A Flexible and Distributed Home Agent Architecture for Mobile IPv6-Based Networks

Albert Cabellos-Aparicio and Jordi Domingo-Pascual*

Technical University of Catalonia,
Department of Computer Architecture,
Advanced Broadband Communications Center,
c/Jordi Girona, 1-3, 08034 Barcelona, Spain
{acabello,jordid}@ac.upc.edu.com
http://www.ccaba.upc.edu

Abstract. Home Agents (HA) represent a single point of failure for Mobile IPv6-based networks. To overcome this problem many solutions have been published providing reliable HA architectures. These solutions require deploying redundant HAs on each sub-network. Although these solutions effectively mitigate this problem, they do not take into account the requirements of large networks with dozens of sub-networks. Deploying several HAs on each sub-network may be too expensive to deploy and to manage. In this paper we present a novel HA architecture that only requires a set of HAs for the whole network. Our basic idea is that the Mobile Node's location can be announced to exit routers, this way re-directing packets can be done without involving the HA. Our solution provides reliability and load balancing as the existing solutions. Finally, we validate our proposal through an analytical model and compare it against other proposals through a simulation.

Keywords: Mobility, Mobile IPv6, Home Agent Architecture.

1 Introduction

Mobile IPv6 [1] is considered to be one of the key technologies to provide mobility to the Internet. With "mobility" a user can move and change his point of attachment to the Internet without losing his network connections.

In Mobile IPv6 a Mobile Node (MN) has two IP addresses. The first one identifies the MN's identity (Home Address) while the second one identifies the MN's current location (Care-of Address). The MN will always be reachable through its Home Address while it will change its Care-of Address according to its movements. A special entity called a Home Agent (HA) placed at the MN's home network maintains bindings between the MN's Home and Care-of Addresses.

* This work was partially funded by IST under contract IST-2006-NoE-0384239 (IST-CONTENT), MEC (Spanish Ministry of Education and Science) under contract TSI 2005-07520-C03-02 and the CIRIT (Catalan Research Council) under contract 2005 SGR 00481.

I.F. Akyildiz et al. (Eds.): NETWORKING 2007, LNCS 4479, pp. 333–344, 2007.

In addition the communications between the MN and its peers (Correspondent Nodes) are routed through the HA. Thus, the MN relies on its HA for its connectivity. However, Mobile IPv6 incorporates a route optimization mechanism where the MN can communicate directly with its Correspondent Nodes (CN). This mechanism avoids triangle-routing through the HA reducing the HA's load. However, this will not be used for short-term communications (e.g a MN accessing a web page).

A HA may be responsible for multiple MNs on a Home Link. The failure of a single HA may then result in the loss of connectivity of numerous MNs. Thus, HAs represent the possibility of a single point of failure for a Mobile IPv6-based network. Moreover, MN's communications through the HA may also lead to either the HA or the Home Link becoming the bottleneck of the system. In addition, the HA's operations such as security check, packet interception and tunneling might not be as optimized in the HA software as plain packet forwarding.

The Mobile IPv6 standard allows the deployment of multiple HAs on the Home Link to provide reliability and load balancing. This is done so that upon the failure of the serving HA another HA can take over the functions of the failed one. This provides continuous service to the MNs registered with the failed HA. However, the transfer of service is problematic [2]. The solution is MN-driven and forces the MN to detect the failure and select a new HA. This causes delayed failure detection, service interruption in the upper layer applications, increased workload on the MN, message overhead over the air interface and IPsec Security Associations re-establishment.

Many research papers have been published that address these problems. The solutions presented in [4][5][6][7][8] increase HA reliability and load balancing by deploying several redundant HAs at the Home Link. In these solutions, all the HAs share the registration state and they define efficient mechanisms for HA recovery. These solutions reduce the service disruption time in front of Mobile IPv6. In addition, the MN's traffic is balanced among the different HAs. The main difference among them is that some [4][5][6] are MN-driven solutions while others [7][8] are transparent to the MN.

