APNIC eLearning:
MPLS L3 VPN

Issue Date: 07 July 2015
Revision: 2.0

APNIC

Agenda

- MPLS VPN
- VRF
- RD & RT
- Control Plane of MPLS L3VPN
- Data Plane of MPLS L3VPN
- Configuration Example

MPLS VPN Models

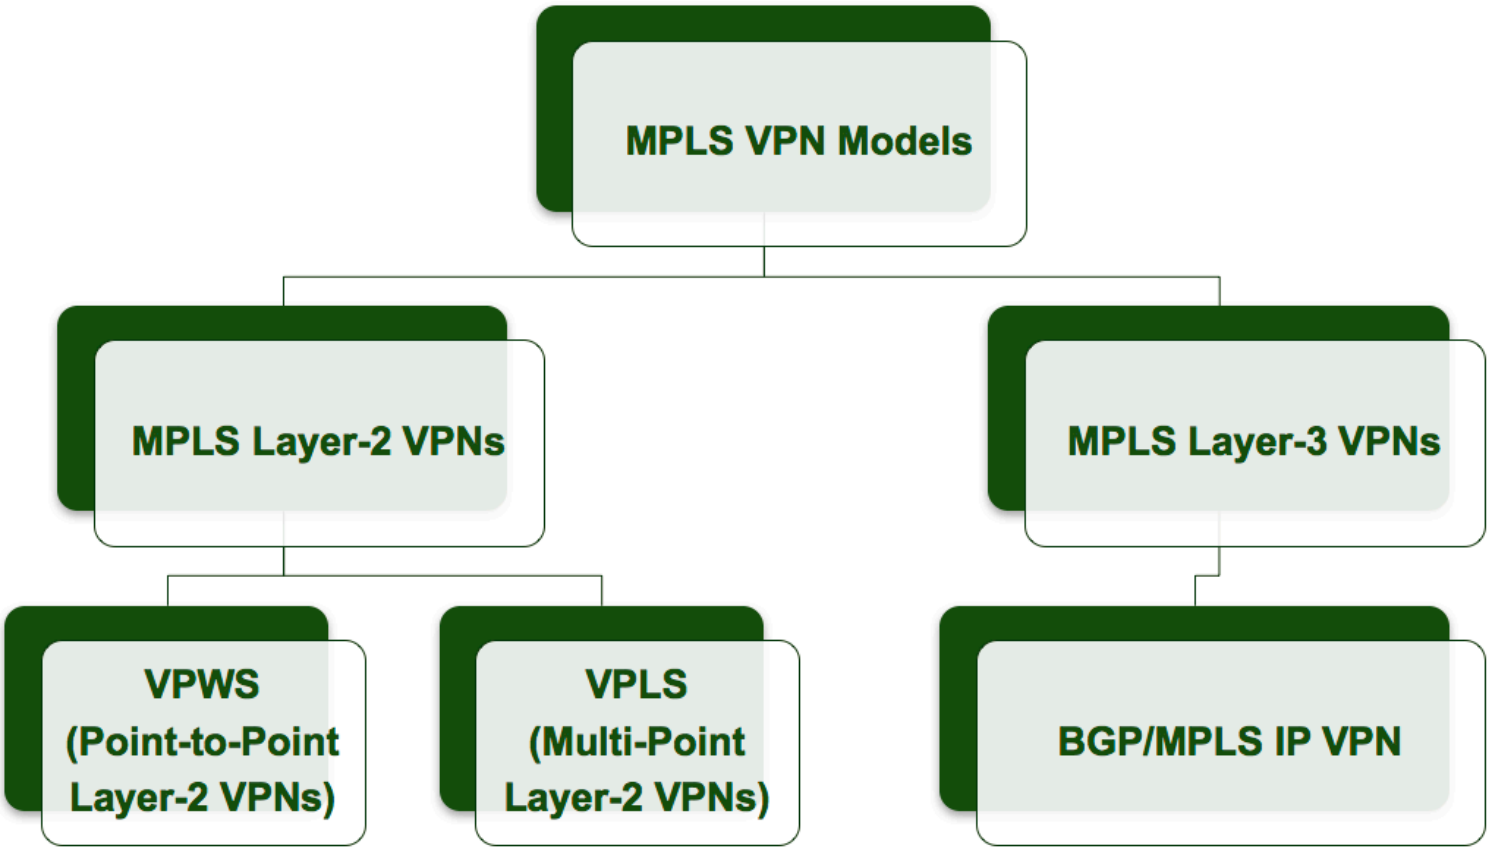

Advantages of MPLS Layer-3 VPN

- Scalability
- Security
- Easy to Create
- Flexible Addressing
- Integrated Quality of Service (QoS) Support
- Straightforward Migration

