

Rendezvous Point Engineering

Last updated: November 2008

Introduction

A Rendezvous Point (RP) is a router in a multicast network domain that acts as a shared root for a multicast shared tree. Any number of routers can be configured to work as RPs and they can be configured to cover different group ranges. For correct operation, every multicast router within a Protocol Independent Multicast (PIM) domain must be able to map a particular multicast group address to the same RP.

This document attempts to describe, compare and contrast the different methods for deploying RPs. For details on how to configure an RP using any of the methods discussed in this document please refer to the following White Paper: [Configuring a Rendezvous Point](#).

This document discusses the following mechanisms:

- Static RP
- Bootstrap Router (BSR)
- Auto-RP
- Anycast-RP
- Phantom RP
- Embedded RP

Static RP

It is possible to statically configure an RP for a multicast group range. The address of the RP must be configured on every router in the domain. Configuring static RPs is relatively easy and can be done with one or two lines of configuration on each router. If the network does not have many different RPs defined and/or they do not change very often, then this could be the simplest method to define RPs. This can also be an attractive option if the network is small.

However this can be a laborious task in a large and complex network. Every router must have the same RP address. This means changing the RP address requires reconfiguring every router. If several RPs are active for different groups, then information regarding which RP is handling which group must be known by all routers. To ensure this information is complete, several configuration commands may be required. If the manually configured RP fails, there is no failover procedure for another router to take over the function performed by the failed RP. This method does not provide any kind of load-balancing. Static RP can be combined with Anycast RP to provide RP load sharing and redundancy.

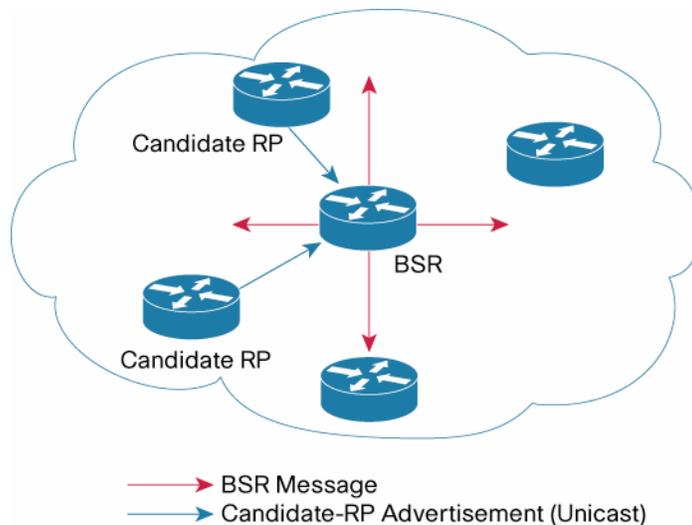
Static RP can co-exist with dynamic RP mechanisms (ie: Auto-RP). Dynamically learned RP takes precedence over manually configured RPs. If a router receives Auto-RP information for a multicast group that has manually configured RP information, then the Auto-RP information will be used.

BSR

The Bootstrap Router (BSR) is a mechanism for a router to learn RP information. It ensures that all routers in the PIM domain have the same RP cache as the BSR. It is possible to configure the BSR to help select an RP set from BSR candidate RPs. The function of the BSR is to broadcast the RP set to all routers in the domain.

Figure 1 shows the BSR mechanism, where the elected BSR sends BSR messages with RP set information out on all enabled interfaces, which are flooded hop-by-hop to all routers in the network. The candidate-RPs send their candidate-RP advertisements directly to the elected BSR.

Figure 1. BSR Mechanism



The elected BSR receives candidate-RP messages from all candidate-RPs in the domain. The bootstrap message sent by the BSR includes information about all the candidate-RPs. Each router uses a common algorithm to select the same RP address for a given multicast group.

For more information about bootstrap routers, see [RFC5059](https://www.rfcs.org/rfcs/rfc5059).

The BSR mechanism is a nonproprietary method of defining RPs that can be used with third-party routers (which support the BSR mechanism). There is no configuration necessary on every router separately (except on candidate-BSRs and candidate-RPs). The mechanism is largely self-configuring making it easier to modify RP information. Information regarding several RPs for different groups is automatically communicated to all routers, reducing administrative overhead. The mechanism is robust to router failover and permits back-up RPs to be configured. In case of RP failure, the secondary RP for the group can take over as the RP for the group.

Auto-RP and BSR protocols must not be configured together in the same network.

