NETEXT	H. Chan	
Internet-Draft	F. Xia	
Intended status: Standards Track	J. Xiang	
Expires: September 10, 2010	H. Ahmed	
	Huawei Technologies	
	March 09, 2010	

TOC

Distributed Local Mobility Anchors
draft-chan-netext-distributed-lma-03

Abstract

This draft proposes a distributed local mobility anchors architecture. It splits the functions of a local mobility anchor into different logical functions: (1) allocation of home network prefixes or home addresses to mobile nodes, (2) location management (LM) which includes managing the IP addresses and locations of the mobile nodes, and (3) mobility routing (MR) which includes intercepting and forwarding packets. The distributed local mobility anchors architecture consists of home local mobility anchors (H-LMA) at the registered networks and visited local mobility anchors (V-LMA) at the visited networks. The V-LMA provides mobility routing function to avoid triangle routing problem in Proxy mobile IP, whereas the H-LMA keeps the location management function. The needed location information of a mobile node is acquired by a V-LMA from the H-LMA only when a packet is first sent to the mobile node via the V-LMA and are then cached at the V-LMA to enable optimized mobility routing for packets subsequently sent to the mobile node.
When either the source or the destination node is a fixed node, bypassing the tunneling role of all LMA's will expose the location information of the mobile node in some cases. Yet unnecessarily long routes are still avoid by having multiple V-LMA's in different networks and serving the packet transport with the nearest V-LMA.

Status of this Memo

time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as “work in progress.”
The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.
The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.
This Internet-Draft will expire on September 10, 2010.

Table of Contents

1. Introduction
2. Motivation
    2.1. Splitting the logical functions of a local mobility anchor
    2.2. Originating local mobility anchor and destination local mobility anchor
    2.3. Home local mobility anchor versus visited local mobility anchor
3. Terminology
4. Overall mechanism
    4.1. Registration
    4.2. Anycast
    4.3. Visited Network
    4.4. Mobility routing
5. Packet flow
    5.1. Sending packets to a mobile node
        5.1.1. Bypassing D-LMA
        5.1.2. Changing MAG without changing D-LMA
        5.1.3. Changing LMA
    5.2. Sending packets from a mobile node
    5.3. Sending packets from a mobile node to another mobile node
    5.4. Bypassing the tunneling role of the LMA
6. Performance
    6.1. Round trip time
    6.2. Call setup delay

6.3. Location update signaling overhead
6.4. Simultaneous move problem
7. Interworking with legacy LMAs
7.1. Reachability from CN outside MN's domain
7.2. Sending packet to CN in a different distributed-LMA domain
7.3. Sending packet to a CN in a legacy PMIP domain
7.4. Sending packet to CN running MIP outside MN's distributed-LMA domain
8. IANA Considerations
9. Security Considerations
10. Acknowledgments
11. References
11.1. Normative References
11.2. Informative References
§ Authors' Addresses

---

1. Introduction

TOC

Proxy mobile IP [RFC5213] as well as mobile IP [RFC3775] support mobility by using a home address for session and a care-of address for routing but has the problem of triangle routing when a mobile node is far from its home agent while being much closer to its correspondent node.

Unneccessarily long routes may be avoided by having multiple home agents in different geographic locations [GHAHA]. These home agents announce the same IP prefixes using anycast. The traffic originating from the mobile node will then be served by the nearest home agent, and the traffic sent from a correspondent node to the mobile node will be intercepted by the home agent nearest to the correspondent node. Therefore both traffic will use the home agent nearest to where the traffic originates, so that triangle routing is avoided. These home agents may possess identical information about the mobile nodes [MHA]. Yet the synchronization of all the home agents will then be a challenge [SMGI]. In addition, the design needs to scale in deployment. Yet the amount of signaling traffic needed in synchronizing the home agents may become excessive when the number of mobile nodes and the number of home agents both increase.

This draft proposes to decouple the logical functions of a local mobility anchor into that of home address allocation, location management, and mobility routing. The mobility routing function may be present in many geographical locations. However, the home address allocation function and the internetwork location management function may be kept only at the network where the mobile node is registered. The individual location management information for a specific mobile node may be acquired whenever needed. Home local mobility anchor and

visited local mobility anchor to a mobile node are then defined in terms of these logical mobility functions, each of which may be implemented in one or multiple instances. These two mobility logical functions do not need to physically co-locate leaving flexibility for the implementation to place them at their most appropriate locations. The concept of proxy home agent and primary home agent has been introduced in [GHAHA], where a proxy home agent closest to a mobile node away from its home agent may perform binding update with the primary home agent on behalf of the mobile node, and also intercept and tunnel messages for the mobile node. This draft extends this work, applies distributed local mobility anchors to proxy mobile IP, and describes mobility routing and its expected performance.
This draft is written using the definitions of Proxy mobile IP, but the proposal works equally well for mobile IP.

---

2. Motivation

TOC

2.1. Splitting the logical functions of a local mobility anchor

TOC

A local mobility anchor, being a home agent, needs to perform the following logical functions: (1) home network prefix or home address allocation function: allocating home network prefix or home address HoA to a mobile node that registers with the network; (2) internetwork location management (LM) function: managing and keeping track of the internetwork location of the mobile node, which include a mapping of the HoA to the mobility anchoring point that the mobile node is anchored to; and (3) mobility routing (MR) function: intercepting packets to/from the home address of a mobile node and forwarding the packets, based on the internetwork location information, either to the destination or to some other network element that knows how to forward to the destination.
When these logical functions are all bundled into one single entity known as the local mobility anchor LMA, having LMA in only one network results in triangle routing problem as shown in Figure 1.

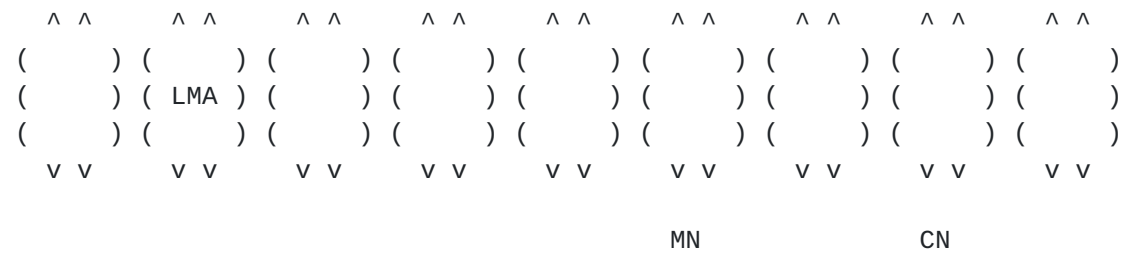

Figure 1. Figure showing the triangle trouting problem with a MN and a CN in networks which may be close to each other but are far from the local mobility anchor (LMA).

The other extreme is to duplicate the LMAs in many networks (Figure 2) to solve triangle routing problem. Yet the location management information will need to be pushed to all these LMAs.

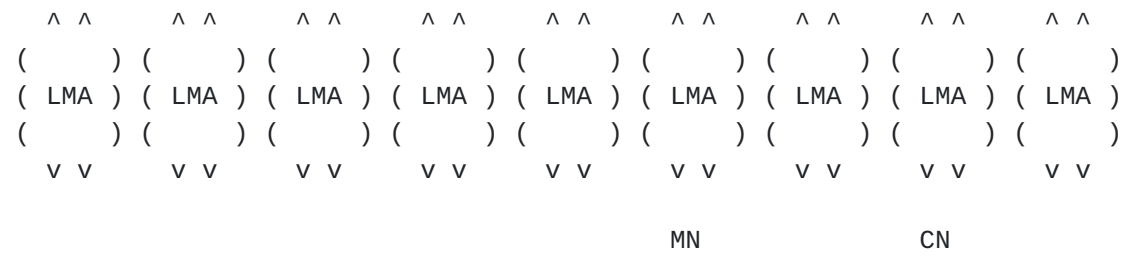

Figure 2. Figure showing the replication of LMAs in multiple networks.