Unfortunately, these proposals are focused on providing HA reliability and load balancing on just a single Home Link but they do not take into account the global requirements of an Autonomous System (AS). An AS that hosts MNs may have dozens of sub-networks. Deploying reliable HAs requires several redundant HAs on each link. It is important to remark that the Mobile IPv6 protocol belongs to the IPv6 standard and, theoretically, any IPv6 node has mobility capabilities. Thus, these approaches are too expensive to deploy and to manage.

A different proposal, which does not require deploying redundant HAs on each Home Link, is the Virtual Mobility Control Domain protocol (VMCD) [9][10]. The VMCD protocol allows multiple HAs to be placed at different domains. A MN may use multiple HAs simultaneously. The basic idea behind this proposal is that each HA advertises, through eBGP, the same home network prefix from multiple routing domains. Each MN then picks the best HA according to its

topological position. The main drawback for this proposal is that the impact on the exterior BGP routing system scalability is unpredictable.

In this paper, we present a novel flexible and distributed HA architecture that takes into account the mobility requirements of an AS and that does not impact the BGP routing system. We consider the HA as an entity that performs several differentiated operations. We have analyzed each operation and we have assigned each of them to an entity of the network. Our basic idea to distribute the operations is that a registration from a MN into a HA can be viewed as an internal route from the network's point of view. That is, when a MN registers a new location into its HA it is actually installing a new route *(Home Address → Care-of Address)*. We believe that this route can be announced throughout the network and thus, it is not necessary to deploy a HA on each link. As we will see, our solution only requires deploying one HA for the whole network. This HA should be reliable and our architecture allows deploying more than one HA to distribute the load. Moreover, our solution can use the redundant mechanisms presented in [4][5][6][7][8]. In addition, our solution reduces considerably the number of MN's data packets transmitted into the network and is compatible with legacy MNs.

2 Introduction

2.1 Design Rationale

In this subsection, we will analyze the different operations of a Home Agent (HA) and how they can be distributed from a network's point of view. In the rest of the paper, we will use the following terminology: we define *Home Network* as the set of Home Links managed by our HA. We define *Exit Routers* (ER) as the routers that connect the Home Network with the rest of the Internet. These ERs may or may not be the AS's border routers and an AS may have several Home Networks.

Home Agents are responsible for maintaining bindings between the MN's identity and its location. The HAs forward the MN's signaling and the MN's data packets as well. MNs send data packets through their HA when communicating with their Home Network or with CNs. Since MNs can communicate directly with its CNs it is expected that communications through the HA are mainly used for short-term connections.

The Mobile IPv6 RFC states that packets sent through the HA may be secured through IPSec [3]. It should be taken into account that the MN can use IPSec with its peers regardless of the IPSec connection with their HA. We believe that it is not useful to secure MN to CN communications because the packets are only secured on half of the path (MN → HA) while the rest of the path (HA → CN) is not secured. Regarding the MN's communications with the Home Network, protecting the path is useful. In this case the HA is acting like a Virtual Private Network (VPN) gateway.

Under these assumptions and following the basic idea that a registration from a MN into a HA can be viewed as an internal route we can distribute the HA's

operations throughout the network. In our architecture, a single HA is required for the whole network; we call it a *flexible* Home Agent (fHA). This fHA will process (using IPSec) the MN's signaling messages and will maintain registration information. It will also distribute this information throughout the network as internal routes. The network will directly process the MN's communications with the CNs while the fHA will process the MN's communications with the Home Network (using IPSec) in the same way as a VPN gateway.

2.2 Overview

Fig. 1 presents an overview of our architecture. Our proposal has only one HA (we call it a fHA) that will serve all of the MNs of the network. Take into account that our proposal allows more than one fHA (section 2.3) to be deployed to distribute the load. This fHA will be identified by an unicast address and the MNs will address its registration messages to it. Upon reception of a registration message, the fHA validates it and sends a routing message announcing the new route towards the MN. This information is then sent to each ER. In addition, the fHA advertises the route to the Home Link's Access Router (AR). At this point, the network knows the location of the MN.

