



EVPN Integrated Routing and Bridging (L2 and L3 Anycast Gateway) and Data Center Interconnect or Border Leaf (Single Homing)

Table 1: Feature History

Feature Name	Release Information	Description
EVPN Integrated Routing and Bridging (L2 and L3 Anycast Gateway) and Data Center Interconnect or Border Leaf (Single Homing)	Cisco IOS XE Bengaluru 17.5.1	This feature allows the devices to forward both layer 2 or bridged and layer 3 or routed traffic providing optimum unicast and multicast forwarding for both intra-subnets and inter-subnets within and across data centers. Data Center Interconnects (DCI) products are targeted at the Edge or Border Leaf (BL) of data center environments, joining data centers to each other in a point-to-point or point-to-multipoint fashion, or at times extending the connectivity to internet gateways or peering points.

Prior to Cisco IOS XE Bengaluru Release 17.5.1, EVPN over MPLS network could only allow the routers to forward layer 2 traffic.

Starting with Cisco IOS XE Bengaluru Release 17.5.1, this feature allows the router in an EVPN over MPLS network to perform both bridging and routing. Integrated Routing and Bridging (IRB) provides the ability to route between a bridge group and a routed interface using a BVI. The BVI is a virtual interface within the router that acts like a normal routed interface. A BVI is associated with a single bridge domain and represents the link between the bridging and the routing domains on the router. An EVPN-based IRB allows the routers to forward both Layer 2 or bridged and Layer 3 or routed traffic providing optimum unicast and multicast forwarding for both intra-subnets and inter-subnets within and across data centers. Data Center Interconnects (DCI) products are targeted at the Edge or Border leaf (BL) of data center environments, joining data centers to each other in a point-to-point or point-to-multipoint fashion, or at times extending the connectivity to internet gateways or peering points.

A bridge domain performs bridging when it forwards traffic to the same subnet. Similarly, a bridge domain interface performs routing when it forwards traffic to a different subnet. The devices in the network forward traffic to each other through the Distributed Anycast Gateways (DAG). The Ethernet VPN over MPLS Integrated IRB Single-Homing (SH) with DAG feature provides support for symmetric IRB model.

In symmetric IRB, both the ingress and egress bridge domain interfaces perform both bridging and routing. A packet first moves through a MAC VRF followed by an IP VRF of the ingress device. It then moves through an IP VRF followed by a MAC VRF on the PE of the egress device. The PEs of ingress and egress devices equally share all the packet processing associated with intersubnet forwarding semantics.

In symmetric IRB, you are required to define only the endpoints on the ingress and egress bridge domain interfaces. Symmetric IRB offers better scalability with the BGP EVPN over MPLS fabric.


- [Distributed Anycast Gateway \(DAG\) with Bridge Domains, on page 2](#)
- [Symmetric IRB with MPLS on Distributed Gateways, on page 3](#)
- [Host MAC-IP Binding on a Single-Homed DAG, on page 4](#)
- [Host MAC-IP Mobility, on page 4](#)
- [ARP and ND Flooding Suppression, on page 5](#)
- [Prerequisites, on page 5](#)
- [Restrictions, on page 6](#)
- [Configuring EVPN Integrated Routing and Bridging \(L2 and L3 Anycast Gateway\) and Data Center Interconnect or Border Leaf \(Single Homing\), on page 6](#)
- [Verification of EVPN Integrated Routing and Bridging \(L2 and L3 Anycast Gateway\) and Data Center Interconnect or Border Leaf \(Single Homing\) Configuration, on page 8](#)
- [EVPN-IRB DHCP v4 and v6 Relay over Segment Routing, on page 14](#)
- [Stitching of Subnet Route from EVPN to L3VPN, on page 21](#)

Distributed Anycast Gateway (DAG) with Bridge Domains

Distributed Anycast Gateway (DAG) is a default gateway addressing mechanism in a BGP EVPN fabric. The feature enables the use of the same gateway IP and MAC address across all the devices in an EVPN over MPLS network with IRB. This ensures that every device functions as the default gateway for the workloads directly connected to it. The feature facilitates flexible workload placement, host mobility, and optimal traffic forwarding across the BGP EVPN fabric.

In the scenario below, the DAGs are directly attached to hosts or network with IP-VRF routing enabled on the IRB (BDI) interfaces on the gateways. To reduce the complexity, only virtual MAC DAGs are supported and the Duplication Address Detection (DAD) for IPv6 on the BDI interfaces on DAG is disabled.

Figure 1: DAG with Bridge Domains




On the DAG, the bridge domain checks if an Address Resolution Protocol (ARP) or Neighbor Discovery Protocol (NDP) packet from a local host is sent to the BDI (Gateway) IP addresses. If the packet is sent to BDI (Gateway) IP addresses, this packet is handled by local BDI and it is not flooded into the bridge domain and sent across the EVPN IRB fabric.

Symmetric IRB with MPLS on Distributed Gateways

Symmetric IRB is a distributed routing model which utilizes direct IP-VRF to IP-VRF connectivity for inter-subnet traffic. To support symmetric IRB, the native IRB needs to be enabled on distributed gateways by creating the BDIs, configuring virtual MAC, IP-VRF, and anycast IP address.

After the native IRB is enabled, BGP allocates the L3 label for the RT-2's and RT-5's per VRF basis and advertises it.


Figure 2: Symmetric IRB with MPLS on Distributed Gateways



Host MAC-IP Binding on a Single-Homed DAG

The Host MAC-IP binding is learned by snooping Address Resolution Protocol (ARP), Neighbor Discovery Protocol, or DHCP packets. After the MAC-IP binding is learned, an age timer (AGE_TIME) is applied to the locally learned binding entry. The binding entry is refreshed whenever the host initiates ARP or ND procedures.

Figure 3: Host MAC-IP Binding on a Single-Homed DAG



Host MAC-IP Mobility

The host MAC-IP mobility helps to handle the following events:

- Host Move Learn from Data Packet and Generic Attribute Registration Protocol (GARP)
- Host Move Detection for Silent Host

Also, the host MAC-IP mobility supports the following scenarios:

- Moving MAC from local to local
- Moving MAC from local to remote
- Moving MAC from remote to local
- Moving IP local to local
- Moving IP from local to remote
- Moving IP from remote to local


ARP and ND Flooding Suppression

The Address Resolution Protocol (ARP) and Neighbor Discovery (ND) Protocol flooding suppression depend on device tracking enabled on the same VLAN or interface. The Switch Integrated Security Feature based (SISF-based) device tracking helps to track the presence, location, and movement of end nodes in the network. The device receives the SISF snoops traffic that extracts the device identity (MAC and IP address) and stores it in a binding table. SISF-based device tracking supports both IPv4 and IPv6.

