
TOMORROW starts here.
cisco
Cisco live!

Deployment Challenges with Interconnecting Data Centres

BRKDCT-3060

David Jansen
DSE
CCIE #5952

Cisco live!

Reference Sessions

- BRKDCT-2049 - Overlay Transport Virtualisation
- BRKDCT-2615 - Active-Active Data Centre Strategies
- BRKDCT-2385 - Cisco Dynamic Fabric Automation Architecture
- BRKRST-3045 - LISP - A Next Generation Networking Architecture

BRKDCT-3060 Abstract

This advanced session discusses the challenges and recommended solutions of extending LAN connectivity between geographically dispersed data centres. Innovations in middleware like 'Virtual-Machines' and 'Servers-clustering are revolutionising Virtualisation of Data Centre, while other key IT processes like 'Disaster Recovery Plan' and 'massive servers migration require optimisation and facilitation. Data-centre is now more and more spreading across multiple sites, and one very difficult point to solve is the extension of VLAN in a large scale with respect to Spanning-Tree stability requirement. The different requirements for providing a robust LAN extension solution will be discussed during this session, including end-to-end loop prevention, multi-homing considerations and optimal bandwidth utilisation. Detailed design guidance will be provided around the deployment of Ethernet based technologies, leveraging Multi Chassis EtherChannel functionalities like VSS and vPC, as well as MPLS based technologies (EoMPLS and VPLS) and an innovative IP based technology called Overlay Transport Virtualisation (OTV). Locator Identity Separation Protocol (LISP) will then be introduced as an emerging technology capable of providing both IP Mobility and Path Optimisation functionalities. This advanced session is intended for network design and operation engineers from Enterprises, Service Providers or Enterprise Hosting Service Providers that are willing to solve this difficult and controversial problem of Data-Centre Interconnect.

Goals of This Session…

- Highlighting the main business requirements driving Data Centre Interconnect (DCI) deployments
- Understand the functional components of the holistic Cisco DCI solutions
- Get a full knowledge of Cisco LAN extension technologies and associated deployment considerations
- Integrate routing aspect induced by the emerging application mobility offered by DCI
- This session does not include:

  Storage extension considerations associated to DCI deployments

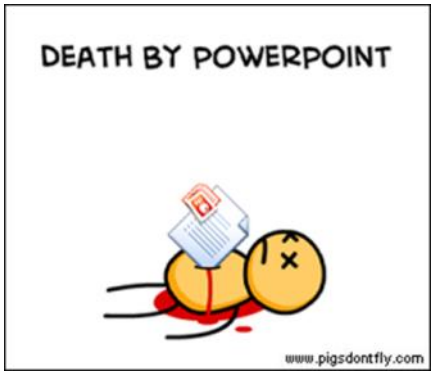

Data Centre Interconnect
Agenda

- Mobility and Virtualisation in the Data Centre
- LAN Extension Deployment Scenarios
  - Ethernet Based Solutions
  - MPLS Based Solutions
    - EoMPLS
    - VPLS
    - A-VPLS
    - EVPN
  - IP Based Solutions
- Encryption
- Path Optimisation
- IP Mobility without LAN Extension
- Fabric Solutions
- Summary and Conclusions
- Q&A

= For your Reference

Mobility and Virtualisation in the Data Centre

Distributed Data Centres
Building the Data Centre Cloud

- Distributed Data Centre Goals
- Seamless workload mobility
- Distributed applications
- Business Continuity
- Operational-models
- Failure containment
- Application Performance
- Application Availability
- Maximise Resource Usage
- Security End-to-End

Geographically Disperse Data Centres

DC Interconnect: End-to-End Requirements

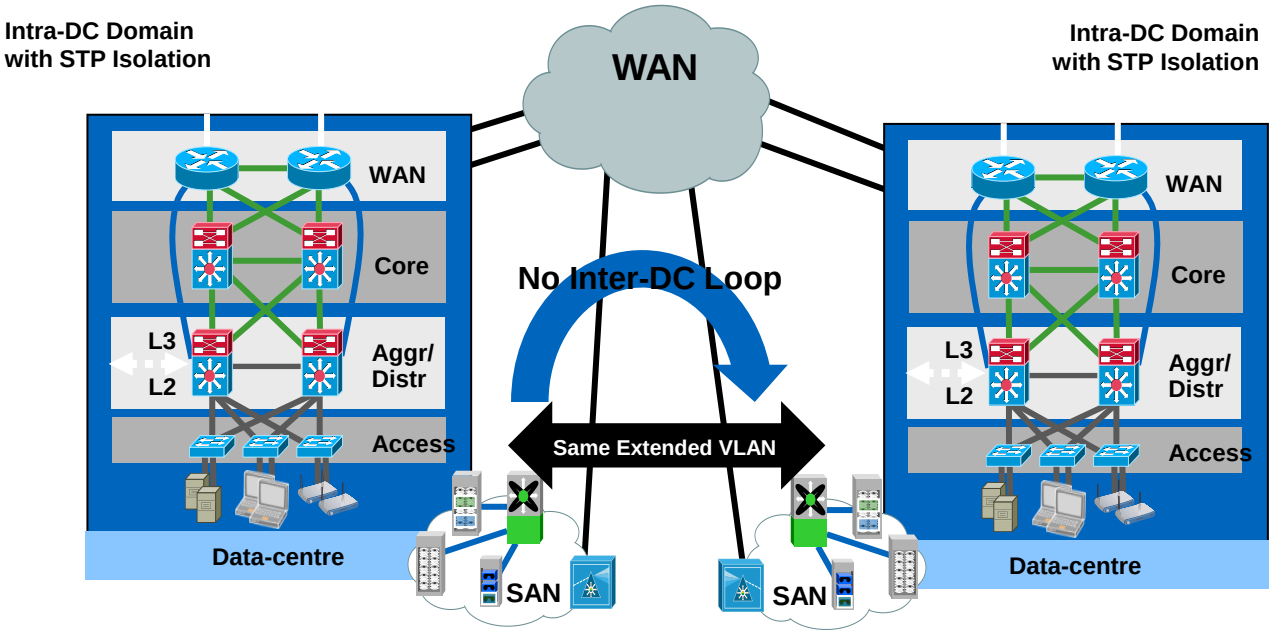

- STP problems across DCI interconnect
- control-plane isolation is desired
- Risk of Inter-DC Loops due to Intra-DC STP Isolation
- End-to-End Loop Avoidance

Solution requirements

- ❑ E2E Loop Prevention
- ❑ STP Isolation
- ❑ Redundant LAN extn.
- ❑ WAN load balancing
- ❑ Core Transparency
- ❑ DC site Transparency
- ❑ Optimal Traffic Handling
- ❑ VLAN Scalability
- ❑ Multi-site Connectivity

Optional

- ❑ Encryption
- ❑ HQOS

Any workload, anywhere, on a high capacity fabric

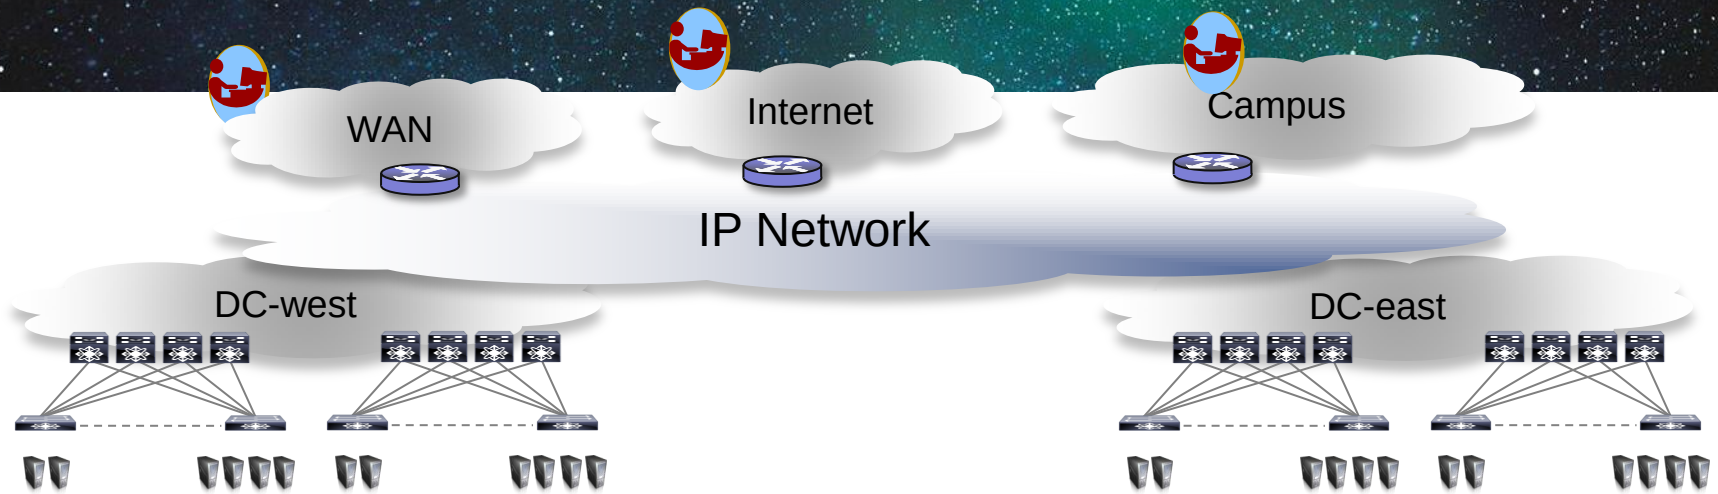

Manageability, Programmability

Workload **Mobility**	Workload **Distribution**	Workload **Segmentation**
Any host anywhere	Stretch clustered apps	Host multiple tenants / cloud

Seamless WAN integration

Network: High Capacity, Multi-Path, Failure Containment

Moving vs. Distributing Workloads

- Why do we really need LAN Extensions?

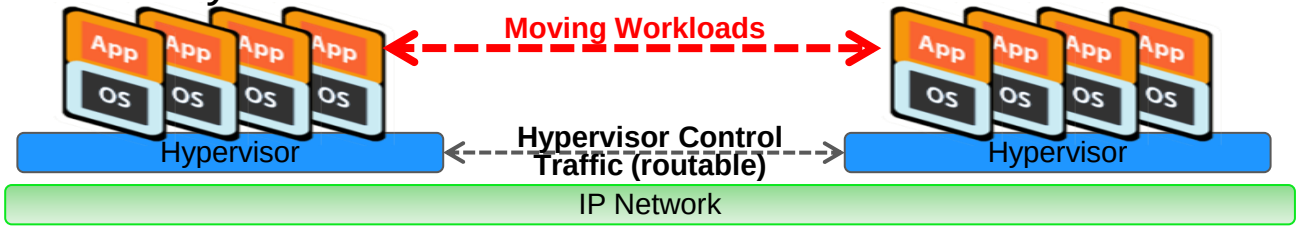

  - **Move workloads** with IP mobility solutions: LISP Host Mobility
    - IP preservation is the real requirement (LAN extensions not mandatory)

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

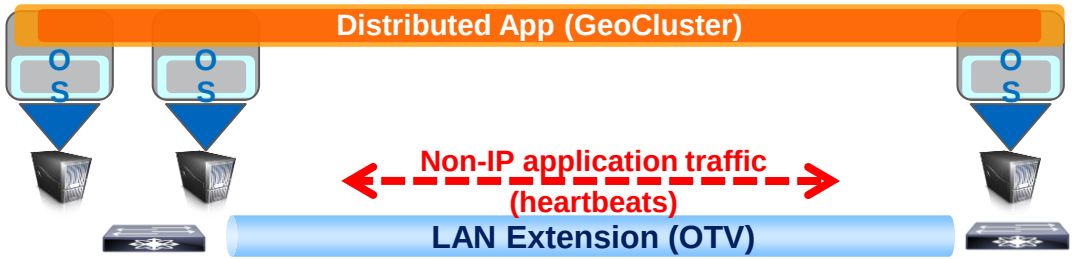

Layer 2 / IP Mobility Use Cases

- Extending Operating System / File System clusters
- Extending Database clusters
- Virtual machine mobility
- Physical machine mobility
- Physical to Virtual (PtoV) Migrations
- Legacy devices/apps with embedded IP addressing
- Time to deployment and operational reasons
- Extend DC to solve power/heat/space limitations
- Data Centre co-location

Layer 2 Risks

- Flooding of packets between Data Centres
- Spanning Tree (STP) is not easily scalable and risk grows as diameter grows
- STP has no domain isolation – issue in single DC can propagate
- First hop resolution and inbound service selection can cause verbose inter-Data Centre traffic
- In general Cisco recommends L3 routing for geographically diverse locations
- This session focuses on making limited L2 connectivity as stable as possible

Maximum Transmission Unit (MTU) Guidance:

- EoMPLS Port Mode: 1522 Bytes
- EoMPLS VLAN Mode: 1526 Bytes
- VPLS: 1526 Bytes (1530 Bytes with control-word)
- A-VPLS: 1530 with flow-label (3rd Label), (1534 with control-word)
- OTV: 1542 Bytes
- LISP
  - IPv4 1536 Bytes
  - IPv6 1556 bytes

Maximum Transmission Unit (MTU) Guidance:

- FabricPath: 1516 Bytes
- VXLAN: 1550 Bytes
- GRE: 1524 Bytes
- 802.1ae: 1540 Bytes
- IPSEC: 1574 Bytes

Data Centre Interconnect
Agenda

- Mobility and Virtualisation in the Data Centre
- LAN Extension Deployment Scenarios
  - Ethernet Based Solutions
  - MPLS Based Solutions
    - EoMPLS
    - VPLS
    - A-VPLS
    - EVPN
- Overlay Transport Virtualisation (OTV)
- Encryption
- Path Optimisation
- IP Mobility without LAN Extension
- Fabric Solutions
- Summary and Conclusions
- Q&A

= For your Reference

Ethernet Based Solutions

Layer 2 Extension Without Tunnels/Tags (vPC/VSS)

- 6500 with Virtual Switching System cluster (Supported distances at 80km (ZR) Dark Fibre)
- Nexus 7000 with Virtual Port-Channels (Supported distances at 80km (ZR-X2) Dark Fibre)
- All traffic flows to a vPC/VSS member node
- Hub-and-spoke topology from a layer 2 perspective
- Dedicated links to vPC/VSS members from each Data Centre aggregation switch
- Can consume lambda or Fibre strands quickly
- Data plane rate limiting in L2 still needs protection
- STP domains are not isolated unless we BPDU-filter at all vPC/VSS aggregation switches

vPC / VSS Design

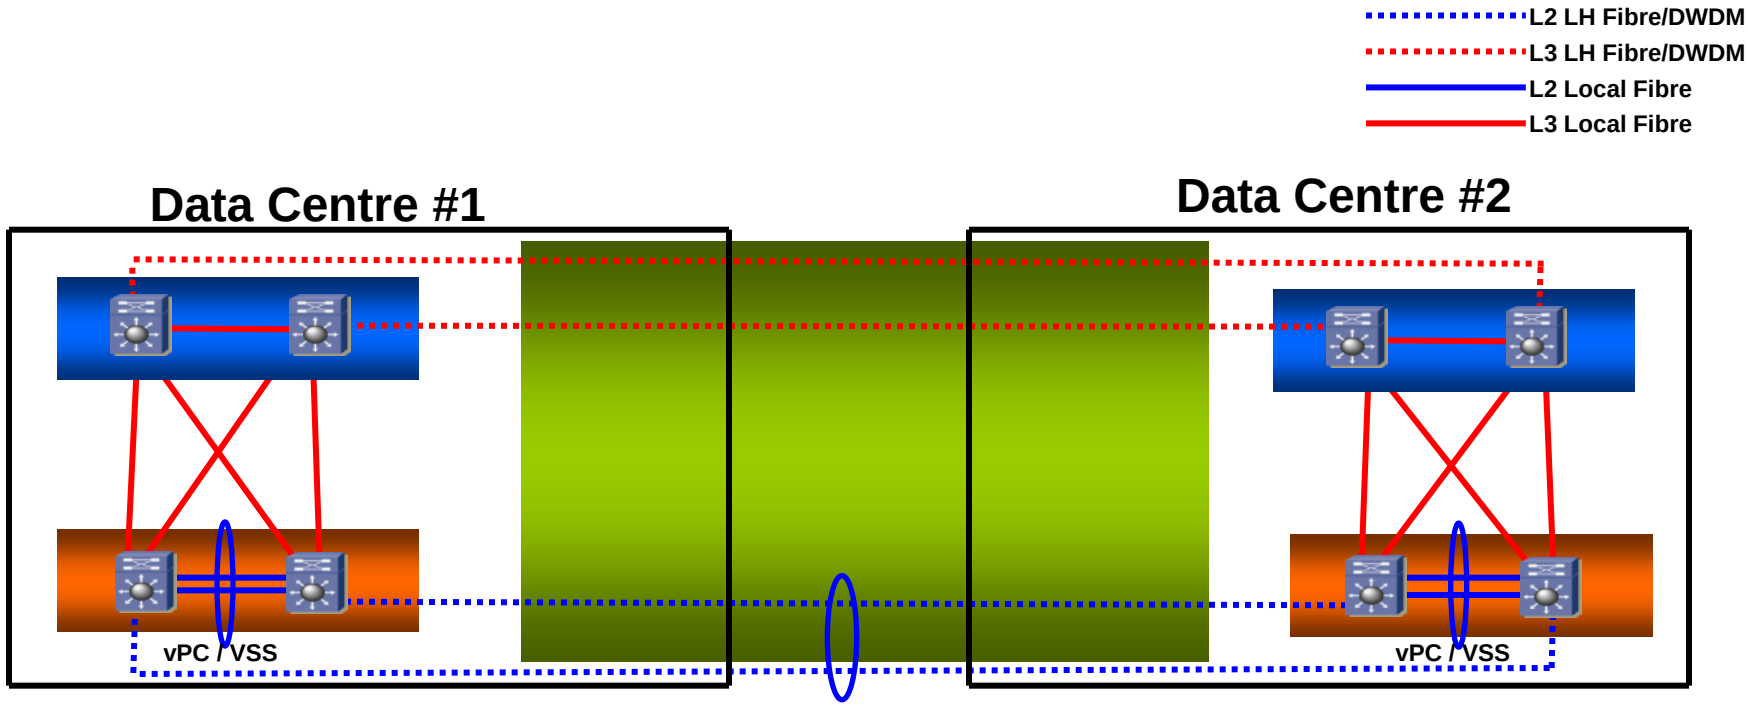

vPC / VSS L2 View

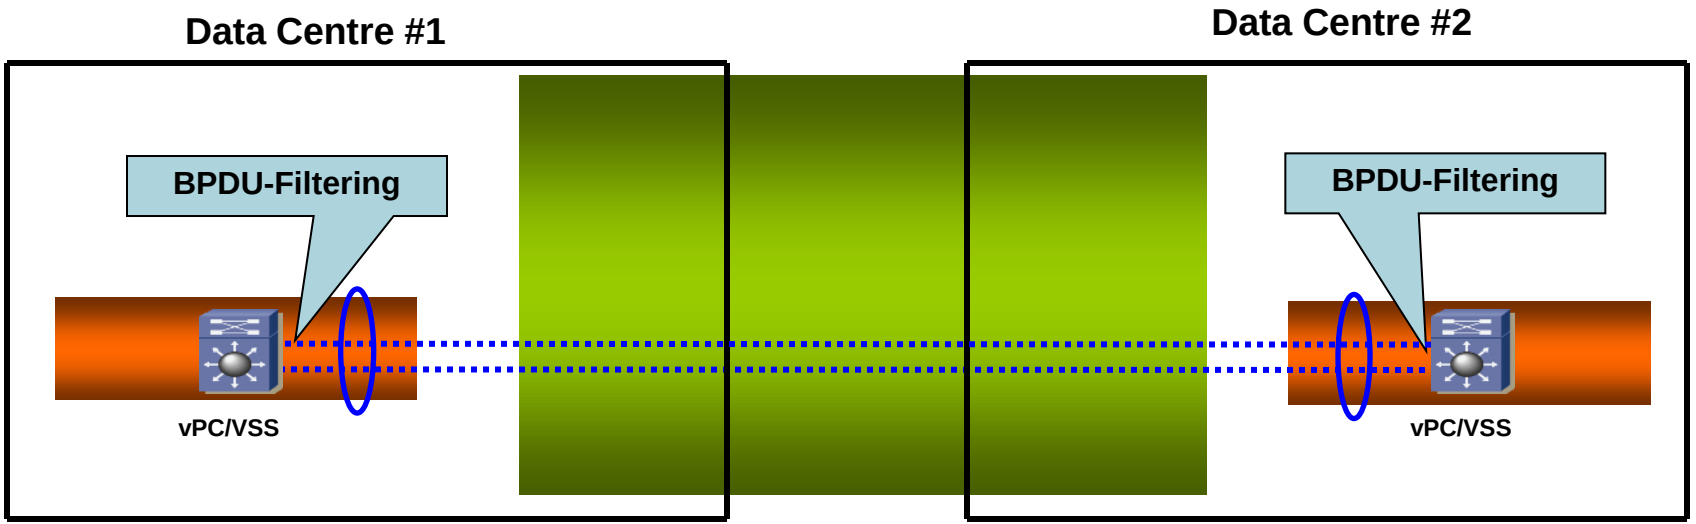

- vPC/VSS Domain ID for facing vPC/VSS layers should be different aggregation
- BPDU Filter on the edge devices to avoid BPDU propagation peer-link
- STP Edge Mode to provide fast failover times
- No Loop must exist outside the vPC/VSS domain
- No L3 peering between Nexus 7000 devices (i.e. pure layer 2)

- Configure root guard on
- bridge-assurance only on vPC

L2 LH Fibre/DWDM

L2 Local Fibre

vPC / VSS Design

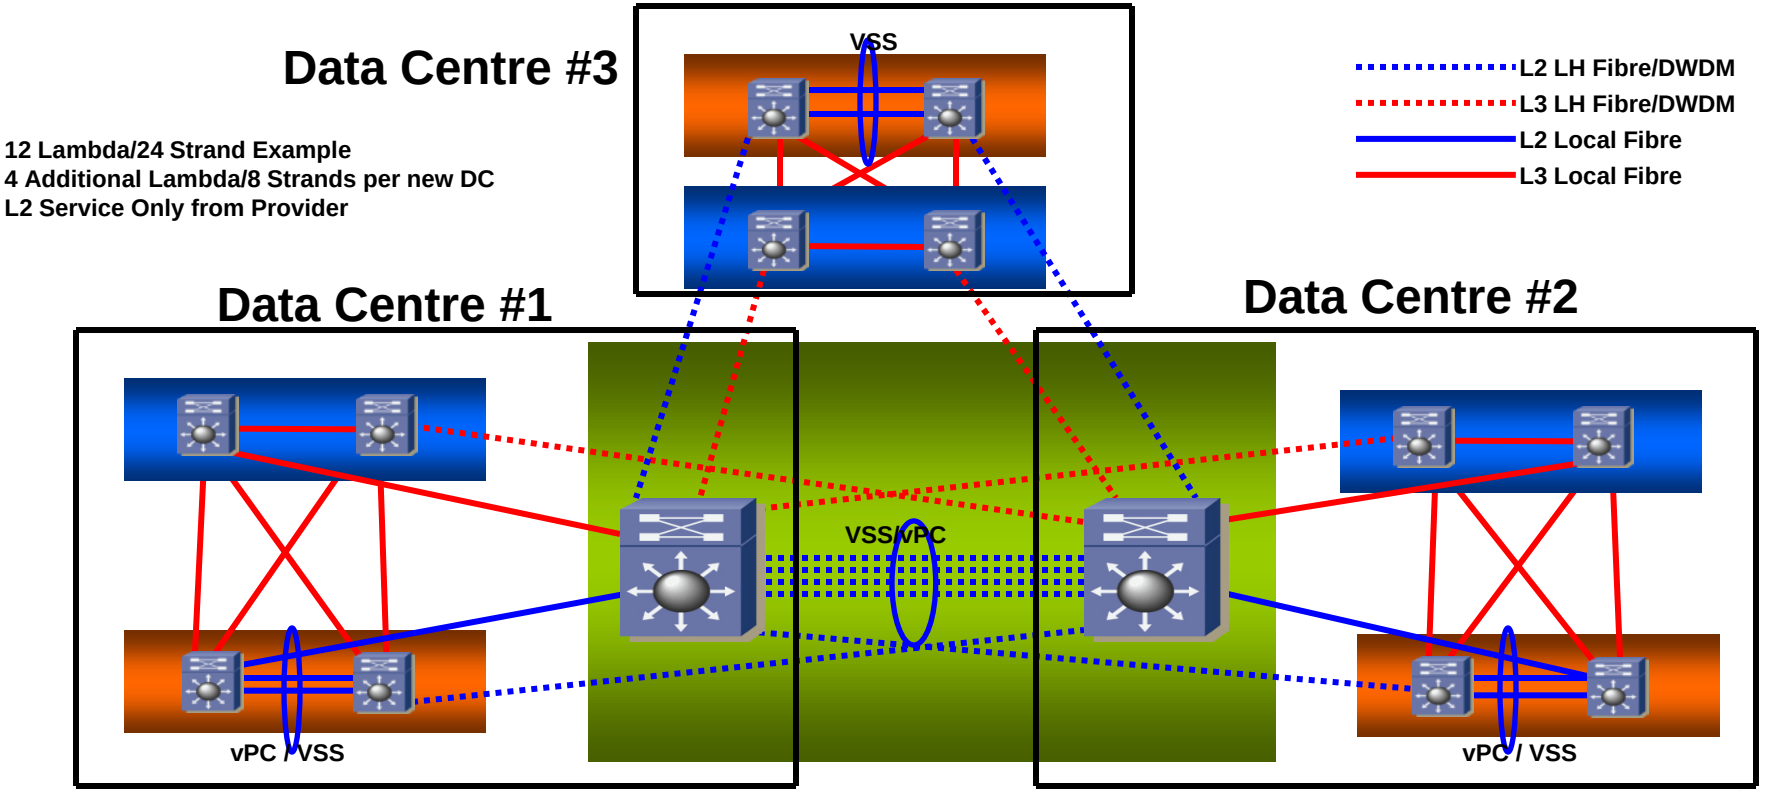

vPC / VSS L2 View

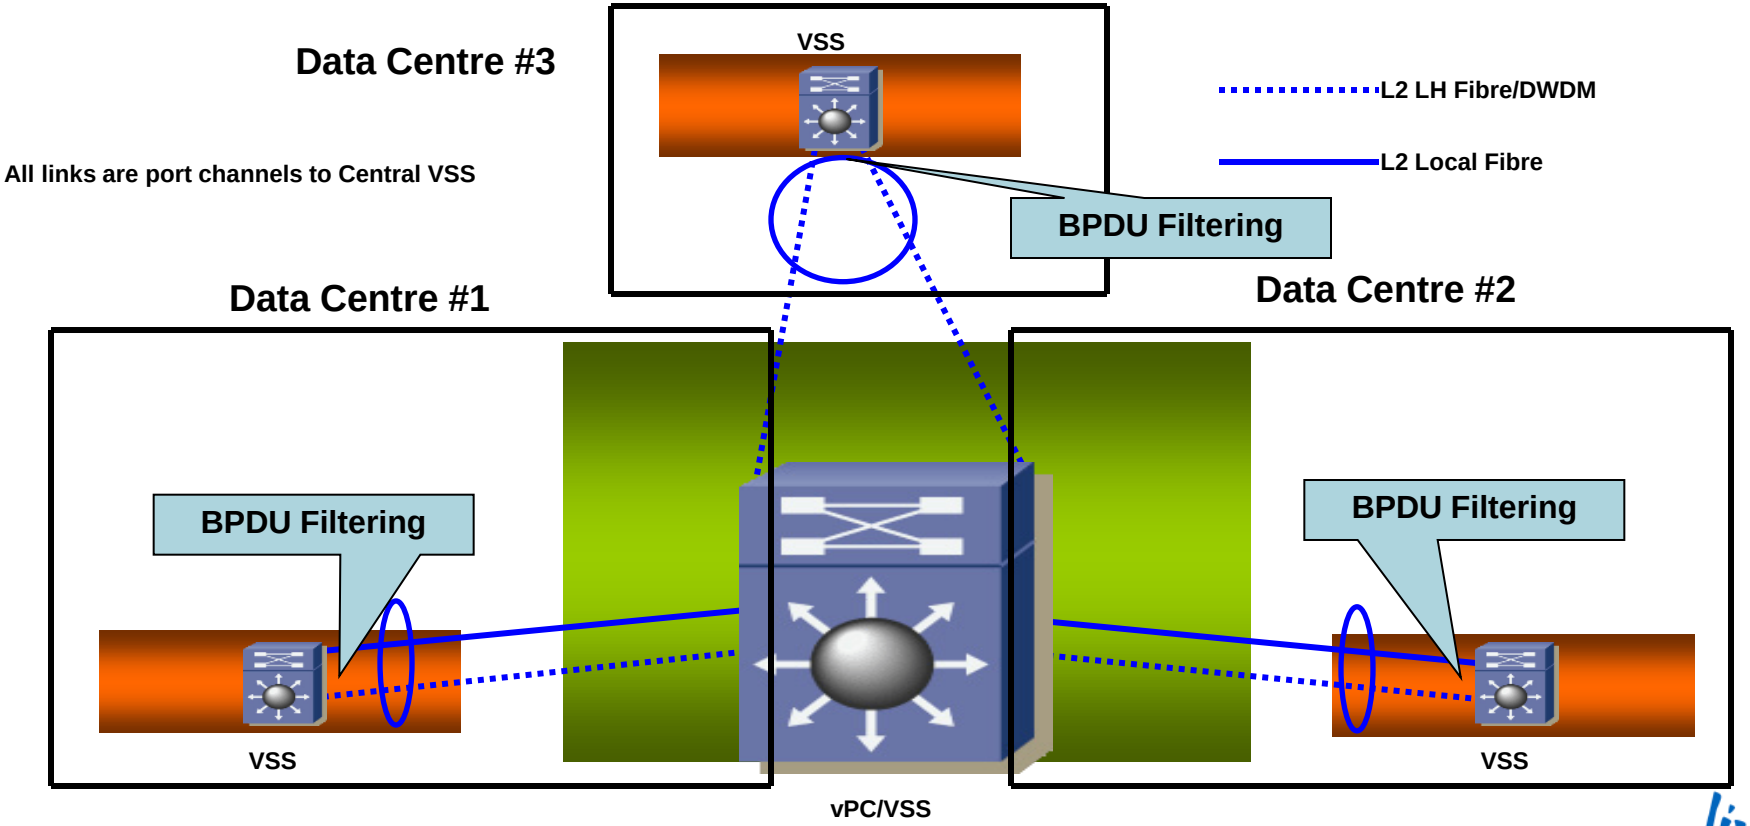

vPC and Layer 3

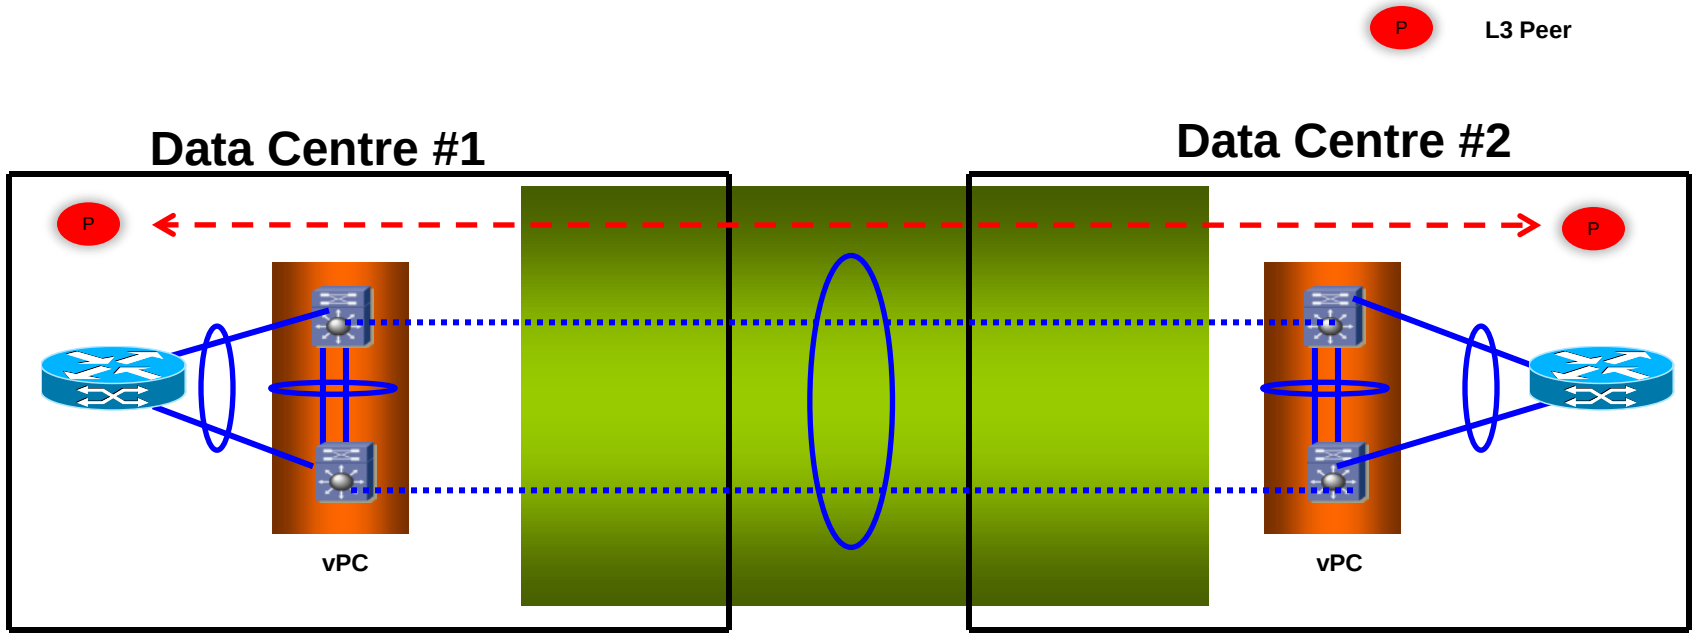

- Nexus 7000 configured for L2 Transport only
- SVI passive-interface (no IGP peering)

