Network Working Group	J. Laganier
Internet-Draft	DoCoMo Euro-Labs
Intended status: Informational	S. Narayanan
Expires: August 16, 2008	iTCD/CSUMB
	P. McCann
	Motorola
	February 13, 2008

TOC

Interface between a Proxy MIPv6 Mobility Access Gateway and a Mobile Node
draft-ietf-netlmm-mn-ar-if-03

Status of this Memo

Abstract

This document describes an interface between mobile nodes (MNs) and a mobility access gateway (MAG) of a network-based localized mobility management (NETLMM) domain. The interface has two functions which are invoked when a MN attaches and detaches from a MAG. The attachment function lets the MAG authenticate the MN identifier, does address(es) and default router configuration for the MN, and informs the MAG about the multicast listener state of the MN. During the attachment function the NETLMM protocol is triggered between the MAG and Localized Mobility Anchor (LMA) to register the MN in the local domain. The detachment

function lets the MAG detect that the MN has left so that it can deregister the MN at the LMA using the NETLMM protocol.

Table of Contents

---

1. Introduction

TOC

It is suggested in [RFC4830] (Kempf, J., “Problem Statement for Network-Based Localized Mobility Management (NETLMM),” April 2007.) that it would be desirable to have a localized mobility management protocol in which the host is not involved. The requirements for such a protocol have been analyzed in [RFC4831] (Kempf, J., “Goals for Network-Based Localized Mobility Management (NETLMM),” April 2007.). Accordingly, a protocol for network-based localized mobility management (NETLMM) of IPv6 nodes is specified by the NETLMM working group [I-D.ietf-netlmm-proxymip6] (Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, “Proxy Mobile IPv6,” May 2008.).

Because the NETLMM protocol is network based, the mobile node (MN) is not required to implement a new mechanism in its IP stack, nor to change its IP address when it attaches to a new mobility access gateway (MAG).
Because the IPv6 MN will use a vanilla IPv6 stack, the interface between an MN and its MAG has to be preserved. This means that standard IPv6 should work seamlessly with the network-based localized mobility support. More specifically, we require the proposed solution to be compatible with the mechanisms specified in:

* Neighbor Discovery for IP version 6 (Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)," December 1998.) [RFC2461]

* IPv6 Stateless Address Autoconfiguration (Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration," December 1998.) [RFC2462]

* Dynamic Host Configuration Protocol for IPv6 (DHCPv6) (Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," July 2003.) [RFC3315]

* Privacy Extensions for Stateless Address Autoconfiguration in IPv6 (Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," January 2001.) [RFC3041]

* Detecting Network Attachment in IPv6 Networks (DNAv6) (Narayanan, S., Kempf, J., Nordmark, E., Pentland, B., Choi, J., Daley, G., and N. Montavont, "Design Alternative for Detecting Network Attachment in IPv6 Networks (DNAv6 Design Alternative)," December 2009.) [I-D.ietf-dna-protocol]

* SEcure Neighbor Discovery (SEND) (Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," March 2005.) [RFC3971]

* Cryptographically Generated Addresses (CGAs) (Aura, T., "Cryptographically Generated Addresses (CGA)," March 2005.) [RFC3972]

This document describes an interface between mobile nodes (MNs) and a mobility access gateway (MAG) of a network-based localized mobility management (NETLMM) domain. The interface has two functions which are invoked when a MN attaches and detaches from a MAG. The attachment function lets the MAG authenticate the MN identifier, does address(es) and default router configuration for the MN, and informs the MAG about the multicast listener state of the MN. During the attachment function the NETLMM protocol is triggered between the MAG and Localized Mobility Anchor (LMA) to register the MN in the local domain. The detachment

function lets the MAG detect that the MN has left so that it can deregister the MN at the LMA using the NETLMM protocol.
In the absence of link-layer specific mechanisms implementing these functions, this document describes which IP protocols should be used to provide the necessary interface between the MN and the MAG.

---

2. Terminology

TOC

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 [RFC2119] (Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.).
The following terms are defined within the scope of this document:

> **Mobile Node (MN)** an IPv6 node moving in the NETLMM domain.
>
> **Mobility Access Gateway (MAG)** a default router that connects the MN to the NETLMM domain.
>
> **Localized Mobility Anchor (LMA)** a router located in the NETLMM domain that handles packet exchanges with nodes in the domain.
>
> **Network-based Localized Mobility Management Domain (NETLMM domain)** an administrative domain spanning links served by a set of LMAs (and their associated MAGs and MNs) that provision addresses from the same IP subnet prefix(es).
>
> **Network-based Localized Mobility Management Protocol (NLMP)** The NETLMM Protocol used in the backhaul of the NETLMM domain between MAGs and LMA.

