Internet Engineering Task Force (IETF)
Request for Comments: 7411
Updates: 5568
Category: Experimental
ISSN: 2070-1721

T. Schmidt, Ed.
HAW Hamburg
M. Waehlisch
link-lab & FU Berlin
R. Koodli
Intel
G. Fairhurst
University of Aberdeen
D. Liu
China Mobile

November 2014

Multicast Listener Extensions for Mobile IPv6 (MIPv6) and Proxy Mobile IPv6 (PMIPv6) Fast Handovers

Abstract

Fast handover protocols for Mobile IPv6 (MIPv6) and Proxy Mobile IPv6 (PMIPv6) define mobility management procedures that support unicast communication at reduced handover latency. Fast handover base operations do not affect multicast communication and, hence, do not accelerate handover management for native multicast listeners. Many multicast applications like IPTV or conferencing, though, comprise delay-sensitive, real-time traffic and will benefit from fast handover completion. This document specifies extension of the Mobile IPv6 Fast Handovers (FMIPv6) and the Fast Handovers for Proxy Mobile IPv6 (PFMIPv6) protocols to include multicast traffic management in fast handover operations. This multicast support is provided first at the control plane by management of rapid context transfer between access routers and second at the data plane by optional fast traffic forwarding that may include buffering. An FMIPv6 access router indicates support for multicast using an updated Proxy Router Advertisements message format.

This document updates RFC 5568, "Mobile IPv6 Fast Handovers".

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7411.

Copyright Notice

Table of Contents

1. Introduction

Mobile IPv6 [RFC6275] defines a network-layer mobility protocol involving participation by Mobile Nodes, while Proxy Mobile IPv6 [RFC5213] provides a mechanism without requiring mobility protocol operations at a Mobile Node (MN). Both protocols introduce traffic disruptions on handovers that may be intolerable in many real-time application scenarios such as gaming or conferencing. Mobile IPv6 Fast Handovers (FMIPv6) [RFC5568] and Fast Handovers for Proxy Mobile IPv6 (PFMIPv6) [RFC5949] improve the performance of handovers for unicast communication. Delays are reduced to the order of the maximum of the link switching delay and the signaling delay between Access Routers (ARs) or Mobile Access Gateways (MAGs) [FMIPv6-Analysis].

No dedicated treatment of seamless IP multicast [RFC1112] data service has been proposed by any of the above protocols. MIPv6 only roughly defines multicast for Mobile Nodes using a remote subscription approach or a home subscription through bidirectional tunneling via the Home Agent (HA). Multicast forwarding services have not been specified in [RFC5213] but are subject to separate specifications: [RFC6224] and [RFC7287]. It is assumed throughout this document that mechanisms and protocol operations are in place to transport multicast traffic to ARs. These operations are referred to as 'JOIN/LEAVE' of an AR, while the explicit techniques to manage multicast transmission are beyond the scope of this document.

Mobile multicast protocols need to support applications such as IPTV with high-volume content streams and allow distribution to potentially large numbers of receivers. They should thus preserve the multicast nature of packet distribution and approximate optimal routing [RFC5757]. It is undesirable to rely on home tunneling for optimizing multicast. Unencapsulated, native multicast transmission requires establishing forwarding state, which will not be transferred between access routers by the unicast fast handover protocols. Thus, multicast traffic will not experience expedited handover performance, but an MN -- or its corresponding MAG in PMIPv6 -- can perform remote subscriptions in each visited network.

This document specifies extensions to FMIPv6 and PFMIPv6 that include multicast traffic management for fast handover operations in the presence of any-source or source-specific multicast. The protocol extensions were designed under the requirements that

- multicast context transfer shall be transparently included in unicast fast handover operations;

- neither unicast mobility protocols nor multicast routing shall be modified or otherwise affected; and
- no active participation of MNs in PMIPv6 domains is defined.

The solution common to both underlying unicast protocols defines the per-group or per-channel transfer of multicast contexts between ARs or MAGs. The protocol defines corresponding message extensions necessary for carrying (*,G) or (S,G) context information independent of the particular handover protocol. ARs or MAGs are then enabled to treat multicast traffic according to fast unicast handovers and with similar performance. No protocol changes are introduced that prevent a multicast-unaware node from performing fast handovers with multicast-aware ARs or MAGs.

The specified mechanisms apply when a Mobile Node has joined and maintains one or several multicast group subscriptions prior to undergoing a fast handover. It does not introduce any requirements on the multicast routing protocols in use, nor are the ARs or MAGs assumed to be multicast routers. It assumes network conditions, though, that allow native multicast reception in both the previous and new access network. Methods to bridge regions without native multicast connectivity are beyond the scope of this document.

Section 5.1 of this memo updates the Proxy Router Advertisements (PrRtAdv) message format defined in Section 6.1.2 of [RFC5568] to allow an FMIPv6 AR to indicate support for multicast.

1.1. Use Cases and Deployment Scenarios

Multicast extensions for fast handovers enable multicast services in domains that operate either of the unicast fast handover protocols: [RFC5568] or [RFC5949]. Typically, fast handover protocols are activated within an operator network or within a dedicated service installation.

Multicast group communication has a variety of dominant use cases. One traditional application area is infotainment with voluminous multimedia streams delivered to a large number of receivers (e.g., IPTV). Other time-critical services, such as news items or stock-exchange prices, are commonly transmitted via multicast to support fair and fast updates. Both of these use cases may be mobile, and both largely benefit from fast handover operations. Mobile operators may therefore enhance their operational quality or offer premium services by enabling fast handovers.

Another traditional application area for multicast is conversational group communication in scenarios like conferencing or gaming as well as in dedicated collaborative environments or teams. Machine-to-machine communication in the emerging Internet of Things is expected to generate various additional mobile use cases (e.g., among cars). High demands on transmission quality and rapidly moving parties may require fast handovers.

Most of the deployment scenarios above are bound to a fixed infrastructure with consumer equipment at the edge. Today, they are thus likely to follow an operator-centric approach like PFMIPv6. However, Internet technologies evolve for adoption in infrastructureless scenarios, for example, disaster recovery, rescue, crisis prevention, and civil safety. Mobile end-to-end communication in groups is needed in Public Protection and Disaster Relief (PPDR) scenarios, where mobile multicast communication needs to be supported between members of rescue teams, police officers, fire brigade teams, paramedic teams, and command control offices in order to support the protection and health of citizens. These use cases require fast and reliable mobile services that cannot rely on operator infrastructure. They are thus expected to benefit from running multicast with FMIPv6.

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 [RFC2119].

This document uses the terminology for mobility entities in [RFC5568], [RFC5949], [RFC6275], and [RFC5213].

A multicast group is any group (*,G) or (S,G) multicast channel listed in a Multicast Listener Report Message.

3. Protocol Overview

This section provides an informative overview of the protocol mechanisms without normative specifications.

The reference scenario for multicast fast handover is illustrated in Figure 1. A Mobile Node is initially attached to the previous access network (P-AN) via the Previous Access Router (PAR) or Previous Mobile Access Gateway (PMAG) and moves to the new access network (N-AN) connected via a New AR (NAR) or New MAG (NMAG).