vPC and Layer 3

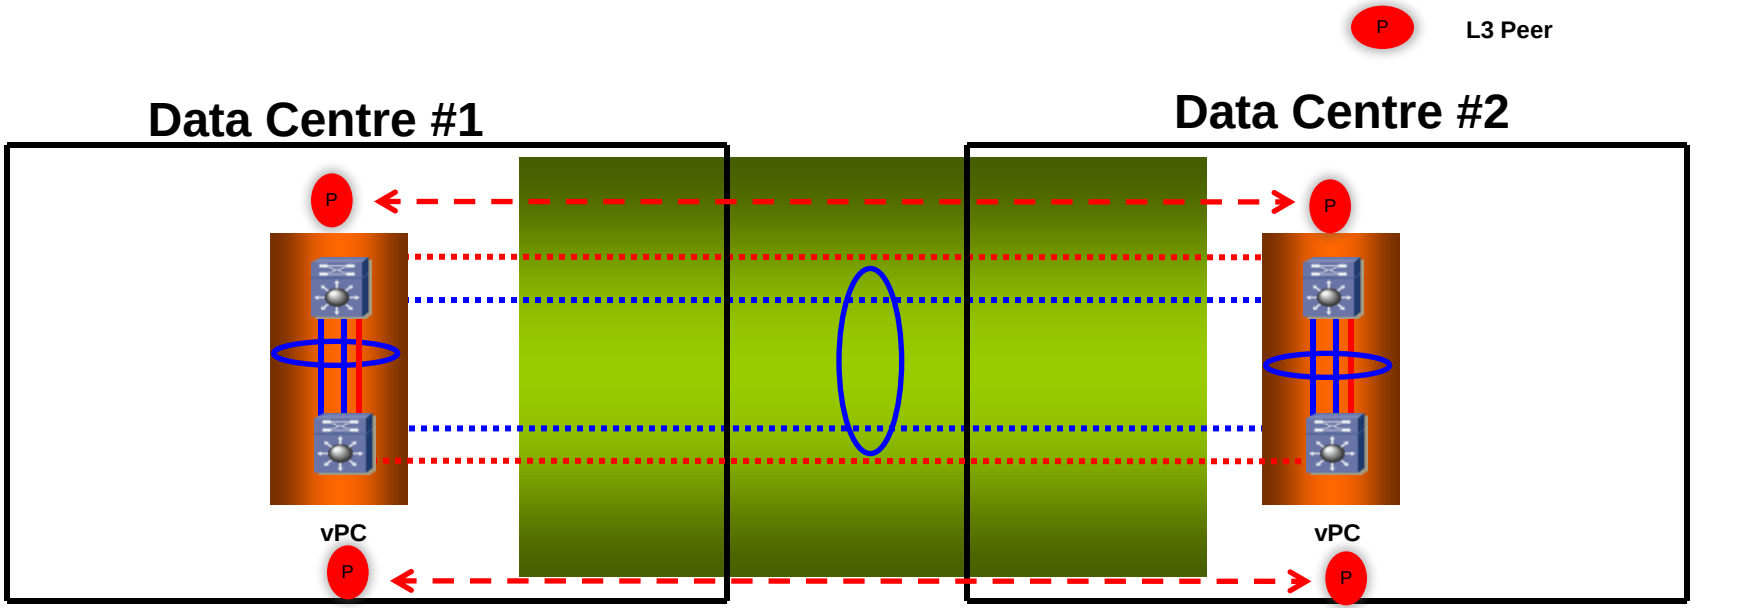

- Peering over a vPC inter-connection on parallel routed interfaces
- SVI passive-interface (no IGP peering)

FabricPath Design (Partial/Full/Ring Topology)

- Leverage vPC+
- Brownfield / Greenfield DC
- STP Integration
- Conversational MAC Learning
- Native VLAN Pruning
- TTL / RPF
- ECMP for L2

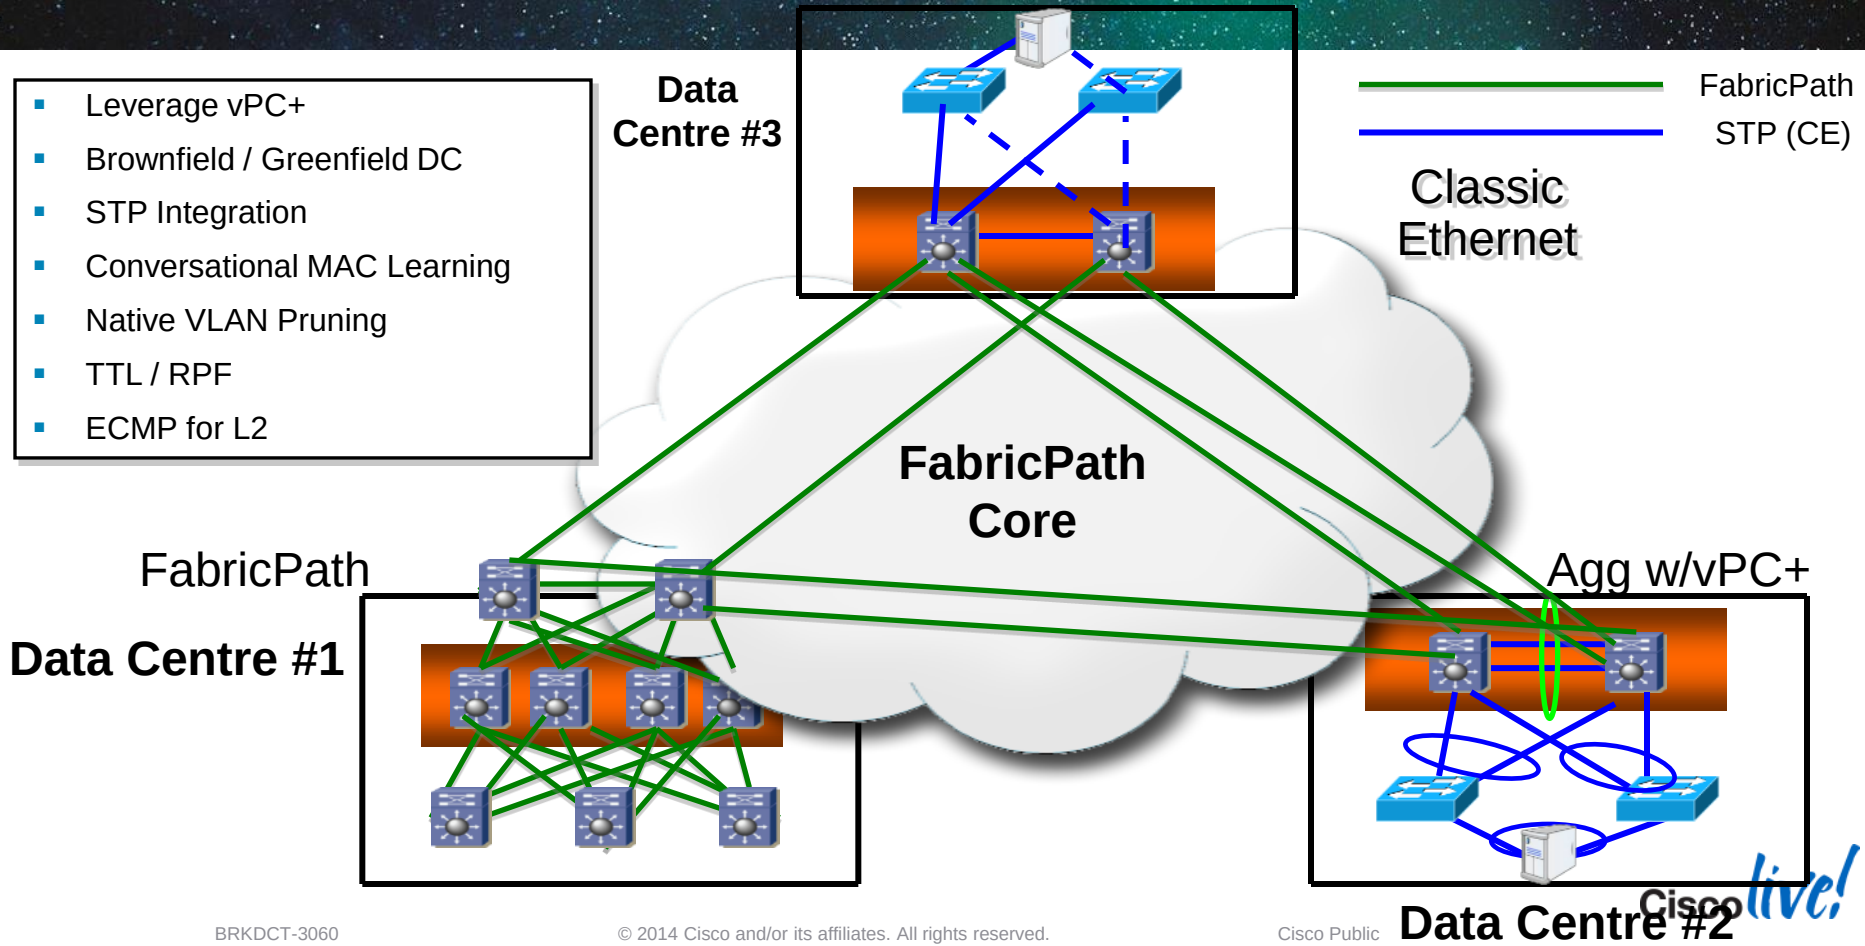

FabricPath for DCI:

- FabricPath L2 ISIS adjacencies are Point to Point
  - Need for direct Point to Point L1 WAN Links
  - FabricPath over VPLS is not supported
  - L2 managed service: Dark Fibre, DWDM, EoMPLS
  - MTU requirements: 16 extra Bytes for FabricPath header
- BFD not supported
- Multi-destination Traffic: Multicast/ARP traffic across DCI can be non-optimal due to MDT (Multi-destination tree)
- FabricPath and HSRP localisation (no solution today)
- STP Integration with FP, need to make FP STP-Root
- Anycast

Data Centre Interconnect
Agenda

- Mobility and Virtualisation in the Data Centre
- LAN Extension Deployment Scenarios
  - Ethernet Based Solutions
  - MPLS Based Solutions
    - EoMPLS
    - VPLS
    - A-VPLS
    - EVPN
- Overlay Transport Virtualisation (OTV)
- Encryption
- Path Optimisation
- IP Mobility without LAN Extension
- Fabric Solutions
- Summary and Conclusions
- Q&A

= For your Reference

MPLS Based Solutions

EoMPLS (Ethernet Over MPLS)

- Encapsulates Ethernet frames inside MPLS packets to pass layer 3 network
- EoMPLS has routing separation from metro core devices providing connectivity – CE flapping routes won't propagate inside MPLS
- Point to point links between locations
- Data plane rate limiting in L2 still needs protection

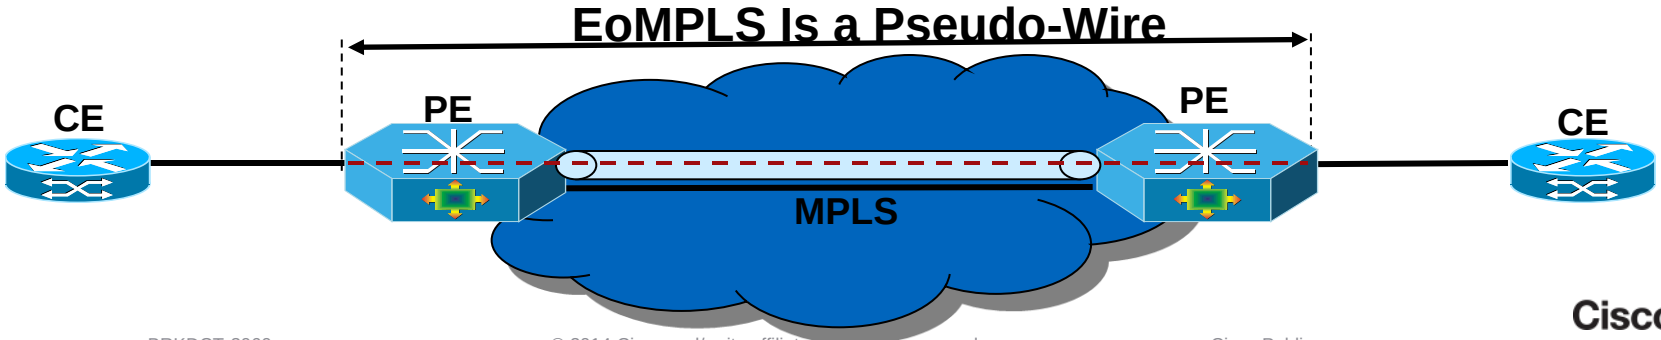

EoMPLS Usage for DCI

End-to-End Loop Avoidance with Active-Active redundancy with fast convergence

On DCI Etherchannel:
- STP Isolation (BPDU Filtering)
- Broadcast Storm Control
- FHRP Isolation

- Port-mode EoMPLS
- EoMPLS remote-port shut down

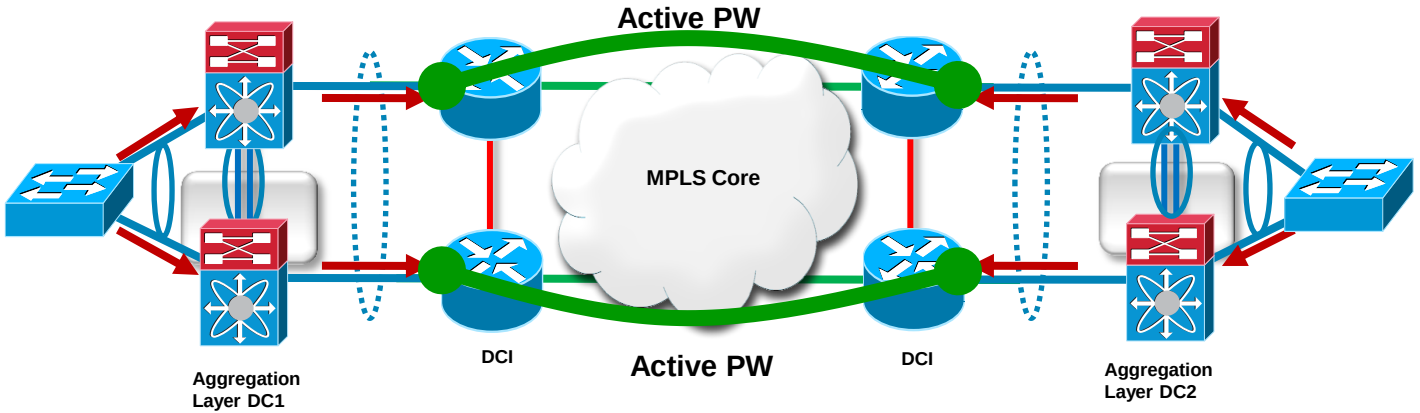

Active/Active vPC or VSS

Active/Active vPC or VSS

Multi-Point Topologies
What is VPLS?

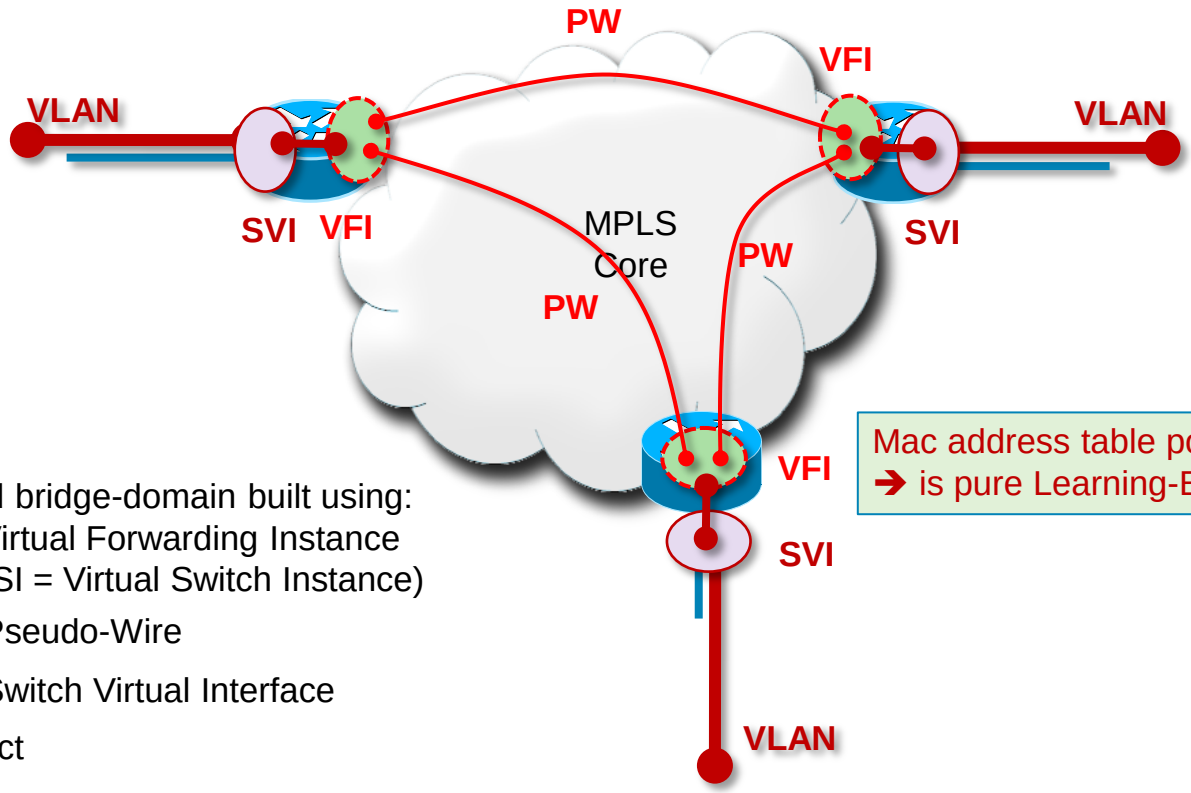

Mac address table population
➔ is pure Learning-Bridge

One extended bridge-domain built using:

- VFI = Virtual Forwarding Instance
  (VSI = Virtual Switch Instance)
- PW = Pseudo-Wire
- SVI = Switch Virtual Interface
- xconnect

VPLS Cluster Solutions

- Using clustering mechanism
  - Two devices in fusion as one
    - VSS Sup720 ES-Modules
    - VSS Sup2T
    - ASR9K nV virtual cluster
    - ➔ One control-plane / two data-planes

- Dual node is acting as one only device
  - Native redundancy (SSO cross chassis)
  - Native load balancing
  - Capability to use port-channel as attachment circuit

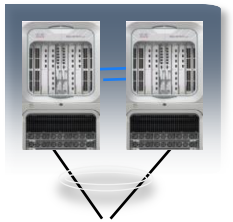

Deployment Example – L2VPN Service

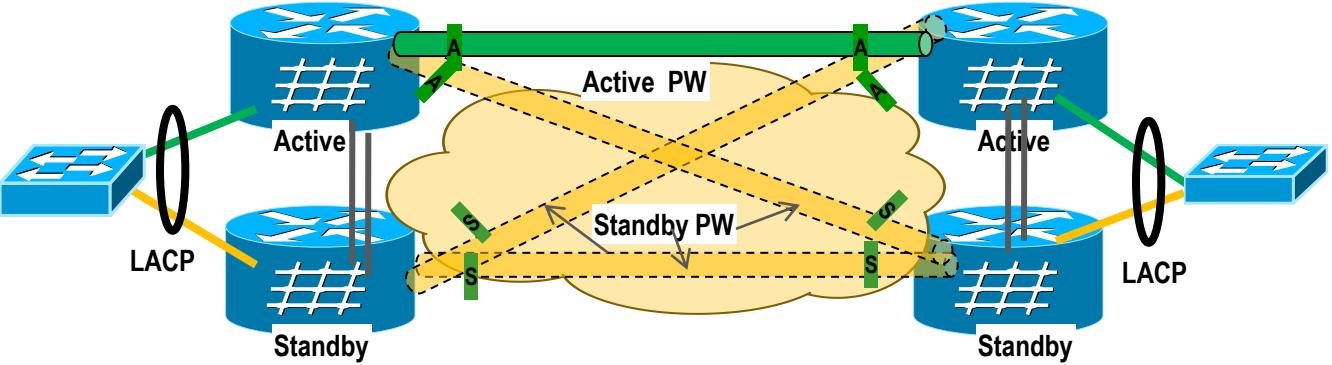

Solution1: MC-LAG + 2-way PW redundancy

- Active/standby MC-LAG → bandwidth inefficiency
- 4 PWs with 3 standby → control plane overhead
- PW failover time depends on the number of PWs → slow convergence
- Require additional state sync (for example, IGMP Snooping table) to speed up service convergence → complex

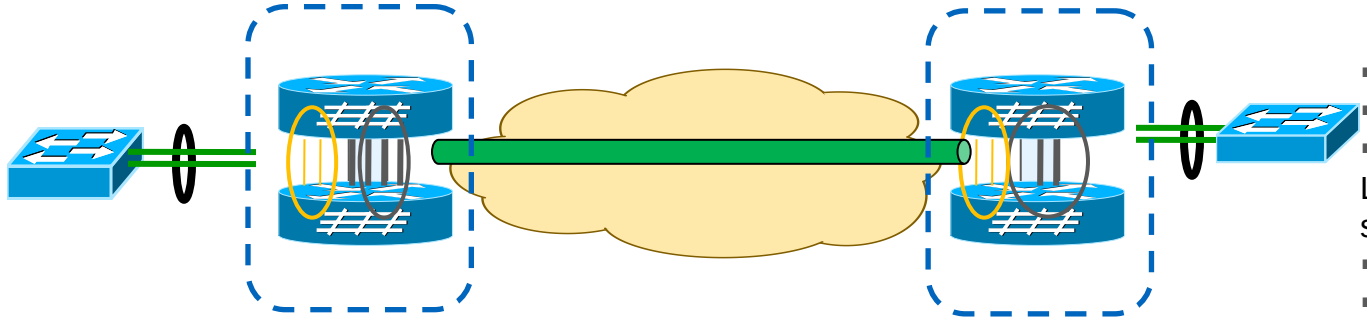

Solution 2: ASR 9000 nV Cluster

- Active/active regular LAG
- Single PW
- Link/Node failure is protected by LAG, PW is even not aware → super fast convergence
- State sync naturally
- Simple, fast solution

VPLS Multi-homing – ASR9K nV Cluster

Simple and faster network convergence

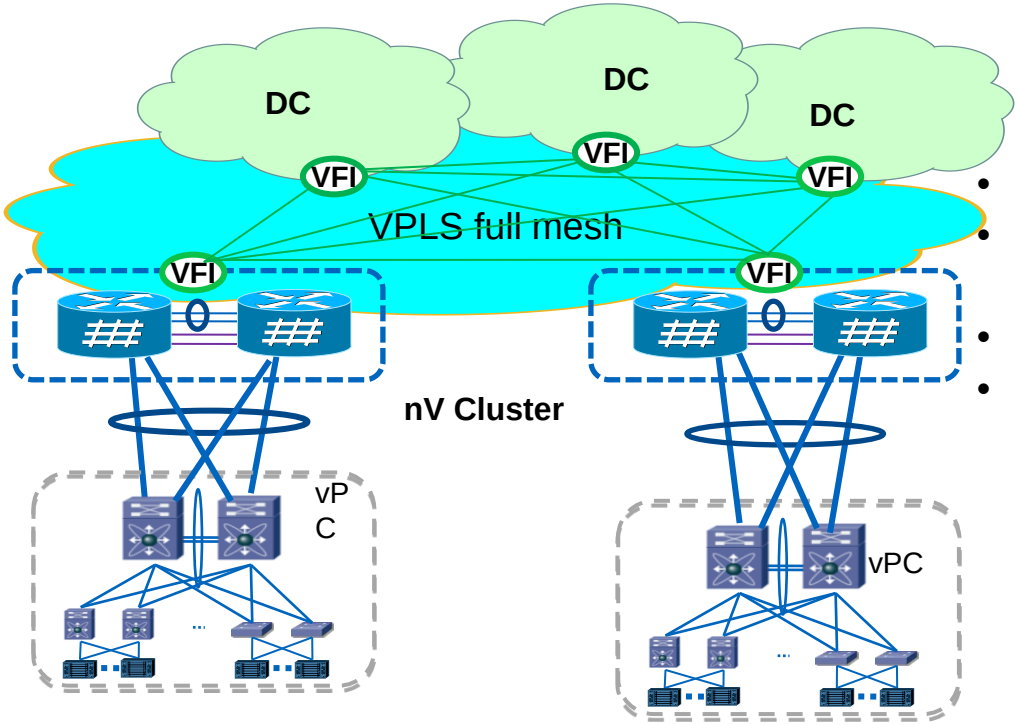

- Reduce the Number of PWs
- Simplify VPLS dual homing with active/active link bundle
- per-flow and per-VLAN load balancing
- Sub-second to 50msec fast convergence

data-plane: port-channel used between the ASR9000 on any 10G or 100G Interfaces.

control-plane: One or two 10G/1G from each RSP this is a Special external EOBC 1G/10G ports on RSP.

Note: Split-brain: keepalive over any L2 cloud Management port or any regular data port or interface or sub-interface.

Multi-Pathing with A-VPLS (6500 and ASR9000)

A-VPLS Pseudowire (FAT-PW) – Single Virtual Ethernet Interface across Multiple Interfaces

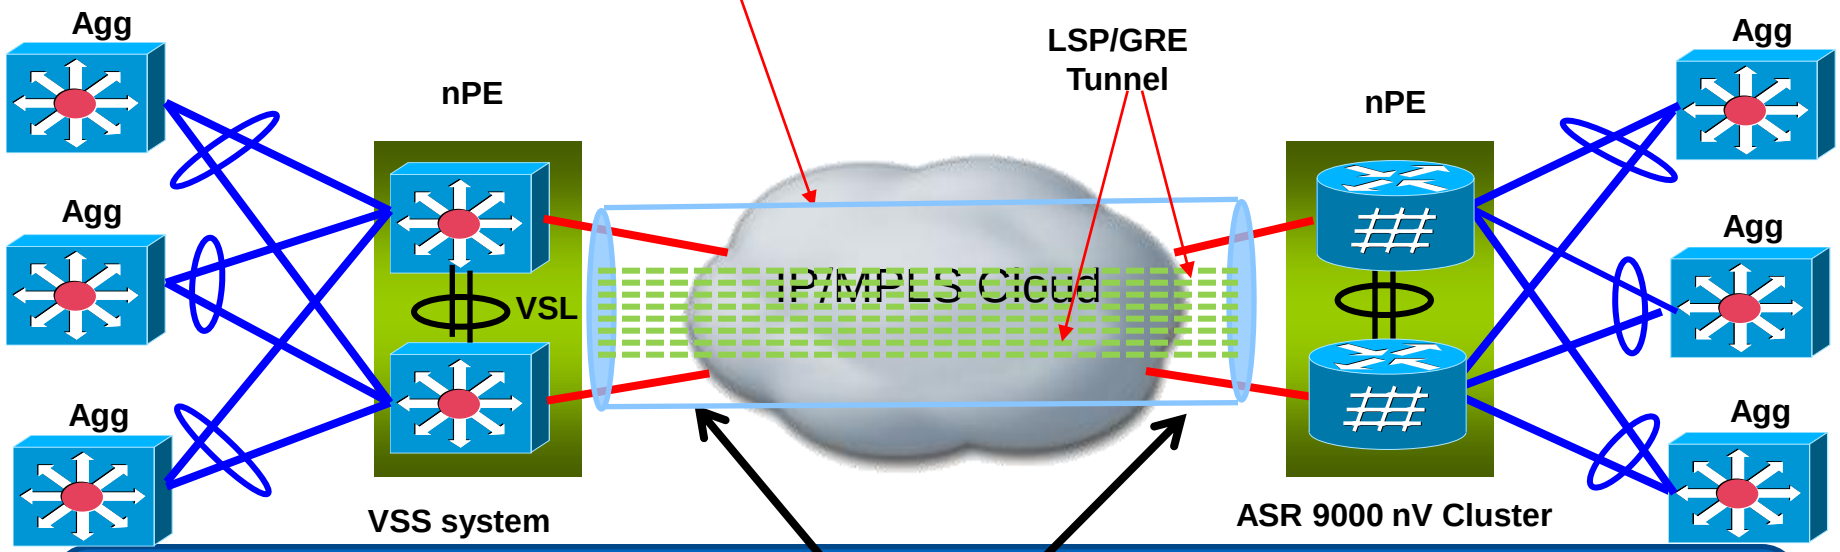

Up to 8 equal cost paths between any two sites
A label is assigned to each equal cost path based on routing reachability of neighbour
Simplified CLI: Virtual Ethernet interface
Loadbalancing at L2/L3/L4

E-VPN – The Principle

From PE1

iBGP L3-NLRI:

- **next-hop: PE1**
- **<C-IP1, L1>**

iBGP L2-NLRI

- **next-hop: PE1**
- **<C-MAC1, L2>**

From PE3

iBGP L3-NLRI:

- **next-hop: PE3**
- **<C-IP5, L1>**

iBGP L2-NLRI

- **next-hop: PE3**
- **<C-MAC3, L1>**

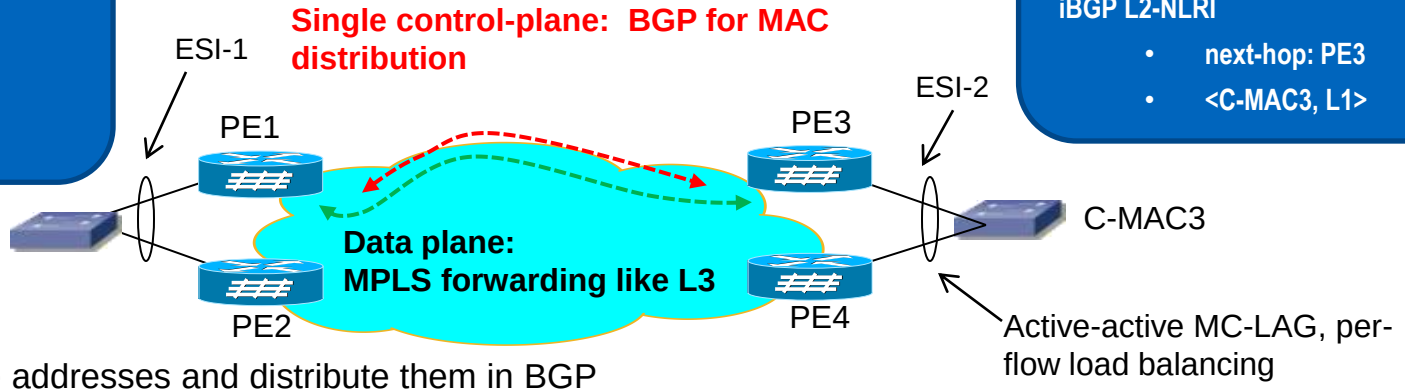

- Treat MAC as routable addresses and distribute them in BGP
- Receiving PE injects these MAC addresses into forwarding table along with its associated adjacency like IP prefix
- When multiple PE nodes advertise the same MAC, then multiple adjacency is created for that MAC address in the forwarding table: multi-path
- When forwarding traffic for a given unicast MAC DA, a hashing algorithm based on L2/L3/L4 header is used to pick one of the adjacencies for forwarding: per-flow load balancing
- PW is not required

Note: Network Layer Reachability Information (NLRI)

E-VPN and Multihoming
Split-horizon

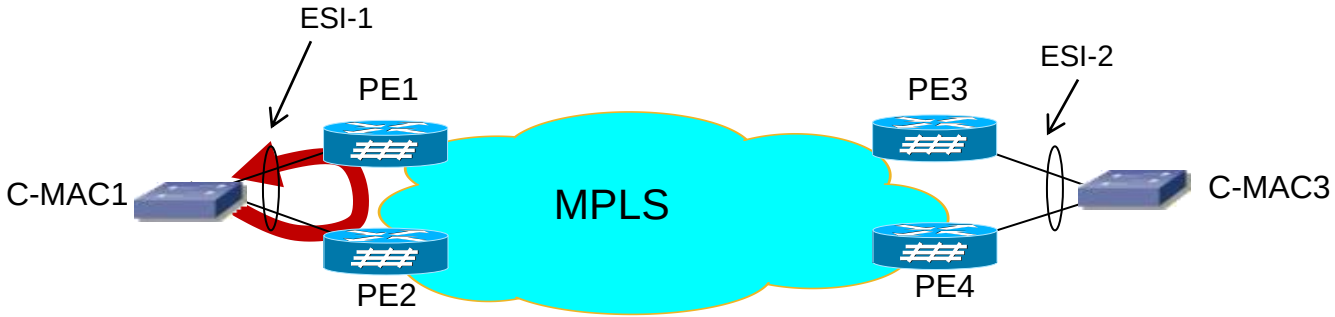

PE1 advertises in BGP a split-horizon label associated with the ESI-1, the Split-horizon label is only used for multi-destination frames (unknown unicast, mcast, bcast). PE2 uses this label to perform split-horizon filtering for frames destined to ESI-1.

