

Network Working Group  
Internet-Draft  
Intended status: Experimental  
Expires: September 12, 2019

S. Venaas  
B. Ganesh  
K. Thiruvengatasamy  
R. Chokkanathapuram  
Cisco Systems, Inc.  
March 11, 2019

PIM Flooding Protocol over Reliable Transport  
draft-venaas-pim-port-pfm-00

## Abstract

The PIM Flooding Protocol (PFM) defined in [RFC8364](#) relies on sending periodic updates as it does not provide for any reliability. If a message is lost, the information will be provided in the next periodic update.

This document extends the Reliable Transport Mechanism for PIM in [RFC6559](#) to allow for sending PFM messages. This significantly reduces the PFM signaling by not requiring frequent periodic updates, and it provides for retransmission, allowing for quick recovery when an IP packet is dropped.

## 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 September 12, 2019.

## Copyright Notice

Copyright (c) 2019 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 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.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Conventions used in this document . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Protocol specification . . . . .	<a href="#">3</a>
<a href="#">4.</a>	PFM over PORT message definitions . . . . .	<a href="#">5</a>
<a href="#">4.1.</a>	PORT PFM Update . . . . .	<a href="#">5</a>
<a href="#">4.2.</a>	PORT PFM Request . . . . .	<a href="#">7</a>
<a href="#">4.3.</a>	PORT PFM Update Option . . . . .	<a href="#">8</a>
<a href="#">4.4.</a>	PORT PFM Request Option . . . . .	<a href="#">8</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">9</a>
<a href="#">6.</a>	IANA considerations . . . . .	<a href="#">9</a>
<a href="#">7.</a>	Normative References . . . . .	<a href="#">9</a>
	Authors' Addresses . . . . .	<a href="#">10</a>

## [1.](#) Introduction

The PIM Flooding Protocol (PFM) defined in [[RFC8364](#)] relies on sending periodic updates as it does not provide for any reliability. If a message is lost, the information will be provided in the next periodic update. With PFM, a router will typically originate a full update every 60 seconds. This ensures that in case of packet drops, one usually will recover in 60 seconds. There is a trade-off between the number of updates and the recovery time.

This document extends the Reliable Transport Mechanism for PIM in [[RFC6559](#)] to allow for sending PFM messages. We will refer to it as PORT (PIM Over Reliable Transport). The use of PORT significantly reduces the PFM signaling by not requiring frequent periodic updates, and it provides for retransmission, allowing for quick recovery when an IP packet is dropped. There will still be some full updates, but they can be sent much more rarely. If there is a packet drop, the reliable transport (TCP/SCTP) will ensure retransmission.

The PORT sessions are established as specified in [[RFC6559](#)] between

PIM neighbors. The sessions may be used to send other PORT messages, or they can be used only for PFM. Unless all the neighbors support PFM over PORT, regular PFM is used. How to signal support and how a router relays a PFM over PORT message as regular PFM and vice versa will be discussed in a later revision.

## [2.](#) 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.

## [3.](#) Protocol specification

PFM messages are sent over PORT by sending PORT PFM Update messages. They contain a PFM message as defined in [[RFC8364](#)]. They also contain a Full ID and a Delta ID that together specifies an ID for the update. Some updates are full updates, they contain all the information an originator is announcing. This would be similar to the periodic updates in regular PFM. Full updates over PORT are sent after some a configurable number of deltas have been sent, or whenever information needs to be withdrawn. Delta updates are used for triggered updates, similar to triggered updates in regular PFM. Each time there is some change a delta update can be triggered.

The Full ID is an unsigned 48 bit value and it is assumed that it is always increasing. That is, any Full Update MUST always have a Full ID larger than any previous updates ever sent using the same Originator address. This MUST also be preserved if the router is reloaded. For the protocol to work, it may also be necessary to ensure this if an address used as an originator address is moved to a different router. It is RECOMMENDED that implementations use the number of seconds since 0h UTC on 1 January 1900 as the ID value. This allows for this protocol to be used for about four million years from the time of publication of this document. If for any reason the clock on a router is adjusted to a value back in time, an implementation would have to ensure that values are still increasing. Since Full Updates do not need to be sent every second, one should in this case be able to catch up.

