
TOMORROW starts here.
cisco
Cisco live!

LISP – A Next-Generation Networking Architecture

BRKRST-3045

Victor Moreno

Distinguished Engineer

Cisco live!

LISP - A Next Generation Routing Architecture

Agenda

- LISP Overview
- LISP Operations
- LISP Deployment Examples
- LISP Status
- LISP Summary
- LISP References

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

Locator/ID Split and LISP

Routing and Addressing Architecture of the Internet Protocol

- Addresses today combine location and identity semantics in a single 32-bit or 128-bit number
- Separating Location and Identity changes this…
  - Provide a clear separation ***at the Network Layer*** between ***what we are looking for*** vs. ***how best to get there***
  - ***Translation vs. Tunnelling is a key question***
- Network Layer Identifier: WHO you are in the network
  - long-term binding to the thing that they name, does not change often at all
- Network Layer Locator: WHERE you are in the network
  - Think of the source and destination "addresses" used in routing and forwarding
- WHERE you are can change! WHO you are should be the same!

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

Original Motivation…

- An IP address “overloads” location and identity
  - Today… “addressing follows topology”
  - Efficient aggregation is only available for Provider Assigned (PA) addresses
  - Ingress Traffic Engineering usually requires Provider Independent (PI) addresses and the injection of “more specifics” :: this limits route aggregation compactness
  - IPv6 does not fix this
- Route scaling issues drive system costs higher
  - Forwarding plane (FIB) requires expensive memory
  - Route scaling “drivers” are also seen in Data Centres and for Mobility :: not just the Internet DFZ

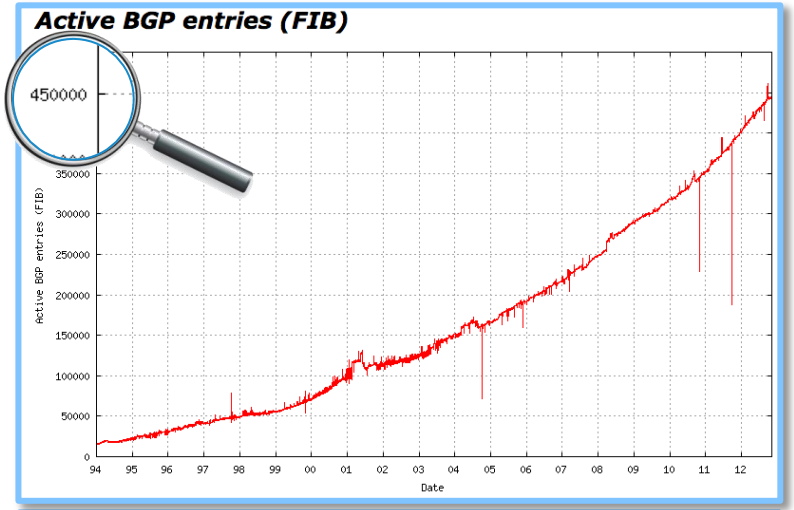

“… routing scalability is the most important problem facing the Internet today and must be solved … ”

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

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

Identity and Location :: an Overloaded Concept in Routing Today…

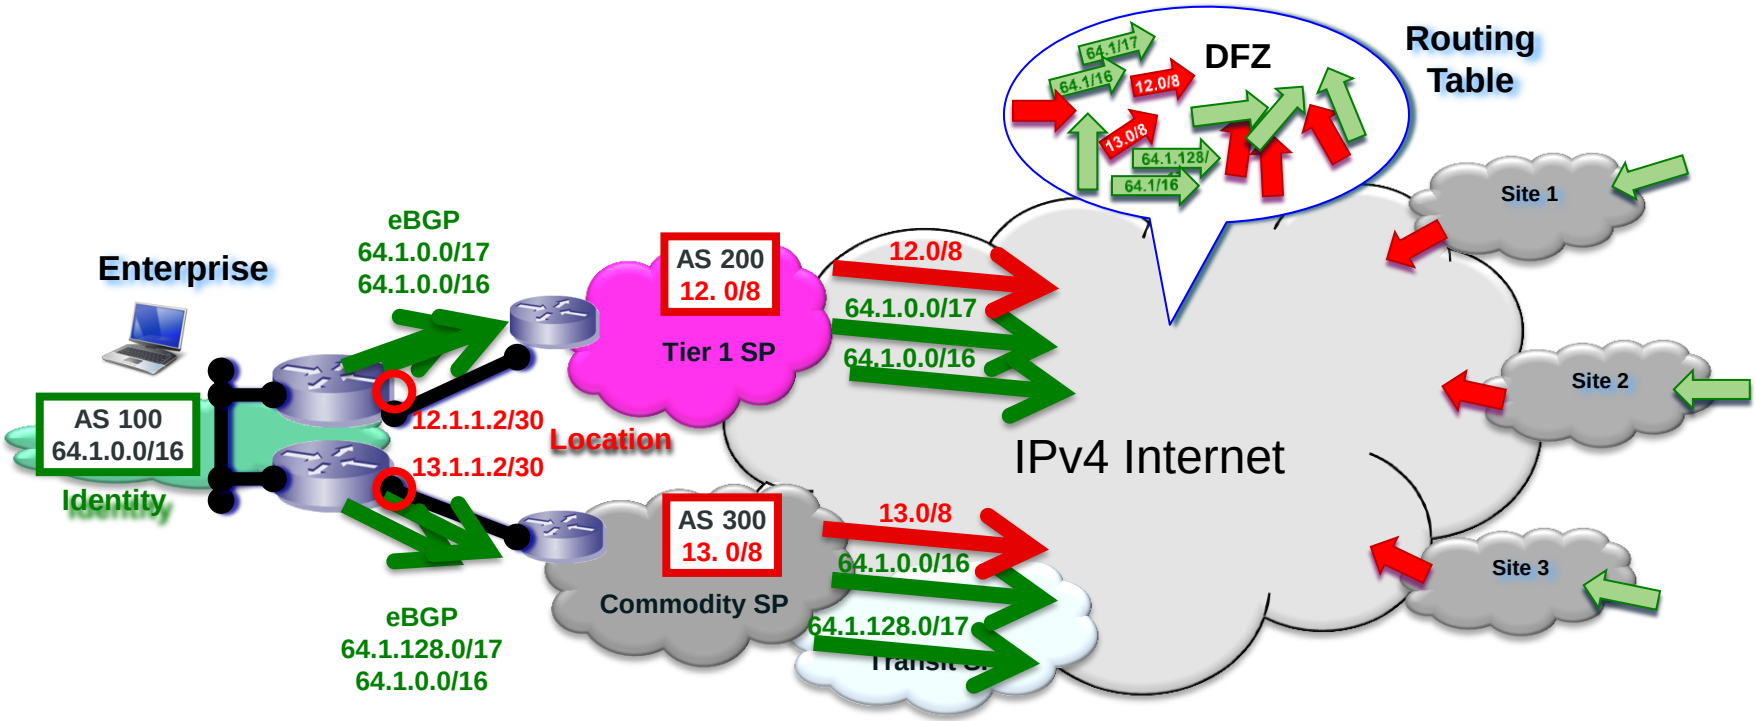

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

Identity and Location :: an Overloaded Concept in Routing Today…

- Let's put **ID address** and **Locator address** in different databases
- Let's create a "**level of indirection**" between **ID** and **LOCATION** in the network!

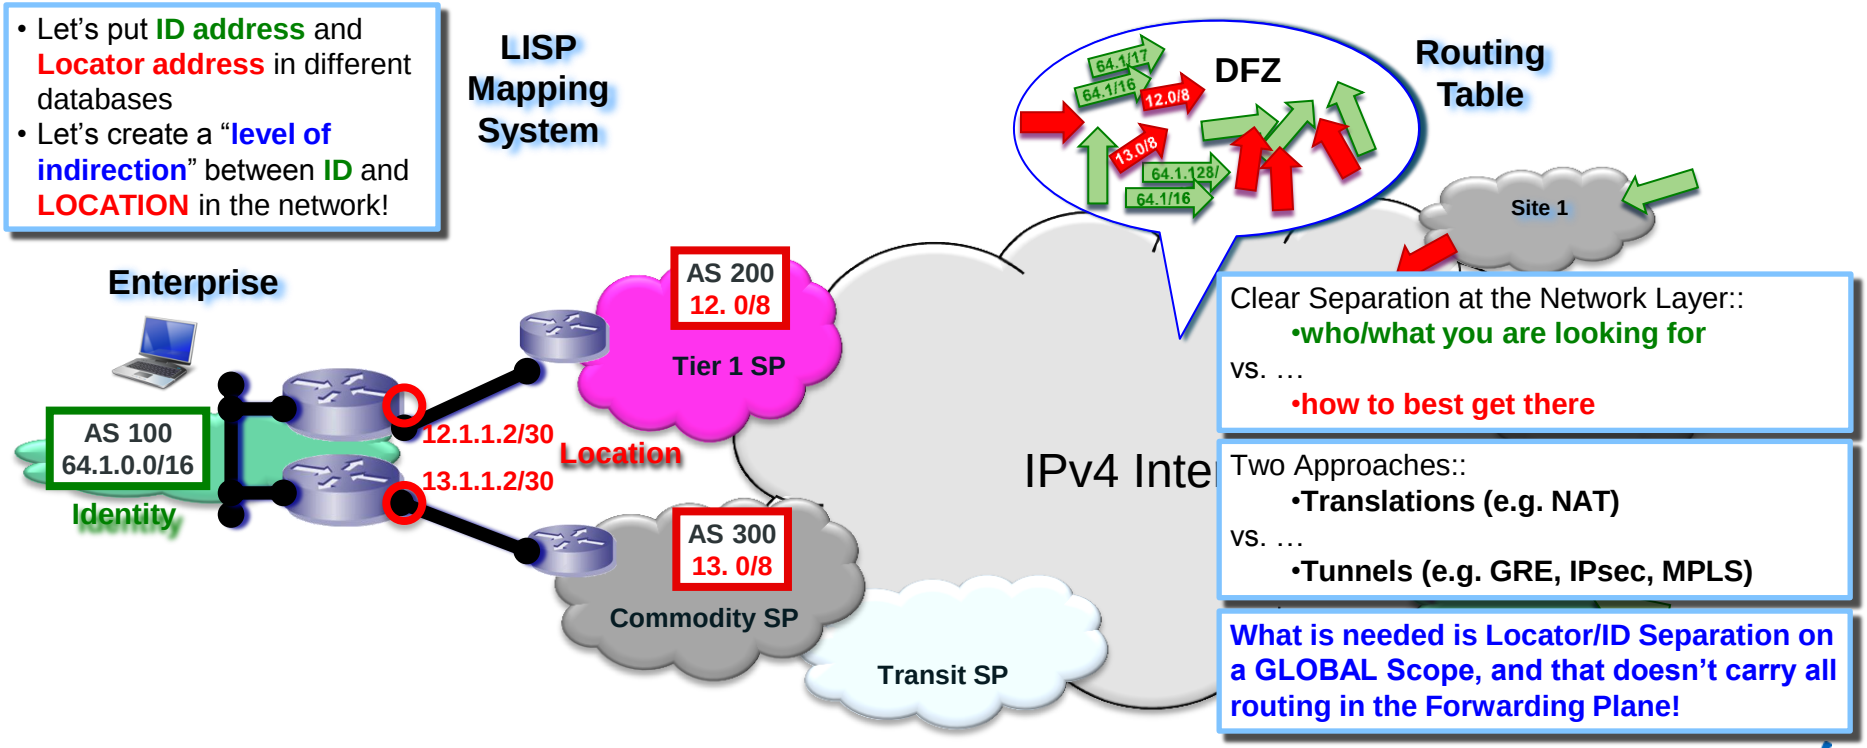

Clear Separation at the Network Layer::
- **who/what you are looking for**

vs. …
- **how to best get there**

Two Approaches::
- **Translations (e.g. NAT)**

vs. …
- **Tunnels (e.g. GRE, IPsec, MPLS)**

What is needed is Locator/ID Separation on a GLOBAL Scope, and that doesn't carry all routing in the Forwarding Plane!

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

Identity and Location :: an Overloaded Concept in Routing Today…

- Let's scale the **ID address** databases to 10^{10} and allow it to hold **any prefix length** (e.g. /32)
- Let's provide a mechanism to provide on-the-fly resolution of **ID** and **locator**
- High scale design, and ability to change **locator** for fixed **ID** enables Mobility!

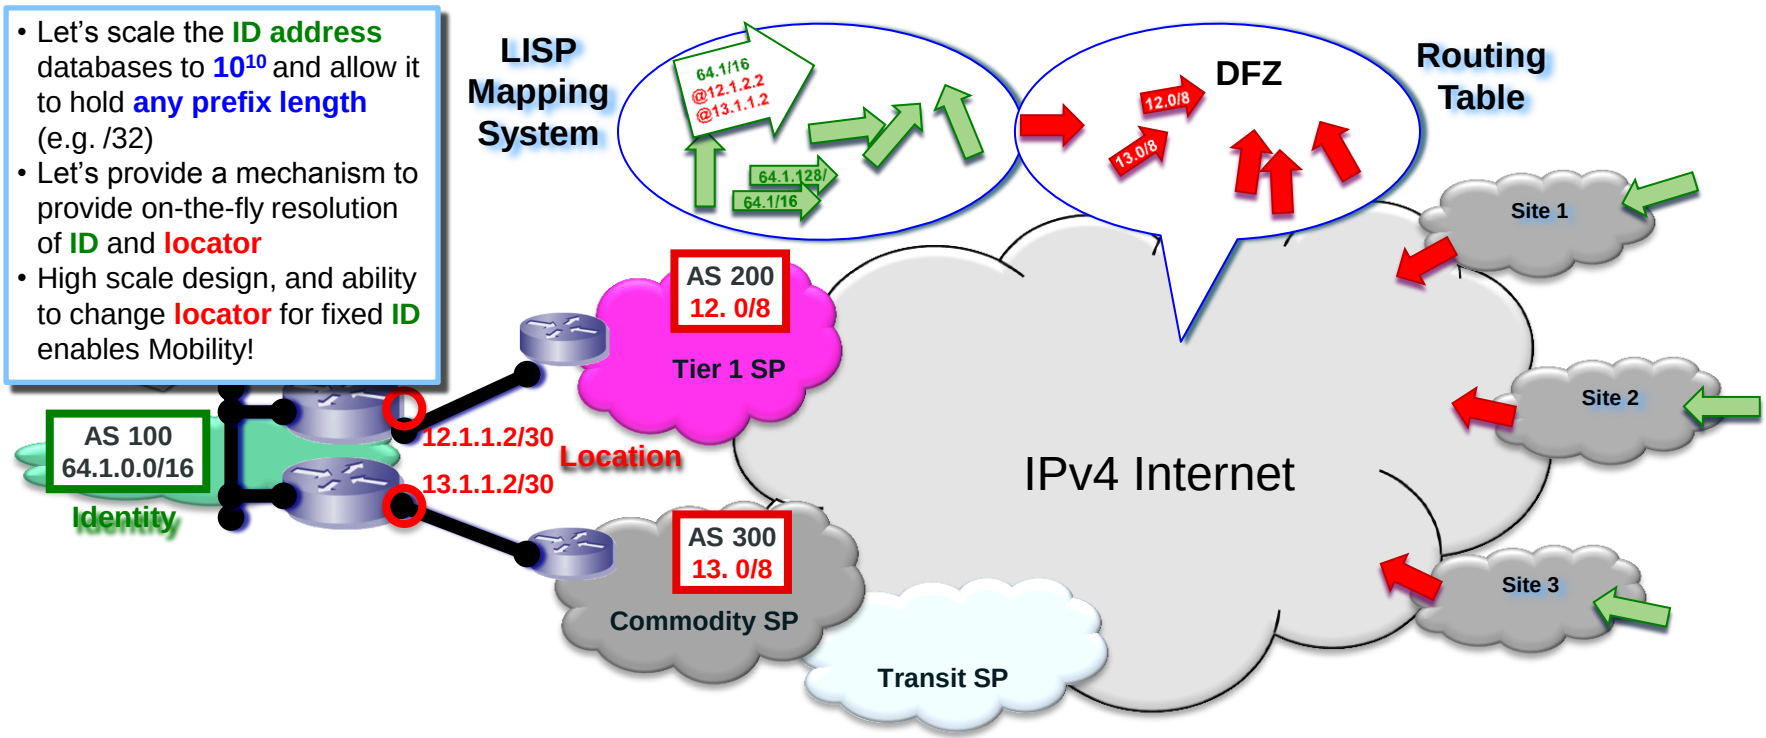

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

LISP :: A Routing Architecture – Not a Feature

level of indirection

IDENTITY **LOCATION**

LISP **changes** the current ROUTING Architecture

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

Locator/ID Separation :: The Mapping System is the Key

- A Mapping Systems is “key” for a Locator/ID separation architecture
  - Mapping systems provide the control plane for the architecture
  - Mapping systems represent the great opportunity for these architecture to excel
    - Most of the time, users/operators think about the data plane
    - The control plane is where the **magic** happens!
- Some general components of a mapping system to be aware…
  These affect how the system scales much differently than routing

  state :: must scale to large numbers (such as $\mathbf{10^{10}}$) of hosts

  rate :: must be small globally; damp reachability and mobility from globally impacting the system

  latency :: must be low enough not to harm existing applications

  scope :: must allow for both a global and a private scope for mapping

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

Locator/ID Separation :: Changing the Routing Architecture

- Locator/ID Separation “architecture” helps solve other current network problems
- IPv4/IPv6 Co-existence at the “ID” and “Locator” spaces
  - IPv4 and IPv6 can be implemented at the “ID” and/or “locator” spaces for simple integration
  - In reality, **anything** can be an “ID” and carried over traditional cores (IPv4 and IPv6)
    - e.g. RFID, VIN#, Geo-Location, MAC-Addr, etc. etc. etc.
- Scaling IP Mobility is very similar to scaling Internet Multihoming
  - Mobility moves an “ID” (unique address) from one network “location” to another network “location”
  - Multihoming connects an “ID” (a unique address) to multiple networks “locations” at the same time
  - With both Mobility and Multihoming, the network must keep more specific state “globally” about where something is located

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

LISP :: A Routing Architecture – Not a Feature

- Uses **pull** vs. **push** routing
  -
  -
- An **over-the-top** technology
  - Address Family agnostic
  - Incrementally deployable
  - End systems can be unaware of LISP
- Deployment **simplicity**
  - No host changes
  - Minimal CPE changes
  - Some new core infrastructure components
- LISP use-cases are **complimentary**
  - Simplified multi-homing with Ingress traffic Engineering; no need for BGP
  - Address Family agnostic support
  - Virtualisation support
  - End-host mobility without renumbering
- Enables **IP Number Portability**
  - Never change host IP's; No renumbering costs
  - No DNS changes; "name == EID" binding
  - Session survivability
- An **Open Standard**
  - Being developed in the IETF (RFC 6830-6836)
  - No Cisco Intellectual Property Rights

© 2014 Cisco and/or its affiliates. All rights reserved.




LISP Operations

LISP Operations

Main Attributes of LISP

- LISP namespaces
  - **EID (Endpoint Identifier)**
  - **RLOC (Routing Locator)** is the IP address of the LISP router for the host
  - **EID-to-RLOC mapping** is the distributed architecture that maps **EIDs** to **RLOCs**

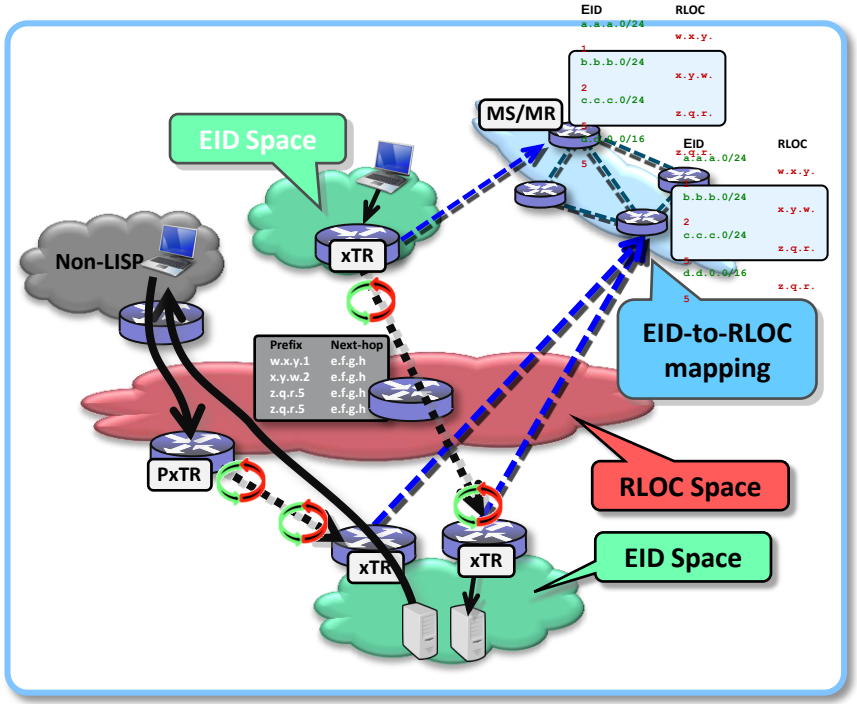

- Network-based solution
- No host changes
- Minimal configuration
- No DNS changes
- Address Family agnostic
- Incrementally deployable (support LISP and non-LISP)
- Support for mobility

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP :: Mapping Resolution “Level of Indirection” DNS Analog

- LISP “Level of Indirection” is analogous to a DNS lookup
  - DNS resolves **IP addresses** for **URL** **Answering the “WHO IS” question**

DNS
Name-to-IP
URL Resolution

  - LISP resolves **locators** for queried **identities** **Answering the “WHERE IS” question**

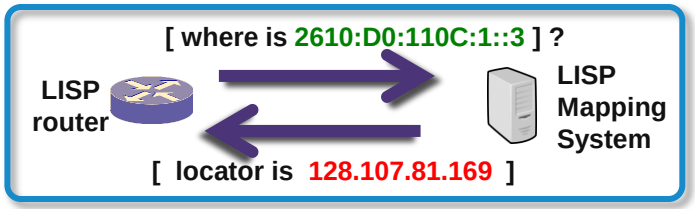

LISP
Identity-to-locator
Mapping Resolution

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP IPv4 EID / IPv4 RLOC Data Packet Header Example

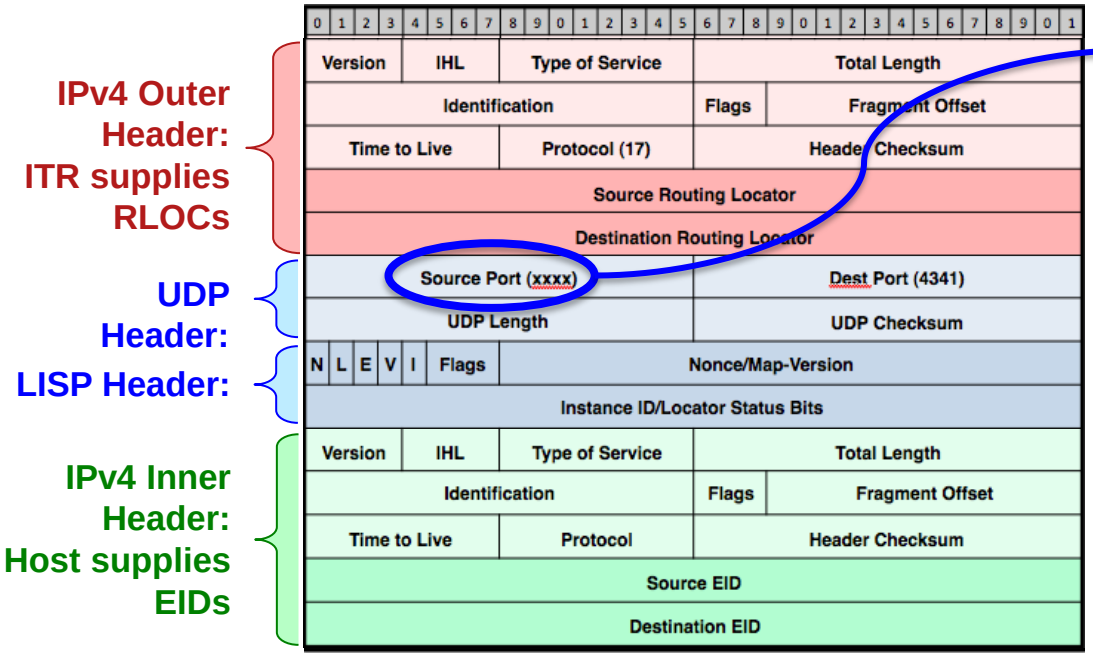

Q: How does the UDP source port get selected ?

A: It's either a 3-tuple or 5-tuple hash of the inner (EID) header

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Encapsulation Combinations – IPv4 and IPv6 Supported

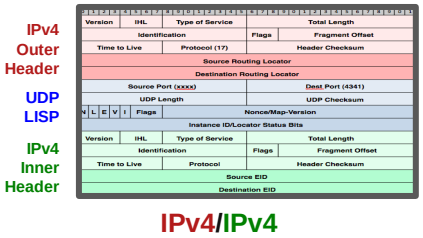

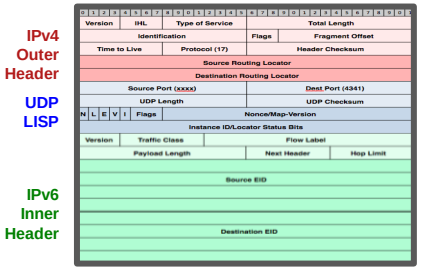

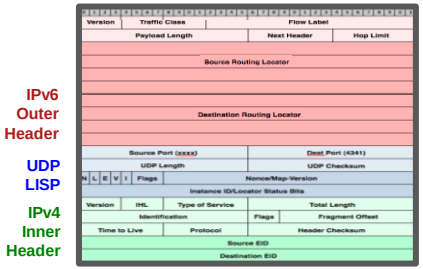

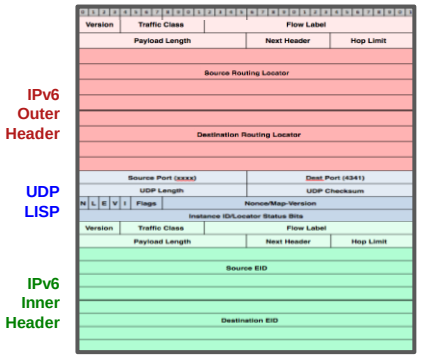

Q: Doesn't encapsulation cause MTU issues?

A: It can… But preparation limits issues…

- Encapsulation overhead is 36B IPv4 and 56B IPv6
- LISP supports "stateful" (PMTUD) and "stateless" (fragmentation) options
- Tunnel/MTU issues are well known (GRE, IPsec, etc.) and are usually operationally tractable

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public

LISP Operations

LISP Data Plane :: 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 and deliver packets to local **EIDs** at site

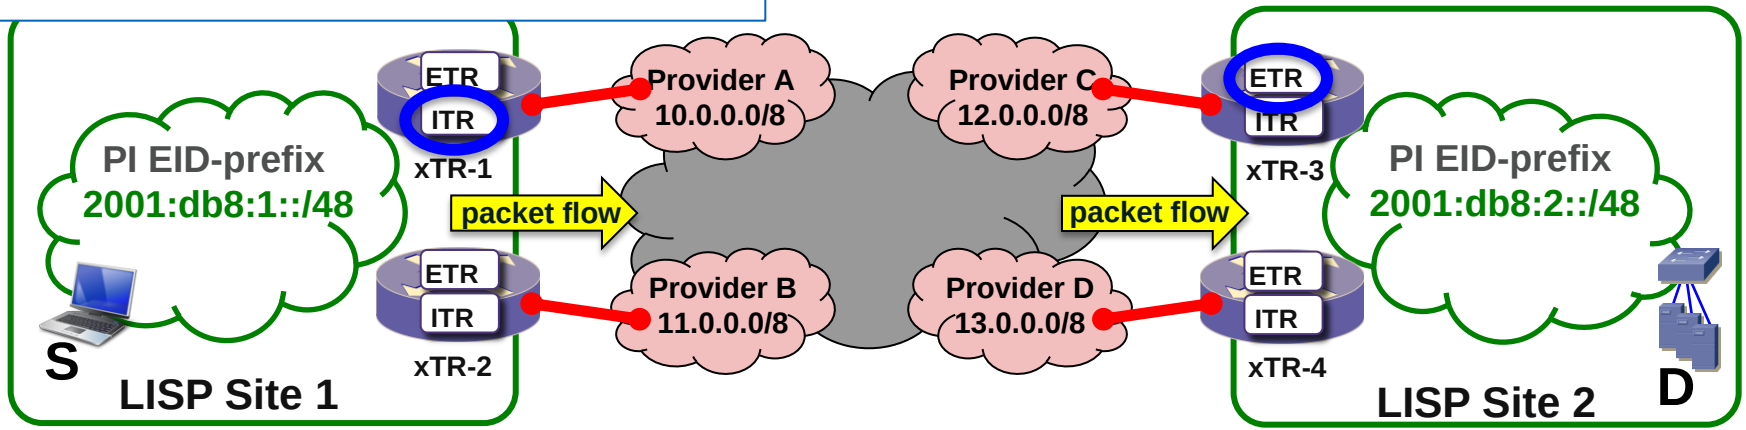

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Data Plane :: Unicast Packet Flow

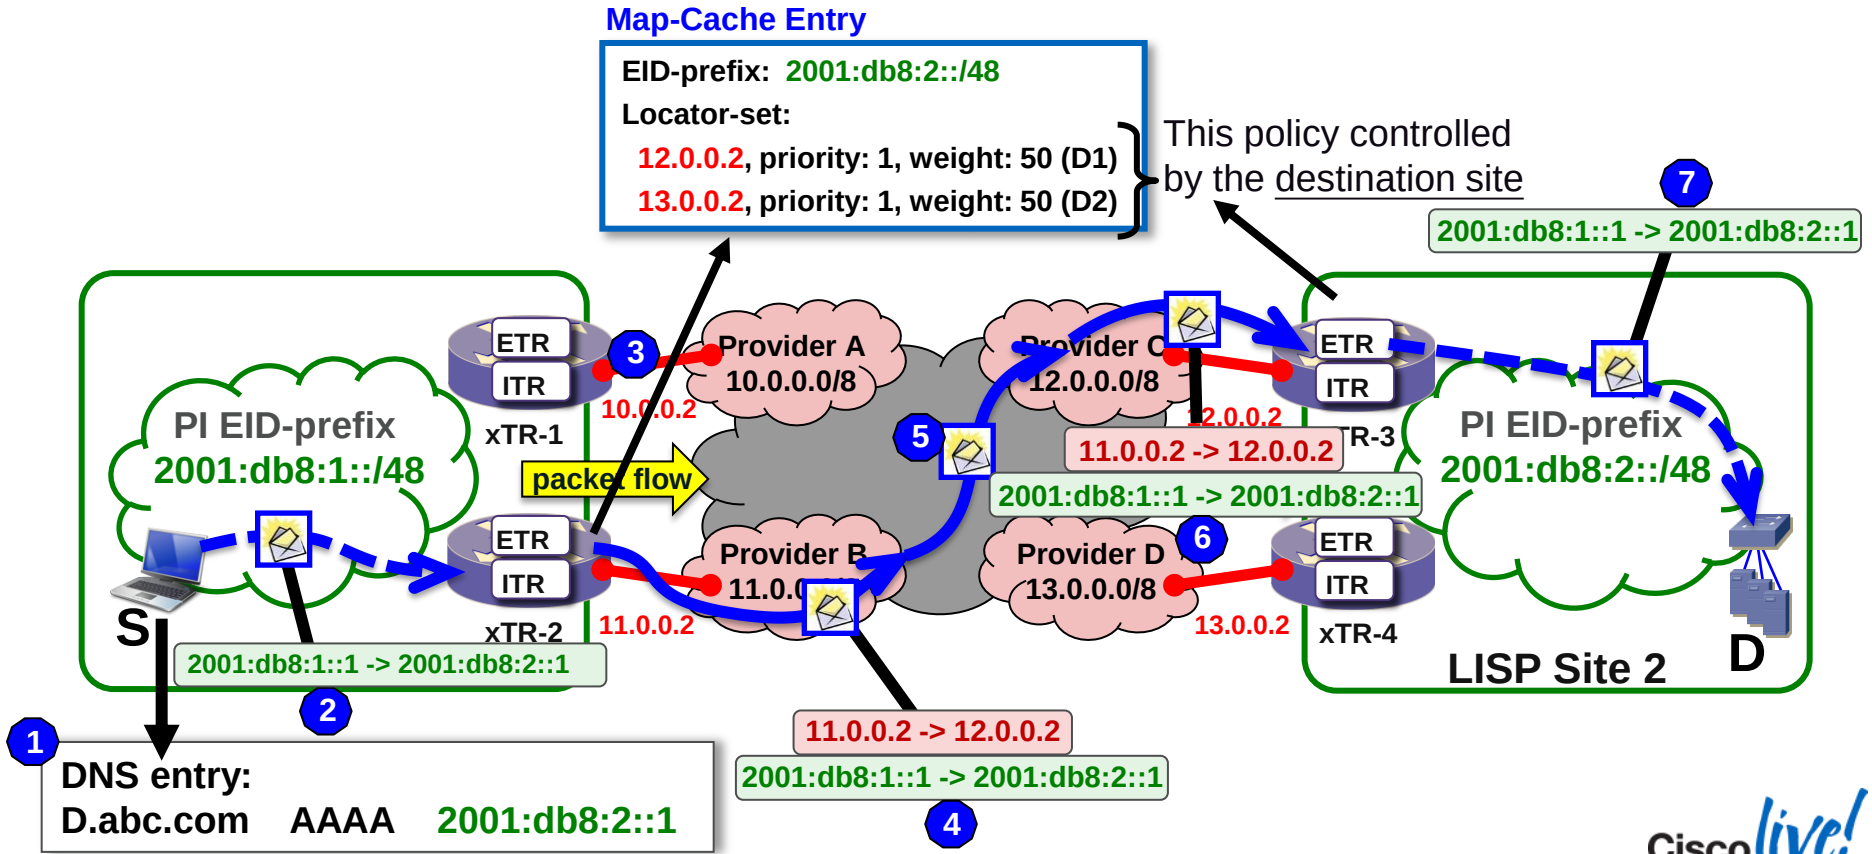

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public

LISP Operations

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

```
!
router lisp
 locator-set SITE2
  12.0.0.2 priority 1 weight 50
  13.0.0.2 priority 1 weight 50
  exit
 !
 eid-table default instance-id 0
  database-mapping 2001:db8:2::/48 locator-set SITE2
  exit
 !
 ipv6 itr map-resolver 66.2.2.2
 ipv6 itr
 ipv6 etr map-server 66.2.2.2 key S3cr3t-2
 ipv6 etr
 exit
!
ip route 0.0.0.0 0.0.0.0 12.0.0.1
ip route 0.0.0.0 0.0.0.0 13.0.0.1
!
```

Identical config on both xTRs!

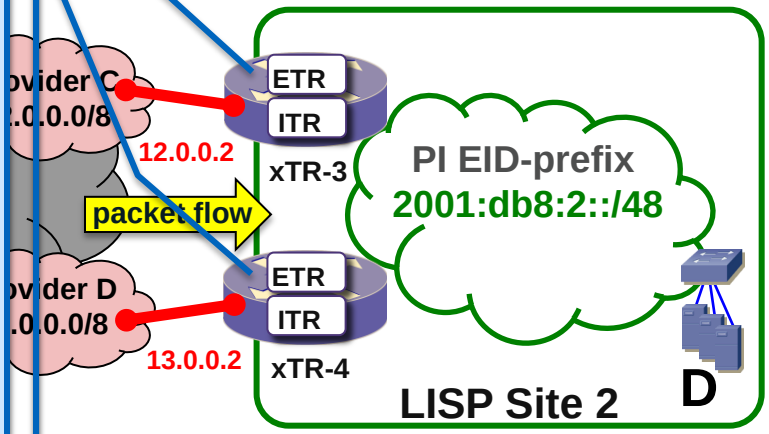

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Control Plane :: Introduction…

- LISP Control Plane Provides On-Demand Mappings
  - Control Plane is separate from the Data Plane (UDP 4342 vs UDP 4341)
  - Map-Resolver and Map-Server (similar to DNS Resolver and DNS Server)
  - LISP Control Plane Messages for EID-to-RLOC
  - Distributed databases and map-caches hold mappings

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Control Plane :: Map-Server/Map-Resolver (MS/MR)

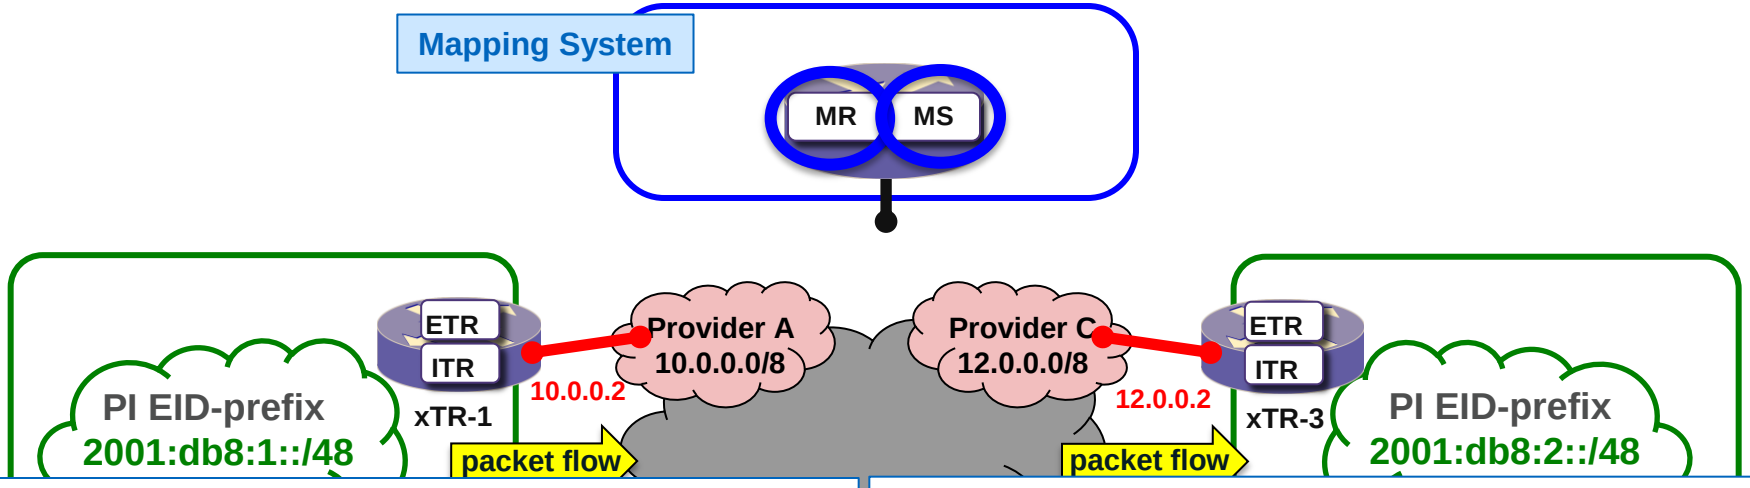

MR – Map-Resolver

- Receives Map-Request from ITR
- Forwards Map-Request to Mapping System
- Sends Negative Map-Replies in response to Map-Requests for non-LISP sites

MS – Map-Server

- LISP site ETRs register their EID prefixes here; requires configured “lisp site” policy, authentication key
- Receives Map-Requests via Mapping System, forwards them to registered ETRs

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Control Plane :: Map-Server/Map-Resolver (MS/MR)

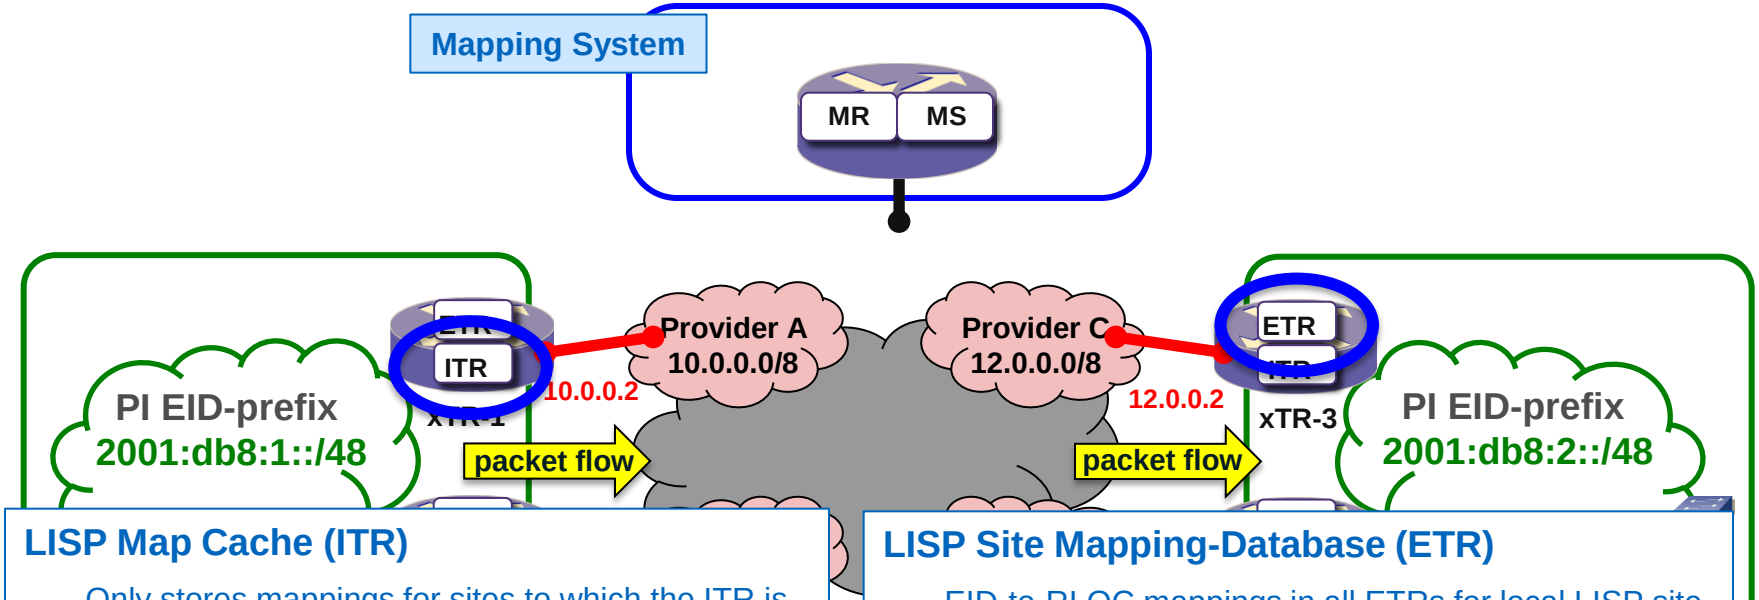

LISP Map Cache (ITR)

