Mobility Through Naming: Impact on DNS

R. Atkinson
Extreme Networks, USA
rja@extremenetworks.com

S. Bhatti
University of St Andrews, UK
saleem@cs.st-andrews.ac.uk

S. Hailes
UCL, UK
s.hailes@cs.ucl.ac.uk

ABSTRACT

An Identifier / Locator addressing scheme can enable a new approach to mobile hosts and mobile networks. Identifier and Locator information is stored in Domain Name System (DNS) Resource Records (RRs). In our on-going work using the Identifier – Locator Network Protocol (ILNP), the DNS would be updated with new Locator values as hosts and/or networks move: new sessions would obtain the correct Locator(s) for a mobile host and/or network from the DNS, in much the same was as currently happens for IP address resolution. However, this use of the DNS is not currently required for mobility using IP. We examine the potential impact on DNS from using a naming approach to mobility.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Network communications; C.2.2 [**Network Protocols**]: Protocol architecture

General Terms

Design

Keywords

Naming, Addressing, Routing, Identifier, Locator, Mobility

1. INTRODUCTION

As interest grows in new approaches to both mobile hosts and mobile networks, we must consider the impact on other parts of the system architecture. There are a number of proposals that change the way that addressing is used in order to provide mobility. If we consider a naming based approach to mobility that uses naming of locations in order to provide part of the end-point addressing, then we must (i) allow the mobile host or network to be discovered for establishing new sessions; and (ii) permit update of existing sessions as new Locator values become used [12]. In this paper, we focus on the first of these issues: a mechanism for allowing hosts names to be resolved to a (set of) valid Locator value(s). In our approach (Section 2), a topologically-significant name, the Locator (L), is updated when a host or network moves. The update of the L value must be applied to the Domain Name System (DNS) to allow correct name resolution for new sessions: for existing sessions, the Locator value must be communicated to the correspondent node.

In this paper, firstly, we summarise our approach and use of the Identifier Locator Network Protocol (ILNP) (Section 2). Then, we discuss the mobility scenarios and present the changes and impact in the use and access of the Domain Name System (DNS) required for ILNP (Section 3). After this, we make an evaluation of the additional (indirect and non-functional) issues related to mobility through naming and the use of the DNS (Section 4). We then list some future work items (Section 5) and conclude (Section 6).

MobiArch'08, August 22, 2008, Seattle, Washington, USA.

2. MOBILITY THROUGH NAMING: ILNP

We describe here the salient features of the Identifier Locator Network Protocol (ILNP): details can be found in [1], and we present here a summary. ILNP describes an architecture, which could be applied to several networking protocols. For pragmatic reasons, we present a specific instance of ILNP, which we call ILNPv6, which has been designed as a backwards-compatible extension to IPv6.

2.1 ILNPv6

ILNPv6 proposes to split the IPv6 Address into two distinct components. The upper 64 bits is the *Locator*, and the lower 64 bits is the *Identifier* (see Figure 1). A Locator names a single subnetwork and is topologically significant, but is never used for identity. An Identifier names a single node (not an interface) and is never topologically significant. This split enables an improved network architecture, particularly with respect to mobility and multi-homing.

In the ILNP architecture, the set of Identifiers used by a node can be very long-lived, but the set of Locators could be very short-lived. Normally, as a node moves from one point of network attachment to another, the Identifier(s) are constant, but the Locator(s) change with each move to a different subnetwork. With ILNP, upper-layer protocols (e.g. TCP and UDP) include the Identifier (I) in their session state (and pseudo header), but never include the Locator (L). Ideally, application protocols (e.g. SSH and HTTP) will use fully-qualified domain names as part of the application-layer name, but they may chose to use the Identifier.

We note that separating Location from Identity to support mobility is not new [2]; certainly the concept has been proposed in NIMROD [3, 11], in the GSE proposal for IPv6 [9], and in HIP [8]. Our proposal differs from those previous concepts in various ways, and has defined the engineering in much more detail than either NIMROD or GSE did.

