IGMP and MLD Extensions for Mobile Hosts and Routers
draft-asaeda-multimob-igmp-mld-mobility-extensions-00

Status of this Memo

Copyright Notice

Abstract

This document describes the IGMP and MLD protocol extensions for mobile hosts and routers. IGMP and MLD are necessary protocols for hosts to request join or leave multicast sessions. While the regular IGMP and MLD protocols support communication between mobile hosts and routers over wireless networks, this document discusses the conditions how mobile hosts and routers use IGMP and MLD in their communication more effectively.

Conventions used in this document

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].

Table of Contents

1. Introduction

The Internet Group Management Protocol (IGMP) for IPv4 and the Multicast Listener Discovery Protocol (MLD) for IPv6 are the necessary protocols for hosts to request to join or leave multicast sessions. These protocols must be also supported by the upstream multicast routers that have the downstream multicast member hosts on the same LAN. By definition, IGMP and MLD work on wireless networks; there is not necessary to change these protocol specs for wireless communication environments. However, when mobile hosts or routers attached on a wireless link multicast IGMP/MLD messages, the transmitted IGMP/MLD messages are flooded to all mobile hosts and routers on the link, where a large amount of flooding data consumes battery power of each mobile host. In addition, it takes higher costs for the upstream router to maintain a large number of IGMP/MLD messages, and in this situation, it takes longer time for the router to converge the membership state (i.e. existence of downstream member hosts), where the longer convergence negatively affects leave latency.

To create the feasible condition to communicate mobile hosts and routers with IGMP/MLD, it is required to "ease processing cost or battery power consumption by eliminating transmission of a large number of IGMP/MLD messages via flooding" and "realize fast state convergence by successive monitoring whether downstream members exist or not". The possible approach to fulfill these requirements is relevant; if the upstream router traces all downstream members by limiting the number of solicited membership reports (by periodical IGMP/MLD Query), the number of transmitted IGMP/MLD messages is effectively reduced, and the upstream router can immediately update the membership information and proceed the fast leave.

The function that enables to trace all downstream members is supported by IGMPv3 [2] and MLDv2 [3]. In the previous version protocols, IGMPv1 [4], IGMPv2 [5], and MLDv1 [6], a host would cancel sending a pending membership reports requested by IGMP/MLD Query if a similar report was observed from another member on the network. This specification effectively reduced a possibility of network congestion or message flooding, but precluded the function for an upstream router to track membership status. On the other hand, in IGMPv3 and MLDv2, the membership report suppression mechanism has been removed, and therefore all downstream member hosts must send their membership reports to an upstream router and the router can keep track of membership status.

If the report suppression mechanism is removed from the host-side protocols, the upstream router supporting IGMPv3/MLDv2 receives all membership reports from the downstream hosts. One may deduce that

the router does not need to periodically send IGMPv3/MLDv2 Query messages to trace membership status. However, IGMPv3/MLDv2 capable routers usually configure to send periodical IGMP/MLD Query messages to seek membership information to the downstream hosts, and disable the function that keeps track of membership status. One of the reasons is that IGMP/MLD message is non-reliable and may be lost in the transmission, and therefore the router would need to confirm the membership by sending query messages. The other reason is that, for keeping track of membership status, the router needs additional processing cost and a possibly large size of the memory to record all member information. The requirement to keep the compatibility mode with older version IGMP/MLD is also the reason, because the router needs to support the downstream hosts that are not upgraded to the latest versions of IGMP/MLD and run the report suppression mechanism.

There is one more important function in IGMPv3 and MLDv2. IGMPv3 and MLDv2 provide the ability for hosts to report interest in receiving packets only from specific source addresses to upstream routers. With IGMPv3/MLDv2, the mobile host specifies the interesting multicast and source addresses with INCLUDE filter mode in the join request. Upon reception, the upstream router establishes the shortest path tree toward the source without coordinating a shared tree. This function is called the source filtering function and required to support Source-Specific Multicast (SSM) [7]. SSM is advantageous to multicast routing tree establishment and satisfies current and future needs.

IGMPv3 and MLDv2 support another operation with EXCLUDE filter mode. When a mobile host specifies multicast and source addresses with EXCLUDE filter mode in the join request, an upstream router forwards the multicast packets sent from all sources *except* the specified sources. In fact, this operation gives the complexity in the host-side procedure. If any application running on a host requests an EXCLUDE filter mode operation, the host sets the interface state to EXCLUDE mode for the requested multicast address, and the source address list of the interface record is the intersection of the source address lists requested by all applications in EXCLUDE mode, minus the source addresses that appear in any application in INCLUDE mode. The state transition that maintains the interface record is complex, and the implementation cost will be relatively high for mobile hosts.

