Cheltenham Research

University of St Andrews

UCL

Evolving the Internet Architecture Through Naming

Ran Atkinson, Cheltenham Research, USA
Saleem Bhatti, University of St Andrews, UK
Steve Hailes, University College London, UK

What’s in a name?

Juliet:

"What's in a name? That which we call a rose
By any other name would smell as sweet.”

Romeo and Juliet (Act 2, Scene 2, 1-2)
William Shakespeare

Juliet was not worried about names, but for the Internet, they can make a difference …

Schedule

1. **Problem space**
2. Introduction to ILNP
3. Using ILNP
4. Issues and related work
5. Wrap-up

Cheltenham Research

University of St Andrews

UCL

PRM Layers for network architecture

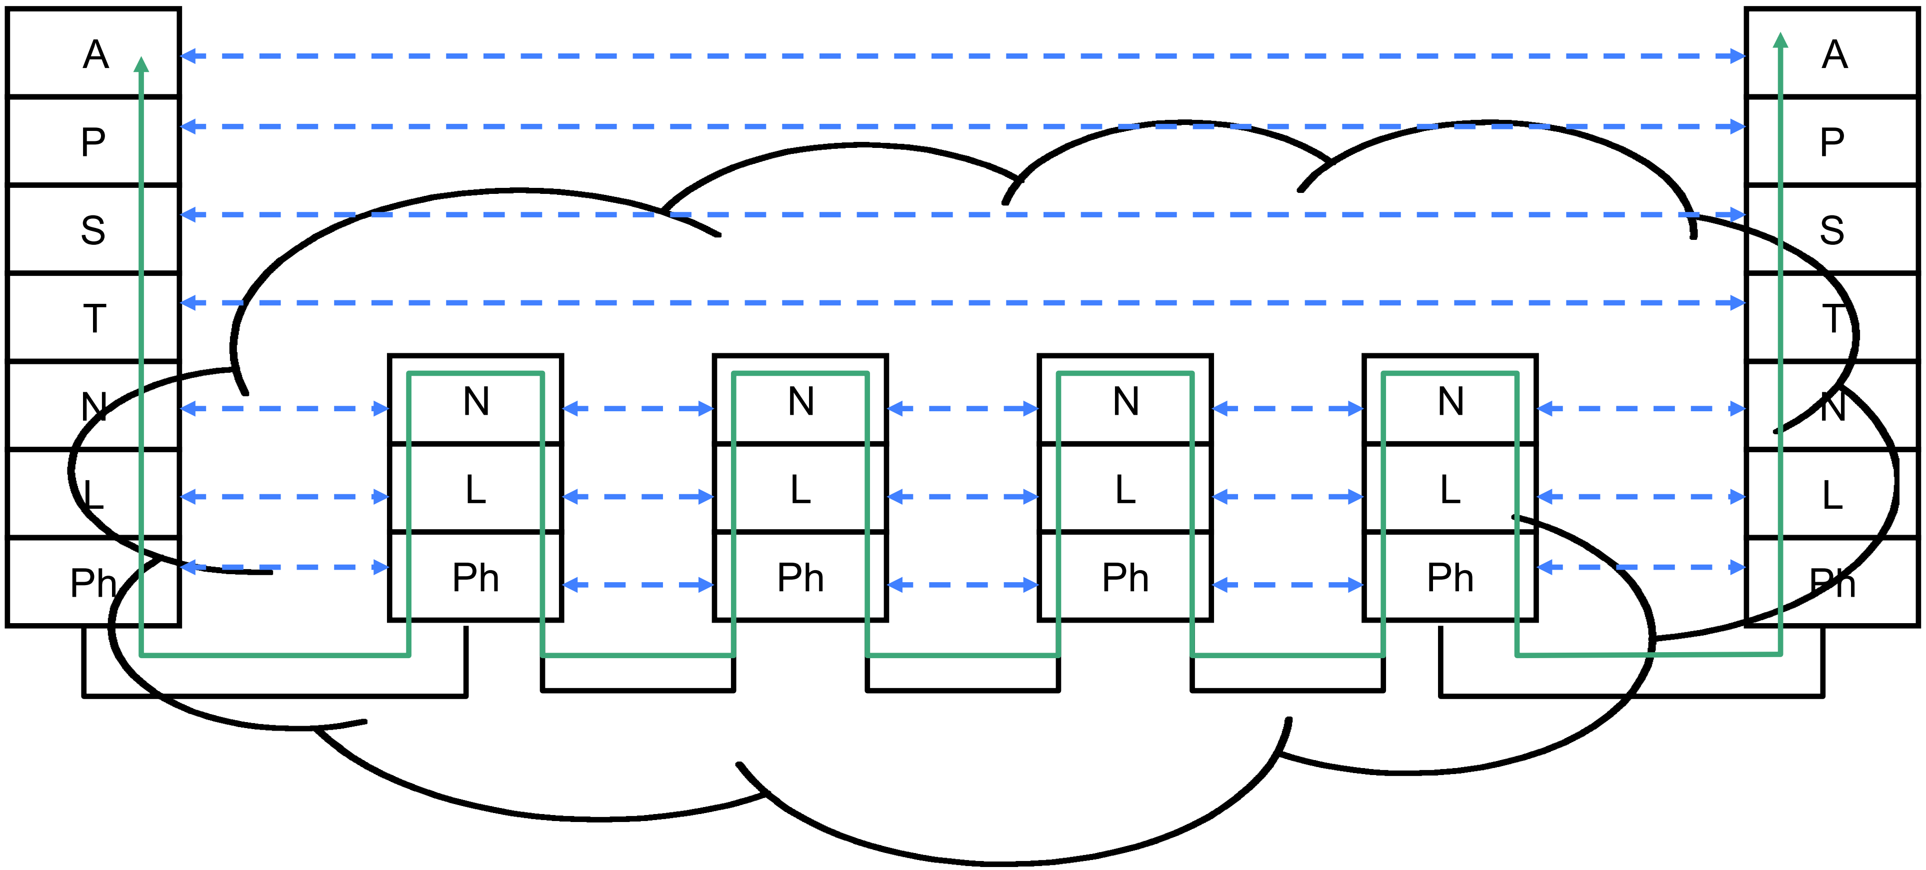

A	application	N	network
P	presentation	L	(data) link
S	session	Ph	physical
T	transport		

- - → logical (peer-to-peer) communication

→ actual end-to-end message path

PRM protocol reference model

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'**

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

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.

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 for IP) is bound to **all** the bits in the local and remote IP address.
- **IP address used as an identifier.**

Network layer

- IP address bits are used in **routing**:
  - IP address prefix, e.g.
    138.251.195.61/24
    means that 138.251.61 (also known as 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**

