



LISP: A Level of Indirection for Routing

*Vince Fuller, Dave Meyer, Darrel Lewis, Dave Oran, Scott Brim,
Eliot Lear, Noel Chiappa & Dino Farinacci*

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
- Ground rules:
 - 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?

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

IPv6: 2001:0102:0304:0506:1111:2222:3333:4444

Locator ID

IPv4: 209.131.36.158.10.0.0.1

Locator ID

Why the Separation?

- The level of indirection allows us to:
 - Keep either ID or Location fixed while changing the other
 - Create separate namespaces which can have different allocation properties

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
 - Some form of mobility?
- Architecturally
 - Create two namespaces: EIDs and Locators

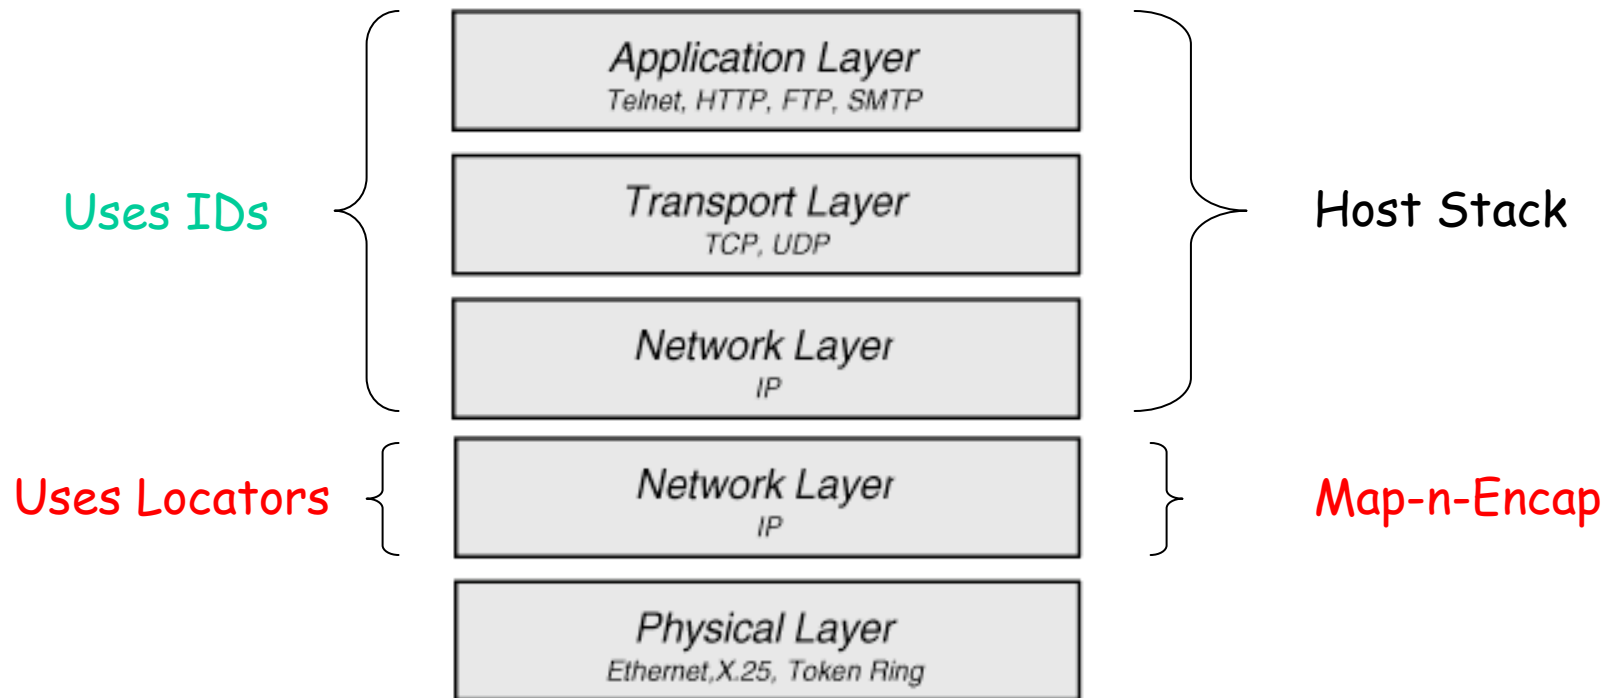
What Provoked This?