When you enable IPv4 or IPv6 flooding suppression, it helps to minimize the flooding of a broadcast or multicast packet over the EVPN IRB fabric and to remote CEs such as host and router. The multicast and broadcast suppression capabilities help to preserve bandwidth in wireless networks.

This helps to suppress the broadcast (ARP) or link-local multicast (NDP) messages circulating in the layer 2 domain and the packets are relayed after converting their L2 addresses to unicast.

Figure 4: ARP and ND Flooding Suppression



Prerequisites

- Host MAC-IP learning
- Symmetric IRB for IP-VRF to IP-VRF inter-subnet traffic over MPLS
- DAG with bridge-domain
- Host MAC-IP mobility
- ARP/ND flooding suppression

Restrictions

- DHCP Snooping and EVPN IRB snooping cannot be enabled on the same bridge domain.
- Stitching on a collapsed Border Leaf (BL) or Spine Leaf is *not* supported.
- EVPN IRB feature and Hot Standby Router Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP) feature are *not* supported simultaneously on a router.
- EVPN Layer 2 is *not* supported on Media Access Control Security (MACsec).
- Asymmetric IRB and centralized IRB are *not* supported.
- Only Virtual MAC is supported for DAG.
- EVPN IRB interface cannot be a part of global Virtual Routing and Forwarding (VRF).
- Virtual Private LAN Service (VPLS) stitching is *not* supported.
- EVPN IRB is supported *only* on BDI interfaces.
- Remote MAC learned via EVPN-BGP cannot be controlled by MACsec/MAC limit features.
- Static MAC should *not* be configured as remote MAC address.
- As MAC scale is limited to 16K MACs at system level, remote MAC of 8000 scale is recommended.
- Remote MAC is learned at the rate of 400 PPS. Beyond this scale, you might encounter stale MACs or MAC stuck scenarios.
- A maximum number of 950 EVI is supported.
- EVPN IRB and MACsec features are *not* supported together.

Configuring EVPN Integrated Routing and Bridging (L2 and L3 Anycast Gateway) and Data Center Interconnect or Border Leaf (Single Homing)

This section shows the configuration of EVPN IRB (L2 and L3 Anycast Gateway) and Data Center Interconnect (DCI) or Border Leaf (BL) Single Homing.

Configure IP VRF

To configure IP Virtual Routing and Forwarding (VRF):

```
vrf definition evpn_vrf
rd 10.0.0.1:1
!
address-family ipv4
 route-target export 1000:1
 route-target import 1000:1
 route-target export 1000:1 stitching
```

```

    route-target import 1000:1 stitching
  exit-address-family
!
address-family ipv6
  route-target export 1000:1
  route-target import 1000:1
  route-target export 1000:1 stitching
  route-target import 1000:1 stitching
exit-address-family

```

Configure Layer 2 Virtual Private Network EVPN

To configure layer 2 virtual private network EVPN:

```

l2vpn evpn
  replication-type ingress
  router-id Loopback0
l2vpn evpn instance 10 vlan-based
l2vpn evpn instance 11 vlan-based
l2vpn evpn instance 12 vlan-based

```

Configure Bridge Domain

To configure bridge domain:

```

bridge-domain 10
  member GigabitEthernet0/3/7 service-instance 10
  member evpn-instance 10
bridge-domain 11
  member GigabitEthernet0/3/7 service-instance 11
  member evpn-instance 11
bridge-domain 12
  member GigabitEthernet0/3/7 service-instance 12
  member evpn-instance 12

```

Configure Bridge Domain IRB Interface

To configure bridge domain IRB interface:

```

interface BDI10
  mac-address 0011.1111.1111
  vrf forwarding evpn_vrf
  ip address 191.168.1.1 255.255.255.0
  ipv6 address 1968:1::1/64
end

interface BDI11
  mac-address 0011.1111.1112
  vrf forwarding evpn_vrf
  ip address 191.168.2.1 255.255.255.0
  ipv6 address 1969:2::1/64
end

```

Configure BGP IRB

To configure BGP IRB:

```

router bgp 1000
  bgp router-id interface Loopback0

```

```

bgp log-neighbor-changes
neighbor 2.2.2.1 remote-as 1000
neighbor 2.2.2.1 ha-mode sso
neighbor 2.2.2.1 update-source Loopback0
!
address-family ipv4
  bgp additional-paths install
  network 10.0.0.1 mask 255.255.255.255
  neighbor 2.2.2.1 activate
  neighbor 2.2.2.1 send-community both
  neighbor 2.2.2.1 send-label
exit-address-family
!
address-family ipv6
  bgp additional-paths install
  no bgp recursion host
  neighbor 2.2.2.1 activate
  neighbor 2.2.2.1 send-community both
  neighbor 2.2.2.1 send-label
exit-address-family
!
address-family l2vpn evpn
  neighbor 2.2.2.1 activate
  neighbor 2.2.2.1 send-community both
exit-address-family
!
address-family ipv4 vrf evpn_vrf
  redistribute connected
exit-address-family
!
address-family ipv6 vrf evpn_vrf
  redistribute connected
exit-address-family

```

Verification of EVPN Integrated Routing and Bridging (L2 and L3 Anycast Gateway) and Data Center Interconnect or Border Leaf (Single Homing) Configuration

This section shows the verification examples of EVPN IRB (L2 and L3 Anycast Gateway) and Data Center Interconnect or Border Leaf (Single Homing) configuration.

Verify Device Tracking Database

Use **show device-tracking database** command to verify device tracking database:

```

Router#show device-tracking database
Binding Table has 16 entries, 7 dynamic
Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Address Resolution Protocol,
DH4 - IPv4 DHCP, DH6 - IPv6 DHCP, PKT - Other Packet, API - API created
Preflevel flags (prlvl):
0001:MAC and LLA match      0002:Orig trunk      0004:Orig access
0008:Orig trusted trunk    0010:Orig trusted access  0020:DHCP assigned
0040:Cga authenticated     0080:Cert authenticated  0100:Statically assigned

      Network Layer Address      Link Layer Address Interface      bd  prlvl  age
      state      Time left
ARP 191.168.3.2                0000.0000.0003      Gi0/3/7      12    0005  174s

