LISP
What Is It,
And How Much Of It Is Real?

AusNOG 02
Sydney, Australia
August 2008

David Meyer

Agenda

- A "Quick" What is LISP?
- Active Internet Drafts
  - And an observation or two…
- Deployment Model
- Numbers and Names
- What The Network Looks Like
  - And how its configured
- Acknowledgements
- Q/A

So What Is LISP?

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

Just Kidding…

What is LISP?
Problem Statement

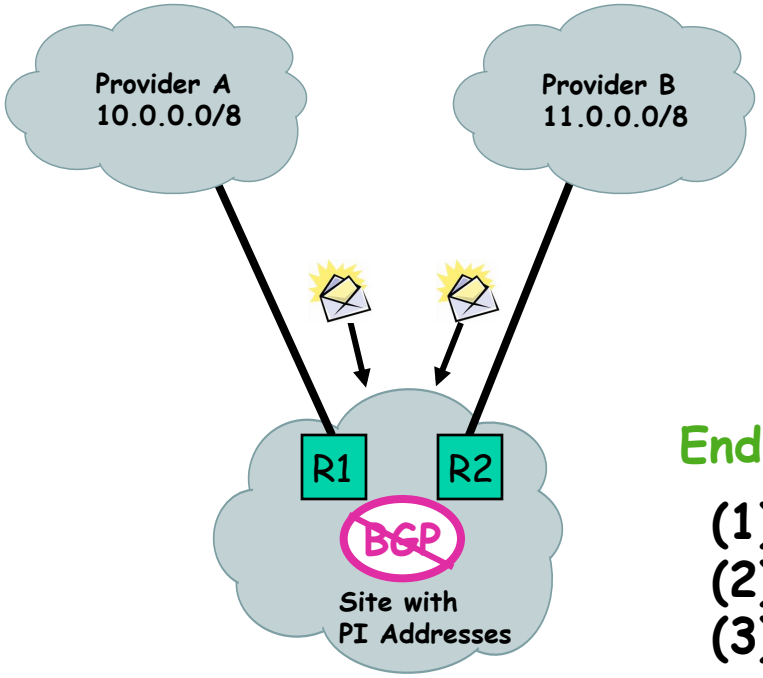

Lower OpEx for Sites and Providers

(1) Improve site multi-homing
(2) Improve provider traffic engineering
(3) Reduce size of core routing tables
(4) Reuse/optimization of PA space

End Site Benefit

(1) Easier Transition to IPv6 (if desired ☺)
(2) Change provider without address change
(3) Active-Active BGP-free multihoming

LISP Concepts

- IPv4 and IPv6 addresses have overloaded semantics
- LISP separates Location from Identity
- Introduces 2 address spaces:
  - Endpoint IDs (EIDs)
  - Routing Locators (RLOCs)
- Use 32-bit EIDs for IPv4 from registry allocation
- Use 128-bit EIDs for IPv6 from registry allocation
- Use topological addresses for Locators from ISP address block allocations
- Two types of Tunnel Routers
  - Ingress Tunnel Router (ITR) – Encaps packets at the sender
  - Egress Tunnel Router (ETR) - Decaps packets at the receiver

What is LISP?

- Locator/ID Separation Protocol
  - Map-and-Encap scheme
  - More in a sec…
- Ground rules for LISP
  - Network-based 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

What is LISP?

- Data plane
  - Design for encapsulation and tunnel router placement
  - Design for locator reachability
  - Data-triggered mapping service
- Control plane
  - Design for a scalable mapping service
  - We've deployed ALT ("Alternate Topology")
    - Documented in draft-fuller-lisp-alt-02.txt

