

Internet Draft

Expires: May 2006

Thomas C. Schmidt
HAW Hamburg
Matthias Waehlich
FHTW Berlin
November 2005

**Seamless Multicast Handover in a
Hierarchical Mobile IPv6 Environment (M-HMIPv6)**
<[draft-schmidt-waehlich-mhmipv6-04.txt](#)>

IPR Statement

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#)[1].

Status of this Memo

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Abstract

This document extends the Hierarchical Mobile IPv6 Internet Draft to include the reception and transmission of Any Source Multicast traffic at the Mobile Node. It introduces handover mechanisms for IPv6 mobile multicast listeners and mobile multicast senders. Operations are based on a Mobile IPv6 environment with local mobility anchor points. These local anchor points are conformal with a Hierarchical Mobile IPv6 proxy infrastructure. Handover latencies in the proposed scheme remain bound to link switching delays with respect to these local proxy points. Thus the M-HMIPv6 achieves

seamless mobility, even though no bicasting of multicast streams is used. Multicast listeners in addition encounter the option to optimize multicast routing by turning to a direct data reception. The mechanisms described in this document do not rely on assumptions of any specific multicast routing protocol in use. The M-HMIPv6 protocol operations utilize the existing HMIPv6, MIPv6 and MLDv2 messages under minor extensions.

Table of Contents

1. Terminology.....	3
2. Introduction.....	3
3. Overview of M-HMIPv6.....	4
3.1 Operations of a multicast listener.....	5
3.2 Operations of a multicast sender.....	5
4. Multicast specific extensions of MIPv6, HMIPv6 and MLDv2.....	6
4.1 M-HMIPv6 flag in MAP option message.....	6
4.2 Use of Home Address Destination Option in mobile multicast.	7
4.3 Binding Cache processing at Correspondent Node.....	7
4.4 Home Agent Multicast Membership control.....	7
4.5 Router attendance control in MLD.....	8
5. Protocol Details.....	9
5.1 Operations of all Mobile Nodes.....	10
5.2 Mobile multicast listener.....	10
5.2.1 Operations of the Mobile Node.....	10
5.2.2 Operations of the MAP.....	11
5.2.3 Buffering.....	12
5.3 Mobile multicast source.....	12
5.3.1 Operations of the Mobile Node.....	12
5.3.2 Operations of the MAP.....	13
5.3.3 Tree initialization procedure.....	13
5.3.4 Buffering.....	14
5.4 Protocol Timer.....	14
6. Security Considerations.....	14
7. IANA Considerations.....	14
8. References.....	15
Acknowledgments.....	16
Author's Addresses.....	16

A. A Note on Tunneling.....	16
Copyright Notice.....	17
Disclaimer of Validity.....	17
Acknowledgement.....	17

[1. Terminology](#)

The terminology used in this document remains conformal to the definitions in MIPv6 [4] and HMIPv6 [6].

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](#) [2].

[2. Introduction](#)

Multicast-based packet distribution plays an important role in real-time applications, as it provides the only efficient, scalable scheme for group communication. However, any source multicasting itself conceals complex mechanisms for group membership management and routing, which both are of slow convergence. To achieve seamless mobility here is one of the most challenging and demanded developments in IP networks today.

In multimedia conference scenarios each member commonly operates as receiver and as sender for multicast based group communication. In addition, real-time communication such as voice or video over IP places severe temporal requirement on mobility protocols: Seamless handover scenarios need to limit disruptions or delay to less than 100 ms. Jitter disturbances are not to exceed 50 ms. Note that 100 ms is about the duration of a spoken syllable in real-time audio traffic.

The fundamental approach to deal with mobility in IPv6 [3] is stated in the Mobile IPv6 RFCs [4,5]. MIPv6 operates address changes on the IP layer transparent to the transport layer as a device moves from one network to the other. MIPv6 involves roundtrip messages for location updates directly with the MNs Home Agent and the Correspondent Node. As these nodes can be far away, MIPv6 may exhibit slow handover performance. The Hierarchical Mobility Management (HMIPv6) Internet Draft [6] introduces a proxy architecture of Mobility Anchor Points (MAPs) to reduce communication delays with

respect to the HA. In addition the Fast Handover for Mobile IPv6 Internet Draft [7] proposes predictive delay hiding techniques to further reduce handover times in unicast data.