- Stimulated from problem statement effort at the Amsterdam IAB Routing Workshop on October 18/19 2006
 - <http://tools.ietf.org/html/draft-iab-raws-report-01>
- More info on problem statement:
 - <http://www.vaf.net/~vaf/apricot-plenary.pdf>

LISP Protocol Details

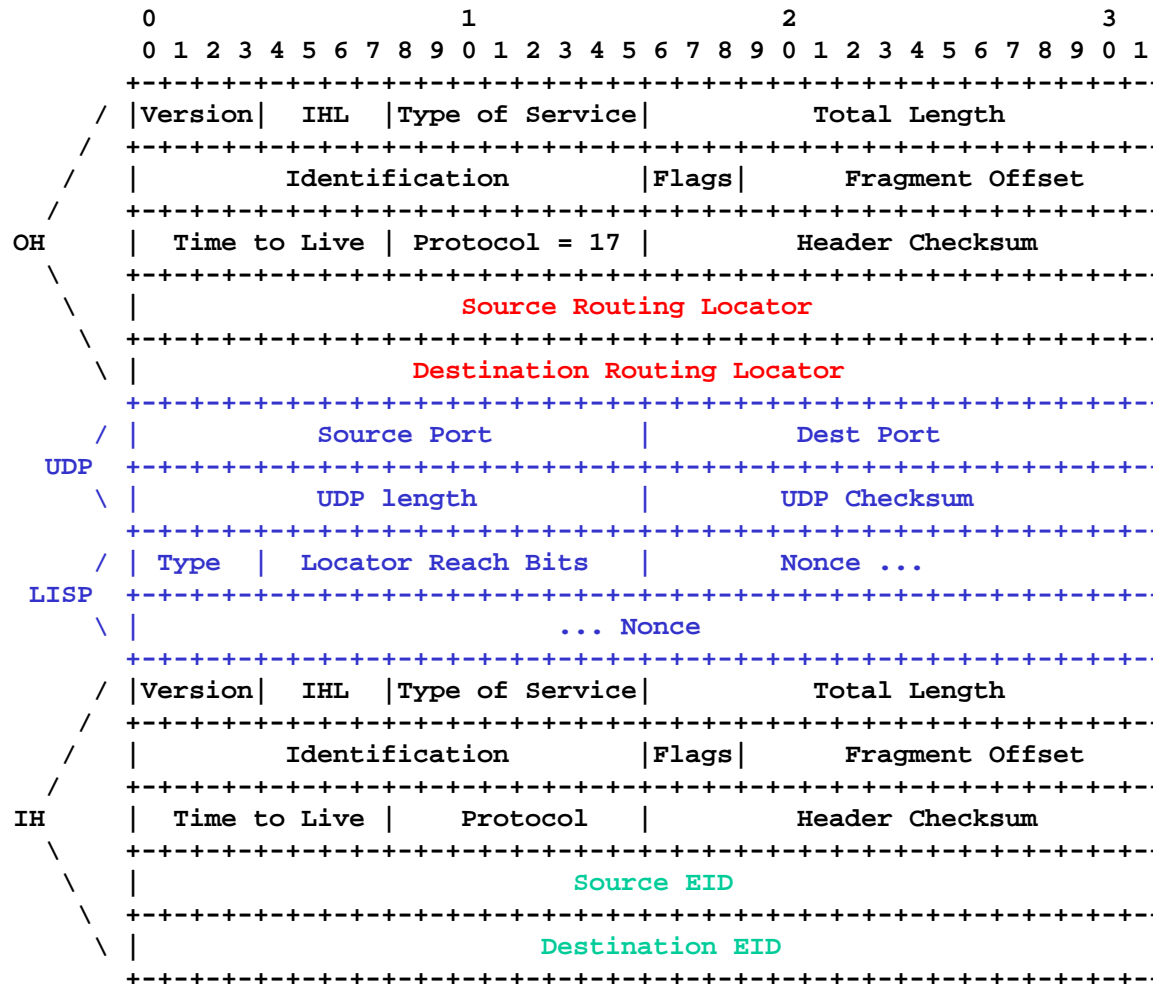
- Data plane
 - Design for encapsulation and tunnel router placement
 - Design for locator reachability
 - Data triggered mapping service
- Control plane
 - Design for a scalable mapping service
 - Examples are: CONS, NERD, and RPMD