- Only stores mappings for sites to which the ITR is currently sending packets
- Populated by receiving Map-Replies from ETRs
- ITRs must respect Map-Reply policy (TTLs, RLOC up/down status, RLOC priorities/weights

LISP Site Mapping-Database (ETR)

- EID-to-RLOC mappings in all ETRs for local LISP site
- ETR is “authoritative” for its EIDs, sends Map-Replies to ITRs
- ETRs can tailor policy based on Map-Request source

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Control Plane :: Control Plane Messages…

- Control Plane Control Plane EID Registration
  - **Map-Register message**

    Sent by ETR to MS to register its associated EID prefixes

    Specifies the RLOC(s) to be used by the MS when forwarding Map-Requests to the ETR
- Control Plane "Data-triggered" mapping services
  - **Map-Request message**

    Sent by an ITR when it needs an EID/RLOC mapping, to test an RLOC for reachability, to refresh a mapping before TTL expiration, or in response to a Solicit Map-Request (SMR)
  - **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
  - **Map-Notify message**

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Control Plane :: Map-Registration Example

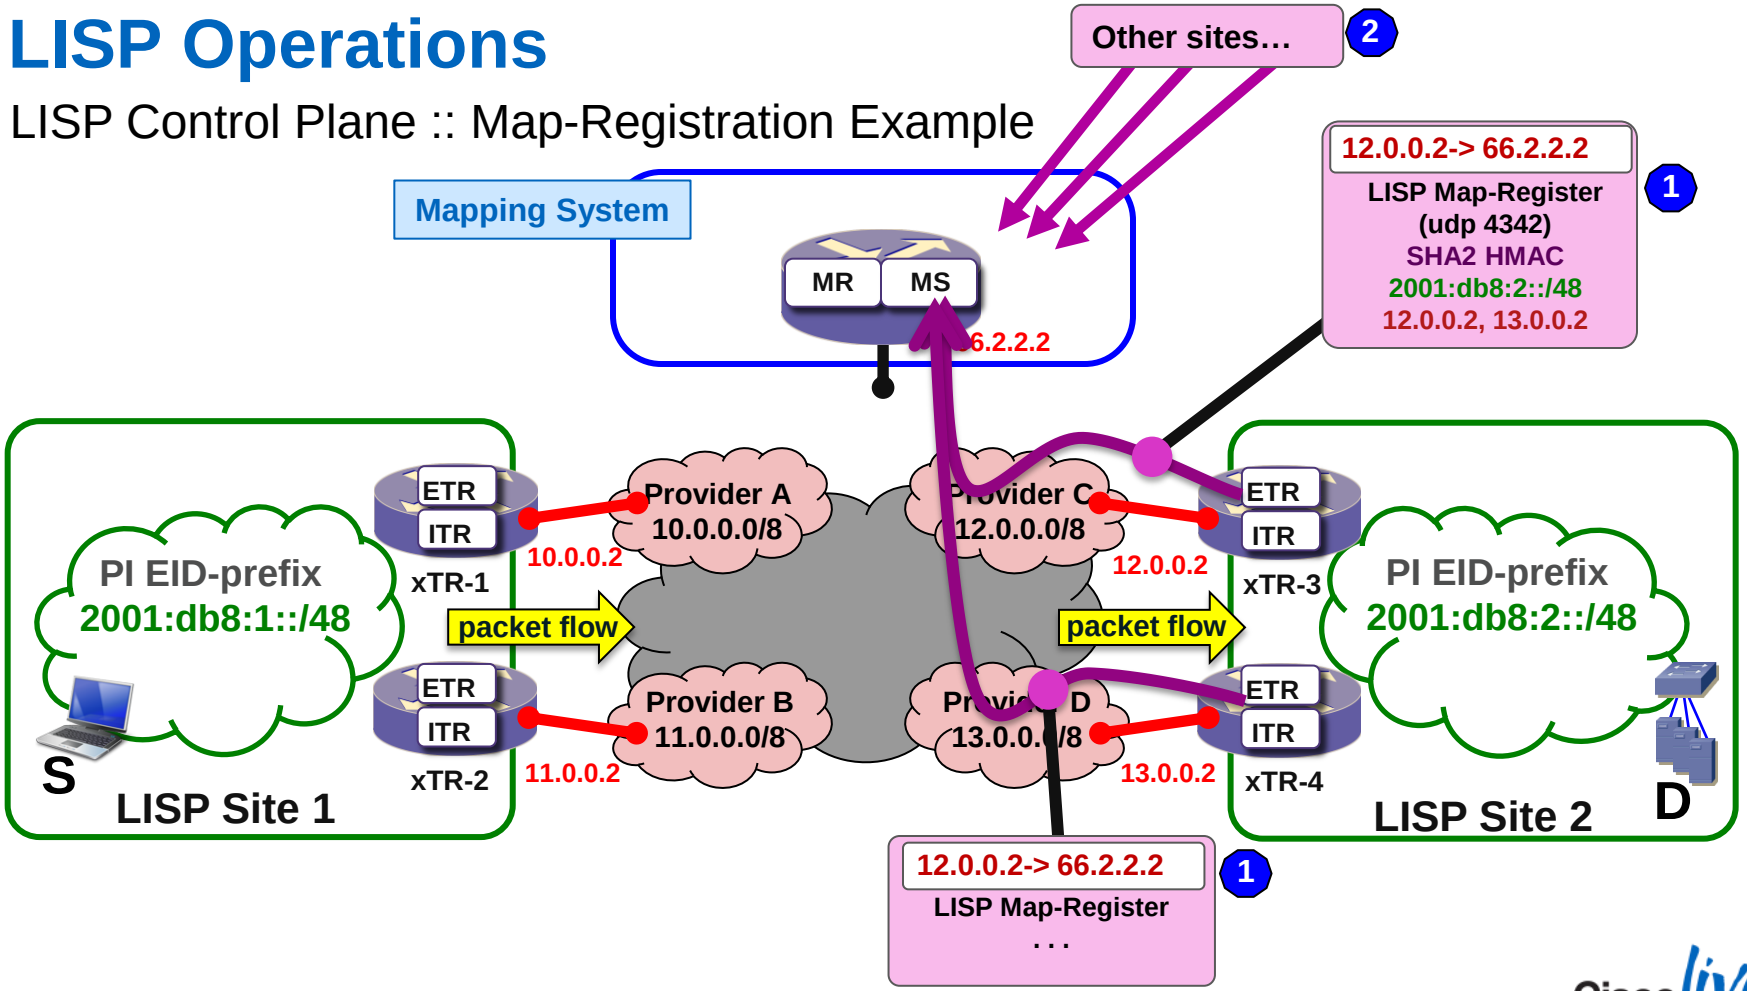

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Control Plane :: Map-Request/Map-Reply Example

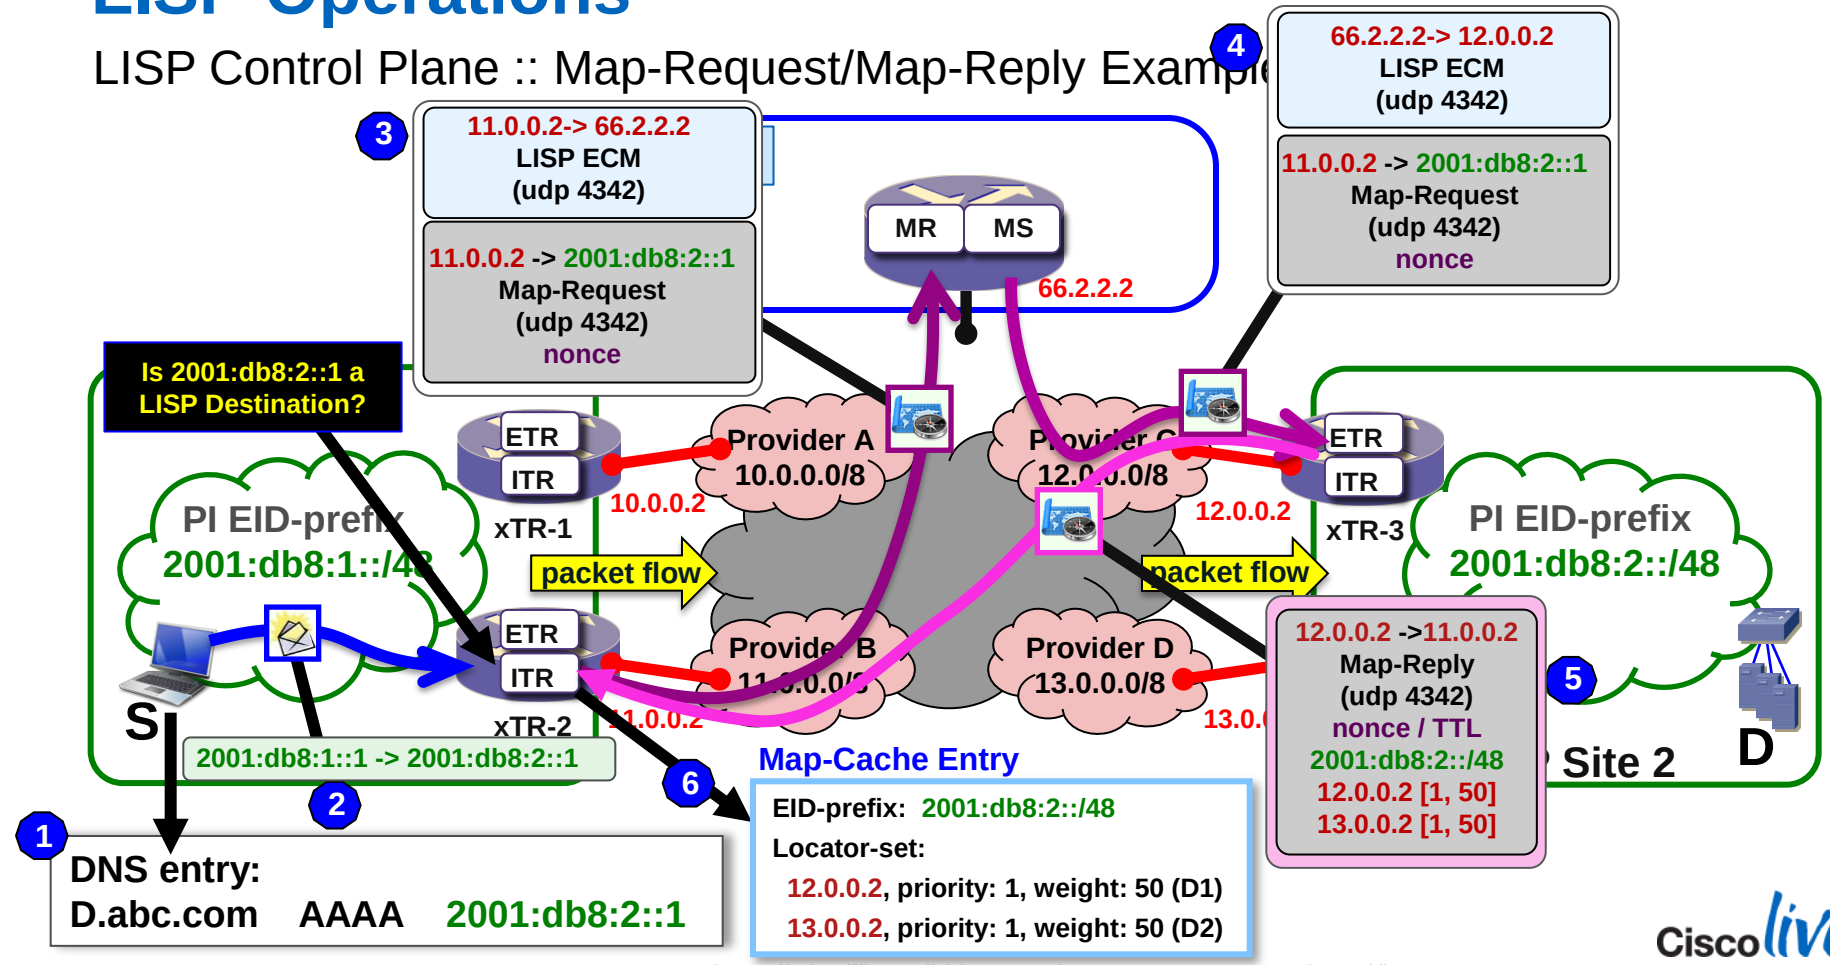

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public

LISP Operations

LISP Control Plane :: Map-Proxy Map-Reply Example

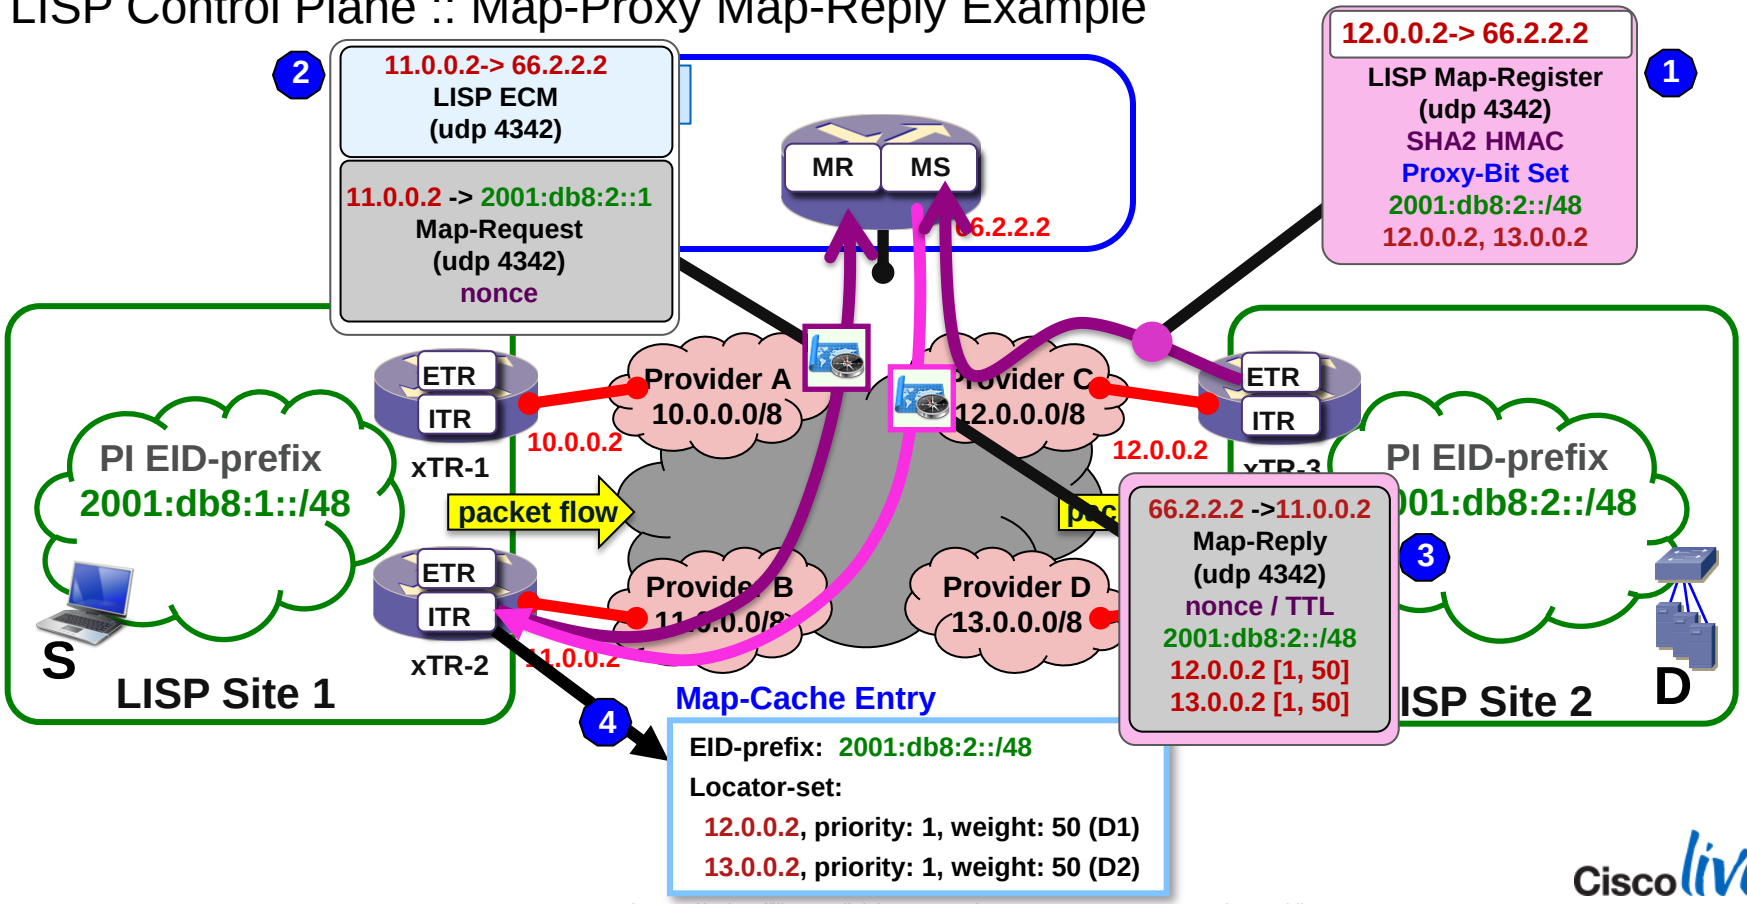

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Control Plane :: Negative Map-Reply Example

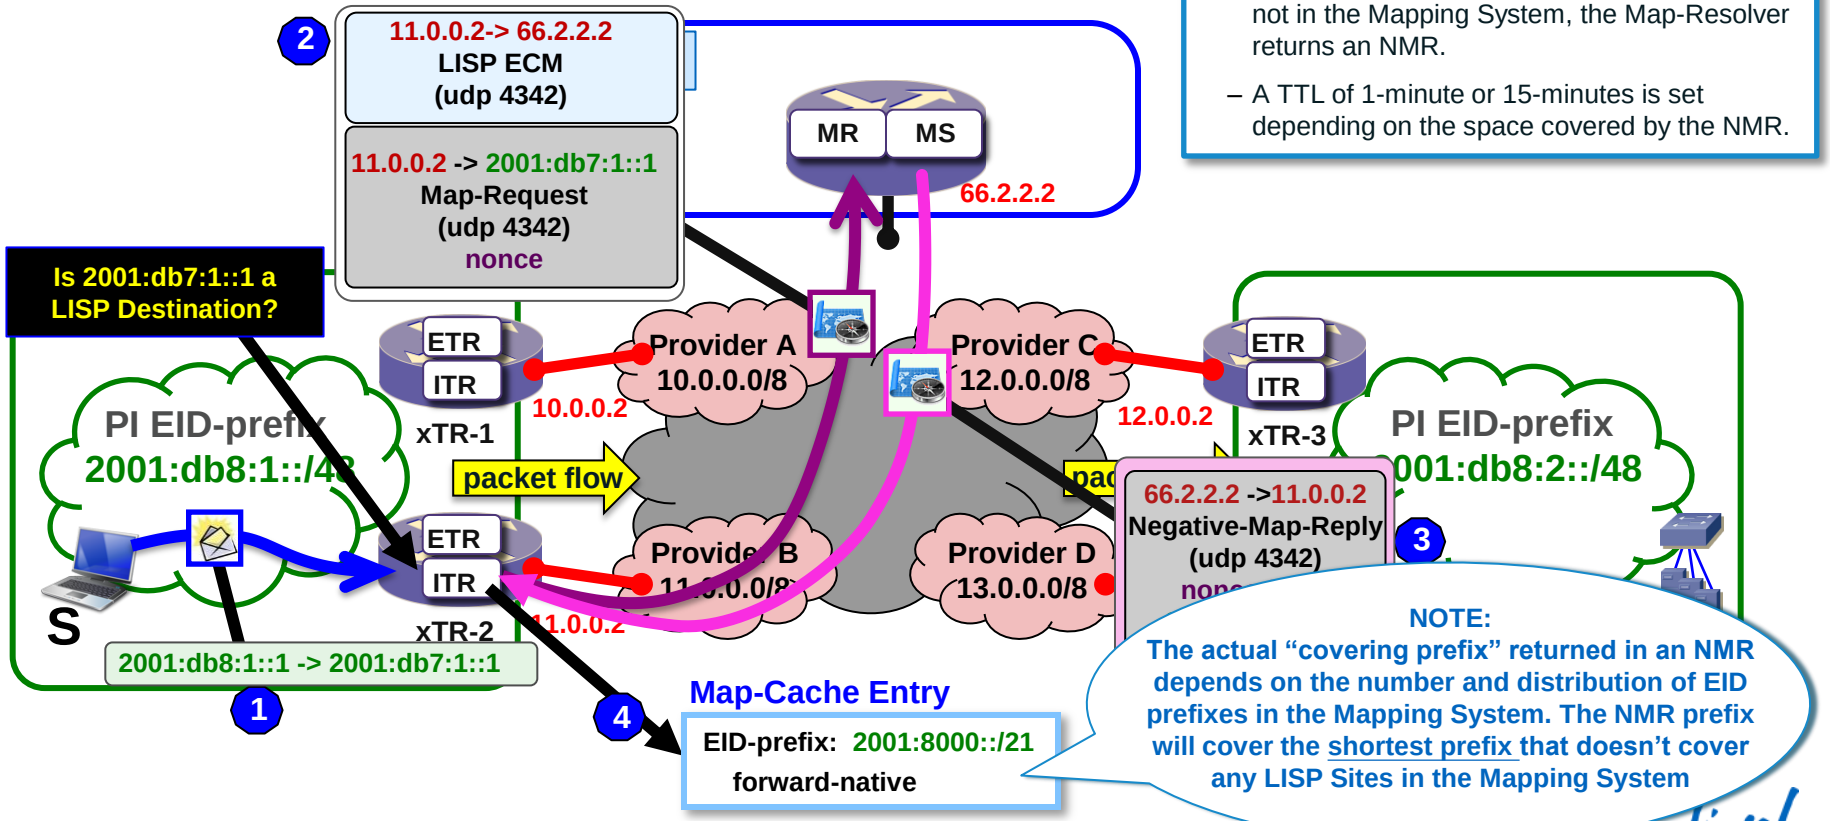

Notes:

- When an ITR queries for a destination that is not in the Mapping System, the Map-Resolver returns an NMR.
- A TTL of 1-minute or 15-minutes is set depending on the space covered by the NMR.

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Control Plane :: Map-Server/Map-Resolver (MS/M

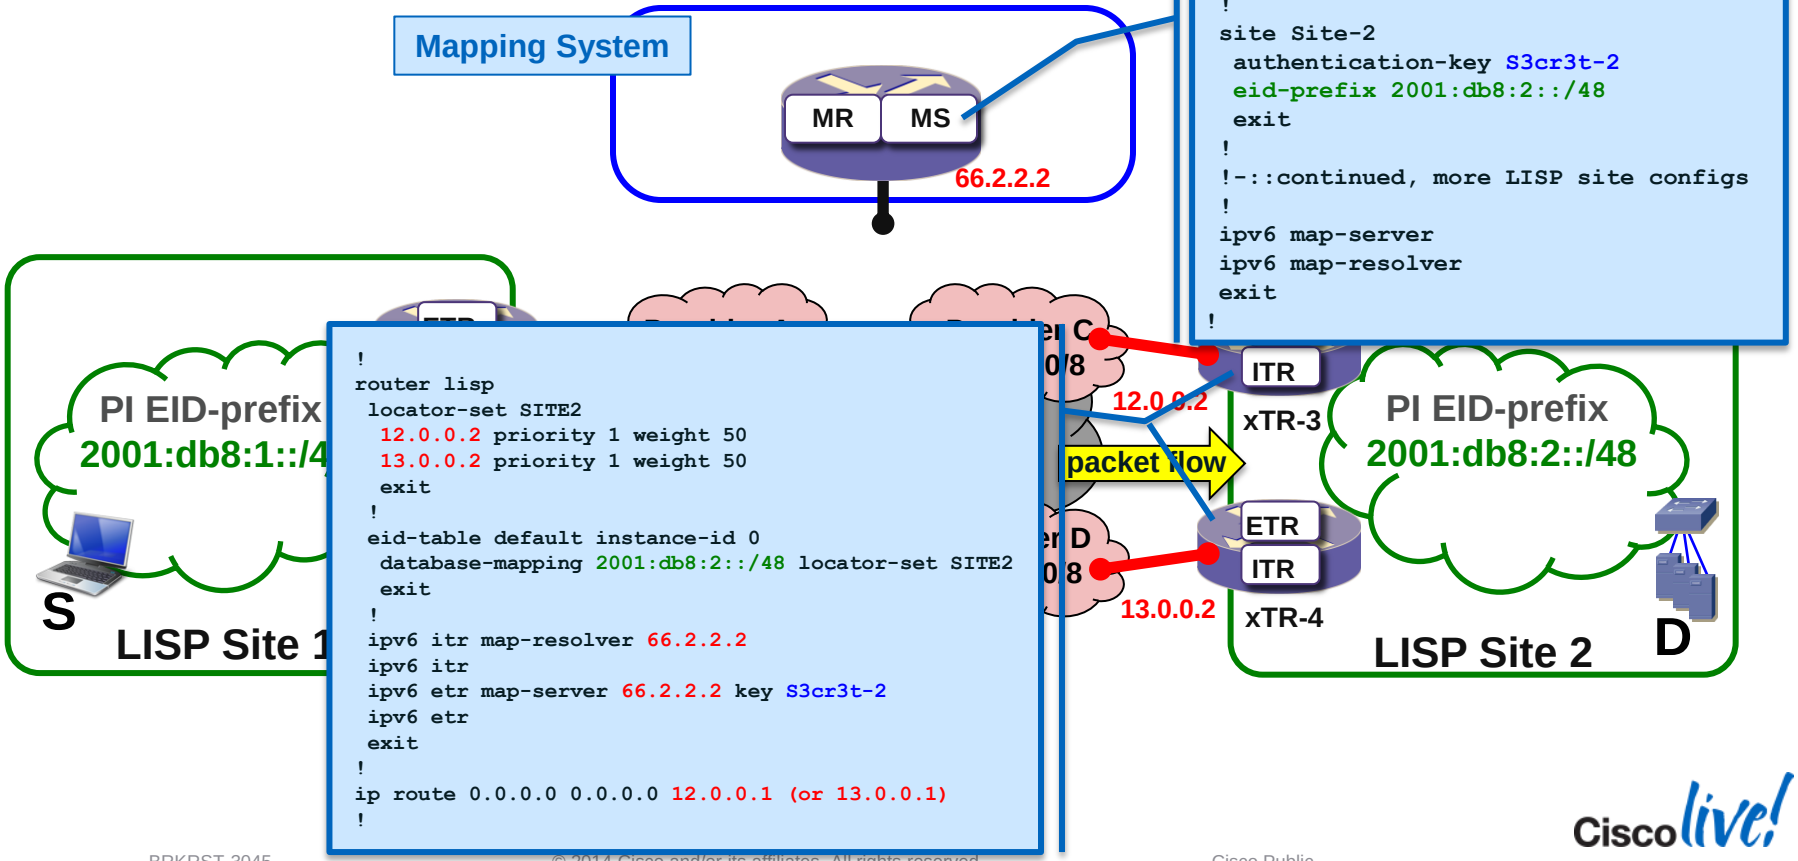


```
!
router lisp
 site Site-1
  authentication-key S3cr3t-1
  eid-prefix 2001:db8:1::/48
  exit
 !
 site Site-2
  authentication-key S3cr3t-2
  eid-prefix 2001:db8:2::/48
  exit
 !
 !-::continued, more LISP site configs
 !
 ipv6 map-server
 ipv6 map-resolver
 exit
!
```

```
!
router lisp
 locator-set SITE2
  12.0.0.2 priority 1 weight 50
  13.0.0.2 priority 1 weight 50
  exit
 !
 eid-table default instance-id 0
  database-mapping 2001:db8:2::/48 locator-set SITE2
  exit
 !
 ipv6 itr map-resolver 66.2.2.2
 ipv6 itr
 ipv6 etr map-server 66.2.2.2 key S3cr3t-2
 ipv6 etr
 exit
!
ip route 0.0.0.0 0.0.0.0 12.0.0.1 (or 13.0.0.1)
!
```


© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

The LISP Beta Network uses DDT today…

LISP Control Plane :: Mapping System Scaling…

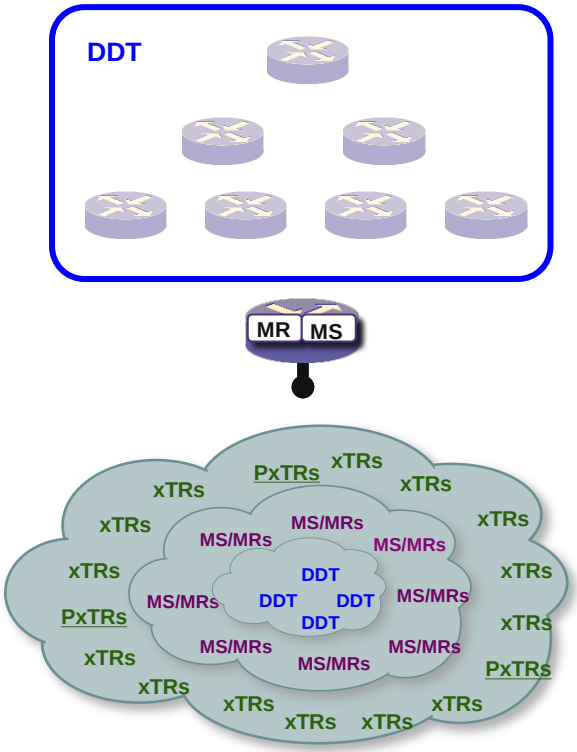

Scaling the LISP Mapping System

- Deploy multiple “stand-alone” Map-Servers” and register each LISP Site to all of them (up to eight)
- Deploy Map-Resolvers in an “Anycast” manner
- Or, deploy a “hierarchical” Mapping System - DDT

DDT – Delegated Distributed Tree

- Hierarchy for Instance IDs and for EID Prefixes
- DDT Map-Resolvers sends (ECM) Map-Requests
- DDT Nodes Return Map-Referral messages
- DDT Resolvers resolve the Map-Server’s RLOC iteratively
- Conceptually, similar to DNS (IN-ADDR hierarchy) but different prefix encoding, messages, etc.

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Deployment Overview

Private and Public LISP Deployment Models…

Private Model

- "Private" LISP deployment support single Enterprises or Entities
- LISP Enterprise deploys:
  - xTRs
  - Mapping System, if required
  - Proxy System, if required

Private Enterprise Examples

Public Model

- "Public" LISP deployment supports the needs of multiple Enterprises
- LISP Service Provider deploys "shared" Mapping System and Proxy System
- LISP Enterprises subscribe to LISP SP, and deploy their own xTRs

Stand-Alone Example

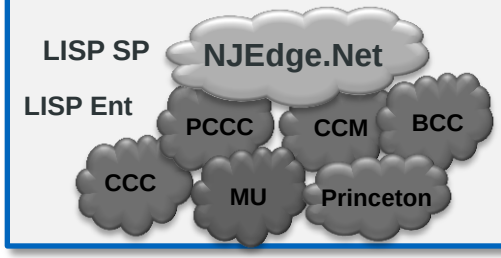

Global Examples

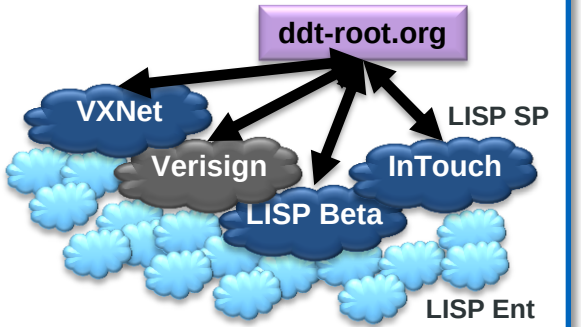

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Internetworking :: Day-One Incremental Deployment

- Early Recognition
  - Up-front recognition of an incremental deployment plan
  - LISP will not be widely deployed day-one
- Interworking for:
  - **LISP-sites** to **non-LISP sites** (e.g. the rest of the Internet)
  - **non-LISP sites** to **LISP-sites**
- Proxy-ITR/Proxy-ETR are deployed today
  - Infrastructure LISP network entity
  - Creates a monetised service opportunity for infrastructure players

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Internetworking :: Day-One Incremental Deployment

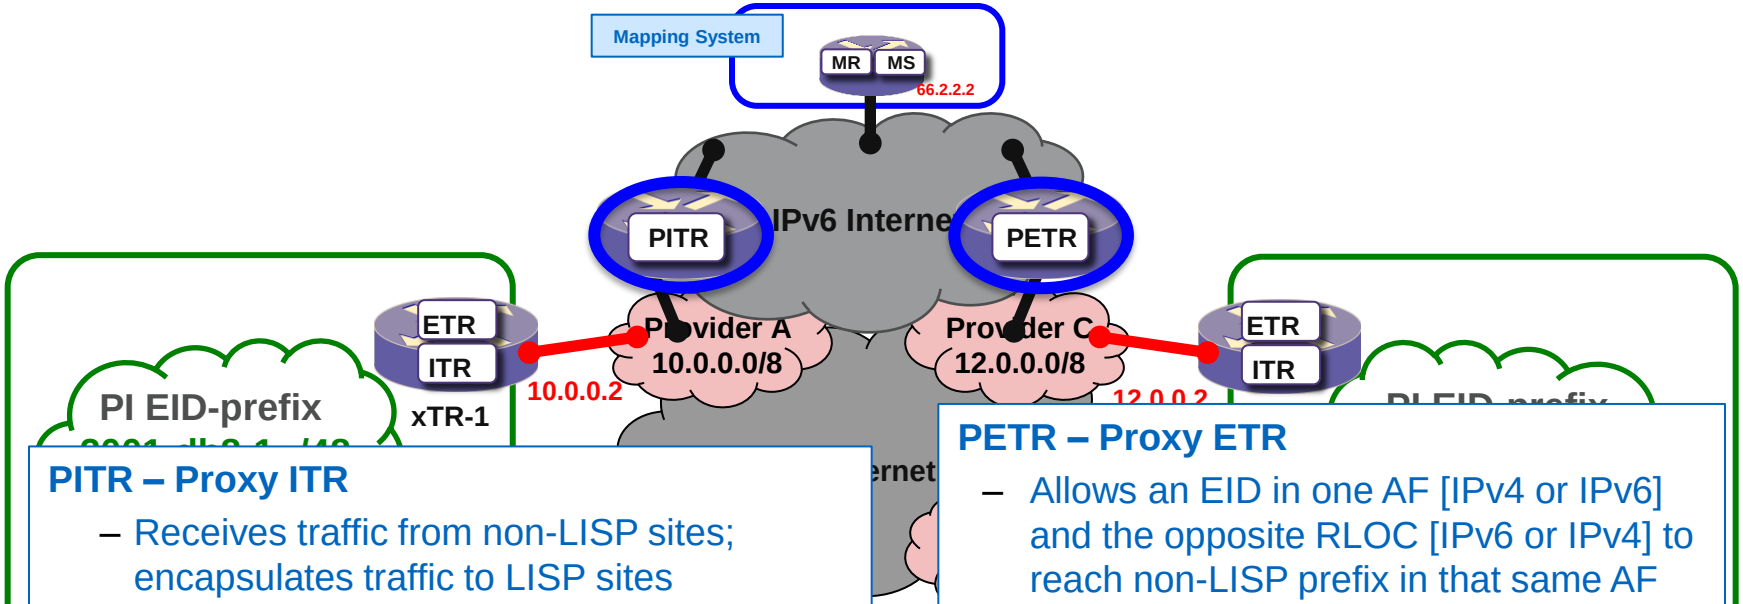

PITR – Proxy ITR

- Receives traffic from non-LISP sites; encapsulates traffic to LISP sites
- Advertises coarse-aggregate EID prefixes
- LISP sites see ingress TE “day-one”

PETR – Proxy ETR

- Allows an EID in one AF [IPv4 or IPv6] and the opposite RLOC [IPv6 or IPv4] to reach non-LISP prefix in that same AF (AF-hop-over)
- Allows LISP sites with uRPF restrictions to reach non-LISP sites

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Operations

LISP Internetworking :: Day-One Incremental Deployment

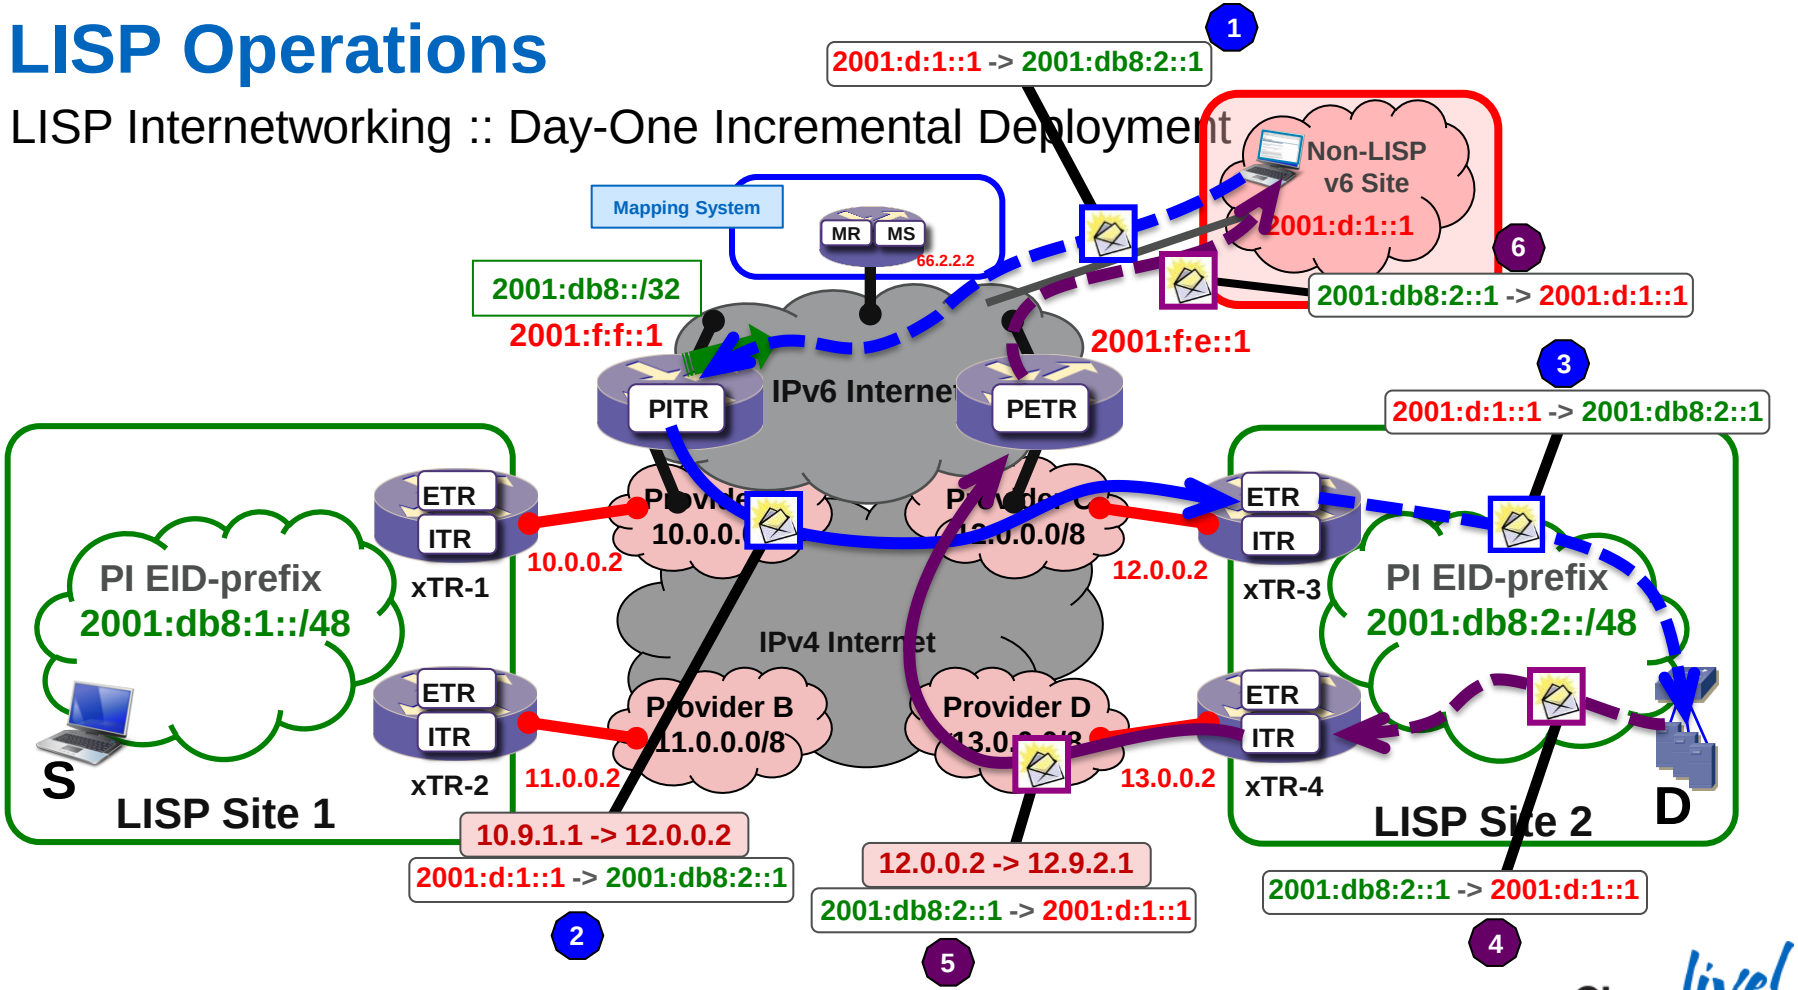

© 2014 Cisco and/or its affiliates. All rights reserved.




