An Update on the Identifier-Locator Network Protocol (ILNP)

Presented by Steve Blake
sblake@extremenetworks.com

Viewgraphs by Ran Atkinson
rja@extremenetworks.com

“Standing on the Shoulders of Giants”

- Computer Science sometimes has been accused of blindly reinventing the wheel.
- We actively tried to avoid that, so credit to:
  - Dave Clark for (c.1995) email to a public mailing list proposing to split the IP address into two pieces.
  - Mike O’Dell for two early proposals (8+8, GSE), in the1990s.
  - The IRTF Name Space RG (NSRG), c. 1999-2002.
- This work extends and enhances those early ideas:
  - Like HIP, this work dates back to the author’s participation in the IRTF NSRG early this decade.

Architectural Claim

If we provide a richer set of namespaces then the Internet Architecture can better support mobility, multi-homing, and other important capabilities:

- provide a broader set of namespaces than at present.
- reduce/eliminate names with overloaded semantics.
- provide crisp semantics for each type of name.

Routing RG Issues

Routing RG Charter

- The Routing RG Charter explicitly lists these four challenges:
  - Scalability
  - Multi-homing
  - Mobility
  - Traffic Engineering

Scalability

- Growth in prefixes inside the Default Free Zone (DFZ) is at least geometric at present.
- Primary cause is growth in site multi-homing, which is also at least geometric at present.
- Primary goal of multi-homed sites is higher availability.
- Important reference for the above data:
    - "IPv4 Address Allocation & the BGP Routing Table Evolution" by X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, & L. Zhang, ACM Computer Communications Review, 2005.

Multi-Homing

- A fundamental issue is that current site multi-homing creates additional entropy in the DFZ RIB/FIB
- Why ?
  - We multi-home sites using Longest Prefix Match
  - Each multi-homed site adds more-specific prefixes to DFZ
- Why this approach for multi-homing ?
  - Transport-layer pseudo-header checksums include location information, not just host identity
- The real fix is to de-couple the transport protocol state from the network location.

Mobility

- Actually, mobility is just highly dynamic multi-homing
  - Want transport-layer session(s) to remain up
  - But want to change the network location of participant(s)
- Again, the cleanest fix is to de-couple the transport session state from the network location(s)
  - Mobile IP{v4, v6} try to hide the real network location through Home Address, Tunnelling, and other mechanisms.
    - Mobile IP WG assumed that one could not change the architecture.
    - ILNP assumes the architecture can be changed.
- Also, consider that mobile nodes/sites might not have any home location.
  - This suggests the use of agent-less mobility approaches

Traffic Engineering

- Traffic Engineering (TE) is another cause for de-aggregated IP routing prefixes.
  - ISPs like to use the routing prefixes to move some traffic away from a congested link or path.
  - Some content providers use the routing prefixes as part of a sophisticated multi-site server load-balancing schema
  - Some sites implement local TE policies in their border routers
    - Site routers prefer lowest-cost (or lowest latency, or some other locally interesting metric) upstream provider for traffic leaving the site.
- TE is an important capability to retain.
- TE is not the dominant source of RIB/FIB entropy.

Heresy

- The Internet's Routing Architecture is just fine.
- The problem is that we are (ab)using routing to work-around limitations in the Internet's Naming Architecture.
- If we can sort out the Naming Architecture, then
  - existing routing protocols don't need to change
  - existing techniques don't need to change.

ILNP:
An 8+8 Approach

What is 8+8 ?

- 1) Name of an addressing architecture that split the IP address into a separate Locator and Identifier.
  - from Mike O'Dell in the middle 1990s.
- 2) An specific proposal on how to enhance IPv6; sometimes this is also called "GSE".
  - Also from Mike O'Dell in the 1990s
- 3) A class of IP architectures that is based on the original concept from (1) above
  - In this talk, we are using definition (3) just above.

The 8+8 Architecture

- Separate the high-order bits ("Routing Prefix") of an IPv6 address into a Locator field, 64 bits wide.
- Separate the low-order bits of an IPv6 address into an Identifier field, 64 bits wide.
- Transport session state contains *only* the Identifier.
- IP packet forwarding/routing uses *only* the Locator.
- One can imagine a range of networking protocols, different in various details, that use this architecture.

