

QuickTime™ and a
decompressor
are needed to see this picture.

LISP: An Architectural Solution to Multi-homing, Traffic Engineering, and Internet Route Scaling

Dave Meyer & Dino Farinacci

LISP Designers:

*Dave Meyer, Vince Fuller, Darrel Lewis, Andrew Partan,
John Zwiebel, Scott Brim, Noel Chiappa & Dino Farinacci*

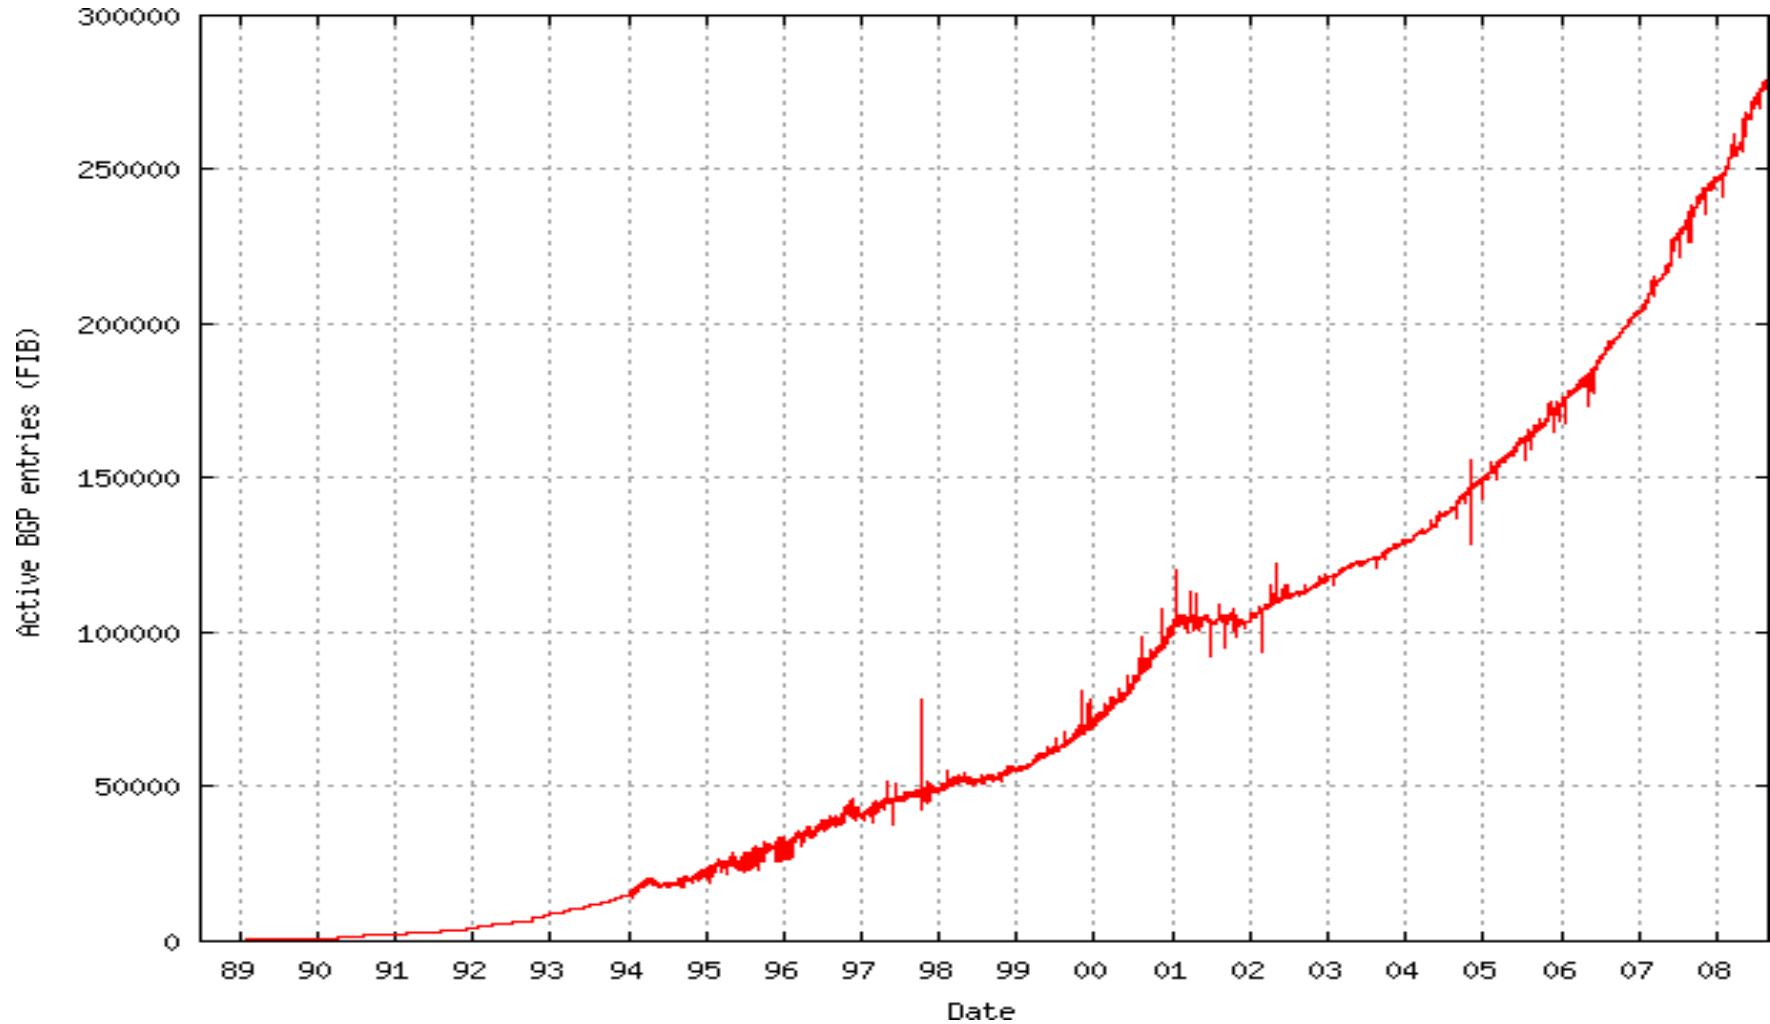
Agenda

- Problem Statement
- Architectural Concepts
 - Separating Identity and Location
 - Map-n-Encap to Implement Level of Indirection
- Unicast & Multicast Data Plane Operation
- Mapping Database Mechanism
 - LISP-ALT
- Locator Reachability
- Interworking LISP Sites and Legacy Sites
- Deployment Status
 - Present Prototype Implementation
 - Present World-Wide LISP Pilot Network
- Q&A

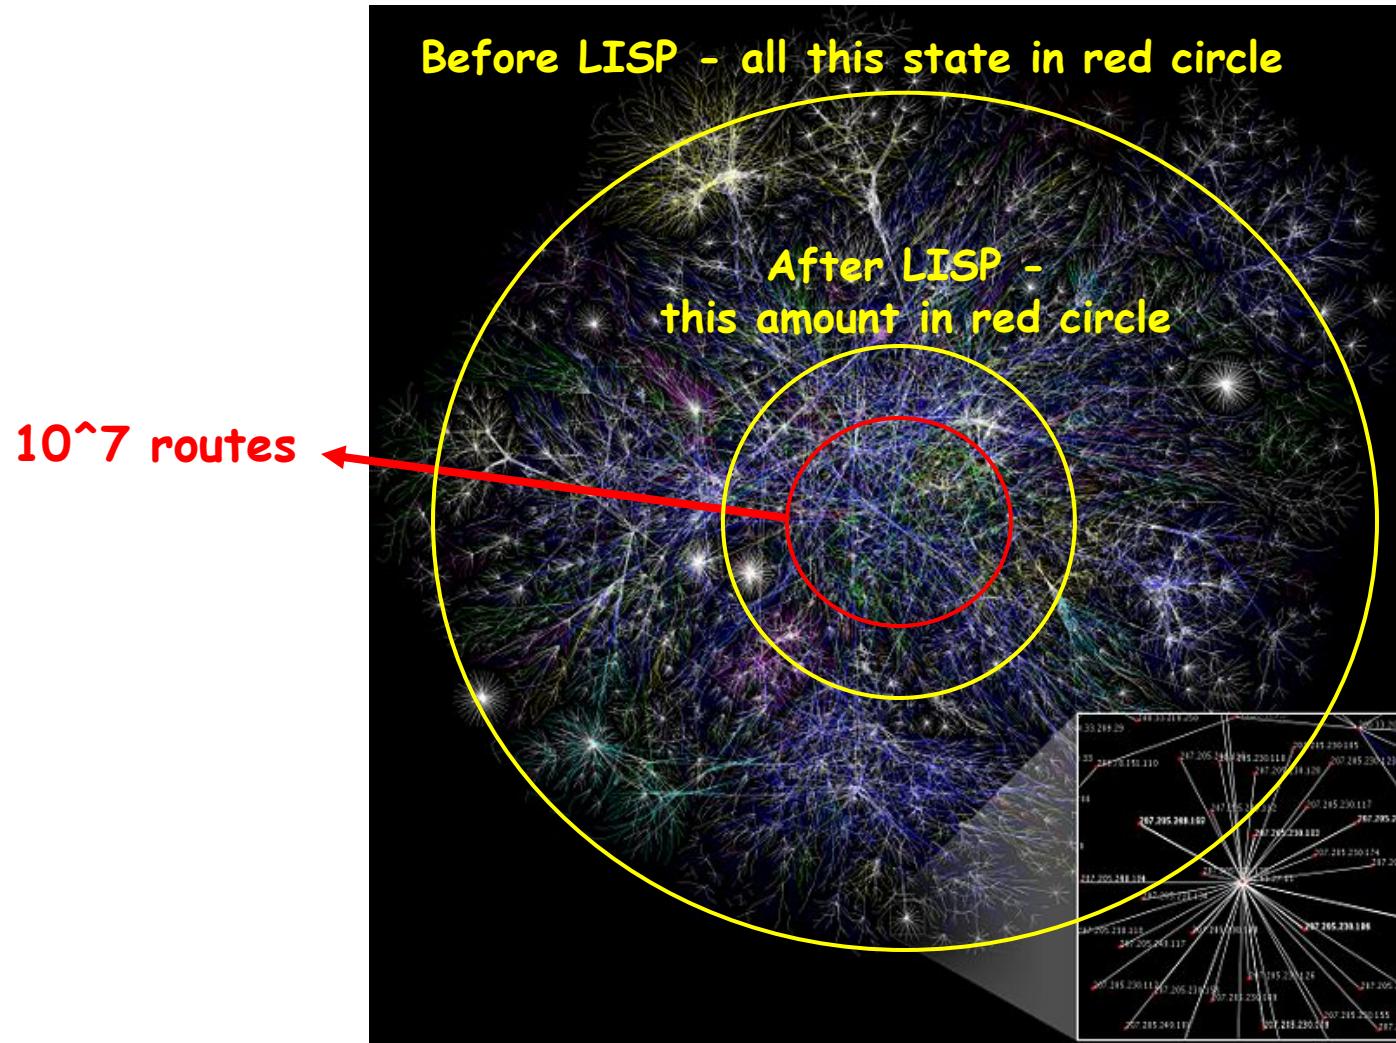
Problem Statement

- What provoked this?
 - Stimulated by problem statement effort at the Amsterdam IAB Routing Workshop on October 2006
 - RFC 4984
 - More info on problem statement:
 - <http://www.vaf.net/~vaf/apricot-plenary.pdf>