The following abbreviations are used throughout this document:
NETLMM: Network-based Localized Mobility Management
ND: Neighbor Discovery.
NS: Neighbor Solicitation.
NA: Neighbor Advertizement.
RS: Router Solicitation.
RA: Router Advertisement.
NDP: Neighbor Discovery Protocol.
SLAAC: StateLess Address AutoConfiguration
DHCP: Dynamic Host Configuration Protocol
SEND: SEcure Neighbor Discovery.
DNA: Detecting Network Attachment.
CGA: Cryptographically Generated Address.
MNID: An authenticated MN identifier (e.g. NAI, a SEND public key used by the MN for generating its CGAs,an IMSI or TMSI, etc.).

3. Operating Environment

TOC

The MN-MAG NETLMM interface is used between an MN and a MAG of a NETLMM domain. It allows the MAG and/or MN to detect network attachment and detachment, causing the MAG to use the NETLMM protocol to update routing at the LMA so that the MN stays reachable when it roams across the NETLMM domain.

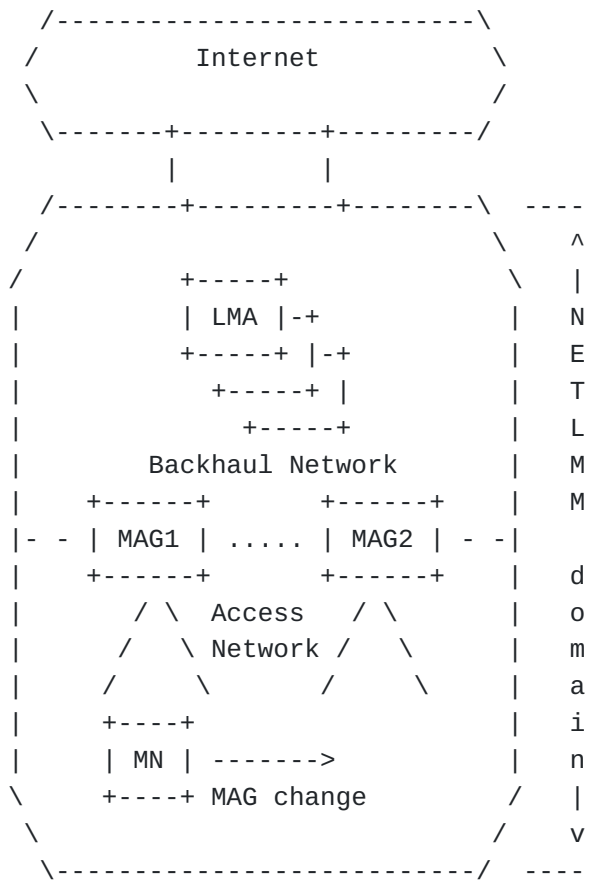

Figure 1: Reference Network Diagram

The deployment scenario is shown in Figure 1 (Reference Network Diagram) above: several MAGs are attached to an IP routing domain connected to the outside Internet via an LMA. The MNs, MAGs, LMAs, and in-between routing fabric constitute the NETLMM domain. Packets arriving at the LMA and destined to a MN are tunneled to the appropriate MAG. Packets from the MN are received by the MAG and tunneled to the LMA and then decapsulated and sent on to the Internet.

In the absence of a link-layer specific mechanisms to implement the MN-MAG interface, it is required to have a common interface defined at the IP layer. Because no NETLMM specific software support is assumed to be present on MNs, this interface has to rely only on standards track IPv6 protocols such as ND, DHCP, SEND, and DNA. Interactions of these components with NETLMM are represented in Figure 2 (NETLMM Component Interactions) below (note that hints received by DNA from other layers are omitted for clarity):

