MPLS in communication networks. (Part 2)

«Computer networks and telecommunications» (Additional chapters).

MPLS VPN Technology

Introducing VPNs

Traditional Router-Based Networks

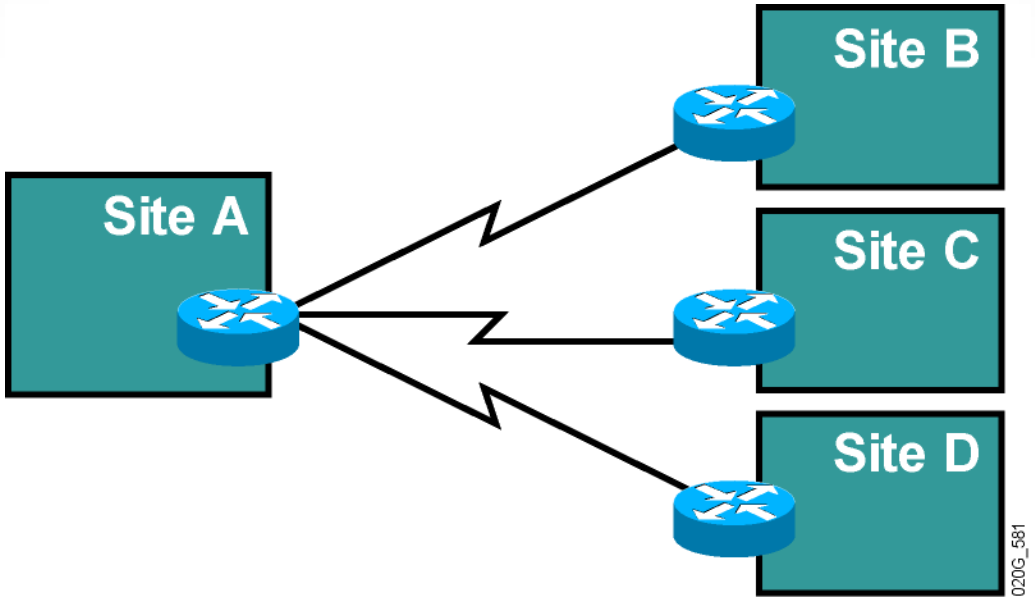

- Traditional router-based networks connect customer sites through routers connected via dedicated point-to-point links.

Virtual Private Networks

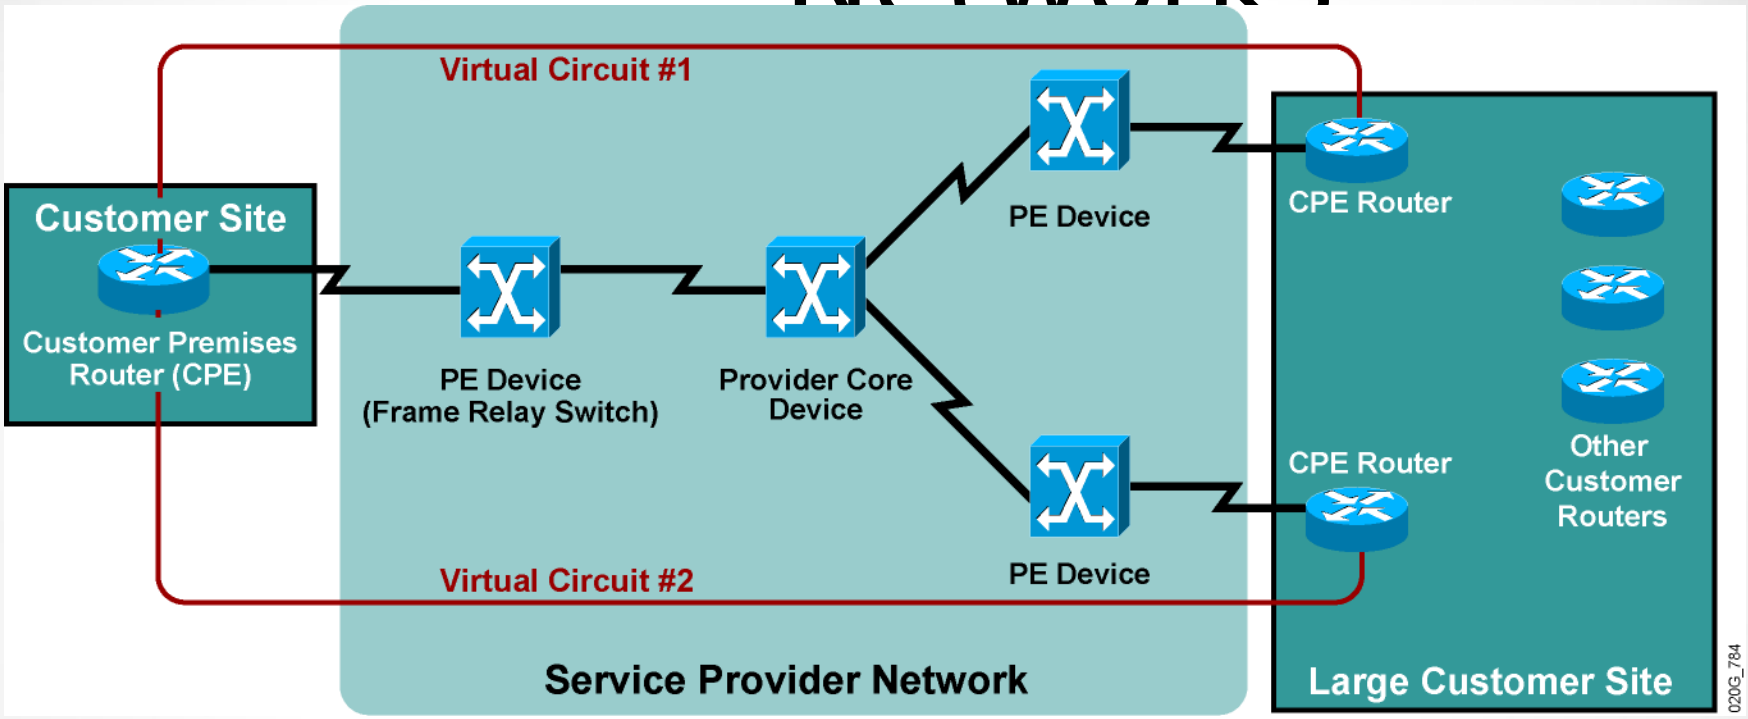

- VPNs replace dedicated point-to-point links with emulated point-to-point links sharing common infrastructure.
- Customers use VPNs primarily to reduce their operational costs.

VPN Terminology

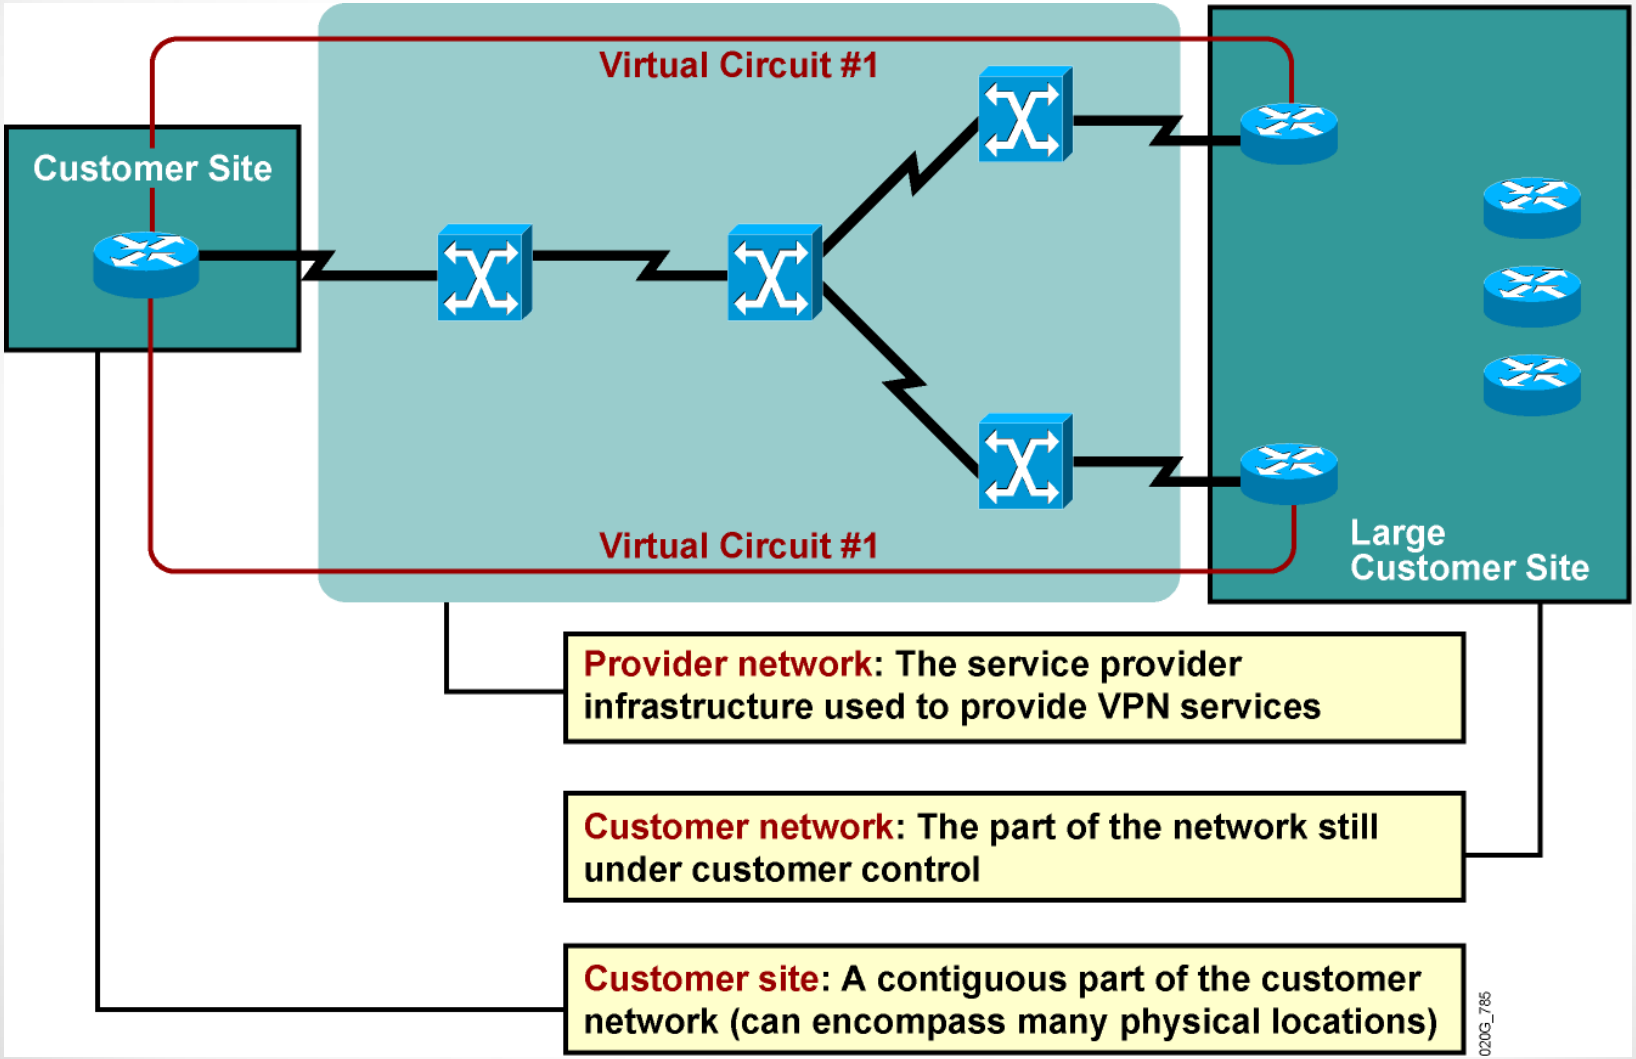

VPN Terminology (Cont.)

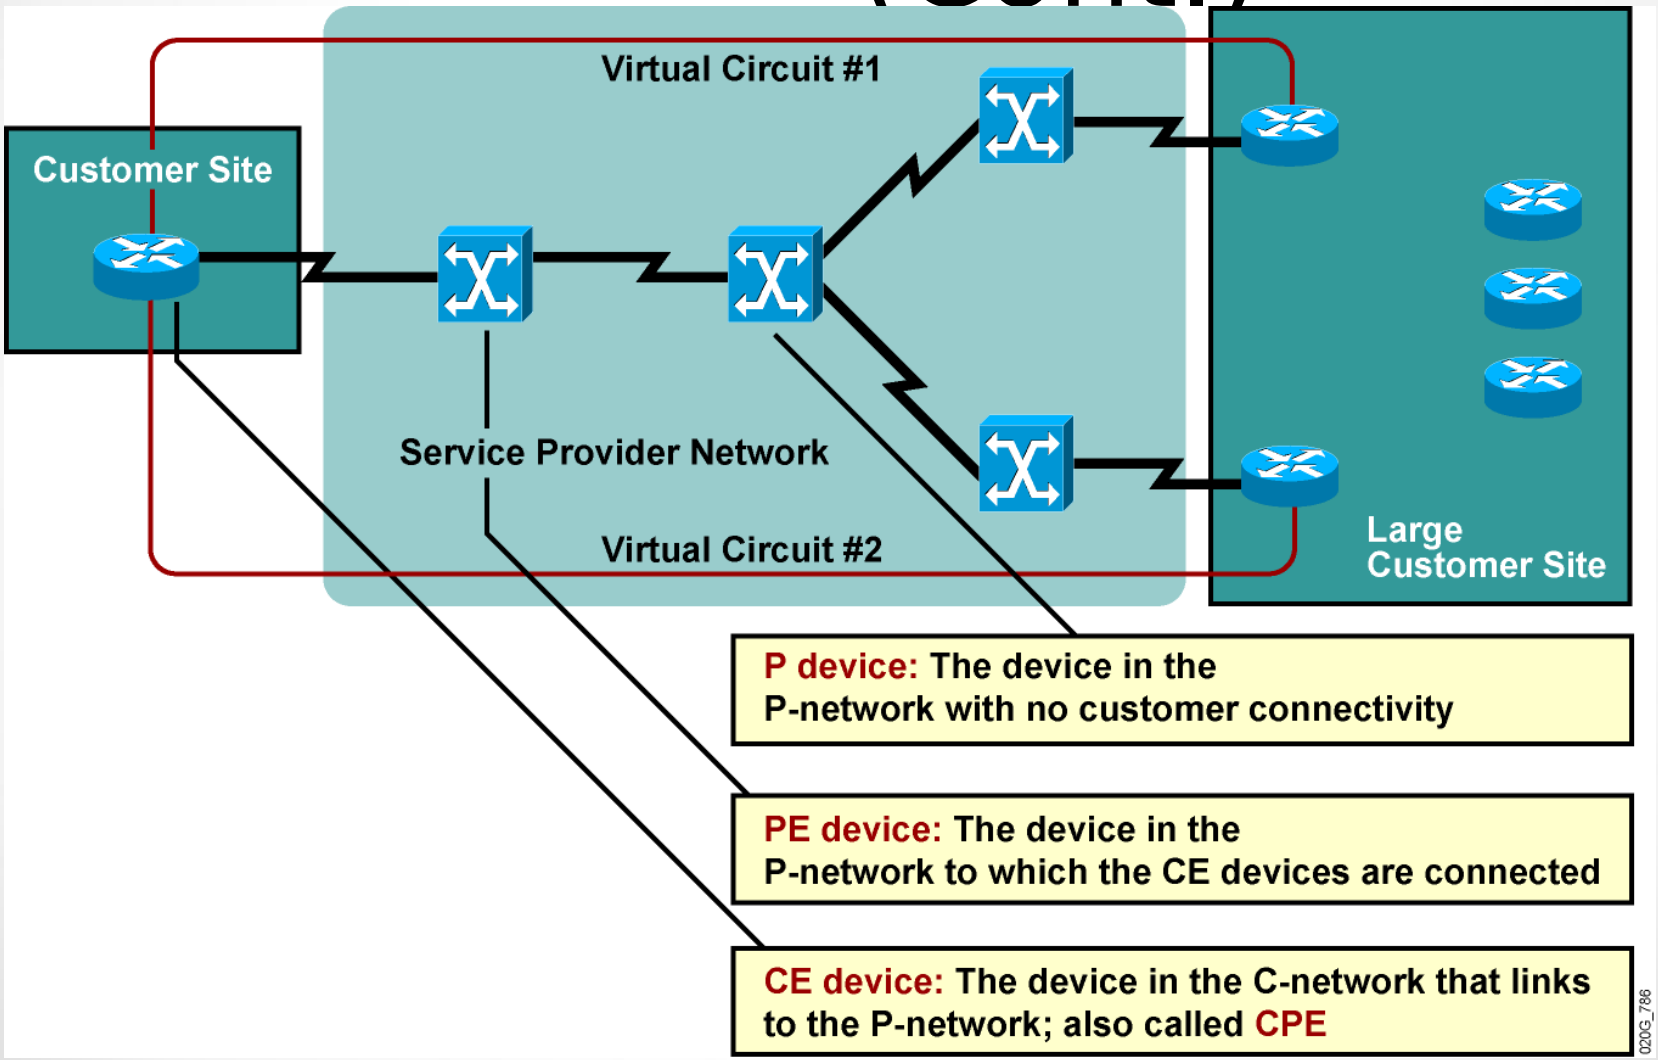

VPN Implementation Models

- VPN services can be offered based on two major models:
  - Overlay VPNs, in which the service provider provides virtual point-to-point links between customer sites
  - Peer-to-peer VPNs, in which the service provider participates in the customer routing

Overlay VPNs: Hub-and-Spoke Topology

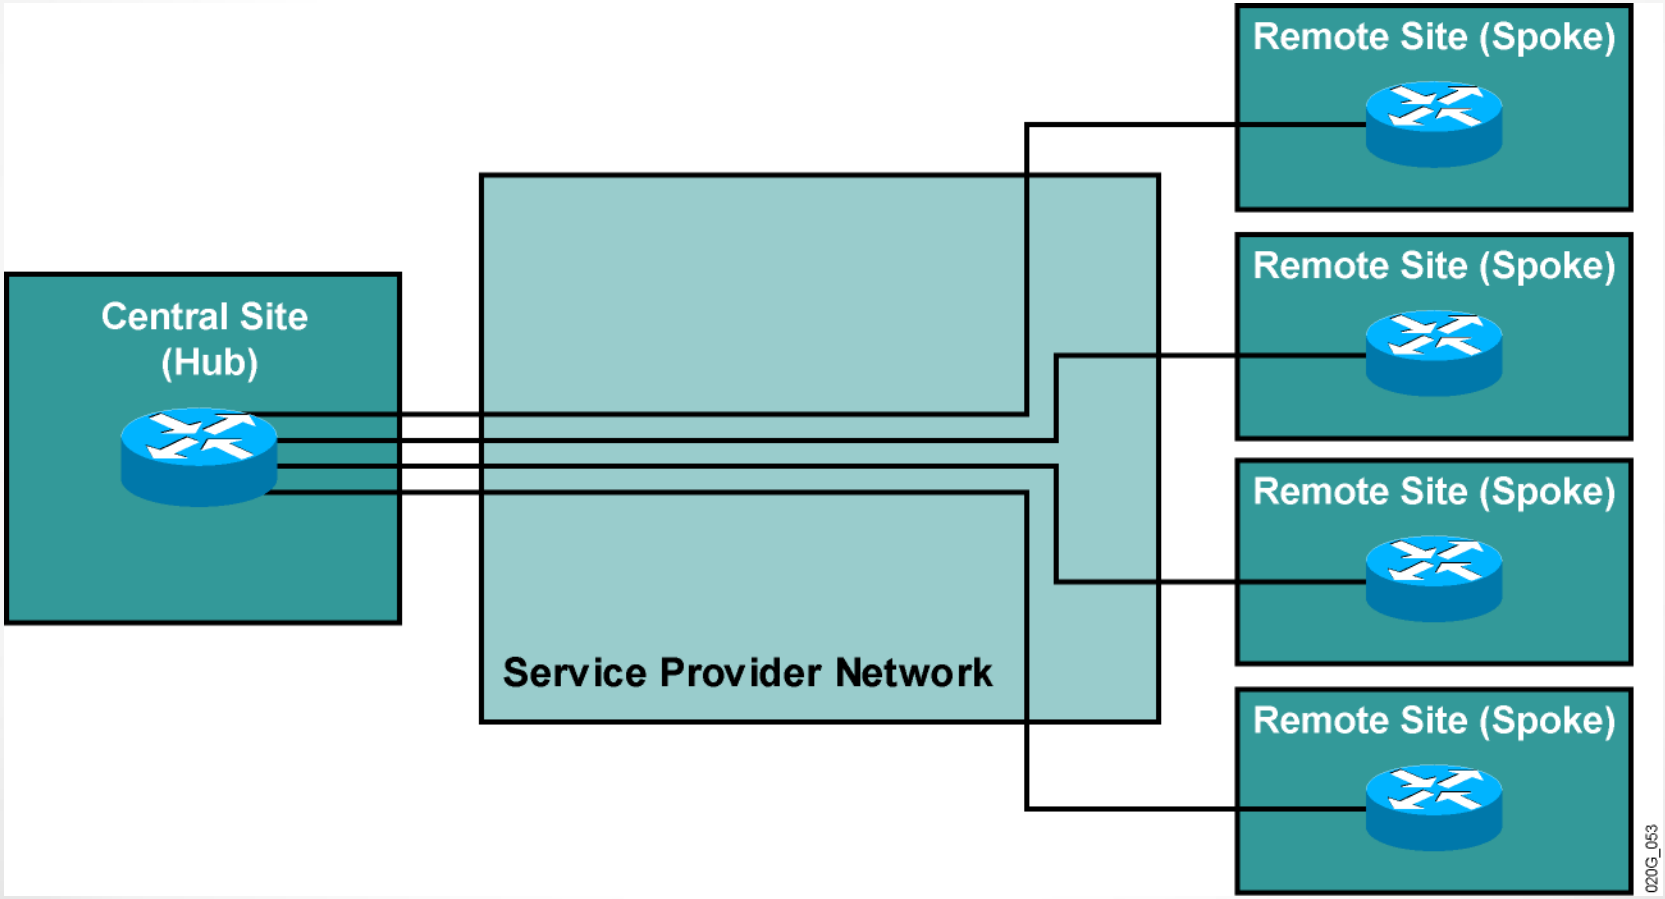

Overlay VPNs: Redundant Hub-and-Spoke Topology

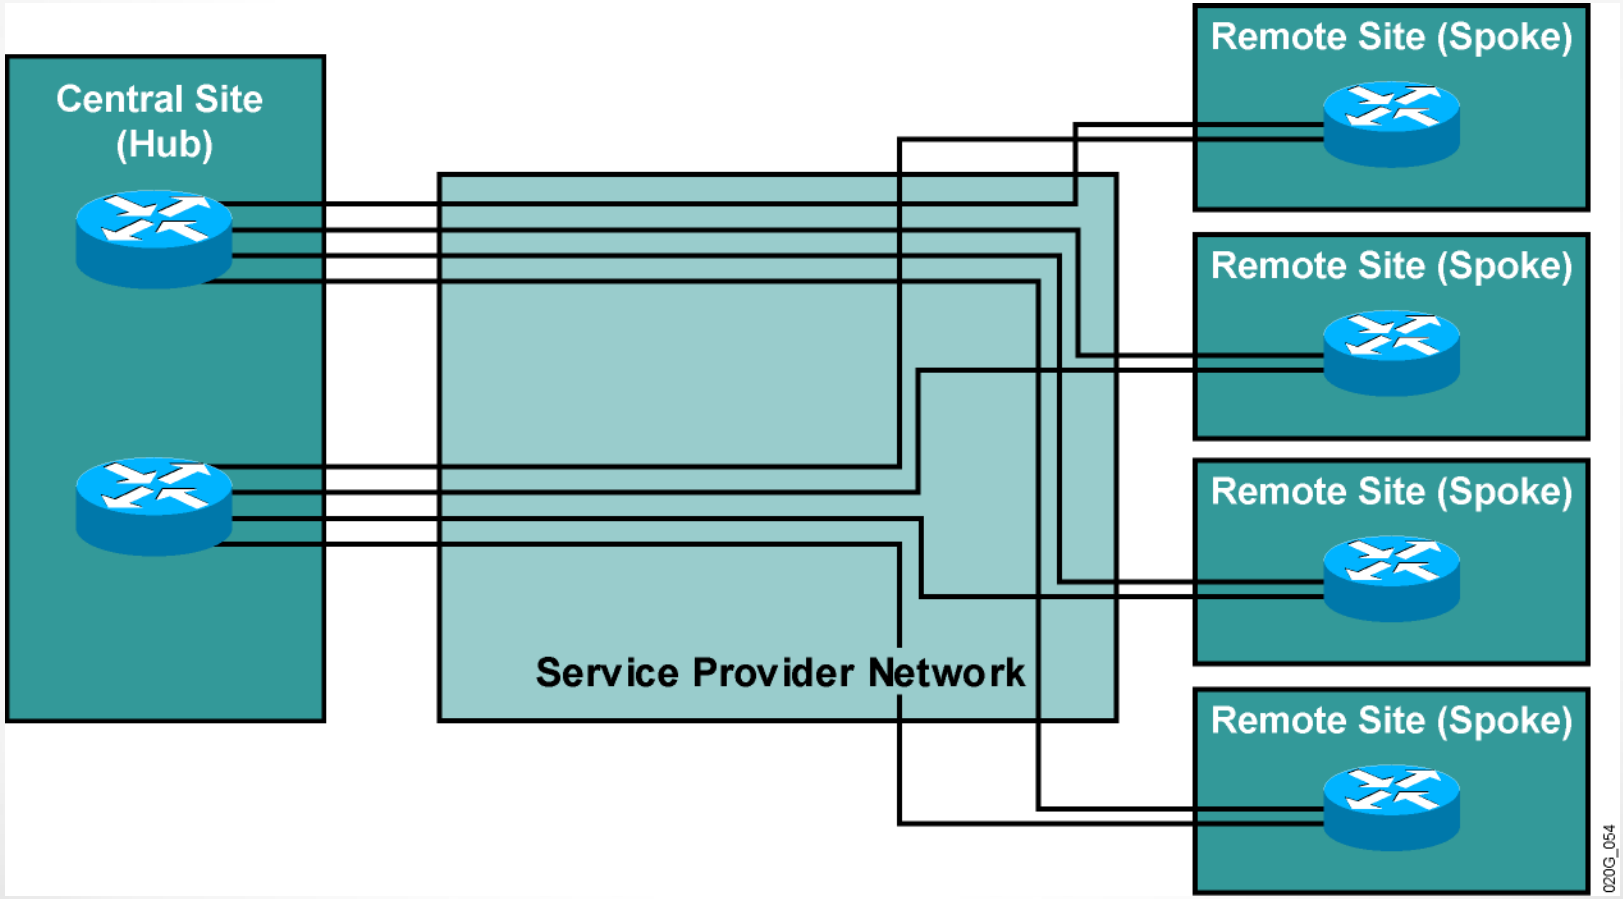

Overlay VPNs: Layer 2 Implementation

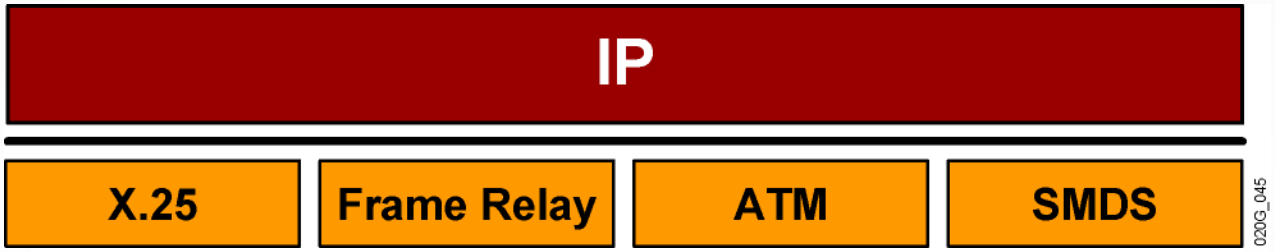

- This is the traditional switched WAN solution:
  - The service provider establishes Layer 2 virtual circuits between customer sites.
  - The customer is responsible for all higher layers.

Overlay VPNs: IP Tunneling

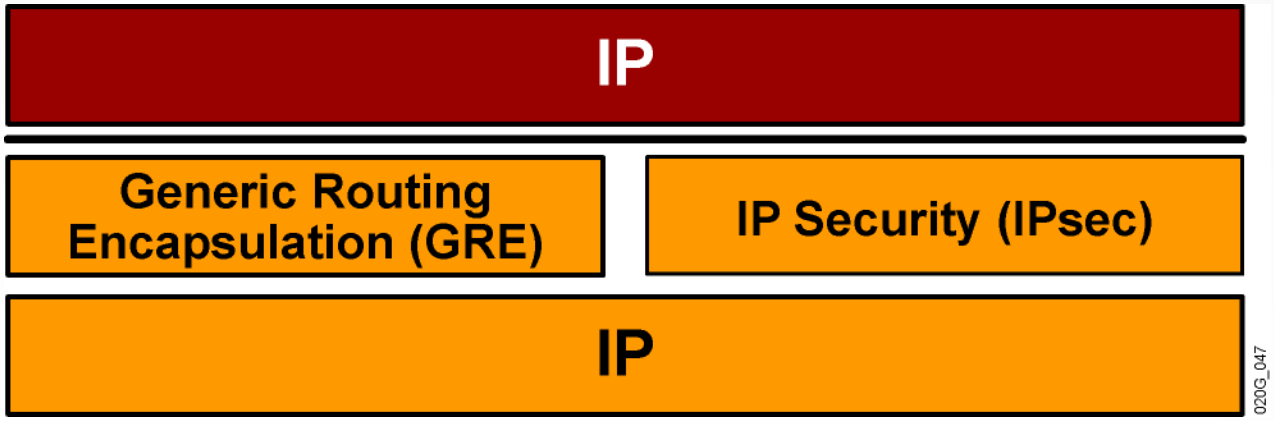

- VPN is implemented with IP-over-IP tunnels:
  - Tunnels are established with GRE or IPsec.
  - GRE is simpler (and quicker); IPsec provides authentication and security.