For Example, a frame originated by a segment must not be received by the same segment

E-VPN and Multihoming
Designated Forwarder (DF)

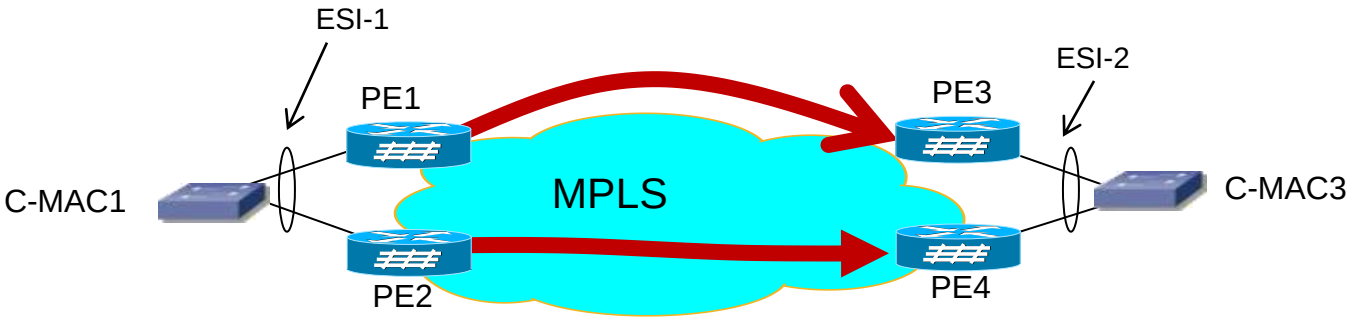

Ensure only one of the remote PE(s) (PE3 and PE4) forward the broadcast frame for the dual-homed devices. There is an DF election between PE3 and PE4 for a given VLAN.

For example, the Non-DF do not transmit any multi-destination frames to that segment (ESI-2)

E-VPN and Multihoming
Aliasing

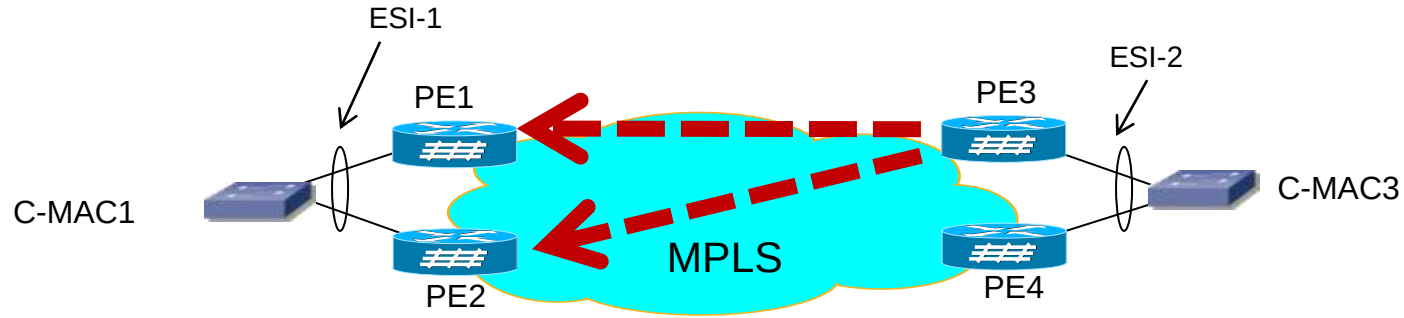

Want to have remote PE(s) to be able to perform load-balancing among the flows destined to C-MAC1 to both PE1 and PE2. In order for PE3 to be able to perform load balancing among the flows destined to C-MAC1 on ESI-1, it needs to know that:

a) ESI-1 sits behind both PE1 and PE2

b) C-MAC-1 is associated with ESI-1

All the remote PEs (PE3 and PE4) use these two routes in combination to associate

a) C-MAC1 to ESI-1

b) Resulting in C-MAC1 behind PE1 and PE2

End-to-End L2 View

Broadcast, Multicast, Unknown Unicast

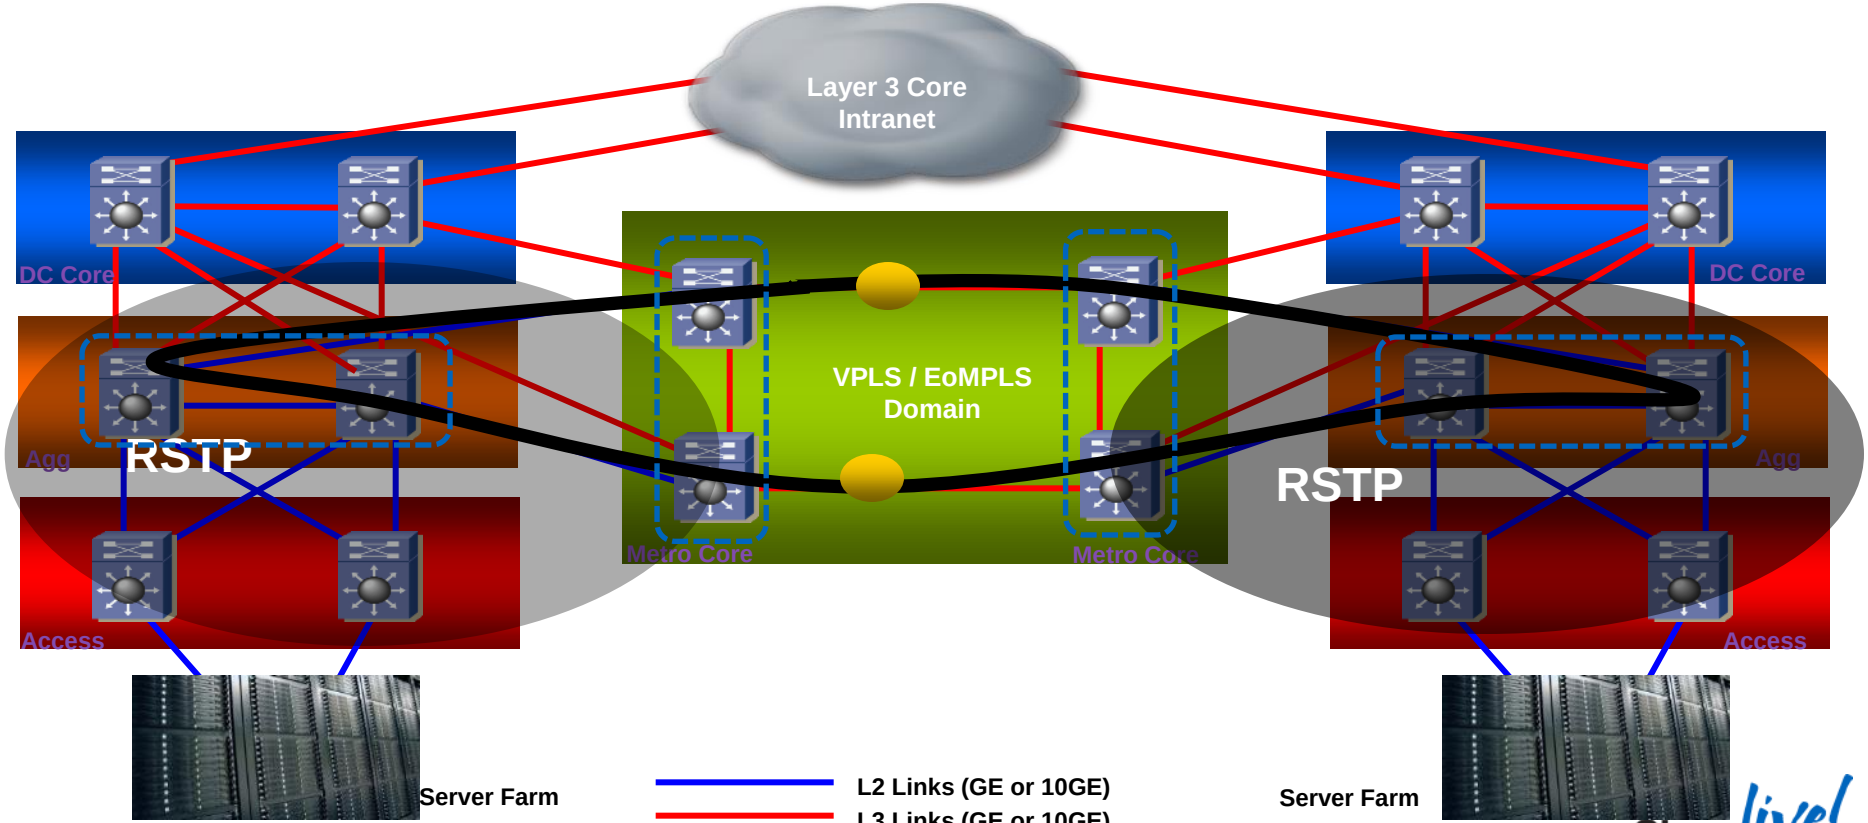

Spanning Tree – Local STP Root Bridges per DC

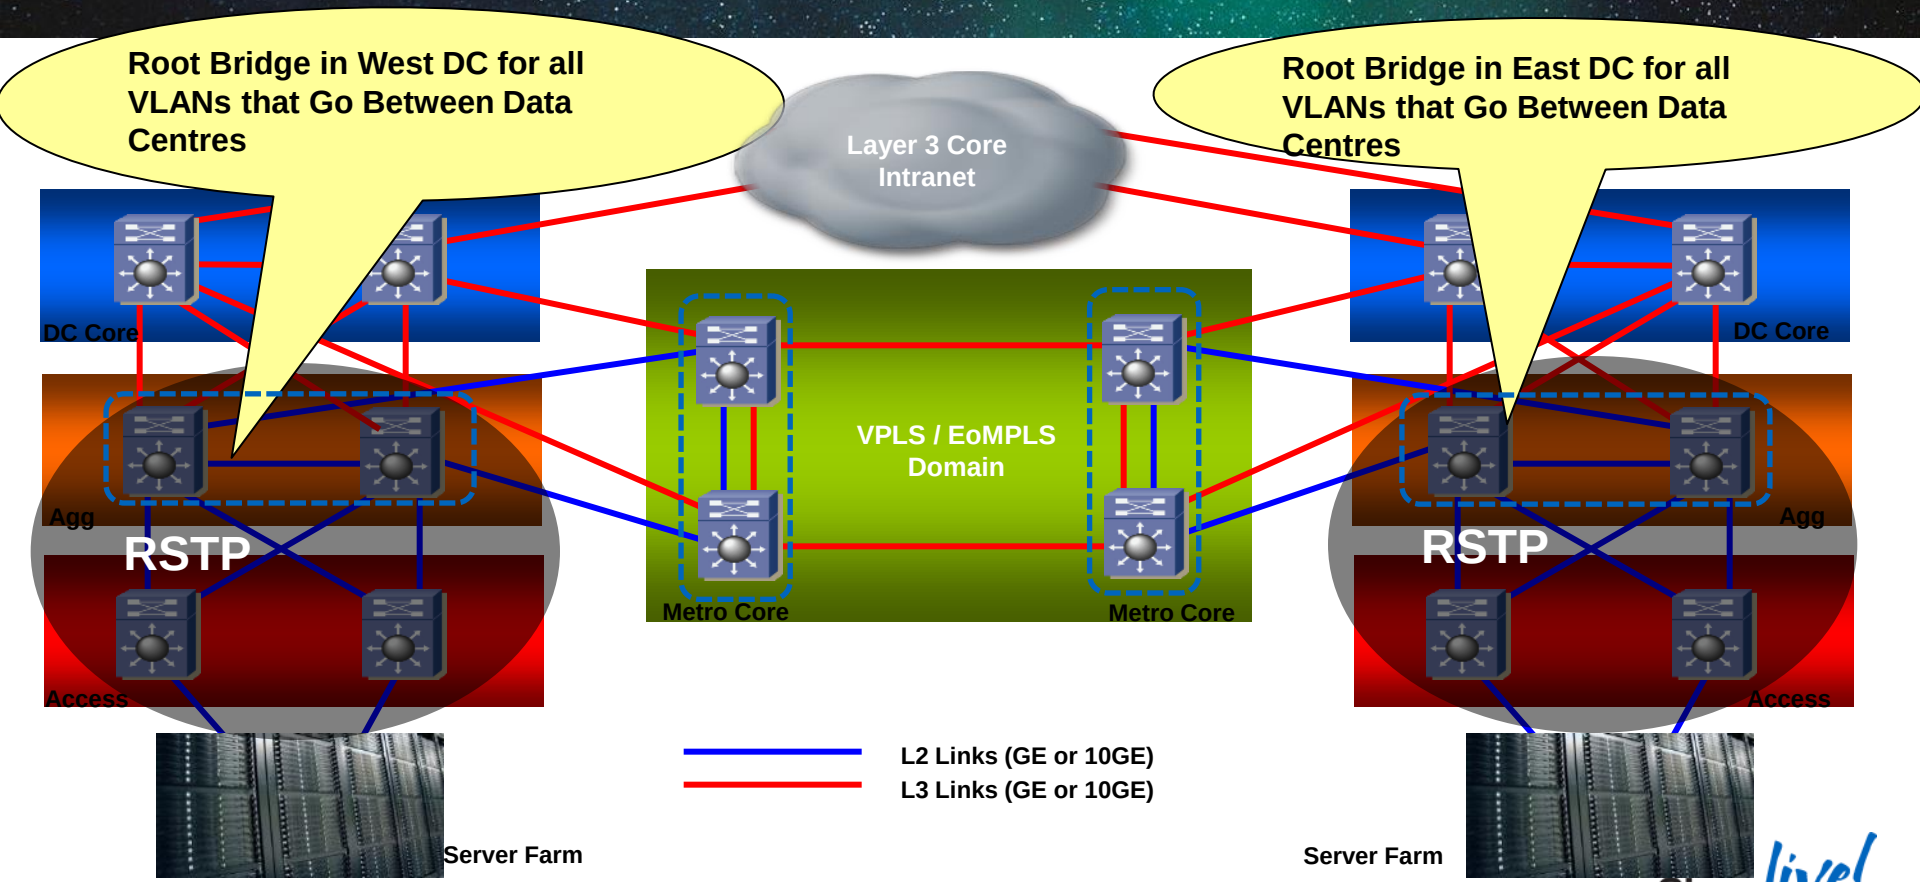

Storm Control

- Traffic storms when packets flood the LAN
- Traffic storm control feature prevents LAN ports from being disrupted by broadcast or multicast flooding
- Rate limiting for unknown unicast (UU) must be handled at Data Centre aggregation; unknown unicast flood rate-limiting (UUFRL):
  - mls rate-limit layer2 unknown rate-in-pps [burst-size]
- Storm Control is configured as a percentage of the link that storm traffic is allowed to use.
  - storm-control broadcast level 1.00 (% of b/w may vary – need to baseline)
  - storm-control multicast level 1.00 (% of b/w may vary – need to baseline)

Summary of MPLS Section

- EoMPLS well suited for Router-Router links
- VPLS well suited for Switch-Switch links
- Straightforward to scale to multiple Data Centre locations
- Leverage Clustering to simplify configuration, multi-homing and achieve active/active deployments.
- A-VPLS
  - Backwards Compatible
  - Load Balancing Enhancements
  - Simplified Configuration
  - Single virtual nPE

Data Centre Interconnect
Agenda

- Mobility and Virtualisation in the Data Centre
- LAN Extension Deployment Scenarios
  - Ethernet Based Solutions
  - MPLS Based Solutions
    - EoMPLS
    - VPLS
    - A-VPLS
    - EVPN
- Overlay Transport Virtualisation (OTV)
- Encryption
- Path Optimisation
- IP Mobility without LAN Extension
- Fabric Solutions
- Summary and Conclusions
- Q&A

Overlay Transport Virtualisation (OTV)

Overlay Transport Virtualisation (OTV)

- OTV is a MAC-in-IP method that extends Layer 2 connectivity
- Ethernet LAN Extension over any Network
- Ethernet in IP “MAC routing”
- Multi-dataCentre scalability
- Simplified Configuration & Operation
- Seamless overlay - no network re-design

- Single touch site configuration
- High Resiliency
- Failure domain isolation
- Seamless Multi-homing
- Maximises available bandwidth
- Automated multi-pathing
- Optimal multicast replication

Introduction
Terminology: Edge Device

- Performs OTV functions
- Support multiple OTV devices per site
- OTV requires the Transport Services (TRS) license
- Creating non default VDC's requires Advanced Services license

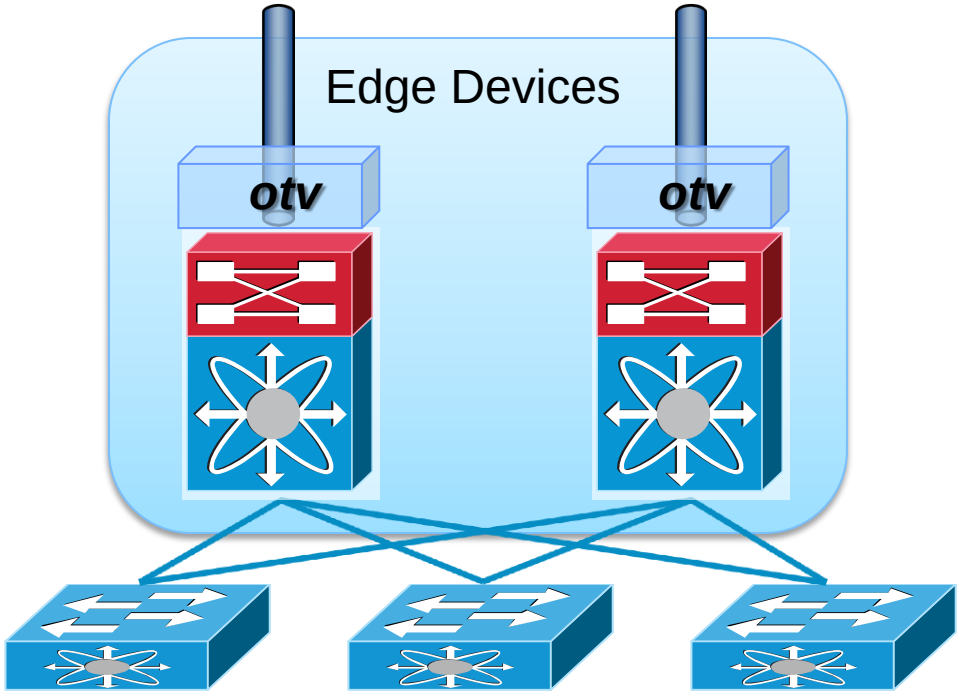

Introduction
Terminology: Internal Interfaces

- Regular layer 2 interfaces facing the site
- No OTV con
- Currently su series modu
- With 6.2.2a as OTV inter
- F1 or F2e m
- F3 Support a

F3 Support Matrix

Join-Interface

Internal Interface	M1	M2	F3
M1	Yes	Yes	No
M2	Yes	Yes	Yes
F1	Yes	Yes	No
F2e	Yes	Yes	No
F3	No	Yes	Yes

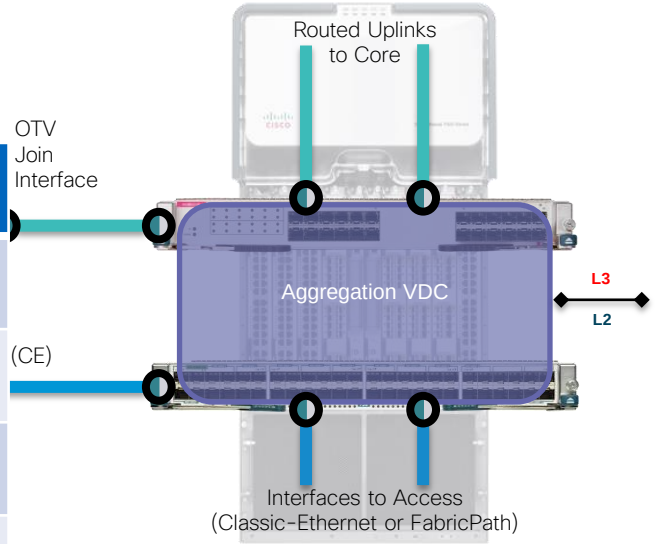

es/F3 interface *

F/M-Series interface

*F3 Support in 6.2(6)

Introduction
Terminology: Join Interface

- Uplink on Edge device that joins the Overlay
- Forwards OTV control and data traffic
- Layer 3 interface
- Currently supported only on M-series modules and F3 with 6.2.6

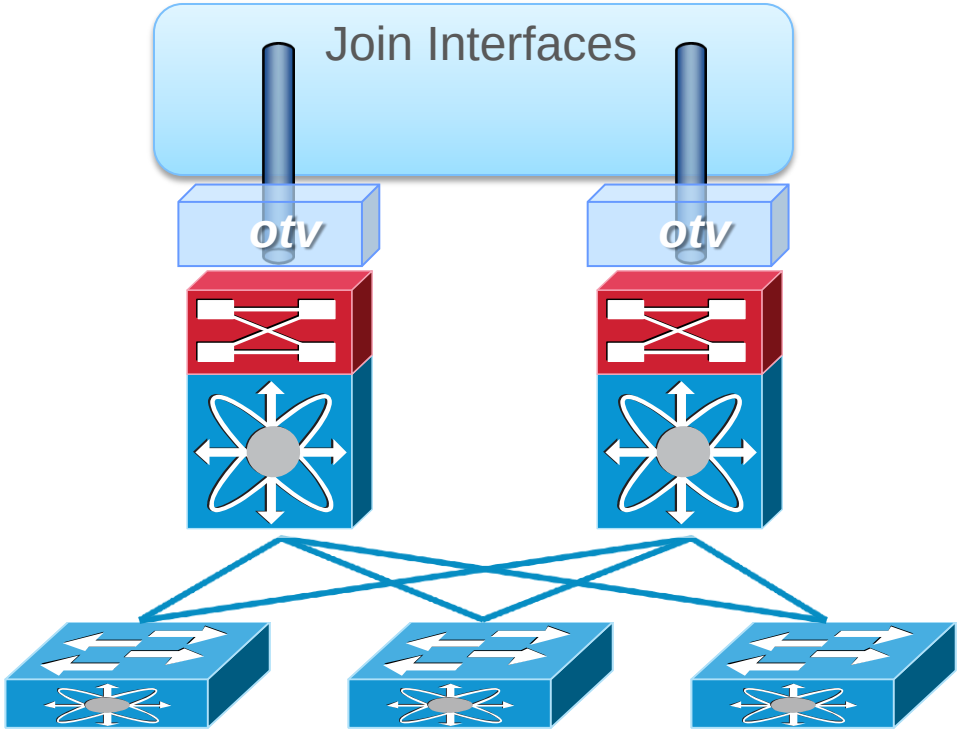

Introduction
Terminology: Overlay Interface

- Virtual Interface where the OTV configurations are applied
- Multi-access multicast-capable interface
- Encapsulates Layer 2 frames

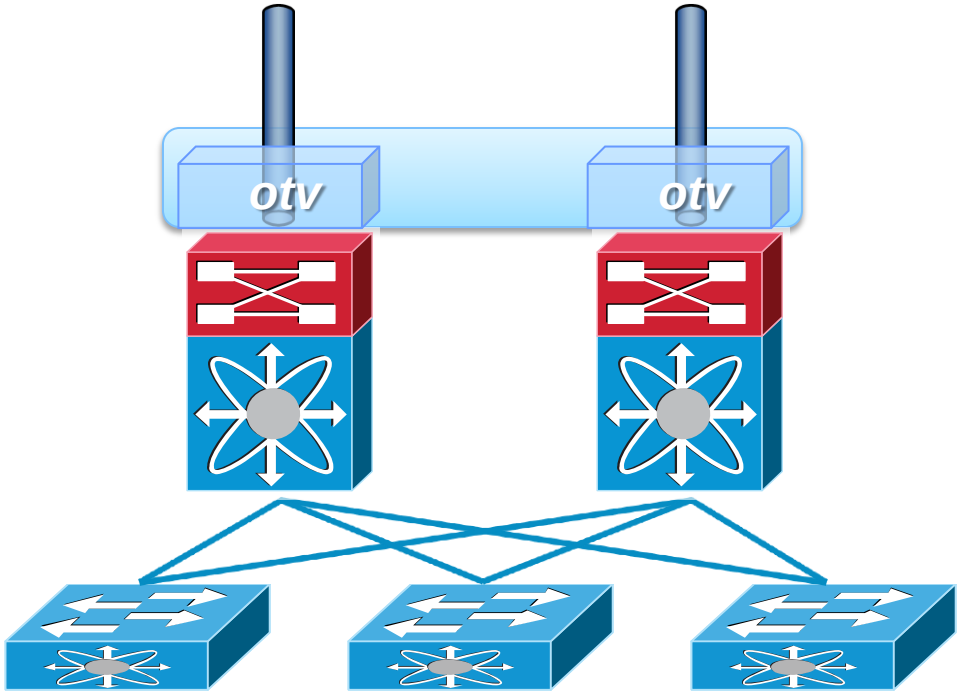

Introduction
Terminology: Authoritative Edge Device

- OTV supports multiple edge devices per site
- A single OTV device is elected as AED on a per-vlan basis
- The AED is responsible for advertising MAC reachability and forwarding traffic into and out of the site for its VLANs

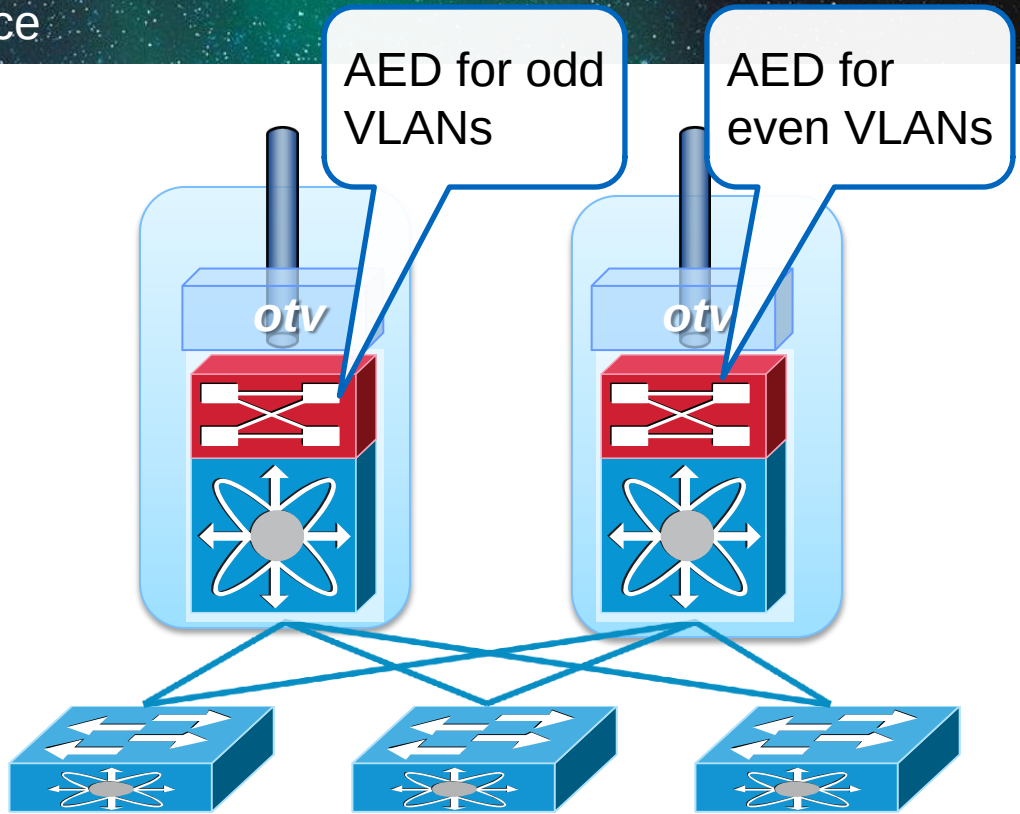

Introduction
Terminology: Site VLAN and Site Identifier

- 5.2(1) added **Dual Site Adjacency**
  1. **Site Adjacency** established across the site vlan
  2. **Overlay Adjacency** established via the Join interface across Layer 3 network