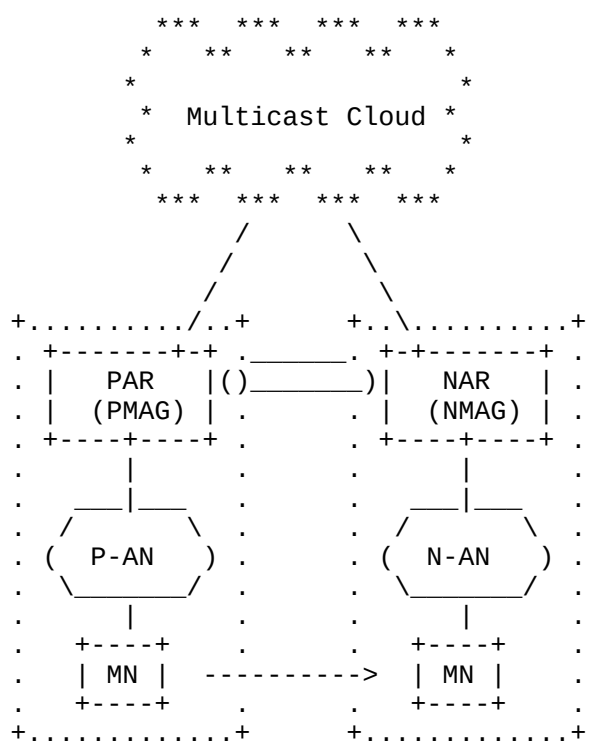

Figure 1: Reference Network for Fast Handover

3.1. Multicast Context Transfer between Access Routers

In a fast handover scenario (see Figure 1), ARs/MAGs establish a mutual binding and provide the capability to exchange context information concerning the MN. This context transfer will be triggered by detecting the forthcoming movement of an MN to a new AR and assists the MN to immediately resume communication on the new subnet using its previous IP address. In contrast to unicast, multicast flow reception does not primarily depend on address and binding cache management but requires distribution trees to adapt so that traffic follows the movement of the MN. This process may be significantly slower than fast handover management [RFC5757]. To accelerate the handover, a multicast listener may offer a twofold advantage of including the multicast groups under subscription in the context transfer. First, the NAR can proactively join the subscribed groups as soon as it gains knowledge of them. Second, multicast flows can be included in traffic forwarding via the tunnel that is established from the PAR to the NAR by the unicast fast handover protocol.

There are two modes of operation in FMIPv6 and in PFMIPv6. The predictive mode allows for AR-binding and context transfer prior to an MN handover, while in the reactive mode, these steps are executed after detection that the MN has reattached to a NAR (NMAG). Details of the signaling schemes differ between FMIPv6 and PFMIPv6 and are outlined in Sections 3.2 and 3.3.

In a predictive fast handover, the access router (i.e., PAR (PMAG) in Figure 1) learns about the impending movement of the MN and simultaneously about the multicast group context as specified in Sections 3.2 and 3.3. Thereafter, the PAR will initiate an AR-binding and context transfer by transmitting a Handover Initiation (HI) message to the NAR (NMAG). The HI message is extended by multicast group states carried in mobility header options, as defined in Section 5.3. On reception of the HI message, the NAR returns a multicast acknowledgement in its Handover Acknowledgement (HAck) answer that indicates its ability to support each requested group (see Section 5.4). The NAR (NMAG) expresses its willingness to receive multicast traffic forwarded by the PAR using standard Multicast Listener Discovery (MLD) signaling for IPv6 or the Internet Group Management Protocol (IGMP) for an IPv4 compatibility case.

Nodes normally create forwarding state for each group requested. There are several reasons why a node may decide not to forward a specific group, e.g., the NAR could already have a native subscription for the group(s) or capacity constraints can hinder decapsulation of additional streams. At the previous network, there may be policy or capacity constraints that make it undesirable to forward the multicast traffic. The PAR can add the tunnel interface obtained from the underlying unicast protocol to its multicast forwarding database for those groups the MN wishes to receive, so that multicast flows can be forwarded in parallel to the unicast traffic.

The NAR implements an MLD proxy [RFC4605] providing host-side behavior towards the upstream PAR. The proxy will submit an MLD report to the upstream tunnel interface to signal the set of groups to be forwarded. It will terminate multicast forwarding from the tunnel when the group is natively received. In parallel, the NAR joins all groups that are not already under subscription using its native multicast upstream interface. While the MN has not arrived at a downstream interface of the NAR, multicast subscriptions on behalf of the MN are associated with a downstream loopback interface. Reception of the Join at the NAR enables downstream native multicast forwarding of the subscribed group(s).

In a reactive fast handover, the PAR will learn about the movement of the MN after the latter has re-associated with the new access network. Also, from the new link, it will be informed about the multicast context of the MN. As group membership information is present at the new access network prior to context transfer, MLD join signaling can proceed in parallel to HI/HAck exchange. Following the context transfer, multicast data can be forwarded to the new access network using the PAR-NAR tunnel of the fast handover protocol. Depending on the specific network topology, multicast traffic for some groups may natively arrive before it is forwarded from the PAR.

In both modes of operation, it is the responsibility of the PAR (PMAG) to properly apply multicast state management when an MN leaves (i.e., to determine whether it can prune the traffic for any unsubscribed group). Depending on the link type and MLD parameter settings, methods for observing the departure of an MN need to be applied (see [RFC5757]). While considering subscriptions of the remaining nodes and from the tunnel interfaces, the PAR uses normal multicast forwarding rules to determine whether multicast traffic can be pruned.

This method allows an MN to participate in multicast group communication with a handover performance that is comparable to unicast handover. It is worth noting that tunnel management between access routers in all modes is inherited from the corresponding unicast fast handover protocols. Tunnels thus remain active until unicast handover operations have been completed for the MN.

3.2. Protocol Operations Specific to FMIPv6

ARs that provide multicast support in FMIPv6 will advertise this general service by setting an indicator bit (’M’ bit) in its PrRtAdv message, as defined in Section 5.1. Additional details about the multicast service support, e.g., flavors and groups, will be exchanged within HI/HAck dialogs later at handover.

An MN operating FMIPv6 will actively initiate the handover management by submitting a Fast Binding Update (FBU). The MN, which is aware of the multicast groups it wishes to maintain, will attach mobility options containing its group states (see Section 5.3) to the FBU and thereby inform ARs about its multicast context. ARs will use these multicast context options for inter-AR context transfer.

In predictive mode, the FBU is issued on the previous link and received by the PAR as displayed in Figure 2. The PAR will extract the multicast context options and append them to its HI message. From the HAck message, the PAR will redistribute the multicast acknowledgement by adding the corresponding mobility options to its

Fast Binding ACK (FBack) message. From receiving the FBack message, the MN will learn about the multicast support for each group in the new access network. If some groups or multicast service models are not supported, it can decide to take actions to overcome a missing service (e.g., by tunneling). Note that the proactive multicast context transfer may proceed successfully, even if the MN misses the FBack message on the previous link.