Overlay VPNs: Layer 3 Routing

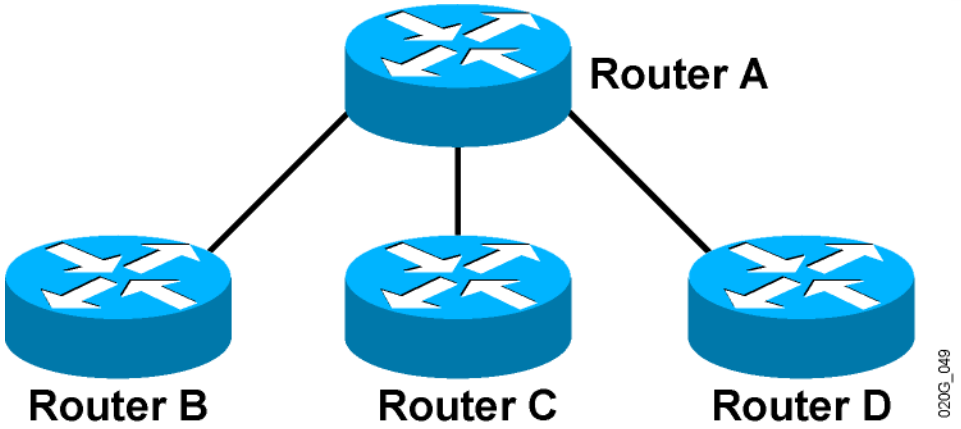

- The service provider infrastructure appears as point-to-point links to customer routes.
- Routing protocols run directly between customer routers.
- The service provider does not see customer routes and is responsible only for providing point-to-point transport of customer data.

Peer-to-Peer VPNs: Implementation Techniques

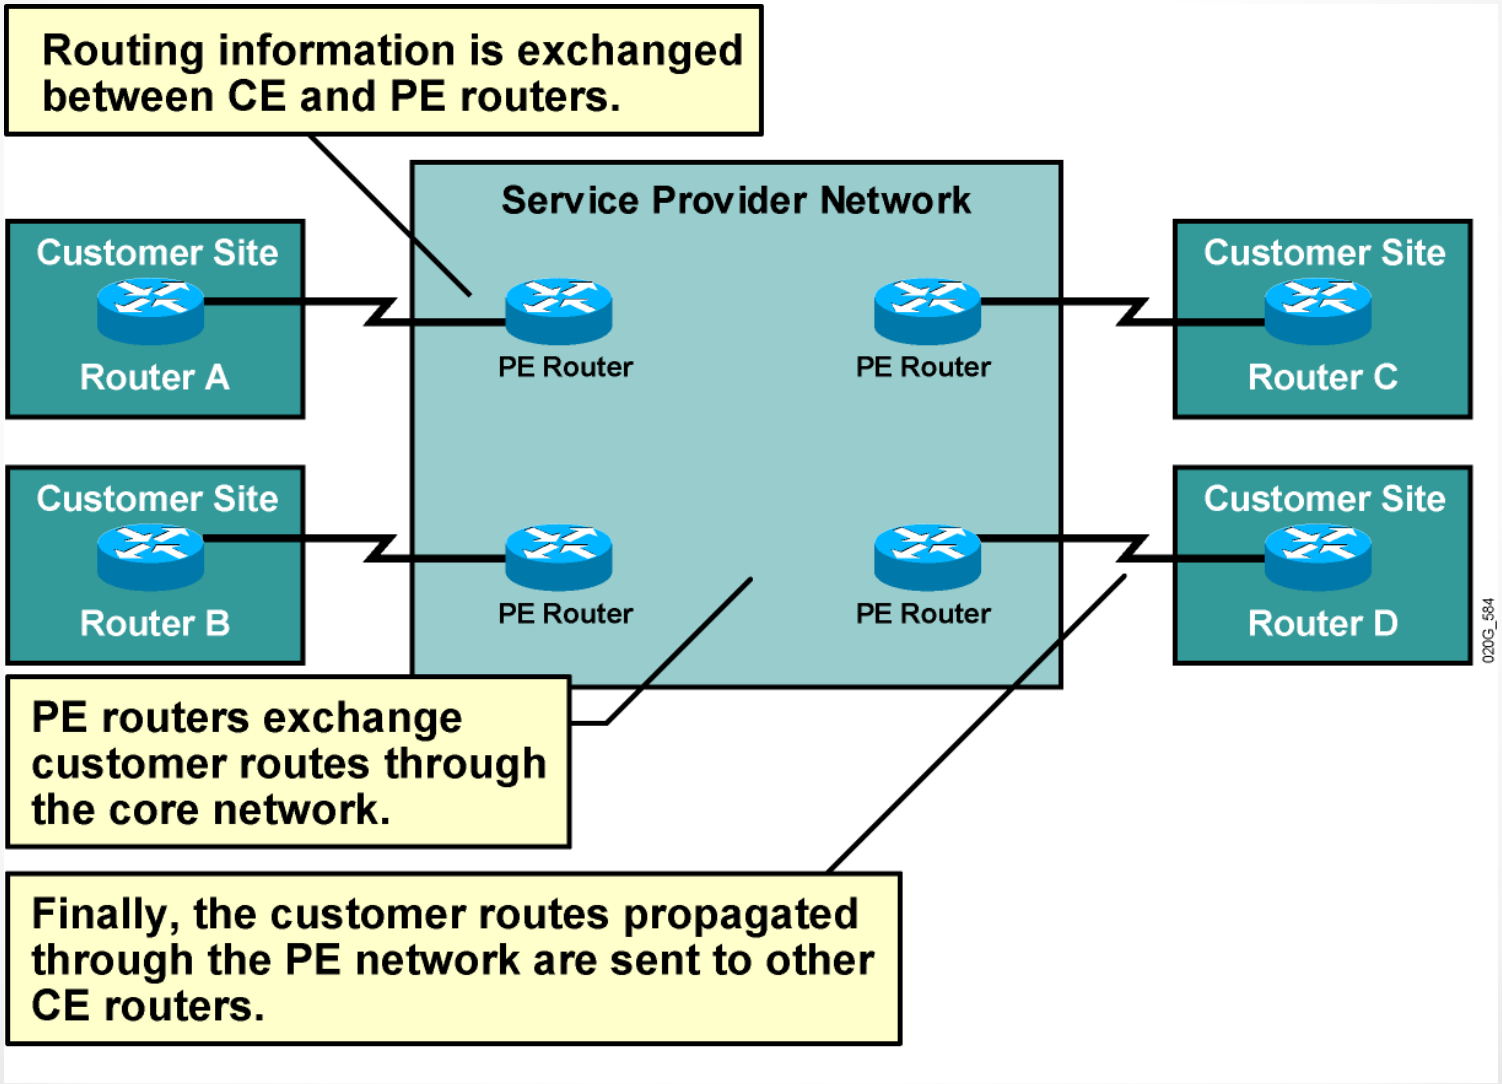

Peer-to-Peer VPNs: Packet Filters

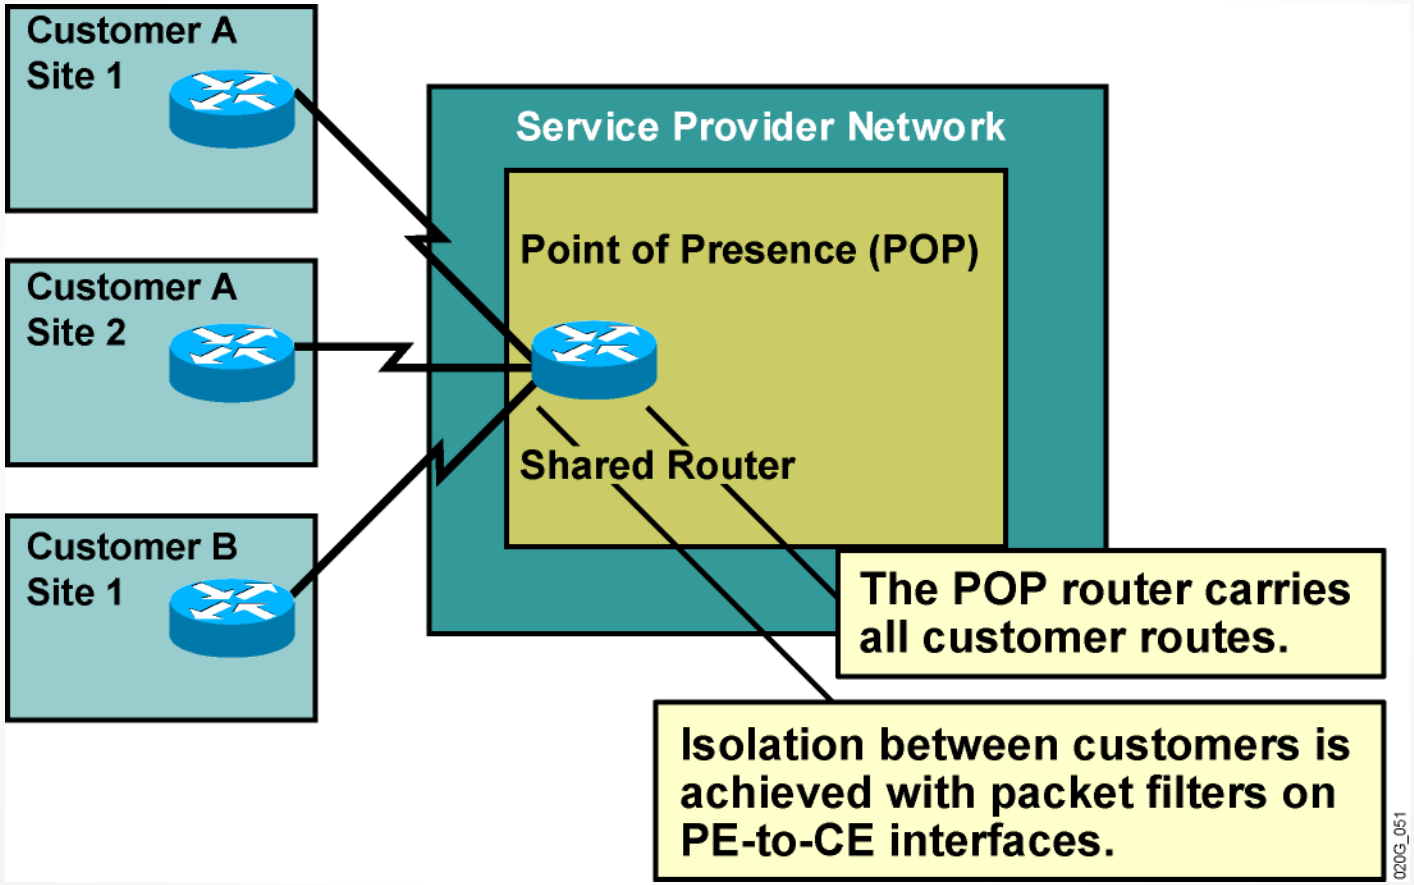

Peer-to-Peer VPNs: Controlled Route Distribution

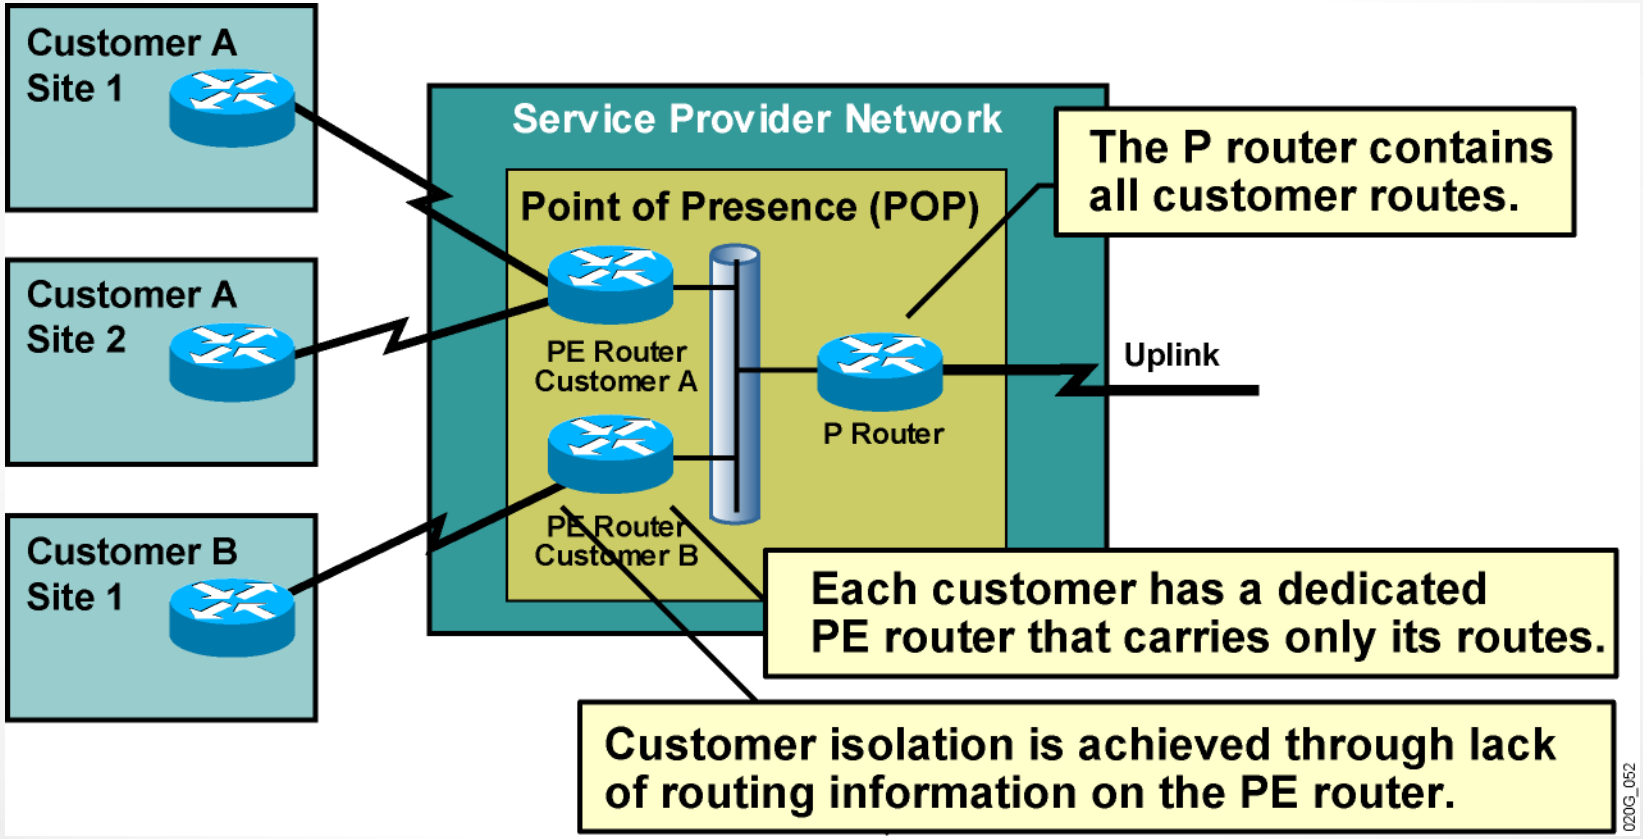

Benefits of VPN Implementations

- Overlay VPN:
  - Well-known and easy to implement
  - Service provider does not participate in customer routing
  - Customer network and service provider network are well-isolated
- Peer-to-peer VPN:
  - Guarantees optimum routing between customer sites
  - Easier to provision an additional VPN
  - Only sites provisioned, not links between them

Drawbacks of VPN Implementations

- Overlay VPN:
  - Implementing optimum routing requires a full mesh of virtual circuits.
  - Virtual circuits have to be provisioned manually.
  - Bandwidth must be provisioned on a site-to-site basis.
  - Overlay VPNs always incur encapsulation overhead.
- Peer-to-peer VPN:
  - The service provider participates in customer routing.
  - The service provider becomes responsible for customer convergence.
  - PE routers carry all routes from all customers.
  - The service provider needs detailed IP routing knowledge.

VPN Connectivity Category

- VPNs can also be categorized according to the connectivity required between sites:
  - Simple VPN: Every site can communicate with every other site.
  - Overlapping VPNs: Some sites participate in more than one simple VPN.
  - Central services VPN: All sites can communicate with central servers but not with each other.
  - Managed network: A dedicated VPN is established to manage CE routers.

MPLS VPN Technology

Introducing MPLS VPN Architecture

Drawbacks of Traditional Peer-to-Peer VPNs

- Shared PE router:
  - All customers share the same (provider-assigned or public) address space.
  - High maintenance costs are associated with packet filters.
  - Performance is lower—each packet has to pass a packet filter.
- Dedicated PE router:
  - All customers share the same address space.
  - Each customer requires a dedicated router at each POP.

MPLS VPN Architecture

- An MPLS VPN combines the best features of an overlay VPN and a peer-to-peer VPN:
  - PE routers participate in customer routing, guaranteeing optimum routing between sites and easy provisioning.
  - PE routers carry a separate set of routes for each customer (similar to the dedicated PE router approach).
  - Customers can use overlapping addresses.

MPLS VPN Architecture: Terminology

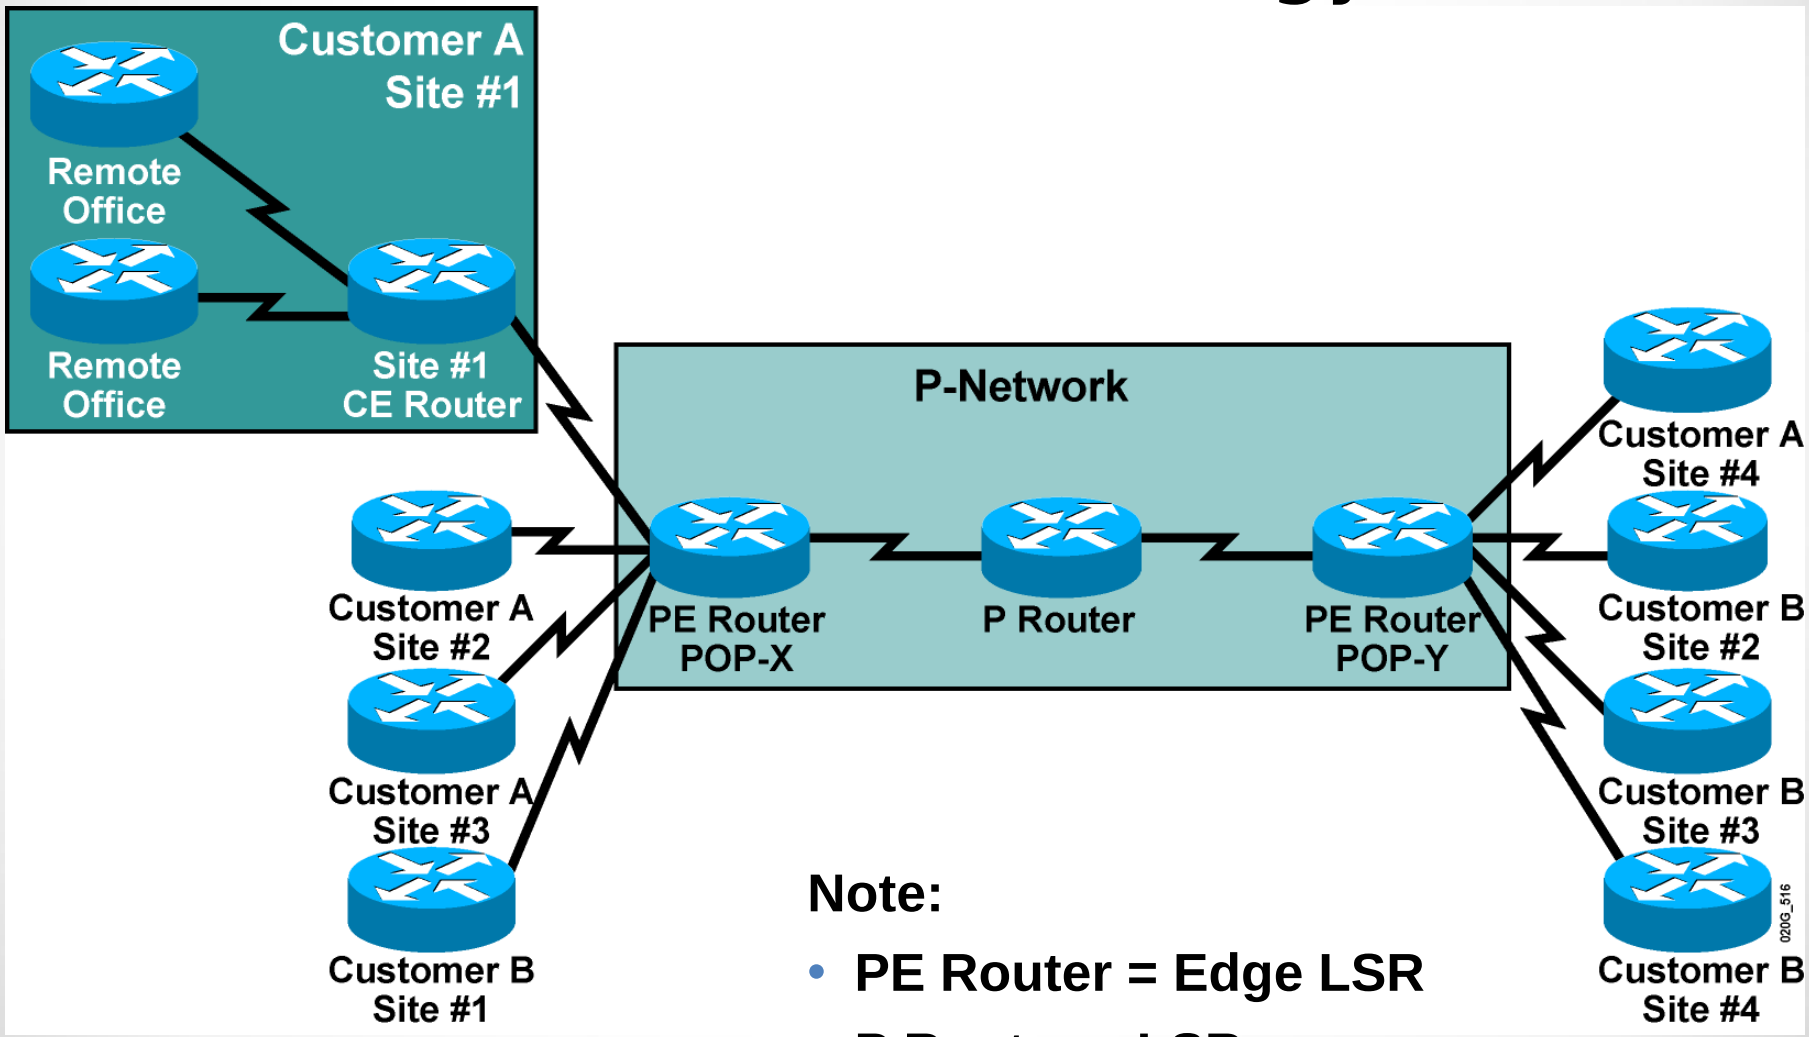

Note:

- **PE Router = Edge LSR**
- **P Router = LSR**

PE Router Architecture

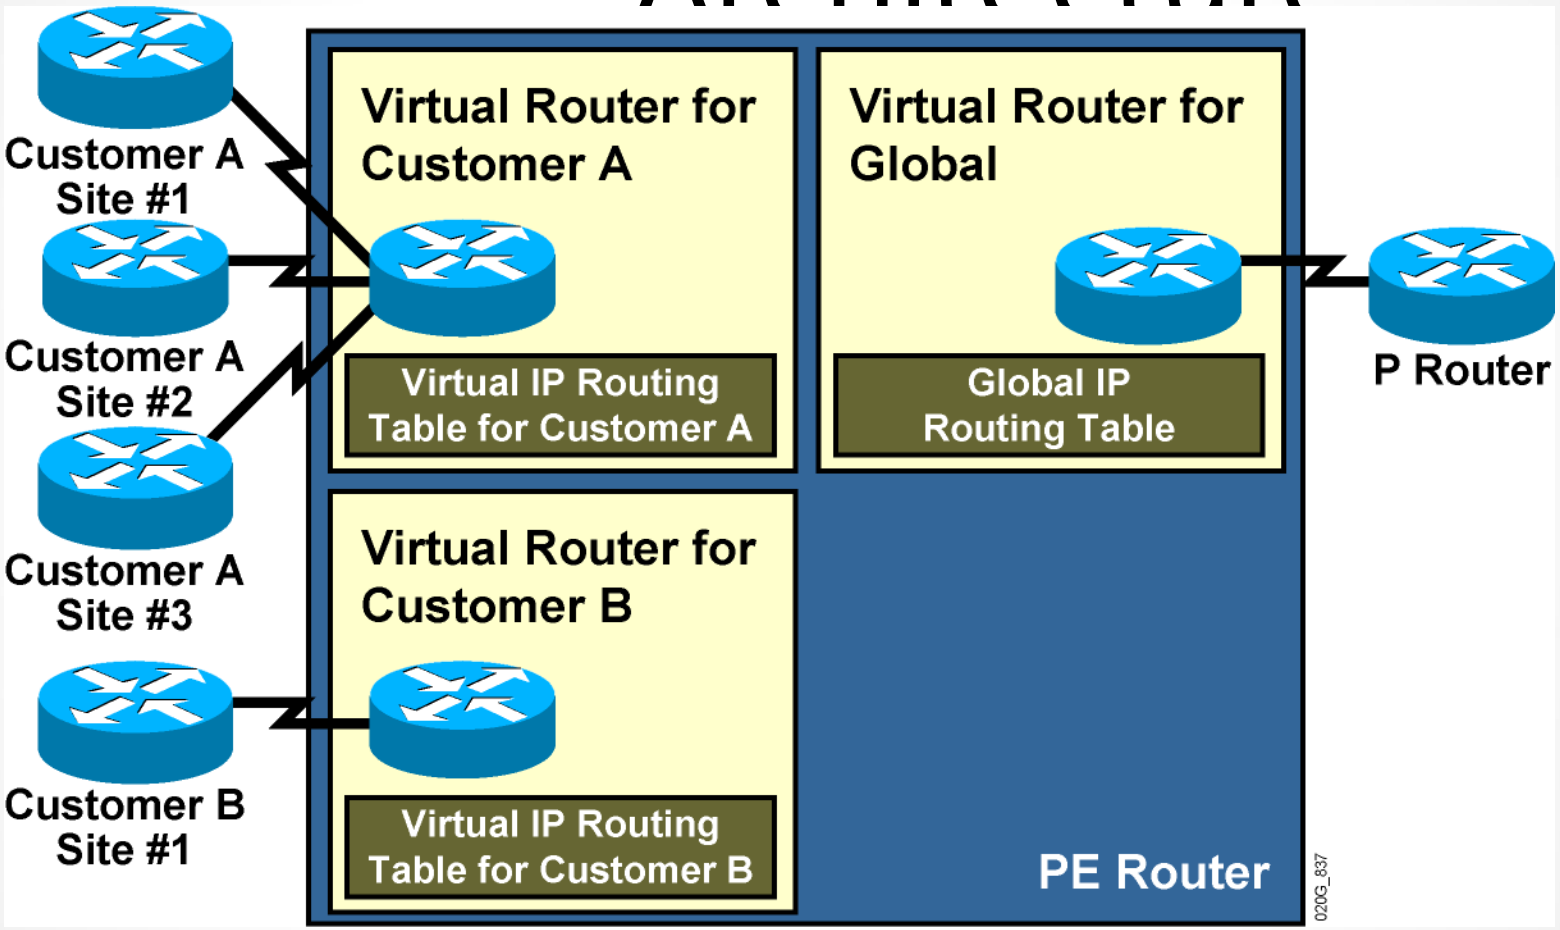

- **PE router in an MPLS VPN uses virtual routing tables to implement the functionality of customer dedicated PE routers.**

Propagation of Routing Information Across the P-Network

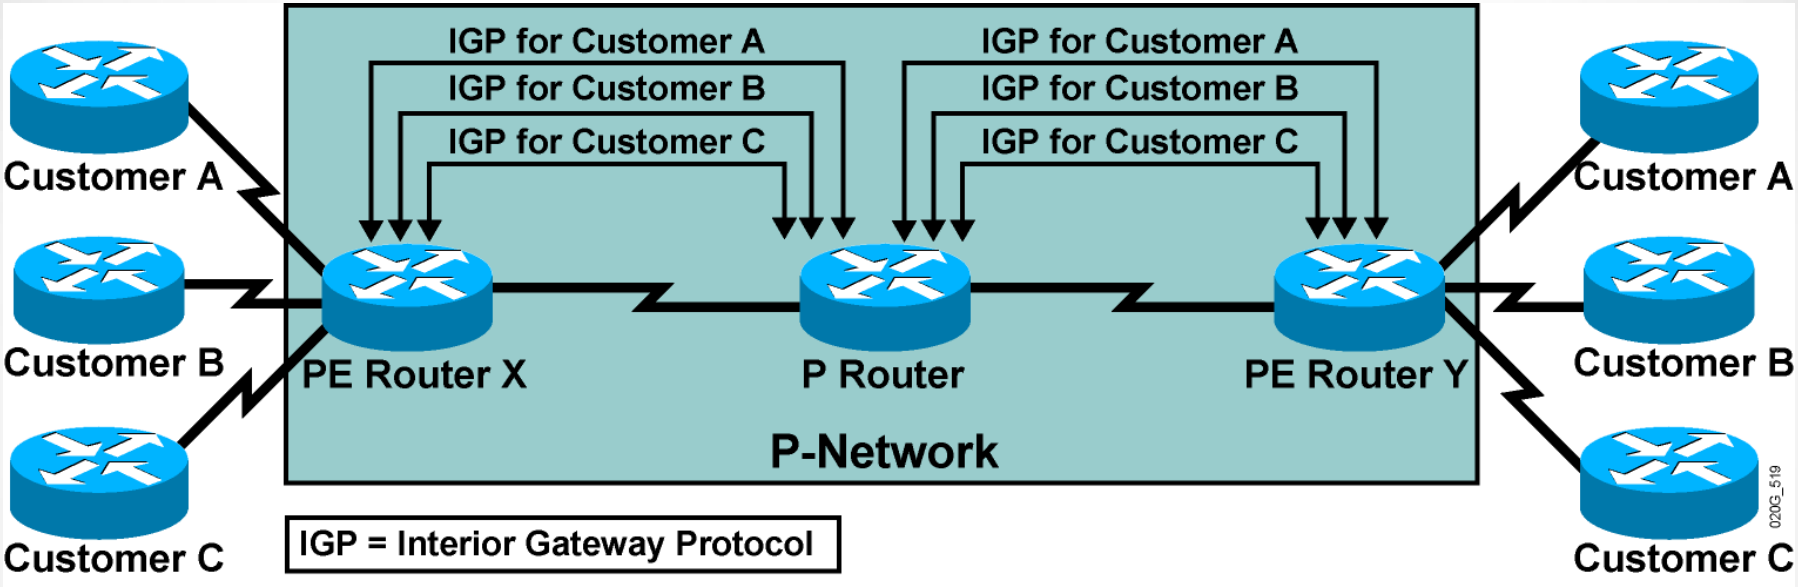

Question: How will PE routers exchange customer routing information?