Cheltenham Research

University of St Andrews

UCL

Interface names

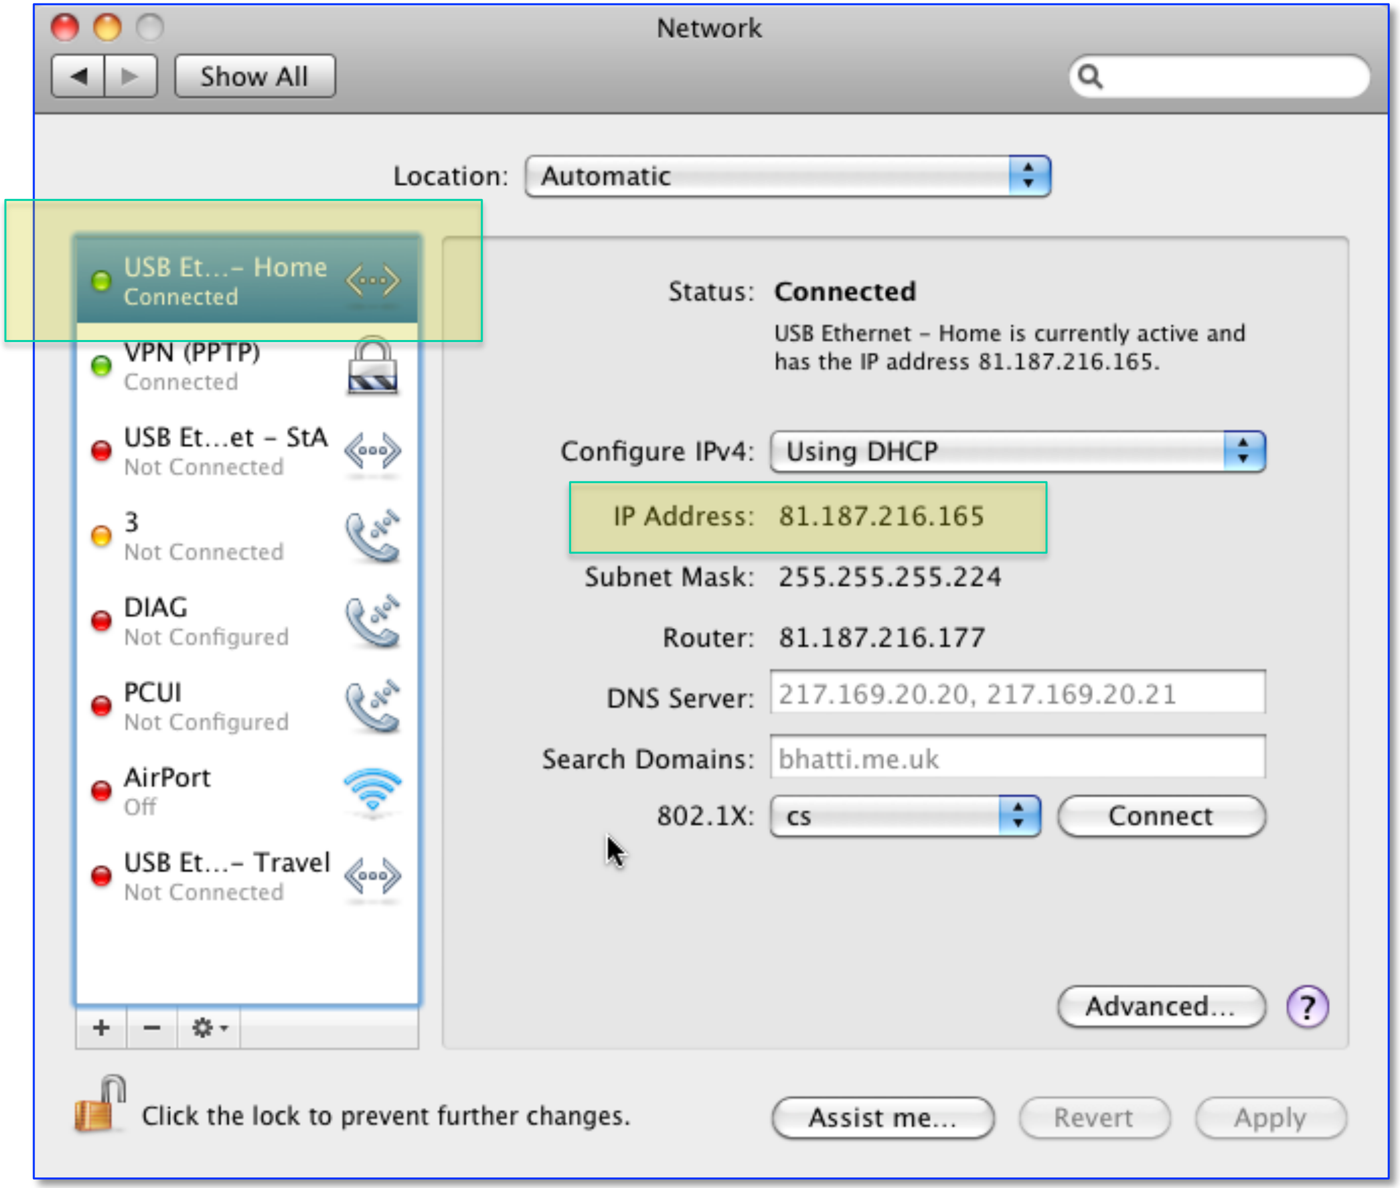

Layers are entangled

Layer	IP
Application	IP address or FQDN
Transport	IP address + port no.
Network	IP address
(Interface)	IP address

This is a serious problem for the future …

(New) Requirements

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

Priorities for ILNP

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 today can be used without modification.
2. Reuse the existing core protocol deployment as much as possible.
3. Try to limit the impact on current applications (but some applications might break).
4. The end system stack will need to change, but changes should run in parallel with current stack.

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.

RFC4984 (Sep 2007) [2]

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.

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 possibly problems with IP address as today.
- Lots of wisdom:
  - IENs 19, 23, 31, 46

Schedule

1. Problem space
2. **Introduction to ILNP**
3. Using ILNP
4. Issues and related work
5. Wrap-up

Identifier / Locator Network Protocol

- This is a work in progress:
  - http://ilnp.cs.st-andrews.ac.uk/
- Focus on network and transport layers (for now)
- This talk - ILNP 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 on 8+8/GSE

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 - another talk!

IPv6 addresses and ILNPv6

IPv6 (as in RFC3587):

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

IPv6 routing (address) prefix

same syntax, different semantics

ILNPv6:

```
|              64 bits                |       64 bits              |
+---+---------------------+-----------+----------------------------+
|              Locator                |    Node Identifier         |
+---+---------------------+-----------+----------------------------+
```

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

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.
- **Upper layer protocols bind only to Identifier.**

Locators and Identifiers [2]

- Locator, L:
  - **Can change** value during the lifetime of a transport session.
  - Multiple Locators can be used simultaneously.
- Identifier, I:
  - **Remains constant** during the lifetime of a transport session.
  - Multiple Identifiers can be used simultaneously by a node, but not for the same session.

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.

Naming: IP vs. ILNP

Protocol Layer	IP	ILNP
Application	FQDN or IP address	FQDN
Transport	IP address (+ port number)	Identifier (+ port number)
Network	IP address	Locator
(Interface)	IP address	(dynamic mapping)
	Entanglement ☹	**Separation ☺**

