

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 8 January 2026

A. Gopal
S. Venaas
Cisco Systems, Inc.
F. Meo
7 July 2025

PIM Flooding Mechanism and Source Discovery Enhancements
draft-ietf-pim-pfm-forwarding-enhancements-02

Abstract

PIM Flooding Mechanism is a generic PIM message exchange mechanism that allows multicast information to be exchanged between PIM routers hop-by-hop. One example is PIM Flooding Mechanism and Source Discovery which allows last hop routers to learn about new sources using PFM messages, without the need for initial data registers, Rendezvous Points or shared trees.

This document defines a new TLV for announcing sources that allows for Sub-TLVs that can be used to provide various types of information. This document also defines methodologies that enhance forwarding efficiency in PFM deployments.

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 8 January 2026.

Copyright Notice

Copyright (c) 2025 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. PFM message handling with Optimizations in effect	7
3.4. PFM message handling with Relaxed-RPF enabled	8
4. Security Considerations	9
5. IANA Considerations	9
6. Acknowledgments	10
7. Normative References	10
Authors' Addresses	10

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 RPF checking at each router. This RPF check is defined in Section 3.4.1 [RFC8364]. Periodic PFM messages are triggered, see Section 3.4.2 [RFC8364] and exchanged to keep the multicast information updated across the PIM domain.

Firstly, the TLV used by PFM [RFC8364] for source discovery only specifies source and group information to announce an active source. There is no convenient way to provide additional information about a flow.

Secondly, a PIM router will flood a PFM message on all its PIM enabled links. It is the recipient's responsibility to perform RPF checks on all received PFM messages and then decide whether to accept or drop a particular message. This means that if two routers have PIM neighborships over more than one link, the same PFM messages are exchanged or dropped over more than one link between the same two routers. This leads to extra processing at each PIM router, periodically, or every time a new source is discovered (in case of a PFM-SD implementation). We can reduce the processing overhead for the router-pair having PIM neighborships over multiple links.

This document discusses two new improvements in PFM message exchanges between PIM routers.

1. This document defines a new TLV for announcing sources that allows for Sub-TLVs that can be used providing various types of information. This enhancement is discussed in detail in Section 2.
2. Utilizing the PIM Router-IDs [RFC6395], PIM routers can limit PFM message exchanges to only on ONE link per router-pair, even though these two PIM routers may maintain PIM neighborships over multiple links. Note that this is applicable in cases where there are only two routers on each of the links between them - either a Point-Point link, or exactly 2 PIM neighbors on a LAN. In such cases, PFM can improve in performance by first identifying the PIM routers in the network using Router Identifiers [RFC6395] (Router-IDs) that are announced via PIM hellos. This enhancement allows PFM to limit message exchanges to only those that are necessary and is discussed in detail in Section 3.

These are independent enhancements and an implementation could support one but not the other, however it is RECOMMENDED to implement both.

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

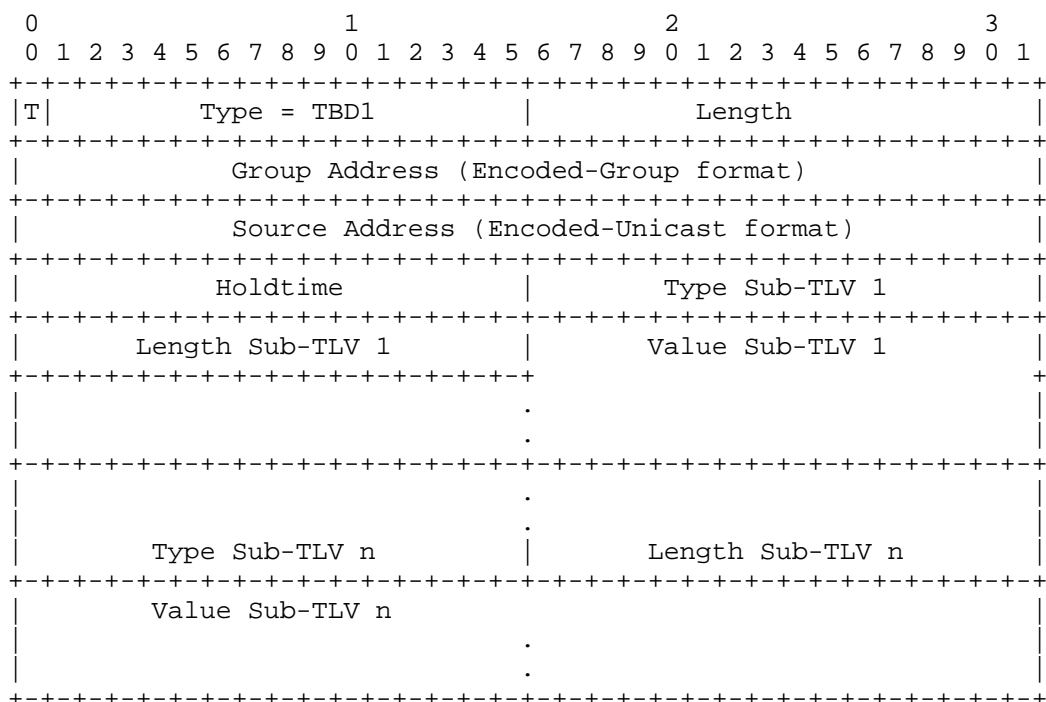
PFM-SD: PIM Flooding Mechanism and Source-Discovery

