned from a neighbour into EVPN, use the following configuration:

```
router bgp
address-family l2vpn evpn
import {vpn4 | vpn6} unicast
```

To add EVPN routes learned from an EVPN neighbor to EVPN the table, do one of the following:

- configure `no bgp default route-target filter` command in the router BGP configuration
- add a `route-target` in EVPN route matching the stitching configuration under VRF address-family

Route Reoriginate

When the EVPN prefix is imported into the VPN VRF table, a new VPN prefix may be reoriginated. Updates may be sent to another VPN neighbor based on the configuration. The reoriginated prefix is the imported VPN prefix where the route target extended communities are replaced with a new set of route target extended communities.

To configure route reorigination, use the following configuration:

```
enable
configure terminal
router bgp 100
address-family vpn4|vpn6
import l2vpn evpn [re-originate [stitching-rt]]
```

When the **import {vpn4 | vpn6} unicast** command is configured under the L2VPN EVPN address-family configuration mode and if the route targets of the VPNv4 or VPNv6 route are within the export stitching route targets, the reoriginated EVPN prefix retains the same route targets. Otherwise, route targets of reoriginated EVPN prefix are swapped using export stitching route target configuration.

When the **import {vpn4 | vpn6} unicast re-originate** command is configured under the L2VPN EVPN address-family, the reoriginated EVPN prefix route target is replaced by the export stitching route targets.

When the **import {vpn4 | vpn6} unicast** command is not configured, then remote VPNv4 or VPNv6 prefix is not imported into EVPN table.

When the **import l2vpn evpn** command is configured under the VPNv4 or VPNv6 address family and if the route targets of the EVPN prefix are within the subset of normal VRF export route targets in VRF, the imported VPNv4 or VPNv6 prefix retains the same route targets. Otherwise, the imported VPNv4 or VPNv6 route targets are swapped with the VRF export route targets and advertised to VPNv4/VPNv6 neighbors.

When the **import l2vpn evpn re-originate** command is configured under the VPNv4 or VPNv6 address-family, the route targets of the imported VPNv4 or VPNv6 prefix are replaced with the VRF export route targets and advertised to VPNv4/VPNv6 neighbors.

When the **import l2vpn evpn re-originate stitching-rt** command is configured, the imported VPN prefix route target is replaced by the VRF import stitching route target and advertised to VPNv4 or VPNv6 neighbors.



Note When the **import l2vpn evpn** command is not configured, the imported VPN prefix is not advertised to VPN neighbors.

EVPN Encapsulation

You can choose the EVPN encapsulation using the neighbour EVPN neighbor configuration by including the encapsulation in the EVPN updates to the neighbor:

```
enable
configure terminal
router bgp 100
address-family l2vpn evpn
neighbor <address> encap {mpls | vxlan}
```



Note

- If VxLan is also configured for the VRF IPv4 address family, EVPN routes imported from VPNv4 table could have both MPLS and VxLAN encapsulation data. EVPN routes imported from VPNv6 table would only have MPLS Encap as VxLAN is not supported for VPNv6 prefixes.
- If EVPN neighbor is configured to only send MPLS encapsulation, then the prefix is advertised to neighbour, only if MPLS encapsulation data with the prefix exists.
- If EVPN neighbor is configured to only send VxLAN encapsulation, then the prefix is advertised to neighbor only if there exists VxLAN encap data with the prefix.
- If EVPN neighbor is configured to only send VxLAN encapsulation, then the prefix is advertised to neighbor only if there exists VxLAN encap data with the prefix.

Howto Configure BGPEVPNoverMultiProtocolLabelSwitching

Configuring BGP over MPLS

Following is sample configuration for configuring BGP over MPLS

```
enable
configure terminal
router bgp 100
template peer-policy policy1
encapsulation mpls ===> template policy config for neighbor encapsulation preference

address-family l2vpn evpn
neighbor 10.1.1.1 encapsulation mpls ===> neighbor encapsulation preference
exit
address-family ipv6 vrf vrf1
advertise l2vpn evpn 2000 ===> import locally sourced or PE-CE vpnv6 routes into
evpn and the limit on the number of prefixes that can be imported (Optional).
network 0::0/0 evpn ===> only import default route into evpn from this vrf ipv6 table
```

Configuring BGP EVPN over MPLS (Inter AS)

