



Introduction to Virtual Private Networks

Outline

Overview

Traditional Router-Based Networks

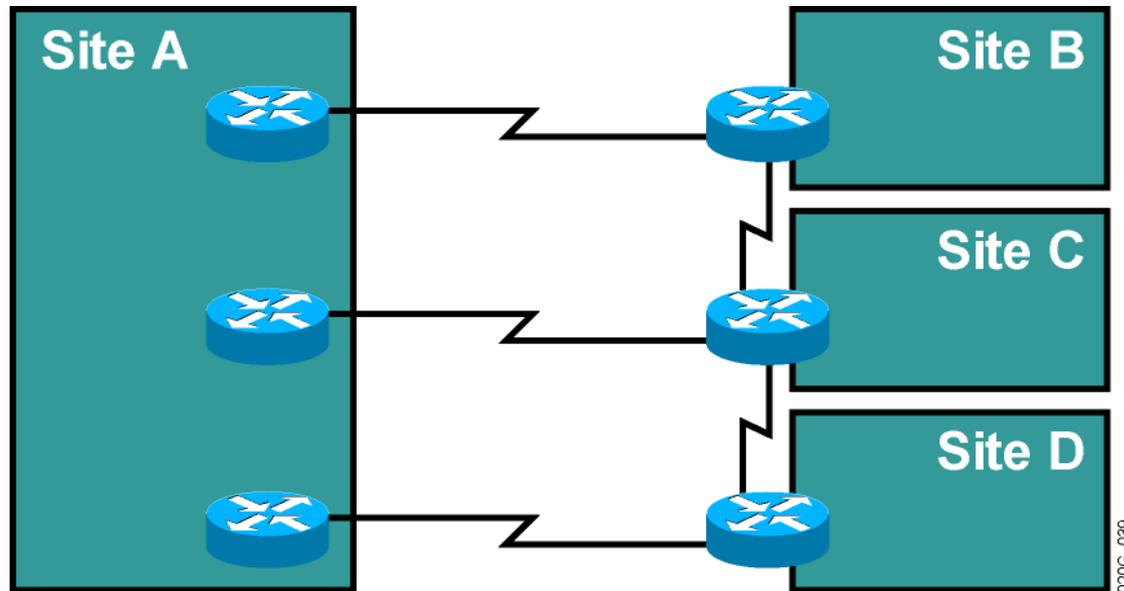
Virtual Private Networks

VPN Terminology

Switched WANs VPN Terminology

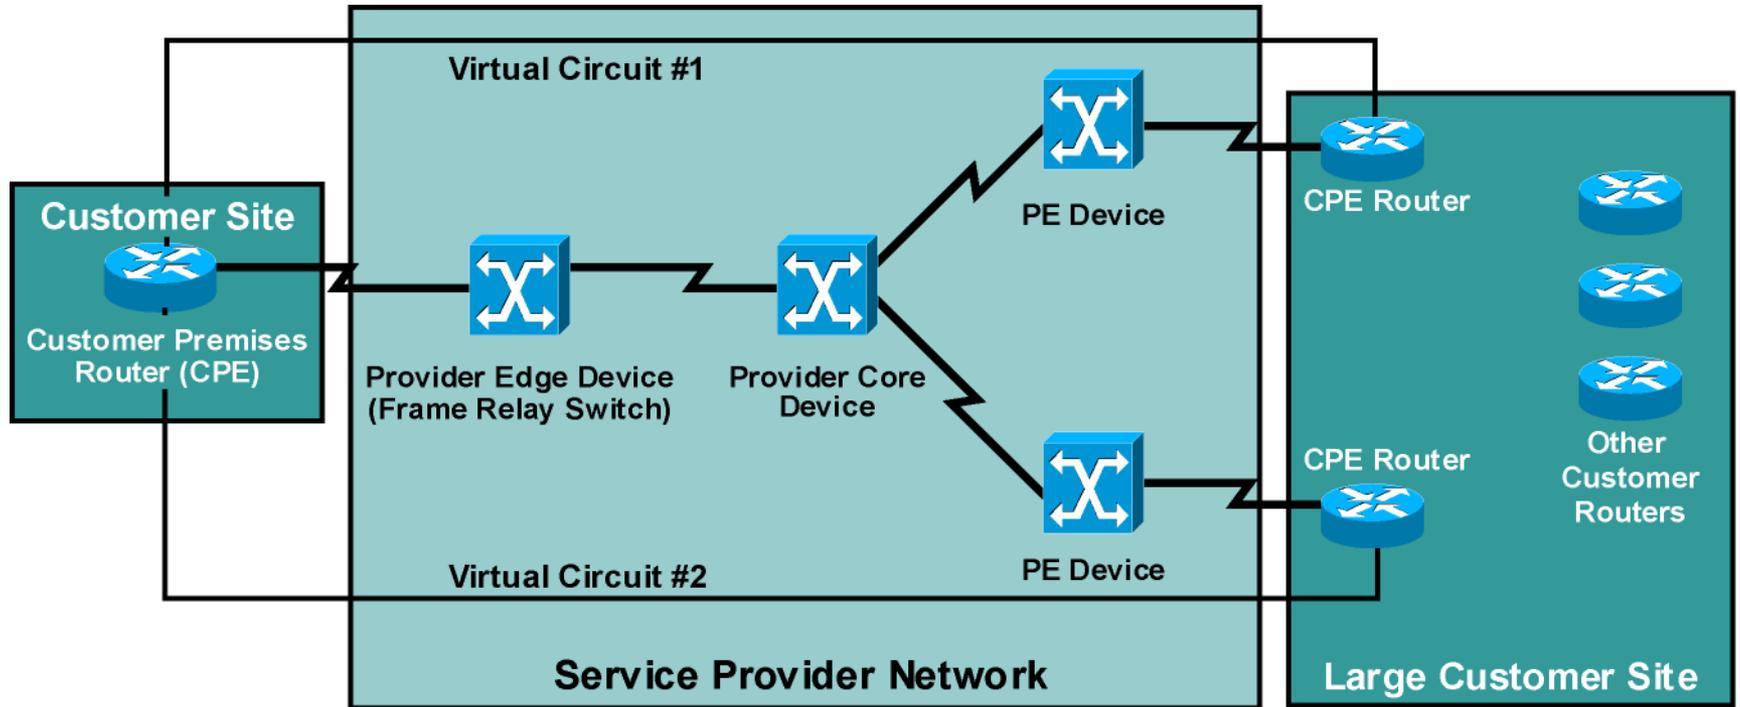
Lesson Summary

Traditional Router-Based Networks



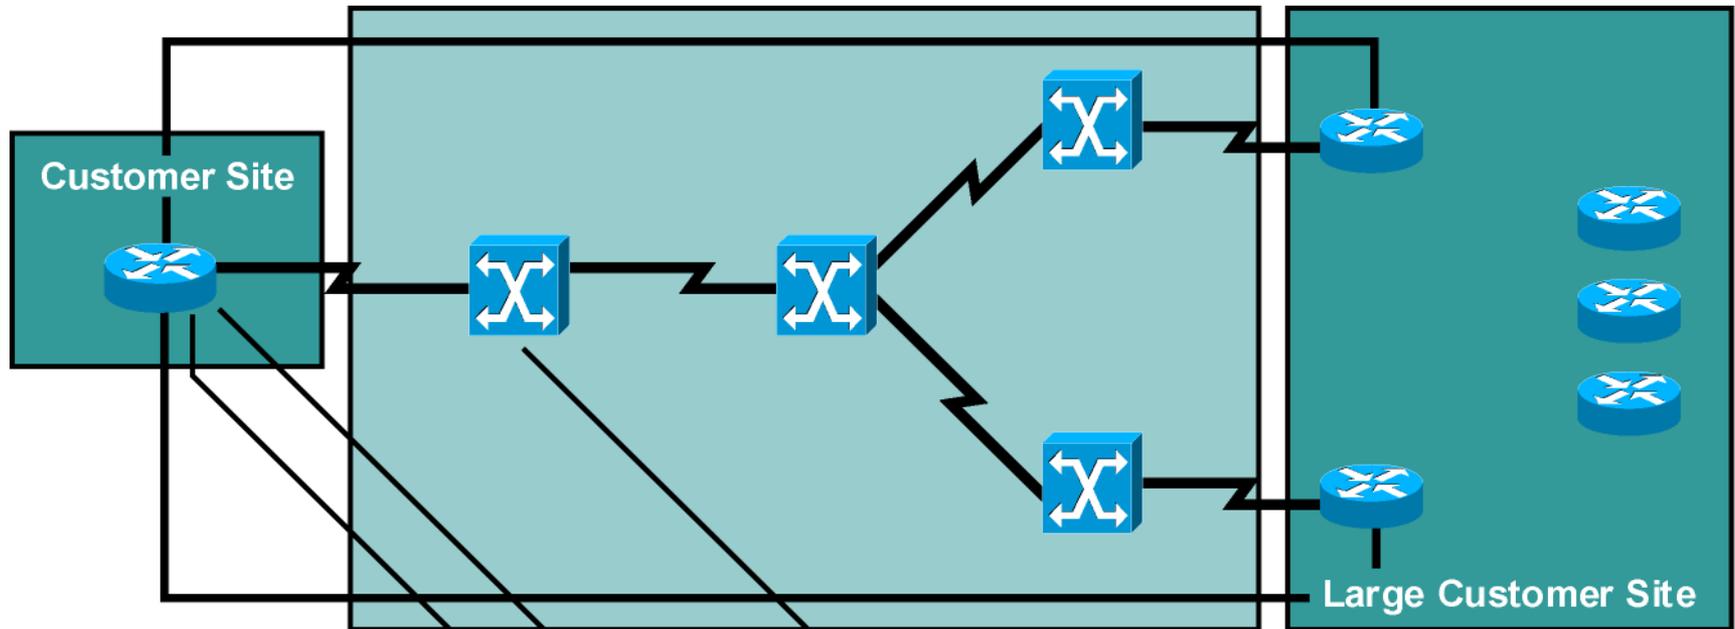
Traditional router-based networks connect customer sites through routers connected via dedicated point-to-point links.

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

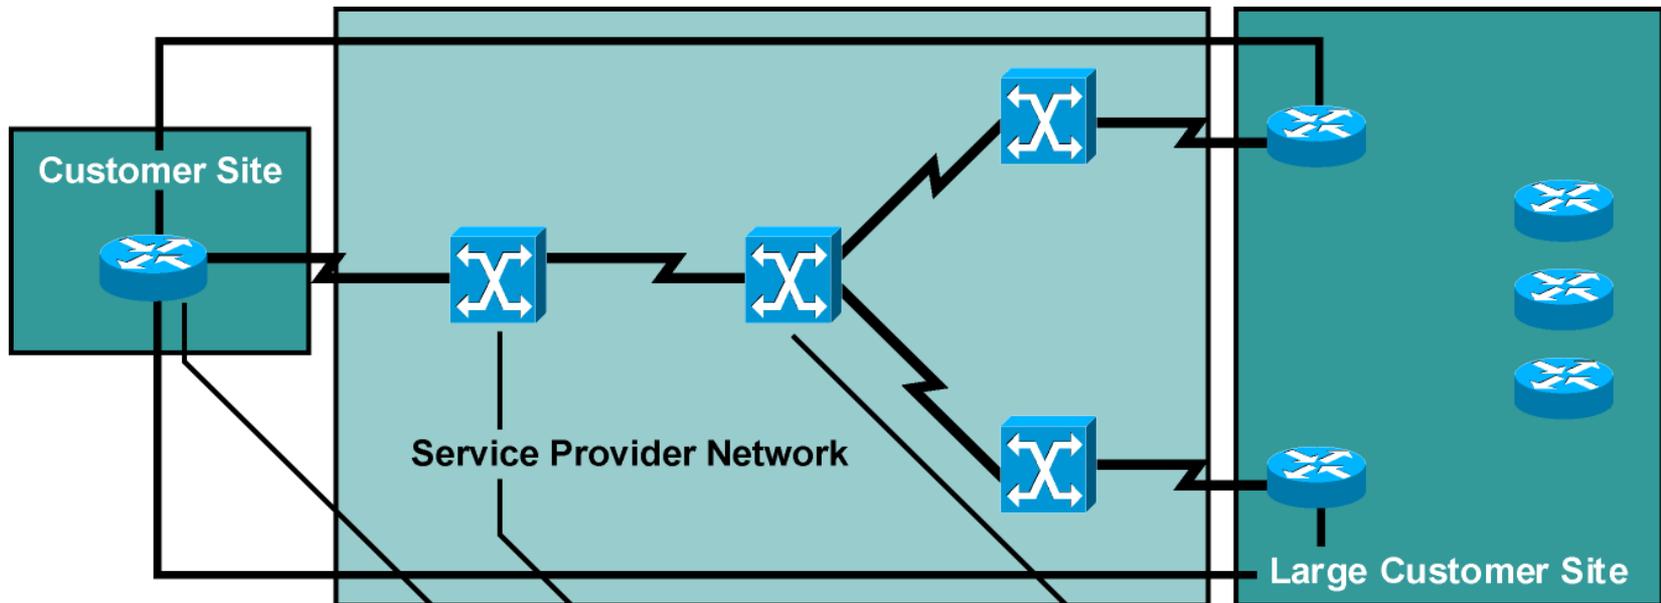


Provider network (P-network): The service provider infrastructure used to provide VPN services

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

Customer site: A contiguous part of the customer network (can encompass many physical locations)

VPN Terminology (Cont.)



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)**

Summary

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

VPNs replaced dedicated point-to-point links with emulated point-to-point links sharing a common infrastructure.

Device names based on their position in the network are as follows:

CE

PE

P

A PVC is established and is always active. An SVC is established through CE-PE signaling on demand from the CE device.



MPLS workshop

Overlay and Peer-to-Peer VPNs

Outline

Overview

VPN Implementation Technologies

Overlay VPNs

Peer-to-peer VPNs

Benefits of VPN Implementations

Drawbacks of Various VPN Implementations

Drawbacks of Traditional Peer-to-Peer VPNs

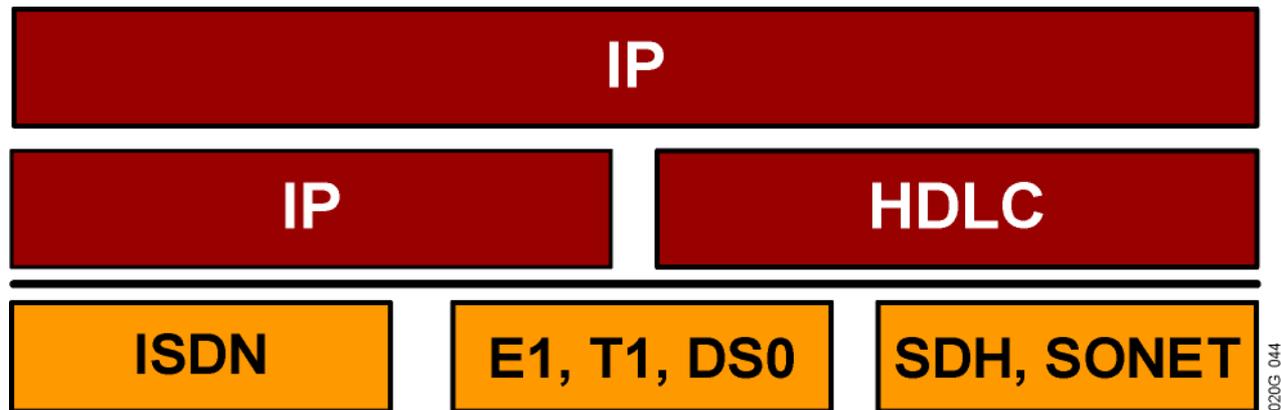
Lesson Summary

VPN Implementation Technologies

- 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

Layer 1 Implementation



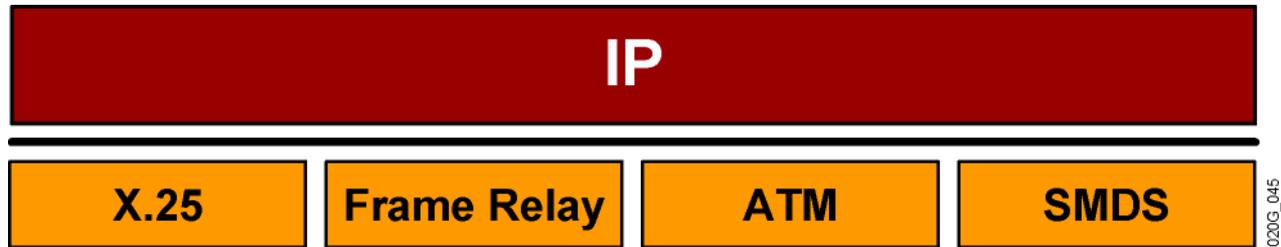
- This is the traditional TDM solution:

Service provider establishes physical-layer connectivity between customer sites.

Customer is responsible for all higher layers.

Overlay VPNs (Cont.)

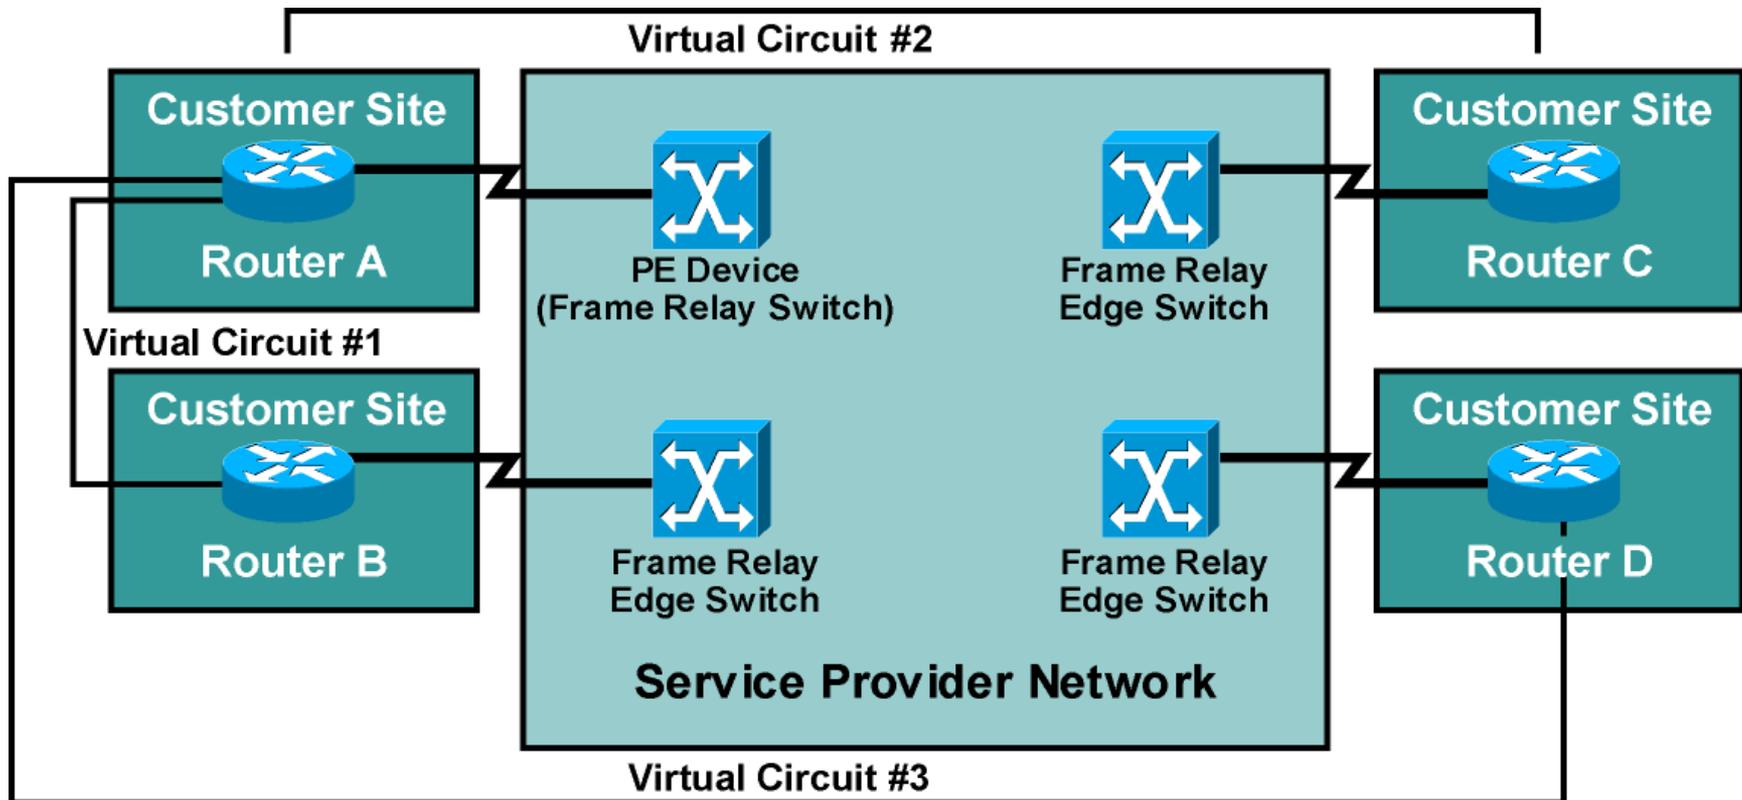
Layer 2 Implementation



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

Overlay VPNs (Cont.)

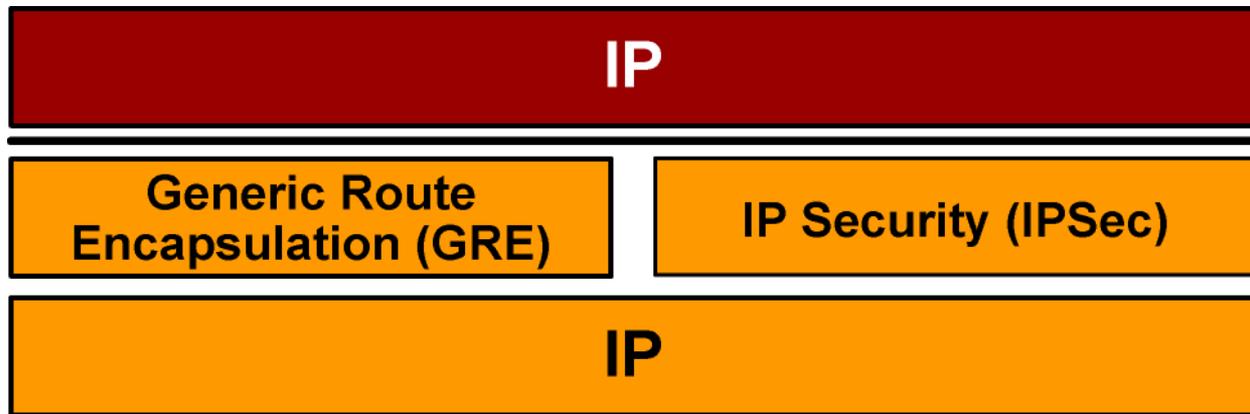
Frame Relay Example



020G_046

Overlay VPNs (Cont.)

IP Tunneling



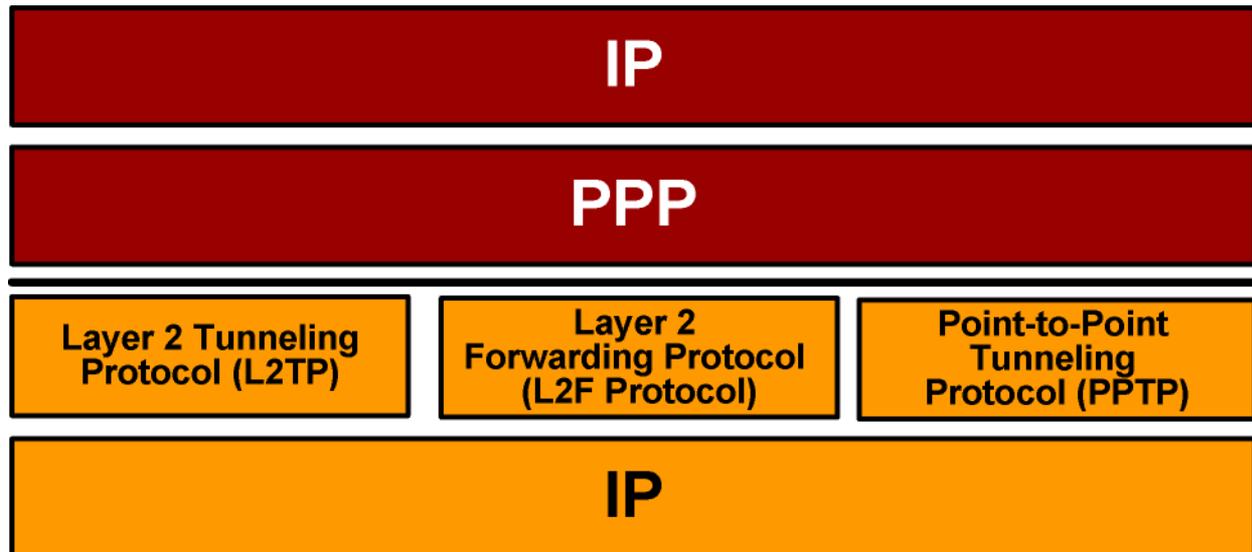
VPN is implemented with IP-over-IP tunnels:

Tunnels are established with GRE or IPsec.

GRE is simpler (and quicker); IPsec provides authentication and security.

Overlay VPNs (Cont.)

Layer 2 Forwarding

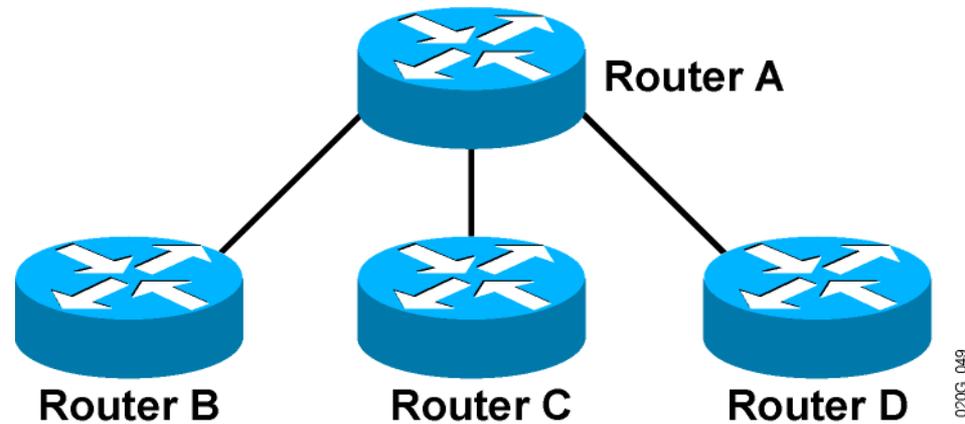


VPN is implemented with PPP-over-IP tunnels.

Usually used in access environments (dialup, digital subscriber line).

Overlay VPNs (Cont.)

Layer 3 Routing

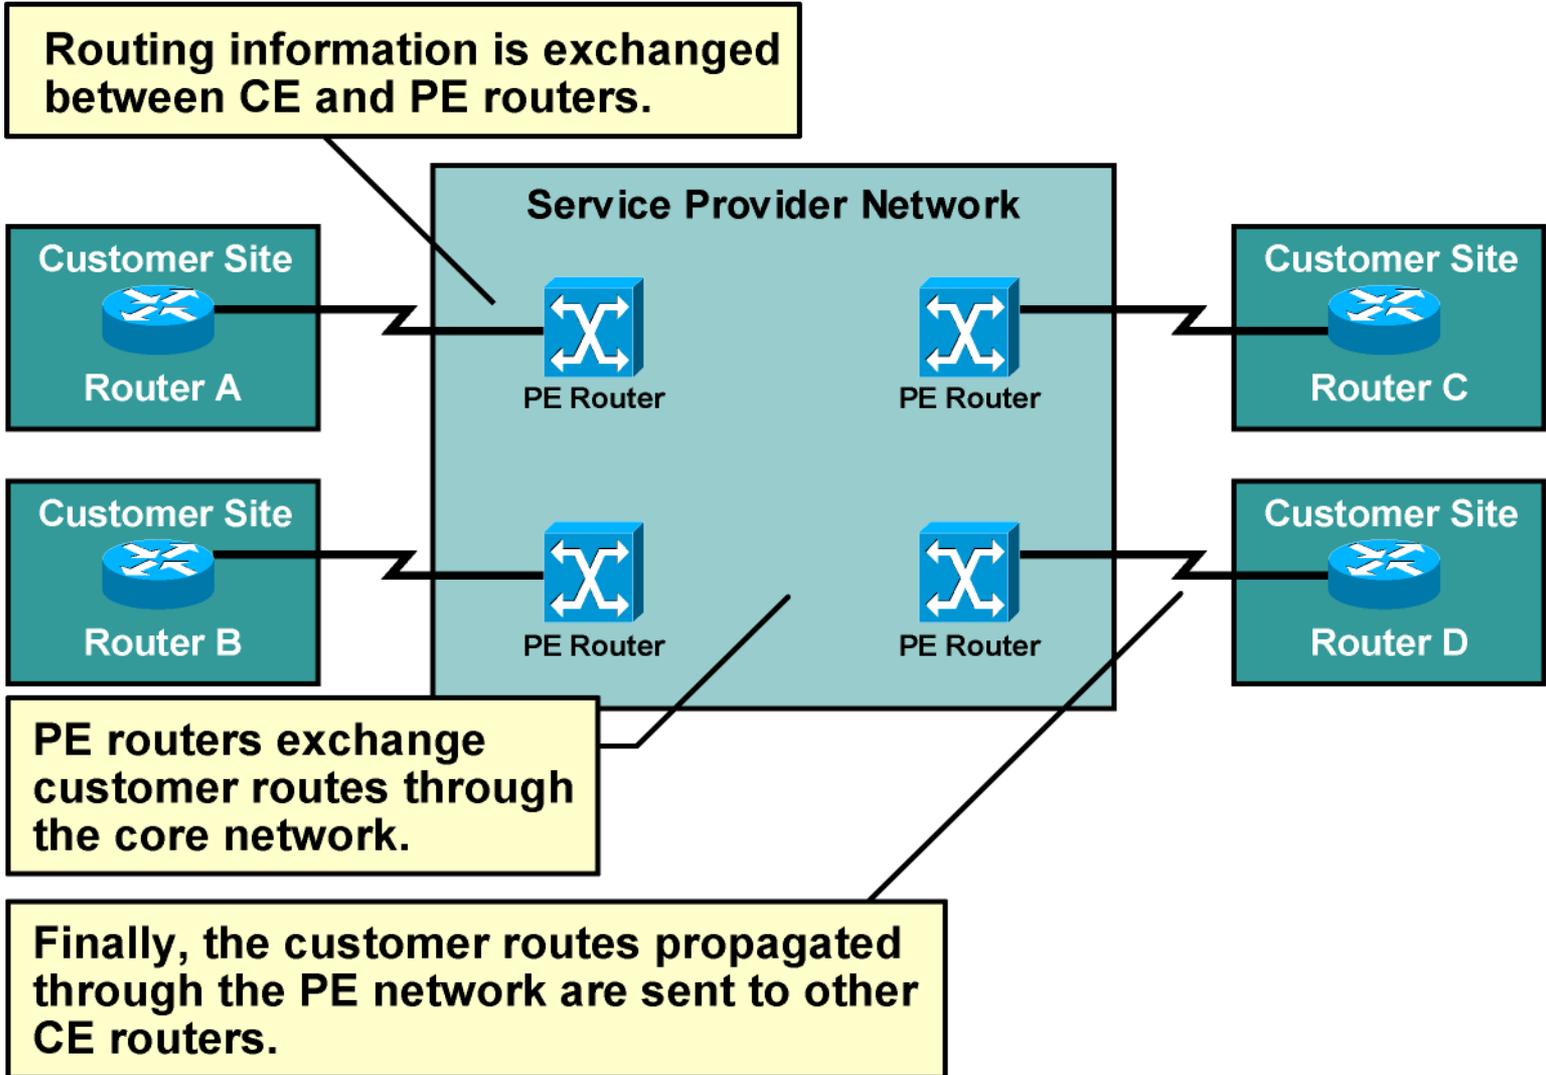


Service provider infrastructure appears as point-to-point links to customer routes.

Routing protocols run directly between customer routers.

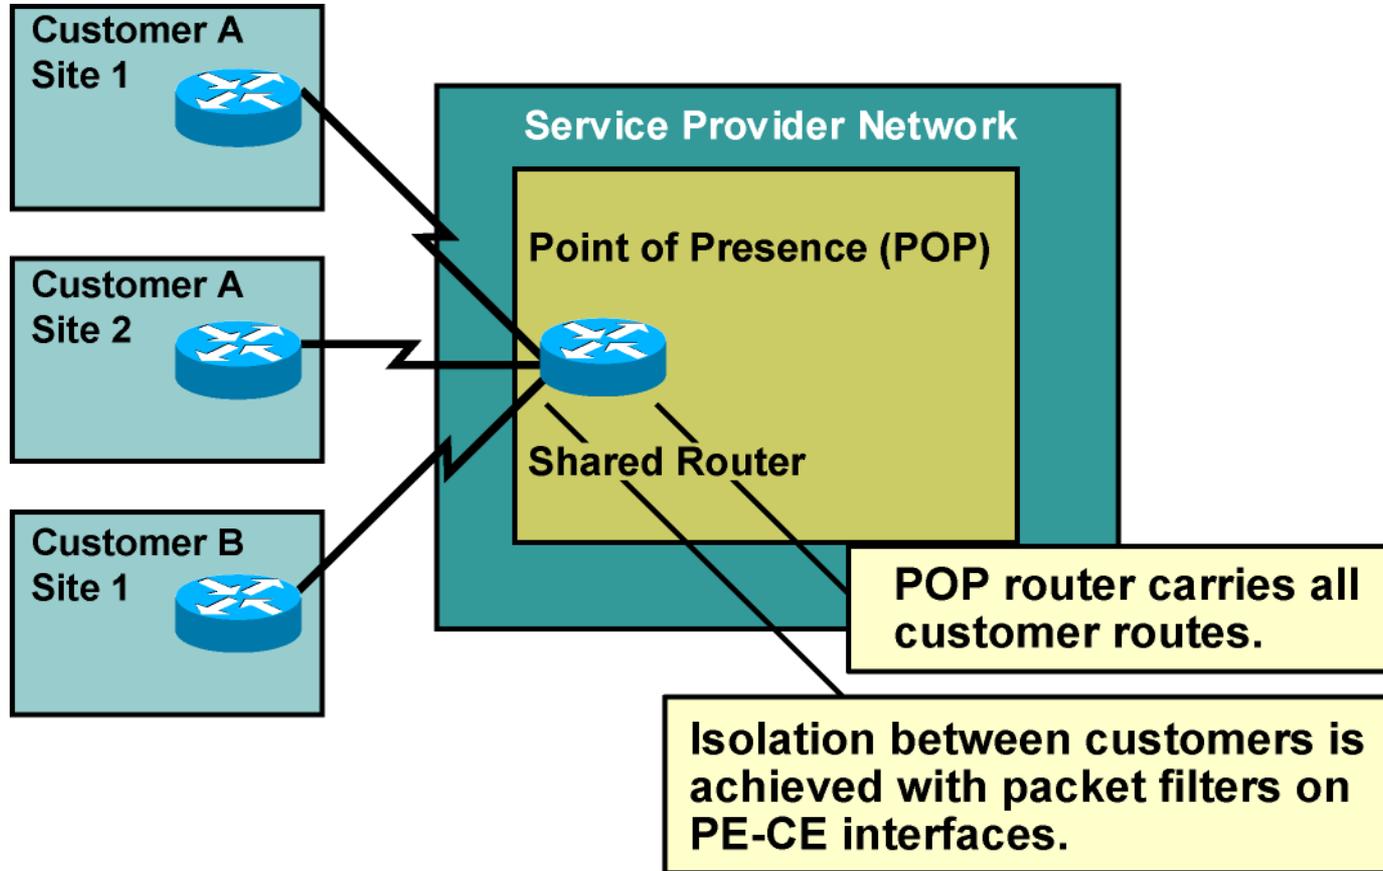
Service provider does not see customer routes and is responsible only for providing point-to-point transport of customer data.

Peer-to-Peer VPNs



Peer-to-Peer VPNs (Cont.)

Packet Filters



0206_051

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.

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.

Summary

The two major VPN models are overlay and peer-to-peer.

Overlay VPNs can be implemented using Layer 1, Layer 2, and Layer 3 technologies.

Traditional peer-to-peer VPNs are implemented using IP routing technology.

Overlay VPNs use well-known technologies and are easy to implement, but require a full mesh of virtual circuits to provide optimum routing.

Summary

Peer-to-peer VPNs guarantee optimum routing between customer sites but require that the service provider participates in customer routing.

Both shared PE router and dedicated PE router implementations of peer-to-peer VPNs require the customers to share a common address space.



MPLS workshop

VPN Types

Outline

Overview

VPN Categorization

Hub-and-Spoke Topology

Partial Mesh Overlay VPN

VPN Business Categorization

Extranet VPN

VPN Connectivity Categorization

Central Services Extranet

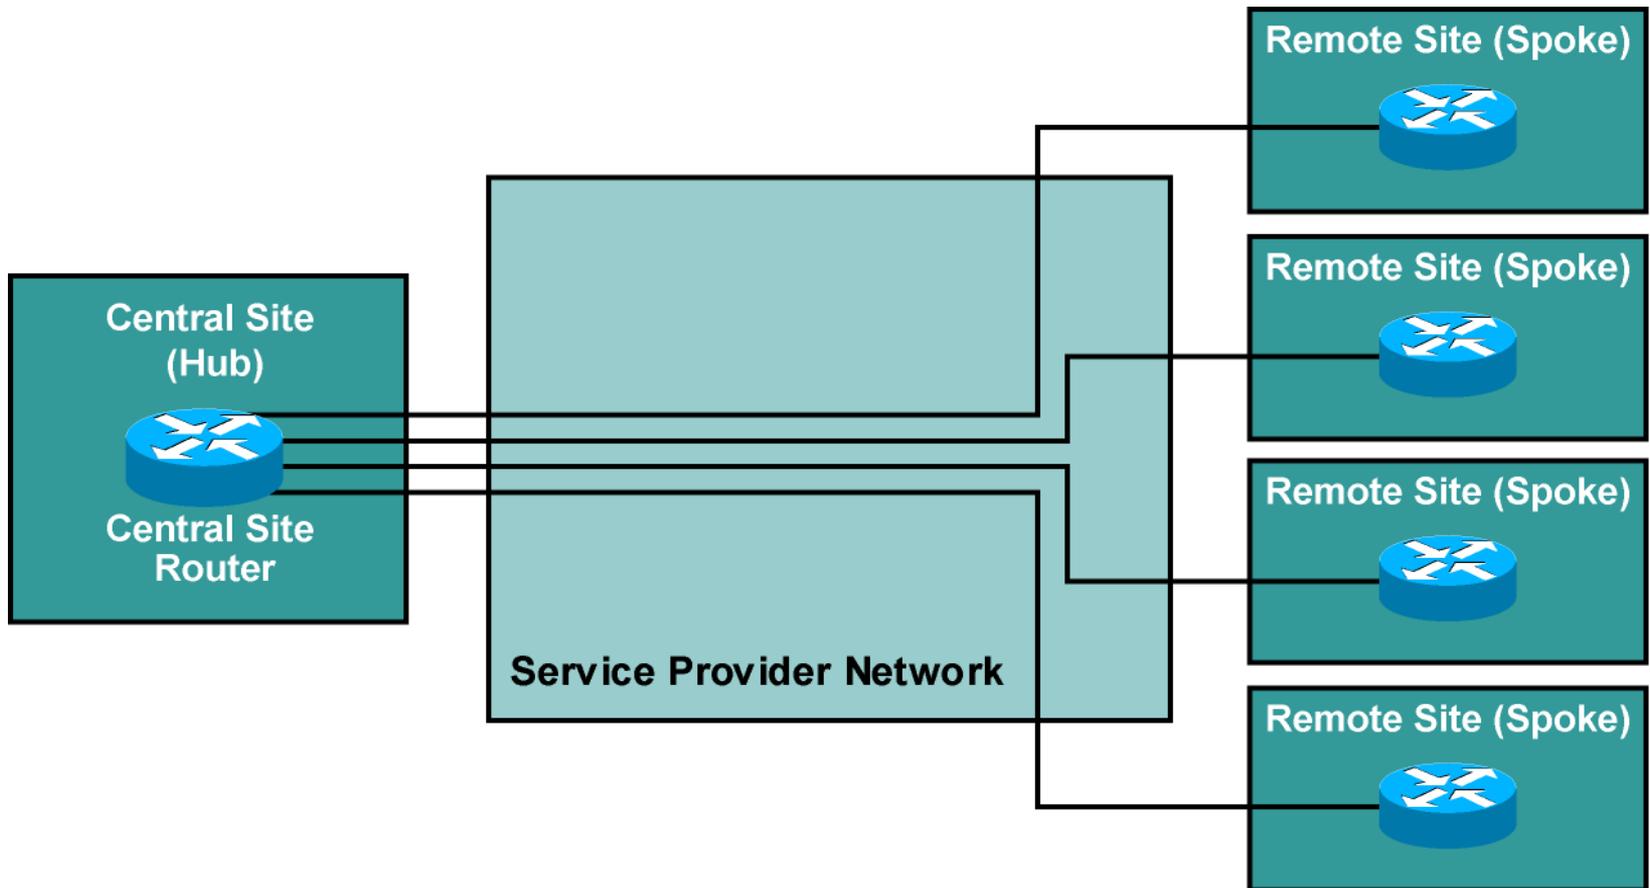
Managed Network Overlay VPN Implementation

Lesson Summary

Overlay VPN Topology Category

- Overlay VPNs are categorized based on the topology of the virtual circuits:
 - (Redundant) hub-and-spoke
 - Partial mesh
 - Full mesh
 - Multilevel—combines several levels of overlay VPN topologies

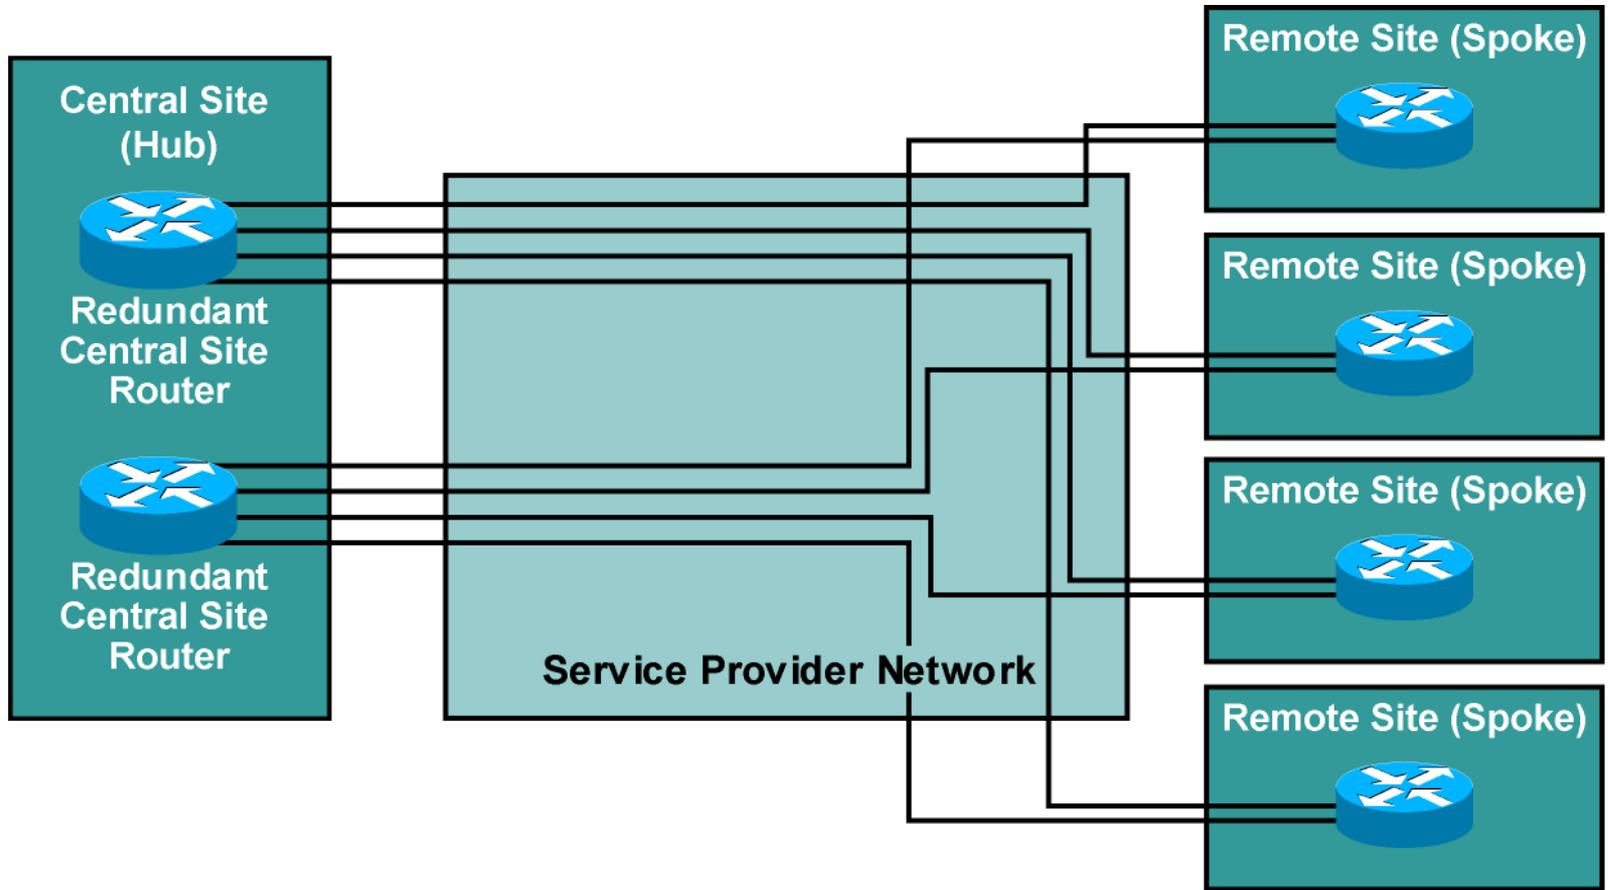
Hub-and-Spoke Overlay VPN Topology



020G_053

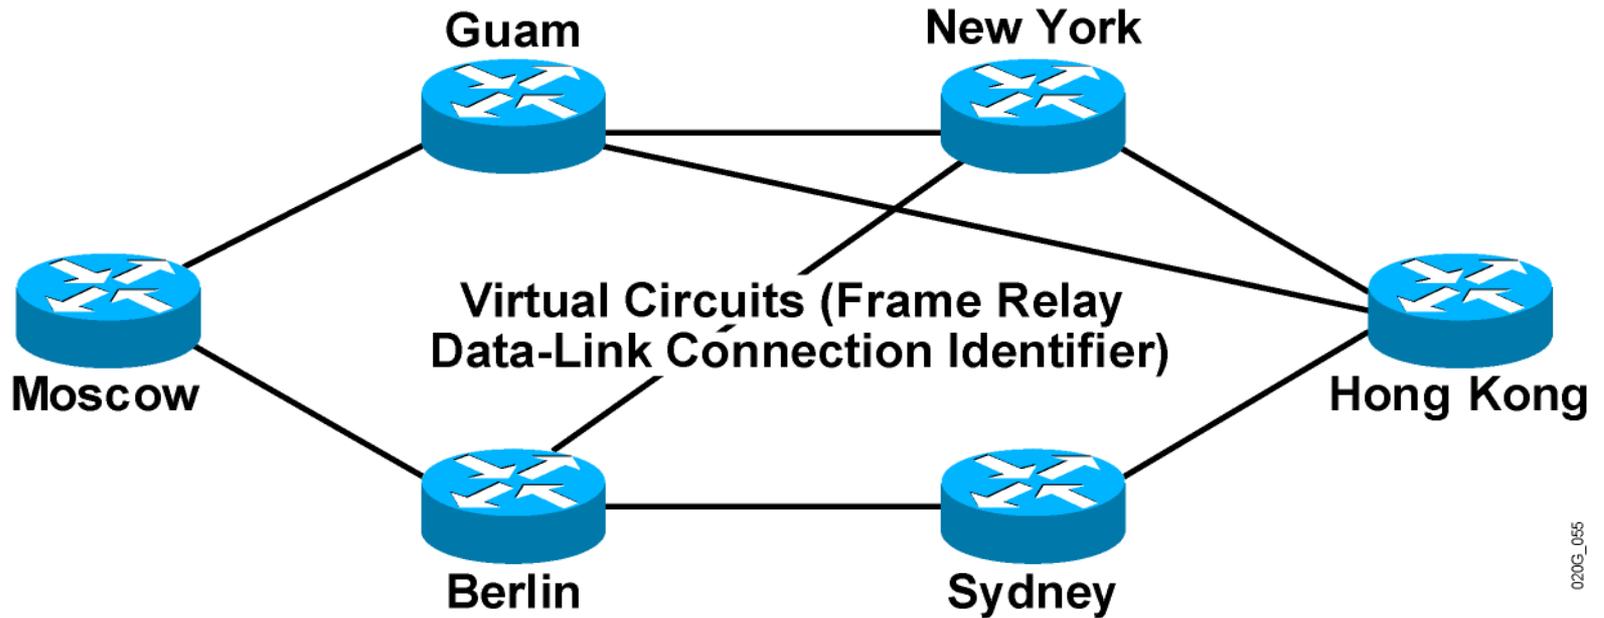
Hub-and-Spoke Overlay VPN Topology (Cont.)

Redundant Hub-and-Spoke Topology



020G_054

Partial Mesh Overlay VPN Topology



020G_055

VPN Business Category

- VPNs can be categorized on the business needs that they fulfill:

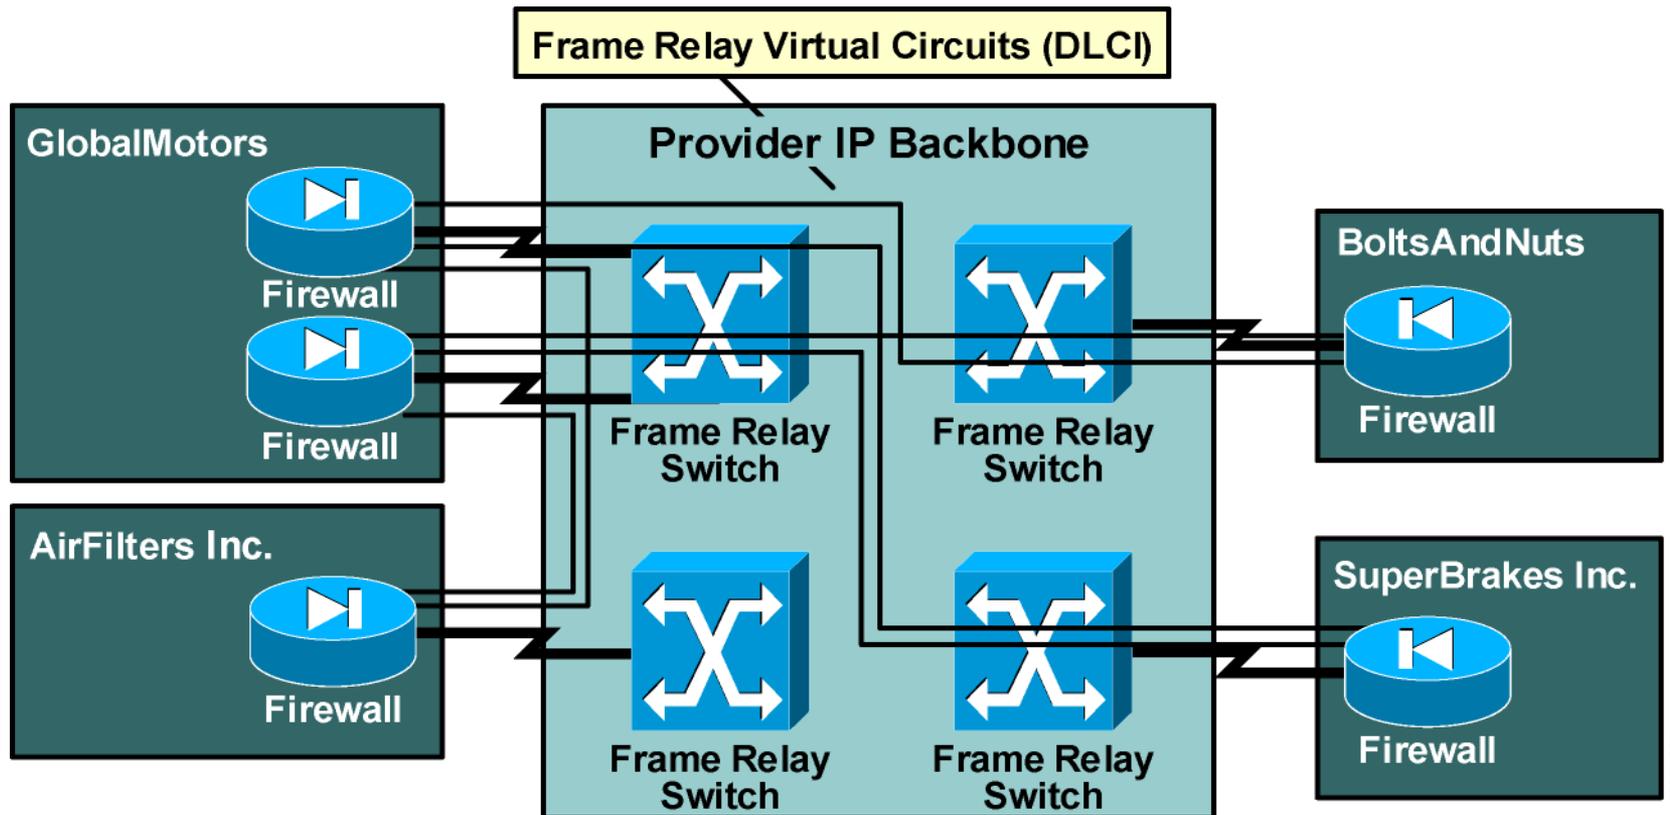
Intranet VPN: Connects sites within an organization.

Extranet VPN: Connects different organizations in a secure way.

Access VPN: VPDN provides dialup access into a customer network.

Extranet VPNs

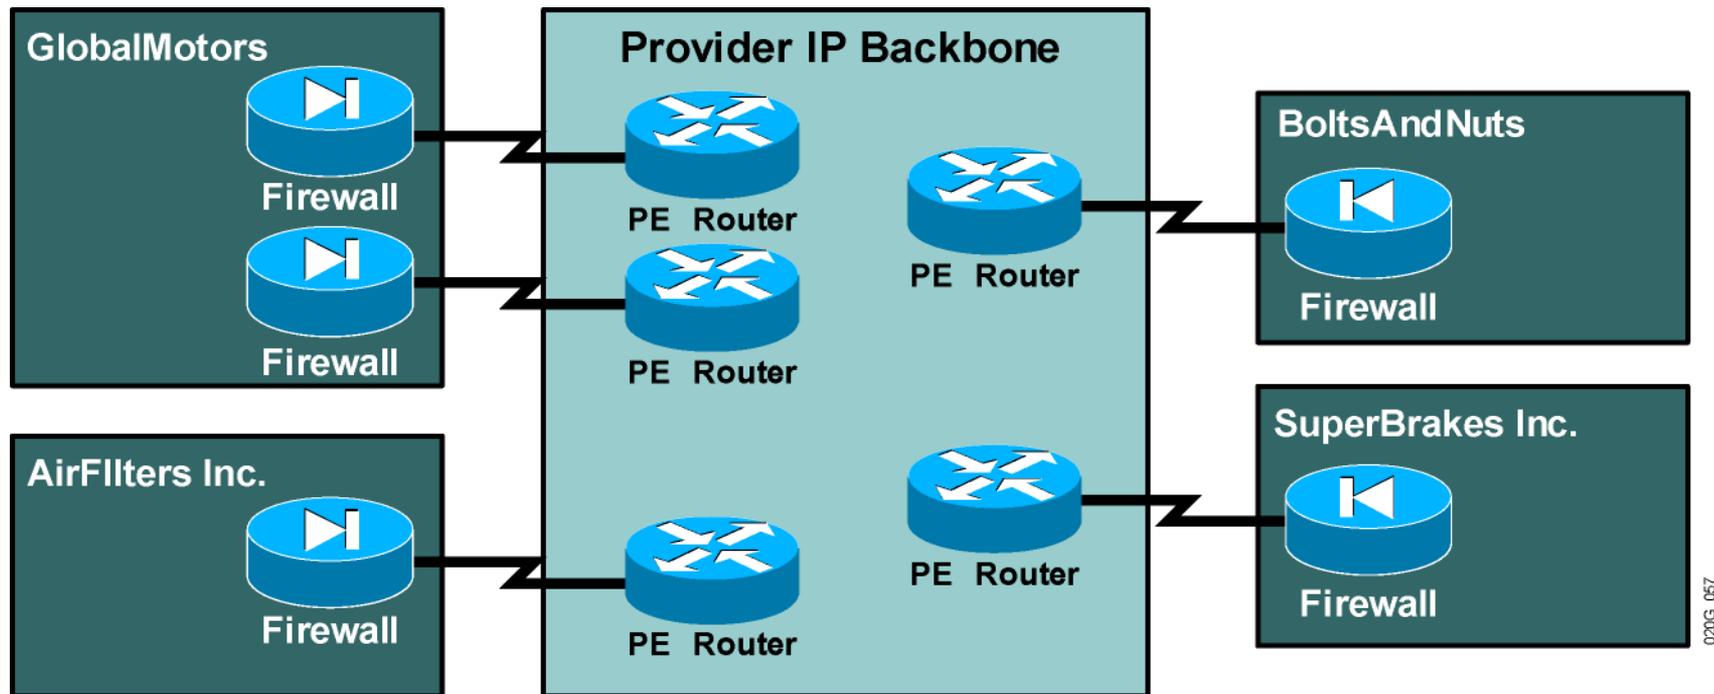
Overlay VPN Implementation



0200_G_056

Extranet VPNs (Cont.)

