Implementing Layer 2 Multicast

Implementing IGMP Snooping

IGMP snooping provides a way to constrain multicast traffic at Layer 2.

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
IGMP snooping	Release 25.4.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100])(select variants only*) *This feature is supported on: • 8011-12G12X4Y-A • 8011-12G12X4Y-D

Feature Name	Release Information	Feature Description
IGMP snooping	Release 25.2.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: K100, P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) * This feature is supported on: • 8212-48FH-M • 8711-32FH-M • 88-LC1-36EH • 8711-48Z-M
IGMP snooping	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100, P100], 8010 [ASIC: A100])(select variants only*) This feature is enhanced to support: • IGMP snooping on BVI, and • IGMP versions IGMPv2 and IGMPv3, providing backward compatibility and enhanced features like source-based filtering. * This feature is supported on: • 8712-MOD-M • 8011-4G24Y4H-I

Feature Name	Release Information	Feature Description
IGMP snooping	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) IGMP Snooping is used in Layer 2 multicast to optimize the distribution of multicast traffic. IGMP membership report messages are examined from hosts to determine which interfaces are connected to devices interested in receiving multicast traffic. This helps in reducing unnecessary traffic by ensuring that multicast data is only sent to ports with interested receivers, rather than flooding the entire VLAN. The benefit of IGMP snooping is bandwidth optimization that limits multicast traffic to only the necessary ports. *This feature is supported on: • 8212-32FH-M • 8711-32FH-M • 8712-MOD-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Internet Group Management Protocol (IGMP) snooping restricts multicast flows at Layer 2 to only those segments with at least one interested receiver. This module describes how to implement IGMP snooping.

Note Multicast traffic without Spanning-Tree protocol is supported at Layer 2 for multicast traffic without snooping enabled.

Prerequisites for IGMP Snooping

Before implementing IGMP snooping, make sure that the network is configured with a Layer 2 VPN (L2VPN).

Supported Features and Restrictions for IGMP Snooping

- EVPN dual-homed Active Active (AA) IGMP State Sync using IGMP snooping profile is not supported.
- Starting with Cisco IOS XR Release 25.1.1, IGMP snooping on BVI is supported.
- Starting with Cisco IOS XR Release 25.1.1, both IGMP versions IGMPv2 and IGMPv3 are supported.
- IGMP snooping is supported only under L2VPN bridge domains.
- Explicit host tracking (an IGMPv3 snooping feature) is not supported.
- IGMPv1 is not supported.
- ISSU is not supported on Layer 2 Multicast.
- IGMPv3-exclude is not supported in EVPN multi-homing or proxy scenarios.
- PIM control packets are supported when snooping is enabled.

These restrictions are applicable for 8712-MOD-M and 8011-4G24Y4H-I routers:

- Only PIM and PIMv6 hello packets are supported when snooping is enabled.
- Explicit tracking should be enabled for IGMPv3.

Information About IGMP Snooping

IGMP Snooping Overview

Description of Basic Functions

IGMP snooping provides a way to constrain multicast traffic at Layer 2. By snooping the IGMP membership reports sent by hosts in the bridge domain, the IGMP snooping application can set up Layer 2 multicast forwarding tables to deliver traffic only to ports with at least one interested member, significantly reducing the volume of multicast traffic.

Configured at Layer 3, IGMP provides a means for hosts in an IPv4 multicast network to indicate which multicast traffic they are interested in and for routers to control and limit the flow of multicast traffic in the network at Layer 3.

IGMP snooping uses the information in IGMP membership report messages to build corresponding information in the forwarding tables to restrict IP multicast traffic at Layer 2. The forwarding table entries are in the form <Route, OIF List>, where:

- Route is a <*, G> route or <S, G> route, where * is any source, G is group and S is the source.
- OIF List comprises all bridge ports that have sent IGMP membership reports for the specified route plus all multicast router (mrouter) ports in the bridge domain.

Implemented in a multicast network, IGMP snooping has the following attributes:

- In its basic form, it reduces bandwidth consumption by reducing multicast traffic that would otherwise flood an entire VPLS bridge domain.
- With the use of some optional configurations, it provides security between bridge domains by filtering the IGMP reports received from hosts on one bridge port and preventing leakage towards the hosts on other bridge ports.
- Using optional configurations, reduces the traffic impact on upstream IP multicast routers by suppressing IGMP membership reports (IGMPv2) or by acting as an IGMP proxy reporter (IGMPv3) to the upstream IP multicast router.

High Availability Features

All high availability features apply to the IGMP snooping processes with no additional configuration beyond enabling IGMP snooping. The following high availability features are supported:

- Process restarts
- RP Failover
- Non-Stop Forwarding (NSF)—Forwarding continues unaffected while the control plane is restored following a process restart or route processor (RP) failover.
- Line card online insertion and removal (OIR)

Bridge Domain Support

IGMP snooping operates at the bridge domain level. When IGMP snooping is enabled on a bridge domain, the snooping functionality applies to all ports under the bridge domain, including:

- Physical ports under the bridge domain.
- Ethernet flow points (EFPs)—An EFP can be a VLAN, VLAN range, list of VLANs, or an entire interface port.
- Ethernet bundles—Ethernet bundles include IEEE 802.3ad link bundles and Cisco EtherChannel bundles. From the perspective of the IGMP snooping application, an Ethernet bundle is just another EFP. The forwarding application in the router randomly nominates a single port from the bundle to carry the multicast traffic.

Multicast Router Port

A Multicast router (Mrouter) port is a port that connects to a Multicast router. The device includes the Multicast router port(s) numbers when it forwards the Multicast streams and IGMP registration messages. This is required so that the Multicast routers can, in turn, forward the Multicast streams and propagate the registration messages to other subnets. The reports would be re-injected over mrouter ports.

Multicast Router and Host Ports

IGMP snooping classifies each port (for example, EFPs, PWs, physical ports, or EFP bundles) as one of the following:

- Multicast router ports (mrouter ports)—These are ports to which a multicast-enabled router is connected. Mrouter ports are usually dynamically discovered, but may also be statically configured. Multicast traffic is always forwarded to all mrouter ports, except when an mrouter port is the ingress port.

- Host ports—Any port that is not an mrouter port is a host port.

IGMP snooping classifies each port (for example, EFPs, physical ports, or EFP bundles) as a host ports, that is, any port that is not an mrouter port is a host port.

Multicast Traffic Handling within a Bridge Domain with IGMP Snooping Enabled

The following tables describe traffic handling behaviors by IGMP snooping mrouter and host ports.

By default, IGMP snooping supports IGMPv2 and IGMPv3. The version of the IGMP querier discovered in the bridge domain determines the operational version of the snooping processes. If you change the default, configuring IGMP snooping to support a minimum version of IGMPv3, IGMP snooping ignores any IGMPv2 queriers.

Table 2: Multicast Traffic Handling for an IGMPv2 Querier

Traffic Type	Received on MRouter Ports	Received on Host Ports
IP multicast source traffic	Forwards to all mrouter ports and to host ports that indicate interest.	Forwards to all mrouter ports and to host ports that indicate interest.
IGMP general queries	Forwards to all ports.	—
IGMP group-specific queries	Forwards to all other mrouter ports.	—
IGMPv2 joins	Examines (snoops) the reports. • If report suppression is enabled, forwards first join for a new group or first join following a general query for an existing group. • If report suppression is disabled, forwards on all mrouter ports.	Examines (snoops) the reports. • If report suppression is enabled, forwards first join for a new group or first join following a general query for an existing group. • If report suppression is disabled, forwards on all mrouter ports.
IGMPv3 reports	Ignores	Ignores
IGMPv2 leaves	Invokes last member query processing.	Invokes last member query processing.

Table 3: Multicast Traffic Handling for an IGMPv3 Querier

Traffic Type	Received on MRouter Ports	Received on Host Ports
IP multicast source traffic	Forwards to all mrouter ports and to host ports that indicate interest.	Forwards to all mrouter ports and to host ports that indicate interest.
IGMP general queries	Forwards to all ports.	—
IGMP group-specific queries	If received on the querier port floods on all ports. Forwards to all other Mrouter ports.	—
IGMPv2 joins	Handles as IGMPv3 IS_EX{} reports.	Handles as IGMPv3 IS_EX{} reports.

Traffic Type	Received on MRouter Ports	Received on Host Ports
IGMPv3 reports	• If proxy reporting is enabled—For state changes or source-list changes, generates a state change report on all mrouter ports. • If proxy reporting is disabled—Forwards on all mrouter ports.	• If proxy reporting is enabled—For state changes or source-list changes, generates a state change report on all mrouter ports. • If proxy reporting is disabled—Forwards on all mrouter ports.
IGMPv2 leaves	Handles as IGMPv3 IS_IN{} reports.	Handles as IGMPv3 IS_IN{} reports.

IGMP Snooping Configuration Profiles

To enable IGMP snooping on a bridge domain, you must attach a profile to the bridge domain. The minimum configuration is an empty profile if BVI is configured. An empty profile enables the default configuration options and settings for IGMP snooping, as listed in the *Default IGMP Snooping Configuration Settings*.

Note You must configure the **system-ip-address** and **internal-querier** when the BVI is not configured, and no other queriers are present in the same domain.

Configuration Example:

```
Router(config)#igmp snooping profile igmpsn
Router(config-igmp-snooping-profile)#system-ip-address 192.0.2.1
Router(config-igmp-snooping-profile)#internal-querier
```

You can attach IGMP snooping profiles to bridge domains or to ports under a bridge domain. The following guidelines explain the relationships between profiles attached to ports and bridge domains:

- Any IGMP Snooping profile attached to a bridge domain, even an empty profile, enables IGMP snooping. To disable IGMP snooping, detach the profile from the bridge domain.
- An empty profile configures IGMP snooping on the bridge domain and all ports under the bridge using default configuration settings.
- A bridge domain can have only one IGMP snooping profile attached to it (at the bridge domain level) at any time.
- Port profiles are not in effect if the bridge domain does not have a profile attached to it.
- IGMP snooping must be enabled on the bridge domain for any port-specific configurations to be in effect.
- If a profile attached to a bridge domain contains port-specific configuration options, the values apply to all of the ports under the bridge, including all mrouter and host ports, unless another port-specific profile is attached to a port.
- When a profile is attached to a port, IGMP snooping reconfigures that port, disregarding any port configurations that may exist in the bridge-level profile.

Creating Profiles

To create a profile, use the **igmp snooping profile** command in global configuration mode.

Attaching and Detaching Profiles

To attach a profile to a bridge domain, use the **igmp snooping profile** command in l2vpn bridge group bridge domain configuration mode. To attach a profile to a port, use the **igmp snooping profile** command in the interface configuration mode under the bridge domain. To detach a profile, use the **no** form of the command in the appropriate configuration mode.

When you detach a profile from a bridge domain or a port, the profile still exists and is available for use at a later time. Detaching a profile has the following results:

- If you detach a profile from a bridge domain, IGMP snooping is deactivated in the bridge domain.
- If you detach a profile from a port, IGMP snooping configuration values for the port are instantiated from the bridge domain profile.

Changing Profiles

You cannot make changes to an active profile. An active profile is one that is currently attached.

- If the active profile is configured under the bridge, you must detach it from the bridge, and reattach it.
- If the active profile is configured under a specific bridge port, you must detach it from the bridge port, and reattach it.

Another way to do this is to create a new profile incorporating the desired changes and attach it to the bridges or ports, replacing the existing profile. This deactivates IGMP snooping and then reactivates it with parameters from the new profile.

Default IGMP Snooping Configuration Settings

Table 4: IGMP Snooping Default Configuration Values

Scope	Feature	Default Value
Bridge Domain	IGMP snooping	Disabled on a bridge domain until an enabling IGMP snooping profile is attached to the bridge domain.
	internal querier	By default Internal Querier is disabled. To enable Internal Querier, add it to the IGMP snooping profile. Internal Querier is not recommended, when BVI and IGMP snooping is configured under a bridge.
	last-member-query-count	2
	last-member-query-interval	1000 (milliseconds)
	minimum-version	2 (supporting IGMPv2 and IGMPv3)
	querier query-interval	60 (seconds) **Note** This is a nonstandard default value.
	report-suppression	Enabled (enables report suppression for IGMPv2 and proxy-reporting for IGMPv3)
	querier robustness-variable	2
	router alert check	Enabled
	tcn query solicit	Disabled
	tcn flood	Enabled
	ttl-check	Enabled
	unsolicited-report-timer	1000 (milliseconds)
Port	immediate-leave	Disabled
	mrouter	No static mrouters configured; dynamic discovery occurs by default.
	static group	None configured

IGMP Snooping Configuration at the Bridge Domain Level

IGMP Minimum Version

The **minimum-version** command determines which IGMP versions are supported by IGMP snooping in the bridge domain:

- When minimum-version is 2, IGMP snooping intercepts IGMPv2 and IGMPv3 messages. This is the default value.
- When minimum-version is 3, IGMP snooping intercepts only IGMPv3 messages and drops all IGMPv2 messages.

IGMPv1 is not supported. The scope for this command is the bridge domain. The command is ignored in a profile attached to a port.

Group Membership Interval, Robustness Variable, and Query Interval

The group membership interval (GMI) controls when IGMP snooping expires stale group membership states. The **show igmp snooping group** command shows groups with an expiry time of 0 until that stale state is cleaned up following the next query interval.

The GMI is calculated as:

GMI = (robustness-variable * query-interval) + maximum-response-time

where:

- maximum-response-time (MRT) is the amount of time during which receivers are required to report their membership state.
- robustness-variable is an integer used to influence the calculated GMI.
- query-interval is the amount of time between general queries.

Values for the components in the GMI are obtained as follows:

- MRT is advertised in the general query, for both IGMPv2 and IGMPv3.
- If the querier is running IGMPv2, IGMP snooping uses the IGMP-snooping-configured values for the robustness-variable and query-interval. These parameter values must match the configured values for the querier. In most cases, if you are interacting with other Cisco routers, you should not need to explicitly configure these values—the default values for IGMP snooping should match the default values of the querier. If they do not, use the **querier robustness-variable** and **querier query-interval** commands to configure matching values.
- IGMPv3 general queries convey values for robustness-variable and query-interval (QRV and QQI, respectively). IGMP snooping uses the values from the query, making the IGMP snooping GMI exactly match that of the querier.

How to Configure IGMP Snooping

The first two tasks are required to configure basic IGMP snooping configuration.

Creating an IGMP Snooping Profile

Procedure

Step 1 **configure**

Step 2 **igmp snooping profile** *profile-name*

Example:

```
RP/0/RP0/CPU0:router(config)# igmp snooping profile default-bd-profile
```

Enters IGMP snooping profile configuration mode and creates a named profile.

The default profile enables IGMP snooping. You can commit the new profile without any additional configurations, or you can include additional configuration options to the profile. You can also return to the profile later to add configurations, as described in other tasks in this module.

Step 3 Optionally, add commands to override default configuration values.

If you are creating a bridge domain profile, consider the following:

- An empty profile is appropriate for attaching to a bridge domain. An empty profile enables IGMP snooping with default configuration values.
- You can optionally add more commands to the profile to override default configuration values.
- If you include port-specific configurations in a bridge domain profile, the configurations apply to all ports under the bridge, unless another profile is attached to a port.

If you are creating a port-specific profile, consider the following:

- While an empty profile could be attached to a port, it would have no effect on the port configuration.
- When you attach a profile to a port, IGMP snooping reconfigures that port, overriding any inheritance of configuration values from the bridge-domain profile. You must repeat the commands in the port profile if you want to retain those configurations.

You can detach a profile, change it, and reattach it to add commands to a profile at a later time.

Step 4 **commit**

Where to Go Next

Attach the profile to bridge domains or ports to complete immediate-leave configuration. See one of the following sections:

Attaching a Profile and Activating IGMP Snooping on a Bridge Domain

To activate IGMP snooping on a bridge domain, attach an IGMP snooping profile to the bridge domain, as described in the following steps.

Procedure

Step 1 **configure**

Step 2 **l2vpn**

Example:

```
RP/0/RP0/CPU0:router(config)# l2vpn
```

Enters Layer 2 VPN configuration mode.