This draft proposes to decouple the logical functions of the local mobility anchor. One may then examine which logical functions should be present in many geographic locations and which logical functions do not.

As illustrated in Figure 3, having mobility routing (MR) function available in multiple geographic locations will solve the triangle routing problem. It is also evident that the home network, which accepts the registration of the mobile node, is responsible for the HoA allocation function. This network may also manage the internetwork location information. Yet pushing the location management (LM) information to the home agents in different networks may be an overkill, especially when the mobile node does not always actually communicate with CNs in all the other networks. Data coherency may be managed using different methods. For example, a distributed database may employ different servers to manage different data. The data in each server is not pushed to all the other servers but the database system only needs to know which data resides in which server. Here, keeping the location management function at the home network will eliminate the need to synchronize the location management information in a timely and scalable manner.

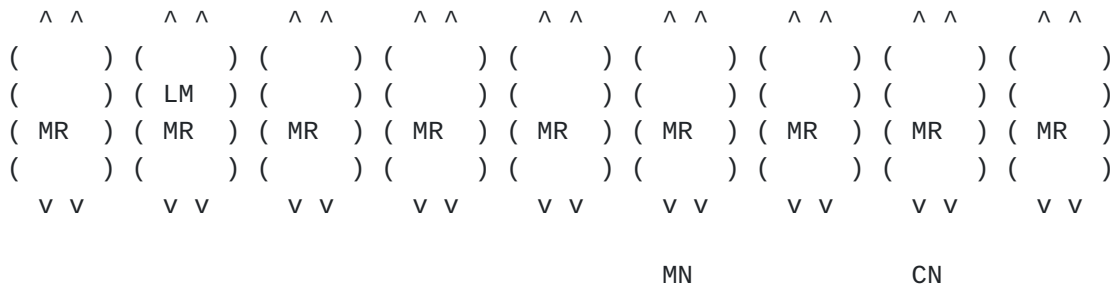

Figure 3. Figure showing the mobility routing (MR) function available in many networks, whereas the dynamic internetwork location management (LM) function resides in one network only.

2.2. Originating local mobility anchor and destination local mobility anchor

TOC

The LMA to which the MN is anchored to is the destination LMA (D-LMA) (Figure 4). It is capable of delivering incoming packets to the MN. When a CN sends a packet to MN, the LMA closest to that CN needs to intercept the packet to avoid triangle routing. This LMA is the originating LMA (O-LMA) that needs to provide mobility routing function for this packet so that the packet may be routed through the internetworks to reach D-LMA.

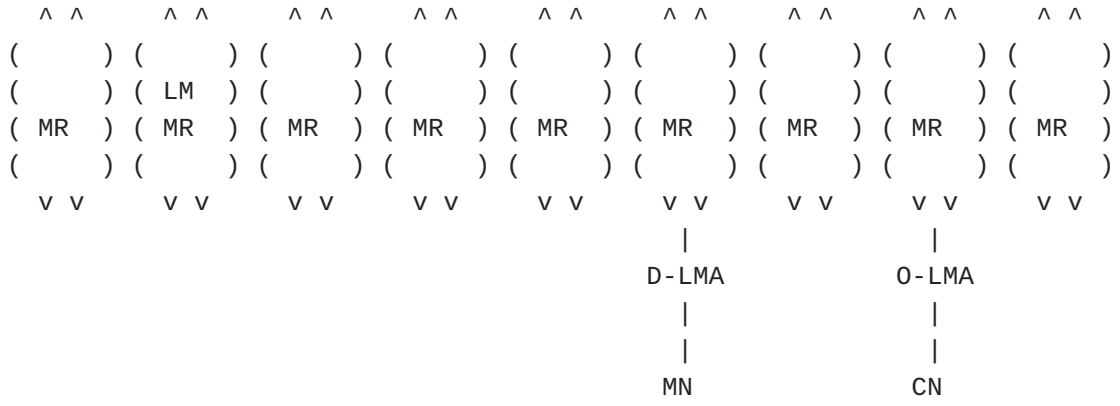

Figure 4. Figure showing the O-LMA and the D-LMA for a packet sent from a CN to a MN.

TOC

2.3. Home local mobility anchor versus visited local mobility anchor

This draft defines home local mobility anchor and visited local mobility anchor as logical functions.
The home local mobility anchor (H-LMA) of a mobile node consists of the logical mobility functions of home-address allocation, location management, and mobility routing, which are provided by the network to which the mobile node is registered. The visited local mobility anchor (V-LMA) is the logical mobility function provided by a visited network. We use the term visited local mobility anchor irrespective of whether the mobile node actually visits that network or not. To the mobile node, V-LMA provides mobility routing function only.
Although the H-LMA performs all the logical mobility functions for a mobile node registered to that network, these logical functions are considered separate and do not need to co-locate. Therefore the local mobility anchor does not need to be one single physical entity. It is possible to have one or multiple physical entities to provide the location management function and one or multiple physical entities to provide mobility routing function. In addition, these different entities do not need to be in one-to-one relationship.
To perform HoA allocation, each H-LMA may use its own block of IP prefixes to allocate IP addresses to the MNs registering to its network. The IP prefixes of all the H-LMAs form a super set of IP prefixes. All the H-LMAs and V-LMAs advertise this same super set of IP prefixes using anycast. Then, no matter where a mobile node is located, the anycast and the routing algorithm will enable the nearest LMA to serve the mobile node.
To perform dynamic internetwork location management function when the MN is in a visited network, H-LMA must know which V-LMA the MN is anhored to. The H-LMAs in different networks provide a distributed database of such records for all the MNs anchored to these networks. The LMA, to which the MN is anchored, delivers incoming packets to the MN; it is the D-LMA for incoming packets. When the MN is in its home network, it is anchored to H-LMA using an HoA address belonging to the H-LMA. When the MN is in a visited network, it is anchored in that network to the nearest V-LMA, and MAG does the MIP signaling on behalf of the MN.
No matter where a correspondent node (CN) is located, any packet sent from the CN to the HoA is intercepted by the nearest LMA. This LMA is the O-LMA. The O-LMA will need to obtain the location information of the MN from the H-LMA in order to route the packets to the D-LMA. Because the HoA of the MN belongs to the IP prefix of its home network, the mapping of the HoA to the H-LMA does not change often and can therefore be known to all V-LMAs. The mobility routing function in the V-LMA before route optimization is simply to forward a packet from the CN to the H-LMA of the MN, and the H-LMA has the dynamic location information about the MN to complete the mobility routing. After route optimization the packets will need to be forwarded directly from the O-LMA to the D-LMA.

3. Terminology

TOC

All the general mobility-related terms and their acronymns used in this document are to be interpreted as defined in the Mobile IPv6 base specification [RFC3775] and in the Proxy mobile IPv6 specification [RFC5213]. These terms include mobile node (MN), correspondent node (CN), home agent (HA), local mobility anchor (LMA), and mobile access gateway (MAG).
In addition, this draft introduces the following terms.

