



ETHERNET VPN (EVPN) NEXT-GENERATION VPN FOR ETHERNET SERVICES

Alastair JOHNSON (AJ)

alastair.johnson@alcatel-lucent.com

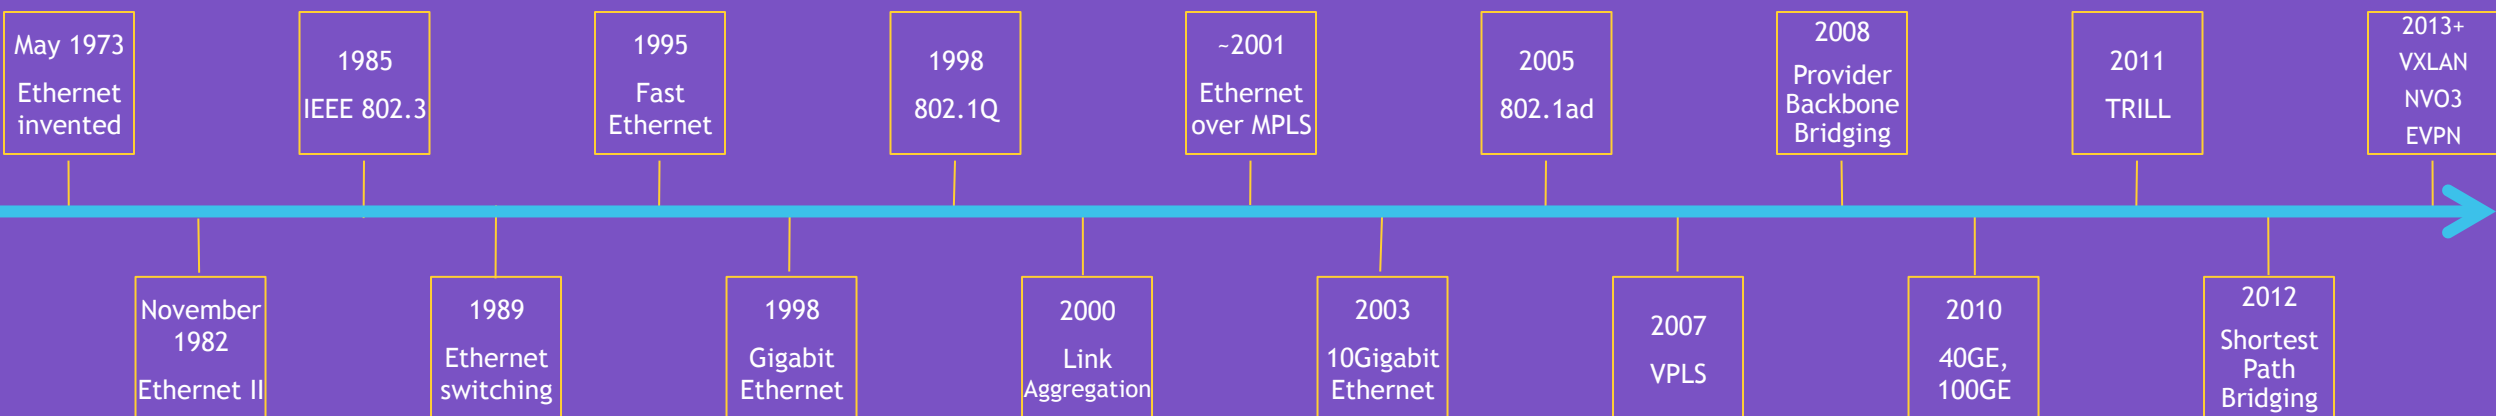
February 2014

AGENDA

1. EVPN Background and Motivation
2. EVPN Operations
3. EVPN Use Cases

ETHERNET CONTINUES TO EVOLVE

“The **widespread adoption of Ethernet L2VPN** services and the advent of **new applications** for the technology (e.g., data center interconnect) have culminated in a **new set of requirements** that are **not readily addressable** by the current VPLS solution”



ETHERNET CONTINUES TO EVOLVE

- EVPN introduces a new model for delivery of L2VPN services
 - Inheriting a decade of VPLS operational experience in production networks
 - Incorporating flexibility for service delivery over L3 networks
 - Abstracting and separating the control and data planes
- Allows operators to meet emerging needs in their networks
 - Datacenter Interconnect
 - Cloud and virtualization services and connectivity management
 - Combination services
 - Use of network overlay and underlay technologies to simplify topology, remove protocols from the network

WHY ANOTHER VPN TECHNOLOGY?

- The popularity of Ethernet L2VPN services and new applications such as DCI (data center interconnect) are creating new requirements for VPNs
 - Multihoming with all-active forwarding
 - Optimizing the delivery of multidestination frames (BUM)
 - Easier provisioning of services
 - L3VPN-like operation for scalability and control
 - Delivering Layer 2 and Layer 3 services over the same interface
- VPLS and L3VPNs are proven technology but cannot meet all of these requirements
- **EVPN supports integrated routing and bridging VPN solutions with MAC/IP mobility over the same VLAN**
 - **More efficient hybrid service delivery over a single VLAN**
 - **Simpler provisioning and management with a single VPN technology**
 - **Multiple data plane encapsulation choices**

KEY OPERATIONAL BENEFITS OF EVPN

HYBRID SERVICE

- L2 and L3 service in a single VPN
- IP-VPN like operation

FLEXIBILITY

- Supports multiple data plane models
- VXLAN encapsulation encourages simple IP network deployment

NETWORK EFFICIENCY

- Optimized multi-destination support
- Load balancing between PEs

OPTIMIZE BCAST/UNKNOWN

- ARP/ND proxy allows PEs to respond to ARP/ND requests on behalf of clients
- Flood-and-learn FDB model changes to pre-sigaled FDB

