

MULTIMOB Group
Internet-Draft
Intended status: Experimental
Expires: August 30, 2013

T C. Schmidt, Ed.
HAW Hamburg
M. Waehlich
link-lab & FU Berlin
R. Koodli
Cisco Systems
G. Fairhurst
University of Aberdeen
Dapeng. Liu
China Mobile
February 26, 2013

Multicast Listener Extensions for MIPv6 and PMIPv6 Fast Handovers
draft-ietf-multimob-fmipv6-pfmipv6-multicast-01

Abstract

Fast handover protocols for MIPv6 and 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, are comprised of delay-sensitive real-time traffic and will benefit from fast handover execution. 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 a management of rapid context transfer between access routers, second at the data plane by an optional fast traffic forwarding that may include buffering.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 30, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
2.	Terminology	5
3.	Protocol Overview	5
3.1.	Multicast Context Transfer between Access Routers	6
3.2.	Protocol Operations Specific to FMIPv6	8
3.3.	Protocol Operations Specific to PFMIPv6	10
4.	Protocol Details	13
4.1.	Protocol Operations Specific to FMIPv6	13
4.1.1.	Operations of the Mobile Node	13
4.1.2.	Operations of the Previous Access Router	14
4.1.3.	Operations of the New Access Router	15
4.1.4.	Buffering Considerations	15
4.2.	Protocol Operations Specific to PFMIPv6	16
4.2.1.	Operations of the Mobile Node	16
4.2.2.	Operations of the Previous MAG	16
4.2.3.	Operations of the New MAG	17
4.2.4.	IPv4 Support Considerations	18
5.	Message Formats	18
5.1.	Multicast Indicator for Proxy Router Advertisement (PrRtAdv)	18
5.2.	Extensions to Existing Mobility Header Messages	19
5.3.	New Multicast Mobility Option	19
5.4.	New Multicast Acknowledgement Option	21
5.5.	Length Considerations: Number of Records and Addresses	22
5.6.	MLD (IGMP) Compatibility Aspects	22
6.	Security Considerations	22
7.	IANA Considerations	23
8.	Acknowledgments	23
9.	References	24
9.1.	Normative References	24
9.2.	Informative References	24
	Appendix A. Change Log	25
	Authors' Addresses	26

1. Introduction

Mobile IPv6 [RFC3775] 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 these handover delays for unicast communication to the order of the maximum delay needed for link switching and signaling between Access Routers (ARs) or Mobile Access Gateways (MAGs) [FMIPv6-Analysis].

No dedicated treatment of seamless multicast data reception 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 bi-directional tunneling via the Home Agent (HA). Multicast forwarding services have not been specified at all in [RFC5213], but are subject to current specification [RFC6224]. 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 serve applications such as IPTV with high-volume content streams to be distributed to potentially large numbers of receivers, and therefore should 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. The solution common to both underlying protocols defines the per-group transfer of multicast contexts between ARs or MAGs. The protocol defines corresponding message extensions necessary for carrying group 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.

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]. The use of the term, "silently ignore" is not defined in RFC 2119. However, the term is used in this document and can be similarly construed.

This document uses the terminology of [RFC5568], [RFC5949], [RFC3775], and [RFC5213]. In addition, the following terms are introduced:

3. Protocol Overview

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

The reference scenario for multicast fast handover is illustrated in Figure 1.

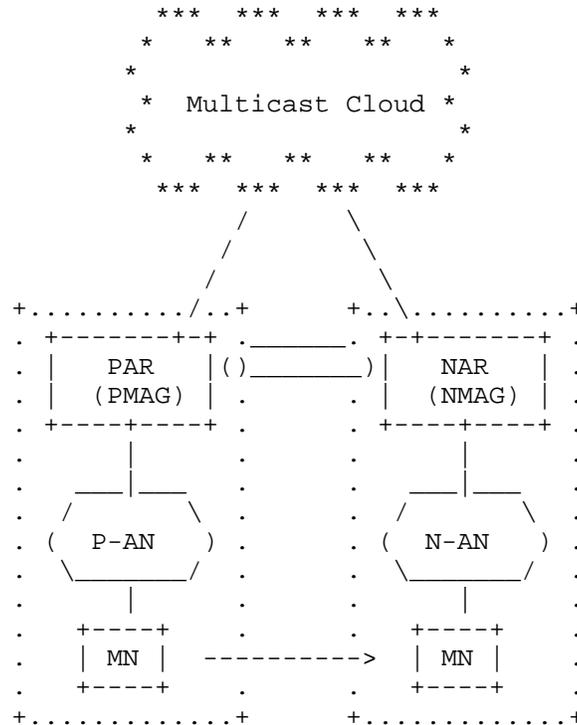


Figure 1: Reference Network for Fast Handover

3.1. Multicast Context Transfer between Access Routers

In a fast handover scenario (cf. 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 assist the MN to immediately resume communication on the new subnet link 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]. Multicast listeners at handover may offer the twofold advantage of including the multicast groups under subscription in 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 established from PAR to NAR.

