Network Working Group
Internet-Draft
Intended status: Experimental
Expires: August 4, 2018

IJ. Wijnands
S. Venaas
Cisco Systems, Inc.
M. Brig
Aegis BMD Program Office
A. Jonasson
Swedish Defence Material Administration (FMV)

January 31, 2018

PIM Flooding Mechanism and Source Discovery
draft-ietf-pim-source-discovery-bsr-12

Abstract

PIM Sparse-Mode (PIM-SM) uses a Rendezvous Point (RP) and shared trees to forward multicast packets from new sources. Once last hop routers receive packets from a new source, they may join the Shortest Path Tree for the source for optimal forwarding. This draft defines a new mechanism that provides a way to support PIM-SM without the need for PIM registers, RPs or shared trees. Multicast source information is flooded throughout the multicast domain using a new generic PIM flooding mechanism. This allows last hop routers to learn about new sources without receiving initial data packets.

Table of Contents

1. Introduction

PIM Sparse-Mode (PIM-SM) [RFC7761] uses a Rendezvous Point (RP) and shared trees to forward multicast packets to Last Hop Routers (LHR). After the first packet is received by a LHR, the source of the multicast stream is learned and the Shortest Path Tree (SPT) can be joined. This draft defines a new mechanism that provides a way to support PIM-SM without the need for PIM registers, RPs or shared trees. Multicast source information is flooded throughout the multicast domain using a new generic PIM flooding mechanism. By

removing the need for RPs and shared trees, the PIM-SM procedures are simplified, improving router operations, management and making the protocol more robust. Also the data packets are only sent on the SPTs, providing optimal forwarding.

This mechanism has some similarities to PIM Dense Mode (PIM-DM) with its State-Refresh signaling [RFC3973], except that there is no initial flooding of data packets for new sources. It provides the traffic efficiency of PIM-SM, while being as easy to deploy as PIM-DM. The downside is that it cannot provide forwarding of initial packets from a new source, see Section 4.4. PIM-DM is very different from PIM-SM and not as mature, Experimental vs Internet Standard, and there are only a few implementations. The solution in this document consists of a lightweight source discovery mechanism on top of the Source-Specific Multicast (SSM) [RFC4607] parts of PIM-SM. It is feasable to implement only a subset of PIM-SM to provide SSM support, and in addition implement the mechanism in this draft to offer a source discovery mechanism for applications that do not provide their own source discovery.

This document defines a generic flooding mechanism for distributing information throughout a PIM domain. While the forwarding rules are largely similar to Bootstrap Router mechanism (BSR) [RFC5059], any router can originate information, and it allows for flooding of any kind of information. Each message contains one or more pieces of information encoded as TLVs (type, length and value). This document defines one TLV used for distributing information about active multicast sources. Other documents may define additional TLVs.

Note that this document is experimental. While the flooding mechanism is largely similar to BSR, there are some concerns about scale as there can be multiple routers distributing information, and potentially larger amount of data that needs to be processed and stored. Distributing knowledge of active sources in this way is new, and there are some concerns, mainly regarding potentially large amounts of source states that need to be distributed. While there has been some testing in the field, we need to learn more about the forwarding efficiency, both the amount of processing per router, and propagation delay, and the amount of state that can be distributed. In particular, how many active sources one can support without consuming too many resources. There are also parameters, see Section 5, that can be tuned regarding how frequently information is distributed, and it is not clear what parameters are useful for different types of networks.

1.1. Conventions Used in This Document

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 RFC 2119 [RFC2119].

1.2. Terminology

RP: Rendezvous Point

BSR: Bootstrap Router

RPF: Reverse Path Forwarding

SPT: Shortest Path Tree

FHR: First Hop Router, directly connected to the source

LHR: Last Hop Router, directly connected to the receiver

PFM: PIM Flooding Mechanism

PFM-SD: PFM Source Discovery

SG Mapping: Multicast source group mapping

2. Testing and Deployment Experiences

A prototype of this specification has been implemented and there has been some limited testing in the field. The prototype was tested in a network with low bandwidth radio links. The network has frequent topology changes, including frequent link or router failures. Previously existing mechanisms like PIM-SM and PIM-DM were tested.