Option #1: Run a dedicated IGP for each customer across the P-network.

This is the wrong answer for these reasons:

- The solution does not scale.
- P routers carry all customer routes.

Propagation of Routing Information Across the P-Network (Cont.)

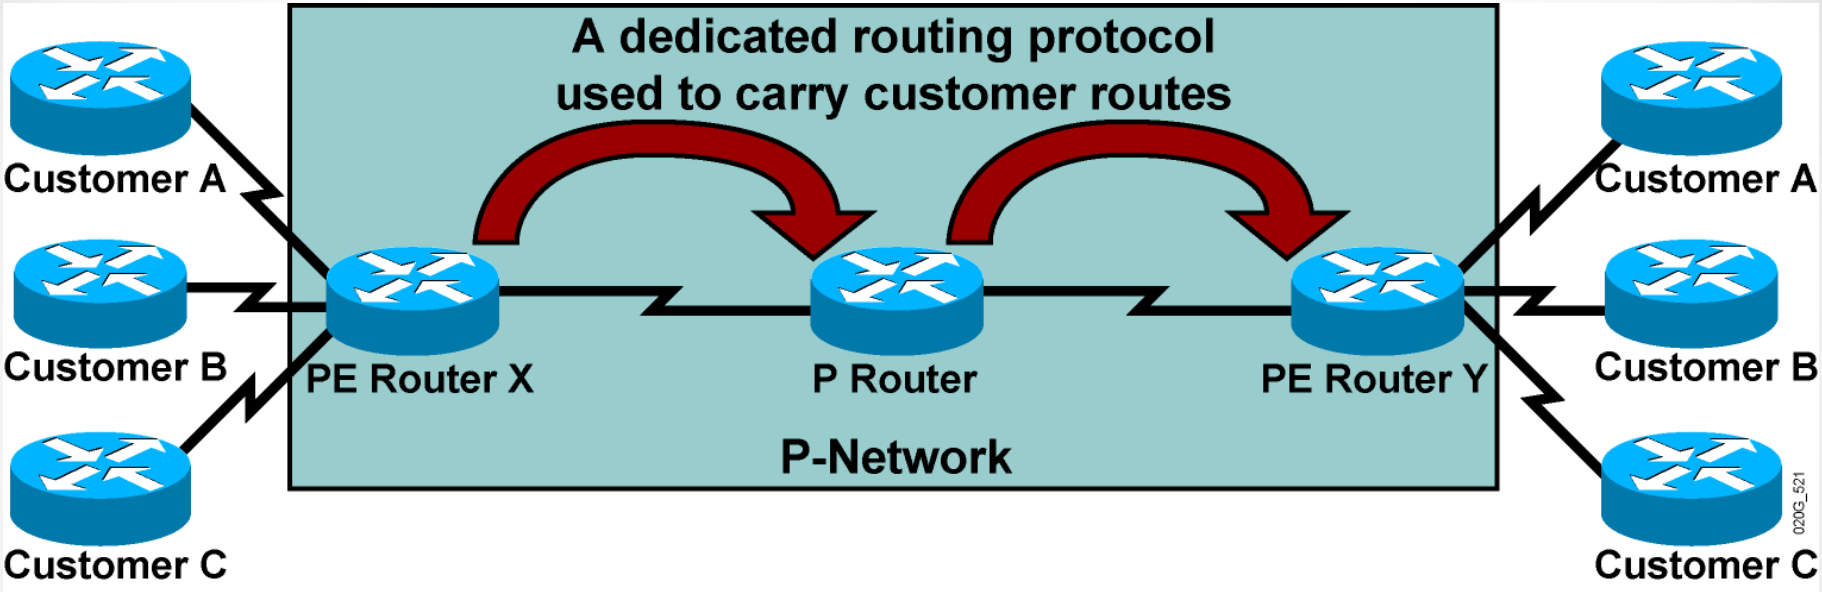

Question: How will PE routers exchange customer routing information?

Option #2: Run a single routing protocol that will carry all customer routes inside the provider backbone.

Better answer, but still not good enough:

- P routers carry all customer routes.

Propagation of Routing Information Across the P-Network (Cont.)

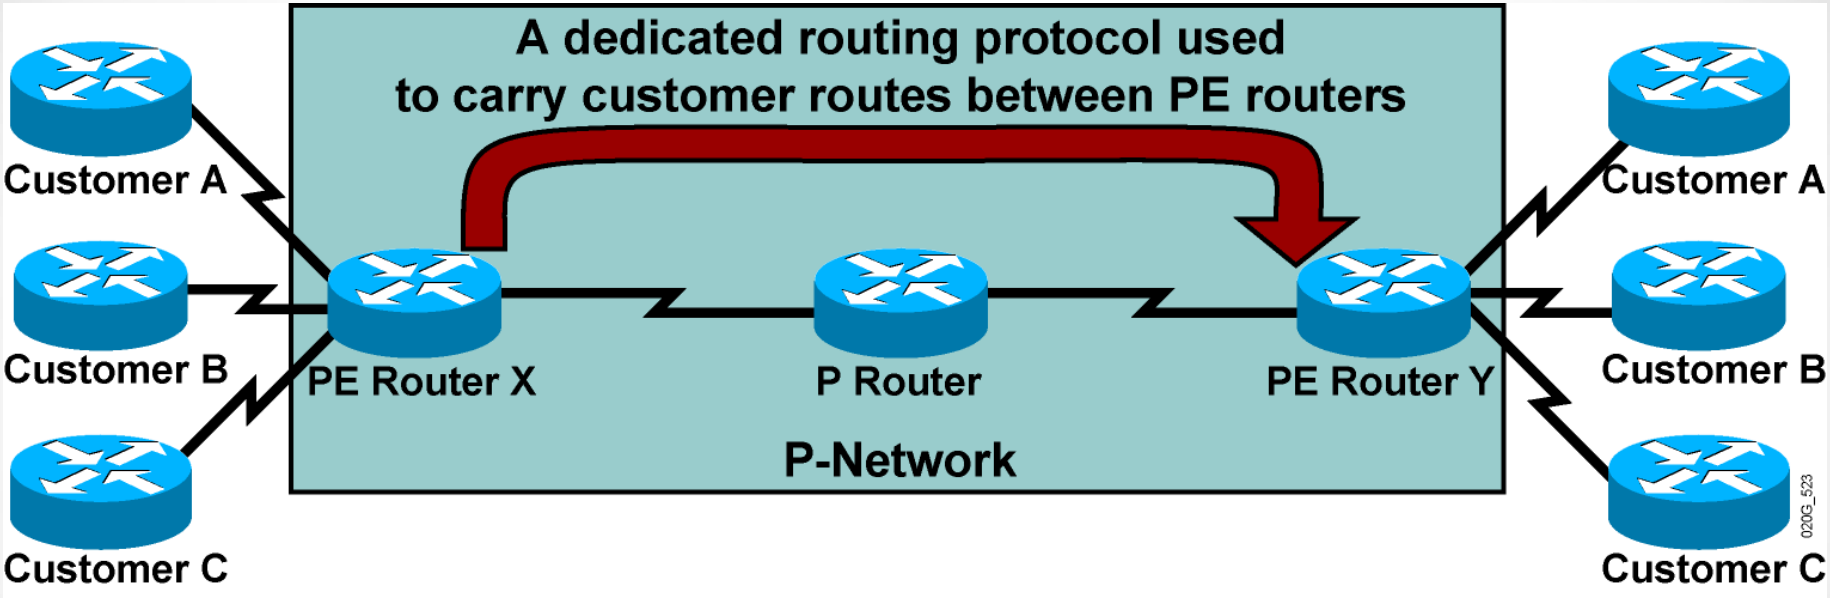

Question: How will PE routers exchange customer routing information?

Option #3: Run a single routing protocol that will carry all customer routes between PE routers. Use MPLS labels to exchange packets between PE routers.

The best answer:

- **P routers do not carry customer routes; the solution is scalable.**

Propagation of Routing Information Across the P-Network (Cont.)

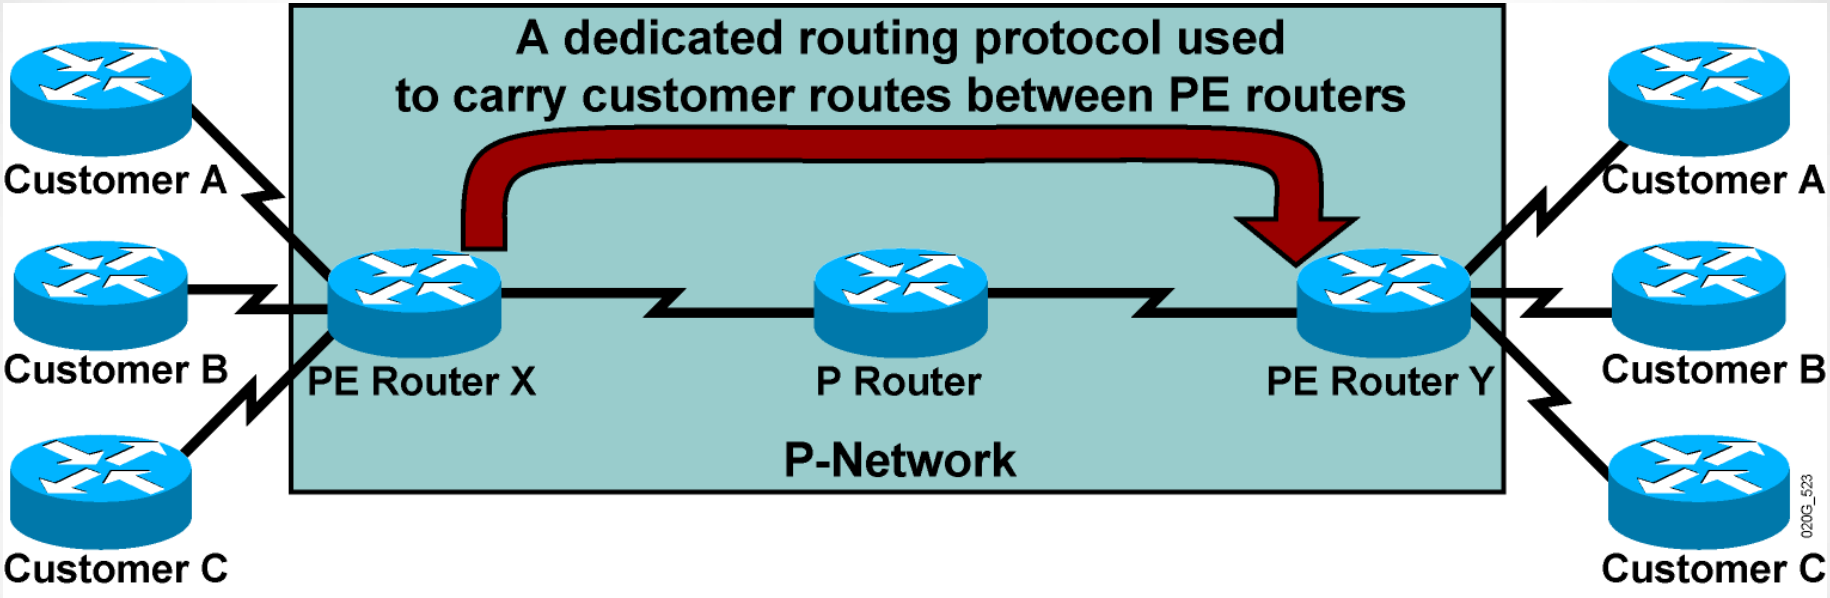

Question: Which protocol can be used to carry customer routes between PE routers?

Answer: The number of customer routes can be very large. BGP is the only routing protocol that can scale to a very large number of routes.

Conclusion:

BGP is used to exchange customer routes directly between PE routers.

Propagation of Routing Information Across the P-Network (Cont.)

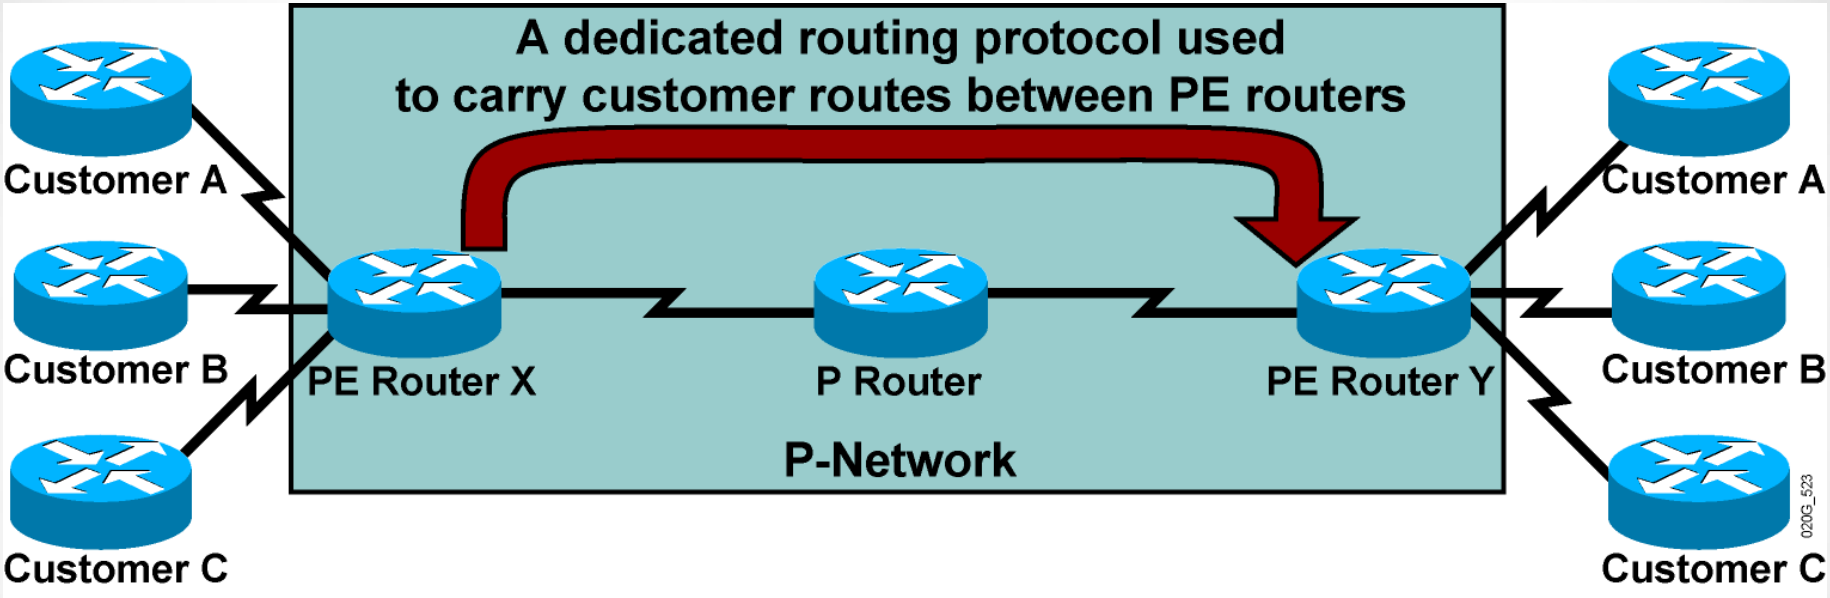

Question: How will information about the overlapping subnetworks of two customers be propagated via a single routing protocol?

Answer: Extend the customer addresses to make them unique.

Route Distinguishers

- The 64-bit route distinguisher is prepended to an IPv4 address to make it globally unique.
- The resulting address is a VPNv4 address.
- VPNv4 addresses are exchanged between PE routers via BGP.
  - BGP that supports address families other than IPv4 addresses is called MP-BGP.
- A similar process is used in IPv6:
  - 64-bit route distinguisher is prepended to a 16-byte IPv6 address.
  - The resulting 24-byte address is a unique VPNv6 address.

Route Distinguishers (Cont.)

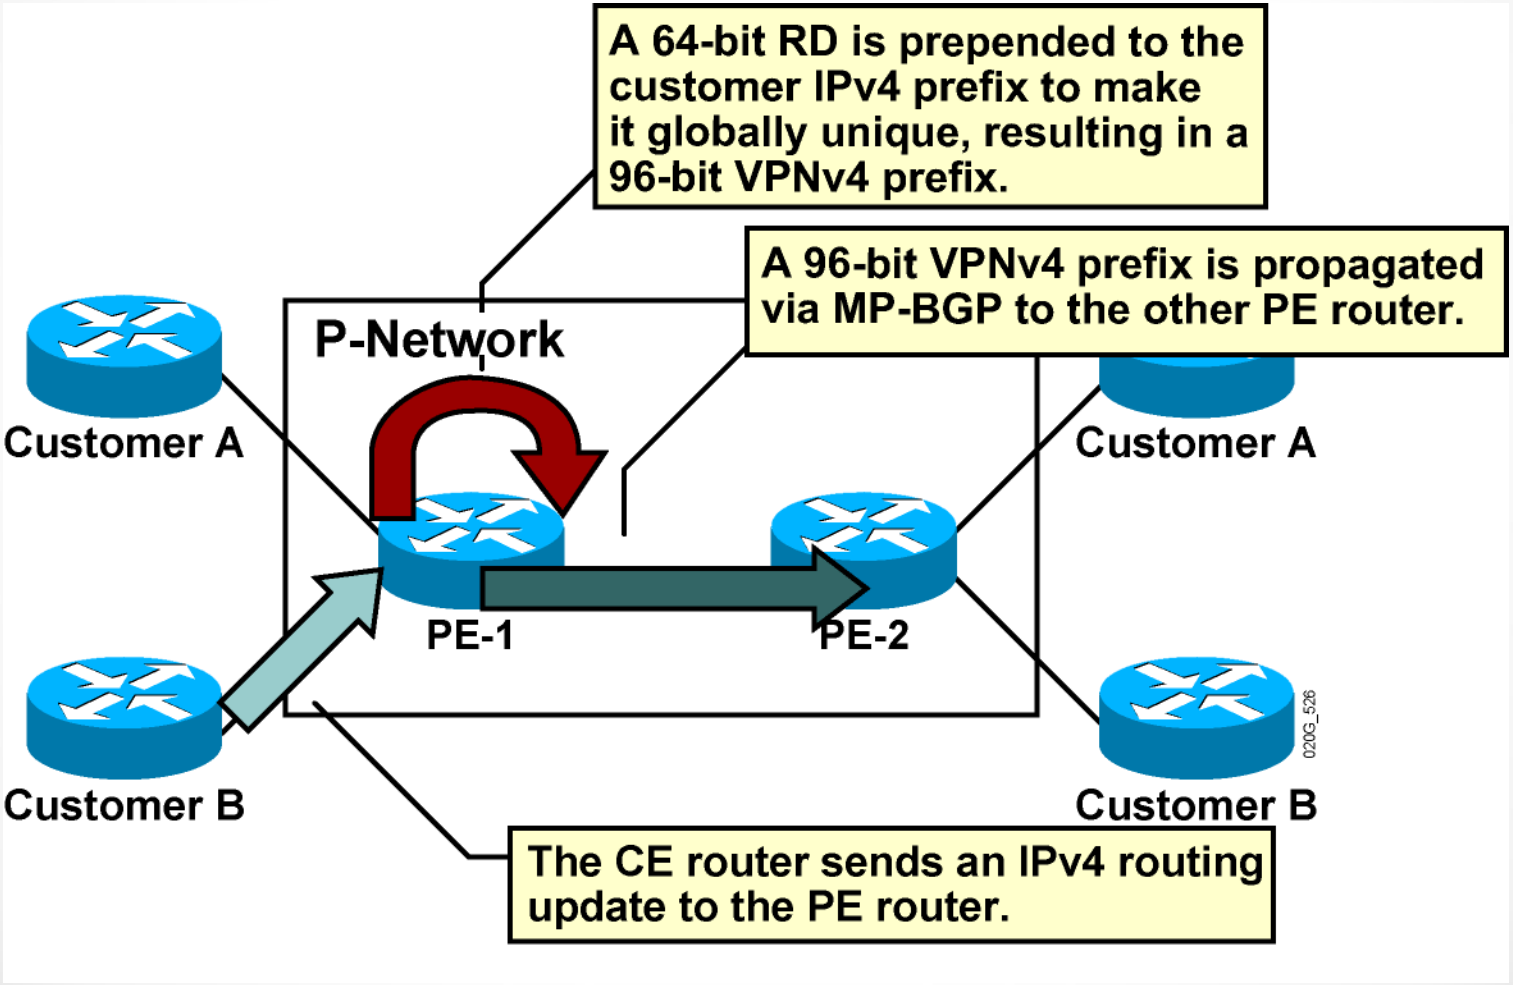

Route Distinguishers (Cont.)

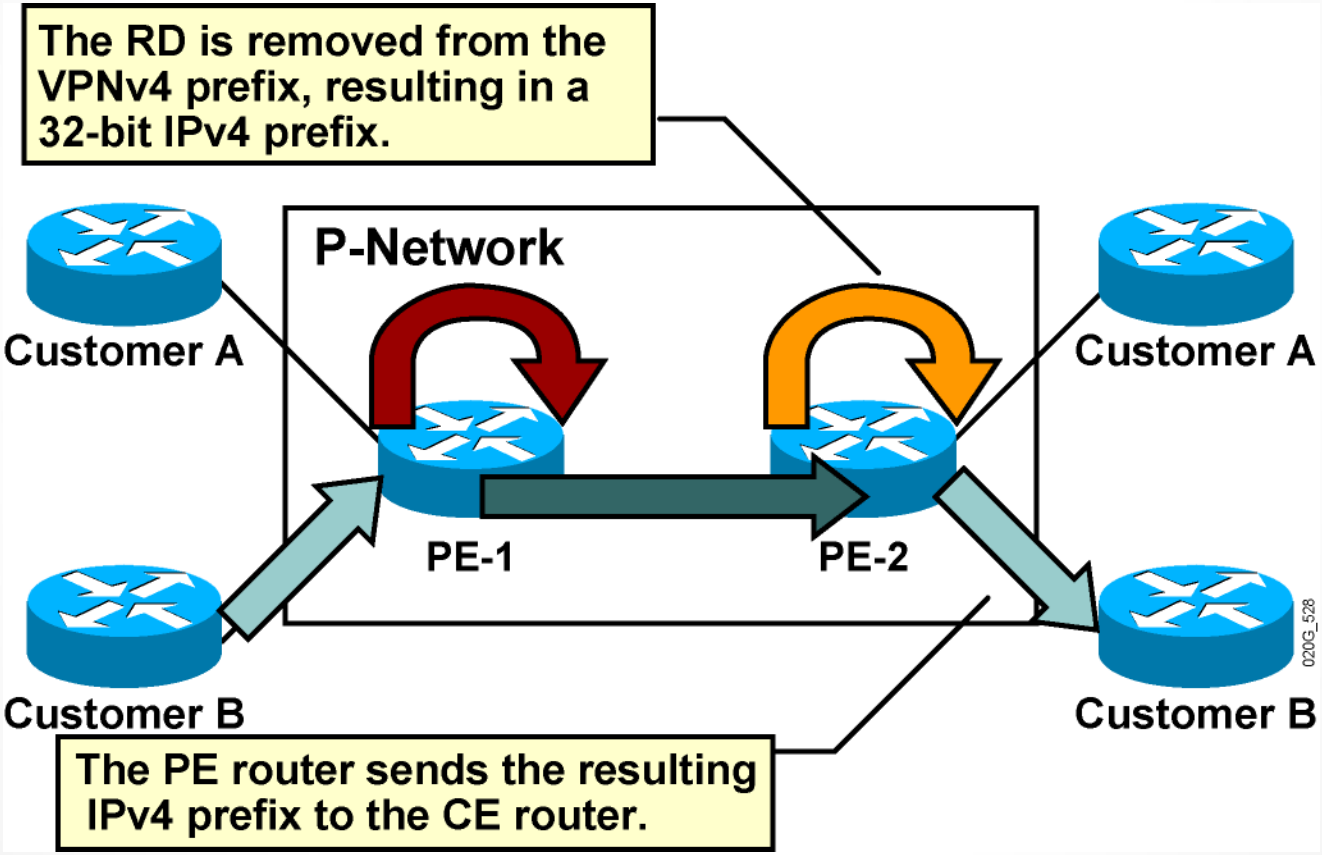

RDs: Usage in an MPLS VPN

- The RD has no special meaning.
- The RD is used only to make potentially overlapping IPv4 addresses globally unique.
- The RD is used as a VPN identifier, but this design could not support all topologies required by the customers.

RTs: Why Are They Needed?

- Some sites have to participate in more than one VPN.
- The RD cannot identify participation in more than one VPN.
- RTs were introduced in the MPLS VPN architecture to support complex VPN topologies.
  - A different method is needed in which a set of identifiers can be attached to a route.

RTs: What Are They?

- RTs are additional attributes attached to VPNv4 BGP routes to indicate VPN membership.
- Extended BGP communities are used to encode these attributes.
  - By marking a set of route provides a mechanism by which to group routes so that routing policies can be applied to all the routes within the same group
  - Extended communities carry the meaning of the attribute together with its value.
- Any number of RTs can be attached to a single route.

RTs: How Do They Work?

- Export RTs:
  - Identifying VPN membership
  - Appended to the customer route when it is converted into a VPNv4 route
- Import RTs:
  - Associated with each virtual routing table
  - Select routes to be inserted into the virtual routing table

RT and RD operation in an MPLS VPN

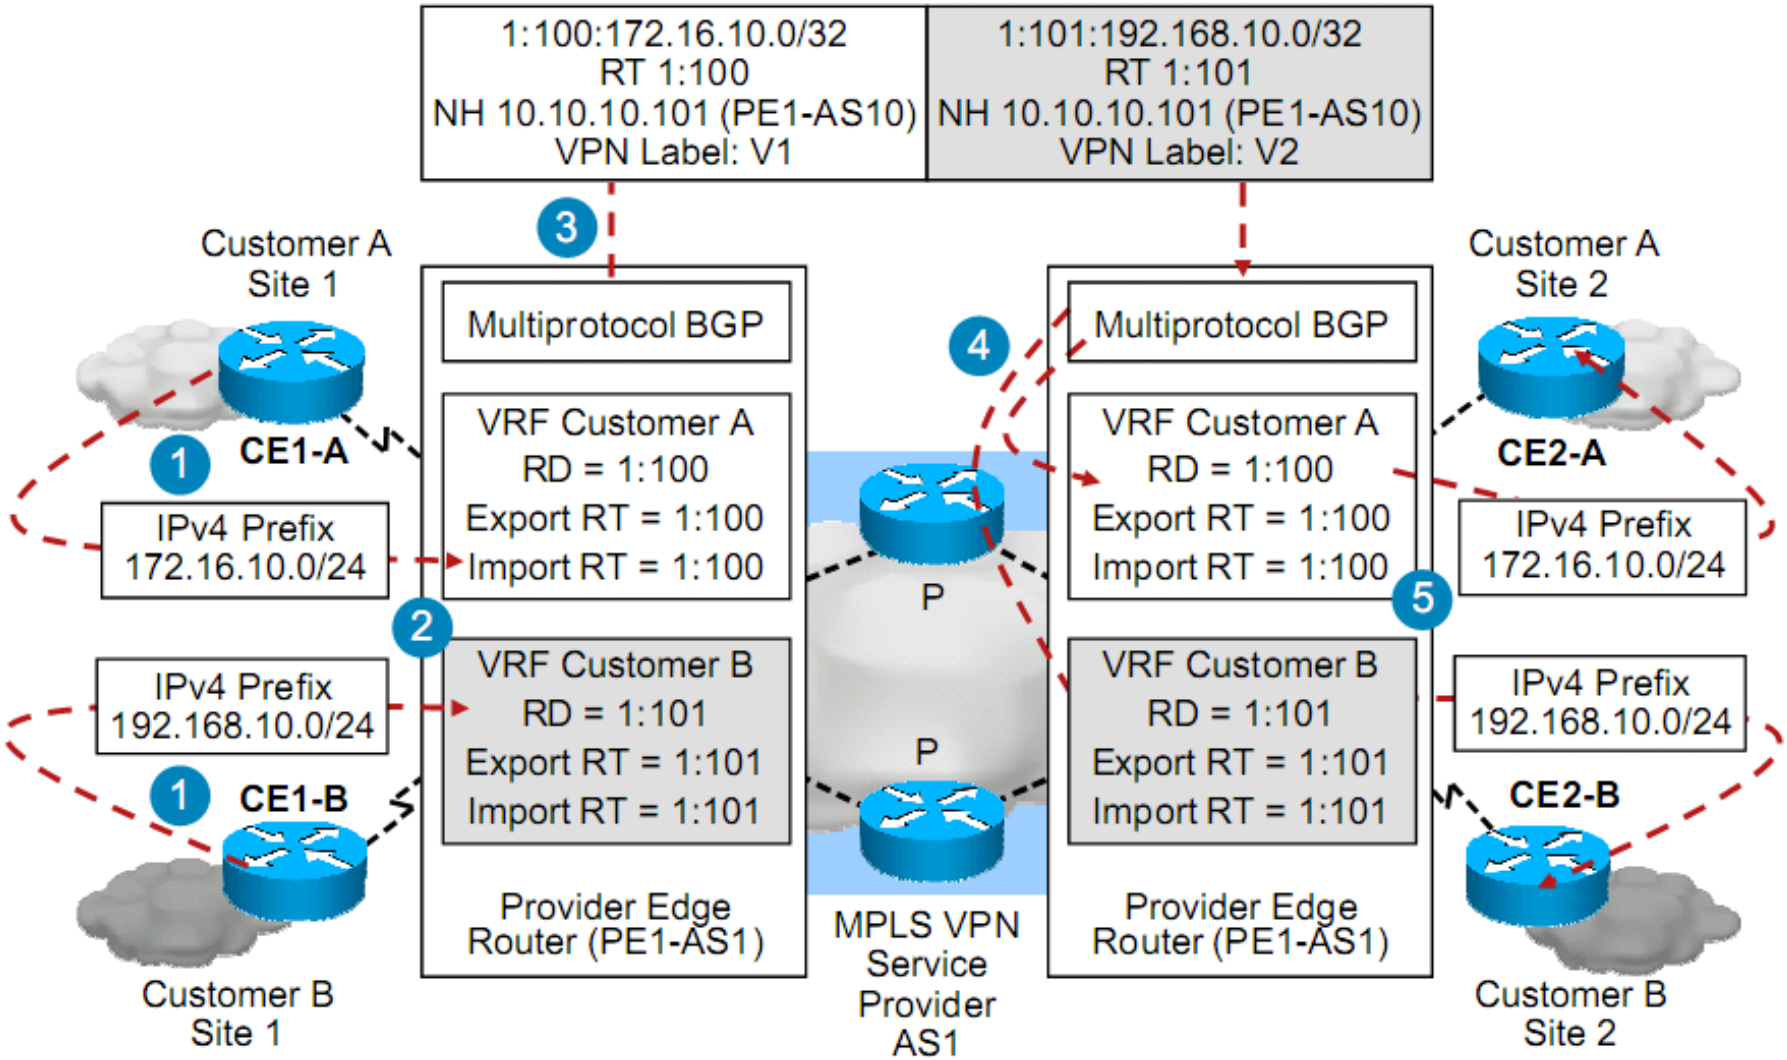

RT and RD operation in an MPLS VPN (cont.)

