Network Working Group A. Gopal
Internet-Draft S. Venaas
Intended status: Experimental Cisco Systems, Inc.
Expires: 18 April 2026 F. Meo
15 October 2025

# PIM Flooding Mechanism and Source Discovery Enhancements
draft-ietf-pim-pfm-forwarding-enhancements-03

## 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



## Copyright Notice

Table of Contents



# 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 for providing various types of information. This enhancement is discussed in detail in Section 2.

2. By leveraging PIM Router-IDs [RFC6395], PIM routers can optimize PFM message exchanges when they maintain multiple neighborships with the same peer router. This optimization is particularly beneficial for router pairs connected via several links. When two routers are the sole neighbors on multiple Point-to-Point links, they need not exchange identical PFM messages across all these links. Instead, PFM can achieve performance improvements by utilizing Router Identifiers [RFC6395] (Router-IDs) announced in PIM Hello messages to identify such scenarios and restrict message exchanges to a subset of available links. This enhancement is detailed in Section 3. Note that PFM message behavior on shared LANs, where there are more than one neighbor on the same link, remains unchanged.

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

P2P: Point-to-Point

## 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

```
       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
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |T|         Type = TBD1         |           Length              |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |             Group Address (Encoded-Group format)              |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |             Source Address (Encoded-Unicast format)           |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |            Holdtime           |        Type Sub-TLV 1         |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |       Length Sub-TLV 1        |       Value Sub-TLV 1         |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               +
      |                               .                               |
      |                               .                               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |                               .                               |
      |                               .                               |
      |        Type Sub-TLV n         |       Length Sub-TLV n        |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      |        Value Sub-TLV n                                        |
      |                               .                               |
      |                               .                               |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

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 multicast group for which the source is being announced.  This address uses the Encoded-Group format specified 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 TBD1. 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 one 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 only on those unsupported interfaces. The conversion mechanism is largely left to the implementation, however, in a nutshell router MUST create and send TLV Type 1 with the source group and holdtime from the Type TBD1 and ignore the sub-TLV. Also, if there are multiple sources for the same group, then they SHOULD be put together in one TLV, and sent as Type 1. However, it MUST still send Type TBD1 TLV on all interfaces where the neighbors do support it.

* A single PFM message MAY contain both Type 1 and Type TBD1 TLVs. If so, when forwarding to neighbors that do not support Type TBD1, the intermediate router MUST convert the PFM message to Type 1 TLV if it has at least one TBD1 TLV, and all instances of TBD1 TLVs MUST be converted to Type 1 TLVs.

# 3. PIM PFM forwarding optimization

## 3.1. RFC 6395 Compliance

For the forwarding optimization in this document to be used, 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 PIM VRF context, if the router has multiple interfaces with only one neighbor per interface, the router SHOULD check if those neighbors announce an [RFC6395] Router-ID. Note that the assumption is that Router-IDs are unique per router in a PIM domain, and each device is advertising its own unique Router-ID in PIM hellos on each of its interfaces, otherwise applying this optimization can cause PFM to break.

### 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 (Option TBD3).

```
 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.

Additionally, for each unique Router-ID received by a PIM router in a PIM domain, the router MUST maintain a set of interfaces where the following two conditions are met: 1. The neighbor with this Router-ID is the only PIM neighbor on this interface and, 2. the neighbor is advertising the PFM optimization option TBD3 on this interface. This set is referred to as the PFM_OPT_IF set for each Router-ID. PFM message exchange is optimized on the interfaces belonging to PFM_OPT_IF for each Router-ID and is discussed in Section 3.4.

### 3.3. Sample PFM Topology

```
                          Router C
                              |
                              |  LAN 1
                              |
        Router A --------------------------------- Router B
           |                                          |
           |----------------- Link L1 ----------------|
           |                                          |
           |----------------- Link L2 ----------------|
           |                                          |
           |----------------- Link L3 ----------------|
           |__________________________________________|
                              |
                              |   LAN 2
                              |
                          Router D
```

Figure 1: Four Router Network Topology Example

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

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 P2P link, or 2 PIM neighbors on a LAN. On P2P links, each router sees only one neighbor, but on shared LANs, routers may see multiple neighbors. An example of such a topology is illustrated in Figure 1. Between Router A and Router B, there are multiple links - 3 P2P links and 2 shared LANs. Traditionally, each of the routers in Figure 1 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. But this can be optimized as follows.

Assume both routers A and B are advertising their respective Router-IDs on all links. When the optimizations specified in Section Section 3.2 are in effect, On both routers A and B, PFM_OPT_IF = {L1, L2, L3}.