Peer-to-Peer VPN Implementation

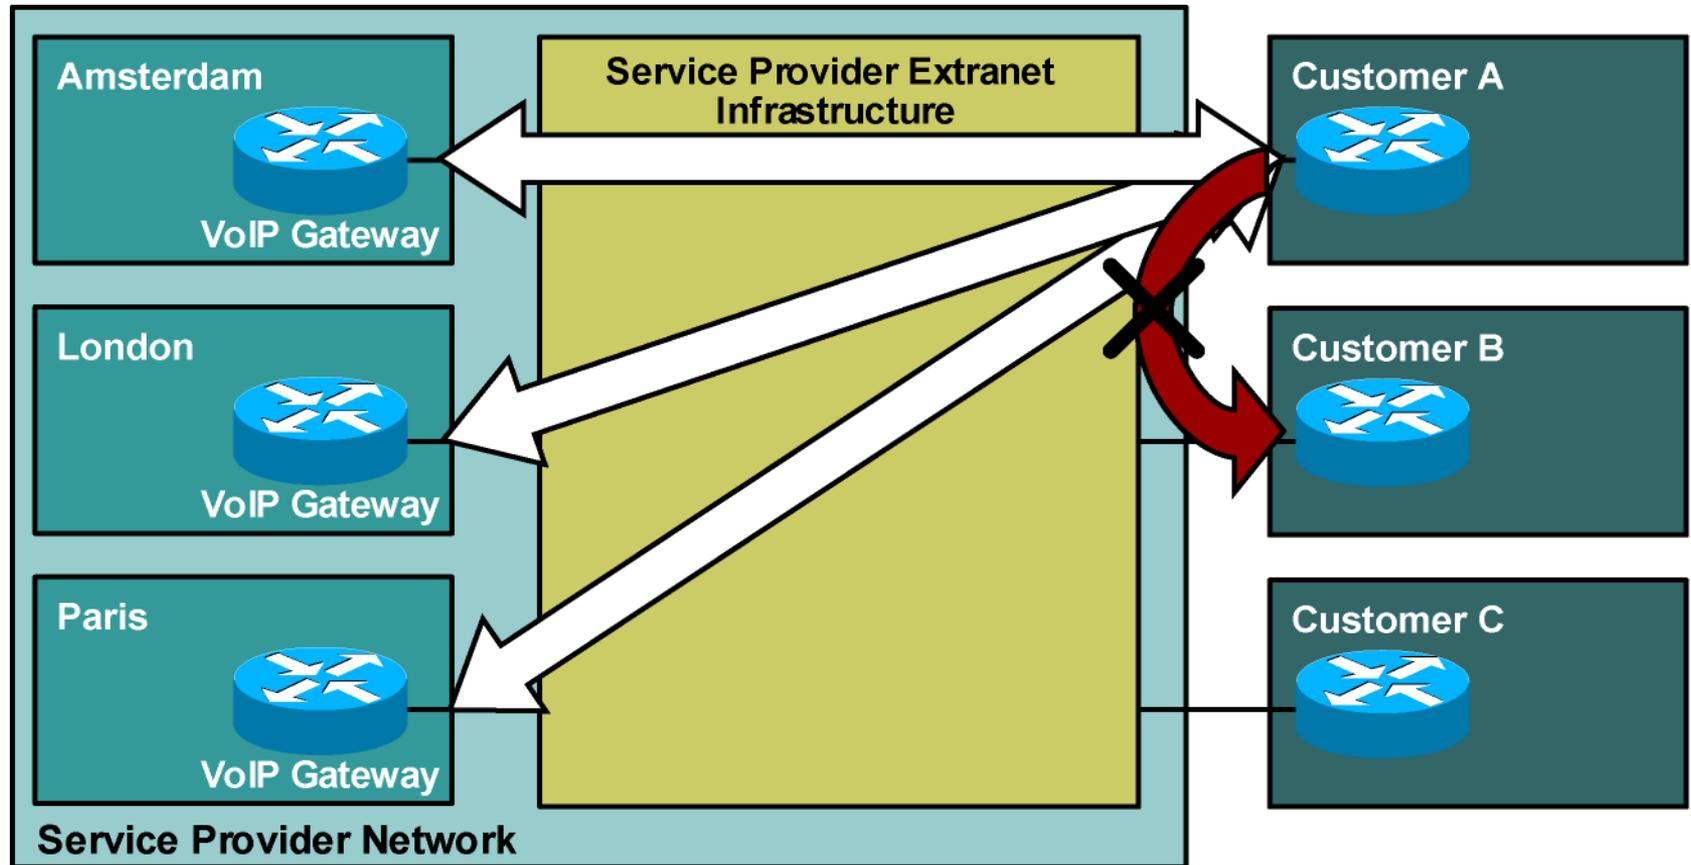


02003_057

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 VPN: 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.

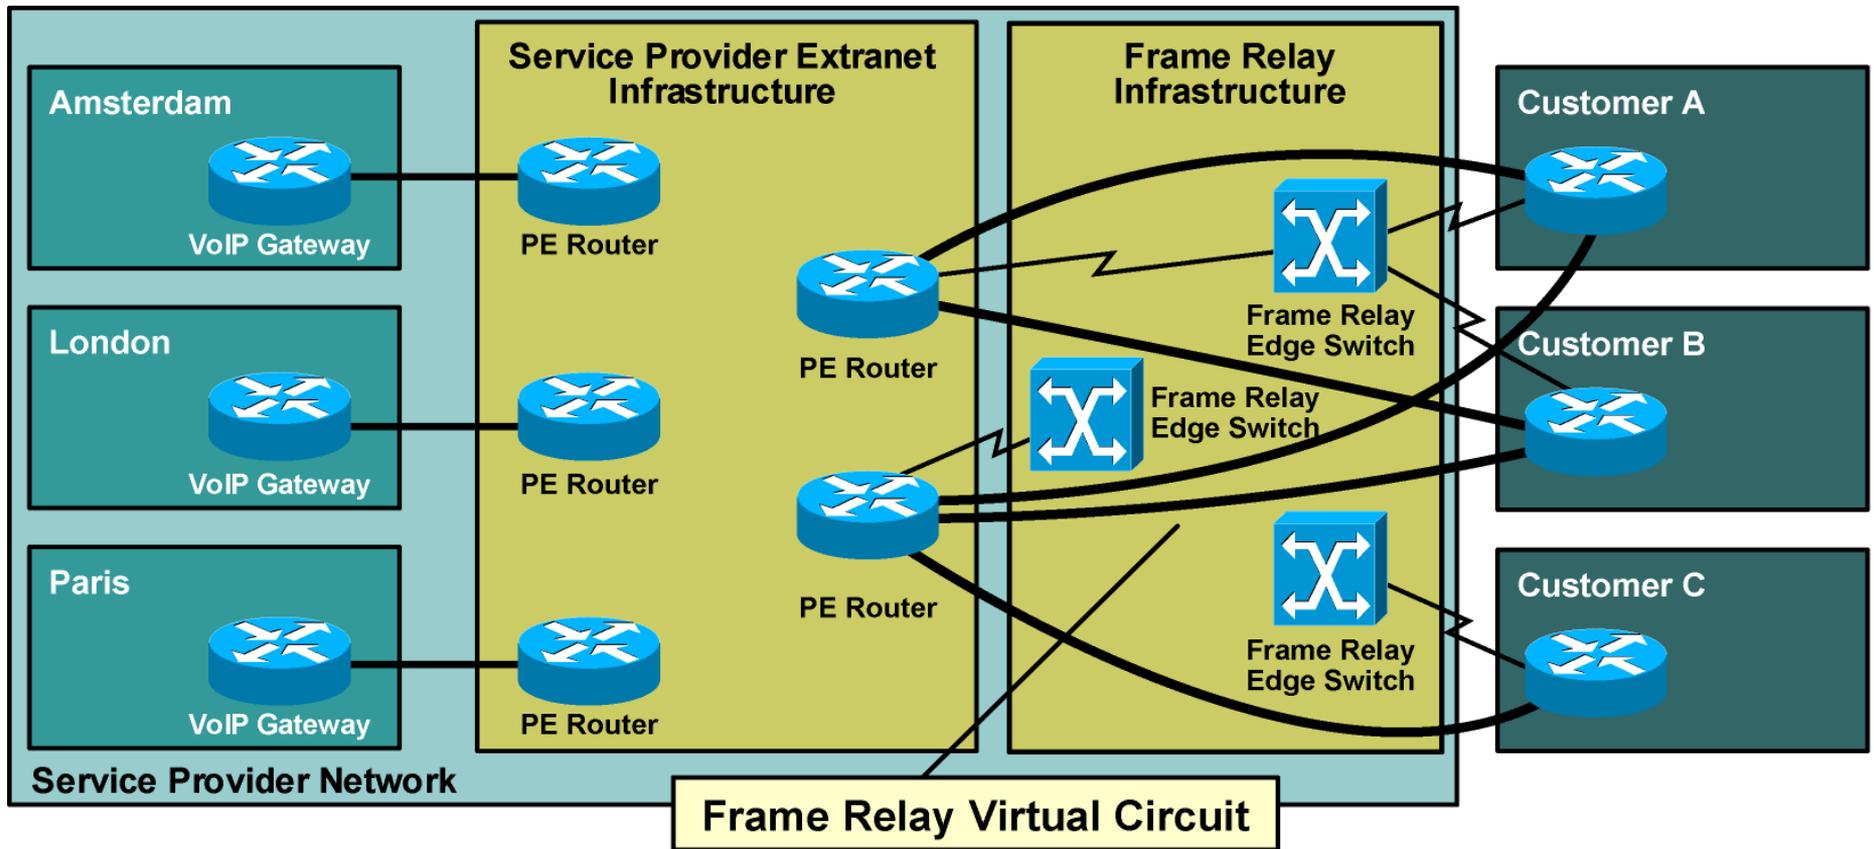
Central Services Extranet



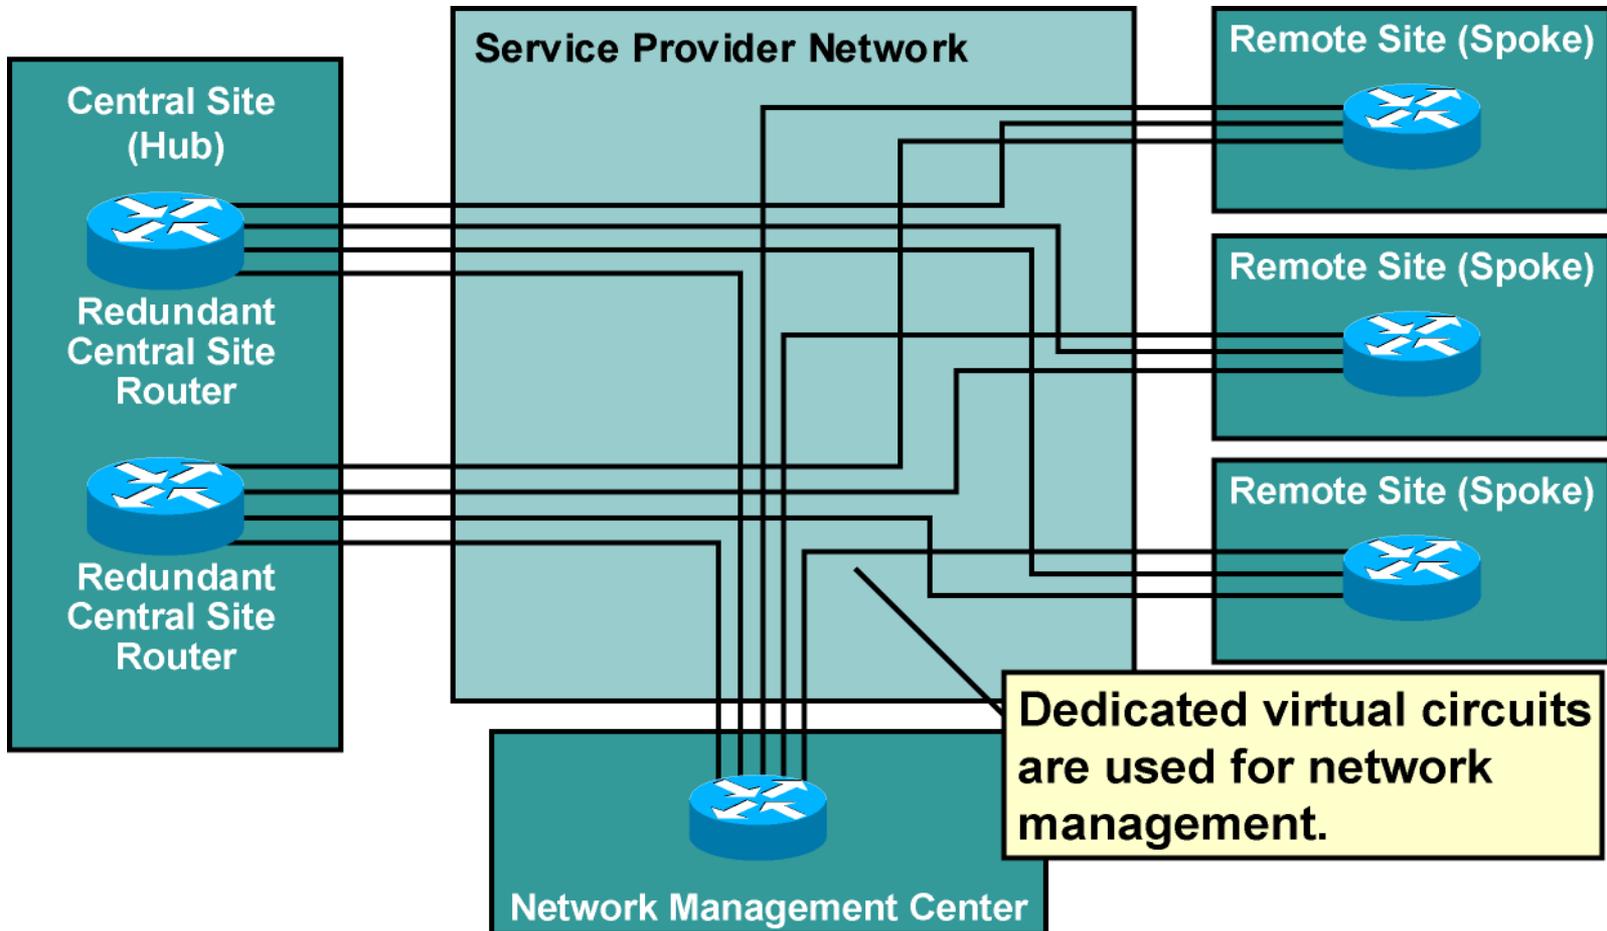
0200_058

Central Services Extranet (Cont.)

Hybrid (Overlay + Peer-to-Peer) Implementation



Managed Network Overlay VPN Implementation



0203_090

Summary

Major VPN topologies consist of the following:

- Hub-and-spoke – simplest topology

- Partial mesh – cost/complexity factors dictate

- Full mesh – connections between all sites

- Multilevel – can be used for large-scale networks

VPNs can be based on business needs:

- Intranet

- Extranet

- Access



MPLS workshop

MPLS VPN Architecture

Outline

Overview

MPLS VPN Architecture

PE Router Architecture

Propagation Routing Information across the P-network

Route Distinguishers

Route Targets

Virtual Private Networks Redefined

Impact of Complex VPN Topologies on Virtual Routing Tables

Lesson Summary

MPLS VPN Architecture

- An MPLS VPN combines the best features of an overlay VPN and a peer-to-peer VPN:

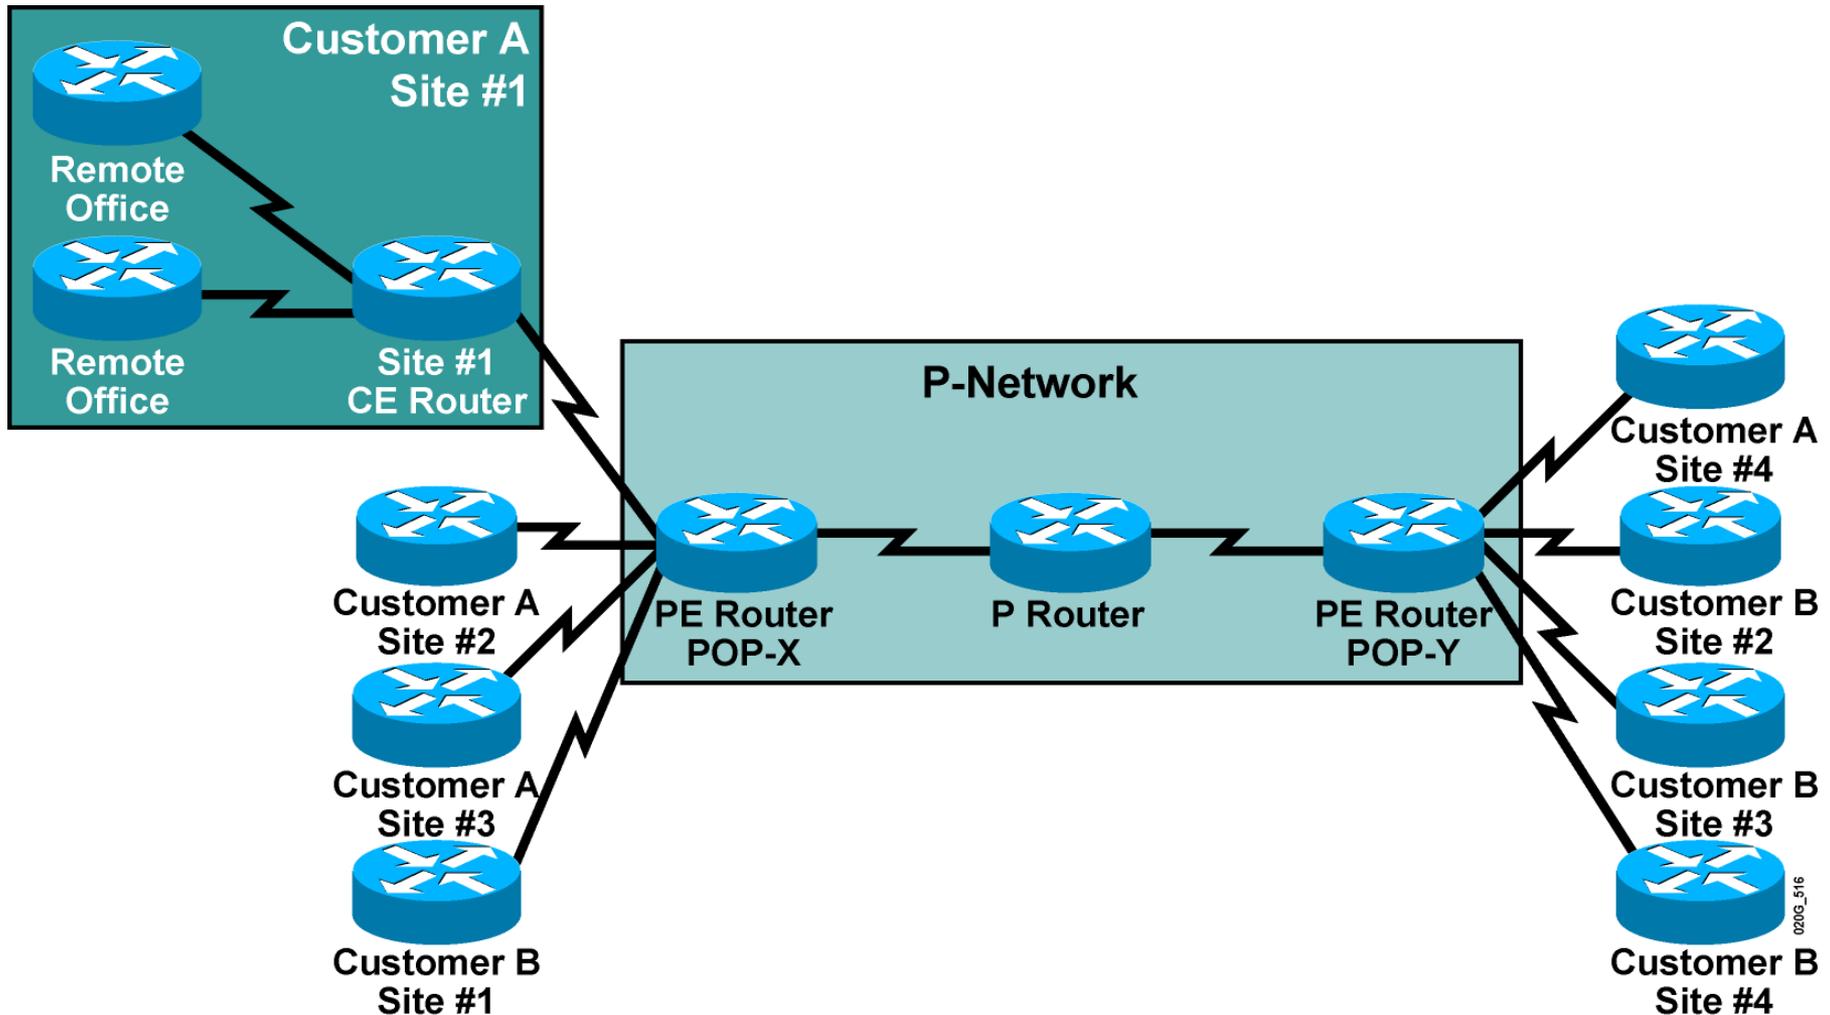
PE routers participate in customer routing, guaranteeing optimum routing between sites and easy provisioning.

PE routers carry a separate set of routes for each customer (similar to the dedicated PE router approach).

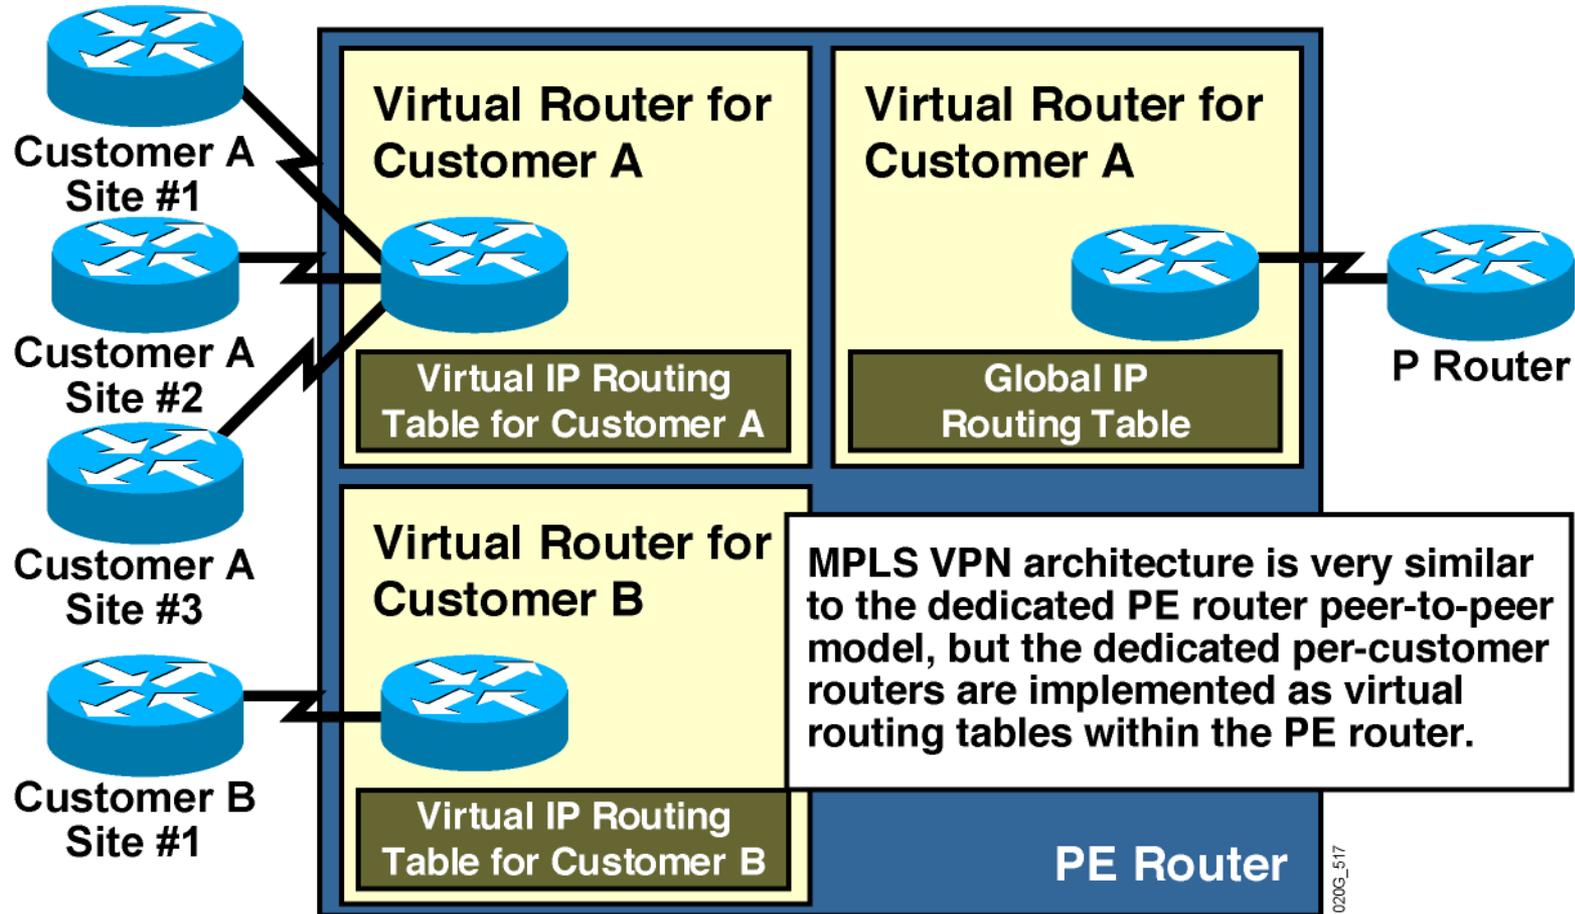
Customers can use overlapping addresses.

MPLS VPN Architecture (Cont.)

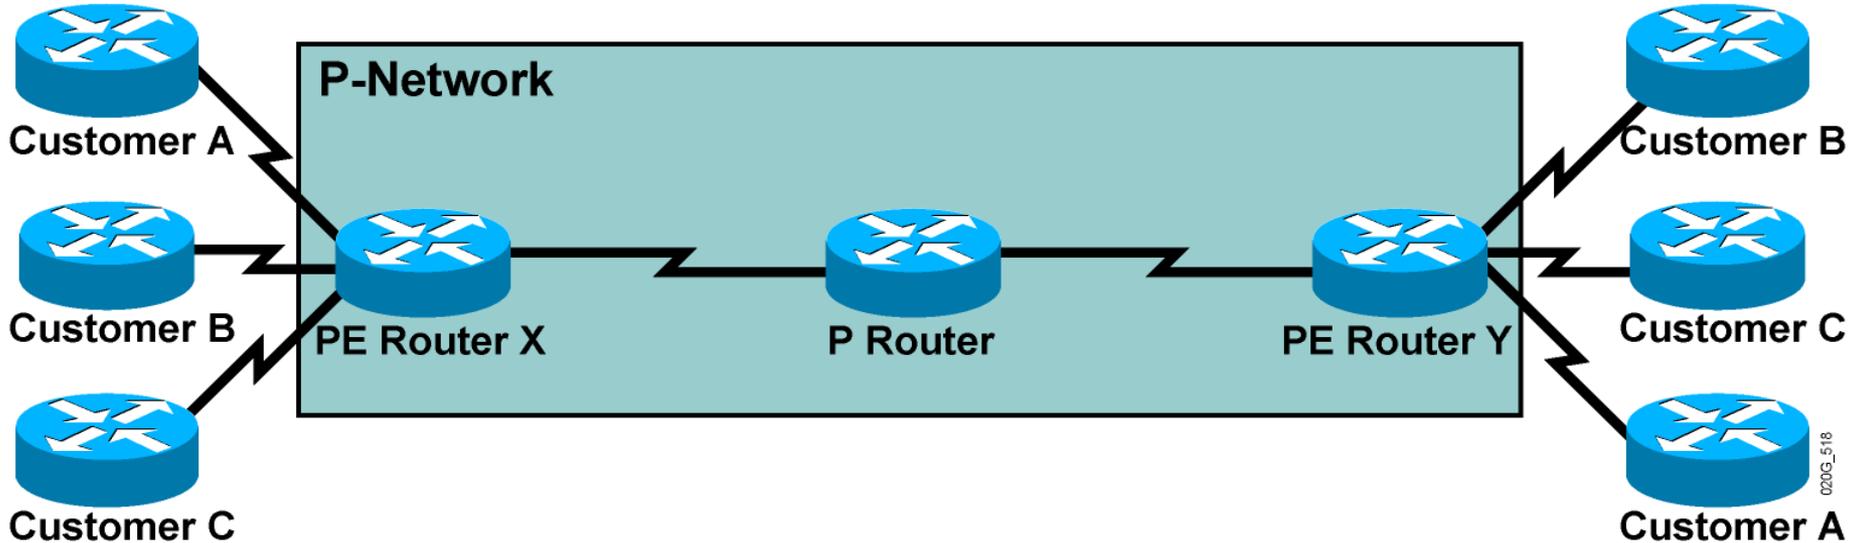
Terminology



PE Router Architecture

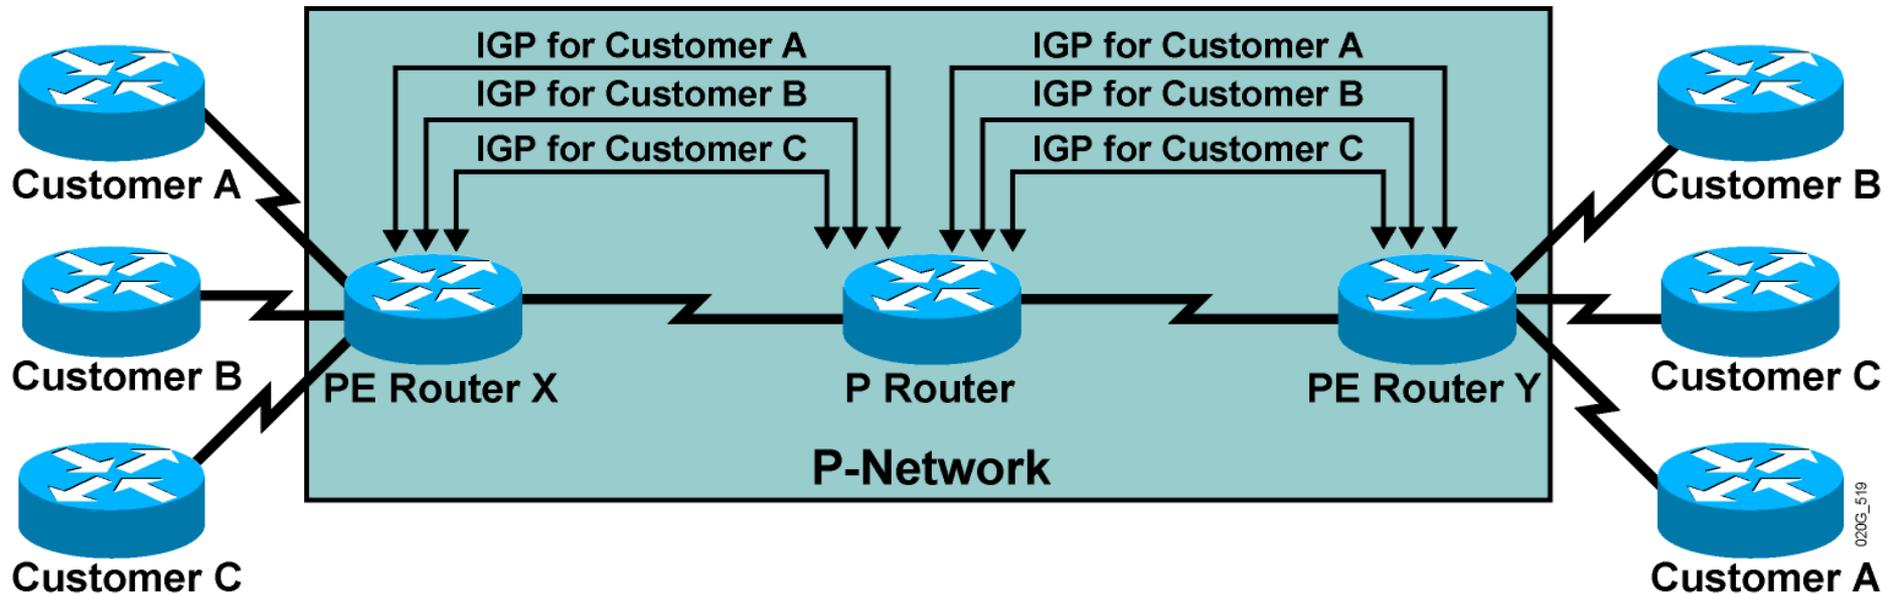


Propagation of Routing Information Across the P-Network



Question: How will PE routers exchange customer routing information?

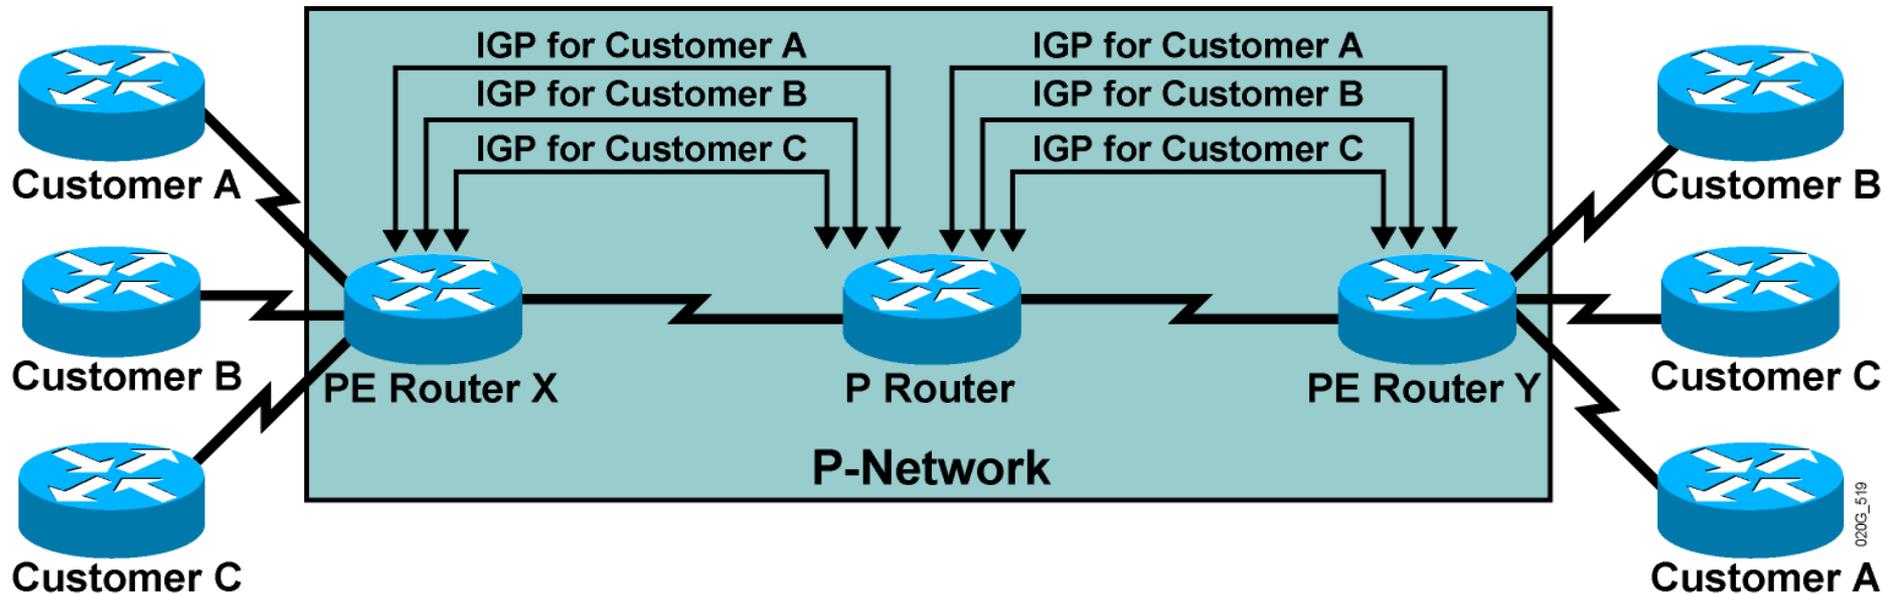
Propagation of Routing Information Across the P-Network



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.

Propagation of Routing Information Across the P-Network



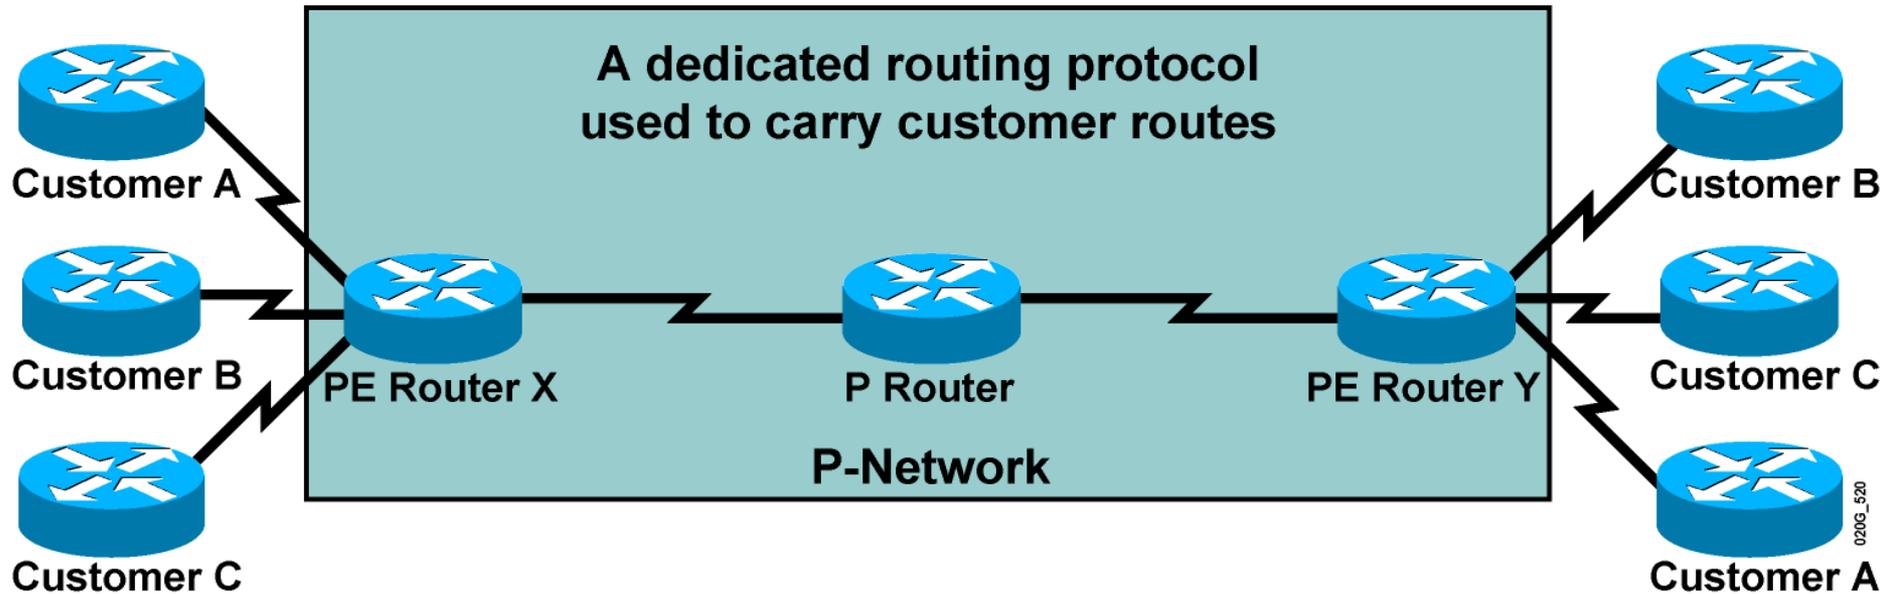
Question: How will PE routers exchange customer routing information?

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:

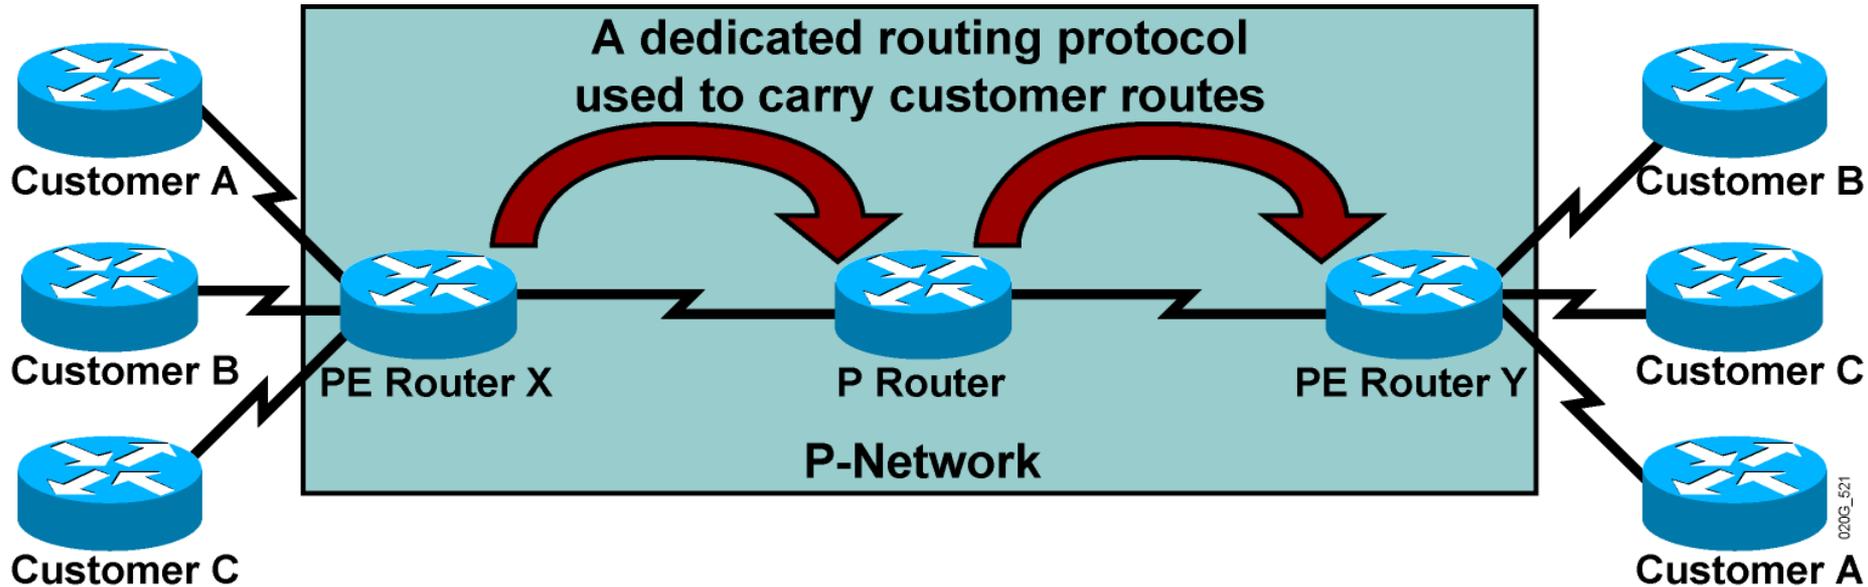
- The solution does not scale.
- P routers carry all customer routes.

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



Question: How will PE routers exchange customer routing information?

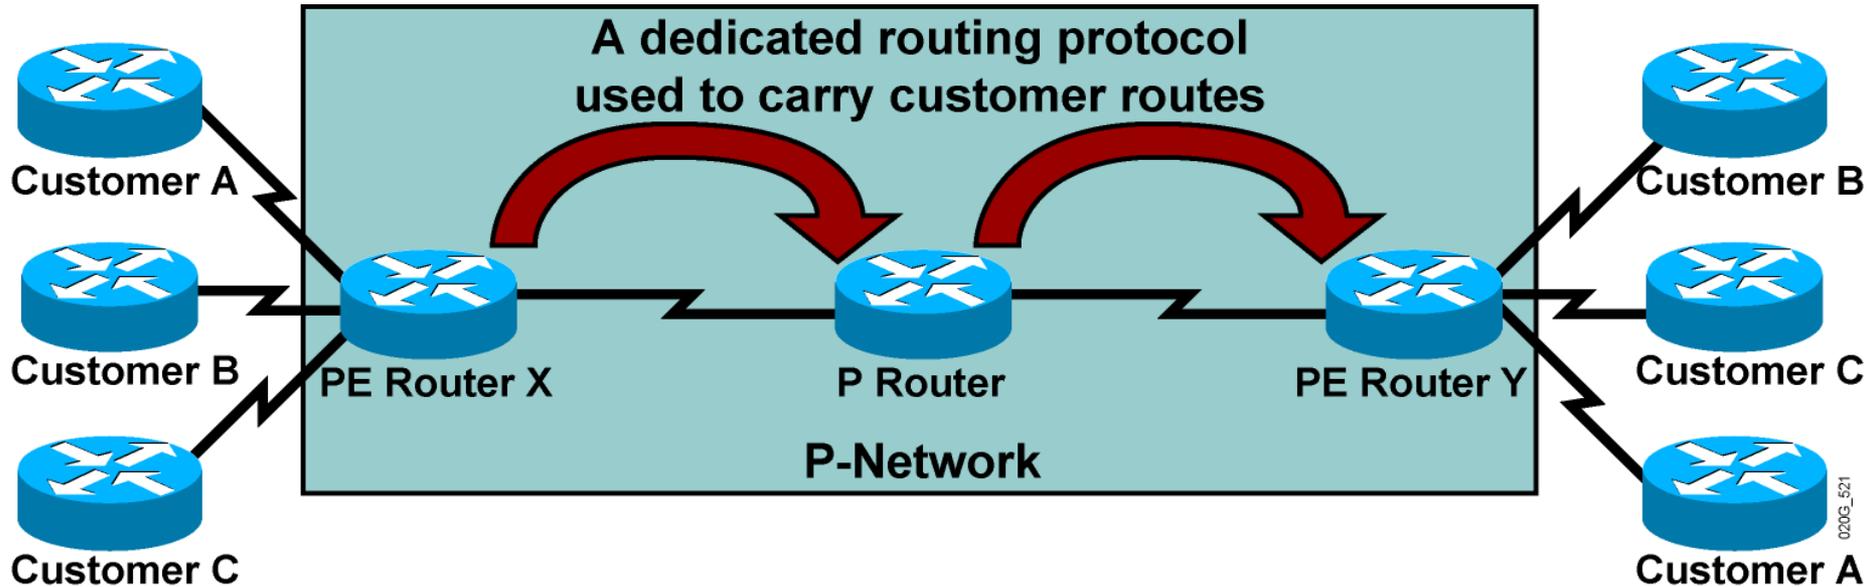
Propagation of Routing Information Across the P-Network (Cont.)



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.

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



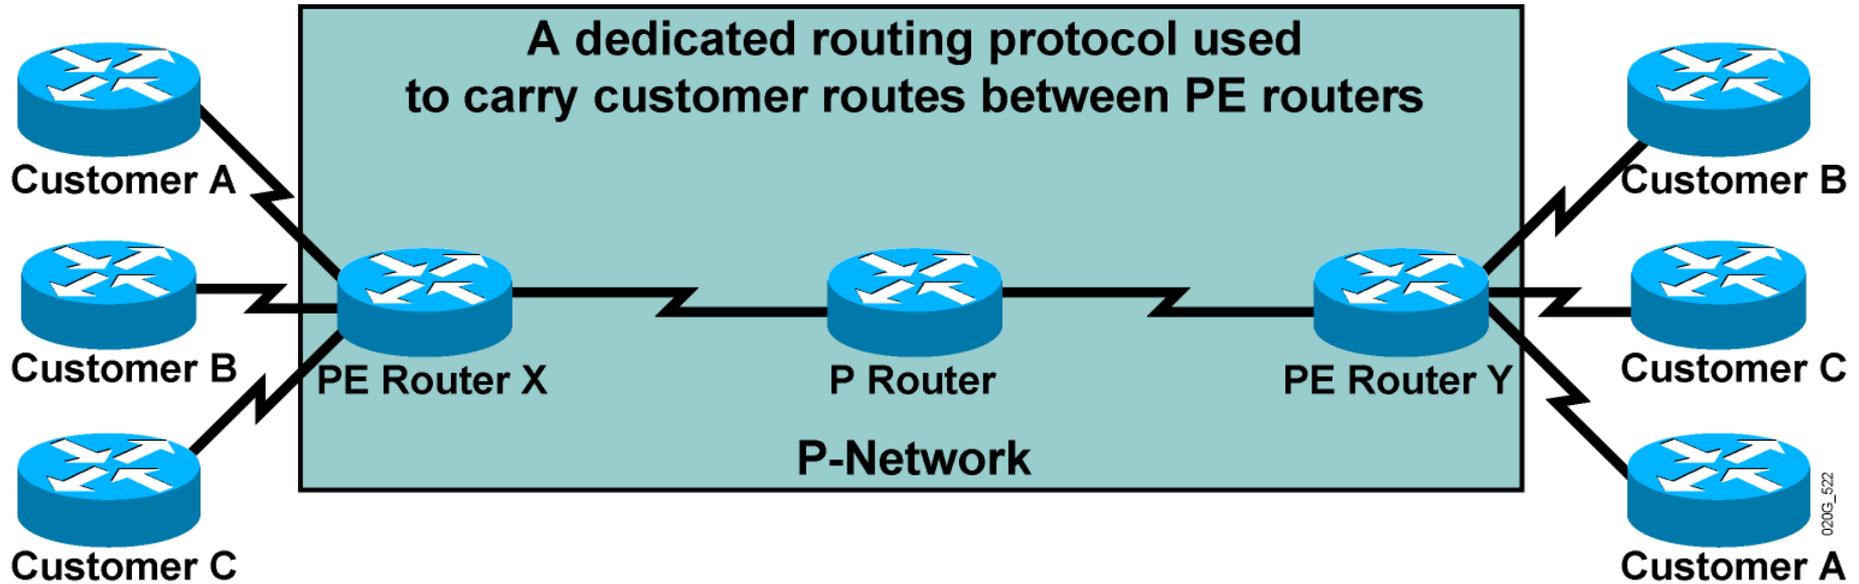
Question: How will PE routers exchange customer routing information?

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

Better answer, but still not good enough:

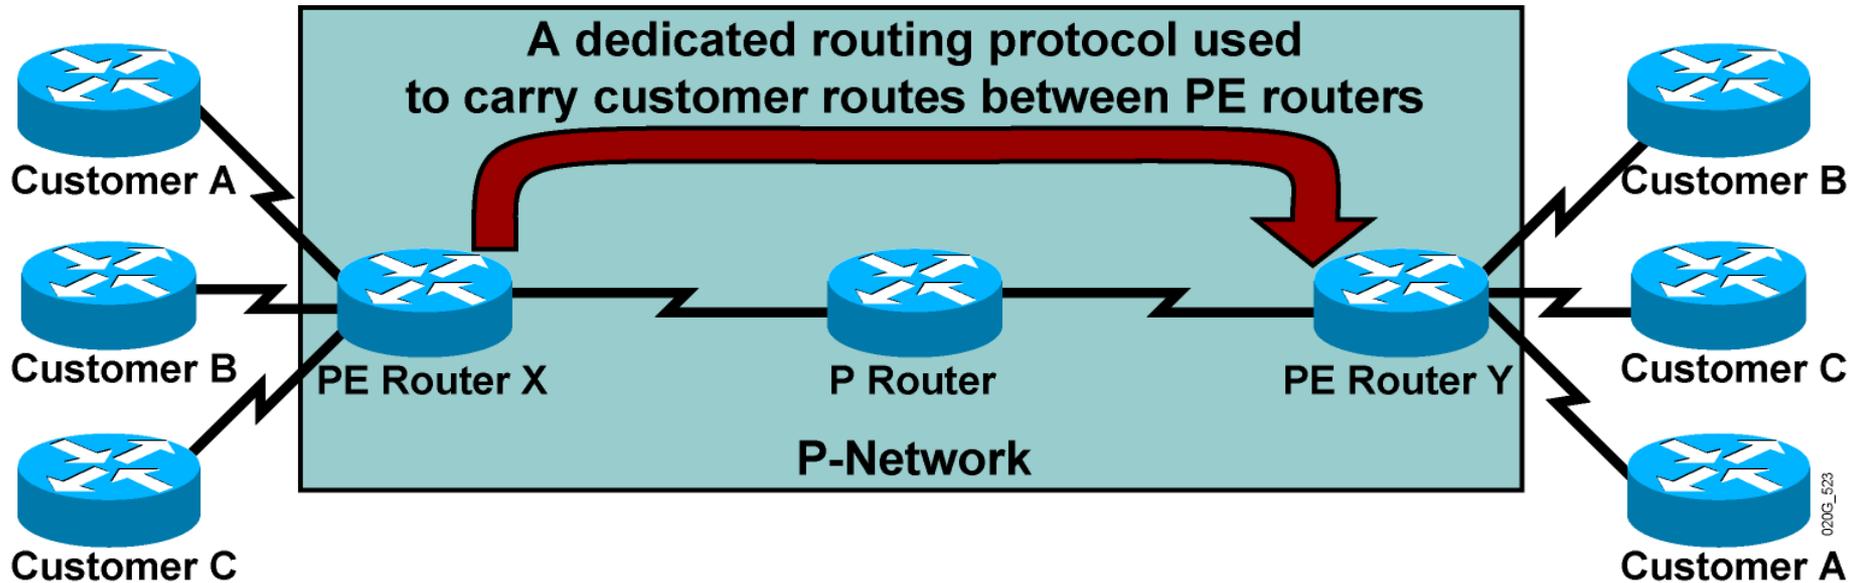
- P routers carry all customer routes.

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



Question: How will PE routers exchange customer routing information?

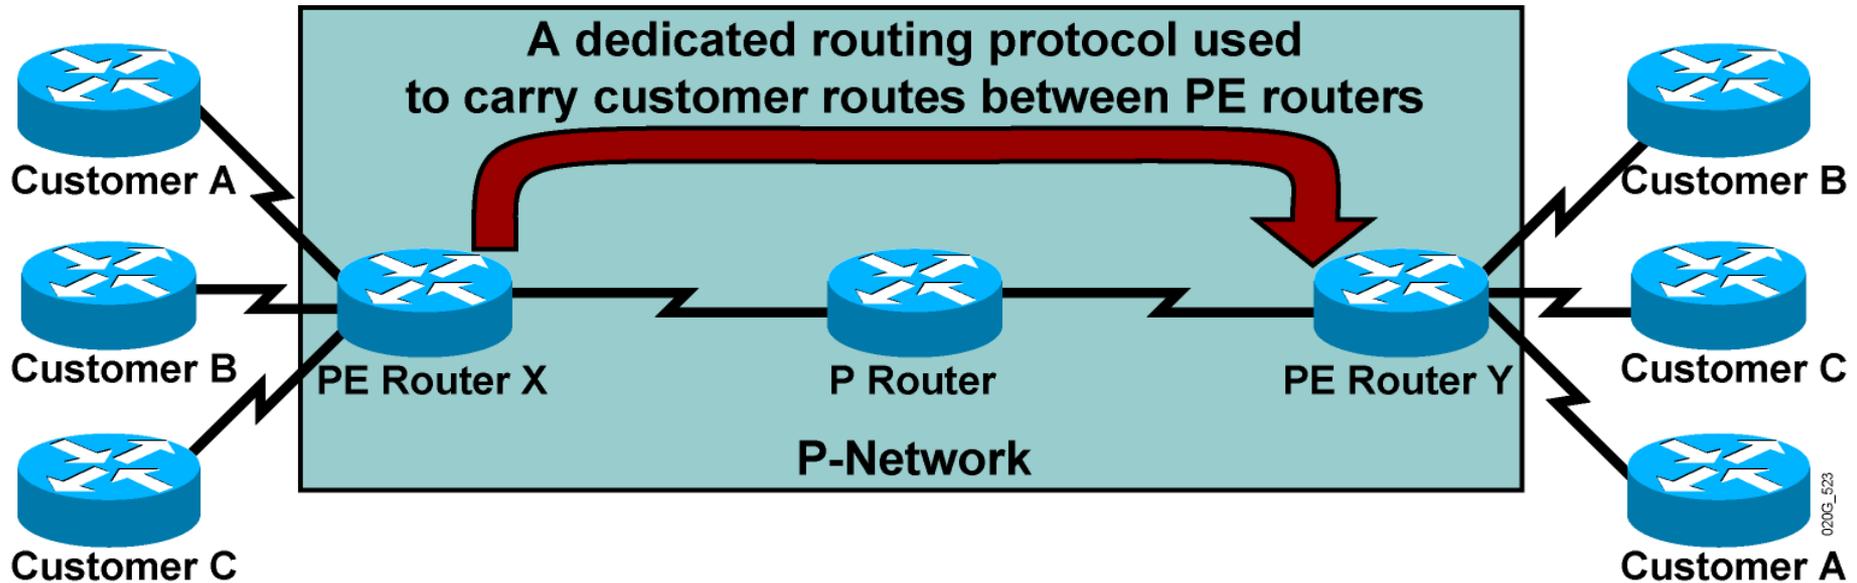
Propagation of Routing Information Across the P-Network (Cont.)



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.

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



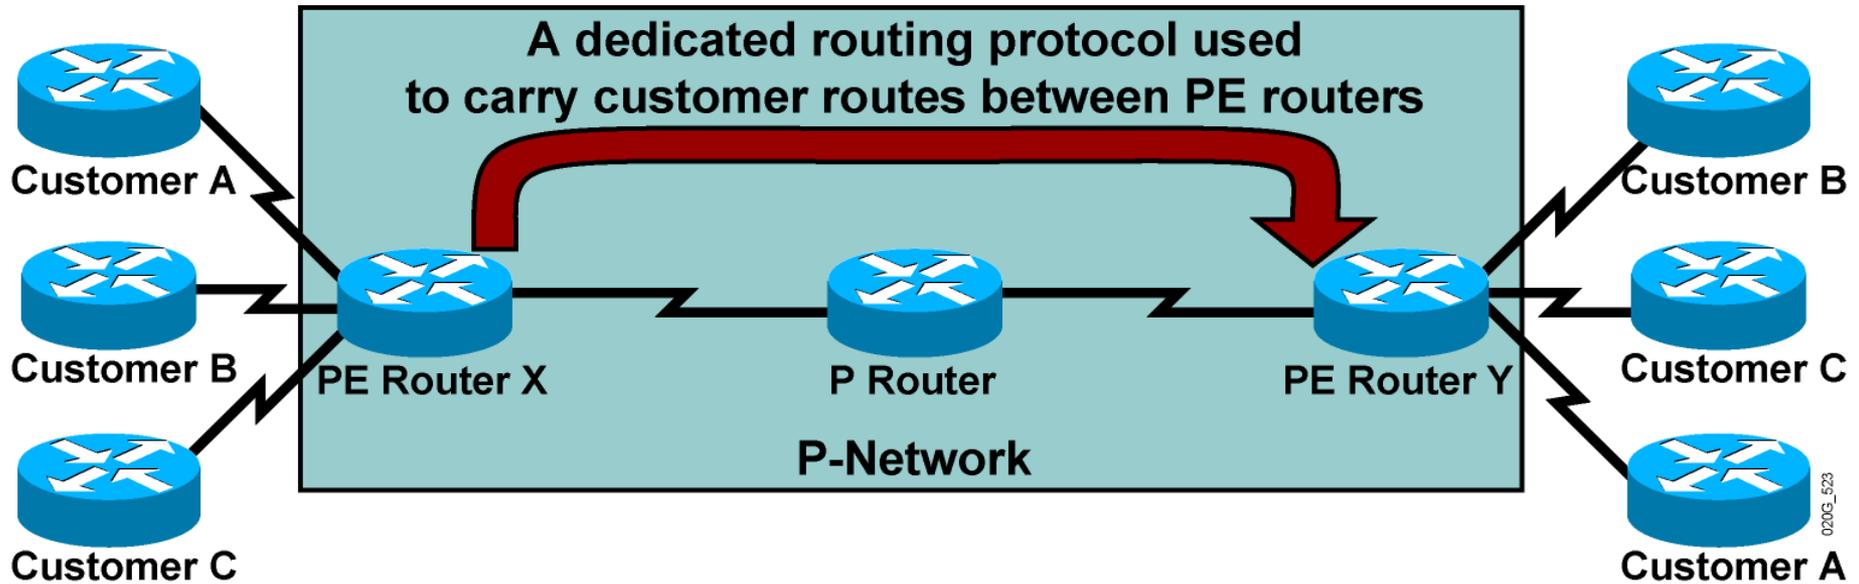
Question: How will PE routers exchange customer routing information?

