

# HIP Location Privacy Framework

Alfredo Matos, Justino Santos,  
Susana Sargento and Rui Aguiar

Instituto Telecomunicações, Universidade de Aveiro  
Campus Universitario de Santiago  
3810-193 Aveiro, Portugal

{alfredo.matos|jsantos}@av.it.pt  
{ssargento|ruilaa}@det.ua.pt

João Girão and Marco Liebsch

NEC Europe Ltd  
Kurfürsten-Anlage 36  
69115 Heidelberg, Germany

{joao.girao|marco.liebsch}@netlab.nec.de

## ABSTRACT

Privacy and security are key aspects in future network architectures. The Host Identity Protocol (HIP) is a new proposal which decouples identifiers from locators and may eventually replace conventional addressing and network transport. In this document we propose an architecture that provides location privacy, based on HIP. We further validate our work by implementation and support the feasibility of our protocol by experimentation.<sup>1</sup>

## Keywords

Location, privacy, Host Identity Protocol, architecture.

## 1. INTRODUCTION

The Internet is evolving to a place where concepts such as mobility and global reachability are common and seamlessly deployed. Nowadays we can already feel the increasing number of access technologies (e.g. WiFi, WiMax, GPRS), user equipments with multiple network interfaces and new trends in services (e.g. VoIP) which require a mobile environment. However, when the current Internet architecture was designed, mobility and multihoming were not taken into consideration. IP addresses are used simultaneously as locators and identifiers for a host mainly because at design time nodes were assumed to be static and trusted. Solutions such as the Host Identity Protocol (HIP) [1, 2] and FARA [3] are today's attempts of addressing these issues.

Location privacy has a large role in new mobile networks: users do not wish to be tracked either by the network provider

<sup>1</sup>The work described in this paper is based on results of IST FP6 Integrated Project DAIDALOS. DAIDALOS receives research funding from the European Community's Sixth Framework Programme. Apart from this, the European Commission has no responsibility for the content of this paper. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiArch '06 San Francisco, California, USA  
Copyright 2006 ACM 1-59593-566-5/06/0012 ...\$5.00.

or by third parties. In [4] location privacy is defined as the capability of preventing other parties from learning one's previous or current location. Location mainly pertains to the topological position of a node, although the topological location can give an accurate geographical position.

For a node to obtain location privacy, there can be no relationship between its identifiers and locators. Location privacy requirements [5] clearly state that the problem is not limited to a single layer. In fact, it concerns all identifiers associated with a node, including MAC and IP addresses. Thus, another important problem is the identifier interdependency where, for instance, a mobile device moving through foreign networks always carries the same unique MAC address. However, location privacy issues below the network layer are considered out of scope for this paper. At the network level, using an IP address for identity and location, as it is currently done, the task of hiding this relationship becomes very hard. In an IPv6 context, a host performing address auto-configuration is implicitly disclosing its MAC address, allowing a direct mapping between MAC address and IPv6 address. Furthermore, the use of Mobile IPv6 [6] discloses the location in the care of address and associates it with the node identifier.

In our work, we address location privacy issues designing a framework that uses HIP as the base protocol. It takes advantage of the identifier/locator separation provided by HIP, that currently does not address location privacy. Our contribution is an architectural solution that provides location privacy, without requiring modifications on the core network, while supporting mobility. The remainder of this paper is structured as follows: In Section 2 we present a survey of related work after which, in Section 3, we introduce HIP and related location privacy issues. In Section 4 we describe the proposed framework. Section 5 describes the protocol operations for registration, packet delivery and mobility of endpoints. In Section 6 we present the deployment of the location privacy framework in an IPv6 scenario, along with a prototype for privacy and performance assessment. We conclude our work in Section 7, where we summarize the advantages of this framework.

## 2. RELATED WORK

The current work in this area focuses on new network architectures or mechanisms that are able to cope with location privacy. IP<sup>2</sup> [7], Turfnet [8], I3 [9] and Blind [10], while not focusing on location privacy, address some of the issues. Furthermore, the technique known as 'onion routing' [11] hides location information between layers of transport.

