



EVPN Solutions/Applications for DC & SP Segments

Ali Sajassi

Distinguished Engineer, Cisco System

MPLS Japan 2015, Nov 9

EVPN History

2006

- Started the project OPEN at Cisco
- OPEN = Optimum Ethernet Network

2008

- Introduced to IETF as Routed-VPLS
- Merged with Juniper's MAC-VPN and was introduced an EVPN

2010

- Following drafts were introduced:
- EVPN
 - IPBB-EVPN
 - EVPN-VPWS
 - EVPN-Overlay
 - EVPN-ETREE

2012


- Following drafts were introduced
- EVPN IRB
 - EVPN DCI

2014/2015

- Enhancements
- Virtual ES
 - Optimized ingress replication
 - IGMP aggregation between PODs
 - mcast tunnels between DCs
 - Inter-AS for IRB
 - L3VPN multi-homing

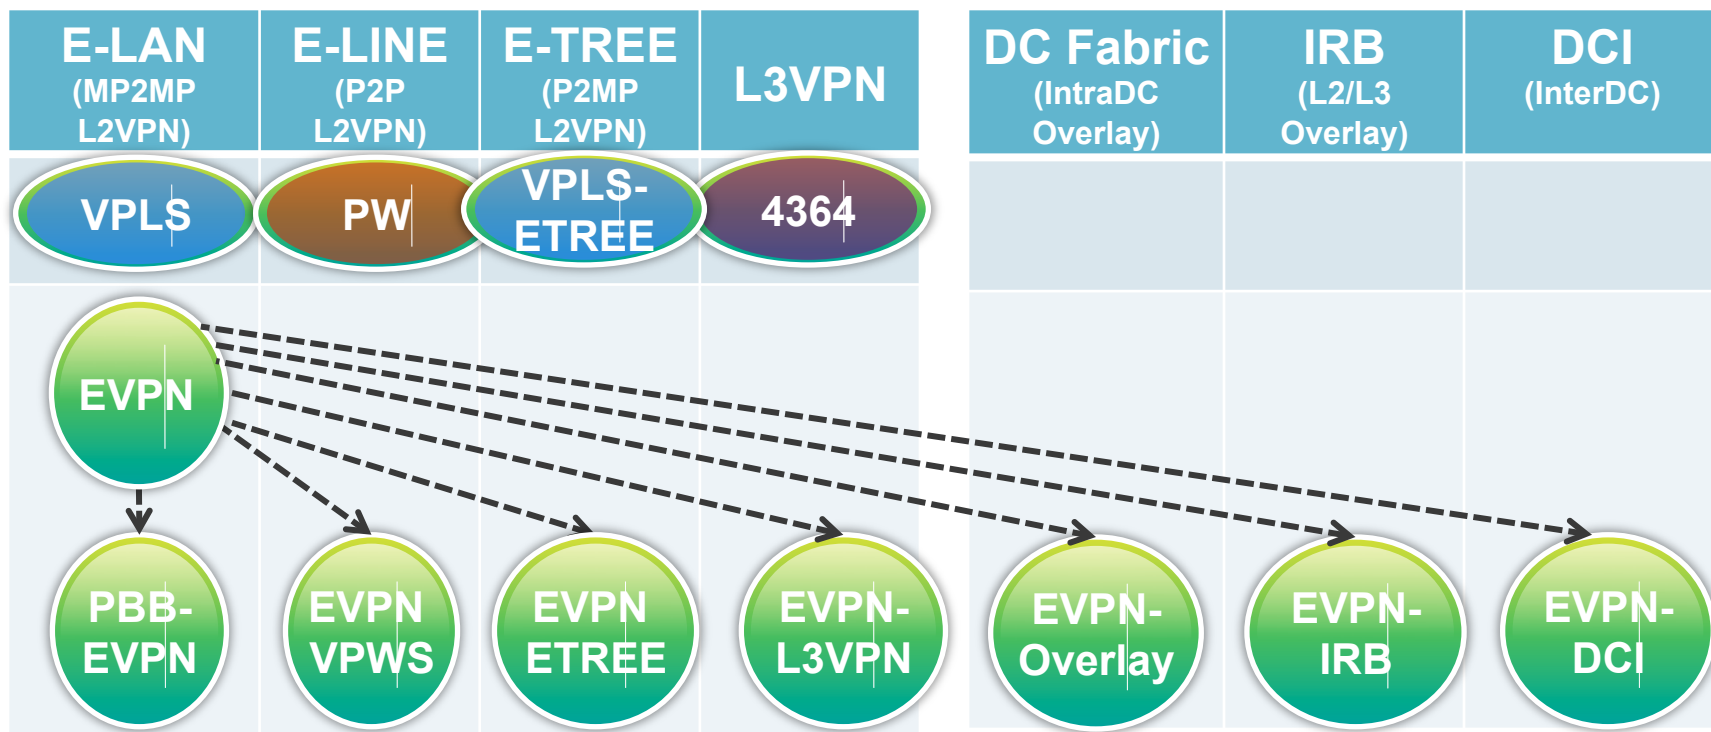
Challenges with Legacy VPN (VPLS)

- Next generation VPN solution need to have
 - Extensive multi-homing capabilities
 - Per-flow load-balancing in the core and access with DF election and loop prevention
 - Fast convergence upon failure
- Existing VPLS solutions do not offer an All-Active per-flow redundancy
- Looping of Traffic Flooded from PE
- Duplicate Frames from Floods from the Core
- MAC Flip-Flopping over Pseudowire
 - E.g. Port-Channel Load-Balancing does not produce a consistent hash-value for a frame with the same source MAC (e.g. non MAC based Hash-Schemes)



What is EVPN ?

EVPN is next generation all-in-one VPN solution



What's the big deal ?

Not only a single technology/solution does the job of many existing solutions but it does it better!

