PIM
Internet-Draft
Intended status: Informational
Expires: 8 January 2026

H. Asaeda
NICT
L. Contreras
Telefonica

7 July 2025

Multipath Support for IGMP/MLD Proxy
draft-ietf-pim-multipath-igmpmldproxy-02

Abstract

This document describes multipath support for Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD) proxy devices. The proposed extension allows IGMP/MLD proxy devices to receive multicast sessions/channels through different upstream interfaces. An upstream interface can be selected on the basis of multiple criteria, such as channel/session IDs and interface priority values. A mechanism for upstream interface takeover that enables switching from an inactive to active upstream interface is also described.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction 2
2. Terminology 4
3. Upstream Selection Mechanism 5
3.1. Channel-Based Upstream Selection 5
3.2. Multiple Upstream Interface Selection for Robust Data Reception 5
4. Candidate Upstream Interface Configuration 5
4.1. Multicast Channel Record 5
4.2. Interface Priority 7
4.3. Default Upstream Interface 7
5. Upstream Interface Takeover 8
6. Dynamic Upstream Interface Configuration 8
6.1. Signaling-based Upstream Interface Configuration 9
6.2. Controller-based Upstream Interface Configuration 9
6.3. Updating YANG Model 9
7. Security Considerations 9
8. Summary of Aspects Requiring Further Discussion 10
9. IANA Considerations 10
10. Acknowledgements 10
11. References 10
11.1. Normative References 10
11.2. Informative References 10
Appendix A. Proof of Concept 12
Authors' Addresses 12

1. Introduction

The Internet Group Management Protocol (IGMP) [RFC3376][RFC5790] for IPv4 and the Multicast Listener Discovery Protocol (MLD) [RFC3810][RFC5790] for IPv6 are the standard protocols for hosts to initiate the joining or leaving of multicast sessions. A proxy device device that performs IGMP/MLD-based forwarding (as known as IGMP/MLD proxy) [RFC4605] maintains multicast membership information using IGMP/MLD protocols on downstream interfaces and sends IGMP/MLD membership report messages via the upstream interface to upstream multicast routers when the membership information changes (e.g., by receiving solicited/unsolicited report messages). The proxy device forwards the appropriate multicast packets received on its upstream interface to each downstream interface based on the subscription of the downstream receiver.

According to the specification of [RFC4605], an IGMP/MLD proxy has _a single_ upstream interface and one or more downstream interfaces. Upstream and downstream interfaces on the IGMP/MLD proxy device must be configured manually, and the upstream interface is expected to be connected to a wider multicast infrastructure. Therefore, IGMP/MLD proxy devices perform the router portion of the IGMP or MLD protocol on their downstream interfaces and the host portion of IGMP/MLD on their upstream interface. They must not perform the router portion of IGMP/MLD on the upstream interface.

Conversely, there is a scenario in which IGMP/MLD proxy devices enable multiple upstream interfaces and receive multicast packets through these interfaces. For example, a proxy device with more than one interface may want to access different networks, such as a global network such as the Internet and local-scope networks; a proxy device with a wired link (e.g., Ethernet) and high-speed wireless link (e.g., 5G) may want to have the capability to connect to the Internet through both links. These proxy devices receive multicast packets from different upstream interfaces and forward them to the downstream interface(s). The applicability of IGMP/MLD proxies with multiple upstream interfaces in Proxy Mobile IPv6 (PMIPv6) [RFC5213] is described in [RFC6224].

This document describes how IGMP/MLD proxy devices can receive multicast sessions/channels through different upstream interfaces. The upstream interfaces can be configured with "channel-based upstream selection." By channel-based upstream selection, IGMP/MLD proxy devices select one or multiple upstream interface(s) from the candidate upstream interfaces "on a per channel/session basis."

Unlike the conventional approach [RFC4605], when a proxy device receives an IGMP/MLD report message on the downstream interface(s), it examines the source and multicast addresses in the records of the IGMP/MLD report message and selects the appropriate upstream interface(s) based on the aforementioned configuration. A dynamic upstream selection mechanism is introduced in another document [I-D.contreras-pim-multiif-config], whereas that document will be integrated into this document.

In general, a proxy device selects "one" upstream interface from the candidate upstream interfaces per session/channel. Furthermore, a proxy device can configure to select "more than two" upstream interfaces from the candidate upstream interfaces per session/channel. In this case, it can receive duplicate (redundant) packets for the session/channel from different upstream interfaces simultaneously, resulting in "robust data reception."

A mechanism for "upstream interface takeover" is also described in this document; when the selected upstream interface is going down or the state of the link attached to the upstream interface is inactive, one of the other active candidate upstream interfaces takes over the upstream interface (if configured).

