

Internet Research Task Force (IRTF)
Request for Comments: 6748
Category: Experimental
ISSN: 2070-1721

RJ Atkinson
Consultant
SN Bhatti
U. St Andrews
November 2012

Optional Advanced Deployment Scenarios for the
Identifier-Locator Network Protocol (ILNP)

Abstract

This document provides an Architectural description and the Concept of Operations of some optional advanced deployment scenarios for the Identifier-Locator Network Protocol (ILNP), which is an evolutionary enhancement to IP. None of the functions described here is required for the use or deployment of ILNP. Instead, it offers descriptions of engineering and deployment options that might provide either enhanced capability or convenience in administration or management of ILNP-based systems.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the individual opinion(s) of one or more members of the Routing Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6748>.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

This document may not be modified, and derivative works of it may not be created, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	4
1.1.	Document Roadmap	5
1.2.	Terminology	6
2.	Localised Numbering	6
2.1.	Localised Locators	7
2.2.	Mixed Local/Global Numbering	9
2.3.	Dealing with Internal Subnets with Locator Rewriting	9
2.4.	Localised Name Resolution with DNS	11
2.5.	Use of mDNS	13
2.6.	Site Network Name in DNS	13
2.7.	Site Interior Topology Obfuscation	14
2.8.	Other SBR Considerations	14
3.	An Alternative for Site Multihoming	16
3.1.	Site Multihoming (S-MH) Connectivity Using an SBR	16
3.2.	Dealing with Link/Connectivity Changes	17
3.3.	SBR Updates to DNS	18
3.4.	DNS TTL Values for L32 and L64 Records	18
3.5.	Multiple SBRs	19
4.	An Alternative for Site (Network) Mobility	20
4.1.	Site (Network) Mobility	20
4.2.	SBR Updates to DNS	22
4.3.	DNS TTL Values for L32 and L64 Records	22
5.	Traffic Engineering Options	22
5.1.	Load Balancing	23
5.2.	Control of Egress Traffic Paths	24
6.	ILNP in Datacentres	26
6.1.	Virtual Image Mobility within a Single Datacentre	27
6.2.	Virtual Image Mobility between Datacentres - Invisible	28
6.3.	Virtual Image Mobility between Datacentres - Visible	29
6.4.	ILNP Capability in the Remote Host for VM Image Mobility	29
7.	Location Privacy	30
7.1.	Locator Rewriting Relay (LRR)	30
7.2.	Options for Installing LRR Packet Forwarding State	31
8.	Identity Privacy	32
9.	Security Considerations	32
10.	References	33
10.1.	Normative References	33
10.2.	Informative References	34
11.	Acknowledgements	37

1. Introduction

This document is part of the ILNP document set, which has had extensive review within the IRTF Routing RG. ILNP is one of the recommendations made by the RG Chairs. Separately, various refereed research papers on ILNP have also been published during this decade. So, the ideas contained herein have had much broader review than the IRTF Routing RG. The views in this document were considered controversial by the Routing RG, but the RG reached a consensus that the document still should be published. The Routing RG has had remarkably little consensus on anything, so virtually all Routing RG outputs are considered controversial.

At present, the Internet research and development community is exploring various approaches to evolving the Internet Architecture to solve a variety of issues including, but not limited to, scalability of inter-domain routing [RFC4984]. A wide range of other issues (e.g., site multihoming, node multihoming, site/subnet mobility, node mobility) are also active concerns at present. Several different classes of evolution are being considered by the Internet research and development community. One class is often called "Map and Encapsulate", where traffic would be mapped and then tunnelled through the inter-domain core of the Internet. Another class being considered is sometimes known as "Identifier/Locator Split". This document relates to a proposal that is in the latter class of evolutionary approaches.

ILNP is, in essence, an end-to-end architecture: the functions required for ILNP are implemented in, and controlled by, only those end-systems that wish to use ILNP, as described in [RFC6740]. Other nodes, such as Site Border Routers (SBRs) need only support IP to allow operation of ILNP, e.g., an SBR should support IPv6 in order to enable end-systems to operate ILNPv6 within the site network for which an SBR provides a service [RFC6741].

However, some features of ILNP could be optimised, from an engineering perspective, by the use of an intermediate system (a router, security gateway or "middlebox") that modifies (rewrites) Locator values of transit ILNP packets. It would also perform other control functions for an entire site, as an administrative convenience, such as providing a centralised point of management for a site. For example, an SBR might manipulate the topological presence of the packet, providing an elegant solution to the provision of functions such as site (network) mobility for an entire end site [ABH09a].

This document discusses several such optional advanced deployment scenarios for ILNP. These typically use an ILNP-capable Site Border Router (SBR).

Nothing in this document is a requirement for any ILNP implementation or any ILNP deployment.

Readers are strongly advised to first read the ILNP Architecture Description [RFC6740], as this document uses the notation and terminology described or referenced in that document.

1.1. Document Roadmap

This document describes engineering and implementation considerations that are common to ILNP for both IPv4 and IPv6.

The ILNP architecture can have more than one engineering instantiation. For example, one can imagine a "clean-slate" engineering design based on the ILNP architecture. In separate documents, we describe two specific engineering instances of ILNP. The term "ILNPv6" refers precisely to an instance of ILNP that is based upon, and backwards compatible with, IPv6. The term "ILNPv4" refers precisely to an instance of ILNP that is based upon, and backwards compatible with, IPv4.

Many engineering aspects common to both ILNPv4 and ILNPv6 are described in [RFC6741]. A full engineering specification for either ILNPv6 or ILNPv4 is beyond the scope of this document.

Readers are referred to other related ILNP documents for details not described here:

- a) [RFC6740] is the main architectural description of ILNP, including the concept of operations.
- b) [RFC6741] describes engineering and implementation considerations that are common to both ILNPv4 and ILNPv6.
- c) [RFC6742] defines additional DNS resource records that support ILNP.
- d) [RFC6743] defines a new ICMPv6 Locator Update message used by an ILNP node to inform its correspondent nodes of any changes to its set of valid Locators.

- e) [RFC6744] defines a new IPv6 Nonce Destination Option used by ILNPv6 nodes (1) to indicate to ILNP correspondent nodes (by inclusion within the initial packets of an ILNP session) that the node is operating in the ILNP mode and (2) to prevent off-path attacks against ILNP ICMP messages. This Nonce is used, for example, with all ILNP ICMPv6 Locator Update messages that are exchanged among ILNP correspondent nodes.
- f) [RFC6745] defines a new ICMPv4 Locator Update message used by an ILNP node to inform its correspondent nodes of any changes to its set of valid Locators.
- g) [RFC6746] defines a new IPv4 Nonce Option used by ILNPv4 nodes to carry a security nonce to prevent off-path attacks against ILNP ICMP messages and also defines a new IPv4 Identifier Option used by ILNPv4 nodes.
- h) [RFC6747] describes extensions to Address Resolution Protocol (ARP) for use with ILNPv4.

1.2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Localised Numbering

Today, Network Address Translation (NAT) [RFC3022] is used for a number of purposes. Whilst one of the original intentions of NAT was to reduce the rate of use of global IPv4 addresses, through use of IPv4 private address space [RFC1918], NAT also offers to site administrators a convenient localised address management capability combined with a local-scope/private address space, for example, [RFC1918] for IPv4.