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:

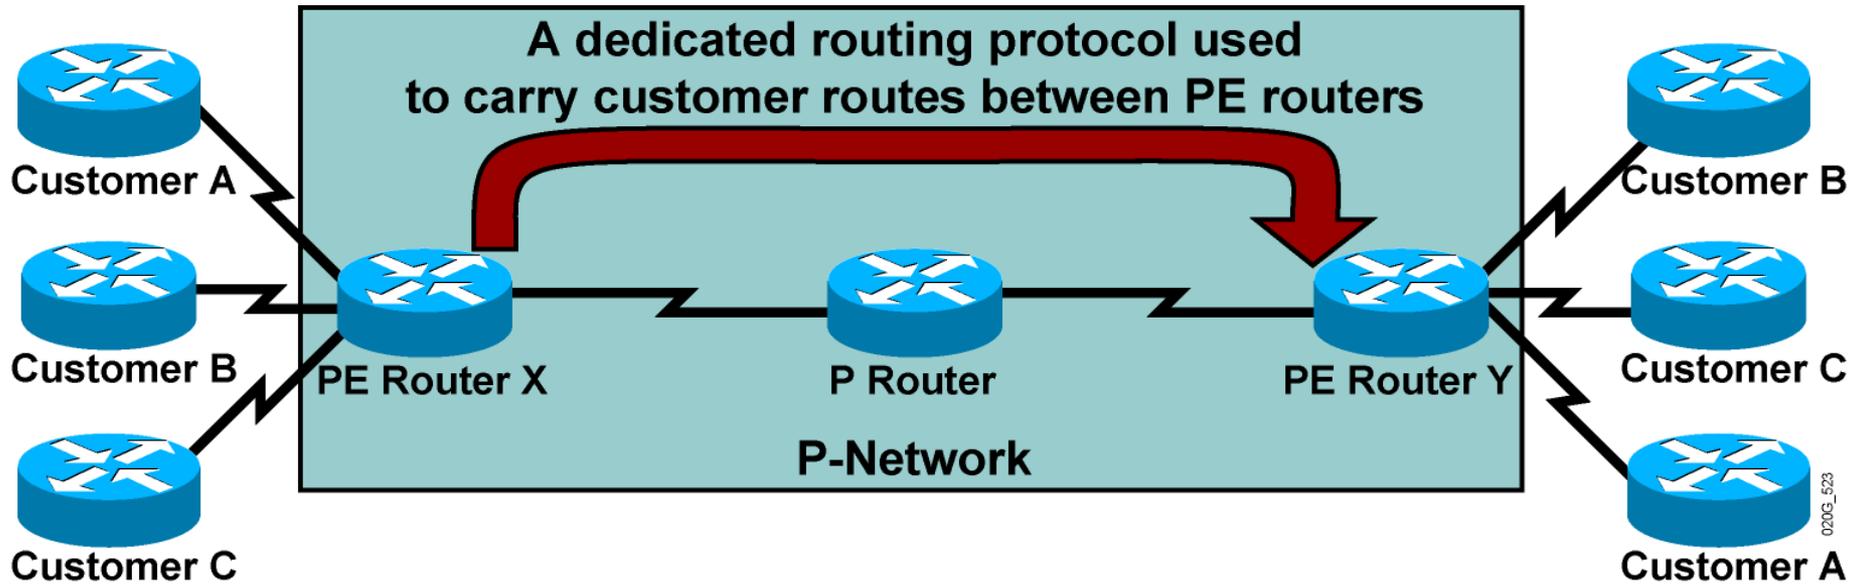
- P routers do not carry customer routes; the solution is scalable.

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



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

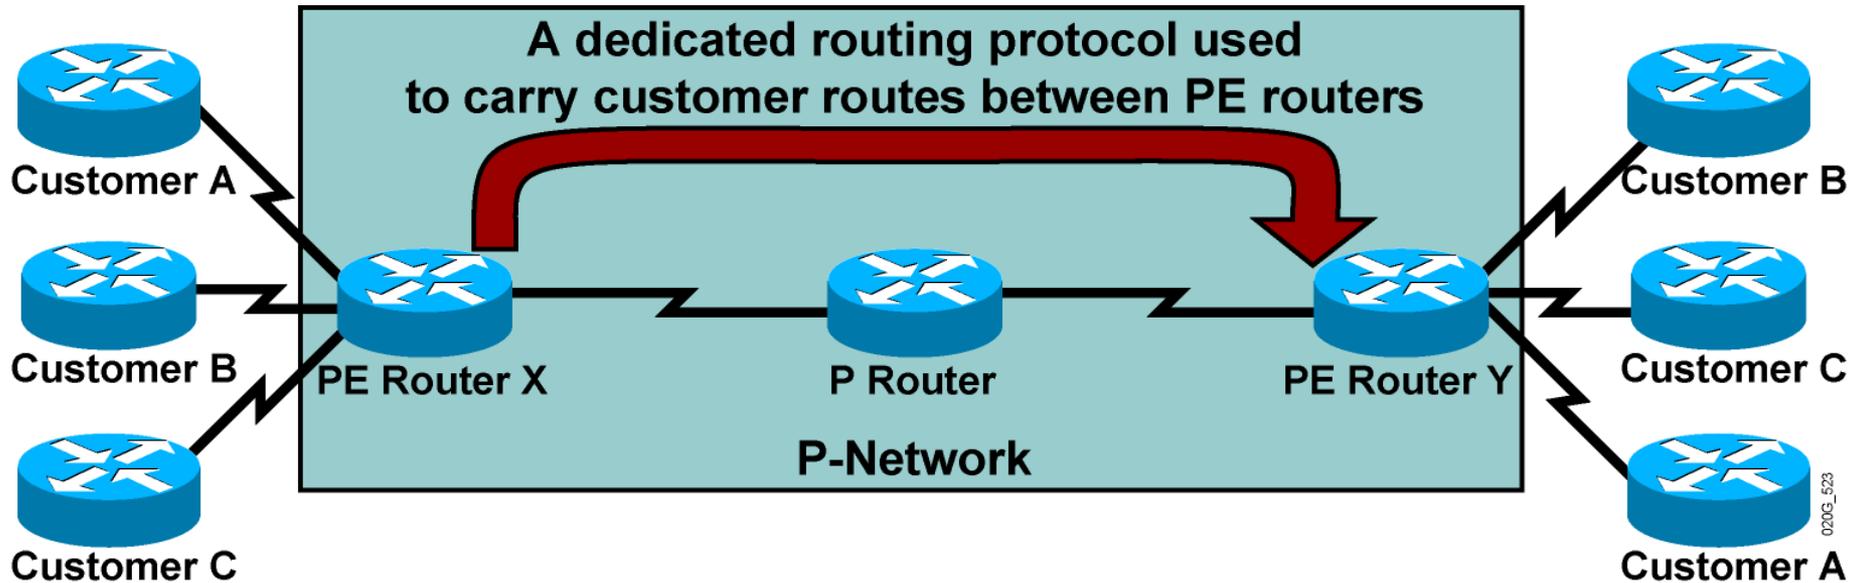
Propagation Routing Information Across the P-Network (Cont.)



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.

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



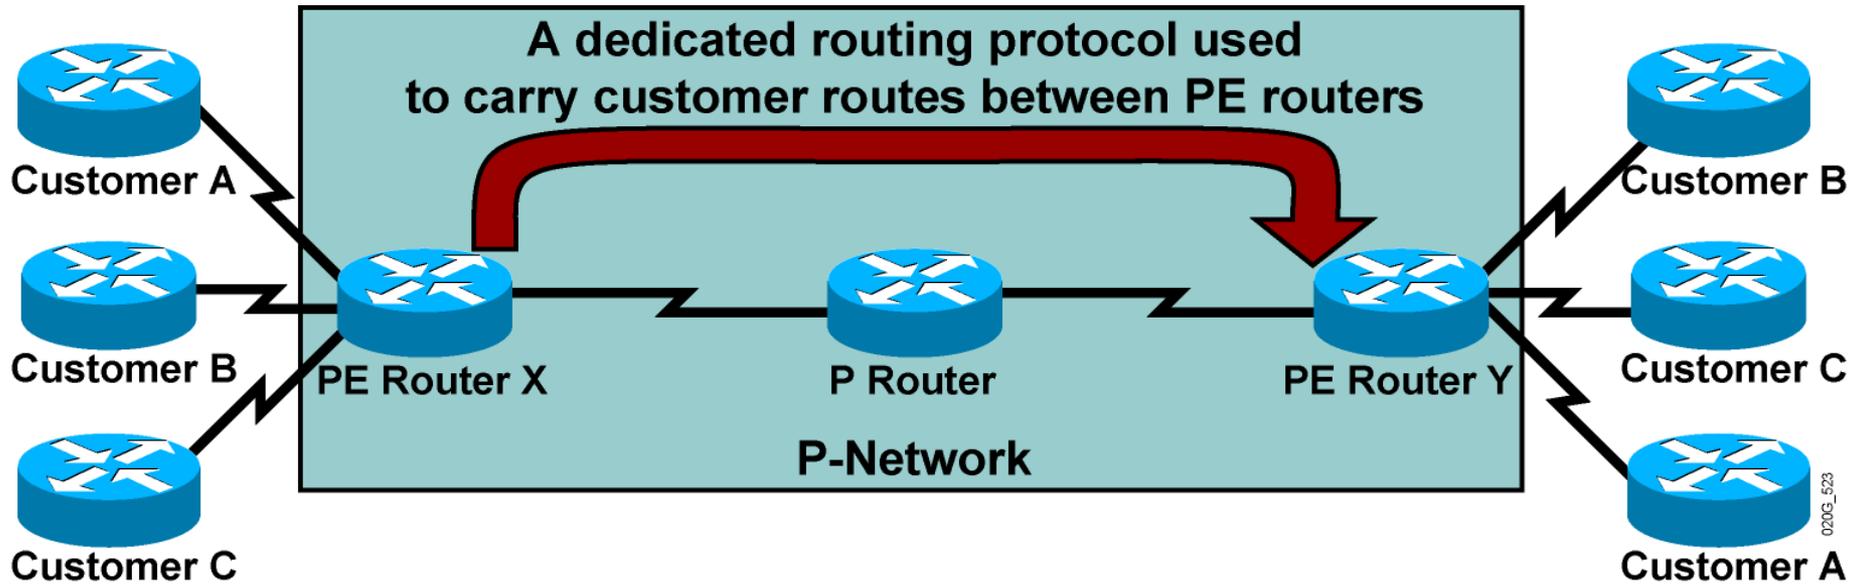
Question: Which protocol can be used to carry customer routes between PE routers?

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:

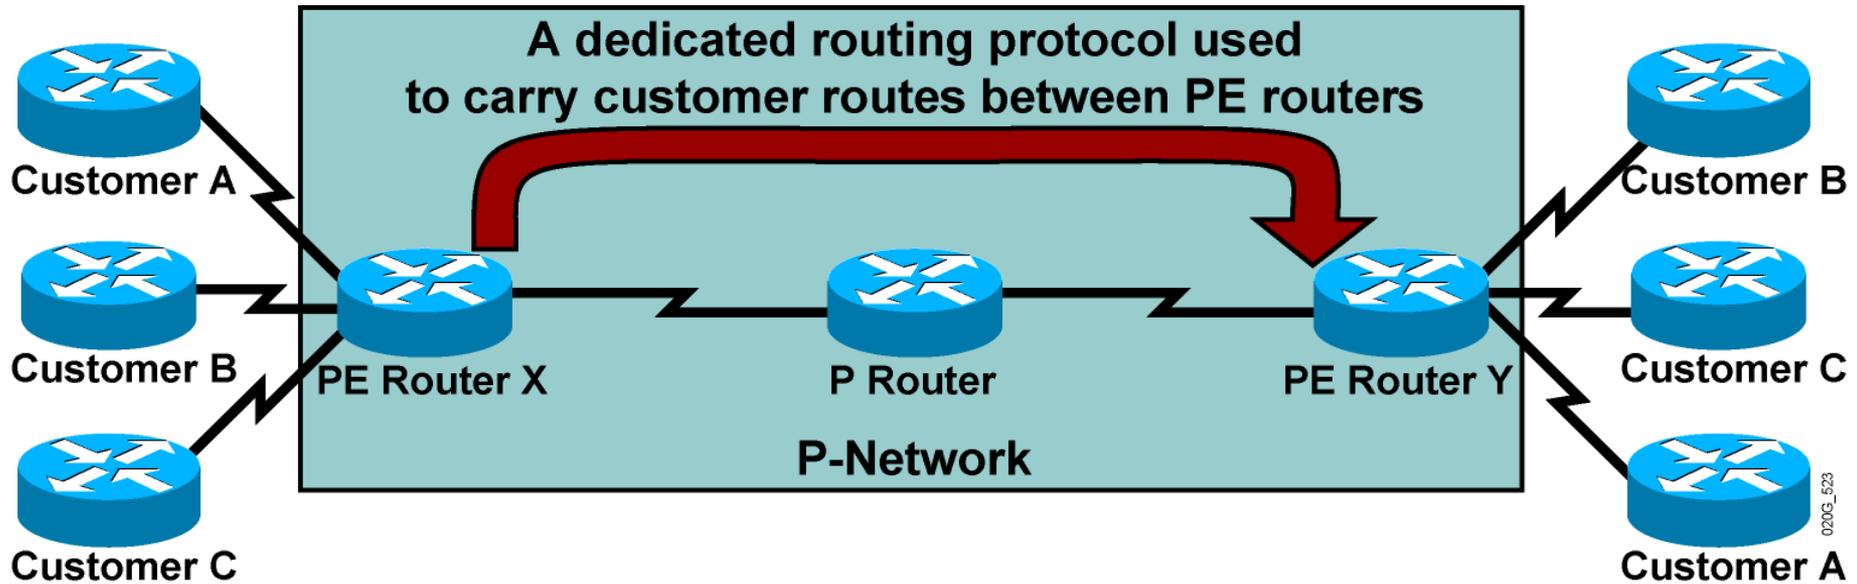
BGP is used to exchange customer routes directly between PE routers.

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?

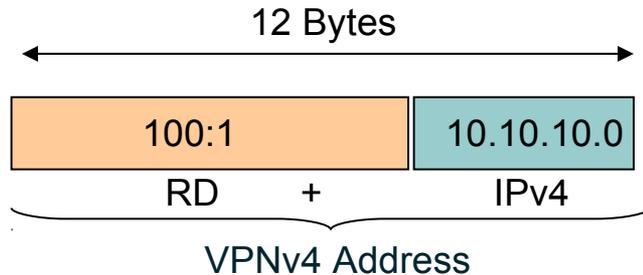
Propagation of Routing Information Across the P-Network (Cont.)



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.

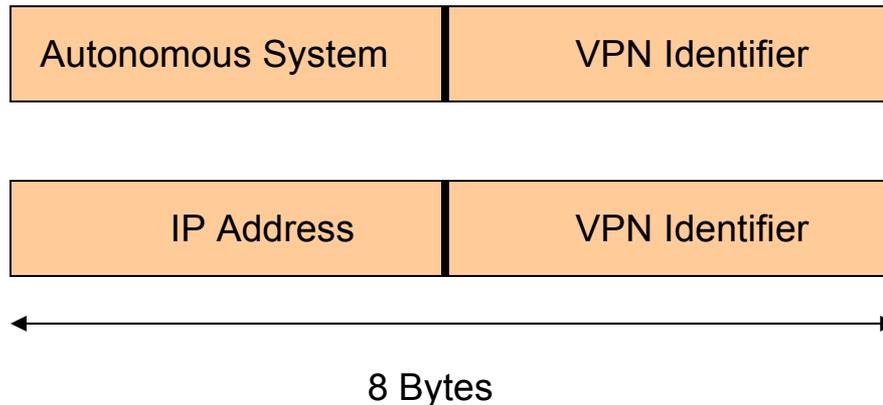
The resulting address is a VPNv4 address.

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.)

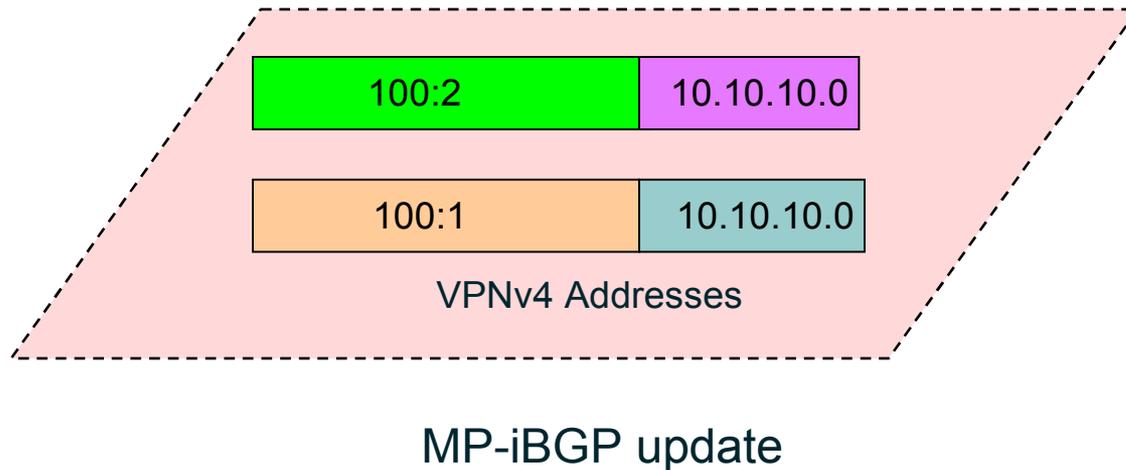
Route Distinguisher Format



Service Providers can use their BGP AS along with VPN customer identifier

Service Provider who do not have BGP AS, can use an IP address

Route Distinguishers (Cont.)

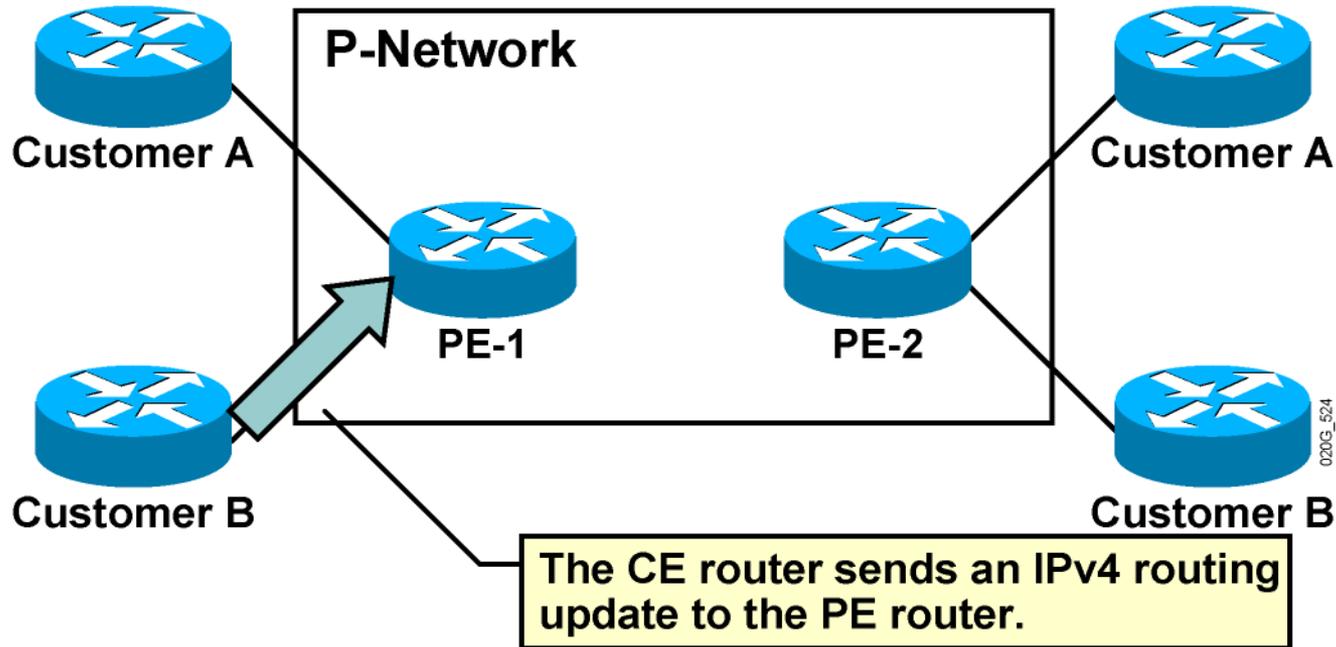


Customer A has RD of 100:1

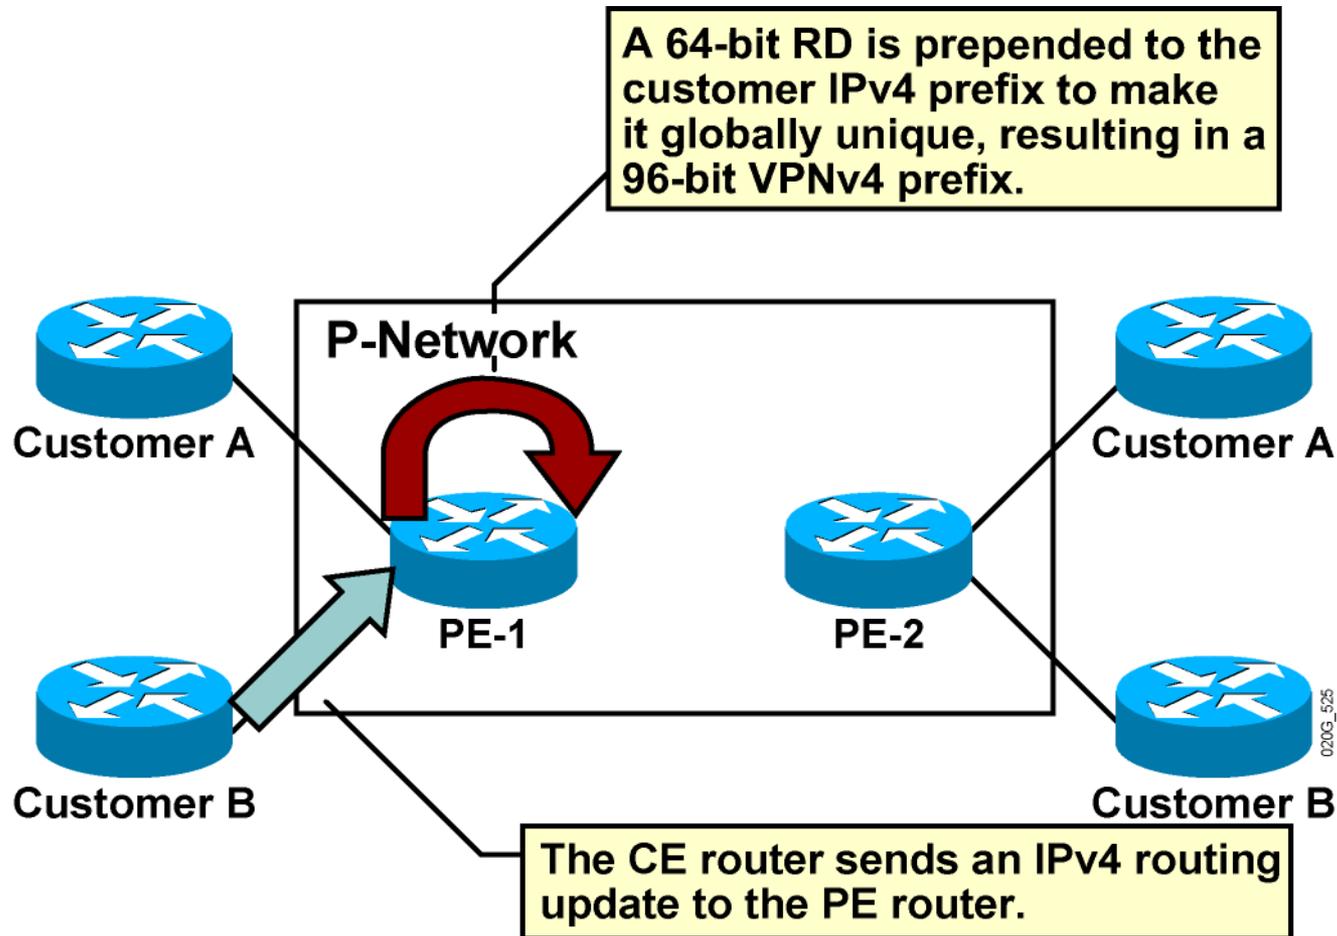
Customer B has RD of 100:2

Route Distinguisher keeps Customer A's update unique from Customer B in the MP-iBGP update, although they use the same IP address

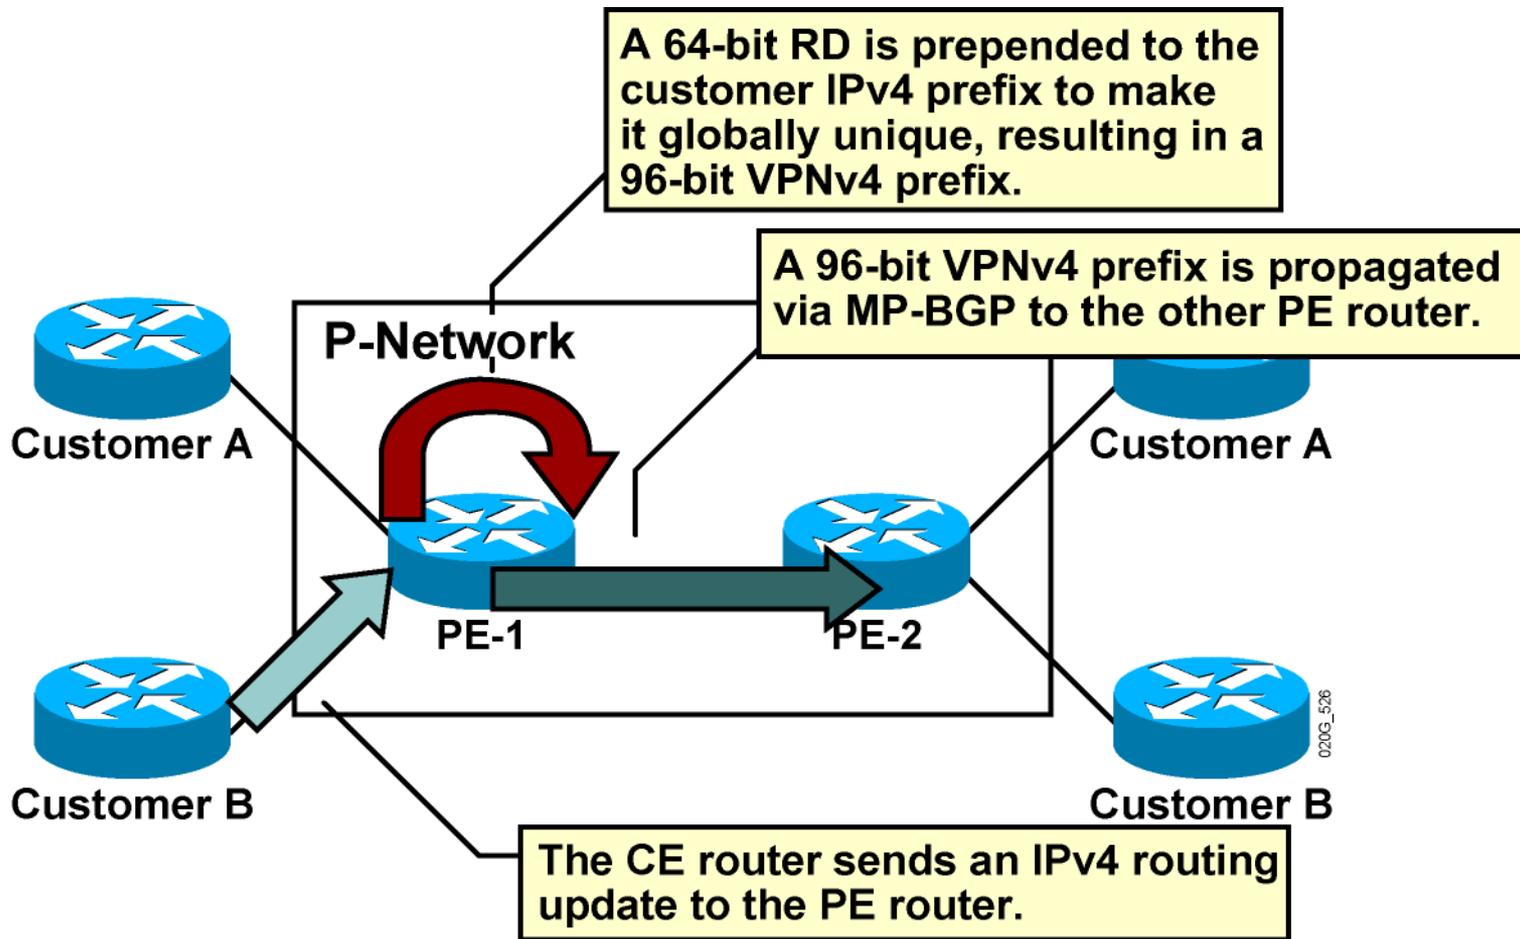
Route Distinguishers (Cont.)



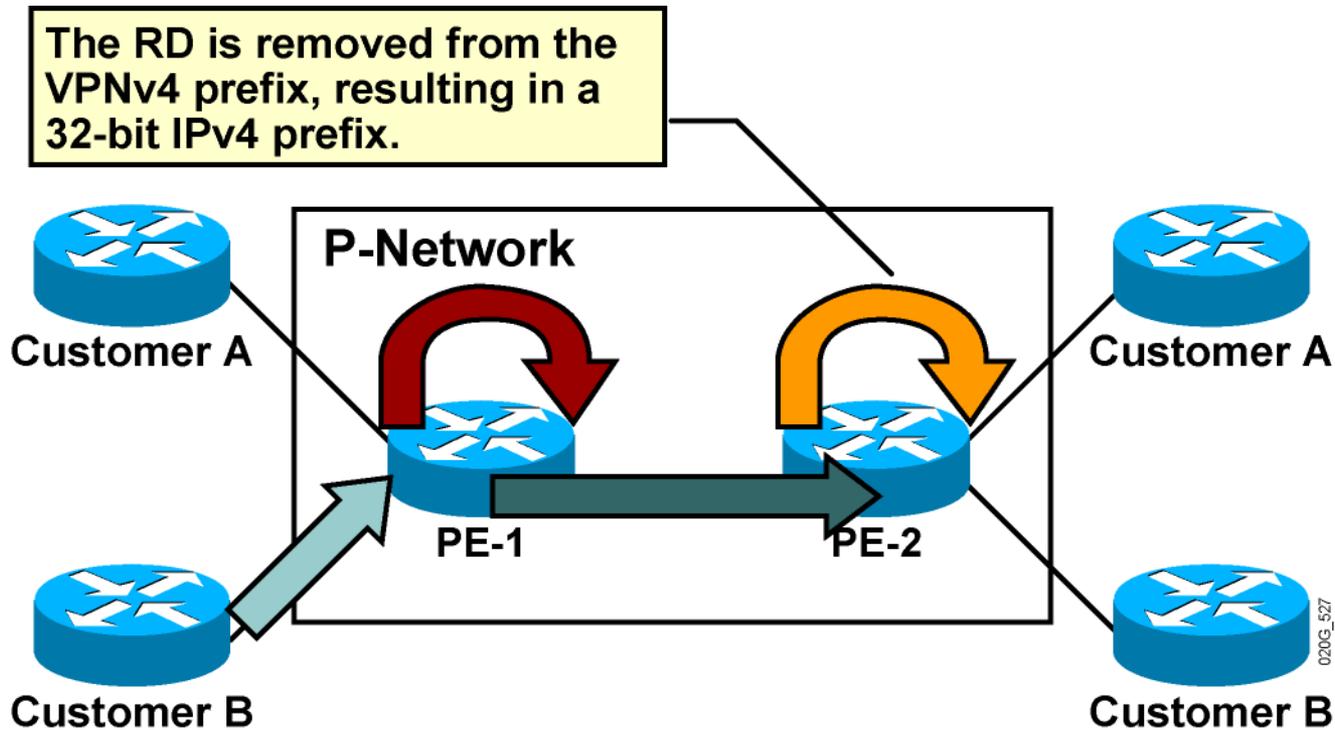
Route Distinguishers (Cont.)



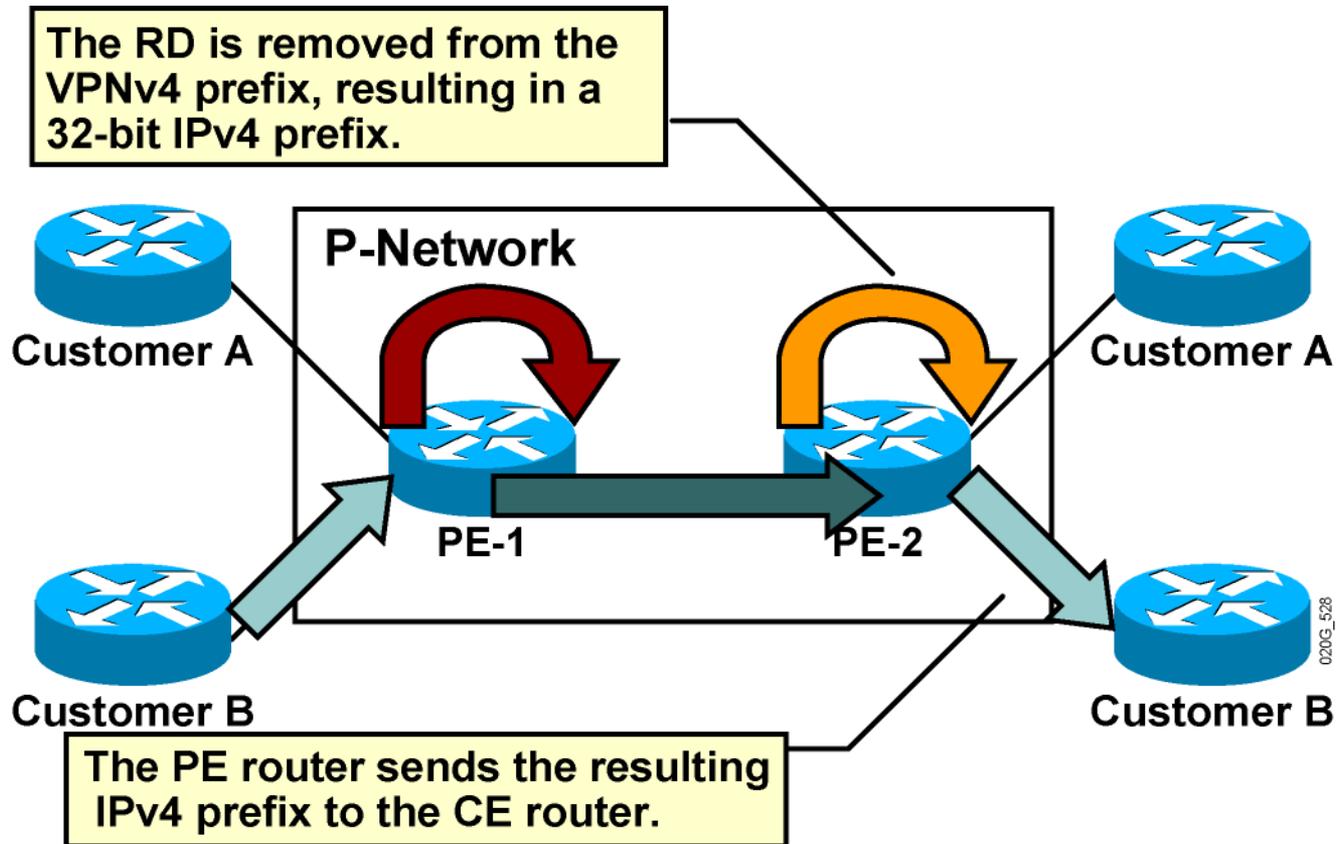
Route Distinguishers (Cont.)



Route Distinguishers (Cont.)



Route Distinguishers (Cont.)



Route Distinguishers (Cont.)

Usage in an MPLS VPN

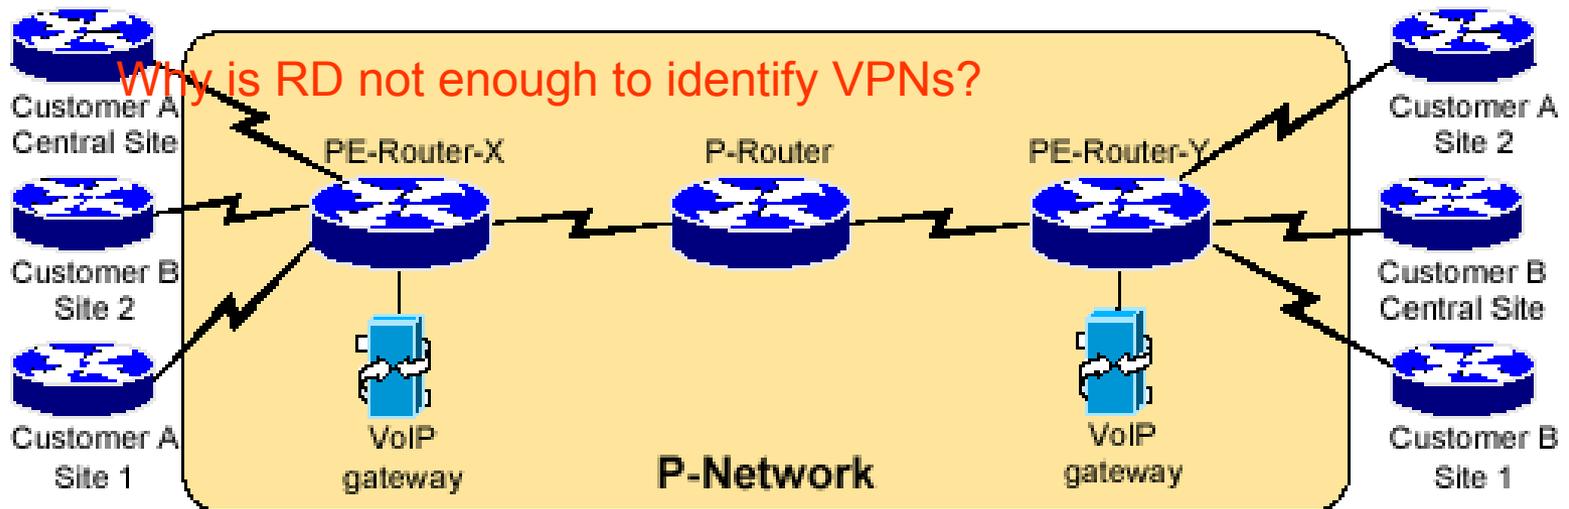
The RD has no special meaning.

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.

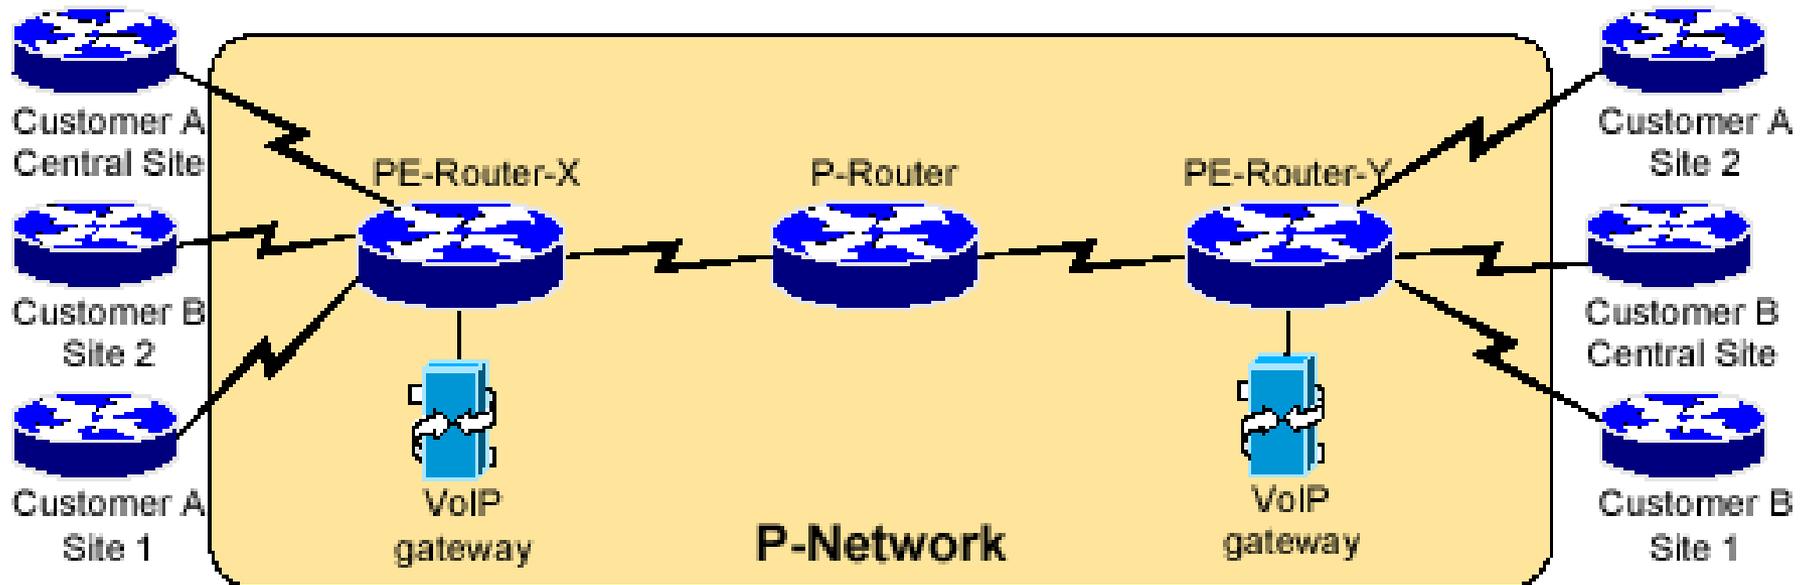
Route Targets

VoIP Service Sample



Route Targets

VoIP Service Sample

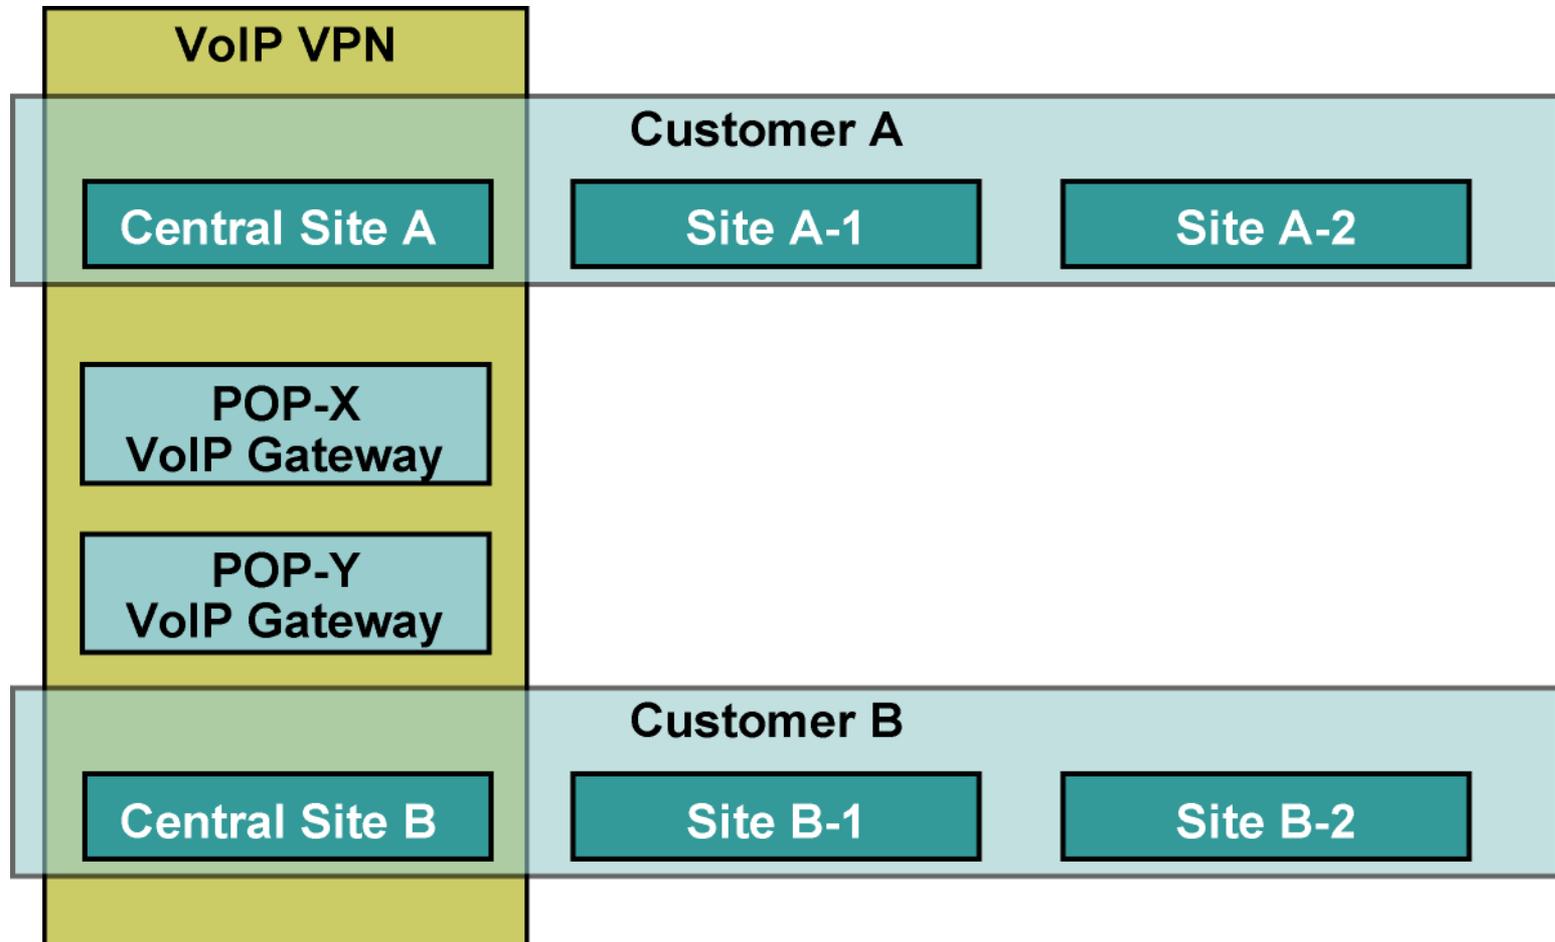


Requirements:

- 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



0206_532

Route Targets (Cont.)

Why Are They Needed?

Some sites have to participate in more than one VPN.

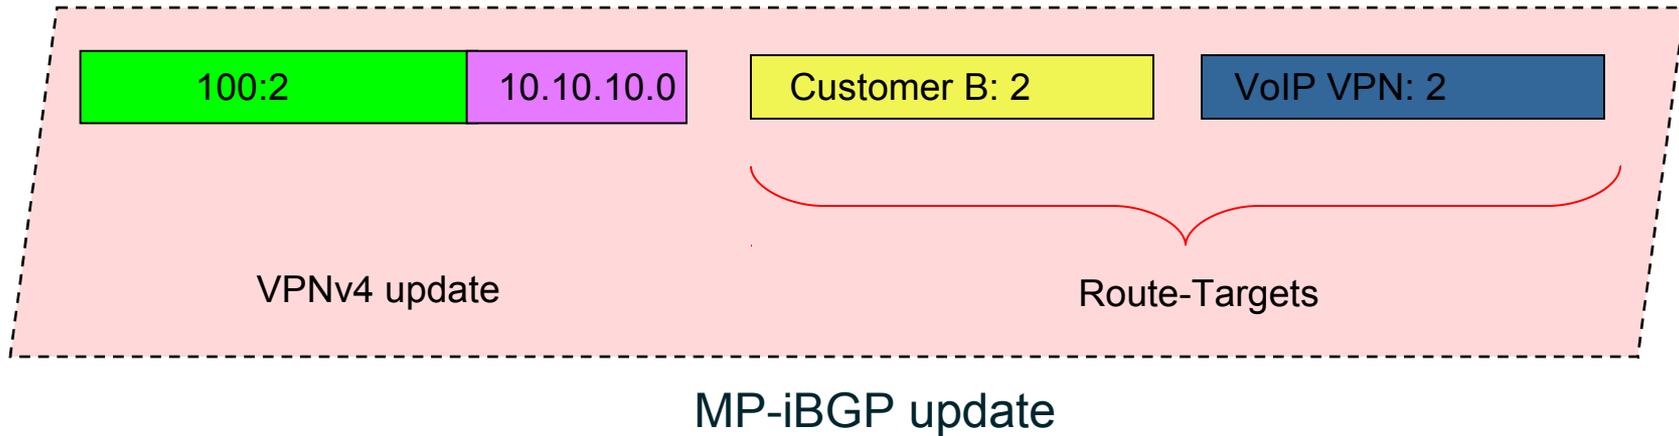
The RD cannot identify participation in more than one VPN.

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.

Format is same as Route Distinguisher

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

Virtual Private Networks Redefined

- With the introduction of complex VPN topologies, VPNs have had to be redefined:

A VPN is a collection of sites sharing common routing information.

A site can be part of different VPNs.

A VPN can be seen as a community of interest (closed user group, or CUG).

Complex VPN topologies are supported by multiple virtual routing tables on the PE routers.

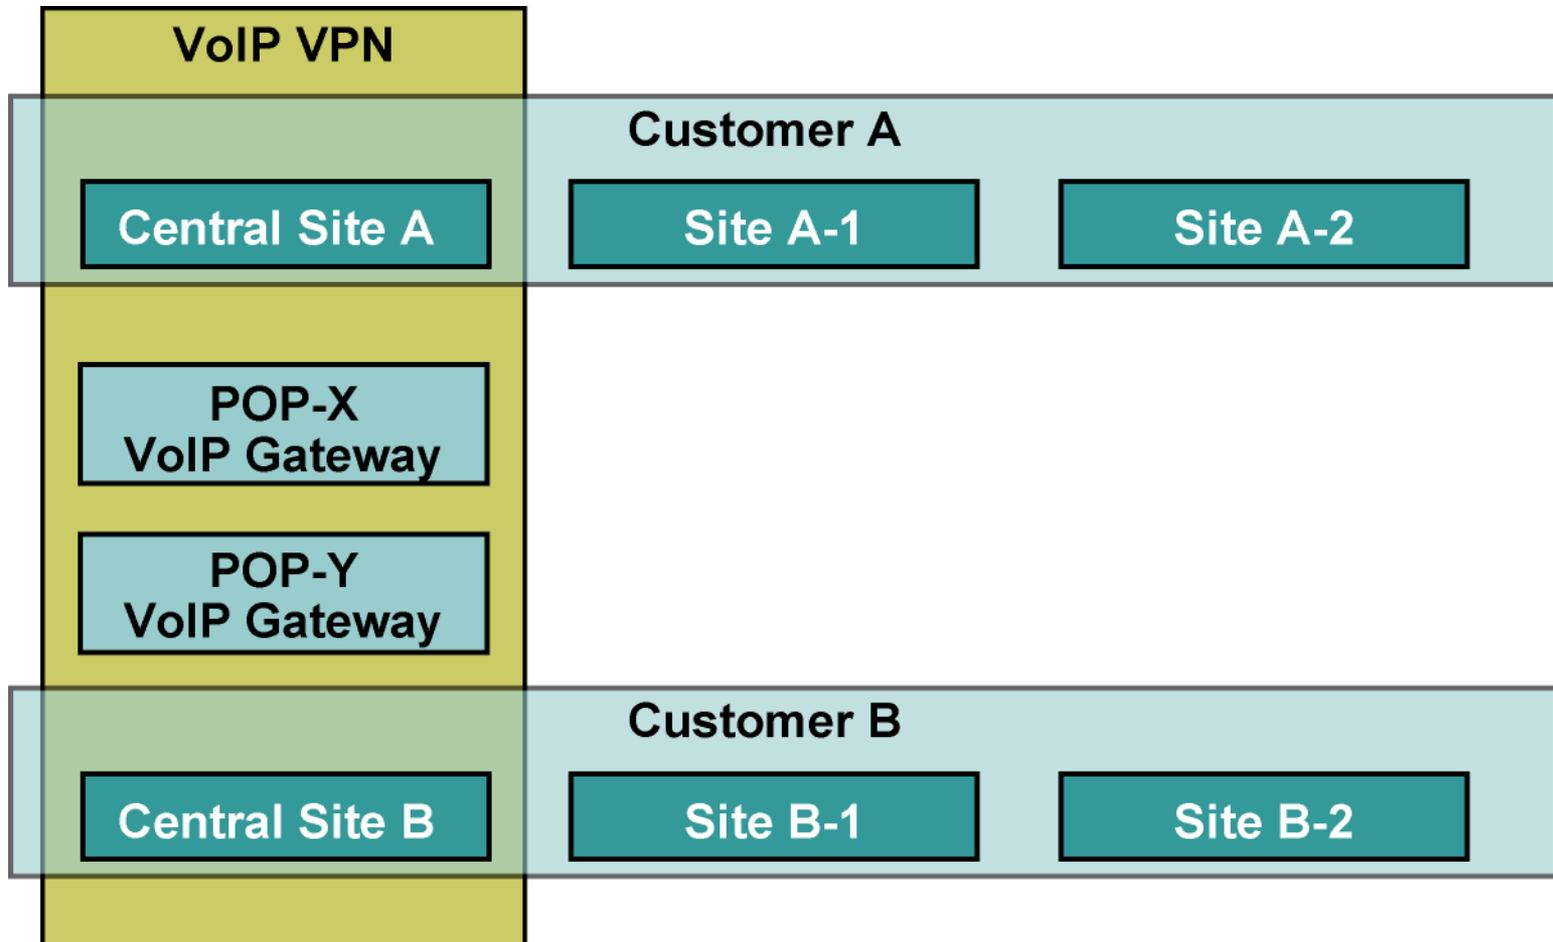
Impact of Complex VPN Topologies on Virtual Routing Tables

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.

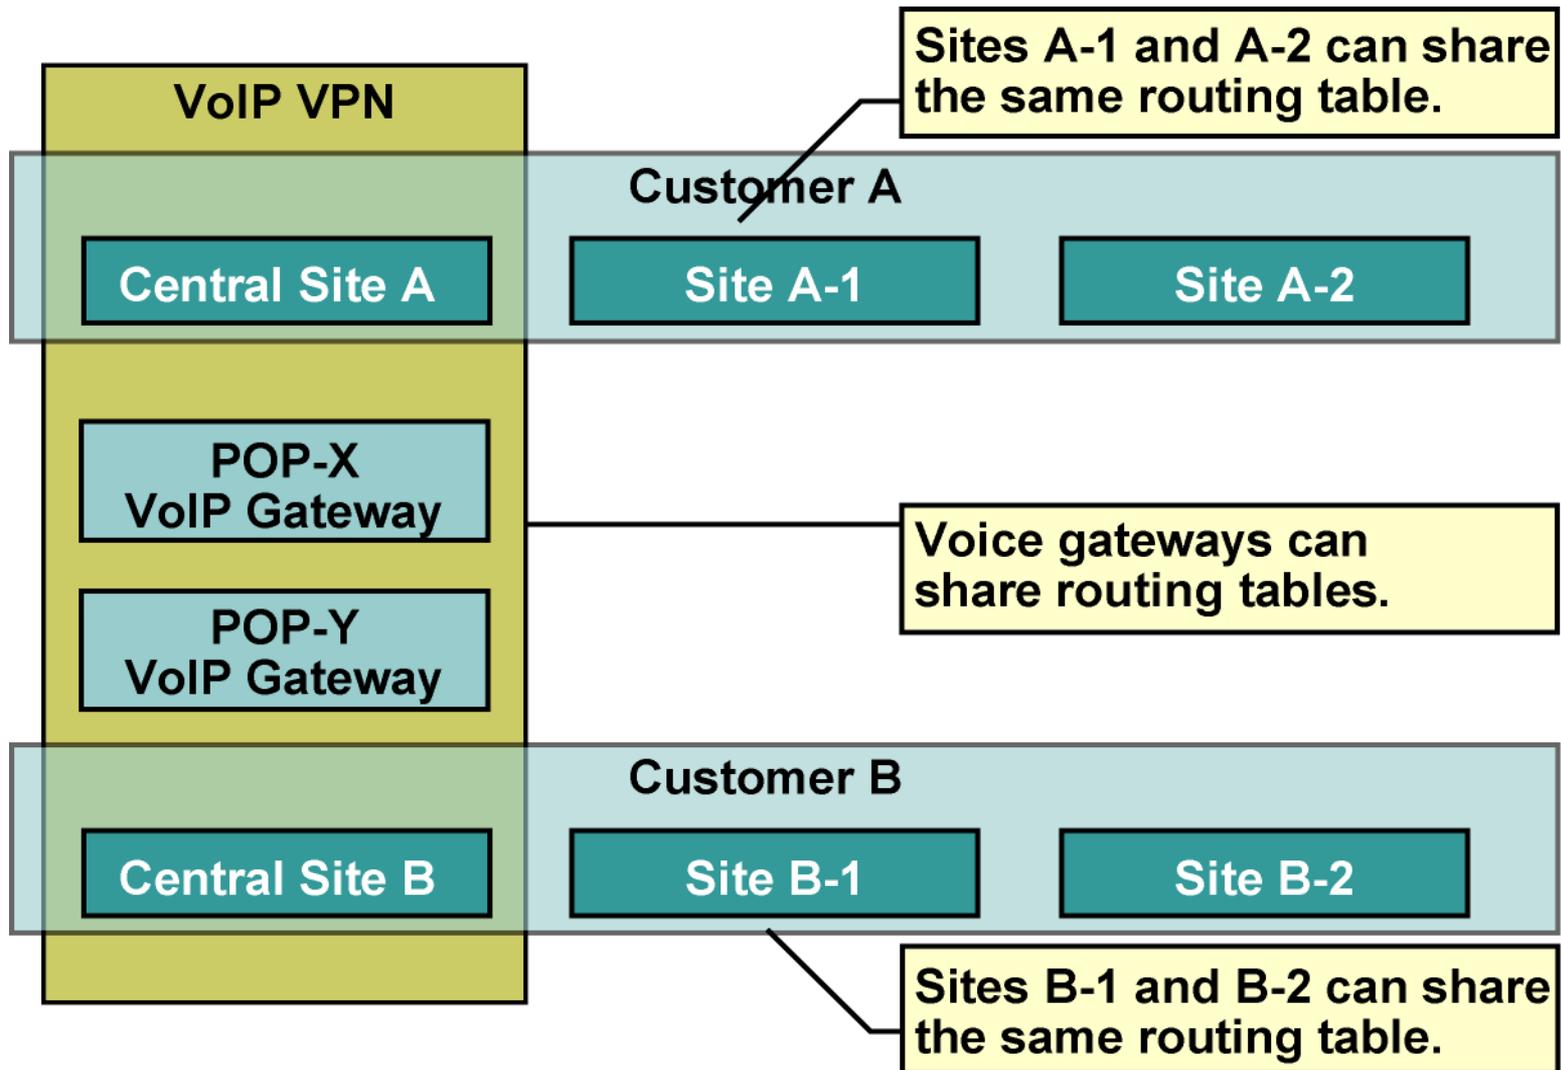
As each virtual routing table requires a distinct RD value, the number of RDs in the MPLS VPN network increases.

Impact of Complex VPN Topologies on Virtual Routing Tables (Cont.)