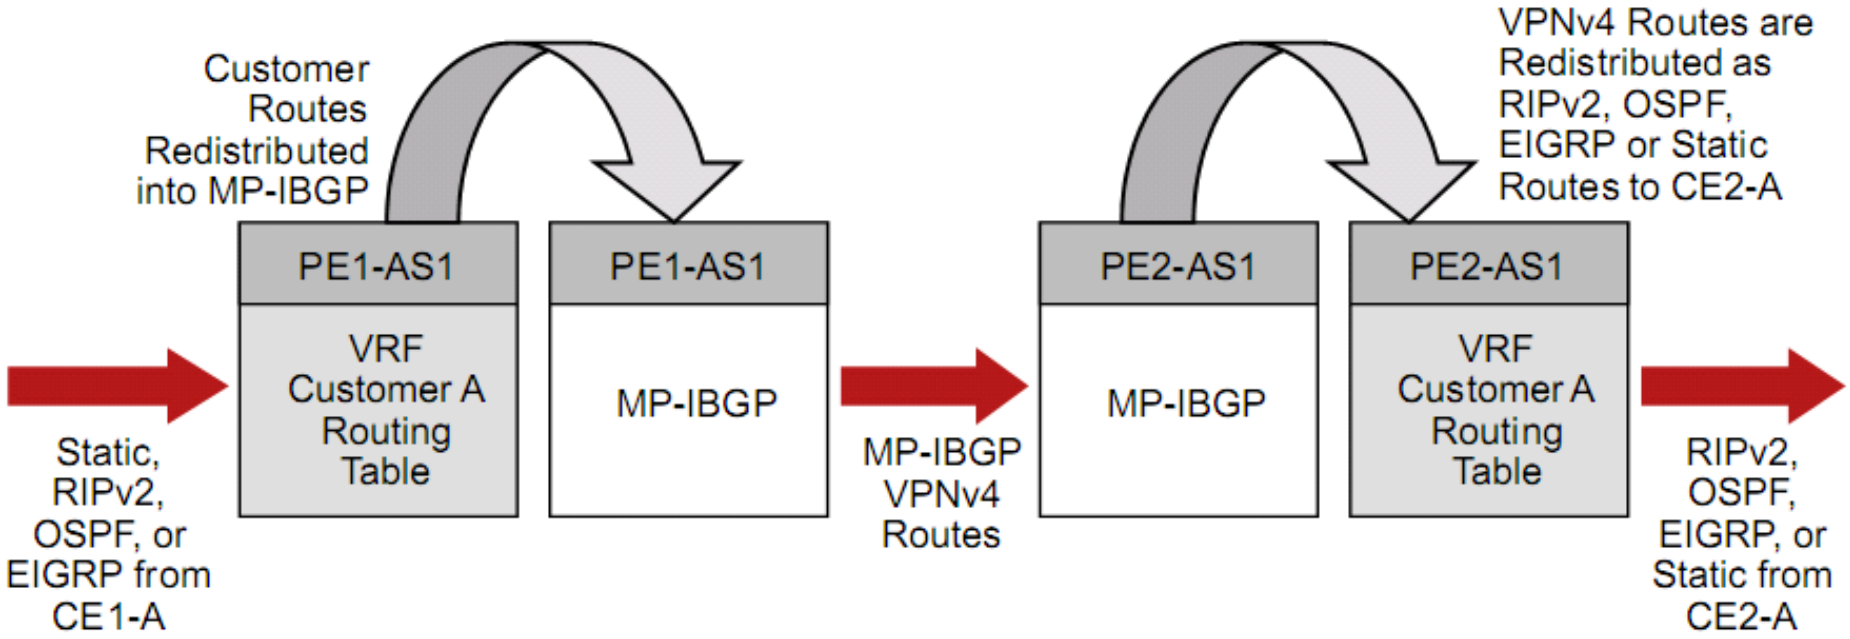

Impact of Complex VPN Topologies on Virtual Routing Tables

- A virtual routing table in a PE router can be used only for sites with identical connectivity requirements.
- Complex VPN topologies require more than one virtual routing table per VPN.
- As each virtual routing table requires a distinct RD value, the number of RDs in the MPLS VPN network increases.

MPLS VPN Technology

Introducing the MPLS VPN Routing Model

MPLS VPN Routing Requirements

- CE routers have to run standard IP routing software.
- PE routers have to support MPLS VPN services and IP routing.
- P routers have no VPN routes.

MPLS VPN Routing: CE Router Perspective

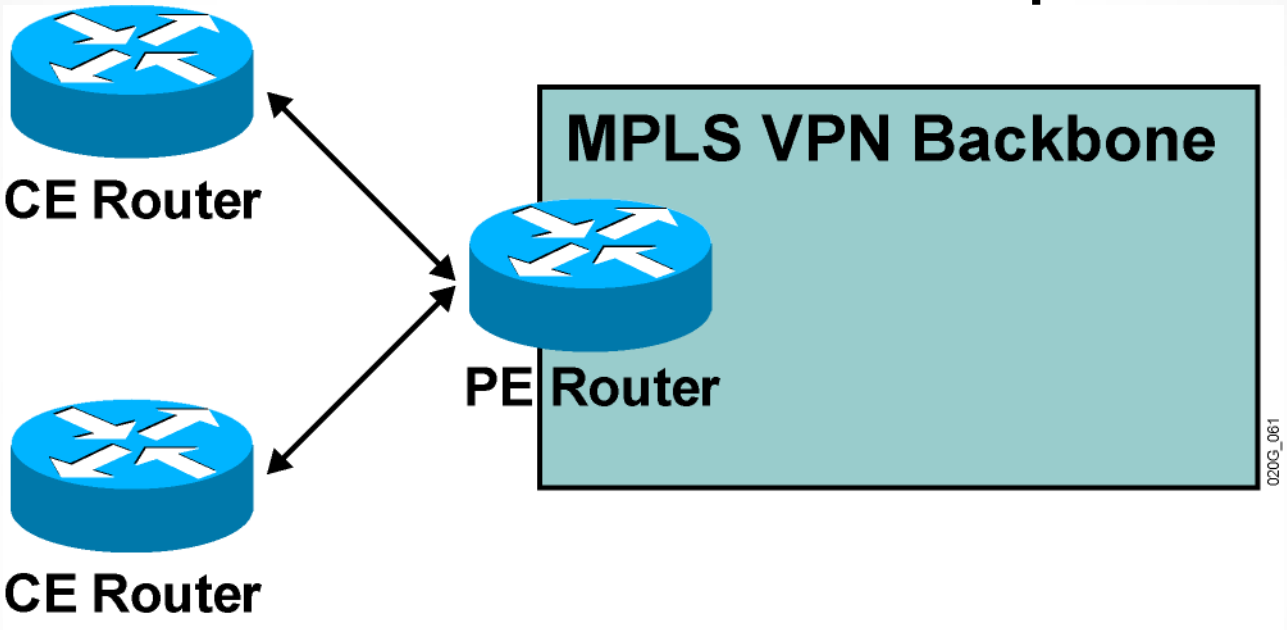

- The CE routers run standard IP routing software and exchange routing updates with the PE router.
  - EBGP, OSPF, RIPv2, EIGRP, and static routes are supported.
- The PE router appears as another router in the C-network.

MPLS VPN Routing: Overall Customer Perspective

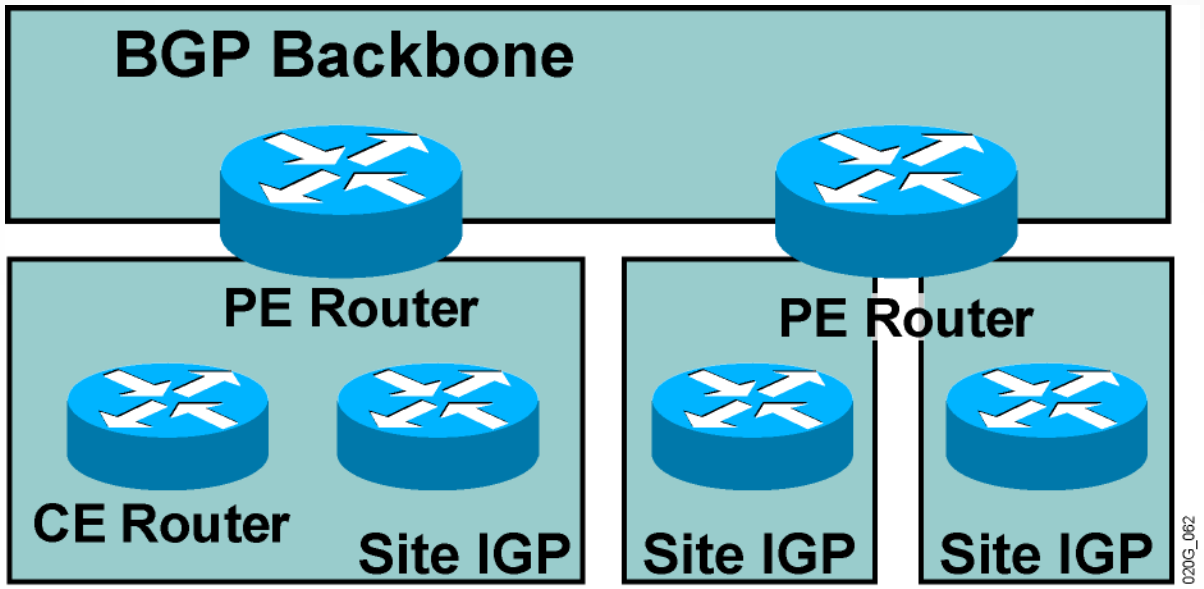

- To the customer, the PE routers appear as core routers connected via a BGP backbone.
- The usual BGP and IGP design rules apply.
- The P routers are hidden from the customer.

MPLS VPN Routing: P Router Perspective

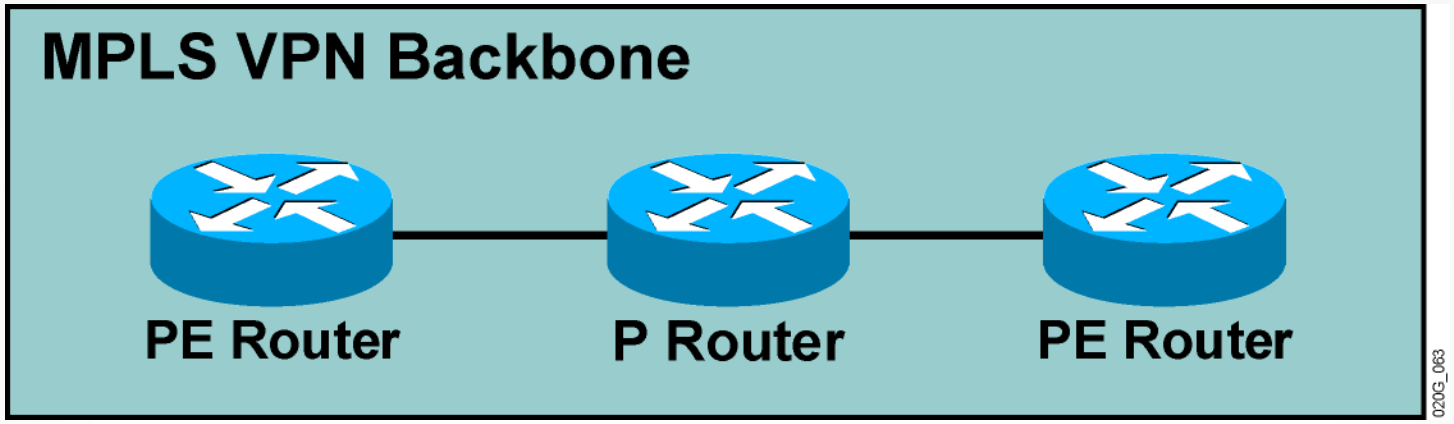

- P routers do not participate in MPLS VPN routing and do not carry VPN routes.
- P routers run backbone IGP with the PE routers and exchange information about global subnetworks (core links and loopbacks).

MPLS VPN Routing: PE Router Perspective

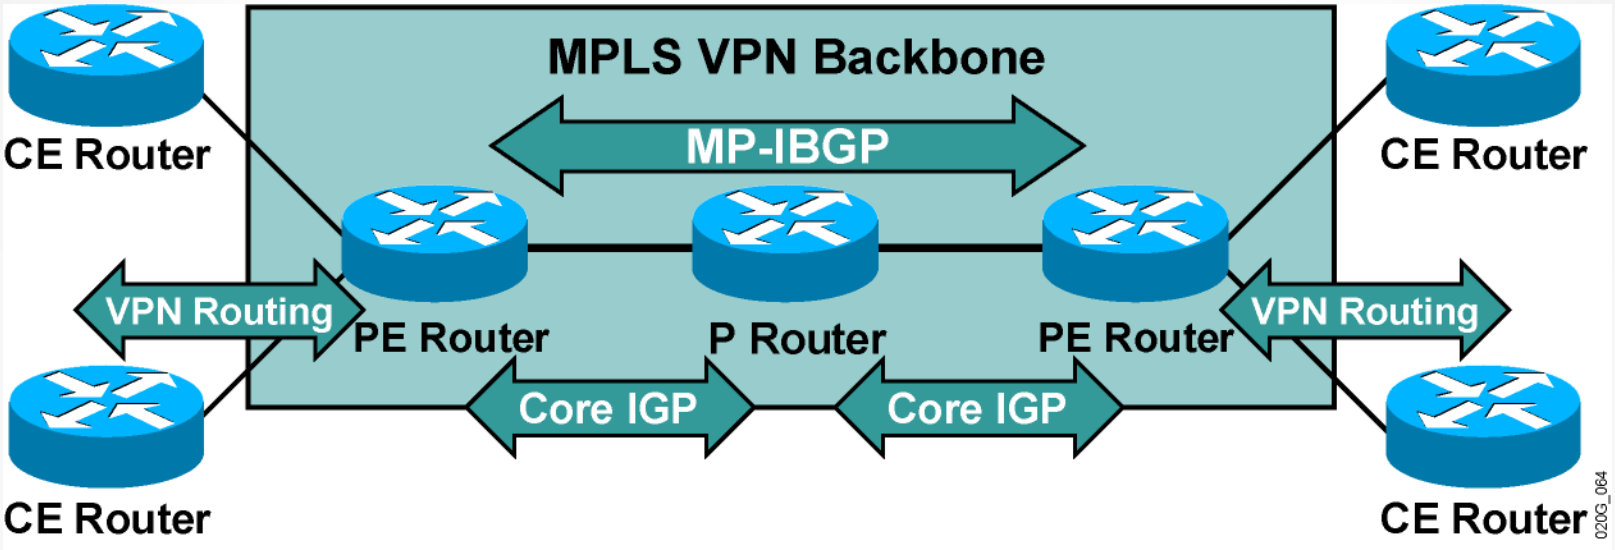

PE routers:

- **Exchange VPN routes with CE routers via per-VPN routing protocols**
- **Exchange core routes with P routers and PE routers via core IGP**
- **Exchange VPNv4 routes with other PE routers via MP-IBGP sessions**

Routing Tables on PE Routers

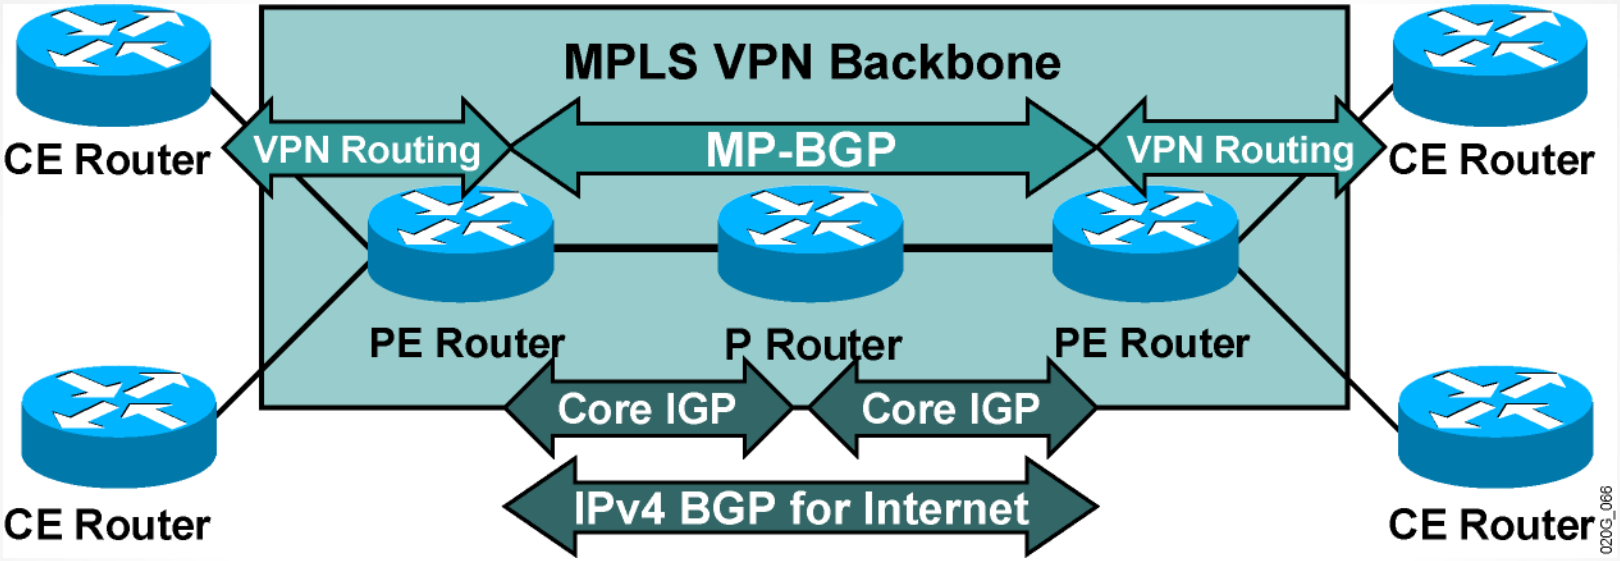

- PE routers contain a number of routing tables:
  - The global routing table contains core routes (filled with core IGP) and Internet routes (filled with IPv4 BGP).

The VPN Instance tables contains routes for sites of identical routing requirements from local (IPv4 VPN) and remote (VPNv4 via MP-BGP) CE routers.

End-to-End Routing Update Flow

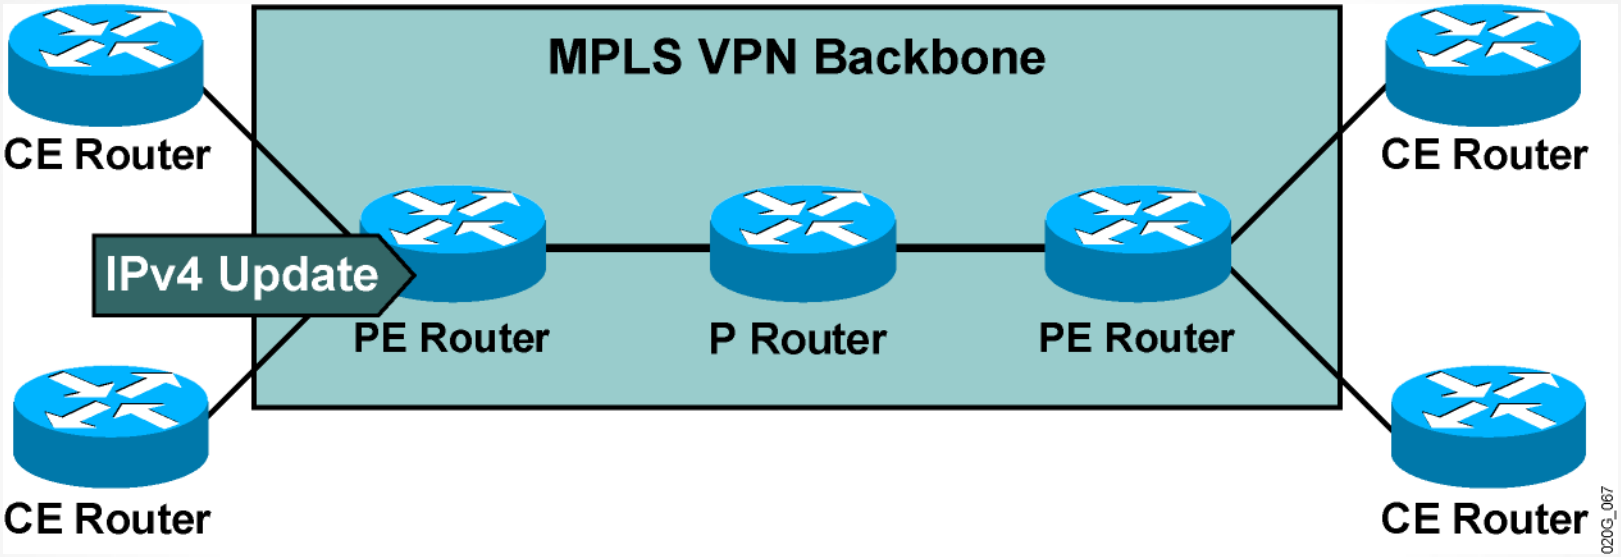

PE routers receive IPv4 routing updates from CE routers and install them in the appropriate VPN Instance table.

End-to-End Routing Update Flow (Cont.)

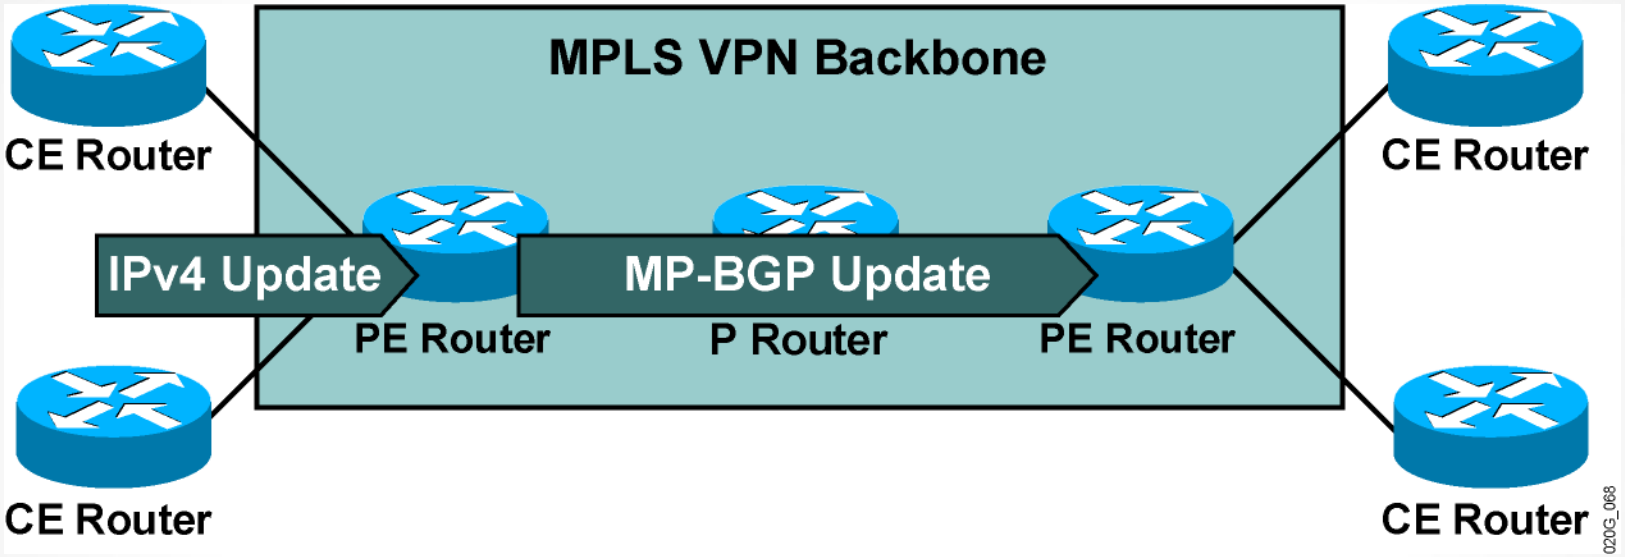

PE routers export VPN routes from VPN Instance tables into MP-BGP and propagate them as VPNv4 routes to other PE routers.

End-to-End Routing Update Flow: MP-BGP Update

- An MP-BGP update contains these elements:
  - VPNv4 address
  - Extended communities (route targets, optionally SOO (Site of Origin))
  - Label used for VPN packet forwarding
  - Any other BGP attribute (for example, AS path, local preference, MED, standard community)

End-to-End Routing Update Flow (Cont.)

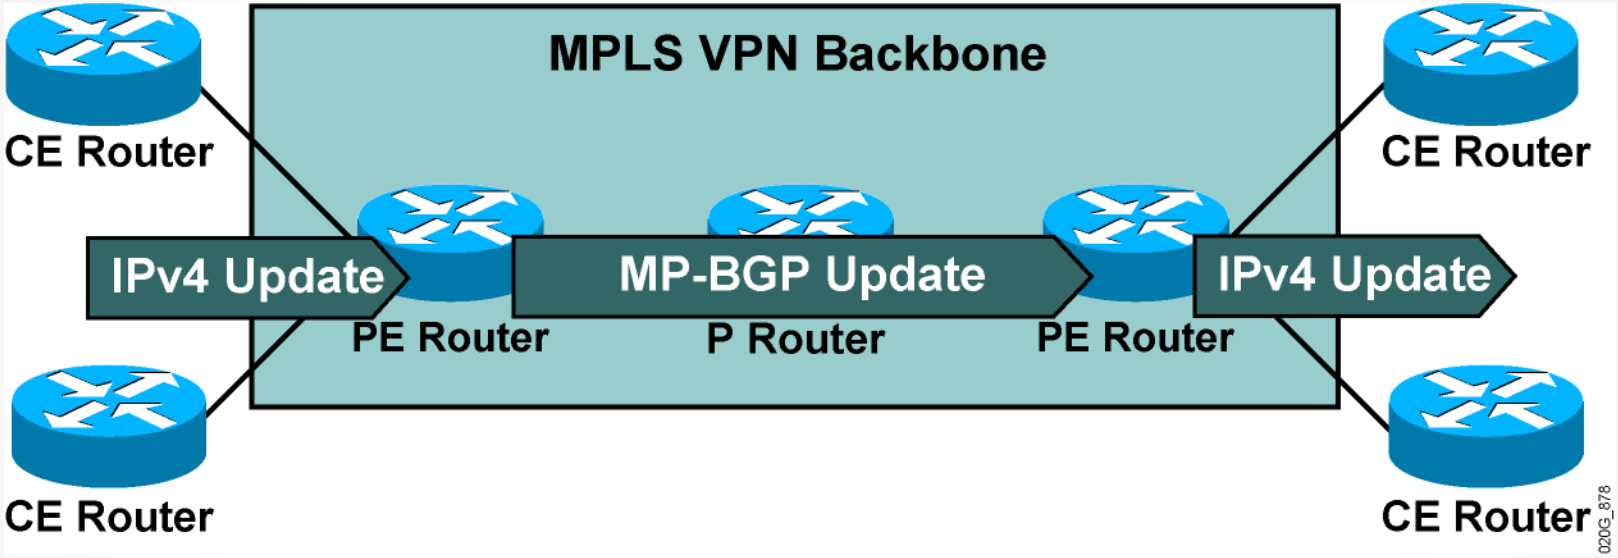

The receiving PE router imports the incoming VPNv4 routes into the appropriate VPN Instance based on route targets attached to the routes.

The routes installed in the VPN Instances are propagated to the CE routers.

Route Distribution to CE Routers

A route is installed in the site VPN Instance if it matches the import route target attribute.

