# IP Multicast Configuration Guide

### Configuration IP Multicast

IP Multicast is not a single protocol, it is IP and as such touches on the same configuration issues we see when configuring unicast IP: routing, forwarding, source identification, addressing, management, reliability, security, etc. As such, configuring and supporting IP Multicast is much easier if approached systematically, taking into account each issue in overall the multicast architecture:

- routing
- tree building and forwarding
- source discovery
- addressing
- layer 2 switching
- media specific issues
- reliability (reliable multicast)
- management
- security

The overall functionality of multicast services will be directly effected by improper or incomplete configuration of any one of these elements. In addition, the order in which some of these elements are configured is important, so the configuration information is presented in the suggested order that they be addressed, routing first.

### Routing

In order for multicast forwarding protocol like PIM to work properly, it must have a valid unicast route to the multicast source for which it is forwarding data. PIM performs what is called the Reverse Path Forwarding (RPF) check: it will accept a multicast packet it receives on an interface, if and only if, it has a preferred unicast route back to the multicast source pointing out that same interface. If the RPF check succeeds, PIM can accept the packet and proceed with tree construction and forwarding of the data down that tree. If the RPF check fails, the packet will be dropped.

For this reason, accurate routing information is absolutely vital to establish a functioning multicast service. This routing can be for interdomain multicast sources in which case MBGP should be utilized, or it can be for intradomain multicast sources, in which case the standard unicast IGP protocols can be utilized.




CISCO SYSTEMS

**Interdomain routing: MBGP**

For the interdomain routing space the obvious protocol is BGP. But in addition to the functionality that BGP4 provides, for multicast we also require that ability to establish a different path or policy for multicast as we do for unicast traffic. We can accomplish this using the multiprotocol extensions defined in BGP4+ (RFC 2283) also known as MBGP. MBGP allows us to specify:

– unicast routes for unicast destinations (for unicast)

– unicast routes for multicast sources (for multicast RPF check)

MBGP allows us to carry both types of routes in the same BGP peering session, and thus we can apply the same path/policy knobs and use the same BGP machinery we are familiar with in standard unicast BGP configurations.

BGP4+ defines two new multiprotocol attributes:

– MP_REACH_NLRI

– MP_UNREACH_NLRI

Further, by applying address family identifiers (AFI 1=IPv4), and subsequent address family identifiers, we can specify whether our prefix is reachable for unicast, multicast, or both: SAFI 1 = nlri unicast, SAFI 2 = nlri multicast, SAFI 3 = both.

By announcing and receiving this nlri reachability information in the MBGP peering session, we can now establish separate routing information bases (RIBs) for unicast and multicast. And because the same IPv4 prefix can be assigned to both lists, we can also apply different policy attributes to the same prefix, based on whether it is in the unicast RIB (and will be used for unicast forwarding), or the multicast RIB (and will be used for multicast RPF check). This information allows us to distinguish different paths and/or policies for unicast and multicast.

The MBGP peering process occurs as follows:

1. Peers are configured to send/receive nlri types

   ```
   [no] neighbor <address> remote-as <asn> [nlri unicast | multicast]
   ```

2. Peers negotiate what nlri type they will be exchanging

   When peers open the MBGP session, they will negotiate the parameters of the session based on the nlri each has configured (step 1), settling on common nlri type. If one peer does not yet offer the capability parameters, the cisco will back off and reopen with no parameters. For some early implementations it was necessary to specify no-negotiation, this command should no longer be necessary.

   ```
   neighbor <address> dont-capability-negotiate
   ```

3. Peers set nlri type on prefixes they will advertise

   There are three basic ways prefixes can be assigned to the Multicast RIB:

   a) prefixes can be received from peers with nlri multicast (AFI/SAFI 1/2 MP_REACH_NLRI)

   b) prefixes can be originated with nlri multicast through the same methods prefixes are originated in unicast bgp, with the exception that now nlri type multicast must be specified as well as, or instead of, unicast nlri.

   ```
   network <address> <mask> [nlri multicast unicast]
   redistribute <igp> route-map <map>
   aggregate-address <address> <mask> [nlri multicast unicast]
   neighbor <address> default-originate [nlri multicast unicast]
   ```

   c) unicast prefixes can be translated from a unicast-nlri-only peer

   CAUTION: This is in essence bgp->mbgp redistribution and should be configured only with extreme caution, but in doing so it can provide a effective method of interfacing with non-mbgp topologies (e.g. when doing initial migration of mbgp into existing networks).

