



MPLS workshop

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.

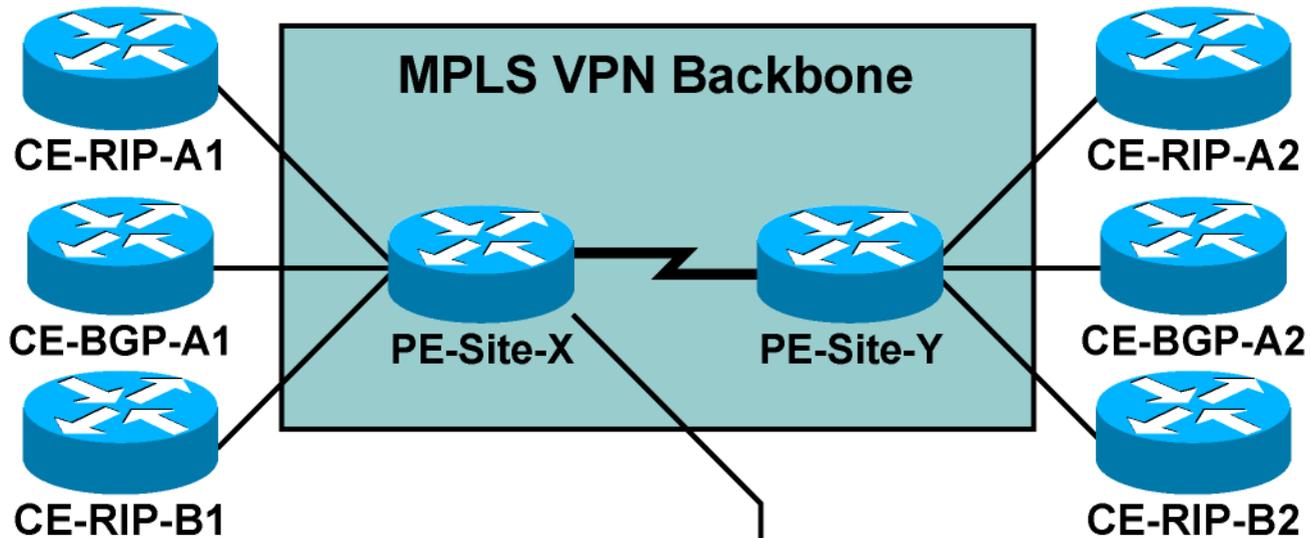
For end-to-end RIP networks, the following applies:

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
```

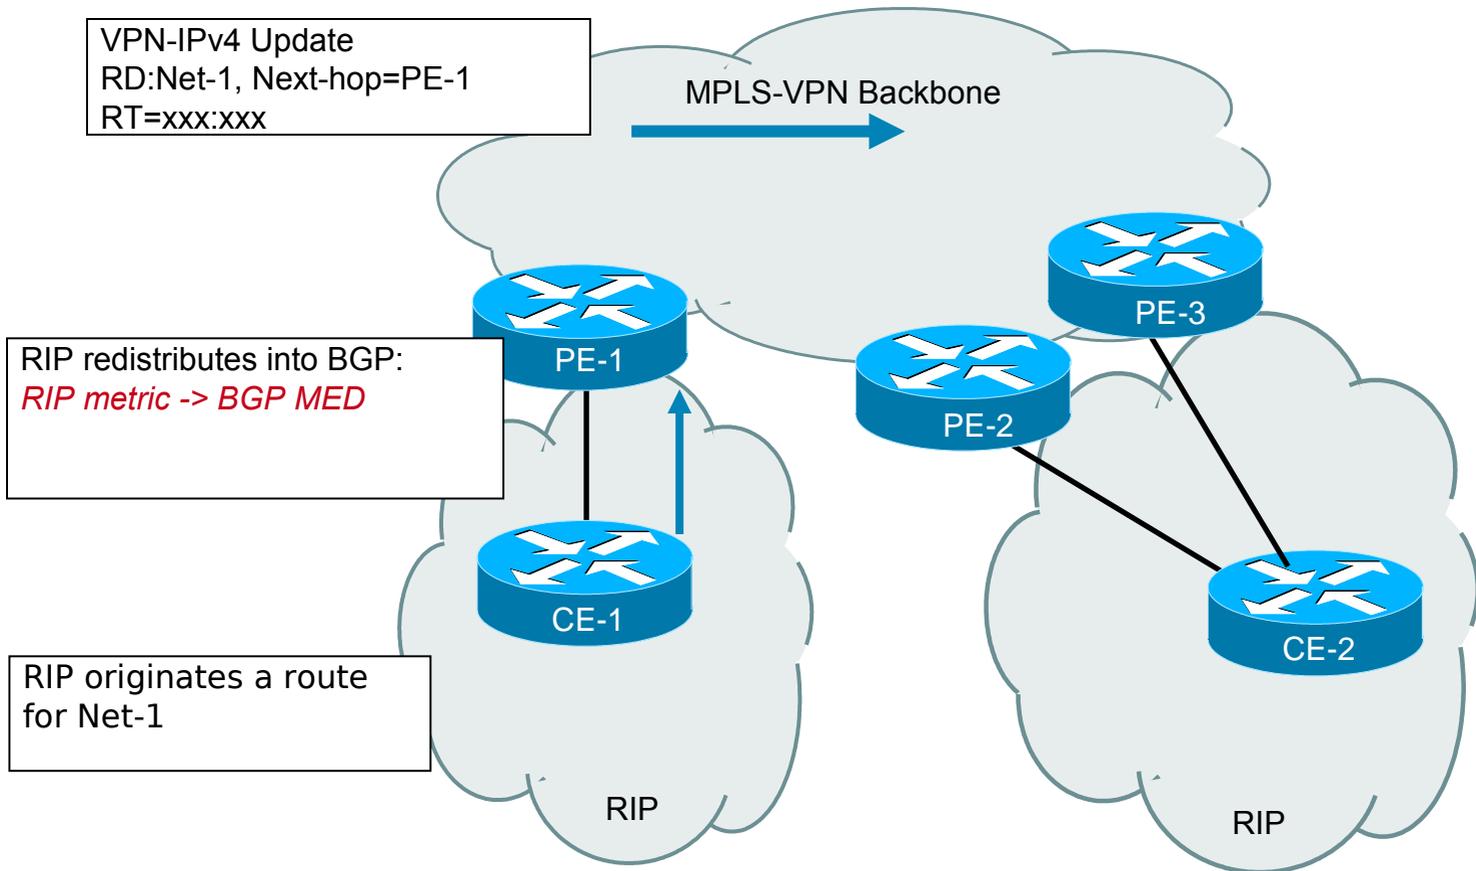
0206_073

Loop Detection with RIP as PE-CE

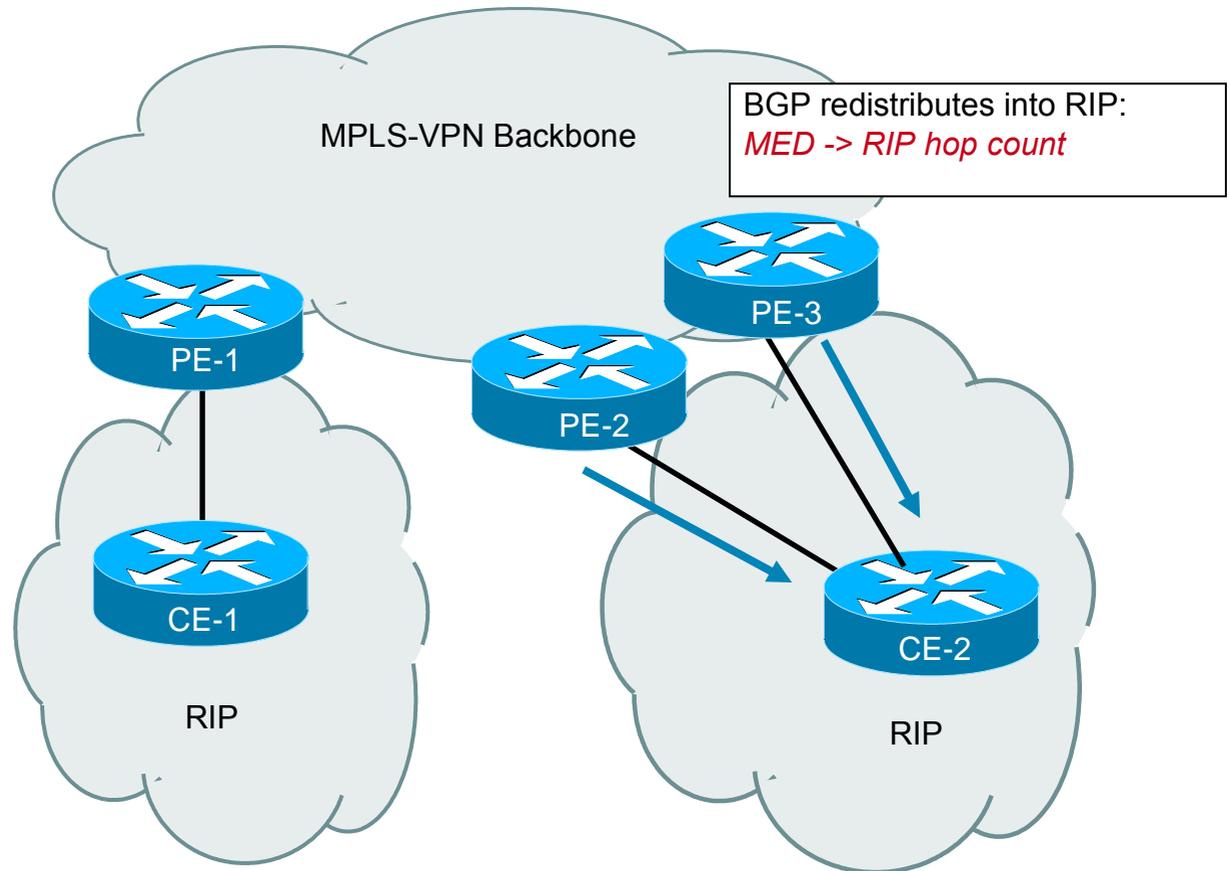
RIP works with the following mechanisms for loop detection:

- Split Horizon
- Site Of Origin (SOO)

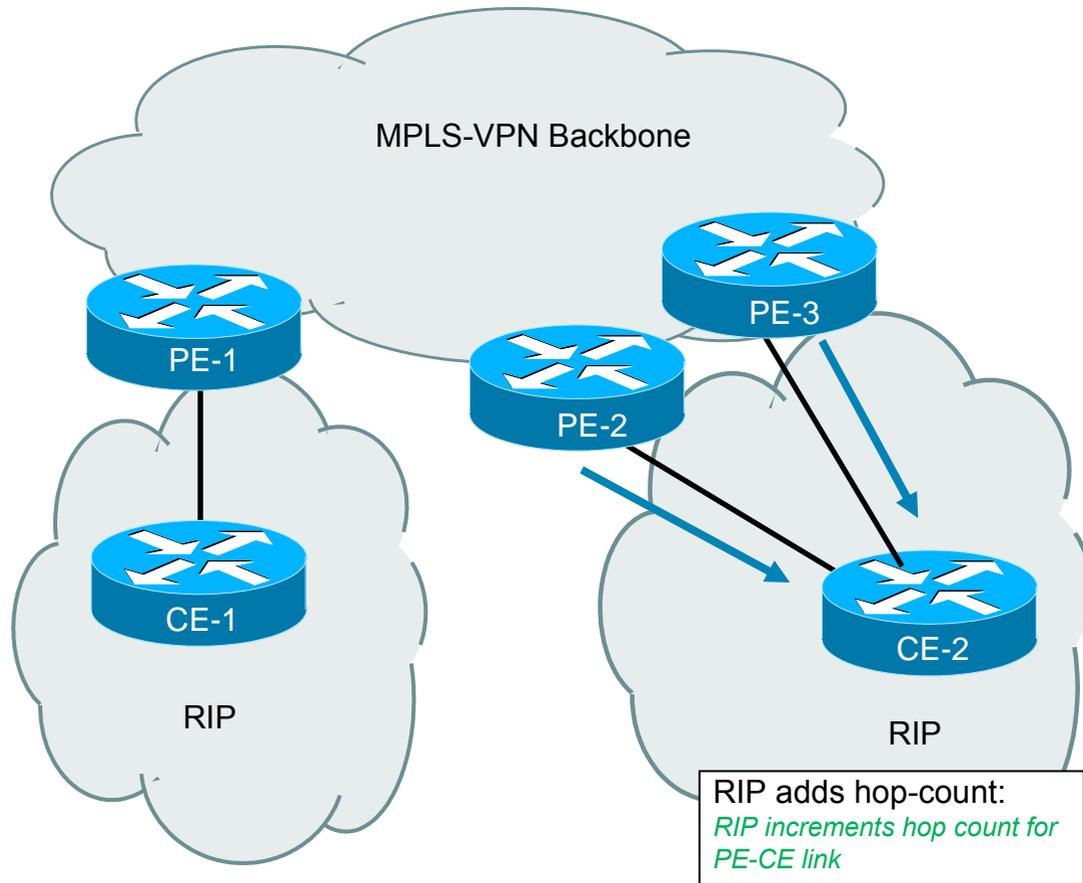
Avoiding Routing Loops: Split-horizon



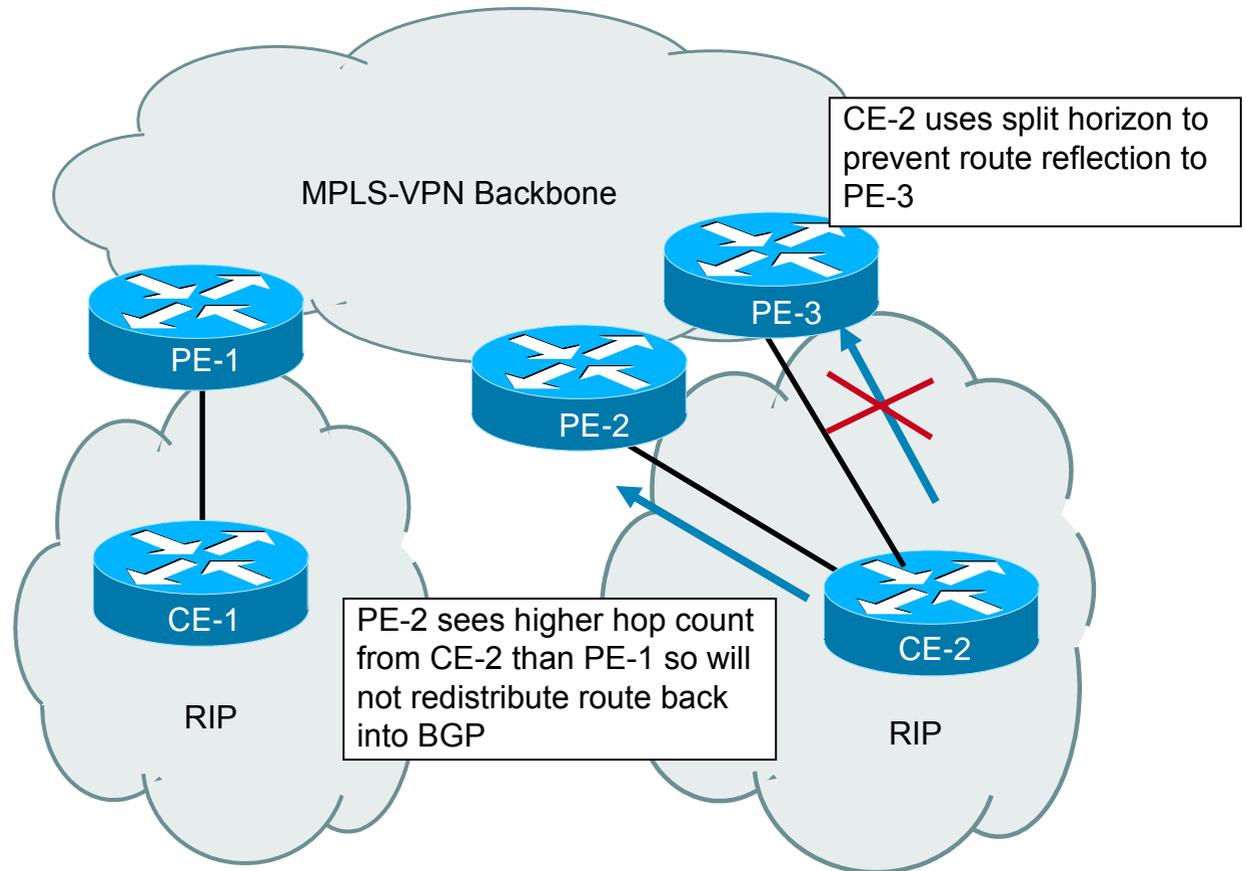
Avoiding Routing Loops: Split-horizon



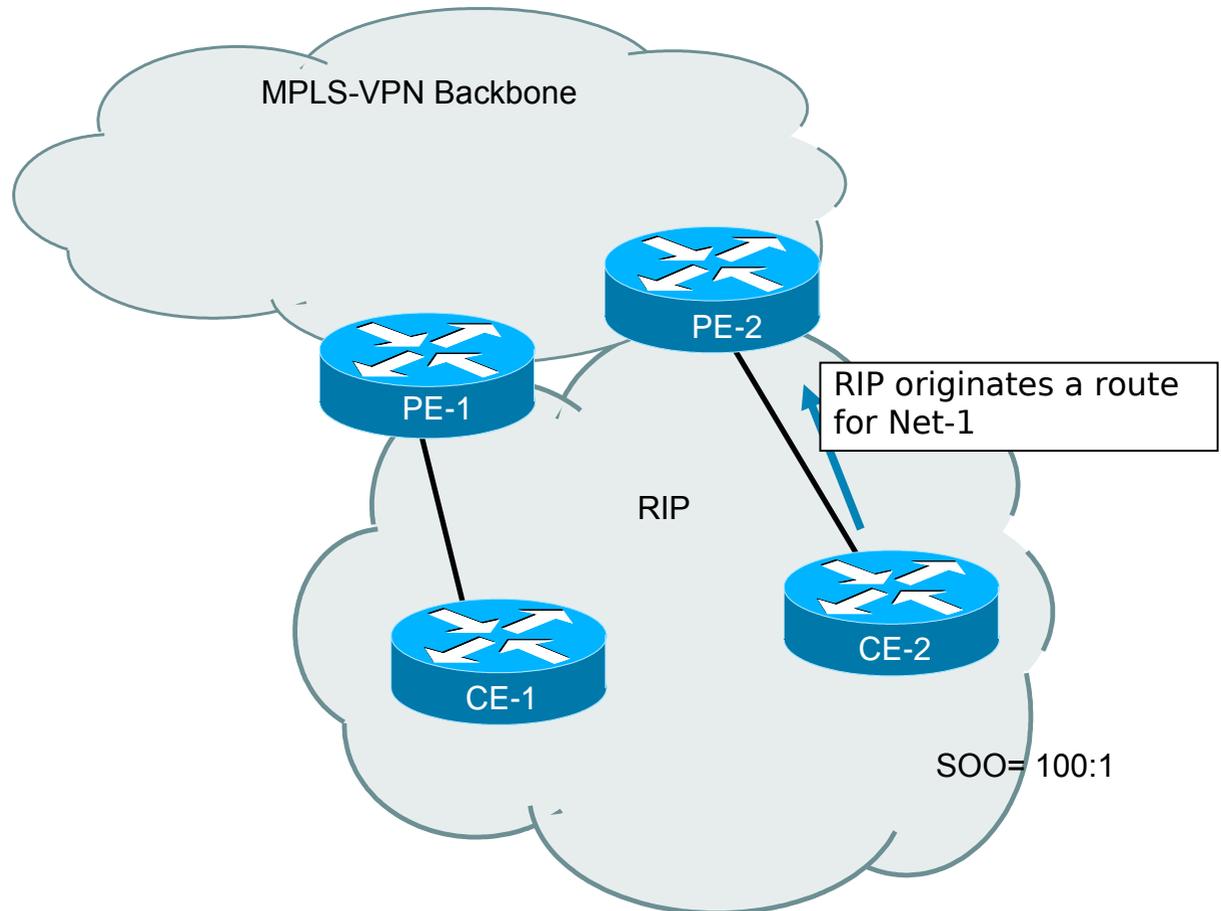
Avoiding Routing Loops: Split-horizon



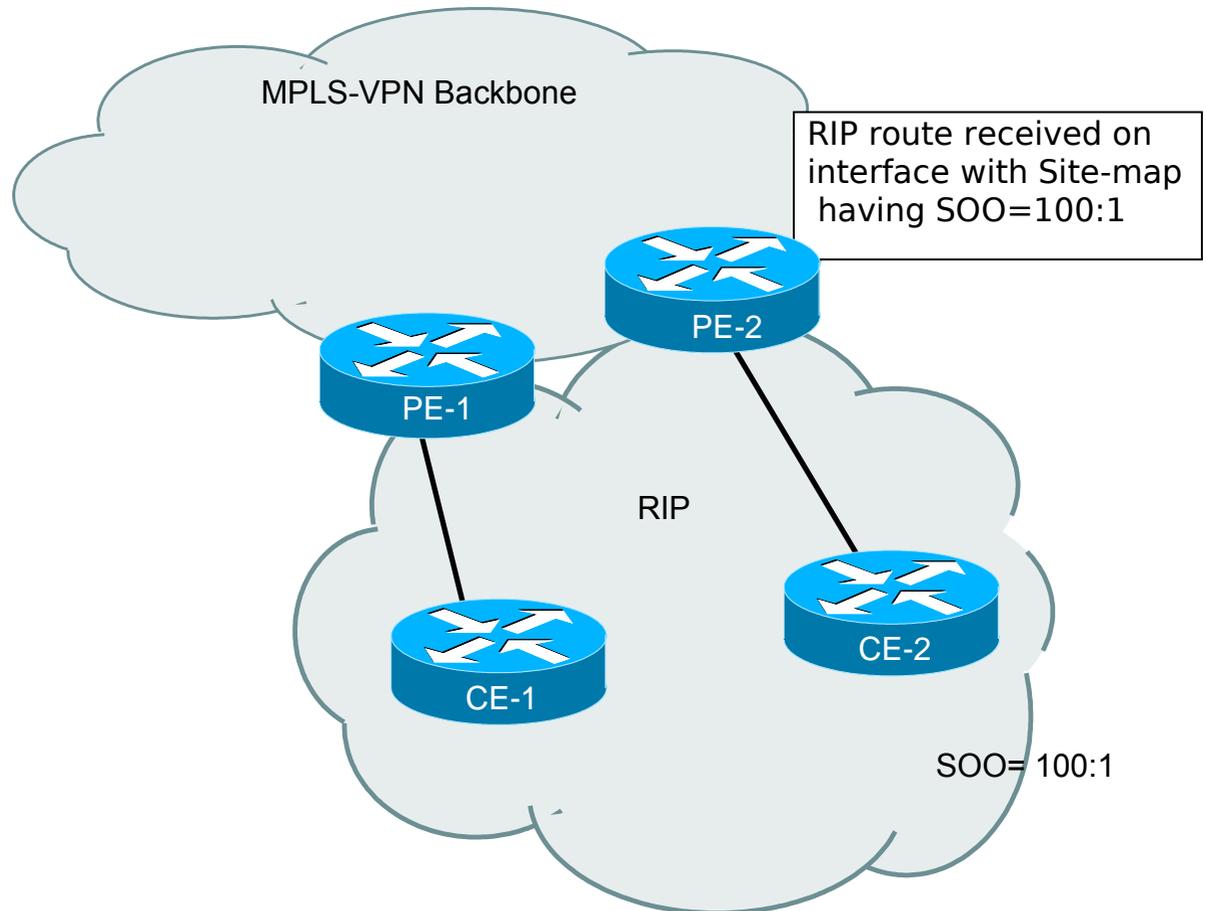
Avoiding Routing Loops: Split-horizon



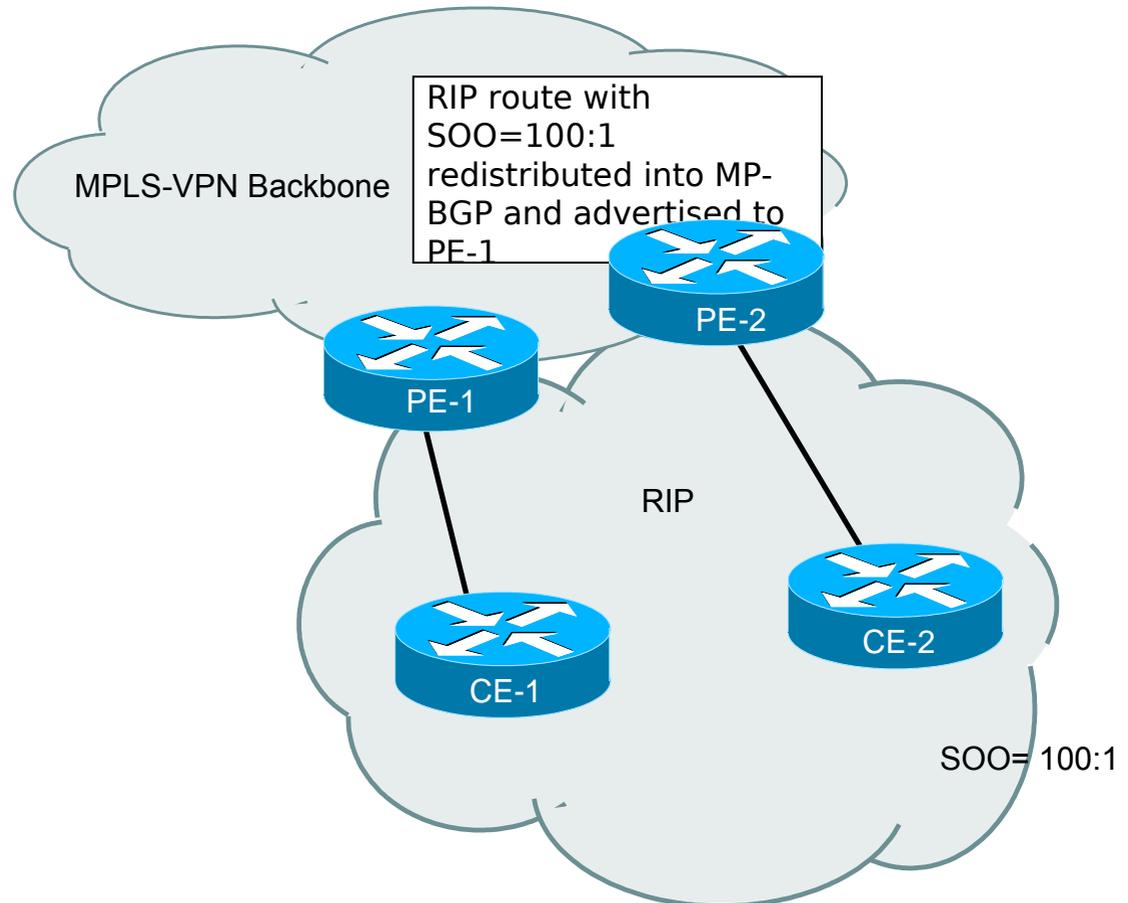
Avoiding Routing Loops: Site Of Origin



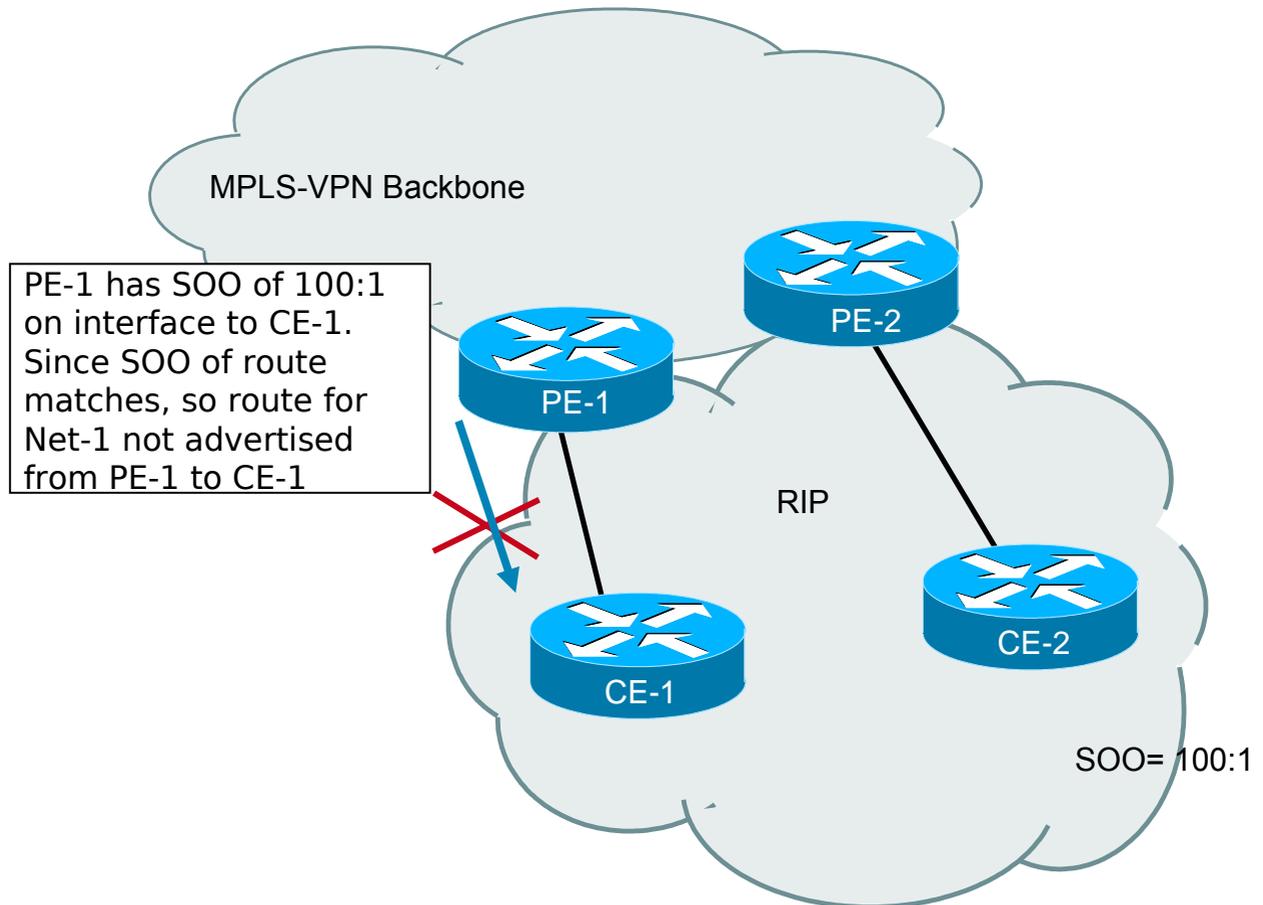
Avoiding Routing Loops: Site Of Origin



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 VPN Implementation

Troubleshooting MPLS VPN

Outline

Overview

MPLS VPN Troubleshooting Preliminary steps

Verify the Routing Information Flow

Validating CE to PE Routing Information Flow

Validating PE to PE Routing Information Flow

Validating PE to CE Routing Information Flow

Verifying the Data Flow

Validating CEF Status

Validating the End-to-end Label Switched Path

Validating the LIB status

Lesson Summary

Preliminary steps in MPLS VPN Troubleshooting

- Perform basic MPLS troubleshooting:

Is CEF enabled?

Are labels for IGP routes generated and propagated?

Are large labeled packets propagated across the MPLS backbone (maximum transmission unit issues)?

Verifying the Routing Information Flow

Verify the routing information flow:

Are CE routes received by a PE?

Are routes redistributed into MP-BGP with proper extended communities?

Are VPNv4 routes propagated to other PE routers?

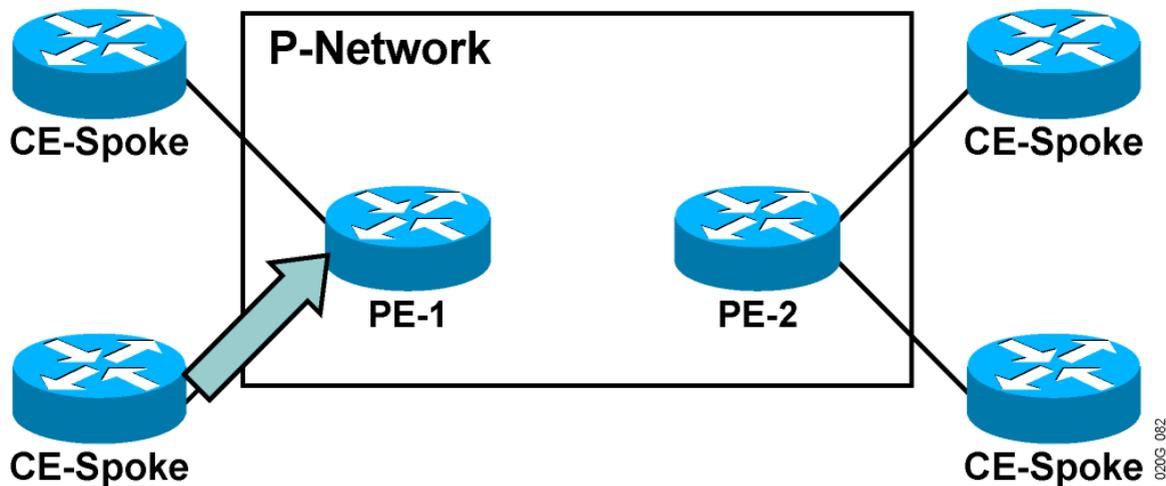
Is the BGP route selection process working correctly?

Are VPNv4 routes inserted into VRFs on other PE routers?

Are VPNv4 routes redistributed from BGP into the PE-CE routing protocol?

Are IPv4 routes propagated to other CE routers?

Validating CE-to-PE Routing Information Flow

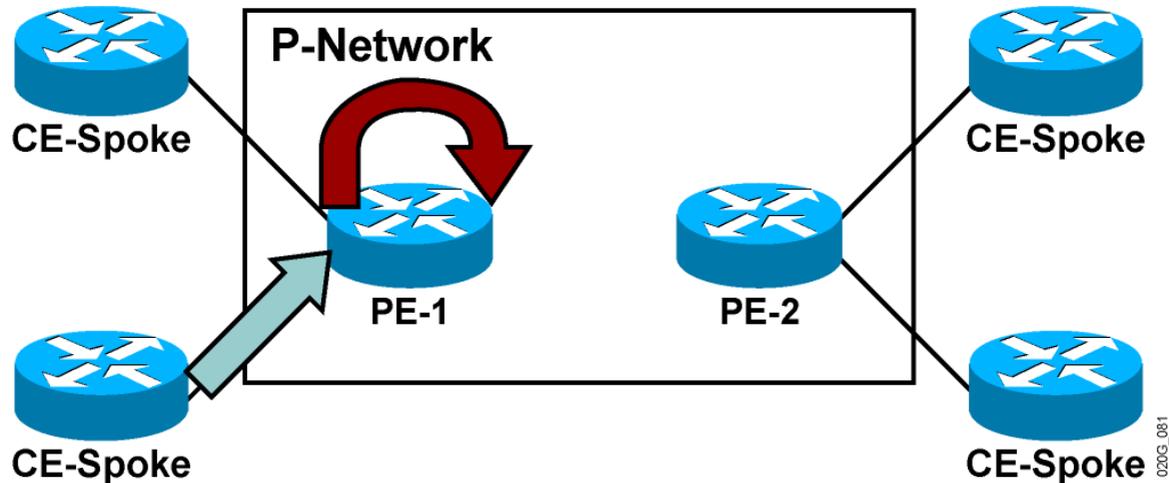


- Are CE routes received by PE?

Verify with `show ip route vrf vrf-name` on PE-1.

Perform traditional routing protocol troubleshooting if needed.

Validating PE-to-PE Routing Information Flow

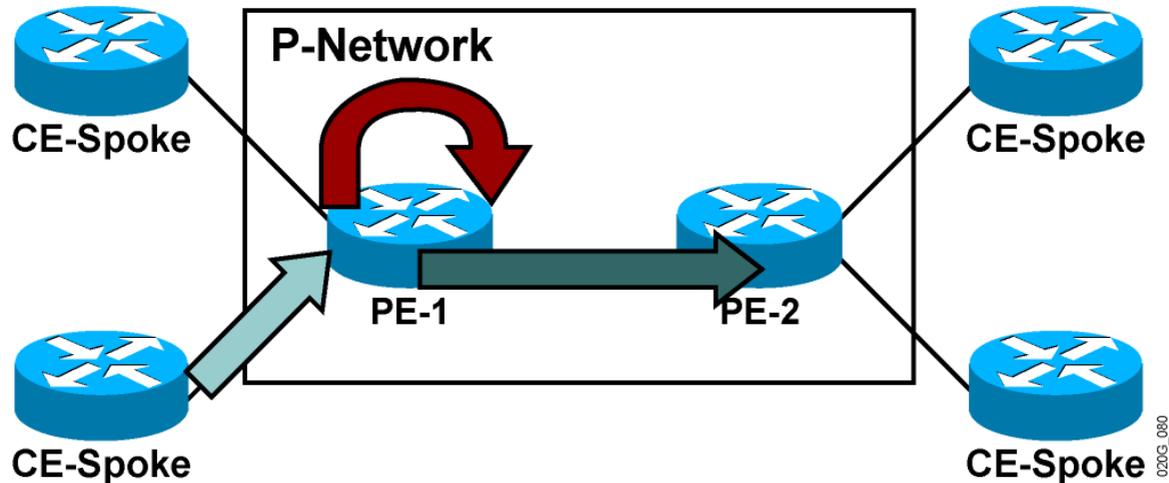


- Are routes redistributed into MP-BGP with proper extended communities?

Verify with `show ip bgp vpnv4 vrf vrf-name ip-prefix` on PE-1.

Troubleshoot with `debug ip bgp` commands.

Validating PE-to-PE Routing Information Flow (Cont.)

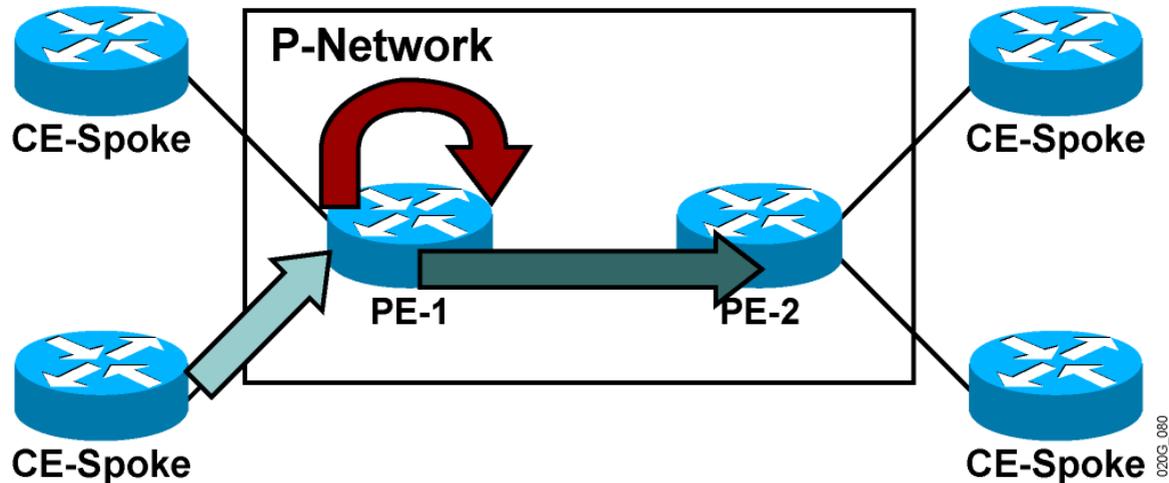


- Are VPNv4 routes propagated to other PE routers?

Verify with `show ip bgp vpnv4 all ip-prefix/length`.

Troubleshoot PE-to-PE connectivity with traditional BGP troubleshooting tools.

Validating PE-to-PE Routing Information Flow (Cont.)



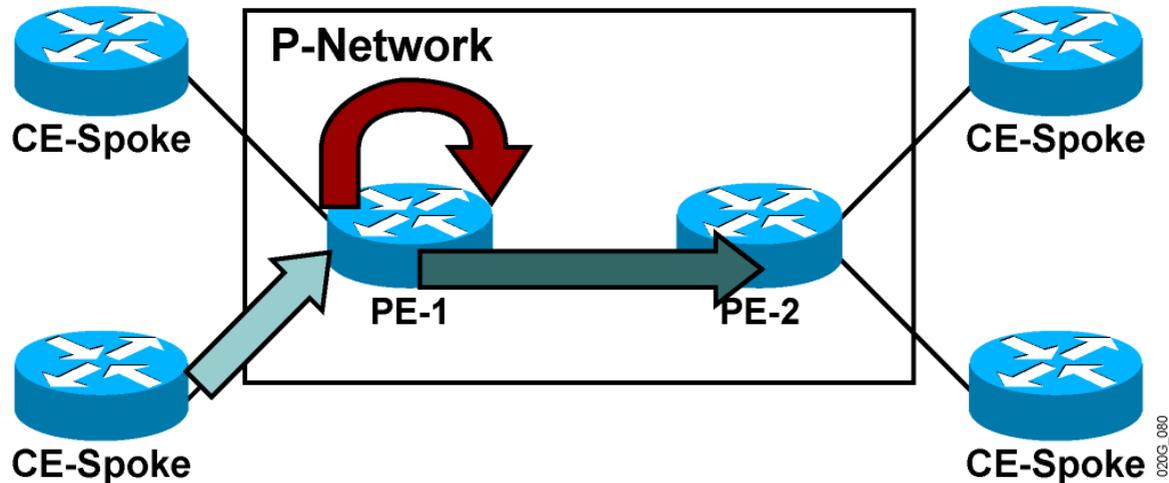
- Is the BGP route selection process working correctly on PE-2?

Verify with `show ip bgp vpnv4 vrf vrf-name ip-prefix`.

Change local preference or weight settings if needed.

Do not change MED if you are using IGP-BGP redistribution on PE-2.

Validating PE-to-PE Routing Information Flow (Cont.)



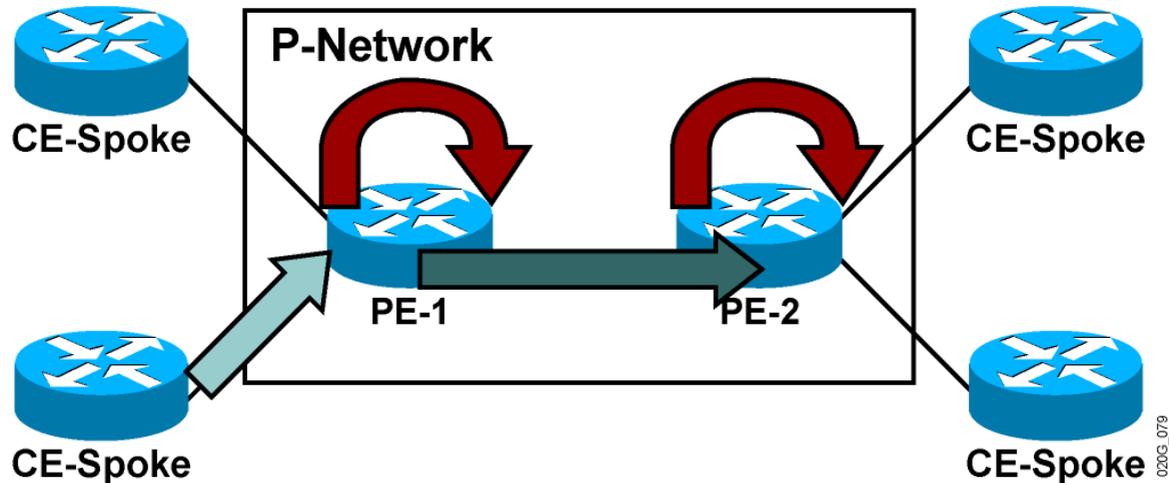
- Are VPNv4 routes inserted into VRFs on PE-2?

Verify with `show ip route vrf`.

Troubleshoot with `show ip vrf detail`.

Perform additional BGP troubleshooting if needed.

Validating PE-to-PE Routing Information Flow (Cont.)

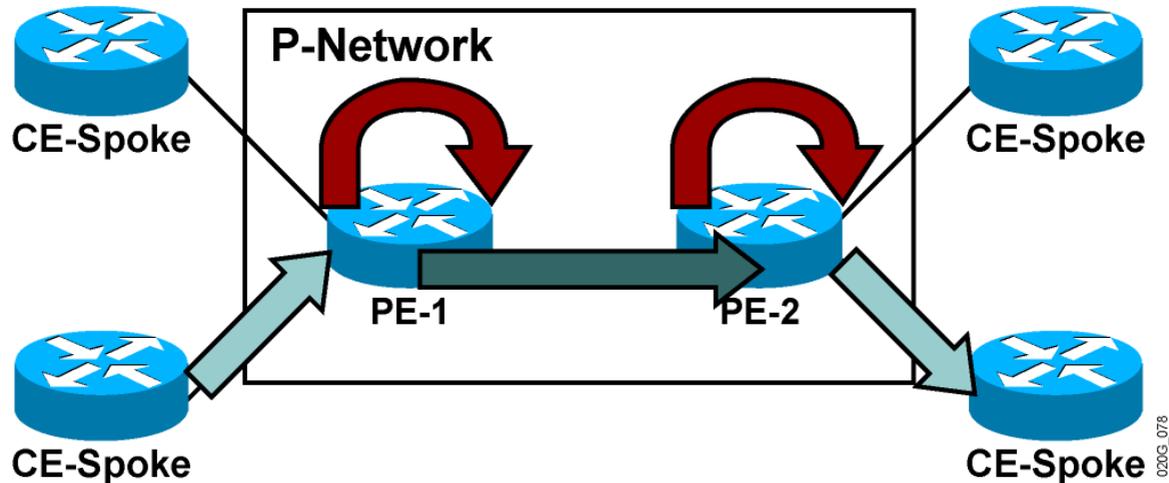


- Are VPNv4 routes redistributed from BGP into the PE-CE routing protocol?

Verify redistribution configuration—is the IGP metric specified?

Perform traditional routing protocol troubleshooting.

Validating PE-to-CE Routing Information Flow



- Are VPNv4 routes propagated to other CE routers?

Verify with `show ip route` on CE Spoke.

Alternatively, does CE Spoke have a default route toward PE-2?

Perform traditional routing protocol troubleshooting if needed.

Verifying the Data Flow

Verify proper data flow:

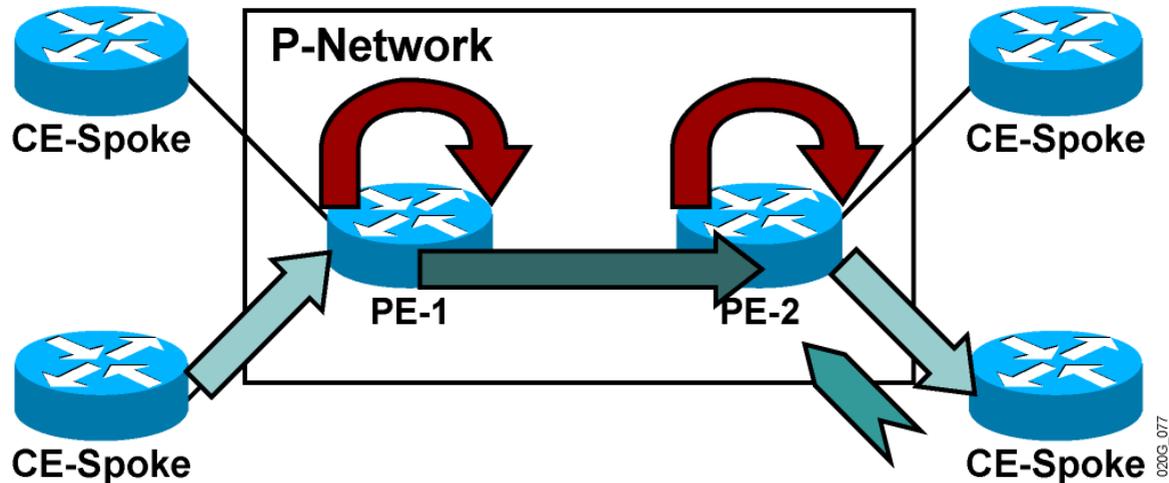
Is CEF enabled on the ingress PE router interface?

Is the CEF entry correct on the ingress PE router?

Is there an end-to-end label switched path tunnel (LSP tunnel) between PE routers?

Is the LFIB entry on the egress PE router correct?

Validating CEF Status



- Is CEF enabled on the ingress PE router interface?

Verify with `show cef interface`.

MPLS VPN needs CEF enabled on the ingress PE router interface for proper operation.

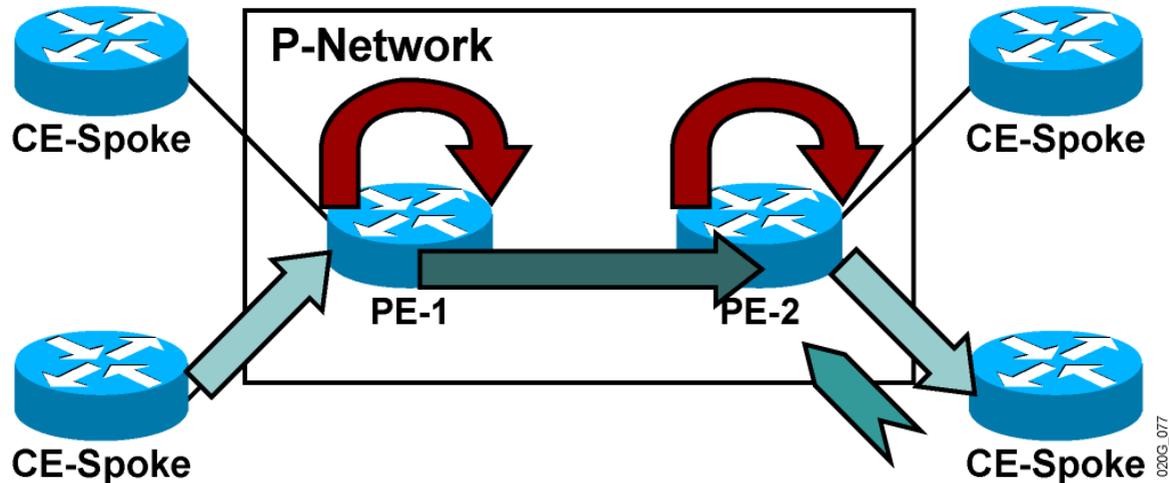
CEF might become disabled because of additional features deployed on the interface.

Validating CEF Status (Cont.)

show cef interface

