Protocol Independent Multicast-Sparse Mode (PIM-SM): Deployment Guidelines

Deborah Estrin
Computer Science Dept/ISI
University of Southern Calif.
Los Angeles, CA 90089
estrin@usc.edu

Ahmed Helmy
Computer Science Dept/ISI
University of Southern Calif.
Los Angeles, CA 90089
ahelmy@catarina.usc.edu

David Thaler
EECS Department
University of Michigan
Ann Arbor, MI 48109
thalerd@eecs.umich.edu

Liming Wei
Cisco Systems Inc.
170 West Tasman Drive,
San Jose, CA 95134
lwei@cisco.com

draft-ietf-idmr-PIM-SM-guidelines-00.ps

Dec 10, 1996

Status of This Memo

Abstract

This document provides guidelines and recommendations for the incremental deployment of Protocol Independent Multicast-Sparse Mode (PIM-SM) [1, 2].

Failure to comply with the *requirements* in this document, may affect the proper operation of the multicast routing protocol. Failure to follow the *recommendations* presented in the document, does not affect the proper operation (i.e. correctness) of the multicast protocol.

Familiarity with PIM terminology is assumed throughout the document.

For description of the terms used, refer to section 4 at the end of the document.

1 PIM-SM Multicast Domain

A *multicast domain* (or simply *domain*) in this document refers to a region of multicast routers implementing the same multicast protocol, and bounded by multicast border router(s) (MBRs), as described in [3, 4].

For a *PIM-SM multicast domain* (or *PIM domain*), the multicast routers implement PIM-SM as described in [2].

1.1 Contiguity Requirement

A PIM domain consists of an "all PIM routers" domain, meaning that all the routers within the domain must implement (the same version of) PIM-SM. This way a PIM domain will have robust multicast connectivity.

For example, a Join/Prune message will not be sent upstream if the upstream router is PIM non-capable (i.e. not a PIM neighbor).

A domain that contains some PIM non-capable routers will not have robust multicast connectivity. Valid multicast paths in such a domain are those having contiguous PIM routers congruent with the underlying unicast routing paths.

1.2 Convexity Requirement

In general, within administrative bounds, domains should be convex with respect to unicast routing.

A PIM domain must be convex with respect to unicast routing. By *convex* we mean that all routers within the PIM domain must have PIM routers within the domain as reverse path forwarding (RPF) neighbors towards any other PIM router within the same domain.

If the convexity requirement is not satisfied, PIM Bootstrap messages may not reach all routers within the domain, and group partitions are likely to occur. A *group partition* occurs when different group participants (i.e. senders and receivers) do not map to the same RP, and hence fail to rendezvous.

In addition, the multicast distribution paths built by PIM Join/Prune messages must form a single tree, for any (S,G), (*,G) or (*,*,RP). Violation of convexity requirement within a domain, may cause a path for PIM Join/Prune messages to cross domain boundaries, creating either multicast tree overlaps (i.e. a graph) or discontinuities within the multicast distribution paths.

1.3 Administrative Boundaries

Administrative boundaries provide a means to control multicast traffic. Such control is imposed based upon the multicast group address. Currently, multicast addresses are classified as either local or non-local. Local groups are typically allocated from the administratively scoped IP multicast address space, as defined by [5]. These groups typically belong to the range 239.0.0.0 to 239.255.255.255.

Different administratively scoped multicast regions/domains are connected through physical interfaces using the multicast border routers, or logical interfaces (i.e. tunnels) leading to regions outside of the administratively scoped region. These routers are configured to prevent local multicast traffic from being forwarded outside of the administrative region. For PIM domains, these routers must coincide with the PIM multicast border routers (PMBRs).

2 PIM-SM Router Configuration

It is assumed that a configuration method is supported by the PIM implementation. The configuration method may be provided by an 'SNMP MIB' through the PIM MIB [6], or a configuration file similar to that described in [7]. In this section, we use the latter method merely to illustrate the configuration by example, and we refer to the former method occasionally.

Configurable parameters considered here pertain to Candidate RPs (C-RPs), Candidate Bootstrap Routers (C-BSRs), and switching-to-the-shortest-path thresholds.

2.1 Configuring C-RPs and C-BSRs

For the proper operation of PIM-SM, each PIM multicast domain (or partition in case of failures), must contain at least a C-RP and a C-BSR. In general the candidacy for RP or BSR is expressed in terms of router setup. Such set up may be done implicitly (i.e. by default), or through explicit configuration, overriding the default behavior.

2.1.1 Configuration Settings

Following is a list of recommended default router behavior, to support proper protocol operation without the need for configuration. In addition, the list includes guidelines for configuring the routers to override the default behavior.

- default configuration of deployed routers is recommended to have the following configuration:

  ```
  ''C-BSR priority 0
   C-RP if BSR and no other C-RPs''
  ```

  In the PIM MIB, the C-BSR priority is pimCandidateBSRPreference, and a C-RP is indicated by a non-zero pimCandidateRPHoldTime.