Map-n-Encap versus NAT

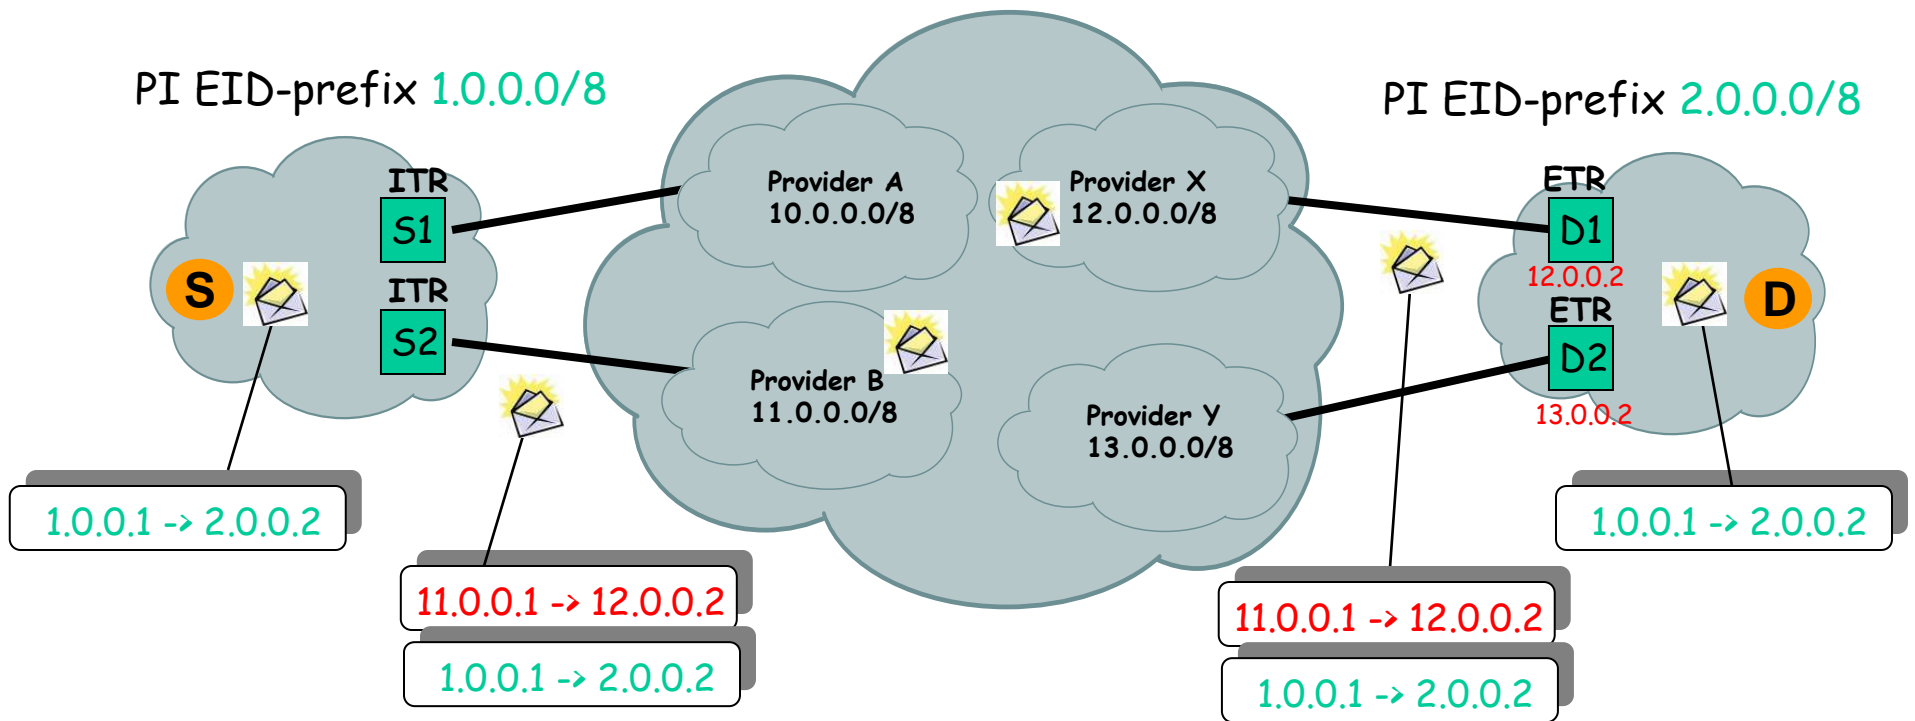


LISP is a Jack-Up Model

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
 - Invoke 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 platform 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

