MULTIMOB Group   H. Asaeda
Internet-Draft   Keio University
Expires: September 9, 2009   T C. Schmidt
                                                          HAW Hamburg
                                                        March 8, 2009

IGMP and MLD Extensions for Mobile Hosts and Routers
draft-asaeda-multimob-igmp-mld-mobility-extensions-02

Status of this Memo

   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/1id-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 9, 2009.

Copyright Notice

Abstract

   This document describes 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. Aside from a modified protocol semantic, optional
   "Notification function" and "Listener Hold function" for the IGMP and
   MLD protocols are introduced.

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) [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 proxy routers [7] that serve multicast member
   hosts 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 (e.g., [12]). Additionally,
   a mobile host may cause initiation and termination of a multicast
   service in the new or the previous network. 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 feasible condition for mobile hosts and routers
   communicating 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". One possible approach to support these requirements is
   explicit tracking. Upstream router traces all downstream members,
   thereby limiting the number of solicited membership reports (by
   periodical multicast 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 explicit tracking function 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
   downstream member hosts speaking IGMPv3 or MLDv2 must send their
   membership reports to an upstream router and the router can keep
   track of their membership status.

   IGMPv3/MLDv2 Query messages are still necessary to maintain

downstream 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.
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. 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.

IGMPv3 and MLDv2 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
with INCLUDE filter mode in its 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) [8].

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
coordination produced by SSM. An upstream router needs to maintain a
shared tree (e.g., RPT in PIM-SM) whenever the router receives join
request with EXCLUDE filter mode from the downstream hosts. This
increases the tree maintenance cost to the multicast routers 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 IGMPv3 and MLDv2 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 downstream hosts' membership status", "varying
IGMP/MLD Query types and values to tune the number of responses", and
"using a source filtering mechanism in a lightweight manner". Aside
from a modified protocol semantic, optional "Notification function"
and "Listener Hold function" for the IGMPv3 and MLDv2 protocols are
introduced. The proposed solutions interoperate with the IGMPv3 and
MLDv2 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.

2.2. IGMP/MLD Query Coordination

2.2.1. Unicasting IGMP/MLD General Query

   IGMP and MLD are non-reliable protocols; to cover the possibility of
   a State-Change Report being missed by one or more multicast routers,
   a host retransmits the same State-Change Report [Robustness Variable]
   - 1 more times, at intervals chosen at random from the range (0,
   [Unsolicited Report Interval]) [2][3]. However, this manner does not
   guarantee that the State-Change Report is reached to the routers.
   The routers therefore need to refresh the downstream membership
   information by receiving Current-State Report periodically solicited
   by IGMP/MLD General Query, in order to be robust in front of host or
   link failures and packet loss. It 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).

   A multicast router periodically transmits IGMP/MLD General Query in
   the [Query Interval] sec. The default value of [Query Interval] is
   125 sec. specified in the standard IGMPv3 and MLDv2 specifications
   [2][3]. Unfortunately, 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. It consumes the bandwidth of a wireless link and may wake all
mobile hosts up by IGMP/MLD General Query. In fact, when mobile
hosts are operating in dormant mode and not communicating with
others, they should 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.

A multicast router attached on a wireless link may want to configure
longer query interval as the [Multicast Query Interval] value
Section 5, in order to reduce the number of IGMP/MLD General Query
messages via multicast. Yet, longer query interval will increase
join latency when an unsolicit Join message with State-Change Record
requesting joining a multicast session is not reached to the router,
or it will increase leave latency when an unsolicit Leave message
with State-Change Record requesting leaving a session is not reached
to the router.

IGMPv3 and MLDv2 specifications [2][3] mention 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, a
router can unicast the message to tracked member hosts in the
[Unicast Query Interval] (described in Section 5), especially when a
multicast router has a small number of mobile hosts that are
listening different multicast sessions. In this case, the router
multicasts IGMP/MLD General Query with longer [Multicast Query
Interval] (described in Section 5) to recognize hosts that were not
tracked.

2.2.2. Multicasting IGMP/MLD Group-Specific Query

In the standard protocols [2][3], an 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. The Group-Specific Query is sent when a router
receives State-Change Records indicating a host is leaving a group,
and never in response to Current-State Records.

A Group-Specific Query builds and refreshes group membership state of
hosts on attached networks. Since a 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 these specific
Queries, and only active group member hosts that have been receiving
multicast contents with the specified address reply IGMP/MLD reports.

