University of St Andrews

Enabling mobile systems with ILNP

Saleem Bhatti, University of St Andrews, UK

ILNP in a nutshell

University of St Andrews

- Identifier Locator Network Protocol:
  - http://ilnp.cs.st-andrews.ac.uk/
  - work-in-progress!
- March 2010: IRTF RRG Chairs recommend 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)

Identifier / Locator Network Protocol

University of St Andrews

- Focus on network and transport layers (for now)
- This talk – 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. **New requirements.**
2. ILNP Rationale.
3. ILNP Operation.
4. Enabling Mobility.
5. DNS with zero TTL.

University of St Andrews

(New) Requirements

- We wish to try and support a ***harmonised solution to many* network functions**:
  - **Mobility (host and network).**
  - **Localised addressing (NAT).**
  - Multi-homing (host and site).
  - Packet-level, end-to-end security.
  - 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

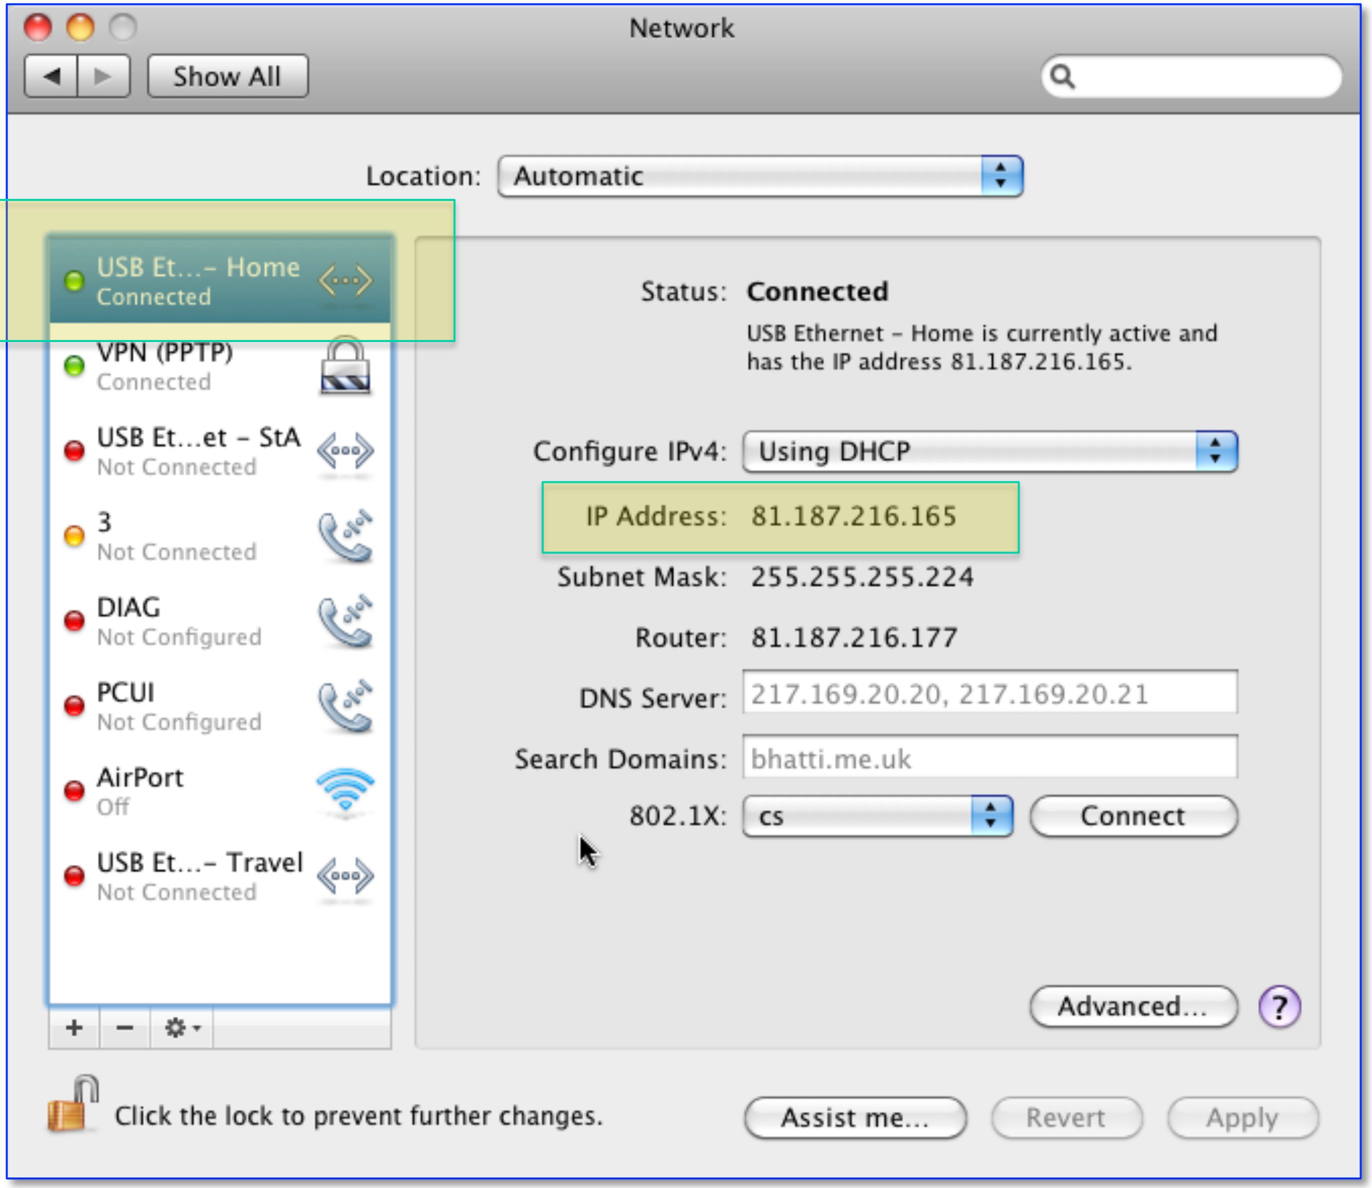

Layers are entangled

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. Enabling Mobility.
5. DNS with zero TTL.

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.

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

University of St Andrews

Outline

1. New requirements.
2. ILNP Rationale.
3. **ILNP Operation.**
4. Enabling Mobility.
5. DNS with zero TTL.

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

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.
- DNS lookups for a FQDN return ID (Identifier) and L64 (Locator) records (possibly LP record).

Using Identifier / Locator values

- Multiple Identifier (I) values:
  - I value must remain constant for a transport session
  - default is EUI-64 (ala RFC3587)
  - can use CGA (ala RFC3972)
  - can support privacy (ala RFC4941)
- Multiple Locator (L) values for a given I value:
  - IPv6 network prefix value is used for L
  - host can be multi-homed
  - **IP-layer soft hand-off for mobility**
  - **multi-path transport protocols (another talk!)**
- Network stack maintains I/L bindings.

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. **Enabling Mobility.**
5. DNS with zero TTL.

University of St Andrews

Mobility functions

- Address allocation + address management
- Routing / forwarding:
  - packets must be sent correctly to a mobile host
- Hand-off (hand-over):
  - maintain existing sessions
- ILNP needs two functions:
  1. rendezvous (for new sessions/connections)
  2. hand-off (maintain existing sessions/connections)

Mobile **hosts** in ILNPv6

- Individual mobile host (MH) obtains new Locator value from new network using IPv6 Router Advertisements.
- Maintain existing sessions/connections:
  - MH sends Locator Update (LU) messages to correspondents for existing sessions.
- Rendezvous:
  - MH updates DNS with new Locator value.
  - **This requires zero TTL for Locator values.**
- If cells overlap, MH can use multiple Locator values simultaneously for **soft hand-off**.
- (Mobility/multi-homing duality.)