LISP is Map-n-Encap

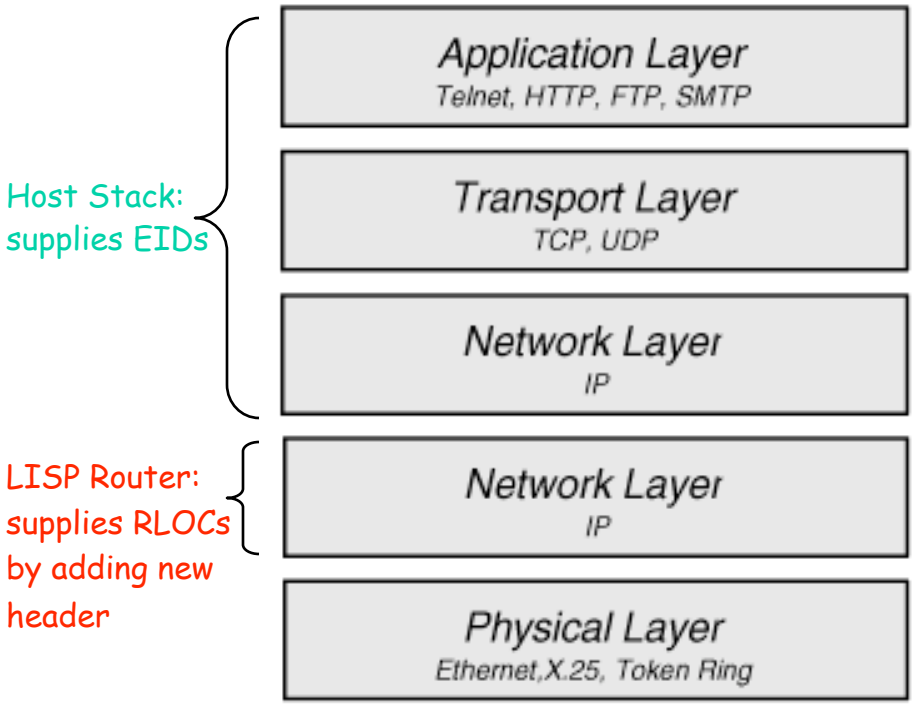

Mapping Entry:

EID-prefix: 2.0.0.0/8

Locator-set (RLOCs):