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 re-attached to NAR. Details of the signaling schemes differ between FMIPv6 and PFMIPv6 and are outlined in Section 3.2 and Section 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 Section 3.2 and Section 3.3. Thereafter, the PAR will initiate an AR-binding and context transfer by transmitting a HI message to NAR (NMAG). HI is extended by multicast group states carried in mobility header options as defined in Section 5.3. On reception of the HI message, NAR returns a multicast acknowledgement in its HACK answer that indicates its ability to support each requested group (see Section 5.4). NAR (NMAG) expresses its willingness to receive multicast traffic from forwarding by PAR using standard MLD signaling. There are several reasons to waive forwarding, 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 of capacity constraints that make it undesirable to forward the multicast traffic. The PAR can add the tunnel interface 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 behaviour on behalf of the upstream PAR. The proxy will submit an MLD report to the upstream tunnel interface to indicate the set of groups to be forwarded. It will terminate multicast forwarding from the tunnel when the group is natively received. In parallel, 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 Loopback as a downstream 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 are 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 though, multicast traffic for some groups may natively arrive before it is forwarded from PAR.

In both modes of operation, it is the responsibility of the PAR (PMAG) to properly apply multicast state management when an MN leaves. Depending on the link type and MLD parameter settings, methods for observing the departure of an MN need to be applied (cf., [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.

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 handovers.

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, FBU is issued on the previous link and received by PAR as displayed in Figure 2. PAR will extract the multicast context options and append them to its HI message. From the HACK message, PAR will redistribute the multicast acknowledgement by adding the corresponding mobility options to its FBACK message. From receiving FBACK, the MN will learn about a per group multicast support in the new access network. If some groups or a multicast flavour are not supported, it MAY decide on taking actions to compensate the missing service. Note that the proactive multicast context transfer may proceed successfully, even if the MN misses the FBACK message on the previous link.

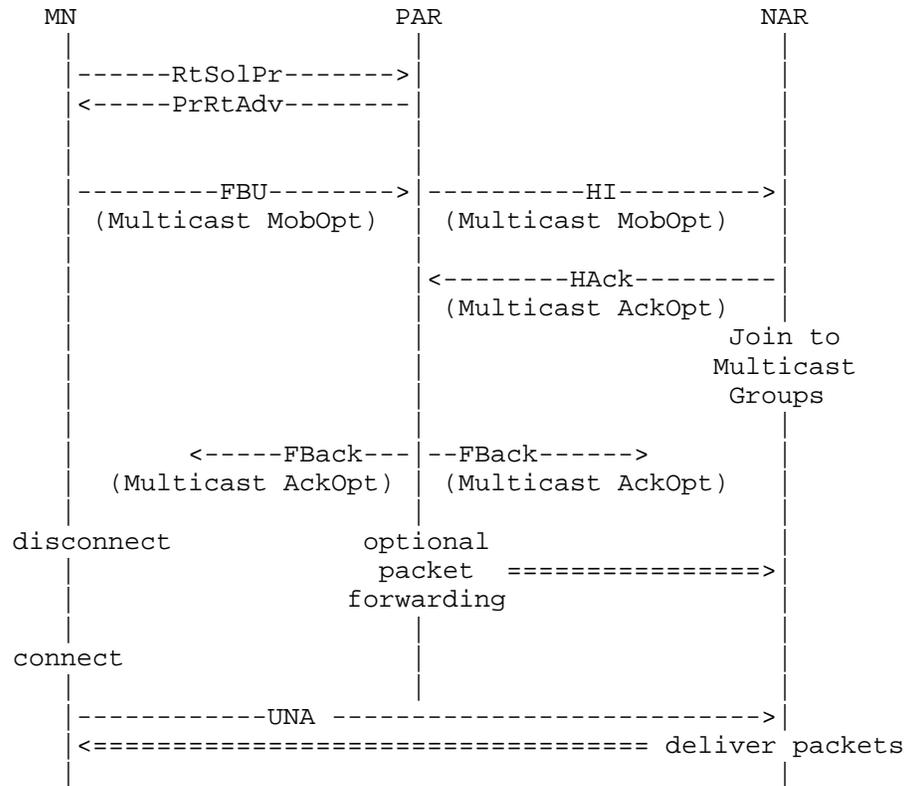


Figure 2: Predictive Multicast Handover for FMIPv6

The call flow for reactive mode is visualized in Figure 3. After attaching to the new access link and performing an unsolicited neighbor advertisement (UNA), the MN issues an FBU which NAR forwards to PAR without processing. At this time, the MN is able to re-join 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 requested flows. Note that group traffic possibly already arrives from a MN's subscription at the time NAR receives the HI message. Such multicast flows may be transparently excluded from forwarding by setting an appropriate multicast acknowledge option. In any case, NAR MUST ensure that not more than one flow of the same group is forwarded to the MN.

