LISP: A Level of Indirection for Routing

Vince Fuller, Dave Meyer, Dave Oran, Scott Brim, Eliot Lear, Noel Chiappa, Dino Farinacci & a cast of thousands…

RIPE 55

Amsterdam

Agenda

- What is LISP?
- Why LISP?
- Protocol Details
  - Data-plane part
  - Control-plane part
- Prototype & Deployment Status
- Summary

What is LISP?

- Locator/ID Separation Protocol
  - Details to follow…
- Design Space:
  - Scalability of the routing/addressing system
- Design Assumptions/Goals:
  - Network-based solution
  - No changes to hosts whatsoever
  - No new addressing changes to site devices
  - Very little configuration file changes
  - Must be incrementally deployable
  - Address family agnostic

Locator/ID Separation?

- The idea here is that the IP address is overloaded
  - It encodes both location in the topology (locator) and the identity of the user of the address
- The locator role is used by the routing system
- The idenity role is used by upper layer protocols
  - e.g., TCP psuedo-header
- Since we want locators to aggregate topologically, and since identities are usually allocated on organizational boundaries, it is difficult to get one number space to efficiently serve both purposes
- One solution is to split the functions -- This is at the heart of the Locator/ID split idea
  - So how might we achieve this?

Implementing a Locator/ID Split

- There are two main ways to implement the Loc/ID split
- Address Rewriting (aka "Network NAT")
  - If you have enough address space (e.g., IPv6), you could use the lower 64 bits as an identifier, and the upper 64 bits as a locator, and rewrite the locator at the border
  - This is the basis of O'Dell's 8+8/GSE scheme
- Map-n-Encap
  - You could also put another header on the packet, and make the inner header carry the IDs and the outer header carry the locators
  - LISP is an instance of this approach

So What is LISP?

- LISP separates out location and identification from an existing IP address semantic