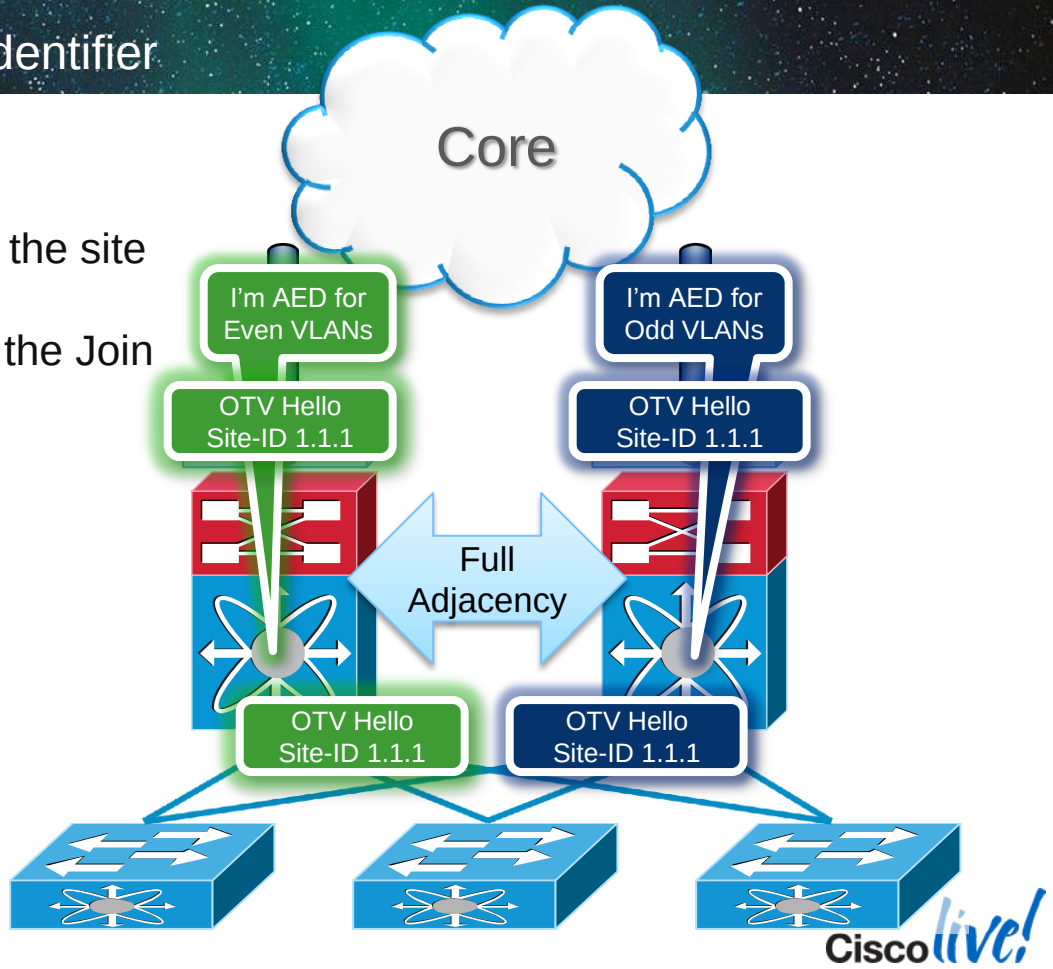

OTV Control Plane
Building the MAC Tables

- **No unknown unicast flooding**
- **Control Plane Learning with proactive MAC advertisement**
- Background process with no specific configuration
- IS-IS used between OTV Edge Devices

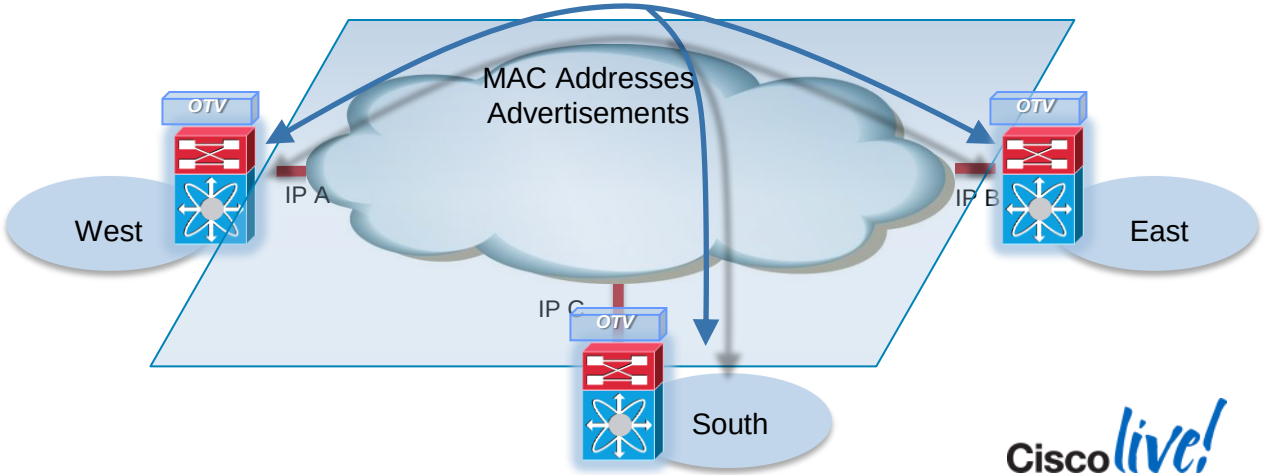

OTV Control Plane

Neighbour Discovery (over Multicast Transport)

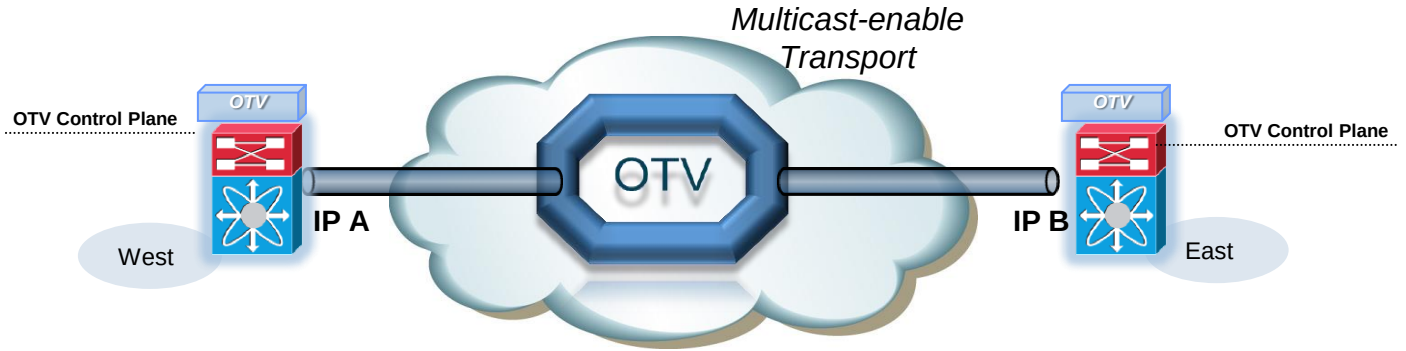

Mechanism

- Edge Devices (EDs) join an multicast group in the transport, as they were hosts (no PIM on EDs)
- OTV hellos and updates are encapsulated in the multicast group

End Result

- Adjacencies are maintained over the multicast group
- A single update reaches all neighbours

OTV Control Plane
Neighbour Discovery (Unicast-only Transport)

Release 5.2 and above

- Ideal for connecting a small number of sites
- With a higher number of sites a multicast transport is the best choice

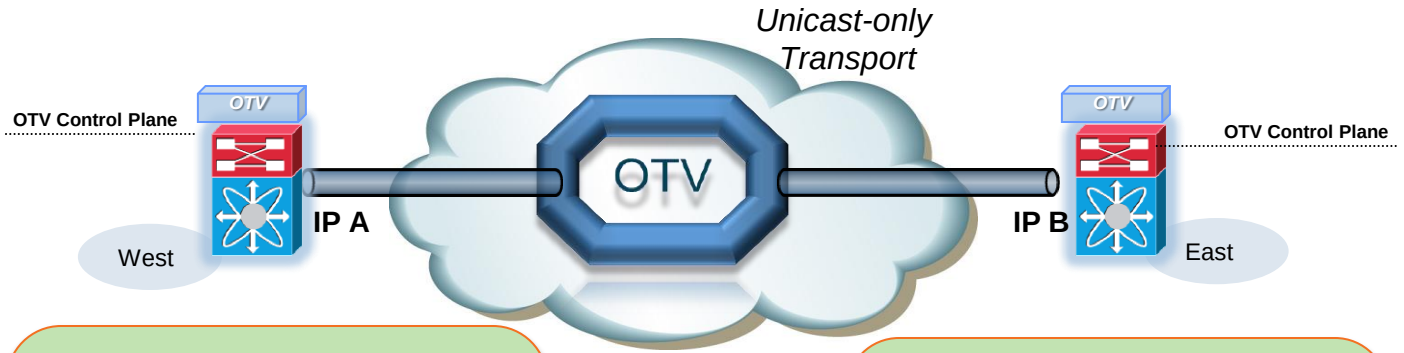

Mechanism

- Edge Devices (EDs) register with an *"Adjacency Server"* ED
- EDs receive a full list of Neighbours (oNL) from the *AS*
- OTV hellos and updates are encapsulated in IP and ***unicast*** to each neighbour

End Result

Neighbour Discovery is automated by the "*Adjacency Server*"

All signalling must be replicated for each neighbour

Data traffic must also be replicated at the head-end

OTV Data Plane
Encapsulation

- **42 Bytes** overhead to the packet IP MTU size
  - Outer IP + OTV Shim - Original L2 Header (w/out the .1Q header)
- 802.1Q header is **removed** and the VLAN field copied over to the OTV shim header
- Outer OTV shim header contains VLAN, overlay number, etc.
- Consider Jumbo MTU Sizing

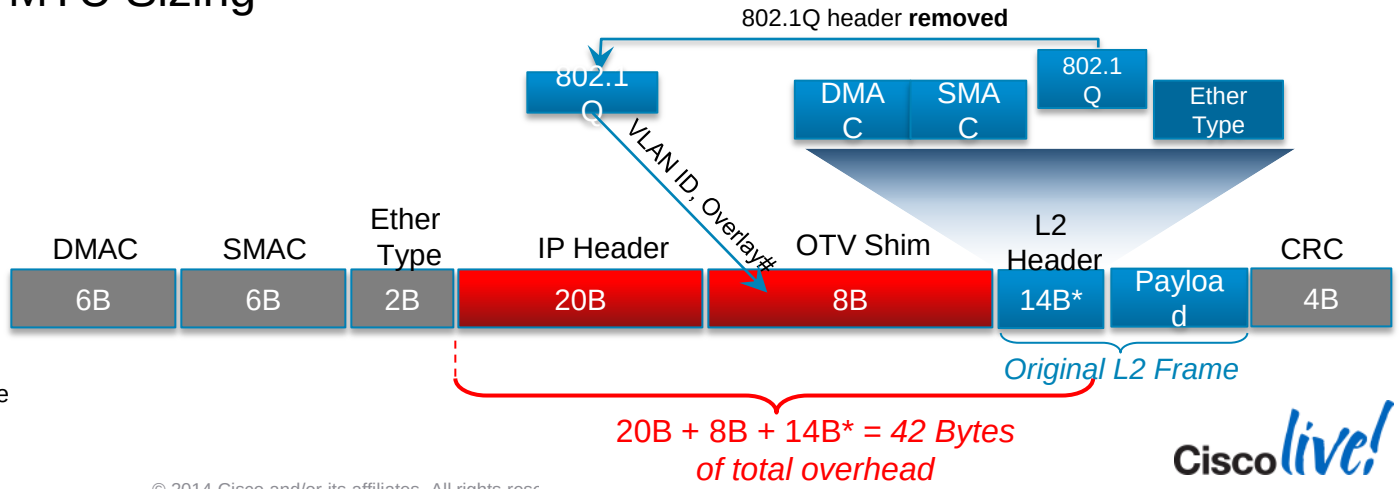

* The 4 Bytes of .1Q header have already been removed

OTV Data Plane
Inter-Site Packet Flow

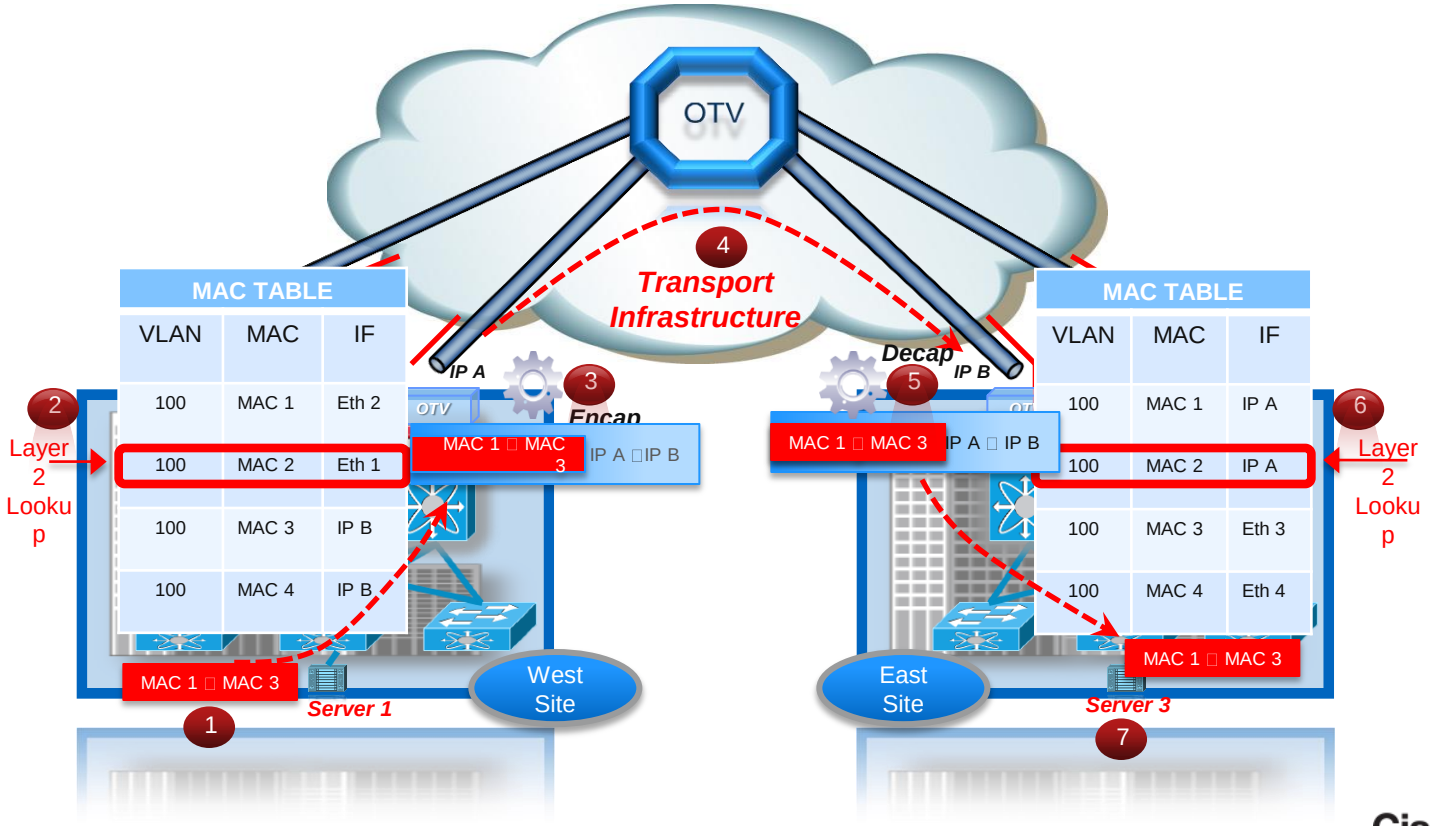

STP BPDU Handling

- When STP is configured at a site, an Edge Device will send and receive BPDUs on the **internal interfaces**.
- An OTV Edge Device will not originate or forward BPDUs on the overlay network.
- An OTV Edge Device can become (but it is not required to) a root of one or more spanning trees within the site.
- An OTV Edge Device will take the typical action when receiving Topology Change Notification (TCNs) messages.

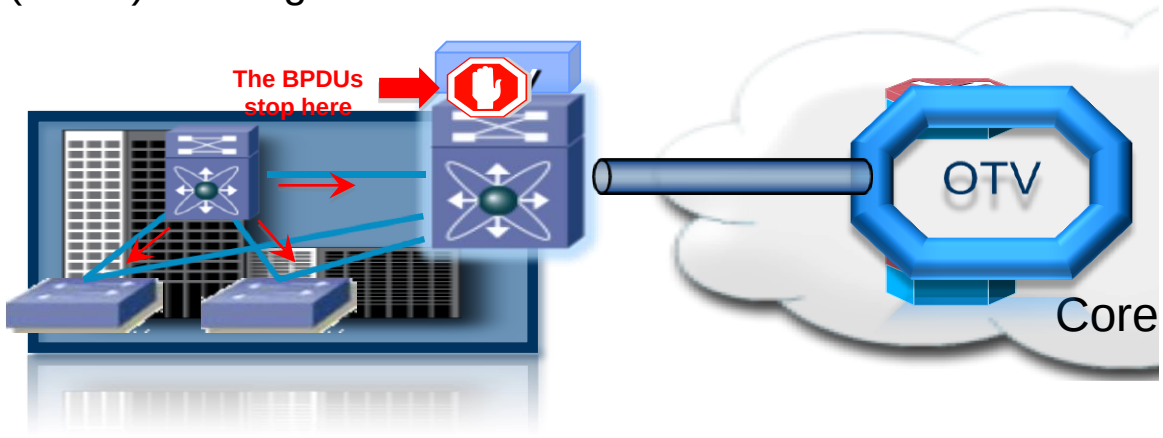

Handling Data-plane Loop Prevention Handling

- Brodcast/M-cast packets reach all Edge Devices within a site.
- **The AED for the VLAN is the only Edge Device that forwards b-cast/ m-cast packets onto the overlay network**
- The b-cast/m-cast packet is replicated to all the Edge Devices on the overlay.
- Only the AED at each remote site will forward the packet from the overlay onto the site.
- Once sent into the site, the b-cast/m-cast packet is replicated per regular switching

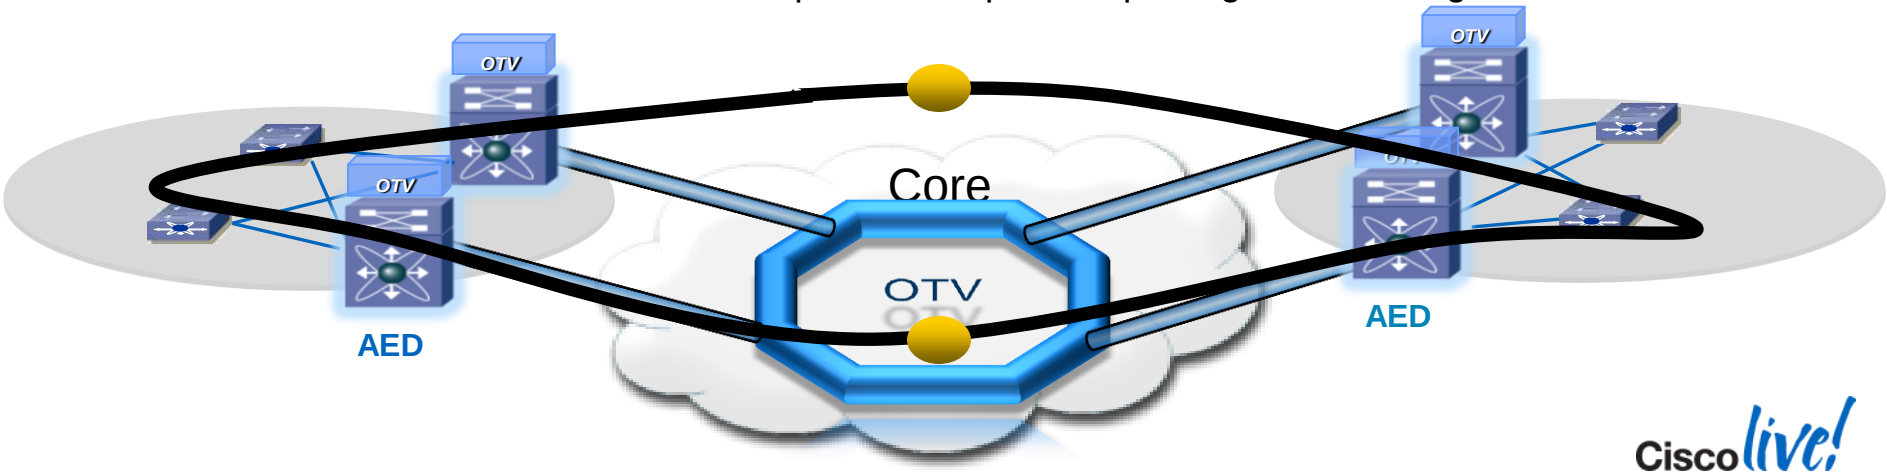

Multi-homing
AED and Broadcast/Multicast Handling

- Broadcast/M-cast packets reach all Edge Devices within a site.
- **The AED for the VLAN is the only Edge Device that forwards b-cast/ m-cast packets onto the overlay network**
- The b-cast/m-cast packet is replicated to all the Edge Devices on the overlay.
- Only the AED at each remote site will forward the packet from the overlay onto the site.
- Once sent into the site, the b-cast/m-cast packet is replicated per regular switching

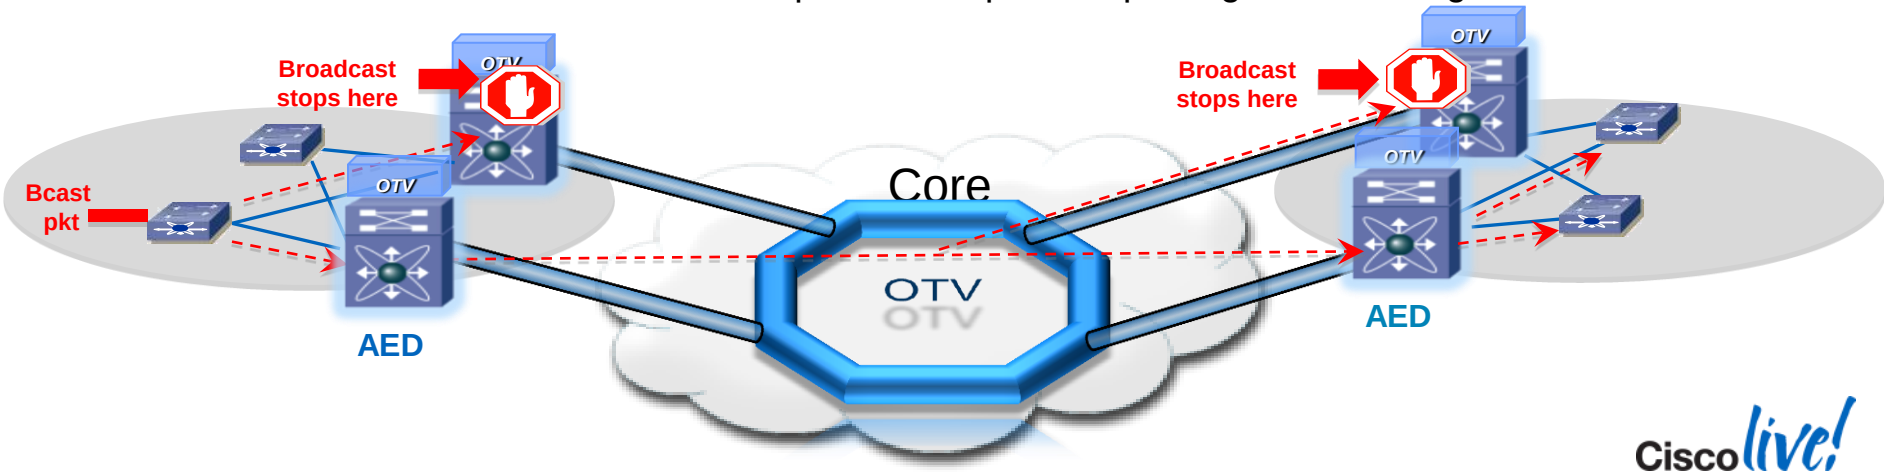

Multi-homing
AED and Unicast Forwarding

- One AED is elected for each VLAN on each site
- Different AEDs can be elected for each VLAN to balance traffic load
- Only the AED forwards unicast traffic to and from the overlay
- Only the AED advertises MAC addresses for any given site/VLAN
- Unicast routes will point to the AED on the corresponding remote site/VLAN

MAC TABLE		
VLAN	MAC	IF
100	MAC 1	IP A
201	MAC 2	IP B

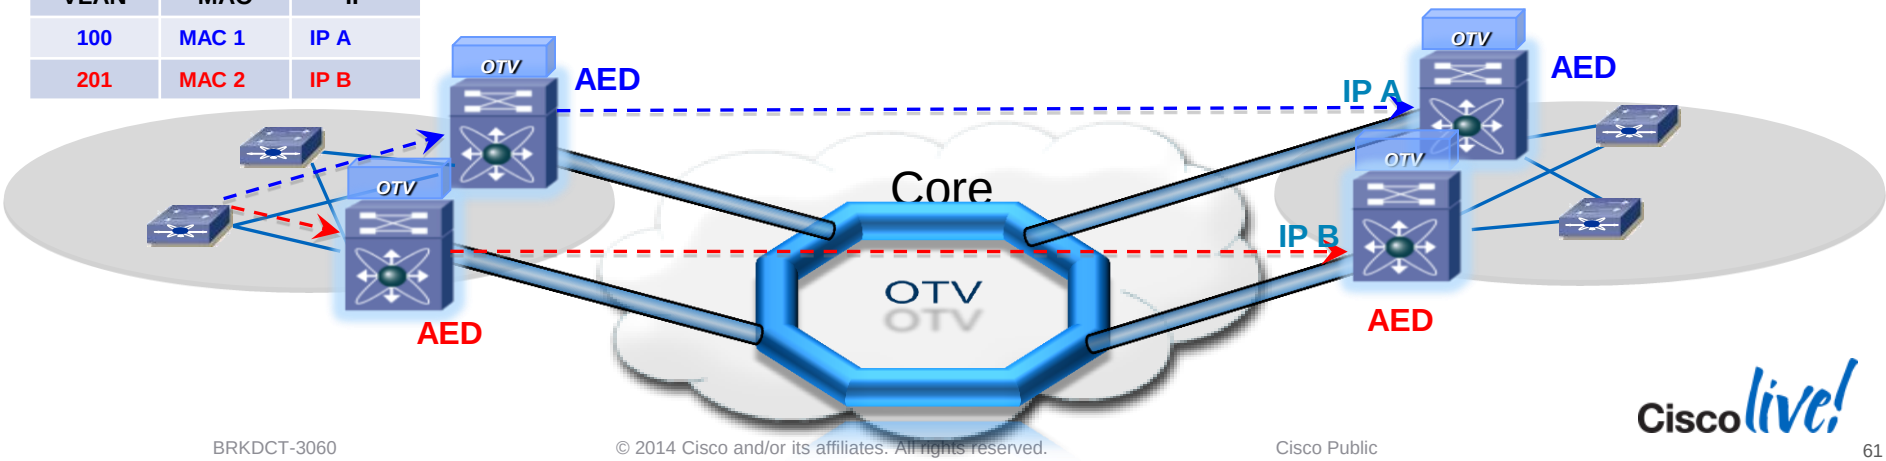

Dedicated Broadcast Group

- Dedicated broadcast group is now configurable (6.2.2+)
  - “otv broadcast-group” configuration line under overlay
  - Optional command

```
interface Overlay3
  otv join-interface loopback11
  otv control-group 224.3.3.0
  otv data-group 232.3.0.0/24
  otv broadcast-group 224.2.2.0
  otv extend-vlan 198-227
  no shutdown
```

```
OTV-a# sh otv

OTV Overlay Information
Site Identifier 0000.0000.0010

Overlay interface Overlay3

 VPN name              : Overlay3
 VPN state             : UP
 Extended vlans        : 198-227 (Total:30)
 Control group         : 224.3.3.0
 Data group range(s)   : 232.3.0.0/24
 Broadcast group       : 224.2.2.0
 Join interface(s)     : Lo11 (172.26.247.125)
 Site vlan             : 99 (up)
 AED-Capable           : Yes
 Capability            : Multicast-Reachable
```

Useful for QoS purposes:
eg. ip multicast rate-limit

OTV VLAN Translation

VLAN Translation: Direct Translation

- Translating local VLAN to remote VLAN
- VLAN in Site West to correspond with a different VLAN in Site East

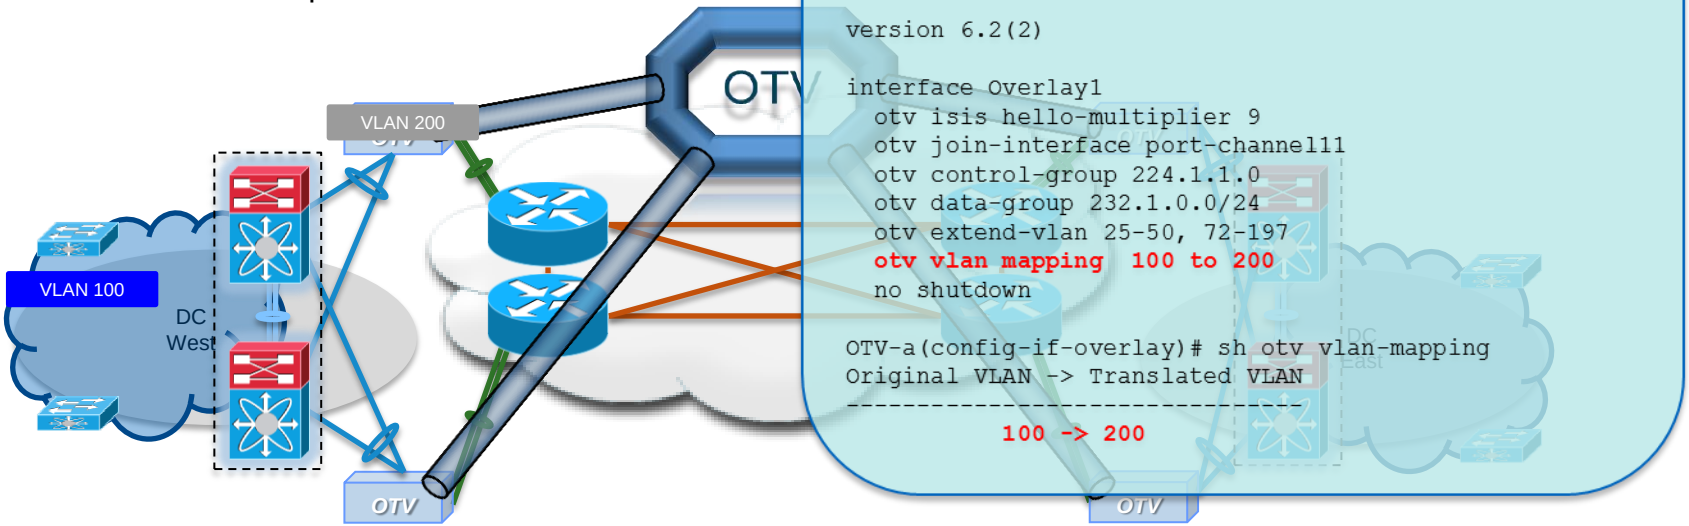


```
OTV-a(config)# int overlay1
OTV-a(config-if-overlay)# otv vlan mapping 100 to 200
OTV-a(config-if-overlay)# sh run int overlay1

!Command: show running-config interface Overlay1
!Time: Fri Mar 29 18:58:12 2013

version 6.2(2)

interface Overlay1
  otv isis hello-multiplier 9
  otv join-interface port-channel11
  otv control-group 224.1.1.0
  otv data-group 232.1.0.0/24
  otv extend-vlan 25-50, 72-197
  otv vlan mapping  100 to 200
  no shutdown

OTV-a(config-if-overlay)# sh otv vlan-mapping
Original VLAN -> Translated VLAN
------------------------------
          100 -> 200
```