When communicating with a CN through the HA, MNs do not address packets to the fHA but to an anycast [13] address owned by the ERs. For instance this anycast address can be configured on a Loopback interface. In this way, a given ER receives the MN's data packets and de-capsulate, lookup and forward packets to the CN. Similarly, CNs send packets to the MN's Home Address. Upon reception, the ER lookups the packet's destination address (the MN's Home Address). Since the fHA has previously installed a route at the ERs they know that the MN it is not at home. Therefore, the ERs encapsulate and forward the packet to the MN's location. Our architecture manages efficiently MN to CN communications because some packets "bounce" at the ERs. This way the network's internal traffic is reduced considerably.

Regarding the communications from the MN to the Home Network, the MNs addresses its IPSec protected packets to the fHA that, in turn, de-capsulate and forward them to the MN's peer. The MN's peers address its data packets to the MN's Home Address. Since the fHA has announced to the Home Link's AR a route for the MN, the AR knows that the MN is away and it encapsulate the packet towards the fHA.

The Home Link's AR also multicasts Neighbor Advertisement messages on behalf the MN. This enables the AR to intercept communications from the Home Link to the MN and forwards them through the tunnel with the fHA.

In the following subsections the detailed operations of our architecture are presented.

2.3 Dynamic fHA Address Discovery

This subsection specifies how the fHA announces their presence. In standard Mobile IPv6 HAs announce their presence through Router Advertisement messages.

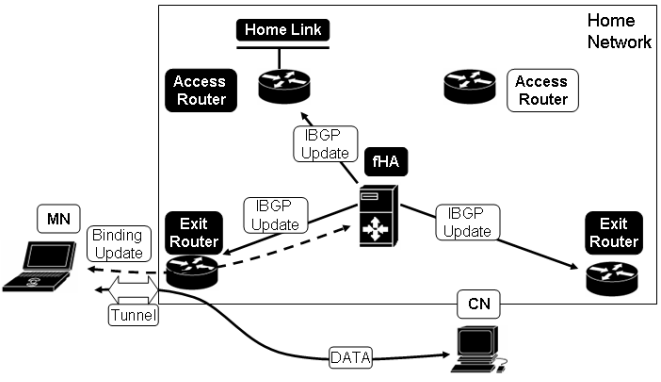

Fig. 1. Overview of our proposal

In this way, the MN's can automatically select a HA. Our architecture implements this functionality in exactly the same way that Mobility Anchor Points (MAP) announce their presence in the Hierarchical Mobile IPv6 (HMIPv6) [12] protocol. Our mechanism is also compatible with legacy MNs.

Each fHA sends Router Advertisement messages announcing its presence to the routers operating in the network. These messages include a preference value. In turn, the routers propagate the fHA's announcements to ARs that then forward them to the Home Link. Each router will decrement the preference value. This way MNs can automatically discover their fHA's address and select the best one according to the preference value.

This mechanism has many benefits. On the one hand, it enables ARs to automatically discover the fHA thus avoiding manual configuration. On the other hand, it allows us to deploy more than one fHA on the network and distribute the load among them.

The fHA's Router Advertisement messages include the prefix(es) of the Home Network that it is serving and the anycast address owned by the ERs. Including the Home Network's prefix enables the MNs to know if its peers are on the Home Network or not. Depending if the peer is on the Home Network or not MNs will address the data packets to the fHA or to the ERs.

Finally, in order to provide compatibility with legacy Mobile IPv6 nodes, MNs may send its traffic to the fHA.

2.4 Signaling Processing

Each MN selects a given fHA through the above-mentioned mechanism. All the fHAs have pre-configured keys with the MNs as the Mobile IPv6 RFC states. Please note that ARs and ERs do not share any keys with the MNs. The fHAs receive registration messages from the MNs as stated by the Mobile IPv6 RFC.

Upon reception of a successful registration message, the fHA has to announce this information (route) to the ERs, to the Home Link's AR and to the rest of the

fHAs. To distribute this type of information we use a routing protocol. Instead of designing a new routing protocol we use an already existing and deployed one. The routing protocol that best fulfills our requirements is the interior Border Gateway Protocol (IBGP) [11]. In our solution the fHAs, the ERs and Home Link's ARs create an IBGP domain. It is very important to remark that this IBGP domain may be an already existing IBGP domain or a separate one. The routes announced through this IBGP domain always have the longest prefix (/128) and never affect regular BGP routes. It should be noted that the routes announced by the fHAs will never be distributed outside the network. Finally, the entities participating in the IBGP domain have pre-configured keys to provide confidentiality, integrity and authentication to the communications.