ILNPv6

- We propose an set of enhancements to IPv6, which we call **ILNPv6**:
  - provides full backwards compatibility with IPv6.
  - provides full support for incremental deployment.
  - **IPv6 routers do not need to change**.
- ILNPv6 “splits” the IPv6 address in half:
  - **Locator (L)**: 64-bit name for the subnetwork
  - **Identifier (I)**: 64-bit name for the host
- Same architecture can work for IPv4 (ILNPv4),
  - but a shortage of bits makes the engineering ugly

IPv6 Packet Header

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version| Traffic Class |           Flow Label                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Payload Length        |  Next Hdr     |   Hop Limit   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                                                               |
+-+-                     Source Address                      -+-+
|                                                               |
+                                                               +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                                                               |
+-+-                   Destination Address                   -+-+
|                                                               |
+                                                               +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

ILNPv6 Packet Header

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version| Traffic Class |           Flow Label                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Payload Length        |  Next Hdr     |   Hop Limit   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                       Source Locator                          +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                     Source Identifier                         +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                     Destination Locator                       +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                   Destination Identifier                      +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Locators vs. Identifiers

- **Locator (L):**
  - uses the existing "Routing Prefix" bits of an IPv6 address.
  - names a single subnetwork (/48 allows subnetting).
  - **topologically significant, so the value of L changes as subnetwork connectivity changes.**
  - only used for routing and forwarding.
- **Identifier (I):**
  - Replaces the existing "Interface ID" bits of an IPv6 address
  - **Names a** (physical/logical/virtual) **host, not an interface.**
  - Remains constant even if connectivity/topology changes.
  - uses IEEE EUI-64 syntax, which is the same as IPv6:
  - only used by transport-layer (and above) protocols.

A Bit More Detail

- All ILNP nodes:
  - have 1 or more Identifiers at a time.
  - Identifiers are independent of the network interface
  - only **Identifiers** are used at the **Transport-Layer** or above.
  - have 1 or more Locators at a time.
  - only **Locators** are used to **route/forward** packets.

- An ILNP "node" might be:
  - a single physical machine,
  - a virtual machine,
  - or a distributed system.

Generating Identifiers

- IEEE EUI-64 format
  - EUI-64 includes 1 bit for multicast/unicast
    - A Group ID sets this bit to “multicast”
    - A Node ID sets this bit to “unicast”
  - EUI-64 includes 1 bit for global-scope/local-scope
    - Global-scope means the other bits were derived from an IEEE MAC.
    - Normally, a node would generate its ID(s) by itself in this way.
    - No need to use IPv6 Duplicate Address Detection (DAD) if Global-scope ID.
- If scope bit is **local**, have 62 bits that can be **anything**:
  - Cryptographically Generated Identifier (a la CGA proposals)
  - Hash of a public-key (a la HIP)
  - Pseudo-randomly generated (a la IPv6 Privacy AutoConf)

Naming Comparison

Protocol Layer	IP	ILNP
Application	FQDN or IP address	FQDN
Transport	IP address (+ port number)	Identifier (+ port number)
Network	IP address	Locator
Link	MAC address	MAC address

ILNP:
Transport Layer Changes

- CRITICAL CHANGE:
  - Transport-layer pseudo-header only includes IDENTIFIER, never the LOCATOR.
- IMPLICATIONS:
  - We can multi-home nodes/sites without impacting routing.
  - Mobility just became a built-in/native capability.
  - Need a way to tell correspondents when we move
  - Historically, IETF concerned about authenticating location changes and providing equivalent security to current IPv6

Security Mechanisms

- IP Security with ILNP:
  - can use IPsec AH and ESP for cryptographic protection.
  - ILNP AH includes I values, but excludes L values.
  - IPsec Security Association (SA) bound to value of I, not L.
- New IPv6 Destination Option - ILNP Nonce:
  - contains clear-text 48-bit or 96-bit unpredictable nonce value
  - protects against off-path attacks on a session (child proof)
    - Existing IPv4/IPv6 without IPsec is vulnerable to on-path attacks
    - Nonce use is both affordable & provides equivalent protection as today
  - primarily used to authenticate control traffic:
    - e.g. ICMP Locator Update (LU) message
- Existing IETF DNS Security mechanisms; no changes.