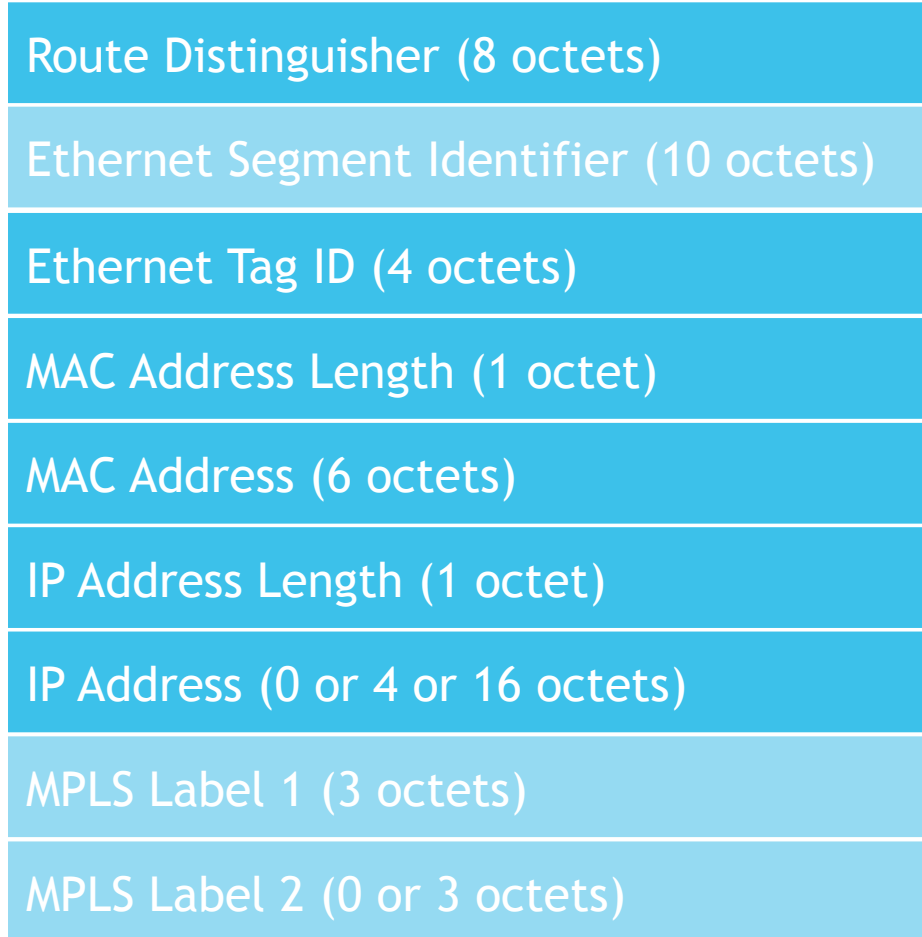
EVPN STATUS

- Hot new technology in the IETF L2VPN WG
- Many mature base I-Ds and new I-Ds
 - draft-ietf-l2vpn-evpn base specification: version -05 expected to go to last call
 - draft-ietf-l2vpn-evpn-req: no more changes expected
 - draft-ietf-l2vpn-pbb-evpn: no more changes expected
- Diverse authors on requirements and base specification
 - Vendors: Alcatel-Lucent, Cisco, Juniper
 - Network operators: Arktan, AT&T, Bloomberg, Verizon
- Shipping implementations
 - Alcatel-Lucent: EVPN, PBB-EVPN, EVPN-VXLAN
 - Cisco: EVPN, PBB-EVPN
 - Juniper: EVPN

draft-allan-l2vpn-mldp-evpn
draft-boutros-l2vpn-evpn-vpws
draft-boutros-l2vpn-vxlan-evpn
draft-ietf-l2vpn-evpn
draft-ietf-l2vpn-evpn-req
draft-ietf-l2vpn-pbb-evpn
draft-ietf-l2vpn-spbm-evpn
draft-ietf-l2vpn-trill-evpn
draft-jain-l2vpn-evpn-lsp-ping
draft-li-l2vpn-evpn-mcast-state-ad
draft-li-l2vpn-evpn-pe-ce
draft-li-l2vpn-segment-evpn
draft-rabadan-l2vpn-dci-evpn-overlay
draft-rabadan-l2vpn-evpn-prefix-advertisement
draft-rp-l2vpn-evpn-usage
draft-sajassi-l2vpn-evpn-etree
draft-sajassi-l2vpn-evpn-inter-subnet-forwarding
draft-sajassi-l2vpn-evpn-ipvpn-interop
draft-sajassi-l2vpn-evpn-vpls-integration
draft-salam-l2vpn-evpn-oam-req-frmwk
draft-sd-l2vpn-evpn-overlay
draft-vgovindan-l2vpn-evpn-bfd
draft-zhang-l2vpn-evpn-selective-mcast
draft-zheng-l2vpn-evpn-pm-framework