Tunnel Depolarisation

- Secondary IP command introduced
  - Configured within interface, not OTV interface
- Introduction of multiple IPs results in tunnel depolarisation
- 3 secondary IPs supported

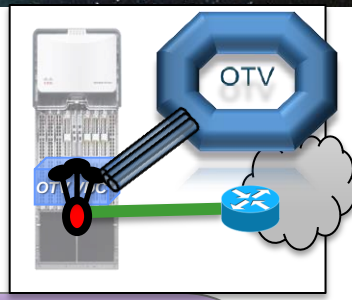


```
OTV-a(config-if)# ip address 2.100.11.1/24 secondary
Disabling IP Redirects on port-channel11 :secondary address
configured.
OTV-a(config-if)# sh run int po11

!Command: show running-config interface port-channel11
!Time: Wed Mar 27 23:05:21 2013

version 6.2(2)

interface port-channel11
  no ip redirects
  ip address 2.100.11.100/24
  ip address 2.100.11.1/24 secondary
  ip ospf network point-to-point
  ip router ospf 1 area 0.0.0.0
  ip igmp version 3
```

```
OTV-a (config-if)# sh otv

OTV Overlay Information
Site Identifier 0000.0000.0011

Overlay interface Overlay1

 VPN name             : Overlay1
 VPN state            : UP
 Extended vlans       : 25-50 72-227 (Total:182)
 Control group        : 224.1.1.0
 Data group range(s)  : 232.1.0.0/24
 Broadcast group      : 224.1.1.0
 Join interface(s)    : Po11 (2.100.11.100)
    Secondary IP Addresses: : 2.100.11.1
 Site vlan            : 1 (up)
 AED-Capable          : Yes1
 Capability           : Multicast-Reachable
```

OTV Use Case
Two Sites Connected

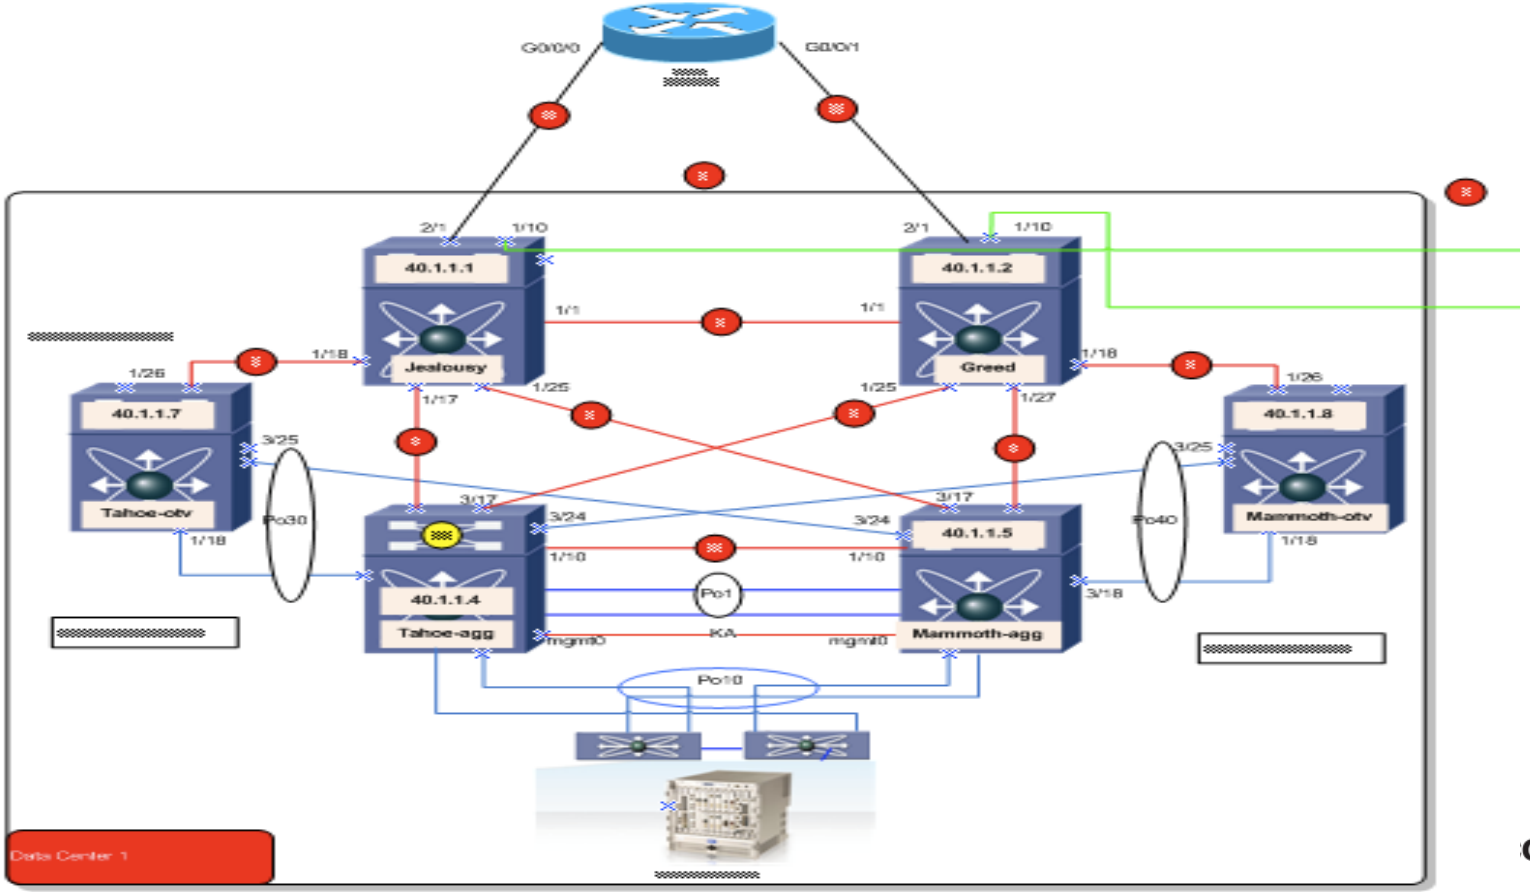

OTV Summary

- STP Isolation: BPDUs are not forwarded over the overlay
- Automated Multi-homing support
- Optimal Multicast Replication
- Control-plane MAC based learning and forwarding
- Simplified Configuration
- Operational Simplicity
- IP Based / Transport Agnostic (IP/MPLS)
- End-to-End loop prevention

Data Centre Interconnect
Agenda

- Mobility and Virtualisation in the Data Centre
- LAN Extension Deployment Scenarios
  - Ethernet Based Solutions
  - MPLS Based Solutions
    - EoMPLS
    - VPLS
    - A-VPLS
    - EVPN
- Overlay Transport Virtualisation (OTV)
- Encryption
- Path Optimisation
- IP Mobility without LAN Extension
- Fabric Solutions
- Summary and Conclusions
- Q&A

= For your Reference




Encryption

Point-to-Point Encryption Solution

Nexus 7000 MACsec can be used to secure data across remote data-Centre if Layer 2 and BPDU transparency is ensured (e.g. dark Fibre or DWDM transport).

Nexus 7000 Supported Hardware:
- F2e (8 ports SFP+)
- F2e 48port 10GBaseT
- M1 Modules
- M2 Modules
- F3 Module interfaces (41 to 48)

Catalyst 6500 Supported Hardware:
- Sup2T 3x1GE interfaces or the 2x10GE interfaces (no Twin-Gig)
- 69xx (10G Only)

Encryption Solution

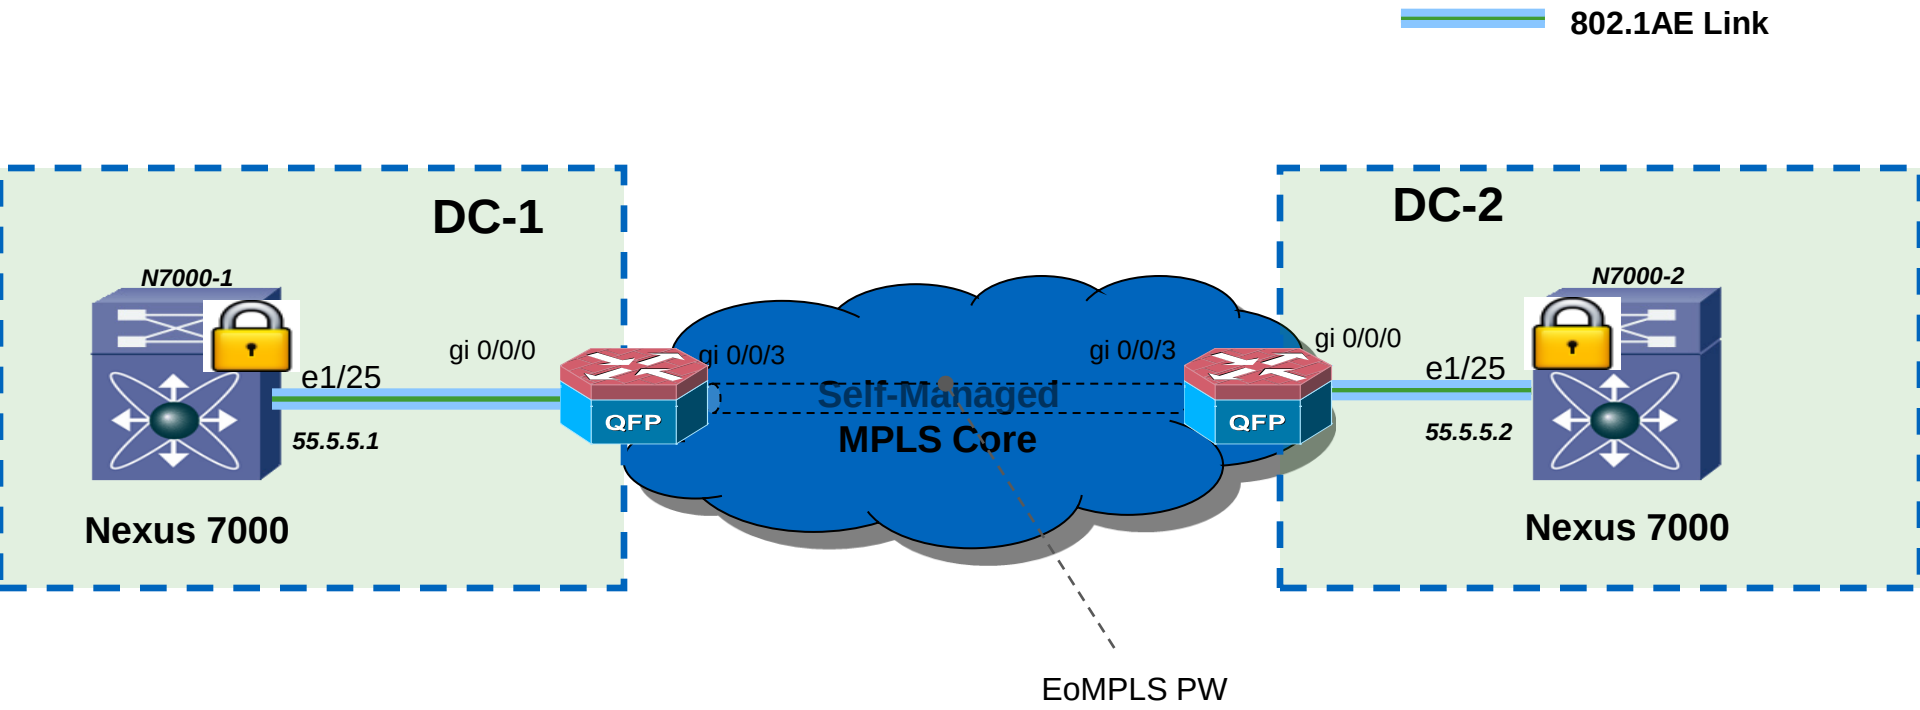

* Remote port shutdown

Nexus 7000 vPC / Catalyst 6500 VSS

Encryption Solution

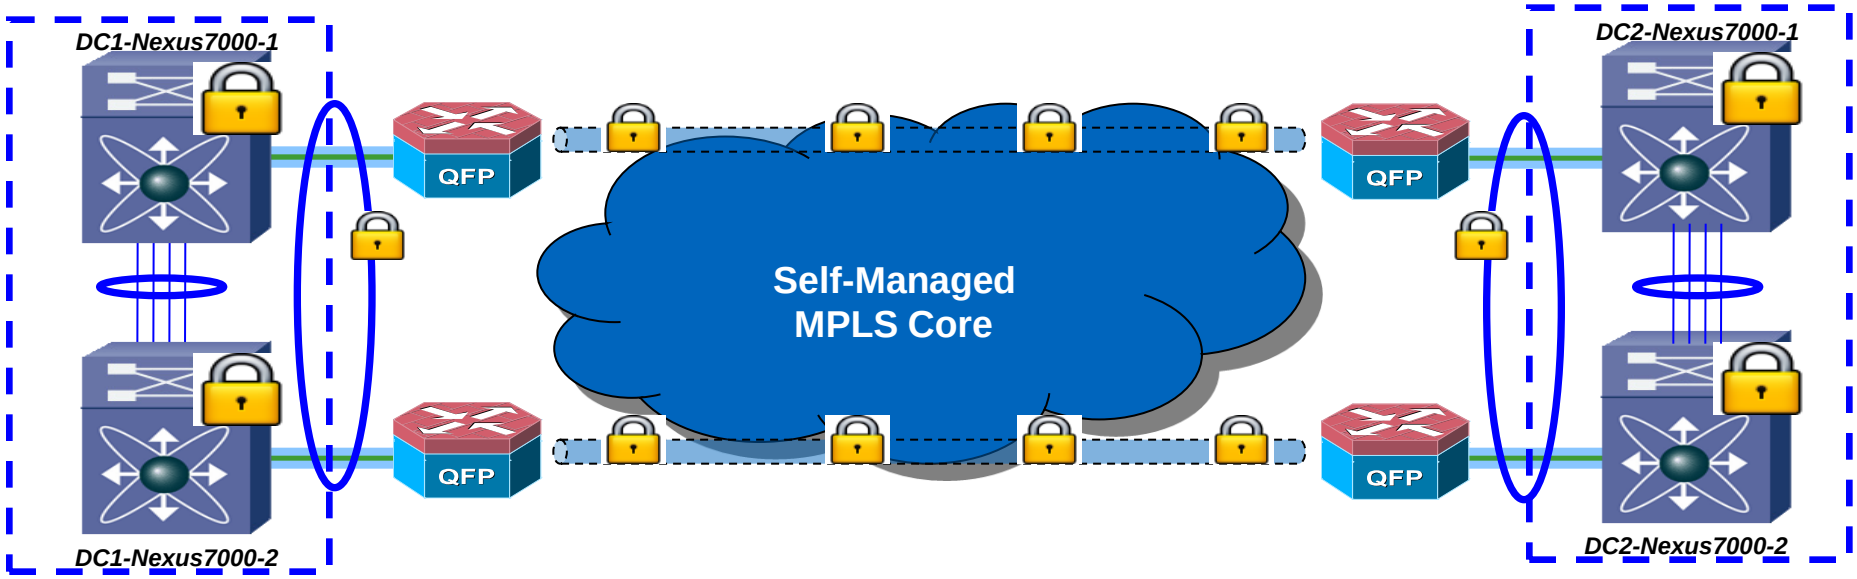

* Remote port shutdown

OTV and 802.1ae Encryption

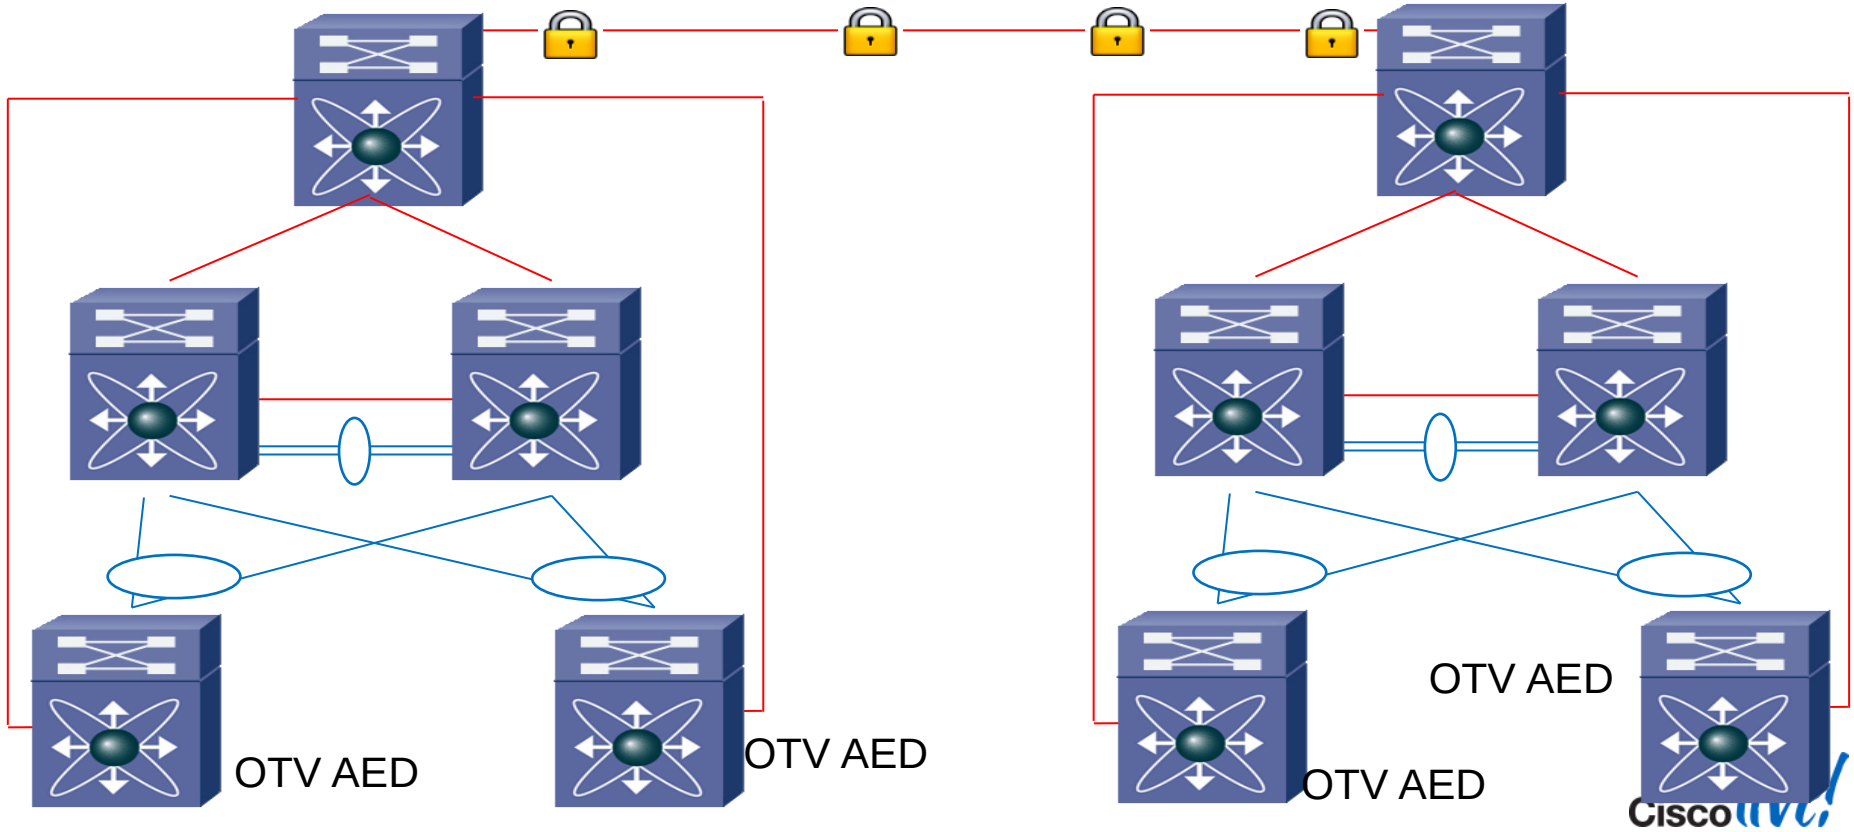

VSPA/ASR1000/ASA Solution Overview

Data Centre Interconnect with MPLSoGREoIPSec

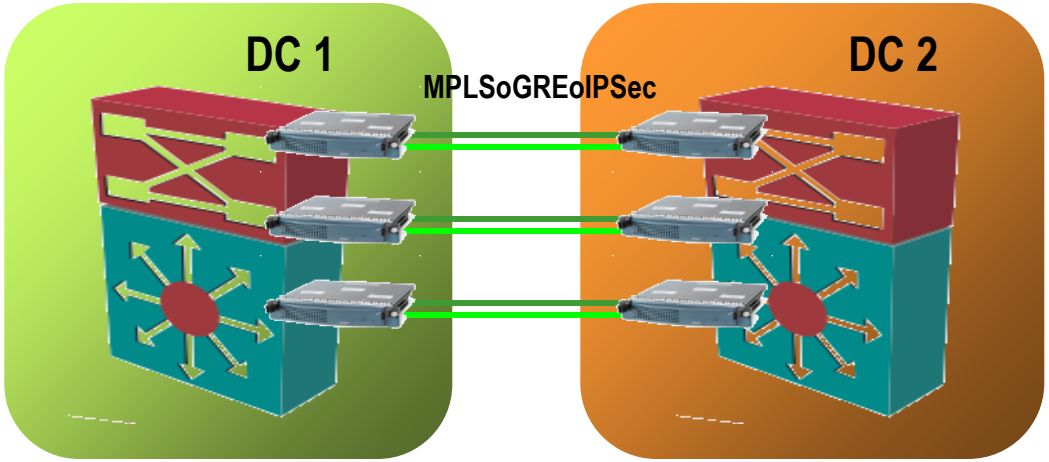

- Leverage ECMP to load balance flows over multiple GRE/IPSec
- Duplicate tunnels per VSPA allow redundant 10GE links to be provisioned
- Inherent crypto engine HA: Traffic will rebalance in the event of a VSPA outage

Solution Objective

- Provide a high speed Layer 2 connection between two or more DCs.. Two or more redundant links are used between the DCs.

VSPA Performance

- Three VSPAs can drive a 10 GE link with IMIX traffic. Single chassis can encrypt three 10 GE links at IMIX rates.

ASR-1000 Performance

- ASR1000-ESP5-1.8Gbps IPSec
- ASR1000-ESP10-4Gbps IPSec
- ASR1000-ESP20-8Gbps IPSec
- ASR1006-2/ESP20-16Gbps IPSec
- ASR1006-2/ESP40 – 25.8Gbps IPSec
- ASA-5585-X Performance
- IPSec 5Gbps

Data Centre Interconnect
Agenda

- Mobility and Virtualisation in the Data Centre
- LAN Extension Deployment Scenarios
  - Ethernet Based Solutions
  - MPLS Based Solutions
    - EoMPLS
    - VPLS
    - A-VPLS
    - EVPN
- Overlay Transport Virtualisation (OTV)
- Encryption
- Path Optimisation
- IP Mobility without LAN Extension
- Fabric Solutions
- Summary and Conclusions
- Q&A

= For your Reference




Path Optimisation

Optimising Traffic Patterns and HA Design

- Many tradeoffs in understanding flows in multi-DC design
- Slides that follow are a specific recommendation that meets the following requirements:
  - Minimise inter-DC traffic to maintenance/failure scenario's
  - Ability to extend clusters between locations (OS, FS, DB, VMware DRS, etc.)
  - Desire to keep flows symmetric in/out of a location for DC services (FW, LB, IPS, WAAS, etc.)
  - Site failure will allow failover, with IP mobility to resolve caching issues
  - Single points of failure in gear won't cause site failover
  - Indicate a location preference for a service to the Layer 3 network
  - If broadcast storm in DC, limit impacts to other DCs
  - If DCI Layer 2 adjacency fails
  - Ability to connect to services in both DC locations (active/active per application)
  - DNS to round-robin clients to DC
  - Allow backup server farms with same service VIP (for backup connections on site fail)
  - Localised HSRP (egress)
  - Inbound traffic draw via LISP (ingress)
- This is a solution in production

Live Moves or Cold Moves

- **Live (hot) Moves** preserve existing connections and state
  - e.g. vMotion, Cluster failover
  - Requires synchronous storage and network policy replication ➔ Distance limitations

- **Cold Moves** bring machines down and back up elsewhere
  - e.g. Site Recovery Manager
  - No state preservation: less constrained by distances or services capabilities

