



What's in a name?

Ran Atkinson, Extreme Networks, USA

Saleem Bhatti, University of St Andrews, UK

Steve Hailes, University College London, UK

Apologies to William Shakespeare

Juliet:

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

“Romeo and Juliet”, W. Shakespeare,
Act 2, Scene 2

**Juliet was not worried about names, but perhaps
for the Internet, we should be ... ?**

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



Schedule

1. Problem space

2. Introduction to ILNP

3. Using ILNP

4. Issues and related work

5. Wrap-up

- ♦ I have planned to talk for ~50mins and then have questions, but I am happy to take questions as I go.

Requirements

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

Priorities

- ◆ 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.



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/>`
 - ◆ So not everything is fully cooked yet ...
- ◆ The output is classed in the body of work that is currently labelled as examining the “identifier / locator split” for network addressing.
- ◆ This talk - ILNP as a possible parallel system on the current Internet infrastructure:
 - ◆ **Heavily influenced by a bottom up approach.**


ILNPv6

- ◆ A set of 'extensions' to IPv6:
 - ◆ Uses same packet format as IPv6.
 - ◆ 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 per RFC3587):

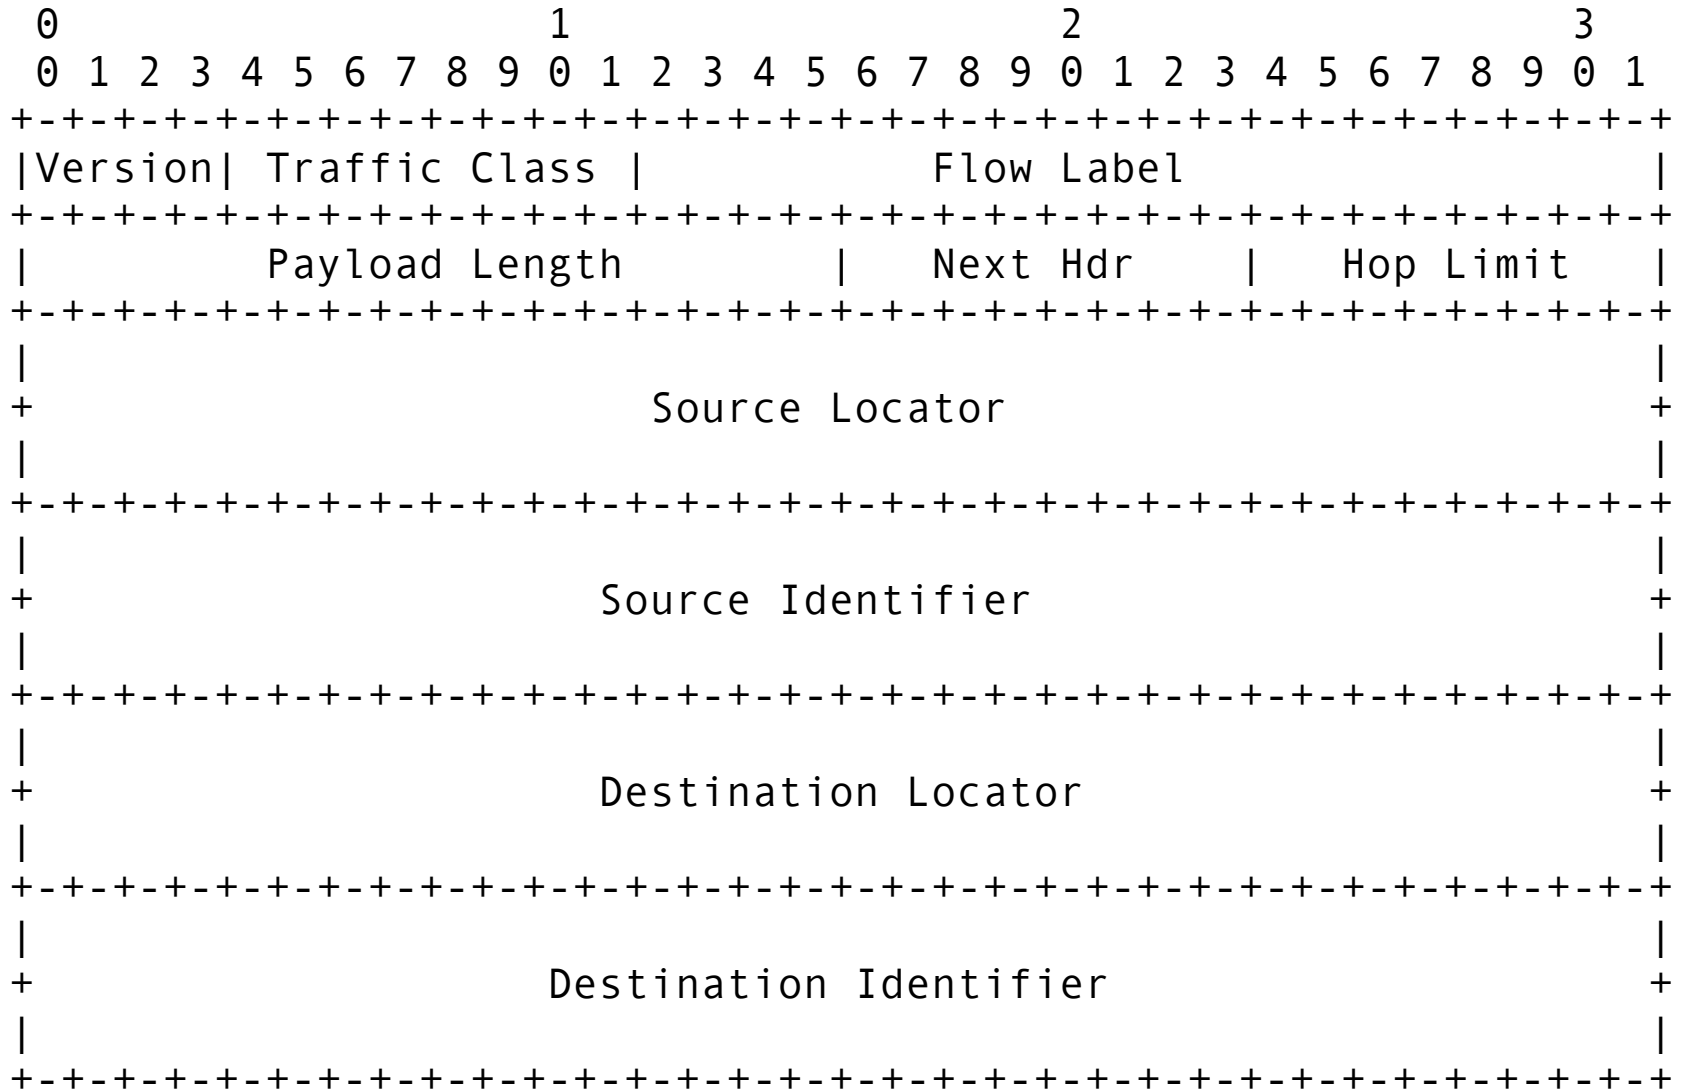


ILNPv6:





ILNPv6 packet header



Locators and Identifiers [1]

- ◆ Locator, L:
 - ◆ Names a (sub)network (as today's network prefix).
 - ◆ **Topologically significant.**
 - ◆ Used only for routing and forwarding in the core.
- ◆ Identifier, I:
 - ◆ Names a logical/virtual/physical node, does **not** name an interface.
 - ◆ **Is not topologically significant.**
- ◆ **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 discovery.
- ◆ Identifier, I:
 - ◆ Default value: a node uses bits from a local interface to form an EUI-64 address which is used as an Identifier for that node.
 - ◆ Other interesting possibilities ... (work in progress) .
 - ◆ Strictly, needs to be unique within the scope of a given Locator value – global uniqueness is nice, 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
Link	MAC address	MAC address

Entanglement :- (Separation :-)

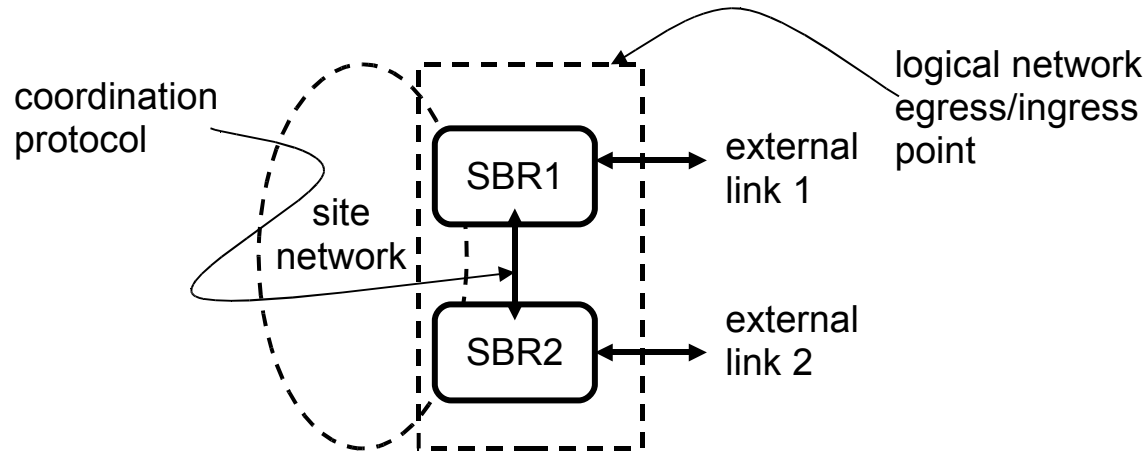
FQDN = fully qualified domain name



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

R. Atkinson, S. Bhatti, and S. Hailes,
“Harmonised Resilience, Multi-homing and Mobility Capability for IP”,
27th IEEE Military Communications Conference (MILCOM 2008),
San Diego, CA, USA, November 2008.


