

EVPN Over MPLS with Integrated Routing and Bridging

EVPN Over MPLS routing and bridging (IRB) allows the device in an EVPN over MPLS network to perform both bridging and routing. IRB allows the devices to forward both Layer 2 or bridged and Layer 3 or routed traffic. 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 Distributed Anycast Gateway feature provides support for symmetric IRB model. This feature is supported only on Cisco ASR 1000 Series Aggregation Services Routers.

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 devicess 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.

To support Ethernet VPN over MPLS Integrated IRB Single-Homing (SH) with Distributed Anycast Gateway, you need to configure the following on the Cisco ASR 1000 Series Aggregation Services Router:

- Host IP-MAC learning in single homing setup
- Symmetric IRB for IP-VRF to IP-VRF inter-subnet traffic over MPLS
- Distributed Anycast Gateway with Bridge-Domain
- Host MAC-IP Mobility
- ARP/ND suppression
- Unknown Unicast Suppression

From Cisco IOS XE Cupertino 17.7.1a, Multi-Homing All-Active hosts is supported.

- Information about EVPN Over MPLS with Distributed Anycast Gatways, on page 2
- ARP and ND Flooding Suppression, on page 7
- MAC-IP Proxy Route for Multi-Homing All-Active Hosts with Symmetric IRB, on page 8
- Prerequisites for EVPN Over MPLS, on page 9
- Restrictions EVPN over MPLS, on page 9
- How to Configure EVPN over MPLS , on page 10

• Verification Examples for EVPN over MPLS, on page 14
• Advertising Proxy MAC-IP Route, on page 20
• Suppressing Unknown Unicast Flooding, on page 20
• Configuring Bridge Domain MAC Age Timer, on page 21
• Configuring ARP and ND Timers, on page 21
• Configuring IP Local Learning, Limits, and Timers, on page 21
• Configuring ARP and ND Flooding Suppression, on page 21
• Additional References for EVPN Single-Homing, on page 22
• Feature Information for EVPN MPLS IRB with Distributed Anycast Gateways, on page 22

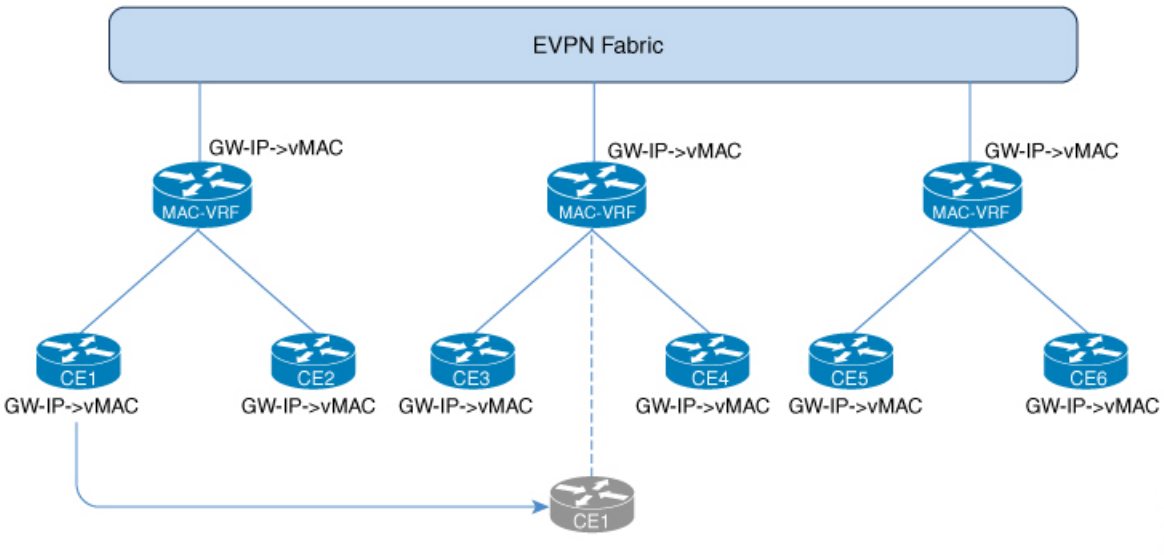
Information about EVPN Over MPLS with Distributed Anycast Gatways

Distributed Anycast Gateway with Bridge Domains

Distributed Anycast Gateway 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. 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 this topology, the Distributed Anycast Gateways 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 is supported and the duplication address detection (DAD) for IPv6 on the BDI interfaces on Distributed Anycast Gateways is disabled.

Figure 1: Distributed Anycast Gateway with Bridge Domains



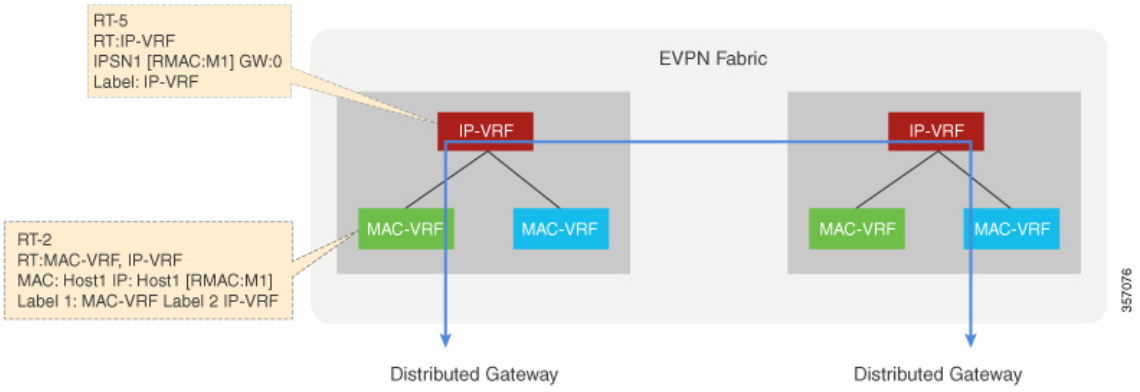
On the DAG, the bridge domain check if an Address Resolution Protocol or Neighbor Discovery Protocol packet from a local host is sent to the BDI (Gateway) IP addresses. If the packet is sent to BDI (Gateway) IP addresses, this packet is handled by local BDI and it is not flooded into the bridge domain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



