

PIM Working Group
Internet Draft
Intended status: Standards Track
Expires: 23 September 2026

Y. Liu
China Mobile
C. Lin
New H3C Technologies
G. Gryszata
Orange
M. Mishra
Cisco Systems, Inc.
23 March 2026

PIM Reverse Path Forwarding (RPF) Vector Conflict Resolution
draft-liu-pim-rpf-vector-conflict-resolution-02

Abstract

[RFC7891] Section 7 defines the handling principles for PIM Join attributes with same type of RFF Vector and Explicit RFP Vector, but with different contents are received.

However, it does not address scenarios where one downstream router includes a RFF Vector in its message while another does not. This leaves the handling of such conflicts unspecified.

This document supplements the existing specifications by defining the processing rules for this specific conflict scenario.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 23 September 2026.

Copyright Notice

Copyright (c) 2026 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

Liu, et al. Expires 23 September, 2026

[Page 1]

(<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction.....	2
2. Terminology.....	3
3. Problem.....	3
3.1. Scene 1.....	3
3.2. Scene 2.....	4
4. Conflicting Resolution Principle.....	5
5. Conflicting Resolution Process.....	5
6. Modifications to RFC 7891.....	6
7. IANA Considerations.....	7
8. Security Considerations.....	7
9. Normative References.....	7
10. Informative References.....	7
Authors' Addresses.....	8

1. Introduction

In the Protocol Independent Multicast - Sparse Mode (PIM-SM) specification [RFC4601], a Join message sent by a node may identify one or more multicast distribution trees that the node wishes to join. Each tree is identified by the combination of a multicast group address and a source address (where the source address may be a wildcard).

[RFC5384] describes the inclusion of PIM Join Attributes in Join messages, associating them with specific multicast trees to be joined.

[RFC5496] describes the scenario where core routers do not maintain external routes, the edge router send PIM join messages to core routers, utilizing PIM Join Attributes to carry RPF Vector TLVs to specify the path for the core routers to send the joins. Core routers select the path for the join based on the addresses carried in the RPF Vector TLV through local unicast routing, hence it is also referred to as "loose" RPF Vector.

[RFC7891] describes the use of PIM Join Attributes to carry Explicit RPF Vector TLVs to strictly specify the join path, in scenarios when network administrators do not want to rely on unicast reachability to the RPF vector address but instead want to build a strict path based on the RPF vector,

A router may receive conflicting RPF vector from different downstream routers. Section 7 of [RFC7891] describes the handling of conflicts when received join attributes contain RPF Vector TLVs or Explicit RPF Vector TLVs that are either of differing types or have mismatched contents. The provided rule might not adequately address scenarios involving multicast backup path architectures.

Moreover, there are cases where one downstream router includes a RPF Vector in its Join message while another does not. [RFC7891] does not specify how to handle such conflicts.

This document updates the existing specifications by defining the processing rules for this specific conflict scenario.

2. Terminology

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].

3. Problem

3.1. Scene 1

The Scenario PIM Join Attribute conflicts occurring when One downstream neighbor includes the Join Attributes in its Join message while Another neighbor sends Join message without the Join Attributes.

[RFC7891] presents a scenario involving PIM with RPF Vector, as depicted in Figure 1. The cost of the links between R5-R6 and R7-R8 is 100, while the cost of all other links is 10. The multicast join path is R4->R3->R6->R5->R2->R1, where the multicast join is routed to the source using the RPF Vector list. This scenario can be used to join along a multicast backup path, allowing multicast traffic forwarding through a backup path in the event of a failure in the primary path (R4->R3->R2->R1).

In this scenario, R8 is connected to Receiver2 and sends a join message to R6 without carrying the join attribute. Intermediate node R6 receives a PIM join with RPF Vector attributes from R3 and another PIM join without RPF Vector attributes from R8.

According to the attribute conflict handling principle outlined in [RFC5384], when both join messages carry RPF Vector attributes, the one with higher priority will be selected. However, in the scenario where one join message carries the RPF Vector attribute and the other does not, there is no specified handling, resulting in uncertain join outcomes.

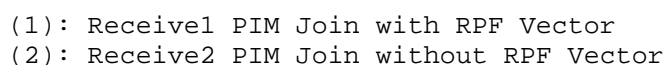


Figure 1

```
[R6] (S,G) Entry 2: Incoming interface: if2(R3-R6), Outgoing
interface: R6-R8 (Join without RPF Vector)
```

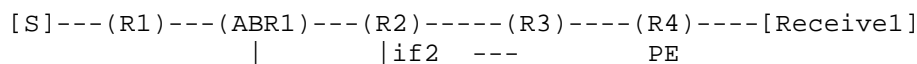
Given these two joins, a priority selection is required. However, the current specifications do not provide guidelines for handling such conflicts, leading to uncertainty in the join status.

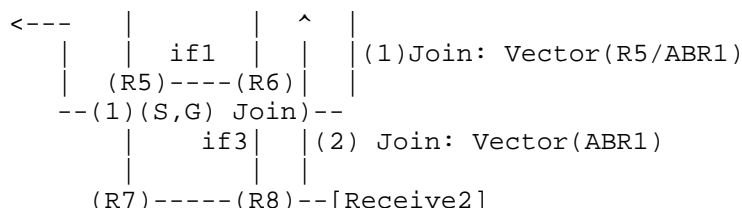
3.2. Scene 2

As shown in the Inter-area scenario in Figure 2, ABR1 advertises the route related to the source address to R4 and R8 via iBGP. Multicast FRR is enabled on R2. When Receiver1 joins, the multicast backup path is R4->R3->R2->R6->R5->ABR1->R1. In the PIM Join message sent by R3 to R6, two type-4 RPF Vectors are carried, with addresses R5 and ABR1 respectively.