With PIM-SM the existing RP election mechanisms were found to be too slow. With PIM-DM, issues were observed with new multicast sources starving low bandwidth links even when there are no receivers, in some cases such that there was no bandwidth left for prune messages.

For the PFM-SD prototype tests, all routers were configured to send PFM-SD for directly connected source and to cache received announcements. Applications such as SIP with multicast subscriber discovery, multicast voice conferencing, position tracking and NTP were successfully tested. The tests went quite well. Packets were rerouted as needed and there were no unnecessary forwarding of packets. Ease of configuration was seen as a plus.

3. A Generic PIM Flooding Mechanism

The Bootstrap Router mechanism (BSR) [RFC5059] is a commonly used mechanism for distributing dynamic Group to RP mappings in PIM. It is responsible for flooding information about such mappings throughout a PIM domain, so that all routers in the domain can have the same information. BSR as defined, is only able to distribute Group to RP mappings. This document defines a more generic mechanism that can flood any kind of information. Administrative boundaries, see Section 3.2, may be configured to limit to which parts of a network the information is flooded.

The forwarding rules are identical to BSR, except that one can control whether routers should forward unsupported data types. For some types of information it is quite useful that it can be distributed without all routers having to support the particular type, while there may also be types where it is necessary for every single router to support it. The mechanism includes an originator address which is used for RPF checking to restrict the flooding, and prevent loops, just like BSR. Like BSR, messages are forwarded hop by hop; the messages are link-local and each router will process and resend the messages. Note that there is no equivalent to the BSR election mechanism; there can be multiple originators. This mechanism is named the PIM Flooding Mechanism (PFM).