Auto-RP

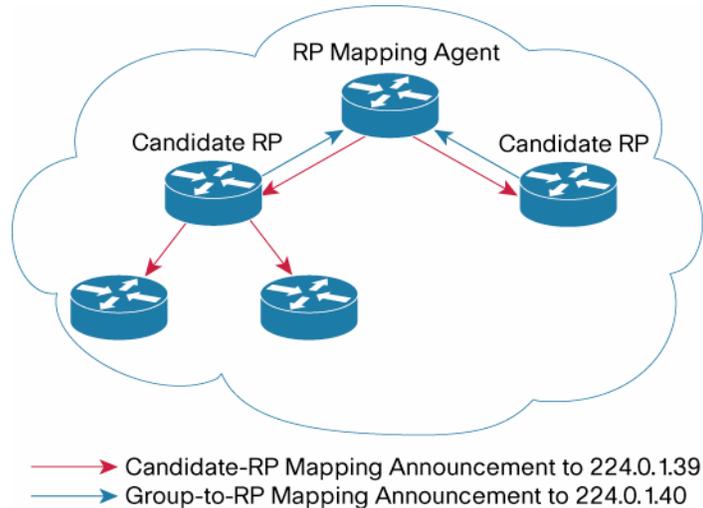
Auto-RP is a mechanism to automate distribution of RP information in a multicast network. The Auto-RP mechanism operates using two basic components, the candidate RPs and the RP mapping agents.

- Candidate RPs advertise their willingness to be an RP via “RP-announcement” messages. These messages are periodically sent to a reserved well-known group 224.0.1.39 (CISCO-RP-ANNOUNCE).

- RP mapping agents join group 224.0.1.39 and map the RPs to the associated groups. The RP mapping agents advertise the authoritative RP-mappings to another well-known group address 224.0.1.40 (CISCO-RP-DISCOVERY). All PIM routers join 224.0.1.40 and store the RP-mappings in their private cache.

Figure 2 shows the Auto-RP mechanism where the RP mapping agent periodically multicasts the RP information that it receives to the Cisco-RP-Discovery group.

Figure 2. Auto-RP Mechanism



All routers automatically learn the RP information making it easier to administer and update RP information. There is no configuration needed on every router separately (except on candidate RPs and mapping agents). Auto-RP permits back-up RPs to be configured enabling an RP failover mechanism.

Auto-RP is not supported for IPv6 Multicast. Auto-RP is a Cisco proprietary mechanism.

BSR and Auto-RP protocols must not be configured together in the same network.

Anycast-RP

Anycast RP is used to define redundant and load-balanced RPs. Anycast-RP has two implementations:

- Using Multicast Source Discovery Protocol (MSDP) described in [RFC3446](#).
- Using Protocol Independent Multicast (PIM) described in [RFC4610](#).

Anycast-RP allows two or more RPs to share the load for source registration and to act as hot back-up routers for each other. Multicast Source Discovery Protocol (MSDP) is the protocol RPs use to share information about active sources. With Anycast RP, the RPs are configured to establish MSDP peering sessions using a TCP connection. Group participants use the closest RP that is favored by the IP unicast route table.

Figure 3. Anycast-RP Mechanism

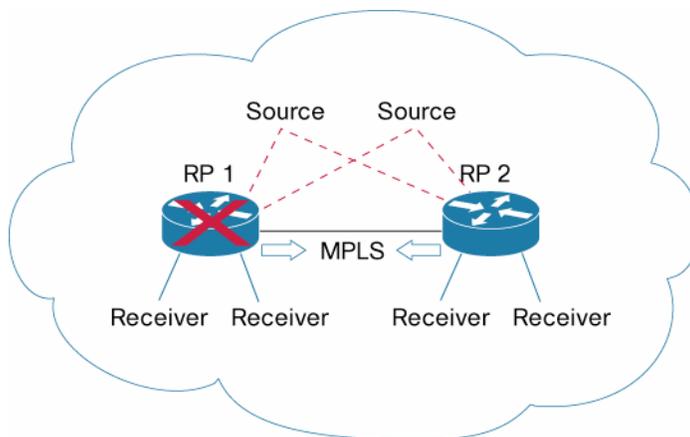


Figure 3 shows the Anycast-RP mechanism. Two or more RPs are configured with the same IP address on loopback interfaces. All the downstream routers are configured to “know” that the Anycast RP loopback address is the IP address of their RP. When an RP fails, IP routing converges and the other RP assumes the RP role. Further details on Anycast-RP can be found in the following White Paper: [Anycast RP Overview](#).

PIM Anycast-RP extends the Register mechanism in PIM so that Anycast-RP functionality can be retained without using MSDP. PIM register messages are sent to the closest RP and PIM join-prune messages are sent in the direction of the closest RP as determined by the unicast routing protocols. If one of the RPs goes down, unicast routing ensures these messages will be sent in the direction of the next-closest RP.