LISP Deployment Examples

LISP Deployment Examples

LISP Deployment Examples…

1. Efficient Multihoming and Multi-AF (IPv4 and IPv6)
2. Efficient Virtualisation and High-Scale VPNs
3. Data Centre/Host Mobility
4. LISP-Mobile Node

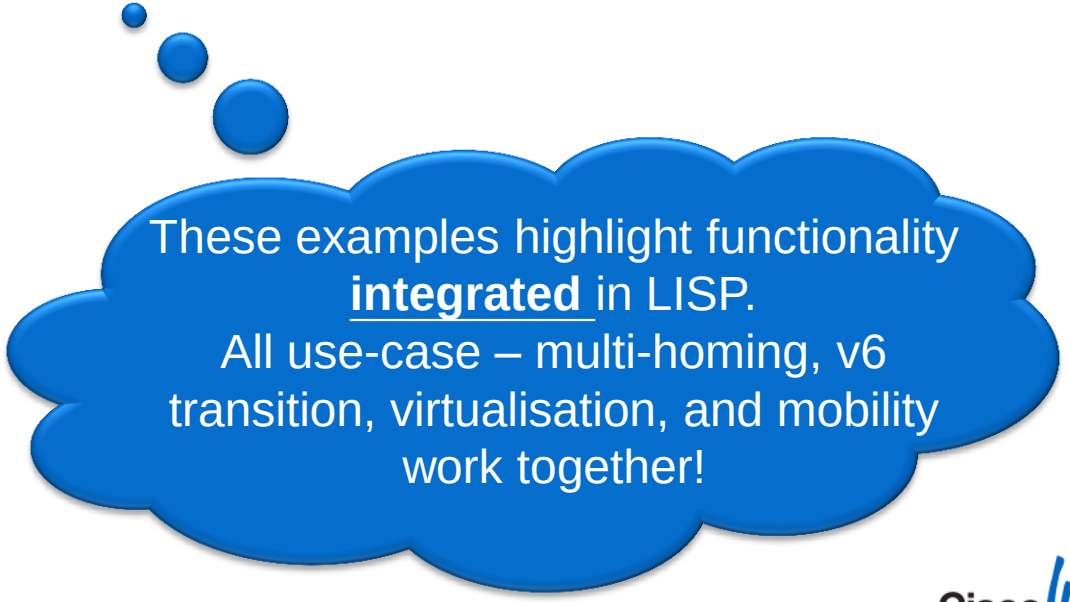

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Deployment Examples

LISP Deployment Examples…

1. Efficient Multihoming and Multi-AF (IPv4 and IPv6)
2. Efficient Virtualisation and High-Scale VPNs
3. Data Centre/Host Mobility
4. LISP-Mobile Node

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming/Multi-AF Details

LISP Operations

LISP Encapsulation – Any IPv4 and IPv6 Combination Supported

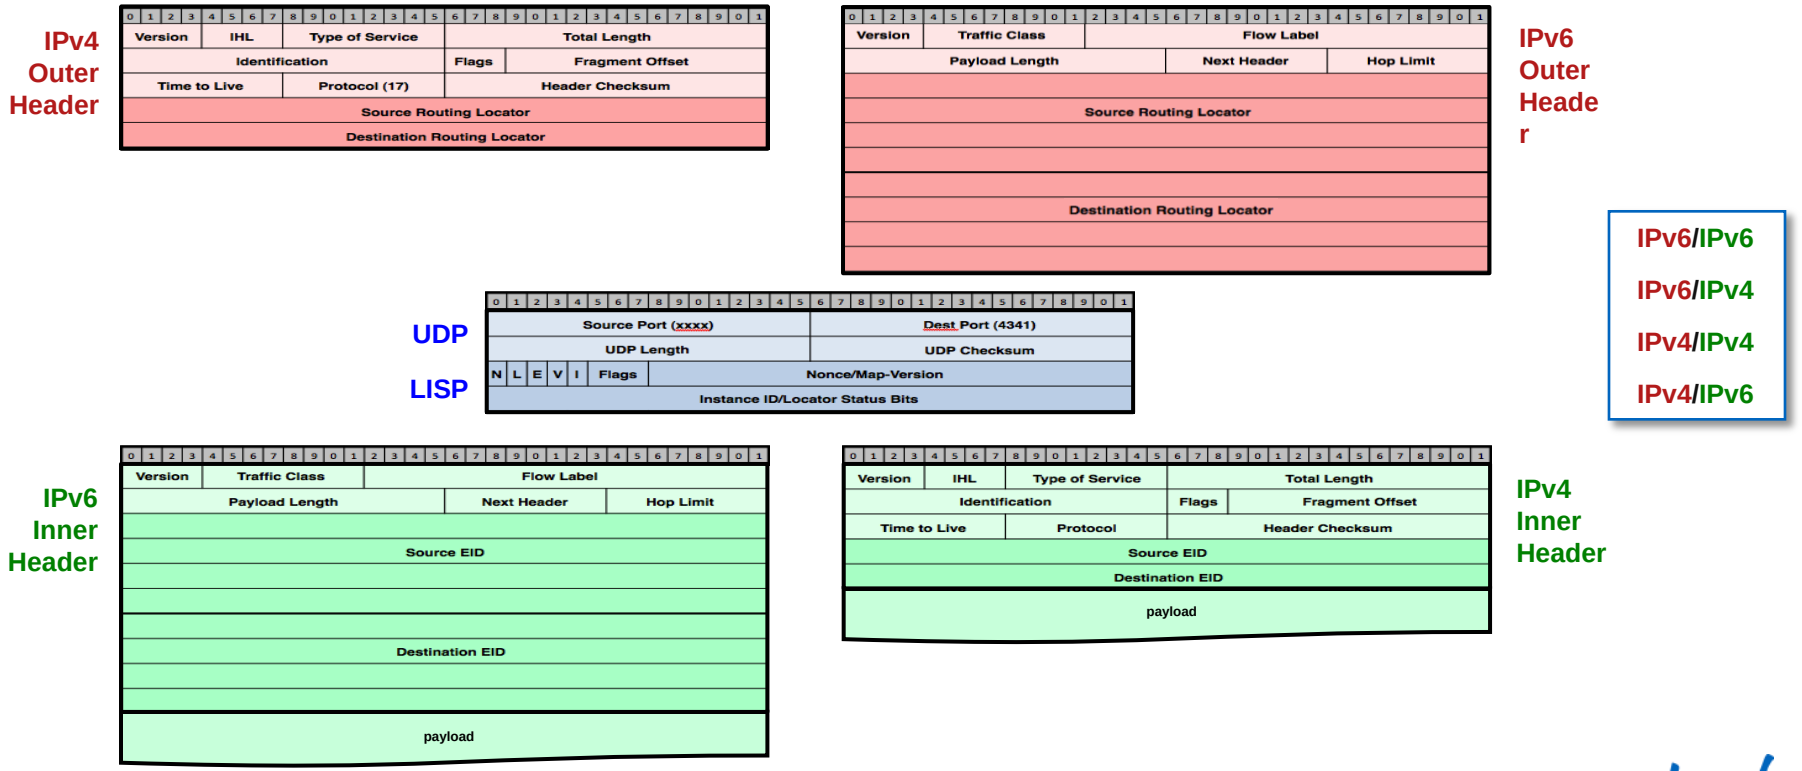

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

Inherent Support for AF-Agnostic Operations

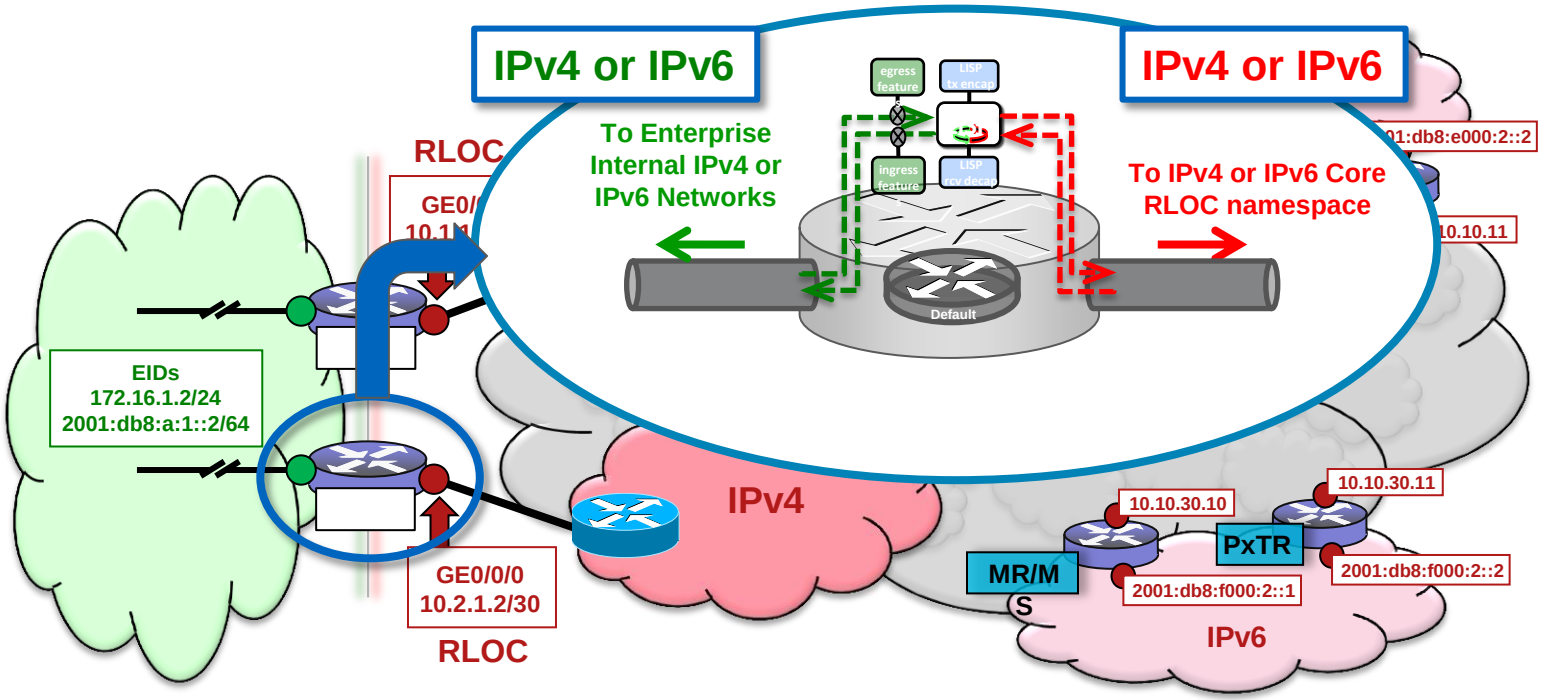

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

Inherent Support for AF-Agnostic Operations

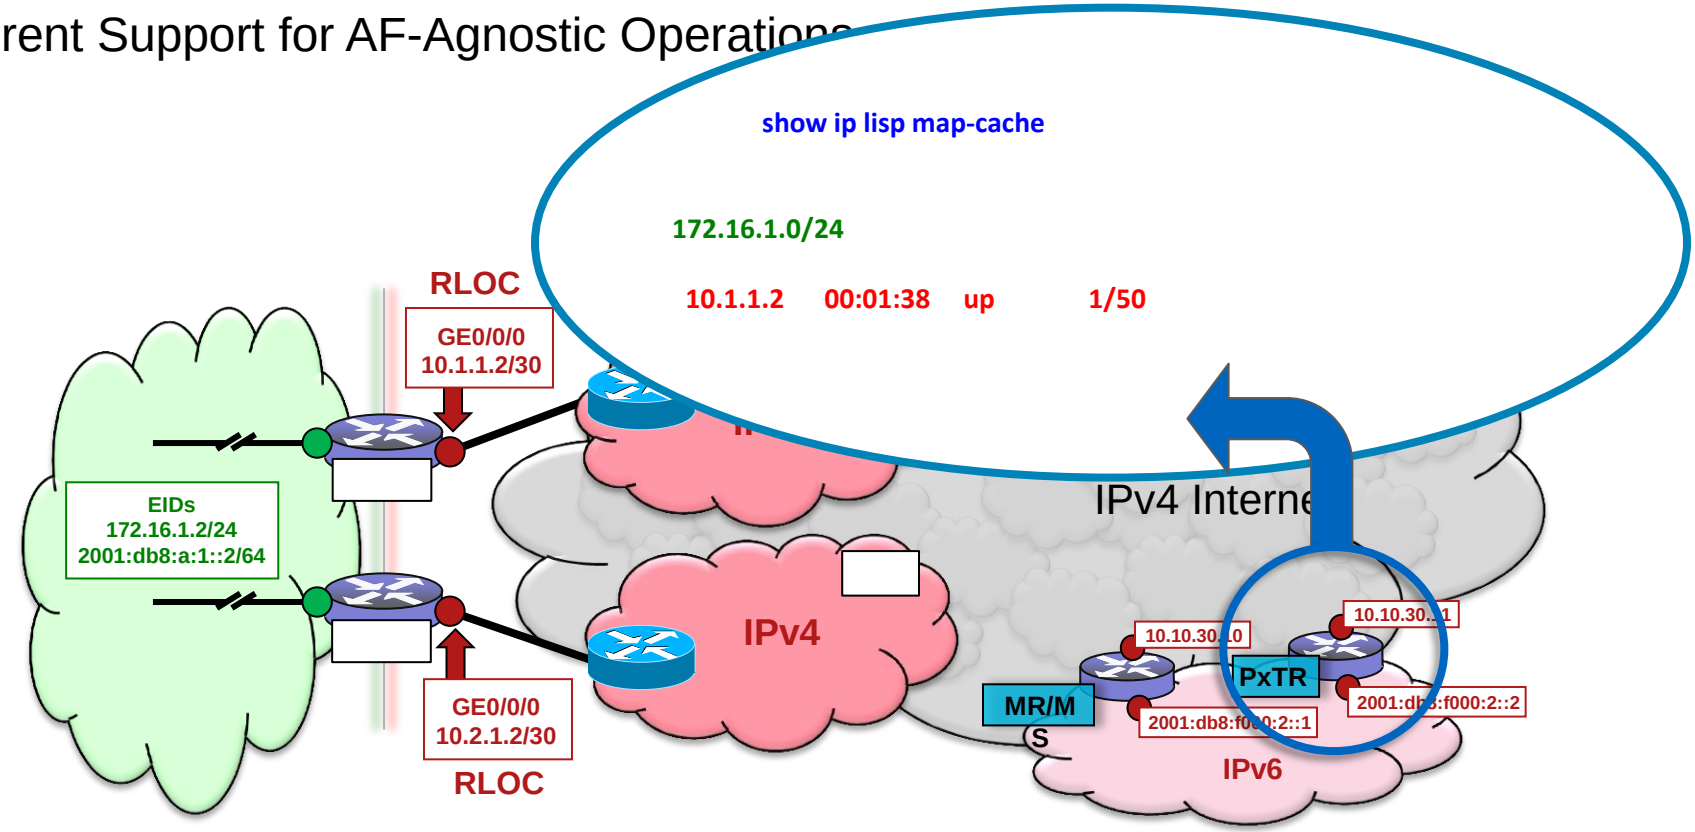

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

Inherent Support for AF-Agnostic Operations

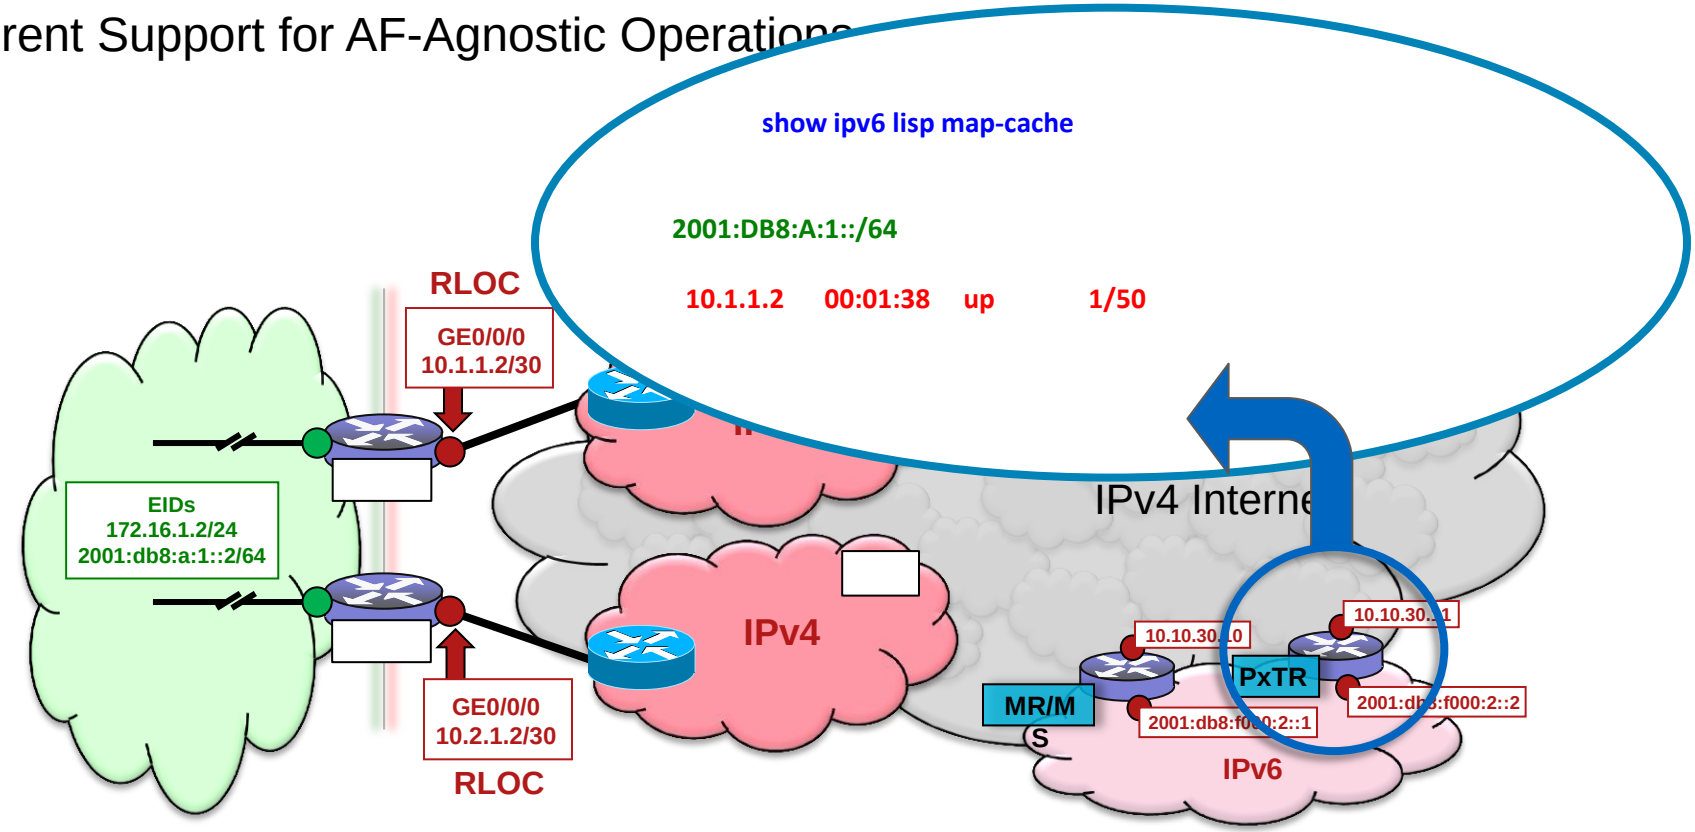

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

Efficient Multi-Homing and Multi-AF Support

- Needs:
  - Site connectivity to multiple providers for resiliency
  - Low OpEx/CapEx solution for Ingress TE
  - Rapid IPv6 deployment, minimal disruption
- LISP Solution:
  - LISP provides a streamlined solution for handling multi-provider connectivity and policy without BGP complexities
  - LISP encapsulation is Address Family agnostic, allowing for IPv6 over an IPv4 core, or IPv4 over an IPv6 core
- Benefits:
  - OpEx-friendly multi-homing across different providers
  - Simple policy management
  - Ingress Traffic Engineering that actually “works”
  - Minimal configuration
  - No core network changes

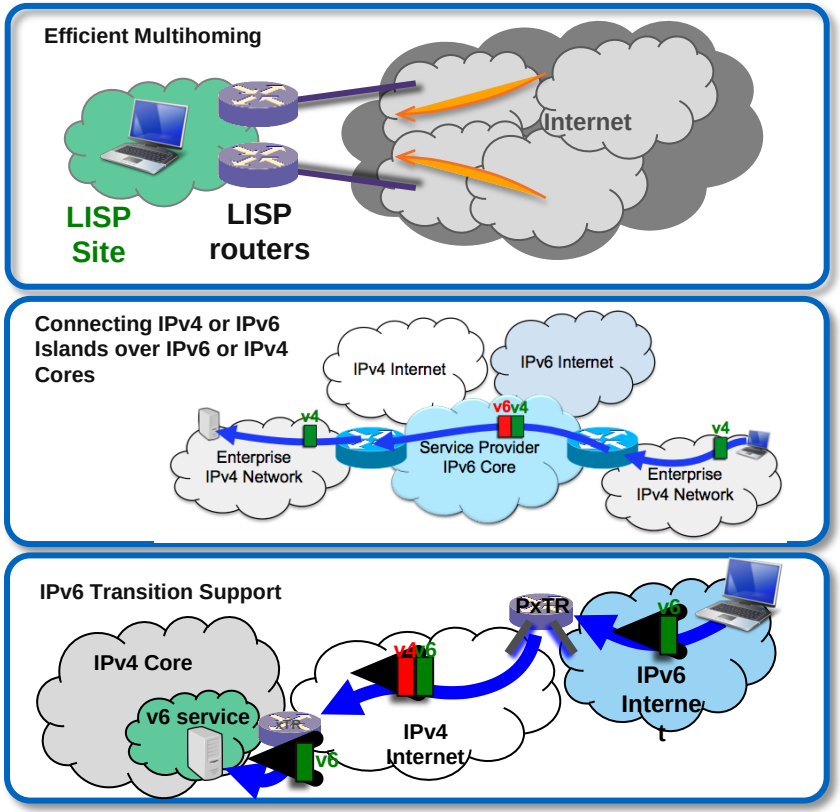

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

Efficient Multi-Homing and Multi-AF -- Customer Example

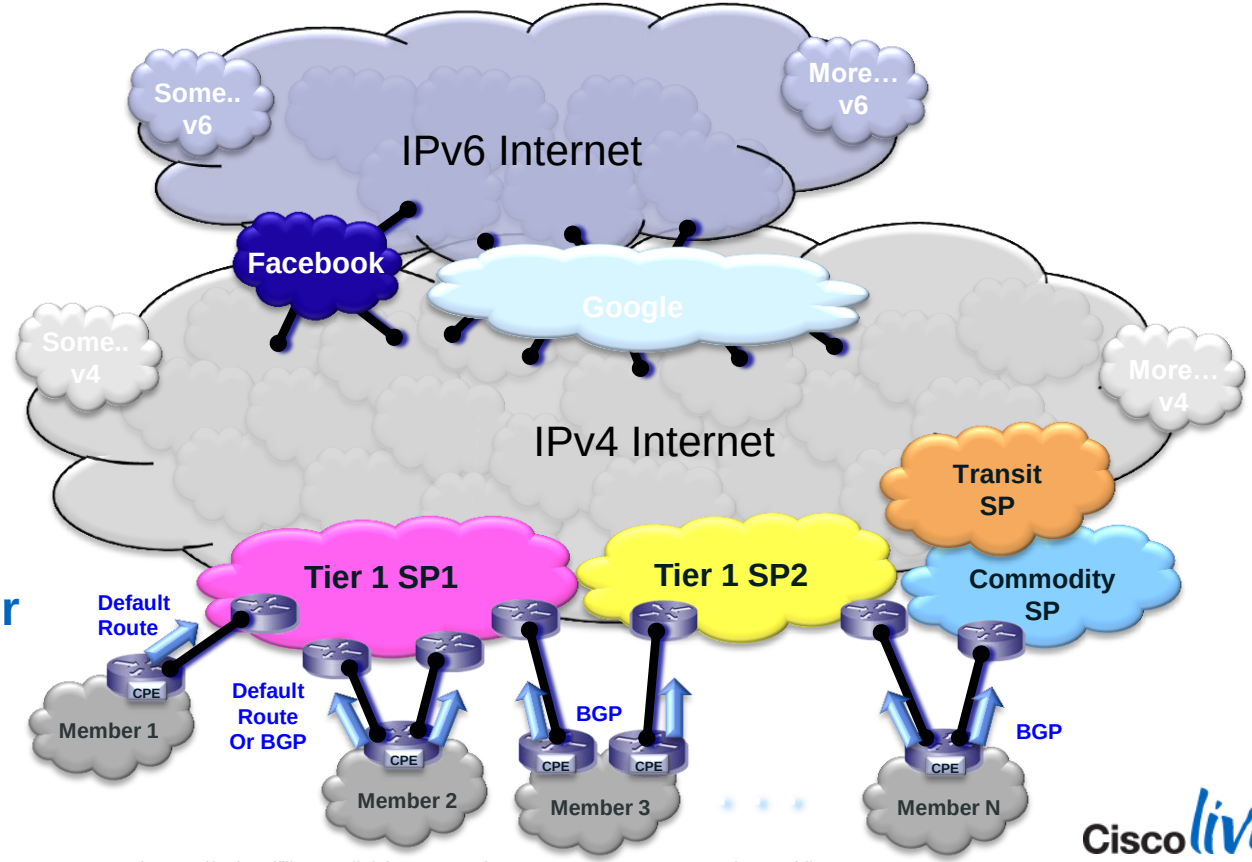

Constituent Member Topologies…

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

NJEDge.Net

Efficient Multi-Homing and Multi-AF -- Customer Example

- **Configuration complexity…**
- **Uneven multihoming load shares…**

Constituent Member Topologies…

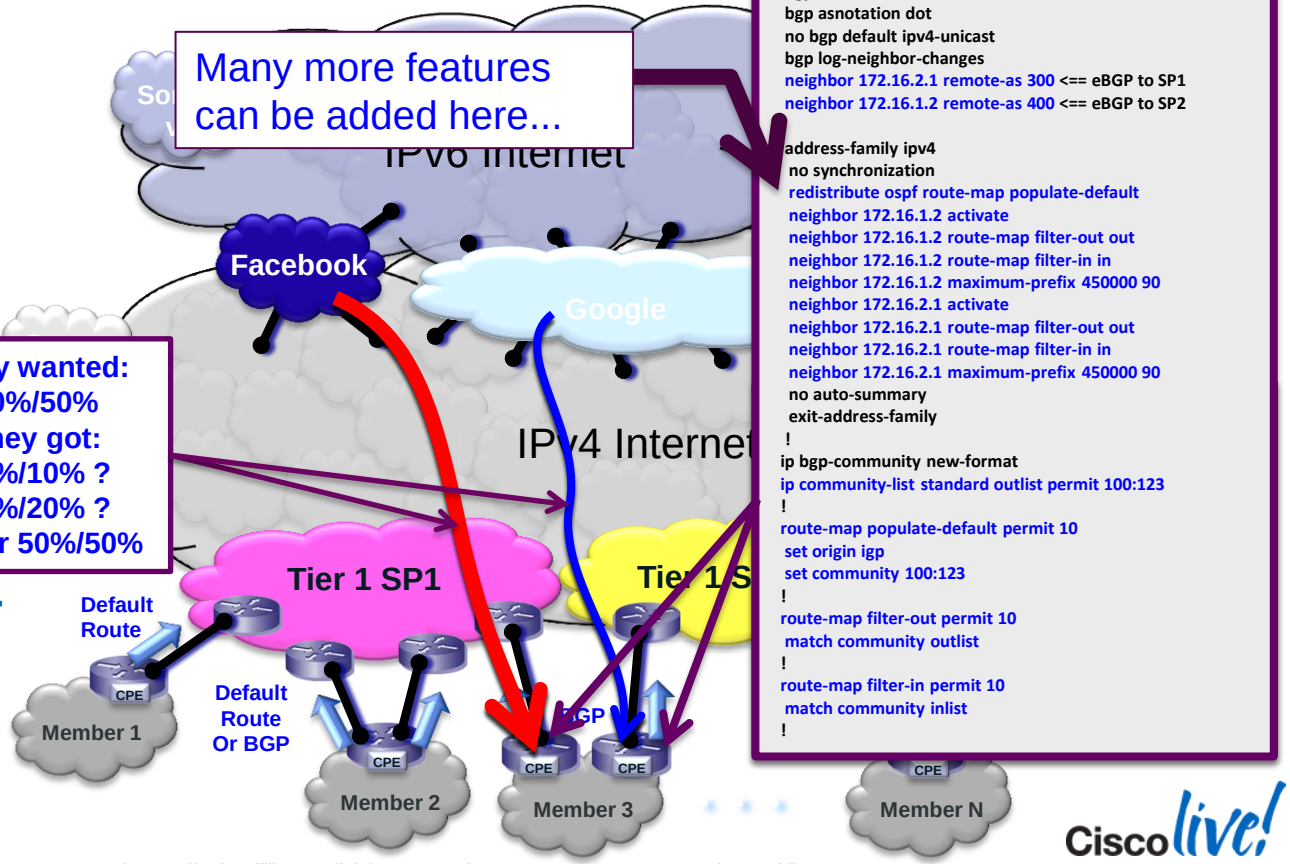

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

Efficient Multi-Homing and Multi-AF -- Customer Example

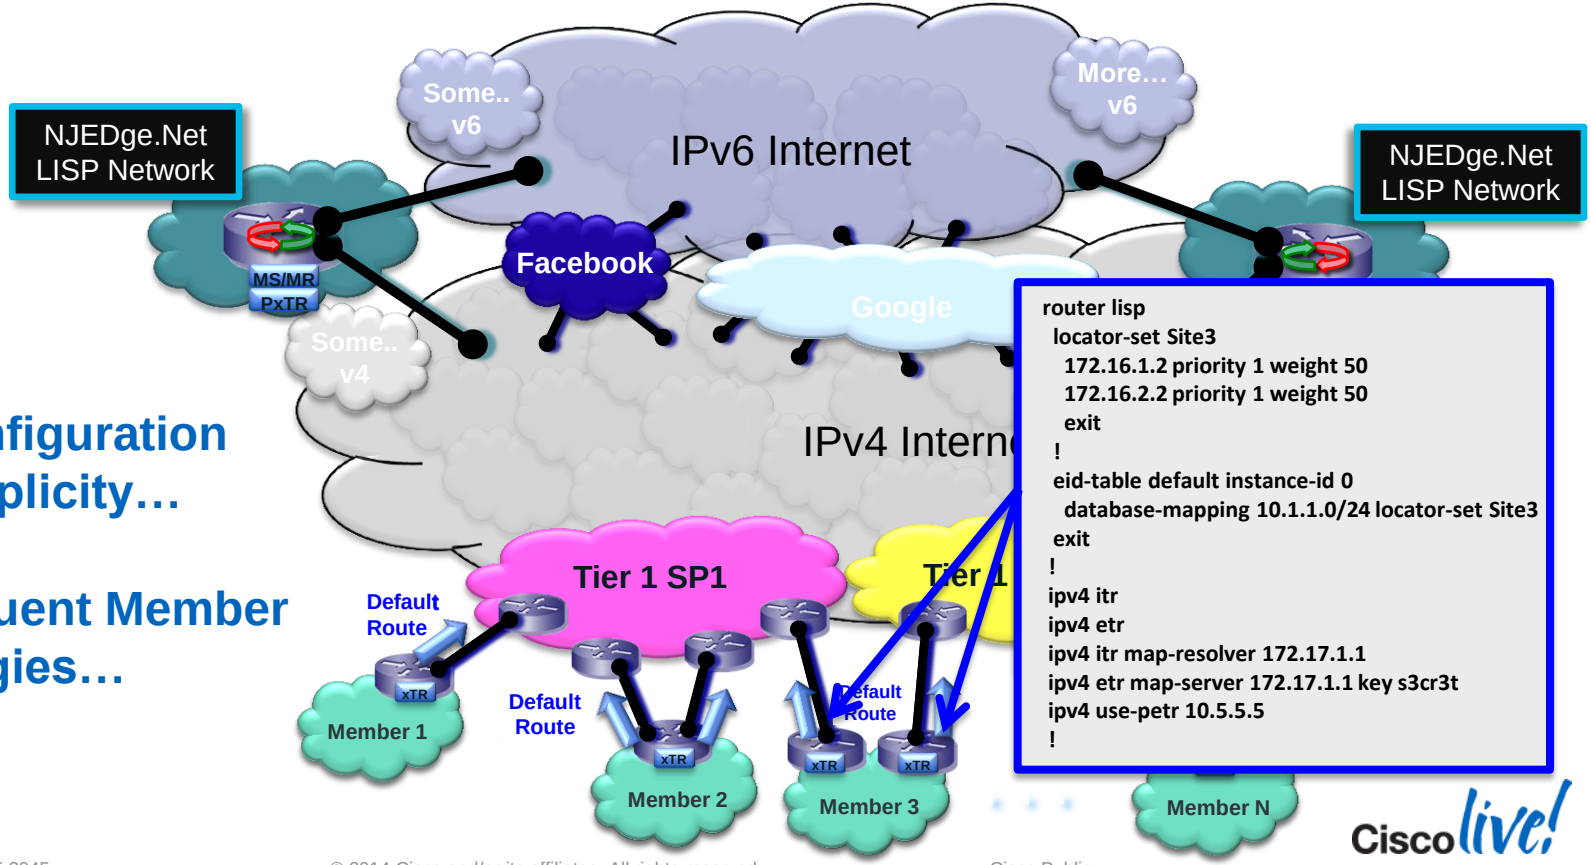

- **Configuration simplicity…**

Constituent Member Topologies…

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

Efficient Multi-Homing and Multi-AF -- Customer Example

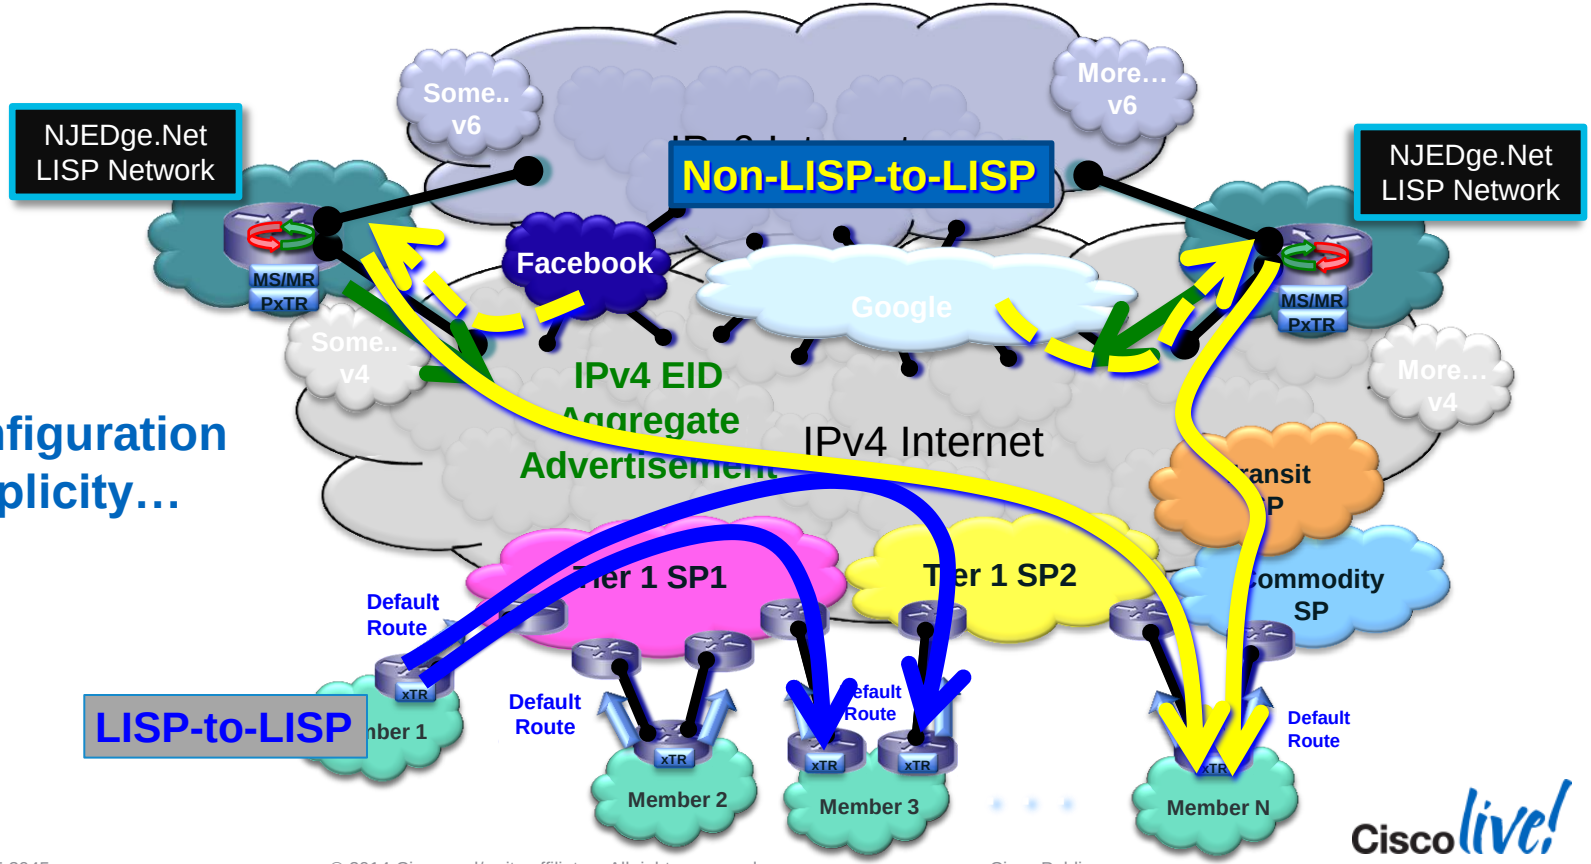

- **Configuration simplicity…**

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

NJEDge.Net

Efficient Multi-Homing and Multi-AF -- Customer Example

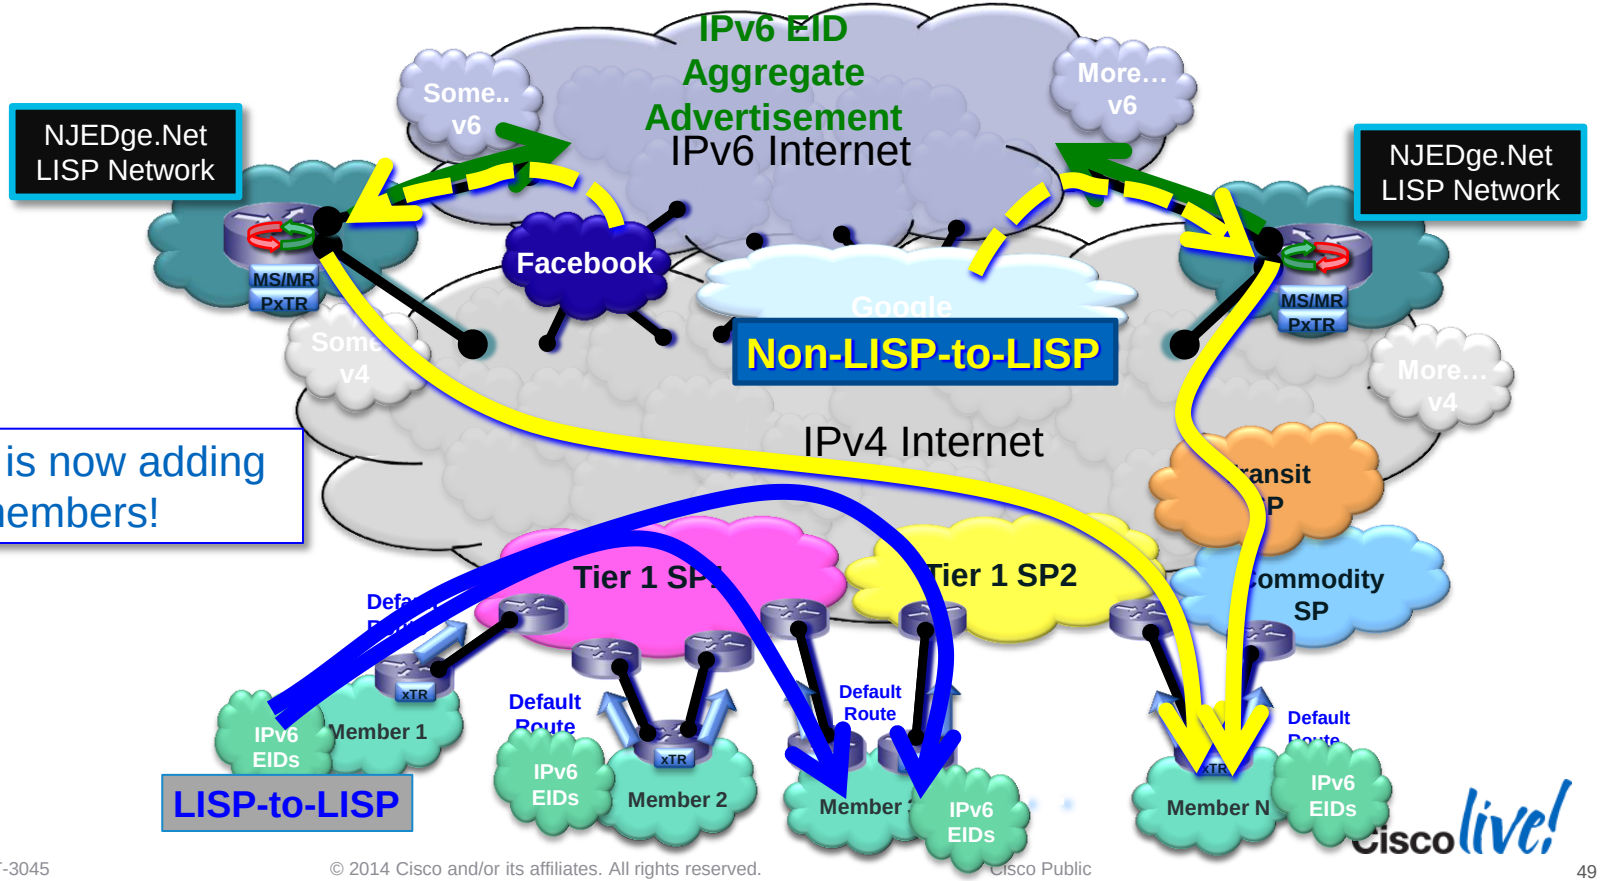

NJEDge.Net is now adding IPv6 for its members!