Furthermore, specifying non-interesting source addresses with EXCLUDE filter mode reduces the advantage of scalable routing tree establishment in an SSM manner, because an upstream router needs to refresh and re-generate some or all of the corresponding routing tree including the RPT whenever the router receives join request with EXCLUDE filter mode from the downstream hosts. This manner increases

the tree maintenance cost to not only the upstream multicast routers but other routers existed on the routing paths. While the mobile multicast communication does not prohibit a traditional (*,G) join request (which is a join request with EXCLUDE filter mode without specifying any source address), all other join requests with EXCLUDE filter mode should be eliminated from the mobile multicast communication.

This document describes the IGMP and MLD protocol extensions for mobile hosts and routers, and discusses the conditions how mobile hosts and routers use IGMP and MLD in their communication over wireless networks effectively. The selective solutions that provide tangible benefits to the mobile hosts and routers are given by "keeping track of membership status by eliminating a report suppression mechanism", "varying IGMP/MLD Query types and values to tune the number of responses", and "using a source filtering mechanism in a lightweight manner". The proposed solutions do not require changing the IGMP and MLD protocols. This condition is advantageous to the deployment.

2. Configurations

2.1. Tracking of Membership Status

Mobile hosts use IGMP and MLD to request to join or leave multicast sessions. When the upstream routers receive the IGMP/MLD reports, they recognize the membership status on the LAN. To update the membership status, the routers send IGMP/MLD Query messages periodically as a soft-state approach does, and the member hosts reply IGMP/MLD report messages upon reception.

IGMP/MLD Query is therefore necessary to obtain the up-to-date membership information, but a large number of the reply messages sent from all member hosts may cause network congestion or consume network bandwidth. To escape from the trouble, a membership report suppression mechanism was proposed in the traditional IGMP and MLD [4][5][6]. By the report suppression mechanism, a host would cancel sending a pending membership reports requested by IGMP/MLD Query if a similar report was observed from another member on the network. However, the report suppression mechanism precluded the function for an upstream router to track membership status. In IGMPv3 and MLDv2, it is hence decided that the membership report suppression mechanism has been removed, and all downstream member hosts must send their membership reports to an upstream router.

By eliminating membership report suppression, an IGMPv3 or MLDv2 capable upstream router could trace all downstream members and track per-host membership status on the interface. This reduces the number of solicited membership reports by periodical IGMP/MLD Query, and finally the total number of transmitted IGMP/MLD messages can be drastically reduced. This is beneficial especially to mobile hosts that do not have enough battery power, since flooding IGMP/MLD messages on a LAN makes all multicast members give significant attention and induces power consumption to the member hosts. This also allows the upstream router to proceed fast leaves, because the router can immediately converge and update the membership information, ideally.

2.2. IGMP/MLD Query Coordination

2.2.1. Unicasting IGMP/MLD General Query

IGMP and MLD are asymmetric and non-reliable protocols; multicast routers still need to solicited membership reports by periodical IGMP/MLD Query, in order to be robust in front of host or link failures and packet loss. Moreover, it happens that mobile hosts may turn off or move from the wireless network to other wireless network managed by the different router without any notification (i.e. leave

report). Therefore, even though multicast routers keep track of the interests of downstream member hosts attached on the same LAN, IGMP/ MLD Query must be sent periodically.

However, periodical message flooding using the all-hosts multicast address (i.e. 224.0.0.1 or ff02::1) as its IP destination address gives the unwilling situation to the mobile hosts. When the mobile hosts are operating in a dormant mode and not communicating with others, they should not be woken up by IGMP/MLD General Query and keep sleeping for saving the battery power. In this case, only the hosts that are receiving multicast contents should make the response to the router.

IGMPv3 and MLDv2 specifications [2][3] say that a host MUST accept and procss any Query whose IP Destination Address field contains any of the addresses (unicast or multicast) assigned to the interface on which the Query arrives. According to the scenario, it unicasts the message to tracked member hosts in the [Unicast Query Interval]. It is happened especially when a multicast router has a small number of mobile hosts that are listening different multicast sessions. In this situation, the router multicasts IGMP/MLD General Query with longer [Query Interval] (described in Section 2.2.3).

[TODO: Define [Unicast Query Interval] value. The value could be same of the default [Query Interval]?]

2.2.2. Multicasting IGMP/MLD Group-Specific Query

In the standard protocols [2][3], IGMP/MLD Group-Specific Query is sent to verify there are no hosts that desire reception of the specified group or to rebuild the desired reception state for a particular group. Group-Specific Queries are sent when a router receives a State-Change record indicating a host is leaving a group.