12.0.0.2, priority: 1, weight: 50

13.0.0.2, priority: 1, weight: 50

LISP Data Plane: How It Works

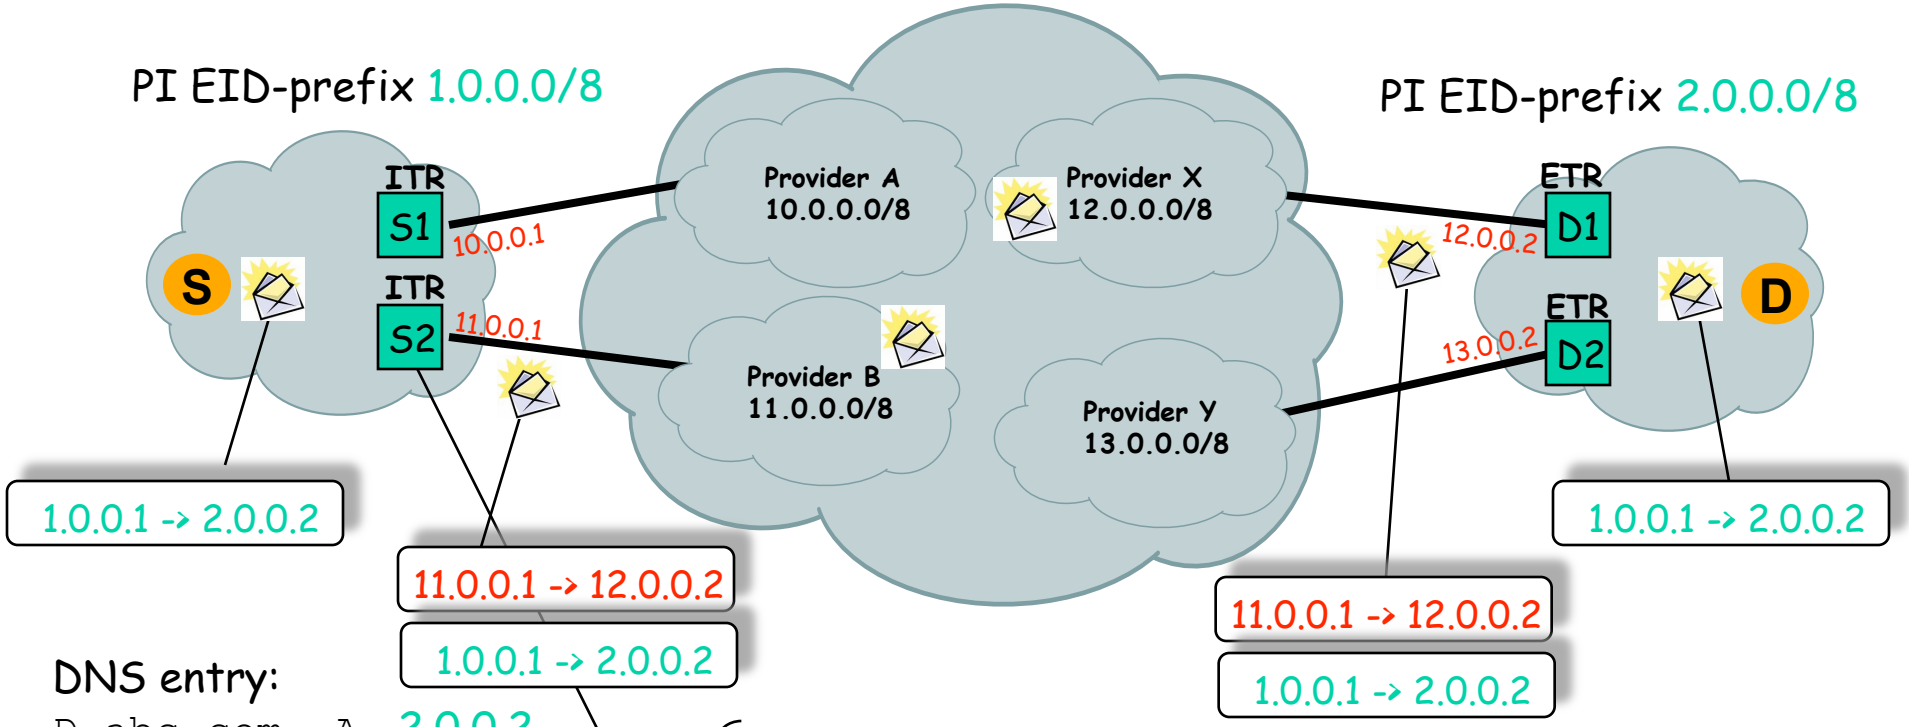

LISP Control Plane (ALT)

- The ALT is just an instance of BGP that runs in a different VRF and carries EID prefixes
  - The ALT typically runs over GRE tunnels, but we also have it running over native and .1q ethernet encapsulations
  - Typically eBGP
- ETRs typically advertise EID-prefixes into the ALT to attract Map-Requests
- ITRs use the ALT to route Map-Requests to the ETRs that are authoritative for an EID prefix

LISP Control Plane (ALT)

- ETRs return Map-Replies on the underlying network to the requesting ITR (in particular, Map-Replies *do not* flow over the ALT)
- The ITR can now LISP-encapsulate packets directly to the destination's ETR
- Its really as simple as that
  - And…a very small amount of new code was written to support this