Figure 2 shows the packet time-sequence diagram for a network-layer handoff using ILNPv6. Since ILNPv6 supports stateless auto-configuration, DHCP is not required. After movement is detected,

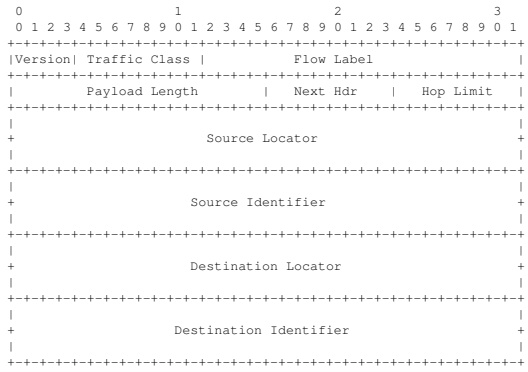

Figure 1: ILNPv6 basic packet header

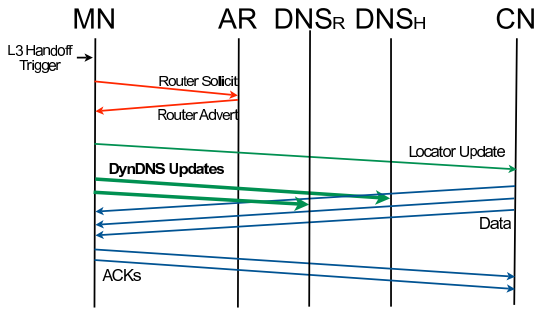

MN	mobile node
AR	router serving MN
DNS_R	DNS Server (reverse)
DNS_H	DNS Server (forward)
CN	correspondent node

Figure 2: ILNP mobility handoff time-sequence diagram

only one RTT delay and one Locator Update are required before data in an existing session can flow from the correspondent node to the mobile node at its new location.

ILNP can provide much lower network-layer handoff latency than either version of Mobile IP (v4 or v6), and Duplicate Address Detection (DAD) is avoided by generating Identifier values by using the bits from any IEEE MAC address associated with a given node. Neighbour Discovery for ILNPv6 works as for IPv6, using the concatenation of the Destination Locator and the Destination Identifier (I:L) in lieu of the IPv6 Destination Address, except that DAD is not required.

As with the current deployed Internet, an initiator of a new session will query the DNS to determine where to send packets for the responder. Unlike Mobile IP, Home Agents or Foreign Agents are not needed or used. Instead, following an earlier proposal by the third author, the Secure Dynamic DNS update ensures that the Locator values stored in the DNS for this node are kept current [10]. Both BIND and Microsoft Windows already support Secure Dynamic DNS Update, so this approach is practical to use at present.

2.2 Mobile networks

An increasing consideration with IP is that of multi-homing (independent of Mobile IP). Additionally, there is a growing interest in mobile networks, i.e. from the former IETF NEMO WG charter: *"The NEMO Working Group is concerned with managing the mobility of an entire network, which changes its point of attachment to the Internet and thus its reachability in the network topology. The mobile network includes one or more mobile routers (MRs) which connect the rest of the mobile network to the global Internet.*
For the purposes of this working group, a mobile network is a leaf network; it does not carry transit traffic. Nonetheless, it could be multihomed, either with a single MR that has multiple attachments to the Internet, or by using multiple MRs that attach the mobile network to the Internet."

Site multi-homing primarily is used to increase Internet availability. So if one of a site's upstream links develops a fault (e.g. due to a fibre cut), then traffic will be able to flow over another upstream link between the same nodes. At present, this is typically enabled by having a site advertise a site-specific routing prefix through all of its upstream links. In turn, this site-specific prefix is propagated throughout the Internet core, thereby increasing entropy of the inter-domain routing system. However, the central issue is how to ensure that an existing IP session will continue to work even if the upstream link changes from one (e.g. primary) provider to another (e.g. backup) provider. So site multi-homing is mainly about maintaining availability when the upstream point of network attachment moves. Node multi-homing is not common in the deployed Internet today.