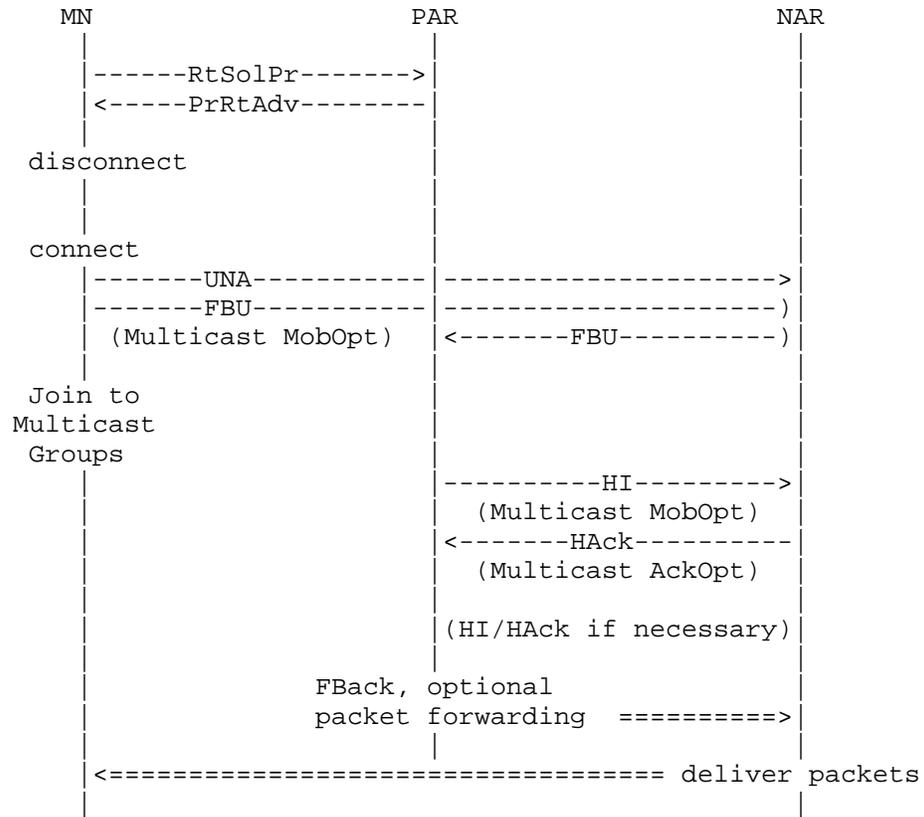


Figure 3: Reactive Multicast Handover for FMIPv6

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. The handover initiation, or the re-association respectively are 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 record of MN's multicast membership. First, MAGs MAY perform an explicit tracking (cf., [RFC4605], [RFC6224]) or extract membership status from forwarding states at node-specific point-to-point links. Second, routers can perform general queries at handovers. Both methods are equally applicable. However, a router that does not operate explicit tracking MUST query its downstream links subsequent to handovers. In either case, the PAR will become knowledgeable about multicast group subscriptions of the MN.

In predictive mode, the PMAG (PAR) will learn about the upcoming movement of the mobile node. Without explicit tracking, it will immediately submit a general MLD query and learn about the multicast groups under subscription. As displayed in Figure 4, it will initiate binding and context transfer with the NMAG (NAR) by issuing a HI message that is augmented by multicast contexts in the mobility options defined in Section 5.3. NAR will extract multicast context information and act as described in Section 3.1.