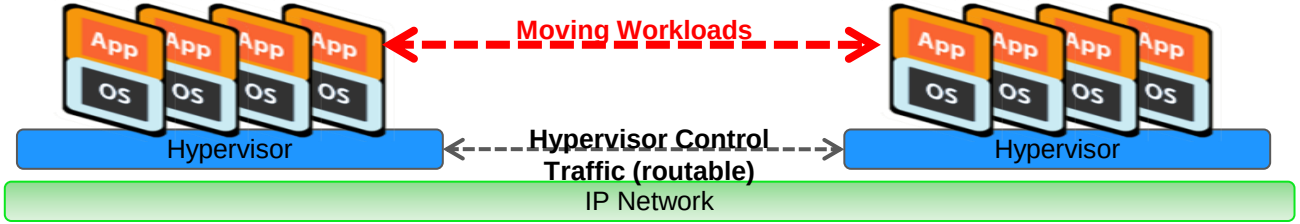

Sample Cluster - Primary Service in Left DC

FHRP localisation – Path Optimisation

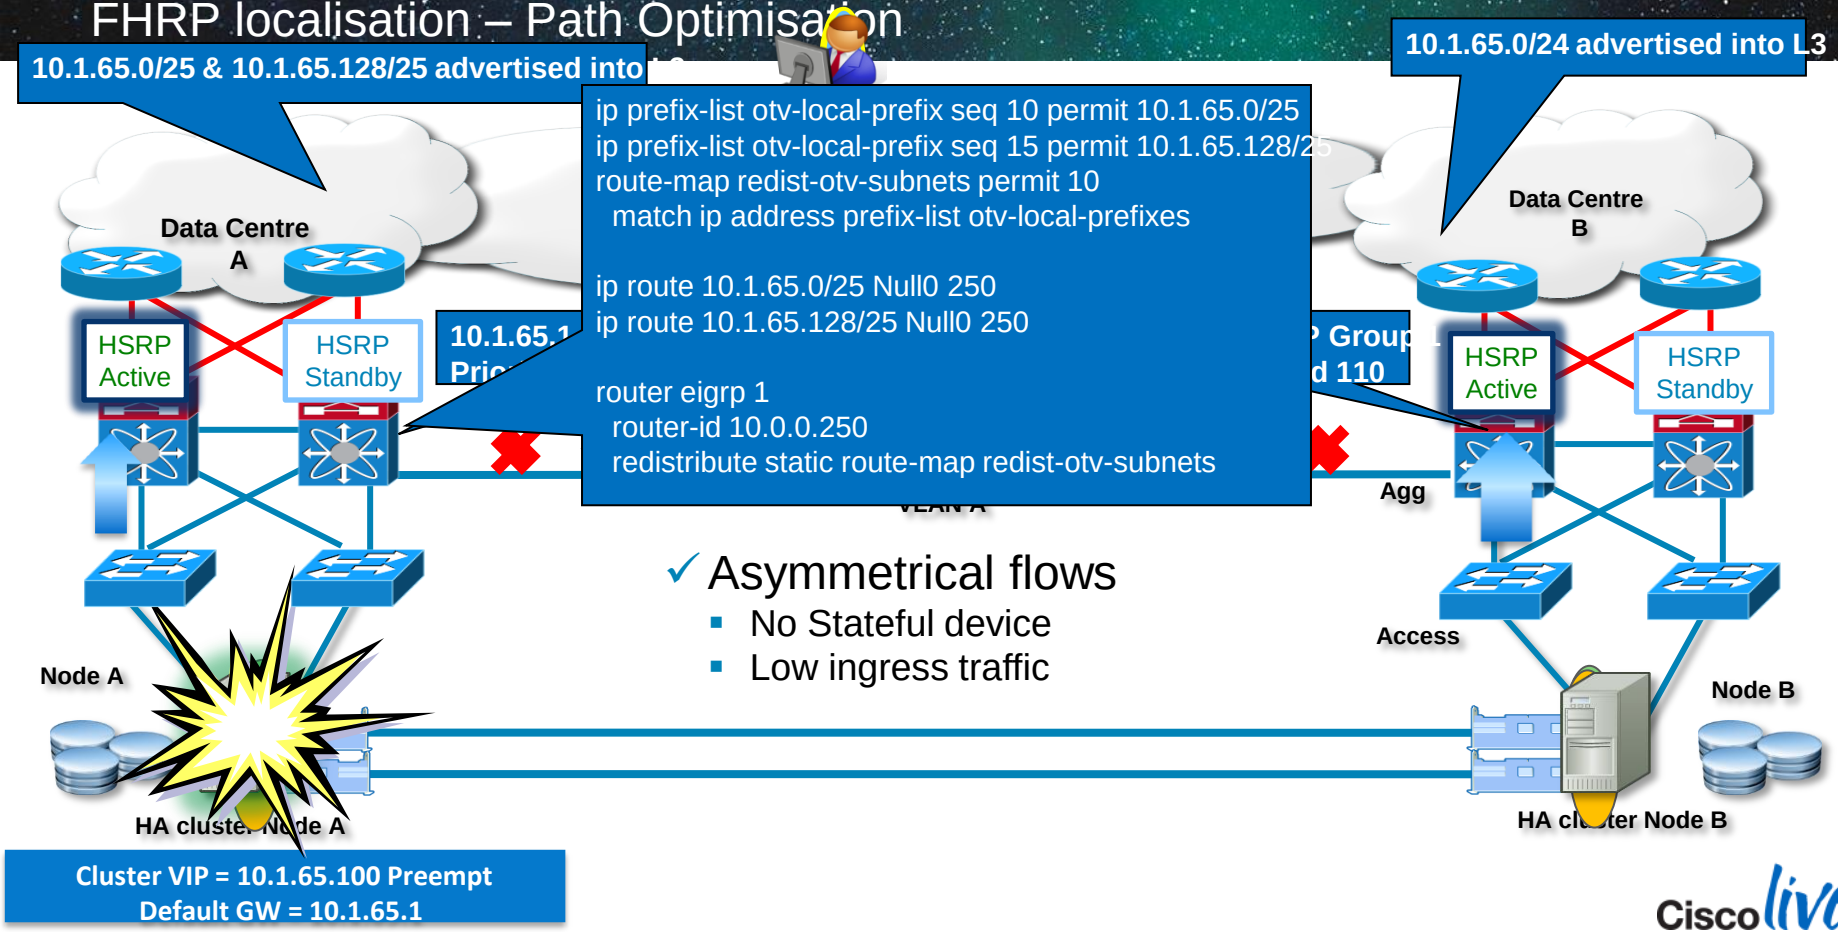

- ✓ Asymmetrical flows
  - No Stateful device
  - Low ingress traffic

Sample Cluster – Active / Active DC

FHRP localisation – Path Optimisation

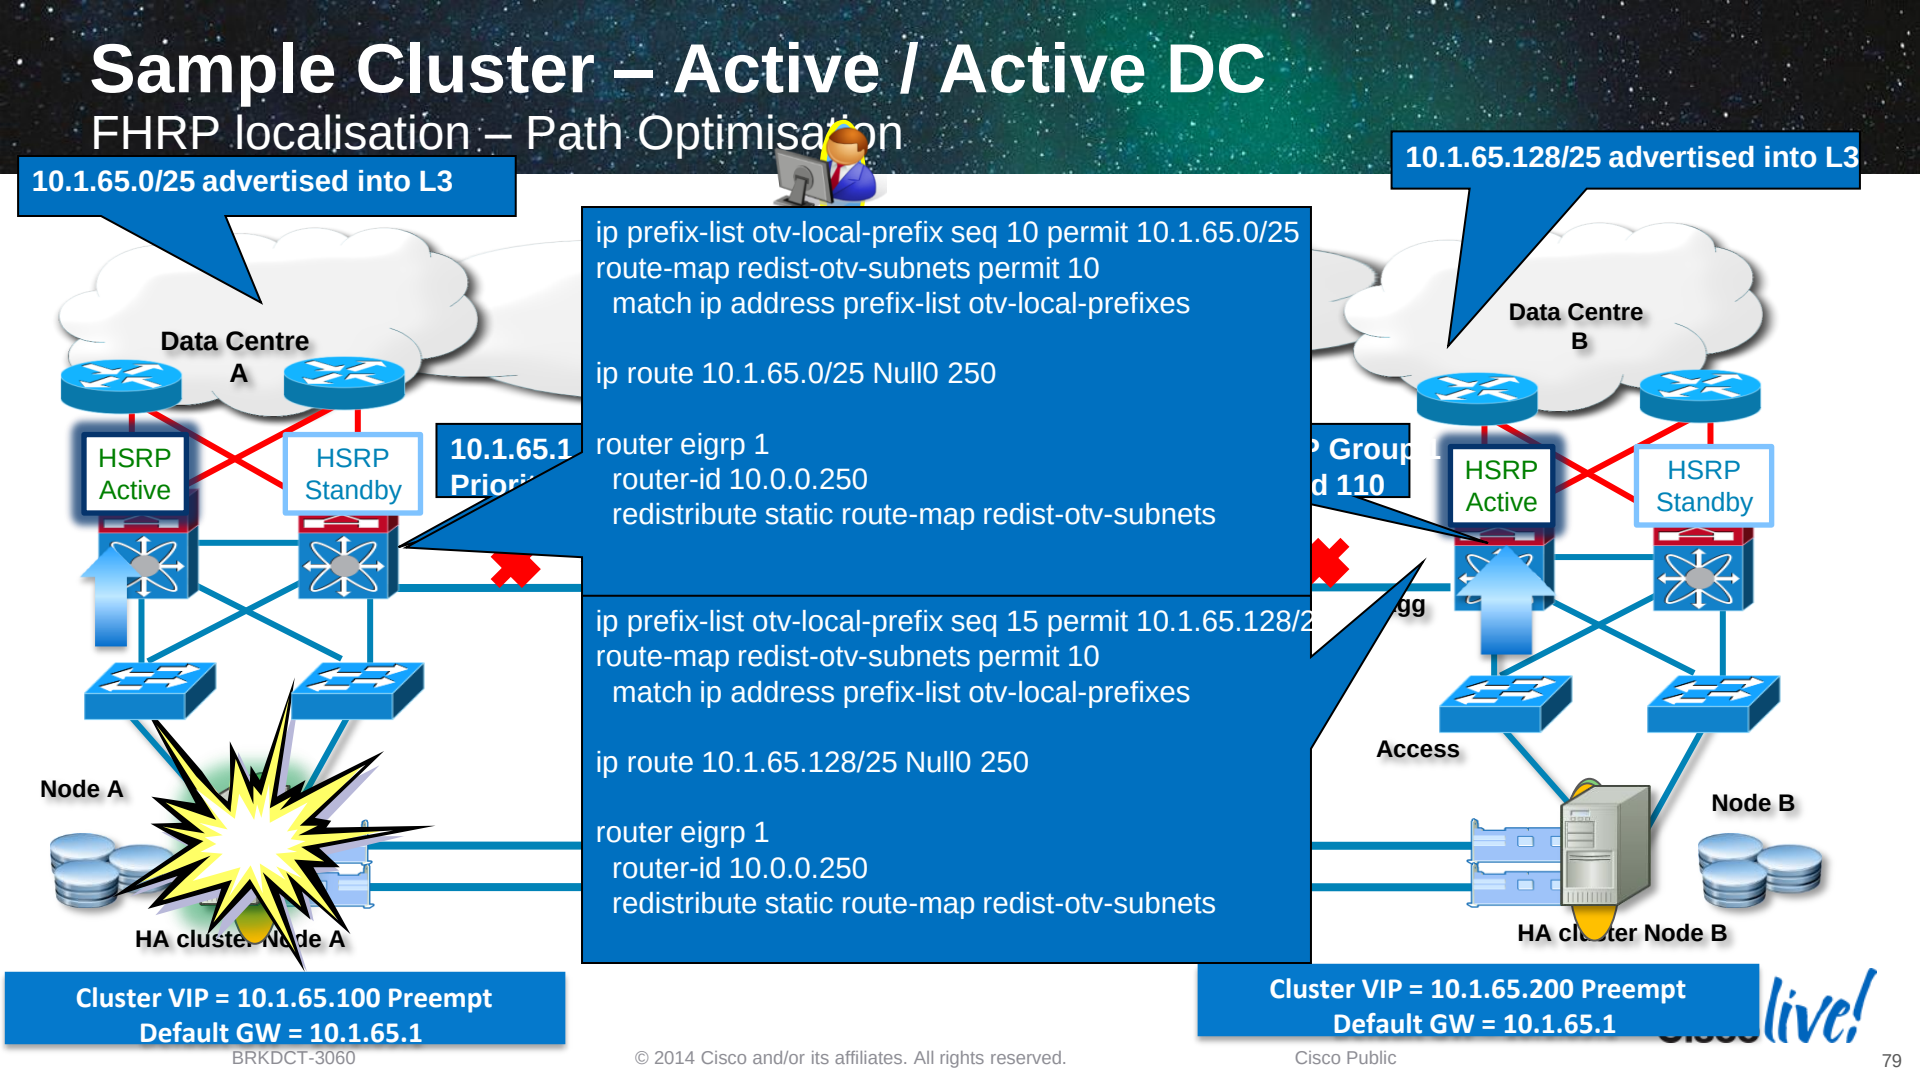

LISP Roles and Address Spaces

What Are the Different Components Involved?

LISP Roles

- **Tunnel Routers - xTRs**
  - Edge devices in charge of encap/decap
  - Ingress/Egress Tunnel Routers (ITR/ETR)
- **EID to RLOC Mapping DB**
  - Contains RLOC to EID mappings
  - Distributed across multiple Map Servers (MS)
  - MS may connect over an ALT network
- **Proxy Tunnel Routers - PxTR**
  - Coexistence between LISP and non-LISP sites
  - Ingress/Egress: PITR, PETR

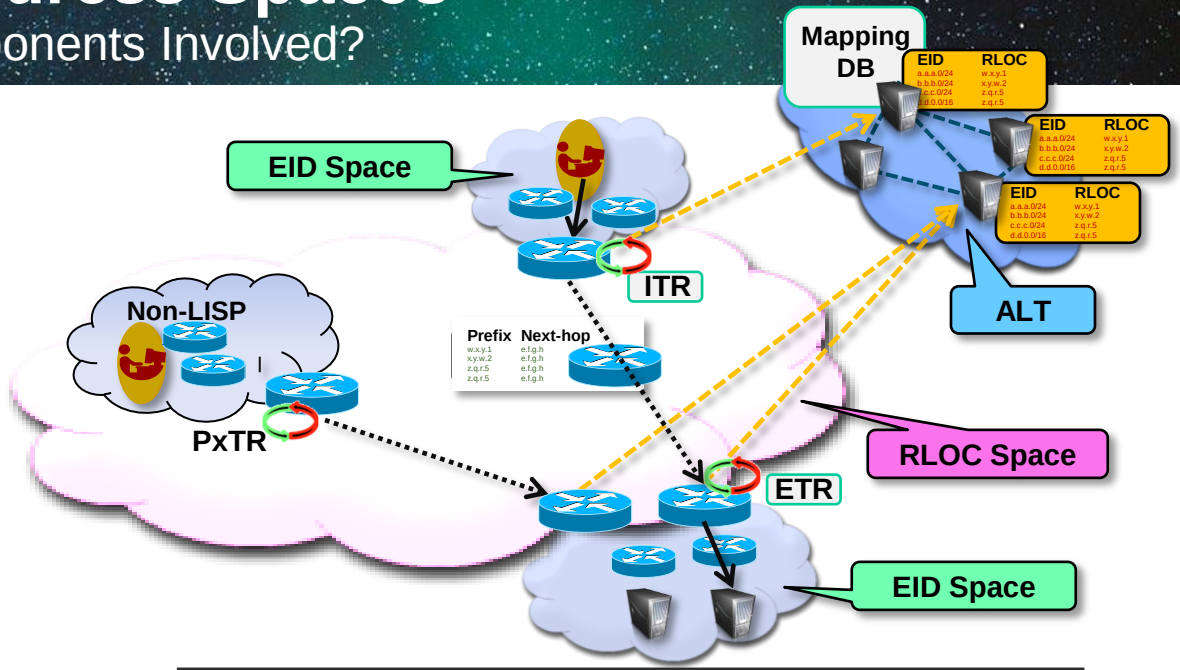

Address Spaces

- **EID = End-point Identifier**
  - Host IP or prefix
- **RLOC = Routing Locator**
  - IP address of routers in the backbone

A LISP Packet Walk

How Does LISP Operate?

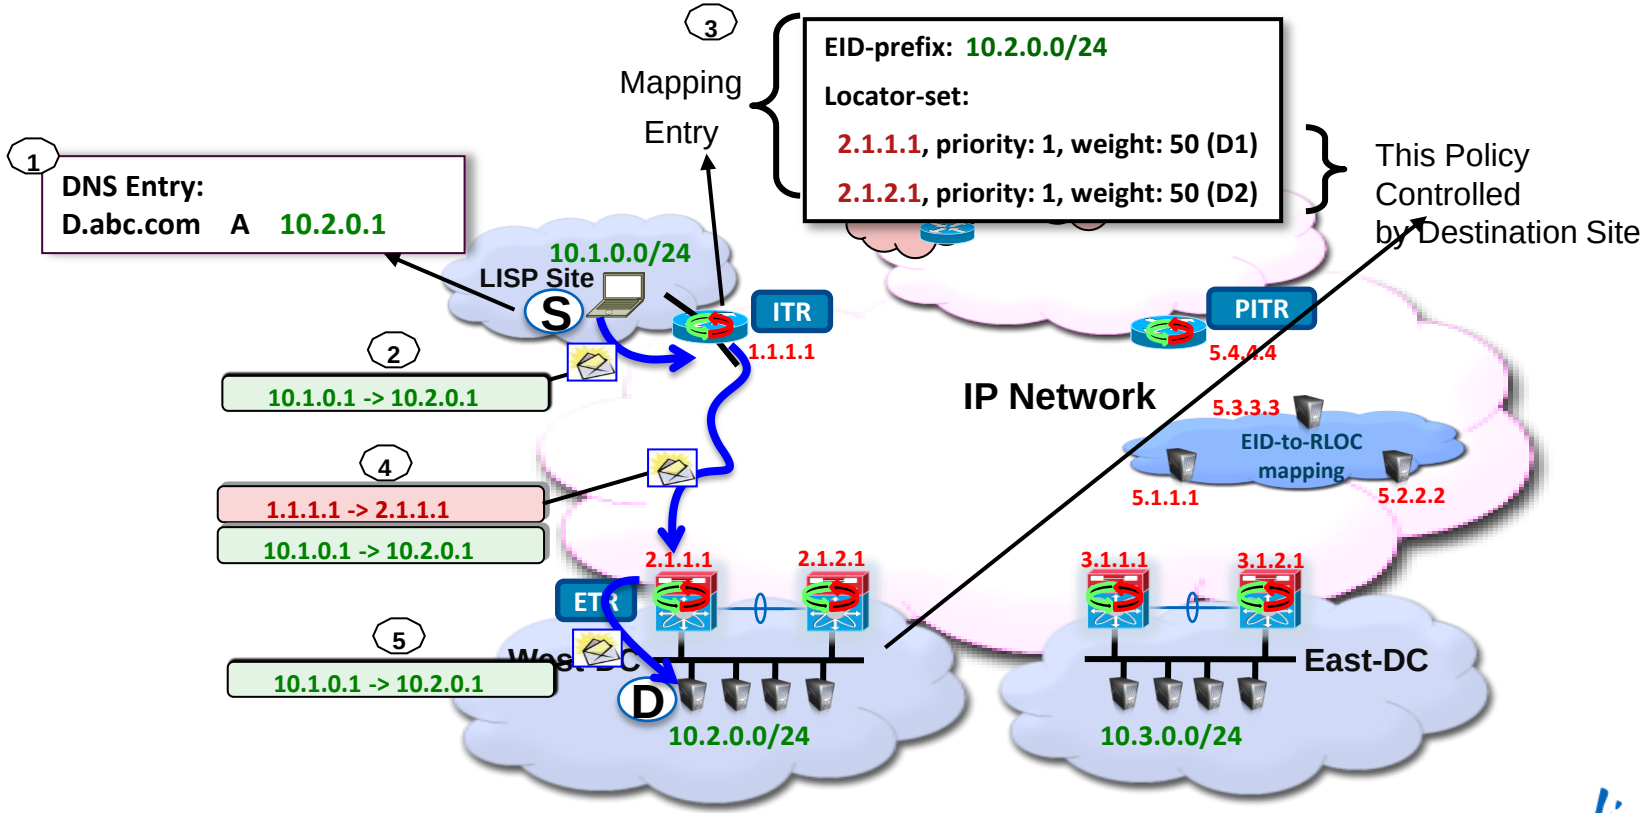

LISP Host-mobility Multi-hop / ESM

Inserting non-LISP Capable HW (FWs and non-LISP M-cards)

- xTR is not the first hop router
- LISP host-mobility functionality is split to two places:
  - SG (Site GWY) → LISP registration/encap/decap (it's an XTR)
  - 1st Hop router (FH) → Move detection and signalling to SG, proxy default GWY
- The SG LISP registers host mappings in the dynamic-eid range (just like XTRs)
- SGs will register the detected hosts based on either:
  - EID-notify messages received from FH (or)
  - Host routes received from FH

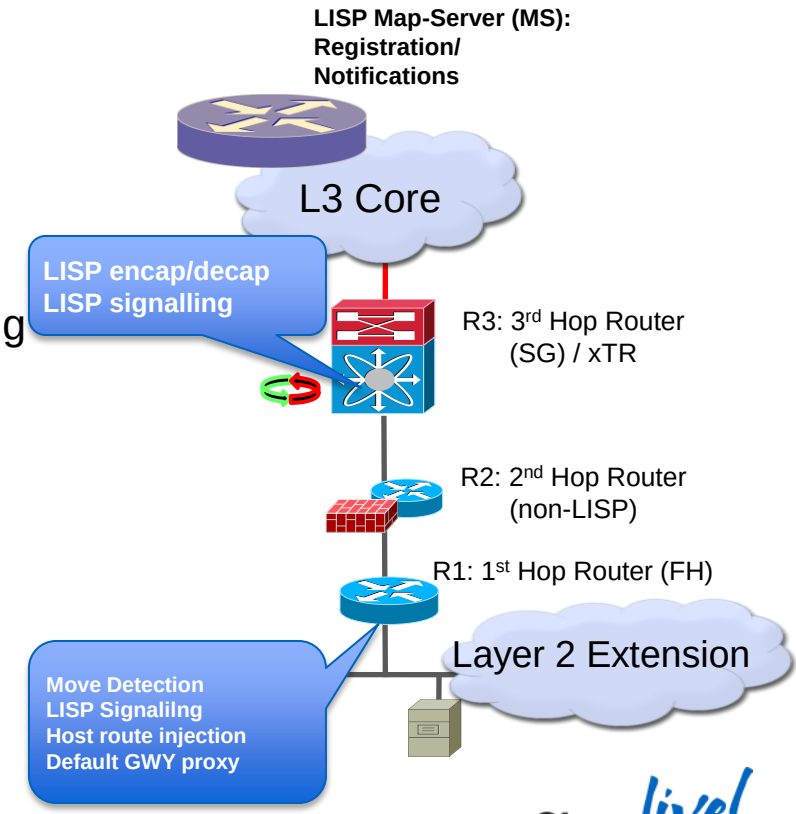

LISP Host-mobility Multi-hop / ESM

@FHR

```
ip lisp etr
lisp dynamic-eid site1
  database-mapping 10.1.1.0/24 100.1.1.1 pr 10 w 50
  eid-notify 12.36.0.3 key 3 75095fe9112836e3
  map-notify-group 225.1.1.1
lisp dynamic-eid site2
  database-mapping 10.2.2.0/24 100.1.1.1 pr 10 w 50
  eid-notify 12.36.0.3 key 3 75095fe9112836e3
  map-notify-group 225.2.2.2

interface Vlan11
  lisp mobility site1
  lisp extended-subnet-mode
  ip address 10.1.1.3/24
  ip router ospf 100 area 0.0.0.1
  hsrp 1
      ip 10.1.1.1

interface Vlan12
  lisp mobility site2
  lisp extended-subnet-mode
  ip address 10.2.2.3/24
  ip router ospf 100 area 0.0.0.1
  hsrp 1
    ip 10.2.2.1
```

```
ip lisp etr
lisp dynamic-eid site1
  database-mapping 10.1.1.0/24 100.1.1.1 pr 10 w 50
  eid-notify 21.24.0.4 key 3 6d018260cf71b07c
  map-notify-group 225.1.1.1
lisp dynamic-eid site2
  database-mapping 10.2.2.0/24 100.1.1.1 pr 10 w 50
  eid-notify 21.24.0.4 key 3 6d018260cf71b07c
  map-notify-group 225.2.2.2

interface Vlan11
  lisp mobility site1
  lisp extended-subnet-mode
  ip address 10.1.1.2/24
  ip router ospf 200 area 0.0.0.2
  hsrp 1
      ip 10.1.1.1

interface Vlan12
  lisp mobility site2
  lisp extended-subnet-mode
  ip address 10.2.2.2/24
  ip router ospf 100 area 0.0.0.2
  hsrp 1
    ip 10.2.2.1
```

LISP Host-mobility Multi-hop / ESM
@SG/XTR

```
ip lisp itr-etr
ip lisp database-mapping 10.1.0.0/16 100.3.3.3 priority 10 weight 50
ip lisp database-mapping 10.2.0.0/16 100.3.3.3 priority 10 weight 50
ip lisp itr map-resolver 100.5.5.5
ip lisp etr map-server 100.5.5.5 key 3 0b50279df3929e28
lisp dynamic-eid site1
  database-mapping 10.1.1.0/24 100.3.3.3 priority 10 weight 50
  eid-notify authentication-key 3 75095fe9112836e3
lisp dynamic-eid site2
  database-mapping 10.2.2.0/24 100.3.3.3 priority 10 weight 50
  eid-notify authentication-key 3 75095fe9112836e3
```

```
ip lisp itr-etr
ip lisp database-mapping 10.1.0.0/16 100.4.4.4 priority 10 weight 50
ip lisp database-mapping 10.2.0.0/16 100.4.4.4 priority 10 weight 50
ip lisp itr map-resolver 100.5.5.5
ip lisp etr map-server 100.5.5.5 key 3 0b50279df3929e28
lisp dynamic-eid site1
  database-mapping 10.1.1.0/24 100.4.4.4 priority 10 weight 50
  eid-notify authentication-key 3 6d018260cf71b07c
lisp dynamic-eid site2
  database-mapping 10.2.2.0/24 100.4.4.4 priority 10 weight 50
  eid-notify authentication-key 3 6d018260cf71b07c
```

LISP Host-mobility IGP Assist / ESM

- Host routing end to end
- LISP provides host mobility detection
- LISP provides signalling to guide IGP convergence
- The IGP propagates host routes received from LISP
- No LISP encapsulation involved

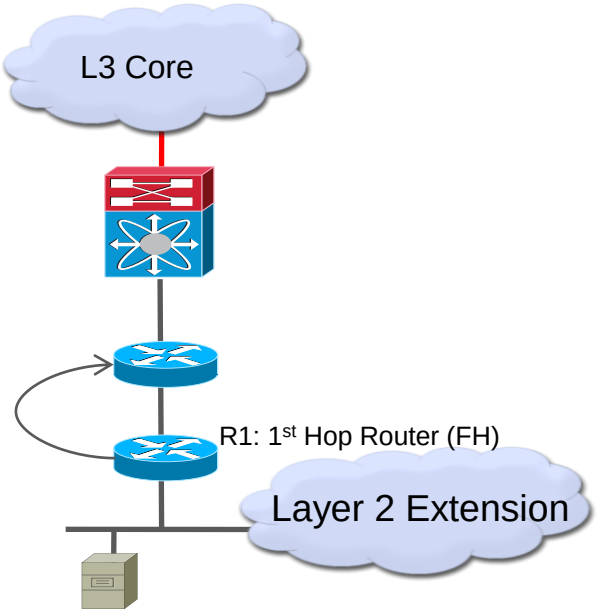

LISP Host-mobility IGP Assist / ESM
@FHR

```
ip lisp etr
lisp dynamic-eid foo
      database-mapping <eid-prefix> <rloc-add-1> p1 w50
      database-mapping <eid-prefix> <rloc-add-2> p1 w50
      map-notify-group 239.1.1.1

router <favorite-routing-protocol> foo
      redistribute lisp route-map <bar>

ip prefix-list <eid-list-name> seq 5 permit <eid-prefix> ge 32
route-map <bar> permit 10
  match ip address <eid-list-name>
```

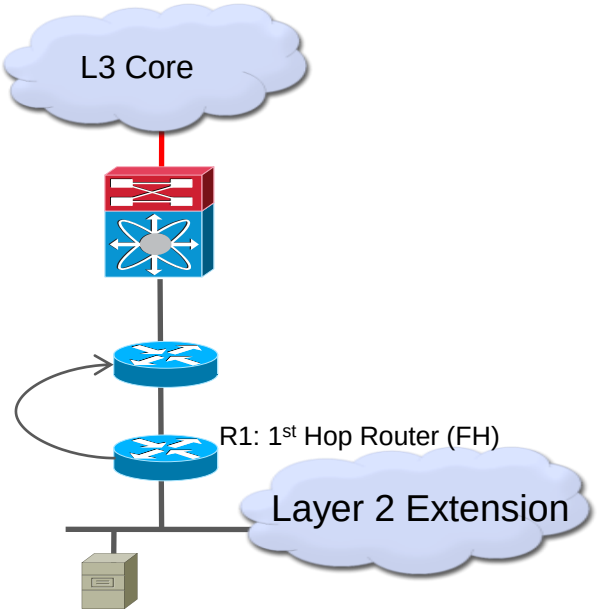

Services - Live Moves

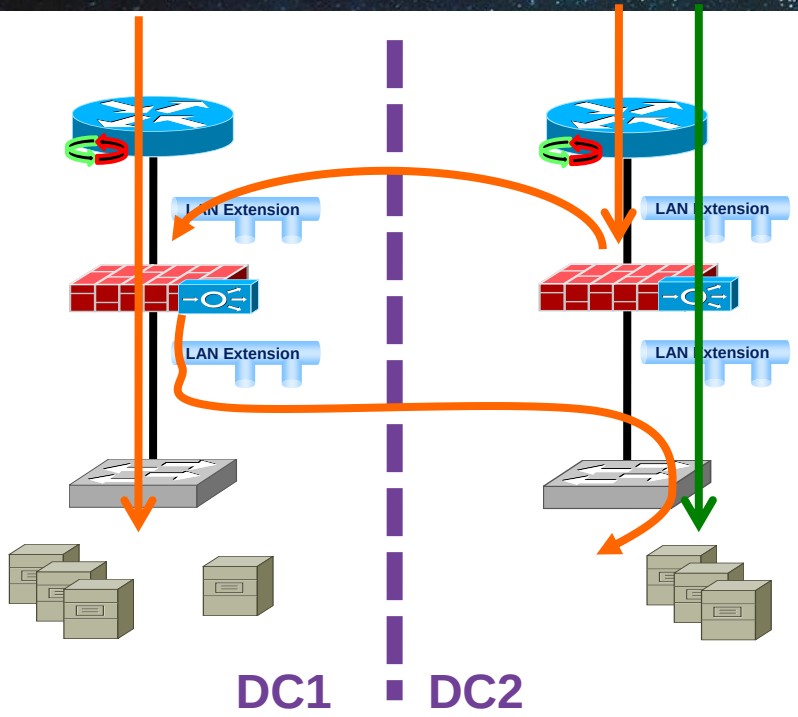

- Redirection of established flows:
  - Extended Clusters
  - Cluster or LISP based re-direction

Services – Cold Moves

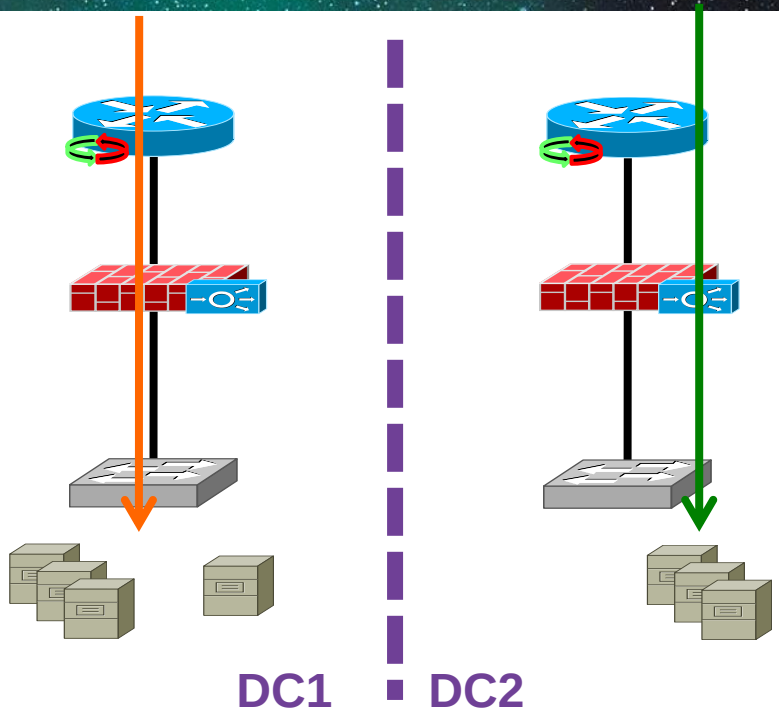

- IP preservation ➔ Uniform Policies

Established before the move

Established after the move

LISP VM-Mobility Configuration

With Extended Subnets → "extended-subnet-mode"

```
ip lisp ITR-ETR
ip lisp database-mapping 10.2.0.0/16 <RLOC-A>
ip lisp database-mapping 10.2.0.0/16 <RLOC-B>

lisp dynamic-eid roamer
    database-mapping 10.2.0.0/24 <RLOC-A> p1 w50
    database-mapping 10.2.0.0/24 <RLOC-B> p1 w50
    map-server 1.1.1.1 key abcd
    map-server 2.2.2.1 key abcd
    map-notify-group 239.10.10.10
    ip lisp itr map-resolver 5.3.3.3
interface vlan 100
    ip address 10.2.0.10 /16
    lisp mobility  roamer
    lisp extended-subnet-mode
    hsrp 101
        ip 10.2.0.1
```

```
ip lisp ITR-ETR

lisp dynamic-eid roamer
    database-mapping 10.2.0.0/24 <RLOC-C> p1 w50
    database-mapping 10.2.0.0/24 <RLOC-D> p1 w50
    map-server 1.1.1.1 key abcd
    map-server 2.2.2.1 key abcd
    map-notify-group 239.10.10.10
    ip lisp itr map-resolver 5.3.3.3
interface vlan 100
    ip address 10.2.0.11 /16
    lisp mobility  roamer
    lisp extended-subnet-mode
    hsrp 101
        ip 10.2.0.1
```

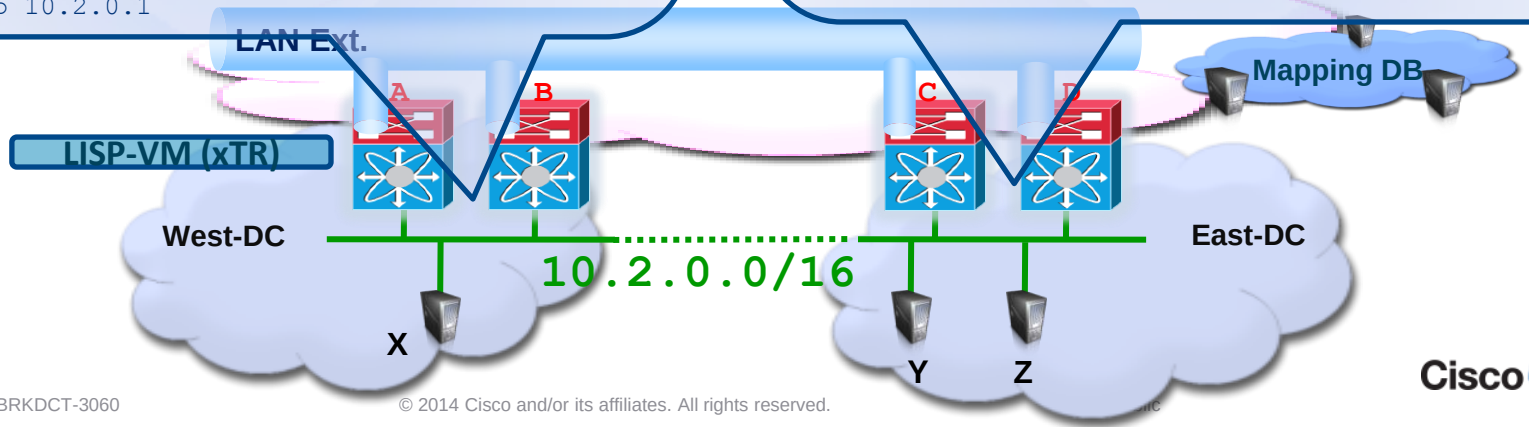

OTV - HSRP Localisation – OTV Edge Device

1) Define HSRPv1 and HSRPv2 to block HSRP Hello Messages

```
ip access-list ALL_IPs
  10 permit ip any any
!
mac access-list ALL_MACs
  10 permit any any
!
ip access-list HSRP_IP
  10 permit udp any 224.0.0.2/32 eq 1985
  20 permit udp any 224.0.0.102/32 eq 1985

vlan access-map HSRP_Local 10
        match ip address HSRP_IP
        action drop
vlan access-map HSRP_Local 20
        match ip address ALL
        action forward
```

OTV - HSRP Localisation - OTV Edge Device

2) Prevent Duplicate HSRP Gratuitous ARP from HSRP VIP

```
arp access-list HSRP_VMAC_ARP
  10 deny ip any mac 0000.0c07.ac00 ffff.ffff.ff00
  20 deny ip any mac 0000.0c9f.f000 ffff.ffff.f000
  30 permit ip any mac any

feature dhcp
ip arp inspection filter HSRP_VMAC_ARP 10,11,600, 601, 700,
701

interface Vlan10
  no shutdown
  no ip redirects
  ip address 192.168.10.3/24
  no ip arp gratuitous hsrp duplicate
  hsrp 10
    priority 110
    ip 192.168.10.1
```

Message without: %ARP-3-DUP_VADDR_SRC_IP: arp [3849] Source address of packet received from 0000.0c07.ac1f on Vlan10(port-channel10) is duplicate of local virtual ip, 192.168.10.1

