MPLS VPN

Prepared by
Eng. Hussein M. Harb

MPLS VPN

Agenda

- Why VPN
- VPN Definition
- VPN Categories
- VPN Implementations
- VPN Models
- MPLS VPN Types
- L3 MPLS VPN
- L2 MPLS VPN

Why VPN?

- VPNs were developed initially to deal with **security** issues of transmitting clear text data across a network.
- Examples of applications that send traffic in a clear text format are Telnet, file transfers via FTP or TFTP.
- VPN - has attracted the attention of many organizations looking to expand their networking capabilities, secure their traffic and reduce their costs.

VPN Definition

The most common definition of a VPN is:

> A data network that utilizes a portion of a shared public network to extend a customer's private network.

VPN Categories

There are three basic VPN categories:

- Intranet
- Extranet
- Internet

Intranet VPN

- An intranet VPN connects resources from the same company across that company's infrastructure.

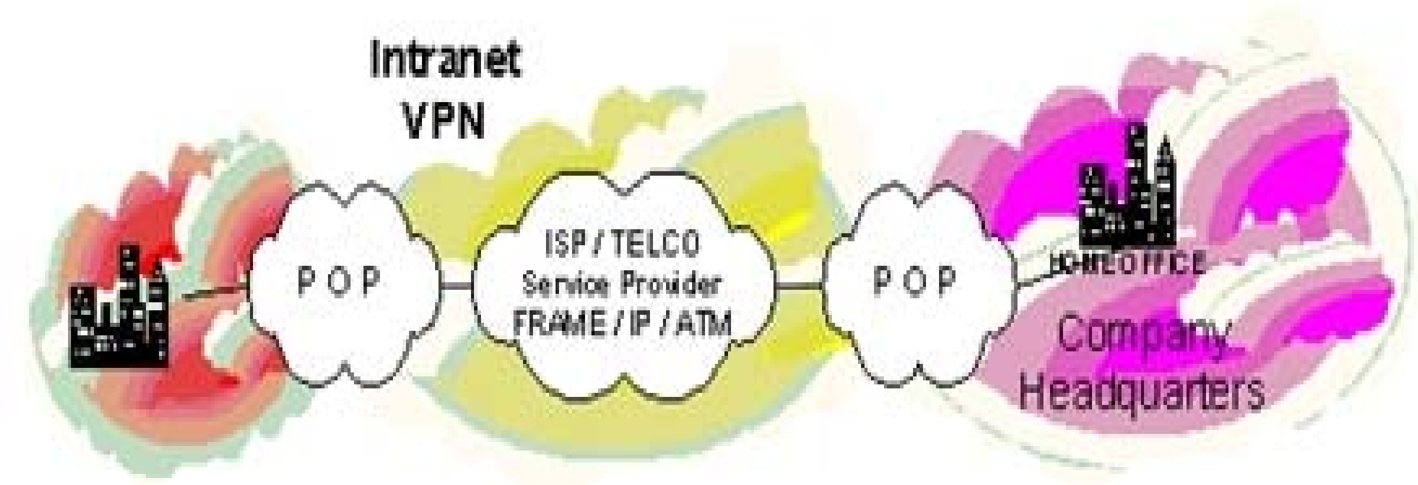

An example of intranet VPN is the connections between different locations within a company's infrastructure, such as VPNs between two offices

Extranet VPN

- An extranet VPN connects resources from one company to another company, such as a business partner.

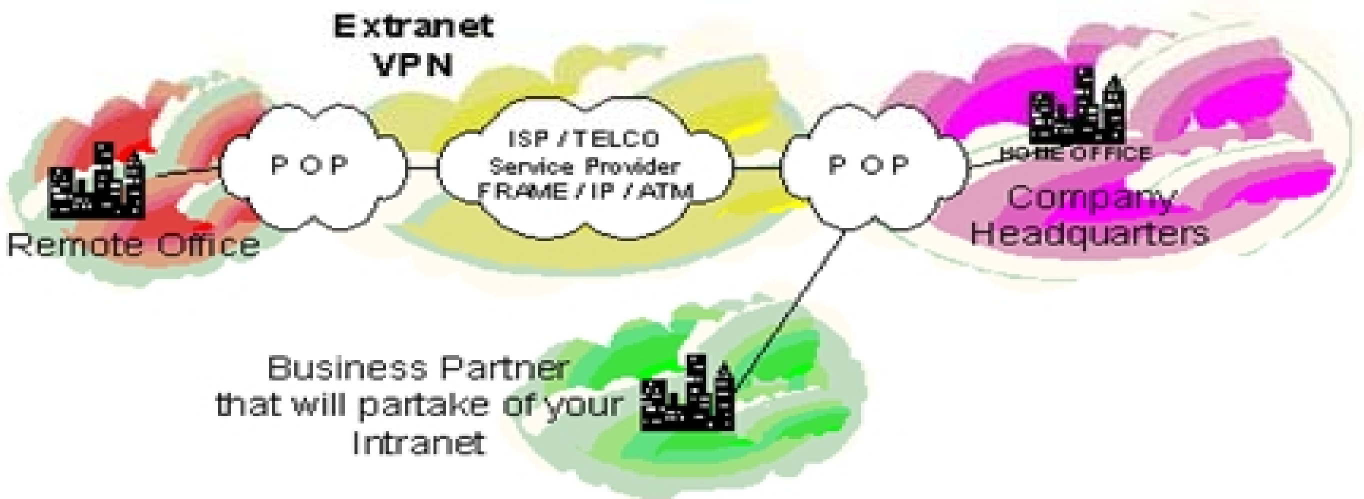

An example of an extranet is a company that has outsourced its help desk functions and sets up a VPN to provide a secure connection from its corporate office to the outsourcing company.

Internet

- An Internet VPN uses a public network as the backbone to transport VPN traffic between devices.

- As an example, you might use the Internet, which is a public network, to connect two sites together or have telecommuters use their local ISPs to set up a VPN connection to the corporate network (remote access connections).

VPN Components

The VPN realm consist of the following regions:

- **Customer network:**
  Consisted of the routers at the various customer sites called customer edge (CE) routers.

- **Provider network:**
  SP devices to which the CE routers were directly attached were called provider edge (PE) routers.
  SP network might consist of devices used for forwarding data in the SP backbone called provider (P) routers.

VPN Implementations

There are many ways for the implementation of VPN such as:

- GRE
- IPsec
- PPTP
- L2TP
- MPLS

MPLS VPN

- MPLS VPNs are enhancement to MPLS
- MPLS uses a virtual circuit (VC) across a private network to emulate the VPN function.
- MPLS alone won't solve security problem; you'll have to complement it with another VPN solution, such as IPsec over MPLS.
- MPLS supports multiple protocols. In other words, you can use MPLS to tag IP packets, Ethernet frames, IPX packets.

VPN Models

The VPN implementations can be classified broadly into one of the following:

- Overlay model
- Peer-to-peer model

Overlay model

- The provider did not participate in customer routing. It provides the customer with transport of data using virtual point-to-point links (PVC or SVC).

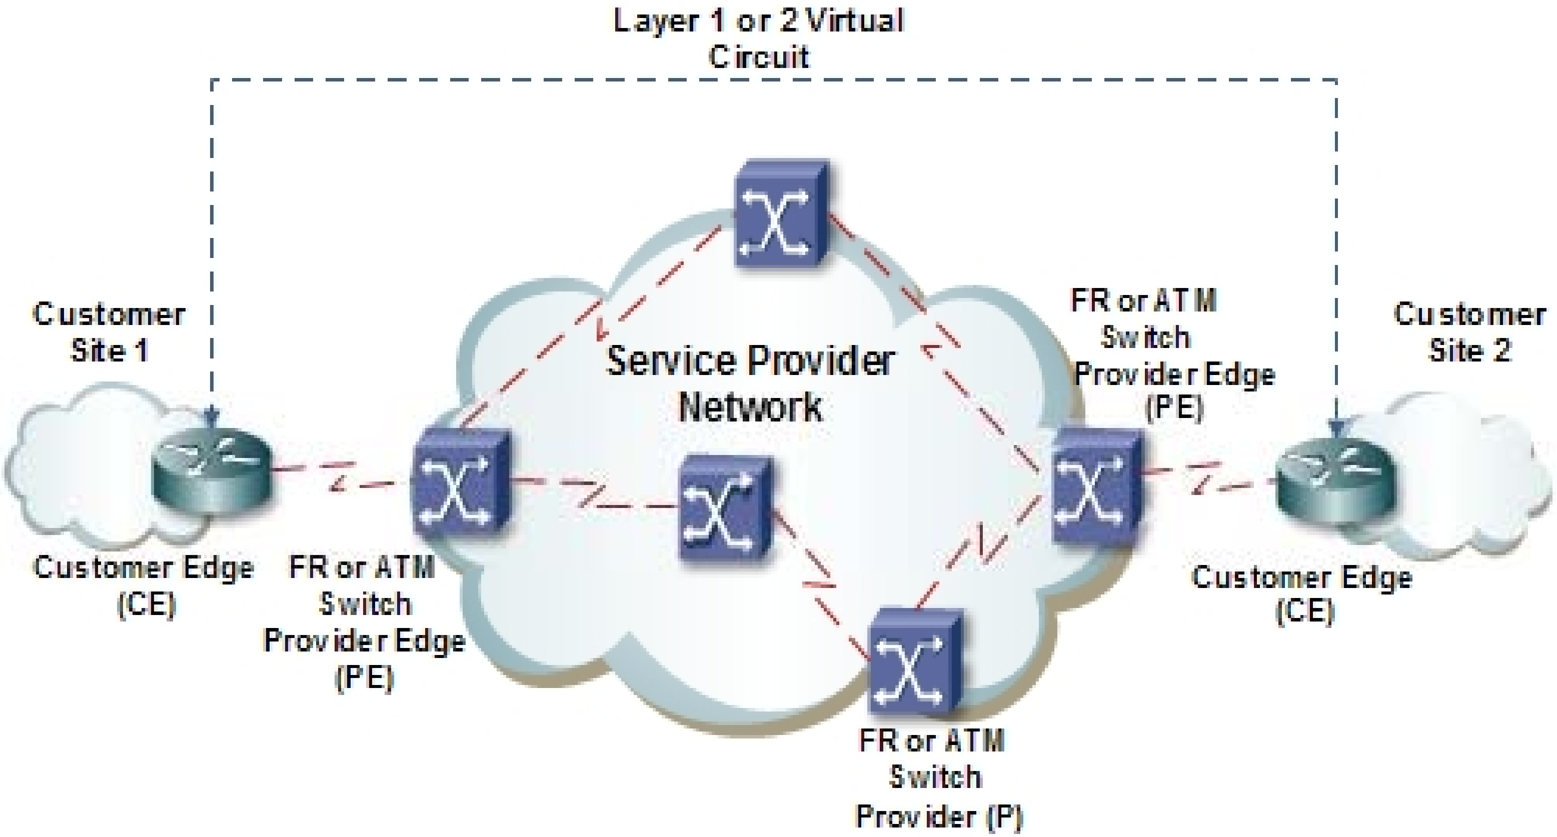
Layer 1 or 2 Virtual Circuit
Customer Site 1
Service Provider Network
FR or ATM Switch Provider Edge (PE)
Customer Site 2
Customer Edge (CE)
FR or ATM Switch Provider Edge (PE)
Customer Edge (CE)
FR or ATM Switch Provider (P)

Overlay model (Continue)

- The drawback of an Overlay model was the full mesh of virtual circuits between all customer sites for optimal connectivity. N sites need N(N-l)/2 circuits.

- Overlay VPNs provides either Layer 1 (physical layer) connectivity or a Layer 2 transport circuit between customer sites for transportation of Layer 2 frames (Or cells) which was traditionally implemented using either Frame Relay or ATM switches .

Peer-to-peer model

- The peer-to-peer model was developed to overcome the drawbacks of the Overlay model
- The service provider would actively participate in customer routing

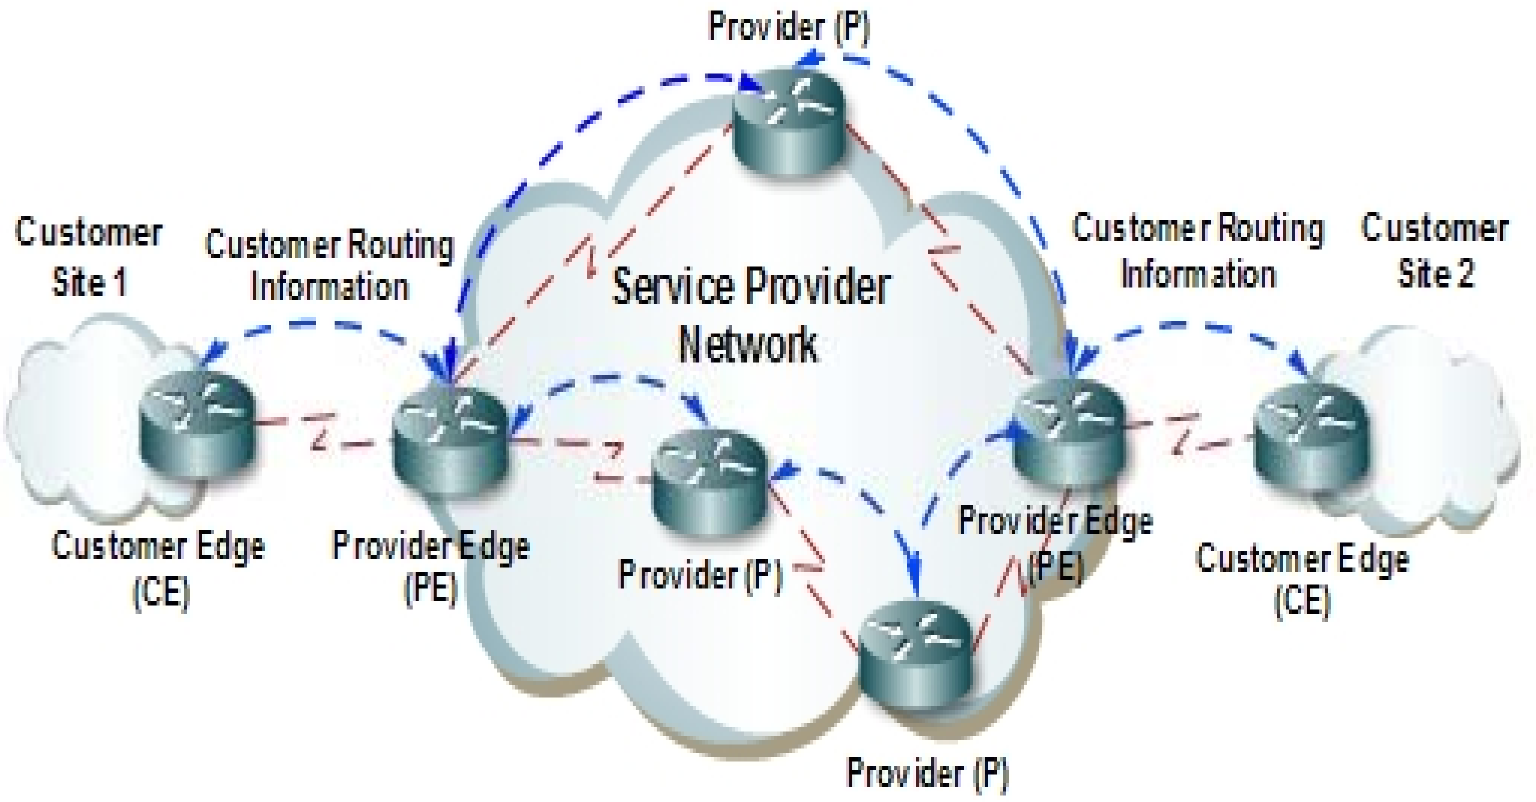

Peer-to-peer model (Continue)

- Routing information is exchanged between the customer routers and the SP routers.
- The peer-to-peer model, consequently, does not require the creation of virtual circuits.
- Separation of customer-specific routing information is achieved by implementing packet filters at the routers connecting to the customer network.

MPLS VPN Types

- **BGP/MPLS VPNs (Layer 3 VPNs):**
  Use extensions to the existing routing protocol of the Internet (BGP-4) to interconnect remote locations, also called RFC 2547bis VPNs.

- **Layer 2 MPLS VPNs:**
  Extends the customer's Layer 2 connectivity across an MPLS infrastructure. Commonly called Martini VPNs. An extension to Layer 2 VPNs also supports Virtual Private LAN Services (VPLS).

L3 MPLS VPN Architecture

- MPLS VPN is an implementation of the peer-to-peer model.
- The MPLS-based VPN model also accommodates customers using-overlapping address spaces.
- However, instead of deploying a dedicated PE router per customer, customer traffic is isolated on the same PE router providing connectivity for multiple customers.
- The MPLS VPN backbone and customer sites exchange Layer 3 customer routing information.

Components of MPLS VPN architecture

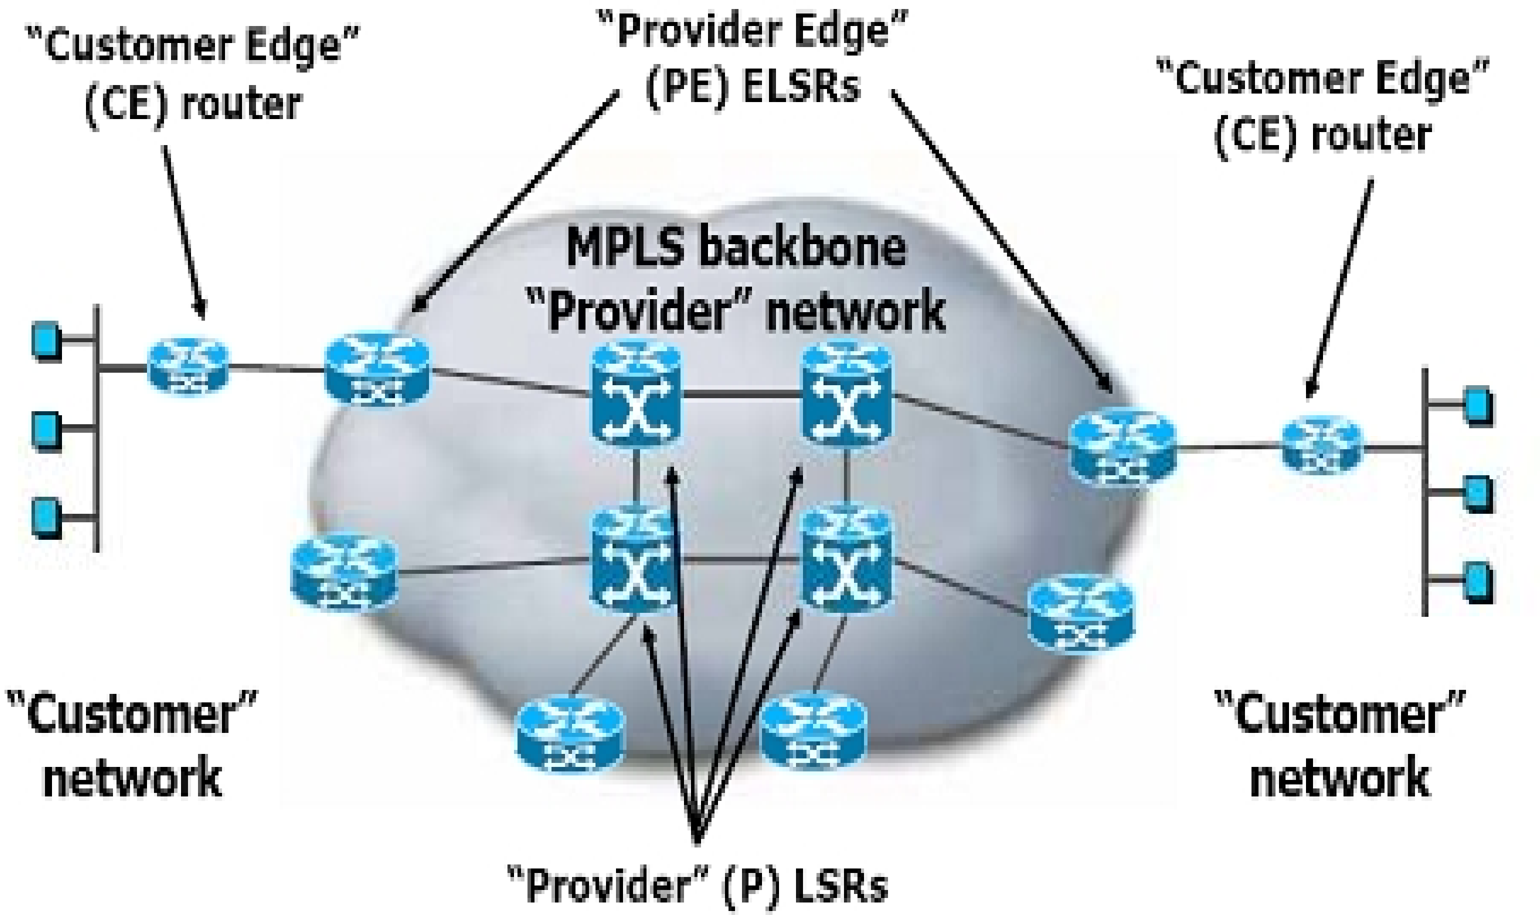
"Customer Edge" (CE) router
"Provider Edge" (PE) ELSRs
"Customer Edge" (CE) router
MPLS backbone "Provider" network
"Customer" network
"Customer" network
"Provider" (P) LSRs