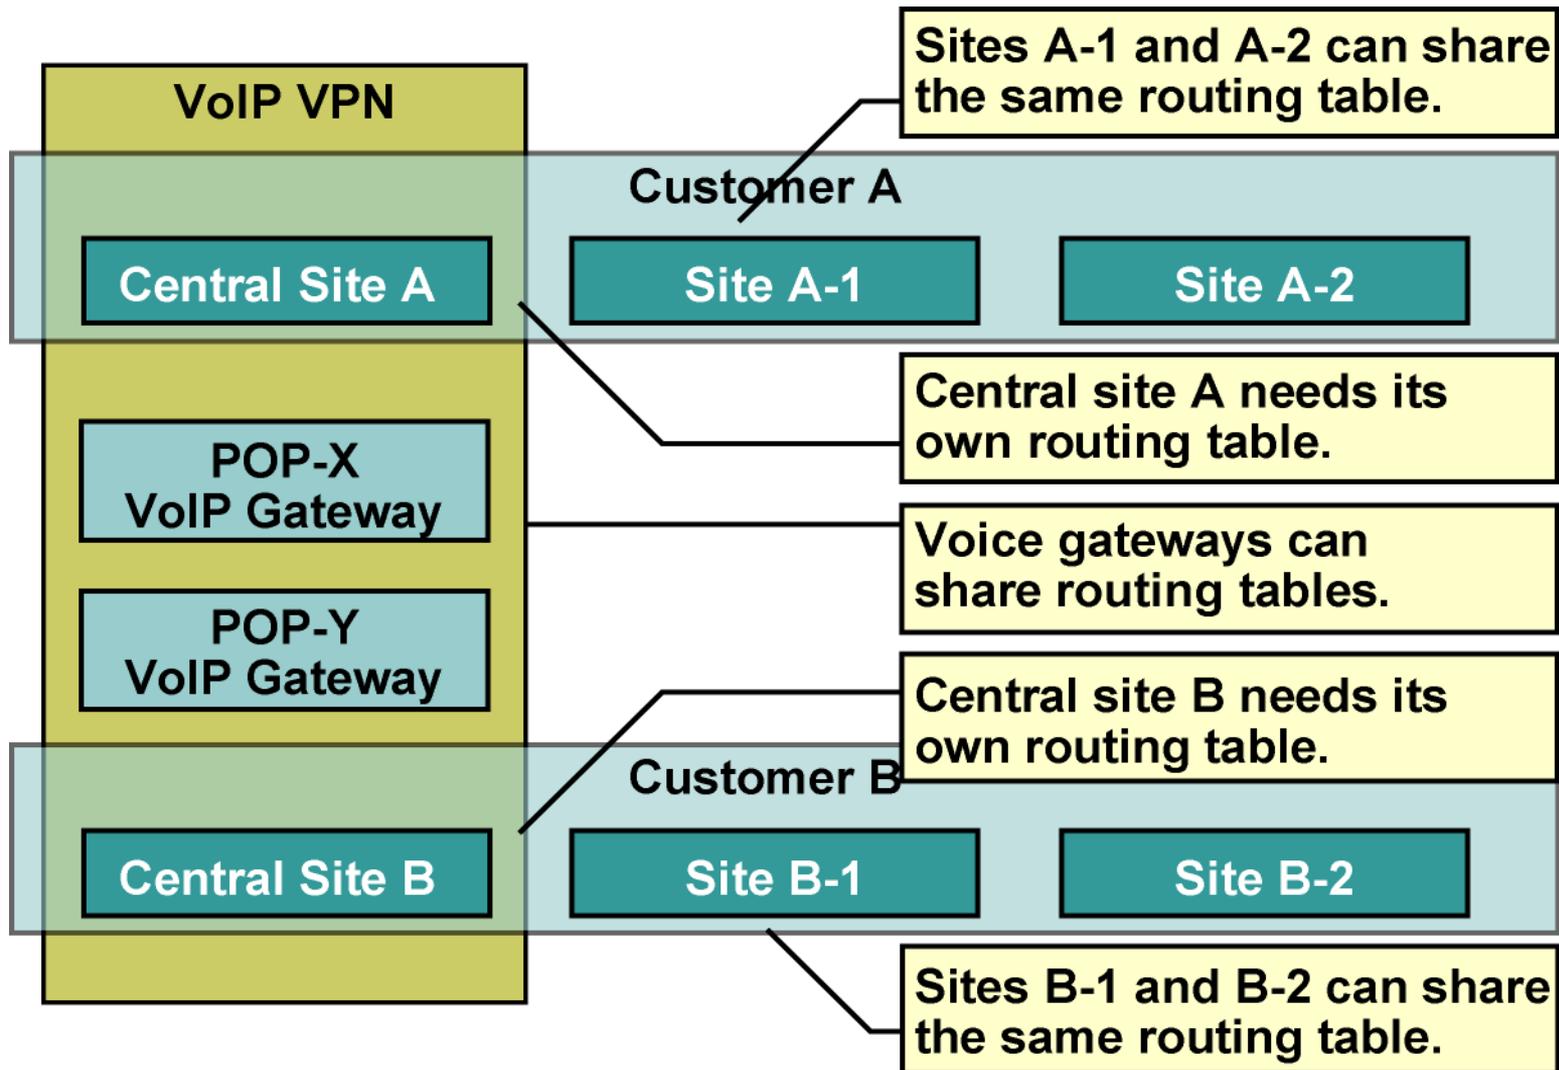


0200_532

Impact of Complex VPN Topologies on Virtual Routing Tables (Cont.)



Impact of Complex VPN Topologies on Virtual Routing Tables (Cont.)



0206_530

Important points to note for RT and RD

Route Distinguishers (RD) are only used to make ipv4 VPN addresses unique when advertising them over MP-iBGP, by making them vpnv4 prefixes

We can have one RD per vrf

Only one vrf can be assigned to an interface

Route Targets (RT) are used for VPN membership, so that complex scenarios can be addressed

VPN is the set of rules for customer connectivity and can be very complex

A VPN may have several RTs

Summary

MPLS VPN architecture combines the best features of the overlay and peer-to-peer VPN models.

Virtual routing tables are created for each customer.

BGP is used to exchange customer routes between PE routers.

Route distinguishers transform non-unique 32-bit addresses into 96-bit unique addresses.

Route targets are used to identify VPN membership in overlapping topologies.

Placing sites with different routing requirements in the same virtual routing table will result in inconsistent routing.



MPLS workshop

MPLS VPN Routing Model

Outline

Overview

MPLS VPN Routing Requirements

MPLS VPN Routing

Support for Existing Internet Routing

Routing Tables on PE Routers

End-to-End Routing Update Flow

Route Distribution to CE Routers

Lesson Summary

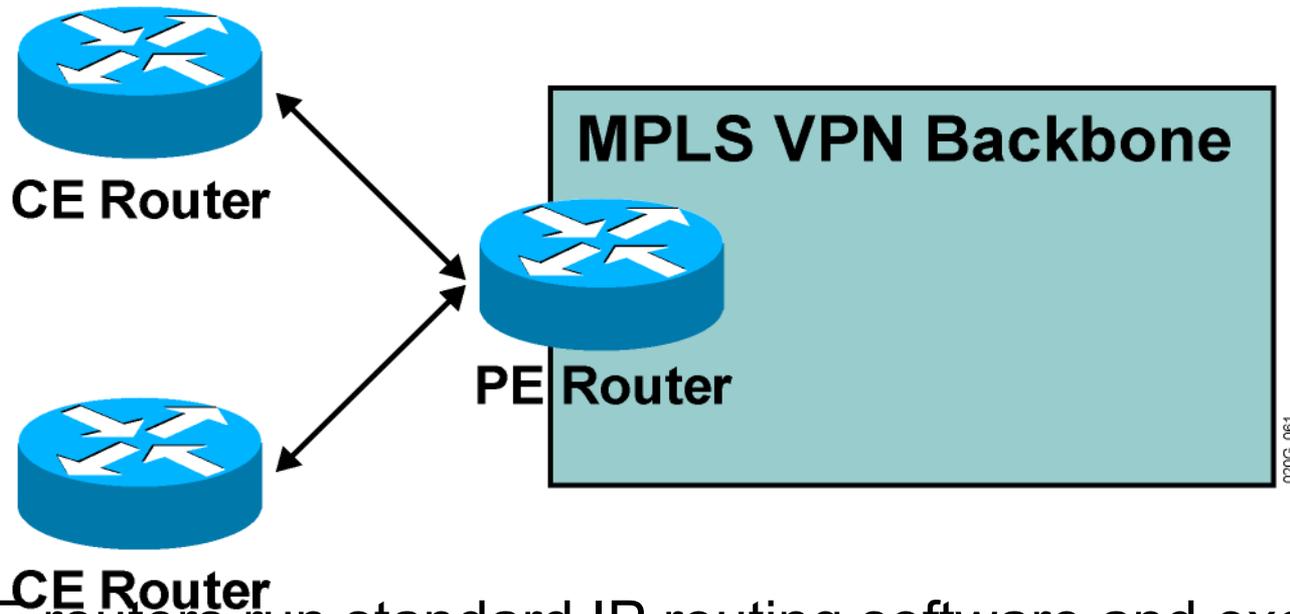
MPLS VPN Routing Requirements

CE routers have to run standard IP routing software.

PE routers have to support MPLS VPN services and Internet routing.

P routers have no VPN routes.

MPLS VPN Routing CE Router Perspective



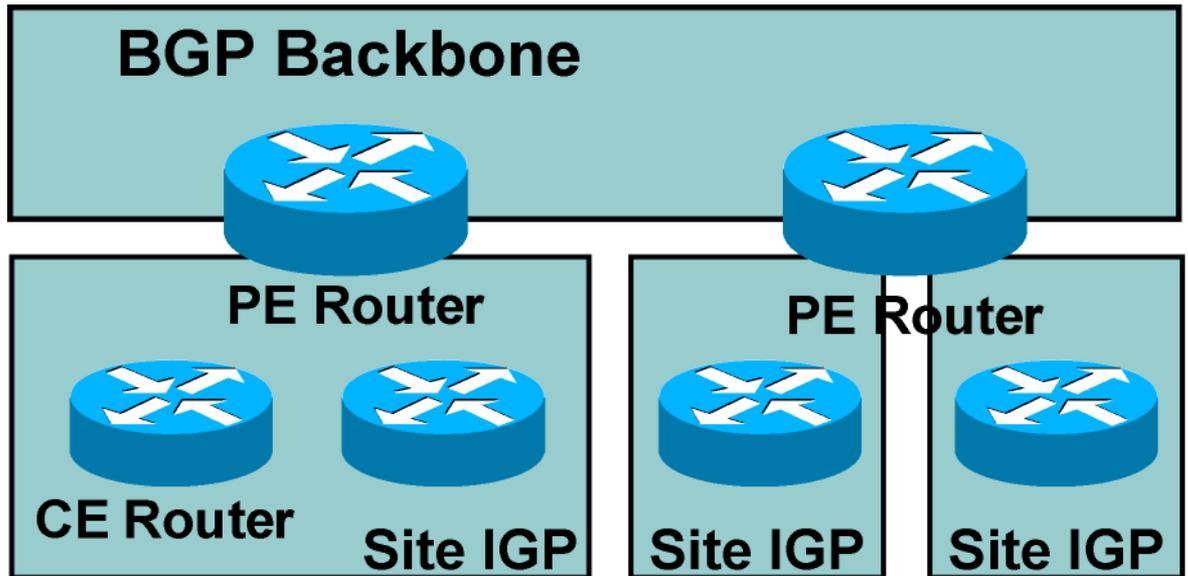
The CE routers run standard IP routing software and exchange routing updates with the PE router.

PE-CE protocols can be EBGP, OSPF, RIPv2, EIGRP, and static routes. ISIS support in the works

The PE router appears as another router in the C-network.

MPLS VPN Routing (cont.)

Overall Customer Perspective



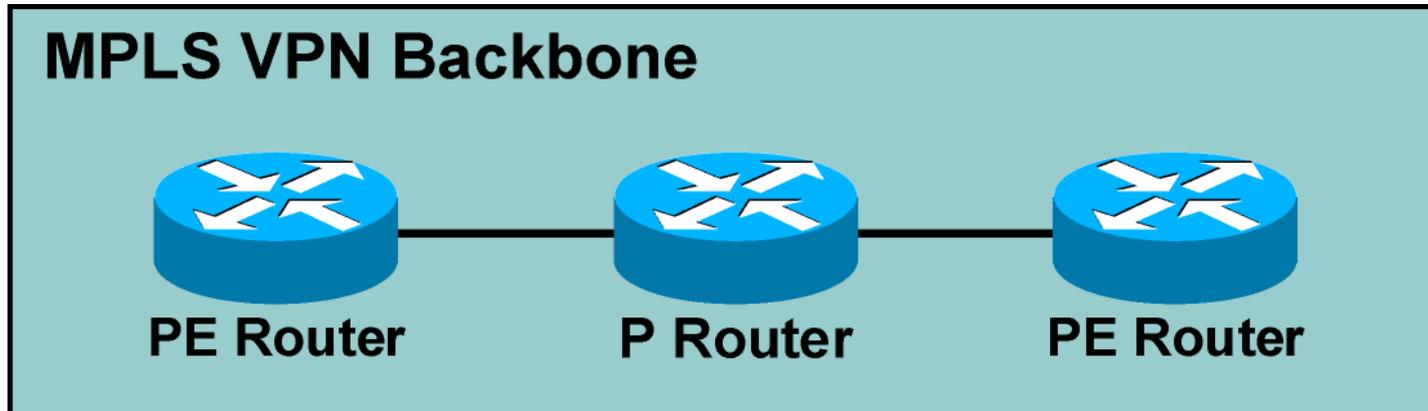
To the customer, the PE routers appear as core routers connected via a BGP backbone.

The usual BGP and IGP design rules apply.

The P routers are hidden from the customer.

MPLS VPN Routing (Cont.)

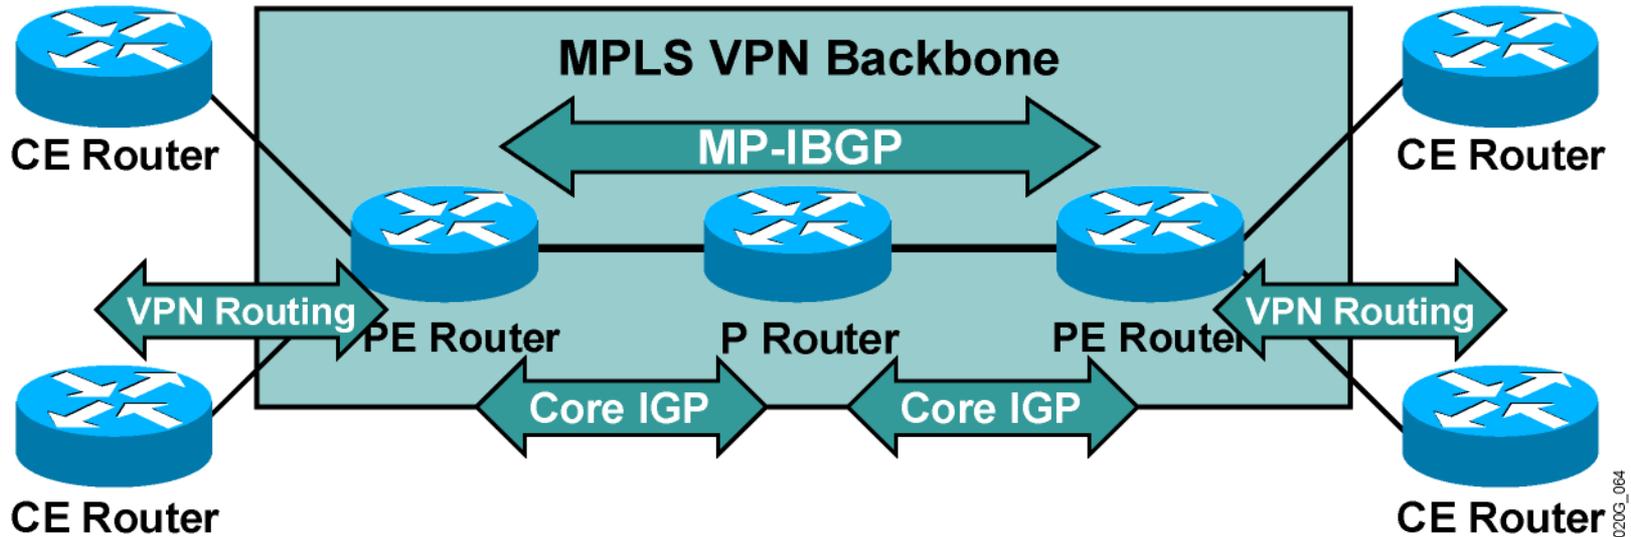
P Router Perspective



- 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.)

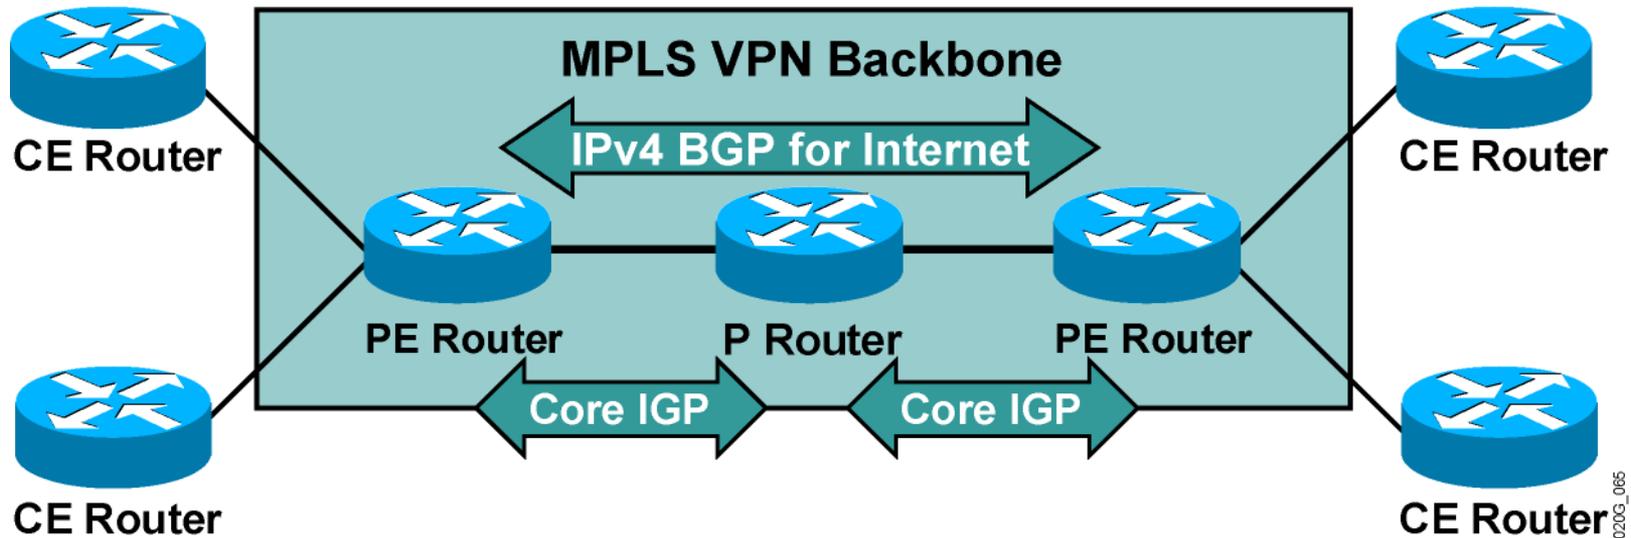
PE Router Perspective



PE routers:

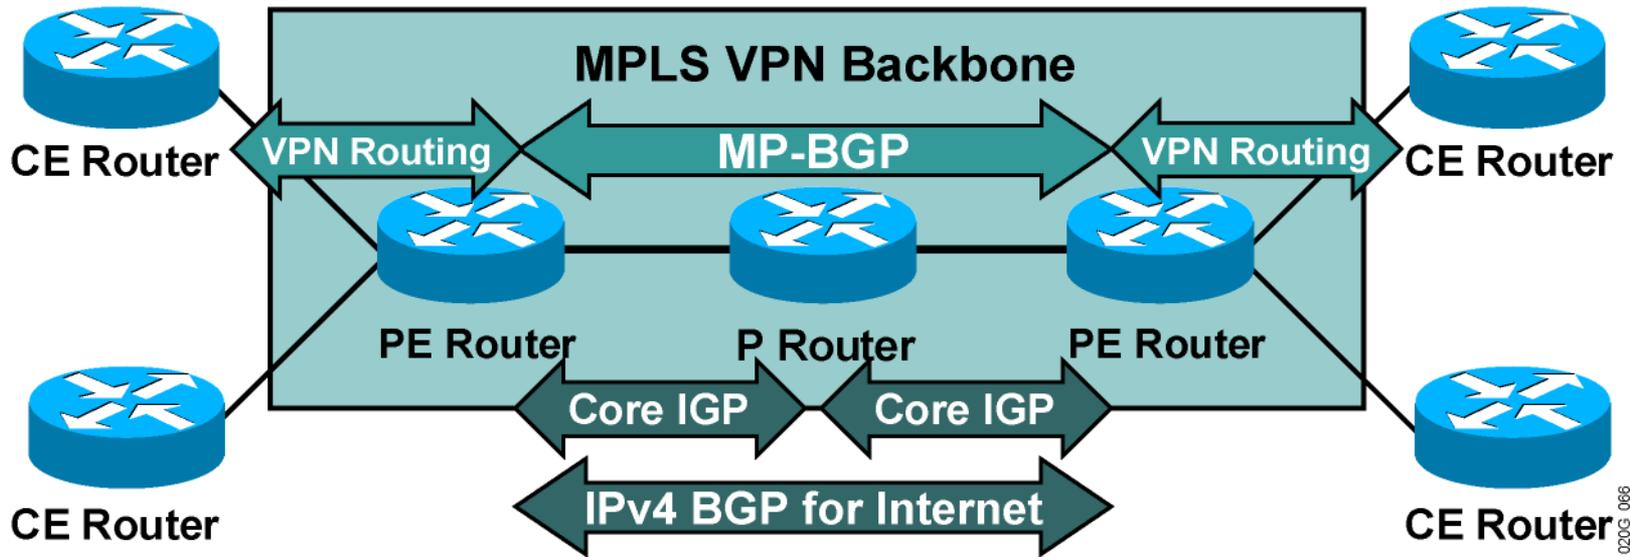
- Exchange VPN routes with CE routers via per-VPN routing protocols
 - Exchange core routes with P routers and PE routers via core IGP
- Exchange VPNv4 routes with other PE routers via MP-IBGP sessions

Support for Existing Internet Routing



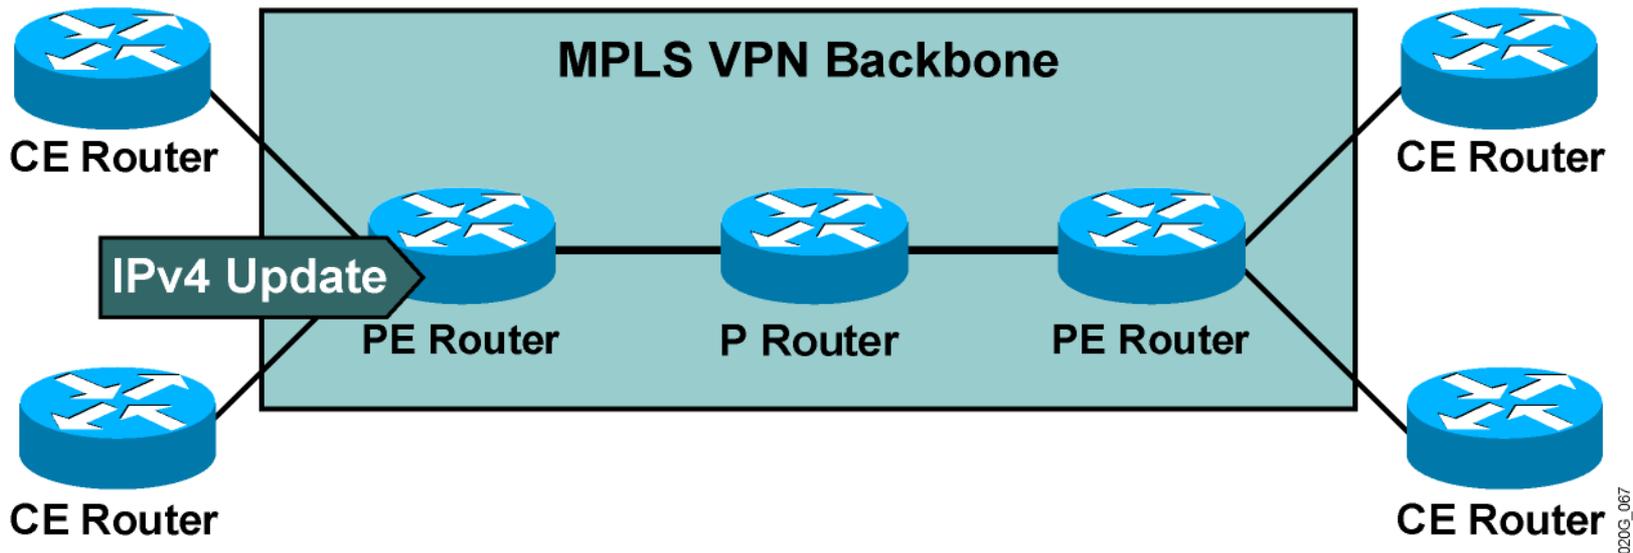
- PE routers can run standard IPv4 BGP in the global routing table:
 - PE routers exchange Internet routes with other PE routers.
 - CE routers do not participate in Internet routing.
 - P routers do not need to participate in Internet routing.

Routing Tables on PE Routers



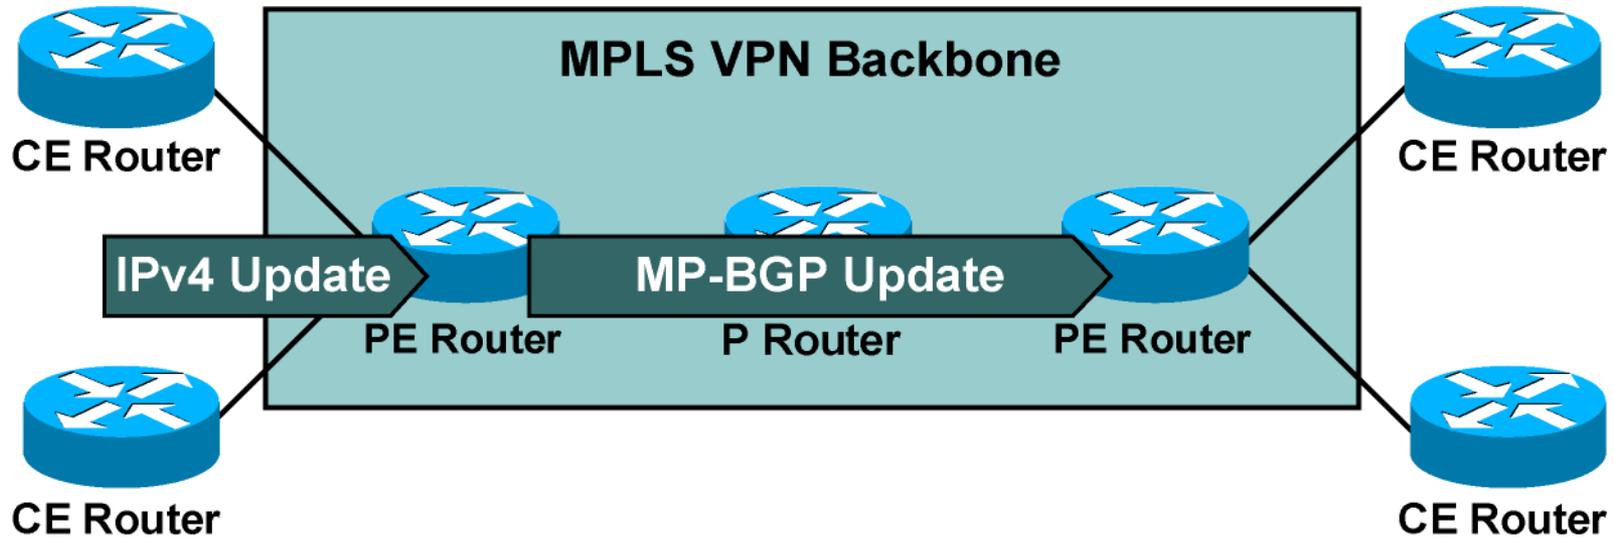
- PE routers contain a number of routing tables:
 - Global routing table**, which contains core routes (filled with core IGP) and Internet routes (filled with IPv4 BGP)
 - VRF tables** for sets of sites with identical routing requirements
 - VRFs** filled with information from CE routers and MP-BGP information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

End-to-End Routing Update Flow



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

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



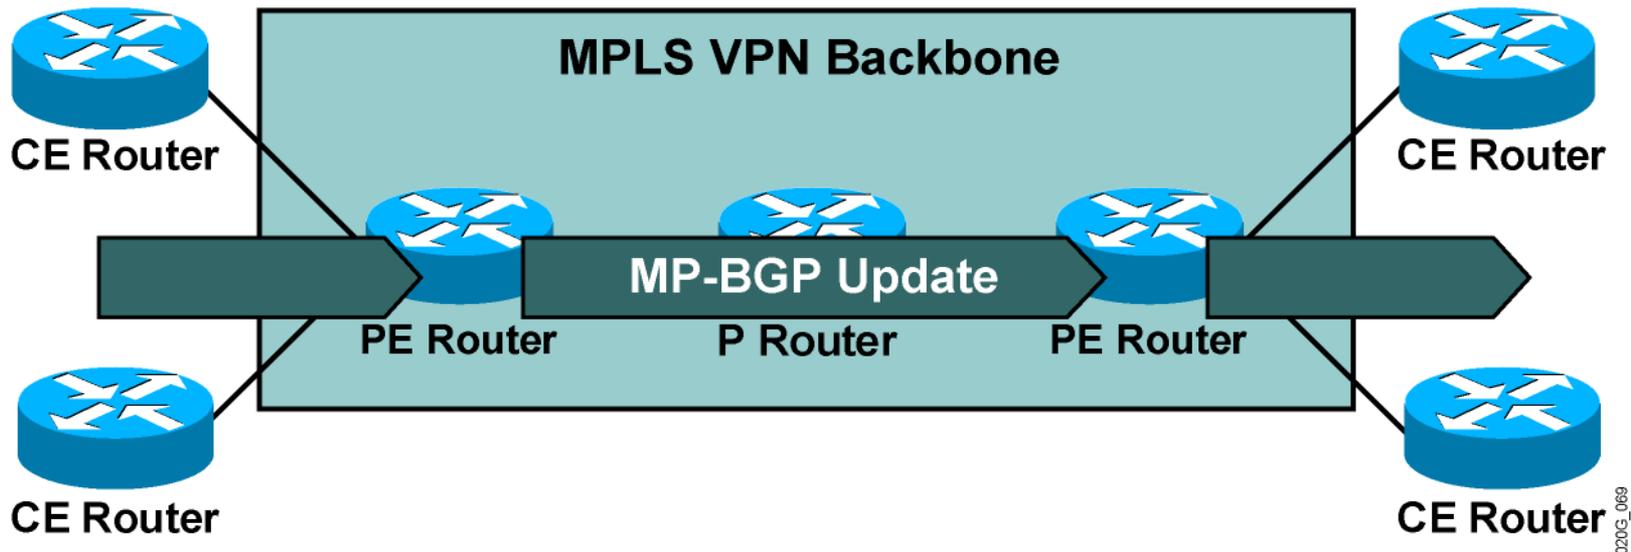
- PE routers export VPN routes from VRF tables into MP-BGP and propagate them as VPNv4 routes to other PE routers.

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

MP-BGP Update

- An MP-BGP update contains the following:
 - VPNv4 address
 - Extended communities
(route targets, optionally SOO)
 - 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.)



- Receiving PE router imports incoming VPNv4 routes into the appropriate VRF based on route targets attached to the routes.
- Routes installed in VRF are propagated to CE routers.

Route Distribution to CE Routers

Route distribution to sites is driven by the following:

SOO

RT BGP communities

A route is installed in the site VRF that matches the RT attribute.

Summary

MPLS VPNs technology does the following:

- Supports the use of standard IP routing between devices

- Provides scalable solutions

- Supports both MPLS VPNs and traditional Internet services

The internal service provider topology is transparent to the customer.

PE routers alone see all routing aspects of the MPLS VPN.

VRF tables contain sets of routes for sites with identical routing requirements.

Routes are transported using the following:

- IGP (internal core routes)

- BGP IPv4 (core Internet routes)

- BGP VPNv4 (PE-to-PE VPN routes)



MPLS workshop

MPLS VPN Packet Forwarding

Outline

Overview

VPN Packet Forwarding Across an MPLS VPN Backbone

VPN Penultimate Hop Popping

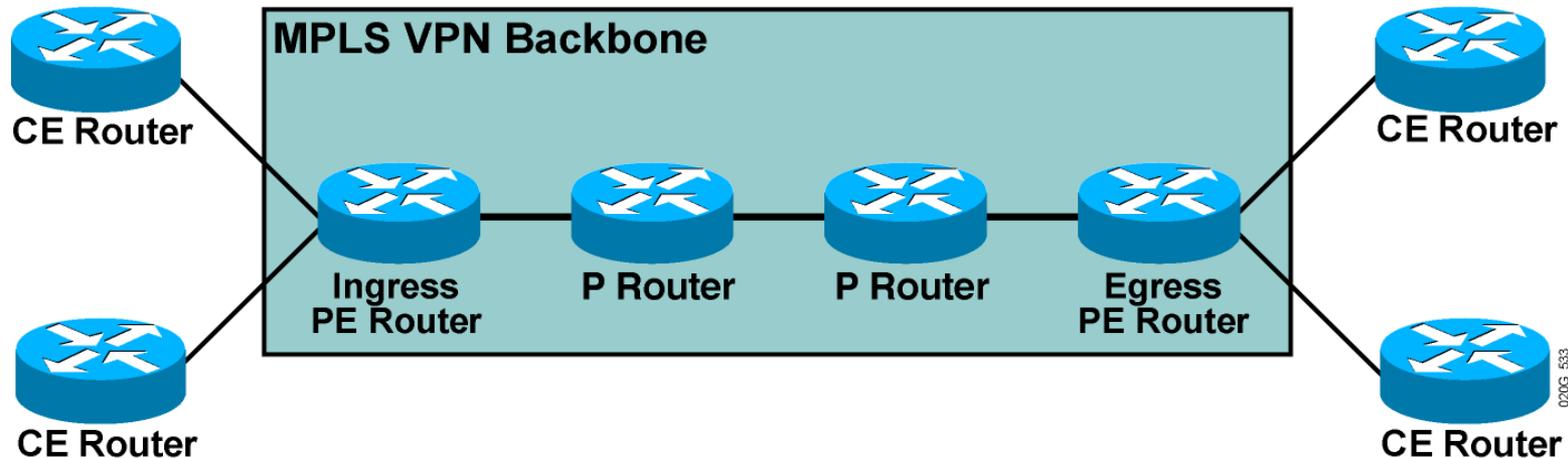
VPN Label Propagation

MPLS VPN and Label Propagation

MPLS VPN and Packet Forwarding

Lesson Summary

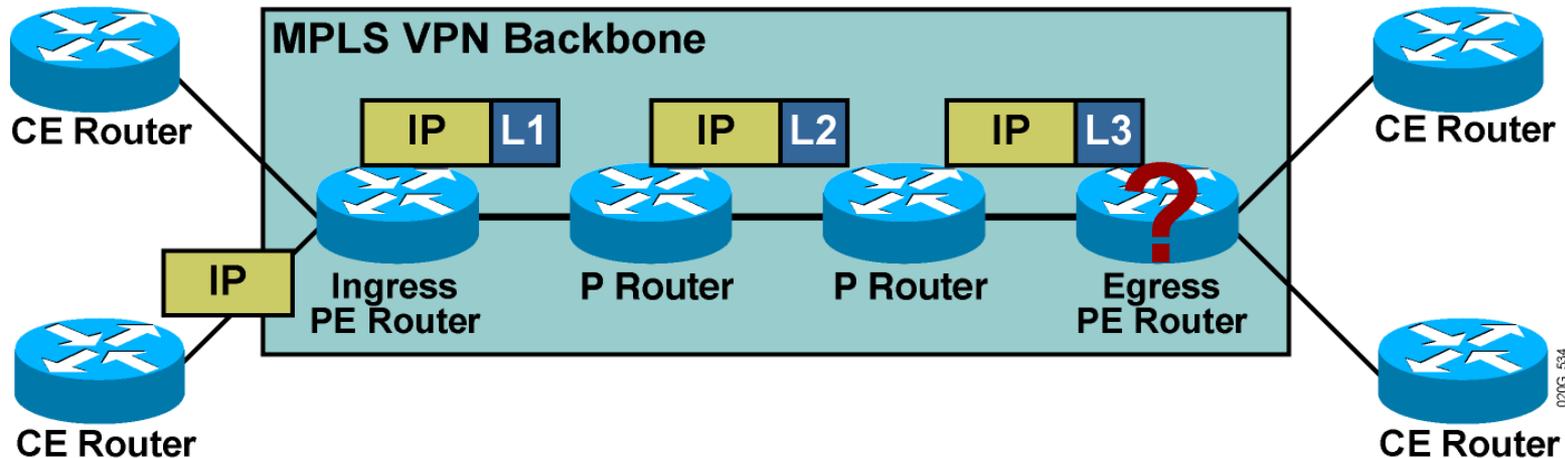
VPN Packet Forwarding Across an MPLS VPN Backbone



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



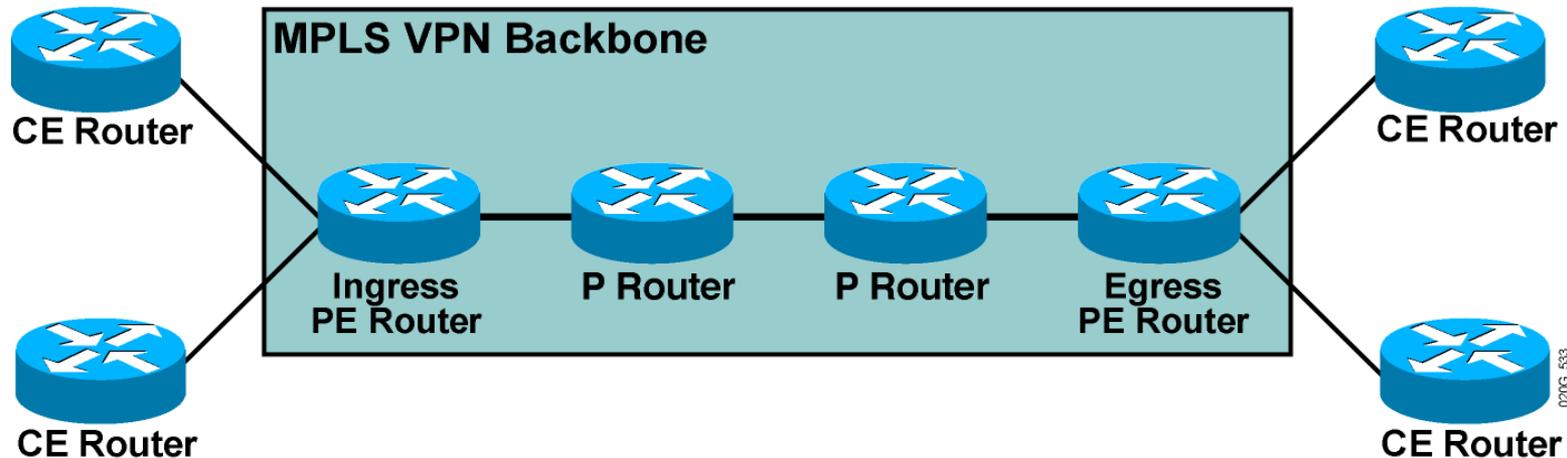
Question: How will the PE routers forward the VPN packets across the MPLS VPN backbone?

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?

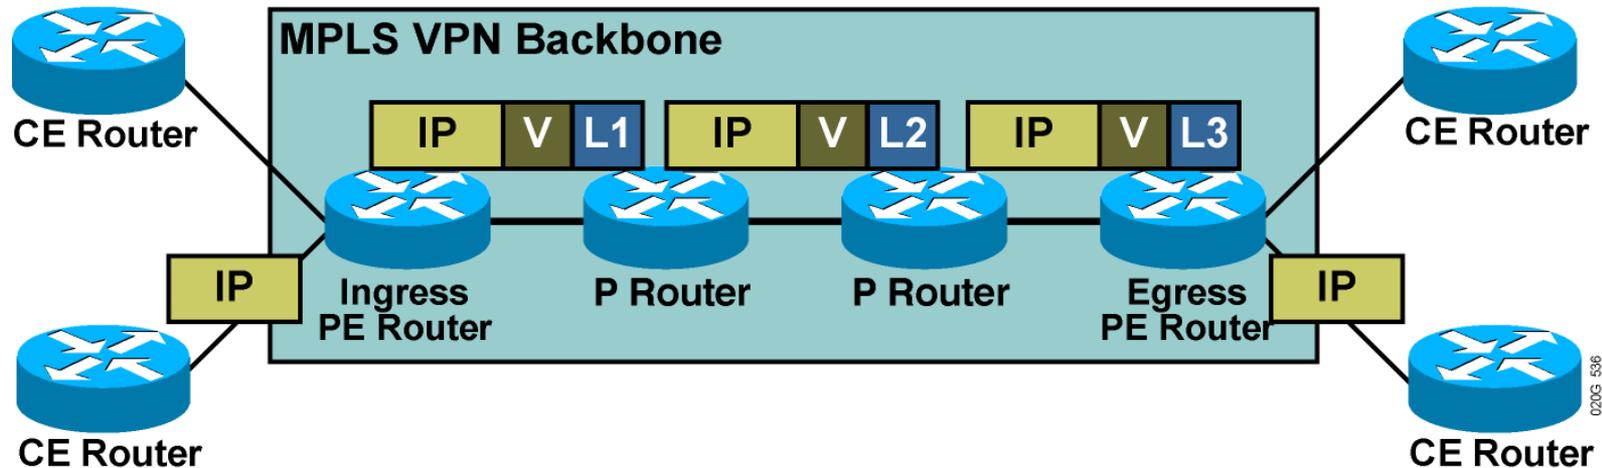
VPN Packet Forwarding Across an MPLS VPN Backbone (Cont.)



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.)



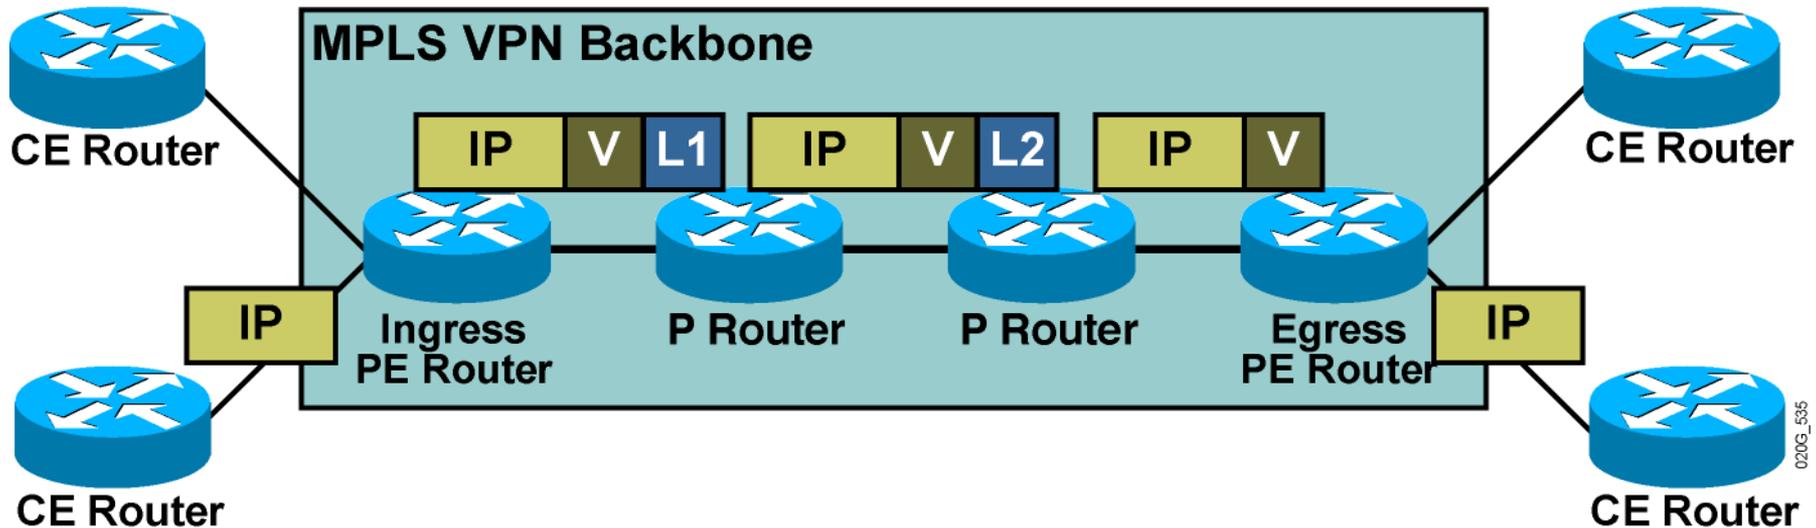
Question: How will the PE routers forward the VPN packets across the MPLS VPN backbone?

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:

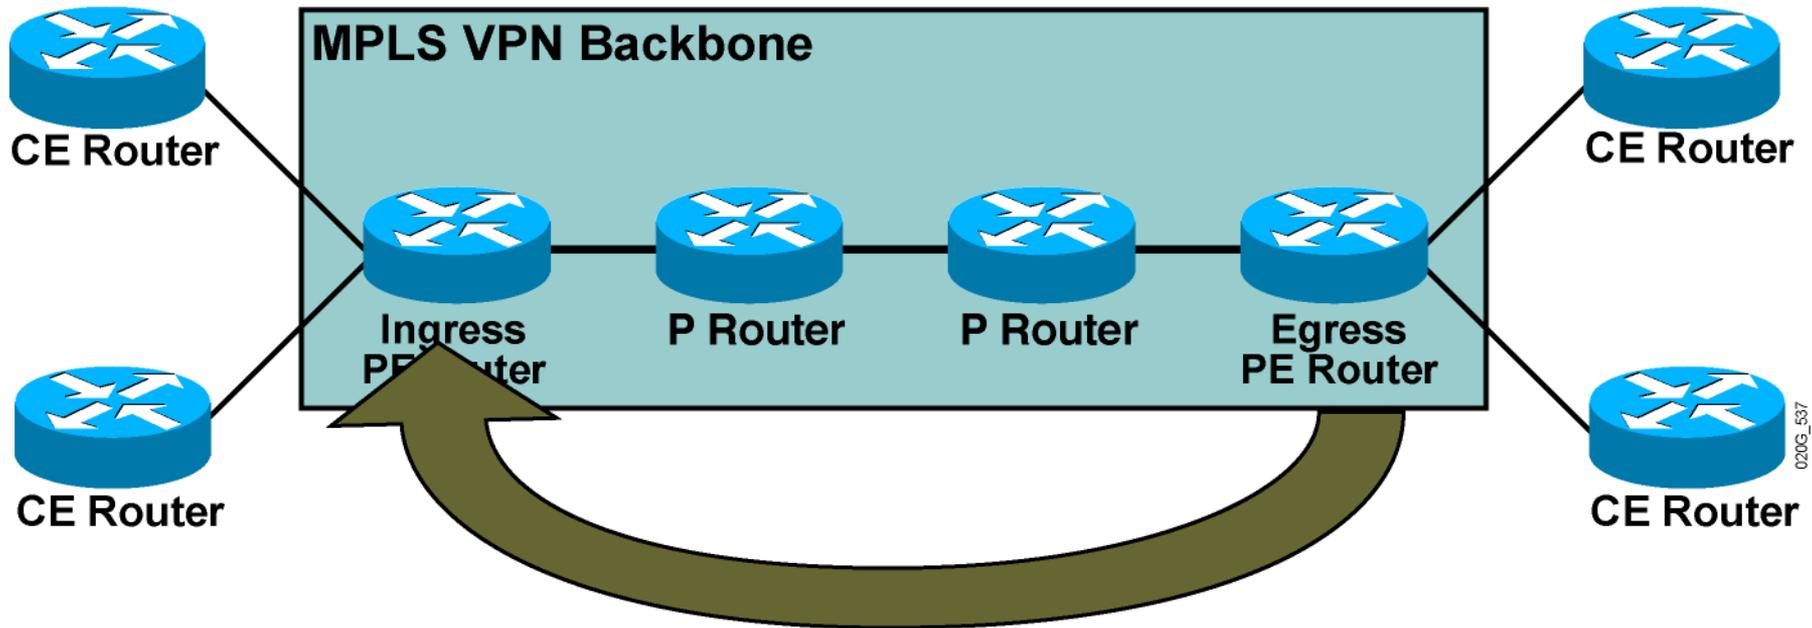
- The P routers perform label switching, and the packet reaches the egress PE router.
- The egress PE router performs a lookup on the VPN label and forwards the packet toward the CE router.

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.

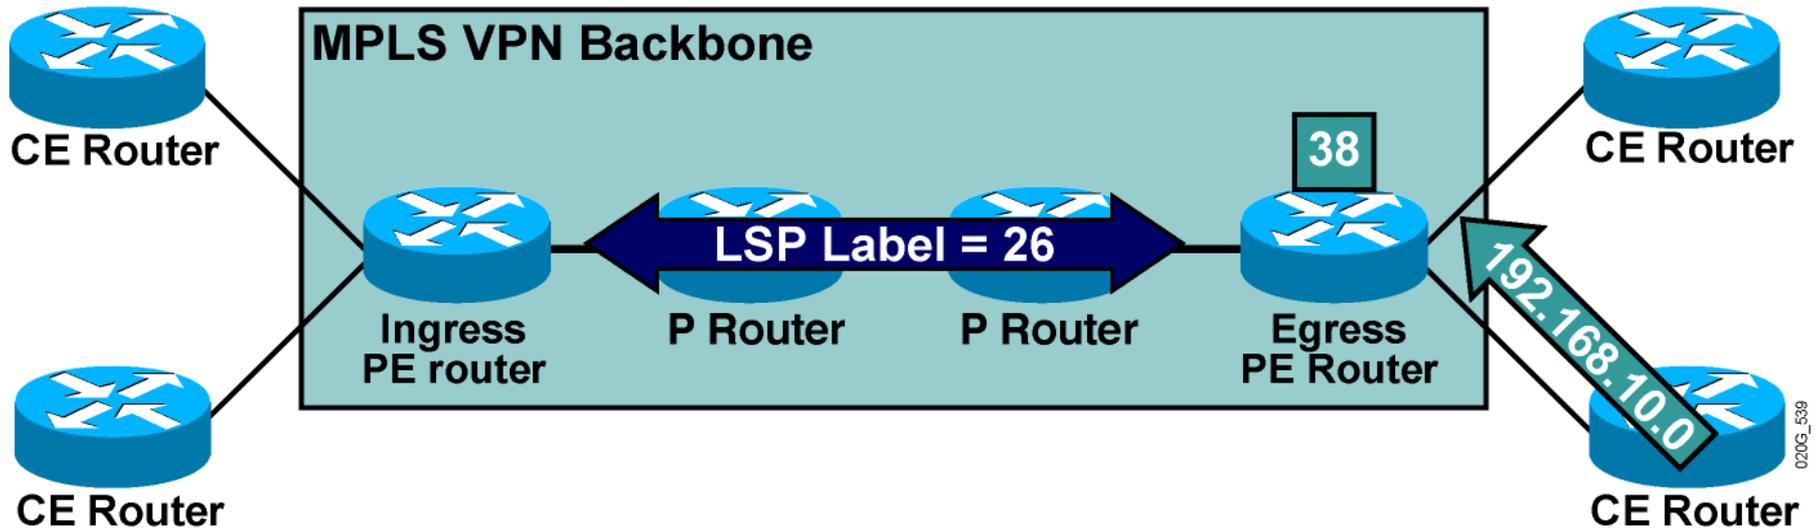
VPN Label Propagation



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

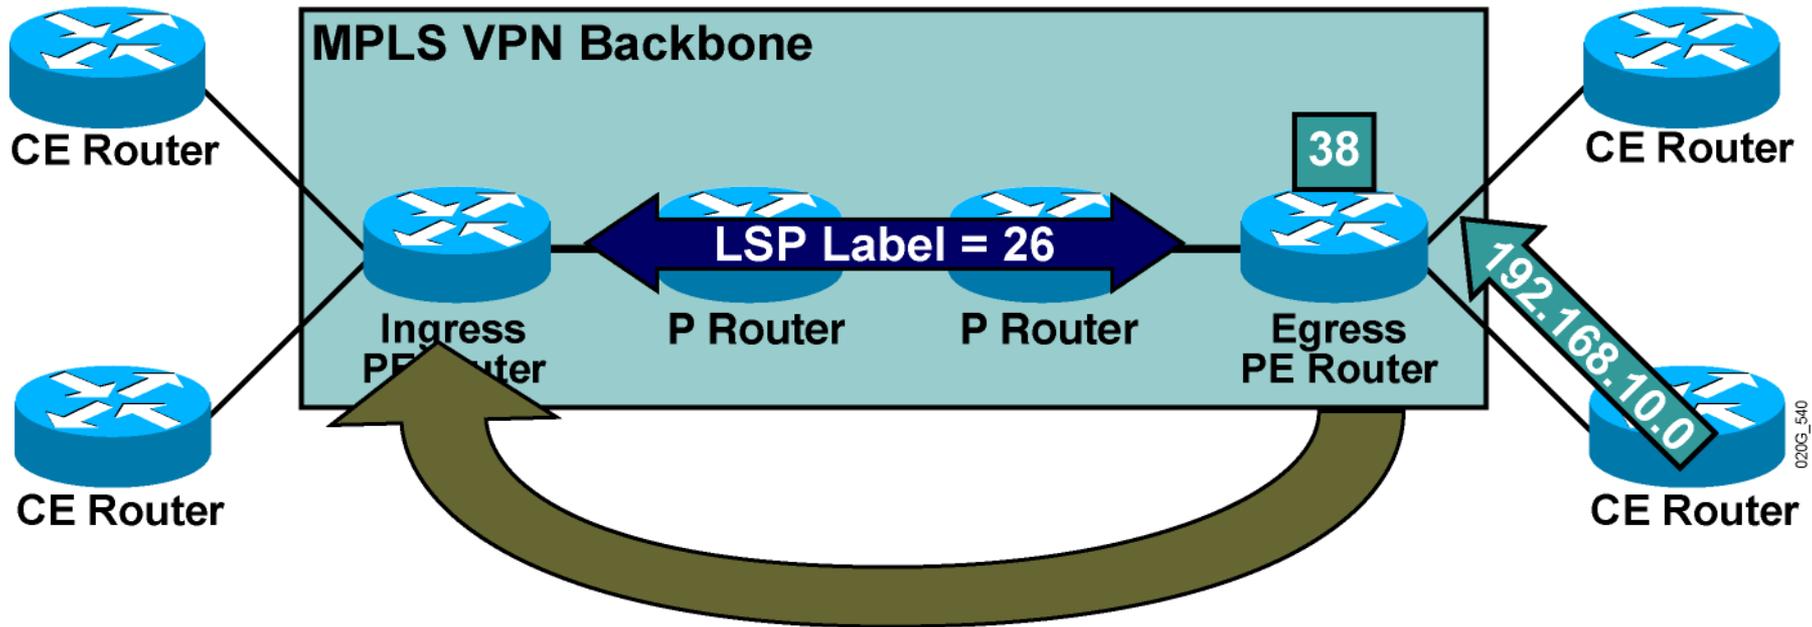
Answer: Labels are propagated in MP-BGP VPNv4 routing updates.

VPN Label Propagation (Cont.)



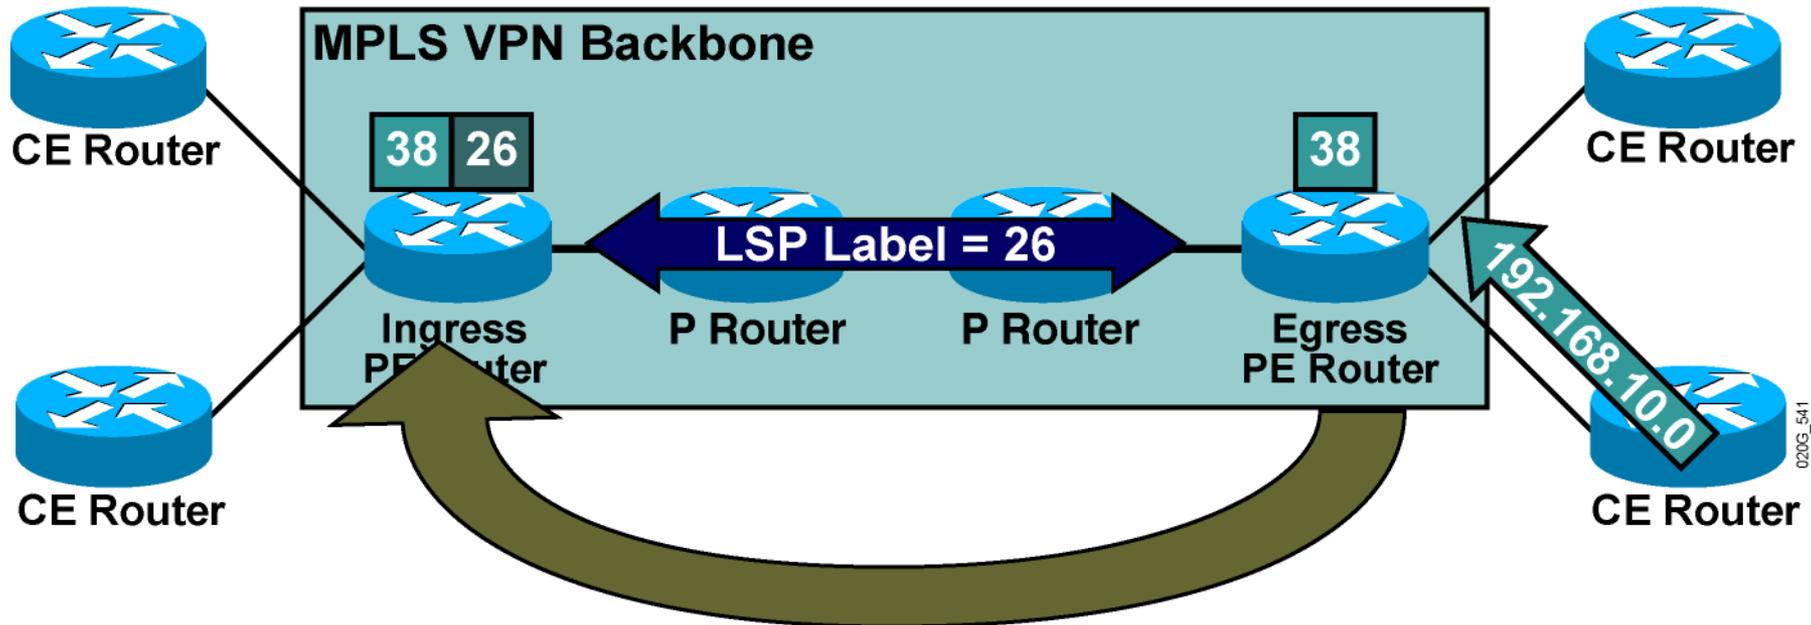
- o 1: A VPN label is assigned to every VPN route by the egress PE router.