   ```
   neighbor <address> translate-update route-map <map>
   route-map <map> permit 20
     match nlri unicast | multicast
   ```

   For more information on translate-update:

   ftp://ftpeng.cisco.com/ipmulticast/translate-update

4. Prefixes are advertised and withdrawn

   As with BGP peering, it is possible to assign an outbound router-map to neighbors or peer-groups.

   ```
   [no] neighbor <address> route-map <map> out
   ```

   Please note that the semantics of outbound route-maps wrt matching nlri type was changed in 12.0(4.0.4)S, 12.0(3.7)S7, and 11.1(26.1)CC. Prior to these versions, nlri type was automatically set to unicast-only by the methods described in step 3 above. So in early images the route-map first had to set the nlri to include multicast, if multicast prefixes were to be announced to peers. The semantics were changed so that if no specific nlri is set in the outbound route-map, then MBGP will send prefixes with the nlri previously negotiated with the peer.

   If it is necessary to change the nlri type for advertised prefixes, this can be done as follows:

   ```
   router bgp 123
   neighbor 123.1.1.1 remote-as 123
   neighbor 123.1.1.1 route-map <map> out

   route-map <map> permit 20
   set nlri multicast (eg, or unicast or unicast multicast)
   ```

   So, in currently recommended images, it is possible through the outbound route-map to assign policy or paths (by assigning next-hop) to prefixes based on matching a particular nlri type. Here is an example of setting next hop differently specifically for multicast prefixes:

   ```
   [no] neighbor <address> route-map <use-mix> out
   route-map use-mix permit 20
      match nlri multicast
      set next-hop <address>
   route-map use-mix permit 30
      match nlri unicast
   ```

5. Policies are applied to received prefixes per normal bgp attributes as defined through in-bound route-maps, matching on nlri before applying the policy.

   ```
   [no] neighbor <address> route-map <map> in
    route-map <map> permit 20
      match nlri unicast | multicast
      do something
   ```

#### Intradomain routing

When using Protocol Independent Multicast (PIM) as the multicast tree building and forwarding protocol, any routing protocol can be used to supply the route for the multicast RPF check. This is significant, as it allows us to use existing unicast routing internally assuming the RPF information defined by our unicast IGP results in the path we want multicast to take. So we can utilize any unicast IGP as long as unicast and multicast are congruent.

In addition, it is possible to utilize multicast-specific routing protocols in the event that we need different, or non-congruent, multicast paths. Currently there are three multicast-specific protocols available in Cisco IOS: MBGP for interdomain routing as previously described; DVMRP-unicast routing; and/or static mroutes for intradomain routing.

##### DVMRP-unicast routing

It is possible to send and receive DVMRP routes between Cisco routers when interfacing with DVMRP environments, or in limited situations where separate routing tables are required to provide valid RPF information for incongruent topologies.

```
[no] ip dvmrp unicast-routing
```

**Note:** Do not configure this command on DVMRP tunnels.

By default, ip dvmrp unicast-routing will allow a router to include in DVMRP reports prefixes of directly connected subnets on multicast enabled interfaces. It is also possible to inject unicast prefixes for subnets that are not directly connected into the DVMRP reports, or to modify the advertised metric for routes that are already in the DVMRP routing table:

```
[no] ip dvmrp metric <metric> [list <access-list>] [<protocol> <process-id>] | dvmrp]
```

You can also apply route maps to the metrics you wish to advertise:

```
[no] ip dvmrp metric <metric> [route-map <map-name>]
```

It is possible, but definitely not recommended, to distribute DVMRP routes into MBGP:

```
[no] ip dvmrp metric <metric> [route-map <map-name>] mbgp
```

It is also possible to generate a DVMRP default route, e.g., for a tail-site customer. This is particularly helpful so tail-sites cannot leak specific dvmrp routes out backdoors to legacy Mbone topologies.

```
[no] ip dvmrp default-information originate | only
```

Through the use of dvmrp-unicast routing it is possible to provide multicast service to dvmrp tail-sites customers.

