

Network Working Group  
Internet-Draft  
Intended status: Experimental  
Expires: 6 November 2026

A. Gopal  
S. Venaas  
Cisco Systems, Inc.  
F. Meo  
5 May 2026

PIM Flooding Mechanism and Source Discovery Enhancements  
draft-ietf-pim-pfm-forwarding-enhancements-05

Abstract

The Protocol Independent Multicast (PIM) Flooding Mechanism (PFM) provides a generic hop-by-hop message exchange framework for distributing multicast information among PIM routers. Existing PFM procedures enable efficient source discovery without reliance on Rendezvous Points, shared trees, or initial data registers.

This document specifies enhancements to PFM forwarding behavior to improve efficiency and scalability. In particular, it introduces mechanisms to reduce redundant message transmission over multiple parallel links and extends the encoding of multicast information through additional Type-Length-Value (TLV) structures and sub-TLVs to convey richer flow-related data. These enhancements optimize control-plane overhead while preserving interoperability with existing PFM procedures, enabling more efficient dissemination of multicast state in PIM networks.

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 6 November 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 (<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
1.1. Conventions Used in This Document . . . . .	3
1.2. Terminology . . . . .	3
2. PIM PFM Sub-TLV . . . . .	4
2.1. Group Source Info TLV . . . . .	4
2.2. Group Source Info TLV Hello option . . . . .	5
2.3. Considerations for using the Group Source Info TLV . . . . .	6
3. PIM PFM forwarding optimization . . . . .	6
3.1. RFC 6395 Compliance . . . . .	7
3.2. PFM optimization Hello option . . . . .	7
3.3. Sample PFM Topology . . . . .	8
3.4. PFM message handling with Relaxed-RPF enabled . . . . .	8
3.5. Maintenance of PFM_OPT_IF . . . . .	9
3.6. PFM message forwarding to sender . . . . .	10
4. Security Considerations . . . . .	10
5. IANA Considerations . . . . .	10
6. Acknowledgments . . . . .	11
7. Normative References . . . . .	11
Authors' Addresses . . . . .	11

## 1. Introduction

PIM Flooding Mechanism [RFC8364] allows a PIM router in the network to originate a PFM message to distribute announcements of active sources to its PIM neighbors [RFC7761]. All PIM neighbors then process this PFM message and flood it further on their PIM-enabled links. To prevent loops, the originator address as defined in Section 3.1 [RFC8364] is used for Reverse Path Forwarding (RPF) checking at each router. This RPF check is defined in Section 3.4.1 [RFC8364]. Periodic PFM messages are exchanged to keep the multicast information updated across the PIM domain (Section 3.4.2 [RFC8364]).

The TLV defined in [RFC8364] for source discovery conveys only source and group information. It does not provide a mechanism to include additional attributes describing a multicast flow.

In addition, PFM messages are flooded on all PIM-enabled links. When two routers maintain multiple PIM adjacencies, identical PFM messages are transmitted across each link. Receivers perform RPF checks and discard duplicates as needed. This behavior introduces unnecessary processing overhead, both periodically and upon source discovery.

This document defines two independent enhancements to PFM message exchange:

1. A new TLV that supports Sub-TLVs, enabling the inclusion of additional flow-related information. This enhancement is specified in Section 2.
2. An optimization for PFM message exchange across multiple PIM adjacencies between the same pair of routers. By leveraging PIM Router-IDs [RFC6395], routers can identify such adjacencies and limit message transmission to a subset of links, reducing redundant processing. This optimization applies to point-to-point links and does not alter behavior on shared LANs. This enhancement is specified in Section 3.

Implementations MAY support these enhancements independently; however, support for both is RECOMMENDED.

### 1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

### 1.2. Terminology

RPF: Reverse Path Forwarding

FHR: First-Hop Router


PFM-SD: PIM Flooding Mechanism and Source-Discovery

## 2. PIM PFM Sub-TLV

PFM-SD [RFC8364] defines a Group Source Holdtime (GSH) TLV for announcing active sources. The GSH TLV conveys only source and group information. This document defines an extension that allows PIM routers to exchange additional information associated with multicast sources.

### 2.1. Group Source Info TLV

This document defines a new Group Source Info (GSI) TLV (Type TBD1). The GSI TLV is functionally similar to the GSH TLV but applies to a single (S,G) entry and supports the inclusion of Sub-TLVs to convey additional flow-specific information. Support for the GSI TLV is advertised using a PIM Hello option (TBD2), as described in Section 2.2



The format of the GSI TLV is as follows:

T-bit (1 bit): Indicates transitivity. If set to 0, a router that does not support the TLV or any contained Sub-TLV MUST NOT forward the message. If set to 1, the message MAY be forwarded even if unsupported TLVs or Sub-TLVs are present.

Type (15 bits): Set to TBD1.

Length (16 bits): The length, in octets, of the Value field.

Group Address: The multicast group address encoded as specified in Section 4.9.1 of [RFC7761]. The length of this field depends on the address family: 64 bits for IPv4 native encoding, 160 bits for IPv6.