Anycast-RP provides RP failover and redundancy. Anycast-RP can be used along with static RP or Auto-RP to provide RP redundancy.

Phantom RP

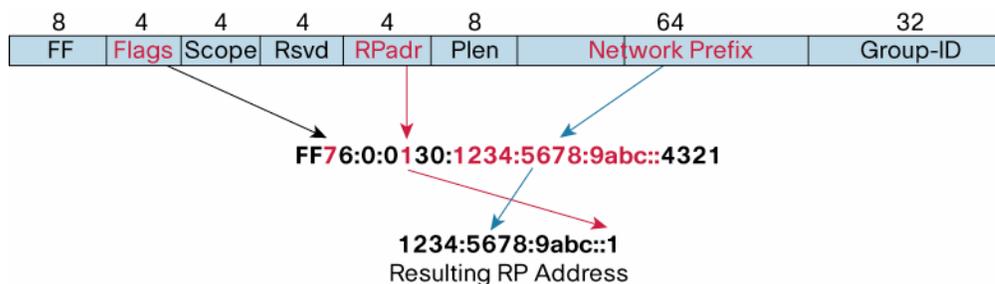
In Bidirectional PIM (Bidir-PIM), the RP does not have an actual protocol function. The RP acts as a routing vector in which all the traffic converges. The RP can be configured as an address that is not assigned to any particular device called a Phantom RP. This means that the RP address does not need to reside on a physical router interface, but can just be an address in a subnet. The RP can also be a physical router, but it is not necessary.

More details on Phantom RP can be found in the [Bidirectional PIM Deployment Guide](#).

Embedded RP

Embedded RP defines an address allocation policy in which the address of the RP is encoded in an IPv6 multicast group address. This allows an easy deployment of scalable inter-domain multicast and simplifies the intra-domain multicast configuration as well. IPv6 Multicast group addresses embedded with RP information start with `ff70::/12` where the flag value of **7** means embedded RP.

Figure 4. Multicast Address with Embedded RP



There is no need to pre-configure routers with the RP address information. Routers can automatically extract and use the RP information from the IPv6 multicast group address. This allows for a large number of RPs to be deployed anywhere in the Internet. Embedded RP requires no change in protocol operations. It can be considered an automatic replacement for static RP configuration.

The router can learn only one RP address for a multicast group using embedded RP. It cannot support RP redundancy. Proposals are being considered to introduce RP redundancy by mechanisms other than BSR for IPv6 multicast. Embedded RP does not support Bidirectional PIM. Embedded RP allows the application to dictate which router is the RP. There is the possibility that a low-end router could end up becoming the RP for hundreds of high data rate sources if the application defines an erroneous RP address (this can be prevented by disabling Embedded RP learning).

For more information on Embedded RP, see [RFC3956](#).

Conclusion

There are various mechanisms for disseminating RP information across a multicast network. Each mechanism has different pros and cons. Table 1-1 compares the approaches discussed in this document.

Table 1. Comparison of RP Mechanisms

	Static RP	BSR	Auto-RP	Embedded RP
Must be configured on every router	Yes	No (except on candidate-BSRs and candidate-RPs)	No (except on candidate RPs and mapping agents)	No (except RP routers)
Supports IPv4 addresses	Yes	Yes	Yes	No
Supports IPv6 addresses	Yes	Yes	No	Yes
RP redundancy	No (unless used with Anycast RP)	Yes	Yes	No

Anycast RP can be used to provide redundancy and load-sharing capabilities. Phantom RP can be an attractive solution for Bidirectional PIM.

Related Documents

IP Multicast Technology Overview, Cisco White Paper:

http://www.cisco.com/en/US/docs/ios/solutions_docs/ip_multicast/White_papers/mcst_ovr.html

Configuring a Rendezvous Point, Cisco White Paper:

http://www.cisco.com/en/US/docs/ios/solutions_docs/ip_multicast/White_papers/rps.html

Anycast RP White Paper:

http://www.cisco.com/en/US/docs/ios/solutions_docs/ip_multicast/White_papers/anycast.html

"Configuring Multicast Source Discovery Protocol," Cisco IOS IP Configuration Guide, Release 12.2:

http://www.cisco.com/en/US/docs/ios/12_2/ip/configuration/guide/1cfmsdp_ps1835_TSD_Products_Configuration_Guide_Chapter.html

Bidirectional PIM Deployment Guide:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6552/ps6592/prod_white_paper090_0aecd80310db2.pdf



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0903R)