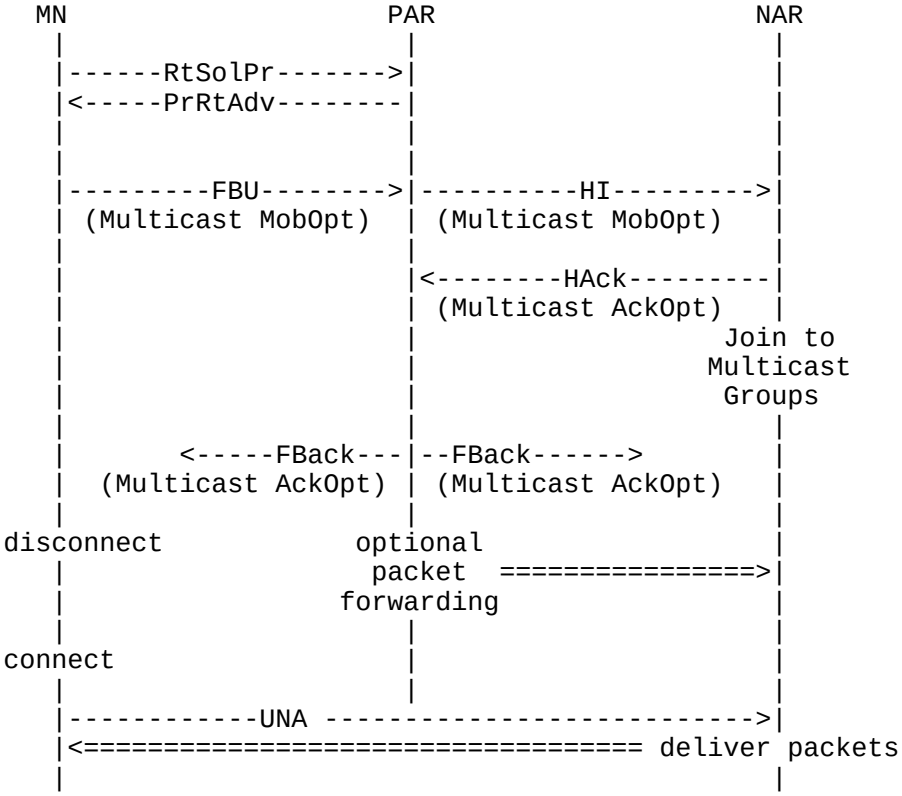

Figure 2: Predictive Multicast Handover for FMIPv6

The flow diagram for reactive mode is depicted in Figure 3. After attaching to the new access link and performing an Unsolicited Neighbor Advertisement (UNA), the MN issues an FBU that the NAR forwards to the PAR without processing. At this time, the MN is able to rejoin all subscribed multicast groups without relying on AR assistance. Nevertheless, multicast context options are exchanged in the HI/HAck dialog to facilitate intermediate forwarding of the requested multicast flows. The multicast traffic could arrive from an MN subscription at the same time that the NAR receives the HI message. Such multicast flows may be transparently excluded from

forwarding by setting an appropriate Multicast Acknowledgement Option. In either case, to avoid duplication, the NAR MUST ensure that not more than one flow of the same group is forwarded to the MN.

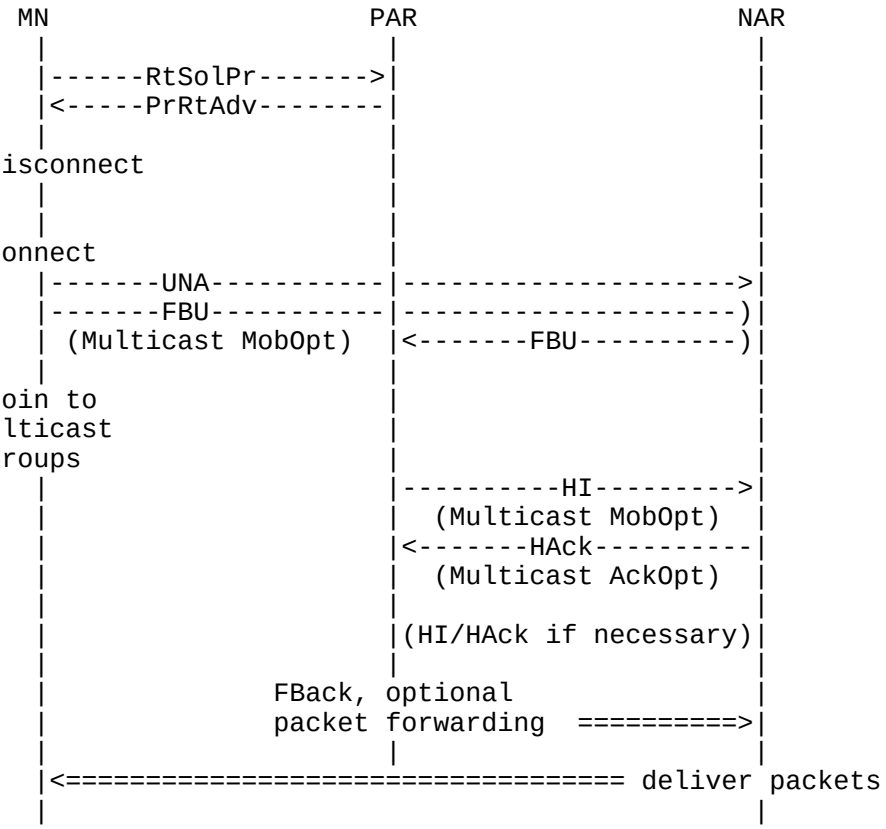

Figure 3: Reactive Multicast Handover for FMIPv6

3.3. Protocol Operations Specific to PFMIPv6

In a proxy mobile IPv6 environment, the MN remains agnostic of network layer changes, and fast handover procedures are operated by the access routers or MAGs to which MNs are connected via node-specific point-to-point links. The handover initiation, or the re-association, is managed by the access networks. Consequently, access routers need to be aware of multicast membership state at the Mobile Node. There are two ways to obtain the multicast membership of an MN.

- MAGs may perform explicit tracking (see [RFC4605] and [RFC6224]) or extract membership status from forwarding states at node-specific links.

- routers can issue a general MLD query at handovers. Both methods are equally applicable. However, a router that does not provide explicit membership tracking needs to query its downstream links after a handover. The MLD membership information then allows the PMAG to learn the multicast group subscriptions of the MN.

In predictive mode, the PMAG will learn about the upcoming movement of the Mobile Node, including its new Access Point Identifier (New AP ID). Without explicit tracking, it will immediately submit a general MLD query and receive MLD reports indicating the multicast address listening state of the subscribed group(s). As displayed in Figure 4, it will initiate binding and context transfer with the NMAG by issuing a HI message that is augmented by multicast contexts in the mobility options defined in Section 5.3. NMAG will extract multicast context information and act as described in Section 3.1.

