An Introduction to the Identifier-Locator Network Protocol (ILNP)

R. Atkinson, Extreme Networks & S. Bhatti, University of St Andrews

13 July 2006

Abstract

Mobility, multi-homing, local addressing and end-to-end security at the network layer remain challenging even with the advent of IPv6. We propose a new network protocol, which can be built upon IPv6 incrementally, that breaks the address into two separate entities, a Locator and an Identifier, with crisp semantics for each, that seeks to solve these issues through an improved naming and addressing architecture.

1 Introduction and Problem Space

Over a decade after its standardisation, Mobile IPv4 is not widely implemented and is not widely deployed, even though mobile devices are now quite common. Similarly, Mobile IPv6 is neither widely implemented nor widely deployed. In fact, there is much active research on ways to optimise Mobile IP so that it might be more widely used. Both forms of Mobile IP implicitly recognise that an IP Address can be used both for identification and for location. The *Home Address* of Mobile IP is used as an identifier that does not change when the mobile node moves, while the *Care-of Address* of Mobile IP is used as a locator that changes each time a mobile node moves.

Similarly, multi-homing remains a significant issue today. The same basic method is used for multi-homing in both IPv4 and IPv6 – advertising a more-specific routing prefix for the site through the global routing system. This practice creates complexity in BGP deployments, complexity in ISP operations, and significantly increases the entropy in the IP routing tables in the Default Free Zone (DFZ). Internet Service Providers are very concerned that IPv6 does not offer improvements over IPv4 in this area.

Local addressing using Network Address Translation (NAT) is increasingly popular amongst edge networks, both for home networks and for enterprise networks. Although originally seen as a way of circumventing the percieved sparsity of IPv4 addresses, NAT is now used as a network administration tool at edge sites to help partition addressing. NAT should not really be needed for IPv6, but anecdotal evidence from vendors suggests that it is being widely requested by IPv6 customers as they find its use convenient. However, because the IP address is changed in packets as they traverse the NAT, end-to-end state is lost. So the NAT must perform some gateway functions at the transport layer and also at the application layer in some cases.

The use of end-to-end security using IPsec relies on the use of the IP address as an identity and froms part of the state for the Security Asscoiation (SA). As we can see from the discussion above, it is clear that if the IP address does not remain constant, we must retro-fit ’fixes’ to IPsec to deal with NATs, mobile IP and multi-homing.

In all of these situations, we believe that the root problem is in the overloaded semantics of the IP address.[**?**, **?**] The IP address is used at the network-layer for routing and for forwarding packets. However, it is also used at the transport-layer as a host identifier. Further, it is used as an application-layer host identifier. We believe that replacing the address with a *Locator*, used only for forwarding packets, and a separate *Identifier*, used only for node identification, can provide significant benefits ov er the current approaches. Additionally, at the application layer, we assume that only the Fully-Qualified Domain Name (FQDN) is used.

Protocol layer	ILNP	IP
Application	FQDN	FQDN, IP Address
Transport	Identifier	IP Address
Network	Locator, Identifier	IP Address
MAC	MAC Address	MAC Address

Our solution is the Identifier Locator Network Protocol (ILNP). This paper provides an introduction to ILNP and gives an overview of how it provides a better solution for mobility and for site multi-homing.

2 Networking Technology Description

This section describes the ILNP and the related networking protocol changes. While ILNP itself is strictly a network-layer protocol, there are a few related changes to other protocols. For example, there are minor enhancements to the User Datagram Protocol, Transmission Control Protocol, and to the Domain Name System. Also, there are changes that provide better data abstraction opportunities to application authors.

2.1 Network-Layer Enhancements

To enable incremental deployment, we find that ILNP can be built upon IPv6. The packet header for ILNP is nearly identical to the packet header for IPv6. However, we replace each 128 bit IPv6 address with a 64-bit Locator followed by a 64-bit Identifier. The Locator names a single subnetwork and is used only for packet forwarding and in routing protocols. Significantly, the Locator is never used above the network-layer. Further, the Identifier always names a node, rather than naming an interface as an IP Address does.

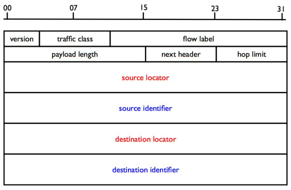

Figure 1: ILNP data packet header format

The Identifier is in IEEE EUI-64 format and is *not* used in forwarding packets. Normally, the Identifier is formed using the IEEE MAC address of some hardware component on the ILNP node. It is not required that an Identifier be globally unique, but it is very helpful if an Identifier is very likely to be globally unique. Existing IEEE processes for MAC address allocation meet this requirement.

