Bidirectional PIM or Bidirectional Multicast

Table 1: Feature History

Feature Name	Release Information	Description
Support of Bidirectional Multicast	Cisco IOS XE Cupertino 17.8.1	Bidirectional PIM or bidirectional multicast (RFC-5015) is an operating mode that enhances PIM by creating bidirectional multicast distribution trees. It helps deploy emerging communication and financial applications that rely on a many-to-many applications model. Use the following command to enable bidirectional PIM: **ip pim bidir-enable**

Bidirectional PIM is a variant of PIM Sparse mode that builds bidirectional multicast trees between sources and receivers without maintaining any source specific state along each node of the tree.

Prior to Cisco IOS XE Cupertino Release 17.8.1, PIM Sparse mode used to only provide loop-free, unidirectional multicast distribution trees.

Starting with Cisco IOS XE Cupertino Release 17.8.1, bidirectional PIM allows multicast devices to keep reduced state information, as compared with unidirectional shared trees in PIM Sparse mode. Bidirectional shared trees convey data from sources to the Rendezvous Point Address (RPA) and distribute them from the RPA to the receivers.

Note This feature is *only* supported on Cisco RSP3 module.

A single Designated Forwarder (DF) exists for each RPA on every link within a bidirectional PIM domain (including multiaccess and point-to-point links). The only exception is the RPL on which no DF exists. The DF is the device on the link with the best route to the RPA, which is determined by comparing Multicast Routing Information Base (MRIB)-provided metrics. A DF for a given RPA forwards downstream traffic onto its link and forwards upstream traffic from its link toward the Rendezvous Point link (RPL). The DF

performs this function for all bidirectional groups that map to the RPA. The DF on a link is also responsible for processing Join messages from downstream devices on the link as well as ensuring that packets are forwarded to local receivers discovered through a local membership mechanism such as IGMPv2.

Bidirectional PIM enables many-to-many applications applications by allowing them to easily scale to a very large number of groups and sources by eliminating the maintenance of source state.

- Advantages of Bidirectional PIM, on page 2
- Restrictions for Bidirectional PIM, on page 2
- Bidirectional Shared Tree, on page 3
- Designated Forwarder, on page 4
- Bidirectional Group Tree Building, on page 4
- Packet Forwarding, on page 5
- Configure Bidirectional PIM, on page 5
- Verification of Bidirectional PIM Configuration, on page 5
- Support for MVPN Bidirectional PIM, on page 7

Advantages of Bidirectional PIM

- Unlike PIM Sparse mode and PIM dense mode, the bidirectional PIM feature offers increased multicast scalability and reduced impact on router CPU and memory resources.
- The feature is well suited for many-to-many type of applications because the multicast routing state does *not* increase as the number of multicast sources increases and the amount of multicast state maintained by the router is reduced significantly compared with the other operating mode.
- Unlike PIM-SM, bidirectional PIM does not switch over to the source tree, and there is no register encapsulation of data from the source to the Rendezvous Point (RP).
- The feature removes the performance cost of maintaining a routing state table for a large number of sources.

Restrictions for Bidirectional PIM

- GRE based MVPN bidirectional PIM is *not* supported.
- MLDP-based MVPN bidirectional PIM is *not* be supported for Cisco IOS XE Cupertino Release 17.8.1.
- IPv6 traffic is *not* supported.
- Per OIF statistics is *not* supported for bidirectional PIM prefixes.
- Additional traffic might be received in some cases after you perform IP PIM bidirectional toggling.
- The maximum TCAM scale for the IPv4 Multicast routing table is 4K. This restriction on scalability is because of the hardware restriction of performing a single RPF check per TCAM entry. The number of DFs and RPAs on the router determine the total TCAM requirements. For a bidirectional prefix, the number of TCAM entries required is equal to “number of DF interfaces (for *, G/m prefix programming)” plus “the number of inherited accept interfaces for the corresponding *, G prefix).

