Cisco Networkers 2011

November 7, 2011 - Hong Kong

LISP Architecture, Protocols, and Product Update

Bruce Pinsky, bep@cisco.com

CCIE #1045/CCDE 2007::0003/CCAr

Session 1.1

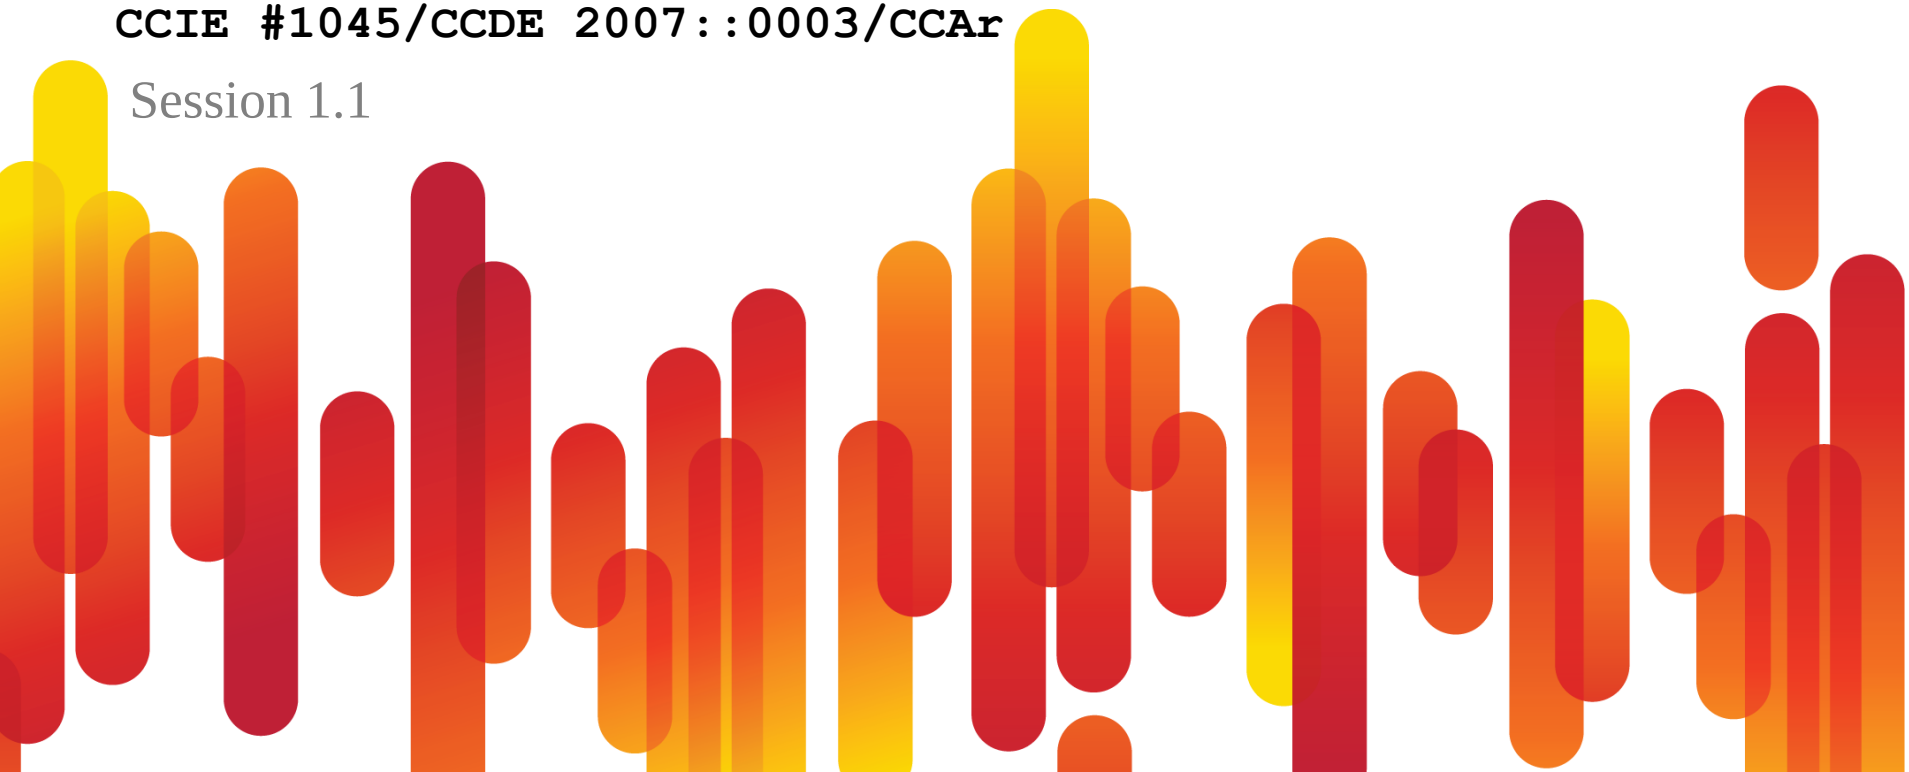

Agenda

- LISP Problem Statement
- LISP Overview
- LISP Product Happenings in 2011
- LISP IETF Happenings in 2011
- What's Next

Today – No ID/Locator Separation

chambers.cisco.com

(1) Is this John at the location 'cisco'?
(2) Is this John at location 'home'?
(3) Is this John at 'Starbucks'?

If I have a connection to John does it break because he changed locations?

198.133.219.25

Application ID

Network ID

Network Location

Future – With ID/Locator Separation

(1) The service binds to an Application Name (DNS)
(2) The Application Name binds to a Endpoint ID (EID)
(3) The EID binds dynamically to a Routing Locator (RLOC)

(1)	Users use DNS names	(the human "who")
(2)	Applications use EIDs	(the network "who")
(3)	Routing uses RLOCs	(the network "where")

DNS Name	EID	RLOC
chambers.cisco.com	198.133.219.25	64.100.100.1

static binding

Provided by the DNS system

dynamic binding

Provided by the LISP routing system

LISP Overview
What's the problem with an "overloaded" semantic?

- The IP address is overloaded on location and identity

Why do current IP semantics cause scaling issues?

- Overloaded IP address semantic makes efficient routing impossible
- Today, "addressing follows topology," which limits route aggregation compactness
- IPv6 does not fix this

Why are route scaling issues bad?

- Routers require expensive memory to hold the Internet Routing Table in forwarding plane
- Your router may have enough memory today; network gear lifetime can be 7 years or more.
- Replacing equipment for the wrong reason (to hold the routing table); gear replacement should be to implement new features and to meet bandwidth requirements

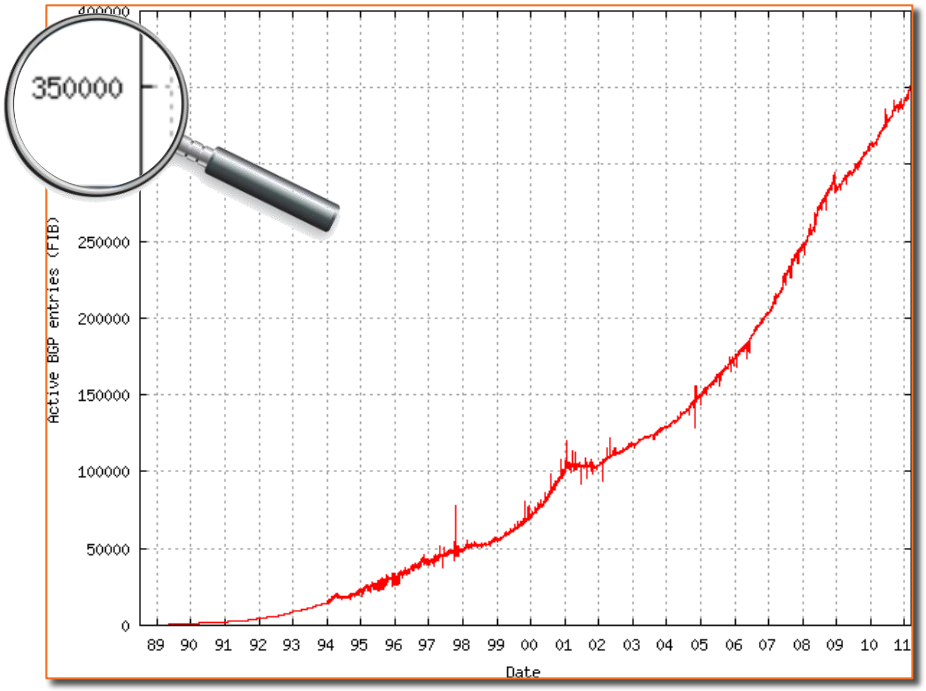

"… routing scalability is the most important problem facing the Internet today and must be solved …"

Internet Architecture Board (IAB)
October 2006 Workshop (written as RFC 4984)

LISP Overview
How does Location/ID Split help solve this problem?

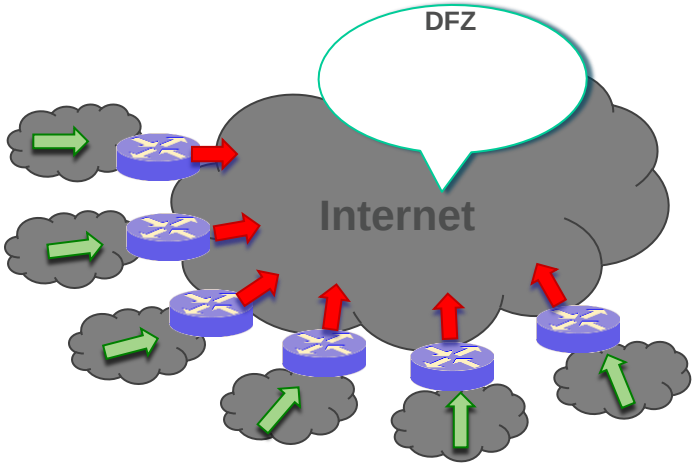

Today's Internet Behavior
Loc/ID "overload"

In this model, everything goes in the Default Free Zone (DFZ)

LISP Mapping System

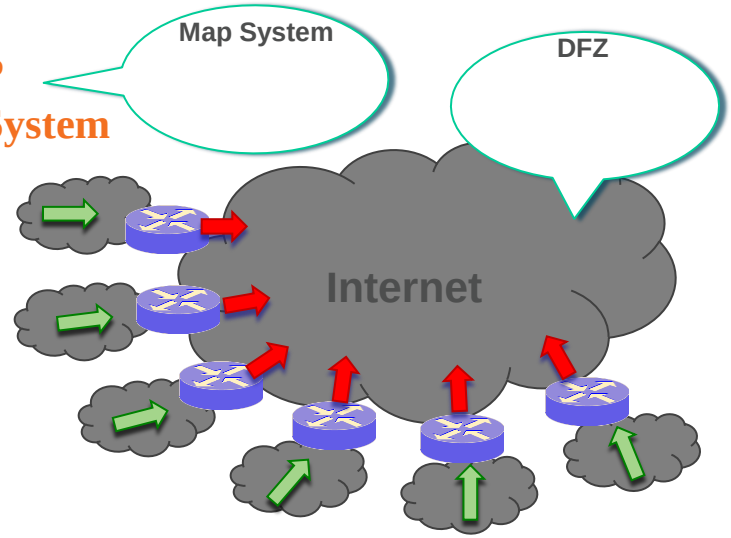

LISP Behavior
Loc/ID "split"

In this model, only RLOCs go in the DFZ; EIDs go in the LISP Mapping System!

What is LISP?

- A new addressing architecture and protocol suite
  - For separating End-point IDs and Locators
- Network-based solution
- No changes to hosts whatsoever
- No addressing changes to site and core devices
- Very few configuration file changes
- Imperative to be incrementally deployable
- Address family agnostic

What is LISP?