When Receive2 joins, the Join message sent by R8 to R6 carries one type-4 RPF Vector with the address ABR1.

According to the attribute conflict handling principle specified in [RFC5384], when both join messages carry the RPF vector attribute, the one with higher precedence will be selected. However, this may result in the join path not being the optimal path.





(1): Receive1 PIM Join with RPF Vector(R5/ABR1)

(2): Receive2 PIM Join with RPF Vector(ABR1)

Figure 2

4. Conflicting Resolution Principle

When two join messages are received:

If one with an RPF Vector and the other without, the join message without the RPF Vector is given priority.

If one join message only contains a type 0 RPF Vector (RPF Vector) and the other contains a type 4 RPF Vector (Explicit RPF Vector), the join message with only the type 0 RPF Vector is given priority.

If one join message only contains a single type 0 RPF Vector (RPF Vector), and the other contains a stack of multiple type 0 RPF Vectors (all the same type), the join message with fewer RPF Vectors is given priority.

This procedure ensures deterministic resolution of attribute conflicts while maintaining backward compatibility with routers that do not support preference attributes.

5. Conflicting Resolution Process

For the scenario in Figure 1, R4 is connected to a local receiver, Receiver1. Enable FRR on R3, the primary path is [R4]->[R3]->[R2]->[R1], while the backup path is [R4]->[R3]->[R6]->[R5]->[R2]->[R1]. R8 is connected to another local receiver, Receiver2. The multicast join path is [R8]->[R6]->[R3]->[R2]->[R1].

To prevent conflicts, the PIM Join without the RPF Vector is prioritized. Thus, R6 retains only the entry with if1(R5-R6) as the incoming interface and R6-R8 as the outgoing interface, stopping R3 from receiving multicast traffic via the backup path.

The entry on R6 is:

[R6] (S,G) Entry 2: Incoming interface: if2(R3-R6), Outgoing interface: if3(R6-R8)

When R8's local receiver departs, R8 sends its prune, prompting R6 to revert to the backup path established by R3, allowing R3 to receive multicast traffic through both paths again.

The final entry on R6 is:

[R6] (S,G) Entry: Incoming interface: if1(R5-R6), Outgoing interface: if2(R3-R6)

For the scenario shown in Figure 2, R3 is connected to a local receiver, Receiver1. FRR is enabled on R2, with the primary path being [R4]->[R3]->[R2]->[ABR1]->[R1], and the backup path being [R4]->[R3]->[R2]->[R6]->[R5]->[ABR1]->[R1]. R8 is connected to another local receiver, Receiver2, and the multicast join path is [R8]->[R6]->[R2]->[ABR1]->[R1].

To avoid conflicts, PIM Join messages with fewer RPF vectors are prioritized. Therefore, R6 only retains the entry with ingress interface if1 (R5-R6) and egress interface R6-R8, thereby preventing R3 from receiving multicast traffic via the backup path.

At this point, the entry on R6 is:

[R6] (S,G) Entry 2: Ingress interface: if2 (R5-R6), Egress interface: if3 (R6-R8)

When the local receiver on R8 leaves, R8 sends a Prune message, prompting R6 to switch back to the backup path established by R3, allowing R3 to again receive multicast traffic via both paths.

The final entry on R6 becomes:

[R6] (S,G) Entry: Ingress interface: if1 (R5-R6), Egress interface: if2 (R2-R6)

6. Modifications to RFC 7891

Revised Section 8 of RFC7891:

Original text:

"If it is determined that there is either a conflict with RPF Vector types or the RPF Vector content, the router uses the RPF Vector stack from the PIM adjacency with the numerically smallest IP address."

Modified to:

"If it is determined that there is either a conflict with RPF Vector types or the RPF Vector content,

If one join carries an RPF Vector while the other doesn't, prefer the join without the RPF Vector;

For joins containing RPF Vectors of identical types, prefer the join with fewer RPF Vectors;

Otherwise, select the RPF Vector stack from the PIM adjacency with the numerically smallest IP address."

7. IANA Considerations

This document has no IANA actions.

8. Security Considerations

TBD

9. Normative References

- [RFC5384] Boers, A., Wijnands, I., and E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", RFC 5384, DOI 10.17487/RFC5384, November 2008, <<http://www.rfc-editor.org/info/rfc5384>>.
- [RFC5496] IJ. Wijnands, A. Boers, E. Rosen, Cisco Systems, Inc., "The Reverse Path Forwarding (RPF) Vector TLV", RFC 5496, DOI 10.17487/RFC5496, March 2009, <<http://www.rfc-editor.org/info/rfc5496>>.
- [RFC7761] Fenner, B., Handley, M., UCL, Holbrook, H., and I. Kouvelas, Arista Networks, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC7761, March 2016.
- [RFC7891] J. Asghar, IJ. Wijnands, Ed., S. Krishnaswamy, A. Karan, Cisco Systems, V. Arya, DIRECTV Inc., "Explicit Reverse Path Forwarding (RPF) Vector", RFC 7891, DOI 10.17487/RFC7891, June 2016, <<http://www.rfc-editor.org/info/rfc7891>>.

10. Informative References

TBD

Authors' Addresses

Yisong Liu
China Mobile
China
Email: liuyisong@chinamobile.com

Changwang Lin
New H3C Technologies
China
Email: linchangwang.04414@h3c.com

Guillaume Gryszata
Orange
France
Email: guillaume.gryszata@orange.com

Mankamana Mishra
Cisco Systems, Inc.
San Jose,
United States of America
Email: mankamis@cisco.com