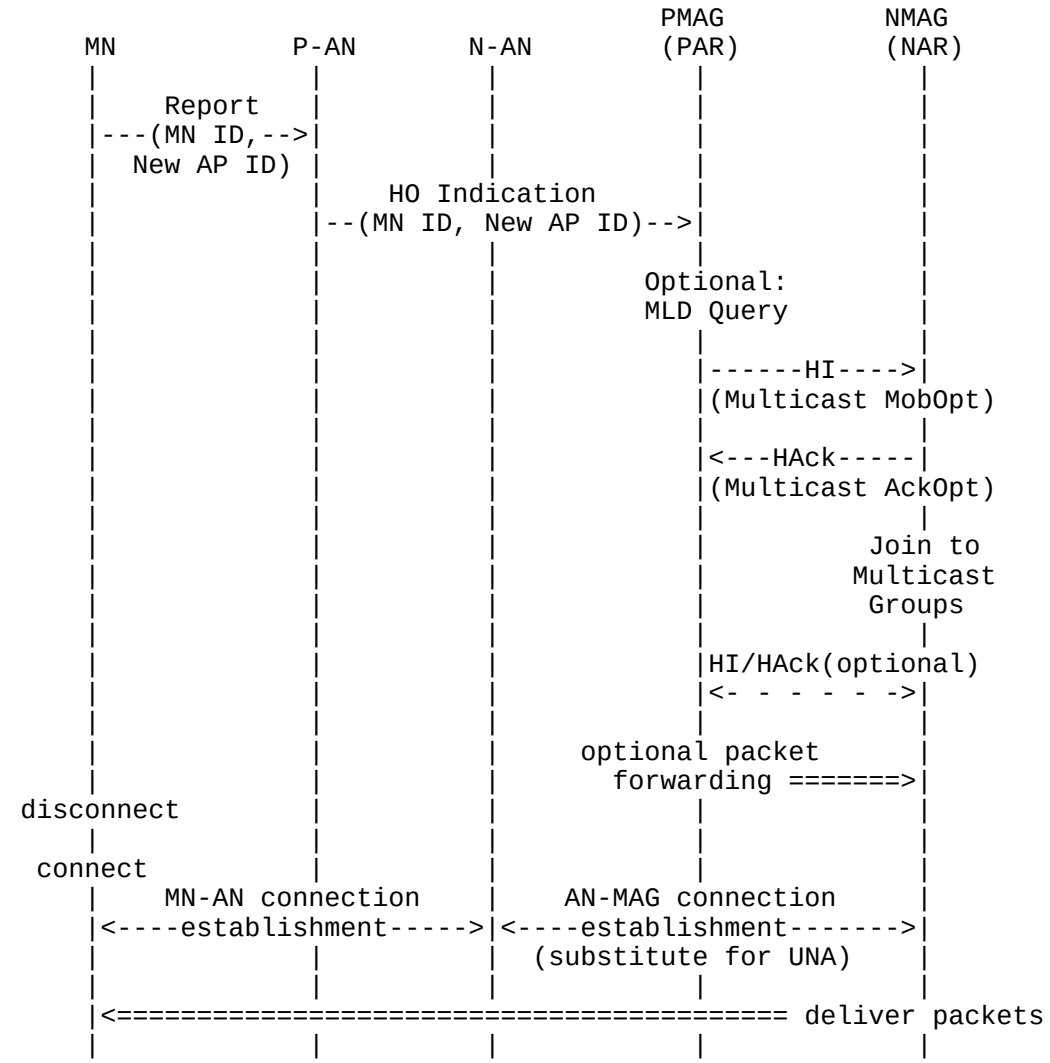

Figure 4: Predictive Multicast Handover for PFMIPv6

In reactive mode, the NMAG will learn the attachment of the MN to the N-AN and establish connectivity using the PMIPv6 protocol operations. However, it will have no knowledge about multicast state at the MN. Triggered by an MN attachment, the NMAG will send a general MLD query and thereafter join the groups for which it receives multicast listener report messages. In the case of a reactive handover, the binding is initiated by the NMAG, and the HI/HAck message semantic is inverted (see [RFC5949]). For multicast context transfer, the NMAG attaches to its HI message those group identifiers it requests to be

forwarded from PMAG. Using the identical syntax in its Multicast Mobility Option headers, as defined in Section 5.4, the PMAG acknowledges the set of requested groups in a HAck answer, indicating the group(s) it is willing to forward. The corresponding call flow is displayed in Figure 5.

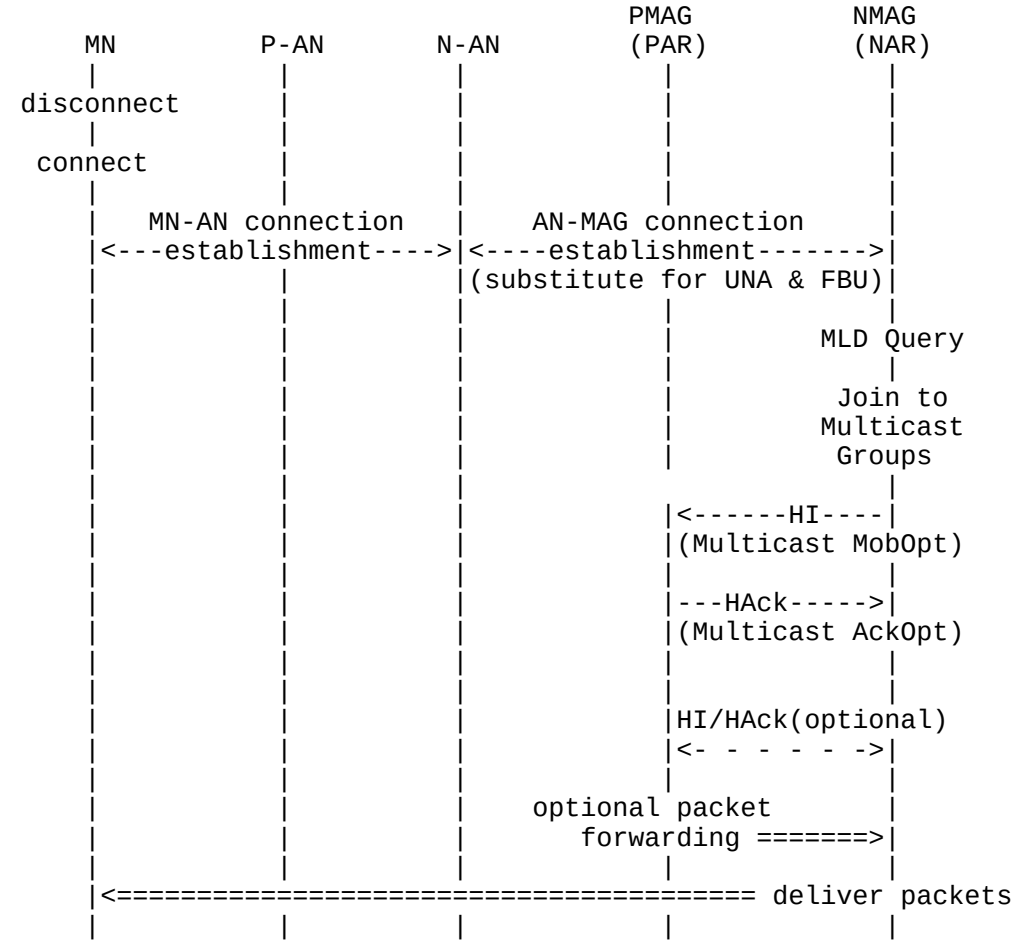

Figure 5: Reactive Multicast Handover for PFMIPv6

4. Protocol Details

This section provides a normative definition of the protocol operations.

4.1. Protocol Operations Specific to FMIPv6

4.1.1. Operations of the Mobile Node

A Mobile Node willing to manage multicast traffic by fast handover operations MUST transfer its MLD listener state records within fast handover negotiations.

When sensing a handover in predictive mode, an MN MUST build a Multicast Mobility Option, as described in Section 5.3, that contains the MLD or IGMP multicast listener state and append it to the Fast Binding Update (FBU) prior to signaling with PAR.