If the Relaxed-RPF optimization is advertised by both routers, the sender MUST choose one link from their PFM_OPT_IF set and send and forward PFM messages to its neighbor using only that link. 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. In other words, the scope of optimization is limited to links present in the PFM_OPT_IF set for each Router-ID.

For example, referring to Figure 1, if Router A is forwarding or originating a PFM message, it MUST send the message on one link out of Links L1, L2, or L3. Router A also MUST send the message on both LAN 1 and LAN 2 to ensure Routers C and D receive the message. This selective behavior reduces PFM message processing overhead on the Point-to-Point links. The mechanism to choose a link from the PFM_OPT_IF set is left to the implementation.

When a router that supports the Relaxed-RPF optimization receives a PFM message, it MUST first determine the expected RPF interface for the message using standard RPF calculations. If the message was received on a link belonging to the PFM_OPT_IF set AND both the sender and receiver support Relaxed-RPF optimization, the receiver MUST accept the message regardless of the RPF check result. In all other cases, the receiver MUST perform normal RPF validation and only accept the message if it arrives on the correct RPF interface.

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: If the neighbor being added is the first neighbor on this link, the router MUST check whether this neighbor supports the optimization and announces a Router-ID. If both conditions hold TRUE, this router MUST check whether PFM_OPT_IF exists for this Router-ID. This means that the newly added neighbor is also the sole neighbor on at least one other link. Therefore, forwarding optimization MUST be enabled on this link by adding it to the existing PFM_OPT_IF set for that Router-ID. If PFM_OPT_IF does not exist for this Router-ID, it MUST be created, and this link MUST be added to the set. If the neighbor being added is the second neighbor on this link, and forwarding optimization was previously enabled for the first neighbor, it MUST now be disabled for that Router-ID on this link. Hence this link MUST be removed from the PFM_OPT_IF set for the first neighbor’s Router-ID.

Neighbor removal on a link: When a PIM neighbor is removed on a link, and there is exactly one remaining neighbor, it MUST be checked whether the remaining neighbor supports the forwarding optimization and is advertising a Router-ID. If all three conditions hold TRUE ((i) sole remaining neighbor that (ii) supports forwarding optimization, and (iii) is advertising a Router-ID), this router must check whether PFM_OPT_IF exists for this Router-ID. If the PFM_OPT_IF set for this Router-ID does not exist, it MUST be created; otherwise, the link MUST be added to the existing set.

Neighbor starts/stops advertising Router-ID: When a PIM neighbor starts advertising a Router-ID on this link, it MUST be checked whether this neighbor also supports the forwarding optimization (TBD3) on this link and whether it is the sole neighbor on this link. If both conditions hold TRUE, this router MUST check whether PFM_OPT_IF exists for this Router-ID. If it does not exist, create PFM_OPT_IF for this Router-ID and this link MUST be added to the set. If PFM_OPT_IF already exists, add this link to the existing set. When a PIM neighbor stops advertising a Router-ID on this link and is still forwarding optimization capable while being the sole neighbor on this link, this link MUST be removed from the PFM_OPT_IF set for this Router-ID. If the PFM_OPT_IF set for this Router-ID becomes empty, it MUST be deleted.

Neighbor starts/stops advertising forwarding optimization: When a PIM neighbor starts advertising the forwarding optimization (TBD3) on this link, it MUST be checked whether this neighbor is the sole neighbor on this link and whether it is advertising its Router-ID on this link. If both conditions hold TRUE, this router MUST check whether PFM_OPT_IF exists for this Router-ID. If it does not exist, create PFM_OPT_IF for this Router-ID and this link MUST be added to the set. If PFM_OPT_IF already exists, add this link to the existing set. When a PIM neighbor stops advertising the forwarding optimization (TBD3) on this link, while it is still advertising a non-zero Router-ID and is the sole neighbor on this link, this link MUST be removed from the PFM_OPT_IF set for this Router-ID. If the PFM_OPT_IF set for this Router-ID becomes empty, it MUST be deleted.

The scenarios described above apply during network and configurations changes as well as software upgrades or downgrades, that could lead to changes in neighbor capabilities. These changes will be reflected in Hello messages with the relevant options. It is essential to consistently maintain the PFM_OPT_IF set for each non-zero Router-ID with any such changes.

## 4. 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.

## 5. 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].

## 6. 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 Message Types

| Type | Name | Reference |
|---|---|---|
| 0 | Reserved | [This document] |
| 1-32767 | Unassigned | |

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

PIM-Hello Options

| Value | Length | Name | Reference |
|---|---|---|---|
| TBD2 | 0 | PFM Group Source Info support | [This document] |
| TBD3 | 0 | PFM optimization support | [This document] |

## 7. Acknowledgments

## 8. 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