```
Router#show cef interface serial 1/0.20
Serial1/0.20 is up (if_number 18)
  Internet address is 150.1.31.37/30
  ICMP redirects are always sent
  Per packet loadbalancing is disabled
  IP unicast RPF check is disabled
  Inbound access list is not set
  Outbound access list is not set
  IP policy routing is disabled
  Interface is marked as point to point interface
  Hardware idb is Serial1/0
  Fast switching type 5, interface type 64
  IP CEF switching enabled
  IP CEF VPN Fast switching turbo vector
  VPN Forwarding table "SiteA2"
  Input fast flags 0x1000, Output fast flags 0x0
  ifindex 3(3)
  Slot 1 Slot unit 0 VC -1
  Transmit limit accumulator 0x0 (0x0)
  IP MTU 1500
```

Validating CEF Status (Cont.)

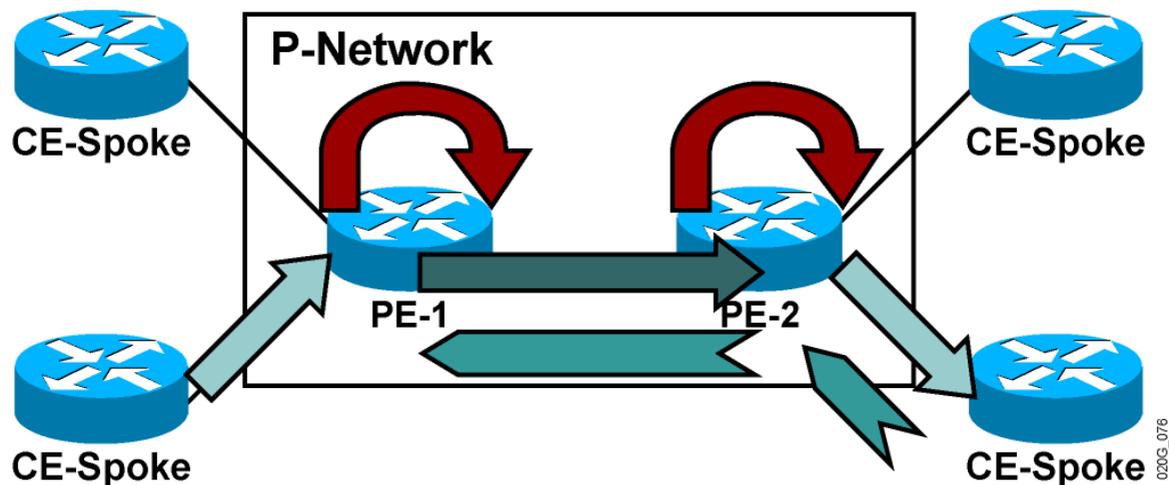


- Is the CEF entry correct on the ingress PE router?

Display the CEF entry with `show ip cef vrf vrf-name ip-prefix/length` detail.

Verify the label stack in the CEF entry.

Validating the End-to-End Label Switched Path



- Is there an end-to-end label switched path tunnel (LSP tunnel) between PE routers?

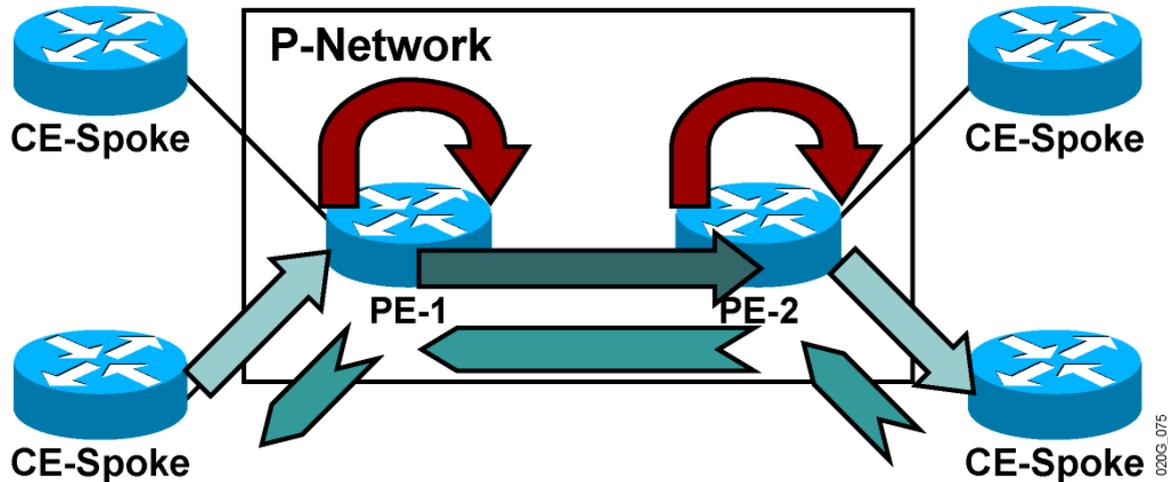
Check summarization issues—BGP next hop should be reachable as host route.

Quick check—if time-to-live (TTL) propagation is disabled, the trace from PE-2 to PE-1 should contain only one hop.

If needed, check LFIB values hop by hop.

Check for MTU issues on the path—MPLS VPN requires a larger label header than pure MPLS.

Validating the LFIB Status



- Is the LFIB entry on the egress PE router correct?

Find out the second label in the label stack on PE-2 with `show ip cef vrf vrf-name ip-prefix detail`.

Verify correctness of LFIB entry on PE-1 with `show mpls forwarding vrf vrf-name value detail`.

Summary

MPLS troubleshooting can be divided into two main steps:

- Verify routing information flow

- Verify proper data flow

Routing information flow troubleshooting requires verification of end-to-end routing information propagation between CE routers.

Verification of the routing information flow should be done systematically, starting at the routing ingress CE and moving to the egress CE.

Verification of the data flow should be done systematically, starting at the data flow ingress CE and moving to the egress CE.



MPLS workshop

Multi-VRF CE (aka VRF-lite)

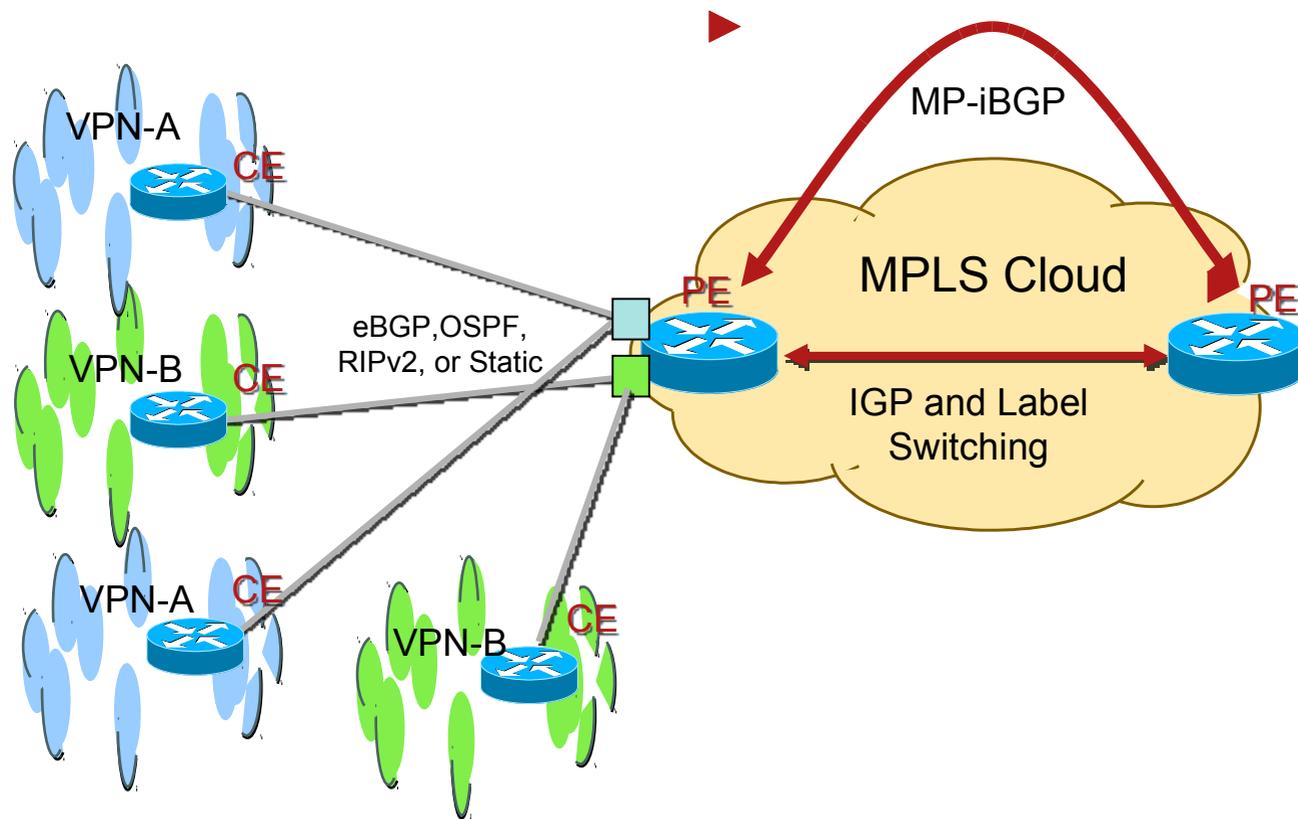
Agenda

- What is Multi-VRF/VRF Lite?
- Applications
- Implementation Example
- Limitations
- OSPF “capability vrf-lite”
- Conclusion

What is Multi-VRF CE?

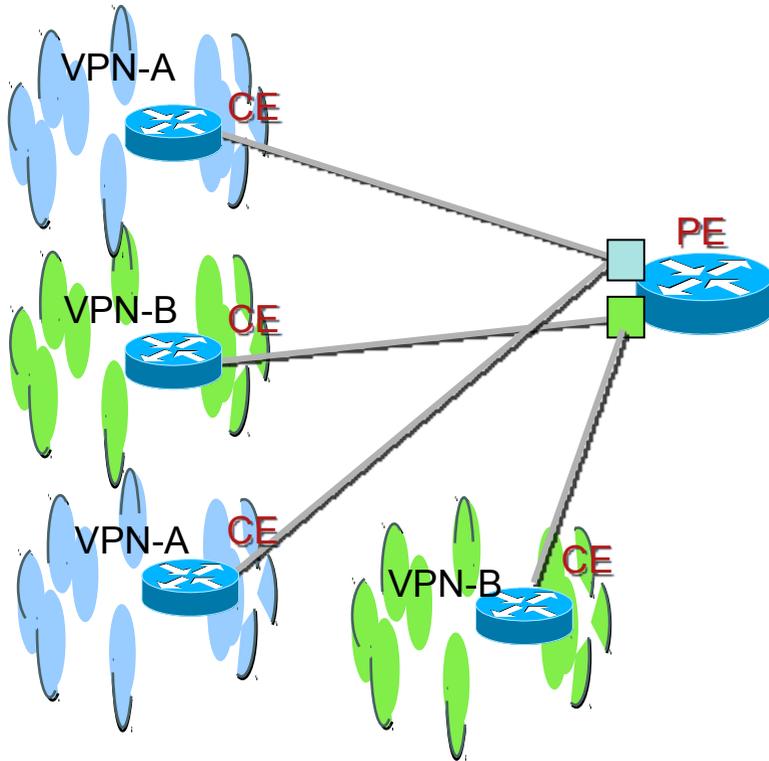
- Multi-VRF CE architecture uses the VRF concept to support multiple (overlapping and independent) routing tables (and forwarding tables) per customer
- Not a feature but an application based on VRF implementation
- Any routing protocol supported by normal VRF can be used in a Multi-VRF CE implementation
- 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



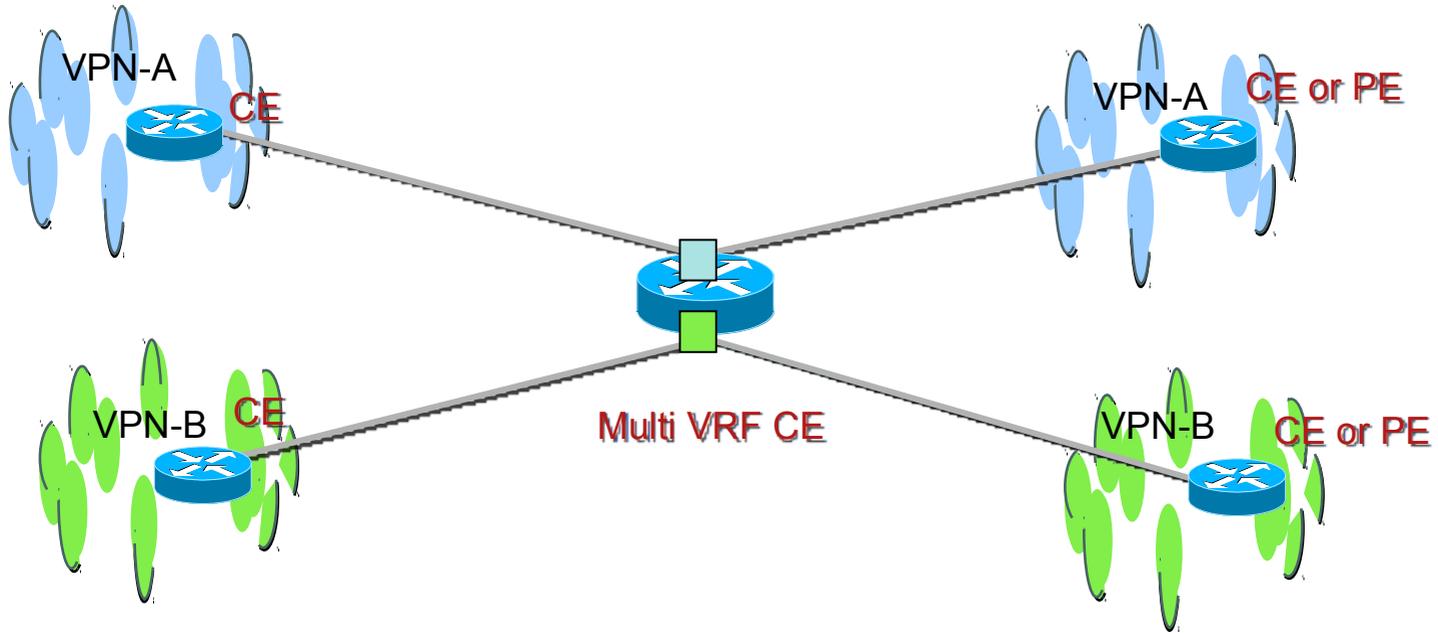
Take the existing PE VRF Functionality...

What is Multi-VRF CE



...And Remove the MPLS cloud

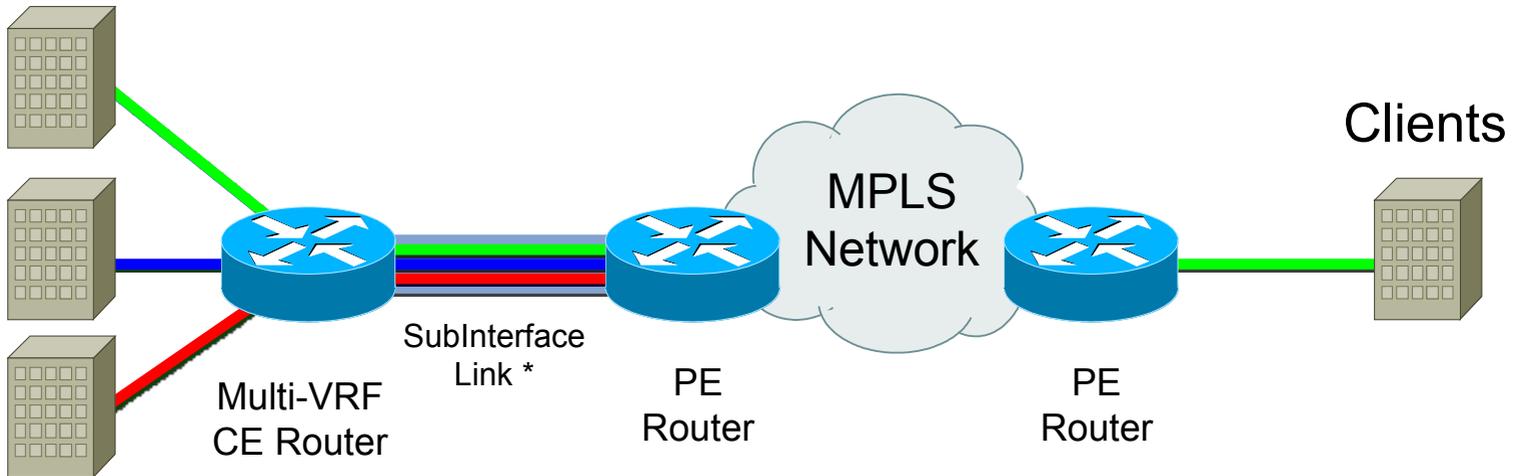
What is Multi-VRF CE



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

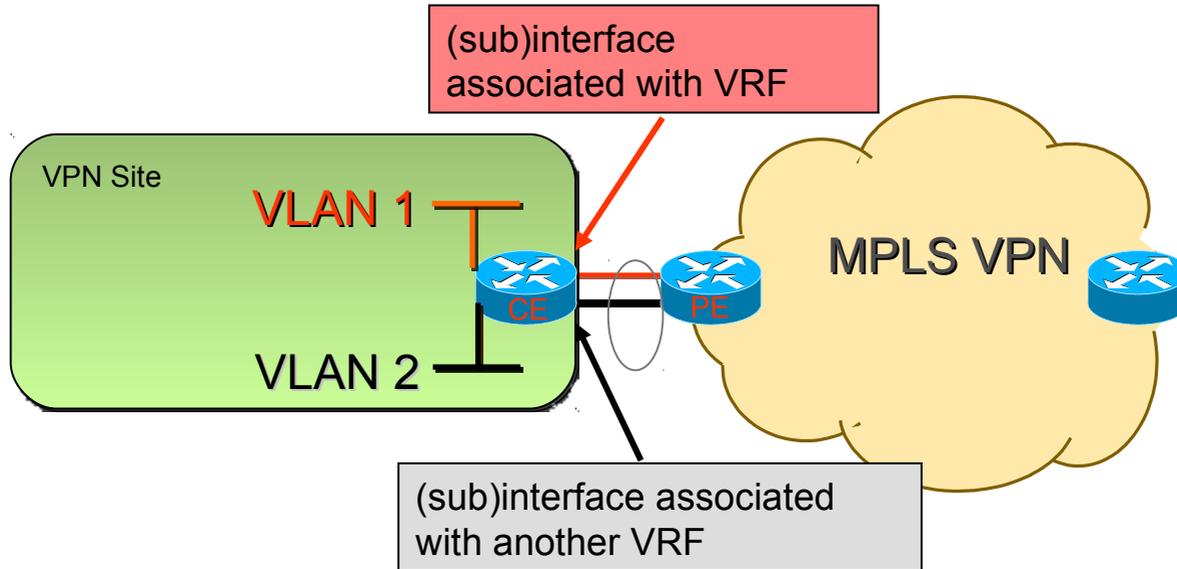
Multi-VRF CE - Extending MPLS-VPN

Clients



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

Multi-VRF CE - a standalone Virtual-router !

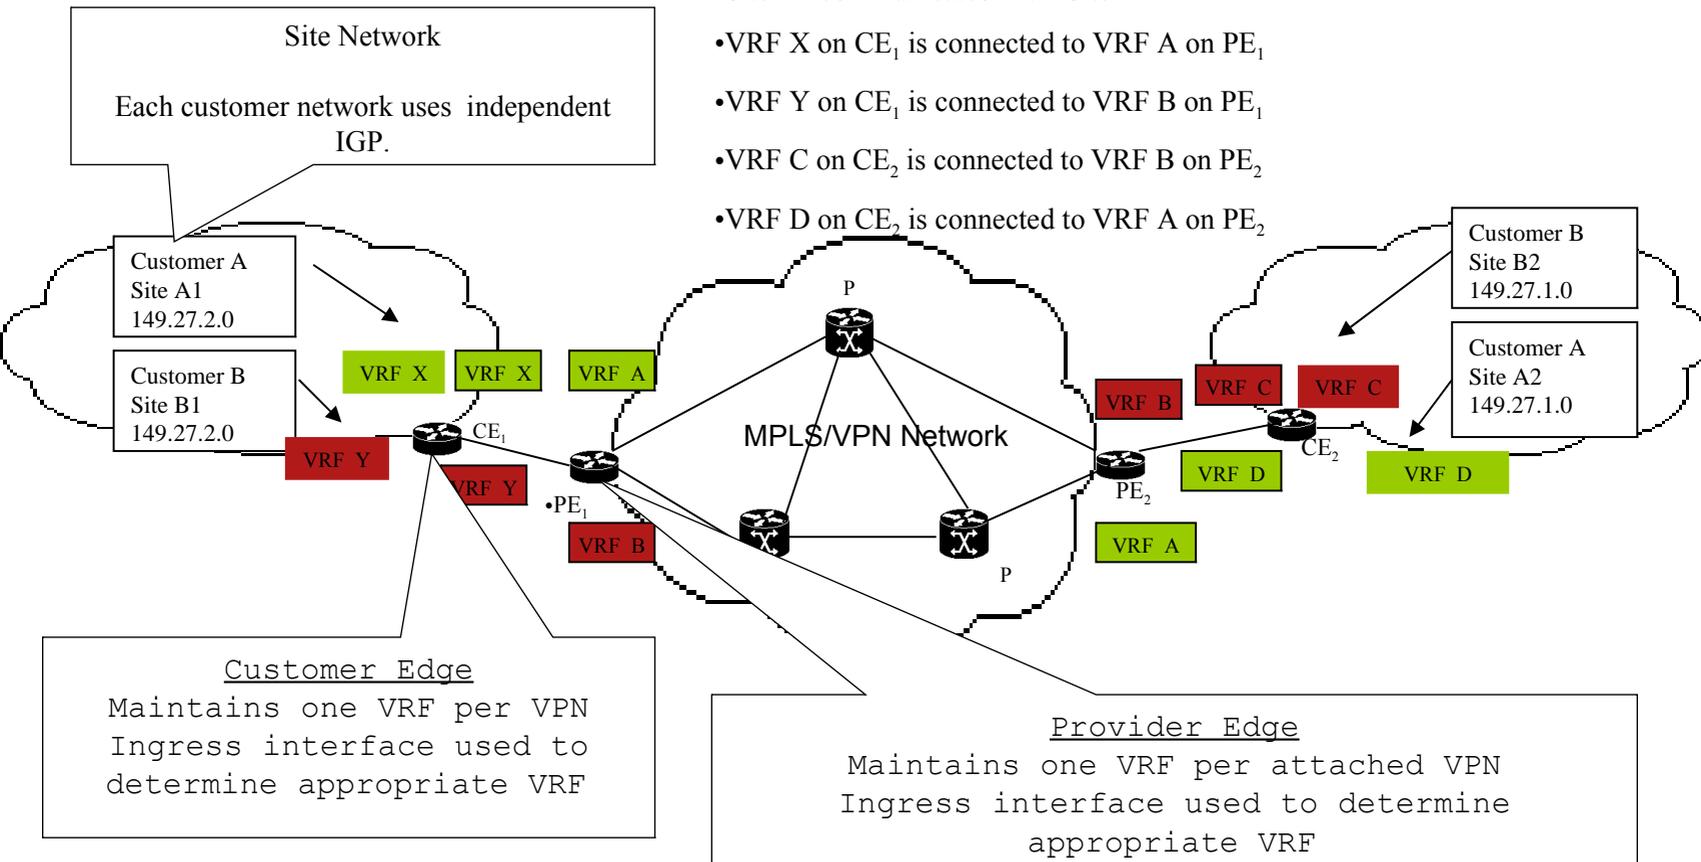


No MPLS, nor MP-iBGP on CE

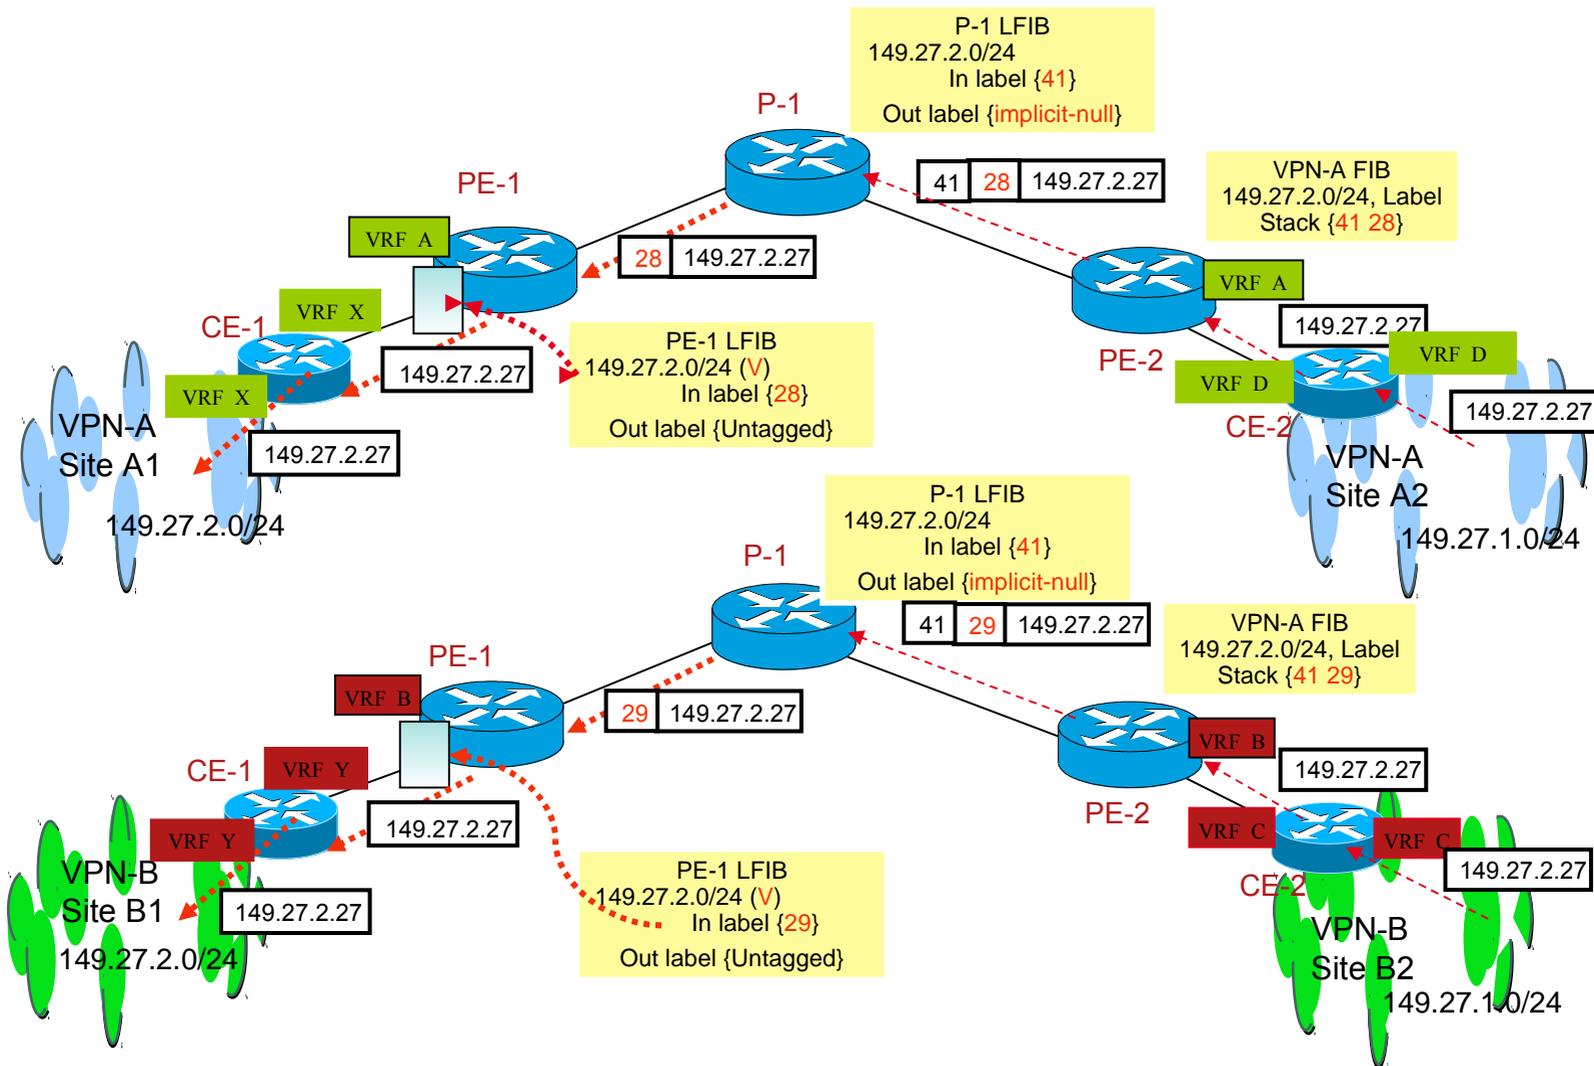
Local Inter-VRF routing is supported

Multi-VRF/VRF-Lite CE Architecture

- Site A1 communicates with Site A2
- Site B1 communicates with SiteB2
- VRF X on CE₁ is connected to VRF A on PE₁
- VRF Y on CE₁ is connected to VRF B on PE₁
- VRF C on CE₂ is connected to VRF B on PE₂
- VRF D on CE₂ is connected to VRF A on PE₂



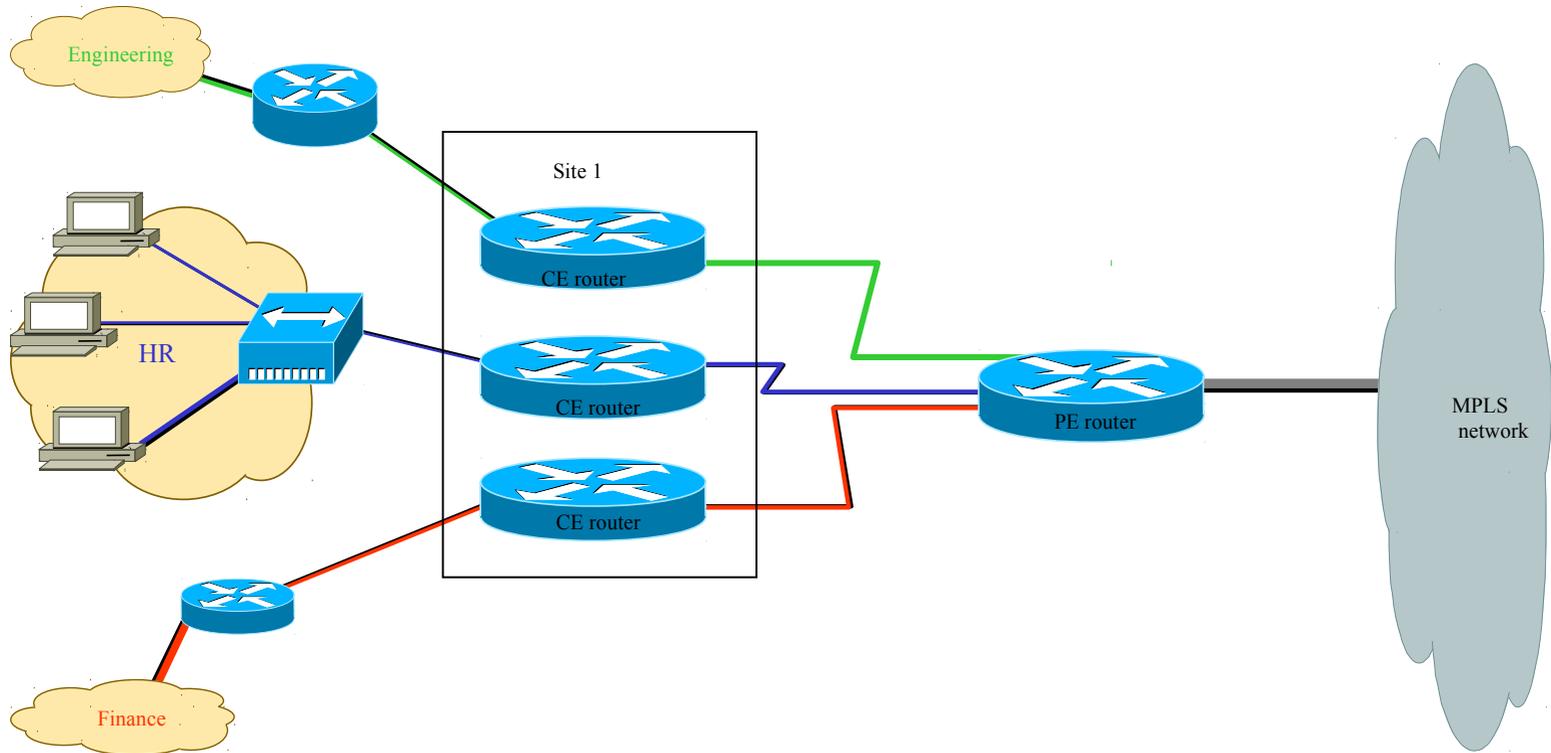
Data Forwarding in MPLS-VPN with Multi-VRF CE



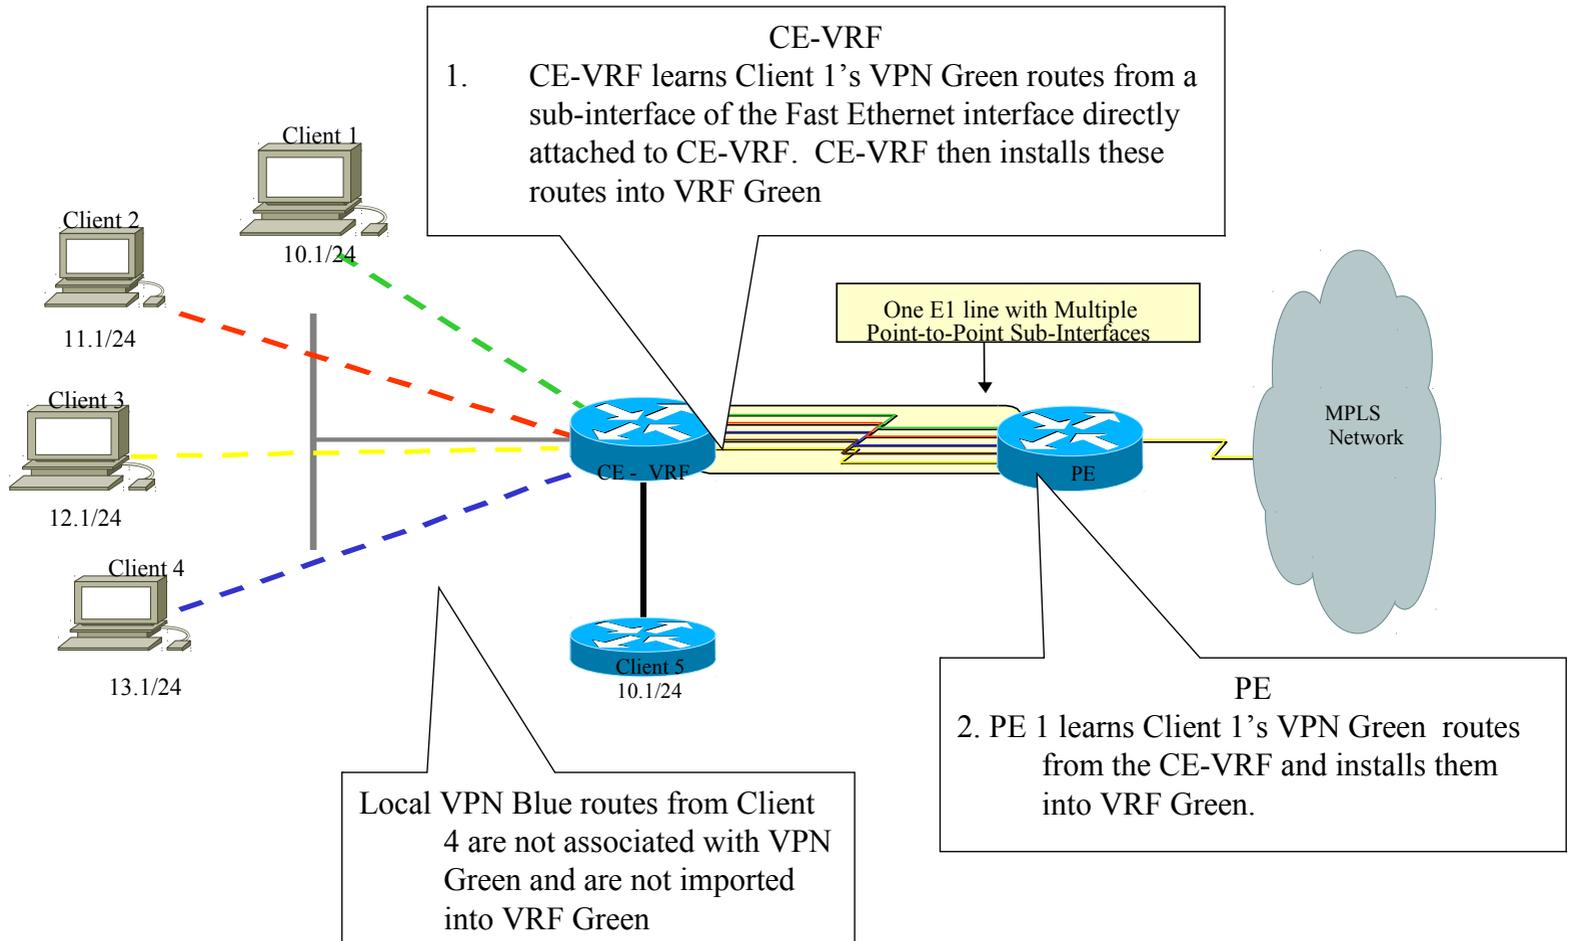
Multi-VRF CE Architecture

- 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



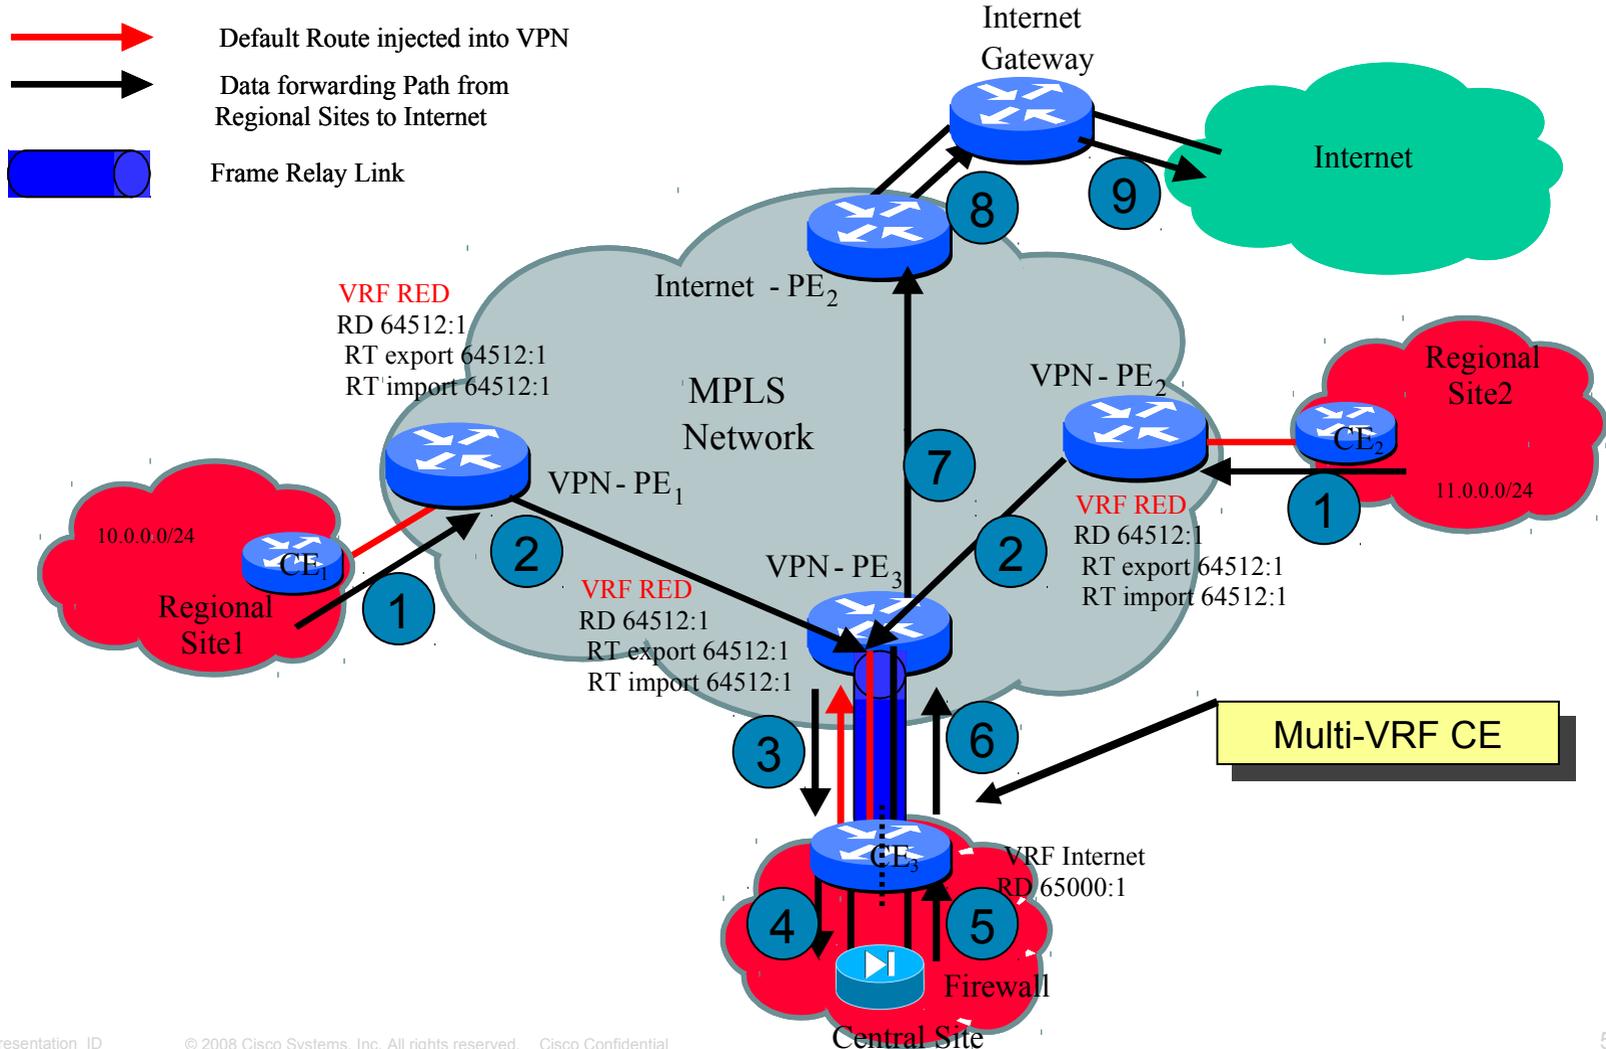
Multi-VRF CE Architecture: Operational Model



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

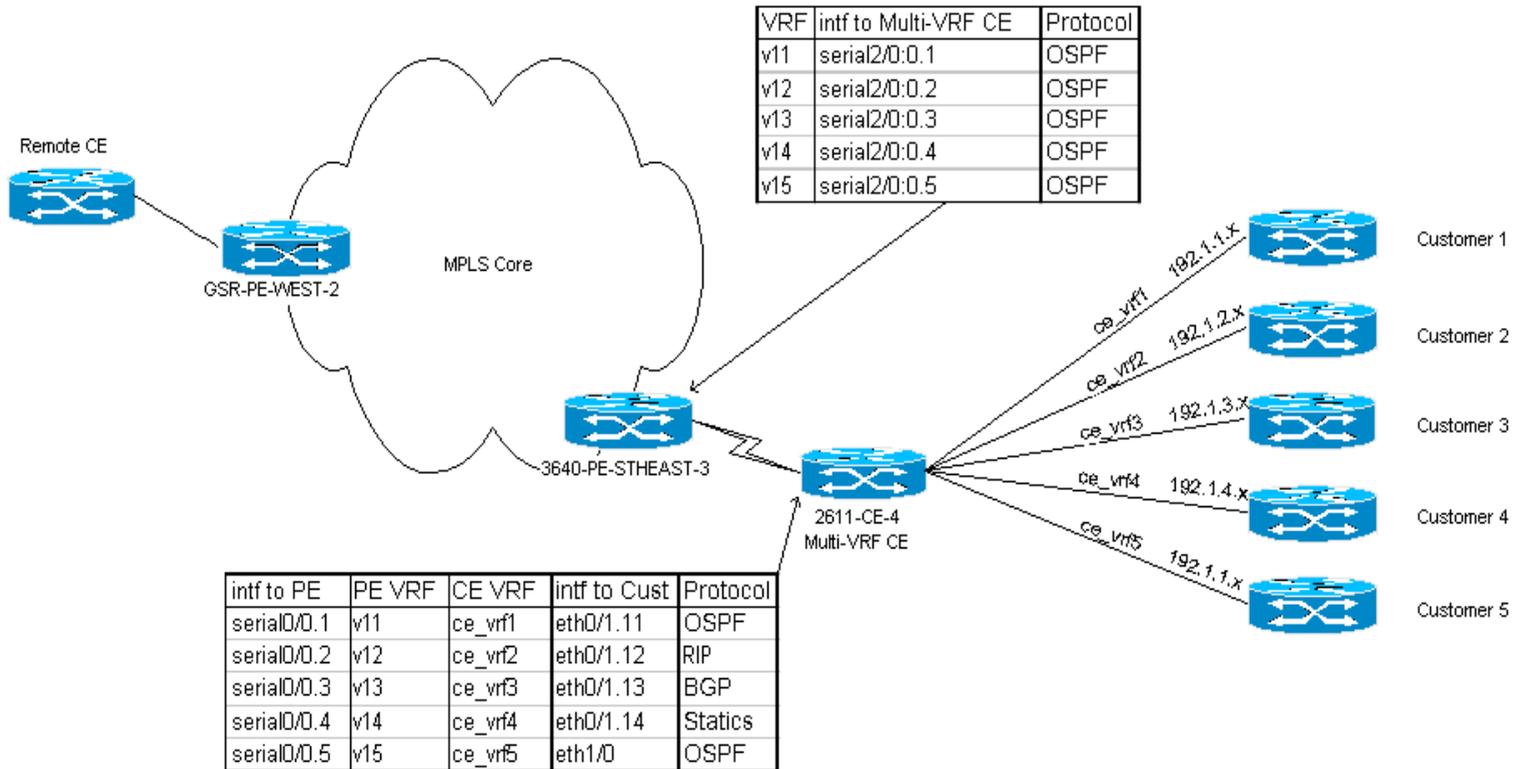
Application 1: Internet Services and VPN Services Using A Single CE



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



Application 2: Basic Setup

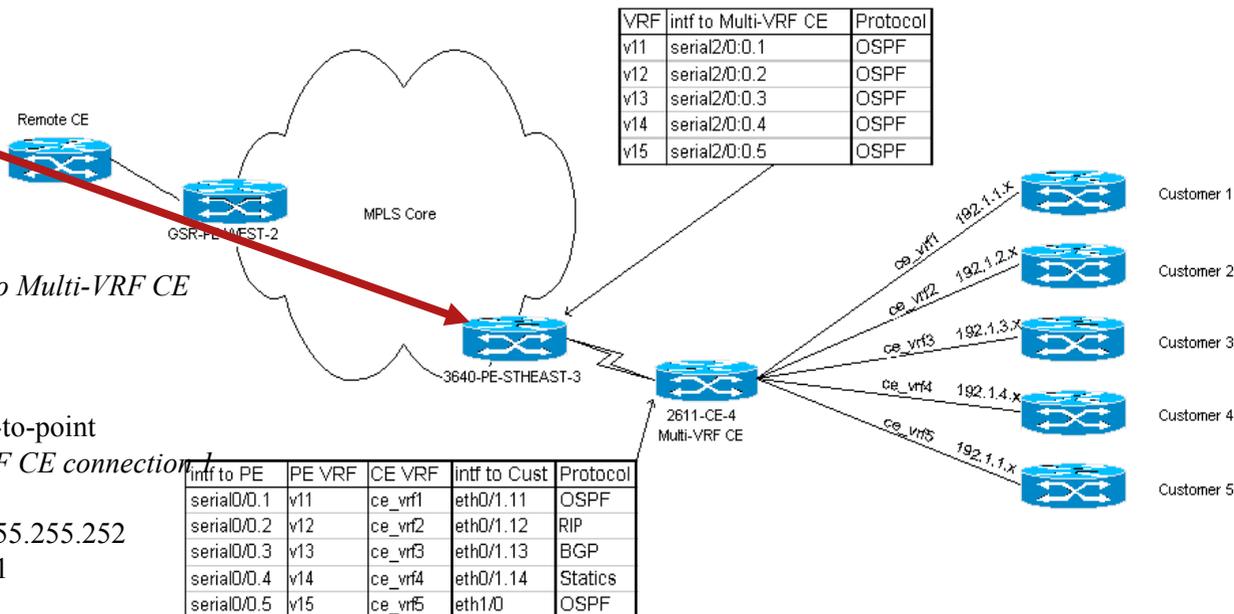
- Inter-site connectivity policies
 - All Customer Routers can communicate with Remote CE's but not with each other.
- All Traffic off 2611-CE-4 is segmented into 5 separate VRFs (labeled ce_vrf1-5)
- 3640-PE-STHEAST-3 uses OSPF as the routing protocol to exchange updates with 2611-CE4, but other routing protocols may be used as well
- All other hosts off 2611-CE4 use a combination of OSPF, EBGP, RIPv2 and static routes

Application 2: Summary Configuration-3640-PE-STHEAST-3