L3 MPLS VPN Routing Model

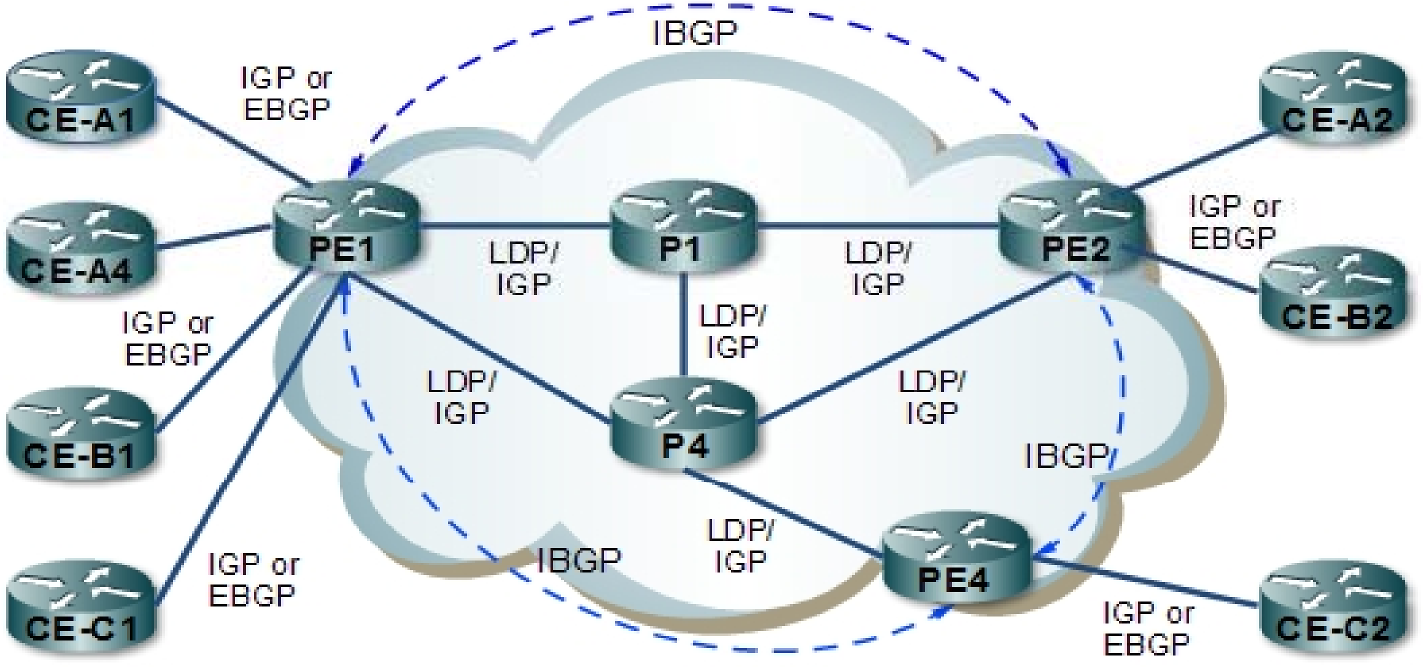

- The only requirement on the CE router is a routing protocol or a static route that enables the router to exchange IPv4 routing information with the connected PE router.

L3 MPLS VPN Routing Model

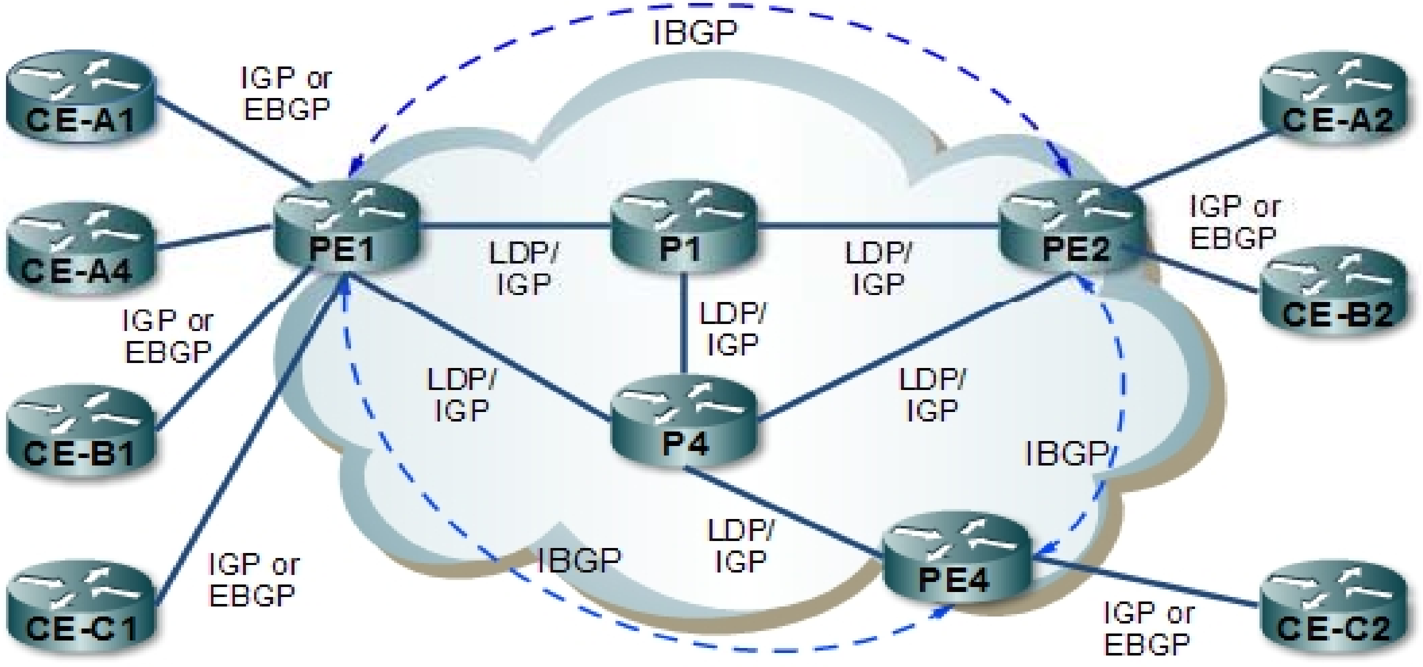

PE routers Perform the following tasks:

- The PE routers exchange IPv4 routes with connected CE routers using individual routing protocol contexts.
- It must isolate customer traffic if more than one customer is connected to the PE router.

L3 MPLS VPN Routing Model

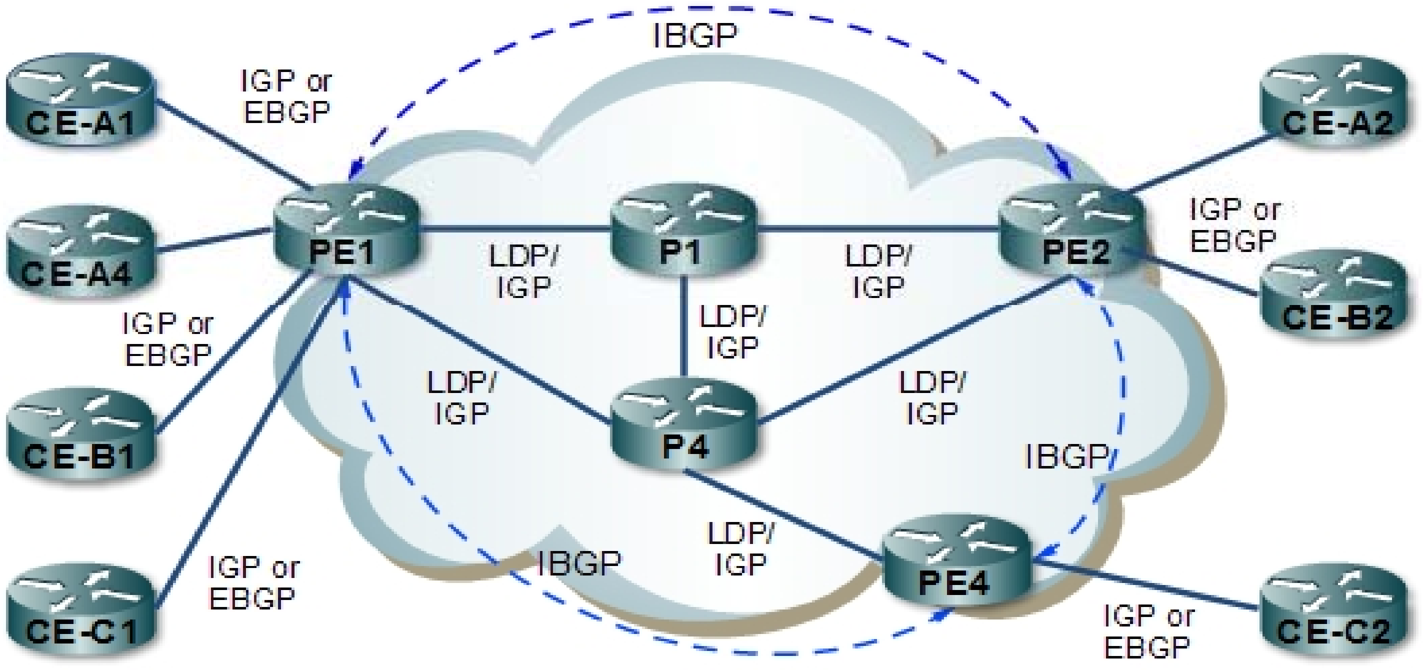

- Multiprotocol BGP is configured between PE routers to carry customer routes.

L3 MPLS VPN Routing Model

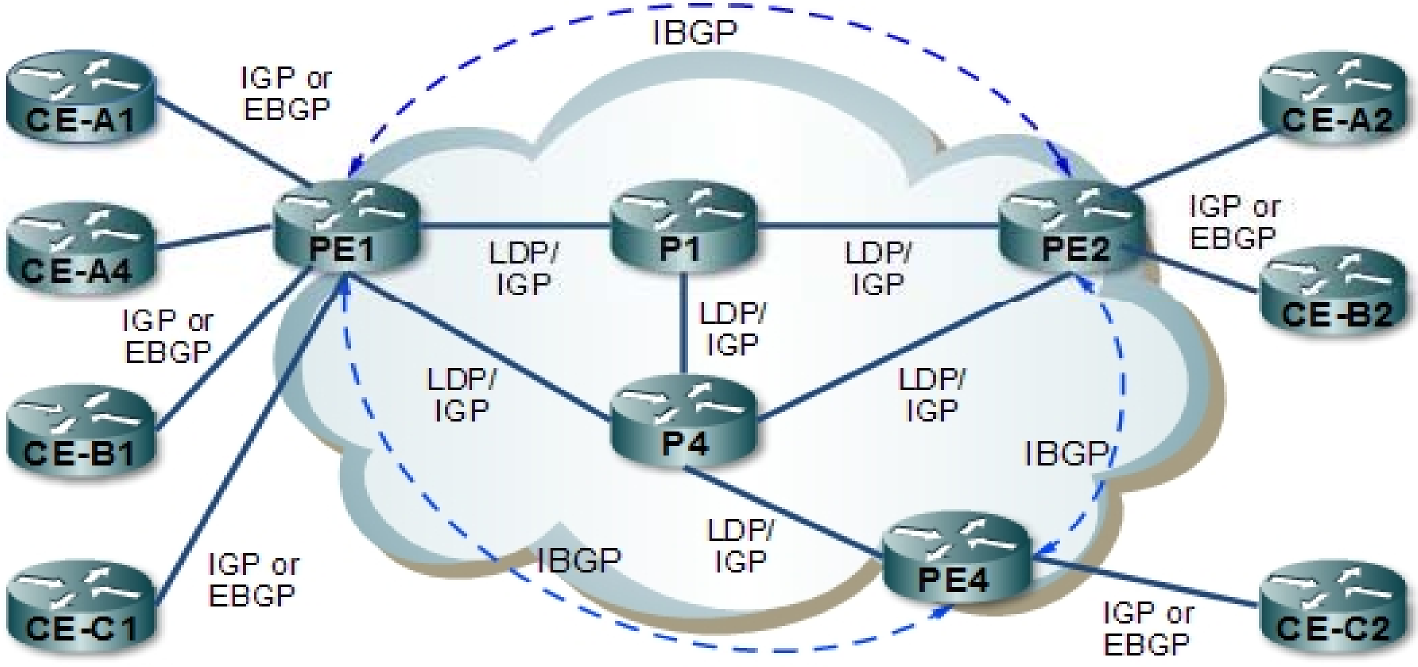

- P routers provide label switching between provider edge routers and are unaware of VPN routes.

Virtual Routing and Forwarding Table (VRF)

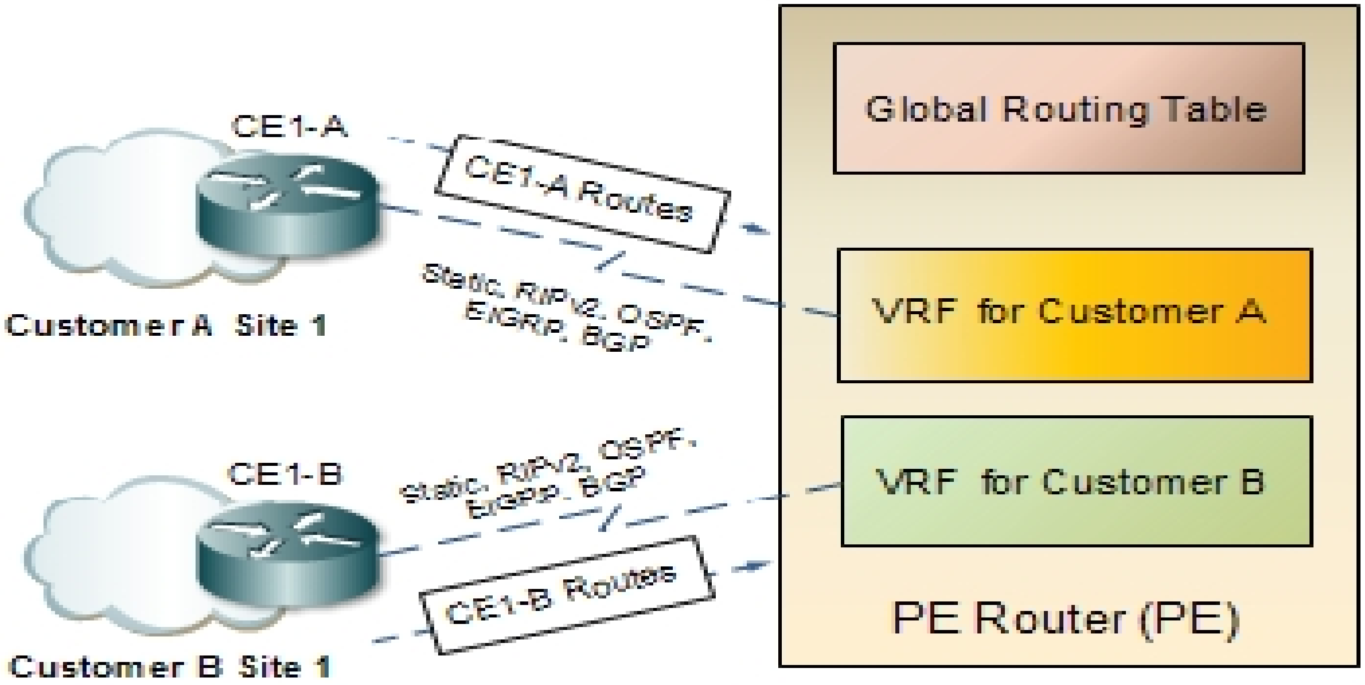

- Customer isolation is achieved on the PE router by the use of virtual routing tables or instances
- The function of a VRF is similar to a global routing table, except that it contains all routes pertaining to a specific VPN versus the global routing table.

Virtual Routing and Forwarding Table (VRF)

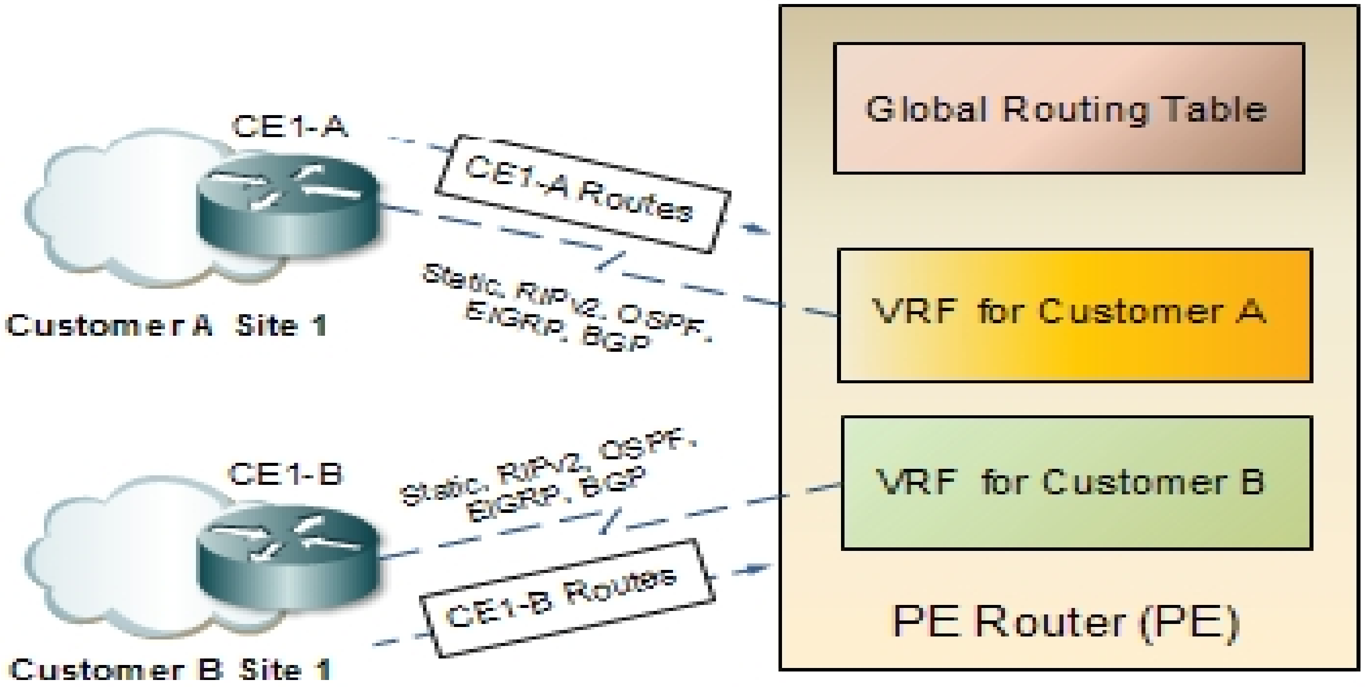

- The VRF also defines the connectivity requirements and protocols for each customer site on a single PE router.
- The VRF defines the interfaces on the local PE router that are part of a specific VPN.