Step 3 **bridge group** *bridge-group-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn)# bridge group GRP1
```

Enters Layer 2 VPN VPLS bridge group configuration mode for the named bridge group.

Step 4 **bridge-domain** *bridge-domain-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg)# bridge-domain ISP1
```

Enters Layer 2 VPN VPLS bridge group bridge domain configuration mode for the named bridge domain.

Step 5 **igmp snooping profile** *profile-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# igmp snooping profile default-bd-profile
```

Attaches the named IGMP snooping profile to the bridge domain, enabling IGMP snooping on the bridge domain.

Step 6 **commit**

Step 7 **show igmp snooping bridge-domain detail**

Example:

```
RP/0/RP0/CPU0:router# show igmp snooping bridge-domain detail
```

(Optional) Verifies that IGMP snooping is enabled on a bridge domain and shows the IGMP snooping profile names attached to bridge domains and ports.

Step 8 **show l2vpn bridge-domain detail**

Example:

```
RP/0/RP0/CPU0:router# show l2vpn bridge-domain
```

(Optional) Verifies that IGMP snooping is implemented in the forwarding plane (Layer 2) on a bridge domain.

Detaching a Profile and Deactivating IGMP Snooping on a Bridge Domain

To deactivate IGMP snooping on a bridge domain, remove the profile from the bridge domain using the following steps.

Note A bridge domain can have only one profile attached to it at a time.

Procedure

Step 1 **configure**

Step 2 **l2vpn**

Example:

```
RP/0/RP0/CPU0:router(config)# l2vpn
```

Enters Layer 2 VPN configuration mode.

Step 3 **bridge group** *bridge-group-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn)# bridge group GRP1
```

Enters Layer 2 VPN VPLS bridge group configuration mode for the named bridge group.

Step 4 **bridge-domain** *bridge-domain-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg)# bridge-domain ISP1
```

Enters Layer 2 VPN VPLS bridge group bridge domain configuration mode for the named bridge domain.

Step 5 **no igmp snooping disable**

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# no igmp snooping disable
```

Detaches the IGMP snooping profile from the bridge domain, disabling IGMP snooping on that bridge domain.

Note
Only one profile can be attached to a bridge domain at a time. If a profile is attached, IGMP snooping is enabled. If a profile is not attached, IGMP snooping is disabled.

Step 6 **commit**

Step 7 **show igmp snooping bridge-domain detail**

Example:

```
RP/0/RP0/CPU0:router# show igmp snooping bridge-domain detail
```

(Optional) Verifies that IGMP snooping is disabled on a bridge domain.

Step 8 **show l2vpn bridge-domain detail**

Example:

```
RP/0/RP0/CPU0:router# show l2vpn bridge-domain
```

(Optional) Verifies that IGMP snooping is disabled in the forwarding plane (Layer 2) on a bridge domain.

Attaching and Detaching Profiles to Ports Under a Bridge

Before you begin

IGMP snooping must be enabled on the bridge domain for port-specific profiles to affect IGMP snooping behavior.

Procedure

Step 1 **configure**

Step 2 **l2vpn**

Example:

```
RP/0/RP0/CPU0:router(config)# l2vpn
```

Enters Layer 2 VPN configuration mode.

Step 3 **bridge group** *bridge-group-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn)# bridge group GRP1
```

Enters Layer 2 VPN bridge group configuration mode for the named bridge group.

Step 4 **bridge-domain** *bridge-domain-name*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg)# bridge-domain ISP1
```

Enters Layer 2 VPN bridge group bridge domain configuration mode for the named bridge domain.

Step 5 **interface** *interface-type interface-number*

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)# interface gig 1/1/1/1
```

Enters Layer 2 VPN VPLS bridge group bridge domain interface configuration mode for the named interface or PW.

Step 6 Do one of the following:

- **igmp snooping profile** *profile-name*
- **no igmp snooping**

Example:

```
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-if)# igmp snooping profile mrouter-port-profile
```

Attaches the named IGMP snooping profile to the port.

Note
A profile on a port has no effect unless there is also a profile attached to the bridge.

The **no** form of the command detaches a profile from the port. Only one profile can be attached to a port.

Step 7 **commit**

Step 8 **show igmp snooping bridge-domain detail**

Example:

```
RP/0/RP0/CPU0:router# show igmp snooping bridge-domain detail
```

(Optional) Verifies that IGMP snooping is enabled on a bridge domain and shows the IGMP snooping profile names attached to bridge domains and ports.

Step 9 **show l2vpn bridge-domain detail**

Example:

```
RP/0/RP0/CPU0:router# show l2vpn bridge-domain
```

(Optional) Verifies that IGMP snooping is implemented in the forwarding plane (Layer 2) on a bridge domain.

Verifying Multicast Forwarding

Procedure

Step 1 **configure**

Step 2 **show l2vpn forwarding bridge-domain** [*bridge-group-name*:*bridge-domain-name*] **mroute ipv4** [**group** *group_IPaddress*] [**hardware** {**ingress** | **egress**}] [**detail**]**location** *node-id*

Example:

```
 RP/0/RP0/CPU0:router#show l2vpn forwarding bridge-domain bg1:bd1 mroute ipv4 group 234.192.4.1
hardware ingress detail location 0/1/cPU0
```

Displays multicast routes as they are converted into the forwarding plane forwarding tables. Use optional arguments to limit the display to specific bridge groups or bridge domains.

If these routes are not as expected, check the control plane configuration and correct the corresponding IGMP snooping profiles.

Step 3 **show l2vpn forwarding bridge-domain** [*bridge-group-name*:*bridge-domain-name*] **mroute ipv4 summary location** *node-id*

Example:

```
RP/0/RP0/CPU0:router# show l2vpn forwarding bridge-domain bg1:bd1 mroute ipv4 summary location 0/3/CPU0
```

Displays summary-level information about multicast routes as stored in the forwarding plane forwarding tables. Use optional arguments to limit the display to specific bridge domains.

Configuration Examples for IGMP Snooping

The following examples show how to enable IGMP snooping on Layer 2 VPLS bridge domains on Cisco 8000 Series Routers:

Configuring IGMP Snooping on Physical Interfaces Under a Bridge: Example

1. Create two profiles.

```
igmp snooping profile bridge_profile
!
igmp snooping profile port_profile
   mrouter
!
```

2. Configure two physical interfaces for L2 transport.

```
interface GigabitEthernet0/8/0/38
   negotiation auto
   l2transport
   no shut
   !
!
interface GigabitEthernet0/8/0/39
   negotiation auto
   l2transport
   no shut
   !
!
```

3. Add interfaces to the bridge domain. Attach bridge_profile to the bridge domain and port_profile to one of the Ethernet interfaces. The second Ethernet interface inherits IGMP snooping configuration attributes from the bridge domain profile.

```
l2vpn
  bridge group bg1
     bridge-domain bd1
     igmp snooping profile bridge_profile
     interface GigabitEthernet0/8/0/38
       igmp snooping profile port_profile
     interface GigabitEthernet0/8/0/39

!
  !
!
```

4. Verify the configured bridge ports.

```
show igmp snooping port
```

Configuring IGMP Snooping on VLAN Interfaces Under a Bridge: Example

1. Configure two profiles.

```
multicast-source ipv4
igmp snooping profile bridge_profile

igmp snooping profile port_profile
   mrouter
!
```

2. Configure VLAN interfaces for L2 transport.

```
interface GigabitEthernet0/8/0/8
   negotiation auto
   no shut
   !
!
interface GigabitEthernet0/8/0/8.1 l2transport
   encapsulation dot1q 1001
   rewrite ingress tag pop 1 symmetric
   !
!
interface GigabitEthernet0/8/0/8.2 l2transport
   encapsulation dot1q 1002
   rewrite ingress tag pop 1 symmetric
   !
!
```

3. Attach a profile and add interfaces to the bridge domain. Attach a profile to one of the interfaces. The other interface inherits IGMP snooping configuration attributes from the bridge domain profile.

```
l2vpn
  bridge group bg1
     bridge-domain bd1
     multicast-source ipv4
     igmp snooping profile bridge_profile
     interface GigabitEthernet0/8/0/8.1
       igmp snooping profile port_profile
     interface GigabitEthernet0/8/0/8.2

     !
  !
!
```

4. Verify the configured bridge ports.

```
show igmp snooping port
```

Configuring IGMP Snooping on Ethernet Bundles Under a Bridge: Example

1. This example assumes that the front-ends of the bundles are preconfigured. For example, a bundle configuration might consist of three switch interfaces, as follows:

```
    interface Port-channel1
    !
interface GigabitEthernet0/0/0/0
    !
interface GigabitEthernet0/0/0/1
!
    interface GigabitEthernet0/0/0/2
       channel-group 1 mode on
    !
    interface GigabitEthernet0/0/0/3
       channel-group 1 mode on
    !
```

2. Configure two IGMP snooping profiles.

```
multicast-source ipv4
       igmp snooping profile bridge_profile
       !
       multicast-source ipv4
       igmp snooping profile port_profile
          mrouter
       !
```

3. Configure interfaces as bundle member links.

```
       interface GigabitEthernet0/0/0/0
         bundle id 1 mode on
         negotiation auto
       !
       interface GigabitEthernet0/0/0/1
         bundle id 1 mode on
         negotiation auto
       !
```

```
interface GigabitEthernet0/0/0/2
  bundle id 2 mode on
  negotiation auto
!
interface GigabitEthernet0/0/0/3
  bundle id 2 mode on
  negotiation auto
!
```

4. Configure the bundle interfaces for L2 transport.

```
interface Bundle-Ether 1
          l2transport
          !
       !
       interface Bundle-Ether 2
          l2transport
          !
       !
```

5. Add the interfaces to the bridge domain and attach IGMP snooping profiles.

```
l2vpn
  bridge group bg1
     bridge-domain bd1
     igmp snooping profile bridge_profile
     interface bundle-Ether 1
       igmp snooping profile port_profile
     interface bundle-Ether 2

     !
  !
!
```

6. Verify the configured bridge ports.

```
show igmp snooping port
```

Configuring Multicast over Integrated Routing Bridging Active/Active Multihome

Configurations performed on peer 1:

1. Layer 2 Base Configuration

```
hostname peer1
!
interface Bundle-Ether2
!
interface Bundle-Ether2.2 l2transport
 encapsulation dot1q 2
 rewrite ingress tag pop 1 symmetric
!
interface TenGigE0/0/0/0
```

```
 bundle id 2 mode on
 no shut
!
```

2. IGMPv2 Snoop Configurations

```
hostname peer1
!
router igmp

  version 2
 !
!
l2vpn
 bridge group VLAN2
  bridge-domain VLAN2
   multicast-source ipv4
   igmp snooping profile 1
   interface Bundle-Ether2.2
   !

   evi 2
   !
  !
 !
multicast-source ipv4
igmp snooping profile 1
!
```

Configurations Performed on Peer 2:

1. Layer 2 Base Configuration

```
hostname peer2
!
interface Bundle-Ether2
!
interface Bundle-Ether2.2 l2transport
 encapsulation dot1q 2
 rewrite ingress tag pop 1 symmetric
!
interface TenGigE0/0/0/0
 bundle id 2 mode on
 no shut
!
```

2. IGMPv2 Snoop Configurations

```
hostname peer2
!
router igmp

  version 2
 !
!
l2vpn
 bridge group VLAN2
  bridge-domain VLAN2
   multicast-source ipv4
   igmp snooping profile 1
   interface Bundle-Ether2.2
   !

   evi 2
   !
  !
 !
multicast-source ipv4
igmp snooping profile 1
!
```

Verifying IGMP Snooping

In this example, the receiver sends an IGMPv2 join for the group 239.0.0.2. On Peer2, this group has a D Flag, that means the actual IGMP joined peer2, but not peer1. On Peer1, this group has a B flag, that means this group is learnt from BGP.

```
RP/0/RP0/CPU0:peer1#show igmp snooping group
Fri Aug 31 22:27:46.363 UTC

Key: GM=Group Filter Mode, PM=Port Filter Mode
Flags Key: S=Static, D=Dynamic, B=BGP Learnt, E=Explicit Tracking, R=Replicated

                            Bridge Domain VLAN10:VLAN10

Group           Ver GM Source          PM Port                        Exp   Flgs
-----           --- -- ------          -- ----                        ---   ----
239.0.0.2       V2  -  *               -  BE2.2                       never B


RP/0/RP0/CPU0:peer2#show igmp snooping group
Fri Aug 31 22:27:49.686 UTC

Key: GM=Group Filter Mode, PM=Port Filter Mode
Flags Key: S=Static, D=Dynamic, B=BGP Learnt, E=Explicit Tracking, R=Replicated

                            Bridge Domain VLAN10:VLAN10

Group           Ver GM Source          PM Port                        Exp   Flgs
-----           --- -- ------          -- ----                        ---   ----
239.0.0.2       V2  -  *               -  BE2.2                       74    D
```

Verifying Dual DR PIM Uplink

In this example, when the source 126.0.0.100 sends traffic to group 239.0.0.2, you see both Peer1 and Peer2 are sending PIM join upstream. The incoming interface for (*,G) and (S,G) should be the interface toward the RP and source respectively. For both Peer1 and Peer2, the outgoing interface should be the BVI interface facing the receiver.

```
RP/0/RP0/CPU0:peer1#show mrib route
:
:

(*,239.0.0.2) RPF nbr: 30.0.0.4 Flags: C RPF
  Up: 00:13:41
  Incoming Interface List
    HundredGigE0/0/0/1 Flags: A NS, Up: 00:13:41
  Outgoing Interface List
    BVI2 Flags: F NS LI, Up: 00:13:41

(126.0.0.100,239.0.0.2) RPF nbr: 30.0.0.4 Flags: RPF
  Up: 00:03:34
  Incoming Interface List
    HundredGigE0/0/0/1 Flags: A, Up: 00:03:34
  Outgoing Interface List
    BVI2 Flags: F NS, Up: 00:03:34
:
:

RP/0/RP0/CPU0:peer2#show mrib route
:
:

(*,239.0.0.2) RPF nbr: 50.0.0.4 Flags: C RPF
  Up: 00:13:33
  Incoming Interface List
    HundredGigE0/0/0/2 Flags: A NS, Up: 00:13:33
  Outgoing Interface List
    BVI2 Flags: F NS LI, Up: 00:13:33