The MN will receive the Multicast Acknowledgement Option(s) as a part of the Fast Binding Acknowledge (FBack) (see Section 5.4) and learn about unsupported or prohibited groups at the NAR. The MN MAY take appropriate actions such as home tunneling to enable reception of groups that are not available via the NAR. Beyond standard FMIPv6 signaling, no multicast-specific operation is required by the MN when reattaching in the new network.

In reactive mode, the MN MUST append the identical Multicast Mobility Option to the FBU sent after its reconnect. In response, it will learn about the Multicast Acknowledgement Option(s) from the FBack and expect corresponding multicast data. Concurrently, it joins all subscribed multicast groups directly on its newly established access link.

4.1.2. Operations of the Previous Access Router

A PAR that supports multicast advertises that support by setting the 'M' bit in the Proxy Router Advertisement (PrRtAdv) message, as specified in Section 5.1 of this document. This indicator exclusively informs the MNs about the capability of the PAR to process and exchange Multicast Mobility Options during fast handover operations.

In predictive mode, a PAR will receive the multicast listener state of an MN prior to handover from the Multicast Mobility Option appended to the FBU. It forwards these records to the NAR within HI messages and will expect Multicast Acknowledgement Option(s) in a HAck, which is itself returned to the MN as an appendix to the FBack. In performing the multicast context exchange, the PAR is instructed

to include the PAR-to-NAR tunnel obtained from unicast handover management in its multicast downstream interfaces and awaits reception of multicast listener report messages from the NAR. In response to receiving multicast subscriptions, the PAR SHOULD forward group data acting as a regular multicast router or proxy. However, the PAR MAY refuse to forward some or all of the multicast flows (e.g., due to administrative configurations or load conditions).

In reactive mode, the PAR will receive the FBU augmented by the Multicast Mobility Option from the new network but continues with an identical multicast record exchange in the HI/HAck dialog. As in the predictive case, it configures the PAR-to-NAR tunnel for the multicast downstream. It then (if capable) forwards data according to the group membership indicated in the multicast listener report messages received from NAR.

In both modes, the PAR MUST interpret the first of the two events -- the departure of the MN or the reception of the Multicast Acknowledgement Option(s) -- as if the MN had sent a multicast LEAVE message and react according to the signaling scheme deployed in the access network (i.e., MLD querying, explicit tracking).

4.1.3. Operations of the New Access Router

A NAR that supports multicast advertises that support by setting the 'M' bit in PrRtAdv as specified in Section 5.1 of this document. This indicator exclusively serves the purpose of informing MNs about the capability of the NAR to process and exchange Multicast Mobility Options during fast handover operations.

In predictive mode, a NAR will receive the multicast listener state of an expected MN from the Multicast Mobility Option appended to the HI message. It will extract the multicast group membership records from the message and match the request subscription with its multicast service offer. Further on, it will join the requested groups using a downstream loopback interface. This will lead to suitable regular subscriptions to a native multicast upstream interface without additional forwarding. Concurrently, the NAR builds a Multicast Acknowledgement Option(s) (see Section 5.4) listing the set of groups that are unsupported on the new access link and returns this list within a HAck. As soon as there is an operational bidirectional tunnel from the PAR to NAR, the NAR joins the groups requested by the MN, which are then forwarded by the PAR using the tunnel link.

In reactive mode, the NAR will learn about the multicast listener state of a new MN from the Multicast Mobility Option appended to each HI message after the MN has already performed local subscriptions of

the multicast service. Thus, the NAR solely determines the intersection of requested and supported groups and issues a join request for each group forwarding this on the PAR-NAR tunnel interface.

In both modes, the NAR MUST send a LEAVE message to the tunnel when it is no longer needed to forward a group, e.g., after arrival of native multicast traffic or termination of a group membership from the MN. Although the message can be delayed, immediately sending the LEAVE message eliminates the need for the PAR and NAR to process traffic that is not to be forwarded.

4.1.4. Buffering Considerations

Multicast packets may be lost during handover. For example, in predictive mode, as illustrated by Figure 2, packets may be lost while the MN is -- already or still -- detached from the networks, even though they are forwarded to the NAR. In reactive mode as illustrated by Figure 3, the situation may be worse, since there will be a delay before joining the multicast group after the MN reattaches to the NAR. Multicast packets cannot be delivered during this time. Buffering the multicast packets at the PAR can reduce multicast packet loss but may then increase resource consumption and delay in packet transmission. Implementors should balance the different requirements in the context of predominant application demands (e.g., real-time requirements or loss sensitivity).

4.2. Protocol Operations Specific to PFMIPv6

4.2.1. Operations of the Mobile Node

A Mobile Node willing to participate in multicast traffic will join, maintain, and leave groups as if located in the fixed Internet. It will cooperate in handover indication as specified in [RFC5949] and required by its access link-layer technology. No multicast-specific mobility actions nor implementations are required at the MN in a PMIPv6 domain.

4.2.2. Operations of the Previous MAG

A MAG receiving a handover indication for one of its MNs follows the same predictive fast handover mode as a PMAG. It MUST issue an MLD General Query immediately on its corresponding link unless it performs explicit membership tracking on that link. After knowledge of the multicast subscriptions of the MN is acquired, the PMAG builds a Multicast Mobility Option, as described in Section 5.3, that contains the MLD and IGMP multicast listener state. If not empty, this Mobility Option is appended to the regular fast handover HI

messages. In the case when a unicast HI message is submitted prior to multicast state detection, the multicast listener state is sent in an additional HI message to the NMAG.

The PMAG then waits until it receives the Multicast Acknowledgement Option(s) with a HAck message (see Section 5.4) and the bidirectional tunnel with the NMAG is created. After the HAck message is received, the PMAG adds the tunnel to its downstream interfaces in the multicast forwarding database. For those groups reported in the Multicast Acknowledgement Option(s), i.e., not supported in the new access network, the PMAG normally takes appropriate actions (e.g., forwarding and termination) according to the network policy. It SHOULD start forwarding multicast traffic down the tunnel interface for the groups indicated in the multicast listener reports received from NMAG. However, it MAY deny forwarding some or all groups included in the multicast listener reports (e.g., due to administrative configurations or load conditions).

After the departure of the MN and on the reception of a LEAVE message, it is RECOMMENDED that the PMAG terminates forwarding of the specified groups and updates its multicast forwarding database. It correspondingly sends a LEAVE message to its upstream link for any group where there are no longer any active listeners on any downstream link.

A MAG receiving a HI message with the Multicast Mobility Option for a currently attached node follows the reactive fast handover mode as a PMAG. It will return a Multicast Acknowledgement Option(s) (see Section 5.4) within a HAck message listing the groups for which it does not provide forwarding support to the NMAG. It will add the bidirectional tunnel with NMAG to its downstream interfaces and will start forwarding multicast traffic for the groups listed in the multicast listener report messages from the NMAG. On reception of a LEAVE message for a group, the PMAG terminates forwarding for the specific group and updates its multicast forwarding database. According to its multicast forwarding state, it sends a LEAVE message to its upstream link for any group where there are no longer any active listeners on any downstream link.

In both modes, the PMAG will interpret the departure of the MN as a multicast LEAVE message of the MN and react according to the signaling scheme deployed in the access network (i.e., MLD querying and explicit tracking).