- LISP is completely open
  - Started in the IRTF
  - Currently has an IETF working group
  - No known IPR
- 100s of Researchers and Operators Contributed to Design
- Multiple Vendors Interested
- Pilot Network up for nearly 4 years
  - 121 nodes in 25 countries
- Building a LISP-MN Pilot Network
  - Testing server capabilities on Android phones
  - Experimenting new mapping database systems and security mechanisms

IETF LISP WG Status

- 7 Internet Drafts past WG last call
- Currently in AD review
  - draft-ietf-lisp
  - draft-ietf-lisp-alt
  - draft-ietf-lisp-interworking
- Currently IESG last call
  - draft-ietf-lisp-multicast
  - draft-ietf-lisp-ms
  - draft-ietf-lisp-map-versioning
- RFC editor queue (for RFC number assignment)
  - draft-ietf-lisp-lig

LISP Data Plane Overview
Ingress/Egress Tunnel Router (xTR)

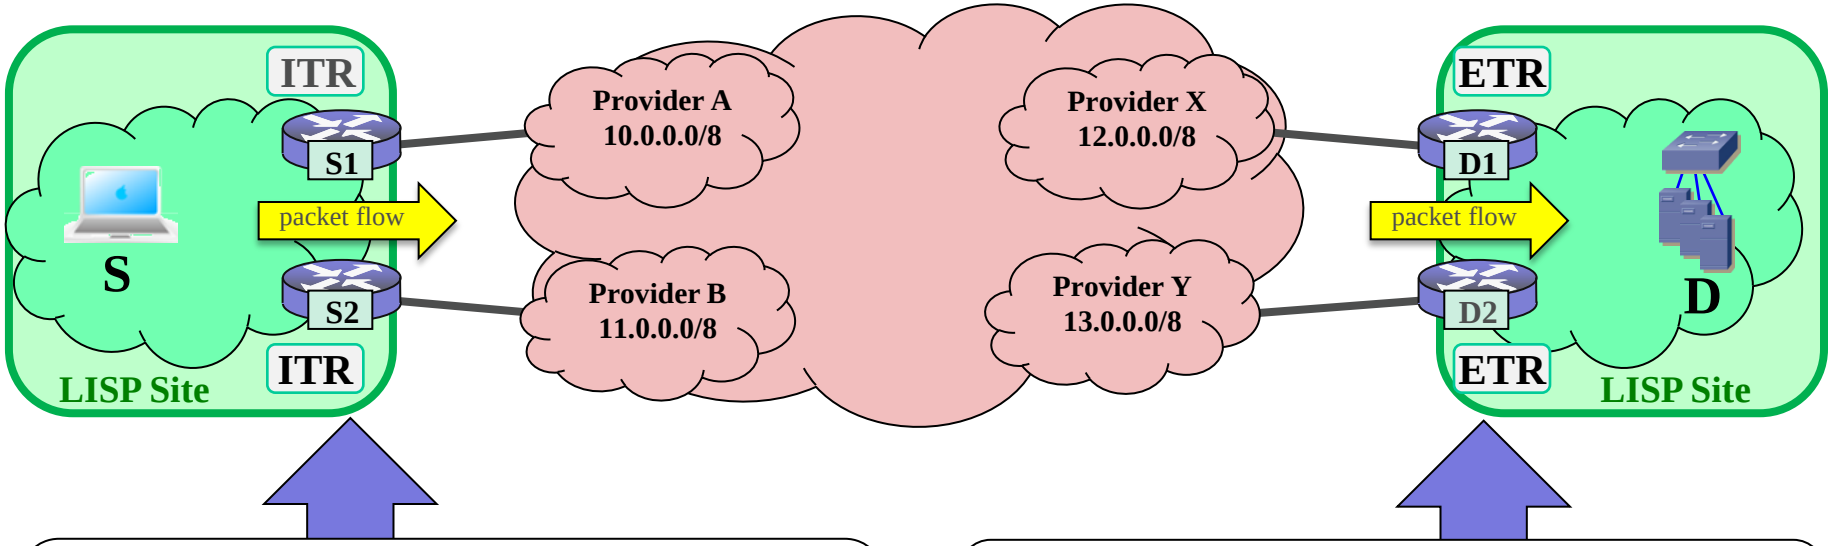

ITR – Ingress Tunnel Router

- Receives packets from site-facing interfaces
- Encap to remote LISP sites, or native-fwd to non-LISP sites

ETR – Egress Tunnel Router

- Receives packets from core-facing interfaces
- De-cap, deliver packets to local **EIDs** at site

LISP Overview
Data-Plane Flow

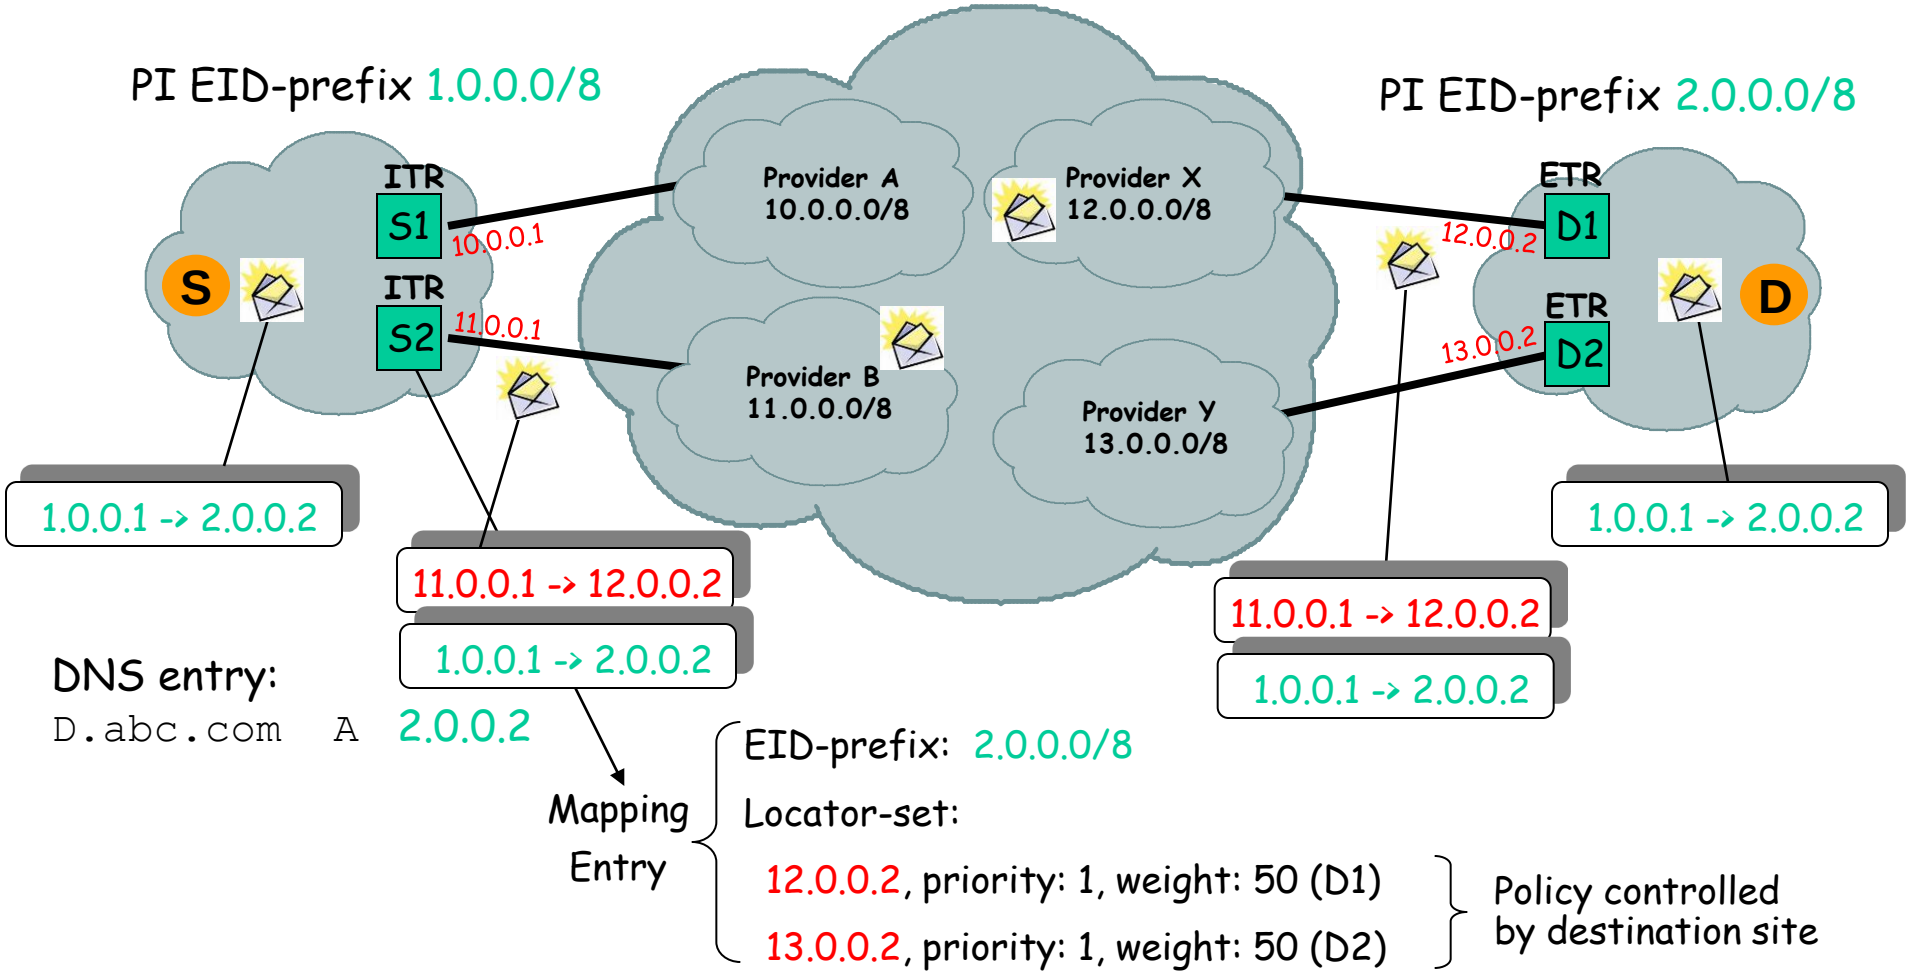

LISP Overview
LISP Mapping Resolution – DNS analog

LISP "Level of Indirection" is analogous to a DNS lookup

- DNS resolves **IP addresses** for **URLs**

- LISP resolves **locators** for queried **identities**

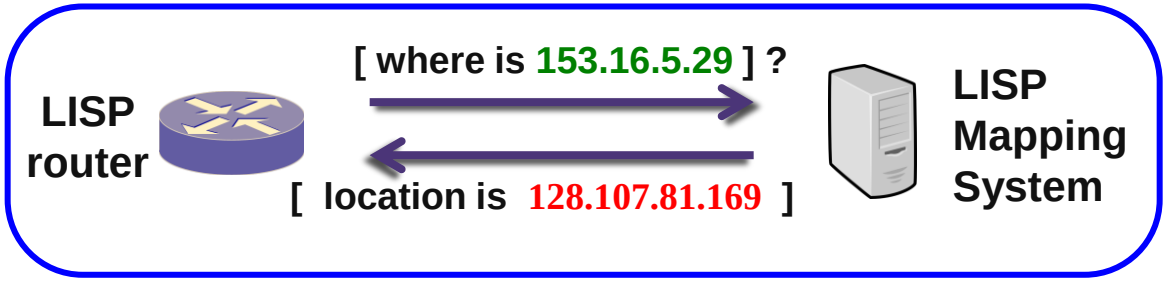

LISP Control Plane Overview
Control Plane Messages

Control Plane EID Registration

Map-Register message

Sent by ETR to Map-Server to register its associated EID prefixes

Specifies the RLOC(s) to be used by the Map-Server when forwarding Map-Requests to the ETR

