MPLS Workshop
Day 3

Jan 15th to Jan 19th 2009

Abdul Rahim , arahim@cisco.com

MPLS Bootcamp

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

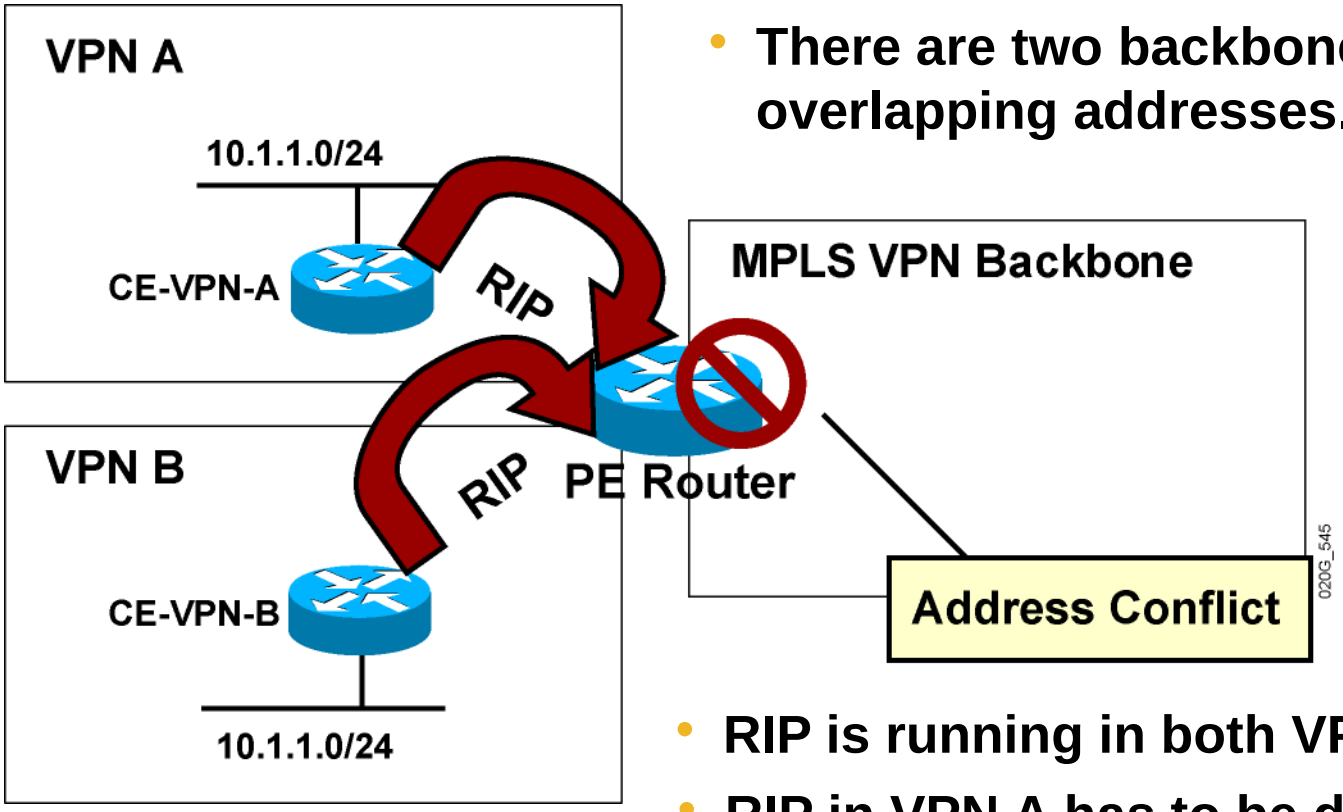

- There are two backbones with overlapping addresses.

- 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

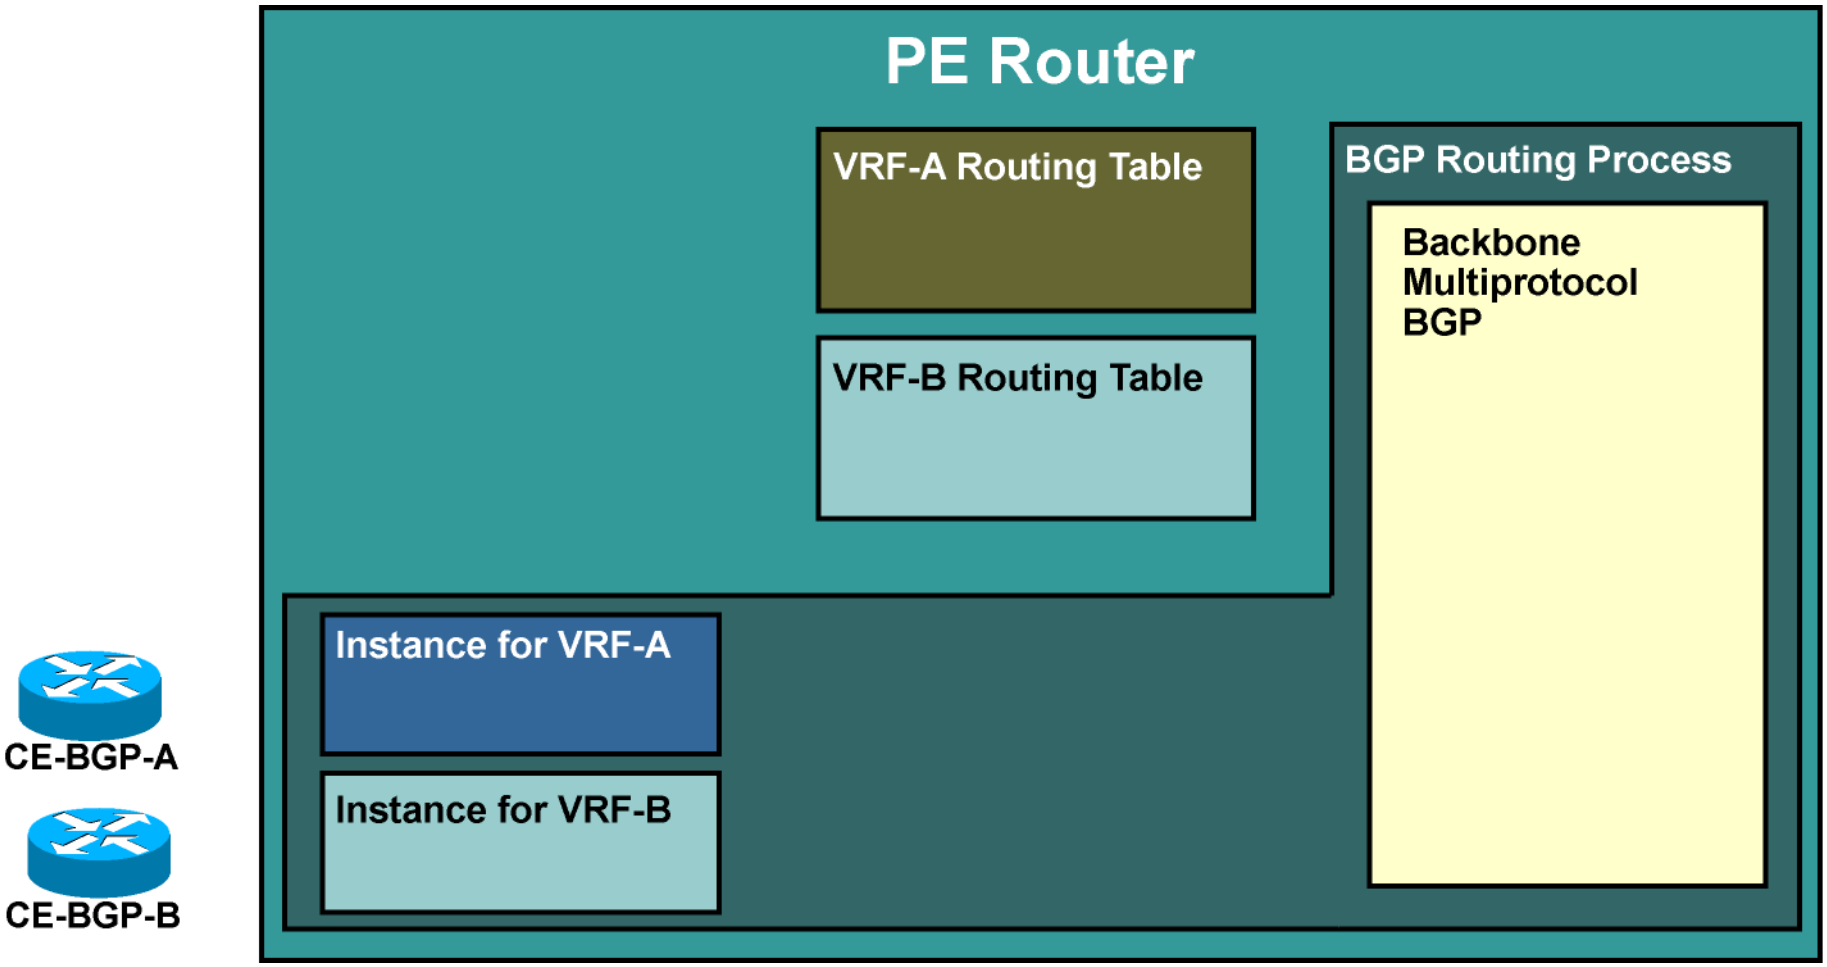

- Two VPNs are attached to the same PE router.
- Each VPN is represented by a VRF.

BGP Route Propagation—Outbound (Cont.)

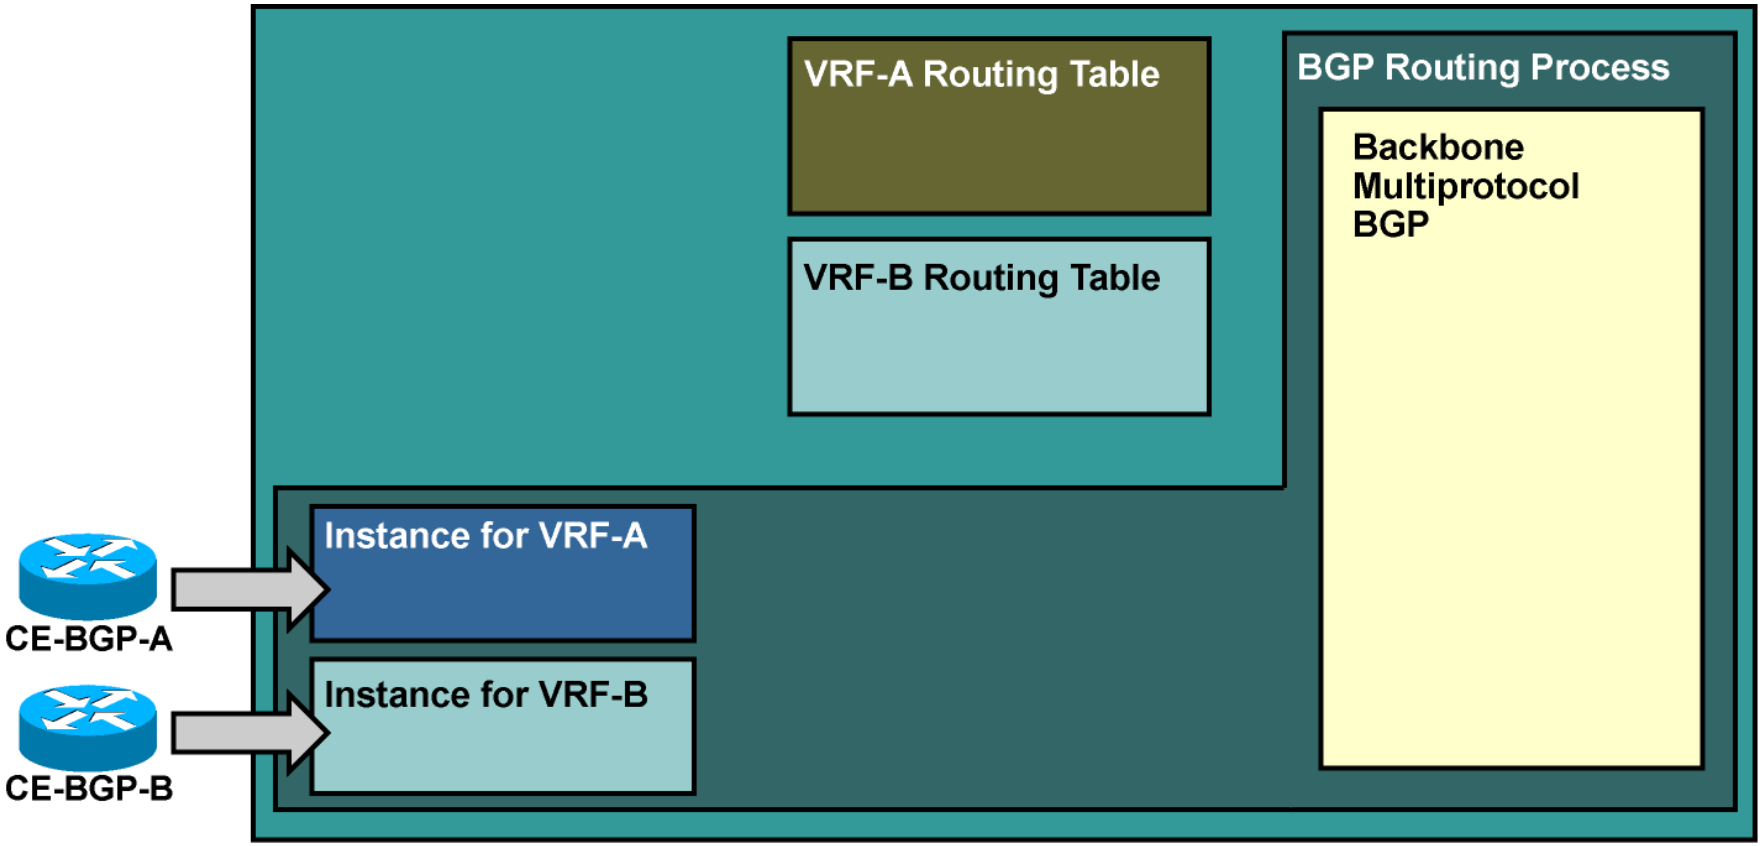

- BGP-speaking CE routers announce their pref xes to the PE router via BGP.

BGP Route Propagation—Outbound (Cont.)

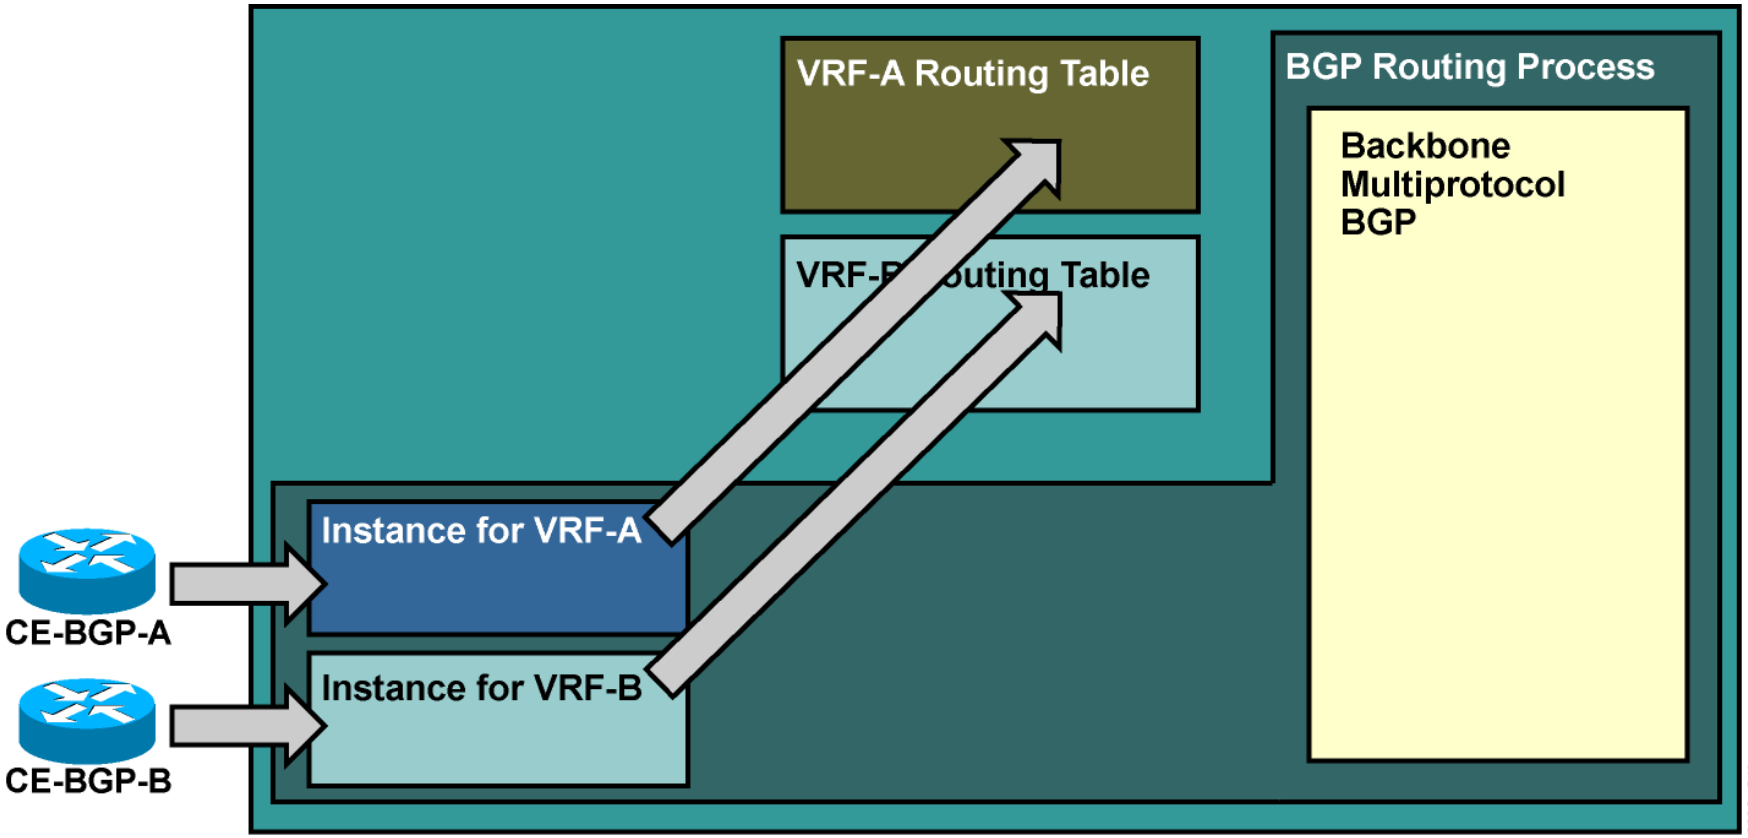

- BGP-speaking CE routers announce their pref xes to the PE router via BGP.
- Instance of BGP process associated with the VRF to which the PE-CE interface belongs collects the routes and inserts them into VRF routing table.

BGP Route Propagation—Outbound (Cont.)

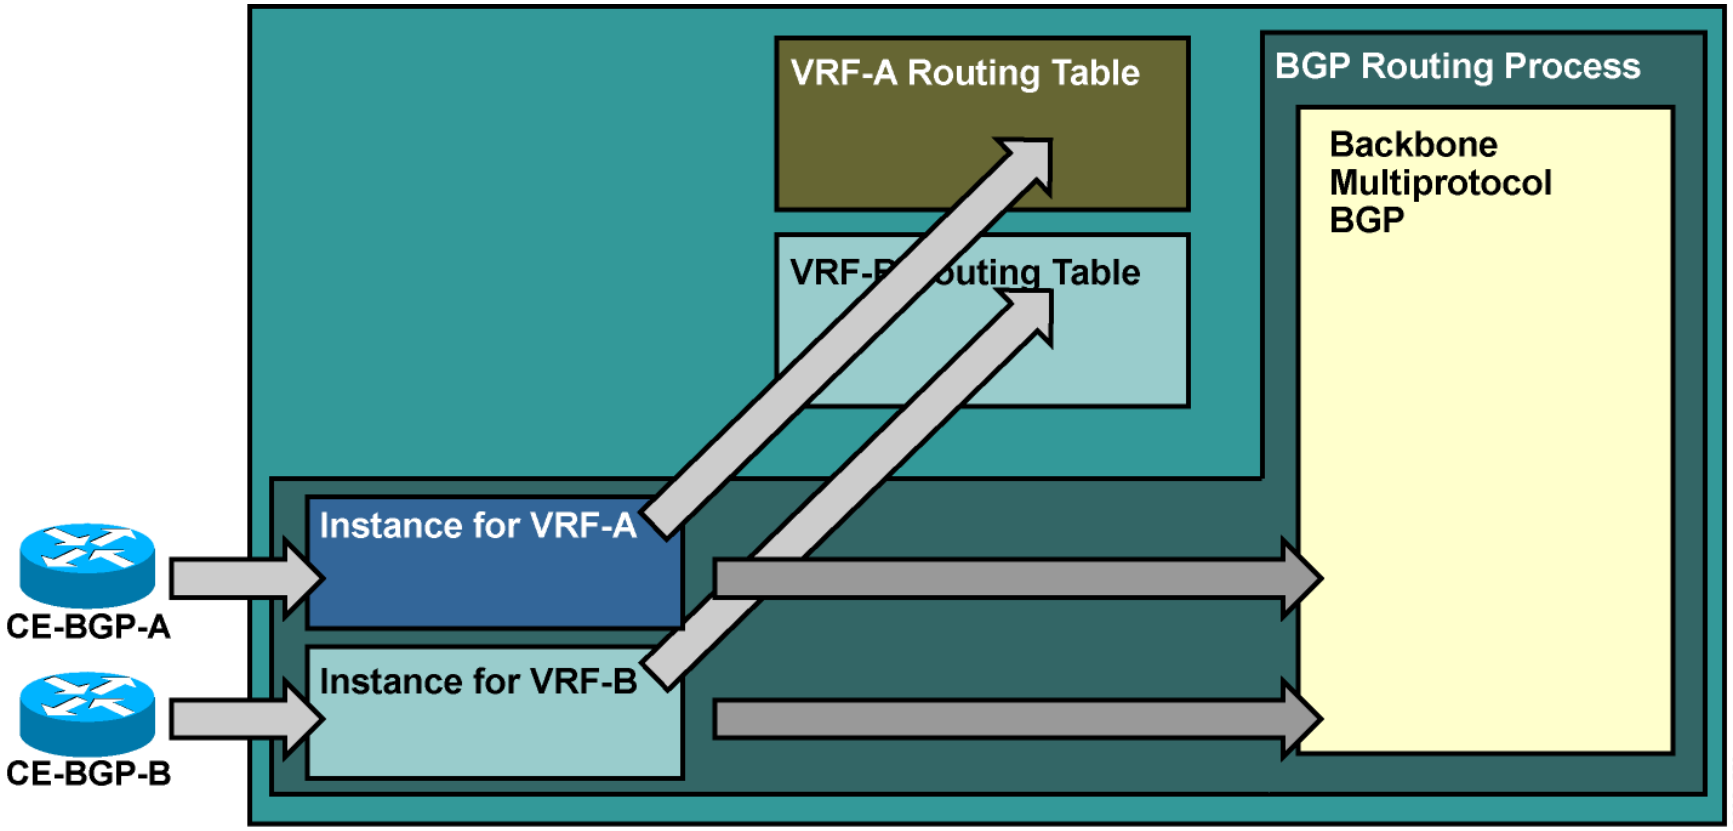

- **Route distinguisher is prepended during route export to the BGP routes from VRF instance of BGP process to convert them into VPNv4 pref xes. Route targets are attached to these pref xes.**

BGP Route Propagation—Outbound (Cont.)

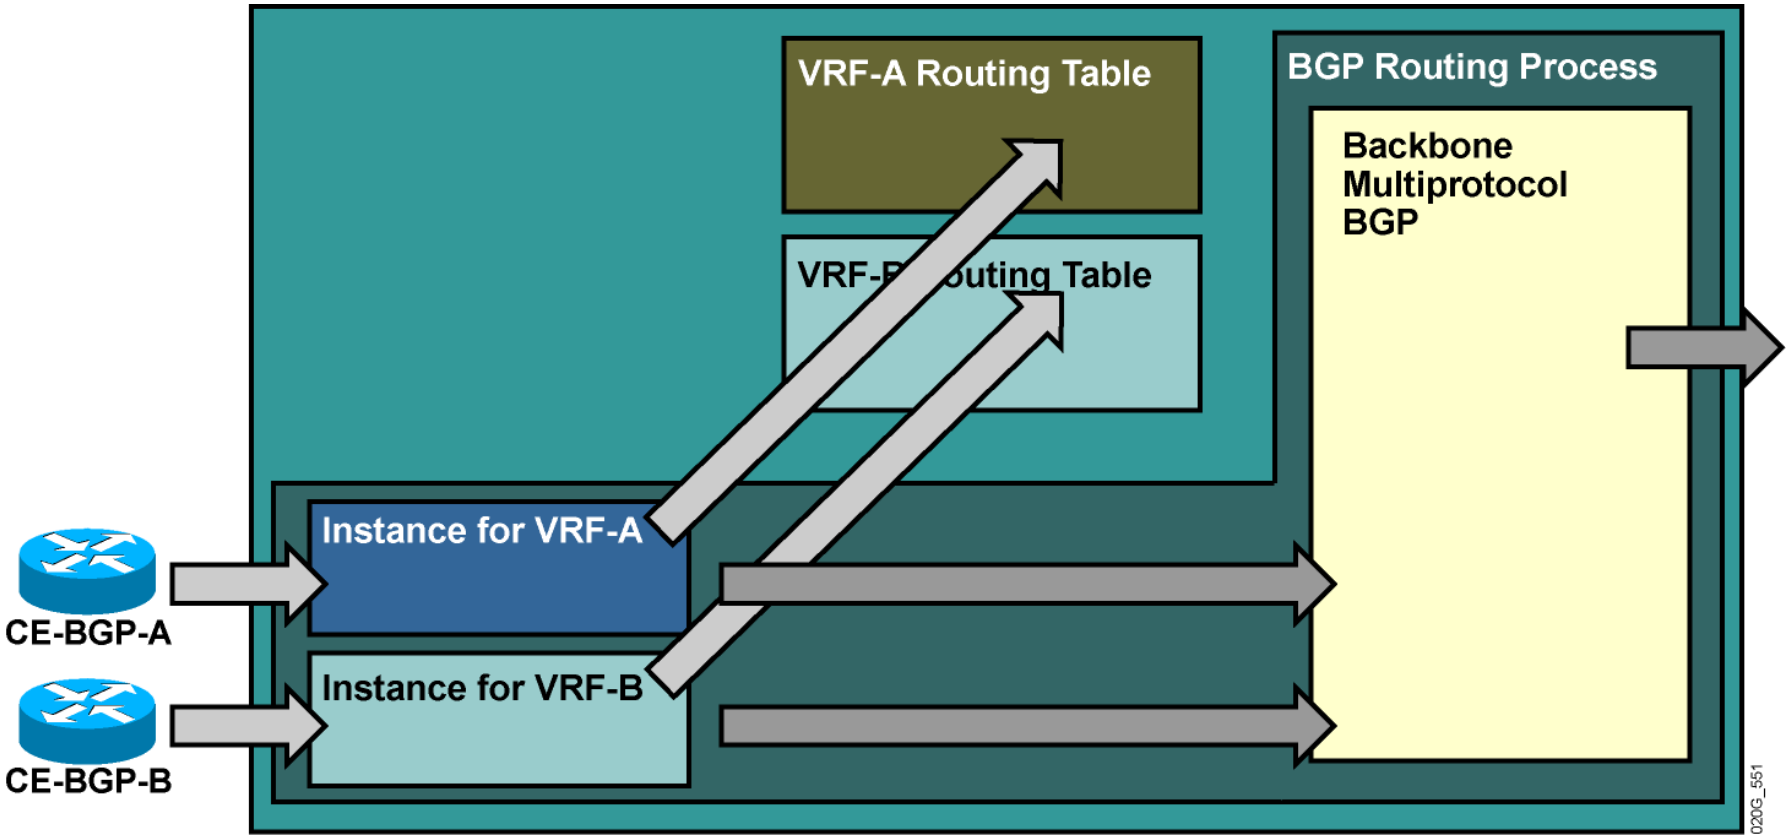

- Route distinguisher is prepended during route export to the BGP routes from VRF instance of BGP process to convert them into VPNv4 pref xes. Route targets are attached to these pref xes.
- VPNv4 pref xes are propagated to other PE routers.

Non-BGP Route propagation - Outbound

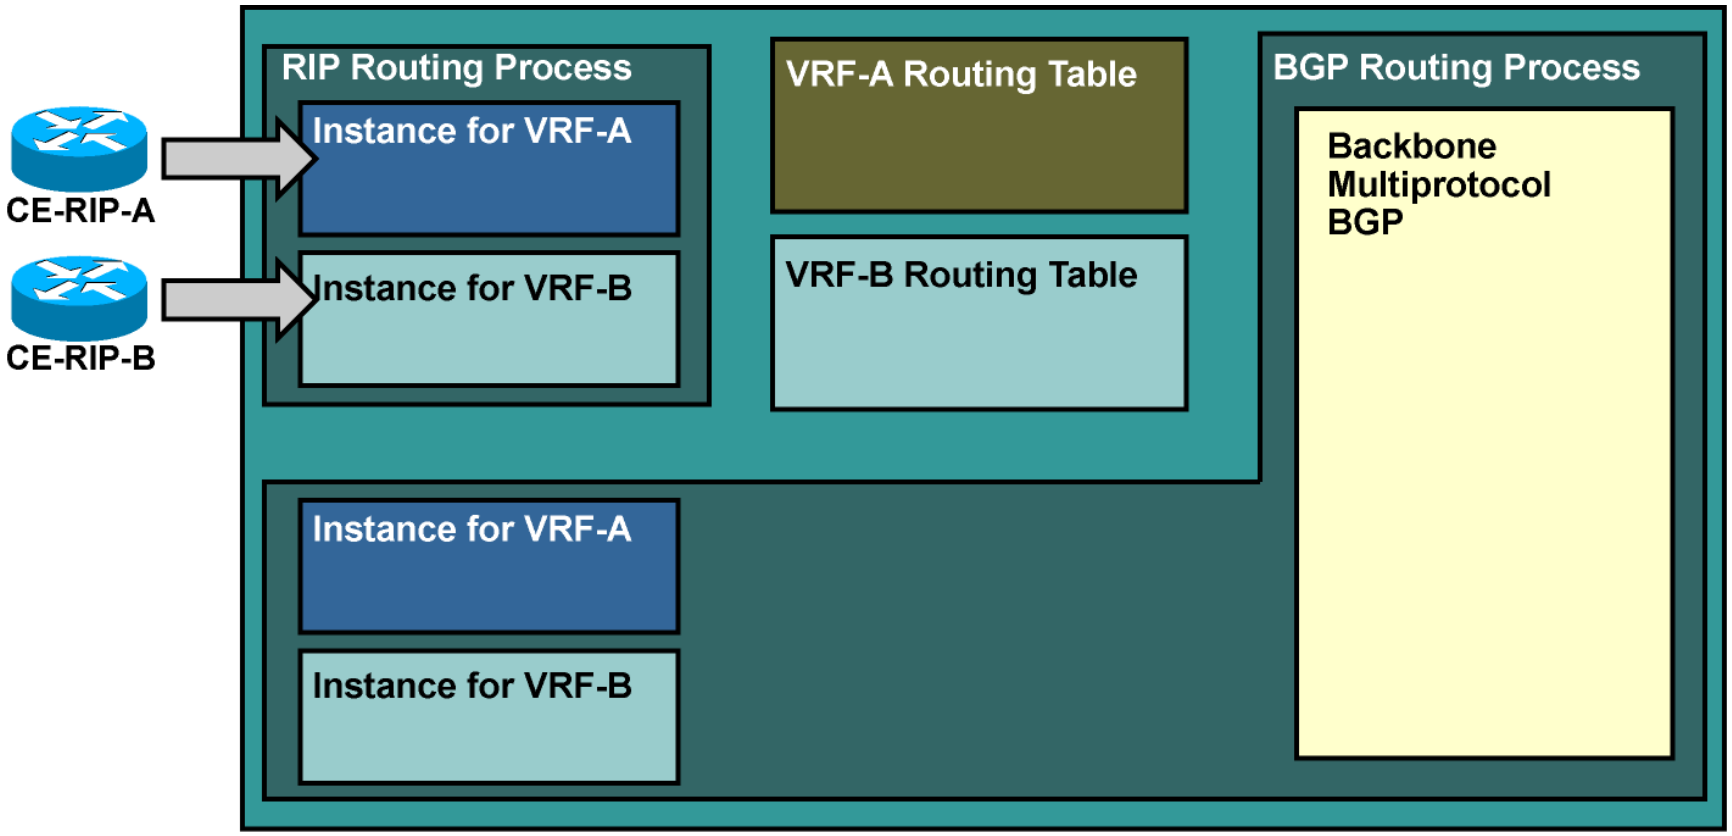

- RIP-speaking CE routers announce their prefixes to the PE router via RIP.

Non-BGP Route propagation—Outbound

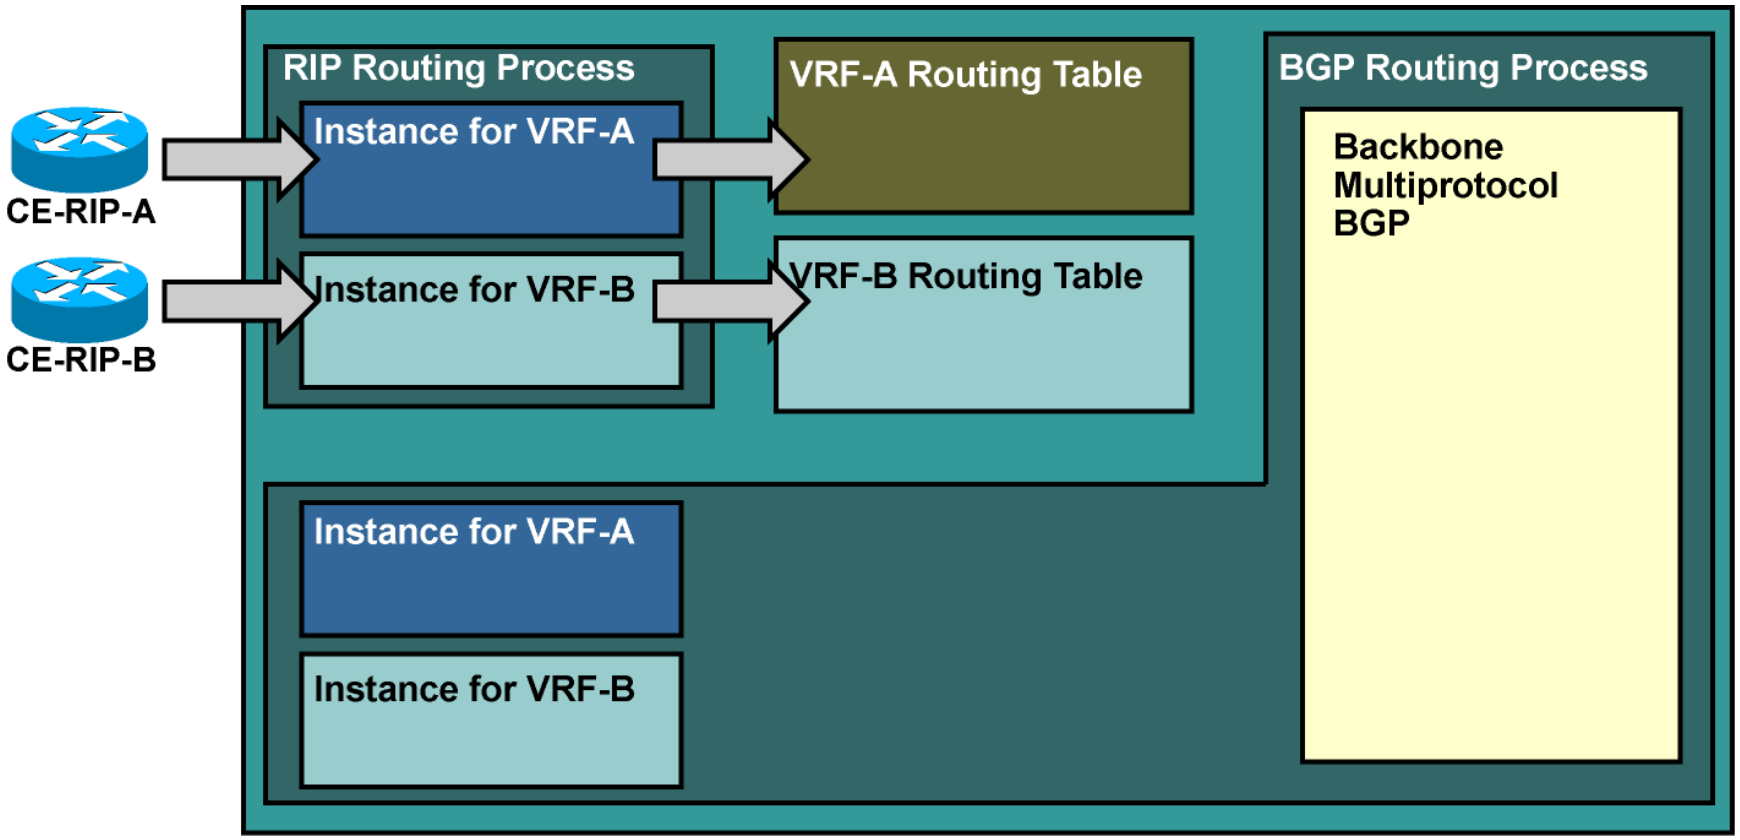

- RIP-speaking CE routers announce their prefixes to the PE router via RIP.
- Instance of RIP process associated with the VRF to which the PE-CE interface belongs collects the routes and inserts them into VRF routing table.

