

PIM
Internet-Draft
Intended status: Informational
Expires: 25 April 2024

H. Asaeda
NICT
L. Contreras
Telefonica
23 October 2023

**Multiple Upstream Interface Support for IGMP/MLD Proxy
draft-asaeda-pim-multiif-igmpmldproxy-06**

Abstract

This document describes the way of supporting multiple upstream interfaces for an IGMP/MLD proxy device. The proposed extension enables an IGMP/MLD proxy device to receive multicast sessions/channels through the different upstream interfaces. The upstream interface can be selected based on multiple criteria, such as the subscriber address prefixes, channel/session IDs, and interface priority values. A mechanism for upstream interface takeover that enables to switch from an inactive upstream interface to an active upstream interface is also described.

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 25 April 2024.

Copyright Notice

Copyright (c) 2023 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 Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	4
3.	Upstream Selection Mechanism	5
3.1.	Channel-Based Upstream Selection	5
3.2.	Subscriber-Based Upstream Selection	5
3.3.	Multiple Upstream Interface Selection for Robust Data Reception	6
4.	Candidate Upstream Interface Configuration	6
4.1.	Address Prefix Record	6
4.2.	Channel/Session ID	7
4.3.	Interface Priority	8
4.4.	Default Upstream Interface	8
5.	Upstream Interface Takeover	8
5.1.	Proxy Behavior	9
5.2.	Active Interval	10
6.	Automatic Upstream Interface Configuration	10
6.1.	Signaling-based Upstream Interface Configuration	10
6.2.	Controller-based Upstream Interface Configuration	10
7.	Security Considerations	11
8.	Consideration for Updating YANG Model	11
9.	Summary of aspects requiring further dicussion	12
10.	IANA Considerations	12
11.	Acknowledgements	12
12.	Informative References	12
Appendix A.	Summary of requirements	14
Appendix B.	Proof of concept	15
	Authors' Addresses	15

[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 joining or leaving of multicast sessions. A proxy device performing IGMP/MLD-based forwarding (as known as IGMP/MLD proxy) [[RFC4605](#)] maintains multicast membership information by IGMP/MLD protocols on the downstream interfaces and sends IGMP/MLD membership report messages via the upstream interface to the upstream multicast routers when the membership information changes (e.g., by receiving solicited/ unsolicited report messages). The proxy device forwards appropriate multicast packets received on its upstream interface to

each downstream interface based on the downstream interface's subscriptions.

According to the specification of [[RFC4605](#)], an IGMP/MLD proxy has `_a_single_` upstream interface and one or more downstream interfaces. The multicast forwarding tree must be manually configured by designating upstream and downstream interfaces on an IGMP/MLD proxy device, and the root of the tree is expected to be connected to a wider multicast infrastructure. An IGMP/MLD proxy device hence performs the router portion of the IGMP or MLD protocol on its downstream interfaces, and the host portion of IGMP/MLD on its upstream interface. The proxy device must not perform the router portion of IGMP/MLD on its upstream interface.

On the other hand, there is a scenario in which an IGMP/MLD proxy device enables multiple upstream interfaces and receives multicast packets through these interfaces. For example, a proxy device having more than one interface may want to access to different networks, such as a global network like the Internet and local-scope networks. Or, a proxy device having 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 shall receive multicast packets from the different upstream interfaces and forward to the downstream interface(s). Several other scenarios and subsequent requirements for the support of multiple upstream interfaces on IGMP/MLD proxy are documented in [[I-D.ietf-pim-multiple-upstreams-reqs](#)].

This document describes the mechanism that enables an IGMP/MLD proxy device to receive multicast sessions/channels through the different upstream interfaces. The mechanism is configured with either "channel-based upstream selection" or "subscriber-based upstream selection", or both of them. By channel-based upstream selection, an IGMP/MLD proxy device selects one or multiple upstream interface(s) from the candidate upstream interfaces "per channel/session". By subscriber-based upstream selection, an IGMP/MLD proxy device selects one or multiple upstream interface(s) from the candidate upstream interfaces "per subscriber/receiver".

When a proxy device transmits an IGMP/MLD report message, it examines the source and multicast addresses in the IGMP/MLD records of the report message. It then transmits the appropriate IGMP/MLD report message(s) from the selected upstream interface(s). Based on this, when a proxy device selects "one" upstream interface from the candidate upstream interfaces per session/channel, it is possible to enable "load balancing" per session/channel. Moreover, when a proxy device selects "more than two" upstream interfaces from the candidate upstream interfaces per session/channel, it potentially receives

duplicate (redundant) packets for the session/channel from the different upstream interfaces simultaneously and provides "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). The potential timer values to switch from an inactive upstream interface to an active upstream interface from a list of candidate upstream interfaces are discussed in this document as well. An "automatic upstream configuration" mechanism that selects an appropriate upstream interface(s) for sessions/channels based on the network and adjacent routers' conditions is also described in this document.