VPN Label Propagation (Cont.)



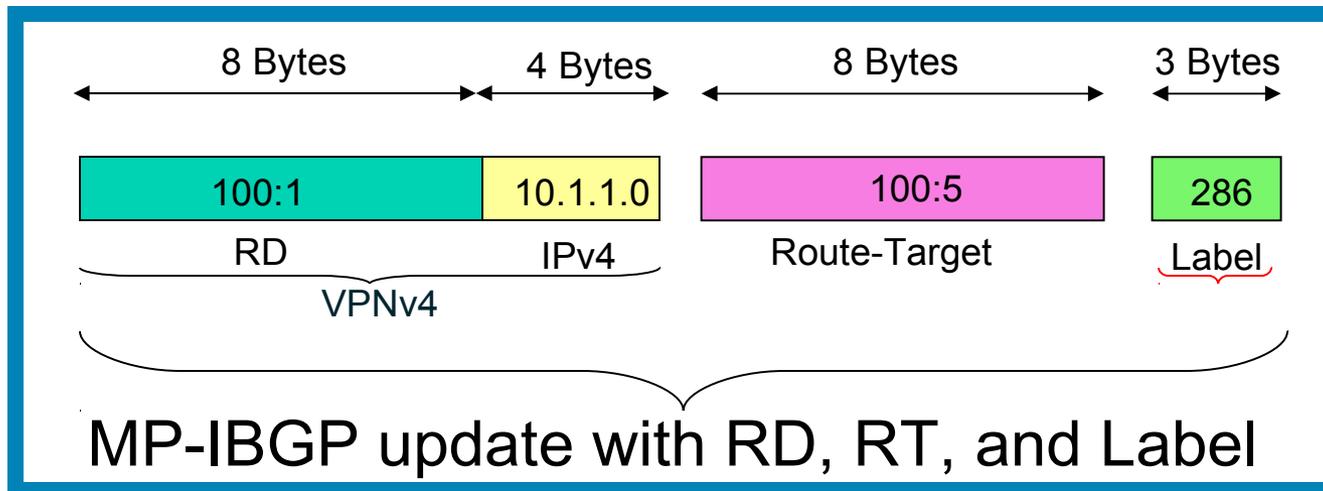
- o 1: A VPN label is assigned to every VPN route by the egress PE router.
- o 2: The VPN label is advertised to all other PE routers in an MP-BGP update.

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.

VPN Label in MP-iBGP update



MPLS VPNs and Label Propagation

The VPN label must be assigned by the BGP next hop.

The BGP next hop should not be changed in the MP-IBGP update propagation.

- Do not use next-hop-self on confederation boundaries.

The PE router must be the BGP next hop.

- Use next-hop-self on the PE router (default on current IOS)

The label must be reoriginated if the next hop is changed.

- A new label is assigned every time that the MP-BGP update crosses the AS boundary where the next hop is changed.

MPLS VPNs and Packet Forwarding

The VPN label is understood only by the egress PE router.

An end-to-end LSP tunnel is required between the ingress and egress PE routers.

BGP next hops must not be announced as BGP routes.

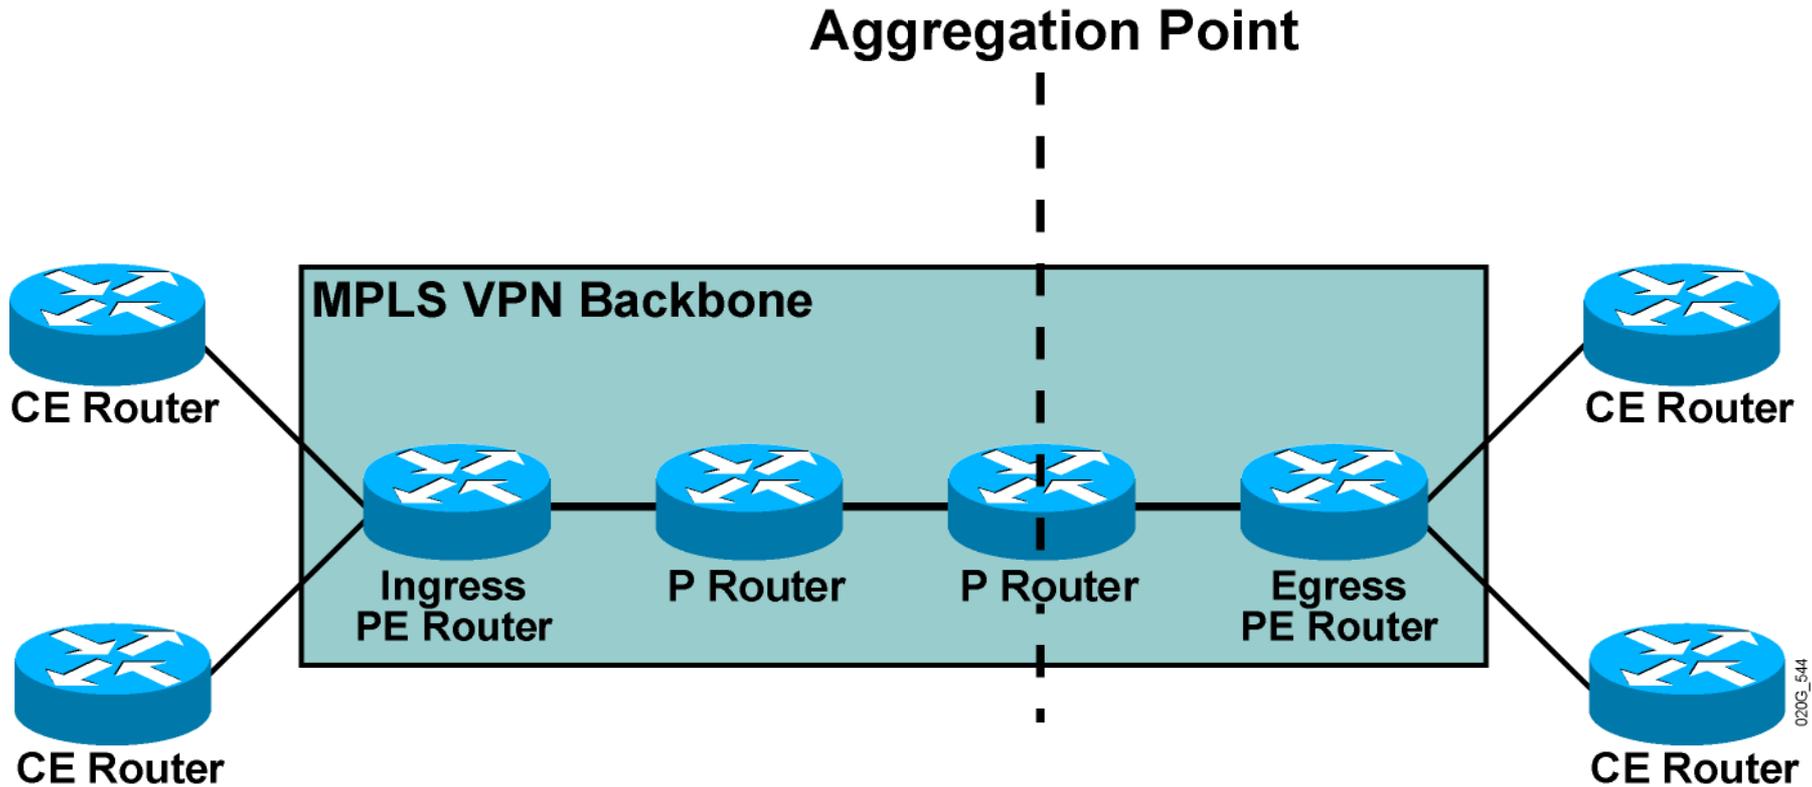
LDP labels are not assigned to BGP routes.

BGP next hops announced in IGP must not be summarized in the core network.

Summarization breaks the LSP tunnel.

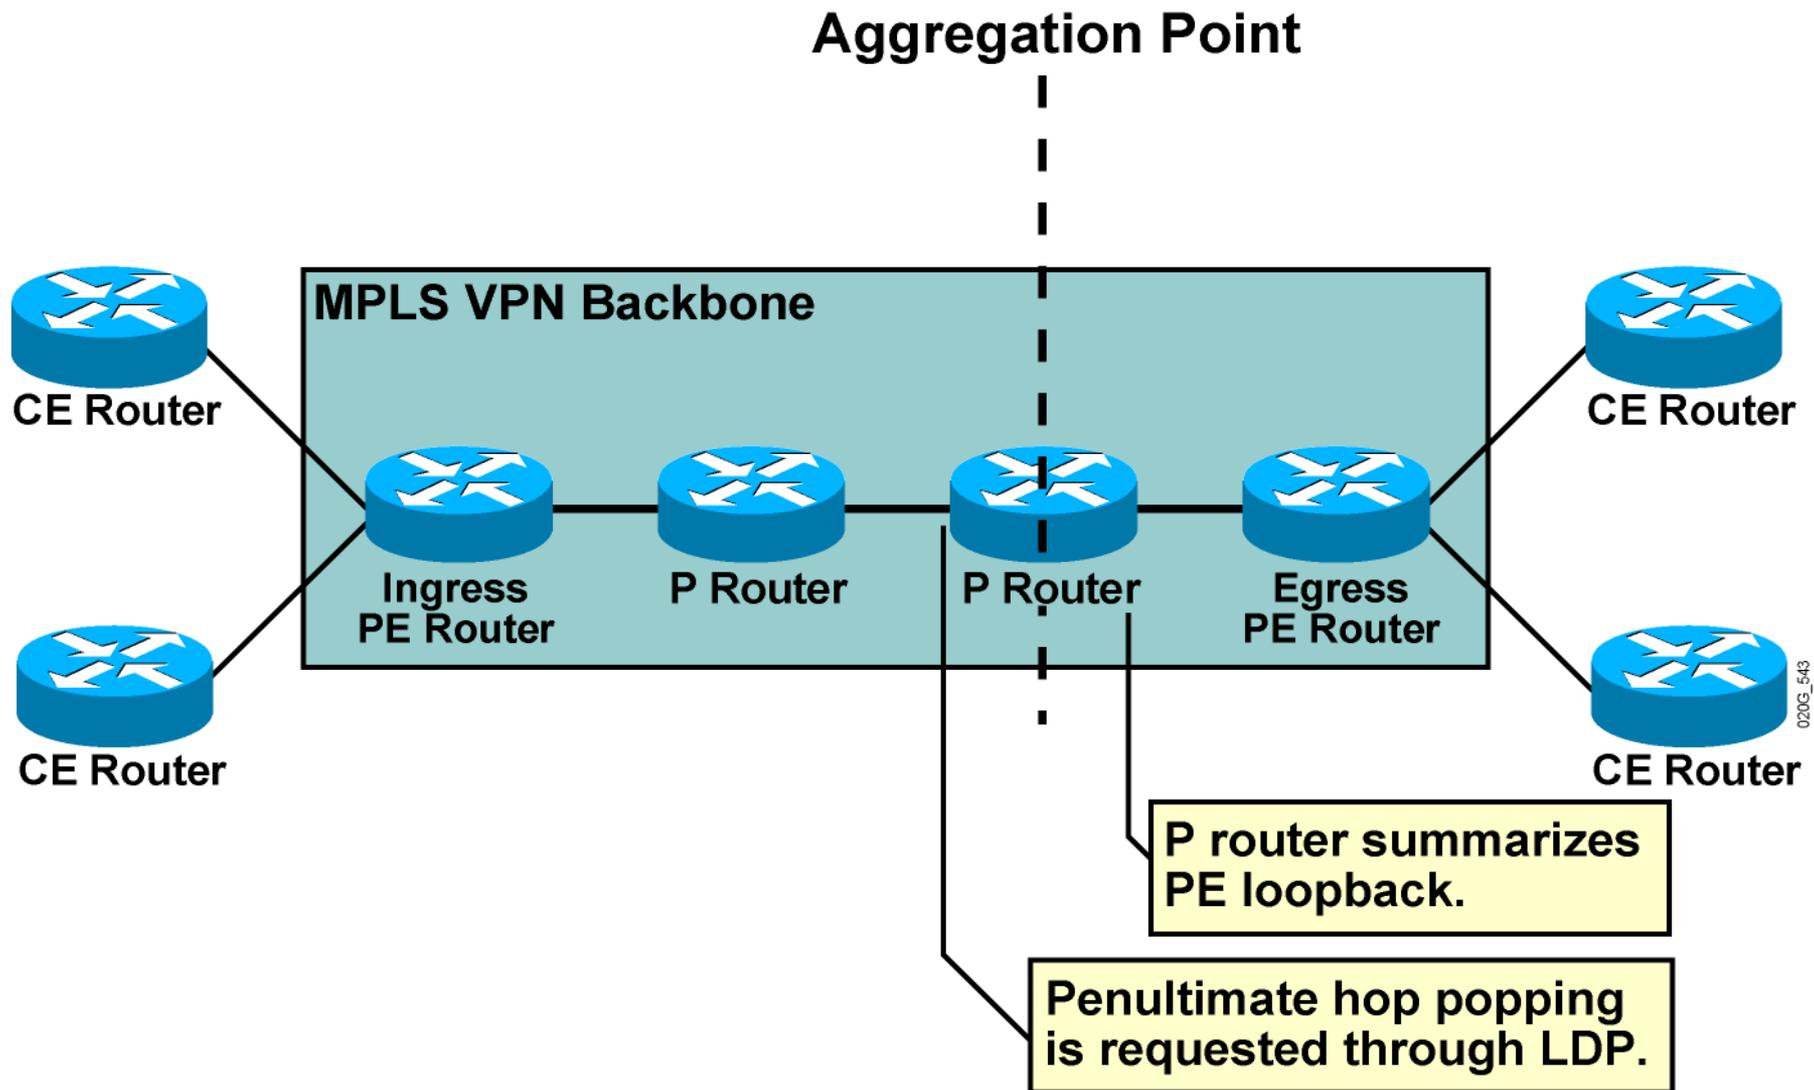
MPLS VPNs and Packet Forwarding (Cont.)

Summarization in the Core



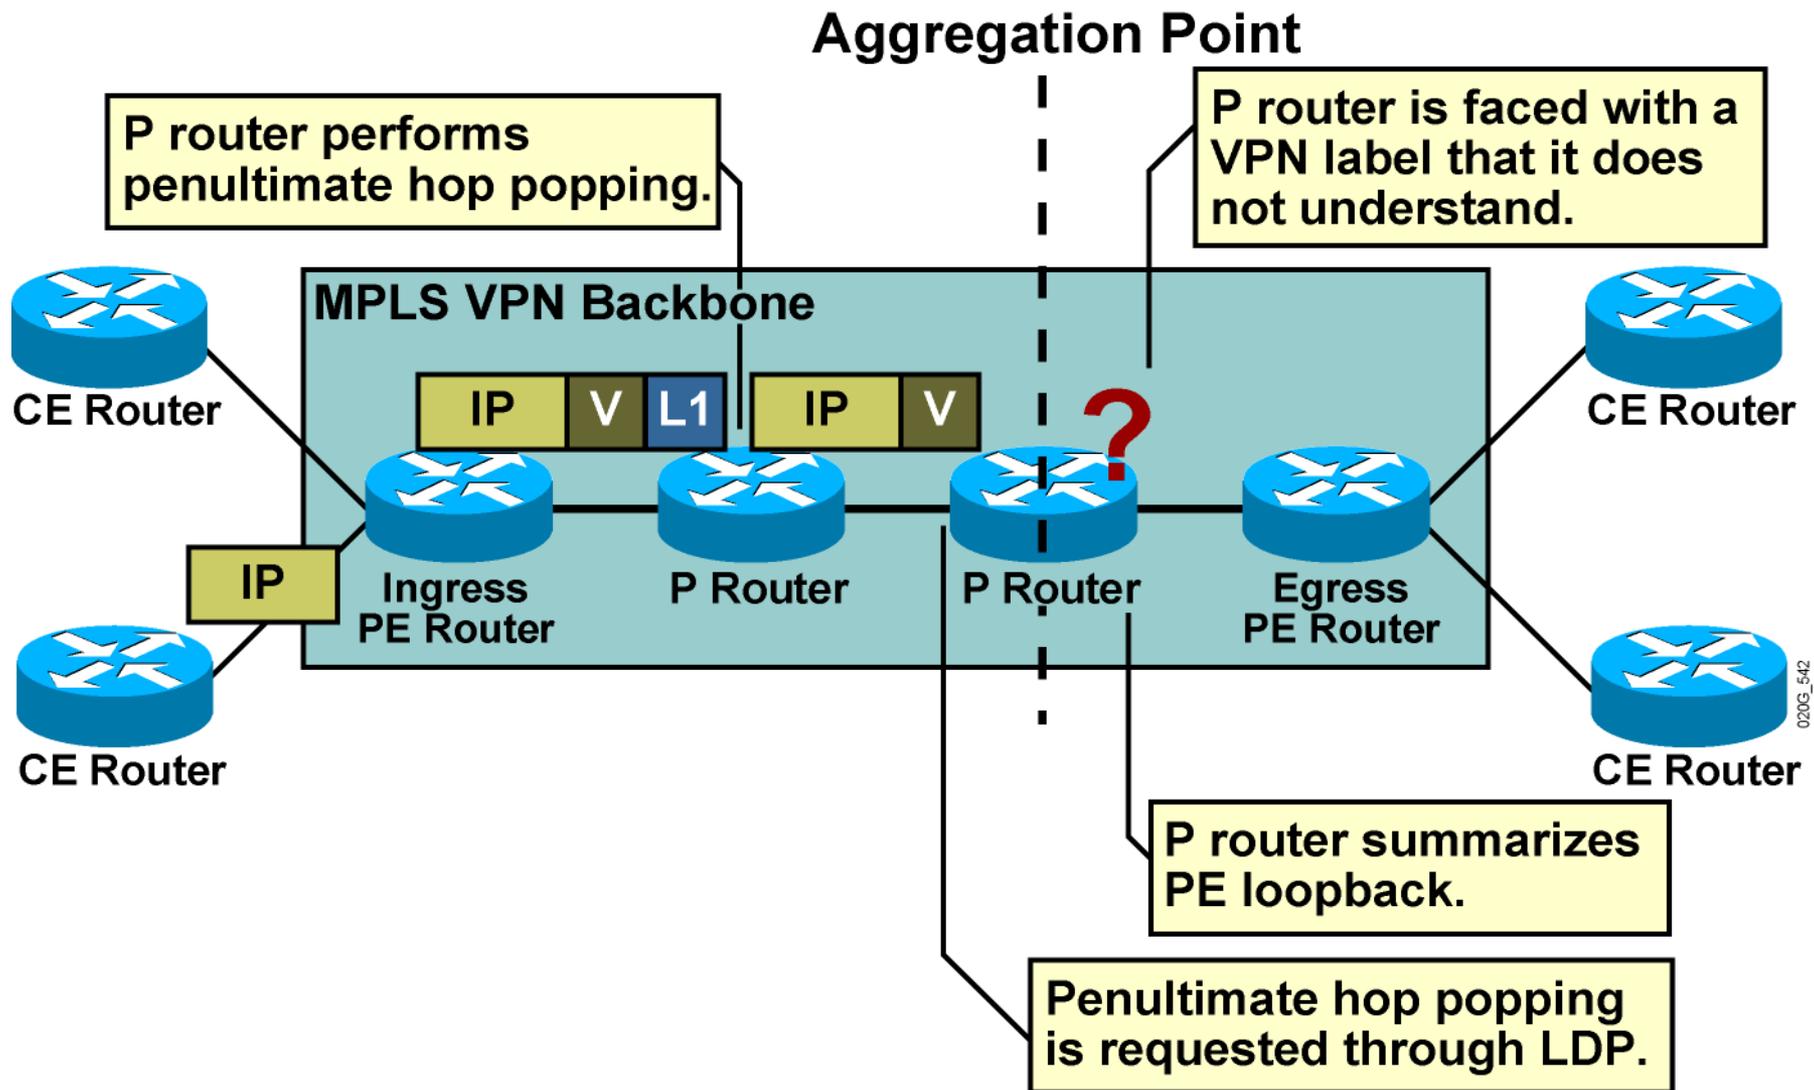
MPLS VPNs and Packet Forwarding (Cont.)

Summarization in the Core



MPLS VPNs and Packet Forwarding (Cont.)

Summarization in the Core



Summary

PE routers forward packets across the MPLS VPN backbone using label stacking.

Labels are propagated between PE routers using MP-BGP.

BGP next hops should not be announced as BGP routes.

LDP labels are not assigned to BGP routes.



MPLS workshop

MPLS VPN Mechanisms of Cisco IOS Platforms

Outline

Overview

Virtual Routing and Forwarding Table

Need for Routing Protocol Contexts

VPN-Aware Routing Protocols

VRF Table

BGP Route propagation - Outbound

Non-BGP Route propagation - Outbound

Route propagation – Inbound

Lesson Summary

Virtual Routing and Forwarding Table

A VRF is the routing and forwarding instance for a set of sites with identical connectivity requirements.

Data structures associated with a VRF are as follows:

- IP routing table

- CEF table

- Set of rules and routing protocol parameters
(routing protocol contexts)

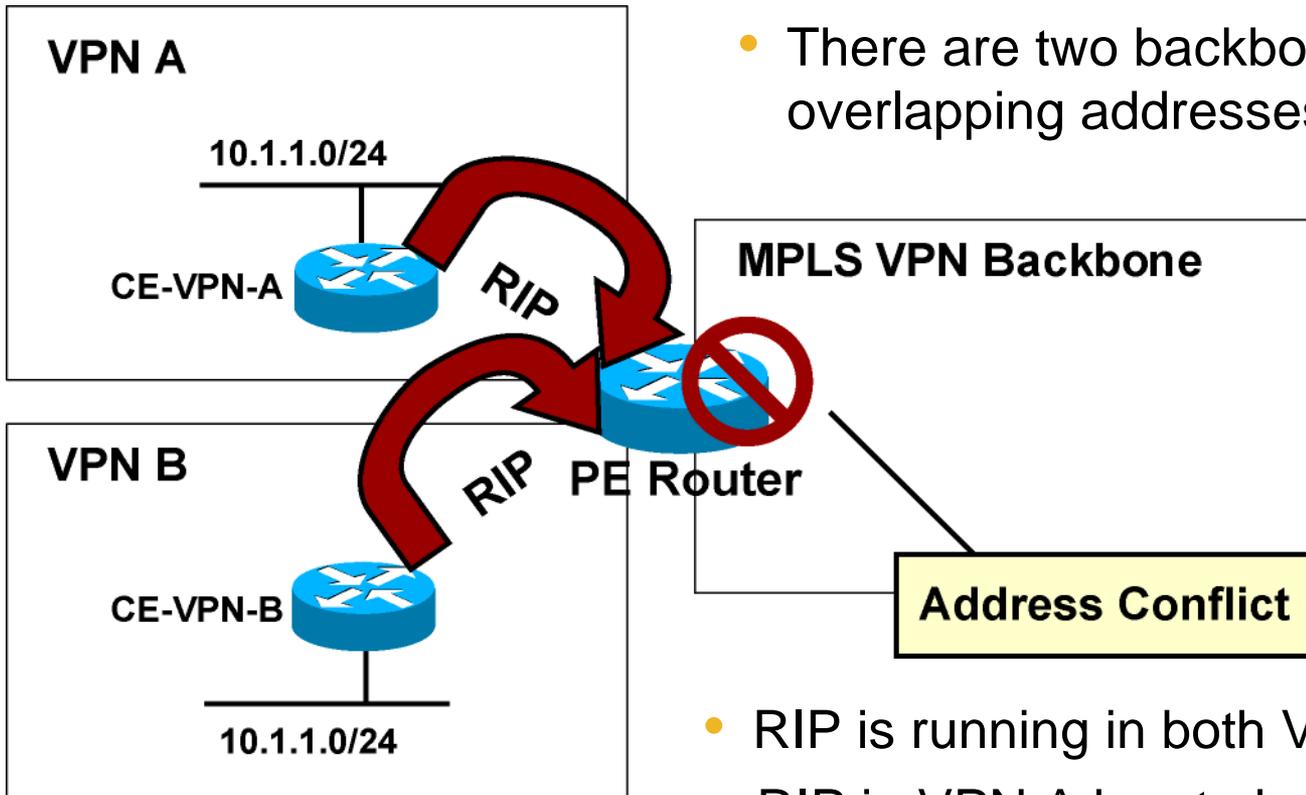
- List of interfaces that use the VRF

Other information associated with a VRF is as follows:

- Route distinguisher

- Set of import and export route targets

Need for Routing Protocol Contexts



- There are two backbones with overlapping addresses.

- RIP is running in both VPNs.
- RIP in VPN A has to be different from RIP in VPN B.
- Cisco IOS software supports only one RIP process per router.

VPN-Aware Routing Protocols

- Routing context = routing protocol run in one VRF:

Supported by VPN-aware routing protocols:

External BGP (EBGP), EIGRP, OSPF, RIP version 2 (RIPv2), static routes

Implemented as several instances of a single routing process (EBGP, RIPv2) or as several routing processes (OSPF)

Independent per-instance router variables for each instance

VRF Table

Contains routes that should be available to a particular set of sites

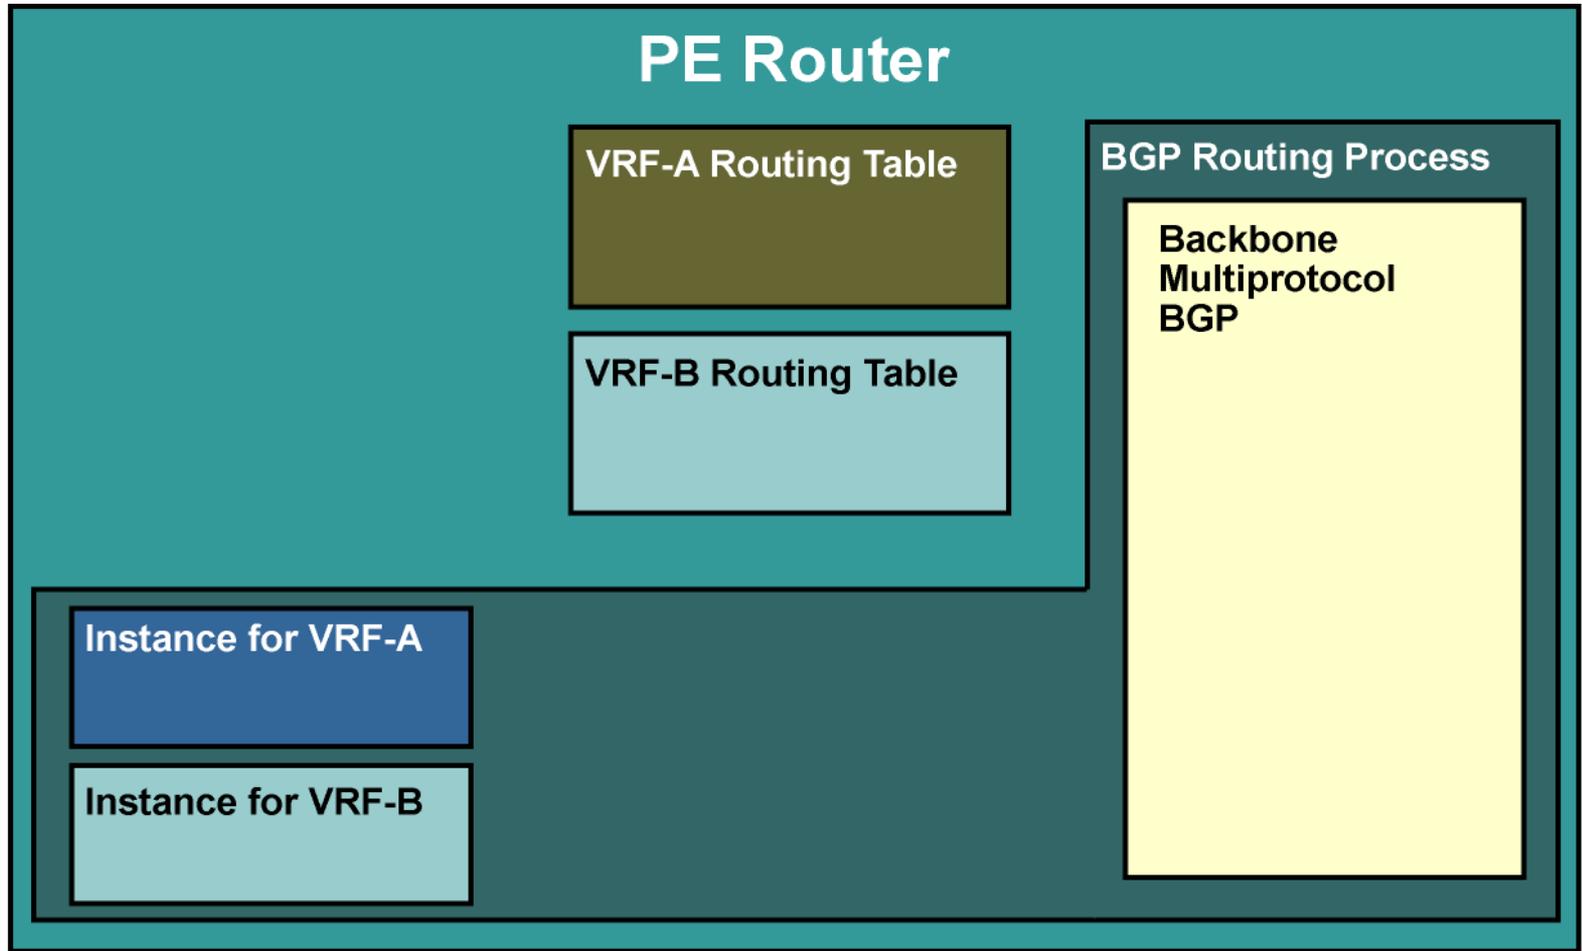
Analogous to standard Cisco IOS software routing table; supports same set of mechanisms

VPN interfaces (physical interface, subinterfaces, logical interfaces) assigned to VRFs:

- Many interfaces per VRF

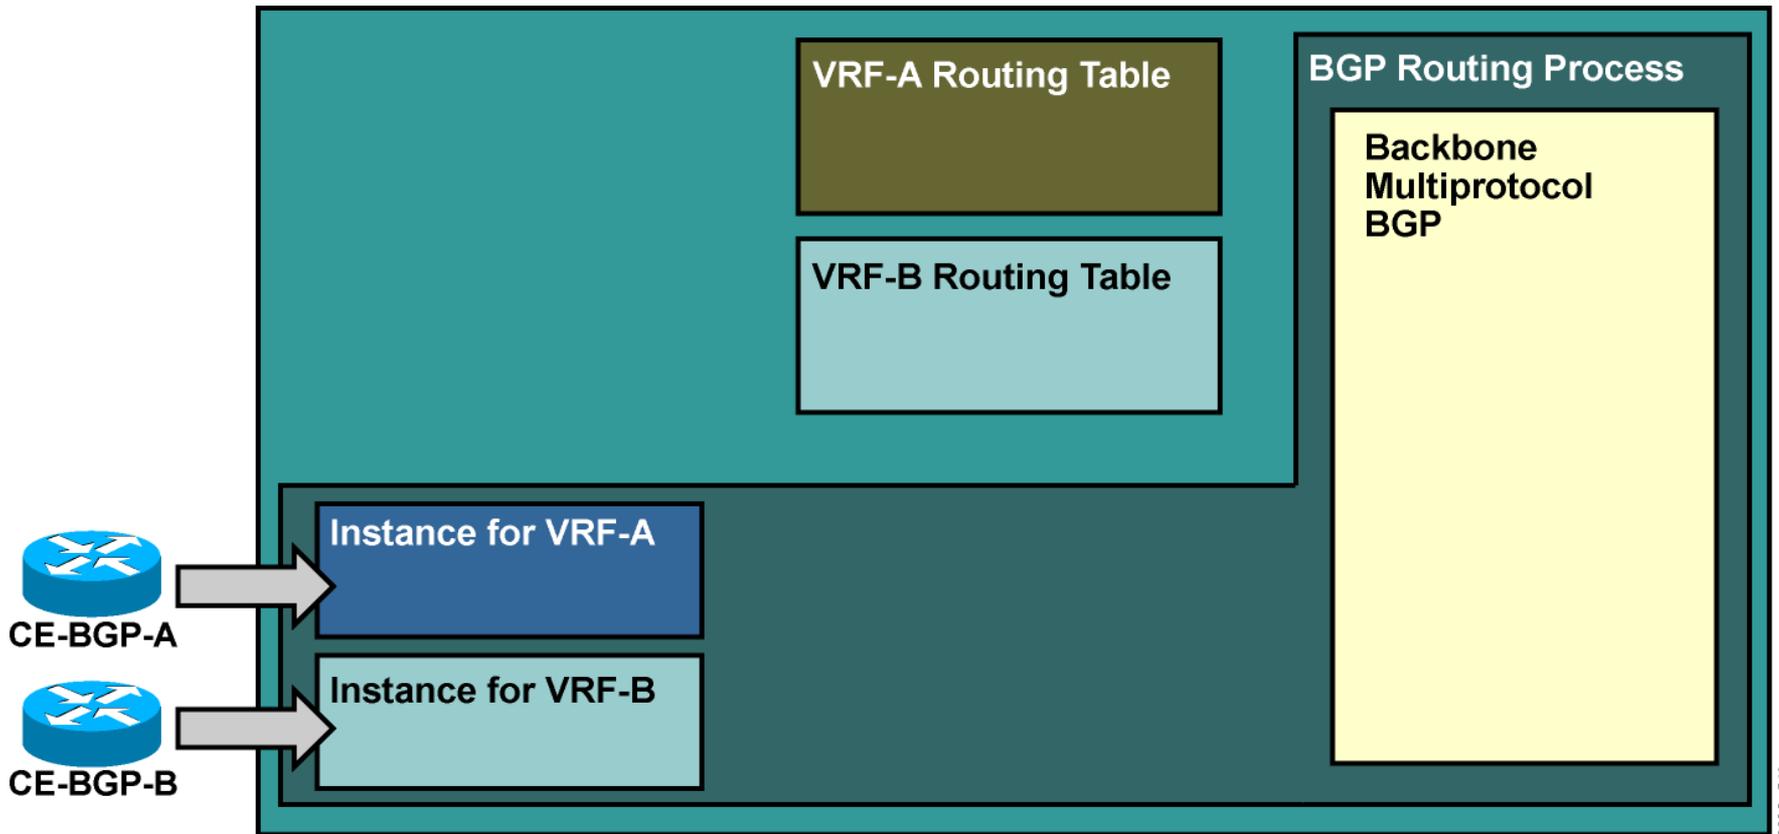
- Each interface assignable to only one VRF

BGP Route Propagation—Outbound



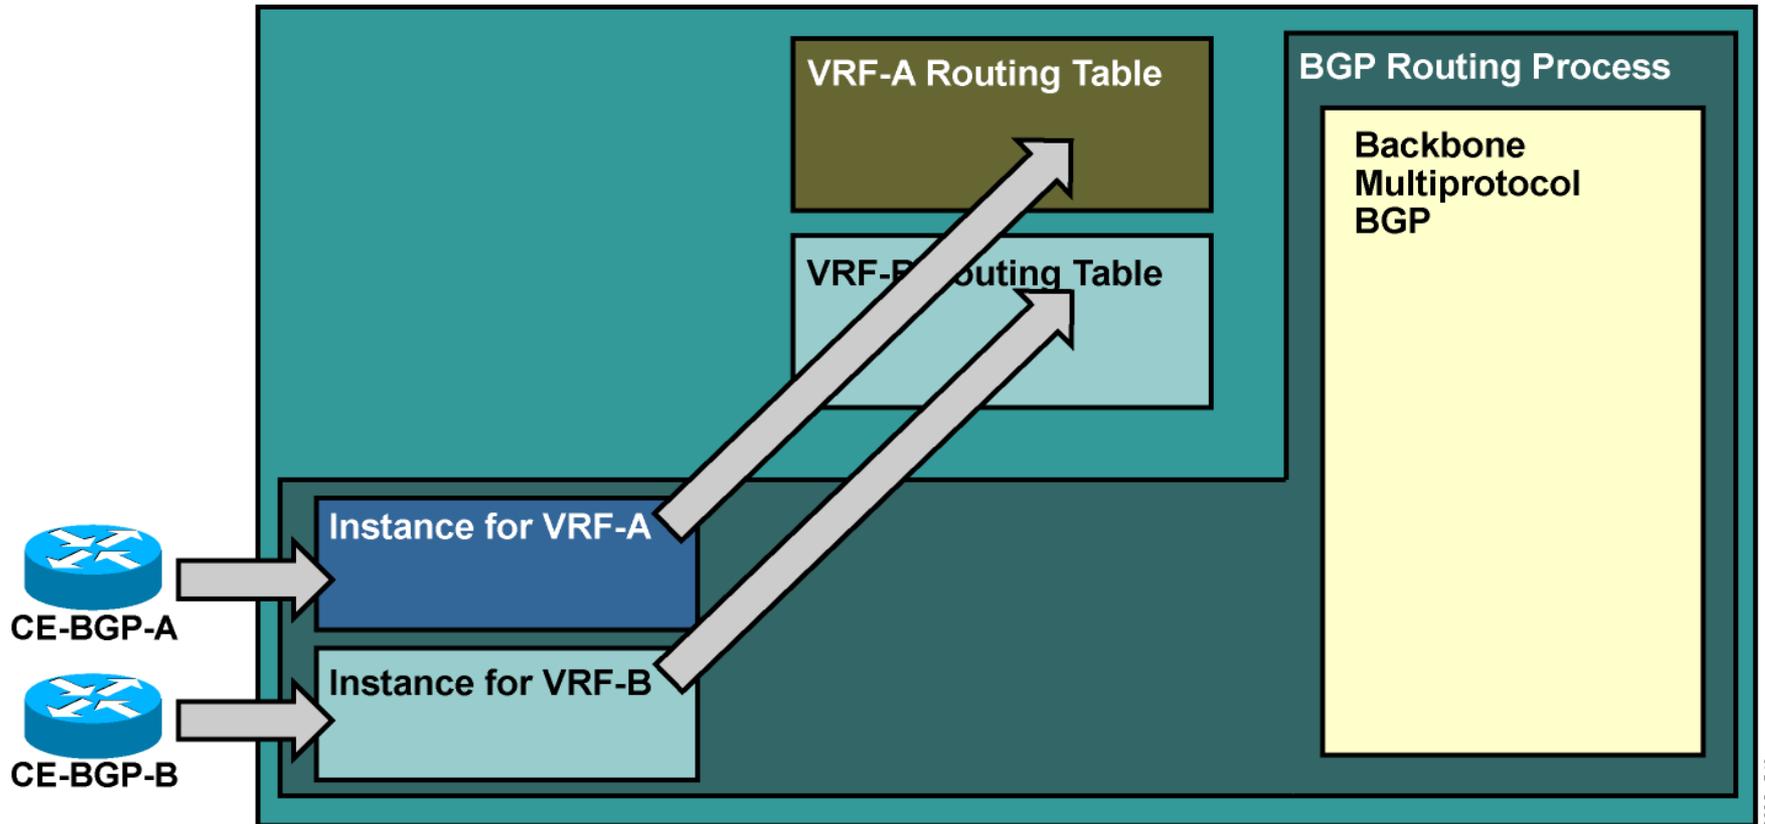
- Two VPNs are attached to the same PE router.
- Each VPN is represented by a VRF.

BGP Route Propagation—Outbound (Cont.)



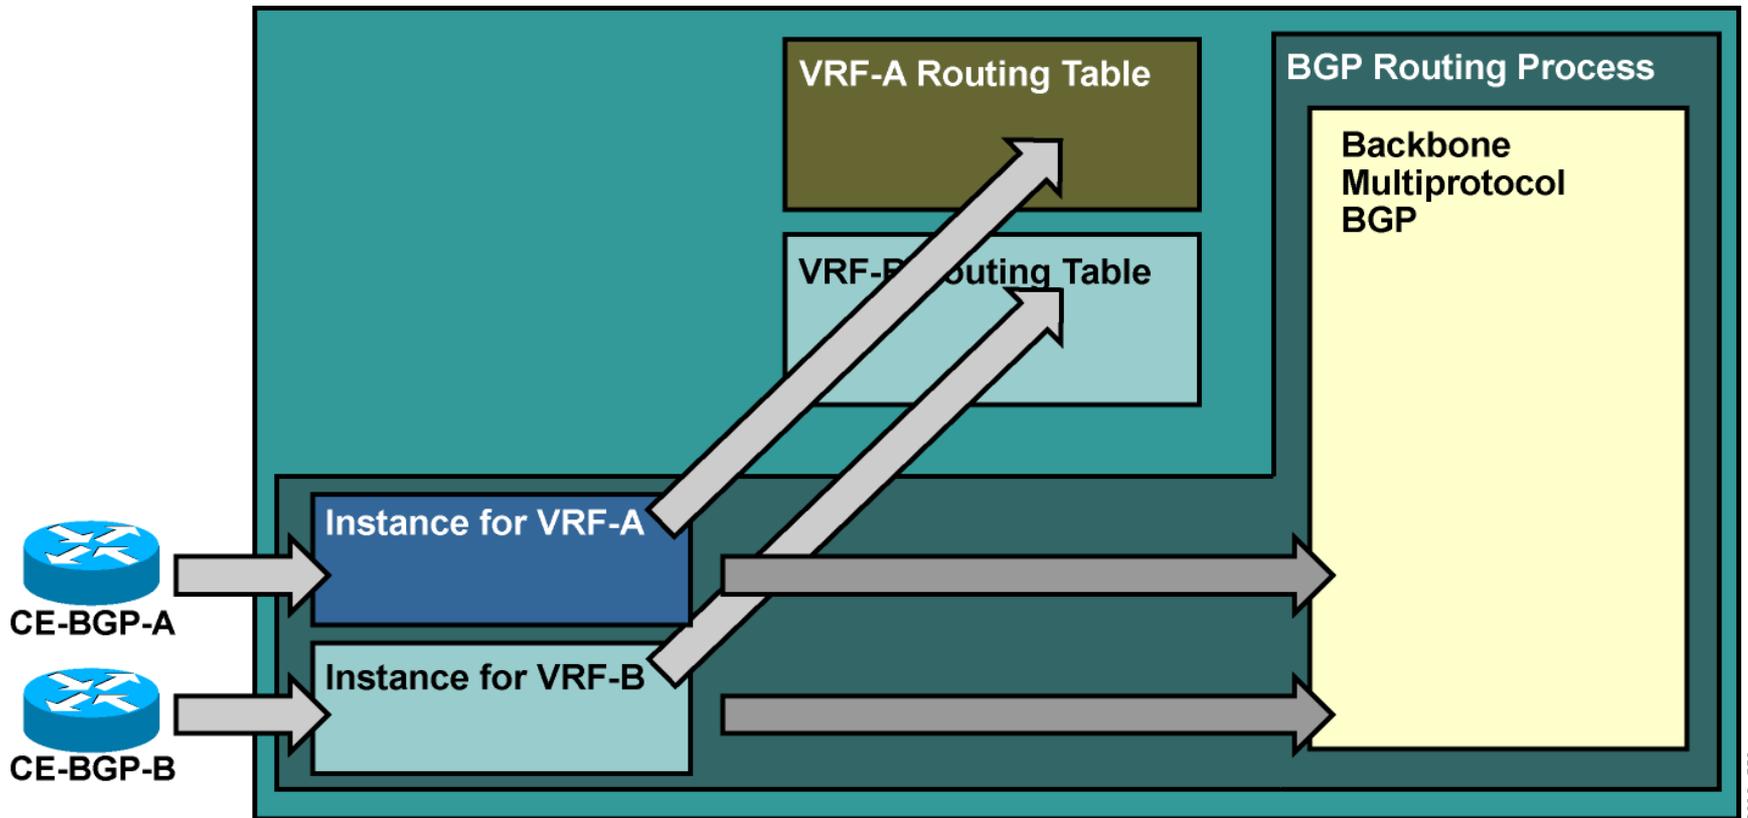
- BGP-speaking CE routers announce their prefixes to the PE router via BGP.

BGP Route Propagation—Outbound (Cont.)



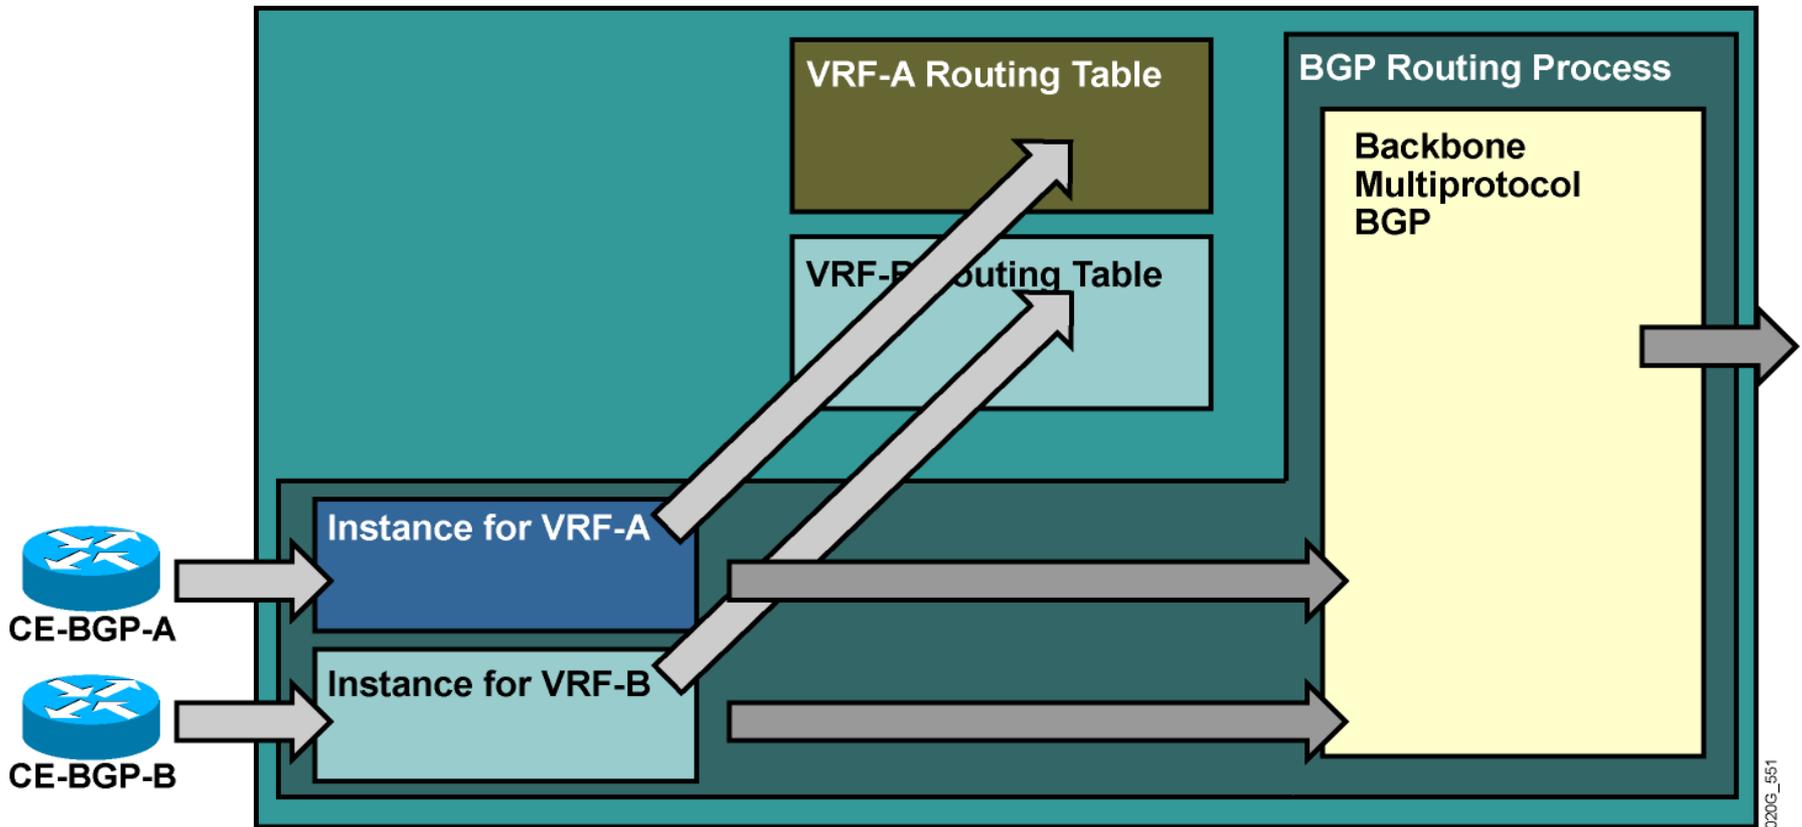
- BGP-speaking CE routers announce their prefixes to the PE router via BGP.
- Instance of BGP process associated with the VRF to which the PE-CE interface belongs collects the routes and inserts them into VRF routing table.

BGP Route Propagation—Outbound (Cont.)



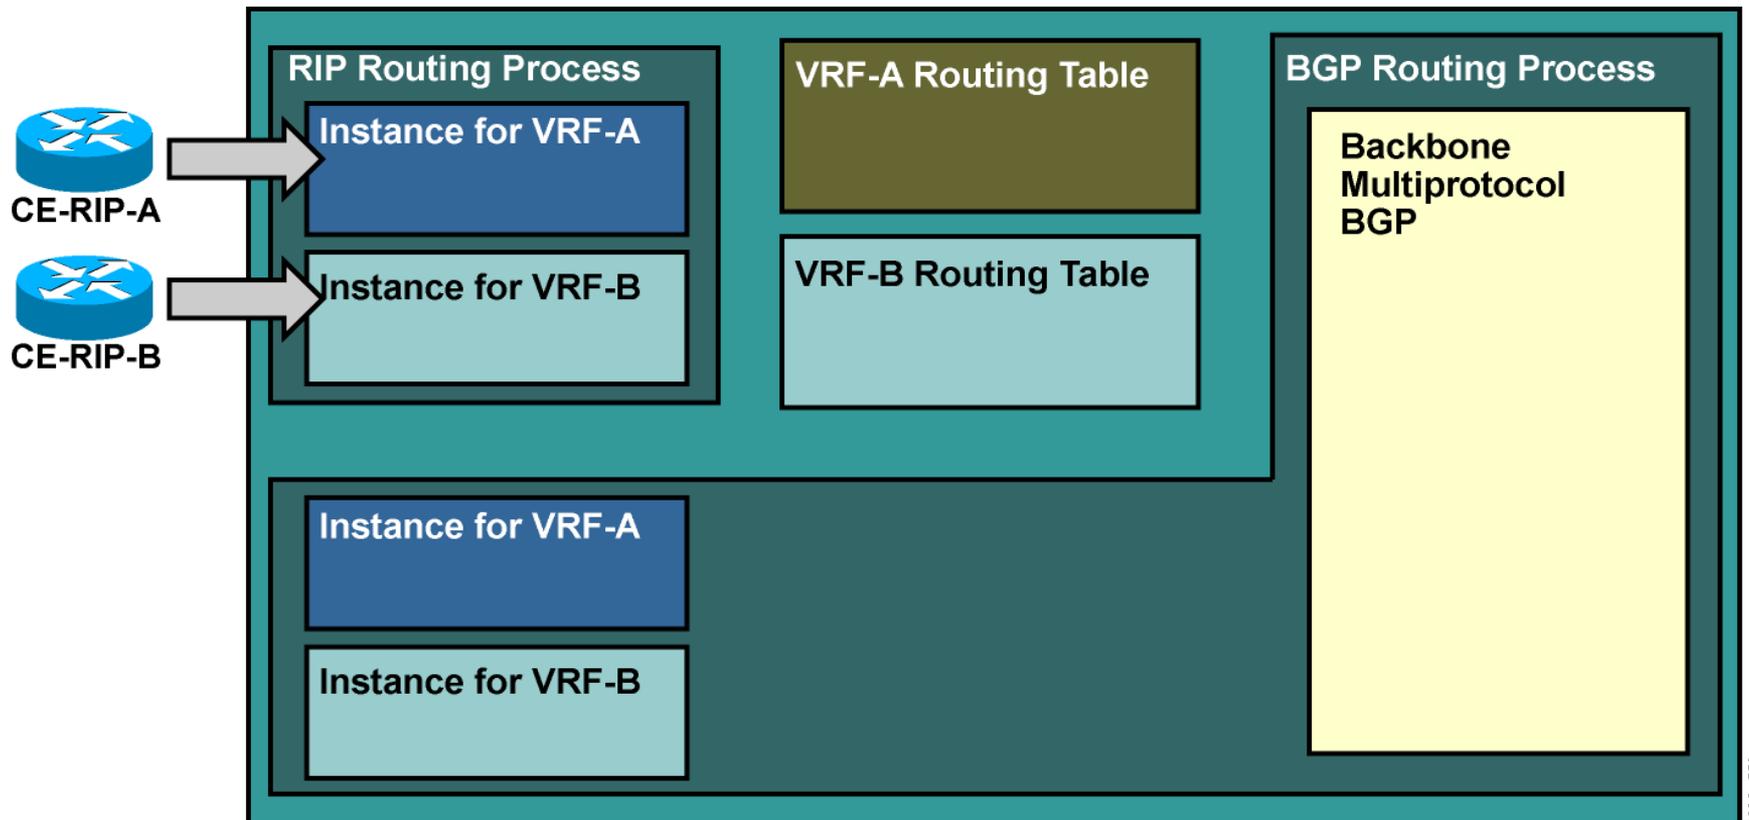
- Route distinguisher is prepended during route export to the BGP routes from VRF instance of BGP process to convert them into VPNv4 prefixes. Route targets are attached to these prefixes.

BGP Route Propagation—Outbound (Cont.)



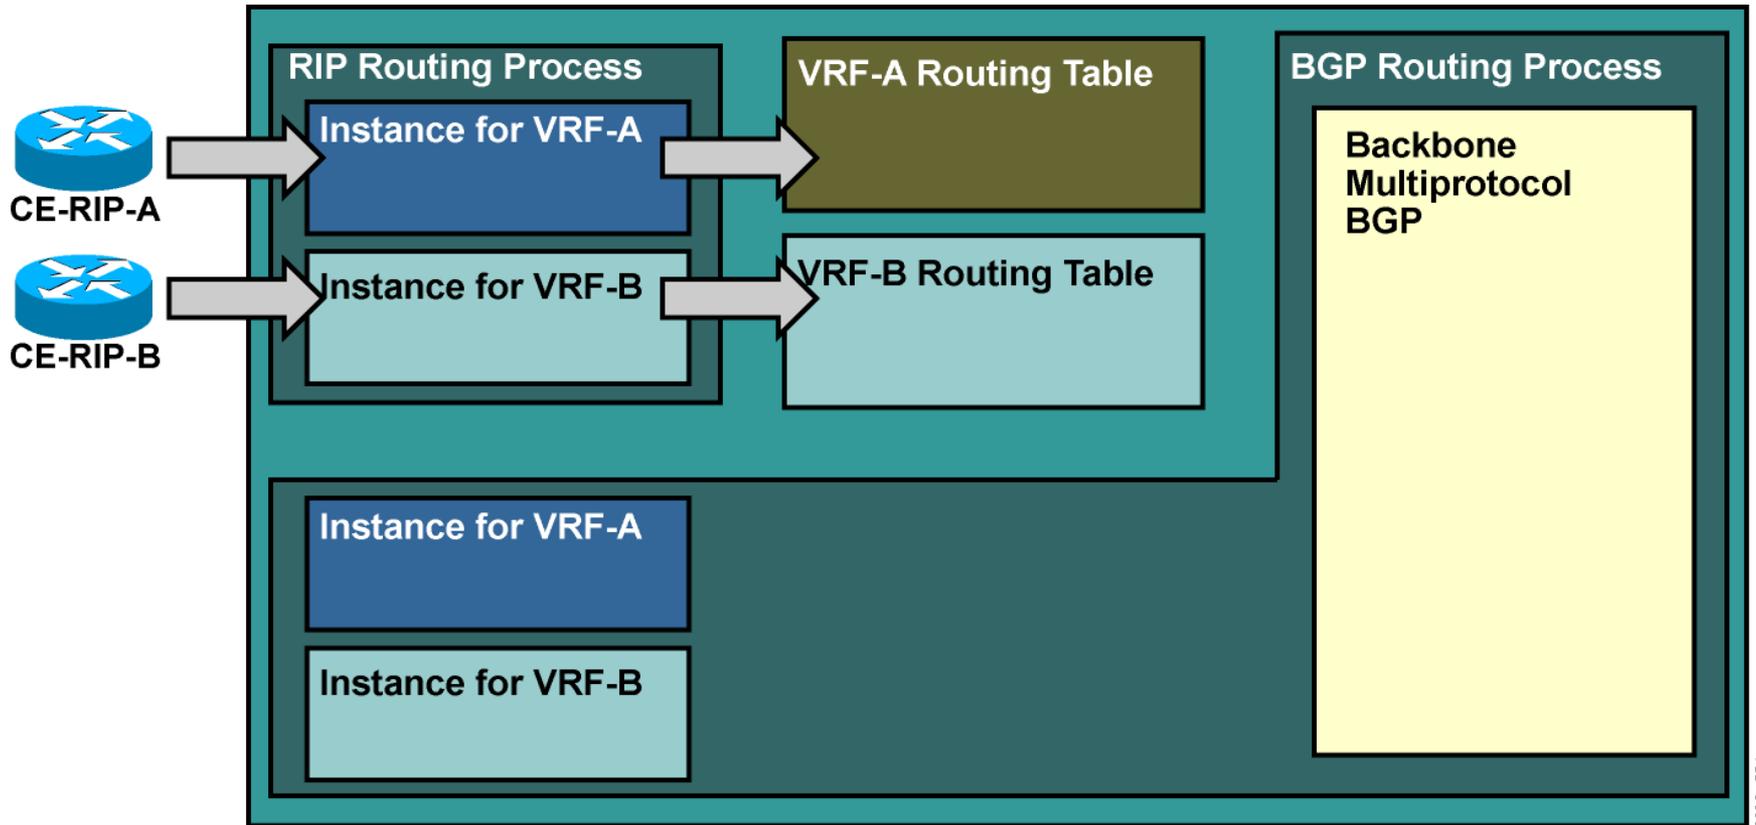
- Route distinguisher is prepended during route export to the BGP routes from VRF instance of BGP process to convert them into VPNv4 prefixes. Route targets are attached to these prefixes.
- VPNv4 prefixes are propagated to other PE routers.