Control Plane "Data-triggered" mapping service

Map-Request message

Sent by an ITR when it needs for EID/RLOC mapping, to test an RLOC for reachability, or to refresh a mapping before TTL expiration

Map-Reply message

Sent by an ETR in response to a valid map-request to provide the EID/RLOC mapping and site ingress Policy for the requested EID

LISP Overview
Mapping Database System

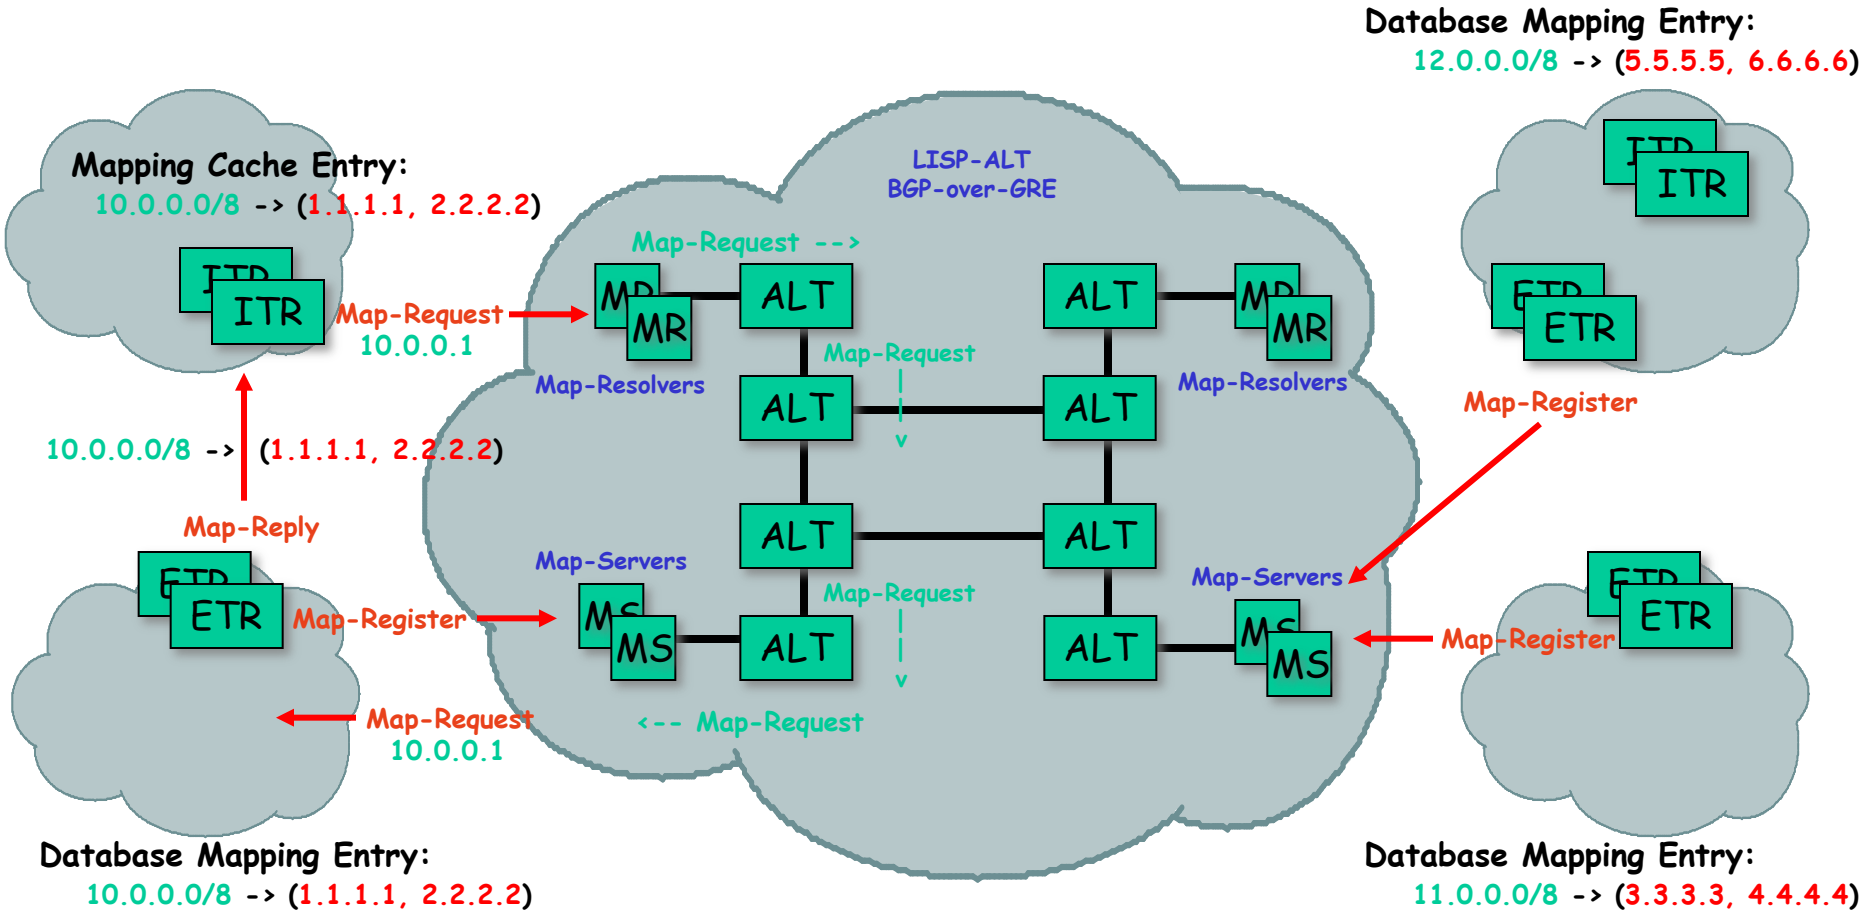

