MULTIMOB Working Group
Internet-Draft
Expires: August 27, 2011

H. Asaeda
Y. Uchida
Keio University
H. Liu
Q. Wu
Huawei Technologies

February 23, 2011

Tuning the Behavior of IGMP and MLD for Mobile Hosts and Routers
draft-asaeda-multimob-igmp-mld-optimization-05

Abstract

IGMP and MLD are the protocols used by hosts to report their IP multicast group memberships to neighboring multicast routers. This document describes the ways of IGMPv3 and MLDv2 protocol optimization for mobility, and aims to become a guideline for query and other timers tuning.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

The Internet Group Management Protocol (IGMP) [2] for IPv4 and the Multicast Listener Discovery Protocol (MLD) [3] for IPv6 are the standard protocols for hosts to initiate joining or leaving multicast sessions. These protocols must be also supported by multicast routers or IGMP/MLD proxies [11] that maintain multicast membership information on their downstream interfaces. Conceptually, IGMP and MLD work on wireless networks. However, wireless access technologies operate on a shared medium or a point-to-point link with limited frequency and bandwidth. In many wireless regimes, it is desirable to minimize multicast-related signaling to preserve the limited resources of battery powered mobile devices and the constrained transmission capacities of the networks. A mobile host may cause initiation and termination of a multicast service in the new or the previous network upon its movement. Slow multicast service activation following a join may degrade reception quality. Slow service termination triggered by IGMP/MLD querying or by a rapid departure of the mobile host without leaving the group in the previous network may waste network resources.

To create the optimal multicast membership management condition, IGMP and MLD protocols could be tuned to "ease a mobile host's processing cost or battery power consumption by IGMP/MLD Query transmission timing coordination by routers" and "realize fast state convergence by successive monitoring whether downstream members exist or not".

This document describes the ways of tuning the IGMPv3 and MLDv2 protocol behavior for mobility, including query and other timers tuning. The selective optimization that provides tangible benefits to the mobile hosts and routers is given by keeping track of downstream hosts' membership status and varying IGMP/MLD Query types and values to tune the number of responses. The proposed behavior interoperates with the IGMPv3 and MLDv2 protocols.

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. Explicit Tracking of Membership Status

Mobile hosts use IGMP and MLD to request to join or leave multicast sessions. When the adjacent upstream routers receive the IGMP/MLD Report messages, they recognize the membership status on the link. 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.

The "explicit tracking function" [9] is the possible approach to reduce the transmitted number of IGMP/MLD messages and contribute to mobile communications. It enables the router to keep track of the membership status of the downstream IGMPv3 or MLDv2 member hosts.

The explicit tracking function reduces the chance of Group-Specific or Group-and-Source Specific Query transmission. Whenever a router that does not enable the explicit tracking function receives the State-Change Report and the router’s membership state is changed to block some source or group, it sends the corresponding Group-Specific or Group-and-Source Specific Query messages to confirm whether the Report sender is the last member host or not. However, if a router enables the explicit tracking function, it does not always need to ask Current-State Report message transmission to the receiver hosts since the router recognizes the (potential) last member host when it receives the State-Change Report. The router can therefore send IGMP/MLD Group-Specific and Group-and-Source Specific Queries LMQC/LLQC times (see Section 4.3 for LMQC/LLQC) only when it recognizes the last member has left from the network. This reduces the transmitted number of Current-State Report messages.

Enabling the explicit tracking function is advantageous for mobile multicast, but the function requires additional processing capability and a possibly large memory for routers to keep all membership status. Especially when a router needs to maintain a large number of receiver hosts, this resource requirement may be potentially-impacted. Therefore, in this document, we propose that adjacent upstream multicast routers SHOULD enable the explicit tracking function for IP multicast communications on wireless networks, if they have enough resources. If operators think that their routers do not have enough resources, they MAY decide to disable this function on their routers. Note that whether routers enable the explicit tracking function or not, they need to maintain downstream membership status by sending IGMPv3/MLDv2 General Query messages as some IGMPv3/MLDv2 messages may be lost during transmission.

4. Tuning IGMP/MLD Timers and Values

4.1. Tuning IGMP/MLD General Query Interval

IGMP and MLD are non-reliable protocols; to cover the possibility of a State-Change Report being missed by one or more multicast routers, "hosts retransmit the same State-Change Report messages [Robustness Variable] - 1 more times", at intervals chosen at random from the range (0, [Unsolicited Report Interval]) [2][3]. Although this behavior increases the protocol robustness, it does not guarantee that the State-Change Report is reached to the routers. Therefore, routers still need to refresh the downstream membership information by receiving Current-State Report periodically solicited by IGMP/MLD General Query sent in the [Query Interval] period, in order to be robust in front of host or link failures and packet loss. It also supports the situation that mobile hosts turn off or move from the wireless network to other wireless network managed by the different router without any notification (e.g., leave request).

The [Query Interval] is the interval between General Queries sent by the regular IGMPv3/MLDv2 querier, and the default value is 125 seconds [2][3]. By varying the [Query Interval], multicast routers can tune the number of IGMP/MLD messages on the network; larger values cause IGMP/MLD Queries to be sent less often.