The first time a router originates a PFM message, it sends a Full update, even though it likely is triggered by some event. Full updates always have the Delta ID set to zero. After that it may send several Delta updates. For each Delta update, the Delta ID is incremented, while the Full ID remains the same. After some time it may decide to send a new full update. The Full ID in the full update MUST be larger than the Full ID in the previous update, and Delta ID is reset to zero. A Full update always has Delta ID zero, and a Delta update always has a non-zero Delta ID.

When a router receives an update it performs RPF check as in regular PFM, boundary processing as in regular PFM. For each interface where

the update would have been forwarded in regular PFM, it will be sent over PORT to all PFM PORT neighbors on the interface. If there are any neighbors on the interface not supporting PFM PORT it MAY revert to sending unreliable PFM messages.

When a router receives a Full update it will remove any stored information from the originator and store the information in the new update. When it receives a Delta update it stores the update and keeps all previous information.

Due to routers being restarted, PORT connections going down etc., some routers MAY have missed some updates, potentially not having received any updates when restarting. In order to receive the most recent data from a neighbor it sends a PORT PFM Request message. For each originator the router has stored information from, it will include an option indicating the Full and Delta IDs of the last message received from that originator. A router receiving the Request compares the IDs of the specified originators with the latest data it has for these originators. If it has a more recent full update, it will first send it to the neighbor. Next, if it has more recent delta updates, it will send all the delta updates in the order they were received. This means that the requesting router receives the messages in order. It will first get a full update if a more recent version exists. The ID of this update may be much larger than the previously seen ID. The first Delta update received, if any, will have ID one if a Full update was received, or one larger than the Delta ID in the request, if not. If multiple Delta updates are received, the Delta ID will increment by one for each update. If the router has stored information for any originators not included in the

request message, it will also send this information. It will first send the stored Full update, and then the Delta updates. As discussed above, the Delta updates MUST be sent in the order they were received, first sending update one, then update two, and so forth.

The Delta ID is an unsigned 16 bits value. It never wraps around. A router MUST send a new full update if the Delta ID value is reaching its maximum value. It is RECOMMENDED having a configurable limit for how many Delta updates can be sent before sending a new Full update. Sending Full updates often is in some ways wasteful, but it limits how many deltas routers need to store, and they are also used to remove information that no longer is needed.

When a router starts up, it is RECOMMENDED that before it originates any messages, it sends a PORT PFM Request message to receive any updates that neighbors may have stored for the originator address it would use. It could simply not include an option with the originator address it would use, and receive any information neighbors may have,

or it could include an option, but with the Full ID set to a value smaller than the Full ID it would use for the next Full Update. E.g., if the ID is based on the number of seconds since the epoch, it could send a request based on the current time. It would then normally get no updates from the neighbors with its own ID. If it does, it is RECOMMENDED to log an error, and ensure that the Full IDs of the next future Full Updates are larger than what was received.