LISP Control Plane Overview
Map-Registration example

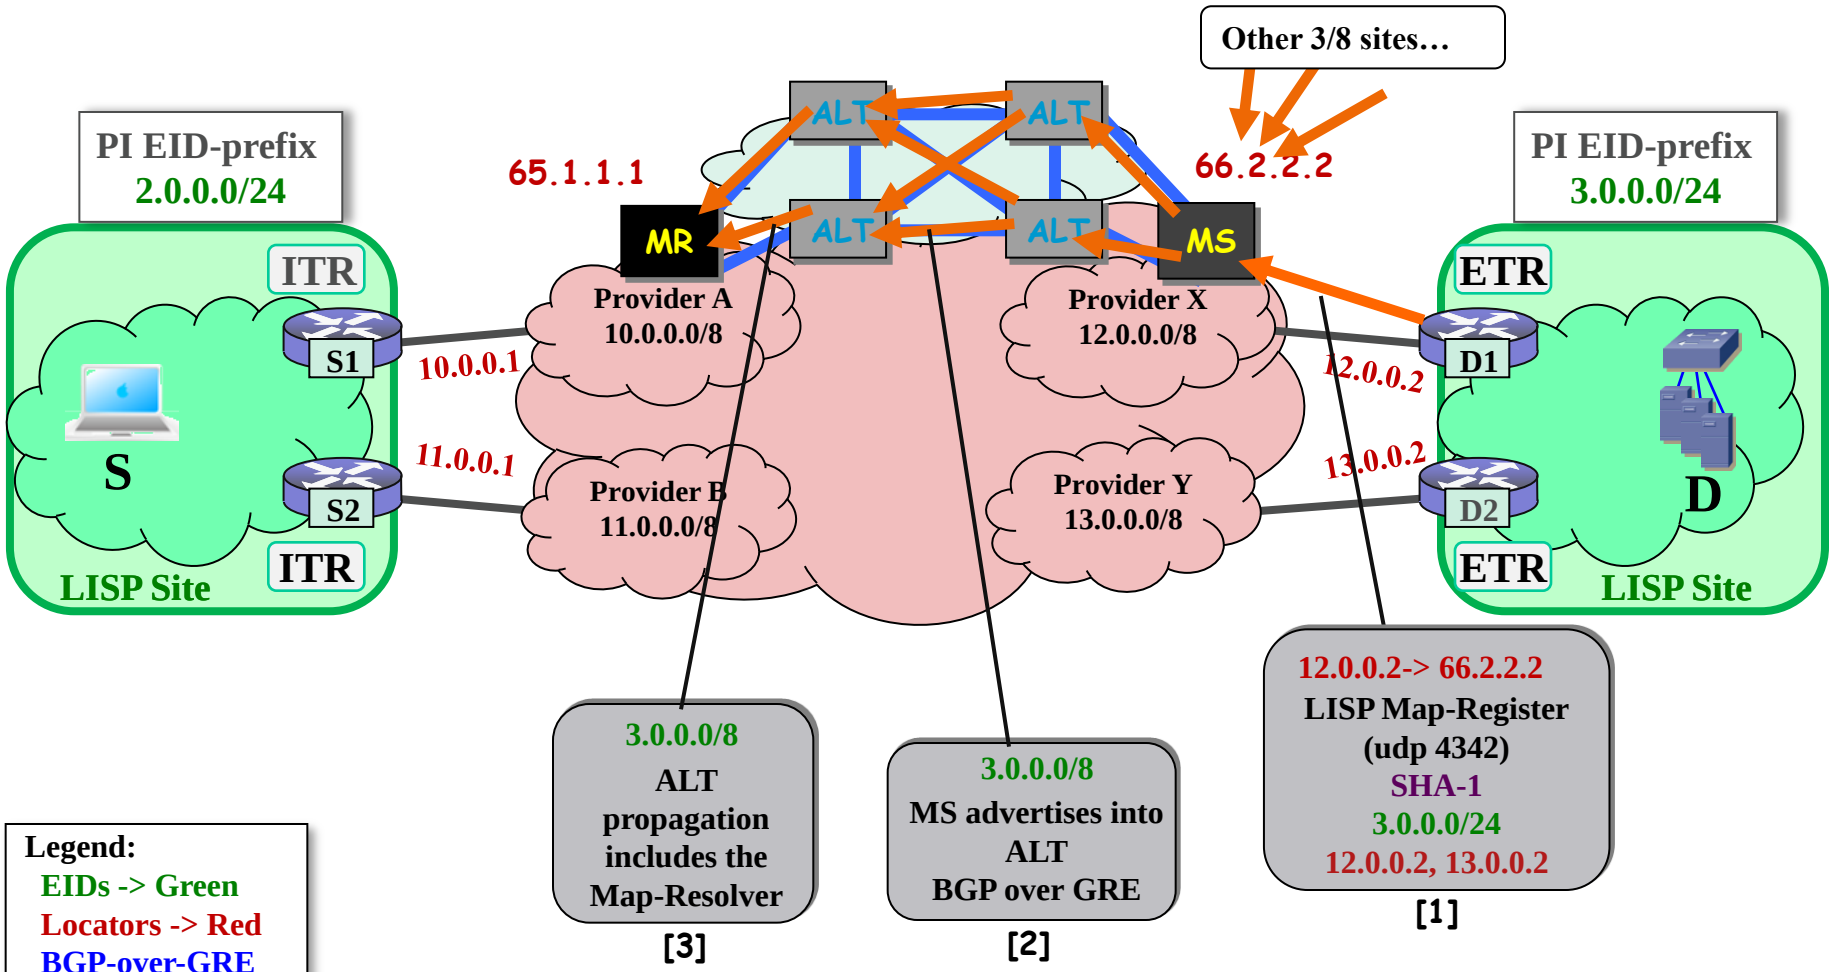

LISP Control Plane Overview
Map-Request example

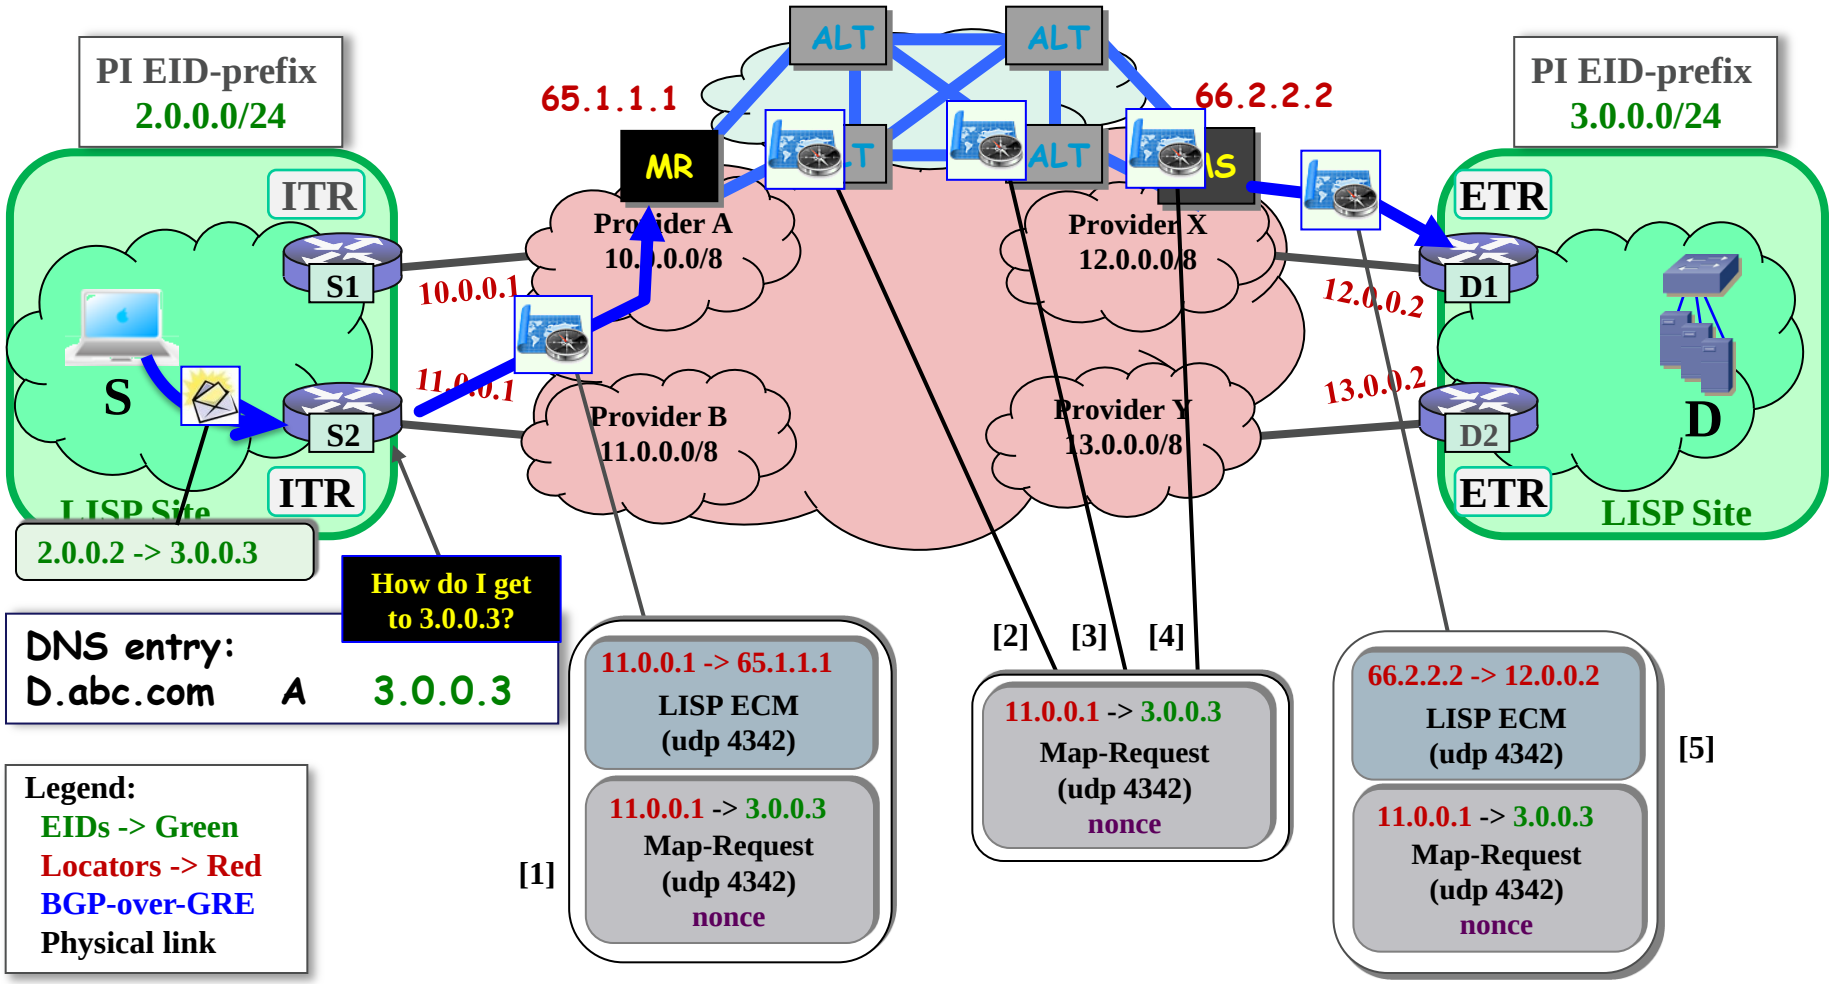

LISP Control Plane Overview
Map-Reply example

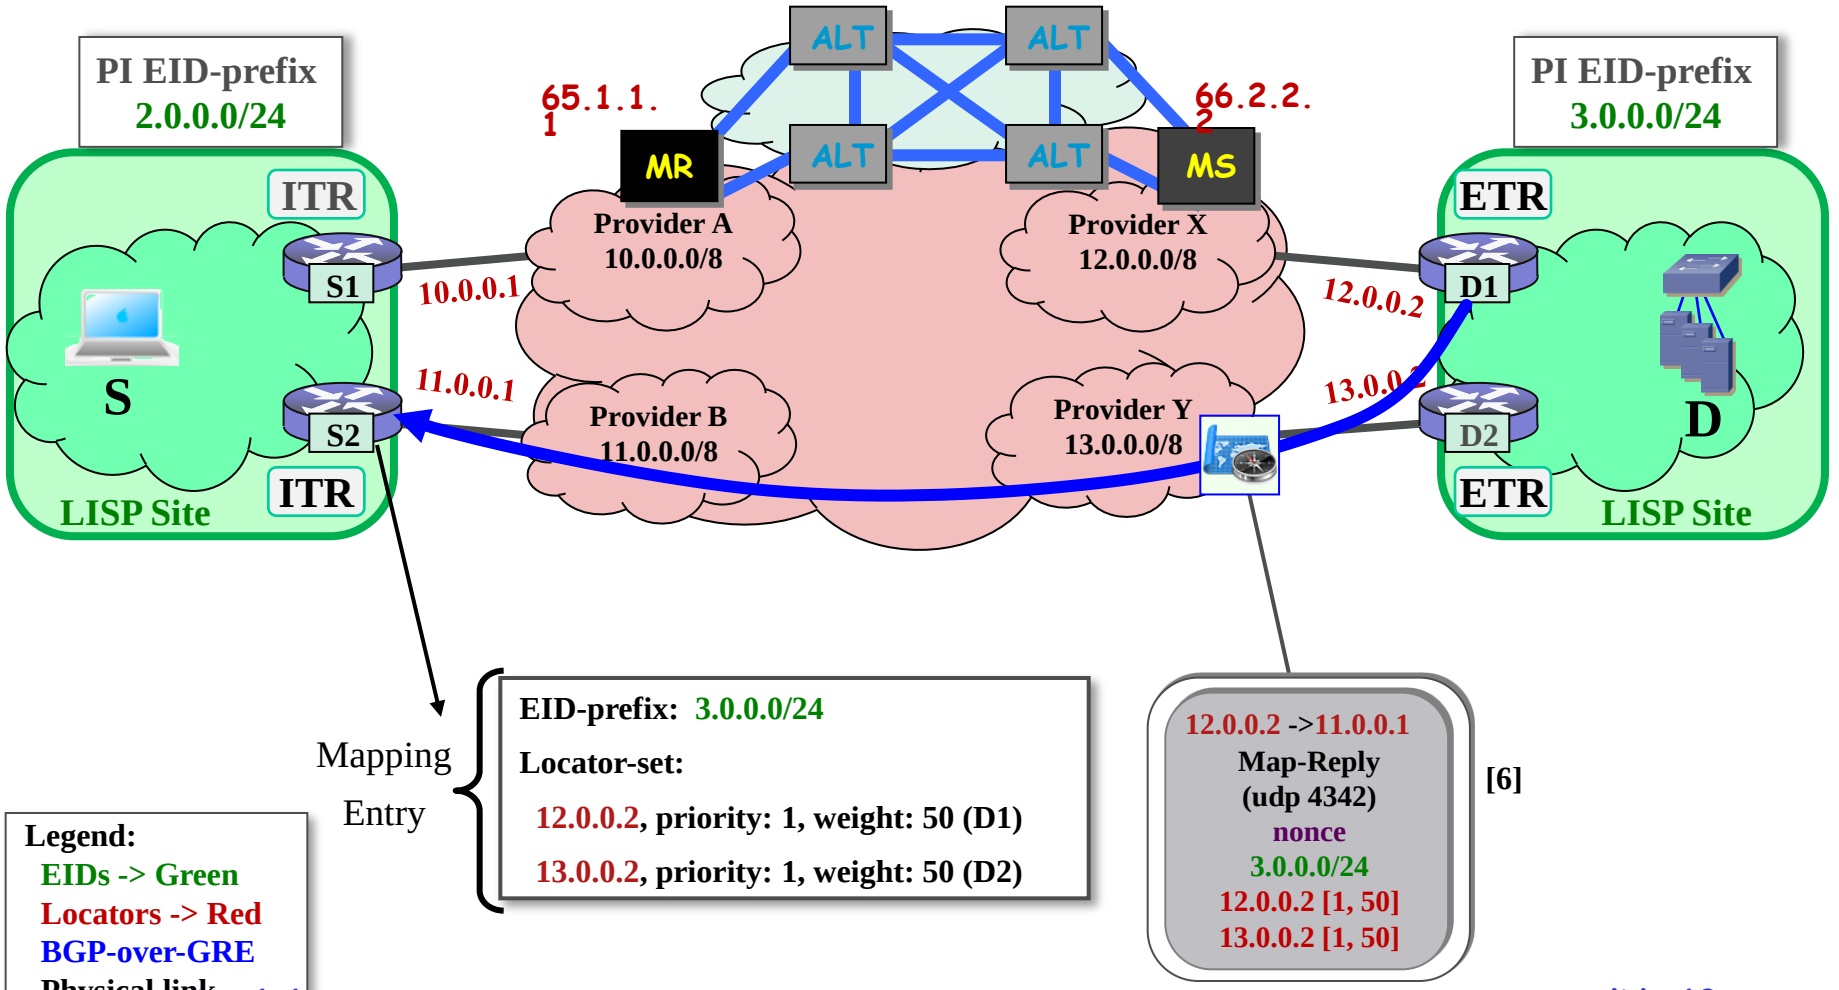

LISP Overview
Mapping Database Modularity

Legend:
- LISP Sites -> green
- 1st layer access infrastructure -> blue
- 2nd layer core infrastructure -> red