```
ip vrf v11
rd 11:1
route-target export 11:1
route-target import 11:1
!
interface Serial2/0:0
description E1 connection to Multi-VRF CE
no ip address
encapsulation frame-relay
!
```

```
interface Serial2/0:0.1 point-to-point
description PE to Multi-VRF CE connection
ip vrf forwarding v11
ip address 220.1.65.5 255.255.255.252
frame-relay interface-dlci 21
!
```

```
router ospf 11 vrf v11
log-adjacency-changes
area 11 virtual-link 220.1.65.6
##VL if CE is not configured with capability-vrflite
redistribute bgp 30000 subnets
network 220.1.65.4 0.0.0.3 area 11
```



Please keep in mind that “capability-vrflite” knob is not needed on the PE.

Application 2: Summary Configuration- Multi-VRF CE

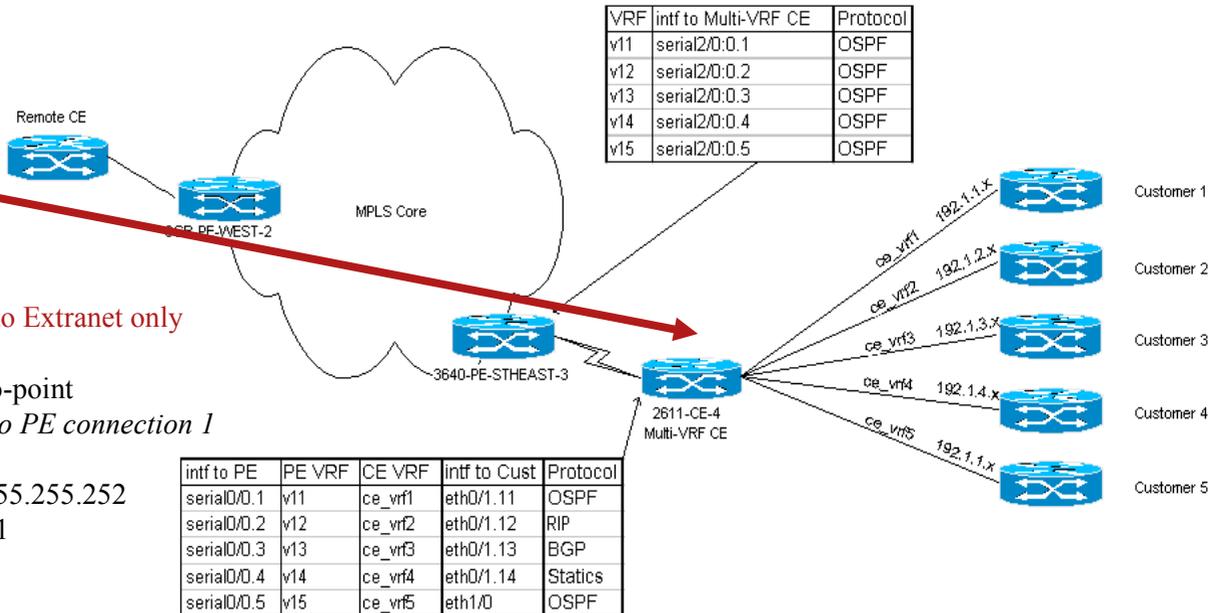
```
ip vrf ce_vrf1
rd 81:81
route-target export 81:1
route-target import 81:1
#Required if you want to do Extranet only
```

```
!
interface Serial0/0.1 point-to-point
description Multi-VRF CE to PE connection 1
ip vrf forwarding ce_vrf1
ip address 220.1.65.6 255.255.255.252
frame-relay interface-dlci 21
```

```
!
interface Ethernet0/1.11
description Multi-VRF CE to host 1 (dup addr)
encapsulation dot1Q 11
ip vrf forwarding ce_vrf1
ip address 192.1.1.1 255.255.255.0
```

```
!
router ospf 11 vrf ce_vrf1
log-adjacency-changes
area 11 virtual-link 220.1.65.5
capability vrf-lite
network 192.1.1.0 0.0.0.255 area 0
network 220.1.65.4 0.0.0.3 area 11
```

OR
[after 12.0(21)ST]



OSPF “Capability vrf-lite”

- To suppress PE-specific checks on a CE-vrf-lite router (OSPF ‘DOWN’ Bit used only in VPNs)
- These checks are required to prevent loops when PE is performing mutual redistribution between OSPF and BGP
- Reference: CSCds82178
- For the Multi-VRF CE these checks may be turned off:

```
router ospf 100 vrf ce_vrf1
capability vrf-lite
```

OSPF “Capability vrf-lite”

- When the OSPF process is associated with the VRF, several checks are performed when LSAs are received:
 - If Type-3 LSA is received, DN bit is checked. If DN bit is set, Type-3 LSA is not considered during the SPF
 - If Type-5/7 LSA is received and the Tag in the LSA is equal to the VPN-tag, Type-5/7 LSA is not considered during the SPF
- These checks are needed to prevent loops when PE is performing a mutual redistribution between OSPF and BGP.

```
3640-PE-STHEAST-3#sh ip route vrf v45
.....<snip>.....
```

Gateway of last resort is not set

```
    220.45.53.0/30 is subnetted, 1 subnets
C      220.45.53.4 is directly connected, Serial2/0:0.5
    200.45.72.0/30 is subnetted, 1 subnets
B      200.45.72.4 [200/0] via 10.13.1.72, 00:39:51
```

After the CE OSPF neighbor comes UP..

```
3640-PE-STHEAST-3#sh ip route vrf v45
.....<snip>.....
```

Gateway of last resort is not set

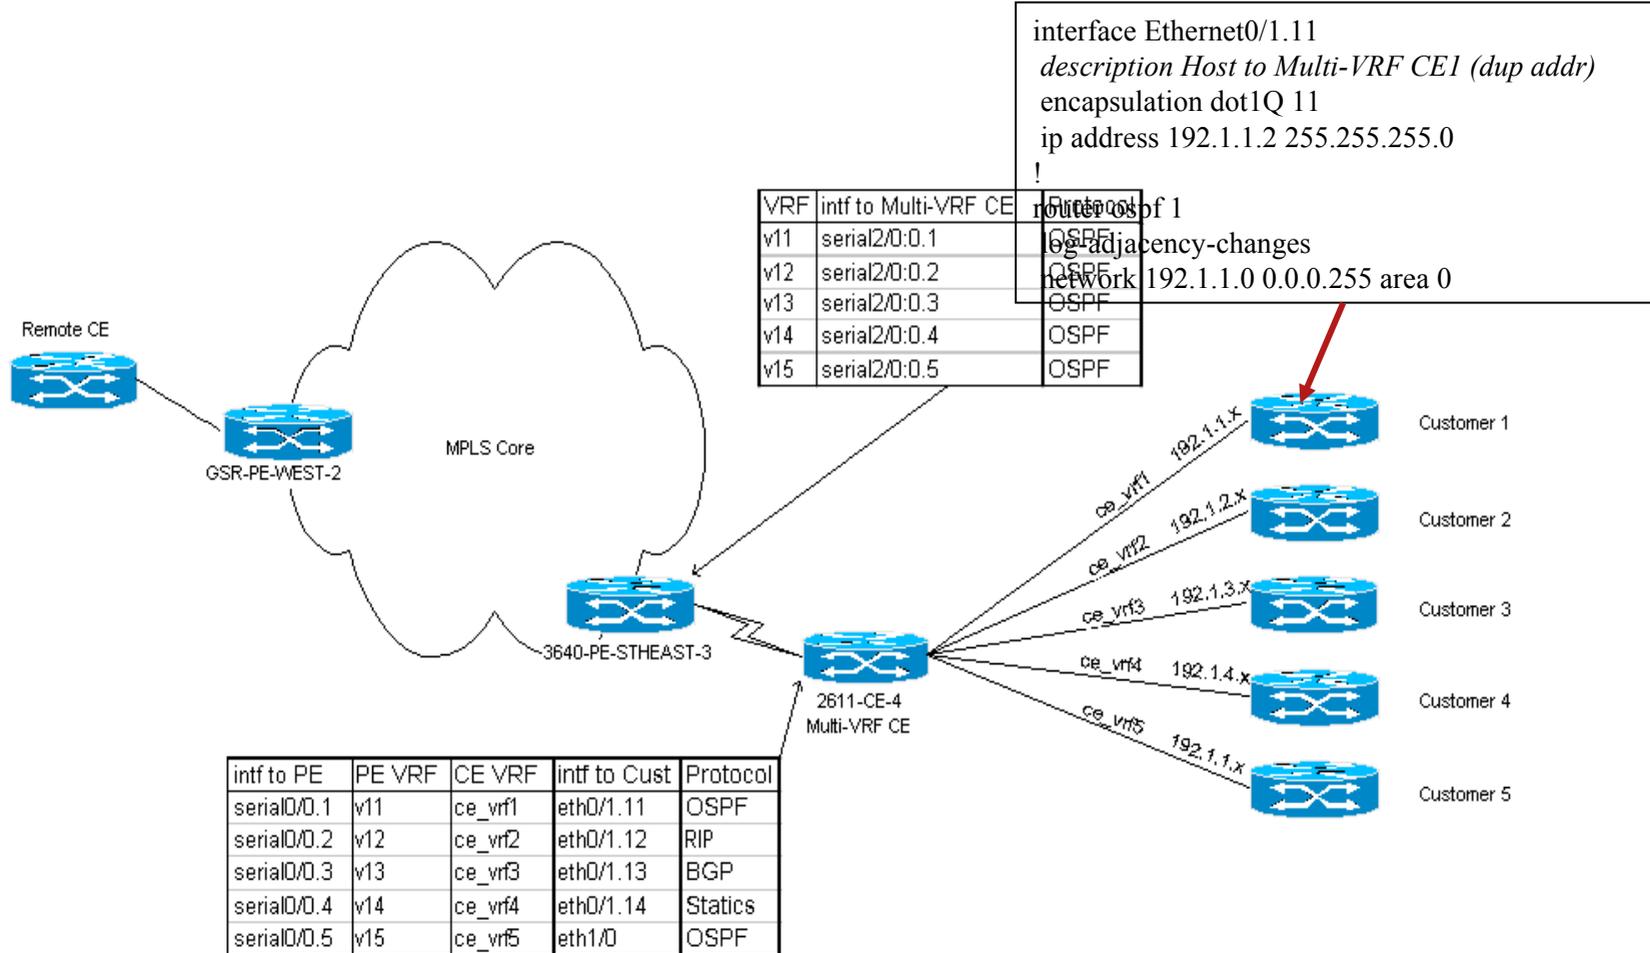
```
O IA 200.41.1.0/24 [110/84] via 220.45.53.6, 00:00:03, Serial2/0:0.5
    220.45.53.0/30 is subnetted, 1 subnets
C      220.45.53.4 is directly connected, Serial2/0:0.5
    200.45.72.0/30 is subnetted, 1 subnets
B      200.45.72.4 [200/0] via 10.13.1.72, 00:40:28
    30.0.0.0/24 is subnetted, 1 subnets
O E2  30.45.106.0 [110/20] via 220.45.53.6, 00:00:03, Serial2/0:0.5
```

```
2611-CE-4#sh ip ospf 45 database summary
      OSPF Router with ID (200.45.72.4) (Process ID 45)
        Summary Net Link States (Area 45)
Routing Bit Set on this LSA
LS age: 637
Options: (No TOS-capability, DC, Downward) <<< Downward => Down (DN) bit set by PE
LS Type: Summary Links(Network)
Link State ID: 200.45.72.4 (summary Network Number)
Advertising Router: 220.45.53.5
LS Seq Number: 800002DB
Checksum: 0x41DC
Length: 28
Network Mask: /24
      TOS: 0  Metric: 10
```

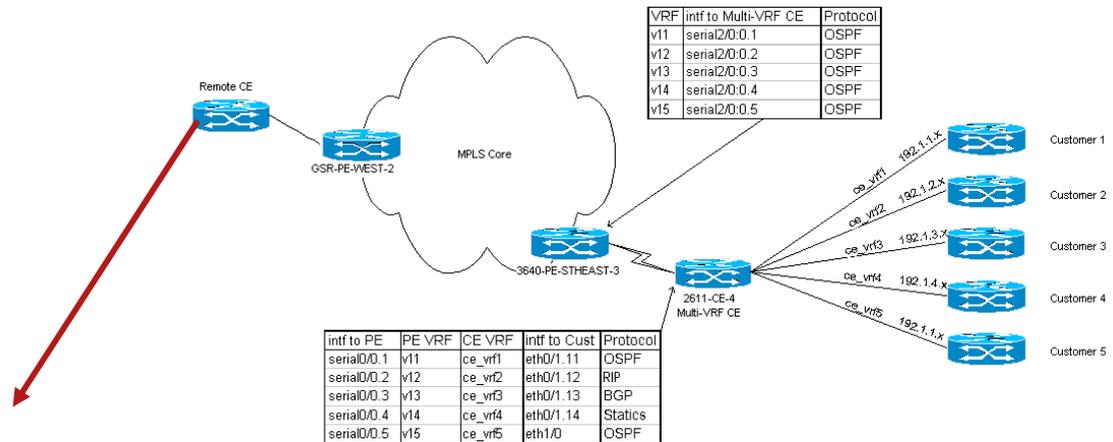
```
3640-PE-STHEAST-3#sh ip os 45 da ex
      OSPF Router with ID (220.45.53.5) (Process ID 45)
        Type-5 AS External Link States
LS age: 430
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 200.41.72.4 (External Network Number )
Advertising Router: 220.45.53.5
LS Seq Number: 80000003
Checksum: 0x5C5F
Length: 36
Network Mask: /30
      Metric Type: 2 (Larger than any link state path)
      TOS: 0
      Metric: 1
      Forward Address: 0.0.0.0
      External Route Tag: 3489690928
```

Application 2: Summary Configuration

Customer 1 Router



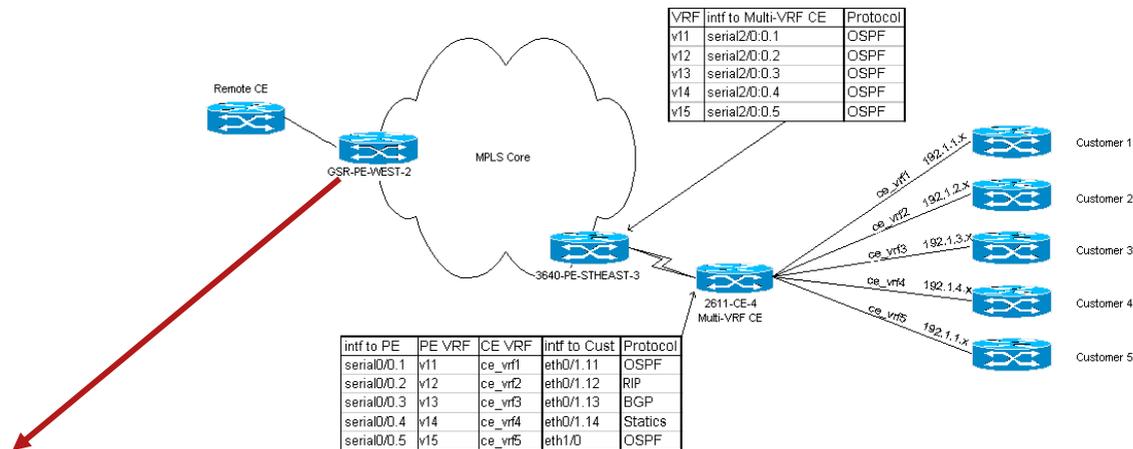
Application 2: Verifying Connectivity- Show Commands Remote CE



```

Remote-CE# sh ip route vrf v15 200.15.44.4
Routing entry for 200.15.44.4/30
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via rip
  Advertised by rip
  Routing Descriptor Blocks:
  * directly connected, via Serial4/3.15
    Route metric is 0, traffic share count is 1
    
```

Application 2: Verifying Connectivity- Show Commands GSR-PE-WEST-2



```
GSR-PE-WEST-2# sh ip route vrf v15 200.15.44.4
```

Routing entry for 200.15.44.4/30

Known via "connected", distance 0, metric 0 (connected, via interface)

Redistributing via bgp 30000

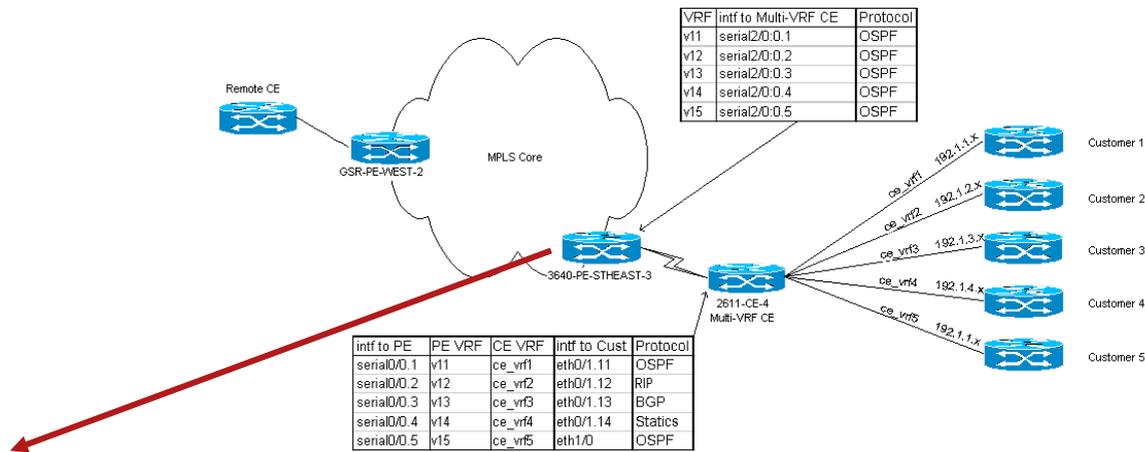
Advertised by bgp 30000

Routing Descriptor Blocks:

* directly connected, via Serial1/0/7.15

Route metric is 0, traffic share count is 1

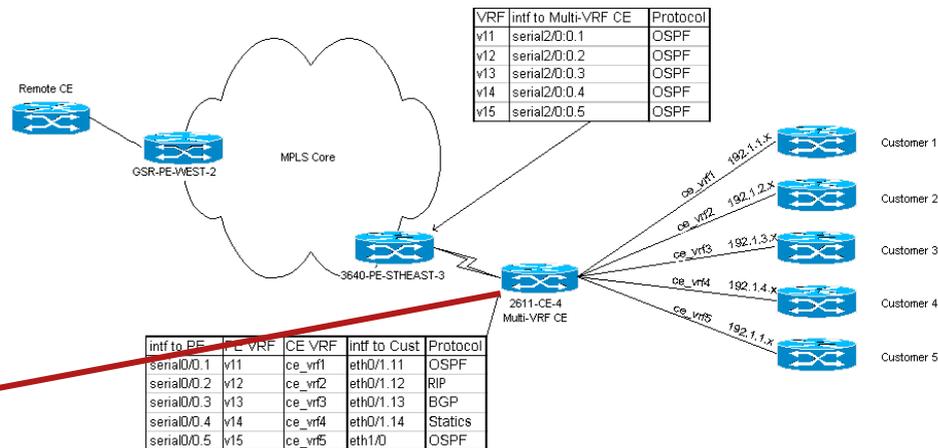
Application 2: Verifying Connectivity- Show Commands 3640-PE-STHEAST3



```

3640-PE-STHEAST-3# sh ip route vrf v15 200.15.44.4
Routing entry for 200.15.44.4/30
  Known via "bgp 30000", distance 200, metric 0, type internal
  Redistributing via ospf 15
  Advertised by ospf 15 subnets
  Last update from 10.13.1.44 00:17:10 ago
  Routing Descriptor Blocks:
  * 10.13.1.44 (Default-IP-Routing-Table), from 10.13.1.48, 00:17:10 ago
    Route metric is 0, traffic share count is 1
    AS Hops 0
    
```

Application 2: Verifying Connectivity- Show Commands Multi-VRF CE



```
2621-CE-4#sh ip route vrf ce_vrf5 200.15.44.4
```

Routing entry for 200.15.44.4/30

Known via "ospf 45" distance 100, metric 1

Tag Complete, Path Length = 1, AS 30000, Type extern 2, forward metric 74

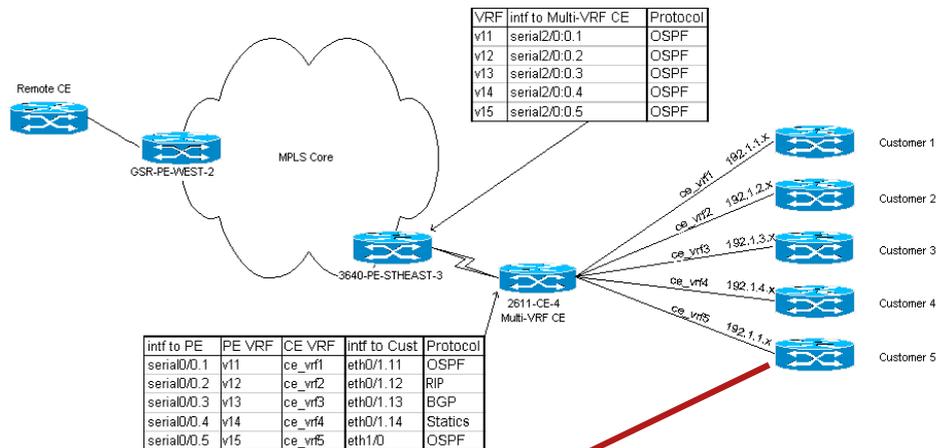
Last update from 220.45.65.21 on Serial 0/0.5, 11:03:35 ago

Routing Descriptor Blocks:

* 200.15.44.4, from 220.45.65.21, 11:03:35 ago, via Serial 0/0.5

Route metric is 1, traffic count is 1

Application 2: Verifying Connectivity- Show Commands Customer 5 Router



```
Customer-5# sh ip route | include 200.15.44.4
O E2 200.15.44.4 [110/1] via 192.1.1.1, 00:16:16, Ethernet1/0
```

