MPLS-VPN

Overview

- Traditional Router-Based Networks
- Virtual Private Networks
- VPN Terminology
- MPLS VPN Architecture
- MPLS VPN Routing
- MPLS VPN Label Propagation

Traditional Router-Based Networks

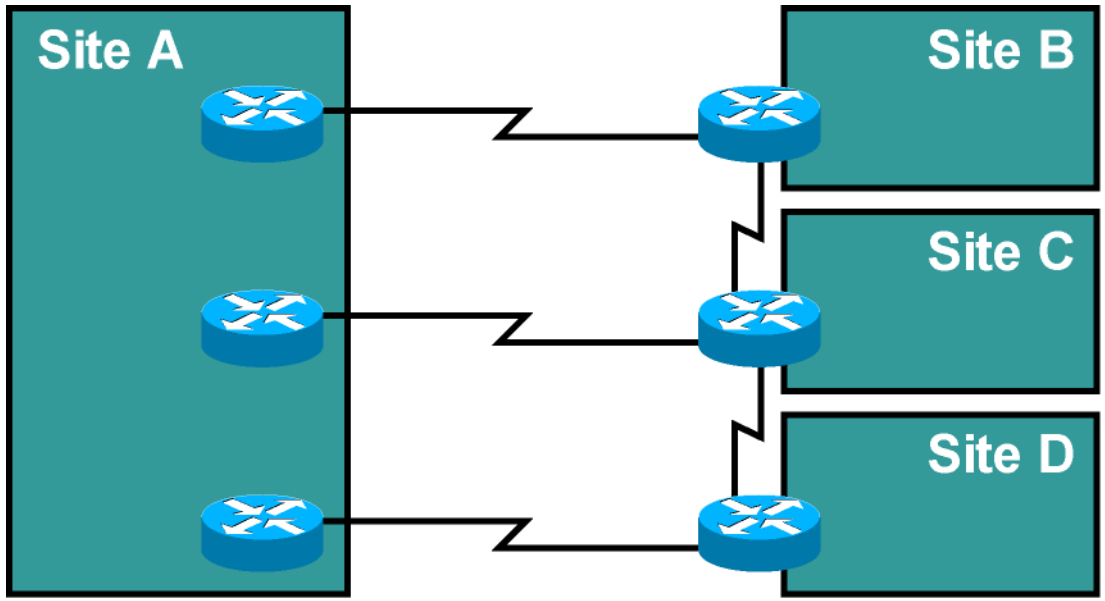

- **Traditional router-based networks connect customer sites through routers connected via dedicated point-to-point links.**

Virtual Private Networks

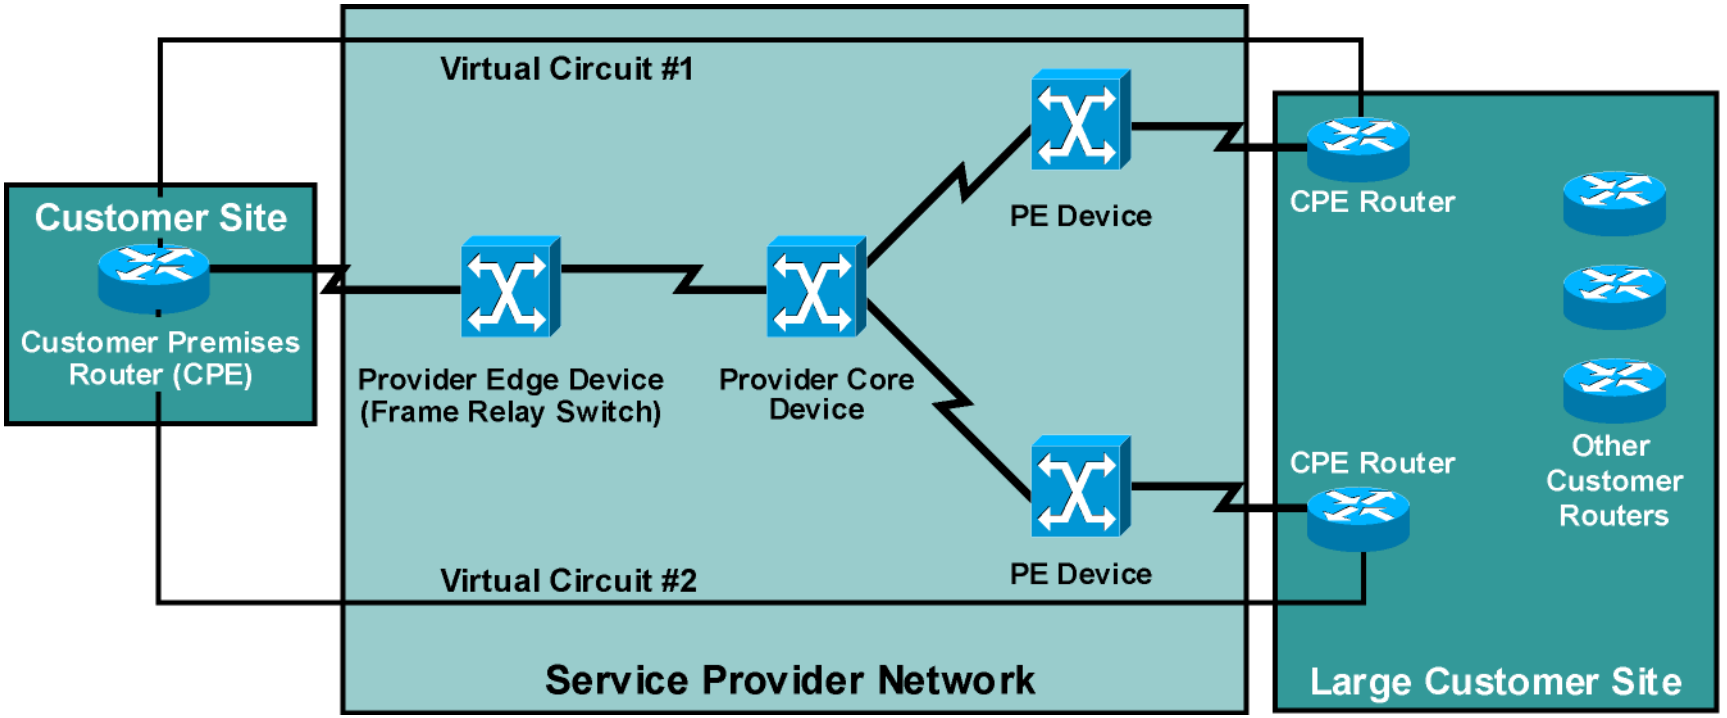

- VPNs replace dedicated point-to-point links with emulated point-to-point links sharing common infrastructure.
- Customers use VPNs primarily to reduce their operational costs.

VPN Terminology

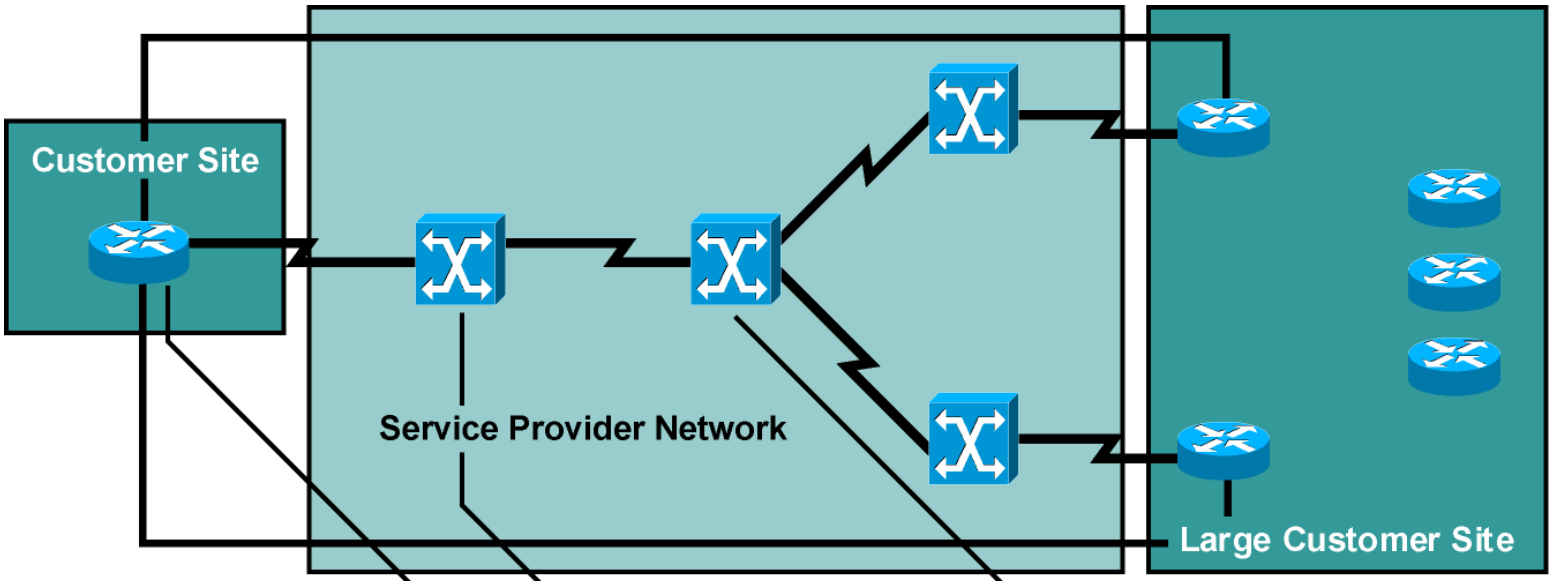

Provider (P) device: The device in the P-network with no customer connectivity

Provider edge (PE) device: The device in the P-network to which the CE devices are connected

Customer edge (CE) device: The device in the C-network that links to the P-network; also called customer premises equipment (CPE)

020G_042

VPN Implementation Technologies

VPN services can be 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

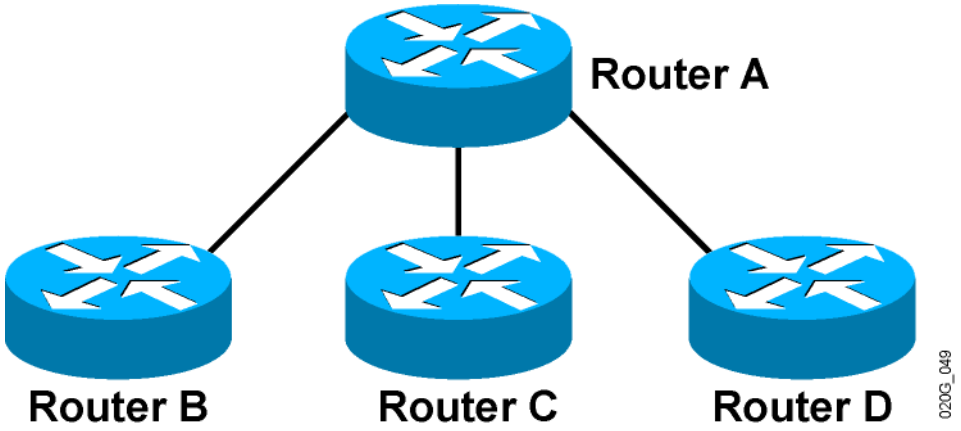

- Service provider infrastructure appears as point-to-point links to customer routes.
- Routing protocols run directly between customer routers.
- Service provider does not see customer routes and is responsible only for providing point-to-point transport of customer data.

Peer-to-Peer VPNs

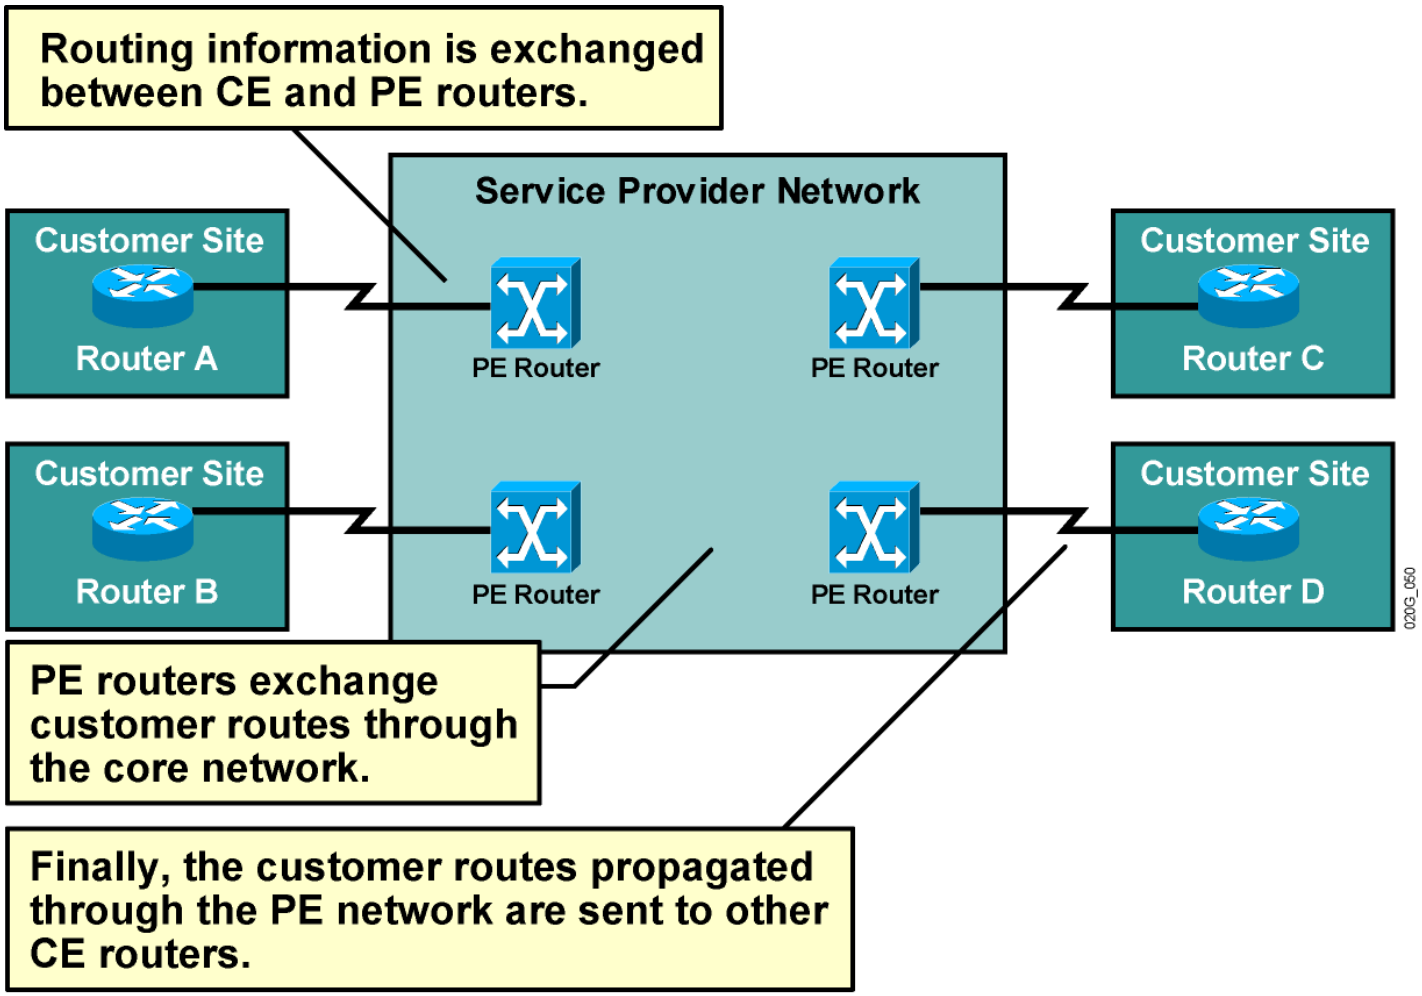

Benefits of VPN Implementations

- **Overlay VPN:**
  - Well-known and is 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 the sites are provisioned, not the links between them.