FQDN = fully qualified domain name

Schedule

1. Problem space
2. Introduction to ILNP
3. **Using ILNP**
4. Issues and related work
5. Wrap-up

Examples of ILNP usage

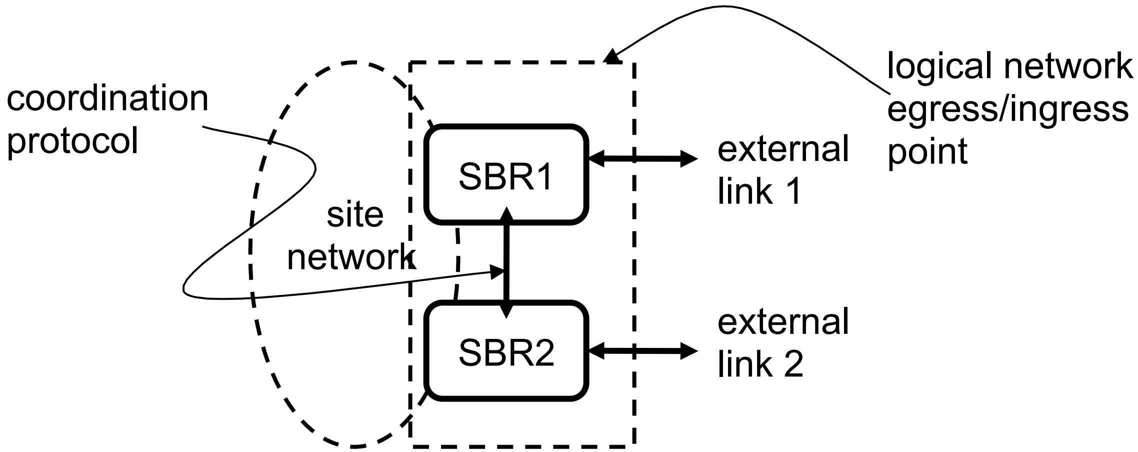

SBR = site border router

NAT in IPv4 and IPv6

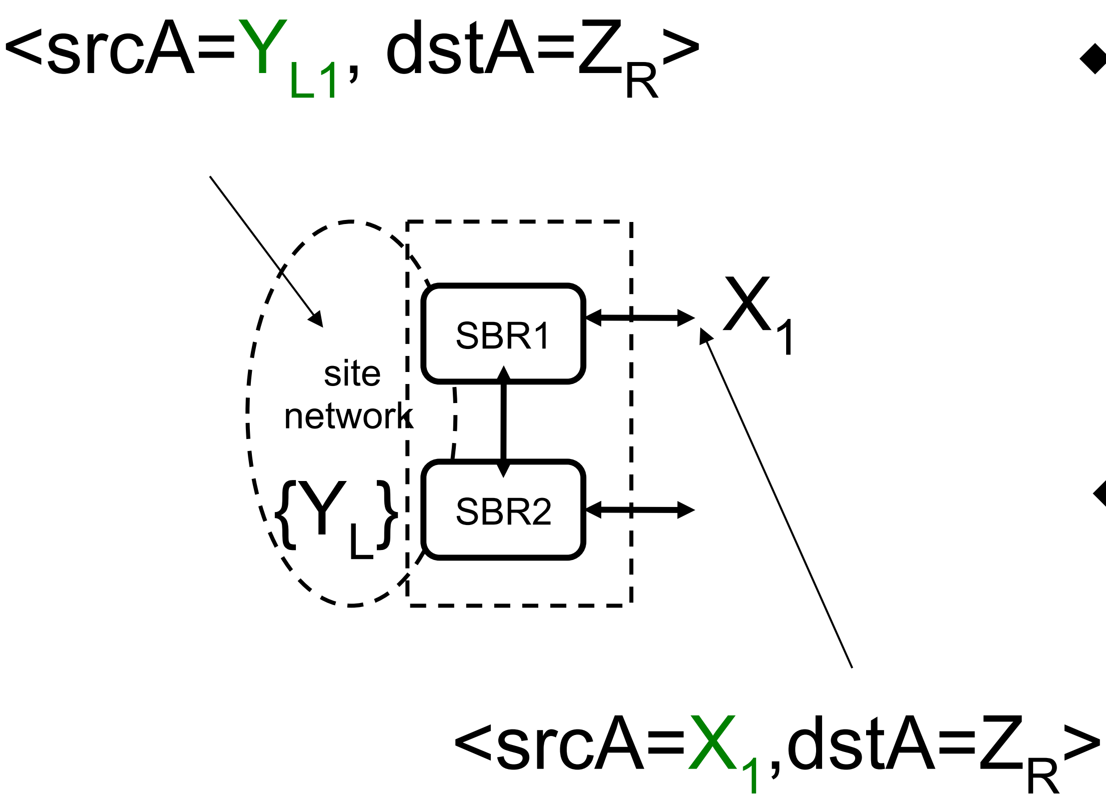

- **NAT allows address reuse for a site**:
  - single address shared amongst many hosts
- End-to-end view is lost, as identity namespace has a discontinuity at the SBR

NAT in ILNPv6

<srcI=I_1, dstI=I_R>
<srcL=L_L,dstL=L_R>

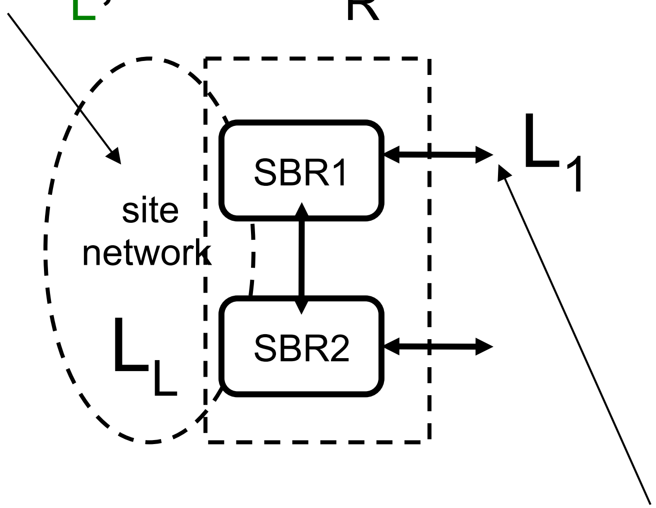

<srcL=L_1,dstL=L_R>
<srcI=I_1, dstI=I_R>

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

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.

