I-AS MPLS Solutions

BRKMPL-2105

The Prerequisites

- Must understand basic IP routing
- Familiar with MPLS Architectures
- Familiar with MPLS Applications
- Some level of MPLS network Design/ Deployment Experience

Agenda

- Inter-AS Networks
  - Inter-AS Connectivity Models
  - Inter-AS L3 VPNs
  - Inter-AS L2VPNs
  - Inter-AS Multicast VPNs
- Carrier Supporting Carrier
  - CSC Service Models
  - MPLS L3 VPNs
  - Multicast VPNs
  - MPLS L2 VPNs
- Inter-AS Traffic Engineering

Inter-Provider MPLS Solutions

- To interconnect multiple independently managed MPLS Domains
  - Fast geographic service coverage expansion
  - Two MPLS VPN Providers peering to cover for a common customer base
- Support original multi-domain network design
  - IGP isolation with service continuity
  - Interconnect BGP confederations with different IGPs in the same AS
- Two models available:
  - 1.Carrier Supporting Carrier (CSC)
  - 2.Inter-Autonomous Systems (I-AS)

Carrier Supporting Carrier vs. Inter-AS

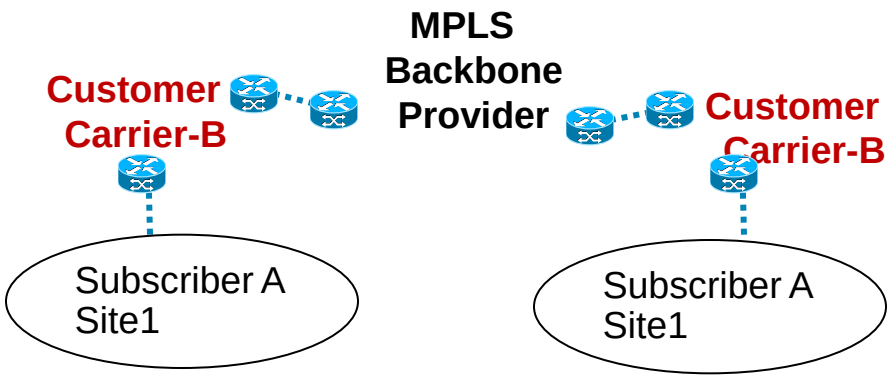

CSC

- Client-Server model
- IP/MPLS Carrier is a customer of another MPLS backbone provider
- IP/MPLS Carrier doesn't want to manage own backbone
- Only the backbone provider is required to have MPLS VPN core
- Customer Carriers do not distribute their subscribers' VPN info to the backbone carrier

I-AS L3 VPNs

Overview

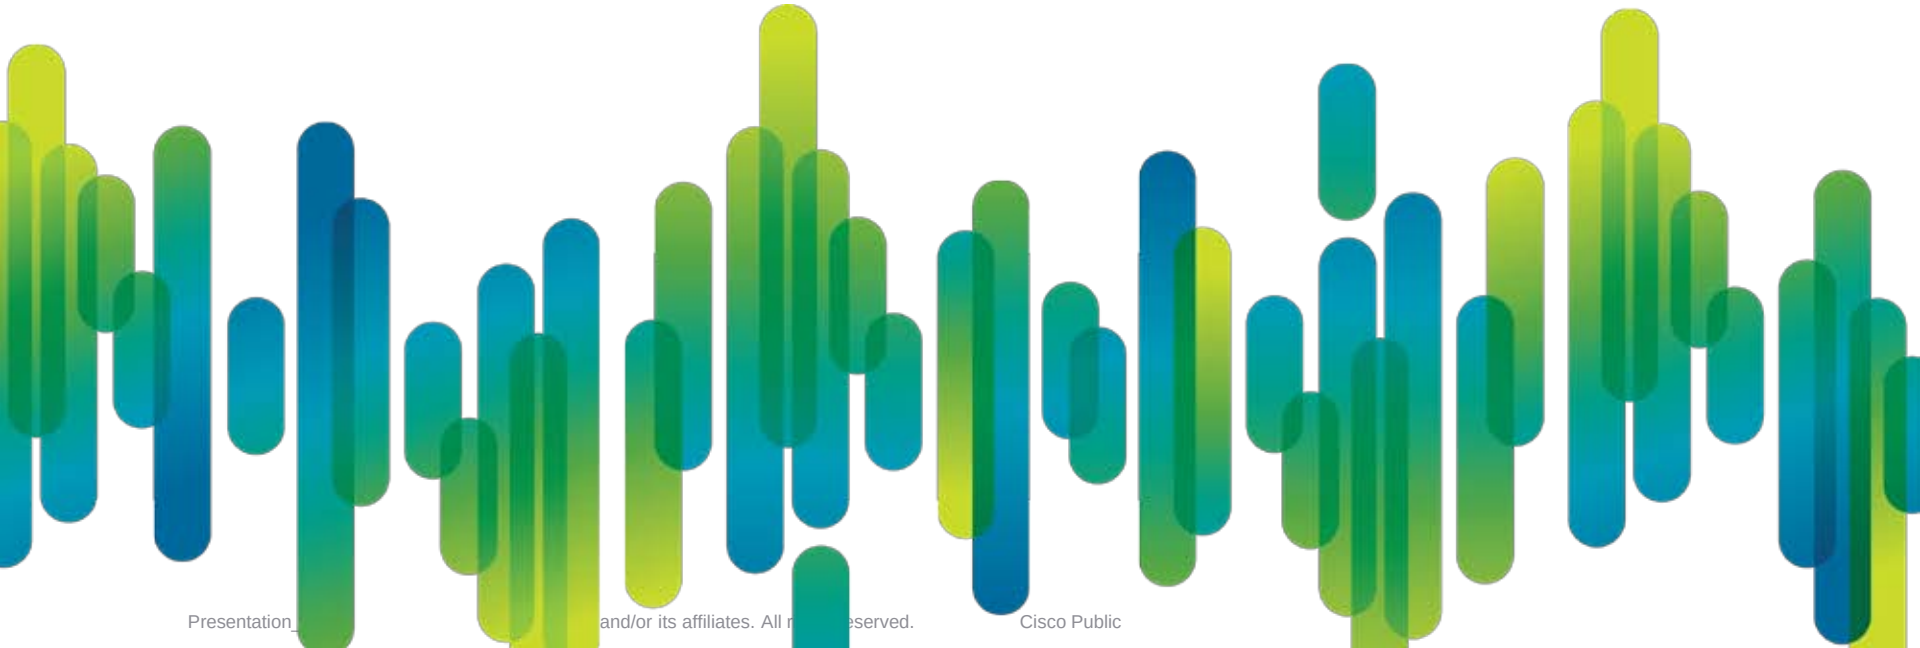

Inter-AS VPNv4 Distribution Options

How to Distribute VPN Routes
between ASBRs?

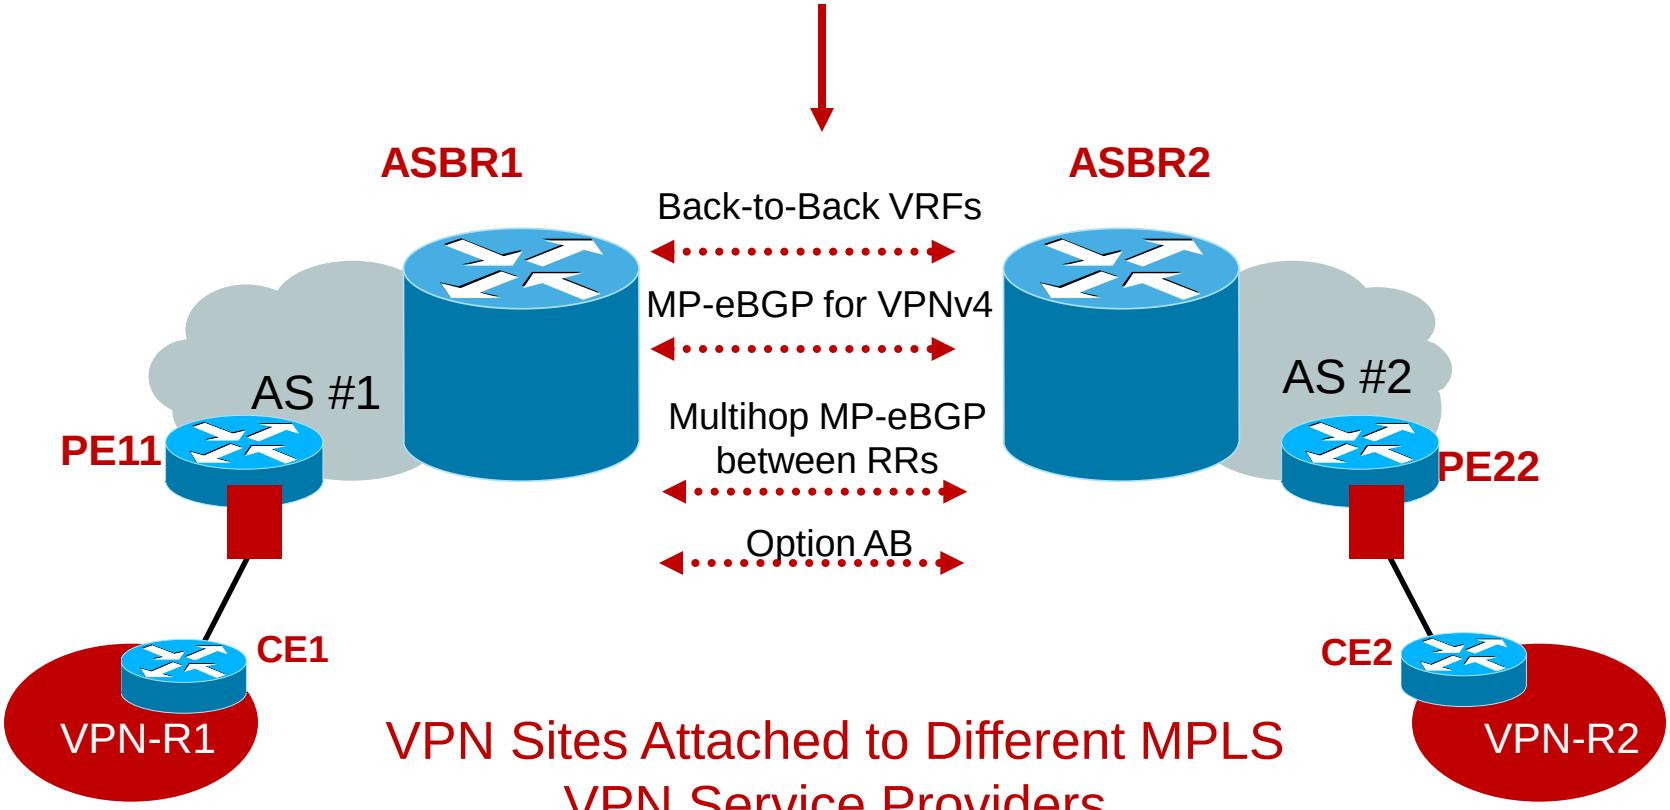

VPN Sites Attached to Different MPLS
VPN Service Providers

Inter-AS VPN—Option A

Back-to-Back VRFs

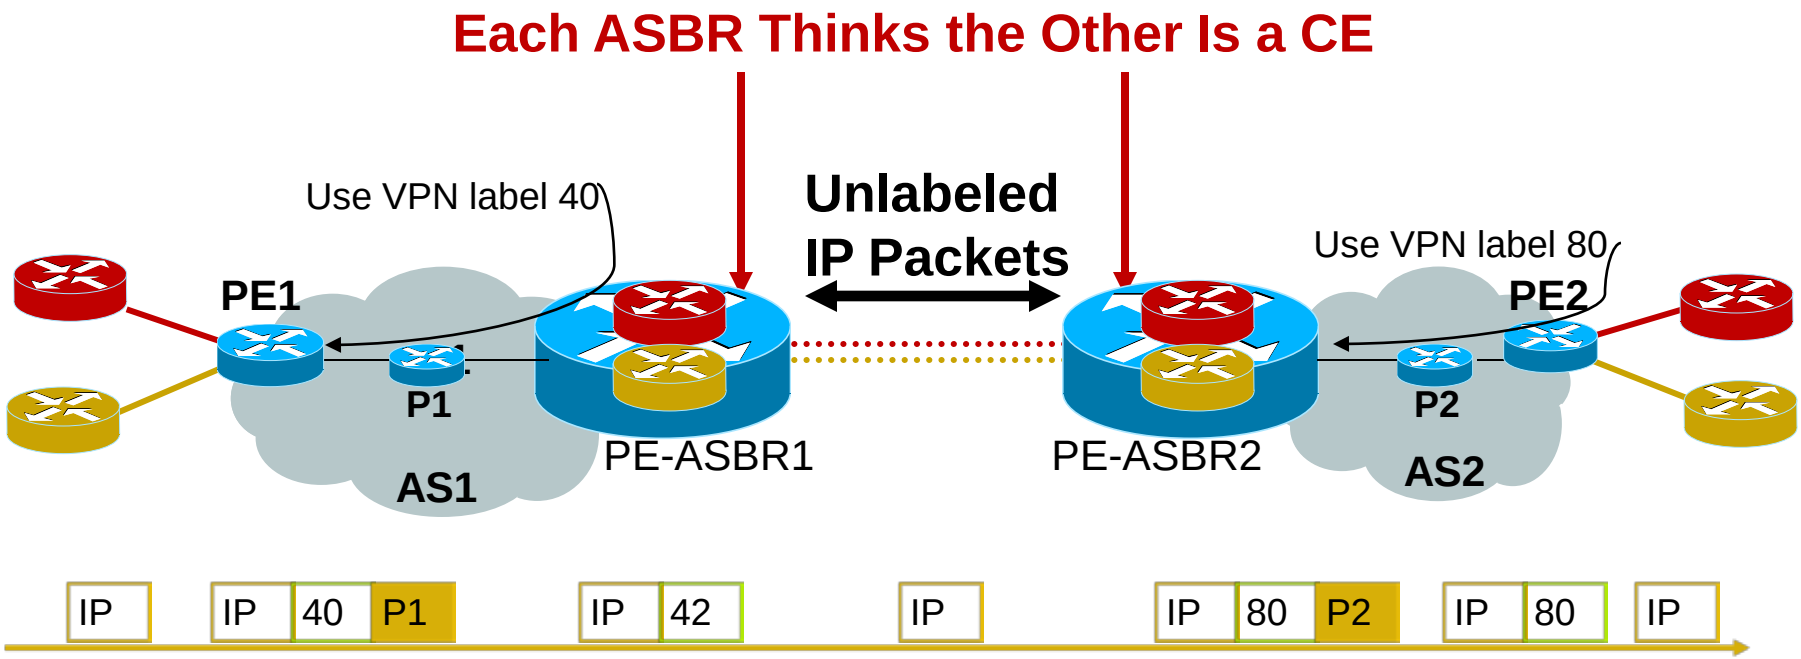

- One logical interface per VPN on directly connected ASBRs
- Packet is forwarded as an IP packet between the ASBRs
- Link may use any supported PE-CE routing protocol
- IP QoS policies negotiated and configured manually on the ASBRs
- Option A is the most secure and easiest to provision
- May not be easy to manage as #s of VPNs grow

Inter-AS VPN—Option B

Setting up Control Plane

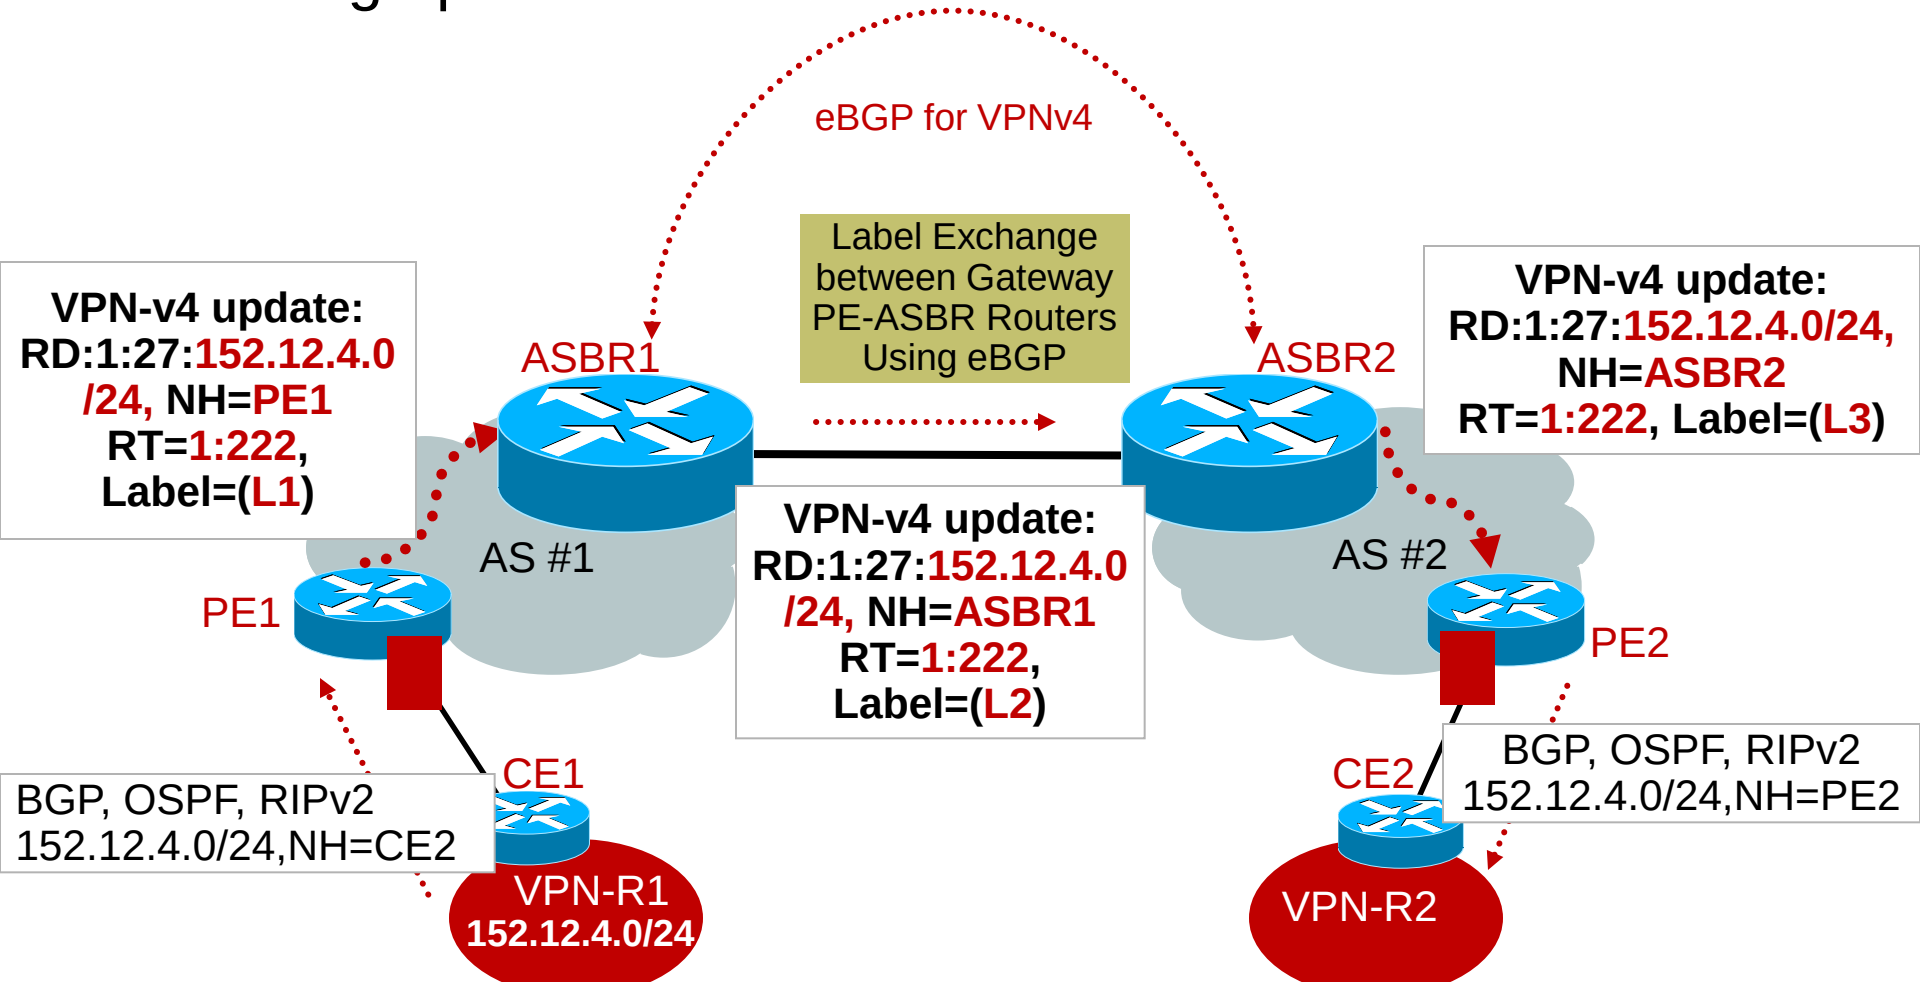

All VPNv4 Prefixes/Labels from PEs Distributed to PE-ASBRs

Inter-AS VPN—Option B
Key Points

- PE-ASBRs exchange routes directly using eBGP
  - External MP-BGP for VPNv4 prefix exchange;
- MP-BGP session with NH to advertising PE-ASBR
  - Next-hop and labels are rewritten when advertised across the inter-provider MP-BGP session
- Receiving PE-ASBR automatically creates a /32 host route to a peer ASBR
  - Which must be advertised into receiving IGP if next-hop-self is not in operation to maintain the LSP
- PE-ASBR stores all VPN routes that need to be exchanged
  - But only within the BGP table
  - No VRFs; labels are populated into the LFIB of the PE-ASBR
- ASBR-ASBR link must be directly connected!!!!!! Could use GRE tunnel-considered directly connected
- Receiving PE-ASBRs may allocate new label
  - Controlled by configuration of next-hop-self (default is off)

Inter-AS VPN—Option B

Packet Forwarding between MPLS VPN AS

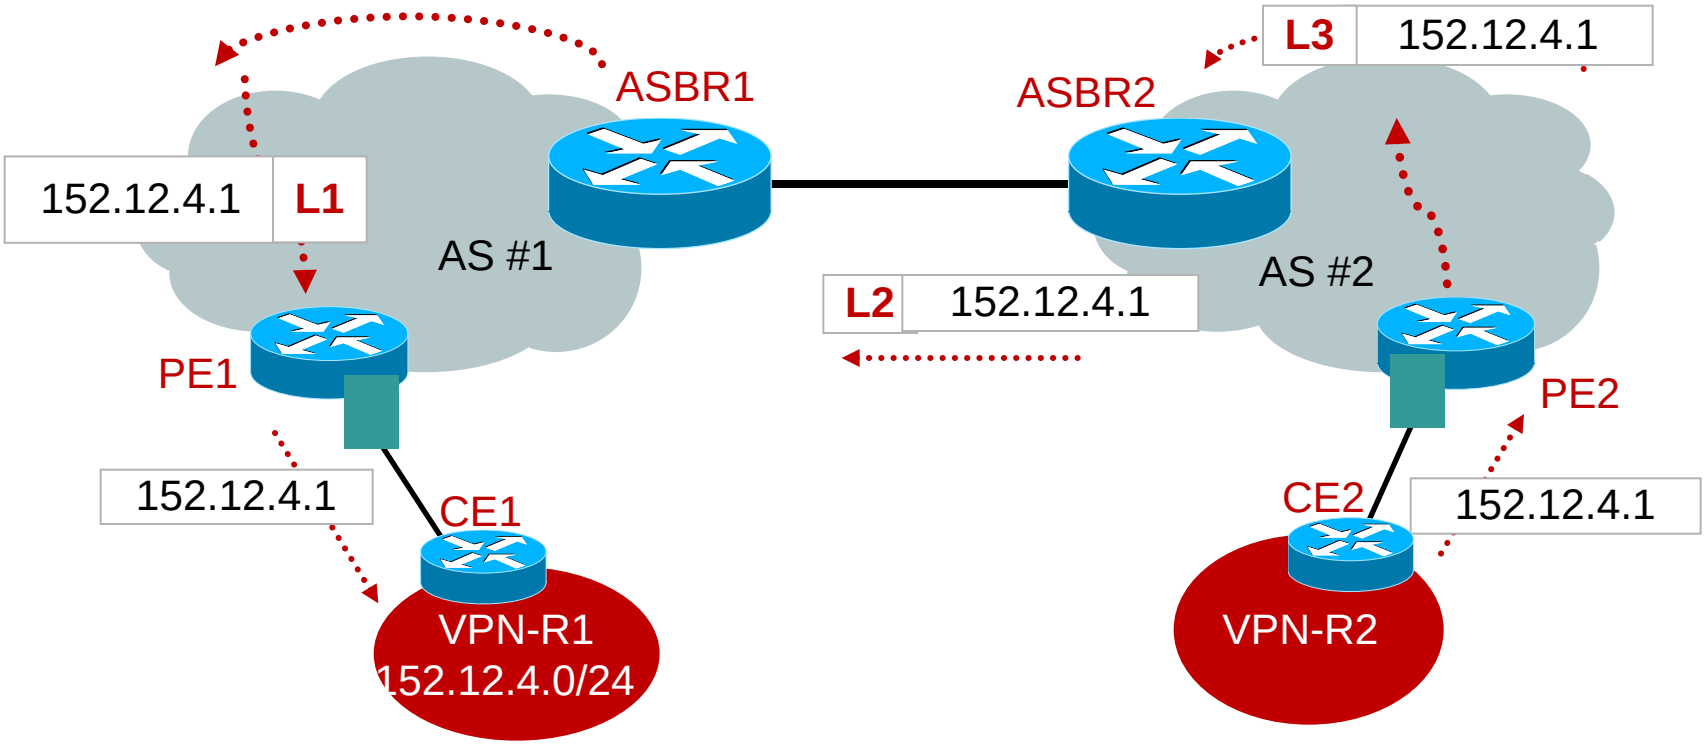

Note: The outer most core (IGP labels in an AS) label is not displayed in this presentation