This document proposes 150 seconds for the [Query Interval] value by changing the Querier's Query Interval Code (QQIC) field specified in the IGMP/MLD Query message, for the case that a router enabling the explicit tracking function sends General Query and potentially operates a large number of member hosts such as more than 200 hosts on the wireless link. This longer interval value contributes to minimizing traffic of Report messages and battery power consumption for mobile hosts.

On the other hand, this document also proposes 60 to 90 seconds for the [Query Interval] value for the case that a router enabling the explicit tracking function attaches to a wireless link having higher capacity of the resource. This shorter interval contributes to quick synchronization of the membership information tracked by the router but may consume battery power of mobile hosts.

If a router does not enable the explicit tracking function, the [Query Interval] value would be its default value, 125 seconds.

4.2. Tuning IGMP/MLD Query Response Interval

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. Its default value is 10 seconds expressed by "Max Resp Code=100" for IGMPv3 [2] and "Maximum Response Code=10000" for MLDv2 [3]. By varying the [Query Response Interval], multicast routers can tune the burstiness of IGMP/MLD messages on the network; larger values make the traffic less bursty as host responses are spread out over a larger interval, but will increase join latency when State-Change Report is missing.

According to our experimental analysis, this document proposes two tuning scenarios for tuning the [Query Response Interval] value in different wireless link conditions; one scenario is for a wireless link with a lower capacity of network resource or a lossy link, and the other scenario is for a wireless link with enough capacity or reliable condition for IGMP/MLD message transmission.

Regarding the first scenario, for instance, when a multicast router attaches to a bursty IEEE 802.11b link, the router configures the longer [Query Response Interval] value, such as 10 to 20 (sec). This configuration will reduce congestion of the Current-State Report messages on a link but may increase join latency and leave latency when the unsolicited messages (State-Change Record) are lost on the router.

The second scenario may happen for a multicast router attaching to a wireless link having higher capacity of the resource or a point-to-(multi-)point link such as an IEEE 802.16e link, because IGMP/MLD messages do not seriously affect the link condition. The router can seek Current-State Report messages with the shorter [Query Response Interval] value, such as 5 to 10 (sec). This configuration will contribute to quickly (at some level) discovering non-tracked member hosts and synchronizing the membership information.

4.3. Tuning Last Member Query Timer (LMQT) and Last Listener Query Timer (LLQT)

Shortening the Last Member Query Timer (LMQT) for IGMPv3 and the Last Listener Query Timer (LLQT) for MLDv2 contributes to minimizing leave latency. LMQT is represented by the Last Member Query Interval (LMQI), multiplied by the Last Member Query Count (LMQC), and LLQT is represented by the Last Listener Query Interval (LLQI), multiplied by the Last Listener Query Count (LLQC).

While LMQI and LLQI are changeable, it is reasonable to use the default values (i.e., 1 second) for LMQI and LLQI in a wireless network. LMQC and LLQC, whose default value is the [Robustness Variable] value, are also tunable. Therefore, LMQC and LLQC MAY be set to "1" for routers enabling the explicit tracking function, and then LMQT and LLQT are set to 1 second. However, setting LMQC and

LLQC to 1 increases the risk of missing the last member; LMQC and LLQC SHOULD be set to 1 only when network operators think that their wireless link is stable enough.

On the other hand, if network operators think that their wireless link is lossy (e.g., due to a large number of attached hosts or limited resources), they MAY set LMQC and LLQC to "2" for their routers enabling the explicit tracking function. Although bigger LMQC and LLQC values may cause longer leave latency, the risk of missing the last member will be reduced.

4.4. Tuning Startup Query Interval

The [Startup Query Interval] is the interval between General Queries sent by a Querier on startup. The default value is 1/4 of [Query Interval]; however, this document recommends the use of its shortened value such as 1 second since the shorter value would contribute to smooth handover for mobile hosts using, e.g., PMIPv6 [12]. Note that the [Startup Query Interval] is a static value and cannot be changed by any external signal. Therefore operators who maintain routers and wireless links must properly configure this value.

4.5. Tuning Robustness Variable

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. The default [Robustness Variable] value defined in IGMPv3 [2] and MLDv2 [3] is "2".

This document proposes "2" for the [Robustness Variable] value for mobility, when a router attaches to a wireless link having lower capacity of the resource or a large number of hosts. For a router that attaches to a wireless link having higher capacity of the resource or reliable condition, it is not required to retransmit the same State-Change Report message; hence the router sets the [Robustness Variable] to "1". Note that whether the explicit tracking function is enabled or not, the [Robustness Variable] value SHOULD NOT be bigger than "2".

5. Destination Address of Specific Query

IGMP/MLD Group-Specific and Group-and-Source Specific Queries defined in [2][3] are sent to verify whether there are hosts that desire reception of the specified group or a set of sources or to rebuild the desired reception state for a particular group or a set of sources. These specific Queries build and refresh multicast membership state of hosts on an attached network. These specific Queries should be sent to each desired hosts with specific multicast address (not the all-hosts/all-nodes multicast address) as their IP destination addresses, because hosts that do not join the multicast session do not pay attention to these specific Queries, and only active member hosts that have been receiving multicast contents with the specified address reply IGMP/MLD reports.