MPLS L3VPN Topology

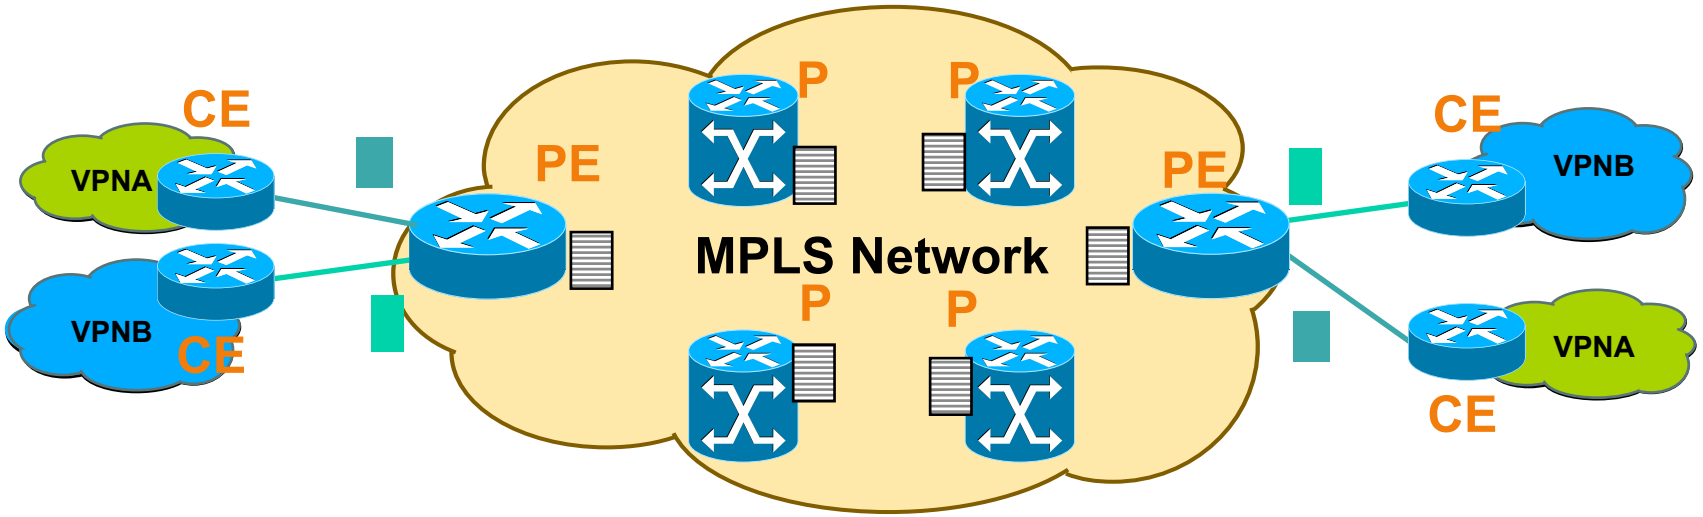

- PE: Provider Edge Router
- P : Provider Router
- CE: Customer Edge Router

Virtual Routing and Forwarding Instance

- Virtual routing and forwarding table
  - On PE router
  - Separate instance of routing (RIB) and forwarding table
- A VRF defines the VPN membership of a customer site attached to a PE device.
- VRF associated with one or more customer interfaces

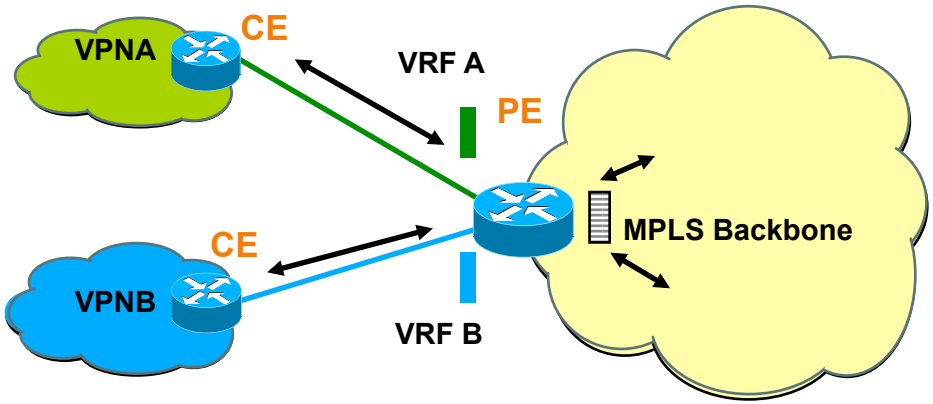

Control Plane: Multi-Protocol BGP

- PE routers use MP-BGP to distribute VPN routes to each other.
- MP-BGP customizes the VPN Customer Routing Information as per the Locally Configured VRF Information at the PE using:
  - Route Distinguisher (RD)
  - Route Target (RT)
  - VPN Label

What is RD

- Route distinguisher is an 8-octet field prefixed to the customer's IPv4 address. RD makes the customer's IPv4 address unique inside the SP MPLS network.
- RD is configured in the VRF at PE

VPNv4 Address:	Route Distinguisher (8 bytes)	IPv4 Address (4 bytes)
Example:	Type 0	100:1	10.1.1.1
	Type 1	192.168.19.1:1	10.1.1.1

Route Advertisement: RD

- VPN customer IPv4 prefix is converted into a VPNv4 prefix by appending the RD to the IPv4 address
- PE devices use MP-BGP to advertise the VPNv4 address

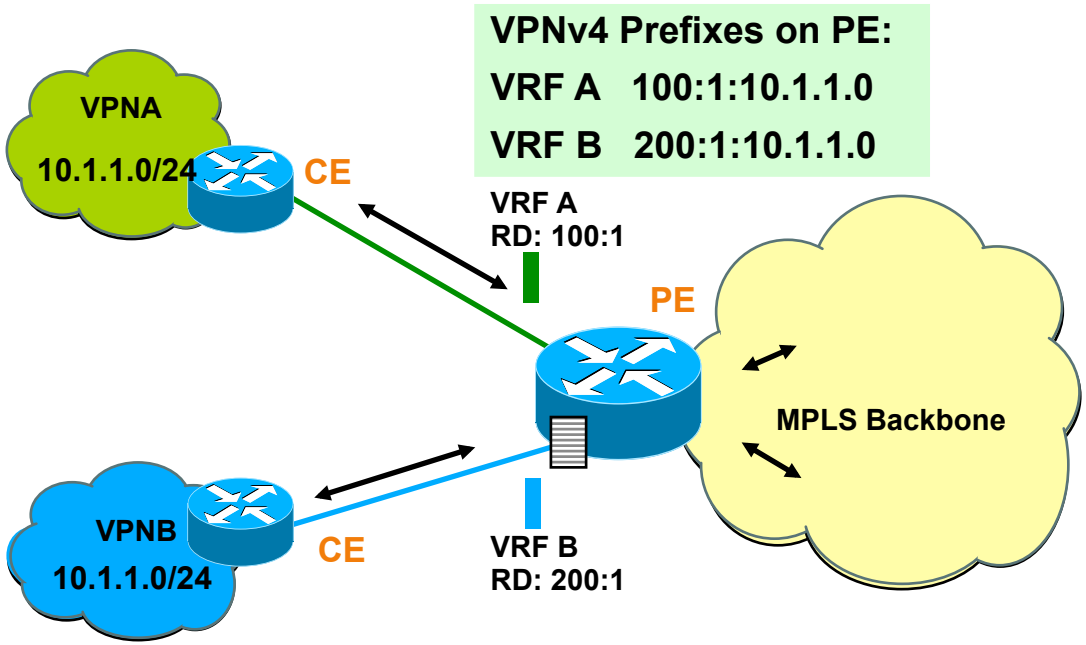

What is RT

- Route Target is a BGP extended community attribute, is used to control VPN routes advertisement.