Source Address: The unicast source address encoded as specified in Section 4.9.1 of [RFC7761]. The length of this field depends on the address family: 48 bits for IPv4 native encoding, 160 bits for IPv6.

Holdtime (16 bits): The lifetime, in seconds, for the advertised (S,G) information.

Sub-TLVs: Zero or more Sub-TLVs MAY be included. Each Sub-TLV consists of:

Type (16 bits): The Sub-TLV Type. Values for this field are assigned by IANA.

Length (16 bits): The length, in octets, of the Value field. The length may be 0 if no value is present.

Value: The content of the Sub-TLV, whose format is determined by the Sub-TLV Type.

## 2.2. Group Source Info TLV Hello option

A PIM router indicates support for the GSI TLV defined in this document by including the Group Source Info TLV Hello option in PIM Hello messages. The format of the Hello option is as follows:

```

      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 = TBD2           |           Length = 0           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

OptionType (16 bits): TBD2

OptionLength (16 bits): 0

The presence of this option signifies that the router supports the GSI TLV.

### 2.3. Considerations for using the Group Source Info TLV

All PIM routers MUST track which neighbors advertise support for the GSI TLV via the Hello option Section 2.2. This tracking is beneficial in heterogeneous networks where only certain routers support the new TLV Type TBD1. If GSI TLV is supported, use of the GSI TLV (Type TBD1) is RECOMMENDED.

A router that supports the GSI TLV MUST:

- \* Advertise its capability by including the Hello option (OptionType TBD2) in PIM Hello messages.
- \* Track, per PIM interface, whether all neighbors support the GSI TLV. The scope and persistence of this state are implementation-specific. An implementation MAY retain this state even if local capability is disabled.
- \* If acting as a First Hop Router (FHR), originate a Type TBD1 TLV when all neighbors on the outgoing interface support Type TBD1.
- \* If acting as an FHR, originate a Type 1 TLV [RFC8364] when any neighbor on the outgoing interface does not support Type TBD1.
- \* Upon receipt of a Type TBD1 TLV, MUST forward the PFM message unchanged on interfaces where all neighbors support Type TBD1.
- \* For interfaces with at least one neighbor that does not support Type TBD1, convert each Type TBD1 TLV to a Type 1 TLV [RFC8364] and forward only on those interfaces. The conversion MUST preserve the group, source, and holdtime fields, and MUST ignore Sub-TLVs. Multiple (S,G) entries for the same group SHOULD be aggregated into a single Type 1 TLV. However, it MUST still send Type TBD1 TLV on all interfaces where the neighbors do support it.
- \* A PFM message MAY contain both Type 1 and Type TBD1 TLVs. When forwarding to neighbors that do not support Type TBD1, all Type TBD1 TLVs MUST be converted to Type 1 TLVs.

### 3. PIM PFM forwarding optimization

### 3.1. RFC 6395 Compliance

To apply the forwarding optimization defined in this document, PIM routers MUST advertise a Router-ID as specified in [RFC6395]. Within a PIM VRF, a router MUST use the same 4-octet Router-ID in PIM Hello messages on all interfaces and MUST cache Router-IDs learned from neighbors. Within a PIM VRF, a router MUST identify interfaces with a single neighbor sharing the same Router-ID, indicating multiple adjacencies to that neighbor. This identification is necessary for applying the forwarding optimization defined in this document. Router-IDs are assumed to be unique within the PIM domain. If this assumption is violated, the optimization defined in this document MUST NOT be applied.

### 3.2. PFM optimization Hello option

A PIM router indicates support for the forwarding optimization by including the PFM Optimization Hello option (OptionType TBD3) in PIM Hello messages.

```

      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 = TBD3           |           Length = 0           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

OptionType (16 bits): TBD3

OptionLength (16 bits): 0

A router that supports this optimization MUST track, per interface, whether all neighbors support the option. This tracking is beneficial in heterogeneous networks where only certain routers support the optimization.

For each learned Router-ID, the router MUST maintain a set of interfaces, denoted as PFM\_OPT\_IF, that satisfy both of the following conditions:

- \* The neighbor with this Router-ID is the sole PIM neighbor on this interface.
- \* The neighbor is advertising the PFM optimization option TBD3 on this interface.

PFM message exchange MAY be optimized on interfaces in the PFM\_OPT\_IF set. This is discussed in Section 3.4.

### 3.3. Sample PFM Topology




Figure 1: Four Router Network Topology Example

### 3.4. PFM message handling with Relaxed-RPF enabled

When two routers maintain multiple adjacencies and are the only neighbors on those links, PFM messages are typically transmitted on all links and filtered by RPF checks at the receiver. This results in redundant processing. If both routers advertise Router-IDs and support the optimization, each router forms a PFM\_OPT\_IF set containing eligible interfaces.