MIPv6 only roughly treats multicast mobility, in a pure remote subscription approach or through bi-directional tunneling via the Home Agent. It thereby suffers from slow handovers and inefficient, triangular forwarding. It is the issue of this document to extend the improved HMIPv6 mobility infrastructure by mechanisms of sending and receiving multicast traffic for the MN. Local MAPs serve as temporary multicast relays to hide partly movement, partly handoff latency of the MN. The multicast support through a MAP infrastructure may significantly reduce the attained handover frequencies [11]. Handover procedures between MAPs solely built on MIPv6 and HMIPv6 signaling are described within this draft. They are designed to limit any disruption or disturbance to the time scale needed for reconnecting to neighboring MAPs. An additional option in multicast data delivery allows for turning to optimal routing, after a receiver handover has been completed. Minor MLD [9,10] extensions are required to operate this optimization option. All mechanisms remain transparent to any specific multicast routing protocol in use.

3. Overview of M-HMIPv6

This multicast mobility scheme is built on a HMIPv6 environment. HMIPv6 introduces Mobility Anchor Points as proxy elements, which may be best viewed as functions on regional routers. For implementing multicast mobility it is advantageous, but not necessary, that these regional routers provide multicast routing functionality.

In M-HMIPv6 a mobile multicast node uses its local MAP as anchor point for multicast communication. All multicast traffic is directed through this MAP using the Regional Care-of Address RCoA as multicast subscriber or source address. Traffic forwarding between MN and its MAP is done using a bi-directional tunnel [8].

If a MN changes location within its MAP domain, it only registers its new LCoA with the MAP as defined in [6]. This does not affect multicast routing trees. When entering a new MAP domain a MN will be eager to sustain multicast connectivity via its previously established MAP. Eventually it will learn of M-HMIPv6 support through router advertisements with MAP option messages, and will then perform a reactive handover. Multicast handover procedures will occur only if the MN changes into a new M-HMIPv6 enabled MAP domain and will then transfer multicast traffic from the previous to the current MAP.

An M-HMIPv6-aware MN SHOULD use the MAP for multicast communication.

However, the MN MAY prefer to use its HA as a multicast anchor point,

e.g. in visited networks within its home site. A mobile node, which is not M-HMIPv6 aware, will not use its MAP as a multicast anchor point, but will use the MIPv6 tunnel through the HA instead. In this sense M-HMIPv6 is simply a smooth extension of HMIPv6, which itself smoothly extends MIPv6.

3.1 Operations of a multicast listener

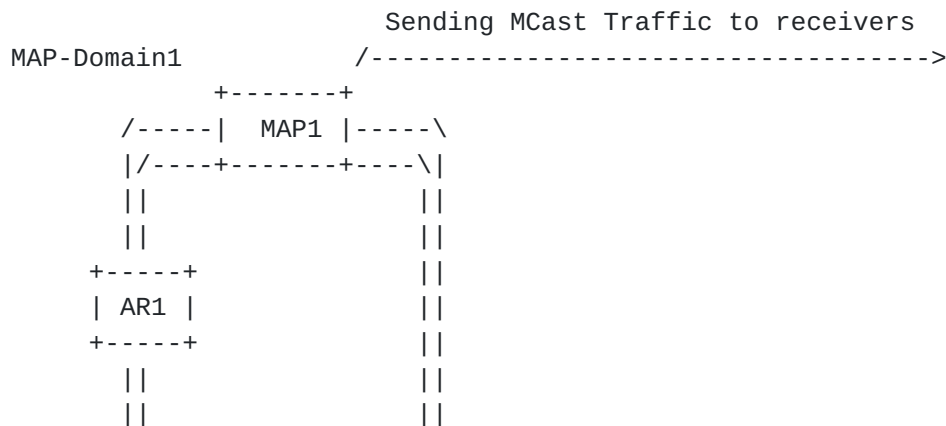
To join a multicast group away from home the MN tunnels the MLD [9,10] listener report to its current MAP using RCoA as source address. The MAP records the group address in its Binding Cache in order to forward multicast packets to the MN and to subscribe for and preserve MNs multicast group membership.