Non-BGP Route propagation—Outbound (Cont.)

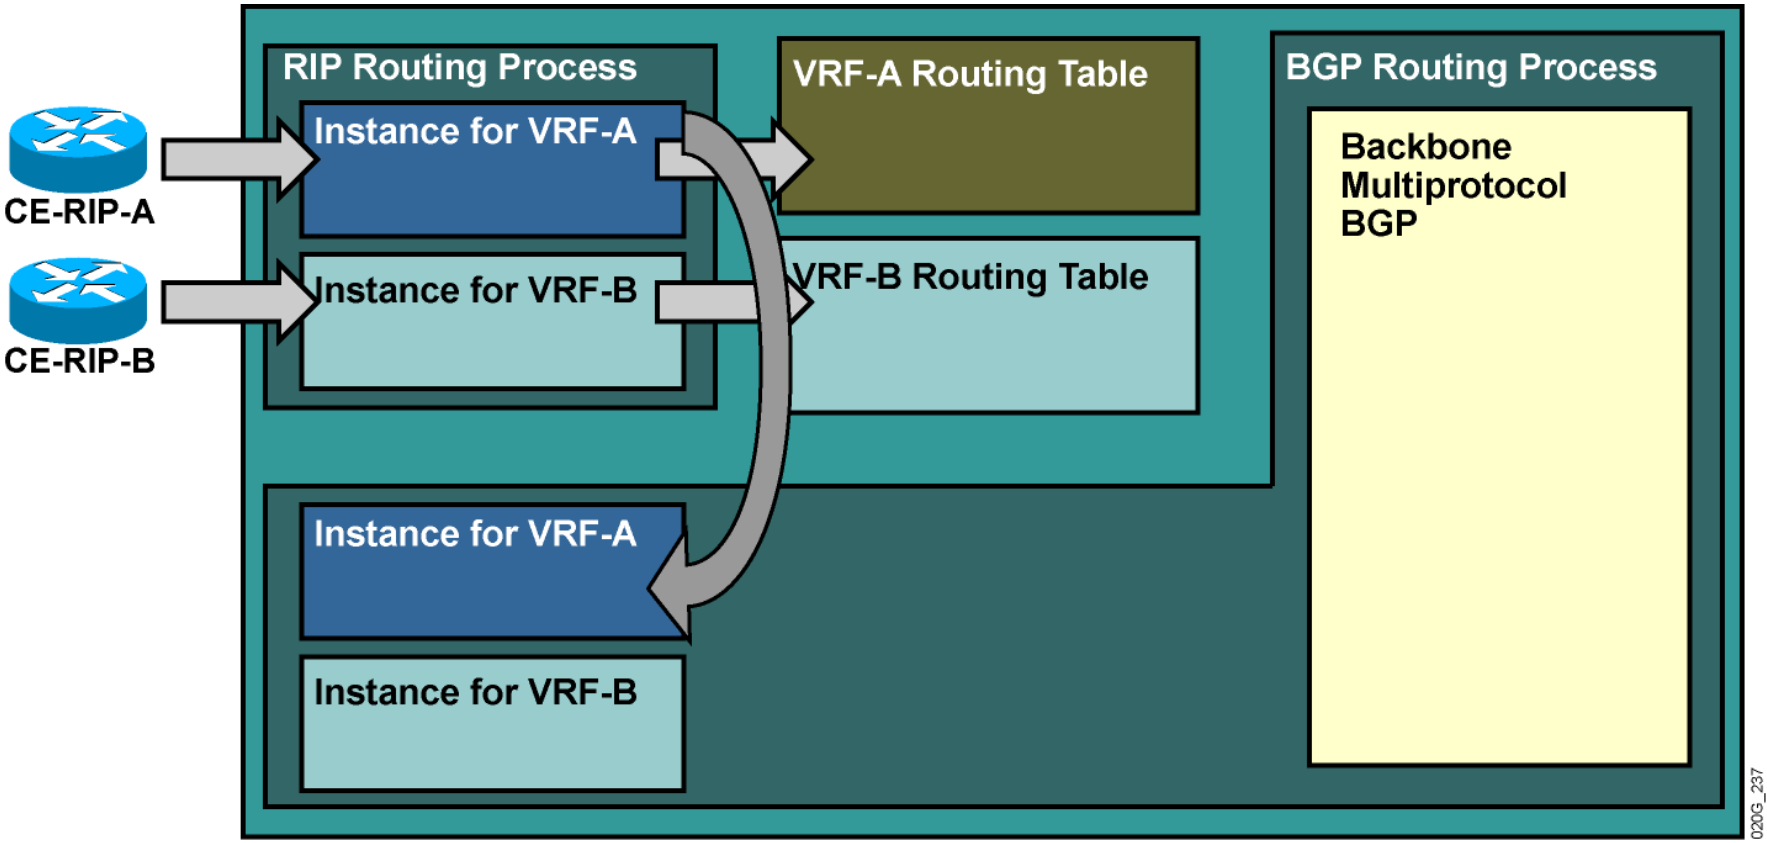

- **RIP routes entered in the VRF routing table are redistributed into BGP for further propagation into the MPLS VPN backbone.**

Non-BGP Route propagation—Outbound (Cont.)

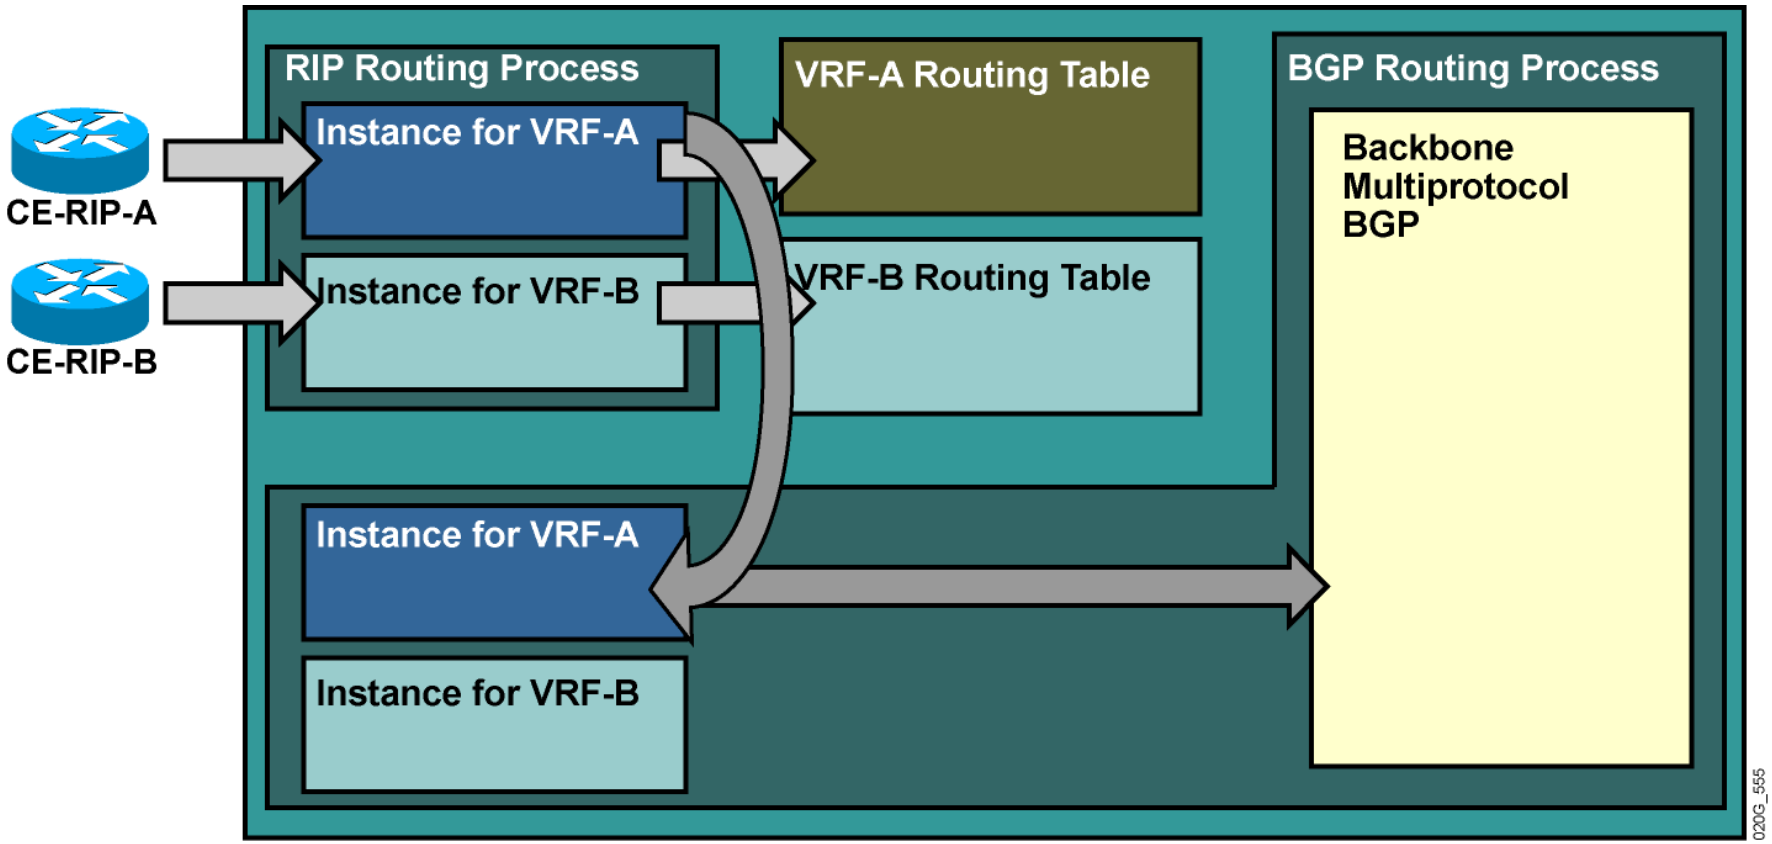

- RIP routes entered in the VRF routing table are redistributed into BGP for further propagation into the MPLS VPN backbone.
- Redistribution between RIP and BGP has to be conf gured for proper MPLS VPN operation.

Route Propagation—Inbound

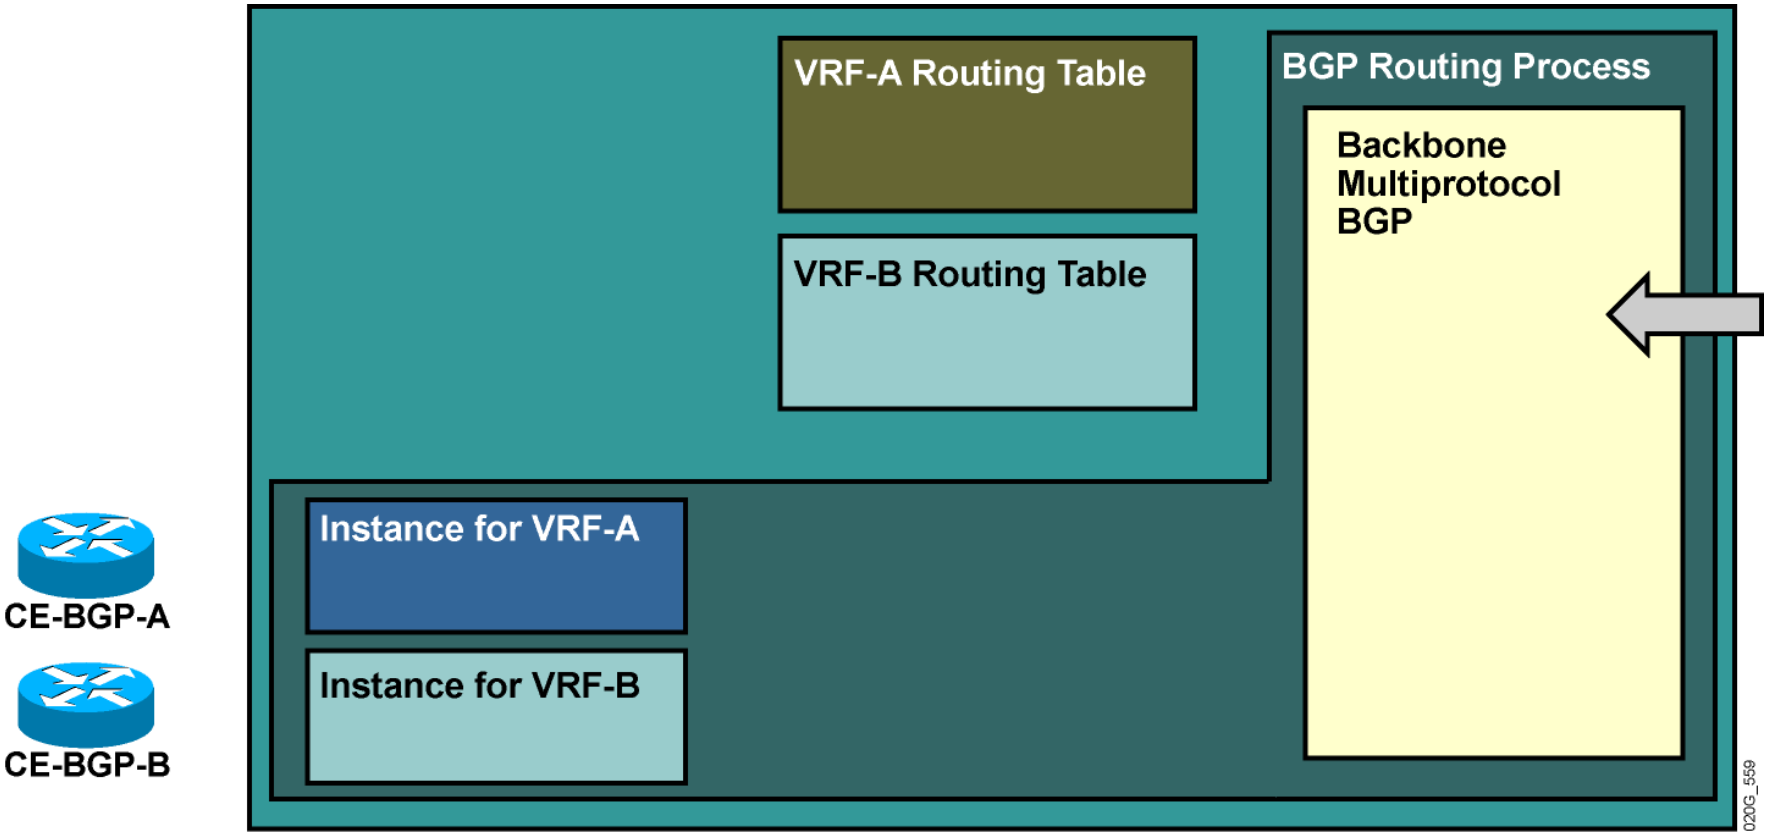

- VPNv4 prefixes are received from other PE routers.

Route Propagation—Inbound (Cont.)

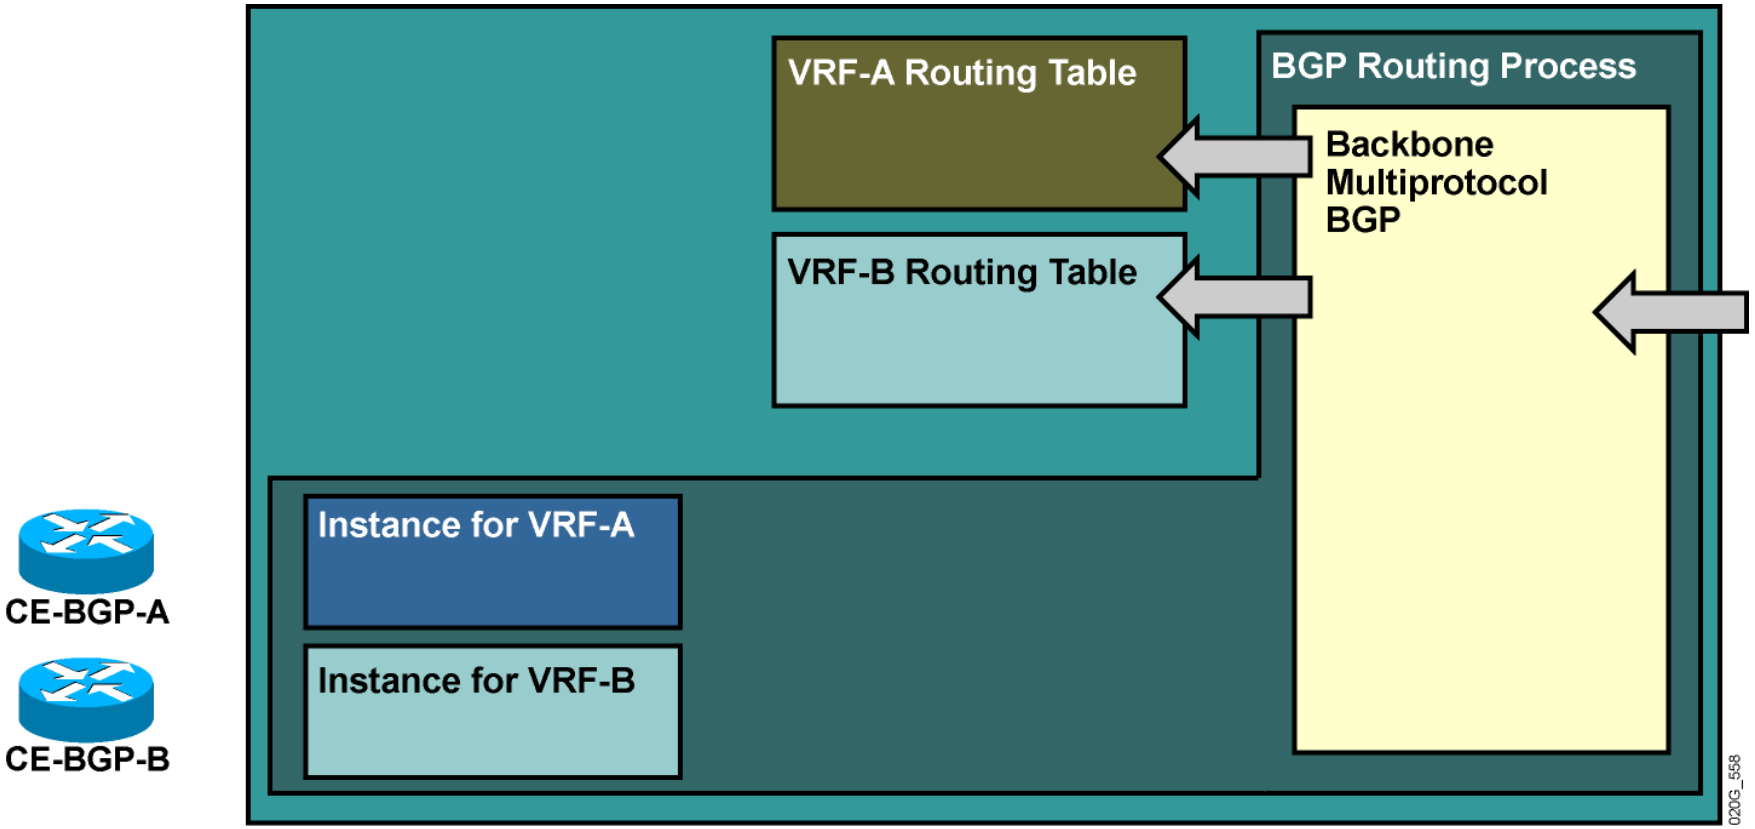

- VPNv4 prefixes are received from other PE routers.
- The VPNv4 prefixes are inserted into proper VRF routing tables based on their route targets and import route targets configured in VRFs.
- Route distinguisher is removed during this process.

Route Propagation—Inbound (Cont.)

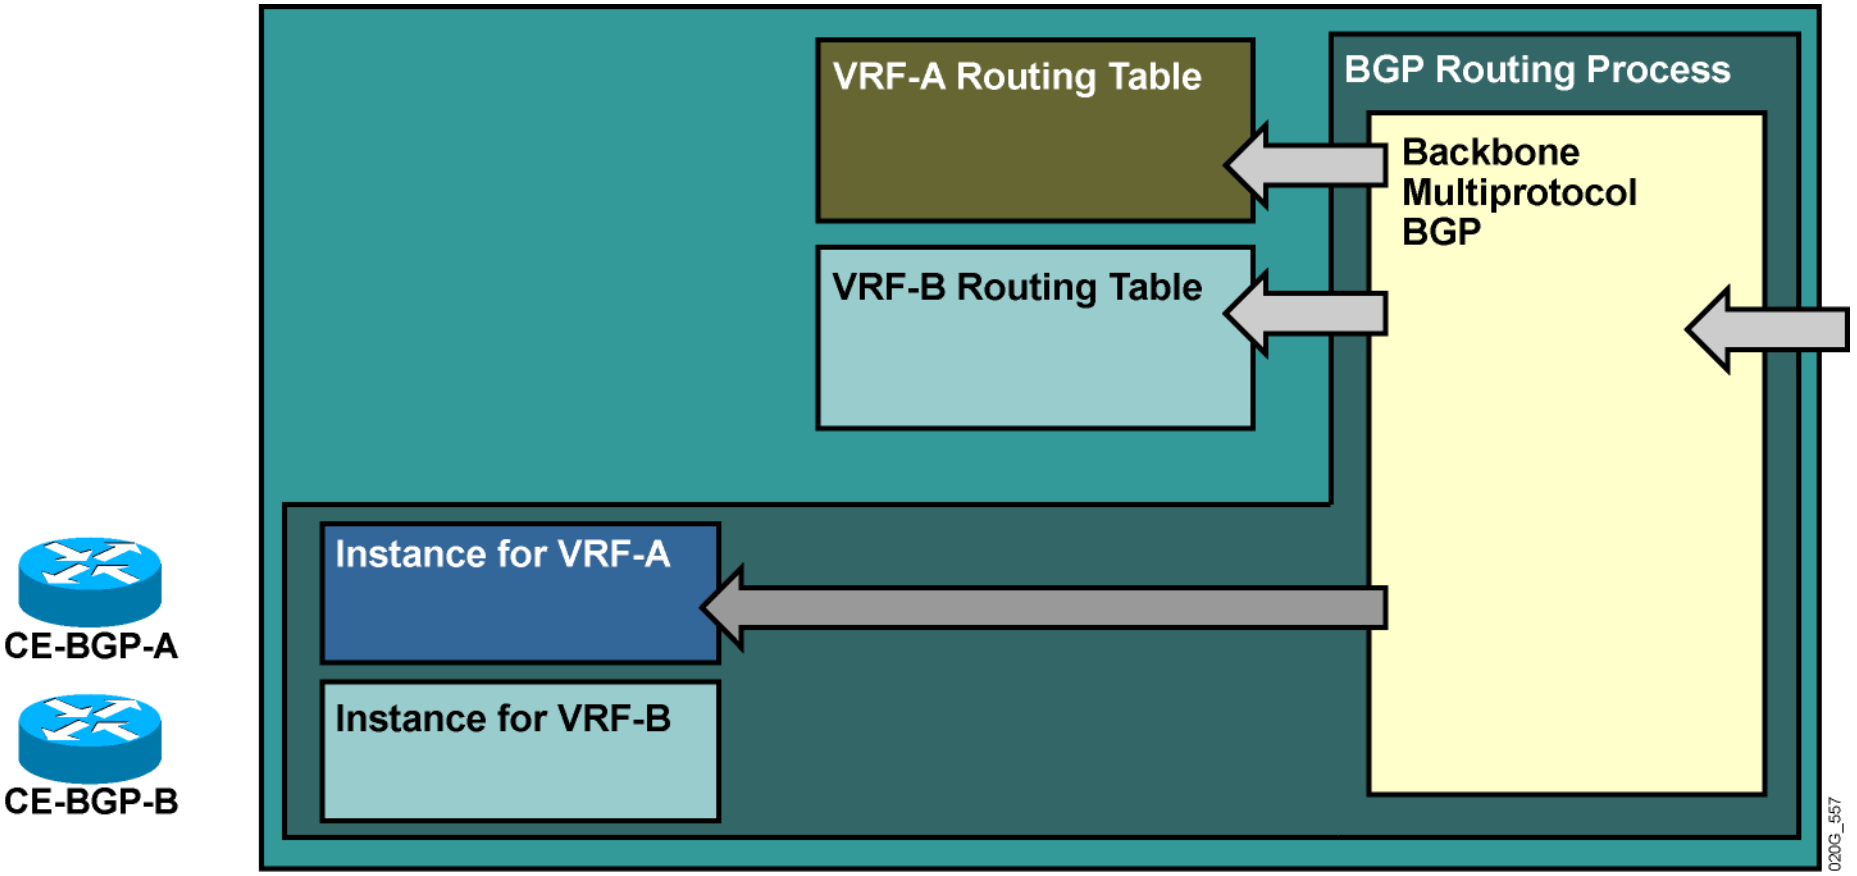

- **Routes are received from backbone MP-BGP and imported into a VRF.**

Route Propagation—Inbound (Cont.)

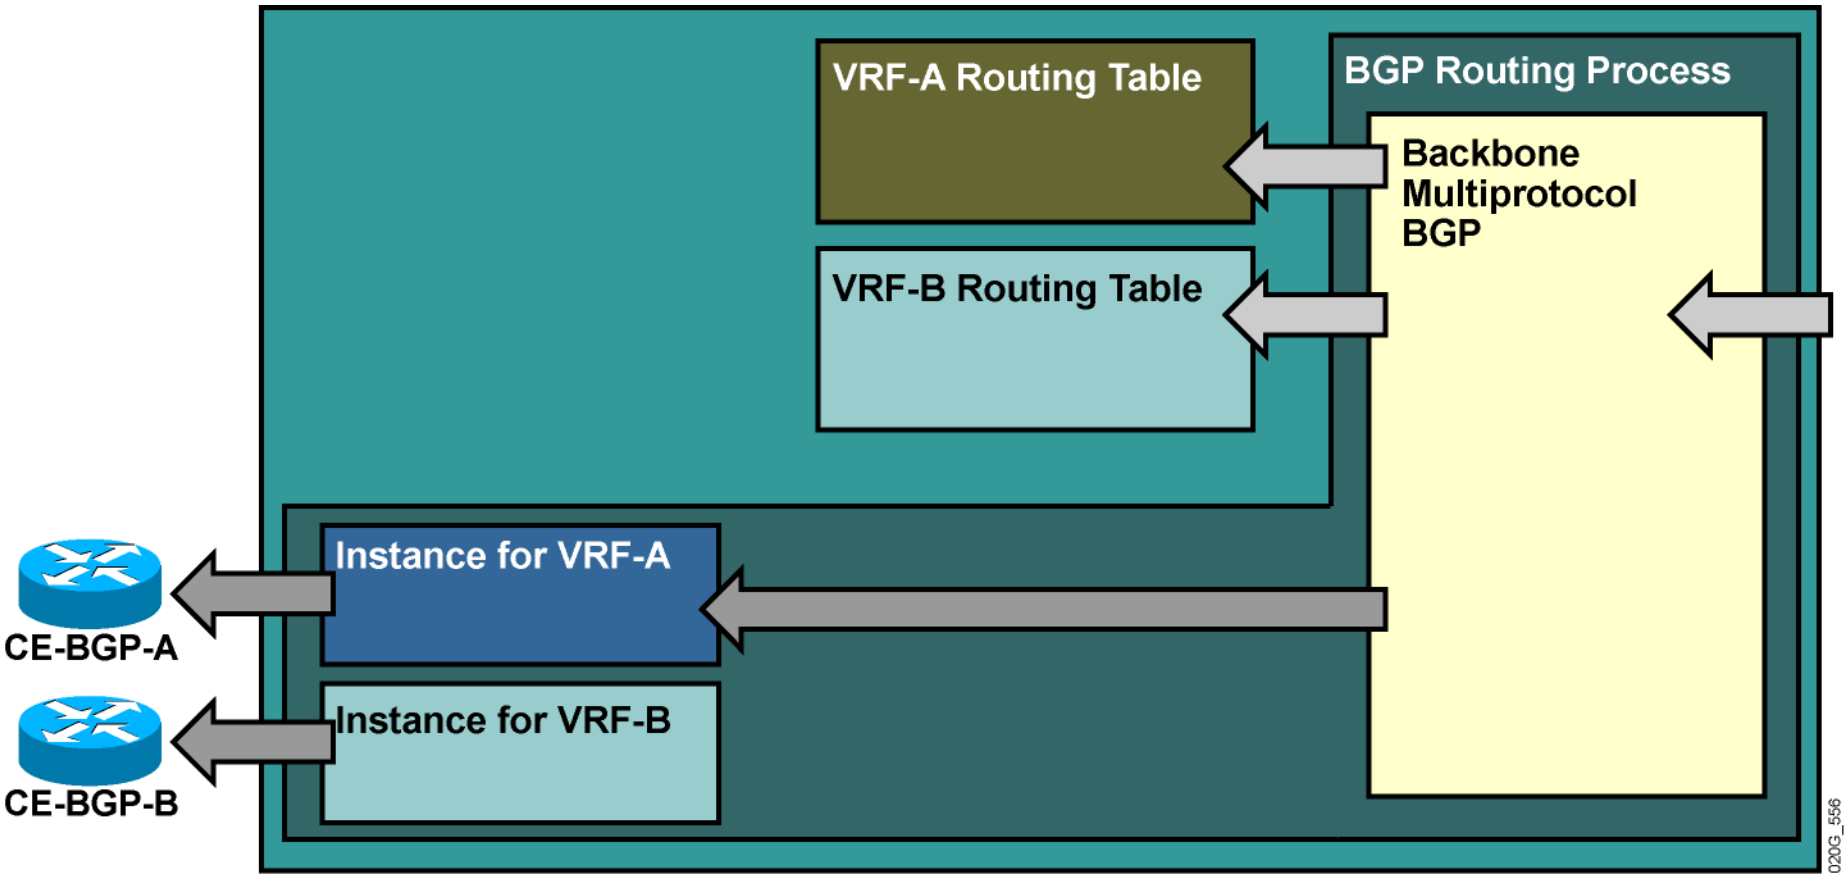

- Routes are received from backbone MP-BGP and imported into a VRF.
- IPv4 routes are forwarded to EBGP CE neighbors attached to that VRF.

Route Propagation—Inbound (Cont.)

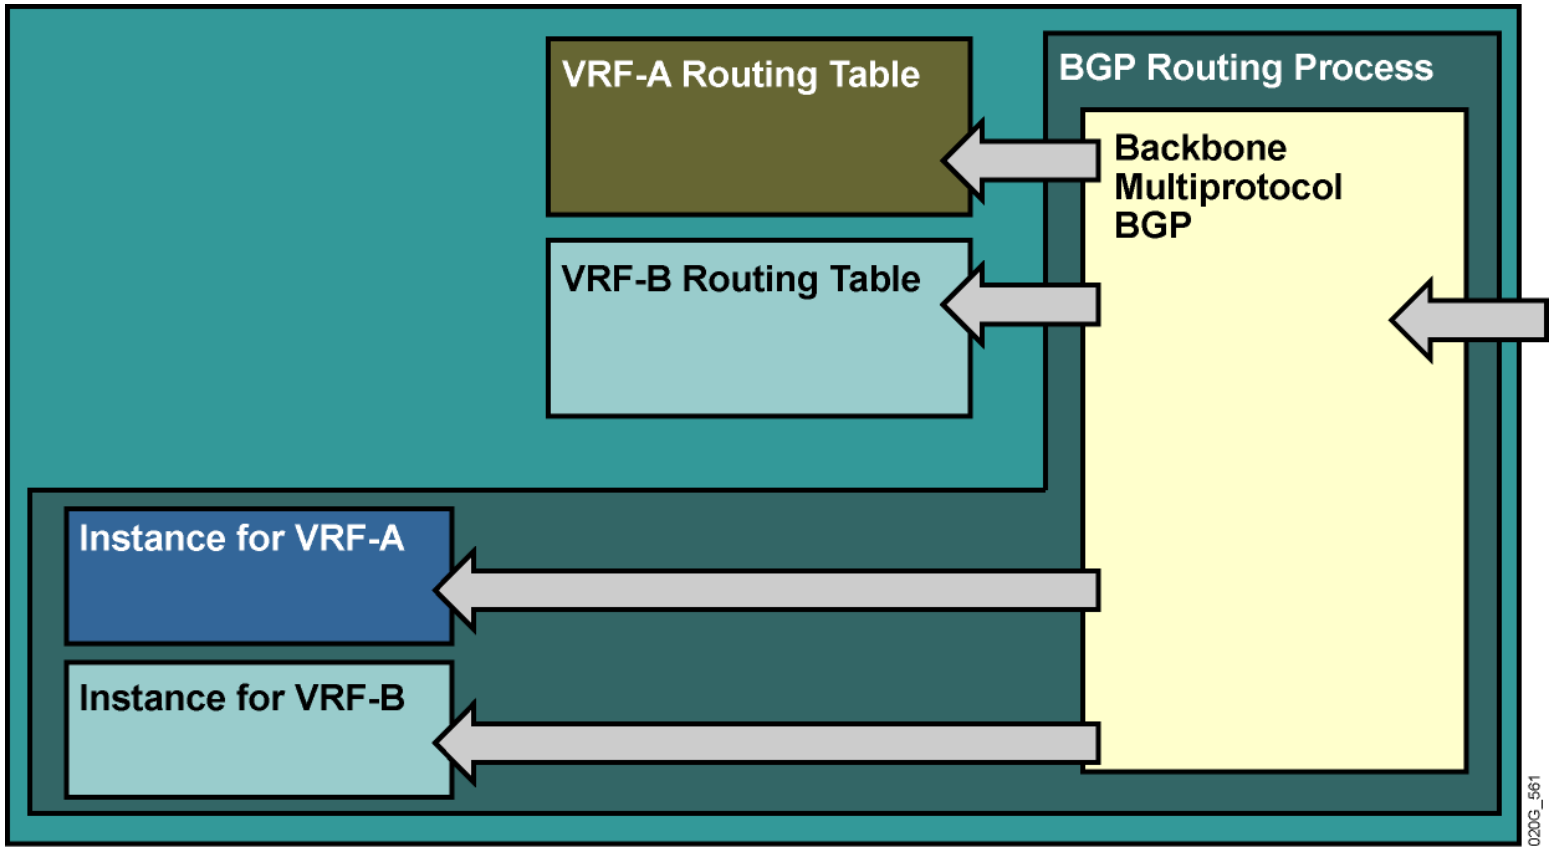

Route Propagation—Inbound (Cont.)