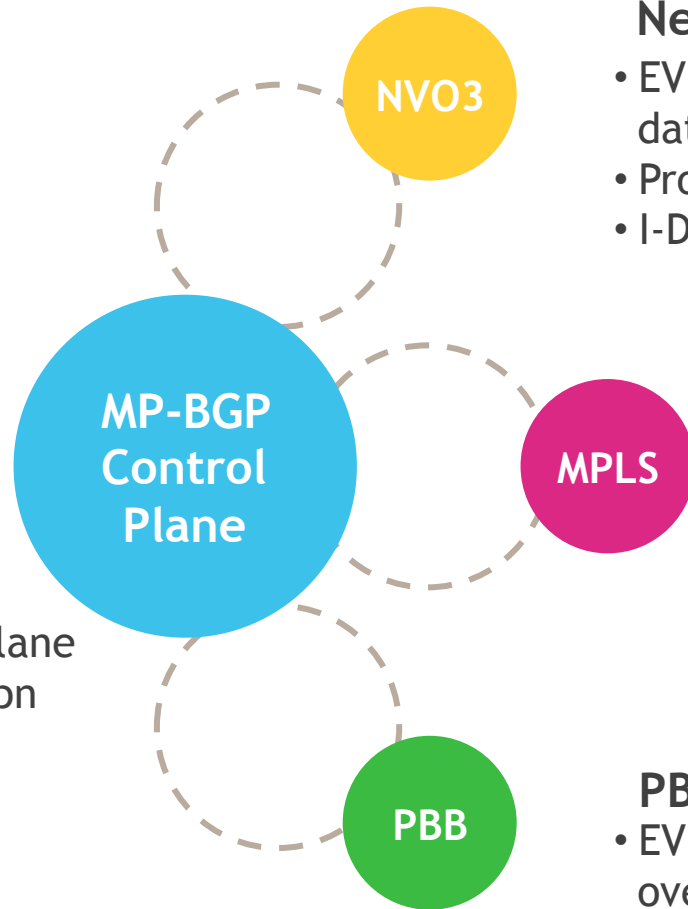
EVPN CONTROL PLANE LEARNING WITH MP-BGP

- BGP advertises MACs and IPs for next hop resolution with EVPN NLRI
 - AFI=25 (L2VPN) SAFI=70 (EVPN)
 - Fully supports IPv4 and IPv6 in the control and data plane
- Offers greater control over MAC learning
 - Who learns what
 - Ability to apply policies
- Maintains virtualization and isolation of EVPN instances
- Enables traffic load balancing for multihomed CEs with ECMP
- Brings inherent BGP scalability to MAC routes
- Even more scalability and hierarchy with route reflectors



MAC/IP Advertisement Route

EVPN DATA PLANES



BGP

- EVPN MP-BGP control plane
- I-D: draft-ietf-l2vpn-evpn

Network Virtualization Over L3 (NVO3)

- EVPN over NVO3 overlay tunnels (VXLAN, NVGRE, MPLSoGRE) for data center fabric encapsulations
- Provides Layer 2 and Layer 3 DCI
- I-D: draft-sd-l2vpn-evpn-overlay