Non-BGP Route propagation - Outbound



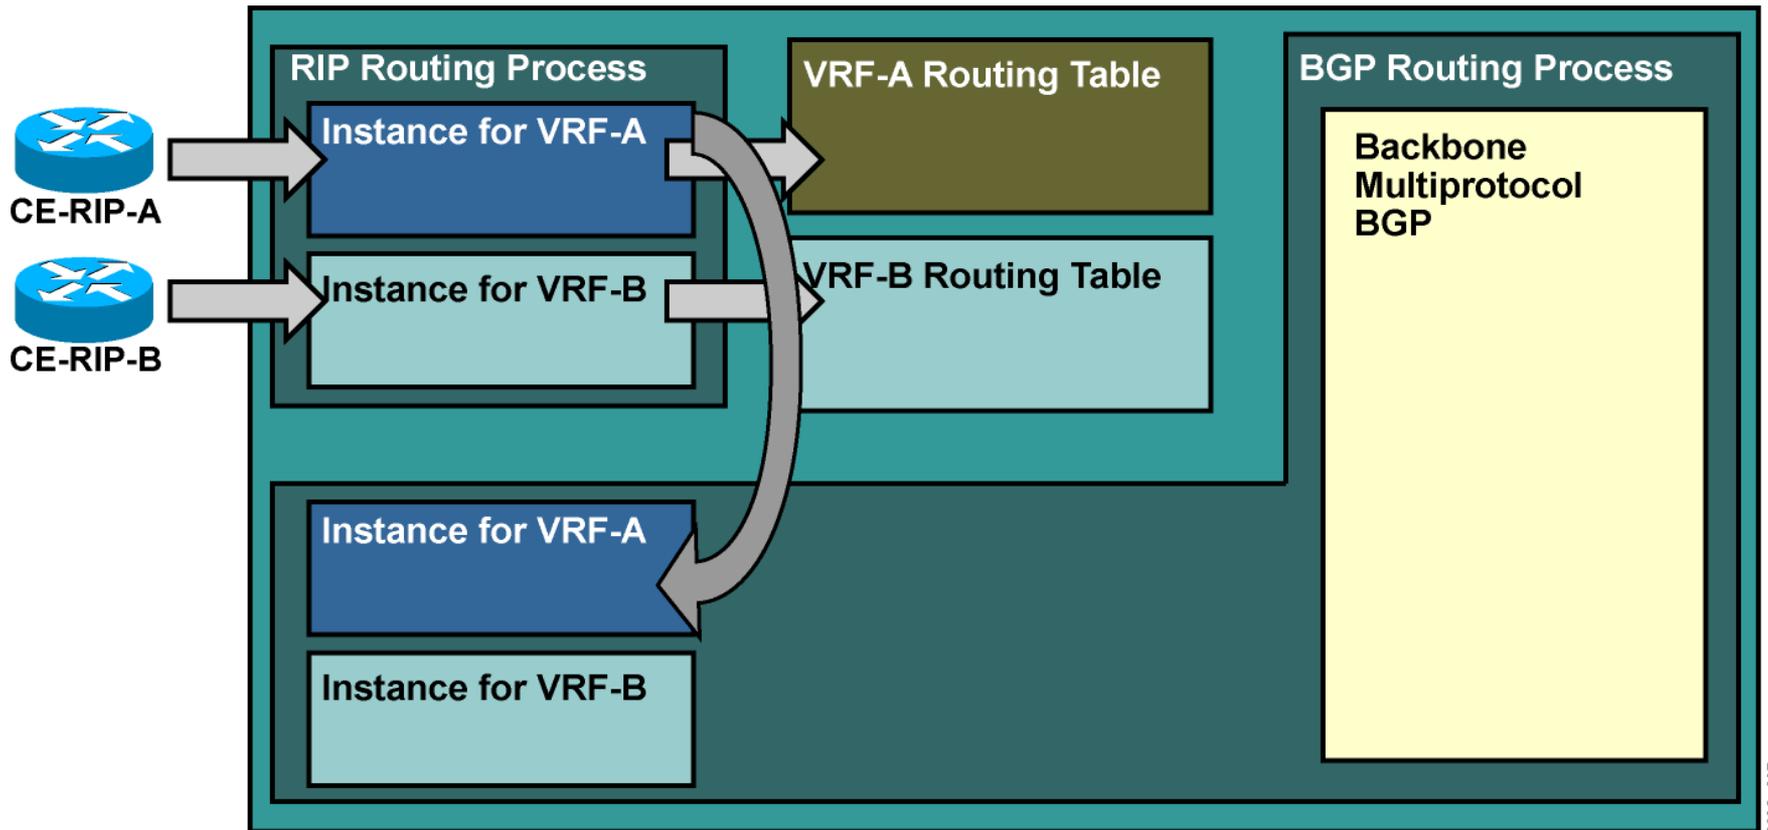
- RIP-speaking CE routers announce their prefixes to the PE router via RIP.

Non-BGP Route propagation—Outbound



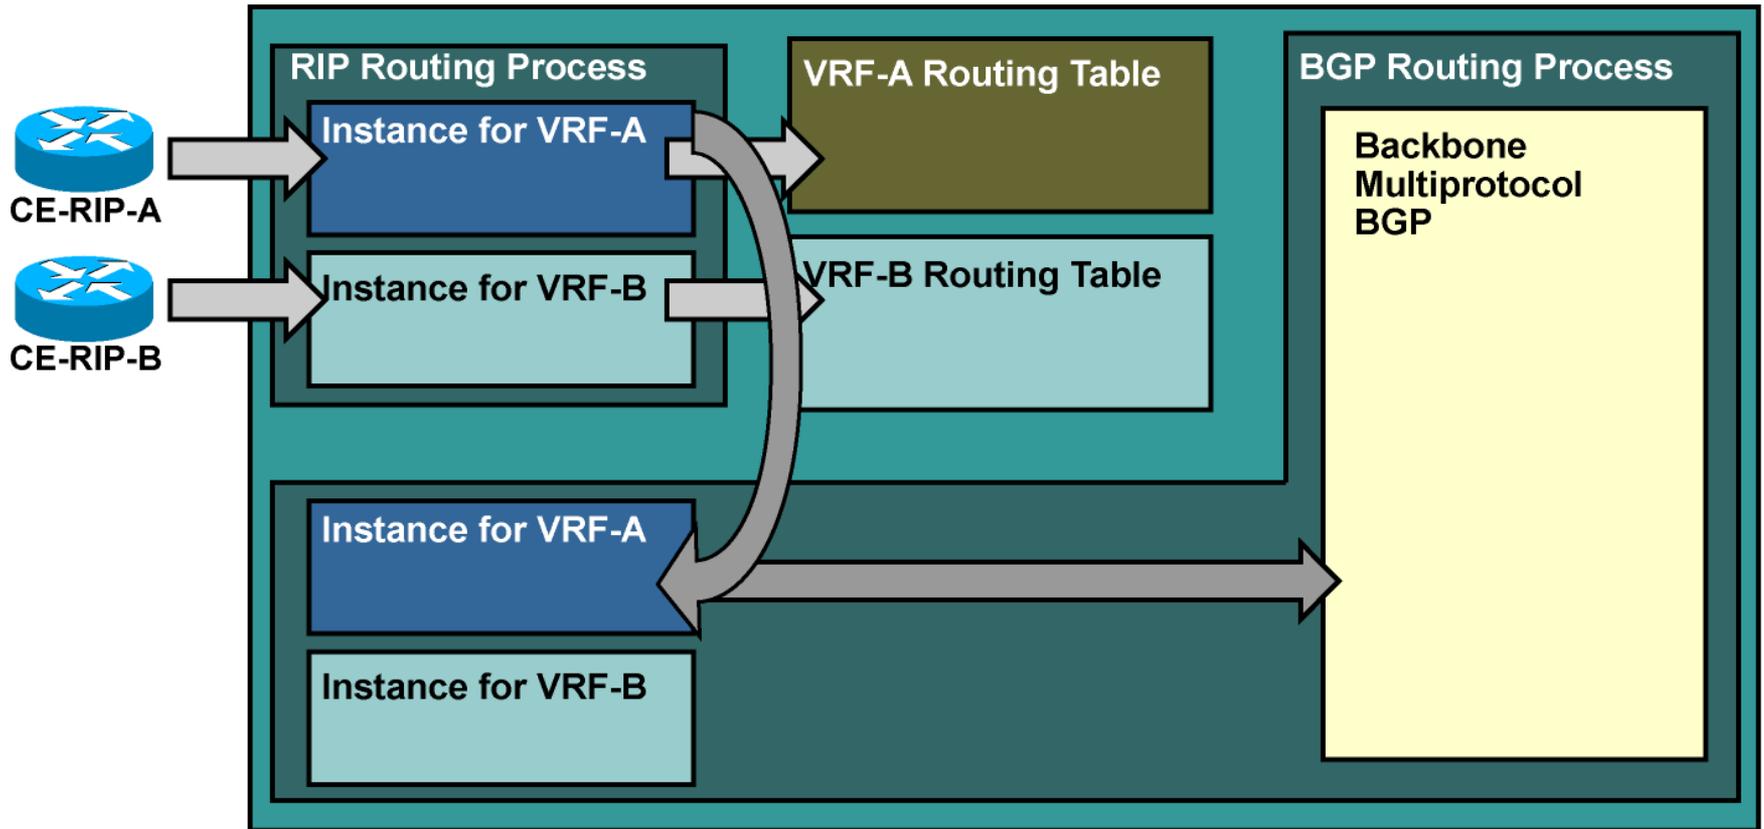
- RIP-speaking CE routers announce their prefixes to the PE router via RIP.
- Instance of RIP process associated with the VRF to which the PE-CE interface belongs collects the routes and inserts them into VRF routing table.

Non-BGP Route propagation—Outbound (Cont.)



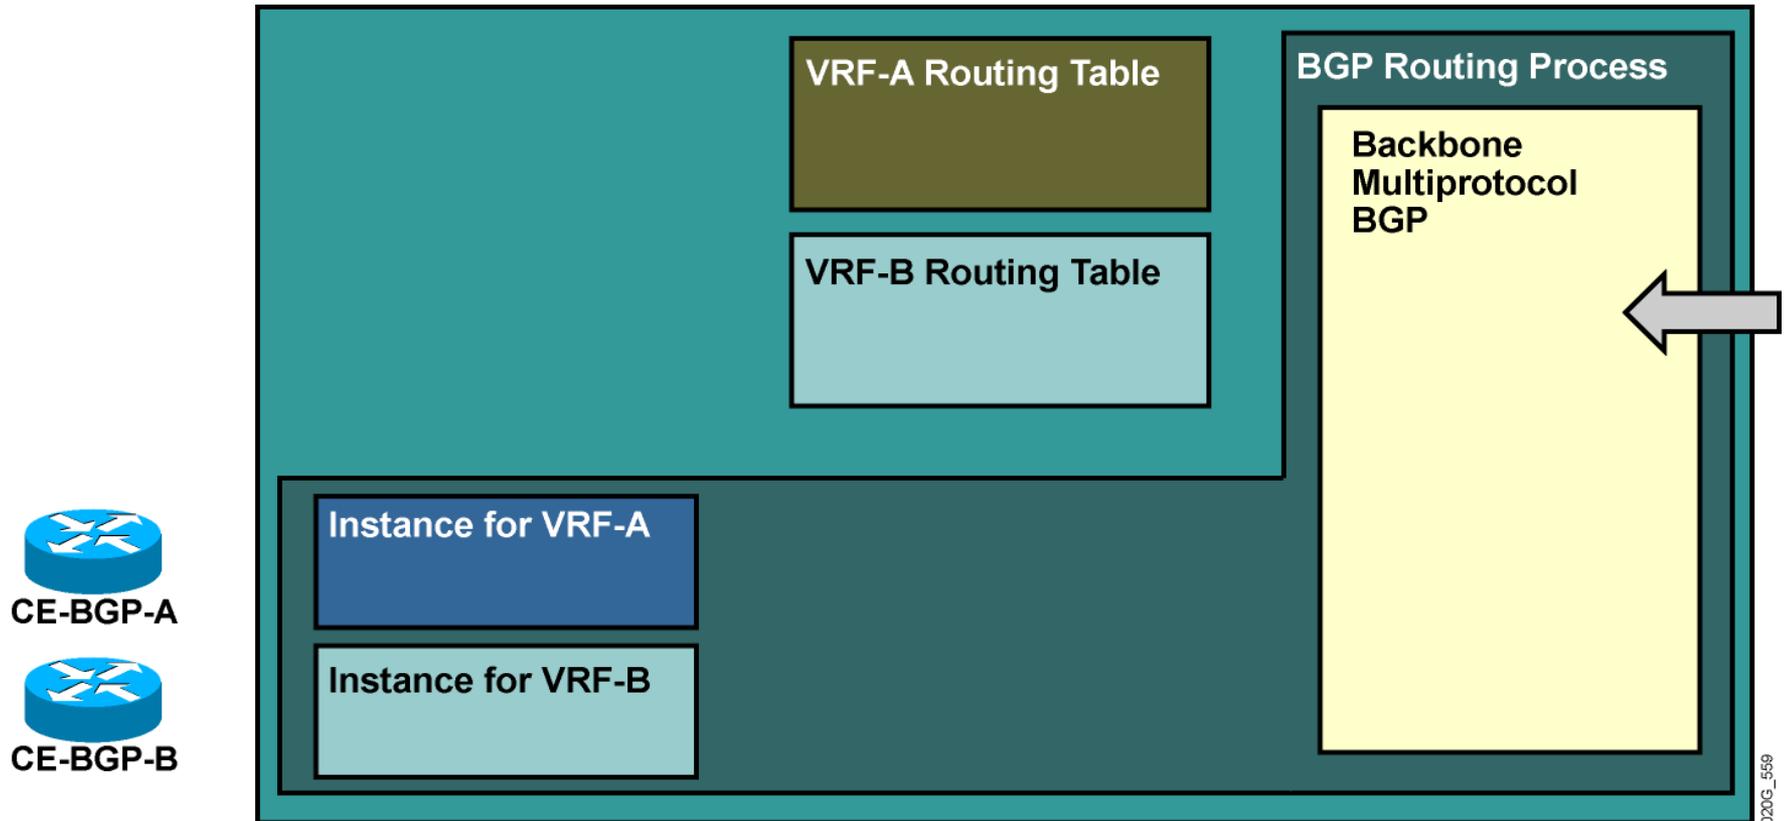
- RIP routes entered in the VRF routing table are redistributed into BGP for further propagation into the MPLS VPN backbone.

Non-BGP Route propagation—Outbound (Cont.)



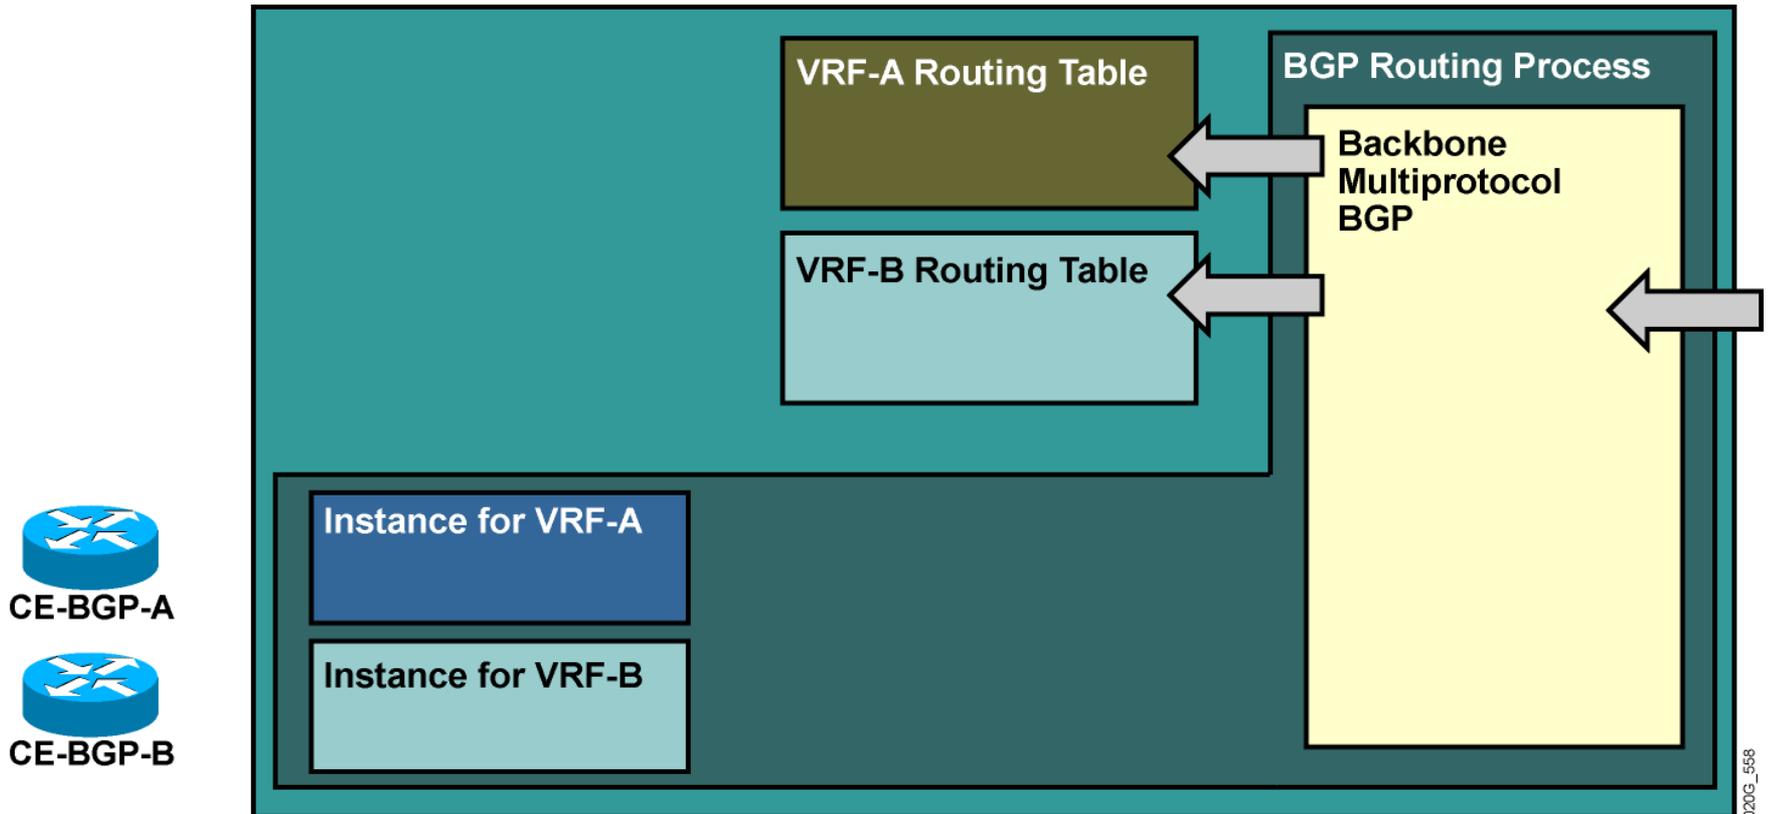
- RIP routes entered in the VRF routing table are redistributed into BGP for further propagation into the MPLS VPN backbone.
- **Redistribution between RIP and BGP has to be configured for proper MPLS VPN operation.**

Route Propagation—Inbound



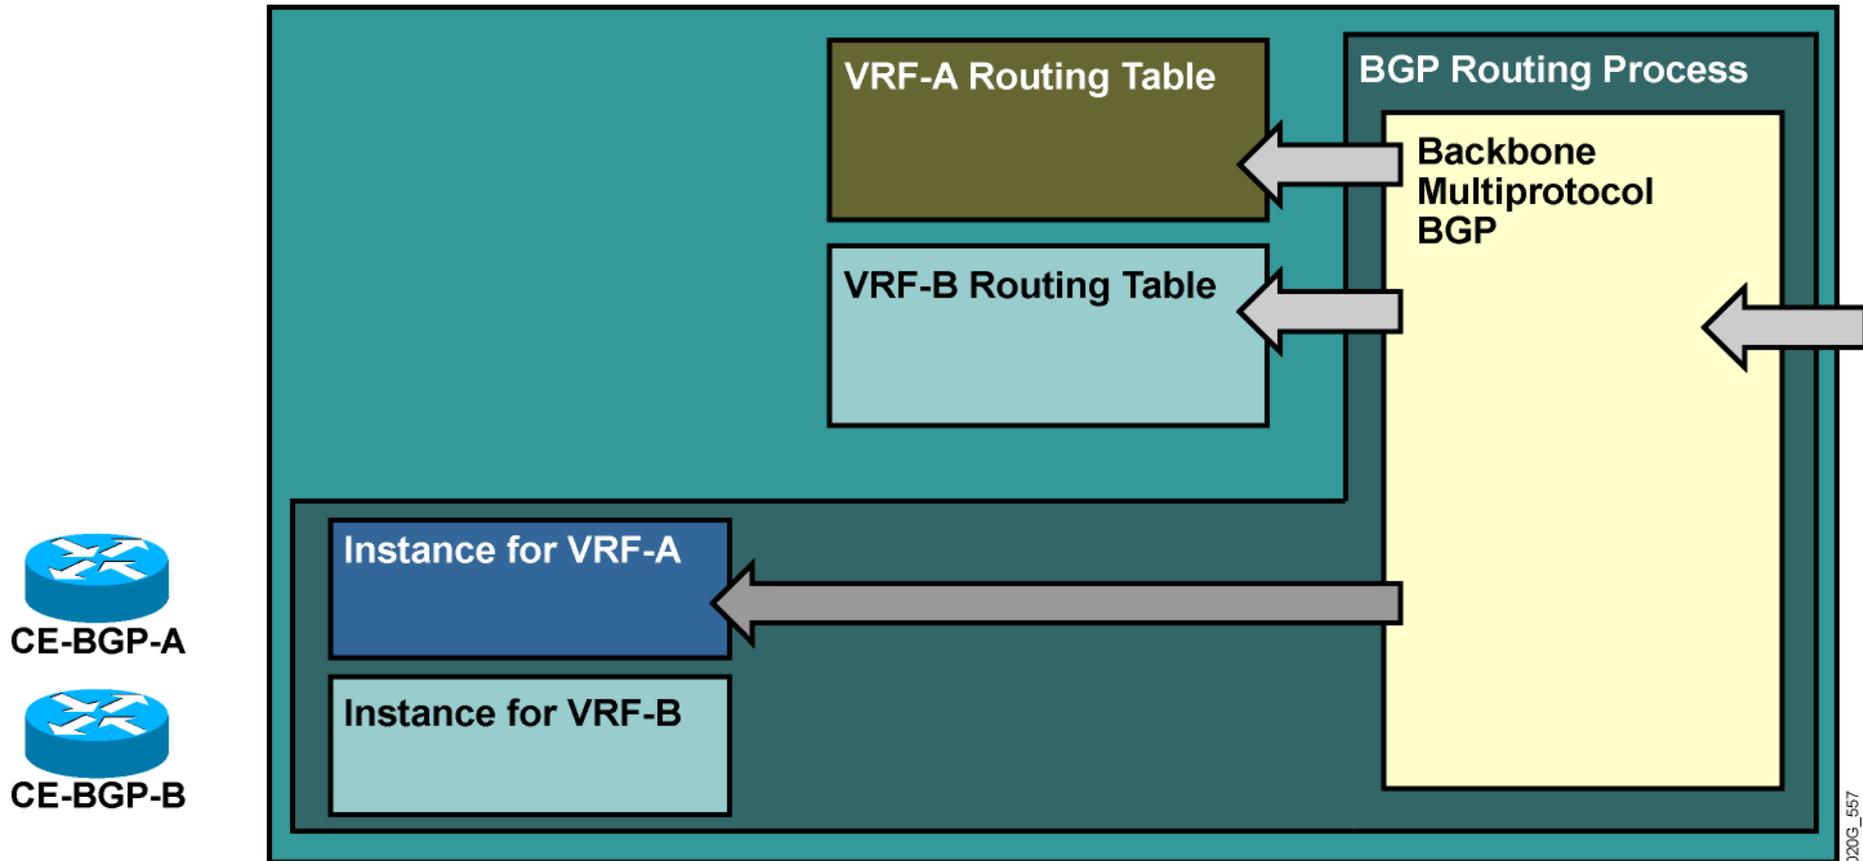
- VPNv4 prefixes are received from other PE routers.

Route Propagation—Inbound (Cont.)



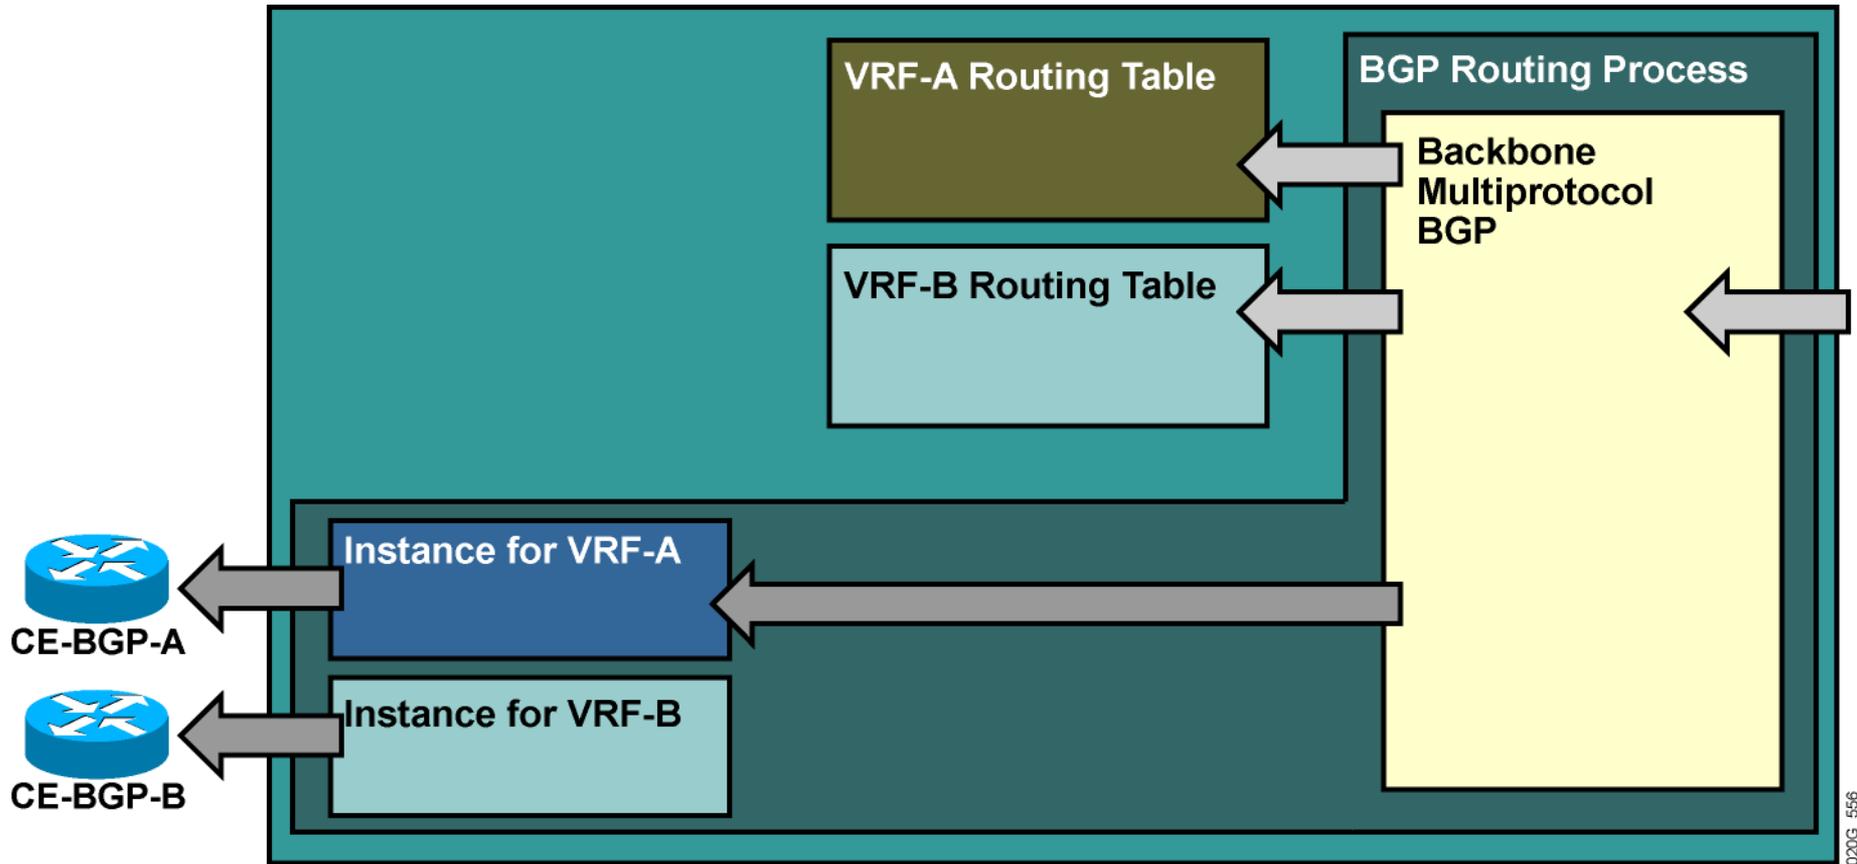
- VPNv4 prefixes are received from other PE routers.
- The VPNv4 prefixes are inserted into proper VRF routing tables based on their route targets and import route targets configured in VRFs.
- Route distinguisher is removed during this process.

Route Propagation—Inbound (Cont.)



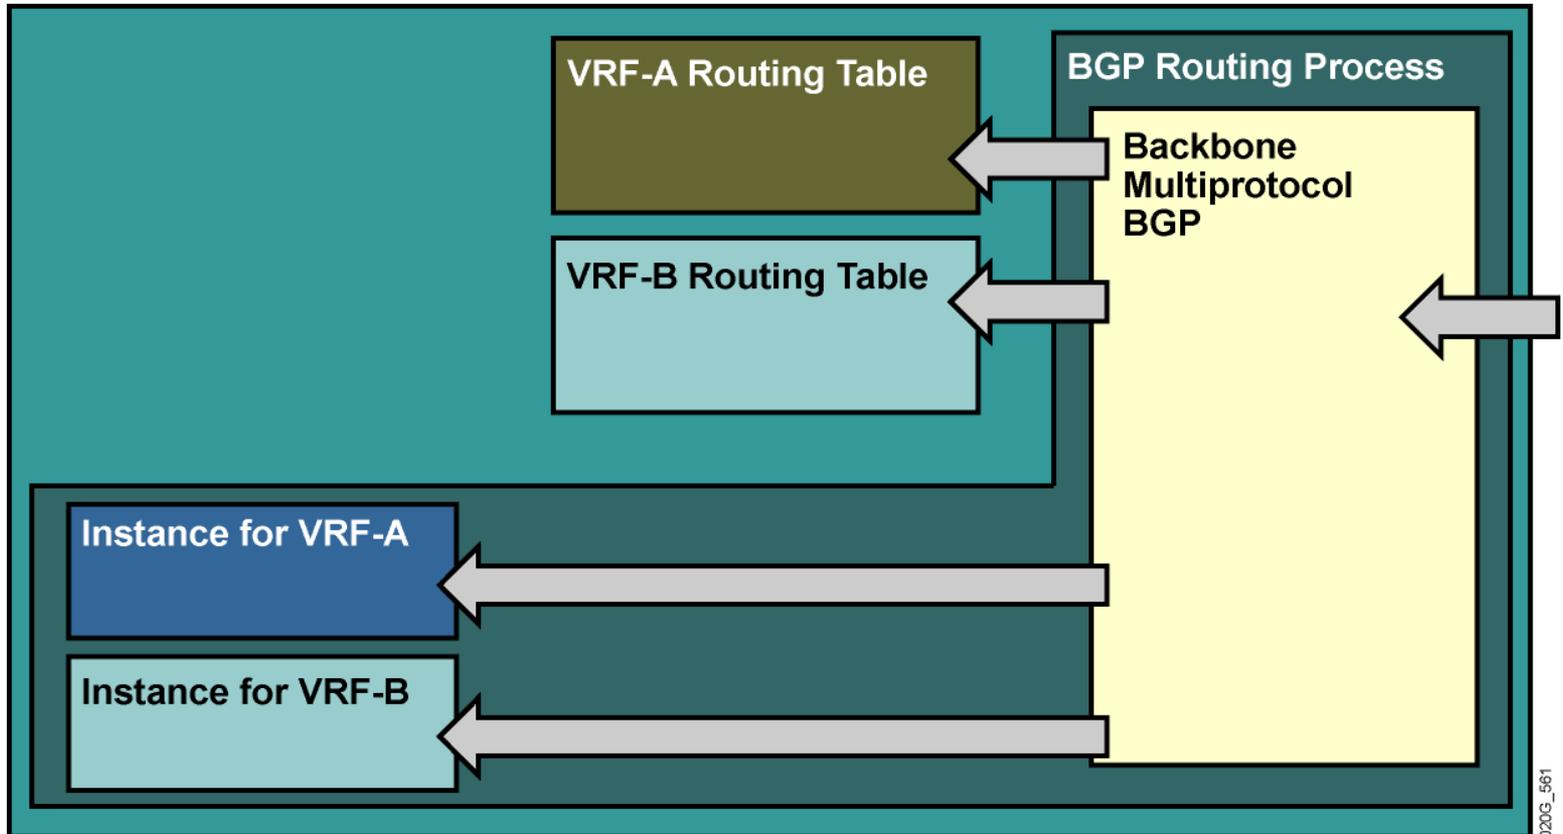
- Routes are received from backbone MP-BGP and imported into a VRF.

Route Propagation—Inbound (Cont.)



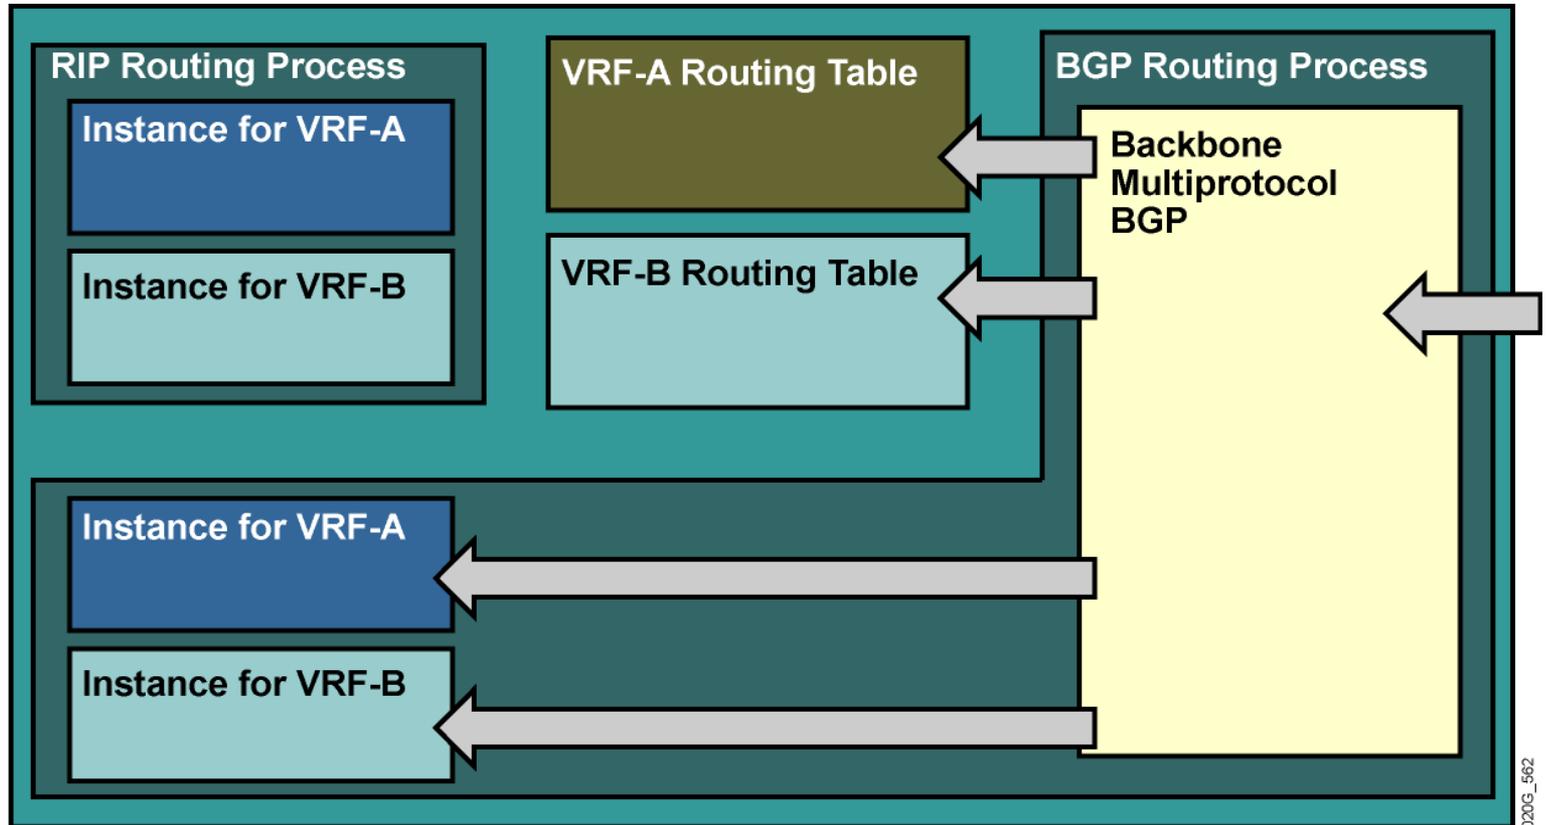
- Routes are received from backbone MP-BGP and imported into a VRF.
- IPv4 routes are forwarded to EBGP CE neighbors attached to that VRF.

Route Propagation—Inbound (Cont.)



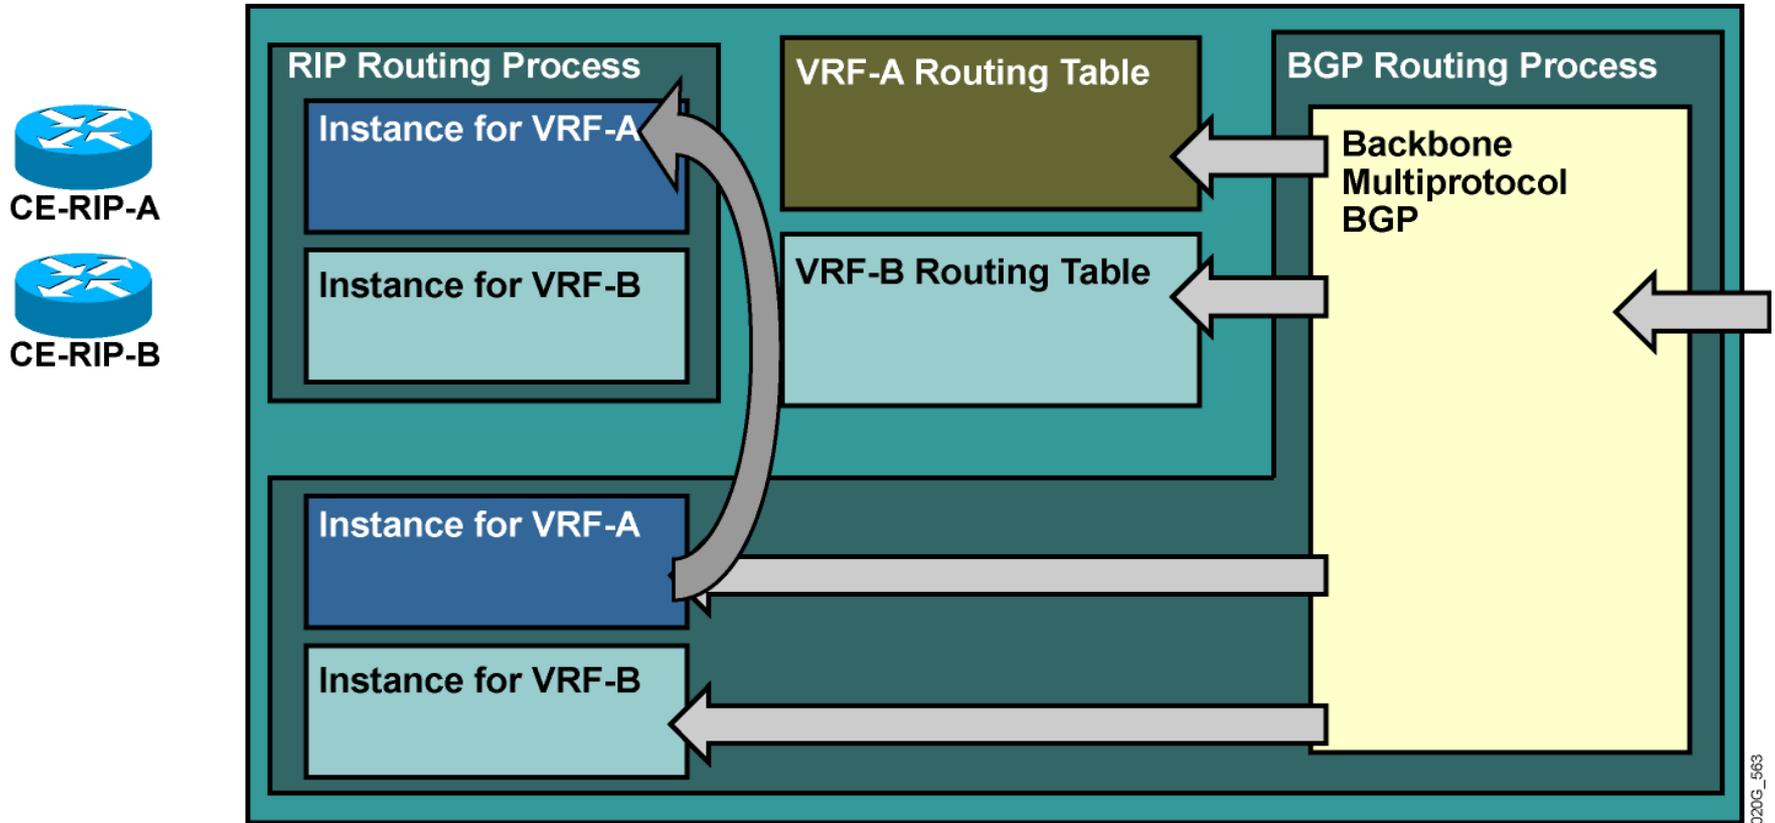
020G_561

Route Propagation—Inbound (Cont.)



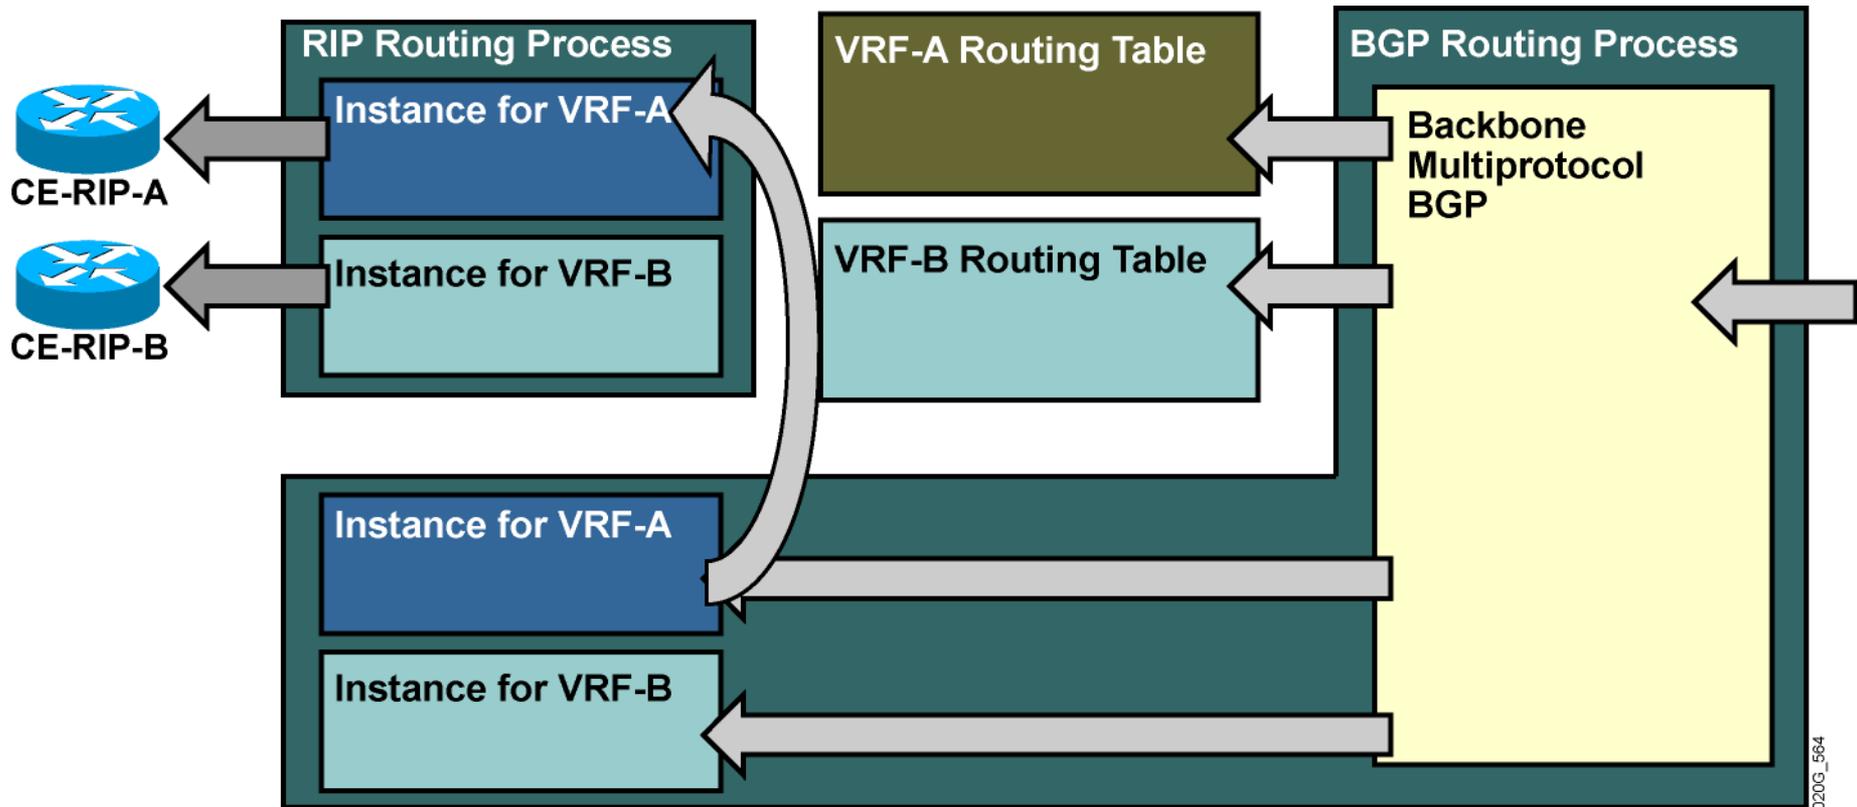
- MP-IBGP routes imported into a VRF are redistributed into the instance of RIP configured for that VRF.

Route Propagation—Inbound (Cont.)



- MP-IBGP routes imported into a VRF are redistributed into the instance of RIP configured for that VRF.
- Redistribution between BGP and RIP has to be configured for end-to-end RIP routing between CE routers.

Route Propagation—Inbound (Cont.)



- Routes redistributed from BGP into a VRF instance of RIP are sent to RIP-speaking CE routers.

Summary

A VRF is a routing and forwarding instance that you can use for a single VPN site or for many sites connected to the same PE router.

Routing contexts were introduced in Cisco IOS software to support the need for separate isolated copies of VPN routing protocols.

No limit to the number of interfaces associated with one VRF, but in practice, each interface can be assigned to only one VRF.



MPLS workshop

Configuring VRF Tables

Outline

Overview

VRF Configuration Tasks

Creating VRF Tables and Assigning RDs

Specifying Export and Import RTs

Assigning an Interface to VRF Table

Sample VPN Network Example

Lesson Summary

VRF Configuration Tasks

- VRF configuration tasks:

- Create a VRF table

- Assign RD to the VRF

- Specify export and import route targets

- Assign interfaces to VRFs

Creating VRF Tables and Assigning RDs

Router(config)#

```
ip vrf name
```

- Creates a new VRF or enters configuration of an existing VRF.
- VRF names are case-sensitive.
- VRF is not operational unless you configure RD.
- VRF names have only local significance.

Router(config-vrf)#

```
rd route-distinguisher
```

- Assigns a route distinguisher to a VRF.
- You can use ASN:nn or A.B.C.D:nn format for RD.
- Each VRF in a PE router has to have a unique RD.

Specifying Export and Import RTs

```
Router(config-vrf) #
```

```
route-target export RT
```

- Specifies an RT to be attached to every route exported from this VRF to MP-BGP
- Allows specification of many export RTs—all to be attached to every exported route

```
Router(config-vrf) #
```

```
route-target import RT
```

- Specifies an RT to be used as an import filter—only routes matching the RT are imported into the VRF
- Allows specification of many import RTs—any route where at least one RT attached to the route matches any import RT is imported into the VRF

Due to implementation issues, at least one export route target must also be an import route target of the same VRF in Cisco IOS Release 12.0 T.

Specifying Export and Import RTs (Cont.)

Router(config-vrf)#

```
route-target both RT
```

- In cases where the export RT matches the import RT, use this form of route-target command.

Sample router configuration for simple customer VPN:

```
ip vrf Customer_ABC  
rd 12703:15  
route-target export 12703:15  
route-target import 12703:15
```

Assigning an Interface to VRF Table

```
Router(config-if) #
```

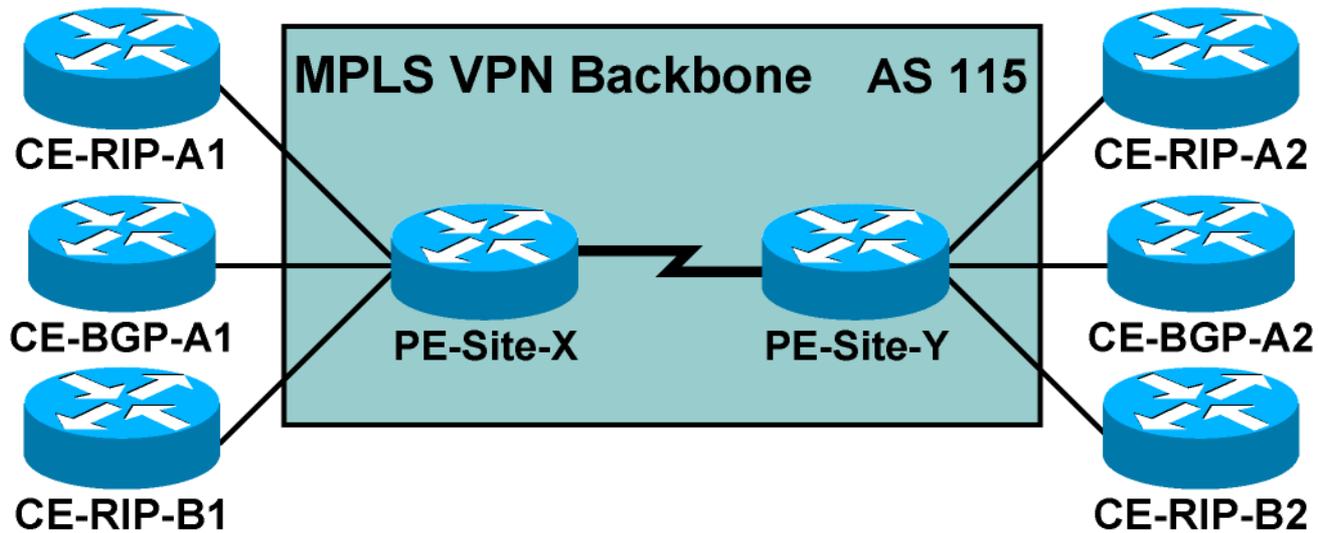
```
ip vrf forwarding vrf-name
```

- Associates an interface with the specified VRF.
- Existing IP address removed from the interface when interface is put into VRF—IP address must be reconfigured.
- CEF switching must be enabled on the interface.

Sample router configuration:

```
ip cef
!
interface serial 0/0
ip vrf forwarding Customer_ABC
ip address 10.0.0.1 255.255.255.252
```

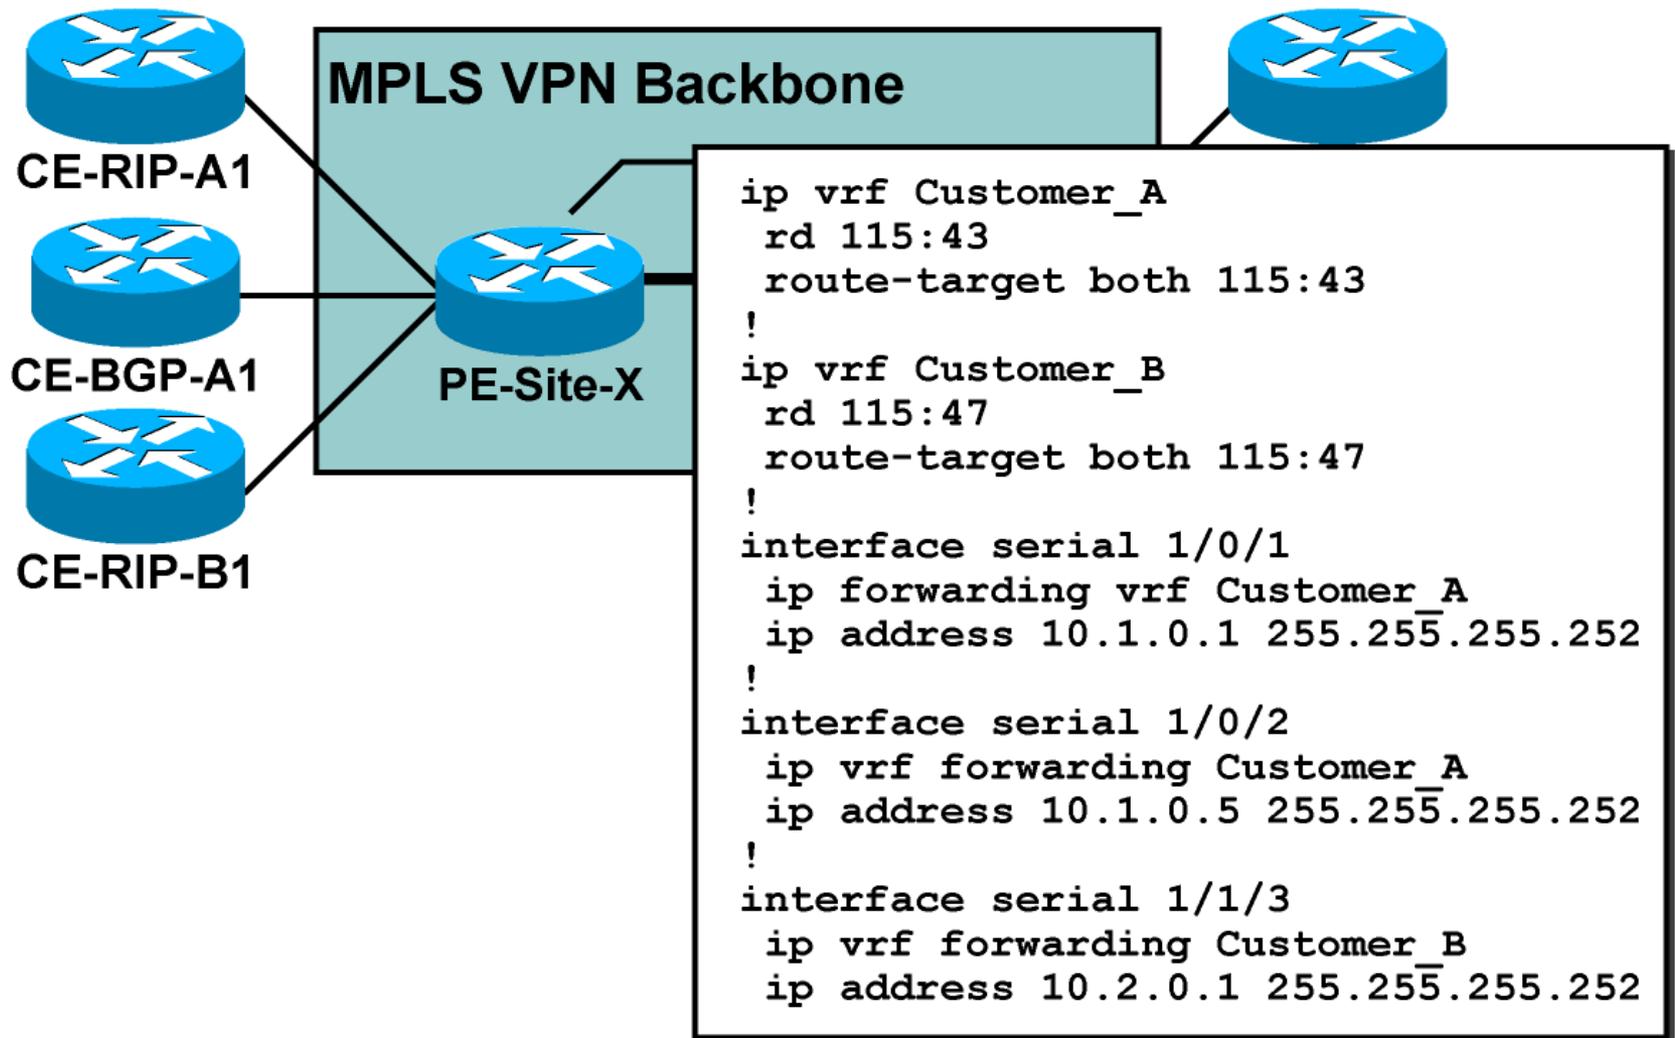
MPLS VPN Network Example



02016_070

- The network supports two VPN customers.
- Customer A runs RIP and BGP with the service provider; customer B uses only RIP.
- Both customers use network 10.0.0.0.

MPLS VPN Network Example (Cont.)



020G_071

Summary

A unique RD must be assigned to every VRF created in a PE router.

The same RD could be used on all PEs for simple VPN service.

For simple VPN service, import and export RT values should be the same.

Two formats for RD and RT are as follows:

ASN:nn

A.B.C.D:nn



MPLS workshop

Configuring an MP-BGP Session Between PE routers

Outline

Overview

Configuring BGP Address families

BGP Neighbors

Configuring MP-BGP

Configuring MP-IBGP

MP-BGP BGP Community Propagation

Disabling IPv4 Route Exchange

Verifying Configurations

Lesson Summary

Configuring BGP Address Families

The BGP process in an MPLS VPN-enabled router performs three separate tasks:

- Global BGP routes (Internet routing) are exchanged as in traditional BGP setup.

- VPNv4 prefixes are exchanged through MP-BGP.

- VPN routes are exchanged with CE routers through per-VRF EBGP sessions.

Address families (routing protocol contexts) are used to configure these three tasks in the same BGP process.

Configuring BGP Address Families (Cont.)

Router(config)#

```
router bgp as-number
```

- Selects global BGP routing process

Router(config-router)#

```
address-family vpnv4
```

- Selects configuration of VPNv4 prefix exchanges under MP-BGP sessions

Router(config-router)#

```
address-family ipv4 vrf vrf-name
```

- Selects configuration of per-VRF PE-CE EBGP parameters

BGP Neighbors

MP-BGP neighbors are configured under the BGP routing process:

- These neighbors need to be activated for each global address family that they support.

- Per-address-family parameters can be configured for these neighbors.

VRF-specific EBGP neighbors are configured under corresponding address families.

Configuring MP-BGP

- MPLS VPN MP-BGP configuration steps:

Configure MP-BGP neighbor under BGP routing process.

Configure BGP address family VPNv4.

Activate configured BGP neighbor for VPNv4 route exchange.

Specify additional parameters for VPNv4 route exchange (filters, next hops, and so on).

Configuring MP-IBGP

Router(config)#

```
router bgp as-number  
  neighbor ip-address remote-as as-number  
  neighbor ip-address update-source loopback-type interface number
```

- All MP-BGP neighbors have to be configured under global BGP routing configuration.
- MP-IBGP sessions have to run between loopback interfaces.

Router(config-router)#

```
address-family vpnv4
```

- Starts configuration of MP-BGP routing for VPNv4 route exchange.
- Parameters that apply only to MP-BGP exchange of VPNv4 routes between already configured IBGP neighbors are configured under this address family.

