     Identifier-Locator Network Protocol (ILNP) Architectural Description

Abstract

   This document provides an architectural description and the concept
   of operations for the Identifier-Locator Network Protocol (ILNP),
   which is an experimental, evolutionary enhancement to IP. This is a
   product of the IRTF Routing Research Group.

Status of This Memo

Copyright Notice

Table of Contents

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.

There has been substantial research relating to naming in the
Internet through the years [IEN1] [IEN19] [IEN23] [IEN31] [IEN135]
[RFC814] [RFC1498] [RFC2956]. Much of that research has indicated
that binding the end-to-end transport-layer session state with a
specific interface of a node at a specific location is undesirable,
for example, creating avoidable issues for mobility, multihoming, and
end-to-end security. More recently, mindful of that important prior
work, and starting well before the Routing RG was re-chartered to
focus on inter-domain routing scalability, the authors have been
examining enhancements to certain naming aspects of the Internet
Architecture. Separately, the Internet Architecture Board (IAB)
recently considered the matter of Internet evolution, including
naming [RFC6250].

Our ideas and progress so far are embodied in the ongoing definition
of an experimental protocol that we call the Identifier-Locator
Network Protocol (ILNP).

Links to relevant material are all available at:
   http://ilnp.cs.st-andrews.ac.uk/

At the time of writing, the main body of peer-reviewed research from
which the ideas in this and the accompanying documents draw is given
in [LABH06], [ABH07a], [ABH07b], [ABH08a], [ABH08b], [ABH09a],
[ABH09b], [RAB09], [ABH10], [RB10], [BA11], [BAK11], and [BA12].

In this document, we:

   a) describe the architectural concepts behind ILNP and how various
      ILNP capabilities operate: this document deliberately focuses
      on describing the key architectural changes that ILNP
      introduces and defers engineering discussion to separate
      documents.

Other documents (listed below):

   b) show how functions based on ILNP would be realised on today's
      Internet by proposing an instance of ILNP based on IPv6, which
      we call ILNPv6 (there is also a document describing ILNPv4,
      which is how ILNP could be applied to IPv4).

   c) discuss salient operational and engineering issues impacting
      the deployment of ILNPv6 and the impact on the Internet.

      d) give architectural descriptions of optional advanced
         capabilities in advanced deployments based on the ILNP
         approach.

1.1. Document Roadmap

   This document describes the architecture for the Identifier-Locator
   Network Protocol (ILNP) including concept of operations. The authors
   recommend reading and understanding this document as the starting
   point to understanding ILNP.

   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) [RFC6741] describes engineering and implementation considerations
      that are common to both ILNPv4 and ILNPv6.

   b) [RFC6742] defines additional DNS resource records that support
      ILNP.

   c) [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.

   d) [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.

   e) [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.

   f) [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.

   g) [RFC6747] describes extensions to the Address Resolution Protocol
      (ARP) for use with ILNPv4.

   h) [RFC6748] describes optional engineering and deployment functions
      for ILNP. These are not required for the operation or use of ILNP
      and are provided as additional options.

1.2. History

   In 1977, Internet researchers at University College London wrote the
   first Internet Experiment Note (IEN), which discussed issues with the
   interconnection of networks [IEN1]. This identified the inclusion of
   network-layer addresses in the transport-layer session state (e.g.,
   TCP checksum) as a significant problem for mobile and multihomed
   nodes and networks. It also proposed separation of identity from
   location as a better approach to take when designing the TCP/IP
   protocol suite. Unfortunately, that separation did not occur, so the
   deployed IPv4 and IPv6 Internet entangles upper-layer protocols
   (e.g., TCP, UDP) with network-layer routing and topology information
   (e.g., IP Addresses) [IEN1] [RFC768] [RFC793].

   The architectural concept behind ILNP derives from a June 1994 note
   by Bob Smart to the IETF SIPP WG mailing list [SIPP94]. In January
   1995, Dave Clark sent a similar note to the IETF IPng WG mailing
   list, suggesting that the IPv6 address be split into separate
   Identifier and Locator fields [IPng95].

   Afterwards, Mike O'Dell pursued this concept in Internet-Drafts
   describing "8+8" [8+8] and "GSE" (Global, Site, and End-system)
   [GSE]. More recently, the IRTF Namespace Research Group (NSRG)
   studied this matter around the turn of the century. Unusually for an
   IRTF RG, the NSRG operated on the principle that unanimity was
   required for the NSRG to make a recommendation. Atkinson was a
   member of the IRTF NSRG. At least one other protocol, the Host
   Identity Protocol (HIP), also derives in part from the IRTF NSRG
   studies (and related antecedent work). This current proposal differs
   from O'Dell's work in various ways, notably in that it does not
   require deployment or use of Locator rewriting.

The key idea proposed for ILNP is to directly and specifically change
the overloaded semantics of the IP Address. The Internet community
has indicated explicitly, several times, that this use of overloaded
semantics is a significant problem with the use of the Internet
protocol today [RFC1498] [RFC2101] [RFC2956] [RFC4984].

While the research community has made a number of proposals that
could provide solutions, so far there has been little progress on
changing the status quo.

1.3. 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. Architectural Overview

ILNP takes a different approach to naming of communication objects
within the network stack. Two new data types are introduced which
subsume the role of the IP Address at the network and transport
layers in the current IP architecture.

2.1. Identifiers and Locators

ILNP explicitly replaces the use of IP Addresses with two distinct
name spaces, each having distinct and different semantics:

   a) Identifier: a non-topological name for uniquely identifying a
      node.

   b) Locator: a topologically bound name for an IP subnetwork.

The use of these two new namespaces in comparison to IP is given in
Table 1. The table shows where existing names are used for state
information in end-systems or protocols.

```
      Layer       |         IP           |     ILNP
   --------------+----------------------+--------------
    Application   | FQDN or IP Address   | FQDN
    Transport     | IP Address           | Identifier
    Network       | IP Address           | Locator
    Physical i/f  | IP Address           | MAC address
   --------------+----------------------+--------------
```

```
   FQDN = Fully Qualified Domain Name
   i/f = interface
   MAC = Media Access Control
```

        Table 1: Use of Names for State Information in Various
                Communication Layers for IP and ILNP

   As shown in Table 1, if an application uses a Fully Qualified Domain
   Name at the application-layer, rather than an IP Address or other
   lower-layer identifier, then the application perceives no
   architectural difference between IP and ILNP. We call such
   applications "well-behaved" with respect to naming as use of the FQDN
   at the application-layer is recommended in [RFC1958]. Some other
   applications also avoid use of IP Address information within the
   application-layer protocol; we also consider these applications to be
   "well-behaved". Any well-behaved application should be able to
   operate on ILNP without any changes. Note that application-level use
   of IP Addresses includes application-level configuration information,
   e.g., Apache web server (httpd) configuration files make extensive
   use of IP Addresses as a form of identity.

   ILNP does not require applications to be rewritten to use a new
   Networking Application Programming Interface (API). So existing
   well-behaved IP-based applications should be able to work over ILNP
   as is.

   In ILNP, transport-layer protocols use only an end-to-end, non-
   topological node Identifier in any transport-layer session state. It
   is important to note that the node Identifier names the node, not a
   specific interface of the node. In this way, it has different
   semantics and properties than either the IPv4 address, the IPv6
   address, or the IPv6 interface identifier [RFC791] [RFC4291].

   The use of the ILNP Identifier value within application-layer
   protocols is not recommended. Instead, the use of either a FQDN or
   some different topology-independent namespace is recommended.

   At the network-layer, Locator values, which have topological
   significance, are used for routing and forwarding of ILNP packets,
   but Locators are not used in upper-layer protocols.

As well as the new namespaces, another significant difference in
ILNP, as shown in Table 1, is that there is no binding of a routable
name to an interface, or Sub-Network Point of Attachment (SNPA), as
there is in IP. The existence of such a binding in IP effectively
binds transport protocol flows to a specific, single interface on a
node. Also, applications that include IP Addresses in their
application-layer session state effectively bind to a specific,
single interface on a node [RFC2460] [RFC6724].

In ILNP, dynamic bindings exist between Identifier values and
associated Locator values, as well as between {Identifier, Locator}
pairs and (physical or logical) interfaces on the node.

This change enhances the Internet Architecture by adding crisp and
clear semantics for the Identifier and for the Locator, removing the
overloaded semantics of the IP Address [RFC1992] [RFC4984], by
updating end-system protocols, but without requiring any router or
backbone changes. In ILNP, the closest approximation to an IP
Address is an I-L Vector (I-LV), which is a given binding between an
Identifier and Locator pair, written as [I, L]. I-LVs are discussed
in more detail below.

Where, today, IP packets have:

- Source IP Address, Destination IP Address

instead, ILNP packets have:

- source I-LV, destination I-LV

However, it must be emphasised that the I-LV and the IP Address are
not equivalent.

With these naming enhancements, we will improve the Internet
Architecture by adding explicit harmonised support for many
functions, such as multihoming, mobility, and IPsec.

2.2. Deprecating IP Addresses

ILNP places an explicit Locator and Identifier in the IP packet
header, replacing the usual IP Address. Locators are tied to the
topology of the network. They may change frequently, as the node or
site changes its network connectivity. The node Identifier is
normally much more static and remains constant throughout the life of
a given transport-layer session, and frequently much longer.
However, there are various options for Identifier values, as
discussed in [RFC6741]. The way that I-LVs are encoded into packet
headers is different for IPv4 and IPv6, as explained in [RFC6741].

Identifiers and Locators for hosts are advertised explicitly in DNS,
through the use of new Resource Records (RRs). This is a logical and
reasonable use of DNS, completely analogous to the capability that
DNS provides today. At present, among other current uses, the DNS is
used to map from an FQDN to a set of addresses. As ILNP replaces IP
Addresses with Identifiers and Locators, it is then clearly rational
to use the DNS to map an FQDN to a set of Identifiers and a set of
Locators for a node.

The presence of ILNP Locators and Identifiers in the DNS for a DNS
owner name is an indicator to correspondents that the correspondents
can try to establish an ILNP-based transport-layer session with that
DNS owner name.

Specifically in response to [RFC4984], ILNP improves routing
scalability by helping multihomed sites operate effectively with
Provider Aggregated (PA) address prefixes. Many multihomed sites
today request provider-independent (PI) address prefixes so they can
provide session survivability despite the failure of one or more
access links or Internet Service Providers (ISPs). ILNP provides
this transport-layer session survivability by having a provider-
independent Node Identifier (NID) value that is free of any
topological semantics. This NID value can be bound dynamically to a
Provider Aggregated Locator (L) value, the latter being a topological
name, i.e., a PA network prefix. By allowing correspondents to
change arbitrarily among multiple PA Locator values, survivability is
enabled as changes to the L values need not disrupt transport-layer
sessions. In turn, this allows an ILNP multihomed site to have both
the full transport-layer and full network-layer session resilience
that is today offered by PI addressing while using the equivalent of
PA addressing. In turn, this eliminates the current need to use
globally visible PI routing prefixes for each multihomed site.

2.3. Session Terminology

To improve clarity and readability of the several ILNP specification
documents, this section defines the terms "network-layer session" and
"transport-layer session" both for IP-based networks and ILNP-based
networks.

