

ENABLING MOBILE NETWORKS THROUGH SECURE NAMING

Devan Rehunathan
University of St Andrews
St Andrews, UK

Randall Atkinson
Extreme Networks
RTP, NC, USA

Saleem Bhatti
University of St Andrews
St Andrews, UK

ABSTRACT

Mobile Networks are increasingly important in land-, sea- and air-based military scenarios. The interest in supporting network mobility for Internet Protocol (IP) networks has led to the Network Mobility (NEMO) protocol extensions being proposed for IP within the IETF. These extensions are based on the work already completed on host mobility for Mobile IP (MIP). The current work is based on the use of software agents: a Home Agent (HA) intercepts packets destined for the addresses in the mobile network and uses an IP-in-IP tunnel to send the packets to the Mobile Router (MR) located at a Care of Address (CoA), which terminates the tunnel. As the mobile network moves to new IP networks, the MR updates the HA with its new CoA. While this tunnelling approach represents a sound engineering solution for backwards compatibility, and is the only one that has been pursued within the IETF, it has seen little deployment, either in support of mobile hosts or mobile networks. We make the case for an alternative approach based on secure naming. We make a comparison in operation with the current tunnelling-based approach, both in architecture and by analysis of protocol operation. Our initial analyses indicate that a naming-based approach shows promise as a viable alternative to a tunnelling-based approach, and could offer other architectural benefits.

I. INTRODUCTION

Today, military networks need maximum flexibility and advanced capabilities in order to deliver different mission solutions. We have outlined previously a proposal for provision of a harmonised set of capabilities for site multi-homing, traffic engineering, end-to-end security, and support for mobile systems and networks [1]. In this paper, we expand on our new approach to mobile networks with the Identifier Locator Network Protocol (ILNP)¹. We present an analytical comparison with the IETF work on Network Mobility (NEMO). We note that, at the time of writing, the work on NEMO is being integrated with Mobile IPv6 under the Mobility Extensions WG², but it seems likely that there will be no major architectural changes.

¹<http://ilnp.cs.st-andrews.ac.uk>

²<http://www.ietf.org/html.charters/mext-charter.html>

We take the position that, with the increased capability in naming functionality within the Domain Name System (DNS), including secure, dynamic update [2], it is now worthwhile considering the use of naming as a key capability to enable mission-critical services within a military context. We support our position by showing the enabling of a secure mobile network capability (itself a technical challenge) through naming.

In our discussion, we chose the abstraction in Figure 1 for the mobile network because it maps to many real scenarios, e.g. a warship with multiple satellite uplinks; or an infantry platoon (mobile network) with multiple radio links. We show only two external links, for simplicity, but a larger number of external links can also be supported. (The ‘coordination protocol’ is not considered in this paper, but commercial systems exist today that provide such functionality and could be adapted for use.)

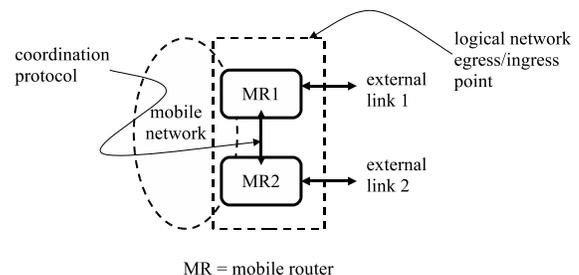


Fig. 1. General scenario: an example mobile network, with two mobile routers (MRs), each providing access to an independent external link.

In Section II we describe the need for secure mobile network protocols in military communications, with an emphasis on the current problems with IP networks today. In Section III, we describe the approach taken by the IETF to support mobile networks, using NEMO. In Section IV we give a broad overview of our proposed solution, ILNPv6. In Section V, we provide a qualitative analysis of both the NEMO and ILNPv6 approach by examining the operation of a mobile network based on each approach. We conclude in Section VI, with a summary of our position and analysis.

II. RATIONALE AND MOTIVATION

For mobile Command and Control (C2), for more complex Command, Control, Communications, Computers, Intelli-

gence, Surveillance, and Reconnaissance (C4ISR), and for the provision of routine mobile communications, support is needed today for *mobile networks*. As depicted in abstract in Figure 1, the model is that of an entire network ‘site’ that is mobile. In theatre, examples include an aircraft carrier, or an infantry platoon using a mobile satellite terminal. For the Command Post Of the Future (COPF) as part of a well-provisioned Tactical Operations Centre (TOC), for robustness and coverage, it would be advantageous to have vertical handover capability across multiple radio uplinks (e.g. SATCOM, MILSTAR, or newer EHF systems). Currently, the military must rely mainly on individual coverage from the various systems, and so mobile networking may be limited to single radio systems: integration across different radio systems is possible, but may not be straightforward or seamless.

