

LISP: Practice and Experience

*NANOG 44
Los Angeles, CA
October 2008*

LISP Team:

*Vince Fuller, Darrel Lewis, Eliot Lear, Scott Brim,
Dave Oran, Elizabeth McGee, David Meyer & Dino Farinacci*

Agenda

- LISP in a Nutshell
- Currently Deployed Network
- Deployment Model
- Numbers and Names
- Configuring LISP
- Futures
- A Few Open Questions
- Active Internet Drafts
- Q/A

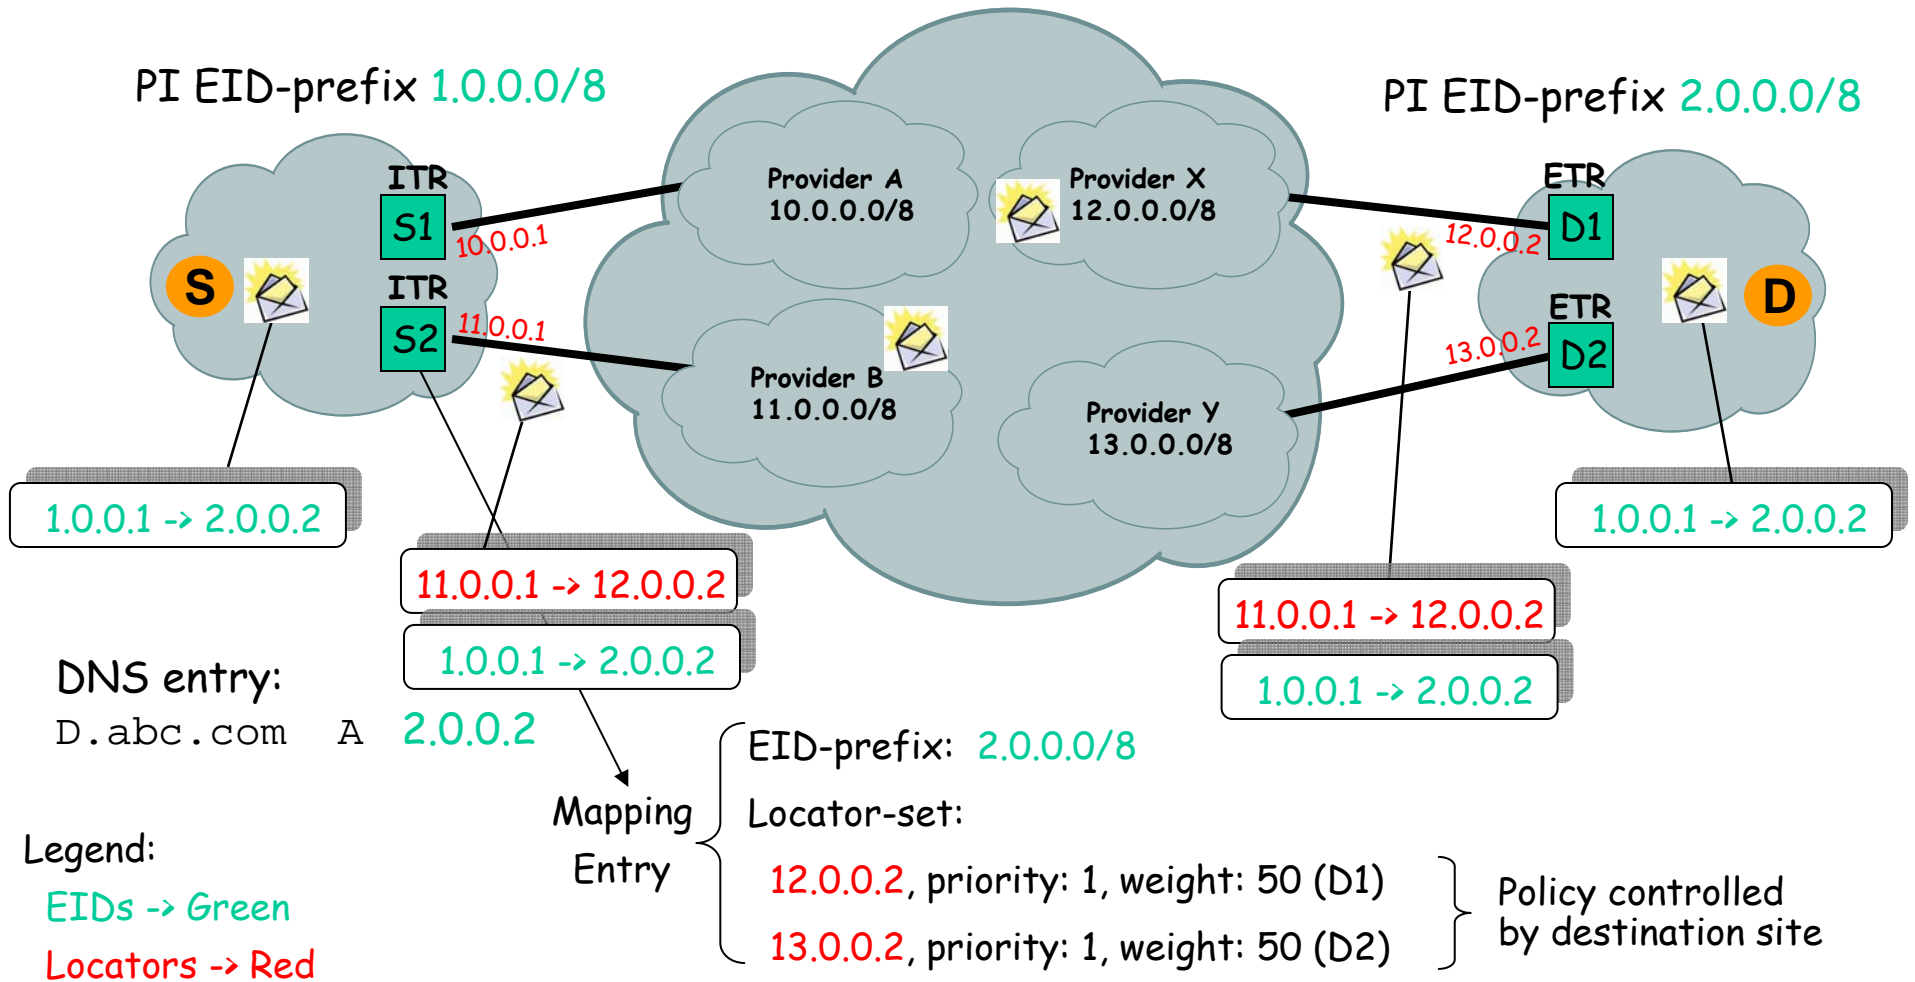
LISP in a Nutshell

- Locator/ID Separation Protocol
 - Endpoint Identifiers (EIDs) to number hosts
 - Topological Routing Locators (RLOCs) for routing
 - Network-based **map-and-encap** solution
 - No changes to hosts whatsoever
 - No new addressing changes to site devices
 - Very few configuration file changes
 - Imperative to be incrementally deployable
 - Address family agnostic
- For more, see tutorials at <http://www.lisp4.net>

New Network Elements

- Ingress Tunnel Router (ITR)
 - Finds EID to RLOC mapping
 - This is the **map** part of map-and-encap
 - Encapsulates to Locators at source site
 - This is the **encap** part of map-and-encap
- Egress Tunnel Router (ETR)
 - Authoritative for its EID to RLOC mapping
 - Decapsulates at destination site

