ILNP: A brief overview

Saleem Bhatti

School of Computer Science, University of St Andrews

http://ilnp.cs.st-andrews.ac.uk/

Outline

1. **What?**
   - **Basic information about ILNP.**
2. Why?
   - The rationale for ILNP.
3. How?
   - Basic operation of ILNP.
4. When?
   - ILNP development.

What is ILNP?

- Identifier Locator Network Protocol:
  - http://ilnp.cs.st-andrews.ac.uk/
- ILNP enhances Internet Protocol functionality through the use of crisp **naming**.
- March 2010: IRTF RRG Chairs recommend ILNP for development within the IETF (RFC6115 Feb 2011)
- People:
  - Ran Atkinson (Cheltenham Research, US)
  - Saleem Bhatti (University of St Andrews, UK)

Identifier / Locator Network Protocol

- This is a work in progress:
  - architectural approach, practical engineering – **ILNPv6**
- Focus on network and transport layers (for now)
- **Realise ILNPv6 as a parallel/concurrent system on the existing Internet infrastructure**:
  - **We take a bottom-up engineering approach.**
  - Initial idea based on Mike O'Dell's 8+8/GSE (1996/7)
    - Many enhancements compared to 8+8/GSE
    - Initial "IPv6 8+8" idea dates from emails posted by Bob Smart (02 Jun 1994) and Dave Clark (11 Jan 1995): http://www.ietf.org/mail-archive/web/rrg/current/msg02455.html

Outline

1. What?
   - Basic information about ILNP.
2. **Why?**
   - **The rationale for ILNP.**
3. How?
   - Basic operation of ILNP.
4. When?
   - ILNP development.

University of St Andrews

(New) Requirements

- We wish to try and support a ***harmonised solution to many* network functions**:
  - Multi-homing (host and site).
  - Mobility (host and network).
  - Multi-path capable transport protocols.
  - Localised addressing (NAT).
  - Traffic engineering capability.
  - Packet-level, end-to-end security.
- Currently, solutions for these functions remain disparate and do not function well together.

University of St Andrews

Separate Semantics

- Semantic overload of IP address:
  - **locator** semantics + **identifier** semantics
  - ease implementation of multi-homing, mobility, etc …
- This is a known problem:
  - RFC4984, IAB, 2007
  - RFC2101, IAB, 1997
  - IEN1, 1977; IEN19, 1978, IEN23, 1978
- Many solutions now proposed:
  - HIP, LISP, SHIM, SixOne
  - **ILNP**

Layers are entangled

Protocol Layer	IP
Application	FQDN or IP address
Transport	IP address (+ port number)
Network	IP address
(Interface)	IP address

Entanglement ☹

A problem for harmonising the new requirements …

Example: Multi-homing in IP

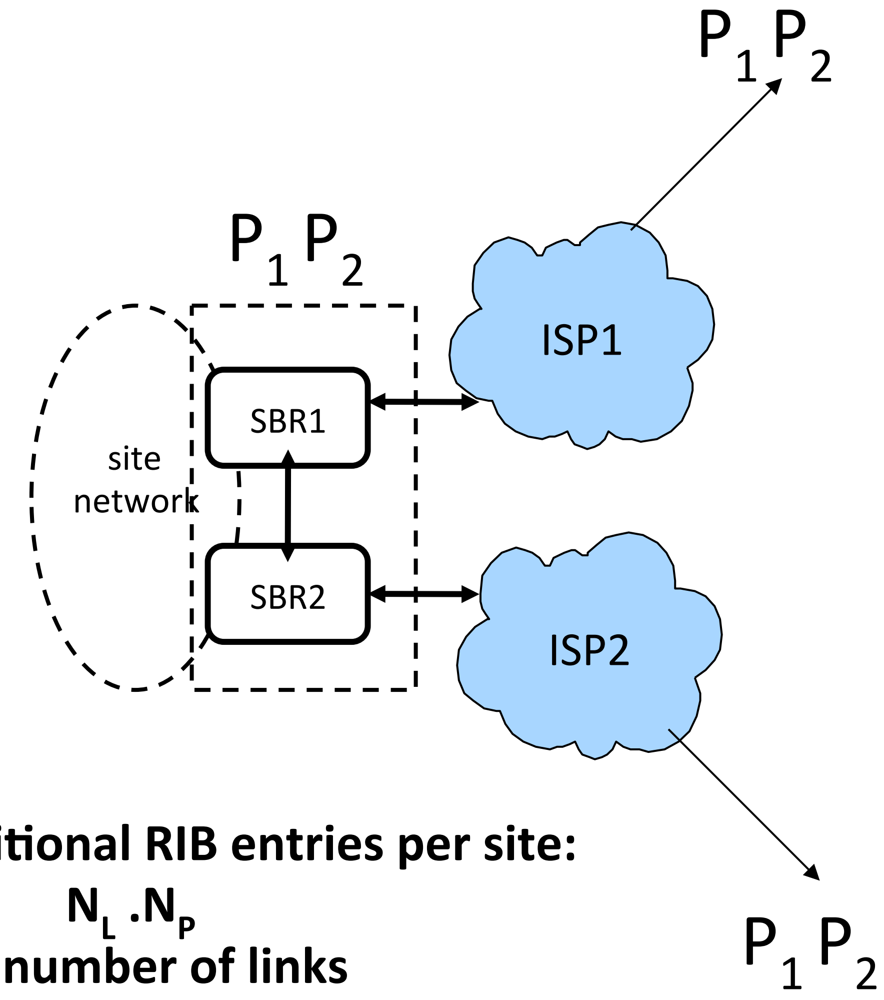

Additional RIB entries per site:

$N_L . N_P$

N_L = number of links

N_P = number of prefixes

- For IP today, Provider Independent (PI) prefixes are popular:
  - Prefix ≡ identity.
  - No renumbering.
- Multi-homing prefixes can lead to bloat in the RIB of the DFZ:
  - **Non-aggregateable prefixes.**

Example: Mobile IP [1]

- Support mobile users without affecting others
- Transparency:
  - to upper layers
  - to remote end-systems
- IPv4 and IPv6:
  - IP address indicates point of attachment to network
  - tied to an **interface**
- Movement of host means:
  - new IPv4 address?
  - update routing information?

University of St Andrews

Example: Mobile IP [2]

- Mobile host (MH):
  - home address (HoA), home network (HN), home agent (HA)
  - care-of-address (CoA), foreign network (FN), foreign agent (FA)
- HoA – identity
- CoA - locator:
- Communication via indirection:
  - HA sends packets to CoA: IP-in-IP enscapsulation
  - HA must reply to ARP for MH

Example: Mobile IP [3]

University of St Andrews

1) MH arrives at FN, and locates FA (using agent advertisements from FA or by solicitation).

2) MH completes registration procedure with FA.

3) MH updates HA with its new CoA (i.e. the FA).

4) Host A now tries to contact MH. Packets for MH are intercepted by HA.

5) HA tunnels the packets from Host A to the CoA for MH (i.e. the FA)

6) The FA de-encapsulates the inner IP packet and transmits the packet locally to MH.

7) The packets from MH to Host A are sent directly from the FN.

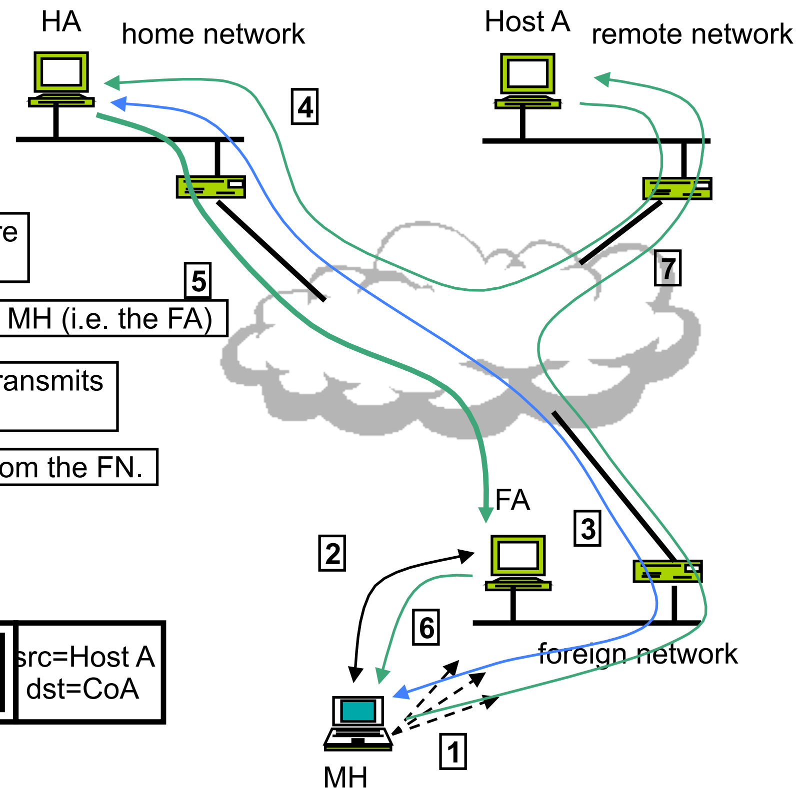

IP-in-IP encapsulation

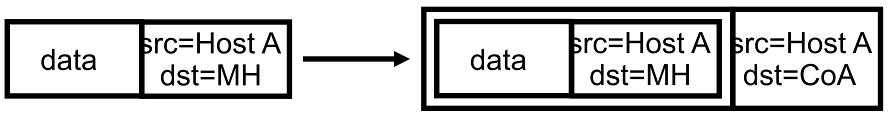

University of St Andrews

What happens if ...

