Implementing Multicast

Implementing Layer-3 Multicast Routing

Multicast routing allows a host to send packets to a subset of all hosts as a group transmission rather than to a single host, as in unicast transmission, or to all hosts, as in broadcast transmission. The subset of hosts is known as group members and are identified by a single multicast group address that falls under the IP Class D address range from 224.0.0.0 through 239.255.255.255.

The multicast environment consists of senders and receivers. Any host, regardless of whether it is a member of a group, can send to a group. However, only the members of a group receive the message.

The following protocols are supported to implement multicast routing:

- IGMP—IGMP is used between hosts on a network (for example, LAN) and the routers on that network to track the multicast groups of which hosts are members.
- PIM SSM— Protocol Independent Multicast in Source-Specific Multicast (PIM-SSM) has the ability to report interest in receiving packets from specific source addresses (or from all but the specific source addresses), to an IP multicast address.

Note MLD Snooping is not supported until Cisco IOS XR Release 6.5.3. The support will be available in future releases.

Prerequisites for Implementing Multicast Routing

- You must install and activate the multicast RPM package.
- You must be familiar with IPv4 multicast routing configuration tasks and concepts.
- Unicast routing must be operational.

Enabling Multicast

Configuration Example

Enables multicast routing and forwarding on all new and existing interfaces.

```
Router#config
Router(config)#multicast-routing
Router(config-mcast)#address-family ipv4
Router(config-mcast-default-ipv4)#interface all enable
*/In the above command, you can also indicate a specific interface (For example, interface
 TenGigE0/0/0/3)
for enabling multicast only on that interface/*
Router(config-mcast-default-ipv4)#commit
```

Running Configuration

```
Router#show running multicast routing
multicast-routing
 address-family ipv4
  interface all enable
  !
```

Verification

Verify that the Interfaces are enabled for multicast.

```
Router#show mfib interface location 0/3/CPU0
Interface : FINT0/3/CPU0 (Enabled)
SW Mcast pkts in : 0, SW Mcast pkts out : 0
TTL Threshold : 0
Ref Count : 2
Interface : TenGigE0/3/0/0/0 (Enabled)
SW Mcast pkts in : 0, SW Mcast pkts out : 0
TTL Threshold : 0
Ref Count : 3
Interface : TenGigE0/3/0/9/0 (Enabled)
SW Mcast pkts in : 0, SW Mcast pkts out : 0
TTL Threshold : 0
Ref Count : 13
Interface : Bundle-Ether1 (Enabled)
SW Mcast pkts in : 0, SW Mcast pkts out : 0
TTL Threshold : 0
Ref Count : 4
Interface : Bundle-Ether1.1 (Enabled)
SW Mcast pkts in : 0, SW Mcast pkts out : 0
TTL Threshold : 0
```

Protocol Independent Multicast

Protocol Independent Multicast (PIM) is a multicast routing protocol used to create multicast distribution trees, which are used to forward multicast data packets.

Proper operation of multicast depends on knowing the unicast paths towards a source or an RP. PIM relies on unicast routing protocols to derive this reverse-path forwarding (RPF) information. As the name PIM implies, it functions independently of the unicast protocols being used. PIM relies on the Routing Information Base (RIB) for RPF information. Protocol Independent Multicast (PIM) is designed to send and receive multicast routing updates.

NCS 5500 series router supports Protocol Independent Multicast in Source-Specific Multicast (PIM-SSM).

PIM on Bundle-Ethernet subinterface is supported.

PIM BFD Overview

The BFD Support for Multicast (PIM) feature, also known as PIM BFD, registers PIM as a client of BFD. PIM can then utilize BFD's fast adjacency failure detection. When PIM BFD is enabled, BFD enables faster failure detection without waiting for hello messages from PIM.

At PIMs request, as a BFD client, BFD establishes and maintains a session with an adjacent node for maintaining liveness and detecting forwarding path failure to the adjacent node. PIM hellos will continue to be exchanged between the neighbors even after BFD establishes and maintains a BFD session with the neighbor. The behavior of the PIM hello mechanism is not altered due to the introduction of this feature. Although PIM depends on the Interior Gateway Protocol (IGP) and BFD is supported in IGP, PIM BFD is independent of IGP's BFD.

Protocol Independent Multicast (PIM) uses a hello mechanism for discovering new PIM neighbors between adjacent nodes. The minimum failure detection time in PIM is 3 times the PIM Query-Interval. To enable faster failure detection, the rate at which a PIM hello message is transmitted on an interface is configurable. However, lower intervals increase the load on the protocol and can increase CPU and memory utilization and cause a system-wide negative impact on performance. Lower intervals can also cause PIM neighbors to expire frequently as the neighbor expiry can occur before the hello messages received from those neighbors are processed. When PIM BFD is enabled, BFD enables faster failure detection without waiting for hello messages from PIM.