Bidirectional Shared Tree

In bidirectional mode, traffic is routed only along a bidirectional shared tree that is rooted at the RP for the group. In bidir-PIM, the IP address of the RP acts as the key to having all routers establish a loop-free spanning tree topology rooted in that IP address. This IP address need not be a router, but can be any unassigned IP address on a network that is reachable throughout the PIM domain. This technique is the preferred configuration method for establishing a redundant RP configuration for bidir-PIM.

Membership in a bidirectional group is signaled by way of explicit Join messages. Traffic from sources is unconditionally sent up the shared tree toward the RP and passed down the tree toward the receivers on each branch of the tree.

The figures below show the difference in state created per router for a unidirectional shared tree and source tree versus a bidirectional shared tree.

Figure 1: Unidirectional Shared Tree and Source Tree

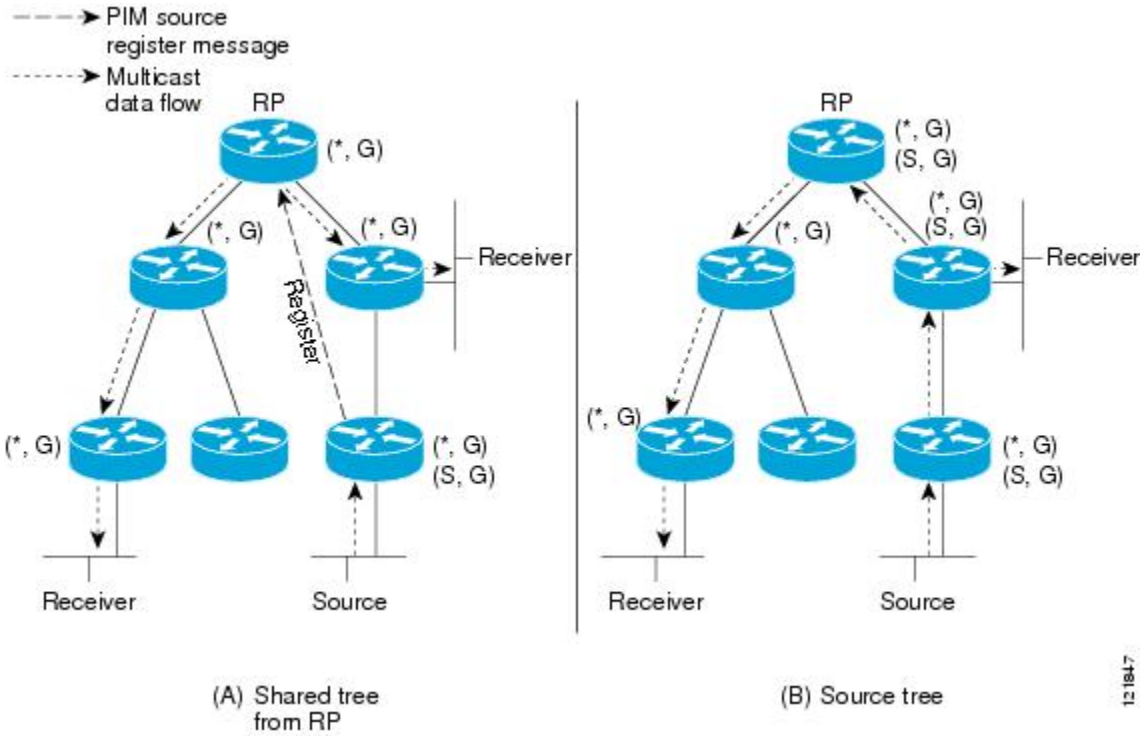

Figure 2: Bidirectional Shared Tree

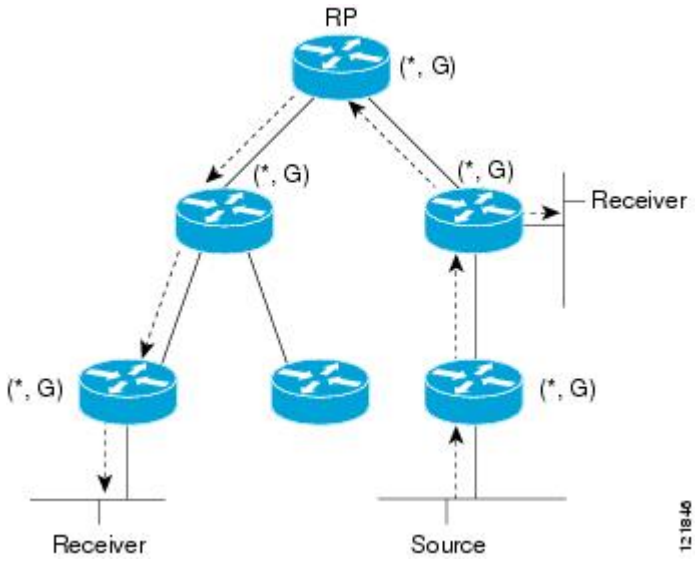

For packets that are forwarded downstream from the RP toward receivers, there are no fundamental differences between bidirectional PIM and PIM Sparse mode. Bidirectional PIM deviates substantially from PIM Sparse mode for traffic that is passed from sources upstream toward the RP.

The feature allows traffic to be passed up the shared tree toward the RP. To avoid multicast packet looping, bidirectional PIM introduces a new mechanism called Designated Forwarder (DF) election, which establishes a loop-free SPT rooted at the RP.

Designated Forwarder

On every network segment and point-to-point link, all PIM routers participate in a procedure called Designated Forwarder (DF) election. The procedure selects one router as the DF for every RP of bidirectional groups. This router is responsible for forwarding multicast packets received on that network.