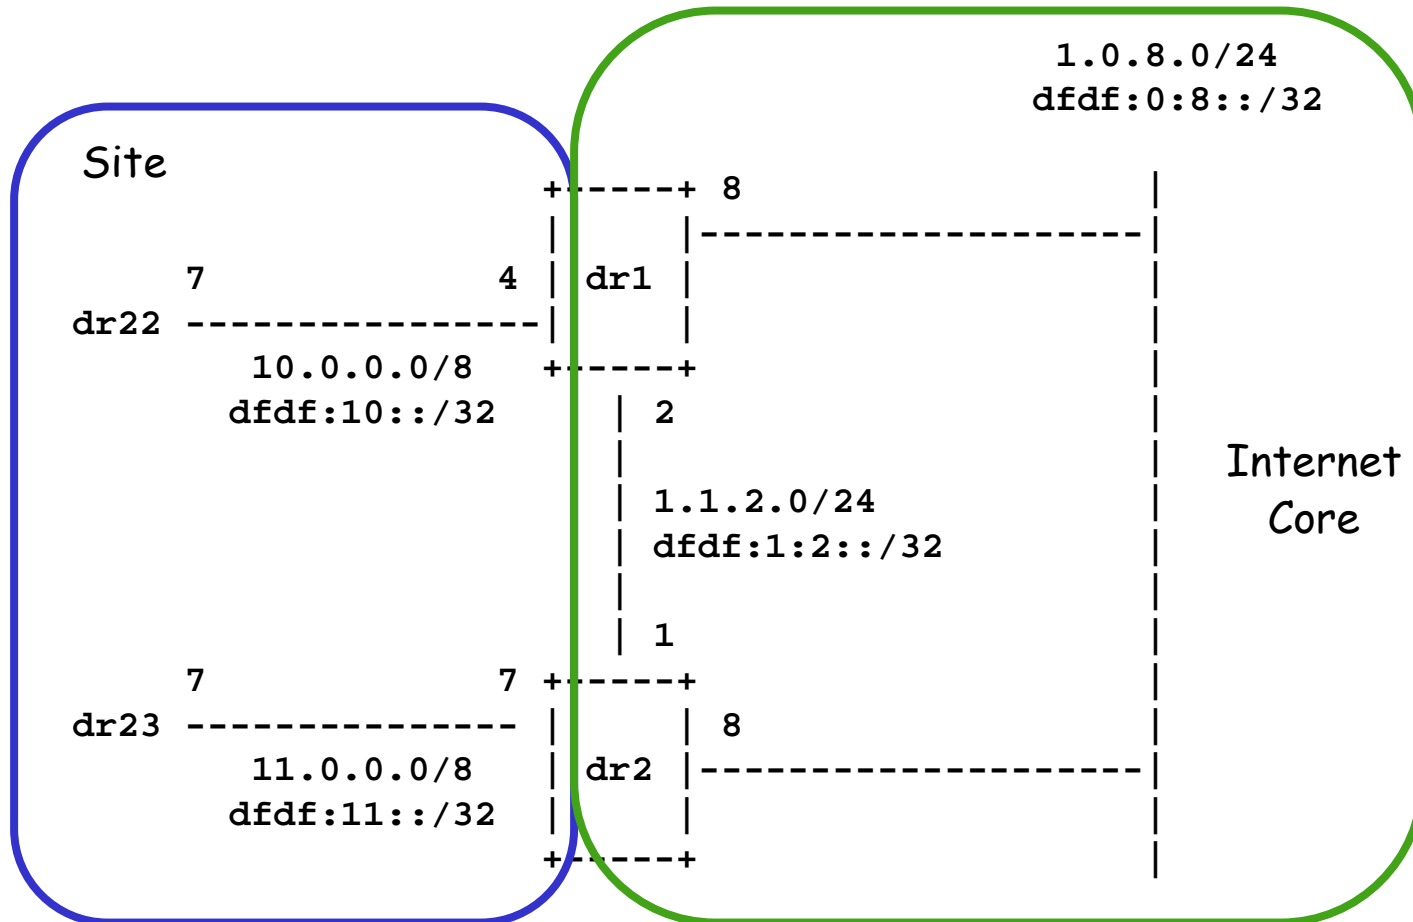
Prototype - What's Next?