```

```

    REACHABLE 125 s
L   191.168.3.1                0011.1111.1113    BD12      12    0100  303mn
    REACHABLE
ARP 191.168.2.2                0000.0000.0002    Gi0/3/7   11    0005   3mn
    REACHABLE 117 s
L   191.168.2.1                0011.1111.1112    BD11      11    0100  303mn
    REACHABLE
ARP 191.168.1.2                0000.0000.0001    Gi0/3/7   10    0005  174s
    REACHABLE 131 s
L   191.168.1.1                0011.1111.1111    BD10      10    0100  303mn
    REACHABLE
L   FE80::211:11FF:FE11:1113    0011.1111.1113    BD12      12    0100  303mn
    REACHABLE
L   FE80::211:11FF:FE11:1112    0011.1111.1112    BD11      11    0100  303mn
    REACHABLE
L   FE80::211:11FF:FE11:1111    0011.1111.1111    BD10      10    0100  303mn
    REACHABLE
ND FE80::200:FF:FE00:3          0000.0000.0003    Gi0/3/7   12    0005  164s
    REACHABLE 146 s
ND FE80::200:FF:FE00:1          0000.0000.0001    Gi0/3/7   10    0005  164s
    REACHABLE 142 s
ND 1970:3::2                    0000.0000.0003    Gi0/3/7   12    0005  174s
    REACHABLE 126 s
L   1970:3::1                    0011.1111.1113    BD12      12    0100  303mn
    REACHABLE
L   1969:2::1                    0011.1111.1112    BD11      11    0100  303mn
    REACHABLE
ND 1968:1::2                    0000.0000.0001    Gi0/3/7   10    0005  174s
    REACHABLE 130 s
L   1968:1::1                    0011.1111.1111    BD10      10    0100  303mn
    REACHABLE

```

Verify L2VPN EVPN Summary

Use **show l2vpn evpn summary** command to verify L2VPN EVPN IRB summary:

```

Router# show l2vpn evpn summary
L2VPN EVPN
  EVPN Instances (excluding point-to-point): 3
    VLAN Aware: 0
    VLAN Based: 3
    VLAN Bundle: 0
  Bridge Domains: 3
  BGP: ASN 1000, address-family l2vpn evpn configured
  Router ID: 10.0.0.1
  Label Allocation Mode: Per-BD
  Global Replication Type: Ingress
  ARP/ND Flooding Suppression: Enabled
  MAC Duplication: seconds 180 limit 5
  MAC Addresses: 3
    Local: 3
    Remote: 0
  Duplicate: 0
  IP Duplication: seconds 180 limit 5
  IP Addresses: 7
    Local: 7
    Remote: 0
    Duplicate: 0
  Advertise Default Gateway: No
  Default Gateway Addresses: 0
    Local: 0
    Remote: 0
  Global IP Local Learn: Enabled

```

```

IP local learning limits
  IPv4: 4 addresses per-MAC
  IPv6: 12 addresses per-MAC
IP local learning timers
  Down:      10 minutes
  Poll:      1 minutes
  Reachable: 5 minutes
  Stale:     30 minutes

```

Verify L2VPN EVPN EVI

Use the **show l2vpn evpn evi detail** command to verify L2VPN EVPN EVI:

```

Router#show l2vpn evpn evi 10 detail
EVPN instance:      10 (VLAN Based)
RD:                 10.0.0.1:10 (auto)
Import-RTs:         1000:10
Export-RTs:         1000:10
Per-EVI Label:      none
State:              Established
Replication Type:   Ingress (global)
Encapsulation:      mpls
IP Local Learn:     Enabled (global)
Adv. Def. Gateway:  Disabled (global)
Bridge Domain:      10
  Ethernet-Tag:      0
  BUM Label:         16
  Per-BD Label:      17
  BDI Label:         none
  State:             Established
  Flood Suppress:    Attached
  Access If:         BDI10
  VRF:               evpn_vrf
  IPv4 IRB:          Enabled
  IPv6 IRB:          Enabled
Pseudoports:
  GigabitEthernet0/3/7 service instance 10
    Routes: 1 MAC, 3 MAC/IP

```

Verify Platform Software Infrastructure Punt Statistics

Use the **show platform software infrastructure punt statistics** command to verify the platform software infrastructure punt statistics:

```

Router#show platform soft infrastructure punt statistics
UEA Punt Statistics

```

Global drops : 0

Queue Name	Rx count	Drop count
SW FORWARDING Q	403038	0
ROUTING PROTOCOL Q	159	0
ICMP Q	0	0
HOST Q	400	0
ACL LOGGING Q	0	0
STP Q	0	0
L2 PROTOCOL Q	0	0
MCAST CONTROL Q	0	0
BROADCAST Q	0	0
REP Q	0	0

BGP LDP Q	0	0
CONTROL Q	0	0
IP MPLS TTL Q	0	0
DEFAULT MCAST Q	0	0
MCAST ROUTE DATA Q	0	0
MCAST MISMATCH Q	0	0
RPF FAIL Q	0	0
ROUTING THROTTLE Q	0	0
MCAST Q	0	0
MPLS OAM Q	0	0
IP MPLS MTU Q	0	0
PTP Q	0	0
LINUX ND Q	0	0
KEEPALIVE Q	9256	0
ESMC Q	0	0
FPGA BFD Q	0	0
FPGA CCM Q	0	0
FPGA CFE Q	0	0
L2PT DUP Q	0	0
TDM CTRL Q	0	0
ICMP UNREACHABLE Q	402918	0
SSFP Q	0	0
MIRROT Q	0	0

Verify Platform Software Infrastructure Inject

Use the **show platform software infrastructure inject** command to verify the platform software infrastructure inject:

```
Router#show platform software infrastructure inject
Statistics for L3 injected packets:
27819 total inject pak, 24 failed
0 sent, 532 prerouted
24 non-CEF capable, 334 non-unicast
4965 IP, 20885 IPv6
0 MPLS, 0 Non-IP Tunnel
0 UDLR tunnel, 0 P2MP replicated mcast
0 Non-IP Fastswitched over Tunnel, 1945 legacy pak path
0 Other packet
0 IP fragmented
25318 normal, 0 nexthop
532 adjacency, 0 feature
0 undefined
0 pak find no adj, 0 no adj-id
4991 sb alloc, 25850 sb local
0 p2mcast failed count 0 p2mcast enqueue fail
0 unicast dhc
0 mobile ip
105 IPv6 NA
101 IPv6 NS
0 Transport failed cases
0 Grow packet buffer
2 Cant-l3-inject-pkts
per feature packet inject statistics
0 Feature multicast
0 Feature Edge Switching Service
0 Feature Session Border Controller
0 Feature interrupt level
0 Feature use outbound interface
0 Feature interrupt level with OCE
0 Feature ICMPv6 error message
0 Feature Session Border Controller media packet injection
0 Feature Tunnel Ethernet over GRE
```

```

0 Feature Secure Socket Layer Virtual Private Network
0 Feature EPC Wireshark injecting packets
0 Feature multicast overlay replication
Statistics for L2 injected packets:
356 total L2 inject pak, 0 failed
3 total BD inject pak, 0 failed
353 total EFP inject pak, 0 failed
0 total VLAN inject pak, 0 failed