IP<sup>2</sup> [7] is able to hide the user location through the use

of anchor points in the network which also deal with mobility. This resembles what happens in HMIPv6 [12] and our proposed framework, but has a big deployment overhead, providing mostly features that are out of scope.

Overlay networks provide also good approaches to hide location information. In Turfnet [8], location privacy is achieved implicitly mainly due to an innovative method of routing and the use of Turfnet Gateways connecting each Turf. However, it is difficult to achieve optimal routing. In I3 [9], a new realm for routing is defined based on names. Using a rendezvous point between communicating partners, it is possible to achieve some degree of location privacy, but it is still an overlay and we aim to achieve location privacy through architectural support.

Blind [10] describes a complete identity protection framework for endpoints. It proposes a Diffie-Hellman authenticated agreement for identity exchange. Regarding location privacy, a solution based on identity aware NATs is proposed: when an endpoint tries to initiate communication with a node in the network, it uses a Forwarding Agent that selects a virtual IP address for it. The peers are able to see only the virtual address, not the real address of the endpoint. Although sharing similarities with our approach, it does not contemplate security between endpoint and Forwarding Agent, or has support for mobility.

Onion routing [11] is particularly interesting: it prevents the transport medium from knowing who is communicating with whom by using multiple onion routers. Each router is responsible for removing one layer of encryption which protects the destination and source in a direct proportion to the number of onion routers. The main drawbacks of this approach are the overhead introduced by the multiple encryptions, non optimal routing, lack of protection of final hops and the endpoints learn the real locators.

The presented approaches lack in either simplicity, performance or in the level of location privacy achieved. Our privacy framework based on HIP aims at an improvement on these points.

### 3. HOST IDENTITY PROTOCOL

HIP introduces a new cryptographic namespace for identification, but most importantly the dual role of IP addresses providing a solution for decoupling the location from the identity. A cryptographic public key, of an asymmetric key pair, the HIP Host Identity (HI), is used and acts as the host's unique identifier. The host's private key serves as an assurance that he owns the identity claimed by the public key. Because the HI is not practical to use on the wire due to its length, it can be represented by a 128-bit Host Identity Tag (HIT). Since the HIT is 128-bits long, the same length as an IPv6 address, it can be used seamlessly by IPv6 applications. When HIP is used, the upper layers, including the applications, are not aware of the IP Address used for routing. In addition to the separation between locator and identifier, HIP defines primitives to negotiate security associations between capable nodes. The core of the HIP protocol [1] is the base exchange (BE) depicted in Fig. 1, which is an authenticated Diffie-Hellman four-way handshake. The BE provides means for two nodes to prove their identity to each other, while optionally establishing cryptographic material used for creating secure communication channels with IPsec Security Associations (SA) [13], bounded to the Host Identities. However, the packets travelling in the network do not contain the actual HI information, but the inbound

packets are identified and mapped to the correct SA using the Security Parameter Index (SPI) value in the IPsec header. Mobility and multihoming are supported by the underlying locator agility provided by HIP, which enables changing the point of attachment without breaking the communication, as shown in [14]. To aid mobility, a rendezvous mechanism was also designed in [15] allowing mobile hosts to be reached without the use of dynamic DNS updates.




Figure 1: HIP Base Exchange

### 3.1 HIP Location Privacy Issues

The current HIP architecture does not take into account location privacy issues, since it requires a node to send its locator to every peer. In the base exchange this occurs when including the locator parameter in R1 and I2 messages. As far as mobility is concerned, when a handover occurs, explicit update messages with the locator parameter must be sent. This procedure is comparable to the Binding Update messages exchanged between MIPv6 [6] enabled Mobile Nodes (MN) and Correspondent Nodes (CN) when performing route optimization. In fact, HIP ultimately suffers from the same location privacy issues as MIPv6 described in [5]. Still regarding mobility, if we consider the presence of a rendezvous server (RVS), the Initiator does not immediately reveal the current locator of the Responder. However, that information is disclosed in the R1 packet. One can learn the current location of a host by simply inspecting the base exchange and update messages, enabling an eavesdropper to learn the involved HITs and IPv6 addresses of both participants, forfeiting the location privacy of the peers.