Inter-AS VPN—Option B

Cisco IOS Configuration

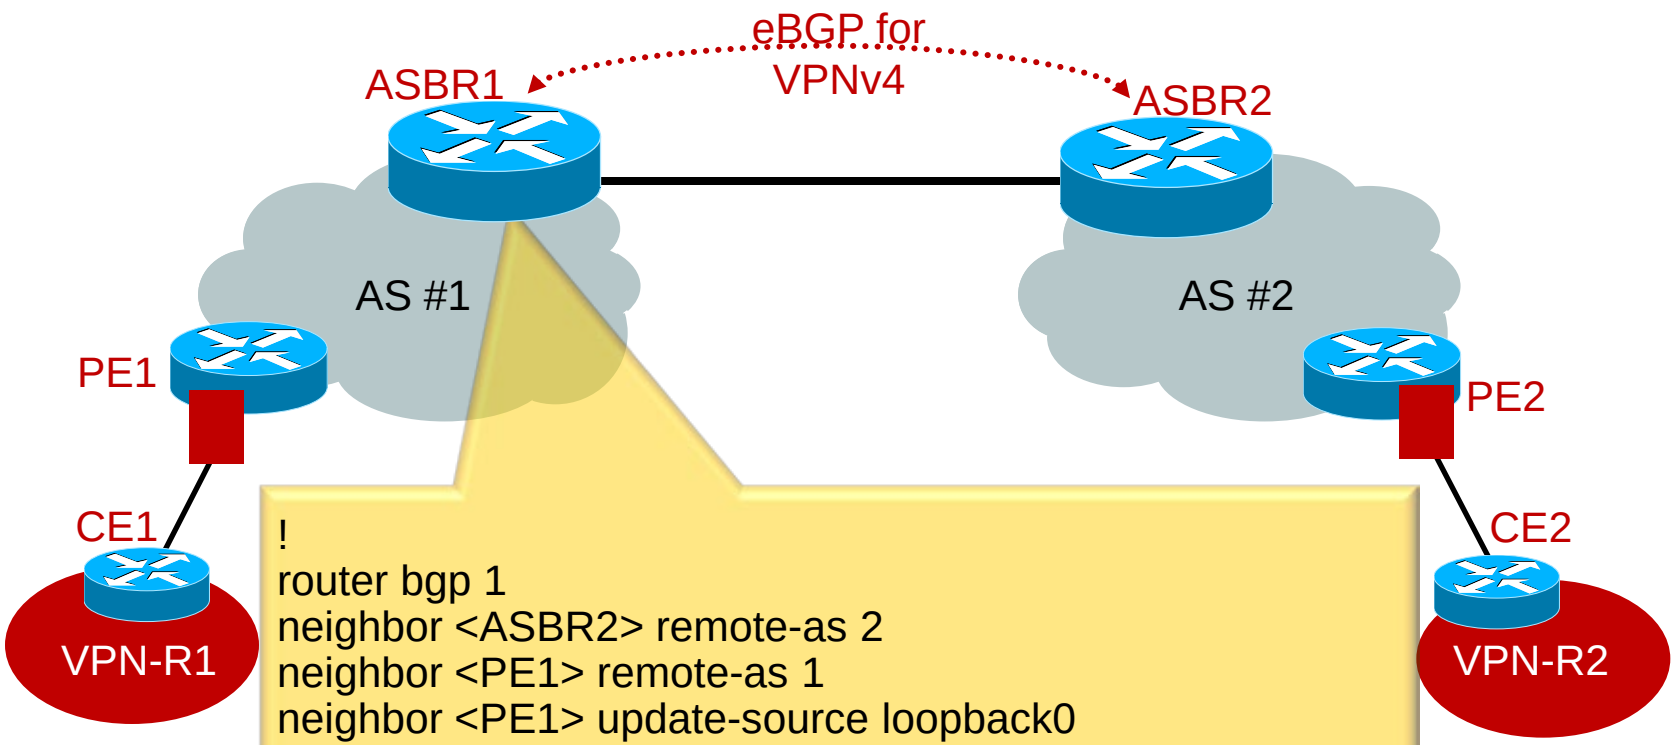

Inter-AS VPN—Option C

Multi-hop eBGP VPNv4 between RRs

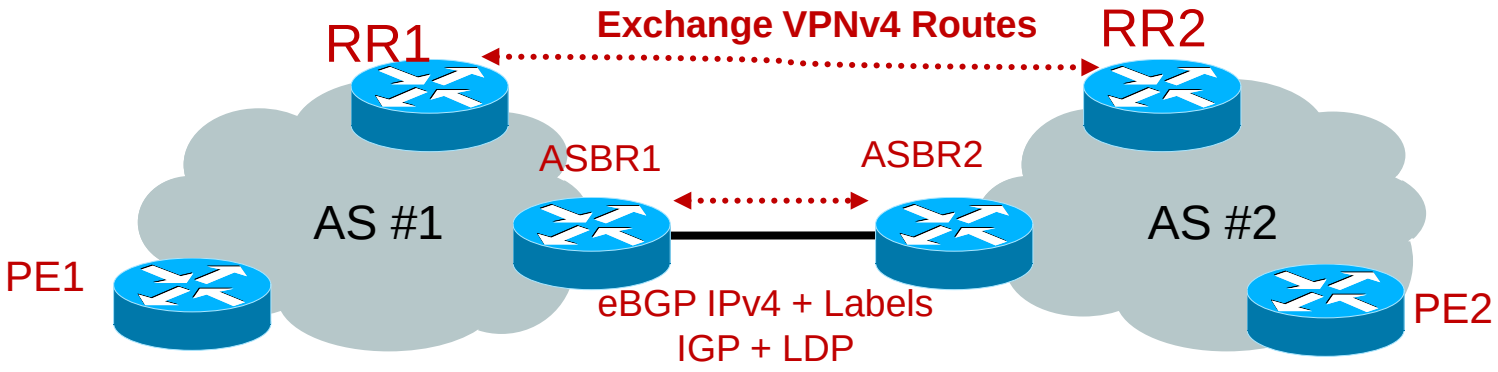

- Eliminates LFIB duplication at ASBRs. ASBRs don't hold VPNv4 prefix/label info.
- ASBRs Exchange PE loopbacks (IPv4) with labels as these are BGP NH addresses
- Two Options for Label Distribution for BGP NH Addresses:

  IGP + LDP OR BGP IPv4 + Labels (RFC3107)
- BGP exchange Label Advertisement Capability - Enables end-end LSP Paths
- Subsequent Address Family Identifier (value 4) field is used to indicate that the NLRI contains a label
- Disable Next-hop-self on RRs

I-AS VPN—Option C

Setting up Control Plane

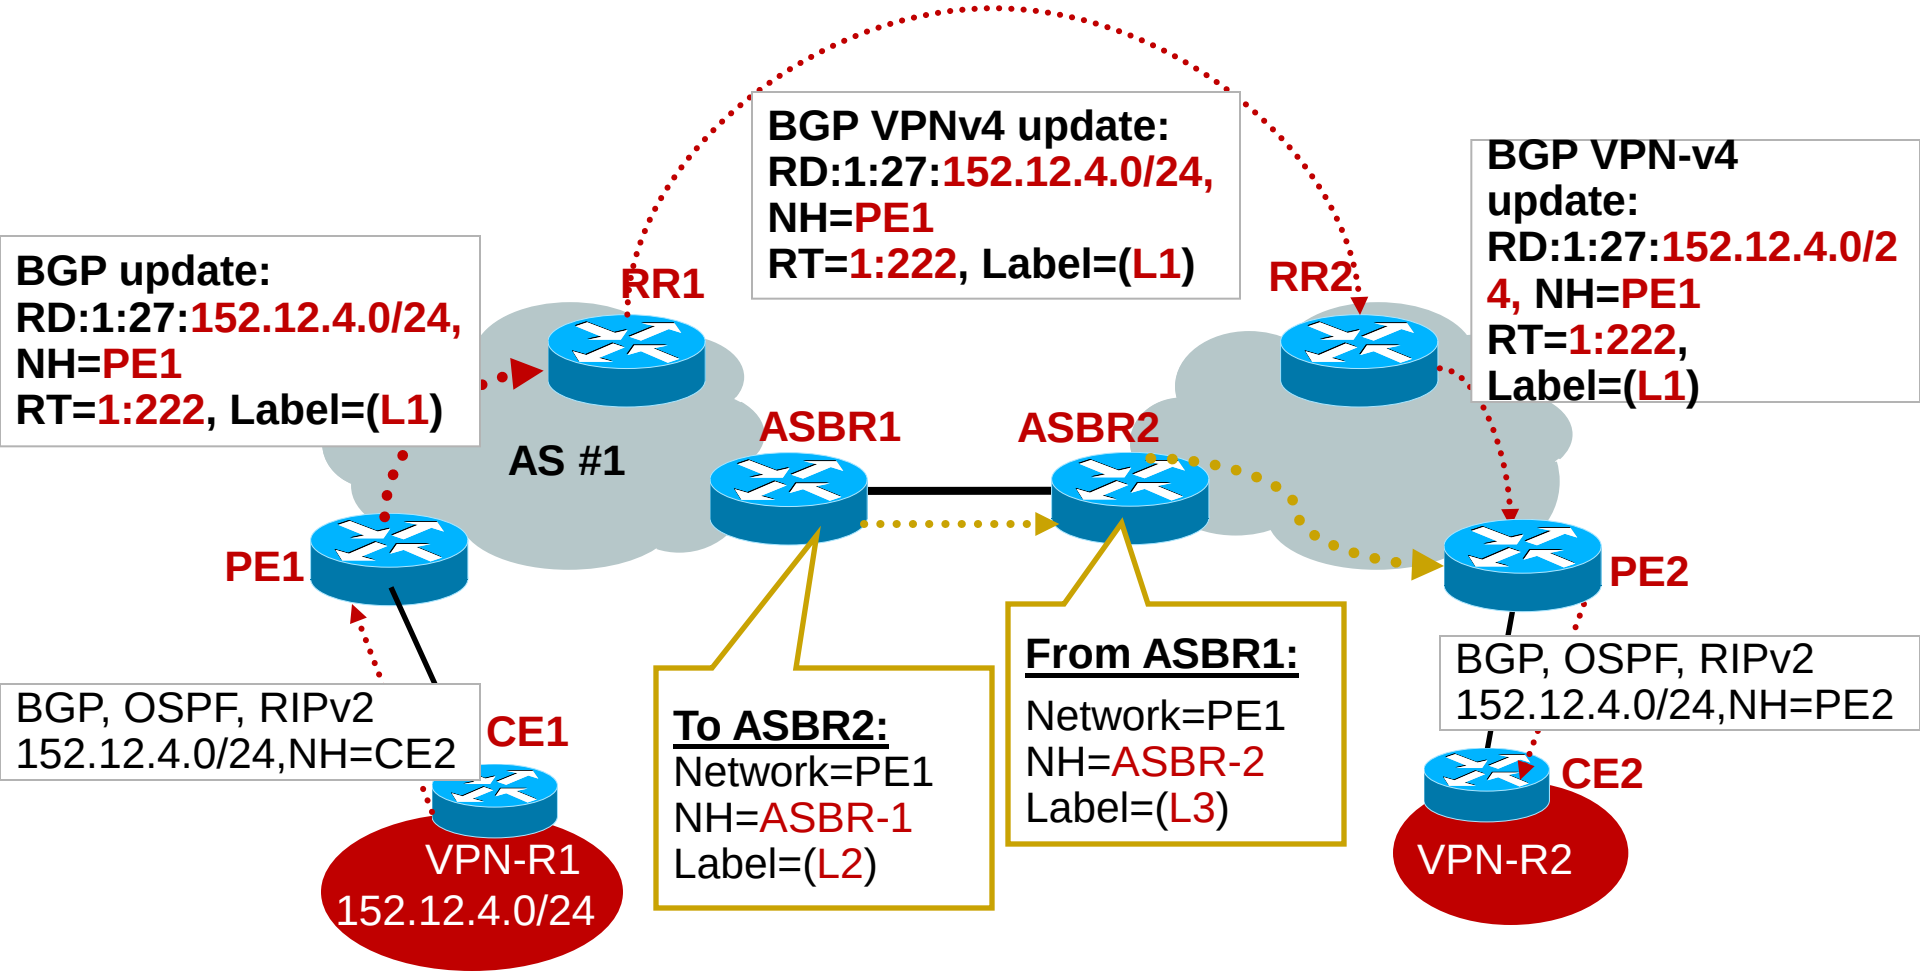

I-AS VPN—Option C

Forwarding Plane

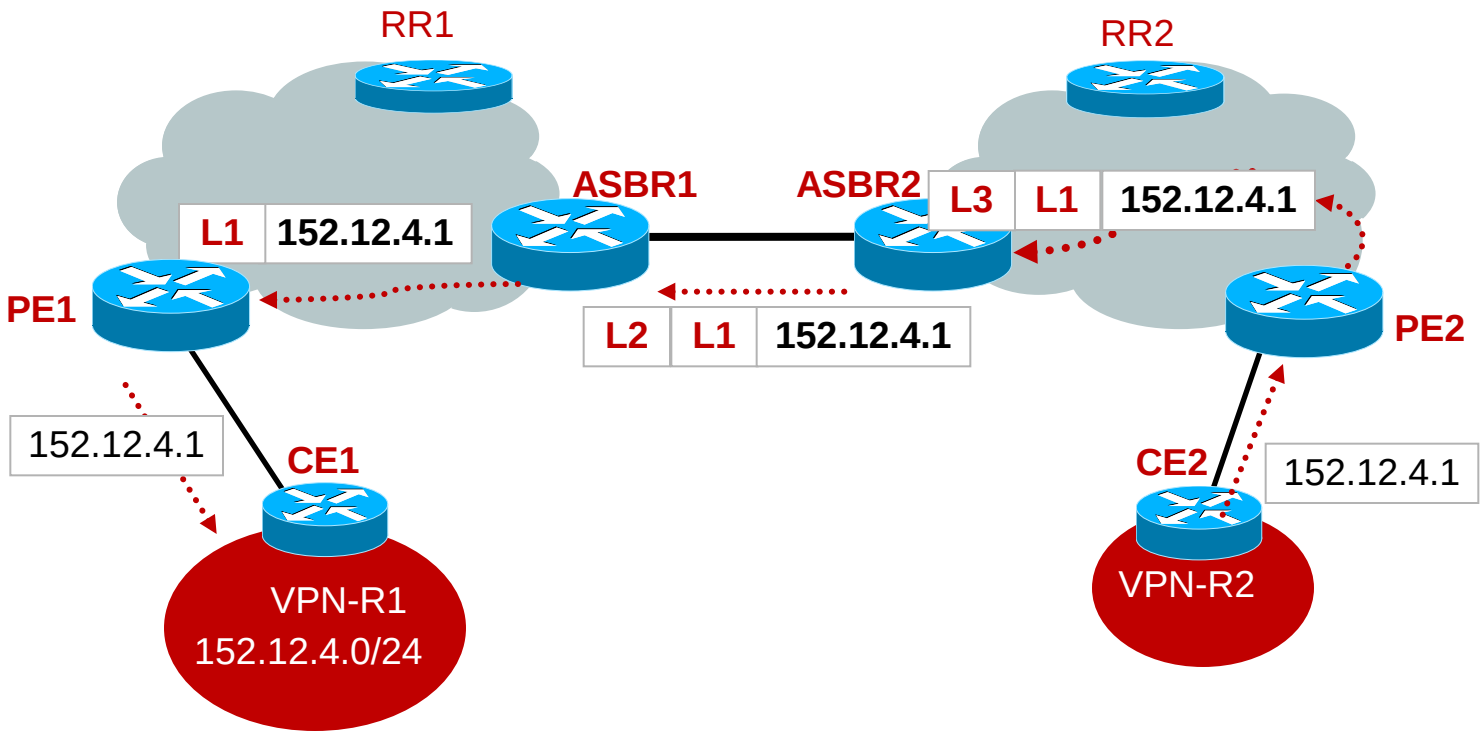

Note: The diagram does not display an outer most core (IGP labels in an AS) label

I-AS VPN—Option C

IPv4+Label, Cisco IOS Configuration

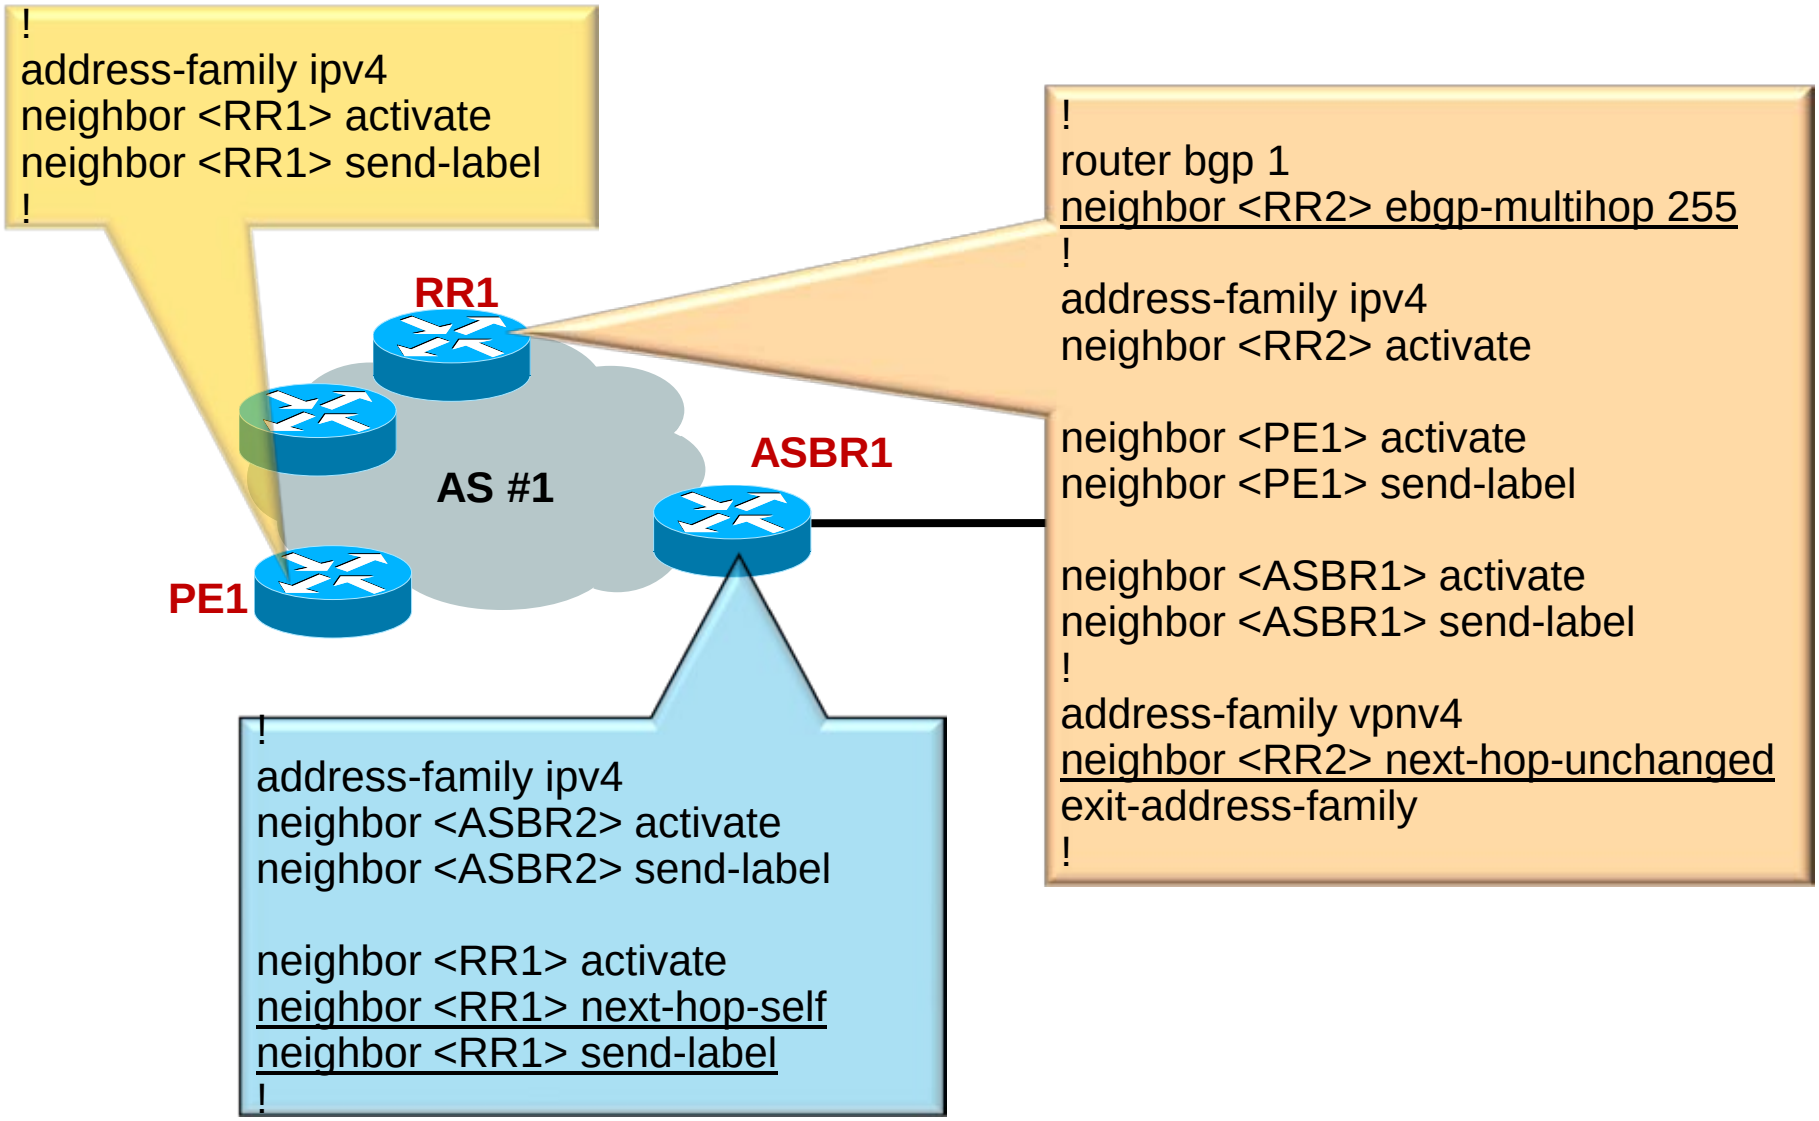

Inter-AS Multipath Load Balance Options

- Support VPNv4 and label negotiated IPv4 eBGP sessions between loopbacks of directly connected routers w/o the use of LDP on the connecting interfaces
- Consider the three topologies – Designated by Topo-1, Topo-2, Topo-3
- Load balancing for Inter-AS sub-cases with:
  1. Interface Peering
  2. Loopback peering
  3. IPv4 + Label
  4. VPNv4 + Label