LISP+ALT Control Plane: How It Works

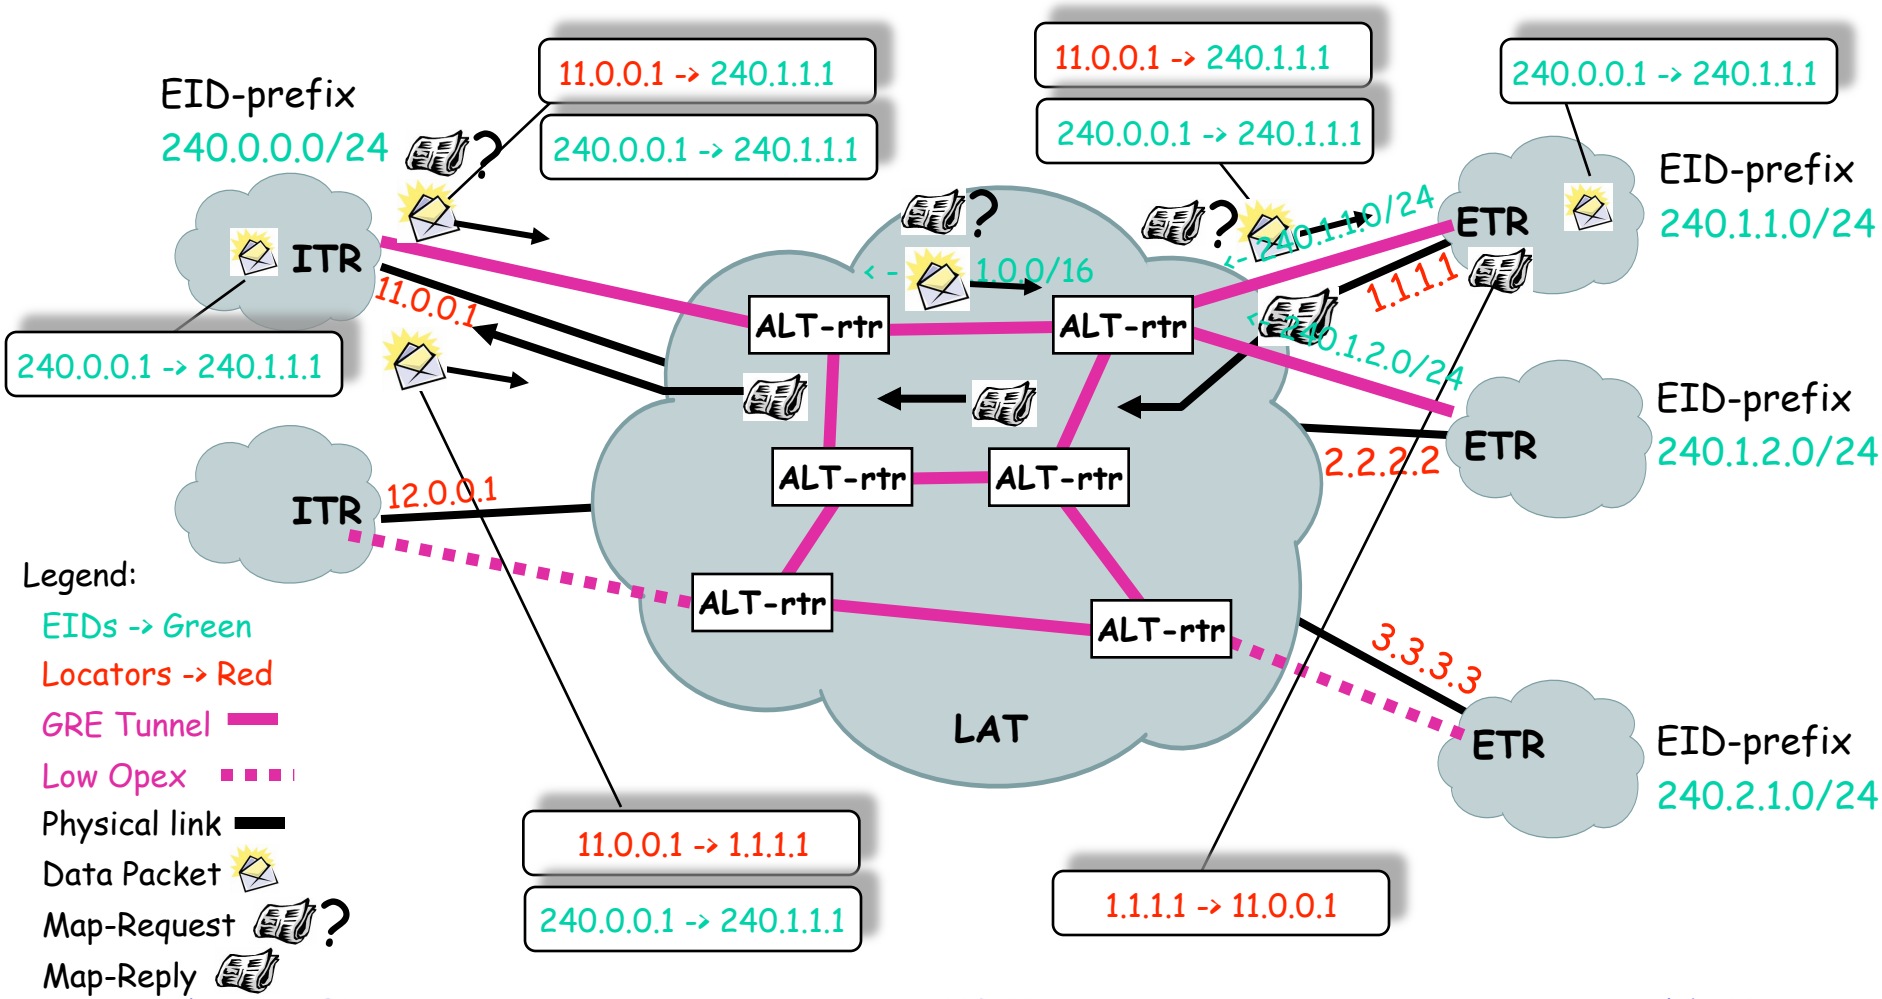

LISP Internet Drafts

draft-farinacci-lisp-08.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

An Observation

- As you'll see (and in contrast to other "experimental" networks such as the 6BONE or the MBONE):
  - The LISP network's data plane is *not* an overlay
  - The LISP network's control plane (ALT) is designed as an overlay
- So since we're not deploying an overlay, the deployed LISP network *could evolve* into the production version of the network
  - This is a significant difference from say, the 6BONE

Deployment Model

- Hardware/Software platform
  - Currently deployed LISP network elements are 1RU PCs ("titanium") running a LISP-capable version of NXOS
  - There are both an IOS and Open Source implementations underway
- 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

- We've built and deployed the interworking mechanisms described in draft-lewis-lisp-interworking-01.txt

- LISP Translation
  - "LISP NAT"
  - http://www.translate.lisp4.net
  - **`ip lisp translate inside 153.16.10.5 outside 128.223.157.65`**
- Proxy Tunnel Router (PTR)
  - Advertises coarsely aggregated EID-prefix(es) into the DFZ
    - Attracts traffic for those prefixes (i.e., Map-Requests)
  - Behaves like an ITR for that traffic
    - tr0.partan.com is a v4 PTR
    - titanium-dmm-alt-only.lisp.uoregon.edu is a v6 PTR