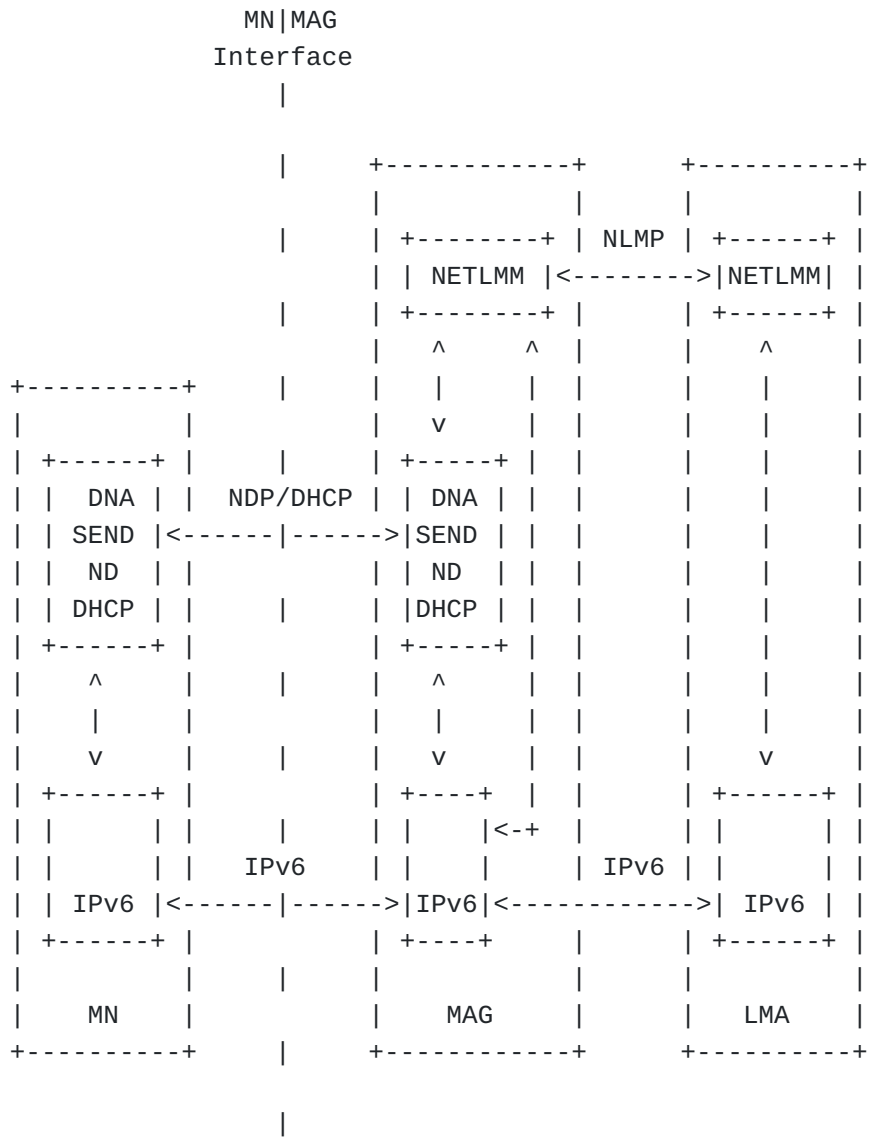

Figure 2: NETLMM Component Interactions

An overview of the interactions between the MN-MAG interface and the NETLMM protocol is shown in Figure 3 (NETLMM MN-MAG Interface Usage Overview).