Static mroutes

It is possible to statically define a route to be used for multicast RPF checks. These static routes are local to the router on which they are configured and they can not be redistributed.

```
[no] ip mroute <source> <mask> [<protocol><as-number>]
        [route-map <map>] <rpf-address> | <interface> [<distance>]
```

**Redistributing IGP into MBGP**

As with unicast BGP it is possible to redistribute an IGP into MBGP, including DVMRP (it is not possible to redistribute static mroutes into MBGP).

```
redistribute <igp> route-map <dvmrp2mbgp>
route-map dvmrp2mbgp permit 20
   set nlri multicast
```

**Note:** As with unicast BGP, redistributing any IGP into MBGP is not recommended.

**Multicast Route Selection**

When a PIM router is running multiple routing protocols, the following rules are used to determine which route will be used to perform the RPF check:

Rule 1) When different routing tables have different distances the one with the smallest value of distance is chosen; (see table of default administrative distances of most protocols)

| Route Source | Default Distance |
|---|---|
| Connected interface | 0 |
| Static mroute | 0 |
| DVMRP | 0 |
| Static route | 1 |
| EIGRP summary route | 5 |
| External M/BGP | 20 |
| Internal EIGRP | 90 |
| IGRP | 100 |
| OSPF | 110 |
| IS-IS | 115 |

| Route Source | Default Distance |
|---|---|
| RIP | 120 |
| EGP | 140 |
| External EIGRP | 170 |
| Internal M/BGP | 200 |
| Unknown | 255 |

Rule 2) When the distances are the same, DVMRP routes is preferred over other unicast routes.

Rule 3) When static mroutes are present, and configured with a distance value no greater than that of the DVMRP routes, the static mroutes are used.

Rule 4) Longest match routing is done among routes within the same routing table, but not across different tables.

It is possible to alter the administrative distances for multicast-specific protocols for MBGP:

```
[no] distance mbgp <dist1> <dist2> <dist3>
```

for dvmrp:

```
[no] ip dvmrp distance <admin-distance>
```

for static mroutes:

```
[no] ip mroute <source> <mask> [<protocol><as-number>]
  [route-map <map>] <rpf-address> | <interface> [<distance>]
```

## Tree Building And Forwarding

#### IP Multicast Forwarding

There are several forms of multicast forwarding available depending on the platform and interface h/w being used:

– process level multicast forwarding
– fast switching
– multicast distributed switching

All forwarding options require as a minimum the following global command:

```
ip multicast-routing [distributed]
```

Process level forwarding/fast switching

Once ip multicast-routing is enabled globally, fast switching of multicast can be enabled or disabled through the following command per interface:

```
[no] ip mroute-cache [distributed]
```

**Note:** On newer platforms ip mroute-cache is enabled by default.

Multicast Distributed Switching (MDS)

Multicast Distributed Switching (MDS) incorporates support for distributed switching of multicast packets at the line cards (VIPs in case of RSP and BFR lc in case of GSR). The switching function is always performed on the line cards so that the RP is used only to do route processing. This feature can work in conjunction with DFIB (cisco express forwarding), DFS (unicast distributed fast switching) or flow switching. Multicast distributed switching is accomplished using a forwarding data structure called MFIB which is downloaded from the RP to all the linecards.

Caveats/Restrictions:

- Prior to IOS versions 11.1(25.2)CC, 12.0(3.7)S3, and 12.0(4.0.4)S there was no MDS support available for switching packets received on subinterfaces.
- Support is available for Ethernet, Fddi, ATM, serial, and POS interfaces MDS is not available for GRE tunnels or Token Ring.
- When running with DFS or FIB, at least 32 meg of memory is a pre-requisite on both the RSP and VIP2 boards.
- MDS will work only on packets received on VIP2 interfaces. If the packet arrives on a VIP1/legacy IP, it will be fast-switched as before. Packets received on VIP2 interfaces and sent on VIP1 or legacy interfaces will be distributed switched.