Mobility routing (MR) is the logical function to intercept and forward packets to/from a mobile node.

Home address allocation is the logical function to allocate home address to a mobile node.

Location management (LM) is the logical function to manage the location of a mobile node, which is in terms of the mapping between the HoA and the internetwork location infromation of the mobile node. There are two different mappings to the internetwork location for a mobile node. The mapping to the H-LMA to which the mobile node is registered is usually a static information. The mapping to the local mobility anchor which is serving the mobile node will change when the mobile node changes it mobility anchoring point; H-LMA needs to know about this mapping.

Home local mobility anchor (H-LMA) to a mobile node is the full set of logical functions of a local mobility anchor to the mobile node. It allocates the home address (HoA) to the mobile node, manages the location of the mobile node, intercepts packets to/ from the mobile node, and forwards these packets. Each mobile node is registered in its home network to a H-LMA, which can download the profile of the mobile node from a home AAA server. If the mobile node is anchored to a visited local mobility anchor (V-LMA), the H-LMA will manage the mapping between the HoA and the V-LMA that the mobile node is currently anchored to. The different logical functions do not need to co-locate, and each of these logical functions may be implemented in one or multiple instances.

Visited local mobility anchor (V-LMA) to a mobile node is a subset of the full logical functions of a local mobility anchor towards the mobile node. It intercepts packets to/from the mobile node and forwards packets using the location management information it acquires from the home local mobility anchor of the mobile node. If the mobile node is anchored to the V-LMA, the V-LMA will

inform the H-LMA of the mobile node and may manage the mapping between the HoA and the proxy-CoA of the mobile node.

Originating local mobility anchor (O-LMA) is the first local mobility anchor that intercepts a packet destined to a mobile node.

Destination local mobility anchor (D-LMA) of a mobile node is the local mobility anchor to which the mobile node is currently anchored. It knows where to deliver packets to reach the mobile node.

Distributed-LMA domain is a domain consisting of networks supporting distributed LMA mechanism. Security association can be set up between the LMA's, and the IP prefixes in each LMA is anycasted by the rest of the LMA's in the same domain.

---

4. Overall mechanism

TOC

The architecture of the distributed LMA is shown in Figure 5.

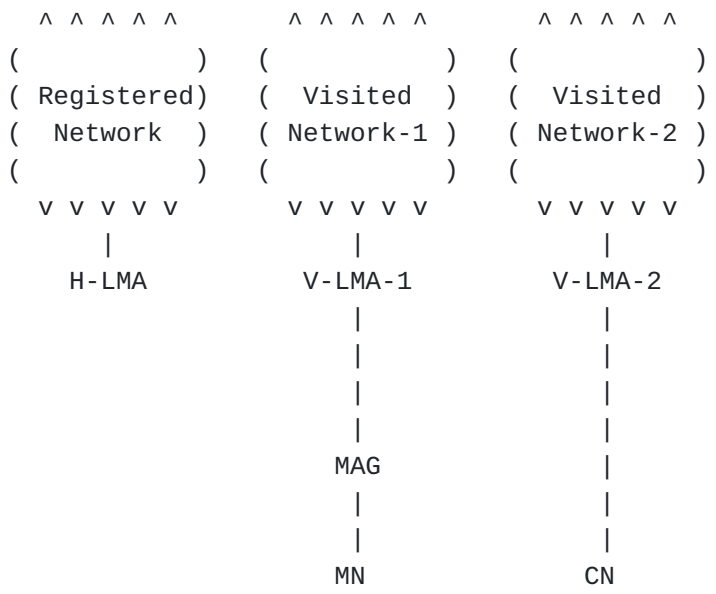

Figure 5. Figure showing a MN registered to its H-LMA in the home network, anchored to the V-LMA (V-LMA-1) in visited network-1, and communicating with a CN served by the V-LMA (V-LMA-2) in visited network-2.

TOC

4.1. Registration

A mobile node MN will register with a H-LMA in its home network.
The H-LMA can download the profile of the MN from the home AAA server.
The H-LMA allocates to the MN a home address HoA belonging to a block of prefixes managed by the H-LMA.
The H-LMA performs mobility routing function for the MN within this home network.
The H-LMA also performs location management for the MN. If the MN has moved to another network and is anchored to a V-LMA of a visited network, the V-LMA needs to update the H-LMN of the new location of the MN, i.e., the new V-LMA the MN is anchored to.

4.2. Anycast

TOC

An example of using anycast for HoA prefixes is shown in Figure 6.

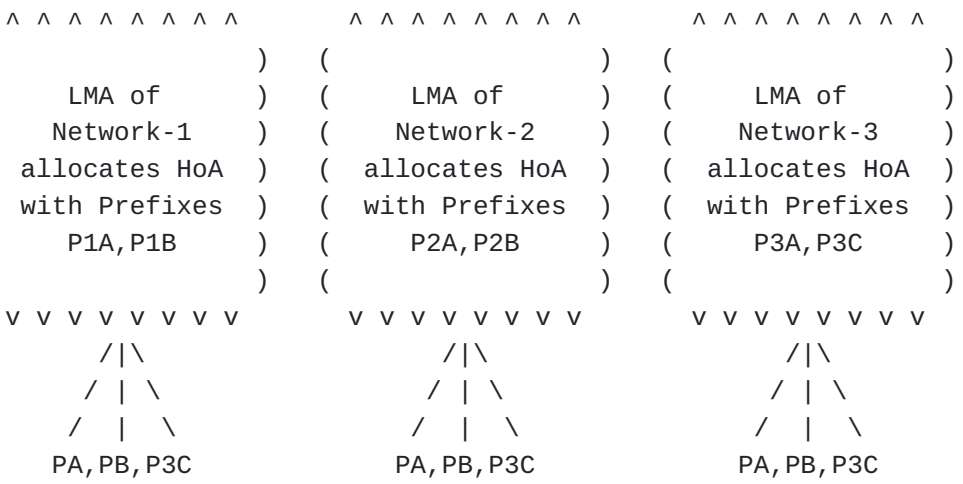

Figure 6. Example of Anycast of HoA Prefixes. The LMA in each network broadcasts the superset of prefixes PA, PB, P3C. Here the PA is the aggregate of P1A, P2A, and P3A; while PB is the aggregate of P1B and P2B.
Each LMA in its network owns a set of IP prefixes which it uses to allocate home network prefixes or HoAs to the MNs registered to that network.
The HoA prefixes of all the LMAs form a superset of HoA prefixes. Some prefixes in this superset may be aggregatable, but it is also possible that some may not be aggregatable.
Each LMA advertises the superset of HoA prefixes. An IP packet sent to any HoA will therefore be intercepted by the LMA nearest to the sender.

TOC

4.3. Visited Network

An MN that has registered with a H-LMA may move to a network other than the home network.
As the MN leaves its home network and enters a visited network, it still receives the prefix advertisement of its HoA from the LMA that uses anycast to advertise the superset of HoA prefixes in the visited network.
A MAG sends binding update to the LMA on behalf of the MN using the HoA of the MN and its proxy CoA.
To ensure robustness, the LMA looks up which H-LMA the MN has registered to, based on the HoA prefix of the MN. It checks with that H-LMA for uniqueness of that HoA address to complete the binding update. If the H-LMA has determined that the HoA is not unique, V-LMA will need to send a registration request to the H-LMA to obtain a valid HoA for the MN.
The visited LMA (V-LMA) in this visited network has become the new mobility anchoring point of MN.
The V-LMA performs mobility routing function for the MN.
The V-LMA informs the H-LMA that it is the current mobility anchoring point for the MN.
After the MN has anchored to a V-LMA (V-LMA-1) in a visited network, it may leave this visited network and move to another visited network. It will then anchor to another V-LMA (V-LMA-2). The H-LMA must again be informed of the new anchoring point.
In addition, the V-LMA-1 is also informed that the MN has anchored to V-LMA-2 so that, for a limited time, if V-LMA-1 receives packets destined to MN, the V-LMA-1 may forward these packets to the V-LMA-2 according to the forwarding mechanism which will be described in the mobility routing section below.