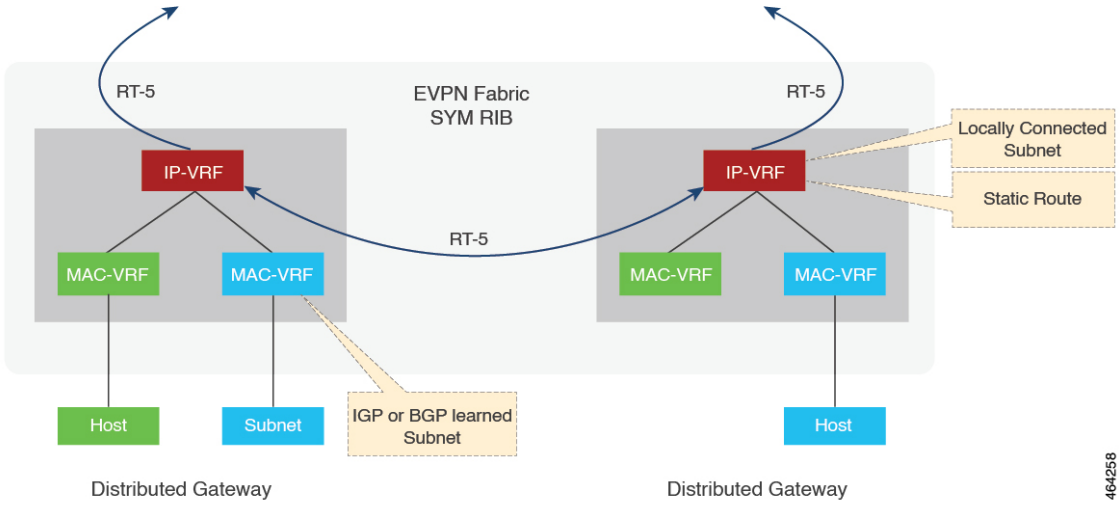
IP Prefix Route on Distributed Gateways

On Distributed Gateways, if the IP-VRF Stitching is configured and IP-VRF to EVPN redistribution is enabled, the IP Prefixes in IP-VRF are advertised as the RT-5 routes. These routes can be a locally connected subnet, a static route, or a route IP Prefix on the CE side.

On the receiving Distributed Gateways, RT-5 is imported into the corresponding IP-VRF and installed as a remote IP route.

In the case of a Distributed Anycast Gateway, many Distributed Gateways could advertise the same subnet prefix route. If the receiving Distributed Gateway has the same local subnet prefix, the local subnet prefix takes precedence. If a remote Distributed Gateway doesn't have the same local subnet prefix, remote routes are chosen or ECMP load balancing is used for forwarding. Host discovery and MAC-IP learning procedures are triggered when traffic is sent to the host before it's learned on any DAGs.

Figure 3: IP Prefix Route on Distributed Gateways



IP Prefix Route on Border Gateways

On Border Gateways that handle the Layer 3 hand-off to Multi-VRF, the IP-VRF Stitching needs to be configured and IP-VRF to EVPN redistribution needs to be enabled. The IP Prefix that is locally connected, a static route, or a learned route from the other side in IP-VRF is advertised as the RT-5 routes to EVPN Fabric. The IP Prefixes (RT-5) and Host Routes (RT-2) imported from EVPN are installed into IP-VRF and redistributed.

On Border Gateways that handle the Layer 3 hand-off to MPLS VPN (L3VPN), the IP-VRF Stitching needs to be configured. The cross-importing between VPN and EVPN Address-Family needs to be enabled as well. The IP Prefixes imported from VPN AF are advertised as RT-5 routes to EVPN Fabric. The IP Prefixes (RT-5) and Host Routes (RT-2) imported from EVPN AF are re-originated and advertised to the VPN side.

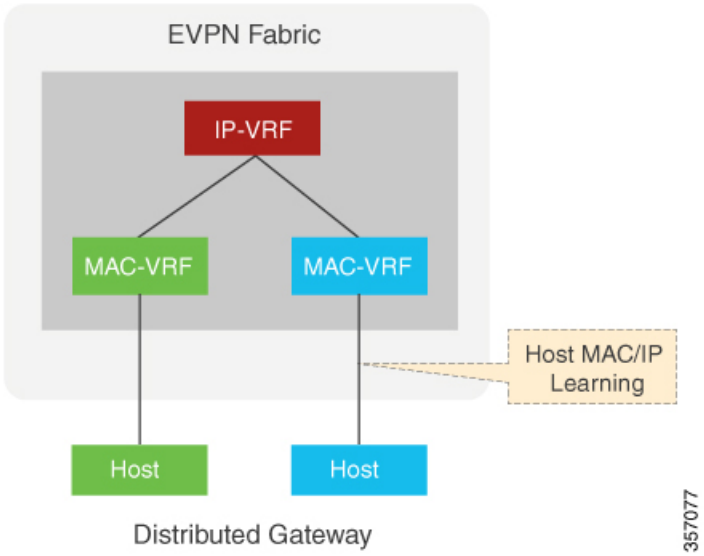
The host routes might be filtered if it's not desired to advertise or redistribute. The IP Prefix route can be used to inject the default route.

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.

When the age timer expires, the MAC-IP binding is deleted and withdrawn from the network. To avoid premature deletion of the MAC-IP bindings, the gateways initiates the ARP or ND procedure to refresh the binding entries. It also initiates the ARP or ND probe procedure with the data plane *MAC_AGE_OUT* event. The local *MAC_AGE_OUT* event triggers the probe of all the MAC-IP bindings derived from that specific MAC to refresh the binding entries.

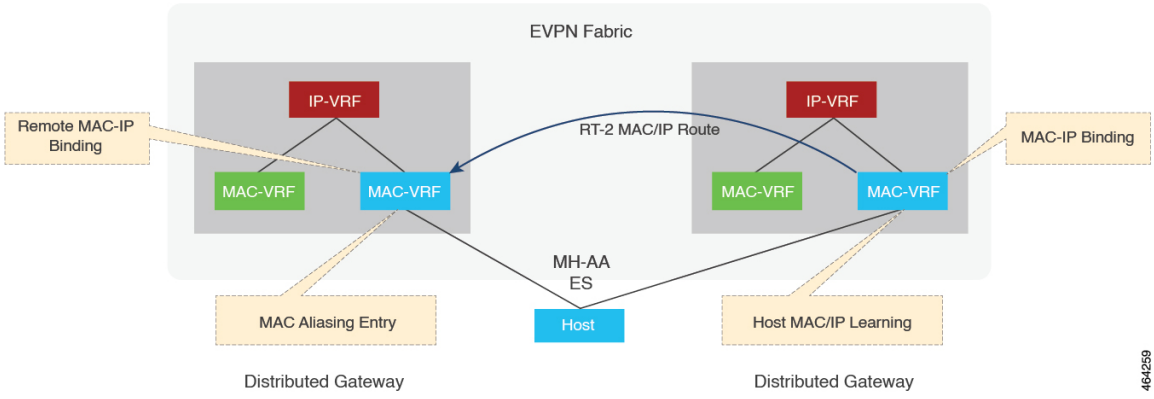
Figure 4: Host MAC-IP Binding on a Single-Homed DAG



Host MAC-IP Binding on Multi-Homing All-Active DAGs

For the hosts on a Multi-Homing All-Active Ethernet Segment, the Host MAC-IP Binding might be initially learned on only one of the multihoming peers. Other peers in the Multi-Homing Group rely on the received remote RT-2 to get the MAC-IP Binding or might locally learn the MAC-IP Bindings.