Configure PIM BFD

This section describes how you can configure PIM BFD

```
Router# configure
Router(config)# router pim address-family ipv4
Router(config-pim-default-ipv4)# interface HundredGige0/1/0/1
Router(config-pim-ipv4-if)#  bfd minimum-interval 10
Router(config-pim-ipv4-if)#  bfd fast-detect
Router(config-pim-ipv4-if)#  bfd multiplier 3
Router(config-pim-ipv4)#  exit
Router(config-pim-default-ipv4)# interface TengigabitEthernet0/0/0/4
Router(config-pim-ipv4-if)#  bfd minimum-interval 50
Router(config-pim-ipv4-if)#  bfd fast-detect
Router(config-pim-ipv4-if)#  bfd multiplier 3
Router(config-pim-ipv4-if)#  exit
```

```
Router(config-pim-default-ipv4)# interface TengigabitEthernet 0/0/0/4.101
Router(config-pim-ipv4-if)#  bfd minimum-interval 50
Router(config-pim-ipv4-if)#  bfd fast-detect
Router(config-pim-ipv4-if)#  bfd multiplier 3
Router(config-pim-ipv4-if)#  exit
Router(config-pim-default-ipv4)# interface  Bundle-Ether 101
Router(config-pim-ipv4-if)#  bfd minimum-interval 50
Router(config-pim-ipv4-if)#  bfd fast-detect
Router(config-pim-ipv4-if)#  bfd multiplier 3
Router(config-pim-ipv4-if)#  exit
Router(config-pim-default-ipv4)#  commit
```

Running Configuration

```
router pim
 address-family ipv4
  interface HundredGige 0/1/0/1
   bfd minimum-interval 10
   bfd fast-detect
   bfd multiplier 3
  !
  interface TengigabitEthernet 0/0/0/4
   bfd minimum-interval 50
   bfd fast-detect
   bfd multiplier 3
  !
  interface TengigabitEthernet 0/0/0/4.101
   bfd minimum-interval 50
   bfd fast-detect
   bfd multiplier 3
  !
  interface Bundle-Ether 101
   bfd minimum-interval 50
   bfd fast-detect
   bfd multiplier 3
  !

  !
 !
!
```

Verification

The show outputs given in the following section display the details of the configuration of the PIM BFD, and the status of their configuration.

```
Router# show bfd session
Wed Nov 22 08:27:35.952 PST
Interface           Dest Addr    Local det time(int*mult)  State      Echo       Async
H/W     NPU
------------------- --------------- ---------------- ---------------- -----      -----
-----    -----
Hu0/0/1/3           10.12.12.2      0s(0s*0) 90ms(30ms*3)     UP                  Yes
0/0/CPU0

Hu0/0/1/2           10.12.12.2      0s(0s*0) 90ms(30ms*3)     UP                  Yes
0/0/CPU0

Hu0/0/1/1           10.18.18.2      0s(0s*0) 90ms(30ms*3)     UP                  Yes
0/0/CPU0
```

```
Te0/0/0/4.101       10.112.112.2   0s(0s*0) 90ms(30ms*3)     UP                  Yes
0/0/CPU0

BE101               10.18.18.2      n/a      n/a              UP                  No    n/a


BE102               10.12.12.2      n/a      n/a              UP                  No    n/a



Router# show bfd client

Name            Node       Num sessions
--------------- ---------- --------------
L2VPN_ATOM      0/RP0/CPU0 0
MPLS-TR         0/RP0/CPU0 0
bgp-default     0/RP0/CPU0 0
bundlemgr_distrib 0/RP0/CPU0 14
isis-1          0/RP0/CPU0 0
object_tracking 0/RP0/CPU0 0
pim6            0/RP0/CPU0 0
pim             0/RP0/CPU0 0
service-layer   0/RP0/CPU0 0
```

Reverse Path Forwarding

Reverse-path forwarding (RPF) is an algorithm used for forwarding multicast datagrams. It functions as follows:

- If a router receives a datagram on an interface it uses to send unicast packets to the source, the packet has arrived on the RPF interface.
- If the packet arrives on the RPF interface, a router forwards the packet out the interfaces present in the outgoing interface list of a multicast routing table entry.
- If the packet does not arrive on the RPF interface, the packet is silently discarded to prevent loops.

PIM uses both source trees and RP-rooted shared trees to forward datagrams; the RPF check is performed differently for each, as follows:

- If a PIM router has an (S,G) entry present in the multicast routing table (a source-tree state), the router performs the RPF check against the IP address of the source for the multicast packet.
- If a PIM router has no explicit source-tree state, this is considered a shared-tree state. The router performs the RPF check on the address of the RP, which is known when members join the group.

Sparse-mode PIM uses the RPF lookup function to determine where it needs to send joins and prunes. (S,G) joins (which are source-tree states) are sent toward the source. (*,G) joins (which are shared-tree states) are sent toward the RP.

Setting the Reverse Path Forwarding Statically

Configuration Example

The following example configures the static RPF rule for IP address 10.0.0.1:

```
Router#configure
Router(config)#multicast-routing
Router(config-if)#static-rpf 10.0.0.1 32 TenGigE 0/0/0/1 192.168.0.2
Router(config-ipv4-acl)#commit
```

Running Configuration

```
multicast-routing
   address-family ipv4
      static-rpf 10.10.10.2 32 TenGigE0/0/0/1 192.168.0.2
```

Verification

Verify that RPF is chosen according to the static RPF configuration for 10.10.10.2

```
Router#show pim rpf
Table: IPv4-Unicast-default
* 10.10.10.2/32 [0/0]
    via GigabitEthernet0/0/0/1 with rpf neighbor 192.168.0.2
```

PIM Bootstrap Router

The PIM bootstrap router (BSR) provides a fault-tolerant, automated RP discovery and distribution mechanism that simplifies the Auto-RP process. This feature is enabled by default allowing routers to dynamically learn the group-to-RP mappings.

PIM uses the BSR to discover and announce RP-set information for each group prefix to all the routers in a PIM domain. This is the same function accomplished by Auto-RP, but the BSR is part of the PIM specification. The BSR mechanism interoperates with Auto-RP on Cisco routers.