In the previous approaches an end-to-end addressing mechanism is used. This means that both Initiator and Responder will always learn each other's current IP address once the BE is completed, since the resolution from Identifier to Locator is performed at the end hosts. If a host present in a DNS query is not registered in an RVS, the DNS resolves to the current IPv6 address of the node. In an architecture which supports location privacy, hosts should never be able to map the identifier to the real locator of the node.

In [16] some considerations and network elements are introduced to shield a HIP node's location. Our proposal is to use the current HIP architecture and introduce new functional units and enhanced protocol operations which solve the above mentioned problems, providing location privacy to the nodes in the network [17].

### 4. LOCATION PRIVACY ARCHITECTURE

As suggested in [16], location privacy is provided by delegating the HIT to IP resolution to a network entity called the Rendezvous Agent (RVA). Moving the resolution upwards in the network topology, from the HIP Mobile Node (HMN) to the RVA has the added benefit that locators do not need to be disclosed in the Access Network. The core feature of our

solution is the concept of RVA protected areas, which are access networks where global locators are either concealed or not used at all. Instead, HITs or local addresses are used to identify the traffic path. The RVA is also responsible for local mobility, i.e. under its protected area.

Rather than defining a specific transport layer for our approach, we define a set of basic requirements which must be met for the protocol to function properly. The only assumption made is that the core network is IP based. We do not specify any particular technology under RVA protected areas. In fact, we consider possible instantiations based on direct IPv6 address translations, tunnels or semantical adaptations (replacing IPv6 addresses with HITs). In Section 6 an IPv6 based solution is described.

An example of the proposed topology is illustrated in Fig. 2. The scenario consists of two RVA protected areas connected to the Internet. An RVA protected area is composed by multiple ARs which are directly connected to an RVA. There are no assumptions on the number of RVA protected areas and its extension, although it is reasonable to think that an RVA covers a large number of ARs. A wider coverage area, geographical or topological, limits the amount of location information revealed to an external eavesdropper. The RVS and DNS servers are located in the core. The AR and the RVA are functional entities, thus they can also be collocated in the same machine, but at the expense of some limitation on location privacy (Section 3.1).




Figure 2: Basic architecture topology example

Depending on the chosen solution for routing in the access networks, already existing HIP elements may require modifications, since hosts and routers depend on the protocol used in the access network. If some form of identity based routing is used, then the amount of information to be kept at each node (eg. AR keeps HIT based neighbor tables) is larger. If the access network remains IPv6 based, then no modification is required other than enhancing the neighbour advertisement protocol.

Another element of the HIP architecture which requires minor modifications is the RVS. The RVS should be capable of performing a double resolution: translating a received HIT of a host into the HIT of its designated RVA and finally translating to the corresponding RVA address.

#### 4.1 Rendezvous Agent

The RVA is an enhanced RVS which performs the IP-HIT

or IP-IP address translations. This functionality split provides location privacy to the HMNs behind it. The mechanism consists of re-addressing packets flowing from and to the core network. To forward packets to a host outside the RVA protected area, the RVA addresses a globally routable IPv6 address previously assigned by another RVA to the destination host. When an RVA receives packets from the outside network to a host belonging to its RVA protected area, it readdresses them to HITs, or local addresses, and forwards the packet to the destination. Note that the RVA is the entity which assigns globally routable IP addresses to the hosts under it, and the only one capable of mapping HIT, or local address, to global addresses. The RVA is capable of forwarding packets based on HITs through maintaining a mapping for every HMN in the protected area to its point of attachment, which is the AR. The RVA is responsible for handling mobility for the nodes in the protected area, raising the possibility that the RVA might have to signal other RVAs or HIP nodes, on behalf of the HMNs, for location updates.

#### 4.2 Privacy Goals

Our proposed architecture presents several improvements over the base protocol: The globally assigned IPv6 addresses limits the amount of location information an eavesdropper in the core network obtains from mapping HITs, or local addresses, to global addresses used in the routing process. If an eavesdropper is on path and able to intercept all messages received by the responder outside the responder's protected area, it does not see local mobility and can only track movement between different RVA protected areas. The size of RVA protected areas determines how much geographical location information an attacker can obtain by using this method.