    - http://www.lisp6.net uses the v6 PTR
    - http://www.lisp4.net uses the v4 PTR
      - round-robins between two mirrors

- More on all of this in a few minutes

Numbers

- EID Prefixes
  - 153.16/16
  - 2610:00d0::/32
    - Note that both of these are advertised into the DFZ for interworking (PTR) purposes
- GRE tunnels numbered out of 240/4
- The ALT uses 4-byte ASNs
  - Format: 32768.X

Names

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

IPv4 EID Assignments

- **NA:** 153.16.0.0/20
  - East US: 153.16.0.0/22
  - Western US: 153.16.8.0/22
  - Western US: 153.16.16.0/22
- **EU:** 153.16.32.0/20
- **Asia:** 153.16.64.0/20
  - Japan: 153.16.64.0/21
- **Africa:** 153.16.96.0/20
- **Latin America:** 153.16.128.0/20
- **Australia:** 153.16.160.0/20
- **Reserved:** 153.16.192.0/20
  - 153.16.224.0/20

IPv6 EID Strategy

2610:D0:/32 -- The LISP IPv6 Universe

2610:D0:x000:/36

|

Continent

2610:D0:xy00:/40

|

Region

2610:D0:xy00:/48

--

Sites

IPv6 EID Assignments

- NA: 2610:D0:1000::/36
  - East US: 2610:D0:1100::/40
  - Western US: 2610:D0:1200::/40
  - Western US: 2610:D0:1300::/40
  - Infrastructure: 2610:D0:1F00::/40
  - Tunnels: 2610:D0:1FFF::/48
- EU: 2610:D0:2000::/36
- Asia: 2610:D0:3000::/36
- Africa: 2610:D0:4000::/36
- Latin America: 2610:D0:5000::/36
- Australia: 2610:D0:6000::/36
- Reserved: 2610:D0:7000::/36 -
- 2610:D0:FFFF::/36