Today, network-layer IP sessions have 2 components:

- Source IP Address (A_S)
- Destination IP Address (A_D)

For example, a tuple for an IP layer session would be:

   <IP: A_S, A_D>

Instead, network-layer ILNP sessions have 4 components:

- Source Locator(s) (L_S)
- Source Identifier(s) (I_S)
- Destination Locator(s) (L_D)
- Destination Identifier(s) (L_S)

and a tuple for an ILNP session would be:

   <ILNP: I_S, L_S, I_D, L_D>

The phrase "ILNP session" refers to an ILNP-based network-layer
session, having the 4 components in the definition above.

For engineering efficiency, multiple transport-layer sessions between
a pair of ILNP correspondents normally share a single ILNP session
(I-LV pairs and associated Nonce values). Also, for engineering
convenience (and to cope with situation where different nodes, at
different locations, might use the same I values), in the specific
implementation of ILNPv6 and ILNPv4, we define the use of nonce
values:

- Source-to-destination Nonce value (N_S)
- Destination-to-source Nonce value (N_D)

These are explained in more detail in [RFC6741], with [RFC6744] for
ILNPv6 and [RFC6746] for ILNPv4.

Today, transport-layer sessions using IP include these 5 components:

 - Source IP Address (A_S)
 - Destination IP Address (A_D)
 - Transport-layer protocol (e.g., UDP, TCP, SCTP)
 - Source transport-layer port number (P_S)
 - Destination transport-layer port number (P_D)

For example, a TCP tuple would be:

   <TCP: P_S, P_D, A_S, A_D>

Instead, transport-layer sessions using ILNP include these 5
components:

- Source Identifier (I_S)
- Destination Identifier (I_D)
- Transport-layer protocol (e.g., UDP, TCP, SCTP)
- Source transport-layer port number (P_S)
- Destination transport-layer port number (P_D)

and an example tuple:

       <TCP: P_S, P_D, I_S, I_D>

2.4. Other Goals

   While we seek to make significant enhancements to the current
   Internet Architecture, we also wish to ensure that instantiations of
   ILNP are:

      a) Backwards compatible: implementations of ILNP should be able to
         work with existing IPv6 or IPv4 deployments, without requiring
         application changes.

      b) Incrementally deployable: to deploy an implementation of ILNP,
         changes to the network nodes should only be for those nodes
         that choose to use ILNP. The use of ILNP by some nodes does
         not require other nodes (that do not use ILNP) to be upgraded.

3. Architectural Changes Introduced by ILNP

   In this section, we describe the key changes that are made to the
   current Internet Architecture. These key changes impact end-systems,
   rather than routers.

3.1. Identifiers

   Identifiers, also called Node Identifiers (NIDs), are non-topological
   values that identify an ILNP node. A node might be a physical node
   or a virtual node. For example, a single physical device might
   contain multiple independent virtual nodes. Alternately, a single
   virtual device might be composed from multiple physical devices. In
   the case of a Multi-Level Secure (MLS) system [DIA] [DoD85] [DoD87]
   [RFC5570], each valid Sensitivity Label of that system might be a
   separate virtual node.

   A node MAY have multiple Identifier values associated with it, which
   MAY be used concurrently.

   In normal operation, when a node is responding to a received ILNP
   packet that creates a new network-layer session, the correct NID
   value to use for that network-layer session with that correspondent
   node will be learned from the received ILNP packet.

   In normal operation, when a node is initiating communication with a
   correspondent node, the correct I value to use for that session with
   that correspondent node will be learned either through the
   application-layer naming, through DNS name resolution, or through

some alternative name resolution system. Another option is an
application may be able to select different I values directly -- as
Identifiers are visible above the network layer via the transport
protocol.

3.1.1. Node Identifiers Are Immutable during a Session

Once a Node Identifier (NID) value has been used to establish a
transport-layer session, that Node Identifier value forms part of the
end-to-end (invariant) transport-layer session state and so MUST
remain fixed for the duration of that session. This means, for
example, that throughout the duration of a given TCP session, the
Source Node Identifier and Destination Node Identifier values will
not change.

In normal operation, a node will not change its set of valid
Identifier values frequently. However, a node MAY change its set of
valid Identifier values over time, for example, in an effort to
provide identity obfuscation, while remaining subject to the
architectural rule of the preceding paragraph. When a node has more
than one Node Identifier value concurrently, the node might have
multiple concurrent ILNP sessions with some correspondent node, in
which case Node Identifier values MAY differ between the different
concurrent ILNP sessions.

3.1.2. Syntax

ILNP Identifiers have the same syntax as IPv6 interface identifiers
[RFC4291], based on the EUI-64 format [IEEE-EUI], which helps with
backwards compatibility. There is no semantic equivalent to an ILNP
Identifier in IPv4 or IPv6 today.

The Modified EUI-64 syntax used by both ILNP Identifiers and IPv6
interface identifiers contains a bit indicating whether the value has
global scope or local scope [IEEE-EUI] [RFC4291]. ILNP Identifiers
have either global scope or local scope. If they have global scope,
they SHOULD be globally unique.

Regardless of whether an Identifier is global scope or local scope,
an Identifier MUST be unique within the scope of a given Locator
value to which it is bound for a given ILNP session or packet flow.
As an example, with ILNPv6, the ordinary IPv6 Neighbour Discovery
(ND) processes ensure that this is true, just as ND ensures that no
two IPv6 nodes on the same IPv6 subnetwork have the same IPv6 address
at the same time.

Both the IEEE EUI-64 specification and the Modified EUI-64 syntax
also has a 'Group' bit [IEEE-EUI] [RFC4291]. For both ILNP node
Identifiers and also IPv6 interface identifiers, this Group bit is
set to 0.

3.1.3. Semantics

Unicast ILNP Identifier values name the node, rather than naming a
specific interface on that node. So ILNP Identifiers have different
semantics than IPv6 interface identifiers.

3.2. Locators

Locators are topologically significant names, analogous to
(sub)network routing prefixes. The Locator names the IP subnetwork
that a node is connected to. ILNP neither prohibits nor mandates in-
transit modification of Locator values.

A host MAY have several Locators at the same time, for example, if it
has a single network interface connected to multiple subnetworks
(e.g., VLAN deployments on wired Ethernet) or has multiple interfaces
each on a different subnetwork. Locator values normally have Locator
Preference Indicator (LPI) values associated with them. These LPIs
indicate that a specific Locator value has higher or lower preference
for use at a given time. Local LPI values may be changed through
local policy or via management interfaces. Remote LPI values are
normally learned from the DNS, but the local copy of a remote LPI
value might be modified by local policy relating to preferred paths
or prefixes.

Locator values are used only at the network layer. Locators are not
used in end-to-end transport state. For example, Locators are not
used in transport-layer session state or application-layer session
state. However, this does not preclude an end-system setting up
local dynamic bindings for a single transport flow to multiple
Locator values concurrently.

The routing system only uses Locators, not Identifiers. For unicast
traffic, ILNP uses longest-prefix match routing, just as the IP
Internet does.

Section 4 below describes in more detail how Locators are used in
forwarding and routing packets from a sending node on a source
subnetwork to one or more receiving nodes on one or more destination
subnetworks.

A difference from earlier proposals [GSE] [8+8] is that, in normal
operation, the originating host supplies both Source Locator and
Destination Locator values in the packets it sends out.

Section 4.3 describes packet forwarding in more detail, while Section
4.4 describes packet routing in more detail.

3.2.1. Locator Values Are Dynamic

The ILNP architecture recognises that Locator values are
topologically significant, so the set of Locator values associated
with a node normally will need to change when the node's connectivity
to the Internet topology changes. For example, a mobile or
multihomed node is likely to have connectivity changes from time to
time, along with the corresponding changes to the set of Locator
values.

When a node using a specific set of Locator values changes one or
more of those Locator values, then the node (1) needs to update its
local knowledge of its own Locator values, (2) needs to inform all
active Correspondent Nodes (CNs) of those changes to its set of
Locator values so that ILNP session continuity is maintained, and (3)
if it expects incoming connections the node also needs to update its
Locator-related entries in the Domain Name System. [RFC6741]
describes the engineering and implementation details of this process.

3.2.2. Locator Updates

As Locator values can be dynamic, and they could change for a node
during an ILNP session, correspondents need to be notified when a
Locator value for a node changes for any existing ILNP session. To
enable this, a node that sees its Locator values have changed MUST
send a Locator Update (LU) message to its correspondent nodes. The
details of this procedure are discussed in other ILNP documents --
[RFC6741], [RFC6743], and [RFC6745]. (The change in Locator values
may also need to be notified to DNS but that is discussed elsewhere.)

3.2.3. Syntax

ILNP Locators have the same syntax as an IP unicast routing prefix.

3.2.4. Semantics

ILNP unicast Locators have the same semantics as an IP unicast
routing prefix, since they name a specific subnetwork. ILNP neither
prohibits nor requires in-transit modification of Locator values.

3.3. IP Address and Identifier-Locator Vector (I-LV)

   Historically, an IP Address has been considered to be an atomic
   datum, even though it is recognised that an IP Address has an
   internal structure: the network prefix plus either the host ID (IPv4)
   or the interface identifier (IPv6). However, this internal structure
   has not been used in end-system protocols; instead, all the bits of
   the IP Address are used. (Additionally, in IPv4 the IPv4 subnet mask
   uses bits from the host ID, a further confusion of the structure,
   even thought it is an extremely useful engineering mechanism.)

   In ILNP, the IP Address is replaced by an "Identifier-Locator Vector"
   (I-LV). This consists of a pairing of an Identifier value and a
   Locator value for that packet, written as [I, L]. All ILNP packets
   have Source Identifier, Source Locator, Destination Identifier, and
   Destination Locator values. The I value of the I-LV is used by
   upper-layer protocols (e.g., TCP, UDP, SCTP), so needs to be
   immutable. Locators are not used by upper-layer protocols (e.g.,
   TCP, UDP, SCTP). Instead, Locators are similar to IP routing
   prefixes, and are only used to name a specific subnetwork.

   While it is possible to say that an I-LV is an approximation to an IP
   Address of today, it should be understood that an I-LV:

      a) is not an atomic datum, being a pairing of two data types, an
         Identifier and a Locator.

      b) has different semantics and properties to an IP Address, as is
         described in this document.

   In our discussion, it will be convenient sometimes to refer to an
   I-LV, but sometimes to refer only to an Identifier value, or only to
   a Locator value.

   ILNP packets always contain a source I-LV and a destination I-LV.

3.4. Notation

   In describing how capabilities are implemented in ILNP, we will
   consider the differences in end-systems' state between IP and ILNP in
   order to highlight the architectural changes.

We define a formal notation to represent the data contained in the
transport-layer session state. We define:

    A = IP Address
    I = Identifier
    L = Locator
    P = Transport-layer port number

To differentiate the local and remote values for the above items, we
also use suffixes, for example:

    _L = local
    _R = remote

With IPv4 and IPv6 today, the invariant state at the transport-layer
for TCP can be represented by the tagged tuple:

    <TCP: A_L, A_R, P_L, P_R>   --- (1)

Tag values that will be used are:

     IP   Internet Protocol
     ILNP Identifier-Locator Network Protocol
     TCP   Transmission Control Protocol
     UDP   User Datagram Protocol

So, for example, with IP, a UDP packet would have the tagged tuple:

    <UDP: A_L, A_R, P_L, P_R>   --- (2)

A TCP segment carried in an IP packet may be represented by the
tagged tuple binding:

    <TCP: A_L, A_R, P_L, P_R><IP: A_L, A_R>   --- (3)

and a UDP packet would have the tagged tuple binding:

    <UDP: A_L, A_R, P_L, P_R><IP: A_L, A_R>   --- (4)

In ILNP, the transport-layer state for TCP is:

    <TCP: I_L, I_R, P_L, P_R>   --- (5)

The binding for a TCP segment within an ILNP packet:

    <TCP: I_L, I_R, P_L, P_R><ILNP: L_L, L_R>   --- (6)

When comparing tuple expressions (3) and (6), we see that for IP, any
change to network addresses impacts the end-to-end state, but for
ILNP, changes to Locator values do not impact end-to-end state. This
provides end-system session state invariance, a key feature of ILNP
compared to IP as it is used in some situations today. ILNP adopts
the end-to-end approach for its architecture [SRC84]. As noted
previously, nodes MAY have more than one Locator concurrently, and
nodes MAY change their set of active Locator values as required.

