MPLS Implementation
MPLS VPN

Describing MPLS VPN Technology

Objectives

- Describe VPN implementation models.
- Compare and contrast VPN overlay VPN models.
- Describe the benefits and disadvantages of the overlay VPN implementation model.
- Describe the benefits and disadvantages of the peer-to-peer VPN implementation model.
- Describe the features of the MPLS VPN architecture.
- Describe routing in the MPLS VPN architecture.

VPN Taxonomy

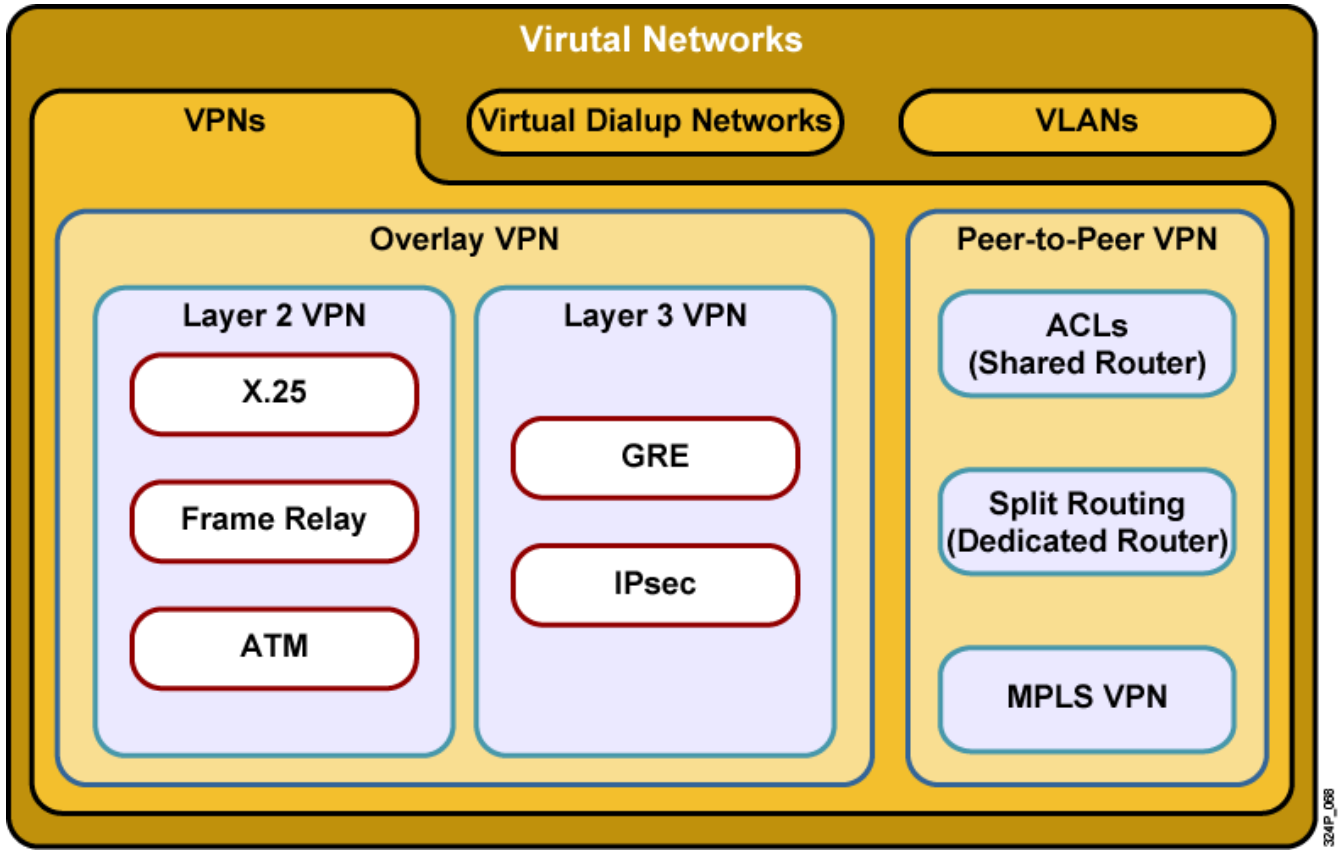

- Overlay VPNs—Service providers provide virtual point-to-point links.
- Peer-to-peer VPNs—Service providers participate in the customer routing.

VPN Terminology

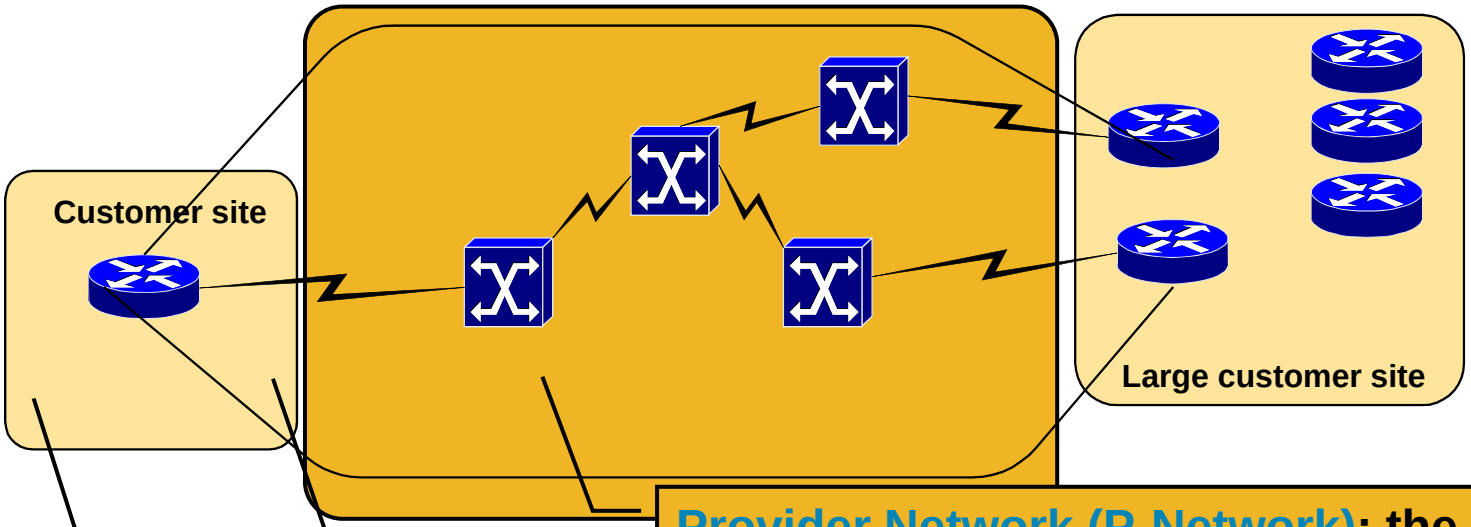

Provider Network (P-Network): the Service Provider infrastructure used to provide VPN services

Customer Network (C-Network): the part of the network still under customer control

Customer Site: a contiguous part of customer network (can encompass many physical locations)

VPN Terminology

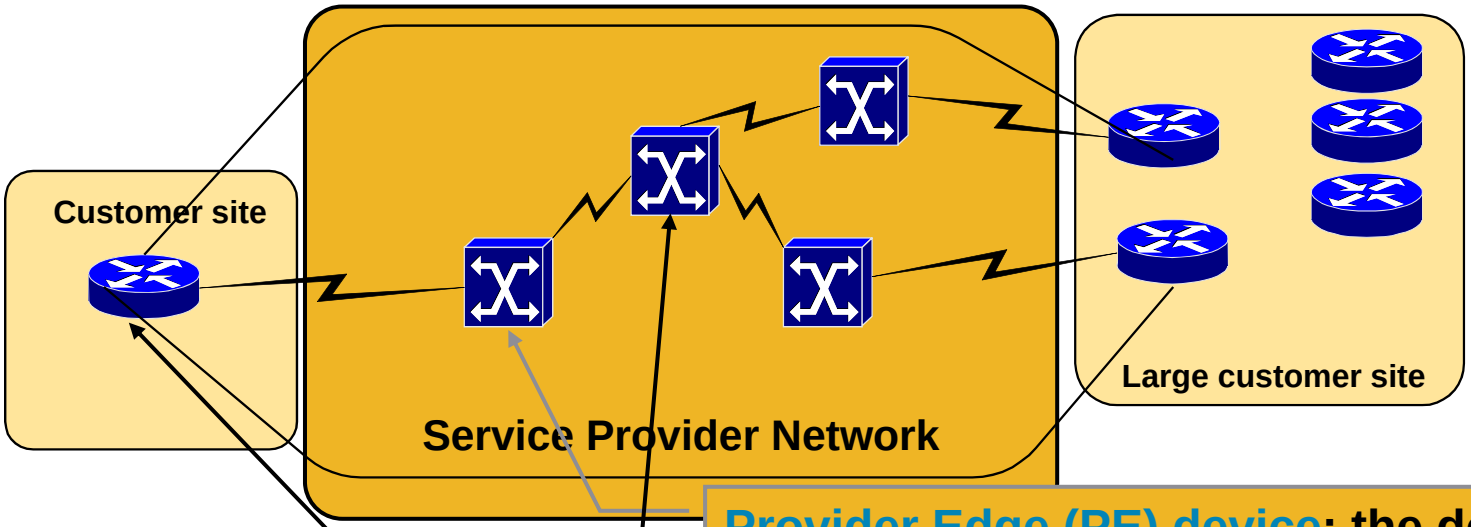

Provider Edge (PE) device: the device in the P-network to which the CE-devices are connected

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

Customer Edge (CE) device: the device in the C-network with link into P-network. Also called Customer Premises Equipment (CPE)

Overlay VPNs

- Layer 1 Overlay VPN
  - Mentioned for historical reasons only.
- Layer 2 Overlay VPN
  - Traditional switched WAN
  - Implmented with X.25, Frame Relay, ATM and SMDS
  - SP is responsible for transport of Layer 2 frames
  - Customer is responsible for all higher layers
- Layer 3 Overlay VPN
  - SP network is invisible to customer routers
  - Uses IP tunneling
  - SP provides point-to-point data transport between customer sites