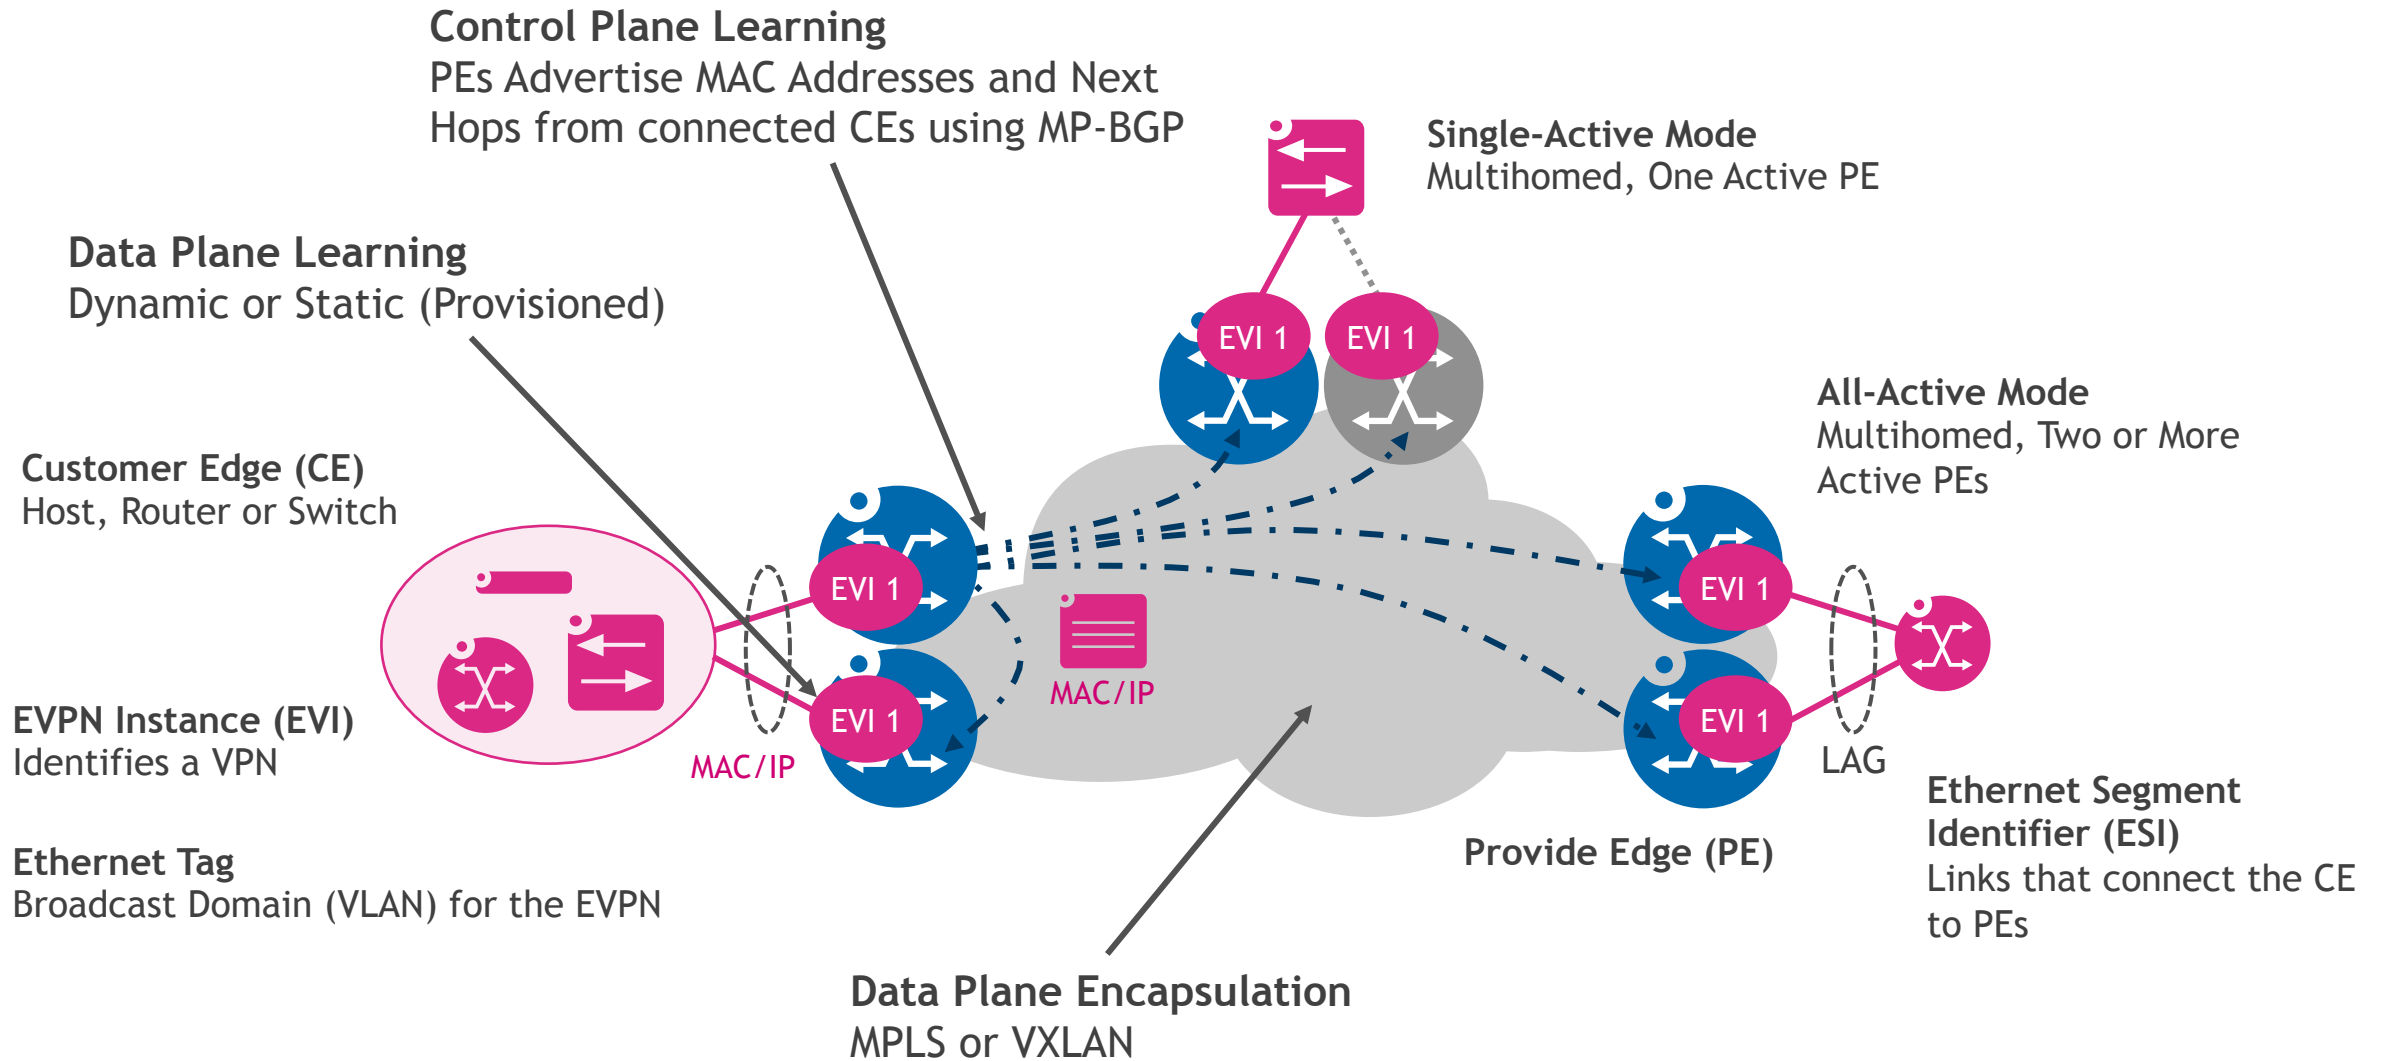
MPLS

- EVPN over MPLS for VLL, VPLS and E-Tree services
- All-active multihoming for VPWS
- RSVP-TE or LDP MPLS protocols
- I-D: draft-ietf-l2vpn-evpn

PBB

- EVPN with PBB PE functionality for scaling very large networks over MPLS
- All-active multihoming for PBB-VPLS
- I-D: draft-ietf-l2vpn-pbb-evpn

EVPN CONCEPTS OVERVIEW



EVPN SERVICES OVERVIEW

VLAN Based Service Interface

- Single broadcast domain (VLAN)
- 1:1 mapping between VLAN ID and EVI
- Single bridge domain for each EVI
- VLAN translation is allowed
- Ethernet tag in route set to 0




VLAN Bundle Service Interface

- Multiple broadcast domains (VLANs)
- N:1 mapping between VLAN ID and EVI
- Single bridge domain for each EVI
- MACs must be unique across VLANs
- VLAN translation is not allowed
- Ethernet tag in route set to 0



VLAN Aware Bundle Service Interface

- Multiple broadcast domains (VLANs)
- N:1 mapping between VLAN ID and EVI
- Multiple bridge domains, one for each VLAN
- VLAN translation is allowed
- Ethernet tag in route set to configured tag



AGENDA


1. EVPN Background and Motivation
2. EVPN Operations
3. EVPN Use Cases

EVPN OPERATION

- Key features control plane features
 - All-Active Multihoming and Designated Forwarder Election
 - All-Active Multihoming and Split Horizon
 - ARP/ND Proxy and Unknown Unicast Flooding Suppression
 - Aliasing
 - MAC Mobility
 - Default Gateway Inter-Subnet Forwarding
 - MAC Mass-Withdraw
- Data planes
 - MPLS
 - PBB-EVPN
 - EVPN-VXLAN

EVPN OPERATION

ALL-ACTIVE MULTIHOMING AND DESIGNATED FORWARDER ELECTION




- Avoids duplicate BUM flooding to all-active CEs
- PEs connected to multihomed CEs discover each other through auto discovery routes
- Elect a designated forwarder (DF) responsible for BUM flooding to the Ethernet segment
- Non-DF PEs block BUM flooding to the CE

