

PIM Working Group
Internet Draft
Intended status: Standards Track
Expires: September 10, 2013

Hong-Ke Zhang
Shuai Gao
Beijing Jiaotong University
T C.Schmidt
HAW Hamburg
Bo-hao Feng
Li-Li Wang
Beijing Jiaotong University
March 11, 2013

Multi-Upstream Interfaces IGMP/MLD Proxy
draft-zhang-pim-muiimp-00.txt

Abstract

In this document, followed by the idea mentioned in [4] and subsequent update in [5], an IGMP/MLD proxy with multiple upstream interfaces called MUIIMP is proposed and analyzed. The MUIIMP inherits the basic rule of the IGMP/MLD proxy but extends with multiple upstream interfaces. To avoid data redundancy, each upstream interface of an MUIIMP device MUST NOT send or subscribe the same data simultaneously.

Requirements Language

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 [RFC2119].

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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".

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on September, 2013.

Copyright Notice

Copyright (c) 2012 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.....	3
3. MUIIMP Behavior.....	4
3.1. The selection of default upstream interface.....	5
3.2. Report of downstream subscriptions to upstream interfaces..	5
3.3. Handover of the upstream interface.....	6
4. Security Considerations.....	6
5. References.....	6
Acknowledgment.....	8

1. Introduction

RFC 4605 [1] specifies an IGMP/MLD proxy mechanism for forwarding based solely upon IGMP/MLD membership information in scenarios where multicast routing is not available. According to [1], an IGMP/MLD Proxy performs the router portion of the IGMP/MLD protocol on its downstream interfaces, and the host portion of the IGMP/MLD protocol on its single upstream interface.

The IGMP/MLD proxy mechanism can effectively extend the multicast scope and greatly simplify the implementation of edge devices. However, the IGMP/MLD proxy may exhibit inefficiency in some specific scenarios due to the limitation of single upstream interface. For example, in PMIPv6 multicast environment, multiple IGMP/MLD proxy instances need to be deployed at the MAG in [6], which may result in tunnel convergence problem. In addition, there are also requirements to extend the IGMP/MLD proxy to support multiple upstream interfaces as the emergence of multi-homing.

One thing to note is the idea about multiple upstream interfaces for IGMP/MLD proxy was firstly proposed in the draft [4] to improve the performance of mobile multicast source. The Multimob working group draft [5] includes the related latest descriptions. Considering the multiple upstream interfaces extension is not only required for mobile multicast sources scenarios, this document is presented here.

In this document, an IGMP/MLD proxy with multiple upstream interfaces called MUIIMP is proposed and described. The MUIIMP inherits the basic rule of the IGMP/MLD proxy but extends with multiple upstream interfaces. To avoid data redundancy, each upstream interfaces of an MUIIMP device MUST NOT send or subscribe the same data simultaneously. The MUIIMP is designed to support local multicast listeners and senders.

2. Terminology

Upstream Interface: A proxy device's interface in the direction of the root of the tree.

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

Default upstream interface: An upstream interface which is by default associated with each downstream node subscribing or sending specific channel (group address prefix) or special multicast state.

3. MUIIMP Behavior

The MUIIMP inherits the basic rule of the IGMP/MLD proxy but extends with multiple upstream interfaces. A MUIIMP device has one or more upstream as well as downstream interfaces, which may be any type interfaces, including physical or logical interfaces.

The MUIIMP performs the router portion of the IGMP/MLD protocol on its downstream interfaces, and the host portion of IGMP/MLD on its upstream interfaces. The MUIIMP device MUST NOT perform the router portion of IGMP/MLD on its upstream interfaces.

The MUIIMP device maintains a database for multicast listeners consisting of the merger of all subscriptions on any downstream interface. In order to avoid the redundant multicast traffic, the proxy device should initiate unique traffic subscriptions. Besides, a policy list that records the default upstream interface for the downstream nodes is held for the selection of upstream interface.

In the following, the MUIIMP device behavior will be discussed according the role of the downstream nodes.

1) Multicast listener on the downstream interface

Multicast listener reports are group-wise aggregated by the MLD proxy. The aggregated report is issued to the upstream interface based on the subscriptions as well as the policy list. When receiving the IGMP/MLD subscriptions on the downstream interface, the MUIIMP checks the membership database to make a decision whether sends IGMP/MLD membership reports on the corresponding default upstream interface or not. Refer to Section 3.2 for the details about membership subscriptions lookup and report decisions.

When receiving packets on its upstream interfaces, the MUIIMP forwards the traffic to all the downstream interfaces based upon the downstream interfaces' subscriptions.