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public
Cisco live!

LISP Use Cases :: Multihoming and Multi-AF

NJEDge.Net

Customer Example :: NJEDge.Net

Key NJEDge.Net LISP Equipment

- ✓ ASR1Ks as MSMRs
- ✓ ASR9Ks as PxTRs (90G Internet capacity)

Key LISP Benefits

- ✓ No BGP to configure or manage
- ✓ No complex configurations
- ✓ Optimised Ingress load balancing
- ✓ Cost Savings by reducing OPEX and CAPEX
- ✓ LISP offers non disruptive transition approach which does not affect end system and allows for incremental deployment
- ✓ Disaster Recovery for Critical Applications introduces Increased Complexity

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

LISP and MPLS Integration

- LISP / MPLS results in an “ideal” deployment environment
  - Locator/ID split idealises a pure “RLOC core” and “EID overlay”
- Opportunities
  - IPv4 over MPLS via LISP

    Use of LISP (v4-over-v4) removes Customer IPv4 Prefixes from MPLS

    PE benefits :: (a) (substantially) improved scaling

    (b) reduced CPU load due to customer route advertisement churn
  - IPv6 over MPLS via LISP

    Use of LISP (v6-over-v4) removes SP from Customer IPv6 config/mgmt

    Immediate support :: even if not running LISP for IPv4

    PE benefits :: (a) no added v6 interface

    (b) no added v6 eBGP peering

    (c) no added IPv6 customer prefixes
  - Permits Inter-Departmental VPNs without additional PE VRFs

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

LISP and MPLS Integration

1: Existing IPv4 MPLS

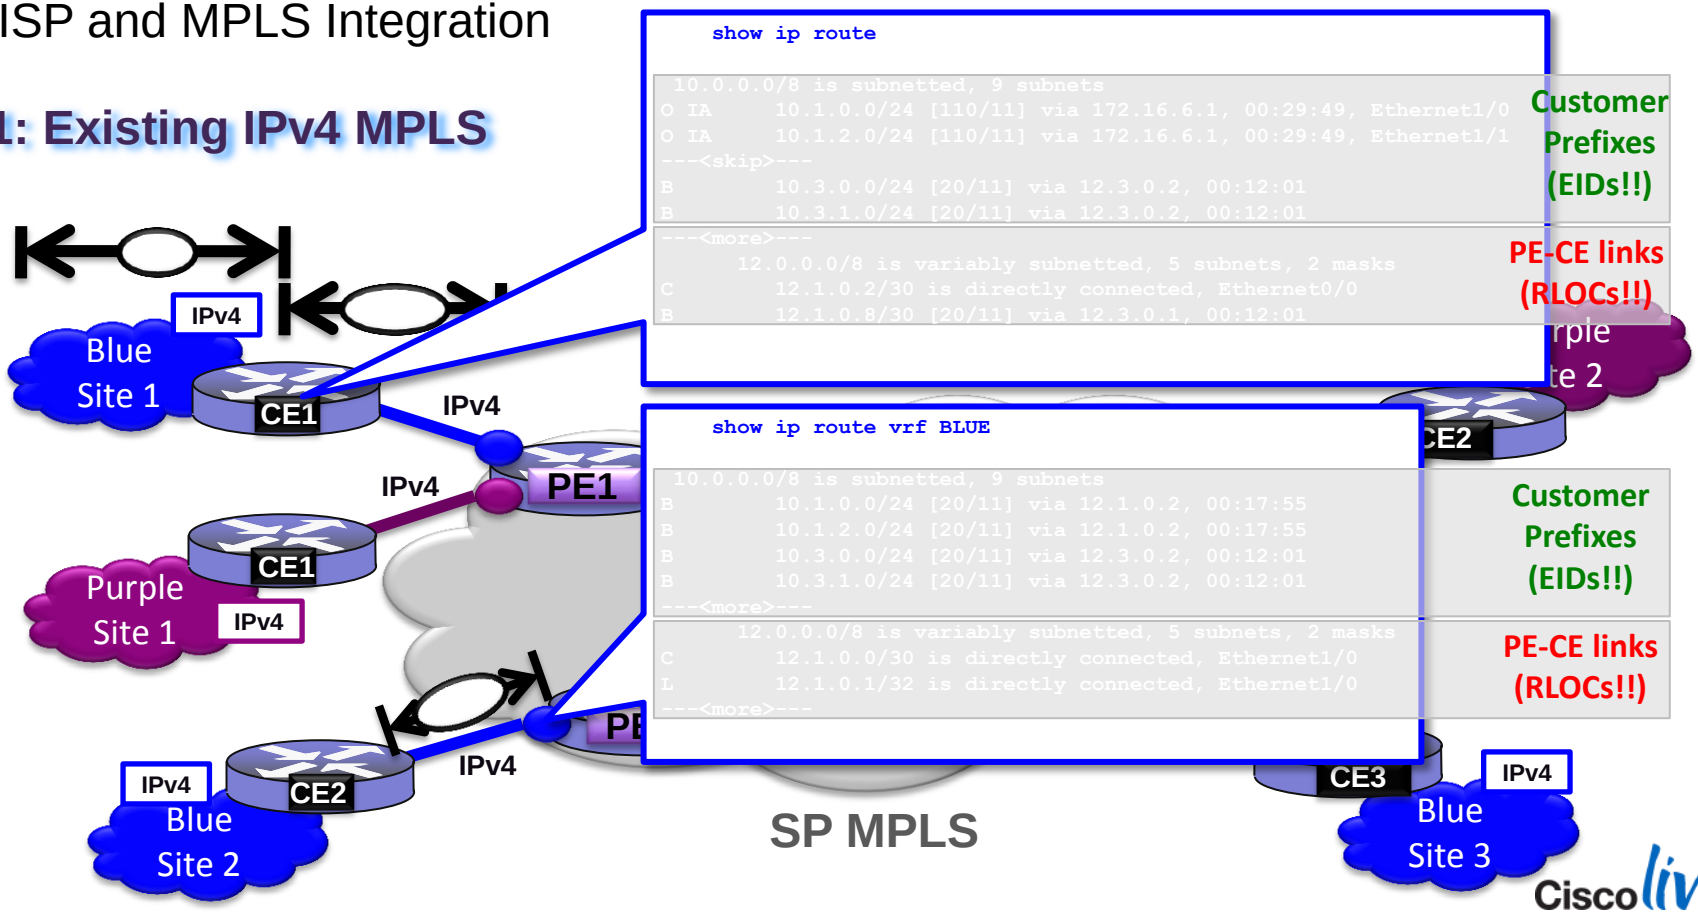

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

LISP and MPLS Integration

1: Existing IPv4 MPLS – Add LISP!

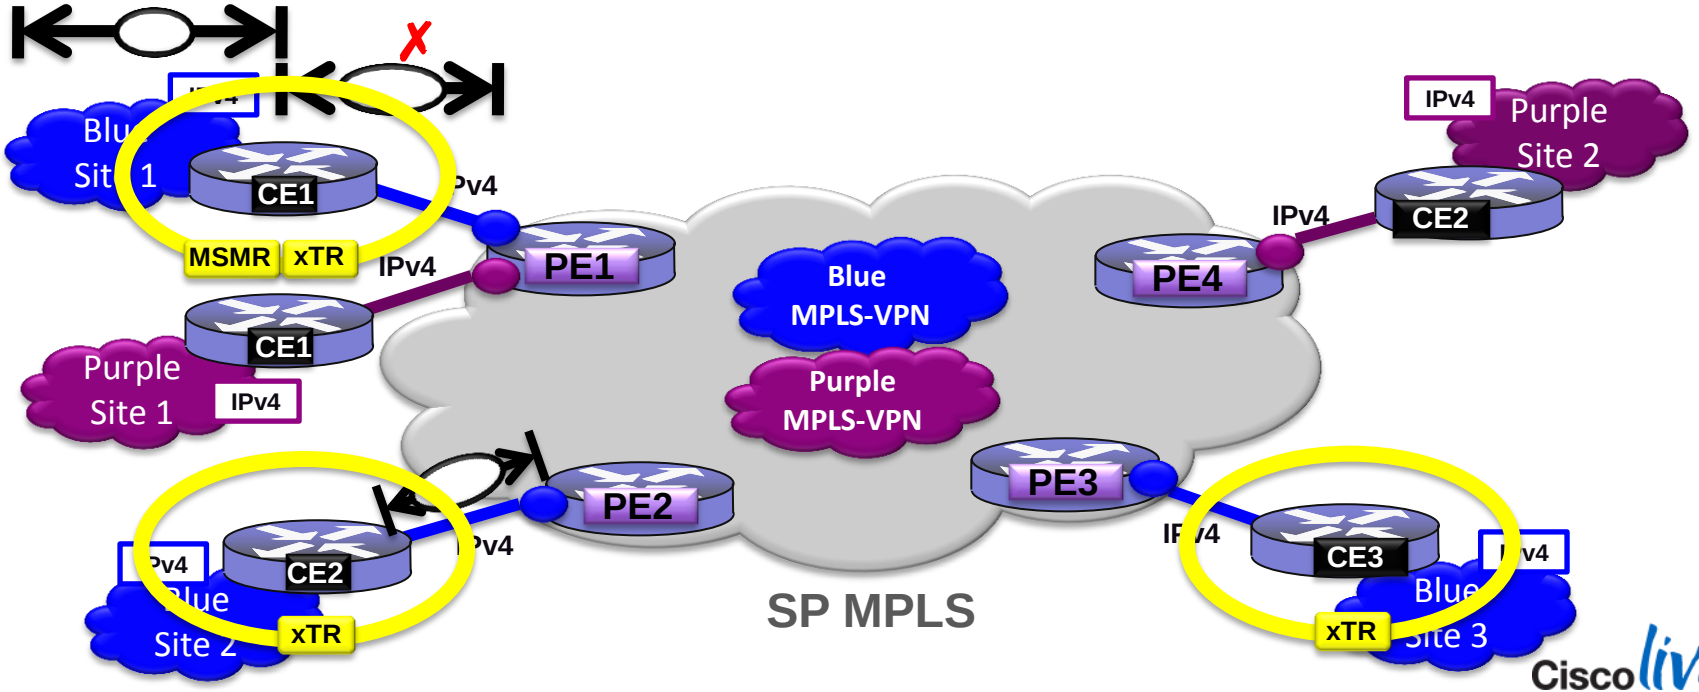

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

LISP and MPLS Integration

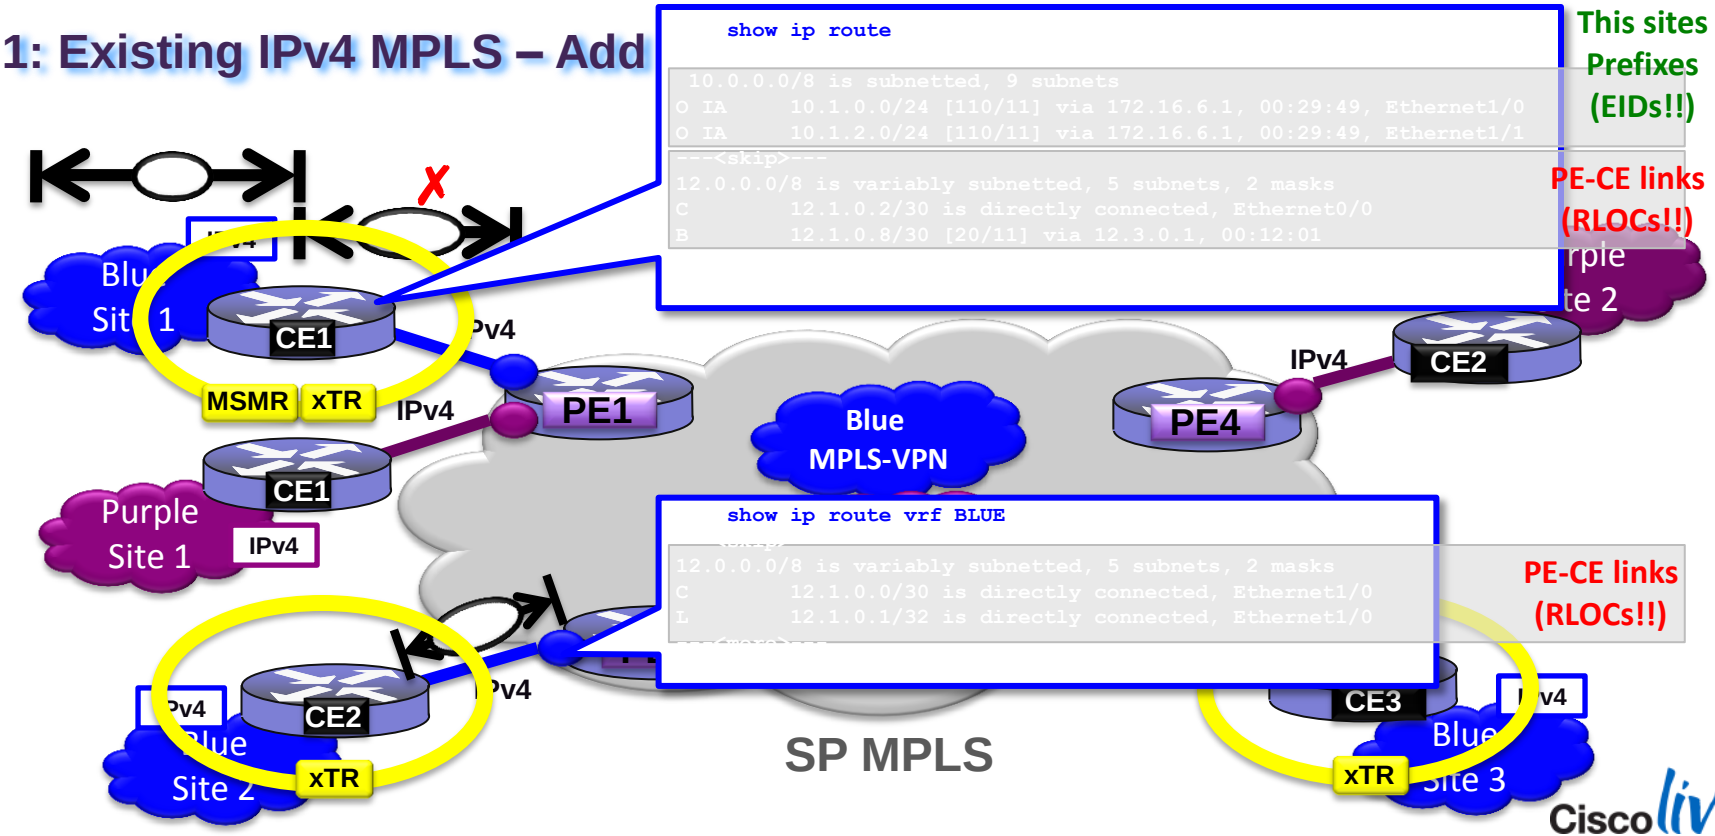

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

LISP and MPLS Integration

1: Existing IPv4 MPLS – Add

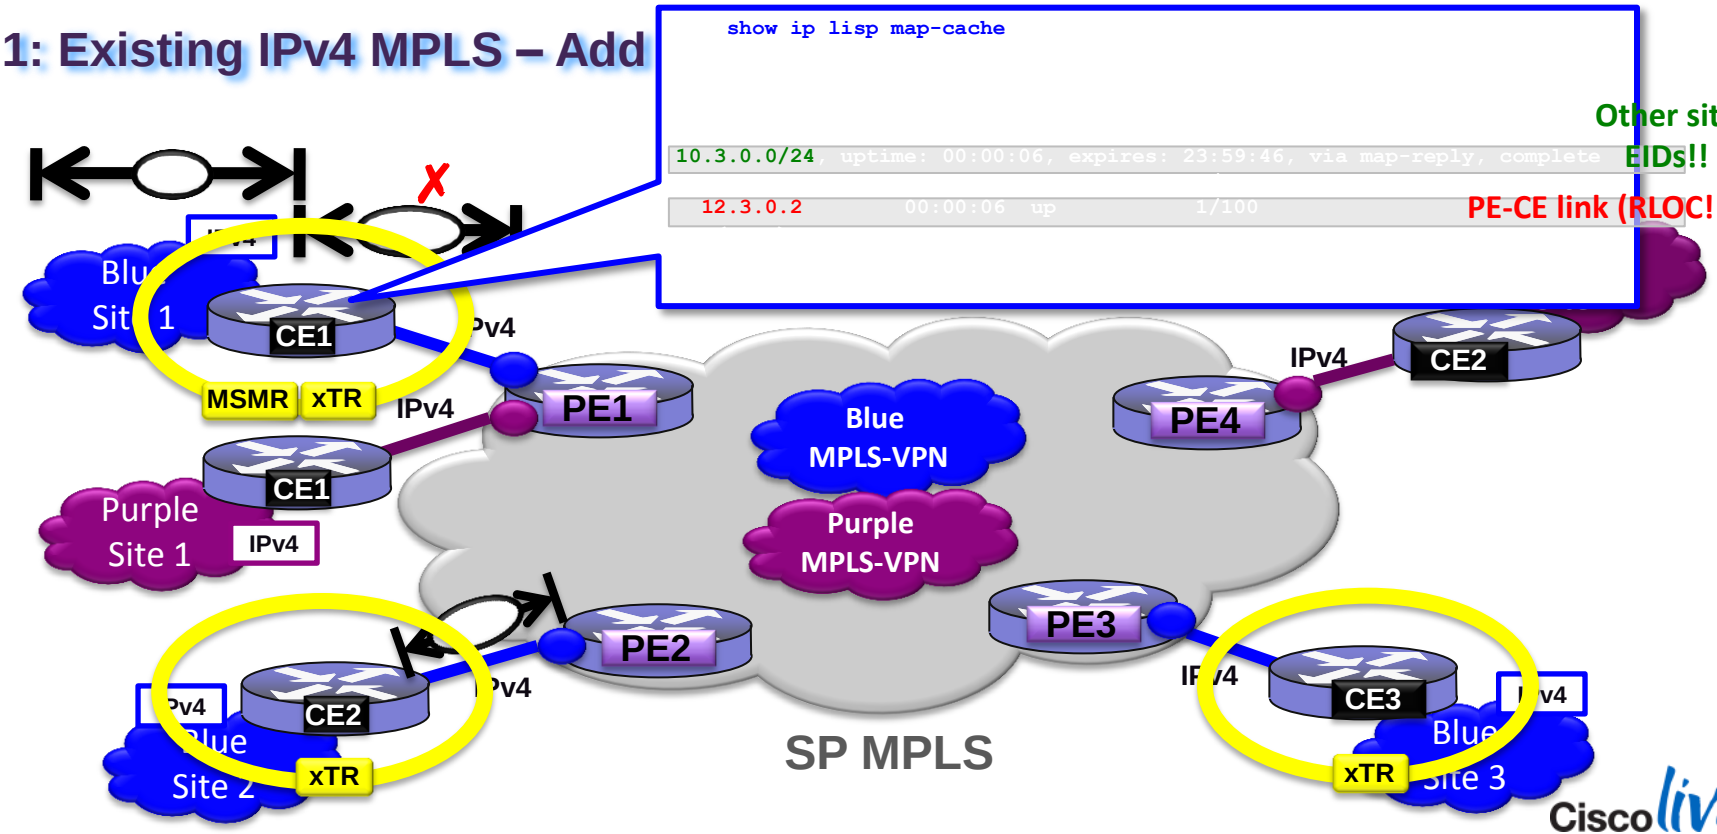

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

LISP and MPLS Integration

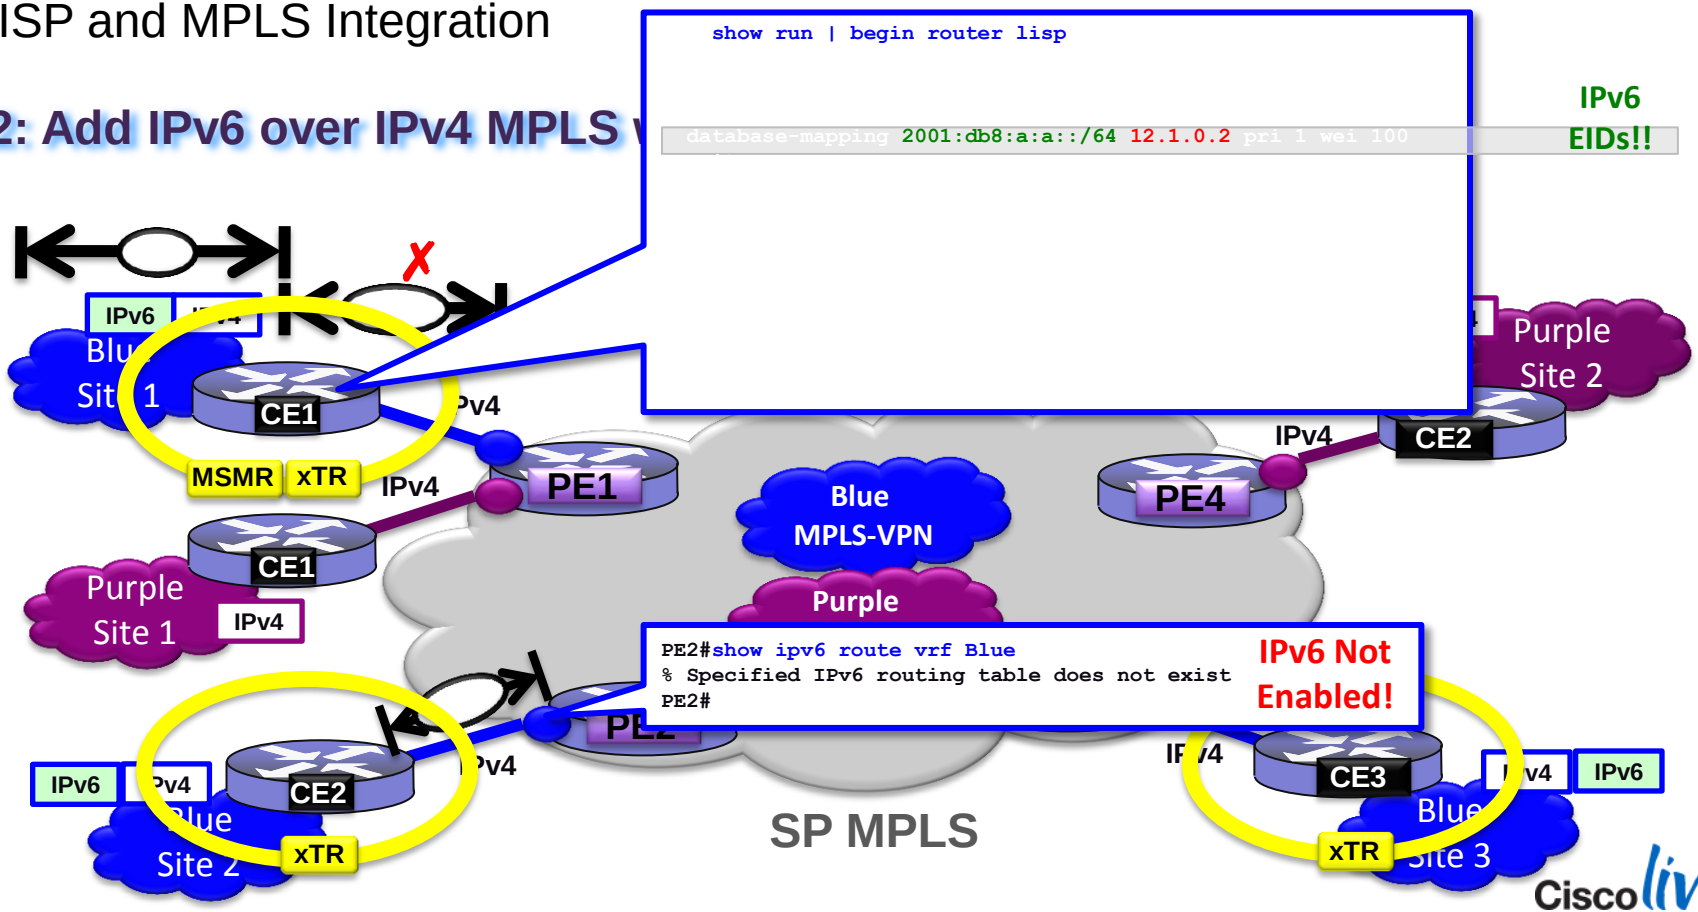

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Multihoming and Multi-AF

LISP and MPLS Integration

2: Add IPv6 over IPv4 MPLS

```
ping 2001:db8:b:b::1 so 2001:db8:a:a::1

show ipv6 lisp map-cache
```

Other site EIDs!!

```
2001:DB8:B:B::/64  uptime: 00:01:17, expires: 23:58:36, via map-reply, complete
  12.3.0.2         00:00:06  up          1/100
```

PE-CE links RLOCs!!

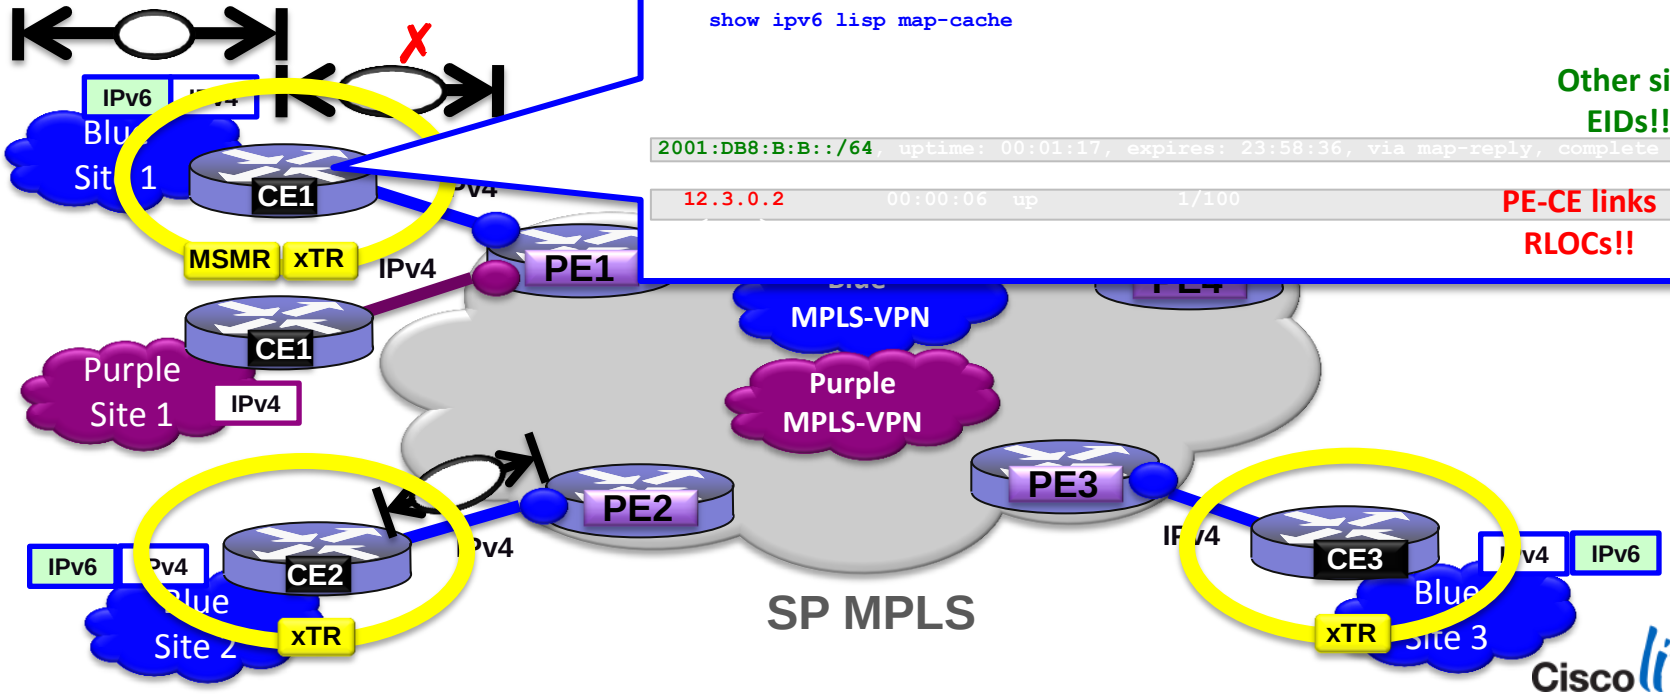

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Use Cases :: Additional “Useful” Features

Disjointed Locator Space Support

- Locator/ID separation creates two namespaces: EIDs and RLOCs
  - EID space is the overlay of Enterprise prefixes
  - RLOC space is the underlay network connectivity
- The fundamental principal of any network is that connectivity must exist between sites
- LISP supports sites being connected to locator spaces that have no connectivity to each other!
  - In LISP, this is known as a “disjointed RLOC set”

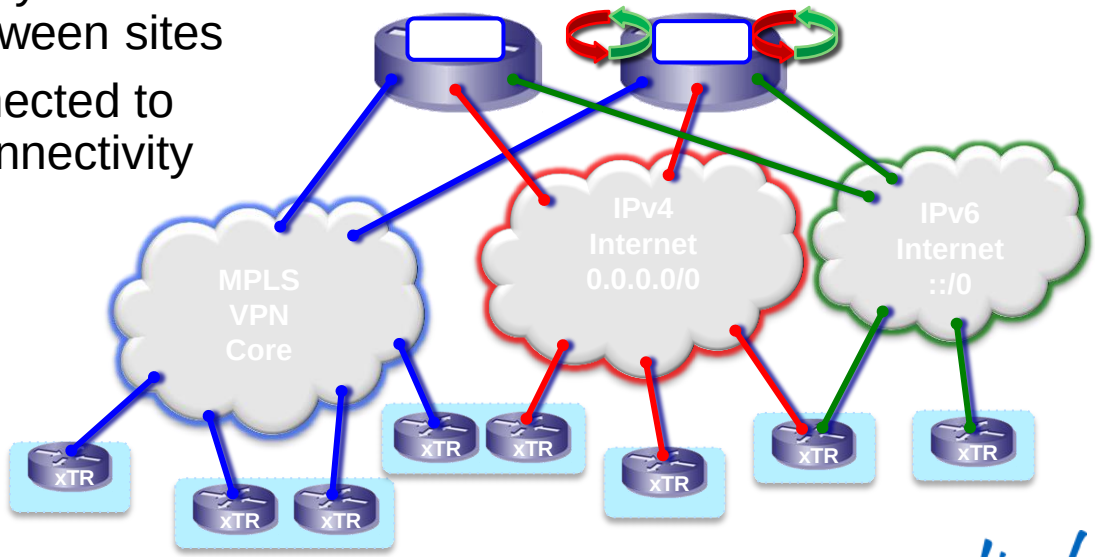

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Use Cases :: Additional “Useful” Features

Disjointed Locator Space Support

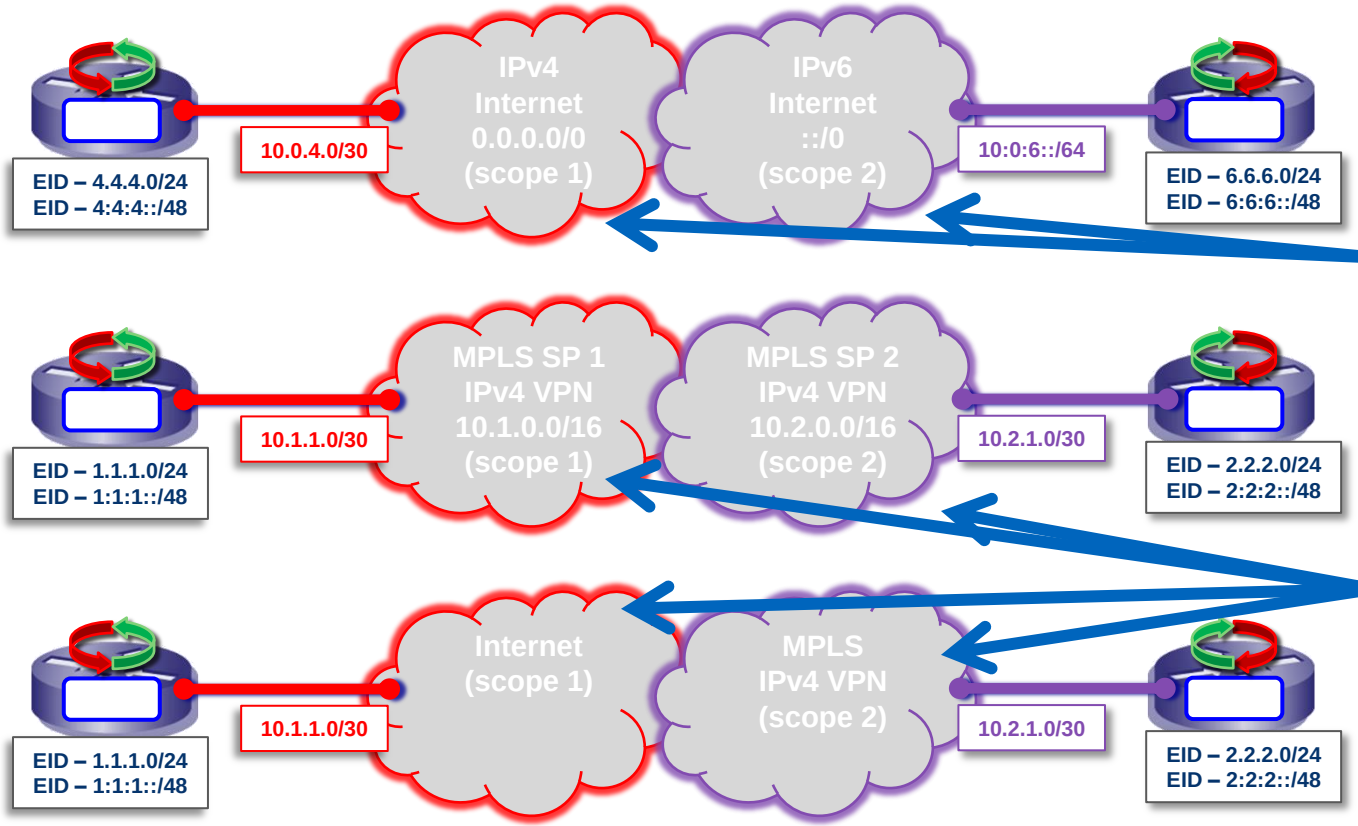

One obvious example of disjointed RLOC spaces is for IPv4 and IPv6 attached sites

The same situation occurs for distinct core networks of the same address family. Two MPLS VPN cores, for example, exhibit disjointed RLOC properties.