In order to enable MDS first use the global configuration command:

```
ip multicast-routing distributed
```

This will also enable distributed fast switching on capable interfaces. MDS can be specifically enabled/disabled on a per interfaces basis by adding or removing the [distributed] option in the following interface command:

```
[no] ip mroute-cache [distributed]
```

For more detailed information about MDS, see:
ftp://ftpeng.cisco.com/ipmulticast/mds.txt

The actual switching mode configured for multicast forwarding on a PIM interface can be checked using the following show command, where D=MDS, *=fast switching, and no entry indicates process-level switching:

```
sho ip pim interface count
Address     Interface     FS  Mpackets  In/Out
10.1.1.1  FastEthernet1/0/0  D  x/y
11.1.1.  1FastEthernet2/0/0  *  x/y
12.2.2.2  FastEthernet3/0/0  x/y
```

#### Before you enable PIM: Bounding the PIM domain

There are various multicast control messages as well as some multicast content that must be prevented from entering or leaving a given PIM domain. For this reason it is strongly recommended that all external interfaces be configured with a multicast boundary prior to enabling PIM internally. The following would be the minimum access list required, stopping auto-rp and administratively scoped data:

```
ip multicast boundary 1
access-list 1 deny   224.0.1.39
access-list 1 deny   224.0.1.40
access-list 1 deny   239.0.0.0 0.255.255.255
access-list 1 permit any
```

If PIMv2 is used, then the following command should also be configured on external interfaces:

```
ip pim border
```

#### PIM Sparse mode

PIM Sparse mode is strongly recommended for most environments. It must be enabled as a minimum on any interface expected to participate in multicast forwarding.

```
ip pim sparse-mode
```

**Note:** Use of ip pim sparse-dense mode is required if auto-rp is to be used; see below.

A PIM Sparse Mode domain is characterized by the presence of a Rendezvous Point (RP). The RP is the point to which a source's first-hop router sends data, and receiver's last-hop router sends PIM join requests so that receivers can discover sources. PIM is the only sparse-mode protocol to perform this function. Other multicast protocols are either dense-mode or depend on a mechanism outside of the protocol itself to advertise sources to receivers. In order for PIM to correctly forward multicast packets

using sparse-mode, all routers in the topology must know the address of the RP for a particular multicast address range and to have a valid route back to the RP. A network manager may configure as many RPs as he likes, but only one RP address will be used for any given multicast group address. In most networks, a single RP can handle the entire multicast group range. There are two methods of specifying the RP address:

– static RP assignment
– auto-rp or bootstrap mechanism (PIMv2 only, IOS after 12.0)

Static RP assignment

To statically configure an RP, which must be consistently configured on all routers within a given administrative scope, configure the following command:

```
[no] ip pim rp-address <ip-address> [<group-access-list>] [override]
```

Using Auto-rp

It is very rarely desirable to statically configure the RP address on each router. Cisco provides an automated solution for establishing the RP called auto-rp. With auto-rp it is possible to configure only a single router and have the RP address announced to all other routers in the topology. There are three steps required to enabling auto-rp :

1. enable pim sparse-dense mode on all internal interfaces:
   ```
   [no] ip pim sparse-dense-mode
   ```
2. have candidate RPs announce which groups they are willing to be an RP for:
   ```
   [no] ip pim send-rp-announce <interface-unit> scope <ttl> group-list <acl  [interval
   <num-seconds>]
   ```
3. establish a RP mapping agent (router) which can listen for candidate RP announcements, resolve primary RP responsibility, and send the final RP mapping assignments out to all routers:
   ```
   [no] ip pim send-rp-discovery [<interface>] scope <ttl>
   ```
   When using a single RP, the same router can perform both steps 2 and 3. For more detailed information about Auto-RP see: ftp://ftpeng.cisco.com/ftp/ipmulticast/autorp.html

**Note**: To check if a particular router is receiving the final auto-rp mapping information, use:

`sho ip pim rp mapping` (the mapping keyword is important)

It should return something like this:

```
router#sho ip pim rp mapping
pim group-to-rp mappings
Group(s) 224.0.0.0/4
  RP 198.9.200.65 (rp.cisco.com), v2v1
    Info source: 198.9.200.65 (?), via Auto-RP
         Uptime: 2d08h, expires: 00:02:21
```

**Note**: The command “show ip pim rp” no longer returns useful information and can be misleading if you do not understand that it is showing you only the RPs that are in the multicast routing table, whether or not these RPs are correct. Be sure to configure ‘ip pim accept-rp auto-rp’ on all routers in your network.