Figure 5: MAC-IP Binding on MH-AA Distributed Gateways



After the MAC-IP Binding is learned locally on any peer, an age timer (AGE_TIME) is applied to the binding entry, and the binding entry is refreshed whenever the host initiates ARP/ND to this peer. When the age timer expires, this locally learned MAC-IP Binding is deleted, and RT-2 is withdrawn or updated if RT-2 Proxy Route is enabled. When the age timer expires, the MAC-IP binding is deleted and withdrawn from the network.

To avoid premature deletion of the MAC-IP bindings, the gateways initiate the ARP or ND procedure to refresh the binding entries. This is triggered by a refresh timer (SEND_REFRESH_TIME). It also initiates the ARP or ND probe procedure with the data plane MAC_AGE_OUT event. The data plane usually has a

shorter age timer. The local MAC_AGE_OUT event triggers the probe of all the MAC-IP bindings derived from that specific MAC to refresh the binding entries across the Multi-Homing Group Peers.

During fast convergence and slow convergence, gateways can initiate the ARP/ND Refresh procedure using local Bindings or the previously received remote MAC-IP Bindings.

Host MAC-IP Mobility

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

- Host Move Learn from Data Packet and GARP
- Host Move Detection for Silent Host

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

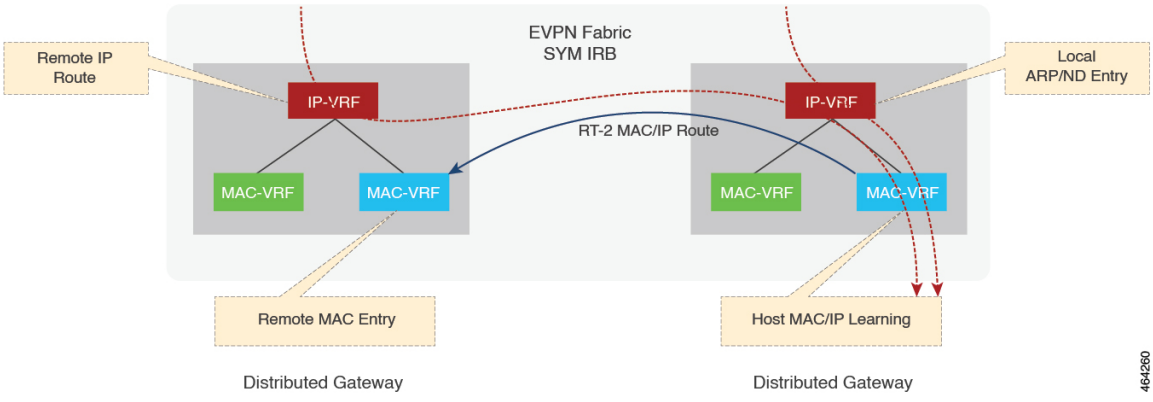
- Moving MAC from local to local
- Moving MAC from local to remote
- Moving MAC from remote to local
- Moving IP local to local
- Moving IP from local to remote
- Moving IP from remote to local

From Cisco IOS XE Cupertino 17.7.1a, the host MAC-IP mobility is supported for hosts on a Multi-Homing All-Active Ethernet Segment.

Host MAC-IP Synchronization

When the MAC-IP Route is imported on remote Distributed Gateways, it can be used for inter-subnet routing. Because Symmetric IRB uses the IP-VRF to IP-VRF Layer 3 Label or VNI for the routing, it doesn't need ARP/ND entries on remote Distributed Gateways, and a remote IP Route with Layer 3 Label (VNI) is installed.

Figure 6: MAC-IP Sync on Remote DAGs (SYM IRB)

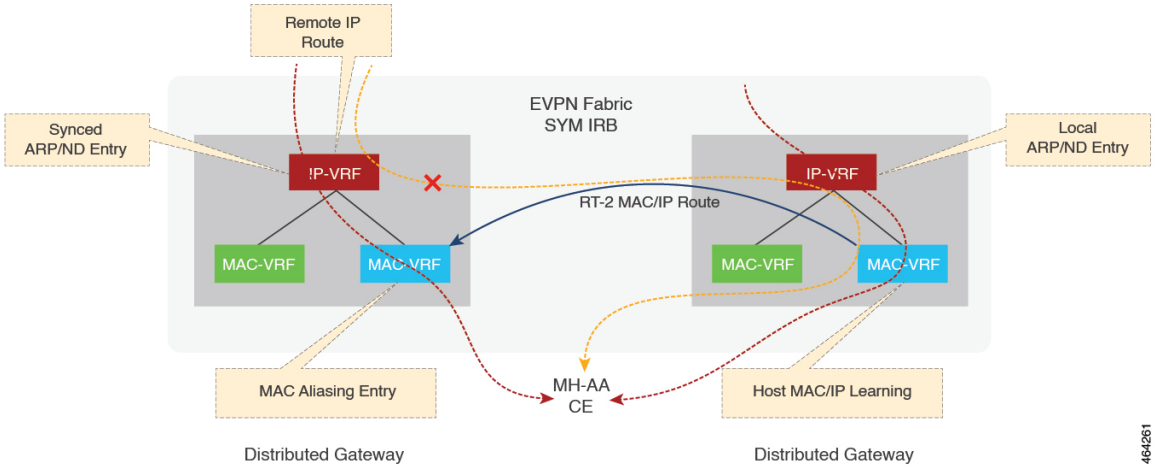


Considering the Multihoming Distributed Gateways, the remote Distributed Gateways could receive multiple MAC-IP routes from that multihoming group members with the same MAC and IP Binding for Layer 3 Load

Balancing. The fast convergence on the multihomed Gateways changes the Layer 2 Bridging in the MAC-VRF on the remote Gateways as it does, but it won't affect the remote IP Route entries and Layer 3 Load Balancing which are sourced from MAC-IP Routes. The withdrawal of the contributing remote MAC-IP route changes Layer 3 Load Balancing on the Remote Distributed Gateways. The withdrawal of all the contributing remote MAC-IP routes triggers the deletion of the remote MAC-IP Route on the Remote Distributed Gateway.