```

Verify BGP L2VPN EVPN Detail

Use the **show bgp l2vpn evpn detail** to verify BGP L2VPN EVPN detail:

```
Router#show bgp l2vpn evpn detail
```

```

Route Distinguisher: 10.0.0.1:10
BGP routing table entry for [2][10.0.0.1:10][0][48][000000000001][0][*]/20, version 1027
  Paths: (1 available, best #1, table evi_10)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.0.0.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 17
      Extended Community: RT:1000:10
      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 22:39:43 IST
BGP routing table entry for [2][10.0.0.1:10][0][48][000000000001][32][191.168.1.2]/24,
version 1021
  Paths: (1 available, best #1, table evi_10)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.0.0.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 17, Label2 22
      Extended Community: RT:1000:1 RT:1000:10
      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 22:39:42 IST
BGP routing table entry for [2][10.0.0.1:10][0][48][000000000001][128][1968:1::2]/36, version
1022
  Paths: (1 available, best #1, table evi_10)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.0.0.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 17, Label2 24
      Extended Community: RT:1000:1 RT:1000:10
      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 22:39:42 IST
BGP routing table entry for
[2][10.0.0.1:10][0][48][000000000001][128][FE80::200:FF:FE00:1]/36, version 1023
  Paths: (1 available, best #1, table evi_10)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.0.0.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 17, Label2 24
      Extended Community: RT:1000:1 RT:1000:10
      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 22:39:42 IST

```

```

Route Distinguisher: 10.0.0.1:11
BGP routing table entry for [2][10.0.0.1:11][0][48][000000000002][0][*]/20, version 1012
  Paths: (1 available, best #1, table evi_11)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.0.0.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 19
      Extended Community: RT:1000:11
      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 22:39:13 IST
BGP routing table entry for [2][10.0.0.1:11][0][48][000000000002][32][191.168.2.2]/24,
version 1011
  Paths: (1 available, best #1, table evi_11)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.0.0.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 19, Label2 22
      Extended Community: RT:1000:1 RT:1000:11
      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 22:39:12 IST
Route Distinguisher: 10.0.0.1:12
BGP routing table entry for [2][10.0.0.1:12][0][48][000000000003][0][*]/20, version 1028
  Paths: (1 available, best #1, table evi_12)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.0.0.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 21
      Extended Community: RT:1000:12
      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 22:39:43 IST
BGP routing table entry for [2][10.0.0.1:12][0][48][000000000003][32][191.168.3.2]/24,
version 1024
  Paths: (1 available, best #1, table evi_12)
  Not advertised to any peer
Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.0.0.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 21, Label2 25
      Extended Community: RT:1000:2 RT:1000:12
      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 22:39:42 IST
BGP routing table entry for
[2][10.0.0.1:12][0][48][000000000003][128][FE80::200:FF:FE00:3]/36, version 1026
  Paths: (1 available, best #1, table evi_12)
  Not advertised to any peer
Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.0.0.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 21, Label2 25
      Extended Community: RT:1000:2 RT:1000:12
      rx pathid: 0, tx pathid: 0x0
      Updated on Dec 16 2020 22:39:42 IST
Route Distinguisher: 10.0.0.1:10
BGP routing table entry for [3][10.0.0.1:10][0][32][10.0.0.1]/17, version 9
  Paths: (1 available, best #1, table evi_10)
  Not advertised to any peer

```

```

Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
  Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
  Extended Community: RT:1000:10
  PMSI Attribute: Flags:0x0, Tunnel type:IR, length 4, label:16 tunnel identifier: 0000
0000
  rx pathid: 0, tx pathid: 0x0
  Updated on Dec 16 2020 17:33:58 IST

```

EVPN-IRB DHCP v4 and v6 Relay over Segment Routing

Table 2: Feature History

Feature Name	Release Information	Description
EVPN-IRB DHCP v4 and v6 Relay over Segment Routing	Cisco IOS XE Bengaluru 17.6.1	This feature introduces a specialised implementation of DHCP packets to support DHCPv4 and DHCPv6 in an EVPN Fabric with Distributed Anycast Gateways (DAGs) on the same Virtual Routing and Forwarding (VRF). It also avoids DHCP discovery packet floods across the fabric. The flooding suppression feature is also enhanced to intercept multicast or broadcast DHCP packets when DHCP relay is configured on the DAG to perform the required action and localize the scope of the service. This feature is not supported with Cisco ASR RSP3 module. It is only supported with Cisco ASR RSP2 module.

Prior to Cisco IOS XE Bengaluru Release 17.5.1, DHCP relay agent was not supported on EVPN fabric solution. EVPN IRB extends a link (layer 2 segment) across a routed backbone (fabric). As a result, multicast or broadcast packets reach all layer 2 segments that introduces problems like black holing etc and security breaches. In large EVPN deployments, the amount of multicast or broadcast traffic in the fabric can be overwhelming and can also overload DHCP server. In addition to scalability and performance impact, there is also a functional problem with DHCPv6 relay agents, which can lead to traffic drop and host reachability problems.

Starting with Cisco IOS XE Bengaluru Release 17.6.1, specialised handling of DHCP packets is implemented to support DHCPv4 and DHCPv6 in an EVPN Fabric with Distributed Anycast Gateways (DAGs) on the same Virtual Routing and Forwarding (VRF). It also avoids DHCP discovery packet floods across the fabric.

If the packets are flooded in the stretched layer 2, all the DHCP relay enabled DAGs relay the packet, which generates unnecessary workload on the DHCP servers. For certain specific scenarios, this also installs route on the wrong DAGs and causes outage.

To avoid the above-mentioned situation, First-Hop device handles the DHCP services and DHCP relay is enabled on IPv4 and IPv6 traffic for the same VRF.


How DHCP Relay Agent Works

DHCP relay agent is used to forward request and replies between DHCP server and client when the server is present on different link. The role of DHCP relay is to intercept the DHCP requests (broadcast or multicast) on a layer 2 network, convert it to layer 3 unicast traffic, and send to DHCP server. DHCP relay can be on any device in the network that has a layer 3 connectivity to the DHCP server. It is usually configured on the default gateway that faces the DHCP client. In an EVPN setup, the DHCP relay is configured on each DAG.