2. PIM PFM Sub-TLV

PFM-SD [RFC8364] defines a Group Source Holdtime (GSH) TLV for announcing active sources. However, it could be beneficial for PIM routers to exchange additional data about these sources.

2.1. Group Source Info TLV

This document defines a new Group Source Info (GSI) TLV that is used similarly to the GSH TLV except that it only provides info for a single source, and includes additional information about the flow in Sub-TLVs. Note that the support for this TLV Type TBD1 is advertised by PIM routers using the PIM Hello Option TBD2 and is discussed in detail in Section 2.2



T: If the Transitive bit is set to 0, a router MUST NOT forward the

message unless it supports this TLV and all the Sub-TLVs that are present in the TLV in this message. If the transitive bit is set to 1, it is forwarded even if the router does not support the TLV or all the Sub-TLVs present.

Type: This TLV has type TBD1.

Length: The length of the value in octets.

Group Address: The group that sources are to be announced for. The format for this address is given in the Encoded-Group format in [RFC7761].

Source Address: The source address for the corresponding group. The format for this address is given in the Encoded-Unicast address in [RFC7761].

Holdtime: The Holdtime (in seconds).

Type Sub-TLV 1..n: The TLV contains n Sub-TLVs, n MAY be 0. The total length of the TLV (the Length field) is used to derive how many octets are used for Sub-TLVs. It will be at least 4 * n octets if n Sub-TLVs are present. Type Sub-TLV indicates the type of the Sub-TLV. The allowed types are Sub-TLV types that are specifically defined for use in the Group Source Info TLV.

Length Sub-TLV 1..n: The length of the Sub-TLV Value field in octets.

Value Sub-TLV 1..n: The value of the Sub-TLV associated with the type and of the specified length.

2.2. Group Source Info TLV Hello option

A PIM router indicates that it supports the Group Source Info TLV specified in this document by including the new Group Source Info TLV Hello option in PIM hellos.

```

      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 = TBD2

OptionLength = 0

2.3. Considerations for using the Group Source Info TLV

All PIM routers MUST track which neighbors announce this option. This tracking is beneficial in heterogeneous networks where only certain routers support the new TLV Type TBD. Additionally, it is RECOMMENDED that only Type TBD1 be used if support is available.

Consider a router capable of exchanging PFM Type TBD1 TLVs. It MUST do the following:

- * It MUST advertise its capability by sending PIM Hello with OptionType TBD2.
- * It MUST track whether all neighbors on each of its PIM interfaces support this new TLV. Scope of this tracking is left to the implementation. It MAY track this information even if the capability on itself is removed.
- * If this router is a First Hop Router (FHR), while originating a PFM message, it MUST originate a Type TBD1 TLV if all neighbors on the PIM interface support Type TBD1.
- * If this router is an FHR, while originating a PFM message, it MUST originate a Type 1 TLV [RFC8364] if at least 1 neighbor on the PIM interface does not support Type TBD1.
- * On the receipt of a Type TBD1 TLV on a Type TBD1-capable intermediate router, this router MUST forward the PFM message as is on the PIM interfaces where all neighbors support this new type.
- * If there are PIM interfaces where at least one router does not support the new TLV, an intermediate router that supports Type TBD1 MUST convert the Type TBD1 TLV to Type 1 TLV [RFC8364] and forward it on those interfaces. The conversion mechanism is largely left to the implementation, however in a nutshell, the router MUST create and send TLV Type 1 with the source group and holdtime from the Type TBD1 and ignore the sub-tlvs. Also, if there are multiple sources for the same group, then they SHOULD be put together in one TLV, and sent as Type 1.

3. PIM PFM forwarding optimization

3.1. RFC 6395 Compliance

For the forwarding optimization in this document to be used, all PIM routers MUST announce a Router-ID as specified in [RFC6395]. A PIM router announces the same 4-byte Router-ID in PIM hellos that it sends to all neighbors on all links. It also caches the Router-IDs of its neighbors, when it receives Hellos from [RFC6395] Compliant PIM neighbors. This can be used to determine that different PIM neighbors are really the same router. In a VRF context, if the router has multiple interfaces with only one neighbor per interface, the router SHOULD check if those neighbors announce an RFC 6395 Router-ID. If the router can see the same Router-ID for multiple neighbors, PFM message exchange is optimized.

3.2. PFM optimization Hello option

A PIM router indicates that it supports enhancement mechanisms specified in this document by including the new PFM optimization Hello option. When this optimization is included in the PIM hello, the router MUST also include the Router-ID Hello Option defined in [RFC6395] with a non-zero Router-ID.