ILNP: DNS Enhancements

- New resource records (forward lookups)
  - **I**: Identifier(s), unsigned 64-it value, EUI-64 syntax.
  - **L**: Locator(s), unsigned 64-bit value, topological.
  - Each of these has a preference value, as with MX records.
  - Nota Bene: DNS permits per-resource-record TTL values.
    - Expect I values to be relatively longer-lived in all cases.
    - Expect L values to be relatively shorter-lived if mobile/multihomed.
- One (optional) performance optimisation
  - **LP**: Locator Pointer; points to an L record.
  - Also has a preference value.
  - Can have a longish DNS TTL, so value can be cached.
- Reverse lookups can work as they do today

DNS Locator Pointer Record

- When an entire network moves together, there might be many L record updates for the DNS at once.
- As a DNS optimisation, we add the Locator Pointer (LP) record:
  - LP record points to the FQDN associated with an L record
  - If DNS lookup yields an LP record, then one needs to perform L record lookup using FQDN provided by LP record response.
  - FQDN/LP associated with a subnetwork, not a single host
  - LP record just adds one additional level of indirection. :-)
- DNS Security works as usual.
- Entirely optional to deploy.

DNS Enhancements

NAME	DNS Type	Definition
Identifier	**I**	**Names a Node**
Locator	**L**	**Names a subnetwork**
Locator Pointer	**LP**	**Forward pointer from FQDN to an L Record**

Generating a Packet

- Source performs DNS lookup on destination's FQDN.
- Source learns the set of I and L values for destination.
  - Like MX records, I and L records have preference values.
  - All valid I and L records are stored in local session cache
- Source selects the Source Locator and the Source ID to use for its own packet(s) to this destination.
- Source selects the Destination Locator and Destination ID to use.
- Source creates the packet and sends it out.

Mobility Approach

Naming and Mobility

- With MIP (v4 and v6), IP addresses retain their dual role, used for both **location** and **identity**:
  - overloaded semantics creates complexity, since all IP addresses are (potentially) topologically significant.
- With ILNP, identity and location are separate:
  - **new Locator used as node moves:**
    - reduces complexity: only Locator changes value.
  - **constant Identifier as node moves:**
    - agents not needed and triangle routing never occurs.
  - **upper-layer state (e.g. TCP, UDP) only uses Identifier.**
    - Recall that an Identifier names a node, not an interface.

Mobility has 2 Primary Aspects

- 1) Rendezvous
  - How initially to find a node's location to start a new session
- 2) Location Updates
  - How to maintain existing communications sessions as one or more end nodes for that session change location
- ILNP uses DNS for initial rendezvous, just as today.
- ILNP primarily uses control traffic for updates,
  - can fall back to DNS if necessary.

Mobility Implementation

- Implementation in correspondent node:
  - uses DNS to find MN's set of Identifiers and Locators.
  - only uses Identifier(s) in transport-layer session state.
  - uses Locator(s) only to forward/route packets.
- Implementation in mobile node (MN):
  - accepts new sessions using currently valid I values.
  - With ILNPv6, when the MN moves:
    - MN uses ICMP Locator Update (LU) to inform other nodes of the revised set of Locators for the MN.
    - LU can be authenticated via IP Security (or Nonce).
    - MN uses Secure Dynamic DNS Update (RFC-3007) to revise the affected DNS resource records in its Authoritative DNS server(s).