Example (mobile) network

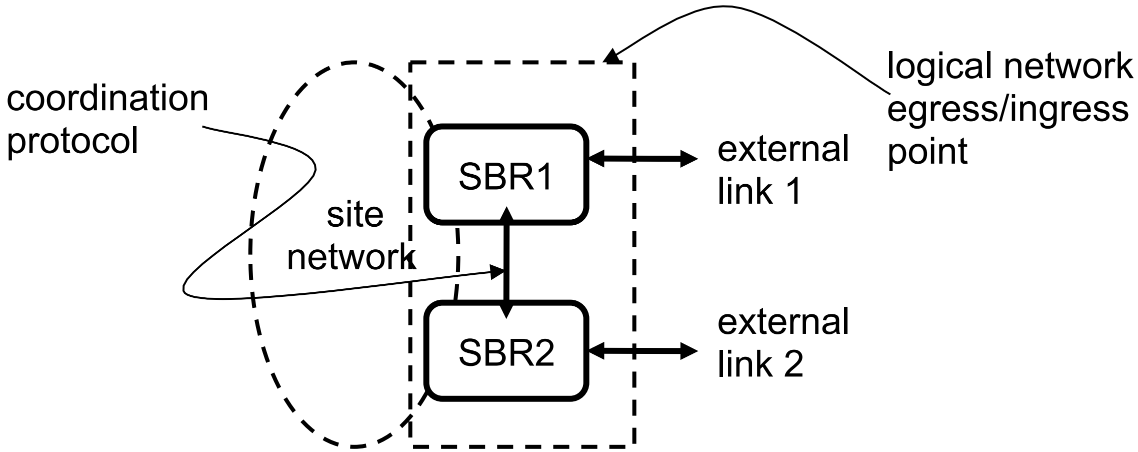

SBR = site border router

NAT in IPv4 and IPv6

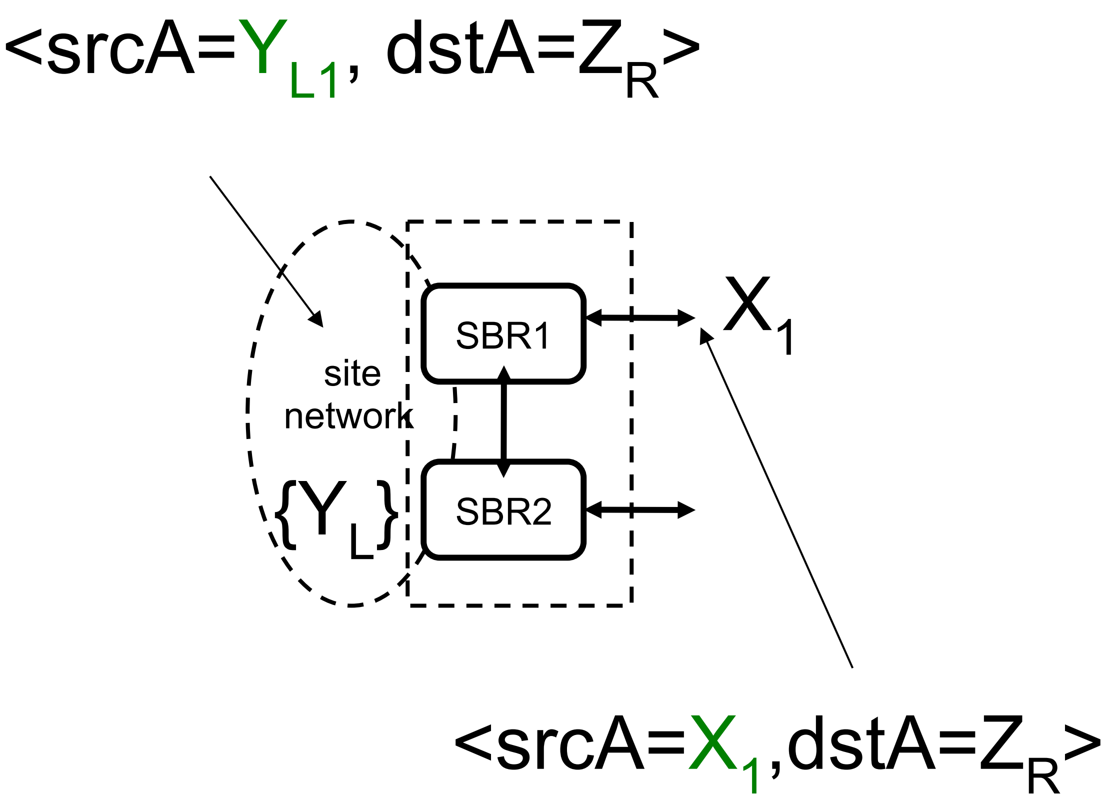

- **NAT allows address reuse for a site**:
  - single address shared amongst many hosts
- End-to-end view is lost: namespace has a discontinuity at the SBR for identity
- ($\{Y_L\}$ ala RFC1918 for IPv4 and RFC4193 for IPv6)

NAT in ILNPv6

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

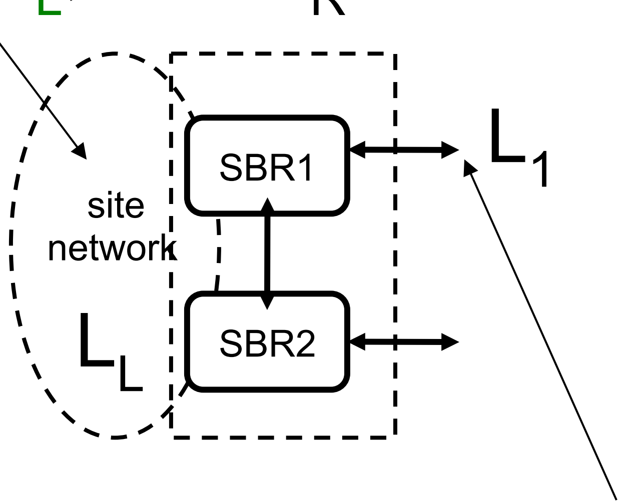

<srcL=L_1,dstL=L_R>
<srcI=I_1, dstI=I_R>

- **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.
- (L_L ala RFC4193)

Mobile **networks** in ILNP [1]