```
Customer-5# sh ip route 200.15.44.4
Routing entry for 200.15.44.4/30
  Known via "ospf 5", distance 110, metric 1
  Tag Complete, Path Length == 1, AS 1, , type extern 2, forward metric 84
  Last update from 192.1.1.1 on Ethernet1/0, 00:02:12 ago
  Routing Descriptor Blocks:
  * 192.1.1.1, from 220.1.65.21, 00:02:12 ago, via Ethernet1/0
    Route metric is 1, traffic share count is 1
```

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



MPLS workshop

MPLS VPN Inter-Provider Solutions

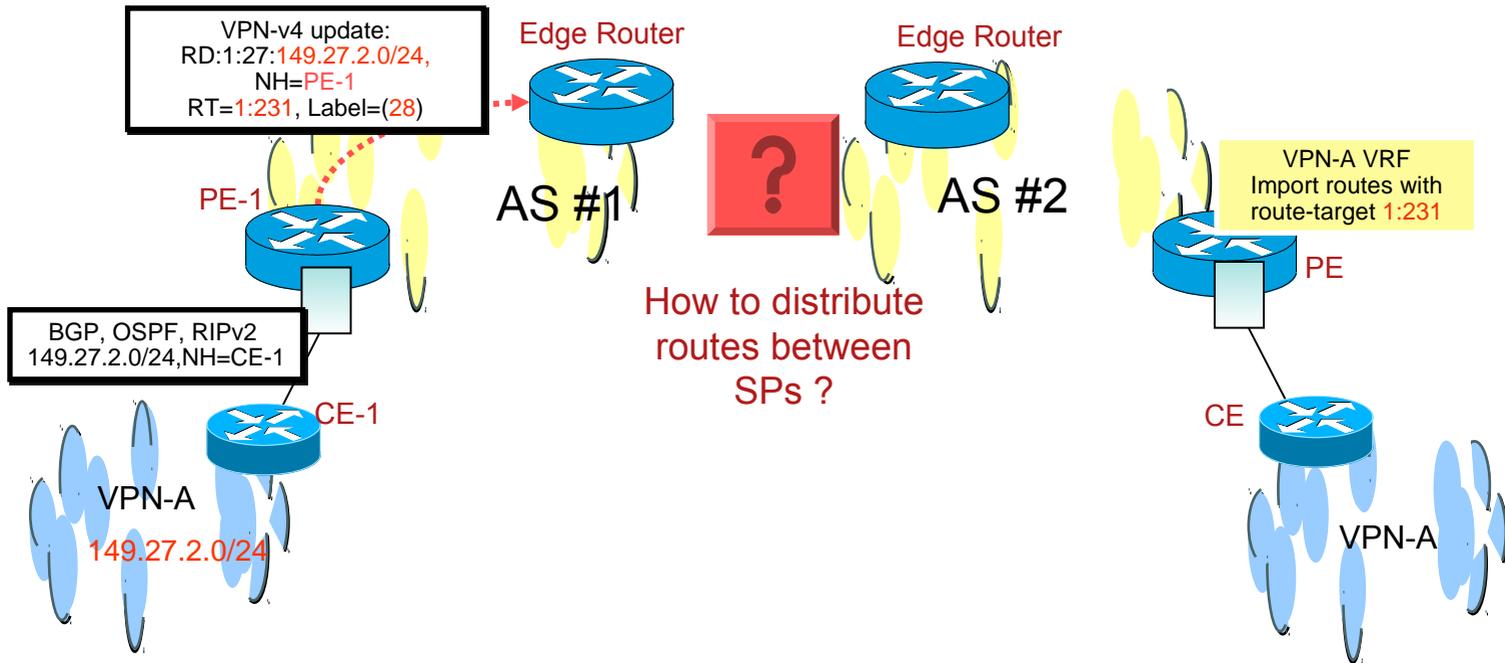
Agenda

- Inter-Provider Connectivity Options
- Scaling Inter-Provider Solutions
- Filtering & Route Distribution Mechanisms
- Distribution of Traffic Load between Providers

VPN Client Connectivity

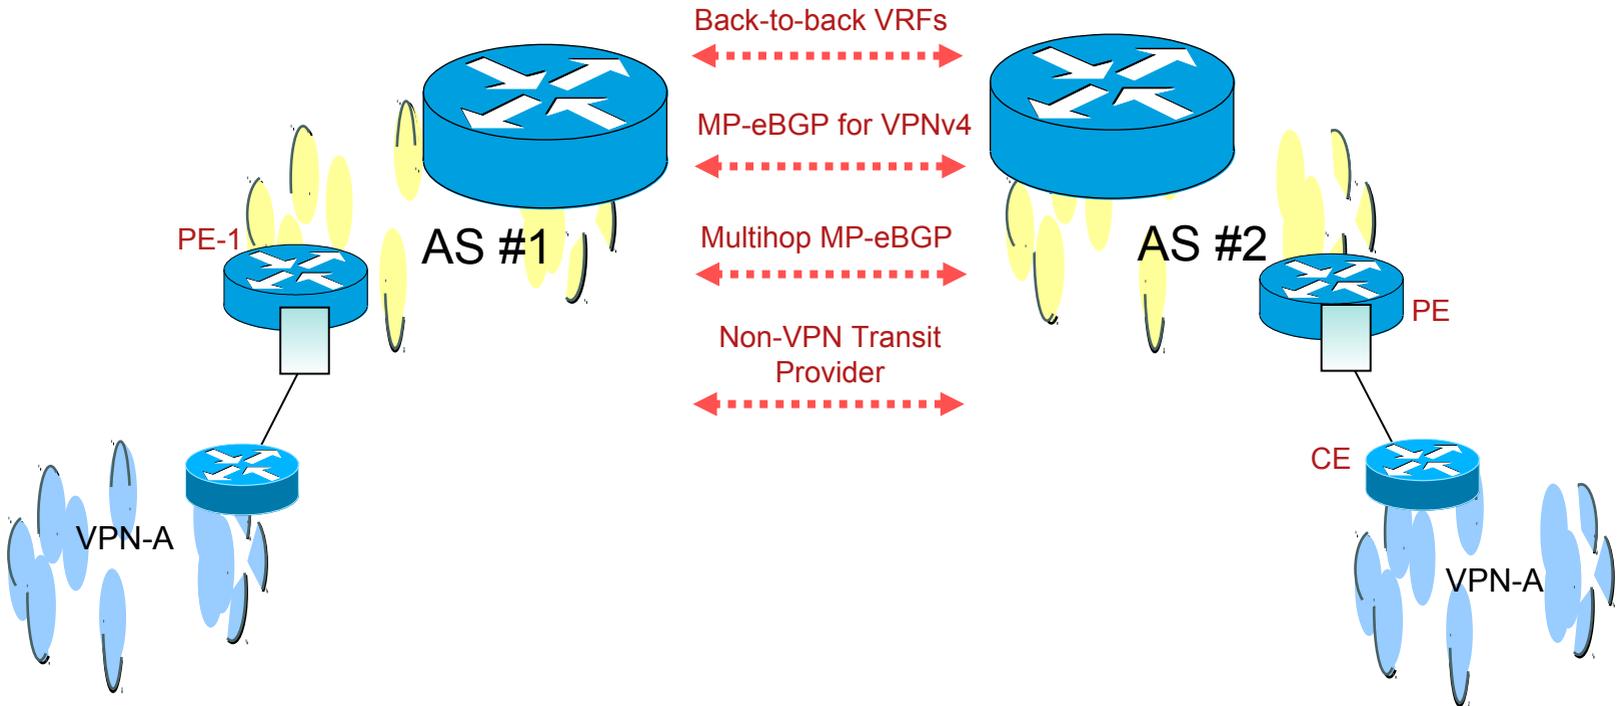
- VPN sites may be geographically dispersed
 - requiring connectivity to separate MPLS VPN Service Providers
- Transit between VPN sites may pass through multiple providers MPLS VPN backbones
 - this implies exchange of VPN routing information between providers
- Referred to as **Multi-Provider** or **Inter-Provider** VPN

VPN Client Connectivity



VPN Sites attached to different Service Providers

VPNv4 Distribution Options

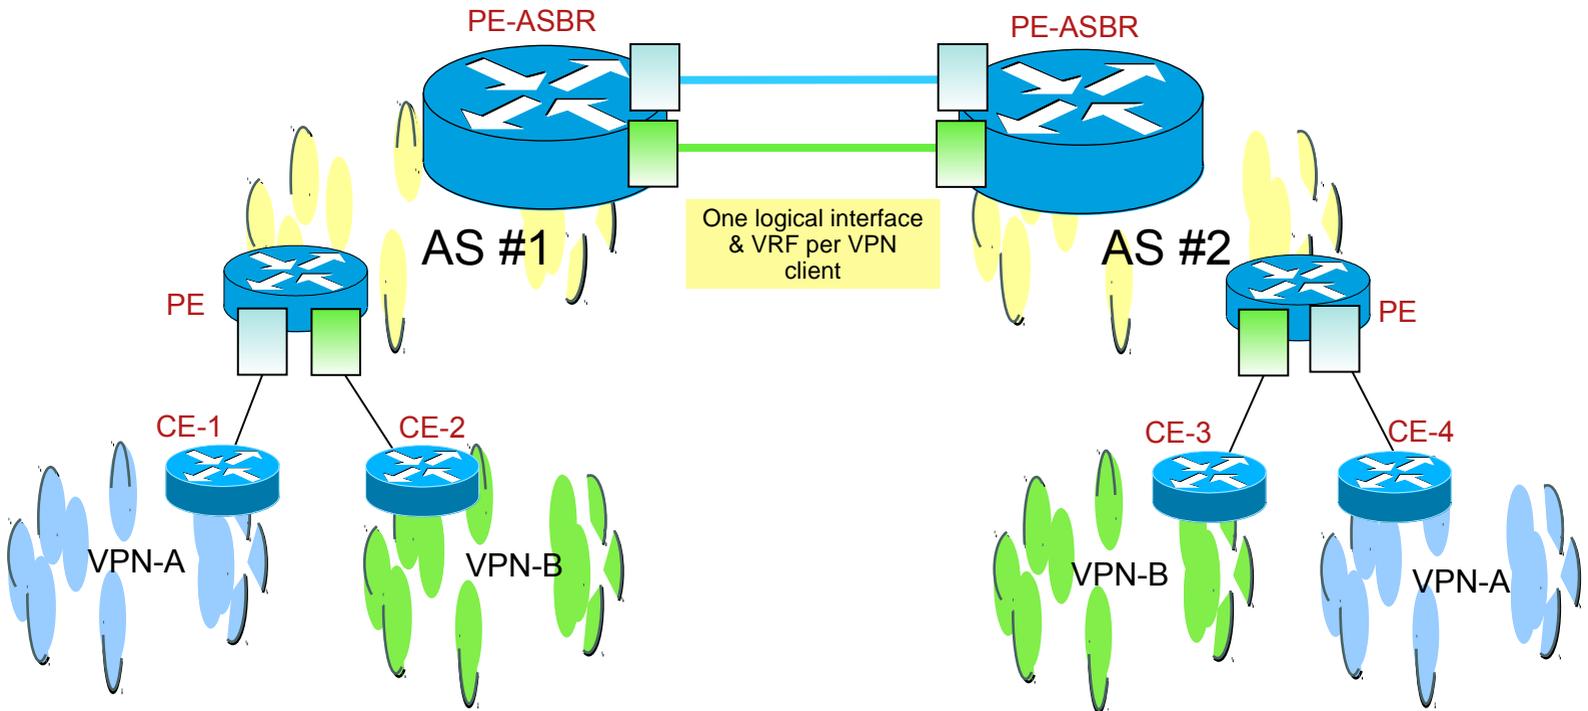


Many options for distribution of VPNv4 prefix information

Back-to-back VRF Connectivity

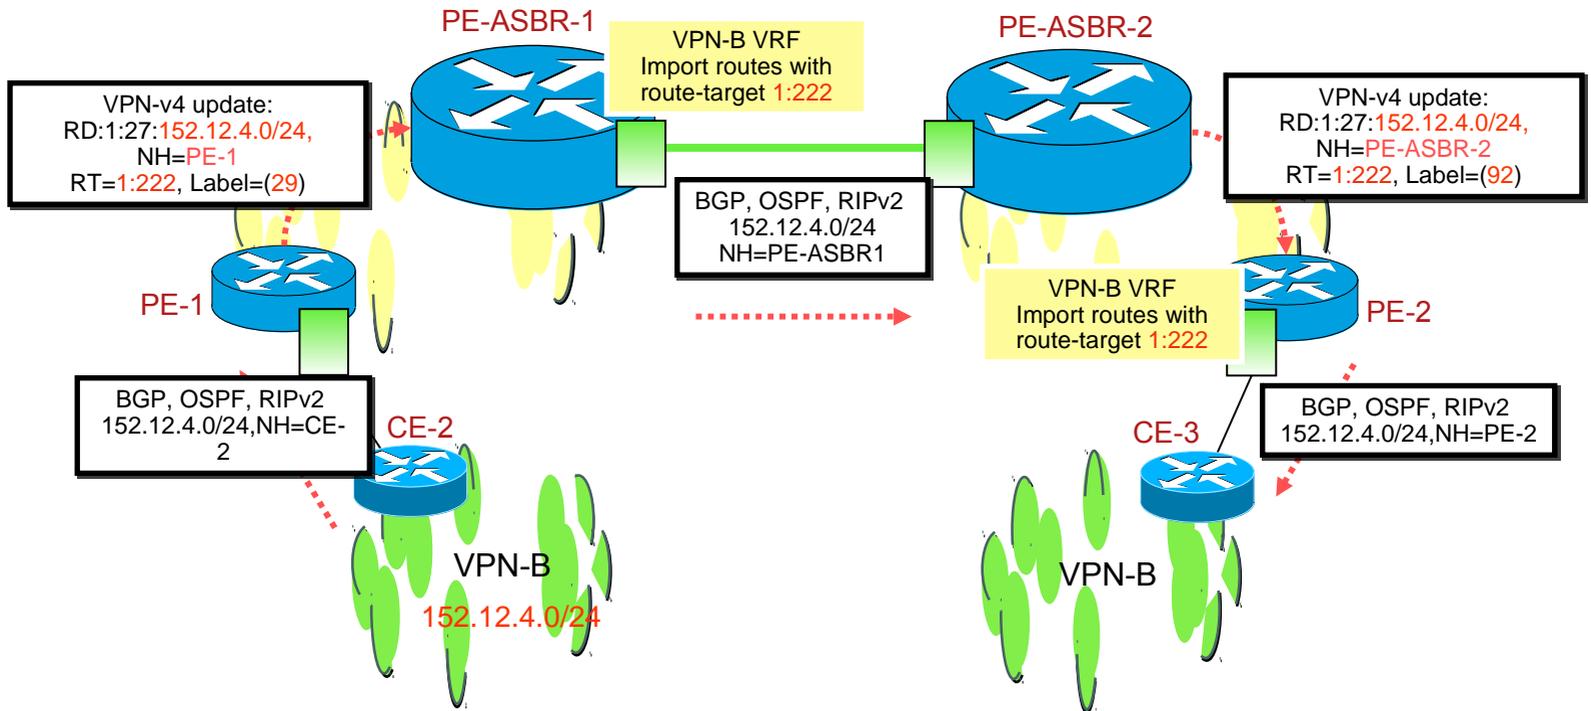
- MPLS VPN providers exchange routes across VRF interfaces
 - VRF represents a particular VPN client
- Each provider PE router treats the other as a CE
 - although both provider interfaces associated with a VRF
- PE routers are gateways used for VPNv4 route exchange
- PE-ASBR to PE-ASBR link may use any supported PE-CE routing protocol
 - currently OSPF, BGP-4, RIPv2 and static

Back-to-back VRF Connectivity



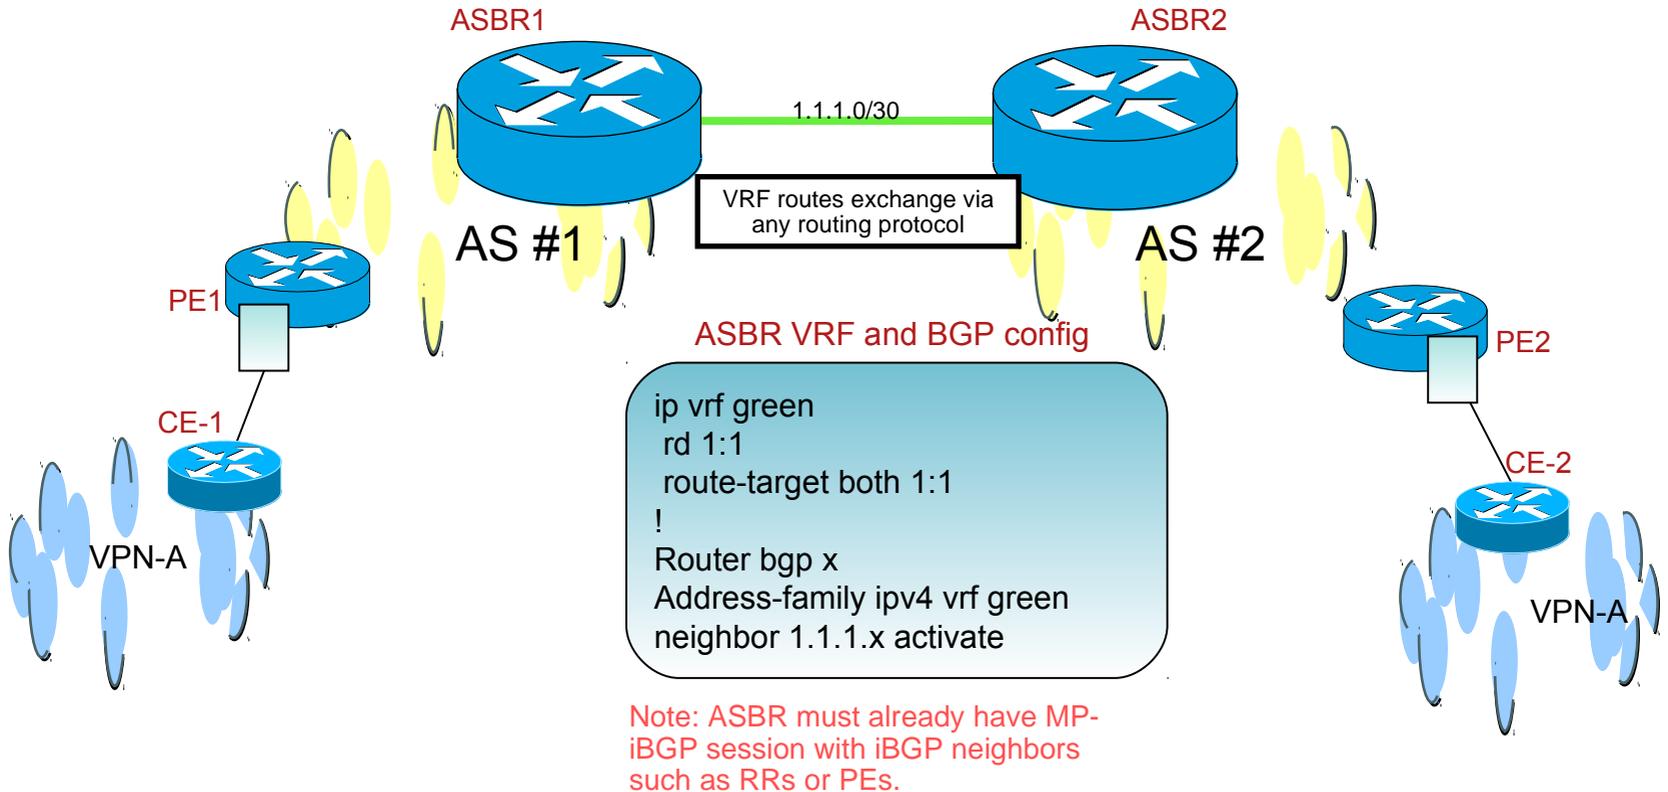
VRF to VRF Connectivity between PE-ASBRs

Back-to-back VRF Connectivity



VRF to VRF Connectivity between PE-ASBRs

Back-to-Back VRF Connectivity



Back-to-back VRF Connectivity

- Scalability is an issue with many VPNs
 - One VRF & logical interface required per VPN client;
 - Gateway PE-ASBR must hold ALL routing information
- PE-ASBR must filter & store VPNv4 prefixes
 - Plus import into VRFs thus increasing MPLS, CEF & routing table memory
- No MPLS required between providers
 - Standard IP between gateway PE-ASBRs;
 - No exchange of routes using MP-eBGP

MP-eBGP between ASBRs for VPNv4

- New CLI “no bgp default route-target filter” is needed on the ASBR to accept VPNv4 prefixes in the absence of VRFs
- PE-ASBRs exchange routes directly using BGP VPNv4 AF MP-eBGP for VPNv4 prefix exchange. No LDP required
- eBGP session with next-hop set to advertising PE-ASBR
Next-hop and labels are rewritten when advertised across the Inter-Provider MP-eBGP session
- PE-ASBR stores all VPN routes which must be exchanged
But only in the BGP table
Labels are populated into the LFIB of the PE-ASBR

MP-eBGP between ASBRs for VPNv4

- Receiving Gateway PE-ASBRs may allocate new label if desired

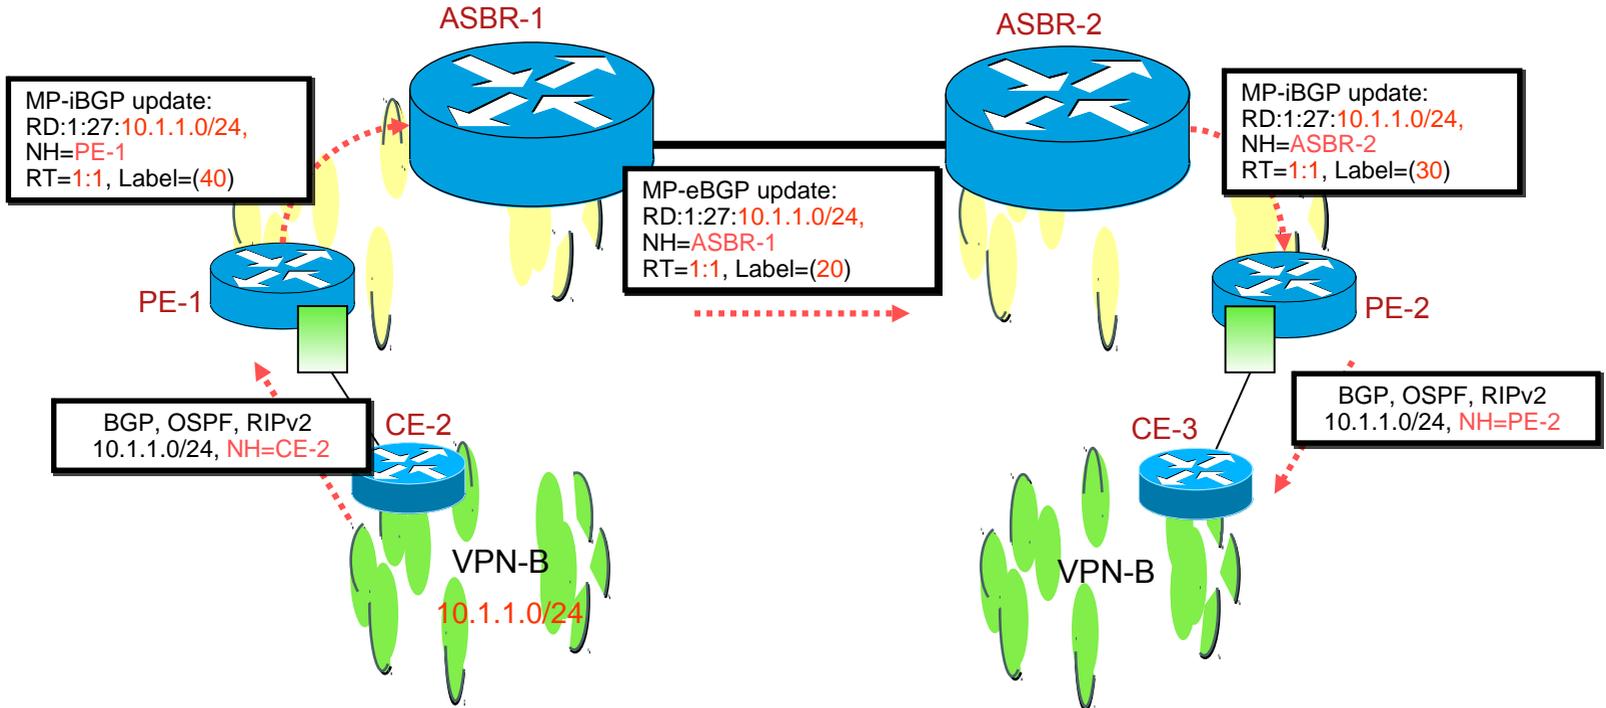
Controlled by configuration of next-hop-self (default is on)

- Receiving PE-ASBR will automatically create a /32 host route for its PE-ASBR neighbor

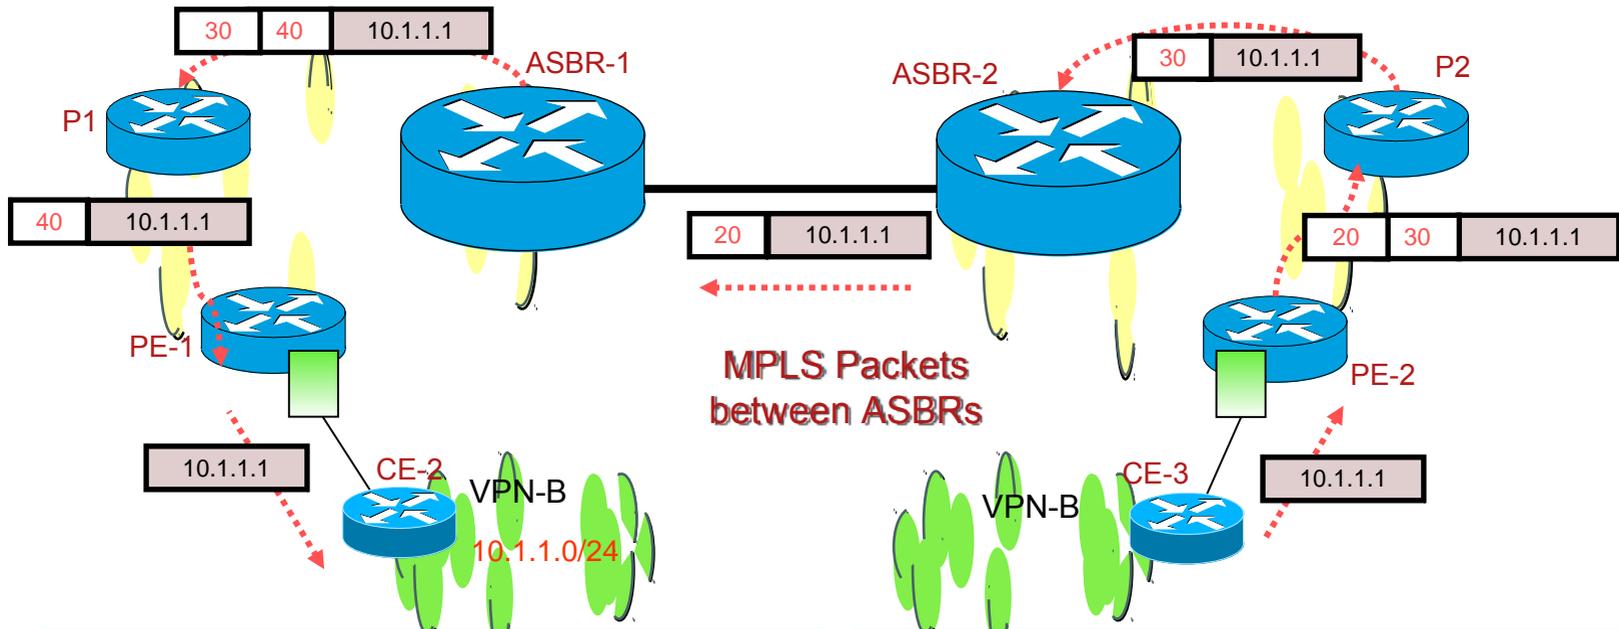
Which must be redistributed into receiving IGP if next-hop-self is NOT in operation;

/32 not created if iBGP session, eBGP multihop or if MP-eBGP exchange of VPNv4 capability not negotiated with neighbor

MP-eBGP between ASBRs for VPNv4 Control Plane



MP-eBGP between ASBRs for VPNv4 Forwarding Plane



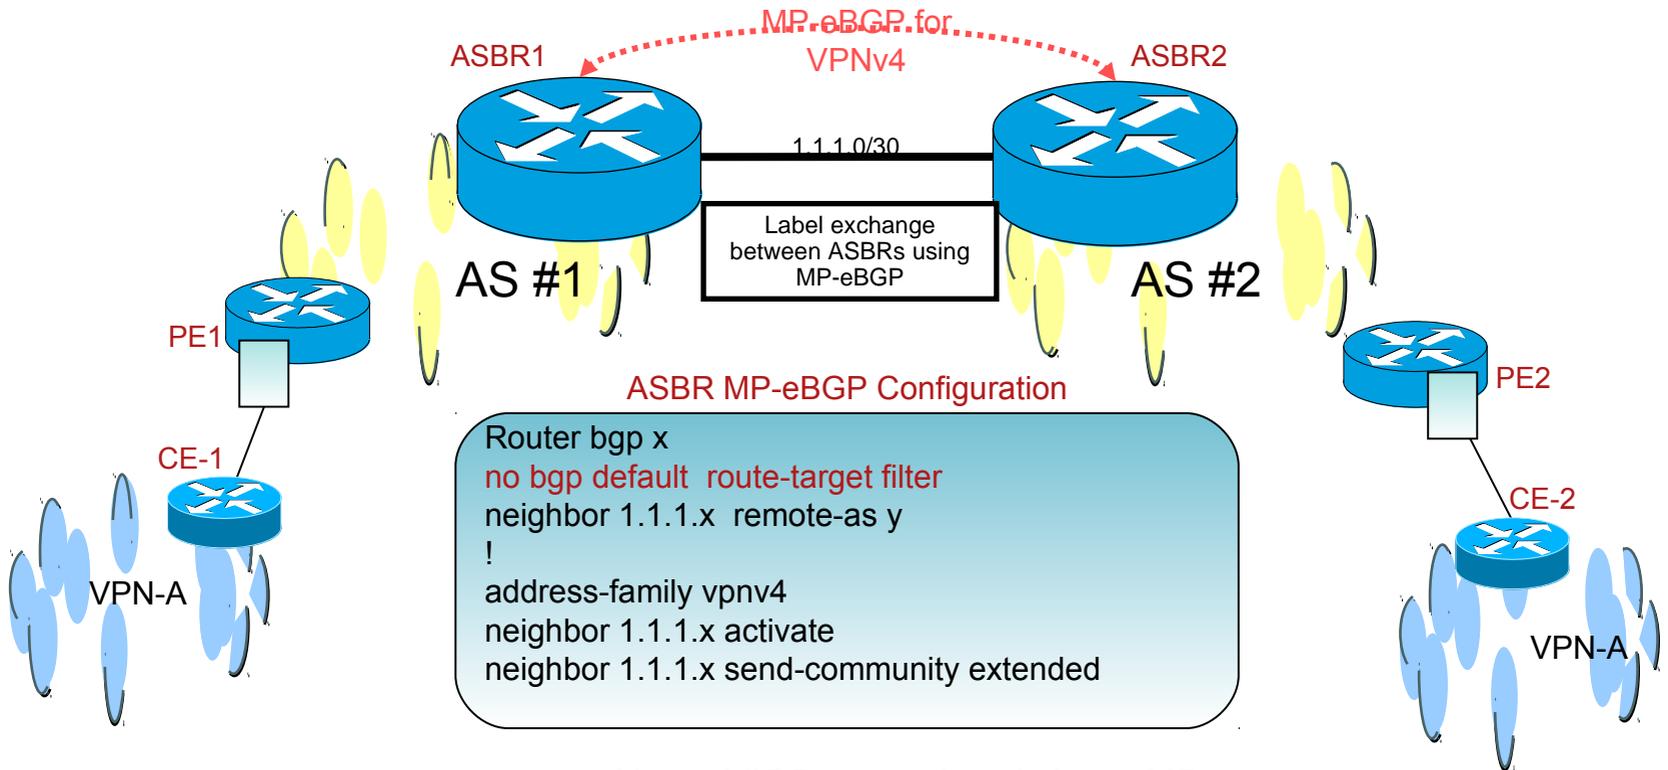
Pros

- More scalable.
Only one interface between ASBRs routers
No VRF configuration on ASBR.
Less memory consumption (no RIB/FIB memory)