For IPv6, NAT would not necessarily be required to reduce the rate of IPv6 address depletion, because the availability of addresses is not such an issue as for IPv4. The IETF has standardised Unique Local IPv6 Unicast Addresses [RFC4193], which provide local-scope IPv6 unicast address space that can be used by end sites. However, localised address management, in a manner similar to that provided by

IPv4 NAT and private address space [RFC1918], is still desirable for IPv6 [RFC5902], even though there is debate about the efficacy of such an approach [RFC4864].

One of the major concerns that many have had with NAT is the loss of end-to-end transport-layer and network-layer session state invariance, which is still considered an important architectural principle by the IAB [RFC4924]. Nevertheless, the use of localised addressing remains in wide use and there is interest in its continued use in IPv6, e.g., proposals such as [RFC6296].

It is possible to have the benefits of NAT-like functions for ILNP without losing end-to-end state. Indeed, such a mechanism -- the use of Locator rewriting in ILNP -- forms the basis of many of the optional functions described in this document. In ILNP, we call this feature "localised numbering".


Recall, that a Locator value in ILNP has the same semantics as a routing prefix in IP: indeed, in ILNPv4 and ILNPv6 [RFC6741], routing prefixes from IPv4 and IPv6, respectively, are used as Locator values.

We note that a deployment using private/local numbering can also provide a convenient solution to centralised management of site multihoming and network mobility by deploying SBRs in this manner -- this is described below.

Please note that with this proposal, localised numbering (e.g., using the equivalent of IP NAT on the ILNP Locator bits) would work in harmony with multihoming, mobility (for individual hosts and whole networks), and IP Security (IPsec), plus the other advanced functions described in this document [BA11] [LABH06] [ABH07a] [ABH07b] [ABH08a] [ABH08b] [ABH09a] [ABH09b] [RAB09] [RB10] [ABH10] [BAK11].

2.1. Localised Locators

For ILNP, the NAT-like function can best be described by using a simple example, based on Figure 2.1.



CN = Correspondent Node
 H = Host
 L_1 = global Locator value
 L_L = local Locator value
 SBR = Site Border Router

Figure 2.1: A Simple Localised Numbering Example for ILNP

In this scenario, the SBR is allocated global locator value L_1 from the upstream provider. However, the SBR advertises internally a "local" Locator value L_L. By "local" we mean that the Locator value only has significance within the site network, and any packets that have L_L as a source Locator cannot be forwarded beyond the SBR with value L_L as the source Locator. In engineering terms, L_L would, for example, in ILNPv6, be an IPv6 prefix based on the assignments possible according to IPv6 Unique Local Addresses (ULAs) [RFC4193].

If we assume that H uses Identifier I_H, then it will use Identifier-Locator Vector (I-LV) [I_H, L_L], and that the correspondent node (CN) uses IL-V [I_CN, L_CN]. If we consider that H will send a UDP packet from its port P_H to CN's port P_CN, then H could send a UDP/ILNP packet with the tuple expression:

<UDP: I_H, I_CN, P_H, P_CN><ILNP: L_L, L_CN> --- (1a)

When this packet reaches the SBR, it knows that L_L is a local Locator value and so rewrites the source Locator on the egress packet to L_1 and forwards that out onto its external-facing interface. The value L_1 is a global prefix, which allows the packet to be routed globally:

<UDP: I_H, I_CN, P_H, P_CN><ILNP: L_1, L_CN> --- (1b)

This packet reaches CN using normal routing based on the Locator value L_1, as it is a routing prefix.

Note that from expressions (1a) and (1b), the end-to-end state (in the UDP tuple) remains unchanged -- end-to-end state invariance is honoured, for UDP. CN would send a UDP packet to H as:

```
<UDP: I_CN, I_H, P_CN, P_H><ILNP: L_CN, L_1>          --- (2a)
```

and the SBR would rewrite the Locator value on the ingress packet before forwarding the packet on its internal interface:

```
<UDP: I_CN, I_H, P_CN, P_H><ILNP: L_CN, L_L>          --- (2b)
```

Again, this preserves the end-to-end transport-layer session state invariance.

As the Locator values are not used in the transport-layer pseudo-header for ILNP [RFC6741], the checksum would not have to be rewritten. That is, the Locator rewriting function is stateless and has low overhead.

(A discussion on the generation of Identifier values for initial use is presented in [RFC6741].)

2.2. Mixed Local/Global Numbering

It is possible for the SBR to advertise both L_1 and L_L within the site, and for hosts within the site to have IL-Vs using both L_1 and L_L. For example, host H may have IL-Vs [I_H, L_1] and [I_H, L_L]. The configuration and use of such a mechanism can be controlled through local policy.

2.3. Dealing with Internal Subnets with Locator Rewriting

Where the site network uses subnets, packets will need to be routed correctly, internally. That is, the site network may have several internal Locator values, e.g., L_La, L_Lb, and L_Lc. When an ingress packet has I-LV [I_H, L_1], it is expected that the SBR is capable of identifying the correct internal network for I_H, and so the correct Locator value to rewrite for the ingress packet. This is not obvious as the I value and the L value are not related in any way.

There are numerous ways the SBR could facilitate the correct lookup of the internal Locator value. This document does not prescribe any specific method. Of course, we do not preclude mappings directly from Identifier values to internal Locator values.

Of course, such a "flat" mapping (between Identifier values and Locators) would serve, but maintaining such a mapping would be impractical for a large site. So, we propose the following solution.

Consider that the Locator value, L_x consists of two parts, L_{pp} and L_{ss} , where L_{pp} is a network prefix and L_{ss} is a subnet selector. Also, consider that this structure is true for both the local identifier, L_L , as well as the global Identifier, L_1 . Then, an SBR need only know the mapping from the values of L_{ss} as visible in L_1 and the values of L_{ss} used locally.

Such a mapping could be mechanical, e.g., the L_{ss} part of L_L and L_1 are the same and it is only the L_{pp} part that is different. Where this is not desirable (e.g., for obfuscation of interior topology), an administrator would need to configure a suitable mapping policy in the SBR, which could be realised as a simple lookup table. Note that with such a policy, the L_{pp} for L_L and L_1 do not need to be of the same size.

From a practical perspective, this is possible for both ILNPv6 [RFC6177] and ILNPv4 [RFC4632]. For ILNPv6, recall that the Locator value is encoded to be syntactically similar to an IPv6 address prefix, as shown in Figure 2.2, taken from [RFC6741].

```

/* IPv6 */
| 3 |      45 bits      | 16 bits |      64 bits      |
+-----+-----+-----+
| 001|global routing prefix| subnet ID | Interface Identifier |
+-----+-----+-----+
/* ILNPv6 */
|      64 bits      |      64 bits      |
+-----+-----+
|      Locator (L64)      | Node Identifier (NID) |
+-----+-----+
+<----- L_pp ----->+<- L_ss -->+

```

L_{pp} = Locator prefix part (assigned IPv6 prefix)


L_{ss} = Locator subnet selector (locally managed subnet ID)

Figure 2.2: IPv6 Address format [RFC3587] as used in ILNPv6, showing how subnets can be identified.