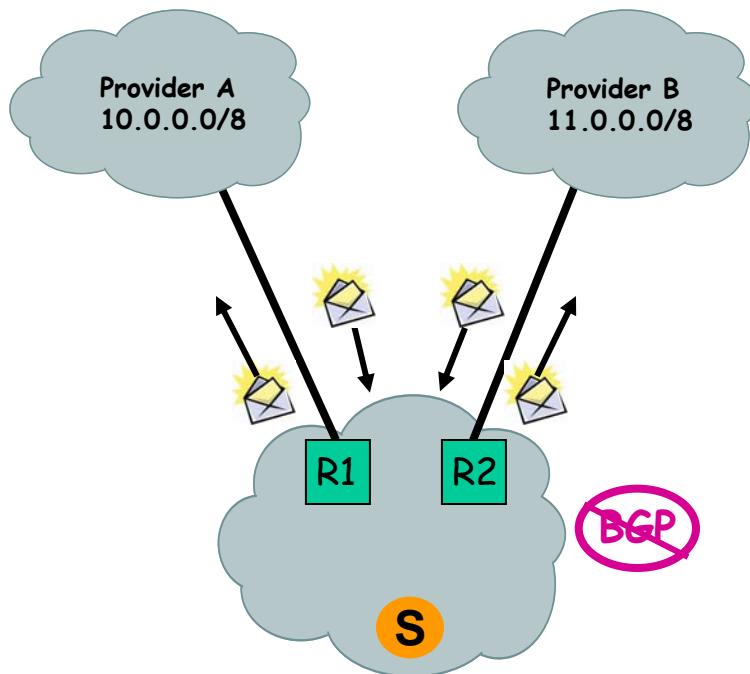
Scaling Internet Routing State



Routing Table Size Problem



Growth in Multi-Homing



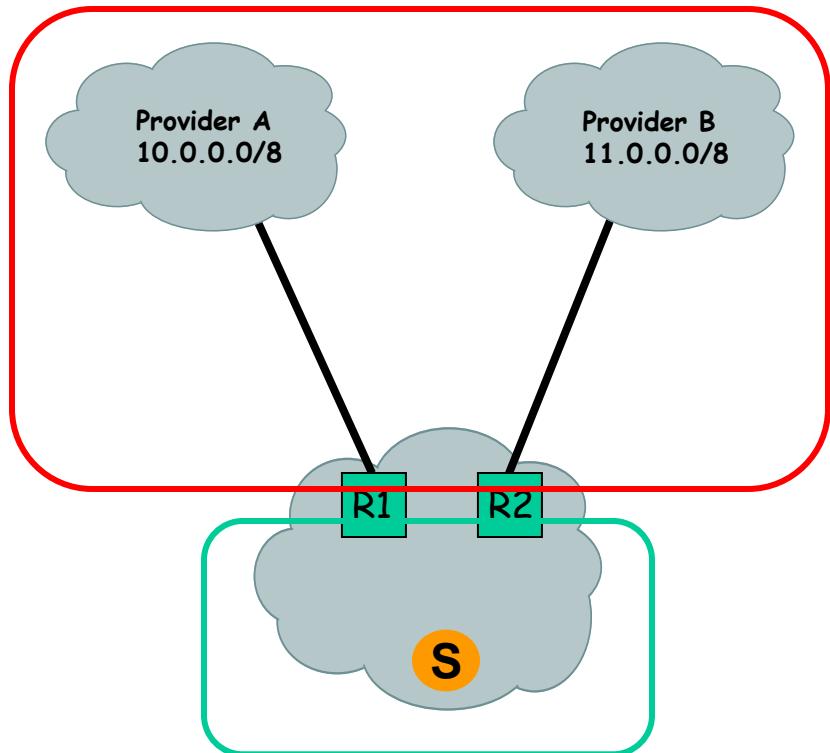
(1) Improve site multi-homing

- a) Can control egress with IGP routing
- b) Hard to control ingress without more specific route injection
- c) Desire to be low OpEx multi-homed (avoid complex protocols, no outsourcing)

(2) Improve ISP multi-homing

- a) Same problem for providers, can control egress but not ingress, more specific routing only tool to circumvent BGP path selection

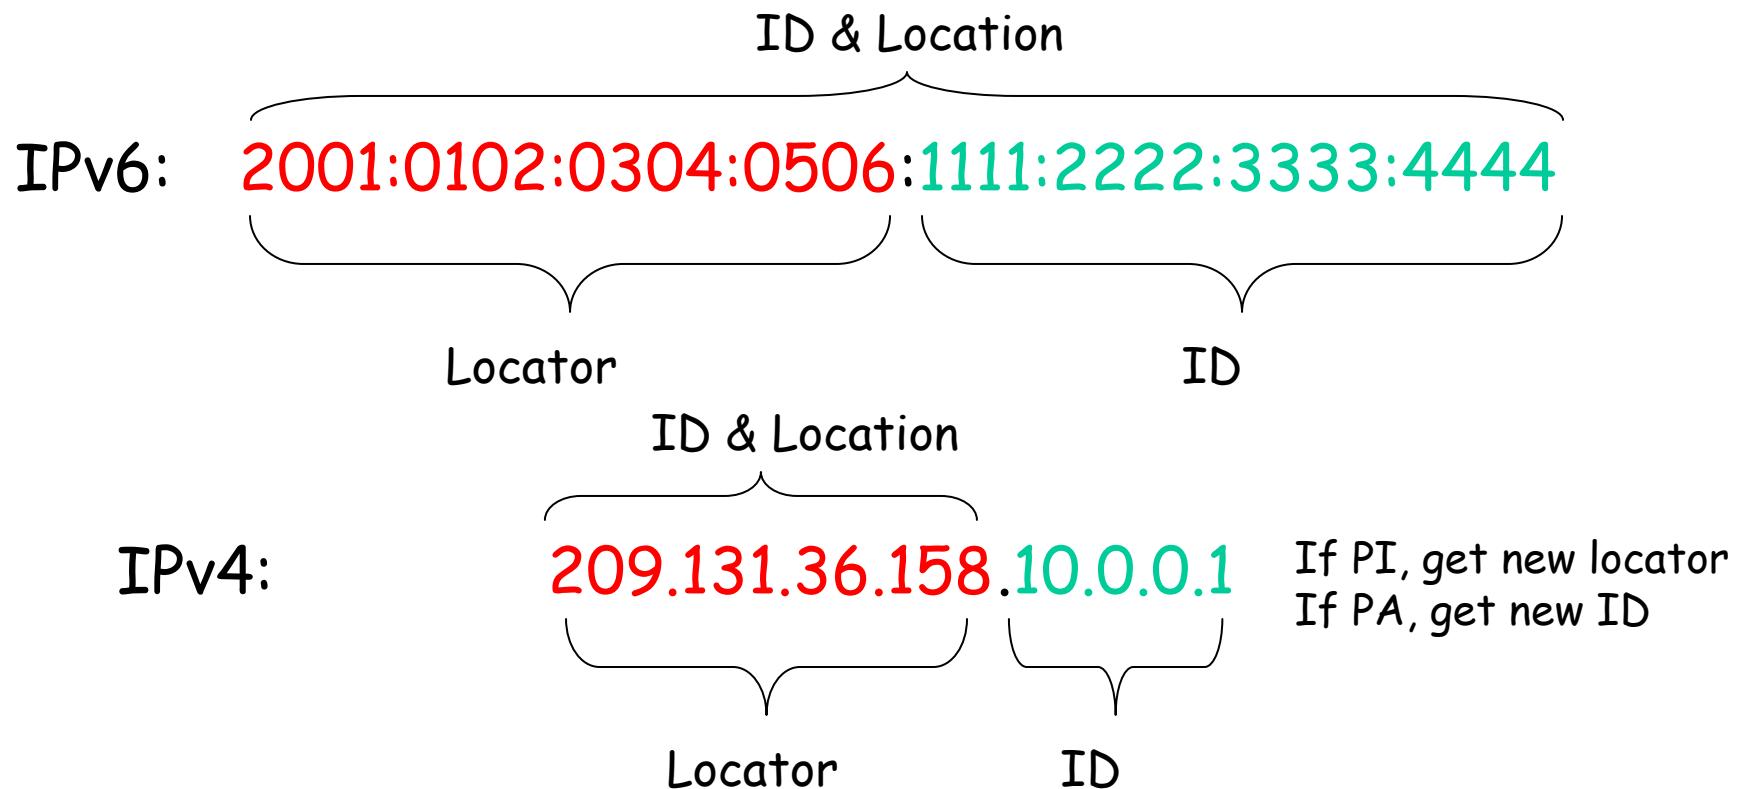
Growth in Multi-Homing



- (3) Decouple site addressing from provider
 - a) Avoid renumbering when site changes providers
 - b) Site host and router addressing decoupled from core topology
- (4) Add new addressing domains
 - a) From possibly separate allocation entities
- (5) Do 1) through 4) and reduce the size of the core routing tables

Separating (or adding) an Address

Changing the semantics of the IP address



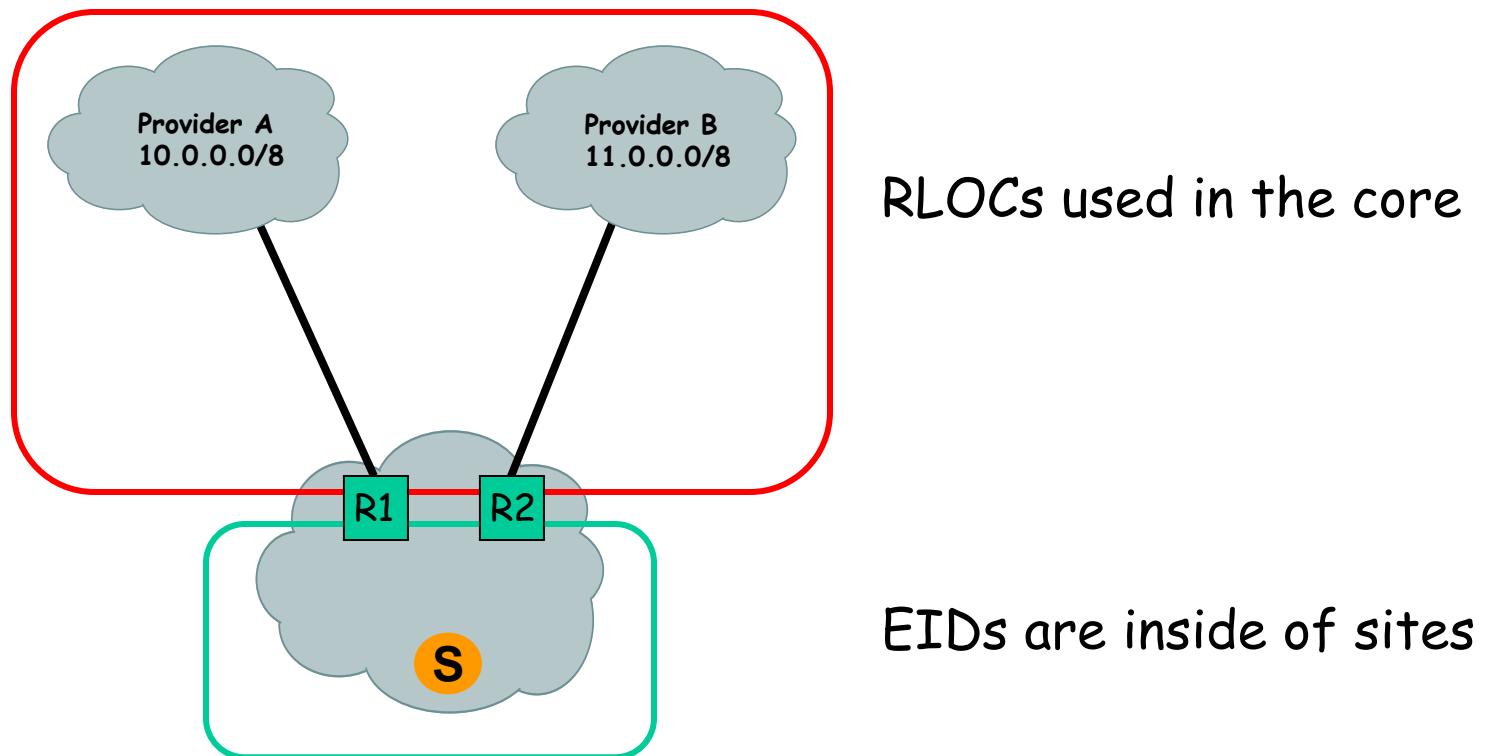
Why the Separation?