On a Multihoming All-Active Ethernet Segment, the Host MAC-IP Binding might be learned on only one of the multihoming peers. To enable the inter-subnet traffic forwarding to the Ethernet Segment LOCALLY on other multihoming peers, ARP/ND entries need to be synced to those peers via RT-2 MAC-IP Route.

Figure 7: MAC-IP Sync on MH All-Active DAGs (SYM IRB)



When a remote RT-2 MAC-IP Route for an MH-AA host is received from a multihoming peer, the RT-2 import on BGP installs a remote IP Route with the L3 Label (VNI). The ARP or ND entry needs to be installed or synced as well. From the forwarding point of view, both RIB (BGP and other routing components) and Adjacency (ARP or others) contribute to forwarding.

When BGP receives a Local or Aliasing MAC-IP Route from L2RIB, it is installed as static into RIB directly with a lower distance, so that RIB would prefer the local route, that is, the best path. During convergence or recovery, if the local or aliasing route is removed, RIB runs best path again.

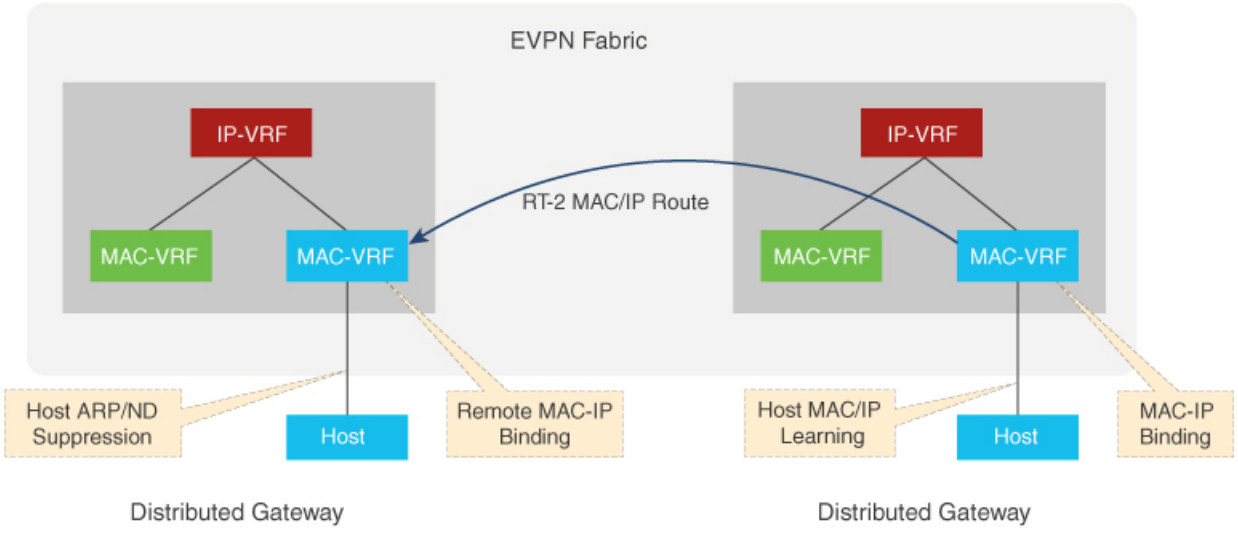
ARP and ND Flooding Suppression

The ARP and ND Flooding Suppression feature depends on device-tracking enabled on the same VLAN or interface. The Switch Integrated Security Features based (SIS-based) device tracking helps to track the presence,location, and movement of end-nodes in the network. SISF snoops traffic received by the device, extracts the device identity (MAC and IP address), and stores them in a binding table. SIS-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 fabric and to remote CEs such as host, router, and switch. For example, Address Resolution packets such as ARP (broadcast) and NS (multicast). The multicast and broadcast suppression capabilities helps to preserve bandwidth in wireless networks.

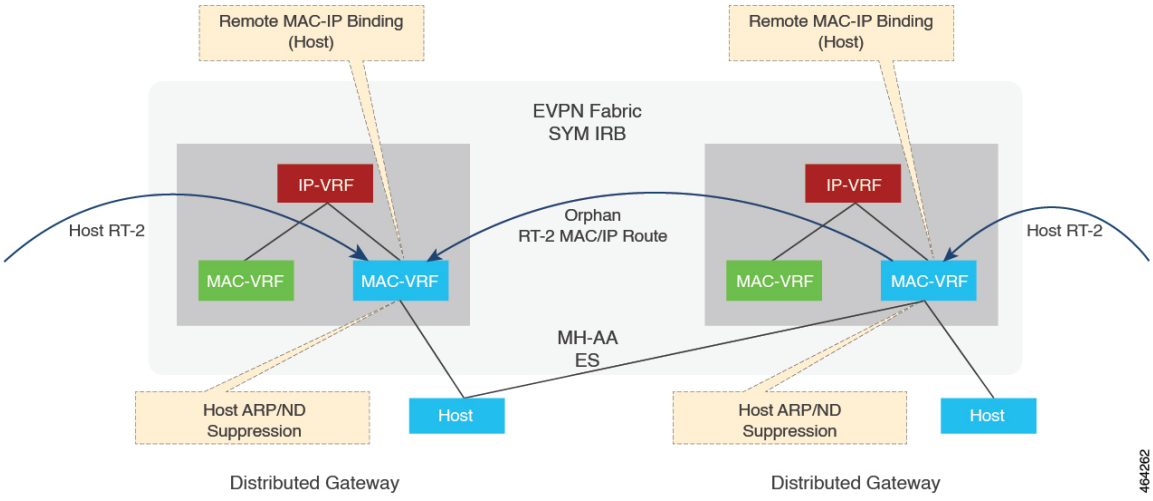
This feature helps to suppress the broadcast (ARP) or link-local multicast (NDP) messages circulating in the layer 2 domain and the packets are relayed after converting their L2 addresses to unicast. By default, this feature is enable and you can use the **disable flooding suppression** command to disable flooding suppression.

Figure 8: ARP and ND Flooding Suppression



From Cisco IOS XE Cupertino 17.7.1a, ARP and ND flooding suppression is supported on Multi-Homing All-Active (MH-AA) hosts. This feature reduces the ARP/ND Flooding traffic across a fabric since the gateway attached to the host might have already learned the MAC-IP Bindings from the control plane. By default, this feature is enabled. For more information, see *Disabling ARP or ND Suppression*.