When the router receives multicast or broadcast DHCP packets from DHCP clients and if DHCP relay is configured on that interface, DHCP relay agent receives a copy of the DHCP packet and sends it to the DHCP server. There can be more than one relay or server on the same vlan for redundancy.

This feature ensures the DHCP packets only get flooded on local ports and not on the fabric ports. This stops other DHCP relay agents in the fabric to receive the DHCP packet. The following network topology is an example to show how DHCP relay agent works.

Figure 5: DHCP Relay Agent Network Topology



Enhancement of Flooding Suppression Feature

Prior to Cisco IOS XE Bengaluru Release 17.6.1, the flooding suppression feature configured on DAGs in an IRB EVPN setup could suppress ARP and NP packets and minimize flooding of these packets into the EVPN fabric. Starting with Cisco IOS XE Bengaluru Release 17.6.1, this feature is enhanced to intercept multicast or broadcast DHCP packets when DHCP relay is configured on the DAG to perform the required action and localize the scope of the service.

Both DHCPv4 and DHCPv6 packets are suppressed across the fabric. All multicast or broadcast DHCP packets are forwarded unchanged in the local domain.

Restrictions

- This feature is *only* supported on DAG deployments.
- This feature is *only* supported on Provider Edge (PE) routers.
- This feature cannot co-exist with DHCP snooping on the PE router.
- BD scale of 4095 is supported for EVPN IRB. However, only 100 of those BDs can be configured to use DHCP relay over EVPN IRB.
- When BDI is disabled or shut down, the DHCP packets remain suppressed. Thus, there is no flooding in the layer 2 domain.

Configure DHCP Relay Flooding Suppression

To configure DHCP relay flooding suppression:

```
PE1(config)#service internal
PE1(config)#l2vpn evpn
PE1(config-evpn)#flooding-suppression
address-resolution Suppress flooding of Address Resolution and Neighbor
Discovery Protocol packets
PE1(config-evpn)#flooding-suppression dhcp-relay
disable Disable flooding suppression
PE1(config-evpn)#flooding-suppression dhcp-relay disable
```



Note

- DHCP relay suppress flooding of DHCP packets to the fabric when DHCP relay is enabled locally.
- YANG model is not supported.
- The DHCP relay flooding suppression feature is enabled by default.

Configure DHCP Relay on DAGs

Global Configuration:

```
ip dhcp relay information option vpn
```

DHCP Server in the Same Tenant IP-VRF:



Note

Same tenant IP-VRFs are the VRFs that are part of the same leaf node.

```
!! Unique Loopback IPv4/IPv6 in the tenant IP-VRF
interface Loopback192
description DHCP-Relay Source Interface
vrf forwarding red
ip address 10.10.10.4 255.255.255.255
ipv6 address 2001:10::4/128
end

interface BDI12
```

```

mac-address 0012.0012.0012
vrf forwarding red
ip dhcp relay source-interface Loopback192
ip address 192.168.12.254 255.255.255.0
ip helper-address 10.10.10.10
ipv6 address 2001:12::254/64
ipv6 dhcp relay destination 2001:10::10
ipv6 dhcp relay source-interface Loopback192
end!

```

DHCP Server in Different Tenant IP-VRF:



Note Different tenant IP-VRFs are not part of the same leaf node and they need the IRB functionality.

```

!! Unique Loopback IPv4/IPv6 in the tenant IP-VRF
interface Loopback192
description DHCP-Relay Source Interface
vrf forwarding green
ip address 10.10.10.4 255.255.255.255
ipv6 address 2001:10::4/128
end
interface BDI12
mac-address 0012.0012.0012
vrf forwarding red
ip dhcp relay source-interface Loopback192
ip address 192.168.12.254 255.255.255.0
ip helper-address vrf green 10.10.10.10
ipv6 address 2001:12::254/64
ipv6 dhcp relay vrf green destination 2001:10::10
ipv6 dhcp relay source-interface Loopback192
end!

```

DHCP Server in Global VRF:

```

!! No need for Unique Loopback IPv4/IPv6
interface BDI12
mac-address 0012.0012.0012
vrf forwarding red
ip address 192.168.12.254 255.255.255.0
ip helper-address global 10.10.10.10
ipv6 address 2001:12::254/64
ipv6 dhcp relay global destination 2001:10::10
end!

```

Verification of DHCPv4 and DHCPv6 Relay Configuration

Use the **show device-tracking policies** command to verify all the SISF feature policies attached to bridge domain.

```

PE3#show device-tracking policies
Target      Type Policy      Feature      Target range
bd 11       bd  evpn-no-device-track Device-tracking bd all
bd 11       bd  evpn-flood-suppress  Flooding Suppress bd all

```

Use the **show flooding-suppression policy** command to verify the settings of a policy.

```

PE3#show flooding-suppression policy
PE3#show flooding-suppression policy evpn-flood-suppress

```

```

Flooding suppress policy evpn-flood-suppress configuration:
  Suppressing  NDP
  Suppressing  DHCPv6
  Suppressing  ARP
  Suppressing  DHCPv4
mode:No-Proxy multicast resolution requests
Policy evpn-flood-suppress is applied on the following targets:
Target          Type  Policy          Feature          Target range
bd 11           BD   evpn-flood-suppress  Flooding Suppr  bd all
bd 12           BD   evpn-flood-suppress  Flooding Suppr  bd all

```

Use the **show device-tracking counters** *bridge-domain-id* to verify the counters.

```

PE-1#show device-tracking counters bd 11
Received messages on bd 11 :
Protocol          Protocol message
NDP               RS[4] RA[4] NS[1777] NA[2685]
DHCPv6
ARP               REQ[12] REP[1012]
DHCPv4
ACD&DAD          --[8]
:
:
Limited Broadcast to Local message on bd 11 :
Type              Protocol message
NDP
DHCPv6            SOL[1] REQ[1] REB[1]
ARP
DHCPv4            DIS[1] REQ4[1]

```

Use the **show l2fib output-list** and **show l2fib bridge-domain** *bd -id* **table unicast** commands to verify the information about Layer 2 Forwarding Information Base (L2FIB).

```

PE3#show l2fib output-list
ID BD Port Flags
-----
1035 11 1 local port list
1036 12 0 local port list
5120 1 0 flood list
5130 11 4 flood list
5131 12 1 flood list
PE3#show l2fib output-list 1035
ID : 1035
Bridge Domain : 11
Reference Count : 3
Flags : local port list
Port Count : 1
Port(s) : BD_PORT Gi0/3/2:11
PE3#show l2fib bridge-domain 11 table unicast
MAC Address Adjacency
-----
7069.5a39.ef8a BD_PORT Gi0/3/2:11
7069.5a39.ef94 MPLS_UC PL:1(1) T:MPLS_UC [MAC]17@99.99.99.1
7486.0bc4.d4d4 BD_PORT Gi0/3/2:11
ffff.ffff.ffff Olist: 1035, Ports: 1