Note that the subnet ID forms part of the Locator value. Note also that [RFC6177] allows the global routing prefix to be more than 45 bits, and for the subnet ID to be smaller, but still preserving the 64-bit size of the Locator overall.

For ILNPv4, the L_{pp} value overall is an IPv4 routing prefix, which is typically less than 32 bits. However, the ILNPv4 Locator value is carried in the 32-bit IP Address space, so the bits not used for the

routing prefix could be used for L_{ss} , e.g., for a /24 IPv4 prefix, the situation would be as shown in Figure 2.3, and L_{ss} could use any of the remaining 8-bits as required.



L_{pp} = Locator prefix (assigned IPv4 prefix)

L_{ss} = Locator subnet selector (locally managed subnet ID)

Figure 2.3: IPv4 address format for /24 IPv4 prefix, as used in ILNPv4, showing how subnets can be identified.

As an example, for the case where the interior topology is not obfuscated, an interior "engineering" node might have an LP record pointing to eng.example.com and eng.example.com might have L32/L64 records for a specific subnet inside the site. Meanwhile, an interior "operations" node might have an LP record pointing at "ops.example.com" that might have different L32/L64 records for that specific subnet within the site. That is, eng.example.com might have Locator value $L_{pp_1}:L_{ss_1}$ and ops.example.com might have Locator value $L_{pp_1}:L_{ss_2}$. However, just as for IPv6 or IPv4 routing today, the routing for the site would only need to use L_{pp_1} , which is a routing prefix in either IPv6 (for ILNPv6) or IPv4 (for ILNPv4).

2.4. Localised Name Resolution with DNS

To support private numbering with IPv4 and IPv6 today, some sites use a split-horizon DNS service for the site [appDNS].

If a site using localised numbering chooses to deploy a split-horizon DNS server, then the DNS server would return the global-scope Locator(s) (L_1 in our example above) of the SBR to DNS clients outside the site, and would advertise the local-scope Locator(s) (L_L in our example above) specific to that internal node to DNS clients inside the site. Such deployments of split-horizon DNS servers are not unusual in the IPv4 Internet today. If an internal node (e.g., portable computer) moves outside the site, it would follow the normal ILNP methods to update its authoritative DNS server with its current Locator set. In this deployment model, the authoritative DNS server for that mobile device will be either the split-horizon DNS server itself or the master DNS server providing data to the split-horizon DNS server.

If a site using localised numbering chooses not to deploy a split-horizon DNS server, then each internal node would advertise the global-scope Locator(s) of the site border routers in its respective DNS entries. To deliver packets from one internal node to another internal node, the site would choose to use either Layer 2 bridging (e.g., IEEE Spanning Tree or IEEE Rapid Spanning Tree [IEEE04], or a link-state Layer 2 algorithm such as the IETF TRILL group or IEEE 802.1 are developing), or the interior routers would forward packets up to the nearest site border router, which in turn would then rewrite the Locators to appropriate local-scope values, and forward the packet towards the interior destination node.

Alternately, for sites using localised numbering but not deploying a split-horizon DNS server, the DNS server could return all global-scope and local-scope Locators to all queriers, and assume that nodes would use normal, local address/route selection criteria to choose the best Locator to use to reach a given remote node ([RFC3484] for older IPv6 nodes, [RFC6724] for newer IPv6 nodes). Hosts within the same site as the correspondent node would only have a ULA configured; hence, they would select the ULA destination Locator for the correspondent (L_L in our example). Hosts outside the site would not have the same ULA configured (L_CN for the CN in our example).

However, ILNP allows use of Locator Preference values [RFC6742] [RFC6743]. These values would indicate explicitly the relative preference value given to Locator values and so result in the selection of the appropriate Locator (and therefore interface) to use for the transmission of an outgoing packet with respect to the value to be inserted into the IPv6 Source Address field (see Section 3 of [RFC6741]). A similar argument, with respect to use of Locator preference values, applies to the value to be inserted into the IPv6 Destination Address field. Certainly, by using appropriate Preference values for a host with multiple Locator values, it would be possible to emulate some level of resemblance to the address selection rules in [RFC3484] and [RFC6724], and this could be controlled via DNS entries for ILNP nodes, for example.

Indeed, with appropriate use of localised or site-wide policy, and appropriate mechanisms in the devices (e.g. in end hosts operating systems or in Site Border Routers), Preference values for Locator values within the DNS could be used for allowing options for multi-homed transport sessions and/or site-controlled traffic engineering [ABH09a]. However, the details for this are left for further study, and overall, the rules defined in [RFC3484] and [RFC6724] cannot be applied directly to ILNPv6 nodes.

Note that for split-horizon operation, there needs to be a DNS management policy for mobile hosts, as when such hosts are away from their "home" network, they will need to update DNS entries so that the global-scope Locator(s) only is (are) used, and these are consistent with the current topological position of the mobile host. Such updates would need to be done using Secure Dynamic DNS Update.

For an ILNP mobile network using LP records, there are likely to separate LP records for internal and external use.

2.5. Use of mDNS

Multicast DNS (mDNS) [mDNS11] is popularly used in many end-system OSs today, especially desktop OSs (such as Windows, Mac OS X and Linux). It is used for localised name resolution using names with a ".local" suffix, for both IPv4 and IPv6. This protocol would need to be modified so that when an ILNP-capable node advertises its ".local" name, another ILNP-capable node would be able to see that it is an ILNP-capable, but other, non-ILNP nodes would not be perturbed in operation. The details of a mechanism for using mDNS to enable such a feature are not defined here.

2.6. Site Network Name in DNS

In this scenario, if H expects incoming ILNP session requests, for example, then remote nodes normally will need to look up appropriate Identifier and Locator information in the DNS. Just as for IP, and as already described in [RFC6740], a Fully Qualified Domain Name (FQDN) lookup for H should resolve to the correct NID and L32/L64 records. If there are many hosts like H that need to keep DNS records (for any reason, including to allow incoming ILNP session requests), then, potentially, there are many such DNS resource records.

As an optimisation, the network as a whole may be configured with one or more L32 and L64 records (to store the value L_1 from our example) that are resolved from an FQDN. At the same time, individual hosts now have an FQDN that returns one or more LP record entries [RFC6742] as well as NID records. The LP record points to the L32 or L64 records for the site. A multihomed site normally will have at least one L32 or L64 record for each distinct uplink (i.e., link from a Site Border Router towards the global Internet), because ILNP uses provider-aggregatable addressing.

More than one L32 or L64 will be required if multiple Locator values are in use. For example, if an ILNPv6 site has multiple links for multihoming, it will use one L64 record for each Locator value it is using on each link.

2.7. Site Interior Topology Obfuscation

In some situations, it can be desirable to obfuscate the details of the interior topology of an end site. Alternately, in some situations, local site policy requires that local-scope routing prefixes be used within the local site. ILNP can provide these capabilities through the ILNP local addressing capability described here, under the control of the SBR.