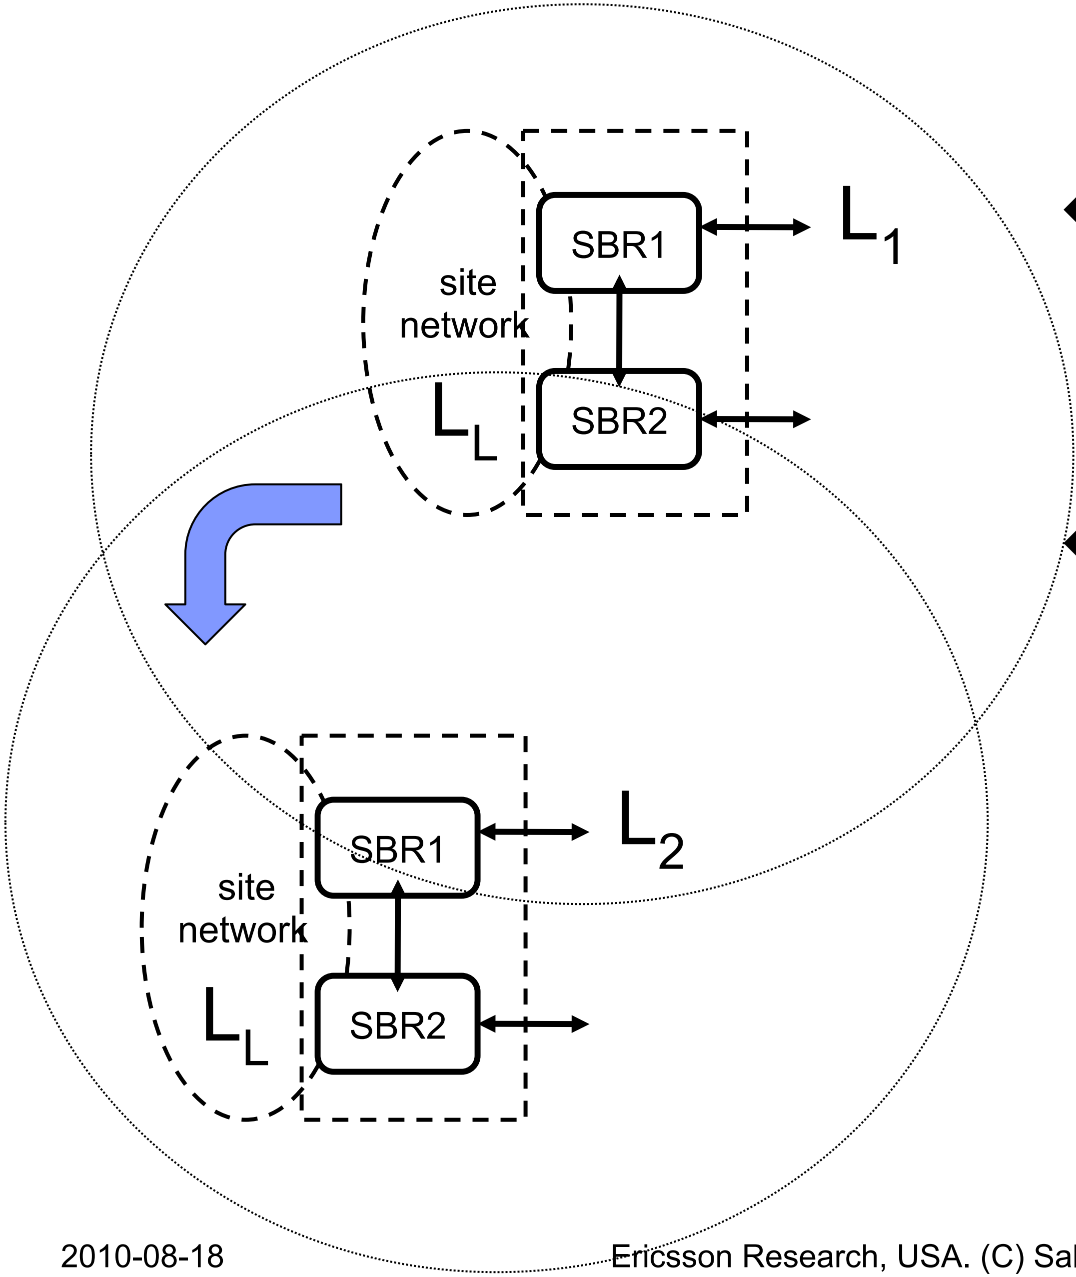

- Use NAT to 'hide' the movement to internal nodes.
- SBR changes Locator value as the mobile network moves:
  - Sends Locator Update (LU) messages to correspondents.
  - Updates DNS.

Mobile **networks** in ILNPv6 [2]

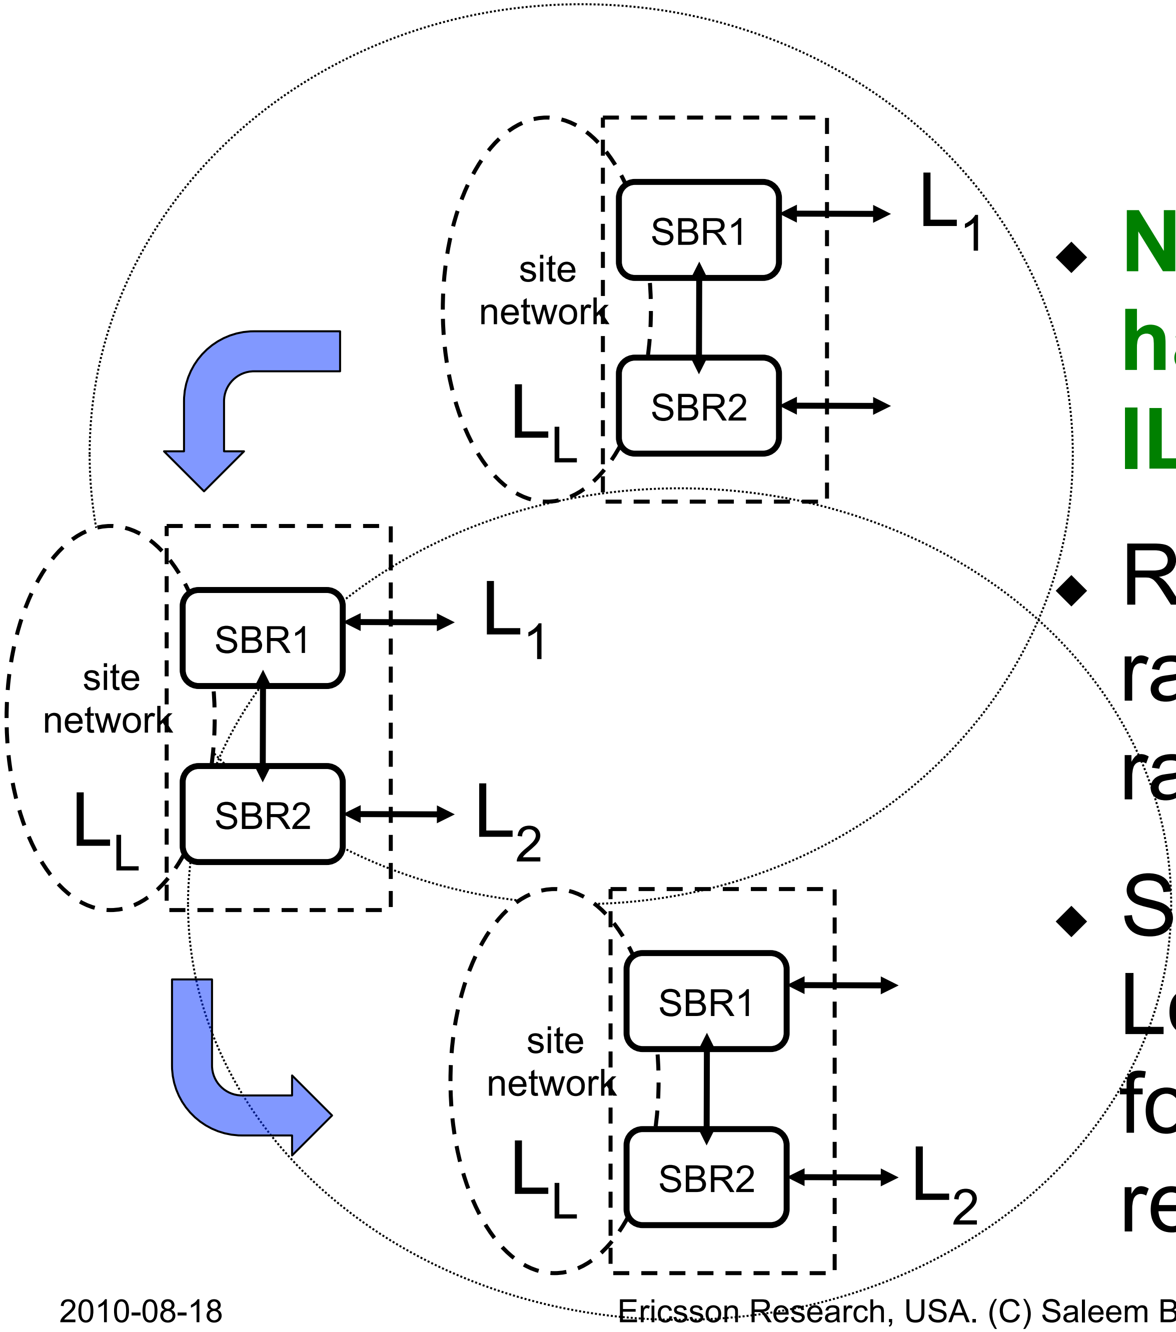

- **Network layer soft-hand-off possible in ILNP.**
- Requires at least 2 radio channels (or 2 radio interfaces).
- SBRs can handle Locator rewriting and forwarding as required.

