Review the IP Multicast Quick-Start Configuration Guide

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

Dense Mode

Sparse Mode with one RP

Sparse Mode with Multiple RPs

Auto-RP with one RP

Auto-RP with Multiple RPs

DVMRP

MBGP

MSDP

Stub Multicast Routing

IGMP UDLR for Satellite Links

PIMv2 BSR

CGMP

IGMP Snooping

PGM

MRM

Troubleshoot

Related Information

Introduction

This document describes the basics of how to configure multicast for various networking scenarios.

Prerequisites

Requirements

Cisco recommends that you have knowledge of this topic:

- Internet Protocol (IP) Multicast.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Background Information

IP multicasting is a bandwidth-conserving technology that reduces traffic because it simultaneously delivers a single stream of information to thousands of corporate recipients and homes.

Applications that take advantage of multicast include video conferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news.

Dense Mode

Cisco recommends that you use Protocol Independent Multicast (PIM) sparse mode, particularly Auto-RP, where possible and especially for new deployments.

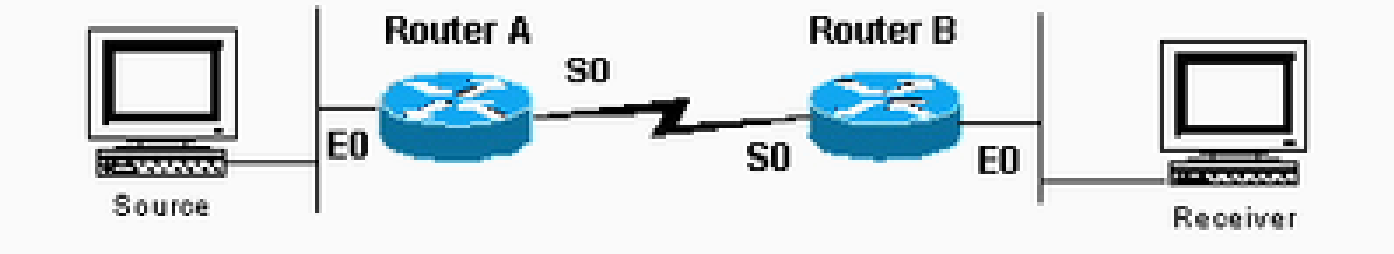
However, if dense mode is desired, configure the global command **ip multicast-routing** and the interface command **ip pim sparse-dense-mode** on each interface that needs to process multicast traffic.

The common requirement, for all configurations within this document, is to configure multicasting globally and configure PIM on the interfaces.

As of Cisco IOS® Sotware Release 11.1, you can configure the interface commands **ip pim dense-mode** and **ip pim sparse-mode** simultaneously with the **ip pim sparse-dense-mode** command.

In this mode, the interface is treated as dense-mode if the group is in dense-mode. If the group is in sparse-mode (for example, if an RP is known), the interface is treated as sparse-mode.

Note: The "Source" in the examples throughout this document represents the source of multicast traffic, and "Receiver" represents the receiver of multicast traffic.



Interface is Treated as Dense-mode if Group is in Dense-mode

Router A Configuration
`ip multicast-routing`

```
interface ethernet0
ip address <address> <mask>
ip pim sparse-dense-mode

interface serial0
ip address <address> <mask>
ip pim sparse-dense-mode
```

Router B Configuration

```
ip multicast-routing

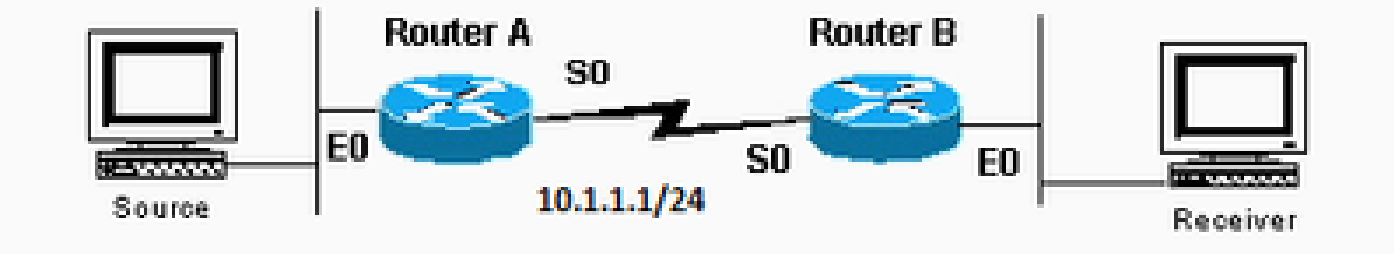
interface serial0
ip address <address> <mask>
ip pim sparse-dense-mode

interface ethernet0
ip address <address> <mask>
ip pim sparse-dense-mode
```

Sparse Mode with one RP

In this example, Router A is the RP which is typically the closest router to the source. Static RP configuration requires that all routers in the PIM domain have the same**i p pim rp-address** commands configured.

You can configure multiple RPs, but there can only be one RP per specific group.



There can be Multiple RPs, but Only One RP Per Specific Group

Router A Configuration

```
ip multicast-routing
ip pim rp-address 10.1.1.1


interface ethernet0
ip address <address> <mask>
ip pim sparse-dense-mode

interface serial0
ip address 10.1.1.1 255.255.255.0
ip pim sparse-dense-mode
```

Router B Configuration

```
ip multicast-routing
ip pim rp-address 10.1.1.1


interface serial0
ip address <address> <mask>
ip pim sparse-dense-mode

interface ethernet0
ip address <address> <mask>
ip pim sparse-dense-mode
```

