

PIM Working Group
Internet-Draft
Intended status: Informational
Expires: June 14, 2013

H. Asaeda
NICT
December 11, 2012

**IGMP/MLD-Based Explicit Membership Tracking Function for Multicast
Routers
draft-ietf-pim-explicit-tracking-03**

Abstract

This document describes the IGMP/MLD-based explicit membership tracking function for multicast routers. The explicit tracking function contributes to saving network resources and fast leaves (i.e. shortening leave latency).

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 June 14, 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	4
3.	Explicit Tracking Function	4
3.1.	Membership State Information	4
3.2.	Specific Query Suppression	5
3.3.	Shortening Leave Latency	6
4.	Lowering the Possibility of Outdated Membership State	7
5.	All-Zero and Unspecified Source Addresses	8
6.	Compatibility with Older Version Protocols	8
7.	IANA Considerations	8
8.	Security Considerations	9
9.	Acknowledgements	9
10.	References	9
10.1.	Normative References	9
10.2.	Informative References	9
	Author's Address	9

1. Introduction

The Internet Group Management Protocol (IGMP) [1] for IPv4 and the Multicast Listener Discovery Protocol (MLD) [2] for IPv6 are the standard protocols used by member hosts and multicast routers. When a host starts/finishes listening to particular multicast channels, it sends IGMP/MLD State-Change Report messages specifying the corresponding channel information as the join/leave request to its upstream router (i.e., an adjacent multicast router or IGMP/MLD proxy device [5]). The "unsolicited" report messages are sent only once upon reception/departure.

IGMP/MLD are non-reliable protocols; the unsolicited report messages may be lost or may not reached upstream routers. To resolve the problem, routers need to update membership information by periodically sending IGMP/MLD General Query messages. Member hosts then reply with "solicited" report messages whenever they receive the query messages.

Multicast routers are capable of periodically maintaining the multicast membership state of downstream hosts attached to the same link by acquiring unsolicited report messages and synchronizing the actual membership state within the General Query timer interval (i.e., [Query Interval] value defined in [1][2].) However, this approach does not guarantee that the membership state is always perfectly synchronized. To minimize the possibility of having outdated membership information, routers may shorten the periodic General Query timer interval. Unfortunately, this increases the number of transmitted solicited report messages and induces network congestion. And the greater the amount of network congestion, the greater the potential for IGMP/MLD report messages being lost and the membership state information being outdated in the router.

The IGMPv3 [1] and MLDv2 [2] protocols can provide the ability to keep track of the downstream (adjacent) multicast membership state to multicast routers, yet the specifications are not clearly given. This document describes the "IGMP/MLD-based explicit member tracking function" for multicast routers and details a way for routers to implement the function. By enabling this explicit tracking function, routers can keep track of the downstream multicast membership state. This function enables the following:

- o Reducing the number of transmitted query and report messages
- o Shortening leave latencies
- o Per-host accounting

- o Maintaining multicast channel characteristics (or statistics)

In addition, the processing of IGMP membership or MLD listener messages consumes CPU resources on individual IGMP/MLD querier and report sender devices. The explicit tracking function therefore reduces not only the network load but also the CPU load on these devices.

The explicit tracking function does not guarantee that the membership state will always be perfectly synchronized; the list of tracked member hosts may be outdated in the router because of host departure from the network without sending State-Change Report messages or loss of such messages due to network congestion. Therefore, a router enabling the function ought to send periodic IGMPv3/MLDv2 General Query messages and inquire about solicited IGMPv3/MLDv2 report messages from downstream member hosts to maintain an up-to-date membership state.

The explicit tracking function potentially requires a large amount of memory so that routers keep all membership states. Particularly when a router needs to maintain a large number of member hosts, this resource requirement could have an impact. Operators may decide to disable this function when their routers have insufficient memory resources, despite the benefits mentioned above.

The explicit tracking function does not change message formats used by the standard IGMPv3 [1] and MLDv2 [2], and their lightweight version protocols [3]; nor does it change a multicast data sender's and receiver's behavior.

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](#) [4].

3. Explicit Tracking Function

3.1. Membership State Information

A router enabling the explicit tracking function maintains the "membership state information". When a multicast router receives a Current-State or State-Change Report message, it creates this membership state information or adds or deletes the receiver IP address to or from it. If there are no more receivers maintained, the router may keep the membership state information with an empty

receiver list.

