

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 30, 2015

P. Pfister
October 27, 2014

Protocol Independent Multicast Border Proxying
draft-pfister-pim-border-proxy-00

Abstract

This document describes an extension to the Protocol Independent Multicast (PIM) multicast routing protocol that enables PIM proxying from PIM domains toward other multicast domains. It relies on PIM Proxies located on domain borders and Proxy Controllers which can be located anywhere. This document describes how subscriptions can be proxied toward domain's border interfaces using MLDv2 and IGMPv3, but other protocols could be used as well. Once multicast traffic is received on a proxied interface, it can be forwarded as if originated by a directly connected source.

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 <http://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 April 30, 2015.

Copyright Notice

Copyright (c) 2014 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 (<http://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

1. Introduction	2
2. Protocol Specifications	3
2.1. Proxy Controller Behavior	3
2.2. Proxy Behavior	3
2.3. PIM Proxy Message format	4
2.3.1. Message Header	4
2.3.2. State Update Message	4
2.3.3. Keep Alive Message	5
3. Different proxy types	6
3.1. IGMP/MLD proxy	6
3.2. IGMP/MLD controller	7
4. Security Considerations	7
5. IANA Considerations	7
6. References	7
6.1. Normative References	7
6.2. Informative References	8
Appendix A. Acknowledgments	8
Appendix B. Discussion	8
Author's Address	8

1. Introduction

The Protocol Independent Multicast (PIM - [RFC4601]) routing protocol initially reacts to two different event types: multicast traffic reception from a connected source, and local multicast subscription (MLD or IGMP) state change on a connected link. This approach works when the network consists of a single PIM multicast domain, but does not when border routers are connected to different multicast domains. In such situations, the border routers need to be told to which group they should subscribe to on their egress interfaces before multicast traffic can be received.

This document defines PIM Proxy's and PIM Proxy Controller's behavior. There may be one or multiple instances of each in the same PIM domain. Each controller opens a TCP connection toward every proxy it wants to interact with and sends updates every time the Proxy's state should be updated.

In PIM-SM domains, one possible application of this extension consists in using the Rendezvous as a PIM Proxy Controller for all border routers, which in turn reflects network-wide subscription

state on domain's external interfaces using MLDv2 [RFC3810] or IGMPv3 [RFC3376]. In PIM-BIDIR domains, the same approach could be used, but would not support Source-Specific multicast. Instead, every router can reflect the local subscription state of links on which it is the Designated Forwarder.

This extension was designed in order to allow ISP originated traffic to reach subscribed hosts located inside a home network [I-D.ietf-homenet-arch]. Input from PIM and Homenet working group regarding other possible solutions enabling multicast routing in home networks are very welcome.

2. Protocol Specifications

This protocol allows controllers to push PIM subscription state toward proxies. The state a given controller pushes toward a given proxy may differ depending on the proxy, the local configuration, and may not reflect the local MLD or IGMP querier state. It is an arbitrary choice and depends on the purpose of the proxy.

The following state is maintained for every established TCP connection and every PIM Group/Source state ((*,G), (S,G,rpt), (S,G)). Although this is implementation specific, the generic behavior would consist in keeping the state for every Group/Source pair in their respective Encoded-Group and Encoded-Source formats (e.g. different states whether [B]idir bit is set or not).

Mode: One of { "NoInfo", "Join", "Prune" }.

If no state is stored for a given Group/Source pair, it is equivalent to the "NoInfo" state. Similarly, after being switched to "NoInfo" state, a stored state may be removed.

2.1. Proxy Controller Behavior

A proxy controller opens a TCP connection with every proxy it wants a peering with. When a new connection is opened, the complete state is first transmitted in order to synchronize the proxy with the controller's state. Then, each time some state is changed, an update is sent through the TCP connection.

2.2. Proxy Behavior

A proxy MUST listen on port PIM_PROXY_TCP_PORT for incoming TCP connections. When a new connection is established, a new subscription state is created.

2.3. PIM Proxy Message format

2.3.1. Message Header

PIM Proxy messages are transported over TCP and make use of the following TLV format.




Figure 1

Type:

One of the defined message types.