3.1. PFM Message Format

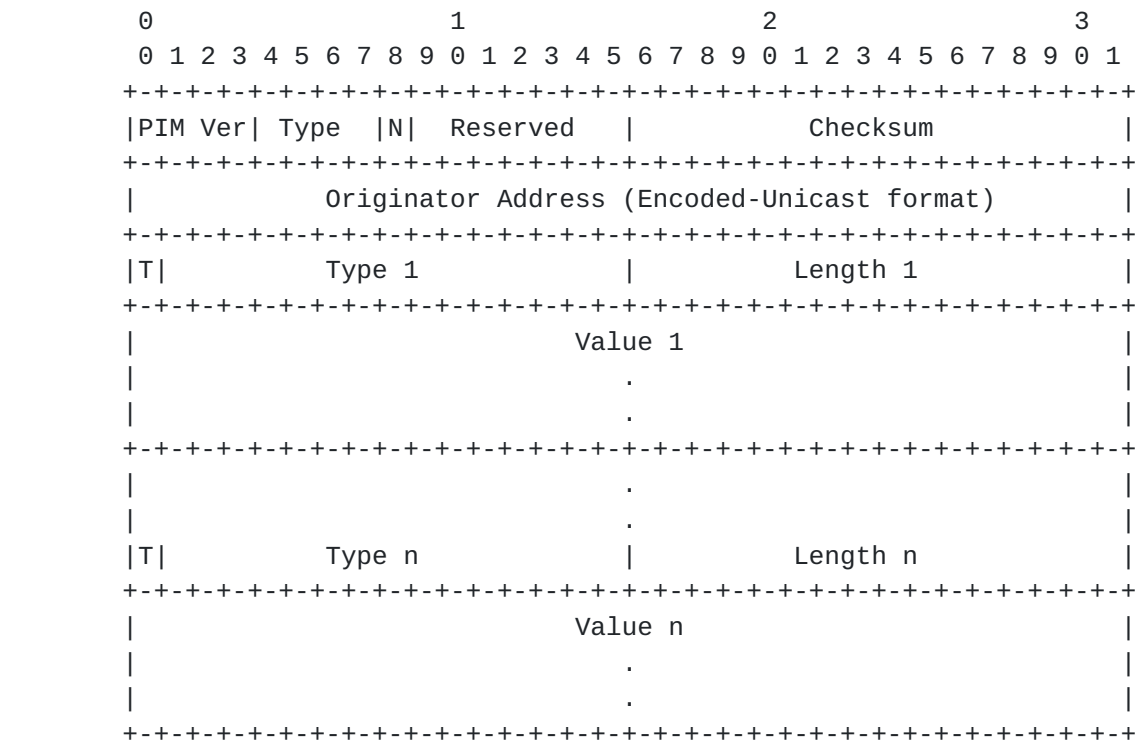

PIM Version, Reserved and Checksum: As specified in [RFC7761].

Type: PIM Message Type. Value (pending IANA) for a PFM message.

[N]o-Forward bit: When set, this bit means that the PFM message is not to be forwarded. This bit is defined to prevent Bootstrap message forwarding in [RFC5059].

Originator Address: The address of the router that originated the message. This can be any address assigned to the originating router, but MUST be routable in the domain to allow successful forwarding. The format for this address is given in the Encoded-Unicast address in [RFC7761].

[T]ransitive bit: Each TLV in the message includes a bit called the Transitive bit that controls whether the TLV is forwarded by routers that do not support the given type. See Section 3.4.2.

Type 1..n: A message contains one or more TLVs, in this case n TLVs. The Type specifies what kind of information is in the Value. The type range is from 0 to 32767 (15 bits).

Length 1..n: The length of the the value field in octets.

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

3.2. Administrative Boundaries

PFM messages are generally forwarded hop by hop to all PIM routers. However, similar to BSR, one may configure administrative boundaries to limit the information to certain domains or parts of the network. Implementations MUST have a way of defining a set of interfaces on a router as administrative boundaries for all PFM messages, or optionally for certain TLVs, allowing for different boundaries for different TLVs. Usually one wants boundaries to be bidirectional, but an implementation MAY also provide unidirectional boundaries. When forwarding a message, a router MUST NOT send it out an interface that is an outgoing boundary, including bidirectional boundary, for all PFM messages. If an interface is an outgoing boundary for certain TLVs, the message MUST NOT be sent out the interface if it is a boundary for all the TLVs in the message. Otherwise the router MUST remove all the boundary TLVs from the message and send the message with the remaining TLVs. Also, when receiving a PFM message on an interface, the message MUST be discarded if the interface is an incoming boundary, including bidirectional boundary, for all PFM messages. If the interface is an incoming boundary for certain TLVs, the router MUST ignore all boundary TLVs. If all the TLVs in the message are boundary TLVs, then the message is effectively ignored. Note that when forwarding an incoming message, the boundary is applied before forwarding. If the message was discarded or all the TLVs were ignored, then no message is forwarded. When a message is forwarded, it MUST NOT contain any TLVs for which the incoming interface is an incoming, or bidirectional, boundary.

3.3. Originating PFM Messages

A router originates a PFM message when it needs to distribute information using a PFM message to other routers in the network. When a message is originated depends on what information is distributed. For instance this document defines a TLV to distribute information about active sources. When a router has a new active source, a PFM message should be sent as soon as possible. Hence a PFM message should be sent every time there is a new active source. However, the TLV also contains a holdtime and PFM messages need to be sent periodically. Generally speaking, a PFM message would typically be sent when there is a local state change, causing information to be distributed with PFM to change. Also, some information may need to be sent periodically. These messages are called triggered and periodic messages, respectively. Each TLV definition will need to define when a triggered PFM message needs to be originated, and also whether to send periodic messages, and how frequent.