Drawbacks of VPN Implementations

- **Overlay VPN:**
  - Implementing optimum routing requires 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:**
  - Service provider participates in customer routing.
  - Service provider becomes responsible for customer convergence.
  - PE routers carry all routes from all customers.
  - Service provider needs detailed IP routing knowledge.

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

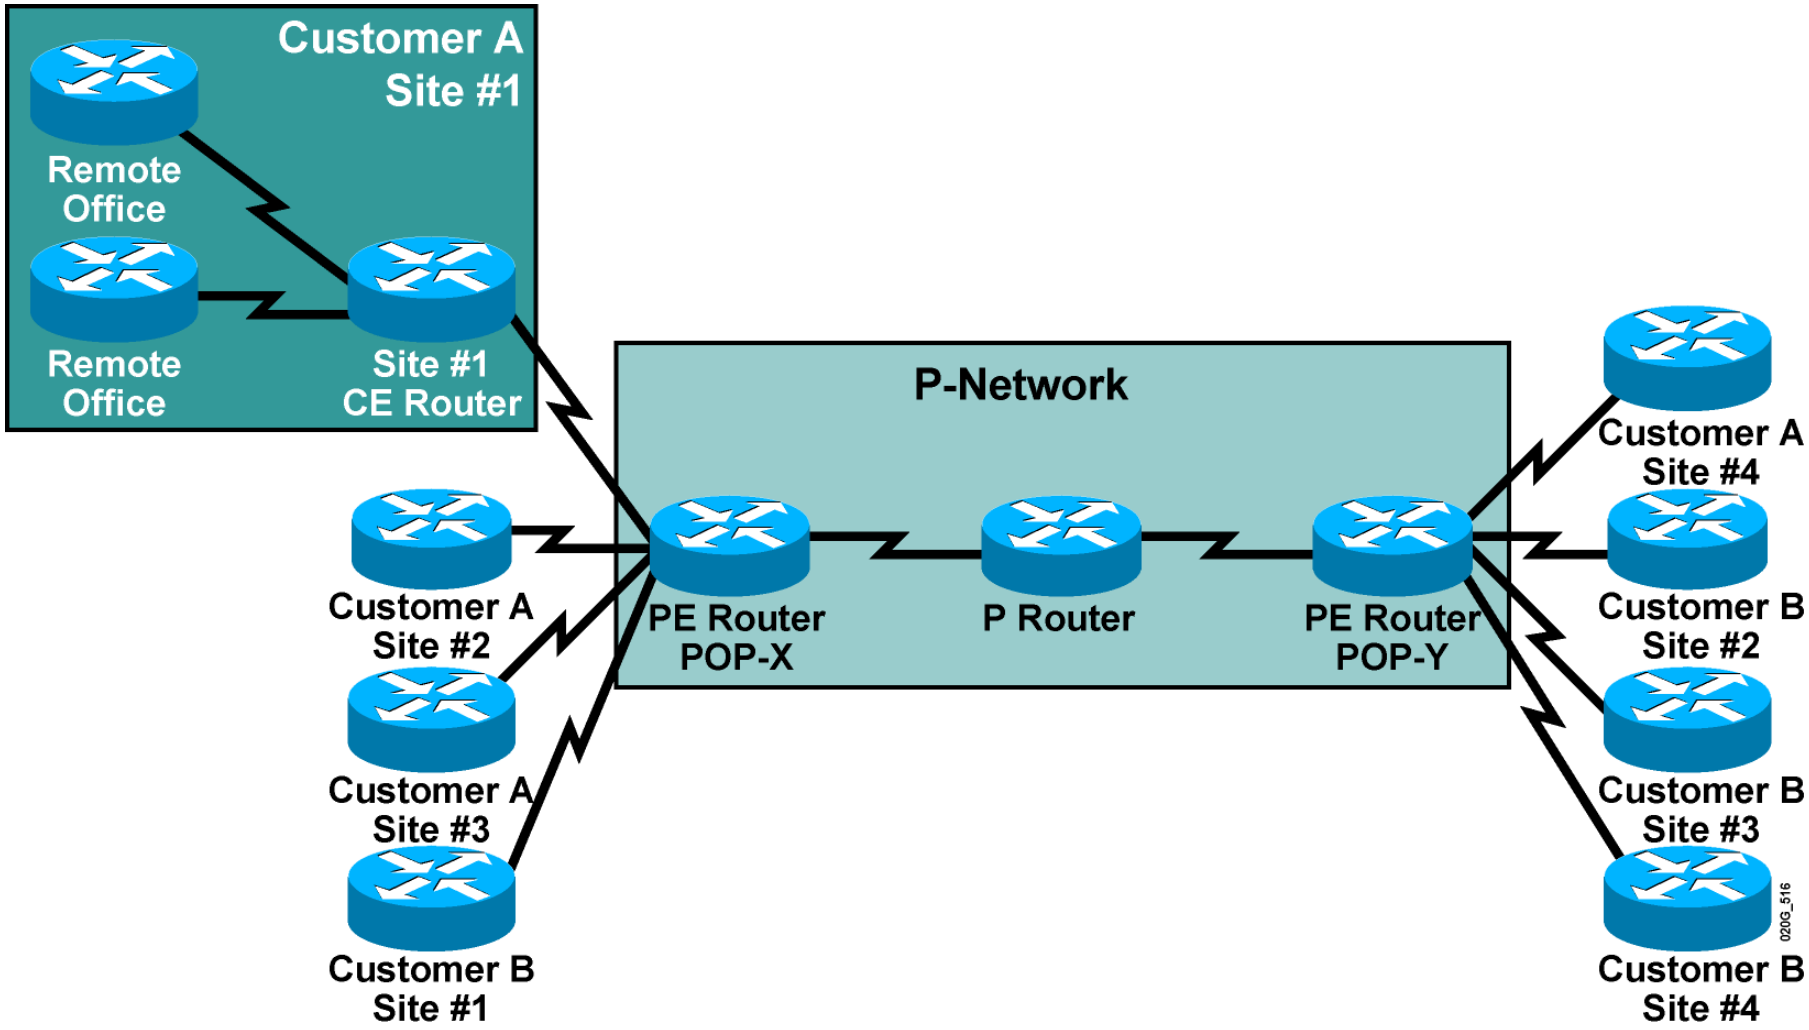

PE Router Architecture

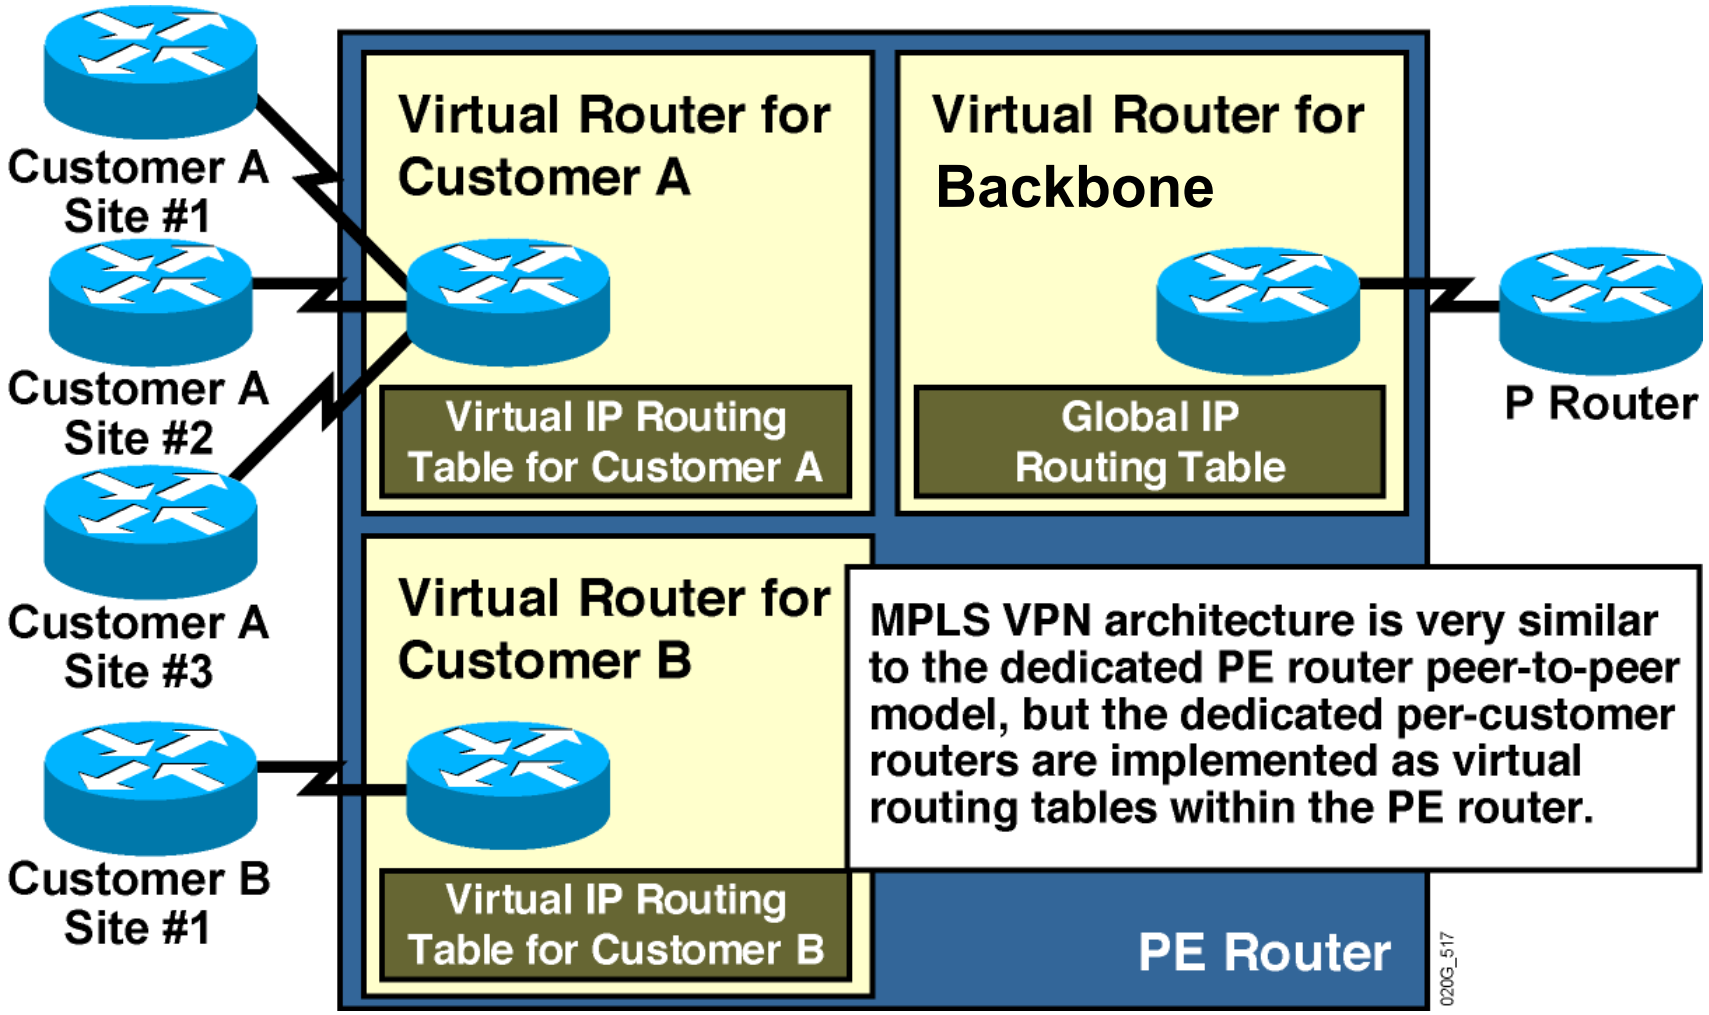

Propagation of Routing Information Across the P-Network

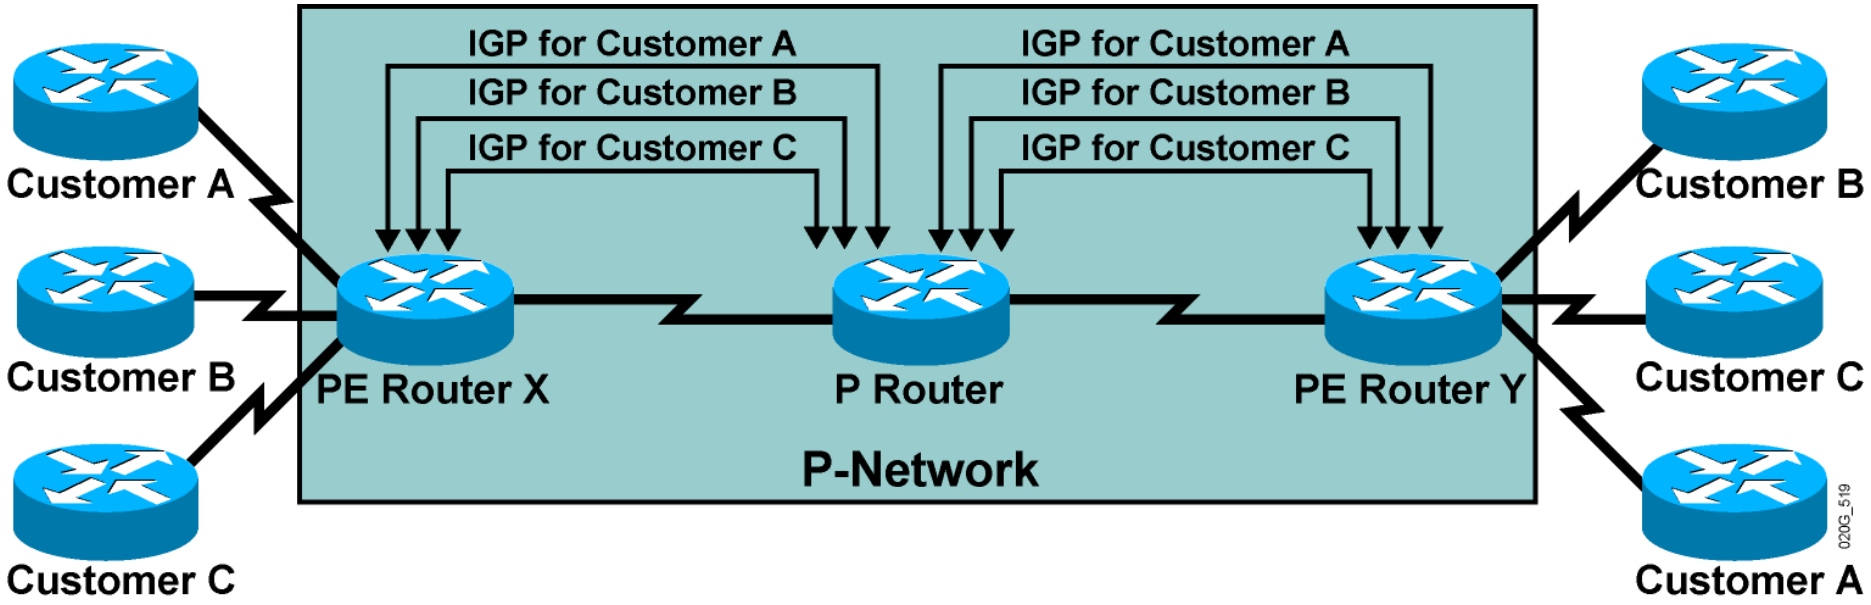

Question: How will PE routers exchange customer routing information?

Answer #1: Run a dedicated Interior Gateway Protocol (IGP) for each customer across the P-network.

This is the wrong answer for the following reasons:

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

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

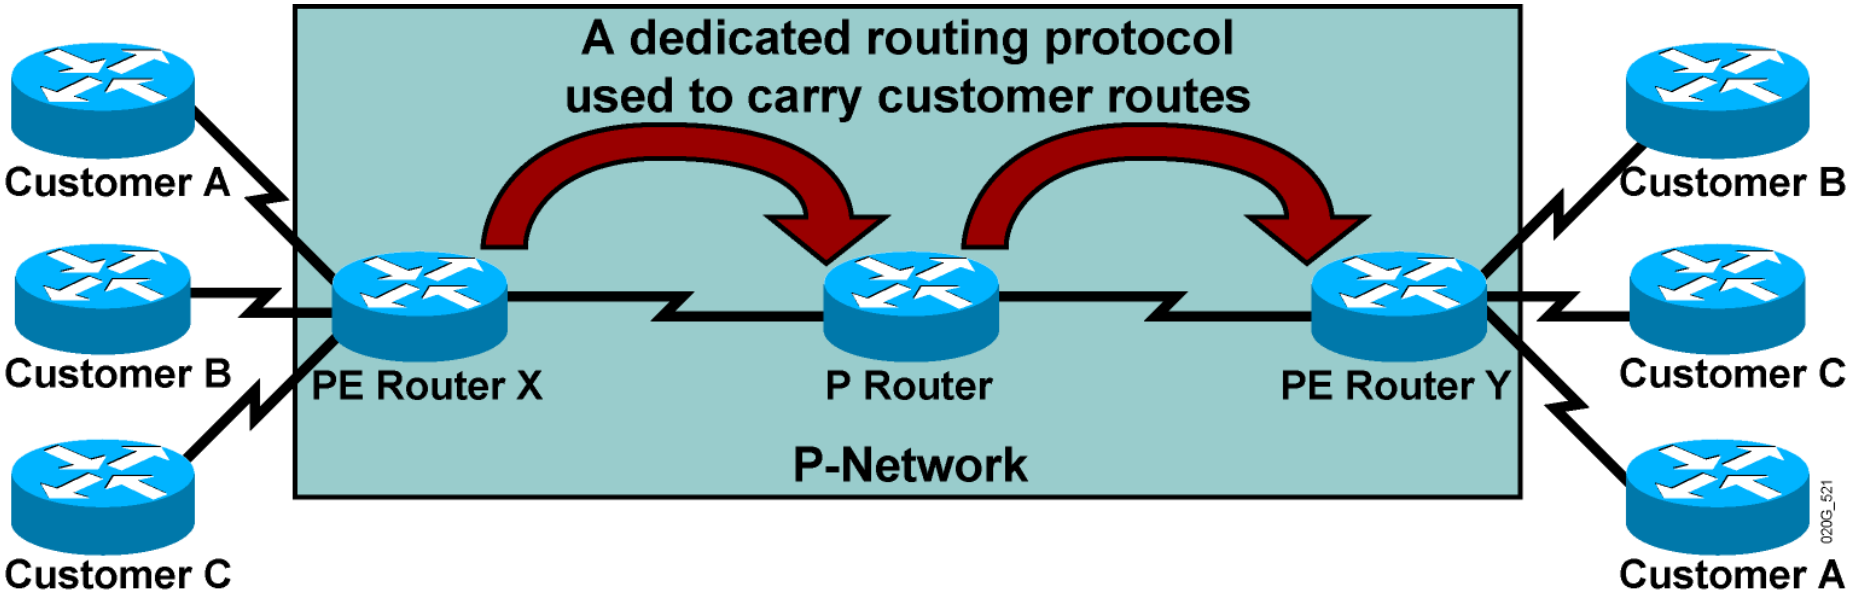

Question: How will PE routers exchange customer routing information?

Answer #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.)

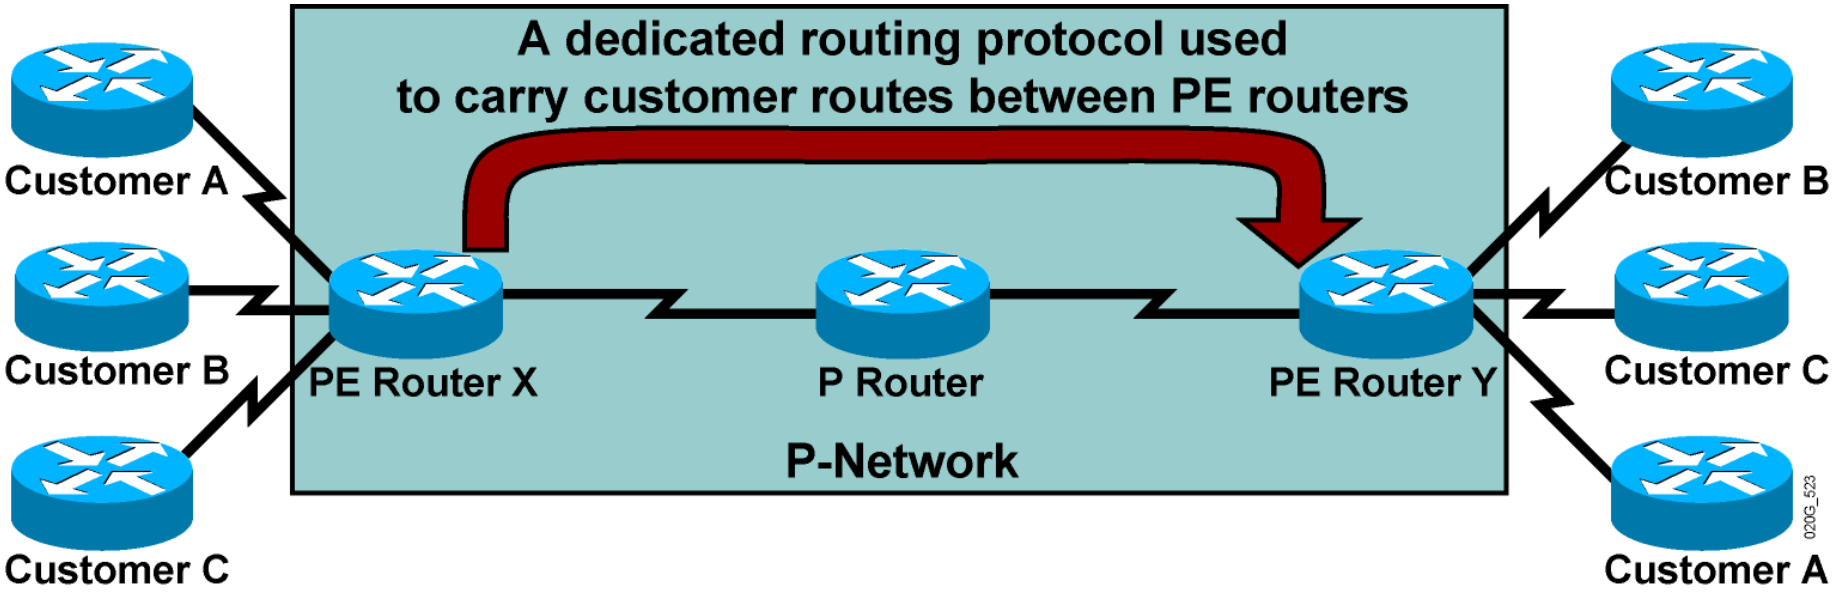

- Question: How will PE routers exchange customer routing information?
- Answer #3: Run a single routing protocol that will carry all customer routes between PE routers. Use MPLS labels to exchange packets between PE routers.
- The best answer:
  - **P routers do not carry customer routes; the solution is scalable.**

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

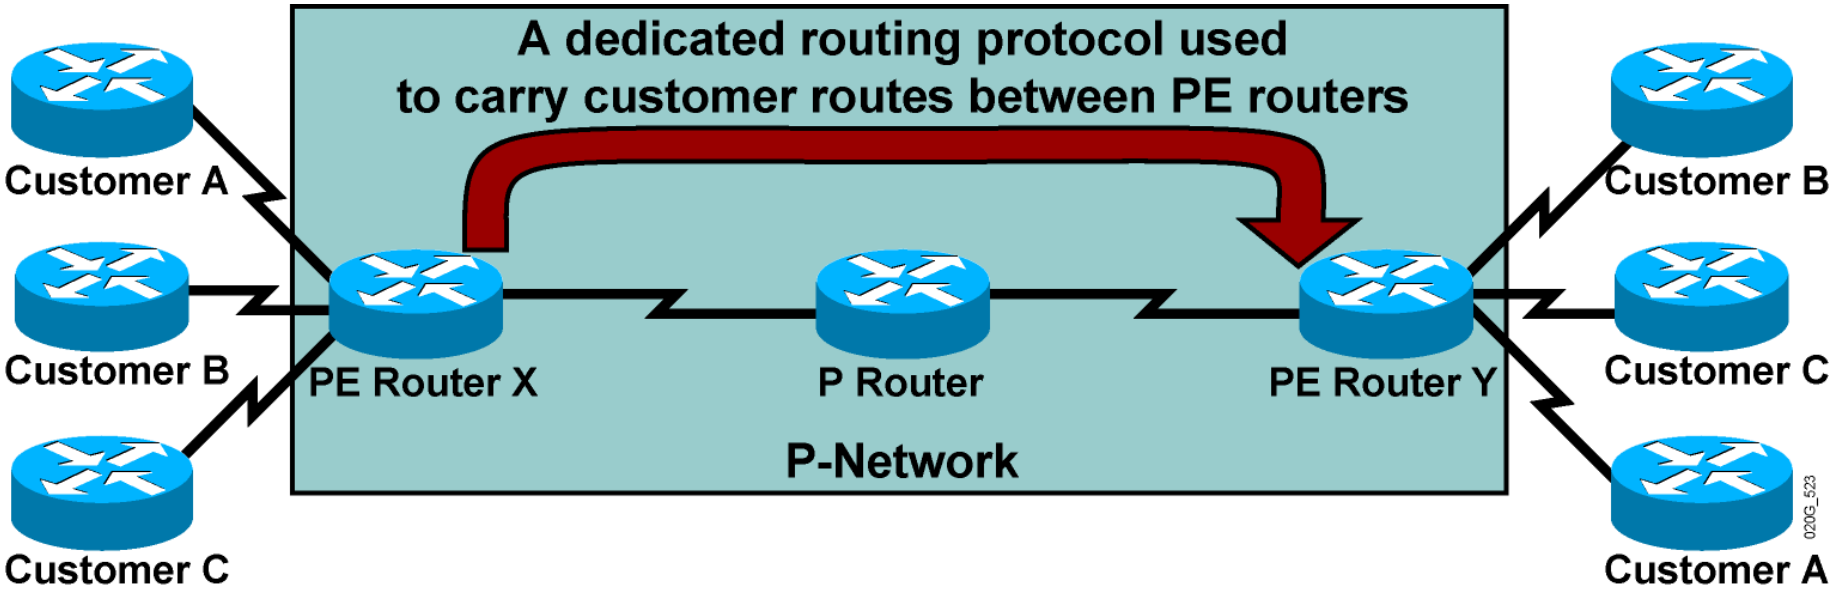

Question: Which protocol can be used to carry customer routes between PE routers?

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.)