The membership state information consists of the following information:

(S, G, number of receivers, (receiver records))

where each receiver record is of the form:

(IGMP/MLD membership/listener report sender's address)

This state information must work properly when a receiver (i.e., report sender) sends the identical report messages multiple times.

In the state information, each S and G indicates a single IPv4/IPv6 address. S is set to "Null" for Any-Source Multicast (ASM) communication (i.e., (*,G) join reception). In order to simplify the implementation, the explicit tracking function does not keep the state of (S,G) joined with EXCLUDE filter mode. If a router receives an (S,G) join/leave request with EXCLUDE filter mode from the downstream hosts, it translates it to a (*,G) join state/leave request and records the state and the receivers' addresses in the maintained membership state information. Note that this membership state translation does not change the routing protocol behavior. The routing protocol must deal with the original join/leave request and translate the request only for the membership state information.

3.2. Specific Query Suppression

In accordance with [1] and [2], whenever a router receives the State-Change Report, it sends the corresponding Group-Specific or Group-and-Source Specific Query messages to confirm whether or not the report sender is the sole member host. All member hosts joining the identical channel send their own Current-State Report messages after acquiring these query messages. Transmitting a large number of Current-State Report messages consumes network resources, and this may pose a particular problem when many hosts joining the identical channel send these reports simultaneously.

The explicit tracking function can reduce the number of Group-Specific or Group-and-Source Specific Query messages transmitted from a router, and reduce the number of Current-State Report messages transmitted from member hosts. If a router enables the explicit tracking function with "specific query suppression", it suppresses specific query transmission and transmits specific query messages only when the router expects that the State-Change Report sender is the sole member of the channel, based on its membership state information (expressed in [Section 3.1](#)).

As standard behavior for [\[1\]](#) and [\[2\]](#), a router also sends a Group-Specific or Group-and-Source Specific Query multiple times when it receives a State Change Report message (e.g., leave request) from a member host. This is in order to confirm whether or not the host is the sole member. However, if the router enabling the explicit tracking function runs specific query suppression and receives one or more replies for the specific query retransmission from the downstream member(s), the router can cancel resending of the identical specific query message(s).

Note that the default behavior of the router that supports the explicit tracking function SHOULD disable this specific query suppression in order to avoid the risk caused by the situation in which multiple multicast routers exist on a LAN and the querier router is not the forwarder router. When the querier suppresses the specific query message transmission, and expects that the State-Change Report sender is not the sole member of the channel, it does not send the specific query and none of the routers on the same LAN receive a Current-State Report message from the corresponding member hosts. The forwarder in this case may prune the routing path though there are other member hosts subscribing to the channel on the LAN.

[3.3.](#) Shortening Leave Latency

A router enabling the explicit tracking function can shorten leave latencies by tuning several timers and values to what it expects whether or not the State-Change Report sender is the channel's sole member.

The [Last Member Query Interval] (LMQI) and [Last Listener Query Interval] (LLQI) values specify the maximum time allowed before sending a responding Report. The [Last Member Query Count] (LMQC) and [Last Listener Query Count] (LLQC) are the number of Group-Specific Queries or Group-and-Source Specific Queries sent before the router assumes there are no local members. The [Last Member Query Time] (LMQT) and [Last Listener Query Time] (LLQT) values are the total time the router should wait for a report after the Querier has sent the first query.

The default value for LMQI/LLQI defined in the standard specifications [\[1\]](#)[\[2\]](#) is 1 second. For a router enabling the explicit tracking function, the LMQI/LLQI MAY be set to 1 second or shorter. The LMQC/LLQC values MAY be set to 1 for the router, whereas their default values are the [Robustness Variable] value whose default value is 2. Smaller LMQC/LLQC values give smaller LMQT/LLQT, which shortens the leave latencies. On the other hand, setting smaller LMQC/LLQC values poses the risk described in [Section 4](#). Operators setting smaller LMQC/LLQC values must recognize

this tradeoff.