Each end system maintains a network-layer mapping table for all current ILNP sessions. This table caches the current Identifier and Locator mappings. The table contents may be updated dynamically by new *ICMP Locator Update* control messages. Those ICMP messages can be authenticated cryptographically (e.g. using the IP Authentication Header) or non-cryptographically, by use of a new IPv6 Destination Option, depending on one's threat environment.

2.2 Transport Layer Enhancements

In each place that an IP address would be used in the Transport layer layer in IPv4 and IPv6, in ILNP, the equivalent Identifier is used. So in this new scheme, only the Locator will change when a node moves or otherwise changes its connectivity. Since the Identifier does not change in such situations, mobility, multi-homing, and NATs have no adverse impact on an ILNP system.

Transport-layer protocols used with ILNP include the Identifier(s) in their transport state (e.g. pseudo-header checksum calculations), but are unaware of the Locator(s) in use for any session. This ensures that changes in location have no impact on transport-layer sessions, whether the location changes were due to a mobile node moving, a multi-homed node losing or gaining an extra uplink, or network-address translation.

2.3 Domain Name System Extensions

ILNP adds four additional Resource Record types to the Domain Name System (DNS). The A6 or AAAA record is replaced by two new records. The *I* record contains a 64-bit Identifier, while the *L* record contains a 64-bit Locator. The process of a reverse lookup is changed slightly and two new records are added to support this function. The *PTRL* record is used to find the authoritative DNS server for a given Locator value, while the *PTRI* record is used to find the fully-qualified domain name of a given Identifier at that location.

As part of the mobility concept, a node will use the Secure Dynamic DNS Update [**?**] protocol to update its location(s) in the DNS whenever its location changes. This helps enable new correspondents to find the node without requiring any special-purpose infrastructure.

The existing DNS Security protocols being standardised by the IETF can also protect the new DNS resource records for ILNP. So use of ILNP does not increase risk or decrease the availability of standard security mechanisms for

DNS information.

Note only DNS server instances providing services for ILNP-enabled nodes need to be modified. There is no need to update DNS software otherwise.

2.4 Application Opportunities

With ILNP, applications should be written to use fully-qualified domain names for all nodes, rather than using IP addresses. Because the BSD Sockets networking API was specified before the creation of the Domain Name System, the Sockets API requires the use of raw IP addresses. This will be addressed by creation of a new networking API that is oriented around domain names instead of using raw IP addresses. With this new API, applications would not need to use the existing DNS library function calls to translate a domain name into an IP address. Instead, the kernel would be responsible for taking the domain name(s) provided by the application and using the DNS protocol to learn the corresponding Identifier and Locator values.

3 ILNP Advantages

3.1 Mobility Concept

Mobile IPv4 and Mobile IPv6 have been specified for years. However, they are not widely implemented, nor are they widely deployed today. Partly this is because the protocols themselves are complex, requiring changes both to routers and to mobile nodes. Partly this is because of key management challenges which have made it impractical to deploy Mobile IP with reasonable cryptographic security. There are so many concerns and issues with aspects of Mobile IP that a veritable cottage industry has arisen proposing various ”fixes” or ”optimisations” to Mobile IP. One of the larger concerns with Mobile IP has been the use of “triangle routing”. With Mobile IP, messages from a correspondent to the mobile node travel via the *Home Agent*, thereby forming two sides of the triangle. Messages from the mobile node back to the correspondent travel directly, forming the third side of the triangle.

ILNP enables mobility using mechanisms that are only in end-systems and do not require any router changes. The key observation is that Mobile IP uses the *Home Address* as if it were an identifier (e.g. in transport-layer session state) and uses the *Foreign Address* as a locator (i.e. to forward packets). Since ILNP provides crisply distinct Identifiers and Locators, it has a better architectural model to support mobility. When a correspondent node wishes to initiate a communications session with a mobile node, the correspondent simply performs a normal DNS lookup on the mobile node’s domain name, thereby obtaining both the mobile node’s Identifier(s) and also its current Locator(s). When a node moves, it sends new *ICMP Locator Update* messages to all existing correspondents. This lets the correspondents learn the new Locator(s) for the mobile node and update the local mapping between Identifiers and Locators. These ICMP messages can be authenticated. Since the Identifier does NOT change when the Locator changes, the upper-layer protocol state (e.g. TCP session state) is not impacted when a Locator changes.

With this new scheme, ILNP never uses ”triangle routing”. Further, since the Identifier normally is formed from a MAC address built into the node’s hardware, there is no need for Duplicate Address Detection (DAD). DAD can be a significant source of network-layer handoff latency in existing Mobile IP mechanisms. Further, there is no need for special router modifications, for Home Agents, or for Foreign Agents. So one can deploy mobile ILNP nodes without as much pre-planning and expense as compared with Mobile IP.