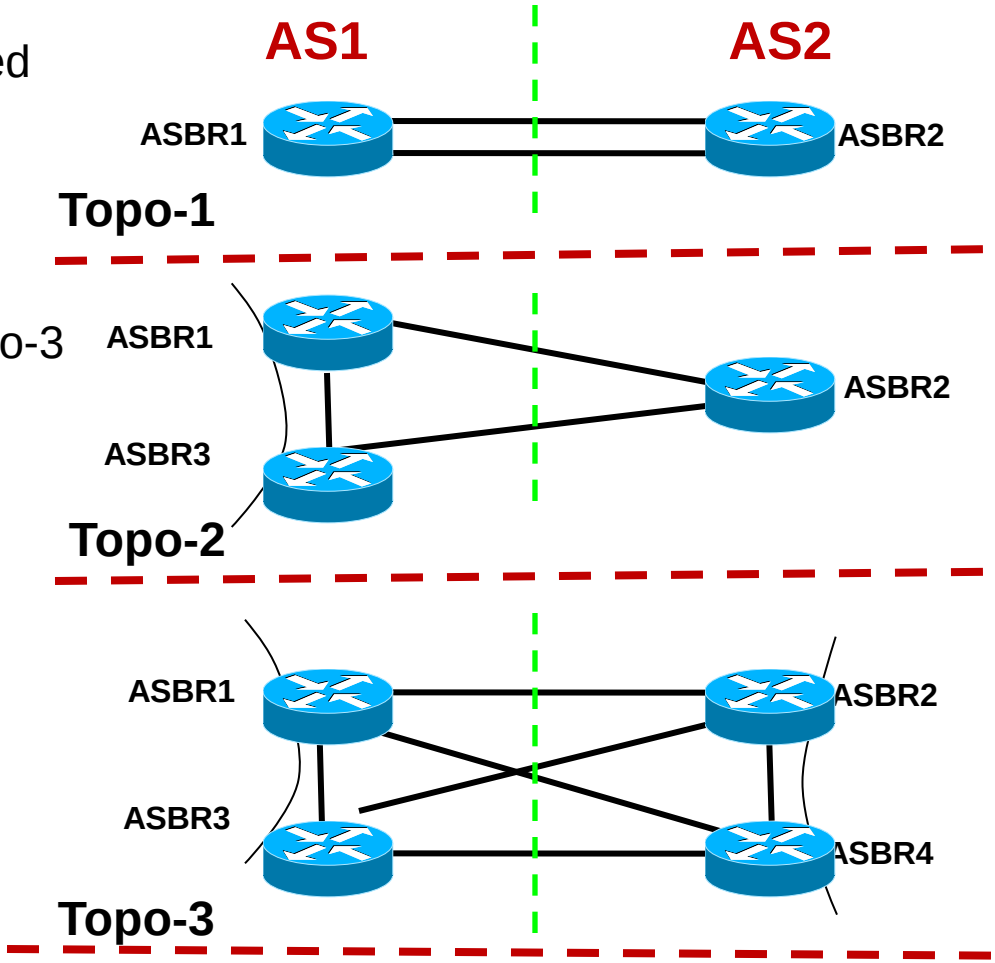

Inter-AS Loopback Peering for Directly Connected ASBRs

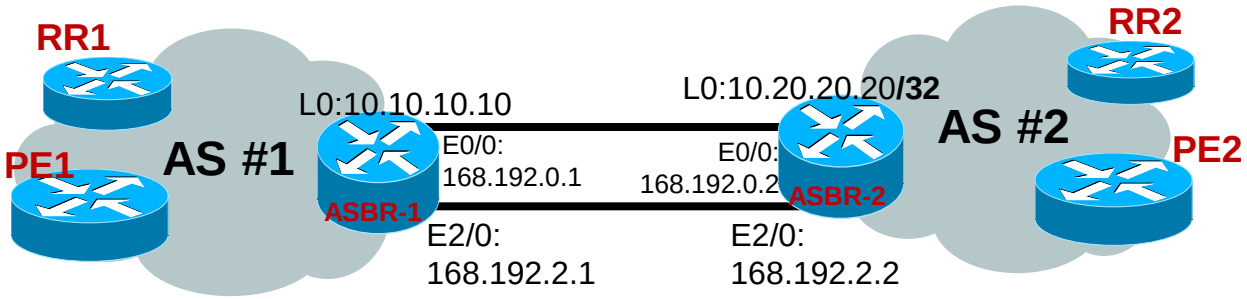

Create loopback interfaces on directly connected ASBRs

```
HOSTNAME ASBR2
!
interface e0/0
ip address 168.192.0.2 255.255.255.252
 mpls bgp forwarding
  ! Enable BGP forwarding on connecting interfaces
!
interface e2/0
ip address 168.192.2.2 255.255.255.252
 mpls bgp forwarding
!
router bgp 2
neighbor 10.10.10.10 remote-as 1
neighbor 10.10.10.10 disable-connected-check
neighbor 10.10.10.10 update-source Loopback0
!
```

```
!
address-family vpnv4
neighbor 10.10.10.10 activate
neighbor 10.10.10.10 send-community extended
!
ip route 10.10.10.10 255.255.255 e0/0 168.192.0.1
ip route 10.10.10.10 255.255.255 e2/0 168.192.2.1
  ! Configure /32 static routes to the eBGP neighbor
loopback address
```

Inter-AS Security Elements

- MD5 Authentication on LDP/BGP Sessions
- Apply max prefix
- Static Labels
- TTL Check to diagnose DoS attacks
- Filtering with BGP attributes ASPATH, ext communities, RDs checks, …etc. Set route-maps to filter and send only the desirable prefixes
- RT Constraint (filtering)
- Customize Route Targets, RT Rewrite

Route Target Rewrite Example

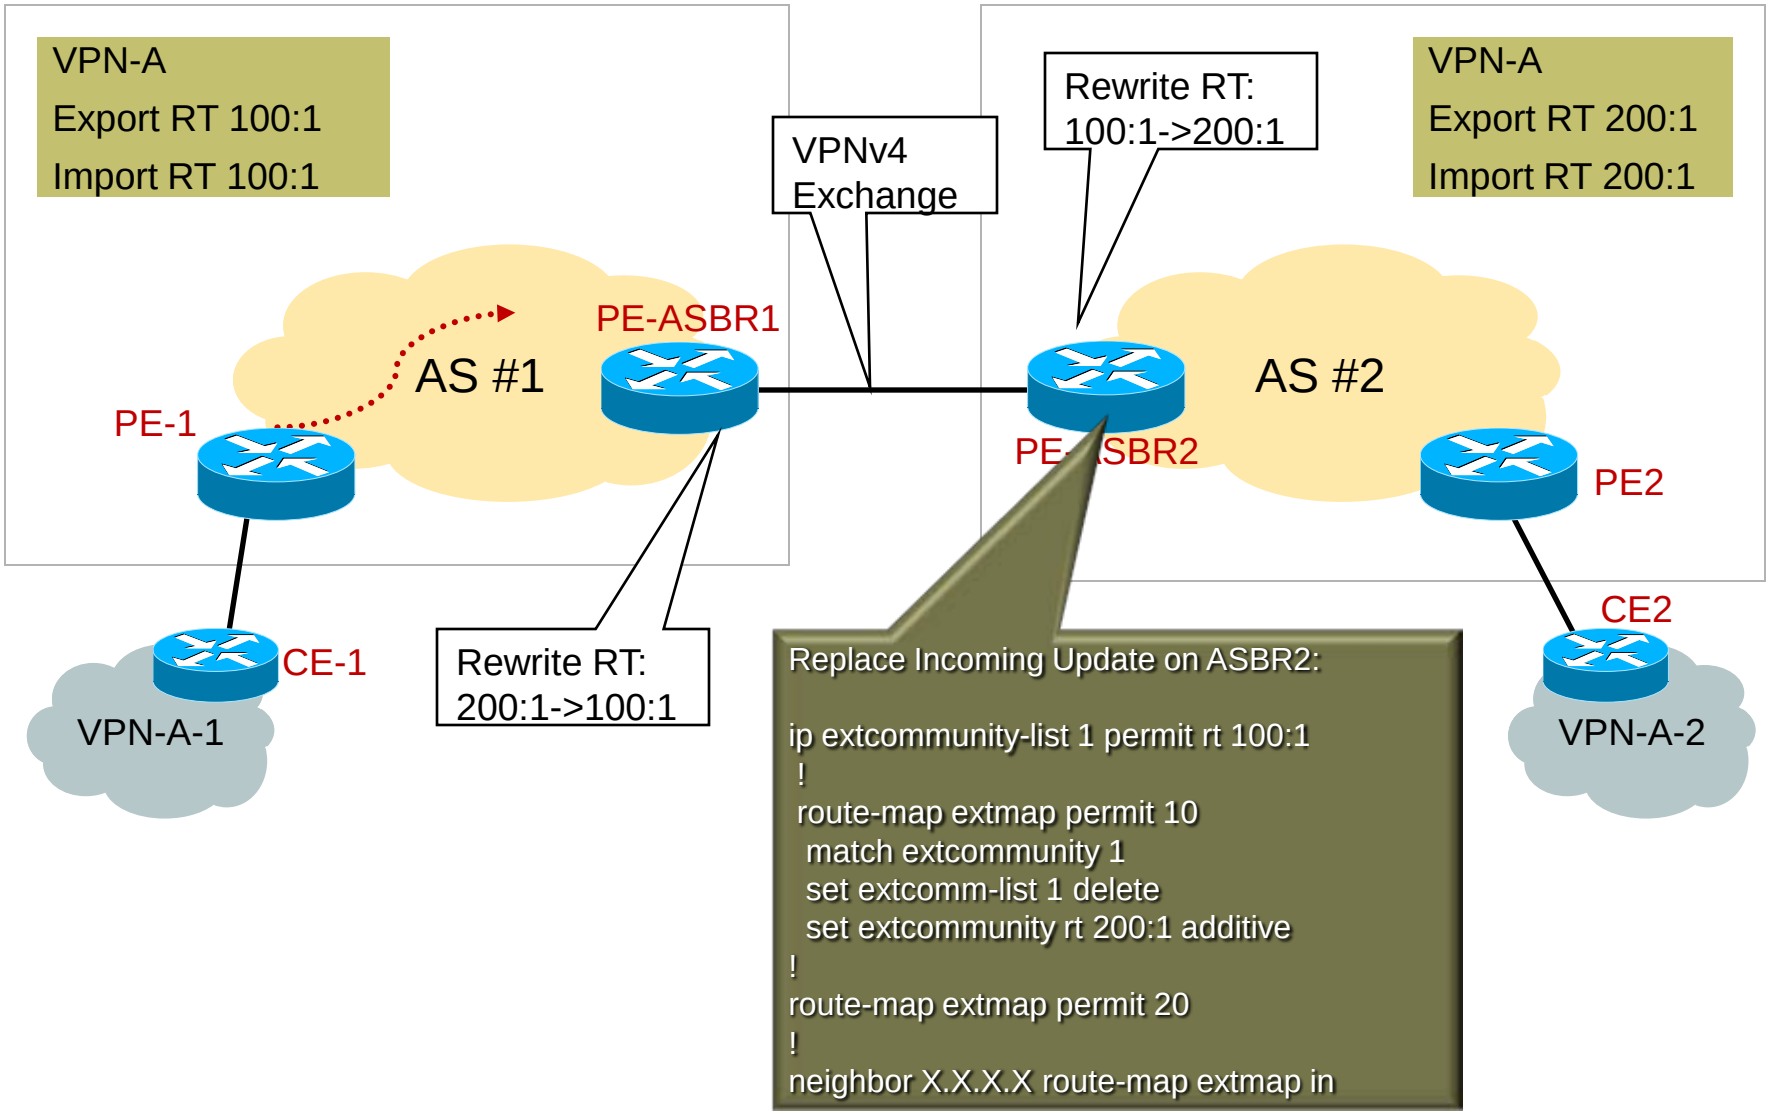

Inter-AS L3VPN Summary

- Three models: Option A, B, and C
- Option A is the most secured. Support granular QoS
- Option B, less invasive
- Option B, only need to know the loopback or interface address of directly connected ASBR
- Option C, most scalable, most invasive, mostly deployed in a single service provider's multi-AS network

I-AS IPv6 VPNs

Overview

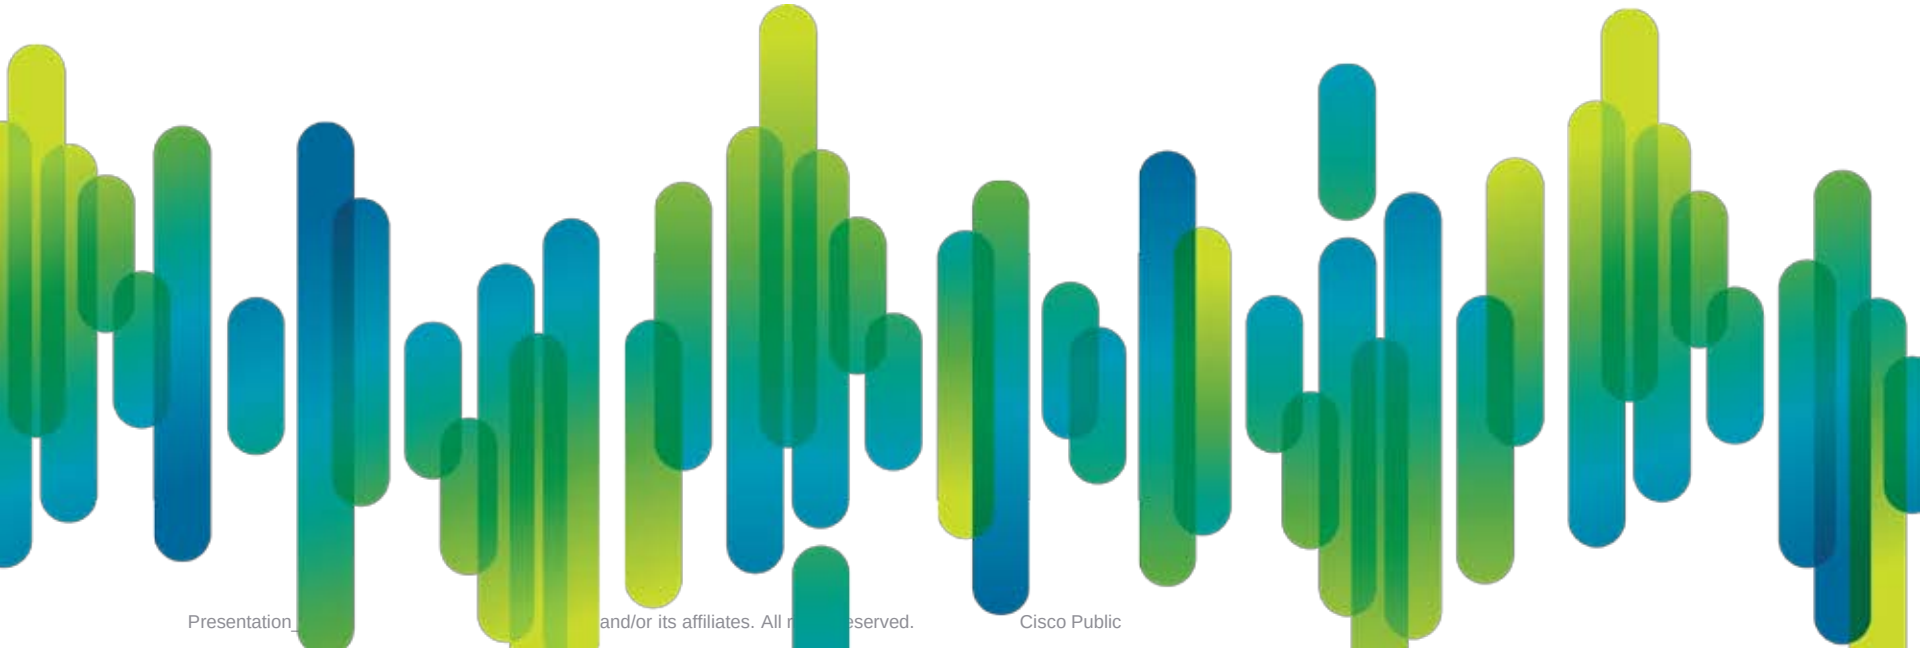

Inter-AS IPv6 VPN

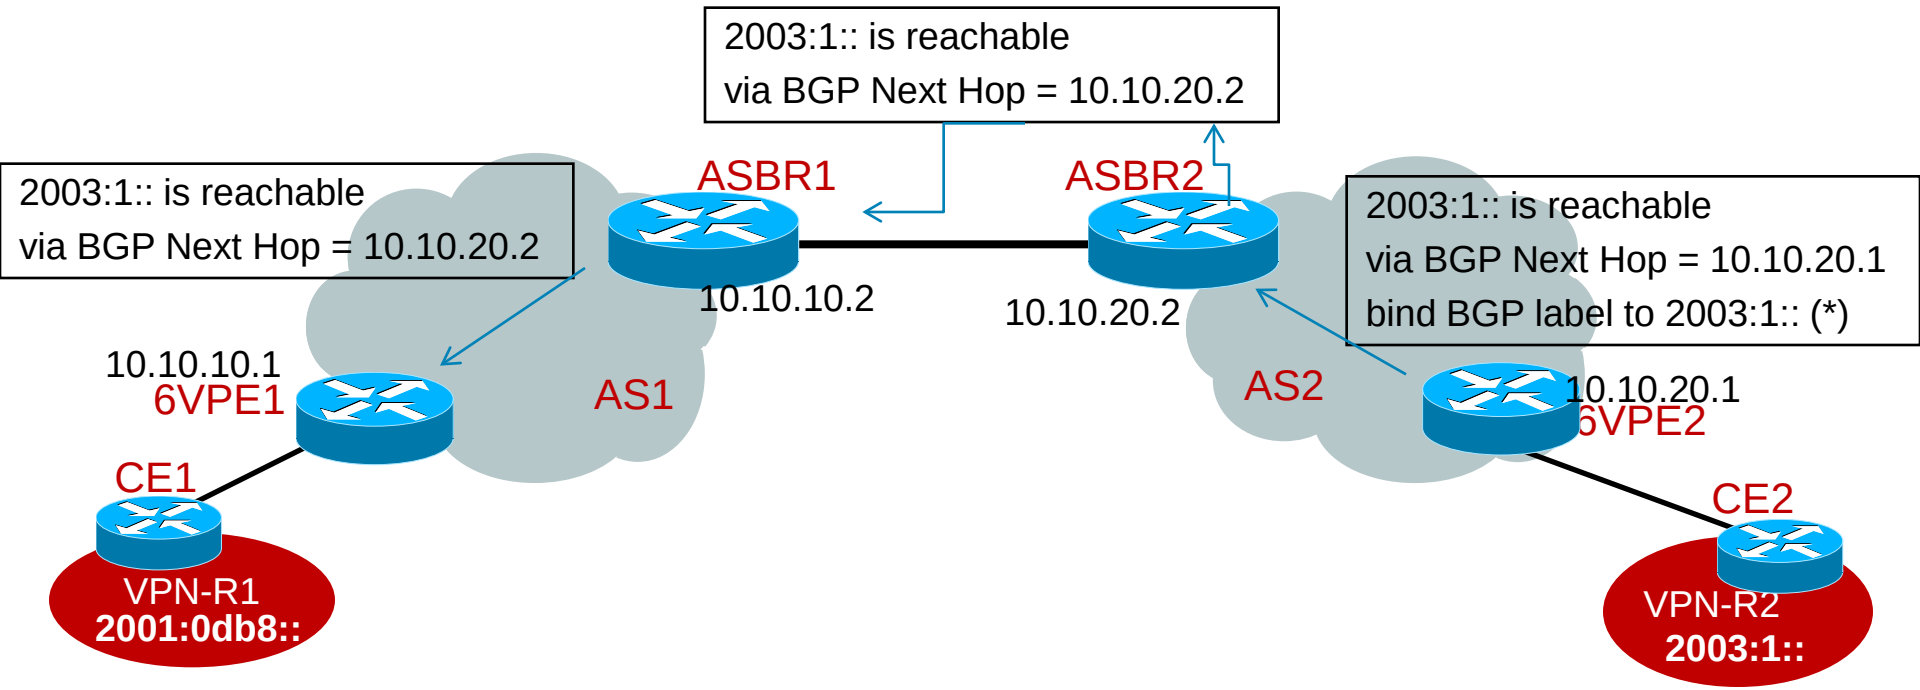

- All three ASBR-to-ASBR connectivity options discussed in earlier sections are supported for
  - -IPv6 Provider Edge Router - 6PE – model (uses vanilla IPv6)
  - -IPv6 VPN Provider Edge - 6VPE – model (uses option A,B,C)
- IPv4 address is used for PE-ASBR and ASBR-ASBR peering

Inter-AS IPv6 VPN Configuration

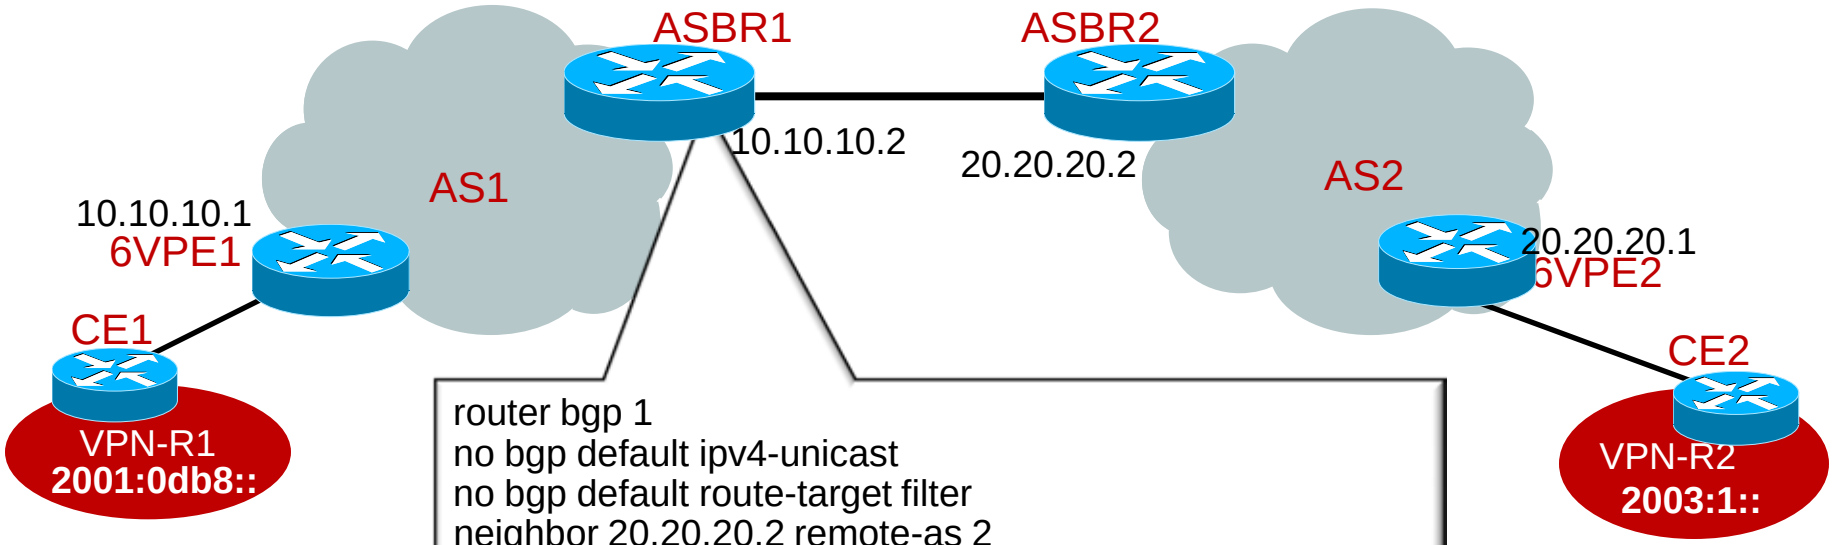


```
router bgp 1
no bgp default ipv4-unicast
no bgp default route-target filter
neighbor 20.20.20.2 remote-as 2
neighbor 10.10.10.1 remote-as 1
neighbor 10.10.10.1 update-source Loopback1
!
address-family vpnv6
!Peering to ASBR2 over an IPv4 link!
neighbor 20.20.20.2 activate
neighbor 20.20.20.2 send-community extended
!Peering to PE1 over an IPv4 link!
neighbor 10.10.10.1 activate
neighbor 10.10.10.1 next-hop-self
neighbor 10.10.10.1 send-community extended
```