- You want mobility and multi-homing to work together?
- Want also to use IPsec?
- And localised addressing?
- And multi-path transport?
- All without increasing the engineering complexity of the network?
- Harmonised functionality is difficult today

Outline

1. What?
   - Basic information about ILNP.
2. Why?
   - The rationale for ILNP.
3. **How?**
   - **Basic operation of ILNP.**
4. When?
   - ILNP development.

Naming: IP vs. ILNP

Protocol Layer	IP	ILNP
Application	FQDN or IP address	FQDN (RFC1958)
Transport	IP address (+ port number)	Identifier (+ port number)
Network	IP address	Locator
(Interface)	IP address	(dynamic binding)

Entanglement ☹ **Separation ☺**

FQDN = fully qualified domain name

Engineering issues for ILNPv6

We wish to have an **incrementally deployable** solution that is also **backwards compatible**:

1. Core network devices and protocols should not need to change, e.g. routers, switches of today can be used without modification.
2. Reuse the existing core protocol deployment as much as possible, e.g. make use of existing IPv6.
3. Try to limit the impact on current applications (but we have to accept some applications might break).
4. The end system stack will need to change, but changes should run in parallel with current stack.

University of St Andrews

ILNPv6

- Can be seen as a set of 'extensions' to IPv6:
  - Uses same packet format as IPv6 in network core.
  - IPv6 core routers do not need to change.
  - Incrementally deployable on IPv6 core.
  - Backwards compatible with IPv6.
- Split 128-bit IPv6 address:
  - **64-bit Locator (L) - network name.**
  - **64-bit Identifier (I) - node name.**
- Could also be retro-fitted to IPv4 (but messy).

IPv6 addresses and ILNPv6

```
IPv6 (as in RFC3587 + RFC3177: RFC6177):

| 3 |     45 bits         |  16 bits  |       64 bits            |
+---+---------------------+-----------+--------------------------+
|001|global routing prefix| subnet ID |   Interface Identifier   |
+---+---------------------+-----------+--------------------------+
```

unicast IPv6 routing (address) prefix

same syntax, different semantics

```
ILNPv6:

|              64 bits                |       64 bits            |
+---+---------------------+-----------+--------------------------+
|              Locator (L64)          | Node Identifier (NID)    |
+---+---------------------+-----------+--------------------------+
```

same syntax and semantics as
IPv6 routing (address) prefix
so IPv6 core routers work as today

these bits only examined and
acted upon by end systems

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                      +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

DNS enhancements required

- New DNS records required:
  - NID (Identifier)
  - L64 (Locator for ILNPv6)
  - L32 (Locator for ILNPv4)
  - LP (Locator Pointer – a little like a CNAME, but not quite)
- Fully qualified domain name (FQDN) lookup:
  - return any NID, L64, L32 and LP records

Mobile **hosts** in ILNPv6

- **Mobility/multi-homing duality.**
- An individual mobile host (MH) picks up a new Locator value as it moves into a new network.
- MH sends Locator Update (LU) messages to correspondents for existing sessions.
- MH updates DNS with new Locator value.
- If cells overlap, MH can use multiple Locator values simultaneously for **soft hand-off**.

Multi-homing in ILNPv6 [1]

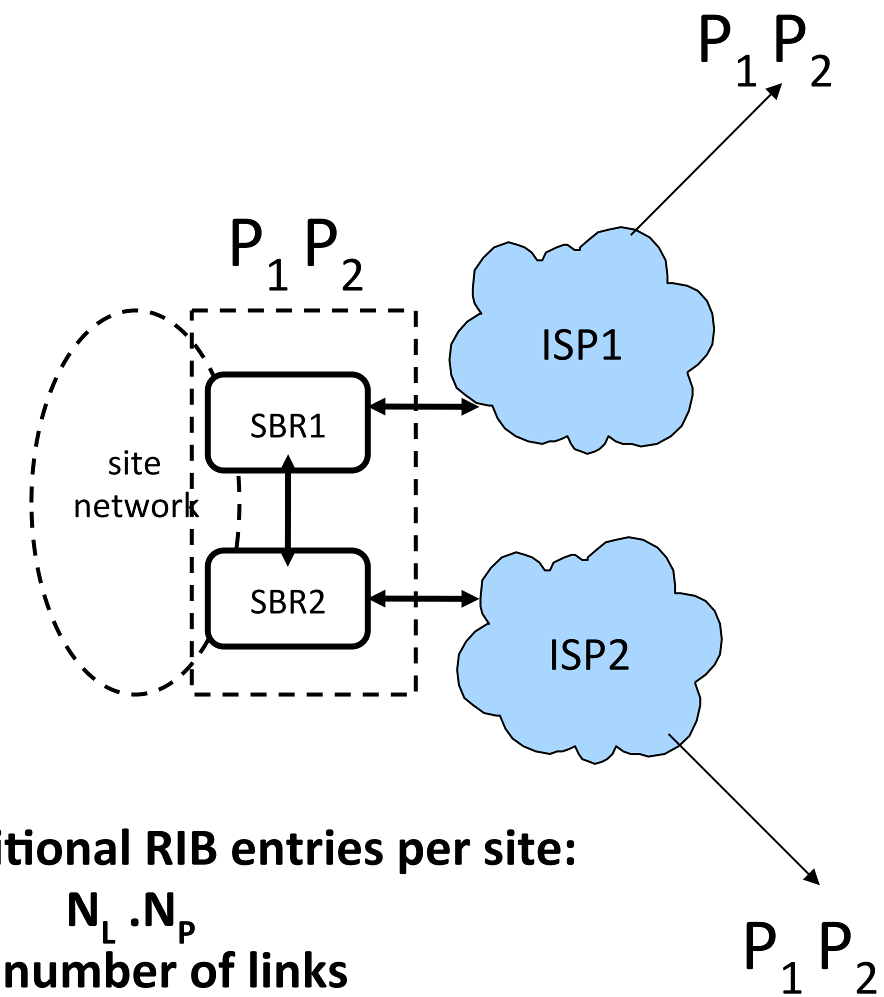

Additional RIB entries per site:

$N_L . N_P$

N_L = number of links

N_P = number of prefixes

- For IP today, Provider Independent (PI) prefixes are popular:
  - Prefix ≡ identity.
  - No renumbering.
- Multi-homing prefixes can lead to bloat in the RIB of the DFZ:
  - **Non-aggregateable prefixes.**

Multi-homing in ILNPv6 [2]

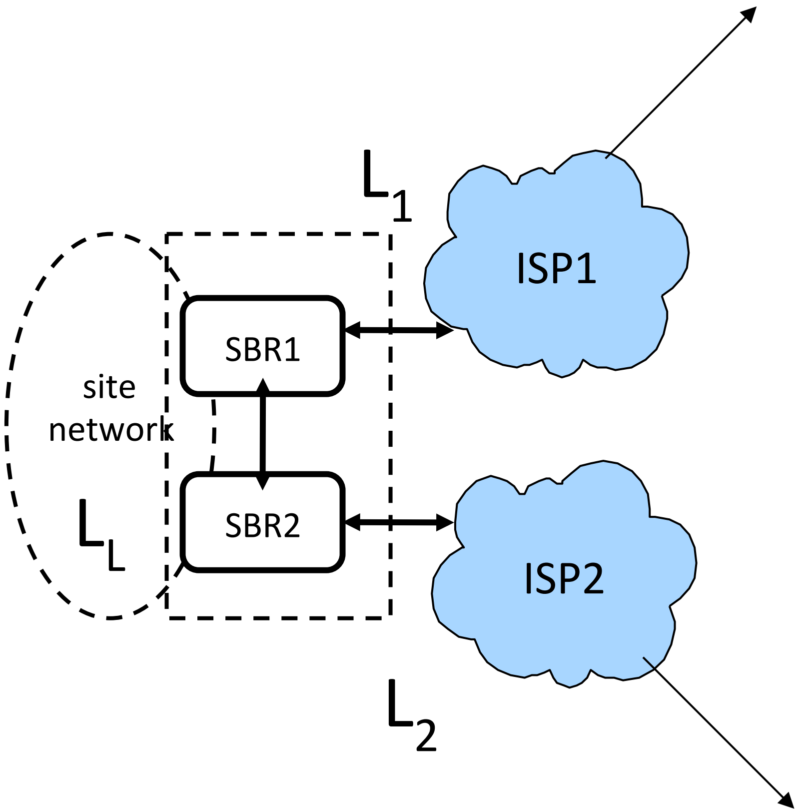

No additional RIB entries

- ILNP, Locator taken from the allocated prefixes of ISP:
  - **Identity not related to Locator.**
  - Renumbering through operation of IPv6.
- **No extra prefixes required**:
  - **All Locator values visible via DNS.**

Outline

1. What?
   - Basic information about ILNP.
2. Why?
   - The rationale for ILNP.
3. How?
   - Basic operation of ILNP.
4. **When?**
   - **ILNP development.**

Thank You!