- *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
- 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

Some Brief Definitions

- IDs or EIDs
 - End-site addresses for hosts and routers at the site
 - They go in DNS records
 - Generally not globally routed on underlying infrastructure
 - New namespace
- RLOCs or Locators
 - Infrastructure addresses for LISP routers and ISP routers
 - Hosts do not know about them
 - They are globally routed and aggregated along the Internet connectivity topology
 - Existing namespace

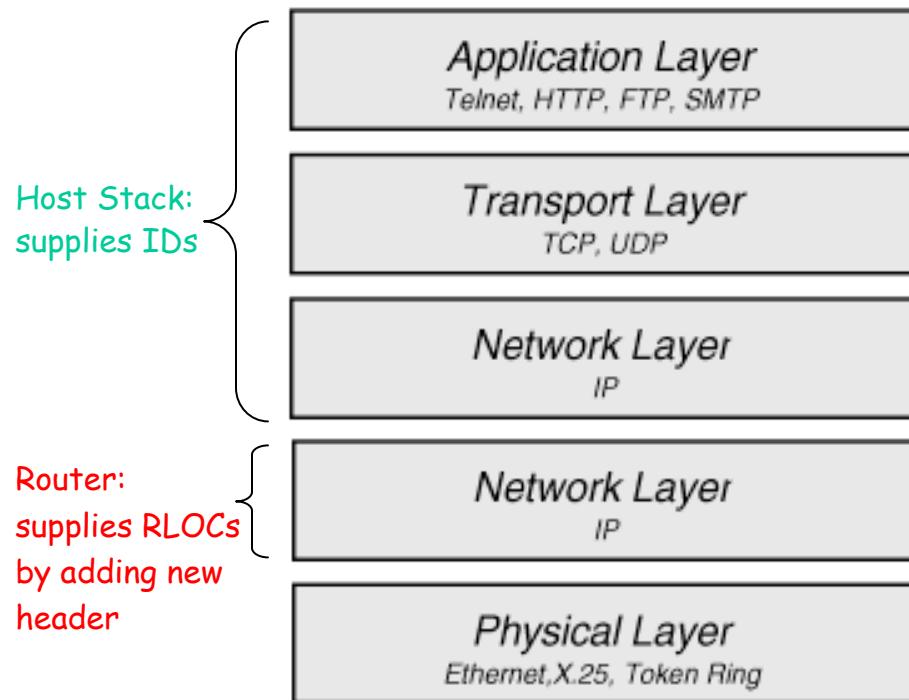
Multi-Level Addressing



What is LISP?

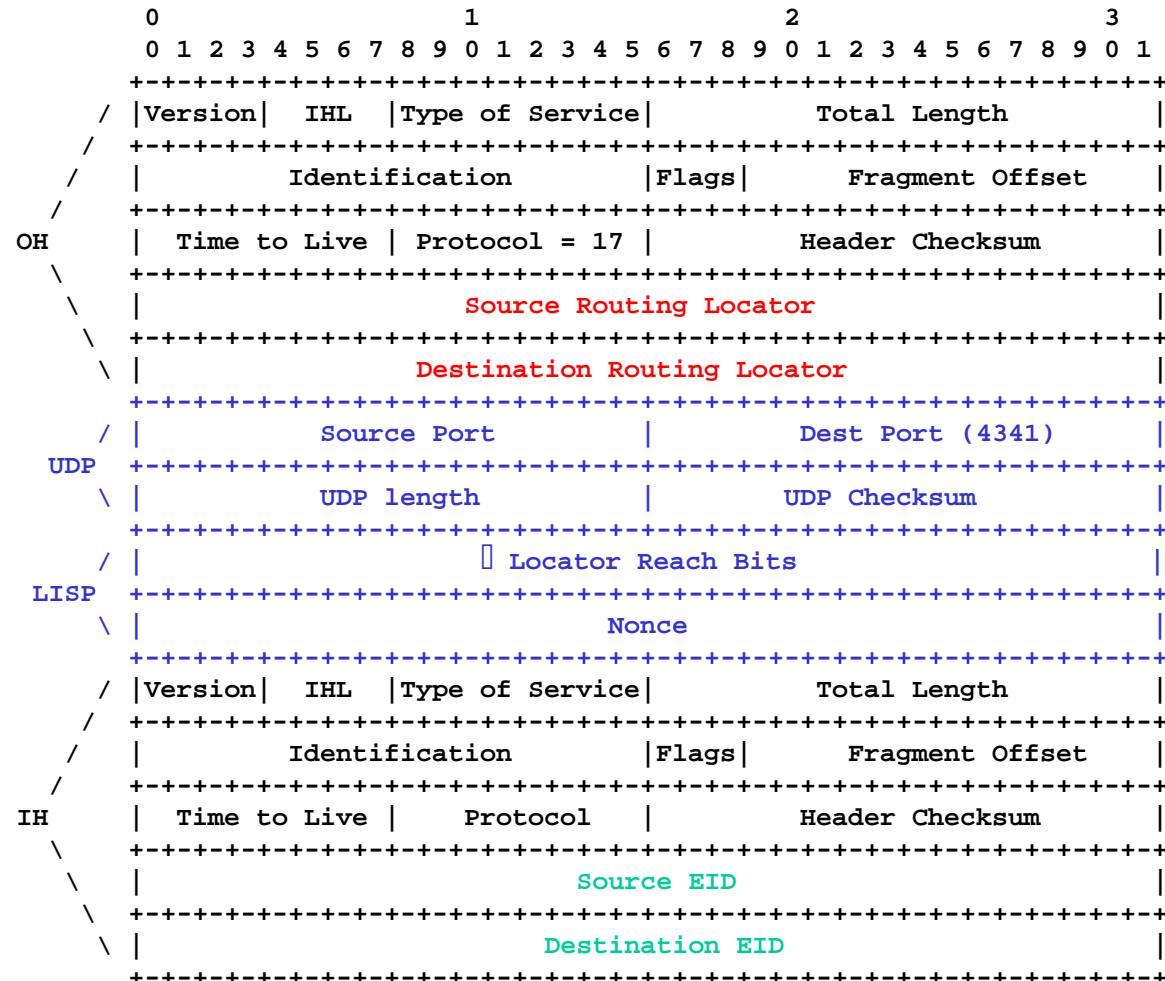
- Locator/ID Separation Protocol
- Ground rules for LISP
 - Network-based solution
 - No changes to hosts whatsoever
 - No new addressing changes to site devices
 - Very few configuration file changes
 - Imperative to be incrementally deployable
 - Address family agnostic

What is LISP?

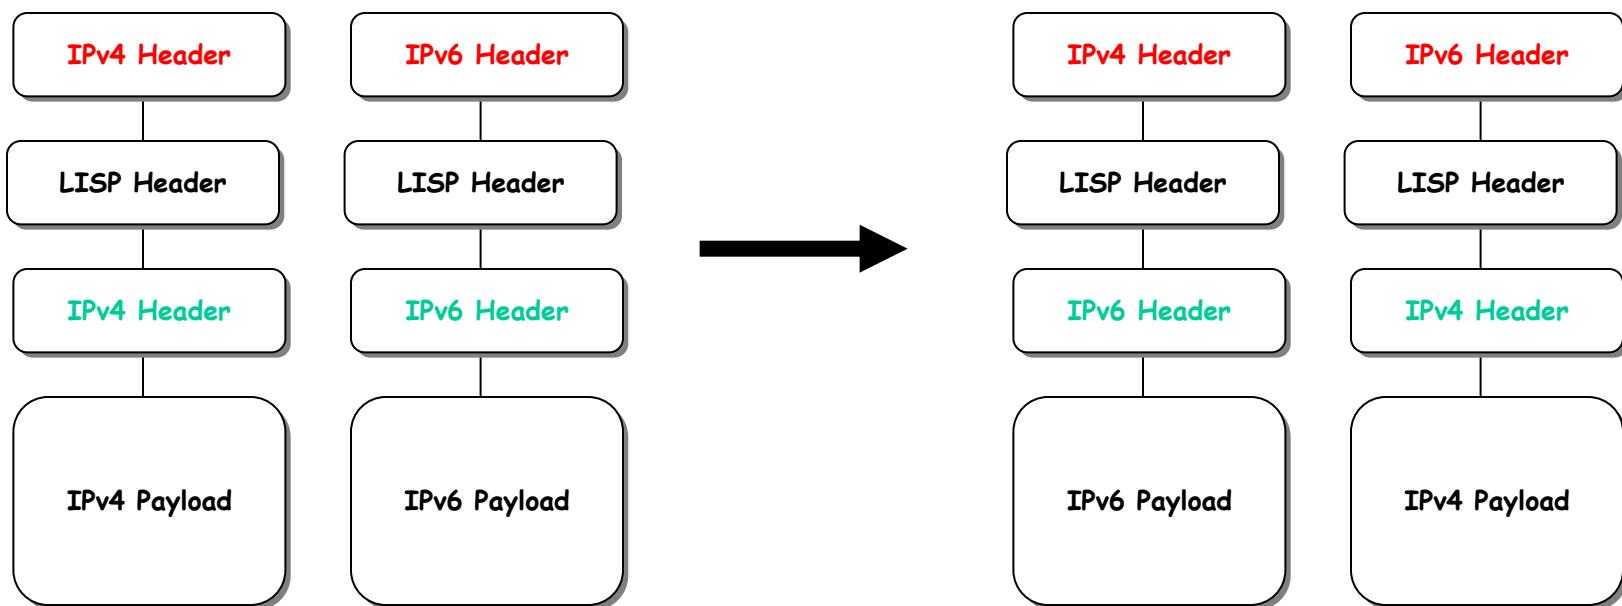


“Jack-Up” or “Map-n-Encap”

[draft-farinacci-lisp-11.txt](#)



LISP for IPv6 Transition



Legend:

EIDs → Green

Locators → Red

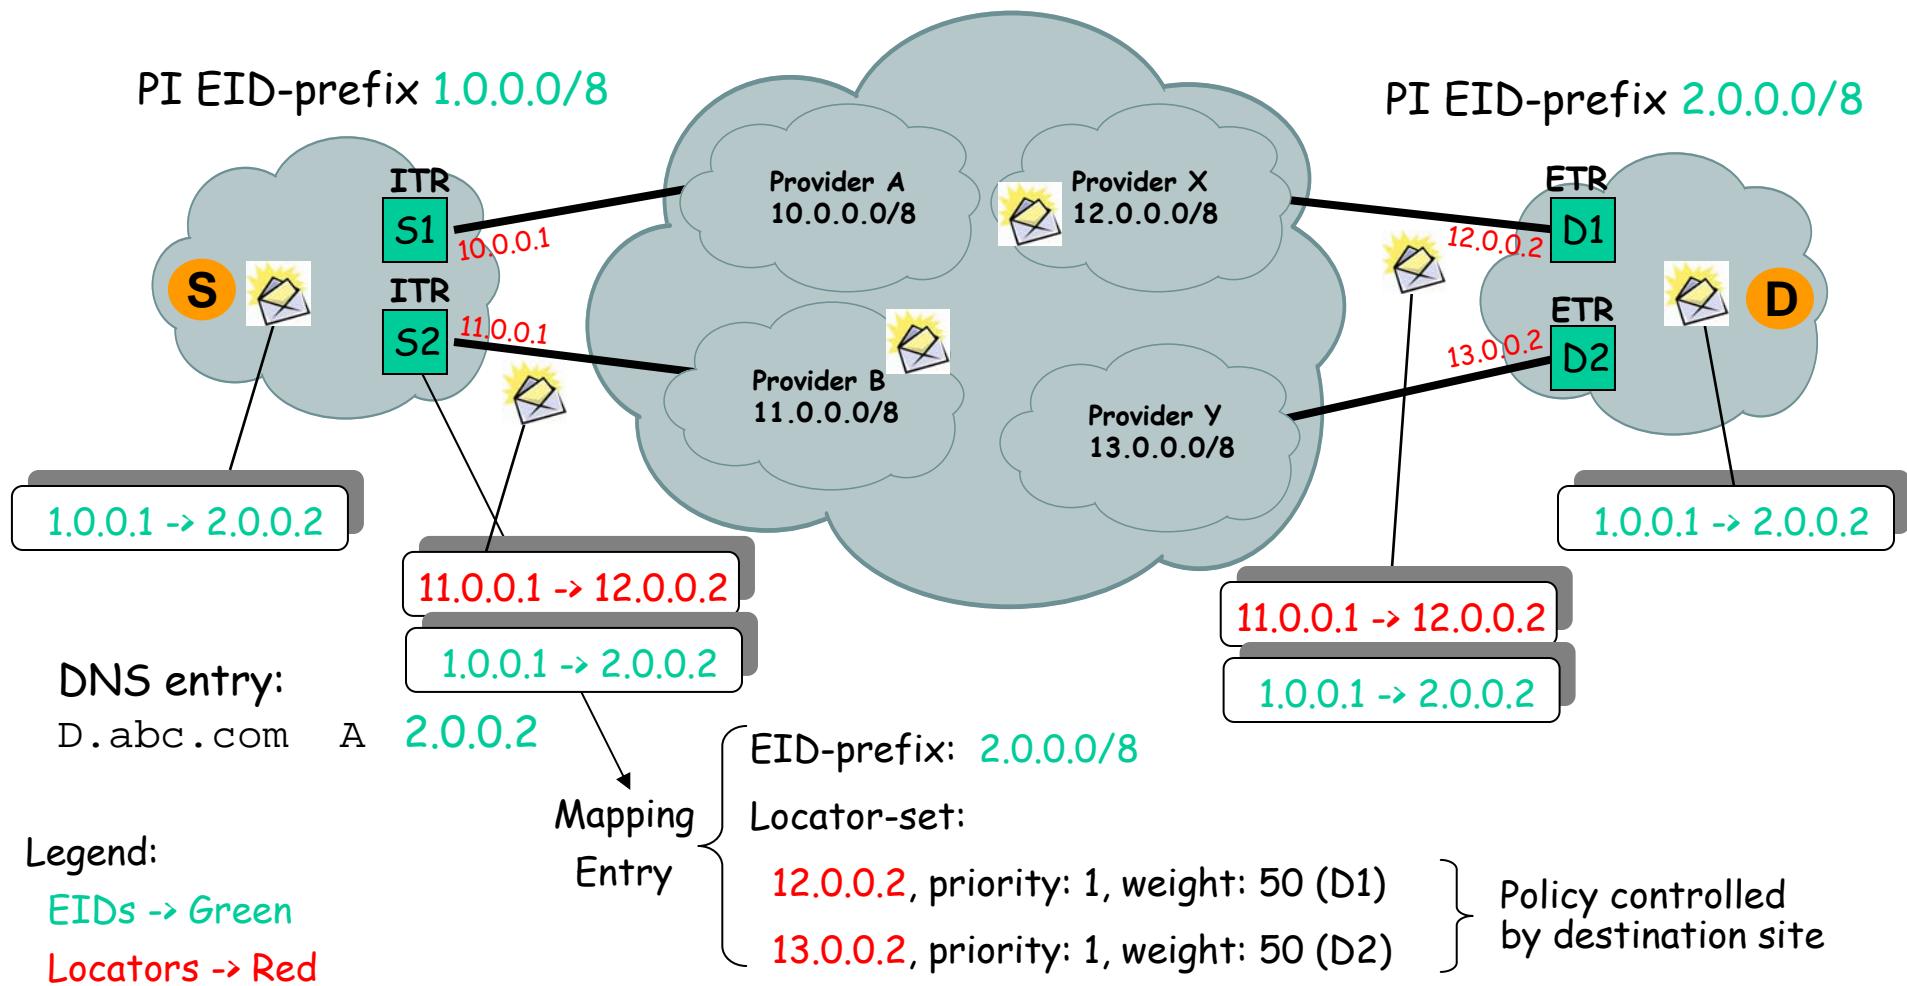
What is LISP?

- 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
 - Examples are: CONS, NERD, ALT, EMACS

LISP Network Elements

- Ingress Tunnel Router (ITR)
 - Finds EID to RLOC mapping
 - Encapsulates to Locators at source site
- Egress Tunnel Router (ETR)
 - Owns EID to RLOC mapping
 - Decapsulates at destination site
- xTR
 - Term used when not referring to directionality
 - Basically a LISP router

Unicast Packet Forwarding



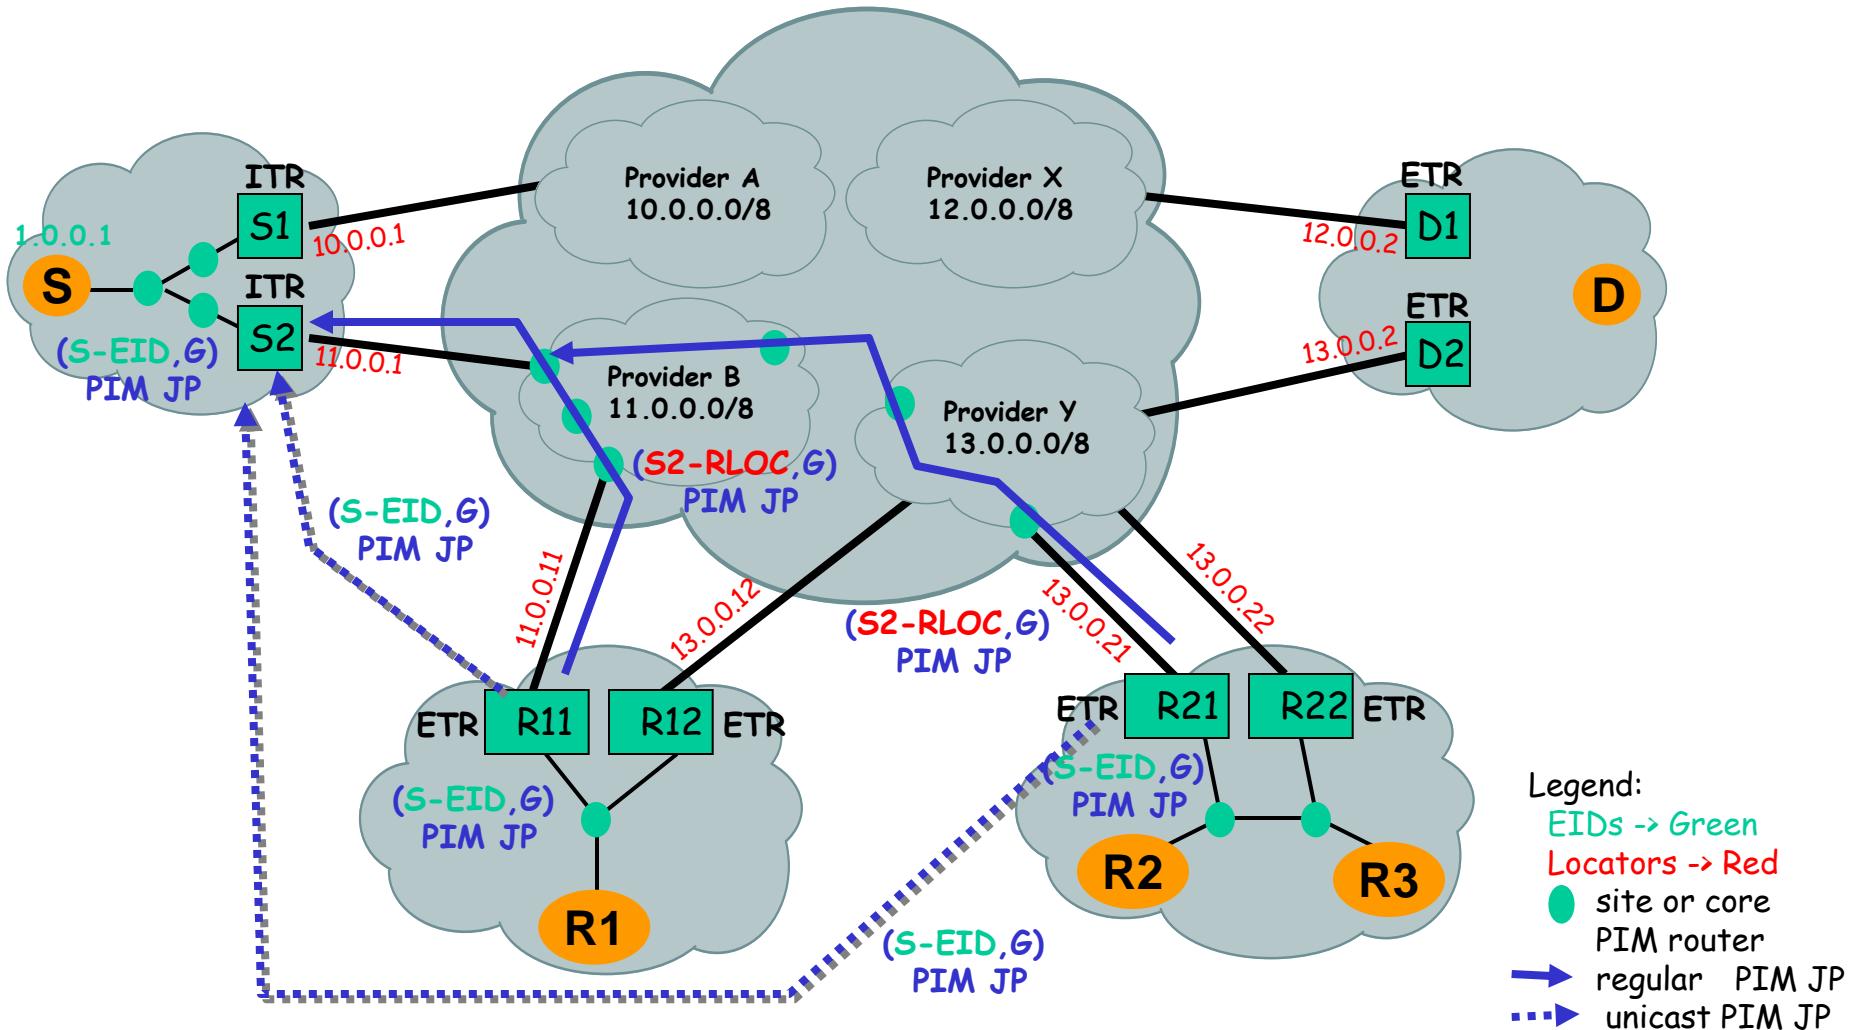
Multicast Packet Forwarding

- Keep EID state out of core network
- No head-end replication at source site
- Packets only go to receiver sites
- No changes to hosts, site routers, core routers
- Use existing protocols
- Support PIM SSM, don't preclude ASM & Bidir
- Have separate unicast and multicast policies