Inter-AS L2 VPNs: VPWS

I-AS Virtual Private Wire Service: Any Transport over MPLS

Overview

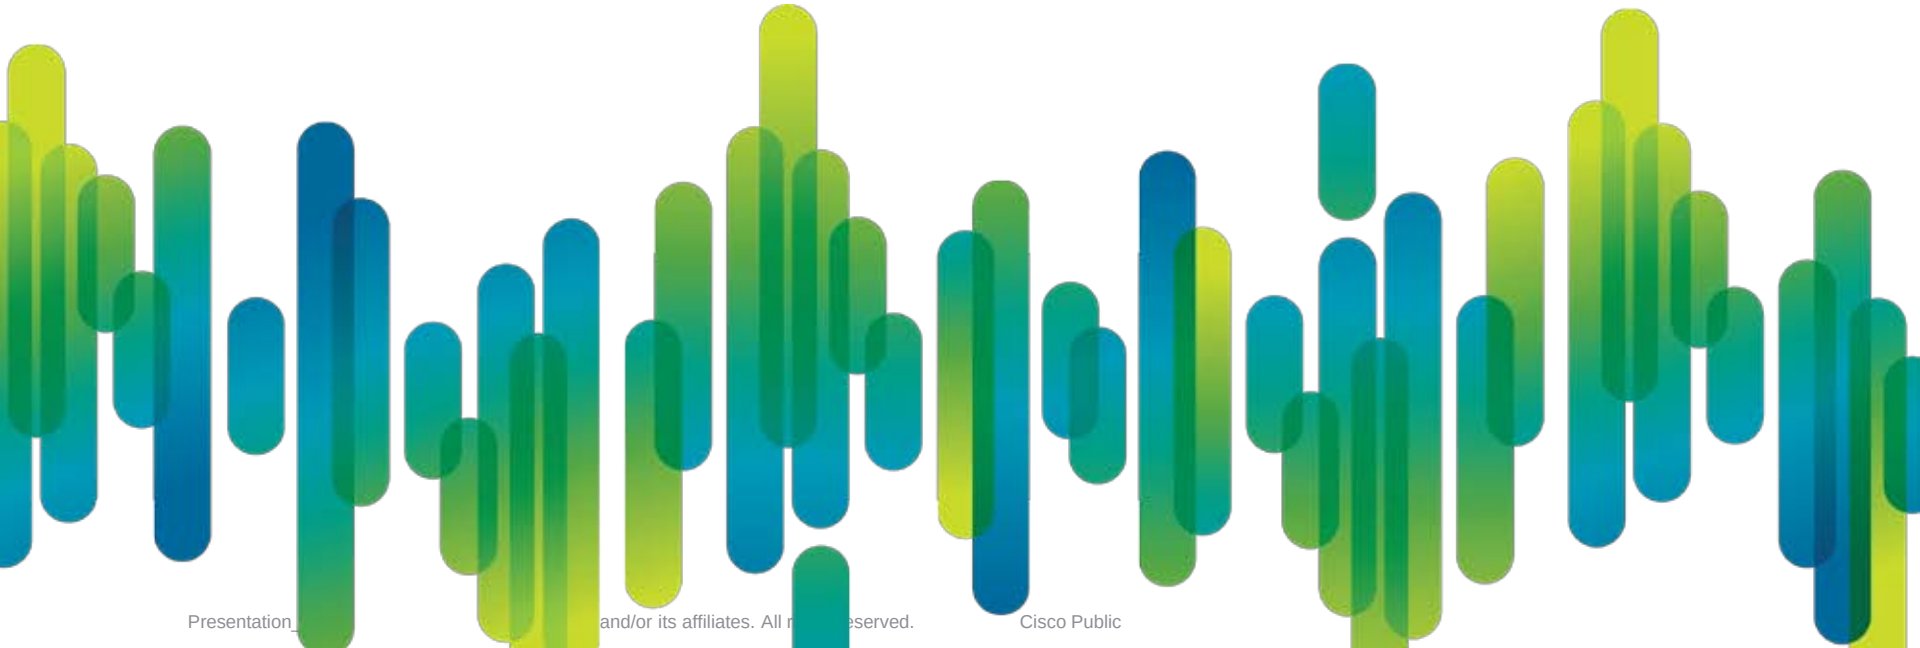

Inter-AS L2VPN
Multiple PW Segments using Option A

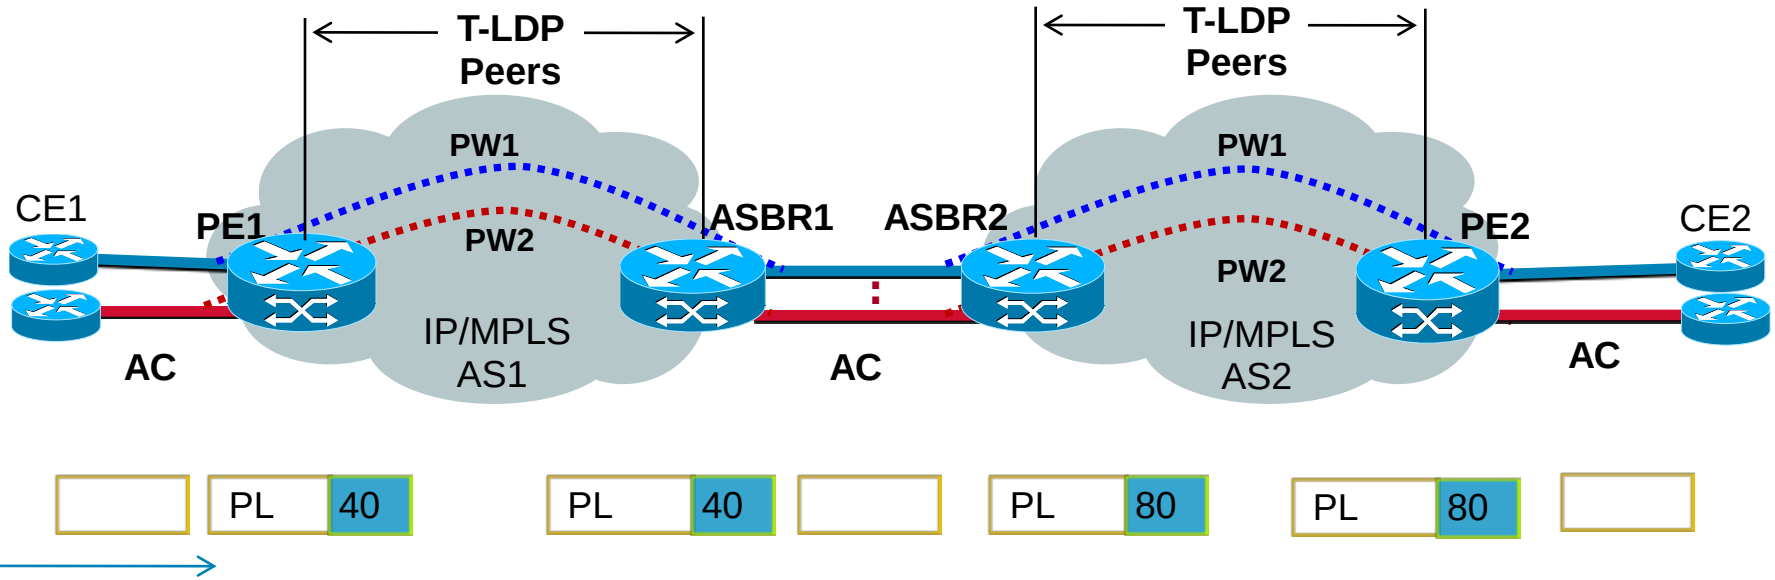

- Any Transport over MPLS is point-to-point L2VPN service
- One PW/AC (AC types: Ethernet, VLAN, PPP, ATM, TDM, FR, HDLC)
- Clear demarcation between ASs
- PE-ASBR exchange PW (VC) label
- Granular QoS control between ASBRs

Pseudowire

PW Label

PL = Payload

Inter-AS L2VPN
Multi-Hop PW using Option B

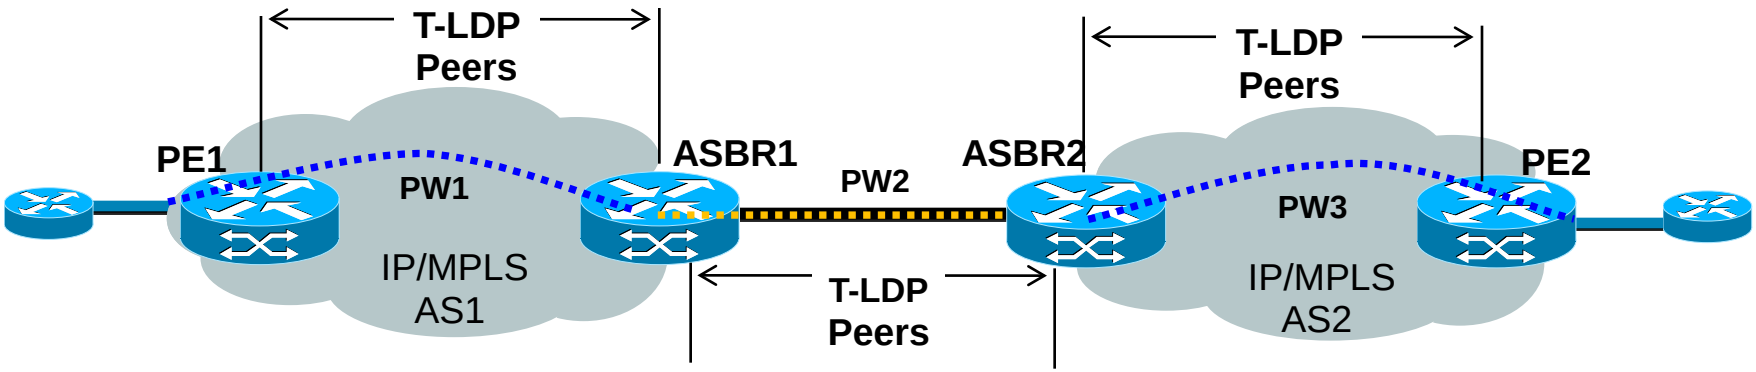

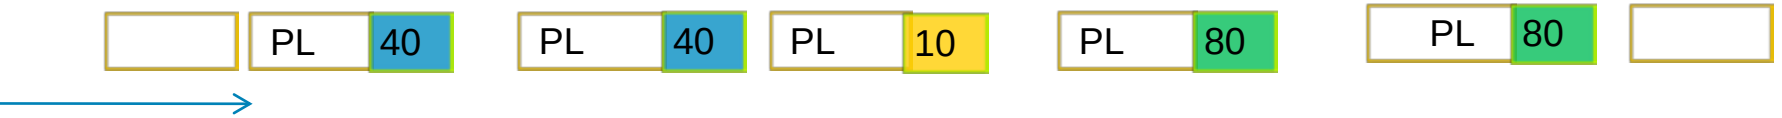

- PE and P devices do not learn remote PW endpoint addresses
- Only PW endpoint address (ASBR) leaked between ASs
- ASBRs swap PW (Virtual Circuit) Label

Inter-AS L2VPN Option B—Configuration

```
!
HOSTNAME PE1
!
interface giga1/0
 xconnect <ASBR1> 10 encapsulation mpls
 !
```

```
!
HOSTNAME PE2
!
interface giga1/0
 xconnect <ASBR2> 20 encapsulation mpls
!
```

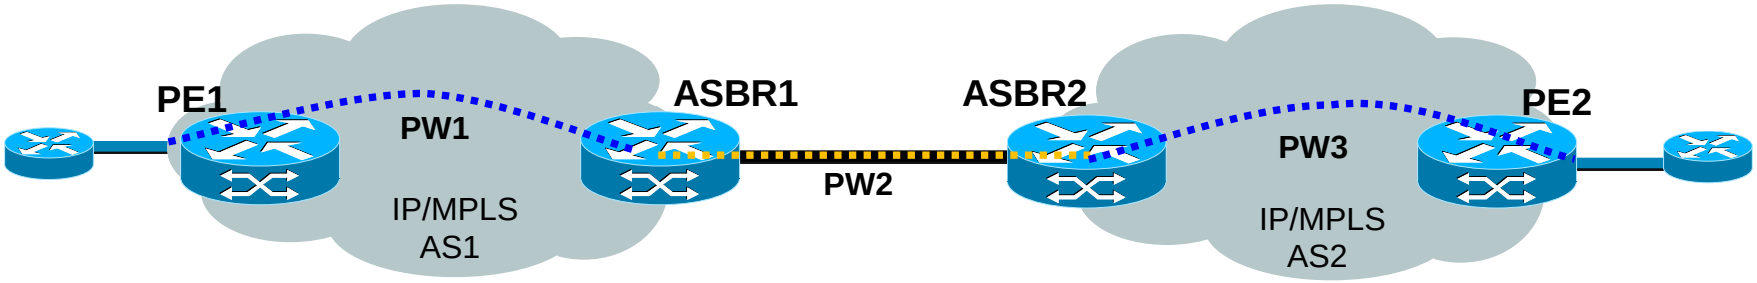


```
HOSTNAME ASBR1
!
pseudowire-class pw-switch
encapsulation mpls
!
l2 vfi pw-switch point-to-point
neighbor <ASBR2> 100 pw-class pw-switch
neighbor <PE3> 10 pw-class pw-switch
!
Interface giga3/0
mpls bgp forwarding
!
! router bgp 1
Neighbor <ASBR2-WAN> remote-as 2
exit-address-family
 !
*Also announce the loopback address (xconnect ID) of ASBR1 in IGP(AS1) and eBGP
```

```
HOSTNAME ASBR2
!
pseudowire-class pw-switch
encapsulation mpls
!
L2 vfi pw-switch point-to-point
neighbor <ASBR1> 100 pw-class pw-switch
neighbor <PE4> 20 pw-class pw-switch
!
Interface giga3/0
mpls bgp forwarding
!
router bgp 2
neighbor <ASBR1-WAN> remote-as1
exit-address-family
 !
*Also announce the loopback address of ASBR2 in IGP(AS2) and eBGP
```

Inter-AS AToM—Option C Single-Hop PW: BGP IPv4+label

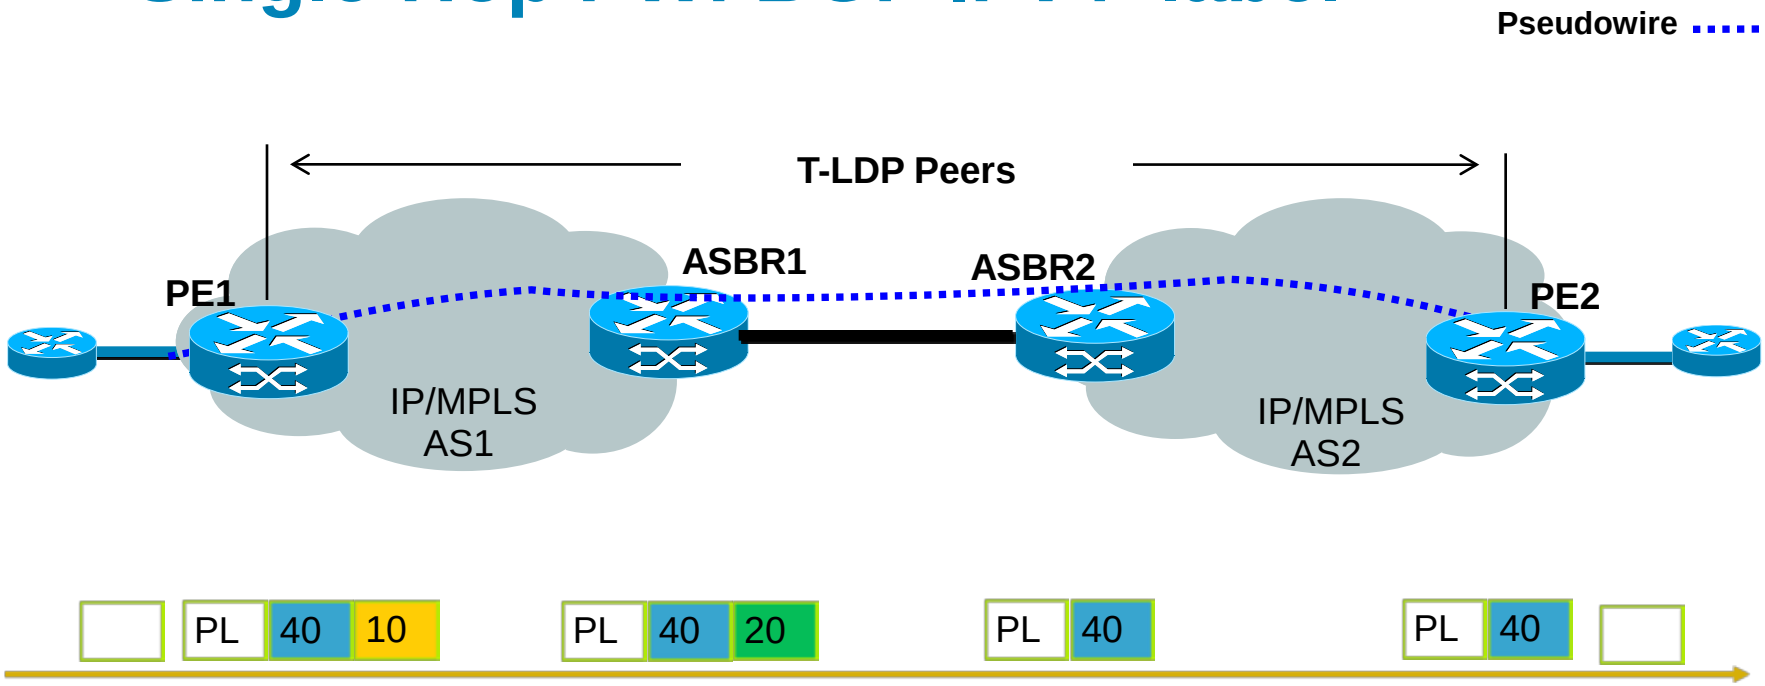

- Single physical interface between ASBRs
- PW endpoint addresses leaked between ASs using eBGP IPv4+label and distributed to PEs using iBGP IPv4+label
- PWs are not terminated on ASBRs

Inter-AS AToM Option C—Configuration

HOSTNAME PE3
!
interface Ethernet1/0
xconnect <PE4> 100
encapsulation mpls
!
! Activate IPv4 label capability !
router bgp 1
!
address-family ipv4
neighbor <ASBR-1> send-label
exit-address-family
!

HOSTNAME PE4
!
interface Ethernet1/0
xconnect <PE3> 100
encapsulation mpls
!
! Activate IPv4 label capability !
router bgp 2
!
address-family ipv4
neighbor <ASBR-2> send-label
exit-address-family
!

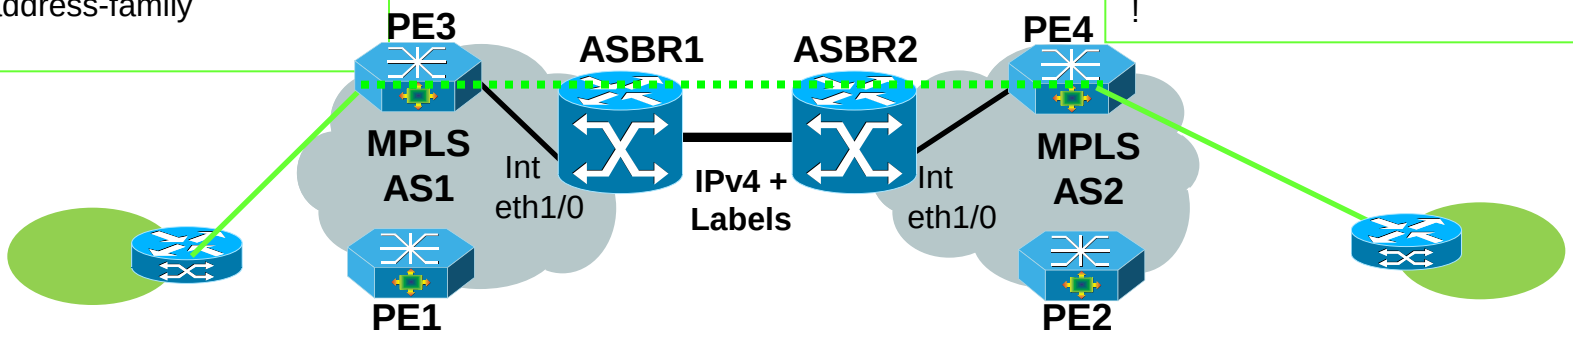

HOSTNAME ASBR1
! Activate IPv4 label capability !
router bgp 1
!
address-family ipv4
neighbor <PE3> send-label
neighbor <ASBR-2> send-label
exit-address-family
!

HOSTNAME ASBR2
! Activate IPv4 label capability !
router bgp 2
!
address-family ipv4
neighbor <PE4> send-label
neighbor <ASBR-1> send-label
exit-address-family
!

I-AS AToM Key Points

- All three I-AS models are supported to carry point-to-point PWs
- Transparently forwarding of data
- The control word negotiation results must match. The control word is disabled for both segments if either side doesn’t support it.
- Per-PW Quality of Service (QoS) is not supported.
- Attachment circuit inter-working is not supported.
- Traffic Engineering (TE) tunnel selection is supported.

Inter-AS L2 VPNs: VPLS

Overview

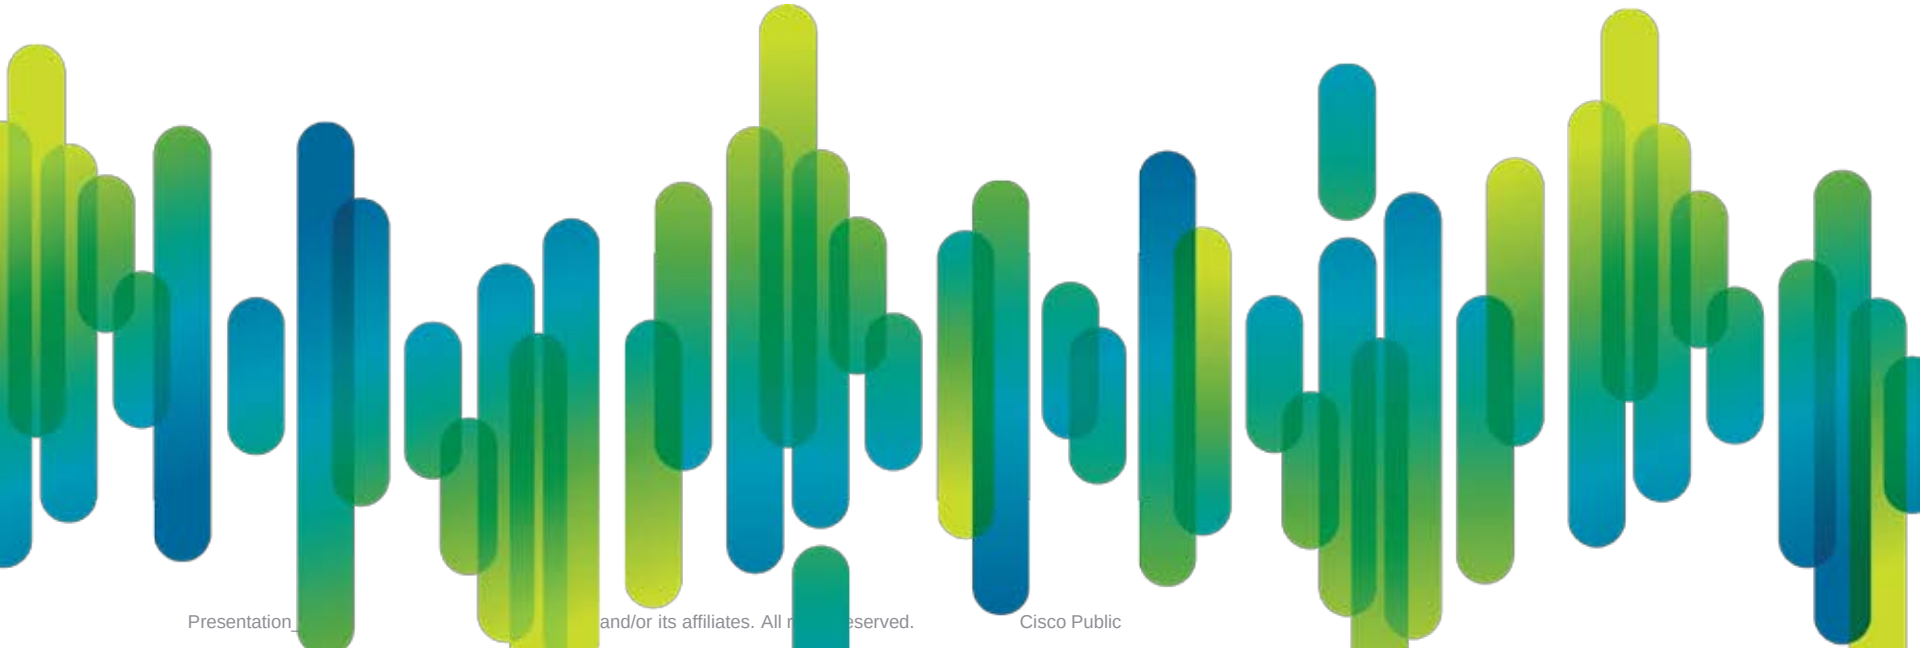

Virtual Private LAN Service Overview

- VPLS provides fully meshed L2 connectivity among VPN Sites
- VPLS VPN sites may span multiple Domains
- PEs aggregating VPN sites in both domains need transparency
- Option A, B and C supported to interconnect ASBRs
- For Option B, use a switching PW on ASBRs as discussed earlier