With this insight, ILNPv6 handles node multi-homing, site multi-homing, node mobility, and network mobility using the same mechanism: through the change in the value of the Locator(s), L. When a mobile node moves to another IP subnetwork, or a multi-homed node changes its set of upstream links, the value of L will change. The affected node or set of nodes discover the change locally from Router Advertisements. An ILNPv6 node may hold and use more than one value of L concurrently if it is multi-homed, whether through multiple interfaces on the node, through a single router that happens to be multi-homed, or through multiple routers, each offering a different value for L. With ILNPv6, a mobile network is merely a special case of site multi-homing: values of L can be changed as site connectivity changes.

2.3 Soft-handoff

Radio engineering support for mobility includes "soft-handoff" where two radio channels are used simultaneously, the current channel that will be released and the new channel that is being acquired, for a smooth transition between channels. For example, the current channel could be held until the completion of update of all correspondents' session state to the new channel. So, the use of soft-handoff methods, is recommended. Using soft-handoffs is beneficial regardless of which network mobility approach is chosen.

Such soft-handoff is also supported by ILNPv6 at the network layer even if radio-layer support is not available. As ILNPv6 supports use of multiple Locator values simultaneously, packets can be sent on the current and new subnetworks during handoff to allow smooth transition. Note that this does not have any impact on DNS other than that as described in Section 3: the *L* record will be updated when the new Locator value is acquired.

3. DNS AND ILNPV6

We first give a summary of the DNS resource records (RRs) used by ILNPv6, and then go onto look at the impact on DNS under given scenarios. Table 1 presents a summary of the DNS RRs that are discussed in this section. A summary of the impact on DNS access is given in Table 2.

3.1 Summary of DNS usage in ILNPv6

ILNPv6 requires the creation of two new Domain Name System (DNS) Resource Records, an *I* record and an *L* record. The *I* record is used to hold the Identifier(s) associated with a domain-name, while the *L* record is used to hold the Locator(s) associated with that same domain-name. Each *L* record has a preference value associated with it, similar to the current DNS practice for Mail Exchanger (*MX*) records. Having this *L* record preference field permits the node associated with the *L* records to inform correspondents of the relative preferences among the several *L* records associated with that node. Normally, if one requests either the *I* or *L* records for a given domain-name, then all *I* and all *L* records associated with that domain-name are returned.

As an optimisation, three other DNS resource records are also added. The *PTRL* and *PTRI* records are used together to perform a reverse lookup, and the *LP* record is used to associate a fully qualified domain name (FQDN) to a (set of) *L* records for a network. Functionally, one could use the existing IPv6 reverse pointer record, but this approach permits the *PTRL* record information to be cached, which improves performance when there are lookups for multiple nodes on the same subnetwork within the specified time-to-live (TTL) value.

When one performs a *PTRL* lookup on a given Locator value, the FQDN of the authoritative DNS server for that subnetwork is returned. In turn, if one then sends a *PTRI* lookup request with some Identifier value to that authoritative DNS server, then the FQDN of the node on that subnetwork with that Identifier is returned.

As a performance optimisation for mobile networks or site multi-homing, ILNP introduces the use of a FQDN to name a *network*. One or more DNS Locator records can be associated with the network name, so that only one DNS update is needed when the named network changes its point of attachment, regardless of how many nodes are on the named network. A node that is connected to such a named network may use a Locator Pointer (*LP*) record in place of an *L* record. Where an *L* record would provide a 64-bit Locator associated with the node, a *LP* record provides the FQDN of the mobile network that the node is connected to. So for a FQDN resolution of a node's domain-name, the *I* records, the *L* records (if any), and the *LP* records are all returned. The correspondent then performs an *L* record lookup in the FQDN found in the *LP* record to learn the actual numeric Locator value(s), i.e. there is a single "site" *L* record pointed to by each node in the network rather than duplicated data in separate *L* records for each host. This helps with managing the DNS data, but also with mobile networks. Alternatively, one could use individual *L* records for hosts, at the cost of numerous, individual DNS updates when a mobile network moves.