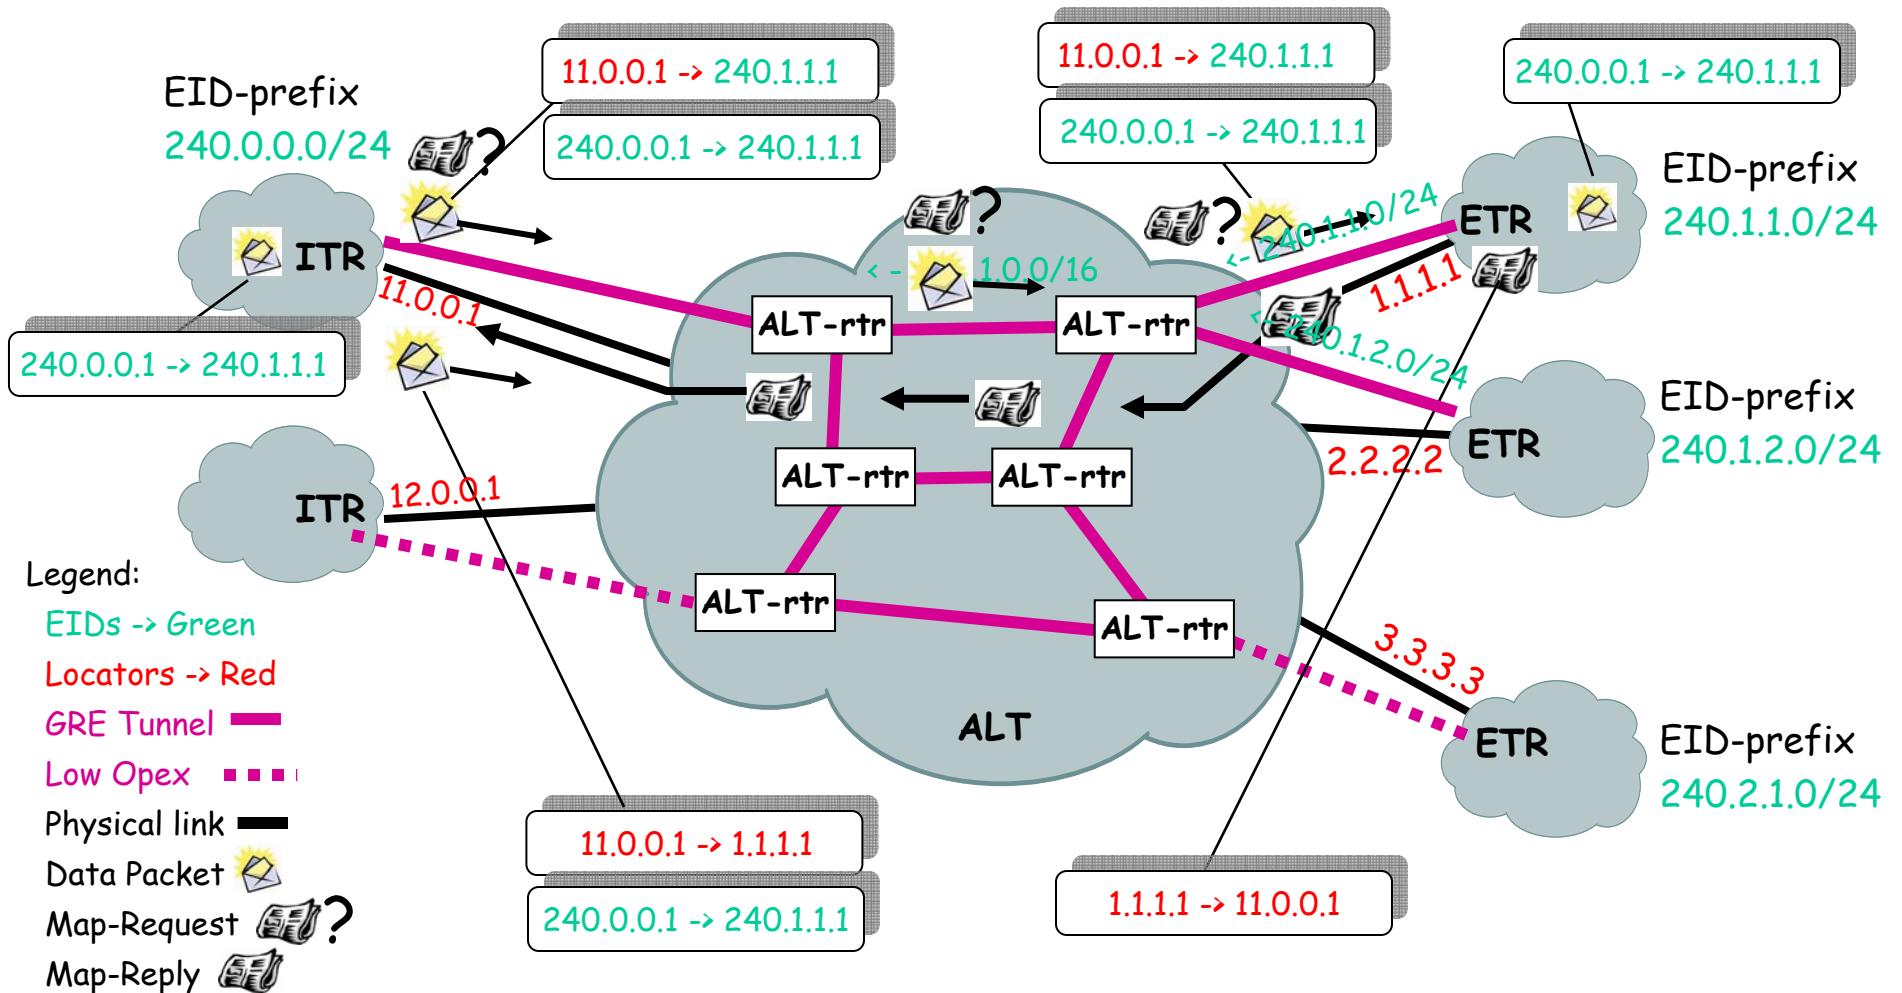
How the LISP Data Plane Works



Finding an ETR: LISP+ALT

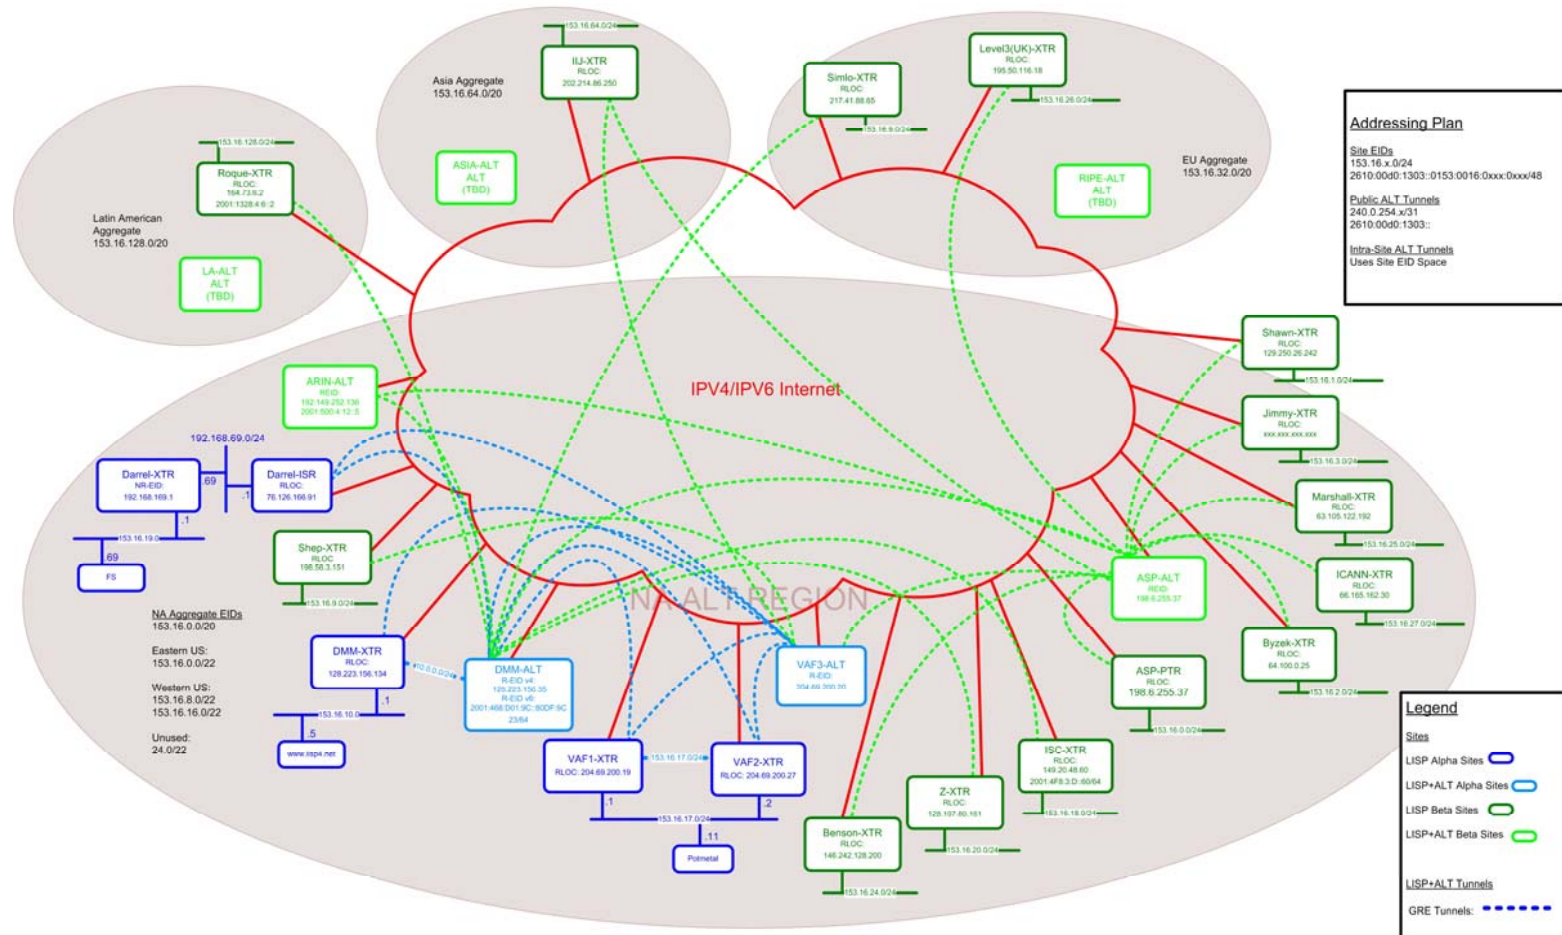
- Hybrid push/pull approach
 - ALT pushes aggregates, LISP pulls specifics
- Hierarchical EID prefix assignment
- Aggregation of EID prefixes
- GRE-based overlay network
- BGP used to advertise EIDs on overlay
- Option for data-triggered Map-Replies