To avoid a single point of failure, you can configure several candidate BSRs in a PIM domain. A BSR is elected among the candidate BSRs automatically.

Candidates use bootstrap messages to discover which BSR has the highest priority. The candidate with the highest priority sends an announcement to all PIM routers in the PIM domain that it is the BSR.

Routers that are configured as candidate RPs unicast to the BSR the group range for which they are responsible. The BSR includes this information in its bootstrap messages and disseminates it to all PIM routers in the domain. Based on this information, all routers are able to map multicast groups to specific RPs. As long as a router is receiving the bootstrap message, it has a current RP map.

Configuring PIM Bootstrap Router

Configuration Example

Configures the router as a candidate BSR with a hash mask length of 30:

```
Router#config
Router(config)#router pim
Router(config-pim-default-ipv4)#bsr candidate-bsr 1.1.1.1 hash-mask-len 30 priority 1
Router(config-pim-default-ipv4-if)#commit
```

Configures the router to advertise itself as a candidate rendezvous point to the BSR in its PIM domain. Access list number 4 specifies the prefix associated with the candidate rendezvous point address 1.1.1.1 . This rendezvous point is responsible for the groups with the prefix 239.

```
Router#config
Router(config)#router pim
Router(config-pim-default-ipv4)#bsr candidate-rp 1.1.1.1 group-list 4 priority 192 interval
 60

Router(config-pim-default-ipv4)#exit
Router(config)#ipv4 access-list 4
Router(config-ipv4-acl)#permit ipv4 any 239.0.0.0 0.255.255.255
Router(config-ipv4-acl)#commit
```

Running Configuration

```
Router#show run router pim
router pim
 address-family ipv4
  bsr candidate-bsr 1.1.1.1 hash-mask-len 30 priority 1
  bsr candidate-rp 1.1.1.1 group-list 4 priority 192 interval 60
```

Verification

```
Router#show pim rp mapping
PIM Group-to-RP Mappings
Group(s) 239.0.0.0/8
  RP 1.1.1.1 (?), v2
    Info source: 1.1.1.1 (?), elected via bsr, priority 192, holdtime 150
      Uptime: 00:02:50, expires: 00:01:54

Router#show pim bsr candidate-rp
PIM BSR Candidate RP Info
Cand-RP          mode  scope priority uptime    group-list
1.1.1.1           BD    16      192       00:04:06    4

Router#show pim bsr election
PIM BSR Election State
Cand/Elect-State          Uptime   BS-Timer     BSR                                  C-BSR
Elected/Accept-Pref     00:03:49 00:00:25 1.1.1.1 [1, 30]          1.1.1.1 [1, 30]
```

PIM-Source Specific Multicast

When PIM is used in SSM mode, multicast routing is easier to manage. This is because RPs (rendezvous points) are not required and therefore, no shared trees (*,G) are built.

There is no specific IETF document defining PIM-SSM. However, RFC4607 defines the overall SSM behavior.

In the rest of this document, we use the term PIM-SSM to describe PIM behavior and configuration when SSM is used.

PIM in Source-Specific Multicast operation uses information found on source addresses for a multicast group provided by receivers and performs source filtering on traffic.

- By default, PIM-SSM operates in the 232.0.0.0/8 multicast group range for IPv4 . To configure these values, use the **ssm range** command.
- If SSM is deployed in a network already configured for PIM-SM, only the last-hop routers must be upgraded with Cisco IOS XR Software that supports the SSM feature.
- No MSDP SA messages within the SSM range are accepted, generated, or forwarded.
- SSM can be disabled using the **ssm disable**command.

- The **ssm allow-override** command allows SSM ranges to be overridden by more specific ranges.

In many multicast deployments where the source is known, protocol-independent multicast-source-specific multicast (PIM-SSM) mapping is the obvious multicast routing protocol choice to use because of its simplicity. Typical multicast deployments that benefit from PIM-SSM consist of entertainment-type solutions like the ETTH space, or financial deployments that completely rely on static forwarding.

In SSM, delivery of data grams is based on (S,G) channels. Traffic for one (S,G) channel consists of datagrams with an IP unicast source address S and the multicast group address G as the IP destination address. Systems receive traffic by becoming members of the (S,G) channel. Signaling is not required, but receivers must subscribe or unsubscribe to (S,G) channels to receive or not receive traffic from specific sources. Channel subscription signaling uses IGMP to include mode membership reports, which are supported only in Version 3 of IGMP (IGMPv3).

To run SSM with IGMPv3, SSM must be supported on the multicast router, the host where the application is running, and the application itself. Cisco IOS XR Software allows SSM configuration for an arbitrary subset of the IP multicast address range 224.0.0.0 through 239.255.255.255.

When an SSM range is defined, existing IP multicast receiver applications do not receive any traffic when they try to use addresses in the SSM range, unless the application is modified to use explicit (S,G) channel subscription.

Benefits of PIM-SSM over PIM-SM

PIM-SSM is derived from PIM-SM. However, whereas PIM-SM allows for the data transmission of all sources sending to a particular group in response to PIM join messages, the SSM feature forwards traffic to receivers only from those sources that the receivers have explicitly joined. Because PIM joins and prunes are sent directly towards the source sending traffic, an RP and shared trees are unnecessary and are disallowed. SSM is used to optimize bandwidth utilization and deny unwanted Internet broad cast traffic. The source is provided by interested receivers through IGMPv3 membership reports.

IGMPv2

To support IGMPv2, SSM mapping configuration must be added while configuring IGMP to match certain sources to group range.