6. Interoperability

IGMPv3 [2] and MLDv2 [3] provide the ability for hosts to report source-specific subscriptions. With IGMPv3/MLDv2, a mobile host can specify a channel of interest, using multicast group and source addresses in its join request. Upon its reception, the upstream router that supports IGMPv3/MLDv2 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) [8].

Recently, the Lightweight-IGMPv3 (LW-IGMPv3) and Lightweight-MLDv2 (LW-MLDv2) [4] protocols have been proposed in the IETF. These protocols provide protocol simplicity for mobile hosts and routers, as they eliminate a complex state machine from the full versions of IGMPv3 and MLDv2, and promote the opportunity to implement SSM in mobile communications.

This document assumes that both multicast routers and mobile hosts MUST be IGMPv3/MLDv2 capable, regardless whether the protocols are the full or lightweight version. And this document does not consider interoperability with older version protocols. The main reason not being interoperate with older IGMP/MLD protocols is that the explicit tracking function does not work properly with older IGMP/MLD protocols.

7. Security Considerations

This document neither provides new functions or modifies the standard functions defined in [2][3][4]. Therefore there is no additional security consideration provided.

8. Acknowledgements

Marshall Eubanks, Gorry Fairhurst, Behcet Sarikaya, Stig Venaas, Jinwei Xia, and others provided many constructive and insightful comments.

9. References

9.1. 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 IGMPv3 and MLDv2 Protocols", RFC 5790, February 2010.

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

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

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

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

[9] Asaeda, H. and Y. Uchida, "IGMP/MLD-Based Explicit Membership Tracking Function for Multicast Routers", draft-asaeda-mboned-explicit-tracking-01.txt (work in progress), October 2010.

9.2. Informative References

[10] Asaeda, H. and T. Schmidt, "IGMP and MLD Protocol Extensions for Mobility", draft-asaeda-multimob-igmp-mld-mobility-extensions-04.txt (work in progress), March 2010.

[11] 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.

[12] Gundavelli, S, Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.

Appendix A. Unicasting General Query

IGMPv3 and MLDv2 specifications [2][3] describe that a host MUST accept and process any Query whose IP Destination Address field contains any of the addresses (unicast or multicast) assigned to the interface on which the Query arrives. In general, the all-hosts multicast address (224.0.0.1) or link-scope all-nodes multicast address (FF02::1) is used as the IP destination address of IGMP/MLD General Query. On the other hand, according to [2][3], a router MAY be able to unicast General Query to tracked member hosts in [Query Interval], if the router keeps track of membership information (Section 3).

Unicasting IGMP/MLD General Query would reduce the drain on battery power of mobile hosts as only the active hosts that have been receiving multicast contents respond the unicast IGMP/MLD General Query messages and non-active hosts do not need to pay attention to the IGMP/MLD messages. This also allows the upstream router to proceed fast leaves (or shorten leave latency) by setting LMQC/LLQC smaller, because the router can immediately converge and update the membership information, ideally.

However, there is a concern in unicast General Query. If a multicast router sends General Query "only" by unicast, it cannot discover potential member hosts whose join requests were lost. Since the hosts do not retransmit the same join requests (i.e., unsolicited Report messages), they loose the chance to join the channels unless the upstream router asks the membership information by sending General Query by multicast. It will be solved by using both unicast and multicast General Queries and configuring the [Query Interval] timer value for multicast General Query and the [Unicast Query Interval] timer value for unicast General Query. However, using two different timers for General Queries would require the protocol extension this document does not focus on. If a router does not distinguish the multicast and unicast General Query Intervals, the router should only use and enable multicast General Query.

Also, unicasting General Query does not remove multicasting General Query. Multicast General Query is necessary to update membership information if it is not correctly synchronized due to missing Reports. Therefore, enabling unicast General Query SHOULD NOT be used for the implementation that does not allow to configure different query interval timers as [Query Interval] and [Unicast Query Interval] (See [10] for the detail). If a router does not distinguish these multicast and unicast General Query Intervals, the router SHOULD only use and enable multicast General Query.

Authors' Addresses

Hitoshi Asaeda
Keio University
Graduate School of Media and Governance
5322 Endo
Fujisawa, Kanagawa 252-0882
Japan

Email: asaeda@wide.ad.jp
URI: http://www.sfc.wide.ad.jp/˜asaeda/

Yogo Uchida
Keio University
Graduate School of Media and Governance
5322 Endo
Fujisawa, Kanagawa 252-0882
Japan

Email: uchida@sfc.wide.ad.jp

Hui Liu
Huawei Technologies Co., Ltd.
Huawei Bld., No.3 Xinxi Rd.
Shang-Di Information Industry Base
Hai-Dian Distinct, Beijing 100085
China

Email: helen.liu@huawei.com

Qin Wu
Huawei Technologies Co., Ltd.
Site B, Floor 12F, Huihong Mansion
No.91 Baixia Rd.
Nanjing, Jiangsu 21001
China

Email: Sunseawq@huawei.com