(126.0.0.100,239.0.0.2) RPF nbr: 50.0.0.4 Flags: RPF
  Up: 00:03:24
  Incoming Interface List
    HundredGigE0/0/0/2 Flags: A, Up: 00:03:24
  Outgoing Interface List
    BVI2 Flags: F NS, Up: 00:03:24
:
:
```

MLD Snooping

Multicast Listener Discovery (MLD) snooping is a technique that uses the MLD protocol to optimize the delivery of multicast traffic.

Table 5: Feature History Table

Feature Name	Release Information	Feature Description
MLD snooping	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100])(select variants only*) *This feature is supported on: • 8711-48Z-M • 8011-12G12X4Y-A • 8011-12G12X4Y-D

Feature Name	Release Information	Feature Description
MLD snooping	Release 25.1.1	Introduced in this release on: Fixed Systems (8010 [ASIC: A100]) This feature is enhanced to support: • MLD snooping on BVI, and • MLD versions MLDv1 and MLDv2. This feature is now supported on 8011-4G24Y4H-I routers.
MLD snooping	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100]), (8700 [ASIC: P100, K100])(select variants only*). Multicast Listener Discovery (MLD) snooping is a technique that uses the MLD protocol to optimize the delivery of multicast traffic. When you enable MLD snooping on the router, it sends multicast data only to network segments with devices that have expressed interest in receiving it. By sending multicast data only to interested devices, the router minimizes unnecessary traffic and conserves bandwidth on the network. *This feature is now supported on: • 8212-48FH-M • 8712-MOD-M • 8711-32FH-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Multicast Listener Discovery (MLD) snooping provides a way to constrain multicast traffic at Layer 2. By snooping the MLD membership reports sent by hosts in the bridge domain, the MLD snooping application can set up Layer 2 multicast forwarding tables to deliver traffic only to ports with at least one interested member, significantly reducing the volume of multicast traffic.

MLD snooping uses the information in MLD membership report messages to build corresponding information in the forwarding tables to restrict IPv6 multicast traffic at Layer 2. The forwarding table entries are in the form <Route, OIF List>, where:

- Route is a <*, G> route or <S, G> route.
- OIF List comprises all bridge ports that have sent MLD membership reports for the specified route plus all multicast router (mrouter) ports in the bridge domain.

For more information regarding MLD snooping, refer the *Multicast Configuration Guide for Cisco 8000 Series Routers*.

Prerequisites for MLD Snooping

- The network must be configured with a layer2 VPN.
- You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Supported Features and Restrictions for MLD Snooping

- Starting with Cisco IOS XR Release 25.1.1, MLD snooping on BVI is supported.
- Receiver behind L2 ACs in the same L2 bridge domain is supported.
- Source behind L2 ACs in the same L2 bridge domain is supported.
- Starting with Cisco IOS XR Release 25.1.1, both MLDv1 and MLDv2 are supported over BVI.
- EVPN MLD sync is not supported.
- VPLS is not supported.
- The **router-alert-check disable** configuration command is not supported.
- EVPN configuration must have the **control-word-disable** configuration.
- PIM control packets (join and hello) processing is not supported when snooping is enabled, so a multicast router selection based on PIM packets won't occur.
- Explicit host tracking.
- Multicast Admission Control.
- Security filtering.
- Report rate limiting.
- Multicast router discovery.
- Starting with Cisco IOS XR Release 25.1.1, IPv6 multicast is supported for a multicast source that is behind the BVI interface.
- In an EVPN dual-home AA scenario:
  - If the multicast source and receiver are in the same bridge domain (BD), the receiver might receive permanent traffic duplication.
  - In an EVPN dual-home receiver AA scenario, transient traffic duplication is expected when the DH node role changes from DF to nDF and vice versa.

- Source=ESI1=BE-X.A, Receiver=ESI1=BE-X.B under the same BD is not supported (where X.A and X.B represent two AC ports for the bundle interface BE).

Advantages of MLD Snooping

- In its basic form, it reduces bandwidth consumption by reducing multicast traffic that would otherwise flood an entire VPLS bridge domain.
- With the use of some optional configurations, it provides security between bridge domains by filtering the MLD reports received from hosts on one bridge port and preventing leakage towards the hosts on other bridge ports.

High Availability (HA) features for MLD

MLD supports the following HA features:

- Process restarts
- Stateful Switch-Over (SSO)

Bridge Domain Support for MLD

MLD snooping operates at the bridge domain level. When MLD snooping is enabled on a bridge domain, the snooping functionality applies to all ports under the bridge domain, including:

- Physical ports under the bridge domain.
- Ethernet flow points (EFPs)—An EFP can be a VLAN, VLAN range, list of VLANs, or an entire interface port.
- Ethernet bundles—Ethernet bundles include IEEE 802.3ad link bundles and Cisco EtherChannel bundles. From the perspective of the MLD snooping application, an Ethernet bundle is just another EFP. The forwarding application in the Cisco 8000 Series Routers randomly nominates a single port from the bundle to carry the multicast traffic.

Note The **efp-visibility** configuration is required when a bridge has attachment circuits as VLAN sub-interfaces from the same bundle-ether or physical interface.

Multicast Router and Host Ports

MLD snooping classifies each port as one of the following:

- Multicast router ports (mrouter ports)—These are ports to which a multicast-enabled router is connected. Mrouter ports are usually dynamically discovered, but may also be statically configured. Multicast traffic is always forwarded to all mrouter ports, except when an mrouter port is the ingress port.
- Host ports—Any port that is not an mrouter port is a host port.

Multicast Router Discovery for MLD

MLD snooping discovers mrouter ports dynamically. You can also explicitly configure a port as an emrouter port.

- Discovery- MLD snooping identifies upstream mrouter ports in the bridge domain by snooping mld query messages and Protocol Independent Multicast Version 2 (PIMv2) hello messages. Snooping PIMv2 hello messages identifies mld nonqueriers in the bridge domain.
- Static configuration—You can statically configure a port as an mrouter port with the **mrouter** command in a profile attached to the port. Static configuration can help in situations when incompatibilities with non-Cisco equipment prevent dynamic discovery.

Multicast Traffic Handling for MLD

The following tables describe the traffic handling behavior by MLD mrouters and host ports.

Table 6: Multicast Traffic Handling for a MLDv1 Querier

Traffic Type	Received on MRouter Ports	Received on Host Ports
IP multicast source traffic	Forwards to all mrouter ports and to host ports that indicate interest.	Forwards to all mrouter ports and to host ports that indicate interest.
MLD general queries	Forwards to all ports.	—
MLD group-specific queries	Forwards to all other mrouter ports.	Dropped
MLDv1 joins	Examines (snoops) the reports. • If report suppression is enabled, forwards first join for a new group or first join following a general query for an existing group. • If report suppression is disabled, forwards on all mrouter ports.	Examines (snoops) the reports. • If report suppression is enabled, forwards first join for a new group or first join following a general query for an existing group. • If report suppression is disabled, forwards on all mrouter ports.
MLDv2 reports	Ignores	Ignores
MLDv1 leaves	Invokes last member query processing.	Invokes last member query processing.

Table 7: Multicast Traffic Handling for a MLDv2 Querier

Traffic Type	Received on MRouter Ports	Received on Host Ports
IP multicast source traffic	Forwards to all mrouter ports and to host ports that indicate interest.	Forwards to all mrouter ports and to host ports that indicate interest.
MLD general queries	Forwards to all ports.	—

Traffic Type	Received on MRouter Ports	Received on Host Ports
MLD group-specific queries	If received on the querier port floods on all ports.	—
MLDv1 joins	Handles as MLDv2 IS_EX{} reports.	Handles as MLDv2 IS_EX{} reports.
MLDv2 reports	• If proxy reporting is enabled—For state changes or source-list changes, generates a state change report on all mrouter ports. • If proxy reporting is disabled—Forwards on all mrouter ports.	• If proxy reporting is enabled—For state changes or source-list changes, generates a state change report on all mrouter ports. • If proxy reporting is disabled—Forwards on all mrouter ports.
MLDv1 leaves	Handles as MLDv2 IS_IN{} reports.	Handles as MLDv2 IS_IN{} reports.

Multicast Listener Discovery over BVI

Multicast IPv6 packets received from core, which has BVI as forwarding interface, is forwarded to access over snooped L2 AC or interface.

Note

- As per MLDv2 RFC recommendation the MLDv2 reports should carry the Hop-by-Hop options header for the reports to get punted up.
- MLDv2 is supported over BVI only when BVI is configured as a forwarding interface.

MLD and BVI Overview

Routers use the Internet Group Management Protocol (IGMP) (IPv4) and Multicast Listener Discovery (MLD) (IPv6) to learn whether members of a group are present on their directly attached subnets. Hosts join multicast groups by sending IGMP or MLD report messages.

MLDv1 and MLDv2 are supported on Cisco 8010 Series Routers (*select varaints). However, MLDv2 is enabled when you configure MLD by default.

MLDv2 shares feature parity with IGMPv3 with respect to all supported interface types with the exception of PPoE and subinterfaces. MLDv2 enables a node to report interest in listening to packets only from specific multicast source addresses.

A BVI interface is a routed interface representing a set of interfaces (bridged) in the same L2 broadcast domain. MLD join messages coming in or out of this broadcast domain passes through the BVI interface.

Multicast Traffic Over Layer 2 IPv6 Network

Table 8: Feature History Table

Feature Name	Release Information	Feature Description
Multicast Traffic over Layer 2 IPv6 Network	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100])(select variants only*) *This feature is supported on: • 8711-48Z-M • 8011-12G12X4Y-A • 8011-12G12X4Y-D
Multicast Traffic over Layer 2 IPv6 Network	Release 25.1.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100], 8010 [ASIC: A100]) This feature allows you to forward the IPv6 multicast packets only to the interested MLD-snooped Access Controllers (AC), whereas in the default case, the bridge floods the IPv6 multicast packets to all AC. Routers use Multicast Listener Discovery (MLD) protocol to discover the devices in a network and create route entries in an IPv6 multicast network. This feature is now supported on: • 8712-MOD-M • 8011-4G24Y4H-I

The Multicast Traffic over Layer 2 IPv6 Network (L2MC IPv6) is an optimized forwarding technique, and it helps in saving the bandwidth. By default, the bridge floods IPv6 multicast packets to all AC, whereas the L2MC IPv6 feature allows you to forward the IPv6 multicast packets only to the interested MLD-snooped AC.

When IPv6 multicast packets are received over Layer 2 AC and interfaces, the lookup gets done for Virtual Switch Interfaces (VSI), Groups (G), and Services (S) or for VSI and G. The VSI details show the VLAN or VXLAN segment to which the packet belongs, while the G and S identify the multicast groups and services to which the packet should be forwarded. Based on this lookup, the traffic is forwarded to the interested receivers connected to the Layer 2 AC.

The MLD control packets received over Layer 2 AC are snooped and punted to create the route entries. This route entries are needed to avail the following supports:

- Layer 2 Multicast IPv6 support.
- EVPN sync support for IPv4 routes.

Limitations and Restrictions

- This feature doesn’t support MLD sync.
- With L2MC IPv6 support, the existing L2MC IPv4 scale reduces proportionally.

Configuration Example

The L2MC IPv6 feature is not enabled by default. Following is a configuration example that shows how to enable the feature.

```
router(config)# l2vpn
 router(config-l2vpn)# bridge group 1
 router(config-l2vpn-bg)#bridge-domain 1
 router(config-l2vpn-bg-bd)#multicast-source ipv6
 router(config-l2vpn-bg-bd)#efp-visibility
 router(config-l2vpn-bg-bd)#mld snooping profile prof1
 router(config-l2vpn-bg-bd)#igmp snooping profile prof1
 router(config-l2vpn-bg-bd)#interface TenGigE0/0/0/0
 router(config-l2vpn-bg-bd-ac)#exit
 router(config-l2vpn-bg-bd)#interface TenGigE0/0/0/4.1
 router(config-l2vpn-bg-bd-ac)#exit
 router(config-l2vpn-bg-bd)#interface TenGigE0/0/0/4.2
 router(config-l2vpn-bg-bd-ac)#exit
 router(config-l2vpn-bg-bd)#routed interface BVI1
 router(config-l2vpn-bg-bd-bvi)#exit
 !
 !


 router(config-l2vpn-bg-bd)#mld snooping profile prof1
 router(config-l2vpn-bg-bd)#internal-querier
!
 router(config-l2vpn-bg-bd)#igmp snooping profile prof1
 router(config-l2vpn-bg-bd)#system-ip-address 1.2.3.4
 router(config-l2vpn-bg-bd)#internal-querier
```

Note With BVI configurations, there is no need to have internal queries address configured MLD snooping profile. It implies that you can make BVI as querier under BVI configuration.

Verification

The following command shows the information about group membership in the Layer 2 Forwarding tables.

```
router# show mld snooping group

Flags Key: S=Static, D=Dynamic, E=Explicit Tracking

                 Bridge Domain bg1:bd1
```

```
Group           Ver GM  Source           PM  Port                     Exp Flg

Ff12:1:1::1     V2  Exc  -               -   GigabitEthernet0/1/1/0   122  DE
Ff12:1:1::1     V2  Exc  2002:1::1       Inc GigabitEthernet0/1/1/1     5  DE
Ff12:1:1::1     V2  Exc  2002:1::1       Inc GigabitEthernet0/1/1/2 never   S
Ff12:1:1::1     V2  Exc  2002:1::1       Exc GigabitEthernet0/1/1/3     -  DE
Ff12:1:1::1     V2  Exc  2002:1::2       Inc GigabitEthernet0/1/1/0   202  DE
Ff12:1:1::1     V2  Exc  2002:1::2       Exc GigabitEthernet0/1/1/1     -  DE
Ff12:1:1::2     V2  Exc  2002:1::1       Inc GigabitEthernet0/1/1/0   145  DE
Ff12:1:1::2     V2  Exc  2002:1::1       Inc GigabitEthernet0/1/1/1     0  DE
Ff12:1:1::2     V2  Exc  2002:1::1       Exc GigabitEthernet0/1/1/2    11  DE


                 Bridge Domain bg1:bd4

Group           Ver GM  Source           PM  Port                     Exp Flg

Ff24:1:1::2     V1  Exc  -               -   GigabitEthernet0/1/1/0   122  DE
Ff28:1:1::1     V1  -    -               -   GigabitEthernet0/1/1/1    33  DE
Ff29:1:2::3     V1  Exc  -               -   GigabitEthernet0/1/2/0   122  DE
Ff22:1:2::3     V2  Exc  2000:1:1::2     Exc GigabitEthernet0/1/2/1     5  DE
```

The following command summarizes the number of bridge domains, mrouter ports, host ports, groups, and sources configured on the router.

```
router#show mld snooping summary
  Bridge Domains:                                      1
  MLD  Snooping Bridge Domains:                        1
  Ports:                                               3
  MLD  Snooping Ports:                                 3
  Mrouters:                                            0
  STP Forwarding Ports:                                0
  ICCP Group Ports:                                    0
  MLD  Groups:                                         0
    Member Ports:                                      0
  MLD  Source Groups:                                  0
    Static/Include/Exclude:                        0/0/0
    Member Ports (Include/Exclude):                  0/0
```

IPv6 Multicast Listener Discovery Snooping over BVI

Multicast Listener Discovery (MLD) snooping provides a way to constrain multicast traffic at L2. By snooping the MLD membership reports sent by hosts in the bridge domain, the MLD snooping application can set up L2 multicast forwarding tables. This table is later used to deliver traffic only to ports with at least one interested member, significantly reducing the volume of multicast traffic.

MLDv2 support over BVI enables implementing IPv6 multicast routing over a L2 segment of the network that is using an IPv6 VLAN. The multicast routes are bridged via BVI interface from L3 segment to L2 segment of the network.

MLDv2 snooping over BVI enables forwarding MLDv2 membership reports received over the L2 domain to MLD snooping instead of MLD.

Restrictions

- You cannot configure `ttl-check` and disable `router-alert-check` on the router for mld messages.
- Static mrouters are not supported for MLD snooping.
- Querier is supported for MLDV2, but it is not supported on MLDV1.

Configuring Internal Querier for MLD Snooping

This configuration enables a multicast router acting as a MLD querier to send out group-and-source-specific query:

```
router# config
RP0/0/RP0/CPU0:router(config)# mld snooping profile grp1
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# system-ip-address fe80::1 link-local
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# internal-querier
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# commit
```

Verification

Use the **show mld snooping profile detail** command to verify the MLD snooping configuration:

```
router# show mld snooping profile detail
Thu Nov 22 13:58:18.844 UTC
MLD  Snoop Profile grp1:
  System IP Address:                    fe80::1
  Bridge Domain References:             2
  Port References:                      12

MLD  Snoop Profile grp10:
  System IP Address:                    fe80::5610
  Bridge Domain References:             0
  Port References:                      0
```

Creating a MLD Snooping Profile

Configuration

```
/* Enter the global configuration mode */
RP/0/RP0/CPU0:router # configure
/* Enters MLD snooping profile configuration mode and creates a named profile. */
RP/0/RP0/CPU0:router(config)# mld snooping profile default-bd-profile
RP/0/RP0/CPU0:router # commit
```

The default profile enables MLD snooping. You can commit the new profile without any additional configurations, or you can include additional configuration options to the profile. You can also return to the profile later to add configurations, as described in other tasks in this module.

If you are creating a bridge domain profile, consider the following:

- An empty profile is appropriate for attaching to a bridge domain. An empty profile enables MLD snooping with default configuration values.
- You can optionally add more commands to the profile to override default configuration values.
- If you include port-specific configurations in a bridge domain profile, the configurations apply to all ports under the bridge, unless another profile is attached to a port.

If you are creating a port-specific profile, consider the following:

- While an empty profile could be attached to a port, it would have no effect on the port configuration.

- When you attach a profile to a port, MLD snooping reconfigures that port, overriding any inheritance of configuration values from the bridge-domain profile. You must repeat the commands in the port profile if you want to retain those configurations.

You can detach a profile, change it, and reattach it to add commands to a profile at a later time.

Running Configuration