Figure 9: ARP and ND Suppression on MH-AA Distributed Gateways

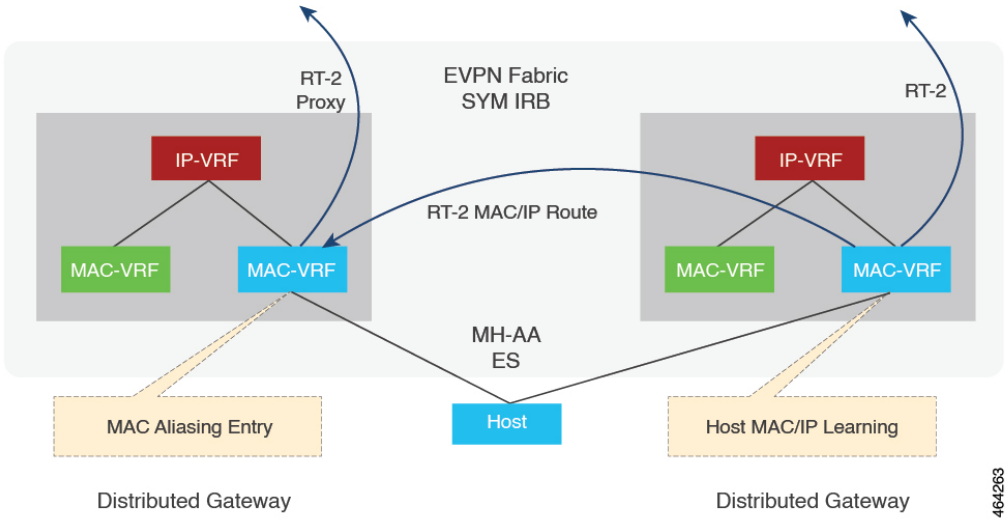


MAC-IP Proxy Route for Multi-Homing All-Active Hosts with Symmetric IRB

Due to load balancing between hosts and the gateways in a Multi-Homing Group, the host MAC-IP binding might be learned and advertised from only one or more gateways. For routing traffic from a remote gateway,

the host is only reachable through the advertising gateways. For Layer 2 Bridging, an aliasing path using the EAD Per EVI route is implemented to achieve Layer 2 load balancing. However, the aliasing path is only valid in a Layer 2 topology (MAC-VRF), and cannot be used in a Layer 3 topology (IP-VRF). To achieve Layer 3 ECMP for a Multi-Homing All-Active Host, a peer in the Multi-Homing Group can choose to be a proxy or readvertise a remote MAC-IP Route that it received for the All-Active host. This feature is enabled by default, and it can be implemented by a peer if it has not locally learned the MAC-IP route.

Figure 10: MAC-IP Proxy on MH All-Active DAGs



Prerequisites for EVPN Over MPLS

This feature requires the following configuration on Cisco ASR 1000 Series Aggregation Services Routers:

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

Restrictions EVPN over MPLS

- Only Dual-Homing (two peers) Active-Active is supported.
- Only Symmetric IRB is supported. Asymmetric IRB and centralized IRB are not supported for BGP EVPN over MPLS.
- Only Gateway virtual MAC address is supported.

- Global VRF is not supported for MPLS IRB.
- VPLS Stitching is not covered and verified for BGP EVPN over MPLS.
- RT-5-only based routing is not applicable to Multi-Homing All-Active Subnets on DAGs if RT-2 is disabled. This routing is not applicable because Multi-Homing peers need RT-2 MAC-IP route for ARP/ND SYNC.
- SISF security feature is not supported on Multi-Homing All-Active DAGs.
- Preserve the original next-hop in iBGP peering for EVPN and do not modify it by default. Use the VTEP IP address as the next-hop for EVPN-VXLAN.
- Use EVPN Type 5 (IP Prefix) routes with the RT5 gateway IP to maintain the same next-hop as in the VPN routing table. Always verify next-hop reachability on the VTEP to ensure proper route installation. If you need to modify the next-hop for non-VXLAN cases, configure **next-hop-self** command on the BGP neighbor or apply a route-map with set ip next-hop. This approach ensures consistent and correct next-hop behavior in EVPN deployments.

How to Configure EVPN over MPLS

Configure Basic EVPN over MPLS

Use the following steps to configure the basic EVPN over MPLS:

Configuring Layer 2 Virtual Private Network EVPN

```
interface Loopback2
 description L2VPN EVPN ROUTER-ID
 ip address 1.1.1.3 255.255.255.255
 ip ospf 1 area 0
 ipv6 address ABCD:1::3/128
end

l2vpn evpn
!! Only Ingress Replication is supported for EVPN MPLS
 replication-type ingress
 router-id Loopback2
end
!
```

Configuring Layer 2 Virtual Private Network EVPN Instance

```
!! EVPN MPLS supports vlan-based, vlan-aware, vlan-bundle types for L2
!! vlan-bundle type won't be supported for IRB
!! default encapsulation is MPLS
l2vpn evpn instance 1 vlan-based
```

Configuring Interface, Bridge-Domain and EFP

```
interface Ethernet0/1
 no ip address
 service instance 11 ethernet
  encapsulation dot1q 11
end
```

```
bridge-domain 11
 !! link point of BD, EFP and EVI
 member Ethernet0/1 service-instance 11
 member evpn-instance 1
end
```

Configuring Layer 2 Virtual Private Network EVPN BGP

```
interface Loopback0
 description BGP UPDATE SOURCE
 ip address 1.1.1.1 255.255.255.255
 ip ospf 1 area 0
 ipv6 address ABCD:1::1/128
end

router bgp 100
 bgp log-neighbor-changes
 bgp graceful-restart
 neighbor 99.99.99.99 remote-as 100
 neighbor 99.99.99.99 update-source Loopback0
 !
 address-family l2vpn evpn
  neighbor 99.99.99.99 activate
  neighbor 99.99.99.99 send-community both
 exit-address-family
end
```

Configure Basic EVPN over MPLS with IRB

Use the following steps to configure the basic EVPN over MPLS with IRB.

