PIM Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 4, 2015

H. Asaeda
NICT

August 31, 2014

IGMP/MLD-Based Explicit Membership Tracking Function for Multicast Routers
draft-ietf-pim-explicit-tracking-10

Abstract

This document describes the IGMP/MLD-based explicit membership tracking function for multicast routers and IGMP/MLD proxy devices supporting IGMPv3/MLDv2. The explicit membership tracking function contributes to saving network resources and 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 March 4, 2015.

Table of Contents

1. Introduction

The Internet Group Management Protocol (IGMP) version 3 [2] for IPv4 and the Multicast Listener Discovery Protocol (MLD) version 2 [3] for IPv6 are the standard protocols used by member hosts and multicast routers. Lightweight IGMPv3 and Lightweight MLDv2 (or LW-IGMPv3 and LW-MLDv2) [4] are subsets of the standard IGMPv3 and MLDv2.

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 [8]). The "unsolicited" report messages are sent only when the host joins/leaves the channels. Since IGMP/MLD are non-reliable protocols, unsolicited report messages may be lost or may not reach upstream routers. To alleviate this problem, unsolicited report messages are retransmitted a number of times according to the value of the [Robustness Variable] defined in [2][3].

In addition, a querier router periodically sends IGMP/MLD General Query messages every General Query timer interval (i.e. [Query Interval] value defined in [2][3]). Upon receiving the query messages, the member hosts reply with "solicited" report messages. Routers then keep their membership state information up to date. However, this approach still 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.

IGMPv3 [2], MLDv2 [3], and these lightweight protocols [4] can provide the ability to keep track of the downstream (adjacent) multicast membership state in multicast routers, yet the specifications are not clearly given. This document describes the "IGMP/MLD-based explicit member tracking function" for multicast routers and a way for routers to implement the function. By enabling this explicit tracking function, routers can keep track of the downstream multicast membership state.

The explicit tracking function is important for the scalability of multicast networks, and might be widely implemented in modern multicast routers. However, it could seriously break IGMP/MLD communications if not implemented or configured correctly. For example, the explicit tracking function is useful for shortening leave latency, while wrong implementations or configurations in routers on a LAN may not work properly that the operators expect.

This document aims to get the explicit tracking function correctly. Regarding the way for shortening leave latency, it specifies the way to do it by tuning the values used by IGMPv3 [2], MLDv2 [3], and their lightweight version protocols [4], or by enabling the mechanism called "specific query suppression" with a robust link state. The latter mechanism does not make the router send any specific query message(s) and immediately leave the group or sources when the sole member has left according to its membership state information.

This document describes the risk of having wrong membership state as well. The explicit tracking function does not change the reliability of the message transmission. 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. This document guides for setting up appropriate values or mechanisms used with the explicit tracking function in routers.

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 might be sensitive. As the security consideration, this document describes that operators may decide to disable this function when their routers have insufficient memory resources, despite the benefits.

The explicit tracking function does not change message formats used by IGMPv3 [2], MLDv2 [3], and their lightweight version protocols [4]; 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 [1].

3. 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 or modifies this membership state information to maintain the membership state up to date.

The membership state information consists of the following information:

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

where "S" denotes source address, "G" denotes group or multicast address, and each receiver record is of the form:

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

In the state information, each S and G indicates a single IPv4/IPv6 address. S is set to "All sources" for Any-Source Multicast (ASM) communication (i.e., (*,G) join reception). In order to simplify the implementation, Lightweight-IGMPv3/MLDv2 [4] do 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, the router translates the request to a (*,G) join state/leave request and records the state and the receivers' addresses in the maintained membership state information.

The membership state information must be identified properly even though a receiver (i.e., IGMP/MLD Report sender) sends the identical report messages multiple times. And the maintained membership state information will be flushed when the router reboots or restarts the multicast routing processes.