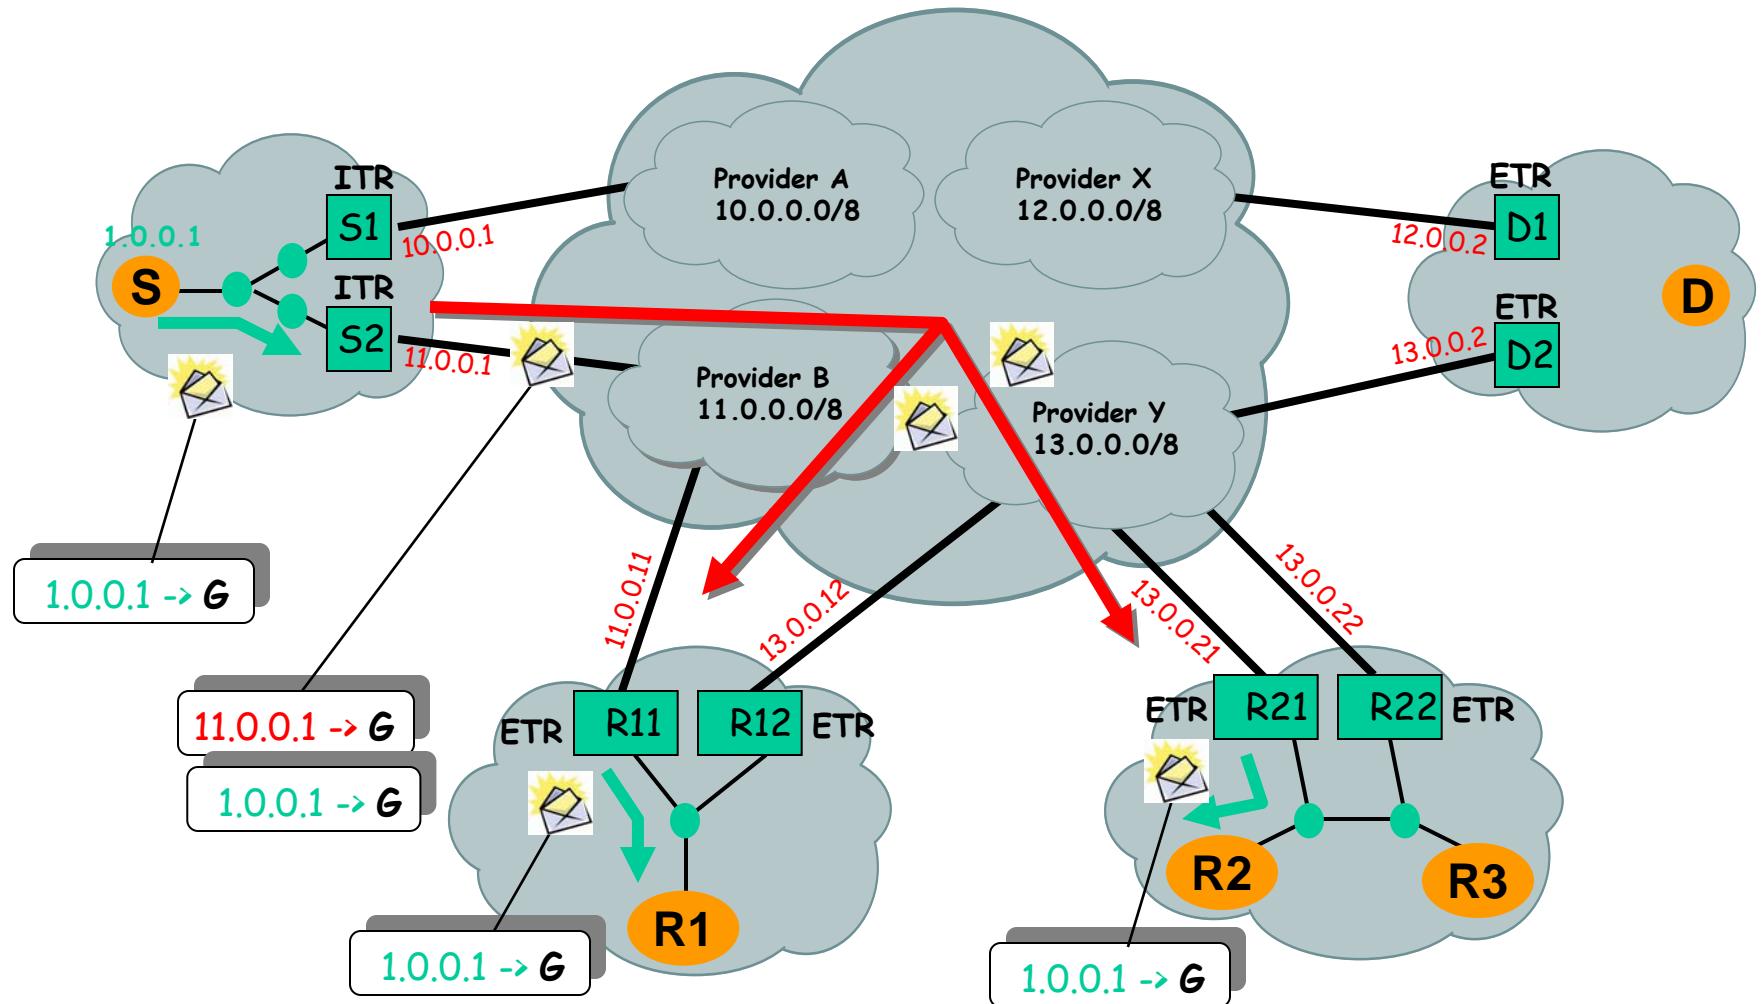
Multicast Packet Forwarding

- Group addresses have neither ID or Location semantics
 - G is topologically opaque - can be used everywhere
- $(S\text{-EID}, G)$
 - $S\text{-EID}$ is source host
 - G is group address receivers join to
 - State resides in source and receiver sites
- $(S\text{-RLOC}, G)$
 - $S\text{-RLOC}$ is ITR on multicast tree
 - G is group address receivers join to
 - State resides in core

Multicast Packet Forwarding



Multicast Packet Forwarding



When the xTR has no Mapping

- Need a scalable EID to Locator mapping lookup mechanism
- Network based solutions
 - Have query/reply latency
 - Can have packet loss characteristics
 - Or, have a full table like BGP does
- How does one design a scalable Mapping Service?

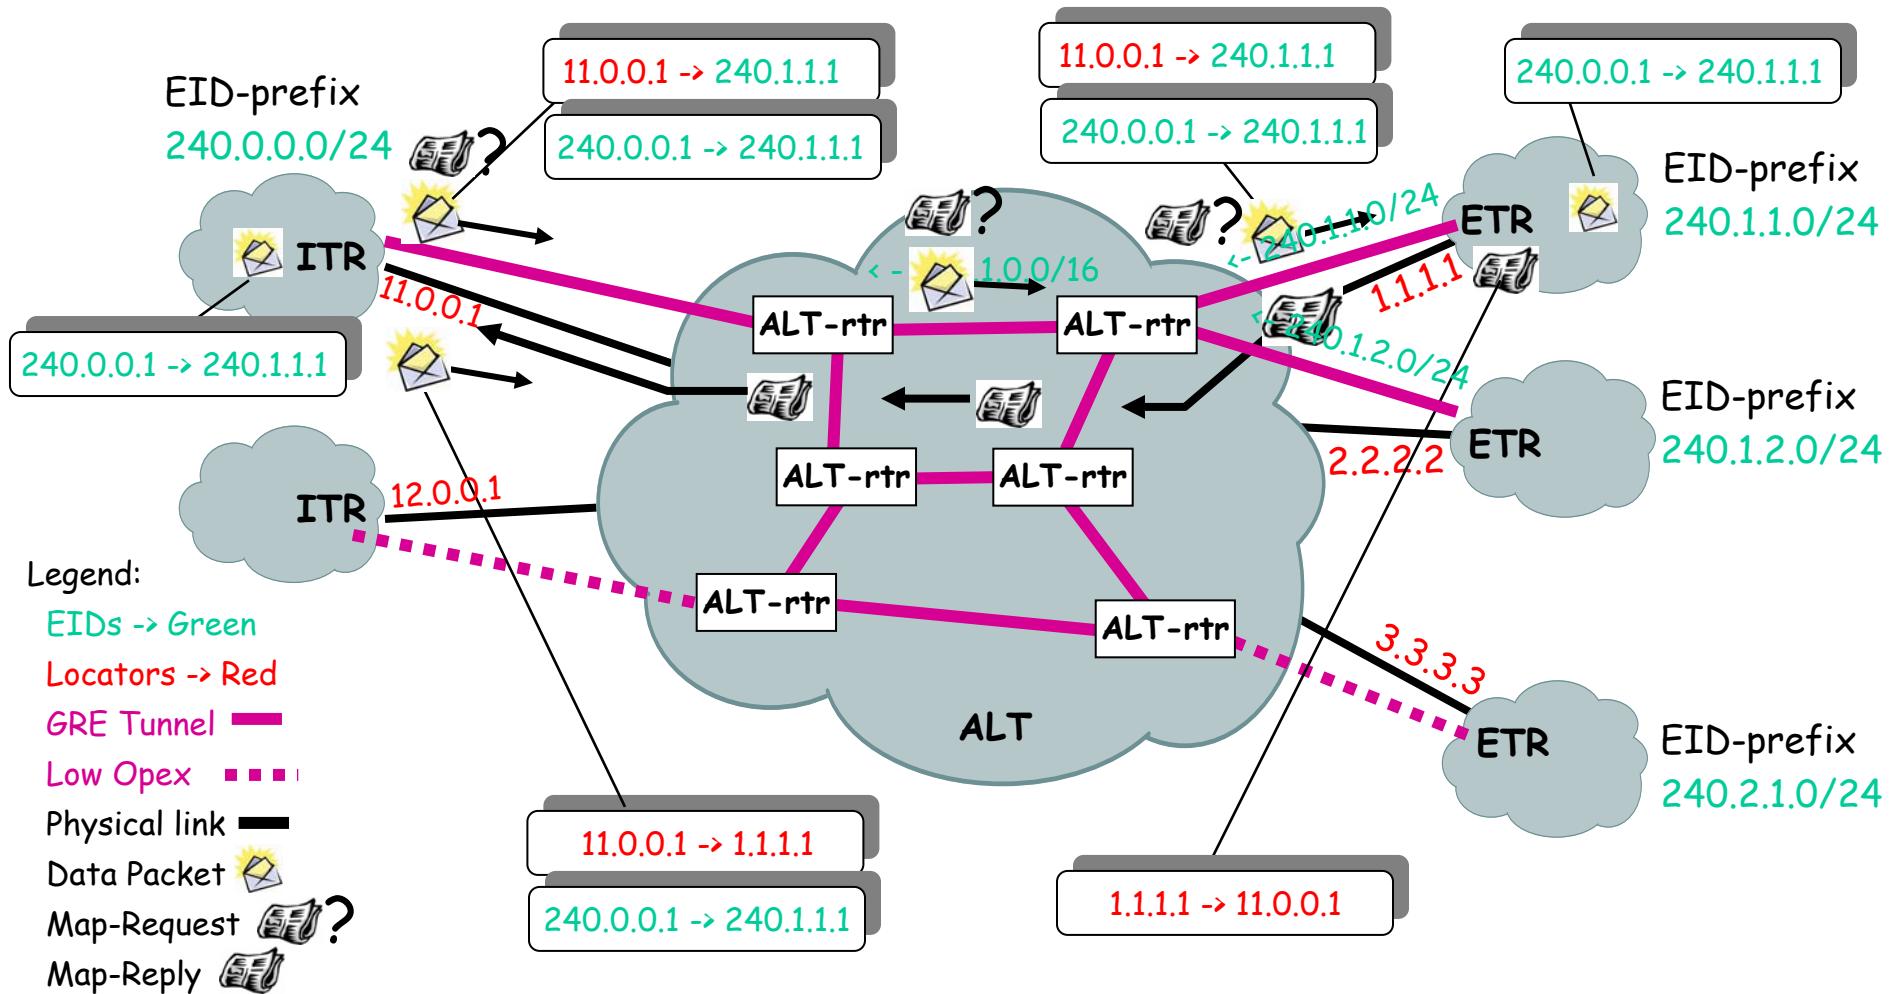
Mapping Database Designs

- You need a “map” before you can “encap”
- We have designed several mapping database protocols
 - CONS, NERD, EMACS, ALT
 - Tradeoff push versus pull benefit/cost
 - Needs to be scalable to 10^{10} entries
- ALT has the most promise
 - We are deploying ALT