3.2 Fixed network and correspondent node

For a fixed node using ILNP, or a node wanting to discover the ILNP address (I:L) for a remote host, the situation is much as it is now for IPv6: a FQDN is used to make a DNS lookup. However, instead of *AAAA* record(s) being returned, *I* and *L* records are returned. For a fixed host (that is not multi-homed or mobile), this normally will be a single I value and a single L value. For such a node, the TTL values for those DNS records can be relatively long (e.g. hours or days), again, much as they are today.

Additionally, whilst we have described the DNS *LP* record mechanism primarily for use in mobile networks, it is also useful for fixed networks using ILNP. In all cases, it acts to reduce the volume of the data in a DNS server for a site, and also to improve manageability of the DNS data.

IMPACT: For a fixed network host, the cost of using DNS with ILNPv6 is no more than it is today for IPv6.

Name	Description	Purpose
I	Identifier Record	Identifier values for a host
L	Locator Record	Locator values for a host or network, including relative preference
PTRI	Reverse Identifier	permits reverse lookup of FQDN from Identifier value
PTRL	Reverse Locator	permits reverse lookup of FQDN from Locator value
LP	Pointer to Locator	names a network using an FQDN, resolves to an FQDN, which in turn resolves to an *L* record, containing the Locator value for a host or network

Table 1: Summary of DNS resource records for ILNP

3.3 Mobile client

If one considers a scenario where some client nodes are mobile, the only affected DNS resource records are those that belong to the mobile clients themselves. Mobile clients will have short TTL values for their DNS *L* records to ensure that stale DNS data is not provided to others.[1] Mobile clients are at the edge or leaf of the DNS tree. Any additional traffic generated is confined to edge DNS servers and does not affect root servers, TLD servers, or even the top-of-user-domain servers.

The ILNP Locator Update (LU) is analogous to the MIPv6 Binding Update, updating existing IP sessions, so there should be little additional DNS traffic due to the use of the *L* record, as new communication sessions would have to start with a DNS lookup anyway in MIPv6. Prior research indicates that there is very little locality in DNS queries for end systems. This means that DNS caching of end node *A* records is ineffective in today's deployed Internet [6]. Even if DNS caching of *L* records is not very successful, it will perform no worse than today's Internet.

IMPACT: If 10% of nodes are mobile (just pulling a number out of the air), then the extra traffic overhead directed at the leaf DNS servers ought not be considerable. It is unclear what the relative numbers of mobile hosts will be so the total impact is hard to estimate. Every mobile node will have to update its *L* record, affecting leaf servers but not the root server.

3.4 Mobile server

The use of changing Locator values can support a mobile server as well. However, in many quarters, the notion of mobile servers is somewhat novel. At present, there are no known large-scale deployments of mobile servers. So the full deployment requirements for mobile servers, i.e. what the numbers are likely to be relative to mobile clients, are not clear. [2]

Potentially, a server could be the end-point of many communication sessions. If it moves, this could result in a large number of Locator Updates (as with Binding Updates in Mobile IPv6). However, only the *L* record(s) of the server needs to be updated in the DNS, which would be a single DNS transaction.

Potentially, a mobile server could result in extra DNS traffic, as the TTL for the *L* record for that server will now be lower than for

[1] Non-mobile clients are not affected and can have much longer TTL values (e.g. hours or days, as today).

[2] A SIP handset is usually deployed with a separate SIP server as intermediary, but handset to handset communications is possible. Since the number of concurrent calls for a single handset is relatively small, scalability ought not be an issue for a SIP handset.

a fixed host. A busy server will then have *L* records that cannot be cached for as long as they could be for a fixed server.

If one considers this in a naive way, one might conclude this would be a problem for widely-used public web servers. However, large public web services usually are not a single server, but instead have a large number of servers, often geographically replicated in various locations, often also located behind load-balancing appliances, such that multiple servers can provide the same content. Also, URL-based referrals are often used by such large public web sites. For this common case, the named web server only makes a decision about which content server should handle a given client request, and then sends a URL referral to the web client, redirecting the client to the chosen content server. In such cases, a wide range of variables, including server load, latency, and network congestion might be included in the server operator's decision of which content server should handle a particular client or client request.