4. Specific Query Suppression

In accordance with [2] and [3], when a router receives the State-Change Report and needs to confirm whether any hosts are still interested in a channel or not, the router sends the corresponding Group-Specific or Group-and-Source Specific Query messages as defined in Section 6.4.2 of [2] and Section 7.4.2 of [3]. The queries sent by actions defined in these sections need to be transmitted [Last Member Query Count] (LMQC) or [Last Listener Query Count] (LLQC) times, once every [Last Member Query Interval] (LMQI) or [Last Listener Query Interval] (LLQI), in order to confirm the sole member. (The default values for LMQI/LLQI defined in [2][3] are 1 second. The default values for LMQC/LLQC are the [Robustness Variable] value whose default value is 2.) All member hosts joining the identical channel then reply with their own states after acquiring these query messages. However, transmitting a large number of IGMP/MLD Report messages consumes network resources, and this may pose a particular problem especially when many hosts joining the identical channel send these reports simultaneously.

The explicit tracking function provides a mechanism called "specific query suppression". With the specific query suppression, regardless of the LMQC/LLQC values, if the router receives one or more replies from the downstream member(s), it SHOULD stop (i.e., cancel) retransmitting the specific query message(s) for the specified source and/or group. It reduces the number of Group-Specific or Group-and-Source Specific Query messages transmitted from a router, and in turn reduces the number of Current-State Report messages transmitted from member hosts. This contributes to saving network resources.

The specific query suppression MAY define an option called "robust link state". A router enabling the specific query suppression with a robust link state does not send any specific query message(s) and immediately leave the group or sources when the sole member has left according to its membership state information. The specific query suppression with a robust link state hence does not rely on LMQC/LLQC and LMQI/LLQI values. This contributes to shortening leave latency described in Section 5. However, this behavior requires that the router perfectly tracks all member hosts. (See a risk of wrong membership expectation described in Section 6.)

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 wrong membership expectation or by the case in which multiple multicast routers exist on a LAN and the querier router is not the forwarder router. The former case is described in Section 6. For the latter case, 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. Then the routers (including the forwarder) on the same LAN do not receive a Current-State Report message from the corresponding member hosts. The forwarder in this case may prune the routing path, although there are other member hosts subscribing to the channel on the LAN.

5. Shortening Leave Latency

A router enabling the explicit tracking function can shorten leave latencies by tuning the following values; [Last Member Query Count] (LMQC), [Last Listener Query Count] (LLQC), [Last Member Query Interval] (LMQI), [Last Listener Query Interval] (LLQI), and [Robustness Variable] values.

The [Last Member Query Interval] (LMQI) and [Last Listener Query Interval] (LLQI) values defined in the standard specifications [2][3] specify the maximum time allowed for a member host to send 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 values for LMQI/LLQI defined in [2][3] are 1 second, yet, for a router enabling the explicit tracking function, the LMQI/LLQI may be set to 1 second or shorter. As well, the default values for LMQC/LLQC are the [Robustness Variable] value whose default value is 2, yet the LMQC/LLQC may be set to 1 for the router. Smaller LMQC/ LLQC values give shorter LMQT/LLQT, which shorten the leave latencies.

Furthermore, if operators are confident that their link is fairly robust (e.g., the [Robustness Variable] value is appropriately configured so that the chances of unsolicited messages being lost are sufficiently low), and if the querier router always acts as the forwarder router for all multicast channels in the LAN, they will set smaller LMQC/LLQC and shorter LMQI/LLQI (and hence shorter LMQT/LLQT) with the specific query suppression, or enable the specific query suppression with a robust link state (Section 4) for their routers.

Note that setting smaller LMQC/LLQC and shorter LMQI/LLQI values or adopting the specific query suppression with a robust link state poses the risk of wrong membership state described in Section 6.

Operators setting these values or enabling that mechanism must recognize this tradeoff.

6. Risk of Wrong Membership State

There are possibilities that a router's membership expectation is 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 one or more than one member host exists for that multicast channel.

The first case may occur in an environment where the sole member host departs the network without sending a State-Change Report message. 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 (and then the group or source timers are expired). However, this situation prolongs leave latency and wastes network resources since the router forwards unneeded traffic for a while.

The second case occurs when a router sends a specific query but does not receive a Current-State 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 loose the data packets until the routing path is reestablished and the data forwarding is restarted.