An attacker tracking the base exchange can learn the SPIs of the IPsec SAs and later on map the SPIs to the assigned IPv6 addresses. Once again, the attacker is limited to information pertaining to the area of the RVA, and not specific locations.

An attacker is only able to learn a HIP node's location if it is in the same access network. In this case, the attacker can track HITs or local IPs, MACs and possibly other access transport information by simply eavesdropping on the physical medium. We believe that this architecture can be extended or combined with other mechanisms to also cover these scenarios.

### 5. PROTOCOL OPERATION

Our location privacy architecture requires extensions to the basic HIP mechanisms. This includes changes to the BE with RVA and RVS, in the mobility signaling, network address translation and signaling between RVAs. Since the simpler case is using IPv6 in the access networks, we generalize by addressing HIT based routing in the next sections.

#### 5.1 Base Exchange with RVA

When a HMN first arrives in a protected area, it has to register with the responsible RVA. The  $HIT_{rva}$  is retrieved from the Advertisement messages sent by the AR. The registration takes the shape of a base exchange with the RVA (Fig. 3) using registration extension [18]. Note that no packet forwarding for the HMN is done until the BE is completed in order to avoid DoS attacks. Once this phase is over, the RVA assigns a global IPv6 address (IPg) that is used for

the registering node in the core network. The IPg should be generated from a pool of addresses assigned to the RVA. In case we are using identity based routing, during this phase the RVA learns the HIT-AR mapping necessary for packet forwarding. Once the BE with the RVA described above is completed, the HMN has to register with its RVS or update it.




Figure 3: Base exchange with Rendezvous Agent

## 5.2 Base Exchange with RVS

The main difference from the normal HIP procedure is the inclusion of a RVA parameter in the I2 packet. It carries the newly discovered RVA identifier informing the RVS of the current RVA in use. The registration is not performed with a locator but with the identifier of the RVA.

## 5.3 Base Exchange between end-points

The HIP base exchange between an initiator and a responder remains unchanged; the key differences are at the network layer. The initiator’s RVA has to perform re-addressing of outgoing packets to globally routable IP addresses (Fig. 4). If the endpoints do not perform resolution, and resort to identity based routing, then the RVA does not know the  $HIT_{responder}$  and has to query the DNS for responder’s address, posing the drawback of performing reverse lookup on HITs, to obtain the IP address of the responder’s RVS. The RVS then relays I1 packet to the IP address of the responder’s RVA. This is done based on the two step mapping previously discussed where the  $HIT_{responder}$  is translated to the  $HIT_{rva}$ , and then  $HIT_{rva}$  is finally translated to the RVA address. Also, the FROM and VIA Parameters are included as described in [15]. Upon receiving the I1, the Responder’s RVA forwards the packet to the destination HIT. The RVA protecting the responder should also store the newly learnt  $HIT_{initiator} - IPg_{initiator}$  mapping for further packet forwarding with HIT routing.

Afterwards, the HMN receives the packet and replies with an R1 packet that is relayed by the RVA, re-addressing the packet based on the learnt mapping. In the remaining BE packets (I2 and R2), the globally assigned IP addresses of both I and R are used between RVAs.

The usage of unpublished HIs by one of the endpoints invalidates HIT based routing in the access networks. Even if another routing solution is used, such as normal IPv6, the usage of unpublished HI means that a local address is leaked outside of the protected area, forfeiting the responder’s location privacy. One possible solution is to use an unpublished HI that is only registered with the RVA, allowing it to be reachable at an assigned address.