Bootstrap Router Mechanism (PIMv2, IOS 12.0 and later)

If all routers in the network are running PIMv2, or if you have routers from another vendor that do not understand auto-rp, then you can configure a BSR instead of Auto-RP. Both are very similar distributing the same information. With BSR configuration, you configure BSR candidates (similar to RP-Announce in Auto-RP) and you configure a bootstrap router, BSR (similar to an Auto-RP Mapping Agent). We have more experience with Auto-RP, which works well. It is not possible to define RPs with different overlapping administrative scopes using a BSR. But, if you want to configure a BSR, this is how to do it:

1. On the candidate bootstrap routers configure:

   **ip pim bsr-candidate <interface> <hash-mask-len> <pref>**

   Where <interface> has the candidate BSR's IP address. It is recommended (but not required) that <hash-mask-len> be the same across all candidate BSRs. A candidate BSR with the highest <pref> value will be elected as the BSR for the entire domain. Example:

   **ip pim bsr-candidate ethernet0 30 4**

   The PIMv2 Bootstrap router (BSR) is used to collect candidate RP information and to disseminate RP-set information associated with each group prefix. To avoid single point of failure, more than one router in a domain can be configured as candidate BSRs. A BSR is elected among the candidate BSRs automatically, based on the preference values configured. The routers to serve as candidate BSRs should be well connected and be in the "backbone" part of the network, as opposed to in the "dial-up" part of the network.
2. Configure candidate RP routers.

   The following example shows a candidate RP, on the interface ethernet0, for the entire admin-scope address range:

   **ip pim rp-candidate ethernet0 group-list 11**

   **access-list 11 permit 239.0.0.0 0.255.255.255**

   For additional information about BSR and PIMv1/PIMv2 interoperability see:
   ftp://ftpeng.cisco.com/ipmulticast/pimv2_config_guide.txt

#### PIM Dense mode

PIM Dense mode may be sufficient for a given group when no multicast exchange is expected outside of the PIM administrative domain and when there is a member of the group on all, or almost all, subnets in the domain. There are few cases where this is true. [NOTE: configuration of dense is only slightly easier but it much harder to debug.] To enable PIM dense-mode:

1) Turn on multicast forwarding as described in section 3.1.

2) Bound the PIM domain as described in section 3.2.

3) Configure PIM sparse-dense mode on each interface to run multicast without the use of an RP and thus be in dense mode. This is recommended, as it will ease any future transition to sparse mode.

**ip pim sparse-dense-mode**

### Multicast Source Discovery Protocol (Msdp)

There is no inherent method to connect PIM sparse-mode domains together unless the RPs for each of the domains is placed on a common subnet configured with 'ip pim dense-mode'. MSDP provides a method of sharing information about active sources among remote multicast domains without having to place the RP for each domain on the domain's border and flood data between the domains.

#### MSDP peering

First configure an RP to MSDP peer with other RPs (or with some common intervening routers):

**ip msdp peer <address> [connect-source] <interface>**

When a source sends to its local designated RP running MSDP, that RP will generate a source-active message (sa-message) which it will send on to its peer(s). The sa-messages contain information about the source, group, and originating RP. Peers will accept the sa-message from an msdp peer, if: a) its is an external peer and its from the next AS in the as-path to the RP, or b) its an internal msdp peer with the valid [M]BGP next-hop to the RP, or c) the peer is configured as a default peer, or d) the peer is part of an MSDP mesh-group. These "peer-rpf checks" provide a method to prevent looping of messages.

**Note:** this means that MSDP peering and connect-source assignment are important in determining acceptance or denial of sa-messages. The address used for MSDP peering must reflect the expected next-hop. You can determine if sa-messages are being accepted or dropped by debugging:

**debug ip msdp peer <address>**

MSDP filtering and redistribution

It is possible to control which sa-messages are originated, which are accepted from msdp peers, and which are forwarded to other msdp peers.

a) controlling which sa-messages are originated

```
[no] ip msdp redistribute [list <acl>] [asn <aspath-acl>] [route-map <map>]
[no] ip msdp ttl-threshold <ip address> <ttl>
```

b) controlling which sa-messages are accepted

```
[no] ip msdp sa-filter in <ip-address-or-name> [list <acl>] [route-map <map>]
```

c) controlling which sa-messages are forwarded

```
[no] ip msdp sa-filter out <ip-address-or-name> [list <acl>] [route-map <map>
```

Caching MSDP messages

By default, MSDP routers do not retain sa-messages. By caching sa-messages it is possible to reduce join latency for local receivers. Typically, sa-caching is enabled at least on the primary RP(s), optionally on intervening routers when useful depending on local memory constraints. If the MSDP cache is not enabled, it may be very difficult to discover MSDP problems.

```
[no] ip msdp cache-sa-state [list <acl>]
```