OTV – HSRP Localisation - OTV Edge Device

3) Filter learning HSRP Virtual MAC address across OTV

```
mac access-list HSRP_VMAC
  10 permit 0000.0c07.ac00 0000.0000.00ff any
  20 permit 0000.0c9f.f000 0000.0000.0fff any
!
vlan access-map HSRP_localisation 10
        match mac address HSRP_VMAC
        match ip address HSRP_IP
        action drop
!
vlan access-map HSRP_localisation 20
        match mac address ALL_MACs
        match ip address ALL_IPs
        action forward
!
vlan filter HSRP_Local vlan-list 10,11,600, 601, 700,
701
```

```
mac-list HSRP_VMAC_Deny seq 5 deny
0000.0c07.ac00 ffff.ffff.ff00
mac-list HSRP_VMAC_Deny seq 10 deny
0000.0c9f.f000 0000.0000.0fff
mac-list HSRP_VMAC_Deny seq 15 permit
0000.0000.0000 0000.0000.0000
!
route-map stop-HSRP permit 10
match mac-list HSRP_VMAC_Deny
!
otv-isis default
vpn Overlay0
redistribute filter route-map stop-HSRP
```

VPLS Localisation

1) Configure virtual port-channel (vPC) on BOTH Nexus 7000 aggregation switches and filter HSRP

```
interface Ethernet2/1
lacp rate fast
switchport
switchport mode trunk
switchport trunk allowed vlan 1,76-80,100-349
channel-group 31 mode active
no shutdown

interface Ethernet2/2
lacp rate fast
switchport
switchport mode trunk
switchport trunk allowed vlan 1200-1449
channel-group 32 mode active
no shutdown
```

```
interface Ethernet2/6
lacp rate fast
switchport
switchport mode trunk
switchport trunk allowed vlan 1,76-80,100-349
channel-group 31 mode active
no shutdown

interface Ethernet2/3
lacp rate fast
switchport
switchport mode trunk
switchport trunk allowed vlan 1200-1449
channel-group 32 mode active
no shutdown
```

VPLS Localisation

2) Access list to filter HSRP hellos configured on both aggregation switches

```
ip access-list HSRP_Deny
statistics per-entry
10 deny udp any 224.0.0.102/32 eq 1985
20 permit ip any any
```

VPLS Localisation

3) Configure port-channel interface on BOTH Nexus 7000 aggregation switches

```
interface port-channel31
switchport
switchport mode trunk
ip port access-group HSRP_Deny in
switchport trunk allowed vlan 1,76-80,100-349
spanning-tree port type edge trunk
spanning-tree bpdufilter enable
vpc 31
```

```
interface port-channel32
switchport
switchport mode trunk
ip port access-group HSRP_Deny in
switchport trunk allowed vlan 1200-1449
spanning-tree port type edge trunk
spanning-tree bpdufilter enable
lacp max-bundle 1
vpc 32
```

Summary State-full Devices Placement with DCI

- Ping-Pong effect might have a bad impact in term of perf with long distances:
  - Greedy bandwidth
  - Latency
- It is commonly accepted to distribute traditional A/S state-full devices between 2 Twin DC for short Metro Distances (+/- 10km max)
  - Keep transparency and easy to operate
  - limited to 2 Active DC
- As of today the preferred method is to deploy Stretch ASA clustering across distributed DC (Metro)
  - All ASA active
  - Not limited to 2 Active DC
- For Geographical Distributed DC
  - if Hot migration is required (i.e. Geo VPLEX), use ASA cluster stretched over multiple sites with LAN extension
  - for Cold migration use ASA cluster distributed per site in conjunction with LISP
- Ingress Path Optimisation
  - LISP Mobility is the preferred choice – It requires LISP Multi-hop
  - GSLB (DNS and KAP-AP) can help to redirect the traffic accordingly, but may face some caveats with proxy DNS and client caching
  - RHI can help but offers App based granularity only for Intranet core (Enterprise owns the L3 core)
- The recommended choice is ASA clustering in conjunction with the traditional DNS and LISP Mobility.
  - Stretched across multiple DC with LAN extension for Hot Migration
  - Confined inside each DC without LAN extension for Cold Migration

Data Centre Interconnect
Agenda

- Mobility and Virtualisation in the Data Centre
- LAN Extension Deployment Scenarios
  - Ethernet Based Solutions
  - MPLS Based Solutions
    - EoMPLS
    - VPLS
    - A-VPLS
    - EVPN
- Overlay Transport Virtualisation (OTV)
- Encryption
- Path Optimisation
- IP Mobility without LAN Extension
- Fabric Solutions
- Summary and Conclusions
- Q&A

= For your Reference

IP Mobility without LAN Extension

Host-Mobility Scenarios

Moves Without LAN Extension (ASM)

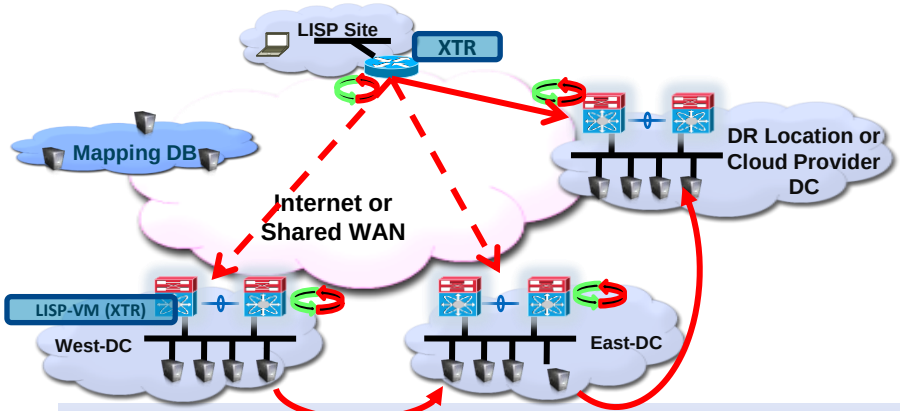

IP Mobility Across Subnets

Disaster Recovery

Cloud Bursting

Application Members in One Location

Moves With LAN Extension (ESM)

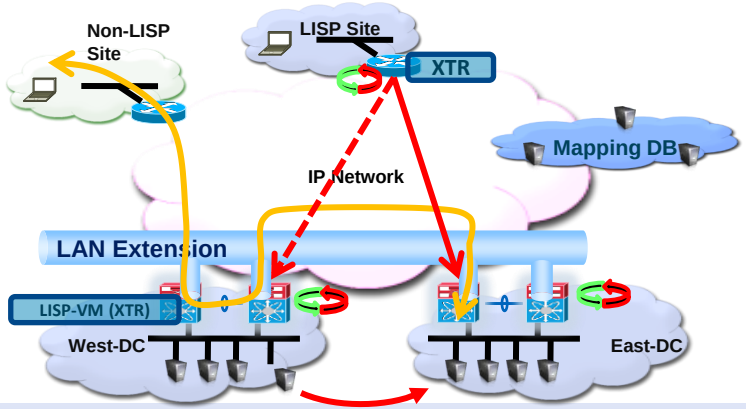

Routing for Extended Subnets

Active-Active Data Centres

Distributed Clusters

Application Members Distributed (Broadcasts across sites)

LISP Host-Mobility Configuration

Without LAN Extensions

```
ip lisp ITR-ETR
ip lisp database-mapping 10.2.0.0/16 <RLOC-A>
ip lisp database-mapping 10.2.0.0/16 <RLOC-B>

lisp dynamic-eid roamer
    database-mapping 10.2.0.0/24 <RLOC-A> p1 w50
    database-mapping 10.2.0.0/24 <RLOC-B> p1 w50
    map-server 1.1.1.1 key abcd
    map-server 2.2.2.1 key abcd
    map-notify-group 239.1.1.1
    ip lisp itr map-resolver 5.3.3.3
interface vlan 100
    ip address 10.2.0.10 /16
    lisp mobility  roamer
    ip proxy-arp
    hsrp 101
        mac-address 0000.0e1d.010c
        ip 10.2.0.1
```

```
ip lisp ITR-ETR
ip lisp database-mapping 10.3.0.0/16 <RLOC-C>
ip lisp database-mapping 10.3.0.0/16 <RLOC-D>

lisp dynamic-eid roamer
    database-mapping 10.2.0.0/24 <RLOC-C> p1 w50
    database-mapping 10.2.0.0/24 <RLOC-D> p1 w50
    map-server 1.1.1.1 key abcd
    map-server 2.2.2.1 key abcd
    map-notify-group 239.2.2.2
    ip lisp itr map-resolver 5.3.3.3
interface vlan 100
    ip address 10.3.0.11 /16
    lisp mobility  roamer
    ip proxy-arp
    hsrp 201
        mac-address 0000.0e1d.010c
        ip 10.3.0.1
```

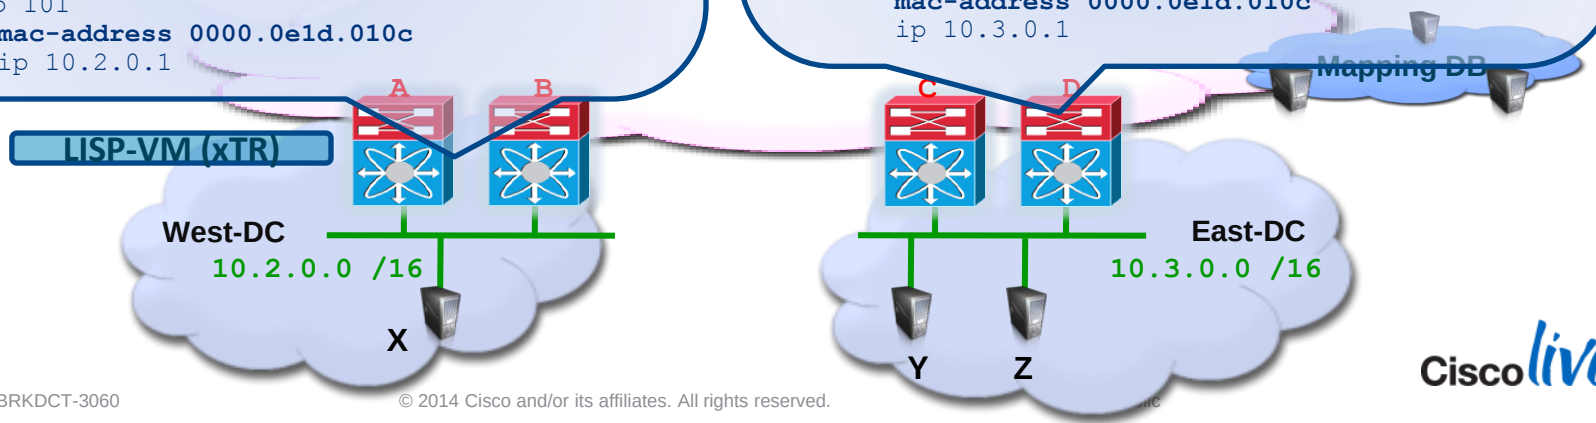

LISP Host-Mobility – Move Detection

Monitor the source of Received Traffic

- The new xTR checks the source of received traffic
- Configured dynamic-EIDs define which prefixes may roam

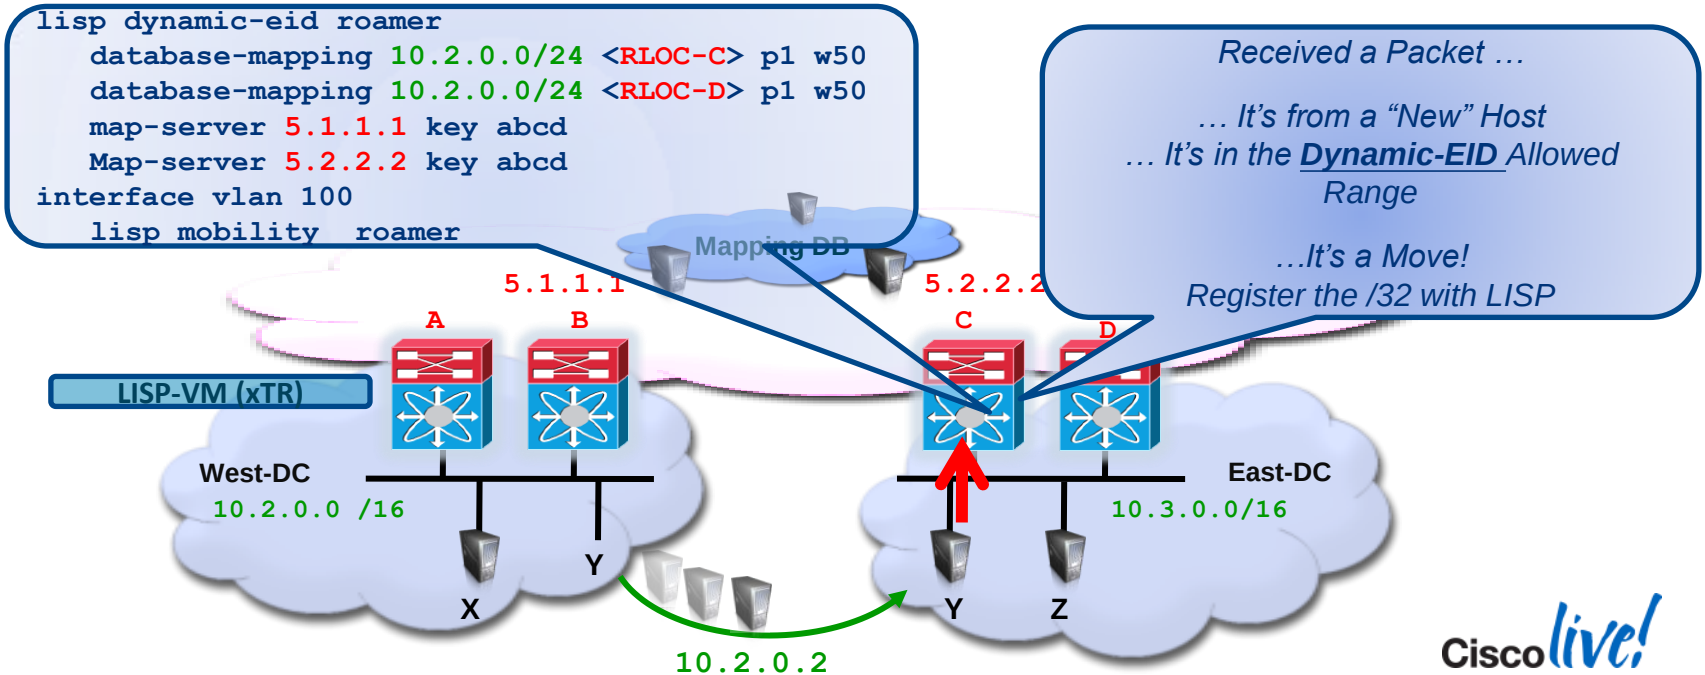

LISP Host-Mobility – Traffic Redirection

Update Location Mappings for the Host System Wide

- When a host move is detected, updates are triggered:
  - The host-to-location mapping in the Database is updated to reflect the new location
  - The old ETR is notified of the move
  - ITRs are notified to update their Map-caches
- Ingress routers (ITRs or PITRs) now send traffic to the new location
- Transparent to the underlying routing and to the host

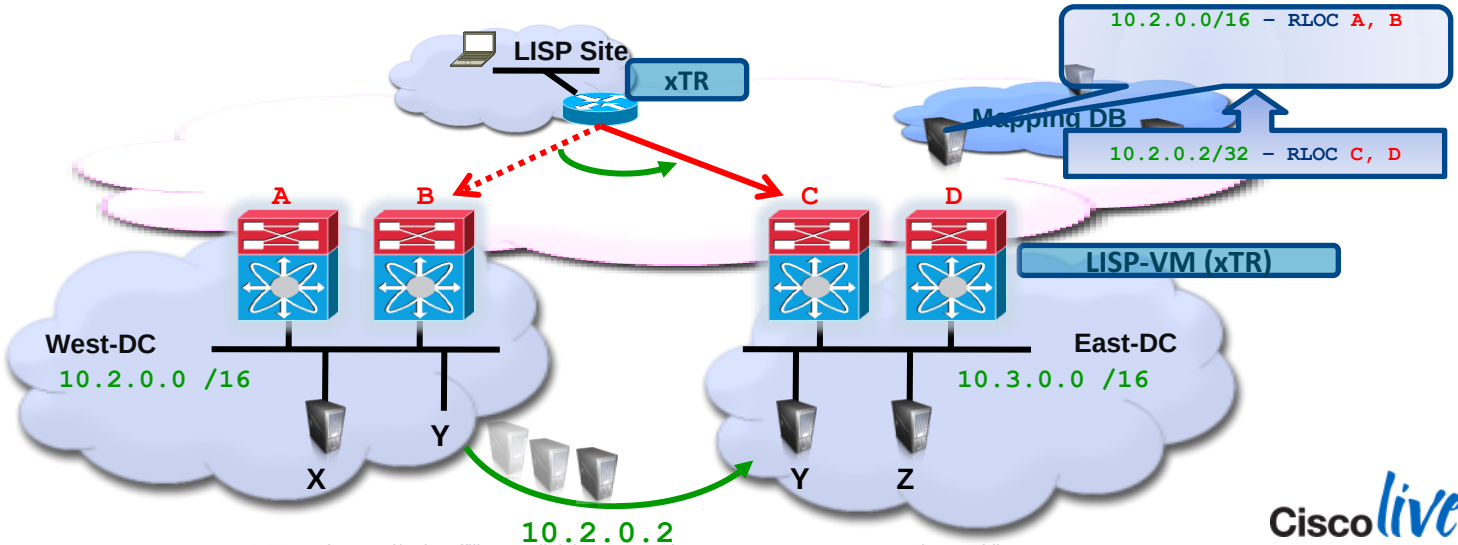

LISP Host-Mobility – First Hop Routing

No LAN Extension – Across subnet mode (ASM)

- 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
- Hosts can move anywhere and always talk to a local gateway with the same MAC
- Totally transparent to the moving hosts

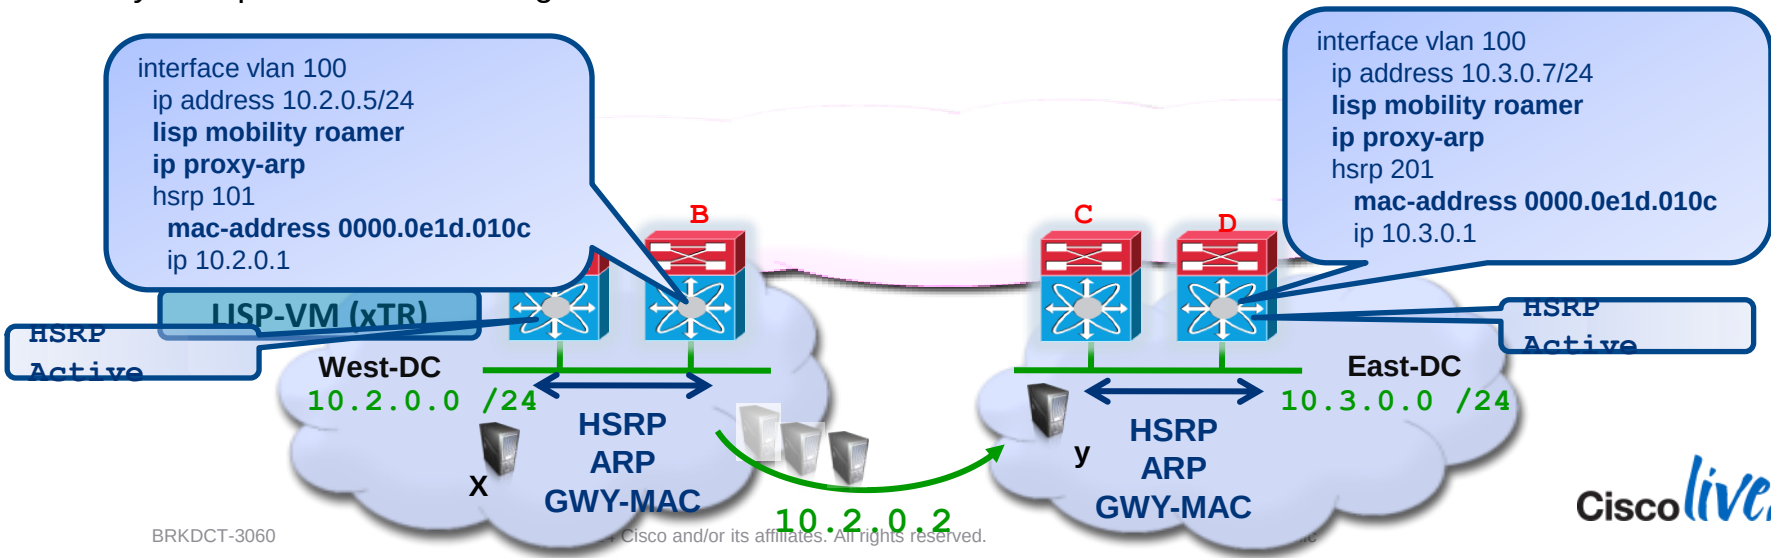

On-subnet Server-Server Traffic

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

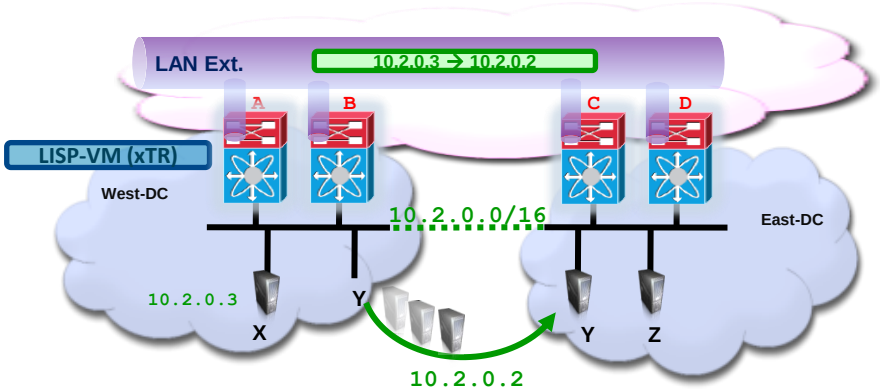

Without LAN Extensions

- 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

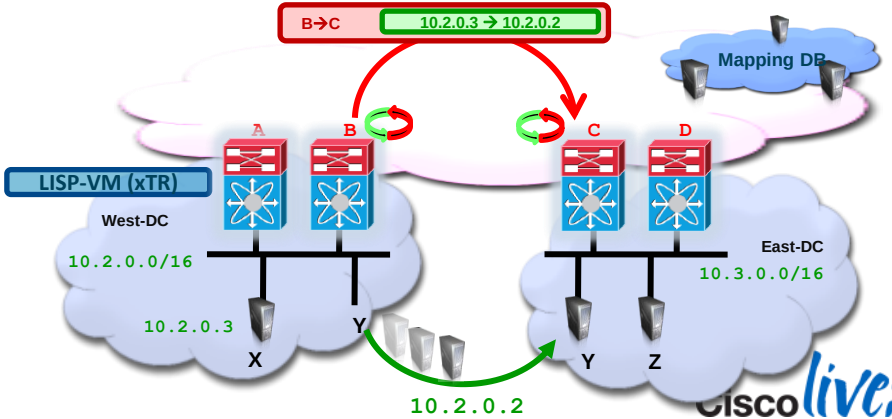

Data Centre Interconnect
Agenda

- Mobility and Virtualisation in the Data Centre
- LAN Extension Deployment Scenarios
  - Ethernet Based Solutions
  - MPLS Based Solutions
    - EoMPLS
    - VPLS
    - A-VPLS
    - EVPN
- Overlay Transport Virtualisation (OTV)
- Encryption
- Path Optimisation
- IP Mobility without LAN Extension
- Fabric Solutions
- Summary and Conclusions
- Q&A

= For your Reference




Fabric Solutions

L2 Host Overlays and Virtualisation – VXLAN

Creating virtual segments

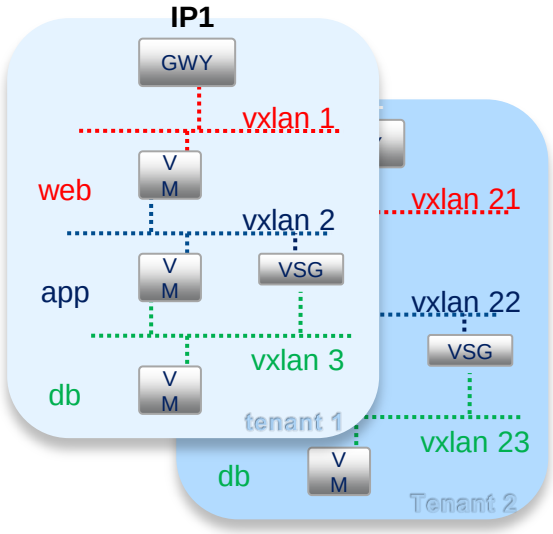

Multi-tier Virtual App = VMs + Segments + Gateway

Application: Cloud Services

VXLAN elastic creation of **virtual Segments**

- Small Segments
  - Usually don't stretch outside of a POD
- Mobile: Can be instantiated anywhere
  - Move along with VMs as necessary
- Very large number of segments
  - Do not consume resources in the network core
- Host overlays are initiated at the hypervisor virtual switch ➔ Virtual hosts only
- Gateway to connect to the non-virtualised world
- VXLAN shipping since 2011 on Cisco Nexus 1000v, other variants: NVGRE, STT

LISP Enables VXLAN to Deliver vApp Mobility

- Move virtual Applications (vApps) among private cloud PODs
  - Move VMs and virtual Segments (VXLANs)
- LISP host mobility allows the vApp to roam
  - Maintain optimal path for Client-Server connectivity
  - Maintain GWY IP address, segmentation and optimal reachability

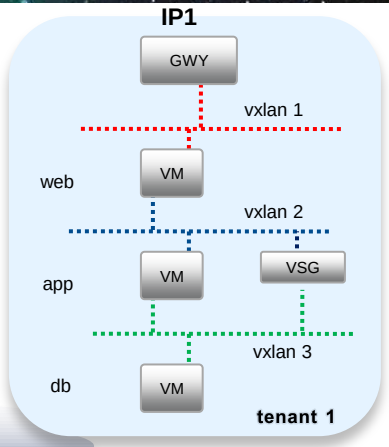

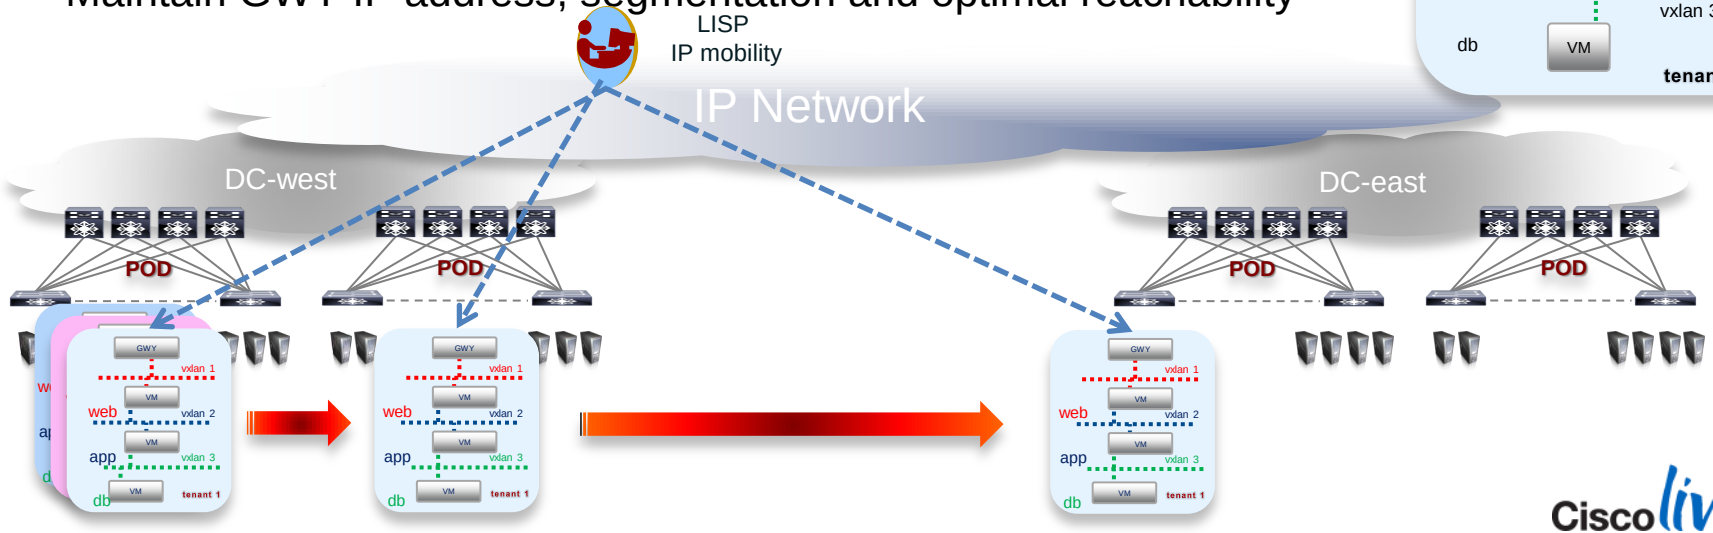

VXLAN IP Fabric

L3-VXLAN & LISP IP Mobility @ DC Core

LISP Mobility:

- LISP encapsulation from client sites
- No host routing in the IP core
- Direct Path Forwarding
- VXLAN Localised to Data Centre

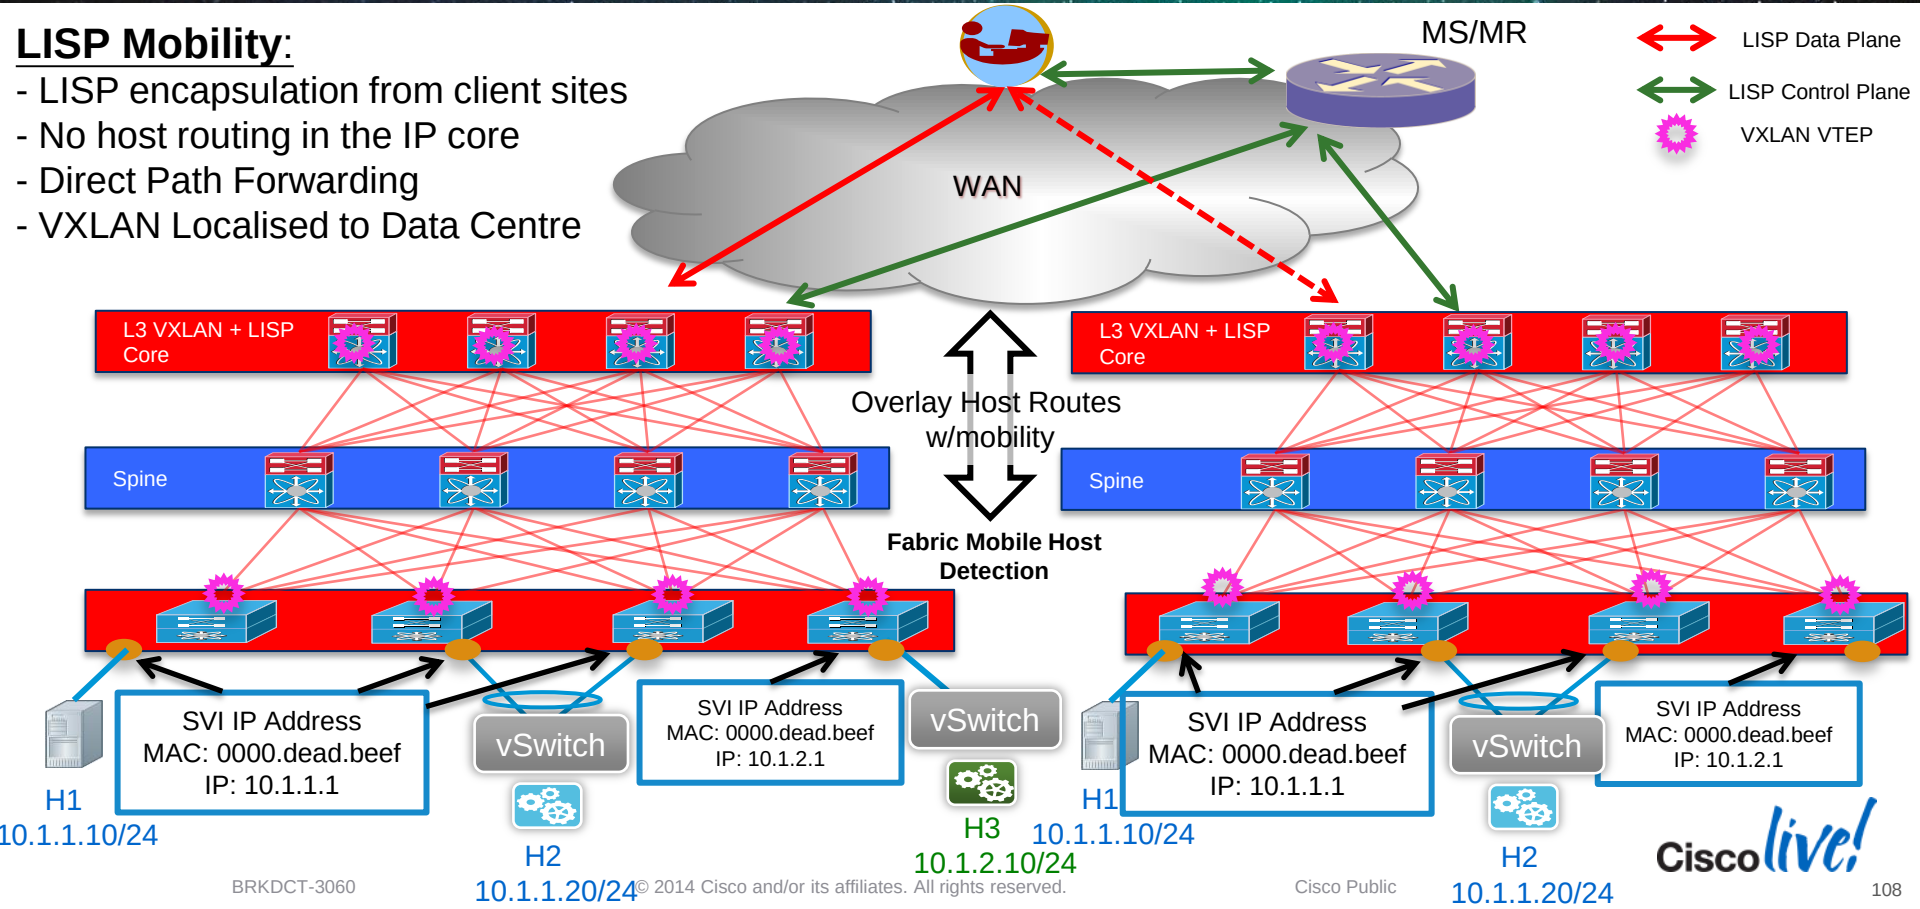

What's Missing from VXLAN for DCI?