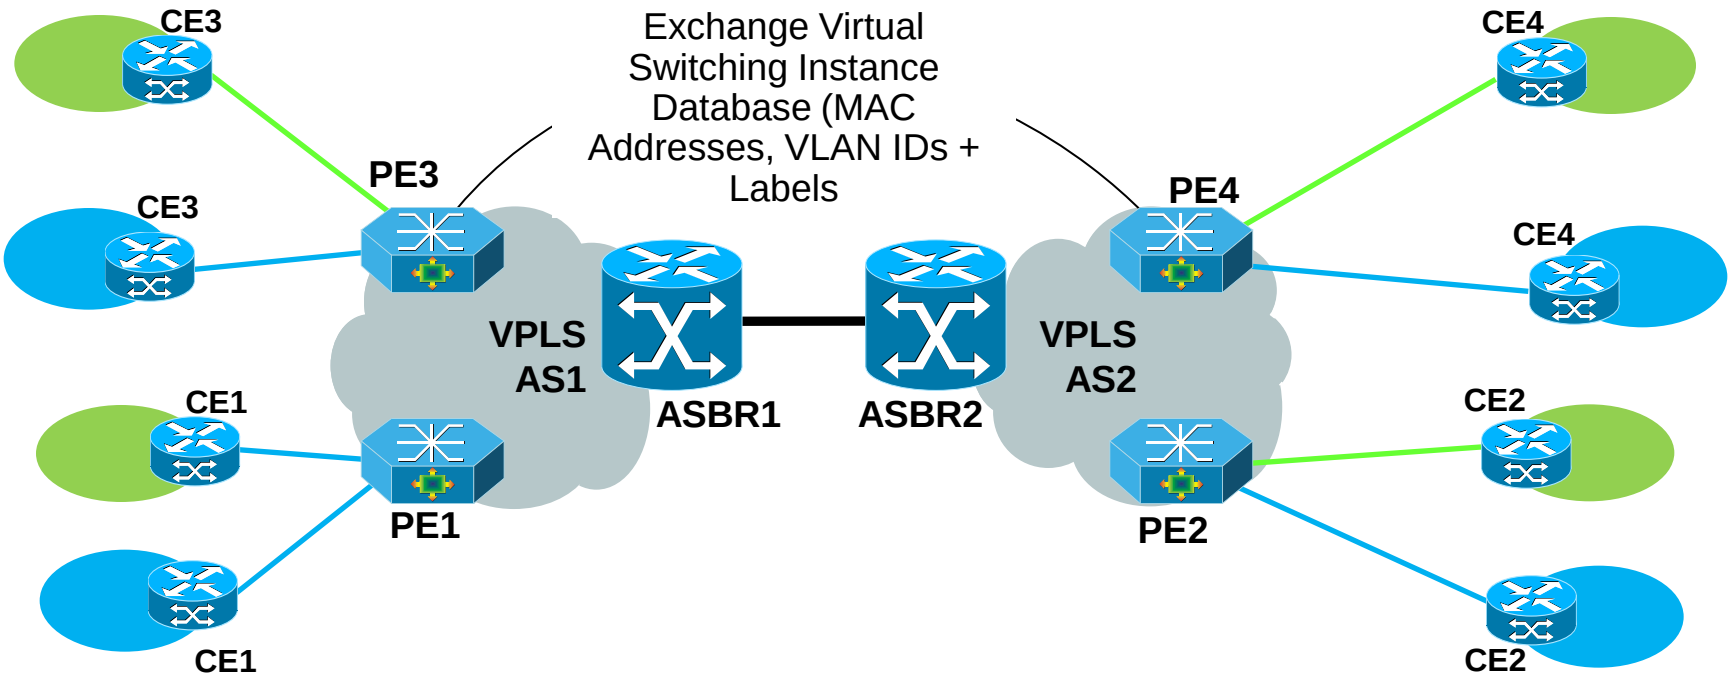

Inter-AS VPLS—Option C Single Hop Pseudowires

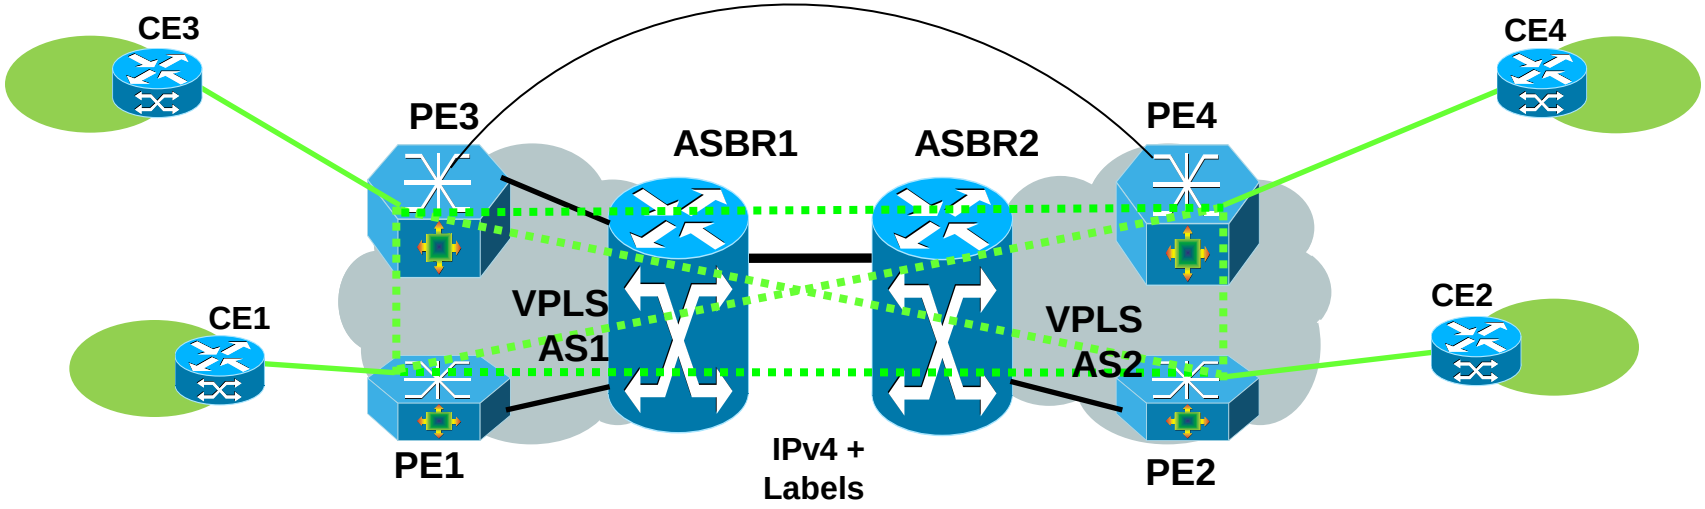

- Reachability between PEs is provided using eBGP+Labels (Option C discussed earlier)
- PWs are transported through ASBRs
- Targeted LDP session is formed between PEs
- Auto discovery of VPLS VPNs is supported using BGP
- Route Distinguisher, Route Target and VPN IDs are used similar way as in MPLS L3 VPNs
- RDs don’t have to match across different domains for the same VPLS VPN sites

VPLS BGP Auto Discovery with Inter-AS Option C—Cisco IOS Configuration

```
! Setup VPLS instance, Define discovery method and set vpn iD !
HOSTNAME PE3
!
l2 vfi customer1 autodiscovery
 vpn id 100

! Activate IPv4 label capability !
router bgp 1
!
address-family ipv4
neighbor <ASBR-1> send-label
exit-address-family
 !
```

```
! Setup VPLS instance, Define discovery method and set vpn iD, vpls-id and RT to match the other side.
HOSTNAME PE4
!
l2 vfi customer1 autodiscovery
 vpn id 100
 rd 1:100
Route-target both 1:100

! Activate IPv4 label capability !
router bgp 2
!
address-family ipv4
neighbor <ASBR-2> send-label
exit-address-family
 !
```

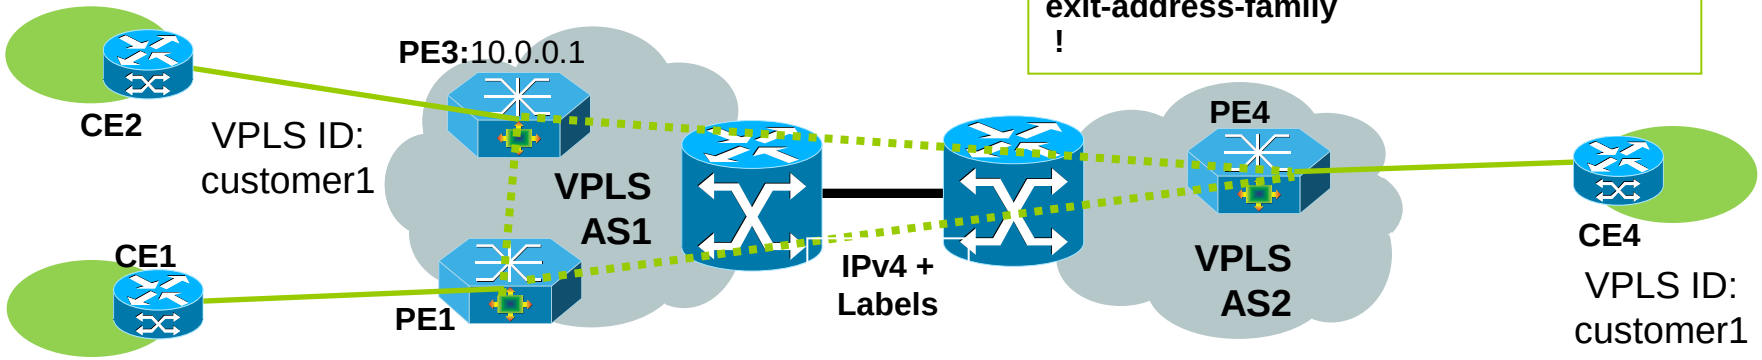

1. PE3 sends this packet to PE4: 1:100:10.0.0.1/96 RT 1:100 VPLS-id 1:100
2. L2 Subsystem on PE4 decodes it: VPN ID:100, Neighbor LDP-ID: 10.0.0.1 (=NH)
3. PWs are setup using directed LDP session among PEs

Inter-AS mVPNs

Overview

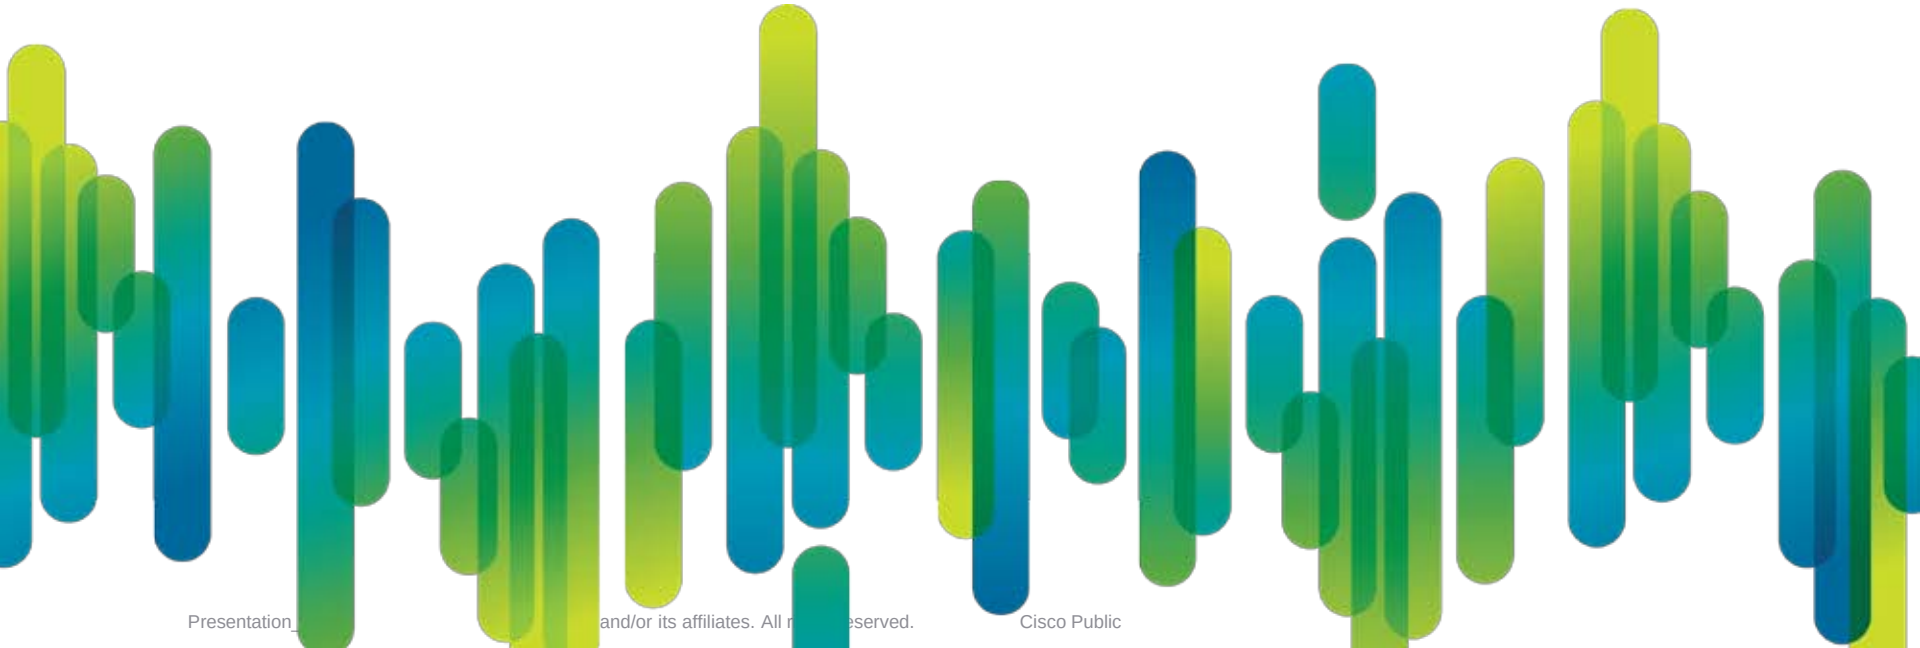

mVPN Concept and Fundamentals—Review

- CEs join MPLS Core through provider's PE devices
- PEs perform RPF check on Source to build Default and Data Trees (Multicast Data Trees – MDT)
- Interfaces are associated with mVRF
- Source-Receivers communicate using mVRFs

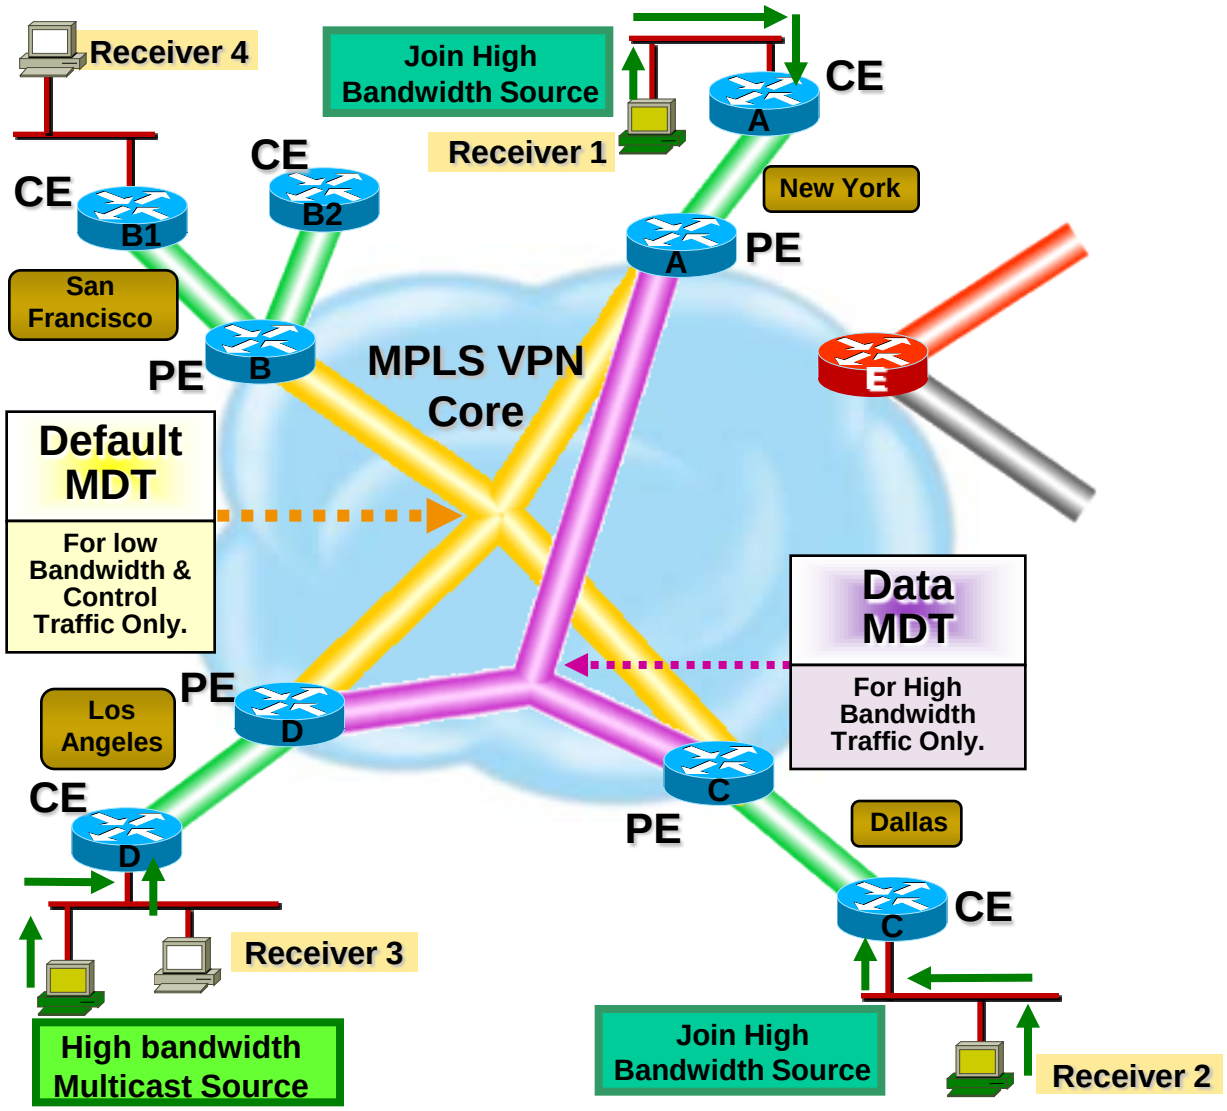

I-AS mVPN Requirements

Challenge: Setup Multicast Data Trees across ASs

- To form the Default MDT, PE routers must perform an RPF check on the source
- The Source address is not shared between ASs

Solution:

- Support reverse path forwarding (RPF) check for I-AS sources – P and PE devices
- Build I-AS MDTs

RPF Check with Option B and Option C

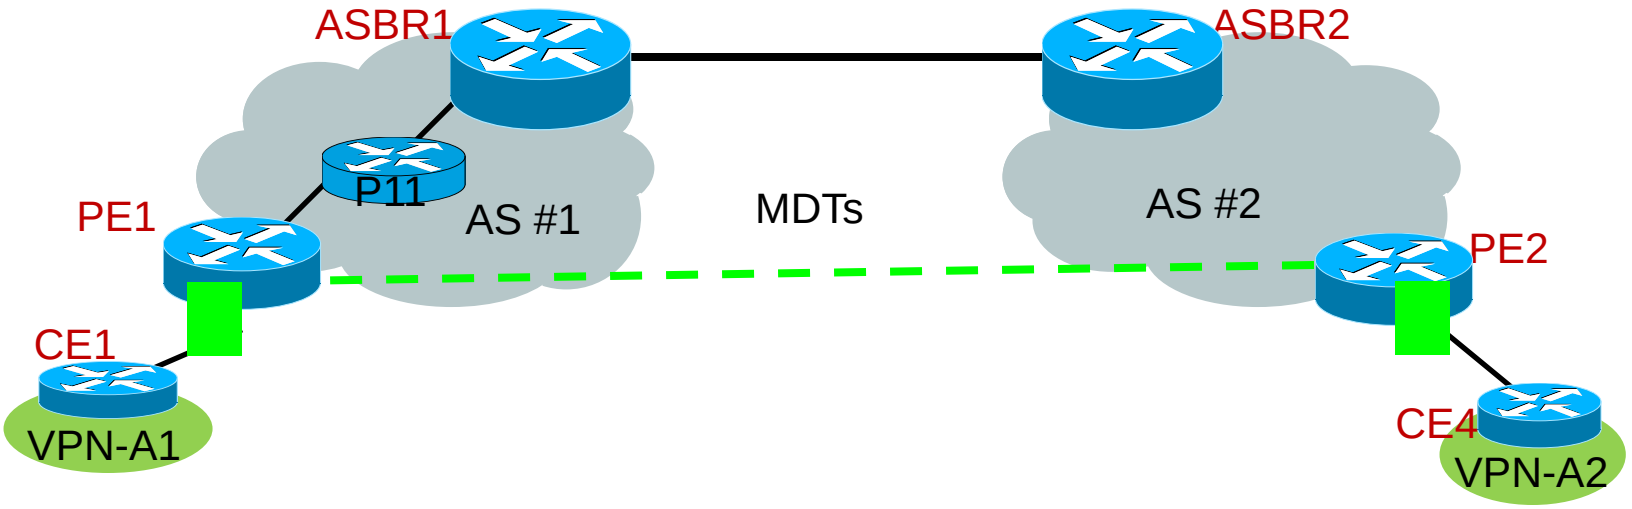

- Two new components:

  BGP Connector Attribute (Originating PE) & PIM RPF Vector (ASBR1 in AS1)

- For Option B(eBGP between ASBRs): Use BGP Connector Attribute to RPF to source that is reachable via PE router in remote AS

  Preserves identity of a PE router originating a VPNv4 Prefix

  Receiving PEs in the remote AS use RPF Connector to resolve RPF

- For Option B and C: Use PIM RPF Vector to help P routers build an I-AS MDT to Source PEs in remote AS

  Leverage BGP MDT SAFI on ASBRs and receiver PEs to insert the RPF Vector needed to build an I-AS MDT to source PEs in remote ASs