When changing its MAP domain, the MN submits a Binding Update with its new LCoA to the previous MAP in addition to regular HMIPv6 handover signaling. On its reception the previous MAP redirects multicast packet forwarding to the MN's new LCoA.

If multicast support is advertised in the new domain the MN immediately SHOULD join the multicast group through the new MAP. Once multicast group traffic arrives the MN SHOULD send a Binding Update with zero lifetime to its previous MAP to eliminate its Binding Cache entry and end packet forwarding.

3.2 Operations of a multicast sender

In a foreign MAP domain a MN initiates multicast-based communication by sending packets through its MAP using RCoA as its source address. As receivers are aware of source addresses and as the mobile RCoA address may change, a Home Address Destination Option MUST be included (s. [section 4.2](#)). By transmitting multicast packets along this path a routing tree originating at the MAP will be constructed. Local movement of the MN within a MAP domain thereby remains transparent to multicast routing.



| Type | Length | Dist | Pref |*|M| Reserved |

Schmidt, Waehlich

Expires - May 2006

[Page 6]

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Valid Lifetime                             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                                                              |
+                                                                              +
|                                                                              |
+          Global IP Address for MAP          +
|                                                                              |
+                                                                              +
|                                                                              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Flags:

*	Used by HMIPv6
M	When set indicates that M-HMIPv6 is supported by the current MAP

4.2 Use of Home Address Destination Option in mobile multicast

Multicast applications normally are aware of source addresses, which MUST NOT change during ongoing communication. A mobile multicast sender therefore MUST include a home address destination option as defined in [4]. This requirement deviates from MIPv6 multicast scheme.

4.3 Binding Cache processing at Correspondent Node

A Correspondent Node receiving multicast packets with Home Address Option in general need not have a Binding Cache Entry for the home address included. A CN therefore SHALL NOT verify multicast packets with respect to its Binding Cache. This requirement deviates from MIPv6 unicast scheme.

4.4 Home Agent Multicast Membership control

To provide multicast connectivity to a mobile multicast listener away from home, a HA needs to take care of the local multicast group management. This essentially can be done by either supplying full multicast routing functionalities at the HA, or by a proxy agent function.

In the first case it suffices for the HA to observe MNs group membership at the (tunnel) interface. For a multicast proxy function

mobile listener and a mobile sender.

Schmidt, Waehlich

Expires - May 2006

[Page 9]

Mobility Anchor Points as defined in [6] attain the role of multicast mobility anchor points (M-MAPs) for mobile group members in M-HMIPv6. All multicast traffic is directed through M-MAPs using RCoA consistently as identifier with respect to the multicast routing tree. M-MAPs may be viewed as HA proxies for multicast streams, just as MAPs in the unicast case.

5.1 Operations of all Mobile Nodes

Being at home the MN may either use its Home Agent or, a possibly distinct, regional M-MAP as multicast anchor point. Away from home the MN will learn about regional M-MAPs through router advertisements (s. [section 4.1](#)). A MN SHOULD register with the regional M-MAP having the highest preference value. If M-HMIPv6 is not supported regionally the MN first SHOULD attempt to employ a previously established M-MAP relation, second register with its HA.

M-MAP presence is advertised via router advertisements with MAP option message as described in [section 4.1](#).

5.2 Mobile multicast listener

Any node on a multicast enabled network may remotely subscribe to multicast group membership by using its link-local address in MLD membership reports. In doing so a MN cannot expect to experience a smooth handover performance while changing from one network to another.

A MN utilizing an HMIPv6 MAP infrastructure can be regarded as eager for decreased handover delays and therefore SHOULD use the M-HMIPv6 M-MAP functionality for other than link locally scoped multicast reception.

5.2.1 Operations of the Mobile Node

A mobile multicast listener registers with its local M-MAP (or HA), where both communicate via a bi-directional tunnel. The MN submits its MLD group membership listener report through this tunnel and answers membership queries of the anchor point.