- Questions
- More information on ILNP:
  - http://ilnp.cs.st-andrews.ac.uk/
- Contact information:
  - Saleem Bhatti <saleem@cs.st-andrews.ac.uk>

Supplementary slides

Development options

- Simulation:
  - Good control, high-scalability, reproducibility of experiments etc.
- Emulation:
  - e.g. use of an overlay network is feasible (Masters student project, 2009), with constraints.
  - OneLab, PlanetLab (control + mgmt + monitoring?)
- **Test-bed – implementation in OS stack:**
  - **prototype for FreebSD in progress (available ~Q4 2012)**
  - **prototype Linux TBA (available ~Q4 2013?)**

Names

- My definition of a "name":
  A set of bits used to label an object. The semantics of the name are defined within the context of use of the object it names.
- Examples:
  - protocol name – 'http'
  - port number – '80'
  - fully qualified domain name (FQDN), e.g. 'marston.cs.st-andrews.ac.uk'
  - **IP address - '138.251.195.61'**

IEN 1 (29 July 1977)

- Section 3 ADDRESSING (pp 6-12):
  - Discusses physical vs. logical addressing
- Section 3.2 Special Topologies (pp 7-8):
  - Specifically discusses "Changes in Topology" (mobility) and "Multiply-Connected Hosts" (multi-homing)
  - Flags problems with use of IP addresses (as today).
- Lots of wisdom:
  - IENs 19, 23, 31, 46

Locators and Identifiers [1]

- **Locator, L:**
  - **Topologically significant.**
  - Names a (sub)network (as today's **network prefix**).
  - Used only for routing and forwarding in the core.
- **Identifier, I:**
  - **Is not topologically significant.**
  - Names a logical/virtual/physical node, does **not** name an interface (value ala RFC4291 Sec 2.5.1).
- **Upper layer protocols bind only to Identifier.**

Locators and Identifiers [2]

- Locator, L:
  - **Can change** value during the lifetime of a transport session (mobility, site-controlled traffic engineering).
  - Multiple Locators can be used simultaneously (multi-homing, multi-path transport protocols).
- Identifier, I:
  - **Remains constant** during the lifetime of a transport session (localised addressing, IPsec.).
  - Multiple Identifiers can be used simultaneously by a node, but not for the same session.

Mapping FQDNs to I/L values

- DNS is used as today:
  - FQDN is used to map to I/L values instead of AAAA
- Need new DNS Resource Records, e.g.:
  - NID – 64-bit Identifier value, EUI-64 syntax
  - L64 – 64-bit Locator value
  - LP – Locator Pointer (like CNAME for L64)
- DNS lookup will return:
  - 1 or more I64 records, 1 or more L64 records
  - For multiple I64 and L64 RRs, use preference bits

Examples of ILNP usage

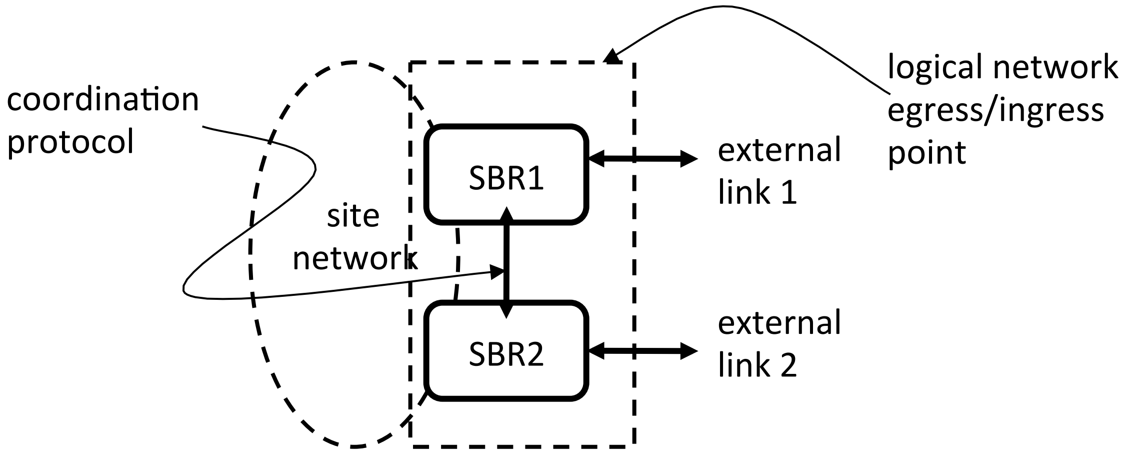

SBR = site border router

NAT in IPv4 and IPv6

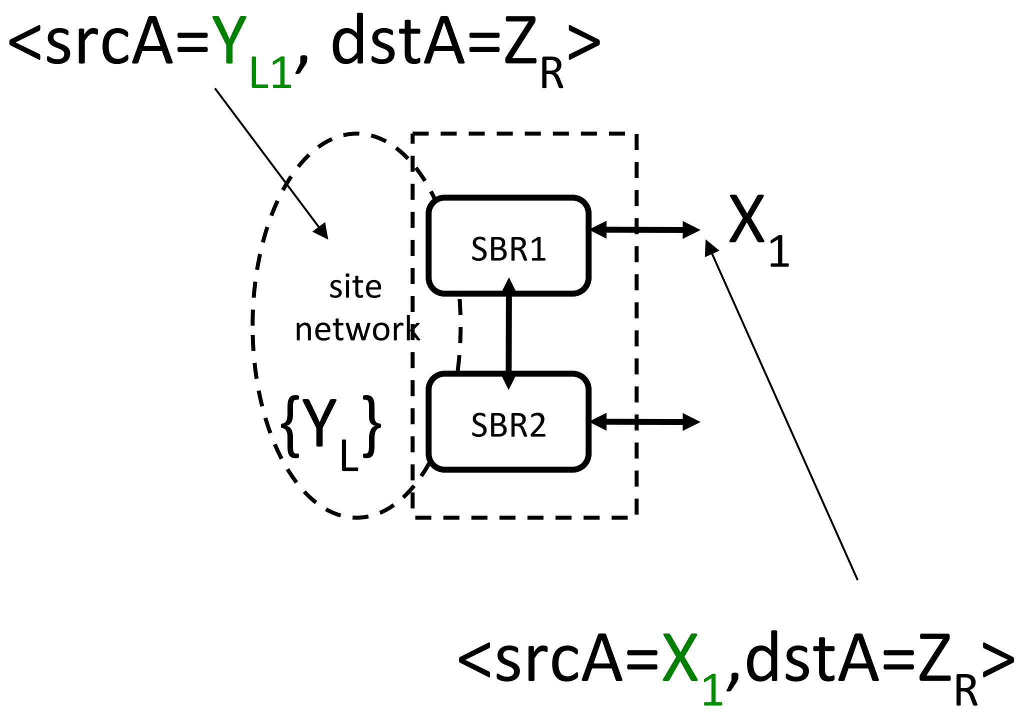

- **NAT allows address reuse for a site**:
  - single address shared amongst many hosts (use of port numbers)
- End-to-end view is lost, as identity namespace has a discontinuity at the SBR

NAT in ILNPv6

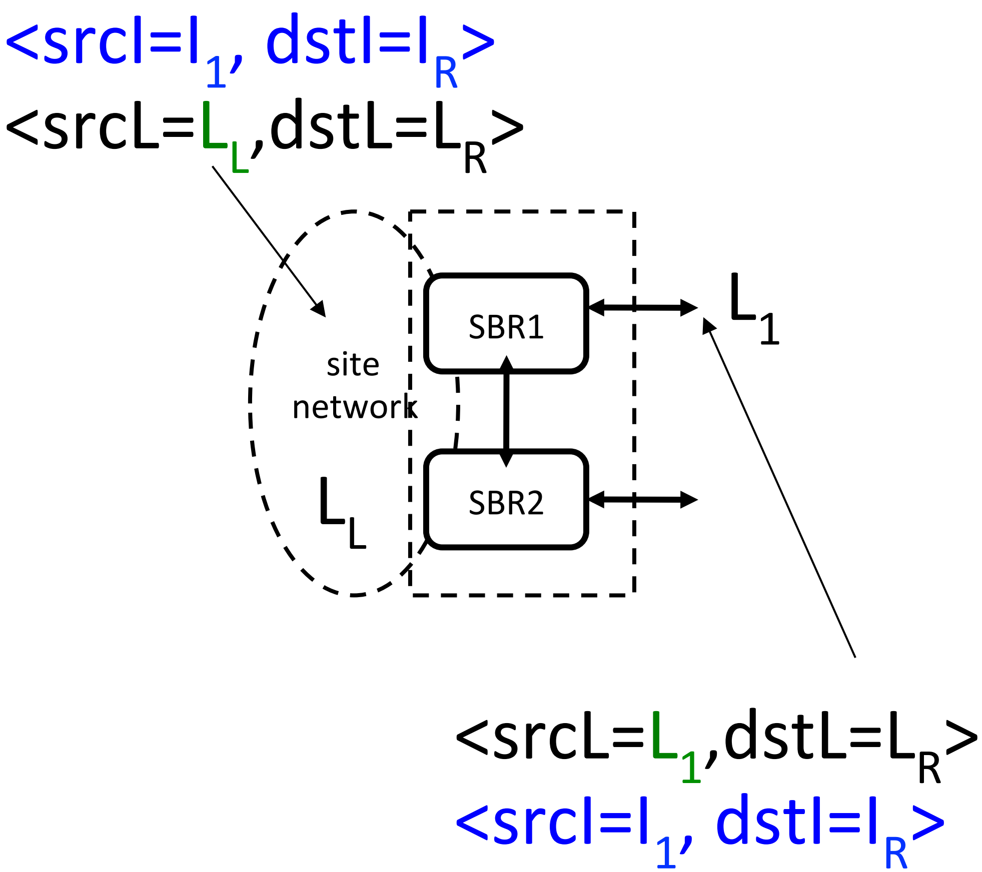

- **NAT is now a feature not a hack**:
  - L is not part of the end system transport session state.
  - L_L value ala RFC4193
  - **end-to-end view**
- SBRs perform **Locator rewriting** without affecting end-to-end state.