For each successful received registration message, the fHAs send an IBGP UPDATE message to the ERs and to the AR responsible of the MN's Home Link. The fHAs are able to determine the appropriate AR by inspecting the MN's Home Address.

We introduce new options in the IBGP UPDATE message. The UPDATE message sent to ERs includes the following information: *(Home Address, Care-of Address, Lifetime)*. Upon reception of this message, the ERs setup a tunnel endpoint with the MN. The tunnel source address is the anycast address while the destination address is the Care-of Address. In addition, each ER adds the following route to its routing table: *(HomeAddress/128 → Tunnel)*. The tunnel and the route are automatically deleted after *"Lifetime"* seconds.

The UPDATE message sent to the AR includes the following information: *(Home Address, Lifetime)*. Upon reception of this message, the AR knows that the MN is away from home (note that the AR does not know the location of the MN). Next, the AR setups a tunnel endpoint towards the fHA that announced the route and adds the following route to its routing table: *(HomeAddress/128 → Tunnel)*. The AR also starts sending multicast Neighbor Advertisement messages on behalf of the MN at the Home Link. If a node of the Home Network (or Home Link) sends a packet to the MN, the AR intercepts it and encapsulates it towards the fHA. Once again, the tunnel and the route are automatically deleted after *"Lifetime"* seconds.

Once the MN returns home it sends a registration message to the fHA. Upon reception, the fHA sends an IBGP WITHDRAWAL message to the ERs and to the corresponding AR to immediately remove all the routes and tunnels related to the MN's Home Address.

2.5 Data Packets Processing

This subsection presents how packets are routed from/to the MNs.

MNs communicating with CNs encapsulate their data packets to the anycast address owned by the ERs (Fig. 2). The packets are received by the "nearest" ER that will de-capsulate and forward them towards the packet's destination address (the CN's address). If the exit point of the CN's address is another ER then the packet traverses the network as a transit packet. It is important to remark that our solution does not require anycast routing. Packets addressed to

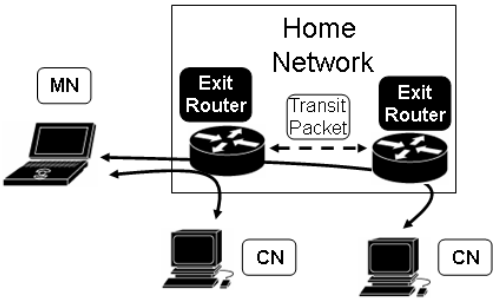

Fig. 2. MNs to CNs communications

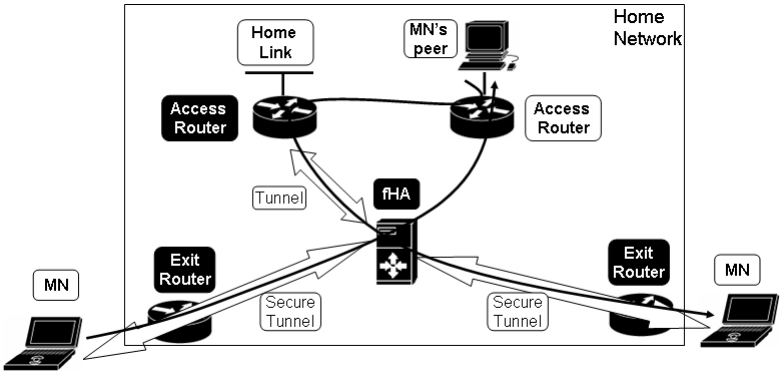

Fig. 3. MNs to Home Network communications

the anycast address are routed normally (like unicast) and delivered to a given ER. We use anycast addresses because it is the standard procedure to assign the same address to different network interfaces.