Configuring Example

Configures the access-list (mc1):

```
Router#configure
Router(config)#ipv4 access-list mc1
Router(config-ipv4-acl)#permit ipv4 any 232.1.1.0 0.0.0.255
Router(config-ipv4-acl)#commit
```

Configures the multicast source (1.1.1.1) as part of a set of sources that map SSM groups described by the specified access-list (mc1):

```
Router#configure
Router(config)#router igmp
Router(config-igmp)#ssm map static 1.1.1.1 mc1
Router(config-igmp)#commit
```

Running Configuration

```
Router#show run router igmp
router igmp
ssm map static 1.1.1.1 mc1
```

Verification

Multipath Option

The multipath option is available under `router pim` configuration mode. After multipath option is enabled, SSM selects different path to reach same destination instead of choosing common path. The multipath option helps load balance the SSM traffic.

Configuring Multipath Option

```
Router#configure
Router(config)#router pim address-family ipv4
Router(config-pim-default-ipv4)#multipath hash source
Router(config-pim-default-ipv4)#commit
```

Running Configuration

```
Router#show running router pim
router pim
 address-family ipv4
  dr-priority 100
  multipath hash source    /*SSM traffic takes different path to reach same destination
based on source hash value.*/
```

Verification

The Bundle-Ether132 and TenGigE0/4/0/18/0.132 are two paths to reach the destination router Turnin-56. Since we have enabled multipath option, the source has two IP addresses 50.11.30.12 and 50.11.30.11. The Multicast traffic from two sources take two different paths Bundle-Ether132 and TenGigE0/4/0/18/0.132 to reach same destination.

This show run output shows that Bundle-Ether132 and TenGigE0/4/0/18/0.132 are connected to same destination router Turnin-56:

```
Router#show run int TenGigE0/1/0/6/3.132
interface TenGigE0/1/0/6/3.132
 description Connected to Turin-56 ten0/0/0/19.132
 ipv4 address 13.0.2.1 255.255.255.240
 ipv6 address 2606::13:0:2:1/120
 encapsulation dot1q 132
!
```

```
Router#show run int be132
interface Bundle-Ether132
 description Bundle between Fretta-56 and Turin-56
 ipv4 address 28.0.0.1 255.255.255.240
 ipv6 address 2606::28:0:0:1/120
 load-interval 30
```

```
Router#show mrib route 50.11.30.11 detail

IP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
    C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
```

```
    IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
    MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
    CD - Conditional Decap, MPLS - MPLS Decap, EX - Extranet
    MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
    MoFB - MoFRR Backup, RPFID - RPF ID Set, X - VXLAN
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
    NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
    II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
    LD - Local Disinterest, DI - Decapsulation Interface
    EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
    EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
    MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
    IRMI - IR MDT Interface

(50.11.30.11,225.255.11.1) Ver: 0x523cc294 RPF nbr: 50.11.30.11 Flags: L RPF, FGID: 11453,
 -1, -1
  Up: 4d15h
  Incoming Interface List
    HundredGigE0/4/0/10.1130 Flags: A, Up: 4d15h
  Outgoing Interface List
    FortyGigE0/1/0/5 Flags: F NS, Up: 4d15h
    TenGigE0/4/0/6/0 Flags: F NS, Up: 4d15h
    TenGigE0/1/0/6/3.132 Flags: F NS, Up: 4d15h
    TenGigE0/4/0/18/0.122 Flags: F NS, Up: 4d15h

Router#show mrib route 50.11.30.12 detail

IP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
    C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
    IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
    MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
    CD - Conditional Decap, MPLS - MPLS Decap, EX - Extranet
    MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
    MoFB - MoFRR Backup, RPFID - RPF ID Set, X - VXLAN
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
    NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
    II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
    LD - Local Disinterest, DI - Decapsulation Interface
    EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
    EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
    MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
    IRMI - IR MDT Interface

(50.11.30.12,226.255.12.1) Ver: 0x5fe02e5b RPF nbr: 50.11.30.12 Flags: L RPF, FGID: 12686,
 -1, -1
  Up: 4d15h
  Incoming Interface List
    HundredGigE0/4/0/10.1130 Flags: A, Up: 4d15h
  Outgoing Interface List
    Bundle-Ether121 Flags: F NS, Up: 4d15h
    Bundle-Ether132 Flags: F NS, Up: 4d15h
    FortyGigE0/1/0/5 Flags: F NS, Up: 4d15h
    TenGigE0/4/0/6/0.117 Flags: F NS, Up: 4d15h
```

Configuring PIM-SSM

Configuration Example

Configures SSM service for the IPv4 address range defined by access list 4.

```
Router#config
Router(config)#ipv4 access-list 4
Router(config-ipv4-acl)#permit ipv4 any 224.2.151.0 0.0.0.255
Router(config-ipv4-acl)#exit
Router(config)#multicast-routing
Router(config-mcast)#address-family ipv4
Router(config-mcast-default-ipv4)#ssm range 4
Router(config-mcast-default-ipv4)#commit
Router(config-mcast-default-ipv4)#end
```

Running Configuration

```
Router#show running multicast-routing
multicast-routing
 address-family ipv4
  ssm range 4
  interface all enable
 !
```

Verification

Verify if the SSM range is configured according to the set parameters:

```
Router#show access-lists 4
ipv4 access-list 4
 10 permit ipv4 any 224.2.151.0 0.0.0.255

*/Verify if the SSM is configured for 224.2.151.0/24/*:

Router#show pim group-map
IP PIM Group Mapping Table
(* indicates group mappings being used)
Group Range          Proto Client   Groups RP address      Info
224.0.1.39/32*       DM    perm     1      0.0.0.0
224.0.1.40/32*       DM    perm     1      0.0.0.0
224.0.0.0/24*        NO    perm     0      0.0.0.0
224.2.151.0/24*      SSM   config   0      0.0.0.0
```

Designated Routers

Cisco routers use PIM-SM to forward multicast traffic and follow an election process to select a designated router (DR) when there is more than one router on a LAN segment.