As described in Section 2.3 above, locator rewriting can be used to hide the internal structure of the network with respect to the subnetting arrangement of the site network. Specifically, the procedure described in Section 2.3 would be followed, with the following additional modification of the use of Locator values:

- (1) Only the aggregated Locator value, i.e., L_{pp} , is advertised outside the site (e.g., in an L32 or L64 record), and L_{ss} is zeroed in that advertisement.
- (2) The SBR needs to maintain a mapping table to restore the interior topology information for received packets, for example, by using a mapping table from I values to either L_{ss} values or internal Locator values.
- (3) The SBR needs to zero the L_{ss} values for all Source Locators of egress packets, as well as perform a Locator rewriting that affects the L_{pp} bits of the Locator value.

Of course, this only obscures the interior topology of the site, not the exterior connectivity of the site. In order for the site to be reachable from the global Internet, the site's DNS entries need to advertise Locator values for the site to the global Internet (e.g., in L32, L64 records).

2.8. Other SBR Considerations

For backwards compatibility, for ILNP, the ICMP checksum is always calculated identically as for IPv6 or IPv4. For ILNPv6, this means that the SBR need not be aware if ILNPv6 is operating as described in [RFC6740] and [RFC6741]. For ILNPv4, again, the SBR need not be aware of the operation if ILNPv4 is operating as it will not need to inspect the extension header carrying the I value.

In order to support communication between two internal nodes that happen to be using global-scope addresses (for whatever reason), the SBR MUST support the "hair pinning" behaviour commonly used in existing NAT/NAPT devices. (This behaviour is described in Section 6 of RFC 4787 [RFC4787].)

In the near-term, a more common deployment scenario will be to deploy ILNP incrementally, with some ordinary classic IP traffic still existing. In this case, the SBR should maintain flow state that contains a flag for each flow indicating whether or not that flow is using ILNP. If that flag indicated ILNP were enabled for a given flow, and ILNP local numbering were also enabled, then the SBR would know that it should perform the simpler ILNP Locator rewriting mapping. If that flag indicated ILNP were not enabled for a given flow and IP NAT or IP NAPT were also enabled, then the SBR would know that it should perform the more complex NAT/NAPT translation (e.g., including TCP or UDP checksum recalculation).

NOTE: Existing commercial security-aware routers (e.g., Juniper SRX routers) already can maintain flow state for millions of concurrent IP flows. This feature would add one flag to each flow's state, so this approach is believed scalable today using existing commercial technology.

Those applications that do not use IP Address values in application state or configuration data are considered to be "well behaved". For well-behaved applications, no further enhancements are required. Where application-layer protocols are not well behaved, for example, the File Transfer Protocol (FTP), then the SBR might need to perform additional stateful processing -- just as NAT and NAPT equipment needs to do today for FTP. See the description in Section 7.6 of [RFC6741].

When the SBR rewrites a Locator in an ILNP packet, that obscures information about how well a particular path is working between the sender and the receiver of that ILNP packet. So, the SBR that rewrites Locator values needs to include mechanisms to ensure that any packet with a new Destination Locator will travel along a valid path to the intended destination node. For ILNPv4, the path liveness will be no worse than IPv4, and mechanisms already in use for IPv4 can be reused. For ILNPv6, the path liveness will be no worse than for IPv6, and mechanisms already in use for IPv6 can be reused.


In the future, the Border Router Discovery Protocol (BRDP) also might be used in some deployments to indicate which routing prefixes are currently valid and which site border routers currently have a working uplink [BRDP11].

3. An Alternative for Site Multihoming

The ILNP Architectural Description [RFC6740] describes the basic approach to enabling Site Multihoming (S-MH) with ILNP. However, as an option, it is possible to leave the control of S-MH to an ILNP-enabled SBR. This alternative is based on the use of the Localised Numbering function described in Section 2 of this document.

3.1. Site Multihoming (S-MH) Connectivity Using an SBR

The approach to Site Multihoming (S-MH) using an SBR is best illustrated through an example, as shown in Figure 3.1.



- CN = Correspondent Node
- H = Host
- L_1 = global Locator value 1
- L_2 = global Locator value 2
- L_L = local Locator value
- SBR = Site Border Router
- sbrN = interface N on SBR

Figure 3.1: Alternative Site Multihoming Example with an SBR

The situation here is similar to the localised numbering example, except that the SBR now has two external links, with using Locator value L_1 and another using Locator value L_2. These could, e.g., for ILNPv6, be separate, Provider Aggregated (PA) IPv6 prefixes from two different ISPs. H has IL-V [I_H, L_L], and will forward a packet to CN as given in expression (1a). However, when the packet reaches the SBR, local policy will decide whether the packet is forwarded on the link sbr1 using L_1 or on sbr2 using L_2. Of course, the correct Locator value will be rewritten into the egress packet in place of L_L.

If only local numbering is being used, then the SBR need never advertise any global Locator values. However, it could do, as described in Section 2.2.

3.2. Dealing with Link/Connectivity Changes

One of the key uses for multihoming is providing resilience to link failure. If either link breaks, then the SBR can manage the change in connectivity locally. For example, assume SBR has been configured to use sbr1 for all traffic, and sbr2 only as backup link. So, SBR directs packets from H to communicate with CN using sbr1, and CN will receive packets as in expression (1b) and respond with packets as in expression (2a).

However, if sbr1 goes down then SBR will move the communication to interface sbr2. As H is not aware of the actions of the SBR, the SBR must maintain some state about IL-V "pairs" in order to hand off the connectivity from sbr1 to sbr2. So, when moving the communication to sbr2, the SBR would firstly send a Locator Update (LU) message [RFC6745] [RFC6743], to CN informing it that L₂ is now the valid Locator for the communication. This operation would not be visible to H, although there might be some disruption to transmission, e.g., packets being sent from CN to H that are in flight when sbr1 goes down may be lost. The SBR might also need to update DNS entries (see Section 3.3). Since ILNP requires that all Locator Update messages be authenticated by the ILNP Nonce, the SBR will need to include the appropriate Nonce values as part of its cache of information about ILNP sessions traversing the SBR. (NOTE: Since commercial security gateways available as of this writing reportedly can handle full stateful packet inspection for millions of flows at multi-gigabit speeds, it should be practical for such devices to cache the ILNP flow information, including Nonce values.)

This approach has some efficiency gains over the approach for multihoming described in [RFC6740], where each hosts manages its own connectivity.

If sbr1 was to be reinstated, now with Locator value L₃, then local policy would determine if the communication should be moved back to sbr1, with appropriate additional actions, such as transmission of LU messages with the new Locator values and also the updates to DNS.

Note that in such movement of an ILNP session across interfaces at the SBR, only Locator values in ILNP packets are changed. As already noted in [RFC6740], end-to-end transport-layer session state invariance is maintained.

3.3. SBR Updates to DNS

When the SBR manages connectivity as described above, the internal hosts, such as H, are not necessarily aware of any connectivity changes. Indeed, there is certainly no requirement for them to be aware. So, if H was a server expecting incoming connections, the SBR must update the relevant DNS entries when the site connectivity changes.

There are two possibilities: each host could have its own L32 or L64 records; or the site might use a combination of LP and L32/L64 records (see Section 2.4). Either way, the SBR would need to update the relevant DNS entries. For our example, with ILNPv6 and LP records in use, the SBR would need to manage two L64 records (one for each uplink) that would resolve from a FQDN, for example, site.example.com. Meanwhile, individual hosts, such as H, have an FQDN that resolves to an NID value and an LP record that would contain the value site.example.com, which then would be used to look up the two L64 records.