While these documents do not include SCTP examples, the same notation
can be used, simply substituting the string "SCTP" for the string
"TCP" or the string "UDP" in the above examples.

3.5. Transport-Layer State and Transport Pseudo-Headers

In ILNP, protocols above the network layer do not use the Locator
values. Thus, the transport layer uses only the I values for the
transport-layer session state (e.g., TCP pseudo-header checksum, UDP
pseudo-header checksum), as is shown, for example, in expression (6)
above.

Additionally, from a practical perspective, while the I values are
only used in protocols above the network layer, it is convenient for
them to be carried in network packets, so that the namespace for the
I values can be used by any transport-layer protocols operating above
the common network layer.

3.6. Rationale for This Document

This document provides an architectural description of the core ILNP
capabilities and functions. It is based around the use of example
scenarios so that practical issues can be highlighted.

In some cases, illustrative suggestions and light discussion are
presented with respect to engineering issues, but detailed discussion
of engineering issues are deferred to other ILNP documents.

The order of the examples presented below is intended to allow an
incremental technical understanding of ILNP to be developed. There
is no other reason for the ordering of the examples listed below.

Many of the descriptions are based on the use of an example site
network as shown in Figure 3.1.

```
       site                             . . . .        +----+
      network                        .        .-----+ CN |
     . . . .        +------+ link1 .         .     +----+
    .       .       |      +------.          .
   .    D    .      |      |       .          .
   .            .----+ SBR |       . Internet  .
   .  H    .       |      |       .          .
   .       .       |      +------.          .
    . . . .        +------+ link2 .         .
                                   .         .
                                    . . . .
```

```
        CN  = Correspondent Node
         D  = Device
         H  = Host
       SBR  = Site Border Router
```

    Figure 3.1: A Simple Site Network for ILNP Examples

   In some cases, hosts (H) or devices (D) act as end-systems within the
   site network, and communicate with (one or more) Correspondent Node
   (CN) instances that are beyond the site.

   Note that the figure is illustrative and presents a logical view.
   For example, the CN may itself be on a site network, just like H or
   D.

   Also, for formulating examples, we assume ILNPv6 is in use, which has
   the same packet header format (as viewed by routers) as IPv6, and can
   be seen as a superset of IPv6 capabilities.

   For simplicity, we assume that name resolution is via the deployed
   DNS, which has been updated to store DNS records for ILNP [RFC6742].

   Note that, from an engineering viewpoint, this does NOT mean that the
   DNS also has to be ILNP capable: existing IPv4 or IPv6 infrastructure
   can be used for DNS transport.

3.7. ILNP Multicasting

   Multicast forwarding and routing are unchanged, in order to avoid
   requiring changes in deployed IP routers and routing protocols.
   ILNPv4 multicasting is the same as IETF Standards Track IPv4
   multicasting [RFC1112] [RFC3376]. ILNPv6 multicasting is the same as
   IETF Standards Track IPv6 multicasting [RFC4291] [RFC2710] [RFC3810].

4. ILNP Basic Connectivity

   In this section, we describe basic packet forwarding and routing in
   ILNP. We highlight areas where it is similar to current IP, and also
   where it is different from current IP. We use examples in order to
   illustrate the intent and show the feasibility of the approach.

   For this section, in Figure 4.1, H is a fixed host in a simple site
   network, and CN is a remote Correspondent Node outside the site; H
   and CN are ILNP-capable, while the Site Border Router (SBR) does not
   need to be ILNP-capable.

```
        site                          . . . .       +----+
      network                       .        .-----+ CN |
      . . . .        +------+       .        .       +----+
      .       .      |  +------.    .              .
      .       .      |  |      .    .              .
      .      .----+ SBR |    . Internet  .
     .  H      .   |  |      .              .
      .       .      |  |      .              .
      . . . .        +------+        .              .
                                      .          .
                                       . . . .
```

          CN = Correspondent Node
           H = Host
         SBR = Site Border Router

       Figure 4.1: A Simple Site Network for ILNP Examples

4.1. Basic Local Configuration

   This section uses the term "address management", in recognition of
   the analogy with capabilities present in IP today. In this document,
   address management is about enabling hosts to attach to a subnetwork
   and enabling network-layer communication between and among hosts,
   also including:

      a) enabling identification of a node within a site.
      b) allowing basic routing/forwarding from a node acting as an end-
         system.

   If we consider Figure 4.1, imagine that host H has been connected to
   the site network. Administratively, it needs at least one I value
   and one L value in order to be able to communicate.

Today, local administrative procedures allocate IP Addresses, often
using various protocol mechanisms (e.g., NETCONF-based router
configuration, DHCP for IPv4, DHCP for IPv6, IPv6 Router
Advertisements). Similarly, local administrative procedures can
allocate I and L values as required, e.g., I_H and L_H. This may be
through manual configuration.

Additionally, if it is expected or desired that H might have incoming
communication requests, e.g., it is a server, then the values I_H and
L_H can be added to the relevant name services (e.g., DNS, NIS/YP),
so that FQDN lookups for H resolve to the appropriate DNS resource
records (e.g., NID, L32, L64, and LP [RFC6742]) for node H.

From a network operations perspective, this whole process also can be
automated. As an example, consider that in Figure 3.1 the Site
Border Router (SBR) is an IPv6-capable router and is connected via
link1 to an ISP that supports IPv6. The SBR will have been allocated
one (or more) IPv6 prefixes that it will multicast using IPv6 Routing
Advertisements (RAs) into the site network, e.g., prefix L_1. L_1 is
actually a local IPv6 prefix (/64), which is formed from an address
assignment by the upstream ISP, according to [RFC3177] or [RFC6177].
Host H will see these RAs, for example, on its local interface with
name eth0, will be able to use that prefix as a Locator value, and
will cache that Locator value locally.

Also, node H can use the mechanism documented in either Section 2.5.1
of [RFC4291], in [RFC3972], [RFC4581], [RFC4982], or in [RFC4941] in
order to create a default I value (say, I_H), just as an IPv6 host
can. For DNS, the I_H and L_1 values may be pre-configured in DNS by
an administrator who already has knowledge of these, or added to DNS
by H using Secure DNS Dynamic Update [RFC3007] to add or update the
correct NID and L64 records to DNS for the FQDN for H.

4.2. I-L Communication Cache

For the purposes of explaining the concept of operations, we talk of
a local I-L Communication Cache (ILCC). This is an engineering
convenience and does not form part of the ILNP architecture, but is
used in our examples. More details on the ILCC can be found in
[RFC6741]. The ILCC contains information that is required for the
operation of ILNP. This will include, amongst other things, the
current set of valid Identifier and Locator values in use by a node,
the bindings between them, and the bindings between Locator values
and interfaces.

4.3. Packet Forwarding

   When the SBR needs to send a packet to H, it uses local address
   resolution mechanisms to discover the bindings between interface
   addresses and currently active I-LVs for H. For our example of
   Figure 3.1, IPv6 Neighbour Discovery (ND) can be used without
   modification, as the I-LV for ILNPv6 occupies the same bits as the
   IPv6 address in the IPv6 header. For packets from H to SBR, the same
   basic mechanism applies, as long as SBR supports IPv6 and even if it
   is not ILNPv6-capable, as IPv6 ND is used unmodified for ILNPv6.

   For Figure 3.1, assuming:

   - SBR advertises prefix L_1 locally, uses I value I_S, and has an
     Ethernet MAC address M_S on interface with local name sbr0

   - H uses I value I_H, and has an Ethernet MAC address of M_H on the
     interface with local name eth0

   then H will have in its ILCC:

      [I_H, L_1]   --- (7a)
      L_1, eth0   --- (7b)

   After the IPv6 RA and ND mechanism has executed, the ILCC at H would
   contain, as well as expressions (7a) and (7b), the following entry
   for SBR:

      [I_S, L_1], M_S   --- (8)

   For ILNPv6, it does not matter that the SBR is not ILNPv6-capable, as
   the I-LV [I_S, L_1] is physically equivalent to the IPv6 address for
   the internal interface sbr0.

   At SBR, which is not ILNP-capable, there would be the following
   entries in its local cache and configuration:

      L_1:I_S   --- (9a)
      L_1, sbr0   --- (9b)

   Expression (9a) represents a valid IPv6 ND entry: in this case, the
   I_S value (which is 64 bits in ILNPv6) and the L_1 values are,
   effectively, concatenated and treated as if they were a single IPv6
   address. Expression (9b) binds transmissions for L_1 to interface
   sbr0. (Again, sbr0 is a local, implementation-specific name, and
   such a binding is possible with standard tools today, for example,
   ifconfig(8).)

4.4. Packet Routing

   If we assume that host H is configured as in the previous section, it
   is now ready to send and receive ILNP packets.

   Let us assume that, for Figure 4.1, it wishes to contact the node CN,
   which has FQDN cn.example.com and is ILNP-capable. A DNS query by H
   for cn.example.com will result in NID and L64 records for CN, with
   values I_CN and L_CN, respectively, being returned to H and stored in
   its ILCC:

      [I_CN, L_CN]   --- (10)

   This will be considered active as long as the TTL values for the DNS
   records are valid. If the TTL for an I or L value is zero, then the
   value is still usable but becomes stale as soon as it has been used
   once. However, it is more likely that the TTL value will be greater
   than zero [BA11] [SBK01].

   Once the CN's I value is known, the upper-layer protocol, e.g., the
   transport protocol, can set up suitable transport-layer session
   state:

      <UDP: I_H, I_CN, P_H, P_CN>   --- (11)

   For routing of ILNP packets, the destination L value in an ILNPv6
   packet header is semantically equivalent to a routing prefix. So,
   once a packet has been forwarded from a host to its first-hop router,
   only the destination L value needs to be used for getting the packet
   to the destination network. Once the packet has arrived at the
   router for the site network, local mechanisms and the packet-
   forwarding mechanism, as described above in Section 4.3, allow the
   packet to be delivered to the host.

   For our example of Figure 4.1, H will send a UDP packet over ILNP as:

      <UDP: I_H, I_CN, P_H, P_CN><ILNP: L_1, L_CN>   --- (12a)

   and CN will send UDP packets to H as:

      <UDP: I_CN, I_H, P_CN, P_H><ILNP: L_CN, L_1>   --- (12b)

   The I value for H used in the transport-layer state (I_H in
   expression (12a)) selects the correct L value (L_1 in this case) from
   the bindings in the ILCC (expression (7a)), and that, in turn,
   selects the correct interface from the ILCC (expression (7b)), as

described in Section 4.2. This gets the packet to the first hop
router; beyond that, the ILNPv6 packet is treated as if it were an
IPv6 packet.