- MPLS label switching between providers
Still simple, more scalable & works today

Cons

- Automatic Route Filtering must be disabled
But we can apply BGP filtering.
- ASBRs are still required to hold VPN routes

MP-eBGP between ASBRs for VPNv4 IOS Configuration

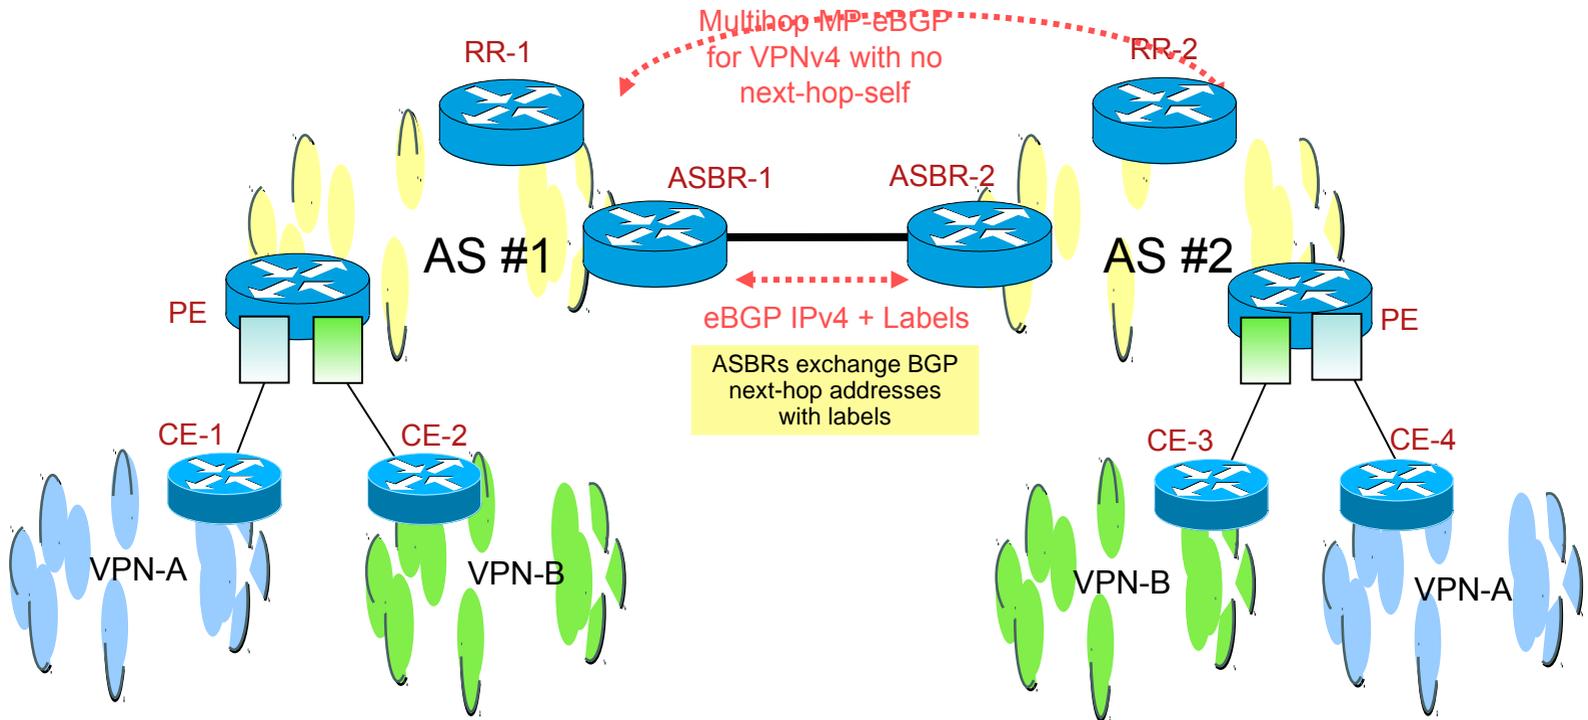


Note: ASBR must already have MP-iBGP session with iBGP neighbors such as RRs or PEs.

Multihop MP-eBGP for VPNv4

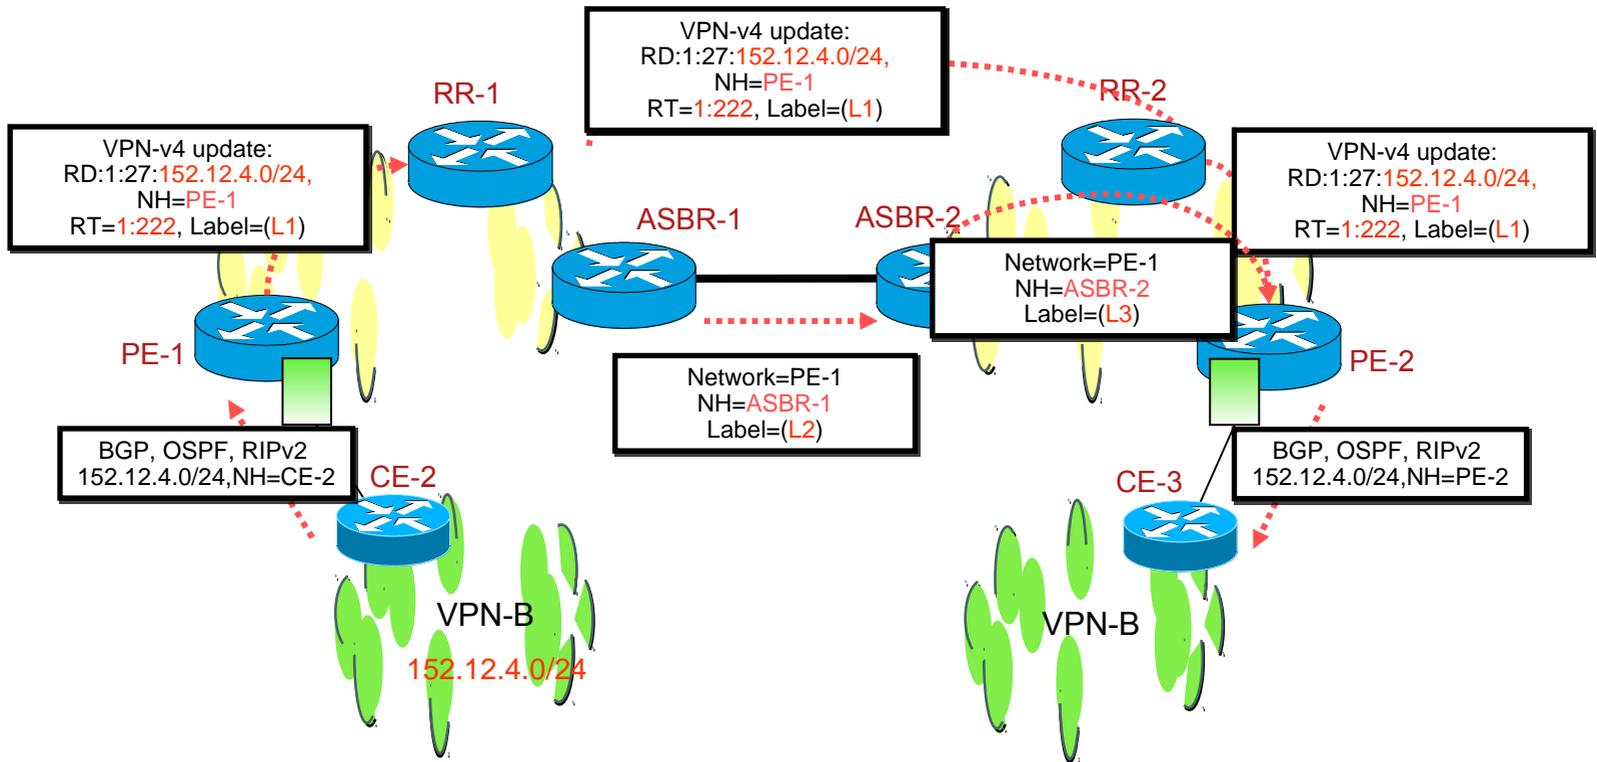
- MPLS VPN providers exchange VPNv4 prefixes via their Route Reflectors
 - Requires Multihop MP-eBGP (VPNv4 routes)
- Next-hop-self MUST be disabled on Route Reflector
 - Preserves next-hop and label as allocated by the originating PE router
- Providers exchange IPv4 routes with labels between directly connected ASBRs using eBGP

Multihop MP-eBGP for VPNv4

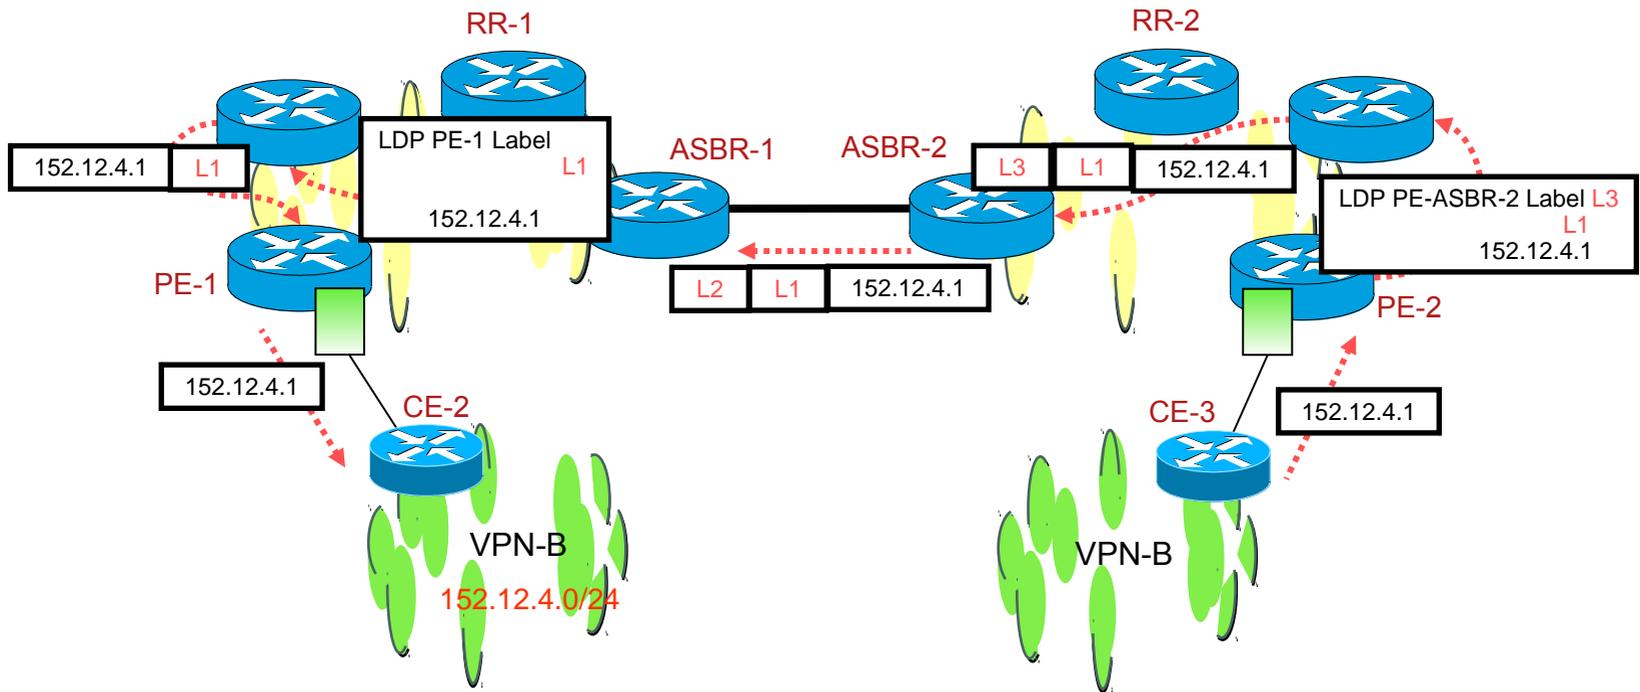


Multihop MP-eBGP VPNv4 prefix exchange between Route Reflectors

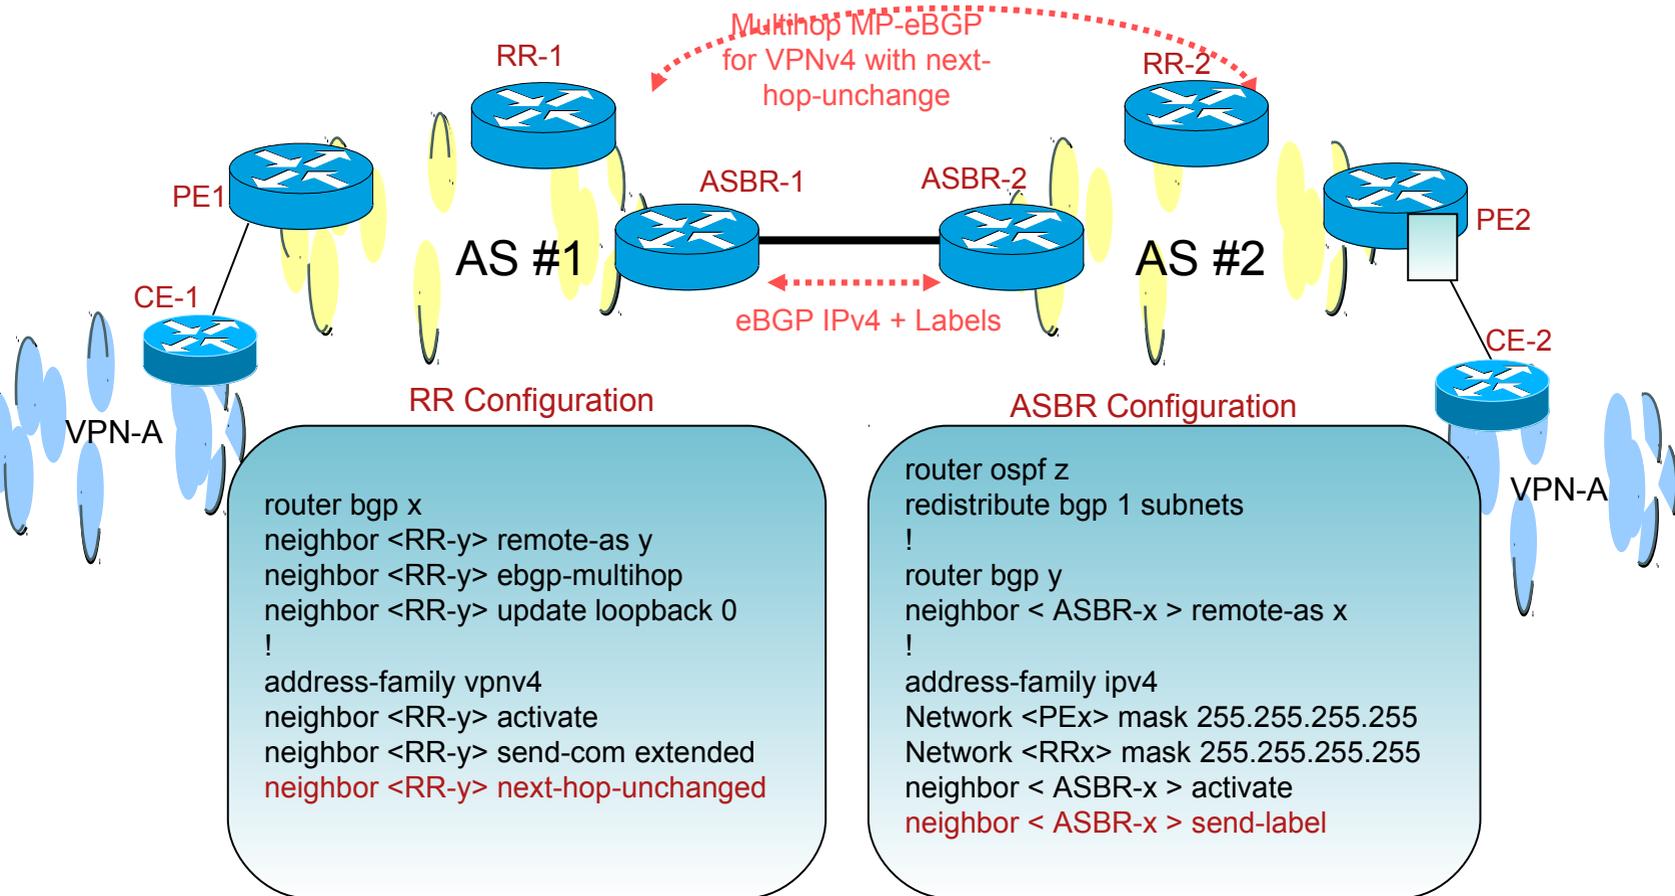
Multihop MP-eBGP for VPNv4 Control Plane



Multihop MP-eBGP for VPNv4 Forwarding Plane



Multihop MP-eBGP for VPNv4 IOS Configuration



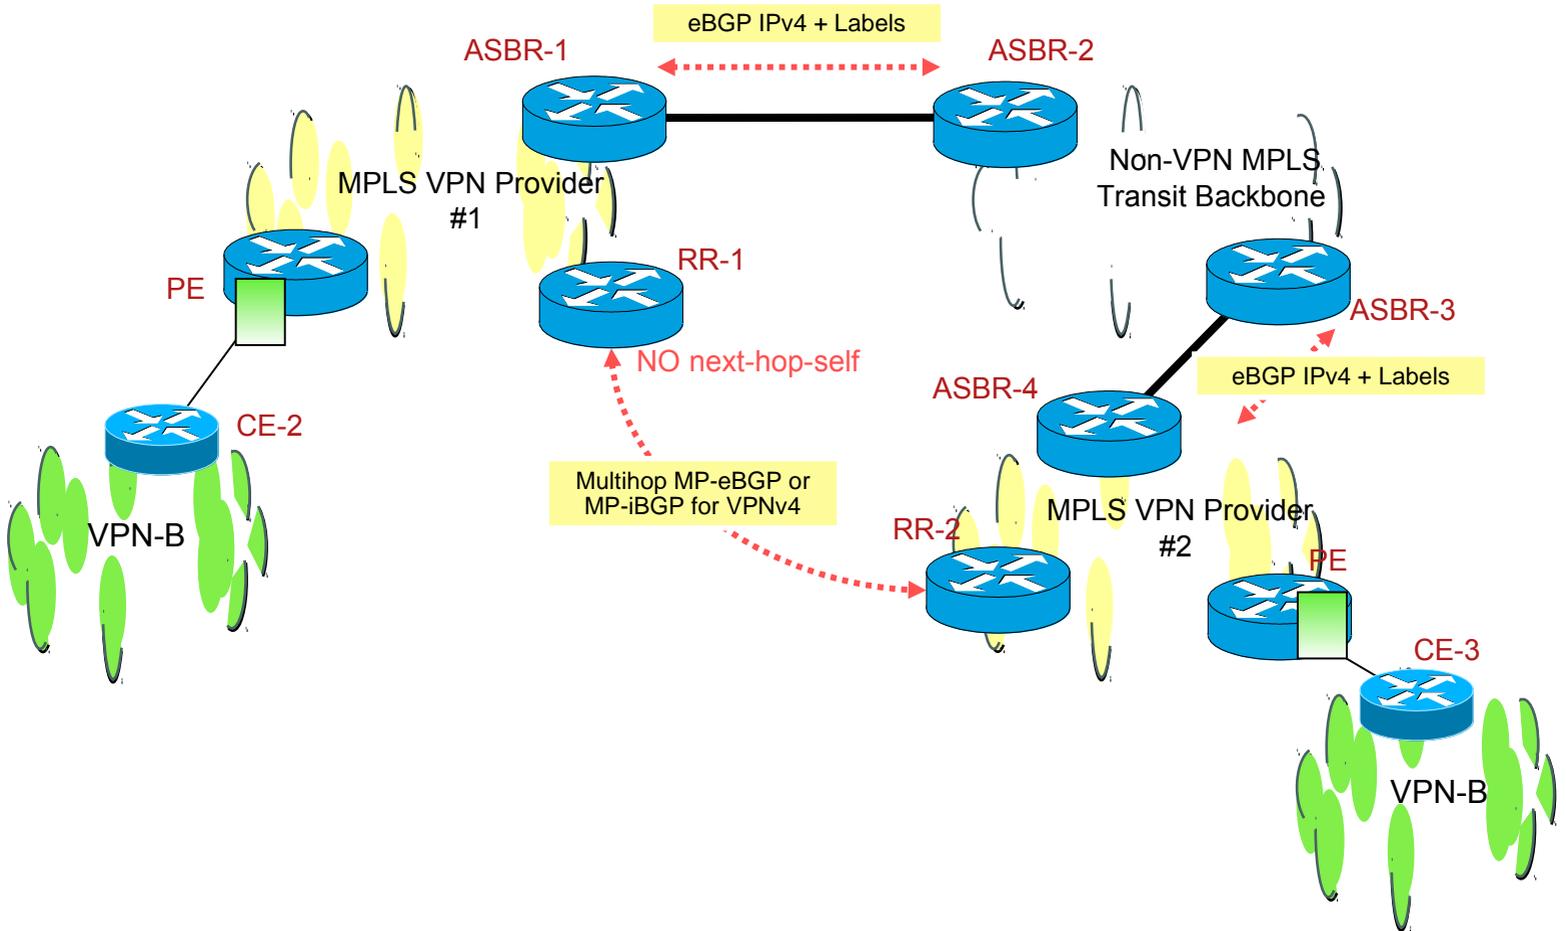
Multihop MP-eBGP for VPNv4

- Improves the scalability of route exchange
 - Eliminates the requirement to hold VPNv4 routes on the ASBRs;
 - Route reflectors already store VPNv4 prefix information
- Packets travel with 3 level label stack
 - <LDP IGP, BGP learnt label for Next-hop, VPN label>
- Advertising PE addresses to another AS may not be acceptable to few providers.

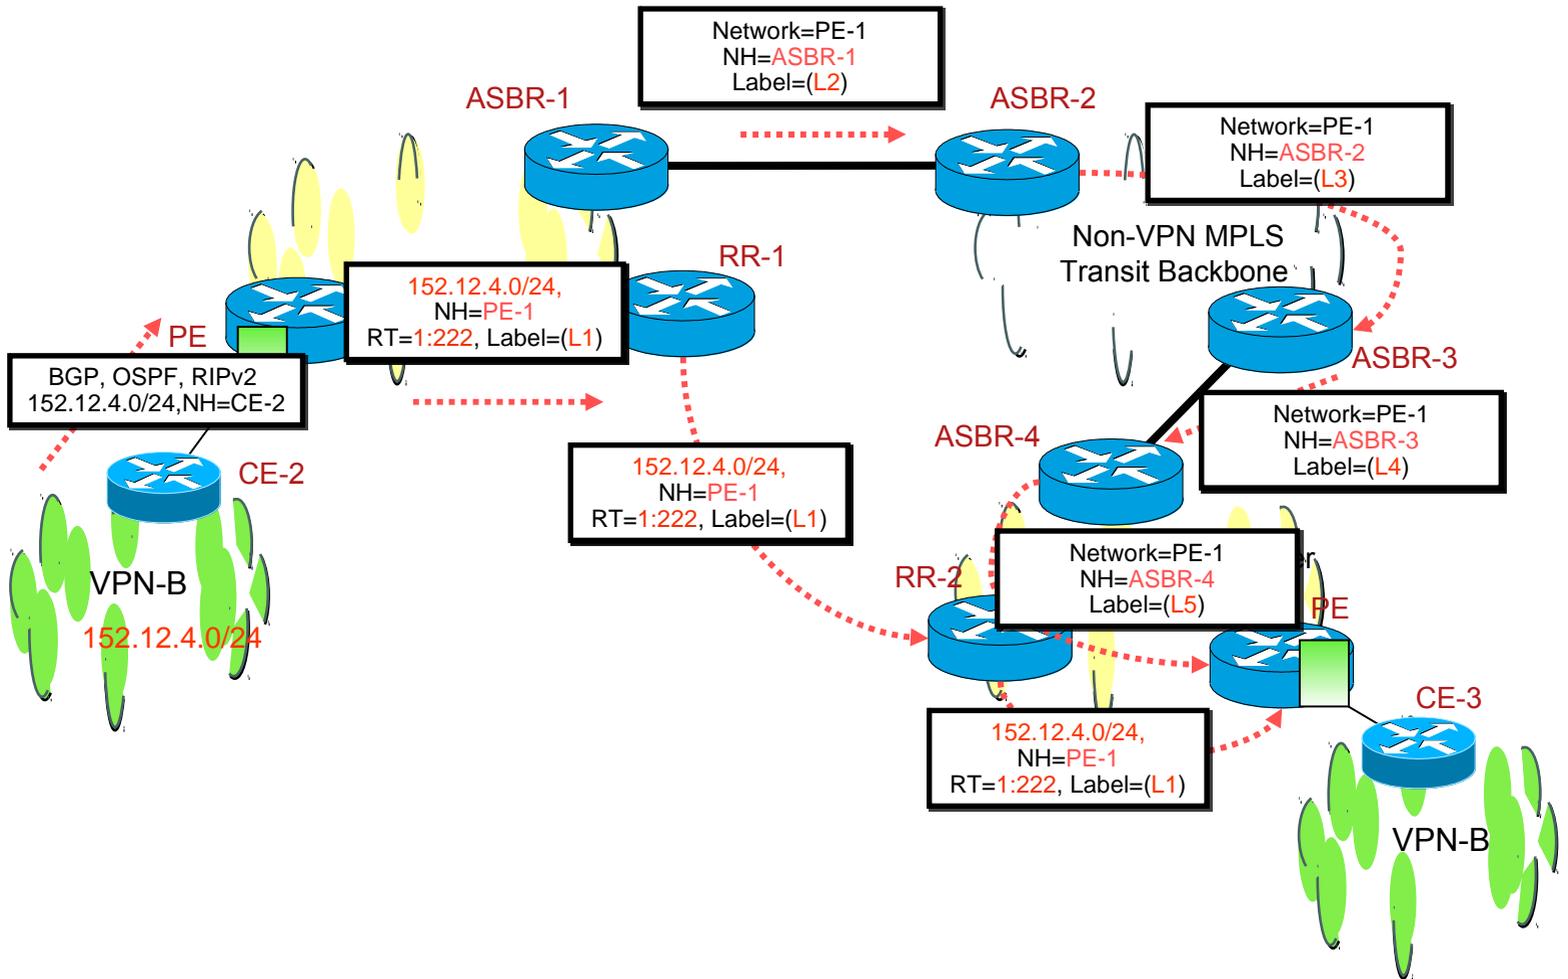
Non-VPN Transit Provider

- Two MPLS VPN providers may exchange routes with one or more third party
 - Which is a non-VPN transit backbone just running MPLS
- Multihop MP-eBGP deployed between edge providers
 - With the exchange of BGP next-hops via the transit provider;
 - BGP-4 + labels required
- Providers may use the same AS# within each region or different AS#
 - Transit network is NOT part of the AS path

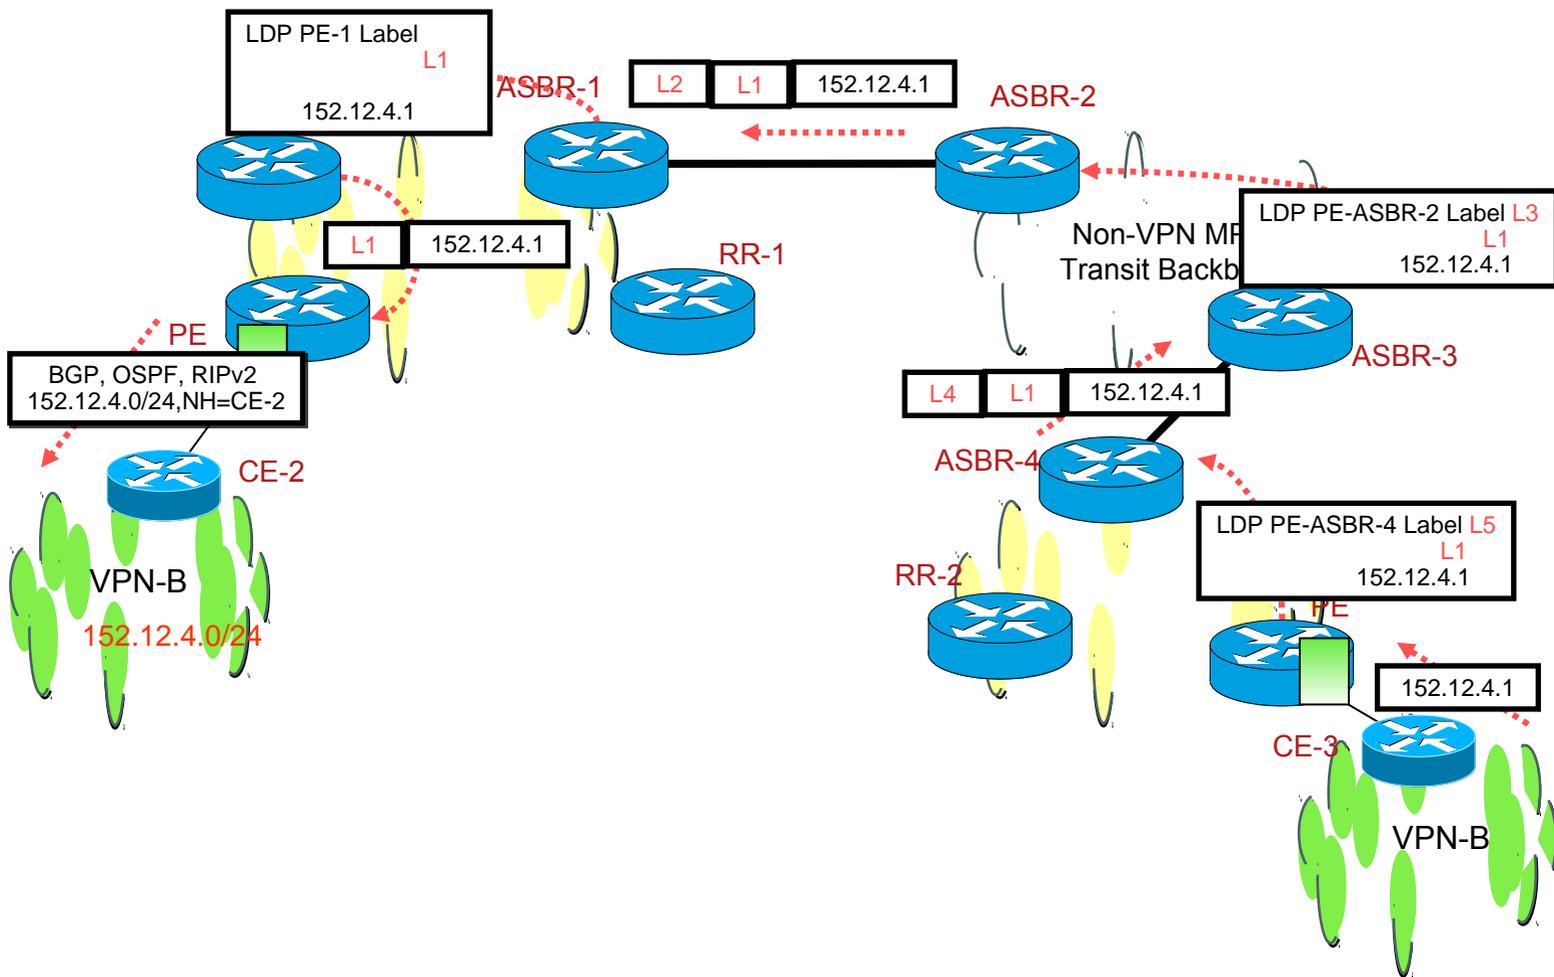
Non-VPN Transit Provider



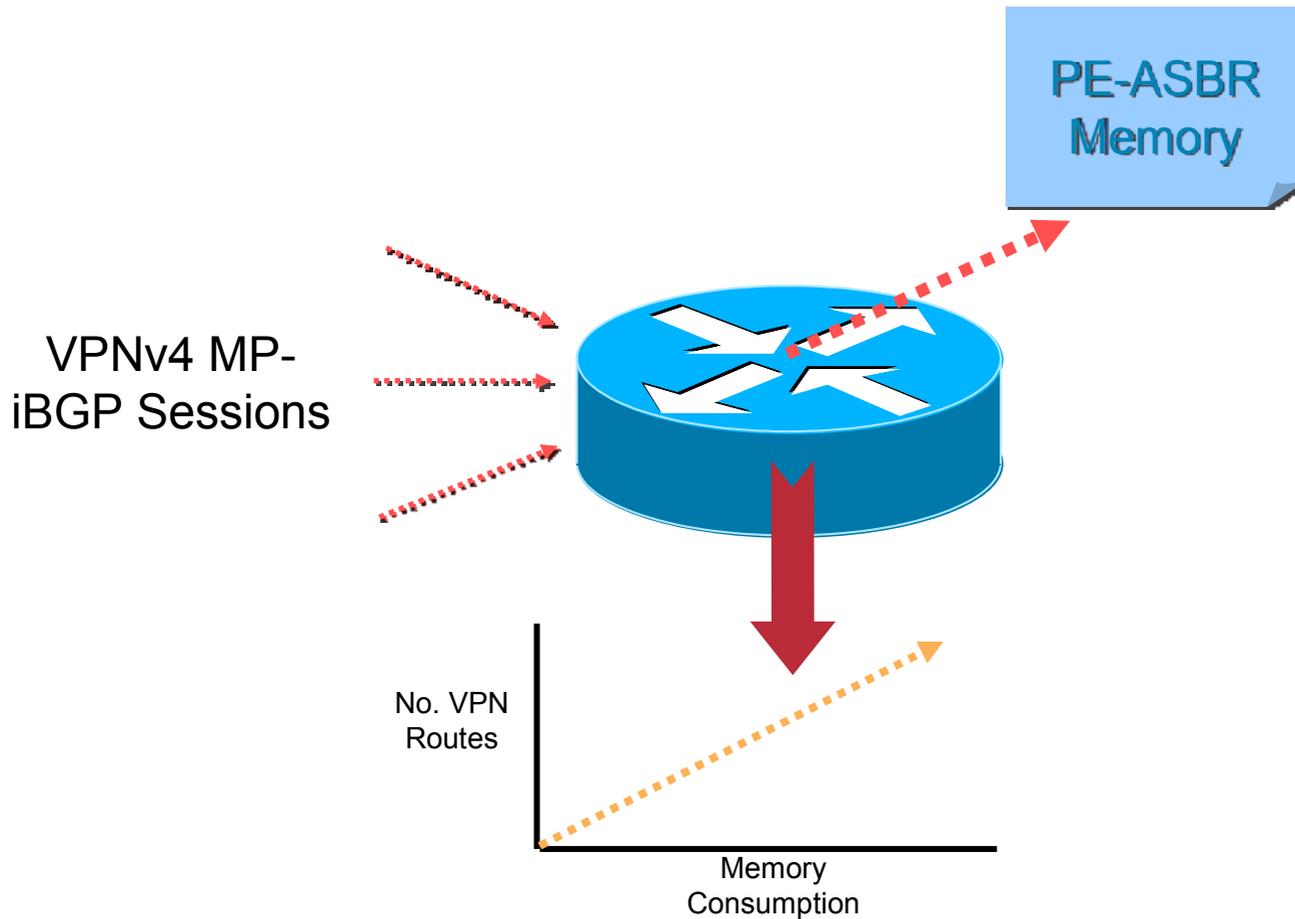
Non-VPN Transit Provider Control Plane



Non-VPN Transit Provider Forwarding Plane



Scaling Inter-Provider Solutions: PE-ASBR Memory Consumption



PE-ASBR Memory Scaling

- Potentially large amounts of VPN routing information
 - That may or may not need to be carried between providers;
 - Large percentage will be local VPN prefixes
 - This is specially true for (1)back-back vrf (2)MP-eBGP on PE-ASBR
- PE-ASBRs must hold relevant VPN routing information
 - But only Inter-AS VPN prefix details
- Two methods available to aid scaling
 - ARF with local VRF import (default)
 - ARF disabled with inbound filtering

ARF with local VRF import

- Automatic Route Filtering (ARF) for non-imported routes

If RT does not match locally configured import statement then drop the route

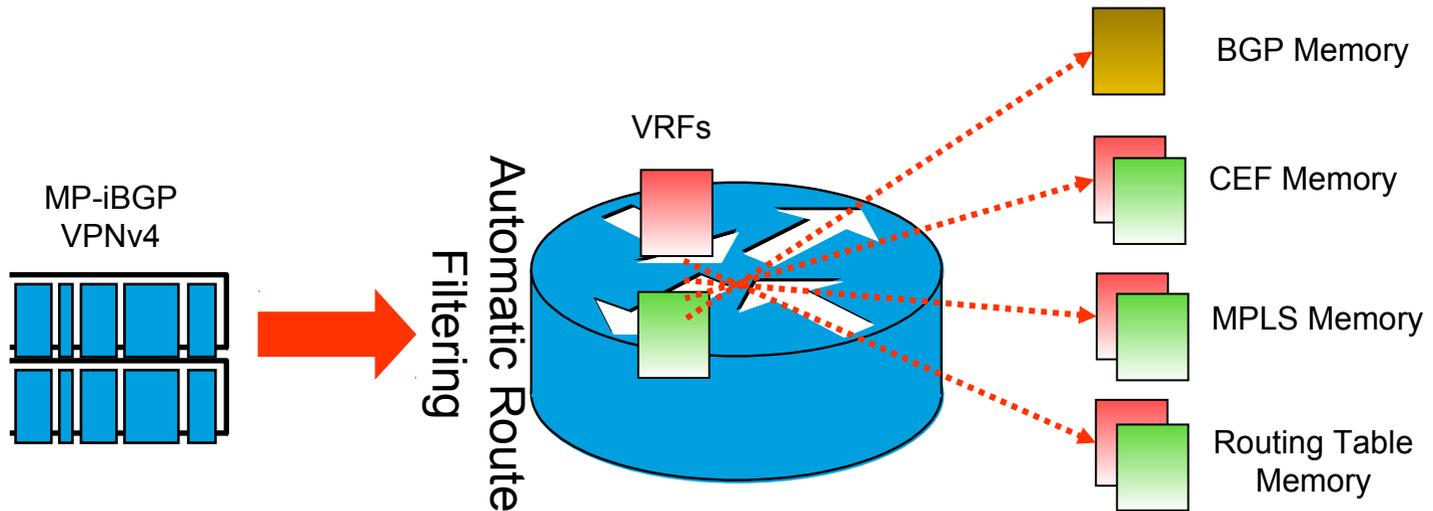
- Each PE-ASBR holds VRFs for Inter-AS VPNs

And imports routes based on route-target values

- PE-ASBR acts like normal PE router

Although also services external MP-BGP sessions

ARF with local VRF import



BGP, CEF, MPLS & RT Memory per-VRF

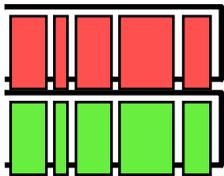
ARF disabled with inbound filtering

- Automatic Route Filtering (ARF) enabled by default
 - Therefore if no VRFs are configured then ALL VPN routes are dropped by the PE-ASBR
- Automatic Route Filtering may be disabled
 - Through use of **no default BGP route-target filter** command within the BGP configuration
- Disabling of ARF will cause ALL routes to be accepted by the PE-ASBR, when it has no VRFs
 - Which implies filtering must occur to drop unwanted routes

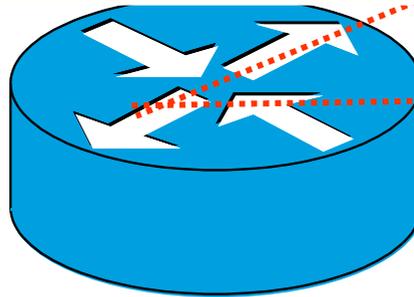
ARF disabled with inbound filtering