Mobile **networks** in ILNP [1]

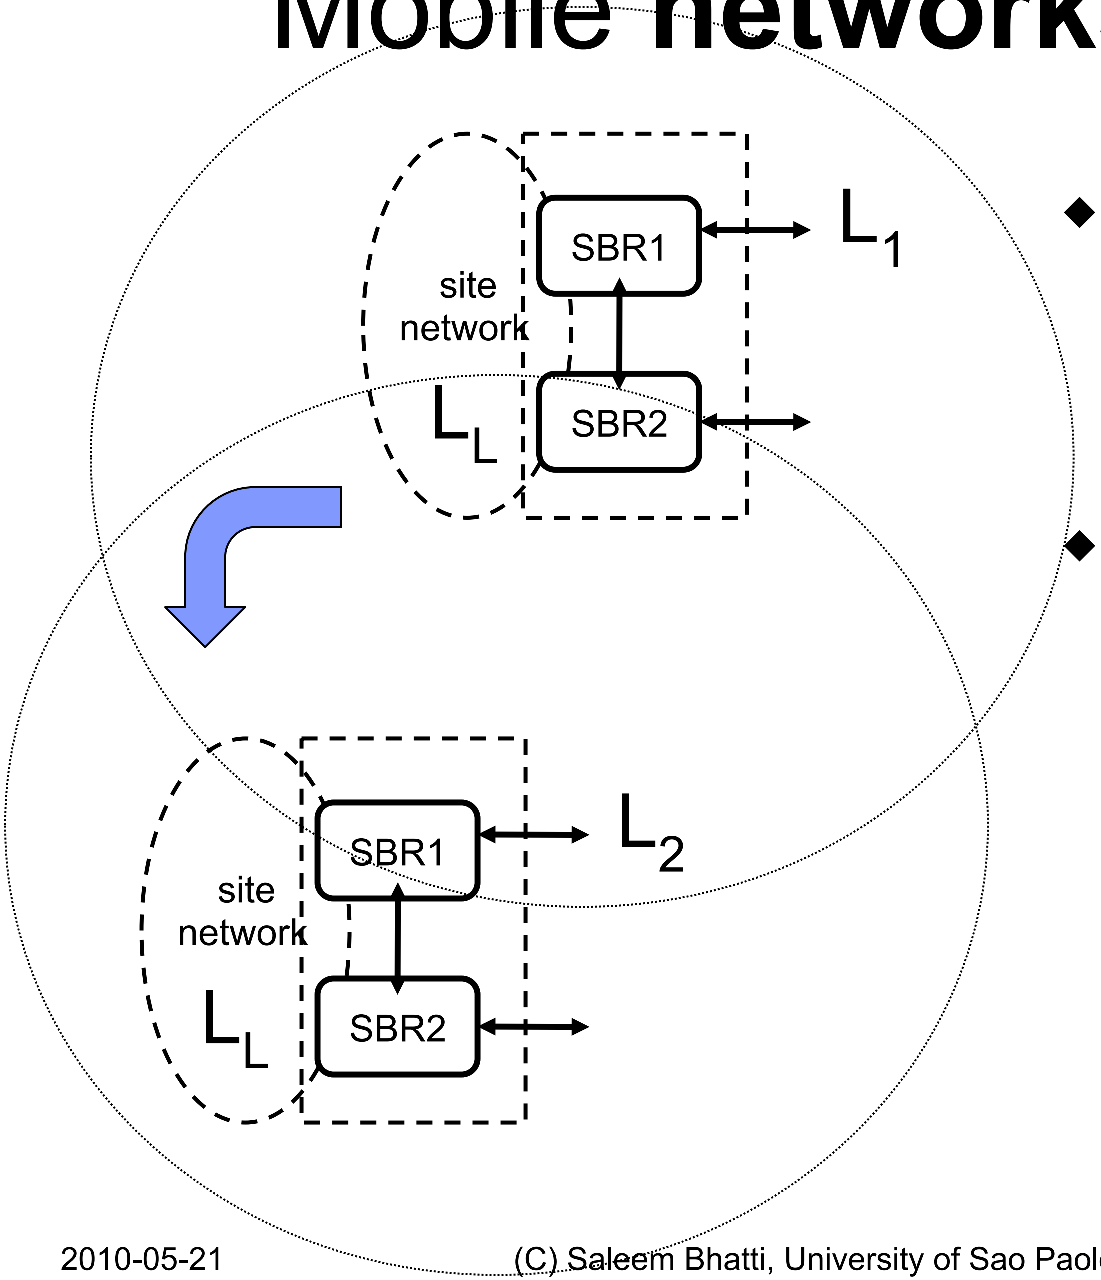

- Use NAT to 'hide' the movement to internal nodes.
- SBR changes Locator value as the mobile network moves:
  - Sends Locator Update (LU) messages to correspondents.
  - Updates DNS.