- 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

- Unit Testing
 - "Dino before Dave" ;-)
 - As in "Make before break"?
- Internal Pilot
 - Meyer, Fuller, Lewis, Shepherd testing since July 2007
- External Pilot
 - Shooting for post Vancouver IETF December 2007

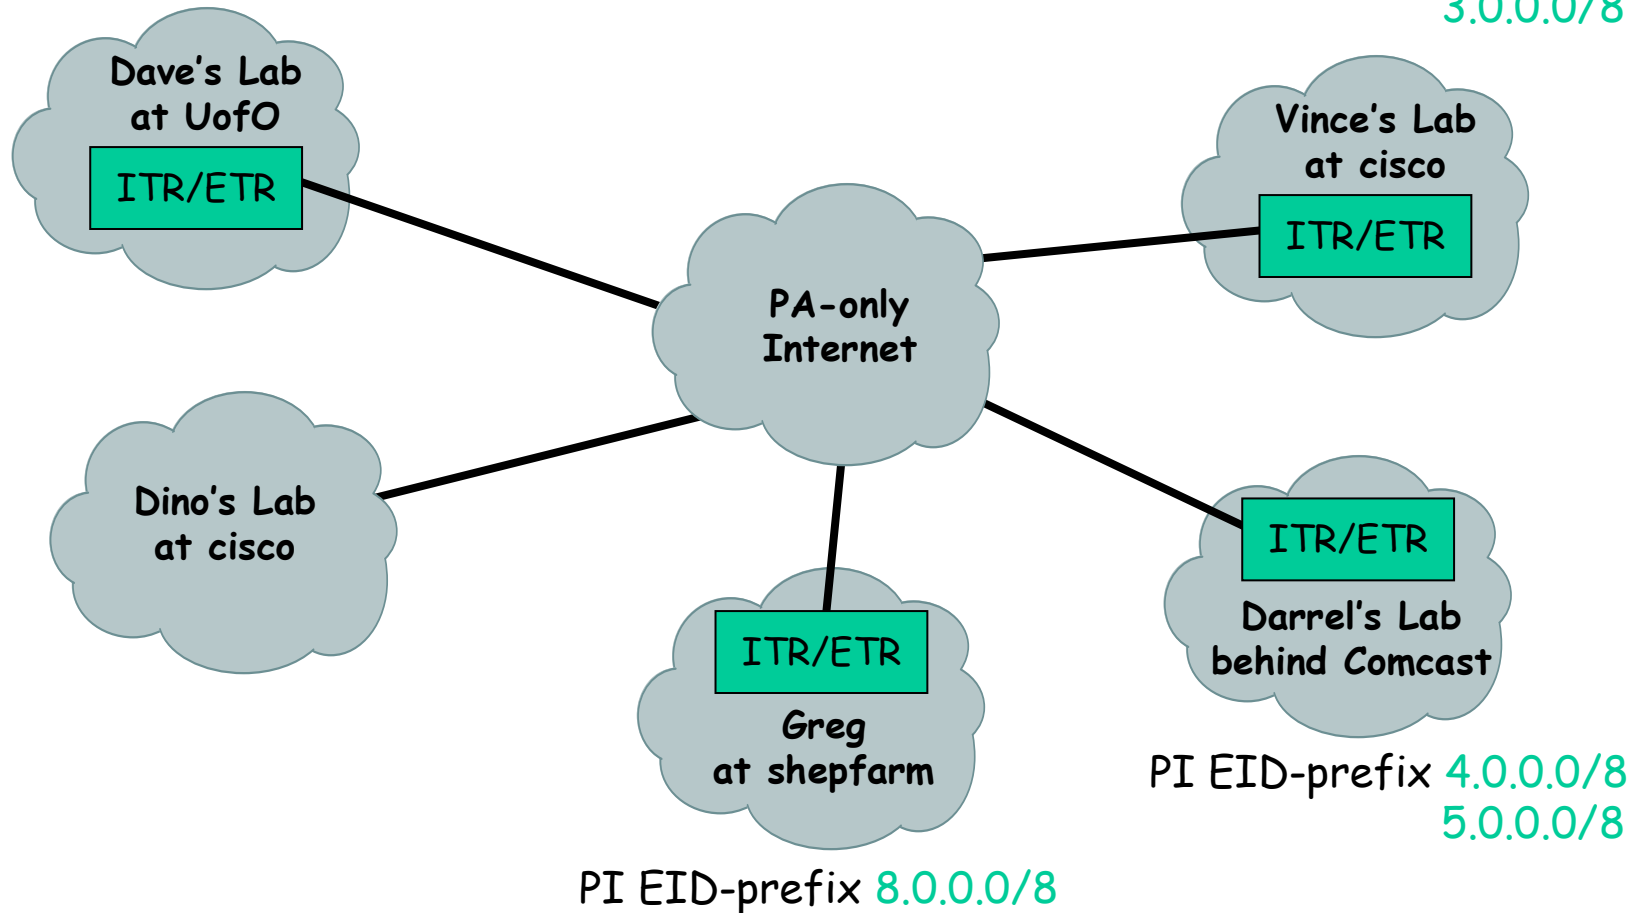
Prototype Unit Testing



Internal Pilot Testing

PI EID-prefix 1.0.0.0/8

PI EID-prefix 2.0.0.0/8
3.0.0.0/8



Summary

- **LISP:** `draft-farinacci-lisp-04.txt`
- **CONS:** `draft-meyer-cons-01.txt`
- **NERD:** `draft-lear-nerd-01.txt`
- Please send us your comments!
- Please state your interest in pilot deployment

`lisppers@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, Ran Atkinson, and Scott Brim.

```
(defun changer-one (failures database)
  (cond ((null (failures database))
        (process-one (car failures) database))
        (t (setf first-part (process-one (car failures) database)
              (setf second-part (process-two (car failures) first-part))
              (cond ((null (cdr failures)) second-part)
                    (t (changer-one (cdr failures) second-part))))))

(defun process-one (pair (car database))
  (cond ((null database) nil)
        (t (equal pair (car database))
           (execute-plan-helper (source path) (parameters) (method) (car database))
           (if (or (null source) (null path) (equal (reverse (car pair)) (car database)))
               (append (list (append (list (car database)
                                       (list 'UNKNOWN)
                                       (remove-pair (cadr pair) (cdr database))))