3.2 Multi-Homing Concept

With IPv4 and IPv6, node multi-homing is not really a supported capability. They do support site multi-homing, by using BGP to advertise the more-specific IP routing prefix that belongs to the multi-homed site. Unfortunately, this means that the more-specific prefix(es) for each multi-homed site must be carried in IP routing tables throughout the default-free zone of the Internet. This causes the total number of IP prefixes in the DFZ to be significantly larger. It also means that the inter-domain routing tables have significant amounts of entropy. At a high-level, the main problem is that the cost of multi-homing a site is borne by the entire global Internet, while the benefits are local to the multi-homed site. It would be preferable to have a solution that did not increase the entropy of the IP routing table and that localised the costs to the site that actually benefits.

ILNP’s separation of Locator from Identifier enables a better approach. With ILNP, each multi-homed node has a Locator for each upstream service provider. Using the same mechanisms as are used for mobility, a node can change its Locators over time as additional upstream connections appear (e.g. due to additional connections being added into a site) or disappear (e.g. due to a back hoe taking out fibre in the ground). So with ILNP, each service provider can advertise only the most aggregated IP prefixes and need not carry any more-specific prefixes.

3.3 Network Address Translation (NAT)

Network Address Translation (NAT) developed widespread deployment starting in the middle 1990s. It is very commonly used between home networks and the global Internet or between corporate networks and the global Internet. Many corporate network managers perceive that NAT devices provide risk reduction and simplify internal address management, so NAT is unlikely to disappear in future. NAT is typically a feature in a router and/or in a network firewall device. When en abled, one side of the NAT normally uses global IP address space, while the other side normally uses private address space.[**?**][**?**] The NAT device changes the IP addresses in the IP packet as the packet moves from one s ide of the device to the other. Since the upper-layer protocols (e.g. TCP pseudo-header checksum) include the IP address, the NAT also needs to recalculate por tions of the transport layer protocols. If the application sends raw IP address es inside the application protocol, either the NAT would need to modify the application protocol or that application will not work properly through the NAT.[**?**] So NAT devices also create problems for the current IP Internet, because the addresses are used above the network-layer.

With ILNP, only the domain name or the Identifier are used above the network-layer. So an ILNP NAT could change the Locators for the packet travelling through the NAT without any impact on the transport-layer protocol or the application itself. This means that one could have all of the perceived benefits of NAT, but without the loss of application transparency and interoperability that exists with IP NAT deployments.

3.4 Incremental Deployment

Because ILNP can be defined as a set of enhancements to IPv6, it should be possible to deploy ILNP incrementally on IPv6 networks. Backbone routers would not require any software changes to support ILNP. Similarly, edge routers would not need to change either IPv6 Neighbor Discovery or IPv6 routing to support ILNP-capable end nodes. Further, it appears that an ILNP-capable end node could concurrently support both ordinary IPv6 and ILNP.

4 Conclusions and Future Work

ILNP is a new network protocol, which can be implemented as a set of IPv6 extensions, that replaces the concept of an IP address with separate Locator and Identifier names. This new set of semantics and the related architectural changes that they enable provide better solutions to mobility and multi-homing for packet networks.

Implementation of ILNP in a POSIX-compatible operating system is planned for the near future. Once an initial implementation exists, then testing will follow. Initial testing will likely be performed within a local network, possibly a sensor network. Once ILNP has been shown to work in a local network, it is likely that wide-area testing over an existing IPv6 research network (e.g. SuperJANET) will be undertaken. Of course the results of such testing might lead to protocol optimisations or corrections.

References

[HS01] M. Holdrege and P. Srisuresh. Protocol Complications with the IP Network Address Translator. RFC 3027, Internet Engineering Task Force, January 2001.

[RMK+96] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de, and E. Lear. Address allocation for private internets. RFC 1918, Internet Engineering Task Force, February 1996.

[Sal93] J. H. Saltzer. On the Naming and Binding of Network Destinations. RFC 1498, Internet Engineering Task Force, August 1993.

[Sho78] J.F. Shoch. Inter-Network Naming, Addressing, and Routing. Internet Experiment Note 19, ARPA Network Working Group, January 1978.

[TS00] G. Tsirtsis and P. Srisuresh. Network Address Translation - Protocol Translation (NAT-PT). RFC 2766, Internet Engineering Task Force, February 2000.

[Wel00] B. Wellington. Secure Domain Name System (DNS) Dynamic Update. RFC 3007, Internet Engineering Task Force, November 2000.