```

Use the **show platform software l2fib f0 bdbdomain idunicast all** command to verify the global bridge domain table for MAC and layer 2 multicast.

```

PE-1#show platform software l2fib f0 bd 10 unicast all
MAC          BD          Nhop type          Nhop Idx          Flags
-----

```

```
ffff.ffff.ffffd      10      olist      1034
static
```

Use the **show platform software l2fib f0 mlist index***Nhop Idx* command to verify the output list of the global bridge domain table for MAC and layer 2 multicast.

```
PE-1#show plat soft l2fib f0 mlist index 1034
L2FIB Mlist entries
```

```
Type Index AOM ID CPP Info
```

```
efp 408011 aom id: 149, CPP info: 0x1808bcc (created)
```

Use the **show platform software dpidb ethernet efp interface***interface-name* command to verify <need information>

```
PE-1#sh platform software dpidb ethernet efp interface gigabitEthernet 0/0/0
DPIDB for interface GigabitEthernet0/0/0 (Ethernet Flow Points)
EFP ID: 10, dpidb index: 0x408011
```

Use the **show platform software dpidb ethernet efp interface show platform hardware pp active bridge-domain id***domain id* command to verify the view the Ethernet Flow Point (EFP) information.

```
PE-1#show platform hard pp ac bridge-domain id 10
Bridge Domain Details
:
Nile Vlan Compression Table
```

type	bridgeD_index	fid_index	floodtype	Entry MET
lp_access	NA	24877	3	184341

Use the **show platform software infrastructure inject** command to verify the platform software infrastructure inject:

```
PE-1#show platform software infrastructure inject
:
Statistics for L2 injected packets:
1 total L2 inject pak, 0 failed
0 total BD inject pak, 0 failed
1 total BD-local inject pak, 0 failed
0 total EFP inject pak, 0 failed
0 total VLAN inject pak, 0 failed
```


Stitching of Subnet Route from EVPN to L3VPN

Table 3: Feature History

Feature Name	Release	Description
Stitching of Subnet Route from EVPN to L3VPN	Cisco IOS XE Bengaluru 17.6.1	This feature introduces the collapsed spine and border leaf node in the network topology of single homing DAGs with symmetric IRB, inter-subnet layer 3 traffic within fabric and inter-subnet layer 3 stitching through layer 3 border gateway. The hosts participating in fabric IRB are directly attached with the collapsed spine and border leaf node. This feature is not supported with Cisco ASR RSP3 module. It is only supported with Cisco ASR RSP2 module.


Prior to Cisco IOS XE Bengaluru Release 17.6.1, the hosts participating in fabric IRB were not attached directly with the spine and border leaf nodes. As a result, the hosts had to pass through the spine and border leaf nodes to reach the network. The following figure shows the network topology that illustrates single homing DAGs with symmetric IRB, inter-subnet layer 3 traffic within fabric and inter-subnet layer 3 stitching through layer 3 border gateway.

Figure 6: Single Homing DAGs with Symmetric IRB Network Topology



Starting with Cisco IOS XE Bengaluru Release 17.6.1, the hosts participating in fabric IRB are directly attached with the collapsed spine and border leaf node. The following figure shows the network topology that illustrates single homing DAGs with symmetric IRB, inter-subnet layer 3 traffic within fabric and inter-subnet layer 3 stitching through spine border gateway.

Figure 7: Collapsed Spine or Leaf or RR on a Single Node Network Topology



Leaf Node: The Provider Edge (PE) node that provides EVPN layer 2 or layer 3 services for Customer Edge (CE) devices (wireless nodes). DAG functionality is enabled for layer 3 services. Each PE may host a management VLAN for all the devices that are co-located with the PE. This VLAN is reachable via the management DAG from aggregation or core layers. This DAG is part of a management VRF. The in-band management interface of PE (statically configured loopback interface) is also a part of the same VRF. DHCPv4 and DHCPv6 relay functionality with centralized DHCP server are located beyond the aggregation domain. DHCP server is a part of the same VRF with the corresponding IRB interfaces.

Spine Node: Performs SR and MPLS forwarding among leaf nodes within a single CRAN hub (RR in EVPN address family).

Border Leaf Node: Forwards traffic between CRAN Hub and the aggregation layer. This node stitches EVPN Layer 3 routes and L3VPN/6VPE BGP routes.

Restrictions

- DHCP Snooping and EVPN IRB snooping cannot be enabled on the same bridge domain.
- Stitching on a collapsed Border Leaf (BL) or Spine or Leaf node is *not* supported.
- EVPN IRB feature and Hot Standby Router Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP) feature are *not* supported simultaneously on a router.
- EVPN Layer 2 is *not* supported on Media Access Control Security (MACsec).
- Only Virtual MAC is supported for DAG.
- Asymmetric IRB and centralized IRB are *not* supported.
- EVPN IRB interface cannot be a part of global Virtual Routing and Forwarding (VRF).
- Virtual Private LAN Service (VPLS) stitching is *not* supported.
- EVPN IRB is supported *only* on BDI interfaces.
- Remote MAC learned via EVPN-BGP cannot be controlled by MACsec/MAC limit features.
- Static MAC should not be configured as remote MAC address.
- As MAC scale is limited to 16K MACs at system level, remote MAC of 8000 scale is recommended.
- Remote MAC is learned at the rate of 400 PPS. Beyond this scale, you might encounter stale MACs or MAC stuck scenarios.
- A maximum number of 950 EVI is supported.
- EVPN IRB and MACsec features are *not* supported together.
- Only 100 BDs can be configured to use DHCP relay over EVPN IRB.

Scale

The following table shows the scale for Cisco ASR RSP2 module.