North-south VXLAN limitations:

- Only one gateway per segment
  - More than one Gateway will lead to loops
  - Traffic is tromboned to the Gateway
  - Defeats the purpose of the geographic dispersion

East-west VXLAN limitations

- No isolation of L2 failure-domain
- VTEP Discovery (Flood to Learn) / Security
- Excessive flood traffic and BW exhaustion
- Large amounts of IP multicast between DCs
- No network resiliency of the L2 overlay

The DCI toolkit solves all these issues in LISP, OTV, MPLS and EVPN

- Loop resolution

- Multi-homing

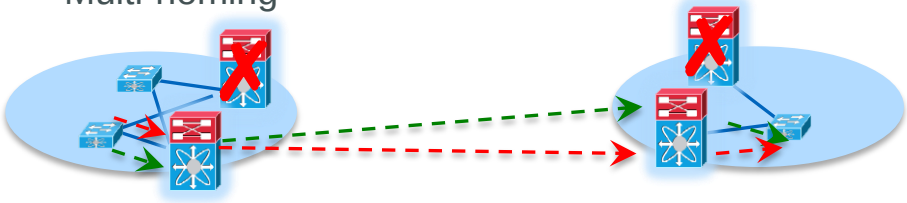

- Broadcast de-duplication

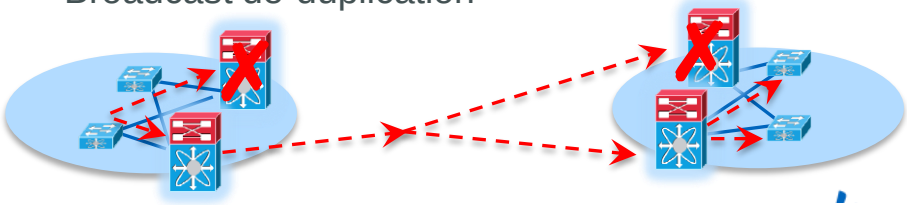

Fabric Terminology

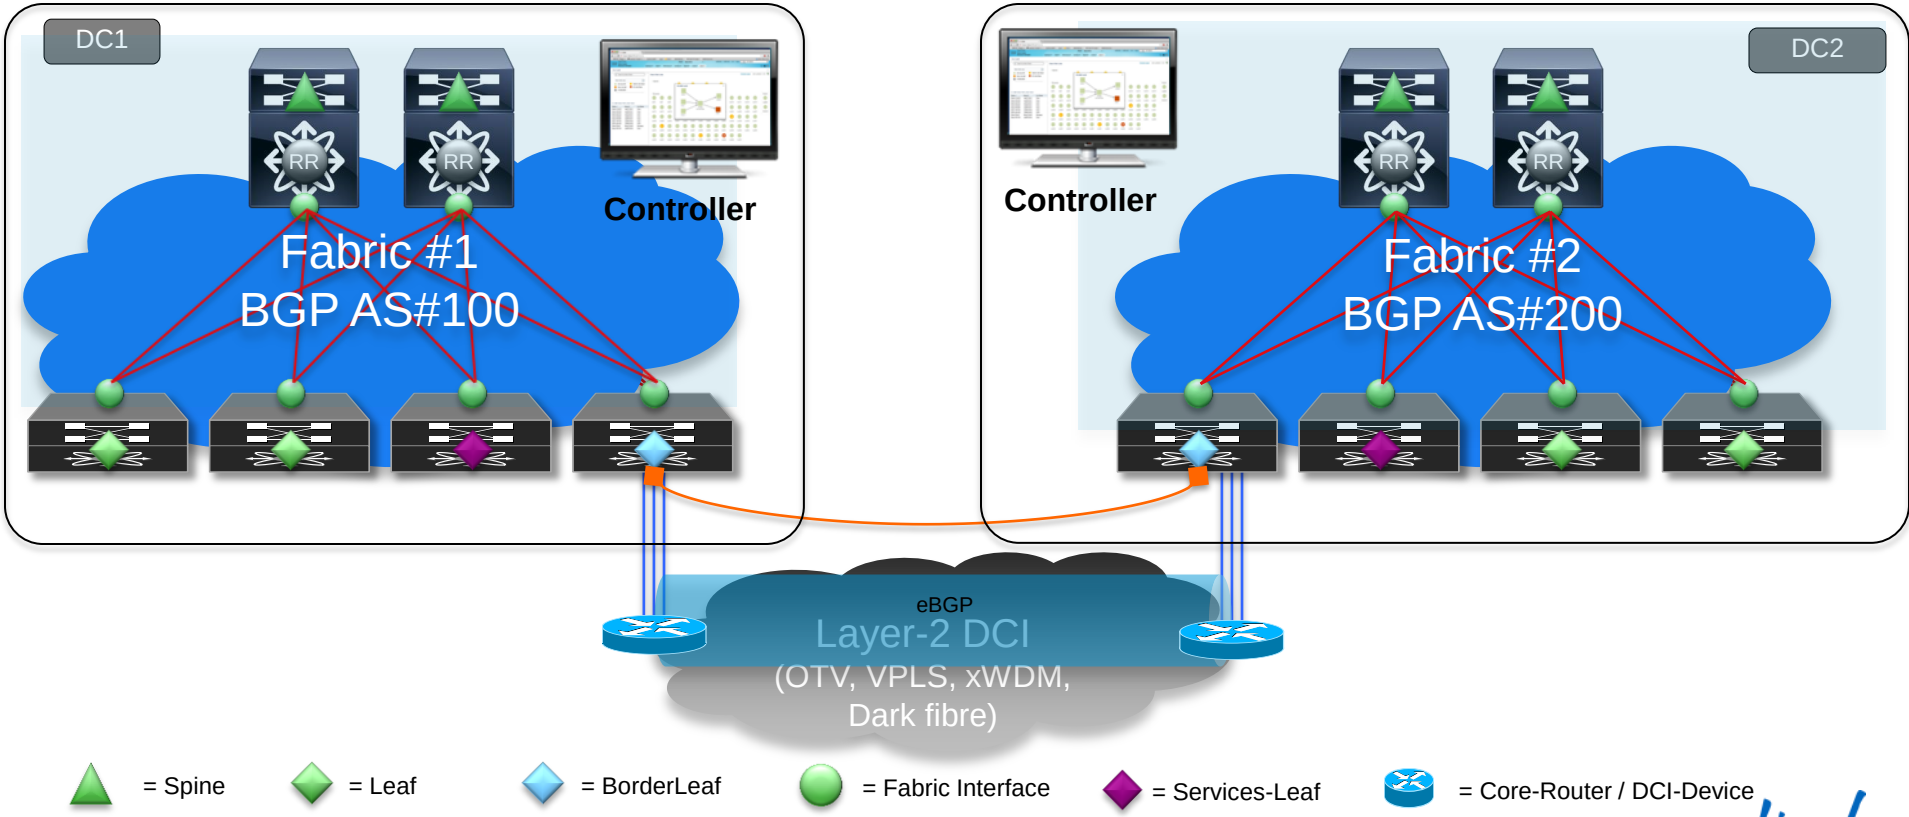

So, do Things Change with Fabrics?

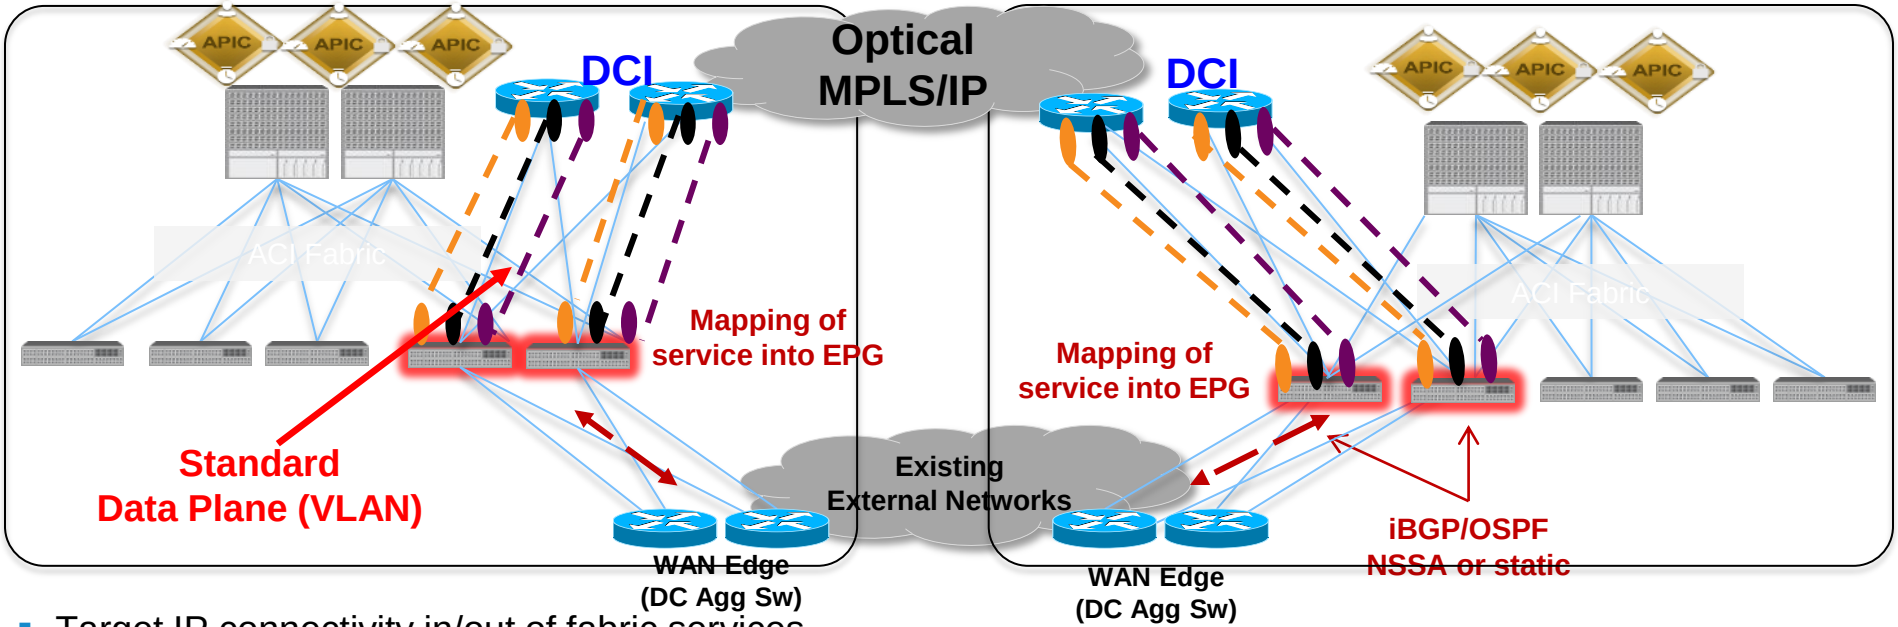

- Target IP connectivity in/out of fabric services
- FCS routing protocol support targets OSPF, iBGP (support for “set community”), and static
- WAN Edge focus: ASR 9000, Nexus 7000, ASR 1000
- Existing principles of Inbound, Outbound traffic flows, DNS/GSS still apply

Dynamic Fabric Automation (DFA) and DCI

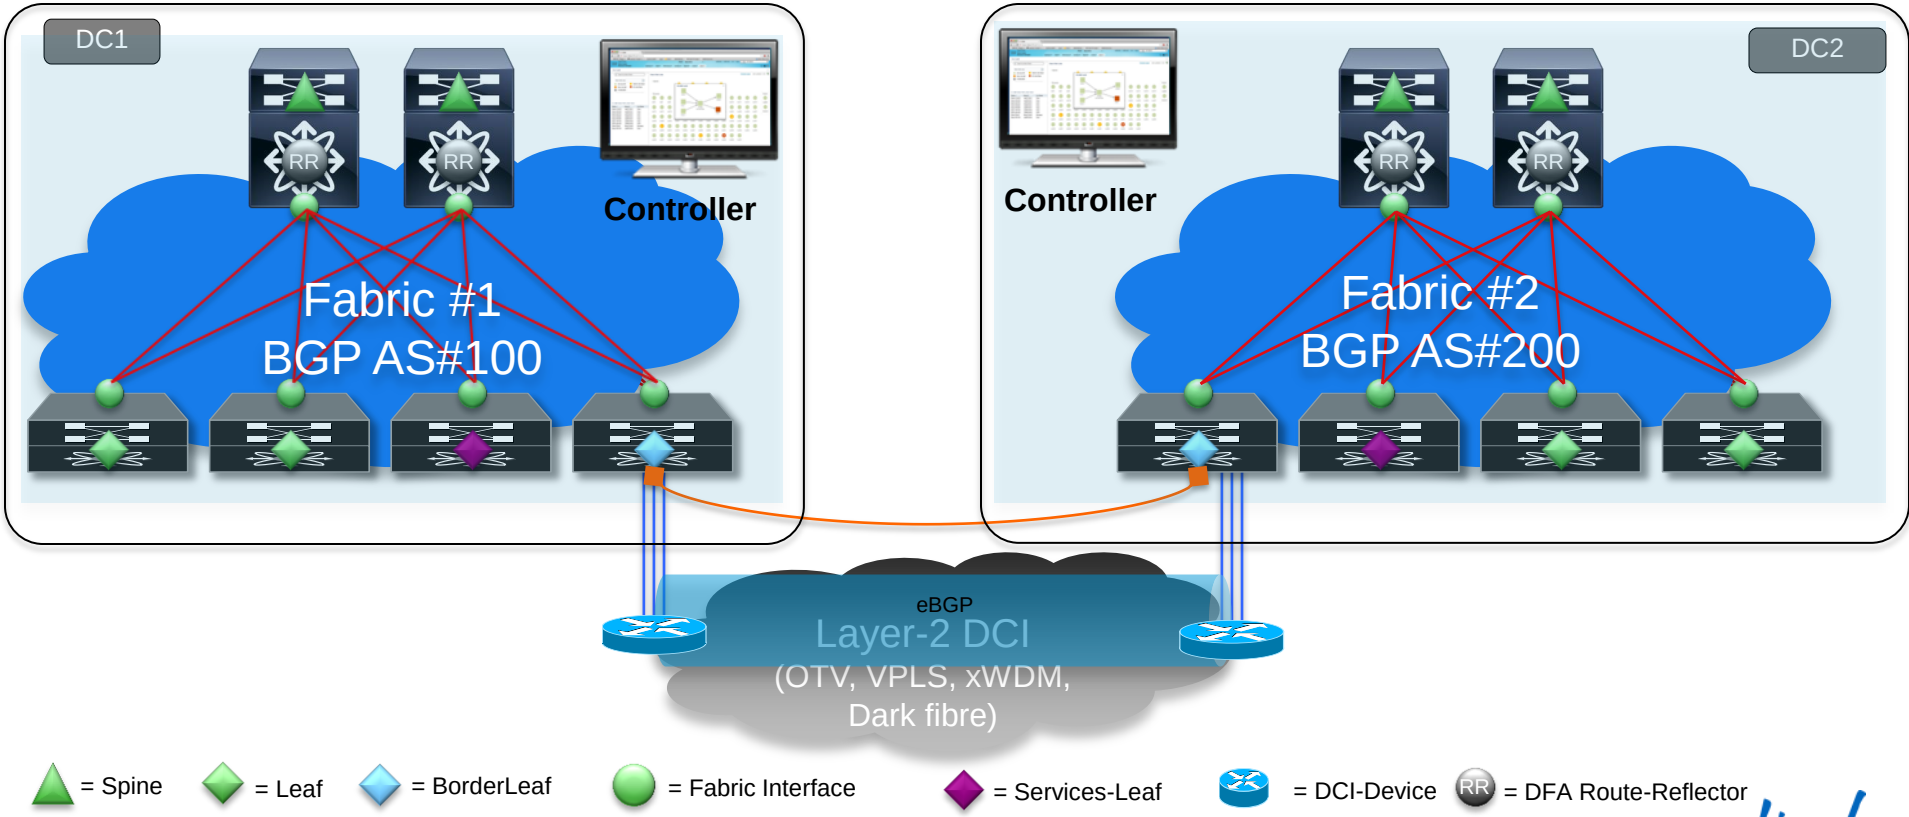

IP Forwarding between Fabrics across L2 based DCI

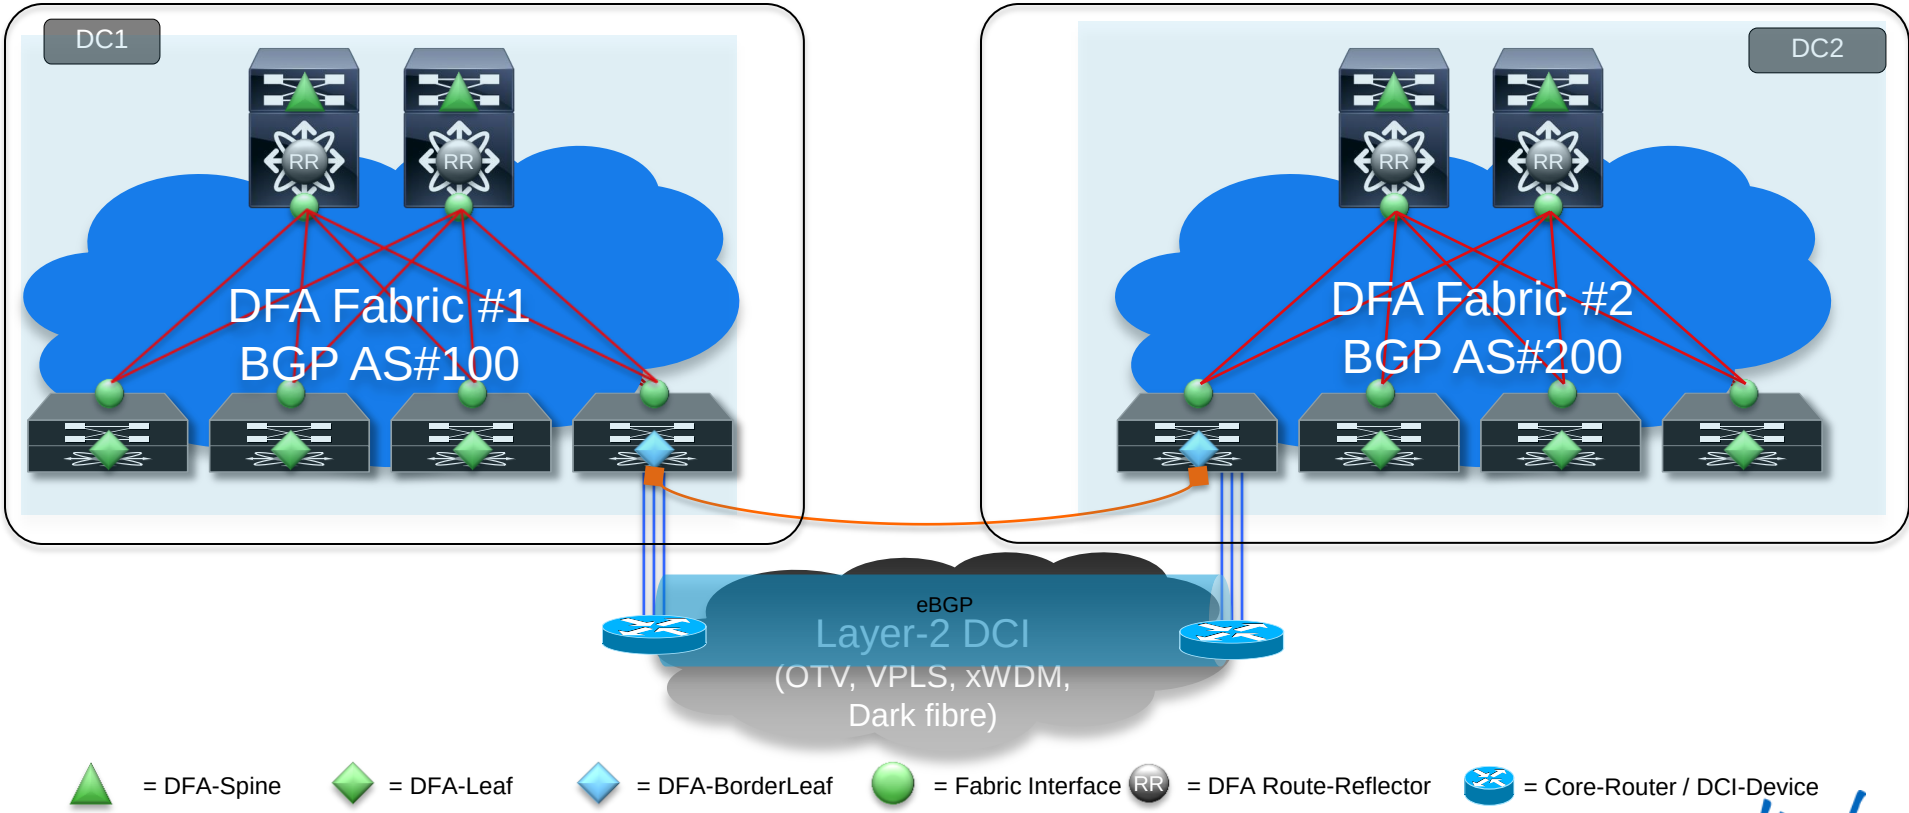

Nexus 9000 IP Transport Split Fabric DCI

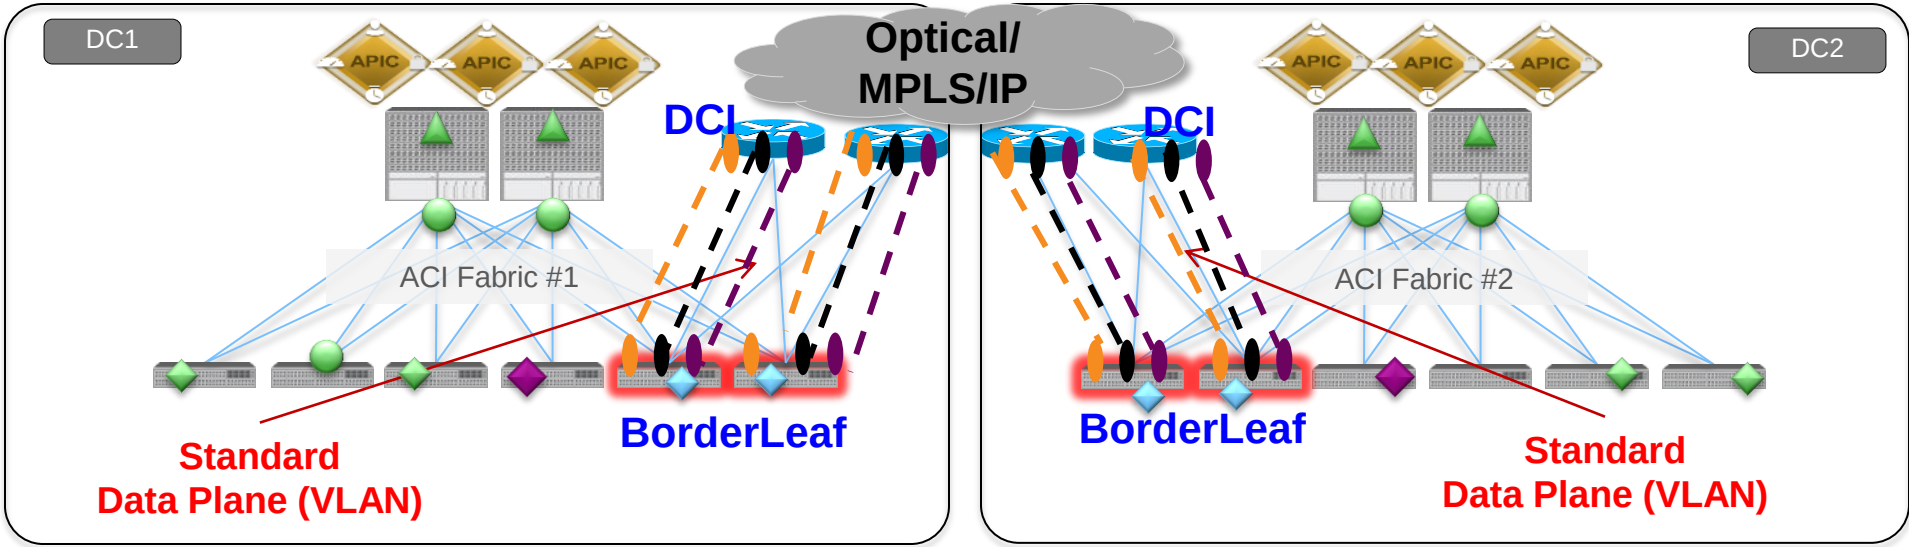

- Independent policy models in each iFabric
- Optical: vPC, FabricPath (P2P)
- Ethernet: EoMPLS, VPLS, PBB-EVPN, A-VPLS
- IP: OTV, EoMPLS/VPLSoGRE (option for encryption), L2TPv3
- Non-APIC Control Dependent L2 Provisioning

Re-Use Existing Nexus7000, ASR9000, ASR1000, Catalyst 6500

Active-Active iFabric Partial Mesh

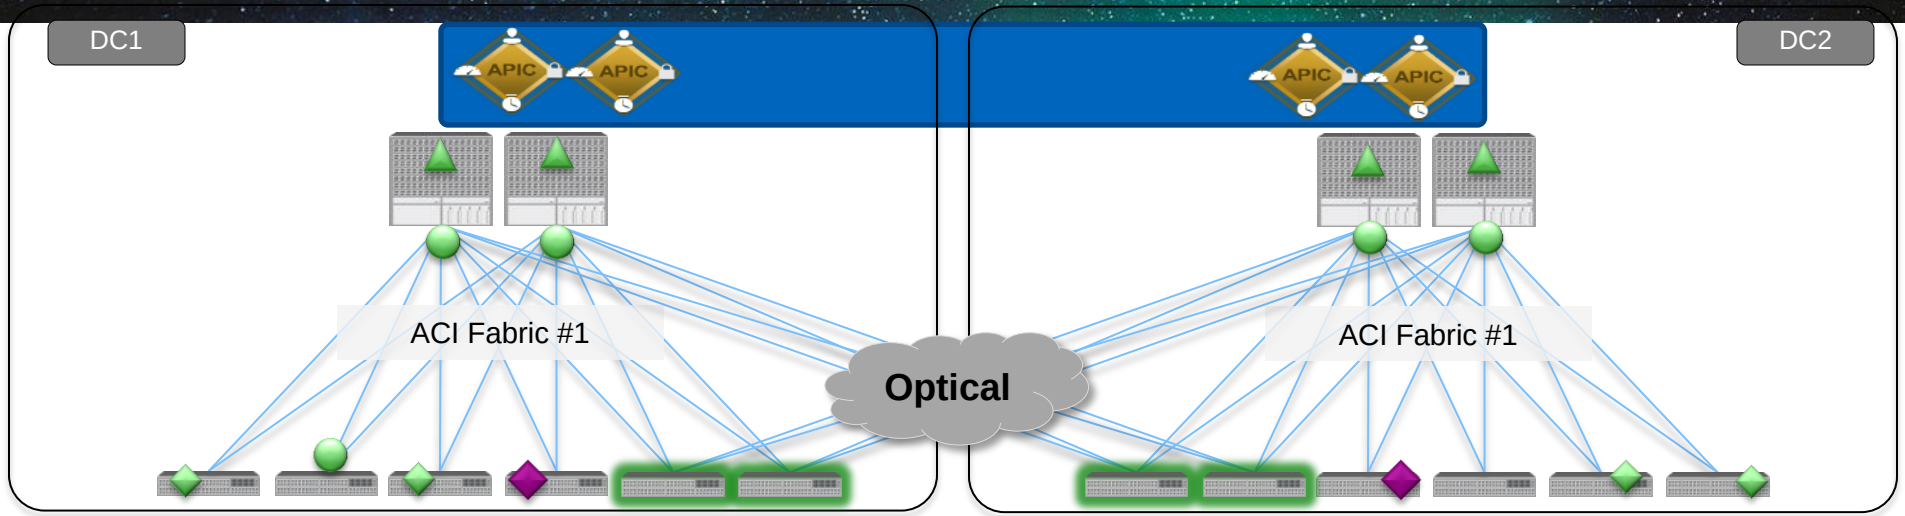

- Interconnect Leafs attach to all spines in remote data centre
- Distance is limited to 40G LR at 10km (Optics Supported)
- Traffic from indirect leafs bounces through Interconnect Leafs (Fast Reroute)
- Devices may be connected to the interconnect leafs (i.e. services, routers, storage, etc.)
- Single iFC Cluster

= Modular or Fixed Leaf

Summary

- Discussed different deployment options and transport options
- Tightly coupled Data Centre with FabricPath
- Spanning-tree isolation
- Traffic Optimisation Egress and Ingress Symmetry
- Encryption Solutions

References

- Cisco Validated Design – DCI Solutions

http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns743/ns749/landing_dci_mpls.html

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.2/collateral/nV_Cluster.pdf

- LISP Host Mobility CVD

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DCI/5.0/LISPmobility/DCI_LISP_Host_Mobility.html

- vPC DCI CVD:

http://www.cisco.com/en/US/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf

- 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
  - IETF LISP Working Group……..…http://tools.ietf.org/wg/lisp/




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


CISCO
TM