A router MUST NOT originate more than Max_PFM_Message_Rate messages per minute. This document does not mandate how this should be implemented, but some possible ways could be having a minimal time between each message, counting the number of messages originated and resetting the count every minute, or using a leaky bucket algorithm. One benefit of using a leaky bucket algorithm is that it can handle bursts better. The default value of Max_PFM_Message_Rate is 6. The value MUST be configurable. Depending on the network, one may want to use a larger value of Max_PFM_Message_Rate to favor propagation of new information, but with a large number of routers and many updates, the total number of messages might become too large and require too much processing.

There MUST be a minimum of Min_PFM_Message_Gap milliseconds between each originated message. The default value of Min_PFM_Message_Gap is 1000 (1 second). The value MUST be configurable.

Unless otherwise specified by the TLV definitions, there is no relationship between different TLVs, and an implementation can choose whether to combine TLVs in one message or across separate messages. It is RECOMMENDED to combine multiple TLVs in one message, to reduce the number of messages, but it is also RECOMMENDED that the message is small enough to avoid fragmentation at the IP layer. When a triggered PFM message needs to be sent due to a state change, a router MAY send a message containing only the information that changed. If there are many changes occuring at about the same time, it might be possible to combine multiple changes in one message. In the case where periodic messages are also needed, an implementation MAY include periodic PFM information in a triggered PFM. E.g., if some information needs to be sent every 60 seconds and a triggered PFM is about to be sent 20 seconds before the next periodic PFM was scheduled, the triggered PFM might include the periodic information and the next periodic PFM can then be scheduled 60 seconds after that, rather than 20 seconds later.

When a router originates a PFM message, it puts one of its own addresses in the originator field. An implementation MUST allow an administrator to configure which address is used. For a message to be received by all routers in a domain, all the routers need to have a route for this address due to the RPF based forwarding. Hence an administrator needs to be careful which address to choose. When this is not configured, an implementation MUST NOT use a link-local address. It is RECOMMENDED to use an address of a virtual interface such that the originator can remain unchanged and routable independent of which physical interfaces or links may go down.

The No-Forward bit MUST NOT be set, except for the case when a router receives a PIM Hello from a new neighbor, or a PIM Hello with a new

Generation Identifier, defined in [RFC7761], is received from an existing neighbor. In that case an implementation MAY send PFM messages containing relevant information so that the neighbor can quickly get the correct state. The definition of the different PFM message TLVs need to specify what, if anything, needs to be sent in this case. If such a PFM message is sent, the No-Forward bit MUST be set, and the message must be sent within 60 seconds after the neighbor state change. The processing rules for PFM messages will ensure that any other neighbors on the same link ignores the message. This behavior and the choice of 60 seconds is similar to what is defined for the No-Forward bit in [RFC5059].

3.4. Processing PFM Messages

A router that receives a PFM message MUST perform the initial checks specified here. If the checks fail, the message MUST be dropped. An error MAY be logged, but otherwise the message MUST be dropped silently. If the checks pass, the contents is processed according to the processing rules of the included TLVs.

3.4.1. Initial Checks

In order to do further processing, a message MUST meet the following requirements. The message MUST be from a directly connected PIM neighbor, the destination address MUST be ALL-PIM-ROUTERS. Also, the interface MUST NOT be an incoming, nor bidirectional, administrative boundary for PFM messages, see Section 3.2. If No-Forward is not set, the message MUST be from the RPF neighbor of the originator address. If No-Forward is set, this system, the router doing these checks, MUST have enabled the PIM protocol within the last 60 seconds. See Section 3.3 for details. In pseudo-code the algorithm is as follows:

```
if ((DirectlyConnected(PFM.src_ip_address) == FALSE) OR
    (PFM.src_ip_address is not a PIM neighbor) OR
    (PFM.dst_ip_address != ALL-PIM-ROUTERS) OR
    (Incoming interface is admin boundary for PFM)) {
     drop the message silently, optionally log error.
}
if (PFM.no_forward_bit == 0) {
    if (PFM.src_ip_address !=
        RPF_neighbor(PFM.originator_ip_address)) {
        drop the message silently, optionally log error.
    }
} else if (more than 60 seconds elapsed since PIM enabled)) {
    drop the message silently, optionally log error.
}
```