Route Distinguisher

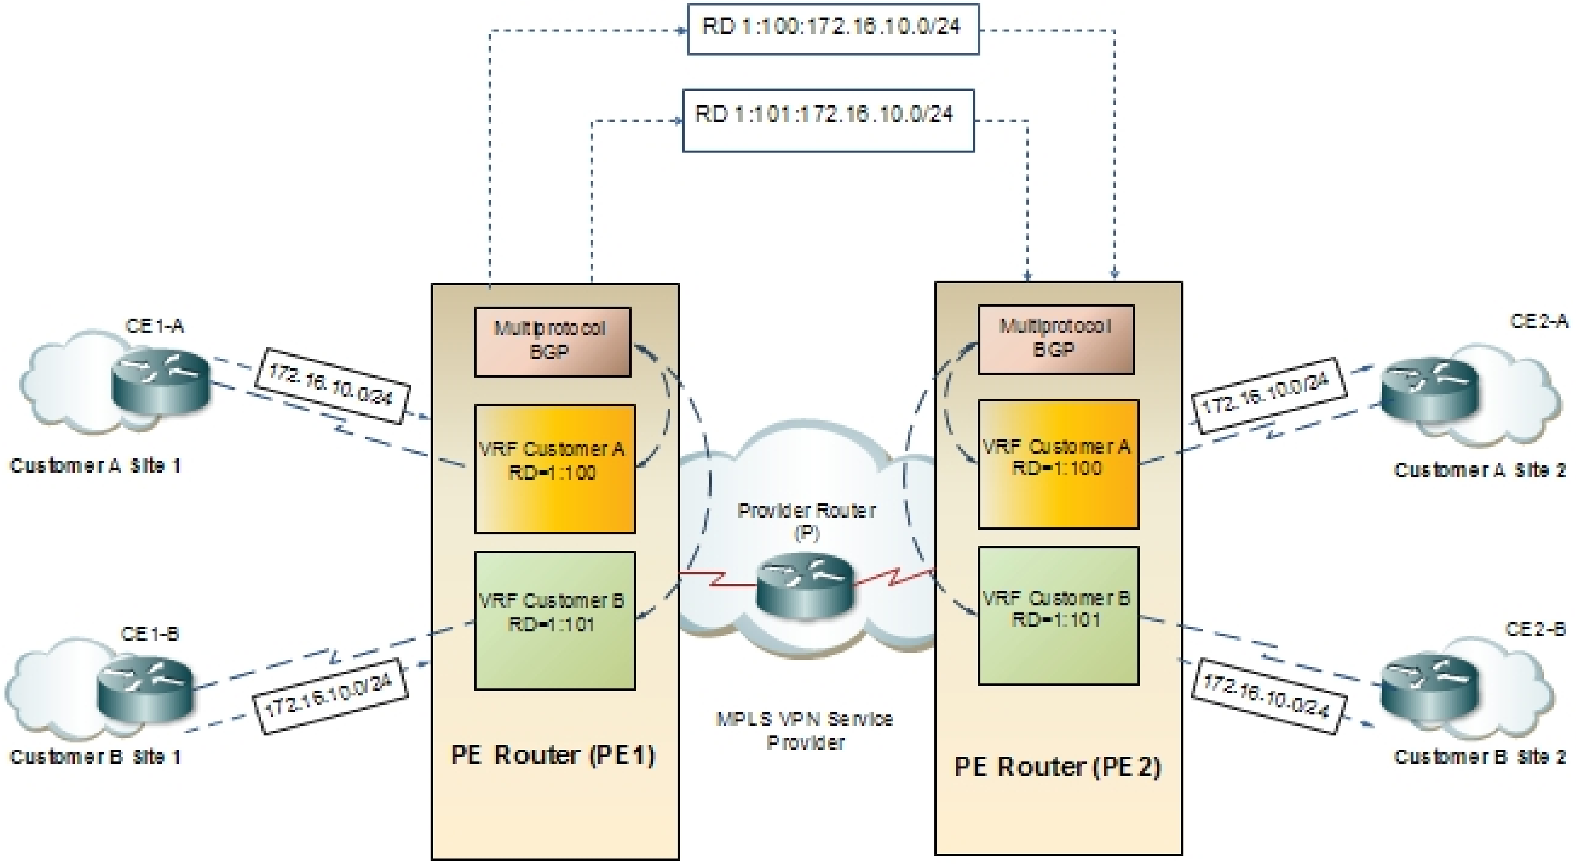

- The RD enable overlapping address spaces in connected customer networks.
- Thus, a unique RD is configured per VRF on the PE router.

Route Distinguisher (Cont.)

- A *RD* is a 64-bit unique identifier that is prepended to the 32-bit customer prefix or route learned from a CE router, which makes it a unique 96-bit address called VPNv4 address that can be transported between the PE routers in the MPLS domain.
- A unique RD is configured per VRF on the PE router.

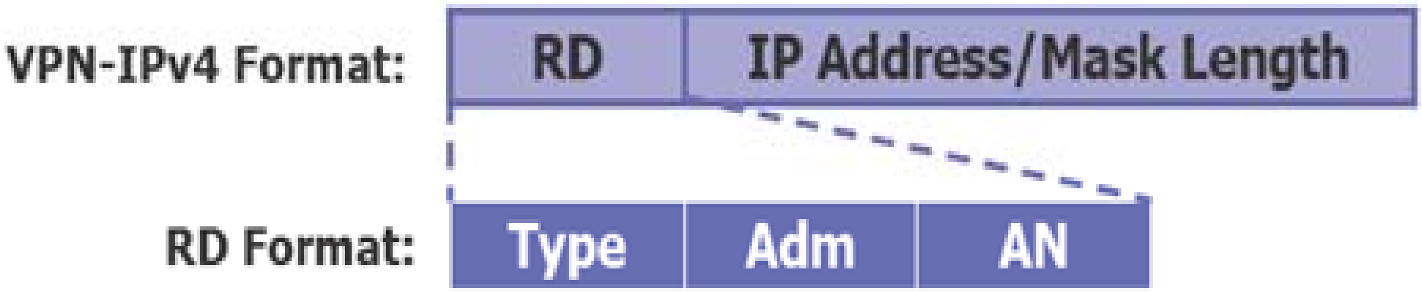

VPN-IPv4 Format: RD | IP Address/Mask Length

RD Format: Type | Adm | AN

Route targets (RT)

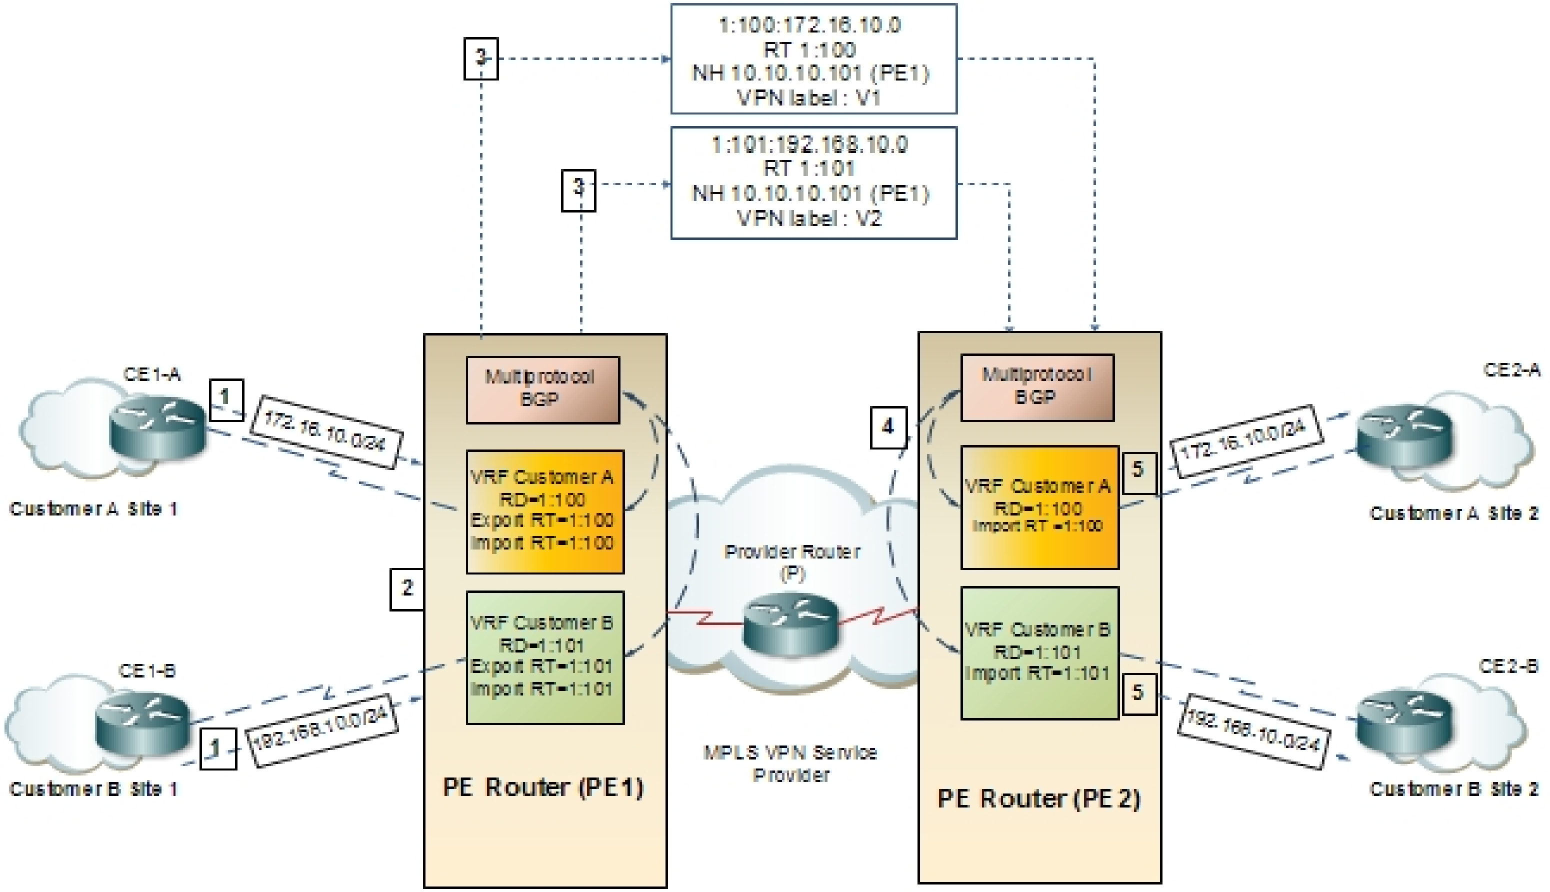

- When a VPN route learned from a CE router is injected into VPNv4 BGP, a list of VPN route target extended community attributes is associated with it.

Route targets (RT)

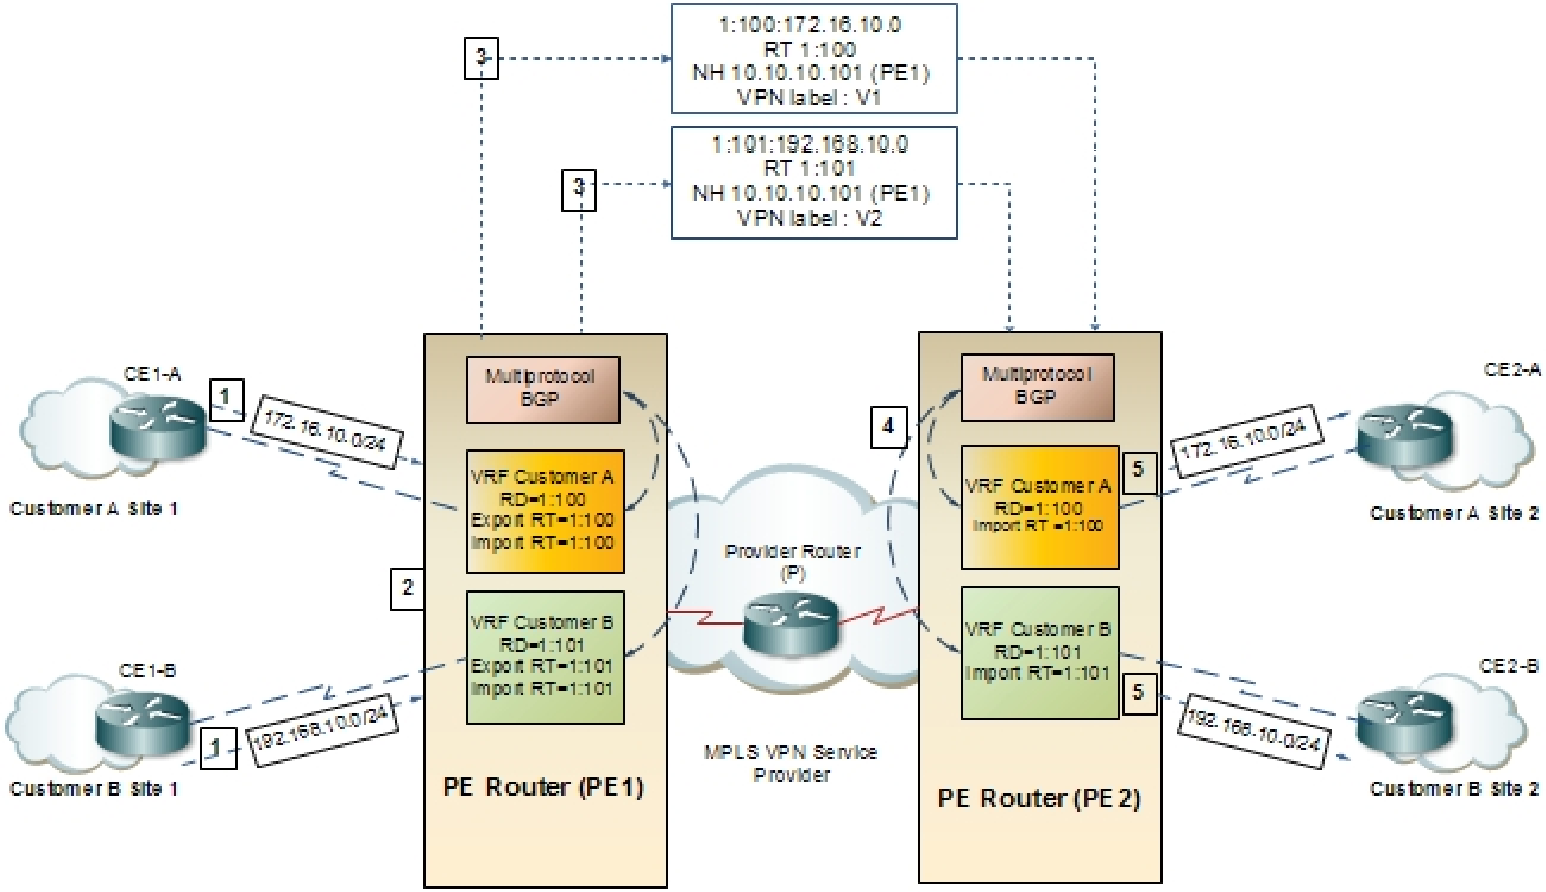

- The export route target is appended to a customer prefix when it is converted to a VPNv4 prefix by the PE router and propagated in MP-BGP updates.

Route targets (RT)

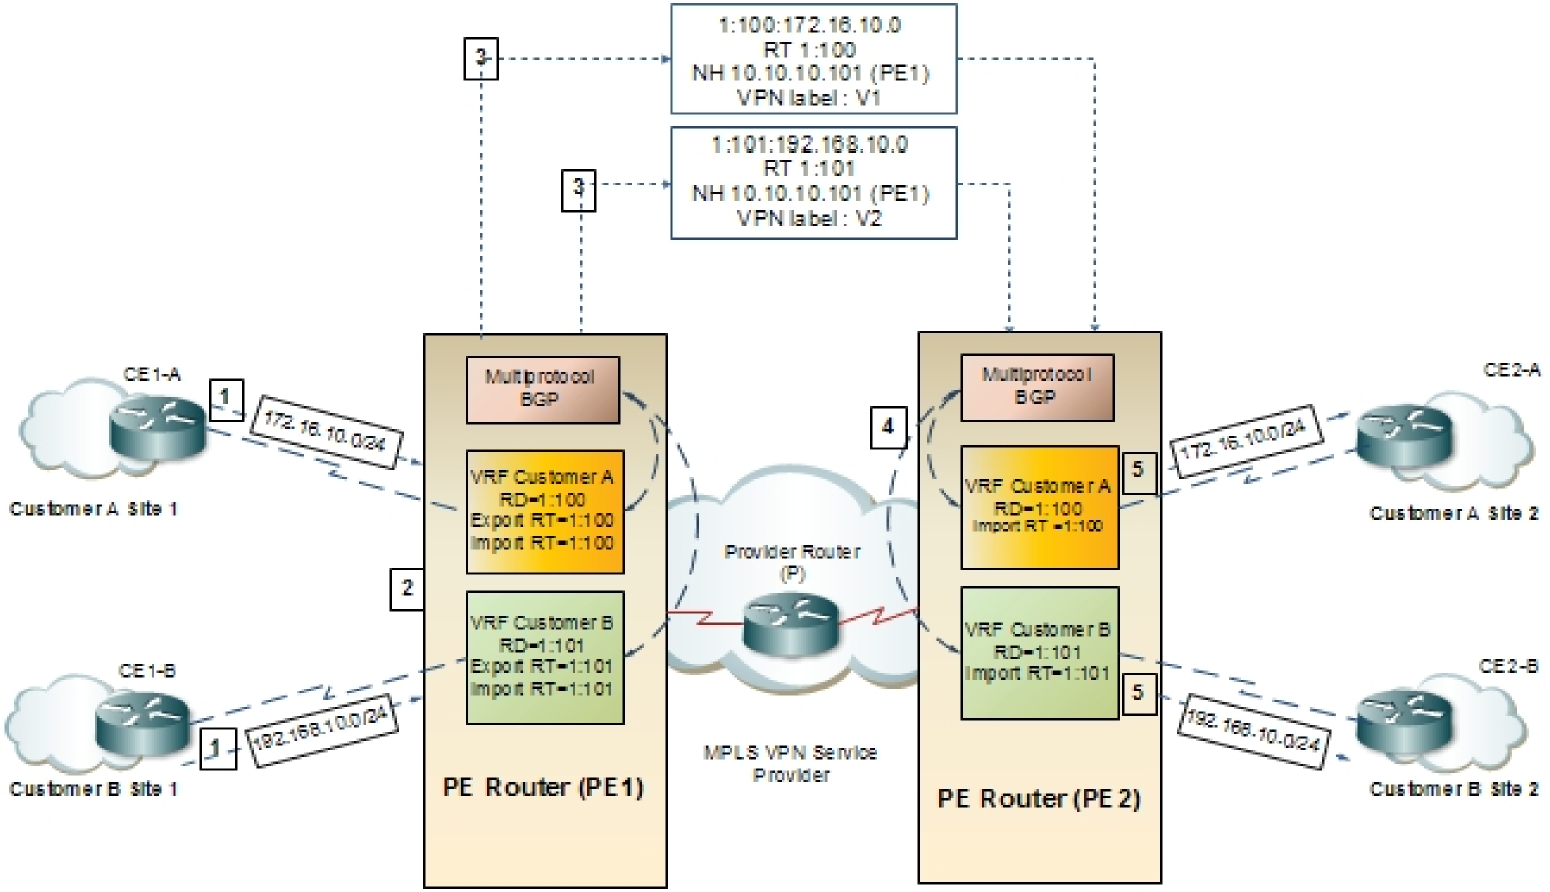

- The import route target is associated with each VRF and identifies the VPN v4 routes to be imported into the VRF for the specific customer.