IMPACT: It is far from clear that much DNS locality exists in the current Internet, at least with respect to *A* or *AAAA* records of edge nodes (including servers). If the current Internet does not have much DNS locality in that respect, then the change to ILNP will not have much DNS impact in this scenario. However, if a server is mobile, when it moves network it will have to update its *L* record, affecting leaf servers but not the root server.

3.5 Mobile network

Mobile networks are handled in the same way as site multi-homing. So our approach is significantly different than that taken by the former IETF Network Mobility (NEMO) Working Group. In ILNP, mobile networks are handled "automatically", using the same mechanisms described above for site multi-homing.

As an engineering optimisation, the Locator Pointer (*LP*) record described above will significantly reduce the number of DNS update transactions that are required when a mobile network changes its point(s) of network attachment. So only one DNS update per subnetwork will be required, rather than one DNS update per node attached to the mobile subnetwork.

IMPACT: Where a *LP* record is configured, there is now an extra stage to a DNS resolution as the initial resolution results in an *LP* record, which then has to be resolved to an *L* record. Again, we can not be sure exactly how much extra DNS traffic this will generate as there are not many deployed mobile networks at present (and we have access to none ourselves). However, potentially the cost for DNS is minimal, requiring an update to the "site" *L* record pointed to by the *LP* record. As the LP record itself can be cached, DNS load will be reduced by any locality in the traffic patterns.

3.6 Simultaneous movement

For the case where two hosts (individual mobile hosts, or as part of a network) move at the same time, there is a synchronisation problem, potentially. ILNPv6 handles this in two ways. Firstly, ILNPv6 layer soft-handoff is recommended, in order to allow Locator Update messages to traverse the network and update session state. Secondly, if the two hosts were unlucky enough to lose synchronisation (all Locator Update messages are lost), then (after whatever Transport protocol timeout applies has occurred) they perform a DNS lookup to find the updated *L* records, and resume the session. Application state should not be adversely affected.

IMPACT: If mobile hosts have to resume the session, to the DNS, this looks like a new session being started and there are no additional mechanisms required.

3.7 DNS usage and applications

For applications using ILNPv6, in general, the current DNS interaction model need not change: stateless, request-response protocol would suffice. However, there may be scenarios in which the way that DNS is used by applications would change. In this discussion, the 'receiver' is the recipient of the response to a DNS query.

If a fixed or correspondent node performs a DNS lookup and receives multiple *L* records (e.g. because the remote host is multi-homed), the node should check the *preference* value in the *L* records to determine which L value to use first in order to send packets to the remote node and establish the session.[3]

In a mobile network, if an *LP* record is returned, the receiver needs to perform another DNS lookup on the FQDN contained in the *LP* record in order to resolve the *L* record(s). However, an *LP* record could also be returned if the site is multi-homed: *LP* records can be used to optimise DNS access by allowing an *LP* record to point to a single *L* record, the latter being the only record updated when the Locater value changes.

IMPACT: Existing applications can rely on the DNS Resolver library to hide the ILNP details from them. To enable this, the DNS Resolver library will need to be updated to support *I* , *L* , and *LP* records. When an *LP* record is encountered, the library must perform an extra DNS lookup to find any correspodning *L* records.

Scenario (mobile entity)	Section	Extra DNS access (mobile host and correspondent)
Fixed	3.2	correspondent: possible extra access required if an *LP* record is used for a multi-homed site
Client	3.3	host: single extra access for update of *L* record(s)
Server	3.4	host: single extra access for update of *L* record(s)
Network	3.5	host: extra access to update multiple *L* records (all hosts in the mobile network), unless an *LP* records is used, and then only a single extra access for the network as a whole to update of the *LP* record is required correspondent: if *LP* record is returned, extra access to resolve network name to *L* record(s)
Simultaneous movement	3.6	same as Client scenario

Table 2: Summary of impact of ILNP on DNS access

4. ADDITIONAL CONSIDERATIONS

As well as the functional considerations for assessing the feasibility and impact of ILNPv6 on the DNS, we must also consider some non-functional issues. In real world systems, non-functional issues may have significant impact on the suitability of a system for use. So, in this section, we discuss the salient issues for ILNPv6. A summary of the issues in this section are given in Table 3.

