NETLMM
Internet-Draft
Expires: December 21, 2006

G. Giaretta
Telecom Italia
K. Leung
Cisco
M. Liebsch
NEC
P. Roberts
Motorola
K. Nishida
NTT DoCoMo Inc.
H. Yokota
KDDI Labs
M. Parthasarathy
Nokia
H. Levkowetz
Ericsson
June 19, 2006

NetLMM Protocol
draft-giaretta-netlmm-dt-protocol-00.txt

Status of this Memo

Copyright Notice

Abstract

This document specifies a network protocol that allows mobile nodes to move around in a localized mobility domain, changing their point of attachment within the domain, but without ever being aware at the IP layer that their point of attachment has ever changed, and maintaining seamless communication in the presence of such mobility events. It defines two protocol entities, a Mobile Access Gateway and a Local Mobility Anchor, and a set of messages between them, that together make these mobility events transparent to the mobile nodes at the IP layer.

Table of Contents

1. Introduction

This document specifies a protocol that allows nodes to move around in an access network attaching to various points of the access network while maintaining an IP layer configuration that does not change as the mobile nodes points of attachment change.

This protocol is not intended to solve all the problems of network-based IP mobility. Over the past decade many companies and forums have provided many, many staff years of research, development, and standardization to realize all IP mobile networks and no doubt many more years of effort are ahead to deliver improvements to realize all the envisioned usage of such technology. Such systems have added technology for specific link layers, and carrying IP packets over those link layers, support for AAA infrastructures, and mobile security to name a few. Challenges still lie ahead in the form perhaps of mobile QoS, power management and paging, and management of changing network conditions in the face of various mobility events.

This protocol is developed in response to the problem statement for network-based localized mobility [I-D.ietf-netlmm-nohost-ps] and this protocol attempts to satisfy the goals in the NetLMM goals document [I-D.ietf-netlmm-nohost-req]. It is intended basically to solve the problem of packet forwarding to nodes that change their point of attachment to the network without the use of any protocol support at the IP layer on the mobile node to support that mobility.

This document defines operation of the protocol for use in an IPv6 infrastructure and in support of IPv6 nodes, but the authors envision that with modifications the protocol could be productively used with an IPv4 infrastructure or to support IPv4 nodes. The document refers conscientiously to mobile nodes rather than mobile hosts because its operation is not limited in any way to host only support.

This protocol is similar and different to various IP mobility protocols the IETF has standardized in the past. It is similar in that it continues the tradition of maintaining address continuity to mobile nodes regardless of the fact that those nodes have changes their points of attachment in the network. It differs in that it does not require any operational changes in protocol operation between the mobile node and the network at the IP layer, in that it supports an infrastructure that embraces network controlled mobility operation, and in that its operation is limited in scope rather than globally applicable.

The differences between this protocol and previous mobility protocols are complementary rather than contradictory. It is quite possible to conceive of deployments in which mobile IP is used in a wide area to

provide mobility services across multiple interface types or separate local mobility domains while NetLMM is used within a single type of access network or a single local mobility domain to facilitate mobility without involving the mobile.

2. Terminology

In this document, several words are used to signify the requirements of the specification. These words are often capitalized. 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].

Mobility terminology in this document follows that in [RFC3753], with the added specification of some terms as they are used in this particular document:

IP Link
: A set of routers, mobile nodes, and wireless access points that share link broadcast capability or its functional equivalent. This definition covers one or multiple access points under one or several access routers. In the past, such a set has been called a subnet, but this term is not strictly correct for IPv6, since multiple subnet prefixes can be assigned to an IP link in IPv6.

Access Network (revised)
: An Access Network consists of following three components: wireless or other access points, access routers, access network gateways which form the boundary to other networks and may shield other networks from the specialized routing protocols (if any) run in the Access Network; and (optionally) other internal access network routers which may also be needed in some cases to achieve a specialized routing protocol.

Local Mobility (revised)
: Local Mobility is mobility over a restricted area of the network topology. Note that, although the area of network topology over which the mobile node moves may be restricted, the actual geographic area could be quite large, depending on the mapping between the network topology and the wireless coverage area.

Localized Mobility Management
: Localized Mobility Management is a generic term for protocols dealing with IP mobility management confined within the access network. Localized Mobility Management signaling is not routed outside the access network, although a handover may trigger Global Mobility Management signaling. Localized Mobility Management

protocols exploit the locality of movement by confining movement related changes to the access network.

Localized Mobility Management Protocol
A protocol that supports localized mobility management.

Intra-Link Mobility
Intra-Link Mobility is mobility between wireless access points within an IP Link. Typically, this kind of mobility only involves Layer 2 mechanisms, so Intra-Link Mobility is often called Layer 2 mobility. No IP link configuration is required upon movement since the link does not change, but some IP signaling may be required for the mobile node to confirm whether or not the change of wireless access point also resulted in a change of IP link. If the IP link consists of a single access point/router combination, then this type of mobility is typically absent.

Mobile Access Gateway (MAG)
A Mobile Access Gateway (MAG) is a router embedded in a device that terminates a specific link layer technology to which mobile nodes attach themselves. It terminates one end of the MAG of the connection to one or more Local Mobility Anchors and participates in the NetLMM protocol exchange.

Local Mobility Anchor (LMA)
A local mobility anchor (LMA) is a router that terminates connections to multiple Mobile Access Gateways, services mobility requests for mobile nodes moving within a NetLMM system, and participates in the NetLMM protocol exchange.

NetLMM Domain
A NetLMM domain is a set of multiple MAGs and a set of one or more LMAs interconnected within an access network that provides mobility operations for attached mobile nodes through the execution of the NetLMM protocol.

NetLMM Address
The invariant IP address on the MN inside the NetLMM domain. For IPv6 it is assumed that this is an invariant routable IP address with global scope.

NetLMM Network Prefix
The NetLMM Network Prefix (NNP) is the IPv6 link prefix of the NetLMM address.

Routing Tag
An opaque identifier that is signaled between MAGs and LMAs and that can be used to distinguish traffic inside packets when the contents inside those packets cannot be inspected due to some operations such as encryption or header compression. It could be used, for example, in a GRE key field if GRE tunneling were being used to distinguish internal packets destined for a mobile when the internal packets headers have been compressed.

3. Functional Entities

The principal functional entities in a NetLMM infrastructure are the Mobile Access Gateway (MAG) and the Local Mobility Anchor (LMA). There are other entities that will make up a mobile access network that are used to support various kinds of functionality (mobile nodes, AAA, routing, DNS, etc.) whose basic functionality may be used by the MAG and the LMA but whose operation is not changed in any way for the proper operation of the NetLMM protocol.

Include a diagram. The diagram should show:
1. multiple MAGs
2. multiple LMAs
3. their interconnectivity

Should it show?
4. mobile nodes moving around the edge
5. the possibility of link layer access technologies beneath MAGs

The Mobile Access Gateway
The Mobile Access Gateway (MAG) is a router that a mobile node is attached to as the first hop router in the NetLMM infrastructure. The MAG is connected to the mobile node over some specific link provided by a link layer but the NetLMM infrastructure is agnostic about the link layer technology that is used. Each MAG has its own identifier used in NetLMM protocol messaging between the MAG and the LMA. The important interfaces between link layer specific functions and the NetLMM function reside on the MAG. There are multiple MAGs in a NetLMM infrastructure.

The Local Mobility Anchor
The local mobility anchor (LMA) is a router that maintains reachability to a mobile node's address while the mobile node moves around within the NetLMM infrastructure. It is responsible to maintain forwarding information for the mobile nodes which includes a set of mappings to associate mobile nodes by their

identifiers with their address information, associating the mobile nodes with their serving MAGs and the relationship between the LMA and the MAGs. There may be one or more LMAs in a NetLMM infrastructure.

4. Protocol Overview

The protocol consists of two major phases, an initiation phase and an operational phase. During the initiation phase the MAGs and an LMA (or multiple LMAs) establish connectivity between them. Although this document describes this phase of the protocol operation as an initiation phase there is no restriction on the ability of adding new MAGs or new LMAs to a NetLMM infrastructure while other nodes are already in operation. During the operational phase the MAGs and LMAs provide mobile connectivity service to mobiles nodes that are attaching to the infrastructure, leaving the infrastructure, and moving around within the infrastructure.

It is not assumed that a MAG is associated with only a single LMA. If there exists multiple LMAs in a NetLMM Domain, each MAG would most likely be associated with, and potentially communicate with all the LMAs rather than only a single LMA.

The NetLMM infrastructure uses 6 messages to establish and maintain associations between the MAGs and the LMAs: Association Request, Associate Reply, Disassociation Request, Disassociate Reply, Heartbeat, and Heartbeat Ack.

A MAG associates itself with an LMA by sending an Association Request message that includes its MAG ID and the supported data forwarding modes (such as IPv6-in-IPv6). In response the LMA creates an association with the MAG and populates state information about the association. The LMA responds, providing its LMA ID and an agreed upon data forwarding mode to the MAG. The MAG can undo the relationship with the LMA through sending a Disassociation Request, to which the LMA responds with a Disassociate Reply. Heartbeat messages are sent between the MAG and LMA to determine the current status of the reachability of the other entity. All of these messages may be sent optionally over an IPsec connection if additional security is desired.

The NetLMM infrastructure uses 14 messages to manage the attachment, departure, mobility, and other activities of mobile nodes within the infrastructure: LMA Allocation Request, LMA Allocation Reply, Location Registration and Acknowledge, Location Deregistration and Acknowledgment, Routing Setup and Acknowledgement, Routing Removal and Acknowledgement, MN Address Setup and Acknowledgement, MN Address

Removal and Acknowledgement.

When a mobile node is to receive service, a policy decision point entity may send an LMA Allocation Request to the LMA creating state for the mobile node. The LMA allocation request authorizes service for a particular Mobile Node ID. This message may contain policy information for the mobile node. The LMA acknowledge the request with an LMA allocation reply. It is possible that an LMA is configured to authorize service for any mobile node and in such a situation the LMA Allocation Request and reply messages are not necessary.