- Flexible DF election and functionality
 - Same DF for all ESIs
 - Different DF per ESI
- Unicast still follows all-active paths

EVPN OPERATION

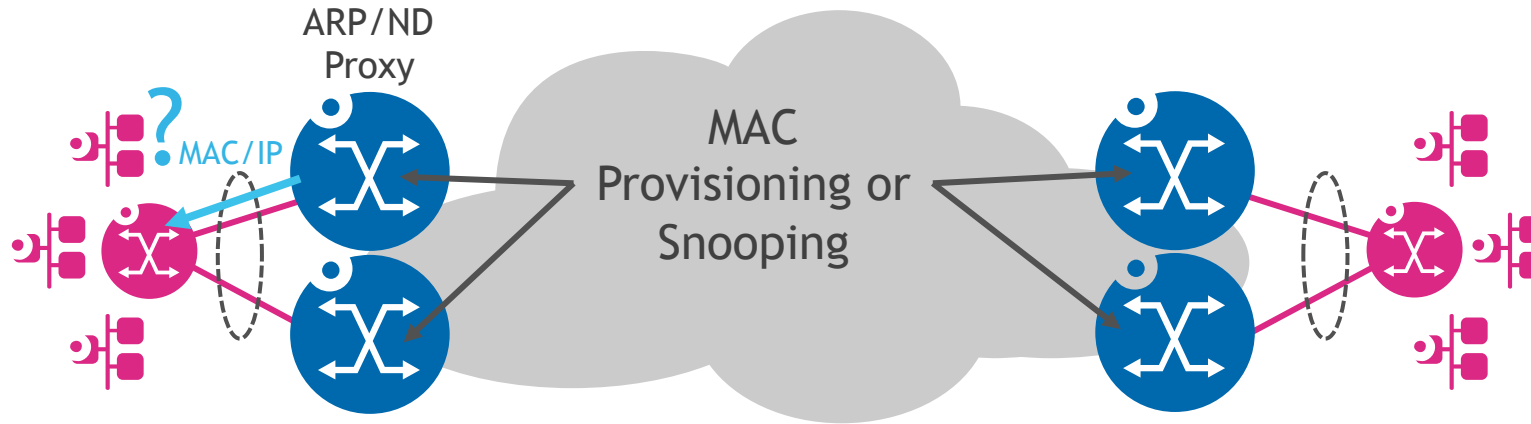
ALL-ACTIVE MULTIHOMING AND SPLIT HORIZON



- Ensures that BUM traffic from an ESI is not replicated back to the same ESI to an all-active CE
- PE advertises a split horizon label for each all-active Ethernet segment
- When an ingress PE floods BUM traffic, it pushes the split horizon label to identify the source Ethernet segment
- Egress PEs use this label for split horizon filtering and drop packets with the label destined to the Ethernet segment
- Implicit split horizon for core, since PEs won't flood received BUM traffic back into core

EVPN OPERATION

ARP/ND PROXY AND UNKNOWN UNICAST FLOODING SUPPRESSION



- ARP/ND is a security issue and a scalability issue in large networks
 - Spoofing and untrusted sources
 - Unknown unicast traffic levels, especially in large data center and IXP networks
 - We really don't need it anymore in orchestrated or provisioned networks where all MACs are known

- EVPN can reduce or suppress unknown unicast flooding since all active MACs and IPs are advertised by PEs
 - PEs proxy ARP/ND based on MAC route table to CEs
 - Snooping optimizes and reduces unknown unicast flooding
 - Provisioning MAC addresses can eliminate unknown unicast flooding entirely

EVPN OPERATION

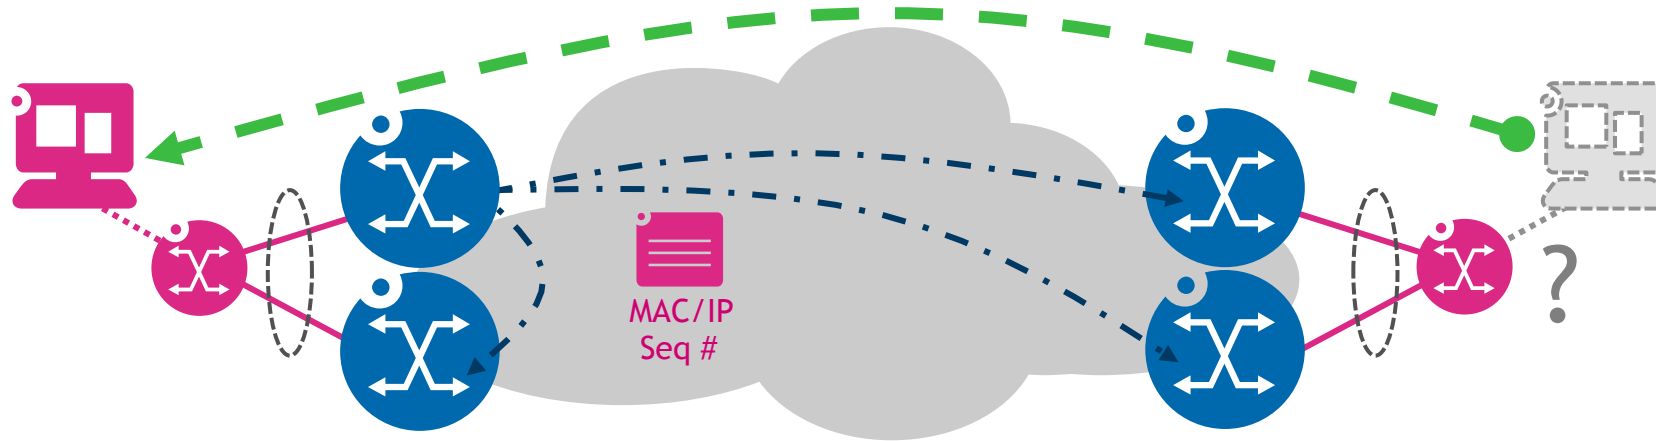
ALIASING



- Provides load-balancing to all-active CE when the MAC address is only learned by one PE
- PEs advertise the ESI in MAC routes with all-active mode
- Remote PEs can load-balance traffic across all PEs advertising the same ESI
- Can also be used for a backup path in single-active mode with standby link