4.1 Overall traffic considerations

Overall, as ILNPv6 does not use a Home Agent (HA) and Foreign Agent (FA) configuration, no extra signalling traffic is required in order to update the HA, as is the case for IPv6. However, ILNPv6 generates DNS traffic which Mobile IPv6 does not. IPv6 Binding

[3]Correspondent nodes may use any valid *L* record of the responder node to send packets to the responder.

Update messages and ILNPv6 Locator Update messages have, almost, a one-to-one correspondence. So, it could be argued that, when comparing Mobile IPv6 to ILNPv6, the overall traffic overhead is comparable: the signalling traffic for each just happens to go to a different entity (HA vs DNS).

For mobile networks, if we compare the IETF NEMO work to ILNPv6, then, at the time of writing, potentially, there is a significant gain in simplicity of the architecture, as ILNPv6 does not require tunnelling to HA entities [13]. However, ILNPv6 sacrifices the transparency goals that have been defined for NEMO for backwards compatibility with normal IPv6 nodes [4]. A more detailed analysis would be required in order to determine the comparative overhead under specific scenarios. Potentially, with ILNPv6, there is the issue of a Locator Update "storm" as session state for many live sessions is updated as a mobile network changes Locator. We consider this an issue for further study.

4.2 System robustness

The use of the HA with most current Mobile IPv6 and NEMO implementations, whilst providing transparency, is potentially a weakness for overall system robustness. It is very likely to be a single-point of failure, and also a potential performance bottleneck. The IETF has undertaken work to address this (e.g. through use of multiple HAs), although the numerous extensions to Mobile IP and NEMO tend also to increase the complexity of Mobile IP and NEMO implementation and deployment.[4]

If there are failures in connectivity between the home network and the correspondent and/or between the home network and the mobile host or network, then the mobile capability is lost. As well as connectivity failures, the home network connectivity listed above now becomes a point of attack for a malicious party wanting to disrupt communications to/from the mobile host/network. It is not possible to have multiple home networks with Mobile IPv6 or NEMO, though multiple home agents may be possible.

This is not the case for ILNP. For initialising a session, the correspondent needs to contact a DNS server, and the mobile host/network needs to keep the DNS entry updated. However, it is possible to have multiple DNS servers that are geographically and/or topologically dispersed to provide the DNS service, helping alleviate potential connectivity problems or malicious intervention. DNS already provides this capability – no additional engineering is required, though here is additional cost in the configuration and management of the DNS service in order to provide this capability. Use of Secure Dynamic DNS Update does not preclude use of redundant DNS servers for the dynamic DNS zone.

4.3 Deployability

BIND and commercial DNS servers (e.g. from Nominum or Microsoft) are widely available, widely deployed, and support DNS Security today. As the proposed DNS enhancements are similar to existing DNS resource records, it should be simple to add support for the new *I* , *L* , *LP* , *PTRI* , and *PTRL* records to existing DNS servers. Further, the only change to DNS processing rules is that a server should return all *I* , *L* , and *LP* records for a fully-qualified domain name when a query is received for any one of those three record types. This is similar to existing processing rules for providing "additional data" in certain DNS responses.

While some might view the use of Secure DNS Dynamic Update as adventurous, BIND and Microsoft implement that specification and can interoperate. So this usage is now quite reasonable.

[4] At the time of writing, the IETF Mobile EXTensions (MEXT) WG has been formed and is considering the harmonisation of MIPv6, NEMO, and IKEv2.

4.4 Authentication

The proposed DNS enhancements do not alter the security properties of the existing DNS. The proposed enhancements create no new vulnerabilities. Further, DNS Security would not need alteration. So the security risks of the DNS are unchanged, and the prospective security solution, DNS Security, is also unchanged.