When a mobile node connects to the NetLMM infrastructure, it first needs to configure an address. Whether it is using stateful or stateless address configuration, the serving LMA needs to be involved in the address allocation process. So when a mobile node connects, the MAG sends a location registration message to the LMA containing its own ID and the Mobile Node's ID. The LMA responds to the message with a a Location Registration Ack message that includes the NetLMM prefix that the MAG is to use in its router advertisement toward the Mobile Node. The MAG in turn sends a router advertisement to the attached mobile. Once address configuration is complete (through either stateful or stateless address configuration) the MAG registers the mobile node's address with the LMA by sending an MN Address Setup message to the LMA including the MAG's ID, the MN's ID, the NetLMM address, and a tunnel ID. The LMA creates forwarding state for packets in response to this message and sends a MN address reply to the MAG acknowledging the packet setup. The MAG, in receiving a successful MN Address Setup Reply, creates forwarding state for packets destined to the mobile node.

When a mobile node then leaves the NetLMM infrastructure, the MAG sends a Location Deregistration message to the LMA including the Mobile Node's ID and the MAG's ID. The LMA cleans up all state for the mobile node identified in the message and sends a Location Deregistration Acknowledgement message.

It is also possible for the LMA to remove a mobile node from the network. This could be done for a number of policy specific reasons in the network. The same two messages are used, Location Deregistration and Local Deregistration Acknowledgement, but they are initiated by the LMA and acknowledged by the MAG in this case. The MAG disconnects the mobile and removes all mobile state in response to this message.

When a mobile node moves from one MAG to another MAG, the new MAG (nMAG) sends a Location Registration Message to the LMA with the MAG ID and the MN ID. The LMA responds by sending a Routing Setup

message to the nMAG that includes the MN ID, the MAG ID, the LMA ID, NetLMM addresses. The new MAG acknowledges this information with a Routing Setup Ack message and the LMA responds with a Location Registration Ack message containing the NetLMM prefix that the nMAG uses in the router advertisement for the MN.

The mobile node can at any time configure new IP addresses for itself using stateless address auto-configuration and this operation is supported in NetLMM. When the mobile issues a new address DAD request, the MAG sends a MN Address Setup Request to the LMA with the mobile node's ID and the new NetLMM address. The LMA validates the usability of the address, updates forwarding state, and acknowledges the request with the MN Address Setup Reply message to the serving MAG which then replies to the MN DAD operation. The means through which the NetLMM infrastructure knows that the mobile node is attached, leaving, or moving is beyond the scope of this protocol specification. It is envisioned that this information is largely access network specific and that the MAG uses an API to trigger much of the operation described herein.

5. Message Types

This document defines a set of control messages and options for the NetLMM protocol. The control messages are carried by the User Datagram Protocol, using a well known port number, as described in Section 5.11. The messages are presented in this section, and the message format and option formats are defined later, in Section 7.

5.1. LMA Allocation Request / Reply messages

The LMA Allocation Request message is used to allocate an LMA for the MN that is initially attached to the network and to validate the Location Registration message coming later from the MAG to which the MN is attached. This message containing the MN ID is sent to the selected LMA and may come from various nodes such as the MAG to which the MN is attached or the PDP that is involved in the authentication of the MN. The LMA Allocation Request is optional and the LMA may be allocated by other means (e.g., static allocation) and can be configured to serve the MN without this message.

The LMA Allocation Reply message is sent from the LMA to the source of the LMA Allocation Request message to notify the status of the request (success or error code).

5.2. Association Request / Reply messages

The Association Request is used to set up the control and data plane

relationship between the MAG and LMA. This message is sent from the MAG to the LMA in the MAGs initiation phase, before it enters the operational phase and handles MN Location Registration and Routing Setup. The message contains the sender's ID, its functional capabilities and supported data forwarding modes. The data forwarding mode specifies the tunnel method of the data plane (e.g., IP-in-IP). The tunnel between the MAG and LMA is bidirectional, which is achieved by establishing two unidirectional tunnels in opposite directions.

The Associate Reply message is sent from the LMA to the MAG to notify the status of the request (success or error code). If the request is successful, the receiver of the request also sends its capabilities (e.g., the receiver's ID, agreed-upon data forwarding mode, etc.) on this message.

5.3. Disassociation Request / Reply messages

The Disassociation Request message is sent from the MAG to the LMA or vice versa to tear down the control and data plane relationship between them. This message contains the MAG ID and LMA ID.

The Disassociate Reply message is sent from the LMN to the MAG or vice versa to inform the status of the request (success or error code).

5.4. Location Registration / Ack messages

The Location Registration message is sent from the MAG to the LMA when the MAG detects the MN having accessed the network without an IP address. By this message, the MAG and LMA create the mobility state of the MN. The IP address of the MN is not known at this point and the MN Address Setup message must follow to update the mobility state and to set up the routing for the MN. This message contains the MN ID, MAG ID and LMA ID.

The Location Registration Ack message is sent from the LMA to the MAG to acknowledge the receipt of the Location Registration message. If the registration is successful, the LMA sends the NetLMM prefix on this message, which in turn is used for the Router Advertisement sent by the MAG.

5.5. Location Deregistration / Ack messages

The Location Deregistration message is sent from the MAG to the LMA or vice versa to delete the mobility state of the MN. The MAG sends this message when the MN is detected to have moved away. On the other hand, the LMA sends this message when it determines that the MN

is at a new location. This message contains the MN ID, MAG ID, LMA ID and the requested revocation time.

The Location Deregistration Ack is sent back to the source of the Location Deregistration message to acknowledge the receipt of the Location Deregistration message. This message contains the agreed revocation time.

5.6. MN Address Setup / Ack messages

When the MAG determines that the MN has obtained its NetLMM address by for example the neighbor advertisement (the DAD procedure) or the DHCP server, the MAG sends he MN Address Setup message to the LMA to update the mobility state and to create a routing entry for the MN on the LMA. This message contains the MAG ID, LMA ID and one or more MN ID(s) and NetLMM address(es) assigned to the MN(s). Optionally, the Routing Tag(s) for the MN(s) may be included.

The MN Address Setup Ack is sent from the LMA to the MAG to acknowledge the receipt of the MN Address Setup message and to notify the MAG if the NetLMM address(es) contained in the MN Address Setup message are accepted (the DAD procedure). If the Routing Tag is contained in the Routing Setup message, a corresponding Routing Tag is returned by the Routing Setup Ack message.

5.7. MN Address Remove / Ack messages

When the MAG determines that one of the MN's NETLMM address(es) is no longer valid or used, the MAG sends the MN Address Remove message to the LMA to delete the corresponding mobility state and routing entry in the LMA. This message contains the MN ID, MAG ID, LMA ID and corresponding NETLMM address.

The MN Address Remove Ack is sent from the MAG to the LMA to acknowledge the receipt of the MN Address Remove message.

5.8. Routing Setup / Ack messages

When the MN moves from the previous MAG to the new MAG (inter-MAG handover), the LMA, which already has the mobility state for the MN, sends the Routing Setup message to the new MAG in response to the Location Registration message from the new MAG. This message is used to update the routing on the new MAG and contains the MN ID, MAG ID, LMA ID and one or more NetLMM address(es) assigned to the MN. Optionally, the Routing Tag for the MN may be included.

The Routing Setup Ack message is sent from the MAG to the LMA to acknowledge the receipt of the Routing Setup message. If the Routing

Tag is included in the Routing Setup message, the corresponding Routing Tag is returned by the Routing Setup Ack message.

5.9. Routing Remove / Ack messages

When the LMA determines that one or more NetLMM address(es) is/are no longer valid by for example the DHCP server, the LMA sends the Routing Remove message to the MAG to delete the corresponding routing entry/entries on the MAG. This message contains the MN ID, MAG ID, LMA ID and NetLMM address(es) of the MN.

The Routing Remove Ack is sent from the MAG to the LMA to acknowledge of the receipt of the Routing Remove message.

5.10. Heartbeat / Ack messages

The Heartbeat message is sent from the MAG to LMA or vice versa to obtain the connectivity status. This message contains the MAG ID, LMA ID and the heartbeat number that is separated from the message sequence number.

The Heartbeat Ack is sent back from the node that received the Heartbeat message to its peer. This message contains the MAG ID, LMA ID and the corresponding heartbeat number.

These messages are suppressed when there is data traffic between the two nodes.

5.11. Message Transport

The new NetLMM control messages defined in this document are carried by the User Datagram Protocol RFC 768 [RFC0768] using well known port number TBD (assigned by IANA).

The message sender SHOULD include a non-zero UDP Checksum. The recipient of the message MUST process and check the UDP checksum. A Zero checksum SHOULD be accepted by the recipient.

The sender and initiator of a message exchange MUST use the following UDP ports:

* Source Port: variable

* Destination Port: TBD (Assigned by IANA)

In case the recipient of a NETLMM message has to reply, the following UDP ports MUST be used:

* Source Port: variable

* Destination Port: Copied from the source port of the received message.

5.12. Message Optimization

In order to minimize routing establishment delays, e.g., handover times, it may be important to achieve routing setup or routing changes with an absolute minimum number of messaging roundtrips between the MAG and the LMA. However, given the messages described earlier in this section, there are several cases where 2 messaging round-trips are needed in order to complete a routing setup.

Future versions of this document will propose optimizations which reduce the number of messages that must be sent in order to achieve routing setup or routing changes. It is however deemed important to have agreement on the base functionality and the base messages before such optimizations are discussed.

6. Protocol Extensibility

NetLMM consists of a set of new control messages, described in Section 5 in this document. Up-to-date values for the message types are maintained in the online IANA registry of assigned numbers.

NetLMM defines a general extension mechanism using options to allow optional information to be carried in the control messages. The options are encoded in the Type-Length-Value format, and are described in detail in Section 7.2. The options carry additional information used for processing the message. Up-to-date values for the option types are maintained in the online IANA registry of assigned numbers.

The Type field in the NetLMM option is split into two ranges: Type values of 0 through 127 (inclusive) for not skippable options and 128 through 255 (inclusive) for skippable options. The recipient of a message with an unrecognized non-skippable option MUST silently discard the message. Otherwise, if no unrecognized non-skippable options are found, the message MUST be processed with any unrecognized skippable option bypassed (i.e. move to next option using the Length field of the unrecognized option) during processing by the receiver. The Sub-Type field provides efficient use of the option type numbering space.