I-AS MVPN MDT Establishment for Option B

From ASBR1 to PE1
RD 1:1, Prefix 152.12.4.0/24
NH ASBR1, CONN PE2

BGP MDT SAFI Update
RD1:1, Prefix PE2, MDT
232.1.1.1, NH ASBR1

From ASBR2 to ASBR1
RD 1:1, Prefix 152.12.4.0/24
NH ASBR2, CONN PE2

BGP MDT SAFI Update
RD1:1, Prefix PE2, MDT
232.1.1.1, NH ASBR2

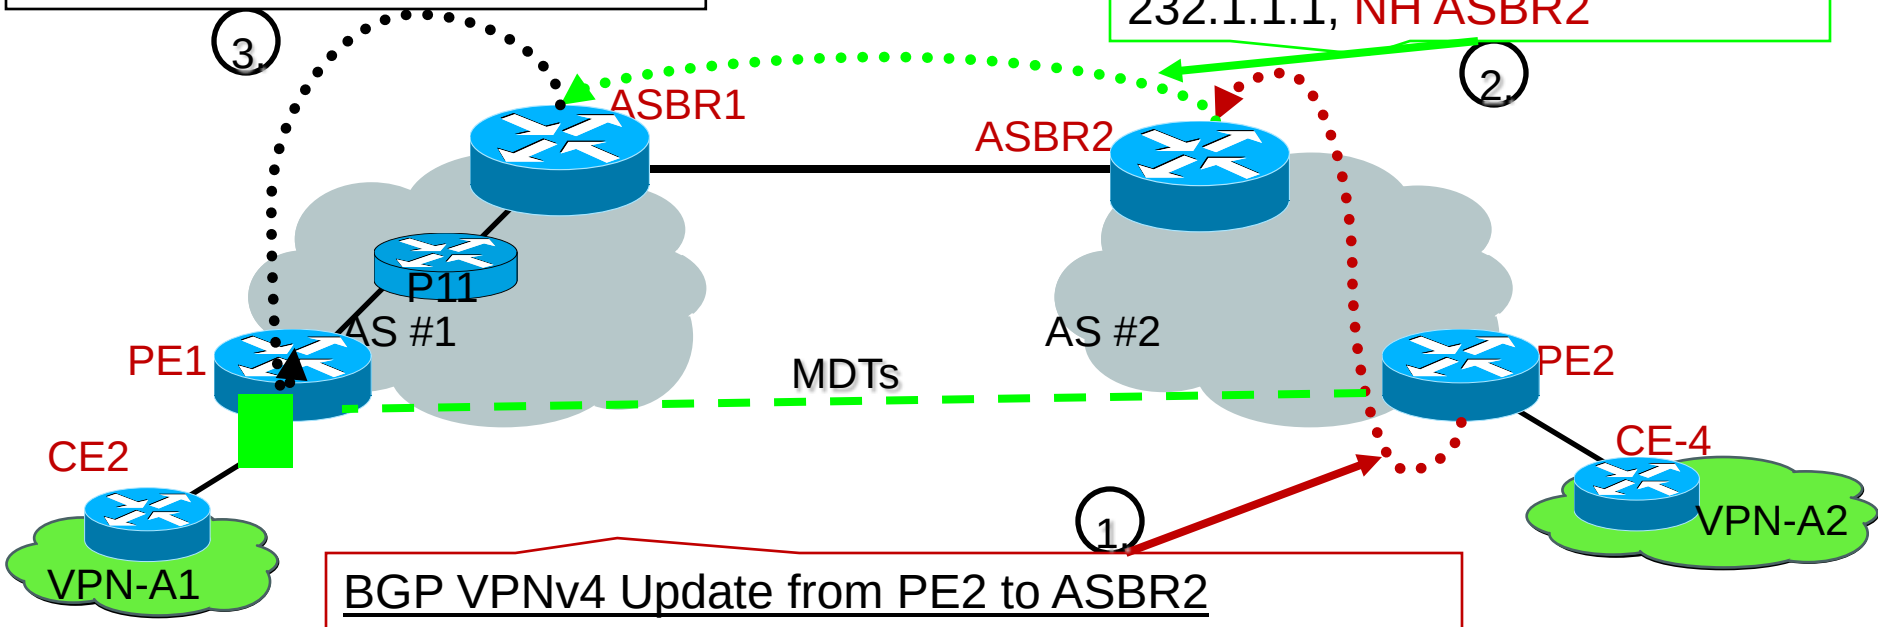

BGP VPNv4 Update from PE2 to ASBR2
RD 1:1, Prefix 152.12.4.0/24 NH PE2, CONN PE2

BGP MDT SAFI Update (Source and Group)
RD1:1, Prefix PE2, MDT 232.1.1.1, NH PE2

Carrier Supporting Carrier

Overview

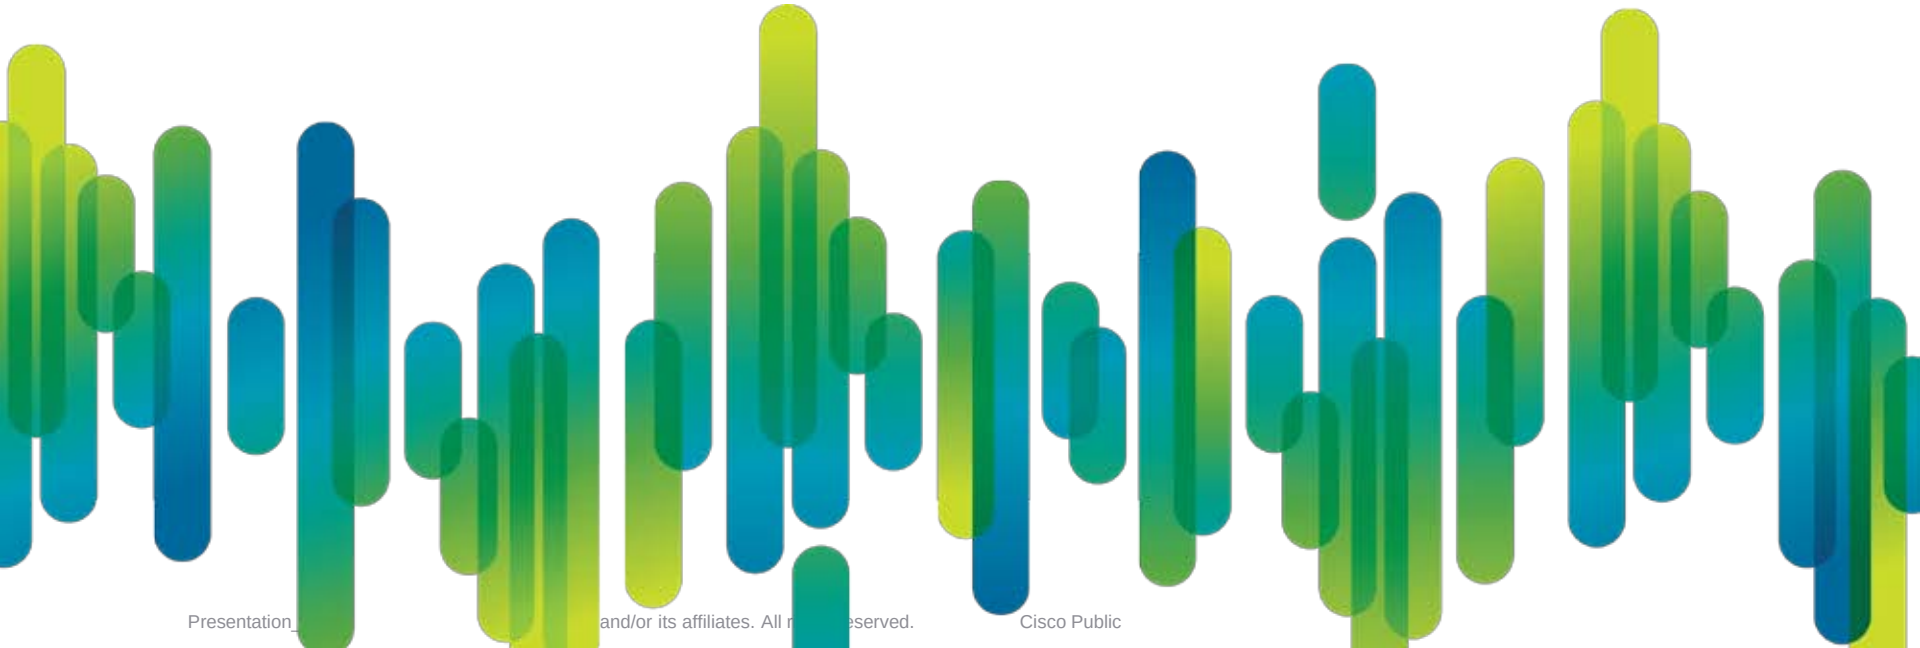

Introducing Carrier Supporting Carrier

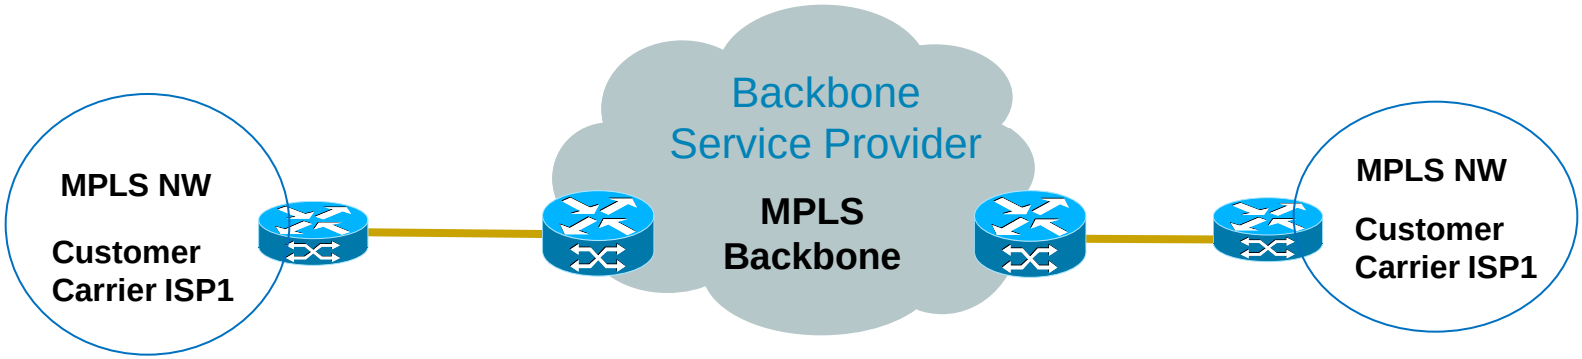

- CSC is one of the VPN services that is applicable in a Multi-AS network environment
- CSC VPN service is a VPN service that provides MPLS transport for customers with MPLS networks
- It is also known as hierarchical MPLS VPN service since MPLS VPN customer carrier subscribes MPLS VPN service from an MPLS Backbone provider
- Defined in RFC 4364. (previously well know by draft 2547biz)

Why Carrier Supporting Carrier?

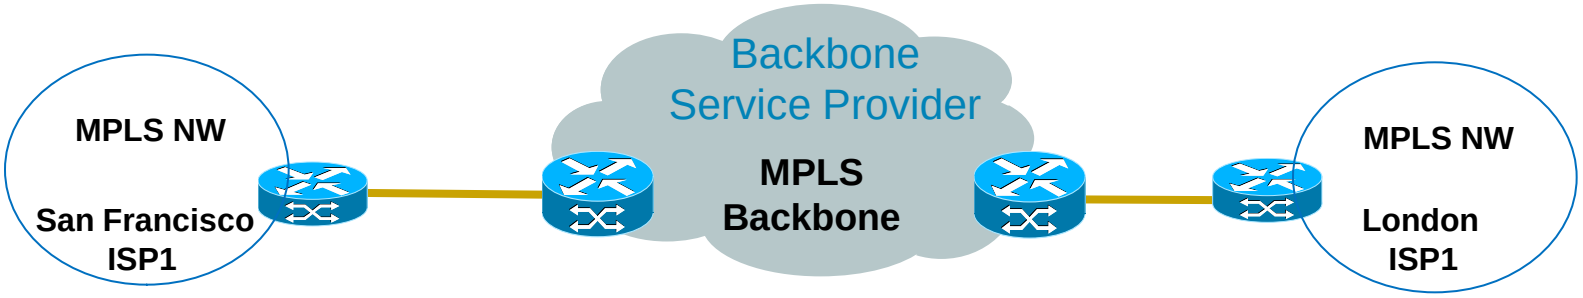

- MPLS VPN services offerings by an MPLS VPN backbone provider to customers with MPLS networks
  - Provide business continuity by extending segmented networks
  - Customer networks include ISP, Carriers, or other enterprise networks
- Address scalability issues at the PEs
  - MPLS-VPN works well for carrying customer IGPs
  - Reduce #s of VPN routes carried by a PE by using hierarchical model
  - Platforms, network scale to N*O(IGP) routes: Internet Routes
  - Separate Carrier's Internal routes from external routes eliminating the need to store customer's external routes

Carrier's Carrier building blocks

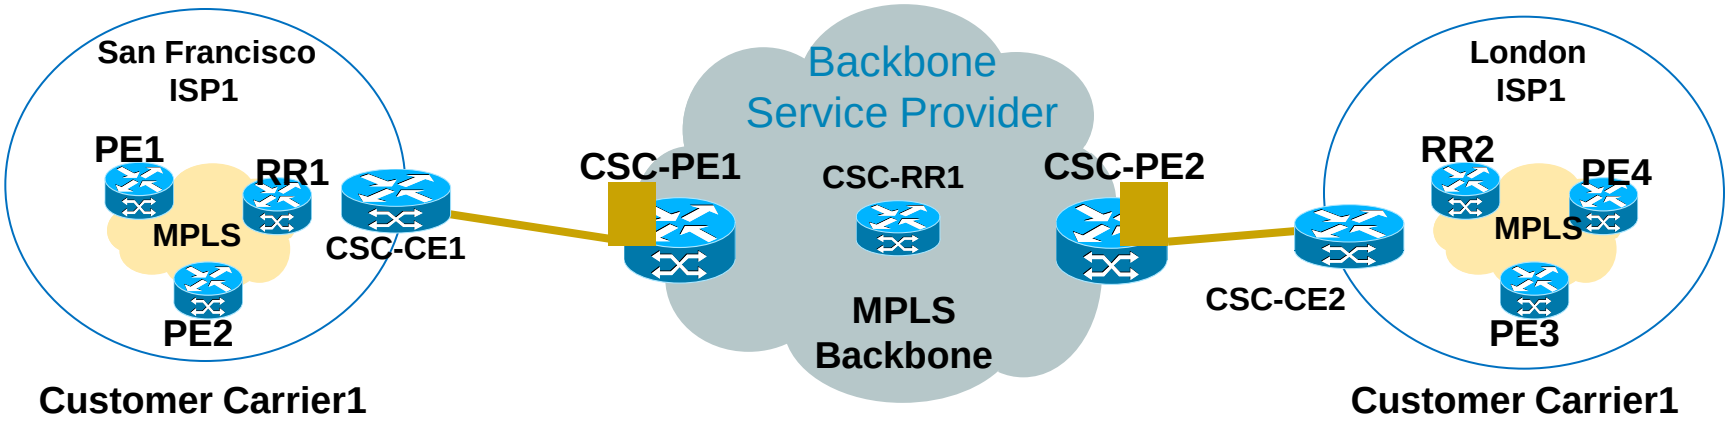

- MPLS MPLS-VPN enabled Carrier's backbone
- CSC-PE: MPLS VPN PEs located in backbone Carrier's Core
- CSC-CE: Located at the Customer Carrier (ISP/SPs/Enterprise) network edge and connects to a CSC-PE
- PE: located in Customer carrier networks & carries customer VPN routes
- CSC-RR: Route Reflectors located in MPLS Backbone provider network
- RR: Route Reflectors located in Customer Carrier Network
- MPLS Label exchange between Carrier's PE & ISP/SPs CE

Carrier's Carrier building blocks (continue)

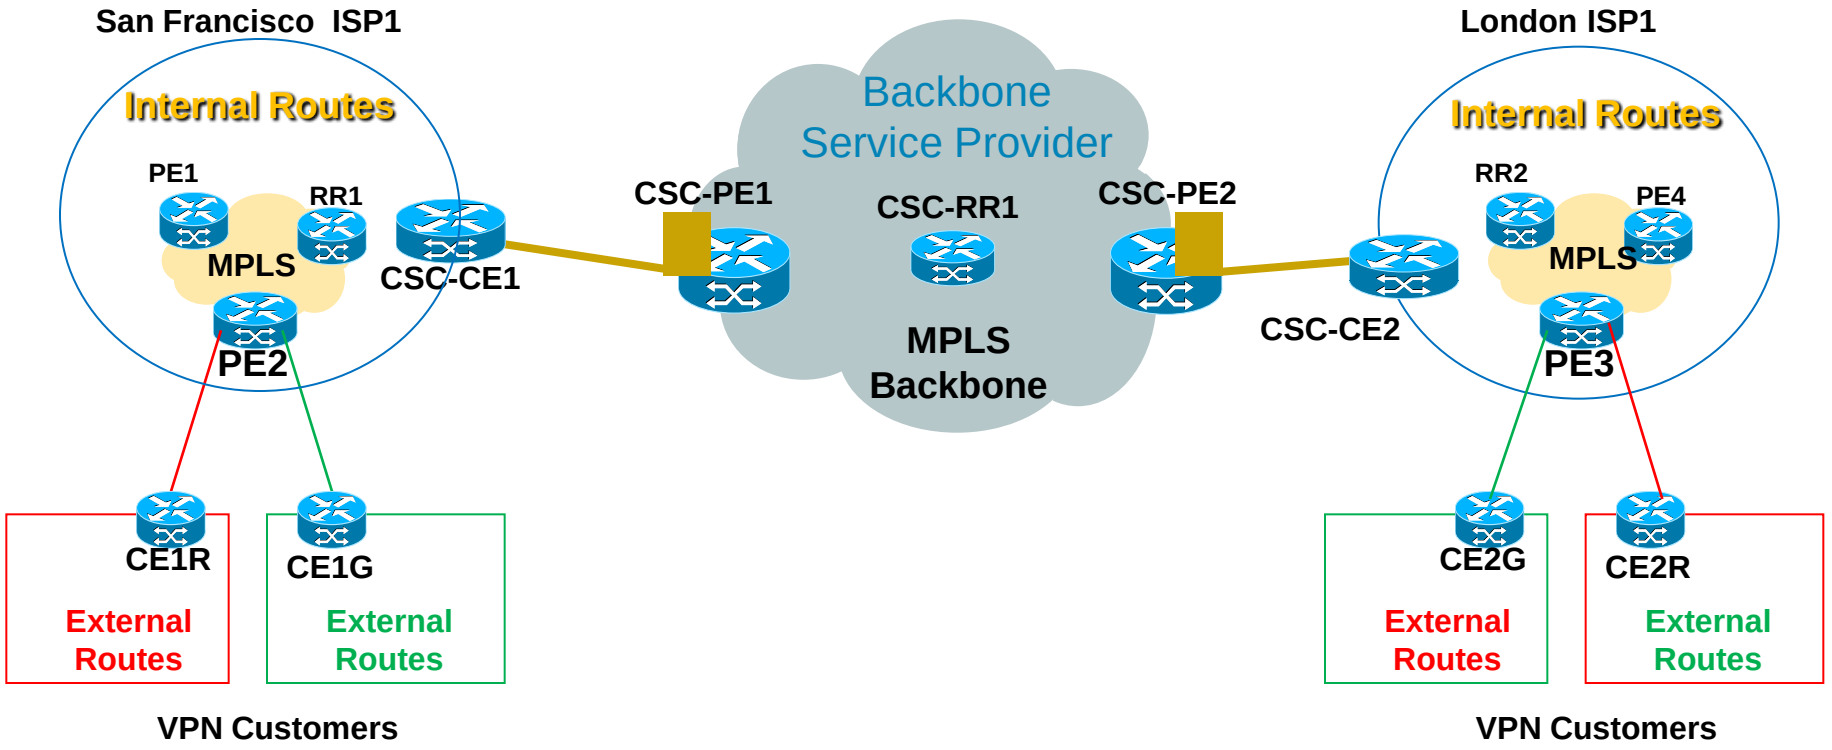

- External Routes: IP routes from VPN customer networks
- Internal Routes: Internal routes (global table) of Customer Carrier network
- External routes are stored and exchanged among Customer Carrier PEs
- MPLS Backbone network doesn't have any knowledge of external routes
- Customer Carrier selectively provides NLRI to MPLS VPN backbone provider

CSC Building Blocks (continue)

- Control Plane configuration is similar to single domain MPLS VPN
- CSC-CE to CSC-PE is a VPN link to exchange Customer Carrier's internal routes. These routes are redistributed into the BSP's CSC-PE using:

  1. Static Routes OR 2. Dynamic IGP OR 3. eBGP

- Customer Carriers don't exchange their Subscribers' (external) VPN routes with the Backbone Service Provider
- CSC-PE-to-CSC-CE links extend Label Switching Path using:
  1. IGP+LDP
  2. eBGPv4 + Labels

Carrier's Carrier building blocks (continue)

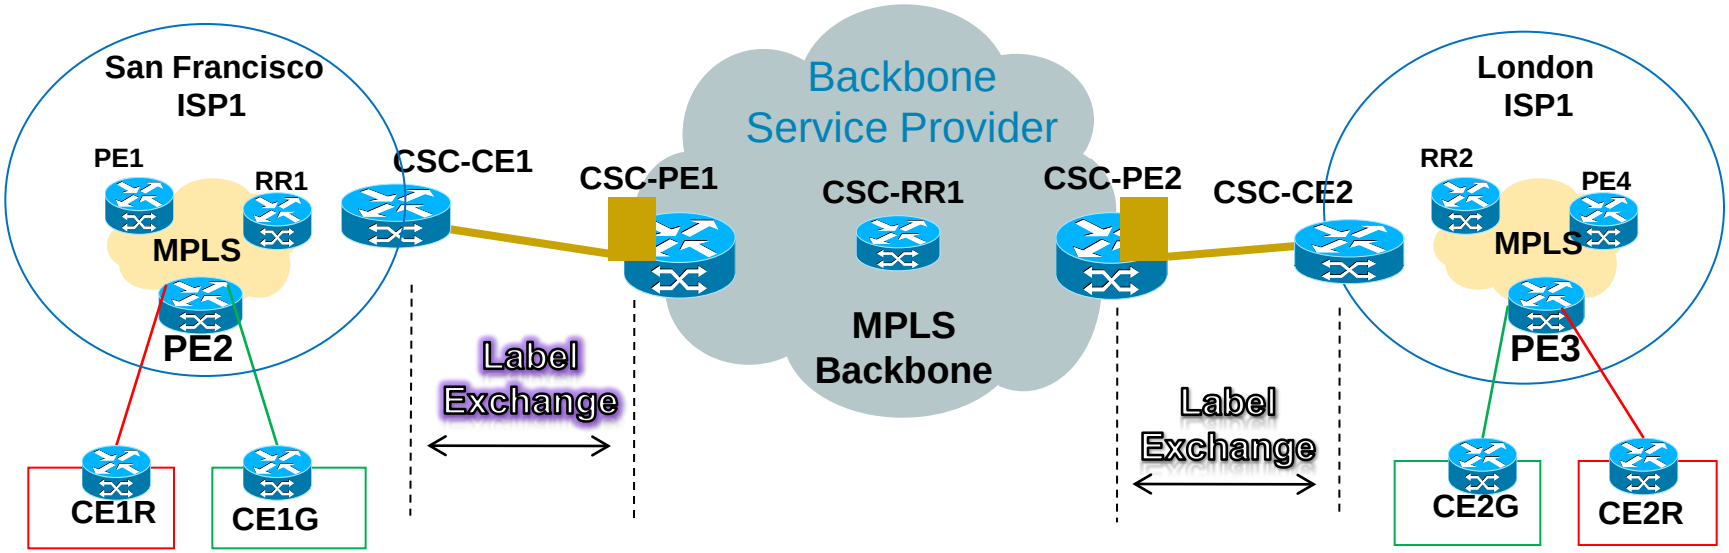

- Label Switched paths between CSC-CE and CSC-PE
- CSC-PE and CSC-CE exchange MPLS Labels

  -this is necessary to transport labeled traffic from a Customer Carrier
- IP between CE and PE for VPN customers

Carrier Supporting Carrier Models

1. Customer Carrier Is Running IP Only

   -similar to basic MPLS L3 VPN environment

2. Customer Carrier Is Running MPLS

   -LSP is established between CSC-CE and CSC-PE

   -Customer carrier is VPN subscriber of MPLS VPN backbone provider

3. Customer Carrier Supports MPLS VPNs

   -LSP is established between CSC-CE and CSC-PE

   -Customer carrier is VPN subscriber of MPLS VPN backbone provider

   -True hierarchical VPN model

CSC Model III
Customer Carrier Supports MPLS VPNs

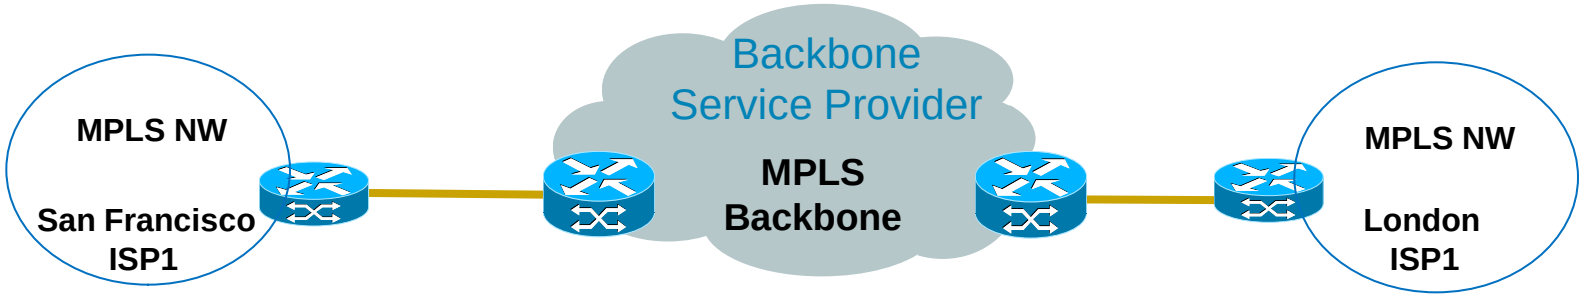

- LSP is extended to CSC-PE, CSC-CE advertises labels for internal routes to CSC-PE; CSC-PE1 performs imposition for site VPN label and IGP label
- PE swaps the site IGP label with a BB VPN label and push IGP label; PHP is now extended to inside of site 2
- External and VPNv4 routes are carried by MP-BGP between customer carrier sites
- CSC-CE and CSC-PE exchange labels using IGP+LDP or eBGP+Label