4.5 Scalability impacts upon DNS

With the current IPv4 Internet, scalability of the DNS depends upon the ability of edge DNS resolvers (closer to the end-user) to cache (1) the *NS* records used to indicate zone delegation[5], and (2) also the *A* records of the upper-level DNS servers associated with those *NS* records (closer to the root). In turn, this depends upon that small set of *NS* and *A* records having moderately long TTLs. Research published earlier this decade [6] indicates that *giving low TTL values to A records will not significantly harm hit rates*.

Secure Dynamic DNS Update is standardised and widely implemented. At least some mobile networks are using it today to optimise initial contact with a mobile node. Early operational experience indicates that this is a reasonable approach.

Additionally, the current authors are running an experiment at St Andrews with relatively short TTL values for *A* and *PTR* records. This will measure the offered DNS load versus the TTL values to provide further confirmation that this is not a problem. The use of relatively short DNS TTL values ensures that other nodes will not be given stale L values for a mobile node or mobile network.

Many top-level DNS servers, for example F-Root, use BGP anycasting with replicated DNS servers, rather than using BGP site-multihoming [5]. BGP anycast techniques work unchanged with ILNP, so DNS servers at, or closer to, the root can continue to have moderately long lifetimes for their *A* (or *AAAA* or *L & I*) records. Also, DNS resolvers will continue to cache learned DNS resource records for the configured TTL values for each learned DNS resource record. So the deployed DNS should continue to scale as well with ILNP as it does for the current IPv4 Internet.

4.6 Link mobility considerations

Mobility research at NATO [7, 14] indicates that there are a number of limitations with existing Mobile IP approaches; that work separately indicates that to be fully successful a mobile node needs to also use use link-layer mobility mechanisms so that the network-layer mobility events (e.g. IP handoff) are not as frequent.

Operationally successful mobility is a multi-layer issue, requiring a multi-layer approach. Network-layer capabilities are important, but are only part of the solution. While some believe that Mobile IP needs to be able to handle extremely rapid (e.g. 1 second) changes in point of network attachment by itself, we believe that IP mobility needs to be combined with link-layer mobility.

The ILNP mechanism also allows multiple Locators to be in use concurrently. So soft handoff is possible, and recommended. Indeed, many nodes will use multiple Locators (albeit with the same Identifier) – mobility and multi-homing are essentially the same in ILNP. Thus, ILNP can provide more flexibility than some other approaches to node mobility or network mobility.

4.7 Other network layer functions

As explained in [1], ILNPv6 can support localised addressing (i.e. Network Address Translation), true end-to-end IPsec, multi-homing and mobility in an integrated fashion, as "first class" services, rather than requiring additional engineering (e.g. tunnelling),

[5] For example, the delegation by a root server of `.COM` to some set of authoritative DNS servers.

middleboxes or entities (e.g. Home Agents) as does Mobile IPv6 and NEMO. For ILNP support of NAT and IPsec with mobility, there is no additional cost or impact with respect to DNS other than those stated in this paper.

Issue	Section	Summary
Traffic	4.1	DNS traffic little impact, but Locator Update traffic may be an issue, e.g. in mobile network
Robustness	4.2	Use of DNS potentially improves system robustness overall compared to use of HAs
Deployability	4.3	Incremental deployability for DNS capability ILNPv6
Authentication	4.4	No impact
Scalability	4.5	Extra DNS traffic not likely to be significant and existing uses of DNS not impacted
Link mobility	4.6	ILNP will can function within a multi-layer approach including support for soft-handoff without affecting DNS
Integration	4.7	ILNP easily integrated with other network functions

Table 3: Summary of issues where ILNP may impact DNS

5. FUTURE WORK

More measurement and experimentation with the actual current DNS impacts of mobile networks and mobile servers is needed in order to more fully understand the current behaviour and to be able to either model or experiment with the behaviour likely when ILNP is in use. Part of the challenge in this area is that very few deployments of either scenario exist in today's Internet. Another challenge is that differing communities have different visions for how they would like to use IP mobility.