L3 MPLS VPN Operation

- **Phase 1:** Propagation of VPN routes and distribution of MPLS labels **(Control Plane)**

- **Phase 2:** Packet forwarding **(Data Plane)**

Control Plane Operation

- Taking the next figure as an example Propagation of VPN routes and distribution of MPLS labels takes place in three different stages

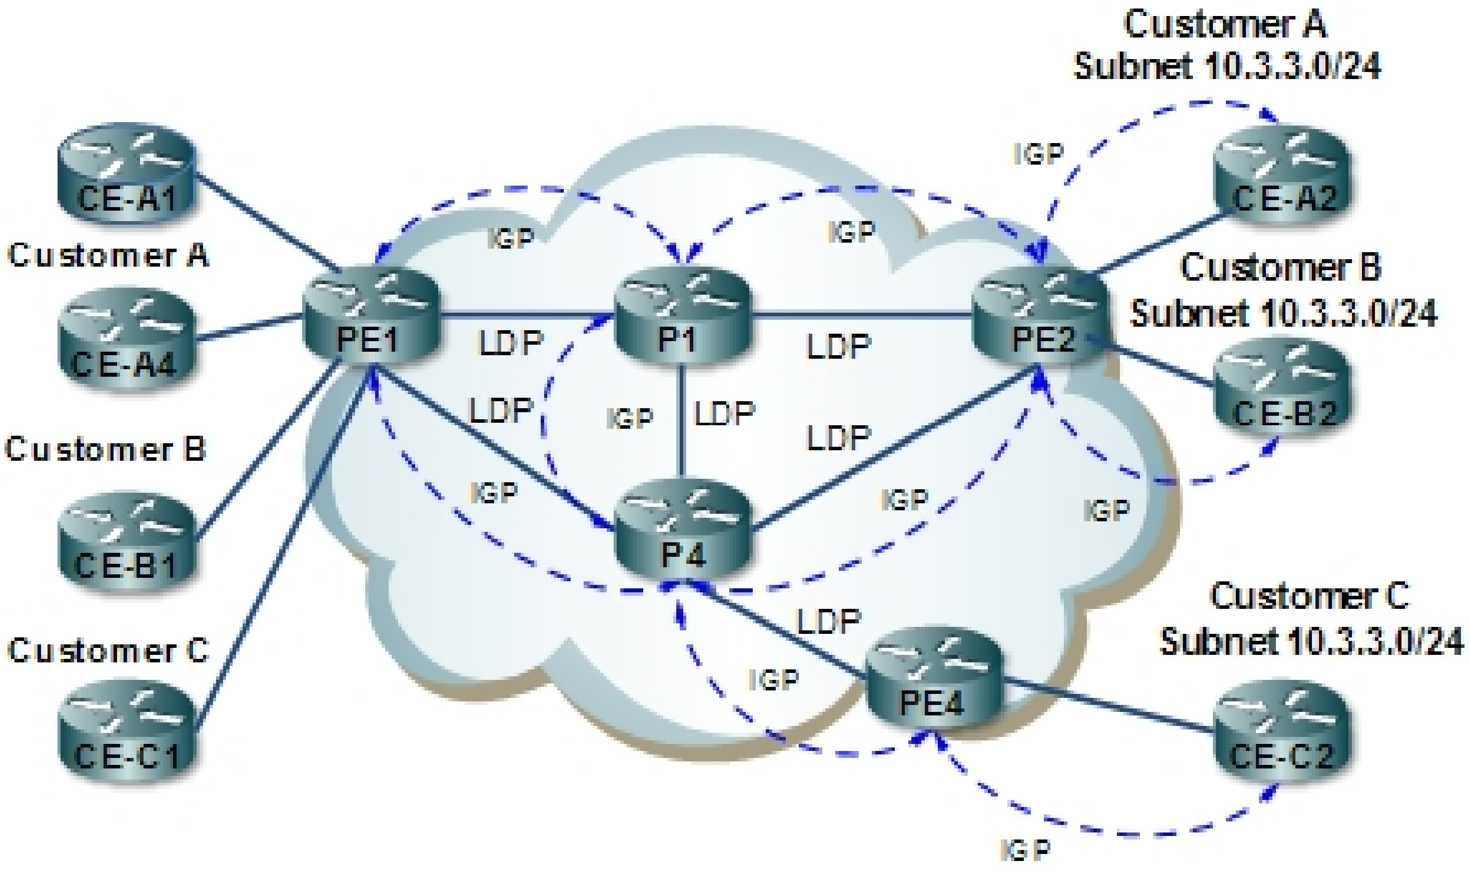

Stage 1

- **Stage 1:** PE routers receive IPv4 routing updates from CE routers and populate these routes into the appropriate VRF table.

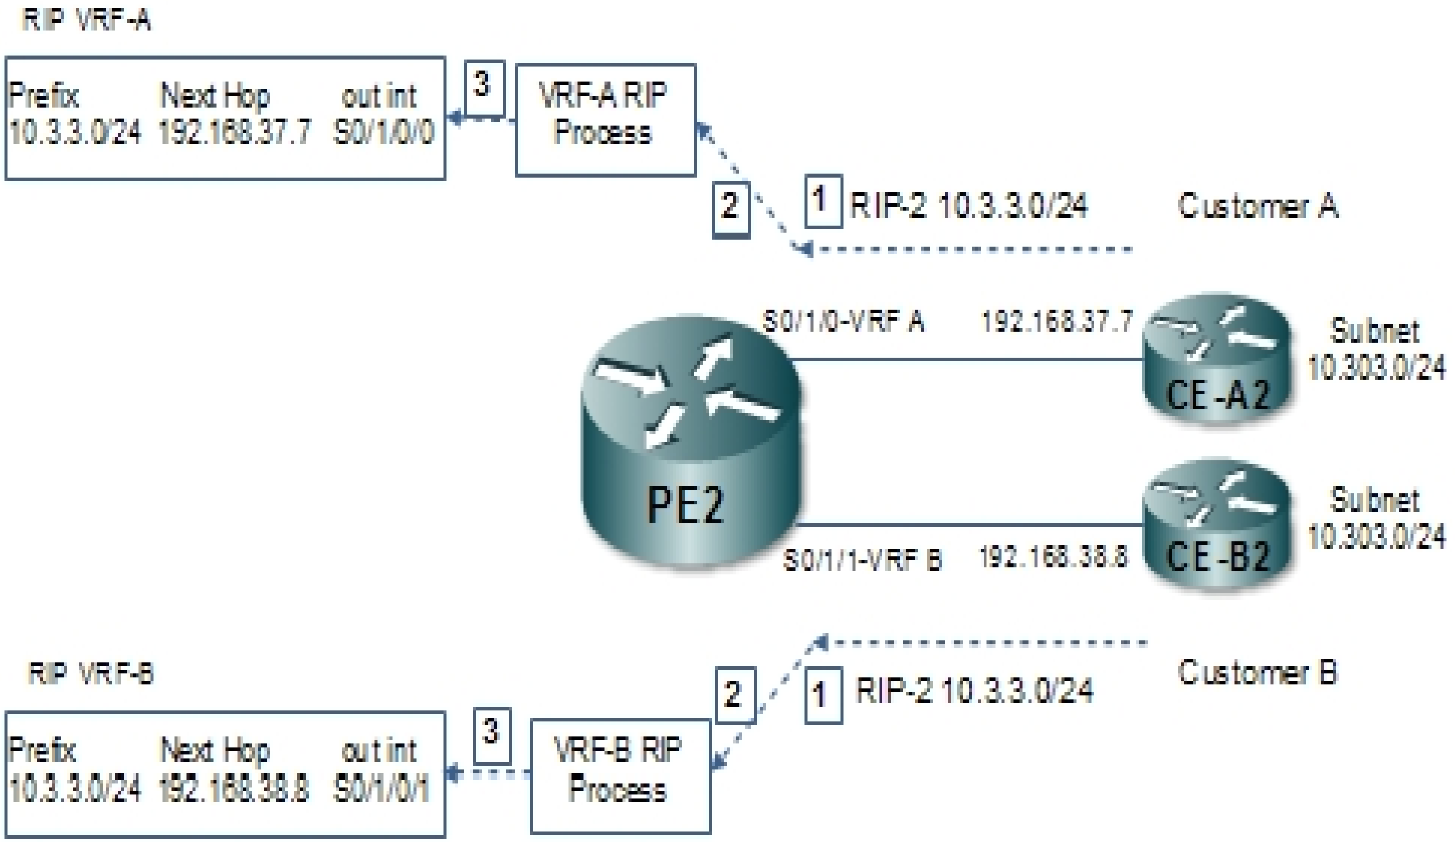

RIP VRF-A

Prefix	Next Hop	out int
10.3.3.0/24	192.168.37.7	S0/1/0/0

RIP VRF-B

Prefix	Next Hop	out int
10.3.3.0/24	192.168.38.8	S0/1/0/1

Stage 2

- **Stage 2:** PE routers export VPN routes from VRF tables into MP-IBGP and propagate them with VPN label as VPNv4 routes via MP-IBGP to other remote PE routers.

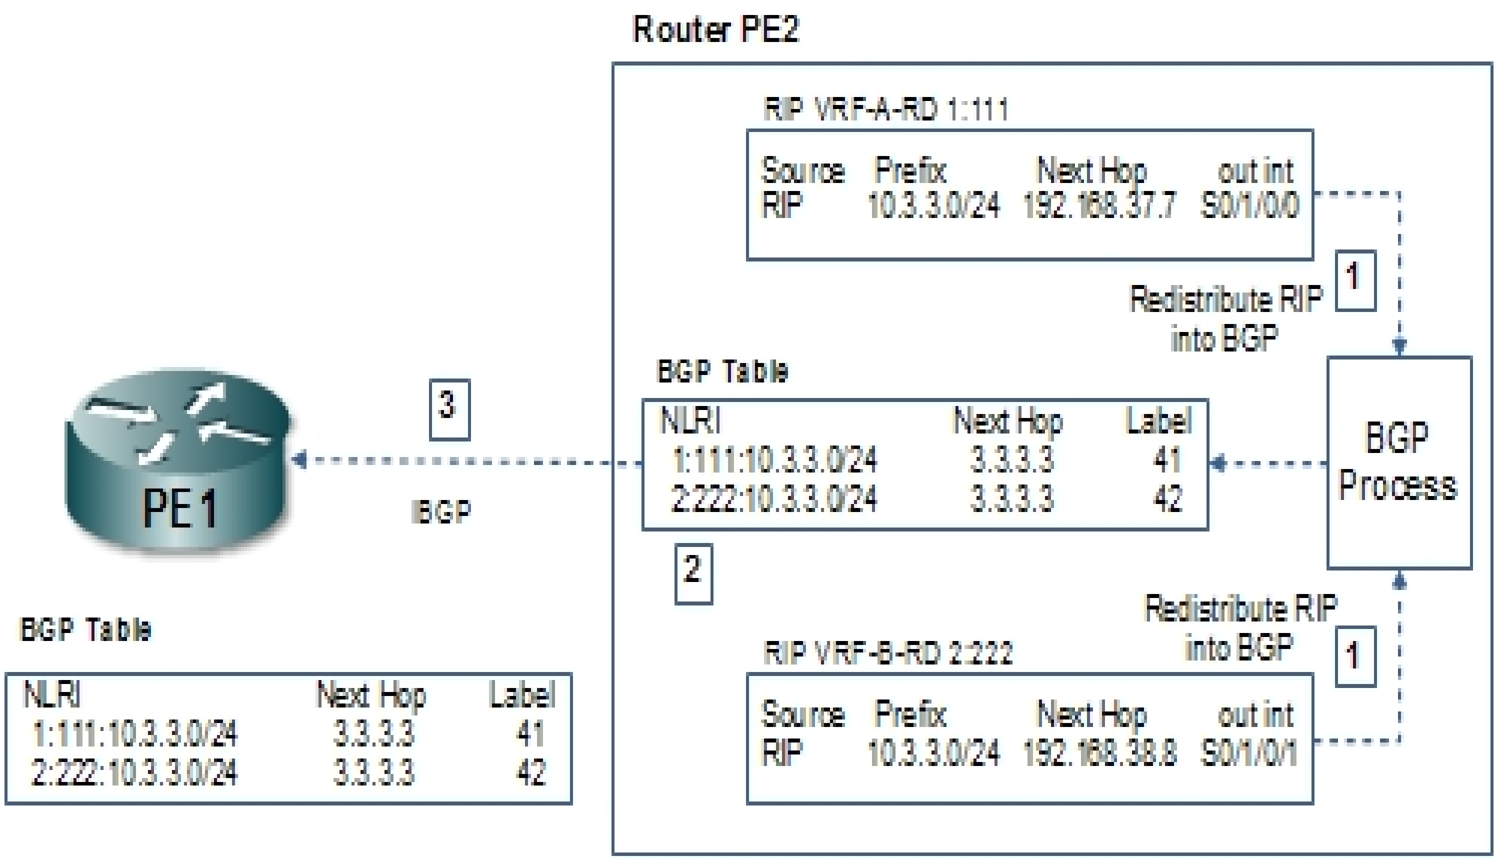

RIP VRF-A-RD 1:111

Source	Prefix	Next Hop	out int
RIP	10.3.3.0/24	192.168.37.7	S0/1/0/0

BGP Table

NLRI	Next Hop	Label
1:111:10.3.3.0/24	3.3.3.3	41
2:222:10.3.3.0/24	3.3.3.3	42

RIP VRF-B-RD 2:222

Source	Prefix	Next Hop	out int
RIP	10.3.3.0/24	192.168.38.8	S0/1/0/1

BGP Table

NLRI	Next Hop	Label
1:111:10.3.3.0/24	3.3.3.3	41
2:222:10.3.3.0/24	3.3.3.3	42

Stage 3

- **Stage 3:** The remote PE routers on receiving MP-IBGP updates will import the incoming VPNv4 routes into their respective VRF tables according to the import RTs. The VPNv4 routes installed in VRF tables are then converted back to IPv4 routes and propagated to the CE routers.

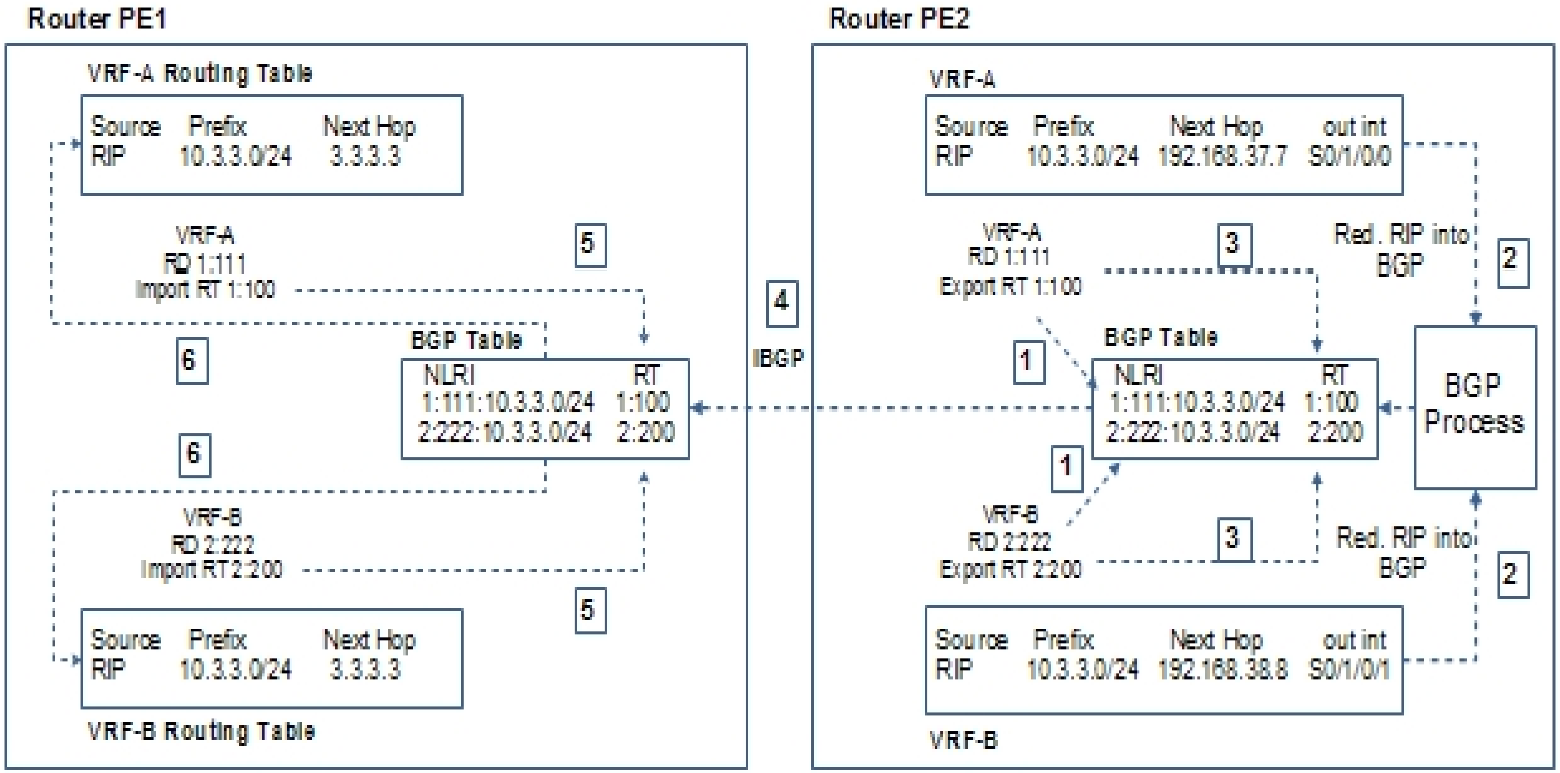

Router PE1, VRF-A Routing Table:

Source	Prefix	Next Hop
RIP	10.3.3.0/24	3.3.3.3

Router PE1, BGP Table:

NLRI	RT
1:111:10.3.3.0/24	1:100
2:222:10.3.3.0/24	2:200

Router PE1, VRF-B Routing Table:

Source	Prefix	Next Hop
RIP	10.3.3.0/24	3.3.3.3

Router PE2, VRF-A:

Source	Prefix	Next Hop	out int
RIP	10.3.3.0/24	192.168.37.7	S0/1/0/0

Router PE2, BGP Table:

NLRI	RT
1:111:10.3.3.0/24	1:100
2:222:10.3.3.0/24	2:200

Router PE2, VRF-B:

Source	Prefix	Next Hop	out int
RIP	10.3.3.0/24	192.168.38.8	S0/1/0/1

Summary for Control Plane Operation

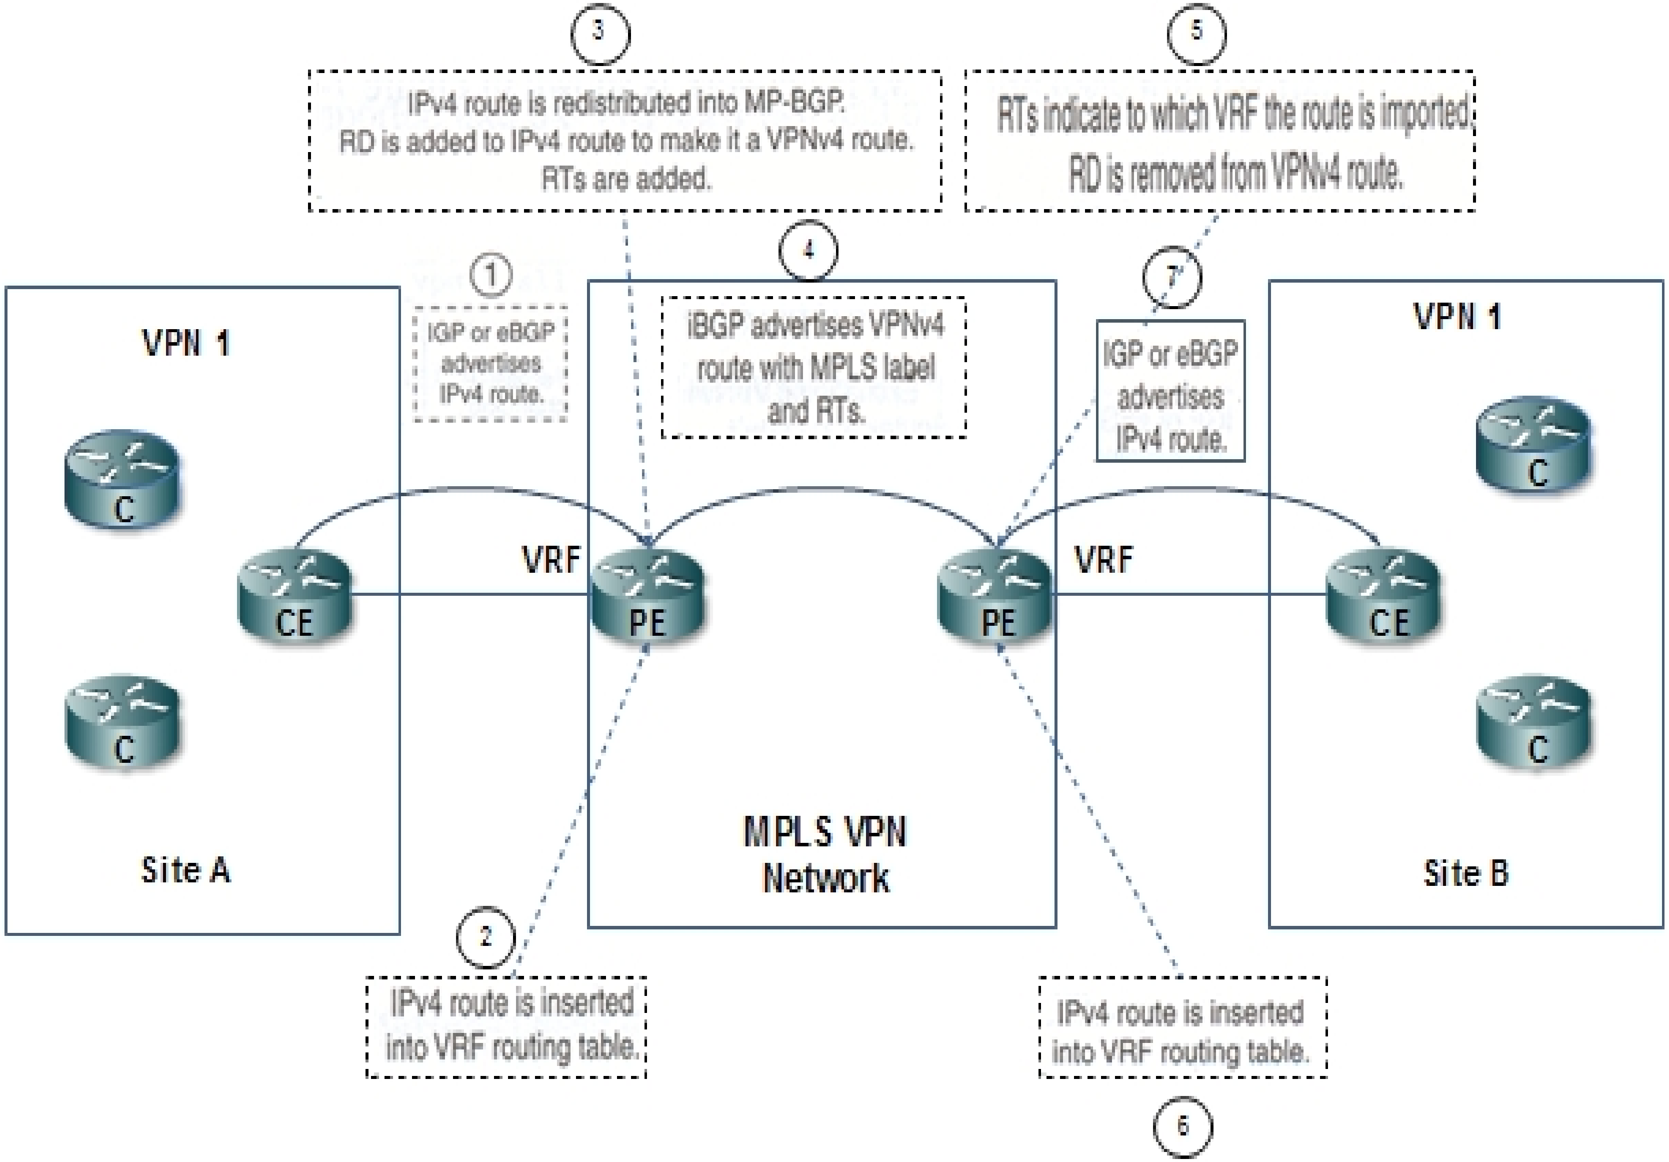

3
IPv4 route is redistributed into MP-BGP.
RD is added to IPv4 route to make it a VPNv4 route.
RTs are added.
5
RTs indicate to which VRF the route is imported.
RD is removed from VPNv4 route.
1
IGP or eBGP advertises IPv4 route.
4
iBGP advertises VPNv4 route with MPLS label and RTs.
7
IGP or eBGP advertises IPv4 route.
VPN 1
C
CE
C
Site A
VRF
PE
PE
VRF
MPLS VPN Network
VPN 1
C
CE
C
Site B
2
IPv4 route is inserted into VRF routing table.
IPv4 route is inserted into VRF routing table.
6

Example-Control Plane Operation

1. IPv4 update for network 172.16.10.0 is received by the egress PE router

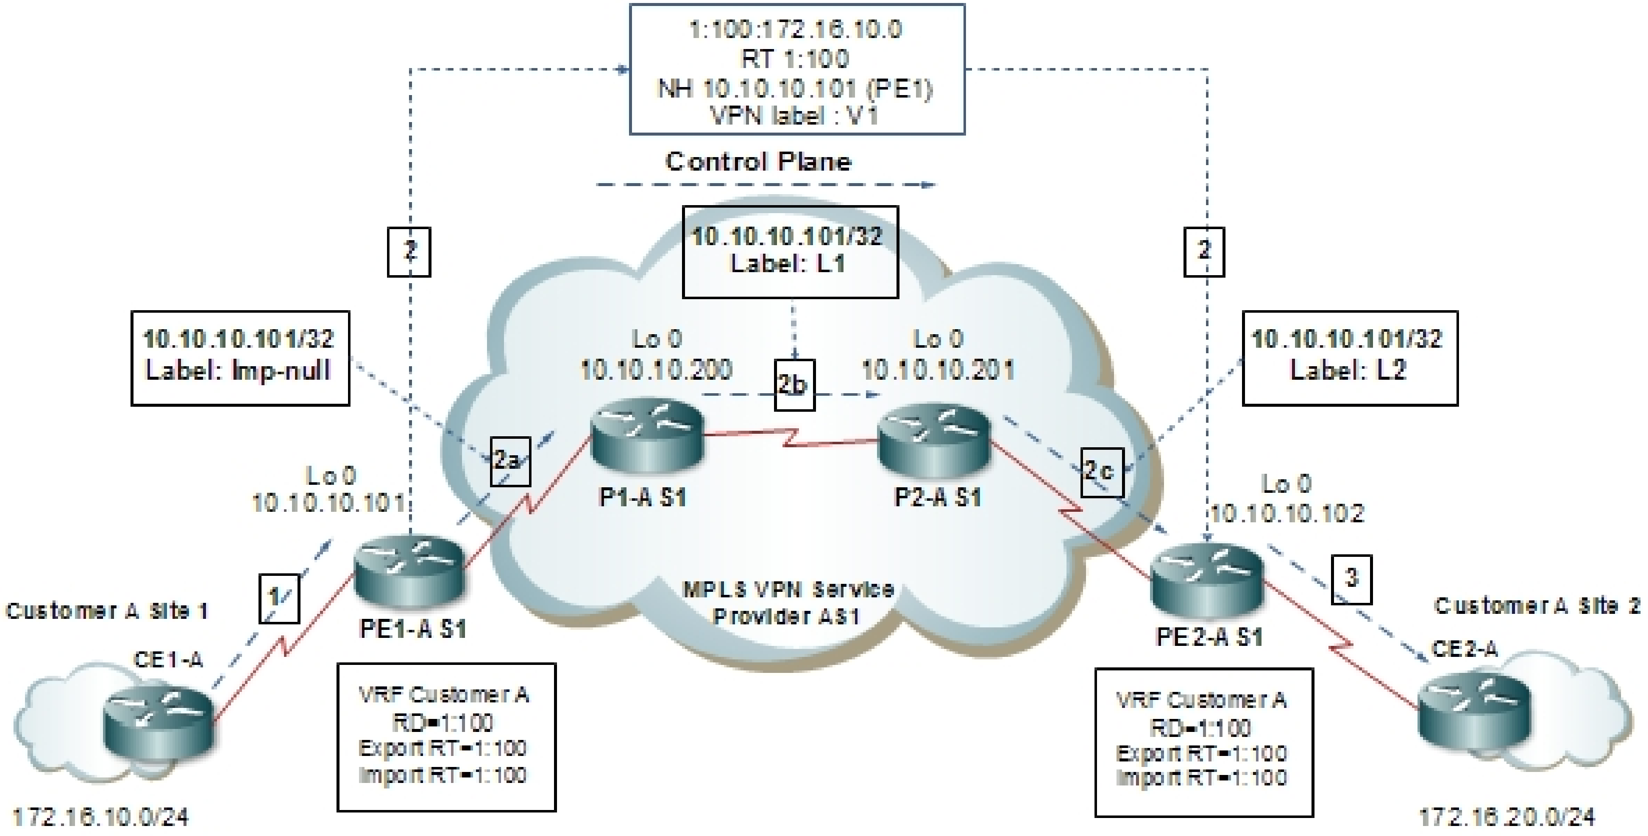

1:100:172.16.10.0
RT 1:100
NH 10.10.10.101 (PE1)
VPN label : V1

Control Plane

10.10.10.101/32
Label: L1

10.10.10.101/32
Label: Imp-null

10.10.10.101/32
Label: L2

VRF Customer A
RD=1:100
Export RT=1:100
Import RT=1:100

VRF Customer A
RD=1:100
Export RT=1:100
Import RT=1:100

Example-Control Plane Operation

2. PE1-AS1 accepts and transforms the IPv4 route, 172.16.10.0/24, to a VPN v4 route by assigning an RD 1:100 and RT 1:100. It allocates a label V1 and rewrites the next-hop attribute to the PE1-AS1 loopback0 IP address 10.10.10.101.

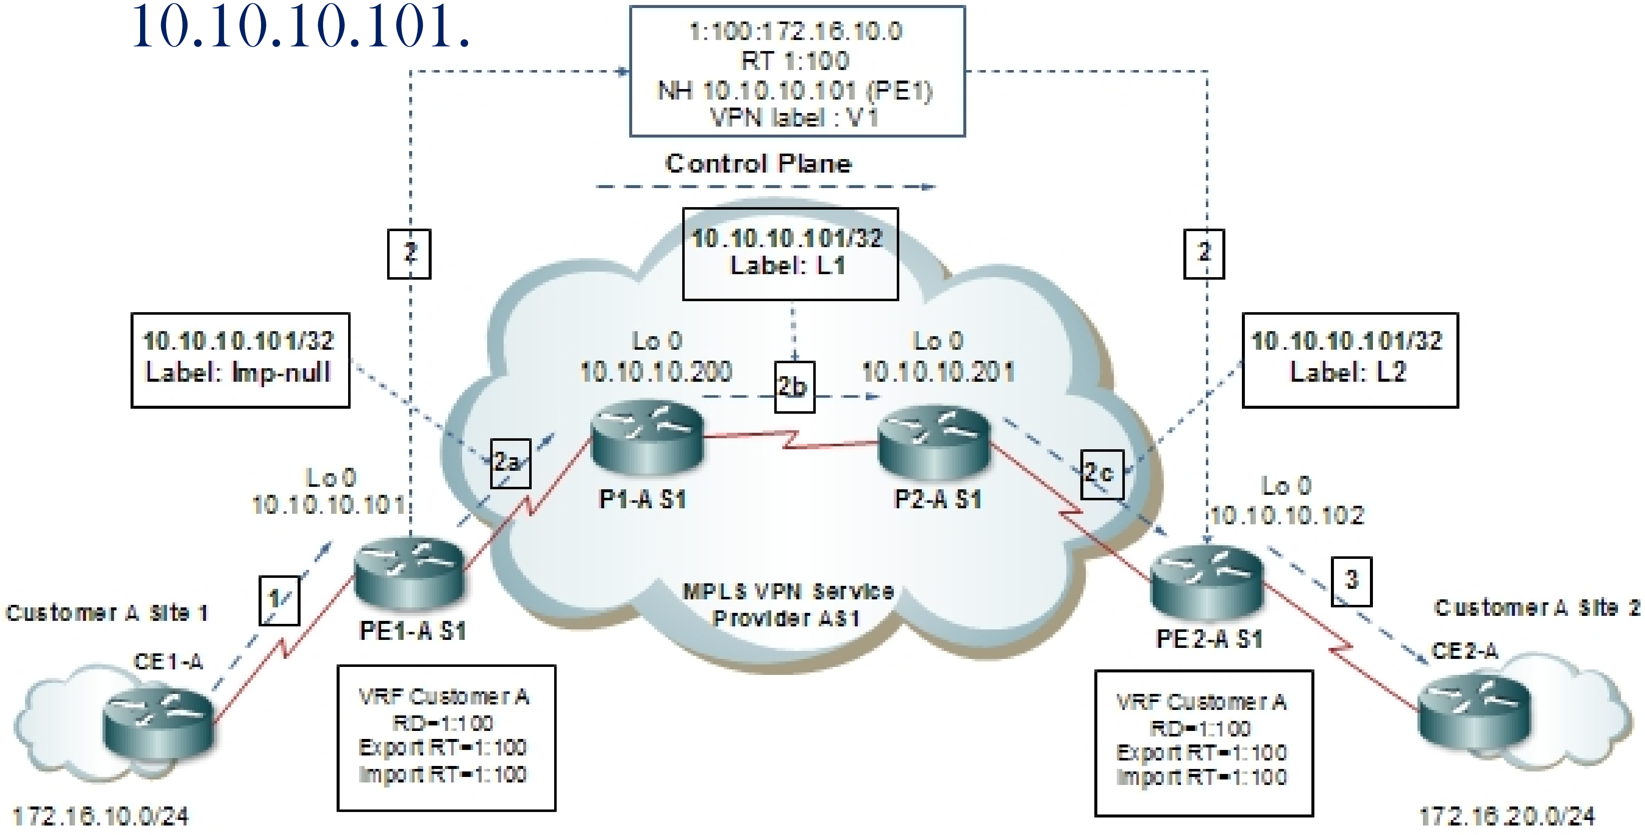

Example-Control Plane Operation

2a. Edge LSR PE2-AS1 requests a label for the 10.10.10.101/32 prefix using LDP from LSR P2-AS1 then from P1-AS1 then from Edge LSR PE1-AS1. Edge LSR PEl-AS1 allocates a label of implicit-null and sends it to P1-AS1.

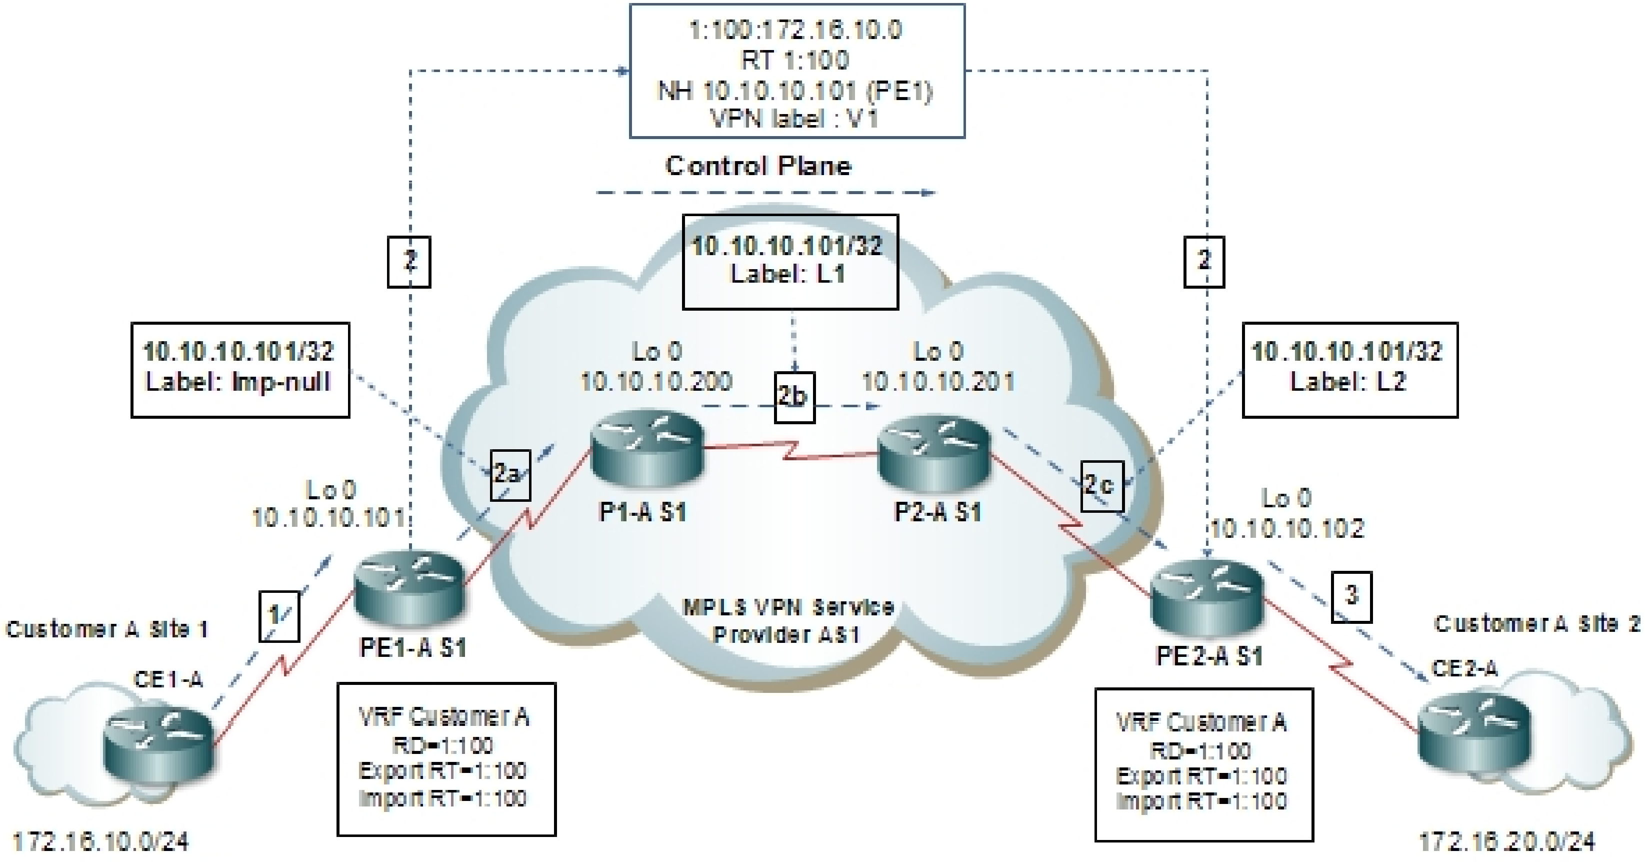

Example-Control Plane Operation

2b. P1-AS1 uses the implicit-null label received from PE1-AS1 as its outbound label value, allocates a label (L1) to prefix 10.10.10.101/32, and sends this label value to P2-AS1 via LDP.