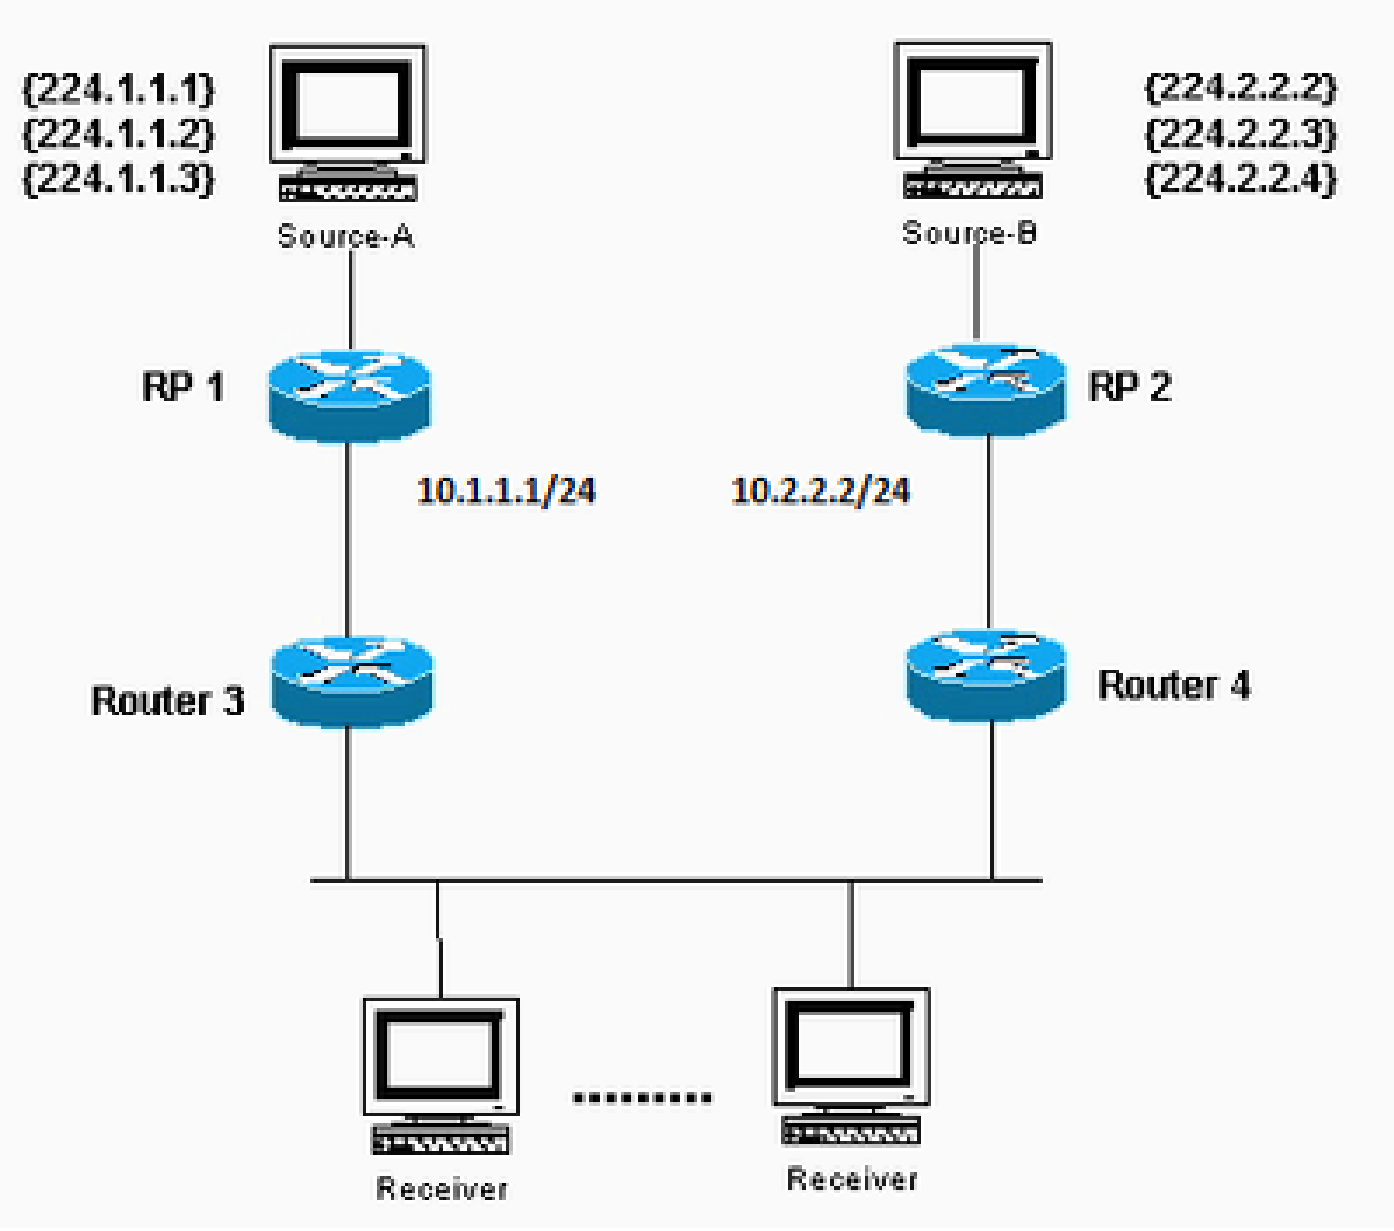
Sparse Mode with Multiple RPs

In this example, Source-A sends to 224.1.1.1, 224.1.1.2, and 224.1.1.3. Source-B sends to 224.2.2.2, 224.2.2.3, and 224.2.2.4. You could have one router, either RP 1 or RP 2, be the RP for all groups.

However, if you want different RPs to handle different groups, you need to configure all routers to include which groups the RPs can serve.

This type of static RP configuration requires that all routers in the PIM domain have the same **ip pim rp-address address acl** commands configured.

You can also use Auto-RP in order to achieve the same setup, which is easier to configure.



Source-A sends to 224.1.1.1, 224.1.1.2, and 224.1.1.3; Source-B sends to 224.2.2.2, 224.2.2.3, and 224.2.2.4.

RP 1 Configuration

```
ip multicast-routing

ip pim RP-address 10.1.1.1 2
ip pim RP-address 10.2.2.2 3

access-list 2 permit 224.1.1.1
access-list 2 permit 224.1.1.2
access-list 2 permit 224.1.1.3
access-list 3 permit 224.2.2.2
access-list 3 permit 224.2.2.3
access-list 3 permit 224.2.2.4
```

RP 2 Configuration

```
ip multicast-routing
```

```
ip pim RP-address 10.1.1.1 2
ip pim RP-address 10.2.2.2  3

access-list 2 permit 224.1.1.1
access-list 2 permit 224.1.1.2
access-list 2 permit 224.1.1.3
access-list 3 permit 224.2.2.2
access-list 3 permit 224.2.2.3
access-list 3 permit 224.2.2.4
```

Configuration for Routers 3 and 4

```
ip multicast-routing
ip pim RP-address 10.1.1.1  2
ip pim RP-address 10.2.2.2  3


access-list 2 permit 224.1.1.1
access-list 2 permit 224.1.1.2
access-list 2 permit 224.1.1.3
access-list 3 permit 224.2.2.2
access-list 3 permit 224.2.2.3
access-list 3 permit 224.2.2.4
```

Auto-RP with one RP

Auto-RP requires that you configure the RPs to announce their availability as RPs and mapping agents. The RPs use 224.0.1.39 to send their announcements.