CSC Model III Routing Exchange

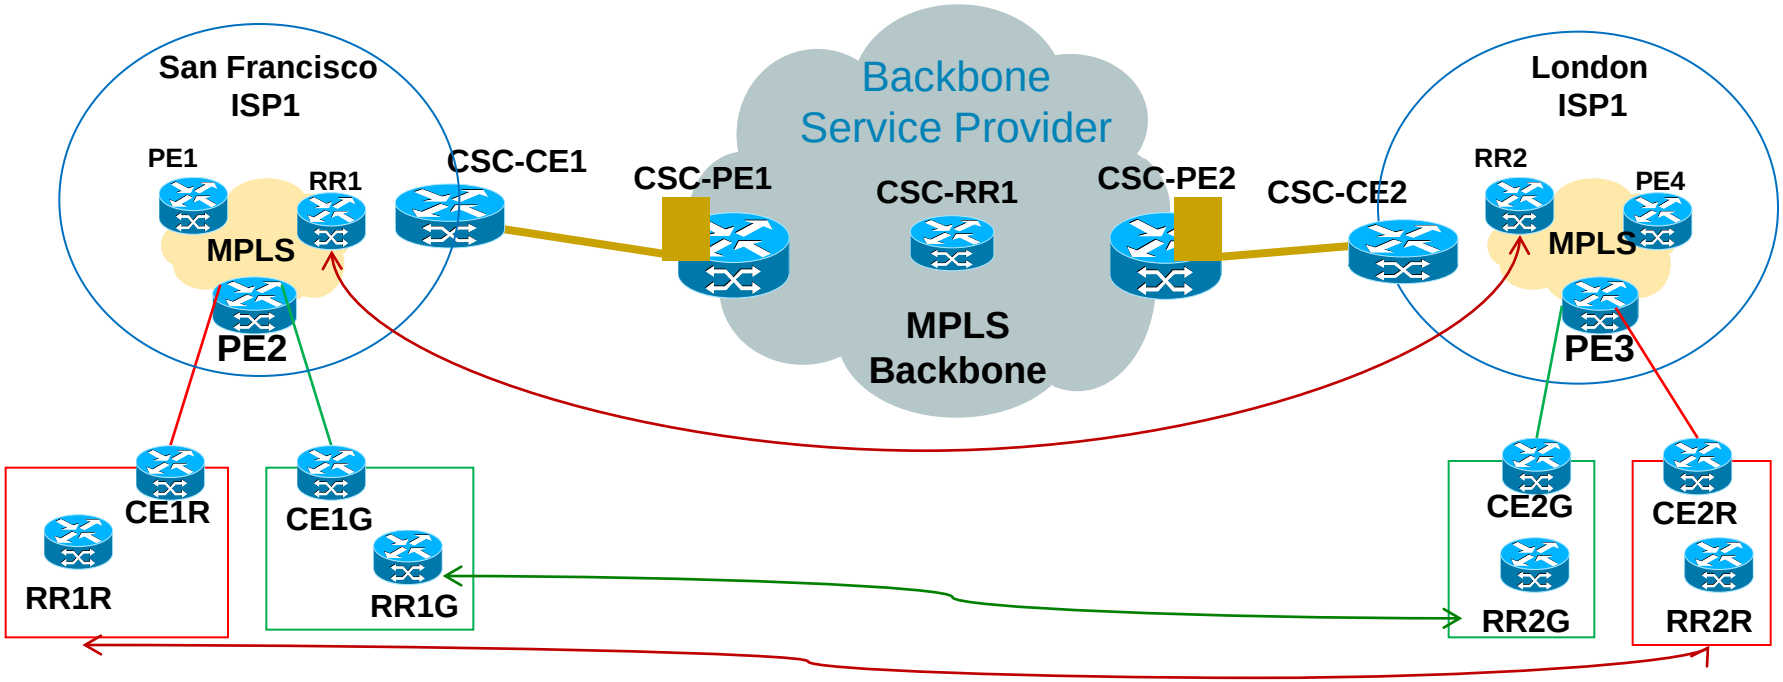

- RR1R and RR2R exchange Red VPN site routes
- RR1 and RR2 exchange ISP1 site routes
- CSC-RR1 updates CSC-PEs
- ISP1 adds Subscriber VPN Label which is removed by the remote ISP1 VPN site
- Backbone CSC-PE1 adds backbone VPN label which is removed by backbone CSC-PE2

CSC Model III—Customer Carrier Supports MPLS VPNs

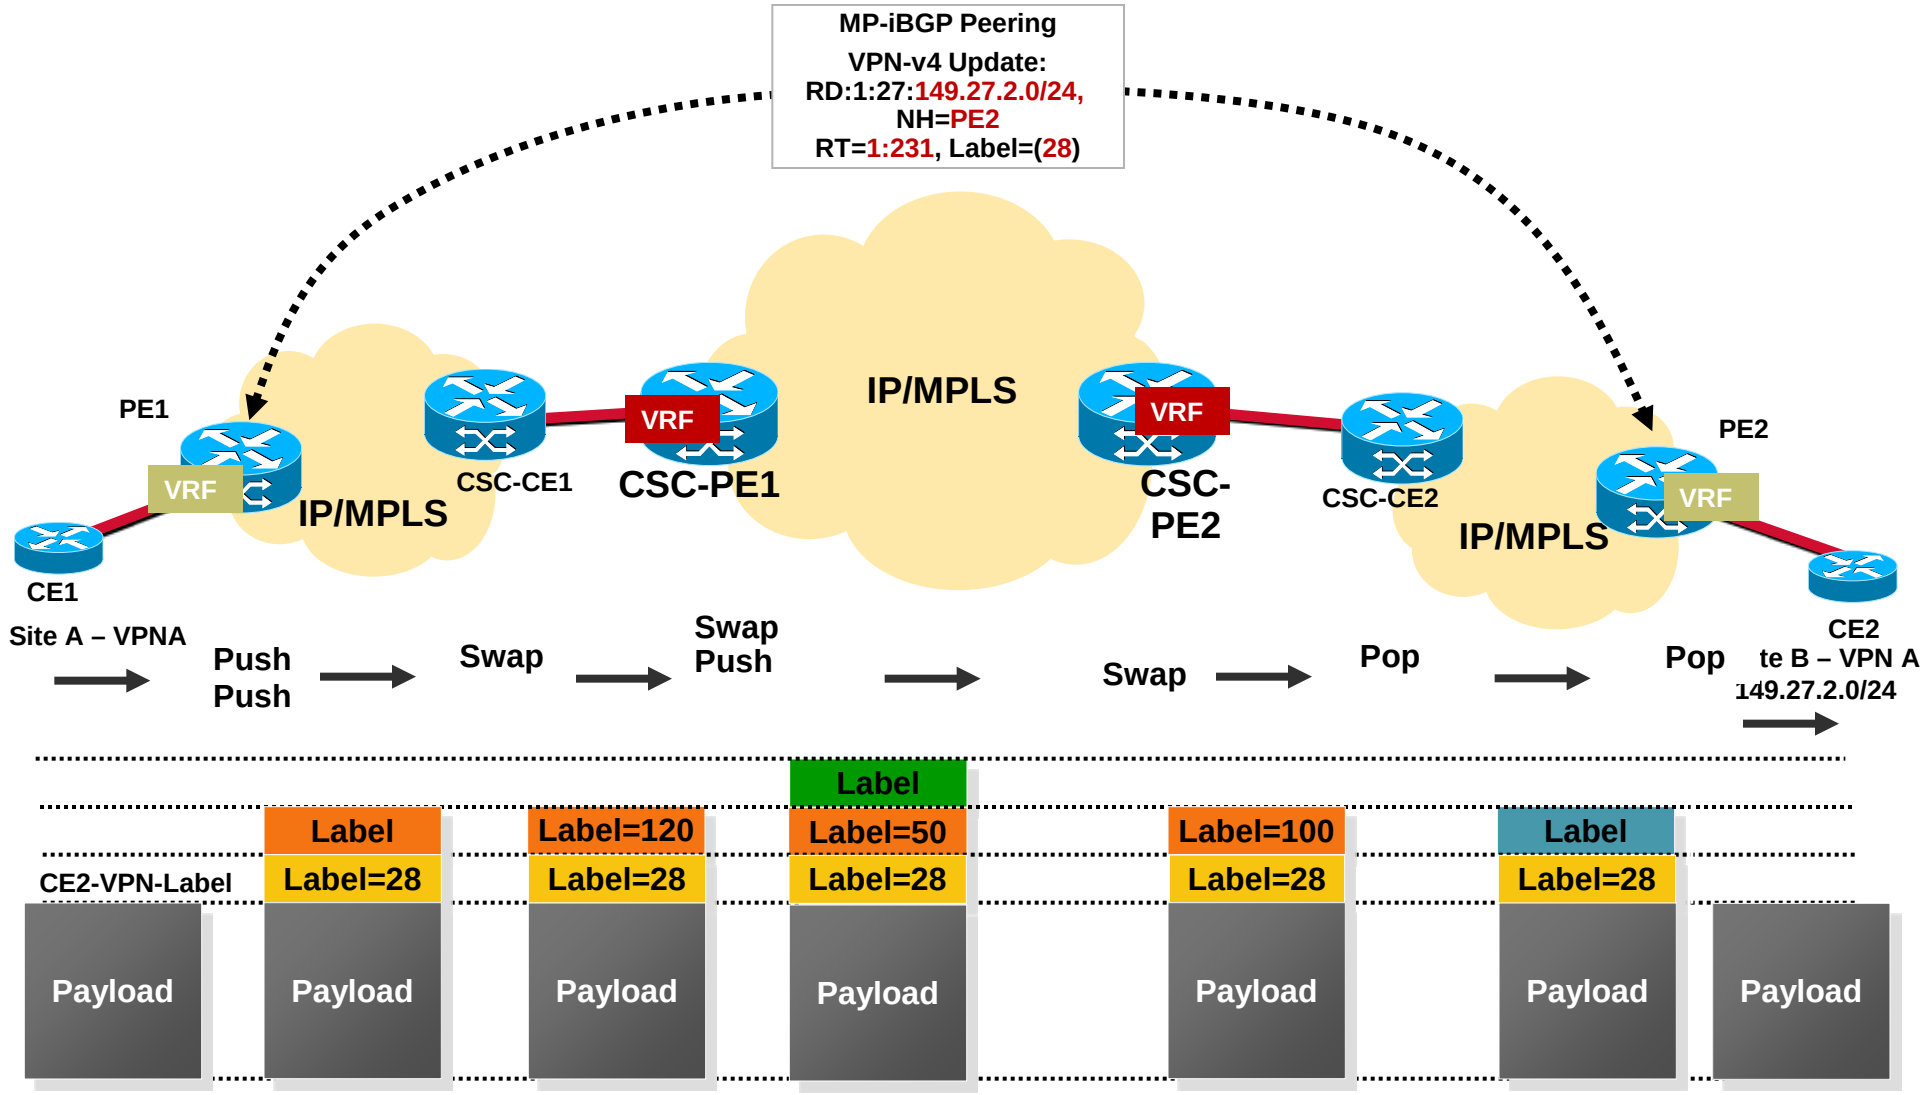

CSC Model III with IPv4+Label Cisco IOS Configuration

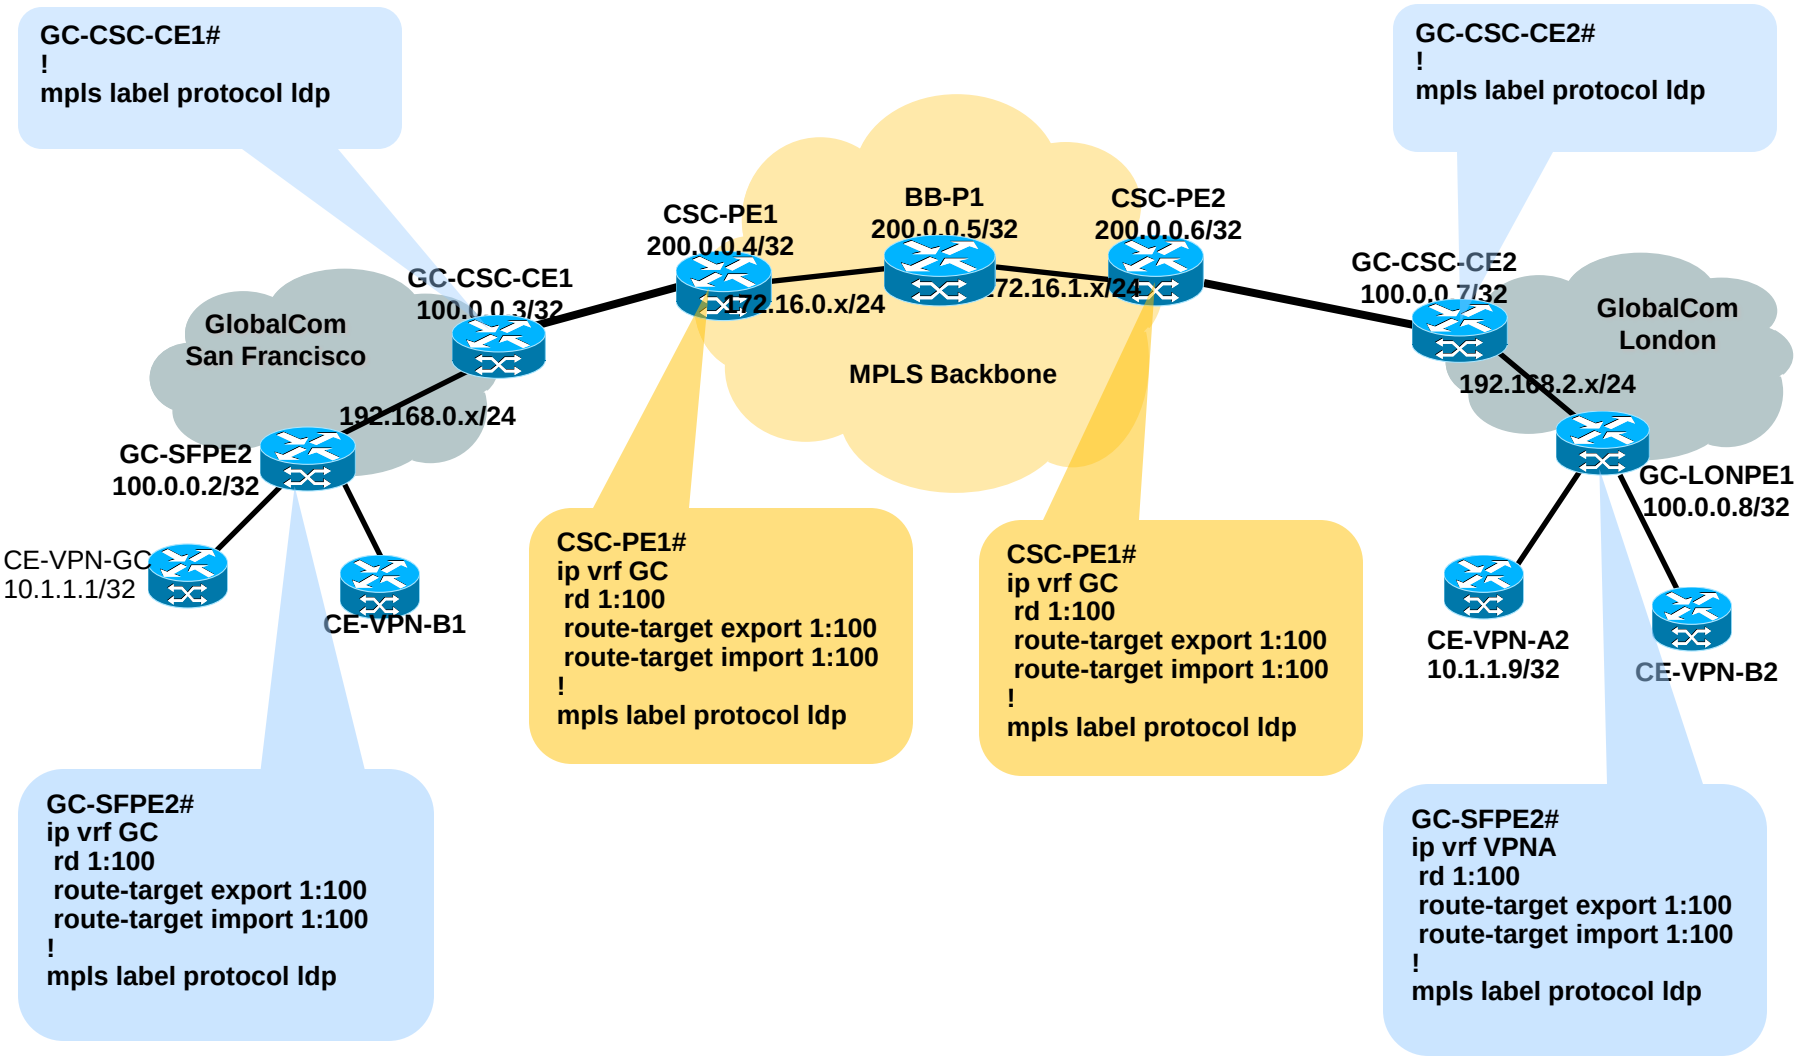

MPLS L2VPNs Across a CSC Network

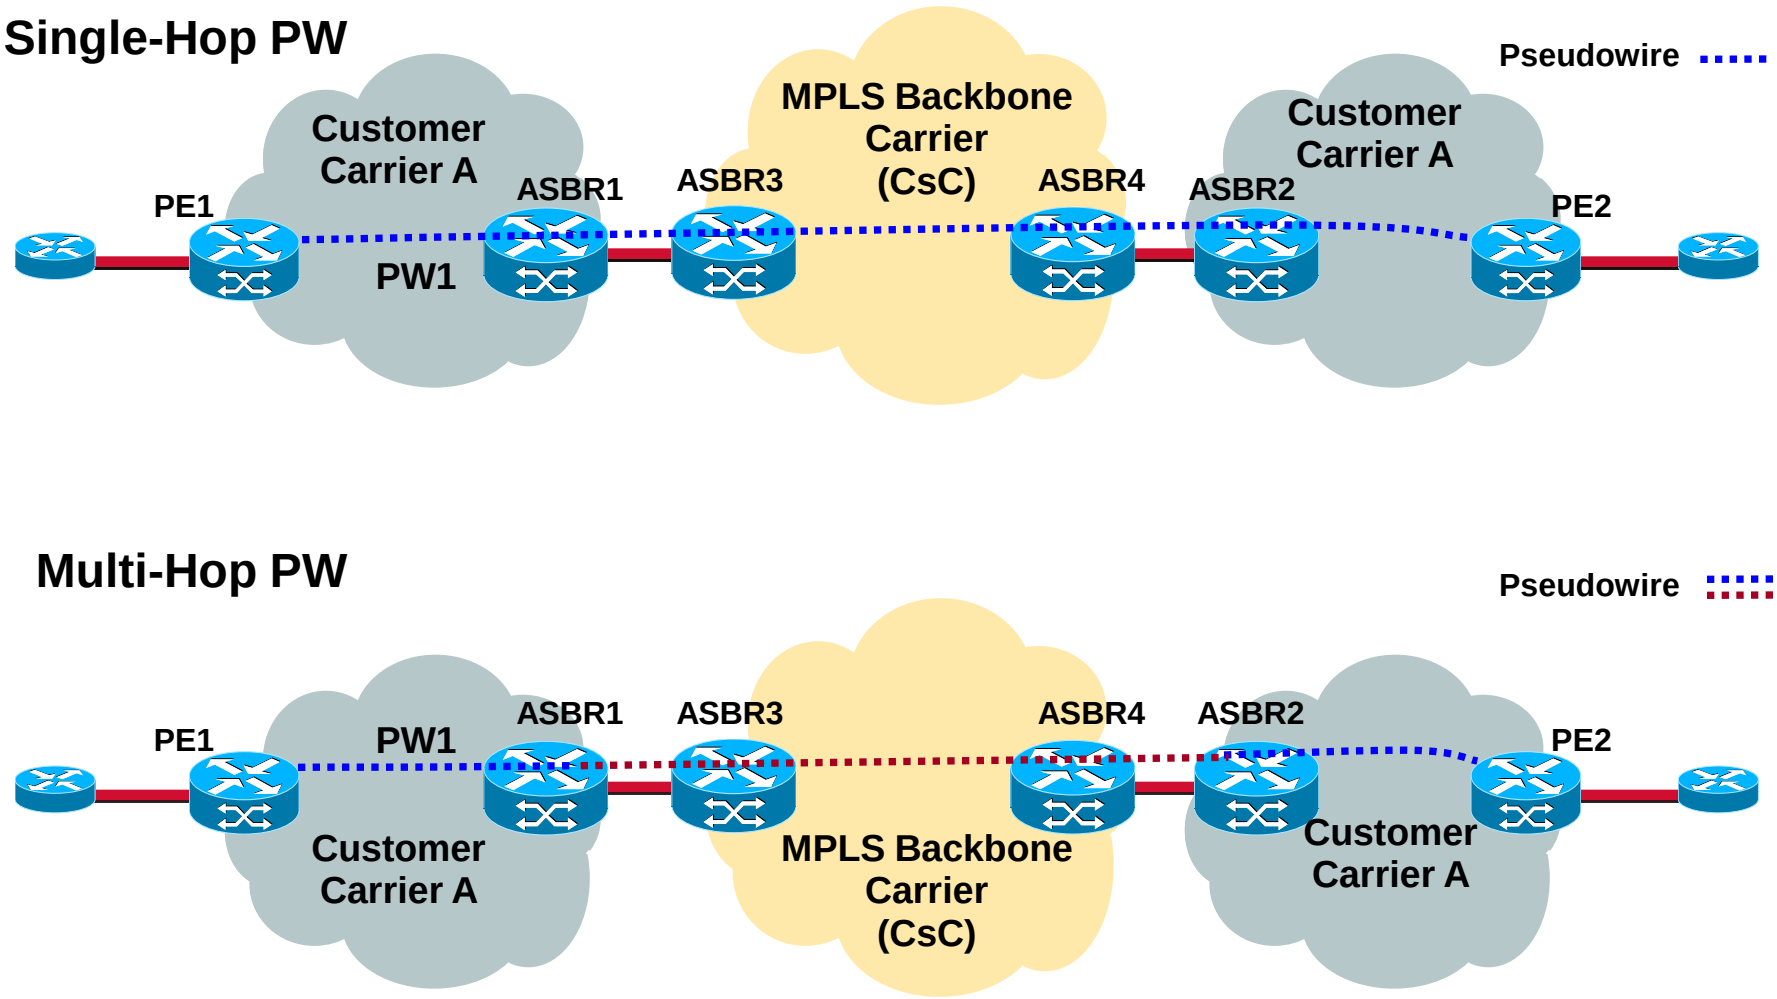

CSC Security Elements

- MD5 authentication on LDP/BGP sessions
- Applying max prefix limits per VRF
- Use of static labels between CSC-CE and CSC-PE
- Route Filtering

  …Customer Carrier may not want to send all the internal routes to MPLS VPN backbone provider…

  Use Route-maps to control route distribution & filter routes

  Use match and set capabilities in route-maps

CSC Summary (1)

- CSC supports hierarchical VPNs
- VPNs inside customer carrier's network are transparent to the backbone MPLS VPN Service Provider
- QoS will be honored based on MPLS EXP bits between CSC-CE and CSC-PE
- Granular QoS policies should be pre-negotiated and manually configured
- Additional supported Services over CSC
  - MPLS IPV6 VPNs
  - Multicast VPNs
  - MPLS L2 VPNs
  - MPLS TE

Best Practice Recommendations

- Do not use Static default routes on CSC-CE

  End-End LSP is required across the VPN and MPLS VPN backbone

- Use dynamic protocol instead of static on CSC-CE – CSC-PE link preferably eBGP+IPv4 Labels
- Set Next-Hop-Self on PEs carrying external routes
- If using IGP on CSC-CE routers, use filters to limit incoming routes from the CSC-PE side
- If using RRs in customer carrier network, set next-hop-unchanged on RRs