Mobile **networks** in ILNP [1]

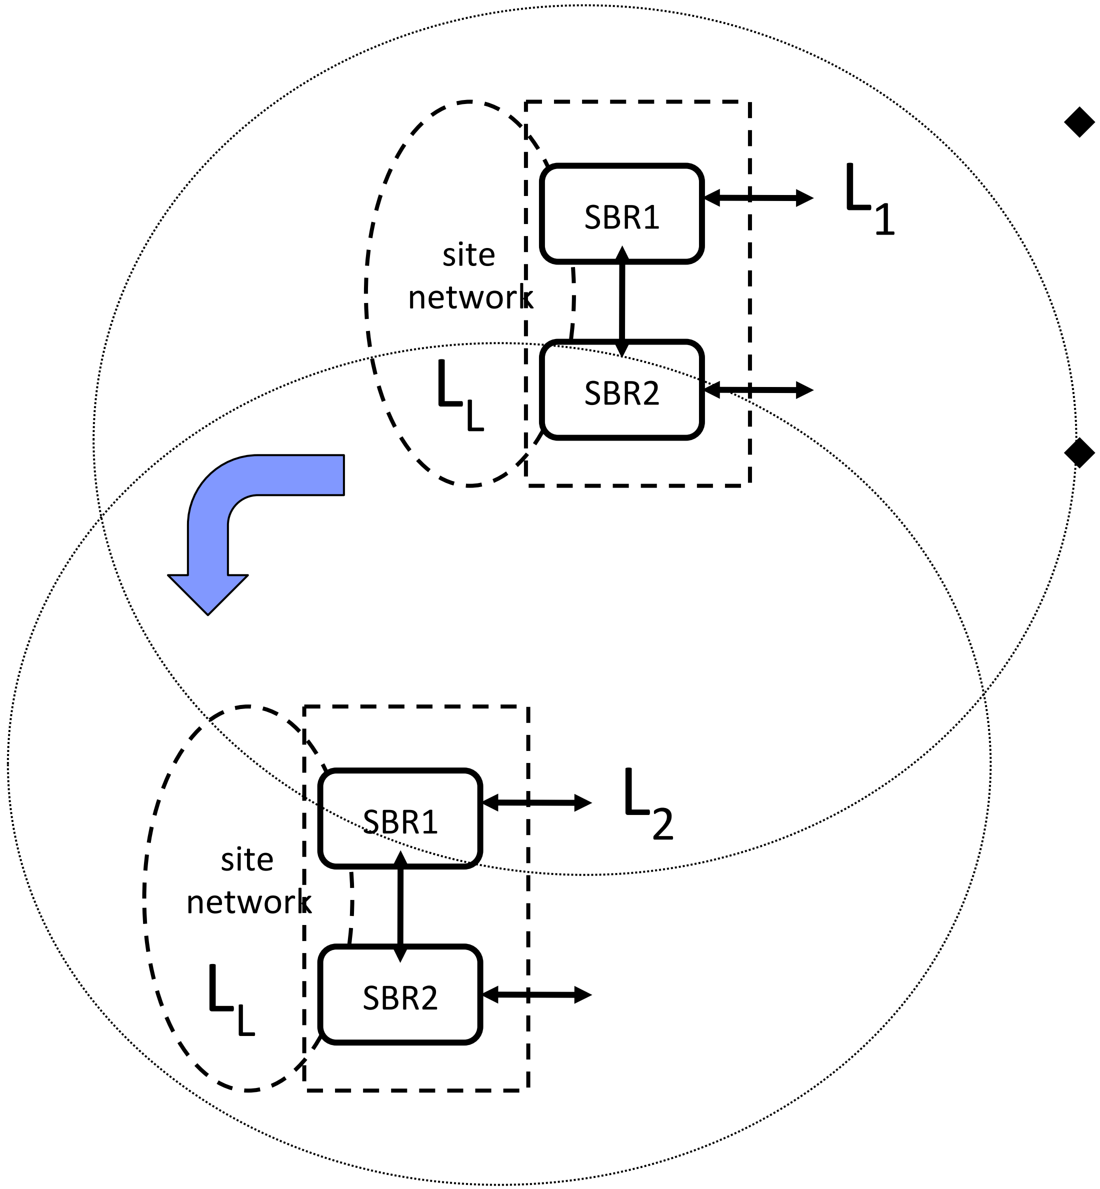

- Use NAT to 'hide' the movement to internal nodes.
- SBR changes Locator value as the mobile network moves:
  - Sends Locator Update (LU) messages to correspondents.
  - Updates DNS.

Mobile **networks** in ILNPv6 [2]