Configuring MP-IBGP (Cont.)

```
Router(config-router-af) #
```

```
neighbor ip-address activate
```

- The BGP neighbor defined under BGP router configuration has to be activated for VPNv4 route exchange.

```
Router(config-router-af) #
```

```
neighbor ip-address next-hop-self
```

- The next-hop-self keyword can be configured on the MP-IBGP session. With current IOS, this is enabled by default

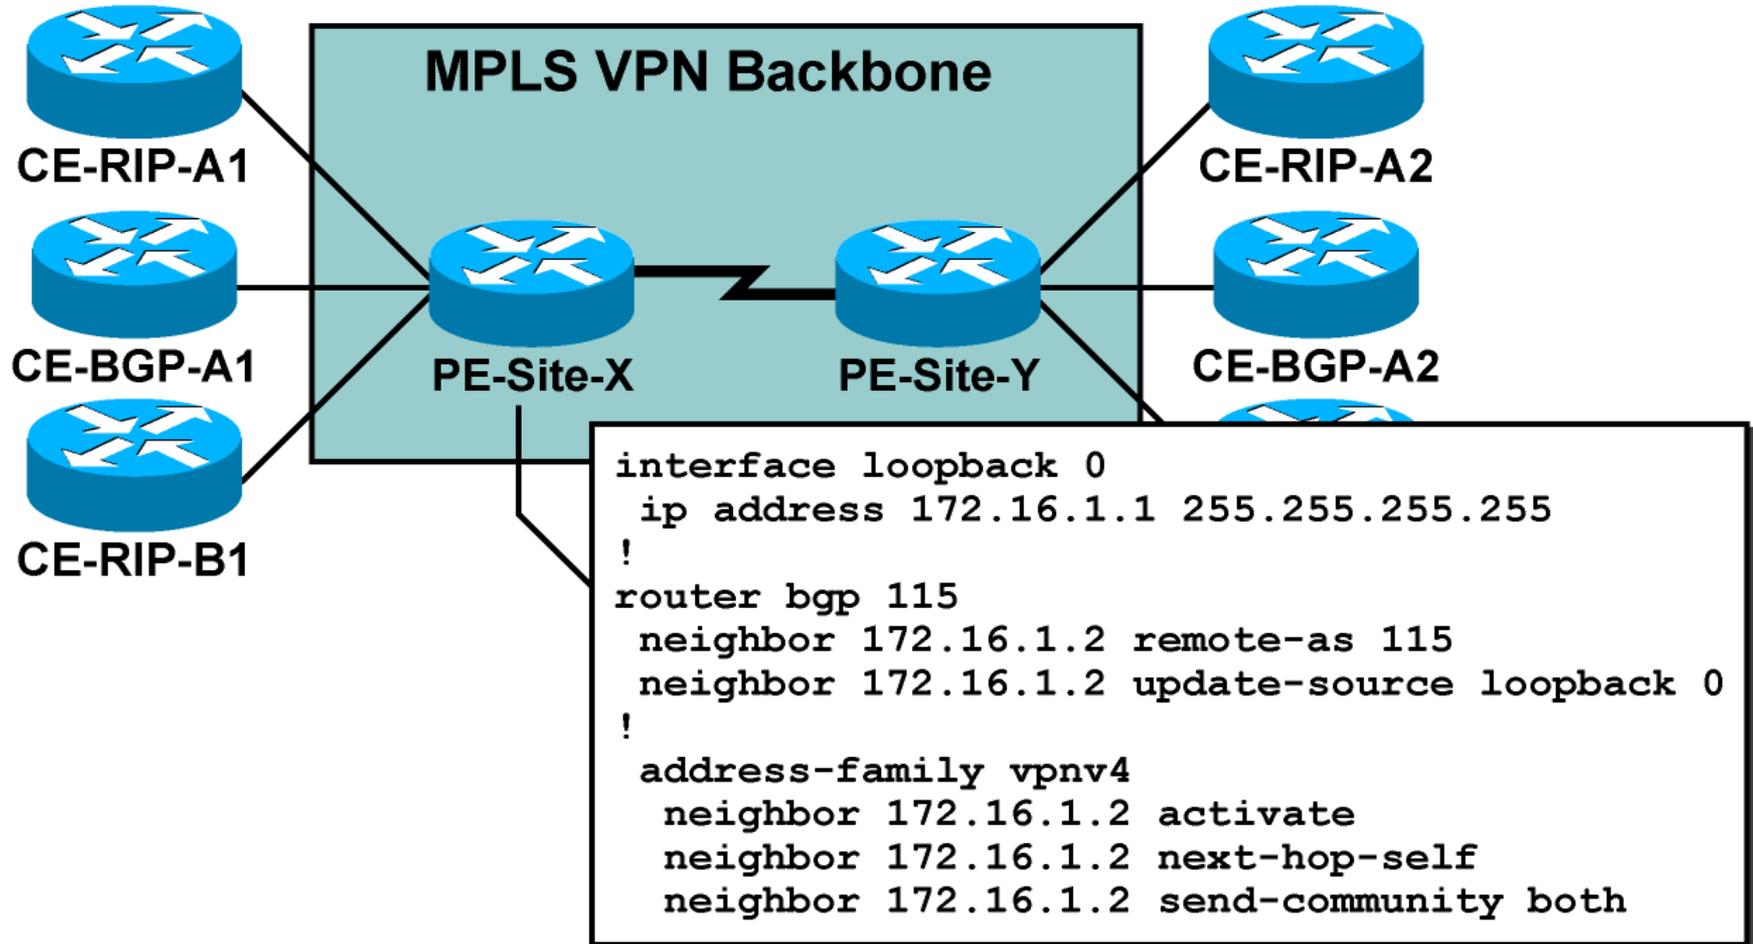
MP-BGP BGP Community Propagation

```
Router(config-router-af) #
```

```
neighbor ip-address send-community [extended | both]
```

- This command configures propagation of standard and extended BGP communities attached to VPNv4 prefixes.
- Default value: only extended communities are sent.
- Usage guidelines:
 - Extended BGP communities attached to VPNv4 prefixes **have to be exchanged** between MP-BGP neighbors for proper MPLS VPN operation.
 - To propagate standard BGP communities between MP-BGP neighbors, use the both option.

MP-BGP BGP Community Propagation (Cont.)



Disabling IPv4 Route Exchange

```
Router(config-router) #
```

```
no bgp default ipv4 unicast
```

- Exchange of IPv4 routes between BGP neighbors is enabled by default—every configured neighbor will also receive IPv4 routes.
- This command disables default exchange of IPv4 routes—neighbors that need to receive IPv4 routes have to be activated for IPv4 route exchange.
- Use this command when the same router carries Internet and VPNv4 routes and you do not want to propagate Internet routes to some PE neighbors.

Disabling IPv4 Route Exchange (Cont.)

- Neighbor 172.16.32.14 receives only Internet routes.
- Neighbor 172.16.32.15 receives only VPNv4 routes.
- Neighbor 172.16.32.27 receives Internet and VPNv4 routes.

```
router bgp 12703
  no bgp default ipv4 unicast
  neighbor 172.16.32.14 remote-as 12703
  neighbor 172.16.32.15 remote-as 12703
  neighbor 172.16.32.27 remote-as 12703

! Activate IPv4 route exchange

neighbor 172.16.32.14 activate
neighbor 172.16.32.27 activate

! Step#2 - VPNv4 route exchange

address-family vpnv4
  neighbor 172.16.32.15 activate
  neighbor 172.16.32.27 activate
```

MPLS/VPN Monitoring Commands

router#

```
telnet host /vrf name
```

- Performs PE - CE telnet through specified VRF

router#

```
ping vrf name ...
```

- Performs ping based on VRF routing table

router#

```
trace vrf name ...
```

- Performs VRF-based traceroute

show ip vrf

```
Router#show ip vrf
```

Name	Default RD	Interfaces
SiteA2	103:10	Serial1/1.1
SiteB	103:20	Serial1/1.2
SiteX	103:30	Ethernet0/0

show ip vrf interfaces

```
Router#show ip vrf interfaces
```

Interface	IP-Address	VRF	Protocol
Serial1/1.1	150.1.31.37	SiteA2	up
Serial1/1.2	150.1.32.33	SiteB	up
Ethernet0/0	192.168.22.3	SiteX	up

Monitoring VRF Routing

router#

```
show ip protocol vrf name
```

- Displays the routing protocols configured in a VRF

router#

```
show ip route vrf name ...
```

- Displays the VRF routing table

router#

```
show ip bgp vpnv4 vrf name ...
```

- Displays per-VRF BGP parameters (PE-CE neighbors ...)

show ip protocol vrf

```
Router#show ip protocol vrf SiteX
```

```
Routing Protocol is "rip"
```

```
  Sending updates every 30 seconds, next due in 10 seconds
```

```
  Invalid after 180 seconds, hold down 180, flushed after 240
```

```
  Outgoing update filter list for all interfaces is
```

```
  Incoming update filter list for all interfaces is
```

```
  Redistributing: rip, bgp 3
```

```
  Default version control: send version 2, receive version 2
```

```
    Interface Send Recv Triggered RIP Key-chain
```

```
    Ethernet0/0 2 2
```

```
  Routing for Networks:
```

```
    192.168.22.0
```

```
  Routing Information Sources:
```

```
    Gateway Distance Last Update
```

```
  Distance: (default is 120)
```

show ip route vrf

```
Router#show ip route vrf SiteA2
```

```
Codes: C - connected,S -static,I -IGRP,R -RIP,M -mobile,B -BGP  
D - EIGRP, EX - EIGRP external,O -OSPF,IA -OSPF interarea  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP  
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,  
* - candidate default, U - per-user static route,o -ODR  
P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
O      203.1.20.0/24 [110/782] via 150.1.31.38, 02:52:13, Serial1/1.1  
      203.1.2.0/32 is subnetted, 1 subnets  
O      203.1.2.1 [110/782] via 150.1.31.38, 02:52:13, Serial1/1.1  
      203.1.1.0/32 is subnetted, 1 subnets  
B      203.1.1.1 [200/1] via 192.168.3.103, 01:14:32  
B      203.1.135.0/24 [200/782] via 192.168.3.101, 02:05:38  
B      203.1.134.0/24 [200/1] via 192.168.3.101, 02:05:38  
B      203.1.10.0/24 [200/1] via 192.168.3.103, 01:14:32
```

```
... rest deleted ...
```

show ip bgp vpnv4 vrf neighbor

```
Router#show ip bgp vpnv4 vrf SiteB neighbor
BGP neighbor is 150.1.32.34, vrf SiteB, remote AS 65032, external link
  BGP version 4, remote router ID 203.2.10.1
  BGP state = Established, up for 02:01:41
  Last read 00:00:56, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received
    Address family IPv4 Unicast: advertised and received
  Received 549 messages, 0 notifications, 0 in queue
  Sent 646 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds

For address family: VPNv4 Unicast
  Translates address family IPv4 Unicast for VRF SiteB
  BGP table version 416, neighbor version 416
  Index 4, Offset 0, Mask 0x10
  Community attribute sent to this neighbor
  2 accepted prefixes consume 120 bytes
  Prefix advertised 107, suppressed 0, withdrawn 63

... rest deleted ...
```

Monitoring MP-BGP Sessions

router#

```
show ip bgp neighbor
```

- Displays global BGP neighbors and the protocols negotiated with these neighbors

Monitoring MP-BGP VPNv4 Table

router#

```
show ip bgp vpnv4 all
```

- Displays whole VPNv4 table

router#

```
show ip bgp vpnv4 vrf name
```

- Displays only BGP parameters (routes or neighbors) associated with specified VRF
- Any BGP show command can be used with these parameters

router#

```
show ip bgp vpnv4 rd value
```

- Displays only BGP parameters (routes or neighbors) associated with specified RD

Monitoring per-VRF CEF and LFIB Structures

router#

```
show ip cef vrf name
```

- Displays per-VRF CEF table

router#

```
show ip cef vrf name prefix detail
```

- Displays details of individual CEF entry, including label stack

router#

```
show tag-switching forwarding vrf name
```

- Displays labels allocated by MPLS/VPN for routes in specified vrf

Summary

MPLS VPN architecture uses the BGP routing protocol in two ways:

- VPNv4 routes are propagated across an MPLS VPN backbone using MP-BGP between the PE routers.

- BGP can be used as the PE-CE routing protocol to exchange VPN routes between the PE routers and the customer edge (CE) routers.

Only one BGP process can be configured per router.

Routing protocol contexts are used to configure independent route exchange mechanisms.



MPLS workshop

Configuring Static routes and BGP as PE-CE routing protocol

Outline

Static Route as PE-CE Protocol

Benefits of BGP as PE-CE protocol

Configuring per-VRF BGP Routing Context

Limiting the Number of Routes in a VRF

Limiting the Number of Prefixes Received from a BGP Neighbor

AS-Override

Hub and Spoke setup in MPLS VPNs

AllowAS-in

Implementing Site of Origin (SOO) for loop prevention

Selective Import

Selective Export

Lesson Summary

Configuring Per-VRF Static Routes

```
router(config)#
```

```
ip route vrf name static route parameters
```

- This command configures per-VRF static routes
- The route is entered in the specified Virtual Routing Table
- You always have to specify outgoing interface, even if you specify the next-hop

Sample router configuration:

```
ip route vrf Customer_ABC 10.0.0.0 255.0.0.0 10.250.0.2  
serial 0/0  
!  
router bgp 12703  
address-family ipv4 vrf Customer_ABC  
redistribute static
```

Benefits of using BGP as PE-CE protocol

- BGP allows continuity of policies between sites
- Use of private AS numbers for VPN sites allows easier configuration and saves AS numbers
- No redistribution involved
- Standard Communities for routing policies between sites
- Route-map and filters based on BGP attributes
- BGP sessions can be authenticated
- PE can limit the total number of prefixes the CE is allowed to announce — Avoids impact of CE mis-configuration

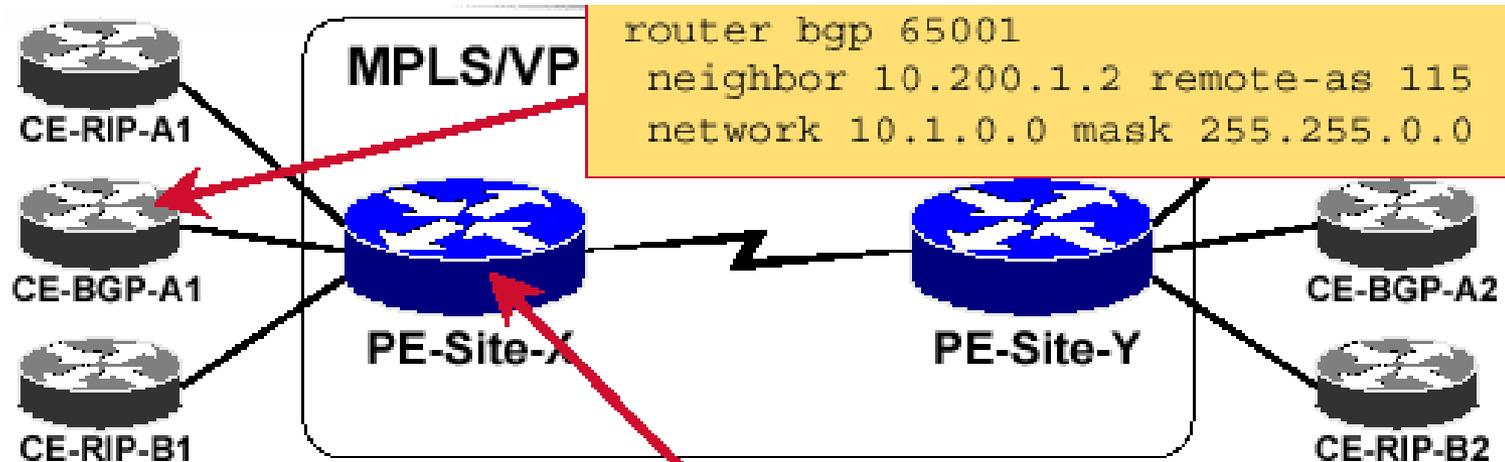
Configuring Per-VRF BGP Routing Context

Router(config)#

```
router bgp as-number  
  address-family ipv4 vrf vrf-name  
  ... Per-VRF BGP definitions ...
```

- There is only one BGP process per router
- Per-VRF parameters are specified in **routing contexts**, which are selected with the **address family** command
- Select per-VRF BGP context with the address-family command.
- Configure CE eBGP neighbors in VRF context, not in the global BGP configuration.
- CE neighbors have to be activated with the neighbor activate command.

Configuring Per-VRF BGP Routing Context (Cont.)



```
router bgp 65001
  neighbor 10.200.1.2 remote-as 115
  network 10.1.0.0 mask 255.255.0.0
```

```
router bgp 115
!
  address-family ipv4 vrf Customer_A
    neighbor 10.200.1.1 remote-as 65001
    neighbor 10.200.1.1 activate
```

Limiting the Number of Routes in a VRF

Service providers offering MPLS VPN services are at risk of denial-of-service attacks similar to those aimed at ISPs offering BGP connectivity:

Any customer can generate any number of routes, using resources in the PE routers.

Therefore, resources used by a single customer have to be limited.

Cisco IOS software offers two solutions:

1. It can limit the number of routes received from a BGP neighbor.
2. It can limit the total number of routes in a VRF.

Limiting the Number of Prefixes Received from a BGP Neighbor

```
Router(config-router-af) #
```

```
neighbor ip-address maximum-prefix maximum [threshold]  
[warning-only]
```

- Controls how many prefixes can be received from a neighbor
- Optional *threshold* parameter specifies the percentage where a warning message is logged (default is 75 percent)
- Optional warning-only keyword specifies the action on exceeding the maximum number (default is to drop peering)

Limiting the Total Number of VRF Routes

- The VRF route limit command limits the number of routes that are imported into a VRF:
 - Routes coming from CE routers
 - Routes coming from other PEs (imported routes)
- The route limit is configured for each VRF.
- If the number of routes exceeds the route limit:
 - Syslog message is generated.
 - The Cisco IOS software can be configured to reject routes (optional).

Limiting the Total Number of VRF Routes (Cont.)

```
Router(config-vrf) #
```

```
maximum routes limit {warn threshold | warn-only}
```

- This command configures the maximum number of routes accepted into a VRF:

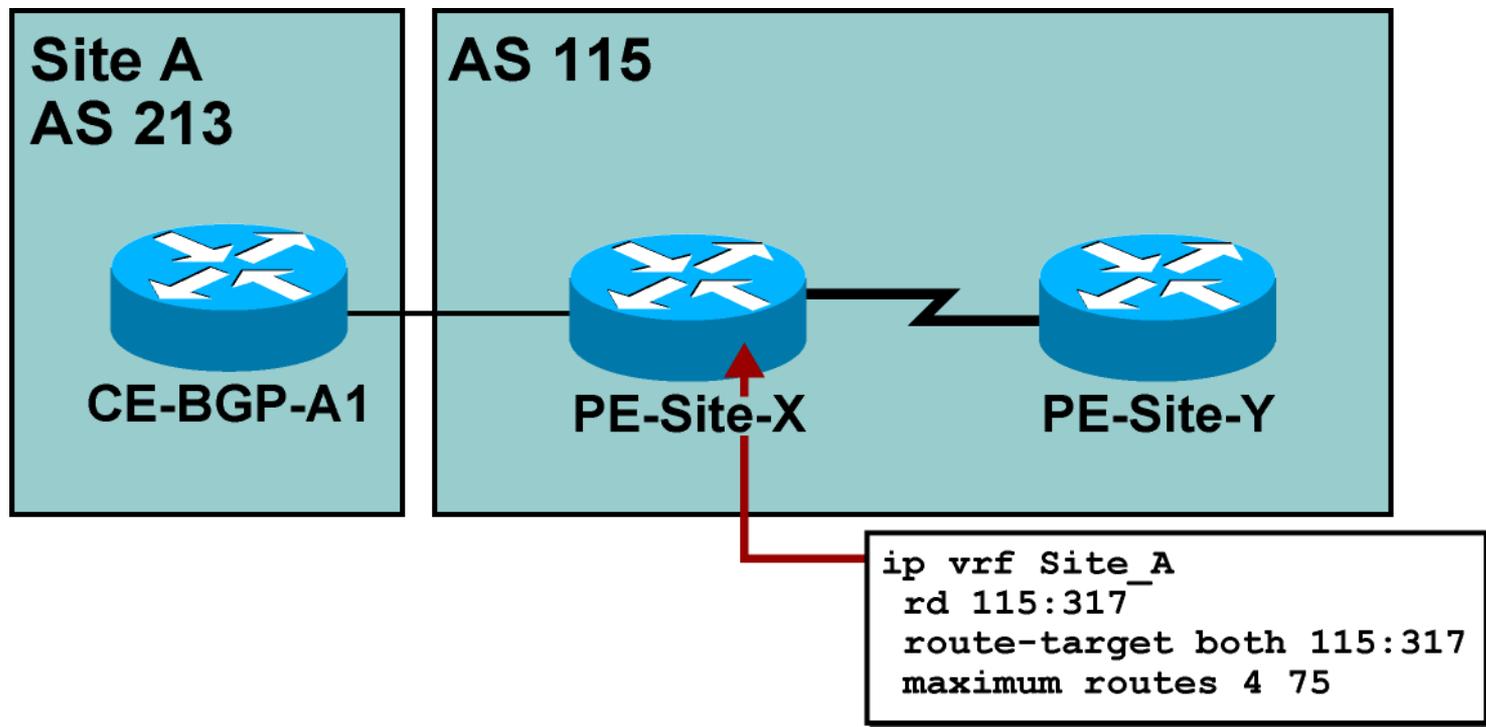
limit is the route limit for the VRF.

warn threshold is the percentage value over which a warning message is sent to syslog.

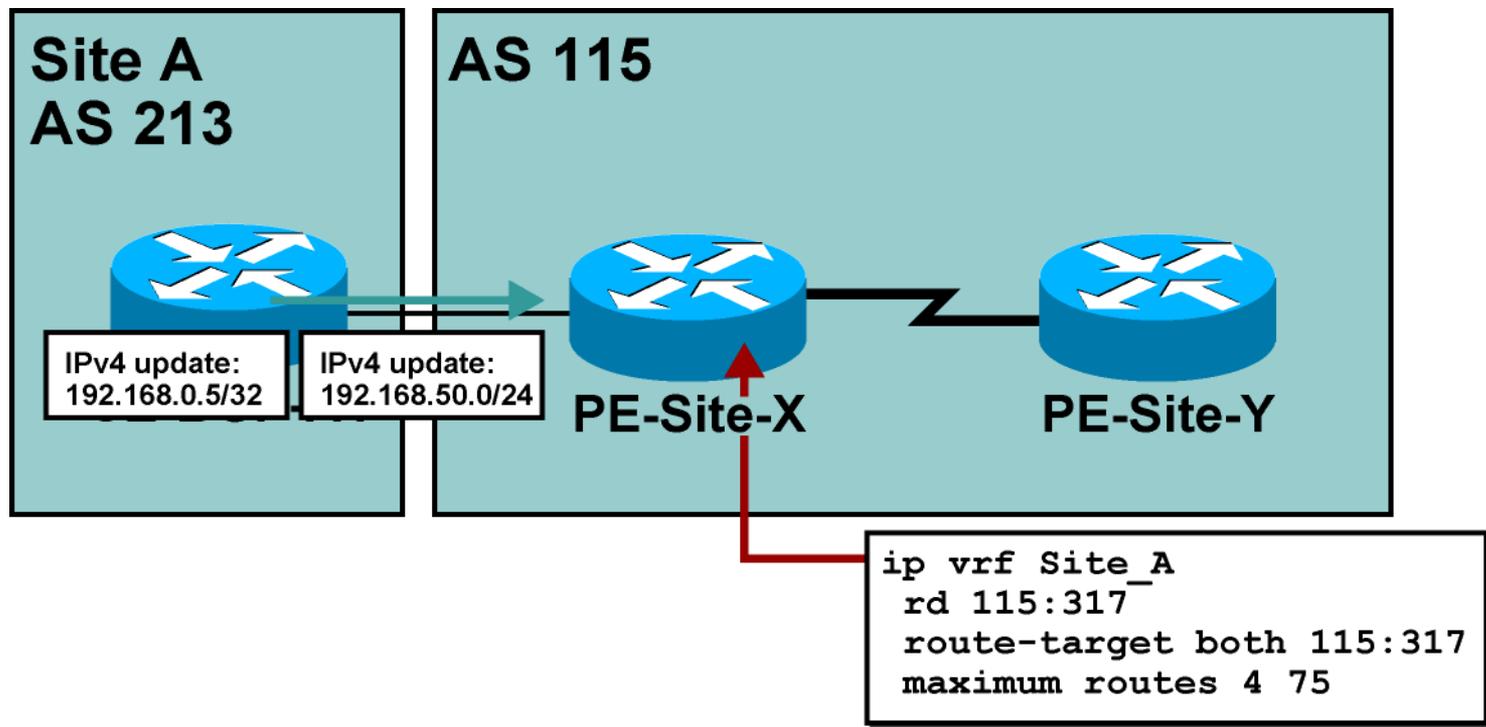
With warn-only the PE continues accepting routes after the configured limit.

- Syslog messages generated by this command are rate-limited.

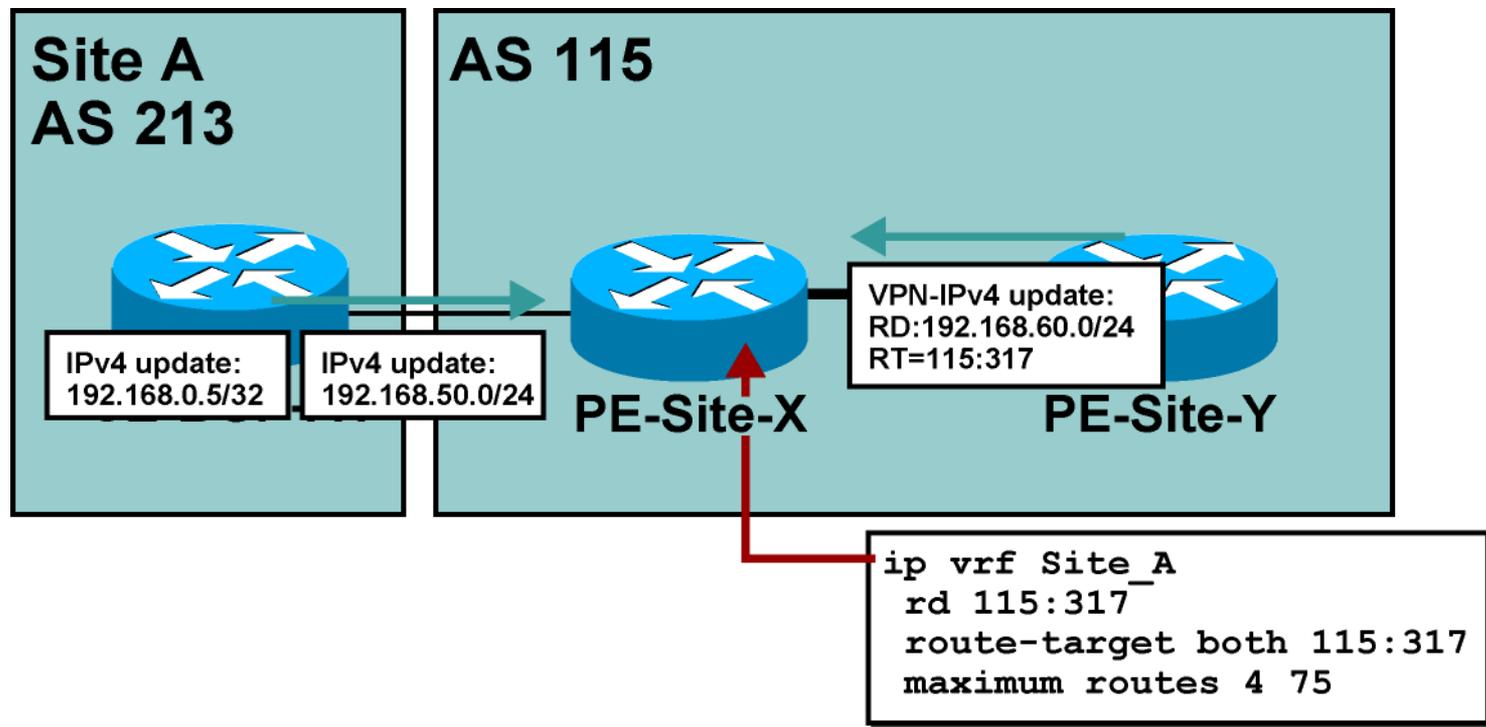
Limiting the Total Number of VRF Routes (Cont.)



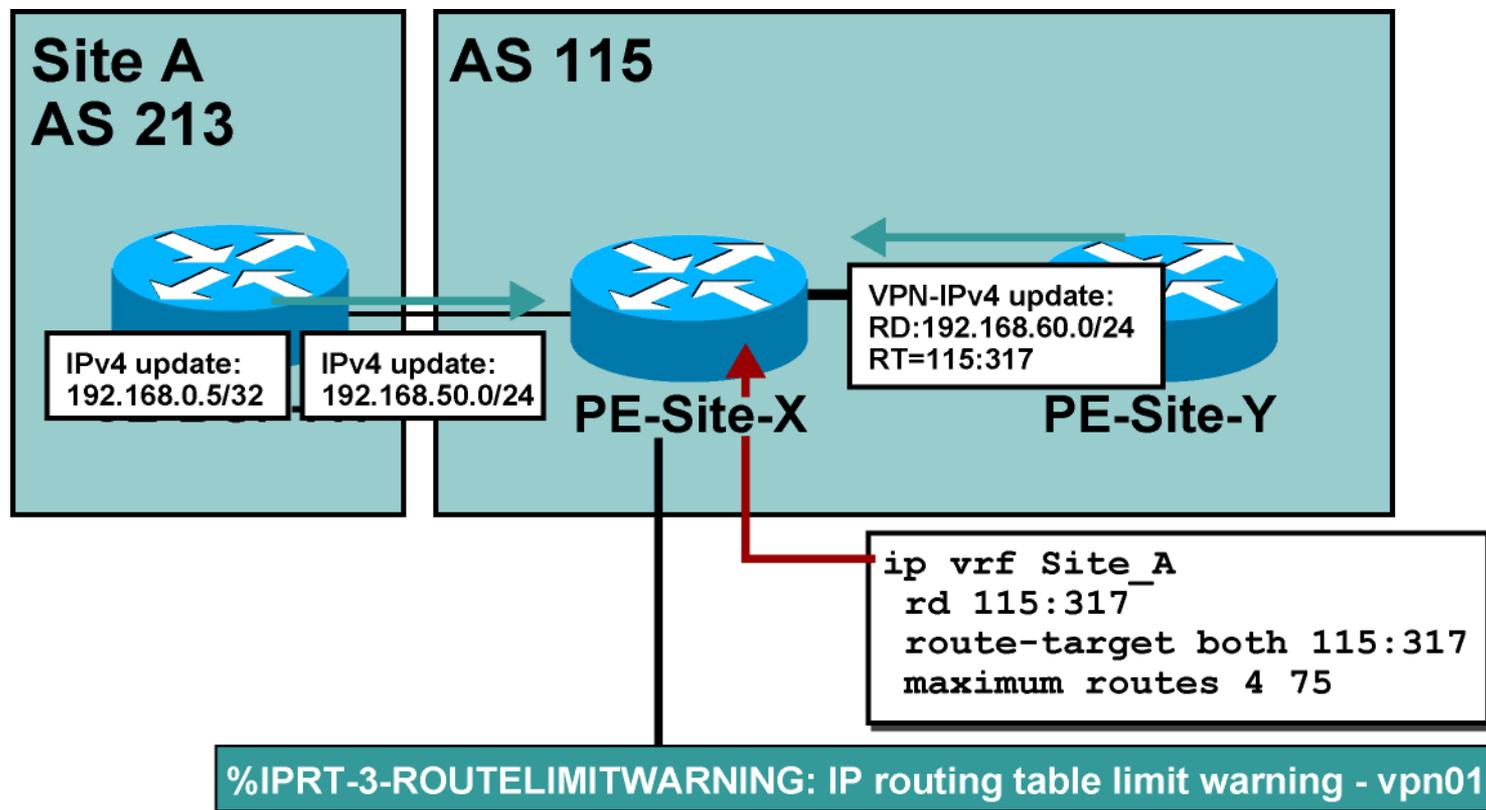
Limiting the Total Number of VRF Routes (Cont.)



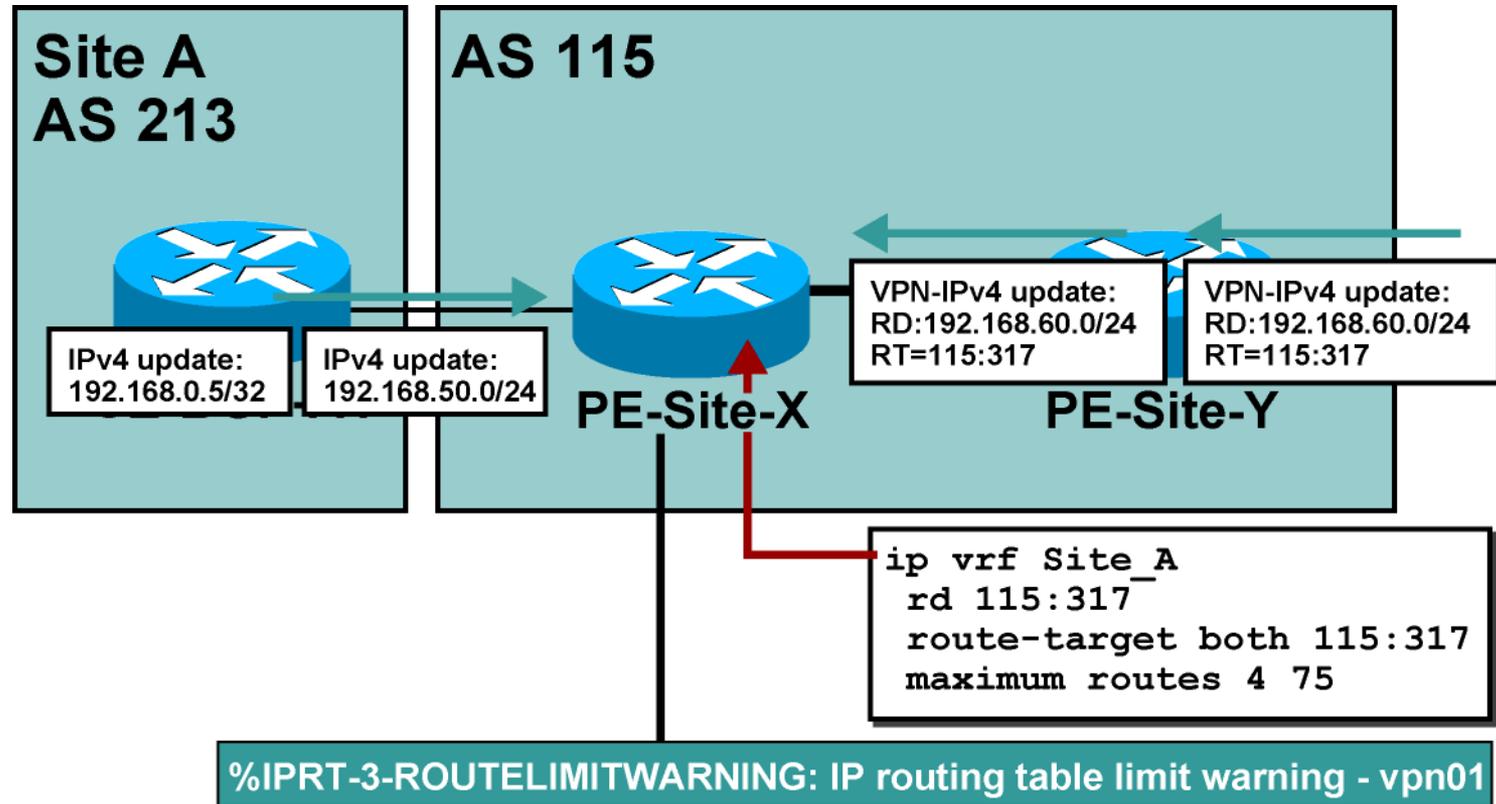
Limiting the Total Number of VRF Routes (Cont.)



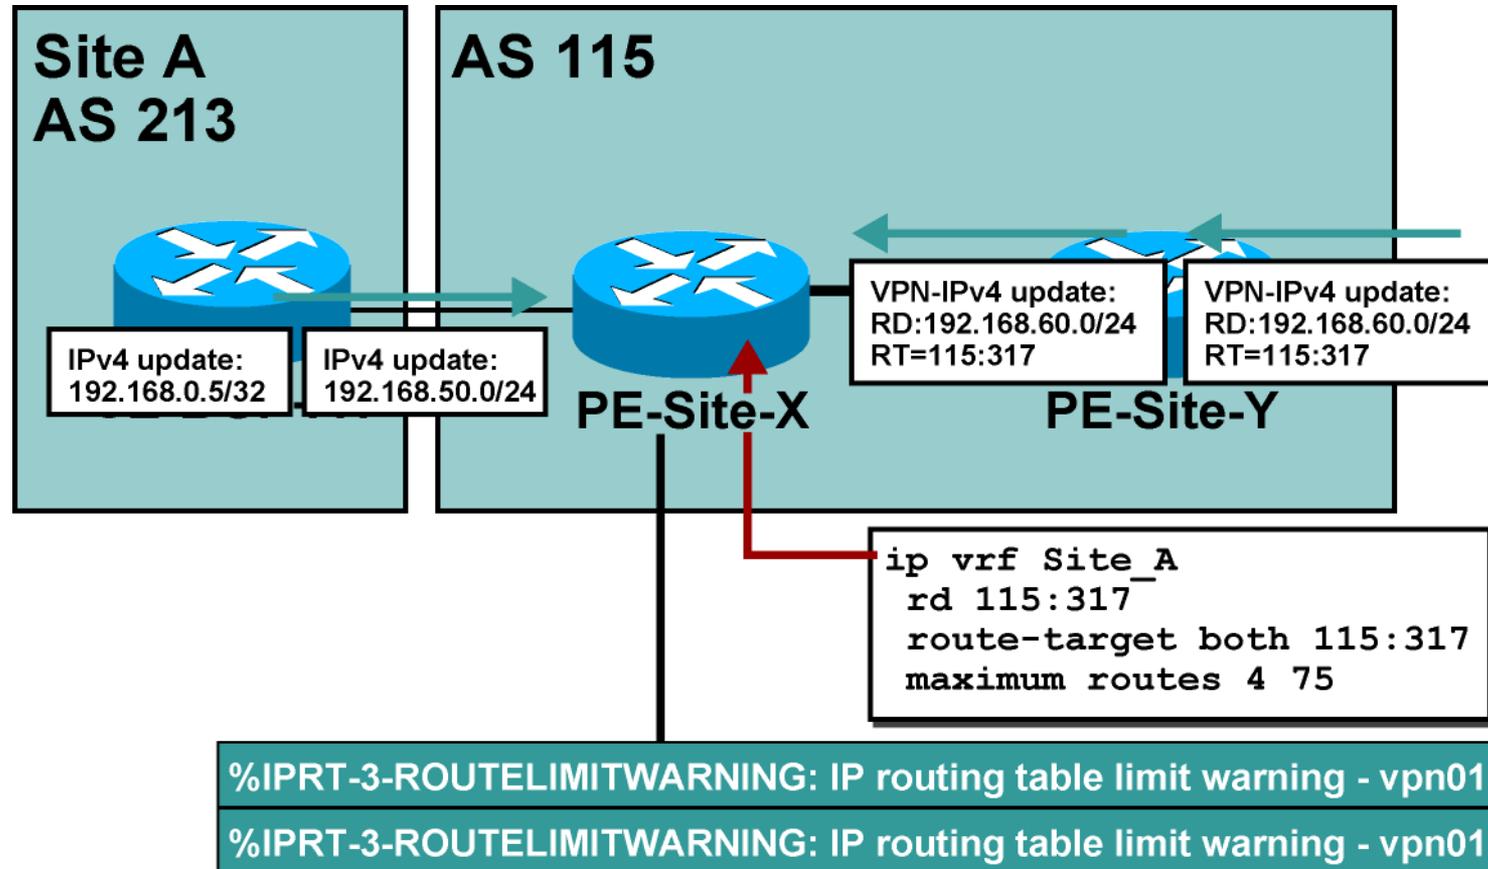
Limiting the Total Number of VRF Routes (Cont.)



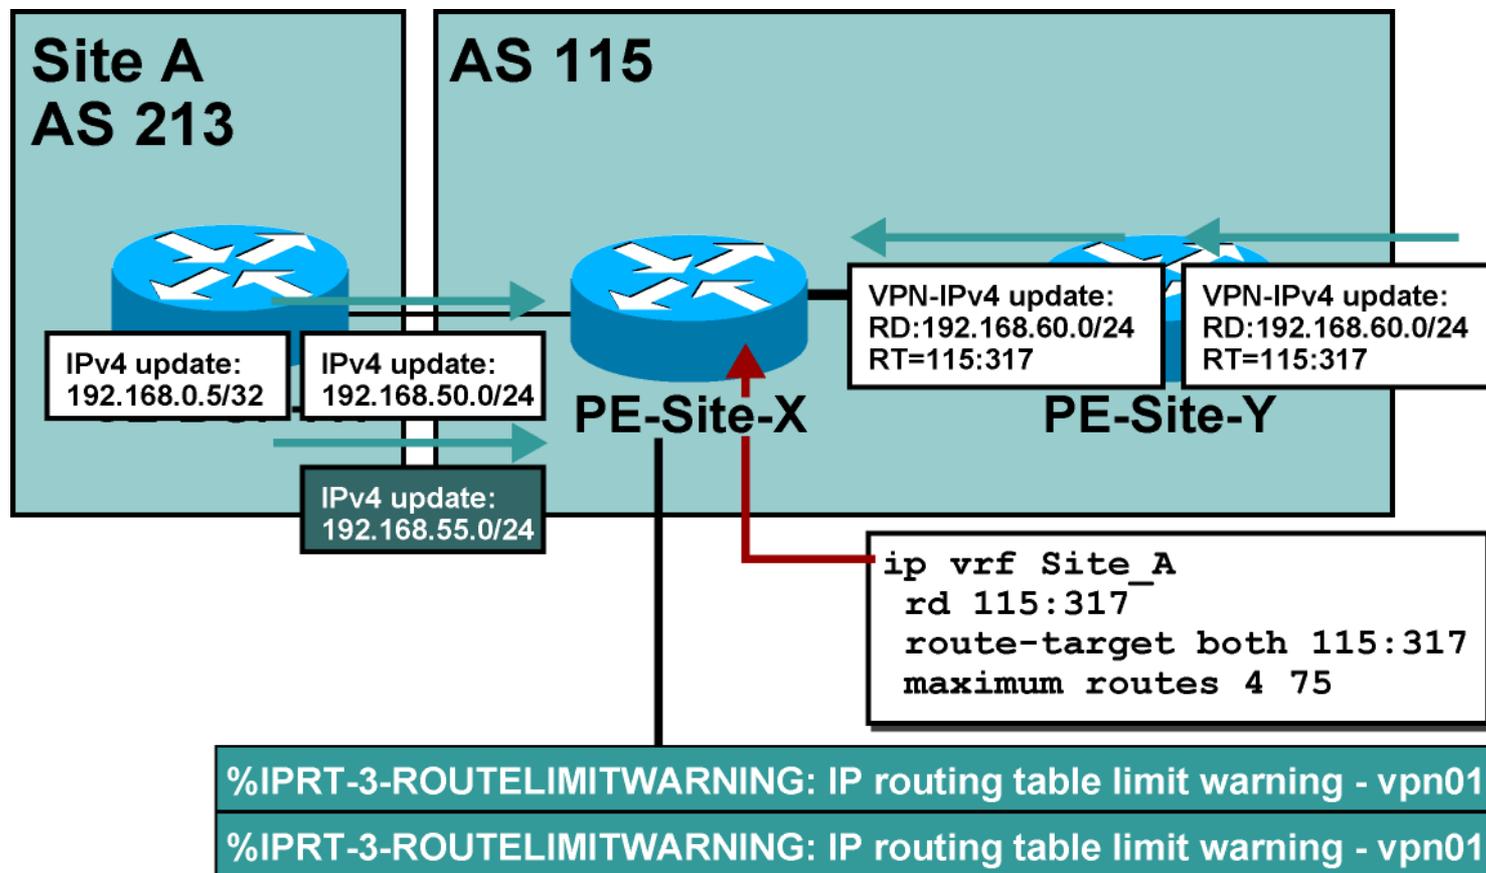
Limiting the Total Number of VRF Routes (Cont.)



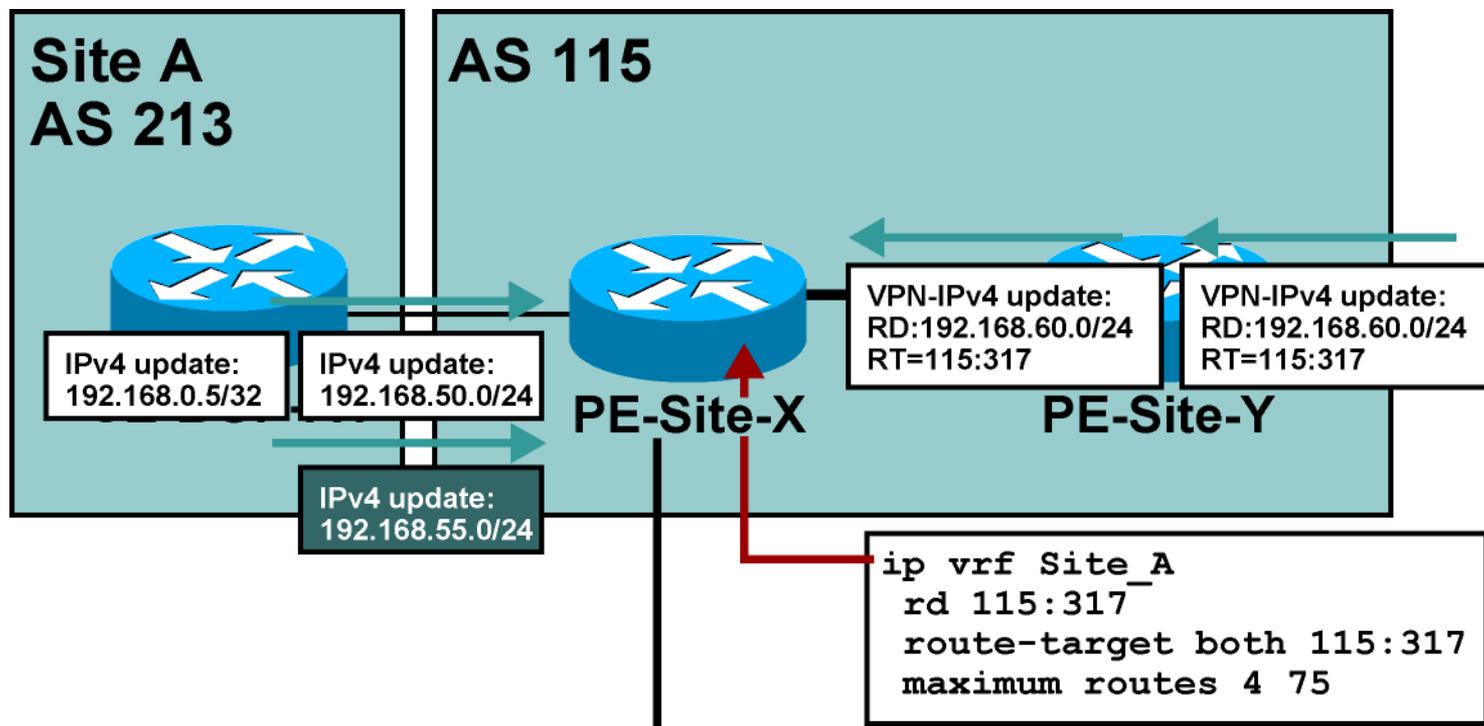
Limiting the Total Number of VRF Routes (Cont.)



Limiting the Total Number of VRF Routes (Cont.)



Limiting the Total Number of VRF Routes (Cont.)



%IPRT-3-ROUTE LIMIT WARNING: IP routing table limit warning - vpn01

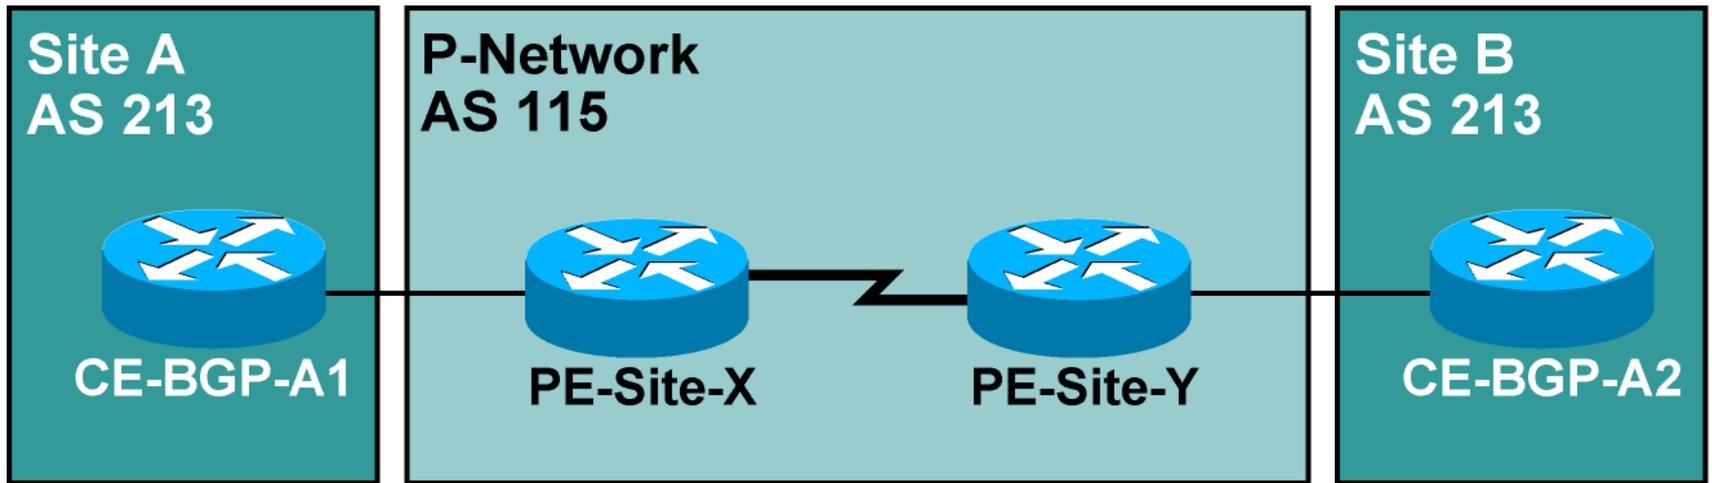
%IPRT-3-ROUTE LIMIT WARNING: IP routing table limit warning - vpn01

%IPRT-3-ROUTE LIMIT EXCEEDED: IP routing table limit exceeded -Site_A, 192.168.55.0/24

0206_375

AS-override

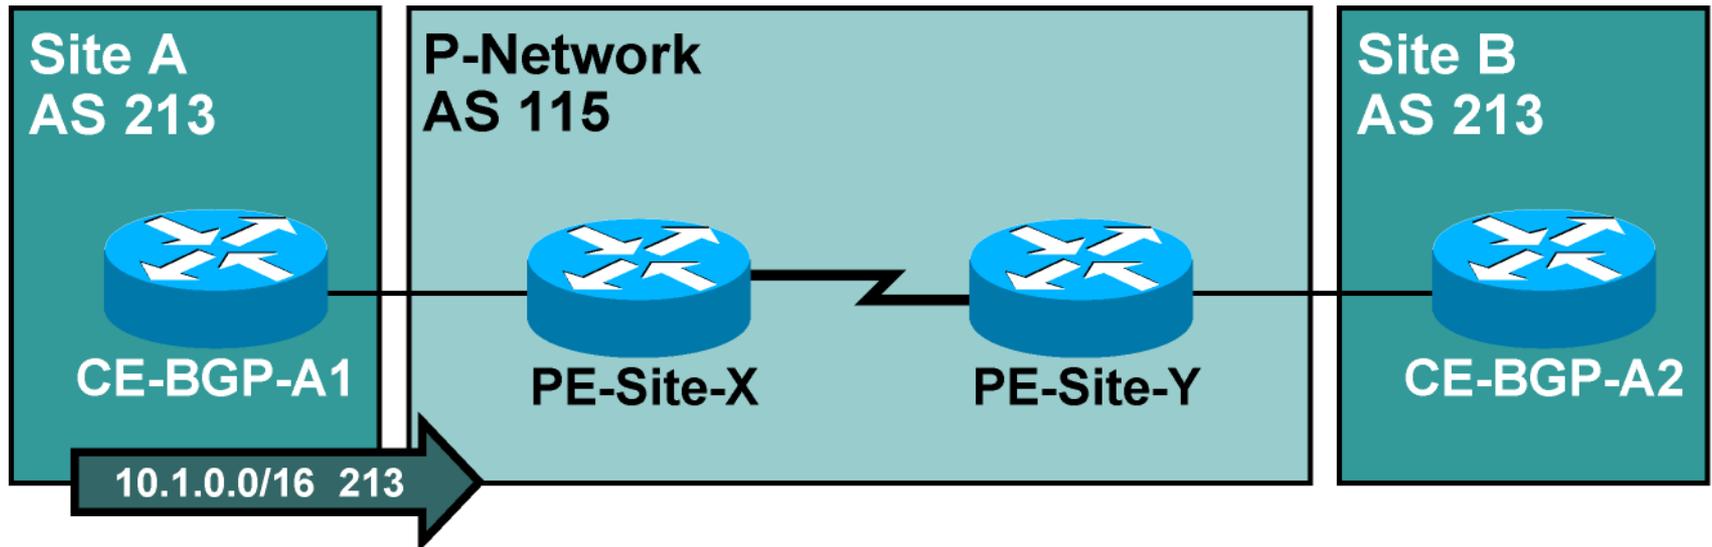
The Issue



- The customer wants to reuse the same AS number on several sites:

AS-override

The Issue

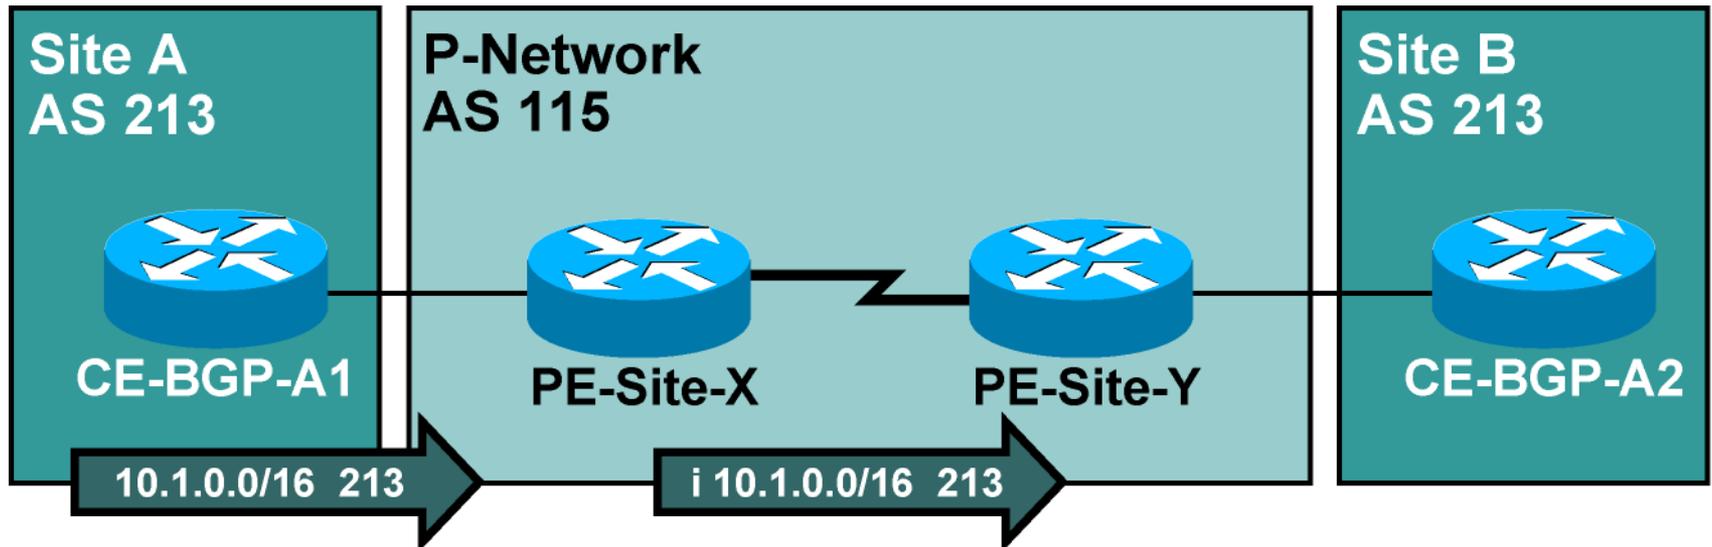


- The customer wants to reuse the same AS number on several sites:
 - CE-BGP-A1 announces network 10.1.0.0/16 to PE-Site-X.

020G_386

AS-override

The Issue

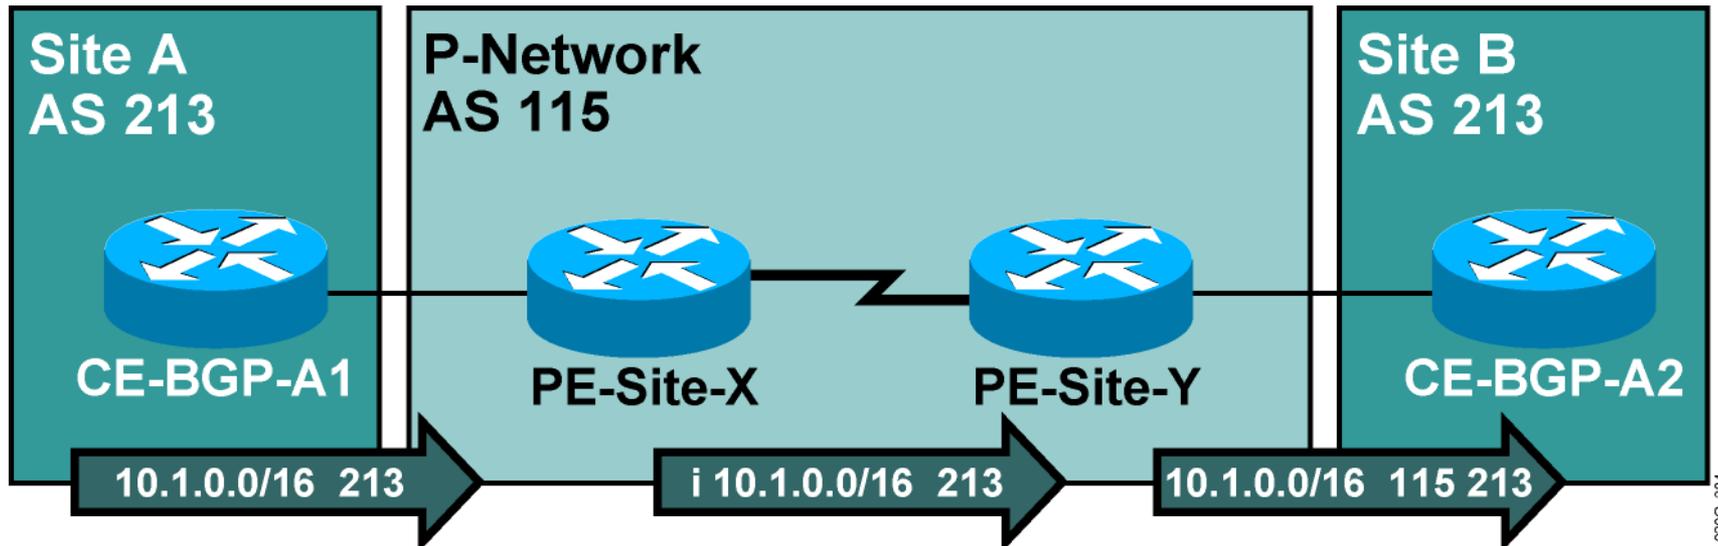


- The customer wants to reuse the same AS number on several sites:
 - CE-BGP-A1 announces network 10.1.0.0/16 to PE-Site-X.
 - The prefix announced by CE-BGP-A1 is propagated to PE-Site-Y as an internal route through MP-BGP.