The RP mapping agent listens to the announced packets from the RPs, then sends RP-to-group mappings in a discovery message that is sent to 224.0.1.40.

These discovery messages are used by the residual routers for their RP-to-group map.

You can use one RP that also serves as the mapping agent, or you can configure multiple RPs and multiple mapping agents for redundancy purposes.

Note that when you choose an interface from which to source RP announcements, Cisco recommends that you use an interface such as a loopback instead of a physical interface.

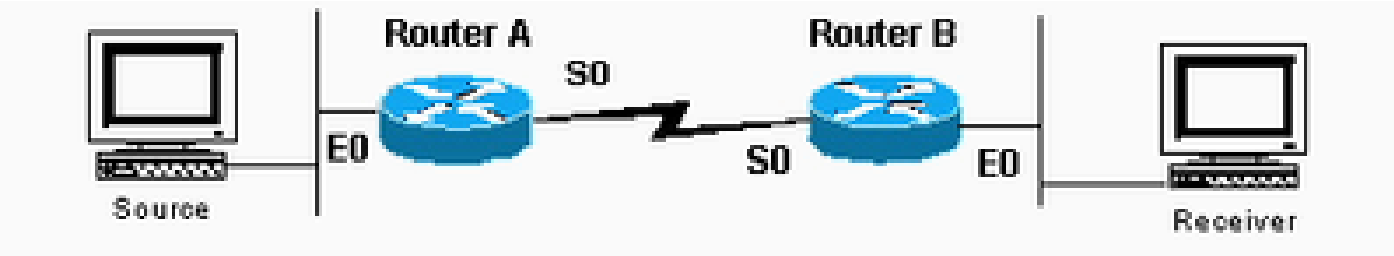
Also, it is possible to use Switched VLAN Interfaces (SVIs). If a VLAN interface is used to announce the RP address, then the **interface-type** option in the **ip pim [vrf vrf-name] send-rp-announce {interface-type interface-number | ip-address} scope ttl-value** command must contain the VLAN interface and the VLAN number. For example, the command looks like **ip pim send-rp-announceVlan500 scope 100** .

If you choose a physical interface, you rely on that interface to always be up. This is not always the case, and the router stops advertising itself as the RP once the physical interface goes down.

With a loopback interface, it is always up and never goes down, which ensures the RP continues to advertise itself through any available interfaces as an RP.

This is the case even if one or more of its physical interfaces fails. The loopback interface must be PIM-

enabled and advertised by an Interior Gateway Protocol (IGP), or it must be reachable with static routing.



The Loopback Interface Must be PIM-Enabled and Advertised by an Interior Gateway Protocol or Reachable with Static Routing

Router A Configuration

```
ip multicast-routing

ip pim send-rp-annouce loopback0 scope 16

ip pim send-rp-discover

 scope 16

interface loopback0
ip address <address> <mask>
ip pim sparse-dense-mode

interface ethernet0
ip address <address> <mask>
ip pim sparse-dense-mode

interface serial0
ip address <address> <mask>
ip pim sparse-dense-mode
```

Router B Configuration

```
ip multicast-routing


interface ethernet0
ip address <address> <mask>
ip pim sparse-dense-mode

interface serial0
ip address <address> <mask>
ip pim sparse-dense-mode
```

Auto-RP with Multiple RPs

The access lists in this example allow the RPs to be an RP only for the groups you want. If no access list is configured, the RPs are available as an RP for all groups.

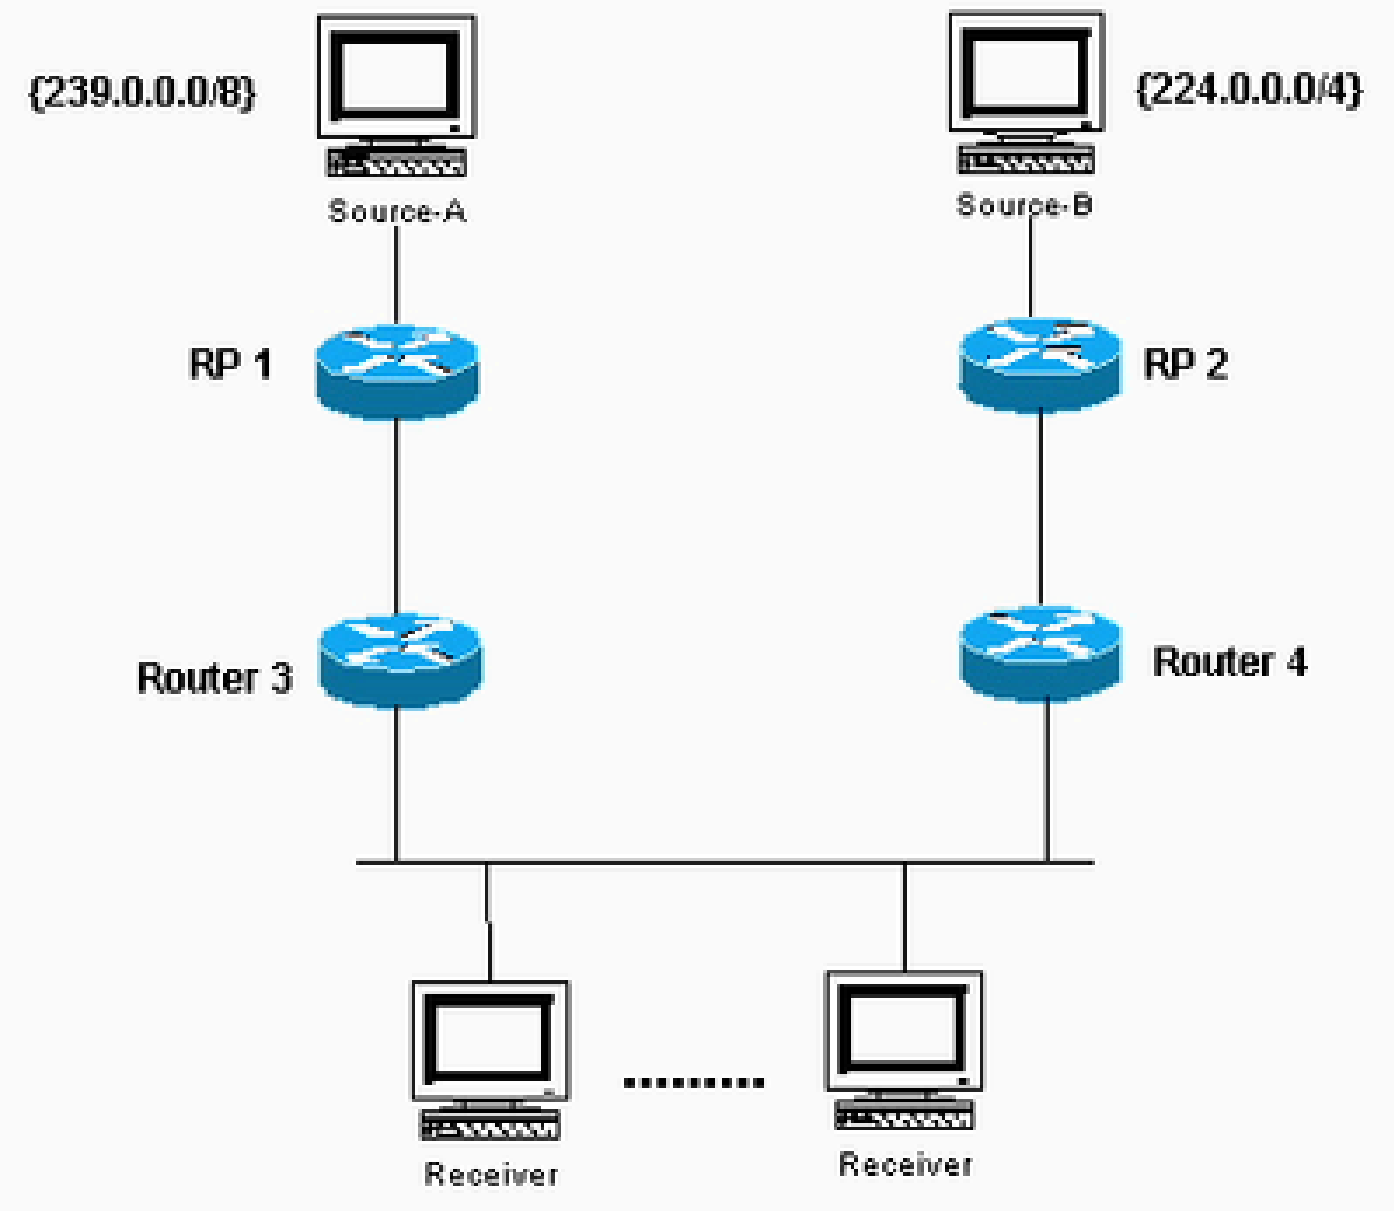
If two RPs announce their availability to be RPs for the same group(s), the mapping agent(s) resolve these

