

PIM Working Group
Internet-Draft
Expires: August 29, 2013

H. Asaeda
NICT
S. Jeon
Institute de Telecomunicacoes
February 25, 2013

Multiple Upstream Interface Support for IGMP/MLD Proxy
draft-asaeda-pim-mldproxy-multif-01

Abstract

This document describes the way of supporting multiple upstream interfaces for an IGMP/MLD proxy device. The proposed extension enables that an IGMP/MLD proxy device receives multicast packets through multiple upstream interfaces. The upstream interface is selected with manually configured supported address prefixes and interface priority value. A take-over operation switching from an inactive upstream interface to an active upstream interface is also considered.

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 August 29, 2013.

Copyright Notice

Copyright (c) 2013 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	3
2. Terminology	4
3. Per-Channel Load Balancing	4
4. Candidate Upstream Interface Configuration	5
4.1. Supported Address Prefix	5
4.2. Interface Priority	7
4.3. Default Interface	7
5. IANA Considerations	8
6. Security Considerations	8
7. Normative References	8
Authors' Addresses	8

1. Introduction

The Internet Group Management Protocol (IGMP) [1][2] for IPv4 and the Multicast Listener Discovery Protocol (MLD) [3][2] 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) [4] 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 [4], 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 Internet and Intranet. Or, a proxy device having wired link (e.g., ethernet) and high-speed wireless link (e.g., WiMAX or LTE) 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).

This document adds the way to manually configure candidate upstream interfaces for an IGMP/MLD proxy device and select "one" single upstream interface from candidate upstream interfaces per session/channel. When the selected upstream interface is down or disabled, one of the other candidate upstream interfaces takes over the upstream interface (if configured). This enables "per-channel load balancing".

Note that this document only specifies the way to configure per-channel load balancing; it does not specify any intelligent

mechanism/algorithm (e.g., based on link or network condition/usage) or threshold value to select an upstream interface from candidate upstream interfaces to improve data reception quality. Also, an IGMP/MLD proxy device does not select multiple upstream interfaces for the same channels/sessions simultaneously; enabling redundant paths to receive duplicate packets via multiple upstream interfaces to improve data reception quality or robustness for a session/channel is out of scope of 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 [5].

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

Upstream interface (or selected upstream interface):

A proxy device's interface in the direction of the root of the multicast forwarding tree. An upstream interface is selected by either manual or automatic configuration.

Downstream interface:

Each of a proxy device's interfaces that is not in the direction of the root of the multicast forwarding tree.

Candidate upstream interface:

An interface that potentially becomes an upstream interface of the proxy device. Candidate upstream interfaces are manually set up on an IGMP/MLD proxy.

Supported address prefix:

The supported address prefix is the address prefix for which a candidate upstream interface supposes to be an upstream interface. The supported source address prefix and the supported multicast address prefix an IGMP/MLD proxy device can configure. The supported address prefix in this document means both source and multicast address prefixes, unless otherwise specified.

3. Per-Channel Load Balancing

An IGMP/MLD proxy device enables "per-channel load balancing" using multiple upstream interfaces to receive different multicast sessions/channel through the different upstream interfaces. Per-channel load balancing makes an IGMP/MLD proxy device select "one" single upstream interface from candidate upstream interfaces per session/channel,

based on the configurations, which will be described in Section 4.

If an IGMP proxy recognizes that an adjacent upstream router is not working, the selected upstream interface attached to that router can be taken over with the different candidate upstream interface. Or, if the selected upstream interface is going down, the proxy would switch from the inactive interface to the other active upstream interface. This "take-over operation" recursively examines the configurations of the candidate upstream interfaces (except the disabled interface) and decides a new upstream interface from them.

Whether the upstream router is active or not would be decided by checking a link condition or IGMP/MLD query message transmission. However, this document does not describe how an IGMP/MLD proxy can detect the upstream router's condition and when it takes that interface over the different candidate upstream interface.

The take-over operation is enabled by default. When it is disabled (by operation), even if no data comes from the selected upstream interface, the IGMP/MLD proxy device keeps using that interface as the upstream interface for the corresponding sessions/channels.