  This way a PIM domain (or partition), after converging on a single elected BSR, will have one C-RP that is the highest addressed router within the domain, supporting all multicast groups (i.e. whole class D address space).

  Once another C-RP is configured and starts sending C-RP-Advs, the BSR refrains from being a C-RP by default, if the C-RP-Advs indicate the whole multicast address space.

- a router configured to be a C-BSR should be configured with a priority higher than '0', to override the default BSR.

- a router configured as a C-RP, defaults to advertising its candidacy for all group prefixes, unless otherwise specified. For example, a PIM router configured as a C-RP, as follows:

  ```
  ''C-RP''
  ```

  is implicitly configured to support all group prefixes (i.e. class-D address).

  On the other hand, a C-RP configured to support only local groups may be specified as follows:

  ```
  ''C-RP prefix 239.0.0.0/8''
  ```

- configured C-BSRs and C-RPs should be well connected stable routers

- in a domain having one MBone connection [i.e. one MBR], configured C-RPs for non-local groups should be located at or near the MBR, to minimize the number of links on the (*,*,RP) tree. Locating the C-RP for non-local groups at the PMBR obviates the need for (*,*,RP) trees.

2.1.2 Number of Configured Routers

This section provides recommended maximum/minimum number of configured C-RPs and C-BSRs.

[This section is incomplete].

2.2 Switching to the shortest path trees (SPTs)

The current PIM-SMv2 spec [2] proposes a heuristic for switching to the SPT based upon data rate. Two thresholds are suggested for switching to the SPT. The first threshold is for the RP, and pertains to the rate of registers received for a specific source-group pair. The second threshold is for the last-hop router, and pertains to the data rate for a specific source-group pair.

In the interoperability context, for robustness reasons, the threshold for switching to the SPT for external sources at the RP is set to '0 Kbps' (i.e. switch to the SPT with the first PIM Register with the 'border bit' set).

For internal sources, and future inter-domain PIM, this threshold is configurable. In this case, a recommended threshold to use is '8 Kbps' sustained for 5 seconds (or 5 KBytes within a 5 second period). This threshold refers to the data encapsulated within PIM Registers for a specific (S,G).

The other threshold is configured at the last-hop routers. The value recommended is '16 Kbps' sustained for 5 seconds (or 10 KBytes within a 5 second period). This refers to the data rate received for a specific (S,G).

3 Incremental Deployment of PIM-SM

3.1 Converting routers to PIM-SM

It is required to convert an entire LAN at a time to PIM-SM. All routers on a LAN must run the same version of PIM-SM.

The Convexity requirement for unicast routing within a domain must not be violated, however. Converting a LAN to PIM-SM such that the new PIM-SM domain is not convex, may affect the proper operation of the protocol.

Also, it is not allowed to have mixed LANs on domain borders. A PMBR interface should run only one multicast protocol. All neighboring routers on that interface must run the same protocol, and only that protocol [3, 4].

4 Terminology

The following terms are used throughout the document:

- **C-BSR** a Candidate BootStrap Router (C-BSR) is a PIM-SM router having one of its interfaces configured to participate in the BSR election. The configuration includes the address of the interface, in addition to a priority. Each PIM domain must have at least one C-BSR.
- **C-RP** a Candidate Rendezvous Point (C-RP) is a PIM-SM router having one (or more) of its interfaces configured as a candidate RP. Associated with the configuration, is one (or more) address prefix(es) conveying the range of multicast addresses for which the router wishes to advertise its candidacy. Each PIM domain must have at least one C-RP. The C-RPs within a PIM domain must cover the whole multicast address space.
- **Elected BSR** A single C-BSR is elected per multicast domain (or partition). The BSR election is based on the priority and the address of the C-BSR, according to [2].

References

[1] S. Deering, D. Estrin, D. Farinacci, M. Handley, A. Helmy, V. Jacobson, C. Liu, P. Sharma, D. Thaler, and L. Wei. Protocol independent multicast - sparse mode (pim-sm): Motivation and architecture. *Proposed Experimental RFC*, September 1996.

[2] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei. Protocol independent multicast - sparse mode (pim-sm): Protocol specification. *Proposed Experimental RFC*, September 1996.

[3] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei. Pim multicast border router (pmbr) specification for connecting pim-sm domains to a dvmrp backbone. *Internet Draft*, September 1996.

[4] D. Thaler. Interoperability rules for multicast routing protocols. *Internet Draft*, November 1996.

[5] D. Meyer. Administratively scoped ip multicast. *Internet Draft*, November 1996.

[6] K. McCloghrie, D. Farinacci, and D. Thaler. Protocol independent multicast mib. *Internet Draft*, November 1996.

[7] A. Helmy. Protocol independent multicast-sparse mode (pim-sm): Implementation document. *Working Draft*, August 1996.