EVPN OPERATION

MAC MOBILITY



- MAC addresses may move between ESIs
- If local learning is used in the data plane, the PE may not detect that a MAC address has moved and won't send a withdraw for it
- New PE sends a new MAC route
- Now there are two routes for the MAC address: an old wrong one and a new correct one
- Each MAC is advertised with a MAC mobility sequence number in an extended community with the MAC route
- PE selects the MAC route with the highest sequence number
- Triggers withdraw from PE advertising MAC route with the lower sequence number

EVPN OPERATION


DEFAULT GATEWAY INTER-SUBNET FORWARDING



- EVPN supports inter-subnet forwarding when IP routing is required
- No additional separate L3VPN functionality, uses EVPN default gateway
- One or more PEs is configured as the default gateway, 0.0.0.0 or :: MAC route is advertised with default gateway extended community
- Local PEs respond to ARP/ND requests for default gateway
- Enables efficient routing at local PE
- Avoids tromboning traffic across remote PEs to be routed MAC move if all default gateways use the same MAC address

EVPN OPERATION


MAC MASS-WITHDRAW



- Provides rapid convergence when a link failure affects many MAC addresses
- PEs advertise two routes
 - MAC/IP address and its ESI
 - Connectivity to ESIs
- If a failure affects an ESI, the PE simply withdraws the route for the ESI
- Remote PEs remove failed PE from the path for all MAC addresses associated with an ESI
- Functions as a MAC mass-withdraw and speeds convergence during link failures
- No need to wait for individual MAC addresses to be withdrawn

EVPN VIRTUAL EXTENSIBLE LAN (VXLAN) DATA PLANE

DRAFT-SD-L2VPN-EVPN-OVERLAY EVPN-VXLAN

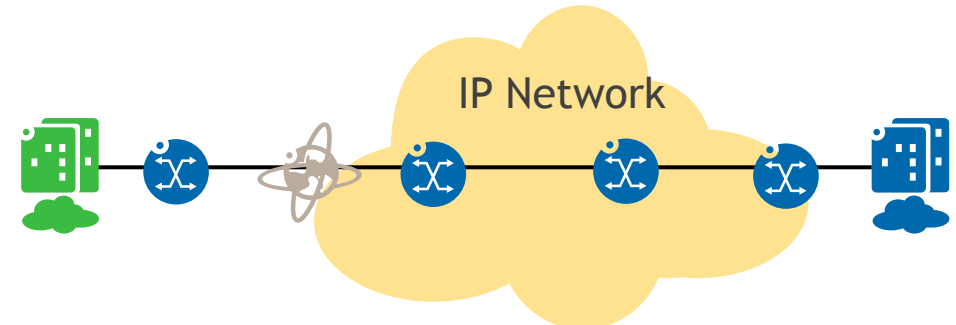




- EVPN-VXLAN uses EVPN over a VXLAN data plane
 - VXLAN is typically used for data center extension over WAN
 - Can also be used as an overlay in any IP network for IP/Ethernet services
 - Useful when MPLS is unavailable or unwanted
 - Alternative to MPLSoGRE

- VXLAN provides the Layer 2 overlay over IP
 - IP reachability is required between PEs
 - EVPN uses BGP control plane for MAC route advertisements
 - VXLAN data plane uses UDP to encapsulate the VXLAN header and Layer 2 frame
- Provides all the benefits of EVPN for DCI and virtualized networks

FLEXIBILITY OF VXLAN DATA PLANE


- VXLAN encapsulates Ethernet in IP
 - IPv4 or IPv6
 - UDP based, where the source port is a hash of MAC or IPs to provide load balancing entropy
 - 8 byte VXLAN header provides VXLAN Network Identifier (24 bits) and flags
 - Total encapsulation is ~50 bytes
- As VXLAN is routable, the underlay network may be any network, making use of existing resiliency and load balancing mechanisms
 - ECMP
 - IGPs/BGP
 - IP Fast Reroute
- VXLAN Tunnel Endpoints may exist on network equipment or computing infrastructure
 - Deliver a L2VPN straight to a hypervisor



AGENDA


1. EVPN Background and Motivation
2. EVPN Operations
3. EVPN Use Cases

INTERNET EXCHANGE POINT PEERING FABRIC



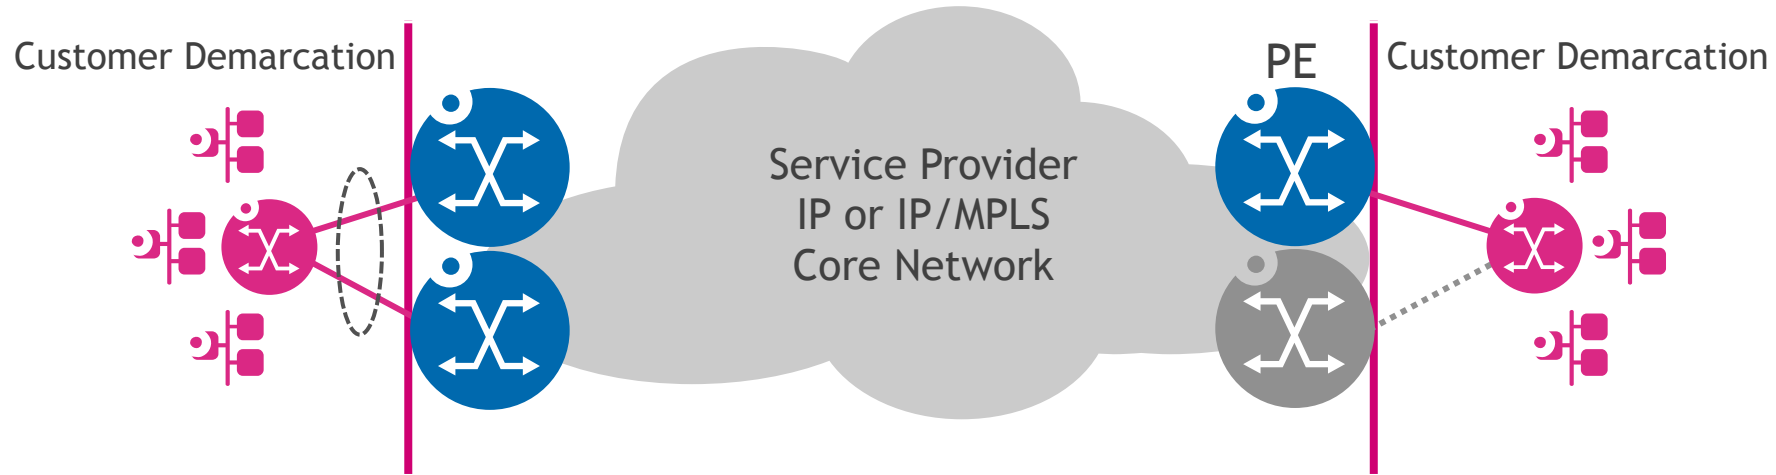
- Provides Layer 2 interconnection over an EVPN peering fabric
 - IP/MPLS core network with MPLS data plane
 - IP core network with VXLAN data plane
- Supports single or all-active multihoming to the peering fabric
- Enables precise fine-grained control over MAC addresses
 - Static MAC provisioning and ARP/ND proxy from PEs can reduce or eliminate unknown unicast
 - Per-MAC loop control vs per-port or per-VLAN isolates potential loops