Length:

The byte length of the value.

Value:

The value.

2.3.2. State Update Message

PIM Proxy State Update messages use the following format.




Figure 2

Type:

State Update message type (1)

Number of Sources:

The number of sources associated with the group.

State:

The state the controller wants the proxy to save for the given group/source pair.

- ```
* NoInfo (0)
* Join (1)
* Prune (2)
```

Upon reception of a State Update message, a proxy will switch the state associated with every Group/Source pairs included in the message to the state specified in the message.

### 2.3.3. Keep Alive Message

Keep Alive messages are used to keep the TCP connection open and have the following format (Most current TCP implementations support TCP keep-alives, but not all. The Keep-Alive TLV is specified because keeping the connection open is a requirement for not losing state).




Figure 3

### 3. Different proxy types

This document does not intend to specify all the possible proxy behavior. Requests could be translated into MLD and IGMP, relayed toward another separated PIM domain, translated into another multicast delivery protocol, or even be used for monitoring purposes.

Any proxy behavior CAN be overridden by local policies. For instance, proxying behavior may depend on the group's scope or firewalling rules.

Once multicast traffic is requested on an egress interface, PIM should be used as usual in order to decide whether incoming traffic should be forwarded on an ingress interface.

#### 3.1. IGMP/MLD proxy

This section describes how to translate a PIM Proxy group's state toward an MLDv2/IGMPv3 listener state. It can be used for both PIM-SM and PIM-BIDIR (states are considered regardless of the BIDIR bit).

According to PIM-SM and PIM-BIDIR specifications, (\*,G) or (S,G) can only be in state "NoInfo" or "Join", and (S,G,rpt) can only be in state "NoInfo" or "Prune".

If (\*,G) is set to "Join", the MLDv2/IGMPv3 group state should be set to "Exclude" and the Exclude Sources List should contain all sources S for which (S, G, rpt) is set to "Prune" and (S,G) is set to "NoInfo".

If (\*,G) is set to "NoInfo", the MLDv2/IGMPv3 group state should be set to "Include" and the Include Sources List should contain all sources S for which (S, G) is set to "Join".

When multiple controllers are pushing state to the same proxy, the algorithm detailed in MLDv2 and IGMPv3 specifications should be used in order to merge the different requests.

### 3.2. IGMP/MLD controller

This section describes how to translate an MLDv2/IGMPv3 querier state into a PIM subscription state.

If the group's MLDv2/IGMPv3 state is "Include", the PIM state consists in (S,G) states set to "Join" for all S in the MLDv2/IGMPv3 source include list.

If the group's MLDv2/IGMPv3 state is "Exclude", the (\*,G) state is set to "Join" and all (S,G,rpt) for S in the MLDv2/IGMPv3 source exclude list are set to "Prune".

### 4. Security Considerations

PIM Proxy messages are sent multiple hops away and are used in order to control other router's behavior. Attackers could open a connection from outside or inside the network and trigger multicast requests and forwarding. TLS or IPsec could be used in order to secure the connection. If not, connections should at least be filtered based on the controller's IP source address.

### 5. IANA Considerations

IANA is kindly requested to:

- o Reserve a TCP port number for PIM Proxies.
- o Create a new registry for PIM Proxy TLVs.
- o Create a new registry for Group/Source states.

### 6. References

#### 6.1. Normative References

- [RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", RFC 5015, October 2007.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

## 6.2. Informative References

[I-D.ietf-homenet-arch]  
Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil,  
"IPv6 Home Networking Architecture Principles", draft-  
ietf-homenet-arch-11 (work in progress), October 2013.

## Appendix A. Acknowledgments

The author would like to thank Steven Barth and Mohammed Hawari for their help in the protocol design and implementation process, as well as Mark Townsley, Stig Venaas, Markus Stenberg and IJsbrand Wijnands for their useful input.

## Appendix B. Discussion

Another message type we would probably need is a "Traffic Present" message. Sent from the Proxy toward the Controller, it would let the controller take actions when the same traffic is provided by two different border routers. But that gets inefficient when there are more than one controller.

## Author's Address

Pierre Pfister  
Paris  
France

Email: pierre@darou.fr