Previous work:

R. Atkinson, S. Bhatti, and S. Hailes,
 A Proposal for Unifying Mobility with Multi-Homing, NAT, and Security,
 5th ACM International Workshop on Mobility Management and Wireless
 Access (MobiWAC), Chania, Crete, October 2007.

IPsec

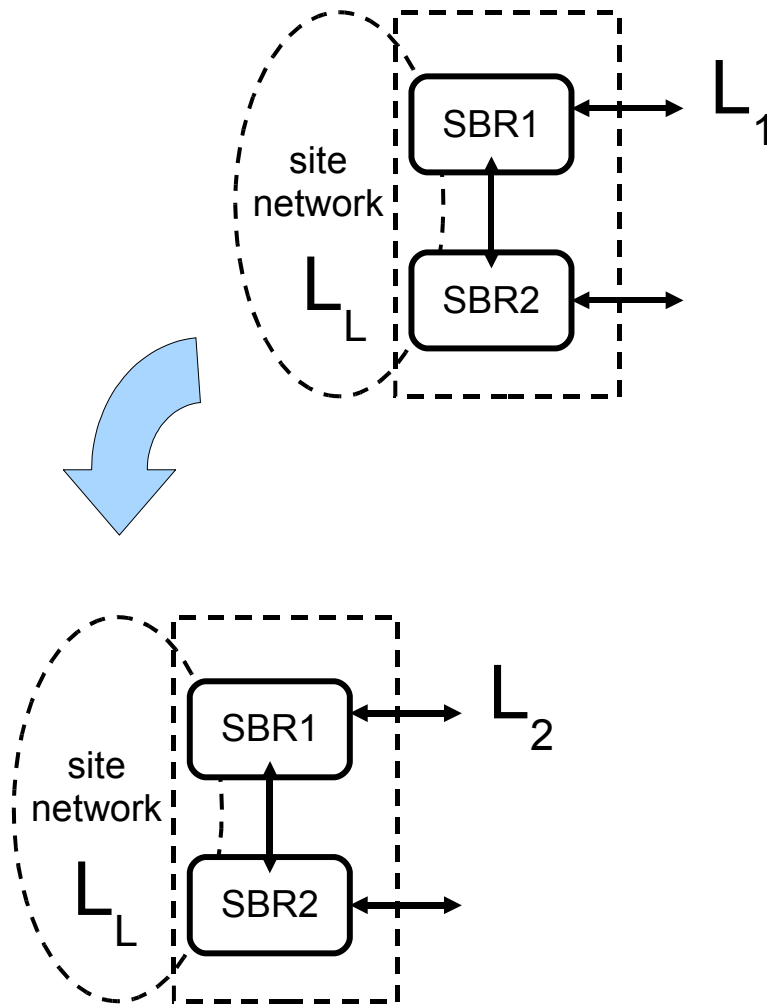
- ♦ IPsec currently uses the whole of the IP address for binding a Security Association (SA).
- ♦ In ILNP, the SA binds only to the Identifier, I:
 - ♦ I remains constant throughout the session.
 - ♦ L value can change (for whatever reason) while the session is in progress.
 - ♦ As long as I does not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