PBB-
EVPN

EVPN
VPWS

EVPN
ETREE

Service	Additional Capabilities
E-LAN	<ul style="list-style-type: none">• Provides All-Active multi-homing• Prevents loop for both all-active & single-active even in transient state• Ability to do per-flow LB & DF
E-Line	<ul style="list-style-type: none">• Both single-segment & multi-segment support• Discovery & signaling via single protocol – BGP• All-active & single-active redundancy support
E-TREE	<ul style="list-style-type: none">• Ingress filtering for traffic destined to egress leaf sites

What's the big deal ? – Cont.

New Applications & Solutions



Service	Features
DC-Fabric (IntraDC Overlay)	<ul style="list-style-type: none">• Geo-redundancy & VM mobility support• ARP suppression & ARP proxy• Support for different encaps such as VxLAN, NVGRE, MPLS, MPLSoUDP• Extensive multi-homing support• Extensive load-balancing capabilities
IRB	<ul style="list-style-type: none">• both L2 & L3 (on a per flow) for a given VLAN/tenant• L3 only mode when needed• Inherent support for anycast GW
DCI (InterDC)	<ul style="list-style-type: none">• Support for any kind of access technology such as EVPN Overlay (VxLAN), Native Ethernet, TRILL, 802.1Qbp, MPLS• Seamless interop with IP-VPN (RFC4364)

SDO Status



Service	IETF drafts
E-LAN	<ul style="list-style-type: none">• RFC7209 – draft-ietf-l2vpn-evpn-req• RFC 7432 - draft-ietf-l2vpn-evpn• RFC 7623 - draft-ietf-l2vpn-pbb-evpn• draft-ietf-bess-evpn-vpls-seamless-integ• draft-rbadan-bess-evpn-optimized-ir• drft-mohanty-bess-evpn-df-election• draft-sajassi-bess-evpn-igmp-mld-proxy• Draft-sajassi-bess-pbb-evpn-anycast-ip-tunnel
E-Line	<ul style="list-style-type: none">• draft-ietf-bess-evpn-vpws• draft-boutros-bess-evpn-service-edge-gw
E-TREE	<ul style="list-style-type: none">• draft-bess-l2vpn-evpn-etree

SDO Status – Cont.