Note that src_ip_address is the source address in the IP header of the PFM message. Originator is the originator field inside the PFM message, and is the router that originated the message. When the message is forwarded hop by hop, the originator address never changes, while the source address will be an address belonging to the router that last forwarded the message.

3.4.2. Processing and Forwarding of PFM Messages

When the message is received, the initial checks above must be performed. If it passes the checks, then for each included TLV, perform processing according to the specification for that TLV.

After processing, the messsage is forwarded. Some TLVs may be omitted or modified in the forwarded message. This depends on administrative boundaries, see Section 3.2, the type specification and the setting of the Transitive bit for the TLV. If a router supports the type, then the TLV is forwarded with no changes unless otherwise specified by the type specification. A router not supporting the given type MUST include the TLV in the forwarded message if and only if the Transitive bit is set. Whether a router supports the type or not, the value of the Transitive bit MUST be preserved if the TLV is included in the forwarded message. The message is forwarded out of all interfaces with PIM neighbors (including the interface it was received on). As specified in Section 3.2, if an interface is an outgoing boundary for any TLVs, the message MUST NOT be sent out the interface if it is an outgoing boundary for all the TLVs in the message. Otherwise the router MUST remove any outgoing boundary TLVs of the interface from the message and send the message out that interface with the remaining TLVs.

4. Distributing Source Group Mappings

The generic flooding mechanism (PFM) defined in the previous section can be used for distributing source group mappings about active multicast sources throughout a PIM domain. A Group Source Holdtime (GSH) TLV is defined for this purpose.