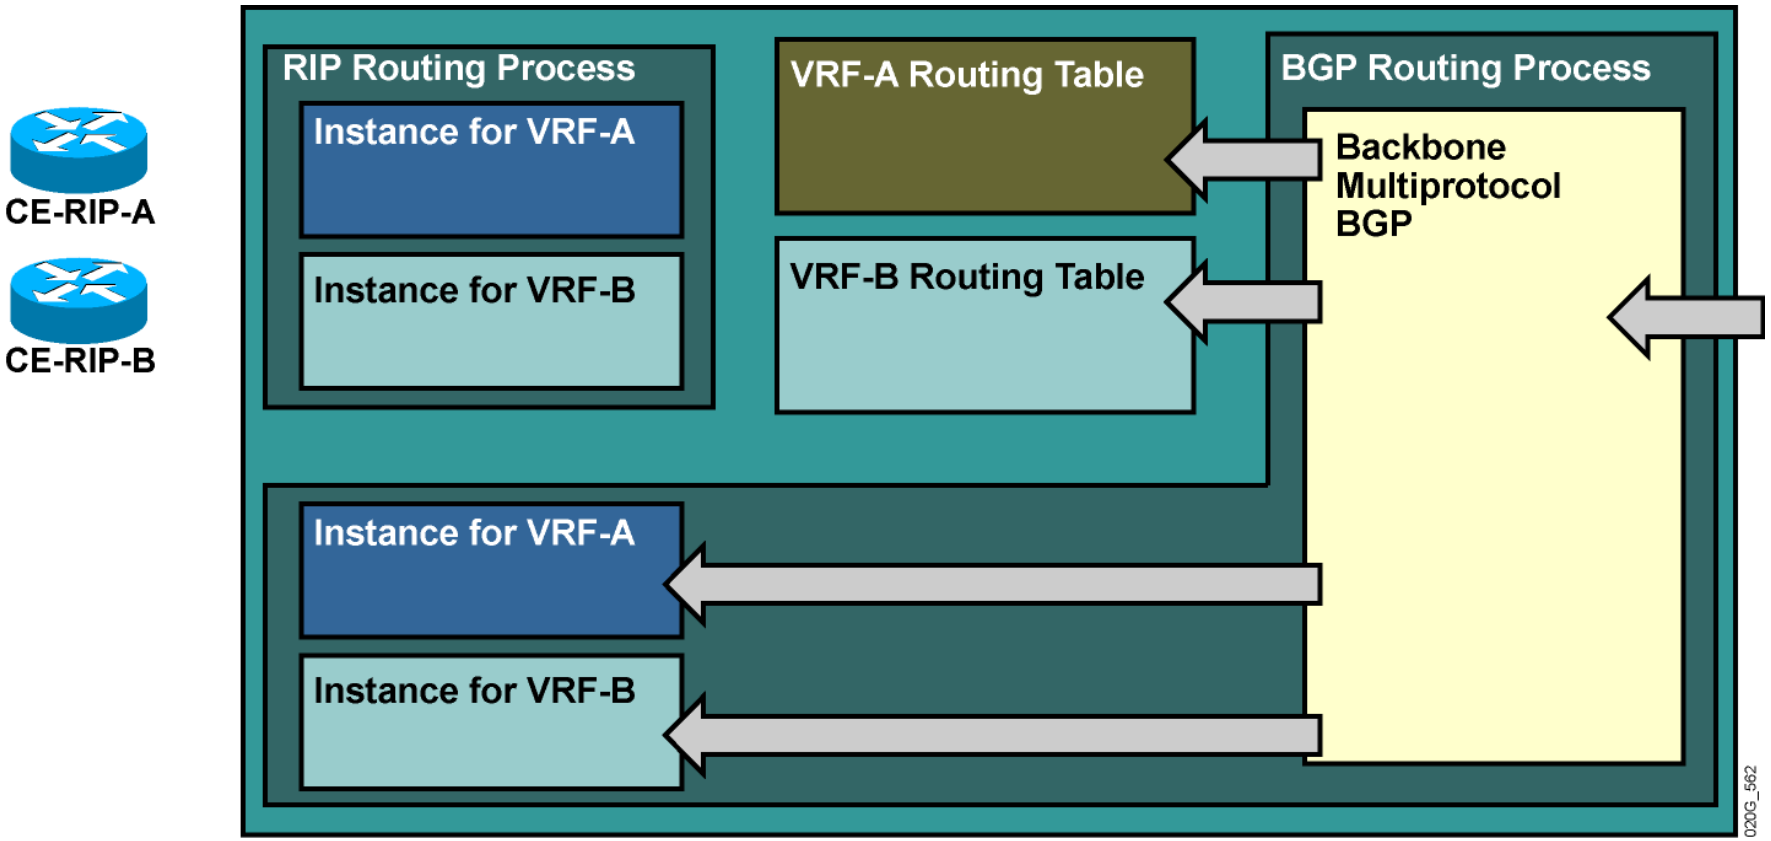

- **MP-IBGP routes imported into a VRF are redistributed into the instance of RIP conf gured for that VRF.**

Route Propagation—Inbound (Cont.)

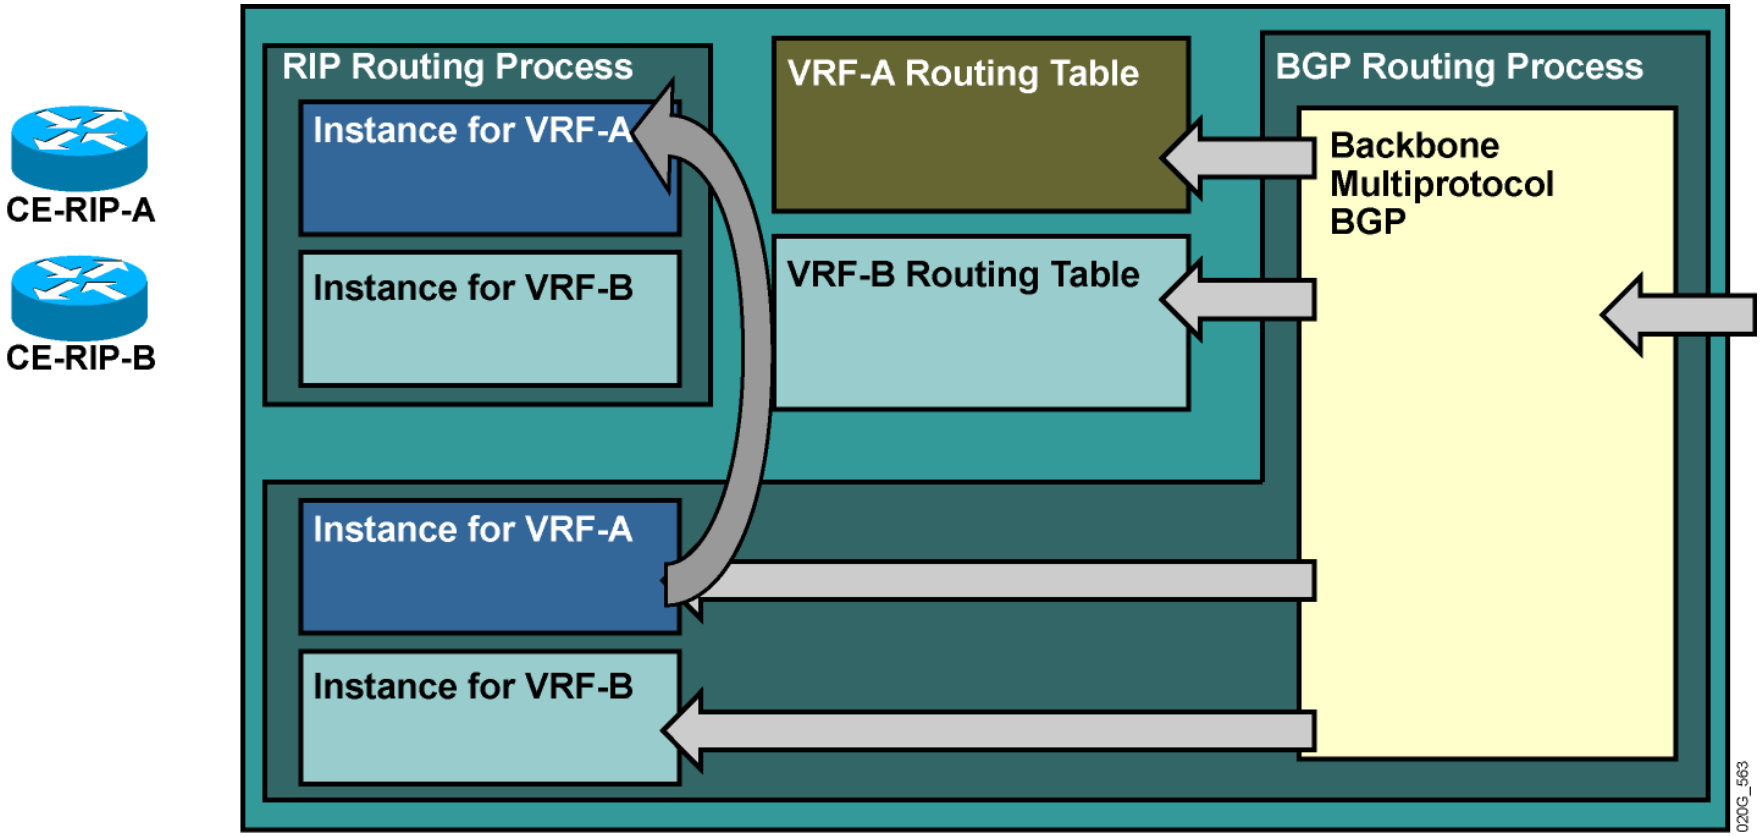

- MP-IBGP routes imported into a VRF are redistributed into the instance of RIP conf gured for that VRF.
- Redistribution between BGP and RIP has to be conf gured for end-to-end RIP routing between CE routers.

Route Propagation—Inbound (Cont.)

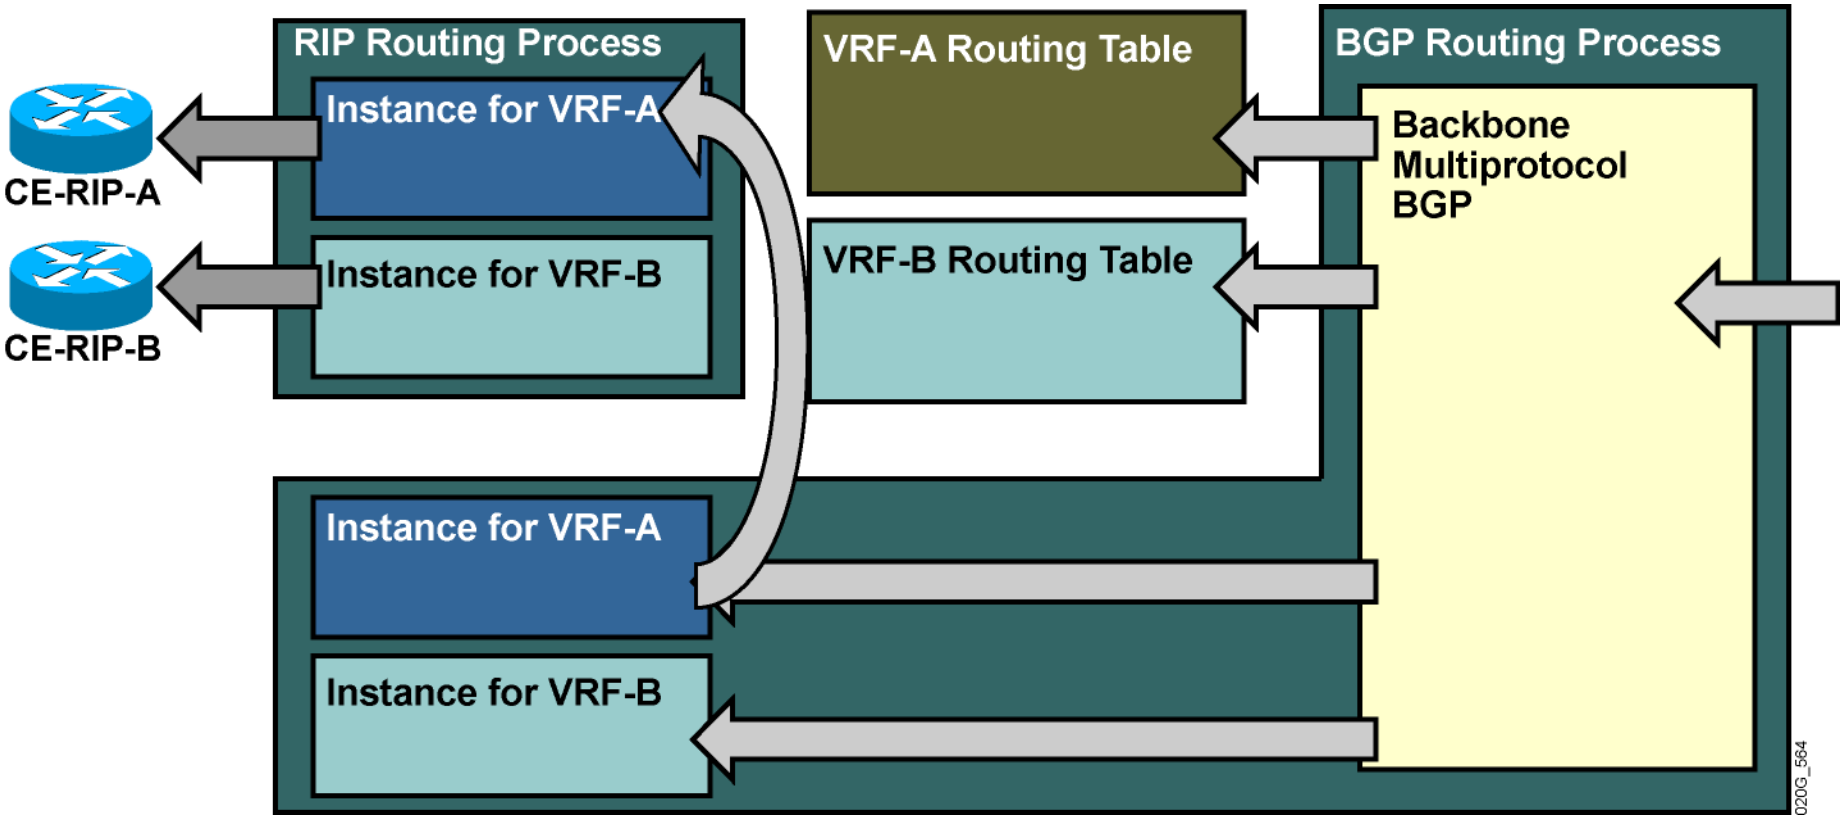

- **Routes redistributed from BGP into a VRF instance of RIP are sent to RIP-speaking CE routers.**

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 Bootcamp

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

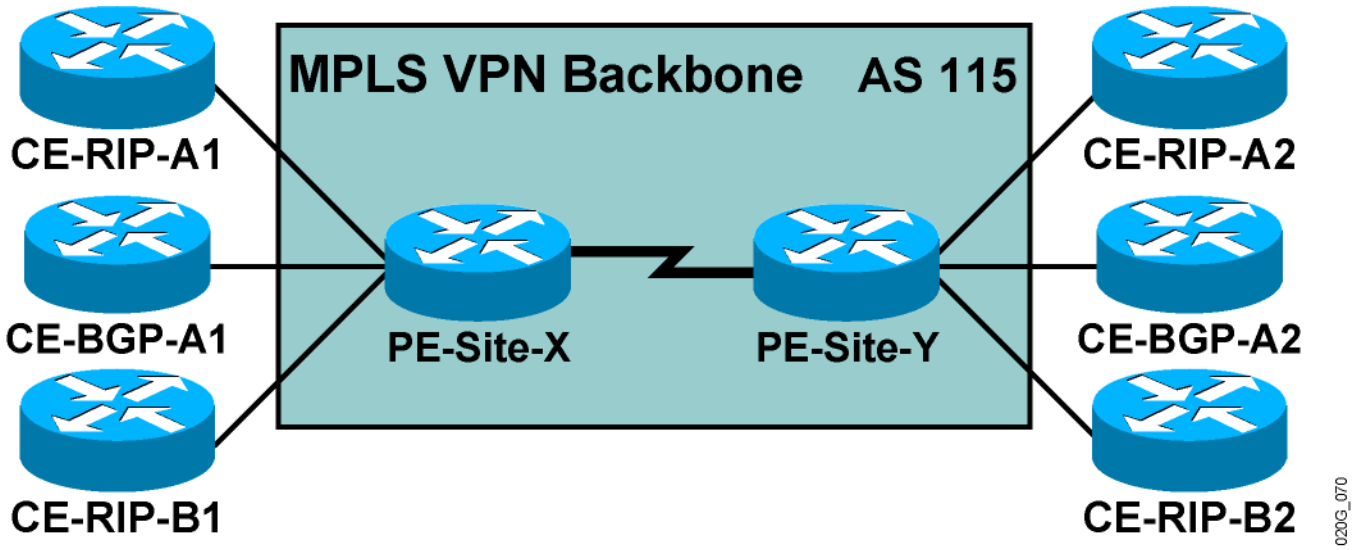

- The network supports two VPN customers.
- Customer A runs RIP and BGP with the service provider; customer B uses only RIP.
- Both customers use network 10.0.0.0.

MPLS VPN Network Example (Cont.)

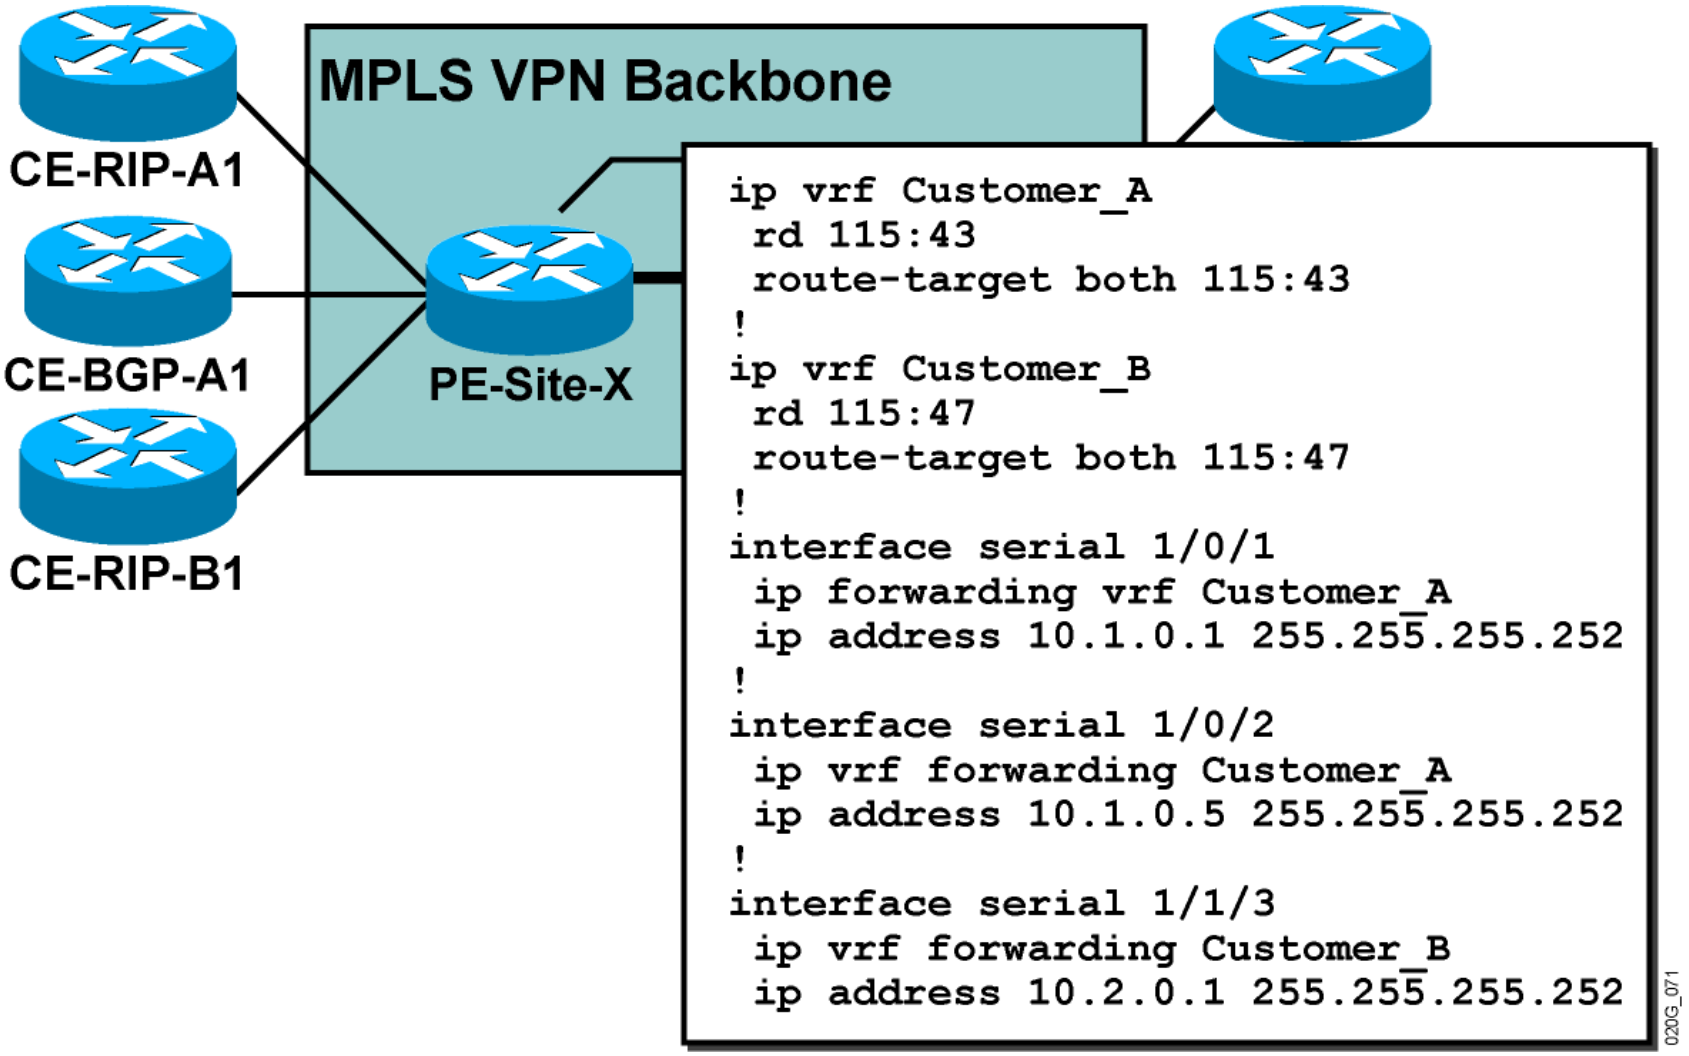

Summary

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 Bootcamp

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.)