**EVPN-
Overlay**

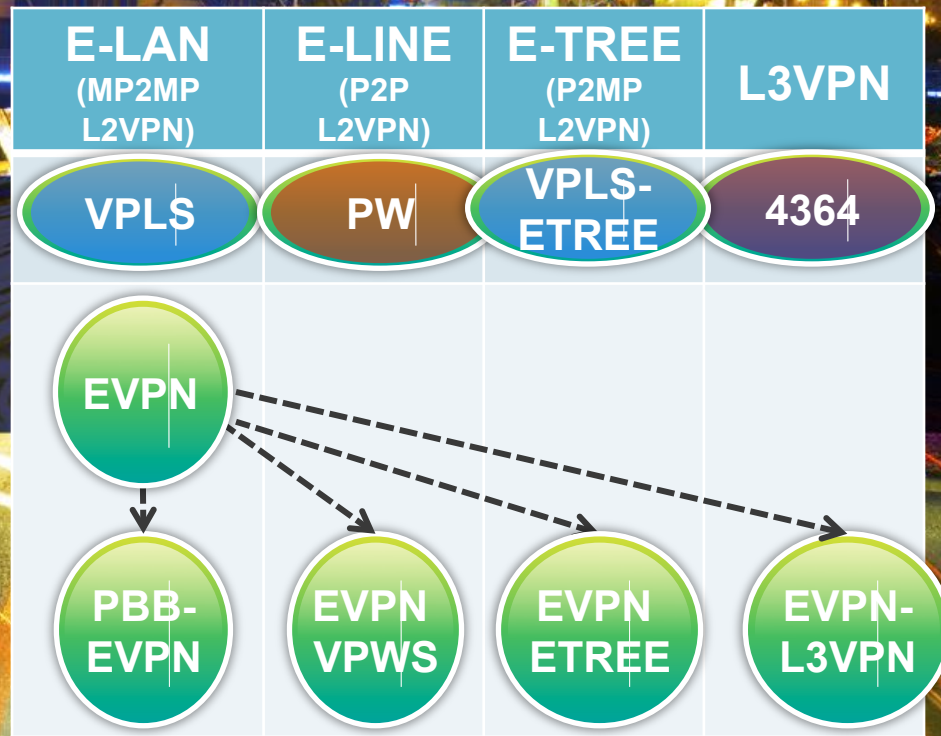
**EVPN-
IRB**

**EVPN-
DCI**

**EVPN-
L3VPN**

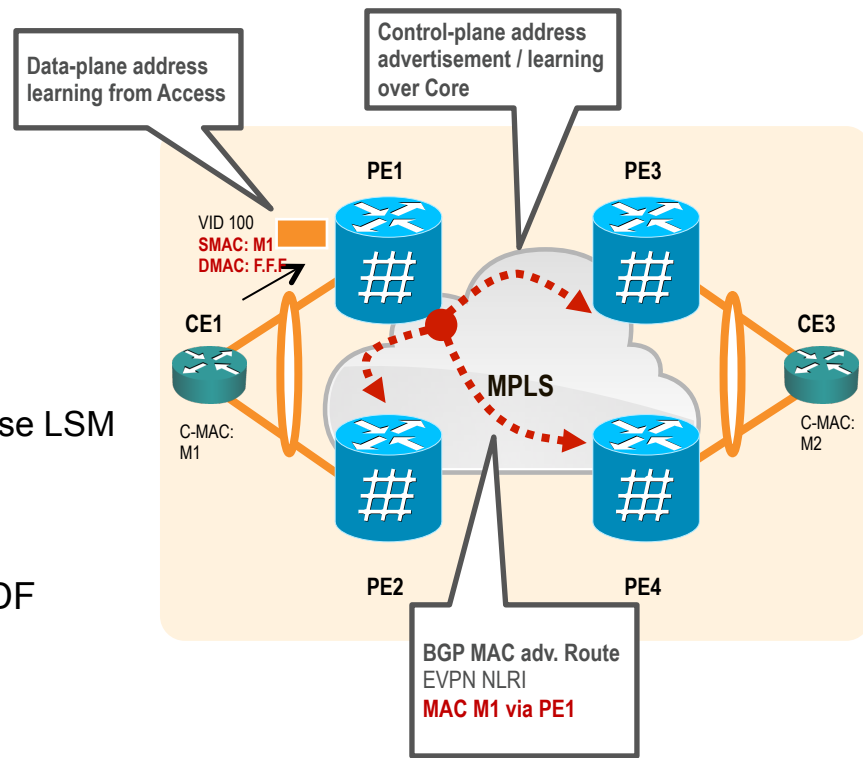
Service	IETF drafts
DC-Fabric (IntraDC Overlay)	<ul style="list-style-type: none">• draft-ietf-bess-evpn-overlay
IRB	<ul style="list-style-type: none">• draft-bess-l2vpn-evpn-inter-subnet-forwarding• draft-bess-l2vpn-evpn-prefix-advertisement
DCI (InterDC)	<ul style="list-style-type: none">• draft-sajassi-l2vpn-evpn-ipvpn-interop• draft-ietf-l2vpn-trill-evpn• Draft-ietf-l2vpn-spb-evpn• draft-boutros-l2vpn-vxlan-evpn
L3VPN	<ul style="list-style-type: none">• draft-sajassi-evpn-l3vpn-multihoming

SP Applications



EVPN - Highlights

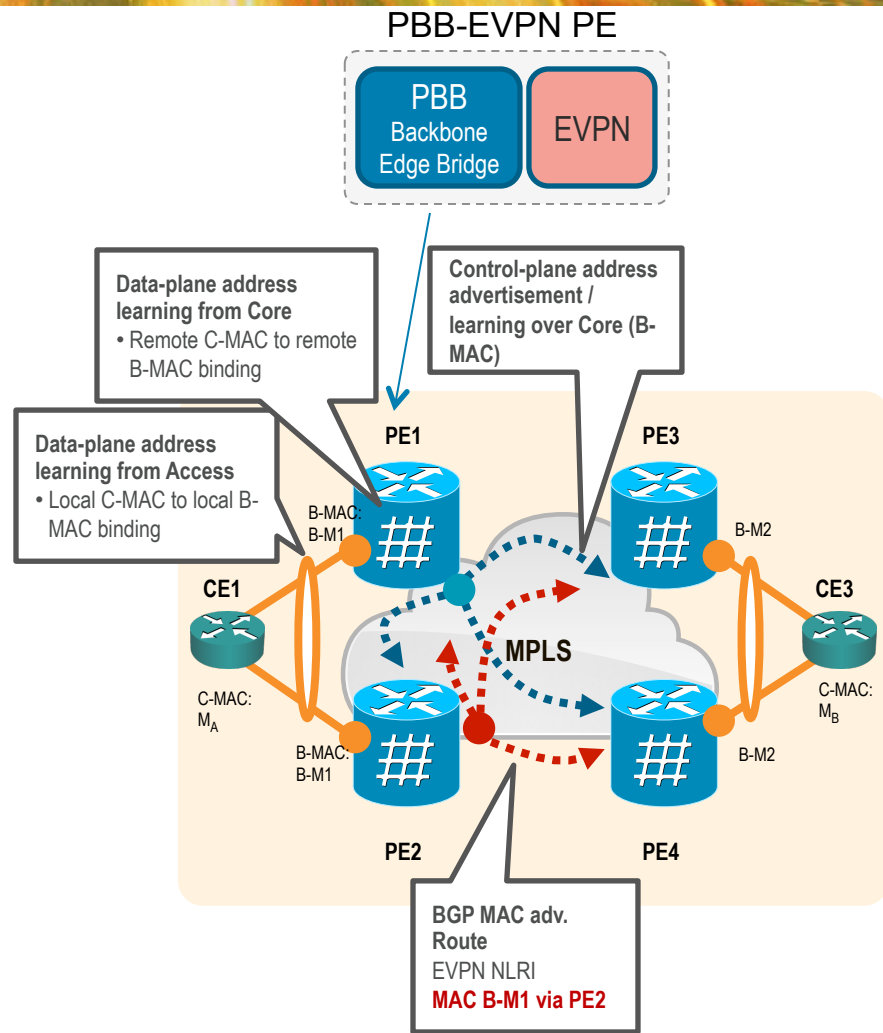
- Next generation solution for Ethernet multipoint (E-LAN) services
- **PEs run Multi-Protocol BGP to advertise & learn Customer MAC addresses (C-MACs) over Core**
 - Same operational principles of L3VPN
- **Learning on PE Access Circuits via data-plane transparent learning**
- No pseudowire full-mesh required
 - Unicast: use MP2P tunnels
 - Multicast: use ingress replication over MP2P tunnels or use LSM
- Provides
 - Extensive multi-homing capabilities
 - Per-flow load-balancing in the core and access with DF election and loop prevention
 - Fast convergence upon failure
- **Multi-vendor Solution – RFC 7432**



PBB -EVPN


Highlights

- Next generation solution for **Ethernet multipoint (E-LAN)** services by combining Provider Backbone Bridging (PBB - IEEE 802.1ah) and Ethernet VPN
- Data-plane learning of local C-MACs and remote C-MAC to B-MAC binding
- **PEs run Multi-Protocol BGP to advertise local Backbone MAC addresses (B-MACs) & learn remote B-MACs**
 - Takes advantage of PBB encapsulation to simplify BGP control plane operation – faster convergence
 - Lowers BGP resource usage (CPU, memory) on deployed infrastructure (PEs and RRs)
 - Scales for very large number of MACs
- Multi-vendor solution – RFC 7623



When to use PBB-EVPN?