A "dynamic upstream selection" is a mechanism that selects an appropriate upstream interface(s) for sessions/channels based on the conditions of the network and adjacent routers. It is briefly introduced in this document, whereas its detailed specifications and IGMP/MLD protocol extension are described in [I-D.contreras-pim-multiif-config].

2. Terminology

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.

In addition, the following terms are used in this document.

* Selected upstream interface (or simply, upstream interface): The interface of a proxy device in the direction of the root of the multicast forwarding tree. A proxy device performs the host portion of IGMP/MLD on its upstream interface. An upstream interface is selected from a list of candidate upstream interfaces.

* Default upstream interface: An upstream interface for multicast sessions/channels for which a proxy device does not select other upstream interfaces. The default upstream interface is configurable.

* Active upstream interface: An upstream interface that has been receiving packets for specific multicast sessions/channels during a predefined active interval.

* Inactive upstream interface: An interface that has not received packets for specific multicast sessions/channels during a predefined active interval.

* Downstream interface: An interface that is not in the direction of the root of the multicast forwarding tree. A proxy device performs the router portion of IGMP/MLD on its downstream interfaces.

* Candidate upstream interface: An interface that potentially becomes an upstream interface of the proxy device. A list of candidate upstream interfaces is configured with channel/session IDs and/or priority values on an IGMP/MLD proxy device.

* Channel/session ID: It consists of source and multicast address prefixes for which a candidate upstream interface is assumed to be an upstream interface for specified multicast sessions/channels. The source or multicast address prefix can be "null".

3. Upstream Selection Mechanism

3.1. Channel-Based Upstream Selection

An IGMP/MLD proxy device selects one or multiple upstream interface(s) from candidate upstream interfaces "per channel/session" based on the "channel/session ID" configuration. This mechanism is known as "channel-based upstream selection". This mechanism enables IGMP/MLD proxy devices to use one or multiple upstream interface(s) from candidate upstream interfaces "per channel/session" based on the "channel/session ID" configuration.

3.2. Multiple Upstream Interface Selection for Robust Data Reception

When more than one candidate upstream interface is configured with the same source and multicast addresses for the "channel/session IDs" and "interface priority values" (this will be described in Section 4.2) are identical, these candidate upstream interfaces act as upstream interfaces for the sessions/channels and receive the packets simultaneously. This multiple upstream interface selection approach implements duplicate packet reception from redundant paths. This may improve the data reception quality or robustness of a session/channel, because the same multicast data packets can come from different upstream interfaces simultaneously. However, robust data reception does not guarantee packets coming from disjoint paths. It only configures the adjacent upstream routers to differ.

4. Candidate Upstream Interface Configuration

Candidate upstream interfaces are a set of interfaces from which an IGMP/MLD proxy device selects as an upstream interface. The upstream interface selection approach works with the configurations of "channel/session ID" and "interface priority value."

4.1. Multicast Channel Record

IGMP/MLD proxy devices can configure the "channel/session ID" in the multicast channel record for each candidate upstream interface.

Channel/session ID consists of source and multicast address prefixes. Source address prefixes MUST be valid unicast address prefixes, and multicast address prefixes MUST be a valid multicast address prefixes. A proxy selects an upstream interface from its candidate upstream interfaces based on the channel/session ID configuration.

The default values of these address prefixes are "null." A null source address prefix represents a wildcard source address prefix, which indicates any host. A null multicast address prefix represents a wildcard multicast address prefix, which indicates the entire multicast address range (i.e., 224.0.0.0/4 for IPv4 or ff00::/8 for IPv6).

The channel/session ID configuration comprises the source and multicast address prefixes. A candidate upstream interface with a non-null source and multicast address configuration is prioritized for upstream interface selection. For example, if a proxy device has two candidate upstream interfaces for the same multicast address prefix G_p but one of them has a non-null source address prefix S_p configuration, that candidate upstream interface is selected for the source and multicast address pair (i.e., (S_p,G_p)). The other candidate upstream interface is selected for the configured multicast address prefix, excluding the source address prefix configured by the prior interface (i.e., (*-S_p,G_p)).

The source address prefix configuration is prioritized over the multicast address prefix configuration. For example, consider the case where an IGMP/MLD proxy device has a configuration with the source address prefix S_p for candidate upstream interface A and the multicast address prefix G_p for candidate upstream interface B. When dealing with an IGMP/MLD record whose source address (S) is in the range of S_p and whose multicast address (G) is in the range of G_p, the proxy device selects candidate upstream interface A, which supports the source address prefix, as the upstream interface and transmits the (S,G) record via interface A.

In summary, the options for selecting an appropriate upstream interface are as follows:

* Association of (S,G) to a specific upstream interface, implying that subscriber requests for specific content delivered from a specific source should be received from a certain upstream interface. This condition is prioritized second.