Configuring IP-VRF

```
vrf definition red
 rd 100:1
 !
 address-family ipv4
  route-target export 100:100
  route-target import 100:100
  route-target export 100:100 stitching
  route-target import 100:100 stitching
 exit-address-family
 !
 address-family ipv6
  route-target export 100:200
  route-target import 100:200
  route-target export 100:200 stitching
  route-target import 100:200 stitching
 exit-address-family
end
```

Configuring Bridge-Domain IRB Interface

```
interface BDI11
 !! virtual MAC for Distributed Anycast Gateway
 mac-address 0011.0011.0011
 vrf forwarding red
 ip address 192.168.11.254 255.255.255.0
 ipv6 address 2001:11::254/64
```

```
 encapsulation dot1Q 11
end
```

Configuring BGP IRB

```
router bgp 100
 bgp log-neighbor-changes
 bgp graceful-restart
 neighbor 99.99.99.99 remote-as 100
 neighbor 99.99.99.99 update-source Loopback0
 !
 address-family l2vpn evpn
  neighbor 99.99.99.99 activate
  neighbor 99.99.99.99 send-community both
 exit-address-family
!! IP Prefix Advertisement
address-family ipv4 vrf red
  advertise l2vpn evpn
  redistribute connected
exit-address-family
address-family ipv6 vrf red
  advertise l2vpn evpn
  redistribute connected
exit-address-family
End
```

EVPN over MPLS with Multi-VRF Hand-off

Use the following steps to configure EVPN-MPLS Multi-VRF:

Configuring BGP Multi-VRF on Border Gateways

```
router bgp 100
 bgp log-neighbor-changes
 bgp graceful-restart
 neighbor 99.99.99.99 remote-as 100
 neighbor 99.99.99.99 update-source Loopback0
 !
 address-family l2vpn evpn
  neighbor 99.99.99.99 activate
  neighbor 99.99.99.99 send-community both
 exit-address-family
address-family ipv4 vrf red
  advertise l2vpn evpn
  redistribute connected
  neighbor 192.168.0.1 remote-as 10
  neighbor 192.168.0.1 activate
  !! If using MPLS between PE and CE
  neighbor 192.168.0.1 send-label
 exit-address-family
 !
 address-family ipv6 vrf red
  advertise l2vpn evpn
  redistribute connected
  neighbor 2001:0::1 remote-as 10
  neighbor 2001:0::1 activate
  !! If using MPLS between PE and CE
  !! neighbor 2001:0::1 send-label <= MPLS is only supported on IPv4 Core (6PE)
 exit-address-family
```

EVPN over MPLS L3VPN Hand-offf

Use the following steps to configure the VPN over MPLS Layer-3 VPN.

Configuring BGP VPNv4/VPNv6 on Border Gateways

```
router bgp 100
 bgp log-neighbor-changes
 bgp graceful-restart
 !! eBGP peering
 neighbor 10.5.0.1 remote-as 10
 neighbor 99.99.99.99 remote-as 100
 neighbor 99.99.99.99 update-source Loopback0
 !
 address-family vpnv4
 !! EVPN to VPNv4
  import l2vpn evpn re-originate
  neighbor 10.5.0.1 activate
  neighbor 10.5.0.1 send-community both
 exit-address-family
 !
 address-family vpnv6
 !! EVPN to VPNv6
  import l2vpn evpn re-originate
  neighbor 10.5.0.1 activate
  neighbor 10.5.0.1 send-community both
 exit-address-family
 !
 address-family l2vpn evpn
 !! VNPv4/VPNv6 to EVPN
  import vpnv4 unicast re-originate
  import vpnv6 unicast re-originate
  neighbor 99.99.99.99 activate
  neighbor 99.99.99.99 send-community both
  neighbor 99.99.99.99 next-hop-self
 exit-address-family
```

Layer 2 Multihoming Configuration for EVPN over MPLS

Use the following steps to configure Layer 2 Multihoming for EVPN over MPLS:

Configuring Ethernet Segment

```
l2vpn evpn ethernet-segment 1
 !! Support both type 3 and type 0
 identifier type 3 system-mac aabb.0000.0001
 !! Only all-active is support for EVPN MPLS
 redundancy all-active
 df-election wait-time 3
end
```

Configuring L2VPN EVPN Instance

```
!! EVPN MPLS supports vlan-based, vlan-aware, vlan-bundle types for L2
!! vlan-bundle type won’t be supported for IRB
!! default encapsulation is MPLS
l2vpn evpn instance 1 vlan-based
```

Configuring Interface, Bridge Domain, and EFP

```
bridge-domain 11
!! link point of BD, EFP, and EVI
member Port-channel1 service-instance 11
member evpn-instance 1

interface Port-channel1
no ip address
no negotiation auto
no mop enabled
no mop sysid
!! link point of interface and ESI
evpn ethernet-segment 1
lacp device-id 0005.0005.0005
service instance 11 ethernet
encapsulation dot1q 11
!
!

!! interface linking PE to MH CE
interface GigabitEthernet0/0/1
no ip address
negotiation auto
lacp rate fast
!! link to Port-channel
channel-group 1 mode active
```

Verification Examples for EVPN over MPLS

Show Device-tracking Policy

Use the following command to verify that all the SISF feature policies are attached to bridge domain:

```
show device-tracking policies
Target               Type  Policy               Feature         Target range
bd 11                 bd  evpn-device-track Device-tracking      bd     all
bd 11                 bd  evpn-flood-suppress  Flooding Suppress bd     all

show device-tracking policy evpn-device-track
Policy evpn-device-track configuration:
  security-level glean
  device-role node
  gleaning from Neighbor Discovery
  gleaning from DHCP
  gleaning from ARP
  gleaning from DHCP4
  NOT gleaning from protocol unkn
Policy evpn-device-track is applied on the following targets:
Target               Type  Policy               Feature        Target range
bd 11                  bd  evpn-device-track Device-tracking bd    all
```

Show MAC and IP Binding Tables on Local PE

Use the following command to verify that the MAC and IP Binding tables are on local PE:

```
show device-tracking database

    Network Layer Address      Link Layer Address Interface   vlan prlvl  age   state
```

```
Time left

ARP 192.168.255.3              0050.56b0.9d09  Te1/0/5         11  0005   34s  REACHABLE
90 s
ARP 192.168.1.8                0050.56b0.c8f5  Te1/0/6         11  0005   14s  REACHABLE
107 s try 0
ND  FE80::250:56FF:FEB0:C8F5   0050.56b0.c8f5  Te1/0/6         11  0005    7s  REACHABLE
116 s
ND  FE80::250:56FF:FEB0:9D09   0050.56b0.9d09  Te1/0/5         11  0005   28s  REACHABLE
95 s
ND  2001:192:168:255::3        0050.56b0.9d09  Te1/0/5         11  0005   32s  REACHABLE
90 s
ND  2001:192:168:1::8          0050.56b0.c8f5  Te1/0/6         11  0005   12s  REACHABLE
111 s
```

```
show l2vpn evpn mac ip
IP Address                EVI   BD.   MAC Address    Next Hop(s)
------------------------- ----- ----- -------------- -------------------------
192.168.1.8               1     11    0050.56b0.c8f5 Te1/0/6:11
192.168.255.3             1     11    0050.56b0.9d09 Te1/0/5:11
2001:192:168:1::8         1     11    0050.56b0.c8f5 Te1/0/6:11
2001:192:168:255::3       1     11    0050.56b0.9d09 Te1/0/5:11
FE80::250:56FF:FEB0:9D09  1     11    0050.56b0.9d09 Te1/0/5:11
```