4.4. Mobility routing

TOC

When an originating LMA (O-LMA) has intercepted a packet with a destination address HoA of an MN, it checks whether or not there is location information for this HoA in its cache.
If the HoA information in the cache memory of the O-LMA indicates that the MN is currently anchored to the destination LMA (D-LMA), it tunnels the packet to the D-LMA.
If the location information is not in the cache memory of the O-LMA, the O-LMA tunnels the packet to the H-LMA based on the HoA prefix. Each H-LMA manages a unique set of HoA prefixes, and each LMA knows which HoA prefix is owned by which H-LMA.
When the H-LMA receives a packet which is destined to an HoA belonging to its HoA prefix and which is tunneled to it by an O-LMA, it checks its location information about this HoA. If the location information

indicates that the MN is currently anchored to a V-LMA, the V-LMA is the D-LMA serving the MN. The H-LMA tunnels the packet to the D-LMA and also sends this location information of HoA to the O-LMA.
When the O-LMA receives the new location information from the H-LMA indicating that the HoA is currently anchored to a D-LMA, it caches this information.
If the O-LMA has no activities related to this HoA, this location information in the cache memory will time out.
If an MN has recently moved from one D-LMA (previous D-LMA) to another D-LMA (new D-LMA), The new D-LMA will send the new location information of the HoA to both the H-LMA and the previous D-LMA.
When the previous D-LMA is informed that the MN has moved to another D-LMA, it caches this location information. This cache memory will time out when its timer expires.
When the previous D-LMA receives packets for the HoA from an O-LMA, it checks its cache memory about the new location information of the MN. If the cache memory has not timed out, it tunnels the packets to the new D-LMA. Meanwhile, it sends this new location information to the O-LMA.
When an LMA receives packets for an HoA which is not anchored to itself, it drops the packet unless it is the previous D-LMA for this HoA and the cache memory of the new location has not expired.

---

5. Packet flow

TOC

There are 3 cases of packet flow to be considered below: (1) sending packet to a mobile node from a non-mobile correspondent node; (2) sending packet from a mobile node to a non-mobile correspondent node; (3) sending packet from a mobile node to a mobile correspondent node.

5.1. Sending packets to a mobile node

TOC

When a correspondent node (CN) first attempts to communicate with the MN using the HoA of that MN, the packet is intercepted by the LMA nearest to that CN because all LMAs are advertising the same superset of IP prefixes using anycast. We call this originating LMA (O-LMA). This O-LMA uses the HoA to look up the H-LMA of the MN and then tunnels the packet to the H-LMA. The H-LMA receives the packet and de-capsulates it to read the HoA of the MN.
If the MN is in a visited network, the H-LMN tunnels the packets to the V-LMA to which the MN is currently anchored. This V-LMA is the D-LMA that will de-capsulate the packet and use the proxy care-of address (proxy CoA) to tunnel the packet to the mobile access gateway (MAG) to deliver to the MN. Figure 7 shows the destination address at the

network layer of the protocol stack of a first packet sent from the CN to the MN.
First packets

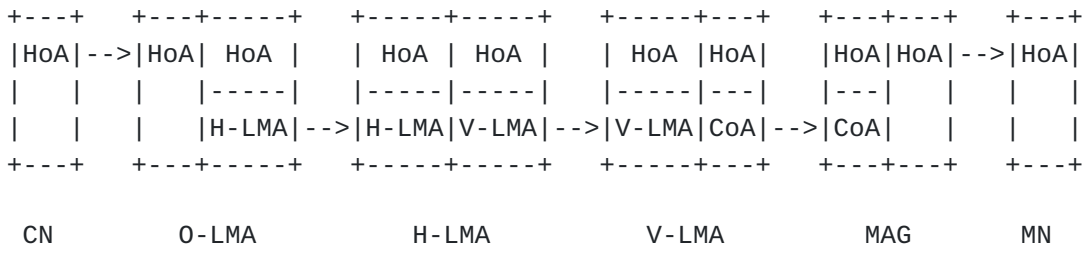

Figure 7. Network layer in the protocol stack of the first packet sent from a CN to a MN in a visited network showing the destination IP address as the packet traverses from the CN to the MN.
Only the first few packets from the CN may encounter triangle routing. When the H-LMA receives this first packet from the O-LMA and forwards this packet to the V-LMA (D-LMA), it also informs the O-LMA that the HoA is currently anchored to the D-LMA. The O-LMA keeps this location management information in a cache memory so that it may forward the packet directly to the D-LMA in future without going through the H-LMA. The D-LMA uses the proxy care-of address (proxy CoA) to tunnel the packets to the MAG to deliver to the MN (Figure 8).
Subsequent packets

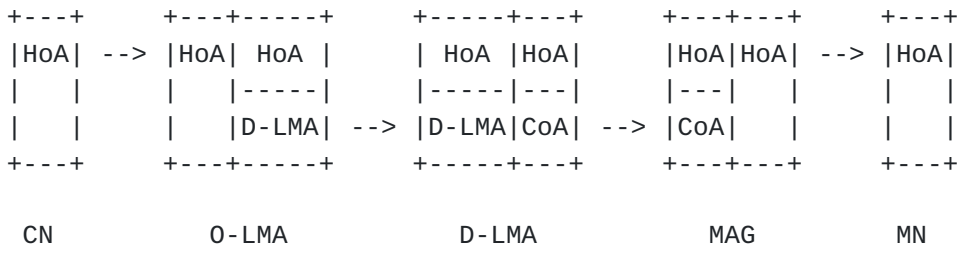

Figure 8. Network layer in the protocol stack of subsequent packets sent from a CN and tunneled to V-LMA in a visited network showing the destination IP address as the packet travses from the CN to the MN.
In the absence of traffic from the O-LMA to the HoA, the cache memory in the O-LMA may time out after a predefined period.

5.1.1. Bypassing D-LMA

TOC

In Figure 8, the packet is first tunneled from the O-LMA to the D-LMA and then tunneled from the D-LMA to the MAG. It is possible for the D-LMA to inform O-LMA the proxy-CoA. The O-LMA will keep this information in its cache memory so that it may tunnel future packets directly to the MAG (Figure 9).
Subsequent packets

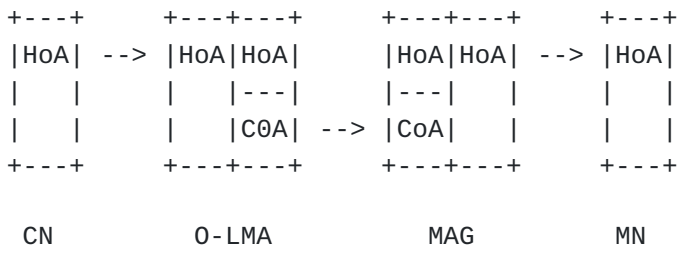

Figure 9. Network layer in the protocol stack of subsequent packets sent from the CN and tunneled to the MAG showing the destination IP address as the packet travses from the CN to the MN.