However, sending many Group-Specific Queries for all corresponding

groups may increase the total number of transmitted IGMP/MLD
messages. [TODO: Therefore, it is necessary to know the condition in
which a Group-Specific Query is used for maintaining the group
membership state of all hosts on a link, instead of a General Query.]

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. [TODO: Define [Multicast Group-Query
Interval]. We currently think this value is same of the default
[Query Interval] value the regular IGMP and MLD define [2][3].]

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

3.1. IGMP/MLD Notification

   This document proposes an IGMP/MLD Notification operation, in which a
   mobile host *periodically* sends Current-State Record messages
   expressing which multicast sessions the host is joining, even the
   host is not requested to report the membership information by its
   upstream router (i.e., no reception of General Query message).

   The IGMP/MLD Notification operation enables the longer [Multicast
   Query Interval] value for IGMP/MLD General Query than the default
   [Query Interval]. If mobile hosts support the IGMP/MLD Notification
   operation, a multicast router can obtain downstream membership
   information without periodical and spontaneous membership
   solicitation by IGMP/MLD General Query. The router only needs to
   refresh downstream membership information by solicit IGMP/MLD General
   Query to the hosts that do not support the IGMP/MLD Notification
   operation or leave from network without sending any message to the
   router.

   If timers are tuned by dynamic nature of membership, the IGMP/MLD
   Notification operation reduces the number of IGMP/MLD General Query
   periodically sent by a router and the total number of IGMP/MLD
   messages. Since a router only needs to refresh downstream membership
   information by solicit General Query to hosts that do not support the
   Notification operation, both [Unicast Query Interval] and [Multicast
   Query Interval] can be set to longer values. This mechanism may
   conserve battery power of dormant mode hosts, as dormant mode hosts
   do not pay attention to the General Query messages at short
   intervals.

   The IGMP/MLD Notification operation also contributes to fast
   handover, because a host receiving data immediately sends unsolicited
   reports without waiting for IGMP/MLD General Query at the new
   network.

   The [Notification Interval] value (described in Section 5) is the
   interval of Current-State Records periodically sent by a member host
   that joins at least one multicast session. Since a mobile host
   periodically unicasts Current-State Record in [Notification Interval]
   that is shorter than the regular General Query interval (i.e. [Query
   Interval] value) and [Multicast Query Interval] and [Unicast Query
   Interval], even if a router tracking membership status misses State-
   Change report that requests a leave operation, the router can operate
   a leave procedure faster than the regular case. When mobile hosts
   receive IGMP/MLD General Query, they reset their [Notification
   Interval] timer value and restart it.

When a multicast router works with the Notification operation, it
must maintain the following information for each multicast session to
recognize receiver host's membership status;

1   Receiver address - indicates an address of a receiver host
         sending the Current-State Report.

2   Last membership report - indicates the time that the router
         receives the last Current-State Report.

3   Filter mode - indicates either INCLUDE or EXCLUDE as defined
         in [2][3].

4   Source addresses and multicast address - indicates the address
         pair that the receiver joins.

3.1.1. Interoperability of IGMP/MLD Notification

An IGMP/MLD Notification operation is a simple extension for mobile
hosts to spontaneously send IGMP/MLD Current-State Report to their
upstream multicast routers. Since a multicast router solicits
downstream membership information by IGMP/MLD General Query, non-
upgraded mobile hosts can coexist in the network. However, join and
leave latency for non-upgraded mobile hosts may become longer due to
the longer [Query Interval] timer configuration for multicast
routers.

An IGMP/MLD Notification operation does not require amy modification
to multicast routers.

3.2. IGMP/MLD Hold and Release

A mobile host sends a join/leave request for particular multicast
session to its upstream router (i.e., a proxy router or an adjacent
router), and the router will maintain IGMP/MLD state of the host and
serve multicast data to the host. According to a fast handover
scenario (e.g., using FMIPv6 [10]), a mobile host wants to accelerate
multicast service termination in the previous network before handoff
and immediately rejoin the session after the movement. As the
router's behavior, when the router remains part of the multicast
delivery tree, it keeps the session active without leaving from the
session, while the data transmission is paused until the host
completes the movement and resumed after the movement.

This document describes an IGMP/MLD Hold operation that enables to
keep the membership state for fast packet forwarding, and an IGMP/MLD
Release operation that ressumes forwarding the multicast session
after the host movement. Originally, an idea of MLD Hold was

proposed in [13].

