

PIM WG
Internet Draft
Intended Status: Proposed Standard
Expires: July 7, 2009

Yiqun Cai
Heidi Ou

Cisco Systems, Inc.

January 7, 2009

PIM Multi-Topology ID (MT-ID) Join-Attribute

[draft-ietf-pim-mtid-00.txt](#)

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on July 7, 2009.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Internet Draft

[draft-ietf-pim-mtid-00.txt](#)

January 2009

Abstract

This document introduces a new type of PIM Join Attribute that extends PIM signaling to identify a topology that should be used when constructing a particular multicast distribution tree.

Table of Contents

1	Specification of Requirements	2
2	Introduction	3
3	Functional Overview	3
3.1	PIM RPF Topology	3
3.2	PIM MT-ID	4
3.3	Applicability	4
4	Protocol Specification of PIM MT-ID	5
4.1	Sending PIM MT-ID Join Attribute	5
4.2	Receiving PIM MT-ID Join Attribute	5
4.3	Validating PIM MT-ID Join Attribute	6
4.4	Conflict Resolution	6
4.4.1	Upstream Routers	6
4.4.2	Downstream Routers	7
5	PIM MT-ID Join Attribute TLV Format	7
6	IANA Considerations	8
7	Security Considerations	8
8	Acknowledgments	8
9	Authors' Addresses	8
10	Normative References	9
11	Informative References	9

[1](#). Specification of Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

Internet Draft

[draft-ietf-pim-mtid-00.txt](#)

January 2009

[2. Introduction](#)

Some unicast protocols, such as OSPF and IS-IS, allow a single network to be viewed as multiple topologies [RFC4915, [RFC5120](#)]. This enables PIM to construct multicast distribution trees using separate network paths even when the roots of the trees are the same.

This capability can be used to improve the resilience of multicast applications. For instance, a multicast stream can be duplicated and transported using different network layer addresses simultaneously. Assuming that two source trees, (S1, G1) and (S1, G2), are used for the stream. By using MT capable unicast routing protocols and procedures described in this document, it is possible to construct two source trees for (S1, G1) and (S1, G2) in such a way that they do not share any transit network segment. As a result, a single network failure will not cause any loss to the stream.

This draft introduces a new type of PIM Join Attribute used to encode the identity of the topology PIM uses for RPF. It is based on [[RFC5384](#)], and specifies additional procedures and rules to process the attribute and resolve conflict.

[3. Functional Overview](#)

[3.1. PIM RPF Topology](#)

PIM RPF topology is a collection of routes used by PIM to perform RPF operation when building shared or source trees. In the rest of the document, PIM RPF topology may be simply referred to as "topology" when there is no ambiguity.

In a multi-topology environment, multiple RPF topologies can be created in the same network. A particular source may be reachable in only one of the topologies, or in several of them via different paths.

To select the RPF topology for a particular multicast distribution tree, one or more of the following can be done.

1. configure a policy that maps a group range to a topology. When RPF information needs to be resolved for the RP or the sources for a group within the range, the RPF lookup takes place in the specified topology. This can be used for PIM-SM/SSM/Bidir.

2. configure a policy that maps a source prefix range to a topology. This can be used for PIM-SM and PIM-SSM.
3. use the topology identified by the Join Attribute encoding in the received PIM packets.

The details of the first two methods are implementation specific and are not discussed in this document. The specification to support the third method is included in this document.

[3.2.](#) PIM MT-ID

For each PIM RPF topology created, a unique numerical ID is assigned. This ID is called PIM MT-ID. PIM MT-ID has the following property,

- this value is not required to be the same as the MT-ID used by the unicast routing protocols that contribute routes to the topology. Although in practice, when only one unicast routing protocol (such as OSPF or IS-IS) is used, PIM MT-ID is typically assigned the same value as the IGP topology identifier.
- this value must be unique and consistent within the network domain for the same topology
- 0 is reserved as the default, and MUST NOT be included in the join attribute encoding.

- how to assign a PIM MT-ID to a topology is decided by the network administrator and is outside the scope of this document

3.3. Applicability

The PIM MT-ID join attribute described in this draft applies to PIM Join/Assert packets used by PIM SM/SSM/Bidir. It is not used in any other PIM packets, such as Prune, Register, Register-Stop, Graft, Graft-ack, DF Election, Candidate-RP, and Bootstrap. As such, it can only be used to build shared or source trees for PIM SM/SSM and PIM-bidir downstream.

When this attribute is used in combination with RPF vectors defined in [[ID.ietf-pim-rpf-vector](#)] [[ID.ietf-l3vpn-2547bis-mcast](#)], they are processed against the topology identified by the PIM MT-ID attribute.

