



Anycast RP

Version History

Version Number	Date	Notes
1	06/18/2001	This document was created.
2	10/15/2001	“Anycast RP Example” section was updated.
3	11/19/2001	Figure 2 was updated.

IP multicast is deployed as an integral component in mission-critical networked applications throughout the world. These applications must be robust, hardened, and scalable to deliver the reliability that users demand.

Using Anycast RP is an implementation strategy that provides load sharing and redundancy in Protocol Independent Multicast sparse mode (PIM-SM) networks. Anycast RP allows two or more rendezvous points (RPs) to share the load for source registration and the ability to act as hot backup routers for each other. Multicast Source Discovery Protocol (MSDP) is the key protocol that makes Anycast RP possible.

The scope of this document is to explain the basic concept of MSDP and the theory behind Anycast RP. It also provides an example of how to deploy Anycast RP.

This document has the following sections:

- [Multicast Source Discovery Protocol Overview](#)
- [Anycast RP Overview](#)
- [Anycast RP Example](#)
- [Related Documents](#)

Multicast Source Discovery Protocol Overview

In the PIM sparse mode model, multicast sources and receivers must register with their local rendezvous point (RP). Actually, the router closest to a source or a receiver registers with the RP, but the key point to note is that the RP “knows” about all the sources and receivers for any particular group. RPs in other domains have no way of knowing about sources located in other domains. MSDP is an elegant way to solve this problem.


MSDP is a mechanism that allows RPs to share information about active sources. RPs know about the receivers in their local domain. When RPs in remote domains hear about the active sources, they can pass on that information to their local receivers. Multicast data can then be forwarded between the domains.

A useful feature of MSDP is that it allows each domain to maintain an independent RP that does not rely on other domains, but it does enable RPs to forward traffic between domains. PIM-SM is used to forward the traffic between the multicast domains.

The RP in each domain establishes an MSDP peering session using a TCP connection with the RPs in other domains or with border routers leading to the other domains. When the RP learns about a new multicast source within its own domain (through the normal PIM register mechanism), the RP encapsulates the first data packet in a Source-Active (SA) message and sends the SA to all MSDP peers. Each receiving peer uses a modified Reverse Path Forwarding (RPF) check to forward the SA, until the SA reaches every MSDP router in the interconnected networks—theoretically the entire multicast internet. If the receiving MSDP peer is an RP, and the RP has a (*, G) entry for the group in the SA (there is an interested receiver), the RP creates (S, G) state for the source and joins to the shortest path tree for the source. The encapsulated data is decapsulated and forwarded down the shared tree of that RP. When the last hop router (the router closest to the receiver) receives the multicast packet, it may join the shortest path tree to the source. The MSDP speaker periodically sends SAs that include all sources within the domain of the RP. [Figure 1](#) shows how data would flow between a source in domain A to a receiver in domain E.

MSDP was developed for peering between Internet service providers (ISPs). ISPs did not want to rely on an RP maintained by a competing ISP to provide service to their customers. MSDP allows each ISP to have its own local RP and still forward and receive multicast traffic to the Internet.

Figure 1 MSDP Example: MSDP Shares Source Information Between RPs in Each Domain



Anycast RP Overview

Anycast RP is a useful application of MSDP. Originally developed for interdomain multicast applications, MSDP used for Anycast RP is an intradomain feature that provides redundancy and load-sharing capabilities. Enterprise customers typically use Anycast RP for configuring a Protocol Independent Multicast sparse mode (PIM-SM) network to meet fault tolerance requirements within a single multicast domain.

In Anycast RP, two or more RPs are configured with the same IP address on loopback interfaces. The Anycast RP loopback address should be configured with a 32-bit mask, making it a host address. All the downstream routers should be configured to “know” that the Anycast RP loopback address is the IP address of their local RP. IP routing automatically will select the topologically closest RP for each source and receiver. Assuming that the sources are evenly spaced around the network, an equal number of sources will register with each RP. That is, the process of registering the sources will be shared equally by all the RPs in the network.

Because a source may register with one RP and receivers may join to a different RP, a method is needed for the RPs to exchange information about active sources. This information exchange is done with MSDP.

In Anycast RP, all the RPs are configured to be MSDP peers of each other. When a source registers with one RP, an SA message will be sent to the other RPs informing them that there is an active source for a particular multicast group. The result is that each RP will know about the active sources in the area of the other RPs. If any of the RPs were to fail, IP routing would converge and one of the RPs would become the active RP in more than one area. New sources would register with the backup RP. Receivers would join toward the new RP and connectivity would be maintained.

Note that the RP is normally needed only to start new sessions with sources and receivers. The RP facilitates the shared tree so that sources and receivers can directly establish a multicast data flow. If a multicast data flow is already directly established between a source and the receiver, then an RP failure will not affect that session. Anycast RP ensures that new sessions with sources and receivers can begin at any time.