An exciting development is application-controlled traffic engineering at the transport-layer through the use of naming: with multiple Locator values. Consider the case that a node is multi-homed and so has multiple Locator values. Providing a session uses the same Identifier value in the session state, the session is free to swap Locator values or use multiple Locator values simultaneously in ILNPv6. So, by setting an equal preference value for each Locator value, a multi-path session can be established: the use of multiple Locators affects the routing of those packets, but the single Identifier value ensures that they are all delivered to the same session. Of course, use of the Locator values in this way deals with the address management for multi-path flows; the important issue of multi-path flow congestion control requires further study.

6. CONCLUSIONS

A naming approach to enabling mobility has great benefits, at the expense of invisibility to nodes that are not mobile-aware. We also note that our proposed approach can be deployed incrementally: ILNPv6 packets look like IPv6 packets to the core network, and no updates are required for edge networks that do not plan to support ILNPv6. As can be seen from Tables 1, 2 and 3, the use of a naming approach to mobility (1) is unlikely to have a significant adverse impact on the DNS; (2) may yield some benefits, for example enabling network-layer soft-handoff, helping increase overall system robustness, and permitting integration of other network functions, without impacting DNS; and (3) can be deployed incrementally and securely into existing infrastructure, because it uses the standard DNS security mechanisms already being rolled out for the deployed IP Internet.

However, some issues require further study. The actual potential impact on DNS needs to be investigated through rigourous experimentation. Such an analysis should include a study on the overall impact on traffic under certain scenarios, particularly for the mobile network scenario and mobile server scenario.

7. REFERENCES

[1] R. Atkinson, S. Bhatti, and S. Hailes. Mobility as an Integrated Service Through the Use of Naming. In *Proceedings of 2nd ACM Workshop on Mobility in the Evolving Internet Architecture (MobiArch)*, Kyoto, Japan, August 2007. ACM.

[2] C. Bennett, S. Edge, and A. Hinchley. Issues in the Interconnection of Datagram Networks. Internet Experiment Note (IEN) 1, ARPA Network Working Group, July 1977.

[3] I. Castineyra, N. Chiappa, and M. Steenstrup. The Nimrod Routing Architecture. RFC 1992, IETF, Aug. 1996.

[4] T. Ernst. Network Mobility Support Goals and Requirements. RFC 4486, IETF, July 2007.

[5] S. Gibbard. Geographic Implications of DNS Infrastructure. *Internet Protocol Journal*, 10(1):12 – 24, Mar. 2007.

[6] J. Jung, E. Sit, H. Balakrishnan, and R. Morris. DNS performance and the Effectiveness of Caching. *IEEE/ACM Transactions on Networking*, 10(5):589 –603, October 2002.

[7] J. Macker. Interoperable Networks for Secure Communications, Task 6, Phase 1. Final Report INSC-TASK6, North Atlantic Treaty Organisation (NATO), Dec. 2003.

[8] R. Moskowitz and P. Nikander. Host Identity Protocol (HIP) Architecture. RFC 4423, IETF, May 2006.

[9] M. O'Dell. GSE - An Alternate Addressing Architecture for IPv6. Internet-Draft draft-ipng-gseaddr-00.txt, IETF, Feb. 1997.

[10] A. Pappas, S. Hailes, and R. Giaffreda. Mobile Host Location Tracking through DNS. In *Proceedings of 2002 London Communications Symposium*, London, England, UK, Sept. 2002. IEEE.

[11] R. Ramanathan. Mobility Support for Nimrod : Challenges and Solution Approaches. RFC 2103, IETF, Feb. 1997.

[12] A. C. Snoeren, H. Balakrishnan, and M. F. Kaashoek. Reconsidering Internet Mobility. In *Proceedings of Eighth Workshop on Hot Topics in Operating Systems (HotOS)*, pages 41 – 46. ACM, May 2001.

[13] P. Thubert, R. Wakikawa, and V. Devarapalli. Network Mobility Home Network Models. RFC 4487, IETF, July 2007.

[14] J. Weston. Interoperable Networks for Secure Communications, Task 3. Final Report INSC-TASK3, North Atlantic Treaty Organisation (NATO), July 2006.