How the ALT Works



What the Network Looks Like

LISP and LISP+ALT Network



Deployment Model

- Currently deployed LISP network elements are 1RU PCs ("titanium") running a LISP-capable version of NXOS
 - There are also IOS and Open Source implementations underway
- Endpoint Identifier (EID) Assignment Strategy
 - The basic idea : Geographic (probably)
 - With "ALT-Aggregators" strategically placed within a geography
- GRE tunnel topology
 - Partially meshed ALT-Aggregators, with sites arranged in a star around one or more ALT-Aggregators
 - ALT-Aggregators are typically "ALT-only"
 - Note the *ALT doesn't require GRE*

Deployment Model: Interworking

- LISP Translation
 - "LISP NAT"
- Proxy Tunnel Router (PTR)
 - Advertises coarsely aggregated EID-prefix(es) into the DFZ to attract traffic for those prefixes
 - Behaves like an ITR for that traffic

Deployment Model: Interworking

- You can also respond to a Map-Request for a v6 EID with a v4 locator (and vice versa)
- We call this "mixed locators"
- This allows you to, for example, connect sites deploying IPv6 EIDs over IPv4 locators without an intervening native IPv6 capable network
- More on Interworking in a minute

Network Numbers

- EID Prefixes
 - 153.16/16, geographically subdivided
 - i.e. 153.16.32.0/20 is EU
 - 2610:00d0::/32, sites get 2610:D0:xyz:/48
 - x is continent, y is region, zz is site
 - Note that both of these are advertised into the DFZ for interworking (PTR) purposes
- GRE tunnels numbered out of 240/4
- ALT uses 4-byte ASNs (32768.x for now)

Network Names

- lisp4.net
 - IPv4 EIDs
 - Exception:
 - www.translate.lisp4.net
 - IPv4 RLOC LISP-translated to an EID
 - More on LISP translation in a moment
- lisp6.net
 - IPv6 EIDs

ITR Configuration

- Enable ITR Functionality
 - `ip lisp itr`
 - `ipv6 lisp itr`
- Use the ALT to resolve mappings
 - `ip lisp alt-vrf lisp`
- Map-Requests vs. Data-Probes
 - `ip lisp itr send-data-probe`
 - Don't use data-probes

ETR Configuration

- Enable ETR Functionality
 - `ip lisp etr`
 - `ipv6 lisp etr`
- Configure an EID-to-RLOC database entry
 - `ip lisp database-mapping <EID-Prefix> <RLOC>`
`priority <p> weight <w>`
 - Priority tells the ETR which mappings to use first
 - Weight is a percentage of traffic (covered by **EID-Prefix**) that should be sent to **RLOC**
 - Weight can be used to implement **active-active BGP-free multihoming** (among other things)

ETR Configuration

- An ETR will typically advertise its EID-prefix into the ALT
 - Attracts Map-Requests to the authoritative ETR
- If you want "Mixed Locators"
 - ipv6 lisp database-mapping 2610:00d0:1200::/48
128.223.156.134 priority 1 weight 100
 - ipv6 lisp database-mapping 2610:00d0:1200::/48
2001:468:D01:9C:80DF:9C86 priority 2 weight 100
- And if you want the Map-Reply to come back over IPv4
 - ipv6 lisp etr send-ip-map-reply