- For SP & DCI applications where E2E L2 services needed
- Lower control-plane overhead than EVPN alone
 - PBB-EVPN uses only a sub-set of EVPN routes
 - Simpler and Faster failure convergence for all-active multi-homing scenarios
 - Faster MAC move convergence handled in data-plane
- Lower control-plane scale requirements than EVPN alone
 - BGP MAC advertisements for smaller Backbone MAC (B-MAC) address space
 - Requires less resources (CPU, memory) on deployed infrastructure (PEs / RRs)



Comparison of L2VPN Solutions

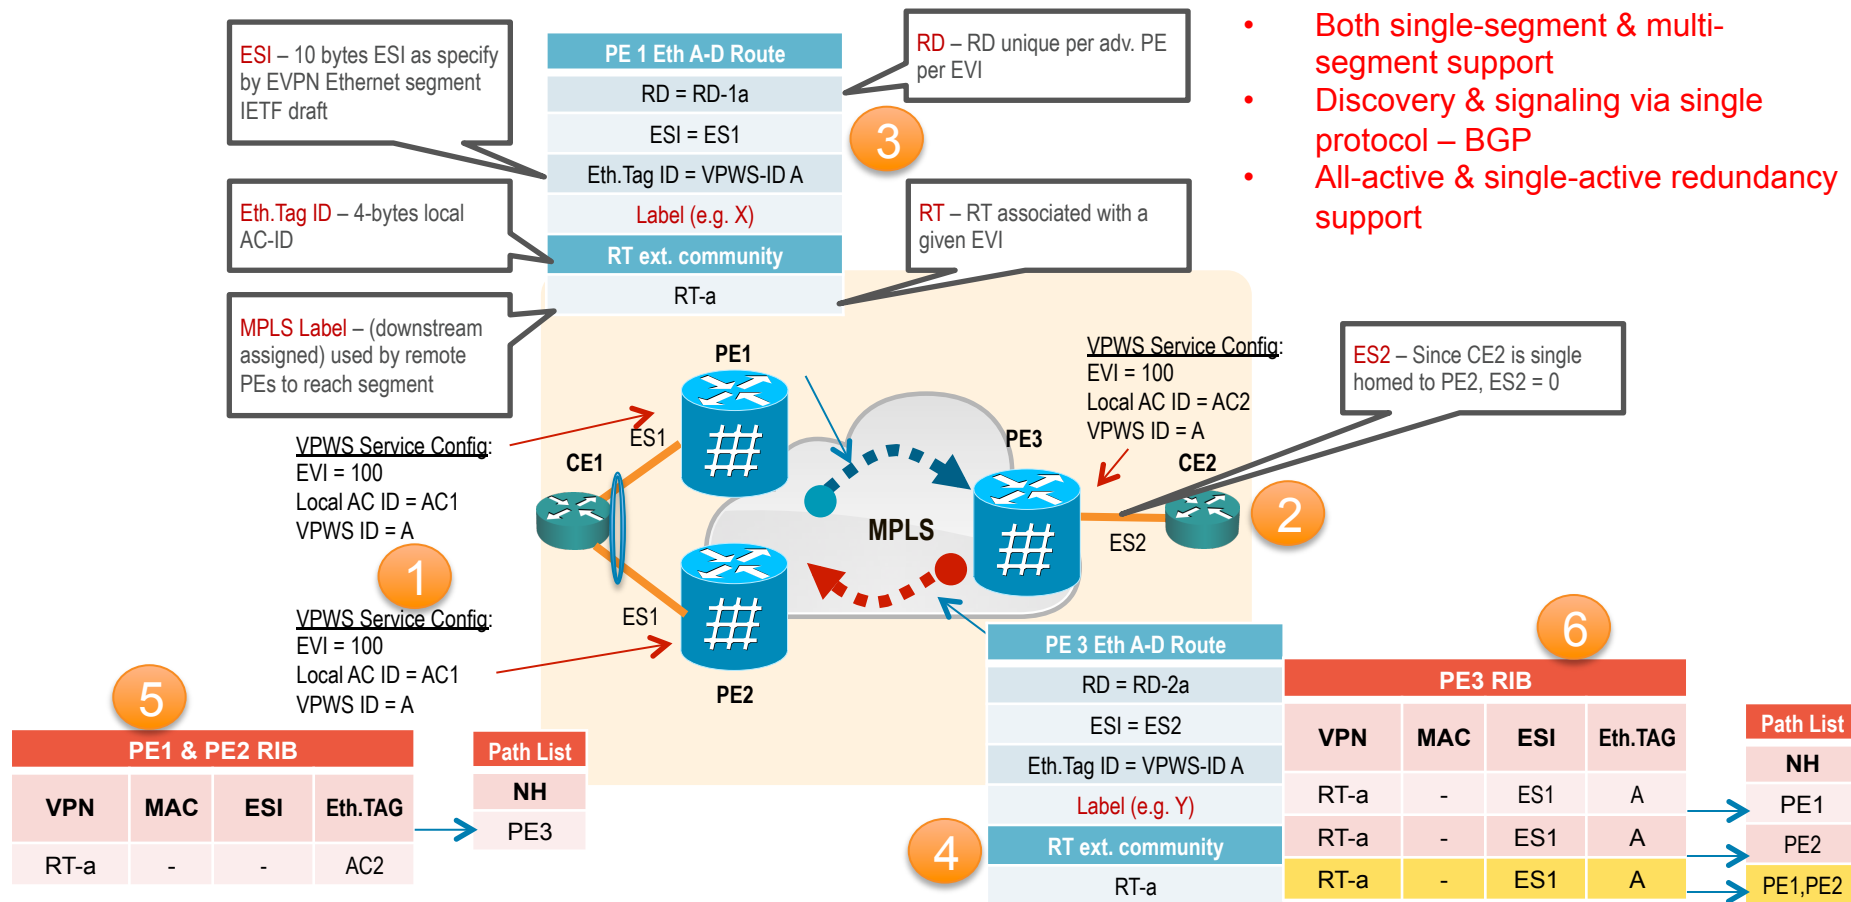
Requirement	VPLS	PBB-VPLS	EVPN	PBB-EVPN
Multi-Homing with All-Active Forwarding				
Service Based Load-balancing CE-to-PE	✓	✓	✓	✓
Flow Based Load-balancing CE-to-PE	✗	✗	✓	✓
Flow Based Load-balancing PE-to-PE	✗	✗	✓	✓
Flow Based Multi-Pathing in the Core	✓	✓	✓	✓
Provisioning Simplicity				
Core Auto-Discovery	✓	✓	✓	✓
Access Auto-Sensing	✗	✗	✓	✓
Redundancy Group Auto-Discovery	✗	✗	✓	✓
Automatic Designated Forwarder election and Service Carving	✗	✗	✓	✓
Service Interfaces				
Port-Based / VLAN-based / VLAN Bundling	✓	✓	✓	✓
VLAN-aware Bundling	✗	✗	✓	✓
Multi-Destination Traffic Forwarding				
Ingress Replication	✓	✓	✓	✓
LSM with P2MP Tree	✓	✓	✓	✓
LSM with MP2MP Tree	✗	✗	✓	✓