2. Terminology

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](#)].

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

- * Selected upstream interface (or simply, upstream interface): A proxy device's interface in the direction of the root of the multicast forwarding tree. A proxy device performs the host portion of IGMP/MLD on its upstream interfaces. An upstream interface is selected from a list of candidate upstream interfaces.
- * Default upstream interface: A default upstream interface is the upstream interface for multicast sessions/channels for which a proxy device cannot choose other interfaces as the upstream interface. A default upstream interface is configured.
- * Active upstream interface: An active upstream interface is the upstream interface that has been receiving packets for specific multicast sessions/channels during the pre-defined active interval.
- * Inactive upstream interface: An inactive upstream interface is the interface that has not received packets for specific multicast sessions/channels during the pre-defined active interval.

- * Downstream interface: Each of a proxy device's interfaces 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 subscriber address prefixes, channel/session IDs, and priority values on an IGMP/MLD proxy device.
- * Channel/session ID: Channel/session ID consists of source address prefix and multicast address prefix for which a candidate upstream interface supposes to be an upstream interface for specified multicast sessions/channels. Both or either source address prefix and/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 the candidate upstream interfaces "per channel/session" based on the "channel/session ID" configuration. This mechanism is called "channel-based upstream selection" whose configuration is explained in [Section 4.1](#) and [Section 4.2](#)). enables an IGMP/MLD proxy device to use one or multiple upstream interface(s) from the candidate upstream interfaces "per channel/session" based on the "channel/session ID" configuration (as will be in [Section 4.1](#) and [Section 4.2](#)).

3.2. Subscriber-Based Upstream Selection

An IGMP/MLD proxy device selects one or multiple upstream interface(s) from the candidate upstream interfaces "per subscriber/receiver". This is called "subscriber-based upstream selection". It enables a proxy device to use one or multiple upstream interface(s) per session/channel from the "candidate upstream interfaces" based on the "subscriber address prefix" configuration (as will be in [Section 4.1](#)).

3.3. 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" (as will be in [Section 4.3](#)) are identical, these candidate upstream interfaces act as the upstream interfaces for the sessions/channels and receive the packets simultaneously. This multiple upstream interface selection implements duplicate packet reception from redundant paths. It may improve data reception quality or robustness for a session/channel, as the same multicast data packets can come from different upstream interfaces simultaneously. However, this robust data reception does not guarantee that the packets come from disjoint paths. It only configures that the adjacent upstream routers are different.

4. Candidate Upstream Interface Configuration

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

4.1. Address Prefix Record

An IGMP/MLD proxy device can configure the "subscriber address prefix" and "channel/session ID" for each candidate upstream interface.

Channel/session ID consists of "source address prefix" and "multicast address prefix". Subscriber address prefix and source address prefix MUST be a valid unicast address prefix, and multicast address prefix MUST be a valid multicast address prefix. A proxy selects an upstream interface from its candidate upstream interfaces based on the configuration of the following address prefix record:

(subscriber address prefix, (channel/session ID))

where channel/session ID includes:

(source address prefix, multicast address prefix)

The default values of these address prefixes are "null". Null source address prefix represents the wildcard source address prefix, which indicates any host. Null multicast address prefix represents the 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 candidate upstream interface having the configuration of subscriber address prefix is prioritized. If network operators want to assign a specific upstream interface for specific subscribers without depending on source and multicast address prefixes, both source and multicast addresses in the address prefix record is configured "null".

If network operators want to select specific upstream interface(s) without depending on subscriber address prefix, subscriber address prefix in the address prefix record is configured "null".

[4.2.](#) Channel/Session ID

Channel/session ID configuration consists of source and multicast address prefixes. Both/either source and/or multicast address may be configured "null". A candidate upstream interface having non-null source and multicast address configuration is prioritized for the upstream interface selection. For example, if a proxy device has two candidate upstream interfaces for the same multicast address prefix and one of them has non-null source address configuration, then that candidate upstream interface is selected for the source and multicast address pair. The other candidate upstream interface is selected for the configured multicast address prefix except the source address configured by the prior interface.

Source address prefix configuration takes priority over multicast address prefix configuration. For example, consider the case that an IGMP/MLD proxy device has a configuration with source address prefix S_p for the candidate upstream interface A and multicast address prefix G_p for the candidate upstream interface B. When it deals with an IGMP/MLD record whose source address, let's say S , is in the range of S_p , and whose multicast address, let's say G , is in the range of G_p , the proxy device selects the candidate upstream interface A, which supports the source address prefix, as the upstream interface, and transmits the (S,G) record via the interface A.