5. Multihoming and Multi-Path Transport

   For multihoming, there are three cases to consider:

      a) Host Multihoming (H-MH): a single host is, individually,
         connected to multiple upstream links, via separate routing
         paths, and those multiple paths are used by that host as it
         wishes. That is, use of multiple upstream links is managed by
         the single host itself. For example, the host might have
         multiple valid Locator values on a single interface, with each
         Locator value being associated with a different upstream link
         (provider).

      b) Multi-Path Transport (MTP): This is similar to using ILNP's
         support for host multihoming (i.e., H-MH), so we describe
         multi-path transport here. (Indeed, for ILNP, this can be
         considered a special case of H-MH.)

      c) Site Multihoming (S-MH): a site network is connected to
         multiple upstream links via separate routing paths, and hosts
         on the site are not necessarily aware of the multiple upstream
         paths. That is, the multiple upstream paths are managed,
         typically, through a site border router, or via the providers.

   Essentially, for ILNP, multihoming is implemented by enabling:

      a) multiple Locator values to be used simultaneously by a node

      b) dynamic, simultaneous binding between one (or more) Identifier
         value(s) and multiple Locator values

   With respect to the requirements for hosts [RFC1122], the multihoming
   function provided by ILNP is very flexible. It is not useful to
   discuss ILNP multihoming strictly within the confines of the
   exposition presented in Section 3.3.4 of [RFC1122], as that text is
   couched in terms of relationships between IP Addresses and
   interfaces, which can be dynamic in ILNP. The closest relationship
   between ILNP multihoming and [RFC1122] would be that certainly ILNP
   could support the notion of "Multiple Logical Networks", "Multiple
   Logical Hosts", and "Simple Multihoming".

5.1. Host Multihoming (H-MH)

   At present, host multihoming is not common in the deployed Internet.
   When TCP or UDP are in use with an IP-based network-layer session,
   host multihoming cannot provide session resilience, because the
   transport protocol's pseudo-header checksum binds the transport-layer
   session to a single IP Address of the multihomed node, and hence to a
   single interface of that node. SCTP has a protocol-specific
   mechanism to support node multihoming; SCTP can support session
   resilience both at present and also without change in the proposed
   approach [RFC5061].

   Host multihoming in ILNP is supported directly in each host by ILNP.
   The simplest explanation of H-MH for ILNP is that an ILNP-capable
   host can simultaneously use multiple Locator values, for example, by
   having a binding between an I value and two different L values, e.g.,
   the ILCC may contain the I-LVs:

      [I_1, L_1]   --- (14a)
      [I_1, L_2]   --- (14b)

   Additionally, a host may use several I values concurrently, e.g., the
   ILCC may contain the I-LVs:

      [I_1, L_1]   --- (15a)
      [I_1, L_2]   --- (15b)
      [I_2, L_2]   --- (15c)
      [I_3, L_1]   --- (15d)

   Architecturally, ILNP considers these all to be cases of multihoming:
   the host is connected to more than one subnetwork, each subnetwork
   being named by a different Locator value.

   In the cases above, the selection of which I-LV to use would be
   through local policy or through management mechanisms. Additionally,
   suitably modified transport-layer protocols, such as multi-path
   transport-layer protocol implementations, may make use of multiple
   I-LVs. Note that in such a case, the way in which multiple I-LVs are
   used would be under the control of the higher-layer protocol.

   Recall, however, that L values also have preference -- LPI values --
   and these LPI values can be used at the network layer, or by a
   transport-layer protocol implementation, in order make use of L
   values in a specific manner.

   Note that, from a practical perspective, ILNP dynamically binds L
   values to interfaces on a node to indicate the SNPA for that L value,
   so the multihoming is very flexible: a node could have a single

interface and have multiple L values bound to that interface. For
example, for expressions (14a) and (14b), if the end-system has a
single interface with local name eth0, then the entries in the ILCC
will be:

    L_1, eth0   --- (16a)
    L_2, eth0   --- (16b)

And, if we assume that for expressions (15a-c) the end-system has two
interfaces, eth0 and eth1, then these ILCC entries are possible:

    L_1, eth0   --- (17a)
    L_2, eth1   --- (17b)

Let us consider the network in Figure 5.1.
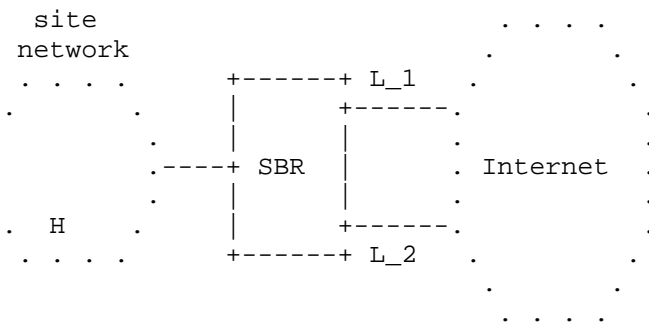
```
      site                                  . . . .
    network                              .         .
    . . . .      +------+ L_1      .              .
    .      .     |      +------.                   .
    .      .     |      |      .                   .
    .          .----+ SBR  |       . Internet    .
    .      .     |      |      .                   .
   .  H    .     |      +------.                   .
    . . . .      +------+ L_2      .              .
                                    .         .
                                      . . . .


        L_1 = global Locator value 1
        L_2 = global Locator value 2
        SBR = Site Border Router
```

       Figure 5.1: A Simple Multihoming Scenario for ILNP

We assume that H has a single interface, eth0. SBR will advertise
L_1 and L_2 internally to the site. Host H will configure these as
both reachable via its single interface, eth0, by using ILCC entries
as in expressions (16a) and (16b). When packets from H that are to
egress the site network reach SBR, it can make appropriate decisions
on which link to use based on the source Locator value (which has
been inserted by H) or based on other local policy.

If, however, H has two interfaces, eth0 and eth1, then it can use
ILCC entries as in expressions (17a) and (17b).

Note that the values L_1 and L_2 do not need to be PI-based Locator
values, and can be taken from ISP-specific PA routing prefix
allocations from the upstream ISPs providing the two links.

Of course, this example is illustrative: many other configurations
are also possible, but the fundamental mechanism remains the same, as
described above.

If any Locator values change, then H will discover this when it sees
new Locator values in RAs from SBR, and sees that L values that were
previously used are no longer advertised. When this happens, H will:

    a) maintain existing active network-layer sessions: based on its
       current ILCC entries and active sessions, send Locator Update
       (LU) messages to CNs to notify them of the change of L values.
       (LU messages are synonymous to Mobile IPv6 Binding Updates.)

    b) if required, update its relevant DNS entries with the new L
       value in the appropriate DNS records, to enable correct
       resolution for new incoming session requests.

From an engineering viewpoint, H also updates its ILCC data, removing
the old L value(s) and replacing with new L value(s) as required.

Depending on the nature of the physical change in connectivity that
the L value change represents, this may disrupt upper-level
protocols, e.g., a fibre cut. Dealing with such physical-level
disruption is beyond the scope of ILNP. However, ILNP supports
graceful changes in L values, and this is explained below in Section
6 in the discussion on mobility support.

5.2. Support for Multi-Path Transport Protocols

ILNP supports deployment and use of multi-path transport protocols,
such as the Multi-Path extensions to TCP (MP-TCP) being defined by
the IETF TCPM Working Group. Specifically, ILNP will support the use
of multiple paths as it allows a single I value to be bound to
multiple L values -- see Section 5.1, specifically expressions (15a)
and (15b).

Of course, there will be specific mechanisms for:
- congestion control
- signalling for connection/session management
- path discovery and path management
- engineering and implementation issues

These transport-layer mechanisms fall outside the scope of ILNP and
would be defined in the multi-path transport protocol specifications.

As far as the ILNP architecture is concerned, the transport protocol
connection is simply using multiple I-LVs, but with the same I value
in each, and different L values, i.e., a multihomed host.

5.3. Site Multihoming (S-MH)

   At present, site multihoming is common in the deployed Internet.
   This is primarily achieved by advertising the site's routing
   prefix(es) to more than one upstream Internet service provider at a
   given time. In turn, this requires de-aggregation of routing
   prefixes within the inter-domain routing system. This increases the
   entropy of the inter-domain routing system (e.g., RIB/FIB size
   increases beyond the minimal RIB/FIB size that would be required to
   reach all sites).

   Site multihoming, in its simplest form in ILNP, is an extension of
   the H-MH scenario described in Section 5.1. If we consider Figure
   5.1, and assume that there are many hosts in the site network, then
   each host can choose (a) whether or not to manage its own ILNP
   connectivity, and (b) whether or not to use multiple Locator values.
   This allows maximal control of connectivity for each host.

   Of course, with ILNPv6, just as any IPv6 router is required to
   generate IPv6 Router Advertisement messages with the correct routing
   prefix information for the link the RA is advertised upon, the SBR is
   also required to generate RAs containing the correct Locator value(s)
   for the link that the RA is advertised upon. The correct values for
   these RA messages are typically configured by system administration,
   or might be passed down from the upstream provider.

   To avoid a DNS Update burst when a site or (sub)network changes
   location, a DNS record optimisation is possible by using the new LP
   record for ILNP. This would change the number of DNS Updates
   required from Order(Number of nodes within the site/subnetwork that
   moved) to Order(1) [RFC6742].

5.3.1. A Common Multihoming Scenario - Multiple SBRs

   The scenario of Figure 5.1 is an example to illustrate the
   architectural operation of multihoming for ILNP. For site
   multihoming, a scenario such as the one depicted in Figure 5.2 is
   also common. Here, there are two SBRs, each with its own global
   connectivity.

```
      site                               . . . .
    network                          .          .
     . . . .        +-------+ L_1  .            .
    .        .      |        +------.            .
   .          .     |        |       .           .
  .          .---+ SBR_A |       .            .
  .          .     |        |       .            .
  .          .     |        |       .            .
  .          .     +-------+       .            .
  .          .          ^          .            .
  .          .          | CP       .  Internet  .
  .          .          v          .            .
  .          .     +-------+ L_2  .             .
  .          .     |        +------.             .
  .          .     |        |       .            .
  .          .---+ SBR_B |       .             .
   .          .     |        |       .            .
    .          .    |        |       .           .
     . . . .        +-------+        .          .
                                      .          .
                                       . . . .
```
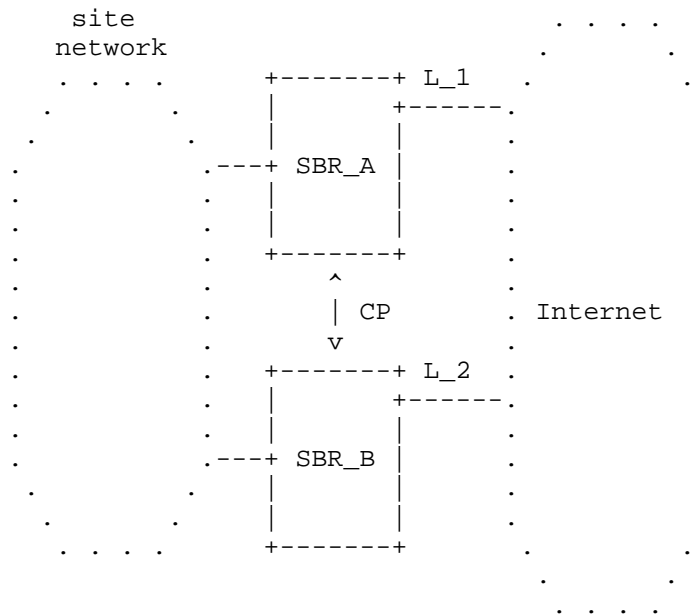
```
        CP      = coordination protocol
        L_1     = global Locator value 1
        L_2     = global Locator value 2
        SBR_A   = Site Border Router A
        SBR_B   = Site Border Router B
```

           Figure 5.2: 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 5.1. 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 capabilities are available today in products. Note
   that, logically, there is little difference between Figures 5.1 and
   5.2, but with two distinct routers in Figure 5.2, the interaction
   using CP is required. Of course, it is also possible to have
   multiple interfaces in each router and more than two routers.