Outline

1. New requirements.
2. ILNP Rationale.
3. ILNP Operation.
4. Enabling Mobility.
5. **DNS with zero TTL.**

University of St Andrews

(Non-)Effectiveness of DNS caching

- Jung, J., Sit, E., Balakrishnan, H., and Morris, R. 2002. *DNS performance and the effectiveness of caching*. IEEE/ACM Trans. on Networking. Vol. 10, No. 5 (Oct. 2002), pp. 589-603.
- DNS caching is ineffective for edge sites:
  - trace-driven emulation (no experiments)
  - A records could have low TTL (e.g. below 1000s)
  - such low TTL would have low impact on DNS load

DNS experiments at StA [1]

- Experiments in Q4/2009
- Modify TTL values of records in operational DNS server at School of CS, St Andrews
  - 4 DNS servers: Windows ActiveDirectory
  - ~400 DNS clients: Windows, Linux, MacOSX, BSD
- TTL values for successive 7-day periods during normal semester:
  - changed DNS TTL on ActiveDirectory
  - used TTL values 1800s, 30s, 15s, 0s
- Configured clients not to cache.

DNS experiments at StA [2]

- Passive collection of packets via port mirror:
  - *tcpdump(8)* targeting *port 53*
  - Captured all DNS packets
- Results shown on following slides are for:
  - A record requests for servers only during the capture period (relevant to ILNP, and less 'noisy' data)
  - using 1 second buckets
- Basic statistics:
  - on time-domain data
- Spectral analysis:
  - examination of request rates
- Analysis: home-brew *python* scripts, NumPy package

2009: Basic dataset meta-data

(awaiting verification)

Data set name	TTL [s]	Duration [s][1]	Total DNS packets captured[2]	Number of A record requests for 67 servers[3]
dns1800	1800	601,200	41,868,522	**2,004,133**
dns30	30	601,200	71,105,247	**2,648,796**
dn15	15	601,200	56,472,027	**3,240,675**
dns0	0	601,200	55,868,573	**4,501,590**

[1] from tcpdump timestamps, rounded to nearest second, 7 days = 604,800 seconds, less 3600s temporal guard band for TTL value changes = 601,200 seconds
[2] includes all request and response packets to/from port 53 (TCP and UDP), including erroneous requests etc
[3] servers that were active during the 4 weeks of data capture