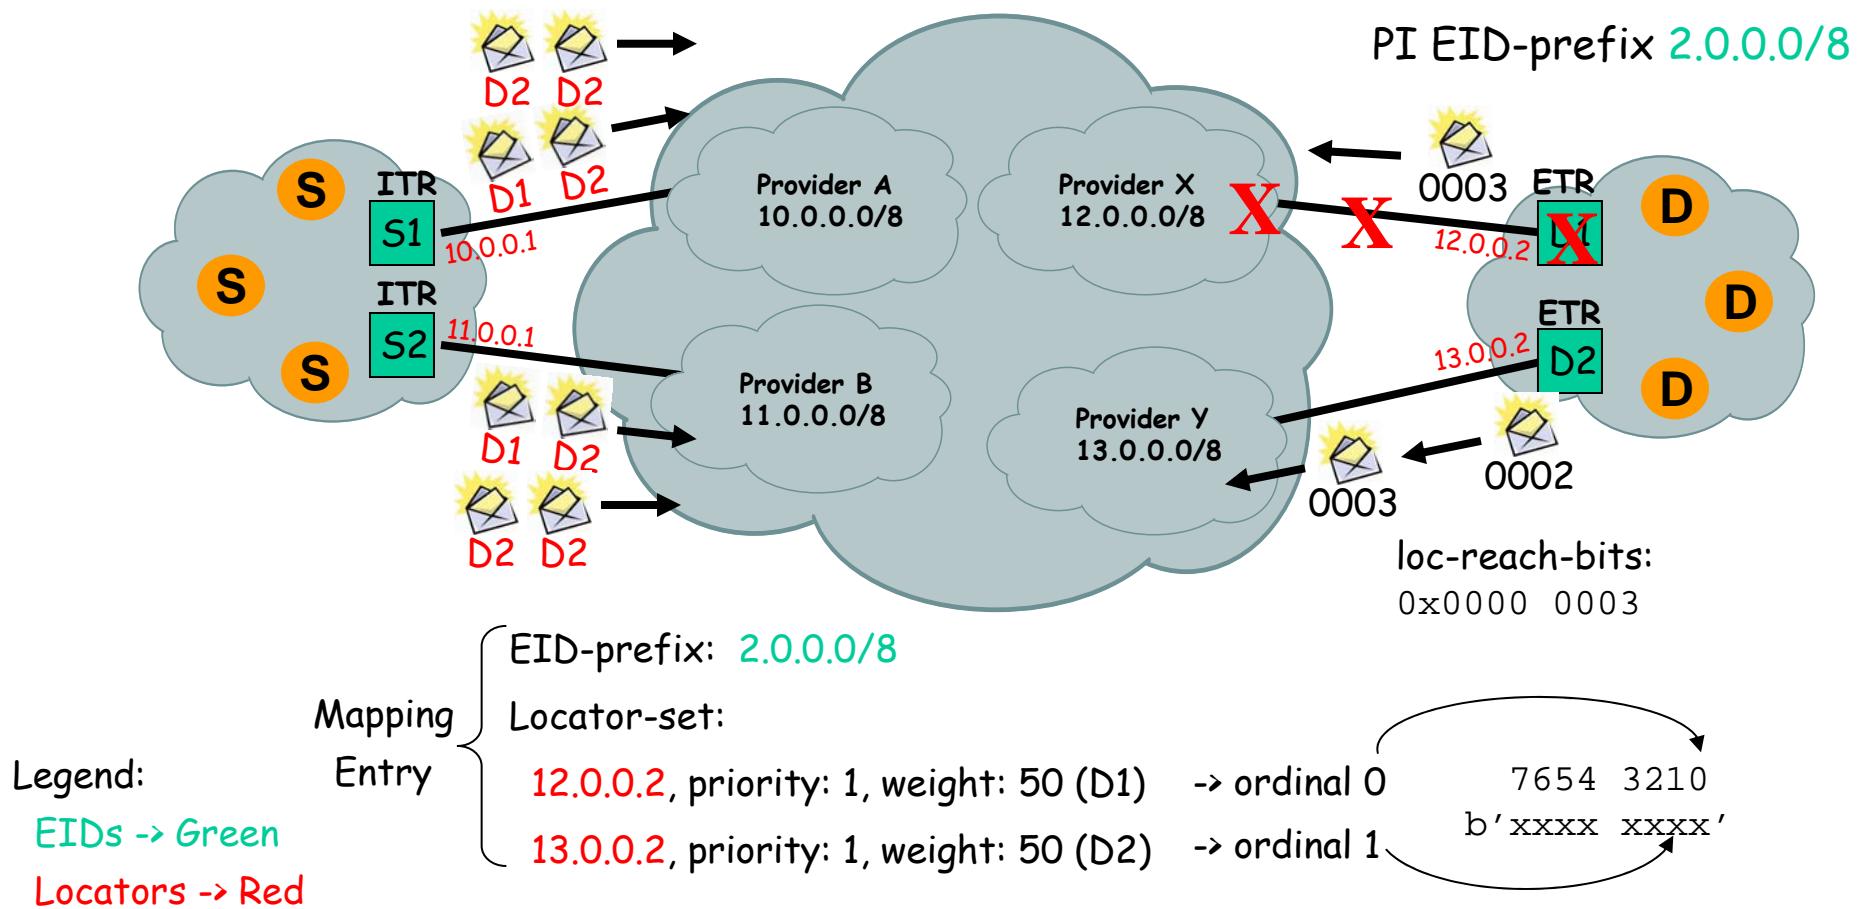
What is LISP+ALT?

- EID namespace is used at the site
- RLOC namespace is used in the Internet core
- Mappings need to be authoritative and reside at site ETRs
- Advertise EID-prefixes in BGP on an alternate topology of GRE tunnels
- ITRs get mappings by routing Map-Requests on ALT topology
- ETRs respond with Map-Replies

How LISP+ALT Works



Locator Reachability



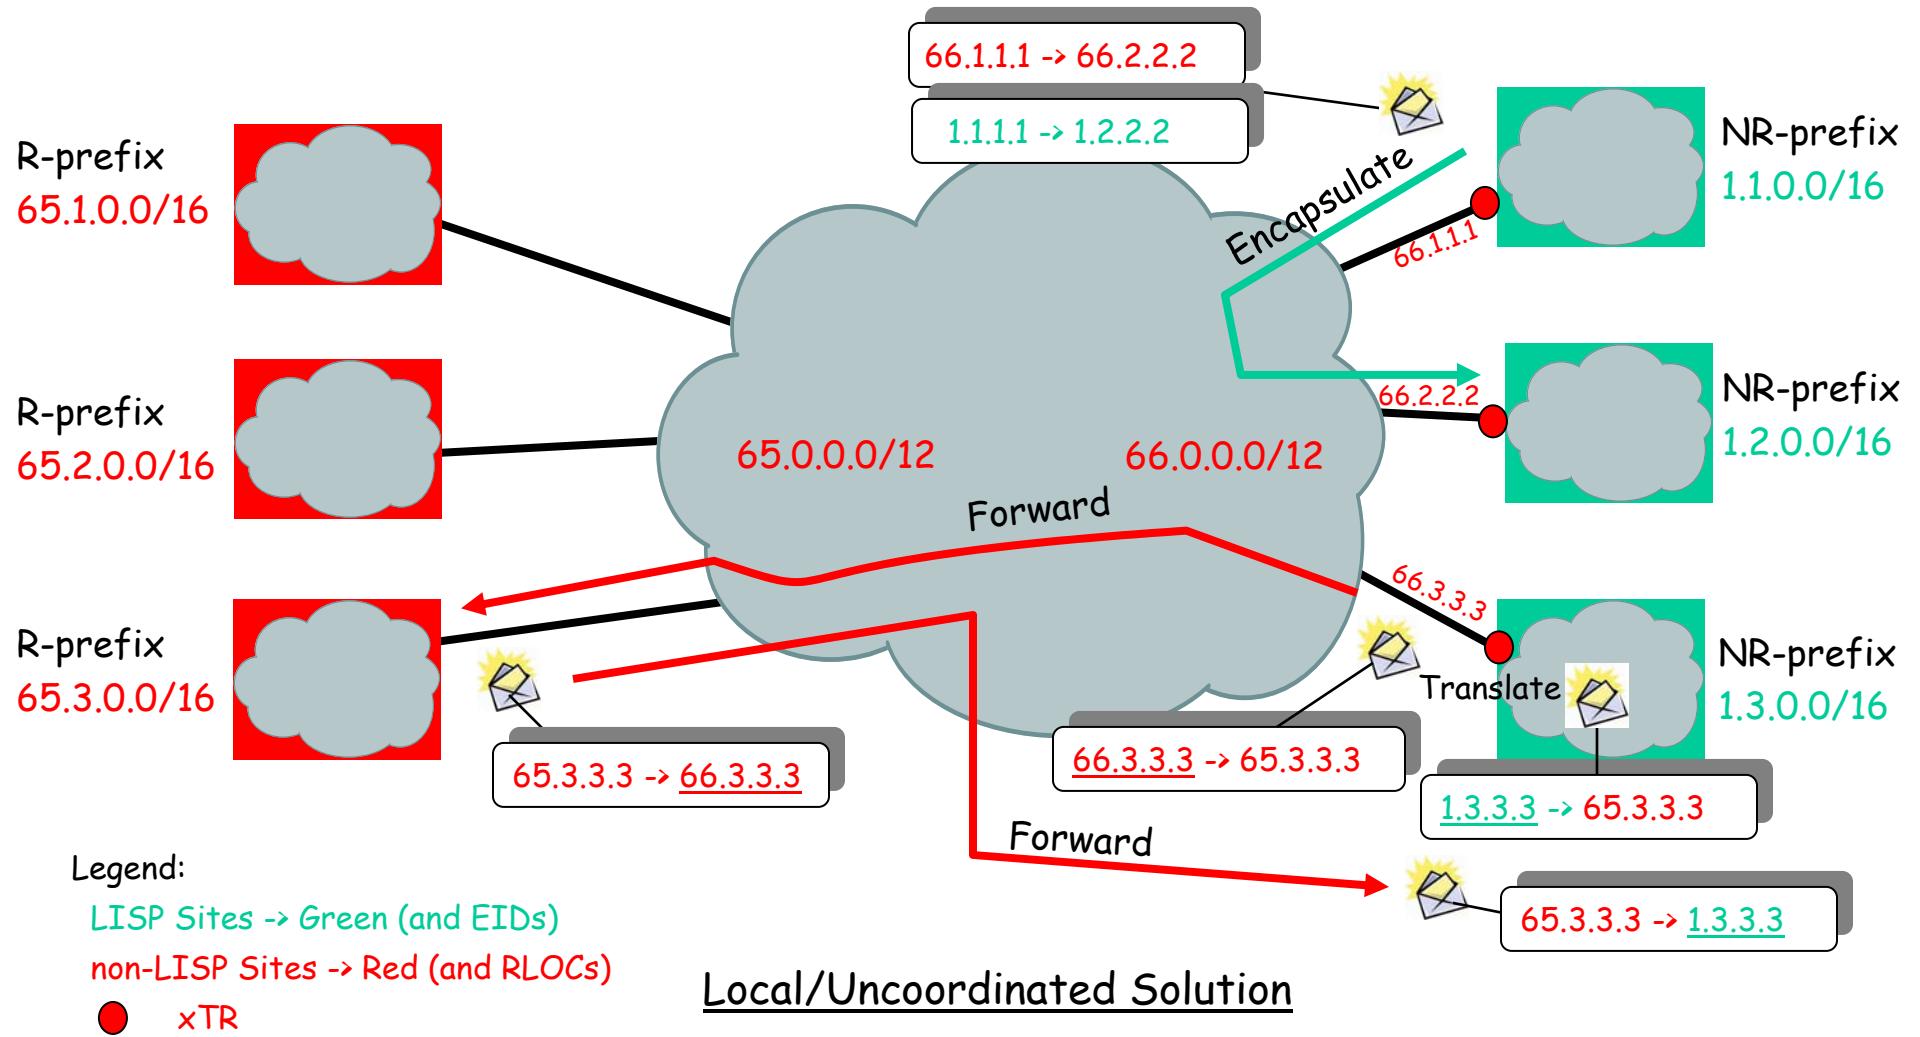
LISP Interworking

- LISP will not be widely deployed day-1
- Need a way for LISP-capable sites to communicate with rest of Internet
- Two basic Techniques
 - LISP Network Address Translators (LISP-NAT)
 - Proxy Tunnel Routers (PTRs)