The DF election is based on unicast routing metrics. The router with the most preferred unicast routing metric to the RP becomes the DF. Use of this method ensures that only one copy of every packet will be sent to the RP, even if there are parallel equal-cost paths to the RP.

A DF is selected for every RP of bidirectional groups. As a result, multiple routers may be elected as DF on any network segment, one for each RP. Any particular router may be elected as DF on more than one interface.

Bidirectional Group Tree Building

On a network that has local receivers, only the router elected as the DF populates the outgoing interface list (olist) upon receiving Internet Group Management Protocol (IGMP) Join messages, and sends (*, G) Join and Leave messages upstream toward the RP. When a downstream router wishes to join the shared tree, the RPF neighbor in the PIM Join and Leave messages is always the DF elected for the interface that lead to the RP.

When a router receives a Join or Leave message, and the router is not the DF for the receiving interface, the message is ignored. Otherwise, the router updates the shared tree in the same way as in sparse mode.

In a network where all routers support bidirectional shared trees, (S, G) Join and Leave messages are ignored. There is also no need to send PIM assert messages because the DF election procedure eliminates parallel downstream paths from any RP. An RP never joins a path back to the source, nor will it send any register stops.

Packet Forwarding

A router creates (*, G) entries only for bidirectional groups. The list of a (*, G) entry includes all the interfaces for which the router has been elected DF and that have received either an IGMP or PIM Join message. If a router is located on a sender-only branch, it will also create a (*, G) state, but the olist will not include any interfaces.

If a packet is received from the RPF interface toward the RP, the packet is forwarded downstream according to the olist of the (*, G) entry. Otherwise, only the router that is the DF for the receiving interface forwards the packet upstream toward the RP; all other routers must discard the packet.