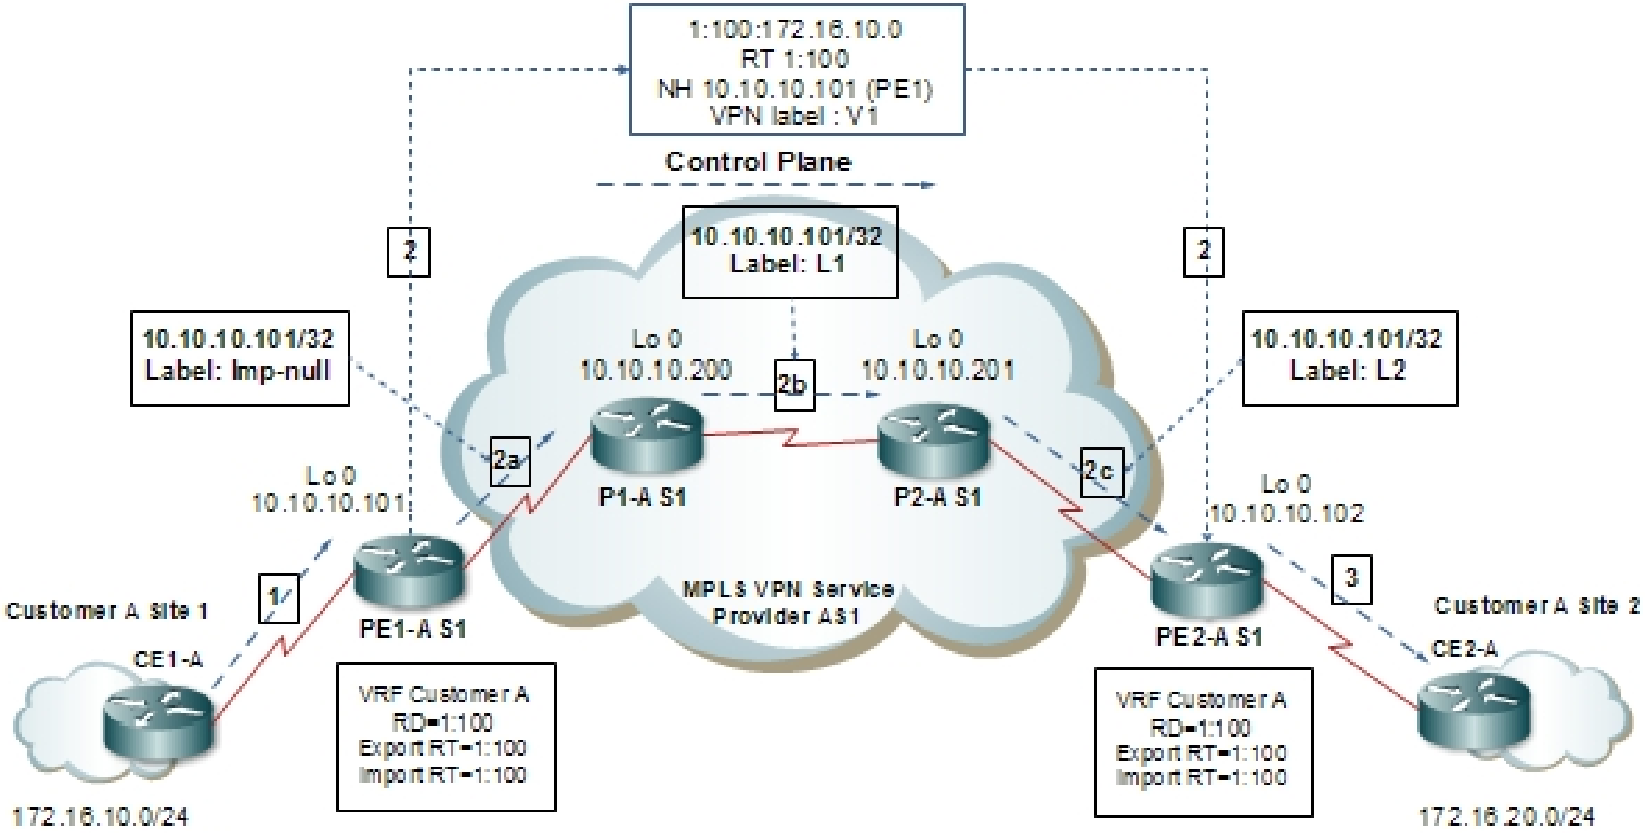

1:100:172.16.10.0
RT 1:100
NH 10.10.10.101 (PE1)
VPN label : V1

10.10.10.101/32
Label: L1

10.10.10.101/32
Label: Imp-null

10.10.10.101/32
Label: L2

VRF Customer A
RD=1:100
Export RT=1:100
Import RT=1:100

VRF Customer A
RD=1:100
Export RT=1:100
Import RT=1:100

Example-Control Plane Operation

2c. P2-AS1 uses the label (L1) received from PI-AS1 as its outbound label value, allocates a label (L2) to prefix 10.10.10.101/32, and sends this label value to PE2-AS1 via LDP.

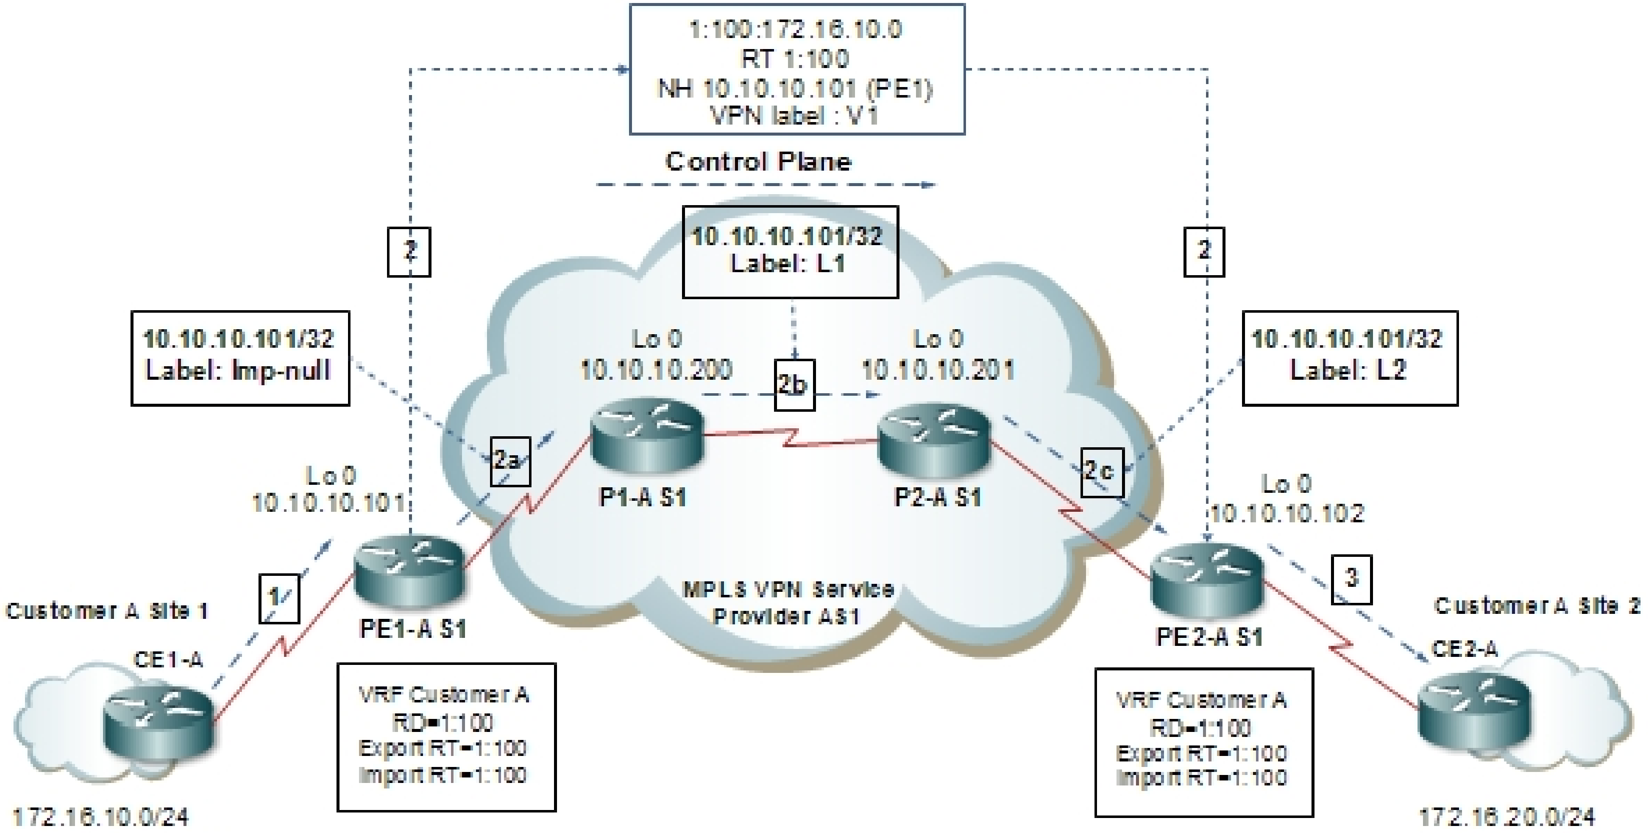

Example-Control Plane Operation

3. PEl-AS1 has the VRF configured to accept routes with RT 1:100 and therefore translates the VPNv4 update to IPv4 and inserts the route in VRF A. It then propagates this route to the CE2-A.