Layer 2 Overlay VPN Using Frame Relay

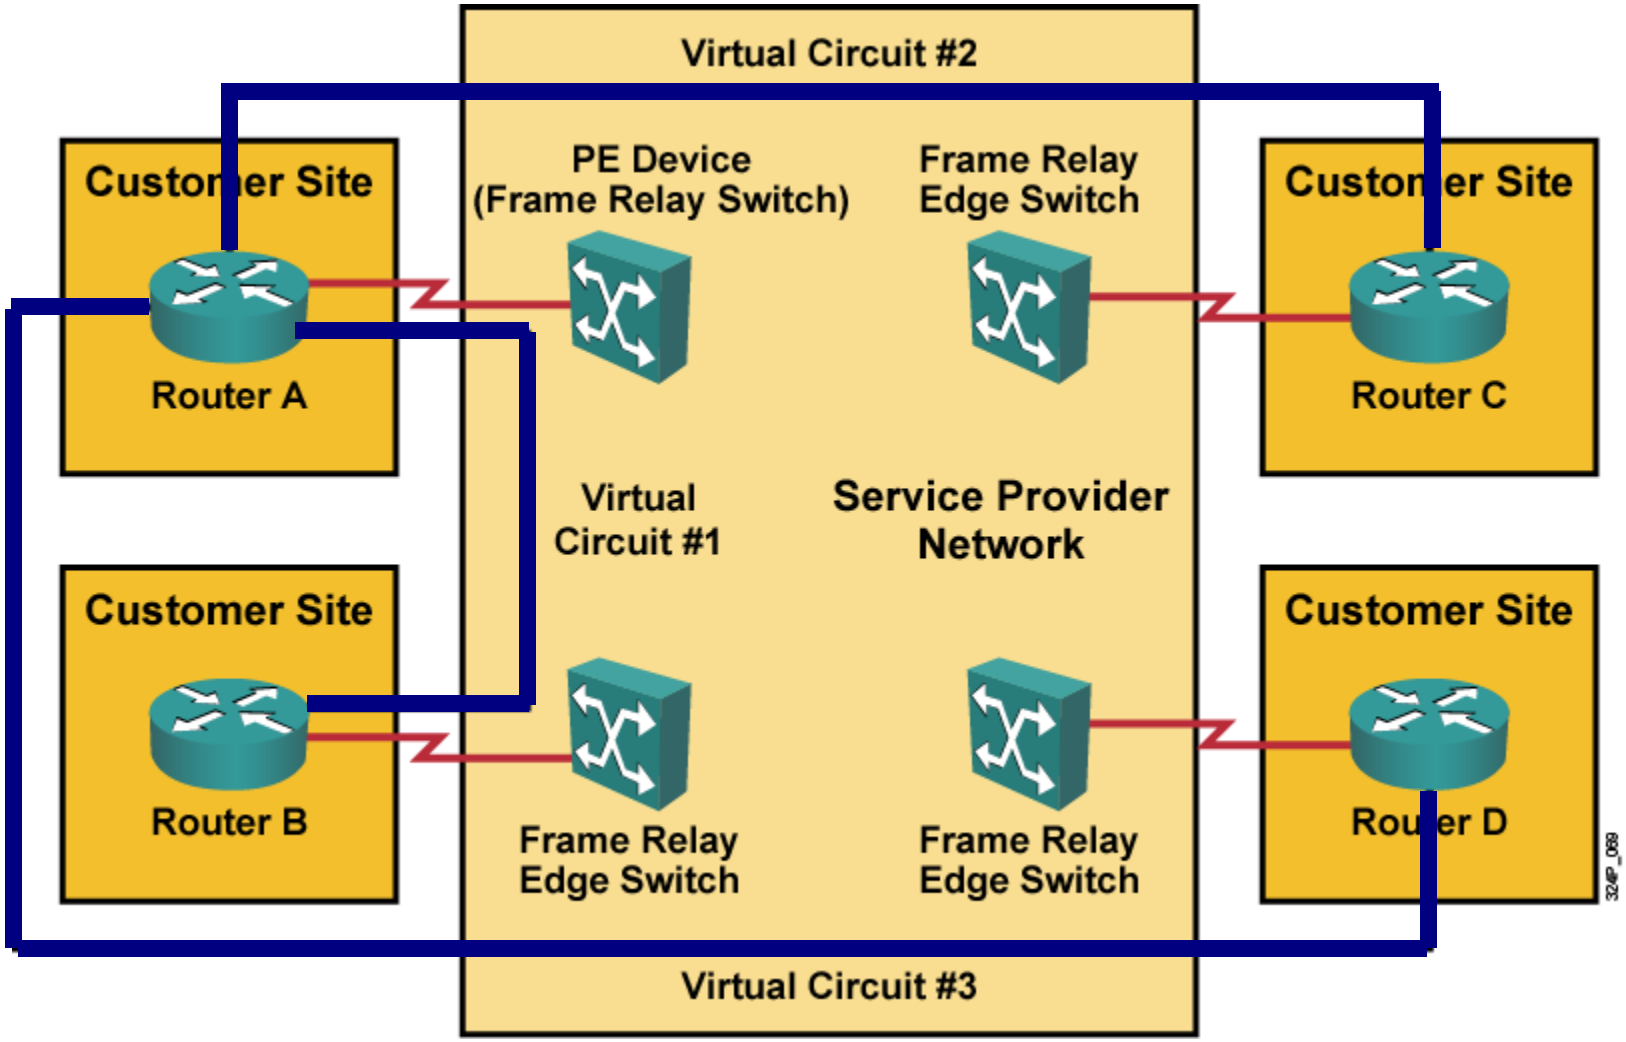

Layer 3 Overlay VPNs

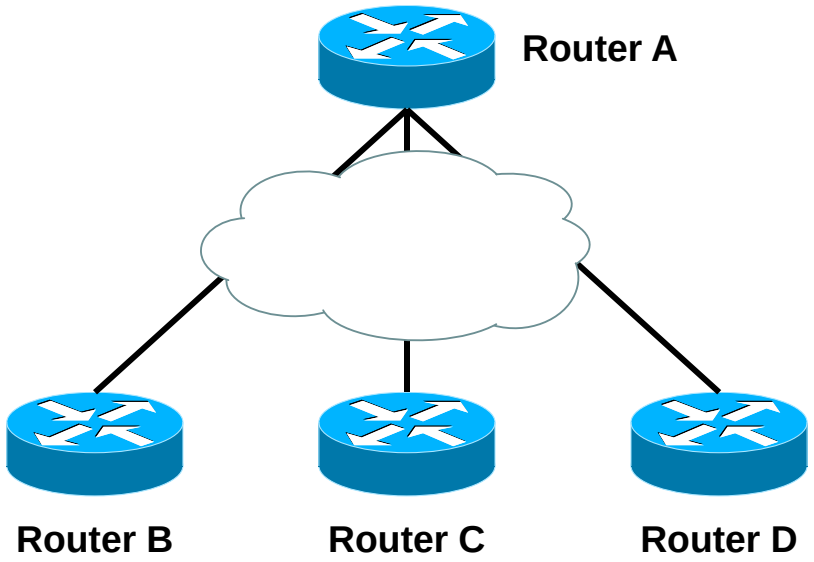

- The service provider infrastructure appears as point-to-point links to customer routes.
- Routing protocols run directly between customer routers.
- The service provider does not see customer routes and is responsible only for providing point-to-point transport of customer data.

Peer-to-Peer VPNs

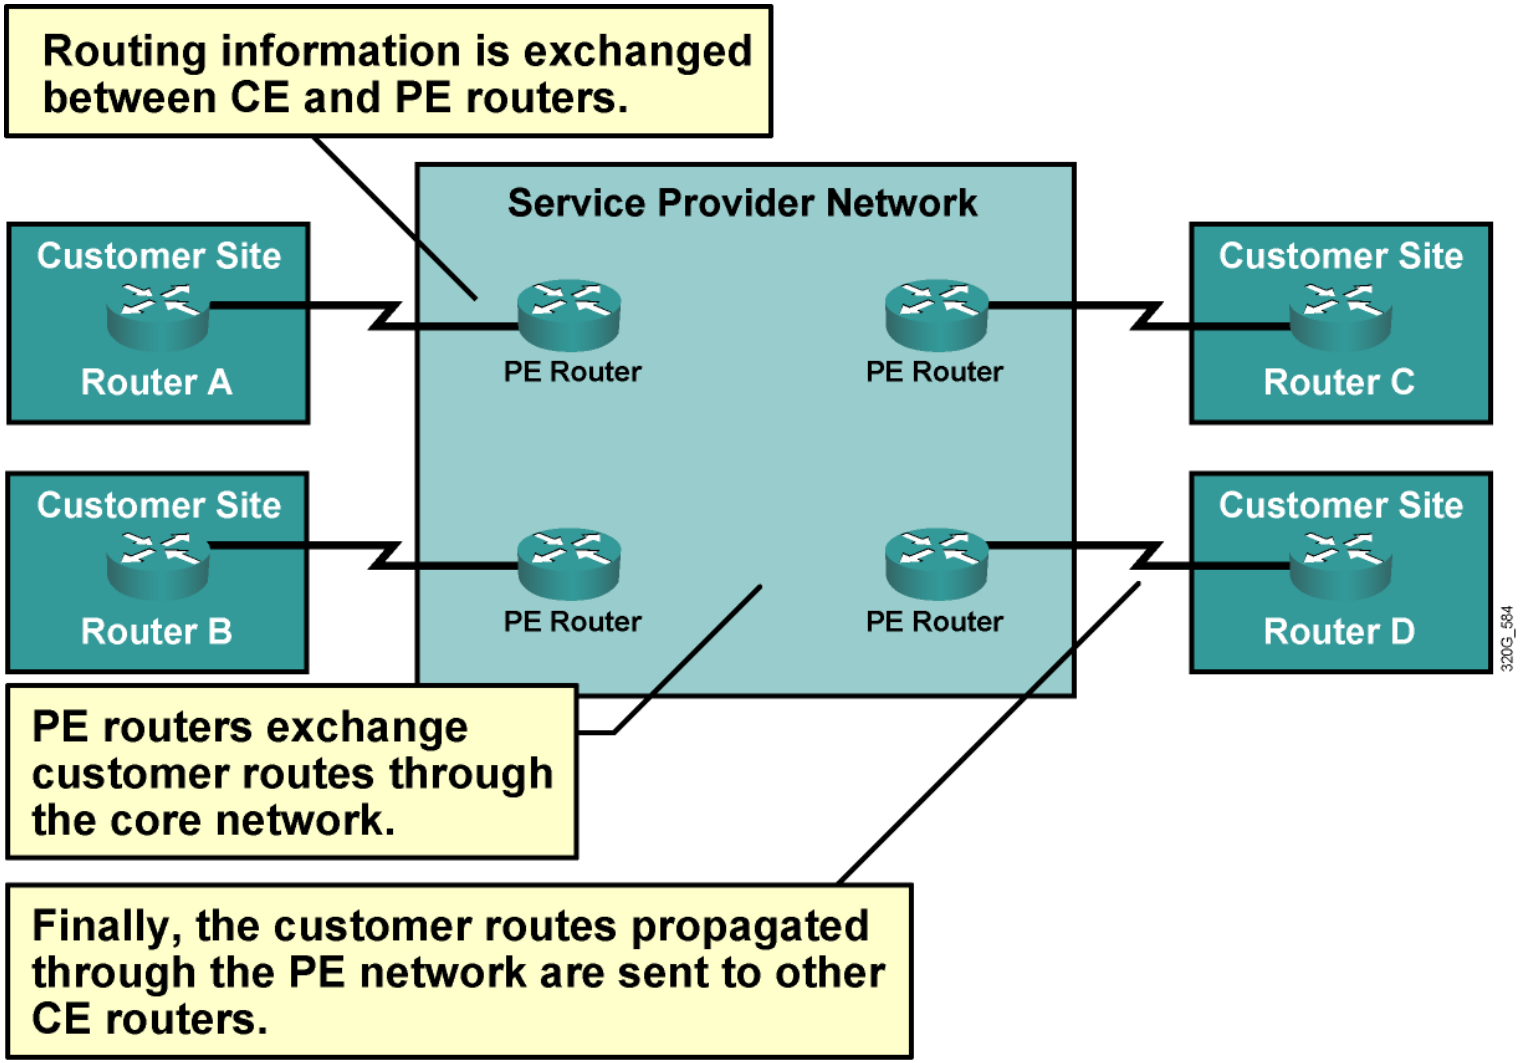

Benefits and Disadvantages of the Overlay VPN Implementation Model

- Benefits:
  - Well-known and easy to implement.
  - Service provider does not participate in customer routing.
  - Customer network and service provider network are well-isolated.
- Disadvantages:
  - Implementing optimum routing requires a full mesh of VCs.
  - VCs have to be provisioned manually.
  - Bandwidth must be provisioned on a site-to-site basis.
  - Overlay VPNs always incur encapsulation overhead (IPsec or GRE).

Benefits and Disadvantages of the Peer-to-Peer VPN Implementation Model

- Benefits:
  - Guarantees optimum routing between customer sites.
  - Easier to provision an additional VPN.
  - Only sites are provisioned, not links between them.
- Disadvantages:
  - 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.

Non-SP Related Drawbacks of 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 overlay VPN and a peer-to-peer VPN models:
  - 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 (Cont.)