dns1800: A record requests TTL=1800s

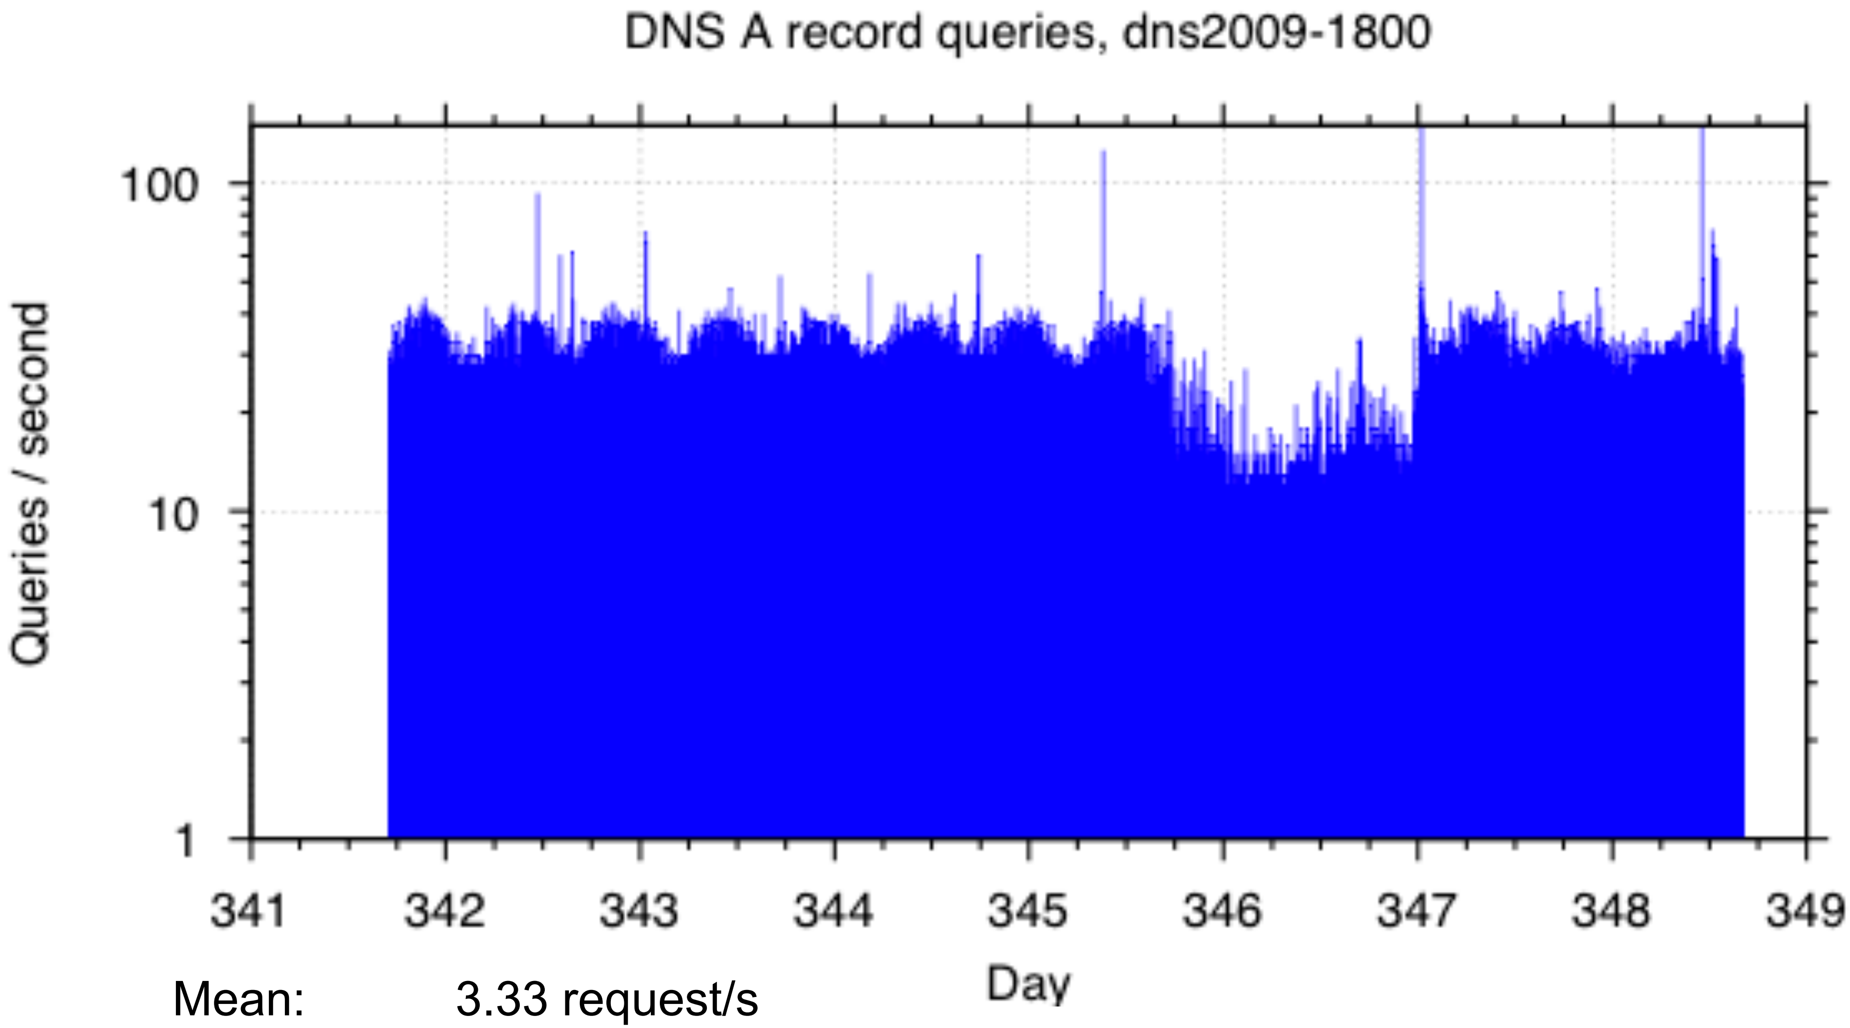

Mean: 3.33 request/s
Std Dev: 3.47 requests/s
Max: 183 requests/s

dns30: A record requests TTL=30s

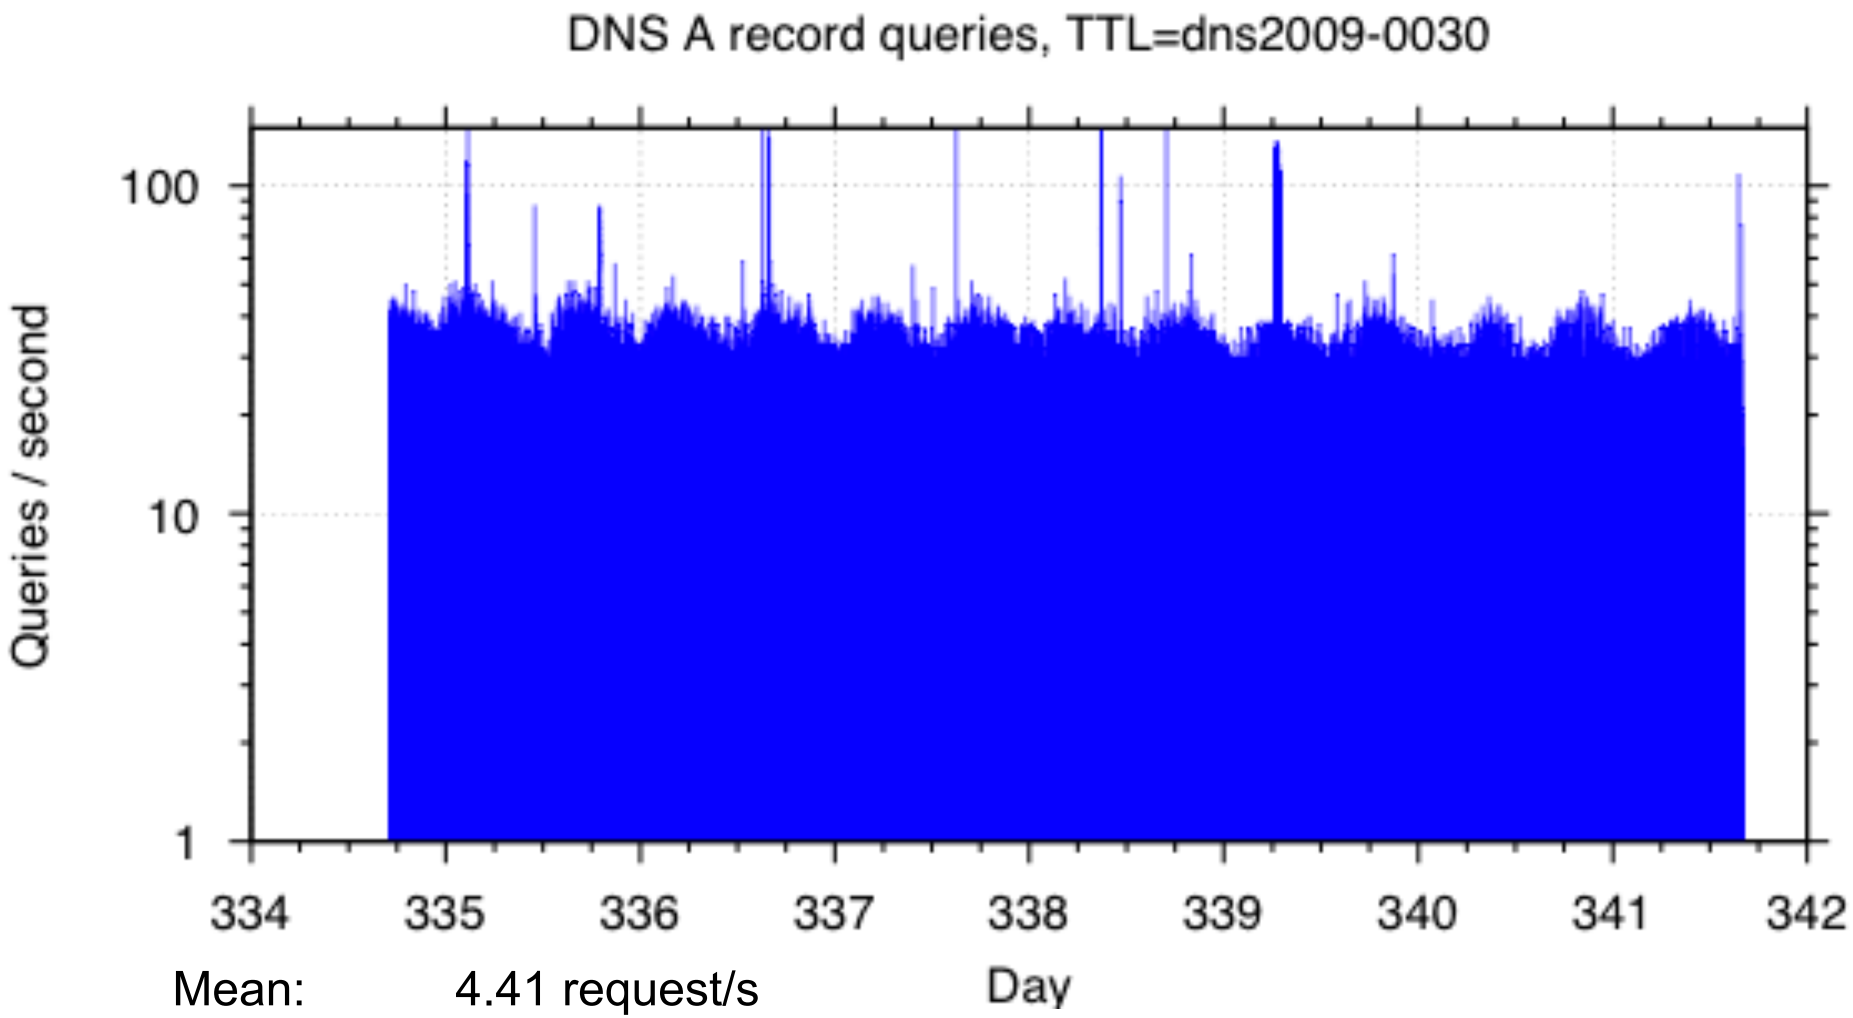

Mean: 4.41 request/s
Std Dev: 4.27 requests/s
Max: 261 requests/s

dns15: A record requests TTL=15s

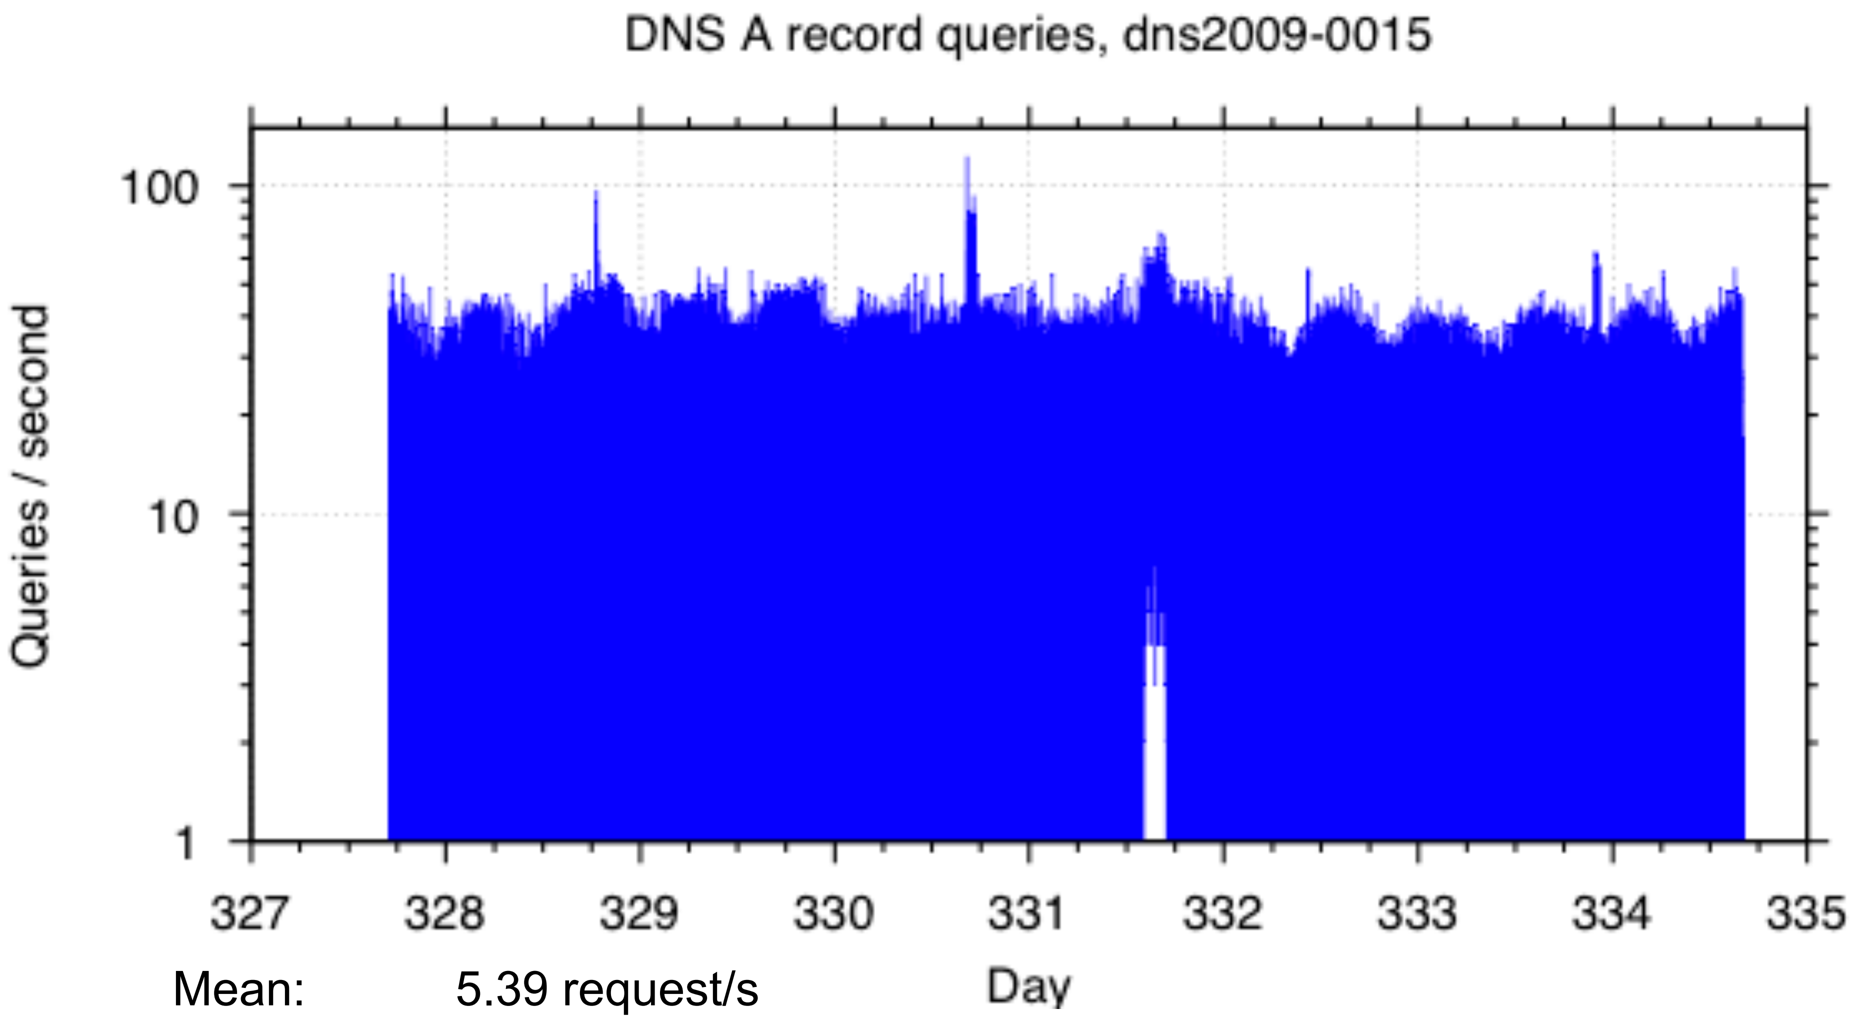

Mean: 5.39 request/s
Std Dev: 4.85 requests/s