Bypassing O-LMA in Figure 9 is possible when the CN is also a mobile node, and this scenario will be discussed in Section 5.3 below. When the CN is a fixed node, bypassing O-LMA by tunneling directly between the CN and the MAG serving the MN using the proxy-CoA will expose the location information of the MN to the CN.

Alternatively, when the first packet from the CN to the MN reaches O-LMA, which tunnels this first packet to H-LMA, it is possible for H-LMA to tunnel the packet to the MAG without going through the D-LMA. Then the H-LMA may also inform the O-LMA to forward future packets to the MAG. It is possible for either the H-LMA or the D-LMA to take the primary responsibility to keep track of which MAG the MN is anchored to. If the D-LMA takes this responsibility, the H-LMA only needs to know which D-LMA the MN is anchored to. It is expected that an MN generally changes its D-LMA much less frequently than its MAG.

5.1.2. Changing MAG without changing D-LMA

TOC

It is possible for an MN to change its mobile access gateway (MAG) and proxy-CoA while anchoring to the same D-LMA. With no change of D-LMA, packets forwarded from the O-LMA to the D-LMA are unaffected.

The MAG may change from a previous MAG to a new MAG. As proxy-CoA subsequently changes, the D-LMA updates the mapping between HoA and proxy-CoA.

If the O-LMA has been tunneling directly to the previous MAG without going through the D-LMA, the previous MAG will need to tunnel the packet to the new MAG. Meanwhile, the previous MAG will inform the O-LMA to tunnel future packets directly to the new MAG.

5.1.3. Changing LMA

TOC

When the movement of a mobile node during an onging session necessitates a change in its local mobility anchor from a previous D-

LMA to a new D-LMA, the H-LMA will be notified to ensure that it has the correct location information. Any other LMA which is serving as an O-LMA may either forward packets to H-LMA or obtain the optimized routing information. Yet some LMA's may have cached the old location information and may continue to tunnel packets to the previous D-LMA. This situation may happen if some CN served by an O-LMA has sent packet to the MN earlier and the cache memory has not yet timed out. This situation may also happen when both the MN and the CN move and change LMAs at the same time.

We add a forwarding mechanism here. When the MN moves from one D-LMA to a new D-LMA, the new D-LMA may notify the previous D-LMA. If any packets destined to the MN reach the prevous D-LMA, the previous D-LMA will forward these packet to the new D-LMA. Meanwhile, the previous D-LMA will inform the O-LMA to tunnel future packets directly to the new D-LMA.

If the O-LMA is already tunneling directly to the previous MAG without going through the previous D-LMA, the previous MAG will need to tunnel the packets to the new MAG. Meanwhile, the previous MAG will inform the O-LMA to tunnel future packets directly to the new MAG.

5.2. Sending packets from a mobile node

TOC

The packets from a mobile node addressed to a correspondent node may go through the LMA to preserve location privacy. Figure 10 shows the source IP address of such a packet, which is tunneled to the O-LMA. This LMA is the closest LMA to which the MN is anchored to and will then send the packet to the correspondent node.

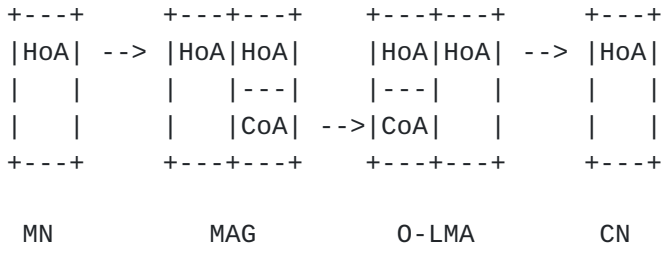

Figure 10. Network layer showing the source IP address as a packet travses from MN to CN.

5.3. Sending packets from a mobile node to another mobile node

TOC

We now consider the case of sending packets from a MN (MN1) to another MN (MN2). Packet sent from MN1 will first be tunneled from MAG to O-LMA as in Figure 10. The route from O-LMA to MN2 will be the same as that

in Figure 7 and Figure 8 for respectively the first packets and subsequent packets.
For the subsequent packets, the traversed route will be: MN1-MAG1-OLMA-DLMA-MAG2-MN2.
It is possible to bypass the O-LMA and/or the D-LMA, and such scenarios will be discussed next.

5.4. Bypassing the tunneling role of the LMA

TOC

With distributed LMA, the sender and receiver are respectivley using their nearest LMA's so that unnecessarily long routes are already prevented. Therefore the purpose of bypassing the an LMA is not in avoiding unnecessarily long route, and bypassing the LMA will likely not reduce the route much further.
The primary purpose of bypassing an LMA is then in removing the need for the LMA to de-capsulate the tunneling header of an incoming packet and to encapsulate the packet again. Here bypassing an LMA in tunneling packets refers to removing the LMA's participation in the tunneling process. The additional steps of ending a tunnel (decapsulation) and starting another tunnel (encapsulation) at the LMA are avoided. Note that the route through the LMA may often be the shortest route, so the route may still pass through the LMA even when LMA is not involved in tunneling.
When both the sender and the receiver are mobile nodes and are attached to networks running PMIP, the packet goes through an MAG (MAG1) at the sending side and through another MAG (MAG2) at the receiving side. It is possible to bypass the O-LMA and/or the D-LMA.
In general, the network needs to know whether bypassing the LMA's will affect location privacy. This capability may reside mainly at the LMA so that the MAG may perform simpler functions only.
Tunneling a packet once between 2 network nodes may be sufficient to protect the location privacy of a mobile node. When the route of a packet involves successive tunneling segments, it is possible to cosolidate the tunneling segments. For a packet traversing from source to destination, if the input side of an LMA is the end of a tunneling segment and the output side of the LMA is the beginning of another tunneling segment, the input tunneling segment and the output tunneling segment may be cosolidated.
In the route: MN1-MAG1-OLMA-DLMA-MAG2-MN2, there are three tunneling segments: between MAG1 and O-LMA, between O-LMA and D-LMA, and between D-LMA and MAG2. The D-LMA is receiving a packet in which the originating IP address is hidden in a tunnel terminating at D-LMA and the destination address is to be hidden in a tunnel it will use from itself to MAG2. The D-LMA may inform the O-LMA to tunnel future packets directly to MAG2. After bypassing the D-LMA this way, the tunnel segments between MAG1 and O-LMA and between O-LMA and MAG2 in the route

MN1-MAG1-OLMA-MAG2-MN2 are still hidding the locations of both the source and destination nodes.
In a similar manner, the O-LMA may determine whether bypassing itself may affect location privacy. If not, it may inform MAG1 the proxy-CoA2 of MAG2 so that MAG1 may tunnel subsequent packets to MAG2. The route then becomes: MN1-MAG1-MAG2-MN2 (Figure 11) which has one tunnel between MAG1 and MAG2 protection the location privacy of the mobile nodes.