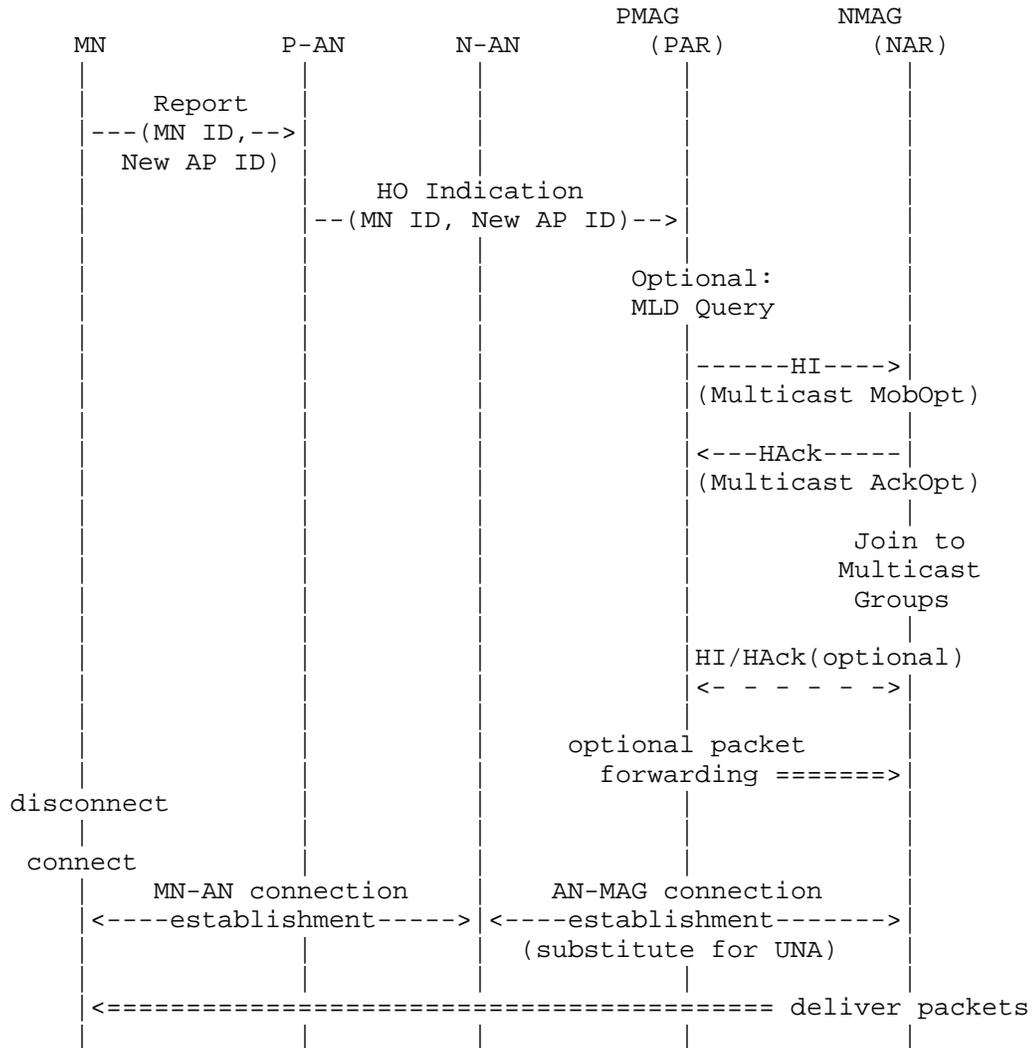


Figure 4: Predictive Multicast Handover for PFMIPv6

In reactive mode, the NMAG (NAR) will learn about MN's attachment to the N-AN and establish connectivity by means of PMIPv6 protocol operations. However, it will have no knowledge about multicast state at the MN. Triggered by a MN attachment, the NMAG will send a general MLD query and thereafter join the requested groups. In the case of a reactive handover, the binding is initiated by 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, PMAG acknowledges those requested groups in its HACK answer that it is willing to forward. The corresponding call flow is displayed in Figure 5.

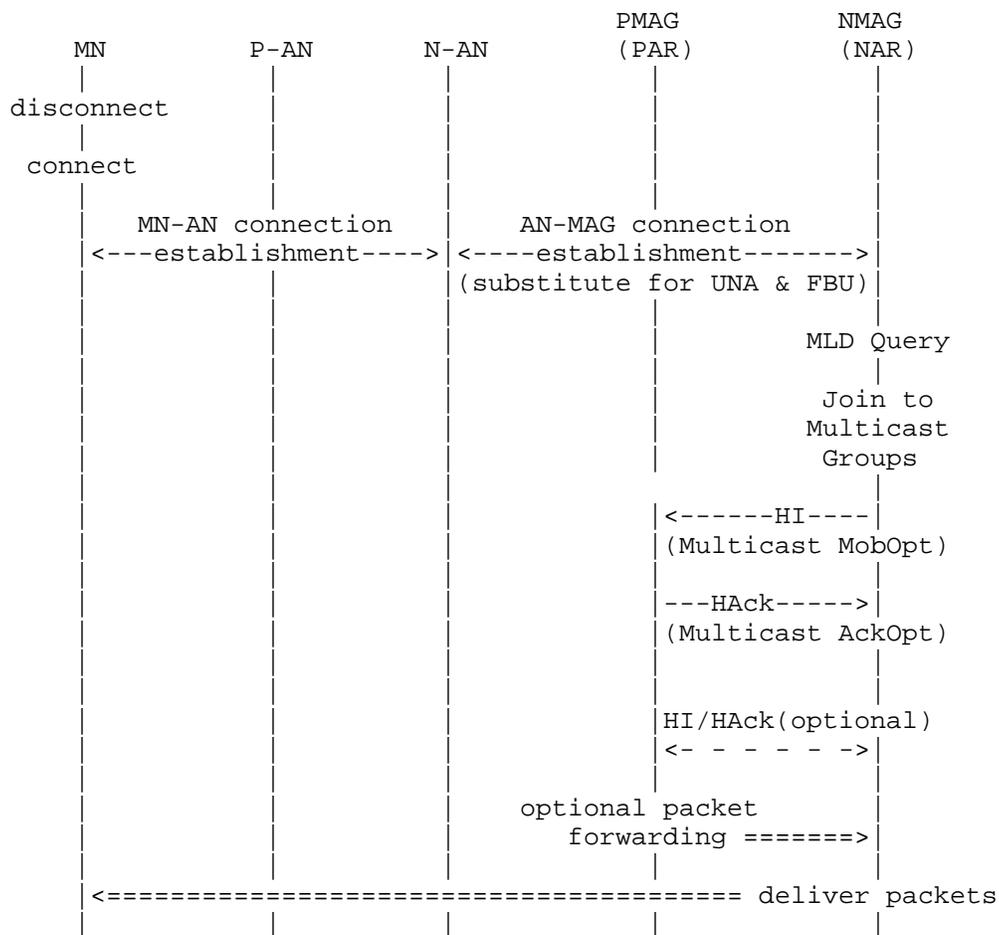


Figure 5: Reactive Multicast Handover for PFMIPv6

4. Protocol Details

4.1. Protocol Operations Specific to FMIPv6

4.1.1. Operations of the Mobile Node

A Mobile Node willing to manage multicast traffic within fast handover operations will inform about its MLD listener state records within handover signaling.

When sensing a handover in predictive mode, an MN will build a

Multicast Mobility Option as described in Section 5.3 that contains the MLD (IGMP) multicast listener state and append it to the Fast Binding Update (FBU) prior to signaling with PAR. It will receive the Multicast Acknowledgement Option(s) as part of Fast Binding Acknowledge (FBack) (see Section 5.4) and learn about unsupported or prohibited groups at the NAR. The MN MAY take appropriate actions like home tunneling to bridge missing multicast services in the new access network. No multicast-specific operation is required by the MN when re-attaching in the new network besides standard FMIPv6 signaling.