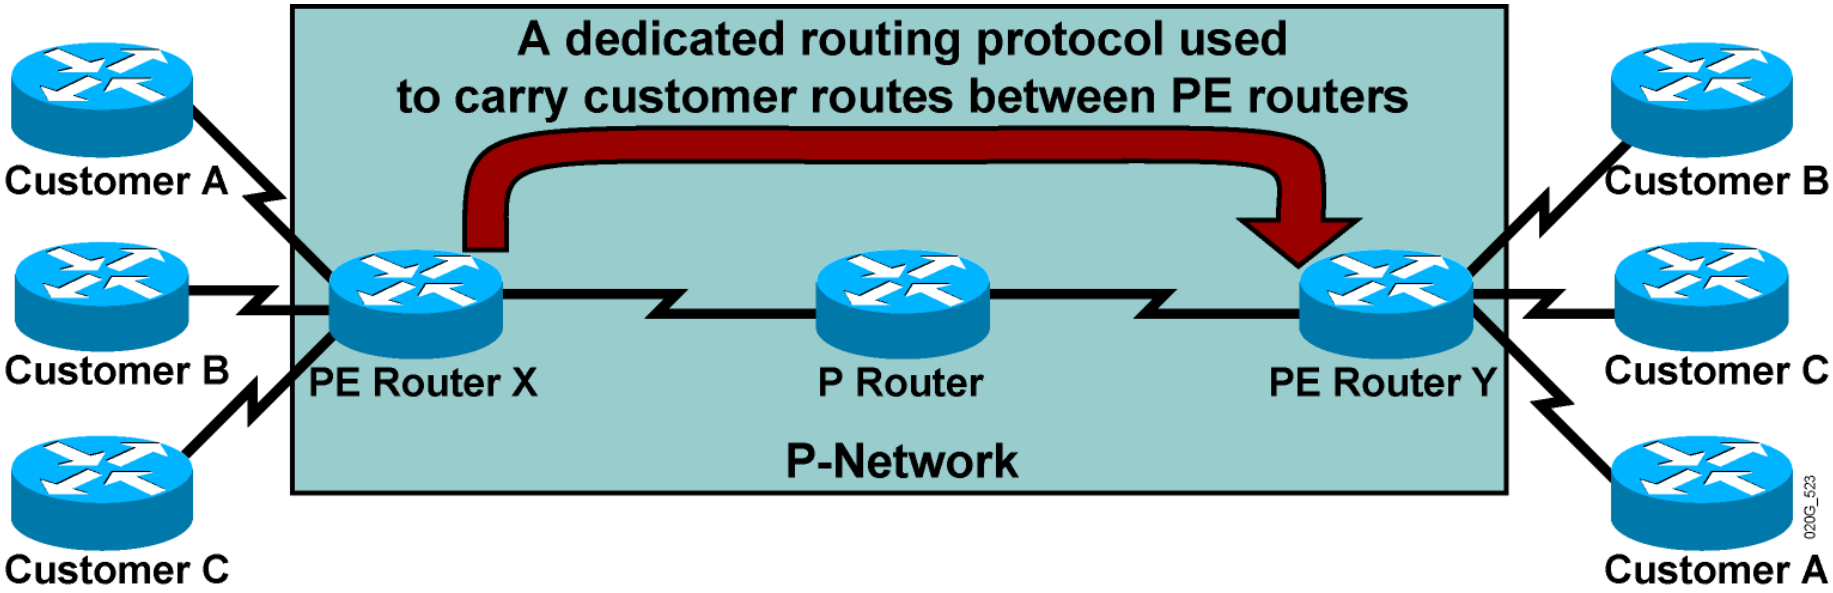

Question: How will information about the overlapping subnets of two customers be propagated via a single routing protocol?

Answer: Extend the customer addresses to make them unique.

Route Distinguishers

- **The 64-bit route distinguisher (RD) 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 Multiprotocol BGP (MP-BGP).

Route Distinguishers (Cont.)

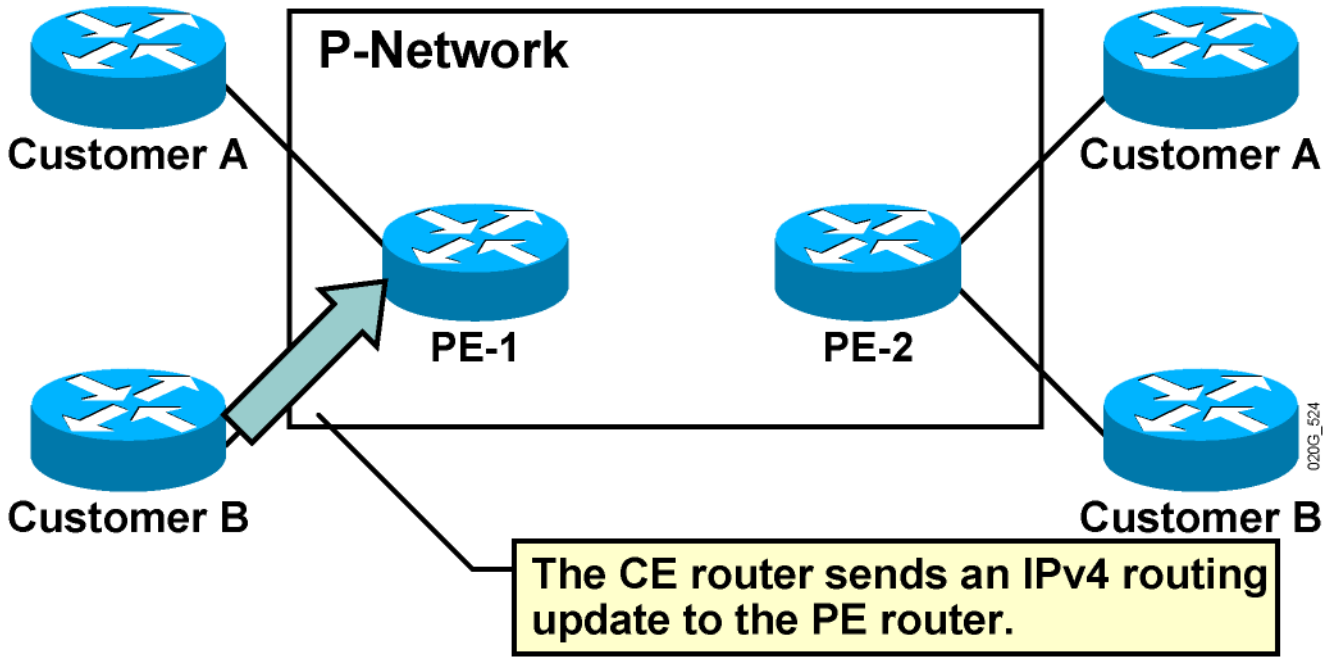

Route Distinguishers (Cont.)

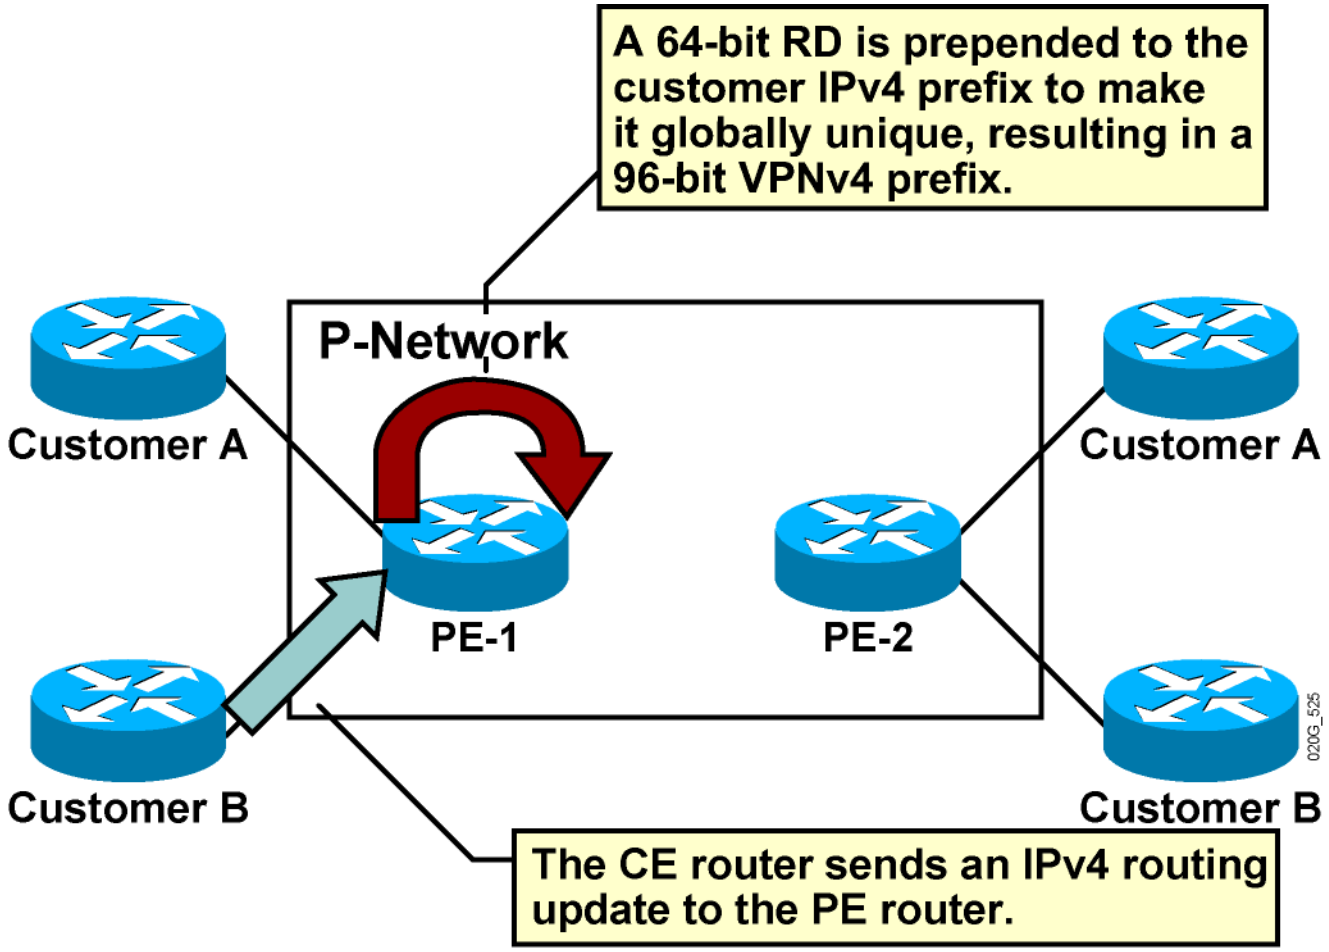

Route Distinguishers (Cont.)

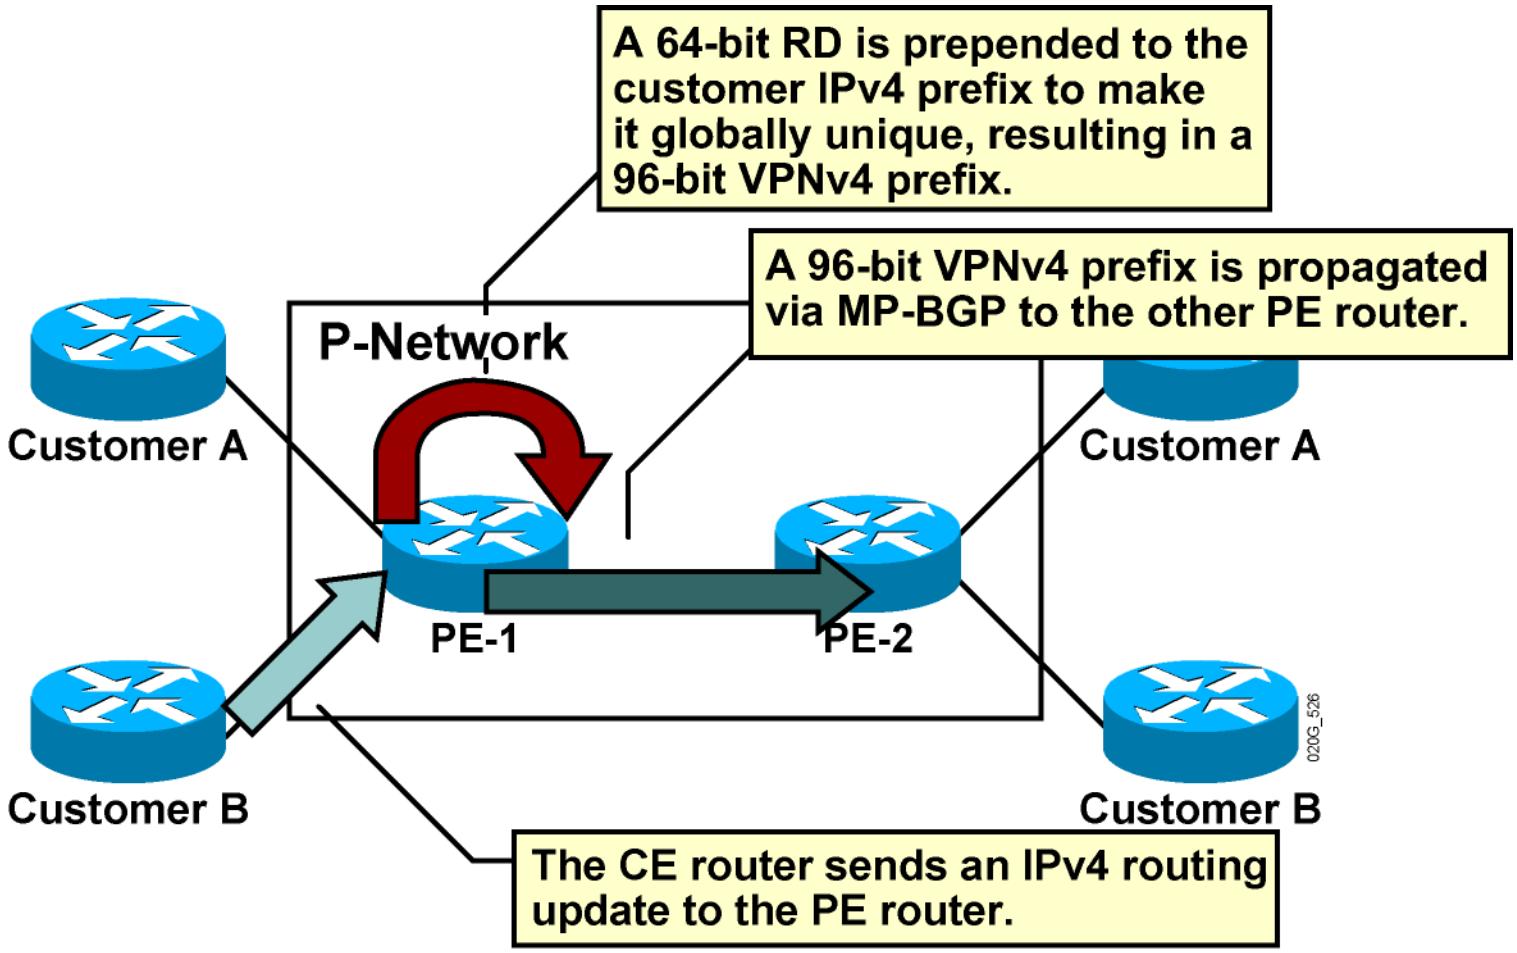

Route Distinguishers (Cont.)