Of course, interoperability across multiple radio systems at the network layer should be possible by use of the Internet Protocol (IP). However, standards for mobile networking above the radio layer, i.e. at the IP layer, are limited, with a tunnelling-based approach, NEMO, being pursued within the IETF. We take the position that, whilst NEMO is likely to provide appropriate functionality, an alternative approach, based on naming, is viable, maybe desirable and should be investigated.

However, first we must consider the relevant problems in IP (both IPv4 and IPv6) today.

A. Naming Problems in IP today

In our discussion below, we use the term *name* in a very general sense, to refer to any label that is attached to a network object. A summary is given in Table I.

TABLE I
TERMINOLOGY USED IN THIS PAPER

Term	DNS Record	Definition
Address	AAAA, A	Name used both for locating and identifying a network entity
Locator	L	Name that locates, topologically, a subnetwork
Identifier	I	Name that identifies a node, within the scope of a given locator
Network Name	LP	Name that names a network, and points to the Locator value(s) for that network

Today, the value of the IP address is currently used for two quite different functions – as a *locator* for naming a specific interface on a node (or a set of node interfaces on a common subnetwork), and as an *identifier* for naming the node itself. The overloaded semantics of the IP address causes

TABLE II
USE OF NAMES IN ILNP AND IP

Protocol layer	ILNP	IP
Application	FQDN	FQDN, IP address
Transport	Identifier, <i>I</i>	IP address
Network	Locator, <i>L</i>	IP address
Link	MAC address	MAC address

entanglement across these functions and across protocol layers. Currently, the IP address values are used within applications, within the transport protocols (e.g. within the TCP pseudo-header checksum), and also within the network layer to route packets between the end nodes – see Table II. The impacts of this usage include:

Localised Addressing: When a site uses Network Address Translation (NAT) to enable private addressing, harmonised use of multi-homing, mobility, traffic engineering, end-to-end security becomes even more difficult, as the NAT introduces a discontinuity in the end-to-end state. These issues result from the use of the IP address (which the NAT modifies) as an identifier both in transport protocols (e.g. TCP/UDP pseudo-header checksum) and also by some applications (e.g. File Transfer Protocol).

Support for Mobility: Both Mobile IP (v4 and v6), and IP Network Mobility (NEMO), require extra IP addresses, known as *Care of Addresses* (CoAs), to be used with a special-purpose overlay-router, known as the Home Agent (HA). The HA uses an IP-in-IP tunnel to forward packets sent by a correspondent from the mobile node’s Home Address. Effectively, with respect to the IP address space, the mobile node or mobile network is located within a fixed topology.

End-to-End Security: IPsec Security Associations, which of course are also used by HAIPE products, include both the source and destination IP addresses. This means that if a node moves, or a network moves, then the existing IPsec Security Associations will cease to be valid. This constraint exacerbates existing concerns about the scalability of key management for IPsec devices. It also means that, regardless of what changes might be proposed for the Internet Key Exchange (IKEv2), support for mobility and multi-homing will remain limited and hard to deploy in the tactical environments where these capabilities are so crucial.

These are all major factors affecting the use of Mobile IP and NEMO today, but for which we have proposed harmonised solutions previously [1].

III. OVERVIEW OF NEMO

NEMO [3] is an extension to Mobile IPv6 that provides continued connectivity for nodes within mobile networks. Currently, work in NEMO is being progressed in the Mobility Extensions for IPv6 (MEXT) working group of the IETF. The main purpose of this group is to create a more complete mobility solution for IPv6.

A. The NEMO approach

The NEMO approach enables network mobility by creating an additional IP address for the MR, the *Care of Address (CoA)*. The CoA can be seen as a temporary address used by the MR as it moves, and used to allow routing to the current location of MR within an IP network. That is the CoA acts as a locator. Meanwhile, the MR maintains another IP address that is available via DNS, its *Home Address (HoA)*, at its ‘home network’ (the IP sub-network to which the HoA belongs), and is used for maintaining session state with correspondent nodes (CNs). That is, the HoA acts as an identity. While the MR is not at its home network, a HA, acts as a proxy for the MR, forwarding packets received at the home network to the MR’s CoA using a bi-directional IP-in-IP tunnel. All traffic from within the mobile network is sent to the MR, is encapsulated through this tunnel back to the HA where it is de-capsulated and forwarded. To the CNs of the mobile network, it appears as though the MR and nodes within the mobile network – the mobile network nodes (MNNs) – are still at the home network.