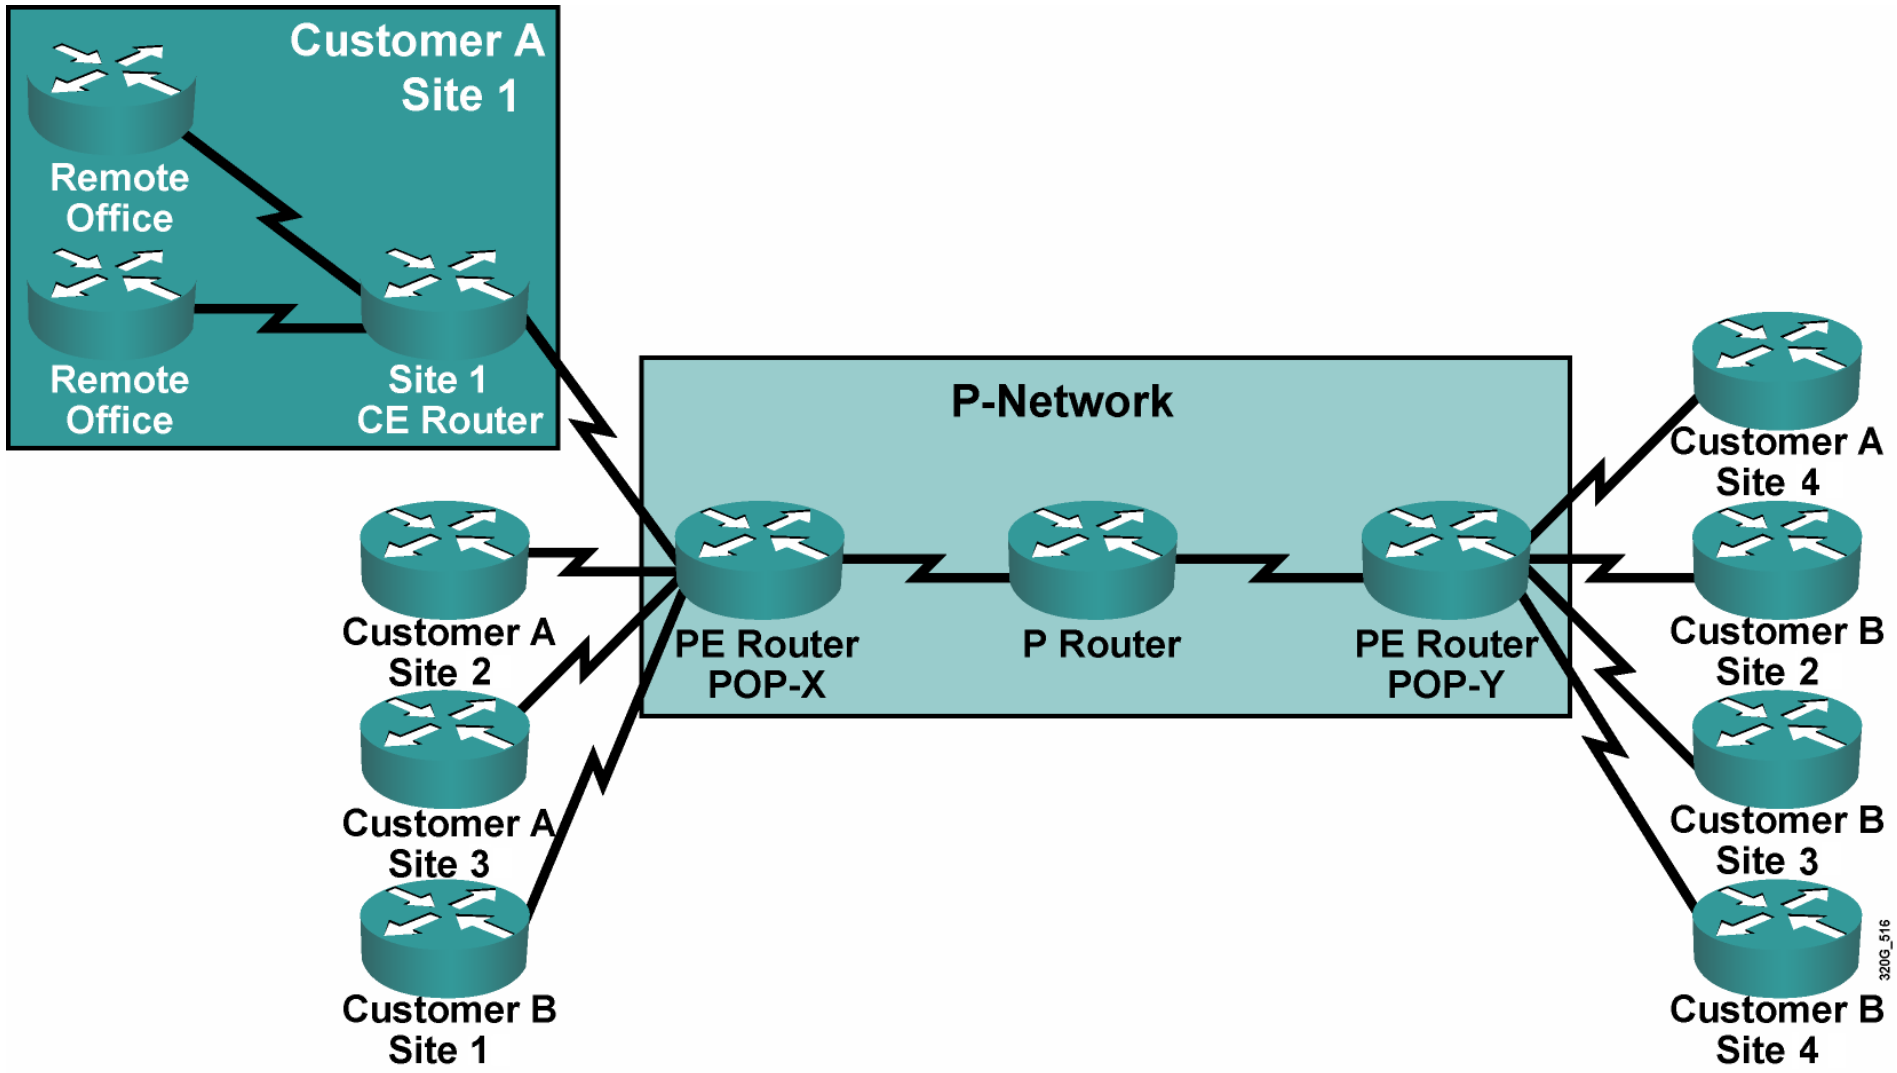

PE Router Architecture

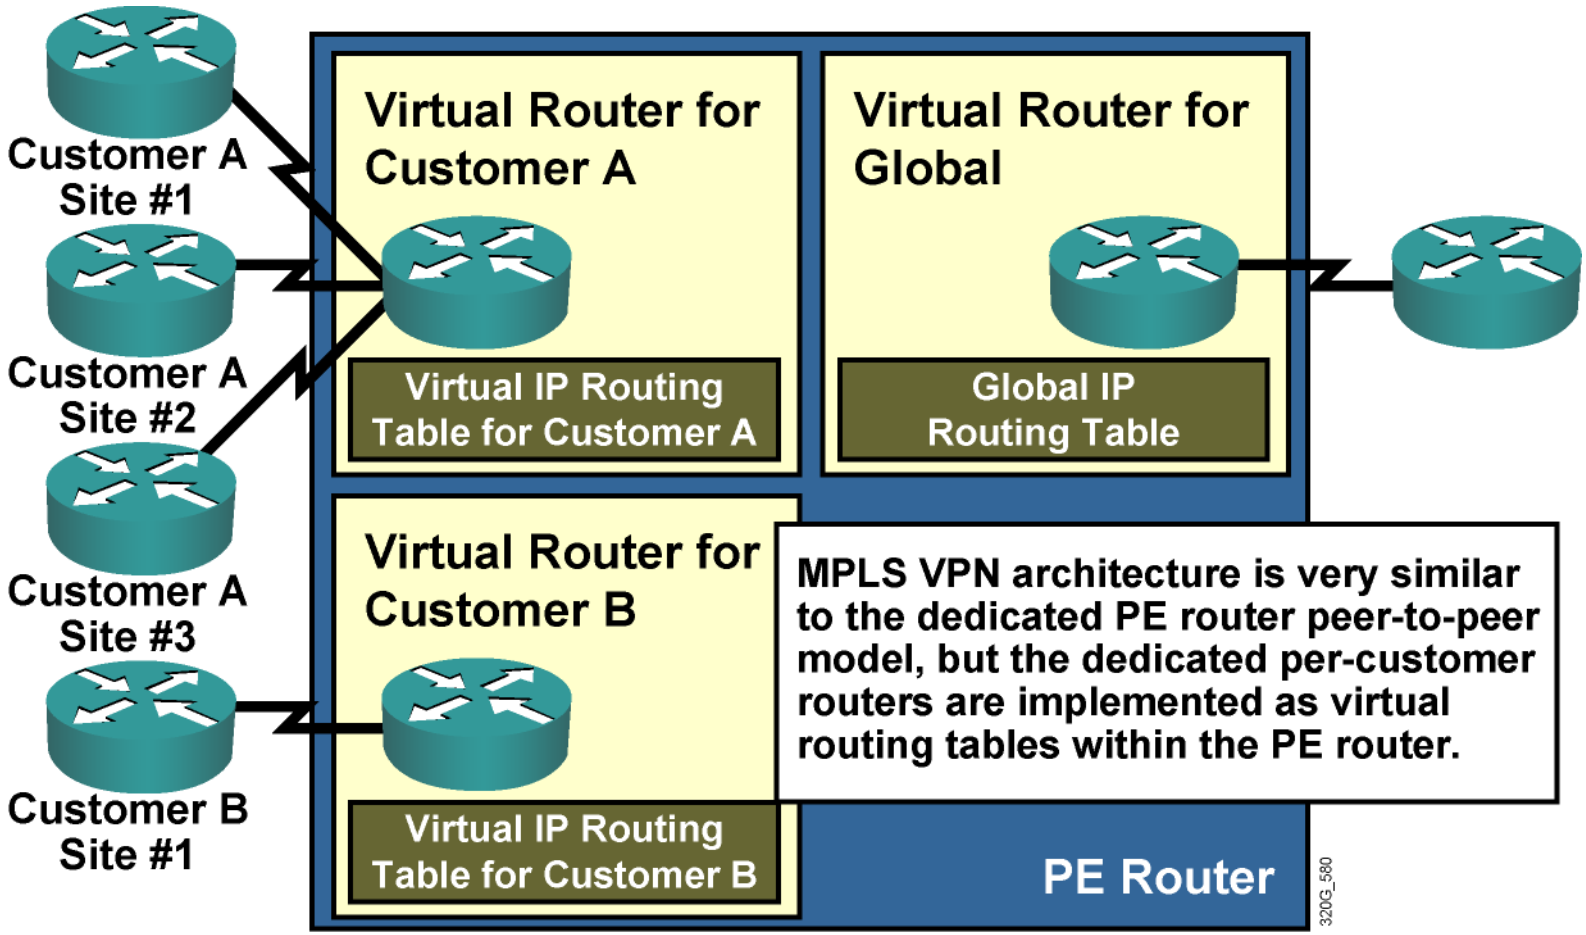

Propagation of Routing Information Across the P-Network

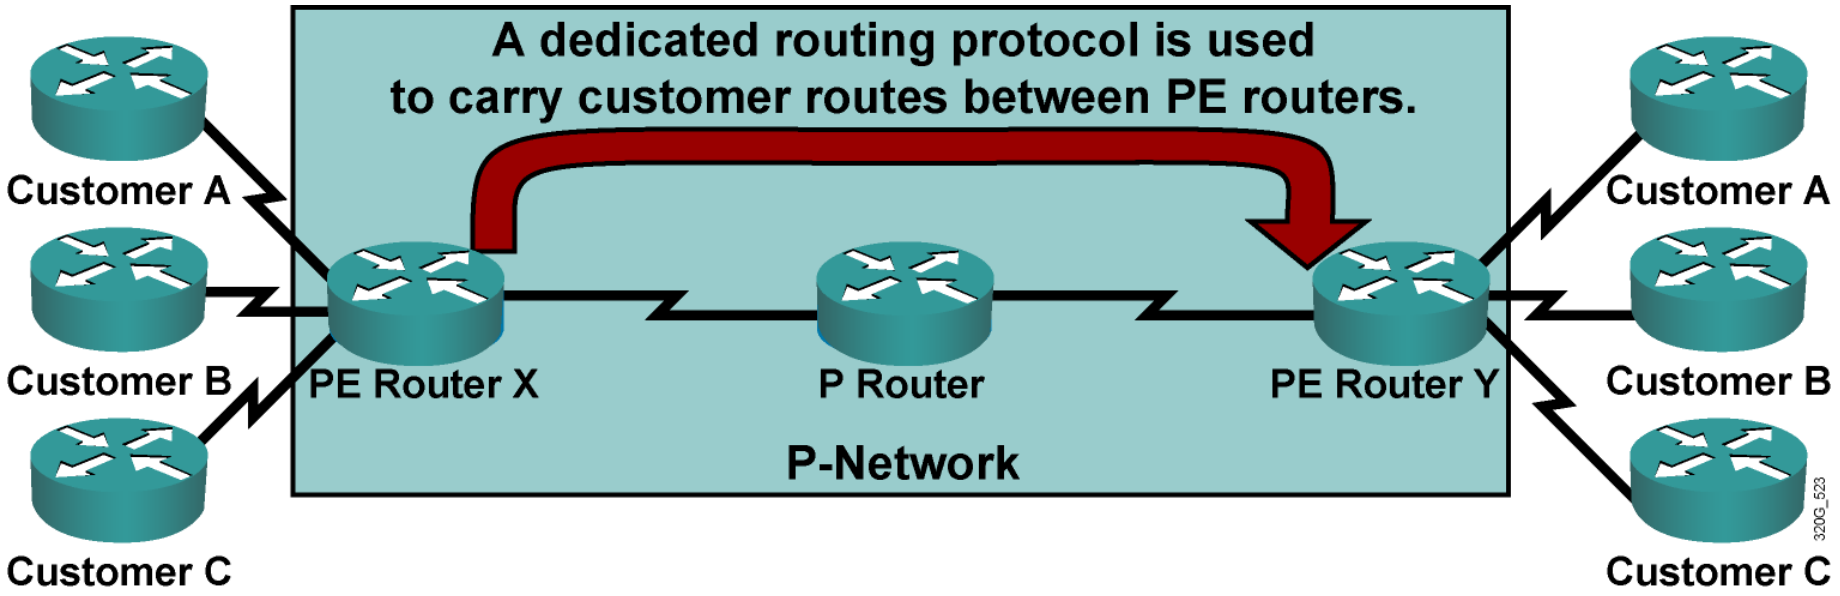

Routing Information Propagation Across P-Network

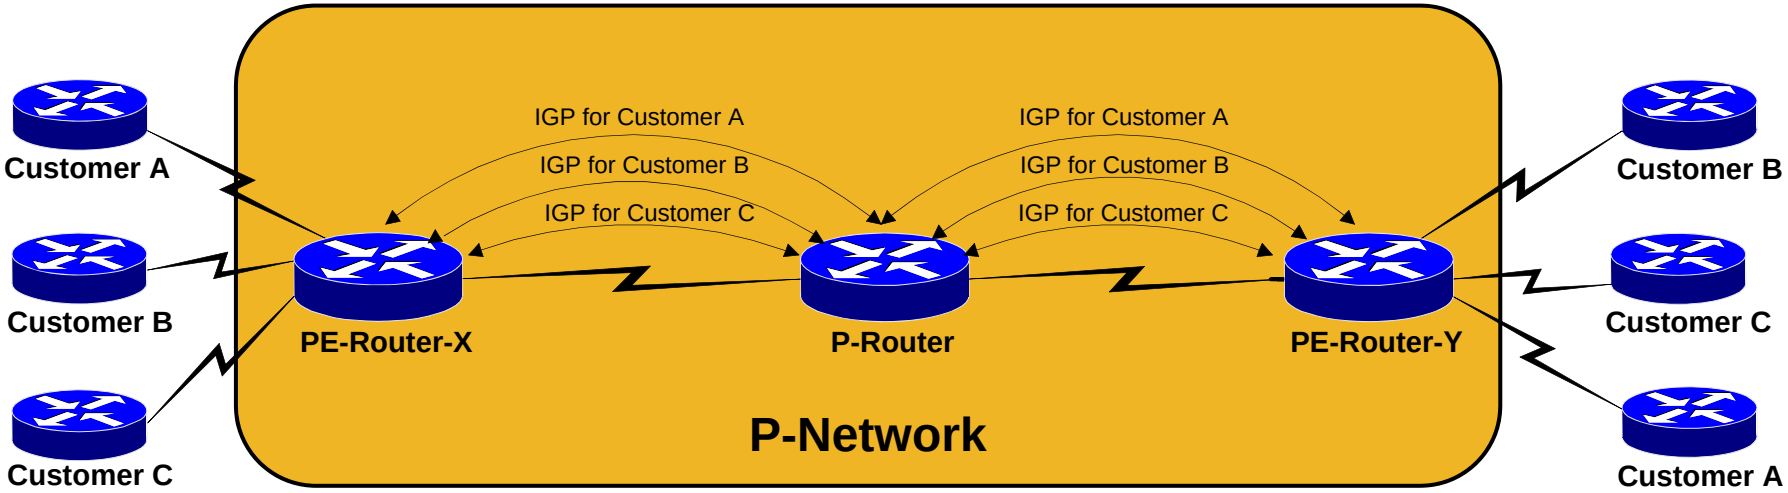

Q: How will PE routers exchange customer routing information?