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

4.1.2. Operations of the Previous Access Router

A PAR will advertise its multicast support by setting the M-bit in PrRtAdv.

In predictive mode, a PAR will receive the multicast listener state of a MN prior to handover from the Multicast Mobility Option appended to the FBU. It will forward these records to NAR within HI messages and will expect Multicast Acknowledgement Option(s) in HACK, which itself is returned to the MN as an appendix to FBACK. In performing multicast context exchange, the AR is instructed to include the PAR-to-NAR tunnel obtained from unicast handover management in its multicast downstream interfaces and await MLD listener reports from NAR. In response to receiving multicast subscriptions, PAR will normally forward group data acting as a regular multicast router or proxy. However, NAR MAY refuse to forward some or all of the multicast flows.

In reactive mode, PAR will receive the FBU augmented by the Multicast Mobility Option from the new network, but will continue with an identical multicast record exchange in the HI/HACK dialog. As in the predictive case, it will configure the PAR-to-NAR tunnel for multicast downstream and forward data according to MLD reports obtained from NAR, if capable of forwarding.

In both modes, PAR will interpret the first of the two events, the departure of the MN or the reception of the Multicast Acknowledgement Option(s) as a multicast LEAVE message of the MN 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

NAR will advertise its multicast support by setting the M-bit in PrRtAdv.

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 MLD/IGMP records from the message and intersect the request subscription with its multicast service offer. Further on it will adjoin the supported groups (channels) to the MLD listener state using loopback as downstream interface. This will lead to suitable regular subscriptions on its native multicast upstream interface without additional forwarding. Concurrently, NAR builds a Multicast Acknowledgement Option(s) (see Section 5.4) listing those groups (channels) unsupported on the new access link and returns them within HACK. As soon as the bidirectional tunnel from PAR to NAR is operational, NAR joins the groups subscribed for forwarding on the tunnel link.

In reactive mode, NAR will learn about the multicast listener state of a new MN from the Multicast Mobility Option appended to HI at a time, when the MN has already performed local subscriptions of the multicast service. Thus NAR solely determines the intersection of requested and supported groups (channels) and issues the join requests for group forwarding on the PAR-NAR tunnel interface.

In both modes, NAR MUST send a LEAVE message to the tunnel immediately after forwarding of a group (channel) becomes unneeded, e.g., after native multicast traffic arrives or group membership of the MN terminates.

4.1.4. Buffering Considerations

Multicast packets may be lost during handover. For example, in predictive mode as illustrated by figure 2, although the NAR can forward the multicast traffic before the MN attaches to it, the multicast packets still will be lost after the MN disconnects from PAR and before it attaches to the NAR. In reactive mode as illustrated by figure 3, the situation may be worse since there will be a delay for joining the multicast group after the MN attaches to the NAR. The multicast packets will be lost during this time. Buffering the multicast packets at the PAR can ease the multicast packet loss problem. It should be noted that many multicast traffic is video/audio which is sensitive to delay, the buffering mechanism at the PAR should be optimized to meet the specific application's delay requirement.

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 predictive fast handover mode as a PMAG. It MUST issue an MLD General Query immediately on its corresponding link unless it performs an explicit tracking on that link. After gaining knowledge of the multicast subscriptions of the MN, the PMAG builds a Multicast Mobility Option as described in Section 5.3 that contains the MLD (IGMP) multicast listener state. If not empty, this Mobility Option is appended to the regular fast handover HI messages, or - in the case of unicast HI message being submitted prior to multicast state detection - sent in an additional HI message to the NMAG. PMAG then waits for receiving the Multicast Acknowledgement Option(s) with HACK (see Section 5.4) and the creation of the bidirectional tunnel with NMAG. Thereafter PMAG will add the tunnel to its downstream interfaces in the multicast forwarding database. For those groups (channels) reported in the Multicast Acknowledgement Option(s), i.e., not supported in the new access network, PMAG normally takes appropriate actions (e.g., forwarding, termination) in concordance with the network policy. It SHOULD start forwarding traffic down the tunnel interface for those groups it receives an MLD listener report message from NMAG. However, it MAY deny forwarding service. After the departure of the MN and on the reception of LEAVE messages for groups/channels, PMAG MUST terminate forwarding of the specific groups and update its multicast forwarding database. Correspondingly it issues a group/channel LEAVE to its upstream link, if no more listeners are present on its downstream links.

A MAG receiving a HI message with Multicast Mobility Option for a currently attached node follows the reactive fast handover mode as a PMAG. It will return Multicast Acknowledgement Option(s) (see Section 5.4) within HACK listing those groups/channels unsupported at NMAG. It will add the bidirectional tunnel with NMAG to its downstream interfaces and will start forwarding multicast traffic for those groups it receives an MLD listener report message from NMAG. At the reception of LEAVE messages for groups (channels), PMAG MUST terminate forwarding of the specific groups and update its multicast

forwarding database. According to its multicast forwarding states, it MAY need to issue a group/channel LEAVE to its upstream link, if no more listeners are present on its downstream links.

In both modes, 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, explicit tracking).