```
RP/0/RP0/CPU0:router(config)# show running-config
configure
   mld snooping profile default-bd-profile
!
```

Verification

Verify that the MLD snooping profile is created:

```
RP/0/RP0/CPU0:router#show mld snooping profile

Profile                                   Bridge Domain       Port
-------                                   -------------       ----
default-bd-profile                                    0          0
grp1                                                  1          2
grp10                                                 1          2
```

Deactivating MLD Snooping on a Bridge Domain

To deactivate MLD snooping from a bridge domain, remove the profile from the bridge domain:

Note A bridge domain can have only one profile attached to it at a time.

Configuration

```
/* Enter the global configuration mode followed by the bridge group and the bridge domain
mode */
RP0/0/RP0/CPU0:router# configuration
RP0/0/RP0/CPU0:router(config)# l2vpn
RP0/0/RP0/CPU0:router(config-l2vpn)# bridge group GRP1
RP0/0/RP0/CPU0:router(config-l2vpn-bg)# bridge domain ISP1

/* Detache the MLD snooping profile from the bridge domain. This disables MLD snooping on
that bridge domain */
/* Note: Only one profile can be attached to a bridge domain at a time. If a profile is
attached, MLD snooping is enabled.
If a profile is not attached, MLD snooping is disabled. */
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd)# no mld snooping profile
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd)# commit
```

Running Configuration

```
RP0/0/RP0/CPU0:router# show running-config
configuration
l2vpn
 bridge-group GRP1
  bridge-domain ISP1
```

```
   no mld snooping profile
!
```

Configuring Static Mrouter Ports (MLD)

Prerequisite

MLD snooping must be enabled on the bridge domain for port-specific profiles to affect MLD snooping behavior.

Note Static mrouter port configuration is a port-level option and should be added to profiles intended for ports. It is not recommended to add mrouter port configuration to a profile intended for bridge domains.

Configuration

```
/* Enter the global configuration mode */
RP0/0/RP0/CPU0:router# configuration

/* Enter the MLD snooping profile configuration mode and create a new profile or accesses
an existing profile.*/
RP0/0/RP0/CPU0:router(config)# mld snooping profile mrouter-port-profile
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# mrouter
/* Configures a static mrouter on a port. */

RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# commit
```

Running Configuration

```
RP0/0/RP0/CPU0:router# show running-config
configuration
  mld snooping profile mrouter-port-profile
    mrouter
!
```

Verification

The below show command output confirms that the mrouter configuration is enabled:

```
RP0/0/RP0/CPU0:router# show mld snooping profile mrouter-port-profile

MLD  Snoop Profile mrouter-port-profile:

  Static Mrouter:                       Enabled

  Bridge Domain References:             0
  Port References:                      0
```

Configuring Immediate-leave for MLD

To add the MLD snooping immediate-leave option to an MLD snooping profile:

Configuration

```
/* Enter the global configuration mode. */
RP0/0/RP0/CPU0:router# configuration

/* Enter MLD snooping profile configuration mode and create a new profile or accesses an
existing profile. */
RP0/0/RP0/CPU0:router(config)# mld snooping profile host-port-profile
/* Enable the immediate-leave option */
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# immediate-leave
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd)# commit
```

If you add the **immediate-leave** option:

- to a profile attached to a bridge domain, it applies to all ports under the bridge.
- to a profile attached to a port, it applies to the port.

Running Configuration

```
RP0/0/RP0/CPU0:router# show running-config
configuration
 mld snooping profile host-port-profile
 immediate-leave
!
```

Verification

Verify that the immediate leave config in the named profile is enabled:

```
RP0/0/RP0/CPU0:router# show mld snooping profile host-port-profile detail

MLD  Snoop Profile host-port-profile:

  Immediate Leave:                     Enabled
  Router Guard:                        Enabled

  Bridge Domain References:            0
  Port References:                     0
```

Configuring Internal Querier for MLD

Prerequisite

MLD snooping must be enabled on the bridge domain for this procedure to take effect.

Configuration

```
/* Enter the global configuration mode. */
RP0/0/RP0/CPU0:router# configuration

/* Enter MLD snooping profile configuration mode and create a new profile or accesses an
existing profile. */
RP0/0/RP0/CPU0:router(config)# mld snooping profile internal-querier-profile

/* Configure an IP address for internal querier use. The default system-ip-address value
(0.0.0.0) is not valid for the internal querier.
You must explicitly configure an IP address. Enter a valid link-local IPv6 address. */
```

```
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# system-ip-address fe80::98 link-local

/* Enable an internal querier with default values for all options.*/
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# internal-querier
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# commit
```

Running Configuration

```
RP0/0/RP0/CPU0:router# show running-config
configuration
 mld snooping profile internal-querier-profile
 system-ip-address fe80::98 link-local
 internal-querier
!
```

Note Internal Querier is not recommended, when BVI and MLD snooping is configured under a bridge.

Verification

Verify that the internal querier config is enabled:

```
RP0/0/RP0/CPU0:router# show mld snooping profile internal-querier-profile detail

MLD  Snoop Profile internal-querier-profile:

  System IP Address:                    fe80::98

  Internal Querier Support:             Enabled

  Bridge Domain References:             0
  Port References:                      0
```

Configuring Static Groups for MLD

To add one or more static groups or MLDv2 source groups to an MLD snooping profile, follow these steps:

Prerequisite

MLD snooping must be enabled on the bridge domain for port-specific profiles to affect MLD snooping behavior.

Configuration

```
/* Enter the global configuration mode. */
RP0/0/RP0/CPU0:router# configuration

/* Enter MLD snooping profile configuration mode and create a new profile or accesses an
existing profile. */
RP0/0/RP0/CPU0:router(config)# mld snooping profile host-port-profile

/* Configure a static group. */
/* Note: Repeat this step to add additional static groups. */
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# static group 239.1.1.1 source 198.168.1.1
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# commit
```

If you add the **static group** option:

- to a profile attached to a bridge domain, it applies to all ports under the bridge.
- to a profile attached to a port, it applies to the port.

Running Configuration

```
RP0/0/RP0/CPU0:router# show running-config
configuration
 mld snooping profile host-port-profile
 static group 239.1.1.1 source 198.168.1.1
!
```

Verification

```
RP0/0/RP0/CPU0:router# show mld snooping bridge-domain f1:100 detail

Bridge Domain        Profile                          Act  Ver  #Ports  #Mrtrs  #Grps
#SGs
-------------        -------                          ---  ---  ------  ------  -----
----
f1:100              grp1                               Y  v2       3       1   1000   1002


  Profile Configured Attributes:
    System IP Address:                    fe80::99
    Minimum Version:                      1
    Report Suppression:                   Enabled
    Unsolicited Report Interval:          1000 (milliseconds)
    TCN Query Solicit:                    Disabled
    TCN Membership Sync:                  Disabled
    TCN Flood:                            Enabled
    TCN Flood Query Count:                2
    Router Alert Check:                   Disabled
    TTL Check:                            Enabled
    nV Mcast Offload:                     Disabled
    Internal Querier Support:             Disabled
    Querier Query Interval:               125 (seconds)
    Querier LMQ Interval:                 1000 (milliseconds)
    Querier LMQ Count:                    2
    Querier Robustness:                   2
    Startup Query Interval:               31 seconds
    Startup Query Count:                  2
    Startup Query Max Response Time:      10.0 seconds
    Mrouter Forwarding:                   Enabled
    P2MP Capability:                      Disabled
    Default IGMP Snooping profile:        Disabled
    IP Address:                           fe80::f278:16ff:fe63:4d81
    Port:                                 BVI1000
    Version:                              v2
    Query Interval:                       125 seconds
    Robustness:                           2
    Max Resp Time:                        10.0 seconds
    Time since last G-Query:              97 seconds
  Mrouter Ports:                          1
    Dynamic:                              BVI1000
  STP Forwarding Ports:                   0
  ICCP Group Ports:                       0
  Groups:                                 1000
    Member Ports:                         0
  V2 Source Groups:                       1002
```

```
Static/Include/Exclude:              0/1002/0
Member Ports (Include/Exclude):      1002/0
```

Configuring MLD Snooping

Configure

```
RP0/0/RP0/CPU0:router# configure
/* Create two profiles. */
RP0/0/RP0/CPU0:router(config)# mld snooping profile bridge_profile
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# mld snooping profile port_profile
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# mrouter
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# exit
RP0/0/RP0/CPU0:router(config)#

/* Configure two physical interfaces for L2 support.*/
RP0/0/RP0/CPU0:router(config)# interface GigabitEthernet0/8/0/38
RP0/0/RP0/CPU0:router(config-if)# negotiation auto
RP0/0/RP0/CPU0:router(config-if)# l2transport
RP0/0/RP0/CPU0:router(config-if)# no shut
RP0/0/RP0/CPU0:router(config-if)# exit
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# interface GigabitEthernet0/8/0/39
RP0/0/RP0/CPU0:router(config-if)# negotiation auto
RP0/0/RP0/CPU0:router(config-if)# l2transport
RP0/0/RP0/CPU0:router(config-if)# no shut
RP0/0/RP0/CPU0:router(config-if)# exit

/* Add interfaces to the bridge domain. Attach bridge_profile to the bridge domain and
port_profile to one of the Ethernet interfaces.
The second Ethernet interface inherits MLD snooping configuration attributes from the bridge
 domain profile.*/
RP0/0/RP0/CPU0:router(config)# l2vpn
RP0/0/RP0/CPU0:router(config-l2vpn)# bridge group bg1
RP0/0/RP0/CPU0:router(config-l2vpn-bg)# bridge-domain bd1
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd)# mld snooping profile bridge_profile
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd-mld-snooping)# interface GigabitEthernet0/8/0/38
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd-mld-snooping-if)# mld snooping profile port_profile
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd-mld-snooping-if)# interface GigabitEthernet0/8/0/39
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd-mld-snooping-if)# exit
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd-mld-snooping)# exit
RP0/0/RP0/CPU0:router(config-l2vpn-bg-bd)# commit
```

Running Configuration

```
RP0/0/RP0/CPU0:router# show running-config
configuration
 mld snooping profile bridge_profile
!
mld snooping profile port_profile
   mrouter
!

interface GigabitEthernet0/8/0/38
   negotiation auto
   l2transport
   no shut
   !
!
```

```
interface GigabitEthernet0/8/0/39
   negotiation auto
   l2transport
   no shut
   !
!

l2vpn
  bridge group bg1
     bridge-domain bd1
     mld snooping profile bridge_profile
     interface GigabitEthernet0/8/0/38
       mld snooping profile port_profile
     interface GigabitEthernet0/8/0/39
     !
  !
!
```

Verification

Verify the configured bridge ports.

```
RP0/0/RP0/CPU0:router# show mld snooping port


                        Bridge Domain f10:109

                                                        State
Port                                               Oper  STP  Red   #Grps   #SGs
----                                               ----  ---  ---   -----   ----
BVI1009                                              Up    -    -       0      0
GigabitEthernet0/8/0/38                              Up    -    -    1000   1000
GigabitEthernet0/8/0/39                              Up    -    -    1000   1000
```

Configuring MLD Snooping on Ethernet Bundles

This example assumes that the front-ends of the bundles are preconfigured. For example, a bundle configuration might consist of three switch interfaces, as follows:

Configure

```
/*  Configure the front-ends of the bundles consisting of three switch interfaces.*/
RP0/0/RP0/CPU0:router# configure
RP0/0/RP0/CPU0:router(config)# interface bundle-ether 1
RP0/0/RP0/CPU0:router(config-if)# exit
RP0/0/RP0/CPU0:router(config)# interface GigabitEthernet0/0/0/0
RP0/0/RP0/CPU0:router(config-if)# exit
RP0/0/RP0/CPU0:router(config)# interface GigabitEthernet0/0/0/1
RP0/0/RP0/CPU0:router(config-if)# exit
RP0/0/RP0/CPU0:router(config)# interface GigabitEthernet0/0/0/2
RP0/0/RP0/CPU0:router(config-if)# channel-group 1 mode on
RP0/0/RP0/CPU0:router(config-if)# exit
RP0/0/RP0/CPU0:router(config)# interface GigabitEthernet0/0/0/3
RP0/0/RP0/CPU0:router(config-if)# channel-group 1 mode on
RP0/0/RP0/CPU0:router(config-if)# exit

/* Configure two MLD snooping profiles. */
RP0/0/RP0/CPU0:router(config)# mld snooping profile bridge_profile
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# exit     !
```

```
RP0/0/RP0/CPU0:router(config)# mld snooping profile port_profile
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# mrouter
RP0/0/RP0/CPU0:router(config-mld-snooping-profile)# exit

/* Configure interfaces as bundle member links. */

RP0/0/RP0/CPU0:router(config)# interface GigabitEthernet0/0/0/0
RP0/0/RP0/CPU0:router(config-if)# bundle id 1 mode on
RP0/0/RP0/CPU0:router(config-if)# negotiation auto
RP0/0/RP0/CPU0:router(config-if)# exit
RP0/0/RP0/CPU0:router(config)# interface GigabitEthernet0/0/0/1
RP0/0/RP0/CPU0:router(config-if)# bundle id 1 mode on
RP0/0/RP0/CPU0:router(config-if)# negotiation auto
RP0/0/RP0/CPU0:router(config-if)# exit
RP0/0/RP0/CPU0:router(config)# interface GigabitEthernet0/0/0/2
RP0/0/RP0/CPU0:router(config-if)# bundle id 2 mode on
RP0/0/RP0/CPU0:router(config-if)# negotiation auto
RP0/0/RP0/CPU0:router(config-if)# exit
RP0/0/RP0/CPU0:router(config)# interface GigabitEthernet0/0/0/3
RP0/0/RP0/CPU0:router(config-if)# bundle id 2 mode on
RP0/0/RP0/CPU0:router(config-if)# negotiation auto
RP0/0/RP0/CPU0:router(config-if)# exit

/* Configure the bundle interfaces for L2 transport. */
RP0/0/RP0/CPU0:router(config)# interface Bundle-Ether 1
RP0/0/RP0/CPU0:router(config-if)# l2transpor
RP0/0/RP0/CPU0:router(config-if)# exit
RP0/0/RP0/CPU0:router(config)# interface Bundle-Ether 2
RP0/0/RP0/CPU0:router(config-if)# l2transpor
RP0/0/RP0/CPU0:router(config-if)# exit

/* Add the interfaces to the bridge domain and attach MLD snooping profiles. */
RP0/0/RP0/CPU0:router(config)# l2vpn
RP0/0/RP0/CPU0:router(config-l2vpn)# bridge group bg1
RP0/0/RP0/CPU0:router(config-l2vpn-bg)# mld snooping profile bridge_profile
RP0/0/RP0/CPU0:router(config-l2vpn-bg-mld-snooping-profile)# interface bundle-Ether 1
RP0/0/RP0/CPU0:router(config-l2vpn-bg-mld-snooping-profile-if)# mld snooping profile
port_profile
RP0/0/RP0/CPU0:router(config-l2vpn-bg-mld-snooping-profile-if)# interface bundle-Ether 2
RP0/0/RP0/CPU0:router(config-l2vpn-bg-mld-snooping-profile-if)# commit
```

Running Configuration

```
RP0/0/RP0/CPU0:router# show running-config
configuration
 interface Port-channel1
    !
interface GigabitEthernet0/0/0/0
    !
interface GigabitEthernet0/0/0/1
!
    interface GigabitEthernet0/0/0/2
       channel-group 1 mode on
    !
    interface GigabitEthernet0/0/0/3
       channel-group 1 mode on
    !
mld snooping profile bridge_profile
       !
       mld snooping profile port_profile
          mrouter
       !
```