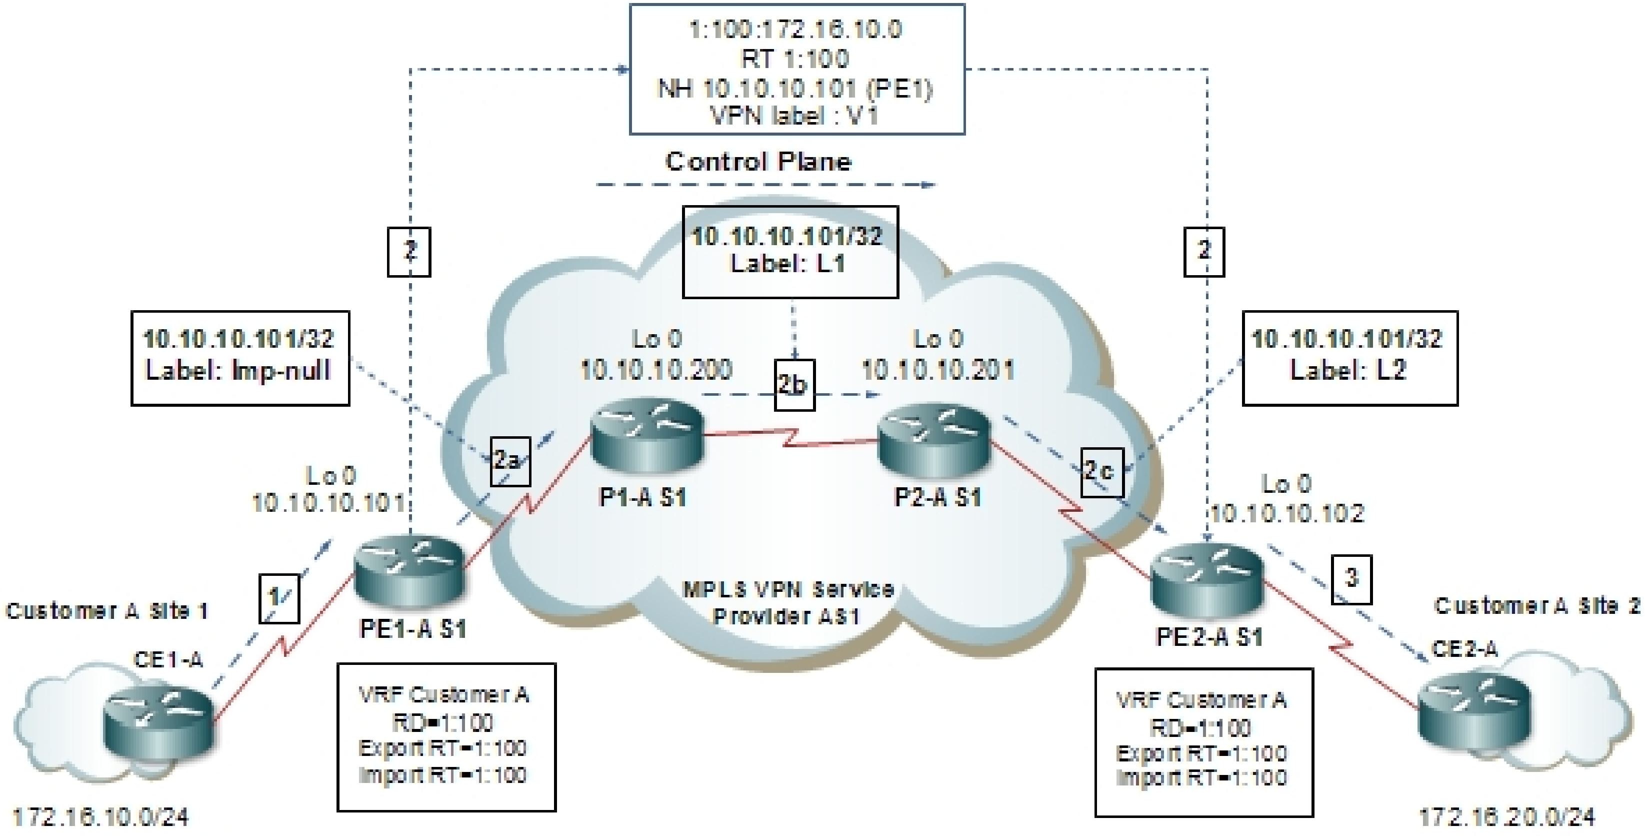

Example-Data Plane Operation

1. CE2-A originates a data packet with the source address of 172.16.20.1 and destination of 172.16.10.1.

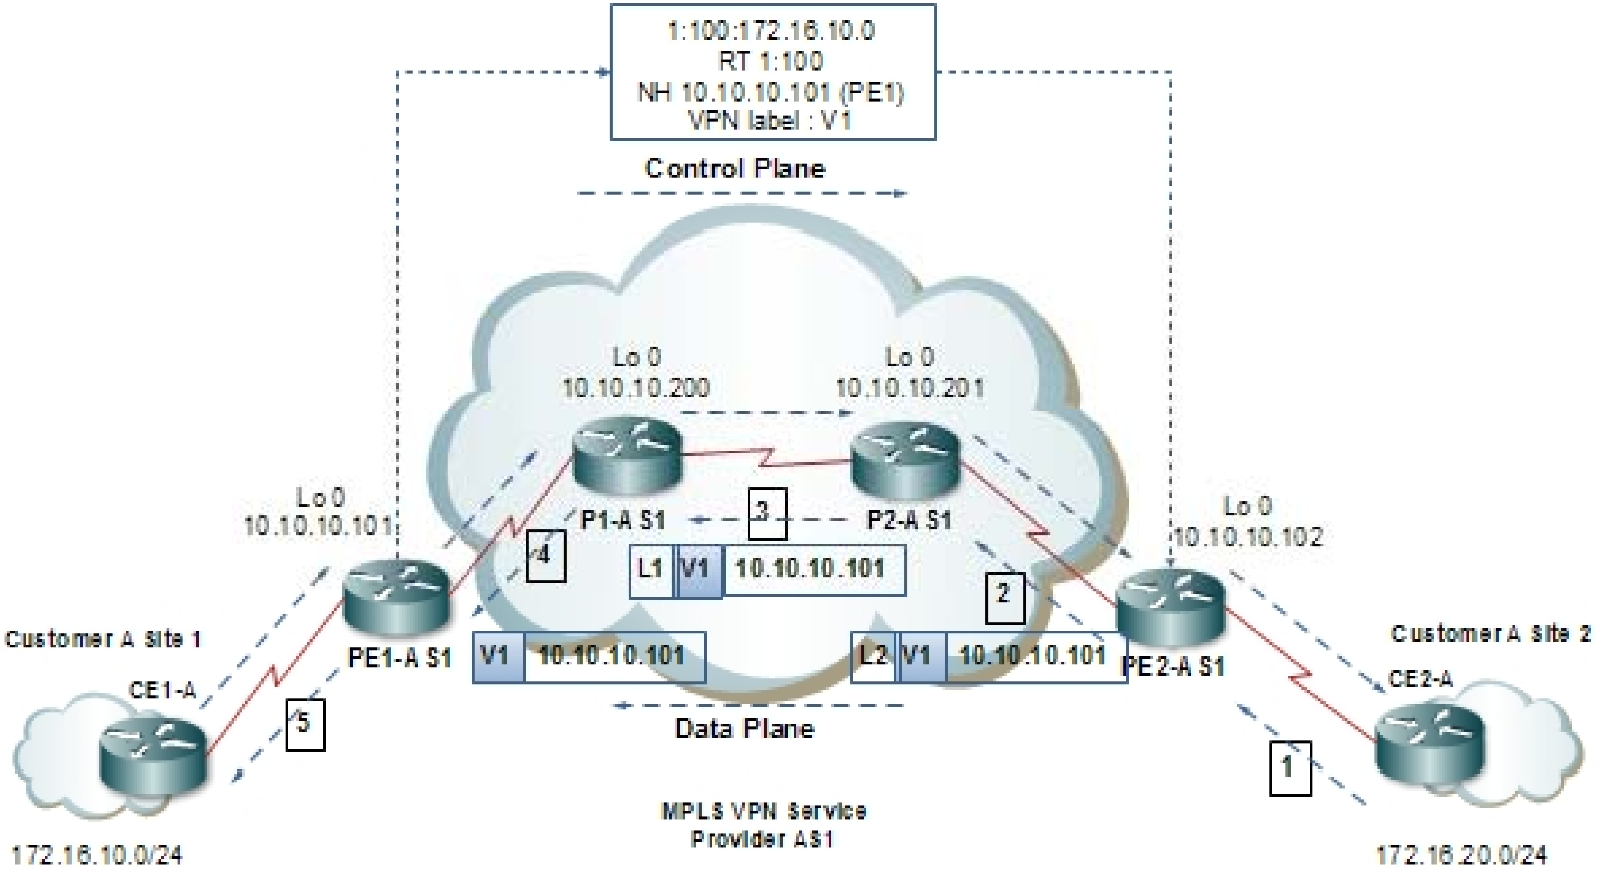

Example-Data Plane Operation

2. PE2-AS1 receives the data packet and appends the VPN label V1 and LDP label L2 and forwards the packet to P2-AS1.

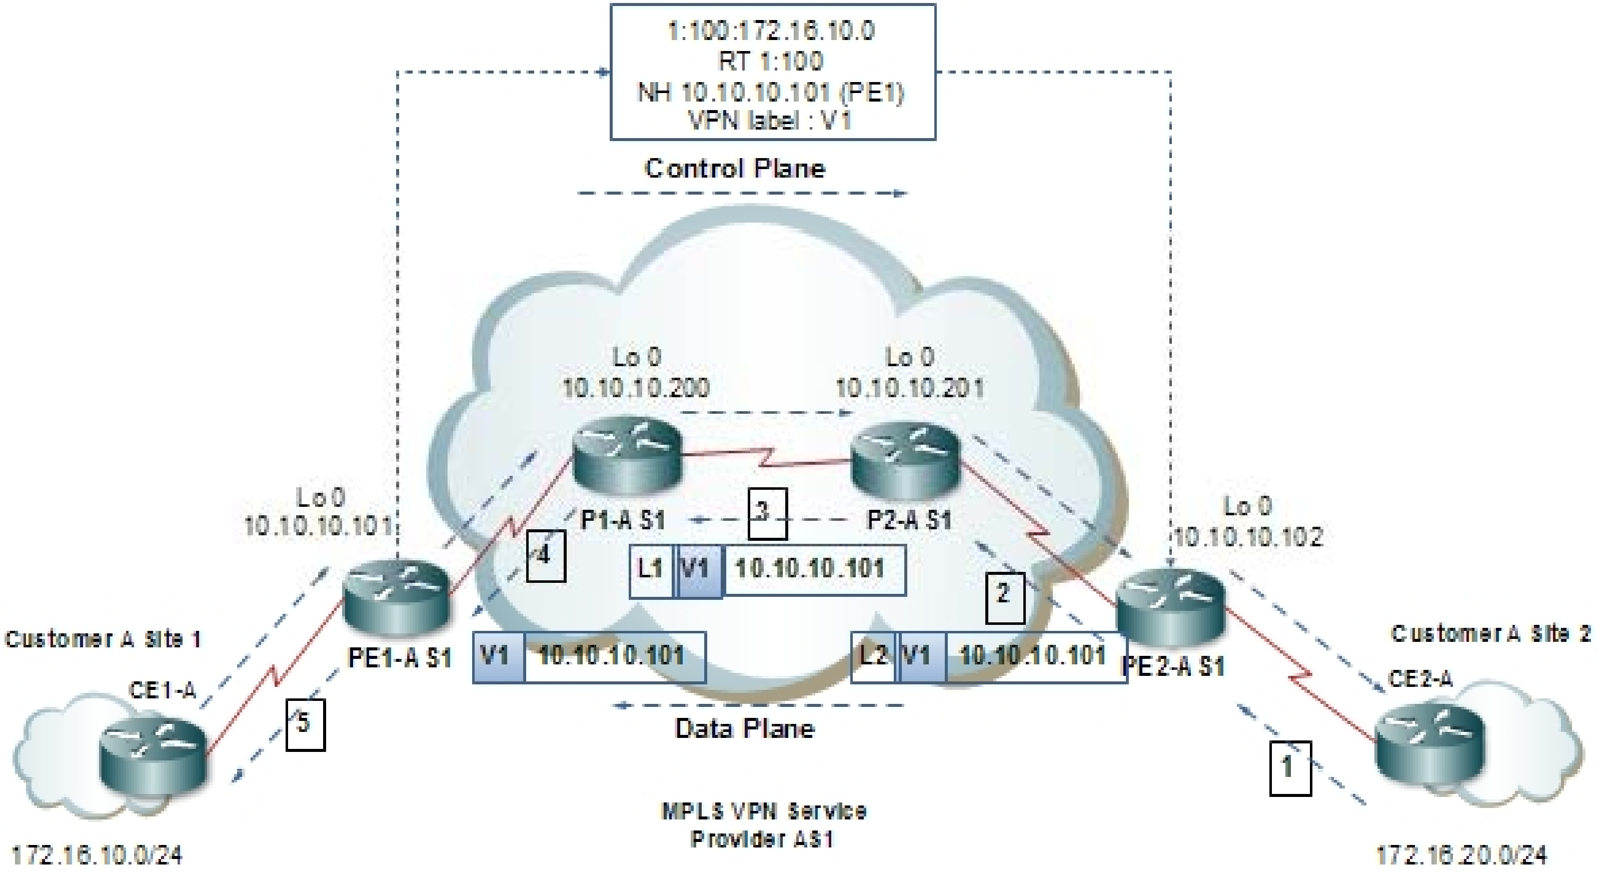

Example-Data Plane Operation

3. P2-AS1 receives the data packet destined to 172.16.10.1 and swaps LDP label L2 with L1.

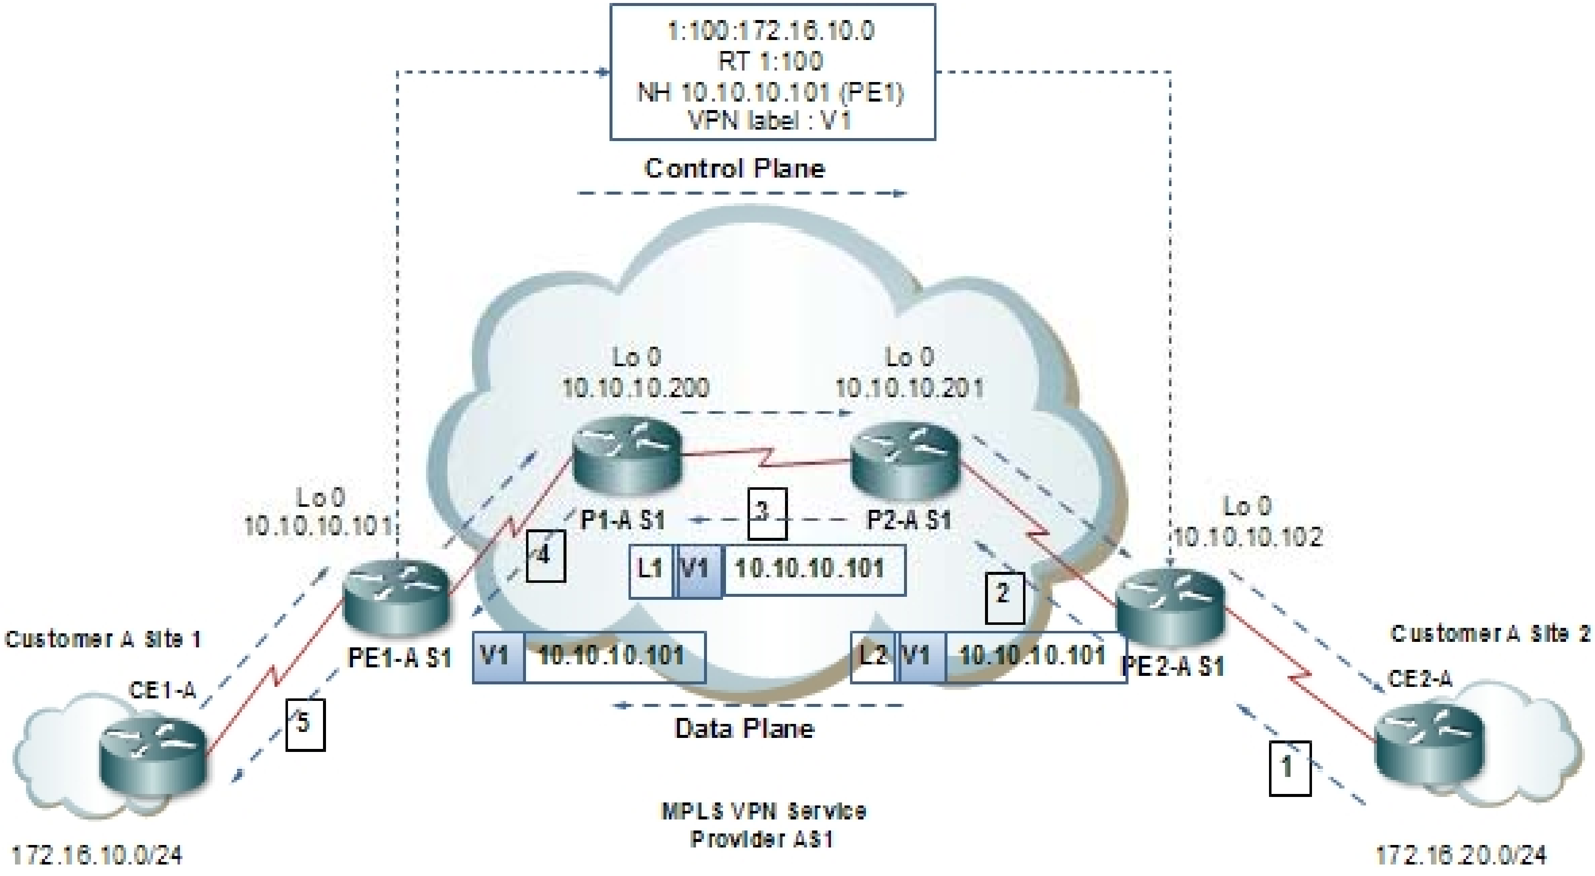

Example-Data Plane Operation

4. P1-AS1 receives the data packet destined to 172.16.10.1 and pops the top label. The resulting labeled packet with VPN Label V1 is forwarded to PE1-AS1.

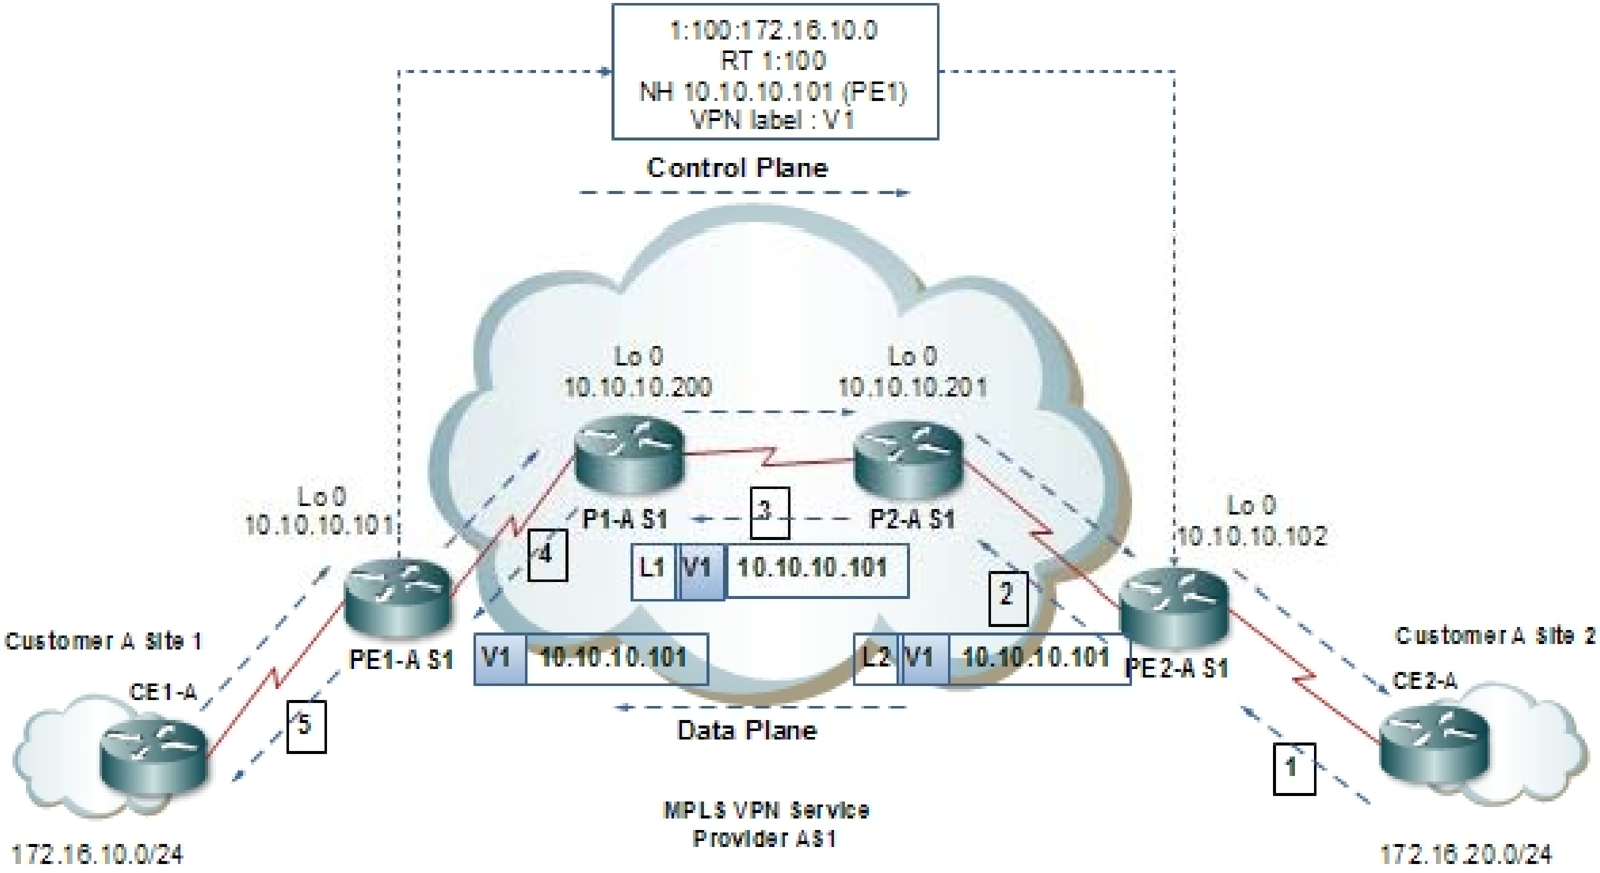

Example-Data Plane Operation

5. PE1-AS1 pops the VPN label and forwards the data packet to CE1-A where the 172.16.10.0 network is located.

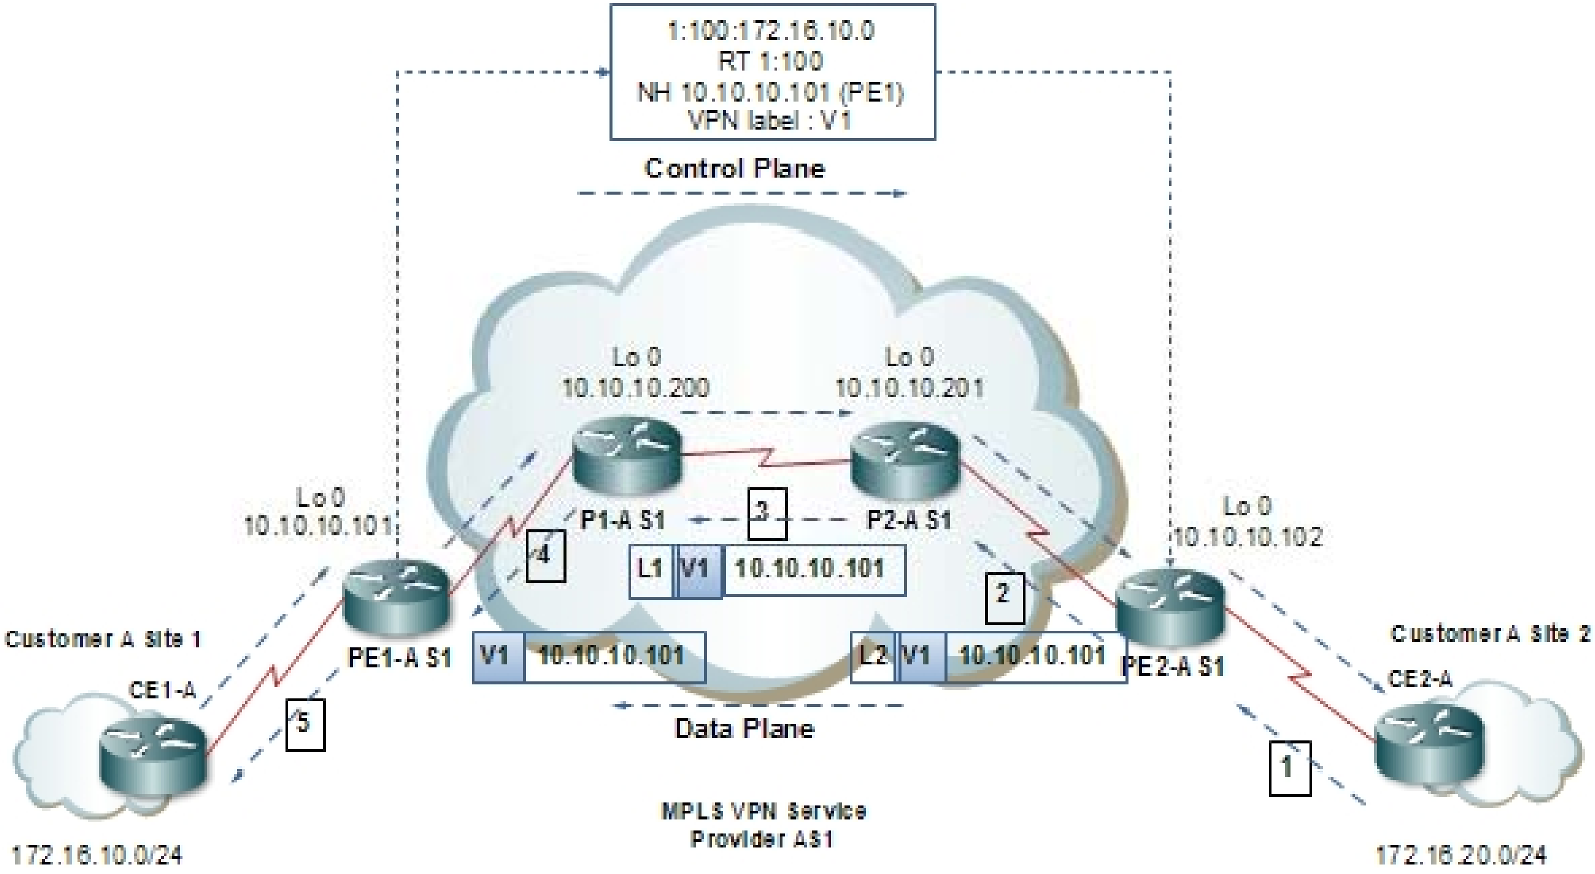

Layer 2 VPN

- Customers may desire to extend their current Layer 2 infrastructure (frame relay, ATM, Ethernet, VLANs, TDM, transparent LAN services, etc.).
- IP-based Layer 3 VPNs will not satisfy any of these requirements; instead, a Layer 2 solution is required.
- MPLS-based Layer 2 VPNs prepends a label to a Layer 2 PDU and then forwarding the packet across the MPLS backbone.

Layer 2 VPN Components

- The Martini draft builds on some fundamental concepts associated with RFC 2547bis VPNs.
- Provider (P) routers still will not be aware of the VPNs. They will continue to forward packets over pre-established LSPs.
- Customer Edge (CE) routers will operate without any knowledge of the existence of MPLS VPNs.
- The PE routers do not participate in the routing algorithms of the end-users, and there are no requirements for the construction of VPN routing and forwarding tables (VRFs).

Martini VPNs (Point-Point Connectivity)

- The Martini drafts introduce the concept of Virtual Circuits (VCs). An LSP acts as a tunnel carrying multiple VCs.

- VCs are uni-directional, for bi-directional communication, a pair of VCs – one in each direction – is needed.