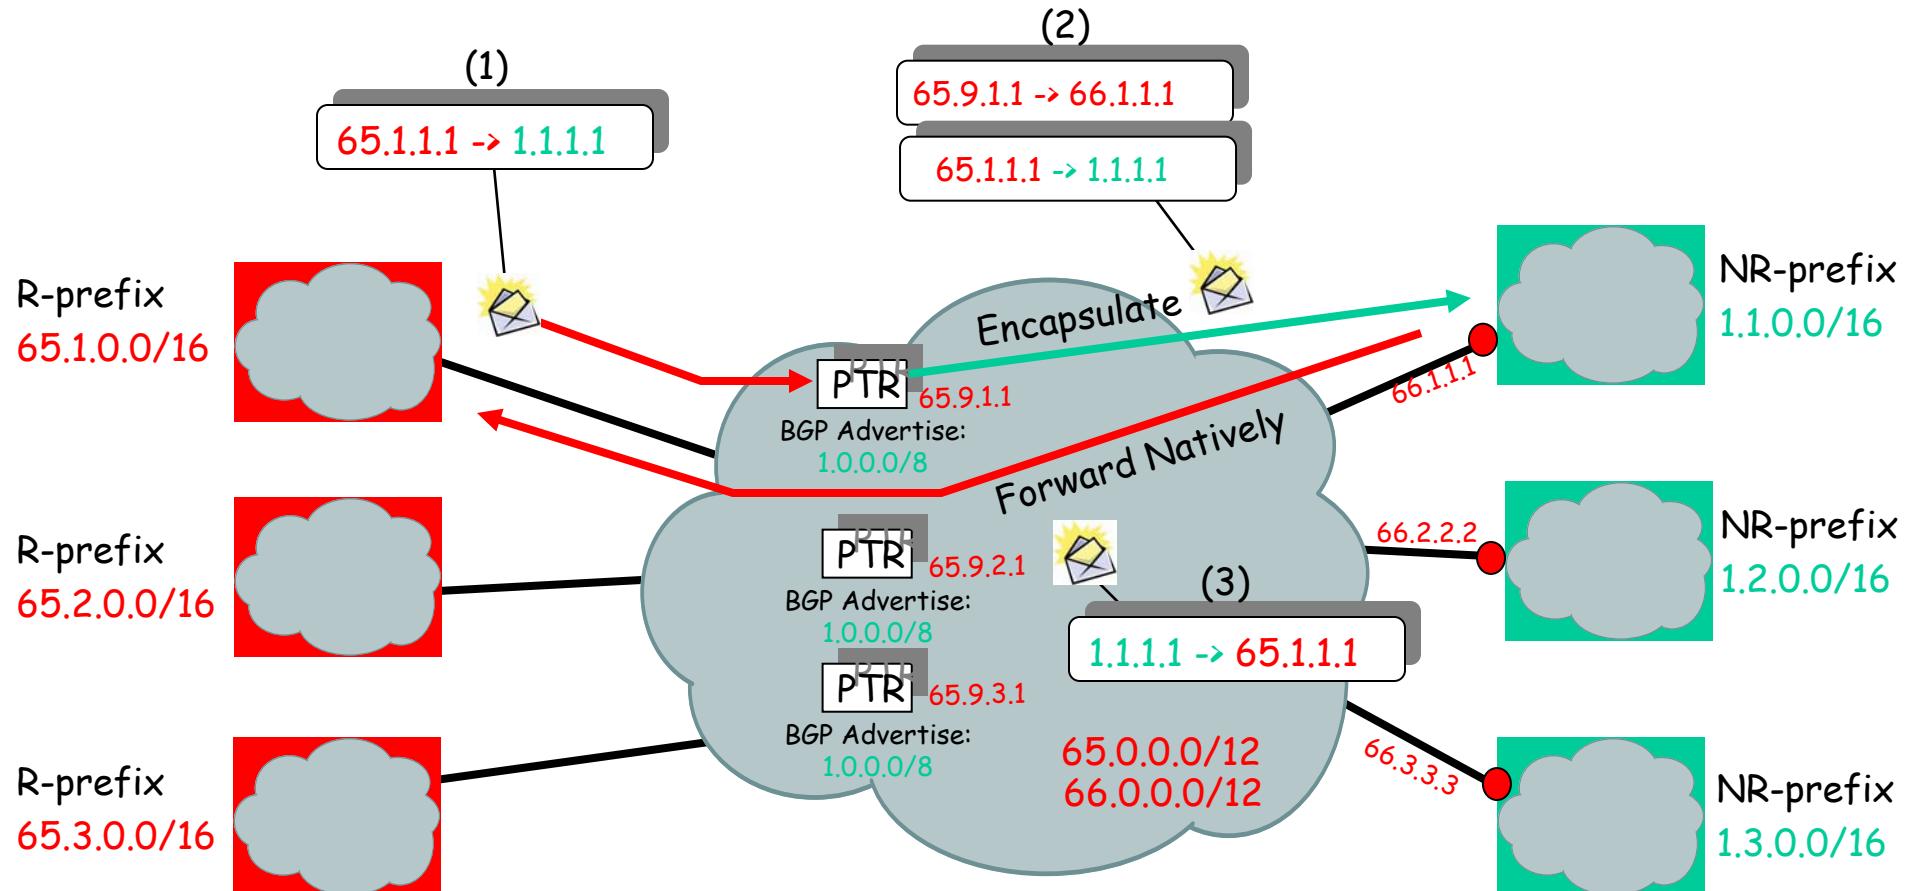
LISP Interworking

- These combinations must be supported
 - Non-LISP site to non-LISP site
 - Today's Internet
 - LISP site to LISP site
 - Encapsulation over IPv4 makes this work
 - IPv4-over-IPv4 or IPv6-over-IPv4
 - LISP-R site to non-LISP site
 - When LISP site has PI or PA routable addresses
 - LISP-NR site to non-LISP site
 - When LISP site has PI or PA non-routable addresses

Interworking using LISP-NAT



Interworking using PTRs



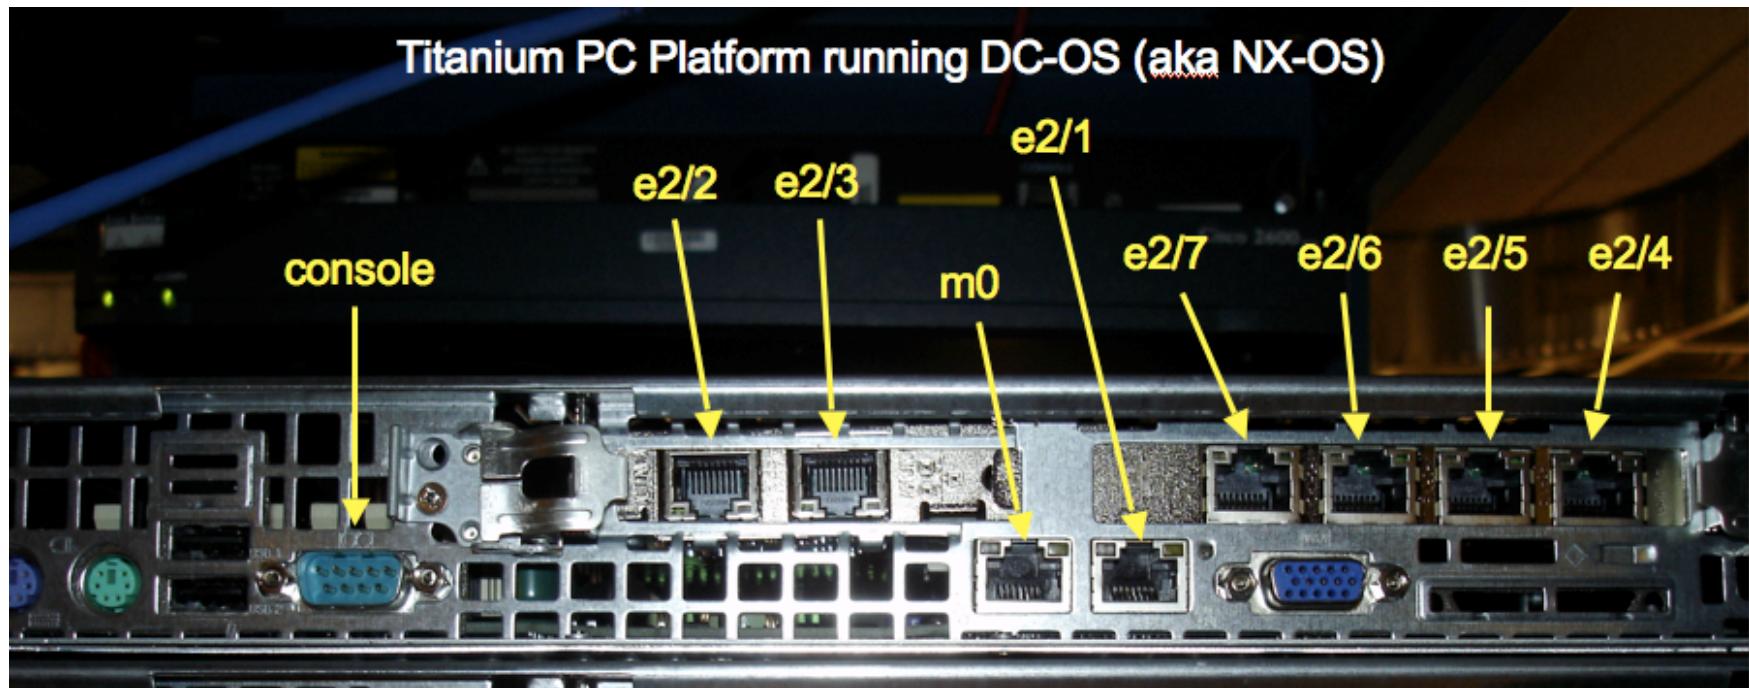
Infrastructure Solution

Prototype Implementation

- cisco has a LISP prototype implementation
- Supports:
 - draft-farinacci-lisp-11.txt
 - draft-fuller-lisp-alt-03.txt
 - draft-lewis-lisp-interworking-02.txt
- Software switching only
- Supports LISP for both IPv4 and IPv6
 - ITR, ETR, and PTR
 - LISP-NAT for IPv4 only

Platform and OS

- NX-OS runs on Linux
- Used as a research platform
- Can run on Nexus 7K/5K & Service Blades



Implementation Schedule

- Started implementation at Prague IETF
 - March 2007
- Implementation put on pilot network
 - July 2007
- Since then released over 90 release builds
 - Releases occur on demand with new features and bug fixes concurrently
- We have phased testing
 - Unit/System Test done in development
 - Alpha test done on pilot network by Dave/Darrel/Vince/John/Andrew
 - Beta test done on pilot network by volunteers

Other Coding Efforts

- IOS implementation under-way
 - Loc/ID split functionality
- Considering IOS-XR implementation
 - TE-ITR/TE-ETR functionality
- OpenLISP implementation been available for FreeBSD a while and being updated
 - For testing the specs
- Considering native Linux implementation
- Any other efforts?