- Route distribution to CE sites is driven by the following:
  - Route targets
  - SOO attribute if defined

MPLS VPN Technology

Forwarding MPLS VPN Packets

VPN Packet Forwarding Across an MPLS VPN Backbone: Approach 1

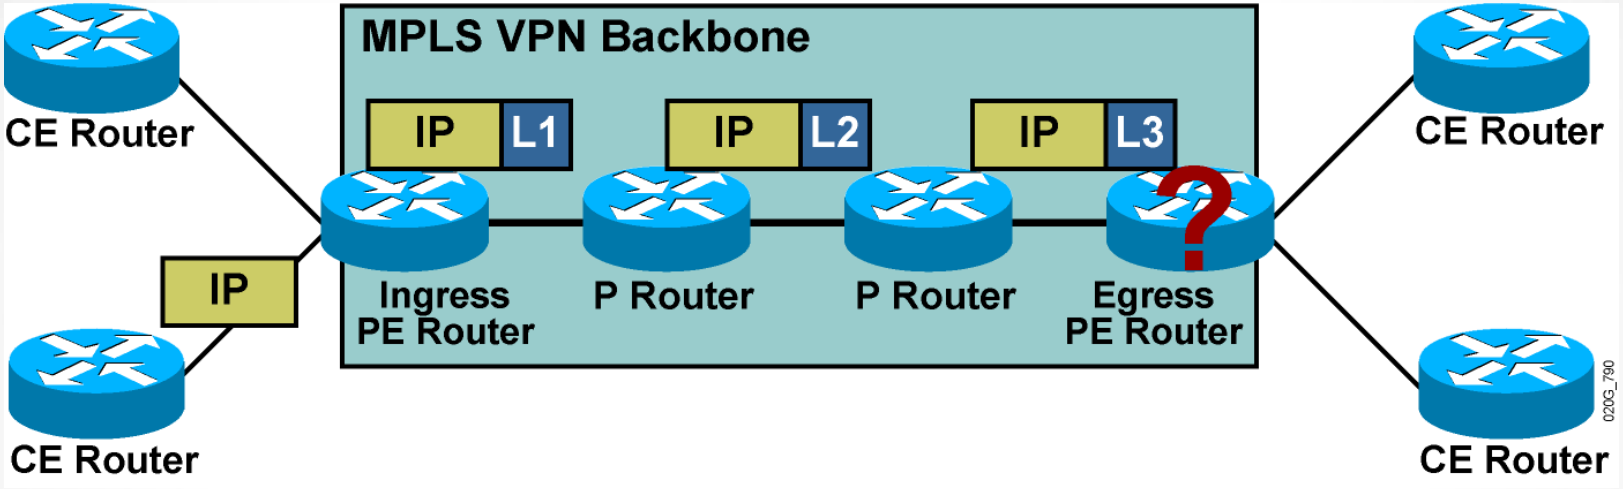

Approach 1: The PE routers will label the VPN packets with an LDP label for the egress PE router, and forward the labeled packets across the MPLS backbone.

Results:

- The P routers perform the label switching, and the packet reaches the egress PE router.

Because the egress PE router does not know which VPN Instance to use for packet switching, the packet is dropped.

VPN Packet Forwarding Across an MPLS VPN Backbone: Approach 2

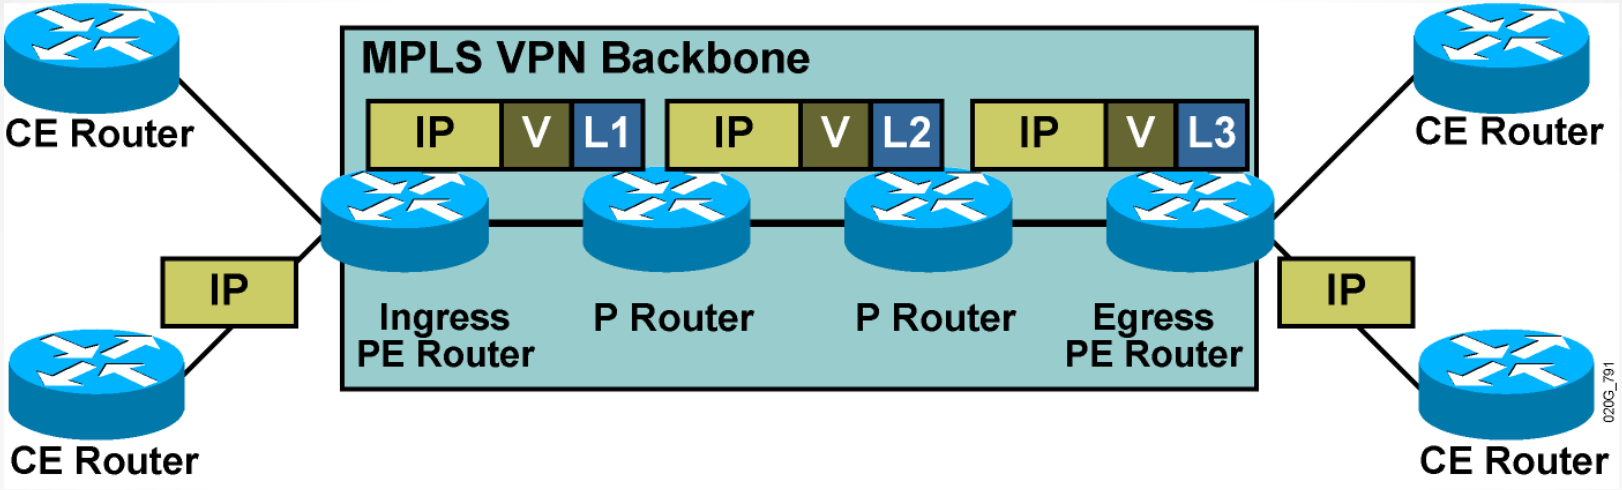

Approach 2: The PE routers will label the VPN packets with a label stack, using the LDP label for the egress PE router as the top label, and the VPN label assigned by the egress PE router as the second label in the stack.

Result:

- The P routers perform label switching using the top label, and the packet reaches the egress PE router. The top label is removed.
- The egress PE router performs a lookup on the VPN label and forwards the packet toward the CE router.

VPN PHP

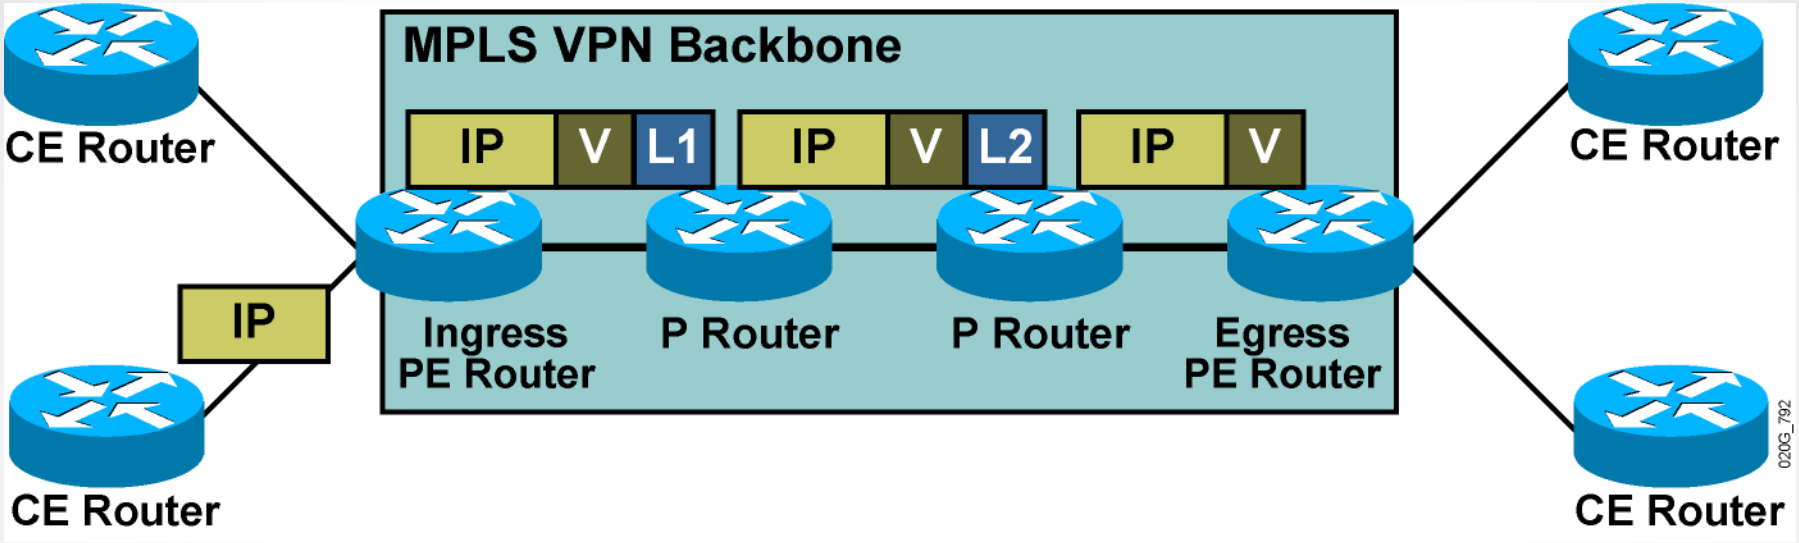

- Penultimate hop popping on the LDP label can be performed on the last P router.
- The egress PE router performs label lookup only on the VPN label, resulting in faster and simpler label lookup.
- IP lookup is performed only once—in the ingress PE router.

VPN Label Propagation

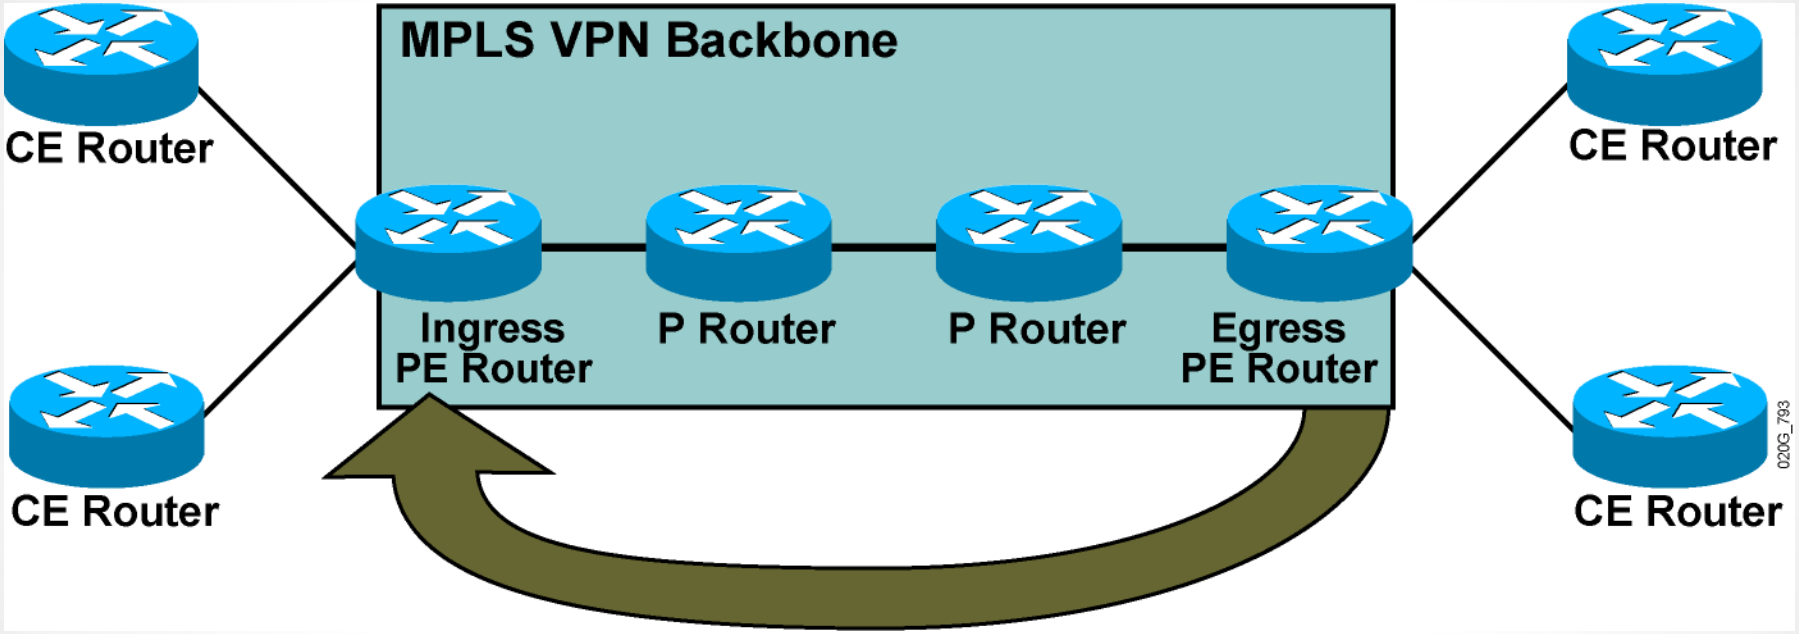

Question: How will the ingress PE router get the second label in the label stack from the egress PE router?

Answer: Labels are propagated in MP-BGP VPNv4 routing updates.

VPN Label Propagation (Cont.)

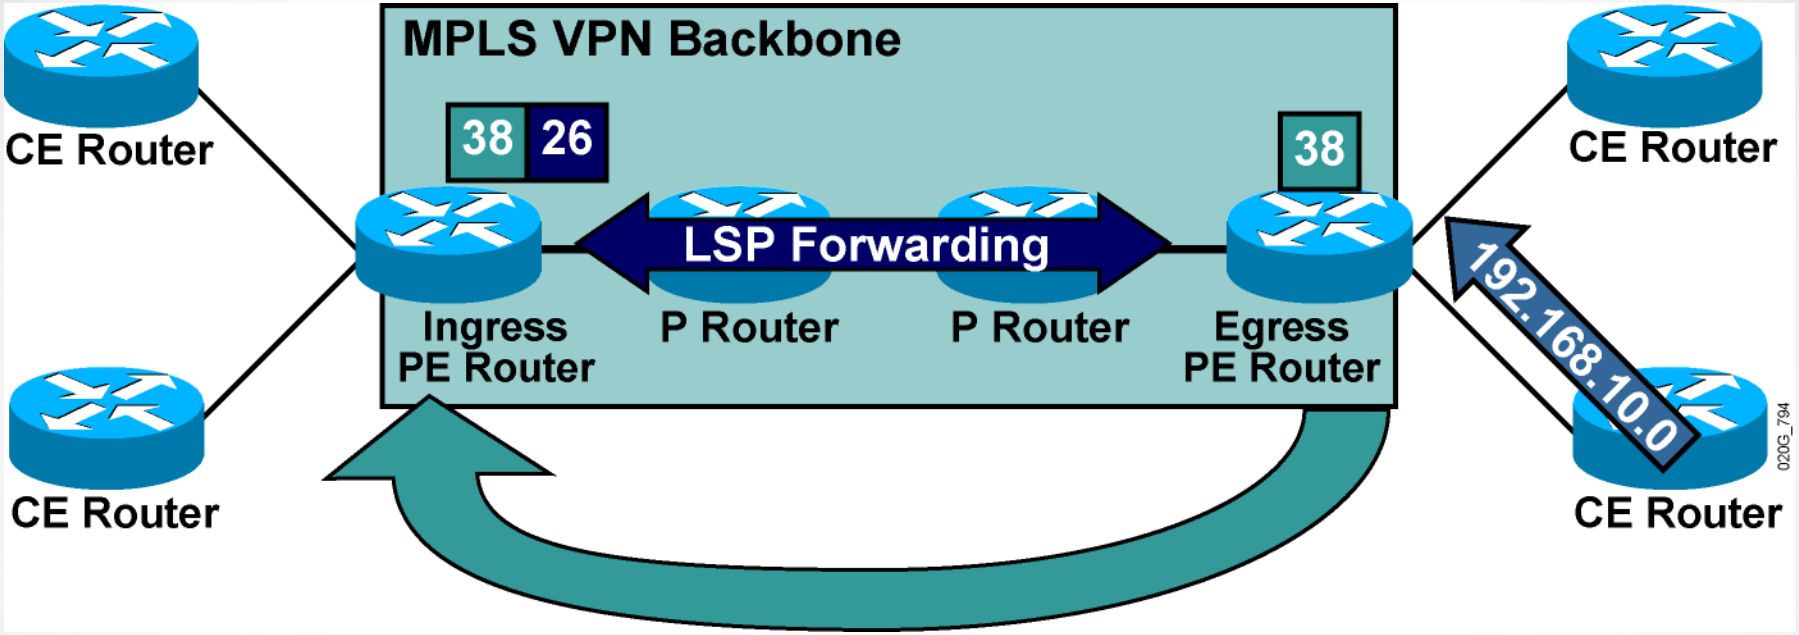

Step 1: A VPN label is assigned to every VPN route by the egress PE router.

Step 2: The VPN label is advertised to all other PE routers in an MP-BGP update.

Step 3: A label stack is built in the VPN Instance table.

MPLS L3 VPN data plane

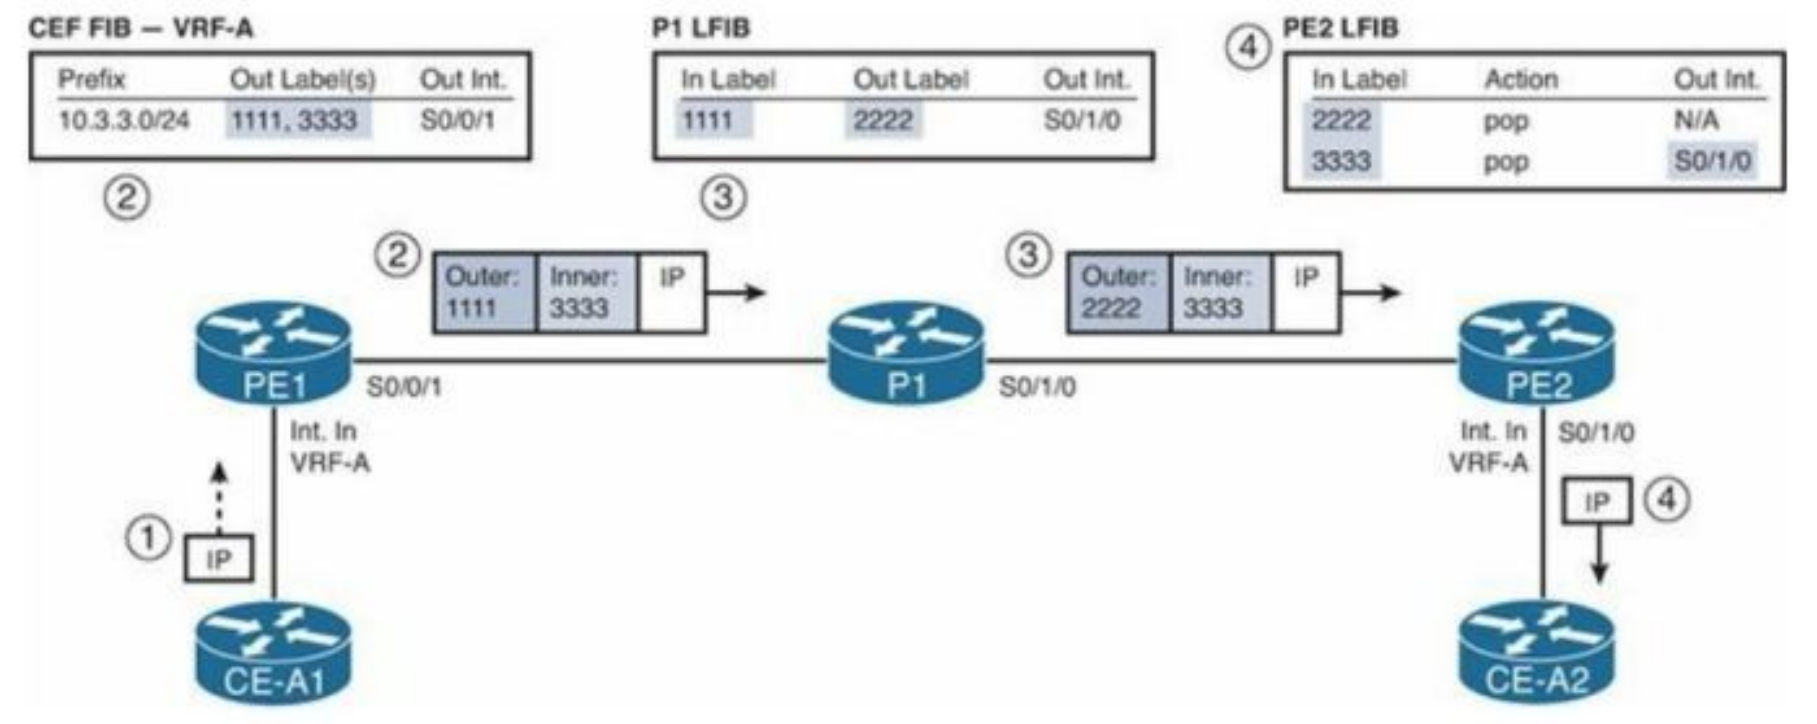

Route Target (RT)

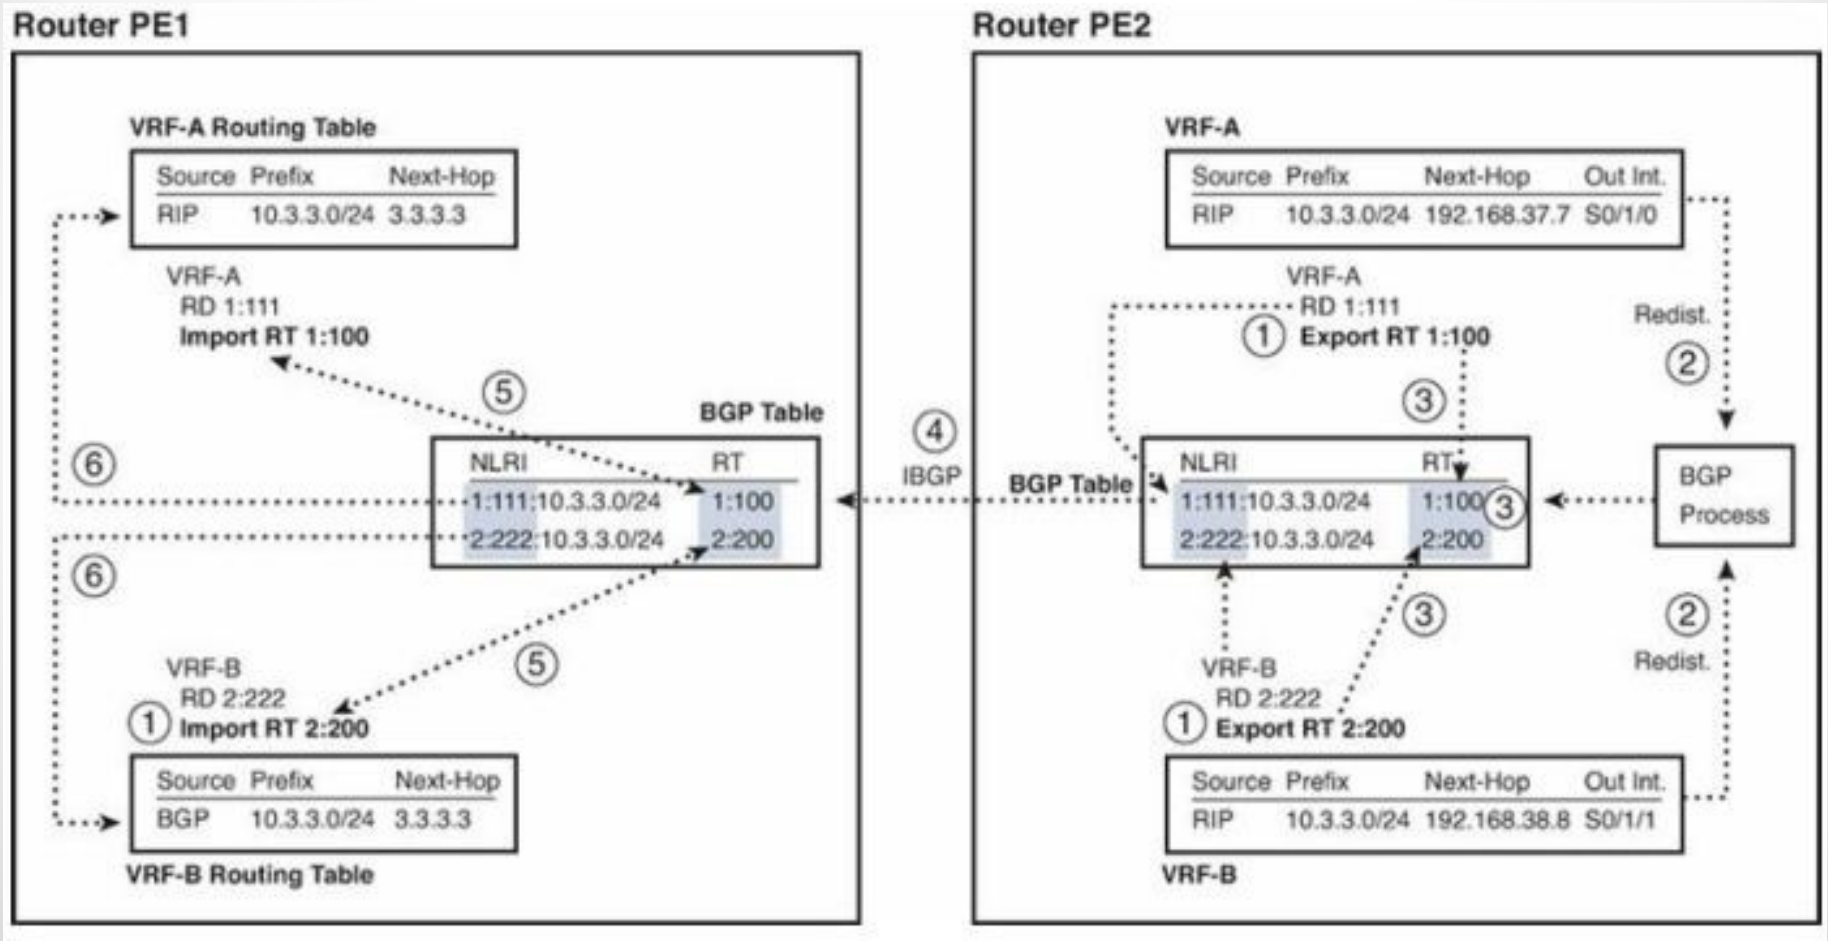

Route Distinguisher (RD)