This approach allows the MR and its MNN(s) to maintain pseudo-end-to-end connectivity despite changing network attachments. The HA achieves this by keeping HoA-to-CoA bindings up-to-date. This approach does not change the way the IP address is used today, and there is no impact on the IP address structure. However, the CoA/HoA duality reflects the locator/identifier semantic overloading of the IP address, as discussed earlier. There are no additional changes required to the rest of the IP architecture. The location of the mobile network is inconsequential as long as the MR and the HA can set-up and maintain the bi-directional tunnel between them.

B. The NEMO Protocol

When a MR running NEMO migrates to a foreign network, it replies to any Neighbour Advertisements it receives from an Access Route (AR), to receive a new CoA on the visited link. The MR then sends a Binding Update (BU) message to the HA, informing it of its change of CoA. The HA updates its Home Address to CoA cache for that MR and replies

with an Binding Acknowledgement (BA). This process sets up and maintains the bi-directional tunnel between them.

All packets meant for the MR are received by the HA, which then uses IP-in-IP encapsulation to forward the packets to the MR at the CoA. All egress packets from the mobile network (e.g. those sent from a MNN to its CN) follow the same return path through the MR-HA tunnel.

A new visiting mobile node (VMN), joining a NEMO mobile network executes a similar process, where it updates its HA with its new CoA by sending a BU: the HA responds with a Binding Acknowledgement (BA). The VMN maintains its own bi-directional tunnel, except that this tunnel is within the MR-HA tunnel. The mobility of the VMNs is obscured twice (once more than is necessary), first by its own tunnel from itself to its HA and second by the tunnel between the MR and the HA of the MR. This leads to inefficiencies in routing (aka ‘ping-pong’ routing).

IV. OVERVIEW OF ILNP

We present a brief overview of our proposed enhancements to the Internet Architecture, and also specifically to IPv6. We use the term *Identifier-Locator Network Protocol for IPv6 (ILNPv6)* to refer to our proposal, as it can be engineered as enhancements to IPv6 [1], [4].

A. IPv6 Enhancements

The IPv6 packet header and the ILNPv6 packet header are deliberately made similar. Essentially, in ILNPv6, the IPv6 address is broken into two separate components, a Locator (L) and an Identifier (I). Significantly, the IPv6 *Interface Identifier* is replaced by an ILNPv6 *Node Identifier* (I), with slightly different semantics. Our approach recognises that an IP address has two very different roles – as a *locator* and as an *identifier*. So we replace the concept of the *address* with the concepts of an *Identifier* combined with a *Locator*. The *Locator* names an IP (sub)network: this is used only in routing, and not by the upper layers (e.g. TCP or UDP). In practise today, the L value in ILNPv6 packets is exactly the same as the top 64 bits of the IPv6 address, and includes the routing information (see Figure 2). The *Identifier* is only used for node identity (e.g. for TCP or UDP session state).

The idea of an *Identifier/Locator* split is not a new idea, but our particular approach is new and is specified in more detail than preceding proposals [6]–[8]. We believe that applications should use fully-qualified domain names (FQDNs), wherever possible, which is consistent with RFC1958 [9]. A summary of the difference between the

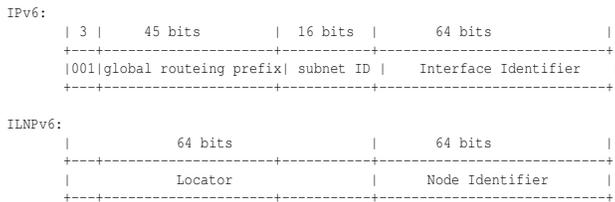


Fig. 2. IPv6 address format (from RFC3587 [5]) as used in ILNPv6

use of names in IP (v4 and v6) and the use in ILNP is given in Table II.