Format:

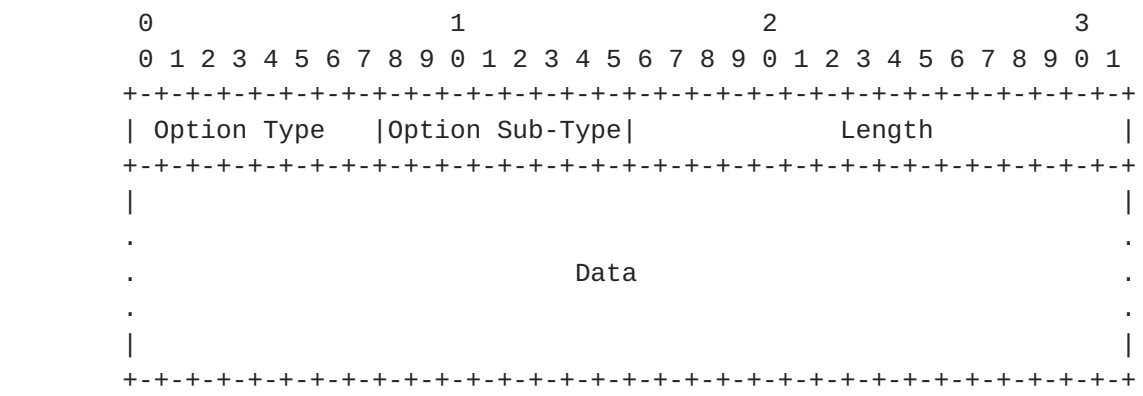

Option Type
This field identifies the particular type of option serving a specified function.

Option Sub-Type
This field indicates the sub-type of the option, and provides for up to 256 related Option Sub-Types with the same Option Type field.

Length
The value represents the length of the "Data" portion of the option, in unit of octets.

7. Message and Message Option Formats

7.1. Message Formats

An NetLMM message consists of a header followed by one or more options. The Message header is common and messages are distinguished by Type field in the header. NetLMM options are in TLV format and 8byte aligned, though Type field divided into two parts; Option Type and Option Sub-Type. The length field indicates the exact length of the following value fields in units of 8 octets. All option payloads whose length is not a unit of 8 octets must be padded to the correct alignment.

7.1.1. Message Header

All NetLMM messages start with the following common header. Parameters for each message are contained in the option format (see following sections).

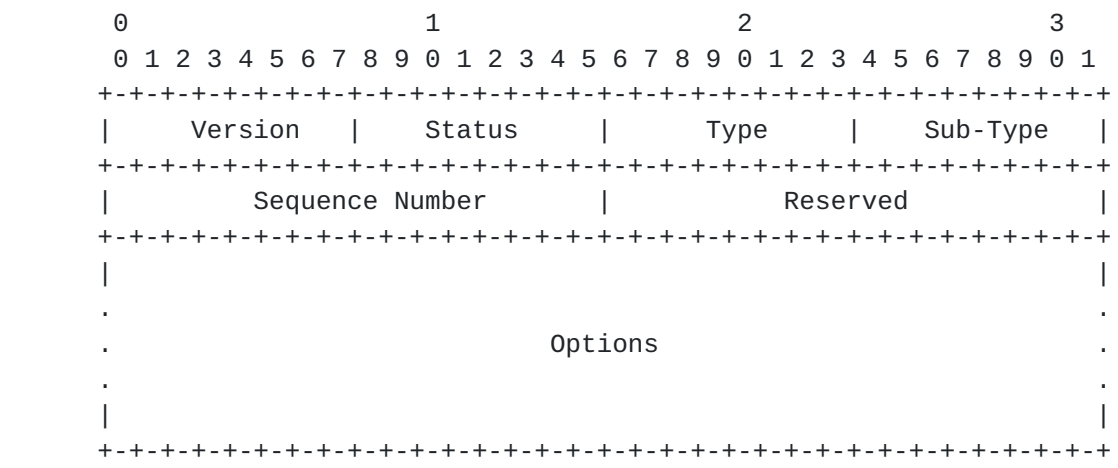

Version
An 8-bit number, indicating the NetLMM protocol version. The version of the NetLMM protocol specified in this document is 1.

Status
An 8-bit number, which indicates the status of the message. When used for request messages, e.g., Location Registration, the status code is normally set to 0, indicating 'Not Applicable'. When used for reply and acknowledgement messages, the status code indicates the result of processing the associated request message. The following values are defined by this document:

0: Not Applicable (N/A)
The status code is not applicable for the message. This is normally the case for request messages.

1: Success Acknowledgement
The associated request message was successfully processed.

2: Administratively Prohibited
An action was refused due to administrative policy reasons.

3: Lack of Resources
The resources needed to provide the requested service was not available.

4: Unauthorized MN
Used by the LMA in response to Location Registration or Routing Setup, to notify the MAG of the MN not being authorized for service.

6: Duplicate Address
Used by the LMA when an MN Address Setup contains an IP address that is duplicated in the same NetLMM domain. The specific invalid addresses MUST be specified in Address Options with Sub-Type TBD, Duplicate Address.

5: Invalid Address
Used by the LMA when an MN Address Setup contains an IP address that is invalid. The specific invalid addresses MUST be specified in Address Options with Sub-Type TBD, Invalid Address.

7: Over IP Address Limit
Used by the LMA on receipt of a Routing Setup or MN Address Setup message, if the maximum number of IP addresses allowed for a MN has been exceeded.

8: Invalid Tunnelling Method
The proposed tunnel method is not supported or unavailable.

9: Invalid Message
The NetLMM Request Message was invalid or malformed.

10: Already Associated
The LMA already had the requesting MAG listed as associated.

Type
8-bit indicator, shows NetLMM message types. The message types are specified in section 5. The values for messages are defined as follows;

0: LMA Allocation Request/Reply

1: Association Request/Reply

2: Disassociation Request/Reply

3: Location Registration/Ack

4: Location Deregistration/Ack

5: MN Address Setup/Ack

6: MN Address Remove/Ack

7: Routing Setup/Ack

8: Routing Remove/Ack

9: Heartbeat/Ack

Sub-Type
8-bit indicator, this indicator is Type dependent field and can use to add extra information for the message.

Sub-types defined in this document, valid for all message types defined in this document:

0: Request
This is the Request message of the given type, with semantics and format as specified in this document. When message variations with other semantics or formats are required in the future, new subtypes should be allocated for them.

1: Ack/Reply
This is the Acknowledgement (or Reply) message to be sent in response to request messages of the given type. It is seen as less likely that it will be necessary to allocate new sub-types for new Ack/Reply messages, but there is no restriction on doing so.

Sequence Number
16-bit length, used to ensure the correspondence of the request and ack/reply pair of the same message between the MAG and the LMA. The sequence number is exchanged between given MAG and LMA, and configured when MAG has joined to a NetLMM domain through the exchange of Association Request/Reply messages. Sequence Number comparisons MUST be performed modulo 2**16, i.e., the number is a free running counter represented modulo 65536. A Sequence Number in a received message is considered less than or equal to the last received number if its value lies in the range of the last received number and the preceding 32768 values, inclusive. For instance, if the last received sequence number was 15, then messages with sequence numbers 0 through 15, as well as 32783 through 65535, would be considered 'less than or equal'.

Reserved
16-bit field reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Options
8 byte aligned field, and can add multiple options. See following sections for options.

7.2. Options

7.2.1. ID Option

The ID option carries various types of identifiers. All messages related to a specific MN must include an ID option providing the MN ID. Multiple ID options can be included in an message. In addition to that for the MN ID there might for instance be ID options for the MAG ID and for the LMA ID in a Location Registration. For the purpose of the ID option, the ID itself is viewed as an octet sequence, but to avoid ID collisions, the ID is prefixed with an ID type. An example for the MN ID is a NAI [RFC4282].

This option can also contain Routing Tag which is an identifier specified by each tunnel method for data transport. The existence or use of the Routing Tag is also dependent on the tunnel method to use.

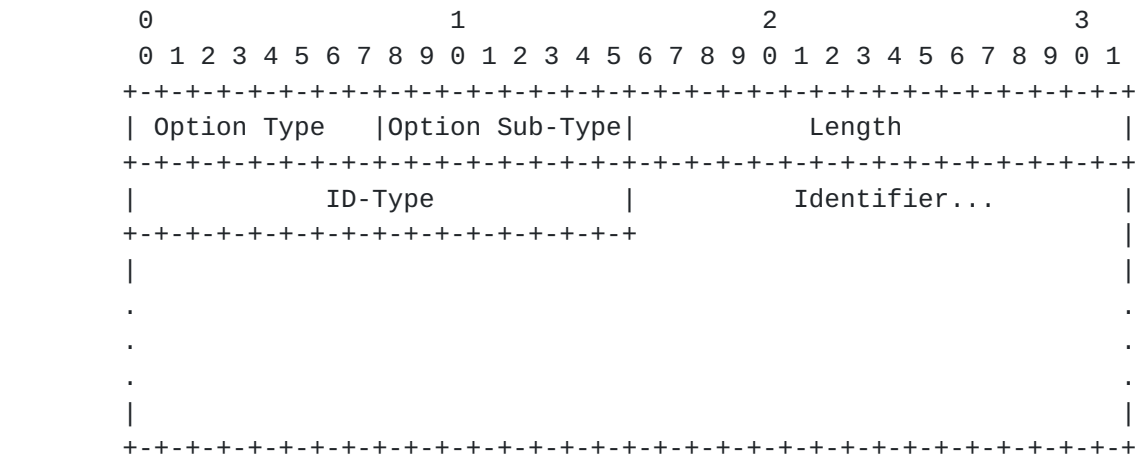

Option Type
0

Option Sub-Type
This field indicates what the ID in this option refers to. It is expected that additional Sub-Types may be defined in the future.

0: MN ID

1: LMA ID

2: MAG ID

3: Routing Tag

Length
Variable. The value indicates exact length of the option in octets, excluding the Type, Sub-Type and Length fields.