4.2.3. Operations of the New MAG

A MAG receiving a HI message with a Multicast Mobility Option for a currently unattached node follows the same predictive fast handover mode as an NMAG. It will decide the multicast groups to be forwarded from the PMAG and build a Multicast Acknowledgement Option (see Section 5.4) that enumerates only unwanted groups. This Mobility Option is appended to the regular fast handover HAck messages or, in the case of a unicast HAck message being submitted prior to multicast state acknowledgement, sent in an additional HAck message to the PMAG. Immediately thereafter, the NMAG SHOULD update its MLD membership state based on the membership reported in the Multicast Mobility Option. Until the MN reattaches, the NMAG uses its Loopback interface for downstream and MUST NOT forward traffic to the potential link of the MN. The NMAG SHOULD issue JOIN messages for those newly selected groups to its regular multicast upstream interface. As soon as the bidirectional tunnel with PMAG is established, the NMAG additionally joins those groups on the tunnel interface requested to be forwarded from the PMAG.

A MAG experiencing a connection request for an MN without prior reception of a corresponding Multicast Mobility Option is operating in the reactive fast handover mode as an NMAG. Following the reattachment, it SHOULD immediately issue an MLD General Query to learn about multicast subscriptions of the newly arrived MN. Using standard multicast operations, the NMAG joins groups not currently forwarded using its regular multicast upstream interface. Concurrently, it selects groups for forwarding from PMAG and builds a Multicast Mobility Option, as described in Section 5.3, that contains the multicast listener state. If not empty, this Mobility Option is appended to the regular fast handover HI messages with the F flag set or, in the case of unicast HI message being submitted prior to multicast state detection, sent in an additional HI message to the PMAG. Upon reception of the Multicast Acknowledgement Option and establishment of the bidirectional tunnel, the NMAG additionally joins the set of groups on the tunnel interface that it wishes to receive by forwarding from the PMAG. When multicast flows arrive, the NMAG forwards data to the appropriate downlink(s).

In both modes, the NMAG MUST send a LEAVE message to the tunnel when forwarding of a group is no longer needed, e.g., after native multicast traffic arrives or group membership of the MN terminates. Although the message can be delayed, immediately sending the LEAVE message eliminates the need for PAR and NAR to process traffic that is not to be forwarded.

4.2.4. IPv4 Support Considerations

An MN in a PMIPv6 domain MAY use an IPv4 address transparently for communication, as specified in [RFC5844]. For this purpose, Local Mobility Anchors (LMAs) can register IPv4-Proxy-CoAs in its binding caches, and MAGs can provide IPv4 support in access networks. Correspondingly, multicast membership management will be performed by the MN using IGMP. For multiprotocol multicast support on the network side, IGMPv3 router functions are required at both MAGs (see Section 5.6 for compatibility considerations with previous IGMP versions). Context transfer between MAGs can transparently proceed in the HI/HAck message exchanges by encapsulating IGMP multicast state records within Multicast Mobility Options (see Sections 5.3 and 5.4 for details on message formats).

The deployment of IPv4 multicast support SHOULD be homogeneous across a PMIP domain. This avoids multicast service breaks during handovers.

It is worth mentioning the scenarios of a dual-stack IPv4/IPv6 access network and the use of Generic Routing Encapsulation (GRE) tunneling as specified in [RFC5845]. Corresponding implications and operations are discussed in the PMIP Multicast Base Deployment document (see [RFC6224]).

5. Message Formats

5.1. Multicast Indicator for Proxy Router Advertisement (PrRtAdv)

This document updates the Proxy Router Advertisements (PrRtAdv) message format defined in Section 6.1.2 of [RFC5568]. The update assigns the first bit of the Reserved field to carry the 'M' bit, as defined in Figure 6. An FMIPv6 AR indicates support for multicast by setting the 'M' bit to a value of 1.

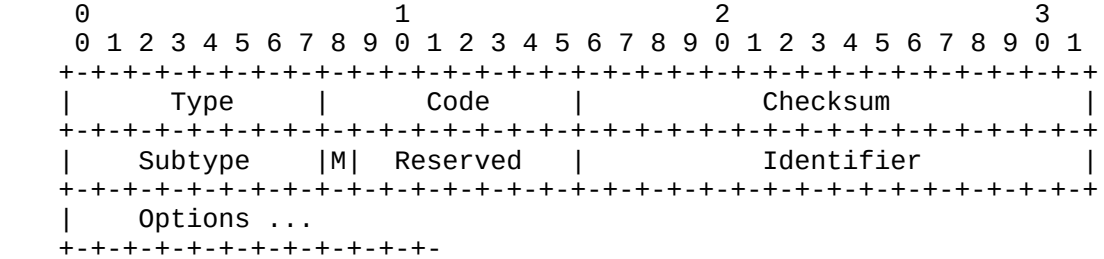

Figure 6: Multicast Indicator Bit for Proxy Router Advertisement (PrRtAdv) Message

This document updates the Reserved field to include the 'M' bit. It is specified as follows.

M = 1 indicates that the specifications of this document apply.

M = 0 indicates that the behavior during fast handover proceeds according to [RFC5568].

The default value (0) of this bit indicates a non-multicast-capable service.

5.2. Extensions to Existing Mobility Header Messages

The fast handover protocols use an IPv6 header type called Mobility Header, as defined in [RFC6275]. Mobility Headers can carry variable Mobility Options.

The multicast listener context of an MN is transferred in fast handover operations from PAR/PMAG to NAR/NMAG within a new Multicast Mobility Option and MUST be acknowledged by a corresponding Multicast Acknowledgement Option. Depending on the specific handover scenario and protocol in use, the corresponding option is included within the mobility option list of HI/HAck only (PFMIPv6) or of FBU/FBack/HI/HAck (FMIPv6).

5.3. New Multicast Mobility Option

This section defines the Multicast Mobility Option. It contains the current listener state record of the MN obtained from the MLD Multicast Listener Report message and has the format displayed in Figure 7.

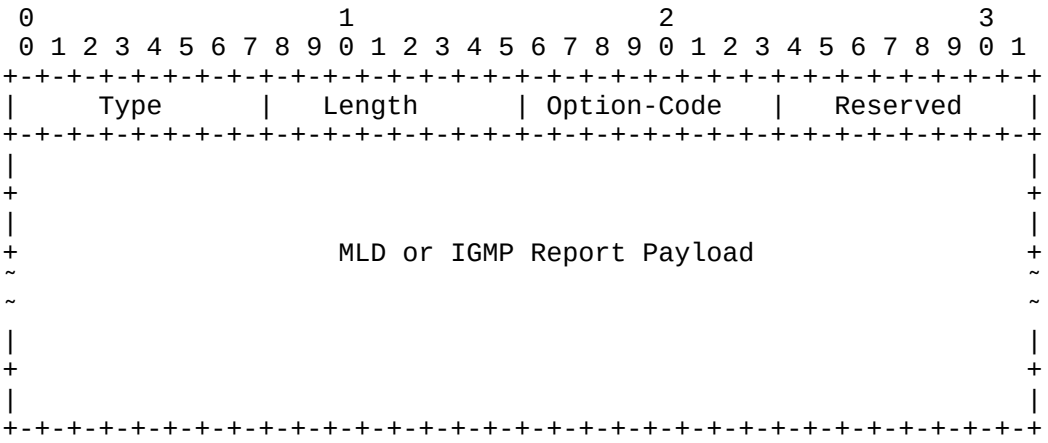

Figure 7: Mobility Header Multicast Option

Type: 60

Length: 8-bit unsigned integer. The length of this option in 32-bit words, not including the Type, Length, Option-Code, and Reserved fields.