The Locator (L) is an unsigned 64-bit value carried in the upper portion of the IPv6 address and is equivalent to an IPv6 address prefix. The (Node) Identifier (I) is an unsigned 64-bit value carried in the lower portion of the IPv6 address. The I value names a (virtual) node itself, rather than the network interface of a node. An end-system may use multiple I values and multiple L values simultaneously. For the duration of a given ILNP session, its I value should remain constant. For practical reasons, the I value is normally formed from one of the MAC addresses associated with the node. This is represented in the IEEE’s EUI-64 syntax, and is very likely to be globally unique as well. This usage is consistent with the IPv6 Addressing Architecture [10]. Strictly, the I value must be unique only within the scope of the L value with which it is used. However, for practical purposes, having an I value that is likely to be globally unique is very useful, and allows us to dispense with IPv6 Duplicate Address Detection (DAD), which in turn greatly reduces the time required for a node to execute a location change.

Current IPv6 address allocation practices provide sites with IPv6 address blocks that are 48-bits long, which leaves 16 bits for intra-site subnetting. As the ILNPv6 network name (ILNPv6 Locator) is the same as an IPv6 routing prefix, ILNPv6 packets can travel across existing deployed IPv6 backbones. Only the host’s IPv6 stack has to be enhanced to enable ILNPv6 on that host (i.e. to deal with Node Identifier values). ILNPv6 Neighbour Discovery (ND) still uses the full 128-bits of the combined L:I value. So IPv6 ND also can be used without change. *In short, already deployed IPv6 routers will support ILNPv6 without any changes.*

B. DNS Enhancements

ILNP builds upon the security and distributed nature of the DNS resolution service. ILNP introduces 5 new DNS records, which extend the information stored within DNS to include the name, location and name of the current mobile network (if a node is within one). For a more in-depth

analysis of the impact of ILNP on DNS, please refer to [11].

To enable ILNPv6, several new DNS resource records are needed, but no changes are required to the DNS protocol. We add the *I* record, which contains the unsigned 64-bit Identifier associated with a domain name. Similarly, the *L* record contains an unsigned 64-bit Locator associated with a domain name. As a node might have multiple Identifiers and multiple Locators, a given domain name also might have multiple *I* and multiple *L* records. The combination of a given *L* record and an associated *I* record is equivalent to the current IPv6 AAAA record.

Reverse lookups can be done as today with IPv6. As a performance optimisation, we also have a pair of new DNS records that could be used for reverse lookups. The *PTRL* record names an authoritative DNS server for an ILNPv6 subnetwork, while the *PTRI* record is used to obtain the name of a node using a given Identifier on a given subnetwork. This usage enables *PTRL* records to be cached, which is beneficial if performing reverse lookups for multiple nodes on the same subnetwork.

Also, a *Locator Pointer (LP)* DNS record adds an extra level of indirection, by pointing to an *L* record, and offers an engineering optimisation for address management. In operational use, for example, many hosts having the same Locator value, i.e. those on the same IP (sub-)network, would ‘share’ a single *L* record. Each hostname would resolve to a unique *I* record (for that host), and an *LP* record, the latter pointing to (naming) the single *L* record that is shared. So, when all the hosts ‘move’, e.g. the uplink for a mobile network changes, only the single *L* record value needs to change.

Some common arguments against the use of DNS for mobility are that (i) DNS will not be able to cope with the additional load of large scale mobility; (ii) DNS is too slow for location update to propagate within DNS in a timely manner; and (iii) that DNS is insecure. However, we take the position that DNS can not only be used to enable host/network mobility, but that it will be extremely capable of doing so:

- 1) DNS is robust, as lookups and updates are distributed across administratively-delegated, replicated DNS servers [12]. Use of DNS for mobility is as secure as regular DNS, since Secure Dynamic DNS Update [2] is standardised and widely implemented [11].
- 2) Traffic caused by mobility will be relatively small, as DNS today deals with a load where close to 50% of DNS traffic is caused by misconfigurations,

aggressive retransmissions and poor caching [13].

- 3) Current implementations of DNS are suitable for use in mobility solutions that require DNS updates at rates as frequent as once per second [14]. Experimental results from [15] show that BIND implementations of DNS with dynamic update can support mobility solutions.
- 4) Findings from [16] suggest that DNS performance will not be degraded with the widespread use of dynamic, low TTL A-record bindings commonly associated with mobility. Large scale mobility will effect leaf DNS servers, and will have little or no effect on root, top-level-domain (TLD), or even the top-of-the-user-domain DNS servers [11].

So, there are strong indications from previous work to suggest that DNS would be suitable for supporting mobility.