If the SBR is multihomed, as in Figure 3.1, then it will have (at least) two Locator values, one for each link, and local policy will need to be used to determine how preference values are applied in the relevant L32 and L64 records.

3.4. DNS TTL Values for L32 and L64 Records

Imagine that in the scenario described above, there was a link failure that resulted in sbr1 going down and sbr2 was used. Existing ILNP sessions in progress would move to sbr2 as described above. However, new incoming ILNP sessions to the site would need to know to use L_2 and not L_1. L_1 and L_2 would be stored in DNS records (e.g., L32 for ILNPv4 or L64 for ILNPv6). If a remote host has already resolved from DNS that L_1 is the correct Locator for sending packets to the site, then that host might be holding stale information.


DNS allows values returned to be aged using Time-To-Live (TTL), which is specified in the time unit of seconds. So that remote nodes do not hold on to stale values from DNS, the L64 records for our site should have low TTL values. An appropriate value must be considered carefully. For example, let us assume that the site administrator knows that when sbr1 fails, it takes 20 seconds to failover to sbr2. Then, 20 s would seem to be an appropriate time to use for the TTL value of an L64 for the site: if a remote node had just resolved the value L_1 for the site, and the link to sbr1 went down, that remote node would not hold the stale value of L_1 for any longer than it takes the site to failover to sbr2 and use L_2.

Our studies for a university school site network show that low TTL values, as low as zero, are feasible for operational use [BA11].

NOTE: From 01 November 2010, the site network of the School of Computer Science, University of St Andrews, UK, has been running operational DNS with DNS A records that have TTL of zero. At the time of writing of this document (November 2012), a zero DNS TTL was still in use at the school.

3.5. Multiple SBRs

For site multihoming, with multiple SBRs, a situation may be as follows (see also Section 5.3.1 in [RFC6740]).



- CP = coordination protocol
- L_1 = global Locator value 1
- L_2 = global Locator value 2
- SBR_A = Site Border Router A
- SBR_B = Site Border Router P

Figure 3.2: A Dual-Router Multihoming Scenario for ILNP

The use of two physical routers provides an extra level of resilience compared to the scenario of Figure 3.1. The coordination protocol (CP) between the two routers keeps their actions in synchronisation according to whatever management policy is in place for the site

network. Such functions are available today in some commercial network security products. Note that, logically, there is little difference between Figures 5.1 and 3.2, but with two distinct routers in Figure 3.2, the interaction using CP is required. Of course, it is also possible to have multiple interfaces in each router and more than two routers.

4. An Alternative for Site (Network) Mobility

The ILNP Architectural Description [RFC6740] describes the basic approach to enabling site (network) mobility with ILNP. However, as an option, it is possible to leave the control of site mobility to an ILNP-enabled SBR by exploiting the alternative site multihoming feature described in Section 3 of this document.

Again, as described in [RFC6740], we exploit the duality between mobility and multihoming for ILNP.

4.1. Site (Network) Mobility

Let us consider the mobile network in Figure 4.2, which is taken from [RFC6740].




Figure 4.1a: ILNP Mobile Network before Handover




Figure 4.1b: ILNP Mobile Network during Handover




Figure 4.1c: ILNP Mobile Network after Handover

H = host
 L_1 = global Locator value 1
 L_2 = global Locator value 2
 L_L = local Locator value
 raN = radio interface N
 SBR = Site Border Router

Figure 4.1: An Alternative Mobile Network Scenario with an SBR

We assume that the site (network) is mobile, and the SBR has two radio interfaces, ra1 and ra2. In the figure, ISP_1 and ISP_2 are separate, radio-based service providers, accessible via interfaces ra1 and ra2.

While the SBR makes the transition from using a single link (Figure 4.1a) to the handover overlap on both links (Figure 4.1b), to only using a single link again (Figure 4.1c), the host H continues to use only Locator value L_L, as already described for Site Multihoming (S-MH). During this time the actions taken by the SBR are the same as already described in [RFC6740], except that the SBR:

- a) also performs that ILNP localised numbering function described in Section 2.
- b) does not need to advertise L_1 and L_2 internally if only local numbering is being used.

As for the case of S-MH above, H need not be aware of the change in connectivity for the SBR if it is only using local numbering, and the SBR would send LU messages for H (for any correspondent nodes, not shown in Figure 4.1), and would update DNS entries as required.

The difference to the S-MH scenario described earlier in this document is that in the situation of Figure 4.1b, the SBR can opt to use soft handover as previously described in [RFC6740].

Again, there is an efficiency gain compared to the situation described in [RFC6740]: the SBR provides a convenient point at which to centrally manage the movement of the site as a whole. Note that in Figure 4.1b, the site is multihomed.

As for S-MH, L_1 and L_2 could be advertised internally, as a local policy decision, for those hosts that require direct control of their connectivity.

Note that for handover, immediate handover will have a similar behaviour to a link outage as described for S-MH. However, as ILNP allows soft-handover, during the handover period, this should help to reduce (perhaps even remove) packet loss.

4.2. SBR Updates to DNS

As for S-MH, a similar discussion to Section 3.3 applies for mobile networks with respect to the updates to DNS. As a mobile network is likely to have more frequent changes to its connectivity than a multihomed network would due to connectivity changes, the use of LP DNS records is likely to be particularly advantageous here.

4.3. DNS TTL Values for L32 and L64 Records

As for S-MH, a similar discussion to Section 3.4 applies for mobile networks with respect to the TTL of L32 and/or L64 records that are used for the name of the mobile network. In the case of the mobile network, it makes sense for the TTL to be aligned to the time for handover.


5. Traffic Engineering Options

The use of Locator rewriting provides some simple yet useful options for traffic engineering (TE) controlled from the edge-site via the SBR, requiring no cooperation from the service provider other than the provision of basic connectivity services, e.g., physical connectivity, allocation of IP Address prefixes and packet forwarding. This does not preclude other TE options that are already in use, such as use of MPLS, but we choose to highlight here the specific options available and controllable solely through the use of ILNP.

When a site network is multihomed, we have seen that the use of the Locator rewriting function permits the SBR to have packet-by-packet control when forwarding on external links. Various configuration and policies could be applied at the SBR in order to control the egress and ingress traffic to the site network.

5.1. Load Balancing

Let us consider Figure 5.1, and assume ILNP local numbering is in use; that H1, H2, and H3 use, respectively, Identifier values, I_1, I_2 and I_3; and all of them use Locator value L_L.



- HN = host N
- L_1 = global Locator value 1
- L_2 = global Locator value 2
- L_L = local Locator value
- SBR = Site Border Router
- sbrN = interface N on sbr

Figure 5.1: A Site Multihoming Scenario for Traffic Control

The SBR could be configured, subject to local policy, to try to control load across the external links. For example, it could be configured initially with the following mappings:

- srcI=I_1, sbr1 --- (3a)
- srcI=I_2, sbr2 --- (3b)
- srcI=I_3, sbr1 --- (3c)

These mappings direct packets matching course Identifier values to particular outgoing interfaces. As load changes, these mappings could be changed. For example, expression (3c) could be changed to:

- srcI=I_3, sbr2 --- (4)

and the SBR would need to send LU message to the correspondents of H3 (sbr to uses L_2 while sbr1 uses L_1). The egress connectivity is totally within control of the SBR under administrative policy, as already seen in the descriptions of multihoming and mobility in this document.