Max: 123 requests/s

dns0: A record requests TTL=0s

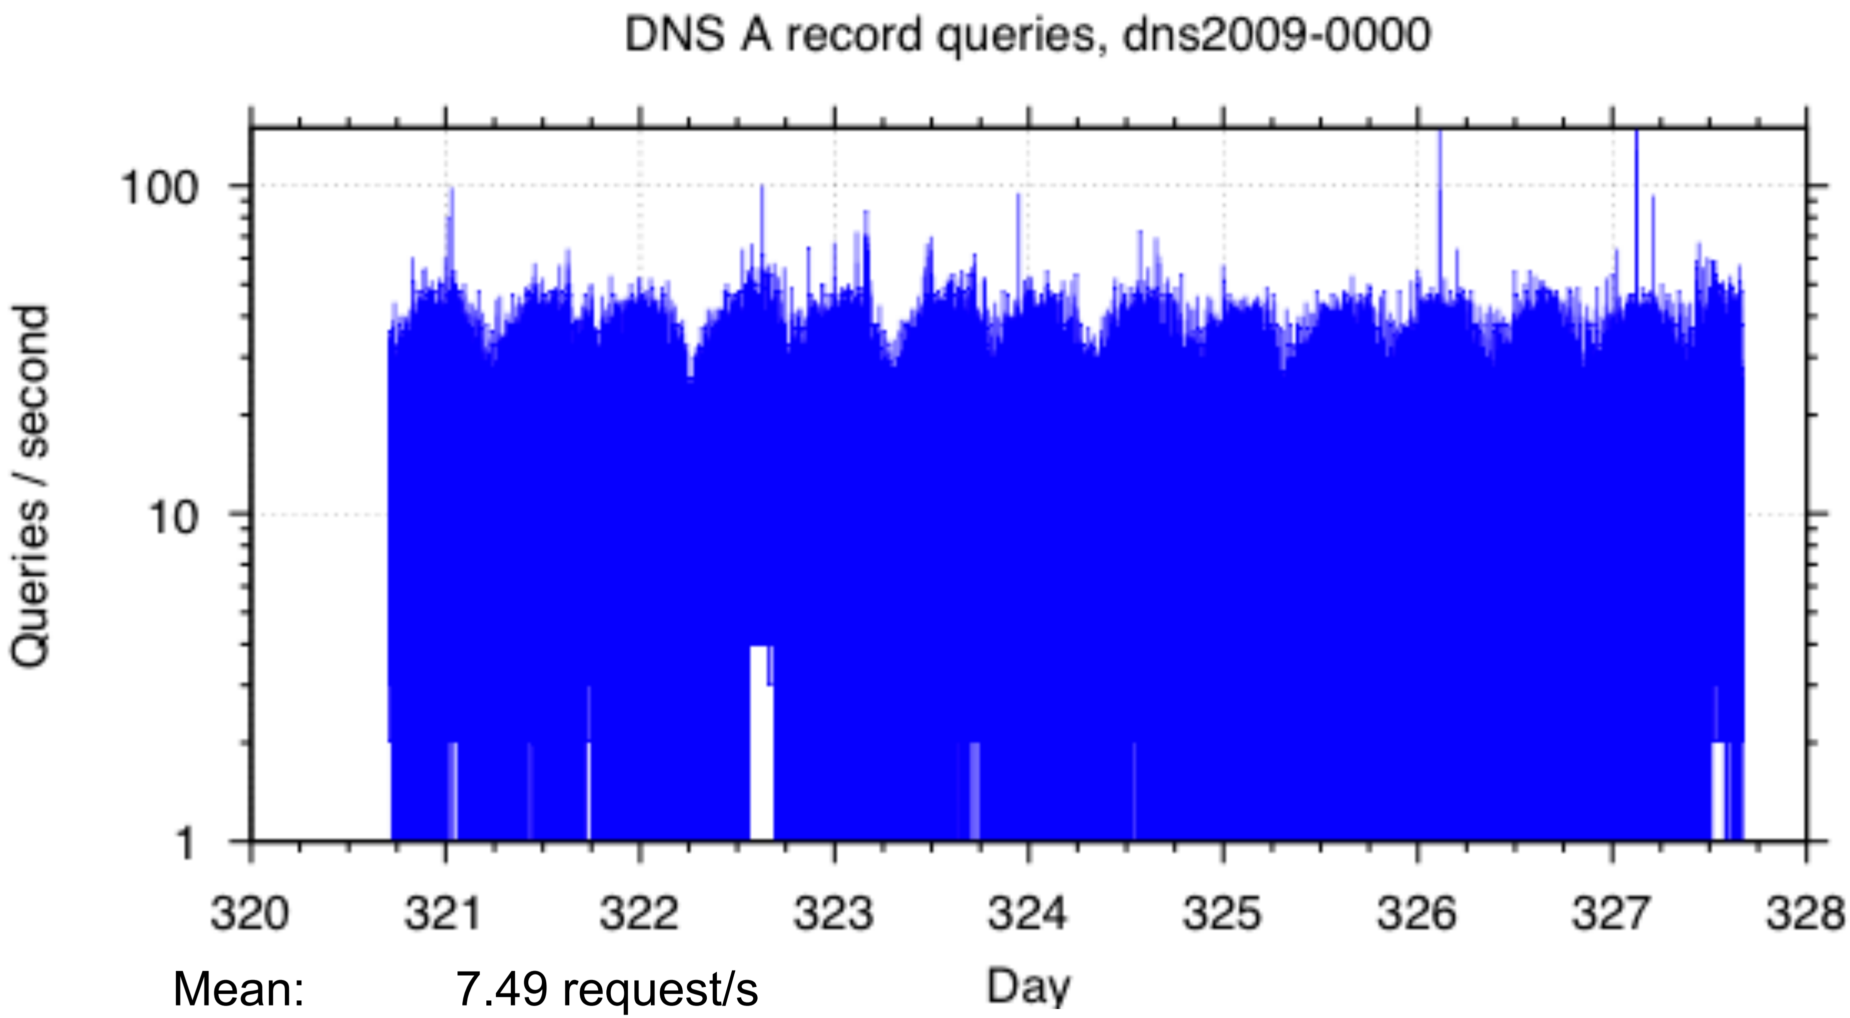

Mean: 7.49 request/s
Std Dev: 4.93 requests/s
Max: 3.69 requests/s

2009 Summary of basic statistics

(awaiting verification)

Data set name	Mean [reqs/s]	Median [reqs/s]	Std Dev [reqs/s]	Maximum [reqs/s]
dns1800	**3.33**	3	3.47	183
dns30	**4.41**	4	4.27	261
dns15	**5.39**	4	4.85	123
dns0	**7.49**	7	4.93	369

60x drop in TTL values results in **⅓ x increase** in A record requests: **0 TTL** gives (only) **2¼ x increase.**

2009 Basic spectral analysis

- Create periodogram by counting frequency of bucket sizes:
  - have used 1s bucket
  - so size of bucket is number of requests/s
- Comparison of periodogram:
  - shows changing dynamics of request rates
  - gives a better understanding of the trends in rates