C. ILNP Security Considerations

The High Assurance IP Encryptor (HAIPE) used to protect existing military IP networks is a US DoD profile of IETF standard IP Security [17], so our discussion of IPsec also addresses military deployments of HAIPE systems. [18] In IPsec today, the IPsec Security Associations (SAs) are bound to full IP addresses at the local and remote sites as a form of end-system identity. So, IPsec requires that the IP addresses at each end-point of the communication remain fixed. For localised addressing, multi-homing, and mobility, this may not remain true, and so IPsec has had to be modified, retrospectively, in order to cope with these functions.

With ILNP, however, IPsec SAs are bound only to the Identifier, never to the Locator. This makes it easy for the IPsec Security Association – and the related secure communications channel – to remain operational even if the end-points move.

For DNS, the existing Secure Dynamic DNS Update standard [2] would permit a mobile node or multi-homed node to update its L record(s) when the node moves or its upstream connectivity changes (e.g. due to a link fault). Widely used systems, such as Microsoft Windows or the BIND software used with UNIX, already include support for Secure Dynamic DNS Update. [19]

Separately, the DNS enhancements for ILNPv6 do not change the fundamental operation of the Domain Name System (DNS). Therefore, the the existing DNS Security (DNSSEC) standards [20] can be used unchanged to authenticate these new DNS records. So our proposed enhancements do not create new security risks.

D. Mobile networks with ILNP

In ILNPv6, the mobile network ‘site’ uses private addressing *internally* (to the site network) and the network’s MR(s) rewrite the Locator values of nodes within the site as packets transit that MR. In this model, nodes that are attached to the mobile network segment normally have DNS LP records that point to a common DNS L record covering the entire mobile sub-network. The common L record would be updated by the MR whenever its uplink moves to a different layer-3 ILNPv6 network.

In Figure 1, let us assume that the network is mobile and has two external links with Locators L_1 and L_2 respectively. These will be held in DNS L records pointed to by a DNS LP record. Within the mobile network, localised addressing is used through Locator rewriting in ILNPv6. That is, a local (private) Locator value, L_L , is used by all nodes in the mobile network, and for all egress packets, the MR rewrites L_L to either L_1 or L_2 depending on the egress links to be used, and performs the complimentary operation for ingress packets. This is the ILNPv6 equivalent of NAT, but unlike IP, does not violate end-to-end state and is completely transparent to all ILNPv6 nodes [1].

Now, let us assume a handover is triggered for the link currently using L_1 . A signal is detected in the new cell and a new Locator value, L_3 is attained. This can be done through normal IPv6 discovery mechanisms, as Locator values are identical to IPv6 network prefixes. We will assume that the radio cells providing L_1 and L_3 overlap. Then, the MR updates the DNS L record currently holding value L_1 , to L_3 (for new sessions), and starts changing the state of existing connections using L_1 , to L_3 by issuing *Locator Update (LU)* messages (synonymous to Binding Update message in IPv6) for correspondents using L_1 . It then transitions sessions from L_1 to L_3 using Locator rewriting. When no more packets arrive from remote locations using L_1 within a given time period (i.e. all sessions have transitioned to L_3), the connection is considered to have completed handover. This is a *soft handover* at the ILNPv6 layer, something that is not currently defined for IPv6 or NEMO. Note that the MR is providing this capability efficiently for the whole mobile network. Note also that during this time, the link using Locator L_2 continues to operate as long as the external link 2 is sound, i.e. multi-homing is possible during mobility.

It is also possible to use ILNPv6 for normal handover, simply by switching to L_3 as soon as possible. Any packets in flight addressed to L_1 may be lost, but can be recovered through the retransmission capability in TCP, for example. albeit this would be inefficient, as it will invoke the congestion control behaviour of TCP (due to missing TCP ACKs).

V. ANALYSIS

We now present an analysis of protocol operation by comparing the control packet exchanges necessary for the operation of NEMO and ILNP. We will compare the number of packets exchanged and qualitative nature of the handover delay in the number of round-trip-times (RTTs) that are required to complete the protocol operations. We will consider the action of a visiting mobile node (VMN) entering a mobile network, and then the process of handover as that VMN's communication is continued as the mobile network moves across the cell boundary.

We chose to model the activity of nodes in the network as two phases:

- *Phase 1 - Registration:* This occurs when an VMN arrives at a mobile network and sets up appropriate configuration so it can communicate with CNs outside the mobile network. The aim of Registration is to permit remote CNs to locate the VMN in its current location so new communication sessions can be established.
- *Phase 2 - Handover:* This occurs when the network containing the VMN moves its connectivity. This is when a mobile network moves from one radio cell to another. The aim of handover is to permit existing communication sessions to be maintained.