ID-Type
This field indicates the type of ID carried in the remainder of the option.
*** Note: Check whether there already exists an applicable IANA registry for ID types which we could use here ***

0: 128-bit opaque cryptographically generated identifier (such as proposed ORCHID IDs

1: NAI according to [RFC4282]

2: Ethernet MAC Address

3: IPv6 Address

Identifier
This is a variable-length octet sequence, which is expected to hold an identifier of the type indicated by the ID-Type field.

7.2.2. NetLMM Address Option

This option conveys NetLMM IP address and the prefix that LMA advertise for MNs. The NetLMM address can be both IPv4 and IPv6. This option can also inform LMA prefix only.

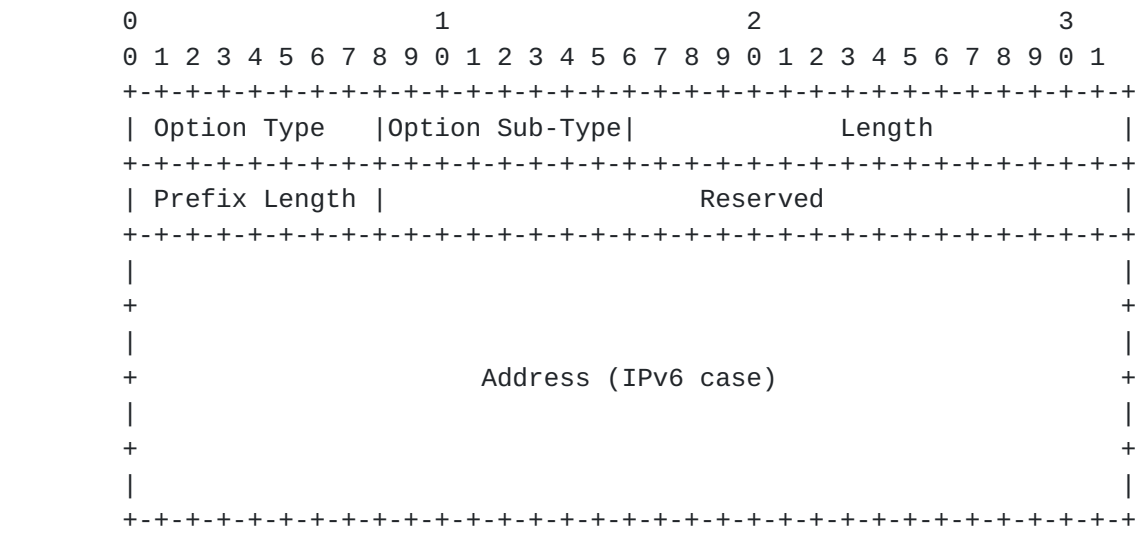

Option Type

1

Option Sub-Type

0: Indicates that the data in the Address field is an IPv6 address or address range. If Prefix Length is 128 this is an address, otherwise it is an address range with span determined by the given prefix length.

1: Indicates that the data in the Address field is an IPv4 address.

2: Indicates that the data in the Address field is an IPv6 network prefix information.

3: Indicates that the data in the Address field is an IPv4 network prefix information.

4: Indicates that the address in the Address field is a duplicate IPv6 address.

5: Indicates that the address in the Address field is a duplicate IPv4 address.

6: Indicates that the address in the Address field is an invalid IPv6 address.

7: Indicates that the address in the Address field is an invalid IPv4 address.

Length
Variable. The value indicates exact length of the option in octets, excluding the Type, Sub-Type and Length fields.

Prefix Length
Variable. The value indicates the number of leadings bits in the Address that are valid (as a network prefix).

If the Sub-Type value indicates an invalid or duplicate address, this field must be set to zero when sending, and ignored on receipt.

Reserved
16-bit field reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Address:
If the Option Sub-Type is 0 or 2, the value in the IPv6 address appears, while the type is 1 or 3, the value in the IPv4 address is presented.

7.2.3. Data Transport Option

This option contains data transport methods capabilities that the MAG or LMA has. This option is used by Association request/reply messages to negotiate the data transport method between MAG and LMA. Multiple methods can be contained in the field with the order of preference. The mandatory transport method is IPv6-in-IPv6, which must be listed.

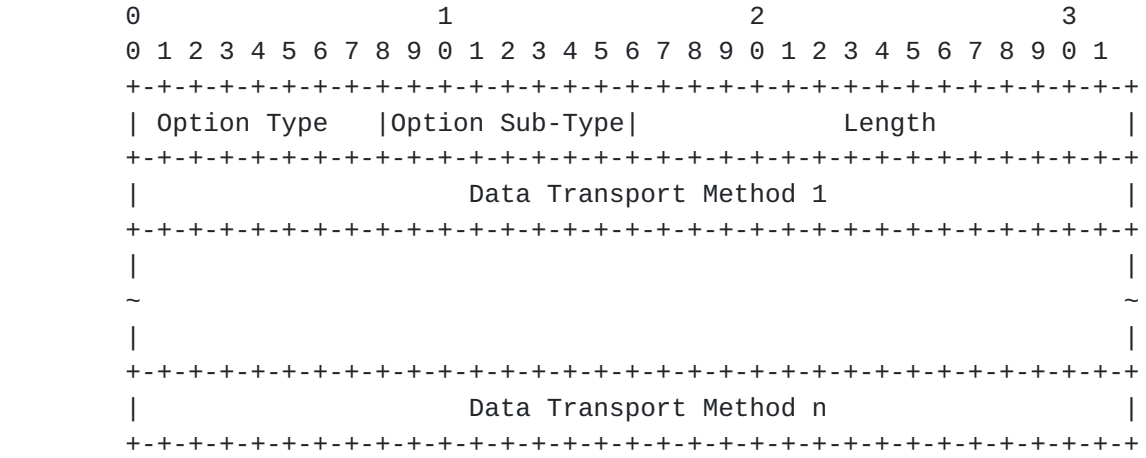

Option Type
2

Option Sub-Type
This field is reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Length
Variable. The value indicates exact length of the option in octets, excluding the Type, Sub-Type and Length fields. The number of Data Transport Method fields is equal to the Length divided by the size, in octets, of the Data Transport Method field.

Data Transport Method
Present the data transport methods available at the sender of the Association request/Reply message, i.e., LMA and MAG. The methods presented in this option are listed in order of the preference.

0: IPv6-in-IPv6

1: GRE

2: IPv4-in-IPv6

7.2.4. Revocation Timer Option

This option indicates the length, by milliseconds, to keep the forwarding state of a given MN at previous MAG in handover scenario. This option is included in Location Deregistration and its acknowledgement. The Revocation Time in the Location Deregistration Ack is copied from that presented in the receipt option of the Location Deregistration. If the MAG can not keep the MN state as LMA has requested for some reason, MAG can input preferable Revocation time or error code instead of copying original value.

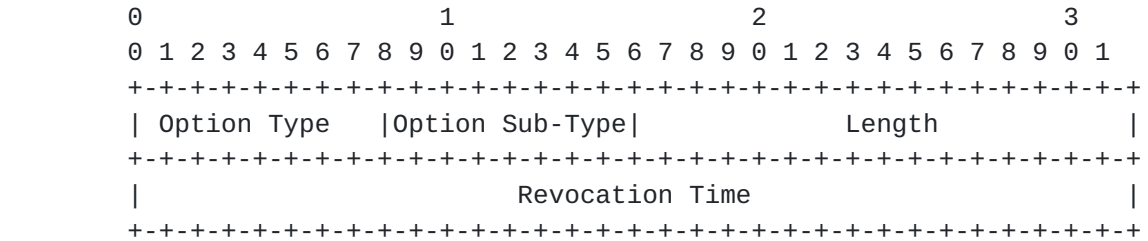

Option Type
3

Option Sub-Type
This field is reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Length
4. The value indicates exact length of data shown in Revocation Time field. Shown in octets, network byte order

Revocation Time
Variable. Indicate the preferred delay to delete the MN forwarding state from previous MAG to LMA in handover. The value shows in millisecond unit.

7.2.5. Timestamp Option

This option contains the timestamp value in the format of NTP timestamp, and records the time when the message is sent. This option can be attached to any message and used to detect an overtaking message in race condition by comparing the timestamp values of messages. Especially in handover scenario, if the network suffers from a sudden propagation delay for some reason or the MN moves rapidly between MAGs, the timestamp may be used to facilitate in-order messages processing regardless of message arrival order. The use of this option is network administrator dependent, and needs some of the time distribution methods, (e.g., NTP or GPS time synchronization system), with the high accuracy enough to support fast moving MNs.

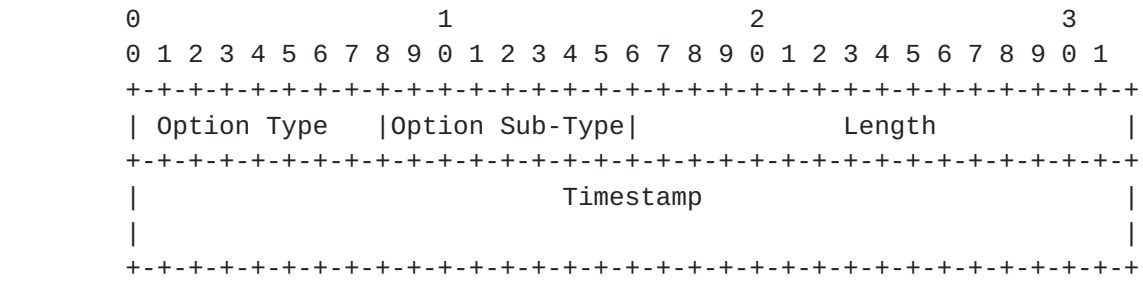

Option Type
4

Option Sub-Type
This field is reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver.

Length
8. The value indicates exact length of Timestamp field Shown in octets, network byte order.

Timestamp
Follows the 64bit length format of the NTP timestamp [RFC4330]

7.2.6. Vendor Specific Option

This option can be used by any vendor or organization that has an IANA-allocated SMI Network Management Private Enterprise Code. Details of the meaning of value field is entirely up to the defining vendor or organization.