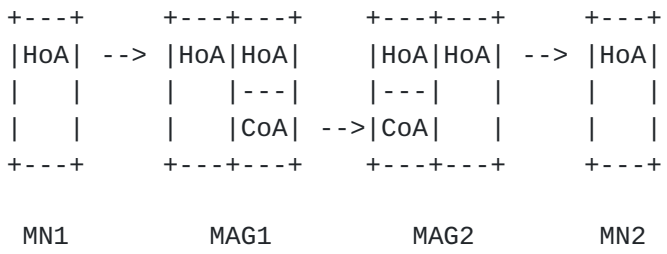

Figure 11. Network layer showing the source IP address as a packet travses from a MN to a CN which is also mobile.
Here, the tunnel from MAG1 to MAG2 is using the destination address proxy-CoA2 and the source address proxy-CoA1. The MAG's are still protecting the location privacy of the end nodes provided the end nodes are not allowed to access these information from their MAG's.
When sending packet from a fixed node to a mobile node using a route: CN-OLMA-DLMA-MAG-MN, there are two tunneling segments: between O-LMA and DLMA, and between D-LMA and MAG. It is possible to bypass the D-LMA as discussed in Section 5.1.1 above. The resulting route, as shown in Figure 9 above, becomes CN-OLMA-MAG-MN. In order to preserve location privacy of MN, it is not desirable to also bypass the O-LMA. Yet it is noted that such a route which does not bypass the tunneling role of the O-LMA is already a short route using the LMA closest to the CN.
In the special case that the O-LMA and the D-LMA are the same LMA, i.e., when the MN and the fixed CN are in the same network, it is usually necessary not to bypass the tunneling role of LMA in order to preverse location privacy of the MN.
When sending packets from a mobile node to a fixed node using a route: MN-MAG-OLMA-CN as shown in Figure 9 above, there is only one tunnel between the MAG and the O-LMA. The O-LMA decapsulated the tunnel header to the send the original packet with the home address (HoA) of MN as the source address to the fixed node. In order to preserve location privacy of MN, it is not desirable to bypass the OLMA. It is noted again that this route which does not bypass the tunneling role of the O-LMA is already a short route using the LMA closest to the MN.

---

6. Performance

TOC

6.1. Round trip time

TOC

In this proposal, the O-LMA will behave like a full functioned LMA for the MN once it has acquired and cached the location management information to optimize routing. The route from the CN to the MN for later packets is the same as for migrating home agents. It is therefore reasonable to say that the round trip time after the first packets is comparable to that of migrating home agents; and experiments with migrating home agents had already reported round trip times approaching that of direct routes between CN and MN [MHA].

6.2. Call setup delay

TOC

Only the first packet or first few packets may, but not always, encounter triangle routing. It is possible to query the H-LMA before sending the first packet. Alternately, a V-LMA may simply route the first packet to the packet's H-LMA.
Here, each H-LMA is responsible for its own block of IP addresses in a network to allocate to the mobile nodes registering to that network. Every LMA may be informed of which H-LMA is responsible for which address block. In other words, the O-LMA does not lack routing information even for the first packets. It lacks only optimized-route informing. Without such information, O-LMA already knows which H-LMA is managing the HoA and may therefore immediately forward the first packets to the H-LMA without waiting for information acquisition. The routing path going through the H-LMA here is comparable to that in the case of mobile IP using the home agent in the home network only.
Triangle routing is encountered only for the first packets and only for certain configurations where the mobile node and the corresponding node are both far from the home network but are close to each other. This is a small price for not pushing the full location management information to all the other home agents for synchronization purpose.
The possible delay for the first packets may affect only the call setup delay. Many communication applications go through a call set up process to begin a communication session. This call setup delay customarily experienced is usually longer than the typical packet delay in an ongoing communication session. Compared with pushing mobility management information to all home agents, the distributed local mobility anchors differ only in a possible triangle routing for the first packets which may be a small overhead added only to the call setup delay.

TOC

6.3. Location update signaling overhead

Consider that there are n mobile nodes and m LMAs. If the location management information is pushed to all the LMAs, the amount of signaling in pushing the location management information to all the LMAs is proportional to n x m.
By keeping the master information at the H-LMA without pushing it to all the LMAs, the amount of signaling is proportional to n only.

6.4. Simultaneous move problem

TOC

Because both the source and the destination nodes may be mobile, they may be changing their LMA's at the same time. In this case, the mobile node has moved from a previous LMA to a new LMA while the correspondent node has also changed its O-LMA. The O-LMA of the correspondent node may be using outdated cache information to route packets to previous D-LMA. The new D-LMA may inform the previous D-LMA to forward these packets to the new D-LMA. Meanwhile, the previous D-LMA may inform the O-LMA to route any future packets directly to the new D-LMA.

---

7. Interworking with legacy LMAs

TOC

The use of anycast and the need for trust relationship among the H-LMA's and V-LMA's in distributed LMA may be less challenging in a domain where the networks belong to one autonomous system or to one service provider. The service provider network may cover large and fragmented geographical areas. The MN in this domain will be reachable from any CN within this same domain. However, the CN may belong to a different domain or to a legacy network not supporting distributed LMA. This section addresses these interworking issues.

7.1. Reachability from CN outside MN's domain

TOC

Figure 12 shows a configuration in which a CN is in a network outside the distributed-LMA domain of a MN. When the CN sends a packet to the MN, the CN's network recognizes that the IP prefix of the destination address belongs to a different domain. The network may find the shortest path for the packet to exit itself (hot potato routing) and let the packet be routed through the transit core network to the MN's domain. As it enters the MN's domain, the use of anycast in this distributed-LMA domain will cause the packet to reach the nearest LMA.

This LMA is now serving as originating LMA (O-LMA) to route the packet to the MN.

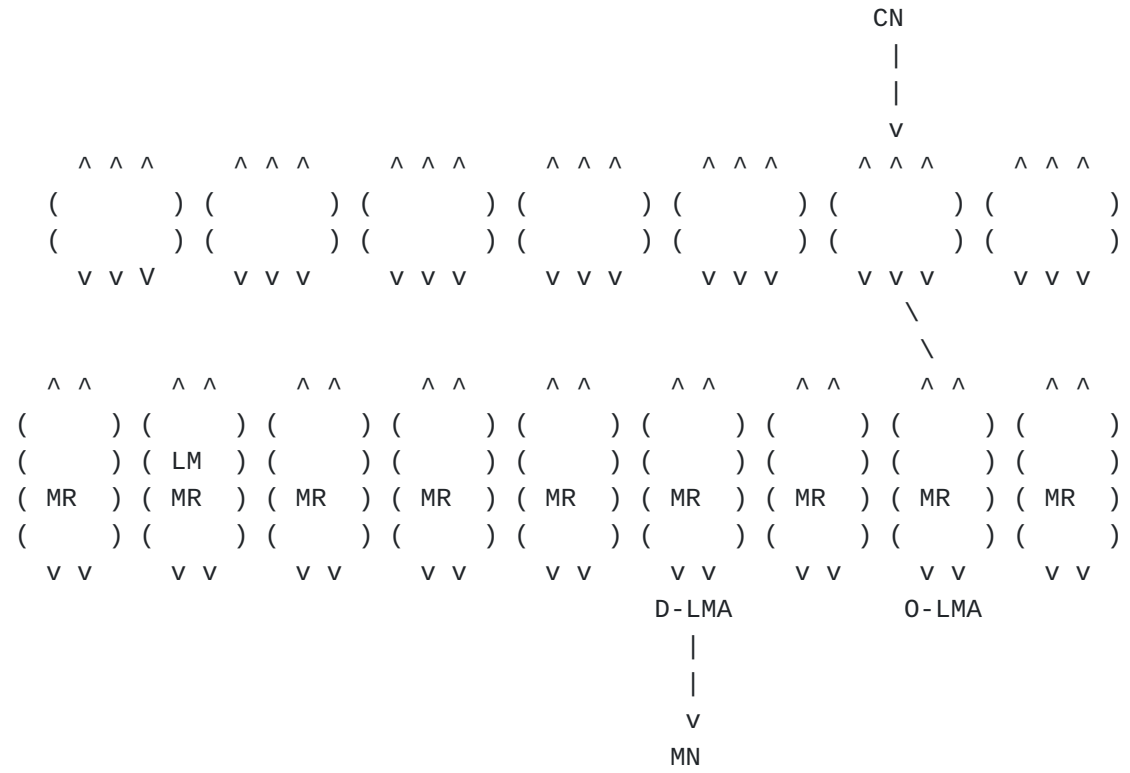

Figure 12. Configuration showing O-LMA and D-LMA for a packet sent by CN from a network outside the domain of MN.
Because routing is based on the destination address and as long as the CN is outside MN's distributed LMA domain, the above process is the same regardless of whether the CN is in a different distributed LMA domain, is in a different PMIP domain, is running MIP, or is a fixed node.