We present the protocol exchanges as timeline diagrams and discuss the operation of each phase for NEMO and then ILNP. For NEMO, we have consulted the relevant parts of the NEMO and Mobile IPv6 specifications in order to determine the protocol behaviour and previous analysis presented in [21]. For ILNP, we have modelled the protocol interaction that would occur by examining the actual protocol interaction for secure, dynamic DNS update between BIND v9.2.0³ and the Linux *nsupdate* utility⁴.

A. Phase 1: Registration

For both NEMO and ILNP, we assume that a VMN enters the mobile network, discovers a suitable router prefix using normal IPv6 procedures, and subsequently forms an address (for NEMO) or obtains a Locator value (for ILNP).

For NEMO, in Figure 3 we see that when a VMN running MIPv6 joins a mobile network running the NEMO protocol, in order to maintain its connectivity, it sends a BU to its own HA and receives a BA in response. This update of location has a duration of 1 RTT and involves 1 packet generated by

the VMN and 1 by the HA. The update may be protected by the use of IPsec, and we assume that the appropriate IPsec control interaction has already taken place. However, if it had not, then this would impose additional overhead.

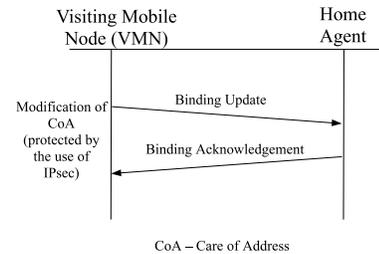


Fig. 3. NEMO Mobile Network Registration for a VMN running MIPv6.

For ILNP, in Figure 4, the VMN updates the DNS entry for its *LP* record by performing a secure, dynamic update. Recall that this does not require any protocol changes to DNS: ILNP would use the normal protocol operations available (delete and add) to update the location of the VMN. This update takes 4 RTTs and results in four packets generated by the VMN and four by the DNS Server.

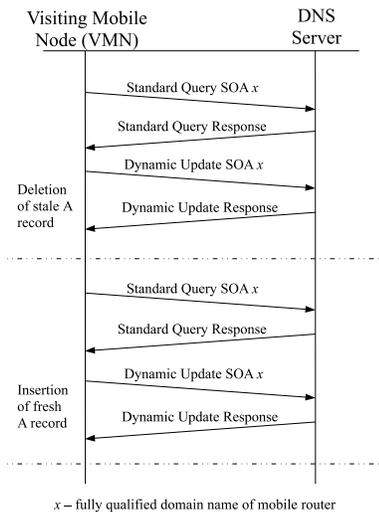


Fig. 4. ILNP Mobile Network Registration for VMN running ILNP.

Overall, the ILNP registration process requires four times as many packets and RTTs, compared to NEMO. The ILNP registration process needs to be re-invoked when the VMN moves to another mobile network. The NEMO handover process for the VMNs does not require an update each time the mobile network moves. For ILNP, once the registration process is complete, and the ILNP VMN has acquired its routing prefix, a local (private) Locator value, it does not need to do anything else until it leaves the ILNP mobile network.

³<https://www.isc.org/software/bind>

⁴<http://linux.yyz.us/nsupdate/>

B. Phase 2: Handover

For handover, we need to consider not only the overhead due to the VMN and its CNs, but also any overhead incurred by the MR in providing ongoing connectivity for the VMN.

For NEMO, we see in Figure 5 the exchange for the MR. When an MR running NEMO changes connectivity, it sends a BU to its own HA and receives a BA. This update of location costs 1 RTT and 2 packets. We note that these packets must traverse the bidirectional IPv6-in-IPv6 tunnel setup between the MR and its HA (we have not considered the tunnelling overhead here).

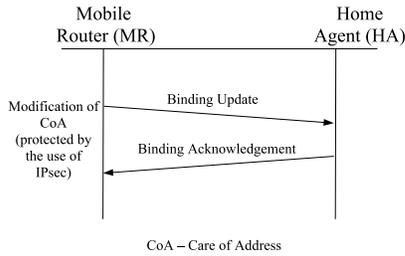


Fig. 5. Network Handover for MR running NEMO.

In Figure 6, we see how the VMN maintains connections already established to a CN. Assuming route optimisation, the VMN starts the Return Routability Procedure with all of its CNs before sending a BU to each CN and receiving a BA. This consists of a testing for reachability of its HA, and for the CoA. This results in the generation of 6 packets, but we will assume this is 2 RTTs, as the ‘Home Test Init’ and ‘Care-of Test Init’ can be completed in parallel.