Example:

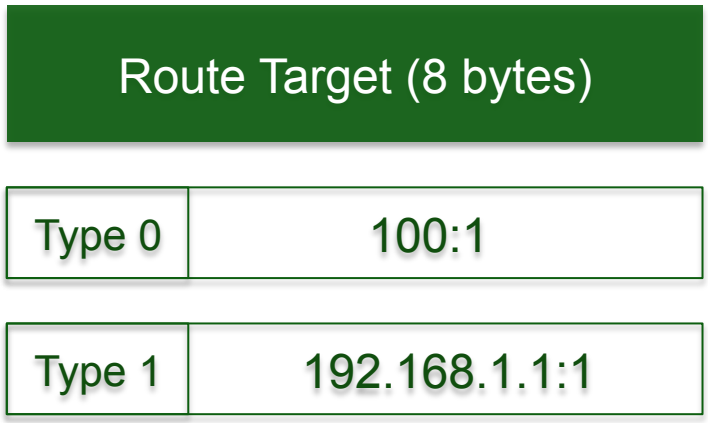

- Two types of RT:
  - Export RT
  - Import RT

Route Advertisement: RT

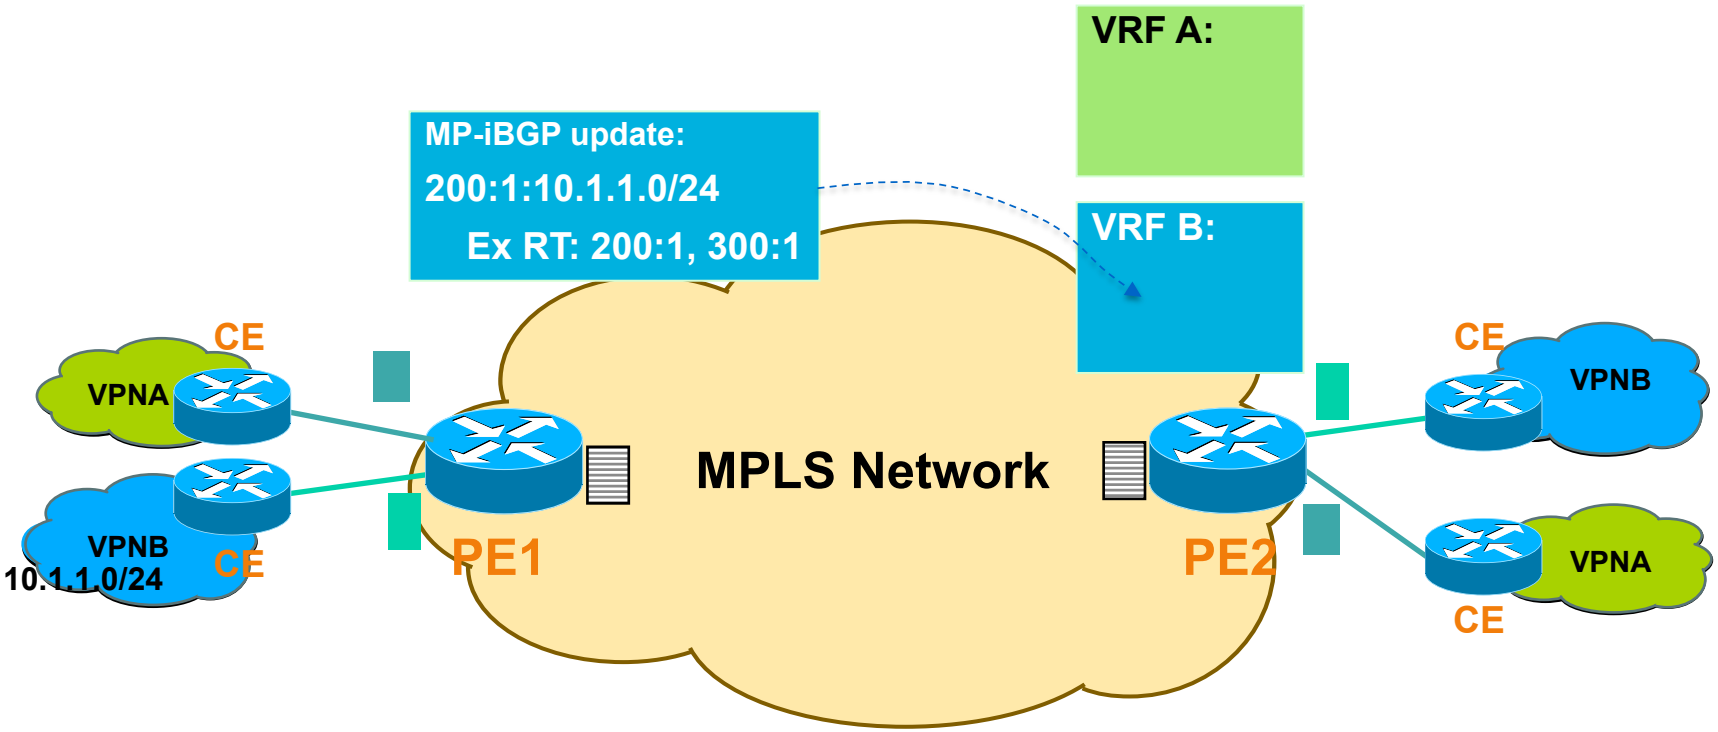

	Import RT	Export RT
VRF A	100:1	100:1
VRF B	200:1 300:1	**200:1** 300:1

	Import RT	Export RT
VRF A	100:1 400:1 500:1	100:1 400:1
VRF B	**200:1**	200:1

Using RT to Configure VPN Topologies