MAC and IP Binding Entries on Remote PE

Use the following command to verify that the MAC and IP Binding tables are on remote PE:

```
show l2vpn evpn mac ip

IP Address                EVI   VLAN  MAC Address    Next Hop(s)
------------------------- ----- ----- -------------- -------------------------
192.168.1.8               1     11    0050.56b0.c8f5 1.1.1.101
192.168.255.3             1     11    0050.56b0.9d09 1.1.1.101
2001:192:168:1::8         1     11    0050.56b0.c8f5 1.1.1.101
2001:192:168:255::3       1     11    0050.56b0.9d09 1.1.1.101
FE80::250:56FF:FEB0:9D09  1     11    0050.56b0.9d09 1.1.1.101
FE80::250:56FF:FEB0:C8F5  1     11    0050.56b0.c8f5 1.1.1.101
```

Displays the device-tracking database on Cisco ASR 1000 Series Aggregation Services Routers.

```
show device-tracking database

 Network Layer Address              Link Layer Address Interface        bd   prlvl  age
 state     Time left
L   192.168.1.100                        aabb.cc00.01ff  BD11             11  0100 2628mn
 REACHABLE
L   FE80::A8BB:CCFF:FE00:1FF               aabb.cc00.01ff  BD11             11  0100 2628mn
 REACHABLE
L   2001:192:168:1::100                    aabb.cc00.01ff  BD11             11  0100 2628mn
 REACHABLE
```

Show MAC and IP Binding Tables on Local PE in a Multi-Homing Setup

Use the following commands to verify the MAC and IP binding tables are on a local PE in a Multi-Homing setup.

```
#show device-tracking database
Binding Table has 6 entries, 3 dynamic (limit 1000000)
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
L   192.168.12.254                           0012.0012.0012         BD12       12
0100       16mn       REACHABLE
ARP 192.168.12.3                             aabb.aabb.0012         Po1        12
0005       8mn        STALE     try 0 1662 s
ND  FE80::A8BB:AAFF:FEBB:12                  aabb.aabb.0012         Po1        12
0005       3mn        REACHABLE  120 s
L   FE80::212:FF:FE12:12                     0012.0012.0012         BD12       12
0100       16mn       REACHABLE
L   2001:12::254                             0012.0012.0012         BD12       12
0100       16mn       REACHABLE
ND  2001:12::3                               aabb.aabb.0012         Po1        12
0005       8mn        STALE     try 0 1618 s
```

The PE that locally learned the MH MAC/IPs (192.168.12.3, 2001:12::3) has a MAC or IP binding in SISF.

```
#show l2vpn evpn mac ip
IP Address                            EVI   BD    MAC Address    Next Hop(s)
------------------------------------- ----- ----- -------------- -----------
192.168.12.3                          2     12    aabb.aabb.0012 Po1:12
                                                                   3.3.3.1
2001:12::3                            2     12    aabb.aabb.0012 Po1:12
                                                                   3.3.3.1
FE80::A8BB:AAFF:FEBB:12               2     12    aabb.aabb.0012 Po1:12
                                                                   3.3.3.1
```

The **Next Hops** column shows a local interface and a next hop to the other MH PE.

```
#show l2vpn evpn mac ip summary
EVI   BD    Ether Tag  Remote IP  Local IP   Dup IP
----- ----- ---------- ---------- ---------- ----------
2     12    0          0          3          0

Total                  0          3          0
```

The MH MAC or IPs are Local in the PE that locally learned these MAC or IPs.

Show MAC and IP Binding Tables on Local Proxy PE in a Multi-Homing Setup

Use the following commands to verify the MAC and IP binding tables are on a local proxy PE in a Multi-Homing setup.

```
#show device-tracking database
Binding Table has 3 entries, 0 dynamic (limit 1000000)
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
L   192.168.12.254                           0012.0012.0012         BD12       12
0100       254mn      REACHABLE
L   FE80::212:FF:FE12:12                     0012.0012.0012         BD12       12
0100       254mn      REACHABLE
L   2001:12::254                             0012.0012.0012         BD12       12
0100       254mn      REACHABLE
```

The PE that is a proxy for the MH MAC or IPs does not have a MAC or IP binding in SISF for the MH MAC or IPs.

```
#show l2vpn evpn mac ip
IP Address                            EVI   BD    MAC Address    Next Hop(s)
------------------------------------- ----- ----- -------------- -----------
192.168.12.3                          2     12    aabb.aabb.0012 Po1:12
                                                                   4.4.4.1
2001:12::3                            2     12    aabb.aabb.0012 Po1:12
                                                                   4.4.4.1
FE80::A8BB:AAFF:FEBB:12               2     12    aabb.aabb.0012 Po1:12
                                                                   4.4.4.1
```

The **Next Hops** column shows a local interface and a next hop to the other MH PE.

```
#show l2vpn evpn mac ip summary
EVI   BD    Ether Tag  Remote IP  Local IP   Dup IP
----- ----- ---------- ---------- ---------- ----------
2     12    0          0          3          0

Total                  0          3          0
```

The MH MAC/IPs are Local in the PE that is proxy for these MAC or IPs.

Show MAC and IP Binding Tables on Remote PE in a Multi-Homing Setup

Use the following commands to verify the MAC and IP binding tables are on a remote PE in a Multi-Homing setup.

```
#show device-tracking database
Binding Table has 3 entries, 0 dynamic (limit 1000000)
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
L   192.168.12.254                           0012.0012.0012         BD12       12
0100       127s       DOWN
L   FE80::212:FF:FE12:12                     0012.0012.0012         BD12       12
0100       127s       DOWN
L   2001:12::254                             0012.0012.0012         BD12       12
0100       127s       DOWN
```

The remote PE does not have a MAC/IP binding in SISF for the MH MAC/IPs.

```
#show l2vpn evpn mac ip
IP Address                            EVI   BD    MAC Address    Next Hop(s)
------------------------------------- ----- ----- -------------- -----------
```

```
192.168.12.3                              2     12    aabb.aabb.0012 3.3.3.1
                                                                     4.4.4.1
2001:12::3                                2     12    aabb.aabb.0012 3.3.3.1
                                                                     4.4.4.1
FE80::A8BB:AAFF:FEBB:12                   2     12    aabb.aabb.0012 3.3.3.1
                                                                     4.4.4.1
```

The **Next Hops** column shows remote next hops to the MH PEs for the MH MAC/IPs.

```
#show l2vpn evpn mac ip summary
EVI   BD    Ether Tag  Remote IP  Local IP   Dup IP
----- ----- ---------- ---------- ---------- ----------
2     12    0          3          0          0

Total                  3          0          0
```

The MH MAC/IPs are Remote in the remote PE.