4.2.3. Operations of the New MAG

A MAG receiving a HI message with Multicast Mobility Option for a currently unattached node follows the predictive fast handover mode as NMAG. It will decide on those multicast groups/channels it wants forwarded from the PMAG and builds a Multicast Acknowledgement Option (see Section 5.4) that enumerates only unwanted groups/channels. This Mobility Option is appended to the regular fast handover HACK messages, or - in the case of unicast HACK message being submitted prior to multicast state acknowledgement - sent in an additional HACK message to the PMAG. Immediately thereafter, NMAG SHOULD update its MLD listener state by the new groups/channels obtained from the Multicast Mobility Option. Until the MN re-attaches, NMAG uses its loopback interface for downstream and does not forward traffic to the potential link of the MN. 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, NMAG additionally joins those groups/channels on the tunnel interface that it wants to receive by forwarding from PMAG. NMAG MUST send a LEAVE message to the tunnel immediately after forwarding of a group/channel becomes unneeded, e.g., after native multicast traffic arrives or group membership of the MN terminates.

A MAG experiencing a connection request for a MN without prior reception of a corresponding Multicast Mobility Option is operating in the reactive fast handover mode as NMAG. Following the re-attachment, it immediately issues an MLD General Query to learn about multicast subscriptions of the newly arrived MN. Using standard multicast operations, NMAG joins the missing groups (channels) on its regular multicast upstream interface. Concurrently, it selects groups (channels) for forwarding from PMAG and builds a Multicast Mobility Option as described in Section 5.3 that contains the MLD (IGMP) 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 upcoming of the bidirectional tunnel, NMAG additionally joins those groups/channels on the tunnel interface that it wants to receive by

forwarding from PMAG. When multicast flows arrive, the NMAG forwards data to the appropriate downlink(s). NMAG MUST send a LEAVE message to the tunnel immediately after forwarding of a group/channel becomes unneeded, e.g., after native multicast traffic arrives or group membership of the MN terminates.

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, 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 HI/HACK message exchanges by encapsulating IGMP multicast state records within Multicast Mobility Options (see Section 5.3 and Section 5.4 for details on message formats).

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

5. Message Formats

5.1. Multicast Indicator for Proxy Router Advertisement (PrRtAdv)

An FMIPv6 AR will indicate its multicast support by activating the M-bit in its Proxy Router Advertisements (PrRtAdv). The message extension has the following format.

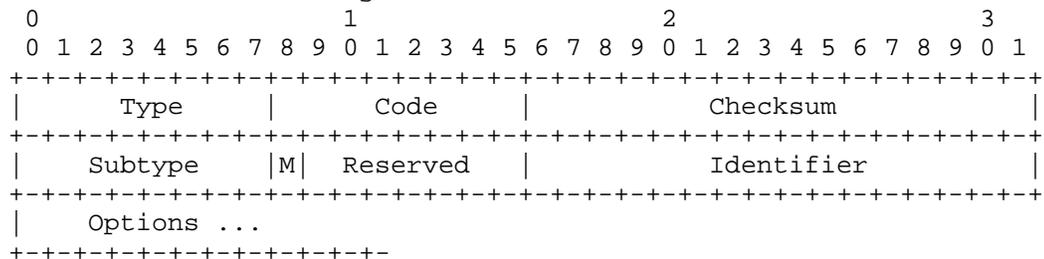


Figure 6: Multicast Indicator Bit for Proxy Router Advertisement (PrRtAdv) Message

5.2. Extensions to Existing Mobility Header Messages

The fast handover protocols use a new IPv6 header type called Mobility Header as defined in [RFC3775]. Mobility headers can carry variable Mobility Options.

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 acknowledged by a corresponding 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

The Multicast Mobility Option contains the current listener state record of the MN obtained from the MLD Report message, and has the format displayed in Figure 7.

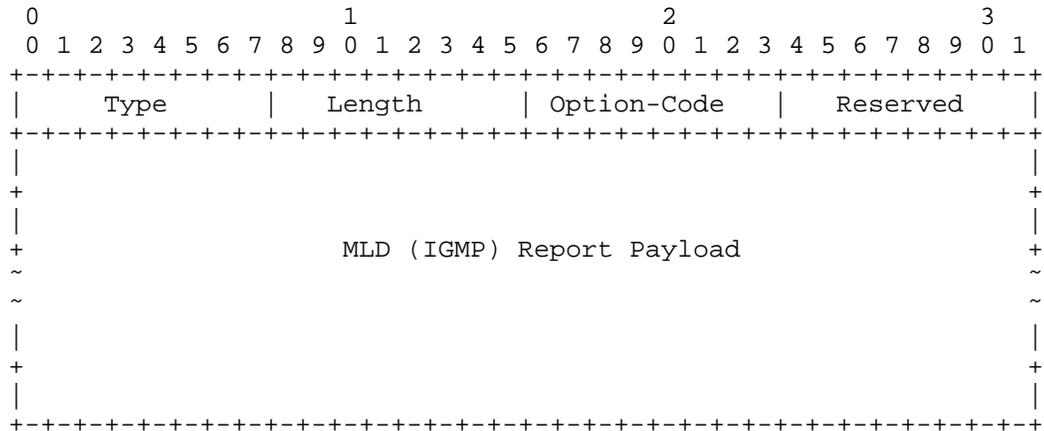


Figure 7: Mobility Header Multicast Option

Type: TBD

Length: 8-bit unsigned integer. The size of this option in 8 octets including the Type, Option-Code, and Length 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 (IGMP) Report Payload: this field is composed of the MLD (IGMP) Report message after stripping its ICMP header. Corresponding message formats are defined for MLDv2 in [RFC3810], and for IGMPv3 in [RFC3376].