Configure Bidirectional PIM

To configure bidirectional PIM:

```
ip multicast-routing distributed

interface Loopback0
 ip address 2.2.2.2 255.255.255.255
 ip ospf 100 area 0


interface GigabitEthernet0/2/1
 ip address 30.1.1.3 255.255.255.0
 ip pim sparse-mode
 ip ospf 100 area 0
 negotiation auto


interface GigabitEthernet0/2/4
 ip address 100.2.1.2 255.255.255.0
 ip pim sparse-mode
 ip ospf 100 area 0
 negotiation auto

router ospf 100
 router-id 2.2.2.2

ip pim bidir-enable
ip pim rp-address 20.1.1.10 BIDIR-GROUPS bidir

ip access-list standard BIDIR-GROUPS
 20 permit 225.0.0.0 0.255.255.255
```

Verification of Bidirectional PIM Configuration

Use the **show ip mroute** command to display the IP multicast routing table details.

```
Router#show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*,225.0.0.0/8), 00:00:54/-, RP 20.1.1.10, flags: B
  Bidir-Upstream: GigabitEthernet0/2/1, RPF nbr: 30.1.1.2
  Incoming interface list:
    GigabitEthernet0/2/4, Accepting/Sparse
    GigabitEthernet0/2/1, Accepting/Sparse

(*, 225.1.1.1), 00:09:01/00:02:59, RP 20.1.1.10, flags: BC
  Bidir-Upstream: GigabitEthernet0/2/1, RPF nbr 30.1.1.2
  Outgoing interface list:
    GigabitEthernet0/2/4, Forward/Sparse, 00:09:01/00:02:59, flags:
    GigabitEthernet0/2/1, Bidir-Upstream/Sparse, 00:09:01/stopped, flags:

(*, 224.0.1.40), 00:14:12/00:02:05, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet0/2/1, Forward/Sparse, 00:14:12/00:02:05, flags:
```

Use the **show ip mfib** command to display the forwarding entries and interfaces in the IPv4 Multicast Forwarding Information Base (MFIB).

```
Router#show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet0/2/1 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet0/2/4 Flags: A
   GigabitEthernet0/2/1 Flags: A F
     Pkts: 0/0/0    Rate: 0 pps
   Null0 Flags: A
 (*,225.1.1.1) Flags: IA HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   530500/1000/64/500, Other: 0/0/0
   GigabitEthernet0/2/4 Flags: F
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet0/2/1 Flags: F
     Pkts: 0/0/0    Rate: 0 pps
```

Use the **show ip pim neigh** command to display the PIM neighbor table details.

```
Router#show ip pim neigh
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface               Uptime/Expires    Ver   DR
Address                                                           Prio/Mode
30.1.1.2          GigabitEthernet0/2/1    00:18:19/00:01:38 v2    1 / B S P G
```

Use the **show ip pim int df** command to display the DF interfaces.

```
Router#show ip pim int df
* implies this system is the DF
Interface               RP              DF Winner        Metric     Uptime
GigabitEthernet0/2/1    20.1.1.10        30.1.1.2         0          00:17:24
GigabitEthernet0/2/4    20.1.1.10       *100.2.1.2        2          00:17:18
```

Support for MVPN Bidirectional PIM

Table 2: Feature History

Feature Name	Release Information	Description
Support for MVPN Bidirectional PIM	Cisco IOS XE Cupertino 17.9.1	This release extends the support of bidirectional PIM over MVPN. This feature is only supported on profile 1 MVPN or default MDT - MLDP MP2MP - PIM C-mcast signaling. This feature is only supported on Cisco RSP3 module. This feature is only supported on NCS 4206 and NCS 4216 routers.

Starting with Cisco IOS XE Cupertino Release 17.9.1, bidirectional PIM is supported over Multicast Label Distribution Protocol (MLDP)-based MVPN. MLDP helps to transport traffic with label encapsulation over an MPLS core network.

An MVPN profile is configured for the global context or per Virtual Routing/Forwarding (VRF). This feature is only supported on profile 1 MVPN or default MDT - MLDP MP2MP - PIM C-mcast signaling.For more information on MLDP-based MVPN configuration, see MLDP-Based MVPN.

Note Data MDTs can only be set up for (S,G) (Shortest Path Tree), and not for (*,G) (Shared Tree).