5.4. Multihoming Requirements for Site Border Routers

   For multihoming, the SBR does NOT need to be ILNP-capable for host
   multihoming or site multihoming. This is true provided the
   multihoming is left to individual hosts as described above. In this
   deployment approach, the SBR need only issue Routing Advertisements

   (RAs) that are correct with respect to its upstream connectivity;
   that is, the SBR properly advertises routing prefixes (Locator
   values) to the ILNP hosts.

   In such a scenario, when hosts in the site network see new Locator
   values, and see that a previous Locator value is no longer being
   advertised, those hosts can update their ILCCs, send Locator Updates
   to CNs, and change connectivity as required.

6. Mobility

   ILNP supports mobility directly, rather than relying upon special-
   purpose mobility extensions as is the case with both IPv4 [RFC2002]
   (which was obsoleted by [RFC5944]) and IPv6 [RFC6275].

   There are two different mobility cases to consider:

      a) Host Mobility: individual hosts may be mobile, moving across
         administrative boundaries or topological boundaries within an
         IP-based network, or across the Internet. Such hosts would
         need to independently manage their own mobility.

      b) Network (Site) Mobility: a whole site, i.e., one or more IP
         subnetworks may be mobile, moving across administrative
         boundaries or topological boundaries within an IP-based
         network, or across the Internet. The site as a whole needs to
         maintain consistency in connectivity.

         Essentially, for ILNP, mobility is implemented by enabling:

      a) Locator values to be changed dynamically by a node, including
         for active network-layer sessions.

      b) use of Locator Updates to allow active network-layer sessions
         to be maintained.

      c) for those hosts that expect incoming network-layer or
         transport-layer session requests (e.g., servers), updates to
         the relevant DNS entries for those hosts.

   It is possible that a device is both a mobile host and part of a
   mobile network, e.g., a smartphone in a mobile site network. This is
   supported in ILNP as the mechanism for mobile hosts and mobile
   networks are very similar and work in harmony.

For mobility, there are two general features that must be supported:

   a) Handover (or Hand-off): when a host changes its connectivity
      (e.g., it has a new SNPA as it moves to a new ILNP subnetwork),
      any active network-layer sessions for that host must be
      maintained with minimal disruption (i.e., transparently) to the
      upper-layer protocols.

   b) Rendezvous: when a host that expects incoming network-layer or
      transport-layer session requests has new connectivity (e.g., it
      has a new SNPA as it moves to a new ILNP subnetwork), it needs
      to update its relevant DNS entries so that name resolution will
      provide the correct I and L values to remote nodes.

6.1. Mobility / Multihoming Duality in ILNP

   Mobility and multihoming present the same set of issues for ILNP.
   Indeed, mobility and multihoming form a duality: the set of Locators
   associated with a node or site changes. The reason for the change
   might be different for the case of mobility and multihoming, but the
   effects on the network-layer session state and on correspondents is
   identical.

   With ILNP, mobility and multihoming are supported using a common set
   of mechanisms. In both cases, different Locator values are used to
   identify different IP subnetworks. Also, ILNP nodes that expect
   incoming network-layer or transport-layer session requests are
   assumed to have a Fully Qualified Domain Name (FQDN) stored in the
   Domain Name System (DNS), as is already done within the deployed
   Internet. ILNP mobility normally relies upon the Secure Dynamic DNS
   Update standard for mobile nodes to update their location information
   in the DNS. This approach of using DNS for rendezvous with mobile
   systems was proposed earlier by others [PHG02].

   Host Mobility considers individual hosts that are individually mobile
   -- for example, a mobile telephone carried by a person walking in a
   city. Network (Site) Mobility considers a group of hosts within a
   local topology that move jointly and periodically change their
   uplinks to the rest of the Internet -- for example, a ship that has
   wired connections internally but one or more wireless uplinks to the
   rest of the Internet.

   For ILNP, Host Mobility is analogous to host multihoming (H-MH) and
   Network Mobility is analogous to site multihoming (S-MH). So,
   mobility and multihoming capabilities can be used together, without
   conflict.

6.2. Host Mobility

   With Host Mobility, each individual end-system manages its own
   connectivity through the use of Locator values. (This is very
   similar to the situation described for H-MH in Section 5.1.)

   Let us consider the network in Figure 6.1.

```
       site                            . . . .
      network A                     .          .
      . . . .         +-------+ L_A   .          .
       .       .      |       +------.            .
      .          .    |       |       .            .
     .           .---+ SBR_A |        .            .
     .         .      |       |        .            .
     .  H(1)      .   |       |        .            .
     .         .      +-------+        .            .
      . . . . .                        .            .
       .  H(2) .                       . Internet   .
      . . . . . .                      .            .
     .          .    +-------+ L_B   .              .
     .  H(3)      .  |       +------.              .
     .         .     |       |       .              .
     .          .---+ SBR_B |        .              .
      .         .     |       |        .            .
       .        .     |       |        .            .
        . . . .       +-------+        .          .
         site                         .          .
        network B                      . . . .
```

        H(X) = host H at position X
        L_A = global Locator value A
        L_B = global Locator value B
        SBR = Site Border Router

     Figure 6.1: A Simple Mobile Host Scenario for ILNP

   A host H is at position (1), hence H(1) in a site network A. This
   site network might be, for example, a single radio cell under
   administrative domain A. We assume that the host will move into site
   network B, which might be a single radio cell under administrative
   domain B. We also assume that the site networks have a region of
   overlap so that connectivity can be maintained; else, of course, the
   host will lose connectivity. Also, let us assume that the host
   already has ILNP connectivity in site network A.

If site network A has connectivity via Locator value L_A, and H uses
Identifier value I_H with a single interface ra0, then the host's
ILCC will contain:

```
   [I_H, L_A]                                           --- (18a)
   L_A, ra0                                             --- (18b)
```

Note the equivalence of expressions (18a) and (18b), respectively,
with the expressions (15a) and (16a) for host multihoming.

The host now moves into the overlap region of site networks A and B,
and has position (2), hence H(2) as indicated in Figure 6.1. As this
region is now in site network B, as well as site network A, H should
see RAs from SBR_B for L_B, as well as the RAs for L_A from SBR_A.
The host can now start to use L_B for its connectivity. The host H
must now:

    a) maintain existing active upper-layer sessions: based on its
       current ILCC entries and active sessions, send Locator Update
       (LU) messages to CNs to notify them of the change of L values.
       (LU messages are synonymous to Mobile IPv6 Binding Updates.)

    b) if required, update its relevant DNS entries with the new L
       value in the appropriate DNS records, to enable correct
       resolution for new incoming network-layer or transport-layer
       session requests.

       However, it can opt to do this one of two ways:

    1) immediate handover: the host sends Locator Update (LU) messages
       to CNs, immediately stops using L_A, and switches to using L_B
       only. In this case, its ILCC entries change to:

```
       [I_H, L_B]                                       --- (19a)
       L_B, ra0                                         --- (19b)
```

       There might be packets in flight to H that use L_A, and H MAY
       choose to ignore these on reception.

    2) soft handover: the host sends Locator Update (LU) messages to
       CNS, but it uses both L_A and L_B until (i) it no longer
       receives incoming packets with destination Locator values set
       to L_A within a given time period and (ii) it no longer sees
       RAs for L_A (i.e., it has left the overlap region and so has
       left site network A). In this case, its ILCC entries change
       to:

```
        [I_H, L_A]                                     --- (20a)
        L_A, ra0                                       --- (20b)
        [I_H, L_B]                                     --- (20c)
        L_B, ra0                                       --- (20d)
```

   ILNP does not mandate the use of one handover option over another.
   Indeed, a host may implement both and decide, through local policy or
   other mechanisms (e.g., under the control of a particular transport
   protocol implementation), to use one or other for a specific
   transport-layer session, as required.

   Note that if using soft handover, when in the overlap region, the
   host is multihomed. Also, soft handover is likely to provide a less
   disruptive handover (e.g., lower packet loss) compared to immediate
   handover, all other things being equal.

   There is a case where both the host and its correspondent node are
   mobile. In the unlikely event of simultaneous motion that changes
   both nodes' Locators within a very small time period, there is the
   possibility that communication may be lost. If the communication
   between the nodes was direct (i.e., one node initiated communication
   with another, through a DNS lookup), a node can use the DNS to
   discover the new Locator value(s) for the other node. If the
   communication was through some sort of middlebox providing a relay
   service, then communication is more likely to disrupted only if the
   middlebox is also mobile.

   It is also possible that high packet loss results in Locator Updates
   being lost, which could disrupt handover. However, this is an
   engineering issue and does not impact the basic concept of operation;
   additional discussion on this issue is provided in [RFC6741].

   Of course, for any handover, the new end-to-end path through SBR_B
   might have very different end-to-end path characteristics (e.g.,
   different end-to-end delay, packet loss, throughput). Also, the
   physical connectivity on interface ra0 as well as through SBR_B's
   uplink may be different. Such impacts on end-to-end packet transfer
   are outside the scope of ILNP.

6.3. Network Mobility

   For network mobility, a whole site may be mobile, e.g., the SBRs of
   Figure 6.1 have a radio uplink on a moving vehicle. Within the site,
   individual hosts may or may not be mobile.

   In the simplest case, ILNP deals with mobile networks in the same way
   as for site multihoming: the management of mobility is delegated to
   each host in the site, so it needs to be ILNP-capable. Each host,

effectively, behaves as if it were a mobile host, even though it may
not actually be mobile. Indeed, in this way, the mechanism is very
similar to that for site multihoming. Let us consider the mobile
network in Figure 6.2.

```
        site                          ISP_1
      network          SBR             . . .
      . . . .        +------+ L_1   .       .
      .       .      |   ra1+------.         .
      .         .----+      |        .         .
      .  H   .        |   ra2+--    .         .
      . . . .        +------+         .       .
                                        . . .
       Figure 6.2a: ILNP Mobile Network before Handover


        site                          ISP_1
      network          SBR             . . .
      . . . .        +------+ L_1   .       .
      .       .      |   ra1+------. . . . .
      .         .----+      |        .         .
      .  H   .        |   ra2+------.         .
      . . . .        +------+ L_2   . . . . .
                                      .       .
                                        . . .
                                      ISP_2
        Figure 6.2b: ILNP Mobile Network during Handover



        site                          ISP_2
      network          SBR             . . .
      . . . .        +------+         .       .
      .       .      |   ra1+--    .         .
      .         .----+      |        .         .
      .  H   .        |   ra2+------.         .
      . . . .        +------+         .       .
                                        . . .
       Figure 6.2c: ILNP Mobile Network after Handover

         H = host
       L_1 = global Locator value 1
       L_2 = global Locator value 2
       SBR = Site Border Router
```

     Figure 6.2: A Simple Mobile Network Scenario for ILNP

In Figure 6.2, we assume that the site network is mobile, and the SBR has two radio interfaces ra1 and ra2. However, this particular figure is chosen for simplicity and clarity for our scenario, and other configurations are possible, e.g., a single radio interface which uses separate radio channels (separate carriers, coding channels, etc.). In the figure, ISP_1 and ISP_2 are separate, radio-based service providers, accessible via ra1 and ra2.

