Developing ILNP

Saleem Bhatti, University of St Andrews, UK

What is ILNP?

- Identifier Locator Network Protocol:
  - http://ilnp.cs.st-andrews.ac.uk/
- ILNP enhances Internet Protocol functionality through the using of crisp **naming**.
- March 2010: IRTF RRG recommends ILNP for development within the IETF: http://www.ietf.org/mail-archive/web/rrg/current/msg06356.html
- People:
  - Ran Atkinson (Cheltenham Research, US)
  - Saleem Bhatti (University of St Andrews, UK)

University of St Andrews

Outline

1. **New requirements.**
2. ILNP Rationale.
3. ILNP Operation.
4. Development Challenges.

(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.

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

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

- 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.**

Interface names

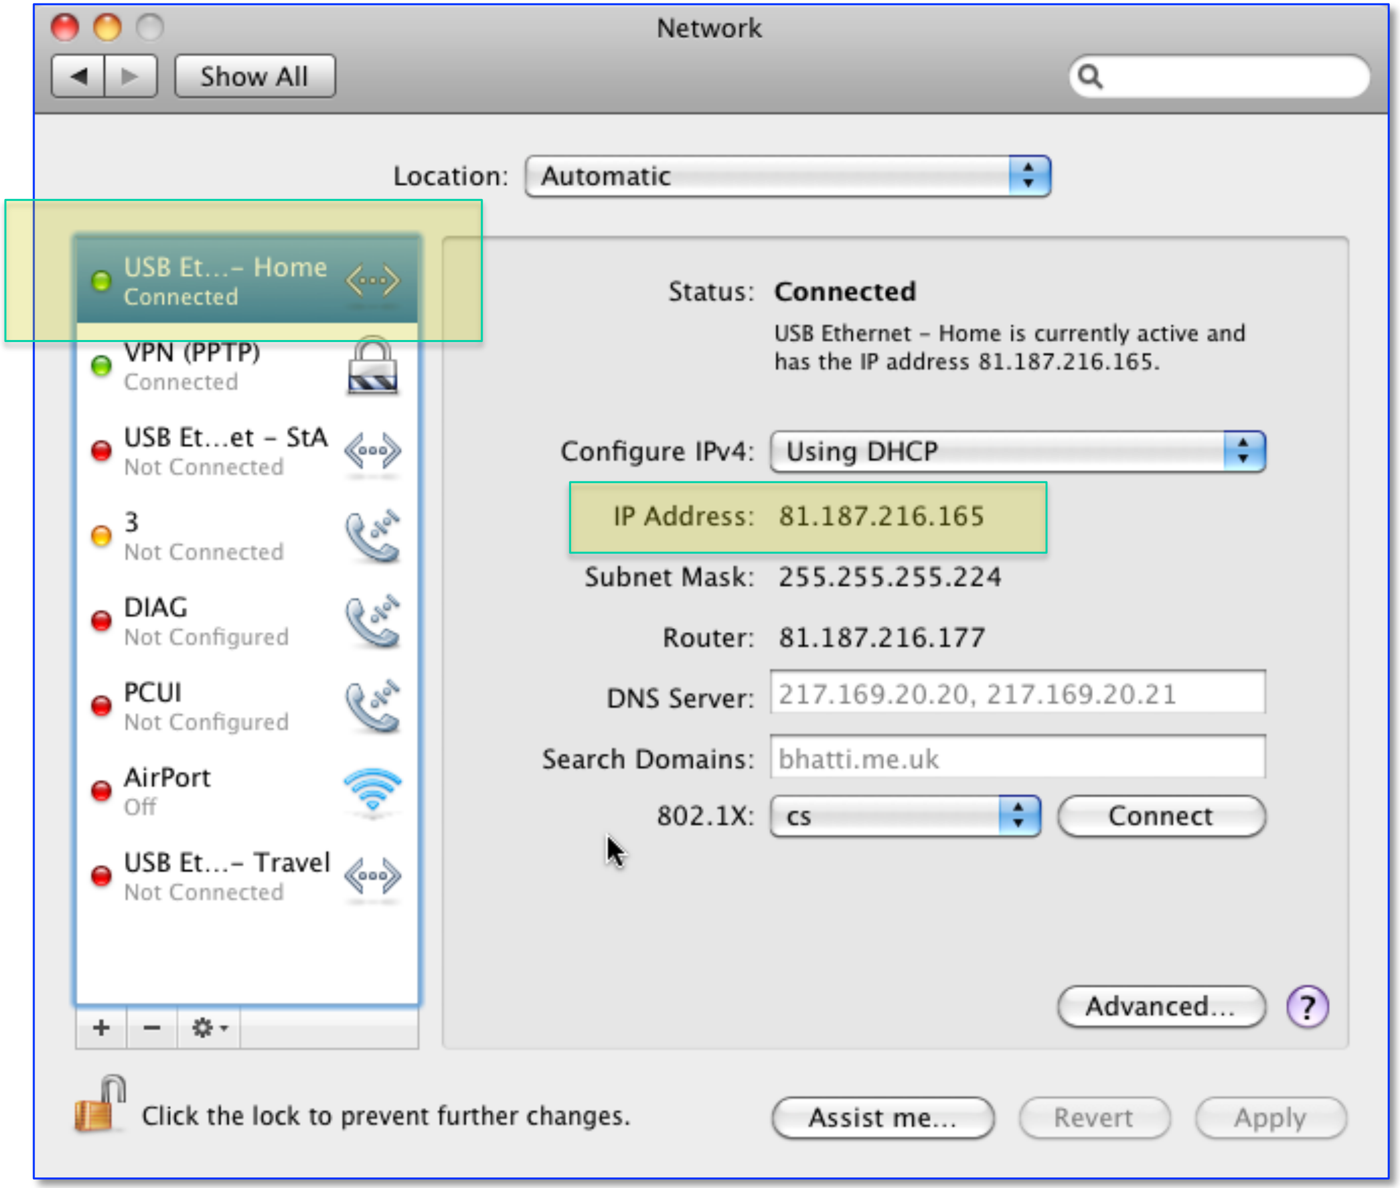

Layers are entangled

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 …

Outline

1. New requirements.
2. **ILNP Rationale.**
3. ILNP Concept of Operation.
4. Development Challenges.

University of St Andrews

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 of today can be used without modification.
2. Reuse the existing core protocol deployment as much as possible, e.g. make sue 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.