conflicts with "the highest IP address wins" rule.

When two RPs announce for that group, you can configure each router with a loopback address in order to influence which router is the RP for a particular group.

Place the higher IP address on the preferred RP, then use the loopback interface as the source of the announce packets; for example, **ip pim send-RP-announceloopback0** .

When multiple mapping agents are used, they each advertise the same group to RP mappings to the 224.0.1.40 discovery group.



Place the Higher IP Address on the Preferred RP

RP 1 Configuration
ip multicast-routing interface loopback0 ip address <address> <mask> ip pim sparse-dense-mode

```
ip pim send-RP-announce loopback0 scope 16 group-list 1
ip pim send-RP-discovery scope 16

access-list 1 permit 239.0.0.0 0.255.255.255
```

RP 2 Configuration

```
ip multicast-routing

interface loopback0
ip address <address> <mask>
ip pim sparse-dense-mode

access-list 1 deny 239.0.0.0 0.255.255.255
access-list 1 permit 224.0.0.0 10.255.255.255
```

DVMRP

Your Internet service provider (ISP) could suggest that you create a Distance Vector Multicast Routing Protocol (DVMRP) tunnel to the ISP in order to gain access to the multicast backbone in the Internet (mbone).

The minimum commands in order to configure a DVMRP tunnel are shown here:

```
interface tunnel0
ip unnumbered <any pim interface>
tunnel source <address of source>
tunnel destination <address of ISPs mrouted box>
tunnel mode dvmrp
ip pim sparse-dense-mode
```

Typically, the ISP has you tunnel to a UNIX machine that runs "mrouted" (DVMRP). If the ISP has you tunnel to another Cisco device instead, use the default GRE tunnel mode.

If you want to generate multicast packets for others on the mbone to see instead of receive multicast packets, you need to advertise the source subnets.

If your multicast source host address is 172.16.108.1, you need to advertise the existence of that subnet to the mbone. Directly-connected networks are advertised with metric 1 by default.

If your source is not directly connected to the router with the DVMRP tunnel, configure this under interface tunnel0:

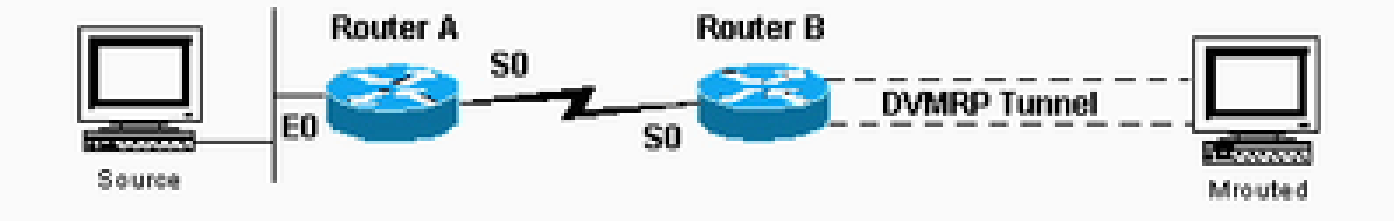
```
ip dvmrp metric 1 list 3
access-list 3 permit  172.16.108.0 0.0.0.255
```

Note: You must include an access list with this command in order to prevent advertising the entire Unicast routing table to the mbone.

If your setup is similar to the one shown here, and you want to propagate DVMRP routes through the domain, configure the**ip dvmrp unicast-routing**command on the serial0 interfaces of Routers A and B.

This action provides the forwarding of DVMRP routes to PIM neighbors who then have a DVMRP routing table used for Reverse Path Forwarding (RPF).

DVMRP learned routes take RPF precedence over all other protocols, except for directly-connected routes.