```
interface GigabitEthernet0/0/0/0
         bundle id 1 mode on
         negotiation auto
       !
       interface GigabitEthernet0/0/0/1
         bundle id 1 mode on
         negotiation auto
       !
       interface GigabitEthernet0/0/0/2
         bundle id 2 mode on
         negotiation auto
       !
       interface GigabitEthernet0/0/0/3
         bundle id 2 mode on
         negotiation auto
       !
interface Bundle-Ether 1
          l2transport
          !
       !
       interface Bundle-Ether 2
          l2transport
          !
       !

l2vpn
         bridge group bg1
            bridge-domain bd1
            mld snooping profile bridge_profile
            interface bundle-Ether 1
              mld snooping profile port_profile
            interface bundle-Ether 2
            !
         !
      !
```

Verification

```
RP0/0/RP0/CPU0:router# show mld snooping port
Bridge Domain BG1:BD1
State
Port Oper STP Red #Grps #SGs
---- ---- --- --- ----- ----
HundredGigE0/0/0/3 Up - - 1 1
HundredGigE0/0/0/7 Up - - 1 1
HundredGigE0/19/0/11 Up - - 1 1
HundredGigE0/19/0/5 Up - - 1 1
RP/0/RP1/CPU0:Router#
```

Multicast IRB

Multicast Integrated Routing and Bridging (IRB) enables the routing of multicast packets into and out of a bridge domain through a Bridge-Group Virtual Interface (BVI). The BVI acts as a normal routed interface within the router, enabling seamless integration of multicast routing with existing network infrastructure. This is particularly useful in scenarios where multicast traffic needs to be efficiently managed across different network segments. For details about BVI, refer *Interface and Hardware Component Configuration Guide for Cisco 8000 Series Routers*.

Table 9: Feature History Table

Feature Name	Release Information	Feature Description
Multicast IRB	Release 25.4.1	Introduced in this release on: Fixed Systems (8700 [ASIC: K100])(select variants only*) *This feature is supported on: The feature supports native multicast for MVPN profiles such as 0, 1, 3, 5, 6, 11, 14, 17, 19, 21, 22, and 25. For more information about the supported MVPN profiles, refer to mVPN Profiles within Cisco IOS XR. • 8712-MOD-M • 8711-48Z-M
Multicast IRB	Release 25.2.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100]); Modular Systems (8800 [LC ASIC: P100]) Multicast IRB enables the routing of multicast packets into and out of a bridge domain through a BVI. The feature supports native multicast for MVPN profiles such as 0, 1, 3, 5, 6, 11, 14, 17, 19, 21, 22, and 25. For more information about the supported MVPN profiles, refer to mVPN Profiles within Cisco IOS XR. *This feature is supported on: • 8212-48FH-M • 88-LC1-36EH • 8711-32FH-M

Feature Name	Release Information	Feature Description
Multicast IRB	Release 25.1.1	Introduced in this release on: Modular Systems (8800 [LC ASIC: P100]) (select variants only*). Multicast Integrated Routing and Bridging (IRB) enables the routing of multicast packets into and out of a bridge domain through a Bridge-Group Virtual Interface (BVI). This feature supports: • Both IPv4 and IPv6 protocols • IGMP snooping and MLD snooping on BVI Interfaces, and • Native multicast for MVPN profiles 0, 6, and 14. *This feature is supported on: • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

BVI interfaces are configured to work with existing VRF routes and are integrated using a replication slot mask. This setup ensures that traffic originating from a VRF BVI is efficiently forwarded to the VPN, enhancing network segmentation and security.

Related topics

- Multicast Listener Discovery over BVI. For information, see Multicast Listener Discovery over BVI.
- IPv6 Multicast Listener Discovery Snooping over BVI. For information, see IPv6 Multicast Listener Discovery Snooping over BVI.

Supported Bridge Port Types

- Bundles
- EFPs (physical, vlans, etc)
- Access Pseudowires

Layer 2 multicast ingress route statistics

Layer 2 multicast ingress route statistics is a core multicast capability that

- provides precise, per-route packet count information at the ingress point for Layer 2 multicast traffic, and

- allows for accurate and efficient traffic measurement, crucial for network performance monitoring, and usage-based billing.

Table 10: Feature History Table

Feature Name	Release Information	Feature Description
Layer 2 multicast ingress route statistics	Release 25.3.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100], 8010 [ASIC: A100]); Modular Systems (8800 [LC ASIC: P100]) The feature introduces statistics collection for Layer 2 multicast routes by programming multicast route counters directly on the ingress line card. This implementation allows for more accurate and efficient traffic measurement and provides detailed per-route statistics at ingress. These statistics are essential for operational tasks such as network performance monitoring, usage-based billing, and troubleshooting multicast forwarding and replication issues. Previously, multicast route statistics was available only for Layer 3 multicast routes.

Enhanced Layer 2 multicast route statistics

Multicast route statistics provides information about the multicast routes. The multicast statistics information includes the total number of packets received. The design aims to provide an ingress stats counter for every multicast route supported. When the counterbank runs out of resources, packet forwarding continues without statistics allocation or notification.

The enhancement allows deeper visibility into multicast traffic patterns.

Cisco IOSXR Software counters are always present. If there is limited number of counters available and you want to enable counters on particular prefixes for troubleshooting purposes, you can configure **hw-module route-stats to enable accounting** for multicast routing for a limited number of routes.

Benefits of Layer 2 multicast ingress route statistics

Layer 2 multicast ingress route statistics offers these benefits:

- Precise packet counts at ingress. Ensures accurate and reliable multicast ingress traffic measurement.
- Faster fault resolution. Improves troubleshooting of the packet forwarding and replication issues.

- Improved performance monitoring and billing. Provides accurate per-route ingress statistics for traffic pattern analysis.

Usage guidelines for Layer 2 ingress route statistics

The usage guidelines listed below apply:

- Global counters are a shared resource. Enable statistics for only a subset of multicast routes when resource contention limits the availability of global counters.
- You can enable the route statistics for both IPv4 and IPv6 routes.

Enable Layer 2 multicast ingress route statistics

Follow these steps to enable Layer 2 multicast ingress route statistics.

Procedure

Step 1 Enable the Layer 2 ingress route statistics for all multicast groups or specific groups in a domain or address family.

- Run the `multicast stats-enable` command to enable Layer 2 multicast route statistics for all multicast groups or routes in a domain or address family.

  You can enable the route statistics for IPv4, IPv6, or both. In the sample configuration, the bridge domain BD1 collects statistics for all IPv4 and IPv6 routes. The configuration enables global statistics collection within the bridge domain.

  ```
  Router(config)#l2vpn
  Router(config-l2vpn)#bridge group BG1
  Router(config-l2vpn-bg)#bridge-domain BD1
  Router(config-l2vpn-bg-bd)#multicast stats-enable ipv4-ipv6
  Router(config-l2vpn-bg-bd)#commit
  ```

- Enable the Layer 2 multicast ingress statistics for specific multicast groups in a domain or address family.

  a. Run the `route-policy` command to create a route policy that matches the specific group.

  In the sample configuration, the L2MC_STATS policy enables ingress multicast statistics only for the multicast groups **209.165.202.129** (IPv4) and **2001:db8::/32** (IPv6).

  ```
  Router#config
  Router(config)#route-policy L2MC_STATS
  Router(config-rpl)#if destination in (209.165.202.129, 2001:DB8::/32) then
  Router(config-rpl-if)#set ingress-statistics-enable
  Router(config-rpl-if)#endif
  Router(config-rpl)#end-policy
  ```

  b. Run the `igmp snooping profile` command to create an IGMP snooping profile that uses the route policy to enable ingress statistics for the required multicast groups.

  ```
  Router#config
  Router(config)#igmp snooping profile igmpv2pro1
  Router(config-igmp-snooping-profile)#group policy L2MC_STATS
  Router(config-igmp-snooping-profile)#system-ip-address 209.165.201.1
  Router(config-igmp-snooping-profile)#internal-querier
  ```

c. Run the `igmp snooping profile` command to apply the previously defined IGMP snooping profile to the bridge domain in bridge group.

In the sample configuration, the command associates the IGMP snooping profile `igmpv2pro1` with the bridge domain, enabling IGMP snooping, and applying any linked route policies for ingress multicast statistics.

```
Router(config)#l2vpn
Router(config-l2vpn)#bridge group BG1
Router(config-l2vpn-bg)#bridge-domain BD1
Router(config-l2vpn-bg-bd)#igmp snooping profile igmpv2pro1
```

Step 2 Run the `show l2vpn forwarding bridge-domain` command to verify the Layer 2 ingress route statistics.

Example:

In the sample output, the **Stats Enabled: Yes** field confirms that statistics collection for the multicast route ingress is active. The **Route Packets/Bytes: 1000/128000** field indicates that the counters have recorded 1000 packets and 128,000 bytes. These counters update every 30 seconds.

```
Router#show l2vpn forwarding bridge-domain bg1:bd1 mroute ipv4 hardware location 0/0/CPU0
Bridge-Domain Name: bg1:bd1
  Prefix: (10.0.0.1,232.0.0.1/32)
  P2MP enabled: N
  IRB platform data: {None}, len: 0
  Bridge Port:
    EVPN, Xconnect id: 0x80000001 NH:3.3.3.3 MOI_ID: 0x2000001 Label:24001
    EVPN, Xconnect id: 0x80000001 NH:4.4.4.4 MOI_ID: 0x2000002 Label:24001
    EVPN, Xconnect id: 0x80000001 NH:2.2.2.2 MOI_ID: 0x2000003 Label:24001
    Bundle-Ether100.1


  Platform MCAST Leaf Details:
    BD_ID:0, MC GID [Current:38, Old:0], MRID: 4, Route created in OFA: Yes, MCGID created in OFA:
Yes,
    State: created
    Mcast Route Ingress Stats Information:
    Previous Route Flags: 0x00000041, Current Route Flags: 0x00000049
    Stats Enabled: Yes, Stats created in OFA: Yes
    Route Packets/Bytes: 1000/128000

    Asyn Flags of Last OFA Action: [action:Update, is_retry:F, is_create_failed:F]
    BVI Enabled: Yes, BVI ole NP: 255, BVI ole local: No, BVI ifh: 0x0, BVI ole action: 0
    EVPN Enabled: Yes, EVPN ole NP: 255, EVPN ole local: No

    NPI data:
      trans_id: 77271, Source: 10.0.0.1, Group: 232.0.0.1
```

Step 3 (Optional) Run the `show controller npu resources counterbank` command to verify how many hardware counters are currently in use and how much capacity remains.

Counter banks store hardware statistics such as ingress or egress counters for routes in NPUs (Network Processing Units). Monitoring these banks helps you identify when the resource limit might prevent assigning counters to new routes.

Example:

In the sample output, the **OOR State** indicates that the counter usage is within the threshold.

```
Router#show controller npu resources counterbank location 0/rP0/CPU0
Thu Aug 1 11:54:03.268 UTC
Hw Resource Information
    Name                         : counter_bank
    Asic Type                    : K100
```

```
NPU-0
OOR Summary
        Estimated Max Entries    : 128
        Red Threshold            : 95%
        Yellow Threshold         : 80%
        OOR State                : Green
        High Water Mark          : 16
        High Water Mark Time     : 2024.Jul.31 12:11:55 UTC


Current Hardware Usage
    Name: counter_bank
        Estimated Max Entries    : 128
        Total In-Use             : 14
        OOR State                : Green
        High Water Mark          : 16
        High Water Mark Time     : 2024.Jul.31 12:11:55 UTC
```

EVPN Layer 2 multicast

EVPN Layer 2 multicast is a core multicast feature that

- supports Layer 2 traffic in EVPN networks, and
- offers a unified EVPN framework for unicast, IRB, and multicast services.

Table 11: Feature History Table

Feature Name	Release Information	Feature Description
EVPN Layer 2 multicast	Release 25.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100], 8010 [ASIC: A100]); Modular Systems (8800 [LC ASIC: P100]) The feature enables Layer 2 multicast forwarding across an EVPN network by extending multicast services into the existing EVPN framework. It builds on EVPN signaling to deliver efficient multicast distribution, reduce unnecessary replication, and improve scalability across the EVPN core.

Multicast handling in EVPN Layer 2

EVPN Layer 2 multicast extends the EVPN framework to handle multicast traffic efficiently. The feature actively distinguishes between Inclusive Multicast Ethernet Tag (IMET) based flooding and Selective Multicast

Ethernet Tag (SMET) based selective forwarding. This reduces unnecessary bandwidth usage across the network.

The feature uses IMET (Type 3) routes to forward basic BUM traffic to all PEs in the EVPN domain, and it uses SMET (Type 6) routes to send multicast traffic only to PEs that have active receivers. IMET provides an inclusive flood-and-learn method, while SMET provides selective forwarding and reduces bandwidth by avoiding multicast delivery to PEs without interested listeners.

The feature supports multihoming by using Designated Forwarder (DF) election to select the active forwarder. It also enables fast failover by signaling state changes through EVPN, ensuring quick traffic recovery when a link or node fails.

The data plane drops multicast control packets arriving on EVI ports at ingress to avoid redundant processing. This process ensures stable operation.

Benefits of EVPN Layer 2 multicast

Benefits of EVPN Layer 2 multicast include:

- Uses control-plane multicast signaling to eliminate unnecessary data-plane flooding and improve overall network efficiency.
- Selectively forwards multicast traffic only to interested receivers, reducing replication and conserving bandwidth.
- Supports large multicast groups and high-density deployments across the EVPN network.
- Simplifies operations by providing deterministic multicast forwarding with clear signaling and reduced troubleshooting complexity.
- Aligns with modern multicast transport needs and serves as the preferred model for scalable, efficient Layer 2 multicast delivery.

Use case: EVPN Layer 2 multicast with multihoming and IGMP snooping

In a network with multiple PE devices operating within a Layer 2 EVPN core, PE1 and PE2 form an EVPN multihoming (MH) All-Active pair connected to a CE device. The network administrator enables IGMP Snooping on the Layer 2 transport interfaces connecting to the CE.

Summary

This use case illustrates how EVPN Layer 2 multicast effectively integrates IGMP snooping, BGP EVPN route types, and Designated Forwarder (DF) election to provide efficient and resilient multicast services within a multihomed environment.

Workflow

Figure 1: Multihomed EVPN Layer 2 multicast with IGMP snooping

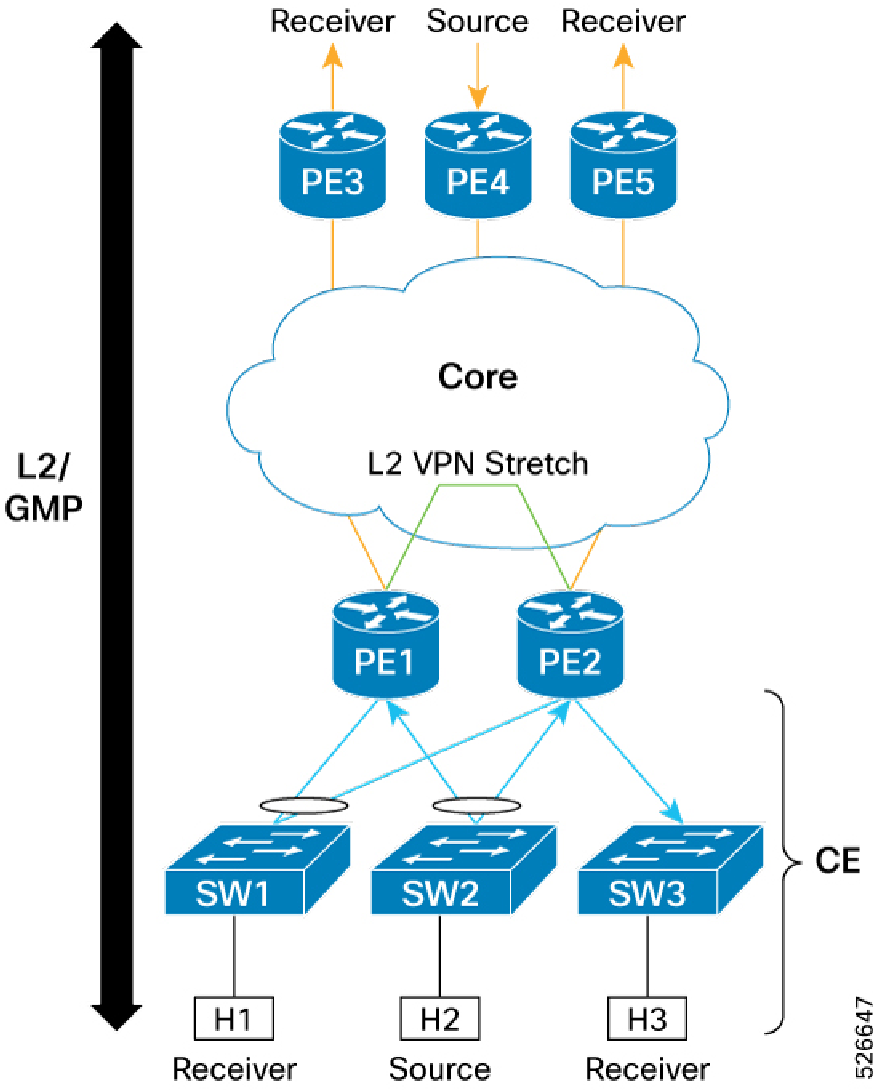

These stages describe the EVPN Layer 2 multicast with multihoming and IGMP snooping use case:

1. EVPN transport setup
    - All PE devices operate as EVPN PEs connected through an EVPN Layer 2 core.
    - PE1 and PE2 form an EVPN multihoming pair in All-Active mode for the CE.
2. IGMP snooping and state creation
    - A receiver connected to the CE sends an IGMP Join message to the multihomed PEs.
    - Based on the CE Link Aggregation Group (LAG) hash, either PE1 or PE2 intercepts the IGMP Join. The receiving PE then establishes the necessary multicast state for the group.

    - The receiving PE actively synchronizes the multicast receiver state with its peer by sending the IGMP-join-sync messages through BGP.

3. Multicast traffic delivery differs based on the forwarding mode:
    - In the IMET (Inclusive Multicast Ethernet Tag) mode, the source PE sends multicast traffic to all PEs in the EVPN domain.
    - In the SMET (Selective Multicast Ethernet Tag) mode, only PEs with active receivers advertise their interest using BGP Route Type 6. The source PE then forwards multicast traffic exclusively to those PEs.

   Regardless of the mode, only the EVPN multihoming DF PE forwards the multicast packet to the receivers behind the CE device. The Non-Designated Forwarder (NDF) PE drops the multicast stream for that multihomed segment to prevent duplicate delivery.

Usage guidelines and restrictions for EVPN Layer 2 multicast

Usage guidelines

Usage guidelines for EVPN Layer 2 multicast include these points:

- Enable Layer 2 EVPN before you configure multicast.
- You can deploy multicast with or without IRB.
- Use Single-active, Port-active, or All-active load-balancing modes when you deploy EVPN Layer 2 multicast with or without BVI. Release 25.4.1 supports EVPN IRB multicast with a distributed anycast gateway.
- The feature supports multicast transport over SR/MPLS, SR-TE, BGP-LU, or LDP over RSVP-TE EVPN core types.
- The feature relies on SMET (Type-6) routes for optimized multicast forwarding and on IMET (Type-3) routes for BUM replication. This approach ensures the correct multicast paths are installed.
- The feature uses IGMP or MLD snooping on access interfaces to dynamically learn receiver membership.

Restrictions

Restrictions for EVPN Layer 2 multicast include these points:

- The feature supports IGMP snooping in both IMET and SMET modes.
- The feature does not support MLD snooping in SMET mode.
- The feature supports IRB in a distributed anycast gateway model.

Configure EVPN Layer 2 multicast

Use the steps in this section to complete the configuration.

Before you begin

Enable Layer 2 EVPN before you configure multicast.

Procedure

Step 1 Run the **interface** command to define the Layer 2 interface.