In Figure 6.2a, the SBR has connectivity via ISP_1 using Locator value L_1. The host H, with interface ra0 and Identifier I_H, has an established connectivity via the SBR and so has ILCC entries as shown in (21):

```
    [I_H, L_1]                                          --- (21a)
    L_1, ra0                                            --- (21b)
```

Note the equivalence to expressions (18a) and (18b). As the whole network moves, the SBR detects a new radio provider, ISP_2, and connects to it using ra2, as shown in Figure 6.2b, with the service areas of ISP_1 and ISP_2 overlapping. ISP_2 provides Locator L_2, which the SBR advertises into the site network along with L_1. As with the mobile host scenario above, individual hosts may decide to perform immediate handover or soft handover. So, the ILCC state for H will be as for expressions (19a) and (19b) and (20a)-(20d), but with L_1 in place of L_A, and L_2 in place of L_B. Finally, as in Figure 6.2c, the site network moves and is no longer served by ISP_1, and handover is complete. Note that during the handover the site is multihomed, as in Figure 6.2b.

6.4. Mobility Requirements for Site Border Routers

As for multihoming, the SBR does NOT need to be ILNP-capable: it simply needs to advertise the available routing prefixes into the site network. The mobility capability is handled completely by the hosts.

6.5. Mobility with Multiple SBRs

Just as Section 5.3.1 describes the use of multiple routers for multihoming, so it is possible to have multiple routers for mobility for ILNP, for both mobile hosts and mobile networks.

7. IP Security for ILNP

IP Security for ILNP [RFC6741] becomes simpler, in principle, than IPsec as it is today, based on the use of IP Addresses as Identifiers.

An operational issue in the deployed IP Internet is that the IPsec
protocols, AH and ESP, have Security Associations (IPsec SAs) that
include the IP Addresses of the secure IPsec session endpoints. This
was understood to be a problem when AH and ESP were originally
defined in [RFC1825], [RFC1826], and [RFC1827] (which were obsoleted
by [RFC4301], [RFC4302], and [RFC4303]). However, the limited set of
namespaces in the Internet Architecture did not provide any better
choices at that time. ILNP provides more namespaces, thus now
enabling better IPsec architecture and engineering.

7.1. Adapting IP Security for ILNP

In essence, ILNP provides a very simple architectural change to
IPsec: in place of IP Addresses as used today for IPsec SAs, ILNP
uses Node Identifier values instead. Recall that Identifier values
are immutable once in use, so they can be used to maintain end-to-end
state for any protocol that requires it. Note from the discussion
above that the Identifier values for a host remain unchanged when
multihoming and mobility are in use, so IPsec using ILNP can work in
harmony with multihoming and mobility [ABH08b] [ABH09a].

To resolve the issue of IPsec interoperability through a Network
Address Translator (NAT) deployment [RFC1631] [RFC3022], UDP
encapsulation of IPsec [RFC3948] is commonly used as of the date this
document was published. This special-case handling for IPsec traffic
traversing a NAT is not needed with ILNP IPsec.

Further, it would obviate the need for specialised IPsec NAT
traversal mechanisms, thus simplifying IPsec implementations while
enhancing deployability and interoperability [RFC3948].

This architectural change does not reduce the security provided by
the IPsec protocols. In fact, had the Node Identifier namespace
existed back in the early 1990s, IPsec would always have bound to
that location-independent Node Identifier and would not have bound to
IP Addresses.

7.2. Operational Use of IP Security with ILNP

Operationally, this change in SA bindings to use Identifiers rather
than IP Addresses causes problems for the use of the IPsec protocols
through IP Network Address Translation (NAT) devices, with mobile
nodes (because the mobile node's IP Address changes at each network-
layer handoff), and with multihomed nodes (because the network-layer
IPsec session is bound to a particular interface of the multihomed
node, rather than being bound to the node itself) [RFC3027]
[RFC3715].

8. Backwards Compatibility and Incremental Deployment

   ILNPv6 is fully backwards compatible with existing IPv6. No router
   software or silicon changes are necessary to support the proposed
   enhancements. An IPv6 router would be unaware whether the packet
   being forwarded were classic IPv6 or the proposed enhancement in
   ILNPv6. IPv6 Neighbour Discovery will work unchanged for ILNPv6.
   ILNPv6 multicasting is the same as IETF standards-track IPv6
   multicasting.

   ILNPv4 is backwards compatible with existing IPv4. As the IPv4
   address fields are used as 32-bit Locators, using only the address
   prefix bits of the 32-bit space, IPv4 routers also would not require
   changes. An IPv4 router would be unaware whether the packet being
   forwarded were classic IPv4 or the proposed enhancement in ILNPv4
   [RFC6746]. ARP [RFC826] requires enhancements to support ILNPv4
   [RFC6747] [RFC6741]. ILNPv4 multicasting is the same as IETF
   standards-track IPv4 multicasting.

   If a node supports ILNP and intends to receive incoming network-layer
   or transport-layer sessions, the node's Fully Qualified Domain Name
   (FQDN) normally will have one or more NID records and one or more
   Locator (i.e., L32, L64, and/or LP) records associated with the node
   within the DNS [RFC6741] [RFC6742].

   When an IP host ("initiator") initiates a new network-layer session
   with a correspondent ("responder"), it normally will perform a DNS
   lookup to determine the address(es) of the responder. An ILNP host
   normally will look for Node Identifier ("NID") and Locator (i.e.,
   L32, L64, and LP) records in any received DNS replies. DNS servers
   that support NID and Locator (i.e., L32, L64, and LP) records SHOULD
   include them (when they exist) as additional data in all DNS replies
   to queries for DNS AAAA records [RFC6742].

   If the initiator supports ILNP, and from DNS information learns that
   the responder also supports ILNP, then the initiator will generate an
   unpredictable ILNP Nonce value, cache that value locally as part of
   the network-layer ILNP session, and will include the ILNP Nonce value
   in its initial packet(s) to the responder [RFC6741] [RFC6744]
   [RFC6746].

   If the initiator node does not find any ILNP-specific DNS resource
   records for the responder node, then the initiator uses classic IP
   for the new network-layer session with the responder, rather than
   trying to use ILNP for that network-layer session. Of course,
   multiple transport-layer sessions can concurrently share a single
   network-layer (e.g., IP or ILNP) session.

If the responder node for a new network-layer session does not
support ILNP and the responder node receives initial packet(s)
containing the ILNP Nonce, then the responder will drop the packet
and send an ICMP error message back to the initiator. If the
responder node for a new network-layer session supports ILNP and
receives initial packet(s) containing the ILNP Nonce, the responder
learns that ILNP is in use for that network-layer session (i.e., by
the presence of that ILNP Nonce).

If the initiator node using ILNP does not receive a response from the
responder in a timely manner (e.g., within TCP timeout for a TCP
session) and also does not receive an ICMP Unreachable error message
for that packet, OR if the initiator receives an ICMP Parameter
Problem error message for that packet, then the initiator concludes
that the responder does not support ILNP. In this case, the
initiator node SHOULD try again to create the new network-layer
session, but this time using IP (and therefore omitting the ILNP
Nonce).

Finally, since an ILNP node also is a fully capable IP node, the
upgraded node can use any standardised IP mechanisms for
communicating with a legacy IP-only node. So, ILNP will not be worse
than existing IP, but when ILNP is used, the enhanced capabilities
described in these ILNP documents will be available.

9. Security Considerations

This proposal outlines a proposed evolution for the Internet
Architecture to provide improved capabilities. This section
discusses security considerations for this proposal.

Note that ILNP provides security equivalent to IP for similar threats
when similar mitigations (e.g., IPsec or not) are in use. In some
cases, but not all, ILNP exceeds that objective and has lower
security risk than IP. Additional engineering details for several of
these topics can be found in [RFC6741].

9.1. Authentication of Locator Updates

All Locator Update messages are authenticated. ILNP requires use of
an ILNP session nonce [RFC6744] [RFC6746] to prevent off-path
attacks, and also allows use of IPsec cryptography to provide
stronger protection where required.

Ordinary network-layer sessions based on IP are vulnerable to on-path
attacks unless IPsec is used. So the Nonce Destination Option only
seeks to provide protection against off-path attacks on an ILNP-based
network-layer session -- equivalent to ordinary IP-based network-
layer sessions that are not using IPsec.

It is common to have non-symmetric paths between two nodes on the
Internet. To reduce the number of on-path nodes that know the Nonce
value for a given session when ILNP is in use, a nonce value is
unidirectional, not bidirectional. For example, for a network-layer
ILNP-based session between nodes A and B, one nonce value is used
from A to B and a different nonce value is used from B to A.

ILNP sessions operating in higher risk environments SHOULD also use
the cryptographic authentication provided by IPsec *in addition* to
concurrent use of the ILNP Nonce.

It is important to note that, at present, a network-layer IP-based
session is entirely vulnerable to on-path attacks unless IPsec is in
use for that particular IP session, so the security properties of the
new proposal are never worse than for existing IP.

9.2. Forged Identifier Attacks

In the deployed Internet, active attacks using packets with a forged
Source IP Address have been publicly known at least since early 1995
[CA-1995-01]. While these exist in the deployed Internet, they have
not been widespread. This is equivalent to the issue of a forged
Identifier value and demonstrates that this is not a new threat
created by ILNP.

One mitigation for these attacks has been to deploy Source IP Address
filtering [RFC2827] [RFC3704]. Jun Bi at Tsinghua University cites
Arbor Networks as reporting that this mechanism has less than 50%
deployment and cites an MIT analysis indicating that at least 25% of
the deployed Internet permits forged Source IP Addresses.

In [RFC6741], there is a discussion of an accidental use of a
duplicate Identifier on the Internet. However, this sub-section
instead focuses on methods for mitigating attacks based on packets
containing deliberately forged Source Identifier values.

Firstly, the recommendations of [RFC2827] and [RFC3704] remain. So,
any packets that have a forged Locator value can be easily filtered
using existing widely available mechanisms.

Secondly, the receiving node does not blindly accept any packet with
the proper Source Identifier and proper Destination Identifier as an
authentic packet. Instead, each ILNP node maintains an ILNP
Communication Cache (ILCC) for each of its correspondents, as
described in [RFC6741]. Information in the cache is used in
validating received messages and preventing off-path attackers from
succeeding. This process is discussed more in [RFC6741].

Thirdly, any node can distinguish different nodes using the same
Identifier value by other properties of their ILNP sessions. For
example, IPv6 Neighbor Discovery prevents more than one node from
using the same source I-LV at the same time on the same link
[RFC4861]. So, cases of different nodes using the same Identifier
value will involve nodes that have different sets of valid Locator
values. A node thus can demultiplex based on the combination of
Source Locator and Source Identifier if necessary. If IPsec is in
use, the combination of the Source Identifier and the Security
Parameter Index (SPI) value would be sufficient to demux two
different ILNP sessions.

Fourthly, deployments in high-threat environments also SHOULD use
IPsec to authenticate control traffic and data traffic. Because
IPsec for ILNP binds only to the Identifier values, and never to the
Locator values, a mobile or multihomed node can use IPsec even when
its Locator value(s) have just changed.