Comparison of L2VPN Solutions (cont.)


Requirement	VPLS	PBB-VPLS	EVPN	PBB-EVPN
Fast Convergence				
CE-PE Link Failures / PE Node Failures	✓	✓	✓	✓
MAC Mobility	✓	✓	✓	✓
CE-PE Link Failures with Local Repair	✗	✗	✓	✓
MAC Scalability				
Scale to Millions of C-MAC Addresses	✗	✓	✗	✓
Confinement of C-MAC entries to PE with active flows	✓	✓	✗	✓
MAC Summarization	✗	✗	✓	✓
MAC Summarization co-existence with C-MAC Mobility	✗	✗	✗	✓
Flexible VPN Policies				
Per C-MAC Forwarding Control Policies	✗	✗	✓	✗
Per-Segment Forwarding Control Policies	✗	✗	✓	✓

EVPN VPWS: All-Active Operation

- Both single-segment & multi-segment support
- Discovery & signaling via single protocol – BGP
- All-active & single-active redundancy support

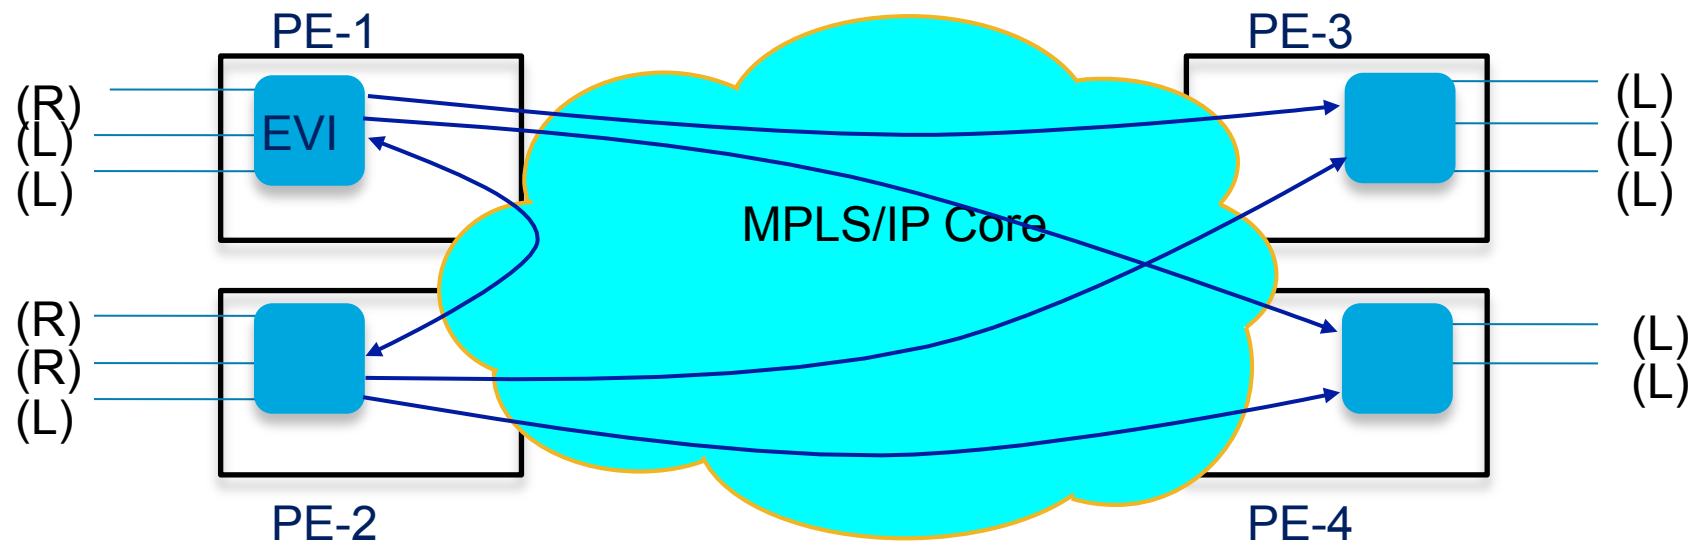


EVPN-VPWS: Service Edge GW



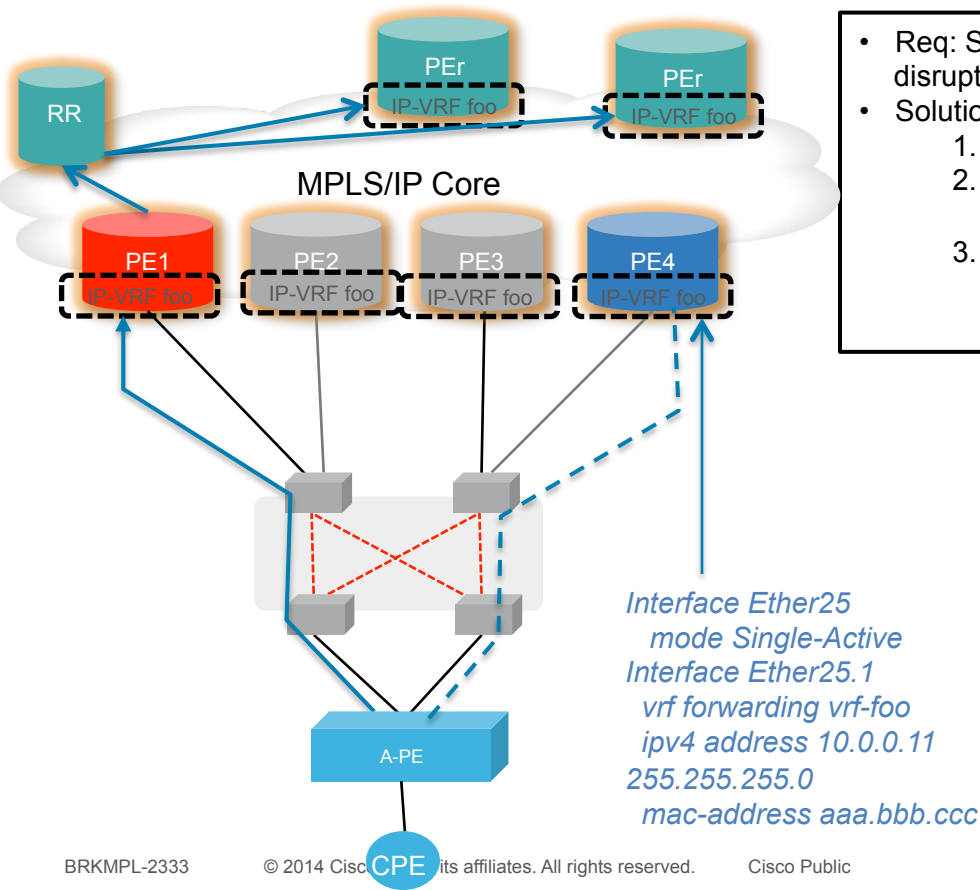
1. Discovery among S-PEs per EVPN/IP-VPN service needing VPWS
2. Origination of single-side signaling by an A-PE for a given VPWS service (identified by the 24-bit Ethernet tag)
3. DF election by S-PEs for Active/Backup
4. Active & Backup S-PEs responding to the single-sided signaling of step (2)
5. Binding of two halves of the EVC (or IPVC) by the A-PE upon receiving the response in step (3) & setting up backup EVC

EVPN-ETREE



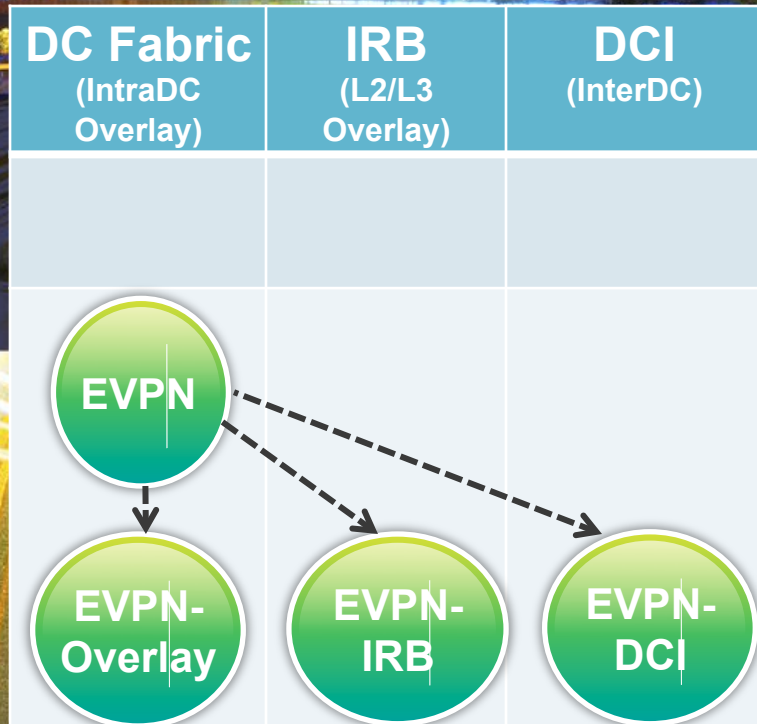
EVPN-ETREE provides ingress filtering for traffic destined to egress leaf sites !!