LAYER 2 OR LAYER 3 DCI



- Provides scalable Layer 2 or Layer 3 DCI services for virtualized data centers
- MAC mobility for VMs that move between data centers
 - Faster moves between DCs, while keeping FDB correct on all nodes with no BUM
- Local IP gateway at each PE optimizes routing
- Provides all the benefits of EVPN for DCI and virtualized networks
 - No BUM for MAC learning
 - Integrate L2 switching and L3 routing in a single service

OVERLAY VPNS SERVICE PROVIDER TO CUSTOMER LAYER 2 AND LAYER 3 SERVICES




- EVPN provides Layer 2 and Layer 3 services
 - Single interface, single VLAN to customer
 - One VPN technology for both services, no need for multiple VPN protocols
 - All-active or single-active PE to CE connection

- EVPN service can be provided over any core network
 - MPLS core can use EVPN
 - IP core can use EVPN-VXLAN

OVERLAY VPNS OVER IP

FLEXIBLE LAYER 2 AND LAYER 3 VPN SOLUTION



- EVPN-VXLAN works over any IP service to provide a flexible Layer 2 and Layer 3 VPN
- Just requires IP connectivity between sites, no MPLS or any special configuration by IP service provider
 - Service provider network is transparent to EVPN
 - EVPN overlay is transparent to service providers
- VPN routing between endpoints can be controlled with BGP and routing policies to service providers
- Routing and MAC/IP advertisement within EVPN controlled via IBGP between PEs

SUMMARY

- EVPN provides next-generation VPN solutions for Layer 2 and Layer 3 services over Ethernet
 - MAC address learning through an MP-BGP control plane instead of the data plane
 - L3VPN-like operation for scalability and control
 - Flow-based load balancing and all-active multipathing
 - Simplified operation with auto discovery and configuration
 - ARP/ND security and MAC provisioning
 - Runs over MPLS or IP data plane
- More information
 - IETF Layer 2 Virtual Private Networks (l2vpn) Working Group
<http://datatracker.ietf.org/wg/l2vpn/>
 - Requirements: draft-ietf-l2vpn-evpn-req
<http://tools.ietf.org/html/draft-ietf-l2vpn-evpn-req>
 - Base specification: draft-ietf-l2vpn-evpn
<http://tools.ietf.org/html/draft-ietf-l2vpn-evpn>

www.alcatel-lucent.com

EVPN REQUIREMENTS AND BENEFITS

	VPN Requirements	VPLS	EVPN	What does it do for me?
Address Learning	Control Plane Address Learning in the Core	✗	✓	Greater Scalability and Control
Provisioning	L3VPN-Like Operation	✗	✓	Simpler Provisioning and Automation
	Auto Discovery and Configuration	PEs Only	✓	Simpler Provisioning and Automation
Resiliency	Active-Standby Multihoming (Service-Based Load Balancing)	✓	✓	Standby Redundancy
	All-Active Multihoming (Flow-Based Load Balancing)	✗	✓	Active Redundancy and Link Utilization
Services	VLAN Based Service Interfaces	✓	✓	Virtualization and Advanced Services
	VLAN Aware Bundling Service Interfaces	✗	✓	Virtualization and Advanced Services
	Inter-Subnet Forwarding	✗	✓	Layer 2 and Layer 3 Over the Same Interface
Flow Optimization	ARP/ND Proxy	✗	✓	Security and MAC Provisioning
	MAC Mobility	✗	✓	Virtualization and Advanced Services