020G_385

AS-override

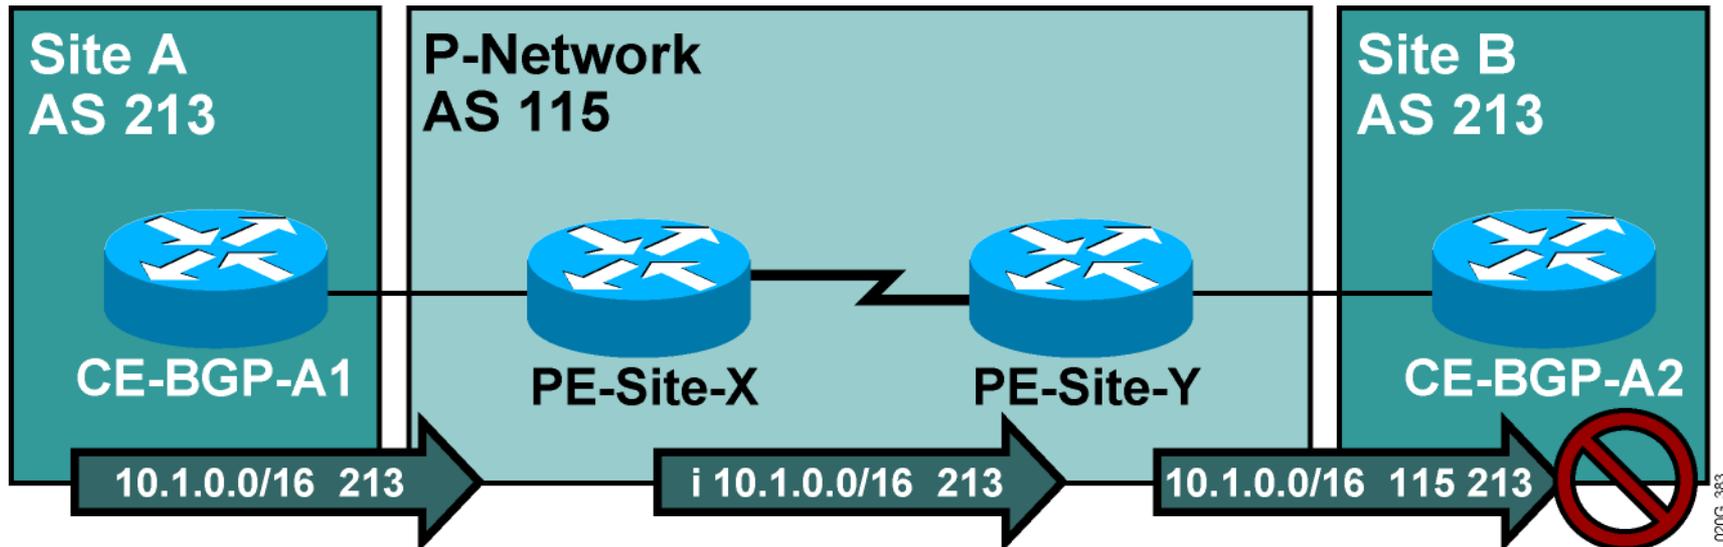
The Issue



- The customer wants to reuse the same AS number on several sites:
 - CE-BGP-A1 announces network 10.1.0.0/16 to PE-Site-X.
 - The prefix announced by CE-BGP-A1 is propagated to PE-Site-Y as an internal route through MP-BGP.
 - PE-Site-Y prepends AS 115 to the AS path and propagates the prefix to CE-BGP-A2.

AS-override

The Issue



- The customer wants to reuse the same AS number on several sites:
 - CE-BGP-A1 announces network 10.1.0.0/16 to PE-Site-X.
 - The prefix announced by CE-BGP-A1 is propagated to PE-Site-Y as an internal route through MP-BGP.
 - PE-Site-Y prepends AS 115 to the AS path and propagates the prefix to CE-BGP-A2.
 - CE-BGP-A2 drops the update because AS 213 is already in the AS path.

AS-override (Cont.)

New AS path update procedures have been implemented in order to reuse the same AS number on all VPN sites.

The procedures allow the use of private as well as public AS numbers.

The same AS number may be used for all sites.

AS-override (Cont.)

Implementation

- With AS-override configured, the AS path update procedure on the PE router is as follows:

If the first AS number in the AS path is equal to the neighbouring AS, it is replaced with the provider AS number.

If the first AS number has multiple occurrences (due to AS path prepend), all occurrences are replaced with the provider AS number.

After this operation, the provider AS number is prepended to the AS path.

AS-override (Cont.)

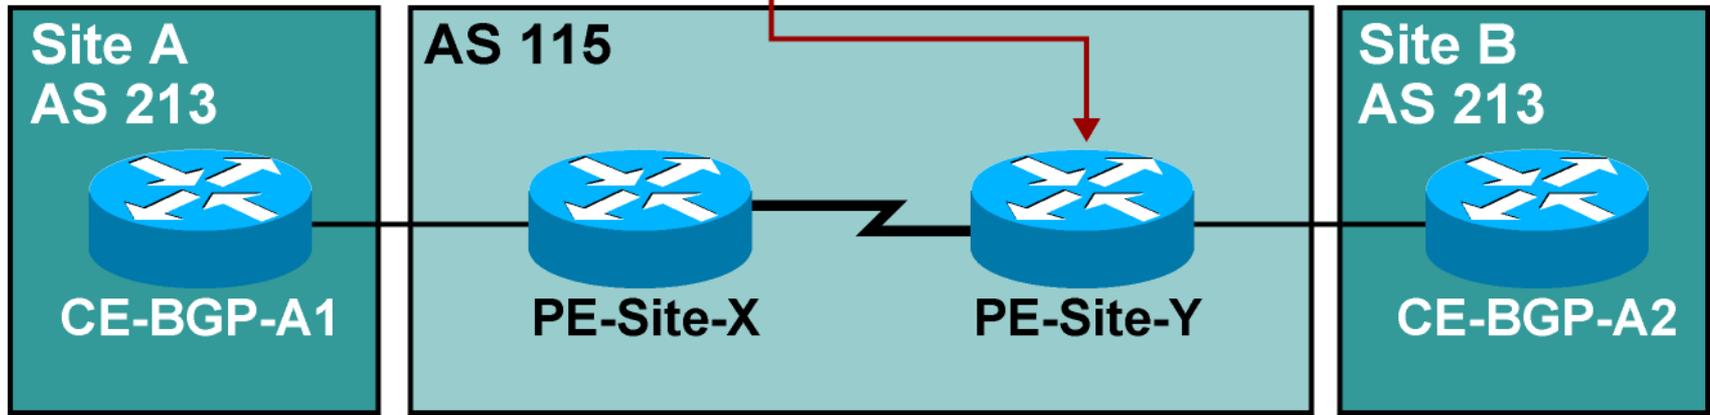
```
Router(config-router-af) #
```

```
neighbor ip-address as-override
```

- Configured on the PE router as an outbound feature
- This command configures the AS-override AS path update procedure for the specified neighbor.
- AS-override is configured for CE EBGP neighbors in the VRF address family of the BGP process.

AS-override (Cont.)

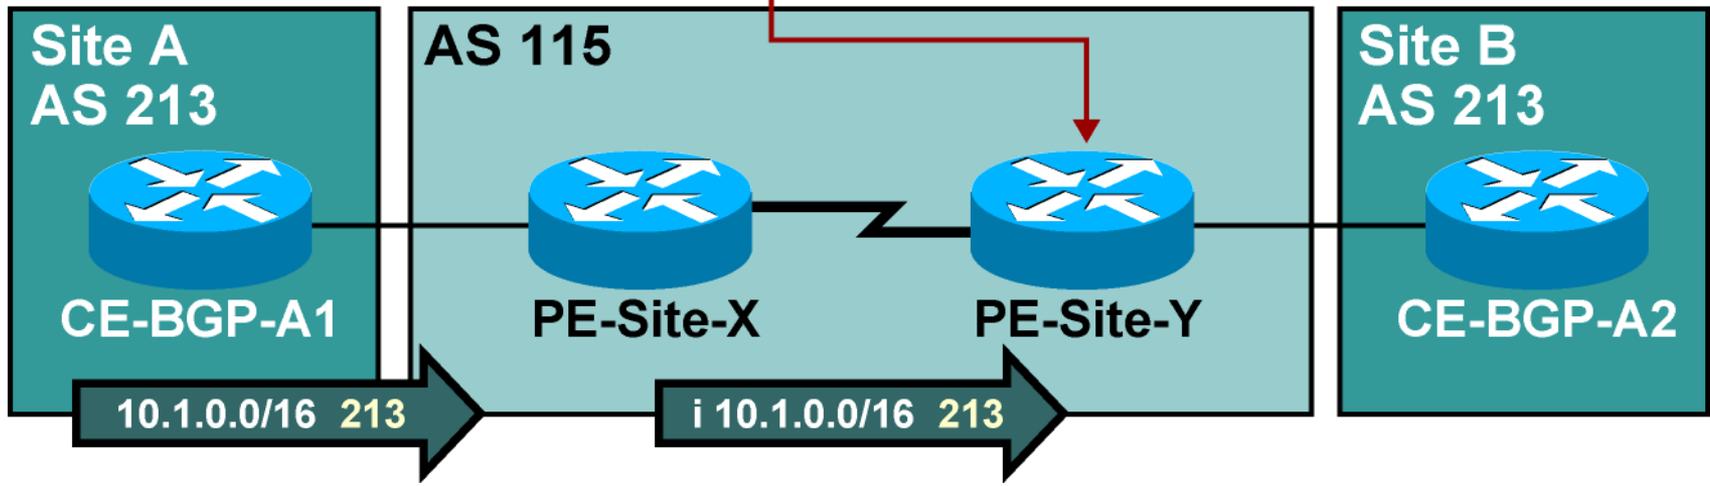
```
router bgp 115
address-family ipv4 vrf Customer_A
neighbor 10.200.2.1 remote-as 213
neighbor 10.200.2.1 activate
neighbor 10.200.2.1 as-override
```



020G_390

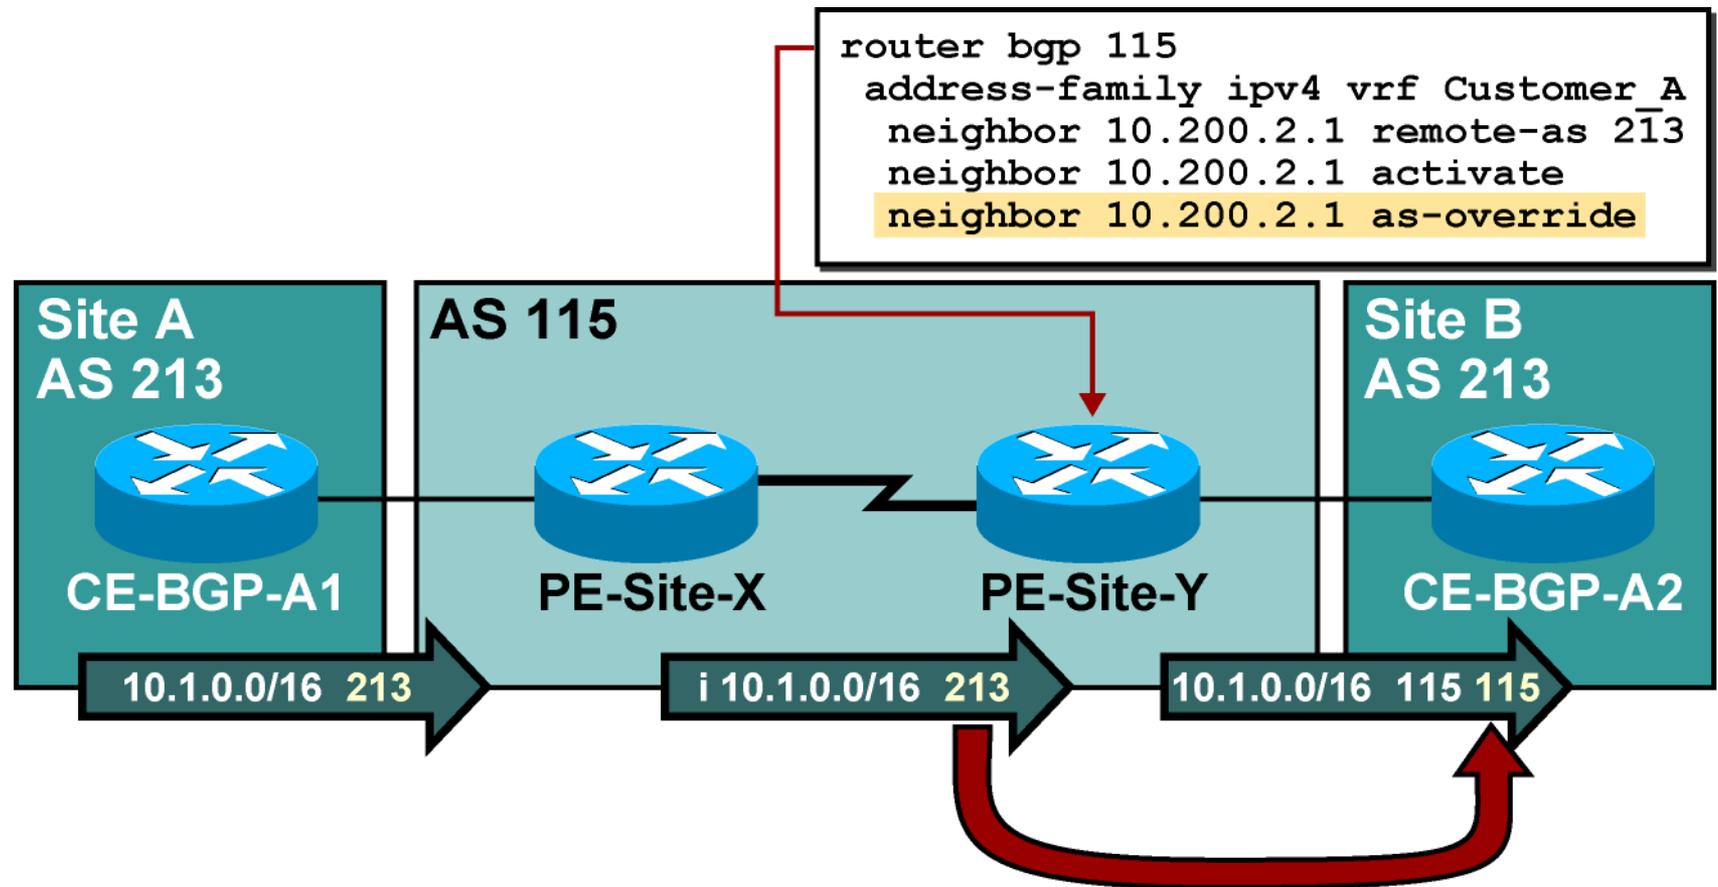
AS-override (Cont.)

```
router bgp 115
address-family ipv4 vrf Customer_A
neighbor 10.200.2.1 remote-as 213
neighbor 10.200.2.1 activate
neighbor 10.200.2.1 as-override
```



02003_389

AS-override (Cont.)

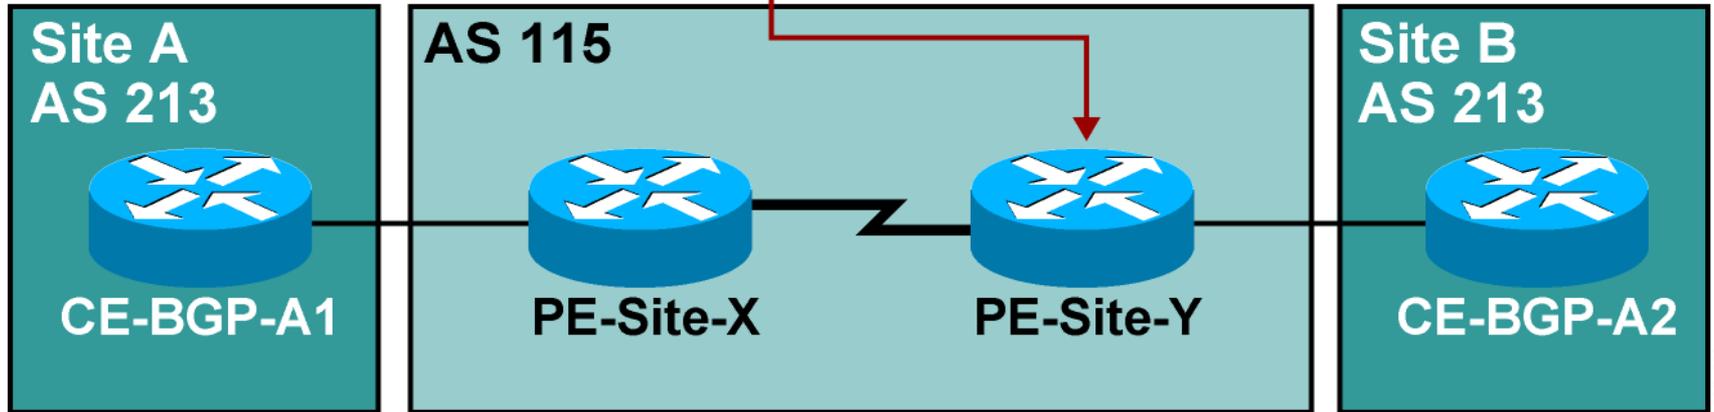


- PE-Site-Y replaces AS 213 with AS 115 in the AS path, prepends another copy of AS115 to the AS path, and propagates the prefix.

AS-override (Cont.)

AS-Path Prepending

```
router bgp 115
address-family ipv4 vrf Customer_A
neighbor 10.200.2.1 remote-as 213
neighbor 10.200.2.1 activate
neighbor 10.200.2.1 as-override
```

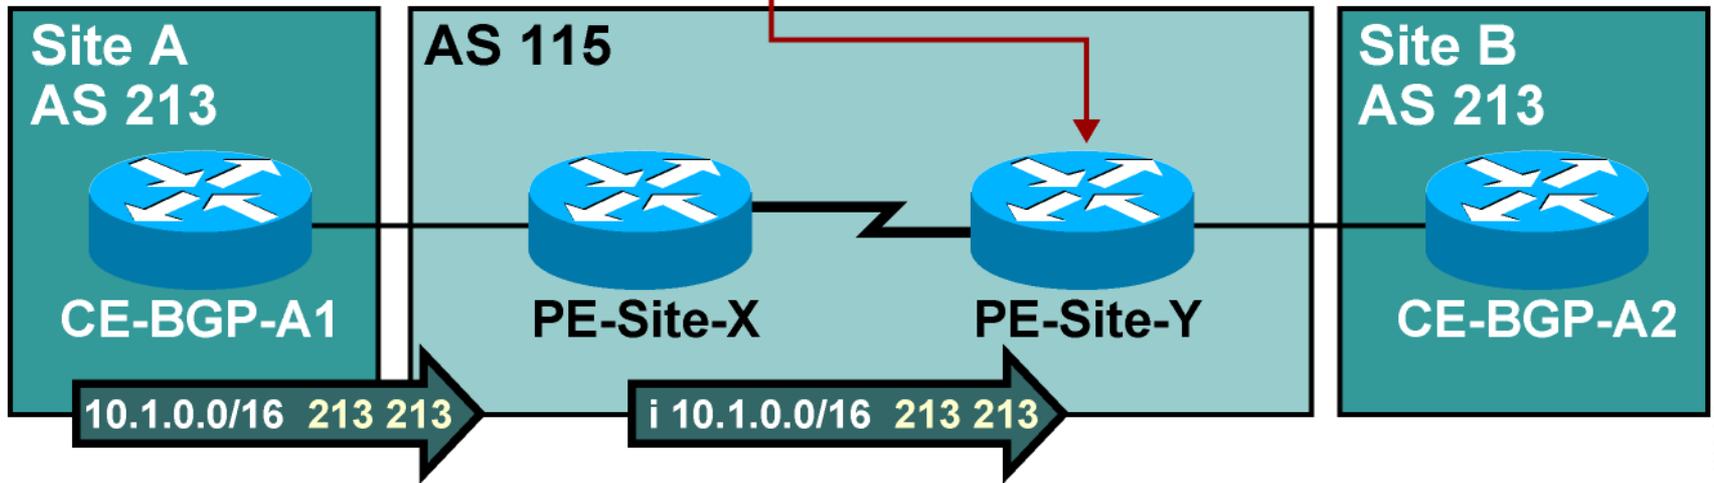


02003_3883

AS-override (Cont.)

AS-Path Prepending

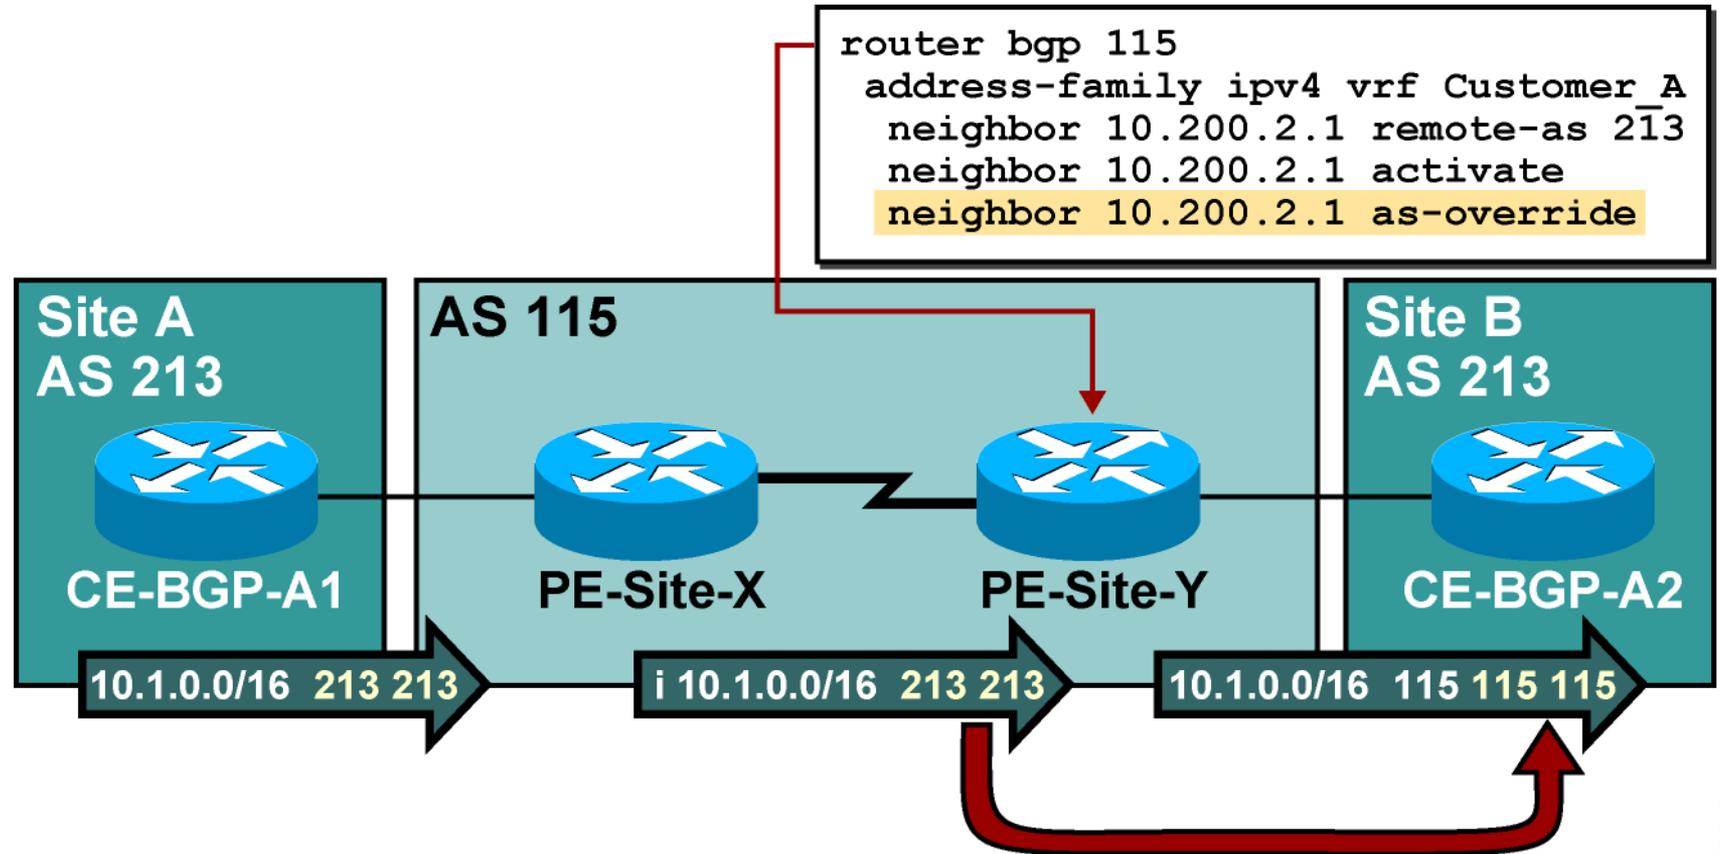
```
router bgp 115
address-family ipv4 vrf Customer_A
neighbor 10.200.2.1 remote-as 213
neighbor 10.200.2.1 activate
neighbor 10.200.2.1 as-override
```



020G_392

AS-override (Cont.)

AS-Path Prepending



- PE-Site-Y replaces all occurrences of AS 213 with AS 115 in the AS path, prepends another copy of AS 115 to the AS path, and propagates the prefix.

Hub & Spoke VPN Topology

- One central site has full routing knowledge of all other sites of the same VPN

Hub-Site

- Other sites will send traffic to the Hub-Site for any destination

Spoke-Sites

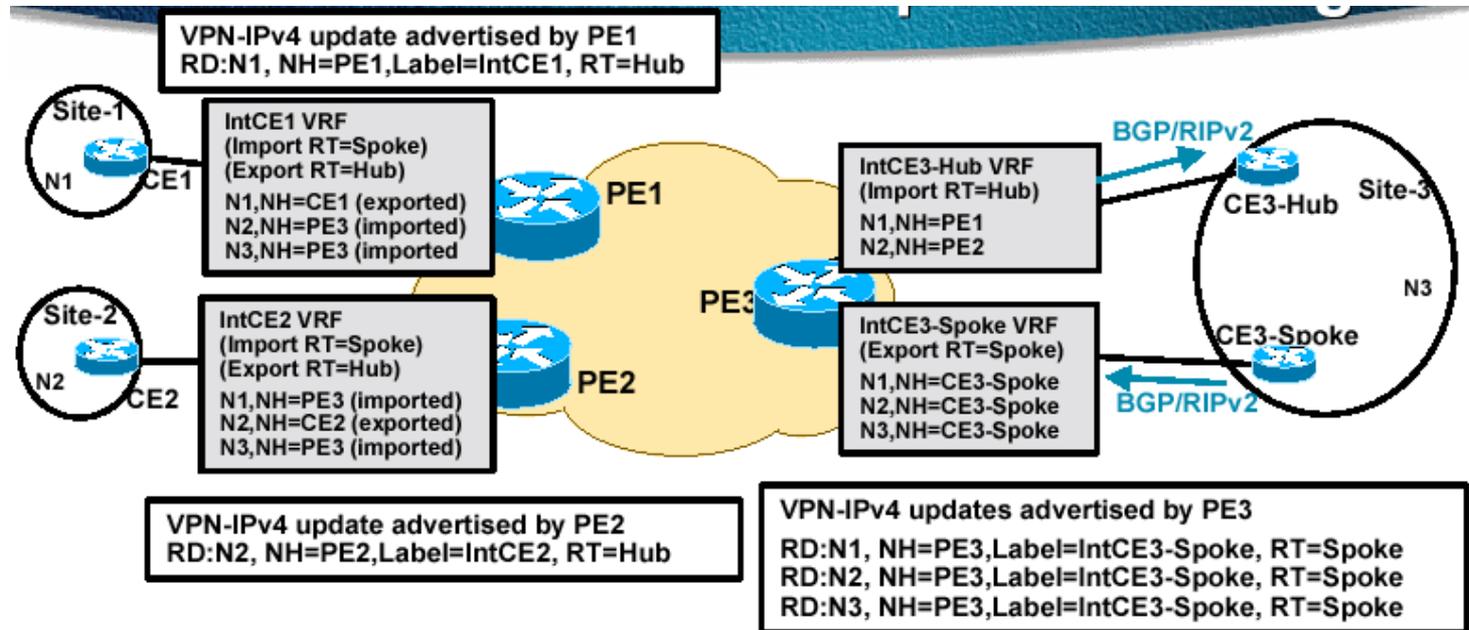
- **The Hub-Site is the central transit point between Spoke-Sites**

Security services (filters)

Traffic logging and/or accounting

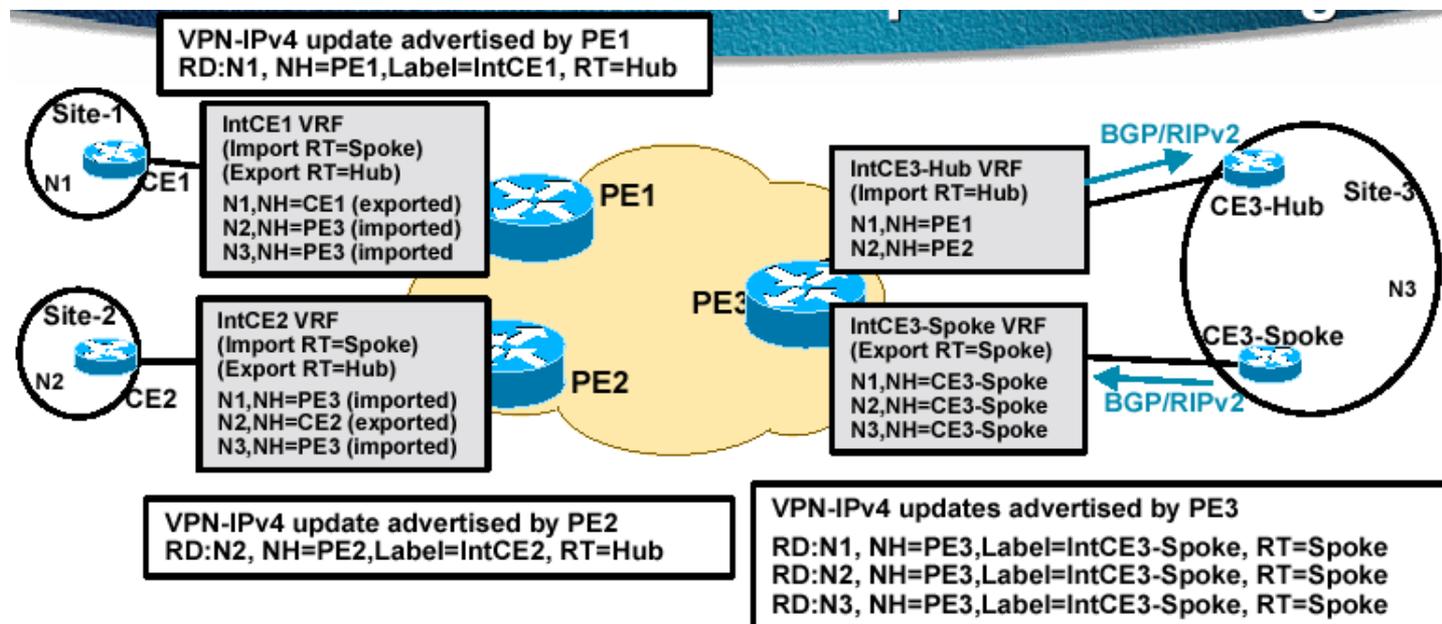
Intrusion Detection systems

VPN Sites with Hub & Spoke Routing



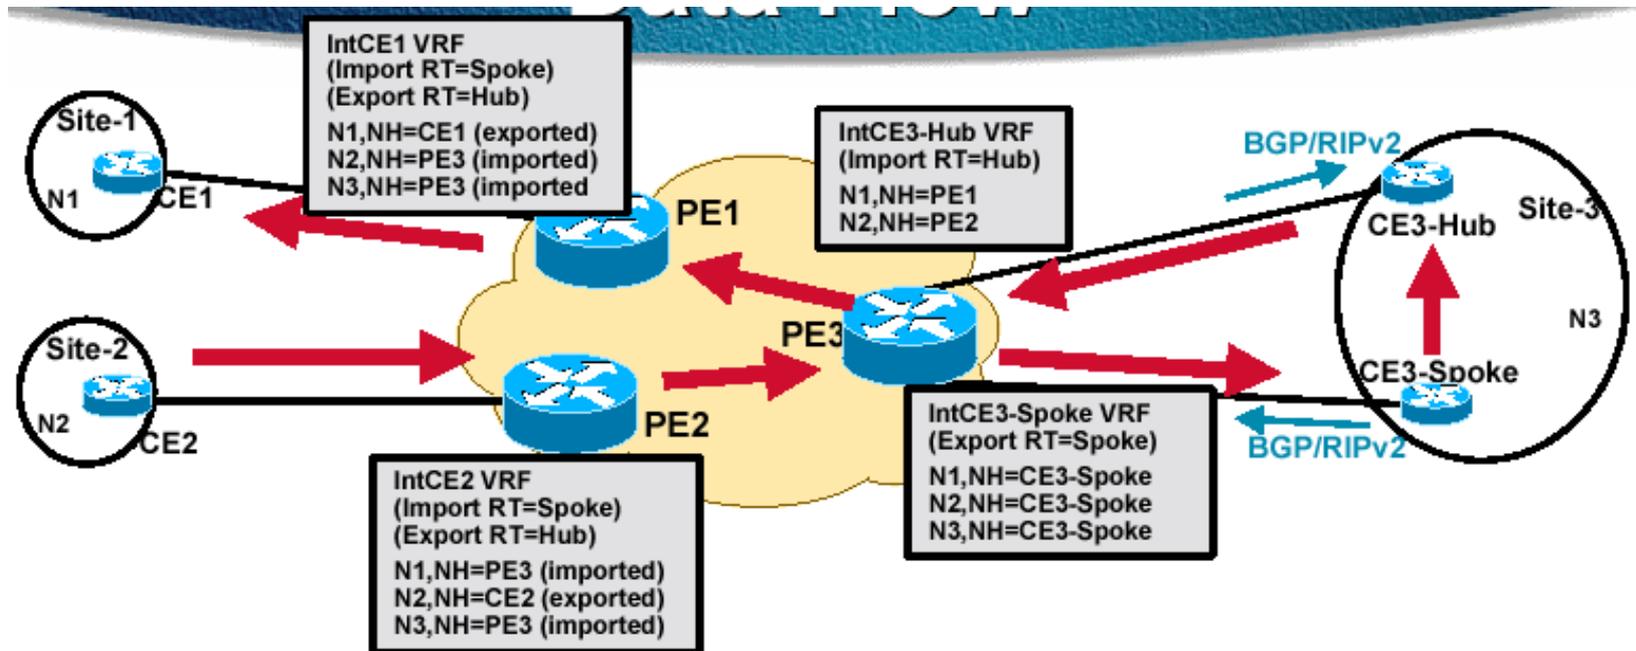
- We need 2 interfaces and 2 unique VRFs on the Hub site. If not, traffic from spokes may just touch PE3 and be forwarded to the spoke site without being processed at the hub site
- Traffic into hub comes in via one VRF (which exports routes, e.g. Spoke vrf) interface and goes out via the other (which imports routes, e.g. Hub vrf).

VPN Sites with Hub & Spoke Routing



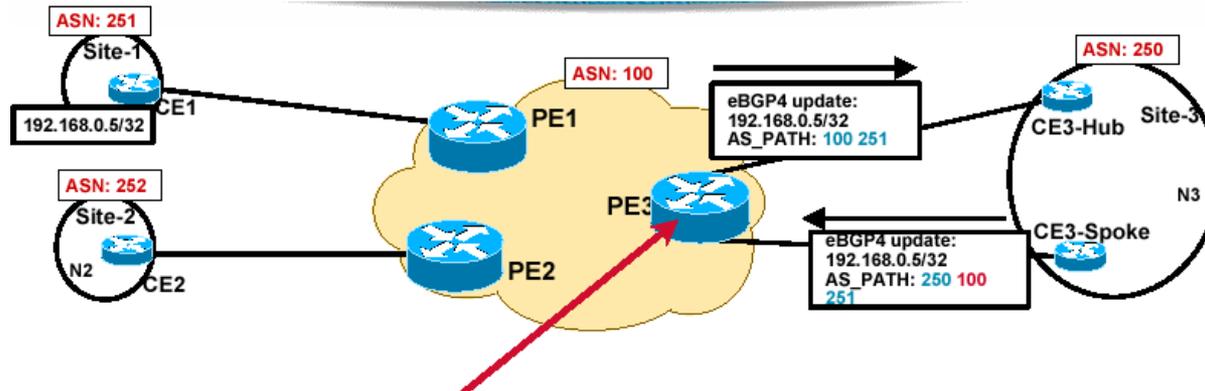
- Spoke routes are imported into Hub VRF on PE-3 from Site-1 and Site-2
- The same Spoke routes are exported to other spokes via the Spoke VRF, in which the next-hop for spoke sites to reach the other spoke site is PE3
- Since Spoke VRF at hub site exports the spoke routes, all the traffic from spokes destined to other spokes will come to this spoke VRF
- Traffic from the spoke VRF will be forwarded to the actual destination spoke via the Hub VRF routing

Hub & Spoke Topology Data Flow



- Traffic from one spoke to another will travel across the hub site
- Allowas-in has to be configured on the PE3 if the Site-3 is using BGP

Allows-in (for BGP updates)



```
router bgp 100
address-family ipv4 vrf Spoke
neighbor 192.168.74.4 remote-as 250
neighbor 192.168.74.4 activate
neighbor 192.168.74.4 allows-in 4
no auto-summary
no synchronization
exit-address-family
```

Allowas-in (Cont.)

- The allowas-in BGP option disables the AS path check on the PE router:

The number of occurrences of the PE router AS number is limited to suppress real routing loops.

The limit has to be configured.

The PE router will **reject** the update only if its AS number appears in the AS path more often than the configured limit.

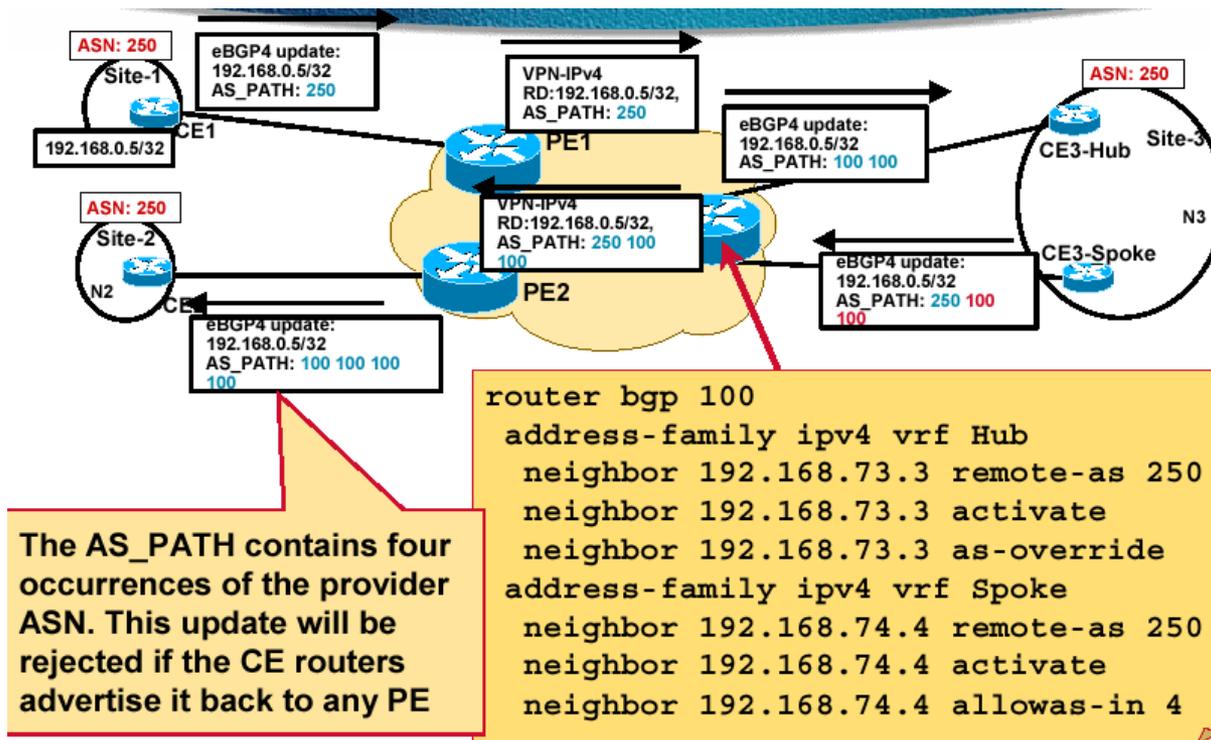
Allowas-in (Cont.)

```
Router(config-router) #
```

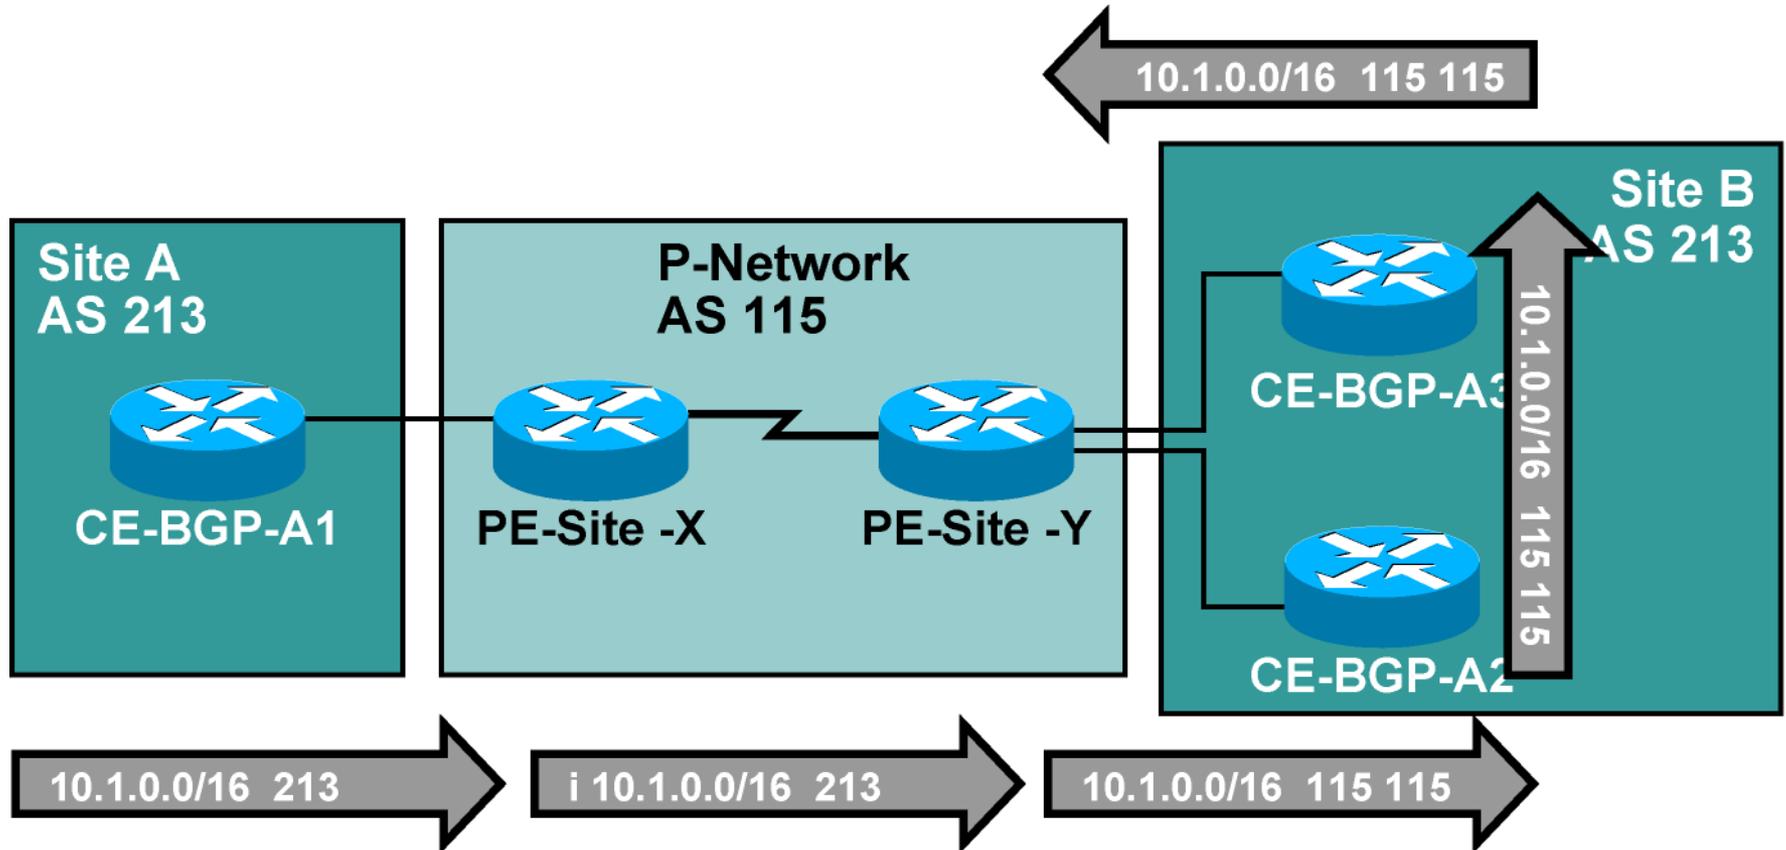
```
neighbor allowas-in number
```

- This command disables the traditional BGP AS path check.
- An incoming update is rejected only if the AS number of the PE router appears in the AS path more often than the configured limit.

Allowas-in in Combination with AS-override



Implementing SOO for Loop Prevention



- AS path-based BGP loop prevention is bypassed with AS-override and allowas-in features.

020G_403

Implementing SOO for Loop Prevention (Cont.)

SOO identifies the Site from which PE router learns a route

The SOO (extended BGP community) can be used to prevent loops in these scenarios.

The SOO is needed only for multihomed sites.

When EBGP is run between PE and CE routers, the SOO is configured through a route map command on a per neighbour basis under address-family ipv4 vrf

For other routing protocols, the SOO can be applied to routes learned through a particular VRF interface

Implementing SOO for Loop Prevention (Cont.)

The same Site of Origin attribute must be used for all CE routers that are at the same site, whether or not those CE routers are attached to the same PE.

Distinct Site of Origin attributes must be used for CE routers, which are at distinct sites.

Note that a route must be associated with at most one attribute of this type.

Implementing SOO for Loop Prevention (Cont.)

Inbound EBGP Update

```
Router(config)#
```

```
route-map name permit seq  
  match conditions  
  set extcommunity soo extended-community-value
```

- Creates a route map that sets the SOO attribute

```
Router(config-router-af)#
```

```
neighbor ip-address route-map name in
```

- Applies inbound route map to CE EBGP neighbor
- Configuring inbound SOO also prevents the PE router from sending any routes outbound on this interface with the same SoO as the one set in the route-map

Implementing SOO for Loop Prevention (Cont.)

Other Inbound Routing Updates

```
Router(config)#
```

```
route-map name permit seq  
  match conditions  
  set extcommunity soo extended-community-value
```

- Creates a route map that sets the SOO attribute

```
Router(config-if)#
```

```
ip vrf sitemap route-map-name
```

- Applies route map that sets SOO to inbound routing updates received from this interface

Selective VRF import/export

- Selective import:

Specify additional criteria for importing routes into the VRF.

- Selective export:

Specify additional RTs attached to exported routes.

Configuring Selective VRF Import

- VRF import criteria might be more specific than just the match on the RT—for example:

Import only routes with specific BGP attributes (community, and so on).

Import routes with specific prefixes or subnet masks (only loopback addresses).

- A route map can be configured in a VRF to make route import more specific.

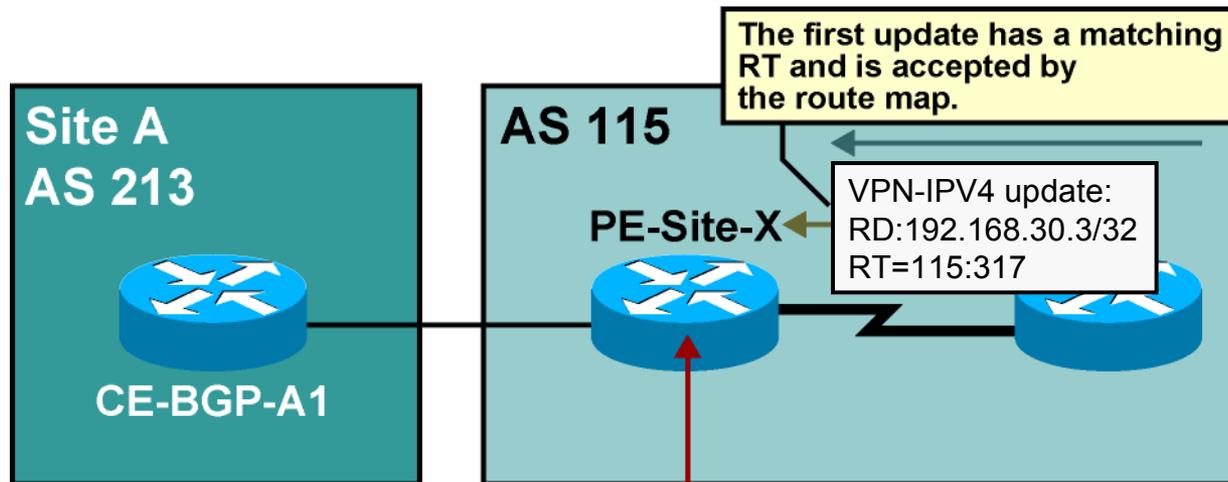
Configuring Selective VRF Import (Cont.)

```
Router(config-vrf) #
```

```
import map route-map
```

- This command attaches a route map to the VRF import process.
- A route is imported into the VRF only if at least one RT attached to the route matches one RT configured in the VRF **and** the route is accepted by the route map.

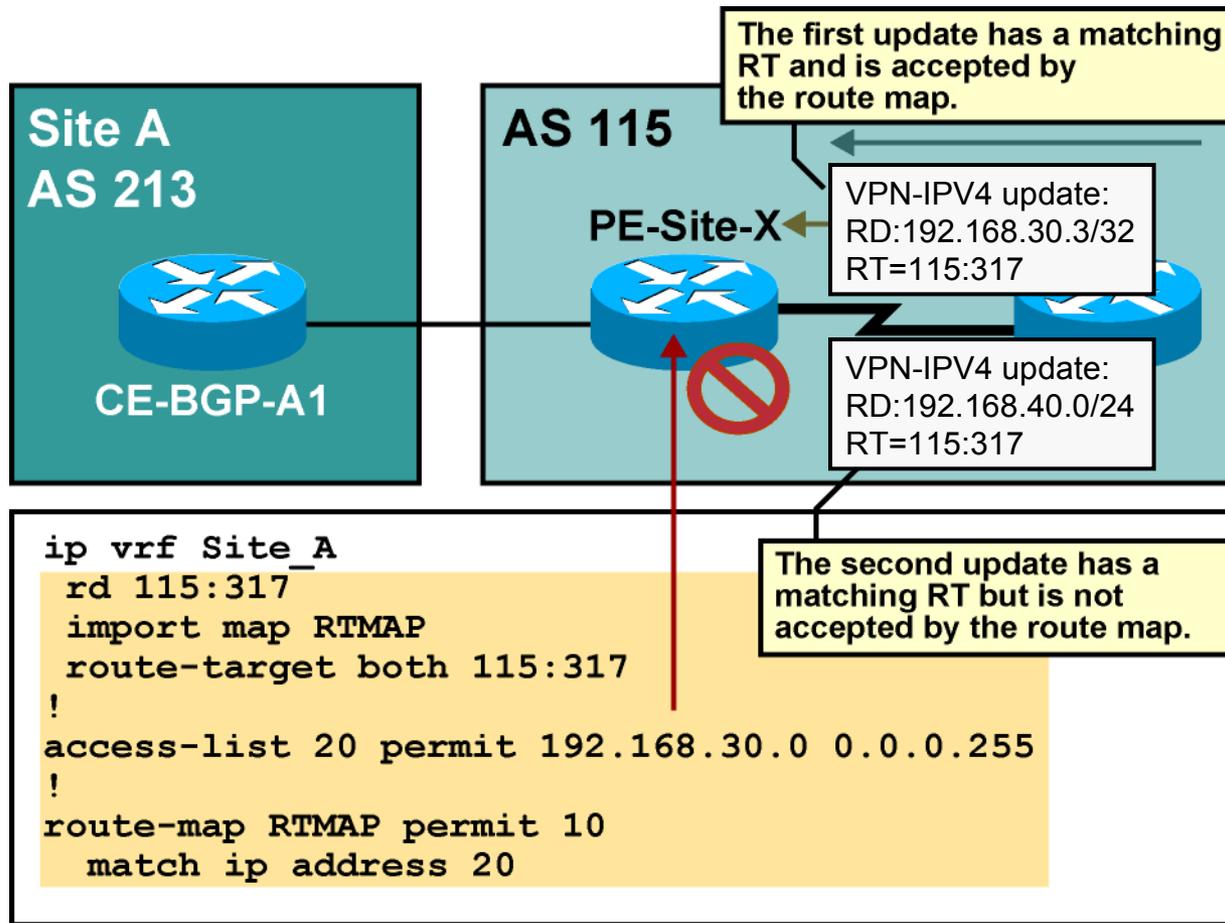
Configuring Selective VRF Import (Cont.)



```
ip vrf Site_A
 rd 115:317
 import map RTMAP
 route-target both 115:317
 !
 access-list 20 permit 192.168.30.0 0.0.0.255
 !
 route-map RTMAP permit 10
  match ip address 20
```

02003_084

Configuring Selective VRF Import (Cont.)



020G_085

Configuring Selective VRF Export

- Routes from a VRF might have to be exported with different RTs:

An example would be export management routes with particular RTs.

- An export route map can be configured on VRF:

This route map can set extended community RTs.

No other set operations can be performed by this route map.

Configuring Selective VRF Export (Cont.)

Router(config)#

```
route-map name permit seq  
  match condition  
  set extcommunity rt extended-community-value [additive]
```

- This command creates a route map that matches routes based on any route map conditions, and sets RTs.

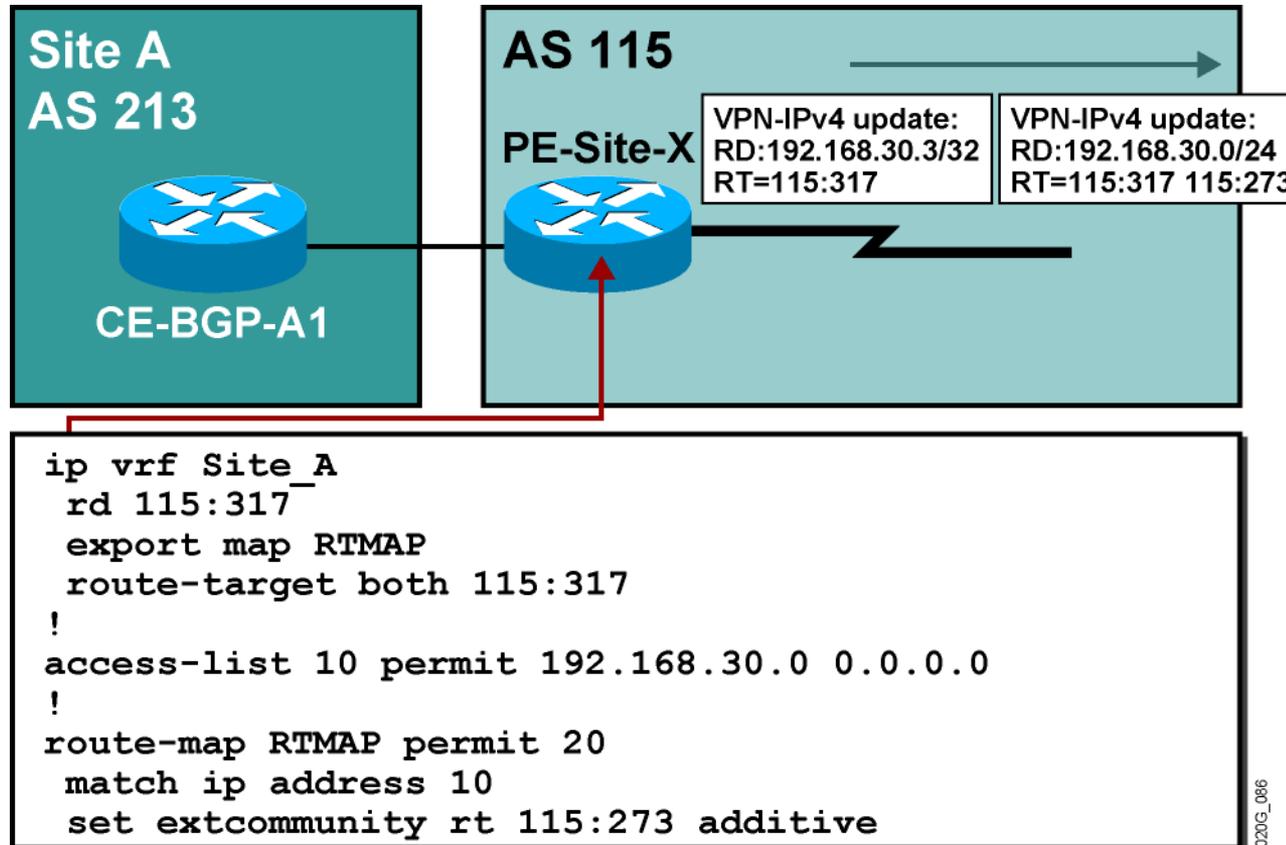
Configuring Selective VRF Export (Cont.)

```
router(config-vrf) #
```

```
export map name
```

- This command attaches a route map to the VRF export process.
- All exported routes always get RTs configured with route-target export in the VRF.
- A route that is matched by the export route map will have additional RTs attached.

Configuring Selective VRF Export (Cont.)



02003_086

Summary

PE-CE routing protocols need to be configured for individual VRFs

Per-VRF routing protocols are configured as individual *address families* belonging to the same routing process

An AS number can be reused using:

- As-override

- Allowas-in

The SOO can be used to provide protection from routing loops.

Route import and export within VRFs can be controlled with import and export route maps.