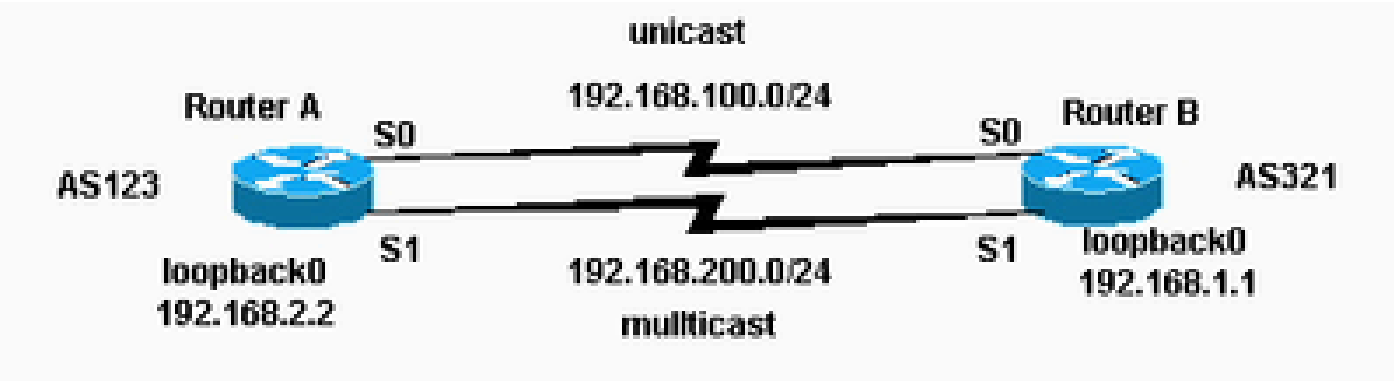
Propagate DVMRP Routes Through the Domain

MBGP

Multiprotocol Border Gateway Protocol (MBGP) is a basic method to carry two sets of routes: one set for unicast routing and one set for multicast routing.

MBGP provides the control necessary to decide where multicast packets are allowed to flow. PIM uses the routes associated with multicast routing in order to build data distribution trees.

MBGP provides the RPF path, not the creation of multicast state. PIM is still needed in order to forward the multicast packets.



PIM is Still Needed in Order to Forward the Multicast Packets

Router A Configuration
ip multicast-routing interface loopback0 ip pim sparse-dense-mode ip address 192.168.2.2 255.255.255.0

```
interface serial0
ip address 192.168.100.1 255.255.255.0


interface serial1
ip pim sparse-dense-mode
ip address 192.168.200.1 255.255.255.0


router bgp 123
network 192.168.100.0 nlri unicast
network 192.168.200.0 nlri multicast
neighbor 192.168.1.1 remote-as 321 nlri unicast multicast
neighbor 192.168.1.1 ebgp-multihop 255
neighbor 192.168.100.2 update-source loopback0
neighbor 192.168.1.1 route-map setNH out


route-map setNH permit 10
match nlri multicast
set ip next-hop 192.168.200.1


route-map setNH permit 20
```

Router B Configuration

```
ip multicast-routing


interface loopback0
ip pim sparse-dense-mode
ip address 192.168.1.1 255.255.255.0


interface serial0
ip address 192.168.100.2 255.255.255.0


interface serial1
ip pim sparse-dense-mode
ip address 192.168.200.2 255.255.255.0


router bgp 321
network 192.168.100.0 nlri unicast
network 192.168.200.0 nlri multicast
neighbor 192.168.2.2 remote-as 123 nlri unicast multicast
neighbor 192.168.2.2 ebgp-multihop 255
neighbor 192.168.100.1 update-source loopback0
neighbor 192.168.2.2 route-map setNH out


route-map setNH permit 10
match nlri multicast
set ip next-hop 192.168.200.2


route-map set NH permit 20
```

If your unicast and multicast topologies are congruent (for example, they go over the same link), the primary difference in the configuration is with the **nlri unicast multicast** command.

An example is shown here:

```
network 192.168.100.0 nlri unicast multicast
```

Congruent topologies with MBGP have a benefit—even though the traffic traverses the same paths, different policies can be applied to unicast BGP versus multicast BGP.

MSDP

Multicast Source Discovery Protocol (MSDP) connects multiple PIM-SM domains.

Each PIM-SM domain uses its own independent RP(s) and does not have to depend on RPs in other domains.

MSDP allows domains to discover multicast sources from other domains. If you are also BGP-peering with the MSDP peer, you must use the same IP address for MSDP as for BGP.

When MSDP does peer RPF checks, MSDP expects the MSDP peer address to be the same address that BGP/MBGP gives it when it performs a route table lookup on the RP in the SA message.

However, you are not required to run BGP/MBGP with the MSDP peer if there is a BGP/MBGP path between the MSDP peers.

If there is no BGP/MBGP path and more than one MSDP peer, you must use the **ip msdp default-peer** command.

The example here shows that RP A is the RP for its domain and RP B is the RP for its domain.



Multicast Source Discovery Protocol (MSDP) connects Multiple PIM-SM domains

Router A Configuration

```
ip multicast-routing


ip pim send-RP-announce loopback0 scope 16 group-list 1
ip pim send-RP-discovery scope 16

ip msdp peer 192.168.100.2
ip msdp sa-request 192.168.100.2
```

```
interface loopback0
ip address <address> <mask>
ip pim sparse-dense-mode

interface serial0
ip address 192.168.100.1 255.255.255.0
ip pim sparse-dense-mode
```