## 5.4 Mobility

In our architecture we can define two types of handover: Intra and Inter RVA. Both these procedures are triggered

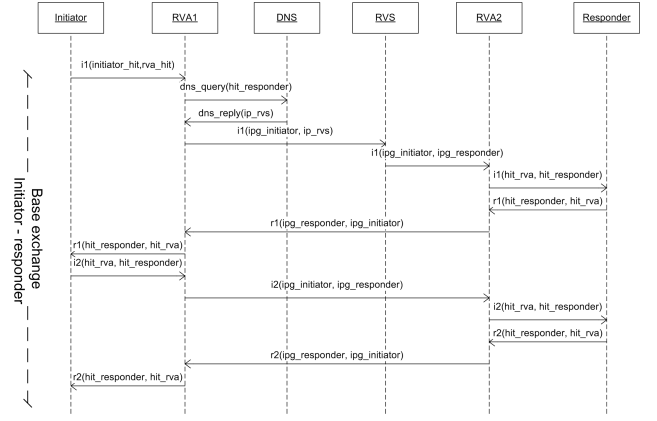


Figure 4: Base Exchange between Initiator and Responder

by the advertisement system when a new access router or a new RVA identifier is detected. Moving within a protected area consists of an intra RVA handover, and requires only updating the RVA binding. The handover is transparent to all communicating peers. If the HIP node changes protected area, then an inter RVA handover occurs. In this scenario the host has to register with the new RVA (performing a Base Exchange), and updating its RVS entry with the new responsible identity (Fig. 5).

The new RVA should also inform the old RVA of the handover, so that packets are forwarded correctly and the connection is not severed. When the new RVA receives a forwarded packet from another RVA, it updates the location to the forwarding RVA. Forwarded packets need to be differentiated from the normal traffic, allowing a RVA to decide whether mobility updates are needed or not.

The peers of the on-going connection should be gradually updated, as the new RVA notices that packets are being forwarded by the old RVA.

As stated before, our framework also requires RVA-to-RVA communication for location updates. The RVA-to-RVA communication should be preceded by a HIP base exchange, allowing secure communication and, more importantly, authentication. But depending on the scenario, the trust relation between the RVAs may be different. For instance, in a network operator scenario, all RVAs may be certified by a common Certification Authority, allowing only trusted RVAs to signal each other. A more flexible solution resides on the HMN providing a certificate to the RVA during the registration process, thus enabling them to prove to each other that they are acting on behalf of the HMN.

## 6. AN EXAMPLE OF IPV6 INSTANCIATION

The framework definition, as described in [17], does not make assumptions on packet transport mechanisms within the RVA protected area, although IPv6 is assumed in the core network. However, the most feasible approach is to also use IPv6 in the protected areas. This means that communication inside the protected area is done with IPv6 addresses with a global format, but with local scope, that are translated by the RVA. The main advantage of this solution is that it requires no changes to routing mechanisms within the access network. With the IPv6 access network, the deployment of the RVA advertisement system consists

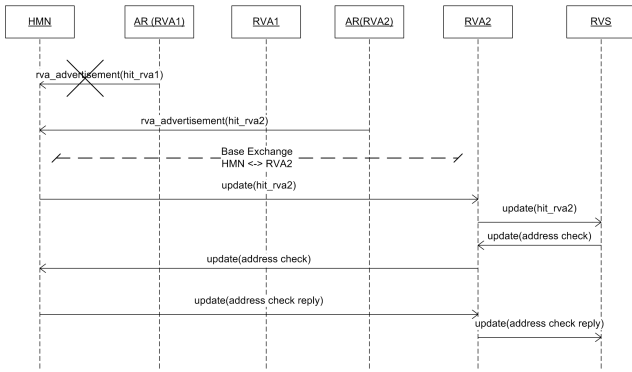


Figure 5: Inter RVA Handover

on enhancing the Router Advertisement [19] messages to carry HIP parameters as options. Just like a HIP parameter, a neighbor discovery option has a type/length/value (TLV) format. The new HIP parameter (RVA\_INFO) is an option that advertises the  $HIT_{rva}$ , along with the advertisement lifetime. Through the advertisement mechanism, the HMN can detect different protected areas. Address configuration and mobility detection should be done according to [19], with the extension of RVA detection by  $HIT_{rva}$  announcements. The registration with the RVA is performed according to Section 5.1, followed by the assignment of a global IPv6 address to the HMN. Later, the HMN registers the acquired  $HIT_{rva}$  with the RVS, for correct I1 packet forwarding. When movement is detected, the HMN updates the binding with the RVA, in case of intra RVA mobility, or registers with the new RVA, in case of inter RVA mobility, updating the RVS afterwards.