Lastly, note well that ordinary IPv4, ordinary IPv6, Mobile IPv4, and
also Mobile IPv6 already are vulnerable to forged Identifier and/or
forged IP Address attacks. An attacker on the same link as the
intended victim simply forges the victims MAC address and the
victim's IP Address. With IPv6, when Secure Neighbour Discovery
(SEND) and Cryptographically Generated Addresses (CGAs) are in use,
the victim node can defend its use of its IPv6 address using SEND.
With ILNP, when SEND and CGAs are in use, the victim node also can
defend its use of its IPv6 address using SEND. There are no standard
mechanisms to authenticate ARP messages, so IPv4 is especially
vulnerable to this sort of attack. These attacks also work against
Mobile IPv4 and Mobile IPv6. In fact, when either form of Mobile IP
is in use, there are additional risks, because the attacks work not
only when the attacker has access to the victim's current IP
subnetwork but also when the attacker has access to the victim's home
IP subnetwork. Thus, the risks of using ILNP are not greater than
exist today with IP or Mobile IP.

9.3. IP Security Enhancements

   The IPsec standards are enhanced here by binding IPsec Security
   Associations (SAs) to the Node Identifiers of the endpoints, rather
   than binding IPsec SAs to the IP Addresses of the endpoints as at
   present. This change enhances the deployability and interoperability
   of the IPsec standards, but does not decrease the security provided
   by those protocols. See Section 7 for a more detailed explanation.

9.4. DNS Security

   The DNS enhancements proposed here are entirely compatible with, and
   can be protected using, the existing IETF standards for DNS Security
   [RFC4033]. The Secure DNS Dynamic Update mechanism used here is also
   used unchanged [RFC3007]. So, ILNP does not change the security
   properties of the DNS or of DNS servers.

9.5. Firewall Considerations

   In the proposed new scheme, stateful firewalls are able to
   authenticate ILNP-specific control messages arriving on the external
   interface. This enables more thoughtful handling of ICMP messages by
   firewalls than is commonly the case at present. As the firewall is
   along the path between the communicating nodes, the firewall can
   snoop on the ILNP Nonce being carried in the initial packets of an
   ILNP session. The firewall can verify the correct ILNP Nonce is
   present on incoming control packets, dropping any control packets
   that lack the correct nonce value.

   By always including the ILNP Nonce in ILNP-specific control messages,
   even when IPsec is also in use, the firewall can filter out off-path
   attacks against those ILNP messages without needing to perform
   computationally expensive IPsec processing. In any event, a forged
   packet from an on-path attacker will still be detected when the IPsec
   input processing occurs in the receiving node; this will cause that
   forged packet to be dropped rather than acted upon.

9.6. Neighbour Discovery Authentication

   Nothing in this proposal prevents sites from using the Secure
   Neighbour Discovery (SEND) proposal for authenticating IPv6 Neighbour
   Discovery with ILNPv6 [RFC3971].

9.7. Site Topology Obfuscation

   A site that wishes to obscure its internal topology information MAY
   do so by deploying site border routers that rewrite the Locator
   values for the site as packets enter or leave the site. This
   operational scenario was presented in [ABH09a] and is discussed in
   more detail in [RFC6748].

   For example, a site might choose to use a ULA prefix internally for
   this reason [RFC4193] [ID-ULA]. In this case, the site border
   routers would rewrite the Source Locator of ILNP packets leaving the
   site to a global-scope Locator associated with the site. Also, those
   site border routers would rewrite the Destination Locator of packets
   entering the site from the global-scope Locator to an appropriate
   interior ULA Locator for the destination node [ABH08b] [ABH09a]
   [RFC6748].

10. Privacy Considerations

   ILNP has support for both:

   - Location Privacy: to hide a node's topological location by
     obfuscating the ILNP Locator information. (See also Section 7 of
     [RFC6748].)

   - Identity Privacy: to hide a node's identity by allowing the use of
     Node Identifier values that are not tied to the node in some
     permanent or semi-permanent manner. (See also Section 11 of
     [RFC6741].)

   A more detailed exposition of the possibilities is given in [BAK11].

10.1. Location Privacy

   Some users have concerns about the issue of "location privacy",
   whereby the user's location might be determined by others. The term
   "location privacy" does not have a crisp definition within the
   Internet community at present. Some mean the location of a node
   relative to the Internet's routing topology, while others mean the
   geographic coordinates of the node (i.e., latitude X, longitude Y).
   The concern seems to focus on Internet-enabled devices, most commonly
   handheld devices such as a smartphone, that might have 1:1 mappings
   with individual users.

   There is a fundamental trade-off here. Quality of a node's Internet
   connectivity tends to be inversely proportional to the "location
   privacy" of that node. For example, if a node were to use a router
   with NAT as a privacy proxy, routing all traffic to and from the

Internet via that proxy, then (a) latency will increase as the
distance increases between the node seeking privacy and its proxy,
and (b) communications with the node seeking privacy will be more
vulnerable to communication faults -- both due to the proxy itself
(which might fail) and due to the longer path (which has more points
of potential failure than a more direct path would have).

Any Internet node that wishes for other Internet nodes to be able to
initiate transport-layer or network-layer sessions with it needs to
include associated address (e.g., A, AAAA) or Locator (e.g., L32,
L64, LP) records in the publicly accessible Domain Name System (DNS).
Information placed in the DNS is publicly accessible. Since the goal
of DNS is to distribute information to other Internet nodes, it does
not provide mechanisms for selective privacy. Of course, a node that
does not wish to be contacted need not be present in the DNS.

In some cases, various parties have attempted to create mappings
between IP Address blocks and geographic locations. The quality of
such mappings appears to vary [GUF07]. Many such mapping efforts are
driven themselves by efforts to comply with legal requirements in
various legal jurisdictions. For example, some content providers
reportedly have licenses authorising distribution of content in one
set of locations, but not in a different set of locations.

ILNP does not compromise user location privacy any more than base
IPv6. In fact, by its nature ILNP provides additional choices to the
user to protect their location privacy.

10.2. Identity Privacy

Both ILNP and IPv6 permit use of identifier values generated using
the IPv6 Privacy Address extension [RFC4941]. ILNP and IPv6 also
support a node having multiple unicast addresses/locators at the same
time, which facilitates changing the node's addresses/locators over
time. IPv4 does not have any non-topological identifiers, and many
IPv4 nodes only support one IPv4 unicast address per interface, so
IPv4 is not directly comparable with IPv6 or ILNP.

In normal operation with IPv4, IPv6, or ILNP, a mobile node might
intend to be accessible for new connection attempts from the global
Internet and also might wish to have both optimal routing and maximal
Internet availability, both for sent and received packets. In that
case, the node will want to have its addressing or location
information kept in the DNS and made available to others.

In some cases, a mobile node might only desire to initiate network-
layer or transport-layer sessions with other Internet nodes, and thus
not desire to be a responder, in which case that node need not be

present in the DNS. Some potential correspondent nodes might, as a
matter of local security policy, decline to communicate with nodes
that do not have suitable DNS records present in the DNS. For
example, some deployed IPv4-capable mail relays refuse to communicate
with an initiating node that lacks an appropriate PTR record in the
DNS.

In some cases (for example, intermittent electronic mail access or
browsing specific web pages), support for long-lived network sessions
(i.e., where network-layer session lifetime is longer than the time
the node remains on the same subnetwork) is not required. In those
cases, support for node mobility (i.e., network-layer session
continuity even when the SNPA changes) is not required and need not
be used.

If an ILNP node that is mobile chooses not to use DNS for rendezvous,
yet desires to permit any node on the global Internet to initiate
communications with that node, then that node may fall back to using
Mobile IPv4 or Mobile IPv6 instead.

Many residential broadband Internet users are subject to involuntary
renumbering, usually when their ISP's DHCP server(s) deny a DHCP
RENEW request and instead issue different IP addressing information
to the residential user's device(s). In many cases, such users want
their home server(s) or client(s) to be externally reachable. Such
users today often use Secure DNS Dynamic Update to update their
addressing or location information in the DNS entries, for the
devices they wish to make reachable from the global Internet
[RFC2136] [RFC3007] [LA2006]. This option exists for those users,
whether they use IPv4, IPv6, or ILNP. Users also have the option not
to use such mechanisms.

11. References

11.1. Normative References

   [RFC768]   Postel, J., "User Datagram Protocol", STD 6, RFC 768,
                August 1980.

   [RFC791]   Postel, J., "Internet Protocol", STD 5, RFC 791,
                September 1981.

   [RFC793]   Postel, J., "Transmission Control Protocol", STD 7, RFC
                793, September 1981.

   [RFC826]   Plummer, D., "Ethernet Address Resolution Protocol: Or
               Converting Network Protocol Addresses to 48.bit Ethernet
               Address for Transmission on Ethernet Hardware", STD 37,
               RFC 826, November 1982.

   [RFC2119]   Bradner, S., "Key words for use in RFCs to Indicate
               Requirement Levels", BCP 14, RFC 2119, March 1997.

   [RFC2460]   Deering, S. and R. Hinden, "Internet Protocol, Version 6
               (IPv6) Specification", RFC 2460, December 1998.

   [RFC3007]   Wellington, B., "Secure Domain Name System (DNS) Dynamic
               Update", RFC 3007, November 2000.

   [RFC4033]   Arends, R., Austein, R., Larson, M., Massey, D., and S.
               Rose, "DNS Security Introduction and Requirements", RFC
               4033, March 2005.

   [RFC4861]   Narten, T., Nordmark, E., Simpson, W., and H. Soliman,
               "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861,
               September 2007.

   [RFC6724]   Thaler, D., Ed., Draves, R., Matsumoto, A., and T.
               Chown, "Default Address Selection for Internet Protocol
               Version 6 (IPv6)", RFC 6724, September 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.

11.2. Informative References

   [8+8]   O'Dell, M., "8+8 - An Alternate Addressing Architecture
                  for IPv6", Work in Progress, October 1996.

   [ABH07a]   Atkinson, R., Bhatti, S., and S. Hailes, "Mobility as an
                  Integrated Service Through the Use of Naming",
                  Proceedings of ACM MobiArch 2007, August 2007, Kyoto,
                  Japan.

   [ABH07b]   Atkinson, R., Bhatti, S., and S. Hailes, "A Proposal for
                  Unifying Mobility with Multi-Homing, NAT, & Security",
                  Proceedings of ACM MobiWAC 2007, Chania, Crete. ACM,
                  October 2007.

   [ABH08a]   Atkinson, R., Bhatti, S., and S. Hailes, "Mobility
                  Through Naming: Impact on DNS", Proceedings of ACM
                  MobiArch 2008, August 2008, ACM, Seattle, WA, USA.

   [ABH08b]   Atkinson, R., Bhatti, S., and S. Hailes, "Harmonised
                  Resilience, Security, and Mobility Capability for IP",
                  Proceedings of IEEE Military Communications (MILCOM)
                  Conference, San Diego, CA, USA, November 2008.

   [ABH09a]   Atkinson, R., Bhatti, S., and S. Hailes, "Site-
                  Controlled Secure Multi-Homing and Traffic Engineering
                  For IP", Proceedings of IEEE Military Communications
                  (MILCOM) Conference, Boston, MA, USA, October 2009.

   [ABH09b]   Atkinson, R., Bhatti, S., and S. Hailes, "ILNP:
                  Mobility, Multi-Homing, Localised Addressing and
                  Security Through Naming", Telecommunications Systems,
                  Volume 42, Number 3-4, pp. 273-291, Springer-Verlag,
                  December 2009, ISSN 1018-4864.

   [ABH10]   Atkinson, R., Bhatti, S., S. Hailes, "Evolving the
                  Internet Architecture Through Naming", IEEE Journal on
                  Selected Areas in Communication (JSAC), vol. 28, no. 8,
                  pp. 1319-1325, IEEE, Piscataway, NJ, USA, Oct 2010.