Router B Configuration

```
ip multicast-routing

ip pim send-RP-announce loopback0 scope 16 group-list 1
ip pim send-RP-discovery scope 16

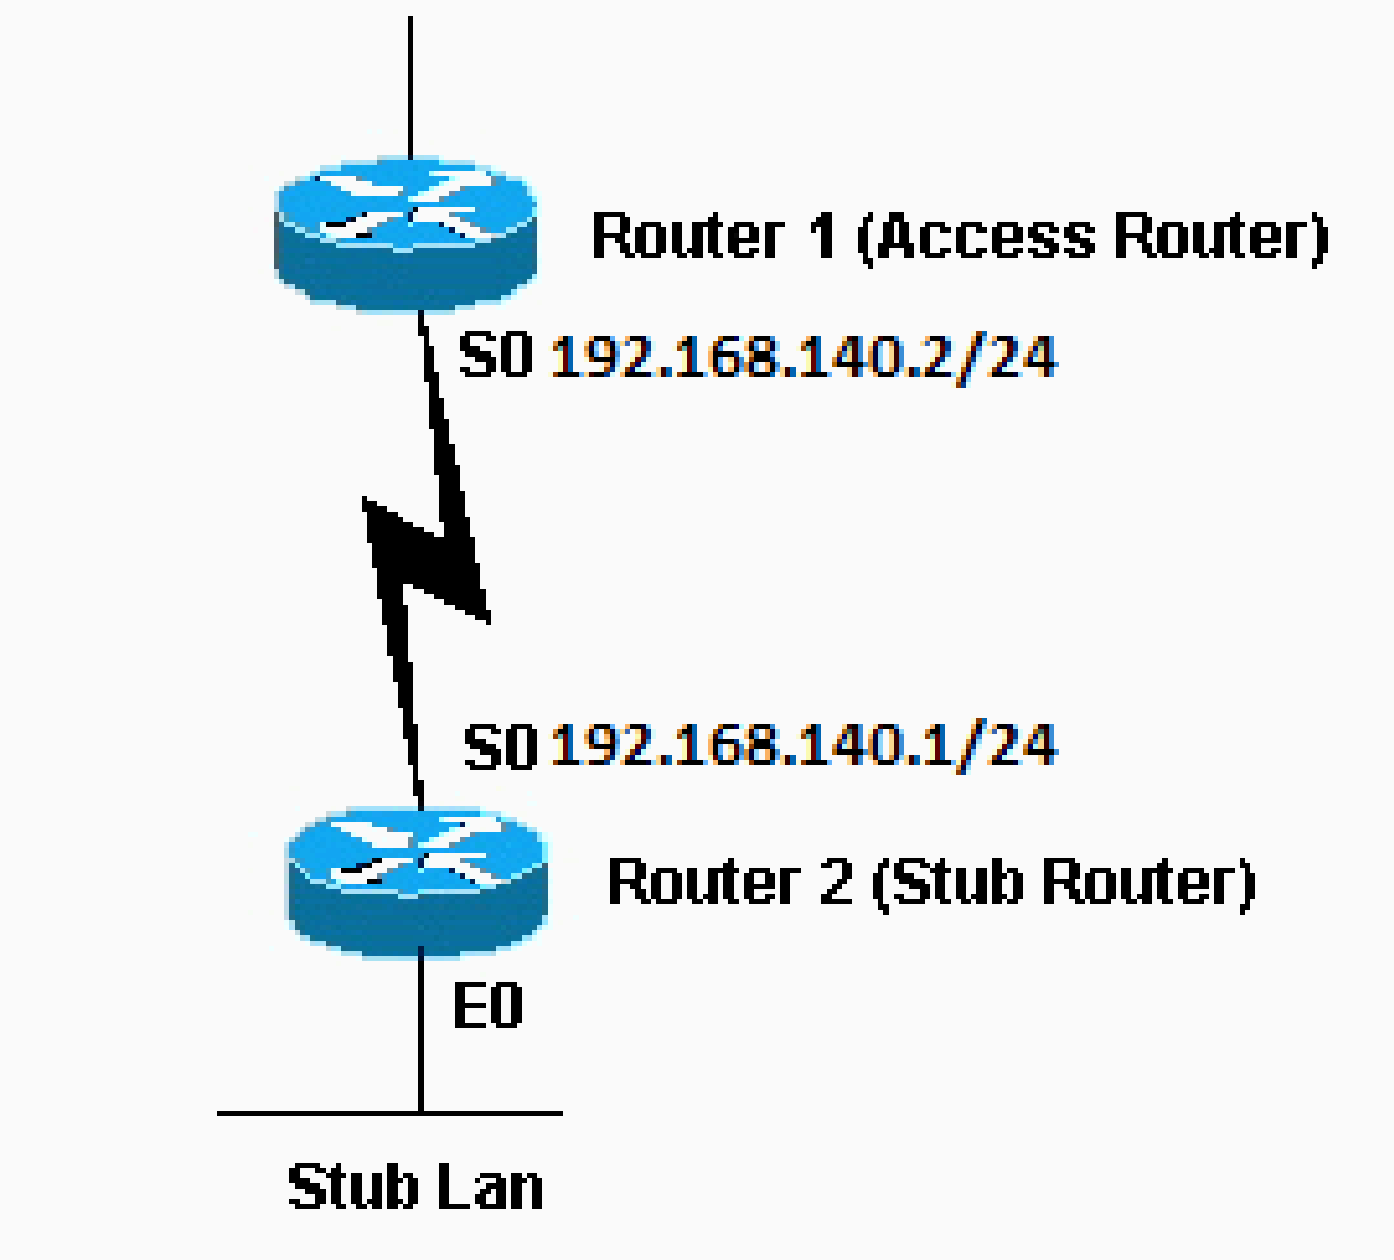

ip msdp peer 192.168.100.1
ip msdp sa-request 192.168.100.1

interface loopback0
ip address <address> <mask>
ip pim sparse-dense-mode

interface serial0
ip address 192.168.100.2 255.255.255.0
ip pim sparse-dense-mode
```

Stub Multicast Routing

Stub multicast routing allows you to configure remote/stub routers as IGMP proxy agents. Rather than fully particpate in PIM, these stub routers forward IGMP messages from the host(s) to the upstream multicast router.



Stub Routers Forward IGMP Messages from the Host(s) to the Upstream Multicast Router

Router 1 Configuration

```
int s0
ip pim sparse-dense-mode
ip pim neighbor-filter 1


access-list 1 deny 192.168.140.1
```

The **ip pim neighbor-filter** command is needed so that Router 1 does not recognize Router 2 as a PIM neighbor.

If you configure Router 1 in sparse mode, the neighbor filter is unnecessary. Router 2 must not run in sparse mode.

When in dense mode, the stub multicast sources can flood to the backbone routers.