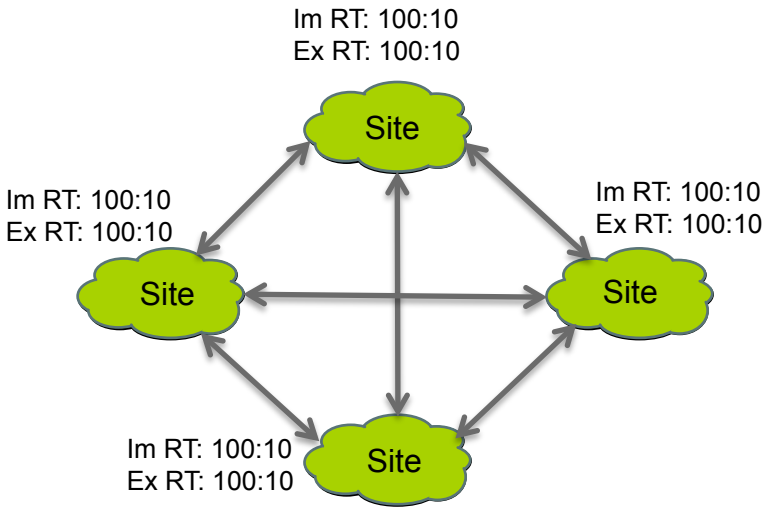

Full Mesh

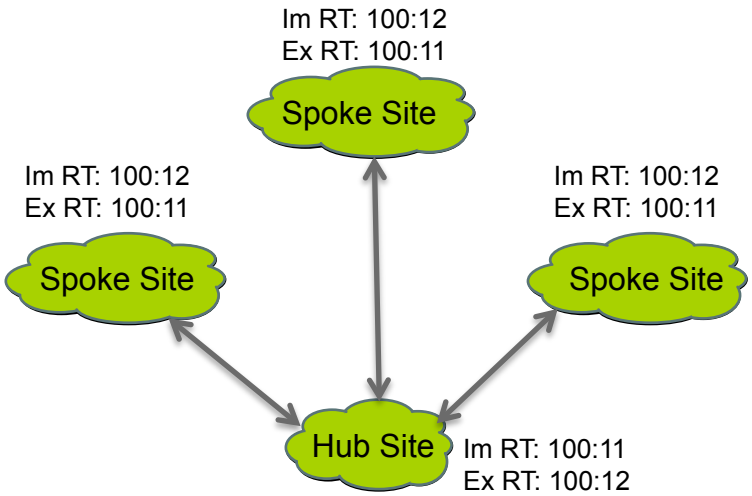

Hub Spoke

In a full-mesh VPN, each site in the VPN can communicate with every other site in that same VPN.

In a hub-and-spoke VPN, the spoke sites in the VPN can communicate only with the hub sites; they cannot communicate with other spoke sites.