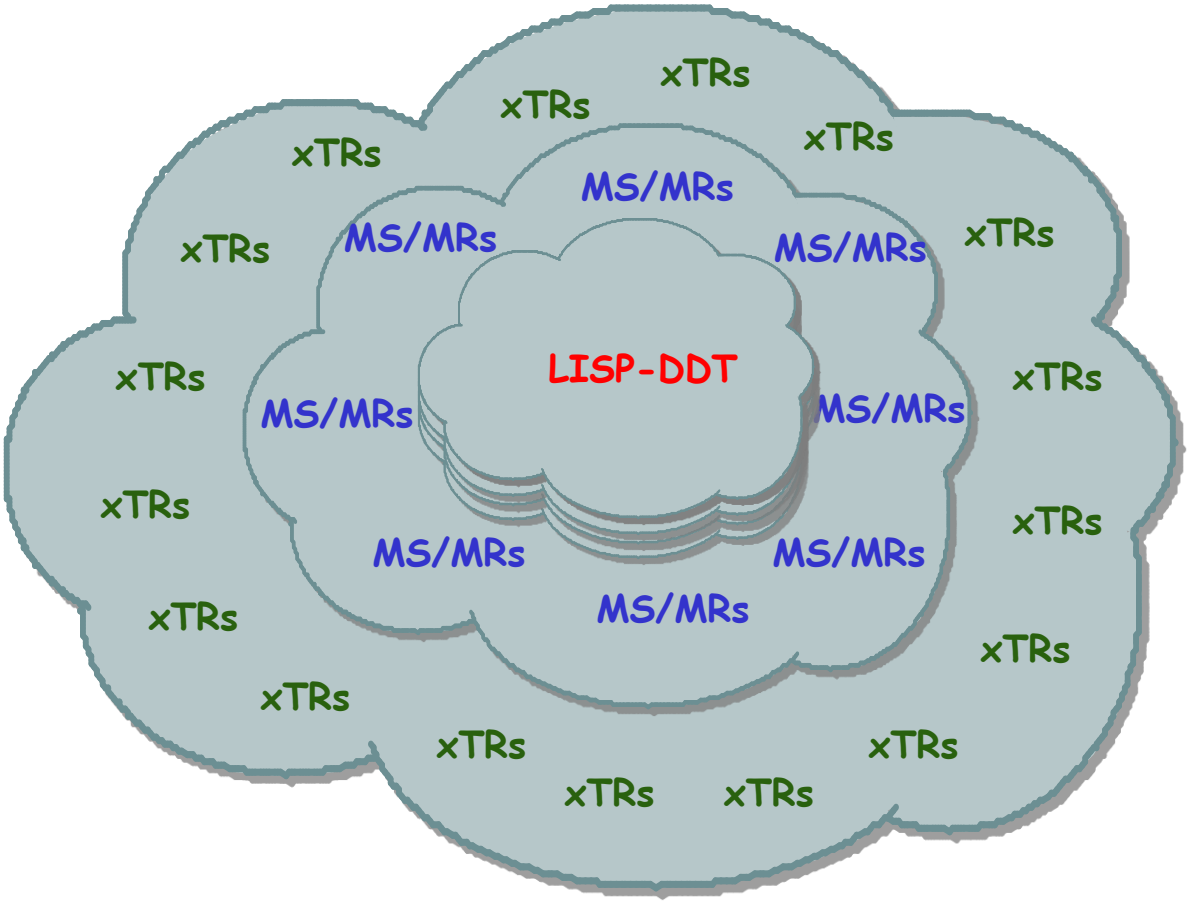

LISP Team Philosophy

- cisco will ship 4 LISP-related releases per year
  - We have been doing this since Dec 2009
  - Engineering releases from LISP team
  - Mainline releases from Bus
- Engineering releases
  - More features – development support
  - Support level scales less so
- Mainline releases
  - Less features – TAC support
  - Support level scales better

Mainline Products

- Shipped summer of 2011
- ISR(s) and ASR 1K
  - IOS 15.1.4M
  - IOS-XE 3.3.0S
- Nexus 7K
  - NX-OS 5.2(1) (mainline delhi)

LISP Engineering-Team Products

- ISR(s) and ASR 1K
  - IOS 15.1(4)XB4
  - IOS-XE 2.5.1XC
- Nexus 7K
  - NX-OS dino-lisp-r4 (delhi branch)
- Titanium and UCS c200 bare metal
  - NX-OS dino-lisp-r4 (delhi branch)
- Titanium and UCS c200 virtualized
  - NX-OS dino-lisp-r4 with VMware and KVM

LISP Shipping Use-Cases

(1) Low OpEx site multi-homing
(2) IPv6 coexistence
(3) VM-mobility (includes cloud mobility)
(4) Multi-tenant VPNs
(5) Mapping System support for LISP-MN

LISP Use Cases
Efficient Multi-Homing

Needs:

- Site connectivity to multiple providers
- Low OpEx/CapEx

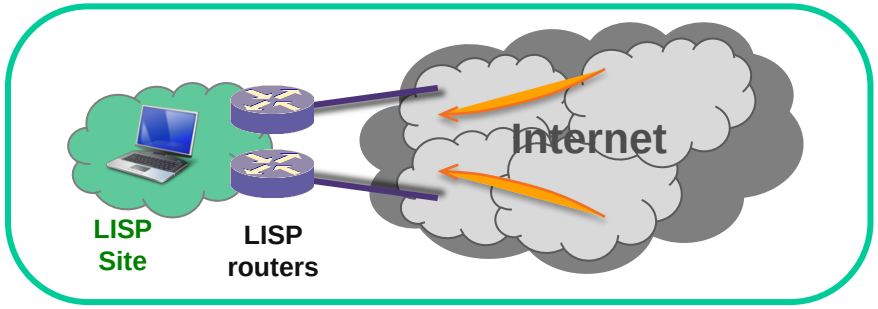

LISP Solution:

- LISP provides a streamlined solution for handling multi-provider connectivity and policy without BGP complexity

Benefits:

- Multi-homing across different providers
- Simple policy management
- Ingress Traffic Engineering
- Egress Traffic Engineering

LISP Use Cases
IPv6 Migration Support

Needs:

- Rapid IPv6 Deployment
- Minimal Infrastructure disruption

LISP Solution:

- LISP encapsulation is Address Family agnostic
  - IPv6 interconnected over IPv4 core
  - IPv4 interconnected over IPv6 core

Benefits:

- Accelerated IPv6 adoption
- Minimal added configurations
- No core network changes
- Can be used as a transitional or permanent solution

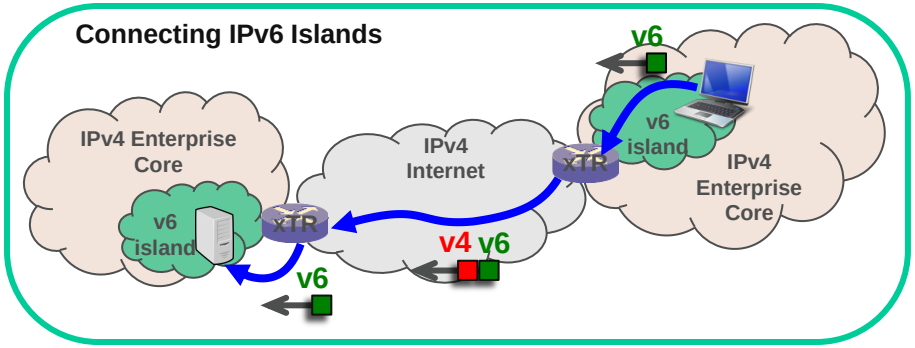

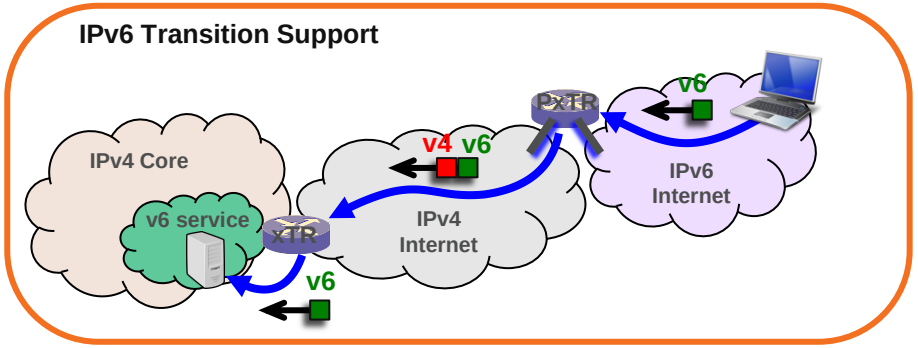

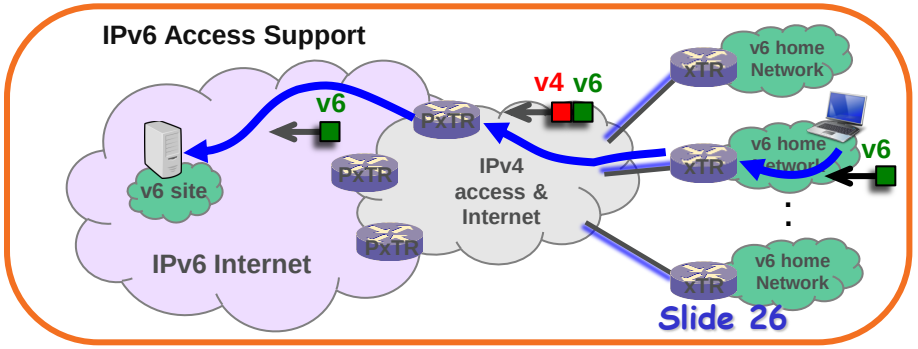

LISP Use Cases
Mobile-Node Mobility

This phone is a LISP site!

EID-prefix: 2610:00d0:xxxx::1/128

Map-Server: 64.1.1.1

64.0.0.1 wifi

65.0.0.1 3G

(1) 2 MNs can roam and stay connected
(2) MNs can be servers
(3) MNs roam without changing DNS entries
(4) MNs can use multiple interfaces
(5) MNs can control ingress packet policy
(6) Faster hand-offs
(7) Low battery use by MS proxy-replying
(8) And most importantly, packets have stretch of 1 so latency is best for delay sensitive applications

LISP-MN can scale to 1 billion hand-sets!

LISP Use Cases
VPNs and Segmentation

Needs:

- Highly-scalable VPNs supporting IPv4 and IPv6
- Remove IGP scaling limitations for Branch WAN aggregation

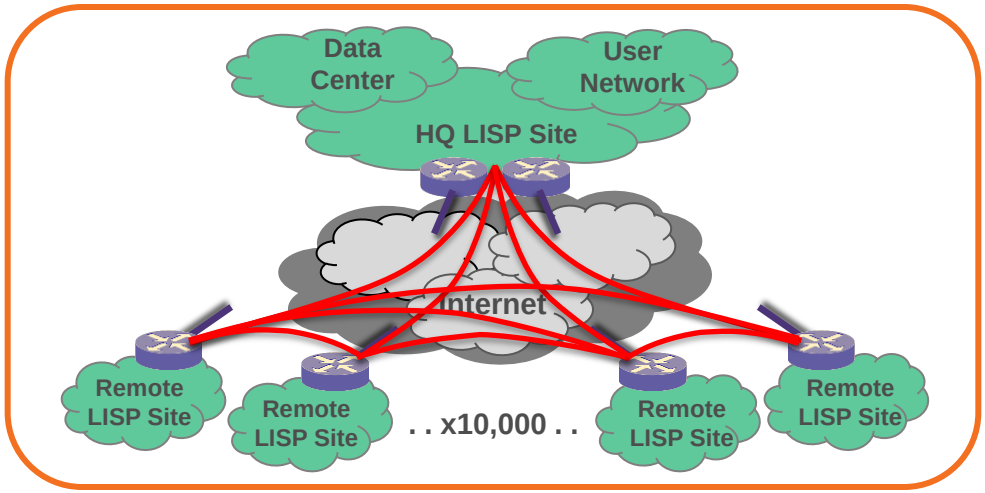

LISP Solution:

- LISP Instance-IDs for Over-the-Top VPNs
- Supports complex topologies including multi-homed branches, partial mesh, etc.
- IPv4/IPv6 co-existence

Benefits:

- No CE-PE coordination required
- LISP Mapping System supports high scalability
- Simplified Management

LISP Use Cases

Separation of Address Spaces

Shared Infrastructure for Scaling, Private VPN sites for Segmentation

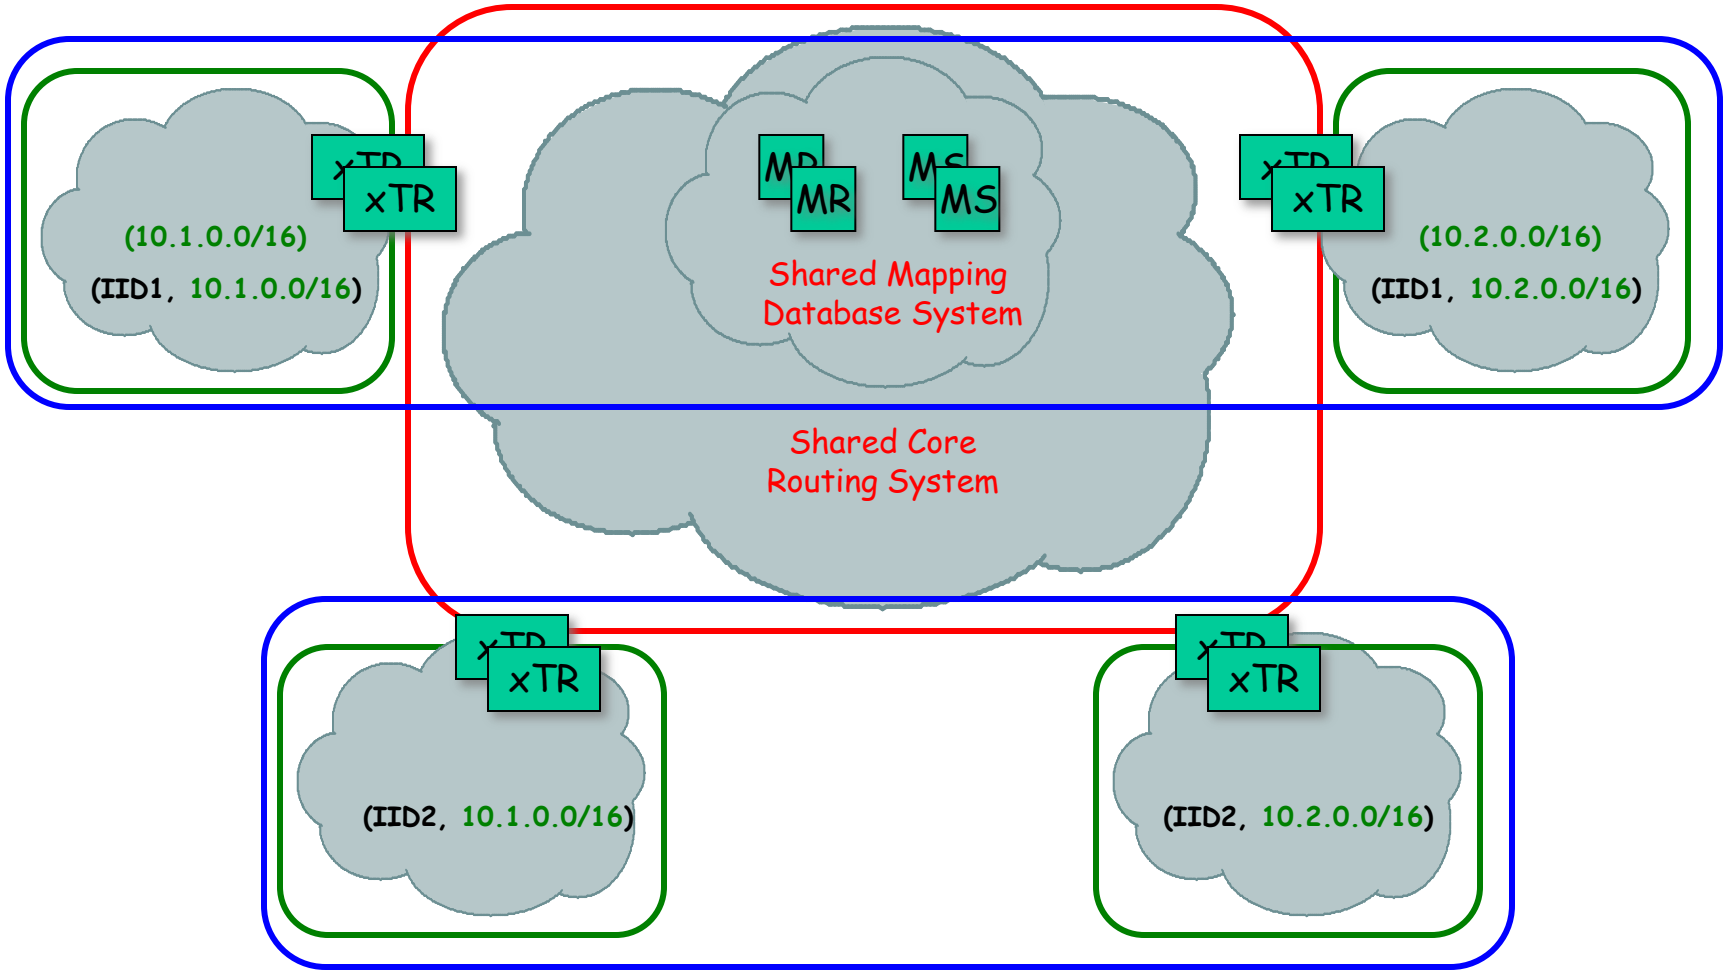

LISP Use Cases

Roaming across Subnets

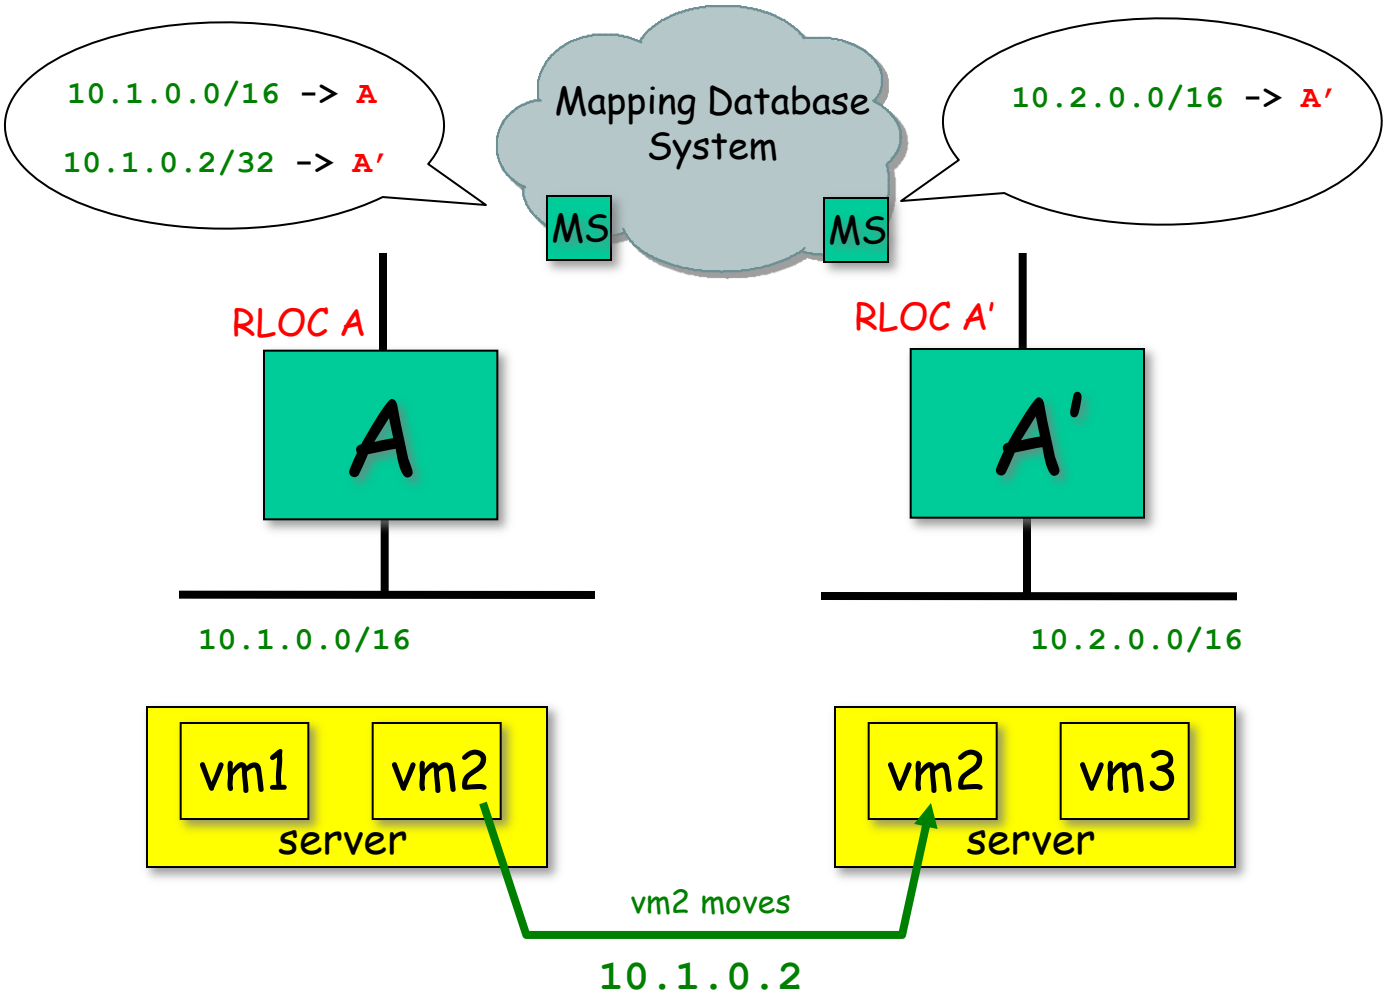

LISP Use Cases
VM-Mobility

Needs:

- VM-Mobility **across subnets**
- Move detection, dynamic EID-to-RLOC mappings, traffic redirection

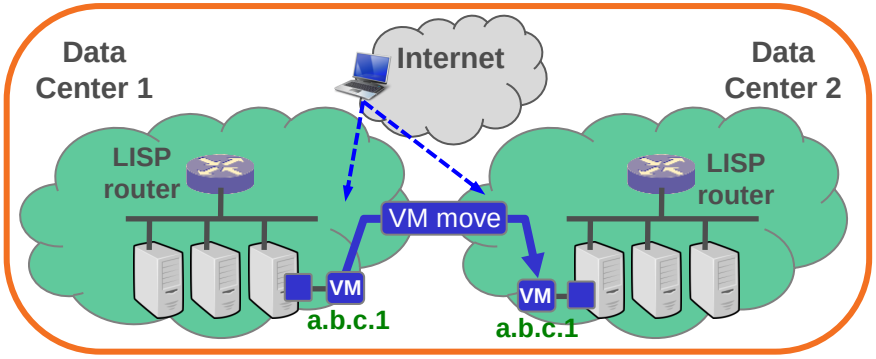

LISP Solution:

- xTR Dynamic EID (VM-Mobility) on access or aggregation switches

Benefits:

- Integrated Mobility
- Direct Data Path (no triangulation)
- Connections maintained across moves
- No routing re-convergence
- Transparent to hosts
- No DNS updates required
- Global Scalability (cloud bursting)
- IPv4/IPv6 Support
- ARP elimination
- Automated move detection