The designated router is responsible for sending PIM register and PIM join and prune messages toward the RP to inform it about host group membership.

If there are multiple PIM-SM routers on a LAN, a designated router must be elected to avoid duplicating multicast traffic for connected hosts. The PIM router with the highest IP address becomes the DR for the LAN unless you choose to force the DR election by use of the **dr-priority** command. The DR priority option allows you to specify the DR priority of each router on the LAN segment (default priority = 1) so that the router with the highest priority is elected as the DR. If all routers on the LAN segment have the same priority, the highest IP address is again used as the tiebreaker.

Note DR election process is required only on multi access LANs. The last-hop router directly connected to the host is the DR.

The figure "Designated Router Election on a Multiaccess Segment", below illustrates what happens on a multi access segment. Router A (10.0.0.253) and Router B (10.0.0.251) are connected to a common multi access Ethernet segment with Host A (10.0.0.1) as an active receiver for Group A. As the Explicit Join model is used, only Router A, operating as the DR, sends joins to the RP to construct the shared tree for Group A. If Router B were also permitted to send (*,G) joins to the RP, parallel paths would be created and Host A would receive duplicate multicast traffic. When Host A begins to source multicast traffic to the group, the DR's responsibility is to send register messages to the RP. Again, if both routers were assigned the responsibility, the RP would receive duplicate multicast packets.

Figure 1: Designated Router Election on a Multiaccess Segment

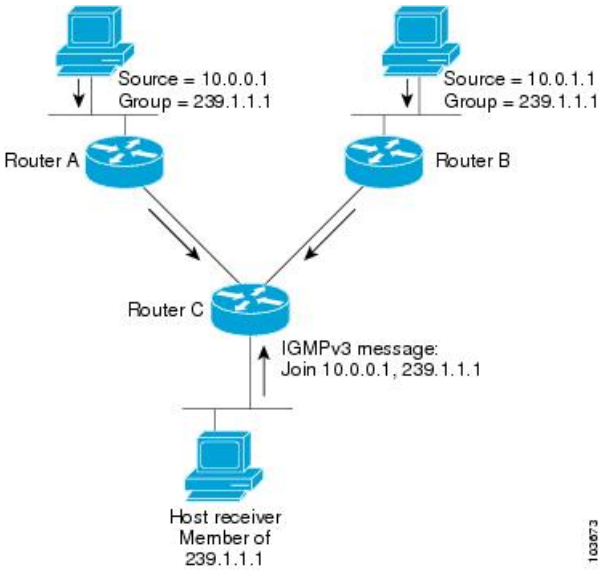

If the DR fails, the PIM-SM provides a way to detect the failure of Router A and to elect a failover DR. If the DR (Router A) were to become inoperable, Router B would detect this situation when its neighbor adjacency with Router A timed out. Because Router B has been hearing IGMP membership reports from Host A, it already has IGMP state for Group A on this interface and immediately sends a join to the RP when it becomes the new DR. This step reestablishes traffic flow down a new branch of the shared tree using Router B. Additionally, if Host A were sourcing traffic, Router B would initiate a new register process immediately after receiving the next multicast packet from Host A. This action would trigger the RP to join the SPT to Host A, using a new branch through Router B.

Note Two PIM routers are neighbors if there is a direct connection between them. To display your PIM neighbors, use the show pim neighbor command in EXEC mode.

- They are not used for unicast routing but are used only by PIM to look up an IPv4 next hop to a PIM source.
- They are not published to the Forwarding Information Base (FIB).
- When multicast-intact is enabled on an IGP, all IPv4 destinations that were learned through link-state advertisements are published with a set equal-cost mcast-intact next-hops to the RIB. This attribute applies even when the native next-hops have no IGP shortcuts.

- In IS-IS, the max-paths limit is applied by counting both the native and mcast-intact next-hops together. (In OSPFv2, the behavior is slightly different.)

Configuration Example

Configures the router to use DR priority 4 for TenGigE interface 0/0/0/1, but other interfaces will inherit DR priority 2:

```
Router#configure
Router(config)#router pim
Router(config-pim-default)#address-family ipv4
Router(config-pim-default-ipv4)#dr-priority 2
Router(config-pim-default-ipv4)#interface TenGigE0/0/0/1
Router(config-pim-ipv4-if)#dr-priority 4
Router(config-ipv4-acl)#commit
```

Running Configuration

```
Router#show run router pim
router pim
 address-family ipv4
  dr-priority 2
  spt-threshold infinity
  interface TenGigE0/0/0/1
   dr-priority 4
   hello-interval 45
```

Verification

Verify if the parameters are set according to the configured values:

```
Router#show pim interface
PIM interfaces in VRF default
Address           Interface           PIM  Nbr   Hello  DR    DR Count Intvl  Prior
100.1.1.1        TenGigE0/0/0/1        on   1     45     4     this system
26.1.1.1         TenGigE0/0/0/26       on   1     30     2     this system
```

Internet Group Management Protocol

Cisco IOS XR Software provides support for Internet Group Management Protocol (IGMP) over IPv4.