```
Router#configure
Router(config)#interface bundle-ether2.1 l2transport
Router(config-subif-l2)#no shutdown
Router(config-subif-l2)#encapsulation dot1q 1
Router(config-subif-l2)#rewrite ingress tag pop 1 symmetric
Router(config-subif-l2)#commit
Router(config-subif-l2)#exit
```

Step 2 Define the L2VPN bridge group, bridge domain, and bridge port.

Example:

```
Router(config)#l2vpn
Router(config-l2vpn)#bridge group bg1
Router(config-l2vpn-bg)#bridge-domain bd1
Router(config-l2vpn-bg-bd)#interface bundle-Ether 2.1
```

Step 3 Enable IGMP or MLD snooping under the bridge domain to monitor multicast group memberships on interfaces for optimal traffic forwarding.

```
Router(config-l2vpn-bg-bd)#igmp snooping profile test
Router(config-l2vpn-bg-bd)#mld snooping profile test
```

Step 4 Run the **evi** command to enable EVPN instance association.

This step indicates that the bridge domain is participating in an EVPN service. The service extends Layer 2 across an MPLS or IP core to other PE devices that are also part of the EVPN instance.

```
Router(config-l2vpn-bg-bd)#evi 1
```

Step 5 Define the IGMP snooping profile to ensure that the CE hosts send IGMP reports and allows the PE to build receiver state.

In the sample configuration, the **internal-querier** command enables an internal IGMP querier. This allows the switch or router to act as the IGMP querier when no Layer 3 multicast router is present.

```
Router(config)#igmp snooping profile test
Router(config-igmp-snooping-profile)#system-ip-address 10.10.10.1
Router(config-igmp-snooping-profile)#internal-querier
Router(config-igmp-snooping-profile)#commit
```

Step 6 (Optional) Configure IGMP snooping proxy for `evi 1` to enable the SMET mode on the associated bridge domain.

If you do not configure IGMP snooping proxy, IMET mode is used for both IGMP and MLD snooping.

```
Router(config)#evpn
Router(config-evpn)#evi 1
Router(config-evpn-instance)#proxy
Router(config-evpn-instance-proxy)#igmp-snooping
Router(config-evpn-instance-proxy)#commit
```

Step 7 Run the **show igmp snooping bridge-domain** command to verify the IGMP snooping configuration, operational status, and various statistics for the specified bridge domain.

In the sample output, **EVPN Selective Mcast: Disabled** field indicates that SMET is not enabled for the bridge domain. This implies that for EVPN multicast, this bridge domain is operating in IMET mode.

EVPN: Enabled field indicates that EVPN is active for the bridge domain. IGMP snooping is integrated with the EVPN framework for multicast.

```
Router#show igmp snooping bridge-domain bg1:bd1 detail statistics
Wed Nov 12 04:30:47.563 UTC

Bridge Domain        Profile                   Act  Ver  #Ports  #Mrtrs  #Grps   #SGs
-------------        -------                   ---  ---  ------  ------  -----   ----
bg1:bd1              test                        Y   v3       5       0      6      0

  Profile Configured Attributes:
    System IP Address:                   1.1.1.1
    Minimum Version:                     2
    Report Suppression:                  Enabled
    Unsolicited Report Interval:         1000 (milliseconds)
    TCN Query Solicit:                   Disabled
    TCN Membership Sync:                 Disabled
    TCN Flood:                           Enabled
    TCN Flood Query Count:               2
    Router Alert Check:                  Enabled
    TTL Check:                           Enabled
    nV Mcast Offload:                    Disabled
    EVPN Selective  Mcast:               Disabled
    EVPN:                                Enabled
    BVI:                                 Disabled
    Internal Querier Support:            Enabled
    Internal Querier Version:            3
    Internal Querier Timeout:            0 (seconds)
    Internal Querier Interval:           60 (seconds)
    Internal Querier Max Response Time:  10.0 (seconds)
    Internal Querier Robustness:         2
    Internal Querier TCN Query Interval: 10 (seconds)
    Internal Querier TCN Query Count:    2
    Internal Querier TCN Query MRT:      0 (seconds)
    Querier Query Interval:              60 (seconds)
    Querier LMQ Interval:                1000 (milliseconds)
    Querier LMQ Count:                   2
    Querier Robustness:                  2
    Startup Query Interval:              15 seconds
    Startup Query Count:                 2
    Startup Query Max Response Time:     10.0 seconds
    Mrouter Forwarding:                  Enabled
    P2MP Capability:                     Disabled
    Default IGMP Snooping profile:       Disabled
    IP Address:                          1.1.1.1
    Port:                                Internal
    Version:                             v3
    Query Interval:                      60 seconds
    Robustness:                          2
    Max Resp Time:                       10.0 seconds
    Time since last G-Query:             52 seconds
  Internal Querier Statistics (elapsed time since last cleared 23:04:52):
    Rx General Queries:                          0
    Rx General Queries When Disabled:            0
    Rx General Queries As Querier:               0
    Rx General Queries As Non Querier:           0
    Rx General Queries As Winner:                0
    Rx General Queries As Loser:                 0
    Rx Global Leaves:                            0
    Rx Global Leaves When Disabled:              0
    Rx Global Leaves As Non Querier:             0
    Rx Global Leaves Ignored:                    0
    Rx Pim Enabled Notifications:                0
    Rx Pim Disabled Notifications:               0
    Rx Local Query Solicitations:                0
    Tx General Queries:                          0
  Ingress-stats Enabled:                  TRUE
  Egress-stats Enabled:                  TRUE
  Mrouter Ports:                         0
  STP Forwarding Ports:                  0
  ICCP Group Ports:                      0
  Groups:                                6
    Member Ports:                        6
  V3 Source Groups:                      0
    Static/Include/Exclude:              0/0/0
    Member Ports (Include/Exclude):      0/0
  Traffic Statistics (elapsed time since last cleared 22:22:10):
                                     Received  Reinjected   Generated
    Messages:                             394           0        1435
      IGMP General Queries:                 0           0        1351
      IGMP Group Specific Queries:          0           0           0
      IGMP G&S Specific Queries:            0           0           0
      IGMP V2 Reports:                    394           0           0
      IGMP V3 Reports:                      0           0          84
      IGMP V2 Leaves:                       0           0           0
      IGMP Global Leaves:                   0           -           0
      PIM Hellos:                           0           0           -
    Rx Packet Treatment:
      Packets Flooded:                                  0
      Packets Forwarded To Members:                     0
      Packets Forwarded To Mrouters:                    0
      Packets Consumed:                               394
    Rx IGMP EVPN Report Group Record Types:
      EVPN Change To Exclude:                           2
      EVPN IGMP V3 reports:                             2
    Rx Errors:
      None
    Rx Other:
      None
    Tx Errors:
      No Querier in BD:                                 1
  EVPN Multihome Statistics (elapsed time since last cleared 23:04:52):
    EVPN Messages Received:                             2
    EVPN Join Add Received:                             2
    EVPN Messages Sent:                                 4
    EVPN Join Add Sent:                                 4
  Startup Query Sync Statistics:
    None
```

Step 8 Run the **show evpn igmp evi** command to display the EVPN Layer 2 multicast receiver state learned through IGMP or MLD snooping and proxy signaling for a specific EVI.

In the sample output, the **Type** field indicates how the entry was learned. The **JOIN** field confirms the EVPN Layer 2 multicast participation. EVPN learns proxy signaling messages from the peer EVPN PE operating in IGMP snooping SMET mode.

```
Router#show evpn igmp evi 1
Wed Nov 12 04:15:52.481 UTC

EVI   Ethernet Segment         (S,G)                                       Source
       Type
----- ------------------------ -------------------------------------------
------------------------------ ------
1     N/A                      (*,225.0.10.1)                             FourHundredGigE0/0/0/1[UR]
     JOIN
```

```
1     N/A                    (*,225.0.10.2)                               FourHundredGigE0/1/0/1[UR]
     JOIN
1     N/A                    (*,226.0.30.1)                               FourHundredGigE0/0/0/1[UR]
     JOIN
1      0001.0101.0101.0101.0101 (*,229.0.0.1)                               2.2.2.2
       JOIN
1      0001.0101.0101.0101.0101 (*,229.0.0.2)                               2.2.2.2
       JOIN
1      0002.0202.0202.0202.0202 (*,230.0.0.1)                               Bundle-Ether200.1
       JOIN
1     N/A                    (*,ff1e:10::1)                               FourHundredGigE0/0/0/1[UR]
     JOIN
1     N/A                    (*,ff1e:10::2)                               FourHundredGigE0/1/0/1[UR]
     JOIN
1     N/A                    (*,ff1e:90::1)                               FourHundredGigE0/1/0/1[UR]
     JOIN
1      0001.0101.0101.0101.0101 (*,ff1e:90::1)                              2.2.2.2
       JOIN
1      0002.0202.0202.0202.0202 (*,ff1e:90::1)                              Bundle-Ether200.1
       JOIN
1      0001.0101.0101.0101.0101 (*,ff1e:90::2)                              2.2.2.2
       JOIN
1      0002.0202.0202.0202.0202 (*,ff1e:90::2)                              Bundle-Ether200.1
       JOIN
1     N/A                     (*,226.0.30.1)                              3.3.3.3
       PROXY
1     N/A                     (*,226.0.30.2)                              3.3.3.3
       PROXY
1     N/A                     (*,226.0.40.1)                              4.4.4.4
       PROXY
1     N/A                     (*,226.0.40.2)                              4.4.4.4
       PROXY
1     N/A                     (10.10.10.1,232.0.0.1)                      3.3.3.3
       PROXY
1     N/A                     (192.0.1.0,230.0.0.1)                       3.3.3.3
       PROXY
```

Step 9 Run the **show l2route evpn imet evi** command to view the IMET route details within the Layer 2 EVPN address family for a specific EVI.