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Use Cases :: Additional “Useful” Features

Disjointed Locator Space Support

```
!
router lisp
 locator-set rtr-set1
  10.1.3.1 priority 1 weight 1
  exit
 !
 locator-set rtr-set2
  10.2.3.1 priority 1 weight 1
  exit
 !
 locator-scope s1
  rtr-locator-set rtr-set1
  rloc-prefix 10.1.0.0/16
  exit
 !
 locator-scope s2
  rtr-locator-set rtr-set2
  rloc-prefix 10.2.0.0/16
  exit
 !
---<etc.>---
```

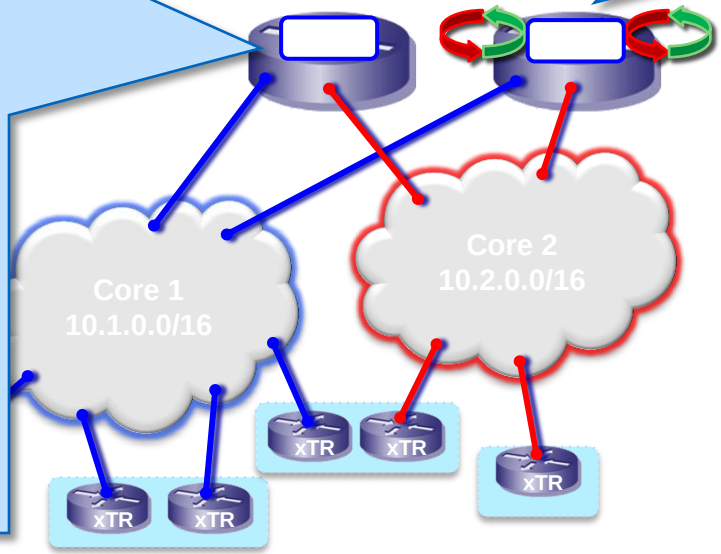


```
!
router lisp
 locator-set setALL
  10.1.3.1 priority 1 weight 1
  10.2.3.1 priority 1 weight 1
  exit
 !
 map-request itr-rlocs setALL
 eid-table default instance-id 0
  map-cache 0.0.0.0/0 map-request
  map-cache ::/0 map-request
  exit
 !
---<etc.>---
```


© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Deployment Examples

LISP Deployment Examples…

1. Efficient Multihoming and Multi-AF (IPv4 and IPv6)
2. Efficient Virtualisation and High-Scale VPNs
3. Data Centre/Host Mobility
4. LISP-Mobile Node

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

LISP+GETVPN Config Guide:
http://lisp.cisco.com

Efficient Virtualisation and High-Scale VPNs – Overview

- Needs:
  - Integrated Segmentation
  - Global scale and interoperability
  - Minimal Infrastructure disruption
- LISP Solution:
  - 24-bit LISP Instance-ID segments control plane and data plane, with VRF binding to the Instance-ID
- Benefits:
  - Very high scale tenant segmentation
  - Global mobility + high scale segmentation integrated in single IP solution
  - IP-based “overlay” solution, transport independent
  - No Inter-AS complexity

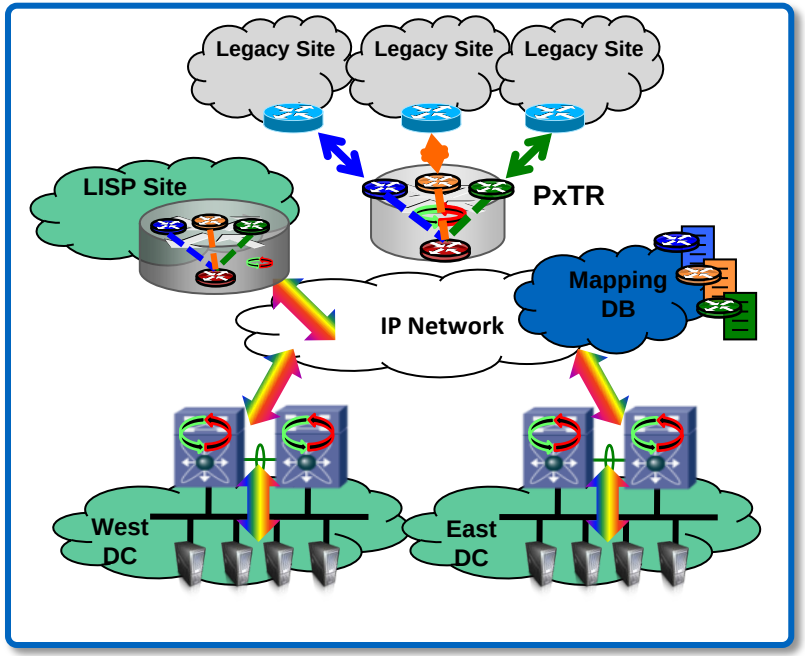

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Virtualisation/VPNs

Efficient Virtualisation/Multi-Tenancy Support – Concepts

- Virtualisation of the DEVICE level
  - Virtual Routing and Forwarding (VRF) tables segment Layer 3 routing tables
  - VRFs are used to virtualise the component resources
  - Virtualisation secures movement of traffic between networks and enhances security policy options

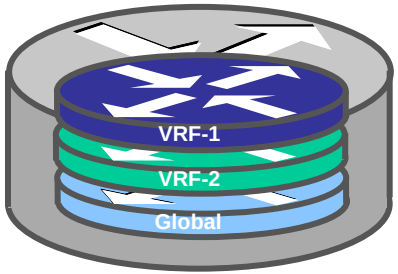

- Virtualisation of the PATH level
  - VRFs assist in path isolation
  - Single-hop (hop-by-hop)
  - Multi-hop (over-the-top)

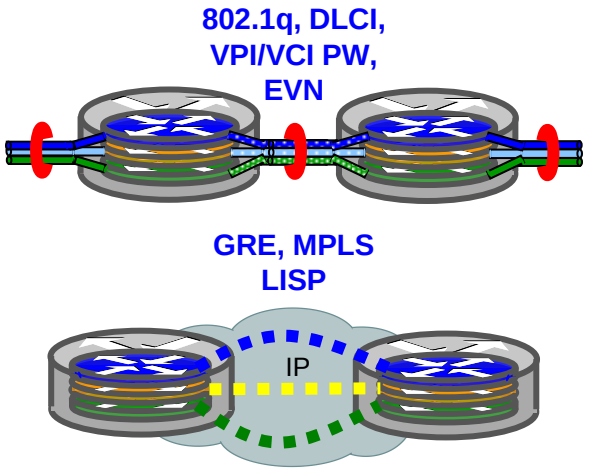

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Virtualisation/VPNs

LISP Virtualisation/Multi-Tenancy Support – Concepts

- Recalling that… LISP is **"Locator/ID" separation**… and creates two namespaces: EIDs and RLOCs… LISP can virtualise both EID and RLOC namespaces, or both!
- Two models of operation are defined: **Shared** and **Parallel**
  - **Shared Model Virtualisation**:
    - Virtualises the EID namespaces
    - Binds an EID namespace privately defined using a VRF to an Instance-ID
    - Uses a common (shared) RLOC (locator) address space
    - The Mapping System is also part of the locator namespaces and is shared
  - **Parallel Model Virtualisation**:
    - Virtualises the RLOC (locator) namespaces
    - One or more EID instances may share a virtualised RLOC namespace
    - A Mapping System must also be part of each locator namespaces

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Virtualisation/VPNs

LISP Virtualisation/Multi-Tenancy Support – Concepts

- **RLOC virtualisation** is enabled in conjunction with **locator table** VRFs
- **EID virtualisation** uses **LISP Instance-IDs** in conjunction with **EID** VRFs
  - **Instance-IDs** maintain address space segmentation in control plane and data plane
  - **Instance-IDs** are numerical tags defined in LISP Canonical Address Format (LCAF)
    - IID: a 24-bit unstructured number
    - Data Plane: IID is included in LISP encapsulation header
    - Control Plane: IID is encoded with the EID in LCAF header

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Virtualisation/VPNs

LISP Virtualisation/Multi-Tenancy Support – Concepts

- Default (non-Virtualised) Model – at the device level
  - Conceptually, the Default Model is just a single Parallel Model instance
  - All EID lookups are also in the same single table – default
  - Thus, EIDs are associated with Instance-ID 0
  - All RLOC lookups are in a single table – default
  - The Mapping System is part of the locator address space

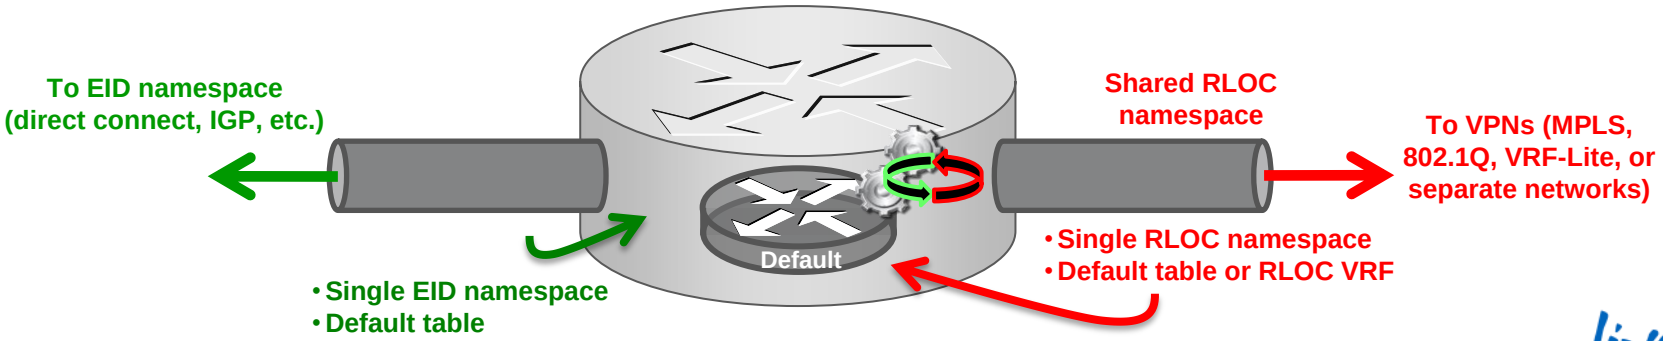

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Virtualisation/VPNs

LISP Virtualisation/Multi-Tenancy Support – Concepts

- Shared Model – at the device level
  - Multiple EID-prefixes are allocated privately using VRFs
  - EID lookups are in the VRF associated with an Instance-ID
  - All RLOC lookups are in a single table – (default/global or RLOC VRF)
  - The Mapping System is part of the locator address space and is shared

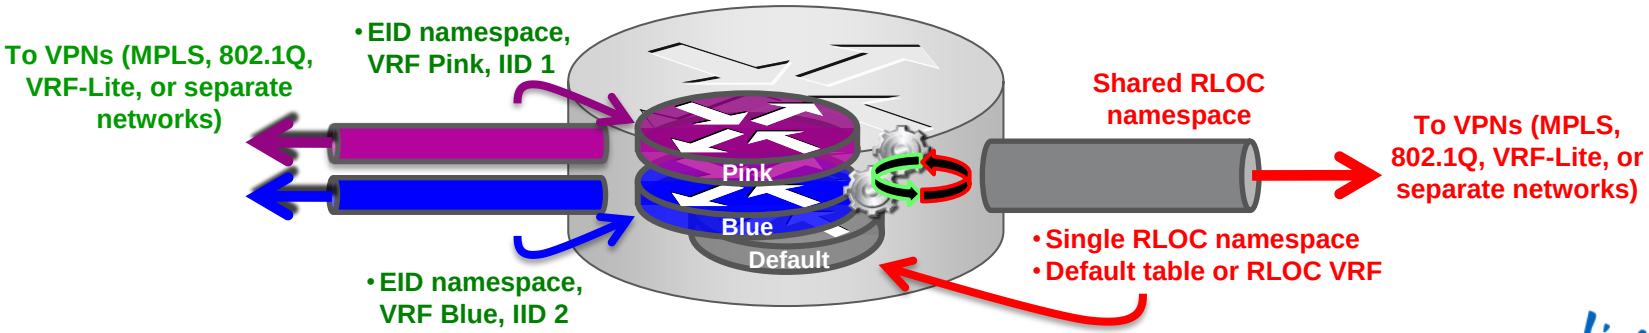

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Virtualisation/VPNs

LISP Virtualisation/Multi-Tenancy Support – Concepts

- Parallel Model – at the device level
  - Multiple EID-prefixes are allocated privately using VRFs
  - EID lookups are in the VRF associated with an Instance-ID
  - RLOC lookups are in the VRF associated with the locator table
  - A Mapping System must be part of each locator address space

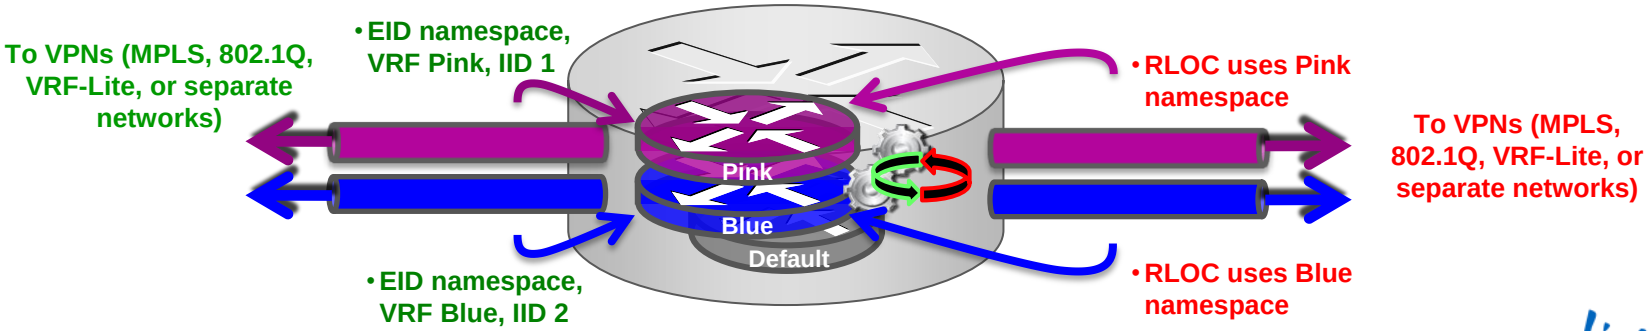

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Virtualisation/VPNs

LISP Virtualisation/Multi-Tenancy Support – Concepts

- Shared and Parallel Models Combined – at the device level
  - Multiple “Shared Model” instantiations combined with Multiple “Parallel Model” instantiations
  - Multiple EID VRFs bound to a single RLOC VRF
  - Multiple RLOC VRFs on the same device

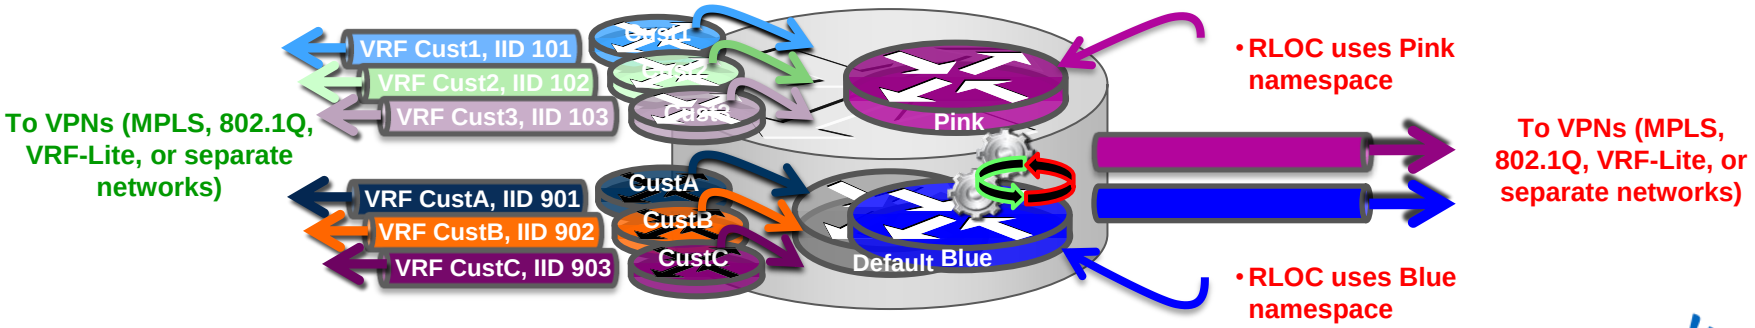

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

Efficient Virtualisation and High-Scale VPNs – Overview

All VPNs share a set of common requirements

1. Encapsulation:

- **– Virtualisation**
  - **EID prefix virtualisation**
    - **Tied to EID VRFs**
  - **Locators can be virtualised too**

2. Site to Site Routing:

- **– Spoke to spoke connectivity**
- **– Optional local Internet offload (split-tunnel)**
- **– No IGP required to branch sites!**

LISP VPNs Routing? or Tunnelling? -- It's BOTH!

3. Security: Built-in and Add-on

- **– Built-in security mechanisms**
- **– LISP Works with any crypto scheme**
  - **Locators or EIDs can be encrypted**
- **– LISP-SEC for control plane security**

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

Efficient Virtualisation and High-Scale VPNs – Overview

LISP – Inherently scalability and virtualisation, rapidly deployable

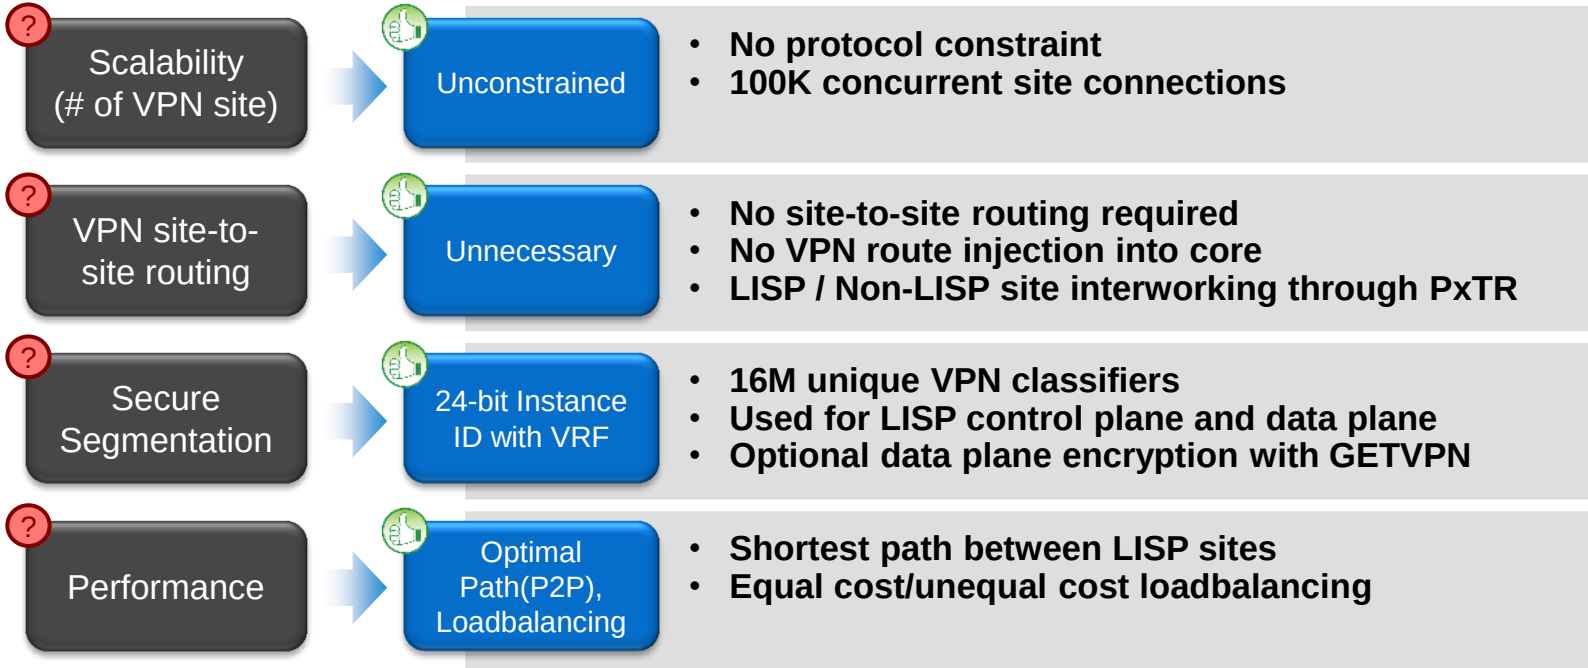

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

Efficient Virtualisation and High-Scale VPNs – Overview

Generalised LISP Shared Model deployment

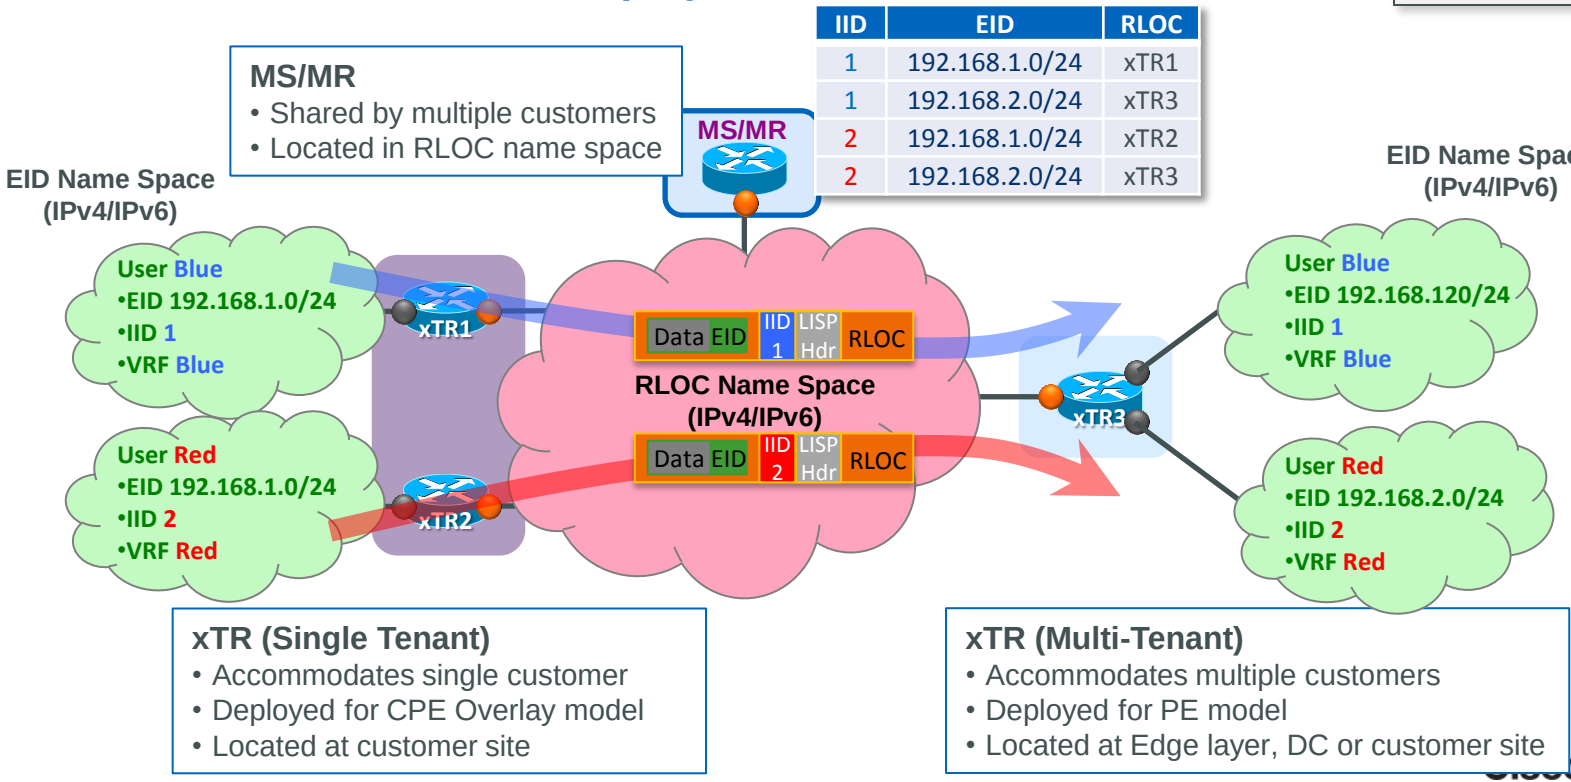

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Use Cases :: Virtualisation/VPNs

Customer Example :: US State Government (Multi-tenancy)

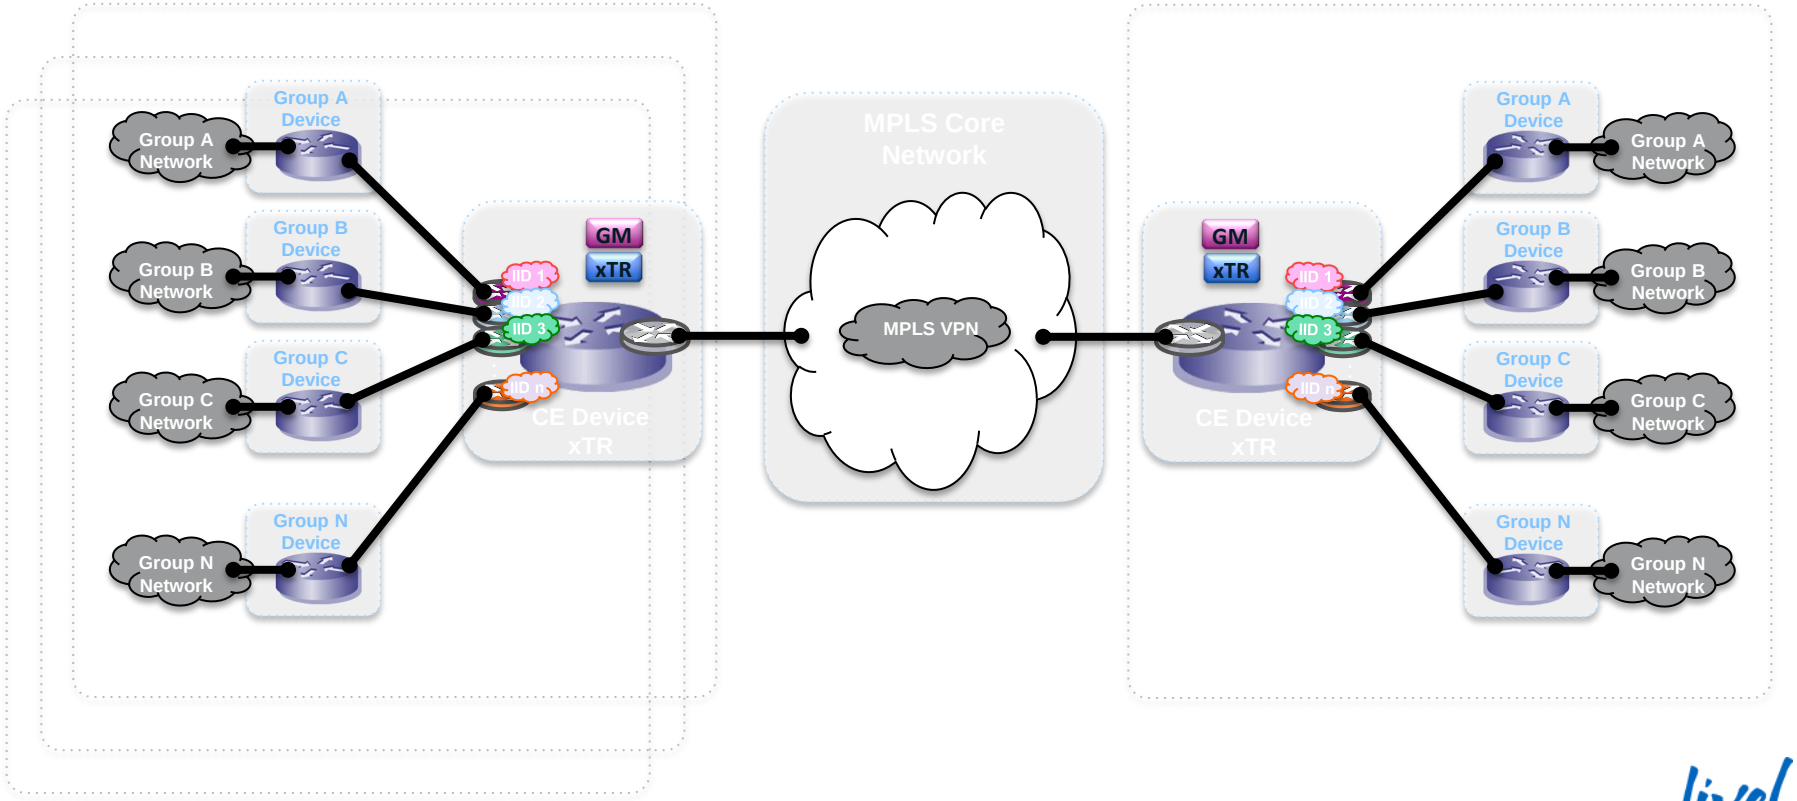

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Use Cases :: Virtualisation/VPNs

Customer Example :: US State Government (Multi-tenancy)

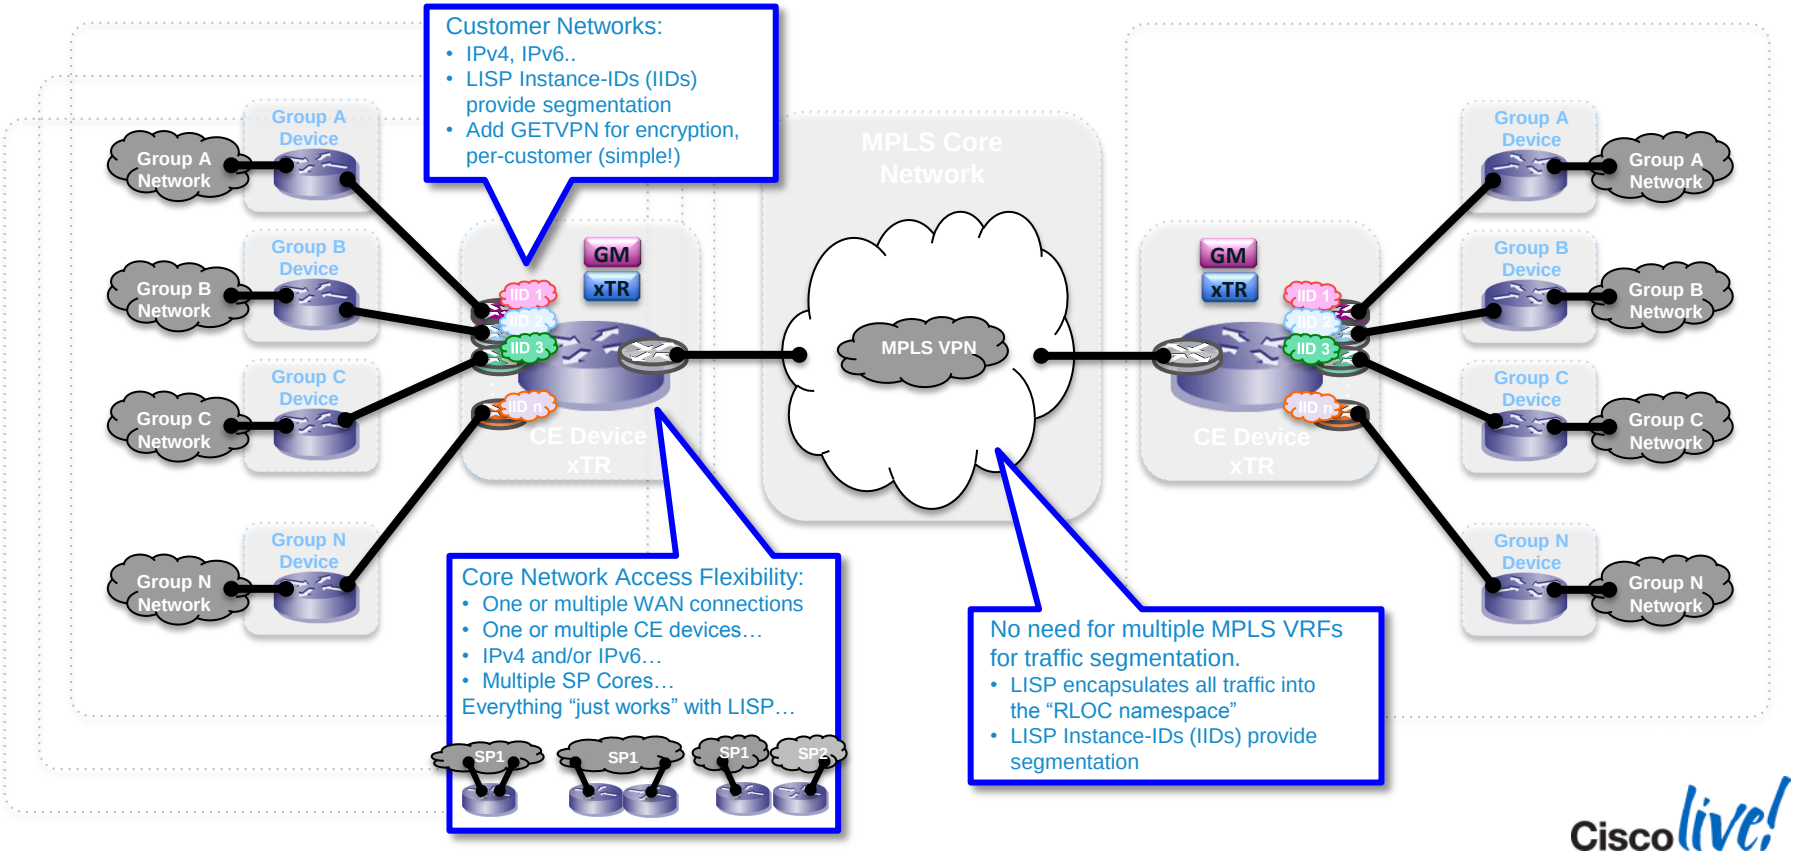

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Use Cases :: Virtualisation/VPNs

Customer Example :: US State Government (Multi-tenancy)

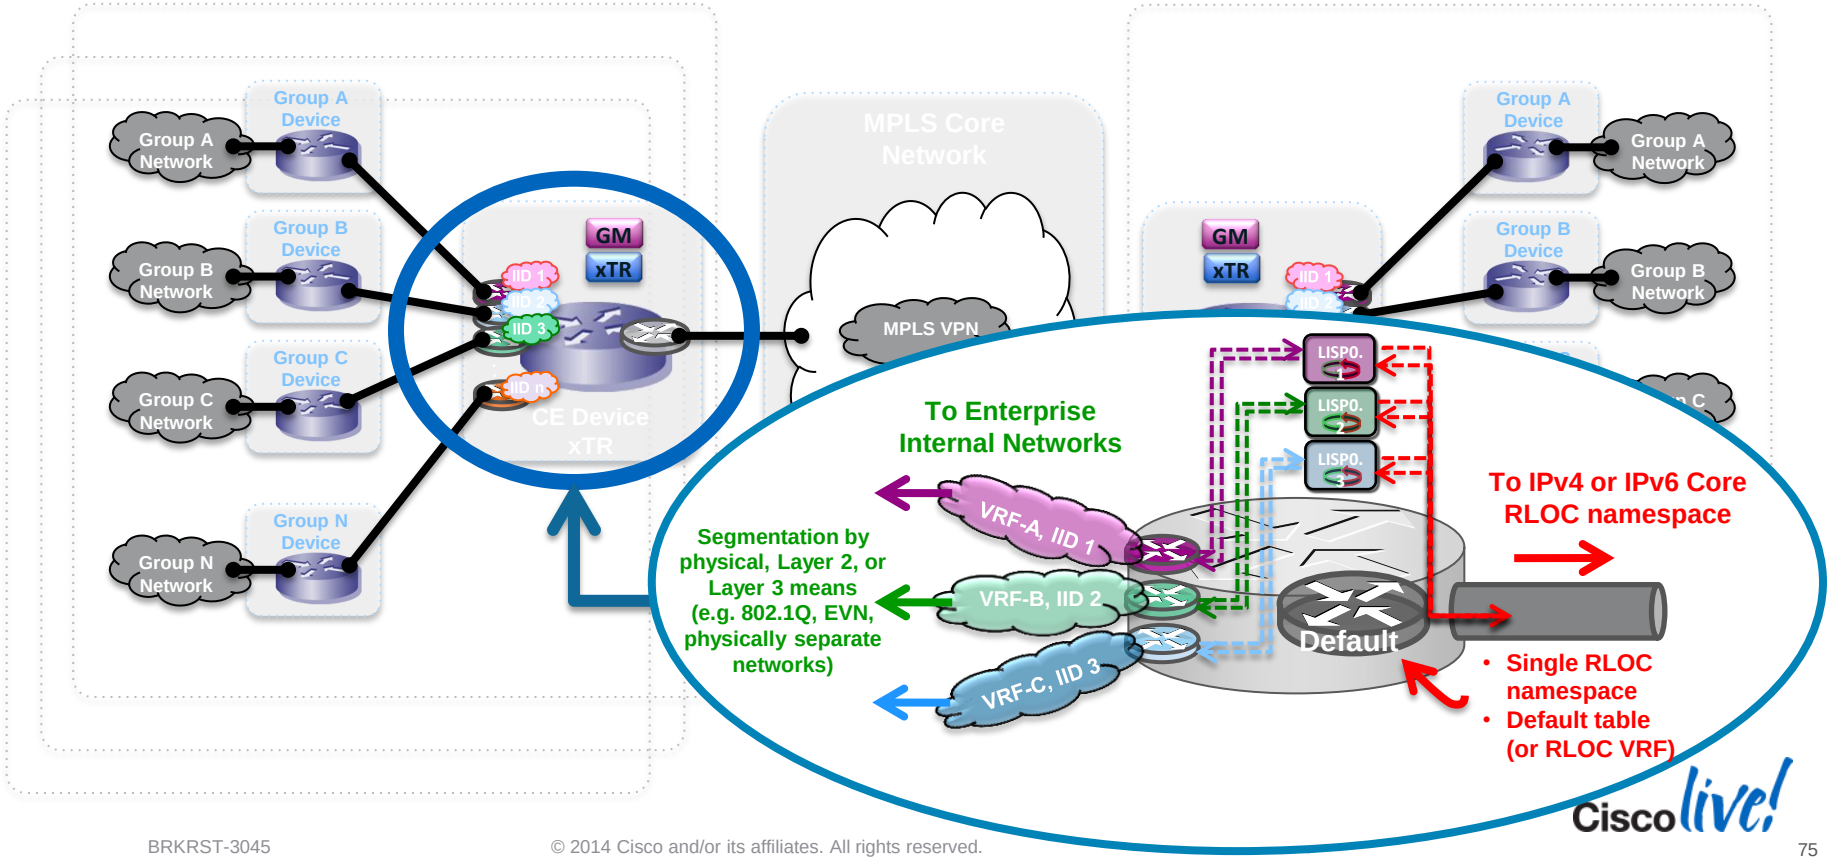

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Use Cases :: Virtualisation/VPNs

Customer Example :: US State Government (Multi-tenancy)

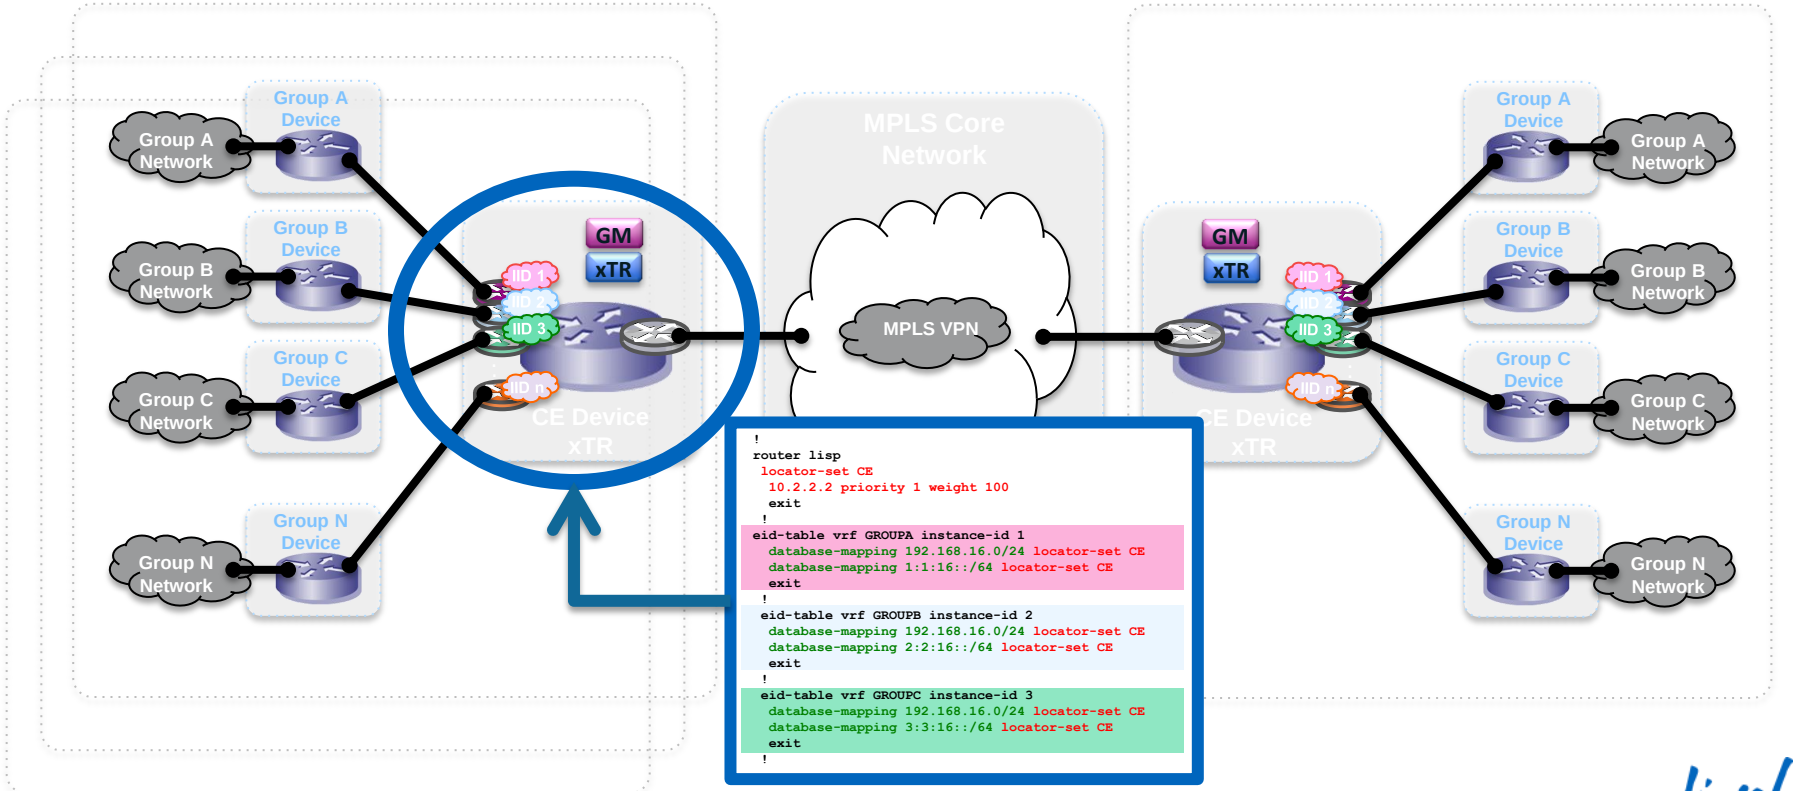

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Virtualisation/VPNs

LISP Virtualisation/Multi-Tenancy Support – Concepts

- LISP and encryption (IOS)
  - Recalling that… LISP is “Locator/ID” separation… and creates two namespaces: **EIDs** and **RLOCs**
  - LISP provides two ways to apply a crypto map

Use-Case		Vanilla IPsec	GETVPN	Comments
LISP Default Model	crypto-map on **RLOC**	✔	✔	LISP encap first, then encryption based on RLOC
	crypto-map on **LISP0**	✔	✔	Encryption first based on EID, then LISP encap
LISP Virtualisation	crypto-map on **RLOC**	✔	✔	LISP encap first, then encryption based on RLOC
	crypto-map on **LISP0.x**	✔	✔	Encryption first based on EID, then LISP encap

lisp.cisco.com

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Virtualisation/VPNs

LISP Virtualisation/Multi-Tenancy Support – Concepts

- Group Domain of Interpretation (GDOI) RFC 6407 – adding encryption

- GDOI
  - RFC 6407
  - “Stateless” IPsec
  - Traffic encryption keys computed on Key Server, distributed to all Group Members
  - Better scaling than vanilla IPsec

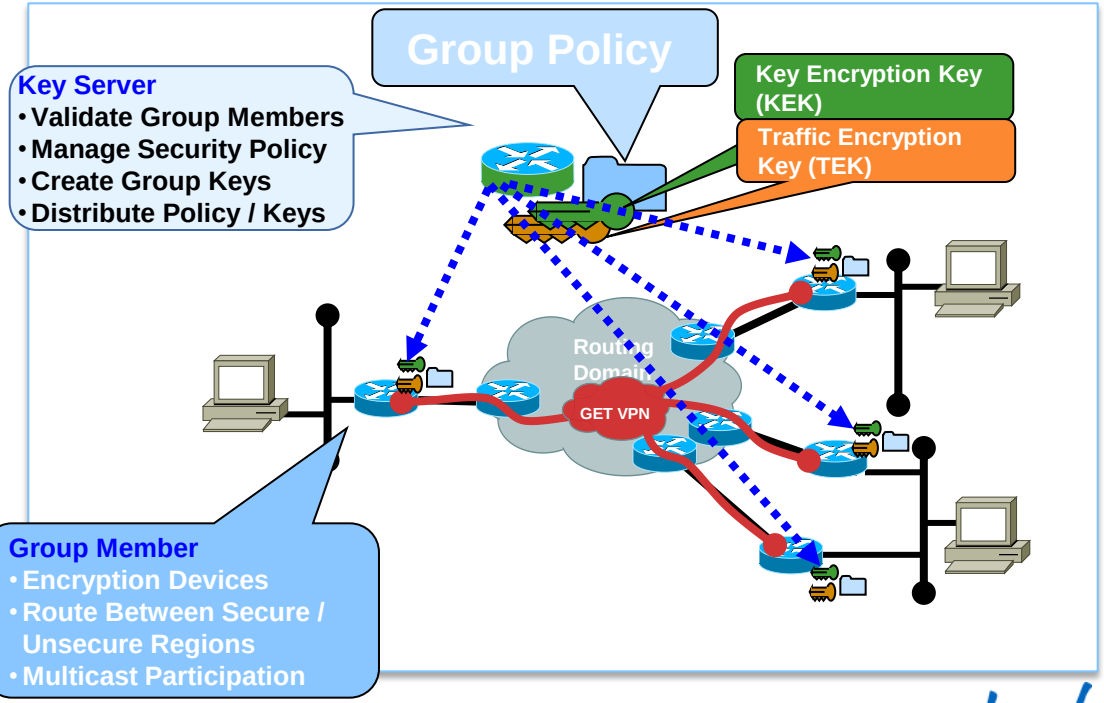

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

Efficient Virtualisation and High-Scale VPNs – Ov

3. Add encryption

- GETVPN Key Servers
- Nothing to do with LISP!