IGMP provides a means for hosts to indicate which multicast traffic they are interested in and for routers to control and limit the flow of multicast traffic throughout the network. Routers build state by means of IGMP messages; that is, router queries and host reports.

A set of routers and hosts that receive multicast data streams from the same source is called a multicast group. Hosts use IGMP messages to join and leave multicast groups.

Note IGMP messages use group addresses, which are Class D IP addresses. The high-order four bits of a Class D address are 1110. Host group addresses can be in the range 224.0.0.0 to 239.255.255.255. The address is guaranteed not to be assigned to any group. The address 224.0.0.1 is assigned to all systems on a subnet. The address 224.0.0.2 is assigned to all routers on a subnet.

NCS 5500 supports IGMPv3 by default. No configuration is required. IGMP Version 3 permits joins and leaves for certain source and group pairs instead of requesting traffic from all sources in the multicast group.

Functioning of IGMP Routing

The following image "IGMP Singaling" , illustrates two sources, 10.0.0.1 and 10.0.1.1, that are multicasting to group 239.1.1.1.

The receiver wants to receive traffic addressed to group 239.1.1.1 from source 10.0.0.1 but not from source 10.0.1.1.

The host must send an IGMPv3 message containing a list of sources and groups (S, G) that it wants to join and a list of sources and groups (S, G) that it wants to leave. Router C can now use this information to prune traffic from Source 10.0.1.1 so that only Source 10.0.0.1 traffic is being delivered to Router C.

Figure 2: IGMP Signaling

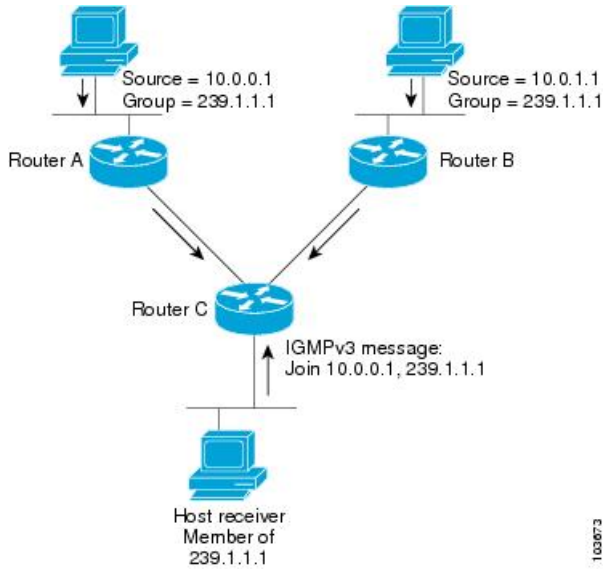

Configuring IGMP Per Interface States Limit

The IGMP Per Interface States Limit sets a limit on creating OIF for the IGMP interface. When the set limit is reached, the group is not accounted against this interface but the group can exist in IGMP context for some other interface.

- If a user has configured a maximum of 20 groups and has reached the maximum number of groups, then no more groups can be created. If the user reduces the maximum number of groups to 10, the 20 joins will remain and a message of reaching the maximum is displayed. No more joins can be added until the number of groups has reached less than 10.

- If a user already has configured a maximum of 30 joins and add a max of 20, the configuration occurs displaying a message that the maximum has been reached. No state change occurs and also no more joins can occur until the threshold number of groups is brought down below the maximum number of groups.

Configuration Example

Configures all interfaces with 4000 maximum groups per interface except TenGigE interface 0/0/0/6, which is set to 3000:

```
Router#config
Router(config)#router igmp
Router(config-igmp)#maximum groups-per-interface 4000
Router(config-igmp)#interface TenGigE0/0/0/6
Router(config-igmp-default-if)#maximum groups-per-interface 3000
Router(config-igmp-default-if)#commit
```

Running Configuration

```
router igmp
 interface TenGigE0/0/0/6
  maximum groups-per-interface 3000
 !
 maximum groups-per-interface 4000
!
```

Verification

```
Router#show igmp summary
Robustness Value 2
No. of Group x Interfaces 37
Maximum number of Group x Interfaces 50000
Supported Interfaces   : 9
Unsupported Interfaces: 0
Enabled Interfaces     : 8
Disabled Interfaces    : 1
MTE tuple count        : 0
Interface                      Number  Max #
                               Groups  Groups
Loopback0                       4       4000
TenGigE0/0/0/0                  5       4000
TenGigE0/0/0/1                  5       4000
TenGigE0/0/0/2                  0       4000
TenGigE0/0/0/3                  5       4000
TenGigE0/0/0/6                  5       3000
TenGigE0/0/0/18                 5       4000
TenGigE0/0/0/19                 5       4000
TenGigE0/0/0/6.1                3       4000
```

SSM Static Source Mapping

Configure a source (1.1.1.1) as part of a set of sources that map SSM groups described by the specified access-list (4).