You can then view this cache to see if you are getting the sa-messages you expect:

```
sho ip msdp sa-cache
```

**Logical RPs**

Logical RPs, as described in draft-ietf-mboned-logical-rp-00.txt provides a simple and effective method of regional load-sharing and redundancy among multiple RPs using the same RP address, serving the same multicast group address space in the same PIM domain. A source's first-hop router will send data to, and receiver's last-hop router will send PIM joins toward, the logical RP nearest to them by virtue of IGP metrics (the route to the RP address). Should a particular logical RP go down, sources and receivers first and last-hop routers simply use the next closest logical RP based on the next best metric to which IGP converges.

Logical RPs are configured as follows:

1. Configure on each RP the logical RP address by using a Loopback with an address common to all participating logical RPs in the domain. CAUTION: this loopback and its address must be configured such that they are not selected as the router-id for protocols like OSPF or BGP. Avoid this by selecting a common loopback for all logical RPs (e.g. loopback 10), and providing it with a lower IP address than any other loopback.
2. Establish each logical RP as if it were the only RP as described in the PIM sparse mode section:

```
[no] ip pim send-rp-announce <interface-unit> scope <ttl> group-list <acl  [interval
<num-seconds>]
```

**Note:** Make sure the scope <ttl> is sufficient to reach all routers throughout the domain.

```
[no] ip pim send-rp-discovery [<interface>] scope <ttl>
```

**Note:** Make sure the scope <ttl> is sufficient to reach all routers in the PIM domain served by the logical RPs.

3. Configure each logical RP to inject the RP address into the IGP for the local domain.
4. Establish full mesh MSDP peering among all logical RPs and include them specifically into an MSDP mesh-group by configuring the following on each, for each logical RP:

```
[no] ip msdp mesh-group <name> <ip-address-or-name>
```

## Multicast Addressing

Multicast sessions which remain within a particular administrative PIM domain should use administratively scoped addresses, while multicast sessions which will cross the interdomain space can obtain addresses from the global address space using methods like sdr, or via static assignment from 233/8 per the GLOP proposal (draft-ietf-mboned-static-allocation-00.txt).

#### Administratively scoped addressing

The key properties of administratively scoped IP multicast are: (i) packets addressed to administratively scoped multicast addresses do not cross configured administrative boundaries, and (ii) administratively scoped multicast addresses are locally assigned, and hence are not required to be unique across administrative boundaries. The administratively scoped IPv4 multicast address space is defined to be the range 239.0.0.0 to 239.255.255.255.

In order to support administratively scoped IP multicast, a router must be configured with the following on the interfaces of the scoped IP multicast boundaries.

```
ip multicast boundary 1
access-list 1 deny   239.0.0.0 0.255.255.255
access-list 1 permit any
```

These multicast scoped boundaries can be set up in such a way as to provide hierarchy of scoped regions. Multicast sessions limited to a particular region should then select addressing from the appropriate scoped address range for that region. A local RP can then be set-up to serve the address range of that scoped region. [Note: this cannot be done using the bootstrap router in PIMv2 unless the region is congruent will all other multicast groups.]

#### Global Addressing

Router configuration for global groups focuses primarily on the use of multicast boundaries as with admin scoping to control portions of the global address space based on policy. Applications can obtain addressing for global groups through various methods, two currently in use include

– dynamically using sdr application
– statically using 233/8 with AS mapping (GLOP)

##### Session Directory (sdr)

Session Directory (sdr) is a multicast application for setting up conferencing sessions. It uses the protocols SDP/SAP to disseminate the names and properties of conferencing sessions over the well known session directory groups 224.2.127.254 for global scope sessions and 239.255.255.255 for administrative scope session. The session properties include contact information, session lifetime and the media being used in the session (audio, video, whiteboard and others) with their specific attributes like ttl-scope, group address and UDP port number.

When **ip sdr listen** is configured on an interface, the software will join the well known directory groups on that interface to receive and store session announcements. The announcements can be displayed with the **show ip sdr** command. IOS uses stored announcements for the **ip multicast rate-limit** command. Use the **ip sdr cache-timeout** command to configure the period of time after which received announcements are expired.