Advertising an EID-Prefix into the ALT (pretty standard stuff)

```
...
vrf context lisp
  ip   route 153.16.10.0/24    null0 tag 1
  ipv6 route 2610:D0:1200::/48 null0 tag 1
...
router bgp 32768.1
  vrf lisp
    address-family ipv4 unicast
      redistribute static route-map static-to-bgp
    address-family ipv6 unicast
      redistribute static route-map static-to-bgp
  vrf lisp
    neighbor FC00:FFFF:FFFF:FFFF::10:0:0:2 remote-as 32768.613
    address-family ipv6 unicast
    route-map my-eid-prefixes out
  vrf lisp
    neighbor 240.0.254.135 remote-as 32768.100
    address-family ipv4 unicast
    route-map my-eid-prefixes out
```

'Low OPEX' xTR

On the Low OPEX xTR (note: BGP-free):

...

```
vrf context lisp
  ip route 153.16.0.0/16 240.0.254.140
  ipv6 route 2610:00d0::/32 2610:00d0:1fff::0240:0000:0254:0140/127
```

On the upstream ALT-Aggregator:

...

```
vrf context lisp
  ip route 153.16.19.0/24 Tunnel3 tag 613
  ipv6 route 2610:00d0:1303::/48 Tunnel3 tag 613
```

This is equivalent to **static routing** a customer

Interworking - LISP Translate

- Essentially "LISP-NAT"
- A router which is upstream from translating ETR advertises the "outside prefix" (usually part of a larger aggregate) into the DFZ, and points the prefix at the ETR doing the translation; standard NAT configuration
- ETR configuration for the translate case:
 - `ip lisp etr`
 - `ip lisp database-mapping 153.16.10.0/24 128.223.156.134 priority 1 weight 100`
 - `ip lisp translate inside 153.16.10.5 outside 128.223.157.65`
- Note that the the "inside" EID (153.16.10.5 in this case) must be covered by the EID prefix in the database-mapping command (153.16.10.0/24 in this case)
- Try <http://www.translate.lisp4.net>

Interworking - LISP PTR

- The PTR advertises the aggregated EID prefix (e.g., 153.16/16 and/or 2610:D0:/32) into the DFZ
 - This attracts traffic addressed to an EID which originates on the Internet to the PTR
- Upon receiving the traffic (addressed to an EID), the PTR functions as an ITR
 - i.e., it queries the ALT to get the EID-to-RLOC mapping and
 - LISP-encapsulates packets to the destination ETR's RLOC
- Note that the PTR doesn't have mapping state since its not really a LISP site

IPv6 LISP PTR Config

```
!  
! Use the LISP VRF for the ALT  
!  
ipv6 lisp alt-vrf lisp  
!  
! Enable the PTR  
!  
ipv6 lisp proxy-itr 2001:0468:0d01:009C::80df:9c23
```

That's really it.

Try <http://www.lisp4.net> or <http://www.lisp6.net>

Futures

- Continue to develop LISP software base
 - NXOS, IOS, OpenLISP,...
 - Recent packet format changes
 - Piggyback mappings on map-requests
 - `draft-farinacci-lisp-09.txt`
- Continue to build out the network
 - New sites: L3 (London), ARIN, UY
 - Several boxes "in-flight"
 - Let us know if you are interested...
- Simplify ALT configuration and operation

Open Questions

- Who runs the mapping system, and what is the business model?
- Complexity of the mapping system?
- Negative Map-Replies?
- Using LISP for IPv4 Address Conservation
- Effects of the mapping system on applications
 - first packet loss/lookup latency
- Scalability of the ALT
- PMTU effects
- "Stretch" effects
- Caching behavior in xTRs
- ...

LISP Internet Drafts

`draft-farinacci-lisp-09.txt`
`draft-fuller-lisp-alt-02.txt`
`draft-lewis-lisp-interworking-01.txt`
`draft-farinacci-lisp-multicast-00.txt`
`draft-meyer-lisp-eid-block-01.txt`

`draft-mathy-lisp-dht-00.txt`
`draft-iannone-openlisp-implementation-01.txt`
`draft-brim-lisp-analysis-00.txt`

`draft-meyer-lisp-cons-04.txt`
`draft-lear-lisp-nerd-04.txt`
`draft-curran-lisp-emacs-00.txt`

Questions/Comments?


Contact us: lisp-interest@lists.civil-tongue.net

Information: <http://www.lisp4.net>

<http://www.lisp6.net>

OpenLISP: <http://inl.info.ucl.ac.be/software/openlisp>

Thanks!

(IP (UDP (LISP (IP (UDP (LISP ())))))))