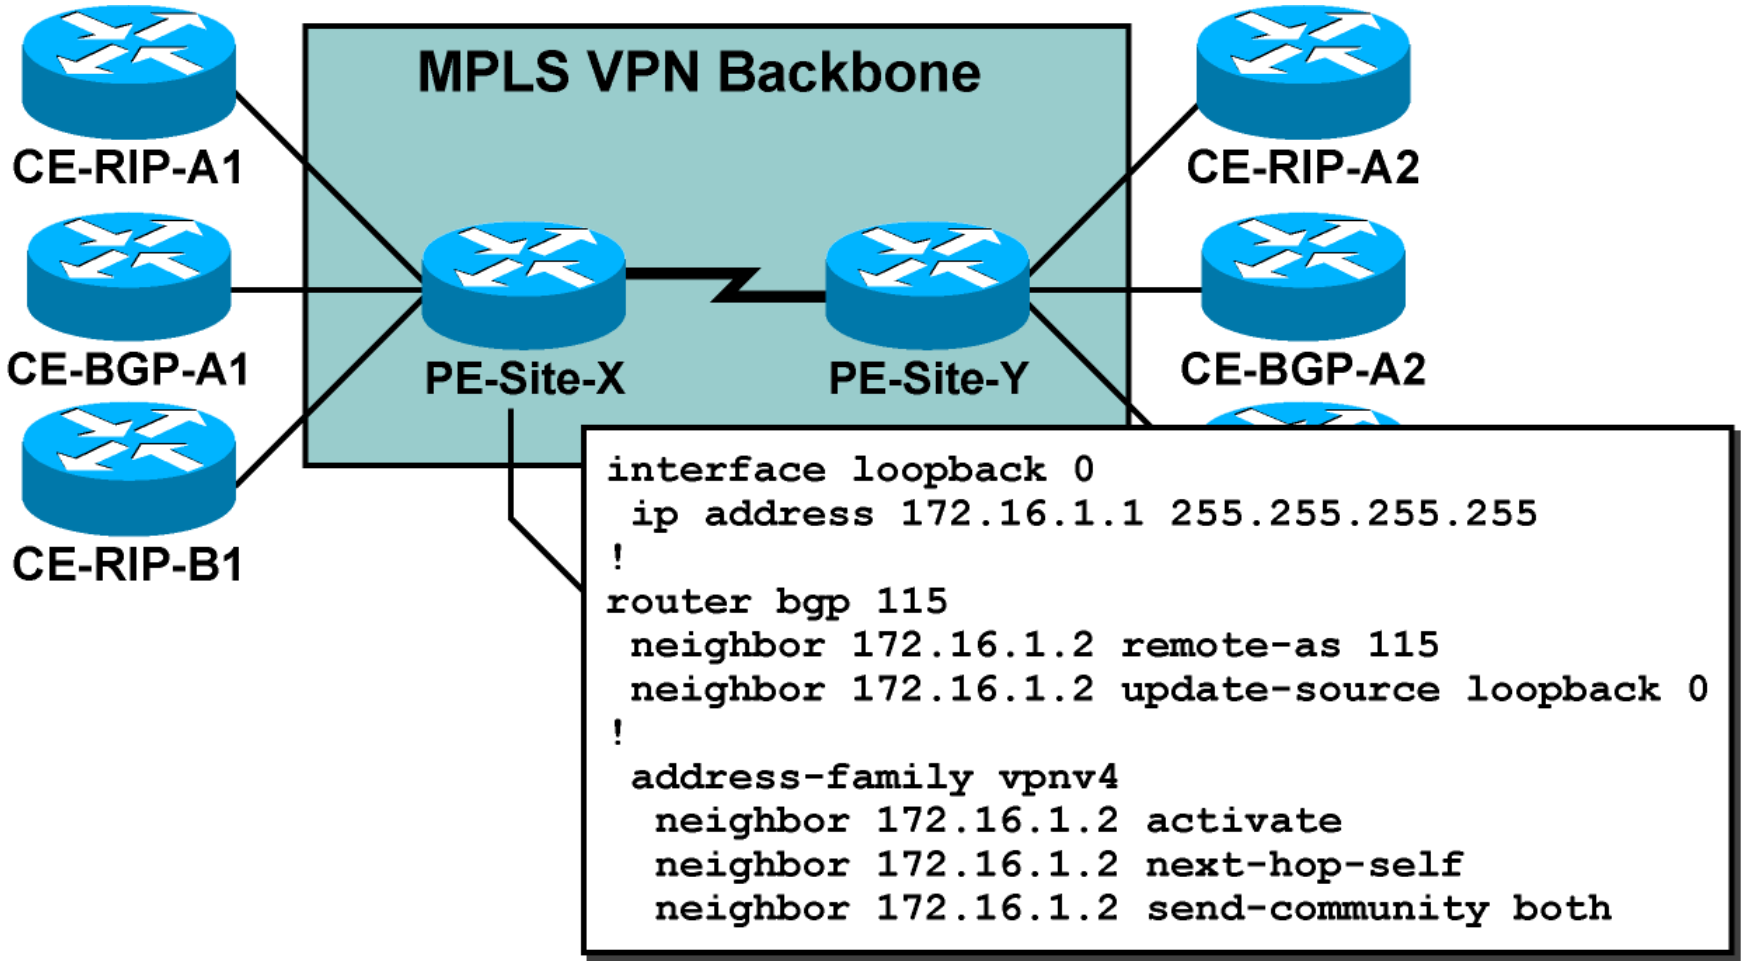

Disabling IPv4 Route Exchange

```
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 Bootcamp

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.)

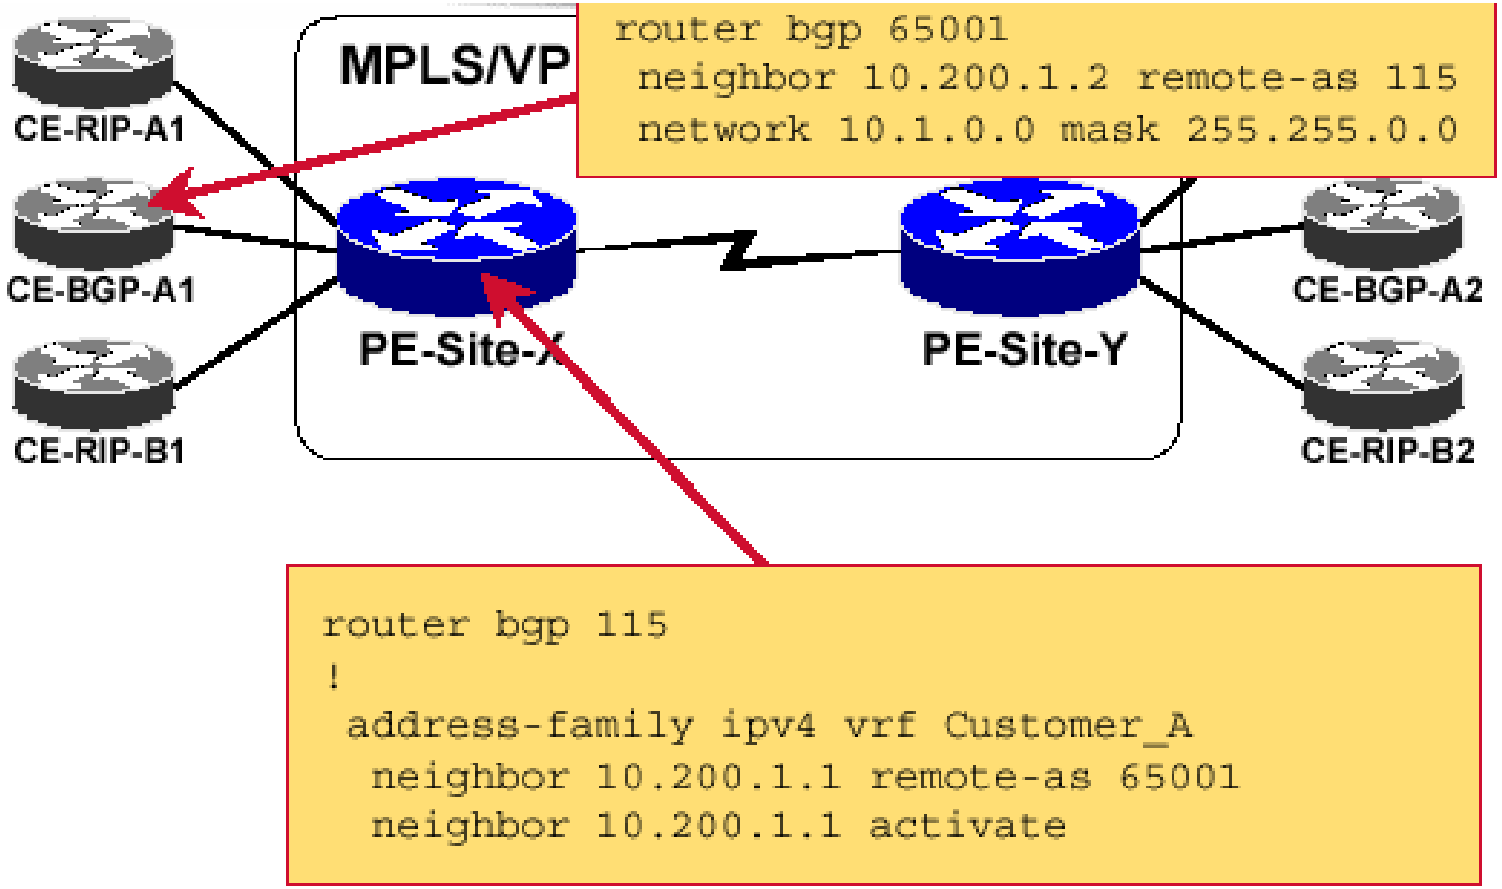

Limiting the Number of Routes in a VRF

- 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:
  - It can limit the number of routes received from a BGP neighbor.
  - 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.)

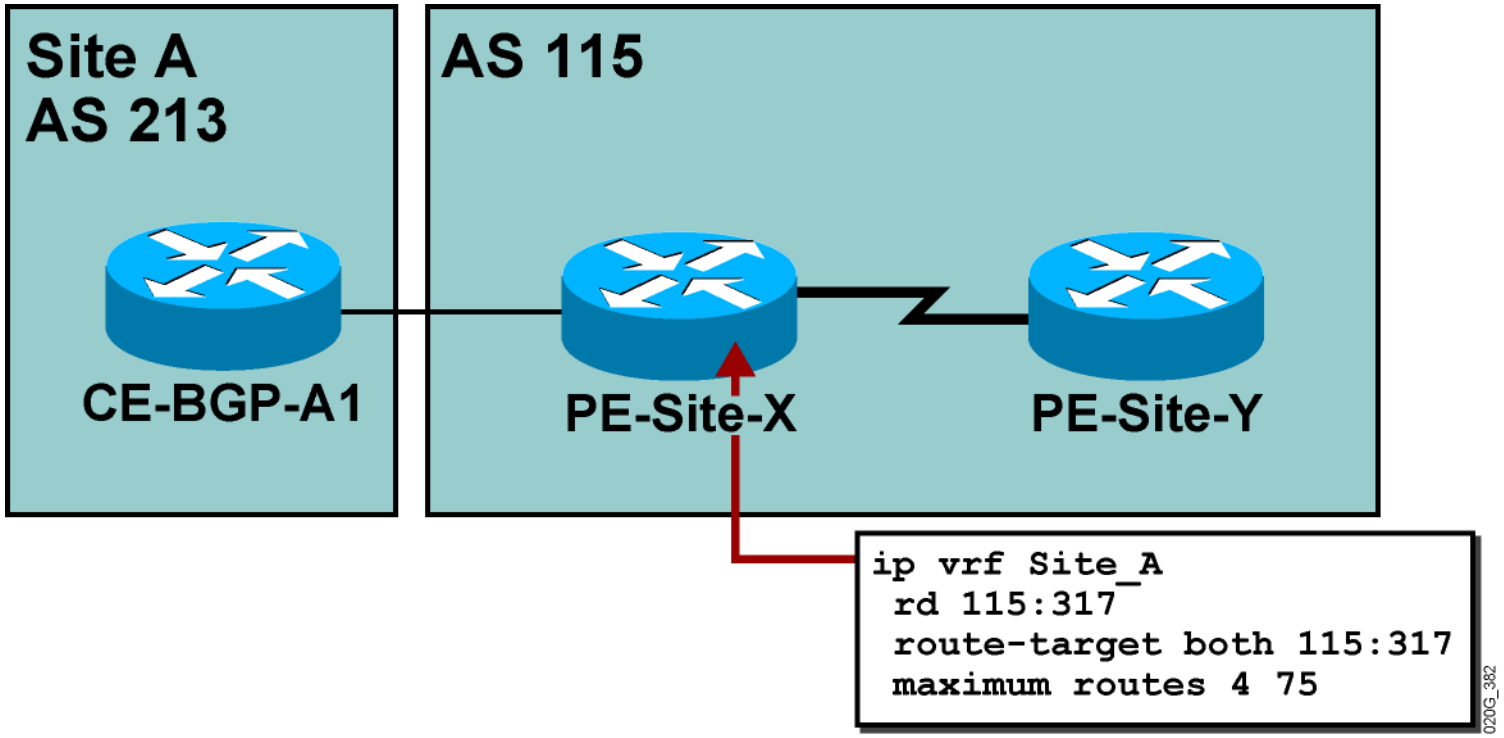

Limiting the Total Number of VRF Routes (Cont.)

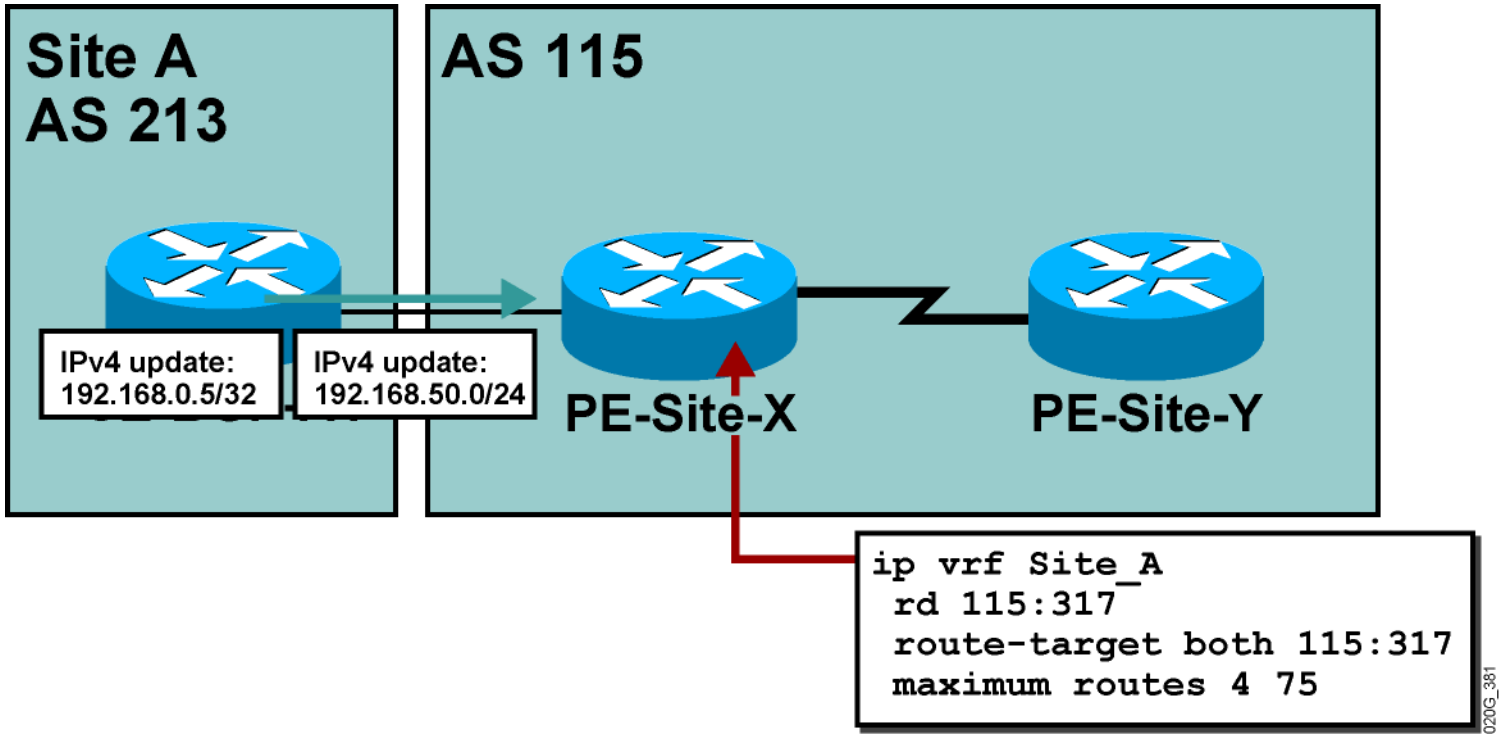

Limiting the Total Number of VRF Routes (Cont.)

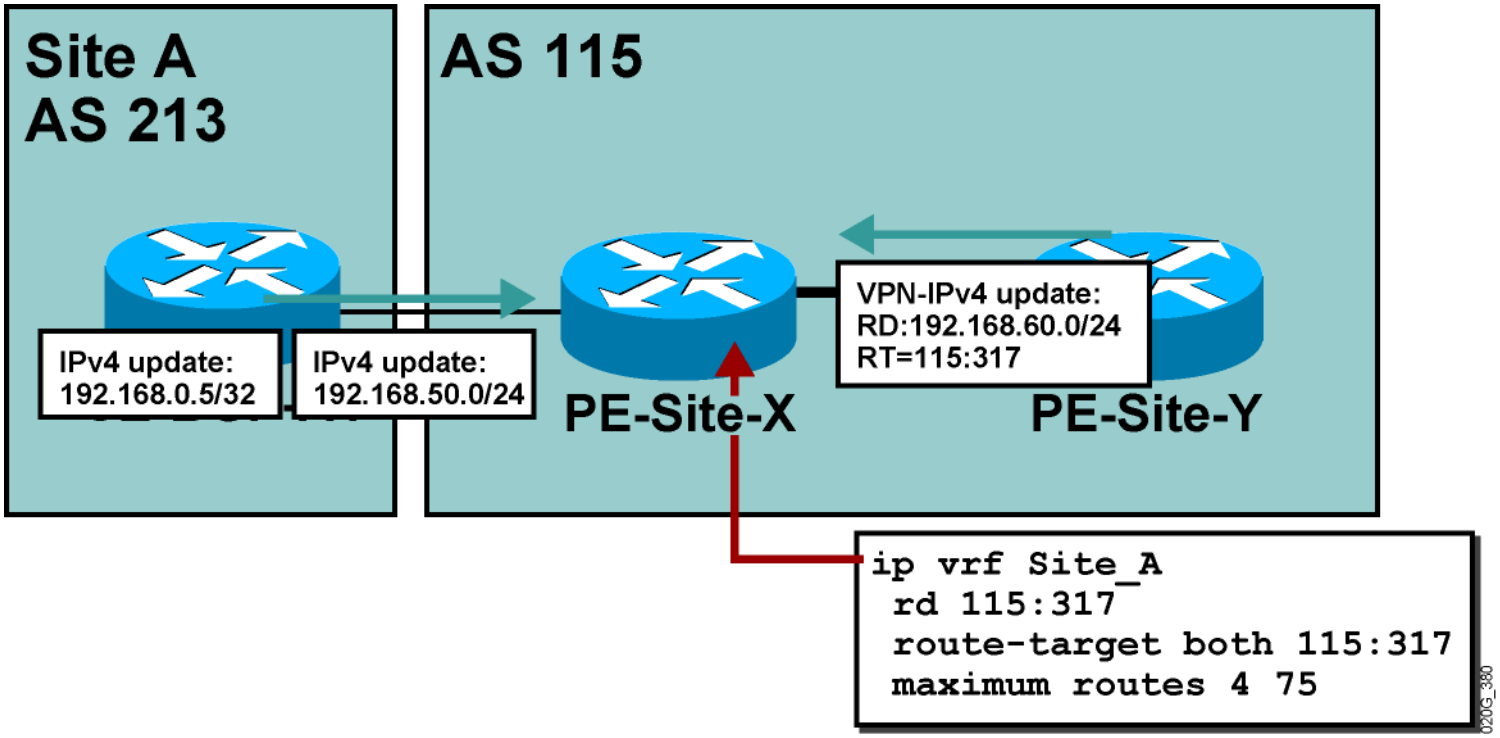

Limiting the Total Number of VRF Routes (Cont.)

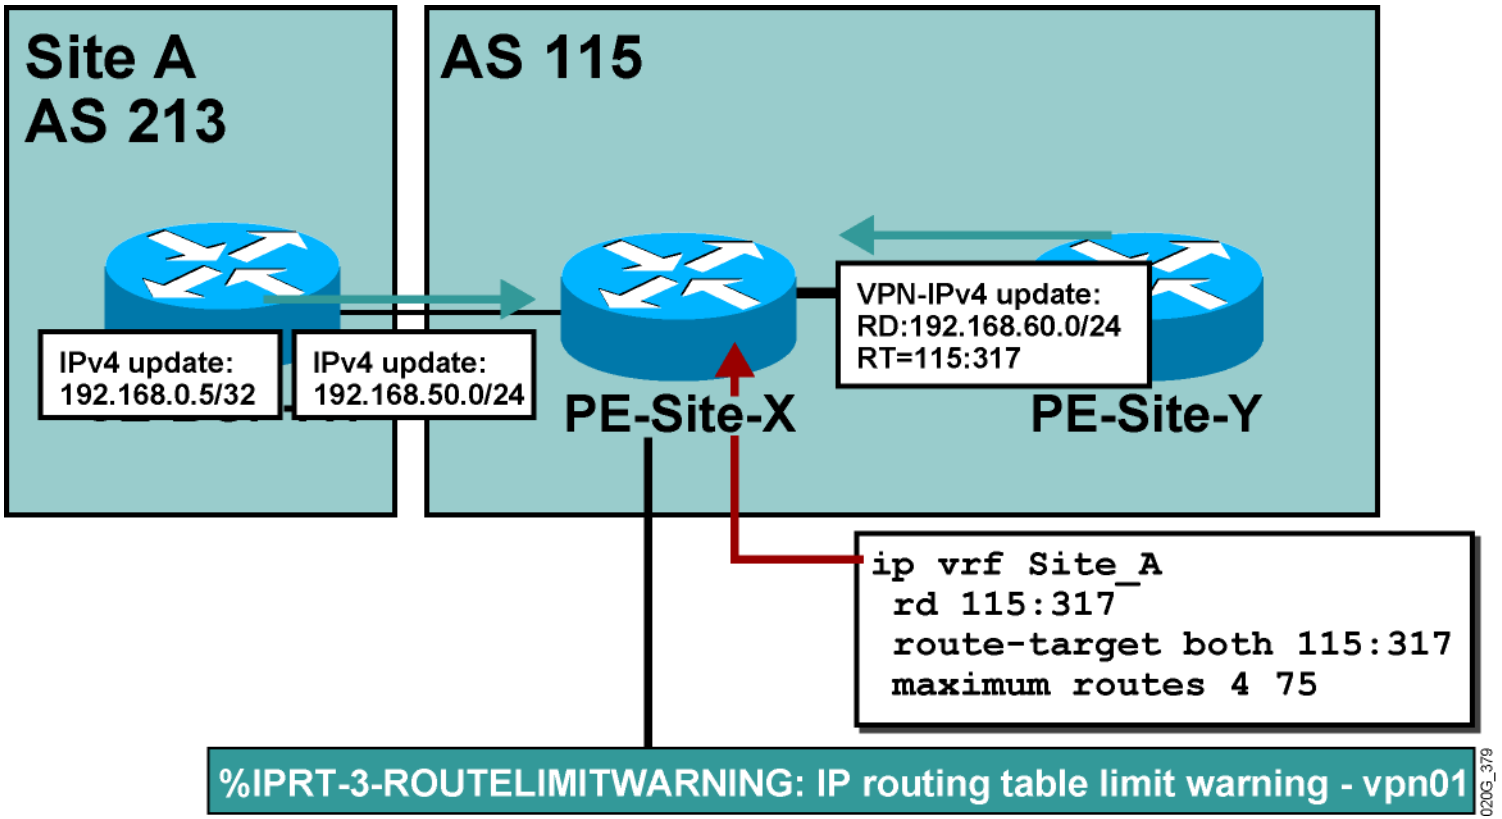

Limiting the Total Number of VRF Routes (Cont.)

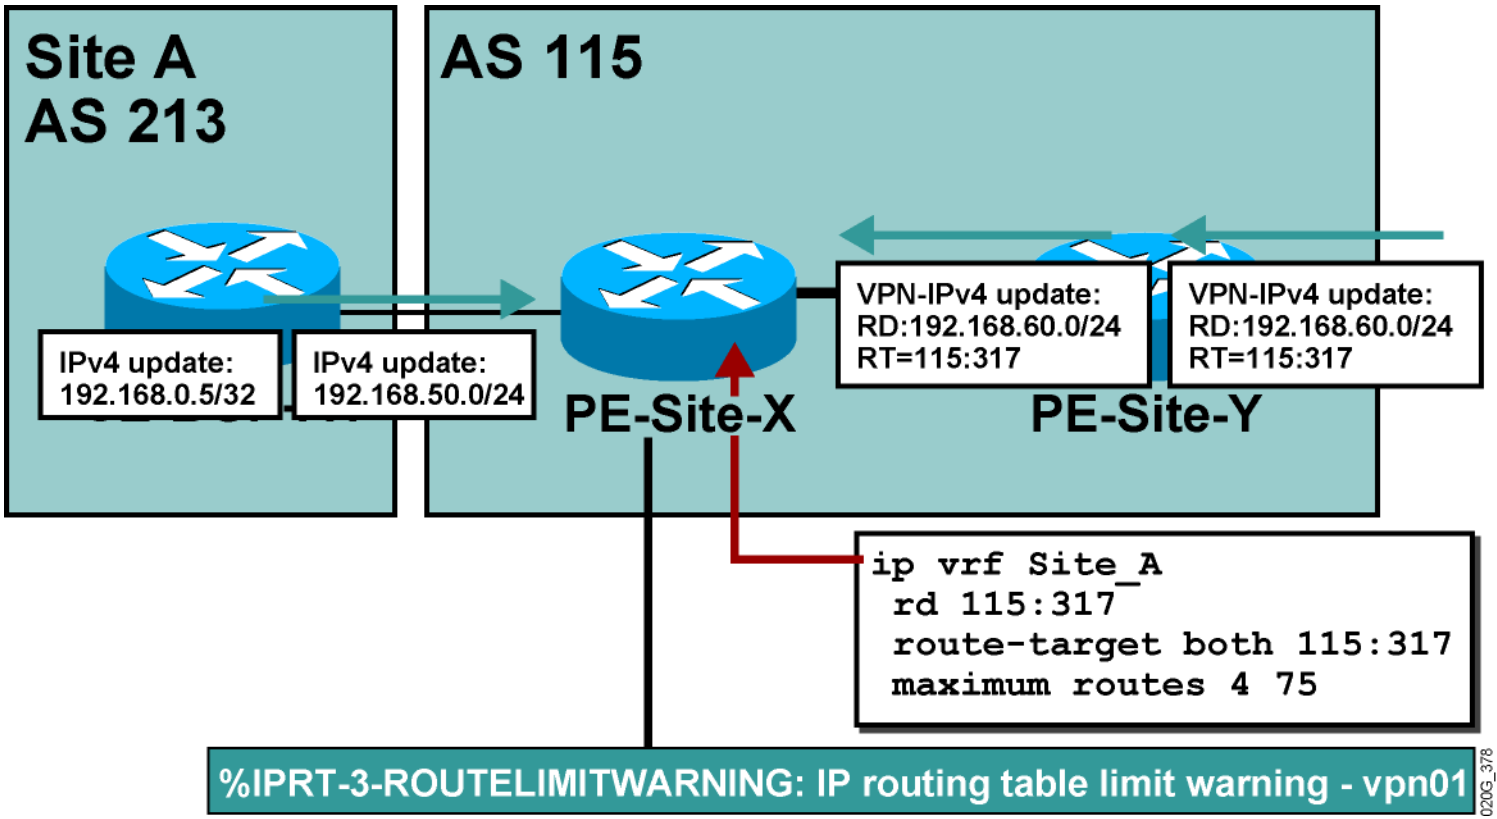

Limiting the Total Number of VRF Routes (Cont.)

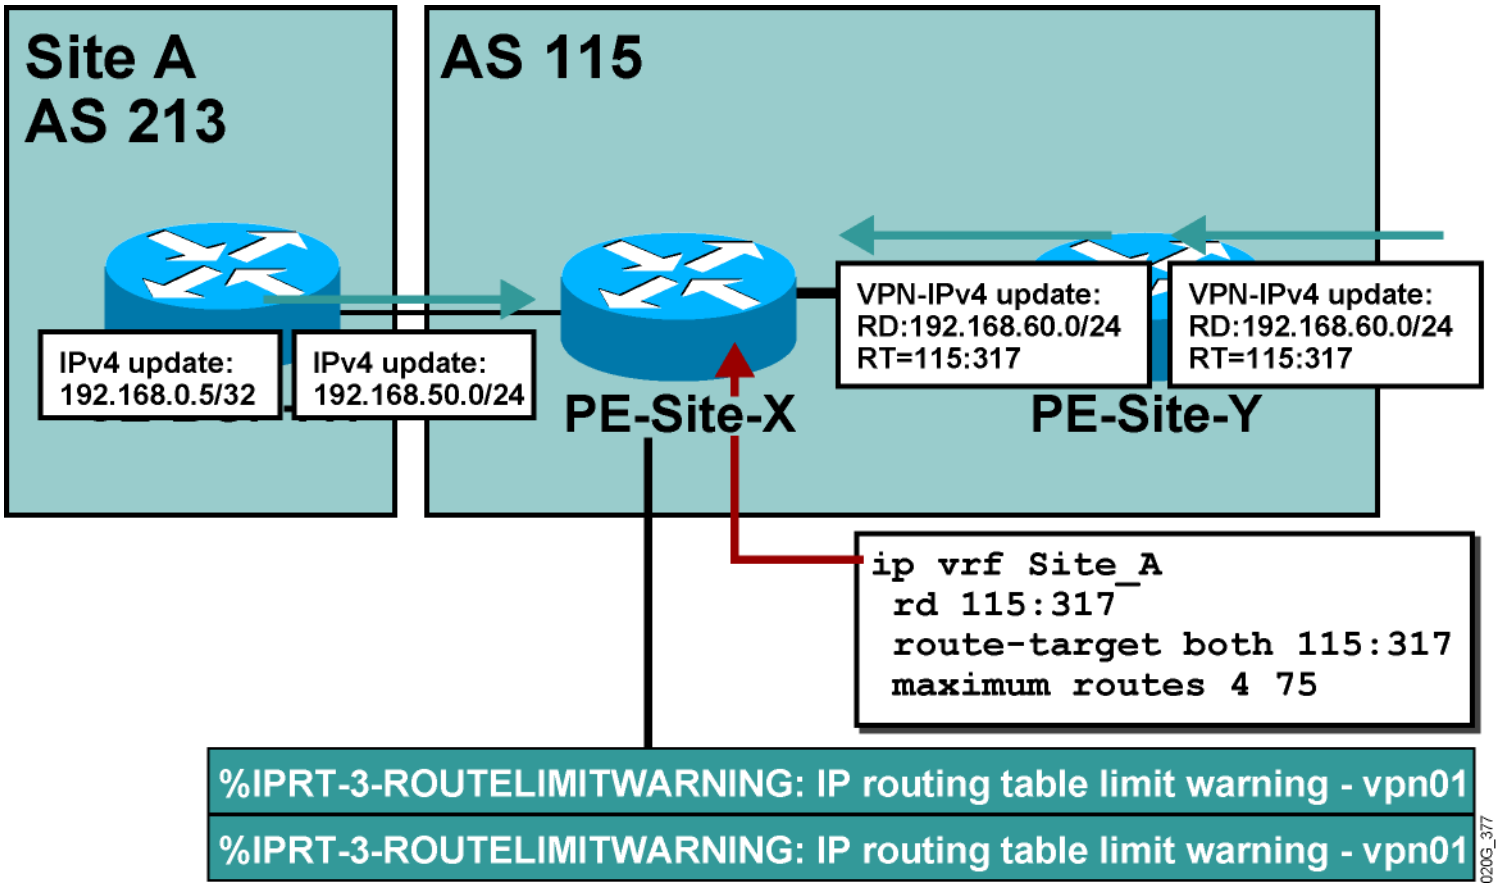

Limiting the Total Number of VRF Routes (Cont.)

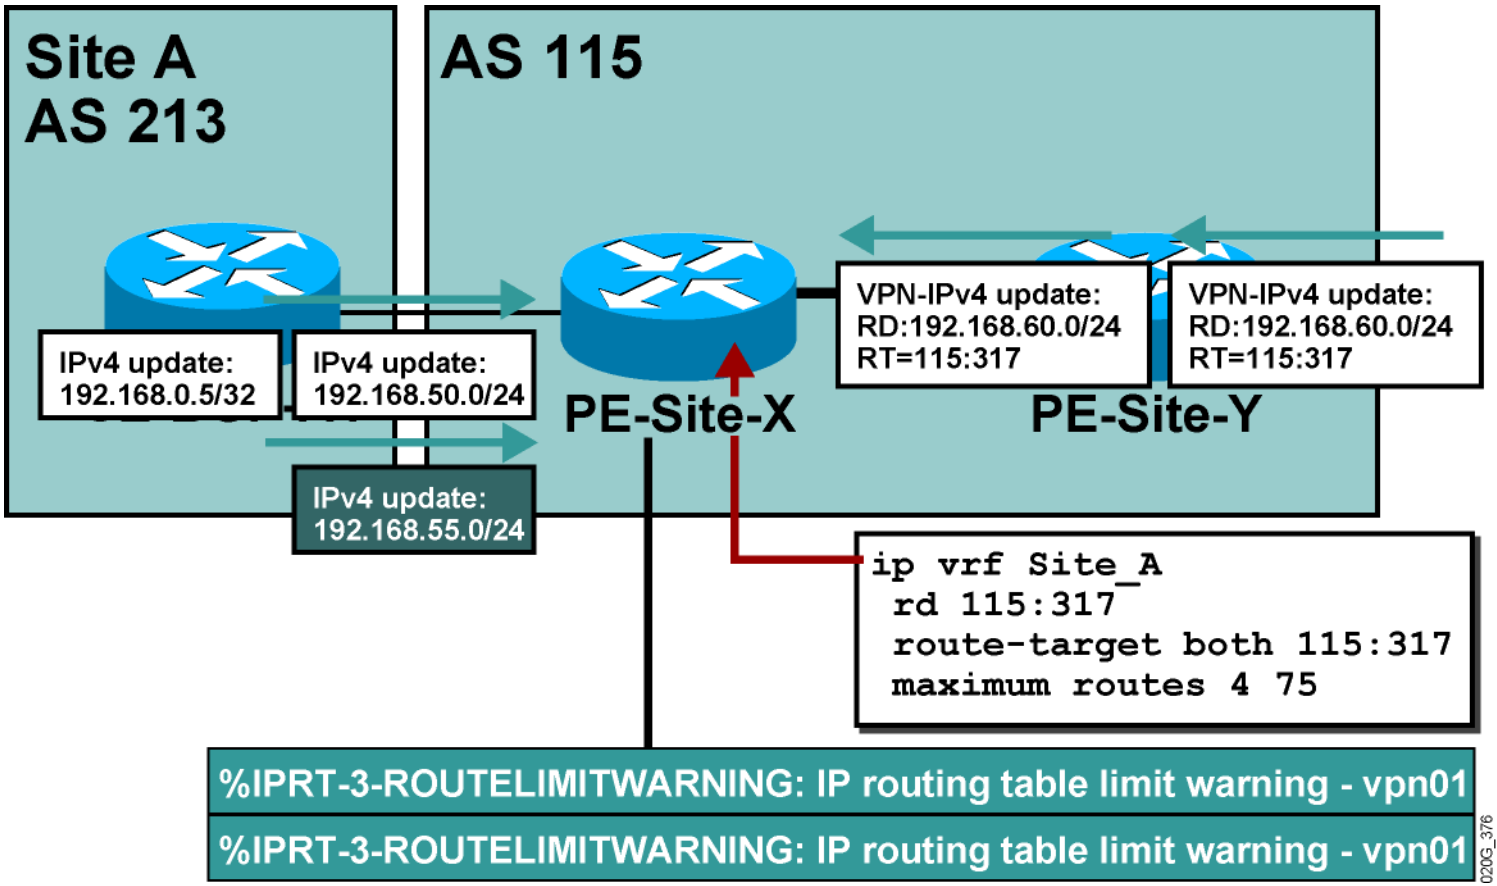

Limiting the Total Number of VRF Routes (Cont.)

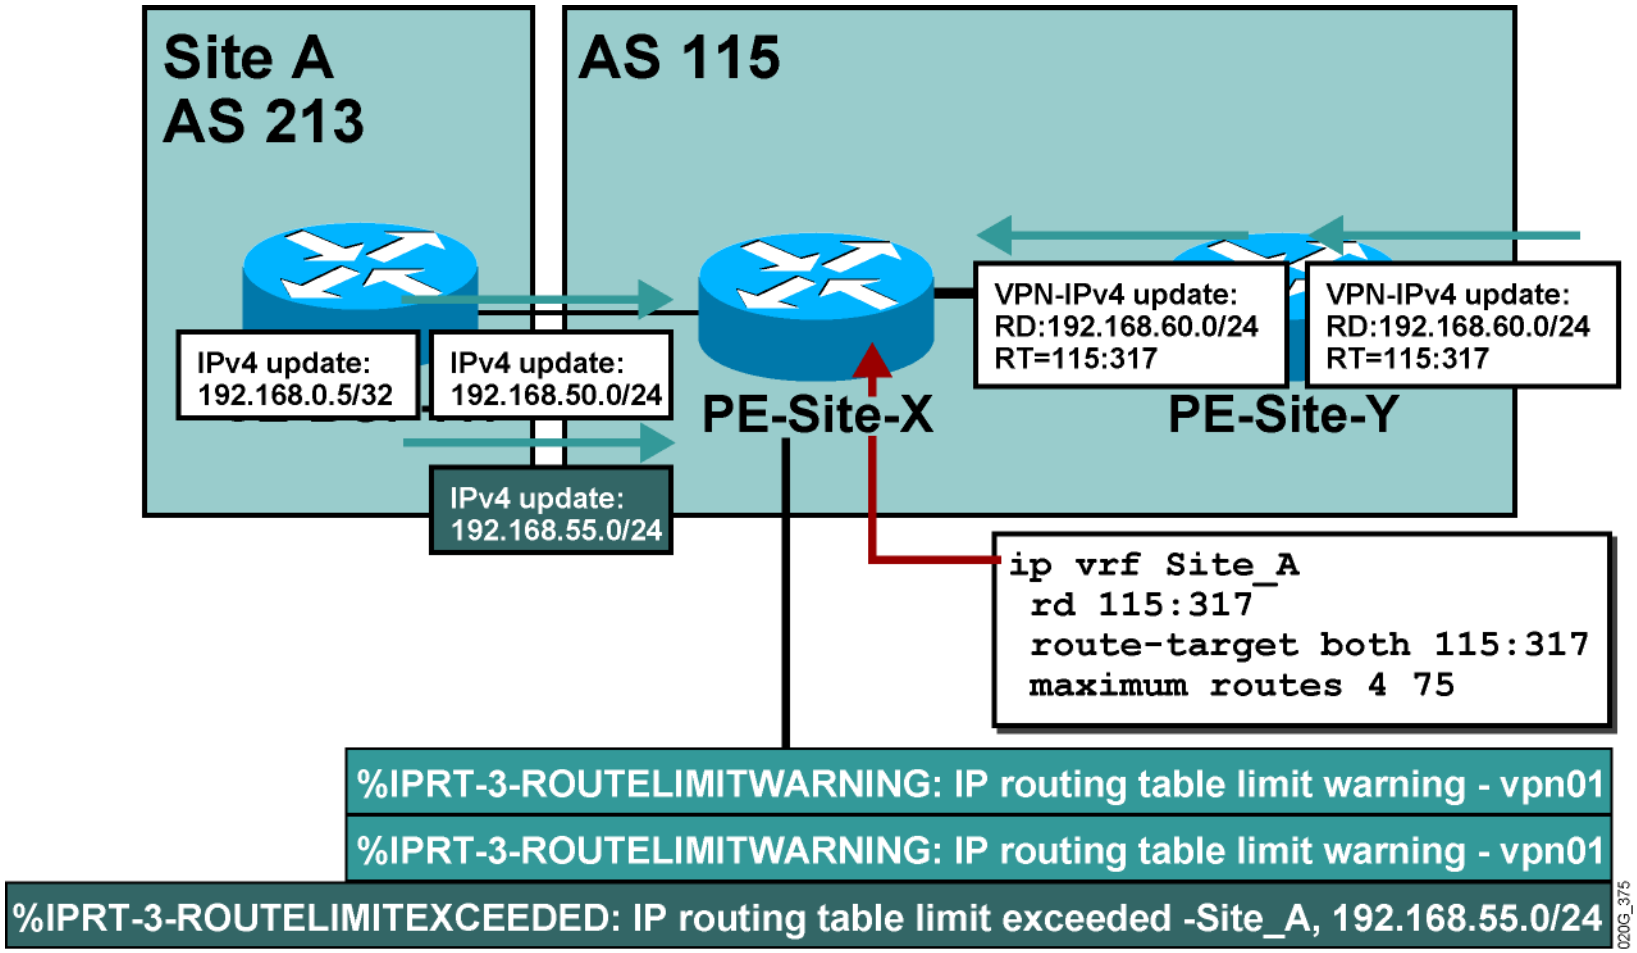

AS-override
The Issue

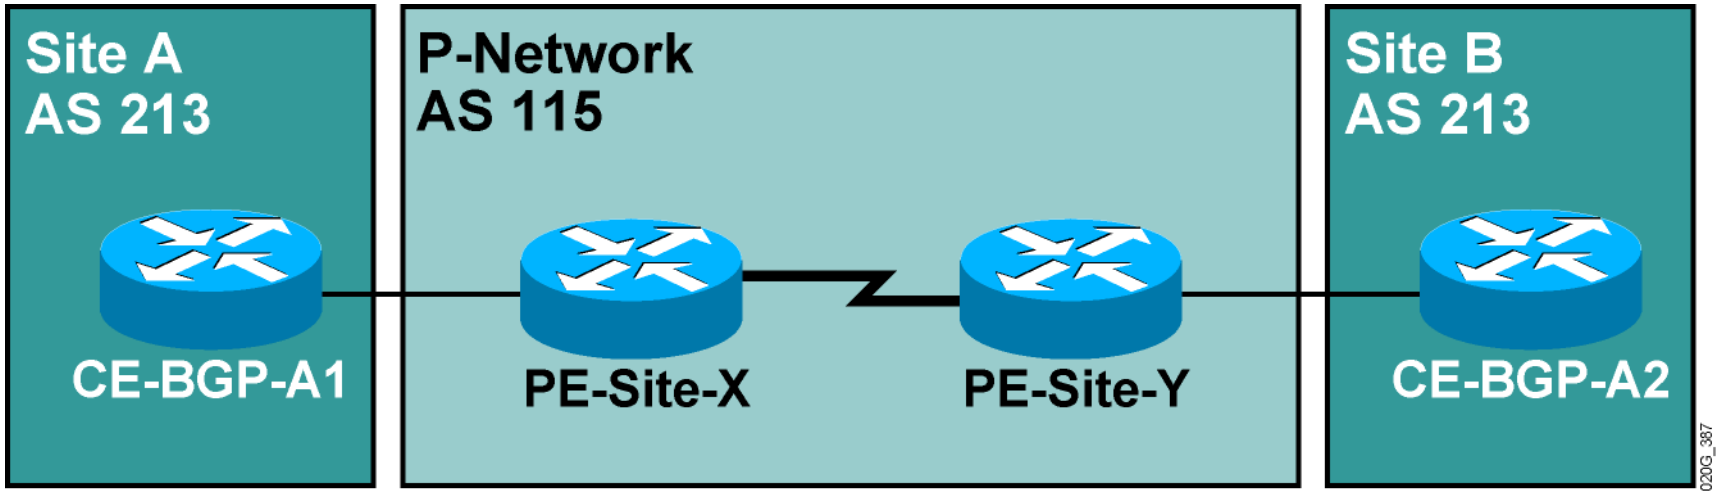

- The customer wants to reuse the same AS number on several sites:

AS-override
The Issue

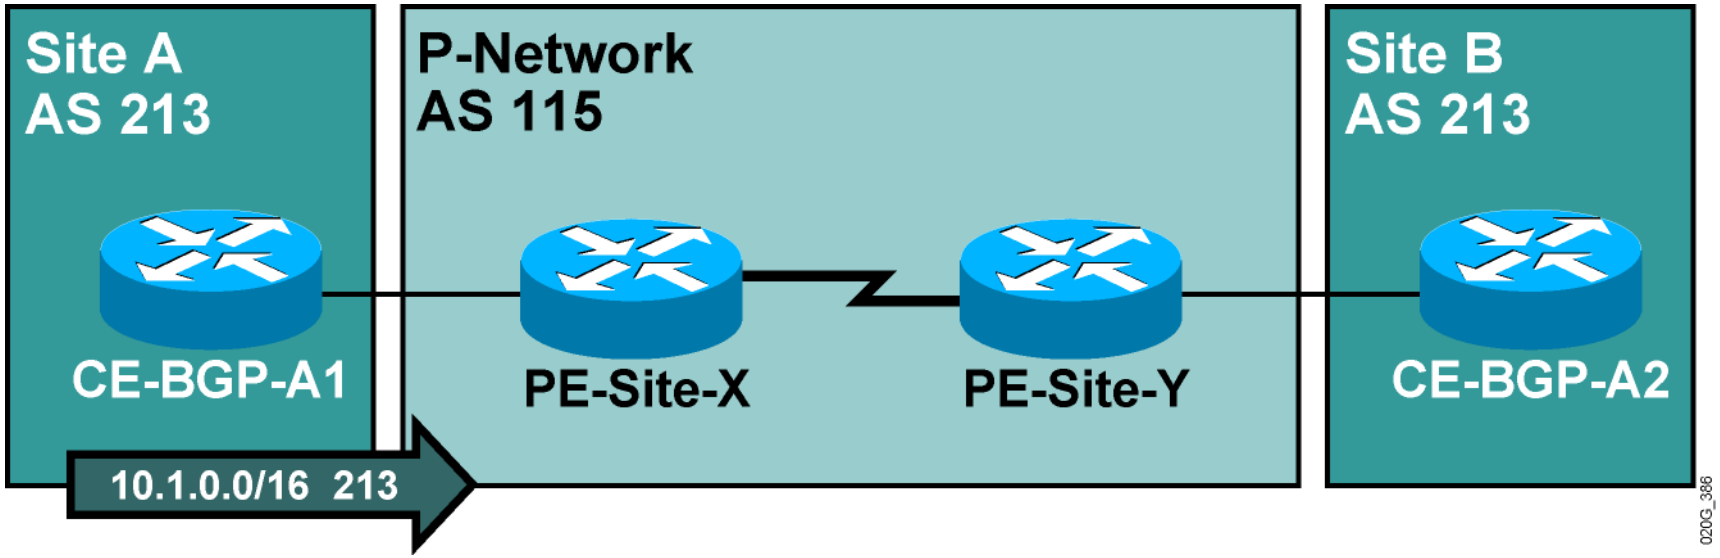

sites:

- **CE-BGP-A1 announces network 10.1.0.0/16 to PE-Site-X.**

AS-override
The Issue

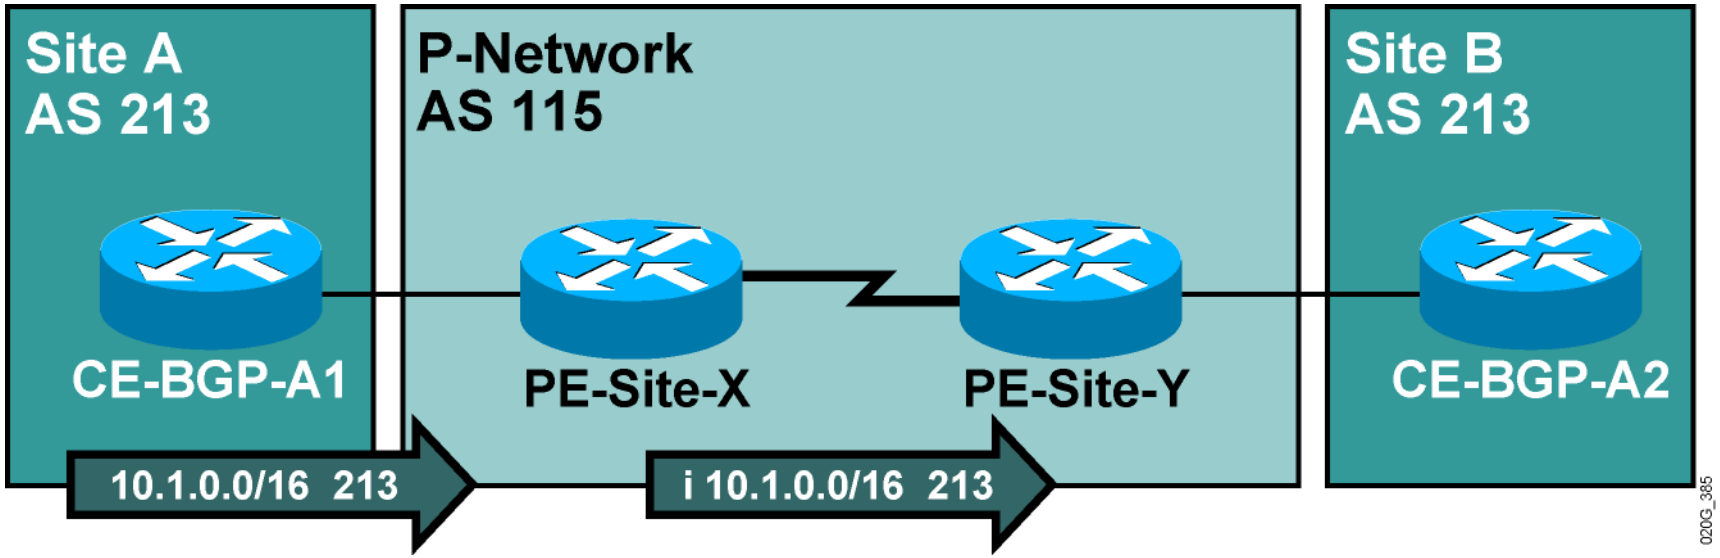

sites:

- **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.**

AS-override
The Issue

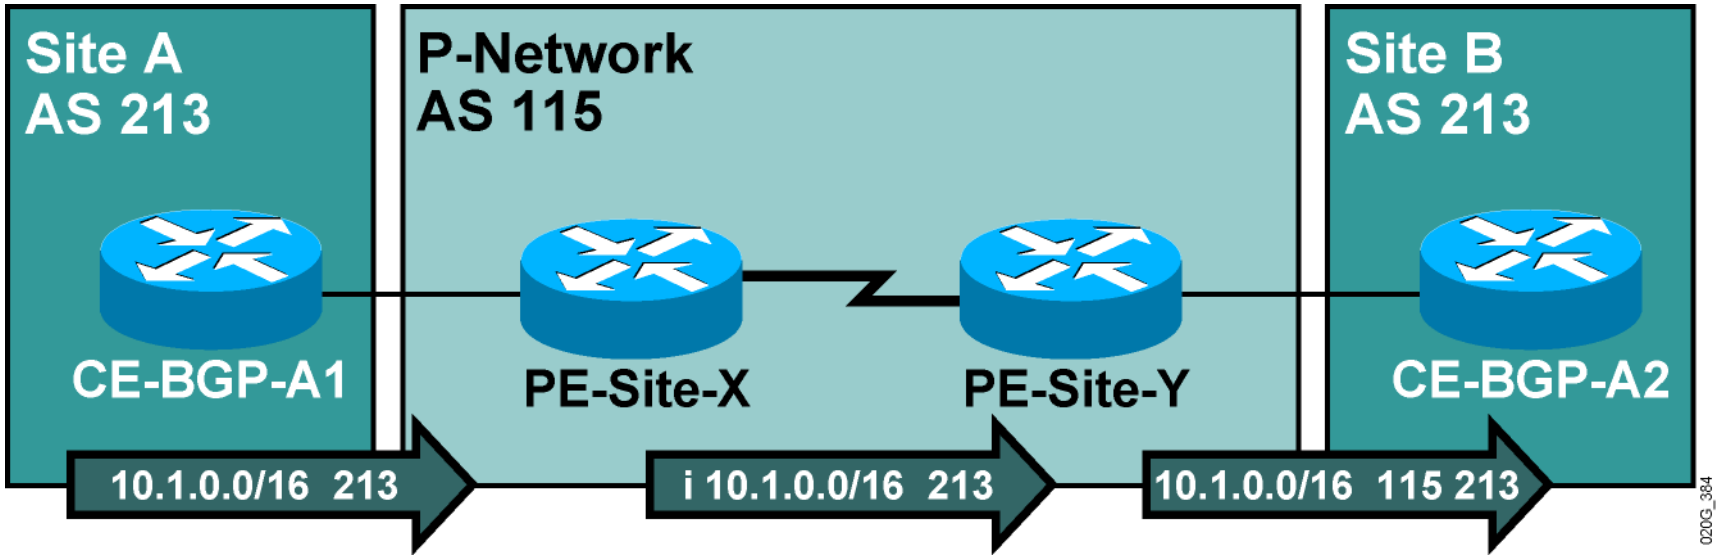

sites:

- **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

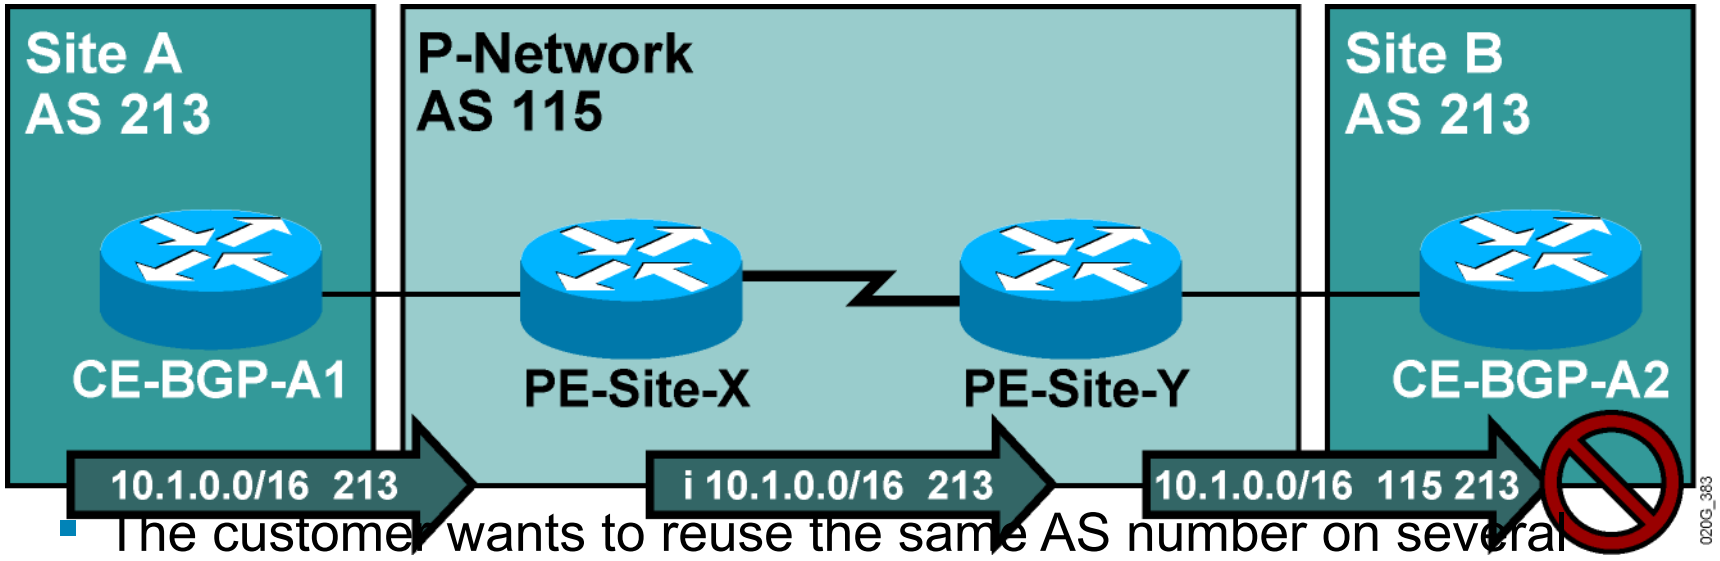

- The customer wants to reuse the same AS number on several sites:

- **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.)