A1: Run a dedicated IGP for each customer across P-network.

Wrong answer:

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

Routing Information Propagation Across P-Network

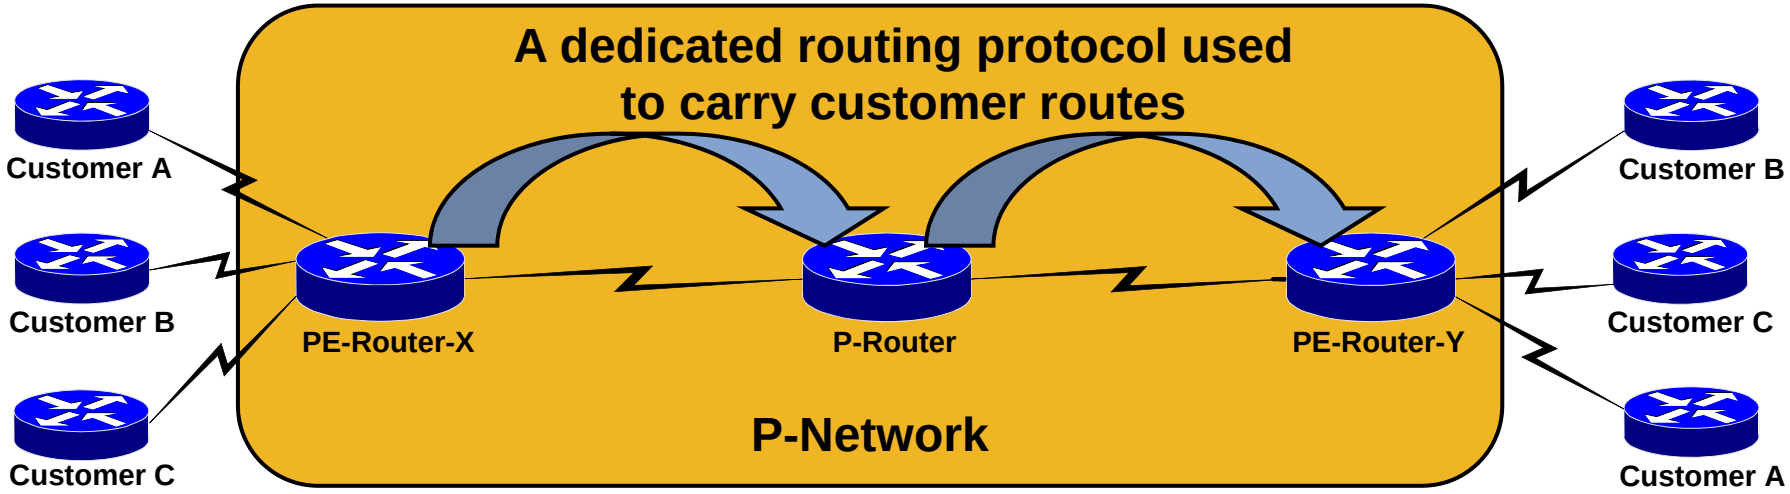

Q: How will PE routers exchange customer routing information?

A2: 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 routers.

Routing Information Propagation Across P-Network

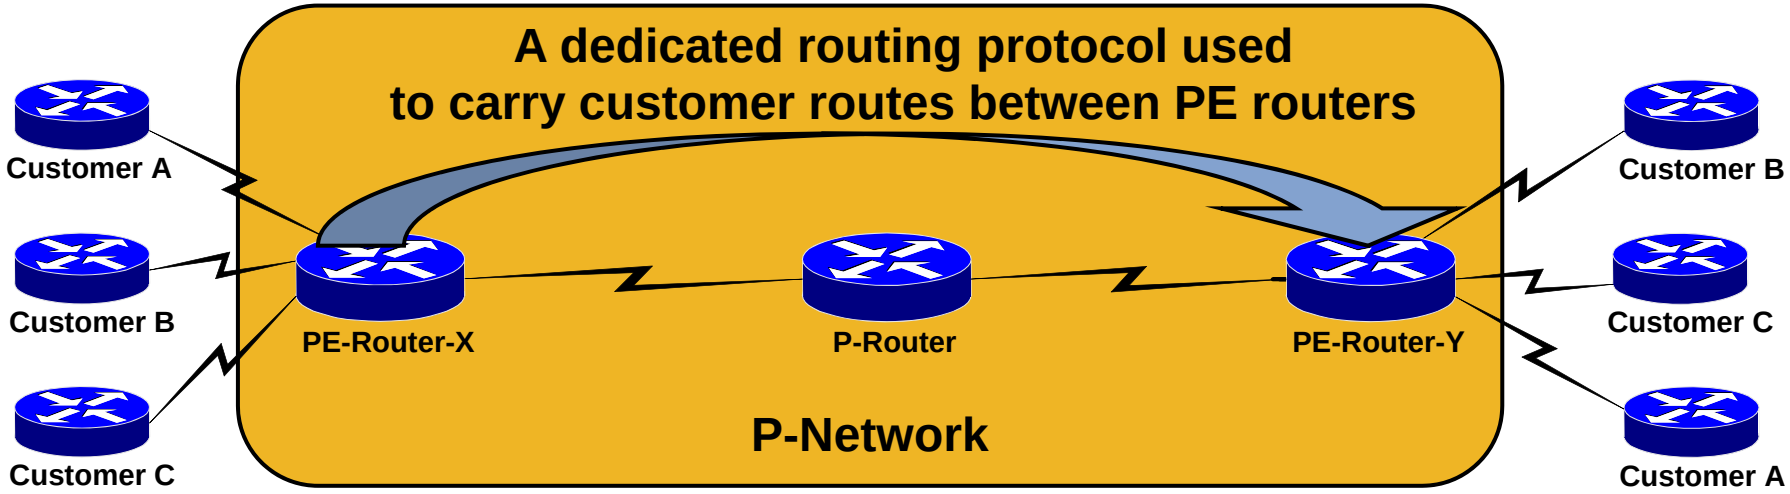

Q: How will PE routers exchange customer routing information?

A3: 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.

Routing Information Propagation Across P-Network

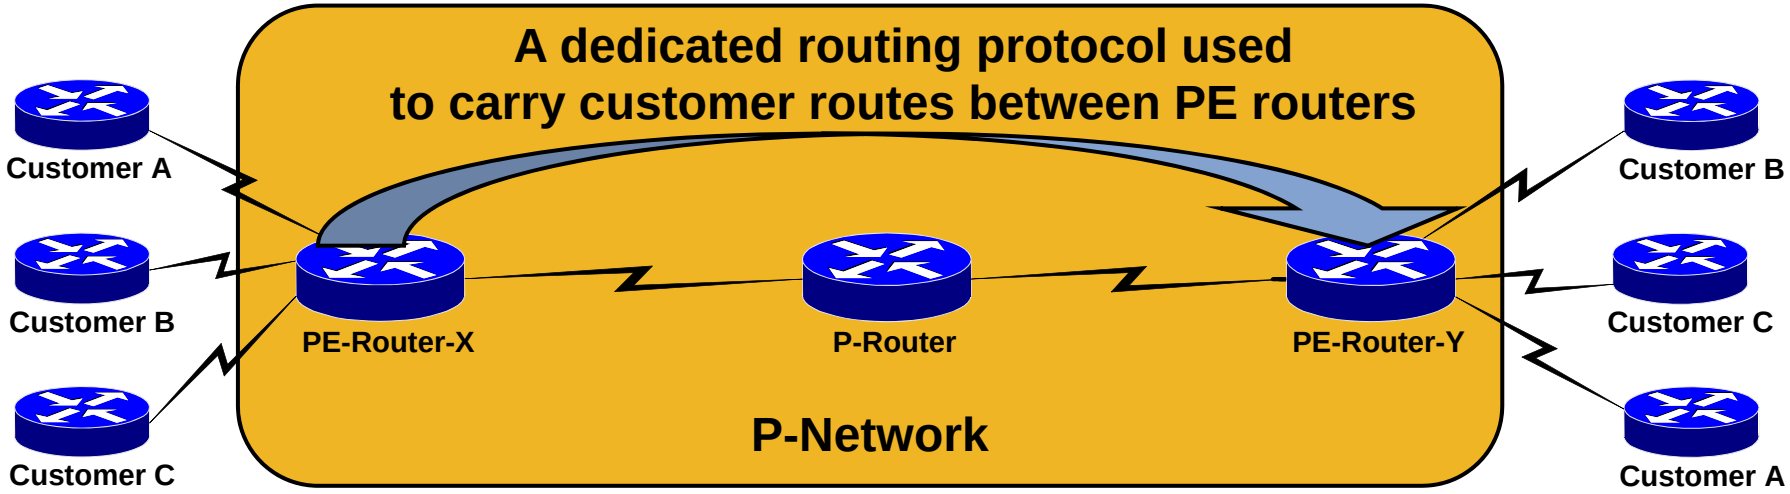

Q: Which protocol can be used to carry customer routes between PE-routers?

A: 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

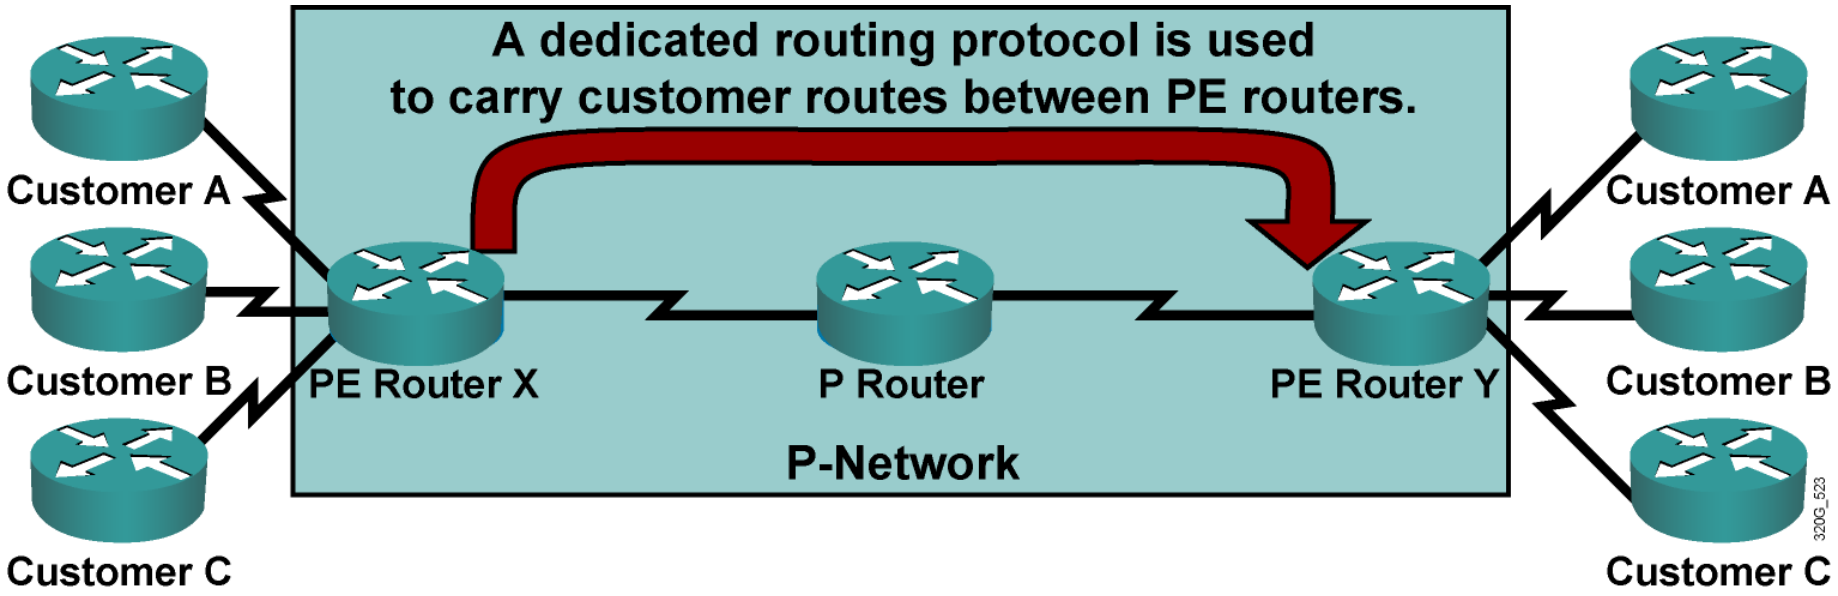

- The number of customer routes can be very large; BGP is the only routing protocol that can scale to such a number.
- BGP is used to exchange customer routes directly between PE routers.