NAT in ILNPv6



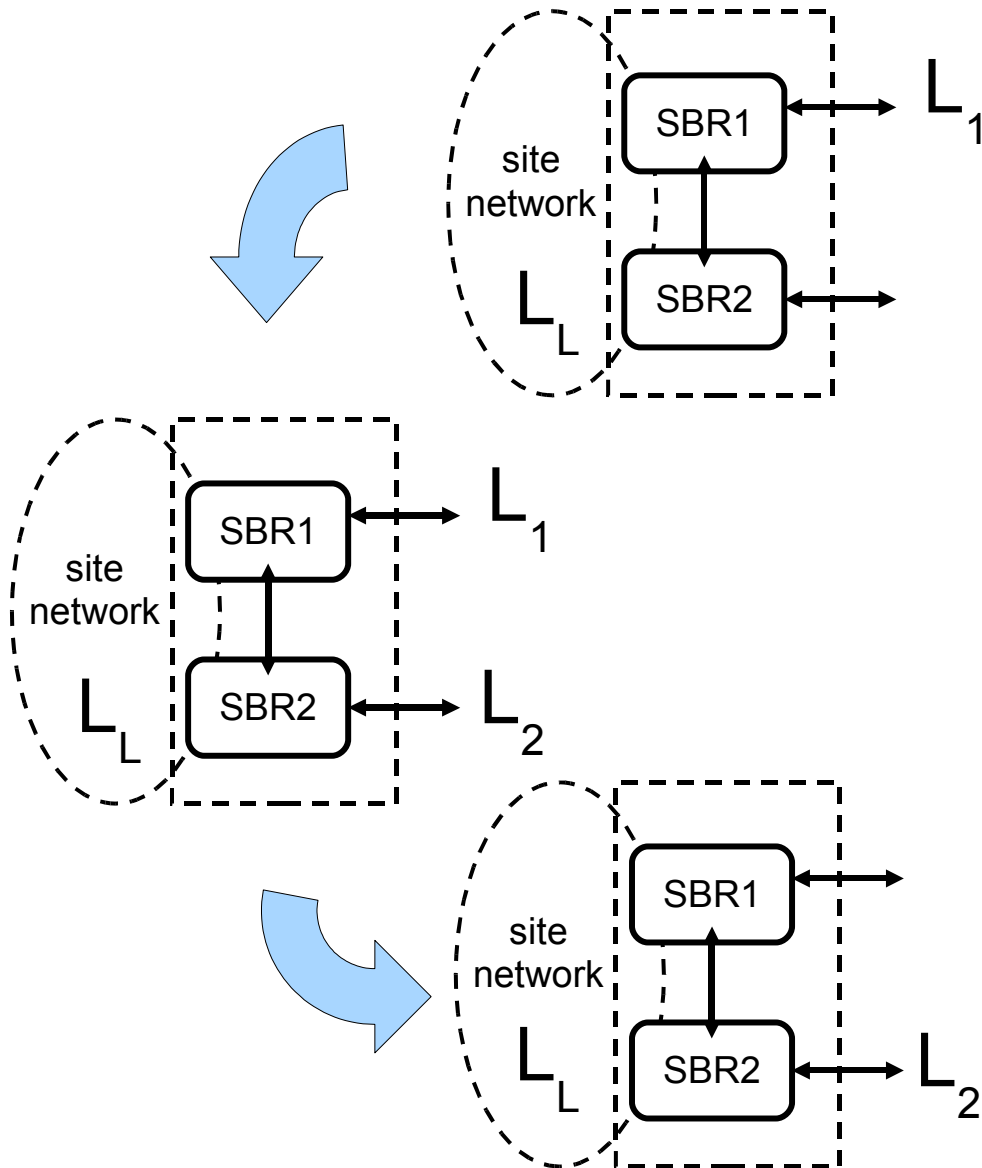
- ♦ NAT is now a feature not a hack:
 - ♦ L is not part of the end system transport session state.
- ♦ SBRs can perform **Locator rewriting** without affecting end-system transport session state.

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/radio-interfaces.
- SBRs can handle Locator rewriting and forwarding as required.

Mobile hosts in ILNPv6

- ♦ 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 update DNS with new Locator value.
- ♦ If cells overlap, individual MH can use multiple Locator values simultaneously for soft hand-off.

Multi-homing in ILNPv6 [1]



Additional RIB entries per site:


$$N_L \cdot N_p$$

N_L = number of links

N_p = number of prefixes

- ♦ For IP today, Provider Independent (PI) prefixes are popular:
 - ♦ Prefix \equiv identity.
 - ♦ 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 thru 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	I	Names a Node
Locator	L	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:
 - ◆ 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:
 - ◆ L 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.

Next steps

- ♦ **Build it.**
BSD stack and Linux stack.
- ♦ **Test it.**
Try it out in the lab and over the national UK academic IPv6 core network IPv6.
- ♦ **Give it away for free.**
We want other people to use.
- ♦ **ILNPv4 ... ?**
Retrofit to IPv4 could result in some engineering and performance ugliness.

Summary

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



Thank you! Questions?

- ◆ ILNP information:
 - ◆ <http://ilnp.cs.st-andrews.ac.uk/>
 - ◆ Papers online, implementation in progress!
- ◆ Partners:
 - ◆ Ran Atkinson <rja@extremenetworks.com>
 - ◆ Saleem Bhatti <saleem@cs.st-andrews.ac.uk>
 - ◆ Steve Hailes <s.hailes@cs.ucl.ac.uk>