Of course, more complex policies are possible, based on:

- whether ILNP sessions are incoming or outgoing
- time of day
- internal subnets

and any number of criteria already in use for control of traffic.

In expressions (3a,b,c) above, source I values are used. However:

- destination I values could be used
- source or destination L values could be used
- mappings could be to L values, not to specific interfaces

and, again, any number of criteria could be used to manipulate the packet path, based on filtering of values in header fields and local policy.

With ILNP, hosts do not need to be aware of the operation of the SBR in this manner.

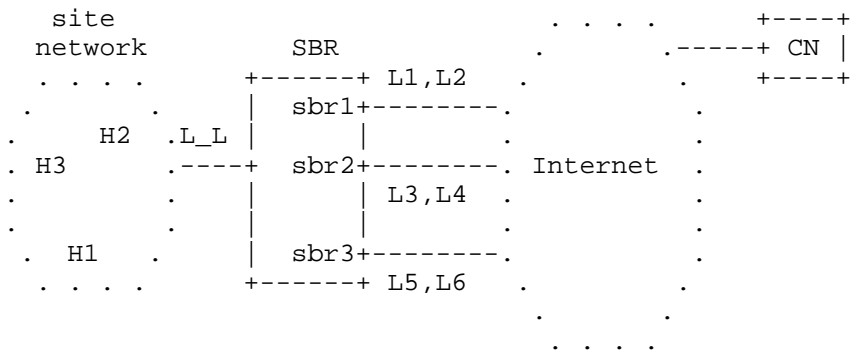
Note, again, that in this scenario, there is nothing to prevent SBR from also advertising L₁ and L₂ into the site network. If required, administrative controls could be used to enable selective hosts in the site network to use L₁ and L₂ directly as described in [RFC6740].

5.2. Control of Egress Traffic Paths

Extending the scenario for load-balancing described above, it is also possible for the ILNP-capable SBR to direct traffic along specific network paths based on the use of different L values, i.e., by using multiple prefixes assigned from upstream providers.

Of course, as previously discussed, these prefixes can be Provider Aggregated (PA) and need not be Provider Independent (PI).

Let us consider Figure 5.2 and assume ILNP local numbering is in use; that H1, H2 and H3 use, respectively, Identifier values, I₁, I₂, and I₃; and all of them use Locator value L_L. Let us also assume that the node CN uses IL-V [I_{CN}, L_{CN}].



CN = correspondent node
HN = host N
LN = global Locator value N
L_L = local Locator value
SBR = Site Border Router
sbrN = interface N on sbr

Figure 5.2: A Site Multihoming Scenario for Traffic Control

Here, many configurations are possible. For example, for egress traffic:

```

srcI=I_2, L2          --- (5a)
srcI=I_3, L3          --- (5b)
dstI=I_CN, L6         --- (5c)
srcI=I_1 dstI=I_CN, L1 --- (5d)

```

Expression (5a) maps all egress packets from H2 to have their source Locator value rewritten to L2 (and implicitly to use interface sbr1). Expression (5b) maps all egress packets from H3 to have their source Locator value rewritten to L3 (and implicitly to use interface sbr2). Expression (5c) directs any traffic to CN to use Locator value L6 as the source Locator (and implicitly to use interface sbr3), and may override (5a) and (5b), subject to local policy, when packets to CN are from H2 or H3.

Meanwhile, in expression (5d), we see a further, more specific rule, in that packets from H1 destined to CN should use Locator value L1 (and implicitly to use interface sbr1).

Note the implicit bindings to interfaces in expressions (5a,b,c,d), compared to the explicit bindings in expressions (3a,b,c). ILNP only requires that the Locator values are correctly rewritten and packets forwarded in conformance with the routing already configured for the Locator values.

Of course, these rules can be changed dynamically at the SBR, and the SBR will migrate ILNP sessions across Locator values, as already described above for mobility.

6. ILNP in Datacentres

As ILNP has first class support for mobility and multihoming, and supports flexible options for localised addressing, there is great potential for it to be used in datacentre scenarios. Further details of possibilities are in [BA12], with a summary presented here.

There are several scenarios that could be beneficial to datacentres, in order to provide functions such as load balancing, resilience and fault tolerance, and resource management:

- Same datacentre, internal Virtual Machine (VM) mobility: This could be beneficial in load balancing, dynamically, where load changes are taking place. The remote user does not see the VM has moved.
- Different datacentres, transparent mobility: This is where the datacentre resources may be geographically distributed, but the geographical movement is transparent to the remote user.
- Different datacentres, mobility is visible: This is where the datacentre resources may be geographically distributed, but the geographical movement is visible to the remote user.


These are three situations that may be supported by ILNP, but they are not the only ones: we provide these here as examples, and they are not intended to be prescriptive. The intention is only to show the flexibility that is possible through the use of ILNP.

This section describes some Virtual Machine (VM) mobility capabilities that are possible with ILNP. Depending on the internal details and virtualisation model provided by a VM platform, it might be sufficient for the guest operating system to support ILNP. In some cases, again depending on the internal details and virtualisation model provided by a VM platform, the VM platform itself also might need to include support for ILNP.

Details of how a particular VM platform works, and which virtualisation model(s) a VM platform supports, are beyond the scope of this document. Internal implementation details of VM platform support for ILNP are also beyond the scope of this document, just as internal implementation details for any other networked system supporting ILNP are beyond the scope of this document.

6.1. Virtual Image Mobility within a Single Datacentre

Let us consider first the scenario of Figure 6.1, noting its similarity to Figure 2.1 for use of localised numbering.



CN = Correspondent Node
 V = Virtual machine image
 Hx = Host x
 L_1 = global Locator value
 L_L = local Locator value
 SBR = Site Border Router

Figure 6.1: A Simple Virtual Image Mobility Example for ILNP

L_L is a Locator value used for the ILNP hosts H1 and H2. Here, the "V*H1" signifies that the virtual machine image V is currently resident on H1. Let us assume that V has Identifier I_V. Note that as H1 and H2 have the same Locator value (L_1), as far as CN is concerned, it does not matter if V is resident on H1 or H2, all transport packets between V and CN will have the same signature as far as CN is concerned, e.g., for a UDP flow (in analogy to (1a)):

<UDP: I_V, I_CN, P_V, P_CN><ILNP: L_1, L_CN> --- (6a)

Now, if V was to migrate to H2, the migration would be an issue purely local to the site network, and the end-to-end integrity of the transport flow would be maintained.

Of course, there are practical operating systems issues in enabling such a migration locally, but products exist today that could be modified and made ILNP-aware in order to enable such VM image mobility.

Note that for convenience, above, we have used localised numbering for ILNP, but if local Locator values were not used and the whole site simply used L_1, the principle would be the same.