Router 2 Configuration

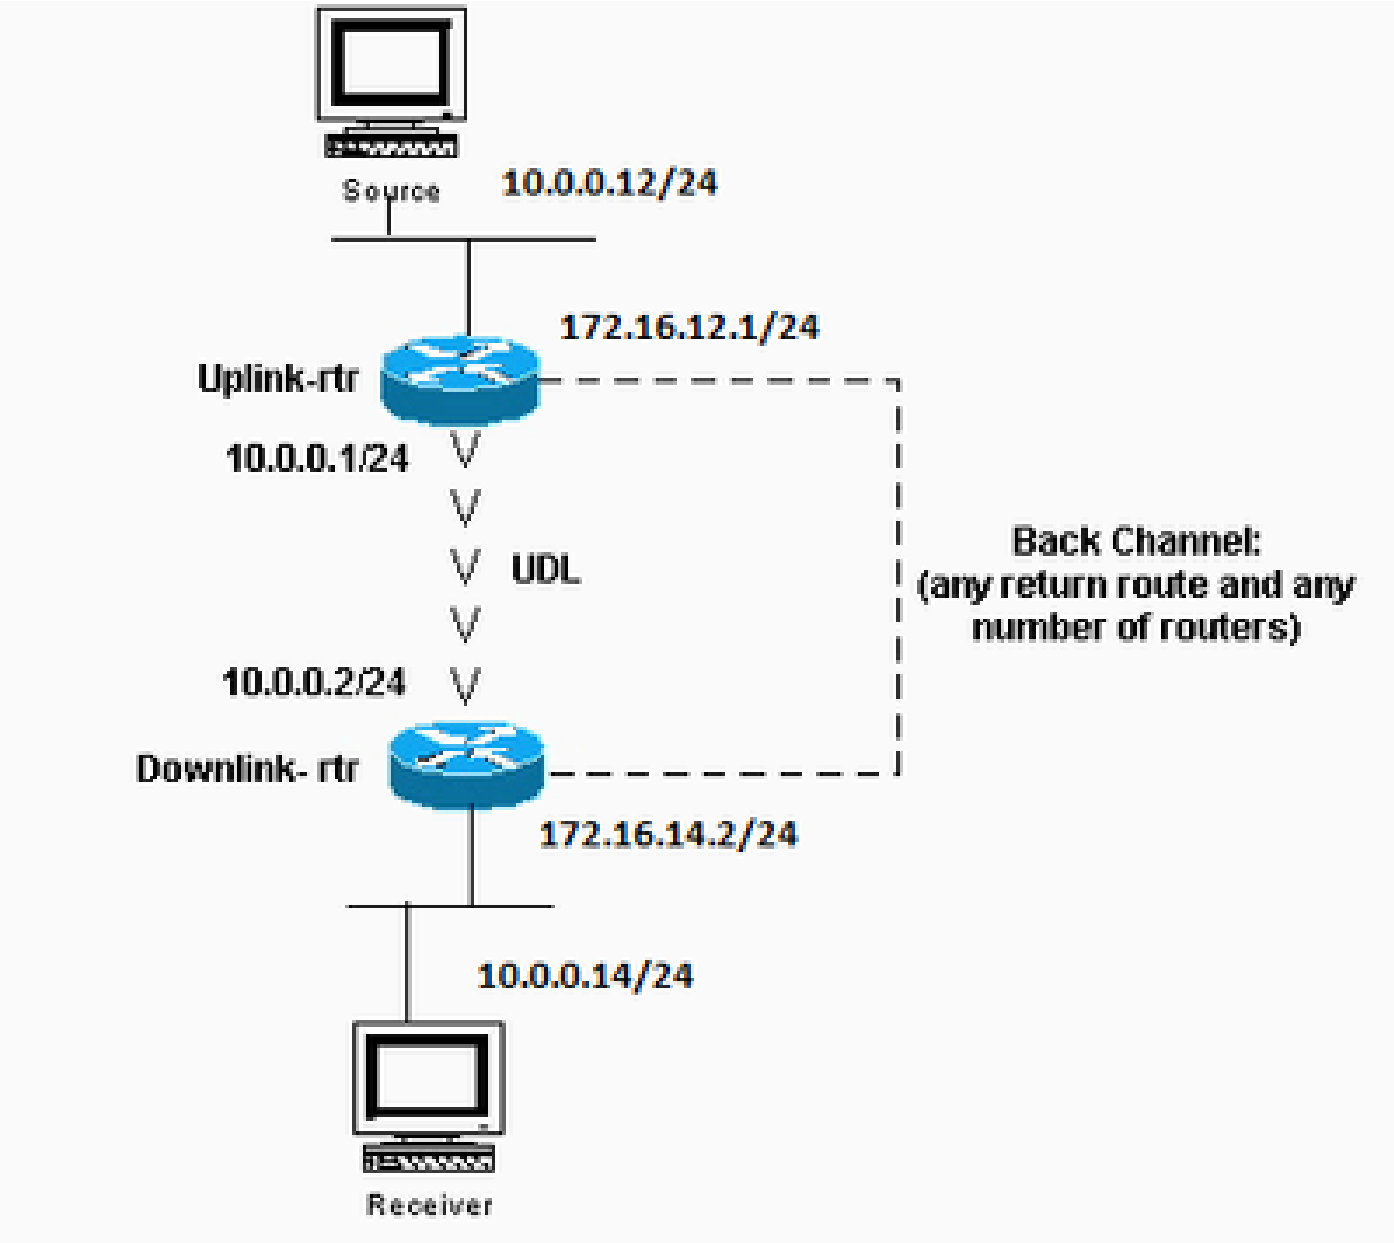
```
ip multicast-routing
int e0
ip pim sparse-dense-mode
ip igmp helper-address 192.168.140.2


int s0
ip pim sparse-dense-mode
```

IGMP UDLR for Satellite Links

Unidirectional Link Routing (UDLR) provides a method for forwarding multicast packets over a unidirectional satellite link to stub networks that have a back channel.

This is similar to stub multicast routing. Without this feature, the uplink router is not able to dynamically learn which IP multicast group addresses to forward over the unidirectional link, because the downlink router cannot send anything back.



Unidirectional Link Routing (UDLR) Provides a Method for forwarding Multicast Packets

Uplink-rtr Configuration
ip multicast-routing interface Ethernet0 description Typical IP multicast enabled interface ip address 172.16.12.1 255.0.0.0 ip pim sparse-dense-mode interface Ethernet1 description Back channel which has connectivity to downlink-rtr ip address 172.16.11.1 255.0.0.0 ip pim sparse-dense-mode interface Serial0 description Unidirectional to downlink-rtr

```
ip address 10.0.0.1 255.0.0.0
ip pim sparse-dense-mode
ip igmp unidirectional-link
no keepalive
```

Downlink-rtr Configuration

```
ip multicast-routing


interface Ethernet0
description Typical IP multicast enabled interface
ip address 172.16.14.2 255.0.0.0
ip pim sparse-dense-mode
ip igmp helper-address udl serial0


interface Ethernet1
description Back channel which has connectivity to downlink-rtr
ip address 172.16.13.2 255.0.0.0
ip pim sparse-dense-mode


interface Serial0
description Unidirectional to uplink-rtr
ip address 10.0.0.2 255.0.0.0
ip pim sparse-dense-mode
ip igmp unidirectional-link
no keepalive
```

PIMv2 BSR

If all routers in the network run PIMv2, you can configure a BSR instead of Auto-RP. BSR and Auto-RP are very similar.

A BSR configuration requires that you configure BSR candidates (similar to RP-Announce in Auto-RP) and BSRs (similar to Auto-RP Mapping Agents). In order to configure a BSR, use these steps:

1. On the candidate BSRs configure:

```
ip pim bsr-candidate interface hash-mask-len pref
```

   Where **interface** contains the candidate BSRs IP address. It is recommended (but not required) that **hash-mask-Len** be identical across all candidate BSRs. A candidate BSR with the largest **pref** value is elected as the BSR for this domain.

   An example of command usage is shown:

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

The PIMv2 BSR collects candidate RP information and disseminates RP-set information associated with each group prefix. In order to avoid single point of failure, you can configure more than one router in a domain as candidate BSRs.