LISP Deployment

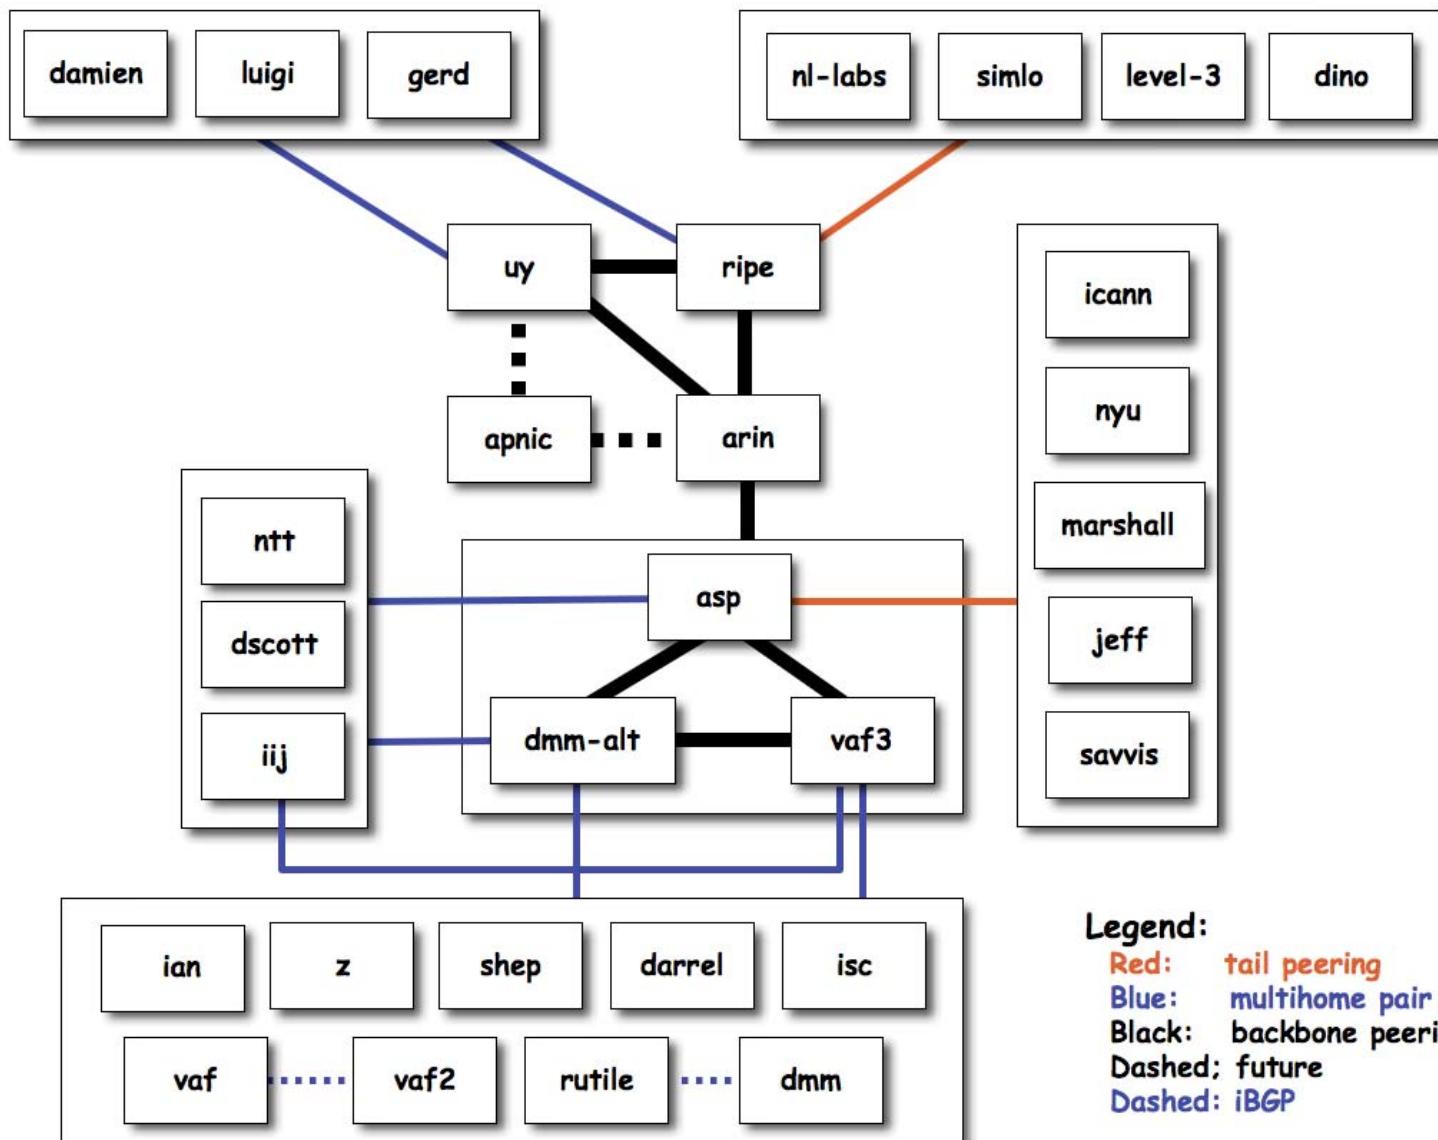
- LISP Pilot Network Operational
 - Deployed for nearly 2 years
 - ~24 sites across 5 countries
 - US, UK, BE, JP, UY
 - Uses the NX-OS Titanium Platform
 - IOS and OpenLISP platforms to be added
 - EID-Prefixes used
 - 153.16.0.0/16 and 2610:00d0::/32
 - RLOCs used
 - Current site attachment points to the Internet

LISP Deployment

- LISP-ALT Infrastructure Built
 - GRE tunnels numbered out of 240.0.0.0/4
 - LISP-ALT uses 32-bit AS numbers
 - Format is 32656.x (ASDot notation)
 - EID-prefixes BGP advertised from 'lisp' VRF
 - RLOCs in default VRF

LISP-ALT Peering Topology

Wed Jan 14 16:04:18 PST 2009



Naming & Addressing

- Domain name lisp4.net
 - Contains hosts and routers from IPv4 EID space
- Domain name lisp6.net
 - Contains hosts and routers from IPv6 EID space

Naming & Addressing

- IPv4 EID Assignments from 153.16.0.0/16
 - North America 153.16.0.0/20
 - /22 for regions in the US
 - Europe 153.16.32.0/20
 - Asia 153.16.64.0/20
 - /21 for regions in Asia
 - Africa 153.16.96.0/20
 - Latin America 153.16.128.0/20
 - Reserved 153.16.{160,192,224}.0/20

Naming & Addressing

- IPv6 EID Assignments from $2610:00d0::/32$
 - $2610:00d0:\text{x}000::/36$
 - x is continent
 - $2610:00d0:\text{xy}00::/40$
 - y is region in continent x
 - $2610:00d0:\text{xy}00::/48$
 - Sites allocate out of /48

LISP Deployment

- LISP Interworking Deployed
 - Have LISP 1-to-1 address translation working
 - <http://www.translate.lisp4.net>
 - Proxy Tunnel Router (PTR)
 - IPv4 PTRs:
 - Andrew, ISC, and UY
 - IPv6 PTRs:
 - Dave (UofO), ISC, and UY
 - <http://www.lisp6.net> reachable through IPv6 PTR
 - <http://www.ptr.lisp4.net> reachable through IPv4 PTR

Other Uses for LISP

- SLBs in Data Centers
 - ETRs directly connected to servers
 - ITRs at Data Center edge
- A/V Mobile Truck Roll
 - Avoid renumber at each event
- BGP-free Core
 - Intra-AS avoiding storing external routes
 - RLOCs are PE routers
- Building topological hierarchy with flat addressing
 - MAC addressing in L2 networks
- MAC address mobility for “extended subnets”
- In an environment of shortage address supply

Open Policy for LISP

- It's been >2 years since the IAB Routing & Addressing WS
 - Some of us committed to working in the IETF and IRTF in an open environment
- This is not a Cisco only effort
 - We have approached and recruited others
 - There are no patents (cisco has no IPR on this)
 - All documents are Internet Drafts
- We need and seek new designers, implementors, and testers
- We need research analysis
- We want this to be an open effort!

Internet Drafts

draft-farinacci-lisp-11.txt
draft-farinacci-lisp-multicast-01.txt
draft-fuller-lisp-alt-03.txt
draft-lewis-lisp-interworking-02.txt
draft-meyer-lisp-eid-block-01.txt
draft-meyer-loc-id-implications-01.txt

draft-mathy-lisp-dht-00.txt
draft-iannone-openlisp-implementation-02.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

References

- Public mailing list:
`lisp@ietf.org`
- More info at:
`http://www.lisp4.net`
`http://www.lisp6.net`

LISP

```

(defun changer-one (failures database)
  (f-p-helper (cond ((null failures) database))
    (failure-points (setf first-part (process-one (car failures) database)))
    (append (list (list index inter-lf)
      (setf (second-part) (process-two (car failures) first-part)))
    (t t)) (cond ((null (cdr failures)) second-part) ; we've reached the conclusion of CBR - return the path
      (defun convert-output-helper-b (pair)
        (format T "CH. ~A~V Path")
      ) (cond ((null path)
        ((ecase (caar pair) (cons (caar path) (rest path)))
          (cons (rest (caadr path)) (first (cadar path)))) ; If we have an exact match between source/destination
          (convert-output-helper-b (rest path)))
        ((ecase (cadar path) (caar (rest path)))
          (cons (list (caar path) (rest path)) (cdadar path))) ; If we have no chosen-path, we can't recurse - we return nil
          (convert-output-helper-b (rest path)))
      ) (cond ((equal (reverse (car pair)) (caar database)) ; If either the two elements of the chosen-path match,
        (equal fail-param 'F) (append (list (list (caar path)) ; If the chosen-path matches the destination
          (append (list (append (list (rest path)) (list 'UNKNOWN) (remove-pair (cadr pair) (cdar database)))) ; chosen-path) (cadr chosen-path))
          (execute-plan-helper source (rest path) 'F)))
        (check-failure (first path) *failure-points* (cdr database)))) ; (equal (cadr chosen-path) destination) nil)
        (append (list (append (list (first path))
          (list 'SUCCESS)))
        (execute-plan-helper source (rest path) (process-one pair (cdar database)))) ; If the chosen-path has more than one element, we
        (t (append (list (append (list (first path)
          (list 'FAILURE)))) ; If the chosen-path has only one element, we're done
        (execute-plan-helper source (append path (list 'UNKNOWN) (remove-pair (cadr pair) (cdar database)))) ; chosen-path) (cadr chosen-path))
        (cond ((eql method :search) ; If the destination method is search, we're done
          ((eql method :cbr) (cbr-super source destination plan))) ; If the destination method is cbr, we're done
        (defun process-two (pair)
          (cond ((null database) nil)
            ((or (equal (cadar pair) (caar database))
              (equal (reverse (cadar pair)) (caar database))) ; If either the two elements of the chosen-path match
              (or (cbr-super source destination plan)
                (t (or (cbr-top (cadr (cbr-evaluate source destination
                  (append path (cbr-get-path chosen-path plan)))
                  (append path (cbr-top (car chosen-path) destination ne
                    (append path (cbr-get-path chosen-path plan))
                    (cbr-new-plan (cbr-evaluate source destination
                      (cbr-retrieve source destination plan))))))) ; If the chosen-path has only one element, we're done
                    (process-one pair (cdar database)))) ; chosen-path) (cadr chosen-path))
              (convert-output-helper-b (rest path)))
            (cbr database)))
        (defvar *database*
          '( (append (list (car database))
            '(((10th Tech-Parkway) (10th Curran) (Tech-Parkway North-Avenue))
              ((10th Curran) (10th McMillan) (10th hemphill) (mcmillan 8th-1)))
            (Curran 8th-1))
            ((10th McMillan) (10th Curran) (10th hemphill) (mcmillan 8th-1))
            ((10th hemphill) (10th mcmillan) (10th dalnev) (hemphill 8th-1))
          )
        )
      )
    )
  )
)

```

RULES