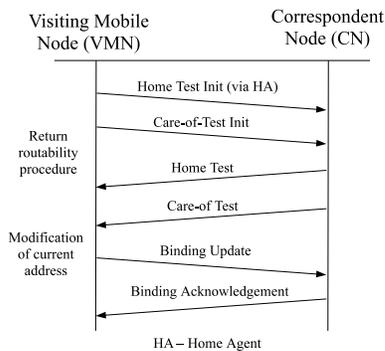


Fig. 6. NEMO Network Handover for VMN running MIPv6 - updates to CN.

For ILNPv6, the MR needs to update its *LP* record in the DNS and then update the Locator values at the CNs. The VMN need not take any action. In Figure 7, we see that 8 packets are generated, requiring 4 RTTs to complete the interaction with the DNS server. Again, this is a normal secure, dynamic DNS protocol exchange. Once the DNS

server has been updated, the MR then updates the Locator values of all relevant CNs, i.e. those using the link that is undergoing handover. In Figure 8, the number of packets generated is dependent on the number of unique CNs maintained by all the VMNs within the mobile network. For each unique CN, the MR sends a LU and waits for an LU Acknowledge from the CN. This results in 2 packets generated in total with a duration of 1 RTT, per CN (which can be sent in parallel, of course).

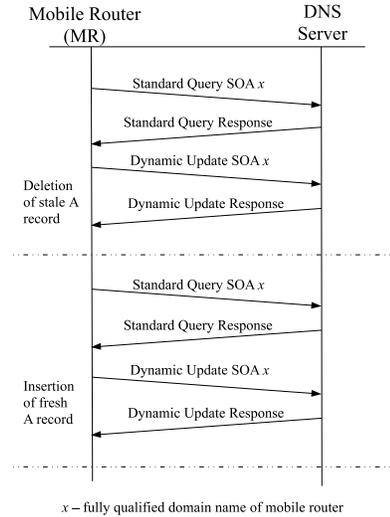


Fig. 7. ILNP Mobile Network Handover for MR running ILNP.

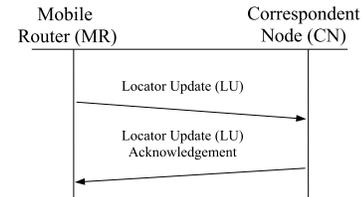


Fig. 8. ILNP Mobile Network Handover for MR running ILNP - updates to CNs.

C. Summary of analysis

From the analysis above, we see that mobile networks with VMN route optimisation has similar protocol overhead, whether ILNP is used or NEMO (with MIP route optimisation, as in Figure 6) is used, specifically:

- ILNP uses eight packets and requires 4 RTTs to allow registration of a VMN, compared to NEMO which requires 2 packets and 1 RTT.
- Handovers (of the VMN and MR) in ILNP are handled completely by the MR. For NEMO, handover and registration is handled separately by each identity.

- In ILNP, the overhead of updating the CNs is handled directly by the MR. With NEMO (and MIP route optimisation), this overhead is handled by each respective VMN, but only occurs once when they first join the network.
- In ILNP, there is the potential for reduced overhead, as the MR will not send duplicate LU messages where VMNs have common CNs.
- The default mechanism for ILNP ensures direct-route paths between the VMN and its CNs, as it uses the normal (optimal) unicast routing from its current point of attachment. Also, this is established when the VMN joins the mobile network for the first time. For NEMO, the VMNs (which in this case runs MIPv6) have to go through additional steps to ensure this (Figure 6).

So, while the protocol overhead of NEMO (with route optimisation) and ILNP is similar, we believe that ILNP offers a simpler protocol architecture compared to NEMO (with and without route optimisation). Additionally, we have described other benefits of ILNP previously, such as harmonised traffic engineering and multi-homing support with mobility for hosts and mobile networks [1]. Also, by leveraging the mature work in secure, dynamic DNS updates, existing infrastructure can be leveraged, rather than having to deploy and maintain Home Agents.

VI. CONCLUSION

Mobile networks are an important capability for current and future military networks. The current IETF work on Network Mobility (NEMO) takes a tunnelling-based approach. We have compared this to a new approach based on the use of a locator/identifier mechanism applied to IPv6 – ILNPv6. We find that protocol overhead is similar for ILNPv6 and NEMO with route optimisation. ILNPv6 always uses the optimal (normal unicast) route(s), offers the potential for a simpler protocol architecture, and has novel features such as soft-handover at the IP level. ILNPv6 also offers easy integration with other capabilities such as IPsec, multi-homing, traffic engineering and localised addressing.