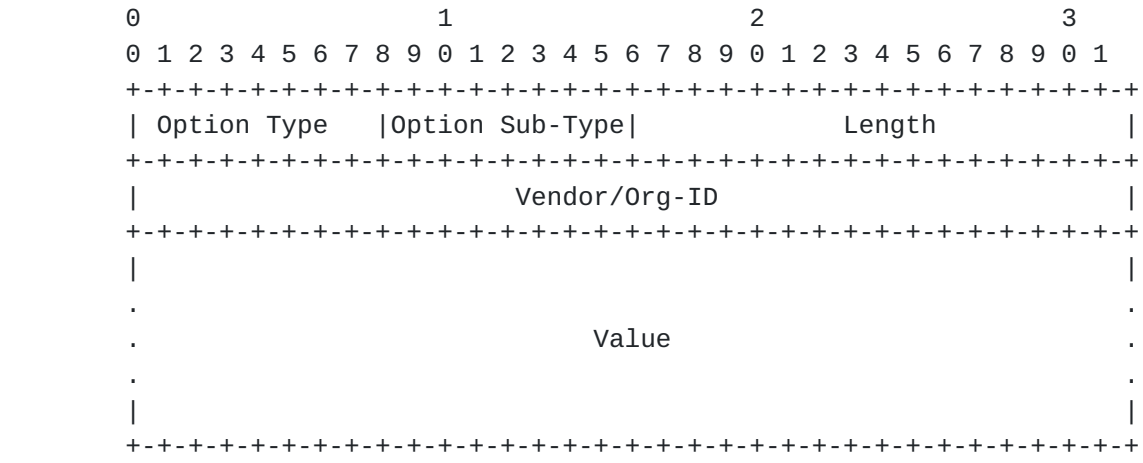

Option Type
5

Option Sub-Type
This field is reserved for future use. The value MUST be initialized to zero by the sender, and MUST be ignored by the receiver. This field may not be assigned any value different from zero by the organizations using the option; only the Value field may be freely used.

Length
Variable. The value indicates exact length of the option in octets, excluding the Type, Sub-Type and Length fields.

Vendor/Org-ID
The high-order octet is 0 and the low-order 3 octets are the SMI Network Management Private Enterprise Number [RFC2578], [ENTERPRISE-NUM], of the Vendor in network byte order.

Value
Variable. Details defined by individual Vendors / Organizations.

8. Protocol Specification

8.1. Mobile Access Gateway Operation

8.1.1. Conceptual Data Structures

Each MAG MUST maintain a NetLMM Routing Cache and an LMA List.

Each MAG Routing Cache entry conceptually contains the following fields for each attached MN:

* MN ID of the attached MN. This identifier is learned from the attach procedure and is used from the MAG to identify the attached MN in the Location Registration message, which is sent to the selected LMA.

* One or more global IP addresses of an attached MN. Each IP address is learned from an LMA through the Routing Setup message from an LMA or my means of local operation. According to the context of the received message or local indication, an IP address is set up, updated or removed from the Routing Cache.

* Routing Tag for the attached MN. Creating the entry and use of this tag is optional and might support identification of the MN in the data plane at the LMA and the MAG. In case the LMA and MAG specify asymmetric tags for the MN, this field MUST draw a distinction.

* LMA ID of the LMA serving an attached MN. The serving LMA and its LMA ID is learned from the LMA selection policy, which is out of scope of this specification.

Each MAG MUST maintain an LMA List, which identifies all LMAs with which the MAG is associated. The LMA List is used to perform heartbeat tests and to map an LMA ID to the associated LMA's IP address(es). The LMA List also supports the procedure of bulk de-registrations at all or a subset of LMAs.

The LMA List conceptually contains the following fields for each LMA

entry:

* LMA ID of the LMA.

* One or more IP address(es) of the LMA. The LMA's IP address information is learned through the LMA selection policy, which is out of scope of this specification. Availability of multiple LMA IP addresses could support operation of multi-homed LMAs. Details about how to handle multiple LMA IP addresses is out of scope of this specification.

* Selected forwarding approach for the association with an LMA. This field is needed in case a single forwarding approach is set up for the association with an LMA.

* Forwarding capabilities of the LMA. In case the LMA attaches a list of its own capabilities to the Associate Reply message, the MAG SHOULD store them in this field of the LMA List.

Each MAG MIGHT maintain a list of available LMAs. Such a list can support the LMA selection procedure and the MAG's association procedure.

The list of available LMAs comprises conceptually the following fields for each LMA:

* LMA ID of the LMA.

* One or more IP address(es) of the LMA. Availability of multiple IP addresses could support the operation of multi-homed LMAs. Details about how to handle multiple IP addresses is out of scope of this specification.

8.1.2. Processing NetLMM Headers

* The Type and Sub-Type fields MUST have a known value (Section 7.1.1). Otherwise, the node MUST discard the message and issue a an Error message with Status field set to 7 ("INVALID MESSAGE").

8.1.3. Processing NetLMM Messages

8.1.3.1. Association Procedure

Each Mobile Access Gateway sends an Association Request message in order to set up the control and data plane relationship with a given local mobility anchor. The actual trigger for this message is out of scope of this document and may depend on network configuration

peculiarities. For example, the Association Request message may be sent during the MAG start up procedure.

The Association Request message MUST include:

* the MAG identifier included in a NetLMM ID option. This identifier is used by the peer to identify the MAG and is included in all subsequent messages.

* the MAG's capabilities in terms of support of data transport methods included in a NetLMM Data Transport Option. The MAG MUST insert in this option all possible tunneling methods that can be used with the peer LMA. Based on configuration, it is possible that some tunneling methods are used only with some LMAs: in this case, the Association Request message MUST contain only the tunneling methods that are administratively permitted with that specific LMA.

When sending an Association Request, the MAG MAY create a tentative entry in its LMA List, including the LMA ID, IP address of the LMA and the proposed forwarding capabilities. However it may be that the MAG does not know these data during the association procedure: in this case, it does not create any tentative entry in the LMA List.

In order to complete the NetLMM association, the MAG MUST receive an Association Reply from the peer LMA with STATUS 1. In this case, the MAG MUST create an entry in its LMA List (or update the tentative entry created earlier), with the messages sent by the LMA in the Association Reply. The MAG MUST also update the forwarding method pre

8.1.3.2. Disassociate Procedure

The Disassociation Request can be sent both by the MAG and by the LMA in order to tear down the control and data plane relationship with the LMA. The event that triggers this message is out of the scope of this specification; for example, the MAG may send a Disassociation Request to all the LMAs present in its LMA List just before shutting down.

In case the Disassociate Procedure is initiated by the MAG, the MAG MUST include an ID Option with the its identity in the Disassociation Request. When sending the Disassociation Request, the MAG MAY set the LMA entry related to the specific LMA as tentative. When it receives a Disassociation Reply with Status 1 "SUCCESS", the MAG MUST delete the correspondent entry in its LMA List.

In case the Disassociate Procedure is initiated by the LMA, when the

MAG receives a Disassociation Request message, it MUST validate it. If it is correct, it MUST delete the related entry in its LMA List and send a Disassociate Reply with Status 1 "SUCCESS". As in all NetLMM messages, the MAG MUST include the ID option with its identity.

8.1.3.3. MN network access procedure

When a new MN attaches to the network, the Mobile Access Gateway receives an indication. This indication can be received by very different means (e.g., L2 mechanisms, AAA infrastructure) that are out of scope of this specification. In any case, regardless how this is accomplished, the MAG receives a MN_Access_Network API that carries the MN identifier (e.g., MAC address of the MN, NAI) and the LMA identifier.

Upon the API notification, the MAG MUST send a Location Registration message to the LMA including three different ID options, containing its own identity, the identity of the MN and the identity of the LMA. How the MAG resolves the LMA ID received in the API into the LMA IP address is out of scope of this specification and is part of the NetLMM bootstrapping procedure. Viable options are pre-configuration or DNS resolution (in case the LMA ID is the FQDN of the LMA). In case there is only one LMA in the local domain, this issue does not exist.

If the location registration is successfully performed, the MAG receives a Location Registration Acknowledge message from the LMA with Status value 1 "SUCCESS". This message includes also a NetLMM Network Prefix that the MAG MUST advertise to the MN. For this purpose, the MAG takes the prefix parameters from the NetLMM Address Option in the Location Registration Acknowledge and creates a Router Advertisement with a correspondent Prefix Option. It then sends a unicast Router Advertisement to the MN. (Note: how about the parameters that are carried in a Prefix Option but are not present in the NetLMM Address Option? e.g., lifetimes?). In the Prefix Option sent to the MN the L flag MUST be unset and the A flag MUST be set or unset, depending on the possibility to perform stateless auto-configuration in the local domain.

8.1.3.4. MN IP address notification procedure

The NetLMM protocol specification is agnostic of how the IP address is assigned to the MN in the local domain. However, the MAG needs to know the IP address assigned or auto-configured by the MN in order to update the routing at the LMA. For this purpose, the MAG needs to play an active role in the DHCP exchange or needs to receive a trigger after the MN has configured an IP address via stateless auto-

configuration. How this is done is described later in this section.

As soon as the MAG knows that a specific IP address has been assigned to a MN, it MUST send a MN Address Setup message to the serving LMA, including three ID options (MN ID, MAG ID, LMA ID), a NetLMM Address option containing the address assigned to (or configured by) the MN. This message is a request to the LMA to start forwarding the packets destined to that IP address to the MAG. The MAG MAY also add the new IP address to the Routing Cache entry related to that MN.

When it receives a MN Address Setup Acknowledgement message with Status 1 "SUCCESS", the MAG MUST add the new IP address to the Routing Cache entry of the MN (or set it as non tentative if previously updated) and update its forwarding state. In case the message contains a different Status value (Values 5 or 6), it MUST reject the assignment of the IP address to the MN, depending on the method used by the MN to request the address (i.e., DHCP or stateless auto-configuration).

In case DHCP is used, the MAG MUST act as DHCP Relay Agent in order to have an active role in the DHCP exchange. The MAG then receives from the DHCP server the IP address assigned to the MN in a DHCP Relay-reply message and updates the routing at the LMA and at the same time can send a DHCP Reply message to the MN with the assigned address. It is important that in the DHCP-Relay-forward message, the MAG includes the identifier of the LMA that is in charge of serving that MN so that the DHCP server can select the IP address accordingly (i.e., from the IP subnet of the LMA). In case the MAG receives a MN Address Setup Acknowledgement message with Status 5 or 6, it MUST send a DHCP Reply message, refusing the IP address assignment to the MN.