EVPN-L3VPN

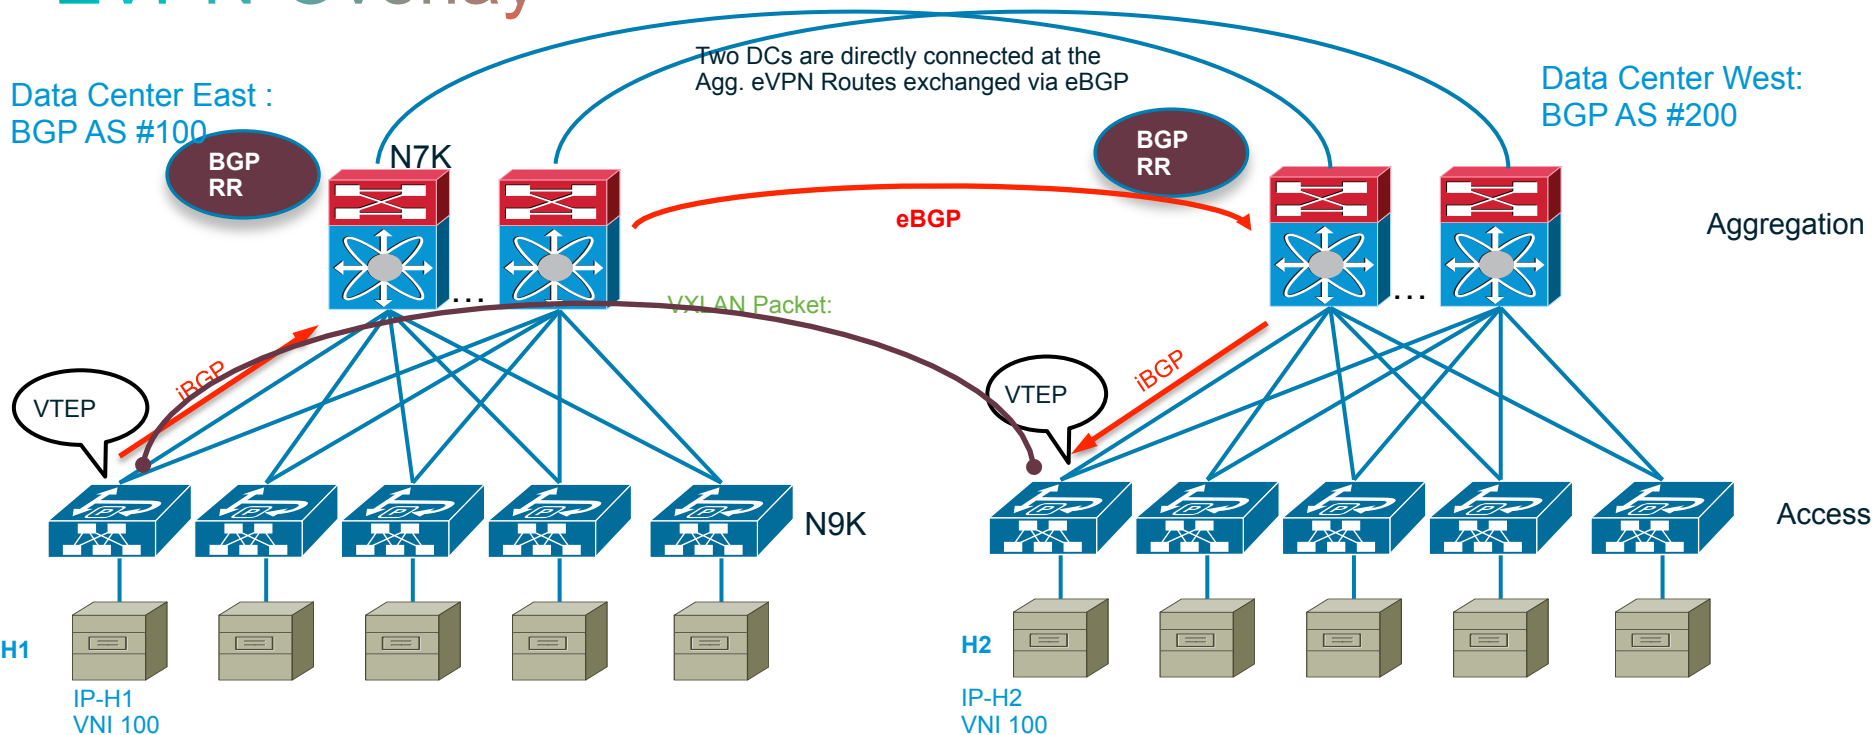


- Req: Single-BGP session from CPE with fast switchover and min. traffic disruption during failure
- Solution has three components:
 1. PIC functionality for fast switchover at PEr devices
 2. Synch up of ARP cache and VRF tables on all PEs in the redundancy group
 3. BGP GR with non-stop forwarding between CPE and PEs in the redundancy group

DC Applications



EVPN-Overlay




- Layer-2 multi-tenancy
- Geo-redundancy & seamless VM mobility support
- ARP suppression & ARP proxy
- Support large scale VMs & policy control