MNs communicating with nodes located into their Home Network (Fig. 3) encapsulate their packets towards the fHA. However, packets sent by MN's peers are addressed to the MN's Home Address. The MN's AR intercepts those packets. Since the AR knows that the MN is away from home, it encapsulates the packet towards the fHA. Since the Mobile IPv6 RFC states that the packets are tunneled through the HA encapsulating the packet from the AR to the fHA does not affect the path's MTU [1]. As has already been mentioned, the MN's communications with the Home Networks are protected with IPSec.

2.6 Flexible Home Agent Location

This subsection discusses the possible locations of the fHAs. Each fHA can be placed anywhere in the network, as a separate server, co-located with an ER/border router or even with a BGP Route Reflector.

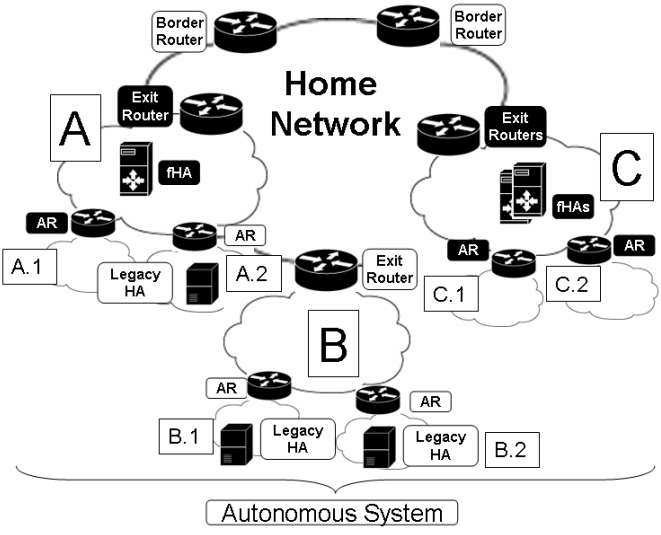

Fig. 4. fHA location example

One of the major benefits of our proposal is its flexibility. On the one hand, our architecture can serve all the MNs of a network with one or more fHAs. If more than one fHA is deployed MNs will select the nearest one based on the preference value. This way the load is distributed among them. Each fHA thus only process signaling messages and communications from/to the Home Network (like a VPN gateway). MNs to CNs communications are then processed by ERs. On the other hand, our architecture is transparent to MNs running with legacy Home Agents and both technologies may co-exist on the same network.

Fig. 4 shows an example of the flexibility of our architecture. This AS has three networks and each one can independently select which approach it deploys. For instance, the *"A"* network can deploy both technologies. The fHA could serve MNs belonging to the *"A.1"* sub-network while MNs belonging to the *"A.2"* sub-network could be served by a legacy HA. The *"B"* network can deploy only legacy Home Agents on each sub-network. Finally, the *"C"* network can deploy two fHAs and all the MNs from *"C.1"* and *"C.2"* could be served by them.

Only routers labeled in black must belong to the IBGP domain with the fHAs of their network. There will be a separate IBGP domain for each Home Network. MNs served by an fHA send its data packets to an anycast address owned by the ERs. Since the prefix of the anycast address belongs to the Home Network's prefix, the AS's border routers knows how to forward the packets and do not need to be aware of our protocol.

3 Evaluation

This section presents an analytical evaluation of our proposed scheme and a comparison with a network running Mobile IPv6 enhanced with existing solutions

[4][5][6][7][8]. We do not consider solutions based on eBGP [9][10] because their impact on the exterior BGP routing system scalability is unpredictable.

3.1 Signaling Overhead

Let N be the number of ERs of a network that is running our proposal, let M be the number of deployed fHAs and let H be the total number of received registration messages per second (including foreign and home network registrations). Then our proposal requires sending $H(N+M)$ IBGP messages per second.

3.2 Transit Traffic Reduction

As has been commented previously, in our proposal some data packets will "bounce" at the network's ER without being transmitted through the network. However in existing solutions [4-8] each packet sent through the HA has to be transmitted twice. One from the ER to the HA and another one in the opposite direction. In this subsection we compare this amount of transit traffic. We only consider the traffic exchanged between MNs and CNs that is routed through the HA.