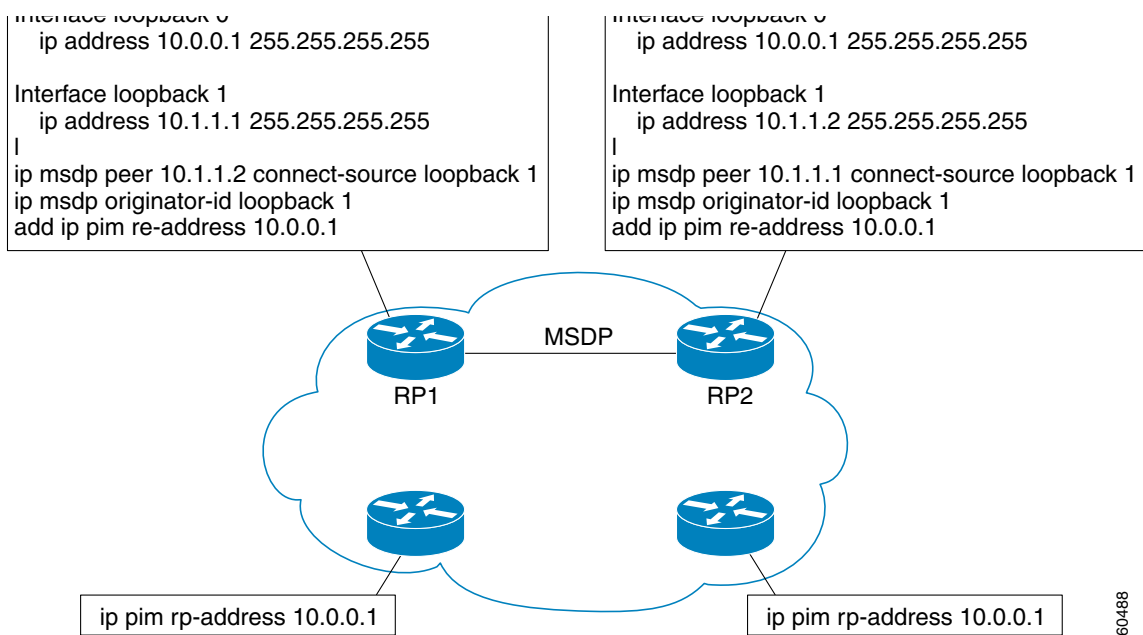
Anycast RP Example

The main purpose of an Anycast RP implementation is that the downstream multicast routers will “see” just one address for an RP. The example given in [Figure 2](#) shows how the loopback 0 interface of the RPs (RP1 and RP2) is configured with the same 10.0.0.1 IP address. If this 10.0.0.1 address is configured on all RPs as the address for the loopback 0 interface and then configured as the RP address, IP routing will converge on the closest RP. This address must be a host route—note the 255.255.255.255 subnet mask.

The downstream routers must be informed about the 10.0.0.1 RP address. In [Figure 2](#), the routers are configured statically with the `ip pim rp-address 10.0.0.1` global configuration command. This configuration could also be accomplished using the Auto-RP or bootstrap router (BSR) features.

The RPs in [Figure 2](#) must also share source information using MSDP. In this example, the loopback 1 interface of the RPs (RP1 and RP2) is configured for MSDP peering. The MSDP peering address must be different than the Anycast RP address.

Figure 2 Anycast RP Configuration



60488

Many routing protocols choose the highest IP address on loopback interfaces for the Router ID. A problem may arise if the router selects the Anycast RP address for the Router ID. We recommend that you avoid this problem by manually setting the Router ID on the RPs to the same address as the MSDP peering address (for example, the loopback 1 address in [Figure 2](#)). In Open Shortest Path First (OSPF), the Router ID is configured using the **router-id** router configuration command. In Border Gateway Protocol (BGP), the Router ID is configured using the **bgp router-id** router configuration command. In many BGP topologies, the MSDP peering address and the BGP peering address must be the same in order to pass the RPF check. The BGP peering address can be set using the **neighbor update-source** router configuration command.

The Anycast RP example in the previous paragraphs used IP addresses from RFC 1918. These IP addresses are normally blocked at interdomain borders and therefore are not accessible to other ISPs. You must use valid IP addresses if you want the RPs to be reachable from other domains.

Related Documents

- *IP Multicast Technology Overview*, Cisco white paper
http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/mcst_sol/mcst_ovr.htm
- *Interdomain Multicast Solutions Using MSDP*, Cisco integration solutions document
http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/mcst_p1/mcstmsdp/index.htm
- *Configuring a Rendezvous Point*, Cisco white paper
http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/mcst_sol/rps.htm

- “Configuring Multicast Source Discovery Protocol,” *Cisco IOS IP Configuration Guide*, Release 12.2
http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcpt3/1cfmsdp.htm
- “Multicast Source Discovery Protocol Commands,” *Cisco IOS IP Command Reference, Volume 3 of 3: Multicast*, Release 12.2
http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fiprmc_r/1rfmsdp.htm