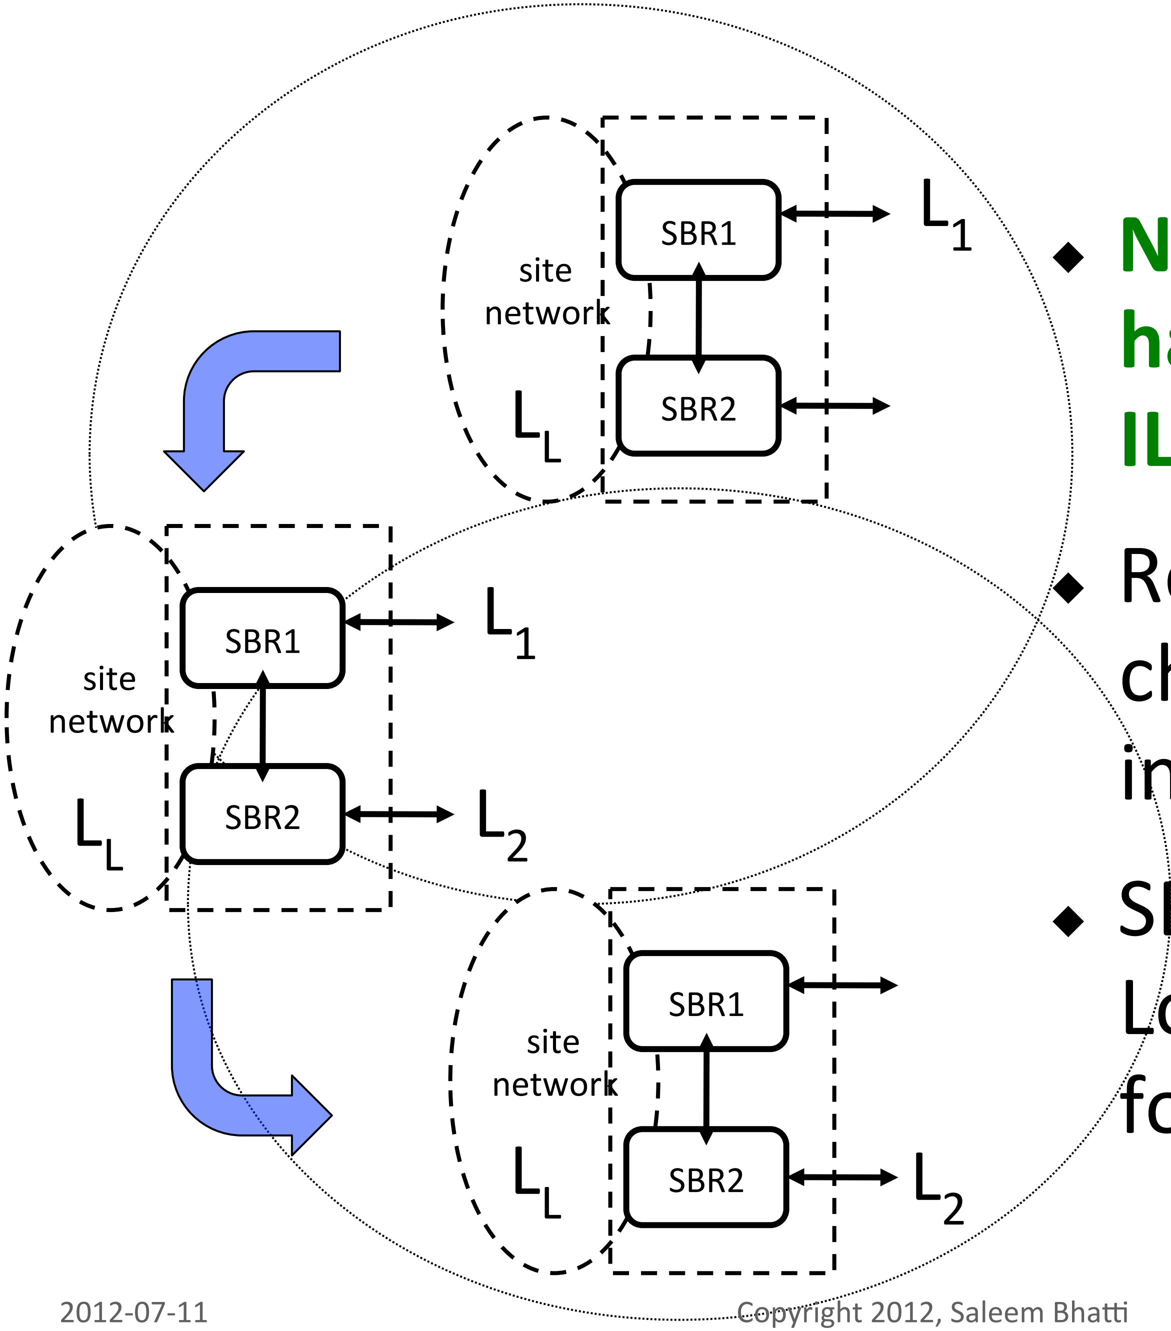

- **Network layer soft-hand-off possible in ILNP.**
- Requires at least 2 radio channels (or 2 radio interfaces).
- SBRs can handle Locator rewriting and forwarding as required.

University of St Andrews

Mobile **hosts** in ILNPv6

- **Mobility/multi-homing duality.**
- An individual mobile host (MH) picks up a new Locator value as it moves into a new network.
- MH sends Locator Update (LU) messages to correspondents for existing sessions.
- MH updates DNS with new Locator value.
- If cells overlap, MH can use multiple Locator values simultaneously for **soft hand-off**.

Multi-homing in ILNPv6 [1]

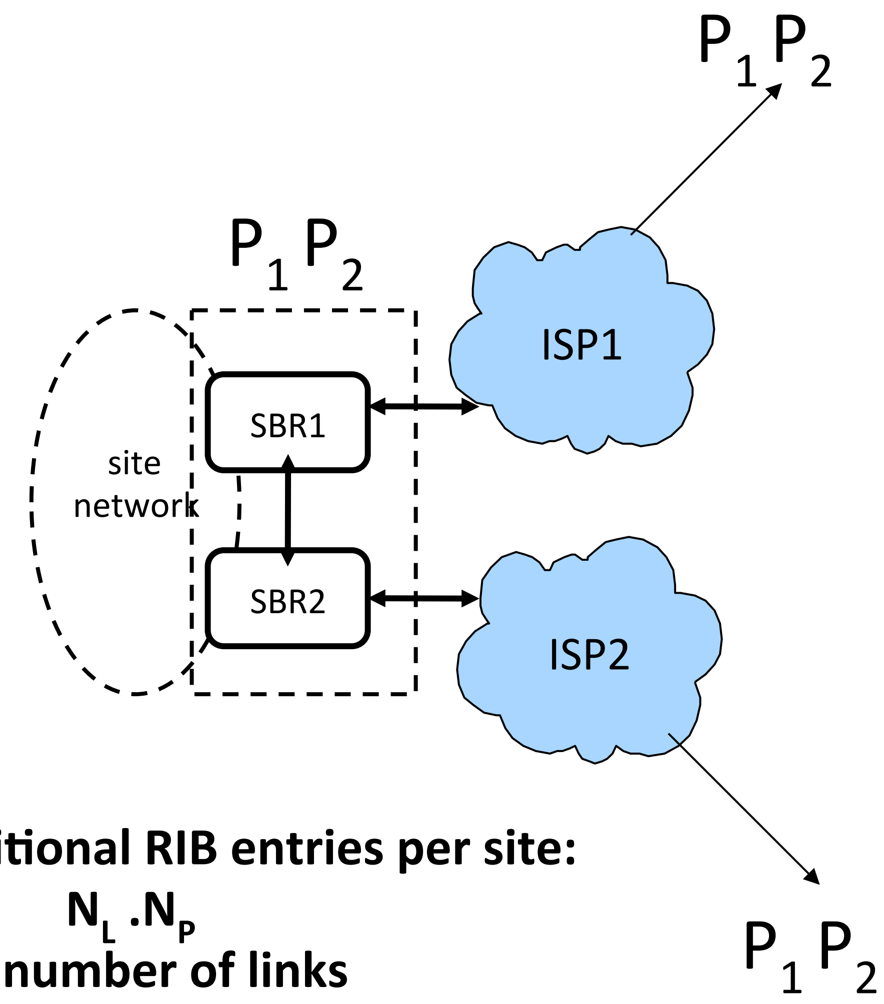

Additional RIB entries per site:

$N_L . N_P$

N_L = number of links

N_P = number of prefixes

- For IP today, Provider Independent (PI) prefixes are popular:
  - Prefix ≡ identity.
  - No renumbering.
- Multi-homing prefixes can lead to bloat in the RIB of the DFZ:
  - **Non-aggregateable prefixes.**

Multi-homing in ILNPv6 [2]

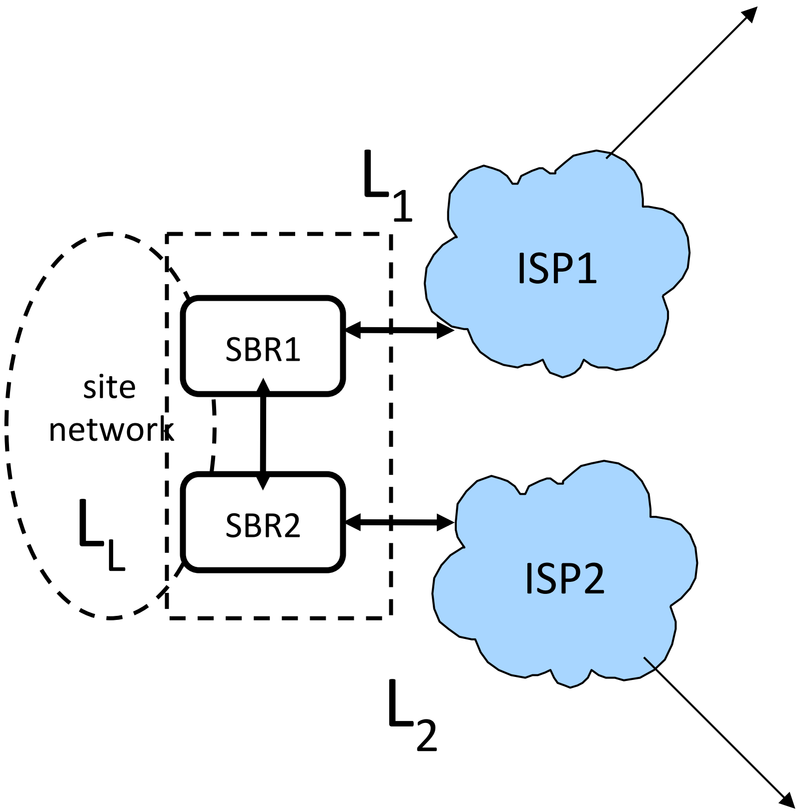

No additional RIB entries

- ILNP, Locator taken from the allocated prefixes of ISP:
  - **Identity not related to Locator.**
  - Renumbering through operation of IPv6.
- **No extra prefixes required**:
  - **All Locator values visible via DNS.**