The same address prefix may 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 based on each interface priority value (as will be in [Section 4.3](#)).

4.3. Interface Priority

An IGMP/MLD proxy device can configure the "interface priority" value for each candidate upstream interface. It is an integer value and is part of the configuration. The default value of the interface priority is the lowest value.

The interface priority value effects only when either of the following conditions is satisfied.

- * None of the candidate upstream interfaces configures both the subscriber address prefix and the channel/session ID.
- * More than one candidate upstream interface configures the same channel/session IDs but does not configure the subscriber address prefix.

In these conditions, the candidate upstream interface with the highest priority is chosen as the upstream interface. And as stated in [Section 3.3](#), if the priority values for candidate upstream interfaces are also identical, all of these interfaces act as the upstream interfaces for the configured channel/session ID and may receive duplicate packets.

4.4. Default Upstream Interface

An IGMP/MLD proxy device SHOULD configure "a default upstream interface" for all incoming sessions/channels. A default upstream interface is selected as the upstream interface, when none of the candidate upstream interfaces configures subscriber address prefix, channel/session ID, or interface priority value, or with either of the following conditions.

- * None of the candidate upstream interfaces configures both the subscriber address prefix, the channel/session ID, and identical interface priority value.
- * More than one candidate upstream interface configures the same channel/session IDs and identical interface priority value, but does not configure the subscriber address prefix.

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

5. Upstream Interface Takeover

5.1. Proxy Behavior

If a selected upstream interface is going down or inactive, or an adjacent upstream router is not working, the upstream interface can be disabled and the other active upstream interface listed in the candidate upstream interfaces covering the same channel/session ID can act as a new upstream interface. It recursively examines the list of the candidate upstream interfaces (except the disabled interface) and decides a new upstream interface from them. If no active candidate upstream interfaces exist, the default upstream interface takes its role.

This function called "upstream interface takeover" is a default function for a proxy device that enables multiple upstream interface support. If a proxy device simultaneously uses more than two upstream interfaces per session/channel, and one or some of these upstream interface(s) is/are inactive, the proxy device acts either of the following behaviors based on the configuration; (1) it only uses the active upstream interface(s) and does not add (i.e., does not complement) other upstream interfaces, (2) it uses the active upstream interface(s) and another candidate upstream interface whose priority is highest among the configured upstream interfaces, or (3) it uses the active upstream interface(s) and the default upstream interface.

The condition whether the upstream adjacent router is active or not can be decided by checking the link/interface condition on the proxy device or detected by monitoring IGMP/MLD Query or PIM [[RFC7761](#)] Hello message reception on the link. There are the cases that PIM is not running on the link or IGMP/MLD Query messages are not always transmitted by the upstream router (e.g., because of enabling the explicit tracking function [[I-D.ietf-pim-explicit-tracking](#)]). Therefore, network operators MUST configure either; (1) the proxy device disables the upstream interface takeover, (2) the proxy device triggers upstream interface takeover by detecting no IGMP/MLD Query message within the active interval, or (3) the proxy device triggers upstream interface takeover by detecting no PIM Hello message within the active interval, for each candidate upstream interface.

Network operators may want to keep out of use for the inactive upstream interface(s). This causes, for example, when subscriber-based upstream selection is configured, according to their accounting policy (because the specific subscribers are planned to use the specific upstream interface and cannot receive packets from other upstream interfaces.) In that case, this upstream interface takeover must be disabled, and the proxy device keeps using that interface as the upstream interface for them (and waits for working the interface later again).

[5.2.](#) Active Interval

Active interval is a period, after which a proxy device recognizes that the selected upstream interface is inactive. Active interval for each candidate upstream interface SHOULD be configured. The active interval values are different in the situation whether the network operators want to trigger by either IGMP/MLD or PIM messages. The default active interval to detect an inactive upstream interface is around twice of IGMP/MLD General Query interval and PIM Hello interval. Further discussion are TBD.

[6.](#) Automatic Upstream Interface Configuration

[6.1.](#) Signaling-based Upstream Interface Configuration

There may be the ways for a proxy device to automatically configure the upstream interface for specific multicast channels/sessions. It works for the case in which there are no static configurations for a candidate upstream interface or operators decide.

[[I-D.contreras-pim-multiif-config](#)] describes a signaling-based configuration method for supporting multiple upstream interfaces in IGMP/MLD proxies.

[6.2.](#) Controller-based Upstream Interface Configuration

A centralized controller can instruct the proxy what upstream interface is the appropriate one to use based on a certain multicast channel or on the user herself.