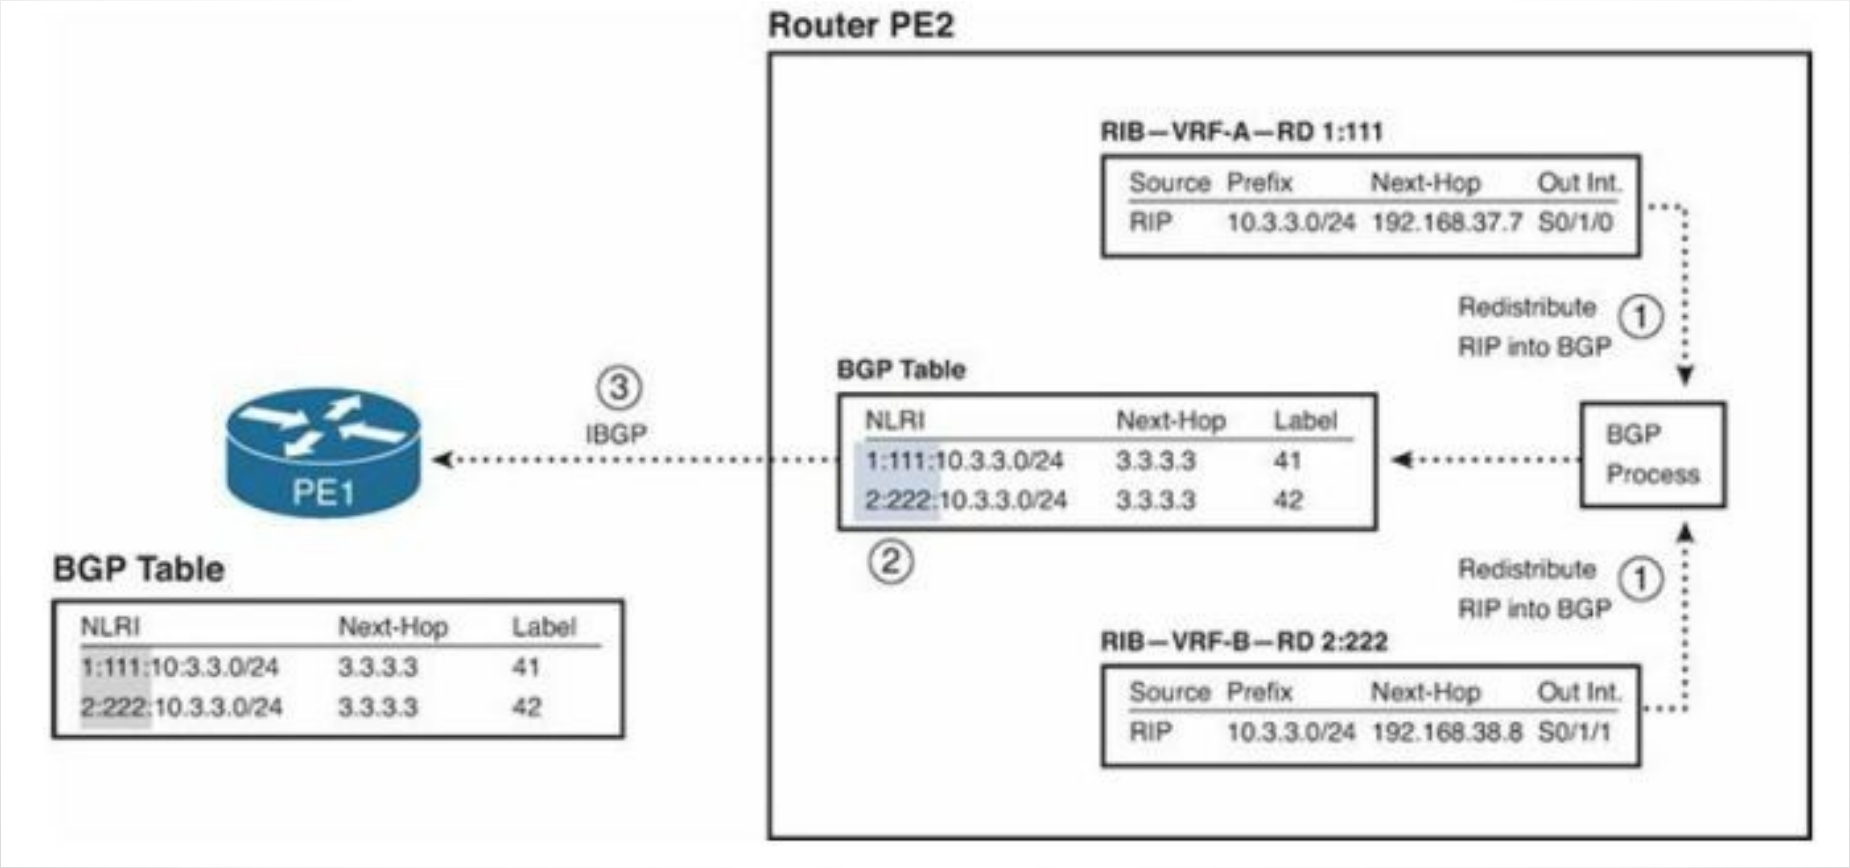

VRF

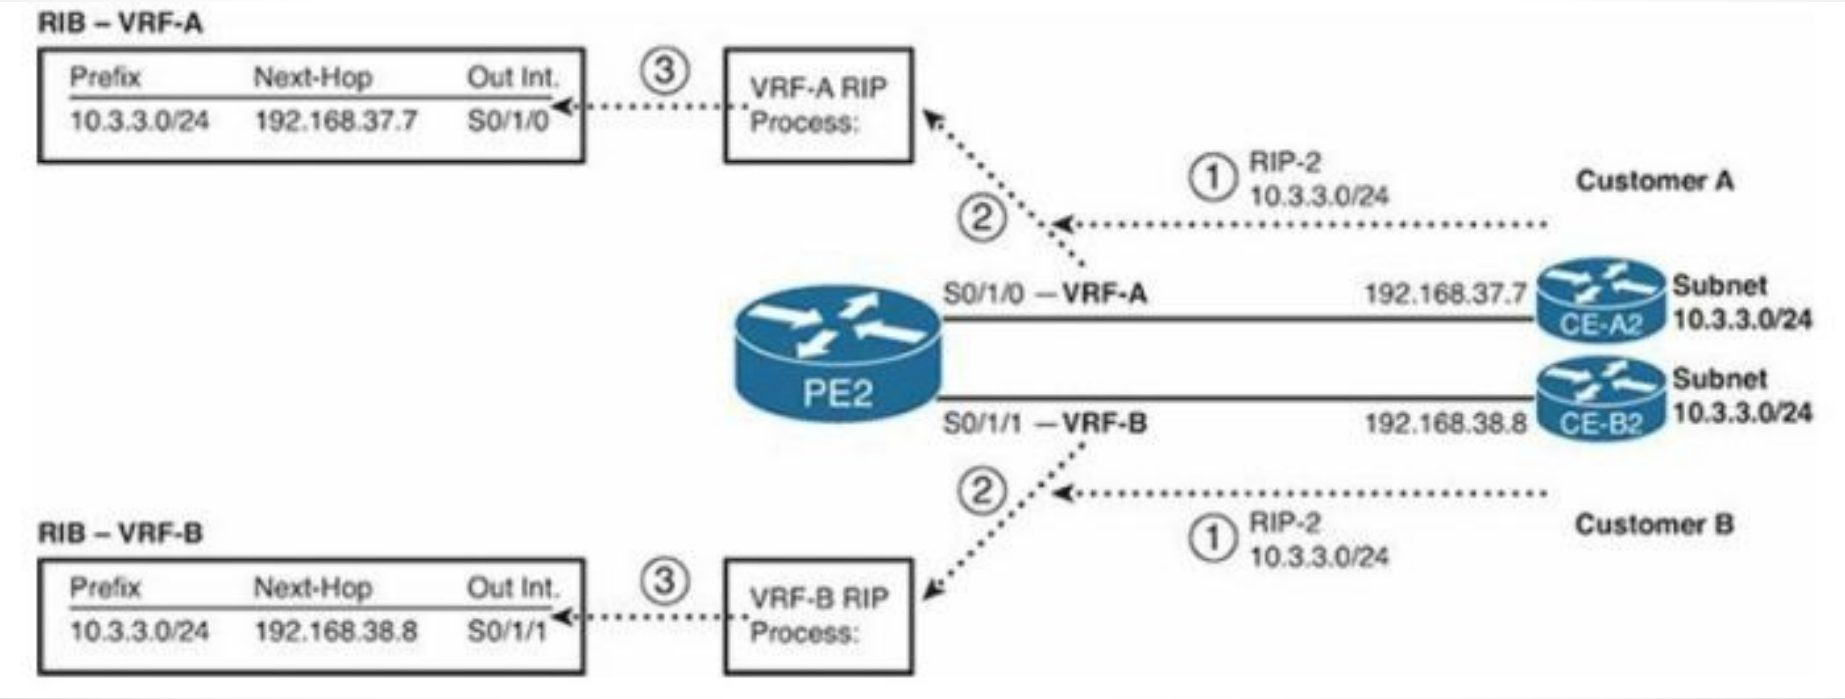
RIB – VRF-A
Prefix
Next-Hop
Out Int.
10.3.3.0/24
192.168.37.7
S0/1/0
3
VRF-A RIP Process:
1
RIP-2
10.3.3.0/24
Customer A
2
S0/1/0 – VRF-A
192.168.37.7
CE-A2
Subnet 10.3.3.0/24
PE2
S0/1/1 – VRF-B
192.168.38.8
CE-B2
Subnet 10.3.3.0/24
2
1
RIP-2
10.3.3.0/24
Customer B
RIB – VRF-B
Prefix
Next-Hop
Out Int.
10.3.3.0/24
192.168.38.8
S0/1/1
3
VRF-B RIP Process:

MPLS VPN Implementation

Configuration example

MPLS VPN Configuration Example

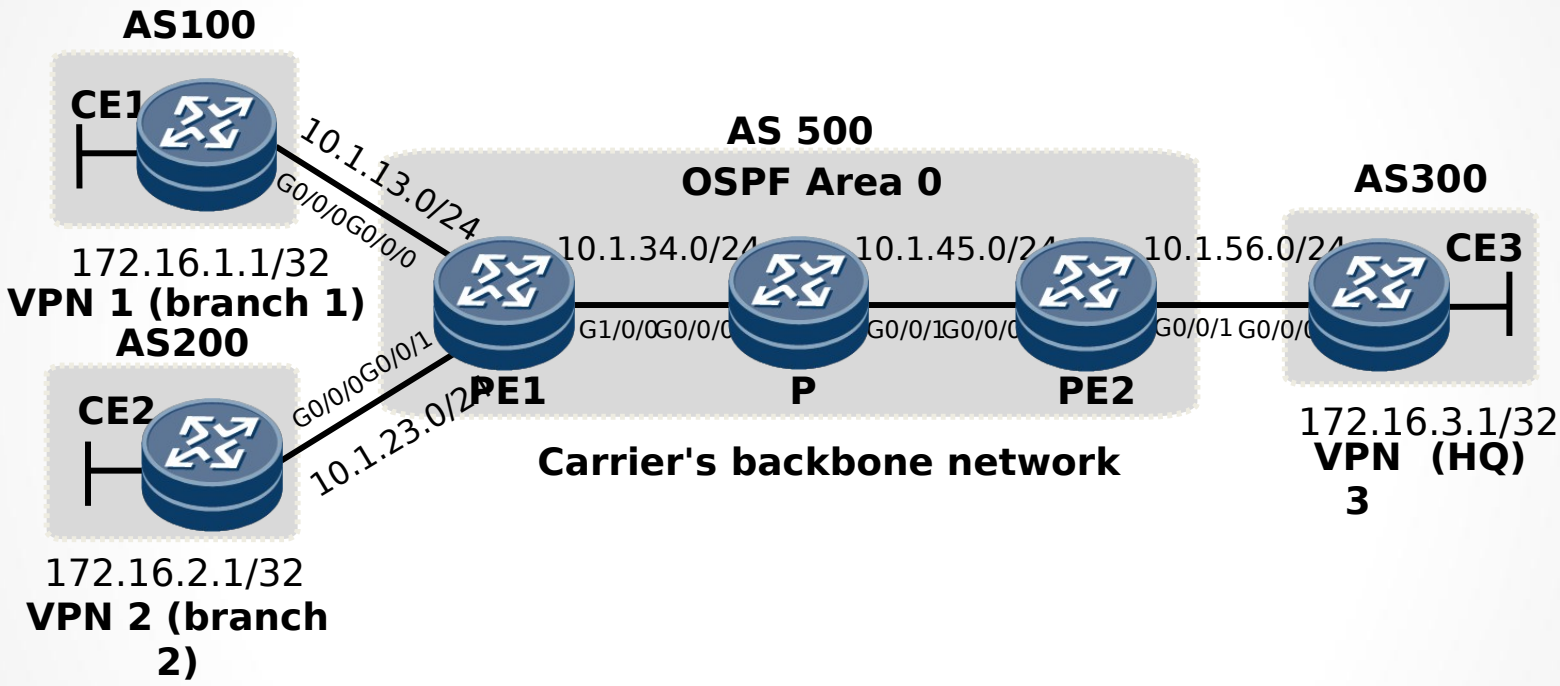

- Both branch 1 and branch 2 can communicate with the headquarters, but branches 1 and 2 cannot communicate with each other. Correctly configure the devices based on information in the figure to allow headquarters users to access the branch users.

Configuring User-side Devices

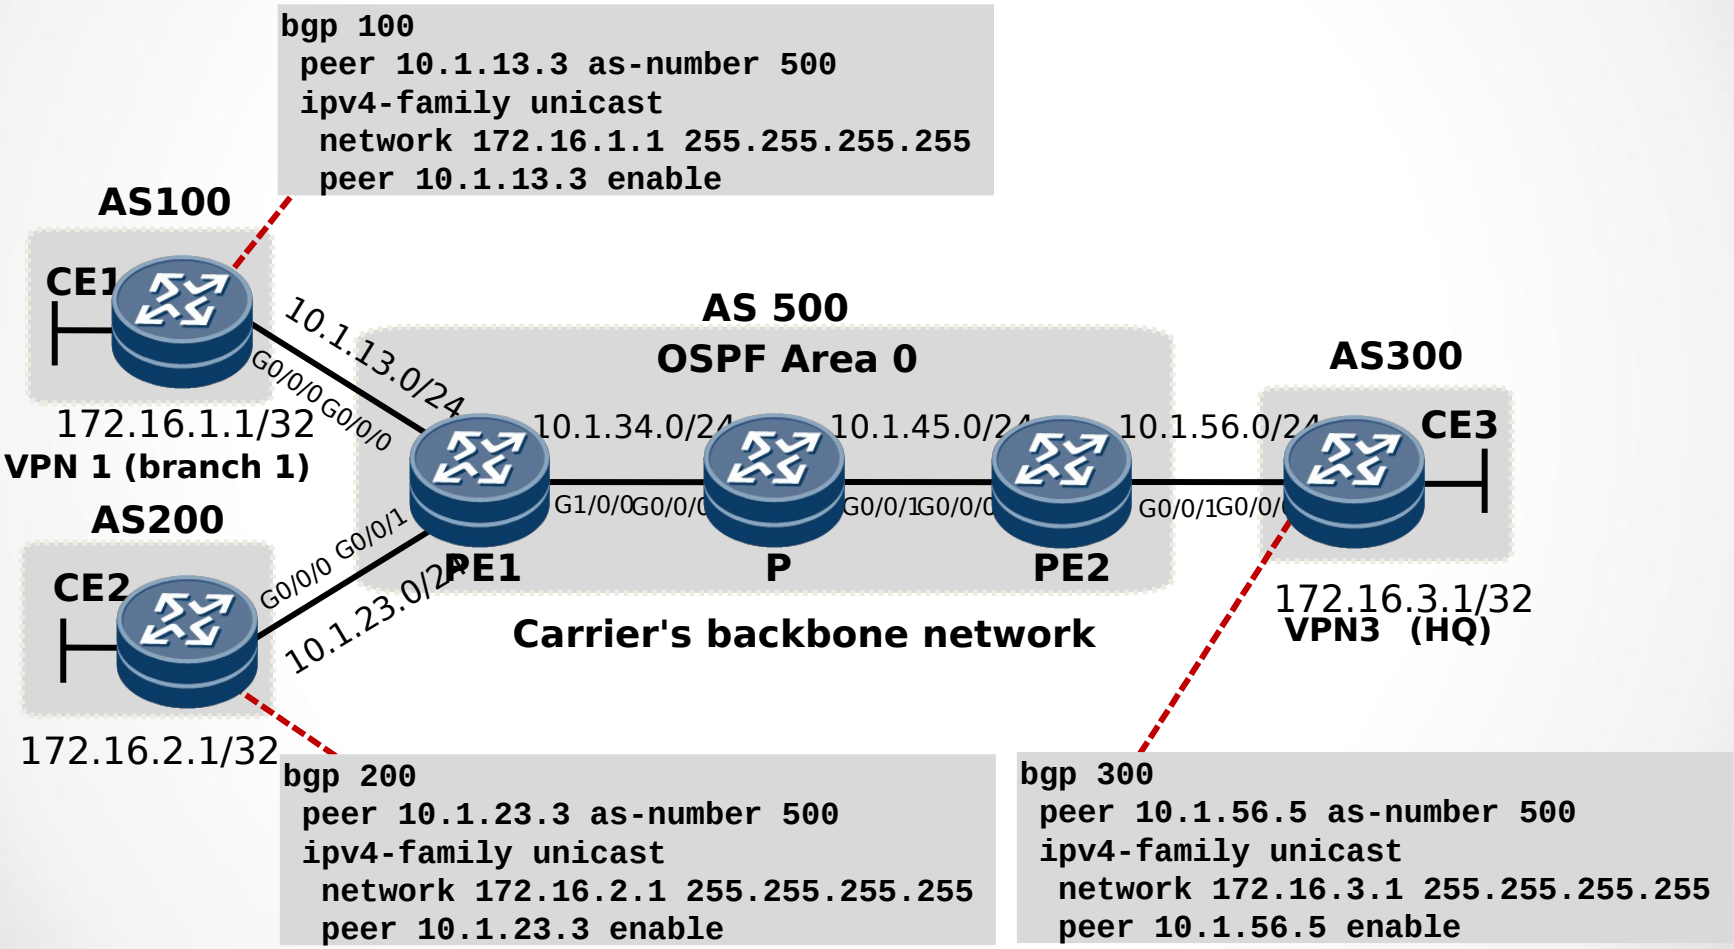

Configuring IGP on the Backbone Network

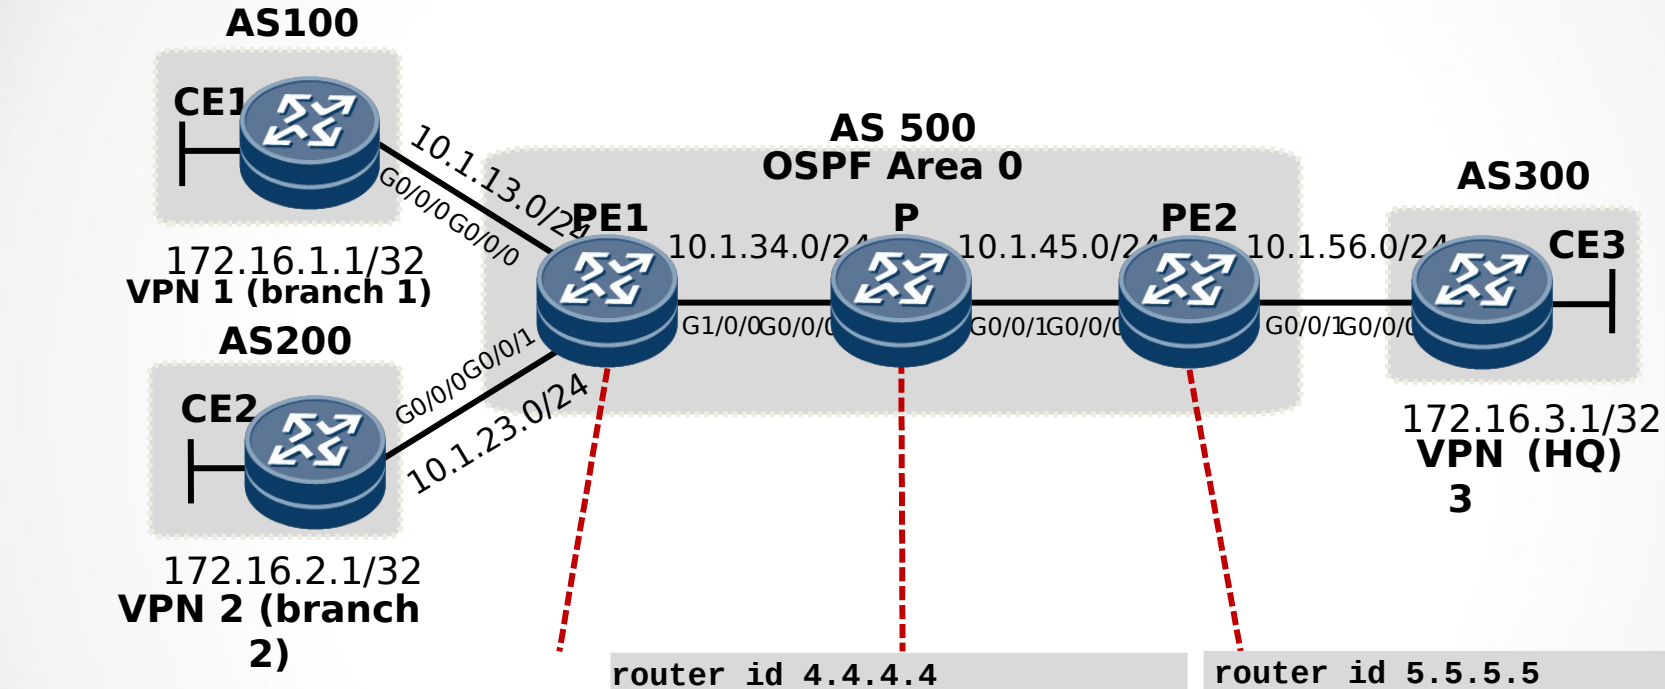


```
router id 3.3.3.3
ospf 1
 area 0.0.0.0
  network 3.3.3.3 0.0.0.0
  network 10.1.34.0 0.0.0.255
```

```
router id 4.4.4.4
ospf 1
 area 0.0.0.0
  network 4.4.4.4 0.0.0.0
  network 10.1.34.0 0.0.0.255
  network 10.1.45.0 0.0.0.255
```

```
router id 5.5.5.5
ospf 1
 area 0.0.0.0
  network 5.5.5.5 0.0.0.0
  network 10.1.45.5 0.0.0.255
```

Configuring VPN Instances

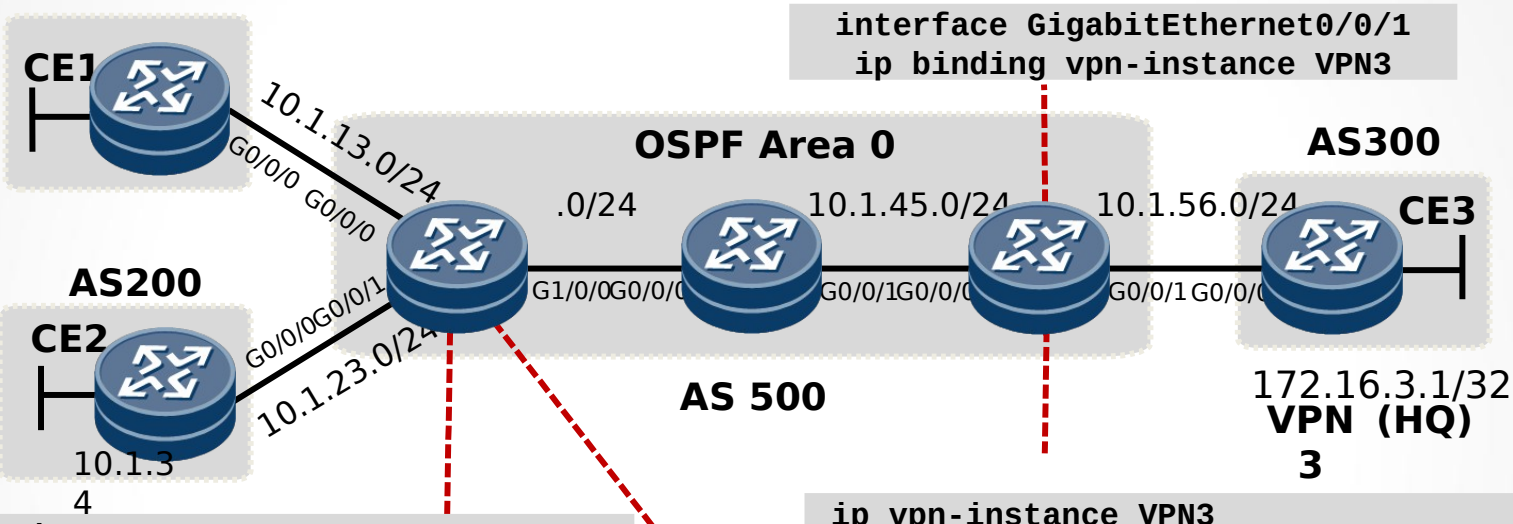


```
interface GigabitEthernet0/0/1
 ip binding vpn-instance VPN3
```

```
ip vpn-instance VPN1
 ipv4-family
  route-distinguisher 1:1
  vpn-target 12:3 export-extcommunity
  vpn-target 3:12 import-extcommunity
#
ip vpn-instance VPN2
 ipv4-family
  route-distinguisher 2:2
  vpn-target 12:3 export-extcommunity
  vpn-target 3:12 import-extcommunity
```

```
ip vpn-instance VPN3
 ipv4-family
  route-distinguisher 3:3
  vpn-target 3:12 export-extcommunity
  vpn-target 12:3 import-extcommunity
```

```
interface GigabitEthernet0/0/0
 ip binding vpn-instance VPN1
#
interface GigabitEthernet0/0/1
 ip binding vpn-instance VPN2
```

Configuring MP-BGP

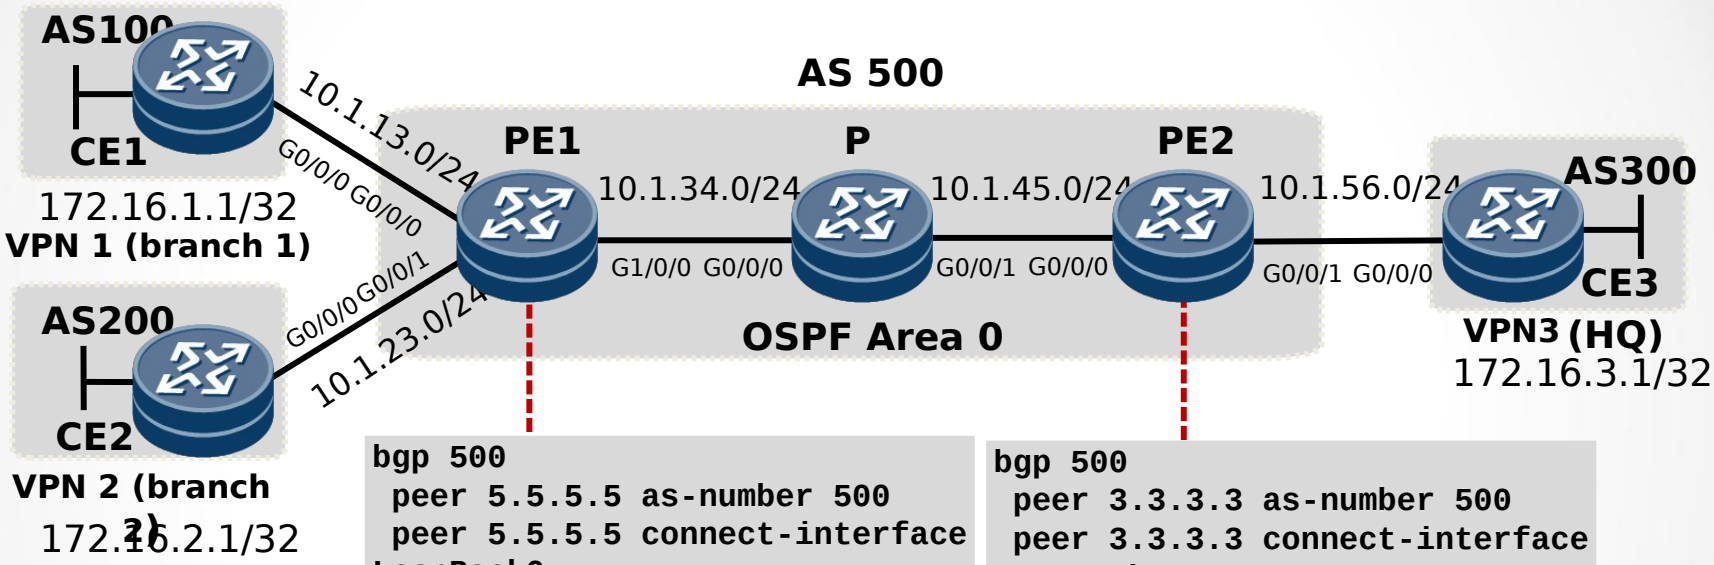


```
bgp 500
 peer 5.5.5.5 as-number 500
 peer 5.5.5.5 connect-interface
LoopBack0
#
 ipv4-family vpnv4
  peer 5.5.5.5 enable
 #
 ipv4-family vpn-instance VPN1
  peer 10.1.13.1 as-number 100
 #
 ipv4-family vpn-instance VPN2
  peer 10.1.23.2 as-number 200
```

```
bgp 500
 peer 3.3.3.3 as-number 500
 peer 3.3.3.3 connect-interface
LoopBack0
 #
 ipv4-family vpnv4
  policy vpn-target
  peer 3.3.3.3 enable
 #
 ipv4-family vpn-instance VPN3
  peer 10.1.56.6 as-number 300
```

Configuring MPLS

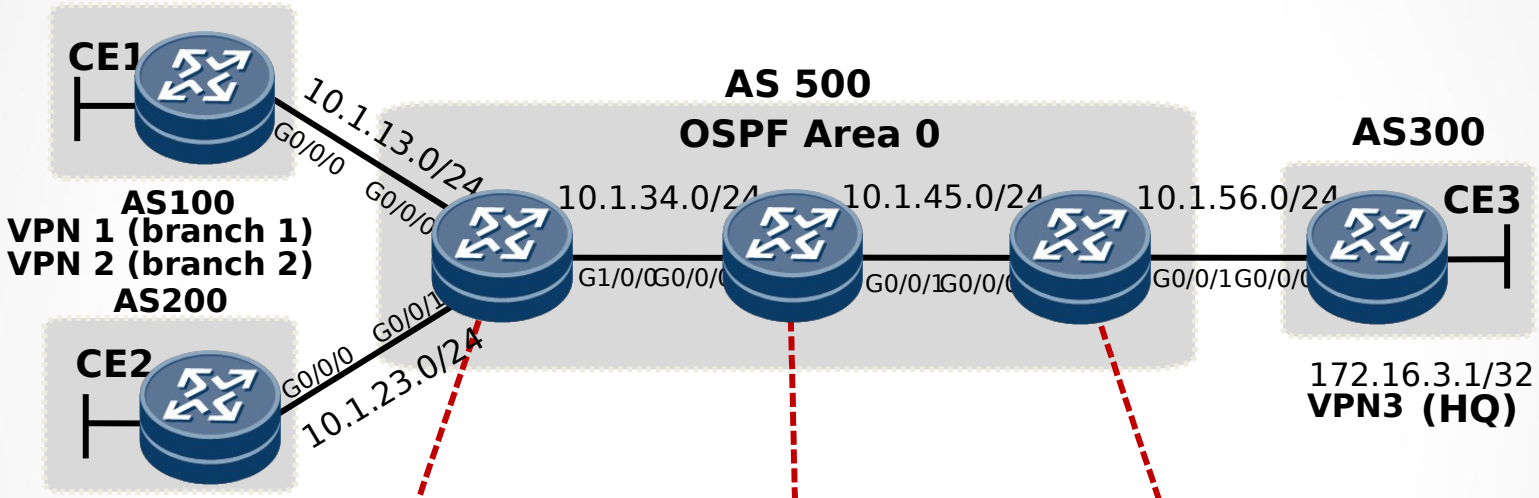


```
mpls lsr-id 3.3.3.3
mpls
mpls ldp
#
interface GigabitEthernet1/0/0
 mpls
 mpls ldp
```

```
mpls lsr-id 4.4.4.4
mpls
mpls ldp
#
interface GigabitEthernet0/0/0
 mpls
 mpls ldp
#
interface GigabitEthernet0/0/1
 mpls
 mpls ldp
```

```
mpls lsr-id 5.5.5.5
mpls
mpls ldp
#
interface GigabitEthernet0/0/0
 mpls
 mpls ldp
```

MPLS VPN Implementation

Monitoring MPLS VPN Operations

Monitoring MPLS VPN

Check the OSPF neighbor relationship on R1 router after the configurations are complete.

[R1]display ospf peer brief

```
        OSPF Process 1 with Router ID 1.1.1.1
                Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                       Neighbor id      State
 0.0.0.0          Serial1/0/0                     2.2.2.2          Full
 ----------------------------------------------------------------------------
 Total Peer(s):     1
```

Monitoring MPLS VPN (Cont.)