Benefits of MVPN Bidirectional PIM

- Enables the use of a single MPLS forwarding plane for both unicast and multicast traffic.
- Enables existing MPLS protection (for example, MPLS Traffic Engineering/Resource Reservation Protocol (TE/RSVP link protection) and MPLS Operations Administration and Maintenance (OAM) mechanisms to be used for multicast traffic.

Restrictions for MVPN Bidirectional PIM

This feature is only supported on MVPN profile 1.

Prerequisites for MVPN Bidirectional PIM

- You should be familiar with bidirectional PIM configurations.
- You should configure basic native multicast configurations including PIM configurations.

Configure MVPN Bidirectional PIM

To configure MVPN bidirectional PIM:

1. **Enable Bidirectional PIM for VRF**:

```
ip pim vrf cu1 rp-address 20.1.1.10 BIDIR-GROUPS bidir
vrf definition cu1
 rd 1:1
 vpn id 1:1
 !
 address-family ipv4
  mdt default mpls mldp 2.2.2.2
  route-target export 1:1
  route-target import 1:1
 exit-address-family
no mpls mldp forwarding recursive
ip multicast-routing vrf cu1 distributed
interface Loopback0
 ip address 5.5.5.5 255.255.255.255
 ip ospf 1 area 0
!
interface Loopback1
 vrf forwarding cu1
 ip address 55.55.55.55 255.255.255.255

ip access-list standard BIDIR-GROUPS
```

```
 10 permit 226.1.1.1
 20 permit 225.0.0.0 0.255.255.255

interface GigabitEthernet0/3/4
 vrf forwarding cu1
 ip address 100.51.1.2 255.255.255.0
 ip pim sparse-mode
 ip ospf 100 area 0
 negotiation auto
router bgp 100
address-family ipv4 vrf cu1
  redistribute connected
  redistribute ospf 100
 exit-address-family
ip pim bidir-enable
```

- **Enable Static RP with ACL**:

```
ip access-list standard BIDIR-GROUPS
permit 226.1.1.1
permit 225.0.0.0 0.255.255.255
```

Verification for MVPN Bidirectional PIM Configuration

Use the **show ip pim vrf** *cu1* **interface** *df* command to verify the Designated Forwarder (DF) configuration.

```
Router#show ip pim vrf cu1 interface df
Interface                RP               DF Winner        Metric     Uptime
GigabitEthernet0/2/5     20.1.1.10         0.0.0.0          0          00:00:00
Lspvif1                  20.1.1.10        *10.0.0.1          0          2d10h
```

Use the **show ip pim vrf** *cu1* **rp mapping** command to verify PIM group to RP mapping configuration.

```
Router#show ip pim vrf cu1 rp mapping

Acl: SPARSE-GROUPS, Static
    RP: 11.11.11.11 (?)
Acl: BIDIR-GROUPS, Static, Bidir Mode
    RP: 20.1.1.10 (?)
ASR907_2039#
```

Use the **show ip mroute vrf** *cu1* command to verify bidirectional PIM configuration for VRF.

```
Router#show ip mroute vrf cu1
(*,226.0.0.0/8), 2d11h/-, RP 20.1.1.10, flags: B
  Bidir-Upstream: GigabitEthernet0/2/5, RPF nbr: 20.1.1.10
  Incoming interface list:
    Lspvif1, Accepting/Sparse
    GigabitEthernet0/2/5, Accepting/Sparse
(*,225.0.0.0/8), 2d11h/-, RP 20.1.1.10, flags: B
  Bidir-Upstream: GigabitEthernet0/2/5, RPF nbr: 20.1.1.10
  Incoming interface list:
    Lspvif1, Accepting/Sparse
    GigabitEthernet0/2/5, Accepting/Sparse

(*, 224.0.1.40), 2d11h/00:02:08, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet0/2/5, Forward/Sparse, 2d11h/00:02:08, flags:
Router#
```