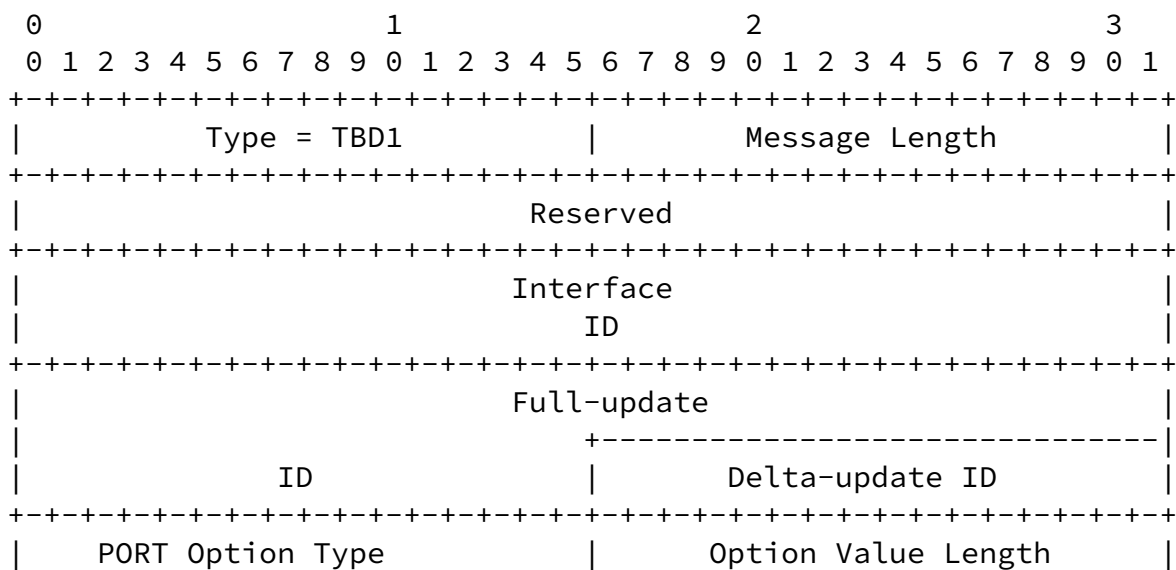
In order to handle extra ordinary cases where a router has originated messages with an erroneously large Full ID, it is RECOMMENDED that implementations provides a way for an administrator to clear the stored PFM state on a router, as well as a way to trigger sending of a Full Update on an originator. This means that as a last resort, an administrator could clear the state for an originator on all the routers, and optionally afterwards trigger a full update by the originator.

#### [4.](#) PFM over PORT message definitions

We define a new PORT message for sending a PFM message. This consists of an update version and a new PORT option containing a PFM message as defined in [[RFC8364](#)]. We also define a new PORT message

for requesting a PFM update from a neighbor. This contains the latest update version that the router has from each originator and requests the neighbor to transmit any information that it is missing.