There are options for selecting the most appropriate upstream interface:

- * Association of membership requests from a specific user, identified by the source IP of the IGMP/MLD message, to a specific upstream interface, meaning that all the multicast traffic for that end user is received from a certain upstream interface.
- * Association of (S,G) to a specific upstream interface, meaning that an end user request for specific content delivered from a specific source should be received from a certain upstream interface.
- * Association of (*,G) to a specific upstream interface, meaning that a user request of given content, independently of the source of that content, should be received from a certain upstream interface.

- * Association of (S,*) to a specific upstream interface, meaning that all the requests from a certain user, independently of the group identifying the content, should be received from a certain upstream interface.

The list above does not show any precedence. Thus precedence should be defined to indicate priority in the selection of the criteria to apply, as a list of ordered actions.

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

To enable this manner of configuration, some control and management interface has to be supported by the proxy in order to receive configuration instructions from the controller.

The controller could interact with a number of proxies in the network. Being a centralized element, it could take coordinated decisions for managing all the multicast traffic in the network in a coordinated manner.

7. Security Considerations

This document neither provides new functions nor modifies the standard functions defined in [[RFC3376](#)][[RFC3810](#)][[RFC5790](#)]; hence there is no additional security consideration provided for these protocols themselves. On the other hand, it may be possible to encounter DoS attacks to make the function for upstream interface takeover stop if attackers illegally sends IGMP/MLD Query or PIM Hello messages on a LAN within a shorter period (i.e., before expiring the active interval for the upstream interface). To bypass such threats, it is recommended to capture the source addresses of IGMP/MLD Query or PIM Hello message senders and check whether the addresses correspond to the correct adjacent upstream routers. These considerations are TBD.

8. Consideration for Updating YANG Model

About the IGMP/MLD YANG model proposed in [[RFC9166](#)], there is a description of interfaces for IGMP (similar for MLD). Once this document becomes officially approved, it will be necessary to update the proposed YANG model to include all the information related to the upstream interfaces defined in this document, and consider the actions related to the selection of the upstream interfaces as mentioned in [Section 6](#).

9. Summary of aspects requiring further discussion

The following is a list of aspects which require further discussion:

- * Value of the default active interval to detect an inactive upstream interface ([section 5.2](#)).
- * Signaling methods (i.e. IGMP/MLD messages) for configuring the upstream interface(s) of interest ([section 6.1](#)).
- * Security threats from potential DoS attacks ([section 7](#)).

These aspects will be analysed in future versions of the document.

10. IANA Considerations

This document has no actions for IANA.

11. Acknowledgements

To be completed.

12. 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-01](#), 23 October 2023, <<https://datatracker.ietf.org/api/v1/doc/document/draft-contreras-pim-multiif-config/>>.

[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.ietf-pim-multiple-upstreams-reqs]

Contreras, L. M., Bernardos, C. J., Asaeda, H., and N. Leymann, "Requirements for the extension of the IGMP/MLD proxy functionality to support multiple upstream interfaces", Work in Progress, Internet-Draft, [draft-ietf-pim-multiple-upstreams-reqs-08](#), 9 November 2018, <<https://datatracker.ietf.org/doc/html/draft-ietf-pim-multiple-upstreams-reqs-08>>.

- [ICIN] "D. Fernández, L. M. Contreras, R. F. Moyano and S. García, "NFV/SDN Based Multiple Upstream Interfaces Multicast Proxy Service," 2021 24th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN), Paris, France, 2021, pp. 159-163.", 2021.
- [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>>.
- [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>>.
- [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>>.
- [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. Summary of requirements

The following Table, from [[I-D.ietf-pim-multiple-upstreams-reqs](#)], summarizes the requirements for the extension of the IGMP/MLD proxy functionality to support multiple upstream interfaces. The solution described in this document is based on those requirements.

Functionality	Multicast Wholesale	Multicast Resiliency	Load Balancing	Network Merging	Network Migration
Upstream Control/Delivery	X	X	X	X	X
Downstr. Control/Delivery	X	X	X	X	X
Active / Standby Upstream		X			
Upstr i/f selection per group			X	X	
Upstr i/f selection all group		X			X
ASM	X	X	X	X	X
SSM	X	X	X		X

Figure 1: Summary of requirements

For a more detailed description on the need for the listed requirements, the reader is referred to [[I-D.ietf-pim-multiple-upstreams-reqs](#)].

[Appendix B](#). Proof of concept

The support of multiple upstream interfaces for IGMP/MLD Proxy has been experimentally implemented following the controller-based configuration approach. The implementation has been based on Linux using an SDN application running over Ryu controller. Such application is able to control by means of OpenFlow from the controller an Open vSwitch which is in charge of relaying the downstream multicast data flows and the upstream IGMP/MLD control traffic. The proof of concept is fully described in [[ICIN](#)].

Authors' Addresses

Hitoshi Asaeda
NICT
Email: asaeda@nict.go.jp

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