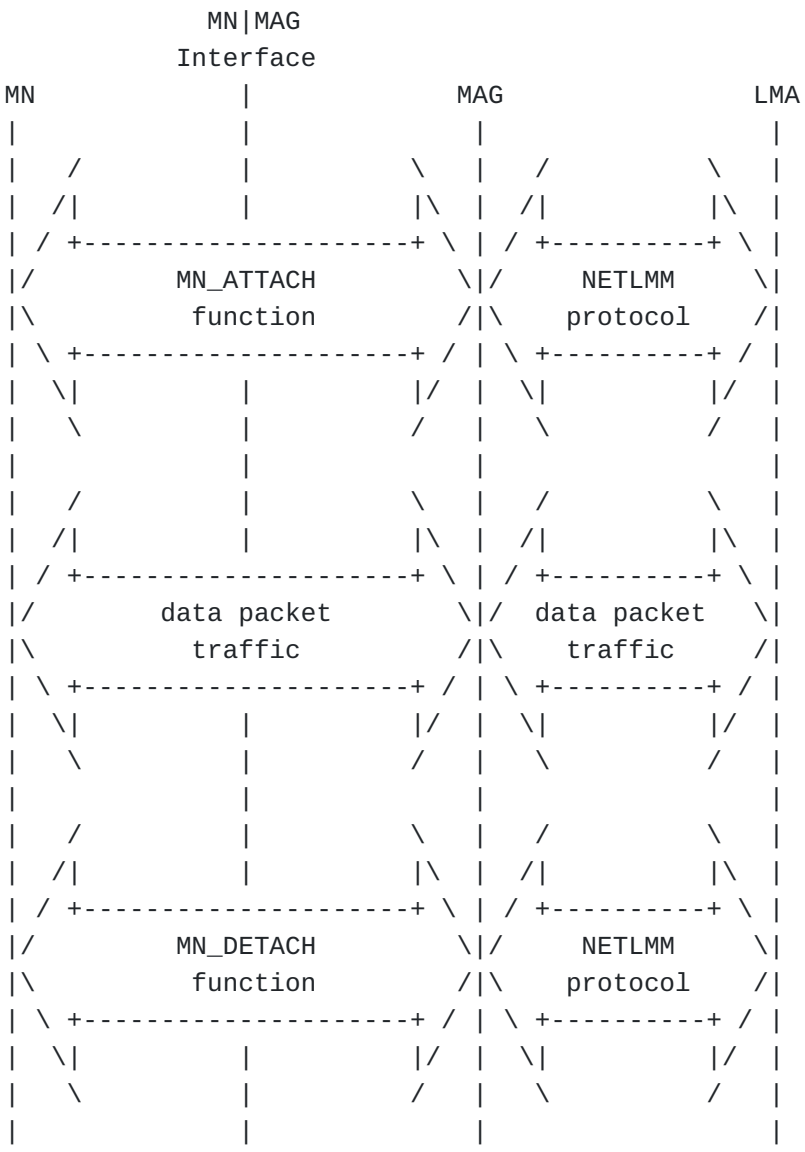

Figure 3: NETLMM MN-MAG Interface Usage Overview

TOC

4. Interactions of NETLMM Architecture with Subnet and Link Models

Within the Internet addressing model, the terms link and subnet have a tight relationship. Their generally admitted definitions are [I-D.thaler-intarea-multilink-subnet-issues] (Thaler, D., “Issues With Protocols Proposing Multilink Subnets,” March 2006.):

> Link: a topological area of an IP network delimited by routers.
>
> Subnet: a topological area of an IP network that uses the same unsubdivided address prefix.

There has recently been protocol proposals achieving multi-link subnets, i.e. the ability for a subnet to span multiple links. However, the consensus in the IETF has been, and remains, that one subnet spans only one link [I-D.thaler-intarea-multilink-subnet-issues] (Thaler, D., “Issues With Protocols Proposing Multilink Subnets,” March 2006.).
A straightforward approach to the design of NETLMM would have been to lay a single subnet on the entire NETLMM domain. That would ensure that the MN does not see layer 3 movements since the subnet would never change. However, such an approach would constitute a multi-link subnet, and is thus not deemed acceptable.
The following subsection will discuss what kind of subnet and link models have been chosen for the NETLMM architecture.

4.1. NETLMM Subnet Model

TOC

Thus, the NETLMM addressing model is subject to the following two constraints:

- The subnet of a MN does not change when the MN changes its attachment point in the domain.
- The subnet of a MN does not span more than one link.

Because of the first constraint, the subnet of a MN must be valid wherever in the domain the MN attaches to. However, because of the second constraint, the subnet cannot be valid at more than one such attachment point. Thus, the subnet of the MN has to follow the movements of the MN. This addressing model is denoted "per-MN subnet model", and satisfies constraints of both the Internet and NETLMM architectures:

> A unique prefix MUST be assigned by the NETLMM fabric to each of the MNs in the domain. The MAG MUST NOT configure a global unicast address based on this prefix.

4.2. NETLMM Link Model

TOC

The choice of the per-MN addressing model is however conflicting with the use of a shared link layer (e.g. Ethernet, 802.11) as a last hop of the NETLMM domain.
In the per-MN subnet model, two MNs always have different subnets. Hence, even though they might be attached to the same shared link layer, they will never communicate directly with global addresses. That happens since on-link determination will always conclude that they are attached to different link because it is based on subnet comparisons. They will however be able to communicate directly with link-local addresses. This is not problematic since link-local addresses are confined to one link and therefore it does not introduce multi-link subnet issues.
There is however one problem that arises due to the use of Solicited-Node and All-Nodes multicast IPv6 addresses [RFC4291] (Hinden, R. and S. Deering, “IP Version 6 Addressing Architecture,” February 2006.) as a destination address for sending unsolicited Neighbor Advertisement (NA) messages [RFC2461] (Narten, T., Nordmark, E., and W. Simpson, “Neighbor Discovery for IP Version 6 (IPv6),” December 1998.). When one MN sends such message, it can be received by other MNs on the same link which will, as a result, create a neighbor cache entry for the sender of the NA. If the NA contained as a target address one of the MN's global unicast address, the receiver is then in a position to communicate directly with this global unicast address, even though it does not share a common subnet prefix (they are per-MN subnet prefixes). This is not a problem as long as these two MNs remain attached on the same link. But if later on one of the MN moves onto a different link, they will no longer be able to communicate directly and this will result in a communication failure, although they were using global addresses whose reachability should be maintained. This is not acceptable.

> Thus, the interface described in this document MUST only be used in deployments where the link between the MN and the MAG is point-to-point. The interface MUST NOT be used in deployments where the link between the MN and the MAG is shared and/or multi-access. Future specifications MAY define interfaces for use with shared and/or multi-access links.