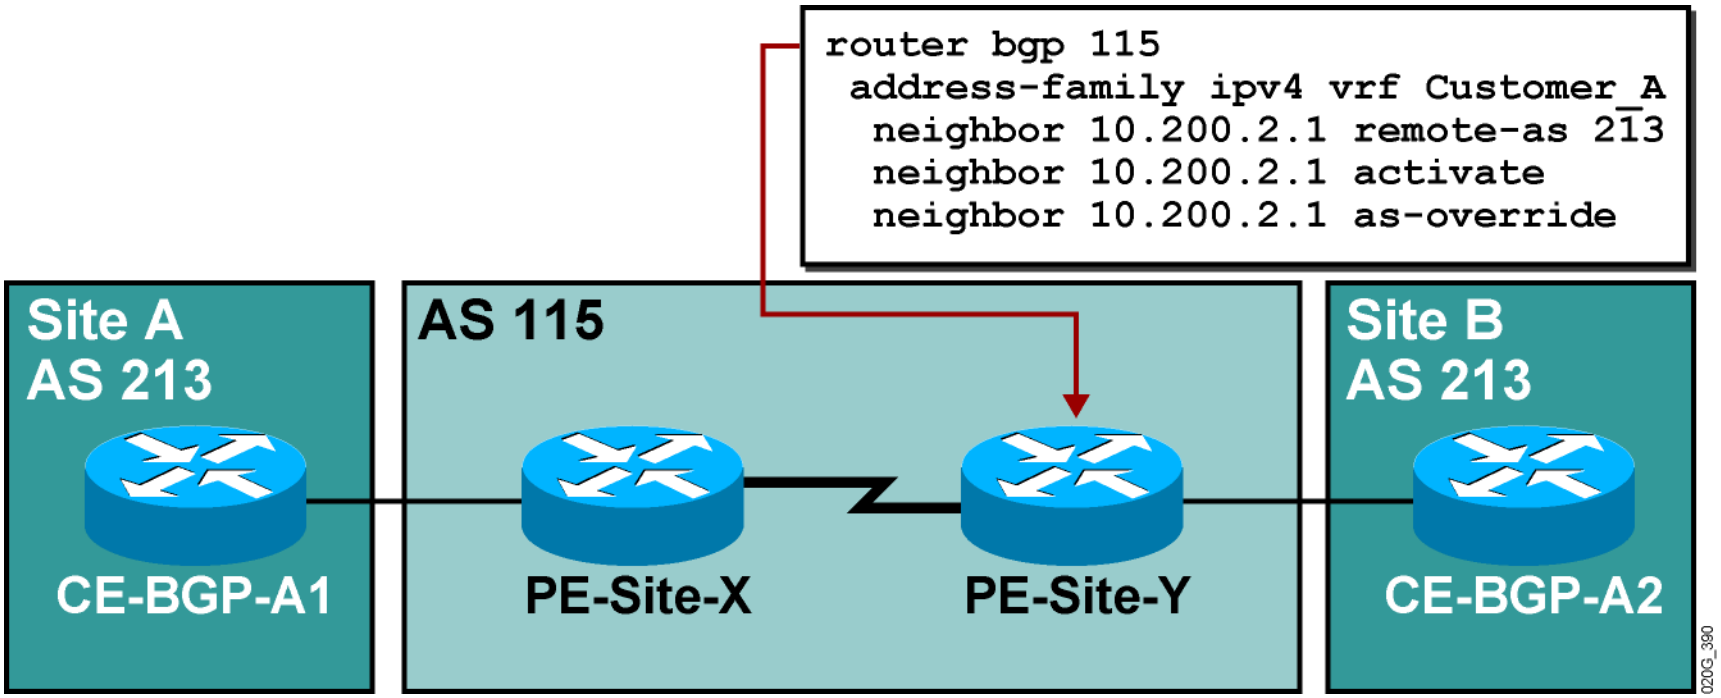

AS-override (Cont.)

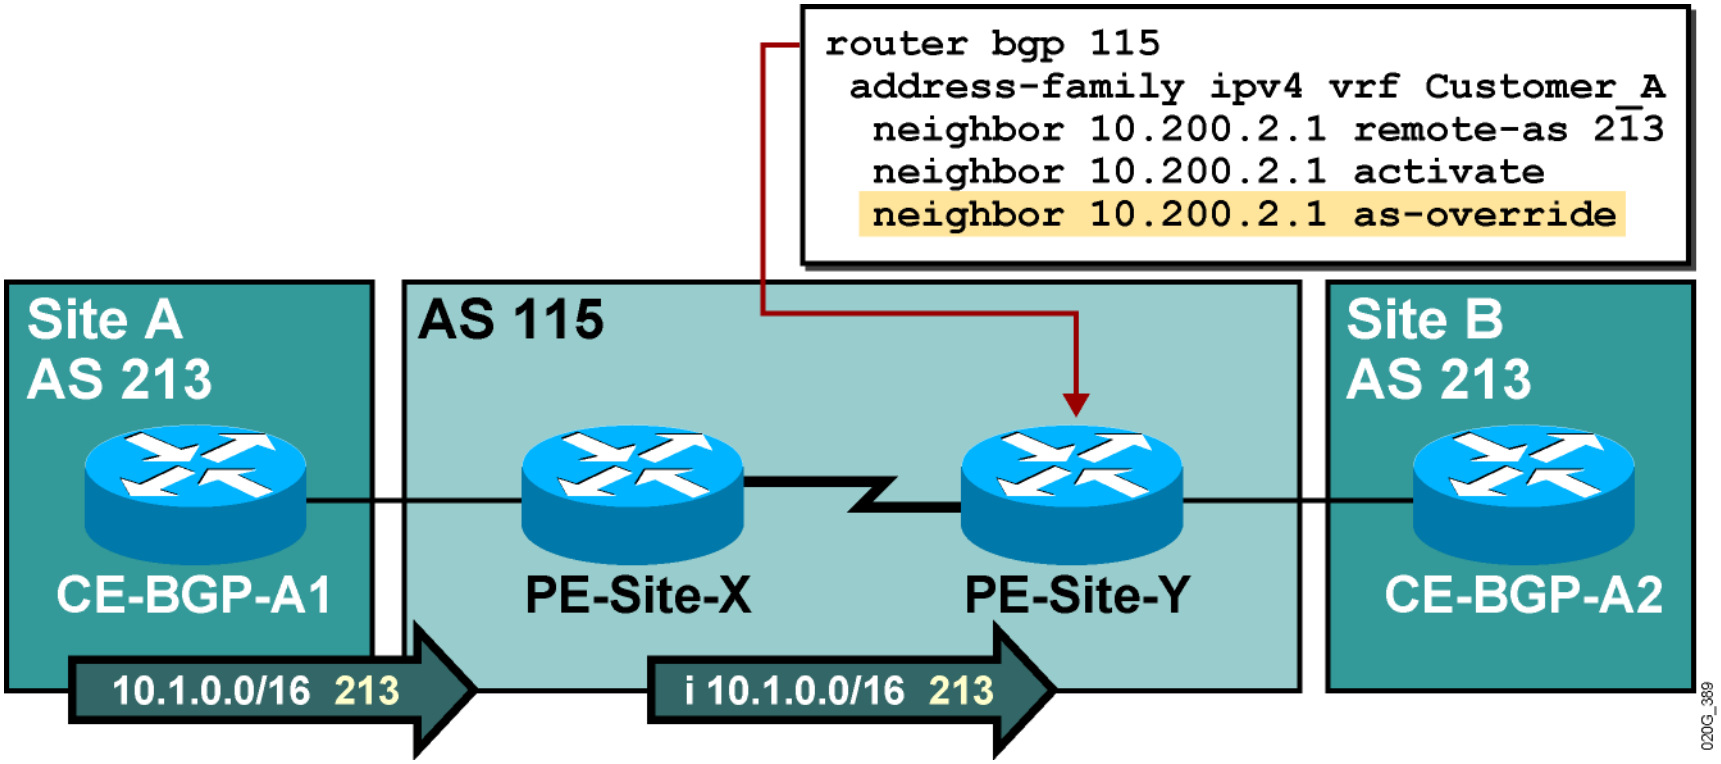

AS-override (Cont.)

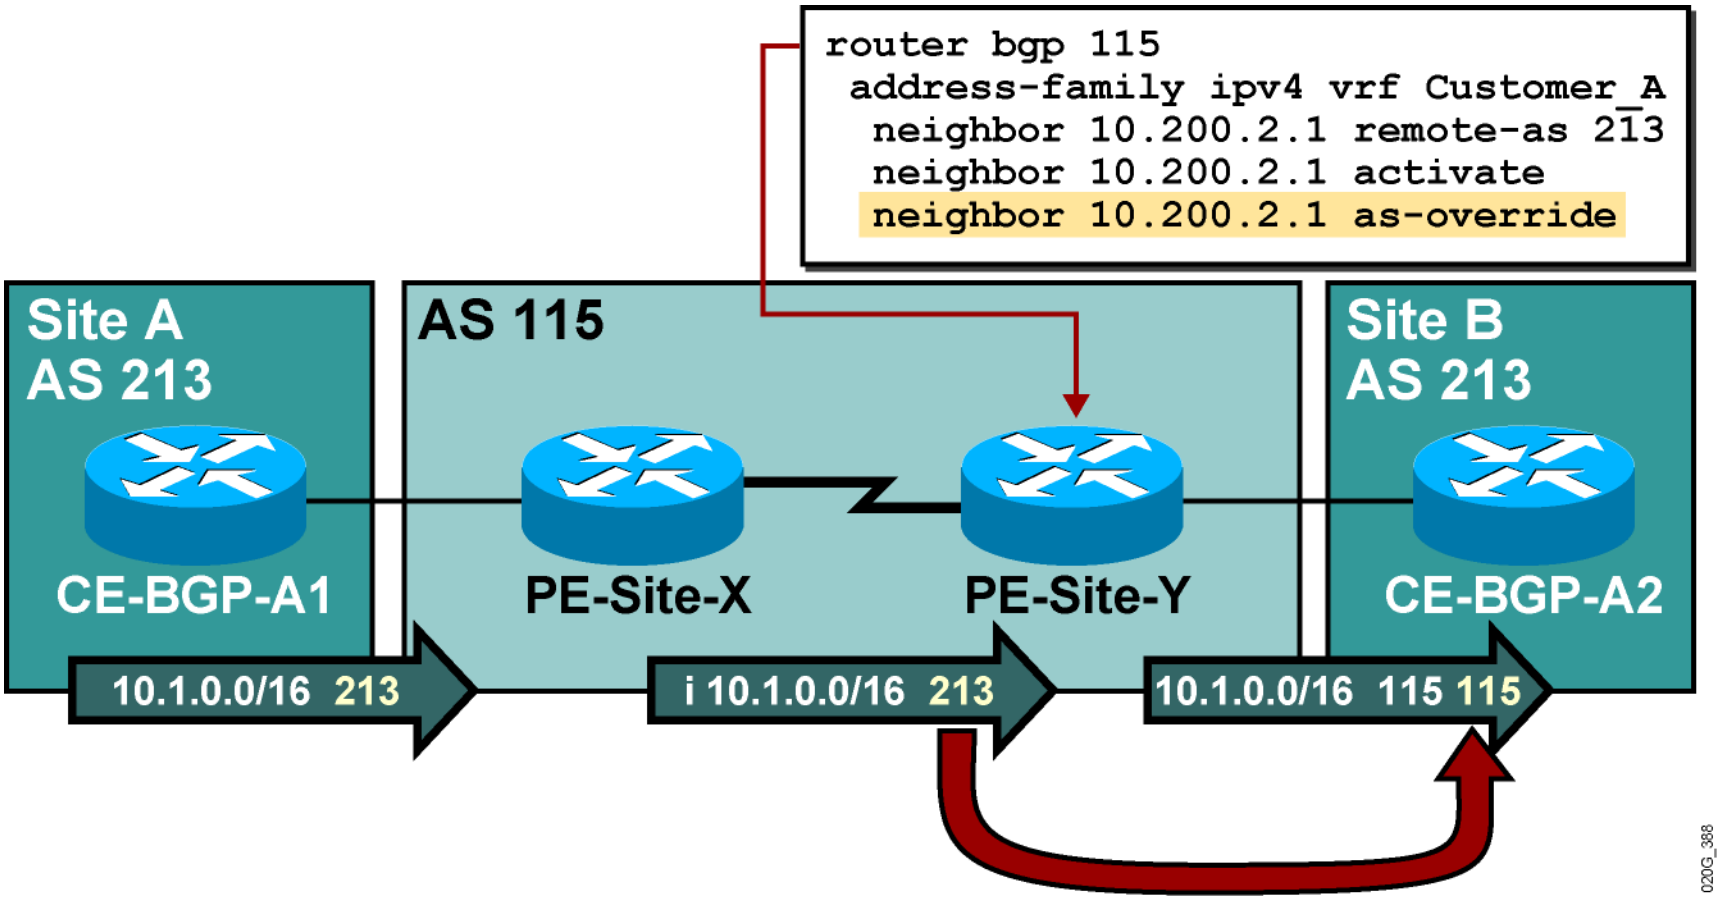

- **PE-Site-Y replaces AS 213 with AS 115 in the AS path, prepends another copy of AS115 to the AS path, and propagates the prefix.**

AS-override (Cont.)
AS-Path Prepending

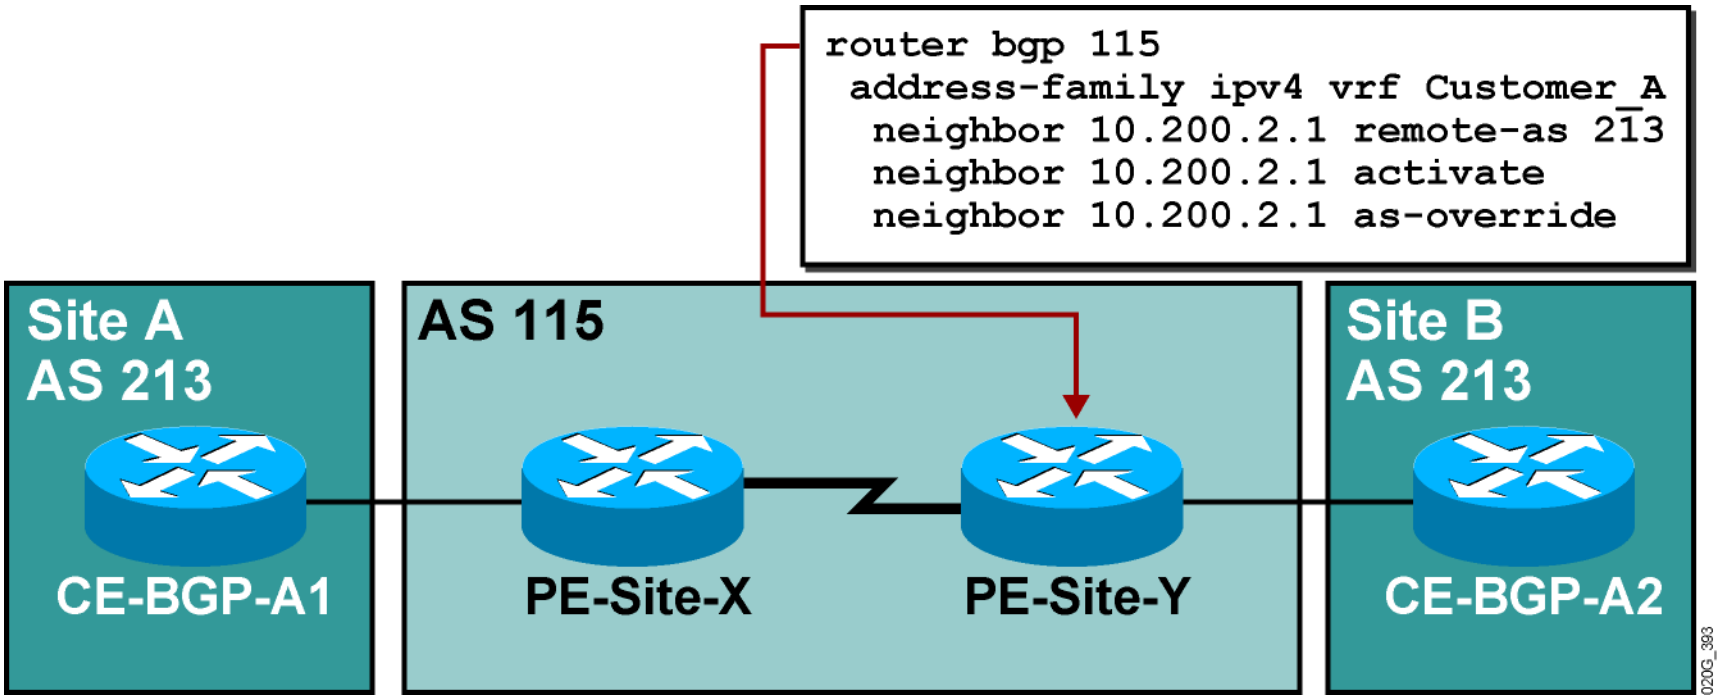

AS-override (Cont.)
AS-Path Prepending

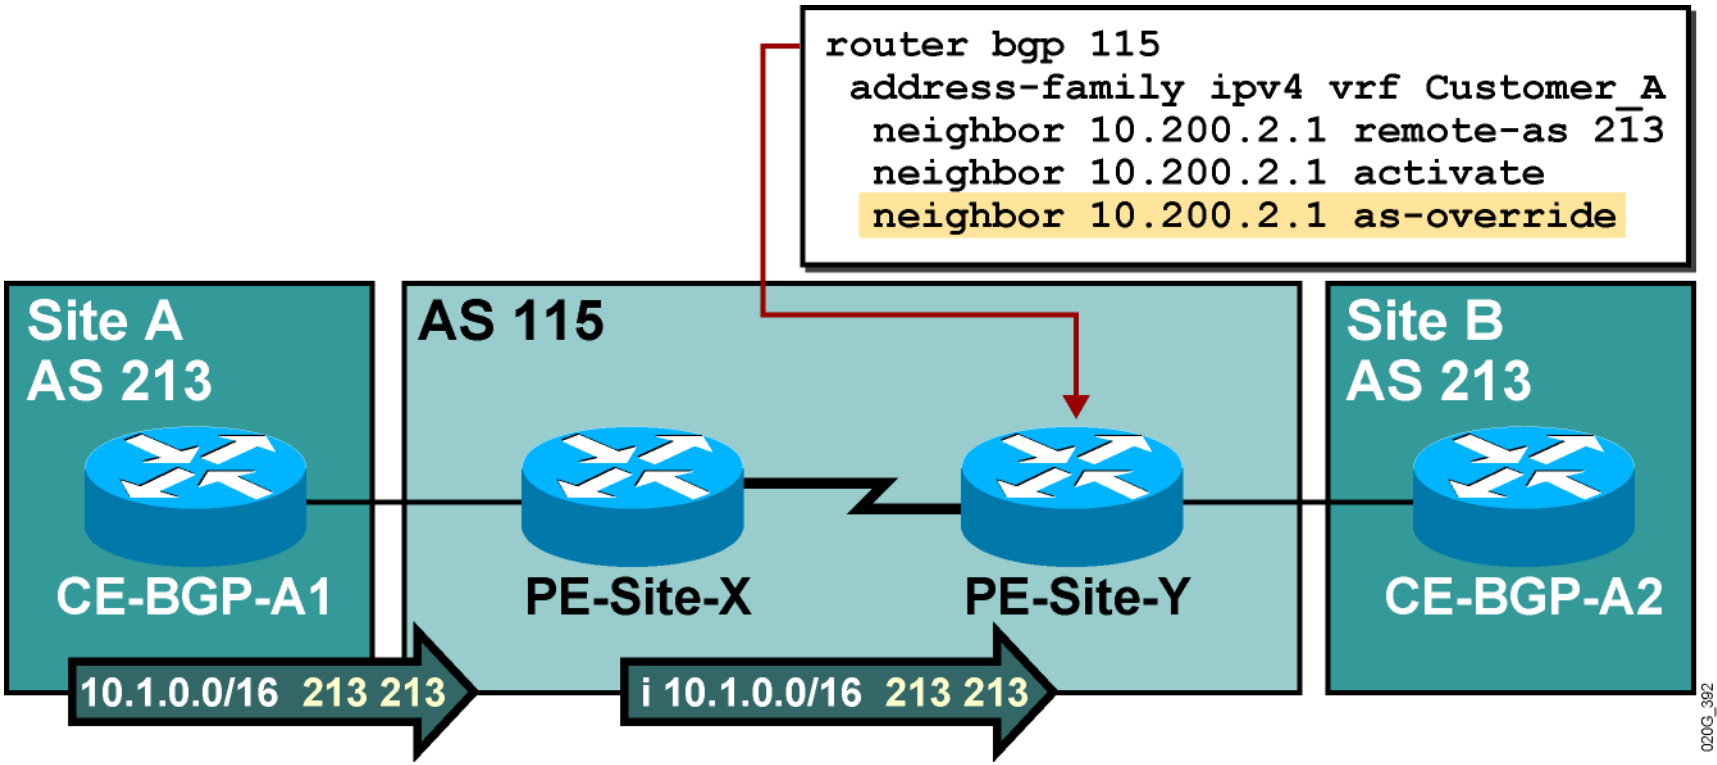

AS-override (Cont.)
AS-Path Prepending

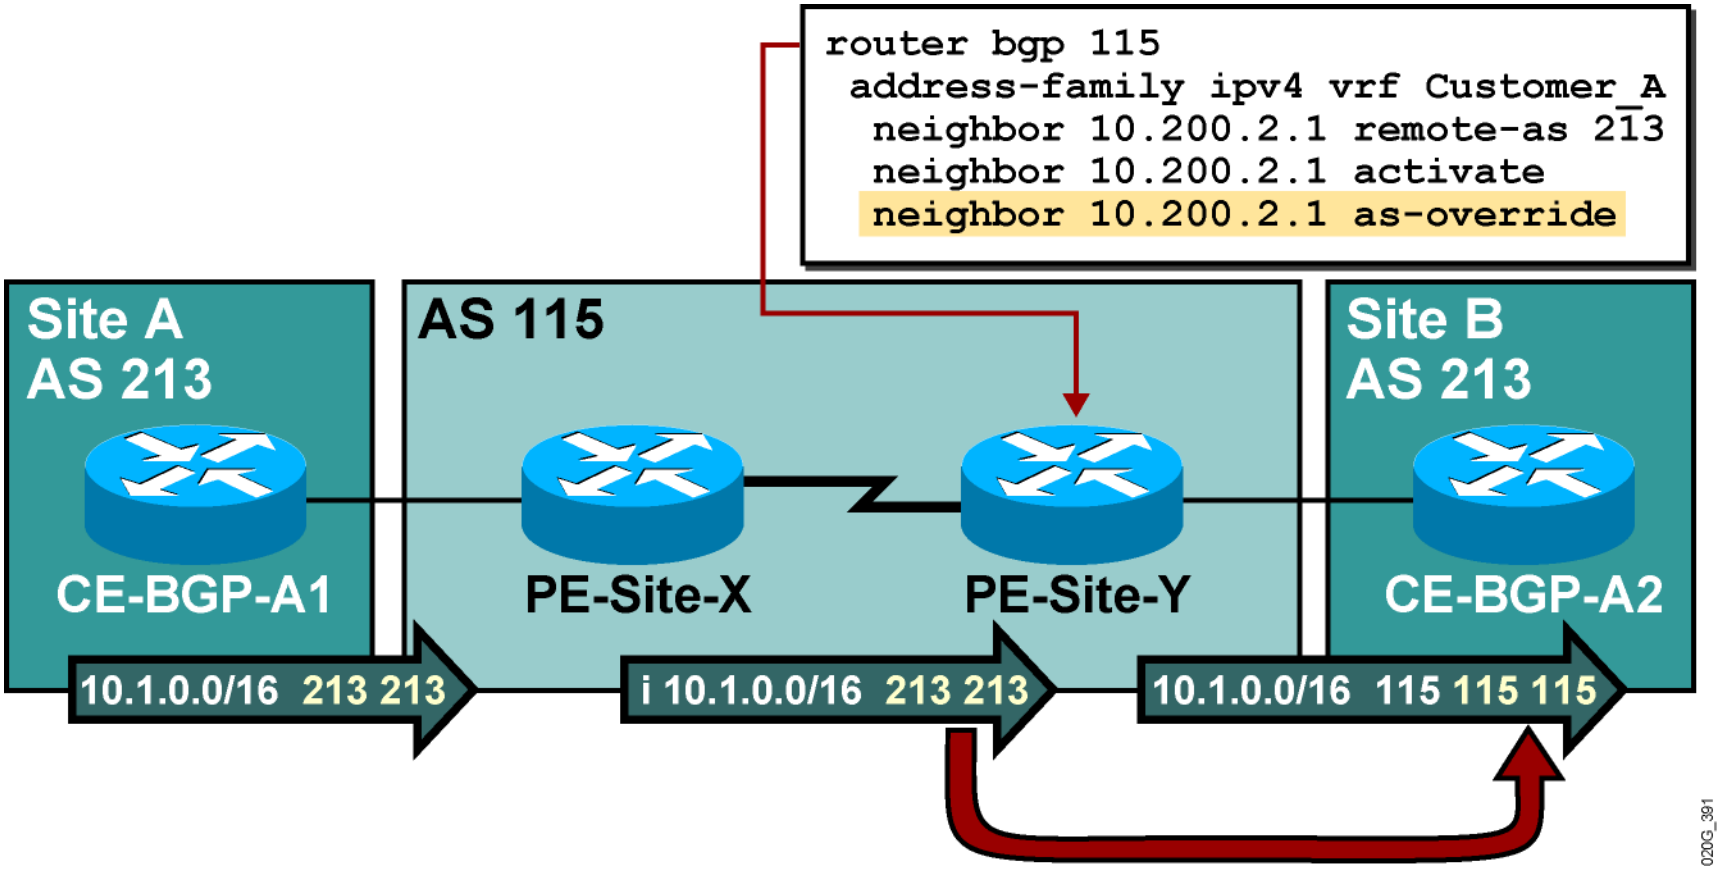

- **PE-Site-Y replaces all occurrences of AS 213 with AS 115 in the AS path, prepends another copy of AS 115 to the AS path, and propagates the prefix.**

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

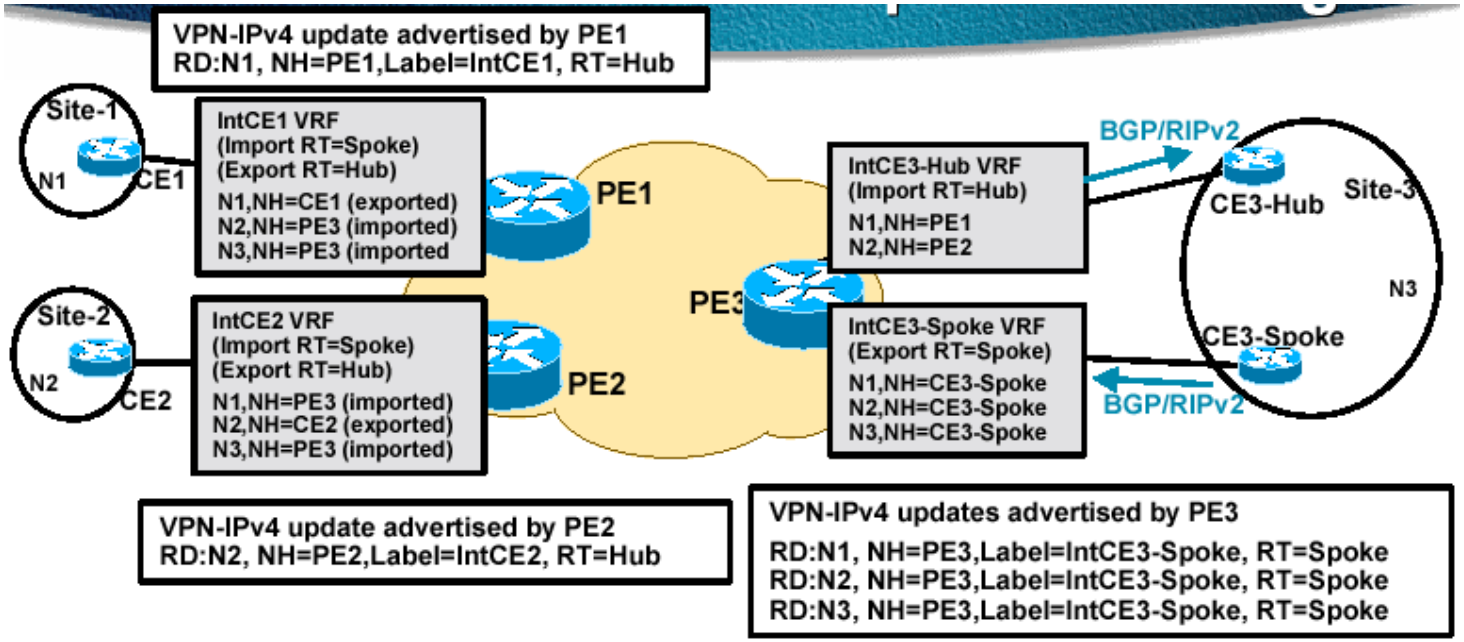

- We need 2 interfaces and 2 unique VRFs on the Hub site. If not, traffic from spokes may just touch PE3 and be forwarded to the spoke site without being processed at the hub site
- Traffic into hub comes in via one VRF (which exports routes, e.g. Spoke vrf) interface and goes out via the other (which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

VPN Sites with Hub & Spoke Routing

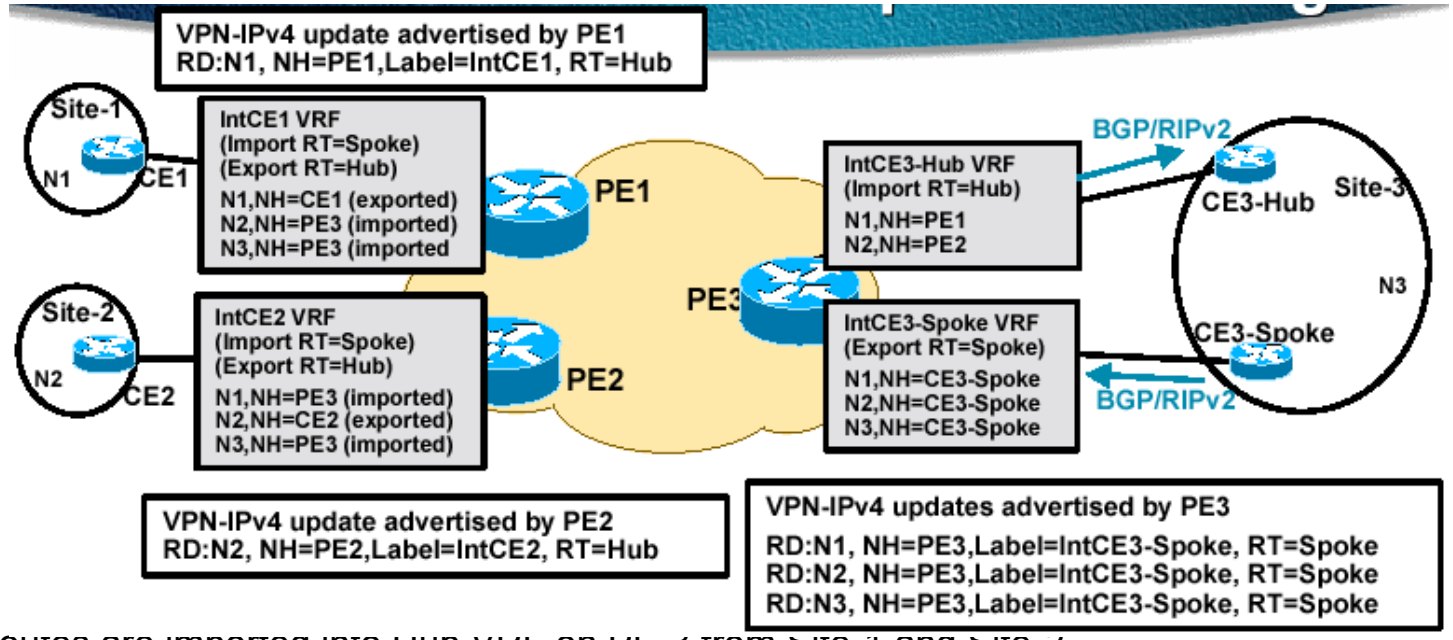

- Spoke routes are imported into Hub VRF on PE-3 from Site-1 and Site-2
- The same Spoke routes are exported to other spokes via the Spoke VRF, in which the next-hop for spoke sites to reach the other spoke site is PE3
- Since Spoke VRF at hub site exports the spoke routes, all the traffic from spokes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

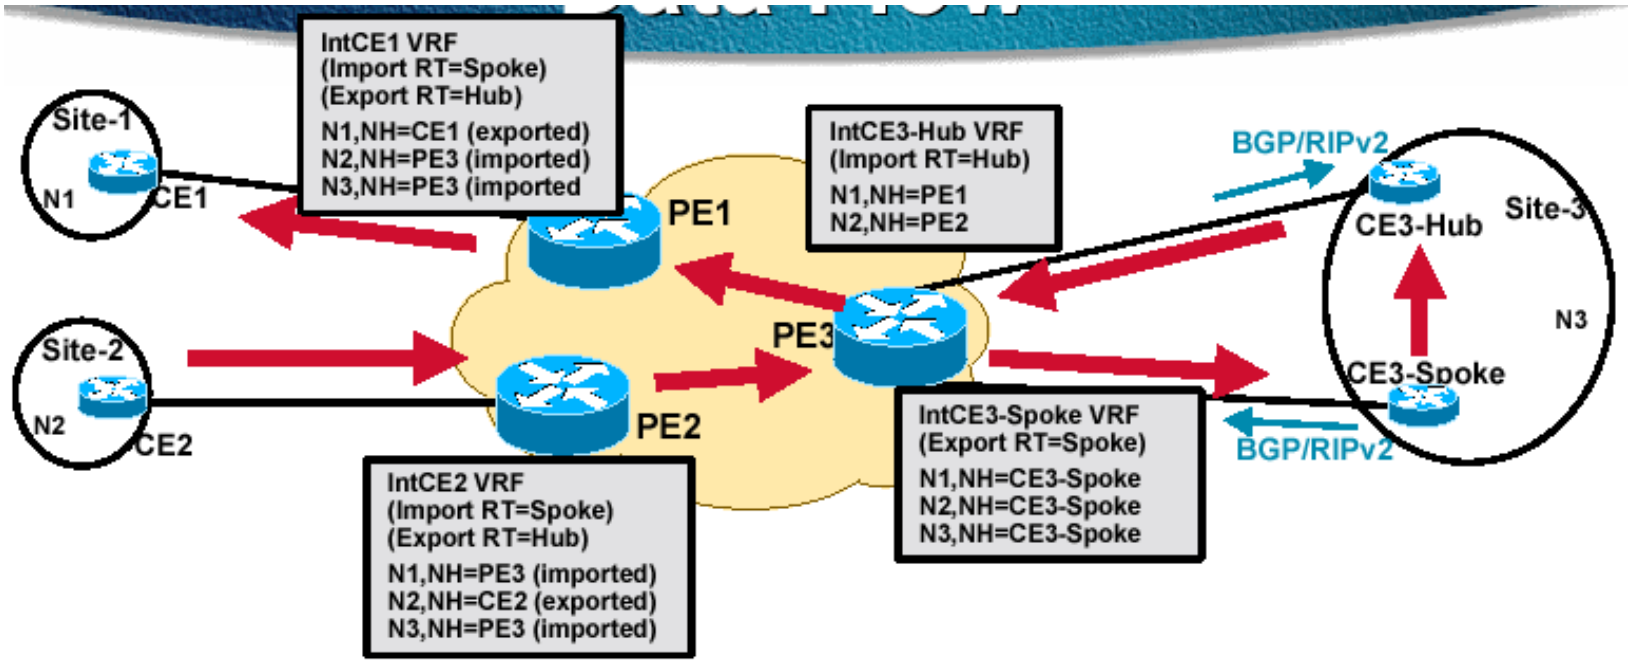

- Traffic from one spoke to another will travel across the hub site
- Allowas-in has to be configured on the PE3 if the Site-3 is using BGP

Allowas-in (for BGP updates)

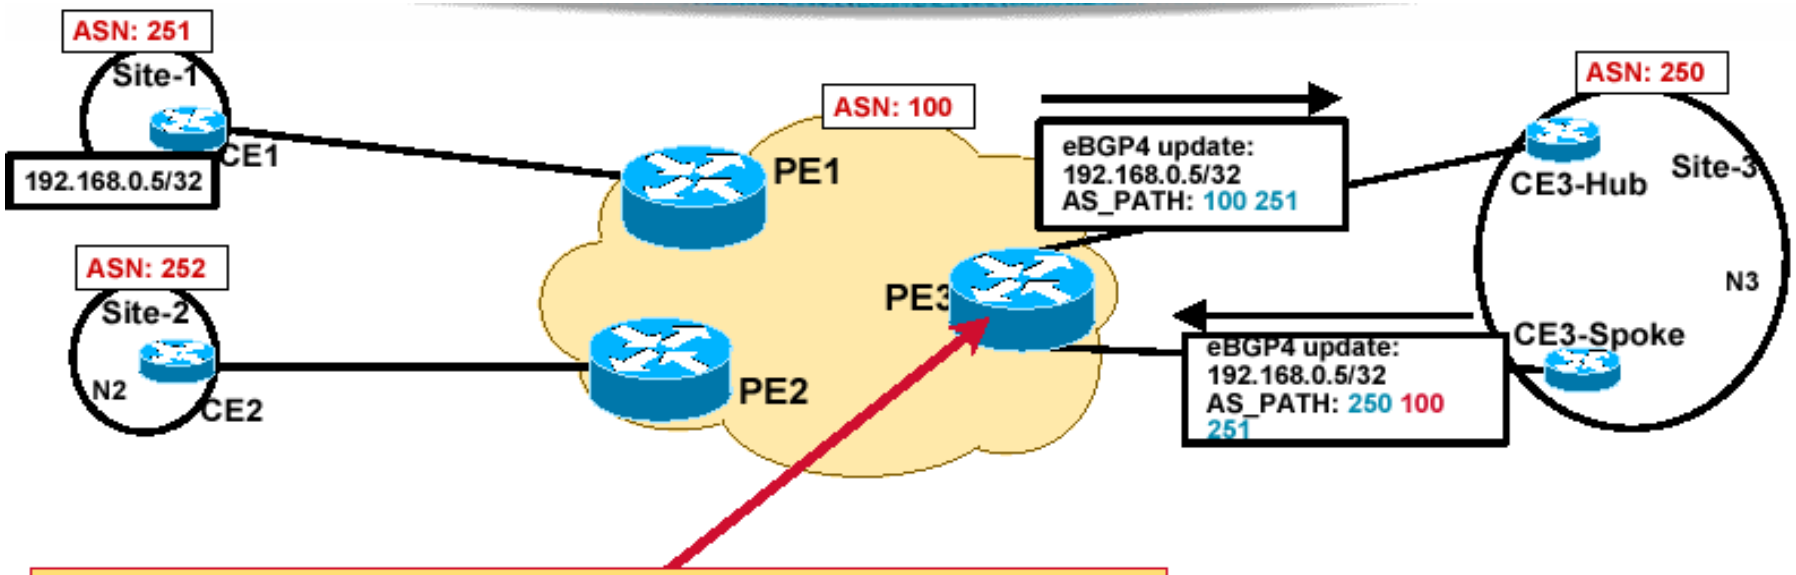