VM migration

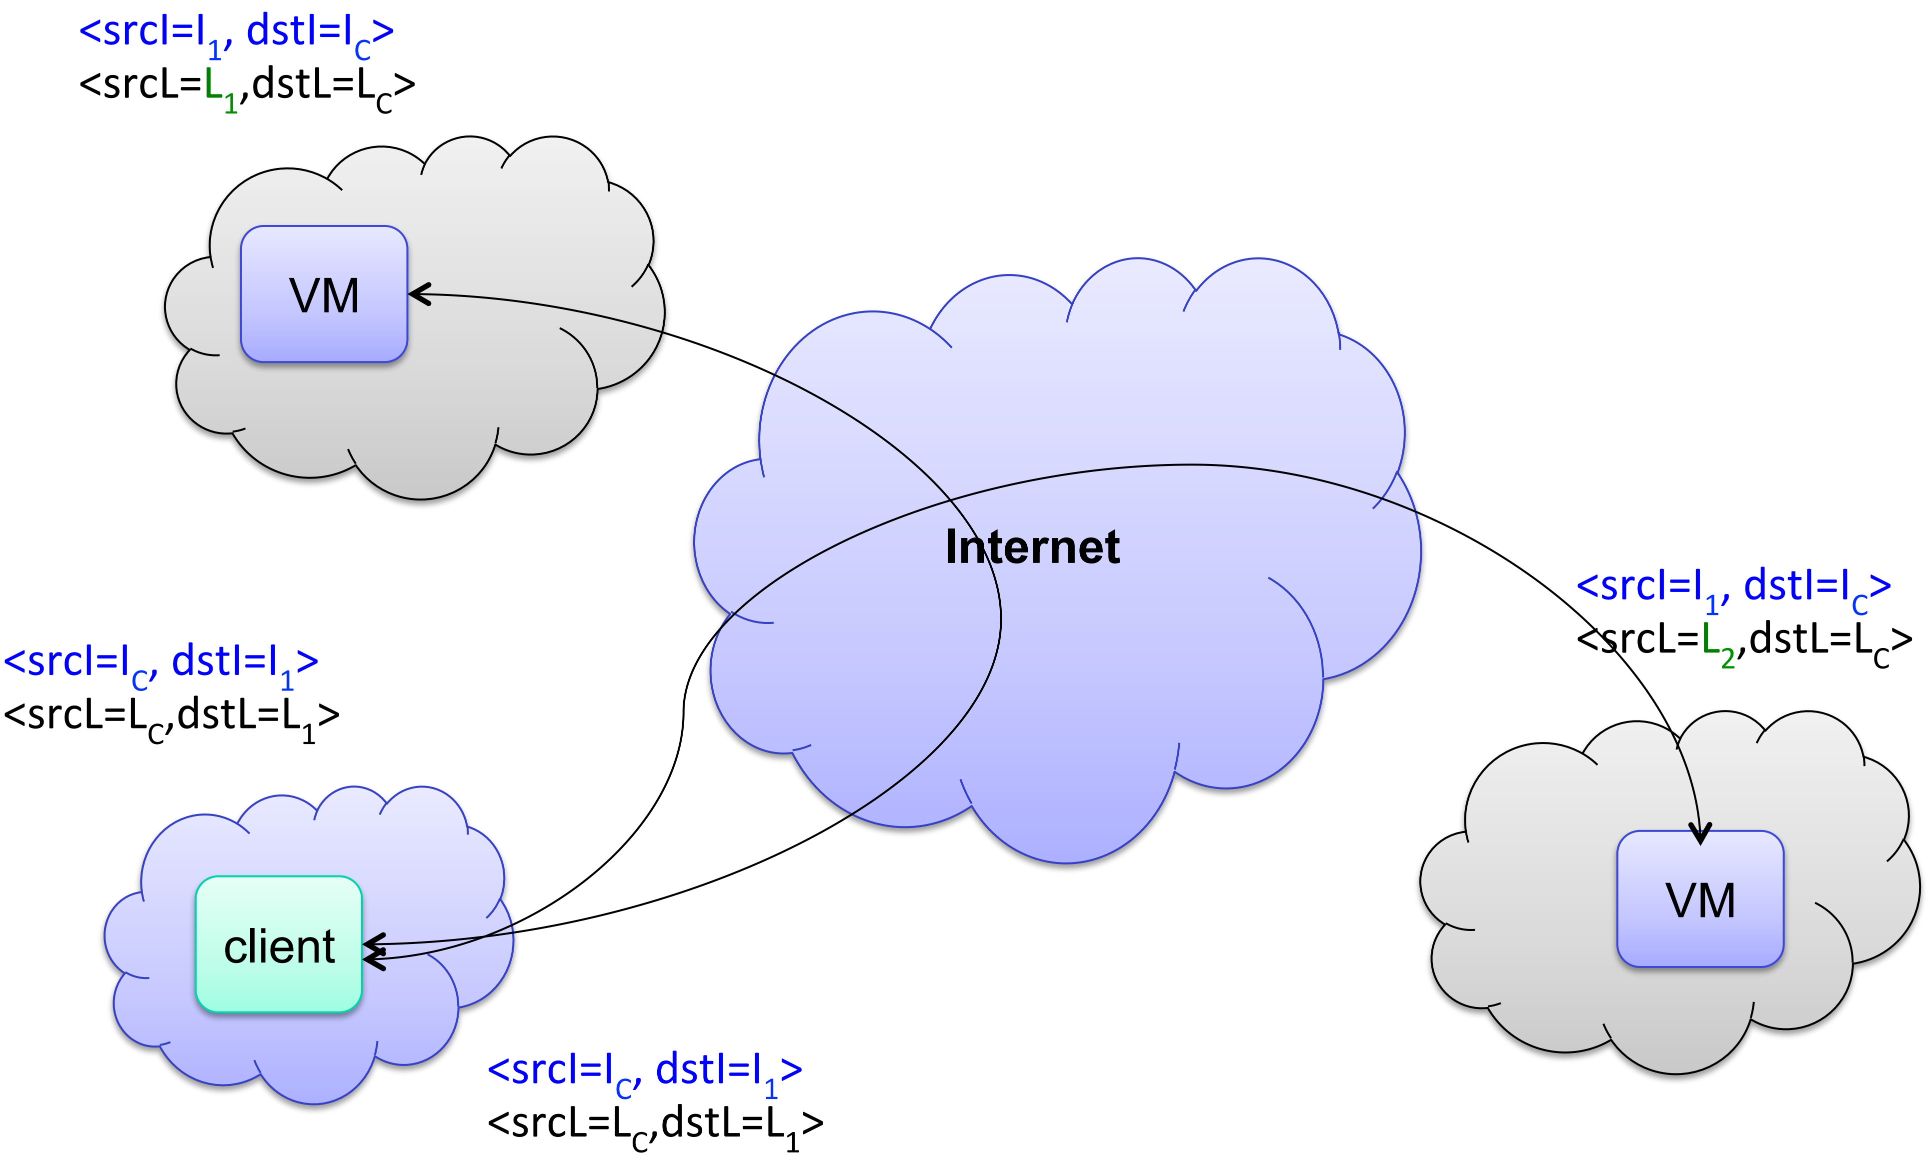

Past relevant work

University of St Andrews

- Our work is based on the following key ideas:
  - IEN1 (1977): separate names for layer 3 & layer 4
  - Bob Smart (c. 1994): email to public SIPP list proposing an 8+8 split for the IPv6 address.
  - Dave Clark (c.1995): email to public IRTF list proposing 8+8 plus ideas on enabling functionality.
  - Mike O'Dell (c.1997): IETF drafts on GSE and 8+8.
  - IRTF NameSpace RG (NSRG)
- We have enhanced and extended those early ideas in order to address a comprehensive set of functionality through naming.

University of St Andrews

Current relevant work

- Host Identity Protocol (HIP) – host-based:
  - IRTF and IETF, RFC4423
  - Research grade implementation available.
  - Uses public-key (non public-key option?)
- SHIM6 – host-based (IETF drafts):
  - Research grade implementation available.
- LISP – network based (IETF drafts):
  - Use of tunnels and additional state/signalling.
- MEXT – host and network mobility (IETF drafts):
  - Aims to combine MIPv6, NEMO and IKEv2.

Other related work on architecture

- NIMROD
- IP Next Layer (IPNL)
- TurfNet
- Internet Indirection Infrastructure (I^3)
- Others ... (see the list of references in the papers on ILNP WWW site)

University of St Andrews

No free lunch

- DNS support – not new, but explicit in ILNPv6:
  - New RRs + **zero TTL for some DNS records**.
  - Secure DNS Dynamic Update for Locator changes.
- Renumbering + address management at sites.
- No globally routeable interface name, which may impact some applications such as SNMP.
- Some legacy applications may break, e.g. FTP.
- Interworking scenarios (IPv6, IPv4).

Application layer protocols

- URLs:
  https://marston.cs.st-andrews.ac.uk/
- Can also use an IP address:
  https://138.251.195.61/
- Notice, the use of **either** a DNS name or an IP address – FQDN and IP address used as synonyms.
- **IP address is overloaded**:
  - used in application protocols as a **session identifier**

Transport protocols

- TCP uses a tuple to identify a TCP connection:
  - local IP address
  - local port number
  - remote IP address
  - remote port number
- TCP state (and the pseudo-header checksum) is bound to **all** the bits in the local and remote IP address.
- **IP address used as an Identifier.**

Network layer

University of St Andrews

- IP address bits are used in **routing**:
  - IP address (network) prefix, e.g.
    138.251.195.61/24
    means that 138.251.61 (the network prefix) is used for routing at the IP layer
- The host part of the address may be further used for sub-netting at the site:
  - IP sub-netting on host bits, e.g.
    138.251.195.61/25
    means 1 bit of the host part of the address is used
- **IP Address used as a Locator.**