```
router bgp 1
!  
no bgp default route-target filter
!  
address-family vpnv4
neighbor 154.27.0.134 activate
neighbor 154.27.0.134 send-community extended
neighbor 154.27.0.134 route-map vpn-routes-filter in
```

MP-iBGP
VPNv4



NO Automatic Route
Filtering



BGP Memory

LFIB Memory

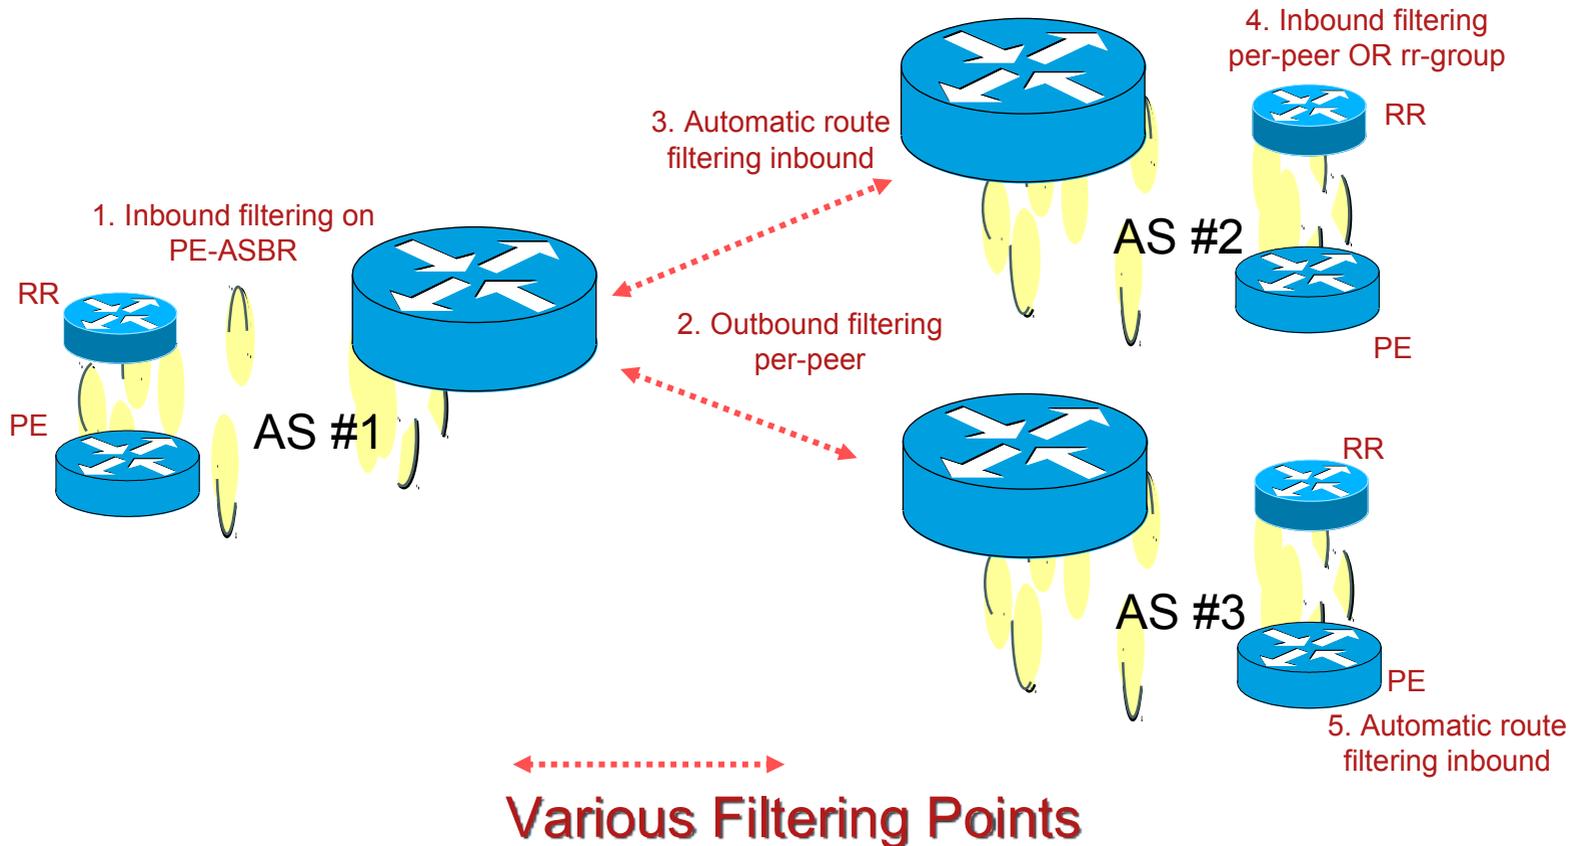
VRF & CEF memory
not required

Routing Table memory
not required

NO per-VRF CEF or RT Memory, only BGP & LFIB

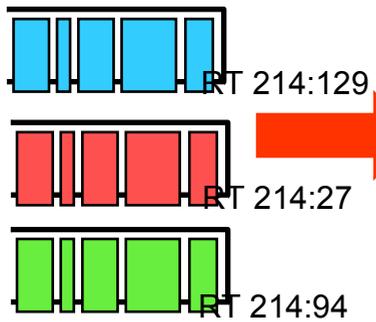
Filtering & Route Distribution Mechanisms

Inter-AS Filtering Points



Inbound filtering on PE-ASBR

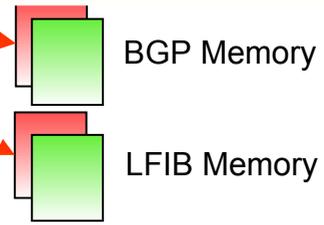
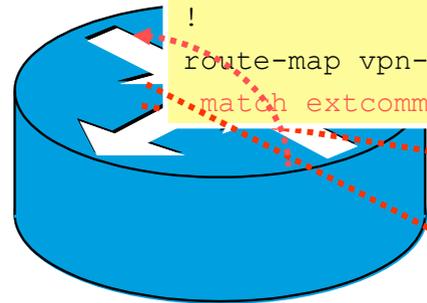
```
router bgp 1
!  
no bgp default route-target filter
!  
address-family vpnv4
  neighbor 154.27.0.134 activate
  neighbor 154.27.0.134 send-community extended
  neighbor 154.27.0.134 route-map vpn-routes-
  filter in
  !  
ip extcommunity-list 1 permit rt 214:27 rt
  214:94
  !  
route-map vpn-routes-filter permit 10
  match extcommunity 1
```



Blue VPN routes
discarded



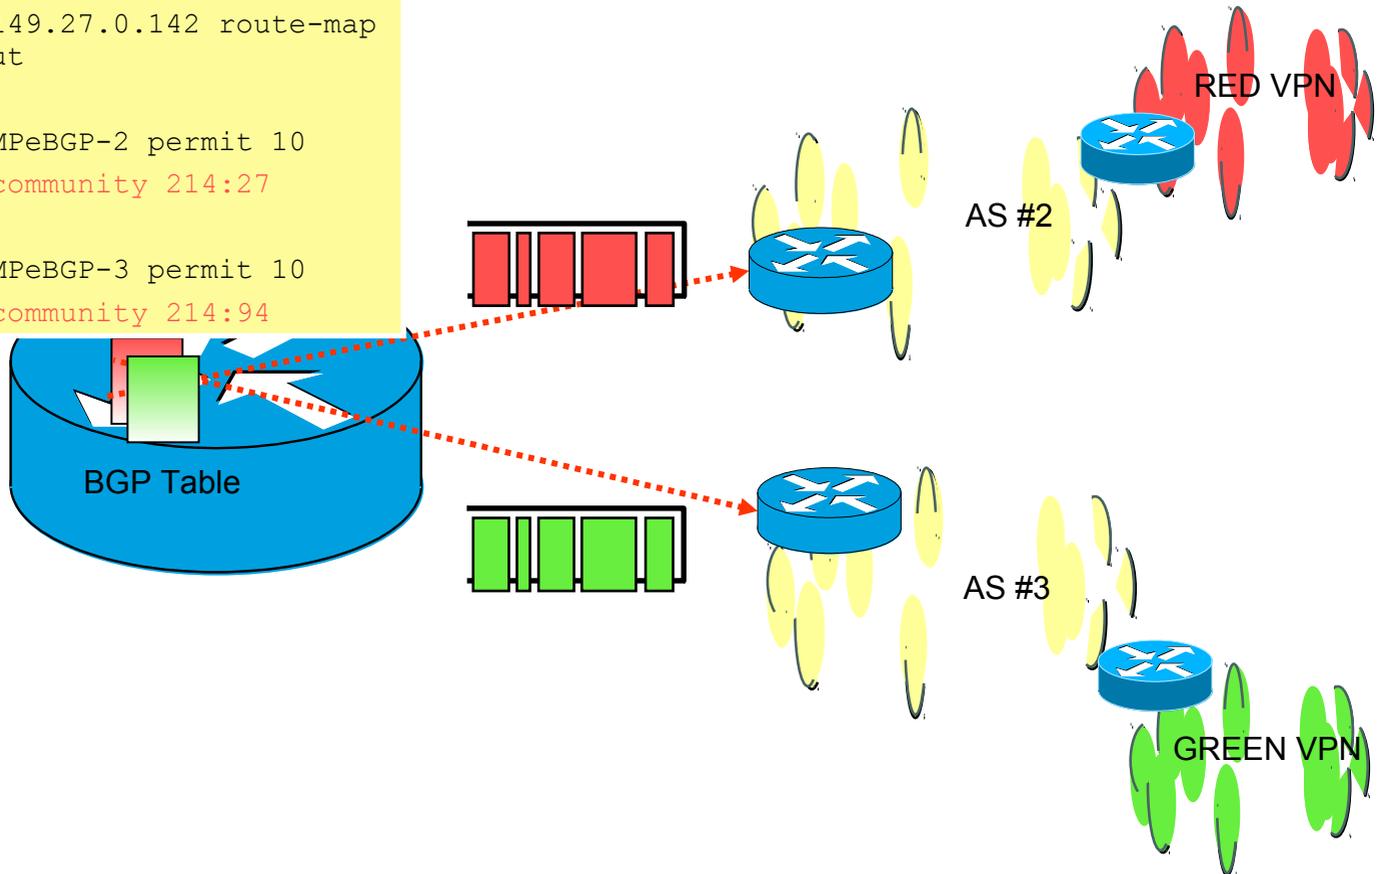
NO Automatic Route
Filtering



NO ARF – Filter inbound on per-peer basis

Outbound filtering on PE-ASBR