7.2. Sending packet to CN in a different distributed-LMA domain

TOC

Figure 13 shows a configuration in which the CN is in a different distributed-LMA domain than that of the MN. For a packet sent from the MN to the CN, the MN's domain recognizes that the IP prefix of the destination address belongs to a different domain. The network may also find the shortest path for the packet to exit itself (hot potato routing) and let the packet be routed through the transit core network to the CN's domain. As it enters the CN's domain, the use of anycast in this distributed-LMA domain will cause the packet to reach the nearest LMA. This LMA is now serving as O-LMA to route the packet to the LMA closest to the CN, which is the D-LMA serving the CN.

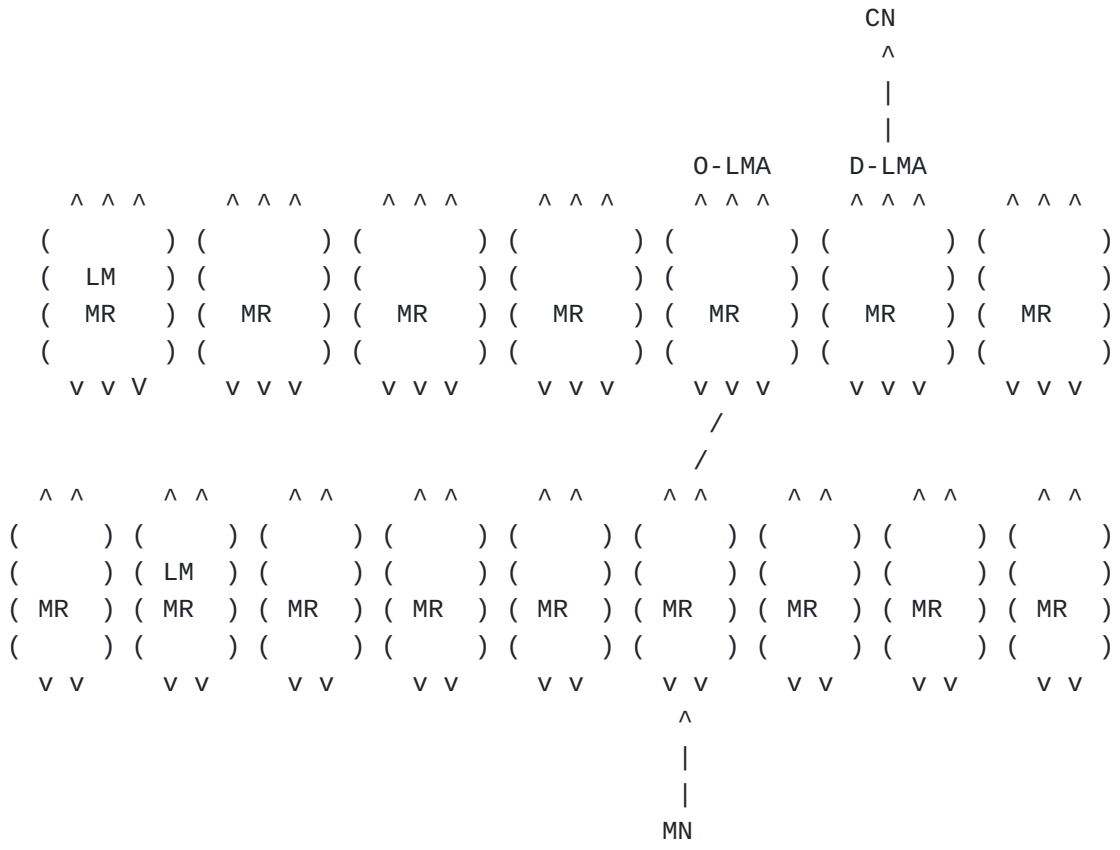

Figure 13. Configuration showing O-LMA and D-LMA for a packet sent by MN to CN in a different distributed LMA domain.

7.3. Sending packet to a CN in a legacy PMIP domain

TOC

Figure 14 shows a configuration in which CN is in a PMIP domain not supporting distributed LMA. The CN may be physically located inside or outside the MN's distributed LMA domain. If the CN is in a network outside MN's distributed LMA domain, for a packet sent from the MN to the CN, the MN's domain recognizes that the IP prefix of the destination address belongs to a different domain. The network may also find the shortest path for the packet to exit itself (hot potato routing) and let the packet be routed to the CN's network domain through the transit core network. As it enters the CN's PMIP domain, the packet will follow the PMIP routing in that domain. It may encounter triangle routing. If it uses route optimization in the PMIP domain, the CN may lose its location privacy to the MN.
It may be possible that the CN's access network is in both the MN's distributed LMA domain and the CN's own PMIP domain. Yet CN's IP prefix address is derived from the legacy LMA in CN's PMIP domain. The network will therefore route the packet to this legacy LMA, which does not

support distributed LMA. It may again encounter triangle routing. If it uses route optimization in the PMIP domain, the CN may not be able to hide its location from the MN.