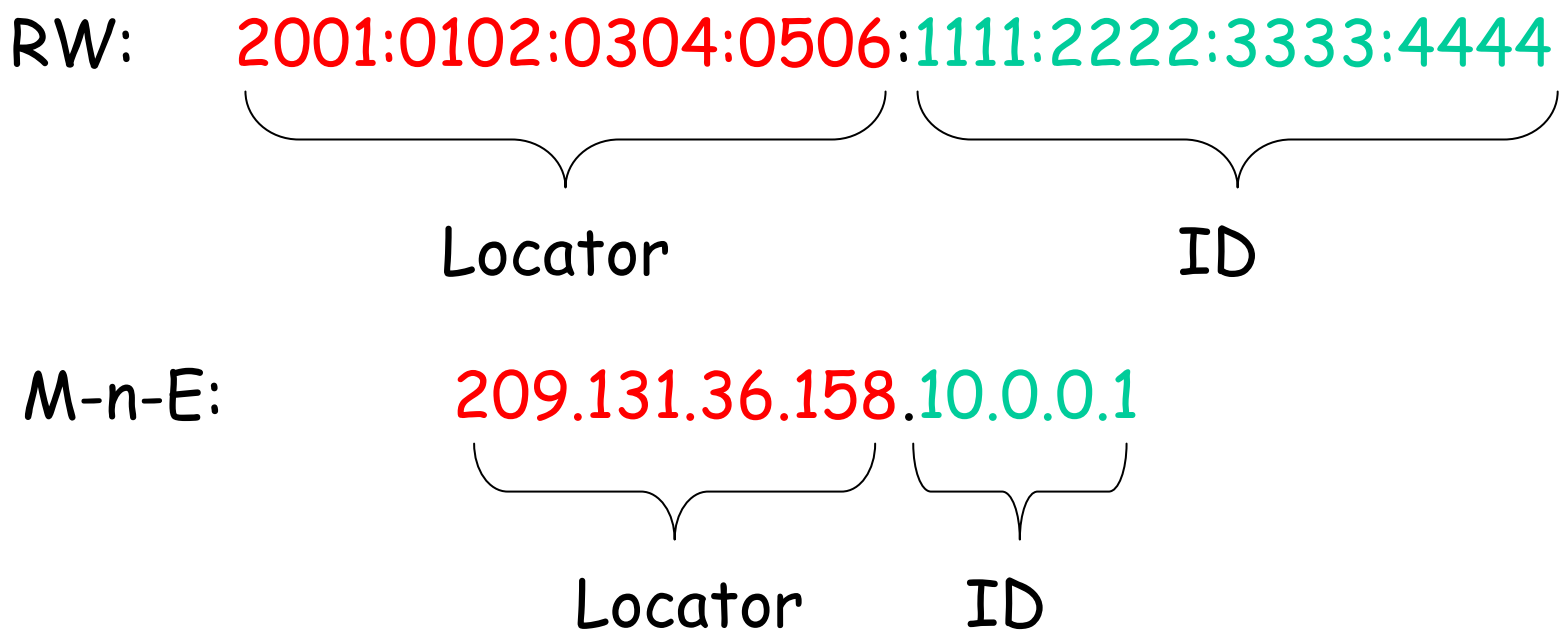

Why the Separation?

- The level of indirection allows us to:
  - Keep either ID or Location fixed while changing the other
  - Creates separate namespaces which can have different allocation properties
- In particular, the EID allocation hierarchy can follow a different topology than the RLOC allocation hierarchy
  - Which must at least be somewhat congruent to network topology if we want to be able to aggregate effectively (Rekhter's Law)

Why the Separation?

- By keeping IDs fixed
  - Assign fixed addresses that never change to hosts and routers at a site
- You can change Locators
  - Now the sites can change providers
  - Now the hosts can move

Why LISP?

- Operationally
  - Improve site multihoming
  - Improve ISP Traffic Engineering
  - Reduce site renumbering costs
  - Reduce size of core routing tables
  - Conserve IPv4 (and IPv6) address space
  - PI for all?
  - Some form of mobility?
- Architecturally
  - Create two namespaces: EIDs and Locators

What Provoked This?

- Stimulated from problem statement effort at the Amsterdam IAB Routing Workshop on October 18/19 2006
  - RFC 4984
- More info on problem statement:
  - http://www.vaf.net/~vaf/apricot-plenary.pdf

LISP Protocol Details

- 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
  - Examples are: CONS, NERD, and RPMD