So, we propose that the use of secure naming should be further investigated as a viable alternative for the provisioning of mobile IP networks.

A. Future work

We are currently building an implementation of ILNPv6 and intend to perform comparative experiments on a testbed in order to verify the analysis presented in this paper. Additionally, we are currently conducting experiments on the use of DNS with very low time-to-live (TTL) values

for DNS records, in order to support high mobility rates through naming. Our initial investigations suggest that low TTL values (as low as zero) are easily tolerable for certain DNS records, such as *A* and *PTR* records, for example.

REFERENCES

- [1] R. Atkinson, S. Bhatti, and S. Hailes, "Harmonised Resilience, Security and Mobility Capability for IP," in *27th IEEE Military Communications Conference*. San Diego, CA, USA: IEEE, Nov. 2008.
- [2] B. Wellington, "Secure Domain Name System (DNS) Dynamic Update," RFC 3007, Nov. 2000.
- [3] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol," RFC 3963, Jan. 2005.
- [4] R. Atkinson, S. Bhatti, and S. Hailes, "A Proposal for Unifying Mobility with Multi-Homing, NAT, and Security," in *5th ACM International Workshop on Mobility Management and Wireless Access - MOBIWAC2007*. Chania, Crete, Greece: ACM, Oct. 2007.
- [5] R. Hinden, S. Deering, and E. Nordmark, "IPv6 Global Unicast Address Format," RFC 3587, August 2003.
- [6] C. Bennett, S. Edge, and A. Hinchley, "Issues in the Interconnection of Datagram Networks," ARPA Network Working Group, Internet Experiment Note (IEN) 1, July 1977.
- [7] I. Castineyra, N. Chiappa, and M. Steenstrup, "The Nimrod Routing Architecture," RFC 1992, August 1996.
- [8] M. O'Dell, "GSE - An Alternate Addressing Architecture for IPv6," IETF, Internet-Draft draft-ipng-gseaddr-00.txt, Feb. 1997.
- [9] B. Carpenter, "Architectural Principles of the Internet," RFC 1958, June 1996.
- [10] R. Hinden and S. Deering, "IP Version 6 Addressing Architecture," RFC 4291, Feb 2006.
- [11] R. Atkinson, S. Bhatti, and S. Hailes, "Mobility Through Naming: Impact on DNS," in *MobiArch 2008*. San Diego, US: ACM, August 2008.
- [12] A. C. Snoeren and H. Balakrishnan, "An End-to-End Approach to Host Mobility," in *MobiCom '00: Proceedings of the 6th annual international conference on mobile computing and networking*. New York, NY, USA: ACM, 2000, pp. 155–166.
- [13] P. Mockapetris and K. J. Dunlap, "Development of the domain name system," *SIGCOMM Comput. Commun. Rev.*, vol. 18, no. 4, pp. 123–133, 1988.
- [14] A. Pappas, S. Hailes, and R. Giaffreda, "Mobile Host Location Tracking through DNS," in *IEEE London Communications Symposium*. London, England: IEEE, September 2002.
- [15] B. Yahya and J. Ben-Othman, "Achieving host mobility using dns dynamic updating protocol," in *Local Computer Networks, 2008. LCN 2008. 33rd IEEE Conference on*, October 2008, pp. 634–638.
- [16] J. Jung, E. Sit, H. Balakrishnan, and R. Morris, "DNS Performance and the Effectiveness of Caching," *IEEE/ACM Transactions on Networking*, vol. 10, no. 5, pp. 589–603, 2002.
- [17] US DoD, "High-Assurance IP Encryption Interoperability Specification (HAIPE IS), Version 1.3.5," May 2004.
- [18] S. Kent and R. Atkinson, "Security Architecture for the Internet Protocol," RFC 2401, Nov. 1998.
- [19] P. Albitz, *DNS and BIND*. Sebastopol, CA, USA: O'Reilly & Associates, Inc., 2001.
- [20] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security Introduction and Requirements," RFC 4033, Mar. 2005.
- [21] H. Petander, E. Perera, K.-C. Lan, and A. Seneviratne, "Measuring and Improving the Performance of Network Mobility Management in IPv6 Networks," *Selected Areas in Communications, IEEE Journal on*, vol. 24, no. 9, pp. 1671–1681, Sept. 2006.