* Association of (S,*) to a specific upstream interface, implying that subscriber requests for specific content, independent of the group identifying the content, should be received from a certain upstream interface. This condition is prioritized third.

* Association of (*,G) to a specific upstream interface, implying that subscriber requests for specific content, independent of the source of the content, should be received from a certain upstream interface. This condition is prioritized fourth.

The same address prefix can be configured on different candidate upstream interfaces. When the same address prefix is configured on different candidate upstream interfaces, an upstream interface for that address prefix is selected on the basis of each interface priority value described in Section 4.2.

4.2. Interface Priority

An IGMP/MLD proxy devices can configure the "interface priority" value for each candidate upstream interface. The priority is indicated by an integer value and is part of the configuration. A lower value indicates a lower priority, and the default value of the interface priority is zero.

The interface priority value is reflected when the channel/session ID is not configured as the candidate upstream interface or when multiple candidate upstream interfaces configure the same channel/session ID. In these cases, the candidate upstream interface with the highest priority is selected as the upstream interface. As stated in Section 3.2, if multiple candidate upstream interfaces have the same priority value, they act as upstream interfaces for the configured channel/session ID in parallel and may receive duplicate packets.

4.3. Default Upstream Interface

Operators can configure "a default upstream interface" for all incoming sessions/channels in the IGMP/MLD proxy devices. A default upstream interface is used as the upstream interface when candidate upstream interfaces are not configured for the channel/session ID or interface priority value. A default upstream interface is also used if the proxy device detects configuration errors.

If a default upstream interface is not configured on an IGMP/MLD proxy device, the candidate upstream interface with the highest IPv4/v6 address is selected as the default upstream interface.

5. Upstream Interface Takeover

"Upstream interface takeover" is a function for proxy devices to realize continuous multicast data reception. A proxy device can simultaneously use more than two upstream interfaces per session/channel. If a proxy device detects that one or some of the selected upstream interface(s) is/are going down or inactive, it disables the interface(s) and only uses the active interface(s). Selection of another active upstream interface is done with the highest priority among the candidate upstream interfaces covering the same channel/session ID. The list of candidate upstream interfaces (except the disabled interface) is recursively examined, and a new upstream interface is selected from the list. If another candidate upstream interface is not configured, the default upstream interface is selected.

The condition of whether the upstream adjacent router is active or inactive can be determined by checking the link/interface conditions on the proxy device or by monitoring the IGMP/MLD Query or PIM [RFC7761] Hello message reception on the link. There are cases where PIM is not running on the link or IGMP/MLD Query messages are not always transmitted by the upstream router (e.g., when the explicit tracking function [I-D.ietf-pim-explicit-tracking] is enabled). [I-D.contreras-pim-multiif-config] describes how to detect link/interface conditions.

An active interval is a period in which the selected upstream interface on the proxy device remains active. The active interval of each candidate upstream interface can be configured. Active interval values vary depending on whether the network operators wish to trigger via IGMP/MLD or PIM messages. The default active interval for detecting an inactive upstream interface MAY be approximately twice the IGMP/MLD General Query interval and PIM Hello interval (TODO). However, defining the optimal timer value for switching from an inactive upstream interface to an active upstream interface from a list of candidate upstream interfaces is out of scope of this document. Note that it SHOULD be possible for operators to change the timer value according to the network conditions or other factors.

6. Dynamic Upstream Interface Configuration

6.1. Signaling-based Upstream Interface Configuration

Operators may want proxy devices to dynamically configure upstream interfaces for specific multicast channels/sessions. [I-D.contreras-pim-multiif-config] describes a signaling-based dynamic upstream interface configuration method to support multiple upstream interfaces for IGMP/MLD proxies. The dynamic upstream interface configuration is enabled when network operators set it up on their proxy devices; however, if upstream interface(s) are statically configured, the static configuration is prioritized.

6.2. Controller-based Upstream Interface Configuration

A centralized controller can instruct a proxy device on the upstream interface to use for specific multicast channels.

The controller should configure a default upstream interface for subscription requests that do not match an explicitly configured behavior. In case of an upstream interface failure, the default upstream interface can take over the failed upstream to provide redundancy.

To enable this type of configuration, some control and management interfaces must be supported by a proxy to receive configuration instructions from the controller.

The controller can interact with multiple proxies in the network. As a centralized element, it can make coordinated decisions to manage multicast traffic in the network in a coordinated manner.

6.3. Updating YANG Model

Regarding the IGMP/MLD YANG model proposed in [RFC9166], there is a description of interfaces for IGMP (similarly for MLD). However, it is necessary to update the proposed YANG model to include all information about the upstream interfaces discussed in this study and to consider actions related to the dynamic upstream interface configuration. [I-D.zcl-pim-multiif-igmp-mld-proxy-yang] is a potential data model proposal used for this purpose.