When a Mobile Node changes its network, it performs a Binding Update with its previous M-MAP and re-establishes the tunnel at its new LCoA. Thereafter it continues to receive multicast group traffic.

On entering a new M-MAP domain a MN - in addition to the above BU - registers with the new M-MAP and establishes a bi-directional tunnel. It immediately sends a MLD listener report through the newly available connection, one for each group/flow to be handed over. Once

multicast group traffic arrives from the new M-MAP, the MN SHOULD submit a BU with zero lifetime to its previous M-MAP and terminate the corresponding tunnel. If previous binding of the MN had been with the HA, the MN MUST NOT terminate its binding, but SHOULD tunnel an MLD listener done message instead.

Note that a MN SHOULD preserve a previously established M-MAP relation until a new multicast forwarding is completely established. In the case of rapid movement this may lead to a previous multicast anchor point persisting through several hops.

As an optional extension to optimize routing a MN MAY be enabled to directly subscribe to multicast groups in its visited subnet. This remote subscription SHOULD be performed, if triggered by M-MAP MLD listener queries marked with attendance code as described in [section 4.5](#). It MAY be performed on MN's own reasons, the recognition of slow handover frequencies or significant M-MAP distance, for instance.

To optimize routing for a specific multicast group the MN undertakes a regular MLD subscription at its link local interface using its LCoA. After receiving multicast data on this link local interface, the MN MUST tunnel an MLD listener report to its M-MAP with attendance code as described in [section 4.5](#). On further MLD listener queries of its M-MAP the MN MUST reply with listener reports. These reports SHOULD carry the attendance code as long as the MN receives multicast streams locally.

[5.2.2](#) Operations of the MAP

M-MAP operations for multicast listener support are completely analog to Home Agent functions as described in [4] and [section 4.4](#). An M-MAP receiving a HMIPv6 BU from a MN will establish a bi-directional tunnel. On reception of a tunneled MLD listener report it will

- o record multicast group membership in its Binding Cache;
- o observe and maintain multicast group membership on its specific tunnel interface;
- o inquire on MNs current group membership as described in [4];
- o forward multicast group traffic to the MN (see [4] on multicast packet forwarding rules).

The M-MAP may control multicast group membership either as a multicast router or as a multicast proxy agent (s. [section 4.4](#)).

As an optional extension to optimize routing the M-MAP MAY be enabled to direct MNs do directly subscribe to multicast groups within their visited subnets by using the MLD attendance extensions described in [section 4.5](#). The M-MAP MAY decide to initiate remote subscriptions of MNs by tunneling MLD queries with attendance code. This decision

could be based on the number of attached MNs subscribed to the same multicast groups or link capacities or forwarding load, for instance. If the M-MAP itself acts as a multicast router, it will operate exactly as described in [section 4.5](#) for each tunnel interface associated with a MN. Otherwise the M-MAP will intercept MLD queries from multicast routers to add attendance codes. Similar it will intercept listener reports from its attached MNs to remove attendance codes.

Regardless of its own queries the M-MAP must continue to deliver multicast streams to any attached MN, which reports on group membership without attendance code.

5.2.3 Buffering

Some L2 technologies imply a noticeable offline period for a MN during handover. To compensate for possible packet loss, buffering mechanisms are needed. In M-HMIPv6 M-MAPs may provide automatic replay buffers at the tunnel entry points, to be played out after a MN's Binding Update occurred.

5.3 Mobile multicast source

A multicast source sending with its link-local address is immobile with respect to multicast application persistence. A mobile multicast sender MAY tunnel multicast traffic through its HA, using its home address as source address [4]. Triangular routing and significant binding update times lead to expected large handover delays, in general.

A MN utilizing a HMIPv6 MAP infrastructure therefore SHOULD use the M-HMIPv6 M-MAP functionality for other than link locally scoped multicast transmissions.

5.3.1 Operations of the Mobile Node

A mobile multicast sender is registered with its local M-MAP, where both communicate via a bi-directional tunnel. The MN submits multicast packets through this tunnel with the RCoA as the source address and the home address included in a home address destination option as defined in [4].