Check VPN instances on R1 router after the configurations are complete.

```
[R1]display ip vpn-instance verbose
 Total VPN-Instances configured      : 1
 Total IPv4 VPN-Instances configured : 1
 Total IPv6 VPN-Instances configured : 0

 VPN-Instance Name and ID : VPN1, 1
  Interfaces : Serial3/0/0
 Address family ipv4
  Create date : 2016/09/20 14:51:08
  Up time : 0 days, 00 hours, 09 minutes and 34 seconds
  Route Distinguisher : 1:1
  Export VPN Targets :  1:2
  Import VPN Targets :  1:2
  Label Policy : label per route
  Log Interval : 5
```

Monitoring MPLS VPN (Cont.)

Check the BGP neighbor relationship between R1 and R4 after the configurations are complete.

```
[R1]display bgp vpnv4 vpn-instance VPN1 peer
 BGP local router ID : 1.1.1.1
 Local AS number : 123
 VPN-Instance VPN1, Router ID 1.1.1.1:
 Total number of peers : 1                 Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.1.14.4       4          14        7        8     0 00:05:21 Established       0

[R4]display bgp peer
 BGP local router ID : 10.1.14.4
 Local AS number : 14
 Total number of peers : 1                 Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.1.14.1       4         123        4        6     0 00:02:56 Established       0
```

Monitoring MPLS VPN (Cont.)

Check VPN routes learned from customer networks in VPN routing table on R1

```
[R1]display ip routing-table vpn-instance VPN1
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: VPN1
         Destinations : 6        Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.1.14.0/24  Direct  0    0           D   10.1.14.1       Serial3/0/0
      10.1.14.1/32  Direct  0    0           D   127.0.0.1       Serial3/0/0
      10.1.14.4/32  Direct  0    0           D   10.1.14.4       Serial3/0/0
    10.1.14.255/32  Direct  0    0           D   127.0.0.1       Serial3/0/0
    192.168.1.0/24  EBGP    255  0           D   10.1.14.4       Serial3/0/0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

Monitoring MPLS VPN (Cont.)

Check the MP-BGP neighbor relationship on R1 after the configurations are complete.

```
[R1]display bgp vpnv4 all peer
 BGP local router ID : 1.1.1.1
 Local AS number : 123
 Total number of peers : 2                 Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  3.3.3.3         4         123        4        7     0 00:02:10 Established       0
```

Monitoring MPLS VPN (Cont.)

Check the MPLS LDP neighbor relationship on R1 after the configurations are complete.