ILNPv6 Network Handoff

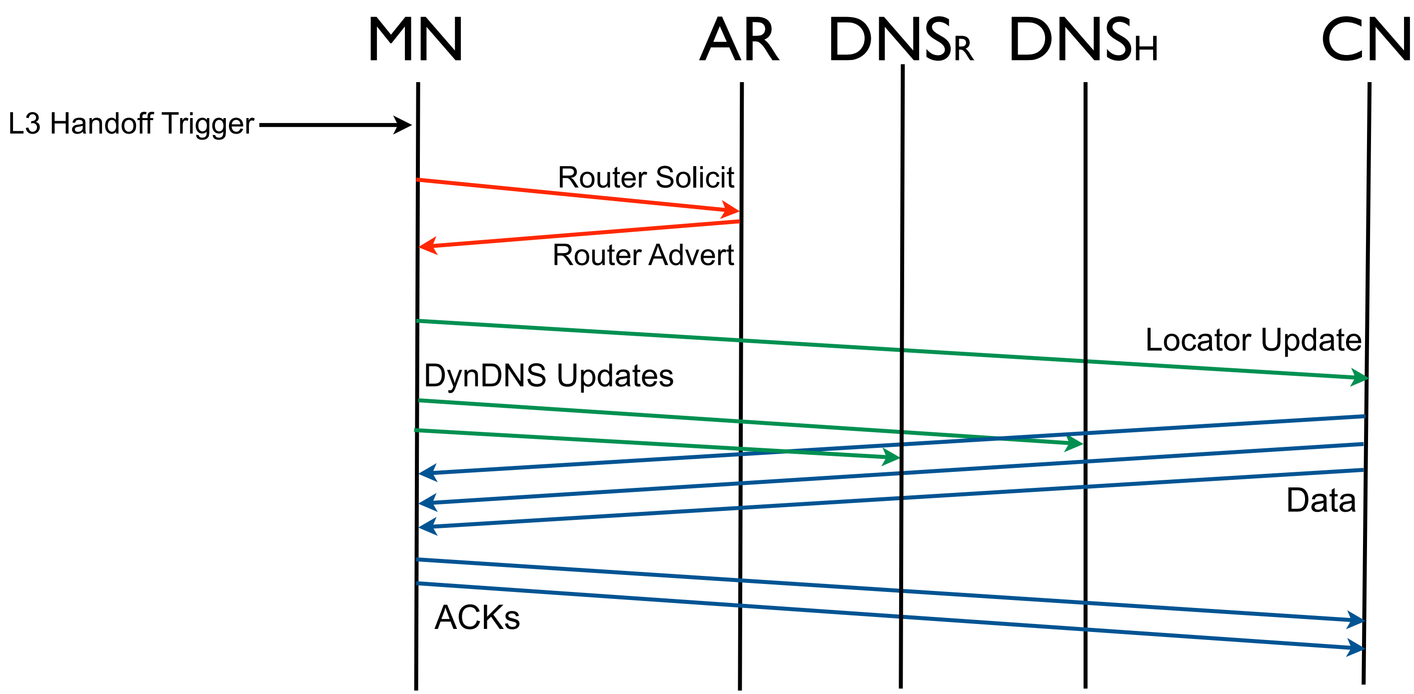

MN	Mobile Node
AR	Router serving MN
DNSR	DNS Server (reverse)
DNSH	DNS Server (forward)
CN	Correspondent Node

Multi-Homing

Multi-Homing Today

- Site Multi-Homing
  - widely used today, growing rapidly in popularity
    - primary driver appears to be network availability
  - handled today by adding more specific prefixes into DFZ
    - each multi-homed site adds 3 or more prefixes into BGP and DFZ RIB
  - main source of DFZ RIB and BGP scaling issues & entropy
- Host Multi-homing
  - traditional deployments can improve initial availability.
  - traditional deployments can't provide invisible session failover to another interface if some fault occurs.
  - Routing prefix length rules (/24 or shorter) limit its usefulness for session continuity and failover/recovery.

Multi-Homing with ILNP

- ILNP supports both site multi-homing & host multi-homing – and provides resilience/availability for both.
- ICMP Locator Update mechanism handles uplink changes (e.g. fibre cut/repair).
- ILNP reduces size of RIB & FIB in DFZ:
  - more-specific routing prefixes are no longer used for this.
- In turn, this greatly helps with BGP scalability.
- New optional DNS Locator Pointer (LP) record can enhance DNS scalability (e.g. for site multi-homing).
- Same approach also supports mobile networks.

ILNPv6: “NAT” Integration

- IP Address Translation (NAT/NAPT) is here to stay:
  - many residential IP gateways use NAT or NAPT.
  - often-requested feature for IPv6 routers is NAT/NAPT.
- ILNPv6 reduces issues with these deployments:
  - With ILNPv6, we have “Locator Translation”, instead.
  - Identifiers don’t change when Locators are translated.
  - Upper-layer protocol state is bound to I only, never to L.
  - Translation is now invisible to upper-layer protocols.
- ILNPv6 IPsec is not affected by NAT:
  - Security Association is bound to Identifiers, not Locators.
  - ILNP AH covers Identifiers, but does not cover Locators.
  - ILNP IPsec and “NAT” work fine together (w/o extra code)

Multicasting

- Multicasting works essentially the same as today
- Implications of the Locator/Identifier split:
  - Destination Identifier in a multicast packet is a Group ID, not a Node ID. Existing EUI-64 “multicast” flag is set.
  - Source Identifier in a multicast packet is sender’s Node ID.
- This change facilitates multicast traversal of NA(P)T boxes, other middle boxes, and also facilitates IPsec.
  - Session state now can be bound to the Sender’s Node ID, Destination’s Group ID (SID, GID), not to a network location.