4.1. Group Source Holdtime TLV

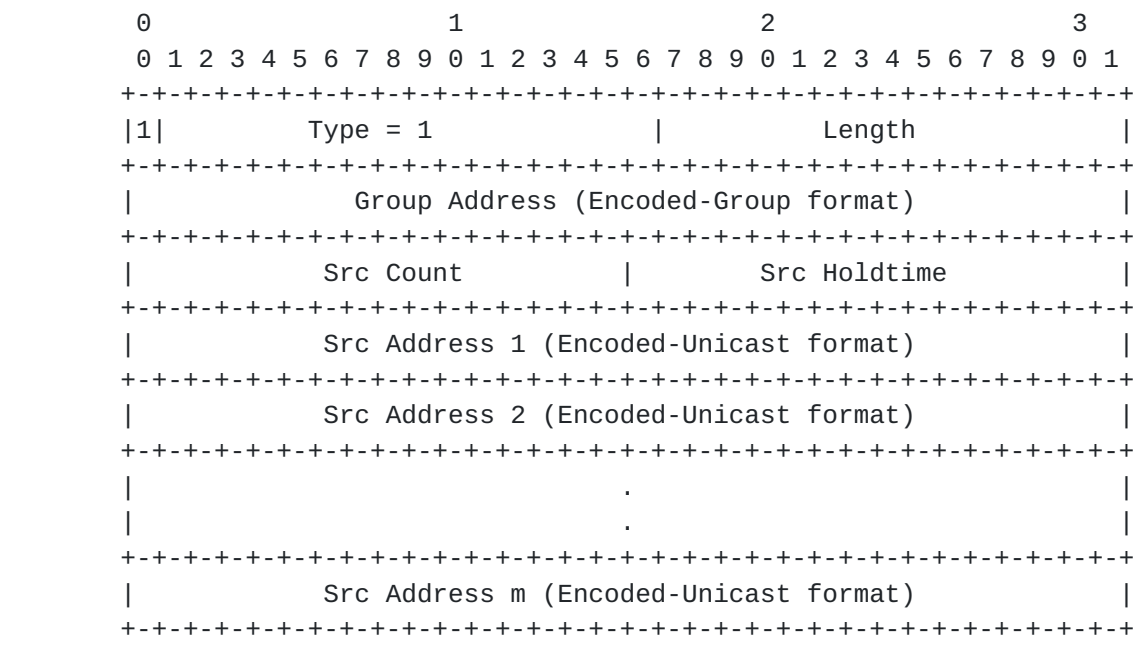

1: The Transitive bit is set to 1. This means that this type will be forwarded even if a router does not support it. See Section 3.4.2.

Type: This TLV has type 1.

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].

Src Count: The number of source addresses that are included.

Src Holdtime: The Holdtime (in seconds) for the included source(s).

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

4.2. Originating Group Source Holdtime TLVs

A PFM message MAY contain one or more Group Source Holdtime (GSH) TLVs. This is used to flood information about active multicast sources. Each FHR that is directly connected to an active multicast source originates PFM messages containing GSH TLVs. How a multicast router discovers the source of the multicast packet and when it considers itself the FHR follows the same procedures as the registering process described in [RFC7761]. When a FHR has decided

that a register needs to be sent per [RFC7761], the SG is not registered via the PIM-SM register procedures, but the SG mapping is included in an GSH TLV in a PFM message. Note, only the SG mapping is distributed in the message, not the entire packet as would have been done with a PIM register.

The PFM messages containing the GSH TLV are sent periodically for as long as the multicast source is active, similar to how PIM registers are sent periodically. This means that as long as the source is active, it is included in a PFM message originated every Group_Source_Holdtime_Period seconds, within the general PFM timing requirements in Section 3.3. The default value of Group_Source_Holdtime_Period is 60. The value MUST be configurable. The holdtime for the source MUST be set to either zero or Group_Source_Holdtime_Holdtime. The value of the Group_Source_Holdtime_Holdtime parameter MUST be larger than Group_Source_Holdtime_Period. It is RECOMMENDED to be 3.5 times the Group_Source_Holdtime_Period. The default value is 210 (seconds). The value MUST be configurable. A source MAY be announced with a holdtime of zero to indicate that the source is no longer active.

If an implementation supports originating GSH TLVs with different holdtimes for different sources, it can if needed send multiple TLVs with the same group address. Due to the format, all the sources in the same TLV have the same holdtime.

When a new source is detected, an implementation MAY send a PFM message containing just that particular source. However, it MAY also include information about other sources that were just detected, sources that are scheduled for periodic announcement later, or other types of information. See Section 3.3 for details. Note that when a new source is detected, one should trigger sending of a PFM message as soon as possible, while if a source becomes inactive, there is no reason to trigger a message. There is no urgency in removing state for inactive sources. Note that the message timing requirements in Section 3.3 apply. This means that one cannot always send a triggered message immediately when a new source is detected. In order to meet the timing requirements, sending of the message may have to be delayed a small amount of time.

When a new PIM neighbor is detected, or an existing neighbor changes Generation Identifier, an implementation MAY send a triggered PFM message containing GSH TLVs for any Source Group mappings it has learned by receiving PFM GSH TLVs as well as any active directly connected sources. See Section 3.3 for further details.

4.3. Processing GSH TLVs

A router that receives a PFM message containing GSH TLVs MUST parse the GSH TLVs and store each of the GSH TLVs as SG mappings with a holdtimer started with the advertised holdtime, unless the implementation specifically does not support GSH TLVs, the router is configured to ignore GSH TLVs in general, or to ignore GSH TLVs for certain sources or groups. In particular, an administrator might configure a router to not process GSH TLVs if the router is known to never have any directly connected receivers.