4. Lowering the Possibility of Outdated Membership State

There are possibilities that the membership expectation performed by a router enabling the explicit tracking function will be inconsistent due to an outdated membership state. For example, (1) a router expects that more than one corresponding member host exists on its LAN, but in fact no member host exists for that multicast channel, or (2) a router expects that no corresponding member host exists on its LAN, but in fact more than one member host exists for that multicast channel. These cases are particularly likely for a router that enables specific query suppression (as in [Section 3.2](#)) and configures small LMQC/LLQC for shortening leave latency (as in [Section 3.3](#)).

The first of these cases may occur in an environment where the sole member host departs the network without sending a State-Change Report message. This is because a router enabling specific query suppression does not send a specific query if it believes the report sender is not the sole member host. The router later detects that there is no member host for the corresponding channels when it does not receive a Current-State Report within the timeout of the response for the periodic General Query. However, this situation prolongs leave latency and wastes network resources since the router forwards unneeded traffic until that point.

The second case occurs when a router sends a specific query but does not receive a State-Change Report from a downstream host within an LMQT or LLQT period. It recognizes that no member host exists on the LAN and might prune the routing path. The router reestablishes the routing path when it receives the solicited report message for the channels. However, the downstream hosts may lose the data packets until the routing path is reestablished and the data forwarding is restarted.

In order to reduce the possibility of the incorrect membership expectation and keep the up-to-date membership state information, when a router enabling the explicit tracking function enables specific query suppression, the router **SHOULD** configure the LMQC/LLQC value to 2 (the default value of the [Robustness Variable] value) or higher; or, when a router enabling the explicit tracking function configures a small LMQC/LLQC, the router **SHOULD NOT** enable specific query suppression.

5. All-Zero and Unspecified Source Addresses

The IGMPv3 specification [1] mentions that an IGMPv3 report is usually sent with a valid IP source address, yet it permits a host to use the 0.0.0.0 source address (since the host has not yet acquired an IP address), and routers must accept a report with this source address. The MLDv2 specification [2] mentions that an MLDv2 report must be sent with a valid IPv6 link-local source address, yet an MLDv2 report may be sent with the unspecified address (::) if the sending interface has not acquired a valid link-local address. [2] also mentions that routers silently discard a message that is not sent with a valid link-local address or sent with the unspecified address, without taking any action, because of security considerations.

When a router enabling the explicit tracking function receives IGMP/MLD report messages with an all-zero or unspecified source address, it deals with the IGMP/MLD report messages correctly as defined in [1][2] and continuously keeps track of the membership state, but SHOULD NOT maintain the host specifying all-zero or an unspecified source address in its membership state information.

6. Compatibility with Older Version Protocols

The explicit tracking function does not work with older versions of IGMP or MLD, IGMPv1 [6], IGMPv2 [7], or MLDv1 [8], because a member host using these protocols enables "membership report suppression" by which the host will cancel sending pending membership reports if a similar report is observed from another member on the network.

To preserve compatibility with older versions of IGMP/MLD, routers need to support downstream hosts that are not upgraded to the latest versions and run membership report suppression. Therefore, if a multicast router enabling the explicit tracking function changes its compatibility mode to the older versions, the router SHOULD disable the explicit tracking function. On the other hand, the router MAY NOT flush the maintained membership state information. When the router changes back to IGMPv3 or MLDv2 mode, it resumes the function with the old membership state information even if the state information is outdated. This provides "smooth state transition" that does not initiate the membership state information from scratch and synchronizes the actual membership state smoothly.

7. IANA Considerations

This document has no actions for IANA.

8. Security Considerations

There is no additional security consideration for [\[1\]](#)[\[2\]](#)[\[3\]](#) provided.

9. Acknowledgements

Toerless Eckert, Sergio Figueiredo, Nicolai Leymann, Stig Venaas, and others provided many constructive and insightful comments.

10. References

10.1. 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] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), June 2004.
- [3] 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.
- [4] Bradner, S., "Key words for use in RFCs to indicate requirement levels", [RFC 2119](#), March 1997.

10.2. Informative References

- [5] 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.
- [6] Deering, S., "Host Extensions for IP Multicasting", [RFC 1112](#), August 1989.
- [7] Fenner, W., "Internet Group Management Protocol, Version 2", [RFC 2373](#), July 1997.
- [8] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", [RFC 2710](#), October 1999.

Author's Address

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