#### [4.1.](#) PORT PFM Update

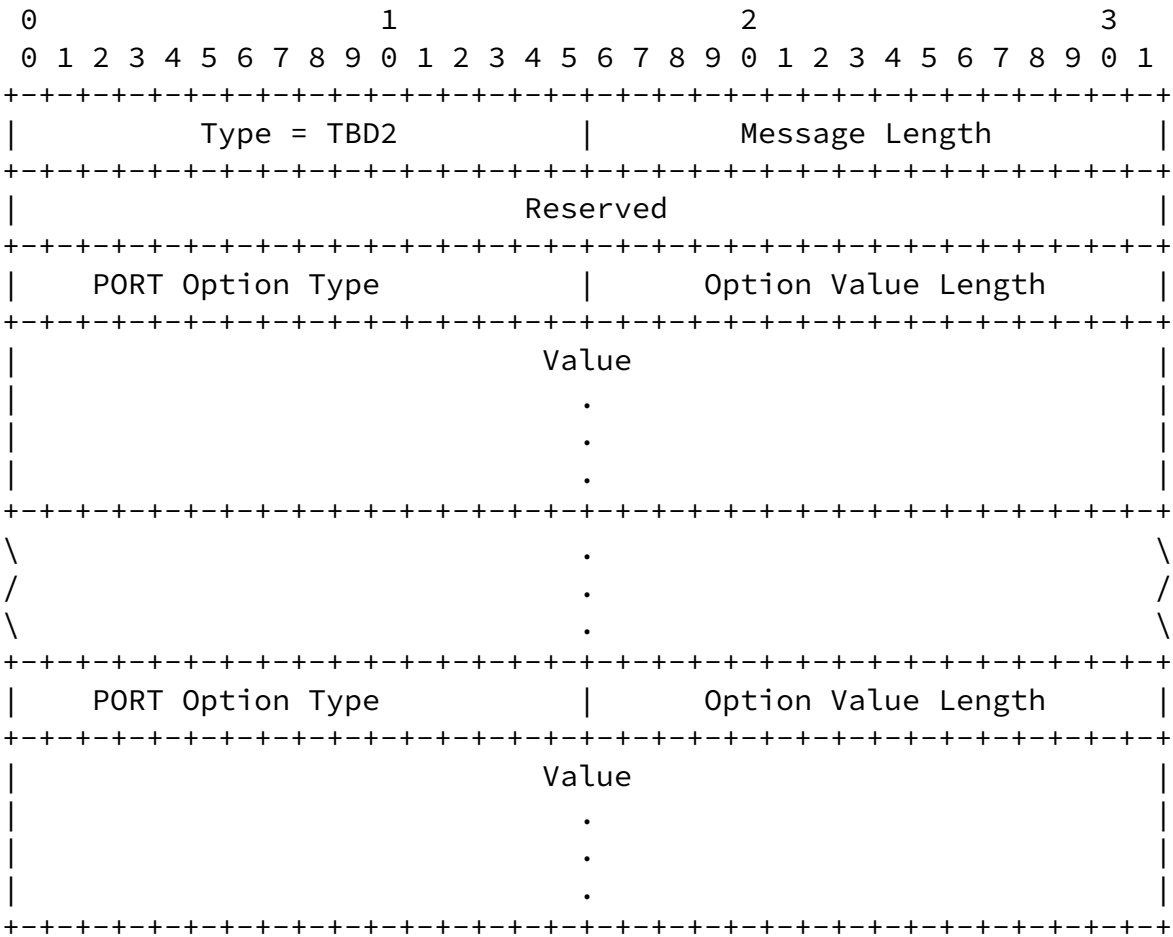




delta. Note that the Full-update ID is also used for a delta. If this is a full update, delta-update is set to 0. This is a 16 bit value.

PORT Options: The general format is defined in [\[RFC6559\] section 5.3](#). This message MUST contain exactly one PFM Update PORT option. The PFM Update PORT option is defined below. It MAY contain other options that are defined for use in a PORT PFM Update message.

4.2. PORT PFM Request



Type: Type is TBD.

Message Length: Length in bytes for the value part of the Type/

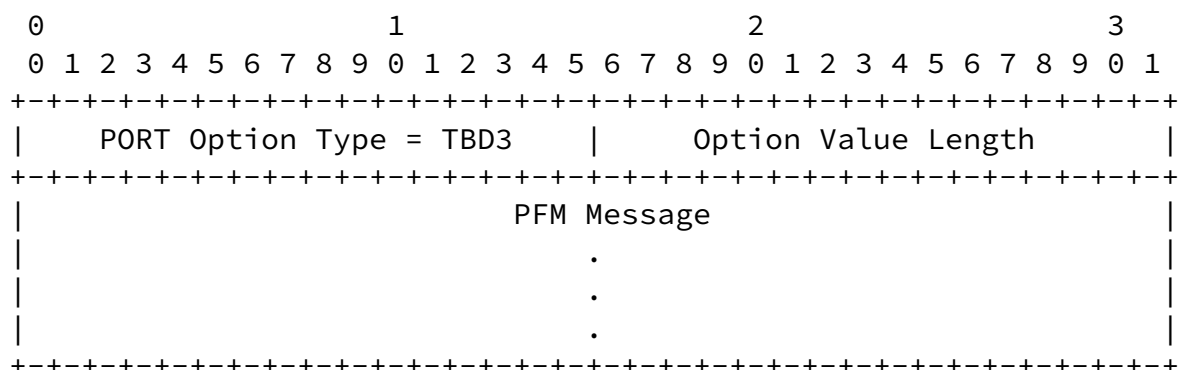


Length/Value encoding. If no PORT Options are included, the length is 12. If n PORT Options with Option Value lengths L1, L2, ..., Ln are included, the message length is 12 + 4\*n + L1 + L2 + ... + Ln.

Reserved: Set to zero on transmission and ignored on receipt.