For each group that has directly connected receivers, this router SHOULD send PIM (S,G) joins for all the SG mappings advertised in the message for the group. Generally joins are sent, but there could for instance be administrative policy limiting which sources and groups to join. The SG mappings are kept alive for as long as the holdtimer for the source is running. Once the holdtimer expires a PIM router MAY send a PIM (S,G) prune to remove itself from the tree. However, when this happens, there should be no more packets sent by the source, so it may be desirable to allow the state to time out rather than sending a prune.

Note that a holdtime of zero has a special meaning. It is to be treated as if the source just expired, and state to be removed. Source information MUST NOT be removed due to the source being omitted in a message. For instance, if there is a large number of sources for a group, there may be multiple PFM messages, each message containing a different list of sources for the group.

4.4. The First Packets and Bursty Sources

The PIM register procedure is designed to deliver Multicast packets to the RP in the absence of a Shortest Path Tree (SPT) from the RP to the source. The register packets received on the RP are decapsulated and forwarded down the shared tree to the LHRs. As soon as an SPT is built, multicast packets would flow natively over the SPT to the RP or LHR and the register process would stop. The PIM register process ensures packet delivery until an SPT is in place reaching the FHR. If the packets were not unicast encapsulated to the RP they would be dropped by the FHR until the SPT is setup. This functionality is important for applications where the initial packet(s) must be received for the application to work correctly. Another reason would be for bursty sources. If the application sends out a multicast packet every 4 minutes (or longer), the SPT is torn down (typically after 3:30 minutes of inactivity) before the next packet is forwarded down the tree. This will cause no multicast packet to ever be forwarded. A well behaved application should be able to deal with

packet loss since IP is a best effort based packet delivery system. But in reality this is not always the case.

With the procedures defined in this document the packet(s) received by the FHR will be dropped until the LHR has learned about the source and the SPT is built. That means for bursty sources or applications sensitive for the delivery of the first packet this solution would not be very applicable. This solution is mostly useful for applications that don't have strong dependency on the initial packet(s) and have a fairly constant data rate, like video distribution for example. For applications with strong dependency on the initial packet(s) using PIM Bidir [RFC5015] or SSM [RFC4607] is recommended. The protocol operations are much simpler compared to PIM SM, it will cause less churn in the network and both guarantee best effort delivery for the initial packet(s).

4.5. Resiliency to Network Partitioning

In a PIM SM deployment where the network becomes partitioned, due to link or node failure, it is possible that the RP becomes unreachable to a certain part of the network. New sources that become active in that partition will not be able to register to the RP and receivers within that partition are not able to receive the traffic. Ideally you would want to have a candidate RP in each partition, but you never know in advance which routers will form a partitioned network. In order to be fully resilient, each router in the network may end up being a candidate RP. This would increase the operational complexity of the network.

The solution described in this document does not suffer from that problem. If a network becomes partitioned and new sources become active, the receivers in that partitioned will receive the SG Mappings and join the source tree. Each partition works independently of the other partition(s) and will continue to have access to sources within that partition. Once the network has healed, the periodic flooding of SG Mappings ensures that they are re-flooded into the other partition(s) and other receivers can join to the newly learned sources.

5. Configurable Parameters

This document contains a number of configurable parameters. These parameters are formally defined in Section 3.3 and Section 4.2, but they are repeated here for ease of reference. These parameters all have default values as noted below.

Max_PFM_Message_Rate: The maximum number of PFM messages a router is allowed to originate per minute, see Section 3.3 for details. The default value is 6.

Min_PFM_Message_Gap: The minimum amount of time between each PFM message originated by a router in milliseconds, see Section 3.3 for details. The default is 1000.

Group_Source_Holdtime_Period: The announcement period for Group Source Holdtime TLVs in seconds, see Section 4.2 for details. The default value is 60.

Group_Source_Holdtime_Holdtime: The holdtime for Group Source Holdtime TLVs in seconds, see Section 4.2 for details. The default value is 210.