Jack-Up Model

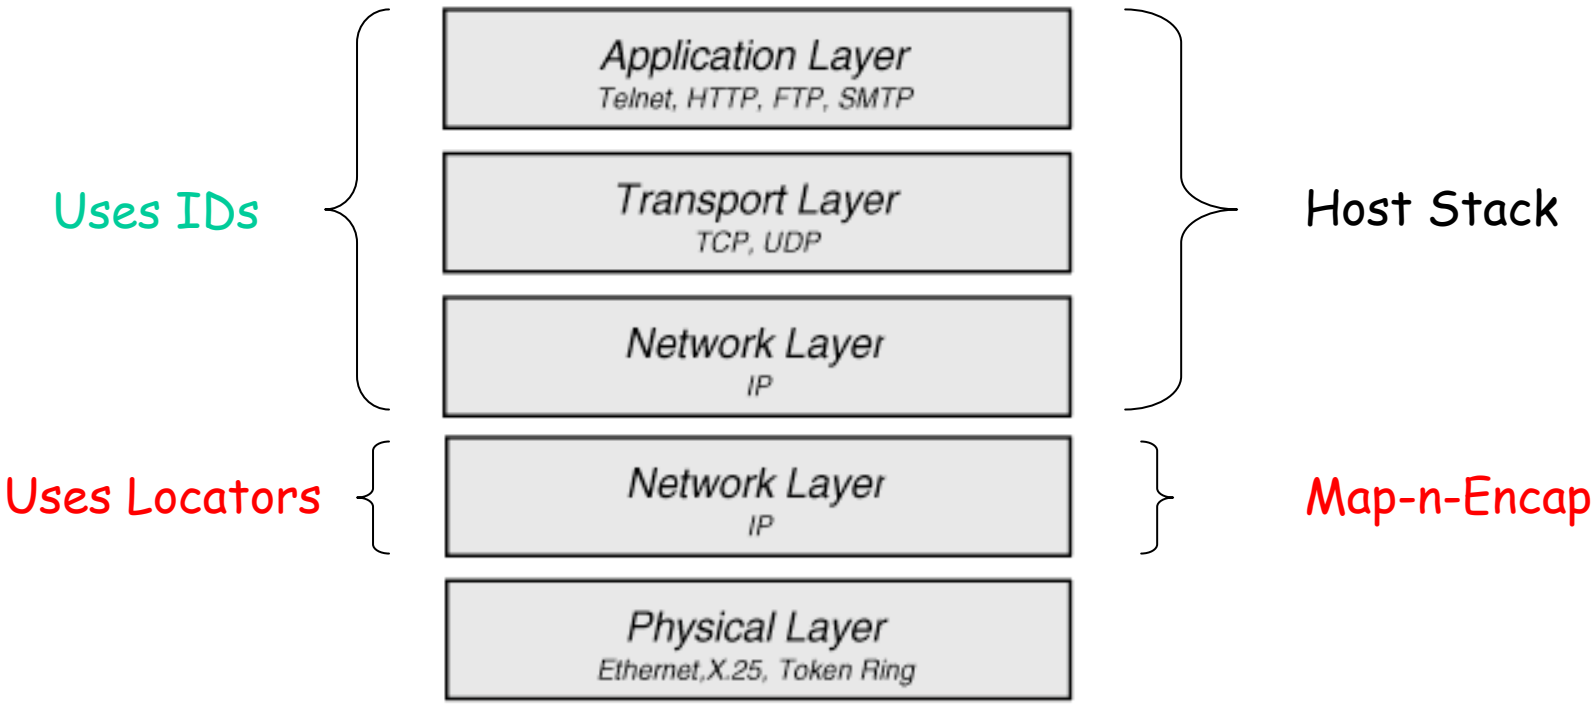

LISP is a Jack-Up Model

LISP Packet Format

LISP Protocol Details

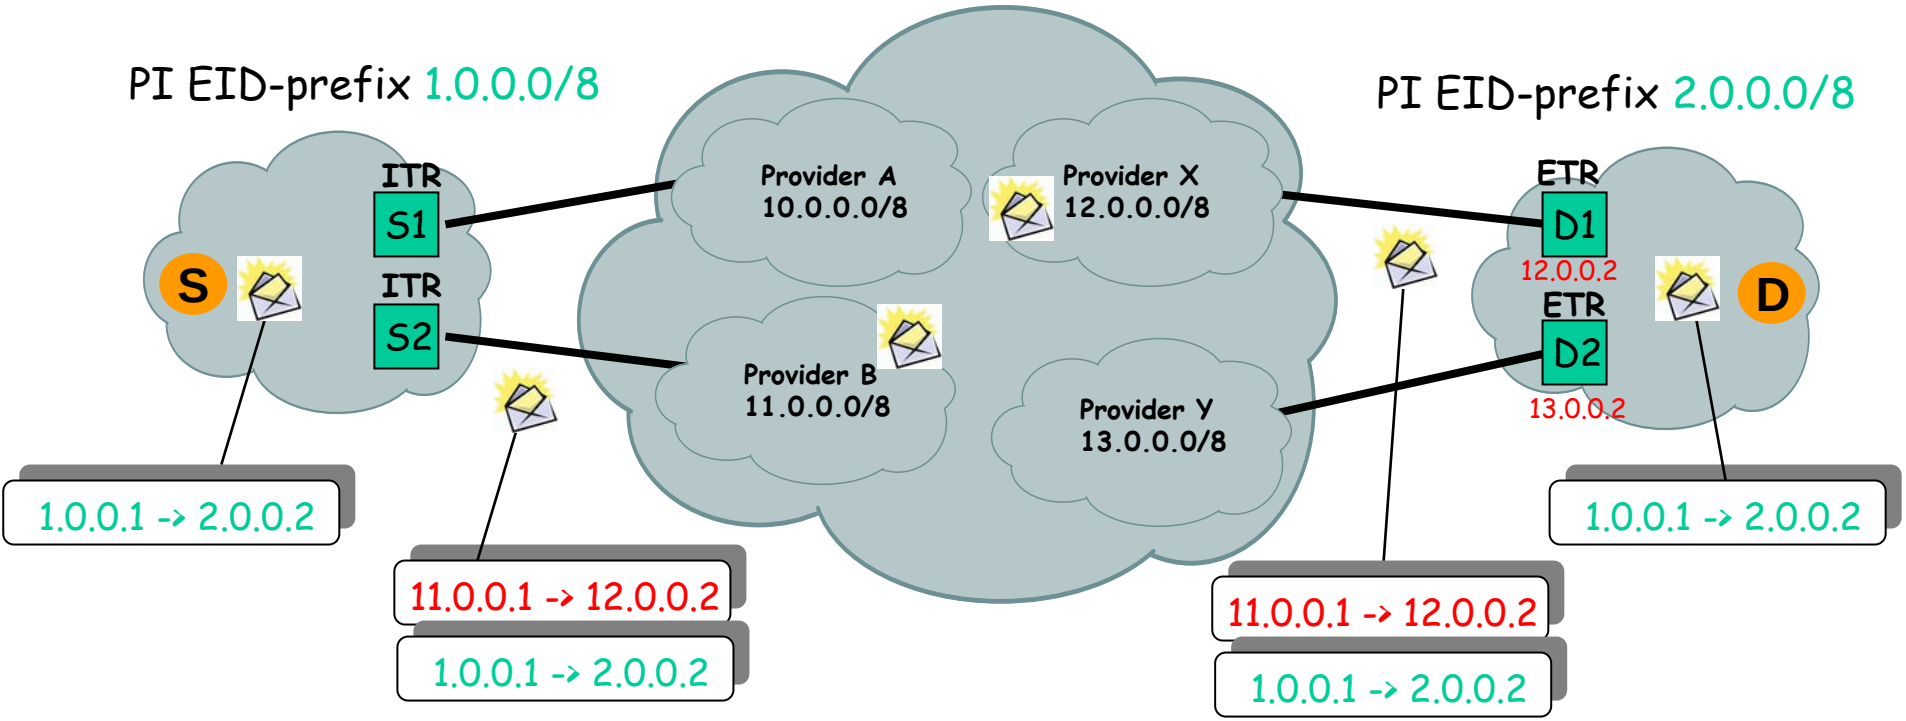

Legend:

EIDs -> Green

Locators -> Red

DNS: D -> 2.0.0.2

Mapping Entry

EID-prefix: 2.0.0.0/8

Locator-set:

12.0.0.2, priority: 1, weight: 50 (D1)

13.0.0.2, priority: 1, weight: 50 (D2)

LISP Protocol Details

- When there is no mapping in the ITR:
  - Use data-triggered UDP Map-Reply
  - Invoking by sending outer DA to inner DA
  - Send on alternative topology
  - BGP-over-GRE using EIDs as NLRI
  - No mapping data in BGP
  - No changes to BGP

LISP Protocol Details

- When there is no mapping in the ITR:
  - ITR sends CONS Map-Request
  - ITR gets Map-Reply
  - Packets get dropped in the meantime
  - Only happens first time when source site talks to destination site
  - Scalable because EID-prefix allocation not tied to underlying topology
  - Pull model