TOC

5. Address Collision Considerations

As per the DNAv6 protocol (Narayanan, S., Kempf, J., Nordmark, E., Pentland, B., Choi, J., Daley, G., and N. Montavont, “Design Alternative for Detecting Network Attachment in IPv6 Networks (DNAv6 Design Alternative),” December 2009.) [I-D.ietf-dna-protocol], the MN will not execute Duplicate Address Detection (DAD)[RFC2462] (Thomson, S. and T. Narten, “IPv6 Stateless Address Autoconfiguration,” December 1998.) after a handoff within the same domain. This is because the MN will always receive the same subnet prefix in the RA and conclude that it did not change link. Hence there seems to be no need for executing DAD again. However, in NETLMM the link did changed. Because the link is point-to-point the only new entity on the link is the MAG, and it is possible that a collision occurs between link local addresses of the MN and the MAG (Note that no collisions are possible with global unicast address(es) since the subnet prefix has been uniquely assigned to the MN).
One solution to this issue would be that in a given domain, each MAG also defends link-local addresses of other MAGs in the domain. This would ensure that when the MN first attaches to the NETLMM domain and executes DAD it is able to pro-actively detect collisions that may happen with any MAG of the domain. Such a solution has however two drawbacks:

- Each MAG needs to know link-local addresses of all other MAGs in the domain.

- If SEND is used, each MAG also need to know private keys of all other MAGs since SEND requires a Neighbor Advertisement (NA) message defending an address to be signed with the SEND public key generating the CGA link-local address.

A much simpler solution is:

> All MAGs in a NETLMM domain MUST configure the same link-local address.

When SEND is used, that means that all MAGs share a single SEND public-private key pair, and hence a single link-local CGA.
Since all MAGs in a domain have the same link local address, if the MN executes DAD at his first attachment and concludes that there is no collision with the link-local address of the first MAG, a collision with any other MAG in the domain is impossible.

TOC

6. MN_ATTACH Function

The MN_ATTACH function is invoked by the MN whenever it attaches to a new MAG, and consists of the following sub-functions:

*MAG_GET_MN_ID: It provides the MAG with the identifier of the MN (MNID). This identifier MUST be securely bound to the MN, and the corresponding binding MUST be verifiable by the MAG. This triggers the MAG to authenticate the MN as the owner of this MNID. If authentication fails the MN_ATTACH function terminates with failure status, otherwise it continues.

*MAG_GET_HI: It provides the MAG with information to put in the Access Technology Type, Mobile Node Interface Identifier, and Handoff Indication fields.

*MN_GET_ADDR_PARMS: It provides the MN with IPv6 addressing configuration parameters, i.e. IPv6 subnet prefix(es) or global address(es). The MAG will then register the MN IPv6 subnet prefix(es) or address(es) with the LMA using the NETLMM protocol.

*MN_GET_DEFAULT_ROUTER: It provides the MN with the link local IPv6 address of its default router (e.g. the MAG).

*MAG_GET_MN_MCAST_GROUPS: It provides the MAG with the multicast group(s) that the MN previously joined (while attached to a previous MAG). This triggers the MAG to subscribe to the multicast tree(s) corresponding to the group(s) joined by the MN.

The MN_ATTACH function will typically be implemented by multiple protocols, some of them possibly non-IP protocols. The following subsections will describe in more details the MAG_GET_MN_ID, MAG_GET_HI, MN_GET_ADDR_PARMS, MN_GET_DEFAULT_ROUTER, and MAG_GET_MN_MCAST_GROUPS subfunctions, in particular what they achieve, and how.

6.1. MAG_GET_MN_ID Sub-function

TOC

The MAG_GET_MN_ID function provides the MAG with the identifier of the MN (MNID). This identifier MUST be securely bound to the MN, and the corresponding binding MUST be verifiable by the MAG [RFC4832] (Vogt, C. and J. Kempf, “Security Threats to Network-Based Localized Mobility Management (NETLMM),” April 2007.). This triggers the MAG to authenticate the MN as the owner of this MNID. If authentication fails the MN_ATTACH function terminates with failure status, otherwise it continues.

When the MN_ATTACH function includes a network access authentication protocol, such as EAP (Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, "Extensible Authentication Protocol (EAP)," June 2004.) [RFC3748], the Network Access Identifier (NAI) authenticated by the network access authentication protocol is a valid MN ID if it satisfies above constraints (freshness of authentication, verifiable by the MAG).
When the mix of protocols implementing the MN_ATTACH does not include a network access authentication protocol, or the network access authentication protocol does not provide a suitable MN identifier, or does not guarantee fresh authentication of the MN, an alternative authentication method based on the DNAv6 protocol (Narayanan, S., Kempf, J., Nordmark, E., Pentland, B., Choi, J., Daley, G., and N. Montavont, "Design Alternative for Detecting Network Attachment in IPv6 Networks (DNAv6 Design Alternative)," December 2009.) [I-D.ietf-dna-protocol] and the SEND protocol (Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," March 2005.) [RFC3971] MUST be used to authenticate the MN, as described below:

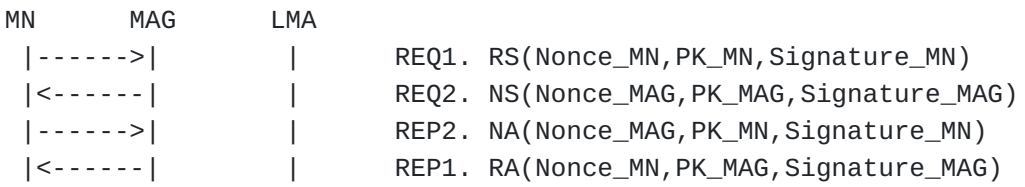

Figure 4: DNAv6/SEND based MNID authentication

*In step REQ1, after attachment occurs, and upon the occurrence of a Layer 2 link-up event notification, the MN initiates self-authentication to the MAG by sending an RS from its link local address to the link-scope all-routers multicast address, as per Section 5.2.5 of the DNAv6 protocol (Narayanan, S., Kempf, J., Nordmark, E., Pentland, B., Choi, J., Daley, G., and N. Montavont, "Design Alternative for Detecting Network Attachment in IPv6 Networks (DNAv6 Design Alternative)," December 2009.) [I-D.ietf-dna-protocol]. Since this RS is not sent from the unspecified address, it contains the MN SEND public key (PK_MN) in a CGA option, as per Section 5.1.1 of the SEND protocol (Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," March 2005.) [RFC3971]. This public key is used as an MN ID by the MAG.

*In step REQ2, after the MAG received from the MN an RS containing the MN ID (PK_MN) and the MN link local address, the MAG MUST

solicit the link-local address of the MN by sending an NS to the link-local address of the MN. This NS contains a fresh nonce (Nonce_MN) as per Section 5.3.3. of the SEND protocol (Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," March 2005.) [RFC3971].

*In step REP2, after the MN received from the MAG a NS containing a fresh nonce, it replies to the MAG with an NA containing the same fresh nonce as per Section 5.3.3 of the SEND protocol (Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," March 2005.) [RFC3971]. This NA is signed with the MN public key (i.e. the MN ID) as per Section 5.2.1 of the SEND protocol (Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," March 2005.) [RFC3971]. The MAG will verify 1) that the Nonce is fresh as per Section 5.3.4.1 of the SEND protocol (Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," March 2005.) [RFC3971], and 2) that the signature is valid for this public key as per Section 5.2.2 of the SEND protocol (Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)," March 2005.) [RFC3971]. If these verifications succeed, the MAG has successfully authenticated the MN as the owner of the MN ID.

*In step REP1, the MAG concludes the DNAv6 protocol (Narayanan, S., Kempf, J., Nordmark, E., Pentland, B., Choi, J., Daley, G., and N. Montavont, "Design Alternative for Detecting Network Attachment in IPv6 Networks (DNAv6 Design Alternative)," December 2009.) [I-D.ietf-dna-protocol] by sending to the MN an RA. This step is not part of the authentication of the MN and is shown here for completeness only. Note that a NETLMM exchange between the MAG and LMA MUST occur between REP2 and REP1 so that the MAG can obtain the proper Home Network Prefix to advertise toward the MN in REP1 (the Router Advertisement).

6.2. MAG_GET_HI Sub-function

TOC

During the MAG_GET_HI function the MAG MUST be given an indication of the link technology in use and MUST populate the Access Technology Type (ATT) appropriately. Usually the MAG will also obtain a link-layer address through link establishment or from the Router Solicitation message in step REQ1. For example, the IPv6CP Interface-Identifier option (S.Varada, Haskins, D., and E. Allen, "IP Version 6 over PPP," September 2007.) [RFC5072] may be negotiated during PPP establishment. The MAG SHOULD place the link-layer address in the Mobile Node Interface Identifier (MNIID) option of the Proxy BU. The MAG MUST set

the Handoff Indicator option to an appropriate value depending on the information it has from link establishment or context transfer signaling. If the MAG knows that there was a previous session for this MN using a different ATT and MNIID, then it SHOULD set the HI field to 1 (Attachment over a new interface). If the MAG knows that there was a previous session using a different ATT but the same MNIID, the MAG SHOULD set the HI field to 2 (Handoff between two different interfaces of the mobile node). If the MAG knows that there was a previous session using the same ATT and the same MNIID, the MAG SHOULD set the HI field to 3 (Handoff between mobile access gateways for the same interface). If the MAG has no information about previous sessions the MAG SHOULD set the HI field to 4 (Handoff state unknown). On subsequent Proxy BUs (sent to refresh the lifetime) the MAG SHOULD always set the HI field to 5 (Handoff state not changed (Re-registration)).

6.3. MN_GET_ADDR_PARMS Sub-function

TOC

The MN_GET_ADDR_PARMS function allows the MN to configure IP addresses. This can be achieved via different means, including:

- Stateless Address Autoconfiguration (SLAAC) (Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration," December 1998.) [RFC2462]: Allows the MN to configure both link local and global unicast address(es).