In a dormant mode operation, IGMP/MLD Group-Specific Query can be also used to build and refresh the group membership state of hosts on attached networks. When more than one mobile host join the multicast contents whose multicast address is same, IGMP/MLD Group-Specific Query can be sent to maintain the group membership state of mobile hosts on attached networks, instead of IGMP/MLD General Query. Since IGMP/MLD Group-Specific Query specifies the corresponding multicast address (not the all-hosts multicast address) as its IP destination address, dormant mode hosts that do not join any multicast session are not woken up by the IGMP/MLD Group-Specific Query and only active group member hosts that have been receiving multicast contents would reply IGMP/MLD reports. This manner contributes to reducing the number of transmitted IGMP/MLD messages.

The [Multicast Group-Query Interval] is the interval between Group-Specific Queries sent by the querier, i.e., the router that sends the Group-Specific Query. This value is same of the default [Query Interval] value the regular IGMP and MLD define [2][3].

[TODO: Define [Multicast Group-Query Interval].]

2.2.3. Values in IGMP/MLD Query

A multicast router operating in a dormant mode keeps track of the membership status and checks the membership status by transmitting unicast IGMP/MLD General Query or multicast IGMP/MLD Group-Specific Query. Cooperating with these scenarios, the message interval between IGMP/MLD General Queries is set to longer than the default [Query Interval] value.

The Query Interval is the interval between General Queries sent by the querier, and the default value is 125 seconds [2][3]. By varying the [Query Interval], multicast routers can tune the number of IGMP messages on the network; larger values cause IGMP Queries to be sent less often.

[TODO: We will provide the appropriate [Query Interval] value that would fit in the mobile communication environment based on some experimental results. In our current sense, this value should be larger than the default value the regular IGMP and MLD define.]

The Query Response Interval is the Max Response Time (or Max Response Delay) used to calculate the Max Resp Code inserted into the periodic General Queries, and the default value is 10 seconds [2][3]. By varying the [Query Response Interval], multicast routers can tune the burstiness of IGMP messages on the network; larger values make the traffic less bursty, as host responses are spread out over a larger interval.

[TODO: We will provide the appropriate [Query Response Interval] value that would fit in the mobile communication environment based on some experimental results. In our current sense, this value should be less than the default value the regular IGMP and MLD define, because, while the larger Query Interval does not reduce the number of transmitted IGMP/MLD messages, it may cause slow leave latency.]

Mobile hosts may receive a variety of Queries on different interfaces and of different kinds (e.g., General Queries, Group-Specific Queries, and Group-and-Source-Specific Queries), each of which may require its own delayed response.

[TODO: The timer management for each queries may or should be

independent. E.g. the timer value for General Query should be longer than the one of other queries. We will investigate this issue.]

To cover the possibility of unsolicited reports being missed by multicast routers, unsolicited reports are retransmitted [Robustness Variable] - 1 more times, at intervals chosen at random from the defined range [2][3]. The QRV (Querier's Robustness Variable) field in IGMP/MLD Query contains the [Robustness Variable] value used by the querier. Routers adopt the QRV value from the most recently received Query as their own [Robustness Variable] value, whose range should be set between "1" to "7". While the default [Robustness Variable] value defined in IGMPv3 [2] and MLDv2 [3] is "2", the [Robustness Variable] value announced by the querier must not be "0" and should not be "1".

[TODO: We will propose the robustness values that would be adjusted according to the number of receivers. In our current sense, this value should not be bigger than "2" especially when the [Query Response Interval] is set to less than its default value.]

2.3. IGMP/MLD Querier Selection

[TODO: Is there any condition or assumption in which multiple multicast routers exist in a single wireless link? If there is the case, do we need to consider IGMP/MLD querier selection mechanism and the corresponding timer values or intervals? The Querier's Query Interval Code (QQIC) field specifies the [Query Interval] used by the querier may be tuned. The actual interval, called the Querier's Query Interval (QQI), is derived from QQIC. Multicast routers that are not the current querier adopt the QQI value from the most recently received Query as their own [Query Interval] value.]

3. Protocol Extensions

IGMPv3 and MLDv2 enable all member hosts to send membership reports to the upstream routers. Not only this function, IGMPv3 and MLDv2 support a source filtering function. An IGMPv3 or MLDv2 capable host can tell its upstream router which group it would like to join by specifying which sources it does (or does not) intend to receive multicast traffic from. IGMPv3 and MLDv2 add the capability for a multicast router to also learn which sources are (and are not) of interest to neighboring hosts, for packets sent to any particular multicast address. This source filtering function is required to invoke Source-Specific Multicast (SSM) [7].

IGMPv3 and MLDv2 introduce antithetic filter modes, INCLUDE and EXCLUDE filter modes, to expand the source filtering function. If a host wants to receive from specific sources, it sends an IGMPv3 or MLDv2 report with the filter mode set to INCLUDE. If the host does not want to receive from some sources, it sends a report with the filter mode set to EXCLUDE. A source list for the given sources shall be included in the report message. INCLUDE and EXCLUDE filter modes are also defined in a multicast router to process the IGMPv3 or MLDv2 reports. When a multicast router receives the report messages from its downstream hosts, it forwards the corresponding multicast traffic by managing requested group and source addresses.