2009 dns1800 periodogram

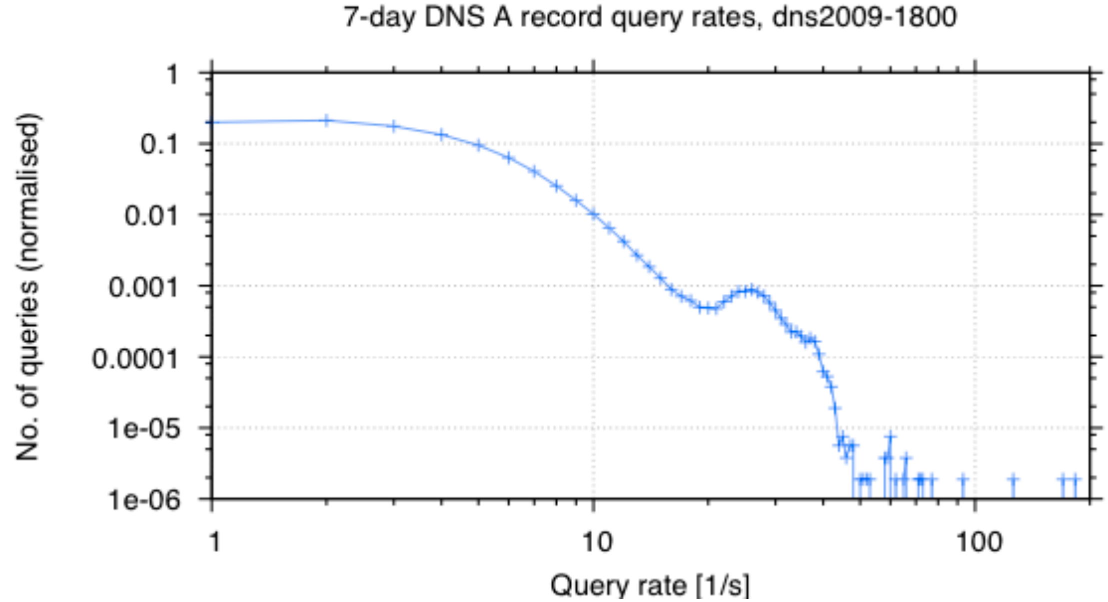

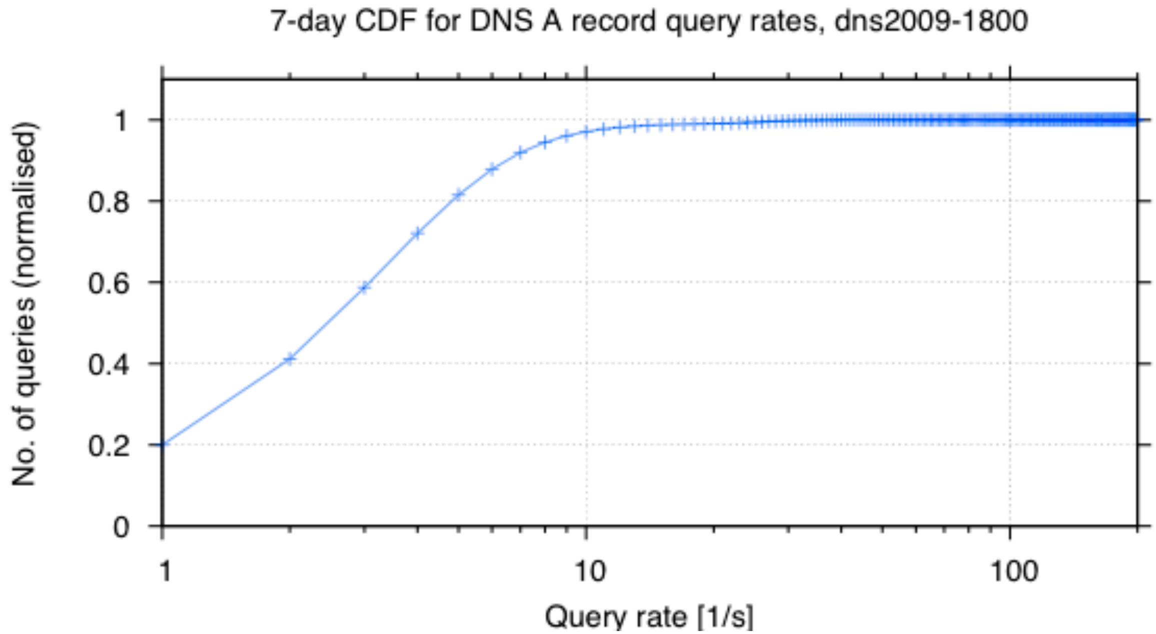

2009 dns30, dns15, dns0 periodograms

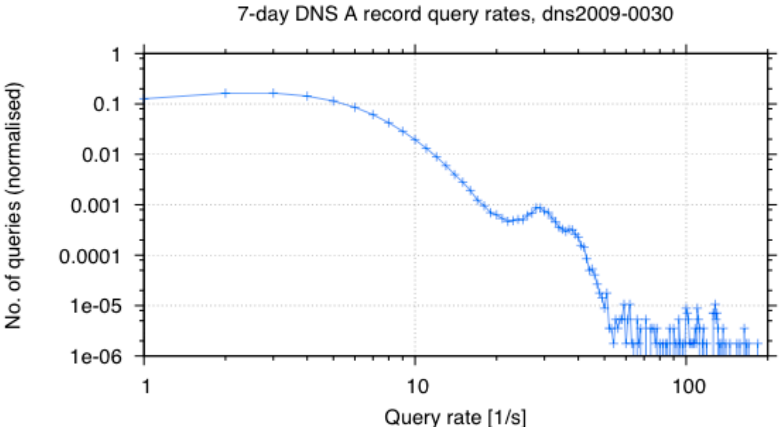

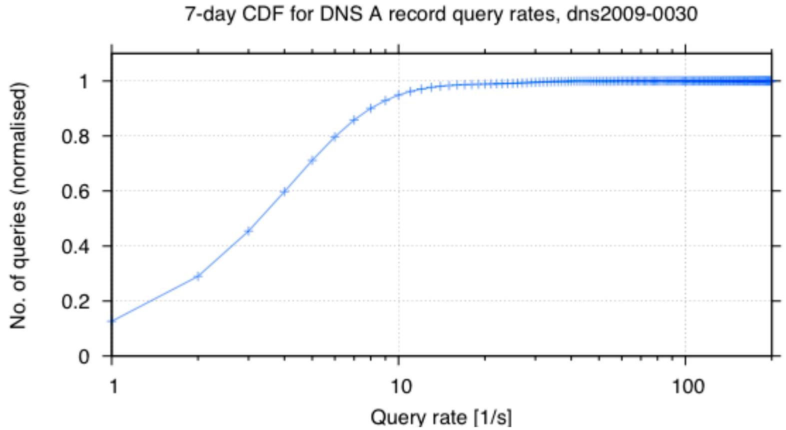

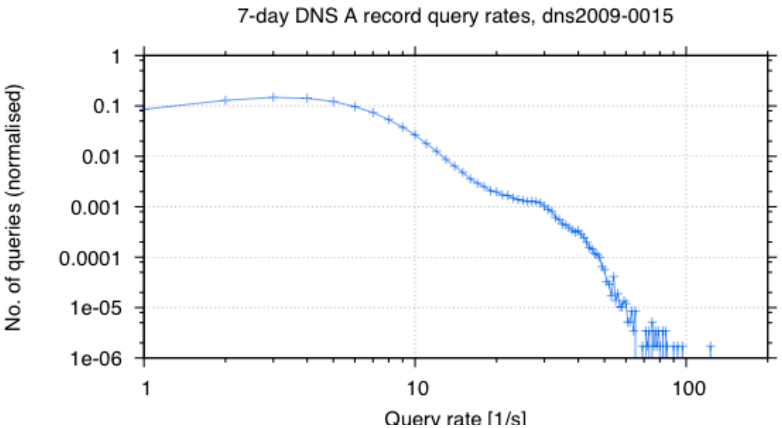

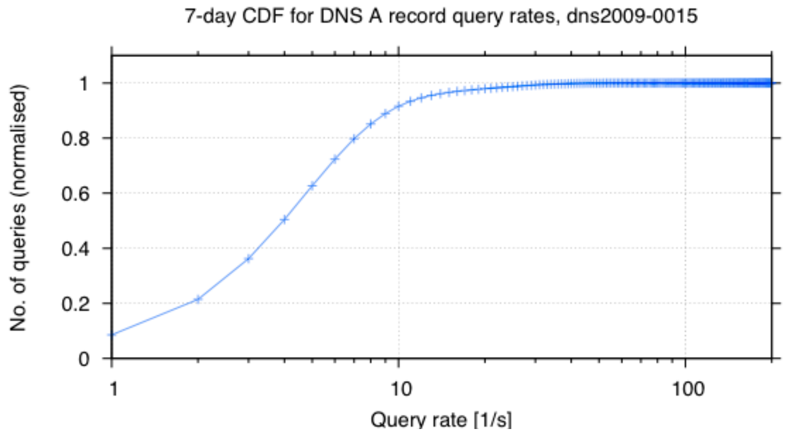

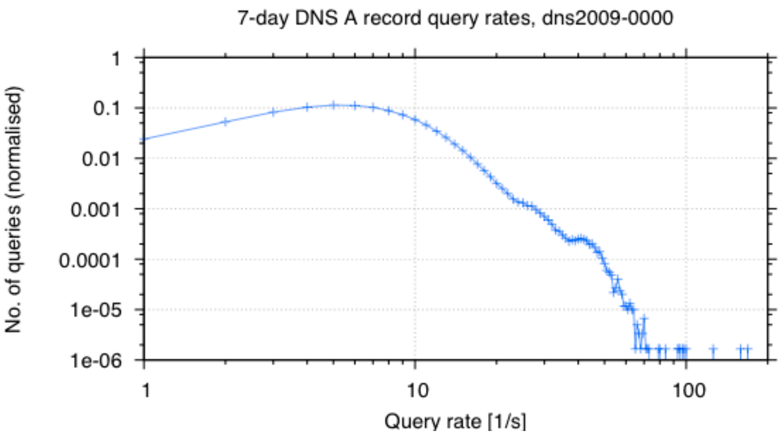

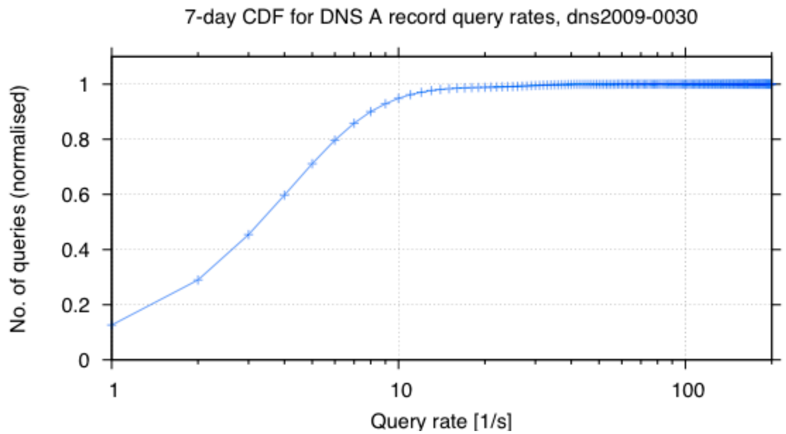

DNS support for ILNP

- We propose it is feasible to have DNS support for mobility by using zero TTL on those DNS records providing address resolution for hosts such as L64 records in ILNPv6.
- Need to evaluate impact of security mechanisms on DNS performance:
  - cryptographic authentication (client- and server-side)
  - Secure DNS Dynamic Update (RFC3007)

University of St Andrews

Thank You!

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