VPN Label

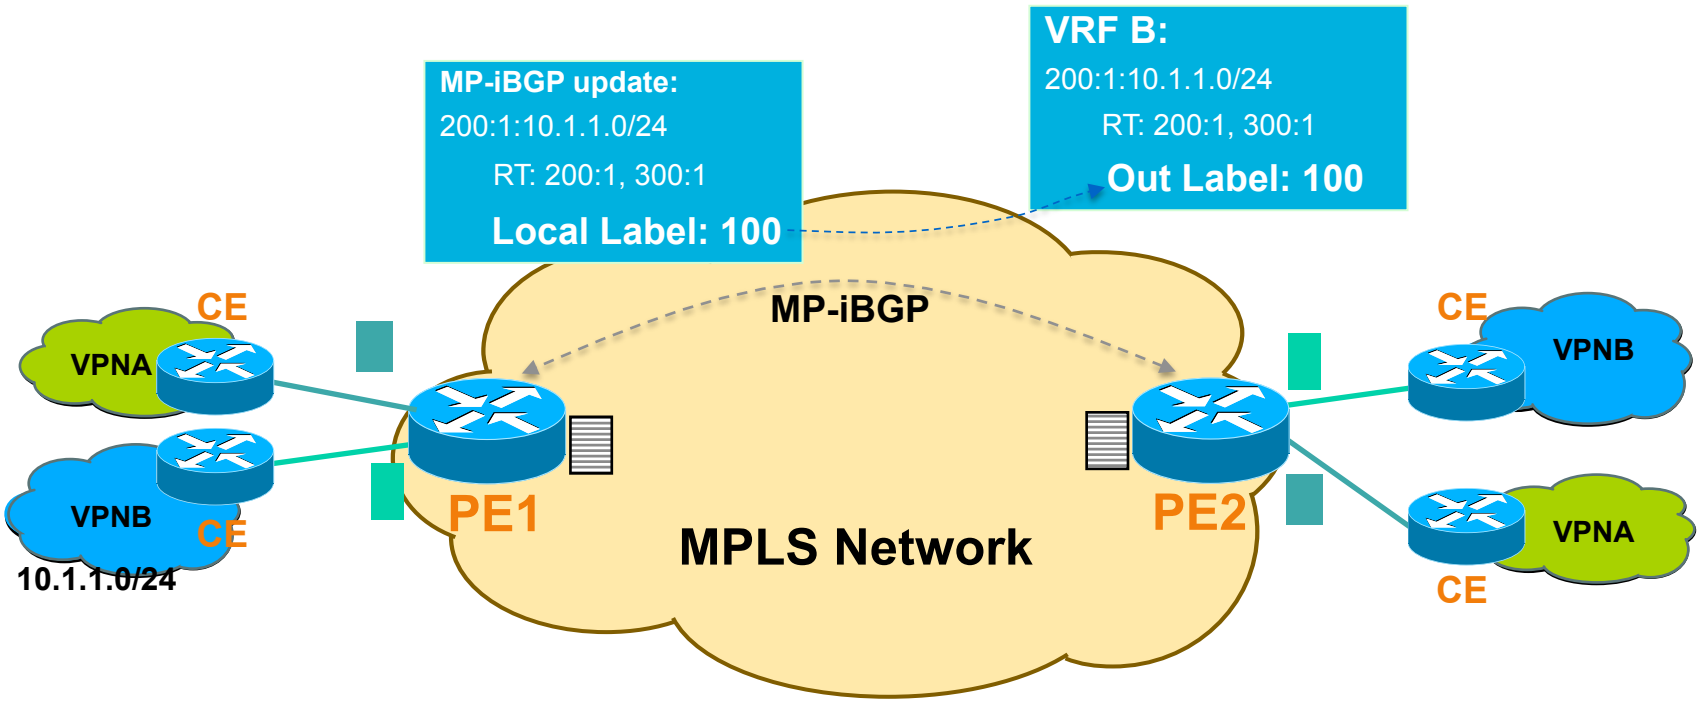

- PE adds the label to the NLRI field.

Control Plane Walkthrough(1/2)

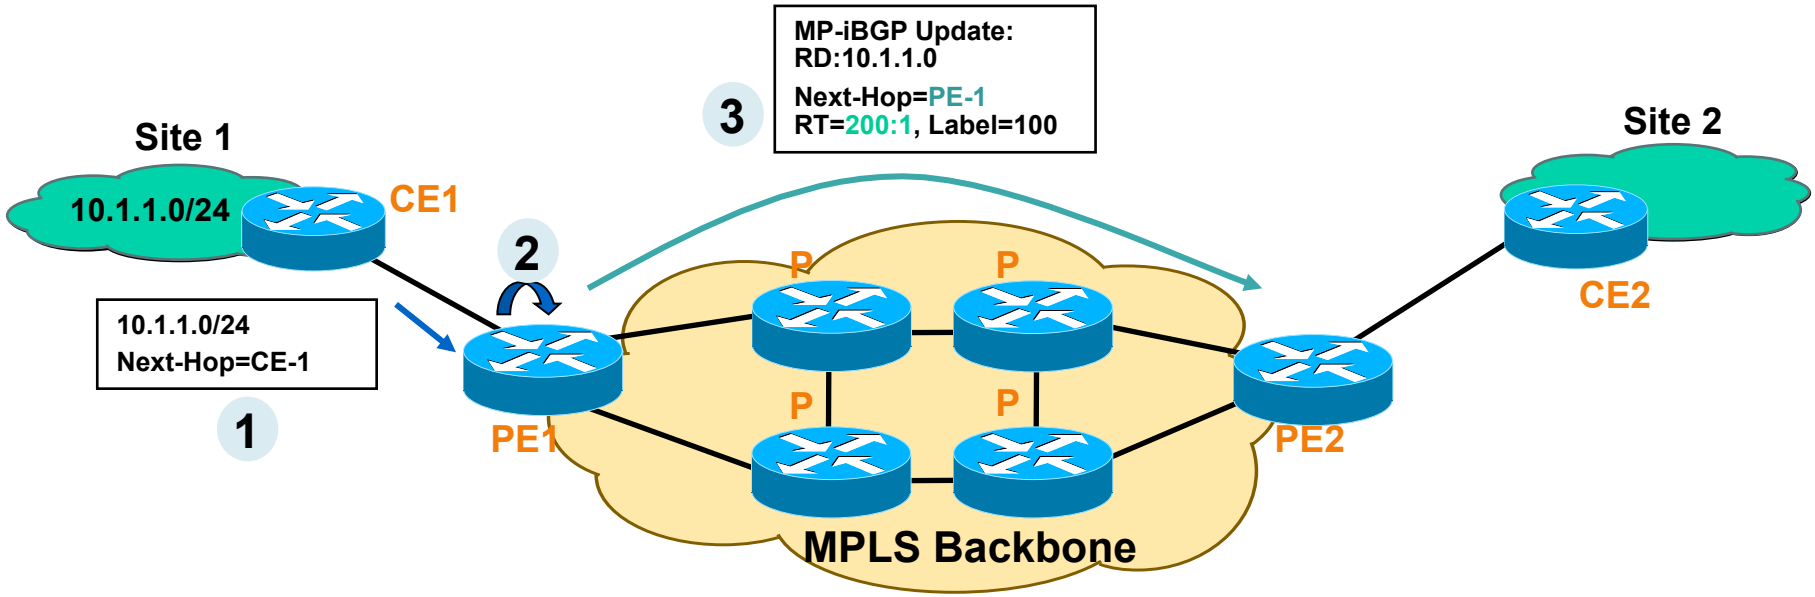

1. PE1 receives an IPv4 update (eBGP/OSPF/ISIS/RIP/EIGRP)
2. PE1 converts it into VPNv4 address and constructs the MP-iBGP UPDATE message
   - Associates the RT values (export RT =200:1) per VRF configuration
   - Rewrites next-hop attribute to itself
   - Assigns a label (100); Installs it in the MPLS forwarding table.
3. PE1 sends MP-iBGP update to other PE routers

Control Plane Walkthrough(2/2)