When Relaxed-RPF is enabled:

- \* A sender MUST select a single interface from its PFM\_OPT\_IF set for PFM transmission to that neighbor. The selection method is implementation-specific.
- \* On shared LANs, the sender MUST send PFM messages as normal since optimization cannot be applied when there are more than two routers on the network segment.
- \* A receiver that supports Relaxed-RPF MUST:
  - Determine the expected RPF interface using standard procedures.
  - Accept a PFM message received on any interface in the PFM\_OPT\_IF set if both sender and receiver support the optimization.



- Otherwise, perform standard RPF validation.

Referring to Figure 1, when Router A originates or forwards a PFM message, it MUST transmit the message on exactly one of links L1, L2, or L3. This behavior reduces processing overhead on point-to-point links. The selection of the interface from the PFM\_OPT\_IF set is implementation-specific. Router A also MUST send the message on both LAN 1 and LAN 2 to ensure Routers C and D receive the message.

### 3.5. Maintenance of PFM\_OPT\_IF

Routers MUST update the PFM\_OPT\_IF set upon neighbor or capability changes:

- \* Neighbor Addition: If the new neighbor is the sole neighbor on the interface and advertises both a Router-ID and the optimization option, the interface MUST be added to the corresponding PFM\_OPT\_IF set. If no set exists, it MUST be created. If a second neighbor appears on the interface, the interface MUST be removed from the PFM\_OPT\_IF set.
- \* Neighbor Removal: If exactly one neighbor remains and it advertises both a Router-ID and the optimization option, the interface MUST be added to the PFM\_OPT\_IF set for that Router-ID. If no set exists, it MUST be created.
- \* Router-ID Changes: If a neighbor starts advertising a Router-ID and satisfies all conditions, the interface MUST be added to the PFM\_OPT\_IF set. If a neighbor stops advertising a Router-ID, the interface MUST be removed from the PFM\_OPT\_IF set for that Router-ID. If the set becomes empty, it MUST be deleted.
- \* Optimization Capability Changes: If a neighbor starts advertising the optimization option and satisfies all conditions, the interface MUST be added to the PFM\_OPT\_IF set. If a neighbor stops advertising the optimization option, the interface MUST be removed from the PFM\_OPT\_IF set for that Router-ID. If the set becomes empty, it MUST be deleted.

These procedures apply during topology changes, configuration updates, and software upgrades or downgrades. Routers MUST maintain accurate PFM\_OPT\_IF state for each Router-ID.

### 3.6. PFM message forwarding to sender

When the TBD3 optimization is enabled on a PIM router, the router MUST NOT forward a PFM message on a link if both of the following conditions are true: (1) the link has only one neighbor, and (2) that neighbor's Router-ID matches the Router-ID of the router that originated the PFM message. It is sufficient for the neighbor to advertise only the Router-ID, without any additional optimization options, since this information alone ensures the message is not sent back to its original sender, thereby reducing unnecessary PFM message forwarding.

## 4. Security Considerations

When it comes to general PIM message security, see [RFC7761]. For PFM security see [RFC8364]. This optimization relies on correct Router-ID and capability advertisement in PIM Hellos, as well as general PIM hello integrity. For the new PFM TLV, the security considerations are the same as for the existing PFM TLV defined in [RFC8364].

## 5. IANA Considerations

This document requires the assignment of a new PFM TLV Type TBD1 in the "PIM Flooding Mechanism Message Types" registry.

### PIM Flooding Mechanism Message Types

Type	Name	Reference
-----		
TBD1	Group Source Info	[This document]

Also, a new registry "PIM Flooding Mechanism Group Source Info Message Types" registry needs to be created. Assignments for the new registry are to be made according to the policy "IETF Review" as defined in [RFC8126]. The initial content of the registry should be:

### PIM Flooding Mechanism Group Source Info Sub-TLV Types

Type	Name	Reference
-----		
0-32767	Unassigned	

This document requires the assignment of two new PIM Hello Options:

## PIM Hello Options

Value	Length	Name	Reference
TBD2	0	GSI TLV support	[This document]
TBD3	0	PFM Optimization	[This document]

## 6. Acknowledgments

## 7. Normative References

- [RFC6395] Gulrajani, S. and S. Venaas, "An Interface Identifier (ID) Hello Option for PIM", RFC 6395, DOI 10.17487/RFC6395, October 2011, <<https://www.rfc-editor.org/info/rfc6395>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8364] Wijnands, IJ., Venaas, S., Brig, M., and A. Jonasson, "PIM Flooding Mechanism (PFM) and Source Discovery (SD)", RFC 8364, DOI 10.17487/RFC8364, March 2018, <<https://www.rfc-editor.org/info/rfc8364>>.

## Authors' Addresses

Ananya Gopal  
Cisco Systems, Inc.  
Tasman Drive  
San Jose, CA 95134  
United States of America

Email: [ananygop@cisco.com](mailto:ananygop@cisco.com)

Stig Venaas  
Cisco Systems, Inc.  
Tasman Drive  
San Jose, CA 95134  
United States of America  
Email: [svenaas@cisco.com](mailto:svenaas@cisco.com)

Francesco Meo  
Email: [fran.meo@gmail.com](mailto:fran.meo@gmail.com)