   [BA11]   Bhatti, S. and R. Atkinson, "Reducing DNS Caching",
                  Proceedings of IEEE Global Internet Symposium (GI2011),
                  Shanghai, P.R. China, 15 April 2011.

   [BA12]   Bhatti, S. and R. Atkinson, "Secure and 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), Baltimore, MD, USA,
               November 2011.

   [CA-1995-01] US CERT, "IP Spoofing Attacks and Hijacked Terminal
               Connections", CERT Advisory 1995-01, Issued 23 Jan 1995,
               Revised 23 Sep 1997.

   [DIA]   Defense Intelligence Agency, "Compartmented Mode
               Workstation Specification", Technical Report
               DDS-2600-6243-87, US Defense Intelligence Agency,
               Bolling AFB, DC, USA.

   [DoD85]   US National Computer Security Center, "Department of
               Defense Trusted Computer System Evaluation Criteria",
               DoD 5200.28-STD, US Department of Defense, Ft. Meade,
               MD, USA, December 1985.

   [DoD87]   US National Computer Security Center, "Trusted Network
               Interpretation of the Trusted Computer System Evaluation
               Criteria", NCSC-TG-005, Version 1, US Department of
               Defense, Ft. Meade, MD, USA, 31 July 1987.

   [GSE]   O'Dell, M., "GSE - An Alternate Addressing Architecture
               for IPv6", Work in Progress, February 1997.

   [GUF07]   Gueye, B., Uhlig, S., and S. Fdida, "Investigating the
               Imprecision of IP Block-Based Geolocation", Lecture
               Notes in Computer Science, Volume 4427, pp. 237-240,
               Springer-Verlag, Heidelberg, Germany, 2007.

   [ID-ULA]   Hinden, R., Huston, G., and T. Narten, "Centrally
               Assigned Unique Local IPv6 Unicast Addresses", Work in
               Progress, June 2007.

   [IEEE-EUI] IEEE, "Guidelines for 64-bit Global Identifier (EUI-64)
               Registration Authority", Piscataway, NJ, USA, March
               1997, <http://standards.ieee.org/regauth/oui/tutorials/
               EUI64.html>.

   [IEN1]   Bennett, C., Edge, S., and A. Hinchley, "Issues in the
                  Interconnection of Datagram Networks", Internet
                  Experiment Note (IEN) 1, INDRA Note 637, PSPWN 76, 29
                  July 1977, <http://www.rfc-editor.org/ien/ien1.pdf>.

   [IEN19]   Shoch, J., "Inter-Network Naming, Addressing, and
                  Routing", IEN 19, January 1978,
                  <http://www.rfc-editor.org/ien/ien19.txt>.

   [IEN23]   Cohen, D., "On Names, Addresses, and Routings", IEN 23,
                  January 1978, <http://www.rfc-editor.org/ien/ien23.pdf>.

   [IEN31]   Cohen, D., "On Names, Addresses, and Routings (II)", IEN
                  31, April 1978,
                  <http://www.rfc-editor.org/ien/ien31.pdf>.

   [IEN135]   Sunshine, C. and J. Postel, "Addressing Mobile Hosts in
                  the ARPA Internet Environment", IEN 135, March 1980,
                  <http://www.rfc-editor.org/ien/ien135.pdf>.

   [IPng95]   Clark, D., "A thought on addressing", electronic mail
                  message to IETF IPng WG, Message-ID:
                  9501111901.AA28426@caraway.lcs.mit.edu, Laboratory for
                  Computer Science, MIT, Cambridge, MA, USA, 11 January
                  1995.

   [LA2006]   Liu, C. and P. Albitz, "DNS & Bind", 5th Edition,
                  O'Reilly & Associates, Sebastopol, CA, USA, May 2006,
                  ISBN 0-596-10057-4.

   [LABH06]   Lad, M., Atkinson, R., Bhatti, S., and S. Hailes, "A
                  Proposal for Coalition Networking in Dynamic Operational
                  Environments", Proceedings of IEEE Military
                  Communications Conference, Washington, DC, USA, Nov.
                  2006.

   [PHG02]   Pappas, A., Hailes, S., and R. Giaffreda, "Mobile Host
                  Location Tracking through DNS", Proceedings of IEEE
                  London Communications Symposium, IEEE, London, England,
                  UK, September 2002.

   [RAB09]   Rehunathan, D., Atkinson, R., and S. Bhatti, "Enabling
                  Mobile Networks Through Secure Naming", Proceedings of
                  IEEE Military Communications Conference (MILCOM),
                  Boston, MA, USA, October 2009.

   [RB10]   Rehunathan, D. and S. Bhatti, "A Comparative Assessment
                  of Routing for Mobile Networks", Proceedings of IEEE
                  International Conference on Wireless and Mobile
                  Computing Networking and Communications (WiMob), IEEE,
                  Niagara Falls, ON, Canada, Oct. 2010.

   [SBK01]   Snoeren, A., Balakrishnan, H., and M. Frans Kaashoek,
                  "Reconsidering Internet Mobility", Proceedings of 8th
                  Workshop on Hot Topics in Operating Systems, IEEE,
                  Elmau, Germany, May 2001.

   [SIPP94]   Smart, B., "Re: IPng Directorate meeting in Chicago;
                  possible SIPP changes", electronic mail to the IETF SIPP
                  WG mailing list, Message-ID:
                  199406020647.AA09887@shark.mel.dit.csiro.au,
                  Commonwealth Scientific & Industrial Research
                  Organisation (CSIRO), Melbourne, VIC, 3001, Australia, 2
                  June 1994.

   [SRC84]   Saltzer, J., Reed, D., and D. Clark, "End to End
                  Arguments in System Design", ACM Transactions on
                  Computer Systems, Volume 2, Number 4, ACM, New York, NY,
                  USA, November 1984.

   [RFC814]   Clark, D., "Name, addresses, ports, and routes", RFC
                  814, July 1982.

   [RFC1112]   Deering, S., "Host extensions for IP multicasting", STD
                  5, RFC 1112, August 1989.

   [RFC1122]   Braden, R., Ed., "Requirements for Internet Hosts -
                  Communication Layers", STD 3, RFC 1122, October 1989.

   [RFC1498]   Saltzer, J., "On the Naming and Binding of Network
                  Destinations", RFC 1498, August 1993.

   [RFC1631]   Egevang, K. and P. Francis, "The IP Network Address
                  Translator (NAT)", RFC 1631, May 1994.

   [RFC1825]   Atkinson, R., "Security Architecture for the Internet
                  Protocol", RFC 1825, August 1995.

   [RFC1826]   Atkinson, R., "IP Authentication Header", RFC 1826,
                  August 1995.

   [RFC1827]   Atkinson, R., "IP Encapsulating Security Payload (ESP)",
                  RFC 1827, August 1995.

   [RFC1958]   Carpenter, B., Ed., "Architectural Principles of the
               Internet", RFC 1958, June 1996.

   [RFC1992]   Castineyra, I., Chiappa, N., and M. Steenstrup, "The
               Nimrod Routing Architecture", RFC 1992, August 1996.

   [RFC2002]   Perkins, C., Ed., "IP Mobility Support", RFC 2002,
               October 1996.

   [RFC2101]   Carpenter, B., Crowcroft, J., and Y. Rekhter, "IPv4
               Address Behaviour Today", RFC 2101, February 1997.

   [RFC2136]   Vixie, P., Ed., Thomson, S., Rekhter, Y., and J. Bound,
               "Dynamic Updates in the Domain Name System (DNS
               UPDATE)", RFC 2136, April 1997.

   [RFC2710]   Deering, S., Fenner, W., and B. Haberman, "Multicast
               Listener Discovery (MLD) for IPv6", RFC 2710, October
               1999.

   [RFC2827]   Ferguson, P. and D. Senie, "Network Ingress Filtering:
               Defeating Denial of Service Attacks which employ IP
               Source Address Spoofing", BCP 38, RFC 2827, May 2000.

   [RFC2956]   Kaat, M., "Overview of 1999 IAB Network Layer Workshop",
               RFC 2956, October 2000.

   [RFC3022]   Srisuresh, P. and K. Egevang, "Traditional IP Network
               Address Translator (Traditional NAT)", RFC 3022, January
               2001.

   [RFC3027]   Holdrege, M. and P. Srisuresh, "Protocol Complications
               with the IP Network Address Translator", RFC 3027,
               January 2001.

   [RFC3177]   IAB and IESG, "IAB/IESG Recommendations on IPv6 Address
               Allocations to Sites", RFC 3177, September 2001.

   [RFC3376]   Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A.
               Thyagarajan, "Internet Group Management Protocol,
               Version 3", RFC 3376, October 2002.

   [RFC3704]   Baker, F. and P. Savola, "Ingress Filtering for
               Multihomed Networks", BCP 84, RFC 3704, March 2004.

   [RFC3715]   Aboba, B. and W. Dixon, "IPsec-Network Address
               Translation (NAT) Compatibility Requirements", RFC 3715,
               March 2004.

   [RFC3810]   Vida, R., Ed., and L. Costa, Ed., "Multicast Listener
                Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June
                2004.

   [RFC3948]   Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and
                M. Stenberg, "UDP Encapsulation of IPsec ESP Packets",
                RFC 3948, January 2005.

   [RFC3971]   Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander,
                "SEcure Neighbor Discovery (SEND)", RFC 3971, March
                2005.

   [RFC3972]   Aura, T., "Cryptographically Generated Addresses (CGA)",
                RFC 3972, March 2005.

   [RFC4193]   Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast
                Addresses", RFC 4193, October 2005.

   [RFC4291]   Hinden, R. and S. Deering, "IP Version 6 Addressing
                Architecture", RFC 4291, February 2006.

   [RFC4581]   Bagnulo, M. and J. Arkko, "Cryptographically Generated
                Addresses (CGA) Extension Field Format", RFC 4581,
                October 2006.

   [RFC4941]   Narten, T., Draves, R., and S. Krishnan, "Privacy
                Extensions for Stateless Address Autoconfiguration in
                IPv6", RFC 4941, September 2007.

   [RFC4982]   Bagnulo, M. and J. Arkko, "Support for Multiple Hash
                Algorithms in Cryptographically Generated Addresses
                (CGAs)", RFC 4982, 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.

   [RFC5061]   Stewart, R., Xie, Q., Tuexen, M., Maruyama, S., and M.
                Kozuka, "Stream Control Transmission Protocol (SCTP)
                Dynamic Address Reconfiguration", RFC 5061, September
                2007.

   [RFC5570]   StJohns, M., Atkinson, R., and G. Thomas, "Common
                Architecture Label IPv6 Security Option (CALIPSO)", RFC
                5570, July 2009.

   [RFC6177]   Narten, T., Huston, G., and L. Roberts, "IPv6 Address
                Assignment to End Sites", BCP 157, RFC 6177, March 2011.

   [RFC6250]   Thaler, D., "Evolution of the IP Model", RFC 6250, May
                2011.

   [RFC6275]   Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility
                Support in IPv6", RFC 6275, July 2011.

   [RFC6748]   Atkinson, R. and S. Bhatti, "Optional Advanced
                Deployment Scenarios for the Identifier-Locator Network
                Protocol (ILNP)", RFC 6748, November 2012.

12. Acknowledgements

Authors' Addresses

   RJ Atkinson
   Consultant
   San Jose, CA 95125
   USA

   EMail: rja.lists@gmail.com

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

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