Configuration Example

```
Router#configure
Router(config)#ipv4 access-list 4
Router(config-ipv4-acl)#permit ipv4 any 229.1.1.0 0.0.0.255
Router(config-ipv4-acl)#exit
```

```
Router(config)# multicast-routing
Router(config-mcast)#address-family ipv4
Router(config-mcast-default-ipv4)#ssm range 4
Router(config-mcast-default-ipv4)#exit
Router(config-mcast)#exit
Router(config)#router igmp
Router(config-igmp)#ssm map static 1.1.1.1 4
*/Repeat the above step as many times as you have source addresses to include in the set
for SSM mapping/*
Router(config-igmp)#interface TenGigE0/0/0/3
Router(config-igmp-default-if)#static-group 229.1.1.1
Router(config-igmp-default-if)#commit
```

Running Configuration

```
Router#show run multicast-routing
multicast-routing
 address-family ipv4
  ssm range 4
  interface all enable
 !
!
Router#show access-lists 4
ipv4 access-list 4
 10 permit ipv4 any 229.1.1.0 0.0.0.255

Router#show run router igmp
router igmp
 interface TenGigE0/0/0/3
 static-group 229.1.1.1
 !
 ssm map static 1.1.1.1 4
```

Verification

Verify if the parameters are set according to the configured values:

```
Router#show mrib route 229.1.1.1 detail
IP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
    C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
    IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
    MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
    CD - Conditional Decap, MPLS - MPLS Decap, EX - Extranet
    MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
    MoFB - MoFRR Backup, RPFID - RPF ID Set, X - VXLAN
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
    NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
    II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
    LD - Local Disinterest, DI - Decapsulation Interface
    EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
    EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
    MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
    IRMI - IR MDT Interface
(1.1.1.1,229.1.1.1) RPF nbr: 1.1.1.1 Flags: RPF
  Up: 00:01:11
  Incoming Interface List
    Loopback0 Flags: A, Up: 00:01:11
  Outgoing Interface List
    TenGigE0/0/0/3 Flags: F NS LI, Up: 00:01:11
```

Information About 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. An empty profile enables the default configuration options and settings for IGMP snooping, as listed in the *Default IGMP Snooping Configuration Settings.*

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 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. Profiles can be attached to ports under the bridge, one profile per port.
- 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.

Note

- IEEE 802.1Q, Q-in-Q, default and untagged encapsulations are supported.
- IEEE 802.1ad and other encapsulations are not supported.
- IEEE 802.1ad is not supported with BVI.

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 **mldp snooping profile** command in l2vpn bridge group bridge domain configuration mode. To attach a profile to a port, use the **mldp 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, MLDP snooping is deactivated in the bridge domain.
- If you detach a profile from a port, MLDP 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 you need to change an active profile, you must detach it from all bridges or ports, change it, 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.

Configuring Access Control

Access control configuration is the configuration of access groups and weighted group limits.

The role of access groups in IGMP v2/v3 message filteringis to permit or deny host membership requests for multicast groups (*,G) and multicast source groups (S,G). This is required to provide black-and-white list access to IPTV channel packages.

Weighted group limits restrict the number of IGMP v2/v3 groups, in which the maximum number of concurrently allowed multicast channels can be configured on a per EFP- and per PW-basis.

IGMP Snooping Access Groups

Although Layer-3 IGMP routing also uses the **igmp access-group** command in support of access groups, the support is not the same in Layer-2 IGMP, because the Layer-3 IGMP routing access group feature does not support source groups.

Access groups are specified using an extended IP access list referenced in an IGMP snooping profile that you attach to a bridge domain or a port.

Note A port-level access group overrides any bridge domain-level access group.

The **access-group** command instructs IGMP snooping to apply the specified access list filter to received membership reports. By default, no access list is applied.

Changes made to the access-list referenced in the profile (or a replacement of the access-list referenced in the igmp snooping profile) will immediately result in filtering the incoming igmp group reports and the existing group states accordingly, without the need for a detach-reattach of the igmp snooping profile in the bridge-domain, each time such a change is made.

IGMP Snooping Group Weighting

To limit the number of IGMP v2/v3 groups, in which the maximum number of concurrently allowed multicast channels must be configurable on a per EFP-basis and per PW-basis, configure group weighting.

IGMP snooping limits the membership on a bridge port to a configured maximum, but extends the feature to support IGMPv3 source groups and to allow different weights to be assigned to individual groups or source groups. This enables the IPTV provider, for example, to associate standard and high- definition IPTV streams, as appropriate, to specific subscribers.

This feature does not limit the actual multicast bandwidth that may be transmitted on a port. Rather, it limits the number of IGMP groups and source-groups, of which a port can be a member. It is the responsibility of the IPTV operator to configure subscriber membership requests to the appropriate multicast flows.

The **group policy** command, which is under igmp-snooping-profile configuration mode, instructs IGMP snooping to use the specified route policy to determine the weight contributed by a new <*,G> or <S,G> membership request. The default behavior is for there to be no group weight configured.

The **group limit** command specifies the group limit of the port. No new group or source group is accepted if its contributed weight would cause this limit to be exceeded. If a group limit is configured (without group policy configuration), a <S/*,G> group state will have a default weight of 1 attributed to it.

Note By default, each group or source-group contributes a weight of 1 towards the group limit. Different weights can be assigned to groups or source groups using the group policy command.

The group limit policy configuration is based on these conditions:

- Group weight values for <*,G> and <S,G> membership are configured in a Route Policy, that is included in an igmp snooping profile attached to a BD or port.
- Port level weight policy overrides any bridge domain level policy, if group-limit is set and route-policy is configured.
- If there is no policy configured, each group weight is counted equally and is equal to 1.
- If policy has been configured, all matching groups get weight of 1 and un-matched groups have 0 weight.

Statistics for Ingress Multicast Routes

Multicast and interface statistics are often used for accounting purpose. By default Multicast Forwarding Information Base (MFIB) does not store multicast route statistics. In order to enable MFIB to create and store multicast route statistics for ingress flows use the **hw-module profile mfib statistics** command in the configuration mode. Use No form of the latter command to disable logging of multicast route statistics. For the configuration to take effect, you must reload affected line cards after executing the command .