Route Distinguishers

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.**

- The 64-bit RD is prepended to an IPv4 address to make the address 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 (MPBGP).

Distinguishing Routes: Steps 1, 2, and 3

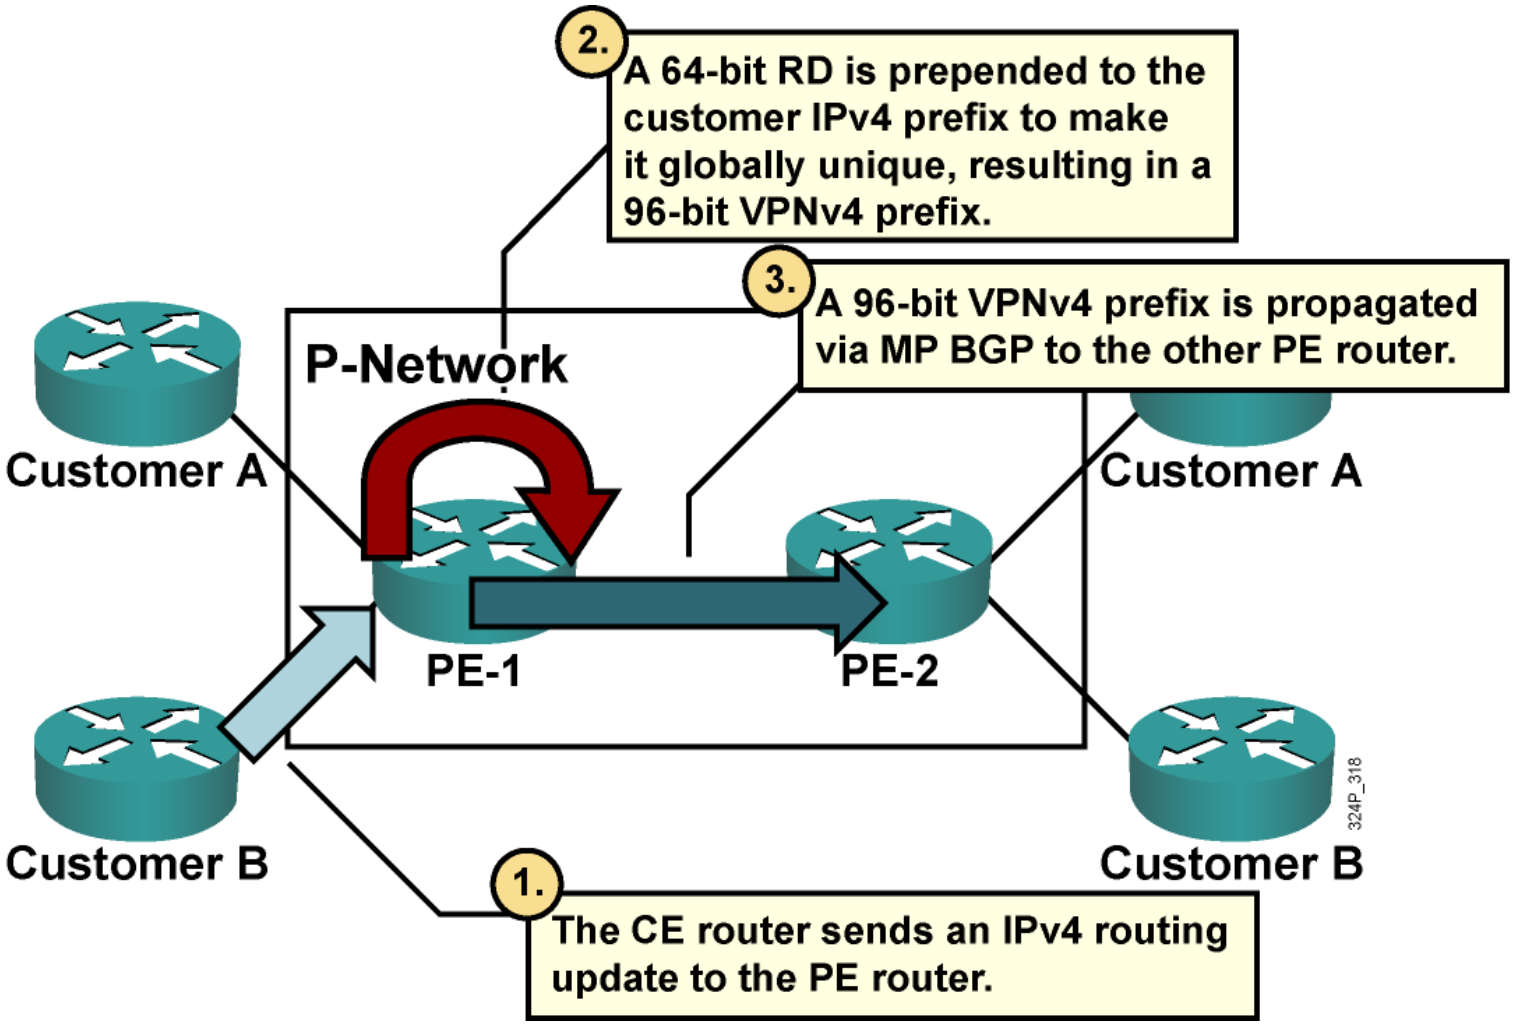

Distinguishing Routes: Steps 4 and 5

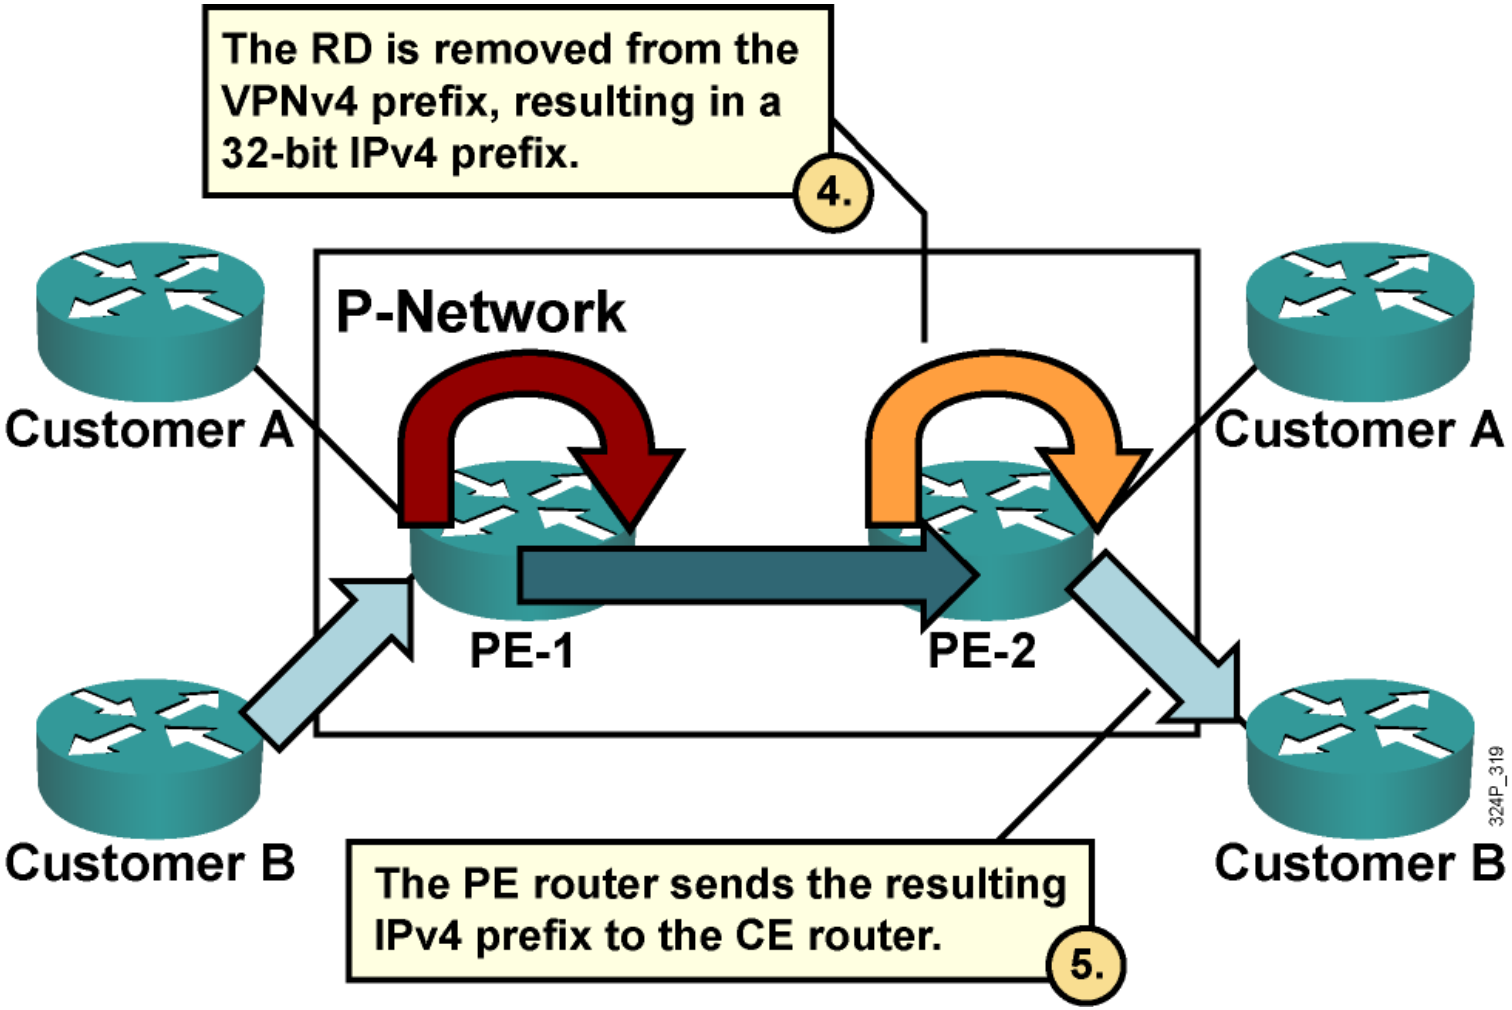

Using RDs in an MPLS VPN

- The RD has no special meaning.
- The RD is used only to make potentially overlapping IPv4 addresses globally unique.
- This design cannot support all topologies that are required by the customer.