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

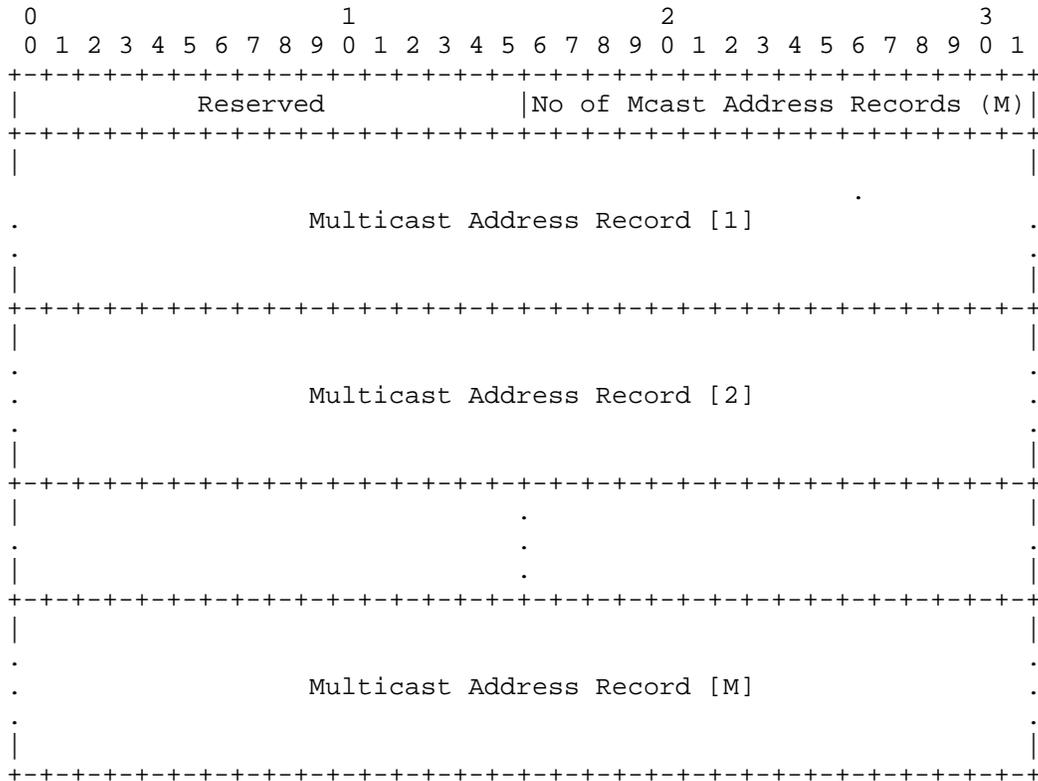


Figure 8: MLDv2 Report Payload

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.

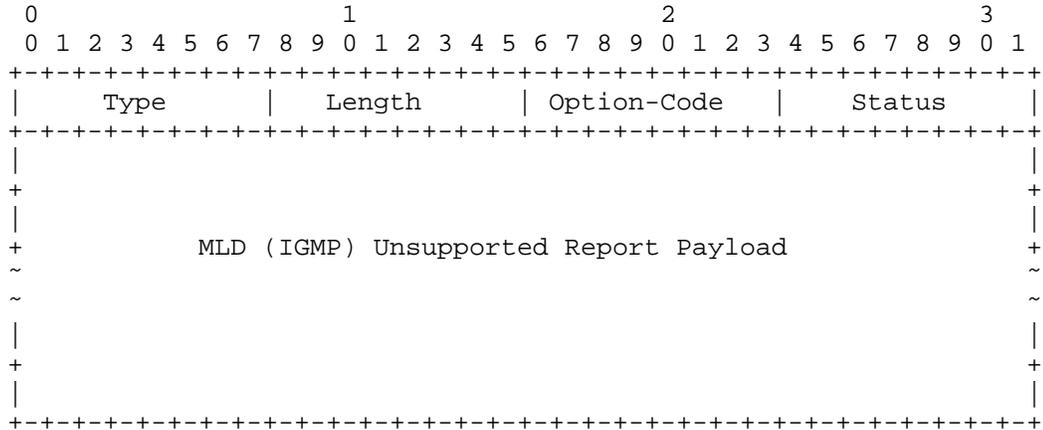


Figure 9: Mobility Header Multicast Acknowledgement Option

Type: TBD

Length: 8-bit unsigned integer. The size of this option is 8 octets. The length is 1 when the MLD (IGMP) Unsupported Report Payload field contains no Mcast Address Record.

Option-Code: 0

Status:

- 1: Report Payload type unsupported
- 2: Requested group service unsupported
- 3: Requested group service administratively prohibited

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

MLD (IGMP) Unsupported Report Payload: this field is syntactically identical to the MLD (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, and thus the composition of multicast address records may differ from initial requests within an MLD (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. The maximal payload length available in FBU/FBACK messages is the PATH-MTU - 40 octets (IPv6 Header) - 6 octets (Mobility Header) - 6 octets (FBU/FBACK Header). For example, on an Ethernet link with an MTU of 1500 octets, not more than 72 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 Report Payloads 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 on an Ethernet link, the number of source addresses is limited to 89.