Device Tracking Counters

Use the following command to verify the device-tracking counters:

```
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]

Received multicast messages on bd 11   :
Protocol        Protocol message
NDP             RS[4] NS[8] NA[8]
DHCPv6
ARP             REQ[6] REP[4]
DHCPv4

Bridged messages from bd 11   :
Protocol        Protocol message
NDP             RS[4] RA[4] NS[2685] NA[1778]
DHCPv6
ARP             REQ[1029] REP[4]
Protocol        Protocol message
NDP
DHCPv6
ARP
DHCPv4
ACD&DAD

Probe message on bd 11   :
Type            Protocol message
PROBE_SEND      NS[908] REQ[1023]
PROBE_REPLY     NA[907] REP[995]

Dropped messages on bd 11   :
Feature             Protocol Msg [Total dropped]
Device-tracking:    NDP      NA  [907]
                    reason:  Silent drop [907]

                    ARP      REP [995]
                    reason:  Silent drop [995]
```

QFP BD SISF Statistics and Snoop Protocols

Use the following command to verify the QFP BD SISF statistics and snoop protocols:

```
show platform hard qfp ac feat bridge data 10
infra-1001x-2#show pla hard qfp ac fea brid da 10
QFP L2BD Bridge Domain information

BD id                  : 10
State enabled          : Yes
Aging timeout (sec)    : 300
……
Unknown unicast olist  : Yes
otv_aed_enabled : No
otv_enabled : No
mcast_snooping_enabled : No
Feature : evpn
SISF snoop protocols   : arp, ndp, dhcpv4, dhcpv6
Mac learned            : 1
……

Bridge Domain statistics

Total bridged               pkts : 577        bytes: 48602
Total unknown unicast       pkts : 7          bytes: 636
Total broadcasted           pkts : 1737       bytes: 181506
Total to BDI                pkts : 0          bytes: 0
Total injected              pkts : 1056       bytes: 105012
……
Total UUF suppression drop  pkts : 0          bytes: 0
Total sisf ctrl punt        pkts : 1577       bytes: 143058
```

Unknown Unicast Flooding Suppression

Use the following command to verify the unknown unicast flooding suppression status:

```
#show bridge-domain 12
Bridge-domain 12 (3 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 5 minute(s)
Unknown Unicast Flooding Suppression: Enabled
```

Debug Commands (EVPN)

Use the following debug command to troubleshoot:

EVPN Debug Commands

- debug l2vpn evpn event
- debug l2vpn evpn event detail

Event Trace Debug Commands

- monitor event-trace sequence-number
- monitor event-trace timestamps datetime msec localtime
- monitor event-trace evpn event size 1000000

- monitor event-trace evpn event include event error major detai
- show tech-support evpn

L2RIB Debug Command

- debug l2rib event
- debug l2rib event detail
- debug l2rib error

Debug Commands (SISF)

- debug device-tracking switcher
- debug device-tracking parser
- debug device-tracking flooding-suppression
- debug device-tracking hw-api
- show device-tracking events
- show device-tracking messages
- show device-tracking counters bd <bd-id>
- show tech-support sisf
- debug platform software fhs all
- Debug ip bgp all update
- debug ip bgp l2vpn evpn evi event [detail]
- debug ip bgp l2vpn evpn evi context [detail]

Advertising Proxy MAC-IP Route

Proxy MAC-IP route is enabled by default. Use the following command to disable proxy MAC-IP route:

```
l2vpn evpn
 multihoming proxy-mac-ip disable
```

Suppressing Unknown Unicast Flooding

Use the following command to suppress unknown unicast flooding:

```
flooding-suppression unknown-unicast
```

This command is supported only at the bridge domain level. By default, suppression of unknown unicast flooding is disabled.

Configuring Bridge Domain MAC Age Timer

Use the following command to configure the bridge domain MAC Age timer:

```
bridge-domain 11
 mac aging-time 10
```

The default aging time is 5 minutes for a bridge domain and 30 minutes for overlay bridge domains. The range is from 1 to 600 minutes.

Configuring ARP and ND Timers

Use the following command to configure the ARP timeout:

```
int BDI11
  arp timeout 600
```

Use the following command to configure the ND cache expiry:

```
int BDI12
  ipv6 nd cache expire 300
```

Configuring IP Local Learning, Limits, and Timers

Use the following command to disable IP local learning from the data plane:

```
l2vpn evpn
 ip local-learning disable
```

Use the following command to limit the number of locally learned IP addresses that can be stored:

```
ip local-learning limit per-mac ipv4
ip local-learning limit per-mac ipv6
```

The default number of IPv4 addresses is 4 and IPv6 addresses is 12.

Use the following command to configure timers:

```
ip local-learning time poll| reachable | stale time
```

The default polling interval is 1 minute, the reachable lifetime is 5 minutes, and the stale lifetime is 30 minutes.

Configuring ARP and ND Flooding Suppression

To configure the ARP and ND flooding suppression, perform the following steps.



Note By default, ARP/ND flooding suppression is enabled.

```
Device(config)#l2vpn evpn
Device(config-evpn)#flooding-suppression address-resolution ?
  disable  Disable flooding suppression
```

Additional References for EVPN Single-Homing

Standards and RFCs

Standard	Title
RFC 7432	BGP MPLS-Based Ethernet VPN

Feature Information for EVPN MPLS IRB with Distributed Anycast Gateways

Feature Name	Releases	Feature Information
EVPN MPLS IRB with Distributed Anycast Gateways	Cisco IOS XE Bengaluru 17.4.	The Ethernet VPN over MPLS Integrated Routing and Bridging (IRB) Single-Homing (SH) with Distributed Anycast Gateway feature provides support for symmetric IRB) model on SH Distributed Anycast Gateways for BGP EVPN over MPLS on Cisco ASR 1000 Series Aggregation Services Routers.
ARP and ND Flooding Suppression	Cisco IOS XE Bengaluru 17.4.	This feature helps to suppress the broadcast (ARP) or link-local multicast (NDP) messages circulating in the layer 2 domain, by either dropping them, or rewriting the layer2 destination from broadcast/multicast to unicast.
Support for EVPN over MPLS IRB Multi-Homing	Cisco IOS XE Cupertino 17.7.1a	This feature enables redundant network connectivity via Multi-homing by allowing a CE device to connect to more than one PE device therefore preventing disruptions in the network. Note that only dual-homing is supported in this release.