VoIP Service on an MPLS VPN

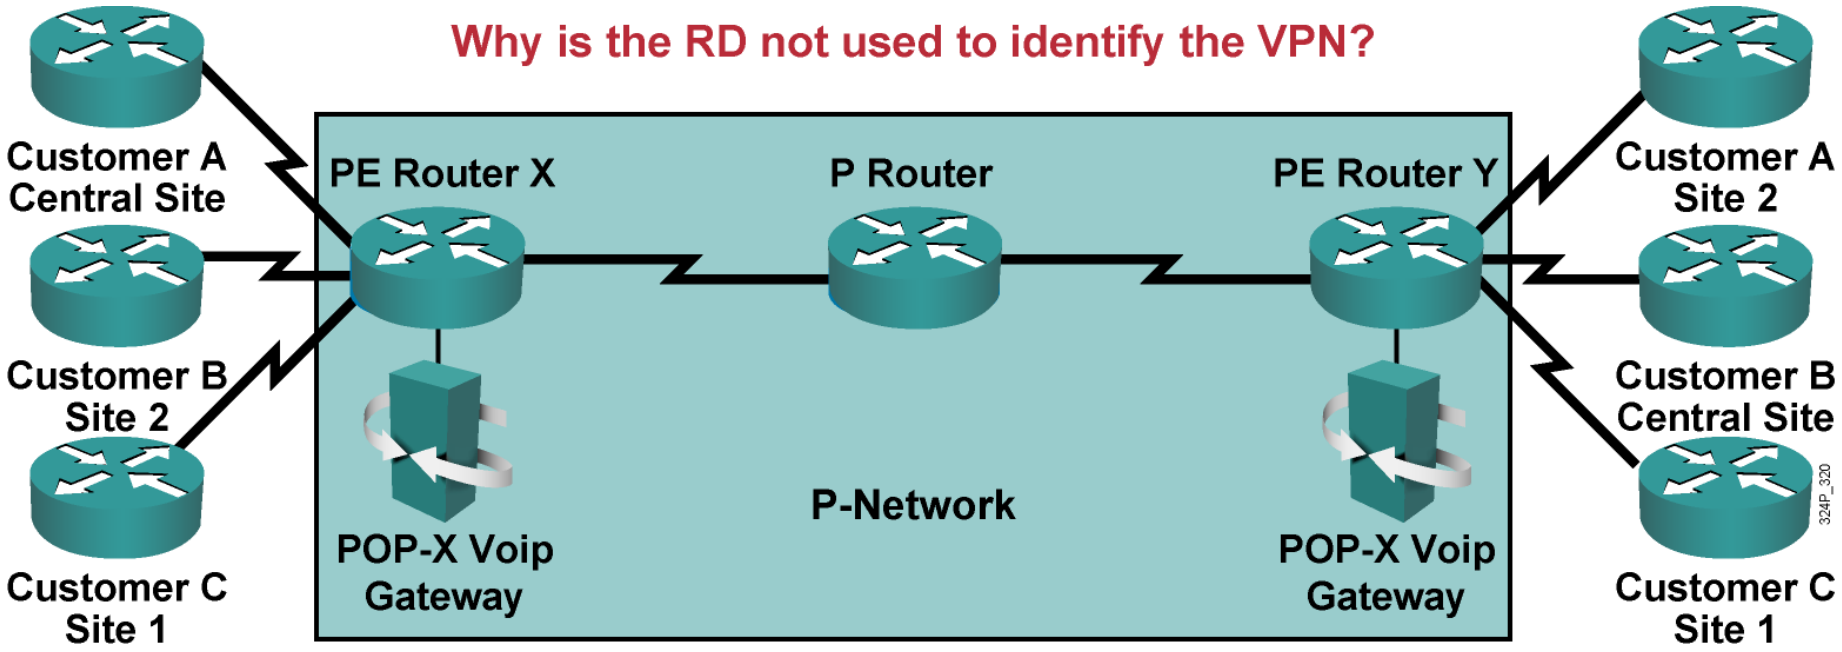

- Requirements:

  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.

Connectivity Requirements for VoIP Service

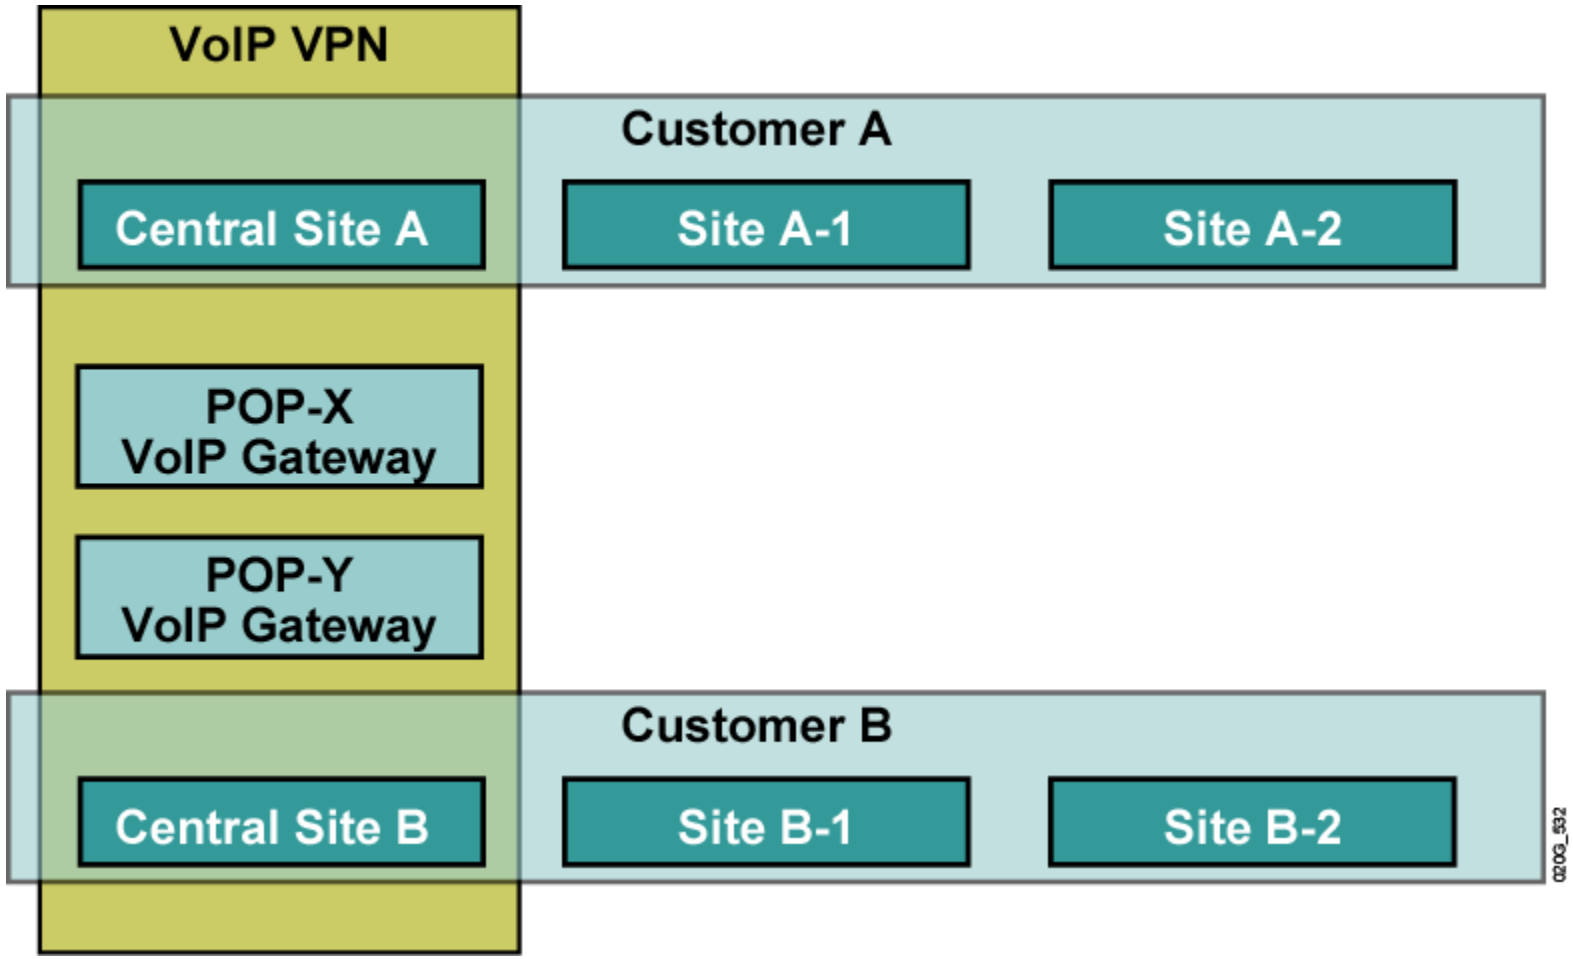

Route Targets

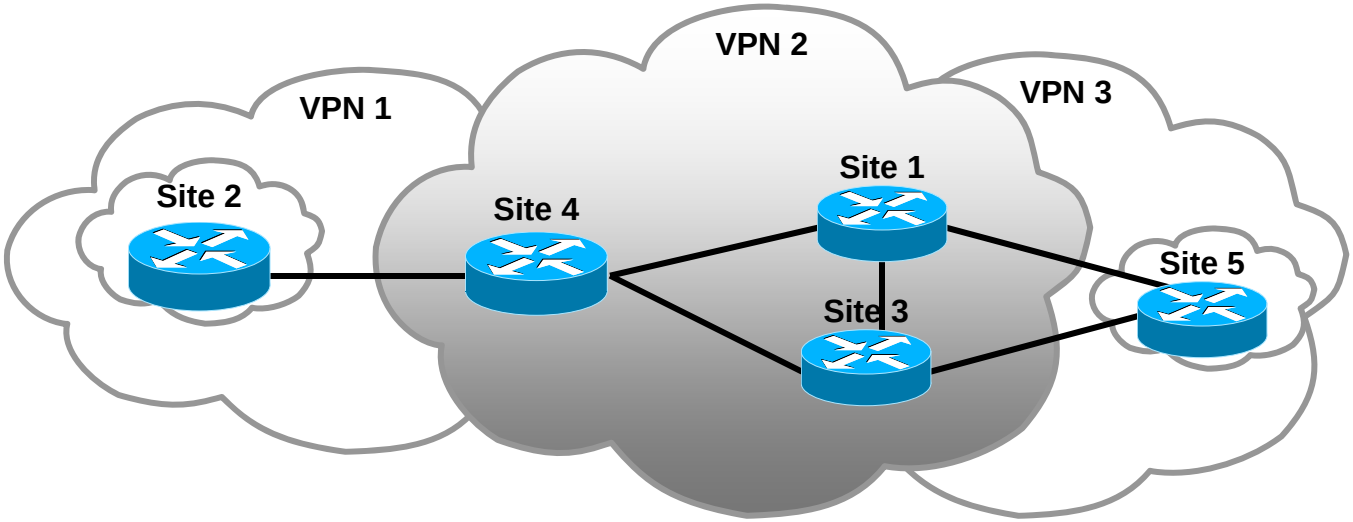

- Some sites participate in more than one VPN.
- The RD cannot identify participation in more than one VPN.
- RTs were introduced in the MPLS VPN architecture to support complex VPN topologies.
- RTs are additional attributes that attach to VPNv4 BGP routes to indicate VPN membership.

How Do RTs Work?

- Export RTs:
  - Identify VPN membership
  - Append to the customer route when the route is converted into a VPNv4 route
- Import RTs:
  - Associate with each virtual routing table
  - Select routes inserted into the virtual routing table

MPLS VPN Routing Criteria

- Designers imposed these criteria on MPLS VPNs:
  - CE routers can only run standard IP routing software.
  - Only PE routers need to support MPLS VPN services and Internet routing.
  - P routers have no VPN routes.

MPLS VPN Routing: CE Router Perspective

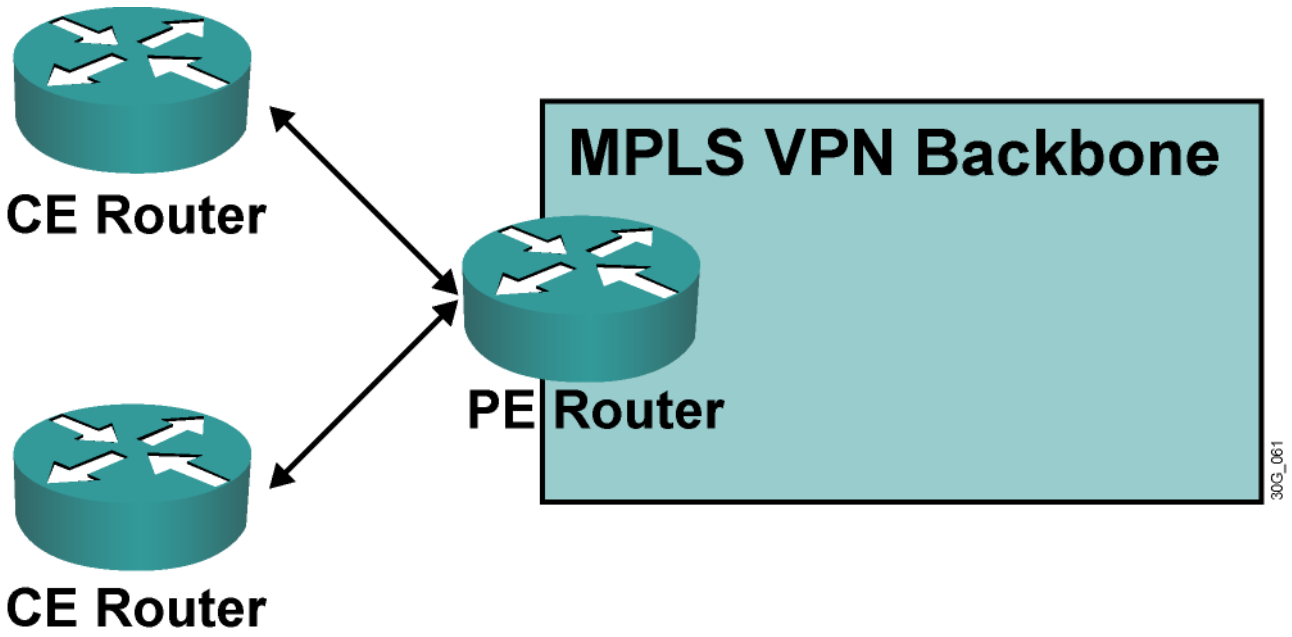

- The CE routers run standard IP routing software and exchange routing updates with the PE router.
- The PE router appears as another router in the C-network.