Note The MFIB counter can store a maximum of 2000 multicast routes, if the count exceeds beyond 2000 routes, then statistics are overwritten.

This table lists commands used to display or reset multicast route statistics stored in MFIB.

Table 1: show and clear commands for Multicast Statistics

Command	Description
show mfib hardware route statistics location <node-id>	Displays platform-specific MFIB information for the packet and byte counters for all multicast route.
show mfib hardware route <source-address> location <node-id>	Displays platform-specific MFIB information for the packet and byte counters for multicast routes originating from the specified source.

Command	Description
Show mfib hardware route <source-address><group-address> location <node-id>	Displays platform-specific MFIB information for the packet and byte counters for multicast routes originating from the specified source and belonging to specified multicast group.
clear mfib hardware ingress route statistics location <node-id>	Resets allocated counter values regardless of the MFIB hardware statistics mode from the designated node.
clear mfib hardware route <source-address> location <node-id>	Resets allocated counter values regardless of the MFIB hardware statistics mode from the designated node for specified multicast route source.
clear mfib hardware route <source-address><group-address> location <node-id>	Resets allocated counter values regardless of the MFIB hardware statistics mode from the designated node for specified multicast route source and multicast group.

Configuring Statistics for Ingress Multicast Routes

Configuration Example

In this example you will enable MRIB route statistics logging for ingress multicast routes for all locations:

```
Router#config
Router(config)#hw-module profile mfib statistics
Router(config)#commit
Router(config)#exit
Router#admin
Router(admin)#reload location all      /*Reloads all line cards. This is required step,
else the multicast route statistcis will not be created.*/
```

Running Configuration

```
Router#show running config
hw-module profile mfib statistics
!
```

Verification

The below show commands display the multicast statistics for source (192.0.2.2), group (226.1.2.1) and node location 0/1/cpu0 for ingress route:

Note The multicast egress statistics per flow (per SG) is not suported. But egress interface level multicast statistics is supported. Also the drop statistics is not supported.

```
Router#show mfib hardware route statistics location 0/1/cpu0
(192.0.2.2, 226.1.2.1) :: Packet Stats 109125794 Byte Stats 23680297298
(192.0.2.2, 226.1.2.2) :: Packet Stats 109125760 Byte Stats 23680289920
(192.0.2.2, 226.1.2.3) :: Packet Stats 109125722 Byte Stats 23680281674
```

```
(192.0.2.2, 226.1.2.4) :: Packet Stats 109125683 Byte Stats 23680273211
(192.0.2.2, 226.1.2.5) :: Packet Stats 109125644 Byte Stats 23680264748
(192.0.2.2, 226.1.2.6) :: Packet Stats 109129505 Byte Stats 23681102585
(192.0.2.2, 226.1.2.7) :: Packet Stats 109129470 Byte Stats 23681094990
(192.0.2.2, 226.1.2.8) :: Packet Stats 109129428 Byte Stats 23681085876
(192.0.2.2, 226.1.2.9) :: Packet Stats 109129385 Byte Stats 23681076545
(192.0.2.2, 226.1.2.10) :: Packet Stats 109129336 Byte Stats 23681065912

Router#show mfib hardware route statistics 192.0.2.2 location 0/1/cpu0
(192.0.2.2, 226.1.2.1) :: Packet Stats 109184295 Byte Stats 23692992015
(192.0.2.2, 226.1.2.2) :: Packet Stats 109184261 Byte Stats 23692984637
(192.0.2.2, 226.1.2.3) :: Packet Stats 109184223 Byte Stats 23692976391
(192.0.2.2, 226.1.2.4) :: Packet Stats 109184184 Byte Stats 23692967928
(192.0.2.2, 226.1.2.5) :: Packet Stats 109184145 Byte Stats 23692959465
(192.0.2.2, 226.1.2.6) :: Packet Stats 109184106 Byte Stats 23692951002
(192.0.2.2, 226.1.2.7) :: Packet Stats 109184071 Byte Stats 23692943407
(192.0.2.2, 226.1.2.8) :: Packet Stats 109184029 Byte Stats 23692934293
(192.0.2.2, 226.1.2.9) :: Packet Stats 109183986 Byte Stats 23692924962

Router#show mfib hardware route statistics 192.0.2.2 226.1.2.1 location 0/1/cpu0
(192.0.2.2, 226.1.2.1) :: Packet Stats 109207695 Byte Stats 23698069815
```

Use Case: Video Streaming

In today's broadcast video networks, proprietary transport systems are used to deliver entire channel line-ups to each video branch office. IP based transport network would be a cost efficient/convenient alternative to deliver video services combined with the delivery of other IP based services. (Internet delivery or business services)

By its very nature, broadcast video is a service well-suited to using IP multicast as a more efficient delivery mechanism to reach end customers.

The IP multicast delivery of broadcast video is explained as follows:

1. Encoding devices in digital master headends, encode one or more video channels into a Moving Pictures Expert Group (MPEG) stream which is carried in the network via IP multicast.
2. Devices at video branch office are configured by the operator to request the desired multicast content via IGMP joins.
3. The network, using PIM-SSM as its multicast routing protocol, routes the multicast stream from the digital master headend to edge device receivers located in the video branch office. These edge devices could be edge QAM devices which modulate the MPEG stream for an RF frequency, or CMTS for DOCSIS.

Figure 3: Video Streaming