Mobile **networks** in ILNPv6 [2]

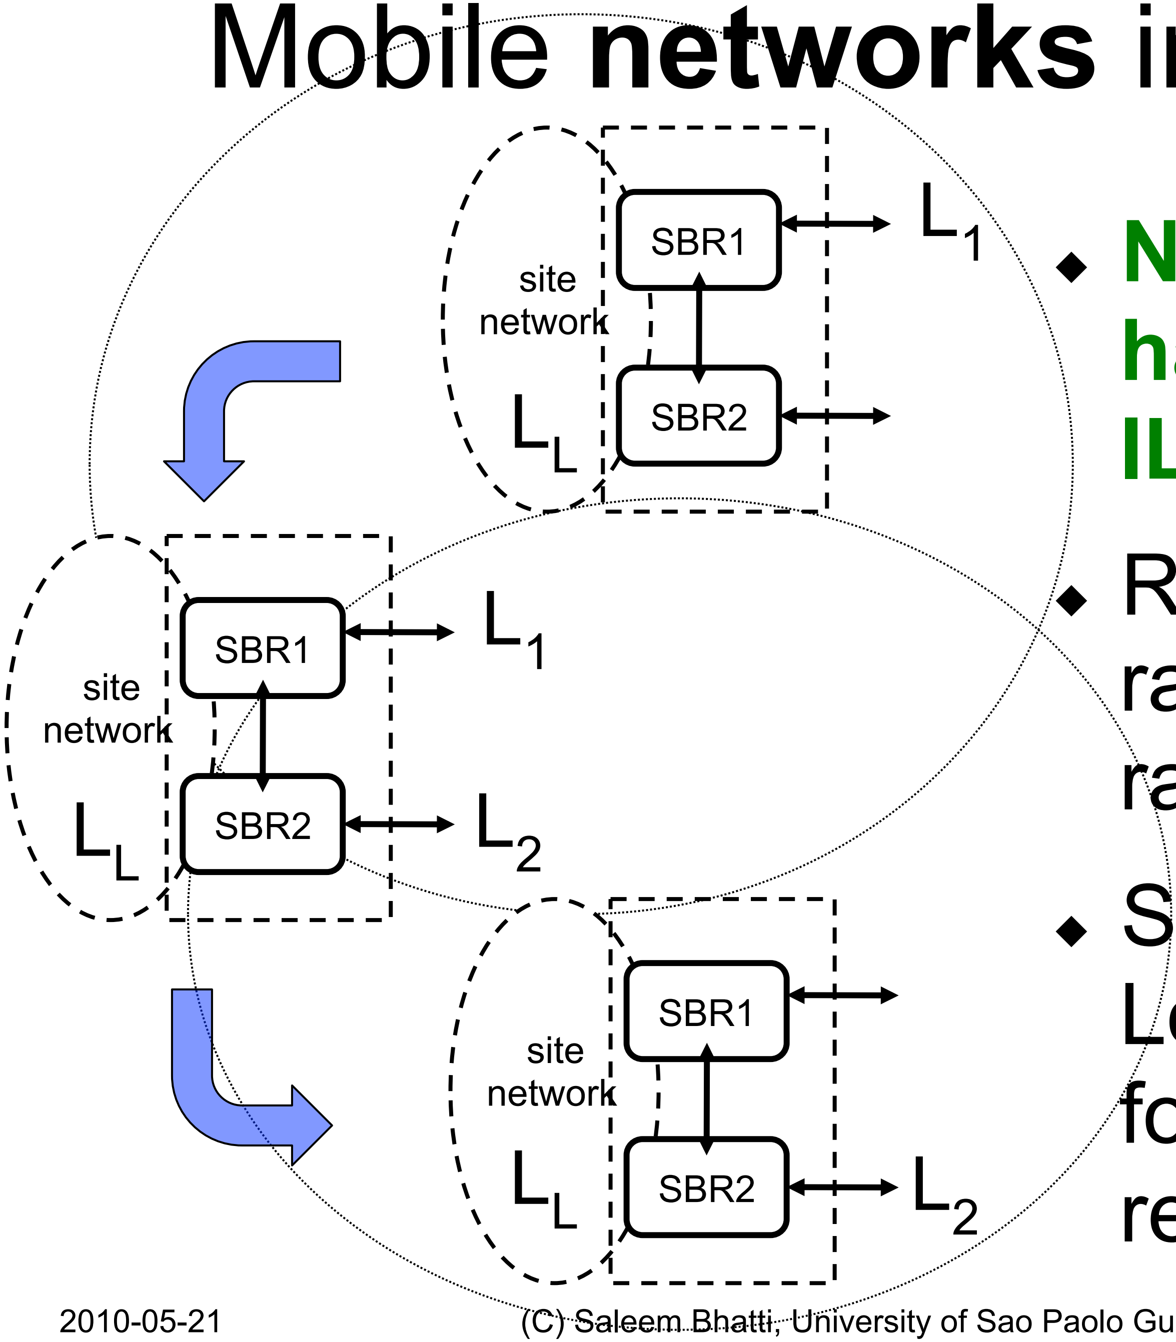

- **Network layer soft-hand-off possible in ILNP.**
- Requires at least 2 radio channels (or 2 radio interfaces).
- SBRs can handle Locator rewriting and forwarding as required.