Interface identifier

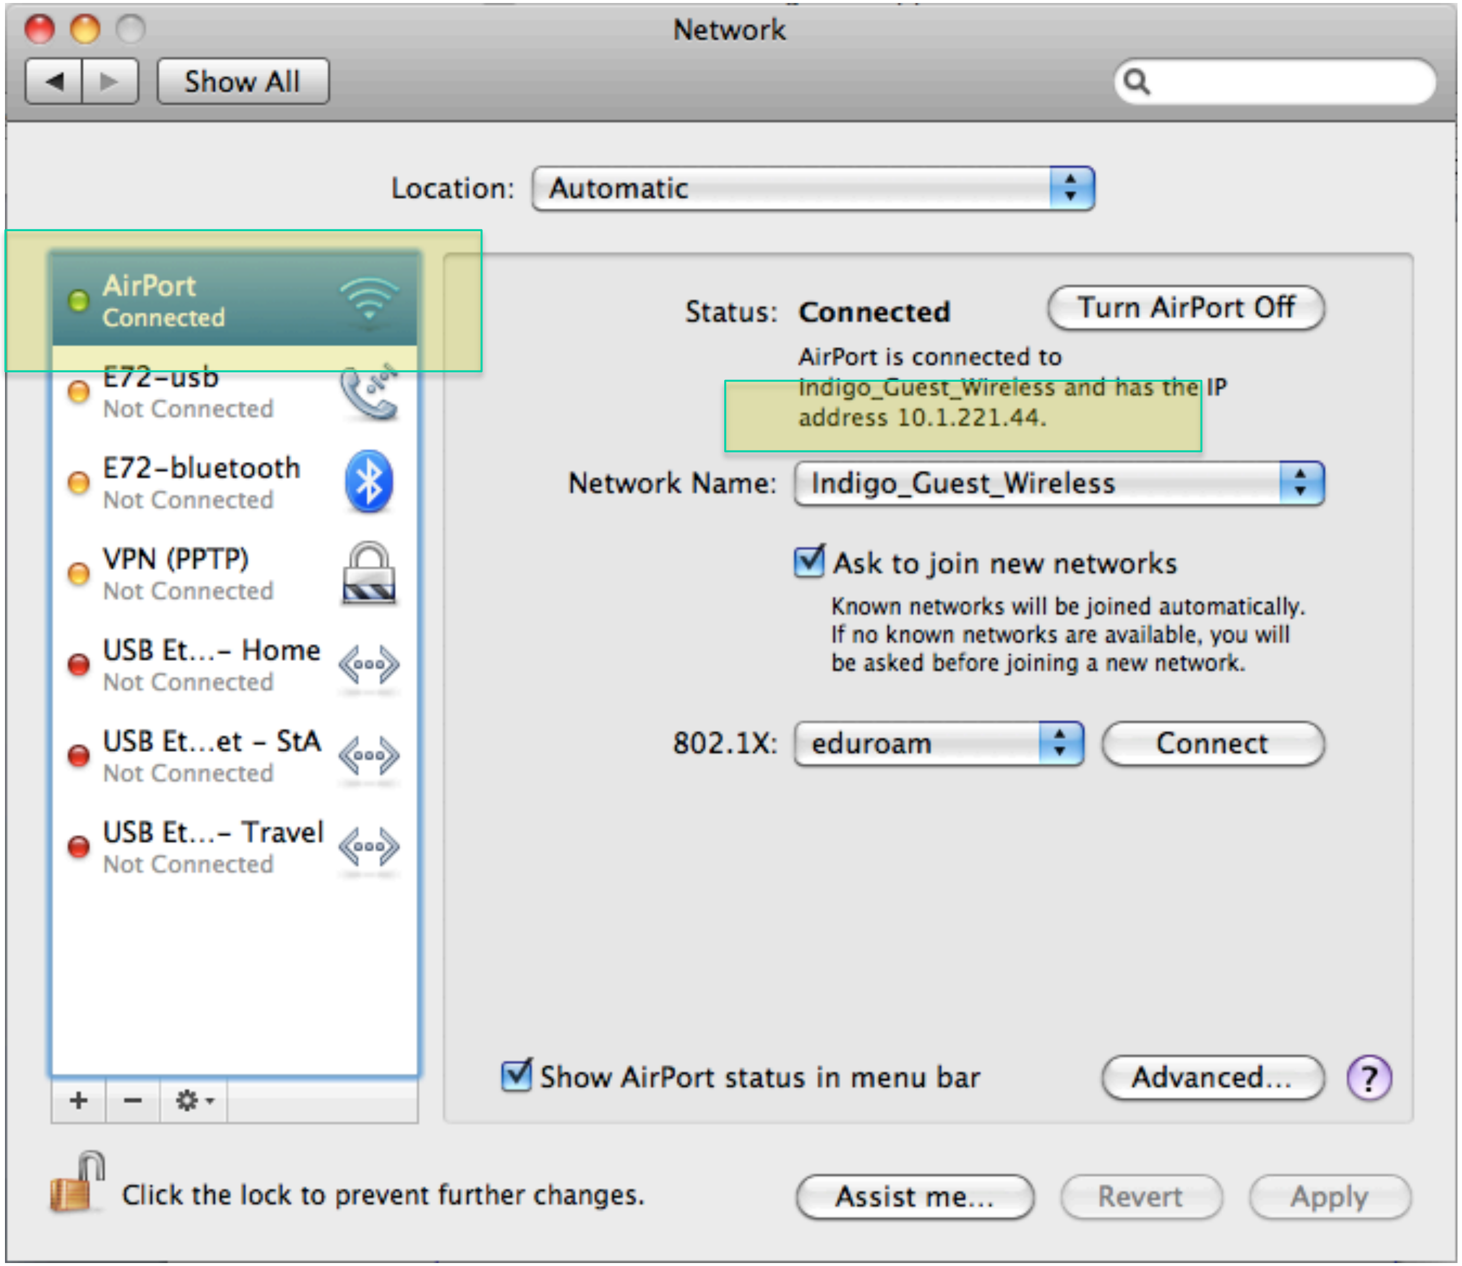

RFC4984 (Sep 2007)

IAB Naming and Addressing Workshop 18-19 October 2006
RFC4984, 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.

RFC2101 (Feb 1997)

IPv4 Address Behaviour Today
RFC2101 pp 3-4

Identifiers should be assigned at birth, never change, and never be re-used. Locators should describe the host's position in the network's topology, and should change whenever the topology changes. Unfortunately neither of the these ideals are met by IPv4 addresses.

RFC4984 (Sep 2007) [1]

IAB Naming and Addressing Workshop 18-19 October 2006
RFC4984 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, and that facilitates a wide spectrum of traffic engineering (TE) requirements.

Problems today

- The growth of the DFZ RIB:
  - currently at ~300,000+ entries and growing
- Multi-homing:
  - increasingly popular for customers
- Traffic Engineering:
  - a useful tool for providers and customers
- Provider Independent Addressing:
  - address is a form of identity; re-numbering is hard

User programs – Java API

- TCP Client:
  ```
  Socket skt = new Socket("srv.blob.com", 1234);
  ```
- Can also use an IP address:
  ```
  Socket skt = new Socket("10.12.14.16", 1234);
  ```
- Notice, the use of **either** a DNS name or an IP address – FQDN and IP address used as synonyms.
- **IP address is overloaded**:
  - may be used in application code in place of FQDN

RFC1958 (June 1996)

Architectural Principles of the Internet
RFC1958 p5, Section 4.1

In general, user applications should use names rather than addresses.

University of St Andrews

Locators and Identifiers [3]

- Locator, L:
  - Network prefix, from normal configuration or using discovery protocol (e.g. IPv6 Router Advertisement).
- Identifier, I:
  - Default value: a node uses bits from a local interface to form an EUI-64 value, which is used as an Identifier for that node.
  - Other interesting possibilities ... (work in progress) .
  - Strictly, needs to be unique within scope of a given Locator value: global uniqueness is good, however.

University of St Andrews

Comparison with LISP [1]

- **LISP**: customer focused, practically-directed engineering solution, with a goal of minimal cost to end sites, employing network upgrades that would be invisible to the end users, and reduce the burden of routing state on "core network". Objective is to provide a product-based solution.

- **ILNPv6**: research vehicle to explore the current use of addressing and examine fundamental architectural issues of how naming can be used to enable new functionality. Objective is to give a proof-of-concept implementation in order to demonstrate that ILNP could be made to work as described.

Comparison with LISP [2]

	LISP	ILNPv6
What changes?	network	host
Architecture	map-and-encap	naming
Site renumbering	no	optional
End-host changes	no	yes
New network entities required	yes	no
Backbone MTU > access MTU	yes	no
BGP & DFZ state reduction	yes	yes
State 'displacement'	EID-RLOC mapping	DNS lookups
Working code	yes	in progress
'Well-behaved' applications work without modification	yes	yes
IPv6	yes	yes
IPv4	yes	possibly[1]

[1] Technically possible, deployability unclear.

Comparison with LISP [3]

	LISP	ILNPv6
Site multi-homing	yes	yes
Host multi-homing	not currently defined[1]	yes
Multicast	yes	yes
Traffic engineering options	yes	yes
Localised addressing (NAT)	in progress[2]	yes
Harmonised functionality	in progress[2]	yes
Mobile hosts	in progress[2]	yes
Mobile networks	not currently defined[1]	yes
Multi-path transport	no	yes

[1] Technically possible, deployability unclear.
[2] Internet draft document available.