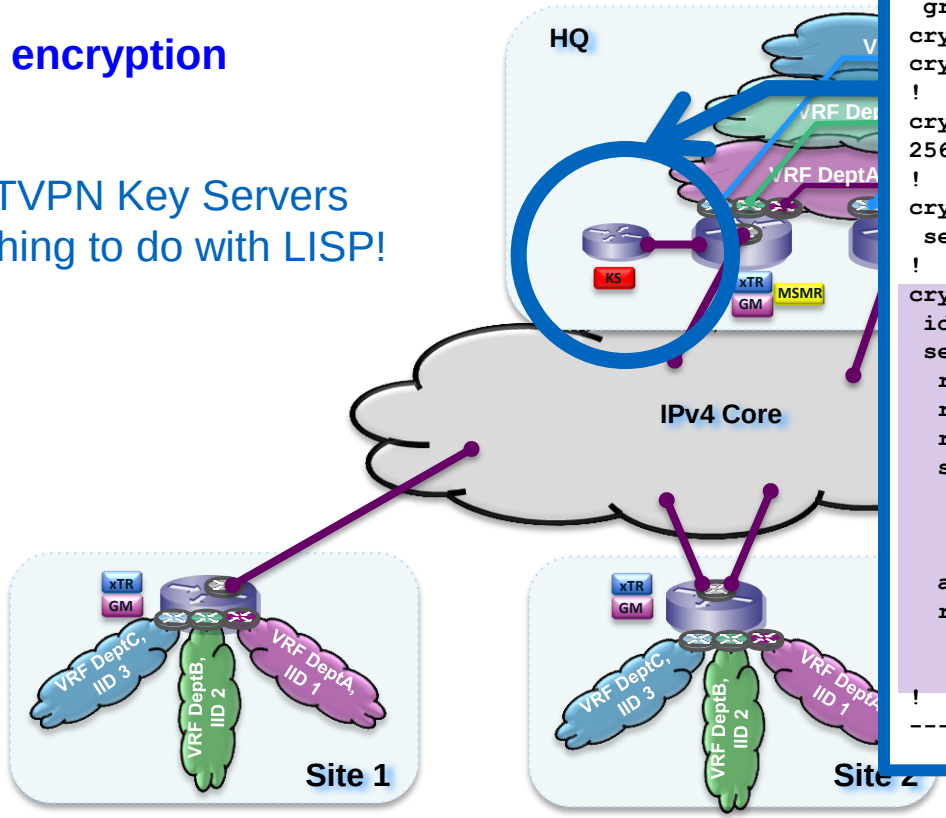

KS1

```
!
crypto isakmp policy 10
 encr aes 256
 authentication pre-share
 group 16
crypto isakmp key FOO address 0.0.0.0
crypto isakmp keepalive 15 periodic
!
crypto ipsec transform-set GDOI-TRANS esp-aes
256 esp-sha512-hmac
!
crypto ipsec profile GDOI-PROFILE
 set transform-set GDOI-TRANS
!
crypto gdoi group V4GROUP-0001
 identity number 10001
 server local
  rekey retransmit 60 number 2
  rekey authentication mypubkey rsa GET-KEYS1
  rekey transport unicast
  sa ipsec 1
   profile GDOI-PROFILE
   match address ipv4 GETVPN-0001
   replay time window-size 5
  address ipv4 192.168.18.2
  redundancy
   local priority 100
   peer address ipv4 192.168.19.2
!
---<cont.>---
```


© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

Efficient Virtualisation and High-Scale VPNs – Ov

3. Add encryption

- GETVPN Key Servers
- Nothing to do with LISP!

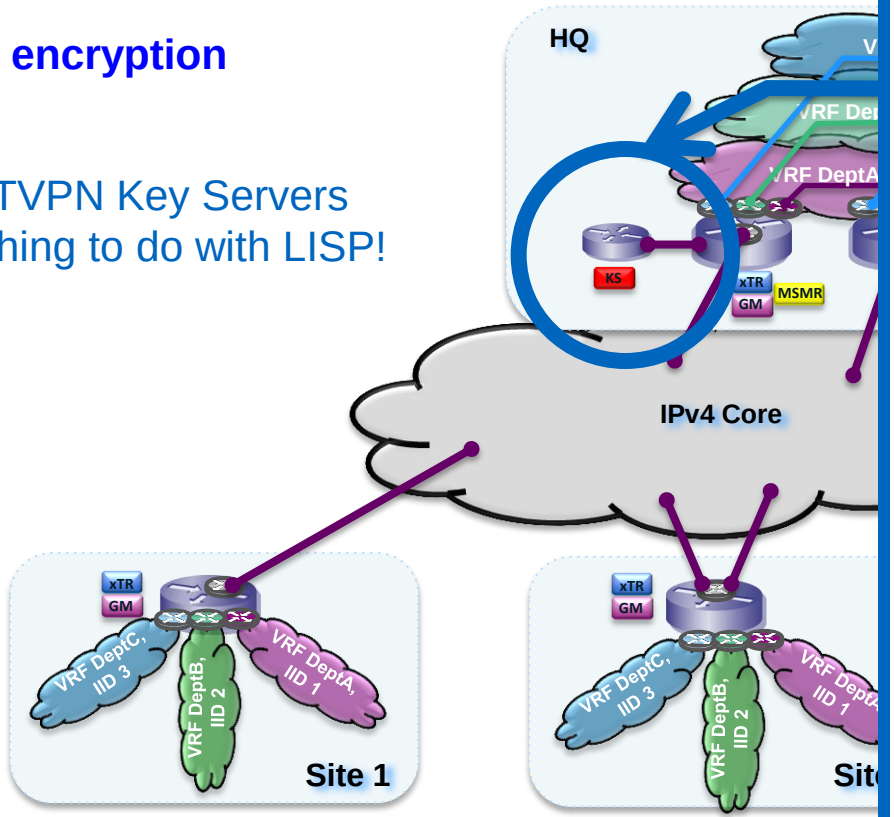

KS1

```
! ---<cont.>---
!
crypto gdoi group ipv6 V6GROUP-0003
 identity number 20003
 server local
  rekey retransmit 60 number 2
  rekey authentication mypubkey rsa GET-KEYS3
  rekey transport unicast
  sa ipsec 1
   profile GDOI-PROFILE
   match address ipv6 GETVPN6-0003
   replay time window-size 5
  address ipv4 192.168.18.2
  redundancy
   local priority 100
   peer address ipv4 192.168.19.2
!
ip access-list extended GETVPN-0001
 permit ip any any
ip access-list extended GETVPN-0002
 permit ip any any
ip access-list extended GETVPN-0003
 permit ip any any
!
ipv6 access-list GETVPN6-0001
 permit ipv6 any any
!
ipv6 access-list GETVPN6-0002
 permit ipv6 any any
!
ipv6 access-list GETVPN6-0003
 permit ipv6 any any
!
```


© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

Efficient Virtualisation and High-Scale VPNs –

3. Add encryption

- GETVPN Group Members
- Add crypto map to LISP0.x

ALL LISP SITES
identical! Cut/Paste!

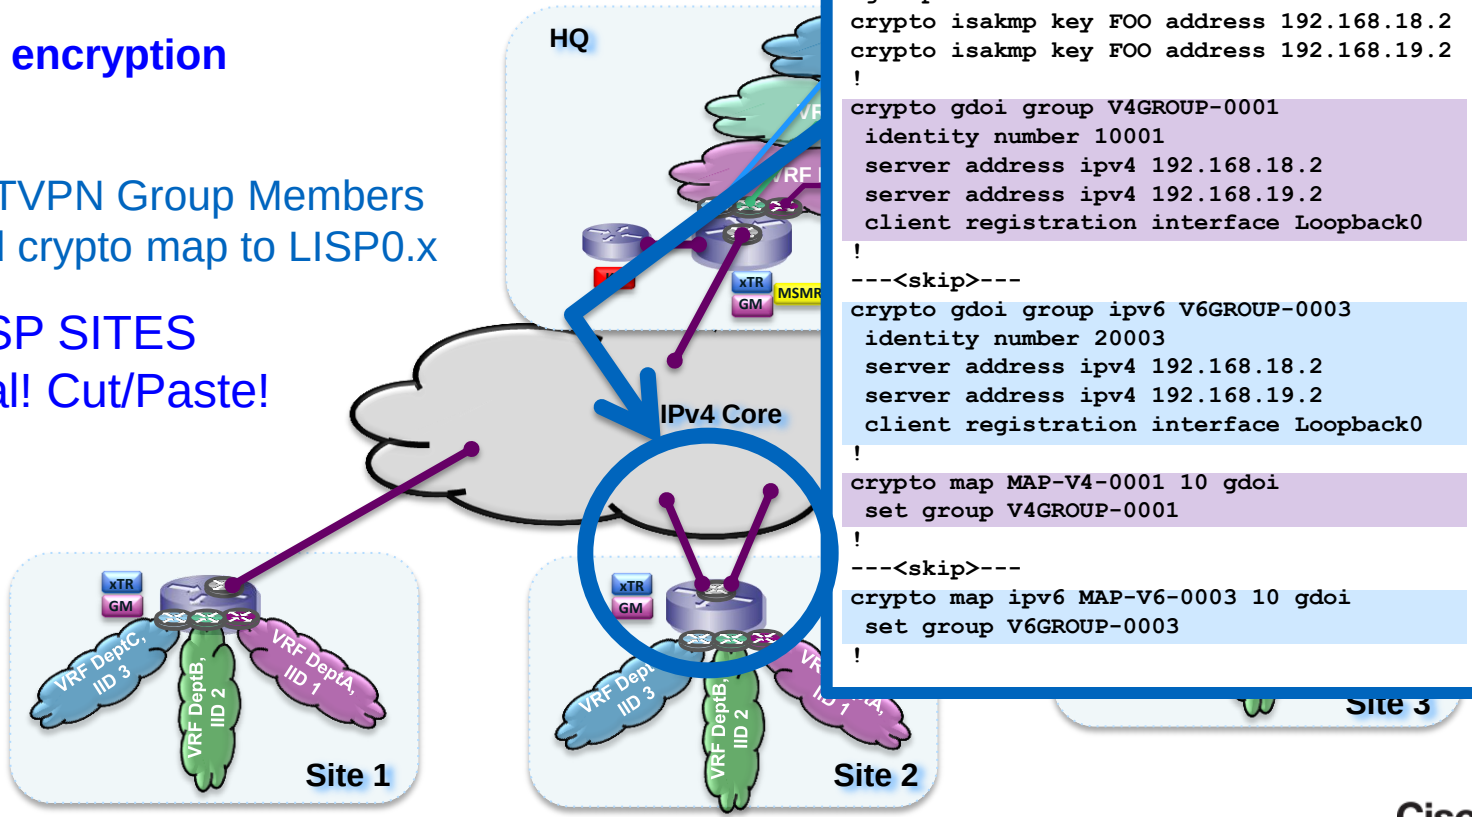

Remote2 xTR/GM

```
!
crypto isakmp policy 10
 encr aes 256
 authentication pre-share
 group 16
crypto isakmp key FOO address 192.168.18.2
crypto isakmp key FOO address 192.168.19.2
!
crypto gdoi group V4GROUP-0001
 identity number 10001
 server address ipv4 192.168.18.2
 server address ipv4 192.168.19.2
 client registration interface Loopback0
!
---<skip>---
crypto gdoi group ipv6 V6GROUP-0003
 identity number 20003
 server address ipv4 192.168.18.2
 server address ipv4 192.168.19.2
 client registration interface Loopback0
!
crypto map MAP-V4-0001 10 gdoi
 set group V4GROUP-0001
!
---<skip>---
crypto map ipv6 MAP-V6-0003 10 gdoi
 set group V6GROUP-0003
!
```


© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

Efficient Virtualisation and High-Scale VPNs –

3. Add encryption

- GETVPN Group Members
- Add crypto map to LISP0.x

ALL LISP SITES
identical! Cut/Paste!

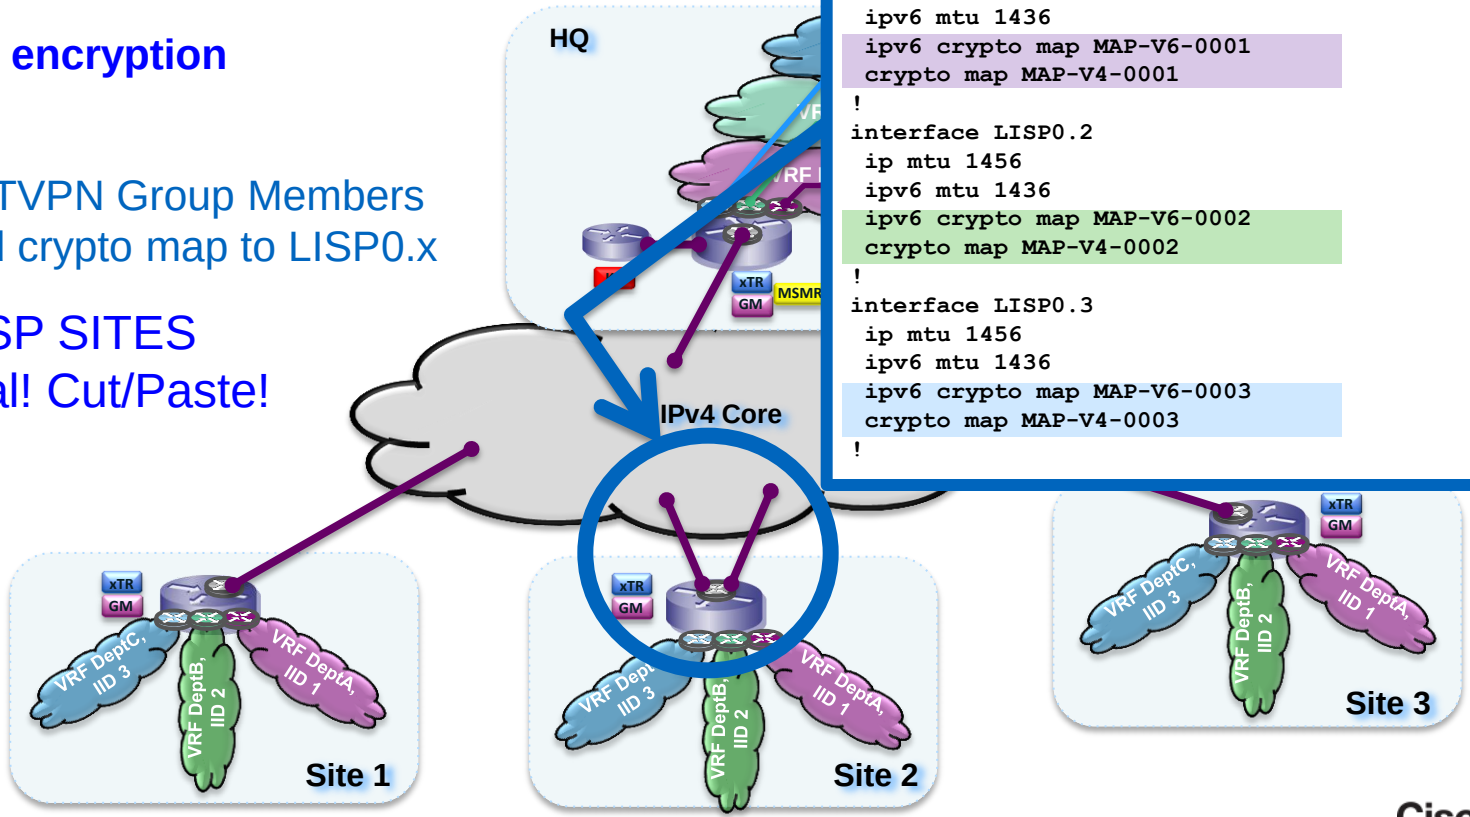

Remote2 xTR/GM

```
!
interface LISP0
!
interface LISP0.1
 ip mtu 1456
 ipv6 mtu 1436
 ipv6 crypto map MAP-V6-0001
 crypto map MAP-V4-0001
!
interface LISP0.2
 ip mtu 1456
 ipv6 mtu 1436
 ipv6 crypto map MAP-V6-0002
 crypto map MAP-V4-0002
!
interface LISP0.3
 ip mtu 1456
 ipv6 mtu 1436
 ipv6 crypto map MAP-V6-0003
 crypto map MAP-V4-0003
!
```


© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

Efficient Virtualisation and High-Scale VPNs – Overview

3. Add encryption

- GETVPN Group Members
- Add crypto map to LISP0.x

Verification…

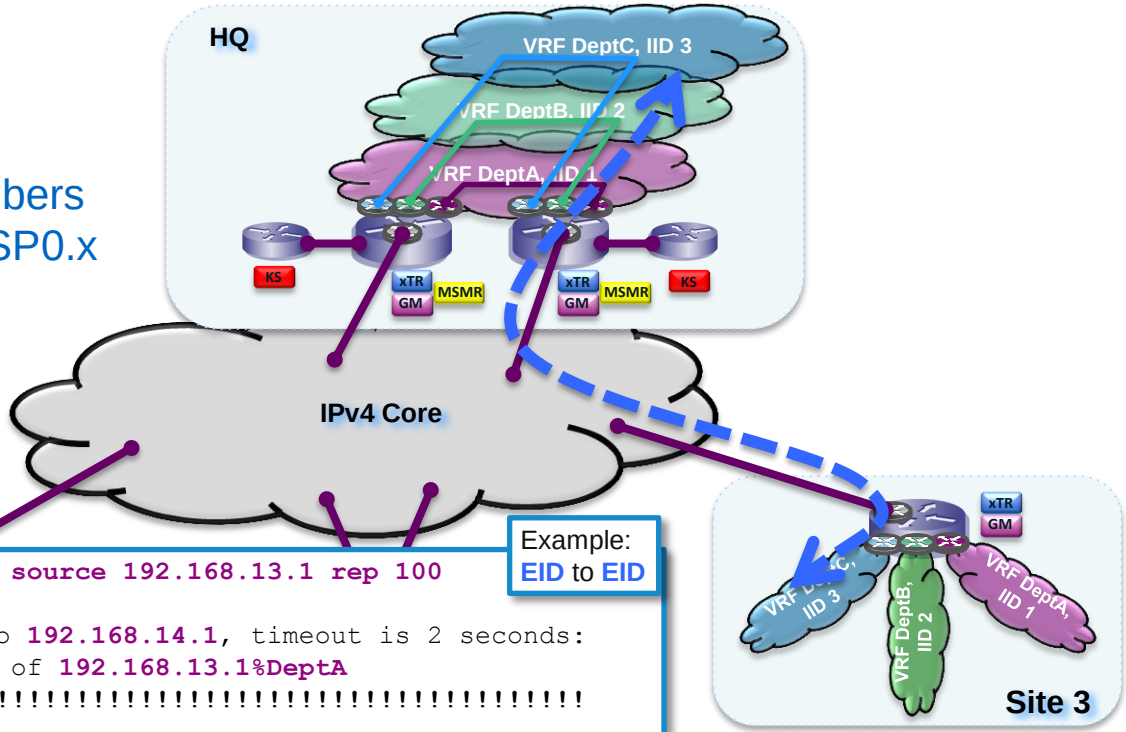

Example: EID to EID

```
Site3#ping vrf DeptA 192.168.14.1 source 192.168.13.1 rep 100
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 192.168.14.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.13.1%DeptA
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (100/100), round-trip min/avg/max = 5/6/12 ms
Site3#
```


© 2014 Cisco and/or its affiliates. All rights reserved.

LISP VPN/Virtualisation

Efficient Virtualisation and High-Scale VPNs – Overview

3. Add encryption

- GETVPN Group Members
- Add crypto map to LISP0.x

Verification…

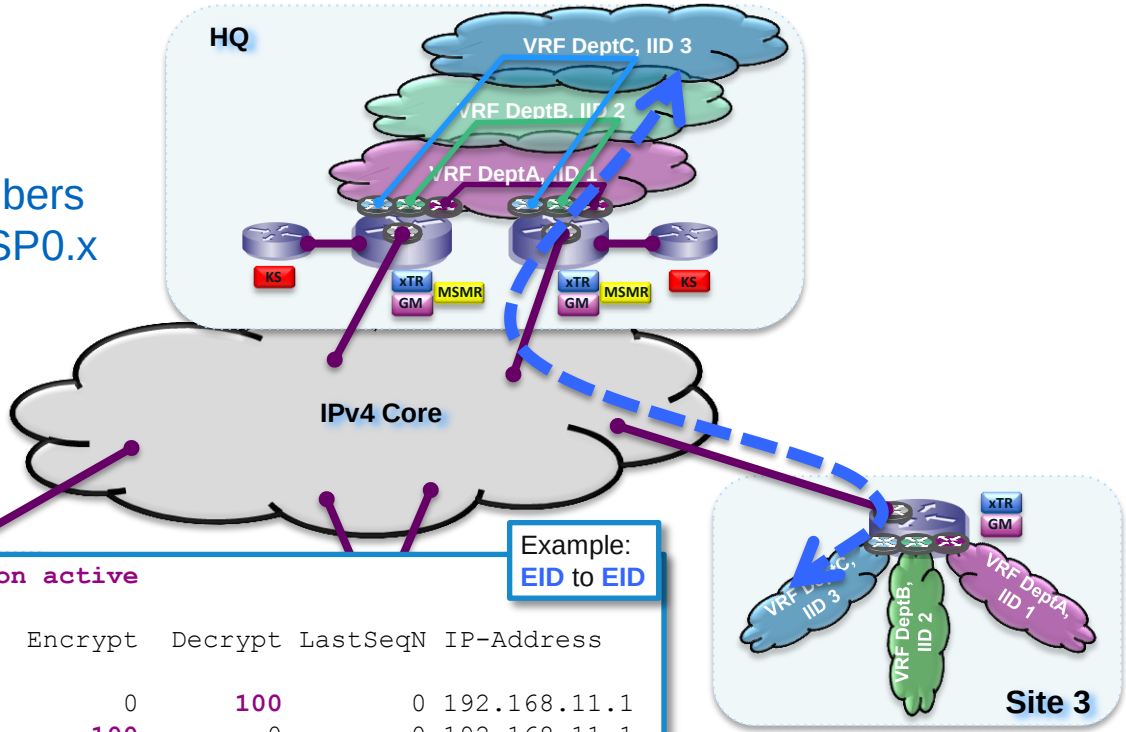

Example: EID to EID

```
Site3#show crypto engine connection active
Crypto Engine Connections
   ID  Type    Algorithm           Encrypt  Decrypt LastSeqN IP-Address
---<skip>---
  143  IPsec   AES256+SHA512             0      100        0 192.168.11.1
  144  IPsec   AES256+SHA512           100        0        0 192.168.11.1
---<skip>---
Site3#
```


© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Deployment Examples

LISP Deployment Examples…

1. Efficient Multihoming and Multi-AF (IPv4 and IPv6)
2. Efficient Virtualisation and High-Scale VPNs
3. Data Centre/Host Mobility
4. LISP-Mobile Node

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Data Centre/Host Mobility

Data Centre/Host Mobility – Overview

- Needs:
  - VM-Mobility extending subnets and across subnets
  - Move detection, dynamic EID-to-RLOC mappings, traffic redirection
- LISP Solution:
  - OTV + LISP for VM-moves in extended subnets
  - LISP for VM-moves across subnets
- Benefits:
  - VM OS agnostic, seamless, integrated, global workload mobility
  - Direct Path (no triangulation)
  - Connections survive across moves
  - No routing re-convergence, no DNS updates
  - Global Scalability (cloud bursting)
  - ARP elimination

Data Centre/Host Mobility

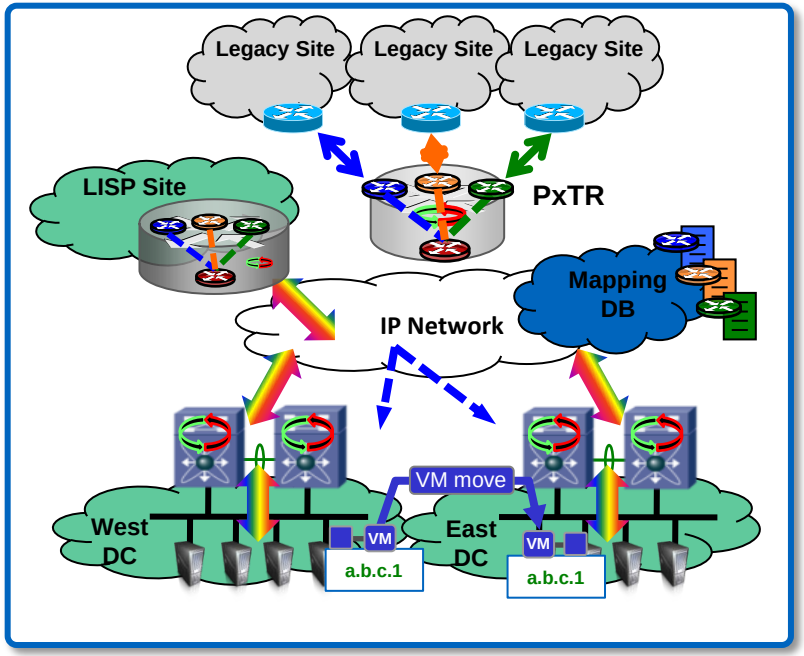

© 2014 Cisco and/or its affiliates. All rights reserved.

Moving vs. Distributing Workloads

Why do we really need LAN Extensions?

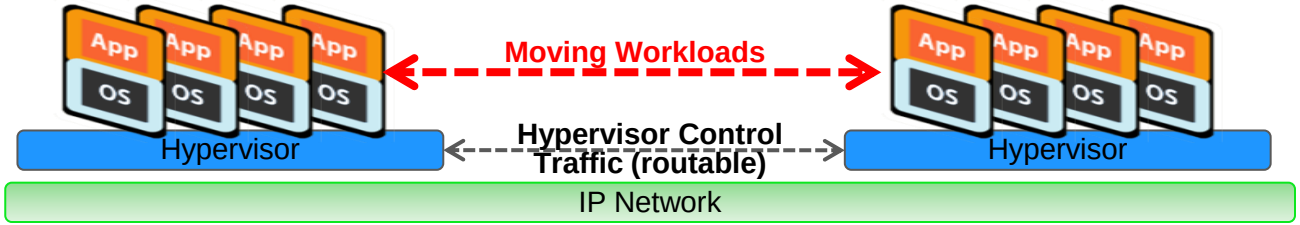

- **Move workloads** with IP mobility solutions: LISP Host Mobility
  - IP preservation is the real requirement (LAN extensions not mandatory)

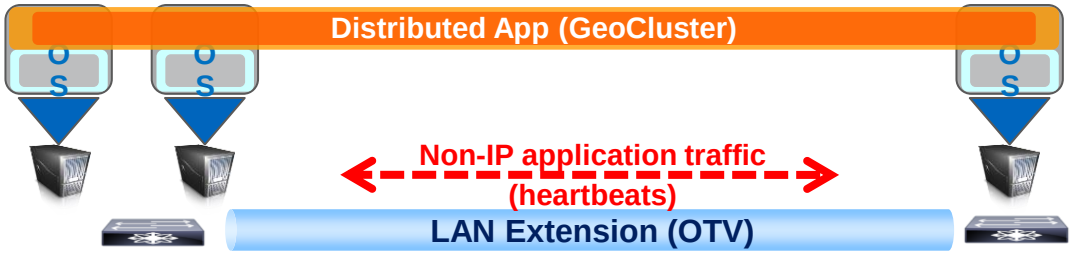

- **Distribute workloads** with LAN extensions
  - Application High Availability with Distributed Clusters

© 2014 Cisco and/or its affiliates. All rights reserved.

Host-Mobility Scenarios

Two Mobility Scenarios

LISP Host Mobility Config Guide:
http://lisp.cisco.com

Moves Without LAN Extension

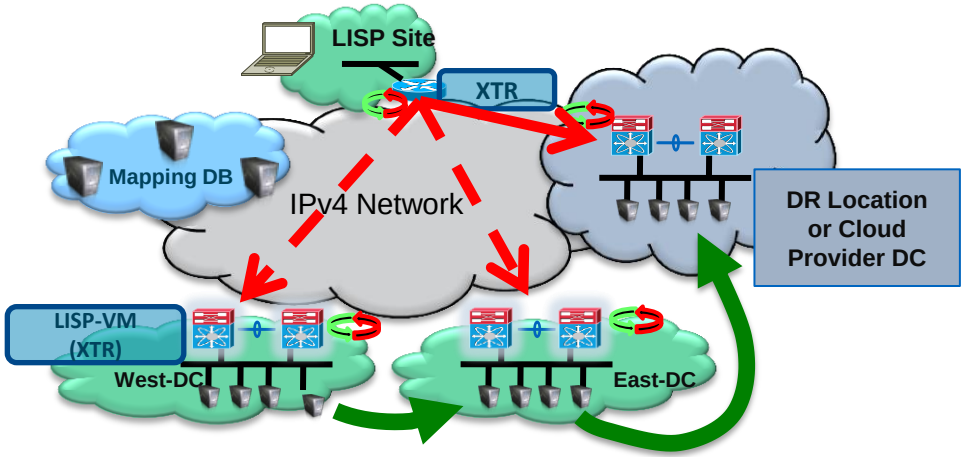

- IP Mobility Across Subnets
  - Disaster Recovery
  - Cloud Bursting
- Application Members In One Location

Moves With LAN Extension

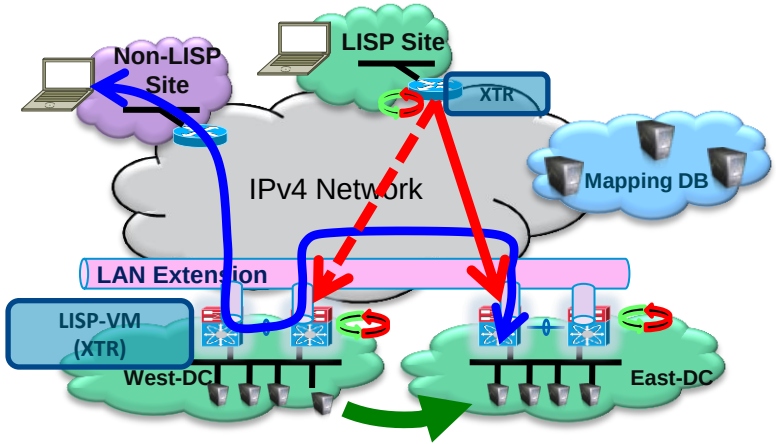

- Routing for Extended Subnets
  - Active-Active Data Centres
  - Distributed Data Centres
- Application Members Distributed
  - Broadcasts across sites

© 2014 Cisco and/or its affiliates. All rights reserved.

A LISP Packet Walk

Before Moving the Host

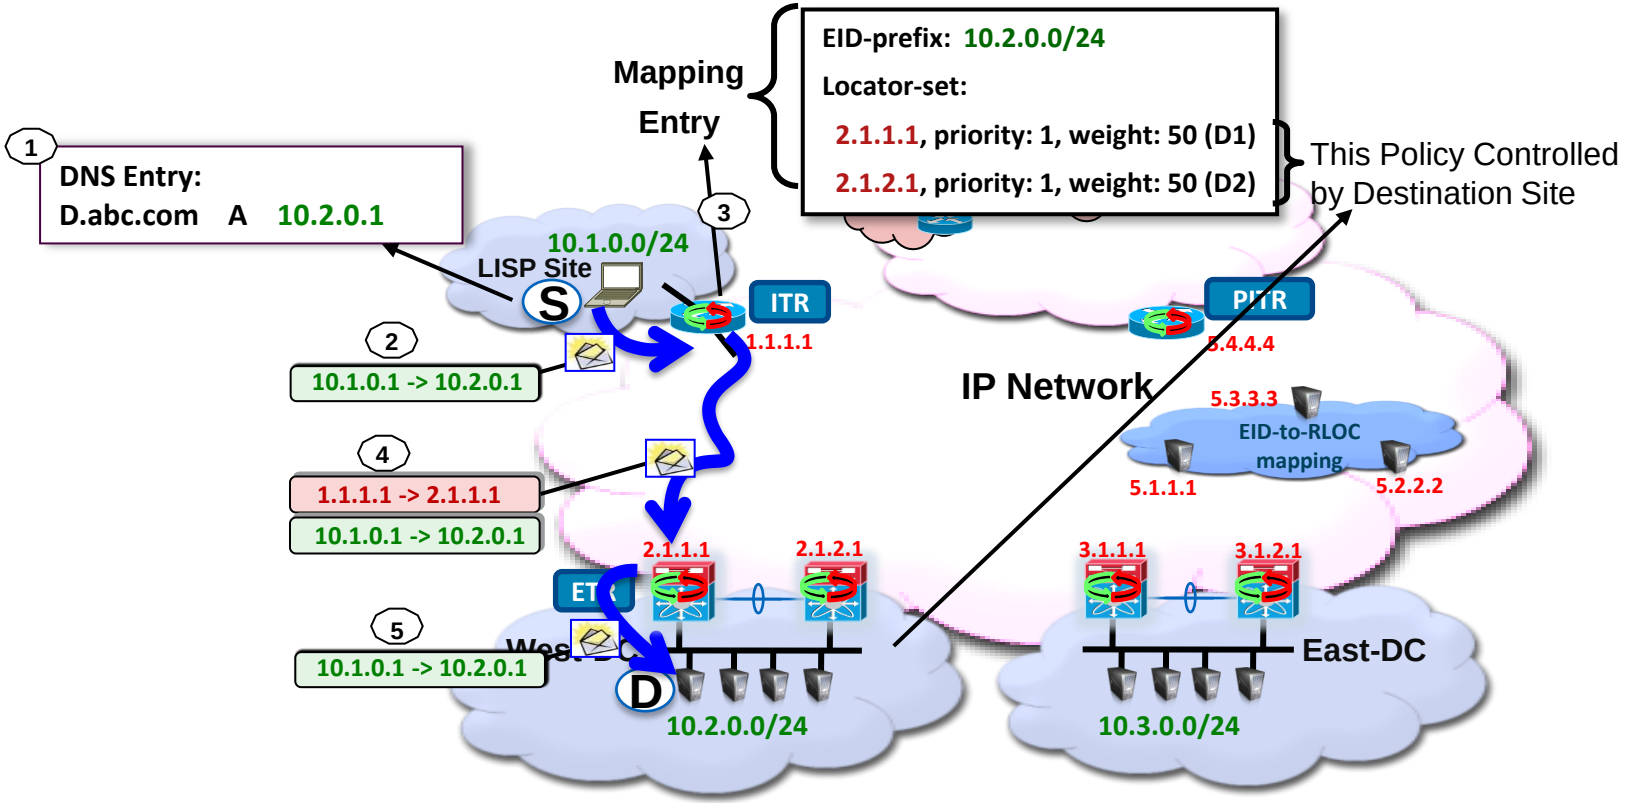

© 2014 Cisco and/or its affiliates. All rights reserved.

A LISP Packet Walk

After Moving the Host

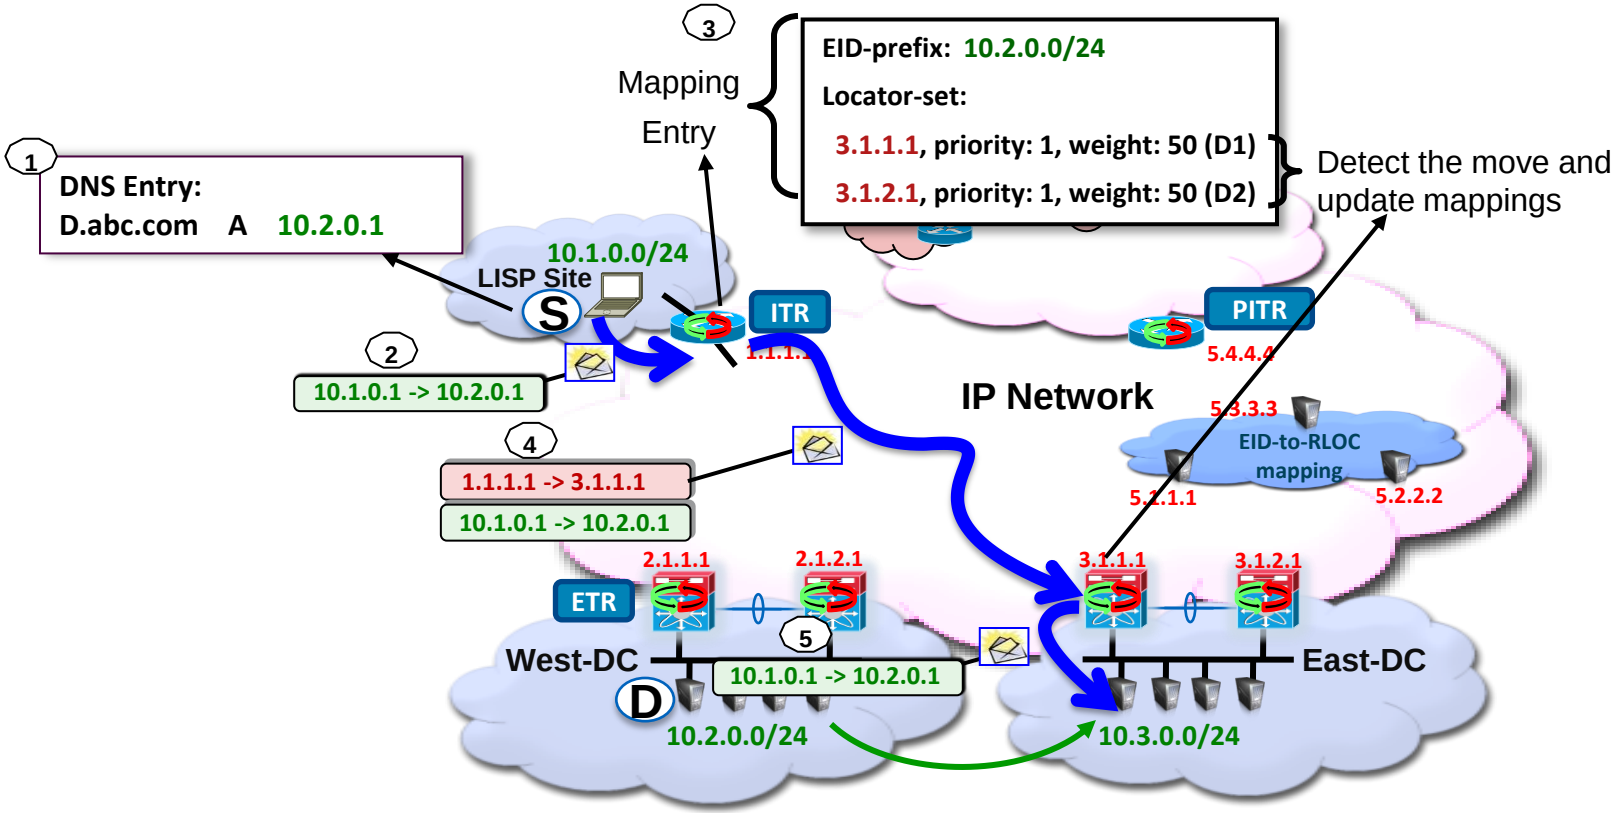

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Host-Mobility – First Hop Routing

No LAN Extension

- SVI (Interface VLAN x) and HSRP configured as usual
  - Consistent GWY-MAC configured across all dynamic subnets
- The lisp mobility <dyn-eid-map> command enables proxy-arp functionality on the SVI
  - The LISP-VM router services first hop routing requests for both local and roaming subnets
- Moving hosts always talk to a local gateway with the same MAC

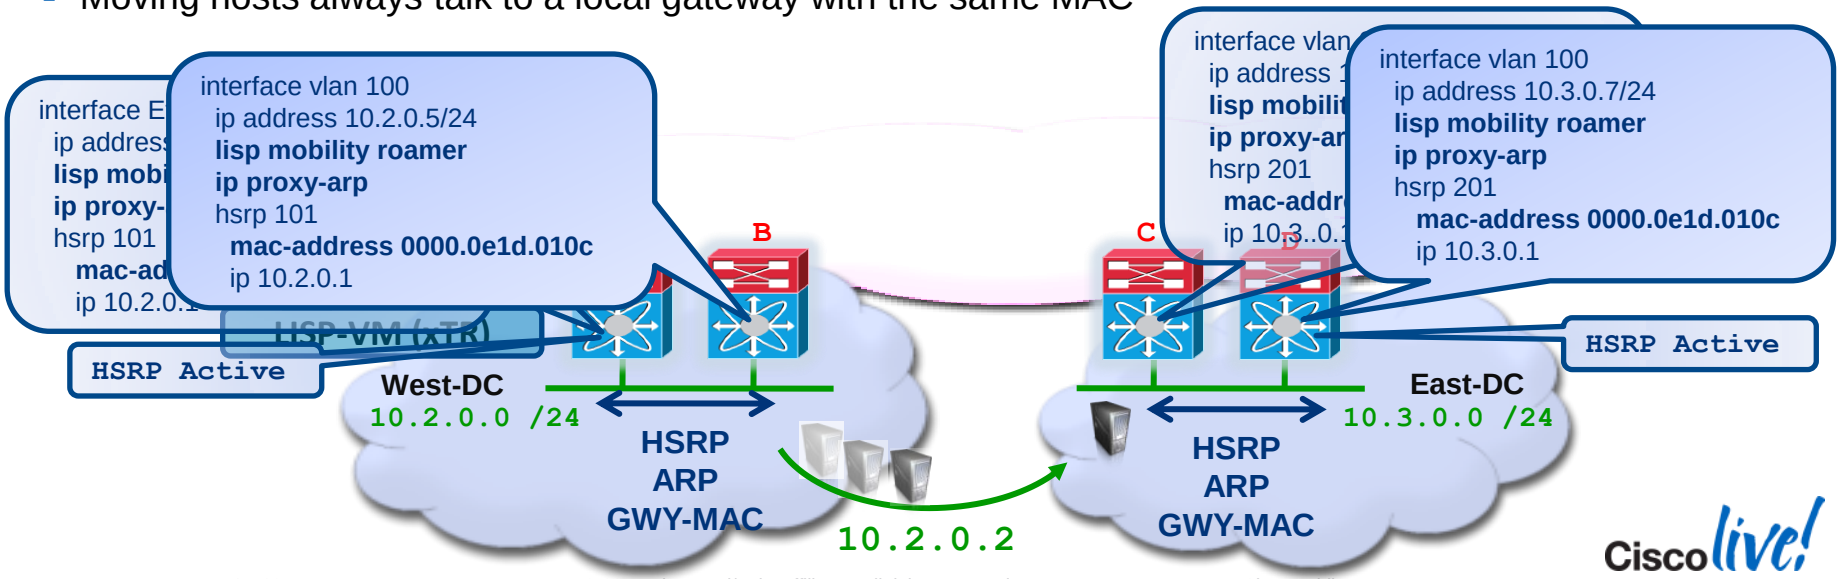

© 2014 Cisco and/or its affiliates. All rights reserved.