In case stateless auto-configuration is performed, the trigger to update the routing at the LMA is the DAD procedure. The DAD procedure is performed on the MAG link, but the MAG needs to verify the address uniqueness at the LMA; for this purpose, it sends the MN Address Setup message. If the LMA replies with a Status value equal to 5 or 6, the MAG MUST send a Neighbor Advertisement in order to fake an IP address duplication.

8.1.3.5. MAG to MAG handover procedure

When a MN hands over from one MAG to another, the new MAG may not know if the event occurred is a handover or a network attach. This is because the base protocol specified in this document is agnostic of any MAG to MAG communication that may be in place. Due to this reason, as for network attach, the MAG will just receive a trigger that a new MN has attached to the link; this trigger, referred as a

MN_Access_Network API carries the MN ID and the LMA ID. As mentioned above, this API can be for example based on an AAA exchange.

After receiving this API, the MAG performs the same procedure described for network access (see section Section 8.1.3.3): it sends a Location Registration message including three different ID options, containing the MN ID, the MAG ID and the LMA ID.

In this handover handover case, the MAG does not receive immediately a Location Registration Acknowledgement message; instead it receives a Routing Setup message from the LMA that includes all IP addresses that are assigned to the MN. These addresses are included in one or more NetLMM Address options. The message contains also some forwarding state, based on the tunneling mechanism used. (Note: how the tunnel ID is carried in the message? Is it really needed at this step?). When the MAG receives this message, it MUST create a new entry in its Routing Cache for the specific MN and MUST update its forwarding state. In case no errors occur, the MAG sends back to the LMA a Routing Setup Acknowledgement message with Status value set to 1 "SUCCESS" and including the forwarding state information.

After that, the MAG receives the Location Registration Acknowledgement message that includes the NetLMM prefix to be delivered to the MN. The procedure is the same as described for network attach in section Section 8.1.3.3.

8.1.3.6. Resource Revocation

If the MAG receives a Location Deregistration message from the LMA, it MUST delete the entry related to the MN specified in the MN ID Option in its Routing Cache. Moreover, the MAG MUST remove any forwarding state for the MN. After doing that, the MAG MUST send a Location Deregistration Acknowledgement to the LMA with Status 1 "SUCCESS".

In case the Location Deregistration contains a Timer attribute (Note: it seems to me the Timer option has not been defined yet), the MAG MAY keep forwarding uplink packets to the LMA for the MN. This may be useful in case of make before break link layer technologies. The adopted timer cannot be greater than the one suggested by the LMA and MUST be sent back to the LMA in the Location Deregistration message acknowledgement.

8.1.3.7. Network Detachment

In case the MAG has an indication that the MN has detached from the network (e.g., via AAA architecture), the MAG MUST (Note: if the protocol is stateful and we do not have a lifetime in the location

registration, this is a MUST, otherwise it can be a SHOULD) send a Location Deregistration message with an ID option including the MN ID. After receiving the Location Deregistration Acknowledgement, the MAG MUST remove any mobility and forwarding state in its Routing Cache related to that MN.

8.1.3.8. IP address no longer in use

In case the MAG has an indication that an IP address is no longer used by the MN, it MUST send a MN Address Remove message to the LMA, including the MN ID and the NetLMM Address that needs to be removed in a NetLMM Address option. How the MAG detects that an IP address is no longer used is out of the scope of this document.

8.1.3.9. Link availability test

MAGs should ensure availability of their link to LMAs. To test link availability, each MAG should periodically send a Heartbeat message to each LMA with which it has associated according to the entries in the LMA List. If the Heartbeat Ack message from the LMA is received within HEARTBEAT_TIMEOUT seconds, availability of the link to the LMA is proven. If the MAG does not receive the Heartbeat Ack message within HEARTBEAT_TIMEOUT seconds, the MAG should send HEARTBEAT_RETRY_MAX further Heartbeat messages with incremented sequence number. In case none of these Heartbeat messages is acknowledged by the LMA, the MAG should raise an alarm. Optionally, the MAG can inform an external OAM function about the broken link.

To avoid superfluous bandwidth consumption, MAGs should send Heartbeat messages to an LMA only in case there was no traffic on the associated link for LINK_ACTIVITY_TIMEOUT seconds. MAGs should perform Heartbeat tests according to the following rule:

In case there was no signaling activity on a link to an LMA for LINK_ACTIVITY_TIMEOUT seconds, the MAG waits for an additional random delay time between HEARTBEAT_DELAY_MIN and HEARTBEAT_DELAY_MAX seconds. After the delay has expired, the MAG sends the Heartbeat message to the LMA. In case the MAG receives a Heartbeat message from the LMA while waiting for the additional random delay, the MAG should reset the delay timer and refrain from sending the Heartbeat message, but MUST acknowledge the Heartbeat message using the same sequence number as in the received message.

8.2. Local Mobility Anchor Operation

8.2.1. Conceptual Data Structures

Each LMA MUST maintain a NetLMM Routing Cache and a MAG List.

Each LMA Routing Cache entry conceptually contains the following fields for each MN:

* MN ID of the registered MN. This Identifier is learned through the Location Registration message, which registers an attached MN. Optionally, the MN ID is learned though the LMA Allocation message beforehand, which could enable service authorization for this particular MN ID at the LMA.

* Routing Tag for the registered MN. Creating the entry and use of this tag is optional and might support identification of the MN in the data plane at the LMA and the MAG. In case the LMA and MAG specify asymmetric tags for the MN, this field MUST draw a distinction.

* MAG ID of the registered MN's serving MAG. This identifier is learned through the Location Registration message, which registers an attached MN. Dependent on the context of this message, the MAG ID entry is either initialized or updated in case of the MN's handover. The MAG ID can be linked to the associated MAG's IP address With the help of the MAG List.

* One or more global IP addresses of a registered MN. Each IP address is learned from an MN Address Setup message. According to the context of the received message, the IP address is set up, updated or removed from the Routing Cache.

Each LMA MUST maintain a MAG List, which refers to associated MAG entities. The list of associated MAGs is used to perform heartbeat tests and to map the Routing Cache's MAG ID entries to the associated MAG's IP address(es). The MAG List also supports the procedure of bulk de-registrations at all or a subset of associated MAGs.

The MAG List conceptually contains the following fields for each associated MAG:

* MAG ID of the associated MAG.

* One or more IP address(es) of the associated MAG. The MAG's IP address information is learned through the Associate message.

* Forwarding capabilities of the associated MAG. This capability list is learned from a particular MAG through the Associate message. From the list of supported forwarding approaches, the LMA enters only these approaches to the capabilities, which are supported by the LMA.

* Forwarding setting for the associated MAG. This field is needed in case the LMA configures a single forwarding approach per MAG association.

* Capability list of the associated MAG. These capabilities are learned from a particular MAG through the Associate message.

8.2.2. Processing NetLMM Headers

All NetLMM local mobility anchors MUST observe the rules described in Section 8.1.2 when processing NetLMM Headers.

8.2.3. Processing NetLMM Messages

8.2.3.1. Associate Procedure

When a LMA receives an Association Request message, it MUST look up into its MAG list based on the MAG identifier contained in the ID option included in the request. If an entry for that MAG ID is already present in the MAG list, the LMA MUST send an Associate Reply message to the MAG with STATUS 10 "Already Associated" (note: an alternative may be that the LMA silently discards the message) (note: another alternative is that the new association request is an association update, e.g., in order to propose a new forwarding method. I would keep things simple and not include this case, not allowing an Association Request if the MAG is already associated).

If an entry for the MAG identifier contained in the ID option does not exist, the LMA MUST create it, including the parameters contained in the Association Request message (MAG ID, MAG IP address). Based on internal policy (e.g., pre-configuration) the LMA MAY accept the data forwarding method proposed by the MAG or MAY propose other methods in Access Reply. After creating the entry, the LMA MUST send an Associate Reply message with STATUS value 1 ("Success"), including its ID in an ID option. Optionally, it MAY include also a Data Transport Option included the data forwarding method to be used. (Note: don't we need a message ID in order to bind the request with the reply?)

The Association Request message MUST include:

* the LMA identifier included in a NetLMM ID option. This identifier is used by the peer to identify the LMA and is included in all subsequent messages.

* the LMA's capabilities in terms of support of data transport methods included in a NetLMM Data Transport Option. The LMA MUST insert in this option all possible tunneling methods that can be

used with the peer MAG Based on configuration, it is possible that some tunneling methods are used only with some MAGs: in this case, the Association Request message MUST contain only the tunneling methods that are administratively permitted with that specific MAG.

8.2.3.2. Disassociate Procedure

In case the Disassociate procedure is initiated by the MAG, when the LMA receives a Disassociation Request message, the LMA MUST validate it. If it is correct, it MUST delete the related entry in its MAG List and send a Disassociate Reply with Status 1 "SUCCESS". As in all NetLMM messages, the LMA MUST include the ID option with its identity.

In case the Disassociate Procedure is initiated by the LMA the LMA MUST include an ID Option with the its identity in the Disassociation Request. When sending the Disassociation Request, the LMA MAY set the MAG entry related to the specific MAG as tentative. When it receives a Disassociation Reply with Status 1 "SUCCESS", the LMA MUST delete the correspondent entry in its MAG List.

8.2.3.3. MN network access procedure

When the local mobility anchor receives a Location Registration message, it MUST validate it (note: what does it mean? should it check if the MAG ID in the message is in the MAG List). If the message is valid, it MUST check if an entry for the MN identifier included in the Location Registration is present. If an entry is already present, it means that a MAG to MAG handover has occurred: the detailed procedure for this event is described in Section 8.2.3.5.

If an entry for that MN identifier is not present, the LMA MUST create a new entry with the MN ID, the MAG ID (?) and the MAG IP address (?). (Note: the two latter parameters are not present in the Routing Cache description. How the protocol works without them?). After creating the entry, it MUST send a Location Registration Acknowledgement with STATUS 1, including three ID options (MN ID, LMA ID, MAG ID) and the NetLMM prefix in a NetLMM Address option. The NetLMM prefix is then forwarded in a Router Advertisement to the MN by the MAG and used by the MN to auto-configure an IP address using stateless auto-configuration and/or to detect a change of local domain.