Traffic Engineering

- For site traffic engineering, several approaches could be used. This describes one approach.
- Can translate Source Locator (and/or Destination Locator) in the site border router
  - Upon egress, router could modify the Source Locator to a value preferred by the site's routing policy.
  - This provides Recipient nodes with a hint about which Locator to use in reply packets.
  - Identifiers are not modified during transit.
  - Rewriting Destination Locator permitted if the router (somehow) knows a better Locator to use.
  - Does not require (or prohibit) use of split-horizon DNS

Transition Considerations

Applications & APIs

- ILNP
  - does not require any API changes.
  - works with existing applications over existing APIs.
- As with SHIM6, location changes can be hidden from the application, and kept below the BSD Sockets API.
  - This preserves the value in dynamic Locator changes.
- For referrals, several options exist:
  - Fully-Qualified Domain Names always work with ILNP.
  - IP Address referrals aren't completely reliable in the current deployed Internet today.
  - 128-bit values (Locator + Identifier) mostly work fine with ILNP, because of server load-balancers and static servers
  - Also: see Brian Carpenter's recent I-D on referrals

Incremental Deployment

- ILNPv6 is a set of extensions to IPv6.
- No changes to IPv6 routers are needed.
- Implications:
    - Existing IPv6 networks already support ILNPv6 packets.
    - No upgrades needed to routers.
- Incentives exist to upgrade host IPv6 stacks to ILNPv6
    - Users gain immediate benefits when they upgrade.
        - Example Benefits: Host Multi-homing, Site Multi-homing, improved Mobility, NAT tolerance, etc.
    - Benefits grow as more nodes upgrade.

Backward Compatibility

- How does an initiating node know whether the remote node is ILNPv6 enabled or not?
  - ILNPv6 DNS records (I, L) also will be returned on DNS lookup for A/AAAA as "Additional Data"
- How does a responding node know whether the remote node is ILNPv6 enabled or not ?
  - ILNPv6 Nonce is present in received packet from remote node that is initiating a new UDP/TCP/SCTP session.
- If either node doesn't support ILNPv6, the other node falls back to using existing ordinary IPv6.
- No loss of connectivity/reachability during evolution.

Deployment Incentives

- Many benefits can be gained incrementally
  - Particularly in mobility, multi-homing, & resilience.
  - So users have incentives to upgrade from IPv6 to ILNPv6
- No changes are needed to IPv6 routers/routing
  - So backbones won't gate/impede user deployments
- ILNP restores the Internet's "smart host" model
  - So OS implementers have incentives to offer upgrades
- Many nodes likely need to be upgraded for a major reduction in DFZ RIB/FIB entropy
  - So, operators have incentives to encourage upgrading

Summary

ILNP: Integrated Solution

- Mobility support is fully integrated, not optional.
  - mobility is native capability.
  - mobility mechanisms are much simpler.
  - authentication is practical to deploy.
- Multi-homing and mobile network supported
  - supports dynamic multi-homing for hosts and networks.
  - supports mobile networks natively.
  - multi-homing also integrated with mobility.
  - routing scalability (BGP, DFZ RIB) is greatly improved.
- Locator translation support ("NAT") is integrated.
- IPsec support is integrated.

Conclusion

- ILNP treats the IP Address as consisting of separate Identifier & Locator values.
- This enables native Mobility (without agents).
- Also, Multi-Homing, NAT, and Security are well integrated with Mobility.
- Incrementally Deployable & Backwards Compatible
- Improvements in the Naming Architecture enable simpler protocol approaches and ILNP is consistent with the wider goals of the future direction of the Internet architecture.

Thank you!

- Several Internet-Drafts exist.
- Updated I-Ds are coming soon.
- Various research papers are also available.
- For more information, please contact:
  - Ran Atkinson rja@extremenetworks.com

Backup Slides

ILNPv6: No Free Lunch

- No globally-routable network interface name:
  - potential impact on SNMP MIBs, e.g. to get interface counters form a particular interface.
- A few legacy apps might remain problematic, not sure yet.
  - Probably should test with FTP
- DNS reliance is not new, but is more explicit:
  - at present, most users perceive “DNS fault” as “network down”.
  - ILNP creates no new DNS security issues.
  - Existing IETF DNS standards work fine without alteration.
  - Both DNSsec and Secure Dynamic DNS Update are widely available in commercial products and also in free software.