Note that discussing the procedure of context transfer (CT) for each
mobility protocol (e.g., FMIPv6, PMIPv6 [11]) and the message format
for CT will be defined in other documents. And also, discussing how
a mobile host or a router detects the movement and triggers to send
an IGMP/MLD Hold message is depend on the mobility protocol, and
hence out of scope of this document.

3.2.1. Action of Mobile Host and Router for IGMP/MLD Hold

An IGMP/MLD Hold operation is used to notify an upstream multicast
router (e.g., HA in case of bidirectional tunneling, adjacent router
in case of remote subscription) to stop forwarding data of the
specified multicast sessions if there is no member joining the same
sessions, but maintain the multicast membership on the router's
interface the message was received.

When a mobile host starts moving from a network to a new network, the
host or the gateway that detects the movement sends an IGMP/MLD Hold
message (mentioned in Section 3.2.3) to its upstream router to make a
fast handover for the specified multicast sessions.

When the multicast router receives an IGMP/MLD Hold message from a
mobile host, it records the host's address and specified group and
source address records and filter-mode, and stops the corresponding
group or source timer until the IGMP/MLD [Hold Interval] timer
(described in Section 5) is expired. The router does not send any
Group-Specific or Group-and-Source Specific Query upon reception of
an IGMP/MLD Hold message.

After the router recognizes the IGMP/MLD Hold message, the router
decides whether it stops forwarding the data or not. The decision is
dependent on whether the router has been keeping track of membership
status. If the router has been tracking all the members, it can
recognize whether the message sender host is the last member joining
the notified session on the interface or not. If the router has not
been tracking, the router waits for receiving the next Current-State
Report and recognizes whether the host is the last member joining the
notified session on the interface or not. And if the message sender
host is the last member, the router stops forwarding data by
disabling the outgoing interface for that session, whereas it keeps
the host's membership state. If the host is not the last member of
the session, the router does not stop forwarding the data even with
reception of this message.

3.2.2. Action of Mobile Host and Router for IGMP/MLD Release

An IGMP/MLD Release operation is used to restart forwarding data of
the specified multicast sessions that are kept with the membership
state at a router, or used to cancel the hold request previously
notified by an IGMP/MLD Hold message.

When a mobile host moves to a new network, the host or the gateway
that detects the movement operates IGMP/MLD Release by sending a
standard IGMP/MLD Current-State Report message as mentioned in
Section 3.2.3) to its upstream router to release the hold request
previously sent.

When a multicast router receives an IGMP/MLD Current-State Report
message, the router examines the message sender and an IGMP/MLD [Hold
Interval] timer (see Section 5). If the router has tracked the
host's membership status, it compares the holding Multicast Address
Records and notified Multicast Address Records specified in the IGMP/
MLD Current-State Report message. Regarding the corresponding
Multicast Address Records, the router will delete the host's
membership state if it has, and restart the given group or source
timer and forwarding the data again if the IGMP/MLD Hold Timer is not
expired.

If either a multicast router has not tracked the corresponding
membership status given by the host that sent the IGMP/MLD Current-
State Report message, or if the IGMP/MLD Hold Timer has expired, then
the router does not take any action.

If a router that did not recognize an IGMP/MLD Hold message (e.g.,
due to packet loss or some troubles in its transmission) receives an
IGMP/MLD Current-State Report message, it will accept the message as
a regular Current-State Report IGMP/MLD message as specified in
Section 3.2.5.