2) Multicast source on the downstream interface

When receiving packets on its downstream interface, the MUIIMP forwards the traffic to the corresponding default upstream interface, as well as all the downstream interfaces other than the incoming interface based upon the downstream interfaces' subscriptions.

The (first) multicast router(s) operating multicast routing protocol like PIM-SM[7] connected to the outside multicast domain should be configured to treat the multicast source inside the MUIIMP domain

being directly connected. Otherwise, it will discard the data due to the failure of the direct connection check.

3.1. The selection of default upstream interface

Typically, the choice of the default upstream interface is based on the policy list which is maintained at the MUIIMP.

The expression of the policy list is like below:

(node prefix, multicast group address/multicast state, upstream interface)

Here node prefix represents the address prefix of the node on the downstream interface that may be a multicast listener or multicast source. And the multicast group address indicates the channel that the multicast listener is subscribing or the multicast source is publishing while the multicast state is only valid for listeners indicating the state about both multicast source and multicast group they are subscribing.

In other word, in the MUIIMP, the multicast group address/multicast state and the node prefix will act as rules to select the default upstream interface. Alternate configurations (e.g., the MAG-LMA tunnel interface in PMIPv6 environment) MAY be applied.

3.2. Report of downstream subscriptions to upstream interfaces

To avoid the redundant multicast traffic, the proxy device MUST NOT send the same multicast subscription record on different upstream interfaces simultaneously. In detail, we recommend the following rules when receiving an IGMP/MLD subscription on the downstream interface.

- 1) If the received IGMP/MLD subscription is new and has not been subscribed by other downstream multicast listeners, the proxy device SHOULD initiate the IGMP/MLD subscription on the corresponding default upstream interface.
- 2) If there exists the same IGMP/MLD subscription which has already been subscribed by other downstream multicast listener, the proxy device SHOULD not initiate extra IGMP/MLD subscription.
- 3) If there exists IGMP/MLD subscriptions which have already included the received IGMP/MLD subscription, the proxy device SHOULD not initiate extra IGMP/MLD subscription.

- 4) If there exists overlapping subsets between the received IGMP/MLD subscription and current IGMP/MLD subscriptions, the proxy device SHOULD initiate the IGMP/MLD subscription on the corresponding default upstream interface excluding the overlapping subsets that have been subscribed before.

All subscriptions sent on the same upstream interface SHOULD be merged according the merging rule in RFC 4605. In addition, the local multicast source should be excluded in the final subscriptions to avoid replicated multicast traffic from outside.

3.3. Handover of the upstream interface

If an upstream interface fails for some reason such as the deletion of the tunnel interface in mobile environment, the handover of the upstream interface is performed. Generally, all the subscriptions sent on the previous invalid upstream interface are transferred to the new valid upstream interfaces which are chosen among the default upstream interfaces of the corresponding downstream nodes. The choice may be made based on the predefined policy (e.g., the interface priority, the number of listeners, the lowest IP address). An alternative may be applied by the MUIIMP device itself according to the traffic monitored or some strategies configured by the operator.

4. Security Considerations

To be done.

5. References

- [1] B. Fenner, H. He, B. Haberman and H. Sandick. "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.
- [2] Cain, B., Deering, S., Kouvelas, I., Fenner, B. and A. Thyagarajan, "Internet Group Management Protocol, Version3", RFC 3376, October 2002.
- [3] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

- [4] Hong-Ke Zhang, Zhi-Wei Yan, Shuai Gao, et al., "Multicast Source Mobility Support in PMIPv6 Network", draft-zhang-multimob-msm-03, July 2011.
- [5] T C. Schmidt, S. Gao, H. Zhang, M. Waehlich, "Mobile Multicast Sender Support in Proxy Mobile IPv6 (PMIPv6) Domains", draft-ietf-multimob-pmipv6-source-02, October 2012.
- [6] T. Schmidt, M. Waehlich and S. Krishnan. "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011.
- [7] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2000.

Authors' Addresses

Hong-Ke Zhang, Shuai-Gao, Bo-Hao Feng, Li-Li Wang
National Engineering Lab for NGI Interconnection Devices
Beijing Jiaotong University, China

Phone: +861051684274
Email: hkzhang@bjtu.edu.cn
shgao@bjtu.edu.cn
11111021@bjtu.edu.cn
liliwang@bjtu.edu.cn

Thomas C. Schmidt
HAW Hamburg
Berliner Tor 7
Hamburg 20099
Germany

Email: schmidt@informatik.haw-hamburg.de
URI: <http://inet.cpt.haw-hamburg.de/members/schmidt>

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.