Option-Code:

1: IGMPv3 Payload Type

2: MLDv2 Payload Type

3: IGMPv3 Payload Type from IGMPv2 Compatibility Mode

4: MLDv2 Payload Type from MLDv1 Compatibility Mode

Reserved: MUST be set to zero by the sender and MUST be ignored by the receiver.

MLD or IGMP Report Payload: This field is composed of the Membership Report message after stripping its ICMP header. This Report Payload always contains an integer number of multicast records. Corresponding message formats are defined for MLDv2 in [RFC3810] and for IGMPv3 in [RFC3376]. This field MUST always contain the first header line (Reserved field and No of Mcast Address Records).

Figure 8 shows the Report Payload for MLDv2 (see Section 5.2 of [RFC3810] for the definition of Multicast Address Records). When IGMPv3 is used, the payload format is defined according to IGMPv3 Group Records (see Section 4.2 of [RFC3376] for the definition of Group Records).

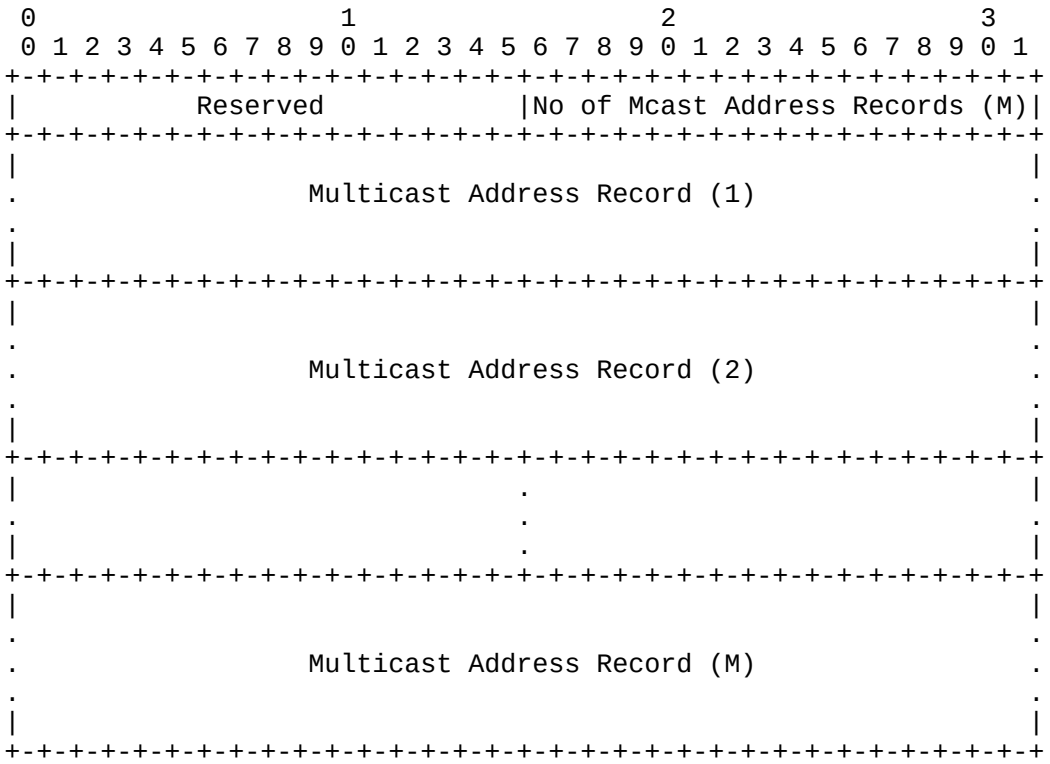

Figure 8: MLDv2 Report Payload

5.4. New Multicast Acknowledgement Option

The Multicast Acknowledgement Option reports the status of the context transfer and contains the list of state records that could not be successfully transferred to the next access network. It has the format displayed in Figure 9.

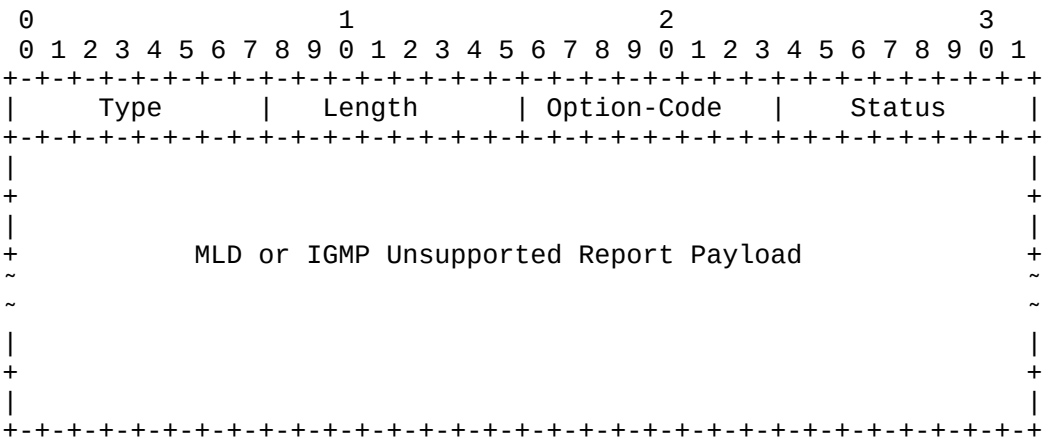

Figure 9: Mobility Header Multicast Acknowledgement Option

Type: 61

Length: 8-bit unsigned integer. The length of this option in 32-bit words, not including the Type, Length, Option-Code, and Status fields.

Option-Code: 0

Status:

1: Report Payload type unsupported

2: Requested group service unsupported

3: Requested group service administratively prohibited

MLD or IGMP Unsupported Report Payload: This field is syntactically identical to the MLD and IGMP Report Payload field described in Section 5.3 but is only composed of those Multicast Address Records that are not supported or prohibited in the new access network. This field MUST always contain the first header line (Reserved field and No of Mcast Address Records) but MUST NOT contain any Mcast Address Records if the status code equals 1.

Note that group subscriptions to specific sources may be rejected at the destination network; thus, the composition of multicast address records may differ from initial requests within an MLD or IGMP Report Payload option.

5.5. Length Considerations: Number of Records and Addresses

Mobility Header messages exchanged in HI/HAck and FBU/FBack dialogs impose length restrictions on multicast context records due to the 8-bit Length field. The maximal payload length available in FBU/FBack messages is 4 octets (Mobility Option header line) + 1024 octets (MLD Report Payload). For example, not more than 51 Multicast Address Records of minimal length (without source states) may be exchanged in one message pair. In typical handover scenarios, this number reduces further according to unicast context and Binding Authorization data. A larger number of MLD reports that exceeds the available payload size MAY be sent within multiple HI/HAck or FBU/FBack message pairs. In PFMIPv6, context information can be fragmented over several HI/HAck messages. However, a single MLDv2 Report Payload MUST NOT be fragmented. Hence, for a single Multicast Address Record, the number of source addresses (S,.) is limited to 62.