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**.

Cheltenham Research

University of St Andrews

UCL

Multi-homing in ILNPv6 [1]

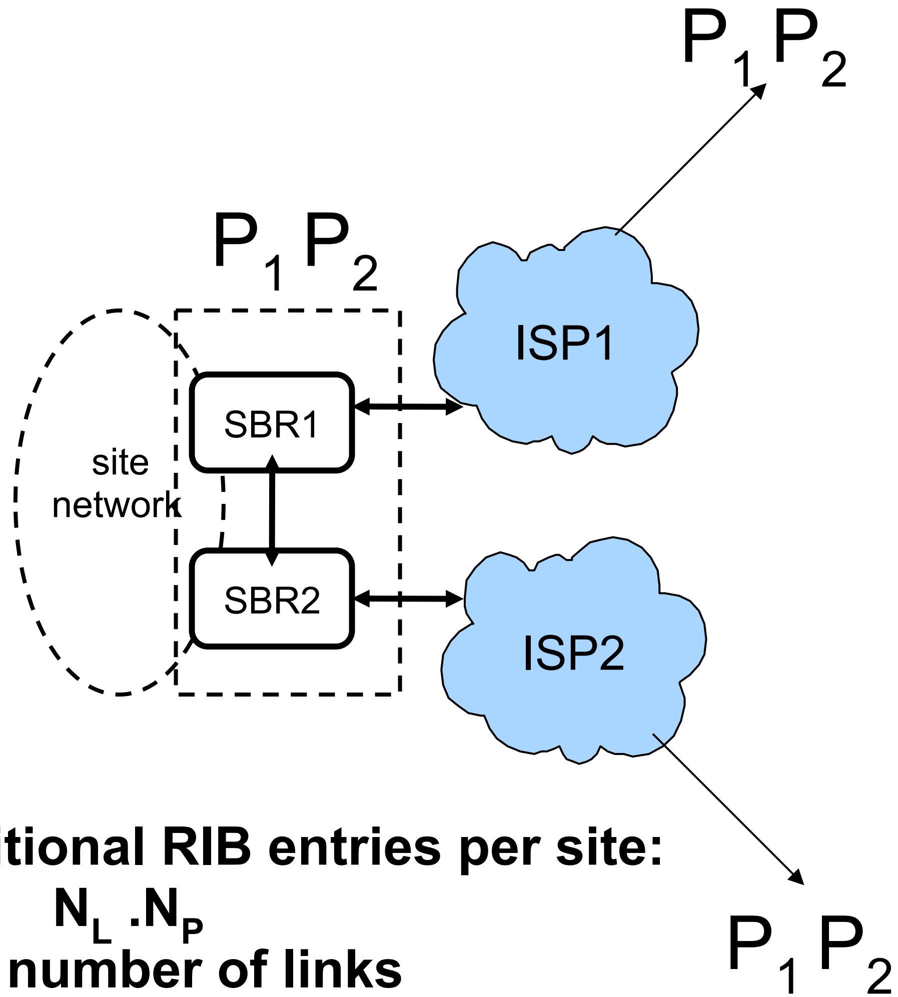

- For IP today, Provider Independent (PI) prefixes are popular:
  - Prefix ≡ identity.
  - No renumbering.
- Multi-homing prefixes can lead to bloat in the RIB of the DFZ:
  - **Non-aggregateable prefixes.**

Additional RIB entries per site:

$N_L \cdot N_P$

N_L = number of links

N_P = number of prefixes

Multi-homing in ILNPv6 [2]