```
address-family vpv4
  neighbor 157.27.0.132 route-map
  MPeBGP-2 out
  neighbor 149.27.0.142 route-map
  MPeBGP-3 out
  !
  route-map MPeBGP-2 permit 10
  match extcommunity 214:27
  !
  route-map MPeBGP-3 permit 10
  match extcommunity 214:94
```



Downstream RT allocation

- Both inbound & outbound filtering restrictive with large number of VPN clients

As each RT must be known and the filters must be established

- Changes to VPN client membership will cause configuration changes on PE-ASBRs

As each filter must be updated to reflect addition/deletion of VPN clients

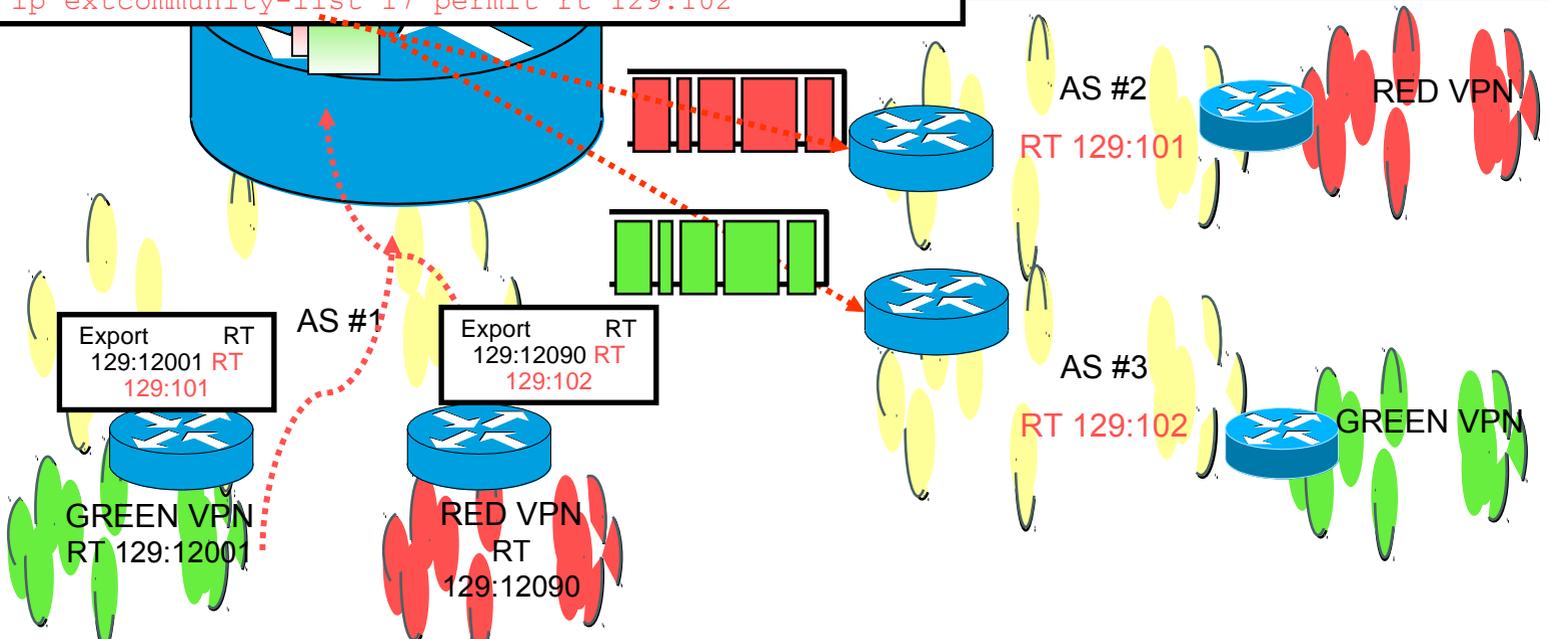
- With large number of clients a simplified filtering scheme is needed

Provided with “Downstream provider RT allocation” scheme

Downstream RT allocation

```
address-family vpnv4
 neighbor 154.27.0.134 activate
 neighbor 154.27.0.134 send-community extended
 neighbor 154.27.0.134 route-map asbr-routes-filter
 in
 neighbor 157.27.0.132 route-map MPeBGP-2 out
 neighbor 149.27.0.142 route-map MPeBGP-3 out
 !
 ip extcommunity-list 1 permit rt 129:101 rt 129:102
 ip extcommunity-list 16 permit rt 129:101
 ip extcommunity-list 17 permit rt 129:102
```

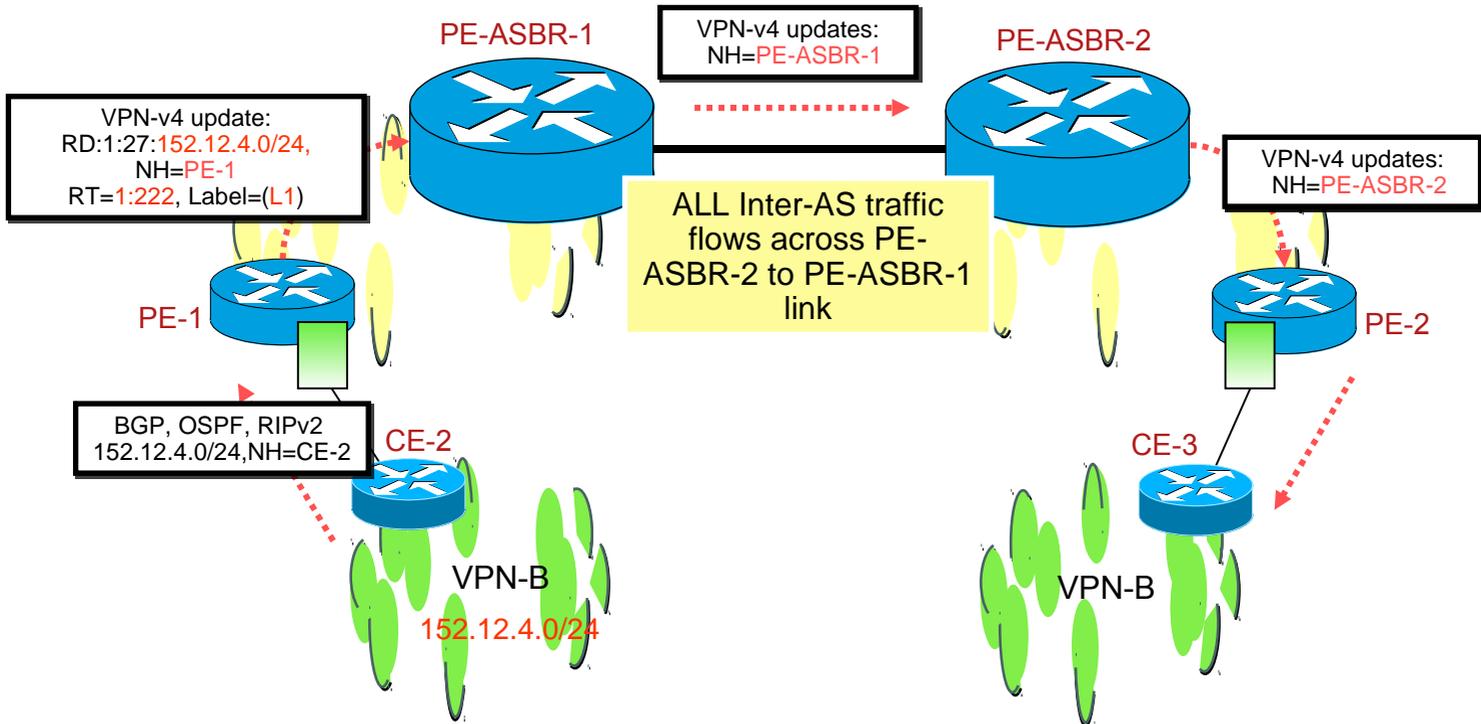
```
route-map asbr-routes-filter permit
 10
 match extcommunity 1
 !
route-map MPeBGP-2 permit 10
 match extcommunity 16
 !
route-map MPeBGP-3 permit 10
 match extcommunity 17
```



Distribution of Traffic Load between Providers

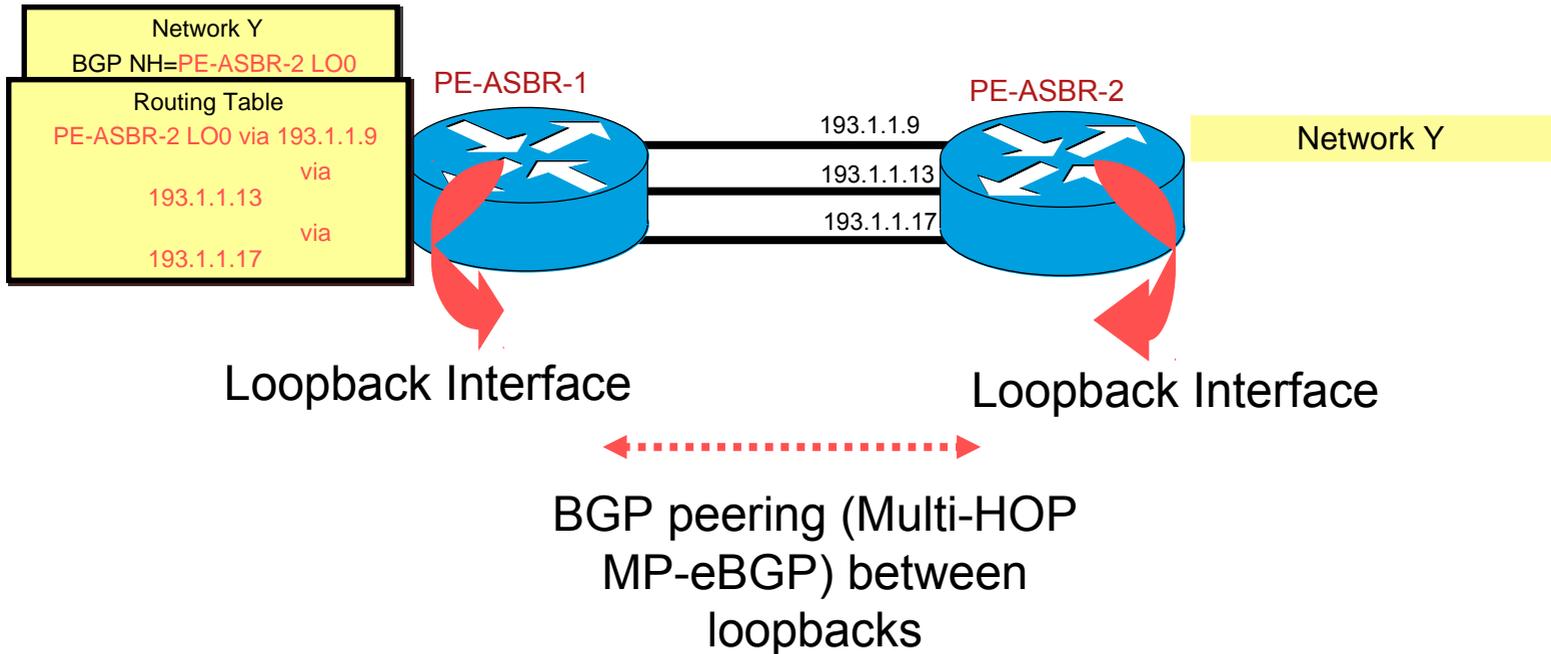
- Balancing of Inter-AS traffic is an important issue
 - For distribution of traffic and redundancy of network design
- All Inter-AS traffic must pass through PE-ASBRs
 - As BGP next-hops are reachable via these routers
- Multiple links provide traffic distribution
 - But do not provide redundancy due to single point of failure of the PE-ASBR

VPN Client Traffic Flow



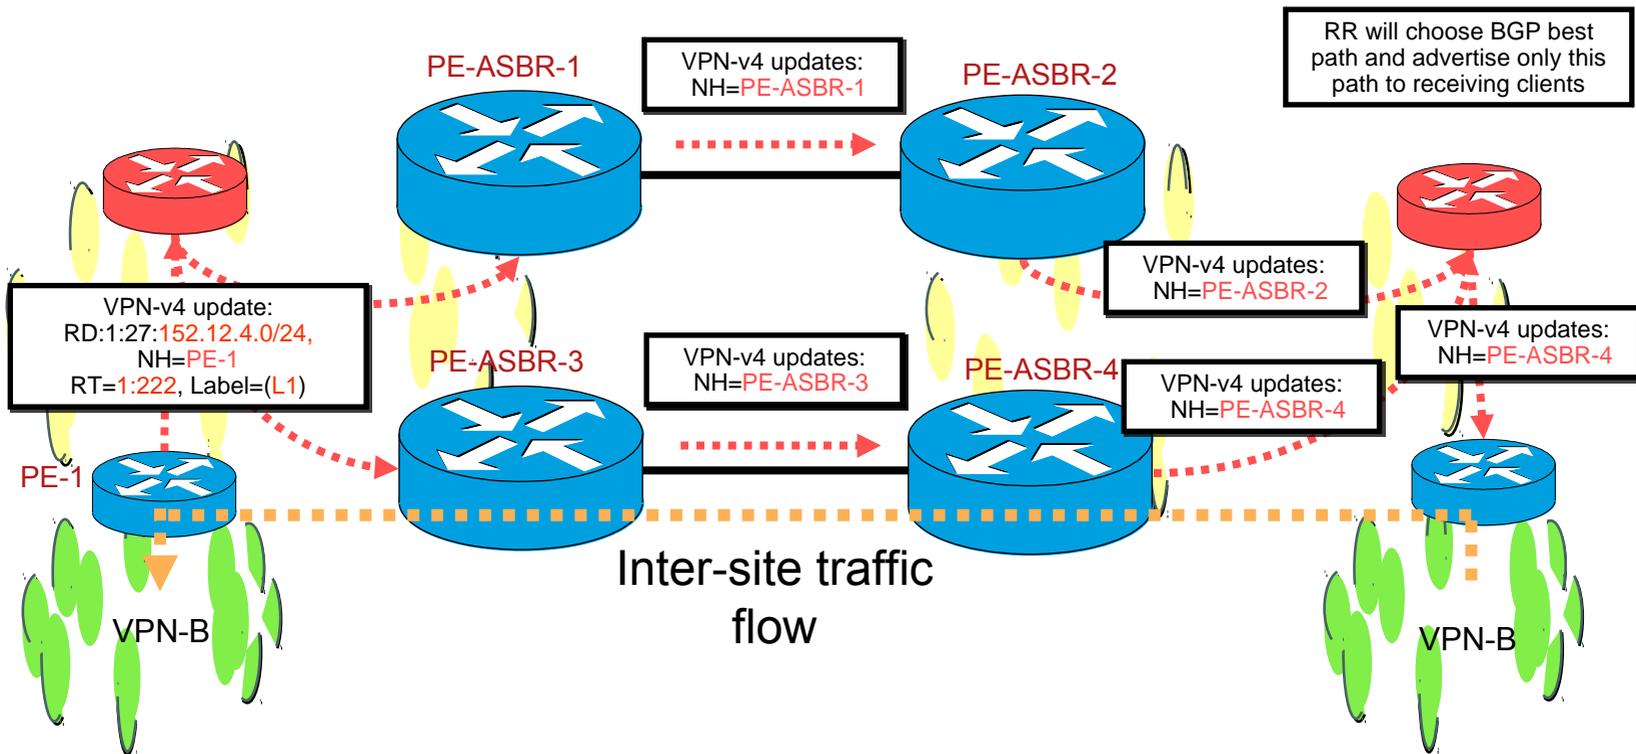
VPN Client to VPN Client traffic flow via Inter-AS Link

Load Balancing between PE-ASBRs



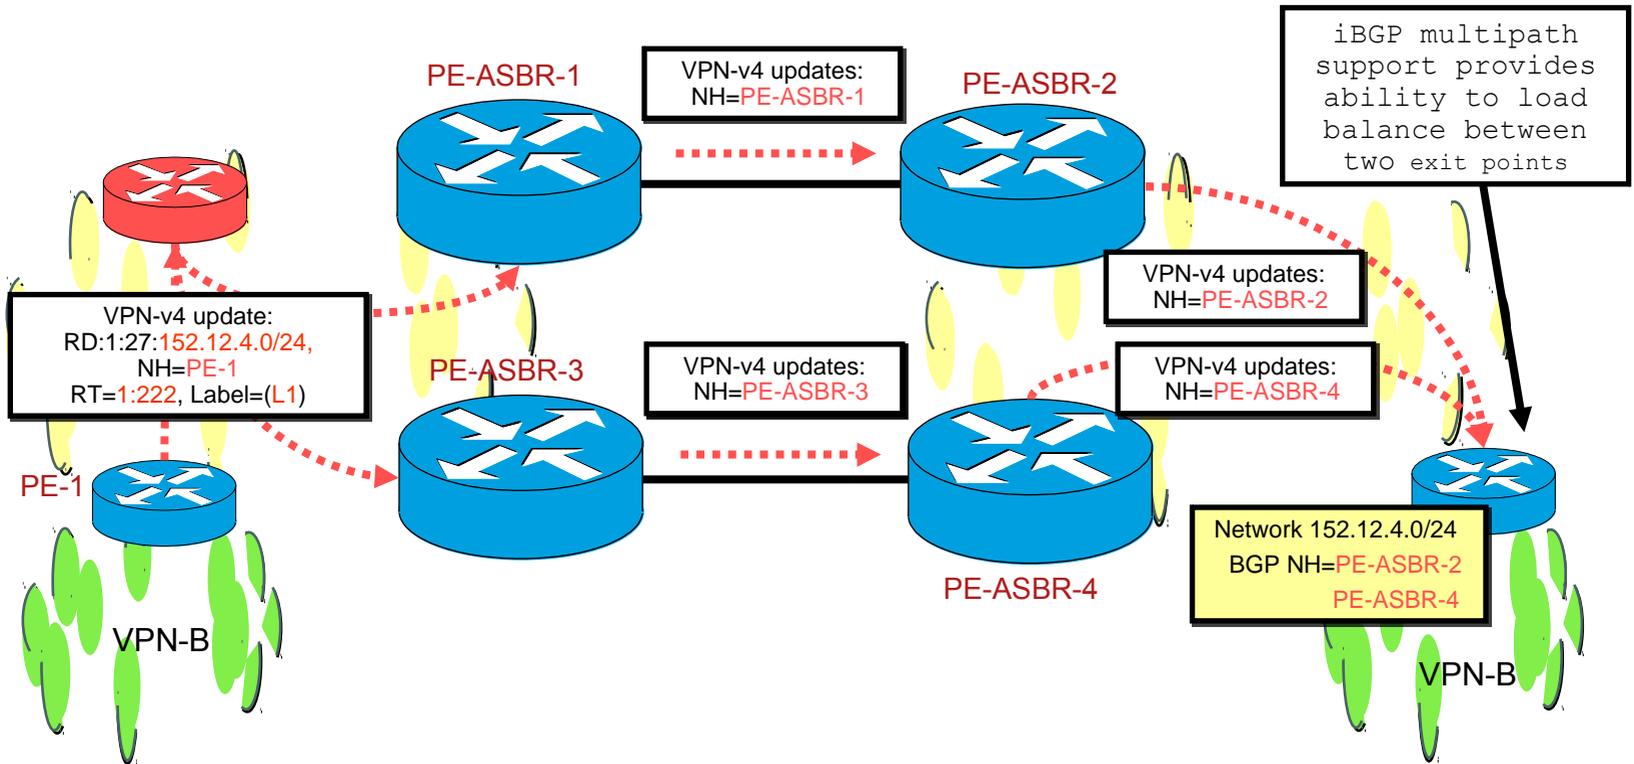
Load Balancing across multiple PE-ASBR links

Redundant PE-ASBR Connections



Redundant PE-ASBR used purely for backup

Redundant PE-ASBR Load Balancing



Load balancing PE-ASBR links without Route Reflectors



Internet Access from a VPN

Overview

Leaking Between VPN and Global Backbone Routing

Separating Internet Access from VPN Service

Internet Access Backbone as a Separate VPN

Internet Access with VRF Aware NAT



Leaking Between VPN and Global Backbone Routing

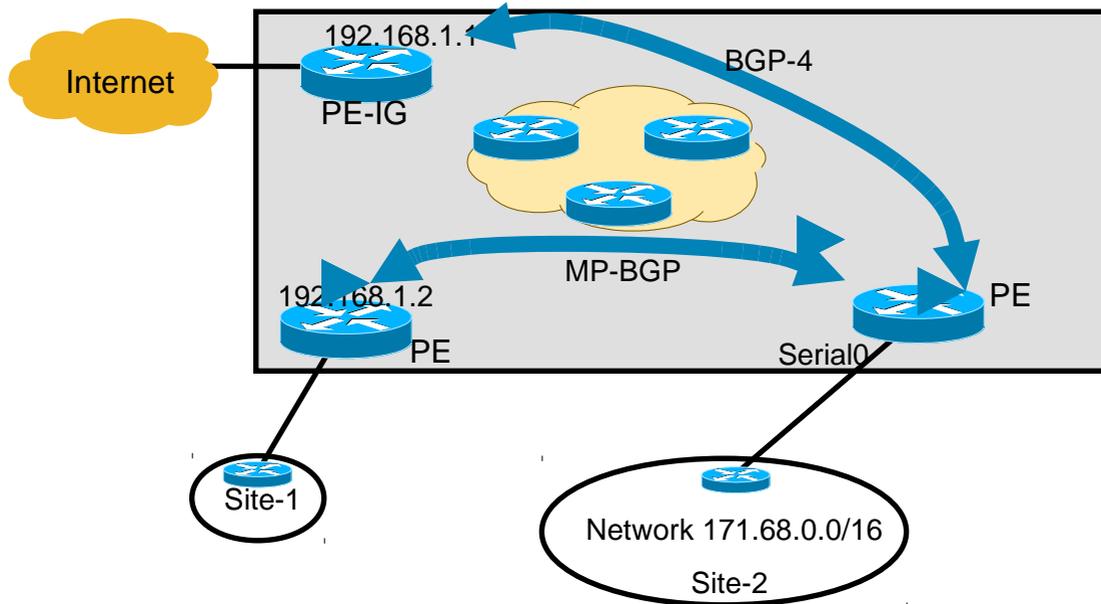
Internet Access Through Global Routing

- Two implementation options:

Internet access is implemented via separate interfaces that are not placed in any VRF (traditional Internet access setup)

Packet leaking between a VRF and the global table is achieved through special configuration commands

Underlying Technology



- Packet leaking between a VRF and a global routing table is based on two IOS features:

A VRF static route can be defined with a global next-hop. This feature achieves leaking from a VRF toward a global next-hop

A global static route can be defined pointing to a connected interface that belongs to a VRF. This feature achieves leaking from a global routing table into VPN space.

Configuring Packet Leaking

Router(config)#

```
ip route vrf name prefix mask next-hop global
```

- Configures a VRF static route with a global next-hop
- Packets matched by this static route are forwarded toward a global next-hop and thus leak into global address space

Router(config)#