A BSR is elected among the candidate BSRs automatically, based on the preference values configured. In order to serve as candidate BSRs, the routers must be connected and be in the backbone of the network, instead of in the dialup area of the network.

2. Configure candidate RP routers. This example shows a candidate RP, on the interface ethernet0, for the entire admin-scope address range:

```
access-list 11 permit 239.0.0.0 0.255.255.255
ip pim rp-candidate ethernet0 group-list 11
```

CGMP

In order to configure Group Management Protocol (CGMP), configure this on the router interface that faces the switch:

```
ip pim sparse-dense-mode
ip cgmp
```

Then, configure this on the switch:

```
set cgmp enable
```

IGMP Snooping

Internet Group Management Protocol (IGMP) snooping is available with release 4.1 of the Catalyst 5000. IGMP snooping requires a Supervisor III card.

No configuration other than PIM is necessary to configure IGMP snooping on the router. A router is still necessary with IGMP snooping to provide the IGMP querying.

The example provided here shows how to enable IGMP snooping on the switch:

```
<#root>

Console> (enable)

set igmp enable
```

```
IGMP Snooping is enabled.
CGMP is disabled.
```

If you try to enable IGMP, but CGMP is already enabled, you see this:

```
<#root>

Console> (enable)

set igmp enable


Disable CGMP to enable IGMP Snooping feature.
```

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 can detect unrecoverable data packet loss.

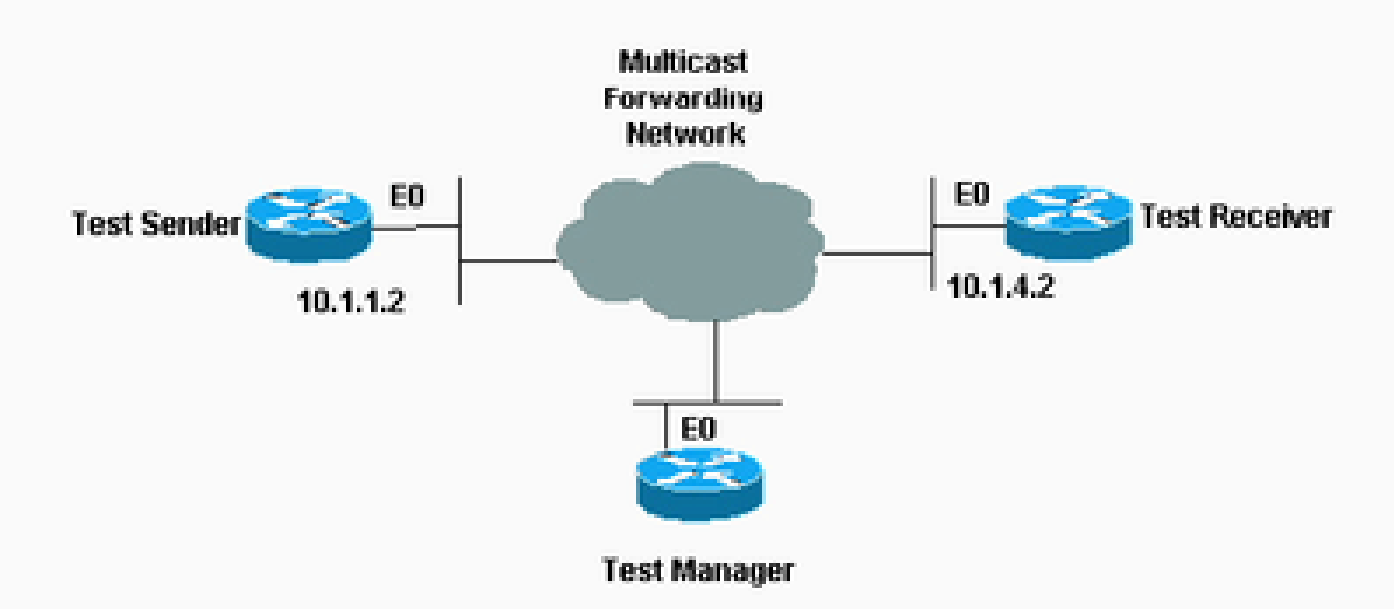
There are no PGM global commands. PGM is configured per interface with the **ip pgm** command. You must enable Multicast routing on the router with PIM on the interface.

MRM

Multicast Routing Monitor (MRM) facilitates automated fault detection in a large multicast routing infrastructure. MRM is designed to alert a network administrator of multicast routing problems near to real-time.

MRM has two components: MRM tester and MRM manager. MRM tester is a sender or receiver.

MRM is available in Cisco IOS Software Release 12.0(5)T and later. Only the MRM testers and managers need to run the MRM-supported Cisco IOS version.



Multicast Routing Monitor (MRM) Facilitates Automated Fault Detection in a Large Multicast Routing Infrastructure

Test Sender Configuration

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

Test Receiver Configuration

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

Test Manager Configuration

```
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
```

Output from the **show ip mrm manager** command on Test Manager is shown here:

```
<#root>

Test_Manager#

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,
```

Start the test with the command shown here. The test manager sends control messages to the test sender and the test receiver as configured in the test parameters.

The test receiver joins the group and monitors test packets sent from the test sender.

```
<#root>

Test_Manager#

mrm start test1


 *Feb  4 10:29:51.798: IP MRM test test1 starts ......
Test_Manager#
```

In order to display a status report for the test manager, enter this command:

```
<#root>

Test_Manager#

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
Test_Manager#
```

The output shows that the receiver sent two status reports (one line each) at a given time stamp. Each report contains one packet loss during the interval window (default of one second).

The "Ehsr" value shows the estimated next sequence number value from the test sender. If the test receiver sees duplicate packets, it shows a negative number in the "Pkt Loss/Dup" column.

In order to stop the test, enter this command:

```
<#root>

Test_Manager#

mrm stop test1


*Feb  4 10:30:12.018: IP MRM test test1 stops
Test_Manager#
```

While the test runs, the MRM sender sends RTP packets to the configured group address at the default interval of 200 ms.

The receiver monitors (expects) the same packets at the same default interval.

If the receiver detects a packet loss in the default window interval of five seconds, it sends a report to the MRM manager.

You can display the status report from the receiver if you issue the **show ip mrm status** command on the manager.

Troubleshoot

Some of the most common problems found when you implement IP multicast in a network are when the router does not forward multicast traffic because of either a RPF failure or TTL settings.

Refer to the IP Multicast Troubleshooting Guide for a detailed discussion about these and other common problems, symptoms, and resolutions.

Related Information

- **IP Multicast Troubleshooting Guide**
- **Troubleshoot Multicast Networks with CLI Tools**
- **IP Multicast Support**
- **Cisco Technical Support & Downloads**