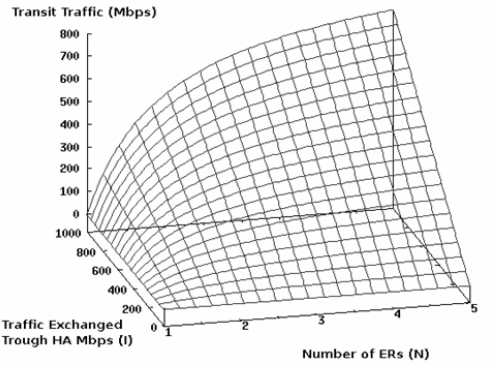

Fig. 5. Transit Traffic in our proposal

Let I be the Kbps of traffic exchanged between all the MNs and its CNs through the HA. Then, existing solutions [4-8] have $2I$ Kbps of transit traffic. If we assume that each ER of the network has the same probability of being the exit point of a given packet then, our proposal has $(1\text{-}1/N)I$ Kbps of transit traffic (Fig. 5).

In addition, transit traffic in existing solutions [4-8] may follow a longer path than in our proposal. While in existing solutions [4-8] transit traffic must be transmitted to the HA in our proposal some transit traffic "bounces" at the ERs and the rest is transmitted from one ER to another. Home Links are usually deployed far away from the ERs while ERs may be close to one another (in terms of number of hops).

3.3 Stored State at the Routers

In this subsection we will analyze the size of the routing tables and the number of tunnels configured at the ERs and ARs of a network running our proposal.

Each ER has 1 tunnel and 1 route for each MN of its Home Network that is away from home. Each route and tunnel requires the Home Address, the Care-of Address and a lifetime, in total 34 bytes. Likewise, each AR will have just 1 tunnel with each fHA and 1 route for each of its nodes away from home.

4 Simulation

In order to validate our proposal we have run a simulation. The simulation is intended to provide realistic values for the equations presented in section 3 and to compare our proposal with existing solutions [4-8].

In order to provide realistic values, we have configured a highly mobile environment by using a Random Trip mobility model [14]. Specifically, we have used the Random Waypoint on Generalized Domain model with a set of 8 domains. Each domain represents a layer-2 network where a MN can move without changing its point of attachment (i.e default router). Only when the MN changes from one domain to another it must register its new location. Please, refer to [14] for further information.

The first domain is considered to be the Home Network while the rest of the domains are foreign networks. The Home Network has 1000 MNs, 2 ERs and 5 sub-networks. When running our proposal the Home Network has 2 fHAs while when running existing solutions [4-8] the Home Network has a set of reliable HAs on each sub-network (5 sets in total). In addition, each MN sends 64Kbps (VoIP) of unidirectional traffic towards its Home Network and 128 Kbps (Data) towards a CN. It should be taken into account that when a MN is at home traffic is sent directly and thus we do not consider it. Similarly, we do not take into account route optimized traffic. Finally, we have simulated this environment during 10000 seconds (roughly 2.7 hours).

Our mobility model produces a mean of 4.68 foreign network registration messages per second (messages/s) and 0.80 Home Network registration messages/s. This means that our proposal requires sending 18.72 IBGP UPDATES messages/s and 3.2 IBGP WITHDRAWAL messages/s. Summarizing, our proposal introduces 21.92 signaling messages/s where each ER must process 5.48 messages/s.

Regarding the transit traffic table 1 presents the results. In our proposal, fHAs have to process 465.04 Mbps. Our simulated network has two fHAs and each one processes 232.52 Mbps of data traffic. In [4-8] HAs process 1412.9 Mbps of traffic, our simulated network has 5 sets of HAs, this means that each set of HA processes 282.58 Mbps. In our proposal, the data traffic destined towards CNs is directly processed by ERs (947.86 Mbps). Regarding the transit traffic, our proposal reduces it by 75% compared to existing solutions [4-8]. It should be taken into account that existing solutions [4-8] must send each data packet twice, one from the ER to the HA and another one in the opposite direction.