When a Mobile Node changes networks, it performs a Binding Update with its previous M-MAP and re-establishes the tunnel at its new LCoA. Thereafter it continues to send its multicast group traffic, using previous RCoA as its source address.

On entering a new M-MAP domain a MN - in addition to the above BU - registers with the new M-MAP and establishes a bi-directional tunnel.

It immediately SHOULD start the tree initialization procedure as defined in [section 5.3.3](#) and start a timer. As soon as this timer exceeds MAX_MCASTTREEINIT_TIMEOUT the MN MUST complete the handover by terminating multicast group forwarding through its previous M-MAP and submit all subsequent traffic to its current M-MAP using its current RCoA as source address. Note that these handover steps can be performed stream wise.

A MN, which moves to a new link within the same M-MAP domain before the timeout is reached, performs a BU with its current M-MAP and continues the handover procedure without resetting its timers.

A MN, which moves into a new M-MAP domain before the timeout occurred, continues to forward multicast traffic through its previously established old M-MAP, discontinues to communicate via its previously not fully established intermediate M-MAP, resets its timer and restarts the tree initialization procedure for its current M-MAP.

Thus in case of rapid movement the MN stays bound with its formerly fully established (or first) M-MAP, serving the last completely erected multicast routing tree.

5.3.2 Operations of the MAP

M-MAP operations for multicast sender support are completely analog to MAP functions for unicast support as described in [6].

5.3.3 Tree initialization procedure

In preparation for a seamless handover of a multicast sender, a distribution tree needs to be constructed by the routers, which originates at the new M-MAP. In general, Any Source Multicast routing trees will be initiated by submitting packets into the appropriate multicast group. Depending on the routing protocol in use, this can be a tardy procedure. A multicast sender MAY initiate a new group tree by bi-casting its packets to its previous and its new point of attachment. Bi-casting in the presence of slow routing protocols, though, may result in a significant amount of duplicate traffic. In many cases it may be highly desirable to proceed with less communication overhead. The tree initialization procedure provides a way for the MN to efficiently bridge the multicast routing convergence gap.

In performing the tree initialization procedure, the source starts to send probe packets, destined to all multicast groups under migration, with complete IPv6 header, but without transport payload. In detail, the next header field of tree initialization packets contains IPv6-NoNxt (59) value. Subsequent packets SHOULD be sent with a random delay uniformly chosen between 0 and MCASTTREEINIT_FREQUENCY.

The tree initialization procedure ends after MAX_MCASTTREEINIT_TIMEOUT is reached with continuous submission of regular traffic to the initiated delivery tree.

5.3.4 Buffering

To prevent or reduce packet loss during handover, the mobile source MAY buffer packets to be sent, while its tunnel to the M-MAP is not established. This buffer should be played out as soon as the tunnel re-establishment to the previous MAP has completed.

5.4 Protocol Timer

MAX_MCASTTREEINIT_TIMEOUT	180 seconds (Default)
	160 seconds (For DVMRP regimes)
	0.5 seconds (For PIM-SM regimes)

MCASTTREEINIT_FREQUENCY	10 seconds (Default)
-------------------------	----------------------

Mobile nodes must allow these variables to be configured by system management.

6. Security Considerations

This specification uses the concepts of Mobile IPv6 and Hierarchical Mobile IPv6 mobility management. All security provisions regarding the relation between the Mobile Node and the Home Agents and between the Mobile Node and the Mobility Anchor Points apply equally to this M-HMIPv6 concept.

Threats of hijacking unicast sessions derive from attempts of a MN to operate binding updates within multicast sessions. Any binding update received within a multicast session therefore MUST be ignored.

7. IANA Considerations

This document defines extension codes for two ICMPv6 messages. For the

Type Multicast Listener Query (Type = decimal 130)

and the

Type Version 2 Multicast Listener Report Message
(Type = decimal 143)

this requires the registration of two codes. The suggested values for these codes are:

Code

- 0: Query on actively forwarded multicast groups
- 1: Query on multicast groups intended for attendance.