3.2.3. IGMP/MLD Hold Message Format

The following figure is the Report message format defined in IGMPv3
[2] and MLDv2 [3]. Type field is either IGMP Version 3 Membership
Report (0x22) or Version 2 Multicast Listener Report (decimal 143);
it is not necessary to change IGMP/MLD Report message format to
support IGMP/MLD Hold messages.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |    Reserved   |            Checksum           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            Reserved           |Nr of Mcast Address Records (M)|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                 Multicast Address Record [1]                  .
.                                                               .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
.                              .                                .
.                              .                                .
.                              .                                .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
.                                                               .
.                 Multicast Address Record [M]                  .
.                                                               .
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Multicast Address Record shown in the following figure defines
new Record Types that express an IGMP/MLD Hold message. For an IGMP/
MLD Release operation, the standard IGMP/MLD Current-State Report
message is used.

```
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |  Record Type  |  Aux Data Len |     Number of Sources (N)     |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    .                                                               .
    .                       Multicast Address                       .
    .                                                               .
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    .                                                               .
    .                      Source Address [1]                       .
    .                                                               .
    |                                                               |
    +-                                                             -+
    .                              .                                .
    .                              .                                .
    .                              .                                .
    +-                                                             -+
    |                                                               |
    .                                                               .
    .                      Source Address [N]                       .
    .                                                               .
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    .                                                               .
    .                         Auxiliary Data                        .
    .                                                               .
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
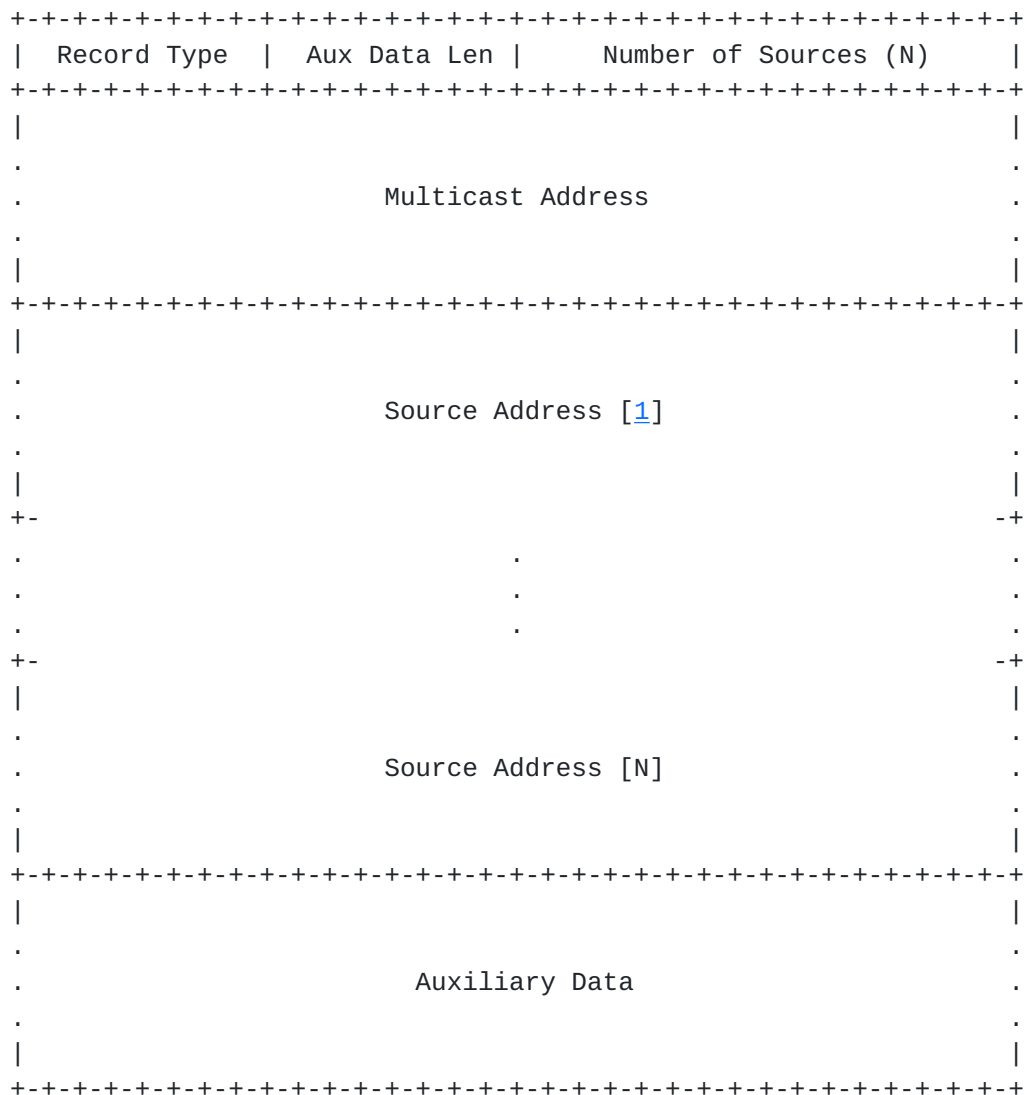
Value   Name and Meaning
-----   ----------------

7   HOLD_INCLUDE - indicates that the interface has a filter mode
        of INCLUDE for the specified multicast address. The Source
        Address [i] fields in this Multicast Address Record contain
        the interface's source list for the specified multicast
        address. A HOLD_INCLUDE Record is never sent with an empty
        source list.

8   HOLD_EXCLUDE - indicates that the interface has a filter mode
        of EXCLUDE for the specified multicast address. The Source
        Address [i] fields in this Multicast Address Record contain
        the interface's source list for the specified multicast
        address, if it is non-empty. When a mobile host works with

LW-IGMPv3/LW-MLDv2 ([Section 3.3](#)), the Multicast Address Record
does not contain the Source Address [i] field with
HOLD_EXCLUDE type value.

[3.2.4](#). **IGMP/MLD Hold Timer**

A router upon receiving an IGMP/MLD Hold message will identify the
corresponding multicast sessions for which the message was received,