```
router bgp 100
 address-family ipv4 vrf Spoke
  neighbor 192.168.74.4 remote-as 250
  neighbor 192.168.74.4 activate
  neighbor 192.168.74.4 allowas-in 4
  no auto-summary
  no synchronization
  exit-address-family
```

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

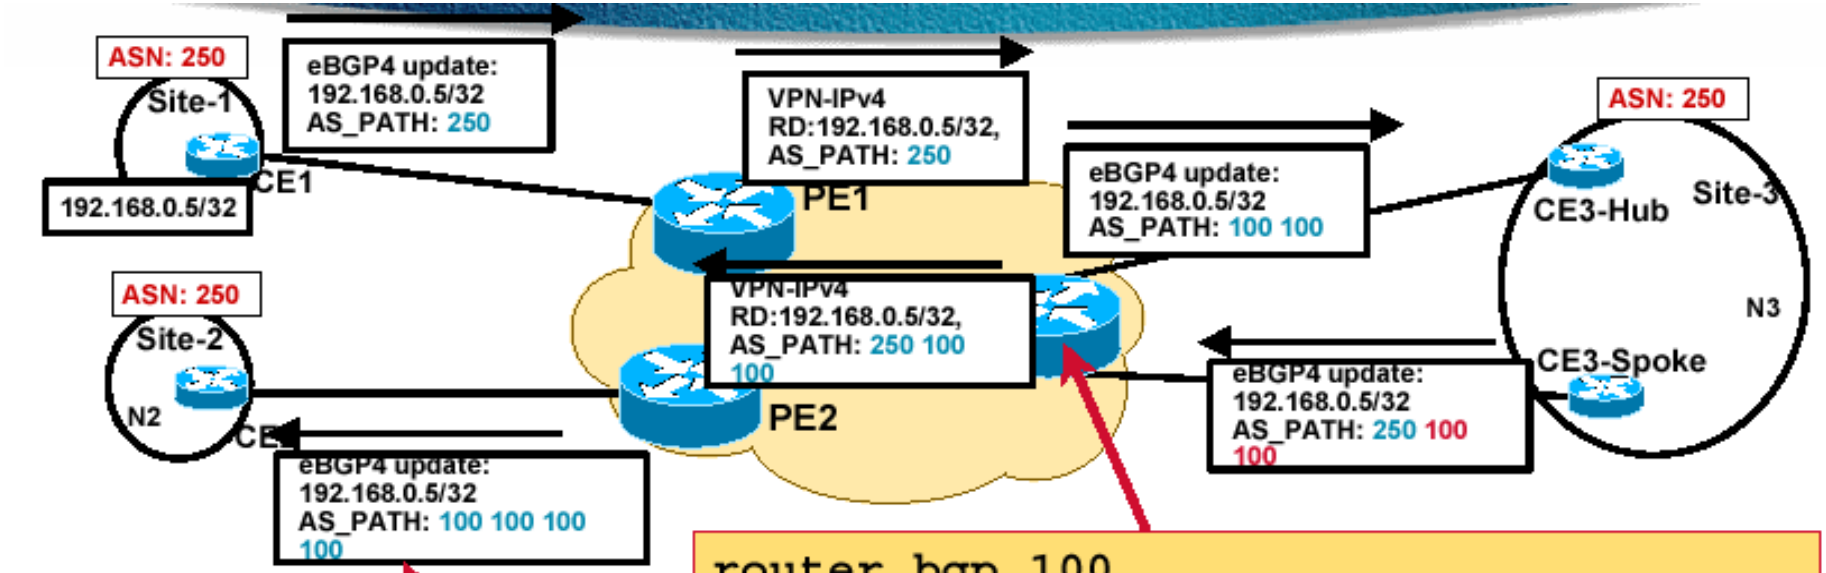

The AS_PATH contains four occurrences of the provider ASN. This update will be rejected if the CE routers advertise it back to any PE

```
router bgp 100
 address-family ipv4 vrf Hub
  neighbor 192.168.73.3 remote-as 250
  neighbor 192.168.73.3 activate
  neighbor 192.168.73.3 as-override
 address-family ipv4 vrf Spoke
  neighbor 192.168.74.4 remote-as 250
  neighbor 192.168.74.4 activate
  neighbor 192.168.74.4 allowas-in 4
```

Implementing SOO for Loop Prevention

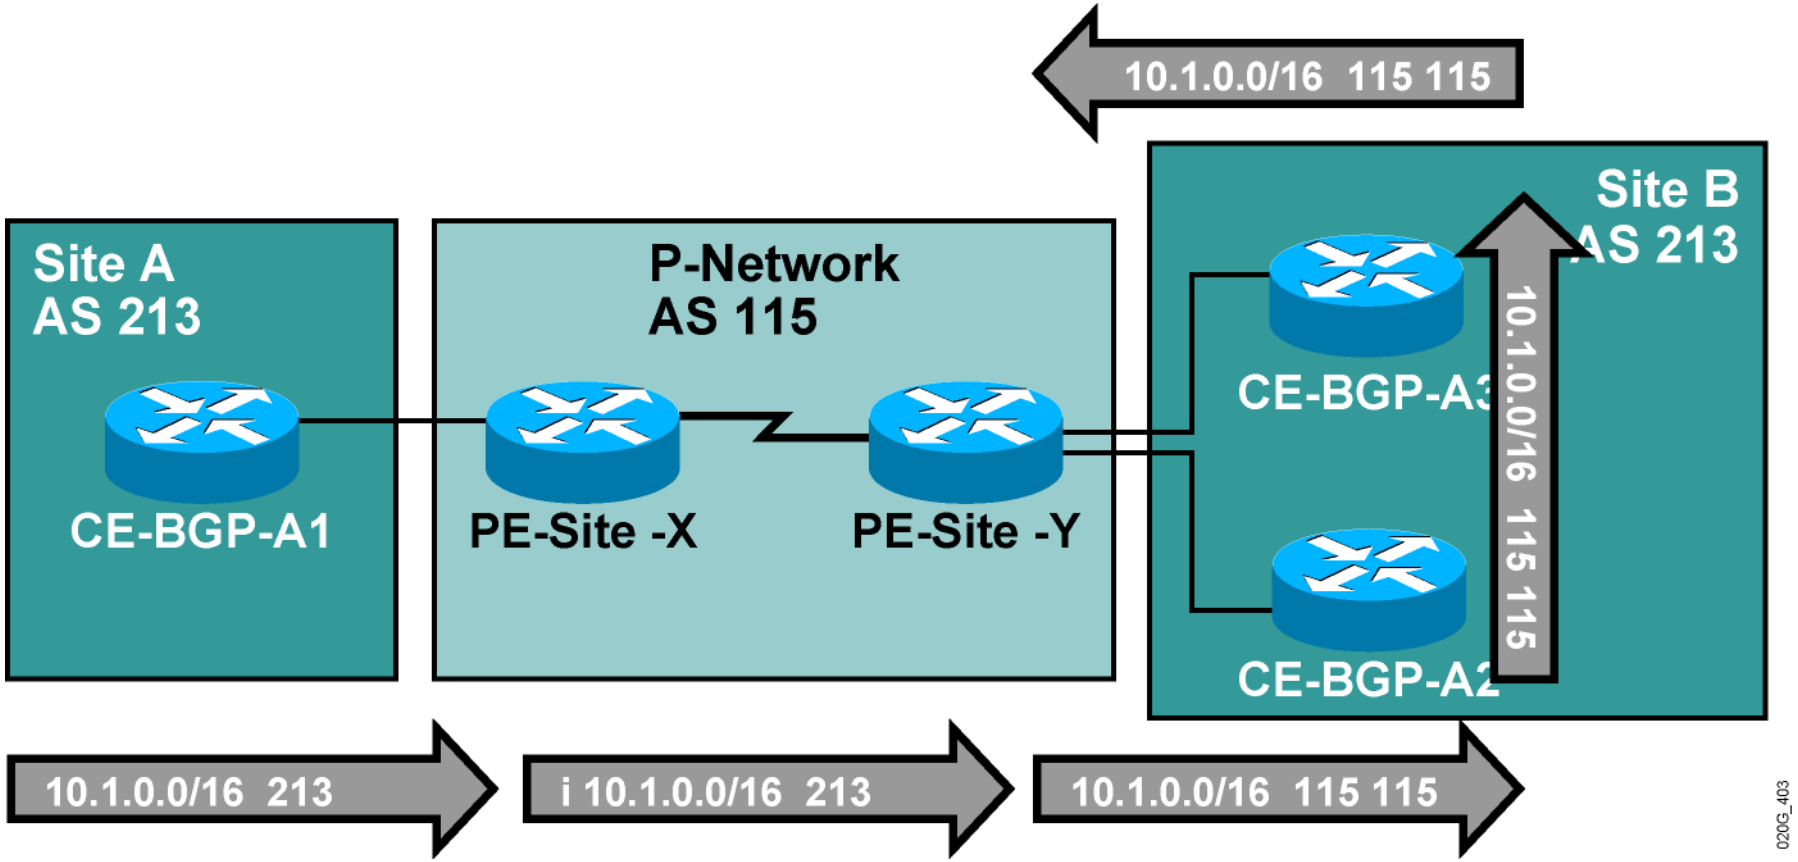

- AS path-based BGP loop prevention is bypassed with AS-override and allowas-in features.

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 conf gured in the VRF **and** the route is accepted by the route map.

Configuring Selective VRF Import (Cont.)

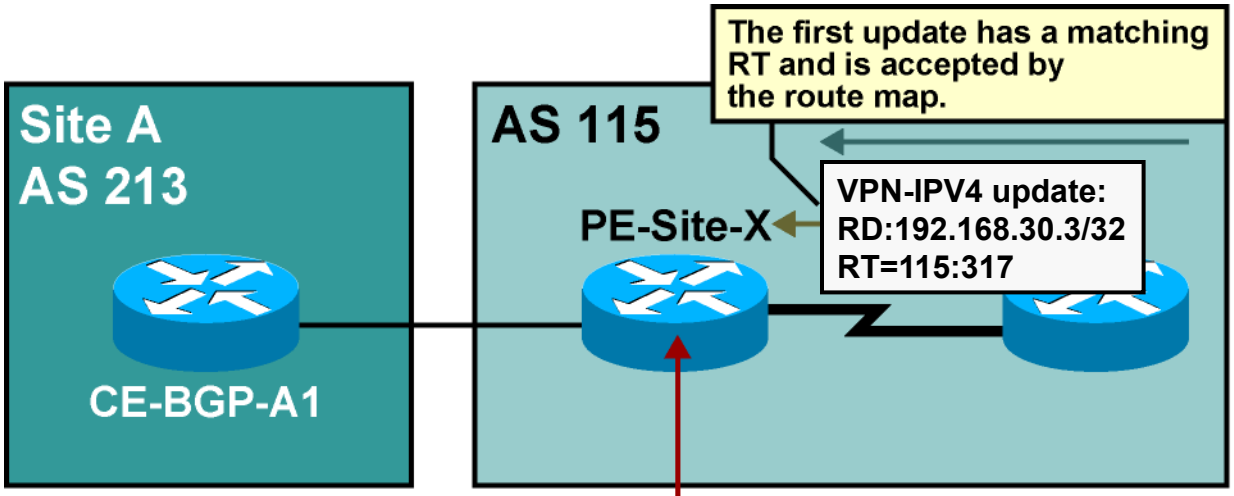


```
ip vrf Site_A
 rd 115:317
 import map RTMAP
 route-target both 115:317
!
access-list 20 permit 192.168.30.0 0.0.0.255
!
route-map RTMAP permit 10
  match ip address 20
```

020G_084

Configuring Selective VRF Import (Cont.)

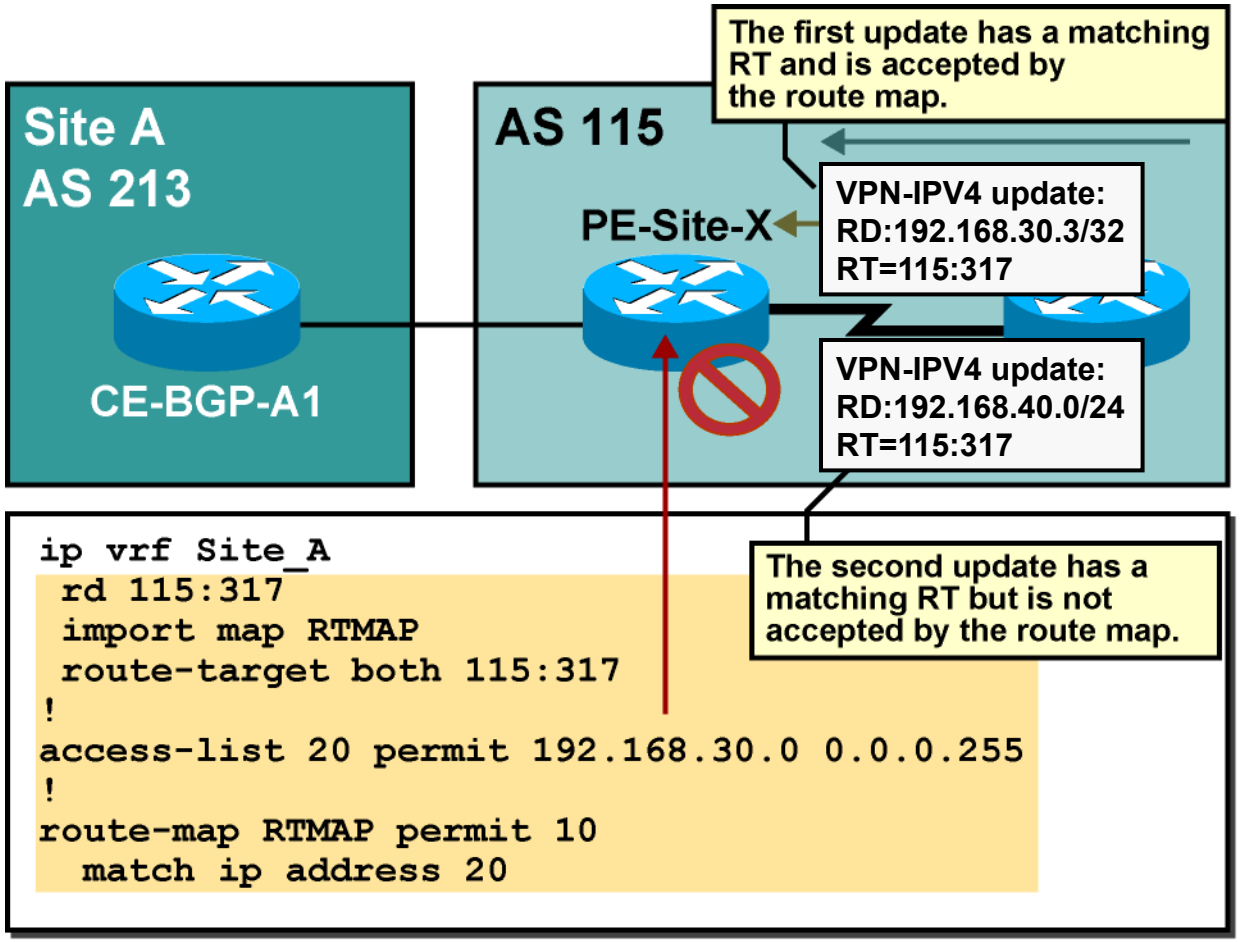

Configuring Selective VRF Export

- 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 conf gured 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.)

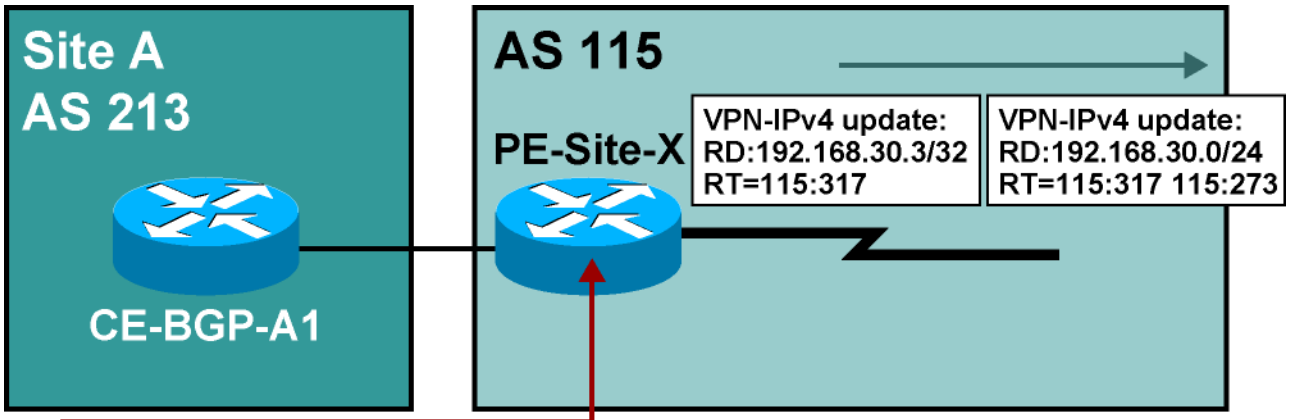


```
ip vrf Site_A
 rd 115:317
 export map RTMAP
 route-target both 115:317
!
access-list 10 permit 192.168.30.0 0.0.0.0
!
route-map RTMAP permit 20
 match ip address 10
 set extcommunity rt 115:273 additive
```

020G_066

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.

MPLS Boot camp

Configuring RIP as a Routing Protocol Between PE and CE Routers

Outline

Configuring RIP PE-CE Routing

Avoiding Routing Loops with RIP as PE-CE Protocol

Configuring RIP PE-CE Routing

- A routing context is configured for each VRF running RIP.
- RIP parameters have to be specified in the VRF.
- Some parameters configured in the RIP process are propagated to routing contexts
  (for example, RIP version).
- Only RIPv2 is supported.
- RIP may work but does not support VLSM (Variable Length Subnet Mask)

Configuring RIP PE-CE Routing (Cont.)
RIP Metric Propagation

```
router rip
 version 2
  address-family ipv4 vrf vrf-name
    version 2
    redistribute bgp as-number metric transparent
```

- BGP routes must be redistributed back into RIP.
- The RIP hop count has to be manually set for routes redistributed into RIP.
- For end-to-end RIP networks, the following applies:
  - On the sending end, the RIP hop count is copied into the BGP multi-exit discriminator attribute (default BGP behavior).
  - On the receiving end, the **metric transparent** option copies the BGP MED into the RIP hop count, resulting in a consistent end-to-end RIP hop count. This hop count does not have the hops traversed via the MPLS VPN backbone
- When you are using RIP with other protocols, the metric must be manually set.

Configuring RIP PE-CE Routing (Cont.)

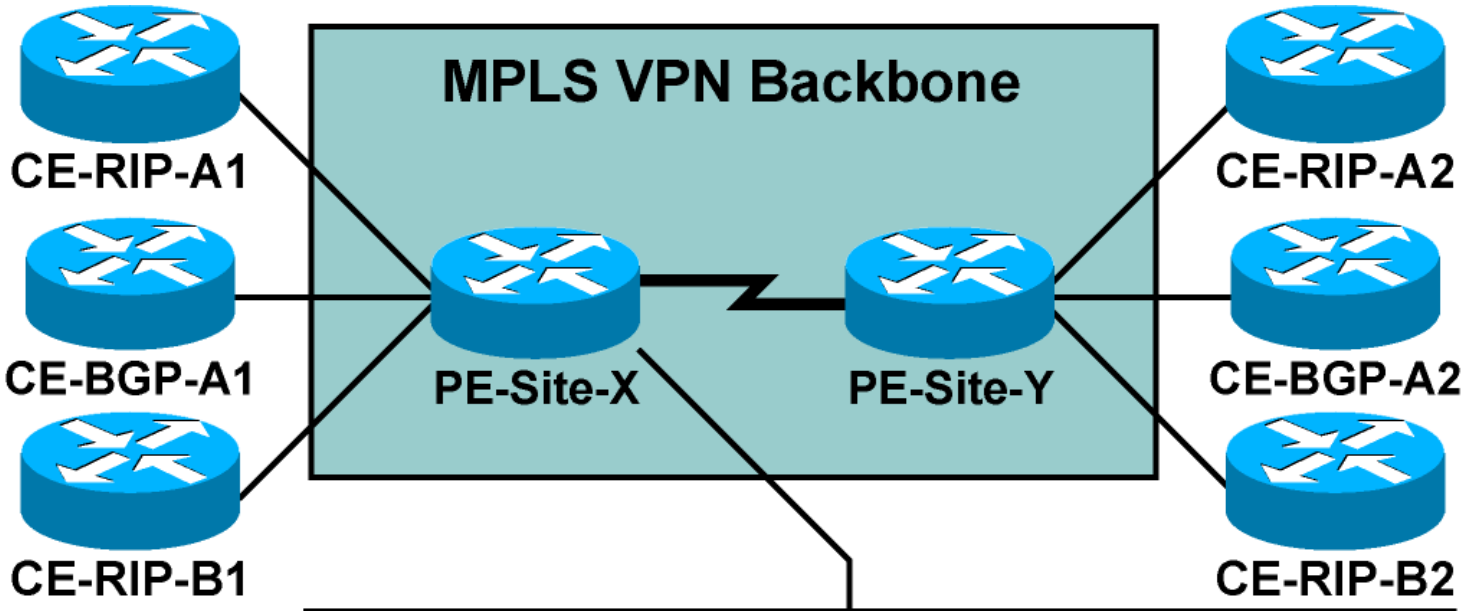