6. Security Considerations

When it comes to general PIM message security, see [RFC7761]. PFM messages MUST only be accepted from a PIM neighbor, but as discussed in [RFC7761], any router can become a PIM neighbor by sending a Hello message. To control from where to accept PFM packets, one can limit which interfaces PIM is enabled, and also one can configure interfaces as administrative boundaries for PFM messages, see Section 3.2. The implications of forged PFM messages depend on which TLVs they contain. Documents defining new TLVs will need to discuss the security considerations for the specific TLVs. In general though, the PFM messages are flooded within the network, and by forging a large number of PFM messages one might stress all the routers in the network.

If an attacker can forge PFM messages, then such messages may contain arbitrary GSH TLVs. An issue here is that an attacker might send such TLVs for a huge amount of sources, potentially causing every router in the network to store huge amounts of source state. Also, if there is receiver interest for the groups specified in the GSH TLVs, routers with directly connected receivers will build Shortest Path Trees for the announced sources, even if the sources are not actually active. Building such trees will consume additional resources on routers that the trees pass through.

PIM-SM link-local messages can be authenticated using IPsec, see [RFC7761] section 6.3 and [RFC5796]. Since PFM messages are link-local messages sent hop by hop, a link-local PFM message can be authenticated using IPsec such that a router can verify that a message was sent by a trusted neighbor and has not been modified. However, to verify that a received message contains correct information announced by the originator specified in the message, one

will have to trust every router on the path from the originator and that each router has authenticated the received message.

7. IANA Considerations

This document requires the assignment of a new PIM message type for the PIM Flooding Mechanism (PFM) with the name "PIM Flooding Mechanism". IANA is also requested to create a registry for PFM TLVs called "PIM Flooding Mechanism Message Types". Assignments for the registry are to be made according to the policy "IETF Review" as defined in [RFC8126]. The initial content of the registry should be:

Type	Name	Reference
0	Reserved	[this document]
1	Source Group Holdtime	[this document]
2-32767	Unassigned	

8. Acknowledgments

The authors would like to thank Arjen Boers for contributing to the initial idea, and David Black, Stewart Bryant, Yiqun Cai, Papadimitriou Dimitri, Toerless Eckert, Dino Farinacci, Alvaro Retana and Liang Xia for their very helpful comments on the draft.

9. References

9.1. 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, <https://www.rfc-editor.org/info/rfc2119>.

[RFC5059] Bhaskar, N., Gall, A., Lingard, J., and S. Venaas, "Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)", RFC 5059, DOI 10.17487/RFC5059, January 2008, <https://www.rfc-editor.org/info/rfc5059>.

[RFC5796] Atwood, W., Islam, S., and M. Siami, "Authentication and Confidentiality in Protocol Independent Multicast Sparse Mode (PIM-SM) Link-Local Messages", RFC 5796, DOI 10.17487/RFC5796, March 2010, <https://www.rfc-editor.org/info/rfc5796>.

[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>.

9.2. Informative References

[RFC3973] Adams, A., Nicholas, J., and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)", RFC 3973, DOI 10.17487/RFC3973, January 2005, <https://www.rfc-editor.org/info/rfc3973>.

[RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, DOI 10.17487/RFC4607, August 2006, <https://www.rfc-editor.org/info/rfc4607>.

[RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", RFC 5015, DOI 10.17487/RFC5015, October 2007, <https://www.rfc-editor.org/info/rfc5015>.

Authors' Addresses

IJsbrand Wijnands
Cisco Systems, Inc.
De kleetlaan 6a
Diegem 1831
Belgium

Email: ice@cisco.com

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

Email: stig@cisco.com

Michael Brig
Aegis BMD Program Office
17211 Avenue D, Suite 160
Dahlgren VA 22448-5148
USA

Email: michael.brig@mda.mil

Anders Jonasson
Swedish Defence Material Administration (FMV)
Loennvaegen 4
Vaexjoe 35243
Sweden

Email: anders@jomac.se