ILNPv4 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|IHL=12 |Type of Service|          Total Length         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Identification        |Flags|      Fragment Offset    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Time to Live |    Protocol   |         Header Checksum       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Source Locator                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Destination Locator                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| OT=ILNPv4_ID  |     OL=5      |         Padding=0x0000        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+-                     Source Identifier                       -+
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+-                  Destination Identifier                     -+
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|OT=ILNPv4_NONCE|     OL=2      |       top 16 bits of nonce    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    lower 32 bits of nonce                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

IHL: Internet Header Length     OT: Option Type   OL: Option Length
```

Traffic Engineering in ILNP

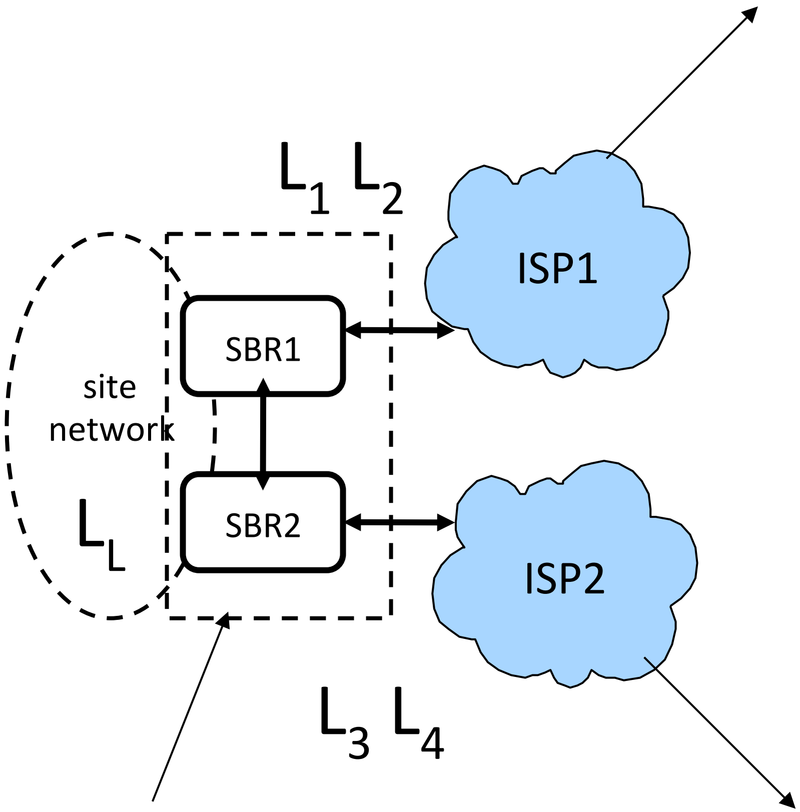

Policy mechanisms to decide on which links packets are forwarded.

- SBR(s) can use today's policy-based approaches for filtering and forwarding with **Locator rewriting**.
- Incoming packets can also be redirected across SBRs.

IPsec

- IPsec currently uses the whole of the IP address for binding a Security Association (SA).
- In ILNP, the SA binds only to the Identifier, I:
  - I remains constant throughout the session.
  - L value can change (for whatever reason) while the session is in progress.
  - As long as I does not change, end-to-end session state is maintained.

No free lunch [1]

- To support mobility and dynamic multi-homing:
  - TTL for DNS records needs to be set as low as possible, ideally to zero.
  - TTL for DNS records for fixed sites can remain as used today.
- To support multi-homing and TE:
  - L64 records could benefit from the use of preference bits to indicate preferred Locator usage.

University of St Andrews

No free lunch [2]

- No globally routeable interface name, which may impact some applications such as SNMP.
- Some legacy applications may break, e.g. FTP.
- DNS reliance in ILNPv6:
  - Not new, but made explicit in ILNPv6.
  - No new security issues created.
  - Can use DNS Security and Dynamic DNS Update, which is already being worked on within the IETF, and already implemented in DNS servers.

Practical issues – initial thoughts

- Portability of applications?
  - What are the range of problems that might exist for porting applications to ILNPv6?
- Optional, enhanced networking API?
  - Use of names, I:L not seen.
  - Exploit ILNP, e.g. signal for change in L.
- **DNS usage impact?**
  - **How might DNS be affected in real use?**
- Adoption in end-system stacks?

ILNPv6 in-house testbed

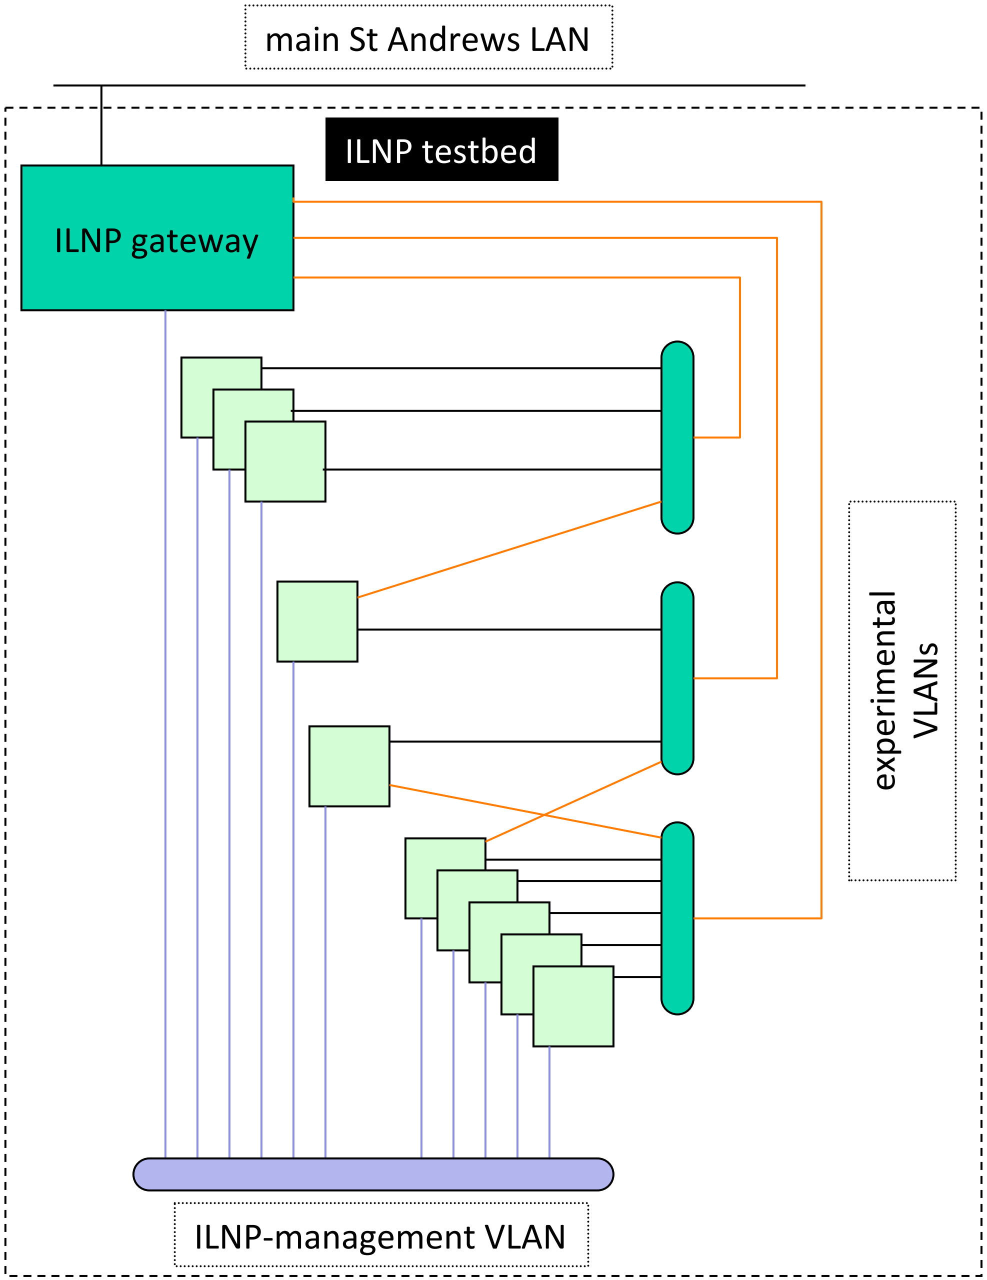

- Use of existing services:
  - e.g. use of deployed DNS and IPv6 routing.
- Exploit VMs when possible.
- Off-the-shelf equipment:
  - easy of use
  - costs
- Open source:
  - leverage existing kernel code
  - make available to community

University of St Andrews

Useful features of a testbed [1]

- Kernel code in practical settings – working on low-level protocols is disruptive:
  - **things will break!**
- Separation of management-, control- and data-(user-) plane functions, logically and physically:
  - **out-of-band** management and control for nodes.
  - separate control of **routing links** and **routing configuration**.
  - data plane connectivity (e.g. via VLANs)

University of St Andrews

Useful features of a testbed [2]

- Control of experimental nodes:
  - **console access** for boot messages and control.
  - **administrator level access**.
  - **power control** for remote power cycling.
- Support services:
  - **Naming** (DNS) configuration and control.
  - **Network monitoring** for troubleshooting and fine-grained operation- and performance-analysis.
- Security: **lots** of issues …

Extending testbed to larger scales

University of St Andrews

- How can we achieve the same level of control and configuration at larger scales?
- Do we need to change the way we undertake low-level protocol development in order to use larger-scale testbeds?
- Are new approaches, such as virtualisation, applicable to such large-scale scenarios with such low-level protocol development?
- Is it possible to conduct such development on distributed, large-scale testbeds?