```
router rip
 version 2
 address-family ipv4 vrf Customer_ABC
  network 10.0.0.0
  redistribute bgp 12703 metric transparent
!
router bgp 12703
 address-family ipv4 vrf Customer_ABC
  redistribute rip
  no auto-summary
```

020G_073

Loop Detection with RIP as PE-CE

RIP works with the following mechanisms for loop detection:

- Split Horizon
- Site Of Origin (SOO)

Avoiding Routing Loops: Split-horizon

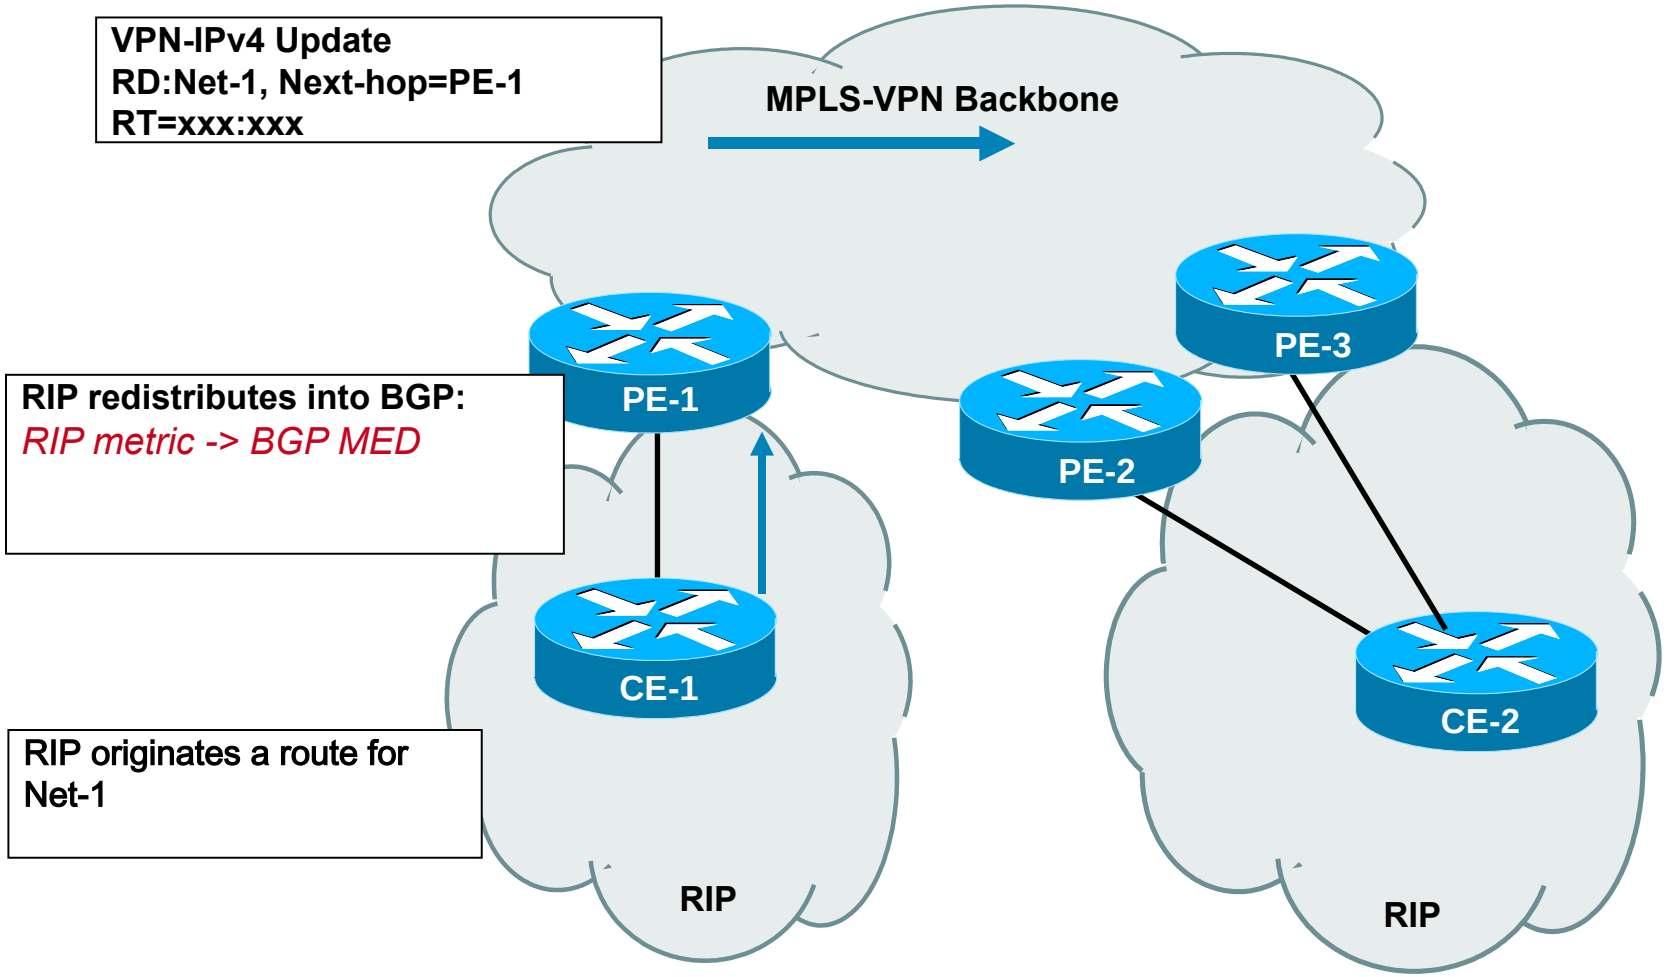

Avoiding Routing Loops: Split-horizon

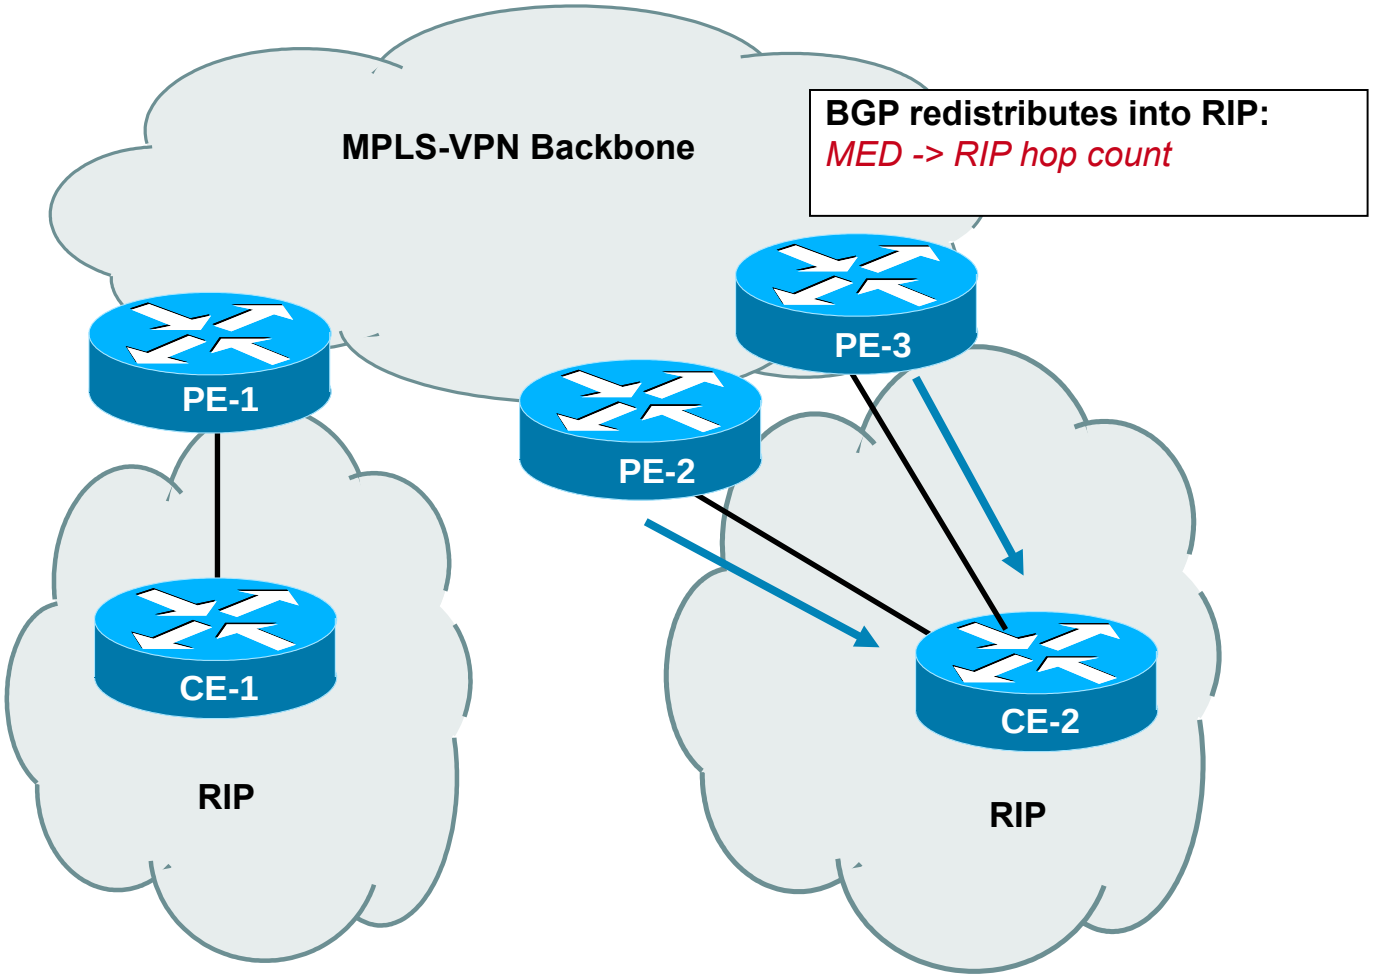

Avoiding Routing Loops: Split-horizon

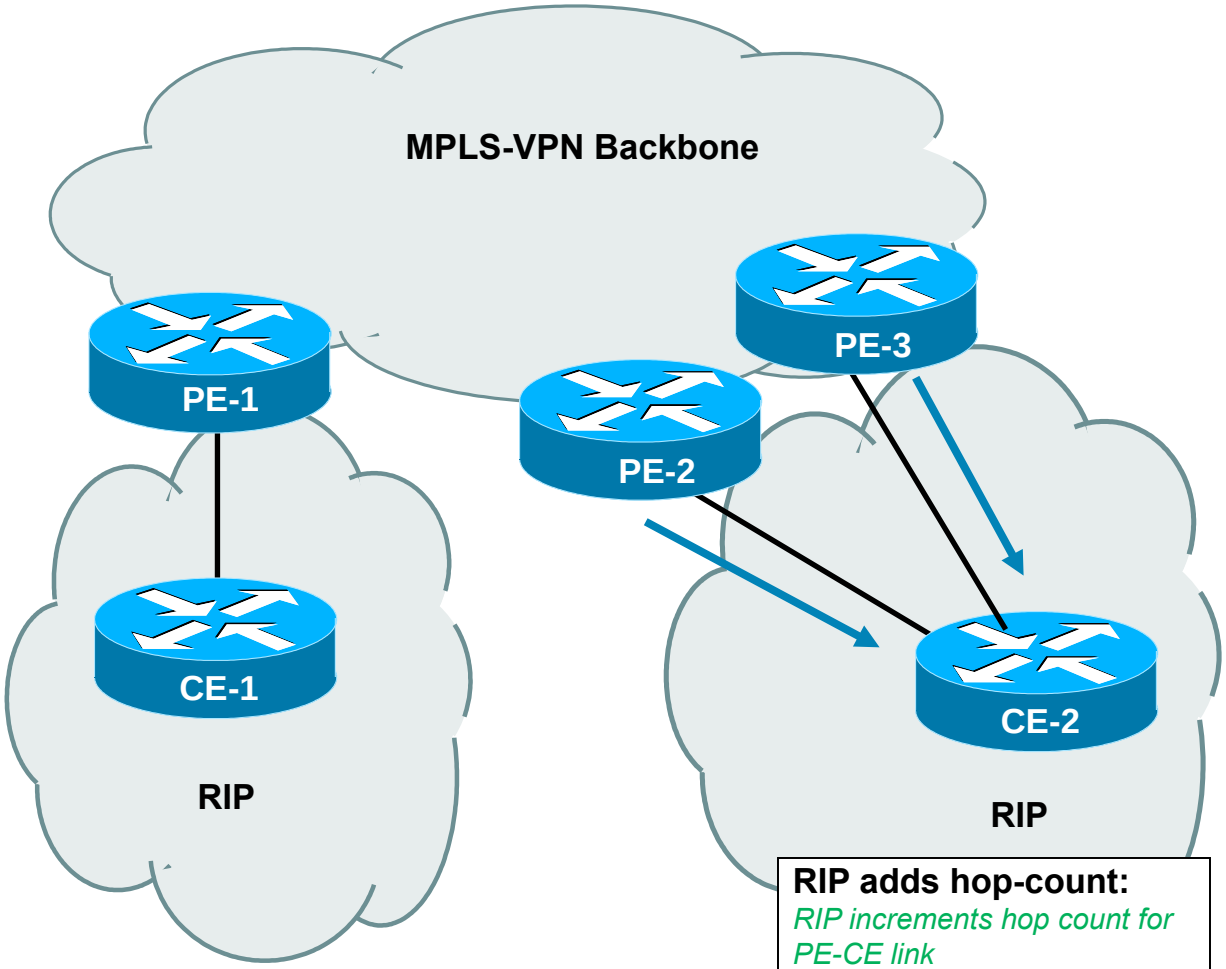

RIP adds hop-count:
RIP increments hop count for PE-CE link

Avoiding Routing Loops: Split-horizon

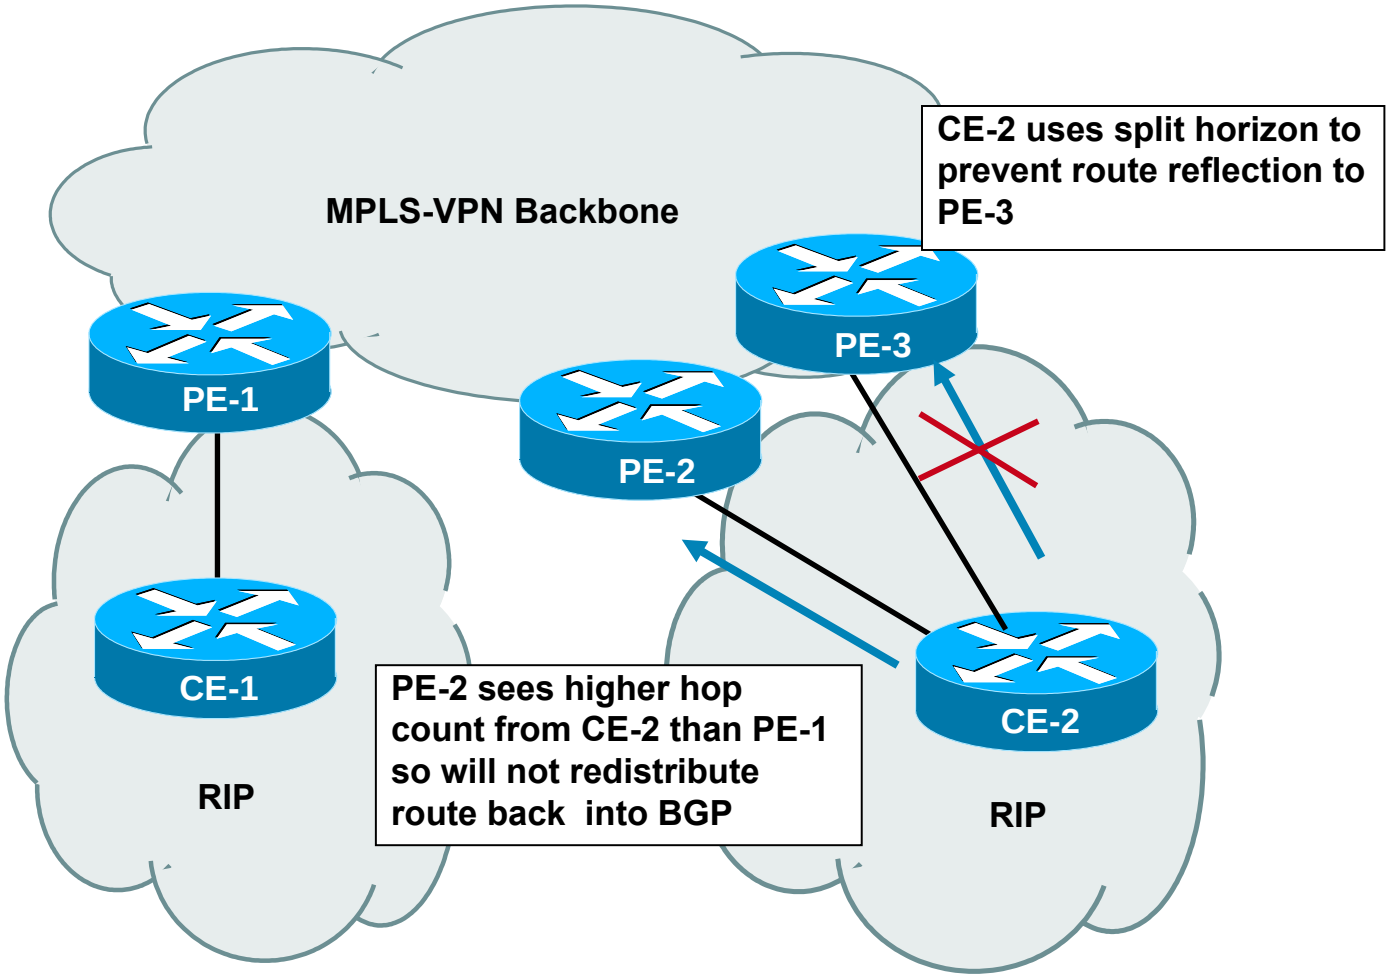

Avoiding Routing Loops: Site Of Origin

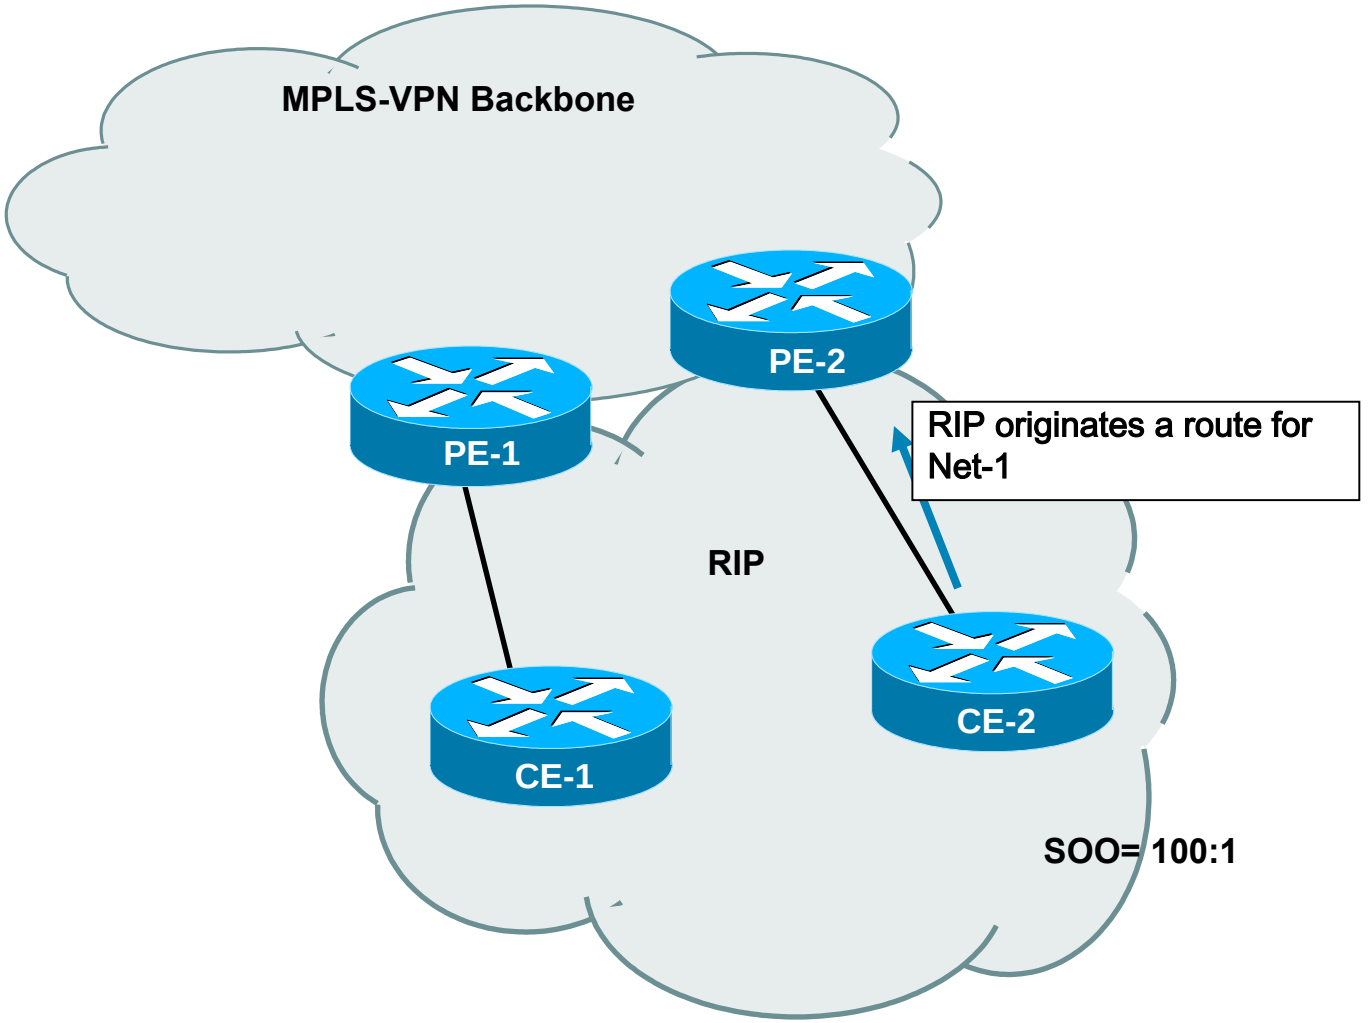

Avoiding Routing Loops: Site Of Origin

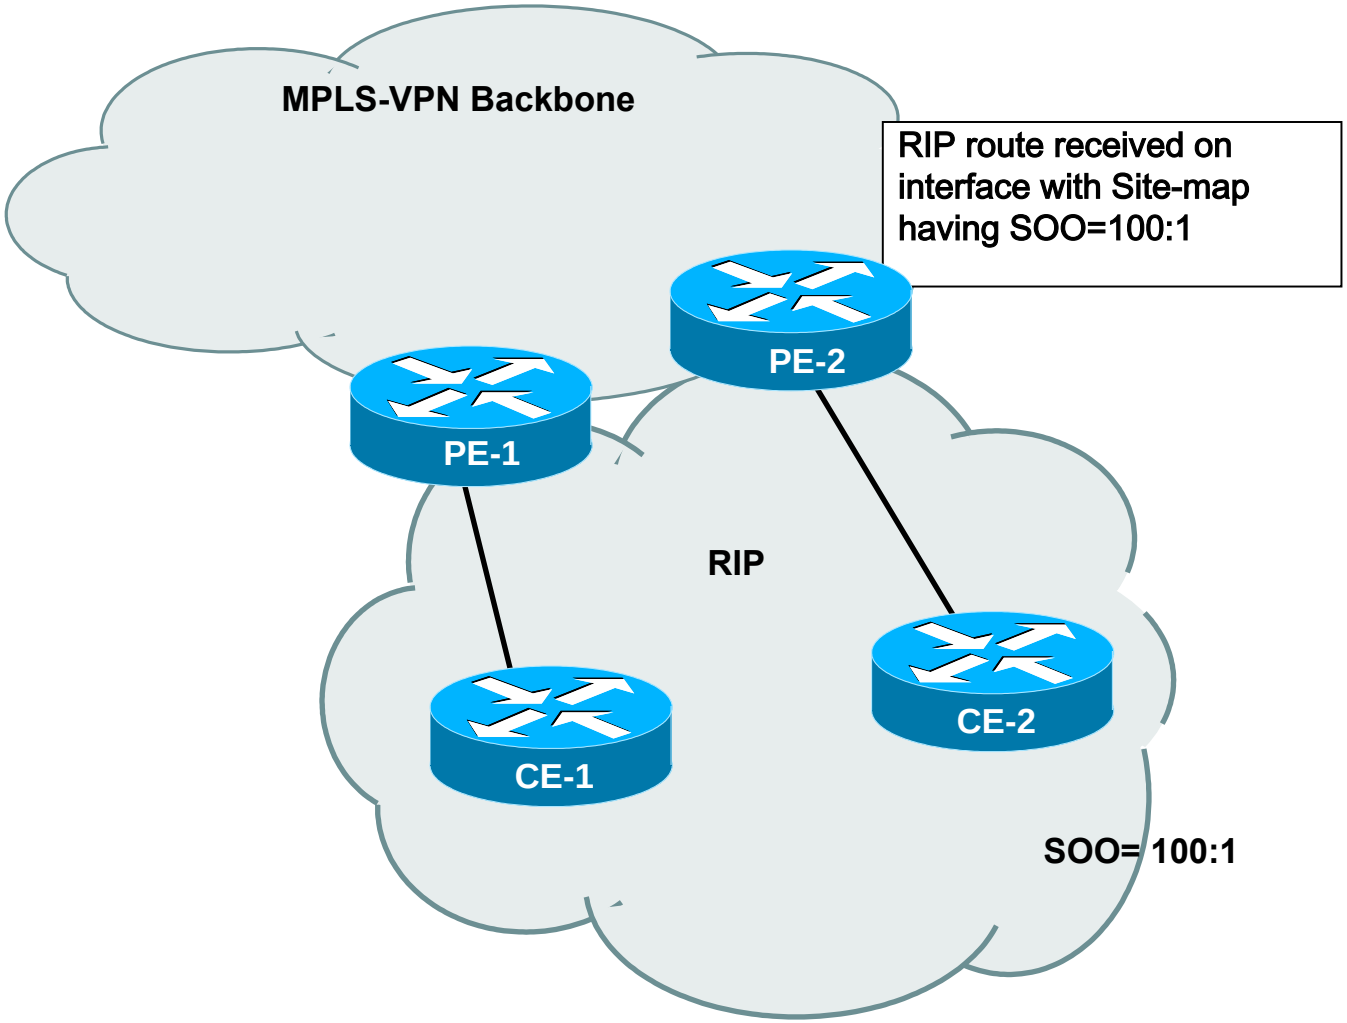

Avoiding Routing Loops: Site Of Origin

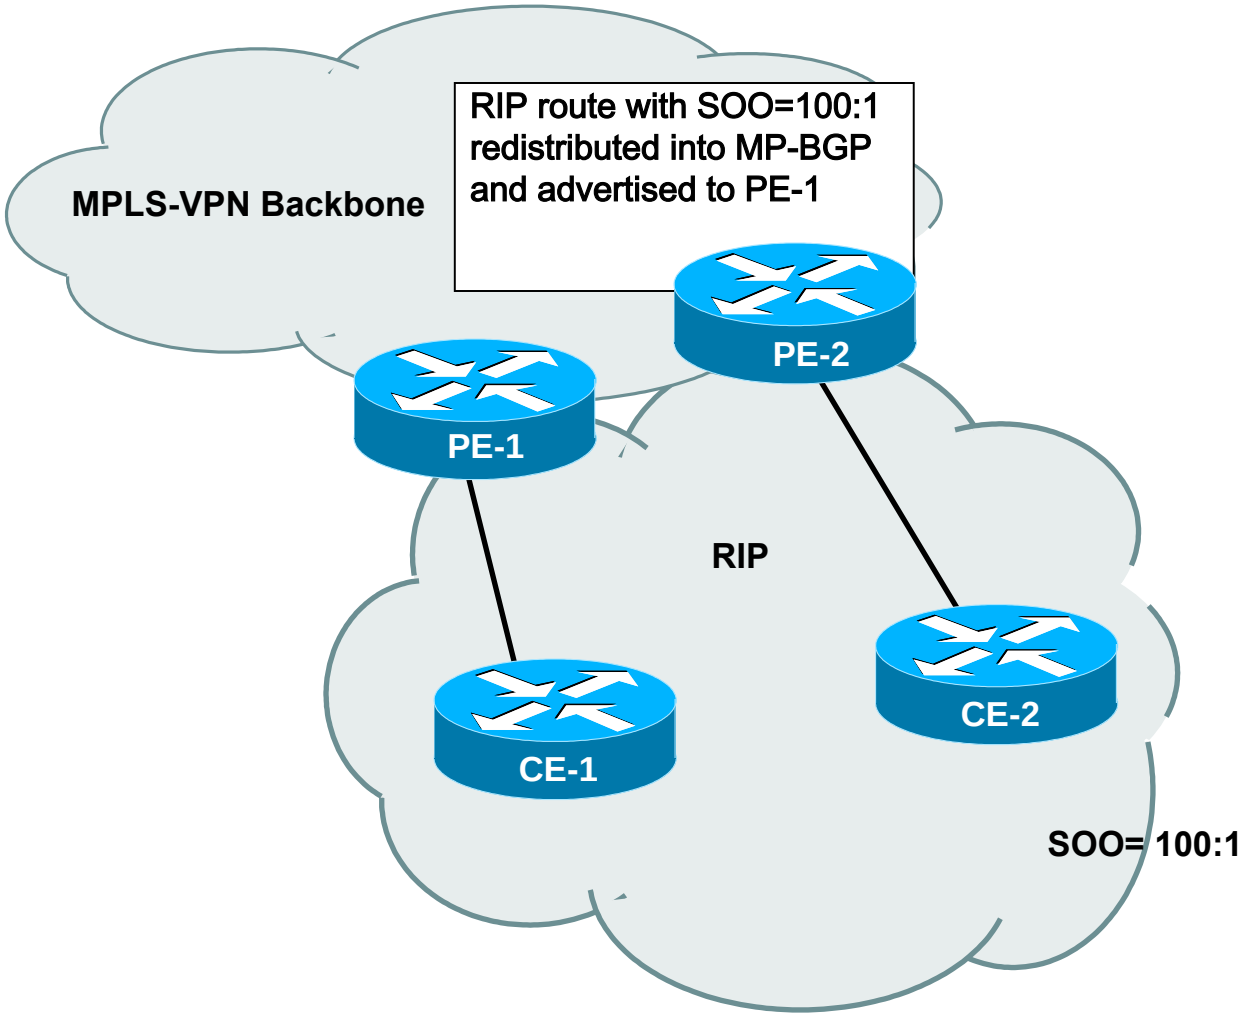

Avoiding Routing Loops: Site Of Origin

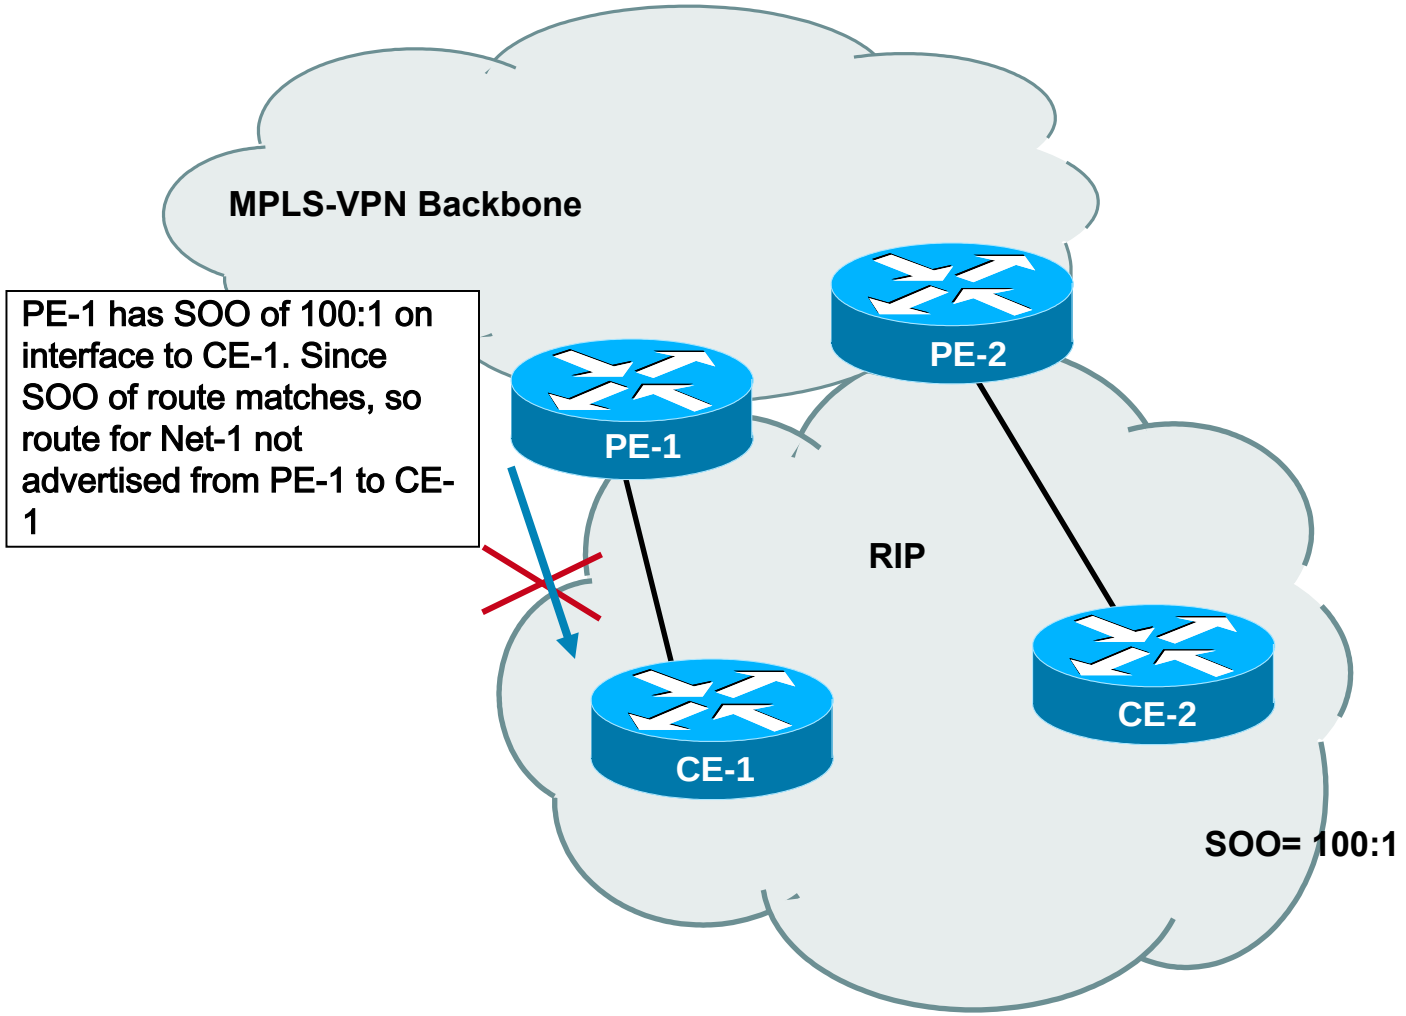

Summary

RIP can be used as a PE-CE routing protocol

RIP v2 should be used as it supports VLSM

RIP has loop detection mechanisms to prevent routing loops with complex connectivity models

MPLS Bootcamp

Multi-VRF CE (aka VRF-lite)

What is Multi-VRF CE?

- Multi-VRF CE architecture uses the VRF concept to support multiple (overlapping and independent) routing tables (and forwarding tables) per customer
- Not a feature but an application based on VRF implementation
- Any routing protocol supported by normal VRF can be used in a Multi-VRF CE implementation
- The CE supports traffic separation between customer networks
- There is no MPLS functionality on the CE, no label exchange between the CE and PE

What is Multi-VRF CE

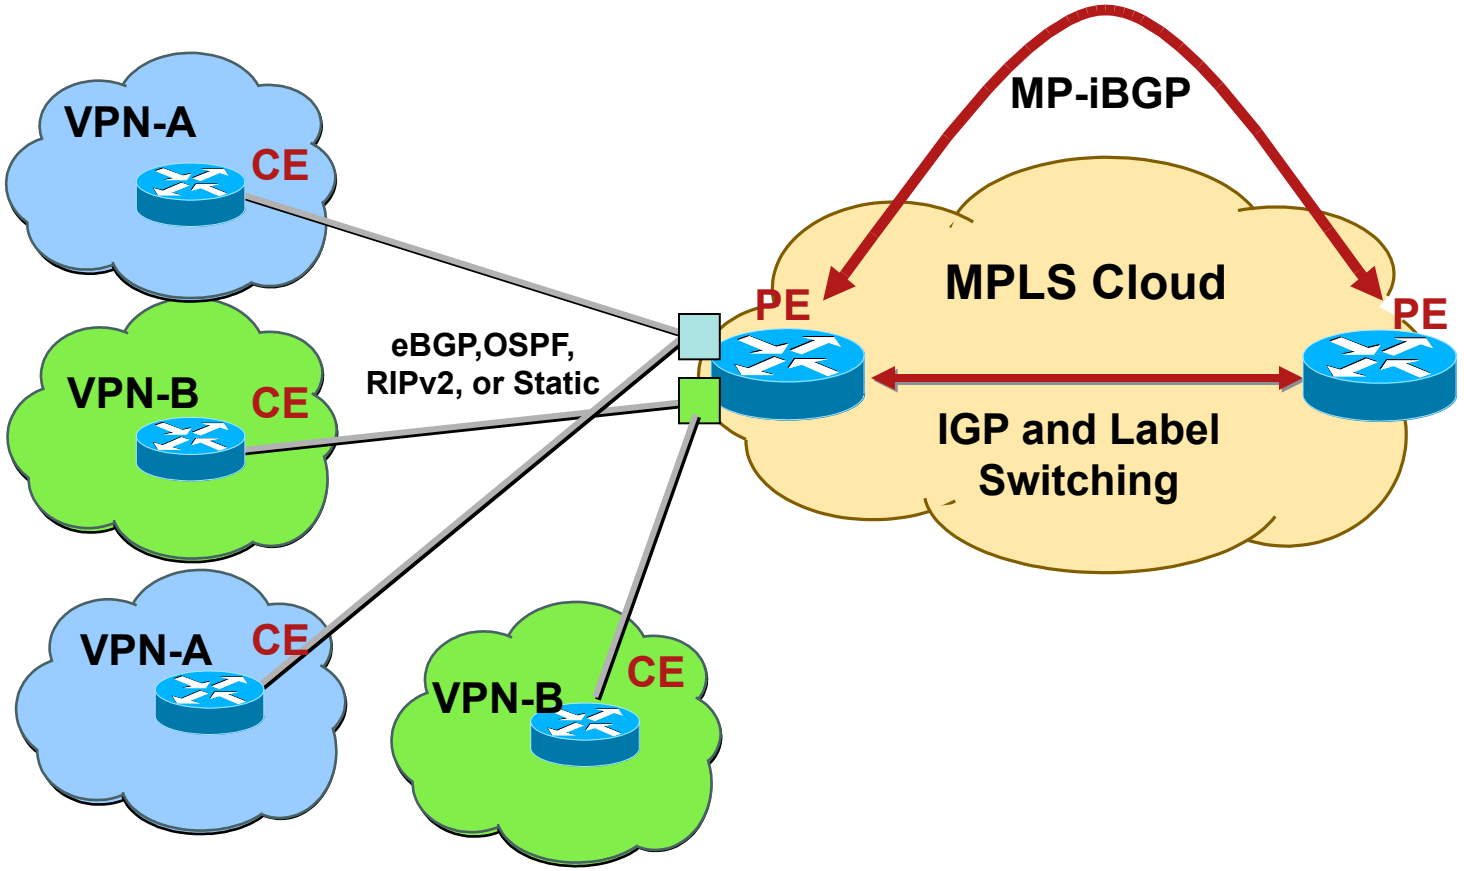

Take the existing PE VRF Functionality…

What is Multi-VRF CE

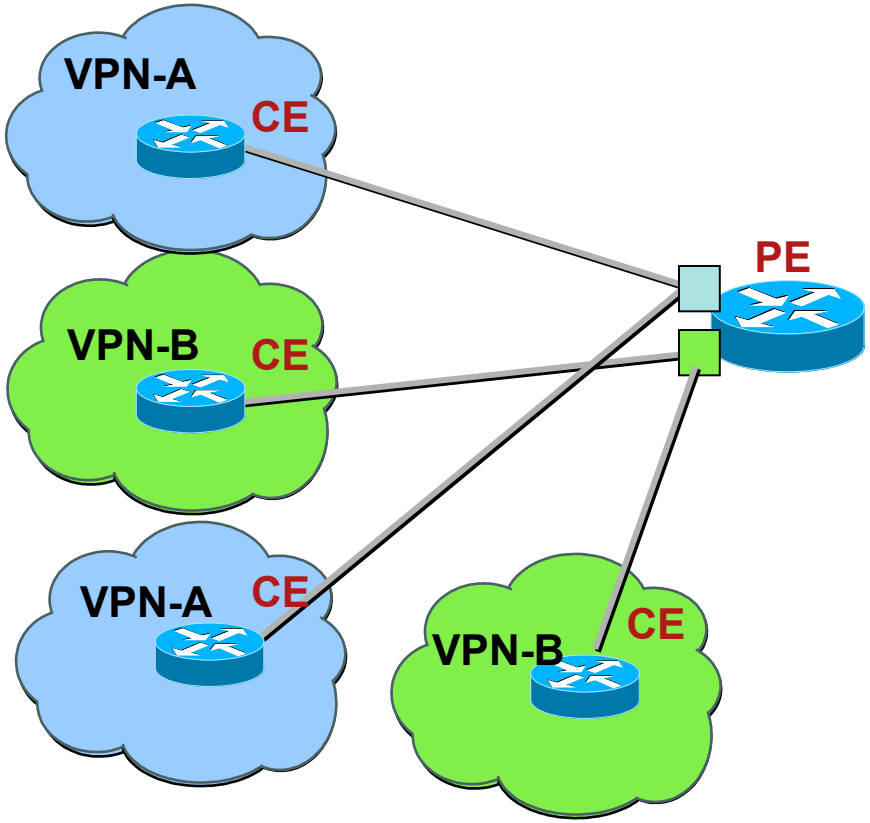

…And Remove the MPLS cloud

What is Multi-VRF CE

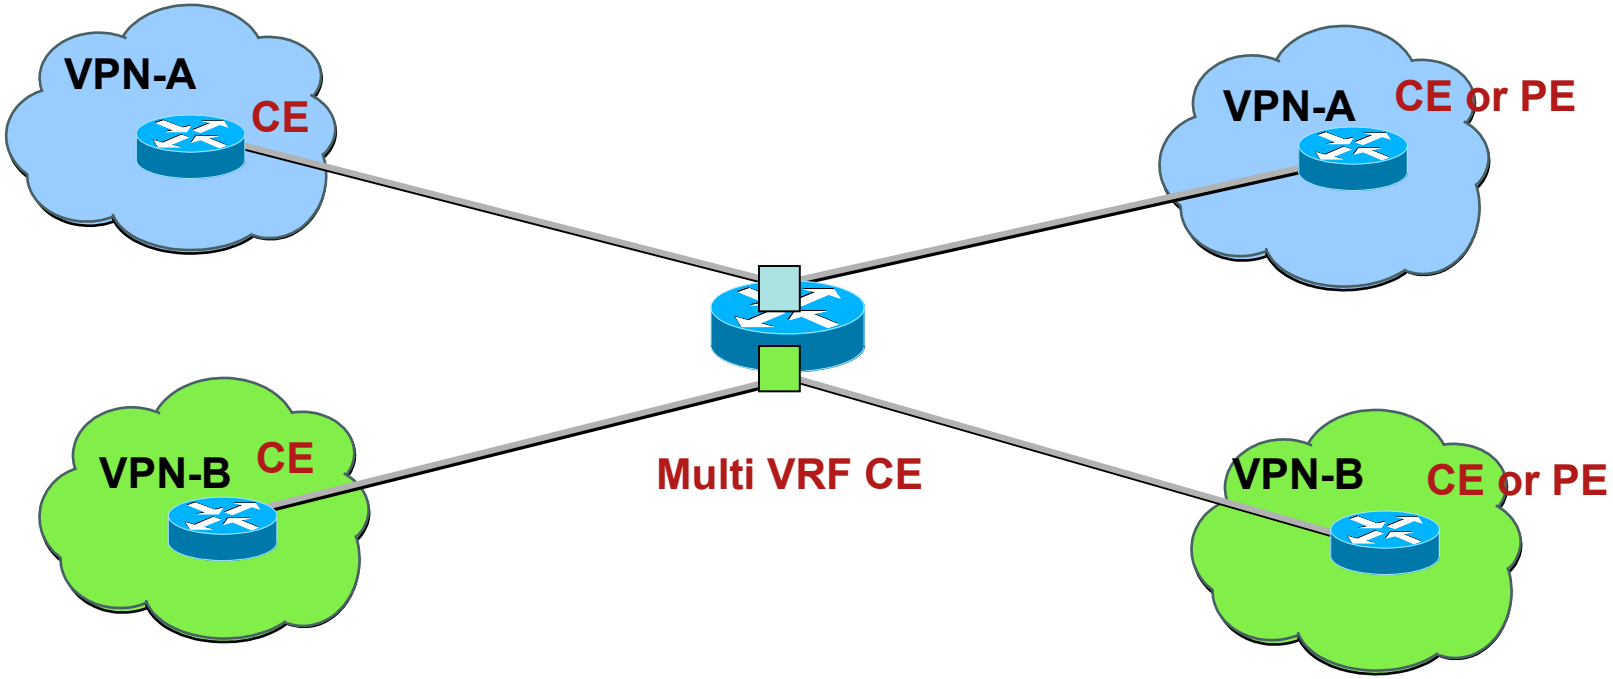

Put it at the customer site and call it a Multi-VRF CE

Multi-VRF CE - Extending MPLS-VPN

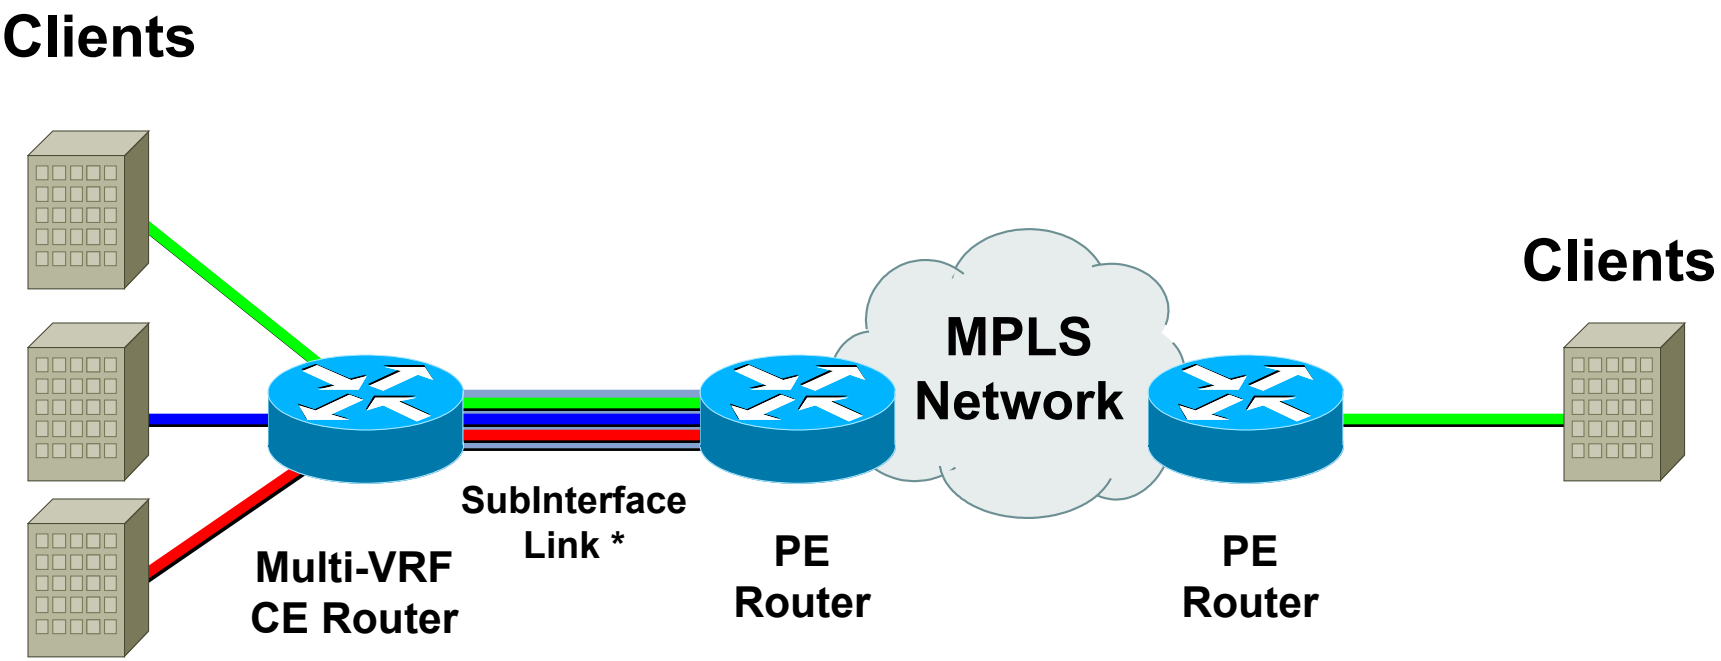

Sub-Interface Link – Any Interface type that supports Sub Interfaces, FE-VLAN, Frame Relay, ATM VC's

Multi-VRF CE - *a standalone Virtual-router !*

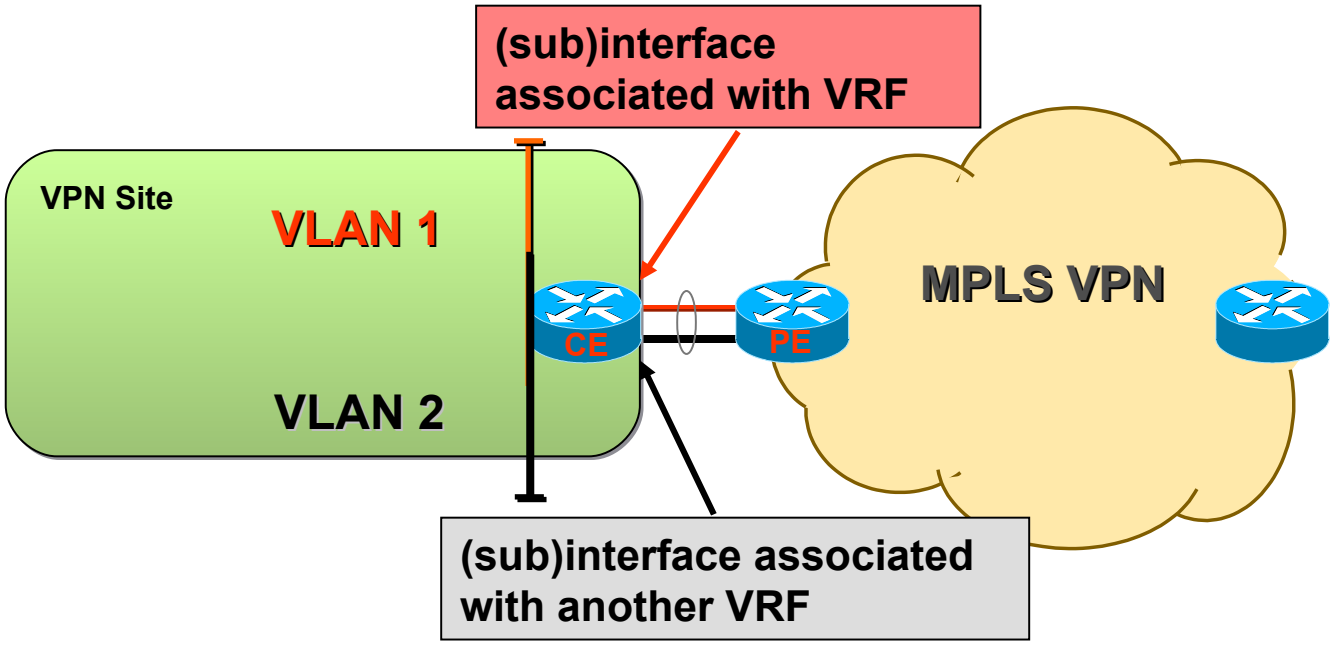

No MPLS, nor MP-iBGP on CE

Local Inter-VRF routing is supported

Multi-VRF CE Architecture

- Enhanced branch office capability