```
Router#show l2route evpn imet evi 1 detail
Wed Nov 12 04:02:19.728 UTC
Topology ID Producer    Originating Router IP Addr  Eth tag ID Flags Type Label   Tunnel ID
           Encap Type
----------- ----------- -------------------------- ---------- ----- ---- -------
--------------------------- -----------
1           L2VPN       2.2.2.2                    0          0     0    24003   N/A
           Mpls
   Last Update: Tue Nov 11 06:10:23.338 UTC
```

EVPN IRB multicast with distributed anycast gateway

EVPN IRB with distributed anycast gateway is a multicast feature that

- enables multicast forwarding in EVPN IRB deployments, and
- improves resiliency and convergence by localizing gateway functions at the PE routers.

Table 12: Feature History Table

Feature Name	Release Information	Feature Description
EVPN IRB multicast with distributed anycast gateway	Release 25.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100]; Modular Systems (8800 [LC ASIC: P100]) The feature enables multicast forwarding in EVPN IRB deployments that use a distributed anycast gateway. Provider Edge (PE) routers use a shared BVI address to operate as local gateways for the same subnet. This shared address allows each router to route and bridge traffic locally.

The feature extends the capabilities of Layer 2 EVPN to support multicast traffic, specifically when IRB is deployed with a distributed anycast gateway model. In this setup, multiple PE devices act as gateways for the same IP subnet, sharing the same BVI address. This arrangement provides optimal forwarding and high availability for both unicast and multicast traffic. The feature improves network resiliency and convergence by localizing gateway functions at each PE.

EVPN IRB distributed anycast gateway overview

The EVPN IRB multicast with distributed anycast gateway scenario enables multicast forwarding in EVPN IRB deployments by localizing gateway functions at each PE router. In this setup, multiple PEs share the same BVI IP and MAC address, allowing each PE to act as a local gateway for the same IP subnet. This shared anycast gateway address facilitates optimal forwarding and high availability for both unicast and multicast traffic.

With active-active multi-homing, all PEs forward traffic and send Protocol Independent Multicast (PIM) join messages, regardless of Designated Forwarder status. This approach provides redundancy and minimizes multicast downtime during failover. Layer 3 traffic between hosts in different subnets is routed symmetrically at both the source and destination PEs over IP or MPLS tunnels, optimizing traffic flow and supporting seamless host mobility.

This distributed anycast gateway model improves network resiliency and convergence by distributing routing capabilities closer to the access network, enabling efficient inter- and intra-subnet traffic flow while maintaining synchronization of multicast group membership and routing state across the EVPN fabric.

Use case: EVPN IRB multicast with all-active multihoming and Layer 3 source

The all-active multihoming and Layer 3 source use case includes the following key components:

- The source connects to the network through PE3, originating the multicast data stream.
- PE routers (PE1, PE2, PE3) are edge devices that form the provider network. PE3 receives the multicast source. PE1 and PE2 send traffic toward the receivers and serve as the BVI anycast gateway.
- Core: The service provider backbone network interconnects the PE routers.

- CE router: Includes the local network. SW1, SW2, and SW3 aggregate traffic from the hosts. H1, H2, and H3 request and receive multicast traffic.

Summary

The use case explains how the multicast state is synchronized across dual-homed PEs using active-active multihoming, ensuring seamless multicast delivery and loop-free forwarding in an EVPN environment.

Workflow

Figure 2: All-active multihoming and Layer 3 source

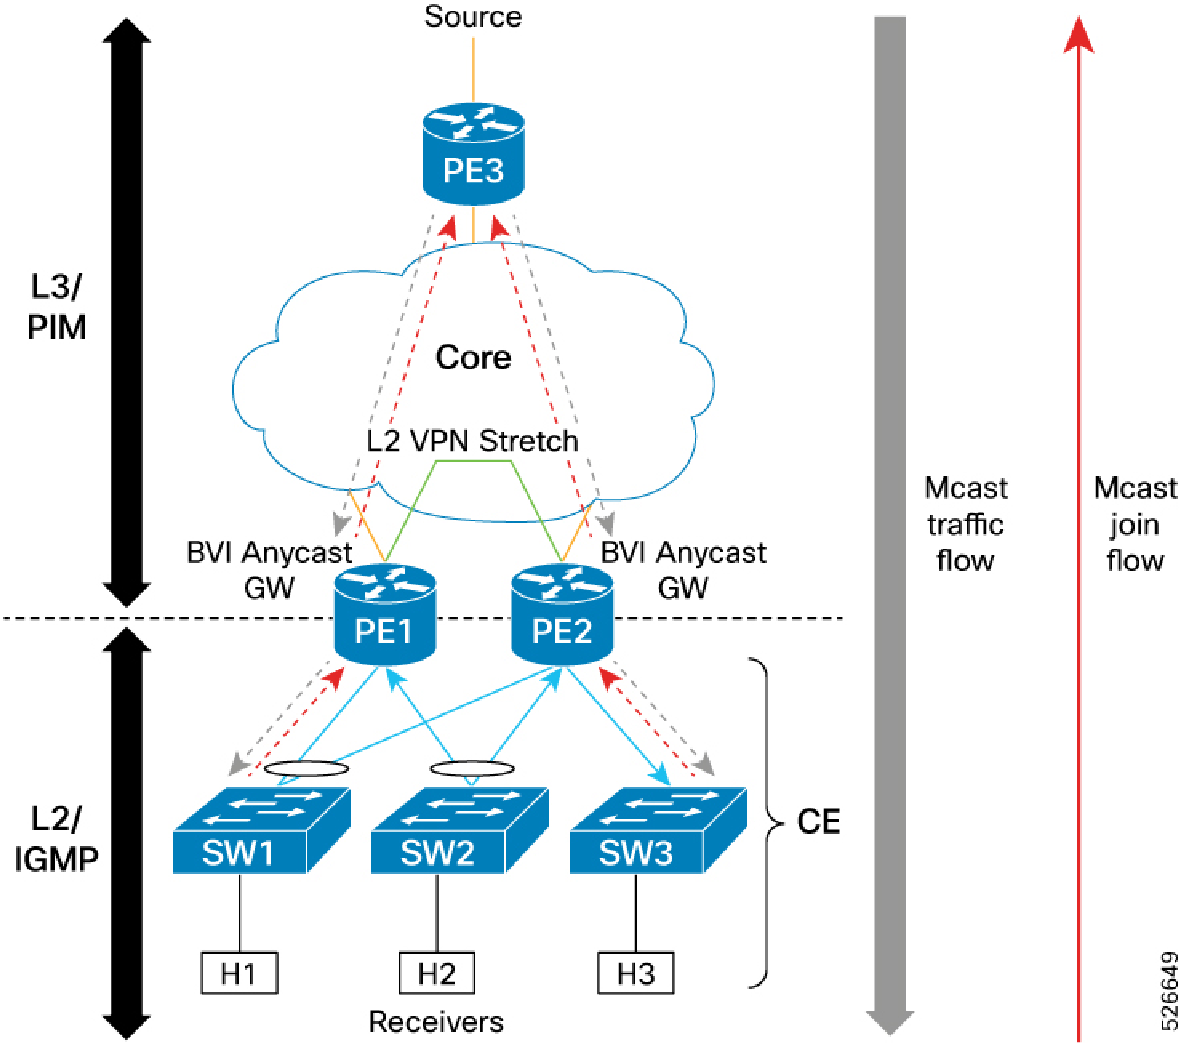

These are the stages of the EVPN IRB multicast with all-active multihoming and Layer 3 source use case:

1. Multihoming and IGMP termination

   PE1 and PE2 operate in EVPN all-active multihoming and terminate IGMP on the anycast BVI. IGMP snooping enabled on the EVPN Layer 2 (EVPN Layer 2 transport). Both PEs host a distributed anycast gateway (BVI) for the multicast VLAN.

2. IGMP join and state synchronization

    - Either PE1 or PE2 receives the IGMP join from the CE, based on the CE LAG hashing.
    - The receiving PE creates the local multicast receiver state and synchronizes it to the peer PE using EVPN IGMP sync routes.

3. PIM and IGMP control plane isolation

    Both PEs act as last-hop routers and independently send PIM Joins toward the Layer 3 multicast source to receive the multicast stream.

4. PIM join and control plane loop prevention

    Both PEs act as last-hop routers and independently send PIM joins to the Layer-3 multicast source to receive the multicast stream. The PEs do not form PIM adjacencies with each other and drop multicast control packets received over the EVPN tunnel to prevent control-plane loops.

5. Multicast data forwarding

    Both PEs receive the multicast traffic, but only the EVPN multihoming Designated Forwarder forwards traffic to the CE.

Use case: Multicast over EVPN all-active multihoming with BVI anycast gateway

The multicast over EVPN all-active multihoming with BVI anycast gateway use case includes these key components:

- Source (H1) or CE (SW1): The multicast source connects to the network through the CE switch.
- PE1 and PE2: These PEs operate in EVPN all-active multihoming and act as first-hop routers for the source.
- BVI anycast gateway: The BVI terminates Layer 3 and IGMP, presenting a single logical gateway.
- L2 EVPN: EVPN provides Layer 2 connectivity and state synchronization between PE1 and PE2.
- Core: The Layer 3 core uses PIM to forward multicast traffic.
- PE3 to PE5: The PEs host the multicast receivers.

Summary

The use case describes multicast traffic in an EVPN all-active multihoming environment using a BVI as the Layer 3 gateway.

Workflow

Figure 3: EVPN all-active multihoming with BVI anycast gateway

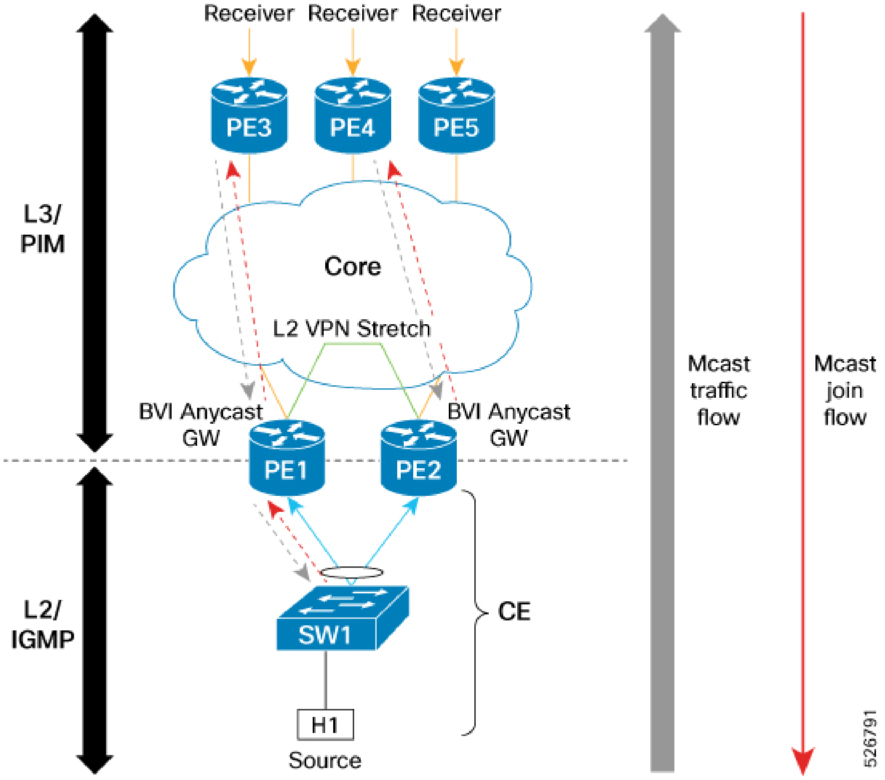

These are the stages of the multicast over EVPN all-active multihoming with BVI anycast gateway use case:

1. Remote receivers that are connected to PE3, PE4, and PE5 send PIM Join messages toward the multicast source. PE1 and PE2 both advertise reachability to the source subnet, so the core delivers PIM joins to both PE1 and PE2.

2. Host H1 sends multicast traffic to the CE switch (SW1). SW1 load-balances the traffic and forwards it to either PE1 or PE2 in the all-active multihoming setup.

3. One PE receives the multicast traffic and forwards a copy across the L2 EVPN network to the peer PE through the BVI. This ensures that both PEs receive the multicast stream.

4. Both PE1 and PE2 forward the multicast traffic through the Layer 3 core to remote receivers. IGMP snooping on the Layer 2 transport also delivers the traffic to any local receivers.

EVPN multihoming active-active

EVPN multihoming access gateway enables redundant network connectivity by allowing a CE device to connect to more than one PE devices. Disruptions to network connectivity are prevented because multihoming allows a CE device to connect to one or multiple PE devices. An Ethernet segment consists of Ethernet links that connect a CE device to multiple PE devices.

Table 13: Feature History Table

Feature Name	Release Information	Feature Description
EVPN multihoming active-active	Release 25.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100], 8010 [ASIC: A100]); Modular Systems (8800 [LC ASIC: P100]) Enable EVPN multihoming access gateway to provide redundant connectivity for CE devices through connections to multiple PE devices. This approach ensures high availability and provides seamless Layer 2 and Layer 3 services in EVPN IRB networks. It supports both VPNv4 and VPNv6 address families and enables efficient host route distribution across data centers.

How EVPN multihoming active-active works

Summary

EVPN IRB multihoming active-active workflow includes these key components:

- **CE or host routers**:
  - CE1 or Host-1 with IP address 10.0.0.1/32
  - CE2 or Host-2 with IP address 10.0.0.2/32
  - CE3 or Host-3 with IP address 10.0.0.3/32
  - CE5 or Host-5 with IP address 20.0.0.1/32
- **PE routers**:
  - PE1 with BVI IP address 10.0.0.5/24 and BVI MAC address 0.0.5
  - PE2 with BVI IP address 10.0.0.5/24 and BVI MAC address 0.0.5
  - PE5
- **Interfaces**:
  - IF1: Connects Host-1 to PE1
  - IF2: Connects Host-3 to PE1
  - IF3: Connects PE1 to Host-3
  - IF4: Connects PE2 to Host-3
  - IF5: Connects Host-3 to PE2
- **Paths**:
  - Path 1: Host-1 communicates with Host-2 using a EVPN single-homing access EVPN gateway interface

  - Path 2: Host-1 communicates with Host-3 using a EVPN single-homing access EVPN gateway interface
  - Path 3: Host-3 communicates with Host-5 using a EVPN single-active multihoming interface
  - Path 4: Host-2 communicates with Host-5 using a EVPN single-active multihoming interface

- Host routing is enabled on PE1 and PE2.
- IRB interfaces are configured as anycast.
- Host-1 and Host-2 connect to PE1 and PE2 using EVPN single-homing interfaces. Host-3 connects to PE1 and PE2 using a EVPN multihoming interfaces.
- IOS XR Software performs the Designated Forwarder (DF) election for the shared Ethernet Segment Identifier (ESI) and elects IF2 as DF and IF5 as Non-Designated Forwarder (NDF).
- Interface IF5 is set to blocked state, which blocks both BUM and unicast traffic.

Figure 4: EVPN IRB multihoming network topology

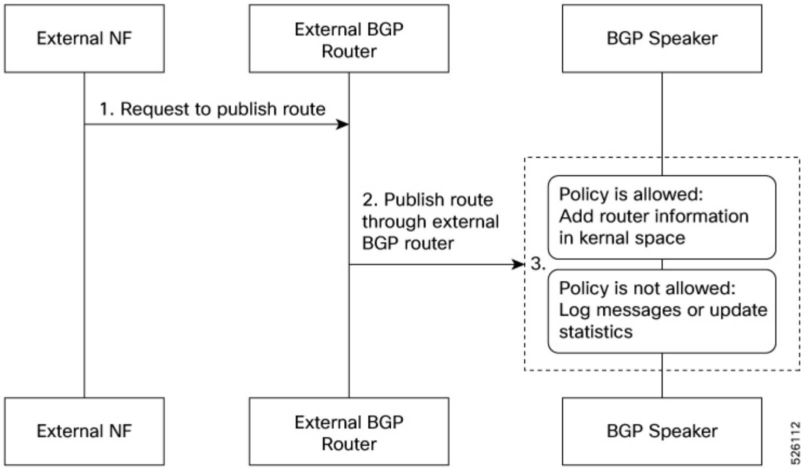
External NF
External BGP Router
BGP Speaker
1. Request to publish route
2. Publish route through external BGP router
3.
Policy is allowed: Add router information in kernal space
Policy is not allowed: Log messages or update statistics
External NF
External BGP Router
BGP Speaker
526112

Workflow

These scenarios describe the EVPN IRB multihoming network topology.

1. Host-2 wants to communicate with Host-5, which is in a different subnet.
   - Host-2 sends an ARP request to its gateway which is IRB interface. It basically ARPs the BVI IP address 10.0.0.5.
   - PE2 learns the Host-2 MAC and IP addresses from these ARP packets and uses this information to program the forwarding adjacency.
   - The BVI interface on PE2 sends an ARP response to Host-2 using its BVI IP address 10.0.0.5 and MAC address 0.0.5.
   - PE2 advertises Host-2 route using EVPN route type-2 to remote PEs. Remote PEs, such as PE5, import and install this route as a remote route.
   - Since Host-5 is directly connected to PE5, it receives the ARP request and responds with a unicast ARP response.

- The ARP process ensures that PE5 learns Host-5 IP address 20.0.0.1/32, enabling communication between Host-2 and Host-5.
- If PE2 does not have a Host-5 specific route, it may use an EVPN route type-5 to forward traffic to PE5, where ARP resolves Host-5, enabling Host-2 and Host-5 to communicate.

2. Host-5 sends a packet to Host-2. If Host-2 has not communicated yet, PE5 might not have Host-2 specific route.

- If PE5 directs traffic to PE1 first, a Generalized Learning (G-LEAN) adjacency process occurs, and traffic is dropped until it is resolved.
- Once PE5 identifies PE2 as the best destination for Host-2, it forwards the packet to PE2, and PE2 performs these steps:
  - PE2 performs an IP lookup, finding the BVI interface as the destination.
  - Destination MAC is set to Host-2 MAC as learned by ARP and source MAC remains as the BVI MAC address 0.0.5.
  - PE2 performs a MAC lookup within the bridge domain and forwards the packet to Host-2.

EVPN single-active multihoming for anycast gateway IRB

EVPN single-active multihoming for anycast gateway IRB is a multicast feature that:

- supports single-active redundancy mode,
- enables PE nodes locally connected to an Ethernet Segment to load balance traffic to and from the Ethernet Segment based on the EVPN service instance (EVI),
- ensures only one PE forwards traffic to and from the Ethernet Segment (ES) within an EVPN service instance, and
- operates exclusively in intersubnet scenarios.

Table 14: Feature History Table

Feature Name	Release Information	Feature Description
EVPN single-active multihoming for anycast gateway IRB	Release 25.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100], 8010 [ASIC: A100]); Modular Systems (8800 [LC ASIC: P100]) Enable EVPN single-active multihoming for anycast gateway IRB to provide single-active redundancy, where only one PE forwards traffic for an Ethernet Segment within each EVPN service instance. This feature supports intersubnet scenarios and balances traffic based on the EVI.

How EVPN single-active multihoming works

Summary

EVPN IRB single-active multihoming workflow includes these key components:

- **CE or host routers**:
  - CE1 or Host-1 with IP address 10.0.0.1/32
  - CE2 or Host-2 with IP address 10.0.0.2/32
  - CE3 or Host-3 with IP address 10.0.0.3/32
  - CE5 or Host-5 with IP address 20.0.0.1/32
- **PE routers**:
  - PE1 with BVI IP address 10.0.0.5/24 and BVI MAC address 0.0.5
  - PE2 with BVI IP address 10.0.0.5/24 and BVI MAC address 0.0.5
  - PE5
- **Interfaces**:
  - IF1: Connects Host-1 to PE1
  - IF2: Connects Host-3 to PE1
  - IF3: Connects PE1 to Host-3
  - IF4: Connects PE2 to Host-3
  - IF5: Connects Host-3 to PE2
- **Paths**:
  - Path 1: Host-1 communicates with Host-2 using a EVPN single-homing access EVPN gateway interface
  - Path 2: Host-1 communicates with Host-3 using a EVPN single-homing access EVPN gateway interface
  - Path 3: Host-3 communicates with Host-5 using a EVPN single-active multihoming interface
  - Path 4: Host-2 communicates with Host-5 using a EVPN multihoming active-active interface
- Switch device: SW1, connects Host-3 to PE1 and PE2.
- Host routing is enabled on PE1 and PE2.
- IRB interfaces are configured as anycast. Interface IF5 is set to blocked state which blocks both BUM and unicast traffic.
- Host-1 and Host-2 connect to PE1 and PE2 using EVPN single-homing interfaces. Host-3 connects to PE1 and PE2 using a EVPN multihoming interfaces.

- IOS XR Software performs the Designated Forwarder (DF) election for the shared Ethernet Segment Identifier (ESI) and elects IF2 as DF and IF5 as Non-Designated Forwarder (NDF).

Workflow

Figure 5: EVPN IRB single-active multihoming network topology

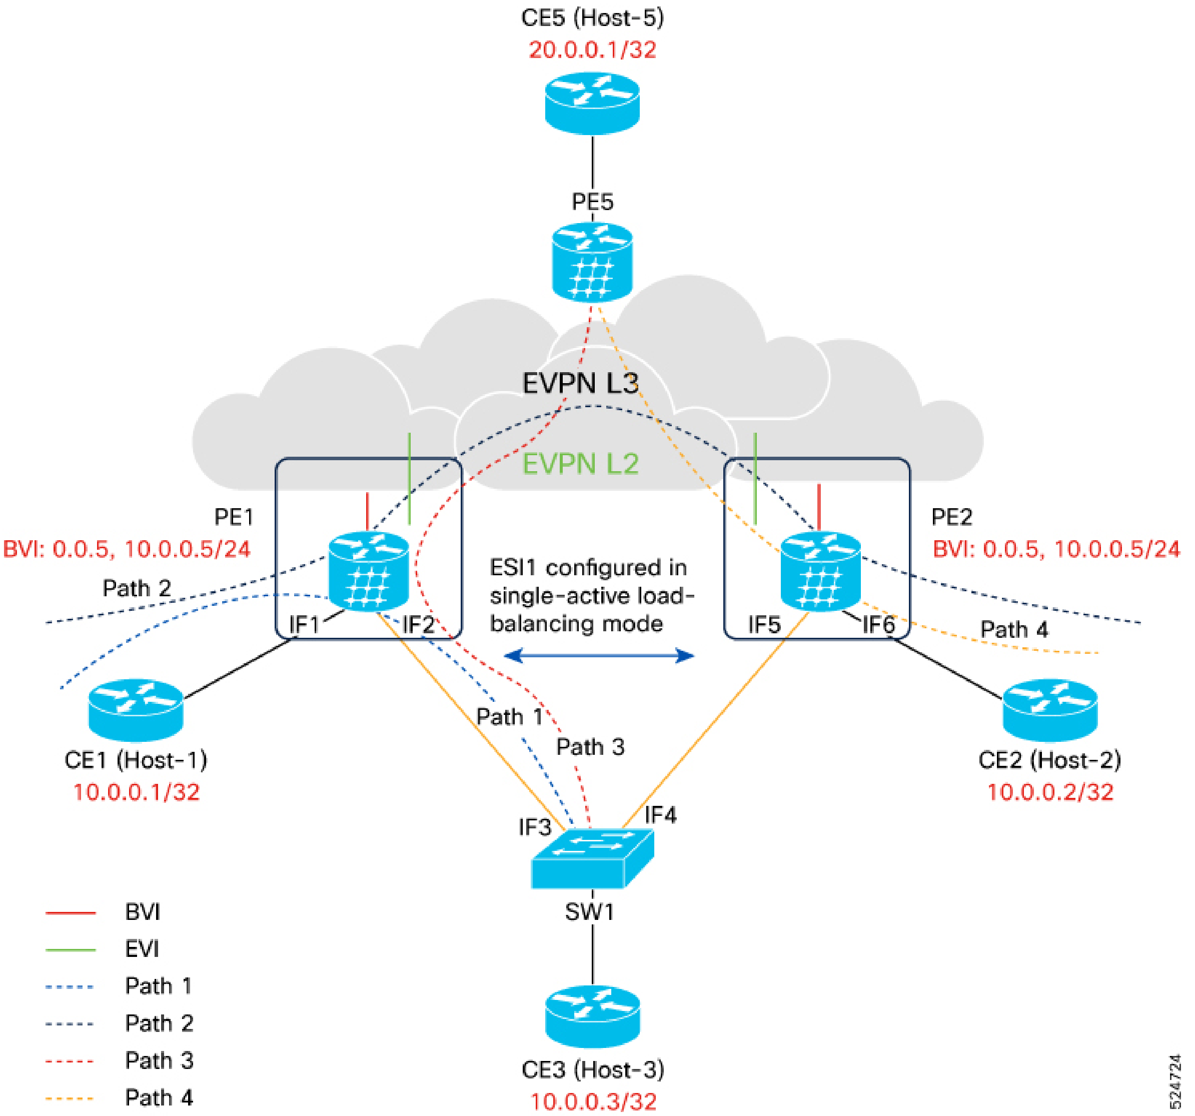

These scenarios describe the EVPN IRB multihoming network topology.

1. Host-3 wants to communicate with Host-5, which are in different subnets.

    - Host-3 sends an ARP request to its IRB gateway, configured with the BVI IP address 10.0.0.5.
    - SW1 learns Host-3 MAC and IP addresses as it forwards the ARP response to PE1, which is the DF. The packet sent to PE2 is dropped since IF5 is in a blocked state as a NDF.

  - PE1 learns Host-3's MAC and IP addresses from the ARP packet and replicates this information across all output interfaces associated with the BVI interface.
  - The BVI interface on PE1 sends an ARP response to Host-3 using its BVI IP address 10.0.0.5 and MAC address 0.0.5.

    SW1 updates its MAC address table with the BVI MAC address of PE1 as it forwards the ARP response to Host-3.
  - PE1 advertises Host-3 host route through EVPN using route type-2. Remote PEs, including PE2 and PE5, learn about Host-3 and install the route in their hardware tables.
  - Since Host-5 is directly connected to PE5, it receives an ARP request and responds with a unicast ARP response.
  - The ARP process ensures that PE5 learns Host-5 IP address 20.0.0.1/32, enabling communication between Host-5 and Host-3.

2. Where Host-5 sends a packet to Host-3. If Host-3 has not communicated yet, PE5 might not have Host-3 specific route.
   - If PE5 directs traffic to PE2, a Generalized Learning (G-LEAN) adjacency process occurs, and traffic is dropped until it is resolved.
   - If PE2 receives, it performs these steps:
     - PE2 floods an ARP request within the bridge domain to resolve Host-3 MAC address. However, as PE2 is the NDF, direct forwarding to Host-3 is not possible due to the blocked state of IF5.
     - A copy of the ARP request is sent to PE1 through the L2 stretch.
     - PE1 forwards the ARP request to Host-3. Once Host-3 responds to the ARP request, PE1 learns Host-3 MAC address.
     - After receiving the ARP response, PE1 updates the route for Host-3 and advertises it as an EVPN route type-2 across the network.
     - This allows packets from Host-5 to reach Host-3 efficiently once the address resolution is complete.

Configure EVPN single-active multihoming

Perform the following tasks on PE1 and PE2 to configure EVPN single-active multihoming:

Procedure

Step 1 Configure EVPN IRB with host routing.

Step 2 Run the **ethernet-segment** command to configure the EVPN ethernet segment.

Example:

```
Router# configure
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether100
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 40.00.00.00.00.00.00.00.01
Router(config-evpn-ac-es)# load-balancing-mode single-active
Router(config-evpn-ac-es)# bgp route-target 4000.0000.0001
Router(config-evpn-ac-es)# commit
```

Step 3 Run the **evi** command to configure the EVI parameters.

Example:

```
Router# configure
Router(config)# evpn
Router(config-evpn)# evi 1
Router(config-evpn-evi)# bgp
Router(config-evpn-evi-bgp)# rd 200:50
Router(config-evpn-evi-bgp)# route-target import 100:6005
Router(config-evpn-evi-bgp)# route-target export 100:6005
Router(config-evpn-evi-bgp)# commit
```

Step 4 Run the **interface** command to define the Layer 2 interface.

Example:

```
Router# configure
Router(config)# interface Bundle-Ether100.1 l2transport
Router(config-subif-l2)# no shutdown
Router(config-subif-l2)# encapsulation dot1q 1
Router(config-subif-l2)# rewrite ingress tag pop 1 symmetric
Router(config-subif-l2)#commit
Router(config-subif-l2)#exit
```

Step 5 Define the bridge domain, bridge group, bridge port, and EVI for the Layer 2 interface.

When deploying EVPN IRB with a multicast-based distributed gateway on Cisco 8000 routers, you must configure a `split-horizon group` on the core interface.

When you configure the `split-horizon group` core under the BVI interface:

- Layer 2 multicast traffic flows from the EVI to the BVI.
- Layer 3 multicast traffic does not flow from the BVI to the EVI.

Example:

```
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)#bridge group bg1
Router(config-l2vpn-bg)#bridge-domain bd1
Router(config-l2vpn-bg-bd)#multicast stats-enable ipv4-ipv6
Router(config-l2vpn-bg-bd)#mld snooping profile test
Router(config-l2vpn-bg-bd)#igmp snooping profile test
Router(config-l2vpn-bg-bd)#interface bundle-Ether 100.1
Router(config-l2vpn-bg-bd-ac)#exit
Router(config-l2vpn-bg-bd)#interface Bundle-Ether 200.1
Router(config-l2vpn-bg-bd-ac)#exit
Router(config-l2vpn-bg-bd)#routed interface BVI1
Router(config-l2vpn-bg-bd-bvi)#split-horizon group core
Router(config-l2vpn-bg-bd-bvi)#exit
Router(config-l2vpn-bg-bd)#evi 1
Router(config-l2vpn-bg-bd-evi)#exit
Router(config-l2vpn-bg-bd)#exit
Router(config-l2vpn-bg)#bridge-domain bd2
Router(config-l2vpn-bg-bd)#igmp snooping profile test
Router(config-l2vpn-bg-bd)#interface Bundle-Ether 100.2
Router(config-l2vpn-bg-bd-ac)#exit
Router(config-l2vpn-bg-bd)#interface Bundle-Ether 200.2
Router(config-l2vpn-bg-bd-ac)#exit
Router(config-l2vpn-bg-bd)#routed interface BVI2
Router(config-l2vpn-bg-bd-bvi)#split-horizon group core
Router(config-l2vpn-bg-bd-bvi)#exit
Router(config-l2vpn-bg-bd)#evi 2
Router(config-l2vpn-bg-bd-evi)#commit
```

Step 6 Run the **bridge group** command to configure the bridge domain on PE1 and PE2.

Example:

```
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# bridge group 6005
Router(config-l2vpn-bg)# bridge-domain 6005
Router(config-l2vpn-bg-bd)# interface Bundle-Ether2.1
Router(config-l2vpn-bg-bd-ac)# evi 6005
Router(config-l2vpnbg-bd-evi)# commit
Router(config-l2vpnbg-bd-evi)# exit
```

Step 7 Run the **vrf** command to configure VRF.

Example:

```
Router# configure
Router(config)# vrf 30
Router(config-vrf)# address-family ipv4 unicast
Router(config-l2vpn-vrf-af)# route-target import 100:6005
Router(config-l2vpn-vrf-af)# route-target export 100:6005
Router(config-l2vpn-vrf-af)# commit
Router(config)#interface BVI1
Router(config-if)#vrf 30
Router(config-if)#ipv4 address 100.2.0.1 255.255.255.0
Router(config-if)#ipv6 address 100:2::1/64
Router(config-if)#mac-address 1002.1111.0002
```

EVPN IRB support

Ethernet Virtual Private Network Integrated Routing and Bridging (EVPN IRB) is a network technology that

- enables Layer 3 (L3) forwarding among hosts across different IP subnets and maintains the multi-homing capabilities of EVPN,
- allows EVPN hosts or subnets to communicate with IP VPNs, which enhances network flexibility and connectivity, and
- enables L3 forwarding among hosts across different IP subnets.

Supported EVPN IRB scenarios

EVPN IRB supports the following scenarios:

- In a single-homing scenario, only physical, VLAN, .1q, .1ad, or QinQ access methods are supported
- In a dual-homing scenario, only two PE gateways in a redundancy group are supported.

- Both IPv4 and IPv6 are supported.

EVPN IRB all-active multihoming

EVPN all-active multi-homing enables a CE device to connect to multiple PE devices, which provides redundant network connectivity and prevents disruptions. The Ethernet segment includes multiple Ethernet links that connect the CE to the PEs, often implemented as a Multi-chassis Link Aggregation Group (MC-LAG) bundle. This mode supports all-active redundancy, allowing the PEs to forward traffic simultaneously.

Distributed anycast gateway

Distributed anycast gateway in EVPN IRB is a routing and bridging feature that:

- enables multiple PE devices to act as gateways for the same IP subnet,
- uses a shared BVI IP and MAC address across the distributed PE devices, and
- optimizes forwarding by distributing routing capabilities closer to the access network, facilitating efficient inter or intra-subnet traffic flow.

EVPN IRB for a subnet is configured on all EVPN PEs hosting that subnet. To facilitate optimal routing and support transparent virtual machine mobility, hosts are configured with a single default gateway address for their local subnet. This anycast gateway address is configured with a single anycast MAC address on all local EVPN PE nodes supporting the subnet. Repeat this process for each locally defined subnet that requires Anycast Gateway support.

The host-to-host Layer 3 traffic, similar to Layer 3 VPN Provider Edge to Provider Edge forwarding, is routed on the source EVPN PE to the destination EVPN PE next-hop over an IP or MPLS tunnel, where it is routed again to the directly connected host. This forwarding approach is also known as Symmetric IRB, because Layer 3 flows are routed at both the source and destination EVPN PEs.

These solutions are part of the distributed anycast gateway feature:

EVPN IRB with Active-Active multihoming with subnet stretch or host-routing across fabric

For a bridge domain or subnet that is stretched across remote EVPN PEs, host routes with a `/32` prefix and MAC routes are distributed in an EVPN overlay control plane to enable Layer 2 and Layer 3 traffic to the endpoints in a stretched subnet.

This type of multihoming has the following characteristics:

- Layer 2 or Layer 3 ECMP for the fabric for dual-homed hosts based on Route Type 1 and Route Type 2.
- Layer 3 unipath over the fabric for single-homed hosts based on Route Type 2.
- Layer 2 subnet stretch over the fabric.
- Layer 2 stretch within redundancy group of leaves with orphan ports.