When in initial communication, the packets have the signature is given in expression (6a). When V moves to H2, it now uses Locator value L_L2, but all communication between V and CN is still routed via SBR1. So, the remote CN still sees that same packet signature as given in expression (6a). L_L1 and L_L2 are, effectively, two internal (private) subnetworks, and are not visible to CN.

However, SBR2 and SBR1 must coordinate so that any further communication to V via SBR1 is routed across the inter-router link. Again, there are commercial products today that could be adapted to manage such shared state.

6.3. Virtual Image Mobility between Datacentres - Visible

Clearly, in the scenario of the section above, once V has moved to site network 2, it may be beneficial, for a number of reasons, for communication to V to be routed via SBR2 rather than SBR1.

When V moves from site network 1 to site network 2, this visibility of mobility could be by V sending ILNP Locator Update messages to the CN during the mobility process. Also, V would update any relevant ILNP DNS records, such as L64 records, for new ILNP session requests to be routed via SBR2.

Indeed, let us now consider again Figure 6.2, and assume now that Local locators L_L1 and L_L2 are not in use on either site network, and each site networks uses its own global Locator value, L_1 and L_2, respectively, internally. In that case, the packet flow signature for V when it is in site network 1 as viewed from CN is, again as given in expression (6a). However, when V moves to site network 2, it would simply use L_2 as its new Locator, send Locator Update messages to CN as would a normal mobile node for ILNP, and complete its migration to H2. Then, CN would see the packet signatures as in expression (6b).

<UDP: I_V, I_CN, P_V, P_CN><ILNP: L_2, L_CN> --- (6b)

In this case, no "special" inter-router link is required for mobility -- the normal Internet connectivity between SBR1 and SBR2 would suffice. However, it is quite likely that some sort of tunnelled link would still be desirable to offer protection of the VM image as it migrates.

6.4. ILNP Capability in the Remote Host for VM Image Mobility

For the remote host -- the CN -- the availability of ILNP would be beneficial. However, for the first two scenarios listed above, as the packet signature of the transport flows remains fixed from the


viewpoint of the CN, it seems possible that the benefits of ILNP VM mobility could be used for datacentres even while CNS remain as normal IP hosts. Of course, a major caveat here is that the application level protocols should be "well behaved": that is, the application protocol or configuration should not rely on the use of IP Addresses.

7. Location Privacy

Extending the Locator rewriting paradigm, it is possible to also enable Location privacy for ILNP by a modified version of the "onion routing" paradigm that is used for Tor [DMS04] [RSG98].

7.1. Locator Rewriting Relay (LRR)

To enable this function, we use a middlebox that we call the Locator Rewriting Relay. The function of this unit is described by the use of Figure 7.1.



LRR = Locator Rewriting Relay

Figure 7.1: Locator Rewriting Relay (LRR) Example

The operation of the LRR is conceptually very simple. We assume that the LRR first has mappings as given in expressions (7b) and (7c) (see next subsection). Expression (7b) says that for packets with src IL-V [I_H, L_1], the packet's source Locator value should be rewritten to value L_X and then forwarded. Expression (7c) has the complimentary mapping for packets with destination IL-V [I_H, L_1] (for the reverse direction).

Expression (6a) is a UDP/ILNP packet as might be sent in Figure 2.1 from H to CN. However, instead of going directly to L_CN, the packet with destination Locator L_1 goes to a LRR. Expression (7d) is the result of the mapping of packet (7a) using expression (7b).

Note that it is entirely possible that the packet of expression (7d) then is processed by another LRR for source Locator value L_X. Effectively, this creates an LRR path for the packet, as an overlay path on top of the normal IP routing.

In this way, there is a level of protection, without the need for cryptographic techniques, for the (topological) Location of the packet. Of course, an extremely well-resourced adversary could, potentially, backtrack the LRR path, but, depending on the LRR overlay path that is created, could be very difficult to trace in reality. For example, the mechanism will protect against off-path attacks, but where the threat regime includes the potential for on-path attacks, cryptographically protected tunnels between H and LRR might be required.

Again, as the Locator value is not part of the end-to-end state, this mechanism is very general and has a low overhead.

7.2. Options for Installing LRR Packet Forwarding State

There are many options for managing the "network" of LRRs that could be in place if such a system was used on a large scale, including the setting up and removal of LRR state for packet relaying, as for expressions (7b) and (7c). We consider this function to be outside the scope of these ILNP specifications, but note that there are many existing mechanisms that could be modified for use, and also many possibilities for new mechanisms that would be specific to the use of ILNP LRRs.

(Note also that the control/management communication with the LRR does not need to use ILNP: IPv4 or IPv6 could be used.)

The host, H, by itself could install the required state, assuming it was aware of suitable information to contact the LRR. The first packet in an ILNP session might contain a header option called a Locator Redirection Option (LRO). The LRO would contain the Locator value that should be rewritten into the source Locator of the packet. When a LRR receives such a packet, it would install the required state. Such a mechanism could be soft-state, requiring periodic use of the LRO in order to maintain the state in the LRR. The LRO could also be delivered using an ICMP ECHO packet sent from H to the LRR, periodically, again to maintain a soft-state update.

It would, of course, be prudent to protect the LRR state control packets with some sort of authentication token, to prevent an adversary from easily installing false LRR state and causing packets

from H or its correspondent to be subject to man-in-the-middle attacks, or black-holing. Again, such attacks are not specific to ILNP or new to ILNP.

It would also be possible to use proprietary application level protocols, with strong authentication for the control of the LRR state. For example, an application level protocol based on XMPP (<http://xmpp.org/>) operating over SSL.

Above, we have offered very brief and incomplete descriptions of some possibilities, and we do not necessarily mandate any one of them: they serve only as examples.

8. Identity Privacy

For the sake of completeness, and in complement to Section 6, it should be noted that ILNP can use either cryptographically verifiable Identifier values, or use Identifier values that provide a level of anonymity to protect a user's privacy. More details are given in Sections 2 and 11 of [RFC6741].

9. Security Considerations

The relevant security considerations to this document are the same as for the main ILNP Architectural Description [RFC6740]. The one additional point to note is that this document describes ILNP capability in the SBR and so those adversaries wishing to subvert the operation of ILNP specifically, have a target that would, potentially, disable an entire site. However, this is not an attack vector that is specific to ILNP: today, disruption of an IPv4 or IPv6 SBR would have the same impact.

The security considerations for Section 7 (Location Privacy) are already documented in [DMS04] and [RSG98]. One possibility is that the LRR mechanism itself could be used by an adversary to launch an attack and hide his own (topological) Location, for example. This is already possible for IPv4 and IPv4 with a Tor-like system today, so is not new to ILNP.

10. References

10.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, August 2006.
- [RFC4787] Audet, F., Ed., and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.
- [RFC4924] Aboba, B., Ed., and E. Davies, "Reflections on Internet Transparency", RFC 4924, July 2007.
- [RFC4984] Meyer, D., Ed., Zhang, L., Ed., and K. Fall, Ed., "Report from the IAB Workshop on Routing and Addressing", RFC 4984, September 2007.
- [RFC5902] Thaler, D., Zhang, L., and G. Lebovitz, "IAB Thoughts on IPv6 Network Address Translation", RFC 5902, July 2010.
- [RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, March 2011.