PE-CE Routing Protocols

- PE-CE routing protocols are configured for individual VRFs.
- Supported protocols include BGP, OSPF, static, RIP, and EIGRP.
- Routing configuration on the CE router has no VRF information.

MPLS VPN Routing: Overall Customer Perspective

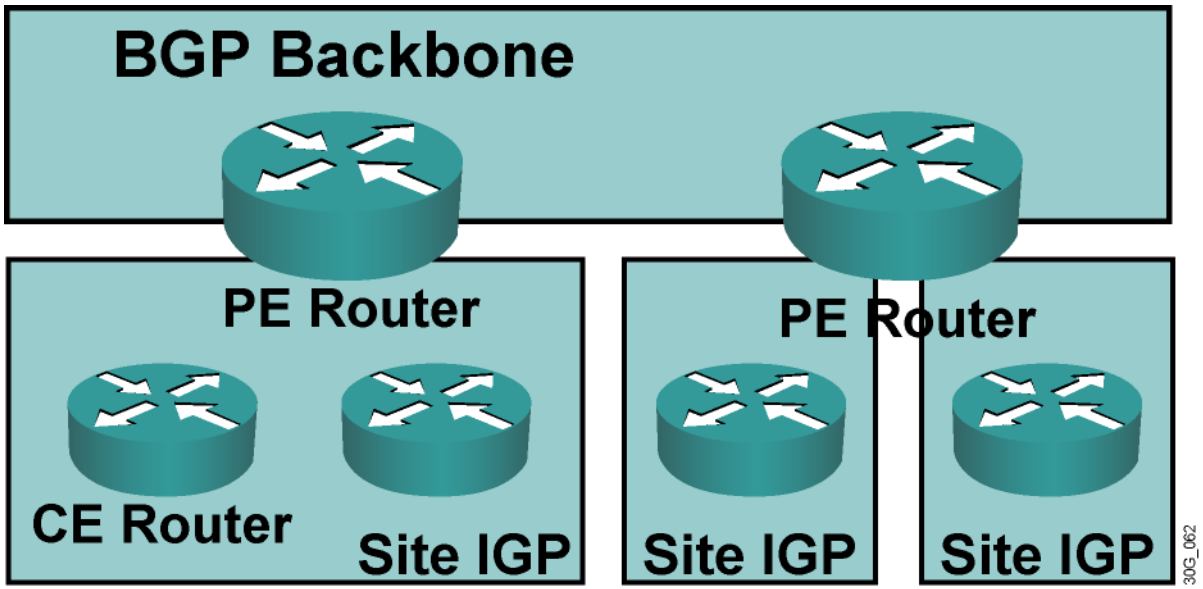

- To the customer, the PE routers appear as core routers that are connected via a BGP backbone.
- The usual BGP and IGP design rules apply.
- The P routers are hidden from the customer.

MPLS VPN Routing: P Router Perspective

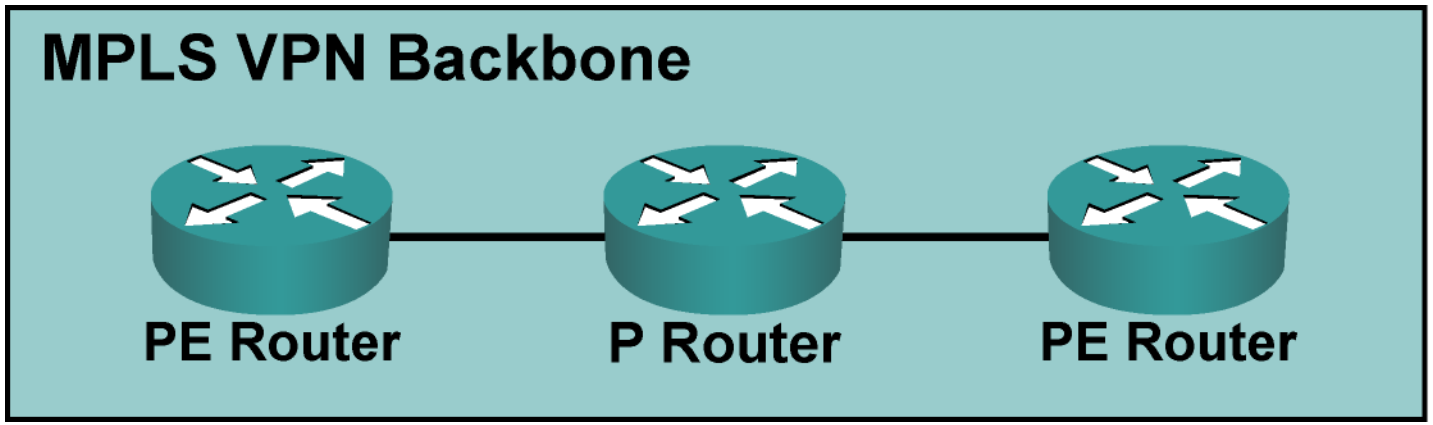

- P routers perform as follows:

  Do not participate in MPLS VPN routing and do not carry VPN routes

  Run backbone IGP with the PE routers and exchange information about global subnetworks (core links and loopbacks)

MPLS VPN Routing: PE Router Perspective

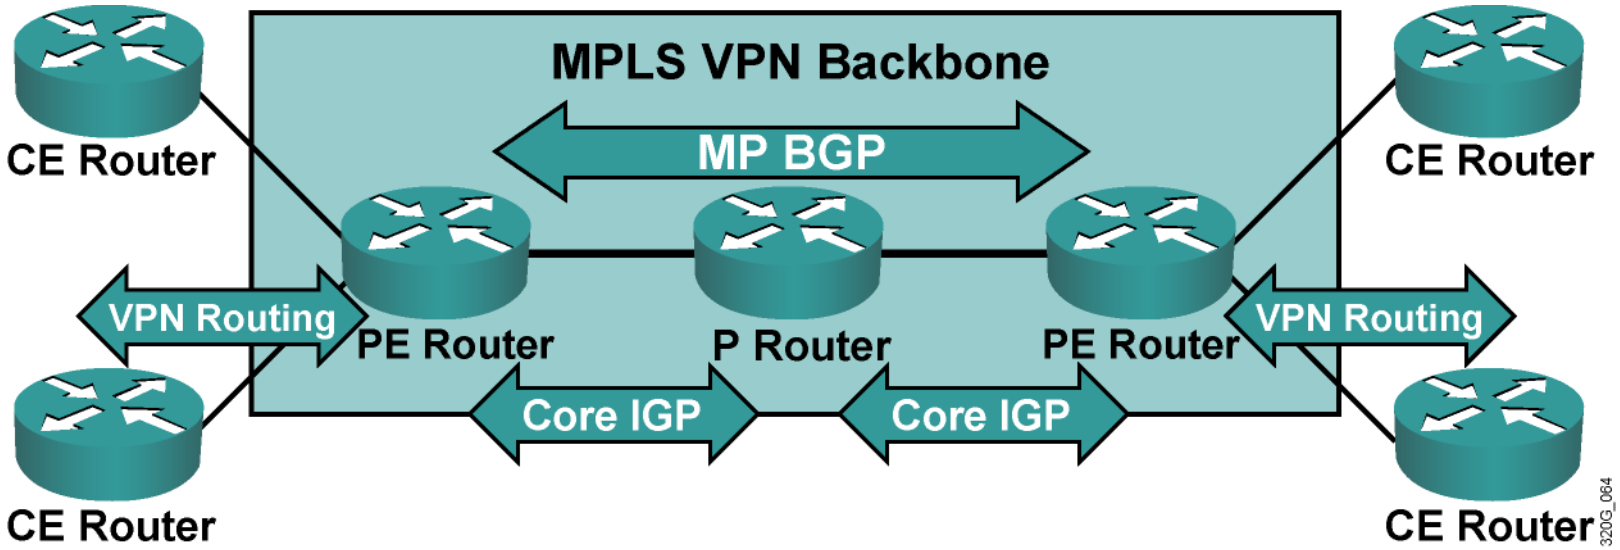

- PE routers exchange the following:
  - VPN routes with CE routers via per-VPN routing protocols
  - Core routes with P routers and PE routers via core IGP
  - VPNv4 routes with other PE routers via MPBGP sessions