When the system has **no ip multicast routing** configured, announcements are only stored if they are received on an interfaces with **ip sdr listen** configured. When the system is configured as a multicast router, it is sufficient to configure **ip sdr listen** on just a single multicast enabled interface. The well known session directory groups are handled as local joined groups after the first **ip sdr listen** is configured (see "L - local" flag in **show ip mroute**). This causes announcements received from all multicast enabled interfaces to be routed within the system and stored.

The following router example enables a router to listen to session directory advertisements:

```
ip routing

interface loopback 0
ip address 10.0.0.51 255.255.255.0
ip pim sparse-dense mode
ip sdr listen
```

#### Static assignments from 233/8

The IETF draft originally titled GLOP (draft-ietf-mboned-static-allocation-00.txt) outlines a method for static assignment of global address space in the 233/8 range. In order to claim and use statically assigned addressing from 233/8 you must have a registered non-private AS. The AS number can be mapped into the middle two octets of 233/8 to obtain the statically assigned /24 for global routing. See:

http://gigapop.uoregon.edu/glop/index.html

#### Multicast NAT

When unicast address translation (NAT) is configured on a Cisco IOS router, multicast sources/receivers or PIM entities, like RP/RP mapping agent, can work on either side of a NAT box without any additional configuration commands. All the routers (inside/outside and NAT box itself) must be fully multicast enabled as usual.

Supported Address Translation:

- Data packet source address translation.
- PIM control packet (PIM payload) address translation, including auto-rp, PIMv2 BSR.
- Mstat/mrinfo/mtrace requests/responses.
- SDR advertisement (SDR application payload).

With the above translations, PIM should just work "fine" in an enterprise domain even if a part of the domain is behind NAT. All sources/receivers behind that NAT box should be able to send/receive to the rest of the PIM cloud and take advantage of the RP/RP mapping agent on either side of the cloud.

Caveats:

- Tunnels must be terminated on the NAT box and "ip nat inside/outside" should be configured on the Tunnel. Tunnel cannot run through the NAT box with end points on either side.
- Addresses in RTP/RTCP or other application payloads are not translated.
- This feature does not translate destination Group addresses.

### Layer 2 Switching Issues

Without some form of layer 2 mechanism, most switches treat multicast packets as broadcast packets, forwarding them to all active ports. With the forwarding capabilities of today's switches, certain amount of multicast traffic can be handled this way without any problem. However, it is more efficient, and at higher multicast rates critical, to be able to forward the multicast packet only on ports with interested hosts. Cisco provides two methods for the switch to perform this layer 2 forwarding of multicast

- Cisco Group Management Protocol (CGMP)
- IGMP snooping

#### Cisco Group Management Protocol (CGMP)

In a switched environment with connected multicast routers, CGMP allows the router to inform the switch of interested hosts based on IGMP information available to the router. So CGMP required configuration of the router and the switches.

On the router interface facing the switch:

```
(pim must be enabled on interface)
ip cgmp
```

On the switch:

```
set cgmp enable
```

#### Internet Group Management Protocol (IGMP) Snooping

IGMP snooping is available with release 4.1 of the Catalyst 5000. IGMP Snooping requires a Supervisor III card. Unlike CGMP, with IGMP snooping no configuration, (other than pim), is necessary on the routers. At least one router must however be connected to the switched environment in order to provide the igmp querying function.

This example shows how to enable IGMP snooping on the switch:

```
console> (enable) set igmp enable
IGMP Snooping is enabled.
CGMP is disabled.
```

This example shows what happens if you try to enable IGMP if CGMP is already enabled:

```
Console> (enable) set igmp enable
Disable CGMP to enable IGMP Snooping feature.
```

### Reliable Multicast

#### Pragmatic General Multicast (PGM)

Pragmatic General Multicast (PGM) is a reliable multicast transport protocol for applications that require ordered, duplicate-free, multicast data delivery from multiple sources to multiple receivers. PGM guarantees that a receiver in the group either receives all data packets from transmissions and retransmissions, or is able to detect unrecoverable data packet loss.

There are no PGM global commands. PGM is configured per interface:

```
[no] ip pgm
```

**Note:** ip multicast-routing must be enabled globally and PIM must be enabled on the interface first.

### Multicast Management

#### Multicast Reachability Monitor (MRM)

Multicast Reachability Monitor facilitates automated fault detection and isolation in a large multicast routing infrastructure. It is designed to alarm a network administrator of multicast routing problems in close to real-time.