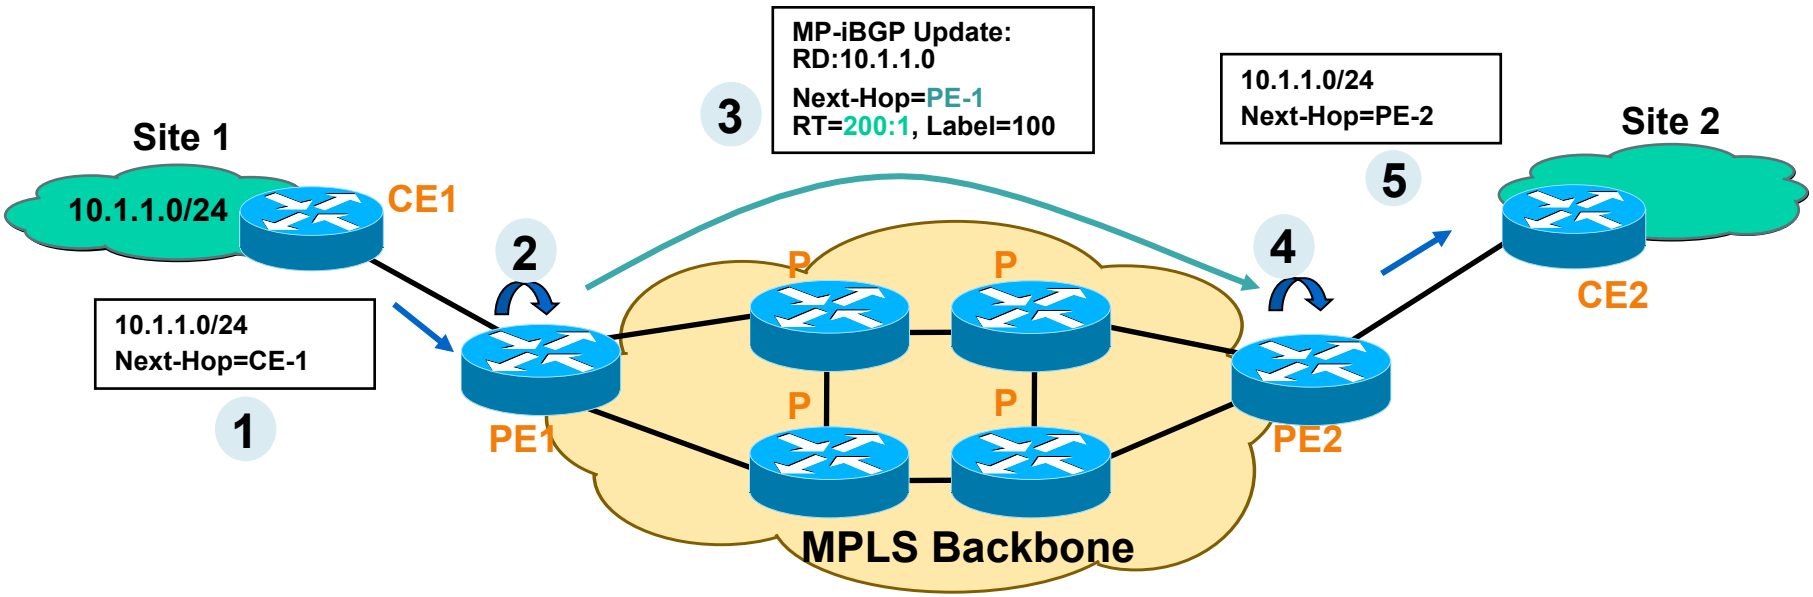

4. PE2 receives and checks whether the RT=200:1 is locally configured as import RT within any VRF, if yes, then
   - PE2 translates VPNv4 prefix back to IPv4 prefix
   - Updates the VRF CEF table for 10.1.1.0/24 with label=100
5. PE2 advertises this IPv4 prefix to CE2 (using whatever routing protocol)

Control Plane: Tunnel Label

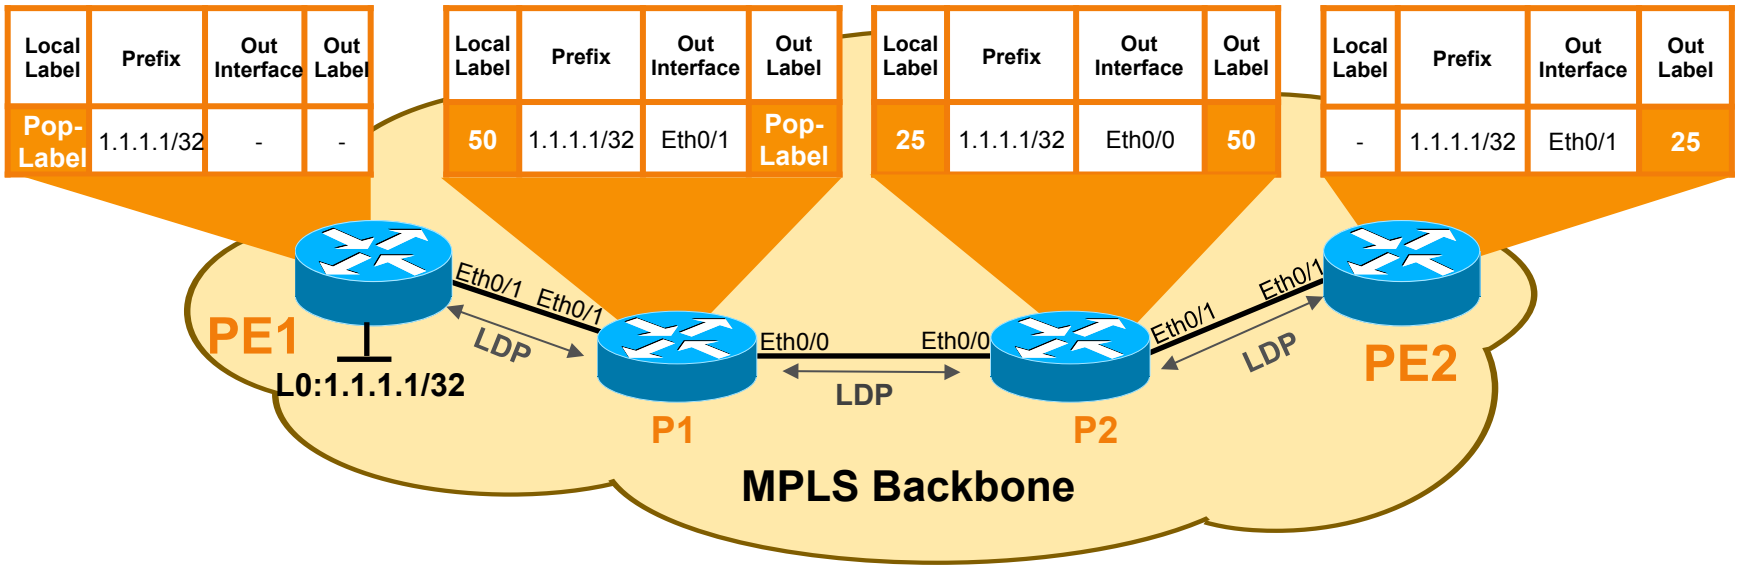

- LDP runs on the MPLS backbone network to build the public LSP. The tunnel label is also called transport label or public label.
- Local label mapping are sent to connected nodes. Receiving nodes update forwarding table.