Table 4: Scale for Cisco ASR RSP2 Module

EVI	BD	MAC per BD	MACs Per System (Local and Remote)	Scale (Default Template)
950	4095 Note BD scale of 4095 is supported for EVPN IRB. However, only 100 of those BDs can be configured to use DHCP relay over EVPN IRB.	16K	16K (A maximum of <= 8K remote MAC is recommended) MACs Per system (local and remote)	IPV4 ROUTES = 20000 IPV6 ROUTES = 4000 Maximum VRF = 128 Maximum BDI = 4095 255 VMAC

Configure BGP L3VPN Subnet Advertisement (Stitching) for Collapsed Spine or Leaf or RR on a Single Node

To configure BGP L3VPN subnet advertisement (stitching) for collapsed spine or leaf or RR on a single node:

```
vrf definition evpn_vrf
rd 10.0.0.1:1
!
address-family ipv4
route-target export 1000:1
route-target import 1000:1
route-target export 1000:1 stitching
route-target import 1000:1 stitching
exit-address-family
!
address-family ipv6
route-target export 1000:1
route-target import 1000:1
route-target export 1000:1 stitching
route-target import 1000:1 stitching
exit-address-family
l2vpn evpn
replication-type ingress
router-id Loopback0
l2vpn evpn instance 10 vlan-based
l2vpn evpn instance 11 vlan-based
l2vpn evpn instance 12 vlan-based
```

Verification of BGP L3VPN Subnet Advertisement (Stitching) for Collapsed Spine or Leaf or RR on a Single Node Configuration

Use the following commands to verify BGP L3VPN subnet advertisement (stitching) for collapsed spine or leaf or RR on a single node configuration.

Use the **show device-tracking database** command to verify device tracking database:

```
PE1#show device-tracking database
Binding Table has 16 entries, 7 dynamic
```

```
Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Address Resolution Protocol,
DH4 - IPv4 DHCP, DH6 - IPv6
DHCP, PKT - Other Packet, API - API created
Preflevel flags (prlvl):
0001:MAC and LLA match 0002:Orig trunk 0004:Orig access
0008:Orig trusted trunk 0010:Orig trusted access 0020:DHCP assigned
0040:Cga authenticated 0080:Cert authenticated 0100:Statically assigned
Network Layer Address Link Layer Address Interface bd prlvl age state Time left
ARP 191.168.3.2 0000.0000.0003 Gi0/3/7 12 0005 174s REACHABLE 125 s
L 191.168.3.1 0011.1111.1113 BD12 12 0100 303mn REACHABLE
ARP 191.168.2.2 0000.0000.0002 Gi0/3/7 11 0005 3mn REACHABLE 117 s
L 191.168.2.1 0011.1111.1112 BD11 11 0100 303mn REACHABLE
ARP 191.168.1.2 0000.0000.0001 Gi0/3/7 10 0005 174s REACHABLE 131 s
L 191.168.1.1 0011.1111.1111 BD10 10 0100 303mn REACHABLE
L FE80::211:11FF:FE11:1113 0011.1111.1113 BD12 12 0100 303mn REACHABLE
L FE80::211:11FF:FE11:1112 0011.1111.1112 BD11 11 0100 303mn REACHABLE
L FE80::211:11FF:FE11:1111 0011.1111.1111 BD10 10 0100 303mn REACHABLE
ND FE80::200:FF:FE00:3 0000.0000.0003 Gi0/3/7 12 0005 164s REACHABLE 146 s
ND FE80::200:FF:FE00:1 0000.0000.0001 Gi0/3/7 10 0005 164s REACHABLE 142 s
ND 1970:3::2 0000.0000.0003 Gi0/3/7 12 0005 174s REACHABLE 126 s
L 1970:3::1 0011.1111.1113 BD12 12 0100 303mn REACHABLE
L 1969:2::1 0011.1111.1112 BD11 11 0100 303mn REACHABLE
ND 1968:1::2 0000.0000.0001 Gi0/3/7 10 0005 174s REACHABLE 130 s
L 1968:1::1 0011.1111.1111 BD10 10 0100 303mn REACHABLE
```

Use the **show l2vpn evpn summary** command to verify L2VPN EVPN IRB summary:

```
PE1# show l2vpn evpn summary
L2VPN EVPN
EVPN Instances (excluding point-to-point): 3
VLAN Aware: 0
VLAN Based: 3
VLAN Bundle: 0
Bridge Domains: 3
BGP: ASN 1000, address-family l2vpn evpn configured
Router ID: 10.0.0.1
Label Allocation Mode: Per-BD
Global Replication Type: Ingress
ARP/ND Flooding Suppression: Enabled
MAC Duplication: seconds 180 limit 5
MAC Addresses: 3
Local: 3
Remote: 0
Duplicate: 0
IP Duplication: seconds 180 limit 5
IP Addresses: 7
Local: 7
Remote: 0
Duplicate: 0
Advertise Default Gateway: No
Default Gateway Addresses: 0
Local: 0
Remote: 0
Global IP Local Learn: Enabled
IP local learning limits
IPv4: 4 addresses per-MAC
IPv6: 12 addresses per-MAC
IP local learning timers
Down: 10 minutes
Poll: 1 minutes
Reachable: 5 minutes
Stale: 30 minutes
```

Use the **show l2vpn evpn evi detail** command to verify L2VPN EVPN EVI:

```

PE1#show l2vpn evpn evi 10 detail
EVPN instance: 10 (VLAN Based)
RD: 10.0.0.1:10 (auto)
Import-RTs: 1000:10
Export-RTs: 1000:10
Per-EVI Label: none
State: Established
Replication Type: Ingress (global)
Encapsulation: mpls
IP Local Learn: Enabled (global)
Adv. Def. Gateway: Disabled (global)
Bridge Domain: 10
Ethernet-Tag: 0
BUM Label: 16
Per-BD Label: 17
BDI Label: none
State: Established
Flood Suppress: Attached
Access If: BDI10
VRF: evpn_vrf
IPv4 IRB: Enabled
IPv6 IRB: Enabled
Pseudoports:
GigabitEthernet0/3/7 service instance 10
Routes: 1 MAC, 3 MAC/IP

```

Use the **show platform software infrastructure punt statistics** command to verify the platform software infrastructure punt statistics:

```

PE1#show platform software infrastructure punt statistics