MRM has two types of components, MRM tester and MRM manager. MRM Tester is a sender and/or receiver.

MRM is available in IOS 12.0(5)T onwards. Only the MRM testers and managers need to be running MRM supported IOS version.

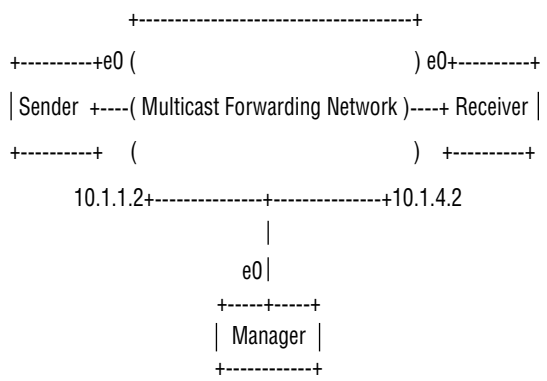


Make sure the “Multicast Forwarding Network” has no access-lists and boundaries that denies MRM data/control traffic. MRM test data is UDP/RTP packets addressed to configured group address.

MRM control traffic between sender, receiver and manager, is addressed to 224.0.1.111 group which is joined by all three.

Test Sender:

```
interface Ethernet0
ip mrm test-sender
```

Test Receiver:

```
interface Ethernet0
ip mrm test-receiver
```

Test Manager:

```
ip mrm manager test1
manager e0 group 239.1.1.1
senders 1
receivers 2 sender-list 1
access-list 1 permit 10.1.1.2
access-list 2 permit 10.1.4.2
Router# show ip mrm manager
Manager:test1/10.1.2.2 is not running
Beacon interval/holdtime/ttl:60/86400/32
Group:239.1.1.1, UDP port test-packet/status-report:16384/65535
Test sender:
10.1.1.2
Test receiver:
10.1.4.2
```

Start the test. Manager sends control messages to test-sender and test-receiver as configured in the test parameters. The test-receiver joins the group and monitors test packets sent from the test-sender.

```
Router# mrm start test1
*Feb  4 10:29:51.798: IP MRM test 'test1' starts ......
Router#
Display Status report on manager:
Router# show ip mrm status
IP MRM status report cache:
Timestamp        Manager         Test Receiver   Pkt Loss/Dup (%)      Ehsr
*Feb  4 14:12:46 10.1.2.2        10.1.4.2        1            (4%)     29
*Feb  4 18:29:54 10.1.2.2        10.1.4.2        1            (4%)     15
Router#
```

The above display shows that the receiver sent two status reports (one line each) at given time stamp. Each report contains 1 packet loss during the interval window (default 1 second). The Ehsr value shows the estimated next sequence number value from the test-sender. If the receiver saw duplicate packets, it would show a negative number in the "Pkt Loss/Dup" column.

Stop the test.

```
Router# mrm stop test1
*Feb  4 10:30:12.018: IP MRM test 'test1' stops
Router#
```

While running the test, MRM sender will start sending RTP packets to configured group address at default interval of 200 ms. The receiver will monitor (expect) the same packets at the same default interval. If the receiver detects a packet loss in default window interval of 5 second, it sends report to MRM manager. The status report from receiver can be seen by "show ip mrm status" command on manager.

For more information on MRM, see:

ftp://ftpeng.cisco.com/ipmulticast/mrm/mrm.guide

**Multicast MIBs**

The following multicast related MIBs are supported:

CISCO-IPMROUTE-MIB

IGMP-MIB

IPMROUTE-MIB

PIM-MIB

For the latest MIB files, see:

ftp://ftpeng.cisco.com/ipmulticast/mibs/


**Corporate Headquarters**
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

**European Headquarters**
Cisco Systems Europe s.a.r.l.
Parc Evolic, Batiment L1/L2
16 Avenue du Quebec
Villebon, BP 706
91961 Courtaboeuf Cedex
France
http://www-europe.cisco.com
Tel: 33 1 69 18 61 00
Fax: 33 1 69 28 83 26

**Americas Headquarters**
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

**Asia Headquarters**
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
http://www.cisco.com
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

**Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Connection Online Web site at http://www.cisco.com/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Singapore Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela



8/99 LW