Host-Mobility and Multi-homing

ETR Updates – Across LISP Sites

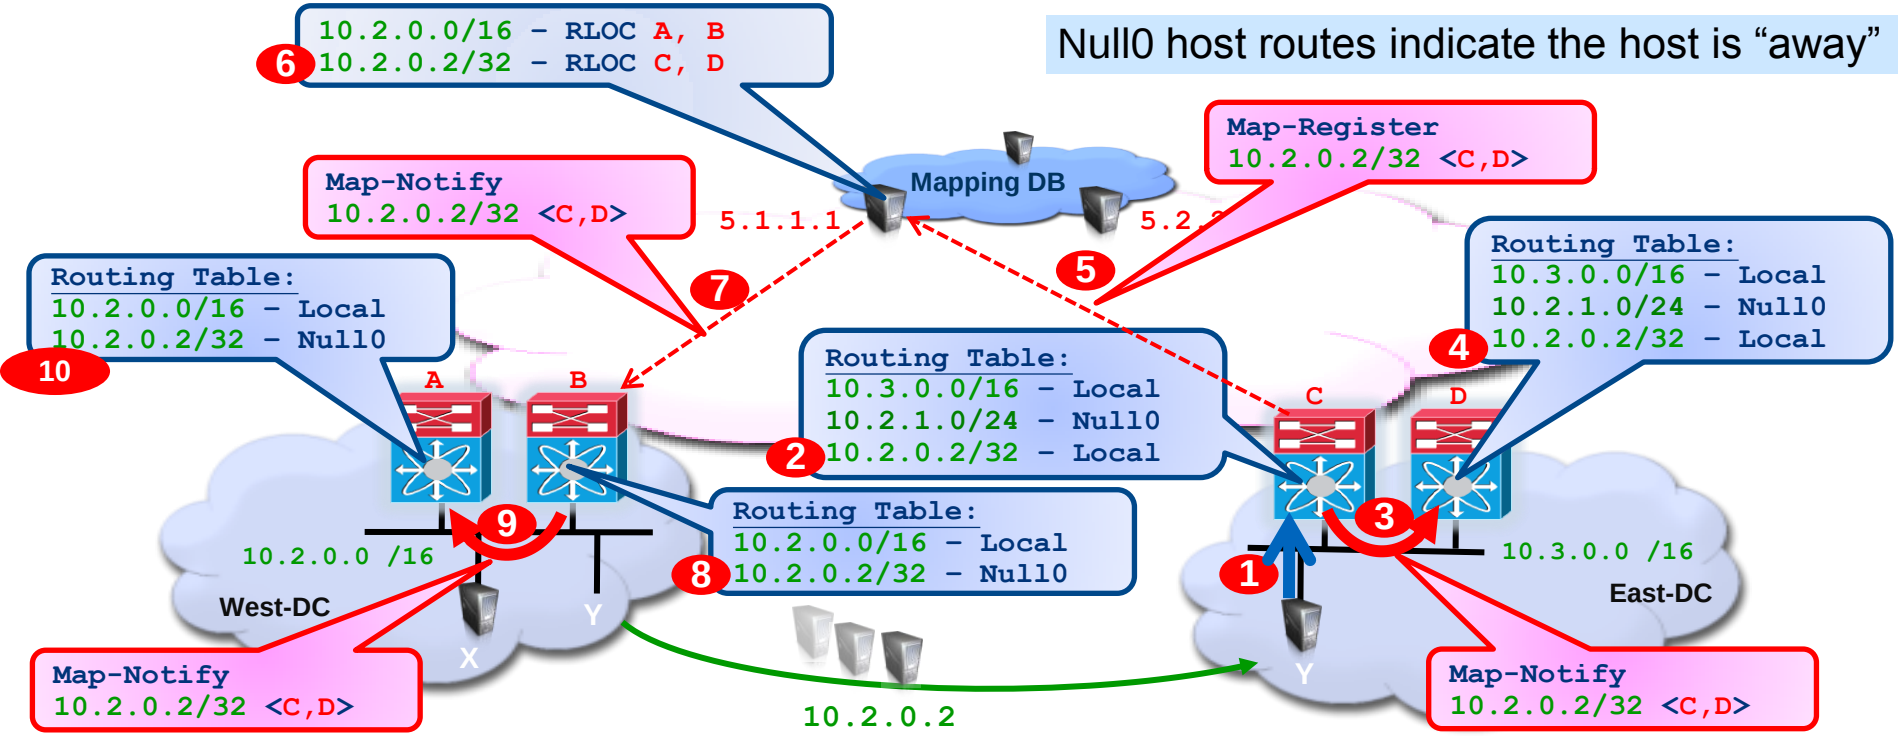

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public

LISP Mobility Across LISP Sites

Client-server communication established without the need to discover the workloads in the “home subnet” in West-DC

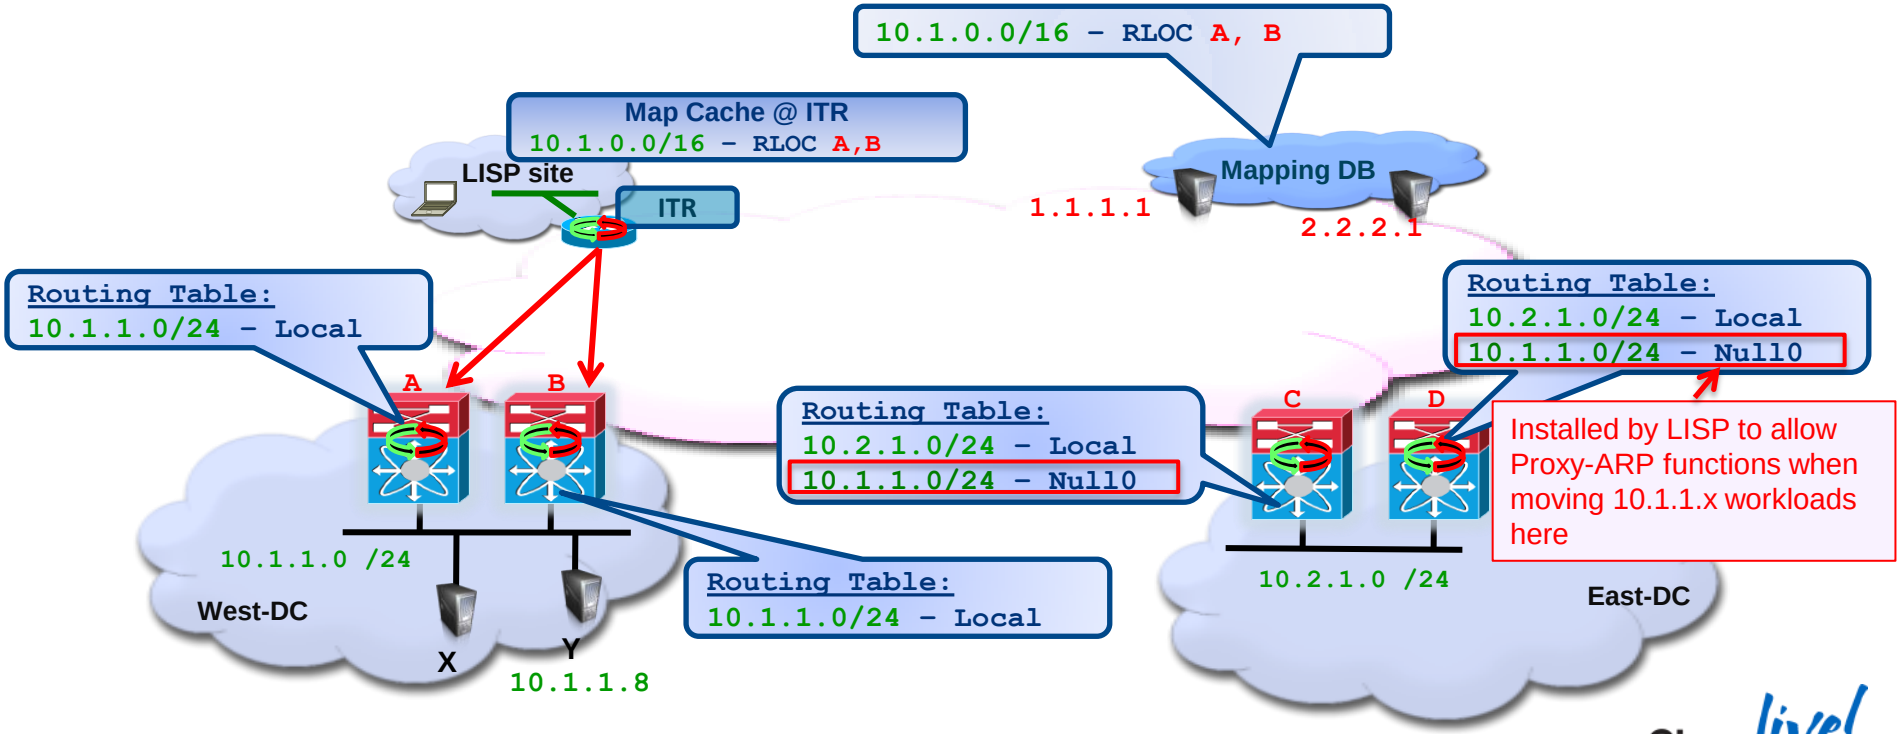

© 2014 Cisco and/or its affiliates. All rights reserved.

Refreshing the Map Caches

1. ITRs and PITRs with cached mappings continue to send traffic to the old locators
   1. The old xTR knows the host has moved (Null0 route)
2. Old xTR sends Solicit Map Request (SMR) messages to any encapsulators sending traffic to the moved host
3. The ITR then initiates a new map request process
4. An updated map-reply is issued from the new location
5. The ITR Map Cache is updated

- Traffic is now re-directed
- SMRs are an important integrity measure to avoid unsolicited map responses and spoofing

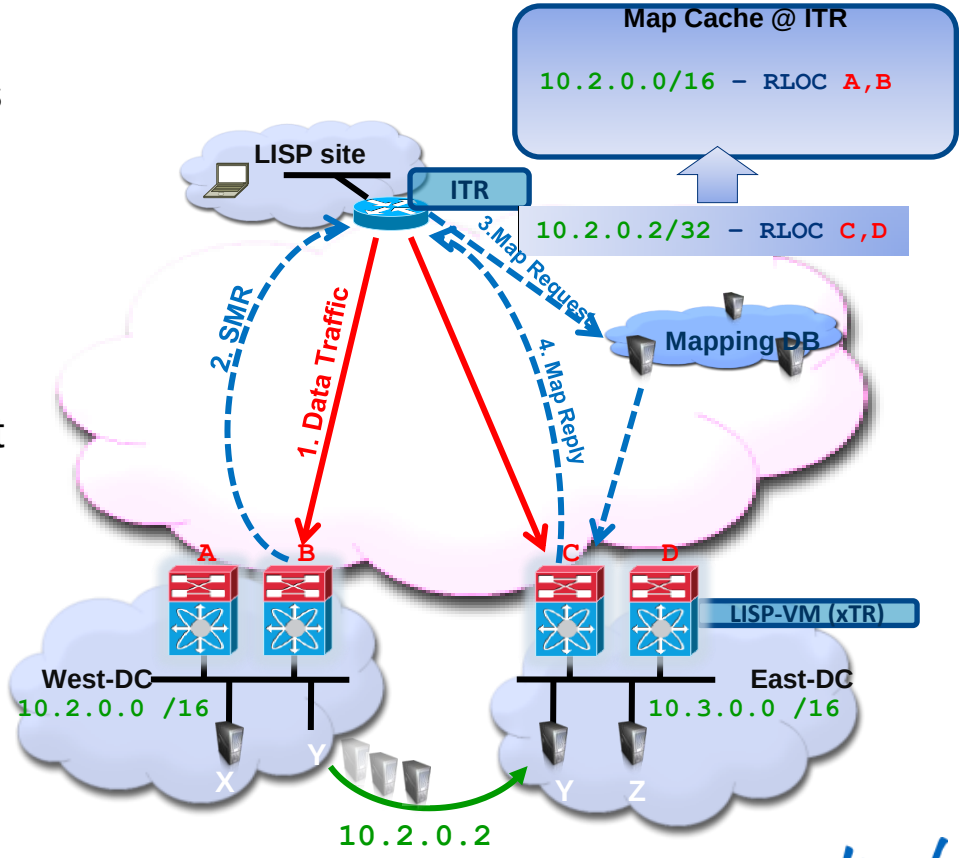

© 2014 Cisco and/or its affiliates. All rights reserved.

On-subnet Server-Server Traffic

West-to-East

- X ARPs for Y, /32 Null0 entry for Y triggers proxy-ARP on West-DC xTRs to ensure traffic is steered there
  - –Note: entry for Y in X ARP cache is cleared by GARP message originated by West-DC XTRs
- Traffic to Y is **LISP encapsulated**

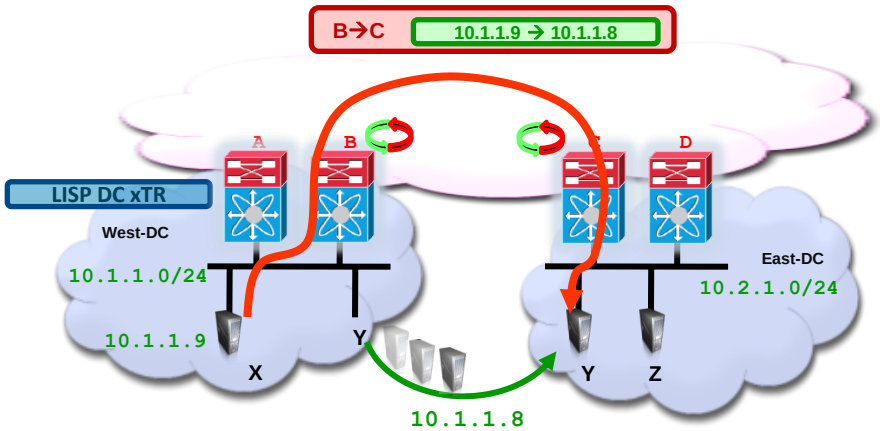

East-to-West

- Y ARPs for X, /24 Null0 entry for the 'home subnet' triggers proxy-ARP on East DC xTRs to ensure traffic is steered there
  - Note: assumption is that ARP cache on Y is refreshed after the move
- Traffic to X is **LISP encapsulated**

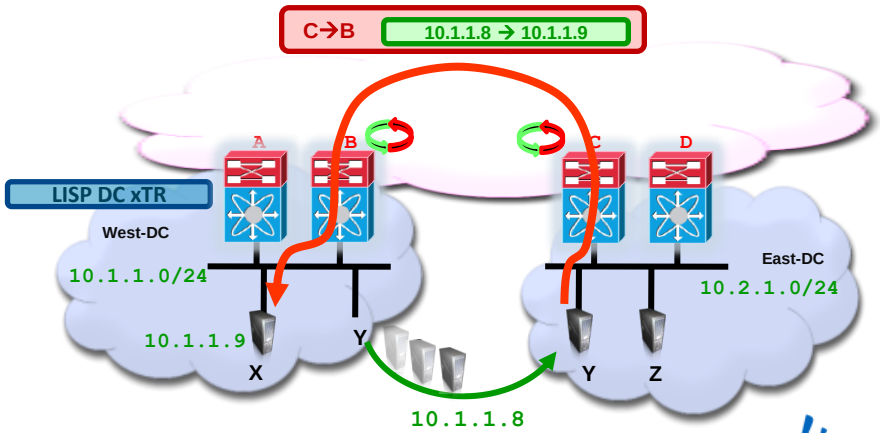

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Host-Mobility – First Hop Routing

With Extended Subnets

- Consistent GWY-IP and GWY-MAC configured across all sites
  - Consistent HSRP group number across sites ➔ consistent GWY-MAC
- Servers can move anywhere and always talk to a local gateway with the same IP/MAC

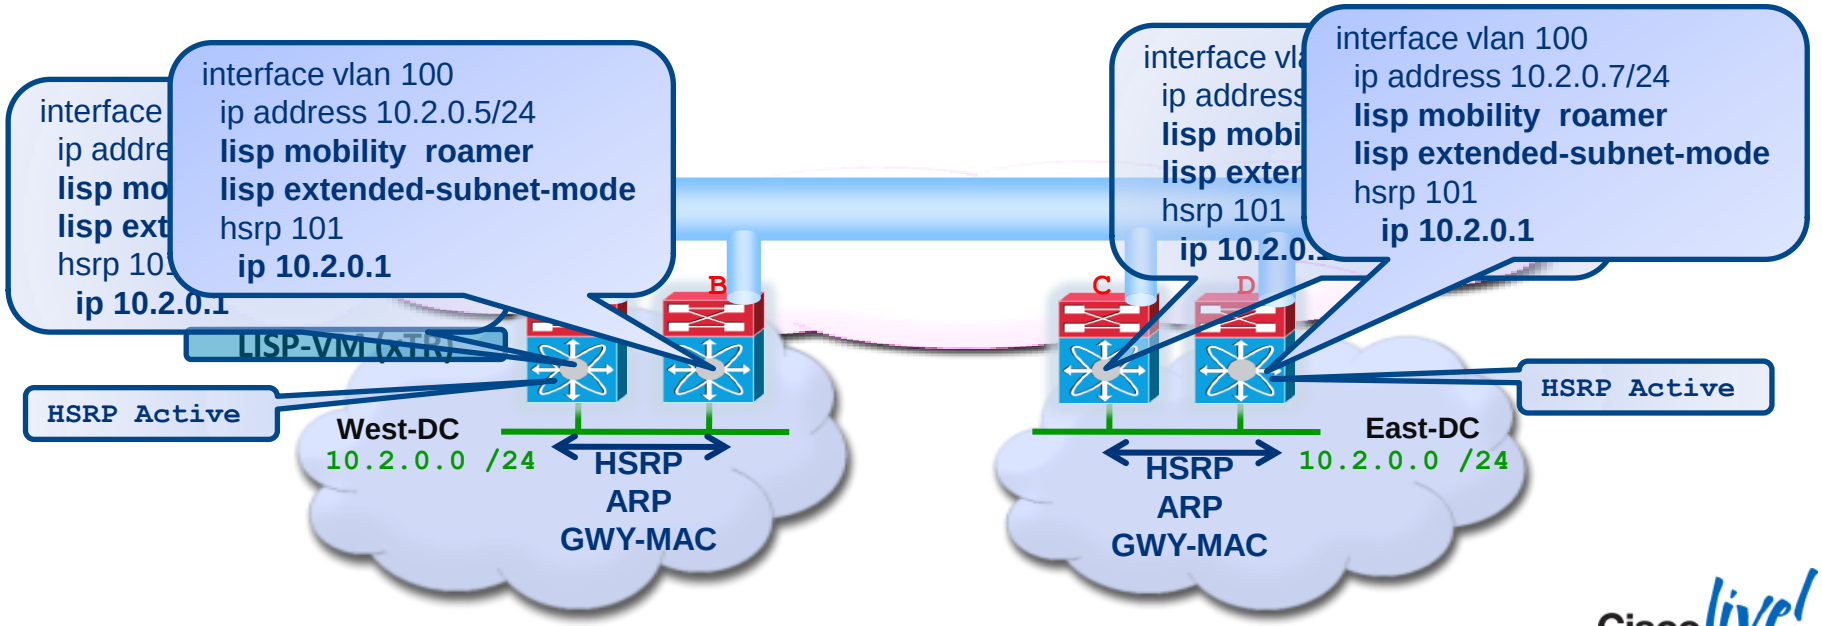

© 2014 Cisco and/or its affiliates. All rights reserved.

Host-Mobility and Multi-homing

ETR Updates – Extended Subnets

Null0 host routes indicate the host is “away”

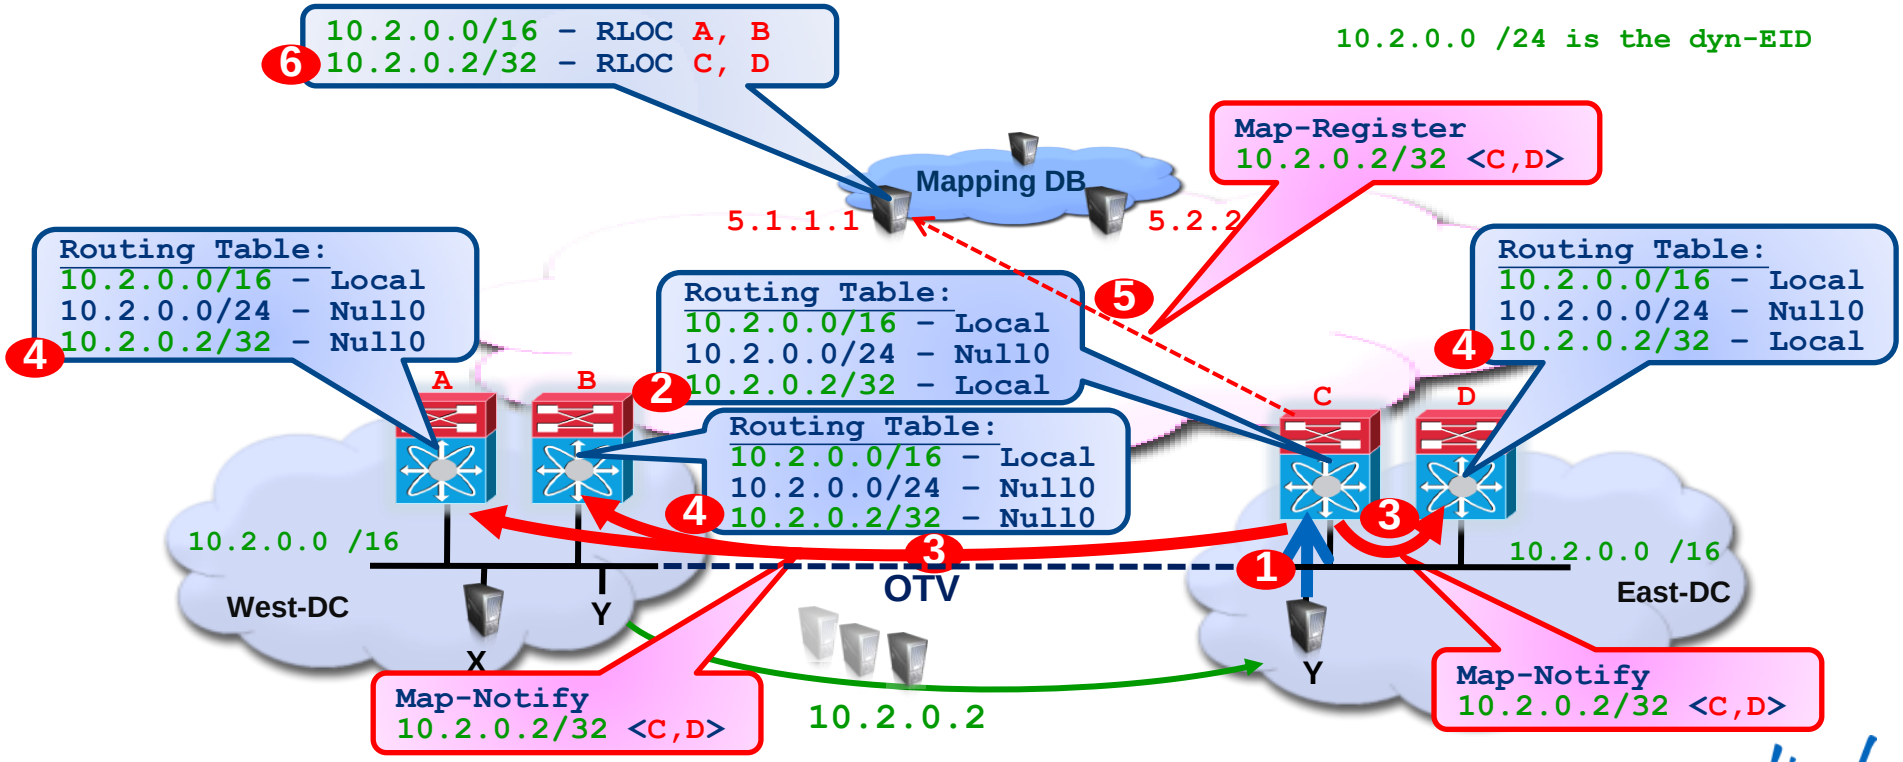

© 2014 Cisco and/or its affiliates. All rights reserved.

Refreshing the Map Caches

1. ITRs and PITRs with cached mappings continue to send traffic to the old locators
   1. The old xTR knows the host has moved (Null0 route)
2. Old xTR sends Solicit Map Request (SMR) messages to any encapsulators sending traffic to the moved host
3. The ITR then initiates a new map request process
4. An updated map-reply is issued from the new location
5. The ITR Map Cache is updated

- Traffic is now re-directed
- SMRs are an important integrity measure to avoid unsolicited map responses and spoofing

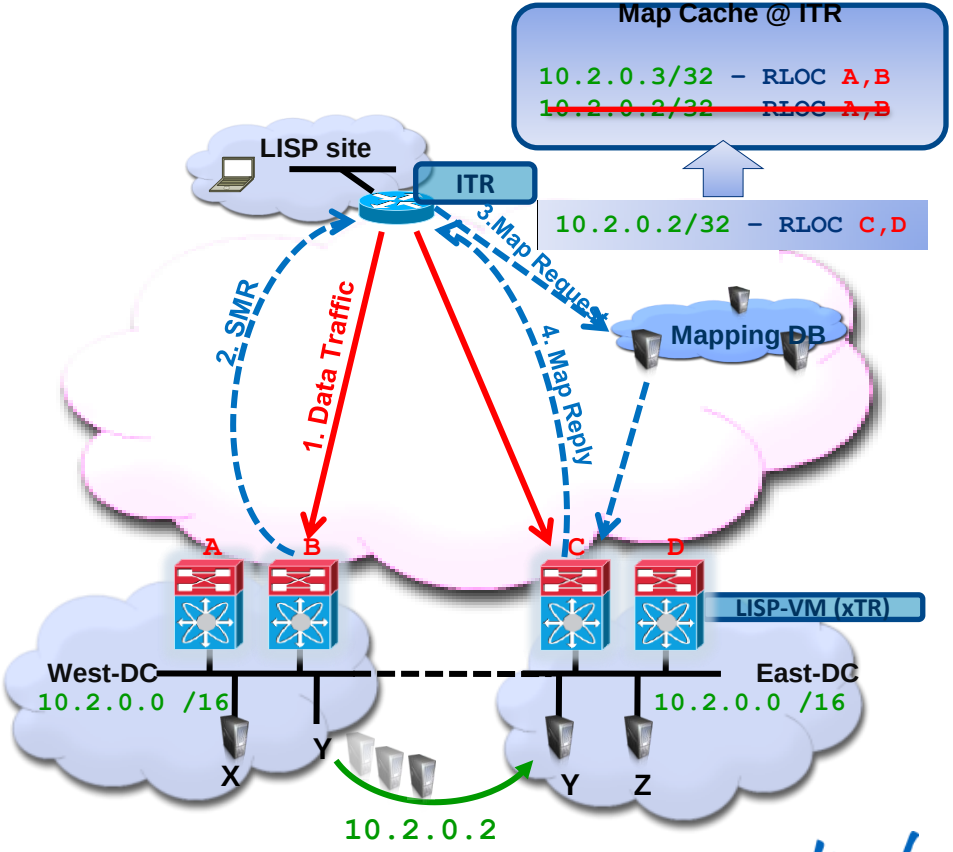

© 2014 Cisco and/or its affiliates. All rights reserved.

Off-subnet Client-Server Traffic

All Off-Subnet/Off-Site Traffic Is LISP Encapsulated

- Clients (192.168.0.1 & 192.168.2.1 communicate with Server 10.2.0.2
- Client-server traffic is LISP encapsulated at the ITRs or PITRs
  - Client-to-server:
    - to ETRs C or D
  - Server-to-client:
    - to ETR (F) for LISP sites
    - to PETR (G) for non-LISP sites
- Server-Server off-subnet traffic across sites is also LISP encapsulated

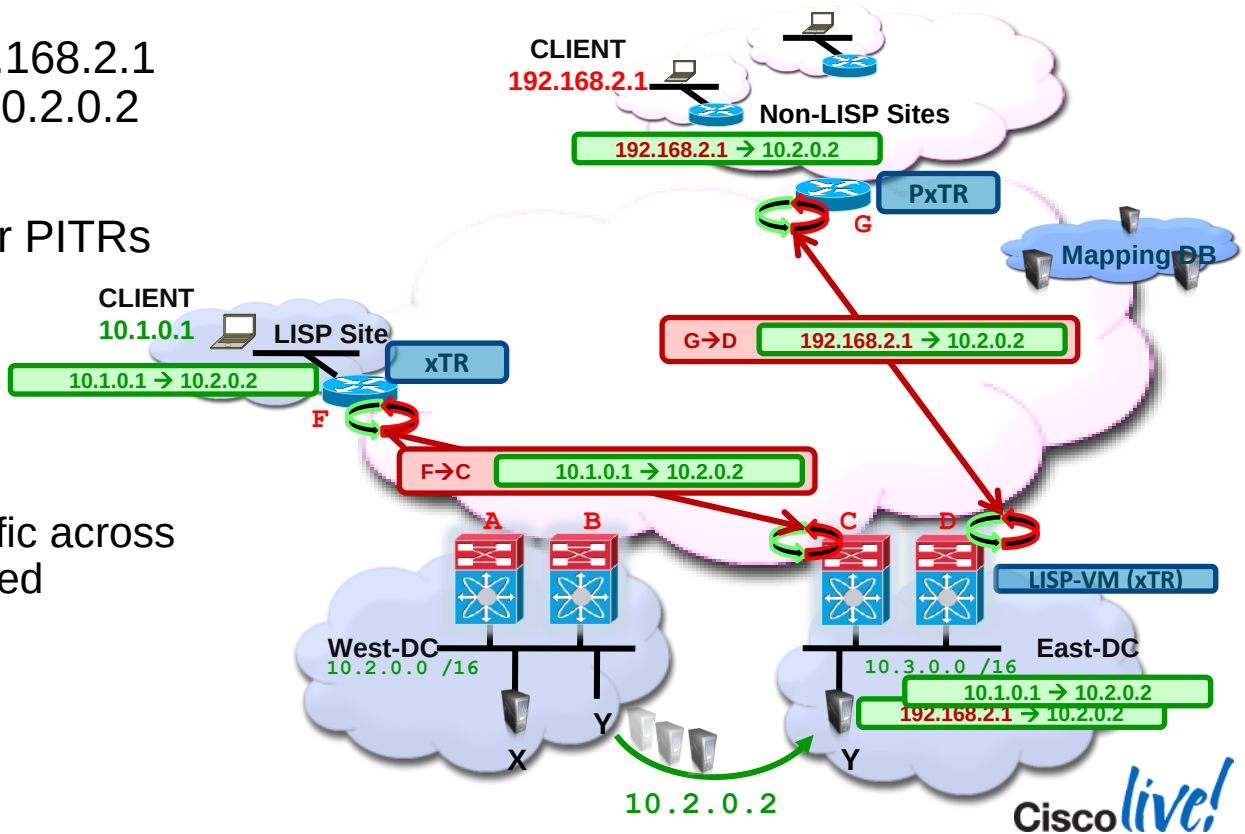

© 2014 Cisco and/or its affiliates. All rights reserved.

On-subnet Server-Server Traffic

On Subnet Traffic Across L3 Boundaries

With LAN Extension

- Live moves and cluster member dispersion
- Traffic between X & Y uses the **LAN Extension**
- Link-local-multicast handled by the LAN Extension

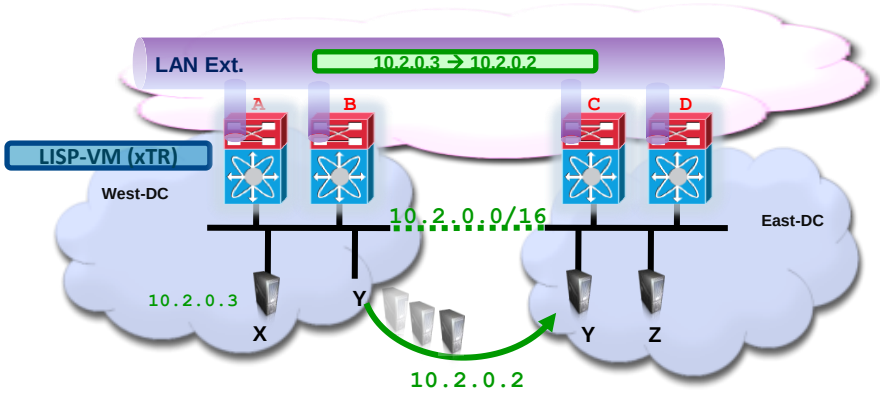

Without LAN Extensions

- Cold moves, no application dispersion
- X- Y traffic is sent to the LISP-VM router & **LISP encapsulated**
- Need LAN extensions for link-local multicast traffic

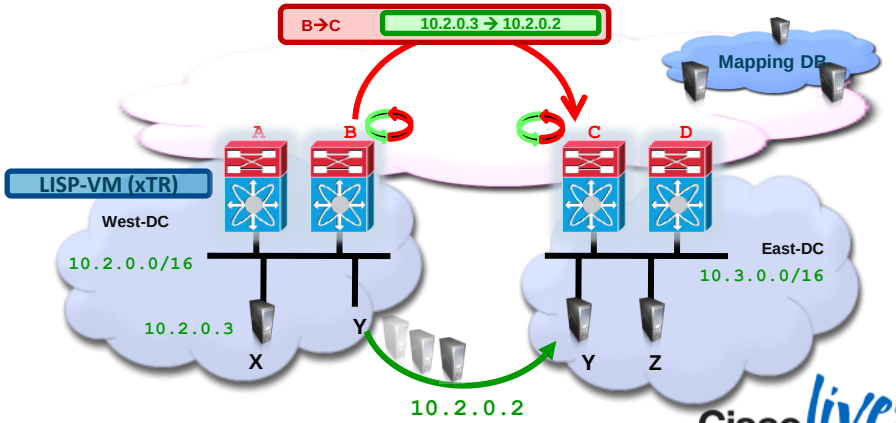

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP – New Features

LISP Multi-Hop Mobility – Extended Subnet Mode (ESM)

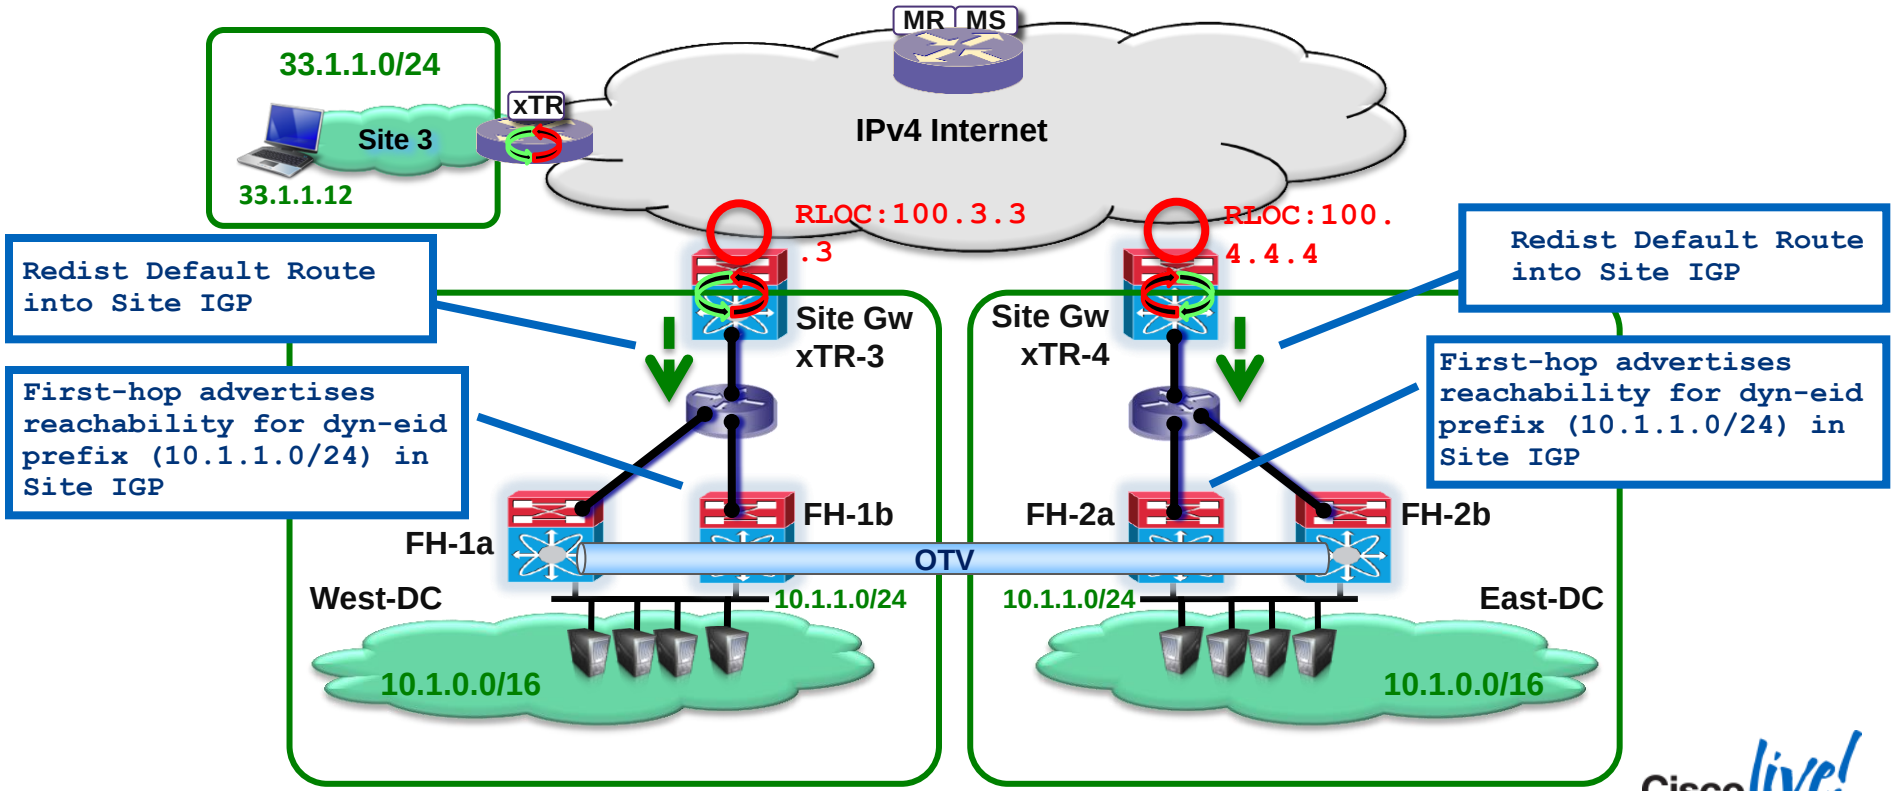

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP – New Features

LISP Multi-Hop Mobility – Extended Subnet Mode (ESM)