End-to-End Routing Information Flow

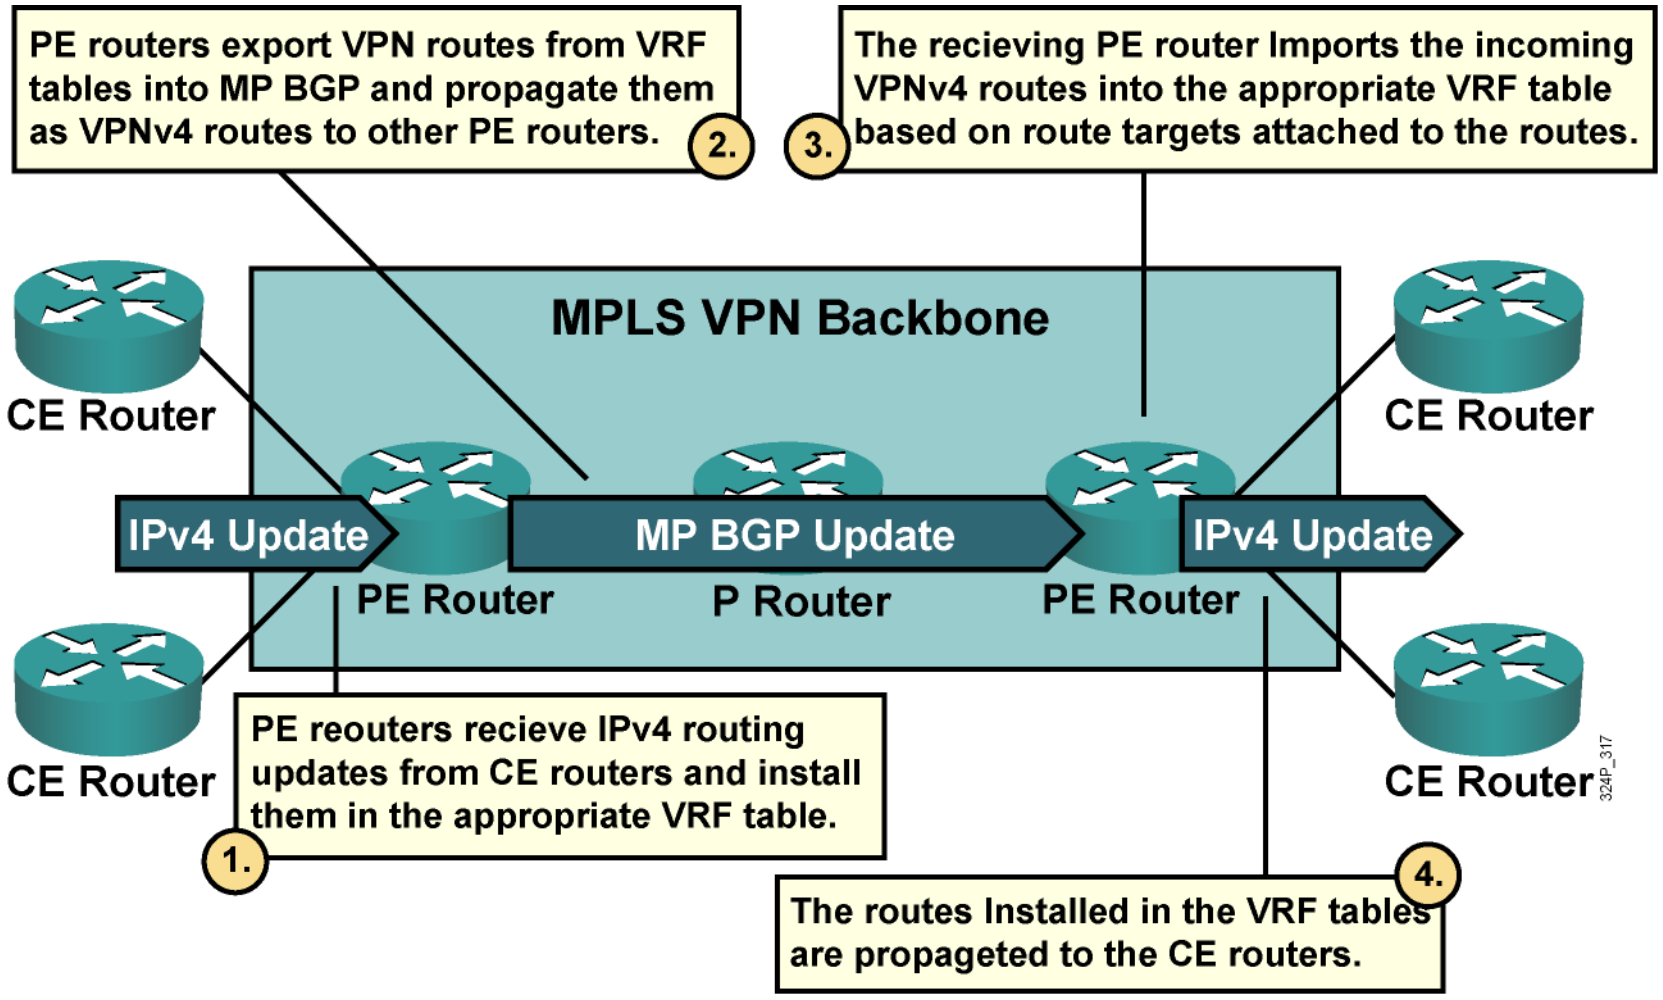

VPN Label Propagation

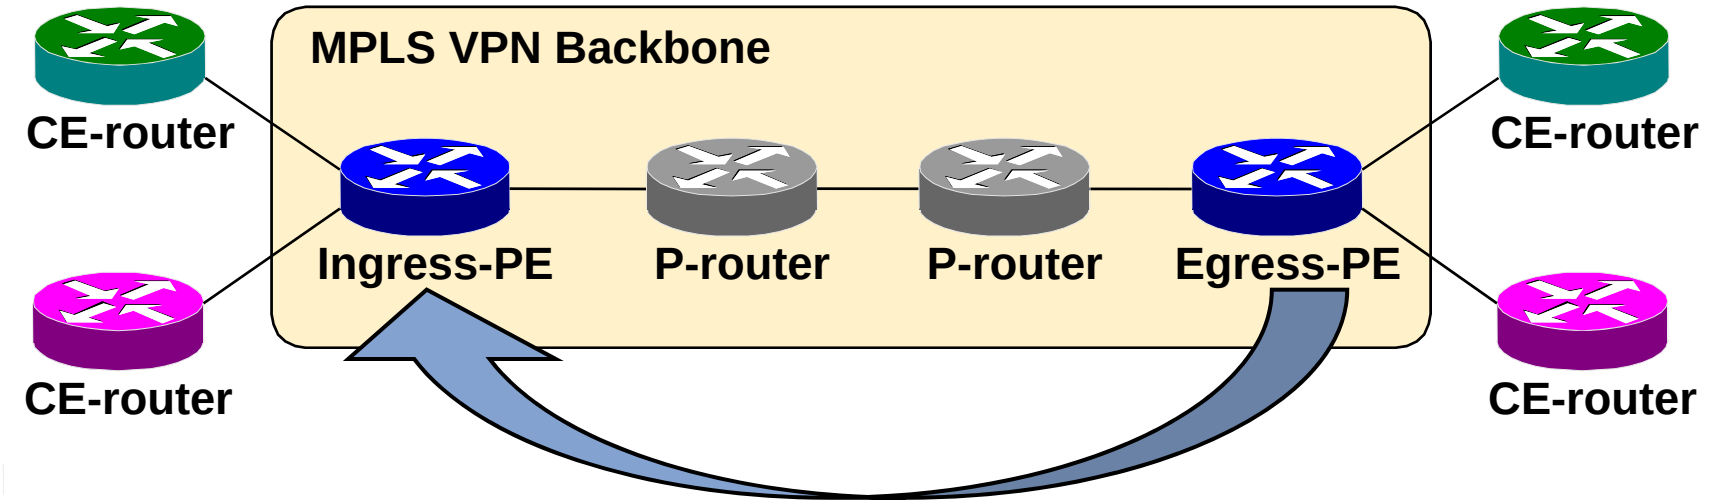

Q: How will the ingress PE-router get the second label in the label stack from the egress PE-router?

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

VPN Label Propagation

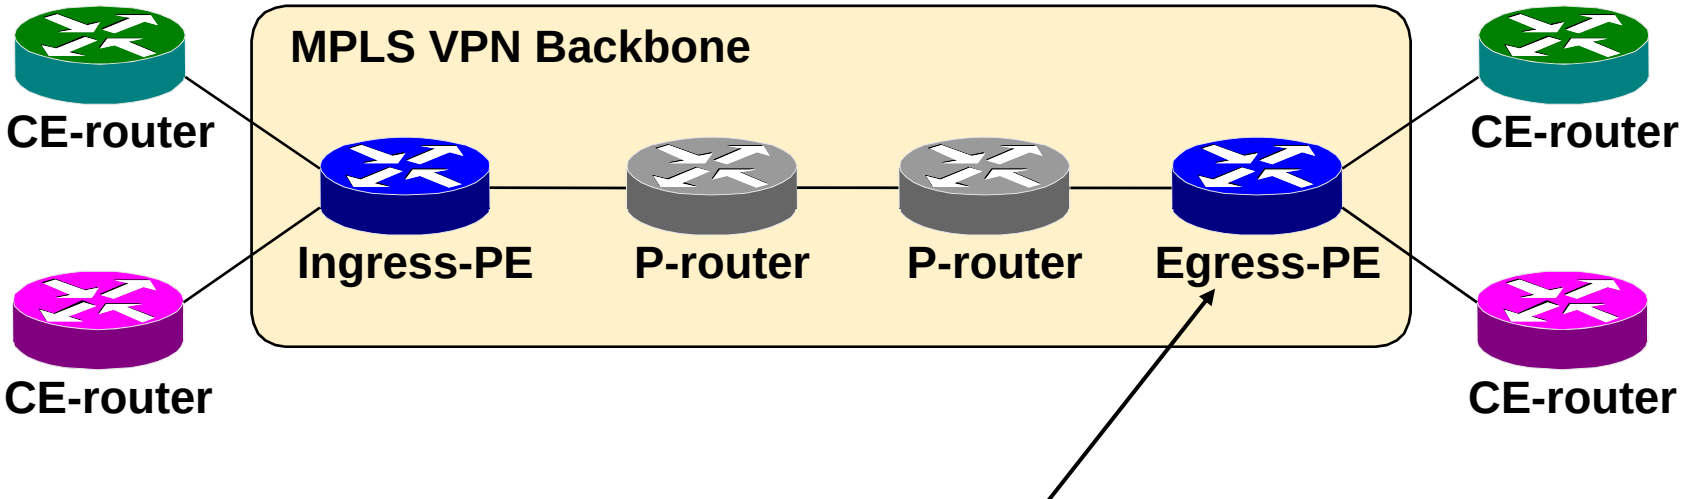

Step #1: VPN label is assigned to every VPN route by the egress PE router

```
Egress-PE#show tag-switching forwarding vrf SiteA2
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
26     Aggregate   150.1.31.36/30[V] 0
37     Untagged    203.1.2.1/32[V]   0          Se1/0.20    point2point
38     Untagged    203.1.20.0/24[V]  0          Se1/0.20    point2point
```

VPN Label Propagation

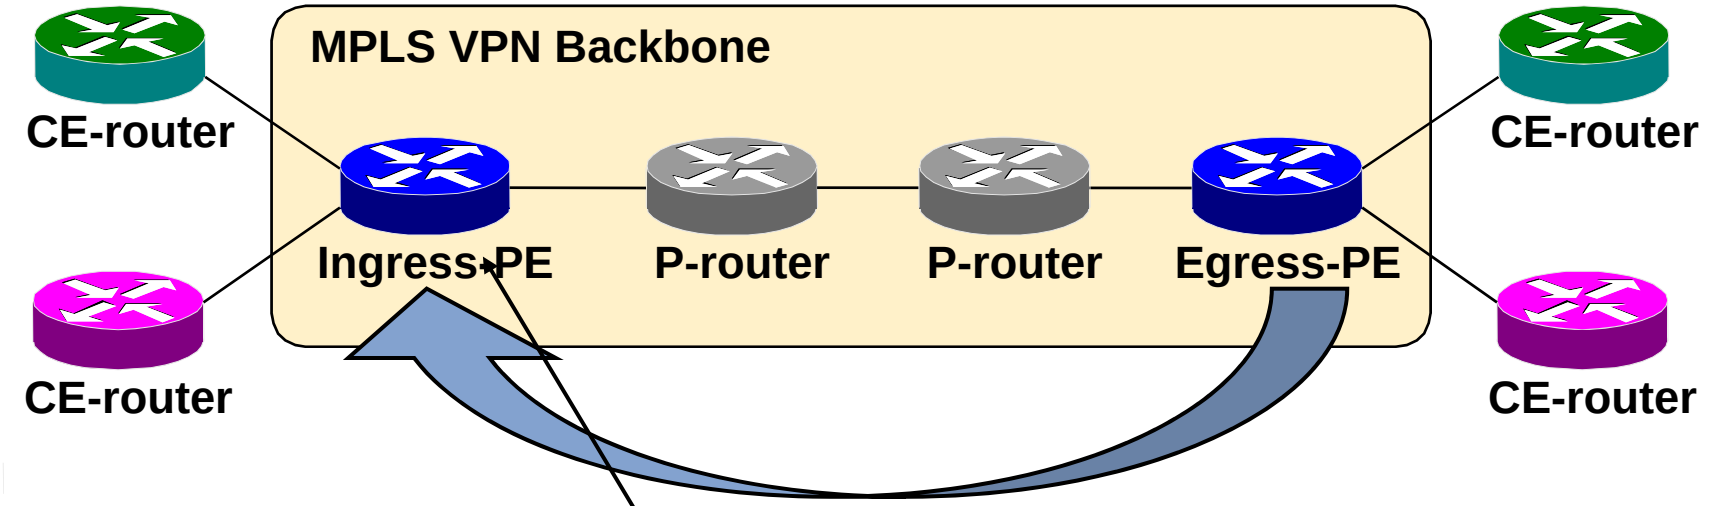

Step #2: VPN label is advertised to all other PE-routers in MP-BGP update

```
Ingress-PE#show ip bgp vpnv4 all tags
   Network                Next Hop      In tag/Out tag
Route Distinguisher: 100:1 (vrf1)
   12.0.0.0          10.20.0.60      26/notag
                     10.20.0.60      26/notag
   203.1.20.0        10.15.0.15      notag/38
```

VPN Label Propagation

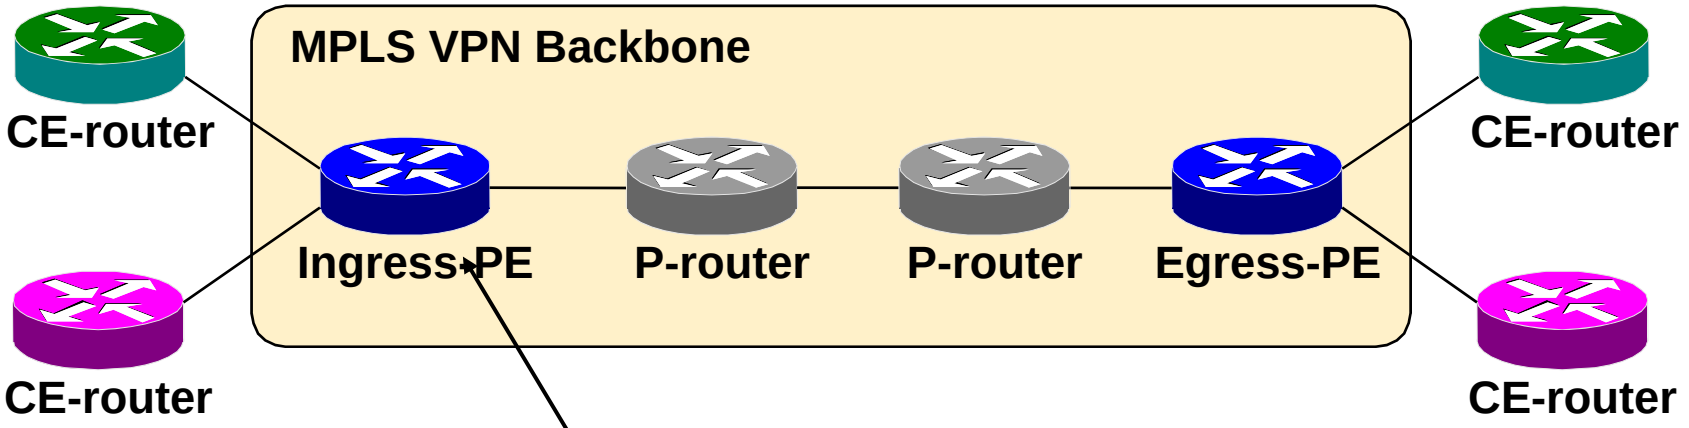

Step #3: Label stack is built in Virtual Forwarding table

```
Ingress-PE#show ip cef vrf Vrf1 203.1.20.0 detail
203.1.20.0/24, version 57, cached adjacency to Serial1/0.2
0 packets, 0 bytes
  tag information set
    local tag: VPN-route-head
    fast tag rewrite with Se1/0.2, point2point, tags imposed: {26 38}
  via 192.168.3.103, 0 dependencies, recursive
    next hop 192.168.3.10, Serial1/0.2 via 192.168.3.103/32
    valid cached adjacency
    tag rewrite with Se1/0.2, point2point, tags imposed: {26 38}
```

Summary

- VPNs allow you to use the shared infrastructure of a SP to implement your private networks. There are two implementation models: overlay and peer-to-peer.
- The MPLS VPN architecture offers SPs a peer-to-peer VPN architecture that combines the best features of overlay VPNs with the best features of peer-to-peer VPNs.
- MPLS VPNs use a 64-bit prefix called the route distinguisher (RD) to convert non-unique 32-bit customer IPv4 addresses into 96-bit unique addresses that can be transported.
- MPLS works by prepending packets with an MPLS header, containing one or more “labels.” This is called a label stack.