What the Network Looks Like

LISP and LISP+ALT Network

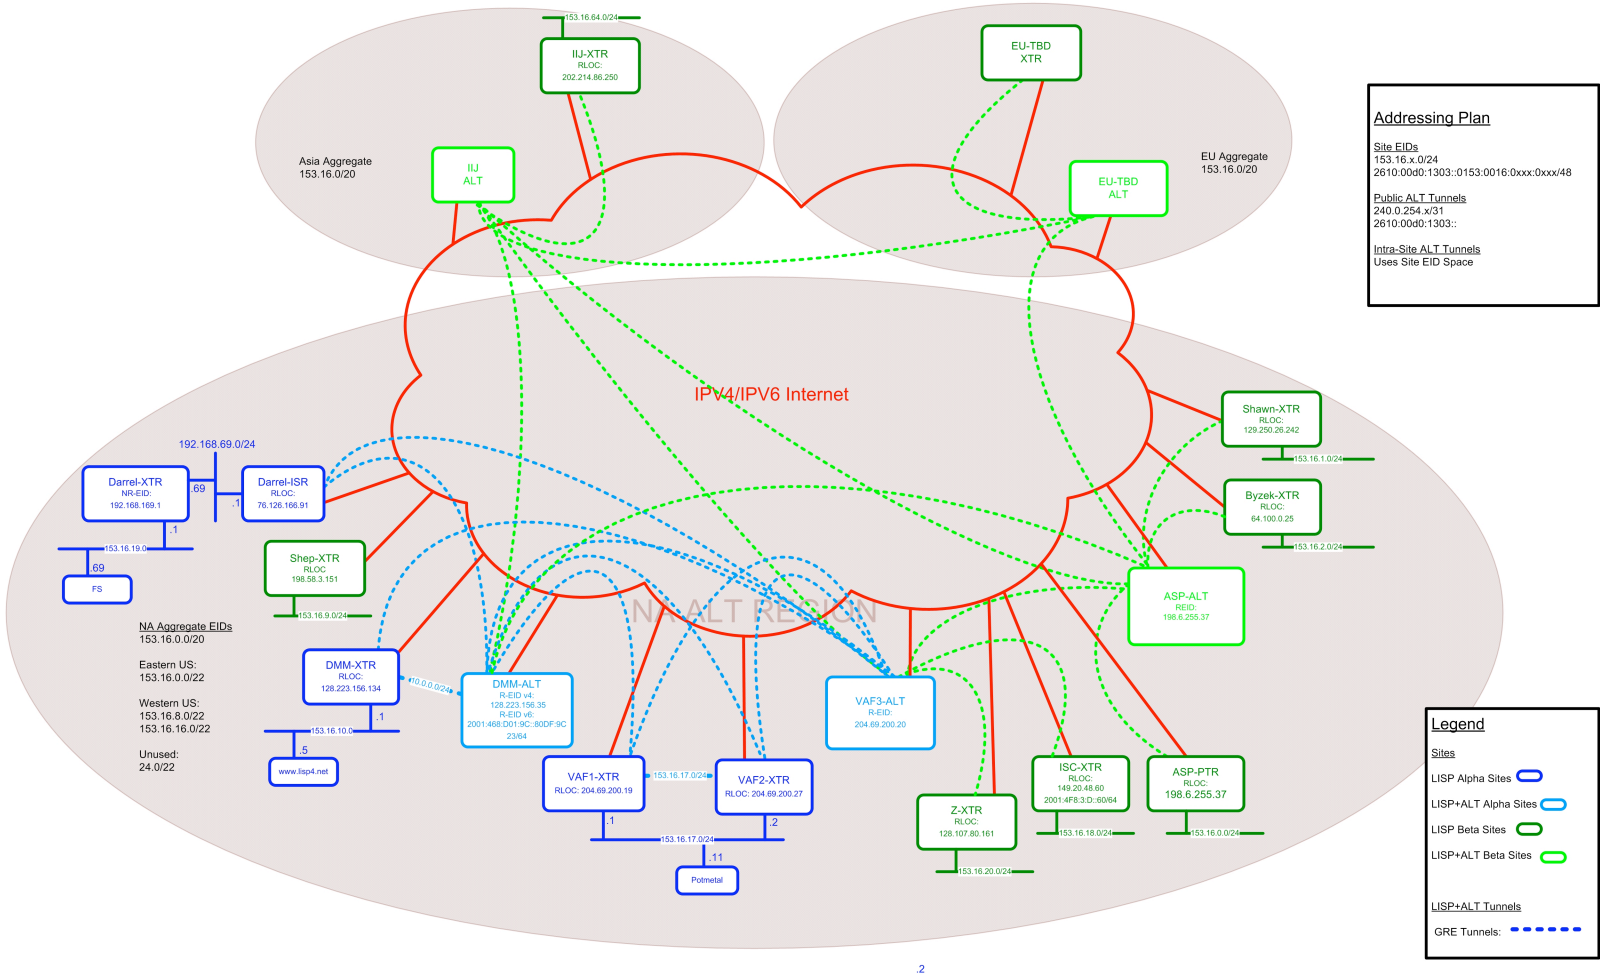

ITR Configuration

- 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 mapping
  - **`ip lisp database-mapping <EID-Prefix> <RLOC> priority <p> weight <w>`**
  - **`ip lisp database-mapping 153.16.10.0/24 128.223.156.134 priority 1 weight 100`**
  - Weight is a percentage of traffic to a given EID (covered by the EID-prefix) that should be sent to the locator
  - Can be used to implement **active-active BGP-free multihoming** (among other things)