```

```

UEA Punt Statistics

```

```

Global drops : 0

```

```

Queue Name | Rx count | Drop count

```

```

-----+-----+-----
SW FORWARDING Q | 403038 | 0

```

```

ROUTING PROTOCOL Q | 159 | 0

```

```

ICMP Q | 0 | 0

```

```

HOST Q | 400 | 0

```

```

ACL LOGGING Q | 0 | 0

```

```

STP Q | 0 | 0

```

```

L2 PROTOCOL Q | 0 | 0

```

```

MCAST CONTROL Q | 0 | 0

```

```

BROADCAST Q | 0 | 0

```

```

REP Q | 0 | 0

```

```

BGP LDP Q | 0 | 0

```

```

CONTROL Q | 0 | 0

```

```

IP MPLS TTL Q | 0 | 0

```

```

DEFAULT MCAST Q | 0 | 0

```

```

MCAST ROUTE DATA Q | 0 | 0

```

```

MCAST MISMATCH Q | 0 | 0

```

```

RPF FAIL Q | 0 | 0

```

```

ROUTING THROTTLE Q | 0 | 0

```

```

MCAST Q | 0 | 0

```

```

MPLS OAM Q | 0 | 0

```

```

IP MPLS MTU Q | 0 | 0

```

```

PTP Q | 0 | 0

```

```

LINUX ND Q | 0 | 0

```

```

KEEPALIVE Q | 9256 | 0

```

```

ESMC Q | 0 | 0

```

```

FPGA BFD Q | 0 | 0

```

```

FPGA CCM Q | 0 | 0

```

```

FPGA CFE Q | 0 | 0

```

```

L2PT DUP Q | 0 | 0

```

```

TDM CTRL Q | 0 | 0

```

```
ICMP UNREACHABLE Q | 402918 | 0
SSFP Q | 0 | 0
MIRROT Q | 0 | 0
```

Use the **show platform software infrastructure inject** command to verify the platform software infrastructure inject:

```
PE1#show platform software infrastructure inject
Statistics for L3 injected packets:
27819 total inject pak, 24 failed
0 sent, 532 prerouted
24 non-CEF capable, 334 non-unicast
4965 IP, 20885 IPv6
0 MPLS, 0 Non-IP Tunnel
0 UDLR tunnel, 0 P2MP replicated mcast
0 Non-IP Fastswitched over Tunnel, 1945 legacy pak path
0 Other packet
0 IP fragmented
25318 normal, 0 nexthop
532 adjacency, 0 feature
0 undefined
0 pak find no adj, 0 no adj-id
4991 sb alloc, 25850 sb local
0 p2mcast failed count 0 p2mcast enqueue fail
0 unicast dhc
0 mobile ip
105 IPv6 NA
101 IPv6 NS
0 Transport failed cases
0 Grow packet buffer
2 Cant-l3-inject-pkts
per feature packet inject statistics
0 Feature multicast
0 Feature Edge Switching Service
0 Feature Session Border Controller
0 Feature interrupt level
0 Feature use outbound interface
0 Feature interrupt level with OCE
0 Feature ICMPv6 error message
0 Feature Session Border Controller media packet injection
0 Feature Tunnel Ethernet over GRE
0 Feature Secure Socket Layer Virtual Private Network
0 Feature EPC Wireshark injecting packets
0 Feature multicast overlay replication
Statistics for L2 injected packets:
356 total L2 inject pak, 0 failed
3 total BD inject pak, 0 failed
353 total EFP inject pak, 0 failed
0 total VLAN inject pak, 0 failed
```

Use the **show bgp l2vpn evpn detail** to verify BGP L2VPN EVPN detail:

```
PE1#show bgp l2vpn evpn detail
Route Distinguisher: 10.0.0.1:10
BGP routing table entry for [2][10.0.0.1:10][0][48][000000000001][0][*]/20, version 1027
Paths: (1 available, best #1, table evi_10)
Not advertised to any peer
Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
EVPN ESI: 00000000000000000000, Label 17
Extended Community: RT:1000:10
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 22:39:43 IST
```

```

BGP routing table entry for [2][10.0.0.1:10][0][48][000000000001][32][191.168.1.2]/24,
version 1021
Paths: (1 available, best #1, table evi_10)
Not advertised to any peer
Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
EVPN ESI: 00000000000000000000, Label1 17, Label2 22
Extended Community: RT:1000:1 RT:1000:10
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 22:39:42 IST
BGP routing table entry for [2][10.0.0.1:10][0][48][000000000001][128][1968:1::2]/36, version
1022
Paths: (1 available, best #1, table evi_10)
Not advertised to any peer
Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
EVPN ESI: 00000000000000000000, Label1 17, Label2 24
Extended Community: RT:1000:1 RT:1000:10
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 22:39:42 IST
BGP routing table entry for
[2][10.0.0.1:10][0][48][000000000001][128][FE80::200:FF:FE00:1]/36, version 1023
Paths: (1 available, best #1, table evi_10)
Not advertised to any peer
Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
EVPN ESI: 00000000000000000000, Label1 17, Label2 24
Extended Community: RT:1000:1 RT:1000:10
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 22:39:42 IST
Route Distinguisher: 10.0.0.1:11
BGP routing table entry for [2][10.0.0.1:11][0][48][000000000002][0][*]/20, version 1012
Paths: (1 available, best #1, table evi_11)
Not advertised to any peer
Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
EVPN ESI: 00000000000000000000, Label1 19
Extended Community: RT:1000:11
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 22:39:13 IST
BGP routing table entry for [2][10.0.0.1:11][0][48][000000000002][32][191.168.2.2]/24,
version 1011
Paths: (1 available, best #1, table evi_11)
Not advertised to any peer
Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
EVPN ESI: 00000000000000000000, Label1 19, Label2 22
Extended Community: RT:1000:1 RT:1000:11
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 22:39:12 IST
Route Distinguisher: 10.0.0.1:12
BGP routing table entry for [2][10.0.0.1:12][0][48][000000000003][0][*]/20, version 1028
Paths: (1 available, best #1, table evi_12)
Not advertised to any peer

```

```

Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
EVPN ESI: 00000000000000000000, Label1 21
Extended Community: RT:1000:12
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 22:39:43 IST
BGP routing table entry for [2][10.0.0.1:12][0][48][000000000003][32][191.168.3.2]/24,
version 1024
Paths: (1 available, best #1, table evi_12)
Not advertised to any peer
Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
EVPN ESI: 00000000000000000000, Label1 21, Label2 25
Extended Community: RT:1000:2 RT:1000:12
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 22:39:42 IST
BGP routing table entry for
[2][10.0.0.1:12][0][48][000000000003][128][FE80::200:FF:FE00:3]/36, version 1026
Paths: (1 available, best #1, table evi_12)
Not advertised to any peer
Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
EVPN ESI: 00000000000000000000, Label1 21, Label2 25
Extended Community: RT:1000:2 RT:1000:12
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 22:39:42 IST
Route Distinguisher: 10.0.0.1:10
BGP routing table entry for [3][10.0.0.1:10][0][32][10.0.0.1]/17, version 9
Paths: (1 available, best #1, table evi_10)
Not advertised to any peer
Refresh Epoch 1
Local
:: (via default) from 0.0.0.0 (10.0.0.1)
Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
Extended Community: RT:1000:10
PMSI Attribute: Flags:0x0, Tunnel type:IR, length 4, label:16 tunnel identifier: 0000 0000
rx pathid: 0, tx pathid: 0x0
Updated on Dec 16 2020 17:33:58 IST

```