LISP Use Cases

LISP Moves Virtual Machines

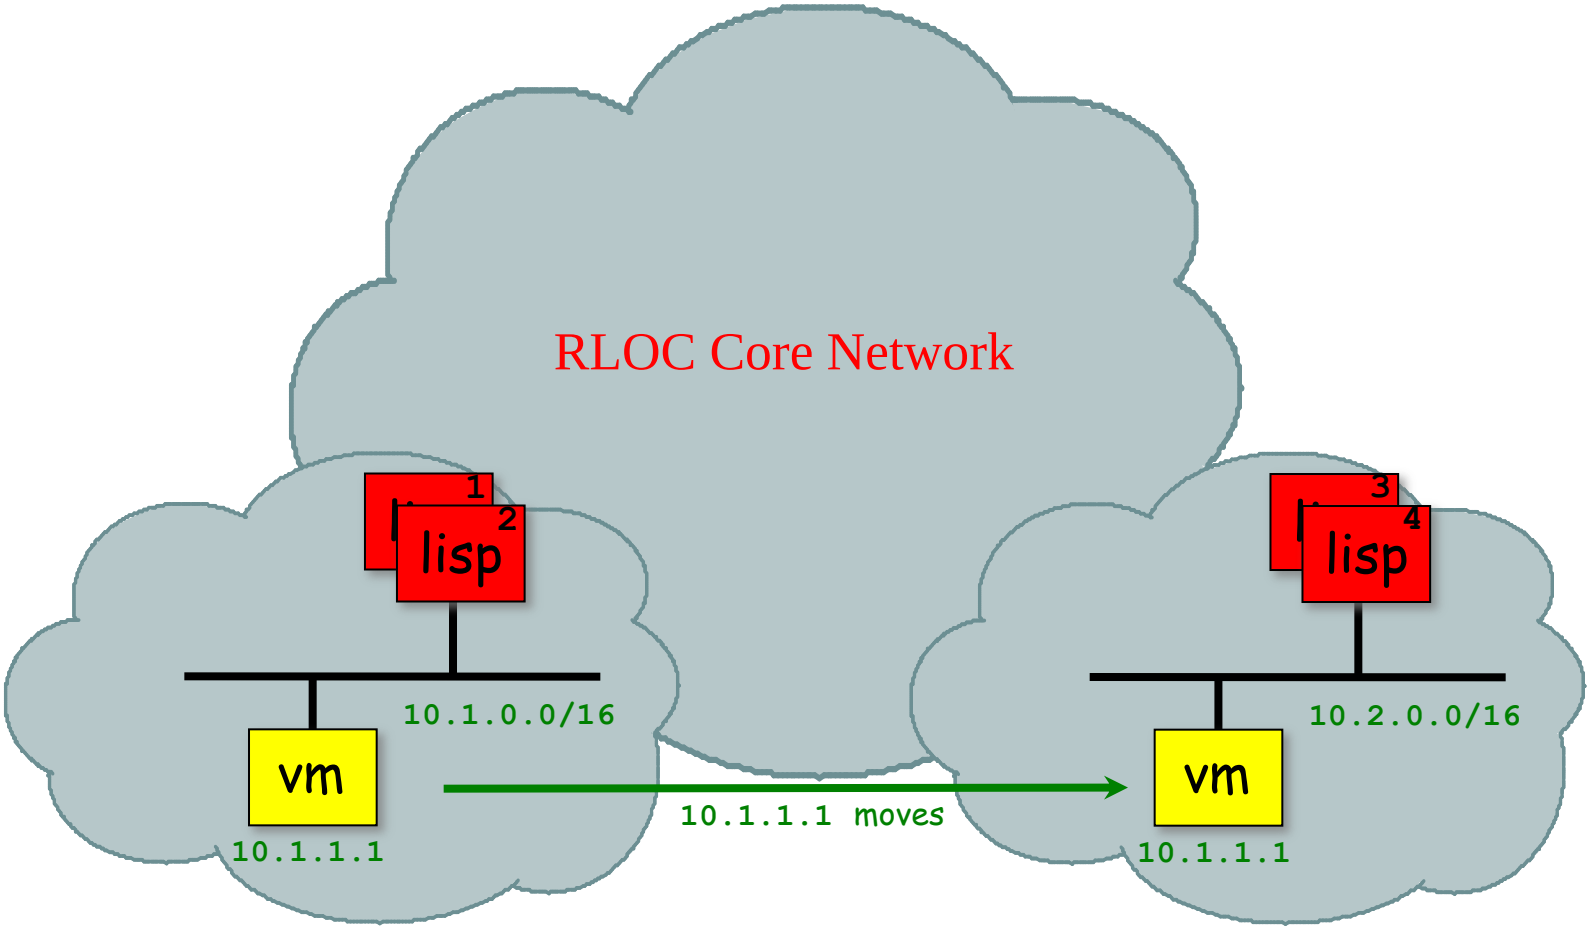

LISP Use Cases

LISP Moves Topologies

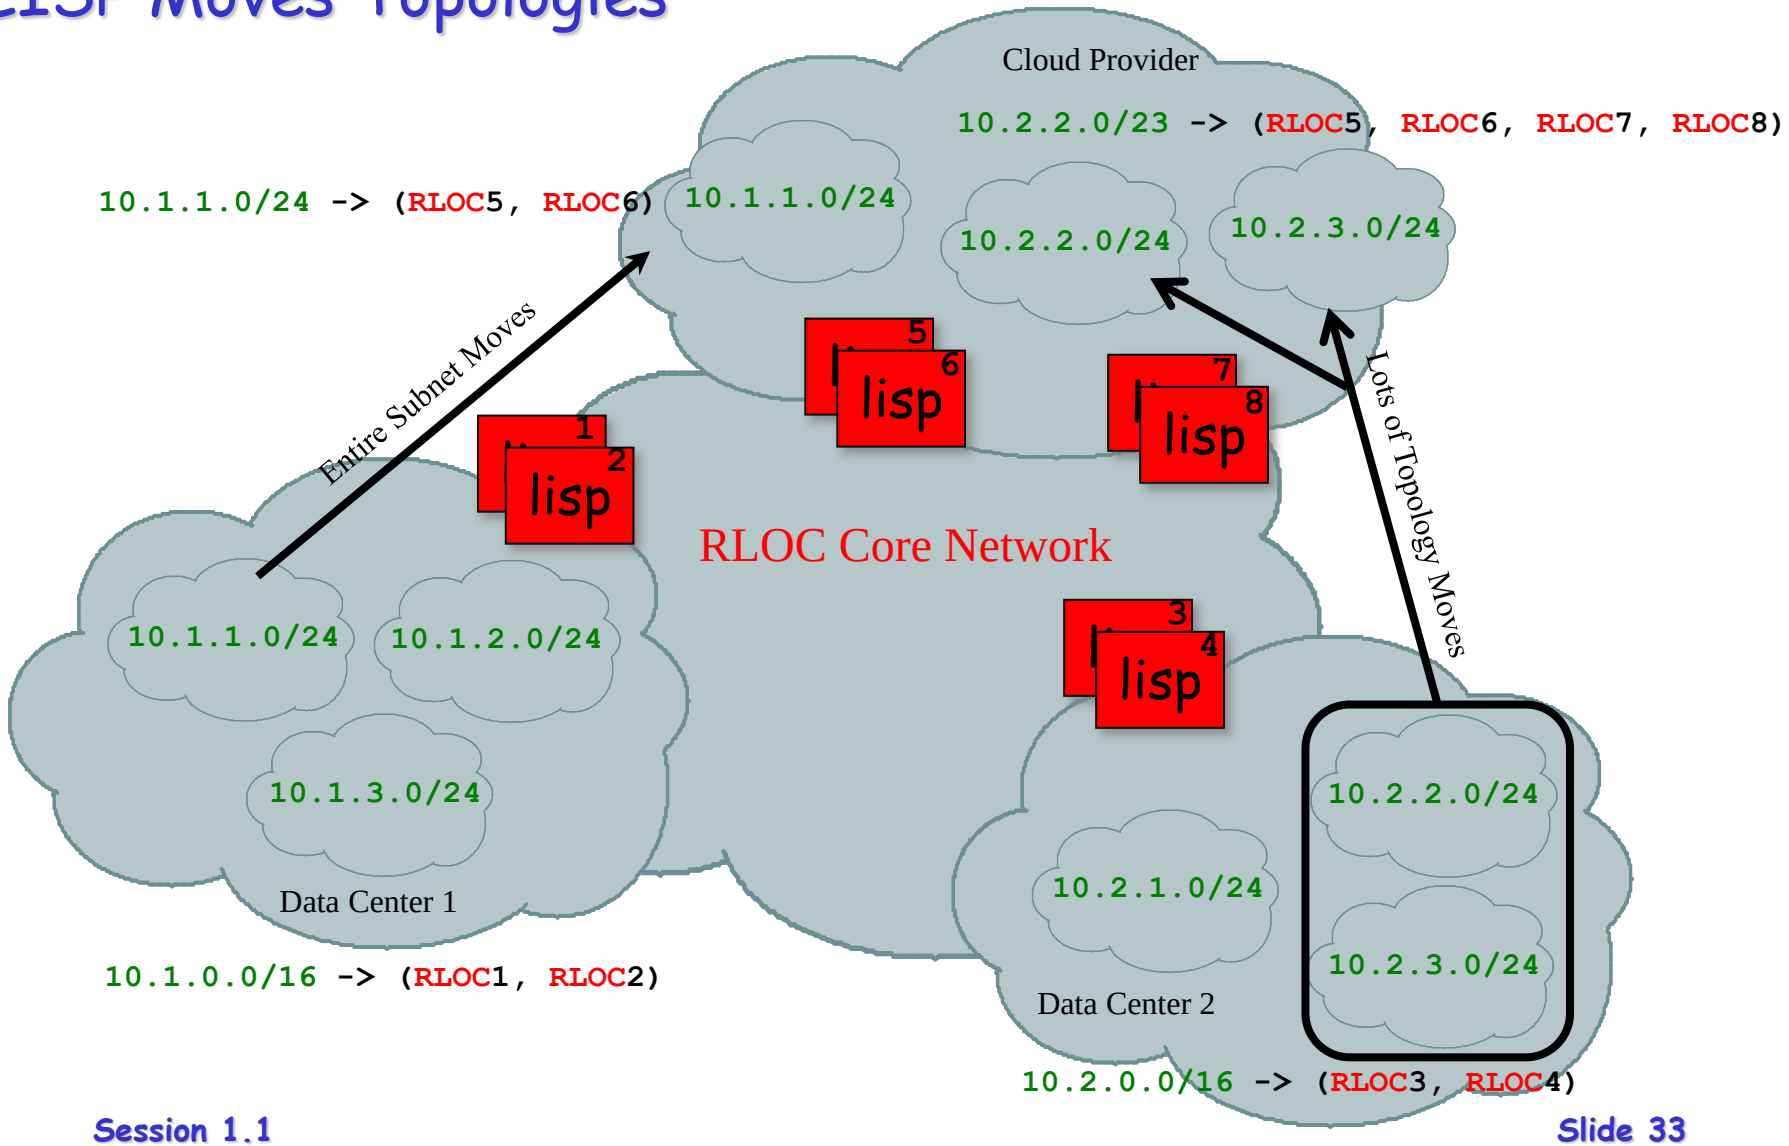

What's Next for Products

- NX-OS 5.2(1) LISP Control-Plane with Platform specific Data-Plane
  - CRS/1 and CRS/3
  - ASR 9K
- IOS 15.1 LISP Control-Plane with Platform specific Data-Plane
  - Cat 6K
- Cat 3K and Cat4K Future
- LISP Mapping Database Services
  - Nexus 1010 Virtual Appliance (NX-OS in a VM)

What's Next for LISP Features

- LISP-SEC and LISP-dSEC
- LISP Nat-Traversal Support (NTR)
- More LISP-ODF Features
- New LISP Mapping Databases
  - Researching ISIS-ALT, LISP-DDT and LISP-LDML
- LISP-Multicast
- Additional LCAF Support
  - LISP-TE
  - Geo Coordinates for RLOC addresses
  - PETR rough nonce security check
- New EID Support (i.e. MAC addresses)
  - Layer-2 LISP

Any other Products

- Well yes, not sellable – but public domain
- LISP-MN Implementation
  - Android (Froyo and Gingerbread)
    - Nexus 1 and Nexus S HTC phones
    - Demoed to Google and Vint Cerf
    - Demoed to AT&T & OnStar
    - Working on port to cisco Cius
  - Linux
    - Open source now on github.com (LISPmob)
    - Working on Linksys OpenWRT release

What did we miss?