Some Existing Namespaces

- IP Address
  - 128.60.80.2
- IP Subnetwork
  - 128.60.80.0/24
- Domain Name
  - itd.nrl.navy.mil
- Communication Endpoint ("Socket")
  - TCP port 25 at itd.nrl.navy.mil
- Mailbox
  - username@itd.nrl.navy.mil
- URL
  - http://www.itd.nrl.navy.mil/index.html

Network Realms

(Scoped Addressing & “NAT”)

NAPT Basics

- Network Address & Port Translation (NAPT)
- Variant of Network Address Translation (NAT)
  - Alters IP addresses
  - Alters TCP/UDP/SCTP port numbers
  - Can multiplex a network behind 1 public IP address
- Question: Does NAPT break ILNP or not ?

NAPT: Rendezvous Issue

- Many sites deploy either NAT or NAPT for perceived security advantages:
  - Primarily: remote notes are blocked from initiating sessions with hosts inside the NAT/NAPT gateway.
  - This can affect some applications (e.g. Video Conferencing, VoIP).
  - ILNP does not change this “security” property, which is good for sites that deploy NAPT for this reason.
- Some sites might deploy NAT or NAPT to get address portability or to conserve addresses:
  - Neither issue exists in an IPv6/ILNPv6 context because of the much larger IPv6 address space & because ILNP handles renumbering/multi-homing natively.
  - So neither reason exists in an IPv6/ILNPv6 context.

NAPT Scenario

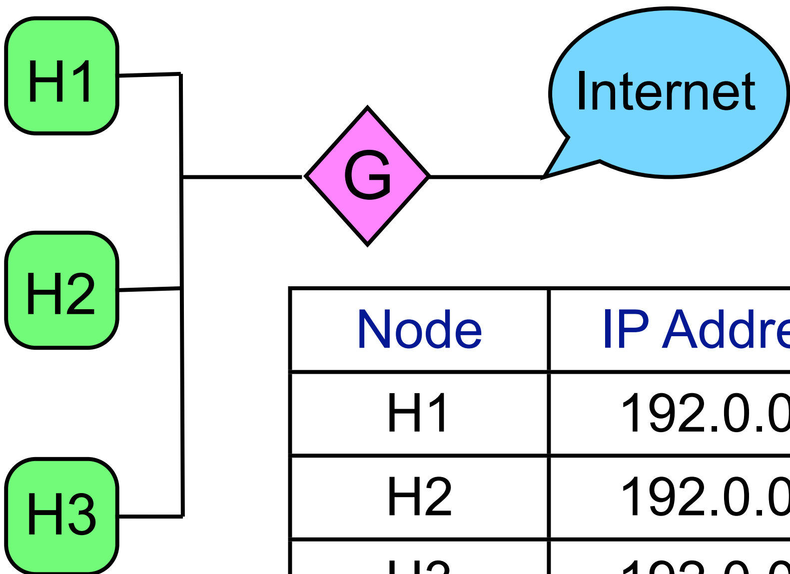

Node	IP Address	Port range
H1	192.0.0.2	5100-5199
H2	192.0.0.3	5200-5299
H3	192.0.0.4	5300-5399
G1	192.0.0.1	5400-5499
G1 (public)	3.1.2.3	-

- G1 uses its 1 public IP address to handle traffic to/from The Internet for itself and hosts H1, H2, & H3 behind G1.
- So, G1 is using NAPT and has different TCP/UDP port numbers in public versus on the private LAN segment.
-

NAPT does not break ILNP

- **IP:** with NAPT, sessions with H1, H2, H3, or G1 all will use the public IP address that belongs to G1:
  - So, ICMP Locator Update messages for sessions to hosts H1, H2, H3 or gateway G1 will be sent to G1's public IP address.
  - So, *all* ICMP Locator Update messages from outside will naturally be sent to G1 by normal ILNP operation:
- **ILNP:** when G1 sees a valid Locator Update message, G1 updates its NAPT lookup table with the new Locator(s):
  - G1 does not need to tell any interior host about the change.
- ILNP can work with NAPT deployments

IAB Naming and Addressing Workshop 18-19 October 2006 [1]

RFC-4984 (Sep 2007), p4

The clear, highest-priority takeaway from the workshop is the need to devise a scalable routing and addressing system, one that is scalable in the face of multihoming, ...

IAB Naming and Addressing Workshop 18-19 October 2006 [2]

RFC-4984 (Sep 2007), p6

.... workshop participants concluded that the so-called "locator/identifier overload" of the IP address semantics is one of the causes of the routing scalability problem as we see today. Thus, a "split" seems necessary to scale the routing system, although how to actually architect and implement such a split was not explored in detail.