- Dynamic Host Configuration Protocol for IPv6 (DHCPv6) (Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," July 2003.) [RFC3315]: Allows the MN to configure global unicast address(es). Typically not used to configure link local unicast address(es).

- IP Version 6 over PPP (S.Varada, Haskins, D., and E. Allen, "IP Version 6 over PPP," September 2007.) [RFC5072]: Allows the MN to configure link local unicast address(es).

Whenever the MN attaches to a new MAG which is in the same domain as its old MAG, the MN_GET_ADDR_PARMS at the new MAG MUST must not change the address(es) that were configured by the MN at the old MAG.

TOC

6.4. MN_GET_DEFAULT_ROUTER Sub-function

The MN_GET_DEFAULT_ROUTER function provides the MN with its default router. This can be achieved via different means, including:

*Router Discovery as specified by the Neighbor Discovery Protocol (Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)," December 1998.) [RFC2461].

*IP Version 6 over PPP (PPPv6) (S.Varada, Haskins, D., and E. Allen, "IP Version 6 over PPP," September 2007.) [RFC5072].

Note that Dynamic Host Configuration Protocol for IPv6 (DHCPv6) (Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," July 2003.) [RFC3315] does not provide a default router. Instead, Router Discovery has to be used.

6.5. MAG_GET_MN_MCAST_GROUPS Sub-function

TOC

The MAG acts as a multicast router for the MN. The MAG_GET_MN_MCAST_GROUPS provides the MAG with the Multicast Address Listening state of the newly attached MN (this state might have been established while attached to a previous MAG). This triggers the MAG to subscribe to the multicast tree(s) corresponding to the source(s) the MN is listening to.
In many system architectures, this can be achieved by having, upon movement of the MN, the old MAG doing context transfer to the new MAG of the Multicast Address Listening state learned via MLDv2 (Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6," June 2004.) [RFC3810] messages.
When the deployment does not offer such context transfer, upon each new MN attachment the MAG MUST send a MLDv2 (Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6," June 2004.) [RFC3810] General Query to the link-scope all-nodes multicast address as per Section 5.1.15 and 7.1 of the MLDv2 protocol (Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6," June 2004.) [RFC3810]. A newly attached MN will then report its Multicast Address Listening state as per Section 6.2 of the MLDv2 protocol (Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6," June 2004.) [RFC3810], thus allowing the MAG to register to the appropriate multicast tree(s).

TOC

7. MN_DETACH Function

When an MN detaches from a MAG, the MAG has to deregister this MN with the LMA.
When the underlying link layer provides a reliable indication of an MN having detached from the MAG, the MAG MUST deregister the MN with the LMA upon reception of such indication.
When the underlying link layer provides no reliable indication of an MN having detached from the MAG, it is necessary to allow the MAG to detect an MN which silently detaches, or crashes, so that it can deregister the MN as a consequence. When such a link layer is used, the MAG MUST periodically execute Neighbor Unreachability Detection as per Section 7.3 of the Neighbor Discovery Protocol (Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)," December 1998.) [RFC2461] with each of the attached MNs, even though it has no traffic to deliver to the MN.
When an MN detaches from a MAG, the MAG MUST conclude that multicast address listening for the MN terminates for all the sources it was listening to.

---

8. Security Considerations

TOC

The security threats to the MN-AR protocol include:

- *Eavesdropping on the MN-AR exchange, where an attacker may learn information such as the MNID that might be confidential.

- *Malicious redirection of packets to a location other than that of the MN, where traffic can be observed more easily by an attacker.

- *Causing denial-of-service by de-registering the MN prematurely.

When the link layer incorporates strong authentication with a secure binding to the MN's link address, these threats are mitigated. A protocol such as EAP can be used in a mode where the NAI is obscured, obviating threat 1. EAP can also generate keys that get securely bound to native link encryption and authentication mechanisms. As long as all MAGs in a domain faithfully authenticate each MN then threats 2 and 3 are also mitigated.
In the absence of strong layer-2 security, the default protocol based on DNAv6 and SEND has somewhat weaker security properties. The MN_PK will be visible to anyone that can eavesdrop on the link. The protocol is vulnerable to a man-in-the-middle attack where the messages are relayed by an attacker to an MN that believes it is attached to a legitimate MAG. This could allow an attacker to redirect traffic.
Finally, if the layer-2 protocol is left vulnerable to spoofing an

attacker may be able to generate a link-down event which would cause the MAG to deregister the MN.

9. IANA Considerations

TOC

There are no IANA considerations.

10. Acknowledgments

TOC

As usual in the IETF, this document is the result of a collaboration between many people. The authors would like to thanks (in alphabetical order) James Kempf, Alexandru Petrescu, Fred Templin and Christian Vogt for discussion and/or comments that helped with first versions of this document.
Ian Chakeres contributed the reference network diagram.
Julien Laganier is partly funded by Ambient Networks, a research project supported by the European Commission under its Sixth Framework Program. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Ambient Networks project or the European Commission.