4. Protocol Specification of PIM MT-ID

4.1. Sending PIM MT-ID Join Attribute

When a PIM router originates a PIM Join/Assert packet, it may choose to encode PIM MT-ID of the topology in which RPF lookup takes place for the corresponding (*,G) or (S,G) entry. The chosen PIM MT-ID MUST be the one decided by local topology selection configuration if it exists, or the one received from downstream routers after conflict resolution procedures are applied.

The following are the exceptions,

- a router MUST NOT attach the attribute if PIM MT-ID is 0. The value of 0 is ignored on reception.
- a router SHOULD NOT do so if the upstream router, or one of the routers on the LAN does not include "PIM Join Attribute" option in its Hello packets.
- a router SHOULD NOT encode PIM MT-ID for pruned sources. If encoded, the value is ignored.

[4.2.](#) Receiving PIM MT-ID Join Attribute

When a PIM router receives a PIM MT-ID join attribute in a Join/Assert packet, it MUST perform the following,

- validate the attribute encoding. The detail is described in the next section.
- if the join attribute is valid, use the rules described in the section "Conflict Resolution" to determine a PIM MT-ID to use.
- use the topology identified by the selected PIM MT-ID to perform RPF lookup for the (*,G)/(S,G) entry unless a different topology is specified by a local configuration. The local configuration always takes precedence.

[4.3.](#) Validating PIM MT-ID Join Attribute

An encoded PIM MT-ID join attribute is valid if all of the following conditions are satisfied,

- there is at most 1 PIM MT-ID attribute encoded.
- the length field must be 2 and the value must not be 0.

If an encoded PIM MT-ID join attribute is deemed invalid, it is ignored and not forwarded further. The packet is processed as if the attribute were not present.

It is important to note that, if the sender is not a PIM neighbor

that has included "PIM Join Attribute" option in its Hello packets, or if the "F" bit in the encoding is reset, the encoding may still be considered valid by an implementation and is allowed to be forwarded.

[4.4. Conflict Resolution](#)

Depending on whether a PIM router is an upstream or a downstream router, the action it takes to resolve conflicting PIM MT-ID attributes differs. The detail is described below.

[4.4.1. Upstream Routers](#)

If an upstream router has a local configuration that specifies a different topology than that from an incoming Join/Assert packet, including the case PIM MT-ID is not encoded in the incoming packet, it is not considered a conflict.

A conflict occurs when a router doesn't have local topology selection policy and it has received different PIM MT-ID from Join packets sent by its downstream routers or Assert packets from another forwarding router on the LAN.

- if an upstream router receives different PIM MT-ID attributes from PIM Join packets, it MUST follow the rules specified in [\[RFC5384\]](#) to select one. The PIM MT-ID chosen will be the one encoded for its upstream neighbor.
- if an upstream router receives a different PIM MT-ID attribute in an ASSERT packet, it MUST use the tie-breaker rules as specified in [\[RFC4601\]](#) to determine an ASSERT winner. PIM MT-ID is not considered in deciding a winner from Assert process.

[4.4.2. Downstream Routers](#)

A conflict is detected by a downstream router when it sees a different PIM MT-ID attribute from other routers on the LAN, regardless of whether the router has local topology selection policy or not.

- if a downstream router sees different PIM MT-ID attributes from

A new PIM Join Attribute type needs to be assigned. 3 is proposed for now.

7. Security Considerations

As a type of PIM Join Attribute, the security considerations described in [[RFC5384](#)] apply here. Specifically, malicious alteration of PIM MT-ID may cause the resiliency goals to be violated.

8. Acknowledgments

The authors would like to thank Eric Rosen, Ice Wijnands, Dino Farinacci, Colby Barth and Les Ginsberg for their input.

9. Authors' Addresses

Yiqun Cai
Cisco Systems, Inc
170 West Tasman Drive
San Jose, CA 95134

E-mail: ycai@cisco.com

Heidi Ou
Cisco Systems, Inc
170 West Tasman Drive
San Jose, CA 95134

E-mail: hou@cisco.com

10. Normative References

[RFC2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", [RFC 4601](#), August 2006.

[RFC5384] A. Boers, I. Wijnands, E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", [RFC 5384](#), November 2008

11. Informative References

[RFC4915] P. Psenak, S. Mirtorabi, A. Roy, L. Nguyen, P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", [RFC 4915](#), June 2007.

[RFC5120] T. Przygienda, N. Shen, N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), February 2008.

[ID.ietf-pim-rpf-vector] I. Wijnands, A. Boers, E. Rosen, "The RPF Vector TLV", [draft-ietf-pim-rpf-vector](#).

[ID.ietf-l3vpn-2547bis-mcast] E. Rosen, R Aggarwal, "Multicast in MPLS/BGP IP VPNs", [draft-ietf-l3vpn-2547bis-mcast](#)