I-AS RSVP TE

Overview

How MPLS TE Works in a Single Domain

1. Head-end learns network topology information using:
   - ISIS-TE
   - OSPF-TE
   - full view of the topology

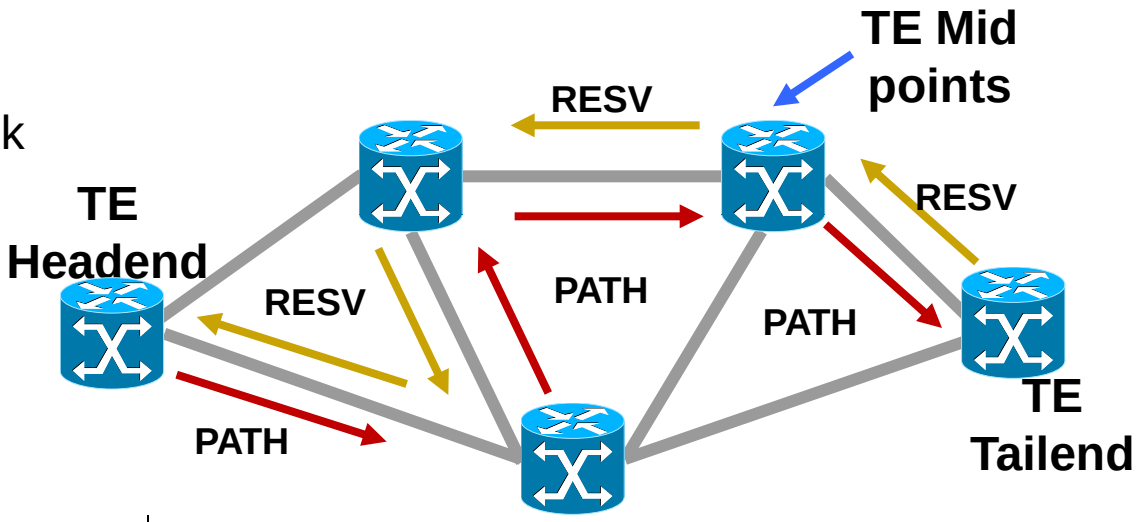

2. Path Calculation (CSPF)
3. Path Setup (RSVP-TE):
   - Label_Request (PATH)
   - Label (RESV)
   - Explicit_Route Object
   - Record_Route (Path/RESV)
   - Session_Attribute (Path)
4. LFIB populated using RSVP labels

5. Packets forwarded onto a tunnel via:
   - Static routed
   - Autoroute
   - Policy route
   - CBTS
   - Tunnel Select
   - Forwarding Adjacency
6. Packets follow the tunnel LSP and Not the IGP LSP

Inter-Domain Traffic Engineering

Challenge:

- Head end and Tail end are located in different domains
- IGP information is not shared between domains
- Head end lacks the knowledge of complete network topology to perform path computation

Solution:

- Use Explicit Route Object (ERO) Loose Hop Expansion, Node-id, and Path re-evaluation request/reply Flags to provide per-domain path computation at the head-end + RSVP Policy Control and Confidentiality

  RFCs: 3209, 4736, 4561, …etc.

  draft-ietf-ccamp-inter-domain-rsvp-te-06.txt an

  draft-ietf-ccamp-inter-domain-pd-path-comp-05.txt

Per-Domain Path Computation Using ERO Loose-Hop Expansion

Head-End Defines the Path with ASBR and the Destination as Loose Hops

Path Computation Completed During TE LSP Setup

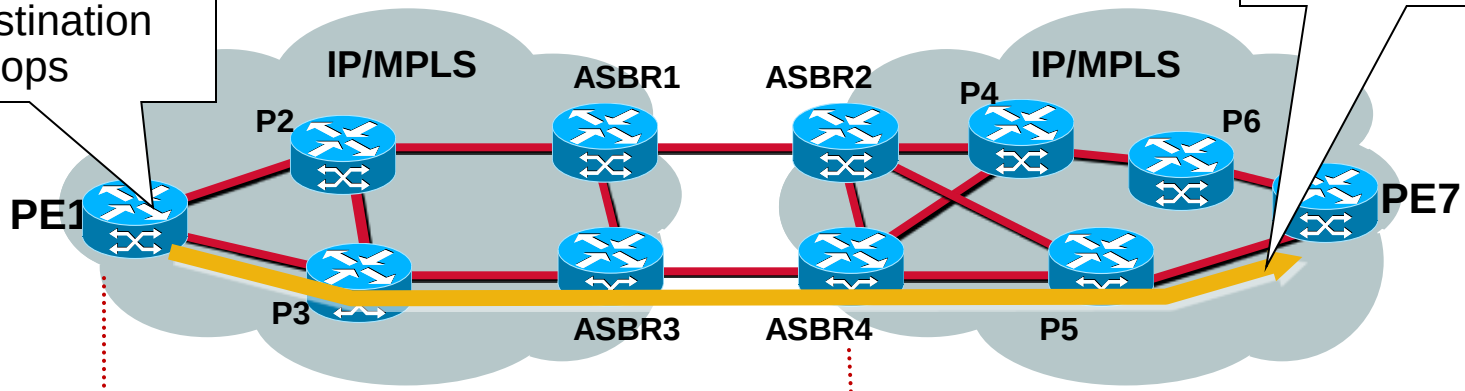

Inter-AS TE LSP

ERO
ASBR4 (Loose)
PE7 (Loose)

Expansion

ERO
R3, ASBR3, ASBR4
PE7 (Loose)

R1 Topology Database

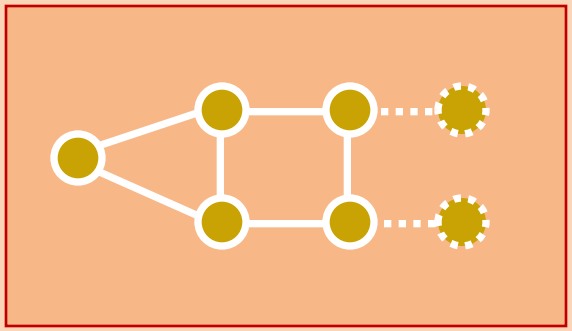

ERO
PE7 (Loose)

Expansion

ERO
P5, PE7

ASBR4 Topology Database

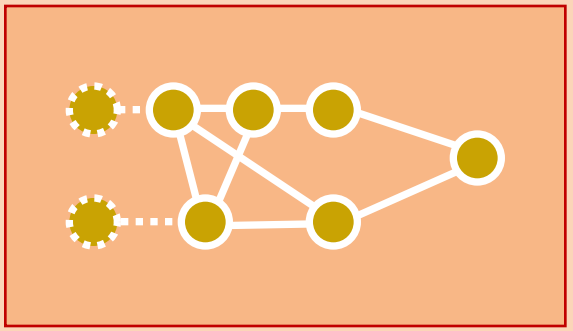

Inter-Domain TE—TE LSP Reoptimization

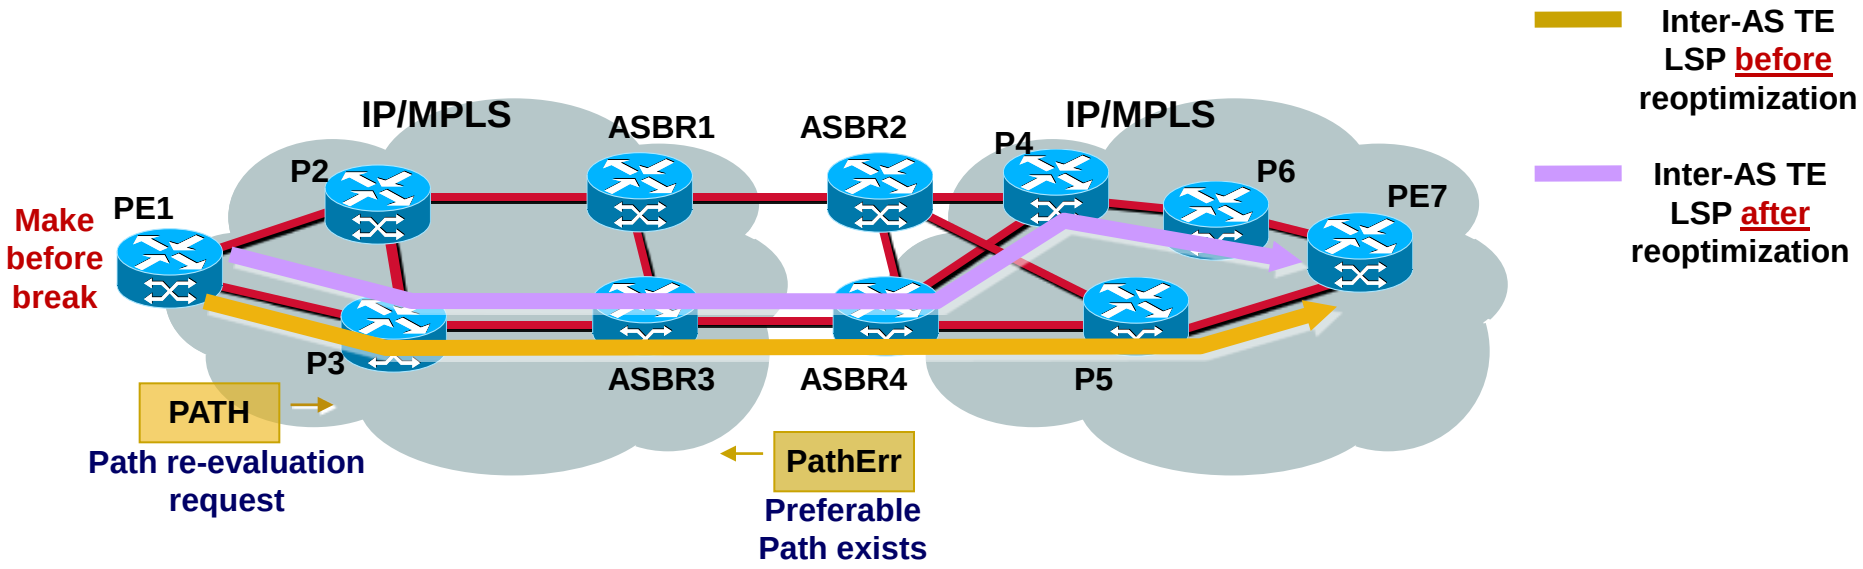

- Reoptimization can be timer/event/admin triggered
- Head end sets ‘path re-evaluation request’ flag (SESSION_ATTRIBUTE)
- Head end receives a PathErr message notification from the boundary router if a preferable path exists
- Make-before-break TE LSP setup can be initiated after PathErr notification

Inter-Domain TE—Fast Re-Route

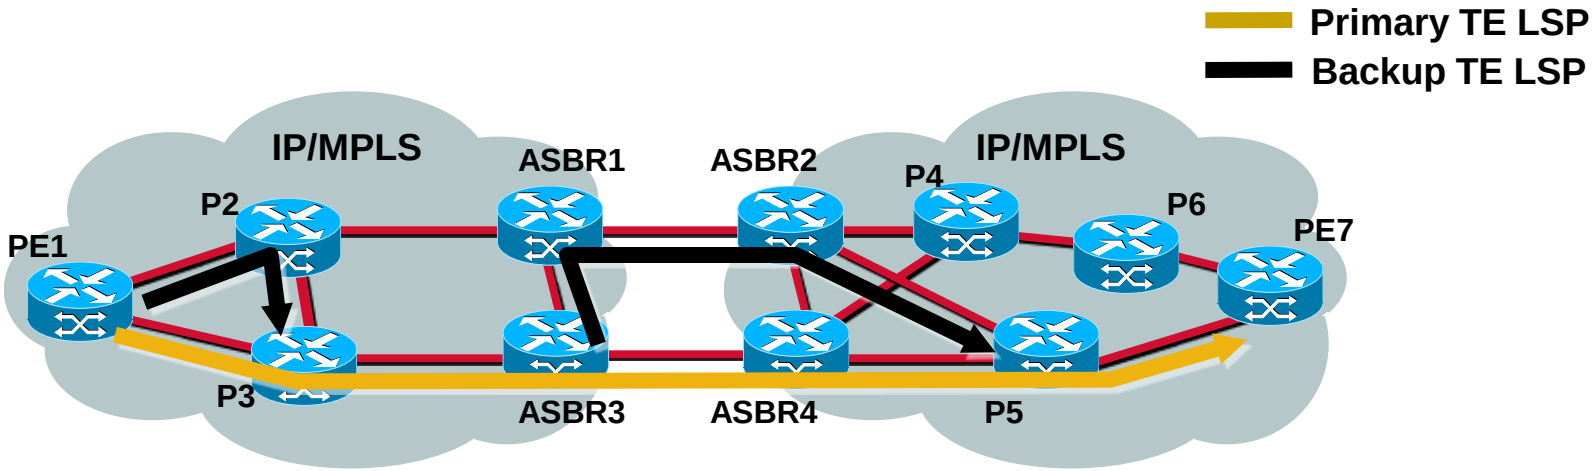

- Same configuration as single domain scenario
- Link and Node protection include ASBRs and ASBR to ASBR links
- Support for Node-id sub-object is required to implement ABR/ASBR node protection
- Node-id helps point of local repair (PLR) detect a merge point (MP)
- Node-id flag defined in draft-ietf-nodeid-subobject

Inter-Domain TE—Policy Control and Confidentiality

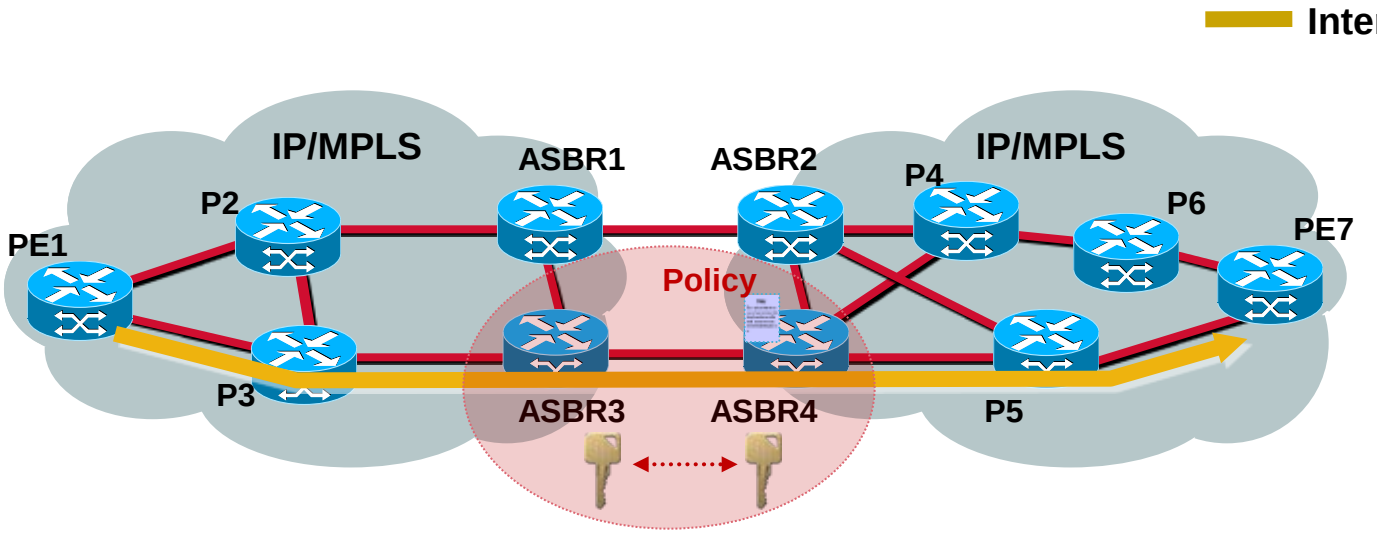

- ASBR may enforce a local policy during Inter-AS TE LSPs setup (e.g. limit bandwidth, message types, protection, etc.)
- Route Recording may be limited
- ASBR may modify source address of messages (PathErr) originated in the AS
- ASBR may perform RSVP authentication (MD5/SHA-1)

Configuring Inter-AS Tunnels (Cisco IOS)

```
mpls traffic-eng tunnels
!
interface Tunnel1
 ip unnumbered Loopback0
 no ip directed-broadcast
 tunnel destination 172.31.255.5
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 10 explicit name LOOSE-PATH
!
ip route 172.31.255.5 255.255.255.255 Tunnel1
!
ip explicit-path name LOOSE-PATH enable
 next-address loose 172.24.255.1
 next-address loose 172.31.255.1
!
```

Loose-hop path

Static route mapping IP traffic to Tunnel1

List of ASBRs as loose hops

Configuring Inter-AS TE at ASBR (Cisco IOS)

```
mpls traffic-eng tunnels
!
key chain A-ASBR1-key
 key 1
 key-string 7 151E0E18092F222A
!
interface Serial1/0
 ip address 192.168.0.1 255.255.255.252
 mpls traffic-eng tunnels
 mpls traffic-eng passive-interface nbr-te-id 172.16.255.4 nbr-igp-id ospf 172.16.255.4
 ip rsvp bandwidth
 ip rsvp authentication key-chain A-ASBR1-key
 ip rsvp authentication type sha-1
 ip rsvp authentication
!
router bgp 65024
 no synchronization
 bgp log-neighbor-changes
 neighbor 172.24.255.3 remote-as 65024
 neighbor 172.24.255.3 update-source Loopback0
 neighbor 192.168.0.2 remote-as 65016
 no auto-summary
!
ip rsvp policy local origin-as 65016
 no fast-reroute
 maximum bandwidth single 10000
 forward all
!
```

Authentication key

Add ASBR link to TE topology database

Enable RSVP authentication

Process signaling from AS 65016 if FRR not requested and 10M or less

Summary

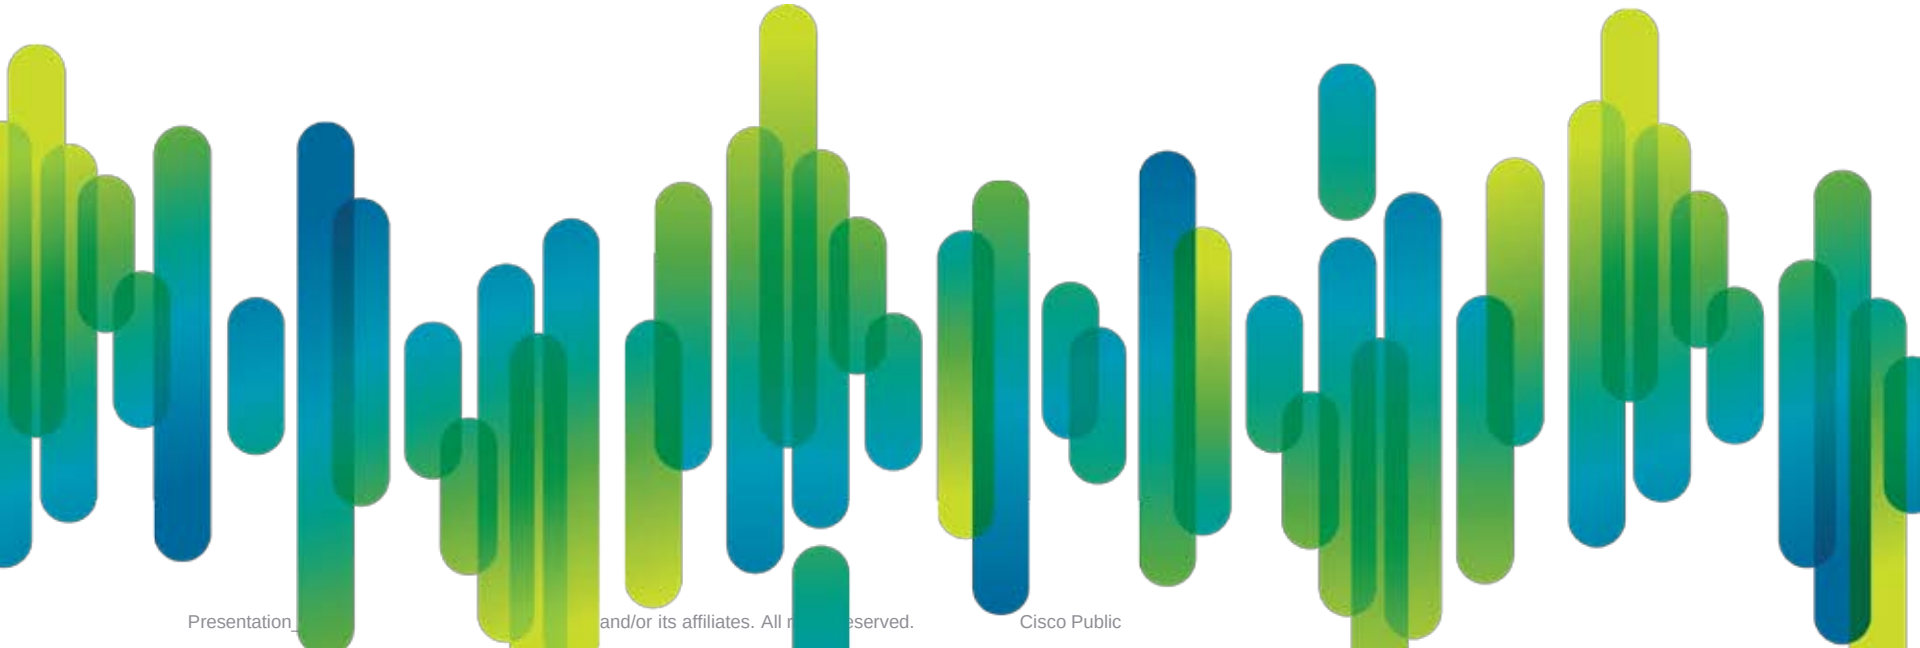

Let's Summarize

CSC: Hierarchical VPNs

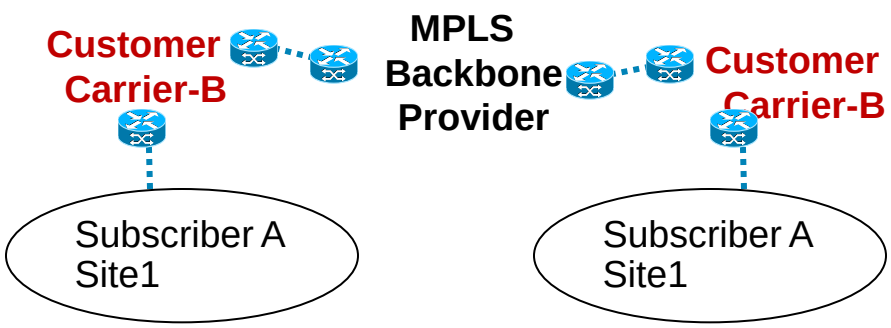

Inter-AS: Extending VPN Boundaries

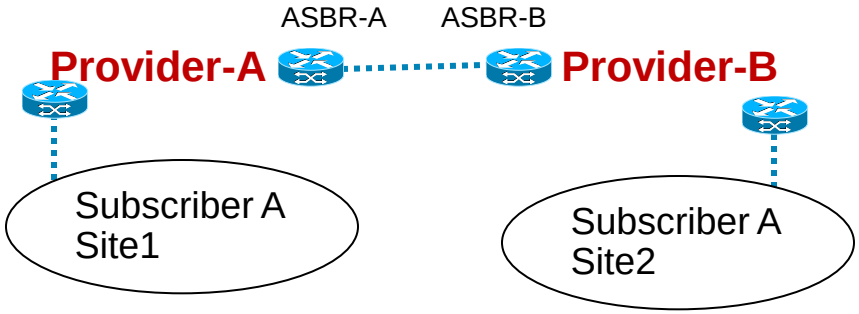

- MPLS VPNs model A, B and C have been deployed to support VPNs among Service Providers and within a single Service Provider's multi-AS networks
- MPLS L2 VPNs, L3VPNs (IPv4, IPv6, and multicast VPNs) are supported in multi-domain environment
- MPLS TE is also supported in multi-area or multi-AS networks
- QoS policies across the ASBRs need to be agreed by the partners and should be configured manually

Meet the Engineer

To make the most of your time at Networkers at Cisco Live 2010, schedule a Face-to-Face Meeting with top Cisco Engineers

Designed to provide a "big picture" perspective as well as "in-depth" technology discussions, these Face-to-Face meetings will provide fascinating dialogue and a wealth of valuable insights and ideas

Visit the Meeting Centre reception desk located in the Meeting Centre in World of Solutions


CISCO

I-AS MVPN Configuration Procedure Option B (SSM)

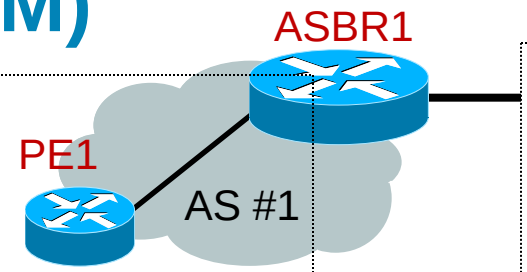


```
! PE1 Configuration:
!
ip multicast-routing
ip multicast routing vrf VPN-A
ip multicast  vrf VPN-A rpf proxy  rd vector
!
router bgp 1
!
address-family ipv4 mdt
neighbor <ASBR1> activate
neighbor <ASBR1> next-hop-self
exit-address-family
!
ip pim ssm default
!
```

```
! ASBR1 Configuration:
!
ip multicast-routing
ip multicast routing vrf VPN-A
!
router bgp 1
!
address-family ipv4 mdt
neighbor <ASBR2> activate
neighbor <PE1> activate
neighbor <PE1> next-hop-self
exit-address-family
!
ip pim ssm default
!
```

Configuration Steps:

1. Enable RPF Vector in the Global table
   ip multicast rpf vector
2. Setup Multicast Address family on ASBRs
   address-family ipv4 mdt
3. Configure PE router to send BGP MDT updates to build the Default MDT
   ip multicast vrf <vrf name> rpf proxy rd vector