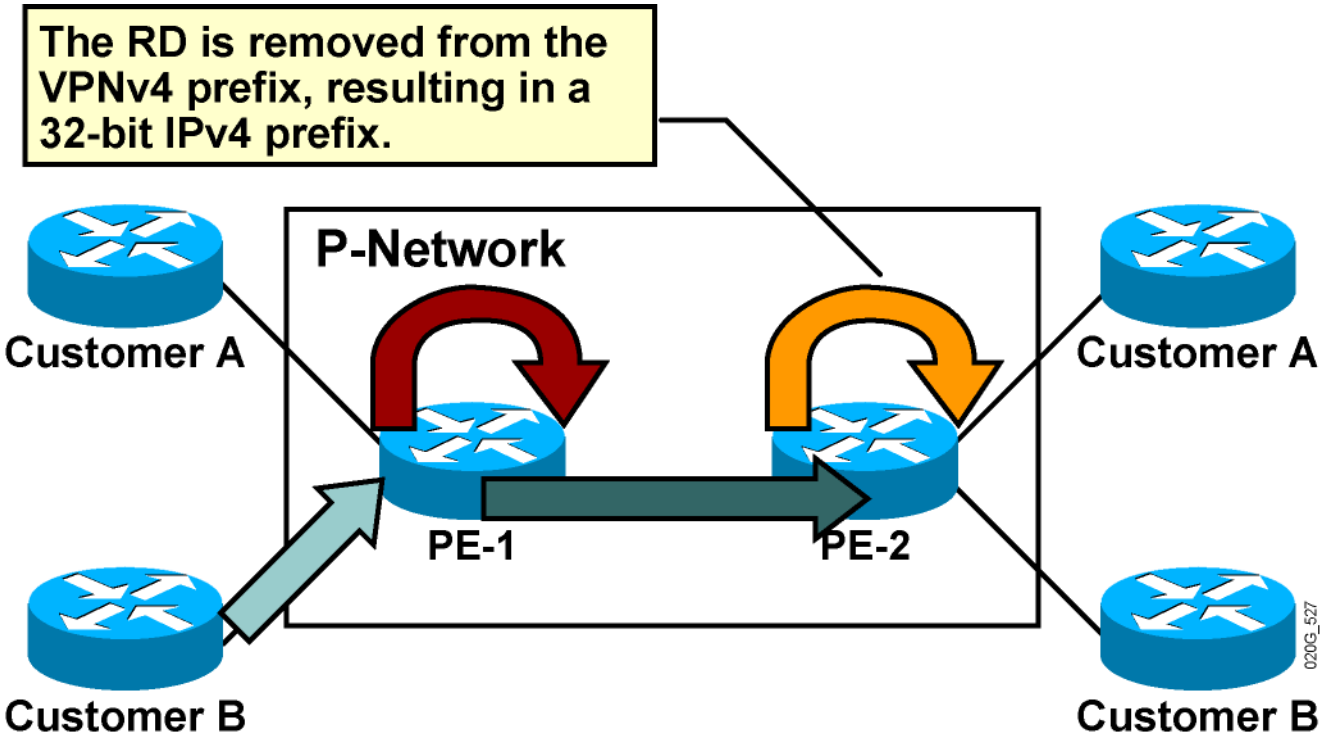

Route Distinguishers (Cont.)

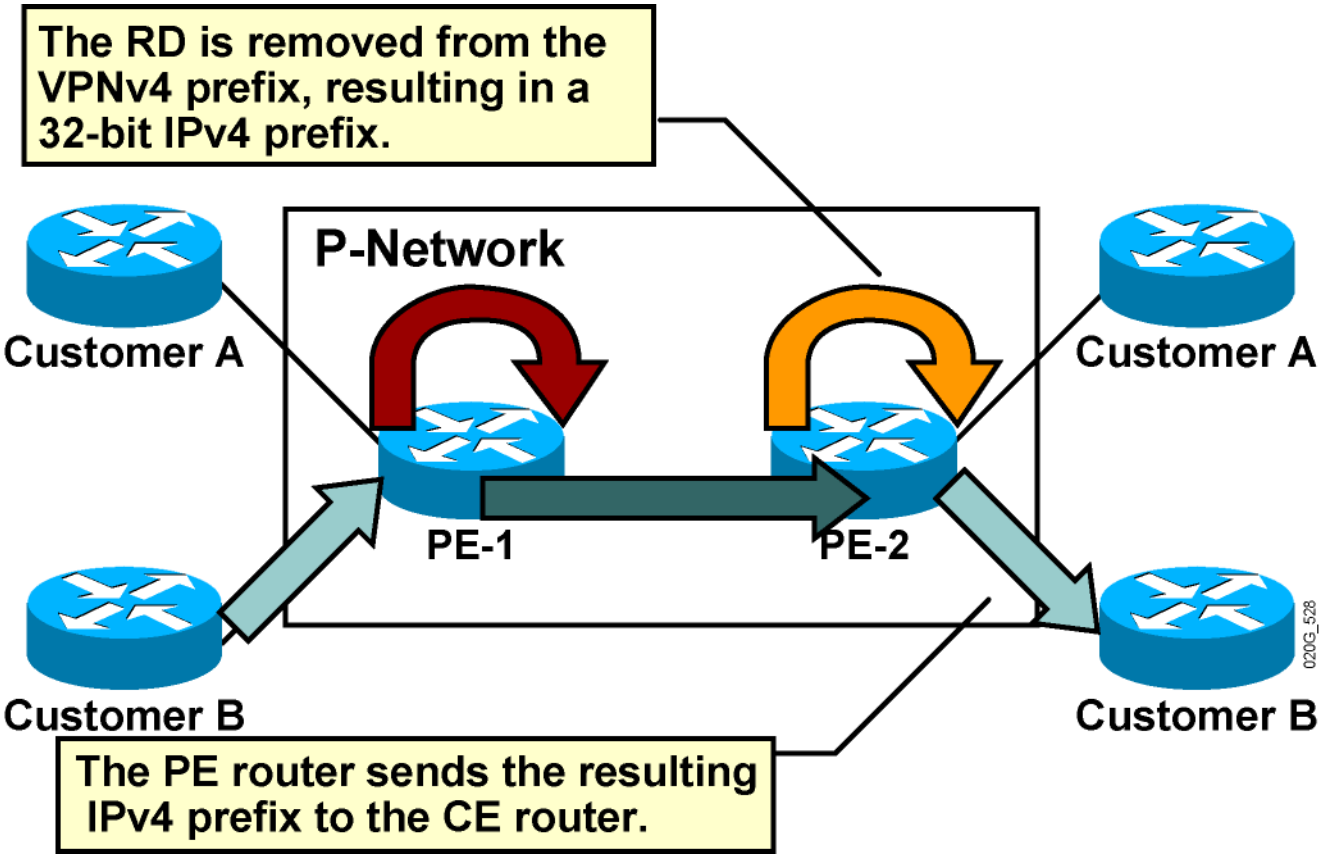

Route Distinguishers (Cont.)
Usage in an MPLS VPN

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

Route Targets - VoIP Service Sample

Why is the RD not used to identify the VPN?

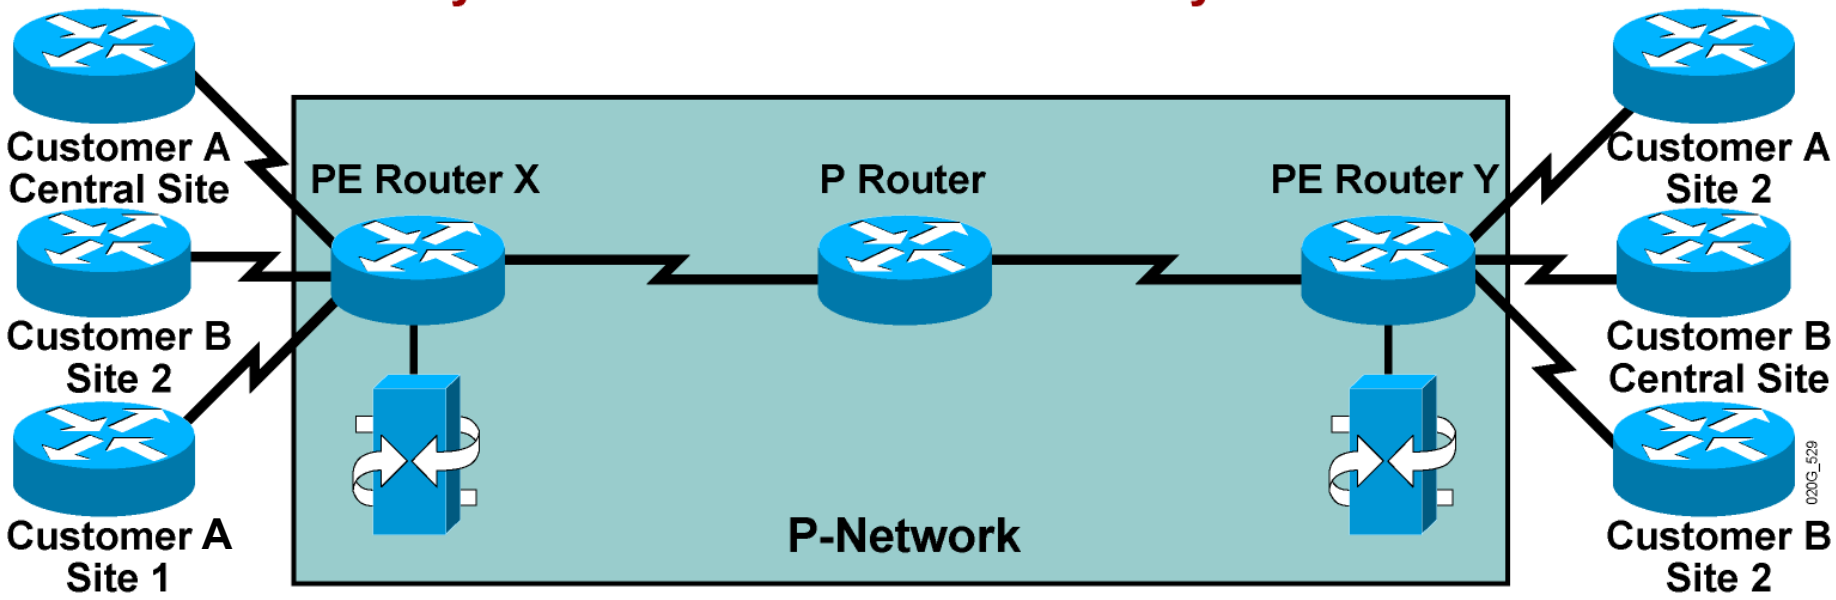

Requirements:

- All sites of one customer need to communicate.
- Central sites of both customers need to communicate with VoIP gateways and other central sites.
- Other sites from different customers do not communicate with each other.

Route Targets (Cont.)
Connectivity Requirements

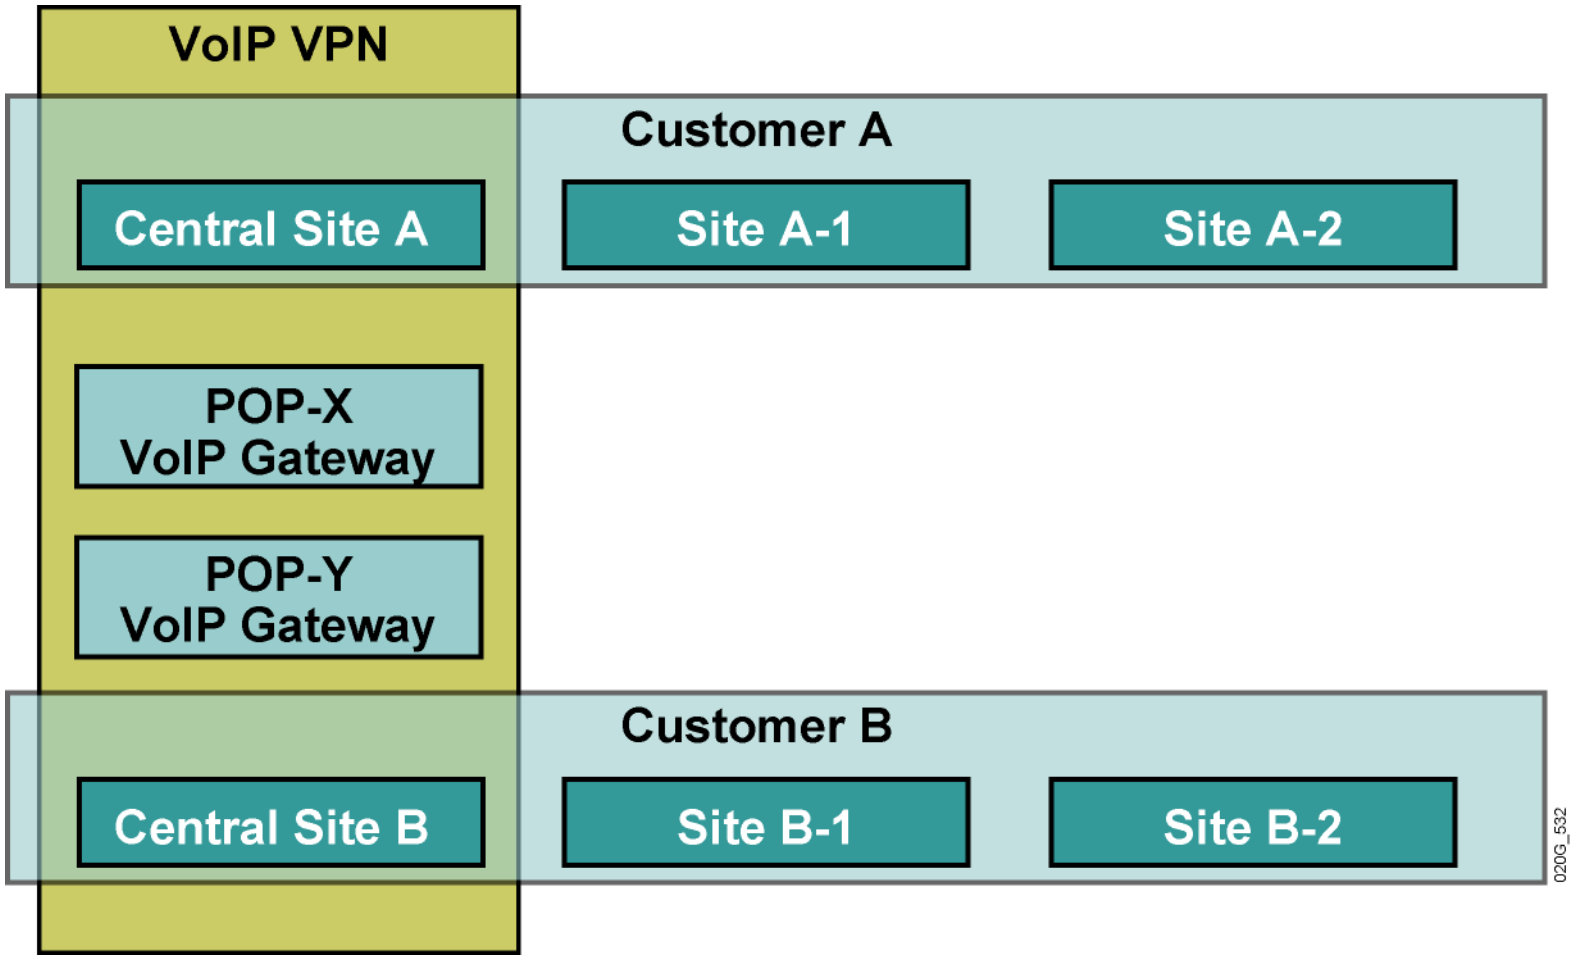

Route Targets (Cont.) - Why Are They Needed?