LISP Protocol Details

- When there is no mapping in the ITR:
  - Lets have mappings always in ITRs
  - NERD pushes a signed file
  - RPMD pushes signed records
  - ITRs never table-miss at expense of compressed data-set sent to every ITR
  - Push model

Prototype

- cisco has a LISP prototype implementation
  - Started the week of IETF Prague (March 2007)
- OS platform is DC-OS
  - Linux underlying OS
- Hardware platforrm is Titanium
  - 1 RU dual-core off-the-shelf PC with 7 GEs
- Based on `draft-farinacci-lisp-04.txt`
- Software switching only
- Supports both IPv4 and IPv6

Prototype

- Supports ITR encap and ETR decap
  - Load-balancing among locators
  - Respects priority & weight per mapping
- Multiple EID-prefixes per site
- Support for locator reachability
- Multi-VRF support for BGP-over-GRE
- Supports both IPv4 and IPv6

Prototype - What's Next?

- Implement "crossover" support
  - IPv6-EIDs over IPv4-Locators
  - IPv4-EIDs over IPv6-Locators
- Implement shortest-path Mobility
  - Use route-returnability check to protect ITR spoofing
- Start CONS implementation

Prototype Testing

- Dino's Unit Testing
- Meyer, Fuller, Lewis, Shepherd testing since July 2007
- External Pilot
  - Shooting for post Vancouver IETF December 2007