```
ip route prefix mask interface
```

- Configures a global static route that can point to an interface in VRF
- Globally-routed packets following this entry will be sent toward a CE router (into a VPN)

Designing Internet Access Through Packet Leaking

- A public address is assigned to an Internet/VPN customer
- A global static route for an assigned address block is configured on the PE router

The static route has to be redistributed into BGP to provide full connectivity to the customer

- A default route toward a global Internet exit point is installed in the customer VRF

This default route is used to forward packets to unknown destinations (Internet) into the global address space

Connectivity from the Customer to the Internet

- A default route is installed into the VRF pointing to a global Internet gateway

Warning: Using a default route for Internet routing does NOT allow any other default route for intra-VPN routing

- The default route is not part of any VPN

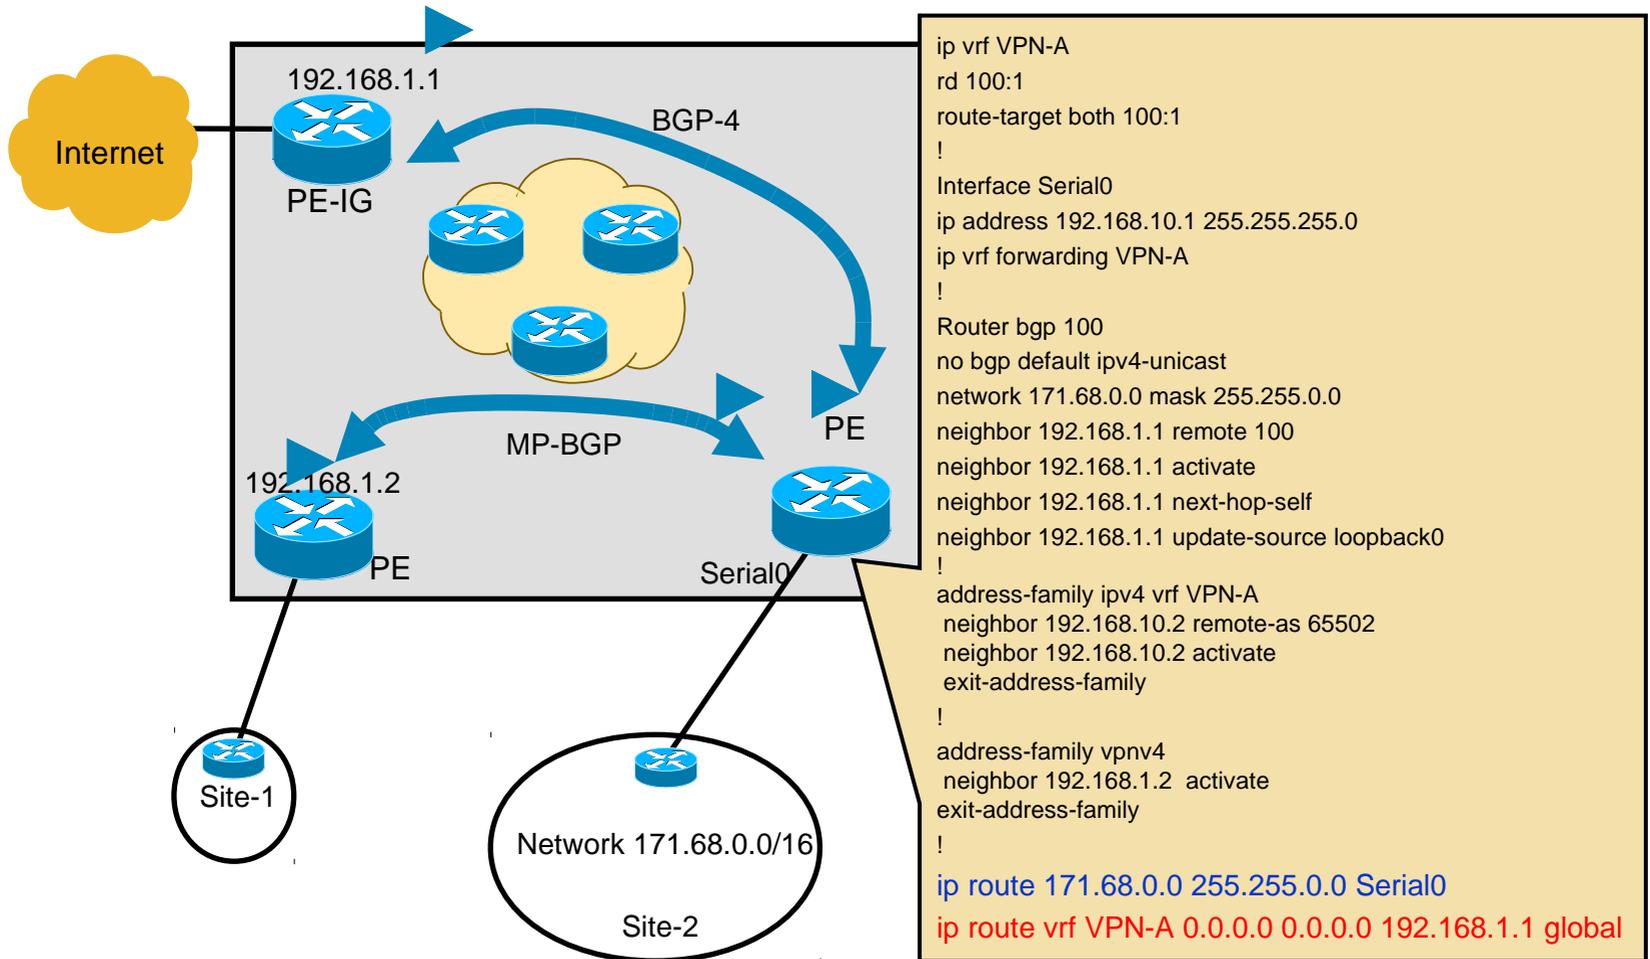
A single label is used for packets forwarded toward the global next-hop

The label used for packet forwarding is the IGP label (TDP/LDP-assigned label) corresponding to the IP address of the global next-hop

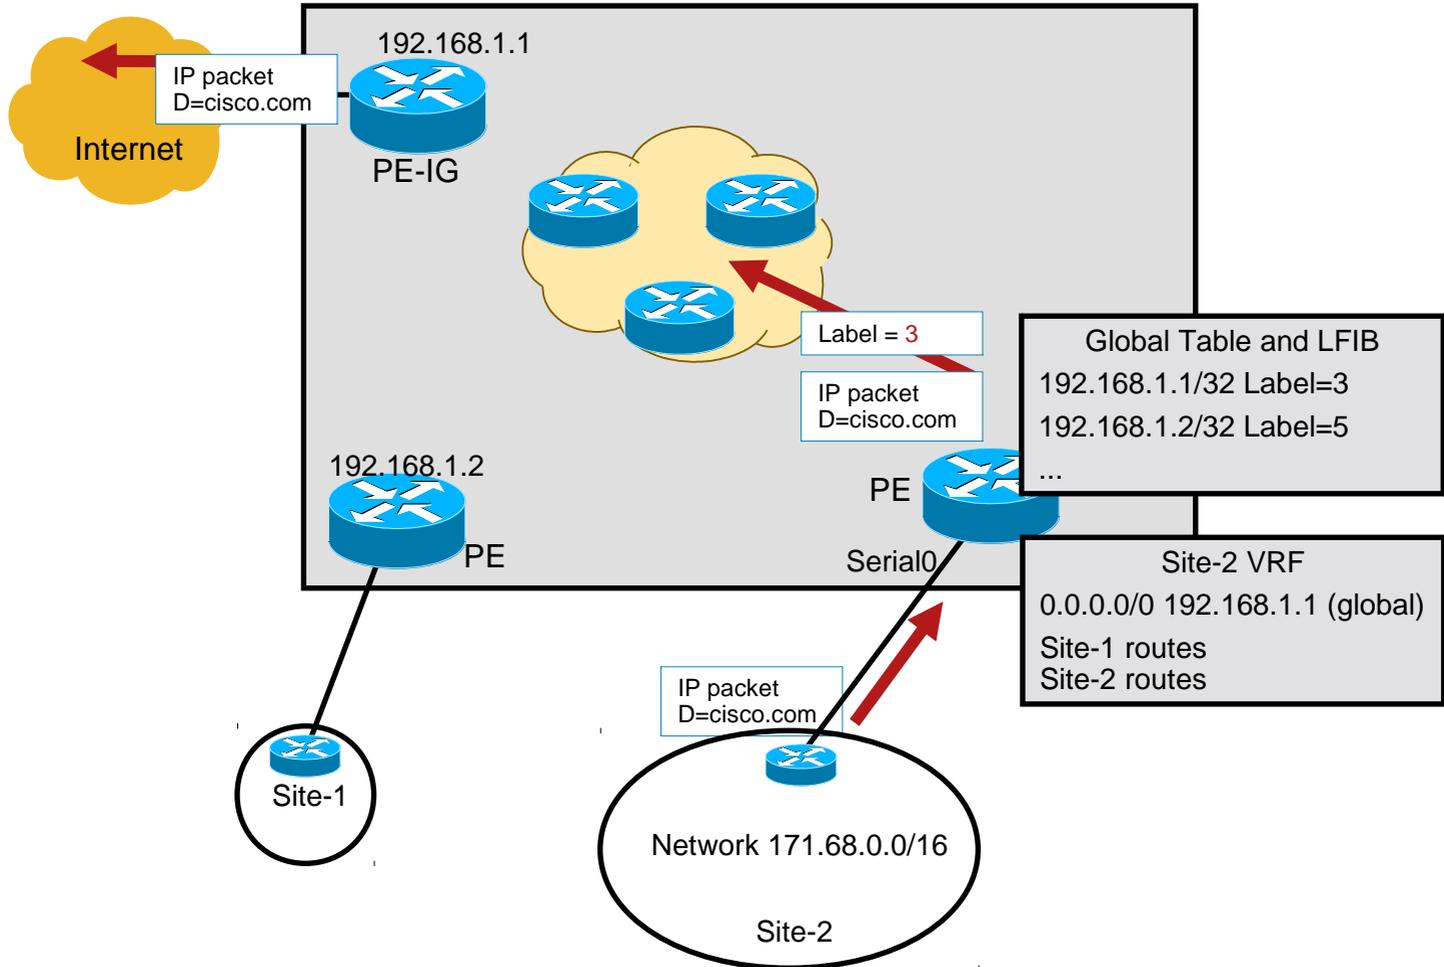
VRF-Specific Default Route

- The Internet gateway specified as the next-hop in the VRF default route need NOT to be directly connected
- The next-hop can be in the upstream AS to achieve redundancy
- Different Internet gateways can be used for different VRFs

An Example of Internet Access Through Packet Leaking



Packet Leaking in Action



Redundant Internet Access with Packet Leaking

- Several VRF default routes can be used with different next-hops
This setup will survive failure of the Internet gateway, not the failure of its upstream link
 - Global next-hop can be in an upstream autonomous system
This setup yields best redundancy because it tests availability of the whole path from PE router to the upstream autonomous system
- Drawback:** local Internet service stops working if the upstream autonomous system is not reachable

Limitations of Packet

- Drawbacks:

Internet and VPN packets are mixed on the same link; security issues arise

Packets moving toward temporarily unreachable VPN destinations might leak into the Internet

A global BGP session between a PE and a CE router needed for full Internet routing exchange is hard to configure

- Benefits:

A PE router does not need Internet routes, only an IGP route toward the Internet gateway



Separating Internet Access from VPN Service

Designing Internet Access Separated from VPN

- Customer Internet access is implemented over different interfaces than VPN access is:

Traditional Internet access implementation model

Requires separate physical links or separate subinterfaces

Maximum design flexibility; Internet access is totally independent from MPLS VPN

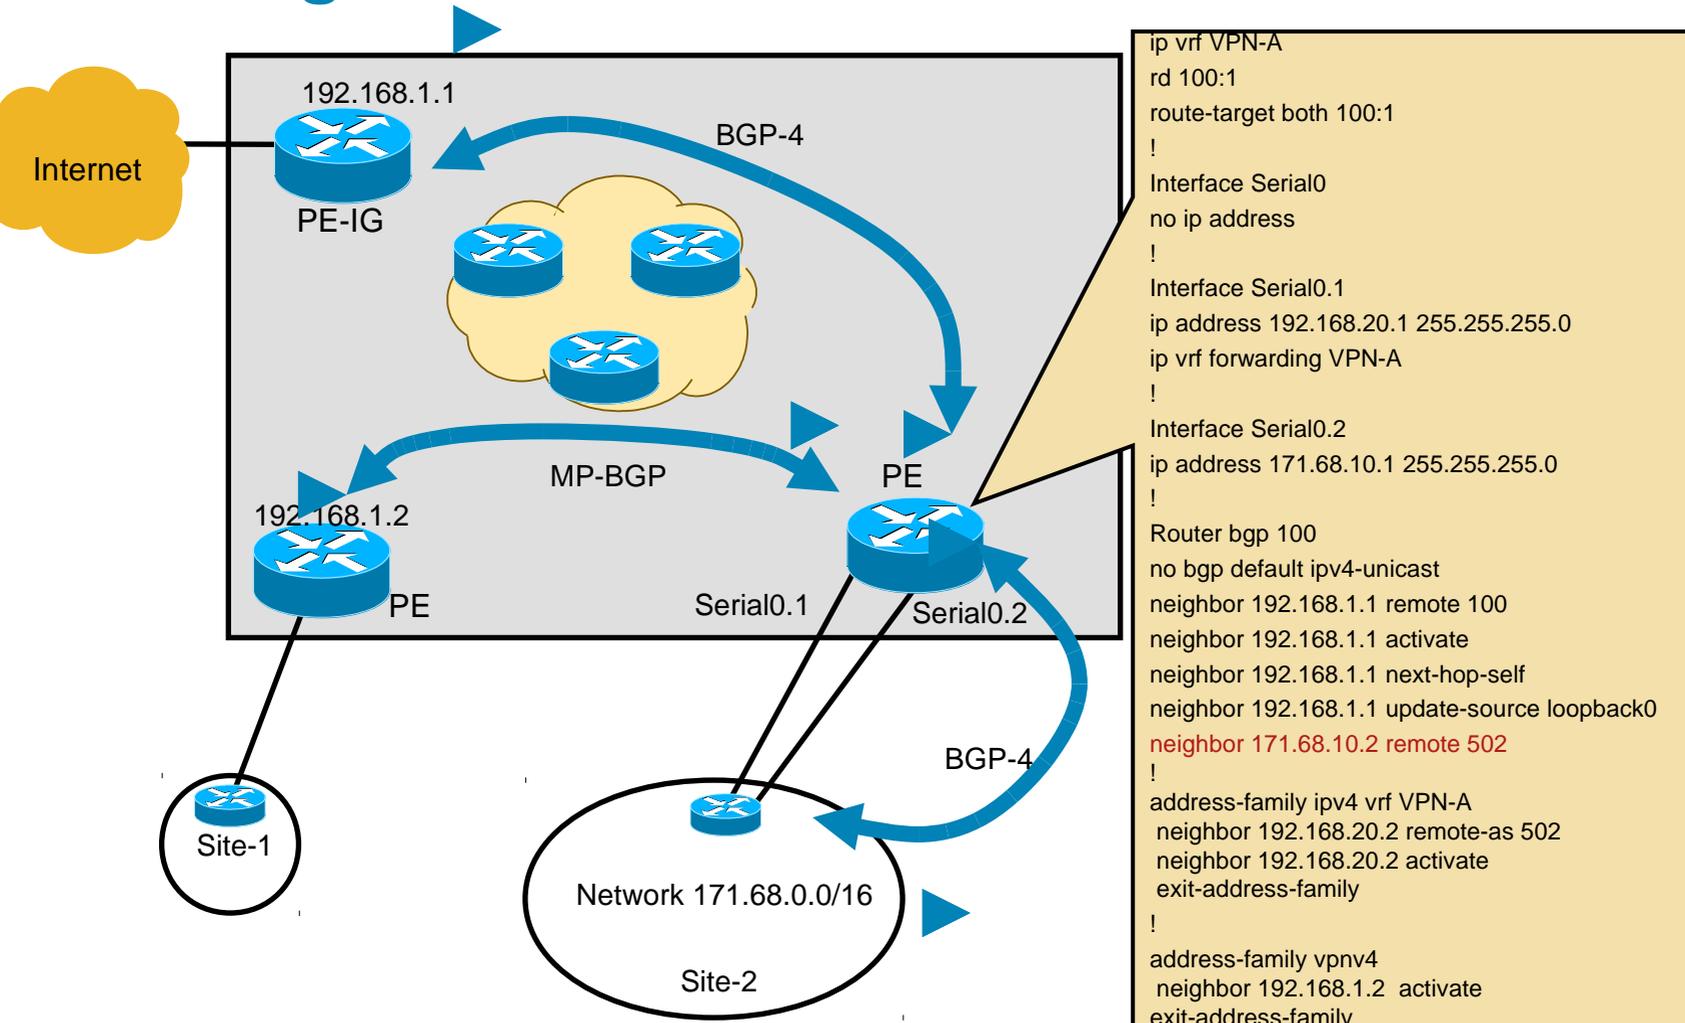
Subinterfaces

- Separate physical links for VPN and Internet traffic are sometimes not acceptable because of high cost
- Subinterfaces can be used over WAN links using Frame Relay or ATM encapsulation (including DSL)
- A tunnel interface could be used; however:

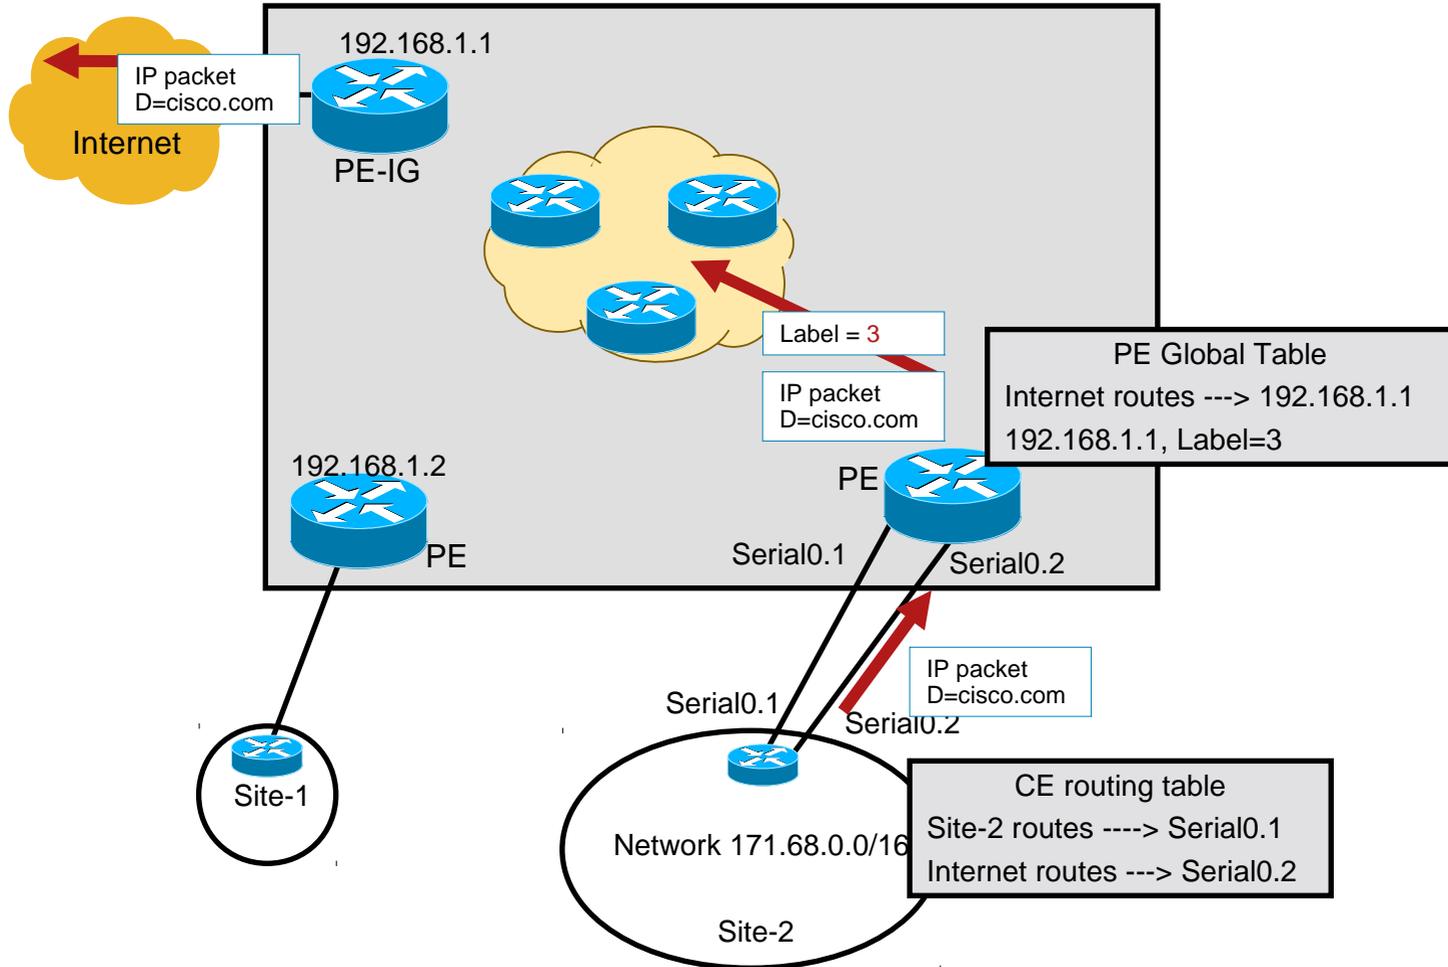
Tunnels are not VRF-aware: VPN traffic must run over a global tunnel

This setup could lead to security leaks because global packets could end up in VPN space

An Example of Internet Access Through a Dedicated Subinterface



Internet Access Through a Dedicated Subinterface - Traffic Flow



Limitations of Separate Internet Access

- Drawbacks:

Requires separate physical link or specific WAN encapsulation

PE routers must be able to perform Internet routing (and potentially carry full Internet routing)

Wholesale Internet access or Central Firewall service cannot be implemented with this model

PE router has internet as well as VPN routes. A lot of ISPs do not like this idea due to security reasons

- Benefits:

Well-known model

Supports all customer requirements

Allows all Internet services implementation, including a BGP session with the customer



Internet Access Backbone as a Separate VPN

Internet Access As a Separate VPN

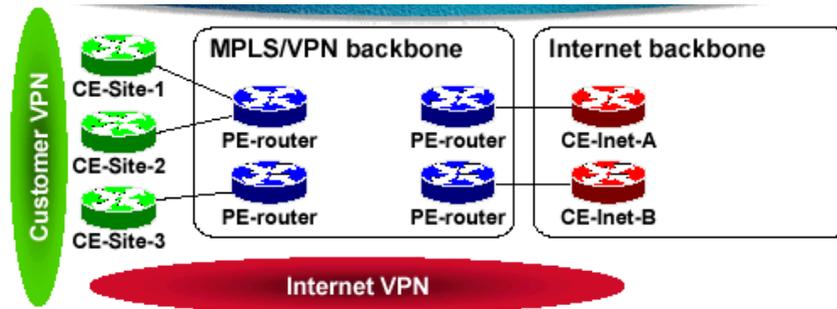
- This design realizes Internet access by using MPLS VPN features:

An Internet gateway is connected as a CE router to the MPLS VPN backbone

An Internet gateway shall not insert full Internet routing into the VPN; only the default route and the local (regional) routes can be inserted

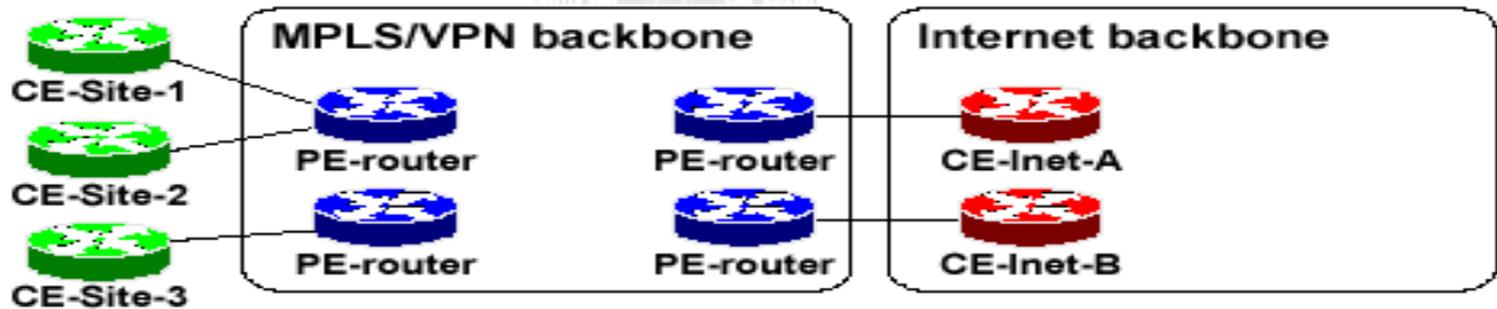
Every customer that needs Internet access is assigned to the same VPN as the Internet gateway

Internet Access as a Separate VPN



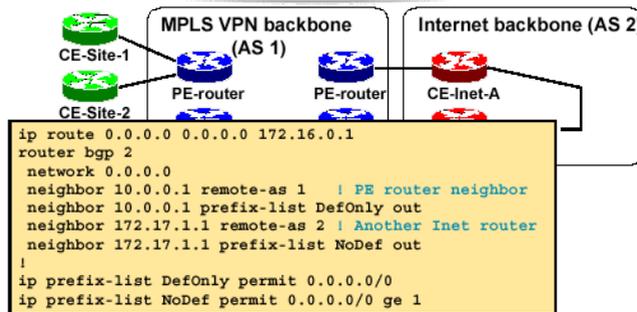
- The Internet backbone is separate from the VPN backbone
- VPN customers are connected to the Internet through a proper VPN/VRF setup

Redundant Internet Access



- Multiple CE-Internet routers can be used for redundancy
All CE-Internet routers advertise default route
Internet VPN will recover from CE-Internet router failure
Preferred default route can be indicated via MED attribute
- Default route should be advertised conditionally to achieve higher resilience

Redundant Internet Access



- Example: CE-Inet-A should advertise default route only if it can reach network 172.16.0.0/16 (upstream ISP core)

Limitations of Running an Internet Backbone in a VPN

- Drawbacks:

- Full Internet routing cannot be carried in the VPN; default routes are needed that can lead to suboptimal routing

- Internet backbones act as CE routers to the VPN backbone; implementing overlapping Internet + VPN backbones is tricky

- Benefits:

- Supports all Internet access service types

- Can support all customer requirements, including a BGP session with the customer, accomplished through advanced BGP setup



Internet Access using VRF Aware NAT

Internet Access using VRF-aware NAT

- If the VPN customers need Internet access without internet routes, then VRF-aware NAT can be used at the Internet-GW i.e. ASBR
- The Internet GW doesn't need to have internet routes either
- **Overlapping VPN addresses is not a problem**

Internet Access using VRF-aware NAT

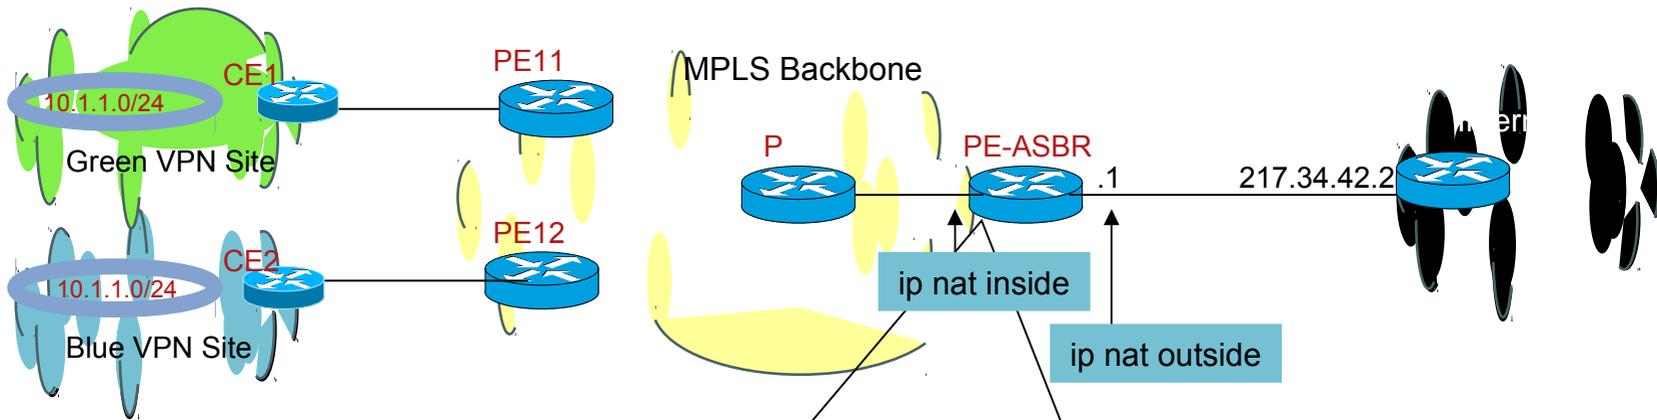
- VPN customers could be using 'overlapping' IP address i.e. 10.0.0.0/8
- Such VPN customers **must NAT** their traffic before using either "extranet" or "internet" or any shared* services
- **PE is capable of NATting the VPN packets** (eliminating the need for an extra NAT device)

* VoIP, Hosted Content, Management etc/

Internet Access using VRF-aware NAT

- Typically, inside interface(s) connect to private address space and outside interface connect to global address space
NAT occurs after routing for traffic from inside-to-outside interfaces
NAT occurs before routing for traffic from outside-to-inside interfaces
- **Each NAT entry is associated with the VRF**
- Works on VPN packets in the following switch paths : IP->IP, IP->MPLS and MPLS->IP

Internet Access using VRF-aware NAT



```

ip vrf green
rd 3000:111
route-target both 3000:1
ip vrf blue
rd 3000:222
route-target both 3000:2

router bgp 3000
address-family ipv4 vrf green
network 0.0.0.0
address-family ipv4 vrf blue
network 0.0.0.0
    
```

VRF specific Config

```

ip nat pool pool-green 24.1.1.0 24.1.1.254 prefix-length 24

ip nat pool pool-blue 25.1.1.0 25.1.1.254 prefix-length 24

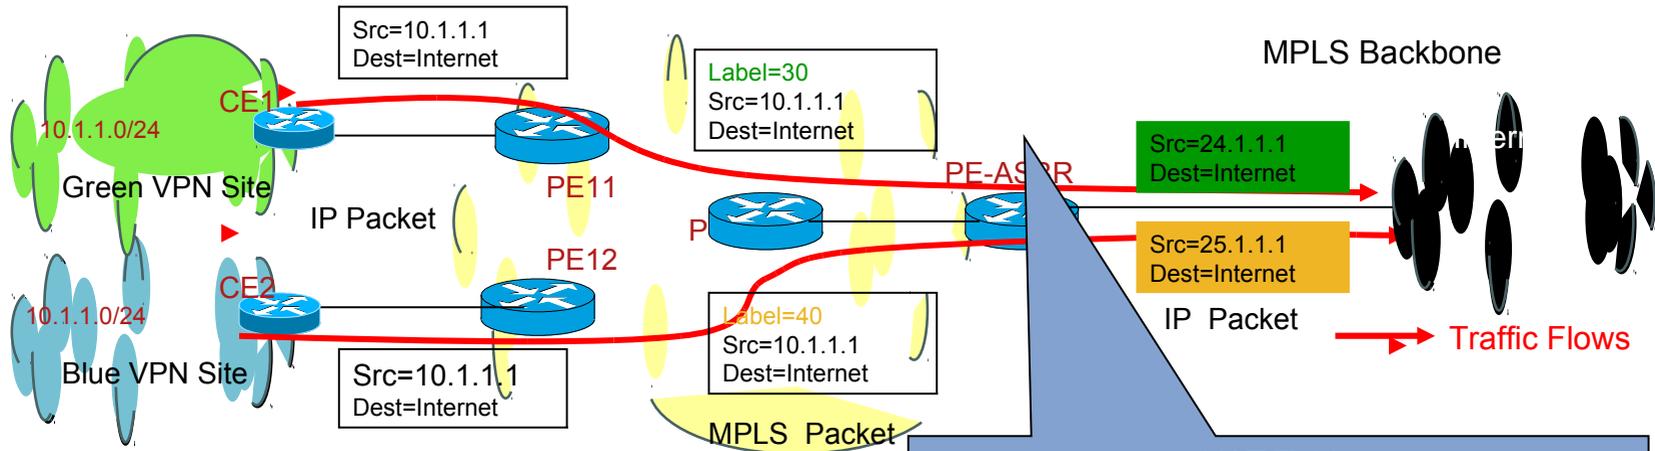
ip nat inside source list vpn-to-nat pool pool-green vrf green
ip nat inside source list vpn-to-nat pool pool-blue vrf blue

ip access-list standard vpn-to-nat
permit 10.1.1.0 0.0.0.255

ip route vrf green 0.0.0.0 0.0.0.0 217.34.42.2 global
ip route vrf blue 0.0.0.0 0.0.0.0 217.34.42.2 global
    
```

VRF-aware NAT Specific Config

Internet Access using VRF-aware NAT



- PE-ASBR removes the label from the received MPLS packets per LFIB
- Performs NAT on the resulting IP packets
- Forwards the packet

<u>VRF IP Source</u>	<u>Global IP</u>	<u>VRF-table-id</u>
10.1.1.1	24.1.1.1	green
10.1.1.1	25.1.1.1	blue