Data Plane

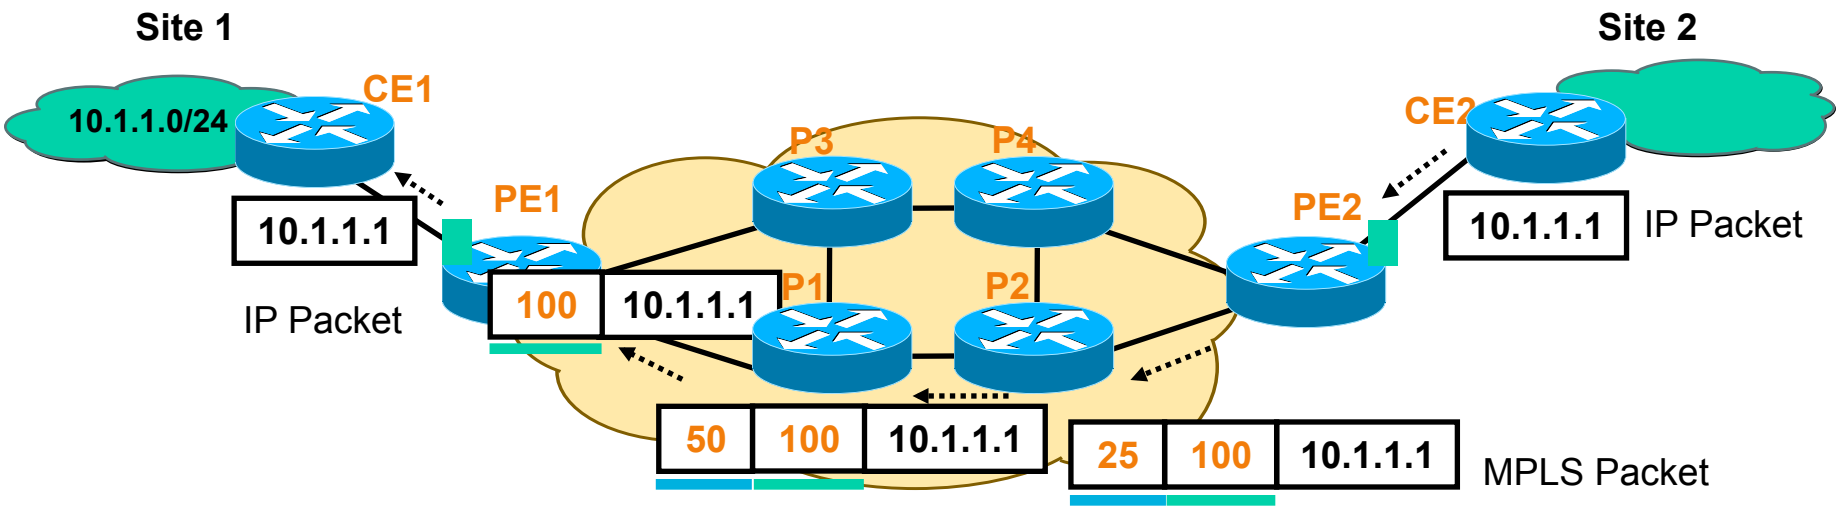

- PE2 imposes two labels for each IP packet going to site2
  - Tunnel label is learned via LDP; corresponds to PE1 address
  - VPN label is learned via BGP; corresponds to the VPN address
- P1 does the Penultimate Hop Popping (PHP)
- PE1 retrieves IP packet (from received MPLS packet) and forwards it to CE1.

Configuration Example

- Task: Configure MPLS L3VPN on Cisco IOS (Version 15.2) to make the following CEs communication with each other.
- Prerequisite configuration:
  - 1. IP address configuration on PE & P routers
  - 2. IGP configuration on PE & P routers
    - Make sure all the routers in public network can reach each other.