```

      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 = TBD3

OptionLength = 0

All PIM routers supporting forwarding optimization MUST track whether it is supported by all PIM neighbors on each PIM interface. This tracking is beneficial in heterogeneous networks where only certain routers support the optimization.

3.3. PFM message handling with Optimizations in effect

Consider a router that is advertising its forwarding optimization capability in the network with Hello Option TBD3 [Section 3.2] to all its PIM neighbors. It MUST simultaneously track whether the optimization TBD3 is advertised by PIM neighbors on each PIM interface. This information, along with each neighbors' Router-IDs will enable this router to make improved forwarding decisions for PFM messages.

3.4. PFM message handling with Relaxed-RPF enabled

When a router sends a PIM Hello with OptionType TBD2 and Flag Bit TBD5 set, it signals to its PIM neighbor that it can optimize forwarding between them if they are the only two neighbors on all connecting links between them.

Consider a topology where two PIM routers maintain multiple PIM neighborships over several links within the same PIM domain — and are the only two routers on these links - either a Point-to-Point link, or 2 PIM neighbors on a LAN. From each router's point of view, there is a single neighbor on each link. Traditionally, each of the routers will send out PFM messages out over all the links to its neighbor. RPF checks are one of the commonly used ways to prevent loops, hence the recipient router performs an RPF check and accepts only on one link, thereby dropping packets from all the others. Since the sender does not know which link will be chosen as the RPF-source on the neighbor, it cannot choose one of the links, without knowing its neighbor's decision.

If the Relaxed-RPF optimization is advertised by both routers, the sender MUST choose one of the links and send and forward PFM messages to its neighbor using only that link. The sender MUST do this only when the receiver is capable of the Relaxed-RPF optimization. Otherwise, the messages may be dropped because of RPF failures. The mechanism to choose a link is left to the implementation.

When a router that supports the Relaxed-RPF optimization receives a PFM message, it MUST first verify if the sender supports Relaxed-RPF optimization. If true, the receiver MUST relax its RPF check and accept the message. Additionally, the receiver MUST record the sender's router ID to prevent forwarding the message back to the sender on any other link. However, if the sender does not advertise the forwarding optimization specified in this document and the receiver supports the optimization, the receiver MUST NOT relax its RPF check, as the sender will still transmit messages across all connecting links.

The optimization mechanism relies heavily on a router's insight into whether all neighbors on each PIM interface support the TLV Type TBD3 and/or Relaxed-RPF optimization. All checks can be done at the time when a PFM message is forwarded, but it is possible to perform most checks when there are neighbor changes, so that the processing at forwarding time can be minimized. The following scenarios MUST be handled:

Adding a new neighbor on any link: When a new neighbor is added, and

it is the only neighbor, one needs to check whether it supports the optimization and announces a router-ID. In that case it may affect optimization if this same router (same router-ID) is also the only neighbor on other links. If this is the second neighbor added, and forwarding optimization was done for the first neighbor, then it now has to be disabled for that first Router-ID. Also, if the new neighbor has a router-ID, optimizations cannot be done for that router-ID either.

Existing neighbor with Router-ID change: If a neighbor changes its router-ID, one will have to check whether the old router-ID, if any, is now the only neighbor on all links it is present and optimizations should be done. The same must be checked for the new router-ID.

Neighbor removal on a link: When a router is removed, and there is exactly one remaining neighbor, it must be checked whether to do forwarding optimization. This means checking whether the remaining neighbor supports the optimization, has a router-ID, and whether that router-ID is now the only neighbor on multiple links. Also, if the neighbor being removed had a router-ID, it should be checked whether the router with that router-ID, is now the only neighbor on all links it is present and optimizations should be done.

4. Security Considerations

When it comes to general PIM message security, see [RFC7761]. For PFM security see [RFC8364].

This document defines a new format allowing for additional flow information. One concern is what happens if wrong information is provided by accident, or intentionally in a spoofed message by an attacker. The impact depends on what information is provided.

TBD any security considerations for forwarding optimizations.

5. IANA Considerations

This document requires the assignment of a new PFM TLV Type TBD1 in the "PIM Flooding Mechanism Message Types" registry. Also, a new registry "PFM Group Source Info Sub-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:

Sub-Type	Name	Reference
0	Reserved	[this document]
1-65535	Unassigned	

This document requires the assignment of two new PIM Hello Options TBD2 and TBD3, both with OptionLength 0.

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

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

Francesco Meo
Email: fran.meo@gmail.com