1. Dyn-EID detection
2. Multicast-map-notify
3. Eid-notify to Site Gw xTR
4. Map-register message

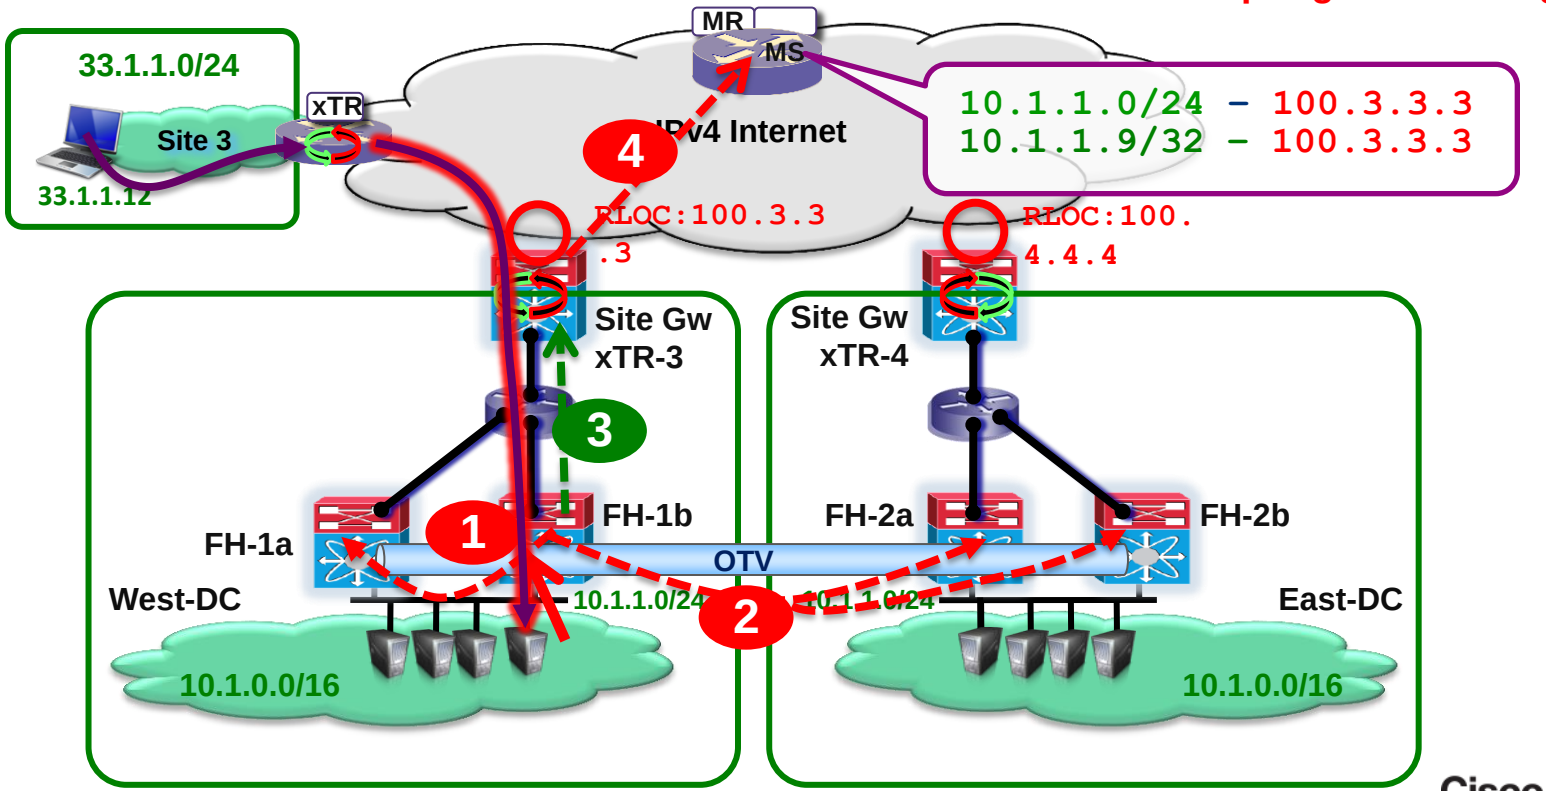

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP – New Features

LISP Multi-Hop Mobility – Extended Subnet Mode (ESM)

1. Dyn-EID detection
2. Multicast-map-notify
3. Eid-notify to Site Gw xTR
4. Map-register message
5. Map-Notify message

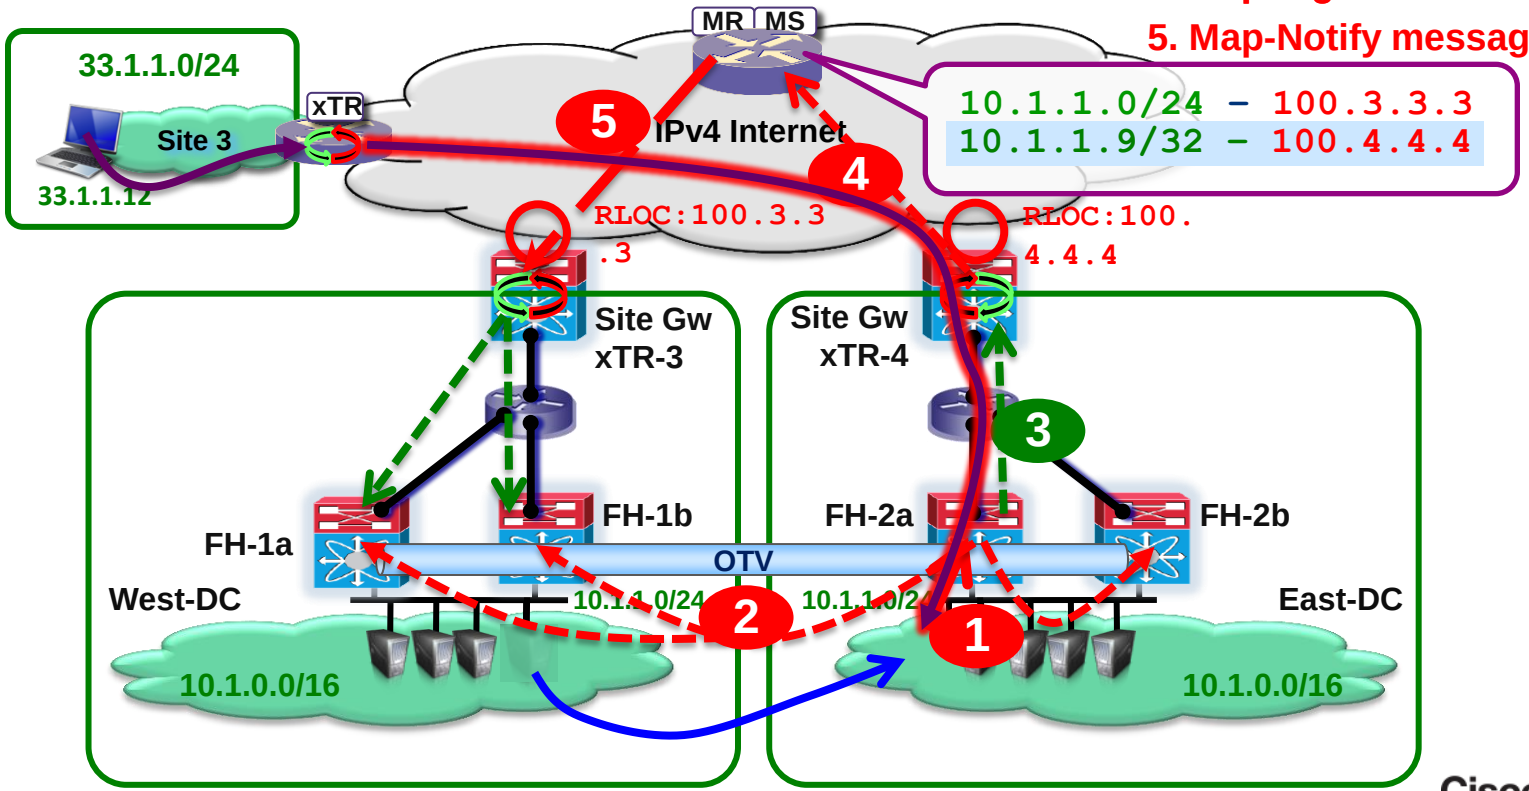

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Mobile Node

Session Continuity While Roaming!

LISP-MN Mobility: Any Network, Anytime, Anywhere…

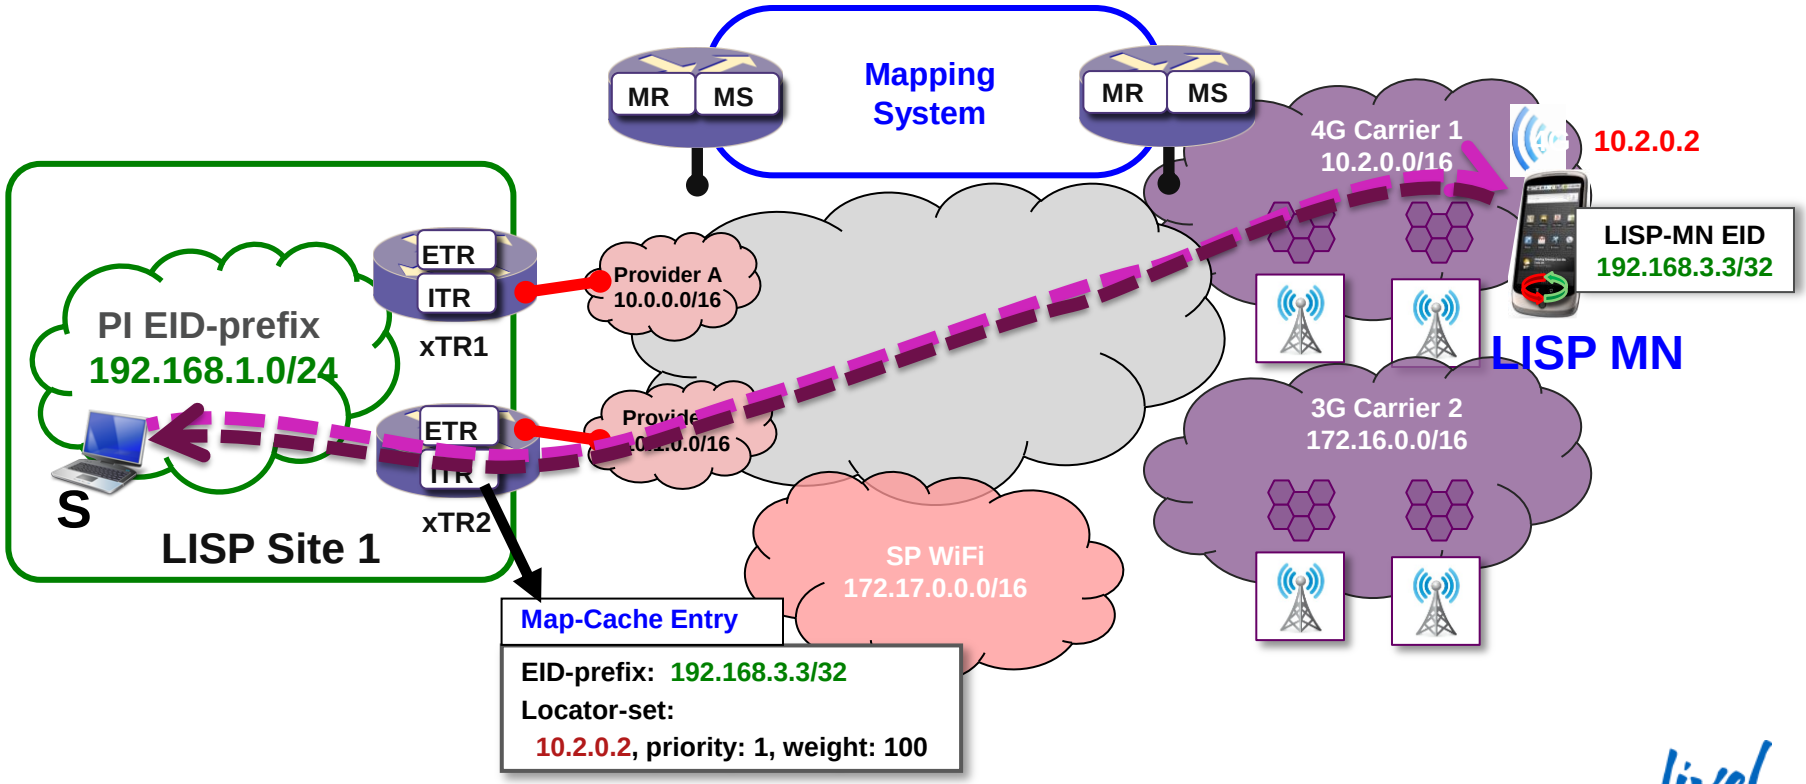

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Mobile Node

Session Continuity While Roaming!

LISP-MN Mobility: Any Network, Anytime, Anywhere…

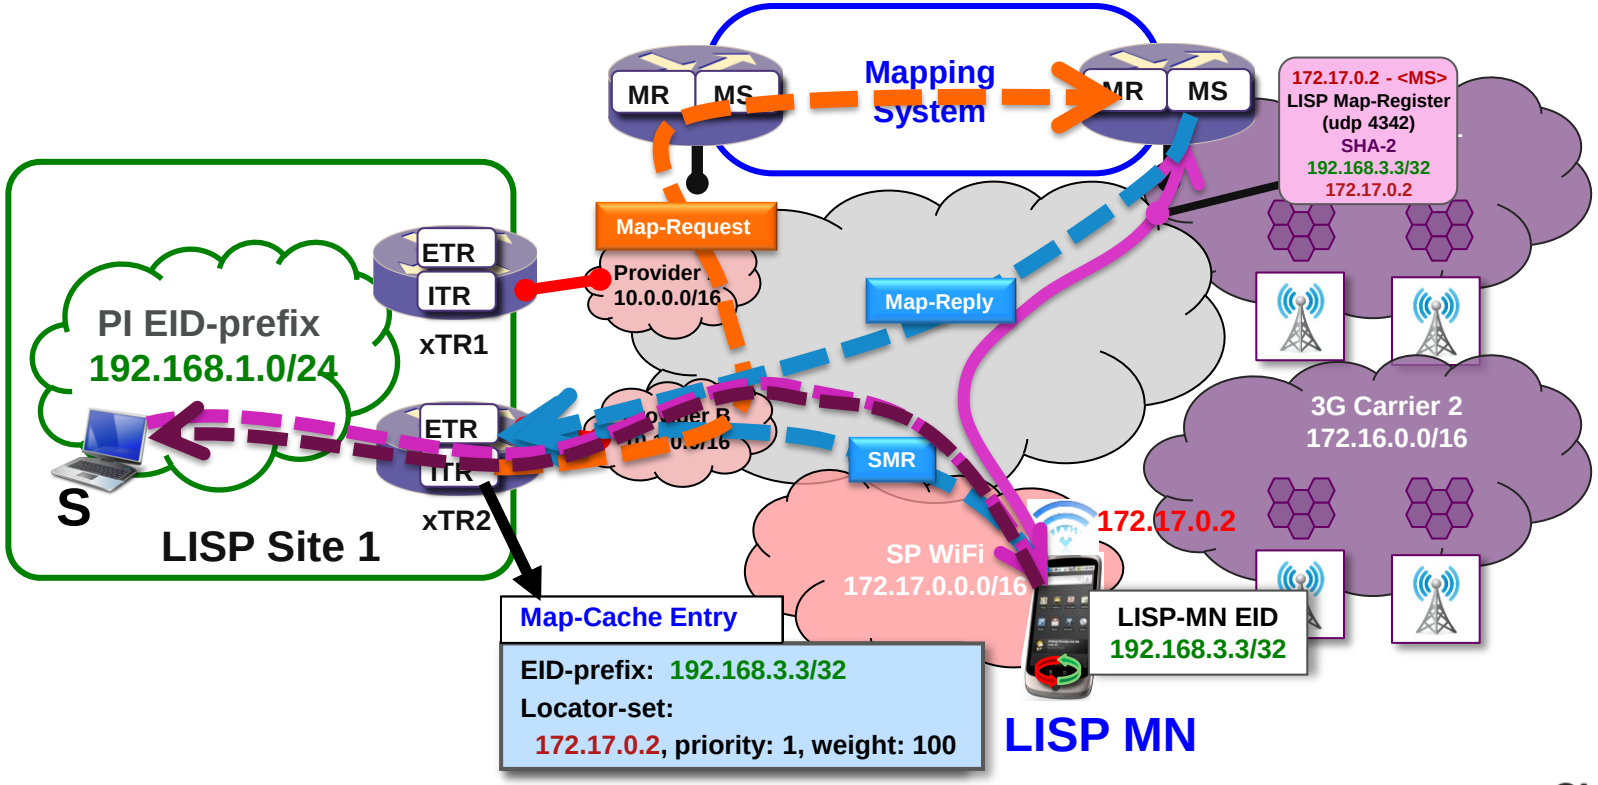

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Use Cases :: LISP Mobility

35 Buses Operational Throughout the Event

Customer Example :: Cisco Live US 2013 Transportation System

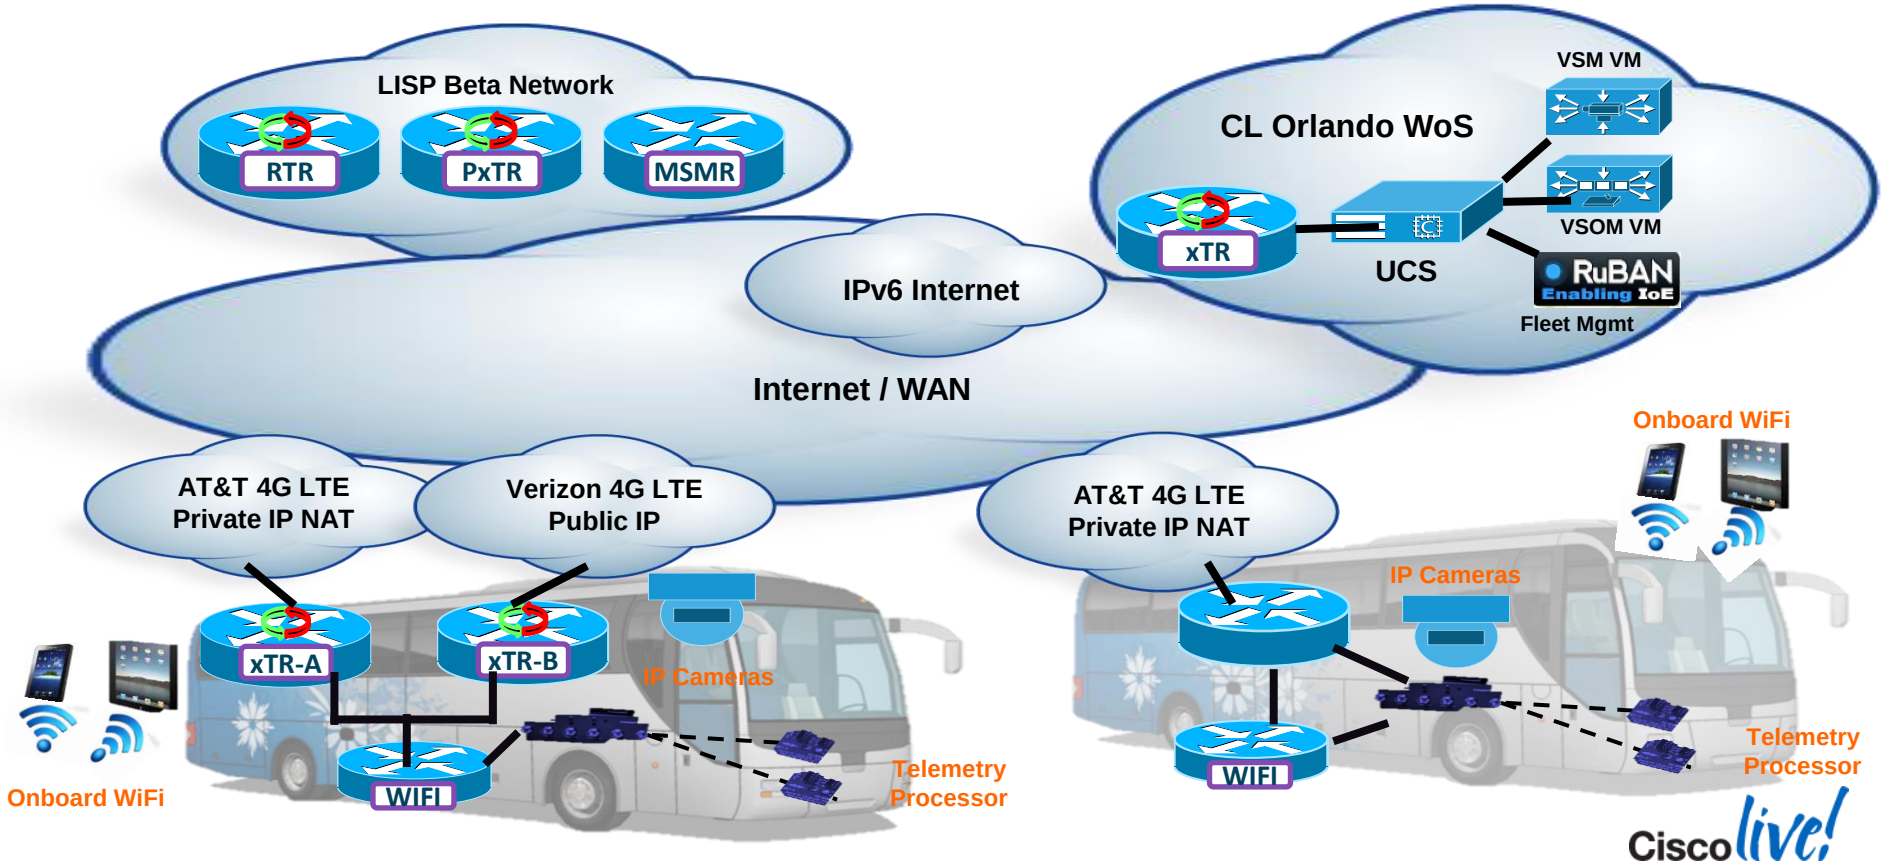

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Mobile Node Embedded Hardware

Open Source LISP Software

Linksys WRT160NL

Architecture	**MIPS** Atheros AP81
CPU	**400 Mhz** Atheros 9130-BC1E
Flash	**8 MB** cFeon EN25P64
RAM	**32 MB** Samsung K4H561638J
Ethernet	100 Mbps RTL8306SD
Wireless	Atheros 9102 802.11 b/g/n (integrated)
Serial / JTAG	Yes / Yes
USB	Yes 1x 2.0

Netgear WNDR3700 v2

Architecture	**MIPS** Atheros AR7161
CPU	**680 Mhz** Atheros 9130-BC1E
Flash	**16 MB** Macronix MX25L12845EWI-10G
RAM	**64 MB** 2 x Nanya NT5DS16M16CS-5T
Ethernet	1 Gbps RTL8366SR
Wireless	Atheros AR9223 802.11b/g/n + Atheros AR9220 802.11a/n
Serial / JTAG	Yes / Yes
USB	Yes 1x 2.0

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Mobile Node

LISP-MN Mobility: NAT Traversal Overview/Data Plane…

- Website: http://lispmob.org/
- GIThub: https://github.com/LISPmob/
- Mailing lists:
  - announce@lispmob.org
  - devel@lispmob.org
  - users@lispmob.org
- IRC: #lispmob channel on Freenode
- Twitter: https://twitter.com/LISPmob

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Status

LISP Status

IETF LISP WG: http://tools.ietf.org/wg/lisp/

LISP RFCs and Drafts…

RFCs	
Locator/ID Separation Protocol (LISP) base document	**RFC 6830**
LISP Map Server	**RFC 6833**
LISP Interworking	**RFC 6832**
LISP Multicast	**RFC 6831**
LISP Internet Groper	**RFC 6835**
LISP Map Versioning	**RFC 6834**
LISP+ALT	**RFC 6836**
LISP MIB	**RFC 7052**

Draft	Target
LISP Canonical Address Format (draft-ietf-lisp-lcaf-03)	Active Working Group Document
LISP Deployment (draft-ietf-lisp-deployment-11)	Active Working Group Document
LISP SEC (draft-ietf-lisp-sec-05)	Active Working Group Document
LISP DDT (draft-fuller-lisp-ddt-01)	Active Working Group Document
LISP Mobile Node (draft-meyer-lisp-mn-09)	Related Working Group Document
LISP NAT-Traversal (draft-ermagan-lisp-nat-traversal)	Related Working Group Document
LISP GPE (draft-lewis-lisp-gpe)	Related Working Group Document

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Status

LISP Deployments - International LISP Beta Network…

- LISP Community Operated:
  - More than 5+ years of operation…
  - More than ~600 Sites, 35 countries…
- Interoperable LISP implementations:
  - Cisco
    - IOS (ISR, ISRG2, 7200) and IOS-XE (ASR1K)
    - Cisco IOS-XR (CRS3, ASR9K (beta))
    - Cisco NX-OS (N7K, C200)
  - **AVM "FRITZ!Box"**
  - **OpenWrt**
  - Open Source
    - FreeBSD: OpenLISP
    - Linux: Aless, LISPmob, OpenWrt
    - Android (Gingerbread)

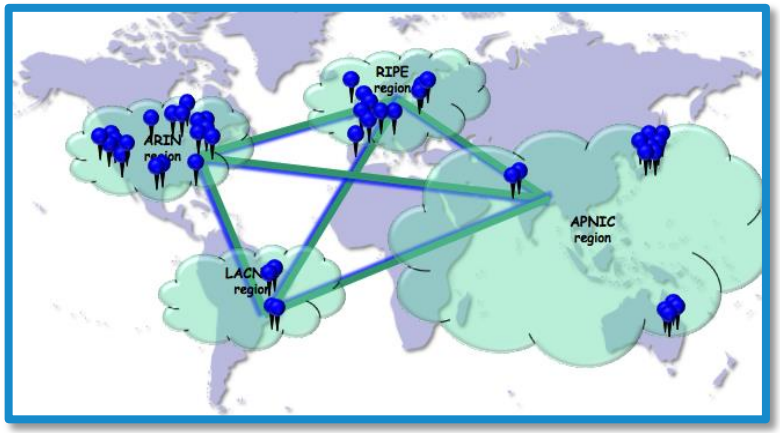

http://lisp.cisco.com

http://www.lisp.intouch.eu/

http://vinciconsulting.com/vxnet

http://lisp.isarnet.net/

Google Microsoft VeriSign at&t NEW YORK UNIVERSITY QUALCOMM SONY Lufthansa free verizon NTT NJEDge.Net UCLA and more…

Plus some others… ;-)

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Status

LISP Software – Available Releases…

- Cisco Releases (http://lisp.cisco.com)

	NX-OS	IOS	IOS-XE	IOS-XR
Software	First Available: 12/2009 Current Main: 6.1(4a) or 6.2(2a)	First Available: 12/2009 Current Main: 15.4(1)T Current Eng: 15.3(3)XB12	First Available: 03/2010 Current Main: 15.3(3)S Current Eng: 15.3(3)S1xb	First Available: 03/2012 Current Main: 4.3.2
Platforms	Nexus 7000 M1-32 linecard	ISR (1800/2800/3800) ISRG2 (800/1900/2900/3900) Catalyst 6500	ASR1K CSR1000V	CRS 3 ASR9k
Features	Roles: ITR/ETR/MS/MR/PITR/PETR AF: EID-v4/v6, RLOC-v4 Virtualisation: Shared/Parallel Mobility: ASM/ESM OTV Multicast: yes	Roles: ITR/ETR/MS/MR/PITR/PETR AF: EID-v4/v6, RLOC-v4/v6 Virtualisation: Shared/Parallel Mobility: ASM/ESM Multicast: roadmap March 2014	Roles: ITR/ETR/MS/MR/PITR/PETR AF: EID-v4/v6, RLOC-v4/v6 Virtualisation: Shared/Parallel Mobility: ASM/ESM OTV Multicast: roadmap Nov 2014	Roles: PITR/PETR AF: EID-v4/v6, RLOC-v4 Virtualisation: Shared/Parallel Mobility: roadmap Multicast: roadmap March 2014

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Summary

LISP References

LISP Information

- **LISP Information**

Cisco LISP Site ……	http://lisp.cisco.com (IPv4 and IPv6)
Cisco LISP Marketing Site ……...	http://www.cisco.com/go/lisp/
LISP Beta Network Site ……	http://www.lisp4.net or http://www.lisp6.net
LISP DDT Root ……...	http://www.ddt-root.org
IETF LISP Working Group ……..…	http://tools.ietf.org/wg/lisp/

- **LISP Mailing Lists**

Cisco LISP Questions ……	lisp-support@cisco.com
IETF LISP Working Group ……	lisp@ietf.org
LISP Interest (public) ……	lisp-interest@puck.nether.net
LISPmob Questions ……...	users@lispmob.org

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Summary

Part of the LISP Solution Space

1. Multihoming
2. IPv6 Transition
3. Virtualisation/VPN
4. Mobility

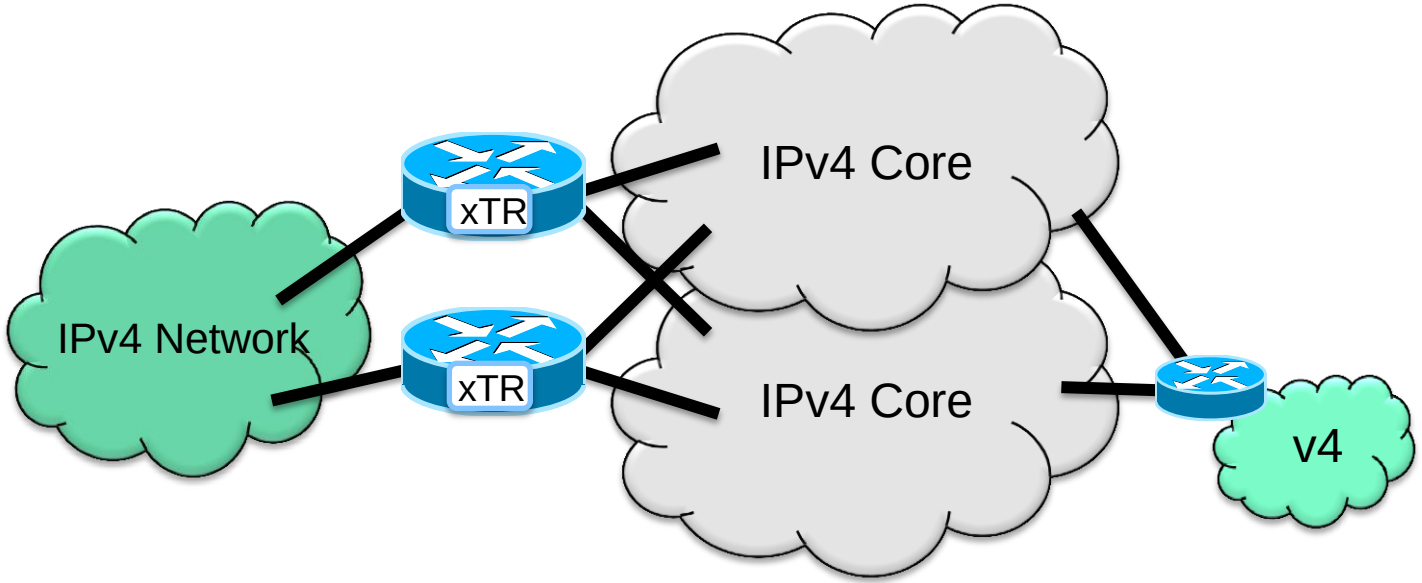

LISP is an Architecture…

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Summary

Part of the LISP Solution Space

1. Multihoming
2. IPv6 Transition
3. Virtualisation/VPN
4. Mobility

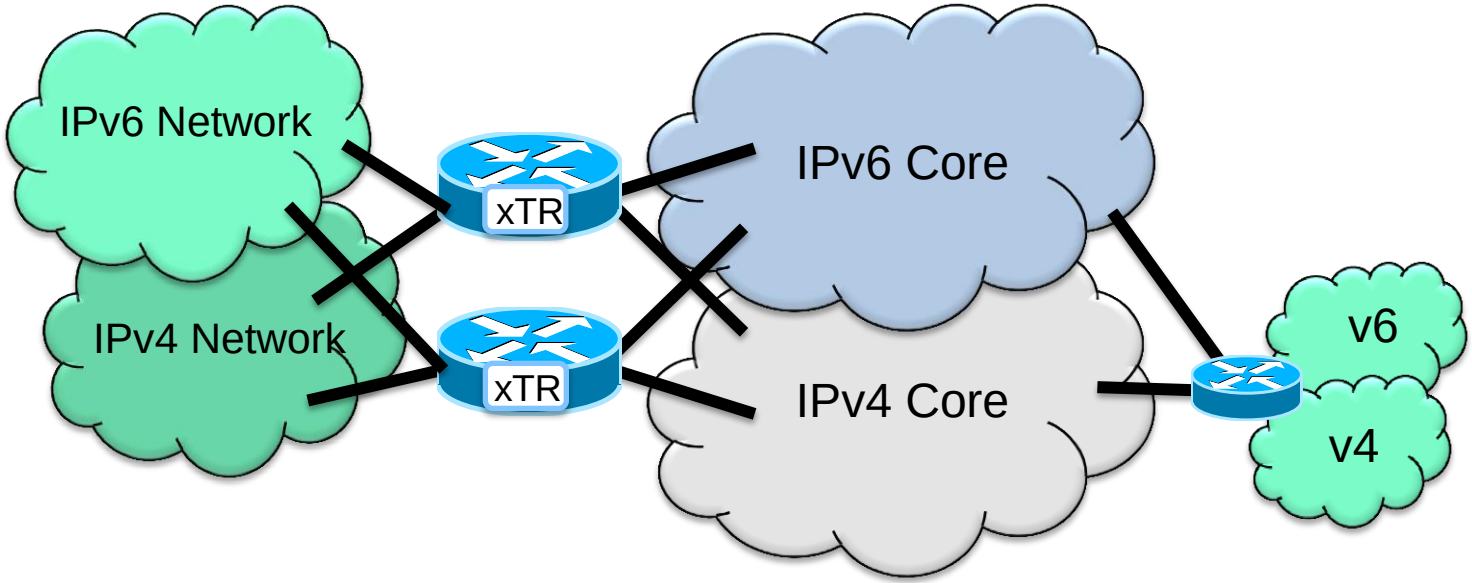

LISP is an Architecture…

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Summary

Part of the LISP Solution Space

1. Multihoming
2. IPv6 Transition
3. Virtualisation/VPN
4. Mobility

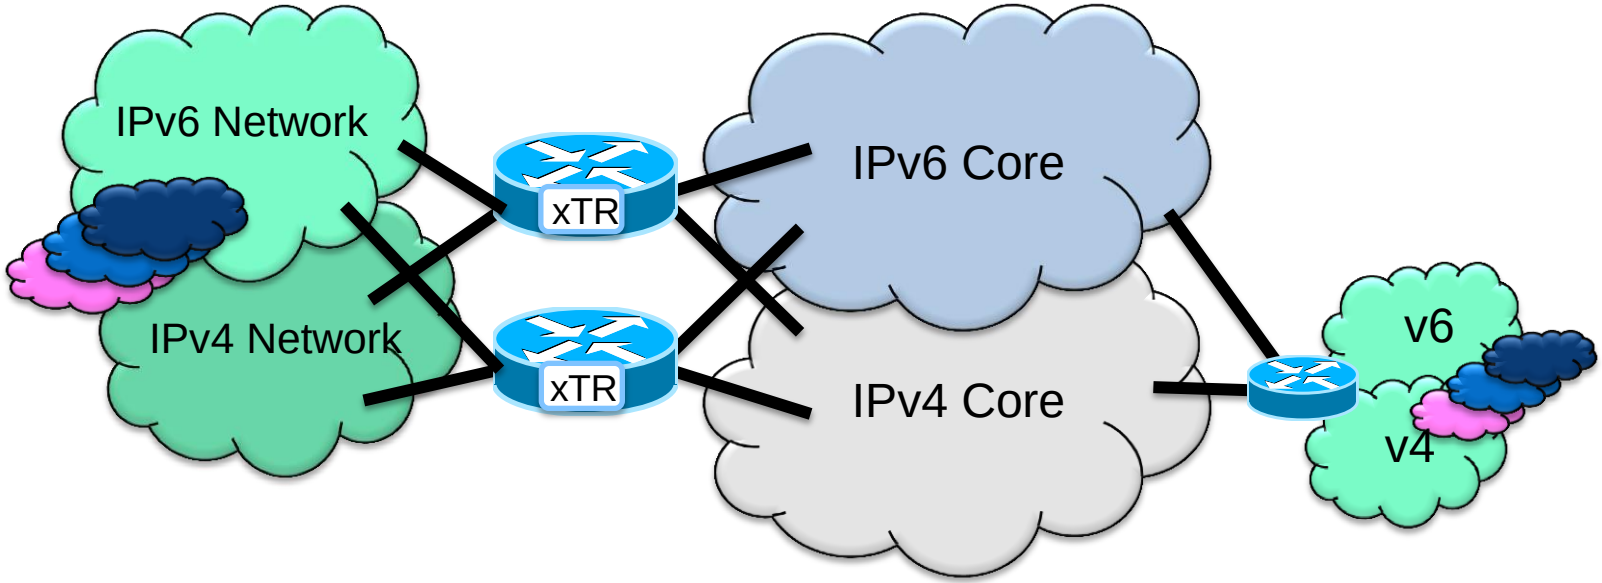

LISP is an Architecture…

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Summary

Part of the LISP Solution Space

1. Multihoming
2. IPv6 Transition
3. Virtualisation/VPN
4. Mobility

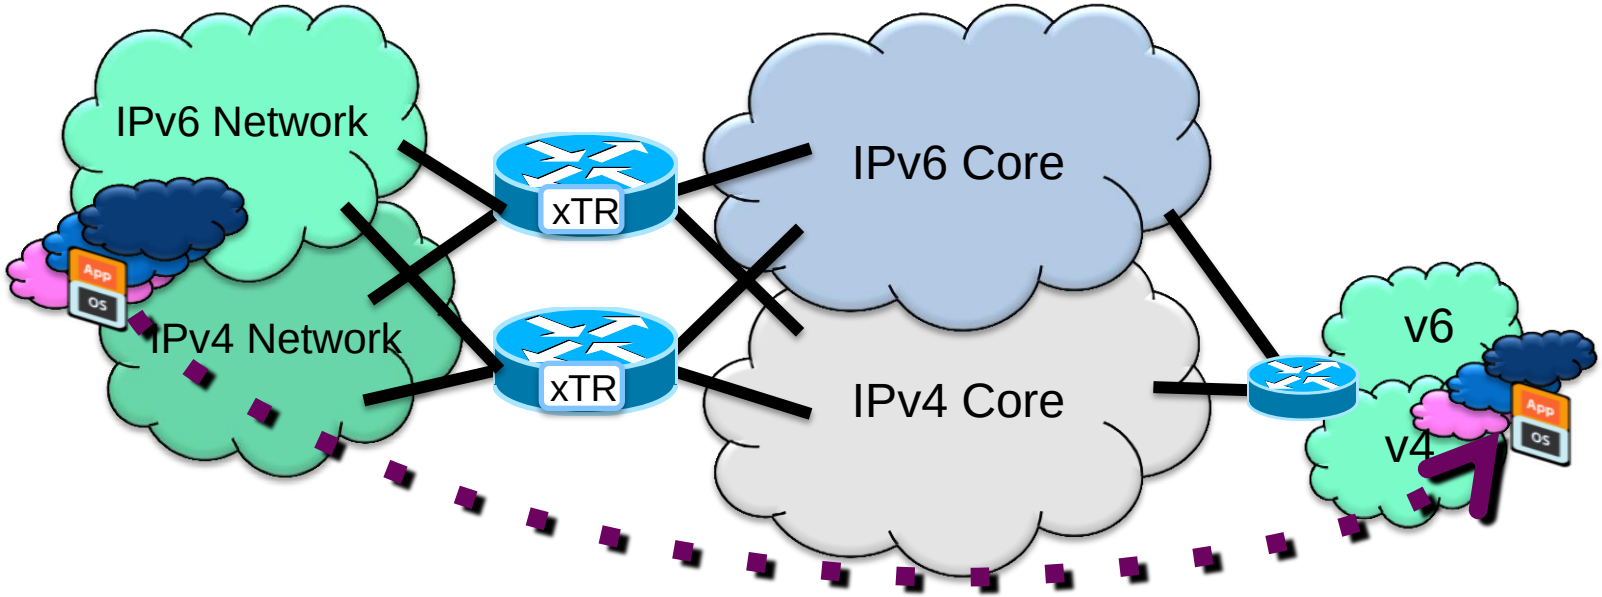

LISP is an Architecture…

© 2014 Cisco and/or its affiliates. All rights reserved.

LISP Overview

LISP :: A Routing Architecture – Not a Feature

- Uses **pull** vs. **push** routing
  -
  -
- An **over-the-top** technology
  - Address Family agnostic
  - Incrementally deployable
  - End systems can be unaware of LISP
- Deployment **simplicity**
  - No host changes
  - Minimal CPE changes
  - Some new core infrastructure components
- LISP use-cases are **complimentary**
  - Simplified multi-homing with Ingress traffic Engineering; no need for BGP
  - Address Family agnostic support
  - Virtualisation support
  - End-host mobility without renumbering
- Enables **IP Number Portability**
  - Never change host IP's; No renumbering costs
  - No DNS changes; "name == EID" binding
  - Session survivability
- An **Open Standard**
  - Being developed in the IETF (RFC 6830-6836)
  - No Cisco Intellectual Property Rights

© 2014 Cisco and/or its affiliates. All rights reserved.




Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm

Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations.
www.CiscoLiveAPAC.com

© 2014 Cisco and/or its affiliates. All rights reserved.


CISCO
TM