PORT Options: The general format is defined in [\[RFC6559\]](#) [section 5.3](#). This message MAY contain zero, one or multiple PFM Request PORT options. The options indicate which versions the requesting router has from which originators; one option per originator. No options, means that the requesting router wants a full update for all known originators. The PFM Request PORT option is defined below. It MAY contain other options that are defined for use in a PORT PFM Request message.

#### [4.3.](#) PORT PFM Update Option

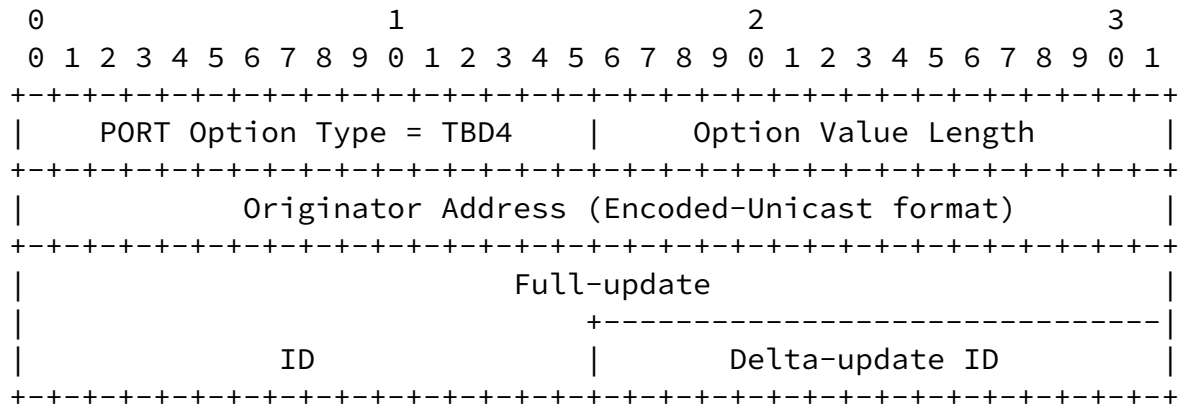


Type: Type is TBD.

Option Value Length: The number of octets that make up the PFM Message.

PFM Message: A PFM Message as defined in [\[RFC8364\]](#).

#### [4.4.](#) PORT PFM Request Option



Type: Type is TBD.

Option Value Length: The length in octets of the originator address plus 6.

Originator Address: The address of an originator as defined in [\[RFC8364\]](#).

Full-update ID: The ID of the last full update that the router has stored. It is requesting getting the most recent newer full update, if any exists. Plus, any deltas after the last full update.

Delta-update ID: The ID of the last delta update that the router has stored. It is requesting getting the most recent newer full update, using the Full-update ID, if it exists plus any deltas after that. If there are no more recent full updates, then it is requesting any delta updates more recent than this ID.

## [5.](#) Security Considerations

To be completed. Largely similar to the considerations for PIM PORT. One may use TCP/SCTP authentication mechanisms.

## [6.](#) IANA considerations

To be completed. IANA would need to assign types for the messages and options defined.

## [7.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997,

Internet-Draft

PIM PORT PFM

March 2019

- [RFC6559] Farinacci, D., Wijnands, IJ., Venaas, S., and M. Napierala, "A Reliable Transport Mechanism for PIM", [RFC 6559](#), DOI 10.17487/RFC6559, March 2012, <<https://www.rfc-editor.org/info/rfc6559>>.
- [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>>.
- [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

Stig Venaas  
Cisco Systems, Inc.  
Tasman Drive  
San Jose CA 95134  
USA

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

Balaji Ganesh  
Cisco Systems, Inc.  
Tasman Drive  
San Jose CA 95134  
USA

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

Kesavan Thiruvengatasamy  
Cisco Systems, Inc.  
Tasman Drive  
San Jose CA 95134  
USA

Email: kethiruv@cisco.com

Venaas, et al.

Expires September 12, 2019

[Page 10]

---

Internet-Draft

PIM PORT PFM

March 2019

Ramakrishnan Chokkanathapuram  
Cisco Systems, Inc.  
Tasman Drive  
San Jose CA 95134  
USA

Email: ramaksun@cisco.com