and once it accepts the request will wait for the [Hold Interval]
timer period (see [Section 5](#)) for the sessions for allowing the router
on the new link to resume the sessions. However, if it does not
receive any IGMP/MLD Current-State Record message for an IGMP/MLD
Release operation within that given amount of time, it will discard
the hold membership state given by the IGMP/MLD Hold message and
proceed as it receives IGMP Leave or MLD Done request.

[3.2.5](#). **Interoperability of IGMP/MLD Hold and Release**

If a multicast router does not support an IGMP/MLD Hold operation, it
will ignore the IGMP/MLD Hold message. In this case, an IGMP/MLD
Current-State Report message sent by a mobile host that previously
requested an IGMP/MLD Hold operation to stop forwarding data is dealt
with as a new join request for the specified multicast sessions (when
the corresponding group or source timer is expired) or simply updates
the corresponding group and/or source timer (when these timers are
not expired and the membership status has been still active).

[3.3](#). **Lightweight-IGMP/MLD**

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) [[8](#)].

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) [9] 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. Interoperability with Older Protocols

   This document assumes multicast routers that deal with mobile hosts
   MUST be IGMPv3/MLDv2 capable (regardless whether the protocols are
   the full version or not). Therefore this document does not consider
   interoprability with older version protocols.

5. Timers, Counters, and Their Default Values

   A multicast router operating in 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 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 messages on the network; larger values
   cause IGMP Queries to be sent less often.

   [TODO: We will provide the appropriate [Multicast 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 [Query Interval] value the regular IGMPv3
   and MLDv2 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.]

The default [Unicast Query Interval] value is 90 sec.

The default [Multicast Query Interval] value is 180 sec.

The default [Notification Interval] value is 60 sec. The [Notification Interval] value must be shorter than [Multicast Query Interval] and [Unicast Query Interval].

The [Hold Interval] is used for the IGMP/MLD Hold Timer Section 3.2.4 to allow a multicast router to resume the multicast session. The multicast router discards the hold membership state if it does not receive any IGMP/MLD Release (or Current-State Record) message for an IGMP/MLD Release operation within the [Hold Interval].

[TODO: We will provide the appropriate [Hold Interval] value that would fit in the mobile communication environment based on some experimental results. Does it depend on mobile protocols?]

6. **Security Considerations**

   TBD.

7. Acknowledgements

   Dave Thaler, Matthias Waehlisch, and others provided many
   constructive and insightful comments.

8. References

8.1. 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]   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.

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

[9]   Liu, H., Cao, W., and H. Asaeda, "Lightweight IGMPv3 and MLDv2
       Protocols",
       draft-ietf-mboned-lightweight-igmpv3-mldv2-04.txt (work in
       progress), September 2008.

8.2. Informative References

[10]   Koodli, Ed., R., "Fast Handovers for Mobile IPv6", RFC 4068,
       July 2005.

[11]   Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy
       Mobile IPv6", draft-ietf-netlmm-proxymip6-16.txt (work in
       progress), May 2008.

[12]   Schmidt, T., Waehlisch, M., and G. Fairhurst, "Multicast
       Mobility in MIPv6: Problem Statement and Brief Survey",
       draft-irtf-mobopts-mmcastv6-ps-04.txt (work in progress),

          July 2008.

   [13] Jelger, C. and T. Noel, "Multicast for Mobile Hosts in IP
          Networks: Progress and Challenges", IEEE Wireless
          Comm. pp.58-64, October 2002.

Authors' Addresses

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

    Email: asaeda@wide.ad.jp

    Thomas C. Schmidt
    HAW Hamburg
    Dept. Informatik
    Berliner Tor 7
    Hamburg,   D-20099
    Germany

    Email: Schmidt@informatik.haw-hamburg.de