11. References

TOC

11.1. Normative references

TOC

[RFC2119]	Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC2461]	Narten, T., Nordmark, E., and W. Simpson, “Neighbor Discovery for IP Version 6 (IPv6),” RFC 2461, December 1998 (TXT, HTML, XML).
[RFC2462]	Thomson, S. and T. Narten, “IPv6 Stateless Address Autoconfiguration,” RFC 2462, December 1998 (TXT, HTML, XML).
[RFC3041]	Narten, T. and R. Draves, “Privacy Extensions for Stateless Address Autoconfiguration in IPv6,” RFC 3041, January 2001 (TXT).

[RFC3315]	Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, “Dynamic Host Configuration Protocol for IPv6 (DHCPv6),” RFC 3315, July 2003 (TXT).
[RFC3748]	Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, “Extensible Authentication Protocol (EAP),” RFC 3748, June 2004 (TXT).
[RFC3810]	Vida, R. and L. Costa, “Multicast Listener Discovery Version 2 (MLDv2) for IPv6,” RFC 3810, June 2004 (TXT).
[RFC3971]	Arkko, J., Kempf, J., Zill, B., and P. Nikander, “SEcure Neighbor Discovery (SEND),” RFC 3971, March 2005 (TXT).
[RFC3972]	Aura, T., “Cryptographically Generated Addresses (CGA),” RFC 3972, March 2005 (TXT).
[RFC4291]	Hinden, R. and S. Deering, “IP Version 6 Addressing Architecture,” RFC 4291, February 2006 (TXT).
[I-D.ietf-dna-protocol]	Narayanan, S., Kempf, J., Nordmark, E., Pentland, B., Choi, J., Daley, G., and N. Montavont, “Design Alternative for Detecting Network Attachment in IPv6 Networks (DNAv6 Design Alternative),” draft-ietf-dna-protocol-09 (work in progress), December 2009 (TXT).
[I-D.ietf-netlmm-proxymip6]	Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, “Proxy Mobile IPv6,” draft-ietf-netlmm-proxymip6-18 (work in progress), May 2008 (TXT).

11.2. Informative references

TOC

[RFC5072]	S.Varada, Haskins, D., and E. Allen, “IP Version 6 over PPP,” RFC 5072, September 2007 (TXT).
[RFC4830]	Kempf, J., “Problem Statement for Network-Based Localized Mobility Management (NETLMM),” RFC 4830, April 2007 (TXT).
[RFC4831]	Kempf, J., “Goals for Network-Based Localized Mobility Management (NETLMM),” RFC 4831, April 2007 (TXT).
[RFC4832]	Vogt, C. and J. Kempf, “Security Threats to Network-Based Localized Mobility Management (NETLMM),” RFC 4832, April 2007 (TXT).
[I-D.thaler-intarea-multilink-subnet-issues]	Thaler, D., “Issues With Protocols Proposing Multilink Subnets,” draft-thaler-intarea-multilink-subnet-issues-00 (work in progress), March 2006 (TXT).

Appendix A. Version history

TOC

A.1. -02 to -04

TOC

*-03 was a tombstone

*Pete McCann added as editor

*Various editorial fixes

*Modified description of REP1 to indicate that Proxy BU/BA must complete before

*Added description of how to set ATT, MNIID, and HI

A.2. -01 to -02

TOC

*revamped document structure to make it agnostic to attachment method (e.g. authentication, address-configuration, etc.).

*specified per-MN subnet prefix, and point-to-point link model.

*specified support for multicast.

*various editorial changes.

A.3. -00 to -01

TOC

*added DHCP access method including DHCP prefix delegation.

*added new network reference diagram.

*added definitions for NETLMM domain and NLMP.

*updated NA proxying method for colliding CGAs.

*added text on sending IP multicast messages to a Layer-2 unicast address.

*added new Section 4.5 text on MNID/IP address binding.

*added new Section 5. on multilink subnet issues.

*various editorial changes.

Authors' Addresses

TOC

	Julien Laganier
	DoCoMo Communications Laboratories Europe GmbH
	Landsberger Strasse 312
	Munich D-80687
	Germany
Phone:	+49 89 56824 231
Email:	julien.ietf@laposte.net
URI:	http://www.docomolab-euro.com/
	Sathya Narayanan
	School of Information Technology and Communications Design
	California State University, Monterey Bay
	3110, Inter-Garrison Road, Building 18, Room 150
	Seaside, CA 93955
	USA
Phone:	+1 831 582 3621
Email:	sathya@njit.edu
	Pete McCann
	Motorola
	MD 2240
	1301 E. Algonquin Rd
	Schaumburg, IL 60196
	USA
Phone:	+1 847 576 3440
Email:	pete.mccann@motorola.com

Full Copyright Statement

TOC

OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.