- The ETR will also typically advertise its EID Prefix into the ALT
  - In the above example, the ETR would advertise 153.16.10.0/24 into the ALT

'Low OPEX' xTR

On the Low-OPEX xTR (no BGP):

…

```
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 ALT-Aggregator:

…

```
vrf context lisp
  ip route   153.16.8.0/22        Null0   tag 613
  ip route   153.16.19.0/24       Tunnel3 tag 613
  ipv6 route 2610:00d0:1303::/48 Tunnel3 tag 613
```

Mixed Locators

- You might want to respond to a Map-Request for a v6 EID with a v4 locator (and vice versa)
  - Allows you to connect sites deploying IPv6 EIDs over IPv4 locators (and vice versa)
  - In particular, without an intervening native IPv6 capable network
    - **Might also be used to implement a sort of NAT-PT**
  - `ipv6 lisp database-mapping 2610:00d0:1200::/48 128.223.156.134 priority 1 weight 100`
- If you want the Map-Reply to come back over IPv4
  - `ipv6 lisp etr send-ip-map-reply`

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 stuff here…

- The translating ETR is configured as follows:
  - `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)
- 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
- The PTR is configured as follows:
  - **`ip lisp alt-vrf lisp`**
  - **`ip lisp proxy-itr <rloc>`**
- Deployed PTRs
  - v4: tr0.partan.com (soon: AS 3943, more sites)
  - v6: titanium-dmm-alt-only.lisp.uoregon.edu
  - http://www.lisp4.net (round-robins between two mirrors)
  - http://www.lisp6.net

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

Debugging The First PTR

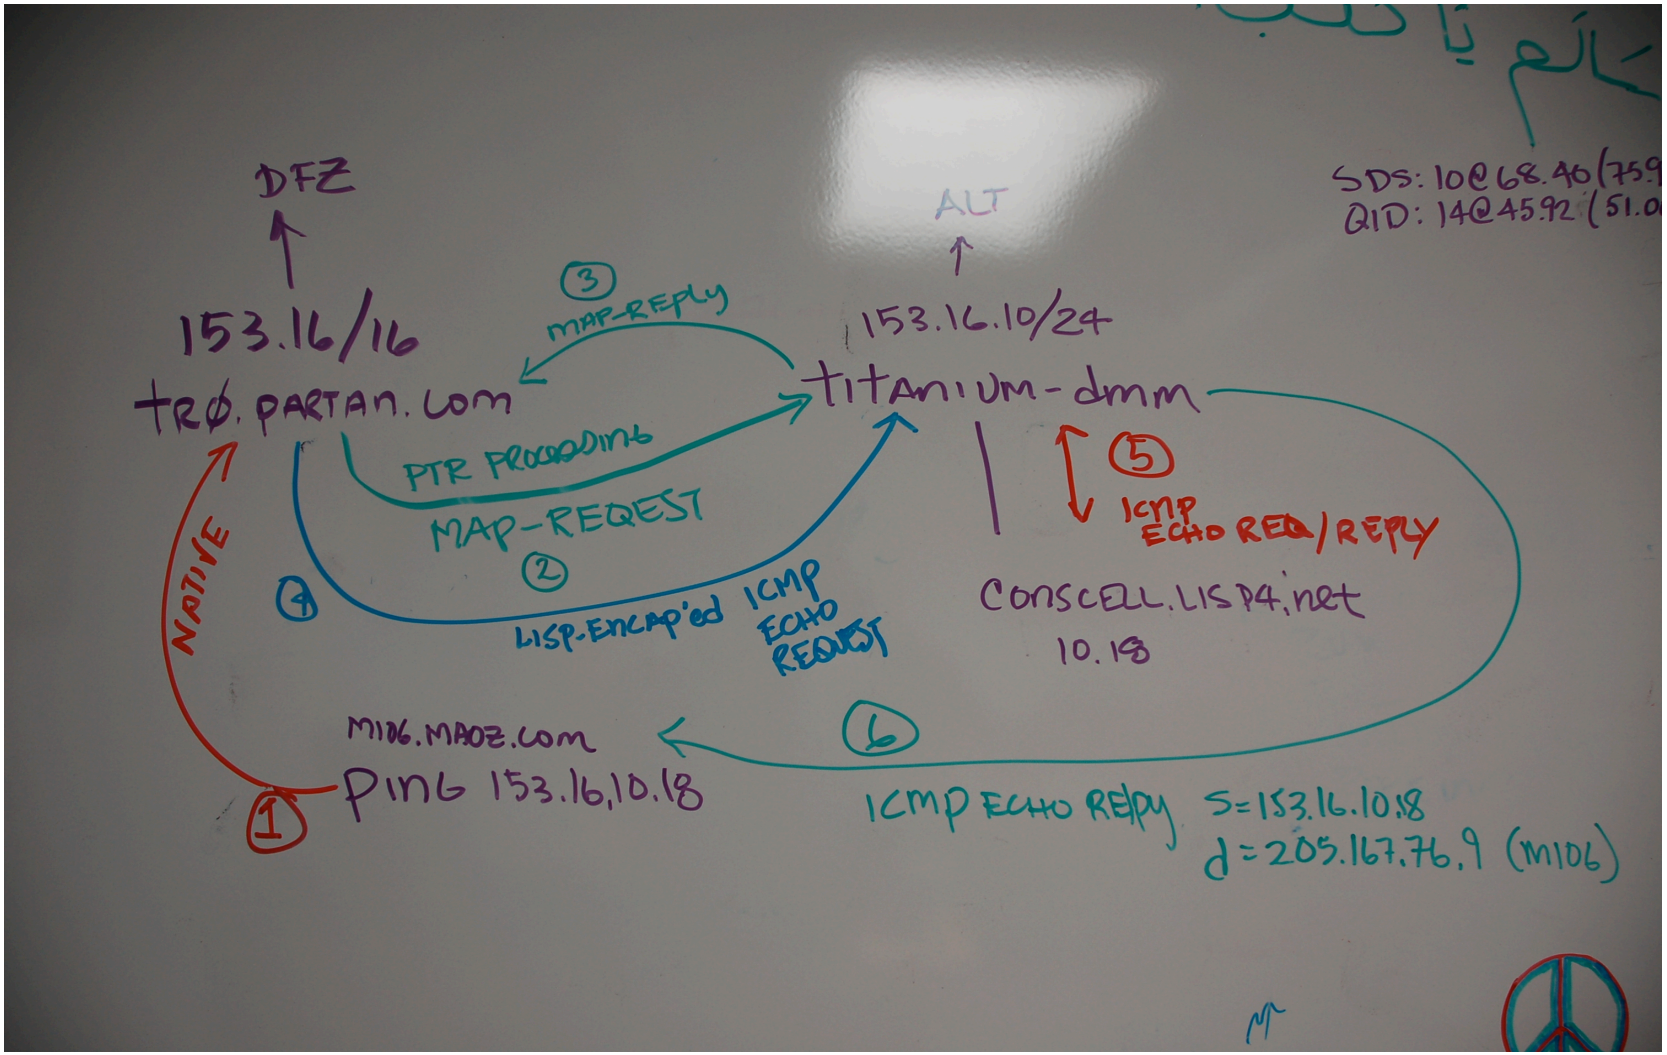

Futures

- Continue to develop LISP s/w base
  - NXOS, IOS, OpenLISP,…
- Continue to build out the network
  - Several boxes "in-flight", working on ARIN, LACNIC, RIPE/NCC, etc
    - Let me know if you are interested….
- Research
  - Topics of study include
    - Effects of the mapping system (first packet loss and/or latency) on applications
    - Scalability of the ALT
    - PMTU (additional IPv4/IPv6 encap)
    - "Stretch" effects
    - Caching behavior in xTRs
    - …

Acknowledgements

- Dino Farinacci
- Vince Fuller
- Darrel Lewis
- Scott Brim
- Eliot Lear
- Noel Chiappa
- And a cast of 1000s

Questions/Comments?

Contact us: lisp-interest@lists.civil-tongue.net
Information: http://www.lisp4.net
OpenLISP: http://inl.info.ucl.ac.be/softwares/openlisp

Thanks!