- 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.

Route Targets (Cont.) - 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.
  - Extended communities carry the meaning of the attribute together with its value.
- Any number of RTs can be attached to a single route.

Route Targets (Cont.) - 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**

MPLS VPN Routing - CE Router Perspective

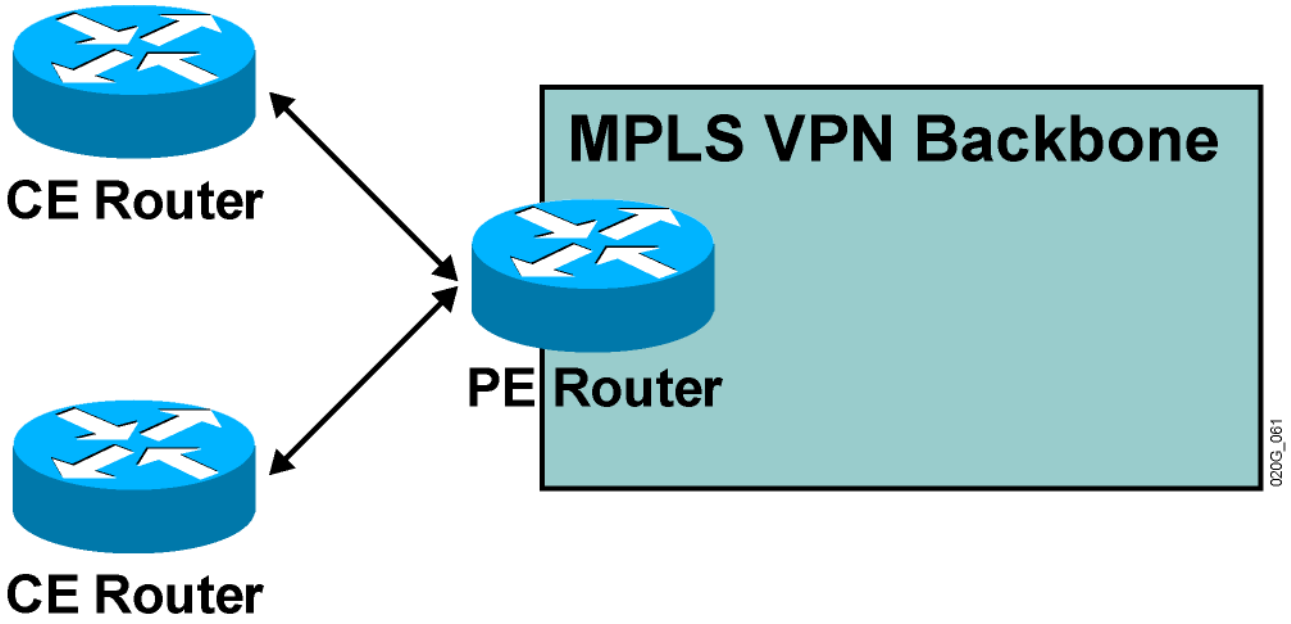

- The CE routers run standard IP routing software and exchange routing updates with the PE router.
  - EBGP, OSPF, RIPv2, EIGRP, and static routes are supported.
- The PE router appears as another router in the C-network.

MPLS VPN Routing (cont.)
Overall Customer Perspective

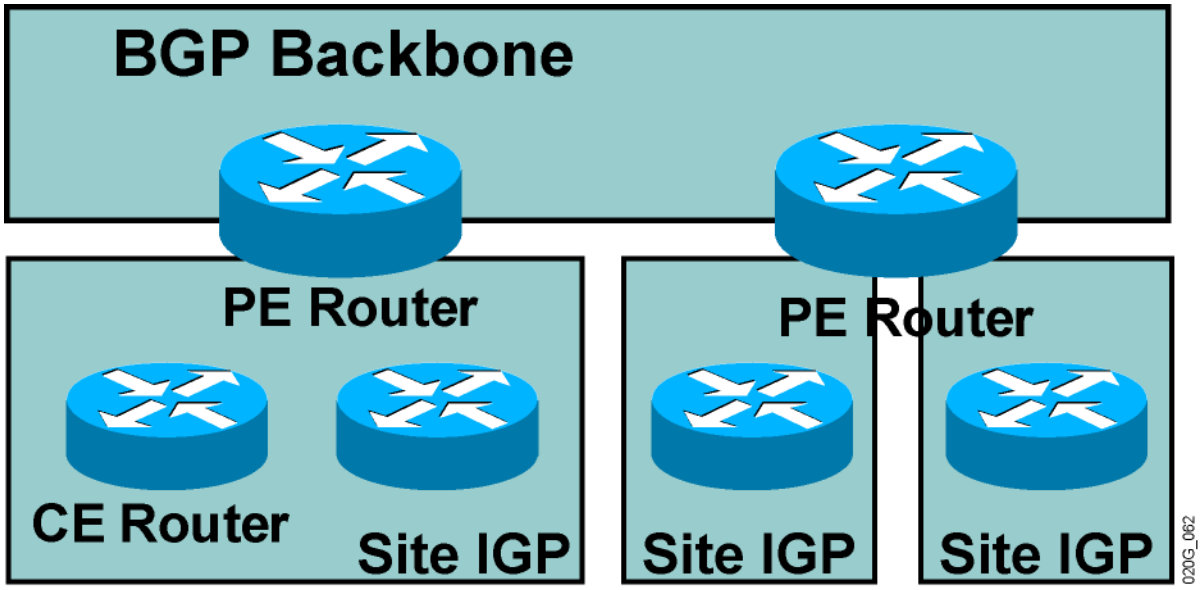

- To the customer, the PE routers appear as core routers connected via a BGP backbone.
- The usual BGP and IGP design rules apply.
- The P routers are hidden from the customer.

MPLS VPN Routing (Cont.)
P Router Perspective

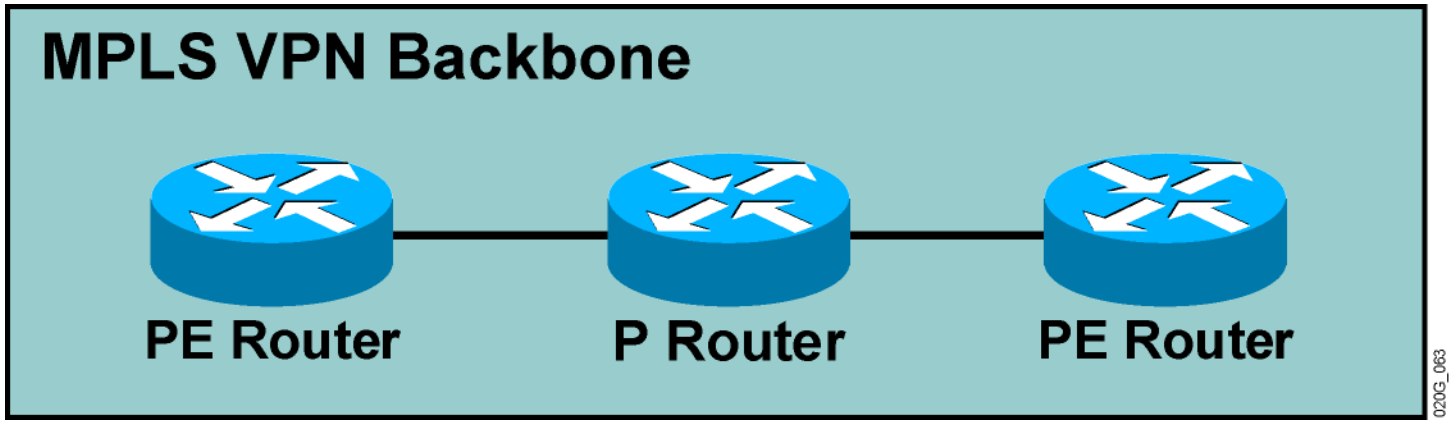

- P routers do not participate in MPLS VPN routing and do not carry VPN routes.
- P routers run backbone IGP with the PE routers and exchange information about global subnets (core links and loopbacks).

MPLS VPN Routing (Cont.)
PE Router Perspective

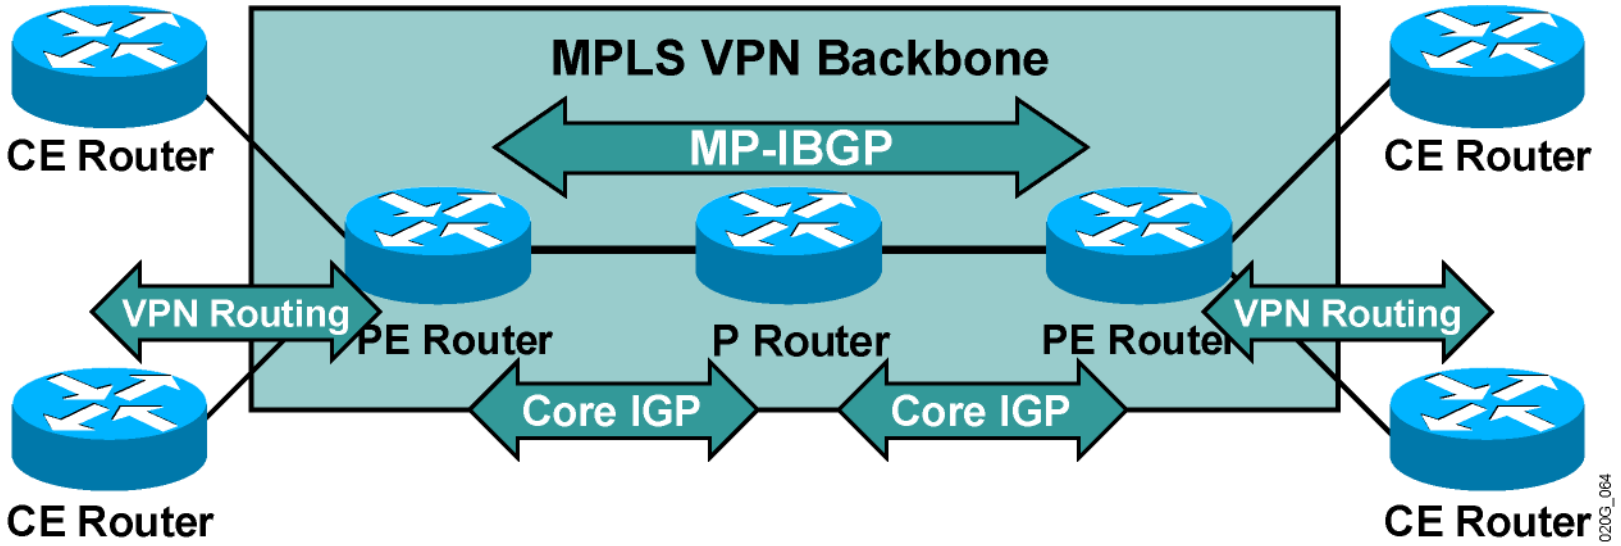

- PE routers:
  - Exchange VPN routes with CE routers via per-VPN routing protocols
  - Exchange core routes with P routers and PE routers via core IGP
  - Exchange VPNv4 routes with other PE routers via MP-IBGP sessions

Routing Tables on PE Routers

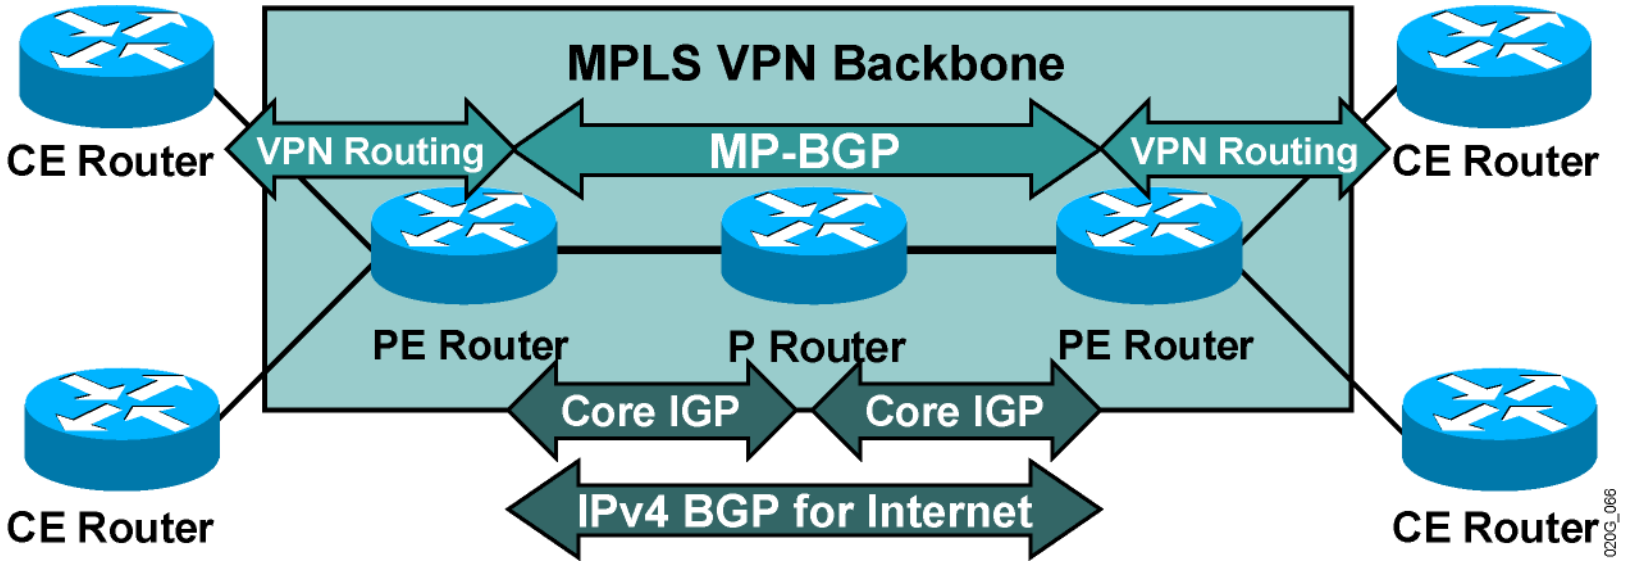

- **PE routers contain a number of routing tables:**
  - Global routing table, which contains core routes (filled with core IGP) and Internet routes (filled with IPv4 BGP)
  - VRF tables for sets of sites with identical routing requirements
  - VRFs filled with information from CE routers and MP-BGP information from other PE routers