                       (cadr chosen-path) (cadr chosen-path))
                     ((equal (cadr chosen-path) destination) nil))
               (append (list (append (list (car database)
                                       (list 'SUCCESS)
                                       (remove-pair (cadr pair) (cdr database))))
                       (cadr chosen-path) (cadr chosen-path))
                       (equal (cadr chosen-path) destination) nil)))))

(defun process-two (pair (car database))
  (cond ((null database) nil)
        (t (or (cbr-super source destination plan)
                (equal (reverse (cadr pair)) (car database))
                (cbr-top (cadr chosen-path) destination)
                (append path (cbr-get-path chosen-path plan))
                (cbr-top (car chosen-path) destination)
                (append path (cbr-get-path chosen-path plan))
                (cbr-top source destination (cbr-new-plan
                                           (cbr-evaluate source destination
                                           (cbr-retrieve source destination plan))))))

(defun convert-output-helper (pair)
  (cond ((null path) nil)
        (t (append (list (cons (first path)
                              (remove-pair (cadr pair) (cdr database))))
                  (convert-output-helper (rest path))))))

(defvar *database*
  '( ((10th Tech-Parkway) (10th Curran) (Tech-Parkway North-Avenue))
    ((10th Curran) (10th Hemphill) (Curran 8th-1))
    ((10th McMillan) (10th Curran) (10th Hemphill) (mcmillan 8th-1))
    ((10th Hemphill) (10th McMillan) (10th Dalnev) (hemphill 8th-1))
  )


```