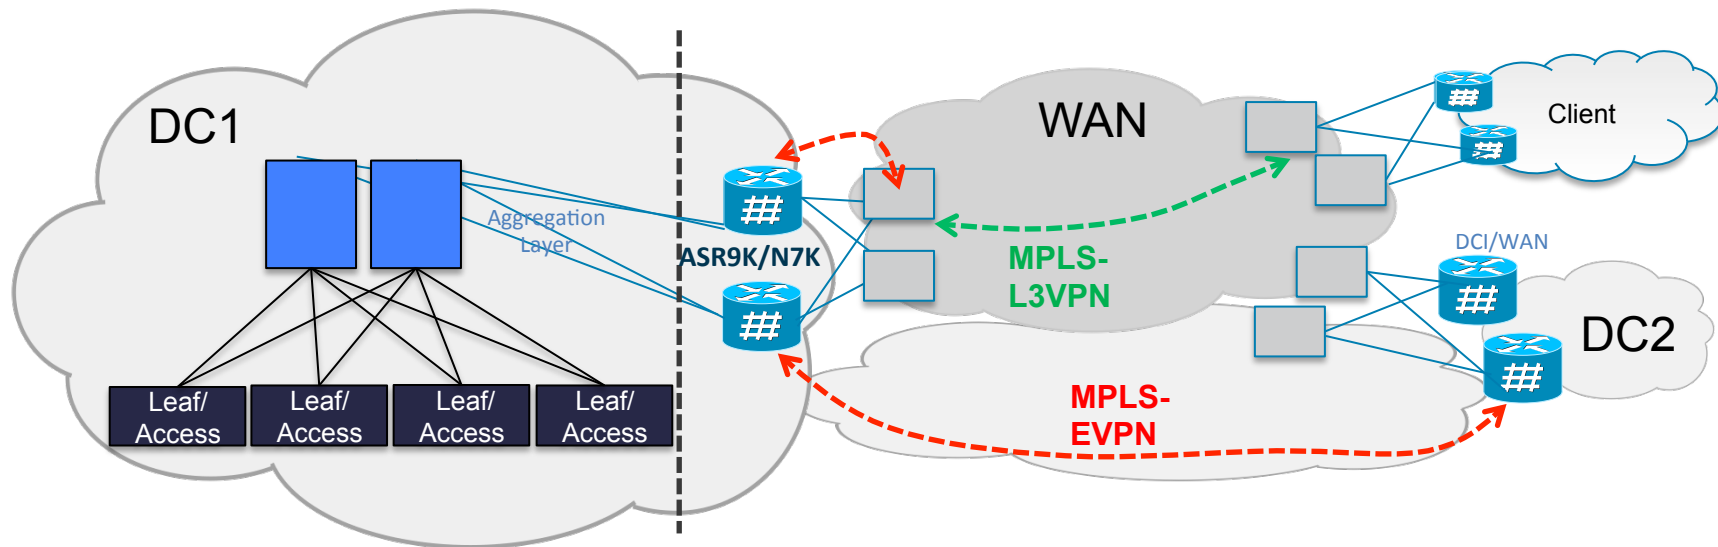
- Support for different encaps such as VxLAN, NVGRE, MPLS, MPLSoUDP
- Extensive multi-homing support
- Extensive load-balancing capabilities

EVPN-IRB

- Multi-tenancy for L2 & L3
- LB (on a per flow) for a given VLAN/tenant
- Inherent support for anycast GW
- ARP suppression
- L3 only mode when needed

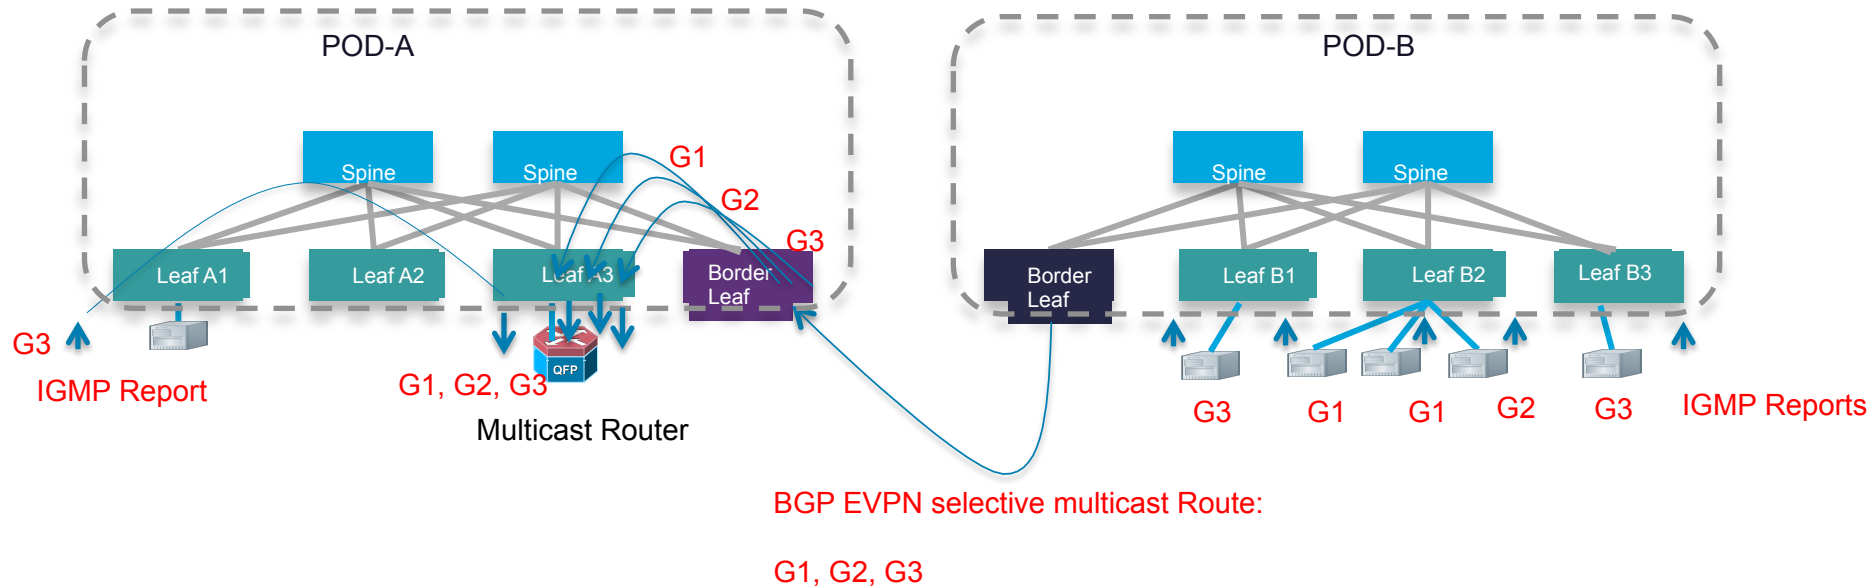


EVPN-DCI



- Seamless interop with IP-VPN (RFC4364)
- Support for any kind of access technology such as EVPN Overlay (VxLAN), Native Ethernet, TRILL, 802.1Qbp, MPLS

IGMP/MLD Proxy



Thank you.