If operators do not believe that their link is fairly robust or that they can configure the [Robustness Variable] value appropriately, they may configure the LMQC/LLQC value to 2 (the default value of the [Robustness Variable] value) or bigger value for their routers. In this case, the routers would enable the explicit tracking function but may want to disable the specific query suppression specified in Section 4. Such configurations will not contribute to saving network resources, but reduce the risk of the incorrect membership expectation.

7. All-Zero and Unspecified Source Addresses

The IGMPv3 specification [2] 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.

When a router enabling the explicit tracking function receives IGMP report messages with an all-zero source address, it deals with the IGMP report messages correctly as defined in [2] and continuously keeps track of the membership state. However, the router SHOULD NOT maintain the host specifying all-zero source address in its membership state information. The router will maintain its membership state information by checking Current-State reports as ordinary routers do.

On the other hand, the MLDv2 specification [3] mentions that routers silently discard a message that is sent with an invalid link-local address or sent with the unspecified address (::), without taking any action, because of security considerations. According to this specification, whether the explicit tracking function is used or not, a router does not deal with a member hosts sending an MLD report message with the unspecified source address.

8. Compatibility with Older Version Protocols

The explicit tracking function does not work with older versions of IGMP or MLD, IGMPv1 [5], IGMPv2 [6], or MLDv1 [7], 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 supporting IGMPv3/MLDv2 enable the host compatibility mode defined in [2][3]. The host compatibility mode of an interface changes the operational protocol version on the LAN whenever an older version query (than the current compatibility mode) is heard or when certain timer conditions occur. The routers can hence support downstream hosts that are not upgraded to the latest versions and run membership report suppression.

Therefore, if a multicast router supporting IGMPv3/MLDv2 and enabling the explicit tracking function changes its compatibility mode to the older versions, the router SHOULD disable the explicit tracking function while it acts as the older version router.

9. Interoperability

There might be various ways to implement the explicit tracking function. Some existing implementations may not implement the mechanisms such as specific query suppression described in this document. Yet, the explicit tracking function does not change on-wire behavior, and the function or mechanisms described in this document do not break the interoperability between the existing implementations and the implementation based on this specification.

On the other hand, for the future implementation for the explicit tracking function, since this document specifies the minimum but effective sets of the explicit tracking function, it is RECOMMENDED to refer and follow this specification as the standard implementation for that function.

10. IANA Considerations

This document has no actions for IANA.

11. Security Considerations

The explicit tracking function potentially requires a large amount of memory so that routers keep all membership states. It gives some impact in the cases where (1) a router attaches to a link or an IGMP/MLD proxy device [8] that has a large number of member hosts, and a router has insufficient memory resources to maintain a large number of member hosts, or (2) a malicious host sends a large number of invalid IGMP/MLD State-Change Report messages without any intent to join the specified channels.

For the first case, operators may disable the explicit tracking function, despite the benefits mentioned above. For the second case, some serious threats may be induced. For instance;

1. Transmitting a large number of invalid IGMP/MLD report messages consumes network resources.

2. Keeping a large number of invalid membership states on a router consumes the router’s memory resources.

3. Dealing with a large number of invalid membership states on a router consumes the router’s CPU resources.

In order to mitigate such threats, a router enabling the explicit tracking function may limit a total amount of membership information the router can store, or may rate-limit State-Change Report messages per host. When the router enables rate-limiting per host, the router MAY ignore the received State-Change Report messages to minimize the processing overhead or prevent DoS attacks. The rate limit is left to the router’s implementation.

12. Acknowledgements

Luis M. Contreras, Toerless Eckert, Adrian Farrel, Sergio Figueiredo, Bharat Joshi, Nicolai Leymann, Magnus Nystrom, Stig Venaas, and others provided many constructive and insightful comments.

13. References

13.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to indicate requirement levels", RFC 2119, March 1997.

[2] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[3] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

[4] 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.

13.2. Informative References

[5] Deering, S., "Host Extensions for IP Multicasting", RFC 1112, August 1989.

[6] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.

[7] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.

[8] 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.

Author’s Address

Hitoshi Asaeda
National Institute of Information and Communications Technology
4-2-1 Nukui-Kitamachi
Koganei, Tokyo 184-8795
Japan

Email: asaeda@nict.go.jp