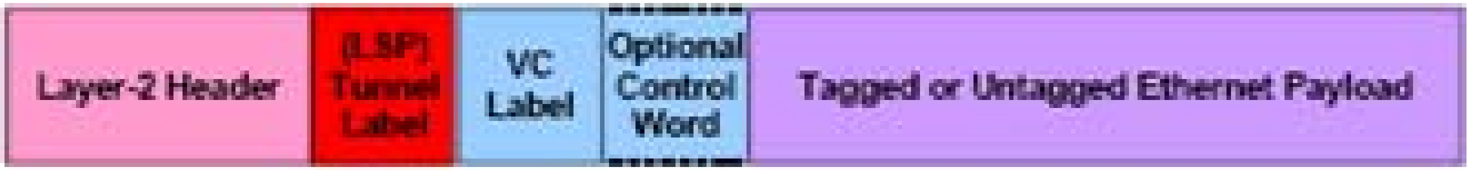

Layer-2 Header	(LSP) Tunnel Label	VC Label	Optional Control Word	Tagged or Untagged Ethernet Payload

L2 VPN Routing Information

- Tunnel LSPs between the PE routers could be created using any protocol like RSVP/TE or LDP.

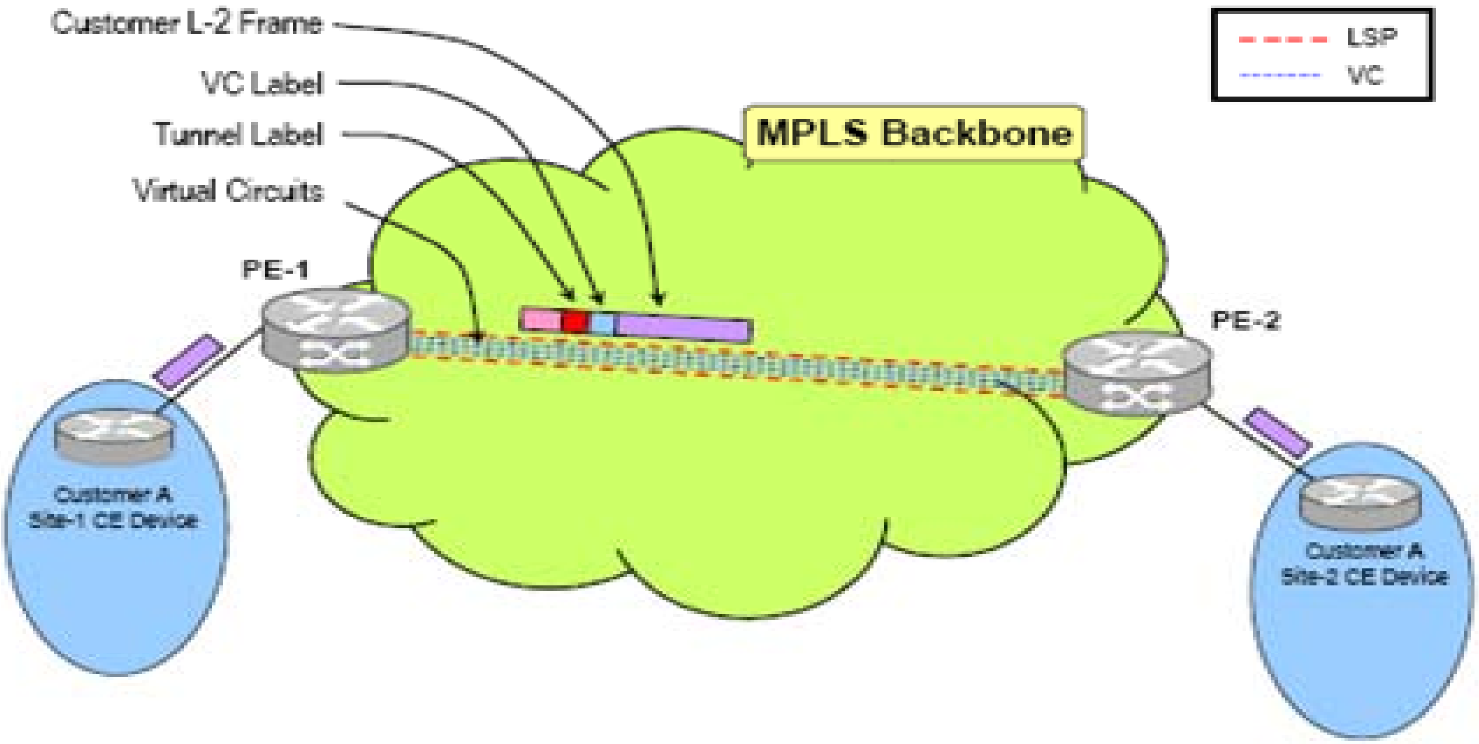

L2 VPN Routing Information

- PE routers exchange the VC labels via LDP. Once the session is established, VC ID data which includes the VC ID, the Group ID, VC Type, the VC Interface Parameters and a Control Word notification can be exchanged.

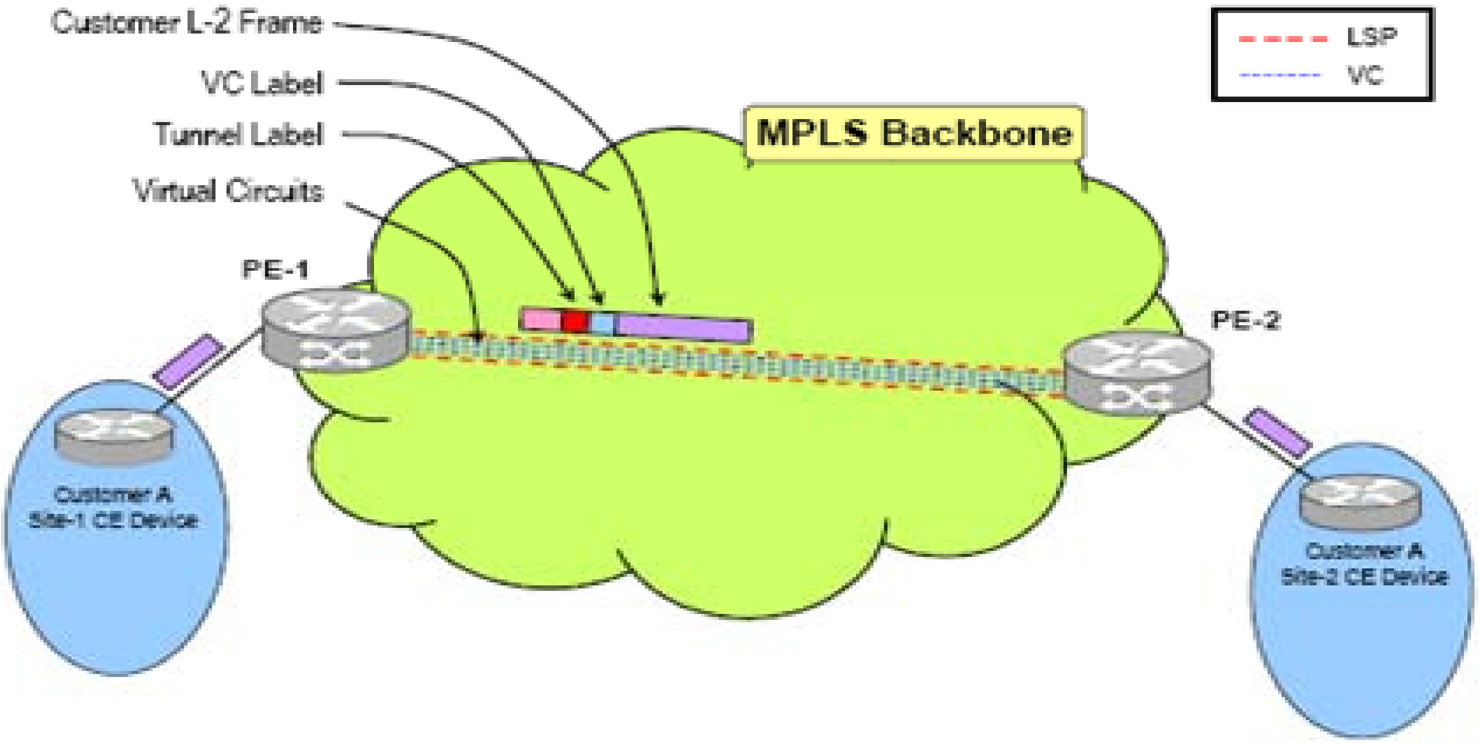

L2 VPN Data Traffic

- The PE router encapsulates the subscriber layer-2 frame and attaches two labels; the top (tunnel label) identifies the destination of the remote PE router.

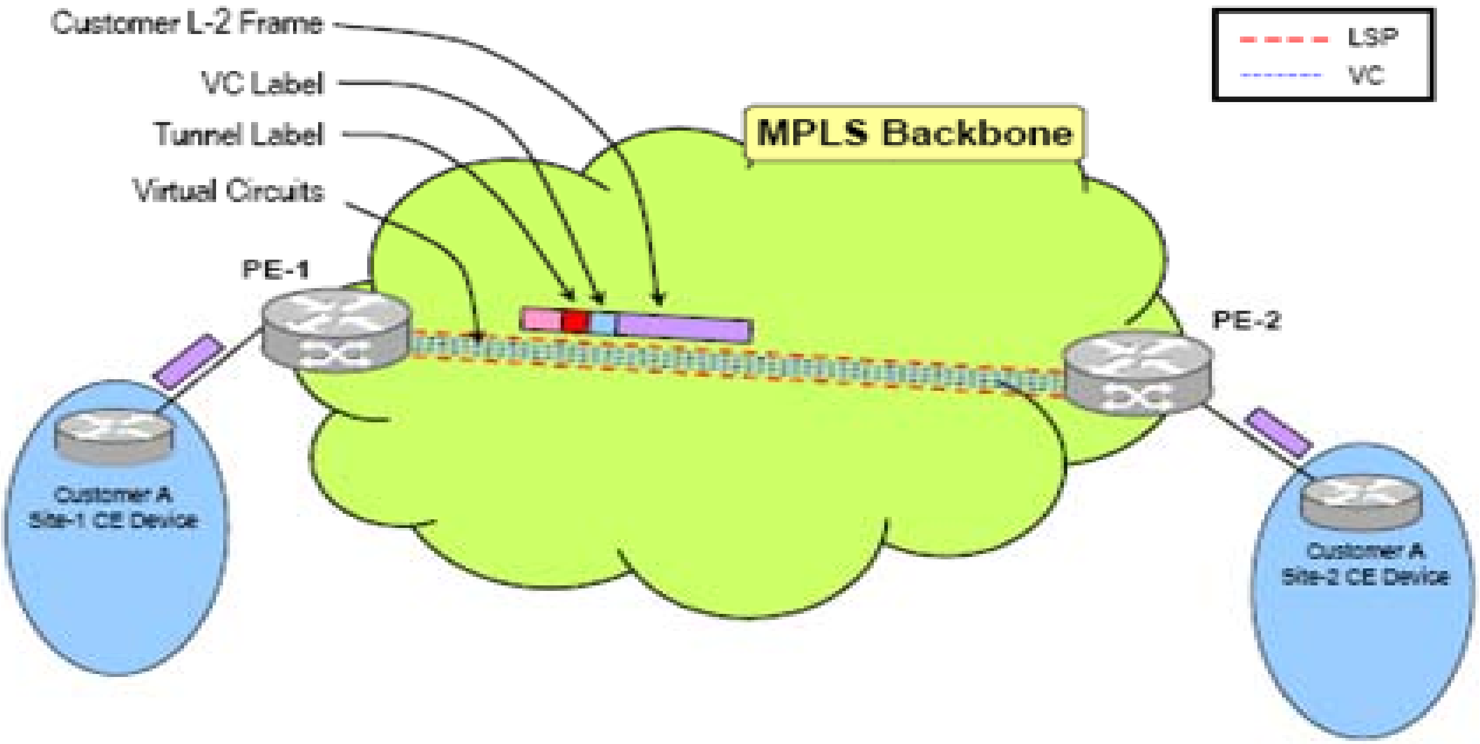

L2 VPN Data Traffic

- The receiving PE router pops the tunnel label, uses the bottom (or inner) label to deliver the packet to the correct end-user (CE router) with the appropriate Layer 2 encapsulation based on the VC label.

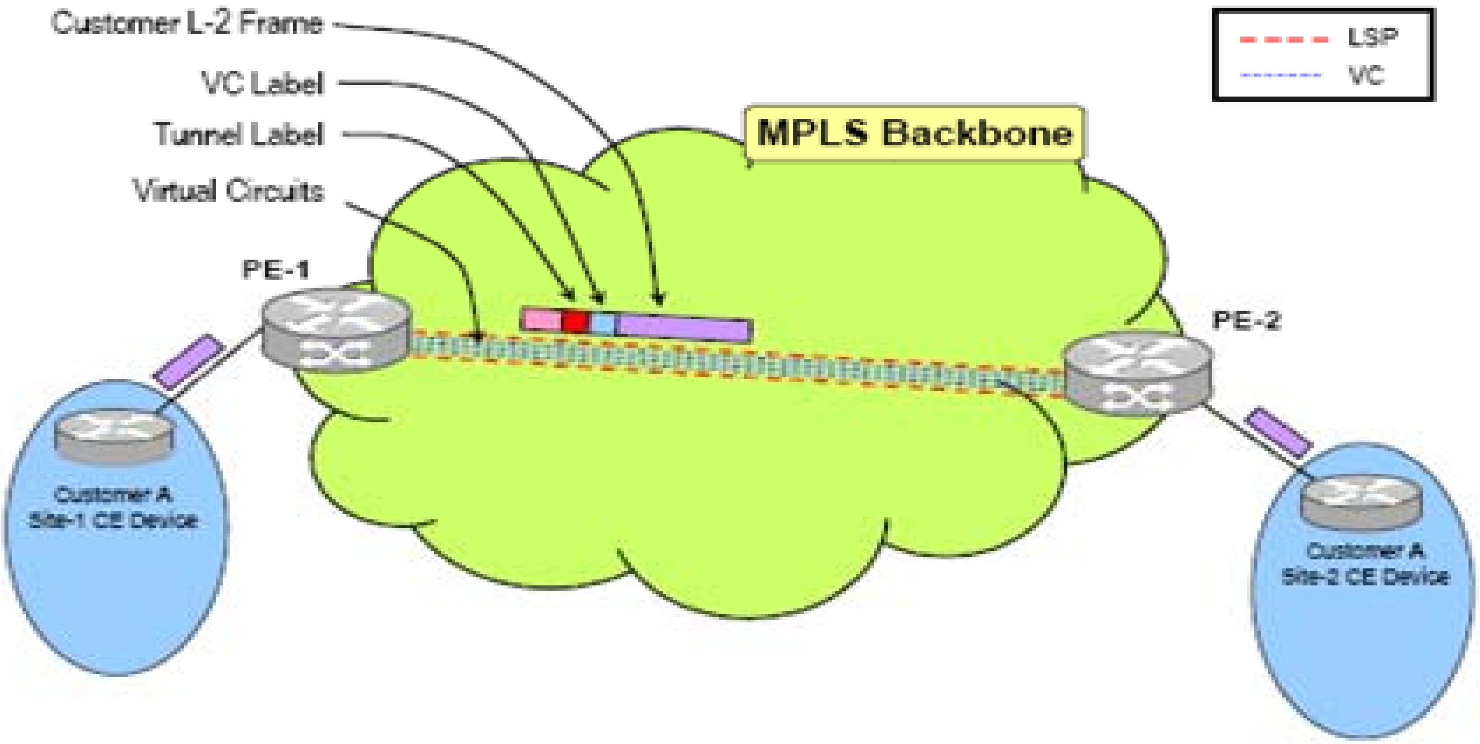

VC Types

- Martini Virtual Circuit Encapsulation Types:
  Frame Relay – Type 01
  ATM AAL5 VCC – Type 02
  ATM Transparent Cell Transport – Type 03
  Ethernet VLAN – Type 04
  Ethernet – Type 05
  HDLC – Type 06
  PPP – Type 07
  CEM – Type 08
  ATM VCC Cell Transport – Type 09
  ATM VPC Cell Transport – Type 10 or Hex. “0A”

Virtual Private LAN Services (VPLS)

- The Kompella draft specification creates a new VC type specifically for Ethernet VPLS frames. This is type eleven (hex B).
- Customer frames are switched based on their destination MAC address.
- VPN is established by creating a full mesh of VCs between the PEs facing the sites that make the VPN.

Virtual Private LAN Services (VPLS)

- PE routers perform source MAC address learning just like a normal transparent switch, except that they perform it on frames received over the VCs.

- A PE router maintains a separate layer-2 forwarding table, called Virtual Forwarding Instance (VFI), for each VPN that it carries.

Virtual Private LAN Services (VPLS)

- PE router does not learn all the MAC addresses in all the VPNs carried by the provider network. A PE router learns MAC addresses related only to the VPNs that it carries. P routers do not learn any MAC addresses, they just perform label switching.

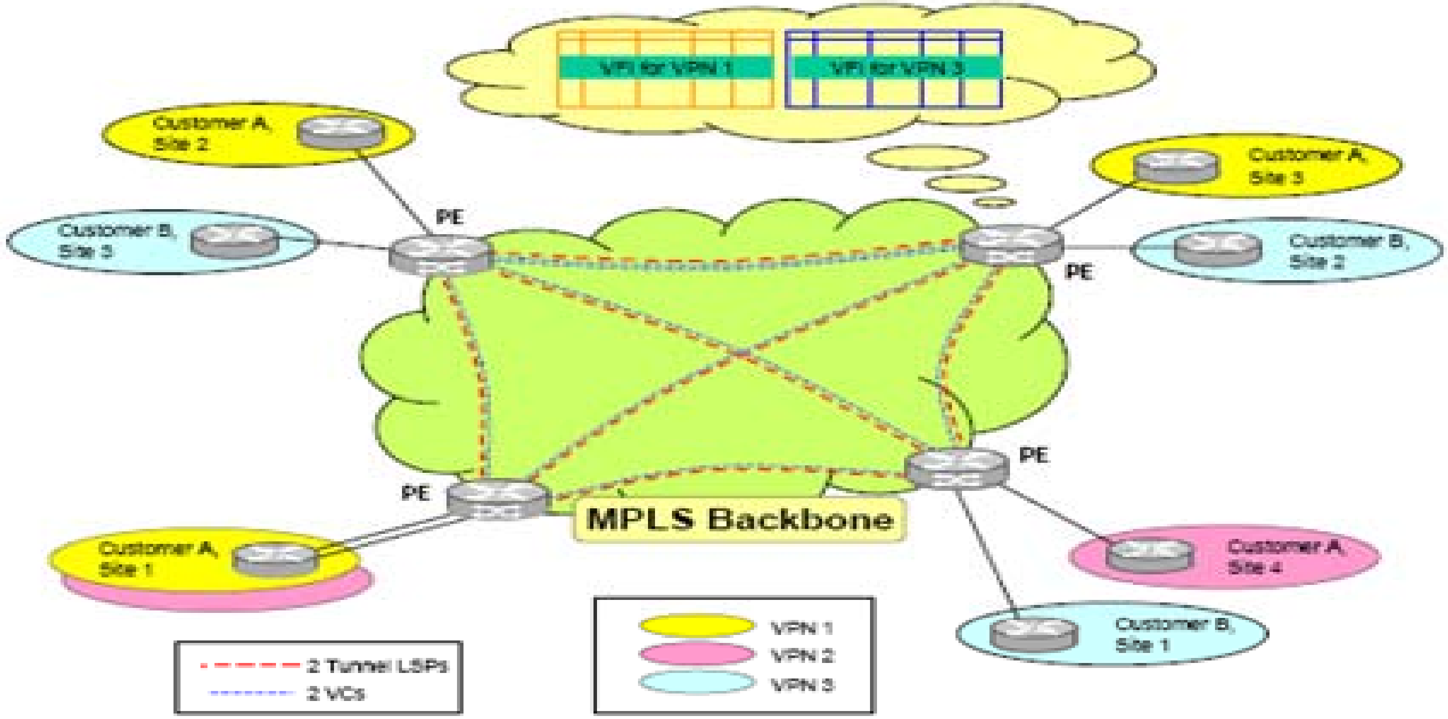

Thank You