Figure 4: BGP over MPLS (Inter AS)



PE1 Configuration

```
vrf def red
  address-family ipv4
    route-target export 1:1 stitching
    route-target import 2:2 stitching

    address-family ipv6
    route-target export 1:1 stitching
    route-target import 2:2 stitching

router bgp 1
  address-family l2vpn evpn
  neighbor 192.0.2.10 activate
  neighbor 192.0.2.10 encap mpls

  address-family ipv4 vrf red
  advertise l2vpn evpn
  address-family ipv6 vrf red
  advertise l2vpn evpn
```

PE2 Configuration

```
router bgp 1
  no bgp default route-target filter
  address-family l2vpn evpn
  neighbor 10.1.1.1 activate
  neighbor 10.1.1.1 encap mpls
  neighbor 10.1.1.1 next-hop-unchanged
  neighbor 203.0.113.10 activate
  neighbor 203.0.113.10 encap mpls
  neighbor 203.0.113.10 next-hop-unchanged
```

PE3 Configuration

```
vrf def red
  address-family ipv4
    route-target export 2:2 stitching
    route-target import 1:1 stitching
  address-family ipv6
    route-target export 2:2 stitching
    route-target import 1:1 stitching

router bgp 1
  address-family l2vpn evpn
  neighbor 192.0.2.10 encap mpls
  address-family ipv4 vrf red
  advertise l2vpn evpn
  address-family ipv6 vrf red
  advertise l2vpn evpn
```

Configuring BGP EVPN over MPLS (InterAS L3VPN)

PE1 Configuration

```
vrf def red
    address-family ipv4
        route-target export 1:1 stitching
        route-target import 2:2 stitching

    address-family ipv6
        route-target export 1:1 stitching
        route-target import 2:2 stitching

router bgp 1
    address-family l2vpn evpn
    neighbor 192.0.2.10 activate
    neighbor 192.0.2.10 encap mpls

    address-family ipv4 vrf red
    advertise l2vpn evpn
    address-family ipv6 vrf red
    advertise l2vpn evpn
```

PE2 Configuration

```
vrf def red
    address-family ipv4
        route-target import 2:2
        route-target import 1:1 stitching
    address-family ipv6
        route-target import 2:2
        route-target import 1:1 stitching

router bgp 1
    address-family l2vpn evpn
    import vpnv4 unicast
    import vpnv6 unicast
    neighbor 10.1.1.1 encap mpls
    neighbor 10.1.1.1 next-hop-unchanged
    address-family vpnv4
    import l2vpn evpn
    neighbor 203.0.113.10 activate
    neighbor 203.0.113.10 next-hop-unchanged

    address-family vpnv6
    import l2vpn evpn
    neighbor 203.0.113.10 activate
    neighbor 203.0.113.10 next-hop-unchanged
```

PE3 Configuration

```
vrf def red
    address-family ipv4
        route-target export 2:2
        route-target import 1:1
    address-family ipv6
        route-target export 2:2
        route-target import 1:1
    router bgp 1
    address-family vpnv4
    neighbor 192.0.2.10 activate
```


Configuration Examples for BGP EVPN over MPLS

Verifying EVPN Neighbor

Use the following command to verify EVPN Neighbour configuration when no encapsulation is configured:

```
PE3# show bgp l2vpn evpn neighbors
For address family: L2VPN E-VPN
  Session: 192.0.2.10
  BGP table version 38, neighbor version 38/0
  Output queue size : 0
  Index 1, Advertise bit 0
  1 update-group member
  Community attribute sent to this neighbor
  Extended-community attribute sent to this neighbor
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
  Prefers VxLAN if VTEP is UP else MPLS
```

Use the following command to verify EVPN Neighbour configuration when MPLS encapsulation is configured:

```
PE3# show bgp l2vpn evpn neighbors 10.1.1.1
For address family: L2VPN E-VPN
  Session: 10.1.1.1
  BGP table version 38, neighbor version 1/38
  Output queue size : 0
  Index 0, Advertise bit 0
  Community attribute sent to this neighbor
  Extended-community attribute sent to this neighbor
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
  Prefers MPLS
```

Additional References for BGP EVPN over MultiProtocol Label Switching

Standards and RFCs

Standard	Title
RFC 7432	BGP MPLS-Based Ethernet VPN