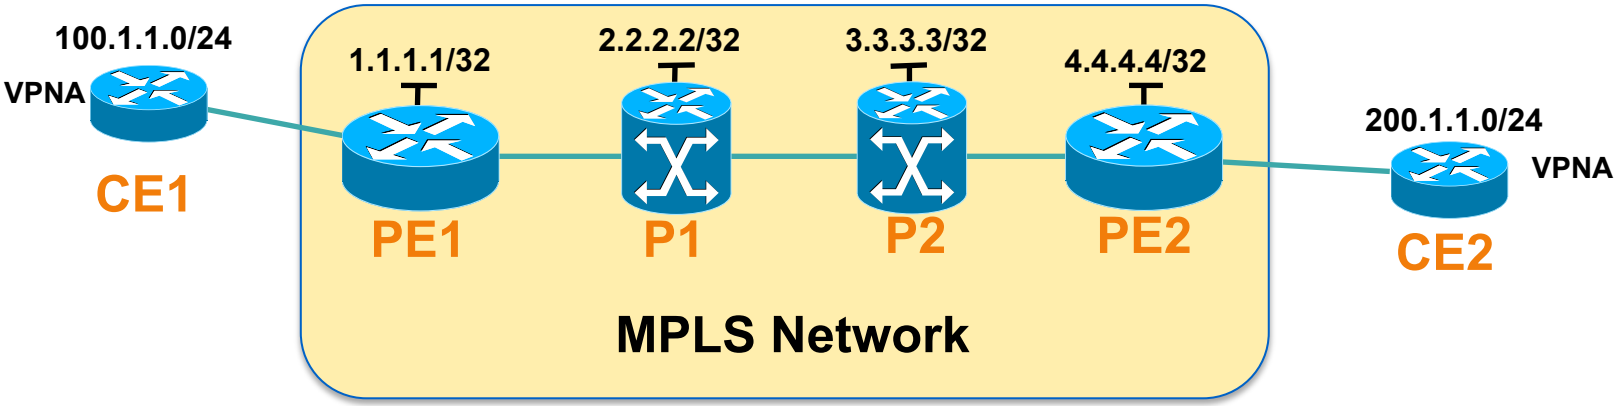

Configure MPLS & LDP

- Configuration steps:
  - 1. Configure MPLS and LDP on PE & P routers

```
ip cef
mpls ldp router-id loopback 0

interface ethernet1/0
mpls ip
mpls label protocol ldp

interface ethernet1/1
mpls ip
mpls label protocol ldp
```

Configure VRF

- Configuration steps:
  - 2. Configure VRF instance on PE routers

```
vrf definition VPNA
 rd 100:10
 route-target export 100:10
 route-target import 100:10
 !
 address-family ipv4
 exit-address-family
!
```

  - bind PE-CE interface under VRF

```
interface FastEthernet0/0
 vrf forwarding VPNA
 ip address 10.1.1.1 255.255.255.252
```

Configure MP-iBGP

- Configuration steps:
  - 3. Enable MP-iBGP neighbors in vpnv4 address-family on PE routers

```
router bgp 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 4.4.4.4 update-source loopback 0
 !
 address-family vpnv4
  neighbor 4.4.4.4 activate
  neighbor 4.4.4.4 send-community both
 exit-address-family
!
```

Configure PE-CE eBGP Neighbour

- Configuration steps:
  - 4. Adding PE-CE eBGP neighbour in VRF context of BGP on PE

```
router bgp 100
 address-family ipv4 vrf VPNA
  neighbor 10.1.1.2 remote-as 65001
  neighbor 10.1.1.2 activate
 exit-address-family
!
```

Adding PE-CE eBGP neighbour in BGP on CE

```
router bgp 65001
 neighbor 10.1.1.1 remote-as 100
 !
 address-family ipv4
  network 100.1.1.0 mask 255.255.255.0
  neighbor 10.1.1.1 activate
 exit-address-family
!
ip route 100.1.1.0 255.255.255.0 null 0
```

Verify Results – VRF Routing Table

- Check the routes of VRF VPNA on PE.

```
PE1#show bgp vpnv4 unicast vrf VPNA
BGP table version is 4, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:10 (default for vrf VPNA)
 *>  100.1.1.0/24    10.1.1.2                 0             0 65001 i
 *>i 200.1.1.0       4.4.4.4                  0    100      0 65002 i
```

Verify Results – VPN Reachability

- CE can learn the routes from each other:

```
CE2#show ip route
....
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.1.2.0/30 is directly connected, FastEthernet0/1
L        10.1.2.2/32 is directly connected, FastEthernet0/1
      100.0.0.0/24 is subnetted, 1 subnets
B        100.1.1.0 [20/0] via 10.1.2.1, 00:38:26
      200.1.1.0/24 is variably subnetted, 2 subnets, 2 masks
S        200.1.1.0/24 is directly connected, Null0
C        200.1.1.1/32 is directly connected, Loopback1
```

Please remember to fill out the feedback form:

- Survey Link

Slides are available for download from APNIC FTP.

APNIC Helpdesk Chat

Helpdesk

Like Share 2

Tweet

APNIC Helpdesk provides assistance to all on matters related to APNIC Services, such as membership and IP address enquiries.

APNIC Helpdesk offers (through prior arrangement) multi-language phone support for the following: Bahasa Indonesia, Bahasa Malaysia, Bengali, Cantonese, English, Filipino (Tagalog), Hindi, Japanese, Malay, Mandarin, Sinhalese, Tamil and Telugu.

You may also find our FAQs helpful with your enquiries.

Contact details

Helpdesk hours 09:00 to 21:00 (UTC +10) Monday - Friday (closed for some public holidays)

Chat APNIC Live Chat Online — Click here to chat

Skype Call — Click here to Skype ID: apnic-helpdesk

Email helpdesk@apnic.net

Phone +61 7 3858 3188

VoIP helpdesk@voip.apnic.net

Fax + 61 7 3858 3199

Service Updates

Service announcement: 10 February 2016

Service disruption: APNIC services were disrupted on Wednesday, 10 February 2016

More announcements

Subscribe to APNIC Service Announcements

Learn more about system maintenance

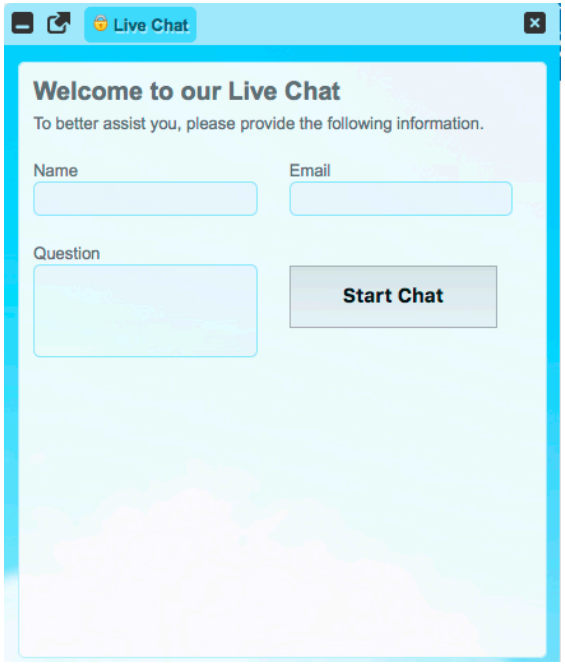

Thank You!

END OF SESSION