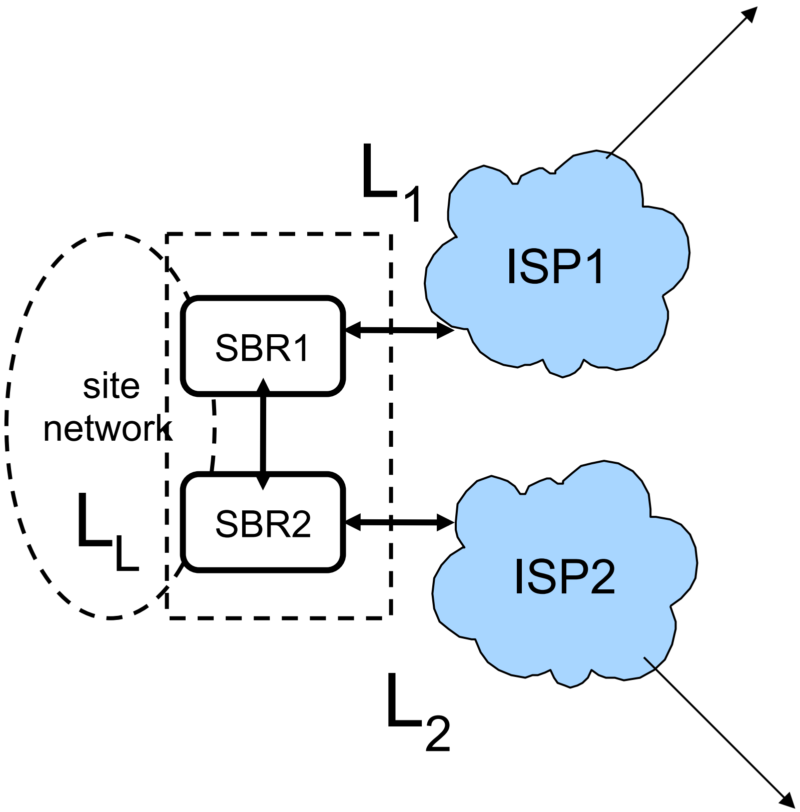

No additional RIB entries

- 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.

Traffic Engineering in ILNP

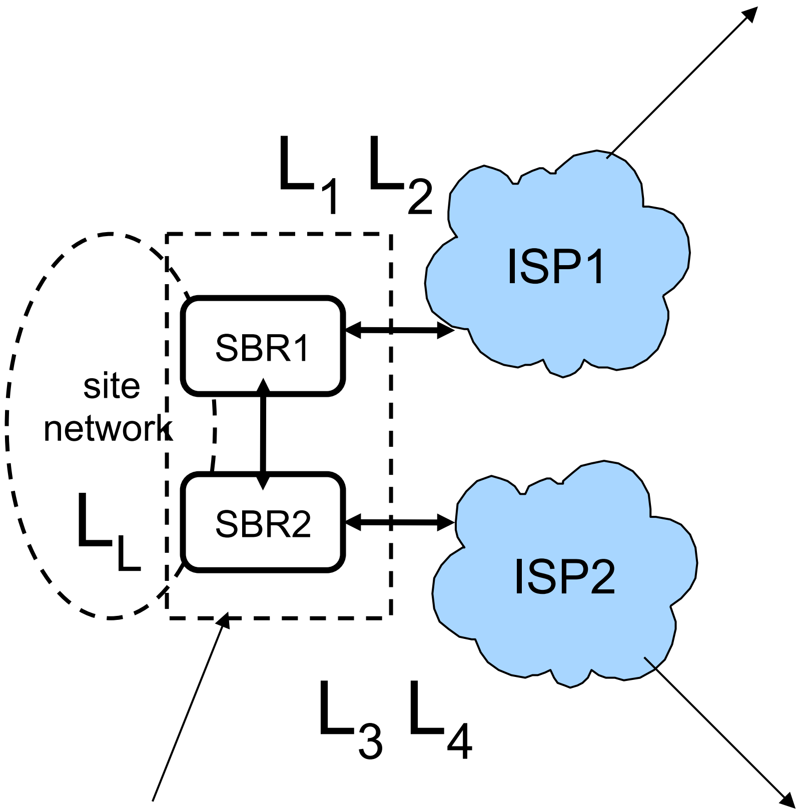

Policy mechanisms to decide on which links packets are forwarded.

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

DNS enhancements required

Name	DNS Type	Definition
Identifier	ID	Names a Node
Locator	L64	Names a subnet
Reverse Locator	PTRL	FQDN for the DNS Server responsible for subnet L
Reverse Identifier	PTRI	FQDN for the I that is present at subnet L
Locator Pointer	LP	Forward pointer from FQDN to an L record

FQDN = fully qualified domain name

Schedule

1. Problem space
2. Introduction to ILNP
3. Using ILNP
4. **Issues and related work**
5. Wrap-up

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.

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?

Past relevant work

- Our work is based on the following key ideas:
  - IEN1 (1977): separate names for layer 3 & layer 4
  - Dave Clark (c.1995): email to public IRTF list proposing to split the IPv6 address into 2 pieces.
  - 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.

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)

Schedule

1. Problem space
2. Introduction to ILNP
3. Using ILNP
4. Issues and related work
5. **Wrap-up**

Next steps

- **Build it.**
  IRTF recommendation: IETF WG for ILNP
  - StA plan to write a BSD stack and Linux stack.
- **Test it.**
  Try it out in the lab and over the national UK academic IPv6 core network.
- **Give it away for free.**
  We want other people to use it. ☺
- ILNPv4 ... ?
  Retrofit to IPv4 is possible but troublesome. ☹

Summary

- **ILNP: separate location and identity.**
- **ILNPv6: can work on existing IPv6 networks.**
- We claim **harmonised** functionality:
  - localised addressing
  - mobility (host and network)
  - traffic engineering capability
  - multi-homing without increased RIB in DFZ
  - end-to-end packet level security
- Now we have to build it!

Thank you! Questions?

- ILNP information:
  - http://ilnp.cs.st-andrews.ac.uk/
  - Papers online, implementation(s) in progress!
- Partners:
  - Ran Atkinson <ran.atkinson@gmail.com>
  - **Saleem Bhatti <saleem@cs.st-andrews.ac.uk>**
  - Steve Hailes <s.hailes@cs.ucl.ac.uk>