### 6.1 Prototype Implementation

In order to validate our approach we have chosen to implement solely the global IPv6 assignment and packet translation, in order to perform our tests. Registration procedures are similar to those performed when registering with the RVS, and therefore are of secondary importance to a prototype. The implementation was developed under Linux, kernel 2.6.15, using as basis the HIPL implementation provided by the InfraHIP project [20]. It consists of a manually triggered registration process, with an IPv6 global address being assigned upon completion for the registering local address. After this, the RVA performs the necessary translations. The work is performed by a module which stores the registered addresses, under a hash table, resorting to Linux IPTables for packet capture.

In Fig. 6 we depict the scenario used for the performance evaluation. This scenario consists of two access networks, both served by the same RVA. This RVA is capable of handling multiple protected areas. After each node registers, the RVA generates addresses from a pool of available prefixes. For simplicity, we use only one prefix for both areas.

The testing procedure is composed of two phases: first we show the results of a base exchange between the two endpoints and the assigned addresses in the several areas and nodes. We then perform measurements on the responsiveness of the system, using Round Trip Time values of ICMPv6 echo request and response packets, and on how our protocol impacts the overall system performance, by measuring the throughput of TCP when our scheme is in place.




Figure 6: Implementation testing scenario.

This prototype aims at showing the flow of addresses through networks between communicating peers, and also how the translation overhead impacts the implementation, allowing us to generalize to protocol operation delays.

### 6.2 Location Leakage Analysis

Since both endpoints are assigned and use global IPv6 addresses for communicating with each other, they cannot determine the actual address, and consequently location, within the RVA protected area. In Table 1 we represent the information gathered at each point in our reference scenario. As an example, we can see that Node A is sending to the global address of B, 2001::6ada:1e65:93f3:ff00, but is unaware of his local address, 3ffa::1, where the packets actually get delivered.

Networks	Node A	Node B
Area 1	3ffe::1	2001::6ada:1e65:93f3:ff00
Core	2001::ded8:ce89:6390:eb00	2001::6ada:1e65:93f3:ff00
Area 2	2001::ded8:ce89:6390:eb00	3ffa::1

Table 1: Summary of the seen addresses on each network

With our approach, we are not concerned with protecting the identity of the nodes themselves. An attacker can, in any point of the network, identify both the HMN and its peers, but not their locations. Furthermore, the RVA is a point of information gathering for the network and, if compromised, reveal the identity and location of registered nodes.

Another mechanism that provides information on the location of the node is hop count. Even if the node is behind an RVA, the hop count, together with the topology of the underlying network, can reveal information on the whereabouts of the node. One mechanism to thwart such attempts is to keep the hop count value between the RVA and the node to 1. This can be achieved by tunneling.

### 6.3 Performance Evaluation

From the prototype we analyze two different performance metrics: the first is the Round Trip Time (RTT) of ICMPv6 Echo Request and Response packets, to evaluate the real impact of packet translation. In fact, we measure the implementation delay to infer the real protocol impact, given the fact that this is not an optimized implementation. We use the RTT to obtain a time synchronization independent value to measure the delay introduced by all translations. The other measurement is the TCP throughput where we note the impact on bandwidth caused by the translation delay. In both cases, we compare the performance with RVA intervention to normal HIP operation.

For each run, regarding the RTT, we perform 100 measurements, and present the averages for each run. Fig. 7 shows the average values, matched against a plain HIP scenario, under the same conditions. As we can see the RVA introduces a slight delay in packet delivery. Using our protocol, the average value for the RTT is of  $1.130 \pm 0.003ms$ , whereas, without RVA, it is of  $1.067 \pm 0.002ms$ . However, the difference is sufficiently small that we can discard the impact on network traffic.

We measure the maximum available bandwidth for a TCP connection in two scenarios with and without RVA processing. Each run consists of starting a TCP traffic generator for 30 seconds, and obtaining the bandwidth of that flow. Table 2 shows the results of the tests, presenting the average of 10 runs. As we can see, the translation delay has very little impact on the TCP throughput. The total bandwidth used with normal HIP is of  $6.43 \pm 0.03Mbit$ , and for HIP with RVA translations is of  $6.44 \pm 0.07Mbit$ . The similarity in these values allow us to neglect the impact on network throughput performance.