8. References

Normative References

- 1 Bradner, S., "Intellectual Property Rights in IETF Technology", [BCP 79](#), [RFC 3979](#), March 2005.
- 2 Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- 3 Hinden, R. and Deering, S. "Internet Protocol Version 6 Specification", [RFC 2460](#), December 1998.
- 4 Johnson, D.B., Perkins, C., Arkko, J. "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- 5 Arkko, J., Devarapalli, V., Dupont, F "Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents", [RFC 3776](#), June 2004.
- 6 Soliman, H., Castelluccia, C., El-Malki, K., Bellier, L. "Hierarchical Mobile IPv6 mobility management", [RFC 4140](#), August 2005.
- 7 Koodli, R. "Fast Handovers for Mobile IPv6", [RFC 4068](#), July 2005.
- 8 Conta, A., Deering, S. "Generic Packet Tunneling in IPv6 Specification", [RFC 2473](#), December 1998.
- 9 S. Deering, W. Fenner and B. Haberman "Multicast Listener Discovery (MLD) for IPv6", [RFC 2710](#), October 1999.
- 10 R. Vida and L. Costa (Eds.) "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC3810](#), June 2004.

Informative References

- 11 Schmidt, T.C. and Waehlich, M. "Predictive versus Reactive - Analysis of Handover Performance and Its Implications on IPv6 and Multicast Mobility", Telecommunication Systems, Vol. 30, No. 1, pp. 123-142, Berlin Heidelberg:Springer, 2005.

- 12 Romdhani, I., Kellil, M., Lach, H.-Y. et. al. "IP Mobile Multicast: Challenges and Solutions", IEEE Comm. Surveys, 6, 1, 2004.
- 13 Suh, K., Kwon, D.-H., Suh, Y.-J. and Park, Y. "Fast Multicast Protocol for Mobile IPv6 in the fast handovers environments", IETF, Internet Draft - (work in progress), February 2004.
- 14 Jelger, C., Noel, T. "Multicast for Mobile Hosts in IP Networks: Progress and Challenges", IEEE Wireless Comm., pp 58-64, Oct. 2002.
- 15 Schmidt, T.C. and Waehlich, M. "Multicast Mobility in MIPv6: Problem Statement", [draft-schmidt-mobopts-mmcastv6-ps-00.txt](#) - (work in progress), October 2005.

Acknowledgments

The authors would like to thank Stefan Zech (FHTW Berlin), Mark Palkow (DaViKo GmbH) and Hans L. Cycon (FHTW Berlin) for valuable discussions and a joyful collaboration.

Author's Addresses

Thomas C. Schmidt
HAW Hamburg, Dept. Informatik
Berliner Tor 7
D-20099 Hamburg
Phone: +49-40-42875-8157
Email: Schmidt@informatik.haw-hamburg.de

Matthias Waehlich
FHTW Berlin, HRZ
Treskowallee 8
D-10318 Berlin
Email: mw@fhtw-berlin.de

A. A Note on Tunneling

Following the concepts of MIPv6 and HMIPv6 the packet forwarding to and from the Mobile Node is organized by means of a tunnel section spanned to a static anchor component such as a MAP or a Home Agent.

Through this technique a MN can hide its movement to CNs or to the routing infrastructure.

However, keeping in mind real-time data requirements it is highly desirable to avoid packet encapsulation. Besides the unwanted overhead, a tunnel may hide QoS information of the original packet headers and may require load and jitter generating packet fragmentation, if the tunnel entry point is distinguished from the sender.

Tunnelling can be avoided by a direct packet forwarding of the static anchor components. Such forwarding requires a change of packet's source or destination address at the forwarder, which usually conflicts with checksums covering IPv6 pseudo headers. In M-MIPv6 multicast communication from a Mobile Node though carries a MIPv6 extension header, the home address destination option header. This header denotes an alternate source address which enters the pseudo header instead of the original IPv6 header address.

Multicast packets sent from the MN therefore could be forwarded by the MAP to the network by replacing source addresses without recalculation of header checksums. Employing such direct packet forwarding would allow a MN to distribute multicast streams without a tunnel.

Copyright Notice

Copyright (C) The Internet Society (2005). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Disclaimer of Validity

"This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

Acknowledgement

Funding of the RFC Editor function is currently provided by the Internet Society