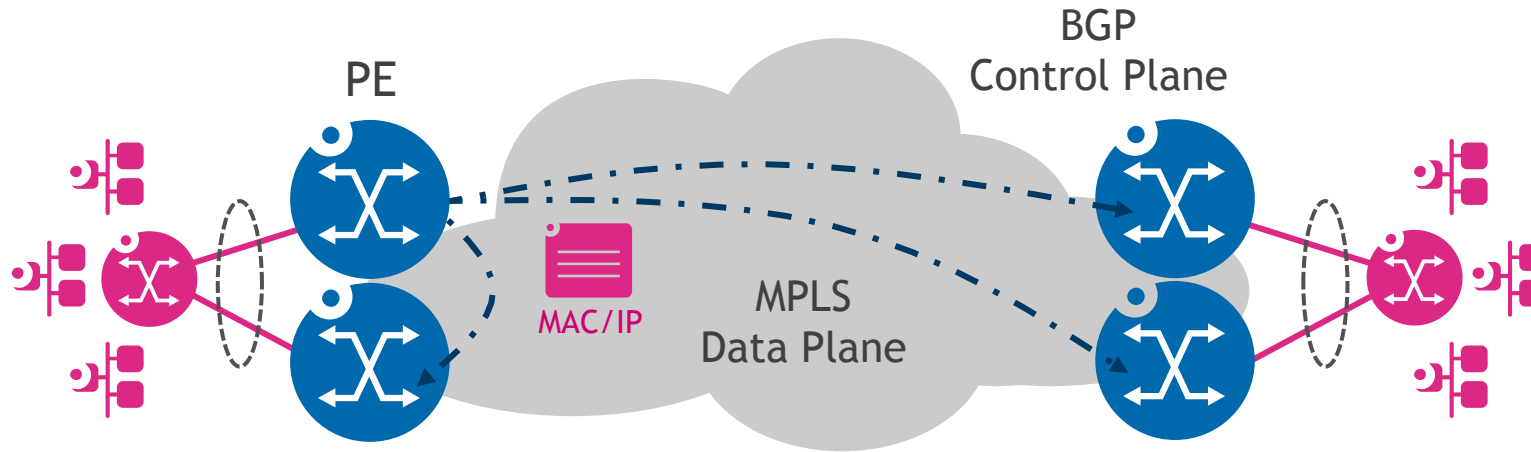
EVPN NLRI ROUTE TYPES AND EXTENDED COMMUNITIES

Route Type	Route Description	Route Usage	Reference
1	Ethernet Auto-Discovery (A-D) Route	Endpoint Discovery, Aliasing, Mass-Withdraw	draft-ietf-l2vpn-evpn
2	MAC Advertisement Route	MAC/IP Advertisement	draft-ietf-l2vpn-evpn
3	Inclusive Multicast Route	BUM Flooding Tree	draft-ietf-l2vpn-evpn
4	Ethernet Segment Route	Ethernet Segment Discovery, DF Election	draft-ietf-l2vpn-evpn
5	IP Prefix Route	IP Route Advertisement	draft-rabadan-l2vpn-evpn-prefix-advertisement

Extended Community Type	Extended Community Description	Extended Community Usage	Reference
0x06/0x01	ESI Label Extended Community	Split Horizon Label	draft-ietf-l2vpn-evpn
0x06/0x02	ES-Import Route Target	Redundancy Group Discovery	draft-ietf-l2vpn-evpn
0x06/0x00	MAC Mobility Extended Community	MAC Mobility	draft-ietf-l2vpn-evpn
0x03/0x030d	Default Gateway Extended Community	Default Gateway	draft-ietf-l2vpn-evpn, bgp-extended-communities

EVPN MULTIPROTOCOL LABEL SWITCHING (MPLS) DATA PLANE

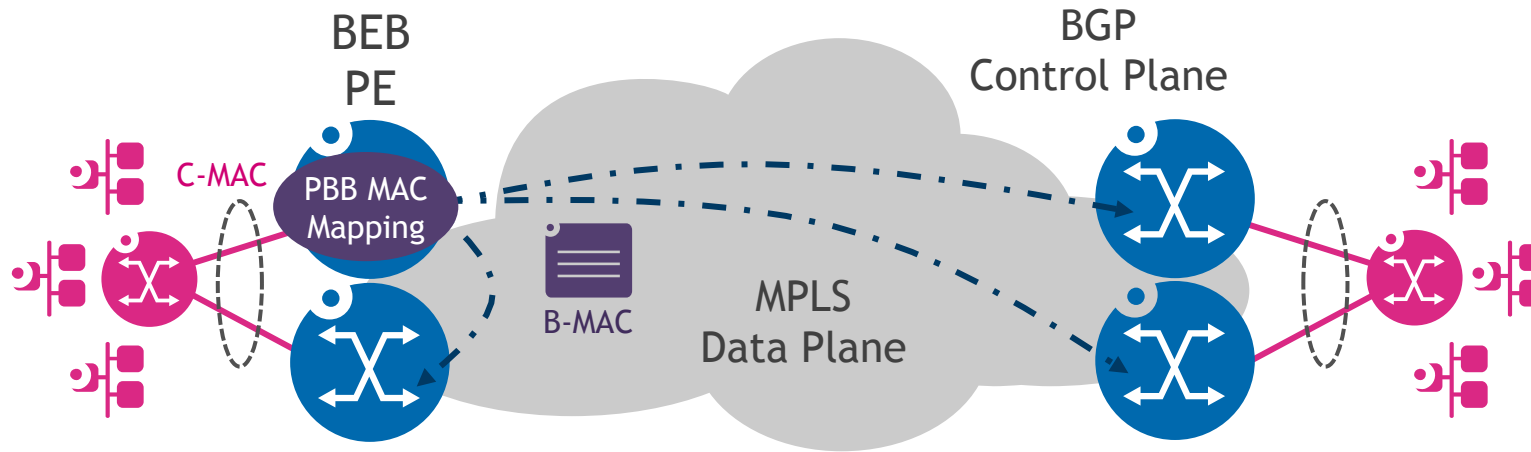
DRAFT-IETF-L2VPN-EVPN



- EVPN over an MPLS data plane is the original EVPN solution in the base specification
- Requires IGP, RSVP-TE or LDP, BGP
- No pseudowires
- MPLS runs in the core network's control plane and data plane
- Core network supports all the MPLS features we know and love, since EVPN uses MPLS as the data plane (TE, FRR, ...)

EVPN PROVIDER BACKBONE BRIDGES (PBB) DATA PLANE

DRAFT-IETF-L2VPN-PBB-EVPN PBB-EVPN



- PBB-EVPN combines IEEE 802.1ah PBB with EVPN
- PEs are PBB Backbone Edge Bridges (BEB)
- Reduces number of MACs in EVPN by aggregating customer MACs with backbone MACs

- Scales very large networks to millions of MACs
 - PEs only advertise backbone MACs with BGP
 - Customer MAC and backbone MAC mapping is learned in the data plane
- MPLS runs in the control plane and data plane