- CE routers use VRF interfaces VLAN-like configuration on the customer side
- CE router can only configure VRF interfaces and support VRF routing tables
- Use using a Multi-vrf CE is an alternative to separate CE routers per each client's organization

Multi-VRF CE Architecture: Replaces Separate CE Routers

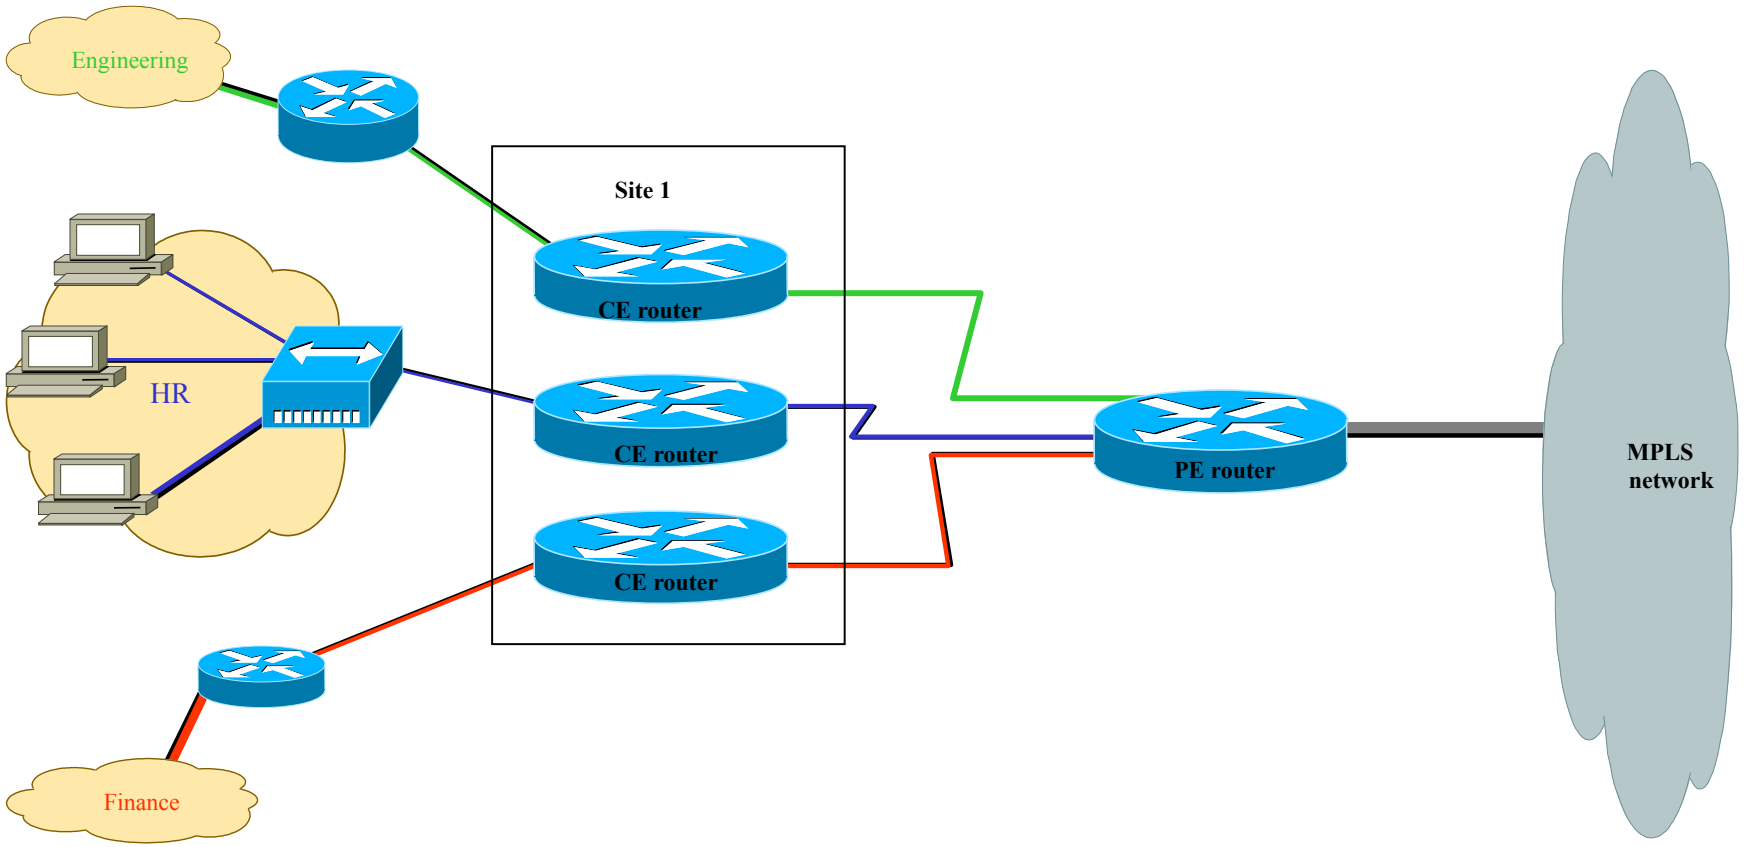

Multi-VRF CE Architecture: Operational Model

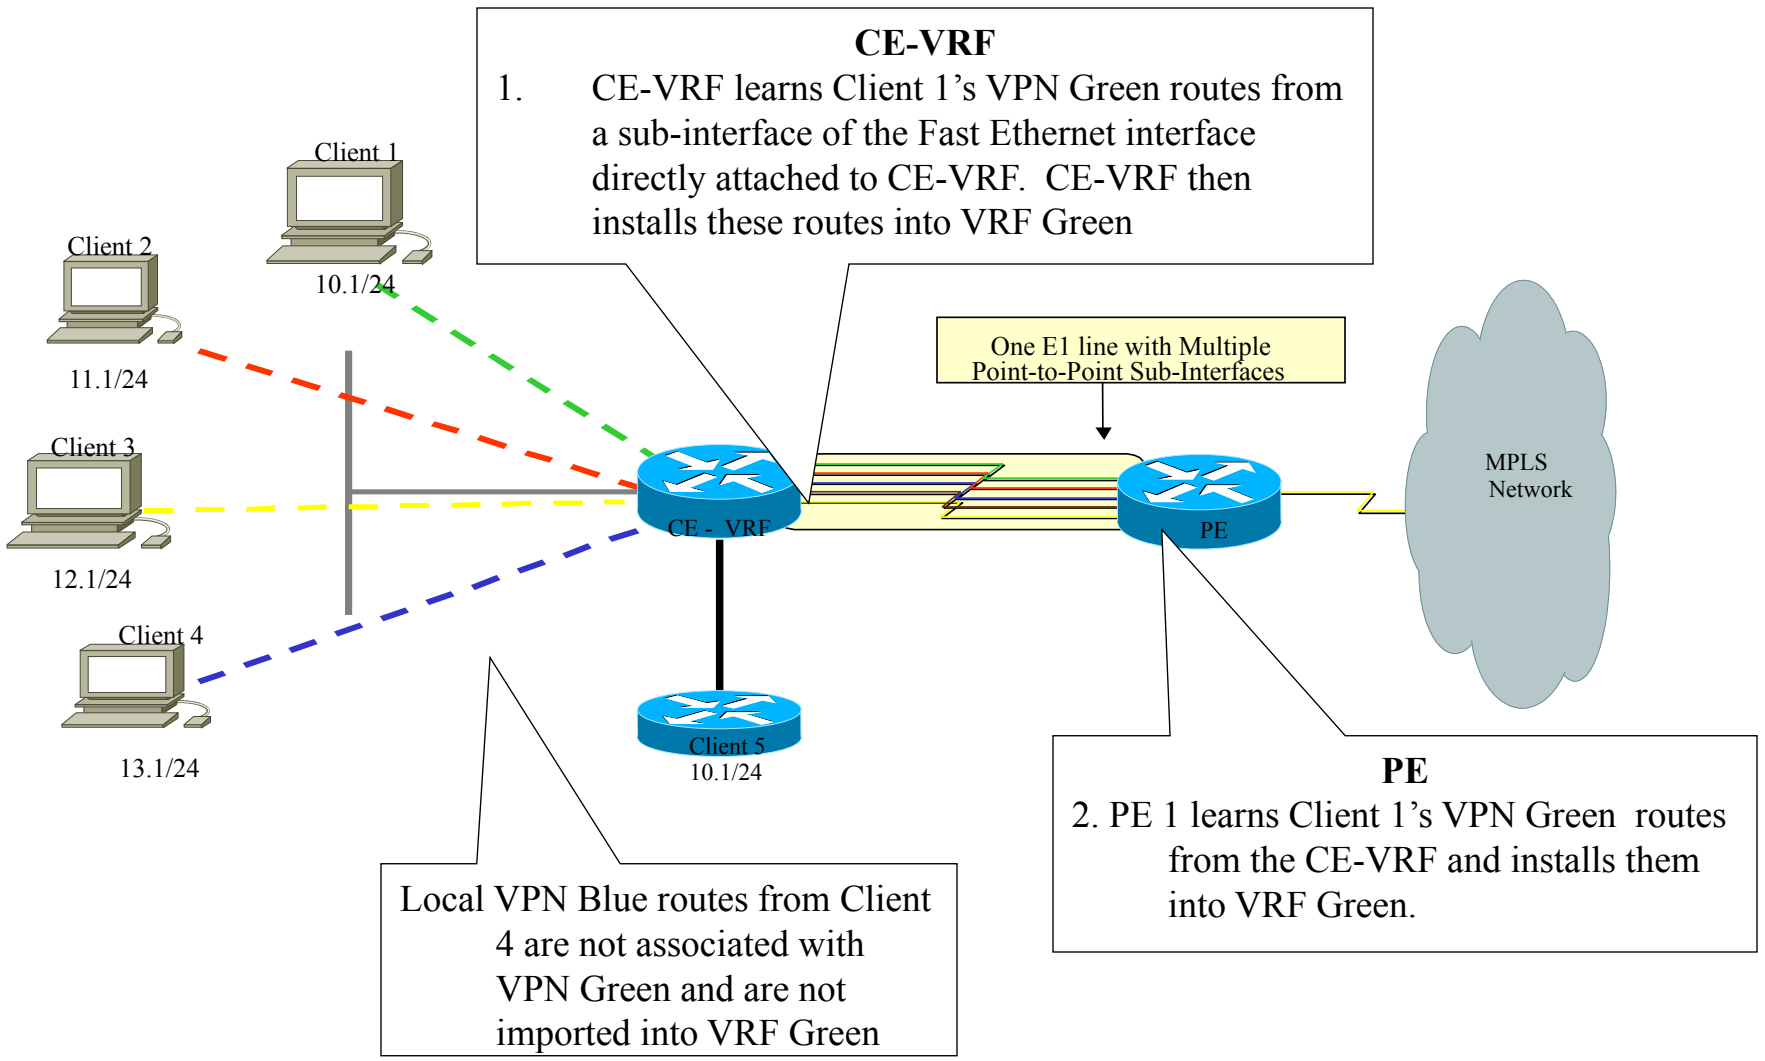

Applications: Two Examples

- Internet and VPN Service Using the Same CE – solution is attractive for small businesses that do not want to install separate CE routers for each service
- Implement Multiple VPNs in a customer site using a single router

Application 1: Internet Services and VPN Services Using A Single CE

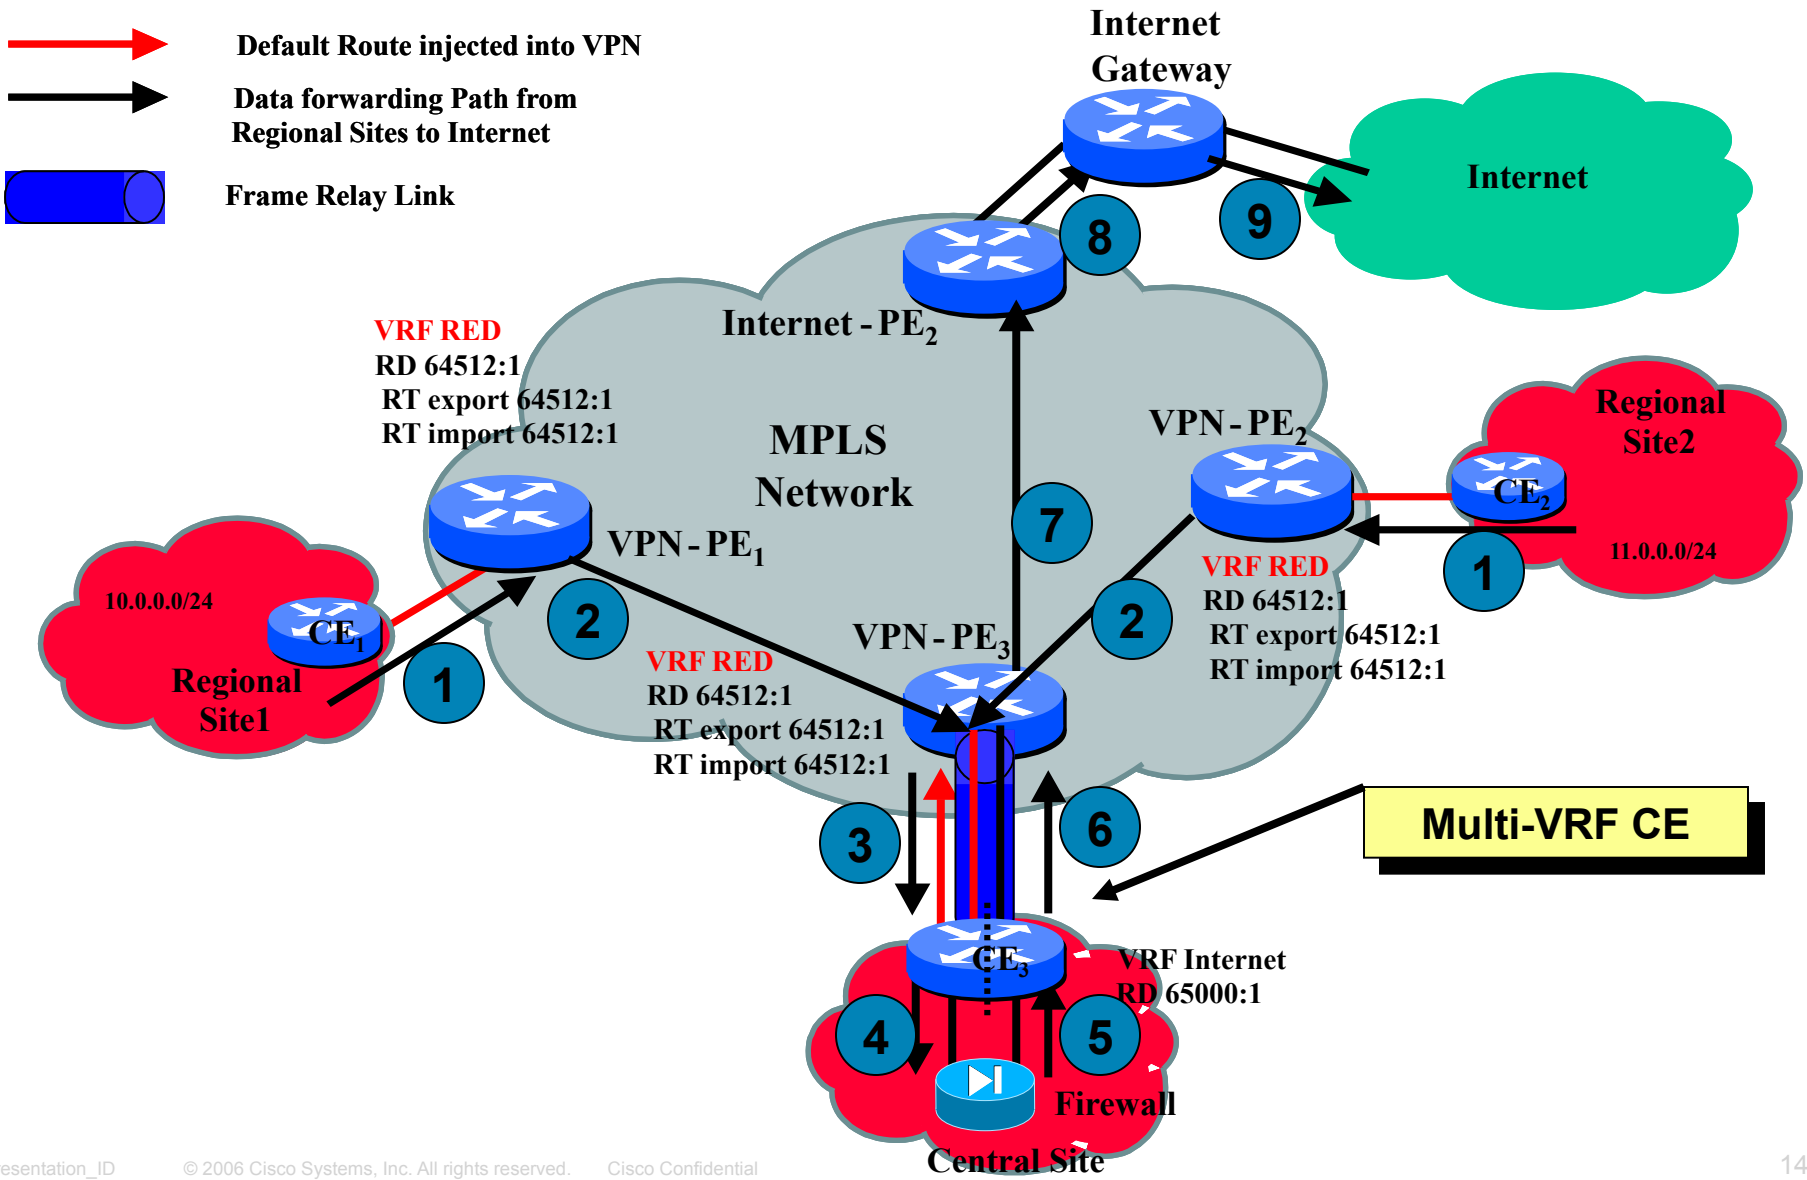

Application 2: Multiple VPNs in a Customer Site Using a Single Router

- Objective: Provide building connectivity via Multi-VRF CE. Multiple departments or companies sharing a building need to be isolated from each other (e.g. financial departments).

Application 2: Overview

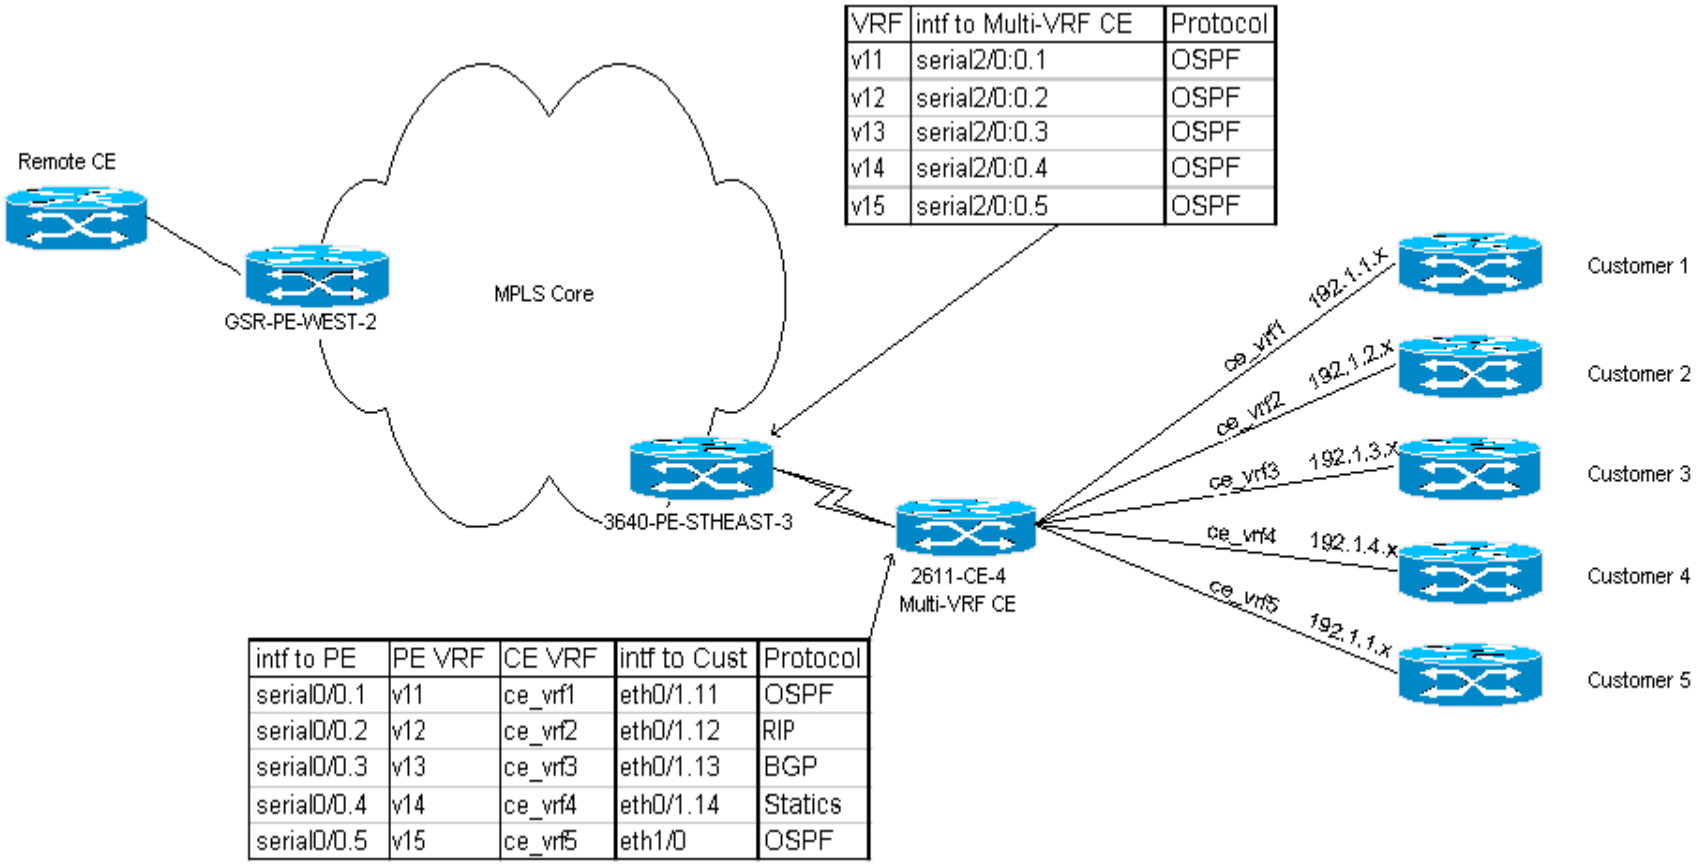

VRF	intf to Multi-VRF CE	Protocol
v11	serial2/0:0.1	OSPF
v12	serial2/0:0.2	OSPF
v13	serial2/0:0.3	OSPF
v14	serial2/0:0.4	OSPF
v15	serial2/0:0.5	OSPF

intf to PE	PE VRF	CE VRF	intf to Cust	Protocol
serial0/0.1	v11	ce_vrf1	eth0/1.11	OSPF
serial0/0.2	v12	ce_vrf2	eth0/1.12	RIP
serial0/0.3	v13	ce_vrf3	eth0/1.13	BGP
serial0/0.4	v14	ce_vrf4	eth0/1.14	Statics
serial0/0.5	v15	ce_vrf5	eth1/0	OSPF

Application 2: Basic Setup

- Inter-site connectivity policies

  All Customer Routers can communicate with Remote CE's but not with each other.

- All Traffic off 2611-CE-4 is segmented into 5 separate VRFs (labeled ce_vrf1-5)
- 3640-PE-STHEAST-3 uses OSPF as the routing protocol to exchange updates with 2611-CE4, but other routing protocols may be used as well
- All other hosts off 2611-CE4 use a combination of OSPF, EBGP, RIPv2 and static routes

Conclusions

- Multi-VRF/VRF-Lite offers the following benefits:
  - Only one CE router is needed facilitating provisioning and network management rather than a multiple CE router solution
  - CE router has VRF functionality without full PE functionality to provide BGP routing tables
    - Note scalability factors
  - Less routing updates to manage
  - Overlapping Customer address spaces
  - Can co-exist with an MPLS-based network but no MPLS enabled on CE
    - Note applicability example for branch offices with multiple networks