VPN Packet Forwarding Across an MPLS VPN Backbone

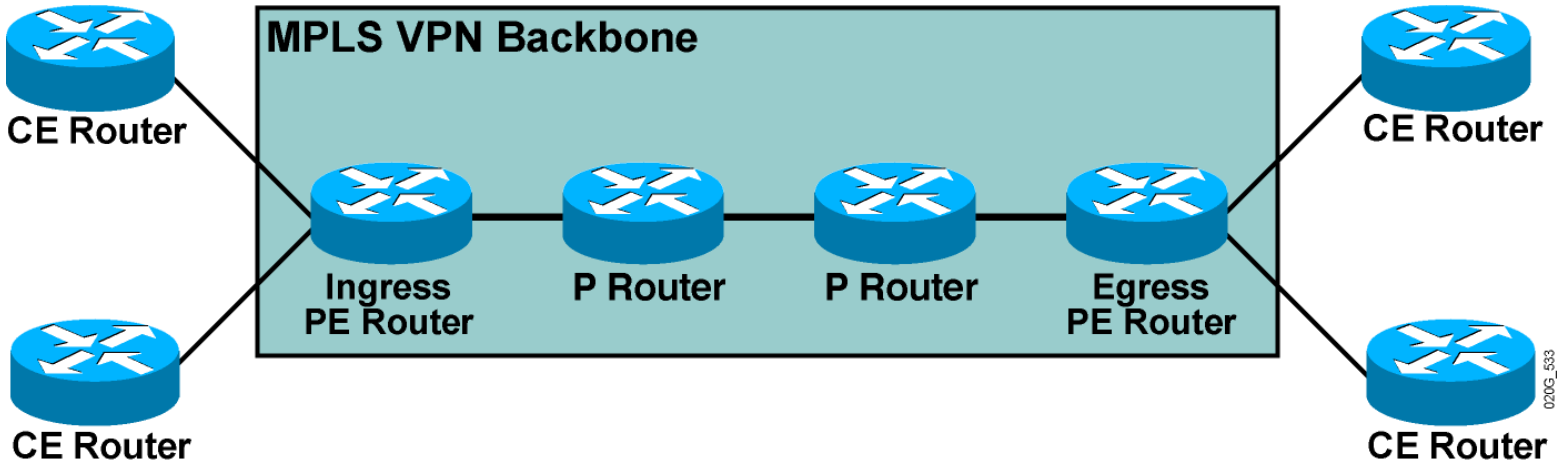

Question: How will the PE routers forward the VPN packets across the MPLS VPN backbone?

Answer #1: They will label the VPN packets with an LDP label for the egress PE router and forward the labeled packets across the MPLS backbone.

VPN Packet Forwarding Across an MPLS VPN Backbone

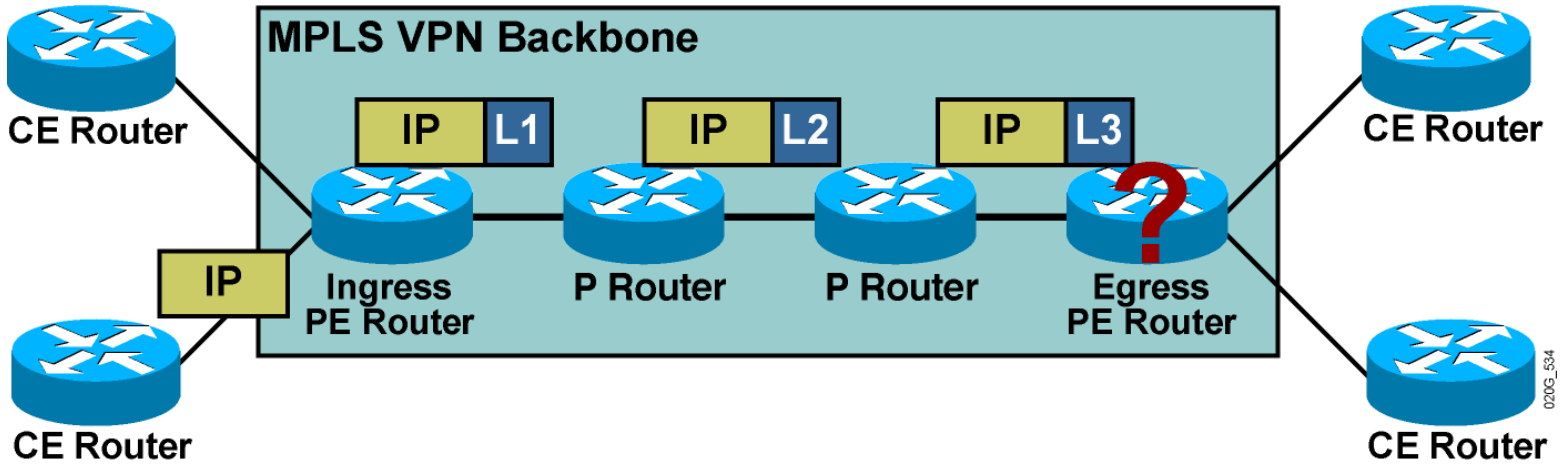

Question: How will the PE routers forward the VPN packets across the MPLS VPN backbone?

Answer #1: They 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.
- However, the egress PE router does not know which VRF to use for packet switching, so the packet is dropped.
- How about using a label stack?

VPN Packet Forwarding Across an MPLS VPN Backbone (Cont.)

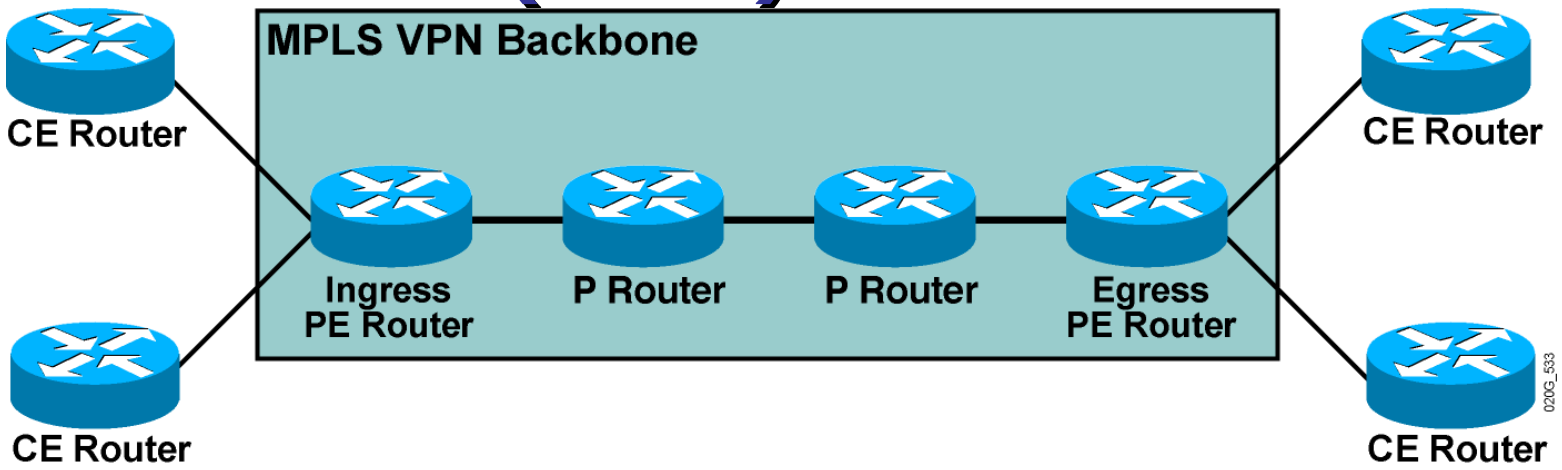

Question: How will the PE routers forward the VPN packets across the MPLS VPN backbone?

Answer #2: They 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.

VPN Packet Forwarding Across an MPLS VPN Backbone (Cont.)

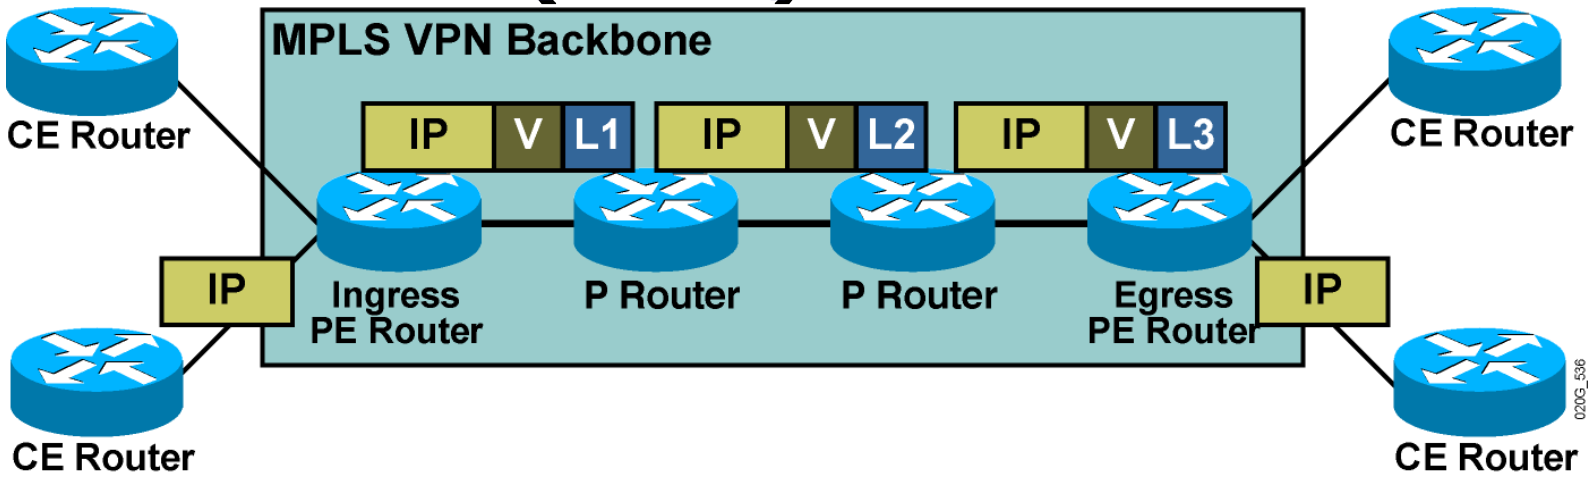

Question: How will the PE routers forward the VPN packets across the MPLS VPN backbone?

Answer #2: They 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, and the packet reaches the egress PE router.
- The egress PE router performs a lookup on the VPN label and forwards the packet toward the CE router.

VPN Penultimate Hop Popping

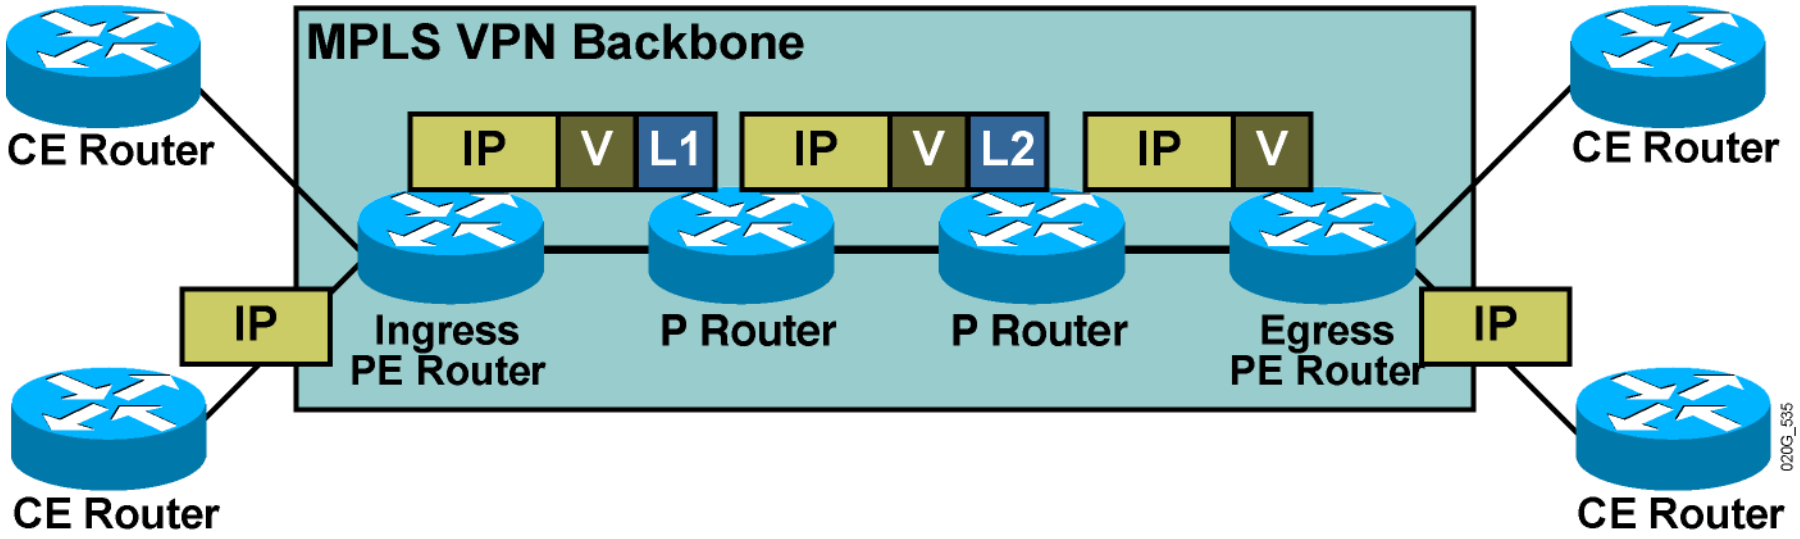

- Penultimate hop popping on the LDP label can be performed on the last P router.
- The egress PE router performs label lookup only on the VPN label, resulting in faster and simpler label lookup.
- IP lookup is performed only once—in the ingress PE router.