```
[R1]display mpls ldp peer
 LDP Peer Information in Public network
 A '*' before a peer means the peer is being deleted.
 ------------------------------------------------------------------------------
 PeerID                 TransportAddress   DiscoverySource
 ------------------------------------------------------------------------------
 2.2.2.2:0              2.2.2.2            Serial1/0/0
 ------------------------------------------------------------------------------
 TOTAL: 1 Peer(s) Found.
```

Layer 2 MPLS VPN

Introduction

Layer 2 VPN Overview

- Many service providers offer Layer 2 transport services.
- Overlay VPNs were built using ATM or Frame Relay PVCs (Permanent Virtual Circuit).
- Building both a Layer 2 and a Layer 3 network is costly.
- Service providers want to unify VPN services and Internet services in a single network: an MPLS- based network.
- However, some customers still want Layer 2 connections (Ethernet VLANs or PVCs).

Layer 2 VPN Overview (Cont.)

- Benefits
  - Incorporates Layer 2 and Layer 3 services over a common infrastructure
  - Scalability: MPLS does not keep state information about virtual circuits inside core
  - Maintains support for existing services while migrating to MPLS
  - Customer sites independent of service provider backbone

Layer 2 VPN Overview (Cont.)

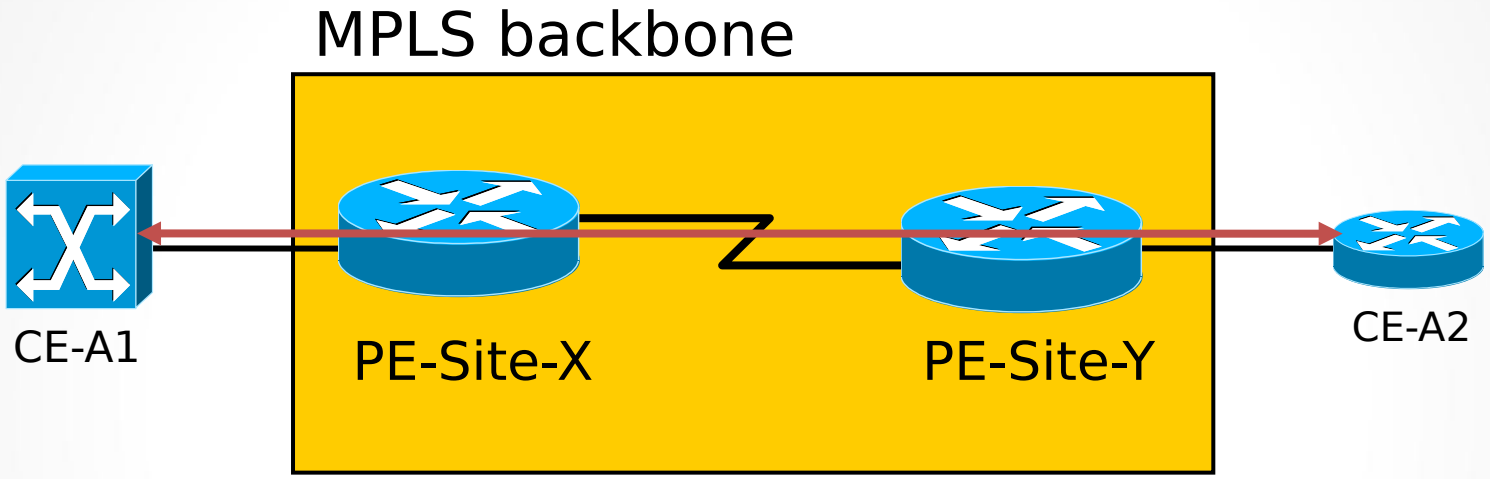

- Layer 2 VPN is not designed specifically for switch replacement.
- It is used for interconnecting and transporting a traffic POP (Point-of-Presence) across an MPLS core network.

Layer 2 VPN Types

- Virtual leased line (VLL)
- **Pseudo-Wire Emulation Edge to Edge (PWE3)**
- **Virtual Private LAN Service (VPLS)**

How Layer 2 VPN Works

- Frames are carried across an MPLS backbone in the following manner:
  - Ingress and egress interfaces are non-MPLS interfaces.
  - Ingress PE encapsulates frame into MPLS; egress PE decapsulates.
  - Label stack of two labels is used.
    - Topmost label (“tunnel label”) used for LSP PE to PE.
    - Second label (“VC label”) identifies outgoing interface in the egress PE.
  - **LDP has been extended** to carry virtual circuit forwarding equivalence class (VC FEC).
    - A directed (multihop) LDP session is used from PE to PE.

Layer 2 VPN Types

Point-to-Point. Applicable to any type of channel layer protocols

It is based on the concept of PW — PseudoWire. Connecting two nodes to each other. The provider's network can be considered as one virtual cable.

The general name of the service: VPWS — Virtual Private Wire Service.

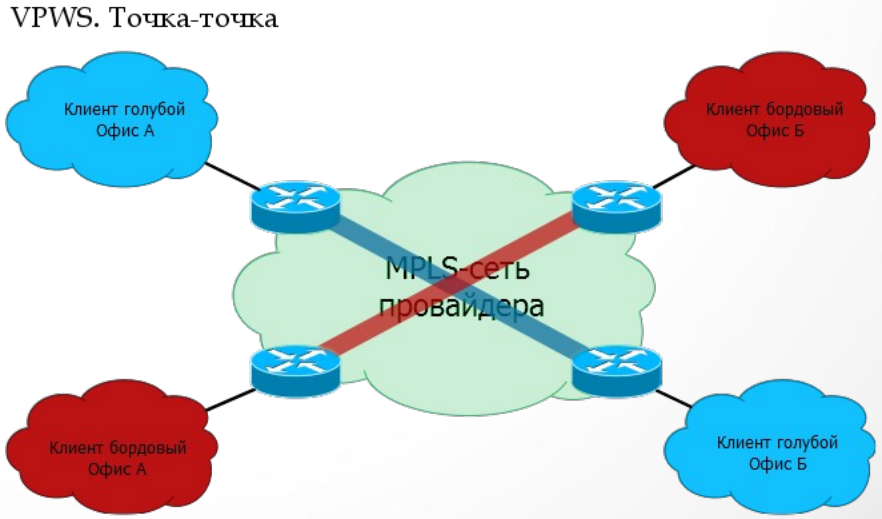

Layer 2 VPN Types

Point-to-Multipoint. The mode is only for the Ethernet network. The client may have several connection points/branches, and they must transmit data to each other, both to one particular branch and to all at once. Can be considered as a virtual Ethernet switch

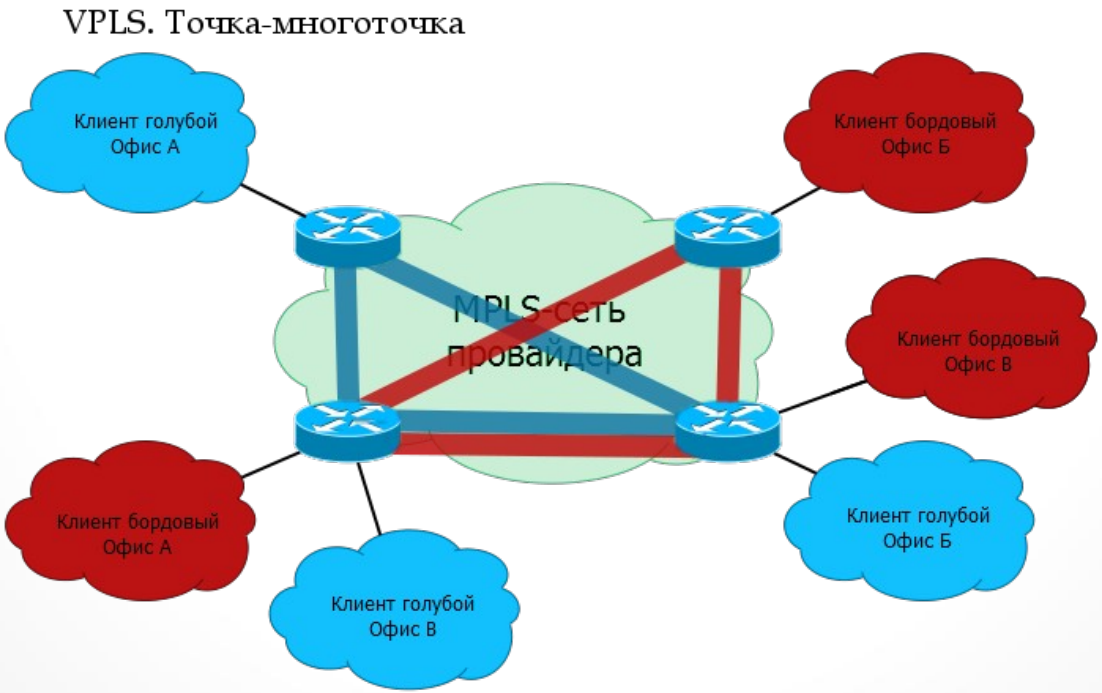

Terminology

PE — Provider Edge — edge routers of the provider's MPLS network, to which client devices (CE) are connected.
CE - Customer Edge is the client's equipment that directly connects to the provider's routers (PE).
AC — Attached Circuit is the interface on the PE for connecting the client.
VC — Virtual Circuit is a virtual unidirectional connection over a shared network that simulates the original environment for the client. Connects AC interfaces of different PE. Together they make up the channel: AC→VC→AC.
PW — PseudoWire — a virtual bidirectional data link between two PE — consists of a pair of unidirectional VC.

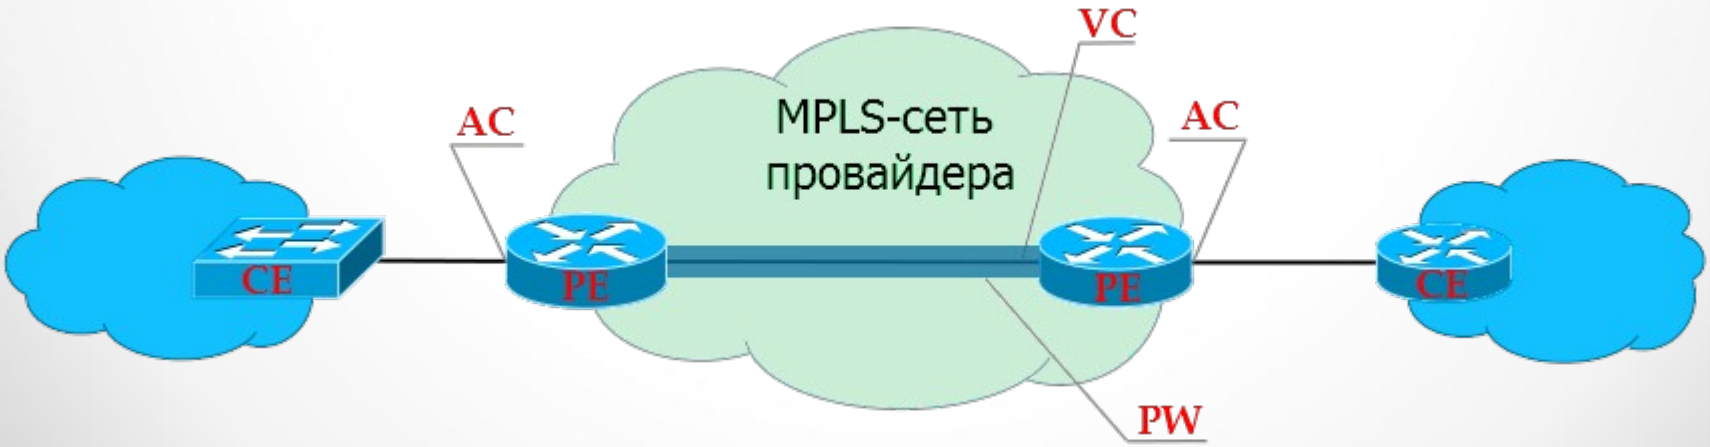

VPWS Data Plane

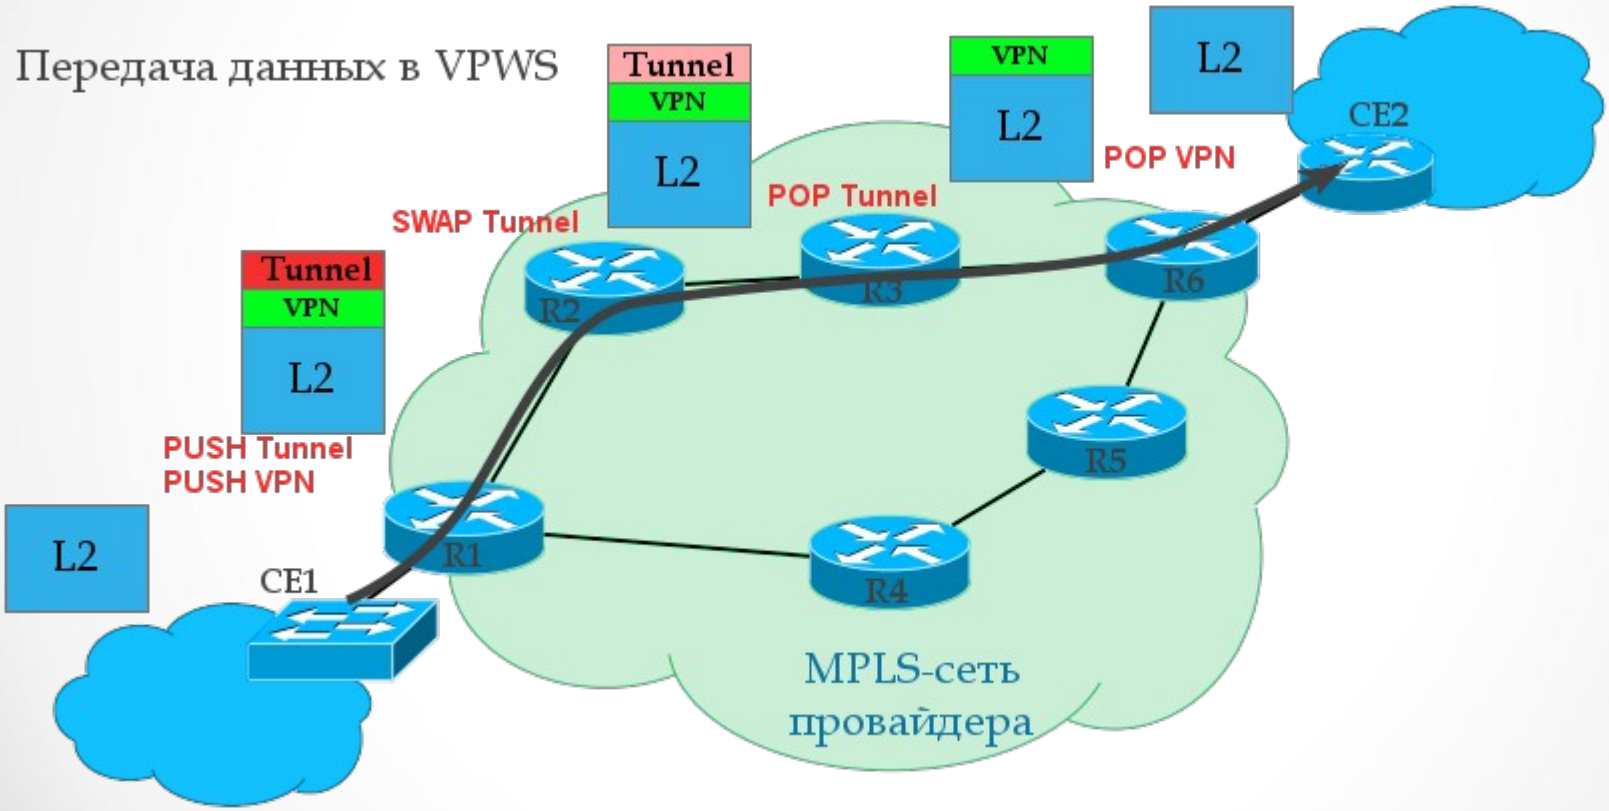

VPWS Data Plane

- R1 receives a frame from the client device (CE1), it already knows the transport label and the output interface on the way to R6 (LSPs are already built. UDP or RSVP-TE protocols are used).
- The interface (AC) to which CE1 is connected must be linked to the client ID — VC ID (analogous to VRF in L3VPN). Based on this information, R1 gives the frame a **service label** that will remain unchanged until the end of the path.
- R1 inserts the **transport label** into the MPLS label stack. This will be an external transport label.
- The MPLS packet is transmitted over the operator's network via P-routers.
- The transport label is removed on the penultimate router — PHP occurs.
- The PE output router (R6) analyzes the service tag and determines which interface to send the unpacked frame to.

VPWS Control Plane

- **LDP** — for each Loopback address of each MPLS router will distribute transport labels over the network.
- In order for the **VC** to switch to the **UP state**, there must be both LSPs — forward and reverse.
- **Targeted LDP** - distribution of service labels between remote routers .
- You need to manually configure a remote LDP session on edge routers. It is not tied to a VPN. The same session can be used to exchange labels with any number of VPNs.
- IP connectivity is sufficient for **tLDP**.
- When AC interfaces with the same VC-ID appear on the edge routers, LDP will help them communicate labels to each other.

tLDP vs LDP

LDP	tLDP
Only directly connected routers can be neighbors	Neighbors can be any routers in the network with which there is IP connectivity
Search for all possible neighbors	The neighbors are already defined by the configuration
Broadcast distribution of Discovery messages	Targeted sending of Discovery messages to specific neighbors
The IP address is used as the FEC	The VC ID is usually used as the FEC

How Layer 2 VPN (Cont.)

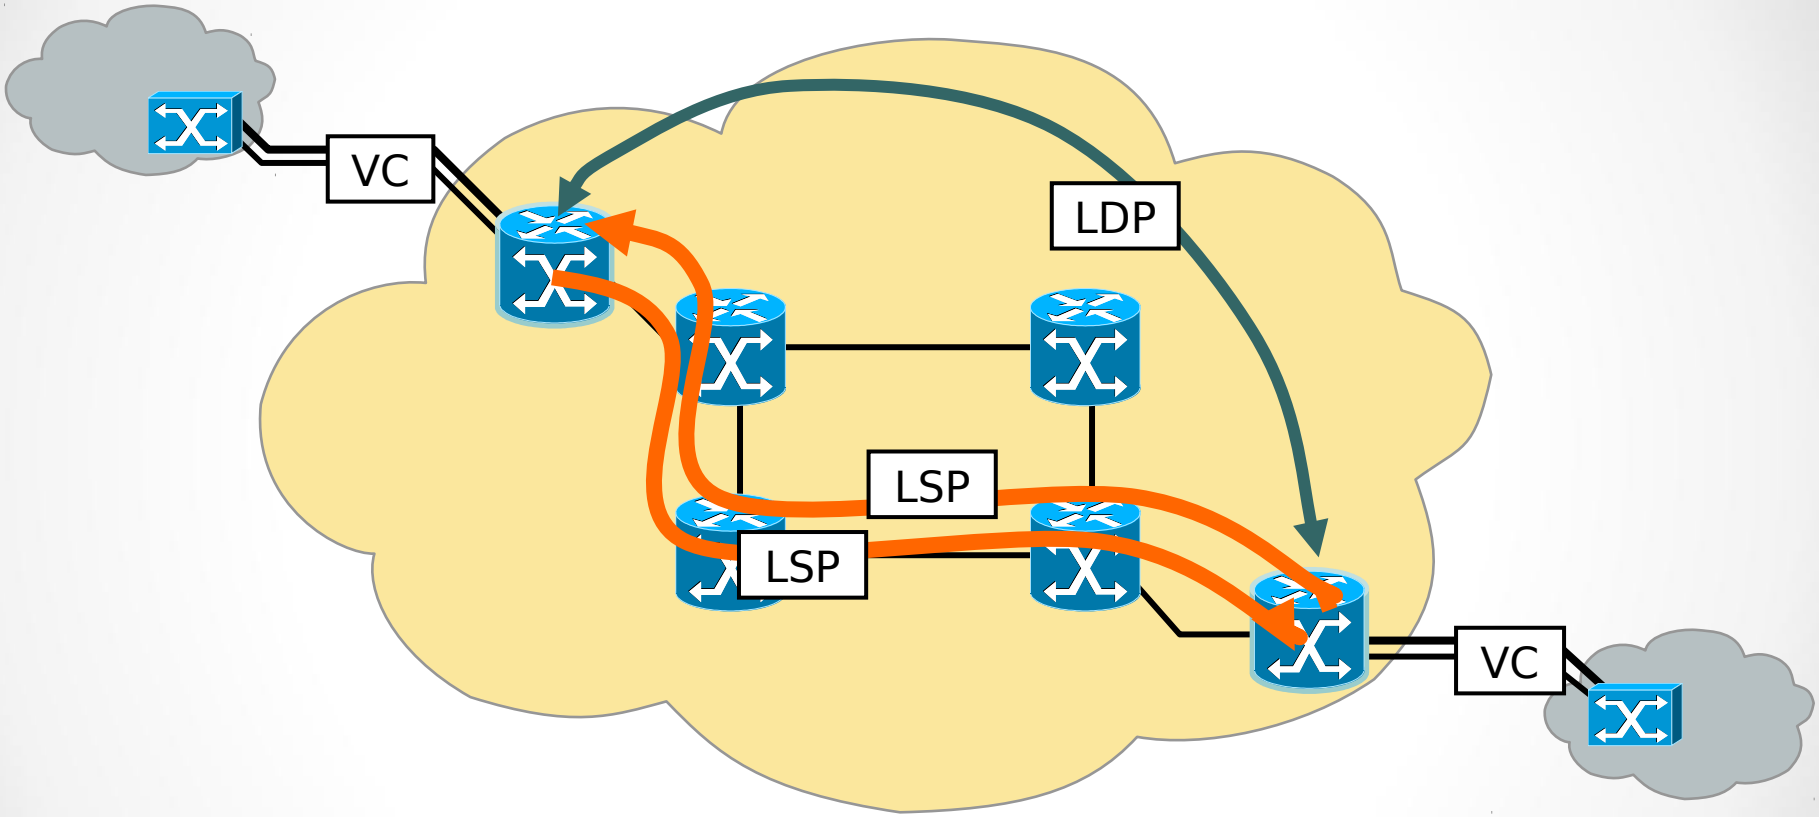

- The IGP and the LDP between directly connected LSRs establish one LSP in each direction.

- **A directed LDP session between PE routers is established.**

How Layer 2 VPN (Cont.)

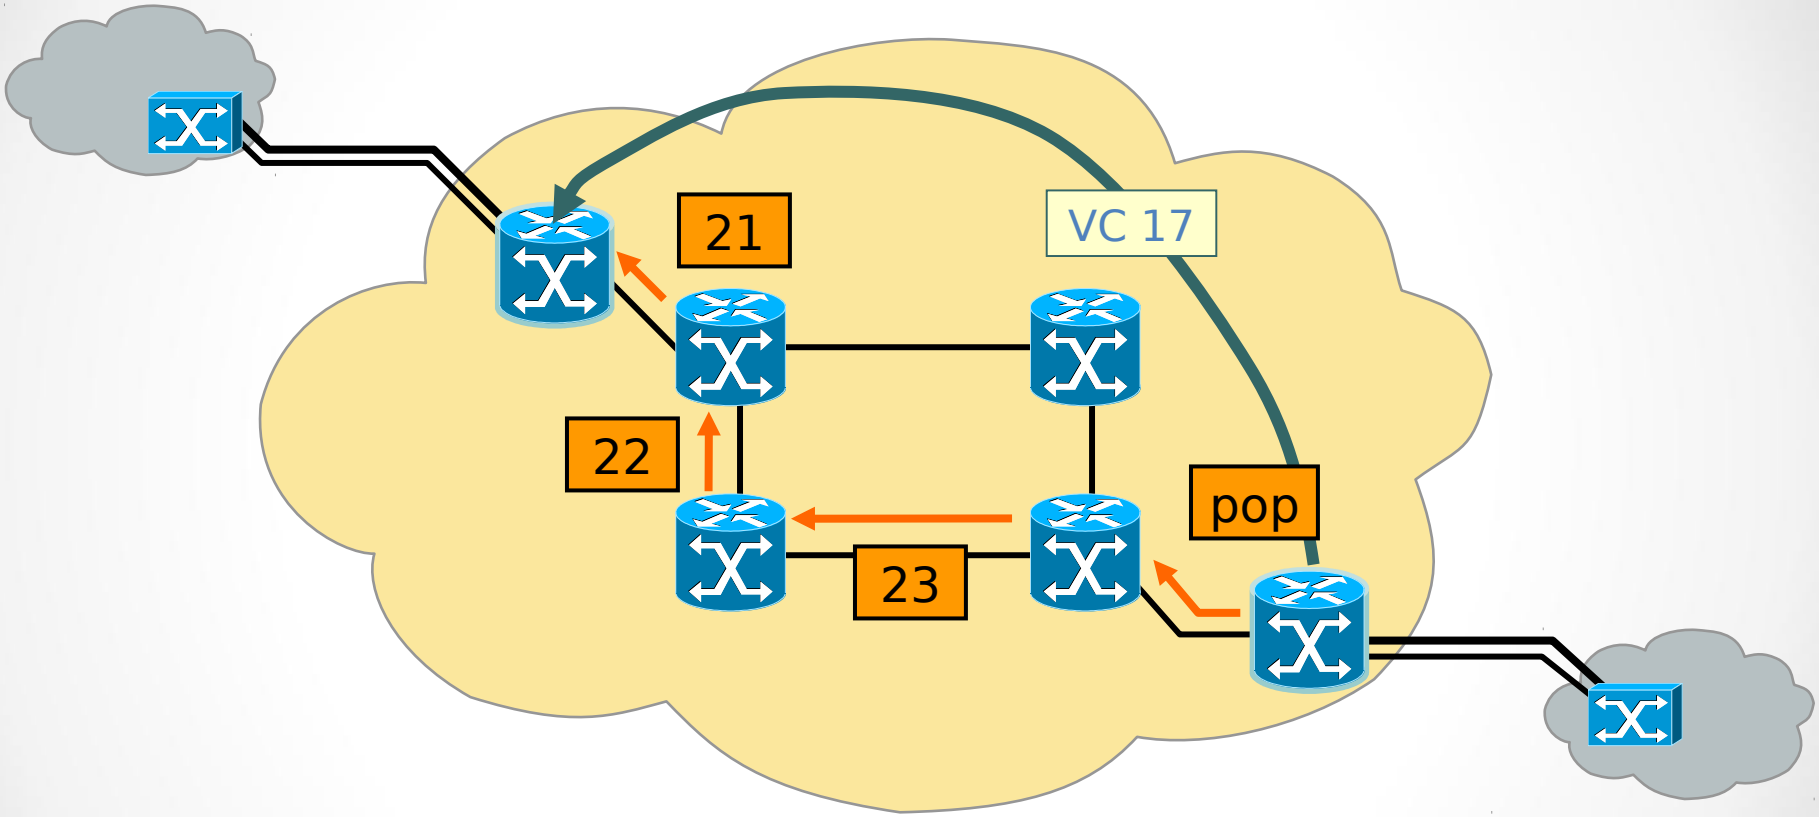

- **LDP between directly connected LSRs generates an LSP.**
  - The egress PE allocates VC label 17.
  - The directed LDP session between PE routers propagates the VC label.

How Layer 2 VPN Works

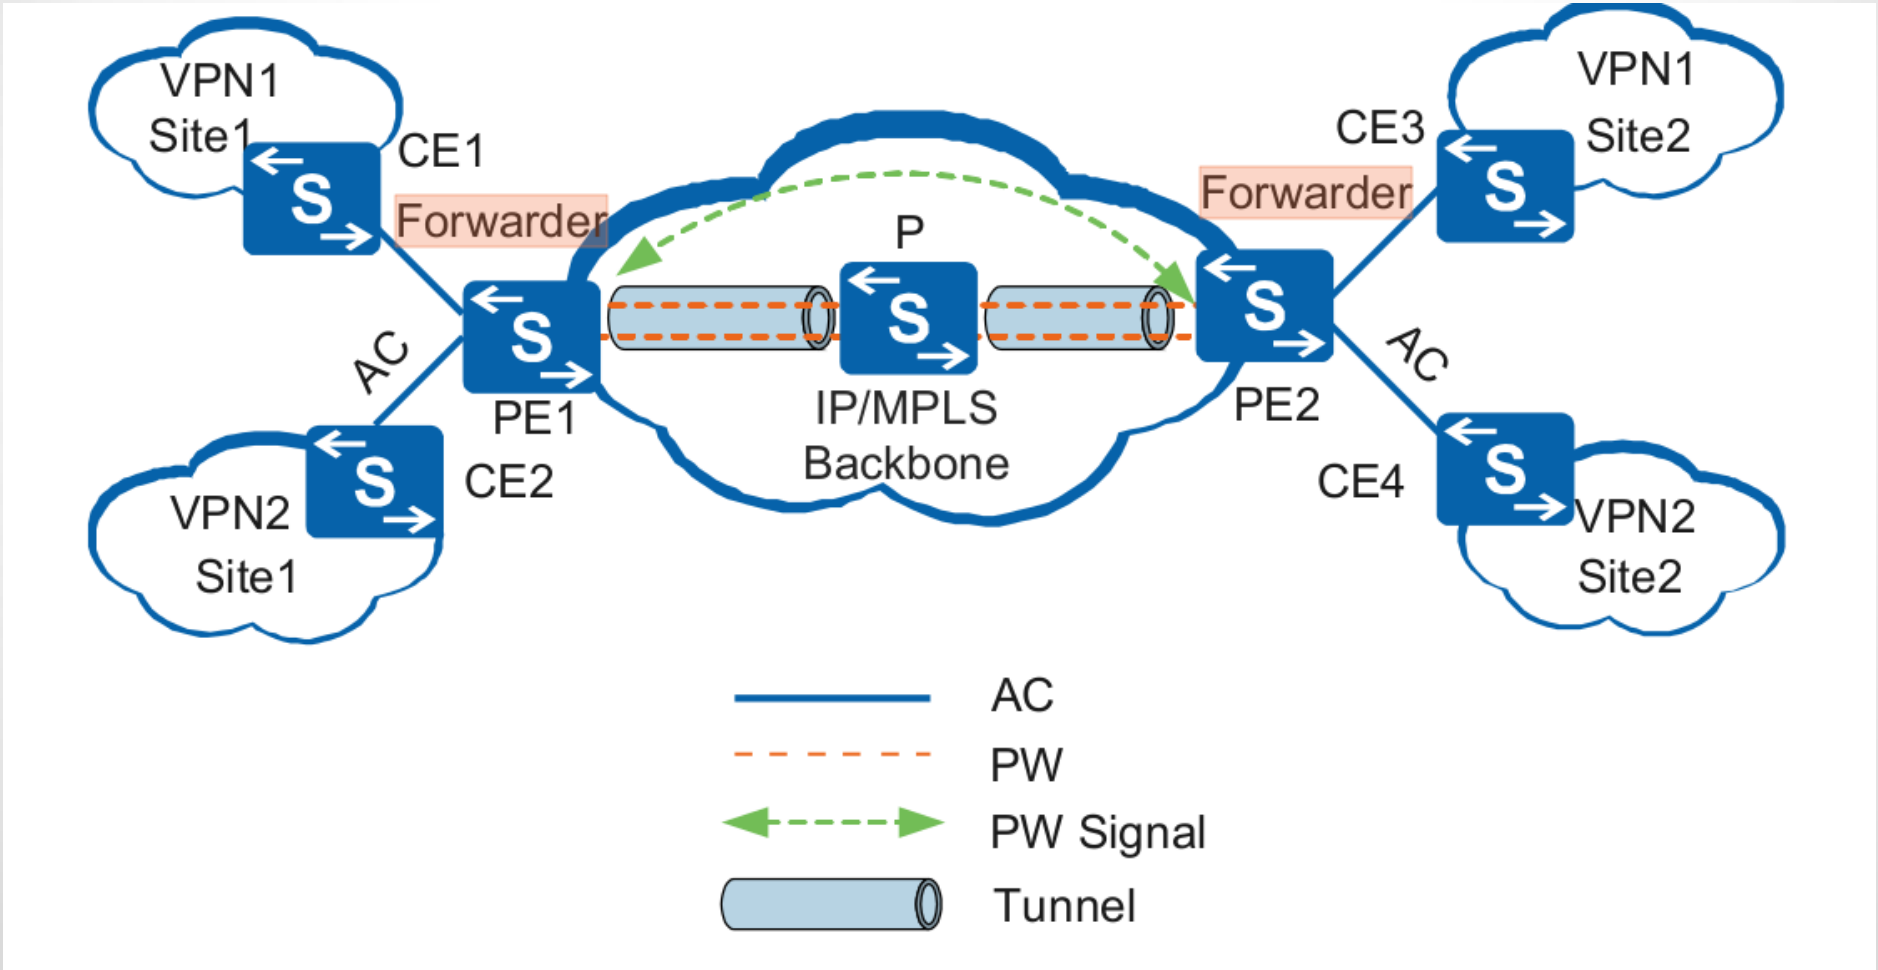

How Layer 2 VPN Works

Forwarder (FWRD) - A PE subsystem that selects the PW to use in order to transmit a payload received on an AC.
Similar to a **forwarding table**. After a PE receives packets from an AC, the forwarder of the PE selects a PW to forward these packets.

As discussed in RFC3985, a pseudowire can be thought of as connecting two "**forwarders**".

Every Forwarder in a PE must be associated with an **Attachment Identifier (AI)**. The AI must be unique in the context of the PE router in which the Forwarder resides. The combination <**PE router IP address, AI**> must be globally unique.

How Layer 2 VPN Works

- PE interprets the **Label Mapping message** (LDP) as a request to set up a PW whose endpoint (at PE) is the Forwarder identified by the **TAI (Target AI).**

- **From the perspective of the signaling protocol, exactly how PE maps AIs to Forwarders is a local matter.** In some VPWS models, the TAI might, for example, be a string that identifies a particular Attachment Circuit (AC)".

- **If PE cannot map the TAI to one of its Forwarders, then PE2 sends a Label Release message** to the sender of Label Mapping message PE, with a Status Code of "Unassigned/Unrecognized TAI", and the processing of the Label Mapping message is complete.

Virtual Private LAN Service

- **VPLS-Virtual Private LAN Service**. It can be considered as a switch emulation.
  - The task of the provider's transport network is to ensure the correct switching of frames, the study of MAC addresses is used.
- **VPLS domain** — isolated virtual L2 network. Two different clients — two different VPLS domains.
- **VSI-Virtual Switching Instance**. A virtual switch within a single node. For each client (or service) he's his own. Traffic from one VSI cannot be transferred to another VSI.
  - Analog of VRF/ VPN-instance in L3VPN.
- **VE-VPLS Edge** - PE node, a member of the VPLS domain.
- **Tunnel** - A connection between a local PE and a remote PE used to transparently transmit data between PEs. A tunnel can carry multiple PWs. Tunnel types can be Label Switched Path (LSP) or MPLS Traffic Engineering

VPLS Control Plane

- Detection of the PE where the clients of this VSI are connected:
  - manual setup (**Martini draft**)
  - automatic detection (**Compella draft**).
- VPLS is a point-to-point PW group.

- The **decision** to transfer the frame is made by **Ingress PE** (selects the desired PW)

Virtual Private LAN Service

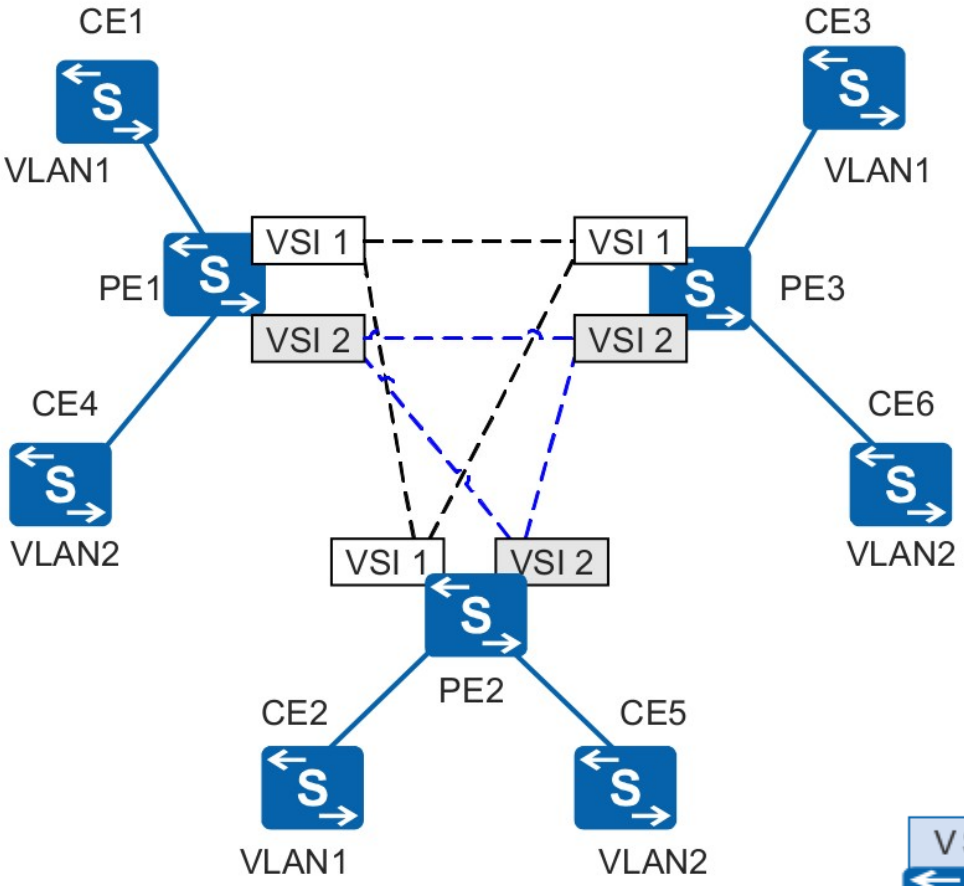

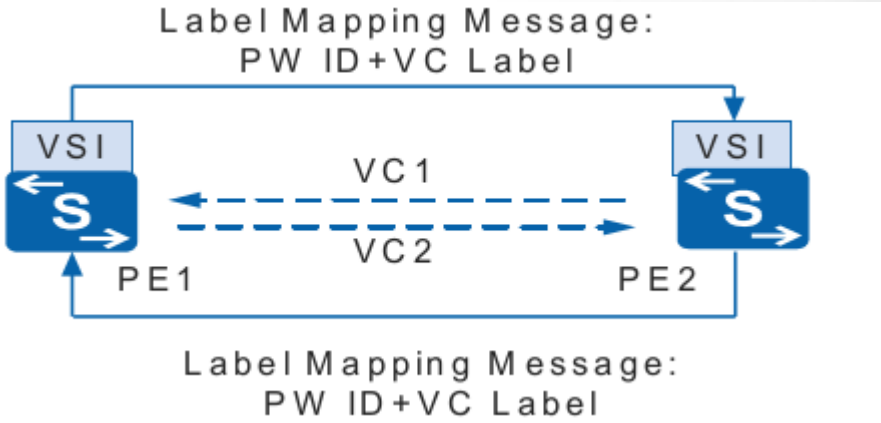

Packet Encapsulation

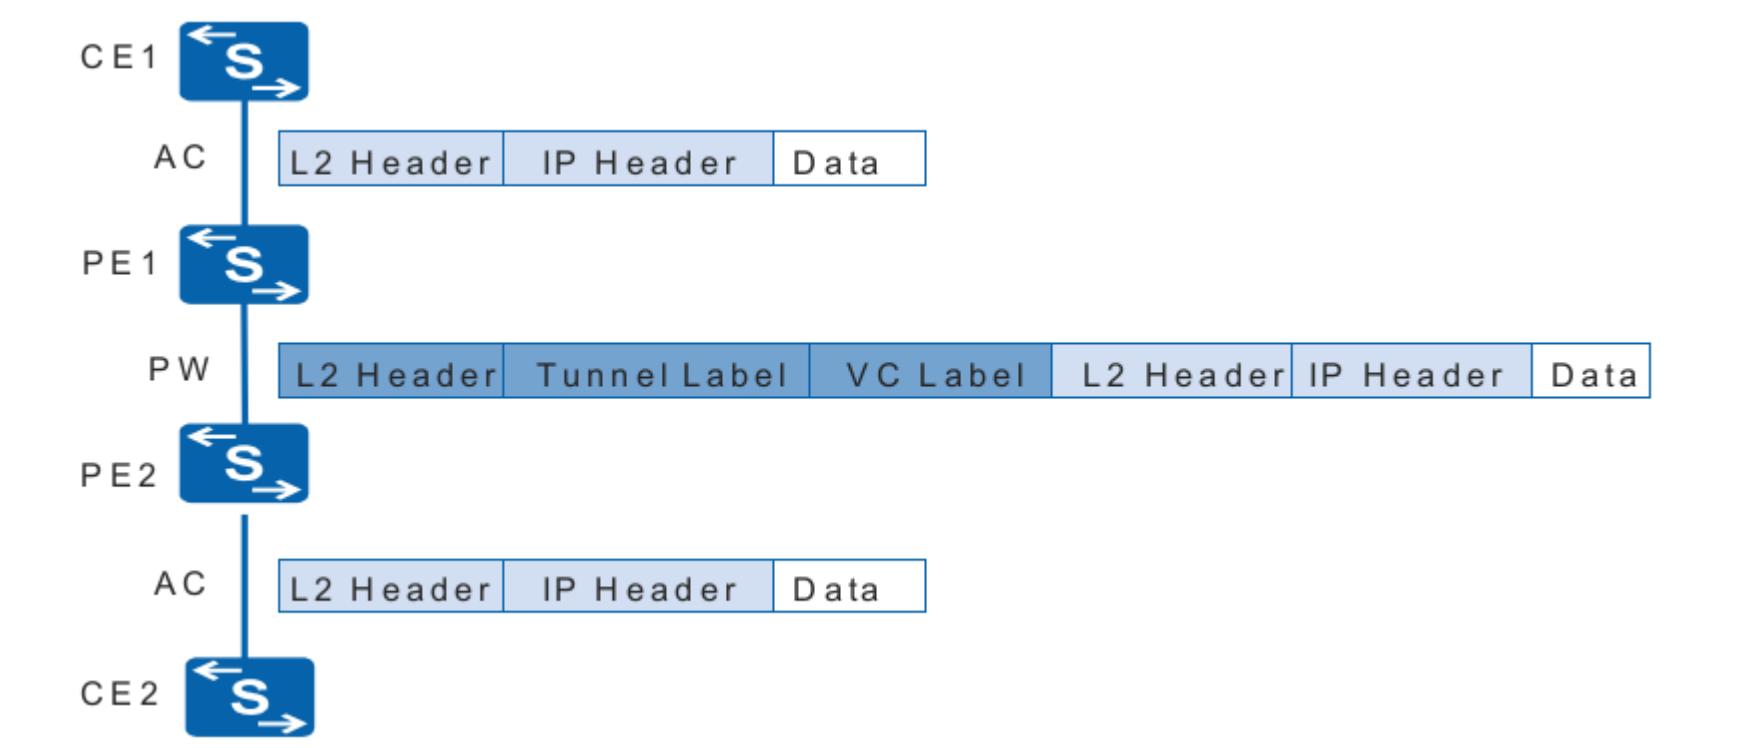

Virtual Private LAN Service

- MAC address learning and flooding on a PE.
- PC1 and PC2 both belong to VLAN10.
- PC1 pings IP address 10.1.1.2.
- PC1 does not know the MAC address corresponding to this IP address and advertises an Address Resolution Protocol (ARP) Request packet.

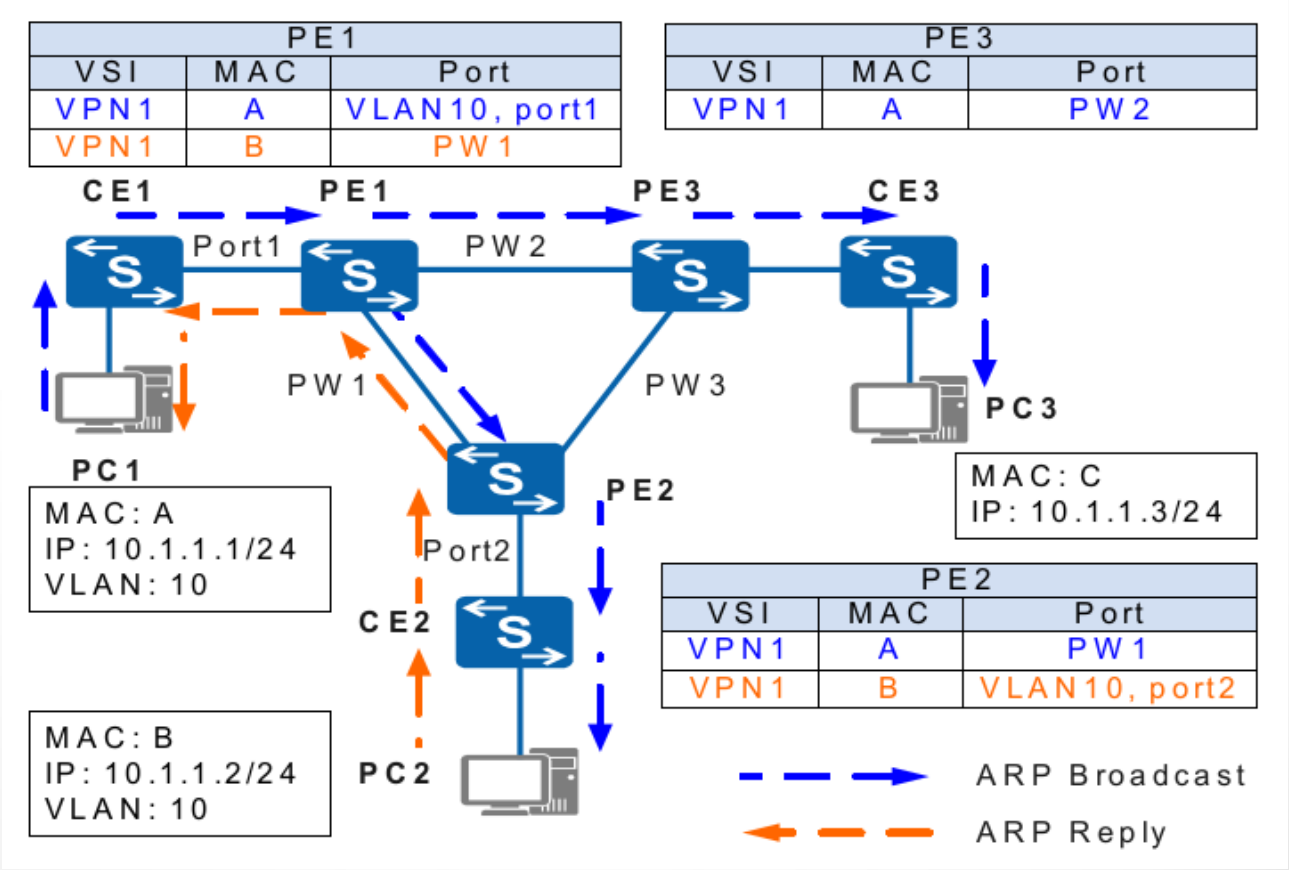

Virtual Private LAN Service

- Dynamic MAC addresses need to be updated and relearned. The VPLS draft defines a **MAC Withdraw message** with an optional MAC type-length-value (TLV) to remove or relearn the MAC address list.
- MAC Withdraw messages enable devices to quickly delete matching MAC addresses when network topology changes.

Virtual Private LAN Service

- Since VPLS emulates a LAN, full mesh connectivity is required.
- Two methods for full mesh establishment for VPLS:
  - Border Gateway Protocol (BGP)
  - Label Distribution Protocol (LDP). The PWs constitute the "data plane", whereby PEs send customer VPN/VPLS traffic to other PEs.
- "**Control plane**"
  - **Auto-discovery** refers to the process of finding other PE routers participating in the same VPN or VPLS.
    - Each PE is configured to participate in a given VPLS. The PE, through the use of BGP, simultaneously discovers all other PEs in the same VPLS, and establishes a full mesh of pseudowires to those PEs.
  - **Signalling** is the process of establishing pseudowires (PW).

- With **LDP**, each PE router must be configured to participate in a given VPLS, and, in addition, be given the addresses of other PEs participating in the same VPLS. A full mesh of LDP sessions is then established between these PEs. LDP is then used to create an equivalent mesh of PWs between those PEs.

VPLS Data Plane

1. The PE router reads the Ethernet frame header and checks the sender's MAC address.
   a) If this address does not exist yet, it writes the corresponding MAC port to the table and proceeds to step 2.
2. The PE router checks the MAC address of the recipient.
   a) If it is present in the MAC address table of this VSI, the PE is looking for an output interface for a frame with a given MAC. Physical interface or a PW
   b) If it is a PW, then it adds the appropriate — service label. It will remain unchanged until the end of the path.
   c) Knowing the IP address of the remote PE, the local PE extracts the transport label from the label table and puts it on top of the stack — it will change on each P-router.
3. If the MAC address is unknown (or if it is a broadcast frame), then the PE must broadcast the frame to all the PE of this VSI
   a) The local PE makes a list of all the remote PE of this VSI, and, having created copies of this frame, inserts service labels into them — each is assigned its own.
   b) A transport label is added to each copy of the frame (also its own for each PE)

VPLS Data Plane

4. After receiving the frame and removing the labels (already determined VSI) the remote PE acts as a switch:

   a) If the **MAC address of the source is NOT known** to him, he enters it into the table. PW to Ingress PE will be specified as the input interface

   b) If the **destination MAC address is known** to him, he sends a frame without MPLS headers to the port for which it is studied.

   c) If this **destination MAC address is NOT known** to him, broadcasting is performed on all AC (Attachment Circuit) ports of this VSI. The PE will not send this frame to the PW of this VSI, because all other PE have already received a copy of this frame from the input PE. That is, the **Split Horizon** rule is applied.

VPLS Kompella Mode

- **RFC-4761**
- The VPN tag distribution protocol is **MP-BGP**. There is an Address Family in BGP for VPLS: L2VPN AFI (25) and VPLS SAFI (65)
- A scheme similar to VPN L3 is used. **Route Target** is used to auto-detect neighbors. **Route-reflectors** - for the implementation of a «full mesh» topology

- Route Target (on thee base of BGP Extended Community), is the main mark of belonging to a particular VSI.
- If the RT of the received announcement matches with the one configured in the VSI, then this VSI accepts information from the announcement.
- To announce the L2 prefix to all the PE of a given VSI, a BGP Update is sent to all configured neighbors (just like in L3VPN, where vpnv4 prefixes are sent to all PE).

VPLS Kompella Mode

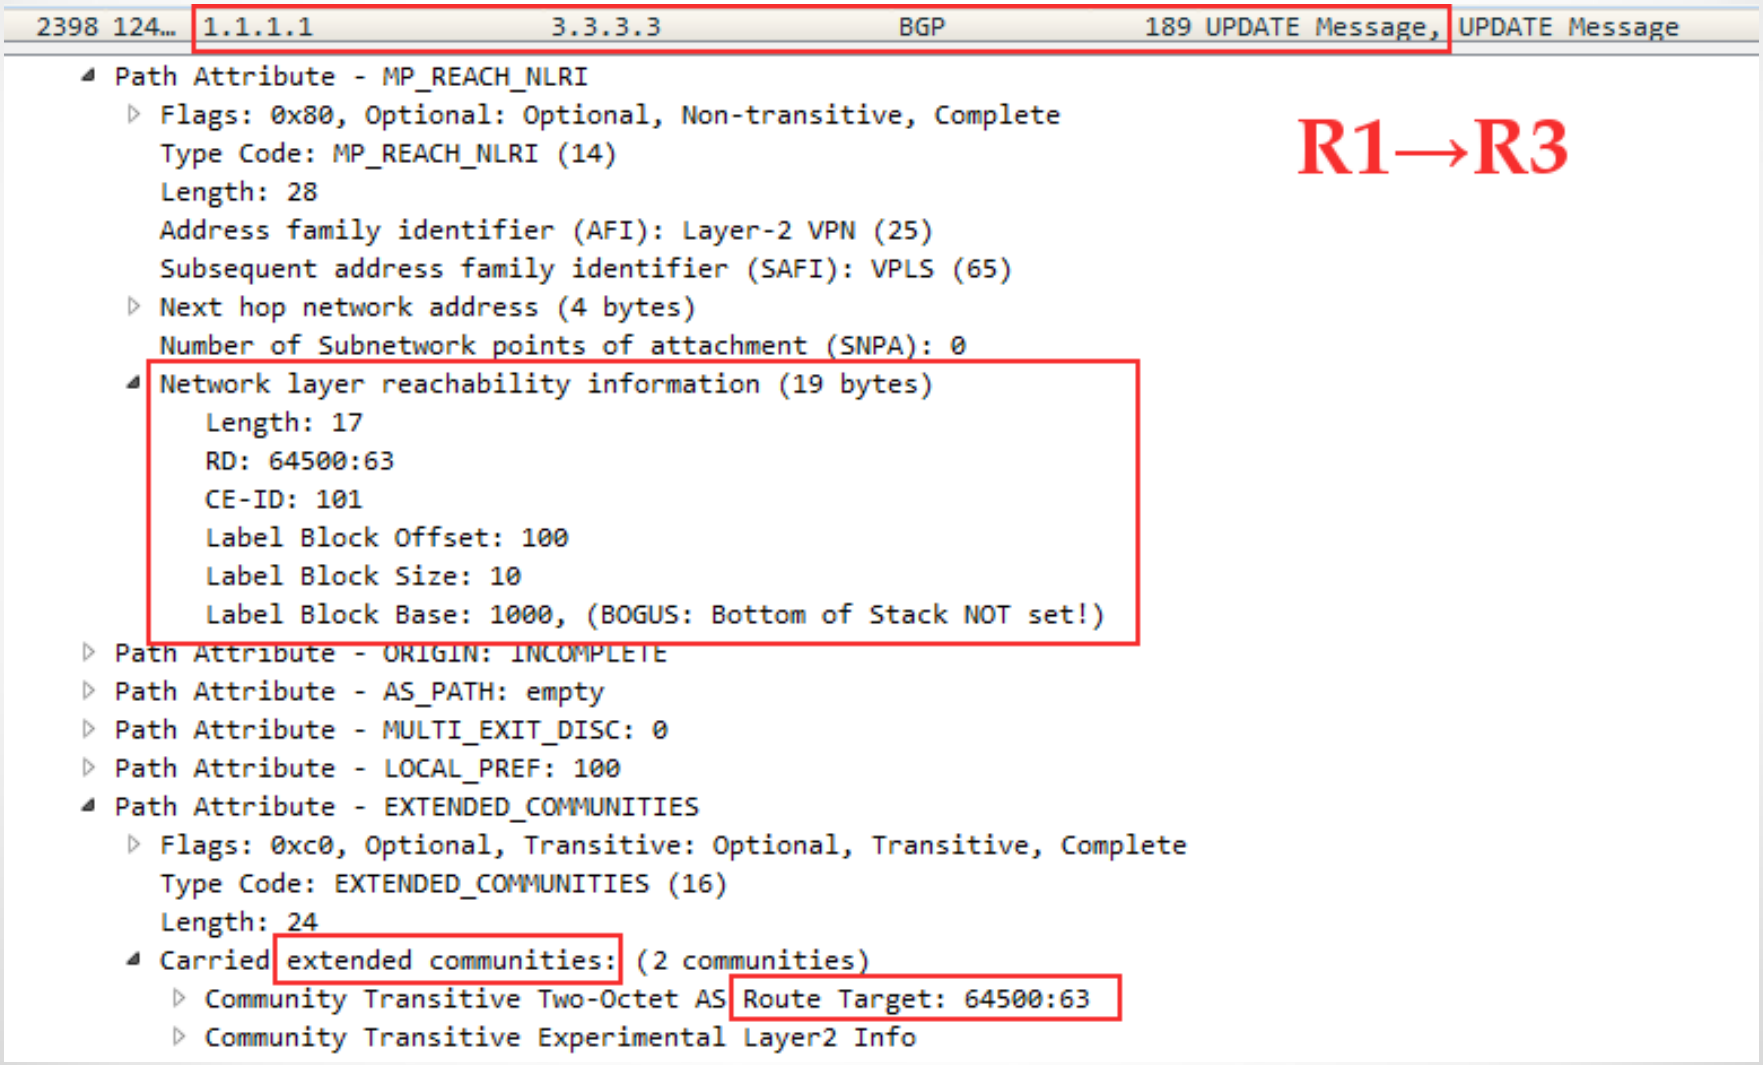

VPLS Kompella Mode

- The VPLS Kompella mode uses the **label block** mechanism. One PE does not tell the exact value of the label — it gives information for its calculation.

- In each VSI, the **RD and RT** are configured. RD allows you to separate the information of different VSI during transmission. RT allows the receiving router to determine which VSI to transmit information to.
- In BGP, a new **Address Family L2 VPN VPLS** is configured, within which the neighborhood with all PE established.
- You must create a full mesh topology. But the **Route-Reflector** mechanism allows you to bypass this requirement by establishing a neighbor relations with only one RR (or several in the case of an RR cluster)?

PW Configuration example

- **Creating a VSI and Configuring LDP Signaling**

- Step 1 Run: **system-view**
- Step 2 Run: **vsi vsi-name [static]**
- Step 3 Run: **encapsulation { ethernet | vlan }**
- Step 4 Run: **pwsignal ldp**
  - The VSI-LDP view is displayed.

- Step 5 Run: **vsi-id** vsi-id
  - The two ends of the VSI must agree on the same VSI ID.
  - The VSI exists only on the PE. One PE can have multiple VSIs. One VPLS on a PE has only one VSI.

- Step 6 Run: **peer peer-address [negotiation-vc-id vc-id] [tnl-policy policy-name] [secondary] [ignore-standby-state]**
  - **peer-address** specifies the IP address of the peer, usually referring to the LSR ID of the peer.

- Step 7 Run: **commit**