Per-channel load balancing does not implement duplicate packet reception from redundant paths using multiple upstream interfaces to improve data reception quality or robustness for a session/channel; therefore IGMP/MLD report messages containing the same IGMP/MLD records are not transmitted from different upstream interfaces simultaneously.

4. Candidate Upstream Interface Configuration

Candidate upstream interfaces are the interfaces from which an IGMP/MLD proxy device selects as an upstream interface. They are manually enabled. The upstream interface selection is done based on "supported address prefix" and "interface priority" value.

4.1. Supported Address Prefix

An IGMP/MLD proxy device MAY configure the "supported address prefix" for each candidate upstream interface. A proxy selects an upstream interface from its candidate upstream interfaces based on the configured supported address prefix. The supported address prefix is manually configured. The supported address prefix consists of the following information:

(source address prefix, multicast address prefix)

When the proxy device transmits an IGMP/MLD report message, it examines the source and multicast addresses in the IGMP/MLD records of the report message and transmits the appropriate IGMP/MLD report message(s) from the selected upstream interface(s) that are configured with the range of the supported source and multicast address prefixes.

The default values of both source and multicast address prefixes are a wildcard. If no address prefix value is configured on a candidate upstream interface, the default value is implicitly set up for the candidate upstream interface. The wildcard multicast address prefix is represented by the entire multicast address range (i.e., '224.0.0.0/4' for IPv4 or 'ff00::/8' for IPv6). The wildcard source address prefix is represented by any host. If the default value is set up on a candidate upstream interface, the decision whether the candidate upstream interface is selected as the upstream interface or not is made by the "interface priority" value described in Section 4.2.

The same address prefix may be configured on different candidate upstream interfaces. As well as the above-mentioned default configuration, 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 described in Section 4.2.

For upstream interface selection, source address prefix takes priority over multicast address prefix. This avoids conflict of upstream interface selection. 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.

Obviously, an IGMP/MLD proxy selects a candidate upstream interface having supported source and multicast address prefixes that include both source and multicast address, rather than the other one whose supported source and multicast address prefixes includes either source or multicast address.

4.2. Interface Priority

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

The interface priority value effects only when the following conditions are satisfied.

- o None of the candidate upstream interfaces configure the supported address prefix.
- o Both source and multicast addresses are included in the supported address prefixes configured by more than one candidate upstream interface.
- o Neither source nor multicast address is included in the supported address prefixes configured by any of the candidate upstream interfaces.
- o The supported source address prefix is not configured or does not include the source address, but (on the other hand) the multicast address is included in the supported multicast address prefix configured by more than one candidate upstream interface.

In these conditions, the candidate upstream interface with the highest priority is chosen as the upstream interface.

4.3. Default Interface

In the following conditions, the candidate upstream interface whose IPv4/v6 address is lowest is selected as the upstream interface for that session/channel.

- o None of the candidate upstream interfaces configure the supported address prefix and interface priority value.
- o Both source and multicast addresses are included in the supported address prefixes configured by more than one candidate upstream interfaces, and these candidate upstream interfaces' priorities are identical.
- o Neither source nor multicast address is included in the supported address prefixes configured by any of the candidate upstream interfaces, and all candidate upstream interfaces' priorities are identical.

- o The supported source address prefix is not configured or does not include the source address, and the multicast address is included in the supported multicast address prefix configured by more than one candidate upstream interface, yet these candidate upstream interfaces' priorities are identical.

5. IANA Considerations

This document has no actions for IANA.

6. Security Considerations

This document neither provides new functions nor modifies the standard functions defined in [1][3][2]. Therefore there is no additional security consideration provided for these protocols.

7. Normative References

- [1] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [2] 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, February 2010.
- [3] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [4] 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, August 2006.
- [5] Bradner, S., "Key words for use in RFCs to indicate requirement levels", RFC 2119, March 1997.

Authors' Addresses

Hitoshi Asaeda

National Institute of Information and Communications Technology (NICT)
Network Architecture Laboratory
4-2-1 Nukui-Kitamachi
Koganei, Tokyo 184-8795
Japan

Email: asaeda@nict.go.jp

Seil Jeon

Institute de Telecomunicacoes
Campus Universitario de Santiago
Aveiro 3810-193
Portugal

Email: seiljeon@av.it.pt