7. Security Considerations

This document neither provides new functions nor modifies the standard functions defined in [RFC3376][RFC3810][RFC5790]; therefore, no additional security considerations are provided for these protocols. Conversely, it is possible to encounter denial-of-service (DoS) attacks to stop upstream interface takeover if attackers illegally send IGMP/MLD Query or PIM Hello messages on a LAN within a shorter period (i.e., before the expiration of the active upstream interface interval). To bypass such threats, it is recommended to capture the source addresses of the IGMP/MLD Query or PIM Hello message senders and examine whether these addresses correspond to the correct adjacent upstream routers. These considerations are TBD.

8. Summary of Aspects Requiring Further Discussion

We have the following open issues.

* Default active interval for detecting an inactive upstream interface (Section 5).

* Interaction with signaling methods (i.e., IGMP/MLD messages) to configure the upstream interface(s) (Section 6).

* Security threats from potential DoS attacks (Section 7).

They will be discussed in the future revisions of this document.

9. IANA Considerations

This document has no IANA actions required.

10. Acknowledgements

TBD.

11. References

11.1. 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>.

[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>.

11.2. Informative References

[I-D.contreras-pim-multiif-config]
Contreras, L. M. and H. Asaeda, "Signaling-based configuration for supporting multiple upstream interfaces in IGMP/MLD proxies", Work in Progress, Internet-Draft, draft-contreras-pim-multiif-config-03, 7 July 2025, <https://datatracker.ietf.org/doc/html/draft-contreras-pim-multiif-config-03>.

[I-D.ietf-pim-explicit-tracking]
Asaeda, H., "IGMP/MLD-Based Explicit Membership Tracking Function for Multicast Routers", Work in Progress, Internet-Draft, draft-ietf-pim-explicit-tracking-13, 1 November 2015, <https://datatracker.ietf.org/doc/html/draft-ietf-pim-explicit-tracking-13>.

[I-D.zcl-pim-multiif-igmp-mld-proxy-yang]
Zhao, H., Contreras, L. M., Liu X., and H. Asaeda, "Yang Data Model for supporting multipath IGMP/MLD proxies", Work in Progress, Internet-Draft, draft-zcl-pim-multiif-igmp-mld-proxy-yang-01, 7 July 2025, <https://datatracker.ietf.org/doc/html/draft-zcl-pim-multiif-igmp-mld-proxy-yang-01>.

[ICIN.xml] Fernandez, D., Contreras, L. M., Moyano, R. F., Garcia, S., and IEEE, "NFV/SDN Based Multiple Upstream Interfaces Multicast Proxy Service", 2021 24th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN), pp. 159-163, DOI 10.1109/icin51074.2021.9385529, 1 March 2021, <https://doi.org/10.1109/icin51074.2021.9385529>.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <https://www.rfc-editor.org/info/rfc3376>.

[RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <https://www.rfc-editor.org/info/rfc3810>.

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, DOI 10.17487/RFC4605, August 2006, <https://www.rfc-editor.org/info/rfc4605>.

[RFC5213] Gundavelli, S., Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, DOI 10.17487/RFC5213, August 2008, <https://www.rfc-editor.org/info/rfc5213>.

[RFC5790] Liu, H., Cao, W., and H. Asaeda, "Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols", RFC 5790, DOI 10.17487/RFC5790, February 2010, <https://www.rfc-editor.org/info/rfc5790>.

[RFC6224] Schmidt, T., Waehlisch, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, DOI 10.17487/RFC6224, April 2011, <https://www.rfc-editor.org/info/rfc6224>.

[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>.

[RFC9166] Zhao, H., Liu, X., Liu, Y., Peter, A., and M. Sivakumar, "A YANG Data Model for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping", RFC 9166, DOI 10.17487/RFC9166, February 2022, <https://www.rfc-editor.org/info/rfc9166>.

Appendix A. Proof of Concept

The support of multiple upstream interfaces for IGMP/MLD proxies was experimentally implemented following a controller-based configuration approach. The implementation was based on Linux using a software-defined networking application running over a Ryu controller. This application uses OpenFlow from the controller to control an Open vSwitch, which relays downstream multicast data flows and upstream IGMP/MLD control traffic. The proof of concept is comprehensively described in [ICIN.xml].

Authors’ Addresses

Hitoshi Asaeda
National Institute of Information and Communications Technology
4-2-1 Nukui-Kitamachi, Koganei,
Tokyo 184-8795
Japan
Email: asaeda@nict.go.jp

Luis M. Contreras
Telefonica
Email: luismiguel.contrerasmurillo@telefonica.com