However, practical applications do not use EXCLUDE mode to block sources very often, because a user or application usually wants to specify desired source addresses, not undesired source addresses. In addition, this scheme leads an implementation cost to mobile hosts and complex procedures to maintain coexisting situation of the interesting source address lists with INCLUDE filter mode or non-interesting source address lists with EXCLUDE filter mode.

Recently, Lightweight-IGMPv3 (LW-IGMPv3) and Lightweight-MLDv2 (LW-MLDv2) [8] are proposed in the IETF MBONED working group. These protocols are the simplified versions of IGMPv3 and MLDv2, and eliminate an EXCLUDE filter mode operation. Not only are LW-IGMPv3 and LW-MLDv2 fully compatible with the full version of these protocols (i.e., the standard IGMPv3 and MLDv2), but also the protocol operations made by hosts and routers are simplified in the lightweight manner, and complicated operations are effectively reduced. LW-IGMPv3 and LW-MLDv2 give the opportunity to grow SSM use.

In the lightweight protocols, EXCLUDE mode on the host part is preserved only for EXCLUDE (*,G) join/leave, which denotes a non-source-specific group report (known as the traditional (*,G) join/leave) and is equivalent to the group membership join/leave triggered

by IGMPv2/IGMPv1/MLDv1.

The aim of LW-IGMPv3 and LW-MLDv2 is not only for contributing to the simpler implementation or reducing the memory size on a host. Another advantage is that it reduces the processing cost on upstream routers by eliminating the EXCLUDE filter mode operations. If both INCLUDE and EXCLUDE filter mode operations are supported in the networks, the routers need to maintain all source addresses joined from their downstream hosts. Even if a Shortest-Path Tree (SPT) is well coordinated, the routers need to refresh (and re-generate) some or all of the corresponding routing paths including the Rendezvous-Point Tree (RPT) whenever the downstream host requests EXLUDE filter mode join. LW-IGMPv3 and LW-MLDv2 preclude the unwilling situation. Since there is no side-effect, this document recommend to adopt LW-IGMPv3 and LW-MLDv2 to mobile hosts and routers, or eliminate EXCLUDE filter mode operation from mobile hosts if IGMPv3 and MLDv2 are adopted to hosts.

4. Implementations

4.1. Host-Side Implementation

Mobile hosts should implement IGMPv3 or LW-IGMPv3 for IPv4 multicast and MLDv2 or LW-MLDv2 for IPv6 multicast. All of these protocols eliminate a membership report suppression mechanism, and make hosts work with the function multicast routers use to trace downstream member hosts. These protocols also support SSM. According to the protocol requirement aforementioned, however, this document recommends to implement LW-IGMPv3 for IPv4 and LW-MLDv2 for IPv6 [8] rather than the full version protocols.

4.2. Router-Side Implementation

To keep track of multicast membership status and cooperate with SSM capable mobile hosts, multicast routers must implement IGMPv3/LW-IGMPv3 or MLDv2/LW-MLDv2. The protocol requirement aforementioned does not require modification of the IGMPv3/LW-IGMPv3 and MLDv2/LW-MLDv2 protocol specifications. This condition is advantageous to the deployment. However, regarding the router-side implementation, the function to trace downstream members requires the hardware requirement that would cost the router additional hardware resources, especially CPU and memory resources.

[TODO: This document assumes that multicast routers are not tiny and non-powerful systems nor battery or power sensitive. Our assumption is correct?]

As well as the host-side implementation, the elimination of the EXCLUDE filter mode will greatly simplify the router behavior, e.g. the action on reception of reports and the setting of the timers. This document therefore recommends to implement LW-IGMPv3 for IPv4 and LW-MLDv2 for IPv6 rather than their full version protocols. The detailed operation being simplified is described in [8].

5. Interoperability

TBD.

[TODO: We believe it would be currently feasible to assume the routers who take care of mobile hosts MUST be IGMPv3/MLDv2 capable (regardless whether the protocols are the full version or not). What we should understand is whether there is the case that mobile hosts may not be IGMPv3/MLDv2 capable or not.]

6. Security Considerations

TBD.

7. References

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

[5] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2373, July 1997.

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

[7] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.

7.2. Informative References

[8] Liu, H., Cao, W., and H. Asaeda, "Lightweight IGMPv3 and MLDv2 Protocols", draft-ietf-mboned-lightweight-igmpv3-mldv2-01.txt (work in progress), June 2007.

Author's Address

Hitoshi Asaeda
Keio University
Graduate School of Media and Governance
5322 Endo
Fujisawa, Kanagawa 252-8520
Japan

Email: asaeda@wide.ad.jp

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).