VPN Label Propagation

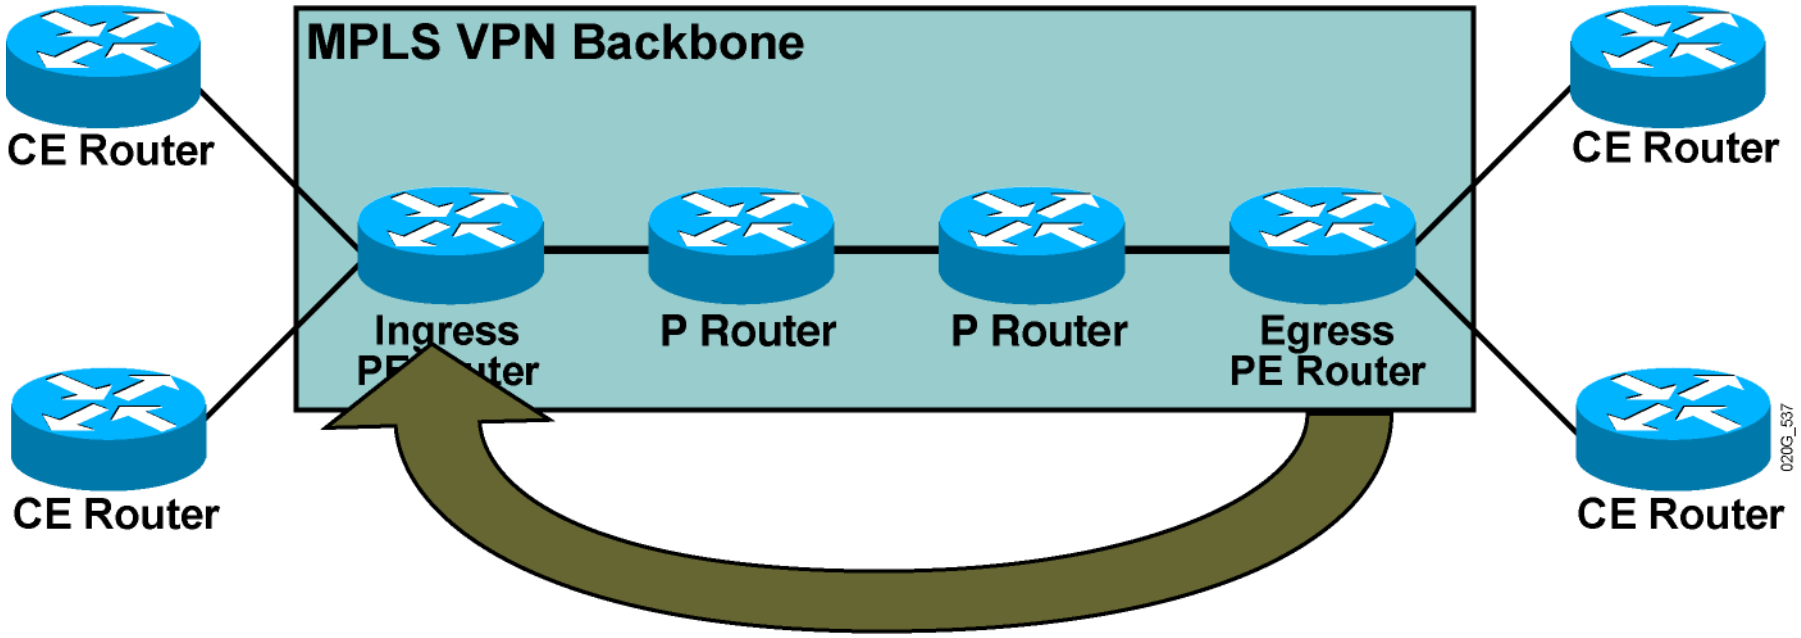

Question: How will the ingress PE router get the second label in the label stack from the egress PE router?

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

VPN Label Propagation (Cont.)

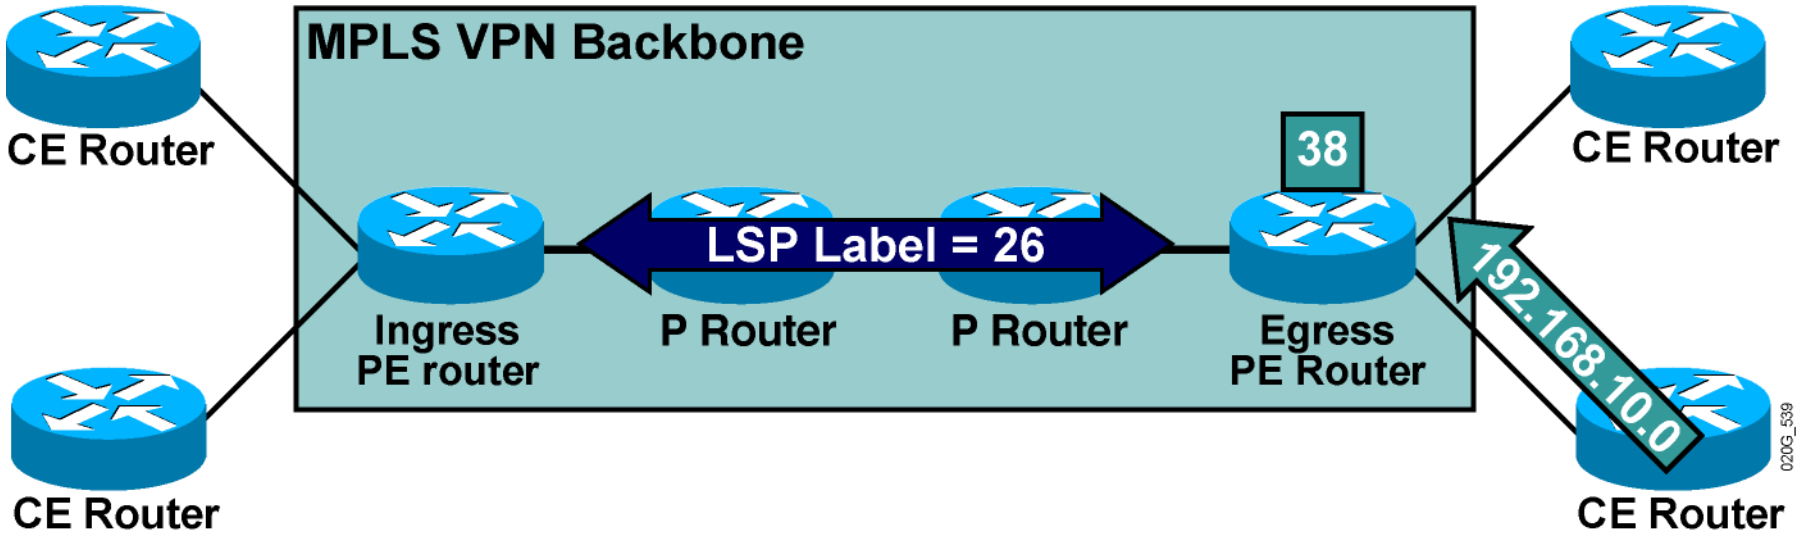

1: A VPN label is assigned to every VPN route by the egress PE router.

VPN Label Propagation (Cont.)

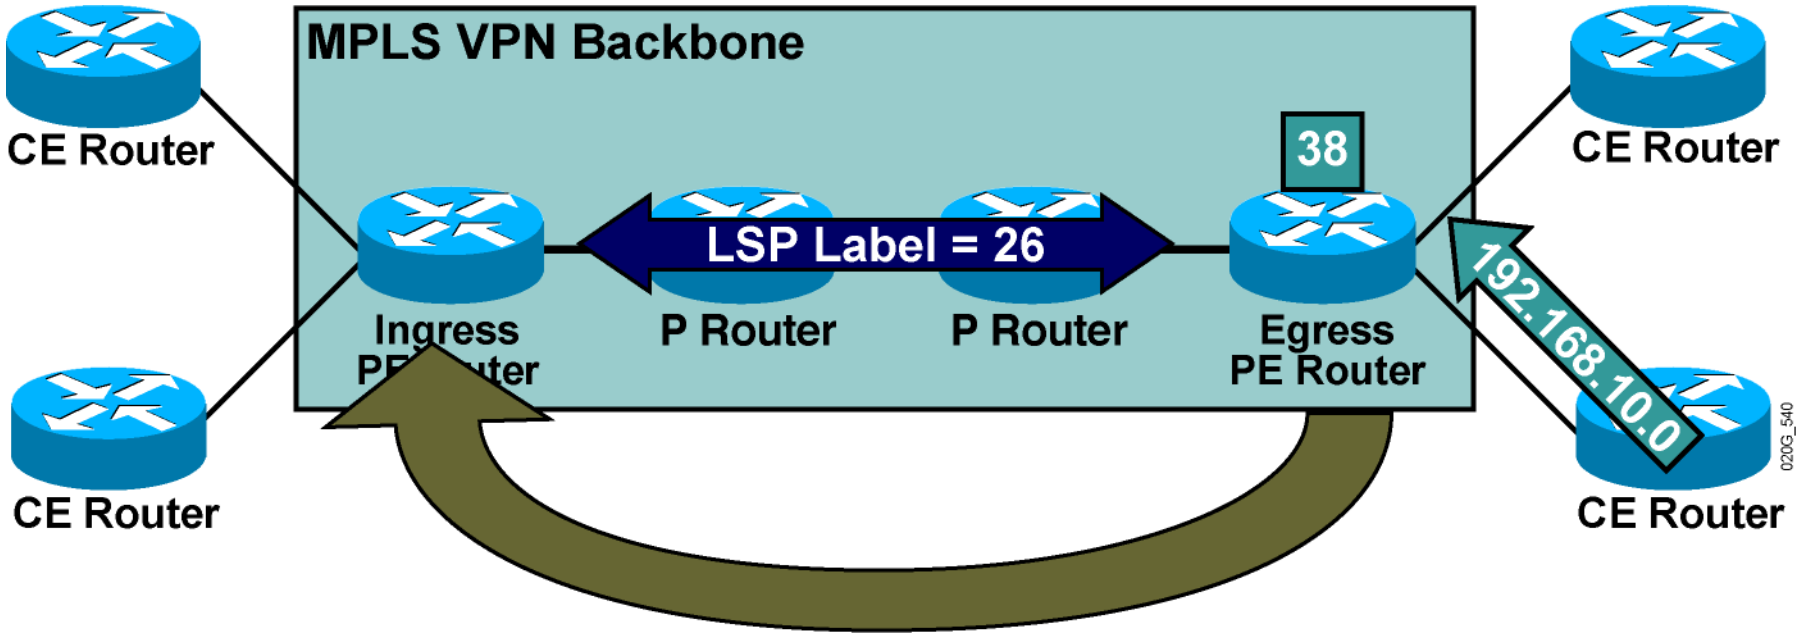

1: A VPN label is assigned to every VPN route by the egress PE router.

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

VPN Label Propagation (Cont.)

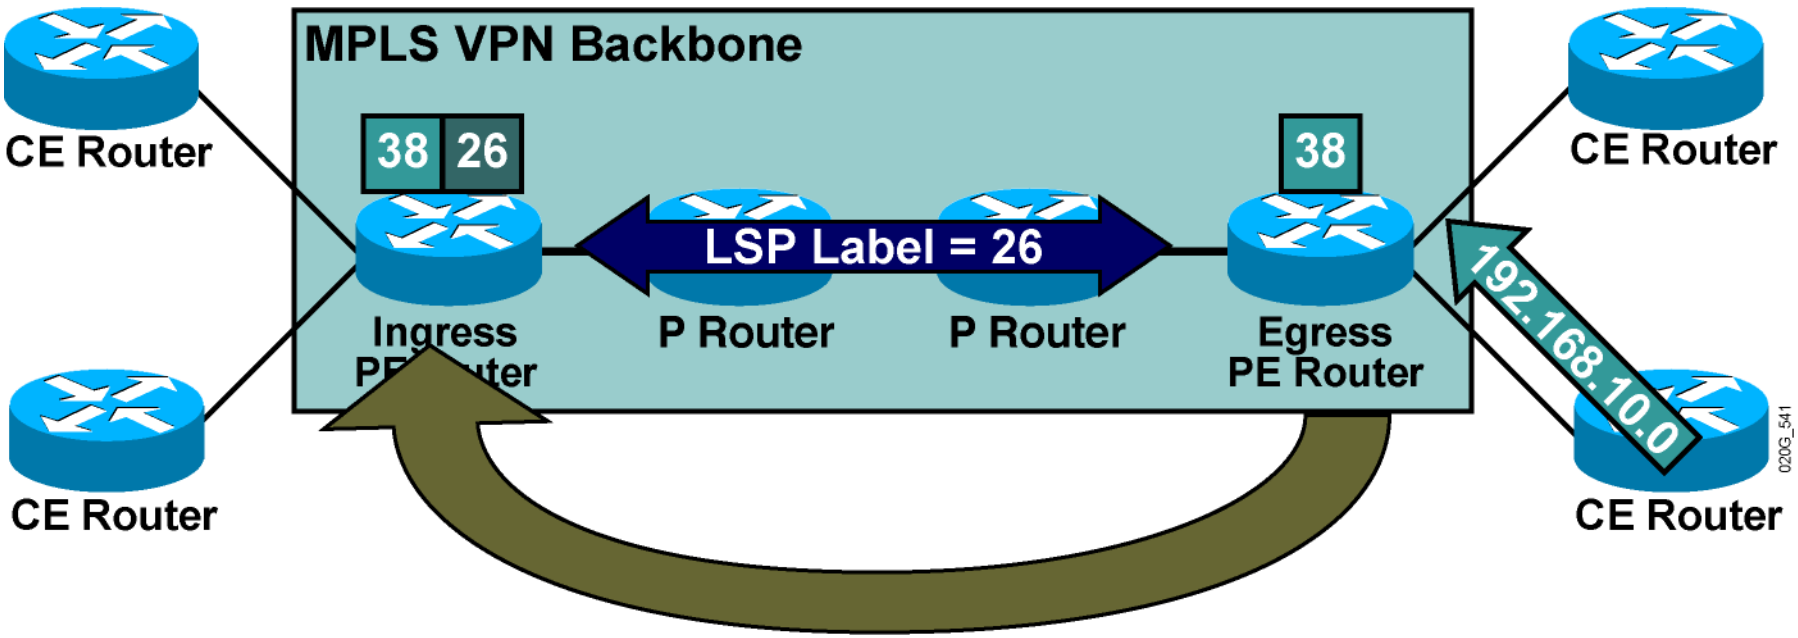

1: A VPN label is assigned to every VPN route by the egress PE router.

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

3: A label stack is built in the VRF table.

Review

- Traditional Router-Based Networks
- Virtual Private Networks
- VPN Terminology
- MPLS VPN Architecture
- MPLS VPN Routing
- MPLS VPN Label Propagation