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.

University of St Andrews

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.

University of St Andrews

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

Outline

1. New requirements.
2. ILNP Rationale.
3. **ILNP Operation.**
4. Development Challenges.

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

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

University of St Andrews

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.**

University of St Andrews

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.

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

University of St Andrews

Outline

1. New requirements.
2. ILNP Rationale.
3. ILNP Operation.
4. **Development Challenges.**

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)
  - OneLab, PlanetLab (control + mgmt + monitoring?)
- **Test-bed – full implementation in OS stack:**
  - **Linux**
  - **FreeBSD**

ILNPv6 in-house testbed

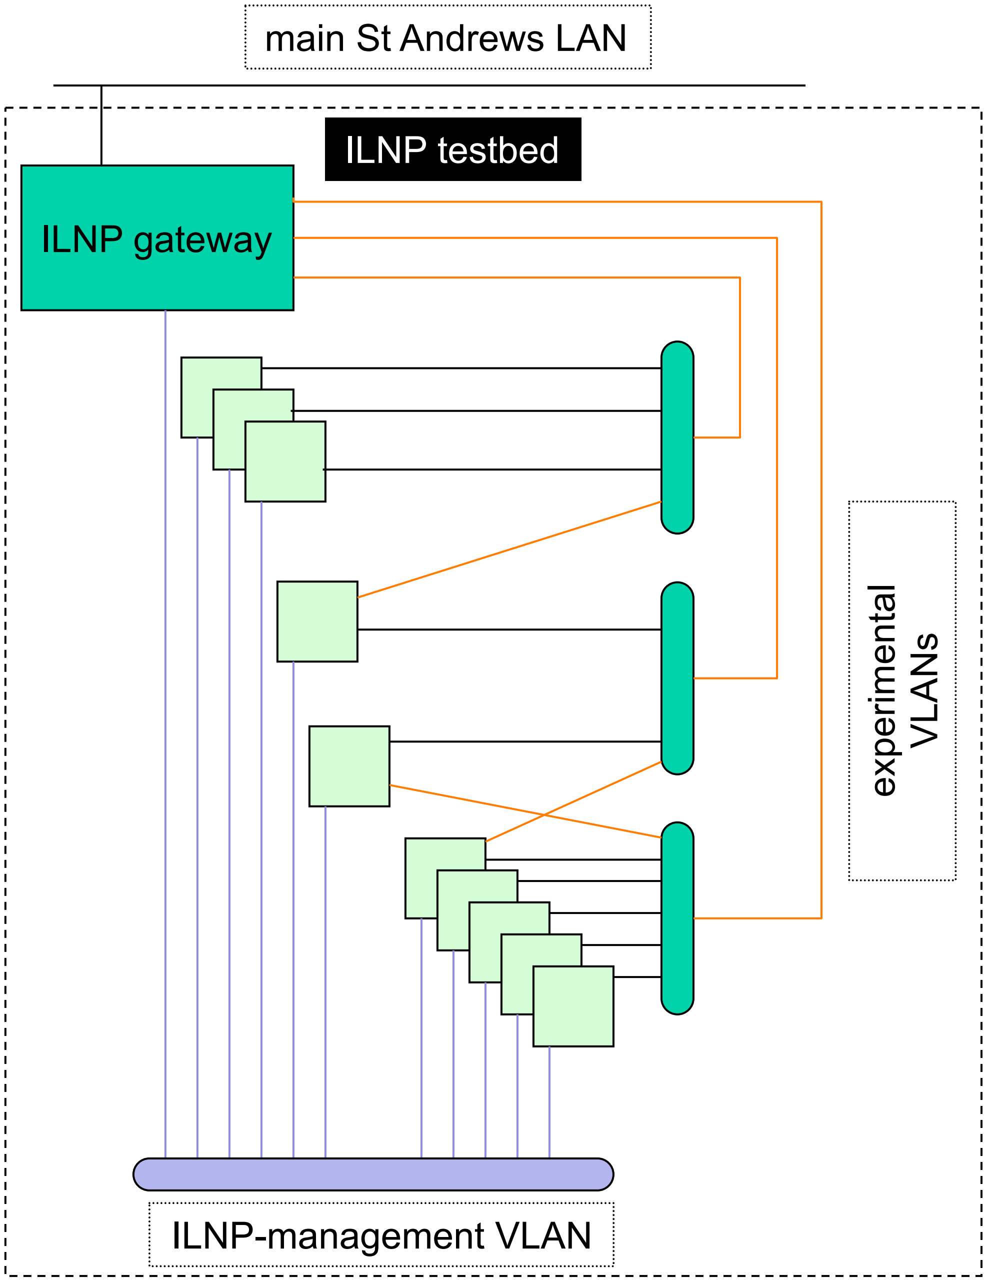

- Use of existing services:
  - e.g. use of deployed DNS and IPv6 routing.
- Exploit VMs when possible.
- Off-the-shelf equipment:
  - easy of use
  - costs
- Open source:
  - leverage existing kernel code
  - make available to community

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)

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

- 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?

Thank You!

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