Table 1. Simulation Results (Values in Mbps)

	Existing Solutions [4-8]	Our Proposal
Traffic sent by MNs through the HA/fHA	1412.9 (465.04 to the HN, 947.86 to CNs)	
Traffic processed by HAs/fHAs	1412.9	465.04
Traffic processed by ERs	N/A	947.86
Transit Traffic	1895.72	473.93

Our solution has to forward transit traffic into the network only if the receiving ER is not the exit point of the packet's destination address

Finally, during the simulation a maximum of 900 nodes were away from home at the same time (average 717, minimum 685). This means that the maximum stored state on each ER is 29.9KB.

This simulation shows that our proposal is viable, and that among other benefits, it can reduce the transit traffic considerably.

5 Conclusion

In this paper, we have presented a flexible and distributed HA architecture. Existing proposals [4-8] provide a reliable HA architecture by deploying redundant HAs on each Home Link. Our proposal has the same benefits but with just one set of fHA for the whole network.

Our solution is reliable: a failure on the MN's ARs will not disconnect the MN. In this case the MN will still be able to communicate with the Home Network (except with the Home Link) and with the rest of the Internet. Since our solution allows deploying several fHA for each network, a failure of a fHA does not disconnect the MN. In this case, our solution can benefit from the proposed efficient failure recovery mechanisms presented in [7][8] because it is fully compatible with them. This way we can minimize the service interruption time. Finally, a failure on a ER does not disconnect the MN. In this case, the network announces the failure of the ER through the exterior routing protocol and the packets will be re-routed. Our solution also provides load balancing because the MN's data packets are processed by ERs or by a set of fHAs. Moreover, it reduces considerably the transit traffic through the network (75% according to our simulation).

Distributing HA's operations requires adding some extra load at the ERs and at the ARs. The ERs have to setup tunnels and configure new routes towards the MNs while the ARs have to configure a tunnel with each fHA and intercept packets destined to its MNs. We believe that routers are hardware machines optimized to perform exactly this type of operations. It is important to remark that signaling and IPSec data packets are processed by fHAs, not by routers. In addition, our simulation of a highly mobile environment shows that each ER would require processing only an average of 5.48 signaling messages per second.

Finally, as future work, we plan to implement Traffic Engineering for MN's traffic in case that the exit routers are the AS's border routers. We also plan to extend our solution to correspondent networks.

References

1. D. Johnson et al: Mobility Support in IPv6. RFC 3775, 2004
2. Jahanzeb Faizan et al: Problem Statement: Home Agent Reliability. IETF Draft (Work in Progress), 2004
3. S. Kent et al: Security Architecture for the Internet Protocol. RFC 2401, 1998
4. F. Heissenhuber, W. Fritsche and A. Riedl: Home Agent Redundancy and Load Balancing in Mobile IPv6. in Proc. 5th Int. Conf. Broadband Communications, 1999
5. H. Deng et al: Load Balance for Distributed Home Agents in Mobile IPv6. in Proc. 14th IEEE PIMRIC 2003
6. H. Deng et al: Load Balance for Distributed Home Agents in Mobile IPv6. IETF Draft (Work in Progress), 2003
7. J. Faizan et al: Efficient Dynamic Load Balancing for Multiple Home Agents in Mobile IPv6 based Networks. in Proc. 5th Int. Conf. Pervasive Services, 2005
8. R. Wakikawa et al: Home Agent Reliability Protocol. IETF Draft (Work in Progress), 2006
9. R. Wakikawa et al: Inter Home Agents Protocol Specifications. IETF Draft (Work in Progress), 2006
10. R. Wakikawa et al: Virtual mobility control domain for enhancements of mobility protocols. INFOCOM 2005
11. Y. Rekhter et al: A Border Gateway Protocol 4 (BGP-4) RFC 1771, 1995
12. H. Soliman et al: Hierarchichal Mobile IPv6 Mobility Management (HMIPv6) RFC 4140, 2005
13. S. Deering et al: Internet Protocol, Version 6 Specification RFC 2460, 1998
14. PalChaudhuri, S. et al: Perfect Simulations for Random Trip Mobility Models 38th Simulation Symposium, 2005