In case the Location Registration is not valid or the registration procedure cannot be completed successfully, the LMA MUST send a Location Registration Acknowledgement with an appropriate Status

value.

8.2.3.4. MN IP address notification procedure

When the MN tries to configure a new IP address on the local domain, the local mobility anchor receives a MN Address Setup message from the MAG which the MN is connected to. This message contains the IP address the MN is trying to configure in a NetLMM Address option. (Note: in the slides there is also a tunnel ID but I am not sure this is actually needed).

When it receives this message, the LMA MUST check if the IP address in the NetLMM Address option is already assigned to another MN. If the address is a duplicated one, it MUST send a MN Address Setup Acknowledgement message with Status 5 "INVALID IP ADDRESS". If the address is not assigned to any other MN, the LMA MUST check if the number of addresses assigned to that MN does not exceed the maximum number of addresses that the MN can configure. This check is performed in the LMA Routing Cache; the maximum number of allowed IP addresses is a per MN variable that is pre-configured at the LMA. If the number of IP addresses exceeds the allowed one, the LMA MUST send a MN Address Setup Acknowledgement message with Status 6 "OVER IP ADDRESS LIMIT". (Note: the first check is mainly performed for the stateless configuration case and may be skipped in case DHCP is used. However the LMA does not know if the address is configured using DHCP or stateless so I would keep this check in any case)

If the IP address is not a duplicated one and the maximum number of allowed IP addresses is not reached, the LMA MUST update the Routing Cache entry indexed by the MN ID contained in the MN Address Setup message with the new IP address and MUST update its forwarding state, depending on the tunneling mechanism used.

8.2.3.5. MAG to MAG handover procedure

When the LMA receives a Location Registration message, it MUST check in its Routing Cache if an entry for the MN ID carried in the message is already present. If it is not, that means that the MN is accessing the network for the first time (see section Section 8.2.3.3). If an entry is already present in the Routing Cache, a handover has occurred. In this case the LMA MUST send a Routing Setup message to the MAG, including three ID options (MN ID, MAG ID, LMA ID), one or more NetLMM Address options containing all NetLMM addresses configured to the MN and some forwarding state information that depends on the tunneling method used (e.g., tunnel ID).(Note: what it the behavior of the LMA during this time interval when packets arrive? Should it already update the Routing Cache entry or should it wait for an ack?)

After that the LMA receives a Routing Setup Acknowledgement message; if the Status of this message is set to SUCCESS, the LMA MUST update the Routing Cache with the new location of the MN, updating the MAG ID. The LMA MUST then send back to the MAG a Location Registration Acknowledgement message with Status value set to 1 "SUCCESS", including in a NetLMM Address option the NetLMM prefix to be sent to the MN.

8.2.3.6. Resource Revocation

There are case (e.g. due to administrative reasons) where the forwarding state of a specific MN must be purged so that the MN is no more able to use the resources provided by the network. In this case, based on a trigger received from the network (e.g. AAA), the LMA MUST send a Location Deregistration message to the peer MAG, including the MN ID, the LMA ID and the MAG ID. Optionally, the LMA MAY include a Timer attribute specifying the remaining time to keep the state of the MN in the MAG. The revocation procedure terminates when the LMA receives a Resource Revocation Acknowledgement with Status 1.

8.2.3.7. Network Detachment

When the LMA receives a Location Deregistration message from the peer MAG, it MUST remove in its Routing Cache the entry of the MN indicated by the MN ID in the Location Deregistration message. After that, it MUST send a Location Deregistration Acknowledgement with Status 1, including the MN ID, the MAG ID and the LMA ID.

8.2.3.8. IP address no longer in use

When the LMA receives a MN Address Remove message, it MUST remove the NetLMM Address included in the NETLMM Address option from the entry in its Routing Cache related to the MN indicated in the message. After that, the LMA MUST respond with a MN Address Remove Acknowledgement with Status 1.

8.2.3.9. Link availability test

LMAs should ensure availability of their link to MAGs. To test link availability, each LMA should periodically send a Heartbeat message to each associated MAG according to the entries in the MAG List. If the Heartbeat Ack message from the MAG is received within HEARTBEAT_TIMEOUT seconds, availability of the link to the MAG is proven. If the LMA does not receive the Heartbeat Ack message within HEARTBEAT_TIMEOUT seconds, the LMA should send HEARTBEAT_RETRY_MAX further Heartbeat messages with incremented sequence number. In case none of these Heartbeat messages is acknowledged by the MAG, the LMA

should raise an alarm. Optionally, the LMA can inform an external OAM function about the broken link.

To avoid superfluous bandwidth consumption, LMAs should send Heartbeat messages to a MAG only in case there was no traffic on the associated link for LINK_ACTIVITY_TIMEOUT seconds. LMAs should perform Heartbeat tests according to the following rule:

In case there was no signaling activity on a link to a MAG for LINK_ACTIVITY_TIMEOUT seconds, the LMA waits for an additional random delay time between HEARTBEAT_DELAY_MIN and HEARTBEAT_DELAY_MAX seconds. After the delay has expired, the LMA sends the Heartbeat message to the MAG. In case the LMA receives a Heartbeat message from a MAG while waiting for the additional random delay, the LMA should reset the delay timer and refrain from sending the Heartbeat message, but MUST acknowledge the Heartbeat message using the same sequence number as in the received message.

9. Data Transport

As soon as a particular MAG has associated with an LMA and an attached MN has been registered with the LMA, the LMA node and the MAG node are responsible for forwarding the MN's data traffic correctly within the NetLMM domain. Associated location and forwarding information is maintained within the LMA's and the MAG's Routing Cache. Different forwarding mechanisms between the LMA node and a particular MAG node might be supported and set up during the MAG's association procedure.

Network entities, which have Version 0 of the NetLMM protocol implemented, MUST support IPv6-in-IPv6 encapsulation to tunnel data packets between an LMA node and an associated MAG node. Support of other forwarding approaches are for future extensions.

9.1. Forwarding of Unicast Data Packets

9.1.1. Handling of hop limit field in IPv6 data packets

According to the NetLMM default mechanism to forward data packets between a particular LMA and MAG by means of encapsulation, LMA nodes and MAG nodes serve as tunnel entry-points and tunnel exit-points respectively. LMAs and MAGs have to decrement the hop limit field of the encapsulated IPv6 header properly. The MAG serves as the default gateway for an attached MN and forwards all packets from the MN into the tunnel, which in turn encapsulates the packet towards the LMA. The LMA on receiving the packet from the MAG decapsulates and forwards the packet using normal forwarding procedures. The packets

destined towards the MN are forwarded in a similar fashion. The procedure of forwarding the packet decrements the hop limit. Hence, the hop limit will get decremented twice for any packet traversing the tunnel between LMA and MAG.

9.1.2. IPv6-in-IPv6 tunnel

LMA and MAG nodes MUST support IPv6-in-IPv6 encapsulation to forward packets within the NetLMM domain. Support of other forwarding schemes is optional. When an LMA node receives an IPv6 packet destined for a registered MN and IPv6-in-IPv6 tunneling has been selected as forwarding approach, it serves as tunnel entry-point. The LMA node decrements the hop limit of the data packet's IPv6 header by one and encapsulates the packet in the tunnel IPv6 header. The tunnel IPv6 header might carry one or more extension headers. The LMA node forwards the tunnel packet to the MAG node, using its own address as source address and the MAG node's address as destination address in the outer IPv6 header. The MAG node terminates the tunnel and MUST process relevant Extension Headers, which might follow the encapsulating IPv6 header. The MAG node forwards the data packet towards the MN after decapsulation.

To forward uplink packets, the MAG node serves as tunnel entry-point and decrements the data packet's hop limit field by one before it encapsulates the packet in the tunnel IPv6 header. The tunnel IPv6 header might carry one or more extension headers. The MAG node forwards the tunnel packet to the LMA node, using its own address as source address and the LMA node's address as destination address in the outer IPv6 header. The LMA node terminates the tunnel and MUST process relevant Extension Headers, which might follow the encapsulating IPv6 header. The LMA node forwards the data packet towards its final destination after decapsulation.

9.1.3. Future extensions

Future extensions might support other approaches to forward data packets between LMA and MAG nodes. Such extensions could include IPv4-in-IPv6 encapsulation to forward IPv4 data packets of dual stack hosts within the NetLMM domain. Operators might prefer the use of other flexible forwarding approaches, such as Generic Routing Encapsulation (GRE) [RFC2784] or Multi-Protocol Label Switching (MPLS), and follow [RFC3031] and associated mechanisms to forward data packets between LMA and MAG nodes. Details about the use of such extensions are out of scope of this specification.

For now, some suggestions of how GRE tunnelling could be used with NetLMM can be found in Appendix B, and similarly use with MPLS is described in Appendix C.

9.2. Forwarding of Multicast Data Packets

9.2.1. Link Local Multicast

The scope of link local multicast packets is confined to the link between MNs and the associated MAG node. MAG nodes process but do not forward link local multicast packets. To support some functions, such as for Duplicate Address Detection (DAD), MAGs might proxy associated Neighbor Discovery messages to perform DAD procedure with LMA nodes.

9.2.2. IP Multicast

Three options have been identified to support IP multicast within a NetLMM domain.

Option 1 implies that MAG nodes are multicast-enabled routers and support for IP multicast is orthogonal to the NetLMM protocol operation. According to native multicast support, access routers terminate a multicast tree and the LMA node does not play any multicast-specific role in forwarding of IP multicast traffic.

Option 2 implies that MAG nodes are multicast-enabled routers and IP multicast traffic is tunnelled via the NetLMM domain's LMA nodes.

Option 3 implies that an NetLMM domain's LMA nodes are multicast-enabled routers. LMA nodes join a multicast-tree and forward IP multicast packets to MAG nodes according to the selected forwarding approach. MAG nodes must coordinate multicast listeners according to IGMP operation [RFC3376] and communicate with LMA nodes using PIM control messages [RFC2362] to control IP multicast forwarding path between LMA nodes and respective MAG nodes, which have multicast listeners attached. LMA nodes coordinate multicast listener registration with other multicast routers, which are outside of an NetLMM domain, by means of PIM control messages. An exemplary IP multicast join procedure is illustrated in Fig. *MCJOIN*

The specification of default operation for IP multicast support and optional enhancements is for further study and t.b.d.