Internal Pilot Testing

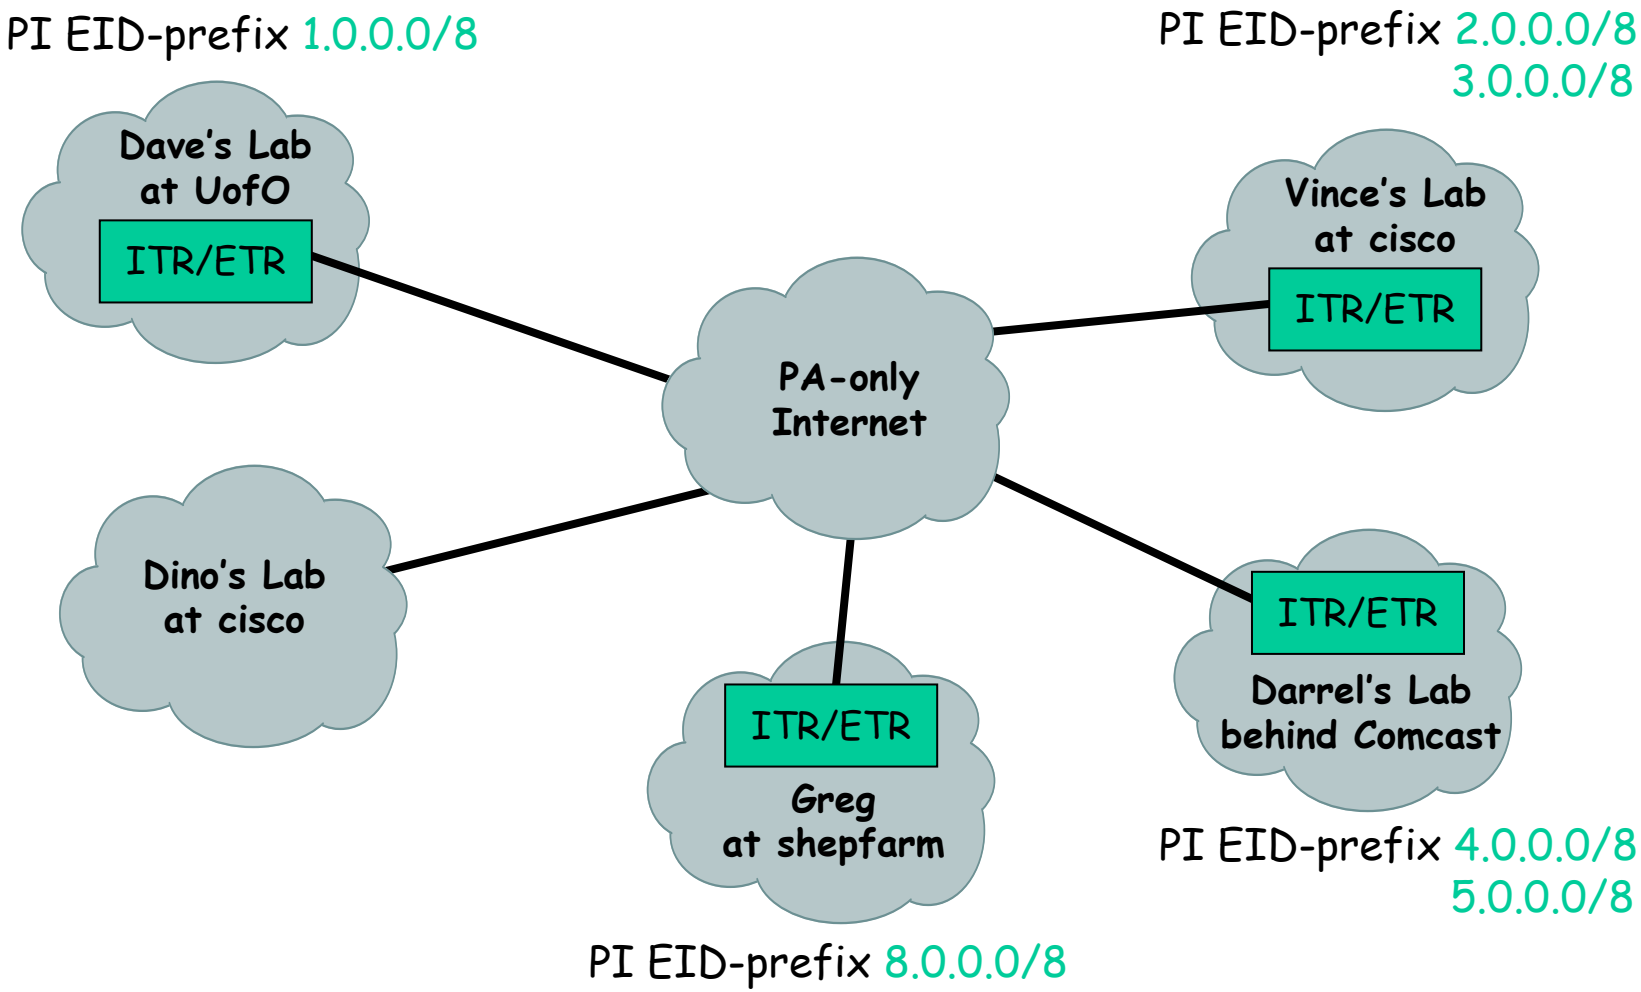

Summary

- LISP: `draft-farinacci-lisp-04.txt`
- CONS: `draft-meyer-cons-02.txt`
- NERD: `draft-lear-nerd-02.txt`
- Please send us your comments!
- Please let us know if you're interested in pilot deployment

`lispers@cisco.com`

ACKs

The authors would like to gratefully acknowledge many people who have contributed discussion and ideas to the making of this proposal. They include Jason Schiller, Lixia Zhang, Dorian Kim, Peter Schoenmaker, Darrel Lewis, Vijay Gill, Geoff Huston, David Conrad, Mark Handley, Ron Bonica, Ted Seely, Mark Townsley, Chris Morrow, Brian Weis, Dave McGrew, Peter Lothberg, Dave Thaler, Eliot Lear, Shane Amante, Ved Kafle, Olivier Bonaventure, Luigi Iannone, Robin Whittle, Brian Carpenter, Joel Halpern, Roger Jorgensen, John Zwiebel, and Ran Atkinson.

Questions/Comments?

Thanks!