5.6. MLD and IGMP Compatibility Requirements

Access routers (MAGs) MUST support MLDv2 and IGMPv3. To enable multicast service for MLDv1 and IGMPv2 listeners, the routers MUST follow the interoperability rules defined in [RFC3810] and [RFC3376] and appropriately set the Multicast Address Compatibility Mode.

When the Multicast Address Compatibility Mode is MLDv1 or IGMPv2, a router internally translates the subsequent MLDv1 and IGMPv2 messages for that multicast address to their MLDv2 and IGMPv3 equivalents and uses these messages in the context transfer. The current state of Compatibility Mode is translated into the code of the Multicast Mobility Option, as defined in Section 5.3. A NAR (NMAG) receiving a Multicast Mobility Option during handover will switch to the lowest level of MLD and IGMP Compatibility Mode that it learned from its previous and new option values. This minimal compatibility agreement is used to allow for continued operation.

6. Security Considerations

Security vulnerabilities that exceed issues discussed in the base protocols mentioned in this document ([RFC5568], [RFC5949], [RFC3810], and [RFC3376]) are identified as follows.

Multicast context transfer at predictive handovers implements group states at remote access routers and may lead to group subscriptions without further validation of the multicast service requests. Thereby, a NAR (NMAG) is requested to cooperate in potentially complex multicast rerouting and may receive large volumes of traffic. Malicious or inadvertent multicast context transfers may result in a significant burden of route establishment and traffic management onto the backbone infrastructure and the access router itself. Rapid rerouting or traffic overload can be mitigated by a rate control at the AR that restricts the frequency of traffic redirects and the total number of subscriptions. In addition, the wireless access network remains protected from multicast data injection until the requesting MN attaches to the new location.

7. IANA Considerations

This document defines two new mobility options that have been allocated from the "Mobility Options" registry at <http://www.iana.org/assignments/mobility-parameters>:

60 Multicast Mobility Option, described in Section 5.3

61 Multicast Acknowledgement Option, described in Section 5.4

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011, <http://www.rfc-editor.org/info/rfc6275>.

[RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008, <http://www.rfc-editor.org/info/rfc5213>.

[RFC5568] Koodli, R., "Mobile IPv6 Fast Handovers", RFC 5568, July 2009, <http://www.rfc-editor.org/info/rfc5568>.

[RFC5949] Yokota, H., Chowdhury, K., Koodli, R., Patil, B., and F. Xia, "Fast Handovers for Proxy Mobile IPv6", RFC 5949, September 2010, <http://www.rfc-editor.org/info/rfc5949>.

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989, <http://www.rfc-editor.org/info/rfc1112>.

[RFC4605] 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, <http://www.rfc-editor.org/info/rfc4605>.

[RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004, <http://www.rfc-editor.org/info/rfc3810>.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002, <http://www.rfc-editor.org/info/rfc3376>.

8.2. Informative References

[RFC5757] Schmidt, T., Waehlisch, M., and G. Fairhurst, "Multicast Mobility in Mobile IP Version 6 (MIPv6): Problem Statement and Brief Survey", RFC 5757, February 2010, <http://www.rfc-editor.org/info/rfc5757>.

[FMCAST-MIP6] Suh, K., Kwon, D., Suh, Y., and Y. Park, "Fast Multicast Protocol for Mobile IPv6 in the fast handovers environments", Work in Progress, draft-suh-mipshop-fmcast-mip6-00, February 2004.

[FMIPv6-Analysis] Schmidt, T. and M. Waehlisch, "Predictive versus Reactive -- Analysis of Handover Performance and Its Implications on IPv6 and Multicast Mobility", Telecommunication Systems, Vol. 30, No. 1-3, pp. 123-142, November 2005, <http://dx.doi.org/10.1007/s11235-005-4321-4>.

[RFC6224] Schmidt, T., Waehlisch, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011, <http://www.rfc-editor.org/info/rfc6224>.

[RFC7287] Schmidt, T., Gao, S., Zhang, H., and M. Waehlisch, "Mobile Multicast Sender Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 7287, June 2014, <http://www.rfc-editor.org/info/rfc7287>.

[RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010, <http://www.rfc-editor.org/info/rfc5844>.

[RFC5845] Muhanna, A., Khalil, M., Gundavelli, S., and K. Leung, "Generic Routing Encapsulation (GRE) Key Option for Proxy Mobile IPv6", RFC 5845, June 2010, <http://www.rfc-editor.org/info/rfc5845>.

Appendix A. Considerations for Mobile Multicast Sources

This document only specifies protocol operations for fast handovers for mobile listeners. In this appendix, we briefly discuss aspects of supporting mobile multicast sources.

In a multicast-enabled Proxy Mobile IPv6 domain, multicast sender support is likely to be enabled by any one of the mechanisms described in [RFC7287]. In this case, multicast data packets from an MN are transparently forwarded either to its associated LMA or to a multicast-enabled access network. In all cases, a mobile source can continue to transmit multicast packets after a handover from PMAG to NMAG without additional management operations. Packets (with a persistent source address) will continue to flow via the LMA or the access network into the previously established distribution system.

In contrast, an MN will change its Care-of Address while performing FMIPv6 handovers. Even though MNs are enabled to send packets via the reverse NAR-PAR tunnel using their previous Care-of Address for a limited time, multicast sender support in such a Mobile IPv6 regime will most likely follow one of the basic mechanisms described in Section 5.1 of [RFC5757]: (1) bidirectional tunneling, (2) remote subscription, or (3) agent-based solutions. A solution for multicast senders that is homogeneously deployed throughout the mobile access network can support seamless services during fast handovers, the details of which are beyond the scope of this document.

Acknowledgments

Protocol extensions to support multicast in Fast Mobile IPv6 have been loosely discussed for several years. Repeated attempts have been made to define corresponding protocol extensions. The first version [FMCAST-MIP6] was presented by Kyungjoo Suh, Dong-Hee Kwon, Young-Joo Suh, and Youngjun Park in 2004.

This work was stimulated by many fruitful discussions in the MobOpts research group. We would like to thank all active members for constructive thoughts and contributions on the subject of multicast mobility. The MULTIMOB working group has provided continuous feedback during the evolution of this work. Comments, discussions, and reviewing remarks have been contributed by (in alphabetical order) Carlos J. Bernardos, Luis M. Contreras, Hui Deng, Shuai Gao, Brian Haberman, Dirk von Hugo, Min Hui, Georgios Karagian, Marco Liebsch, Behcet Sarikaya, Stig Venaas, and Juan Carlos Zuniga.

Funding has been provided by the German Federal Ministry of Education and Research within the projects Mindstone, SKIMS, and SAFEST. This is gratefully acknowledged.

Authors' Addresses

Thomas C. Schmidt (editor)
HAW Hamburg
Dept. Informatik
Berliner Tor 7
Hamburg D-20099
Germany

EMail: t.schmidt@haw-hamburg.de

Matthias Waehlisch
link-lab & FU Berlin
Hoenower Str. 35
Berlin D-10318
Germany

EMail: mw@link-lab.net

Rajeev Koodli
Intel
3600 Juliette Lane
Santa Clara, CA 95054
United States

EMail: rajeev.koodli@intel.com

Godred Fairhurst
University of Aberdeen
School of Engineering
Aberdeen AB24 3UE
United Kingdom

EMail: gorry@erg.abdn.ac.uk

Dapeng Liu
China Mobile

Phone: +86-123-456-7890
EMail: liudapeng@chinamobile.com