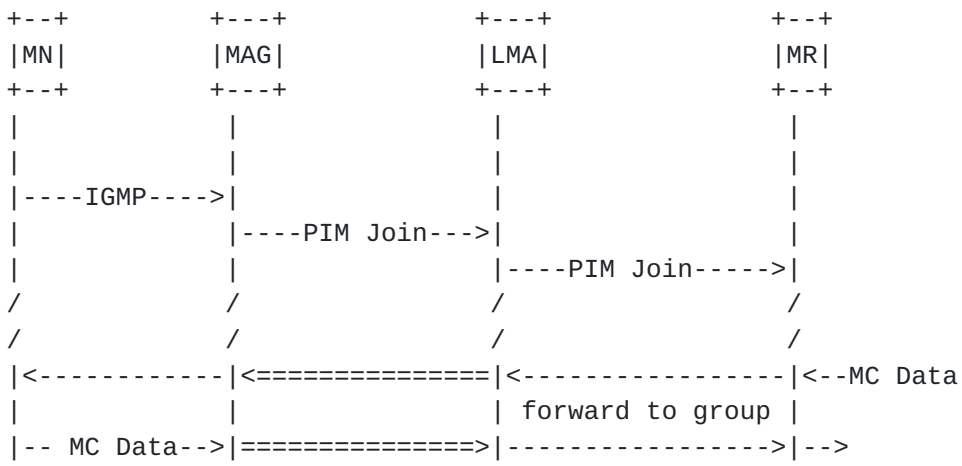

Figure *MCJOIN*: IP multicast join procedure for the case that LMA nodes and MAG nodes are multicast routers.

9.3. Forwarding of Broadcast Data Packets

Version 0 of the NetLMM protocol specification does not consider forwarding of broadcast data packets.

10. Protocol Constants and Configuration Variables

11. Security Considerations

The NetLMM protocol is executed between the MAG and LMA. The messages are used to create, update and delete mobility state in MAG and LMA. If the NetLMM signalling messages are not authenticated, following attacks are possible.

* A malicious node can pretend to be MAG and associate with the LMA. This in itself may not create any harm if subsequent messages are authenticated. But it does allow for the attacker to learn the capabilities of the LMA which in turn may be used to exploit the specific weaknesses.

* A malicious node can pretend to be MAG and send a location registration message to LMA. There are a couple of variants of this attack.

  - The MN is currently attached to MAG-1 and the IP address of the MN is known. A malicious node MAG-2 sends the location registration message to redirect all the traffic destined for the MN to itself. This enables the attacker to steal the MN's

traffic. This attack may be detected by MAG-1 when it receives the location de-registration message from LMA while the MN is still attached. The detection of the attack depends on the security of mechanism used to detect the MN's attachment.

- The MN is currently attached to MAG-1 and the IP address of the MN is not known. A malicious node MAG-2 sends the location registration message to redirect all the traffic destined for the MN to itself. The LMA would update the mobility state for MN with the ID of MAG-2. Later when the MN-address-setup messages comes from MAG-1, it would fail because the MAG IDs don't match.

* A malicious man-in-the-middle node can alter the messages sent between MAG and LMA that can result in random attacks. For example, the attacker can modify the MN Identifier in the location registration message of MN-1 to that of MN-2. When MN-2 moves to a new link, it results in a new location registration message to be sent with MN-2 ID. This will cause the traffic destined to MN-1 address to be destined to the current link causing disruption for both MN-1 and MN-2.

These attacks can be prevented by providing data origin authentication for the messages exchanged between LMA and MAG. This can be achieved by using IPsec ESP [RFC4303] with null-encryption and non-null authentication between MAG and LMA.

It is possible to filter the signalling messages at the edge of the network so that a rogue MN or rogue node on the Internet cannot source such messages. Hence, any messages exchanged between the MAG and LMA can only come from within the network. This level of security may be sufficient for some deployments precluding the need for protecting the signalling messages. In such cases, IPsec may not be used to protect the signalling messages.

Anomalous events should be logged. For example, once an address is assigned to the Mobile node, MN address setup message should appear at the LMA sooner or later. If it does not appear within a certain period of time, it should be considered as an error and logged. [TBD: May be document more of such events].

NetLMM messages are triggered by MN attachment and detachment to the MAG. The threats specific to MN-MAG interface are discussed in [I-D.ietf-netlmm-threats]. When neighbor discovery messages [I-D.ietf-ipv6-2461bis] are used as triggers, it should be secured using [RFC3971].

12. IANA Considerations

13. Contributors

This contribution is a joint effort of the NetLMM design team of the NetLMM WG. The members include Kent Leung, Gerardo Giaretta, Phil Roberts, Marco Liebsch, Katsutoshi Nishida, Hidetoshi Yokota, Henrik Levkowetz, and Mohan Parthasarathy.

14. Acknowledgments

15. References

15.1. Normative References

[I-D.ietf-ipv6-2461bis]
Narten, T., "Neighbor Discovery for IP version 6 (IPv6)", draft-ietf-ipv6-2461bis-07 (work in progress), May 2006.

[I-D.ietf-netlmm-nohost-ps]
Kempf, J., "Problem Statement for Network-based Localized Mobility Management", draft-ietf-netlmm-nohost-ps-04 (work in progress), June 2006.

[I-D.ietf-netlmm-nohost-req]
Kempf, J., "Goals for Network-based Localized Mobility Management (NETLMM)", draft-ietf-netlmm-nohost-req-01 (work in progress), April 2006.

[I-D.ietf-netlmm-threats]
Kempf, J. and C. Vogt, "Security Threats to Network-based Localized Mobility Management", draft-ietf-netlmm-threats-00 (work in progress), April 2006.

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2362] Estrin, D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., and V. Jacobson, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification", RFC 2362, June 1998.

[RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.

[RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[RFC3753] Manner, J. and M. Kojo, "Mobility Related Terminology", RFC 3753, June 2004.

[RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[RFC4282] Aboba, B., Beadles, M., Arkko, J., and P. Eronen, "The Network Access Identifier", RFC 4282, December 2005.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.

[RFC4330] Mills, D., "Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI", RFC 4330, January 2006.

15.2. Informative References

[ENTERPRISE-NUM] IANA, "IANA Enterprise Numbers Registry"", 2006.

[RFC4064] Patel, A. and K. Leung, "Experimental Message, Extensions, and Error Codes for Mobile IPv4", RFC 4064, May 2005.

Appendix A. TODO (Things that remain to be specified...)

This is a short list of things that remain to be specified in future revisions of this document:

* Describe the re-send mechanism for control messages, in order to provide reliable delivery.

* Add an "LMA Announce message" which can be multicast from a newly connected LMA to trigger listening MAGs to send it Association Requests.

* Add the capability to do bulk MN de-registrations and possibly registrations.

* A review to ensure that all aspects of the protocol permits operation with both single /128 addresses and for instance /64 prefix allocations to mobile nodes.

* Add reserved status ranges (Section 7.1.1) for vendor-specific status, LMA status?, MAG status?.

* Add experimental message, option and status values, in accordance with [RFC4064].

* Add a Heartbeat Sequence Number Option, to hold heartbeat-specific sequence numbers needed by algorithms for reacting to missed heartbeats. Write up suggested algorithm.

* Make sure the use of the identity / locator split paradigm is consistent and workable throughout the document. Cover resolution of ID to locator.

* Define how capability exchanges are handled, and how a unique common capability is derived, for instance to find the tunnelling method to be used as a result of the Association Request and Reply.

* Add message and signalling optimization according to Section 5.12.

* Maybe change the range-based skippable or non-skippable nature of Options to be indicated by a bit instead?

* Point out somewhere that although a NetLMM Domain may have multiple LMAs, a MN is always served by the same LMA once an LMA has been assigned.

*

Appendix B. Using GRE Tunnels with NetLMM

** Text to be written **

Appendix C. Using MPLS with NetLMM

** Text to be written **

Appendix D. TTL Handling

** Text to be written **

Appendix E. MN-AR Interface considerations

This document assumes that the MN-AR interface document will describe the following in more detail:

* - A mechanism for indicating duplicate address to the MN

* - No redirects should be sent by MAG to MN even if the destination is directly connected to MAG

* - Trigger for IP address configuration

* - MN Identifier option in the trigger ?

* - If SEND is used, Proxy SEND details are needed for defending the address in the case of a duplicate

* - Router advertisement details : unicast only, what else does it contain etc."

Appendix F. Out of scope

- Inter-MAP handover - Fast handover - Hierarchical MAP

Authors' Addresses

Gerardo Giaretta
Telecom Italia
via Reiss Romoli 274
Torino 10148
Italy

Phone: +39 011 228 6904
Email: gerardo.giaretta@telecomitalia.it

Kent Leung
Cisco
170 W. Tasman Drive
San Jose, CA 95134
USA

Phone: +1 408 853 9580
Email: kleung@cisco.com

Marco Liebsch
NEC Network Laboratories
Kurfuersten-Anlage 36
Heidelberg, 69115
Germany

Phone: +49 6221-90511-46
Email: marco.liebsch@netlab.nec.de

Phil Roberts
Motorola
1301 E Algonquin Rd
Schaumberg, IL 60196
USA

Email: phil.roberts@motorola.com

Katsutoshi Nishida
NTT DoCoMo Inc.
3-5 Hikarino-oka, Yokosuka-shi
Kanagawa,
Japan

Phone: +81 46 840 3545
Email: nishidak@nttdocomo.co.jp

Hidetoshi Yokota
KDDI R&D Laboratories, Inc.
2-1-15 Ohara, Fujimino
Saitama, 356-8502
Japan

Phone: +81 49 278 7894
Email: yokota@kddilabs.jp

Mohan Parthasarathy
Nokia

Email: mohan.parthasarathy@nokia.com

Henrik Levkowetz
Ericsson
Torsgatan 71
Stockholm S-113 37
SWEDEN

Phone: +46 708 32 16 08
Email: henrik@levkowetz.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.