- Other platforms?
- Other LISP features?

LISP is here now!

- With real implementation experience!
- With real deployment experience!
- With real customer engagement!
- Has been that way for several years
- http://www.lisp{4,6}.net
- http://lisp.cisco.com
- lisp@ietf.org
- lisp-support@cisco.com

LISP References
Resources

LISP Information

- IETF LISP WG — http://tools.ietf.org/wg/lisp/
- LISP Beta Network — http://www.lisp4.net or http://www.lisp6.net
- Cisco LISP Site — http://lisp.cisco.com (v4 and v6)
- Cisco LISP Marketing — http://www.cisco.com/go/lisp

Mailing Lists

- IETF LISP WG — lisp@ietf.org
- LISP Interest (public) — lisp-interest@puck.nether.net
- Cisco LISP Questions — lisp-support@cisco.com

LISP Deployment Experience Executive Panel

Thursday, July 14th, 8:00 am, Mandalay Bay C
Session ID: **PNLRST-4000**

Want a frank technical discussion with a Cisco Fellow and Cisco CTO? Dino Farinacci, Cisco Fellow will give an introduction to LISP, a next-generation routing architecture, that is in real cisco product's and deployed by a distinguished set of customers. Padmasree Warrior, SVP and CTO, will share with you Cisco's commitment to LISP and how the technology will change Internet Routing and Roaming. Dino and Padma will be joined by a distinguished group of Internet architects and customers, where they will share their experience deploying LISP. Finally, we will have an open-mic Q&A with Dino, Padma, and the panel members. This intimate, hands-down no nonsense interactive exchange with top industry experts is a must see. This session was a great success at CiscoLive London. There is no telling what will be uncovered during this session.

Hosts:

- Padmasree Warrior, Chief Technology Officer, Cisco
- Dino Farinacci, Cisco Fellow

Panelists:

- John Manville, VP Information Technology, Cisco
- Donn Lee, Senior Network Engineer, Facebook
- Danny McPherson, Chief Security Officer, Verisign
- Parantap Lahiri, Director Network Architecture, Microsoft
- Hwa-Jung Han, Manager, Verizon
- John Feurerherd, Data Center Architect, Wells Fargo

http://www.ustream.tv/recorded/15990101

LISP – A Routing Architecture; Not a Feature
LISP Innovations

LISP enables **IP Number Portability**

- With session survivability
- Never change host IP addresses
  No renumbering costs
- No DNS "name -> EID" binding change

LISP uses **pull** vs. **push** routing

- OSPF and BGP are **push** models;
  routing stored in the forwarding plane
- LISP is a **pull** model; Analogous to
  DNS; massively scalable

LISP is an **over-the-top** technology

- Address Family agnostic
- Incrementally deployable
- No changes in end systems

LISP creates a **Level of Indirection**

- Separates End-Host and Site addresses

LISP deployment **simplicity**

- No host changes
- Minimal CPE changes
- Some new core infrastructure components

LISP enables other **interesting features**

- Simplified multi-homing with Ingress traffic
  engineering – without the need for BGP
- End-host mobility without renumbering
- Address Family agnostic support

LISP is an **Open Standard**

- No Cisco Intellectual Property Rights

Complete Your Session Evaluation

- Please give us your feedback!!

  Complete the evaluation form you were given when you entered the room
- This is session 1.1

Don’t forget to complete the overall event evaluation form included in your registration kit

YOUR FEEDBACK IS VERY IMPORTANT FOR US!!! THANKS


CISCO

LISP - Data Format Example
IPv4 EID/IPv4 RLOC Example

draft-ietf-lisp-12

IPv4 Outer Header: Router supplies RLOCs

UDP

LISP header

IPv4 Inner Header: Host supplies EIDs

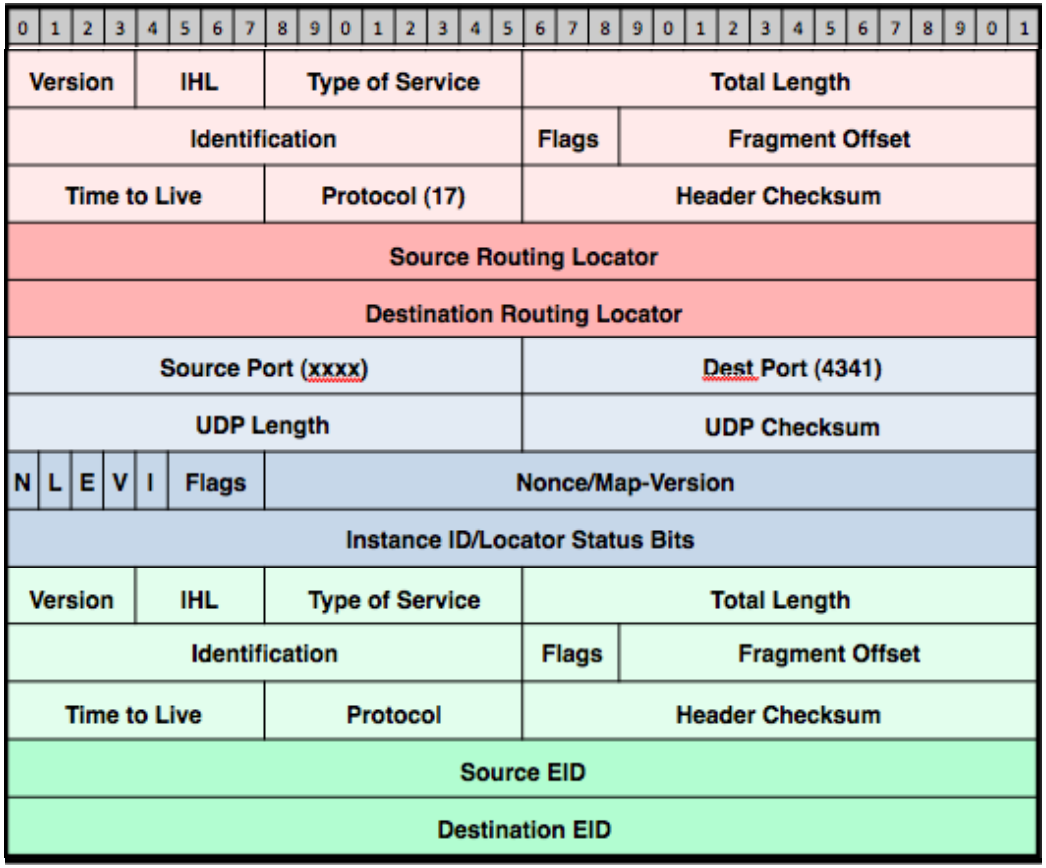

LISP - Data Format Example

All Combinations - IPv4 and IPv6 Supported

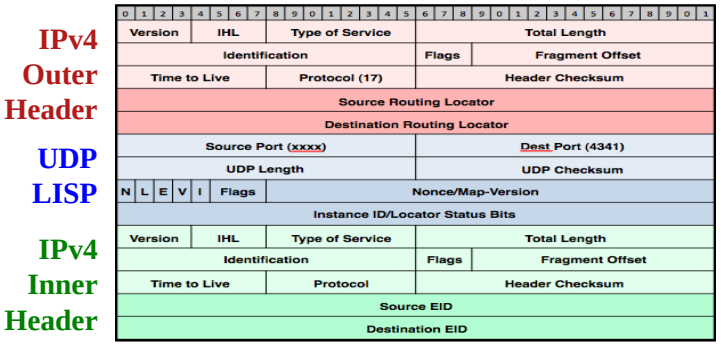

IPv4/IPv4

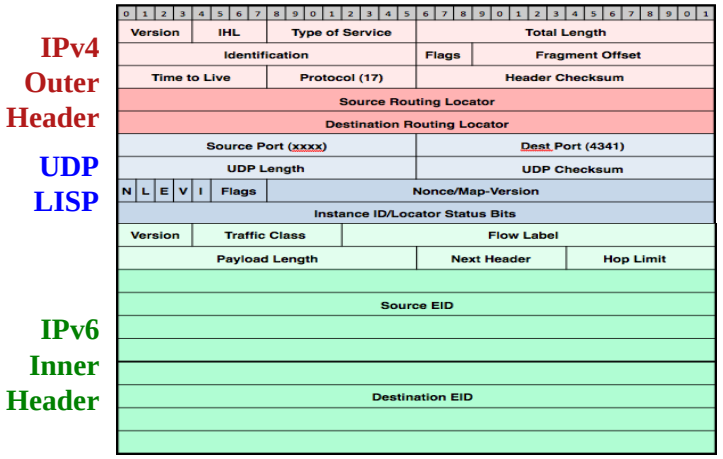

IPv4/IPv6

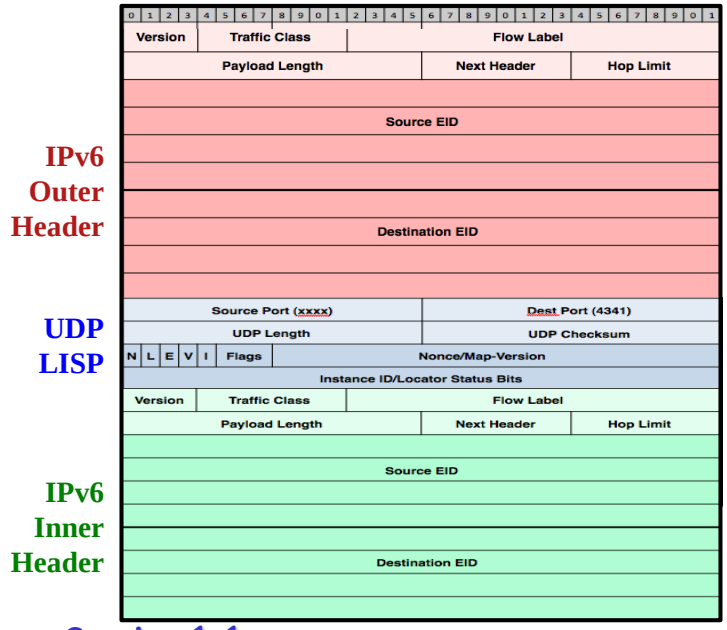

IPv6/IPv6

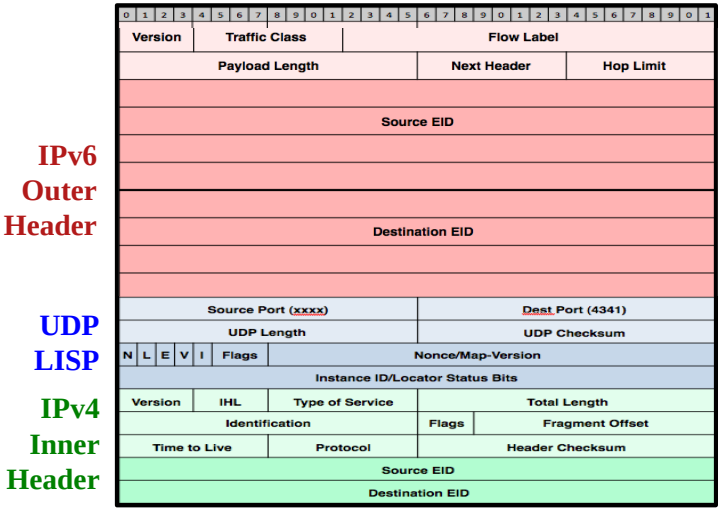

IPv6/IPv4