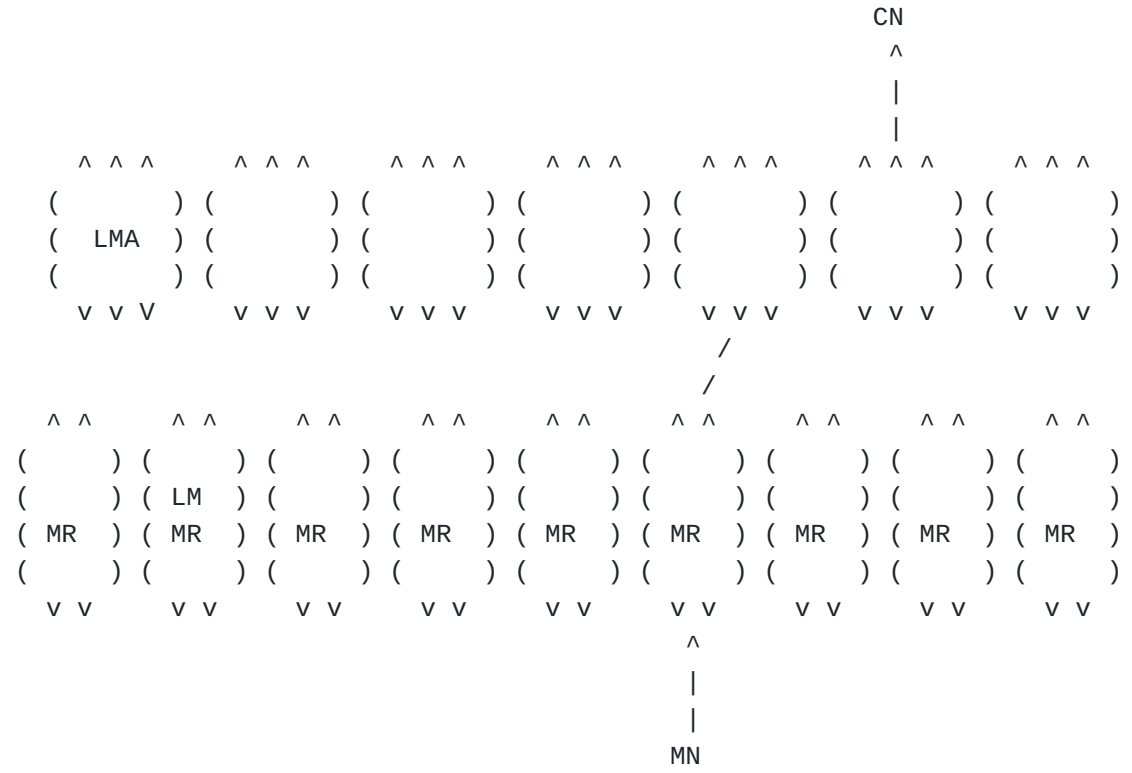

Figure 14. Configuration of sending packet from MN from a distributed-LMA domain to CN belonging to a legacy PMIP domain.

7.4. Sending packet to CN running MIP outside MN's distributed-LMA domain

TOC

Figure 15 shows a configuration in which the CN is running MIP. The IP address of the CN is derived from the CN's home network. The CN may be physically located inside or outside the MN's distributed-LMA domain. In either case, the packet is routed to the CN's home network where it is intercepted by the CN's HA. Again, the packet may encounter triangle routing. If it uses route optimization defined in MIP, the CN may not be able to hide its location from the MN.

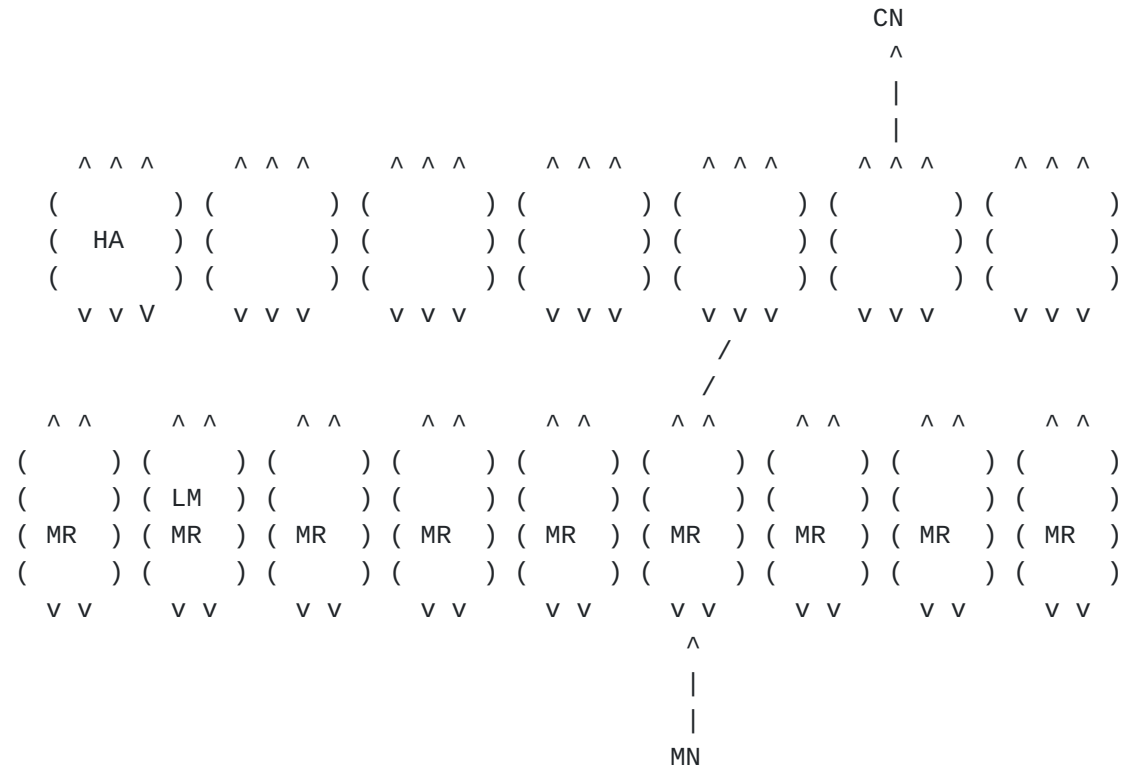

Figure 15. Configuration of sending packet from MN to CN running MIP but not supporting distributed LMA.

---

8. IANA Considerations

TOC

This document does not require any IANA actions.

---

9. Security Considerations

TOC

As in PMIP, a trust relationship is needed among the LMAs. When the O-LMA tunnels packets back to the H-LMA, the security considerations are not different from that in PMIP.
Untrusted LMAs make the network vulnerable to various attacks. An untrusted LMA may tunnel many packets to the D-LMA causing DOS attacks. With route optimization, the H-LMA may send location information to the O-LMA which will use this information to tunnel packets directly to the D-LMA. The trust relationship between the H-LMA and the O-LMA and the protection of the location information messages are important. The protection mechanisms needed are similar to those of proxy binding updates in [GHAHA].

When the MN moves from one D-LMA to a new D-LMA, the lack of secure mechanism in sending location information update from the new D-LMA to the previous D-LMA may enable a rogue LMA to hijack the traffic. Proper trust relationships among LMAs and secured mechanisms are needed to protect these messages. These mechanisms are similar to those needed in [GHAHA].

10. Acknowledgments

TOC

This document has benefited from discussions with Da Peng Liu, Zhen Cao, Xiaoyan Jiang, and others.

11. References

TOC

11.1. Normative References

TOC

[RFC3775]	Johnson, D., Perkins, C., and J. Arkko, “Mobility Support in IPv6,” RFC 3775, June 2004 (TXT).
[RFC5213]	Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, “Proxy Mobile IPv6,” RFC 5213, August 2008 (TXT).

11.2. Informative References

TOC

[GHAHA]	Thubert, P., Wakikawa, R., and V. Devarapalli, “Global HA to HA protocol,” draft-thubert-mext-global-haha-01 (work in progress), July 2009 (TXT).
[MHA]	Wakikawa, R., Valadon, G., and J. Murai, “Migrating Home Agents Towards Internet-scale Mobility Deployments,” Proceedings of the ACM 2nd CoNEXT Conference on Future Networking Technologies, Lisboa, Portugal, December 2006.
[SMGI]	Zhang, L., Wakikawa, R., and Z. Zhu, “Support Mobility in the Global Internet,” Proceedings of ACM Workshop on MICNET, MobiCom 2009, Beijing, China, September 2009.

Authors' Addresses

TOC

	H Anthony Chan
	Huawei Technologies
	1700 Alma Ave
	Plano, TX 75075
	USA
Email:	anthonychan@huawei.com
	Frank Xia
	Huawei Technologies
	1700 Alma Ave
	Plano, TX 75075
	USA
Email:	xiayangsong@huawei.com
	Justin Xiang
	Huawei Technologies
	1700 Alma Ave
	Plano, TX 75075
	USA
Email:	zengjun.xiang@huawei.com
	Hanan Ahmed
	Huawei Technologies
	10180 Telesis Ct. Suite 365
	San Diego, CA 92121
	USA
Email:	ahanan@huawei.com