5.6. MLD (IGMP) Compatibility Aspects

Access routers (MAGs) MUST support MLDv2 (IGMPv3). To enable multicast service for MLDv1 (IGMPv2) listeners, the routers MUST follow the interoperability rules defined in [RFC3810] ([RFC3376]) and appropriately set the Multicast Address Compatibility Mode. When the Multicast Address Compatibility Mode is MLDv1 (IGMPv2), a router internally translates the following MLDv1 (IGMPv2) messages for that multicast address to their MLDv2 (IGMPv2) 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 minimum obtained from its previous and newly learned value of MLD (IGMP) Compatibility Mode for continued operation.

6. Security Considerations

Security vulnerabilities that exceed issues discussed in the base protocols of this document ([RFC5568], [RFC5949], [RFC3810], [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 re-routing 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 re-routing 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 new flags and status codes in the HI and HACK messages as well as two new mobility options. The Type values for these mobility options are assigned from the same numbering space as allocated for the other mobility options defined in [RFC3775]. Those for the flags and status codes are assigned from the corresponding numbering space defined in [RFC5568], or [RFC5949] and requested to be created as new tables in the IANA registry (marked with asterisks). New values for these registries can be allocated by Standards Action or IESG approval [RFC5226].

8. Acknowledgments

Protocol extensions to support multicast in Fast Mobile IPv6 have been loosely discussed since several years. Repeated attempts have been taken to define corresponding protocol extensions. The first draft [fmcast-mip6] was presented by Suh, Kwon, Suh, and Park already 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. Comments, discussions and reviewing remarks have been contributed by (in alphabetical order) Carlos J. Bernardos, Luis M. Contreras, Dirk von Hugo, Marco Liebsch, Behcet Sarikaya, Stig Venaas and Juan Carlos Zuniga.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC5568] Koodli, R., "Mobile IPv6 Fast Handovers", RFC 5568, July 2009.
- [RFC5949] Yokota, H., Chowdhury, K., Koodli, R., Patil, B., and F. Xia, "Fast Handovers for Proxy Mobile IPv6", RFC 5949, September 2010.
- [RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, August 1989.
- [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.
- [RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.
- [RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

9.2. Informative References

- [RFC5757] Schmidt, T., Waehlich, M., and G. Fairhurst, "Multicast Mobility in Mobile IP Version 6 (MIPv6): Problem Statement and Brief Survey", RFC 5757, February 2010.
- [fmcast-mip6] Suh, K., Kwon, D., Suh, Y., and Y. Park, "Fast Multicast Protocol for Mobile IPv6 in the fast handovers environments", draft-suh-mipshop-fmcast-mip6-00 (work in progress), July 2004.

[FMIPv6-Analysis]

Schmidt, TC. and M. Waehlich, "Predictive versus Reactive - Analysis of Handover Performance and Its Implications on IPv6 and Multicast Mobility", Telecommunication Systems Vol 33, No. 1-3, pp. 131-154, November 2005.

[RFC6224] Schmidt, T., Waehlich, M., and S. Krishnan, "Base Deployment for Multicast Listener Support in Proxy Mobile IPv6 (PMIPv6) Domains", RFC 6224, April 2011.

[RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010.

[RFC5845] Muhanna, A., Khalil, M., Gundavelli, S., and K. Leung, "Generic Routing Encapsulation (GRE) Key Option for Proxy Mobile IPv6", RFC 5845, June 2010.

Appendix A. Change Log

The following changes have been made from draft-ietf-multimob-fmipv6-pfmipv6-multicast-00.

1. Buffering text added from new co-author Dapeng.
2. Several editorial improvements.

The following changes have been made from draft-schmidt-multimob-fmipv6-pfmipv6-multicast-04.

1. Following working group feedback, multicast traffic forwarding is now a two-sided option between PAR (PMAG) and NAR (NMAG): Either access router can decide on its contribution to the data plane.
2. Several editorial improvements.

The following changes have been made from draft-schmidt-multimob-fmipv6-pfmipv6-multicast-03.

1. References updated.

The following changes have been made from draft-schmidt-multimob-fmipv6-pfmipv6-multicast-02.

1. Detailed operations on PFMIPv6 entities completed.
2. Some editorial improvements & clarifications.

3. References updated.

The following changes have been made from draft-schmidt-multimob-fmipv6-pfmipv6-multicast-01.

1. First detailed operations on PFMIPv6 added.
2. IPv4 support considerations for PFMIPv6 added.
3. Section on length considerations for multicast context records corrected.
4. Many editorial improvements & clarifications.
5. References updated.

The following changes have been made from draft-schmidt-multimob-fmipv6-pfmipv6-multicast-00.

1. Editorial improvements & clarifications.
2. Section on length considerations for multicast context records added.
3. Section on MLD/IGMP compatibility aspects added.
4. Security section added.

Authors' Addresses

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

Email: schmidt@informatik.haw-hamburg.de

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

Email: mw@link-lab.net

Rajeev Koodli
Cisco Systems
30 International Place
Xuanwu District,
Tewksbury MA 01876
USA

Email: rkoodli@cisco.com

Godred Fairhurst
University of Aberdeen
School of Engineering
Aberdeen AB24 3UE
UK

Email: gorry@erg.abdn.ac.uk

Dapeng Liu
China Mobile

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