Figure 7: Round Trip Time Impact

Average Bandwidth (Mbits/s)	
Without RVA	With RVA
6.43	6.44

Table 2: TCP Bandwidth usage from the performed tests.

## 7. CONCLUSION

We described a framework able to conceal the IP addresses of communicating HIP nodes, using an architectural approach, based on protected areas that provide location privacy. The transport in these areas could be HIT based and therefore no locators are necessary. In case the transport in the access network requires locators for routing, the scope of these names are deemed as local and are never revealed outside the AN. From the prototype implementation, we illustrated the feasibility of the architecture, while still taking advantage of IPv6 routing. We were also able to confirm the performance of the protocol. The results obtained showed that, with the proposed architecture, we are able to hide location information with only minimal impact on network performance.

## 8. REFERENCES

- [1] R. Moskowitz, "Host Identity Protocol." Internet Draft (Work in Progress), July 2006.
- [2] R. Moskowitz, "Host Identity Protocol Architecture." RFC 4423 (Proposed Standard), may 2006.
- [3] D. Clark *et al.*, "Fara: reorganizing addressing architecture," 2003.
- [4] W. Haddad, "Privacy for Mobile and Multi-homed Nodes: MoMiPriv Problem Statement." Internet Draft (Work in Progress), February 2005.
- [5] W. Haddad, "Privacy for Mobile and Multi-homed Nodes: Formalizing the Threat Model." Internet Draft (Work in Progress), February 2005.
- [6] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6." RFC 3775 (Proposed Standard), June 2004.
- [7] T. Okagawa *et al.*, "Ip packet routing mechanism based on mobility management in a ip based network," *8th International Conference on Intelligence in next generation networks*, 2003.
- [8] S. Schmidt *et al.*, "Turfnet: An Architecture for dynamically composable networks," *Proceedings in 1st IFIP TC6 WG6.6 Workshop on Autonomic Communication (WAC 2004)*, 2004.
- [9] I. Stoica *et al.*, "Internet indirection infrastructure," *Proceedings in ACM SIGCOMM Conference (SIGCOMM'02)*, pp. 73–88, August 2002.
- [10] J. Ylitalo and P. Nikander, "Blind: A complete identity protection framework for end-points," *Security Protocols, Twelfth International Workshop*, 2004.
- [11] R. Dingledine *et al.*, "TOR: The second-generation onion router," *Proceedings of 13th USENIX Security Symposium*, 2004.
- [12] H. Soliman *et al.*, "Hierarchical Mobile IPv6 Mobility Management (HMIPv6)." RFC 4140 (Proposed Standard), June 2004.
- [13] P. Jokela, R. Moskowitz, and P. Nikander, "Using ESP transport format with HIP." Internet Draft (Work in Progress), June 2006.
- [14] P. Nikander *et al.*, "End-Host Mobility and Multi-Homing with Host Identity Protocol." Internet Draft (Work in Progress), June 2006.
- [15] J. Laganier and L. Eggert, "Host Identity Protocol (HIP) Rendezvous Extensions." Internet Draft (Work in Progress), June 2006.
- [16] M. Liebsch and L. Eggert, "Host Identity Protocol (HIP) Rendezvous Mechanisms." Internet Draft (Work in Progress), July 2004.
- [17] A. Matos, J. Santos, J. Girao, M. Liebsch, and R. Aguiar, "Host Identity Protocol Location Privacy Extensions." Internet Draft (Work in Progress), March 2006.
- [18] T. Koponen and L. Eggert, "Host Identity Protocol (HIP) Registration Extension." Internet Draft (Work in Progress), June 2006.
- [19] T. Narten *et al.*, "Neighbor Discovery for IP Version 6 (IPv6)." RFC 2461, December 1998.
- [20] "Infrastructure for hip, <http://infrahip.hiit.fi/>, Helsinki Institute for Information Technology."