

PIM WG
Internet Draft
Intended Status: Proposed Standard
Expires: September 29, 2010

Yiqun Cai
Heidi Ou

Cisco Systems, Inc.

March 29, 2010

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

[draft-ietf-pim-mtid-04.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/1id-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 September 29, 2010.

Copyright Notice

Copyright (c) 2010 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. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4](#).e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

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	3
2	Introduction	3
3	Functional Overview	3
3.1	PIM RPF Topology	3
3.2	PIM MT-ID	4
3.3	Applicability	5
4	Protocol Specification of PIM MT-ID	5
4.1	PIM MT-ID Hello Option	5
4.2	PIM MT-ID Join Attribute	5
4.2.1	Sending PIM MT-ID Join Attribute	5
4.2.2	Receiving PIM MT-ID Join Attribute	6
4.2.3	Validating PIM MT-ID Join Attribute	6
4.2.4	Conflict Resolution	7
4.2.4.1	Conflict Resolution Rules For Upstream Routers	7
4.2.4.2	Conflict Resolution Rules For Downstream Routers ...	8
5	Packet Format	8
5.1	PIM MT-ID Hello Option	8
5.2	PIM MT-ID Join Attribute TLV Format	9
6	IANA Considerations	9
7	Security Considerations	9
8	Acknowledgments	10
9	Authors' Addresses	10
10	Normative References	10
11	Informative References	11

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](#)].

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 per PIM domain. 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. In practice, when only one unicast routing protocol (such as OSPF or IS-IS) is used, PIM MT-ID is recommended to be assigned using the same value as the IGP topology identifier. This is for the purpose of reducing management overhead and simplifying troubleshooting.
- 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 [RFC5496] [[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

The change to the PIM protocol includes two pieces, PIM MT-ID Hello Option and PIM MT-ID Join Attribute.

4.1. PIM MT-ID Hello Option

A router MUST include both PIM MT-ID and PIM Join Attribute Hello Option in its PIM Hello packets if it supports functionality described by this document.

4.2. PIM MT-ID Join Attribute

4.2.1. 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 any of the routers on the LAN does not include "PIM Join Attribute" or "PIM MT-ID" option in its Hello packets.

- a router SHOULD NOT encode PIM MT-ID for pruned sources. If encoded, the value is ignored.

4.2.2. 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.

4.2.3. 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. If there are multiple PIM MT-ID Join Attributes included, only the last one is accepted.
- 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 silently ignored. 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" or "PIM MT-ID" option in its Hello packets, the encoding may still be considered valid by an implementation.

4.2.4. Conflict Resolution

It is important to note that the definition of "PIM MT-ID conflict" varies depending on whether it is on an upstream or a downstream router.

On an upstream router, a conflict occurs when the 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. In another word, 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 does not apply the conflict resolution procedures.

On the other hand, when a downstream router sees a different PIM MT-ID attribute from other routers on the LAN it applies rules to resolve the conflicts regardless of whether the router has local topology selection policy or not.

It MUST be noted that the MT-ID value being considered for comparison does not include the four reserved bits. That is, only the lower order 12 bits are used in resolving conflicting attributes.

4.2.4.1. Conflict Resolution Rules For Upstream Routers

- 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.

In order to minimize the chances of potential transient forwarding loops, an upstream router MAY choose to ignore the incoming PIM Join/Prune packets all together if it sees a conflict in PIM MT-ID attributes. This action may also be taken by an upstream router which has locally configured topology selection policy, as an exception to the rules described above.

- 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.

4.2.4.2. Conflict Resolution Rules For Downstream Routers

- if a downstream router sees different PIM MT-ID attributes from PIM Join packets, it MUST follow the specification of [\[RFC4601\]](#) as if the attribute did not exist. For example, the router suppresses its own Join packet if a Join for the same (S,G) is seen.

The router MUST NOT use the rules specified in [\[RFC5384\]](#) to select a PIM MT-ID from Join packets sent by other downstream routers.

- if a downstream router sees its preferred upstream router loses in the ASSERT process, and the ASSERT winner uses a different PIM MT-ID, the downstream router SHOULD still choose the ASSERT winner as the RPF neighbour but it MUST NOT encode PIM MT-ID when sending Join packets to it.

5. Packet Format**5.1. PIM MT-ID Hello Option**

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      OptionType      |      OptionLength      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     OptionValue                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- OptionType: 30.
- OptionLength: 8.
- OptionValue: There is none specified at this moment.

5.2. PIM MT-ID Join Attribute TLV Format

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|F|E| Attr Type | Length           |R R R R| Value                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- F bit: 0 Non-transitive Attribute.
- E bit: As specified by [[RFC5384](#)].
- Attr Type: 3.
- Length: 2.
- R: Reserved bits, 4 in total.
- Value: PIM MT-ID, 1 to 4095. Range 2048 to 4095 are for experimental and proprietary use.

6. IANA Considerations

A new PIM Hello Option type, 30, has been assigned for PIM MT-ID Hello Option. The detail is in [[HELLO](#)].

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, Les Ginsberg, Dimitri Papadimitriou and Thomas Morin 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.

[RFC5496] I. Wijnands, A. Boers, E. Rosen, "The Reverse Path Forwarding (RPF) Vector TLV", [RFC 5496](#), March 2009.

[HELLO] IANA, "PIM-Hello Options",
<http://www.iana.org/assignments/pim-parameters/pim-parameters.xhtml#pim-parameters-1>

[ID.ietf-l3vpn-2547bis-mcast] E. Rosen, R Aggarwal, "Multicast in MPLS/BGP IP VPNs", [draft-ietf-l3vpn-2547bis-mcast-10](#), January 2010