- [RFC6740] Atkinson, R. and S. Bhatti, "Identifier-Locator Network Protocol (ILNP) Architectural Description", RFC 6740, November 2012.
- [RFC6741] Atkinson, R. and S. Bhatti, "Identifier-Locator Network Protocol (ILNP) Engineering and Implementation Considerations", RFC 6741, November 2012.
- [RFC6742] Atkinson, R., Bhatti, S. and S. Rose, "DNS Resource Records for the Identifier-Locator Network Protocol (ILNP)", RFC 6742, November 2012.
- [RFC6743] Atkinson, R. and S. Bhatti, "ICMPv6 Locator Update Message", RFC 6743, November 2012.
- [RFC6744] Atkinson, R. and S. Bhatti, "IPv6 Nonce Destination Option for the Identifier-Locator Network Protocol for IPv6 (ILNPv6)", RFC 6744, November 2012.
- [RFC6745] Atkinson, R. and S. Bhatti, "ICMP Locator Update Message for the Identifier-Locator Network Protocol for IPv4 (ILNPv4)", RFC 6745, November 2012.
- [RFC6746] Atkinson, R. and S. Bhatti, "IPv4 Options for the Identifier-Locator Network Protocol (ILNP)", RFC 6746, November 2012.
- [RFC6747] Atkinson, R. and S. Bhatti, "Address Resolution Protocol (ARP) Extension for the Identifier-Locator Network Protocol for IPv4 (ILNPv4)", RFC 6747, November 2012.

10.2. Informative References

- [ABH07a] Atkinson, R., Bhatti, S., and S. Hailes, "Mobility as an Integrated Service Through the Use of Naming", Proceedings of ACM Workshop on Mobility in the Evolving Internet Architecture (MobiArch), ACM SIGCOMM, Kyoto, Japan. 27 Aug 2007.
- [ABH07b] Atkinson, R., Bhatti, S., and S. Hailes, "A Proposal for Unifying Mobility with Multi-Homing, NAT, & Security", Proceedings of 2nd ACM Workshop on Mobility Management and Wireless Access (MobiWAC), ACM, Chania, Crete, Oct 2007. ISBN: 978-1-59593-809-1

- [ABH08a] Atkinson, R., Bhatti, S., and S. Hailes, "Mobility Through Naming: Impact on DNS", Proceedings of 3rd ACM Workshop on Mobility in the Evolving Internet Architecture (MobiArch), ACM SIGCOMM, Seattle, WA, USA, Aug 2008.
- [ABH08b] Atkinson, R., Bhatti, S., and S. Hailes, "Harmonised Resilience, Security, and Mobility Capability for IP", Proceedings of the IEEE Military Communications Conference (MILCOM), IEEE, San Diego, CA, USA, Nov 2008.
- [ABH09a] Atkinson, R, Bhatti, S., and S. Hailes, "Site-Controlled Secure Multi-Homing and Traffic Engineering For IP", Proceedings of IEEE Military Communications Conference (MILCOM), IEEE, Boston, MA, USA, Oct 2009.
- [ABH09b] Atkinson, R., Bhatti, S., and S. Hailes, "ILNP: Mobility, Multi-Homing, Localised Addressing and Security Through Naming", Telecommunication Systems", vol. 42, no. 3-4, pp 273-291, Springer-Verlag, Dec 2009.
- [ABH10] Atkinson, R., Bhatti, S., and S. Hailes, "Evolving the Internet Architecture Through Naming", IEEE Journal on Selected Areas in Communication (JSAC), vol. 28, no. 8, pp 1319-1325, IEEE, Oct 2010.
- [appDNS] Peterson, J., Kolkman, O., Tschofenig, H., and B. Aboba, "Architectural Considerations on Application Features in the DNS", Work in Progress, July 2012.
- [BA11] Bhatti, S. and R. Atkinson, "Reducing DNS Caching", Proceedings of IEEE Global Internet Symposium (GI2011), Shanghai, P.R. China, 15 Apr 2011.
- [BA12] Bhatti, S. and R. Atkinson, "Secure & Agile Wide-area Virtual Machine Mobility", Proceedings of IEEE Military Communications Conference (MILCOM), Orlando, FL, USA, Oct 2012.
- [BAK11] Bhatti, S., Atkinson, R., and J. Klemets, "Integrating Challenged Networks", Proceedings of IEEE Military Communications Conference (MILCOM), IEEE, Baltimore, MD, USA, Nov 2011.
- [BRDP11] Boot, T. and A. Holtzer, "BRDP Framework", Work in Progress, January 2011.

- [DMS04] Dingledine, R., Mathewson, N., and P. Syverson, "Tor: the second-generation onion router", Proceedings of 13th USENIX Security Symposium, USENIX Association, San Diego, CA, USA, 2004.
- [IEEE04] "IEEE 802.1D - IEEE Standard for Local and Metropolitan Area Networks, Media Access Control (MAC) Bridges", IEEE Standards Association, New York, NY, USA, 9 June 2004. Print: ISBN 0-7381-3881-5 SH95213. PDF: ISBN 0-7381-3982-3 SS95213.
- [LABH06] Atkinson, R., Lad, M., Bhatti, S., and S. Hailes, "A Proposal for Coalition Networking in Dynamic Operational Environments", Proceedings of IEEE Military Communications Conference (MILCOM), IEEE, Washington, DC, USA, Nov 2006.
- [mDNS11] Cheshire, S. and M. Krochmal, "Multicast DNS", Work in Progress, December 2011.
- [RAB09] Rehunathan, D., Atkinson, R., and S. Bhatti, "Enabling Mobile Networks Through Secure Naming", Proceedings of IEEE Military Communications Conference (MILCOM), IEEE, Boston, MA, USA, Oct 2009.
- [RB10] Rehunathan, D. and S. Bhatti, "A Comparative Assessment of Routing for Mobile Networks", Proceedings of 6th IEEE International Conference on Wireless and Mobile Computing Networking and Communications (WiMob), IEEE, Niagara Falls, ON, Canada, Oct 2010.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011.
- [RSG98] Reed, M., Syverson, P., and D. Goldschlag, "Anonymous Connections and Onion Routing", IEEE Journal on Selected Areas in Communications, Vol. 16, No. 4, IEEE, Piscataway, NJ, USA, May 1998.

11. Acknowledgements

Steve Blake, Stephane Bortzmeyer, Mohamed Boucadair, Noel Chiappa, Wes George, Steve Hailes, Joel Halpern, Mark Handley, Volker Hilt, Paul Jakma, Dae-Young Kim, Tony Li, Yakov Rehkter, Bruce Simpson, Robin Whittle, and John Wroclawski (in alphabetical order) provided review and feedback on earlier versions of this document. Steve Blake provided an especially thorough review of an early version of the entire ILNP document set, which was extremely helpful. We also wish to thank the anonymous reviewers of the various ILNP papers for their feedback.

Roy Arends provided expert guidance on technical and procedural aspects of DNS issues.

Authors' Addresses

RJ Atkinson
Consultant
San Jose, CA 95125
USA

EEmail: rja.lists@gmail.com

SN Bhatti
School of Computer Science
University of St Andrews
North Haugh, St Andrews
Fife KY16 9SX
Scotland, UK

EEmail: saleem@cs.st-andrews.ac.uk