Inter-Autonomous System MPLS

This document details the 3 primary methods to using inter-AS MPLS commonly referred as options A, B and C.

To illustrate the concepts, the sample network shown below is used. Customer A has sites connected to ISP 1 and ISP 2, in differently configured VRFs. Each Provider runs IS-IS as an IGP and uses LDP to assign labels. Route reflectors exchange VPNv4 and external routes.

The examples given in this document will follow the prefix and label advertisements and traffic flow involved when PC1 sends a packet to PC2.

All output shown and referenced is based on Cisco IOS. The local PC is represented by loopbacks on the CE routers.

A brief description of each option is given on the next page

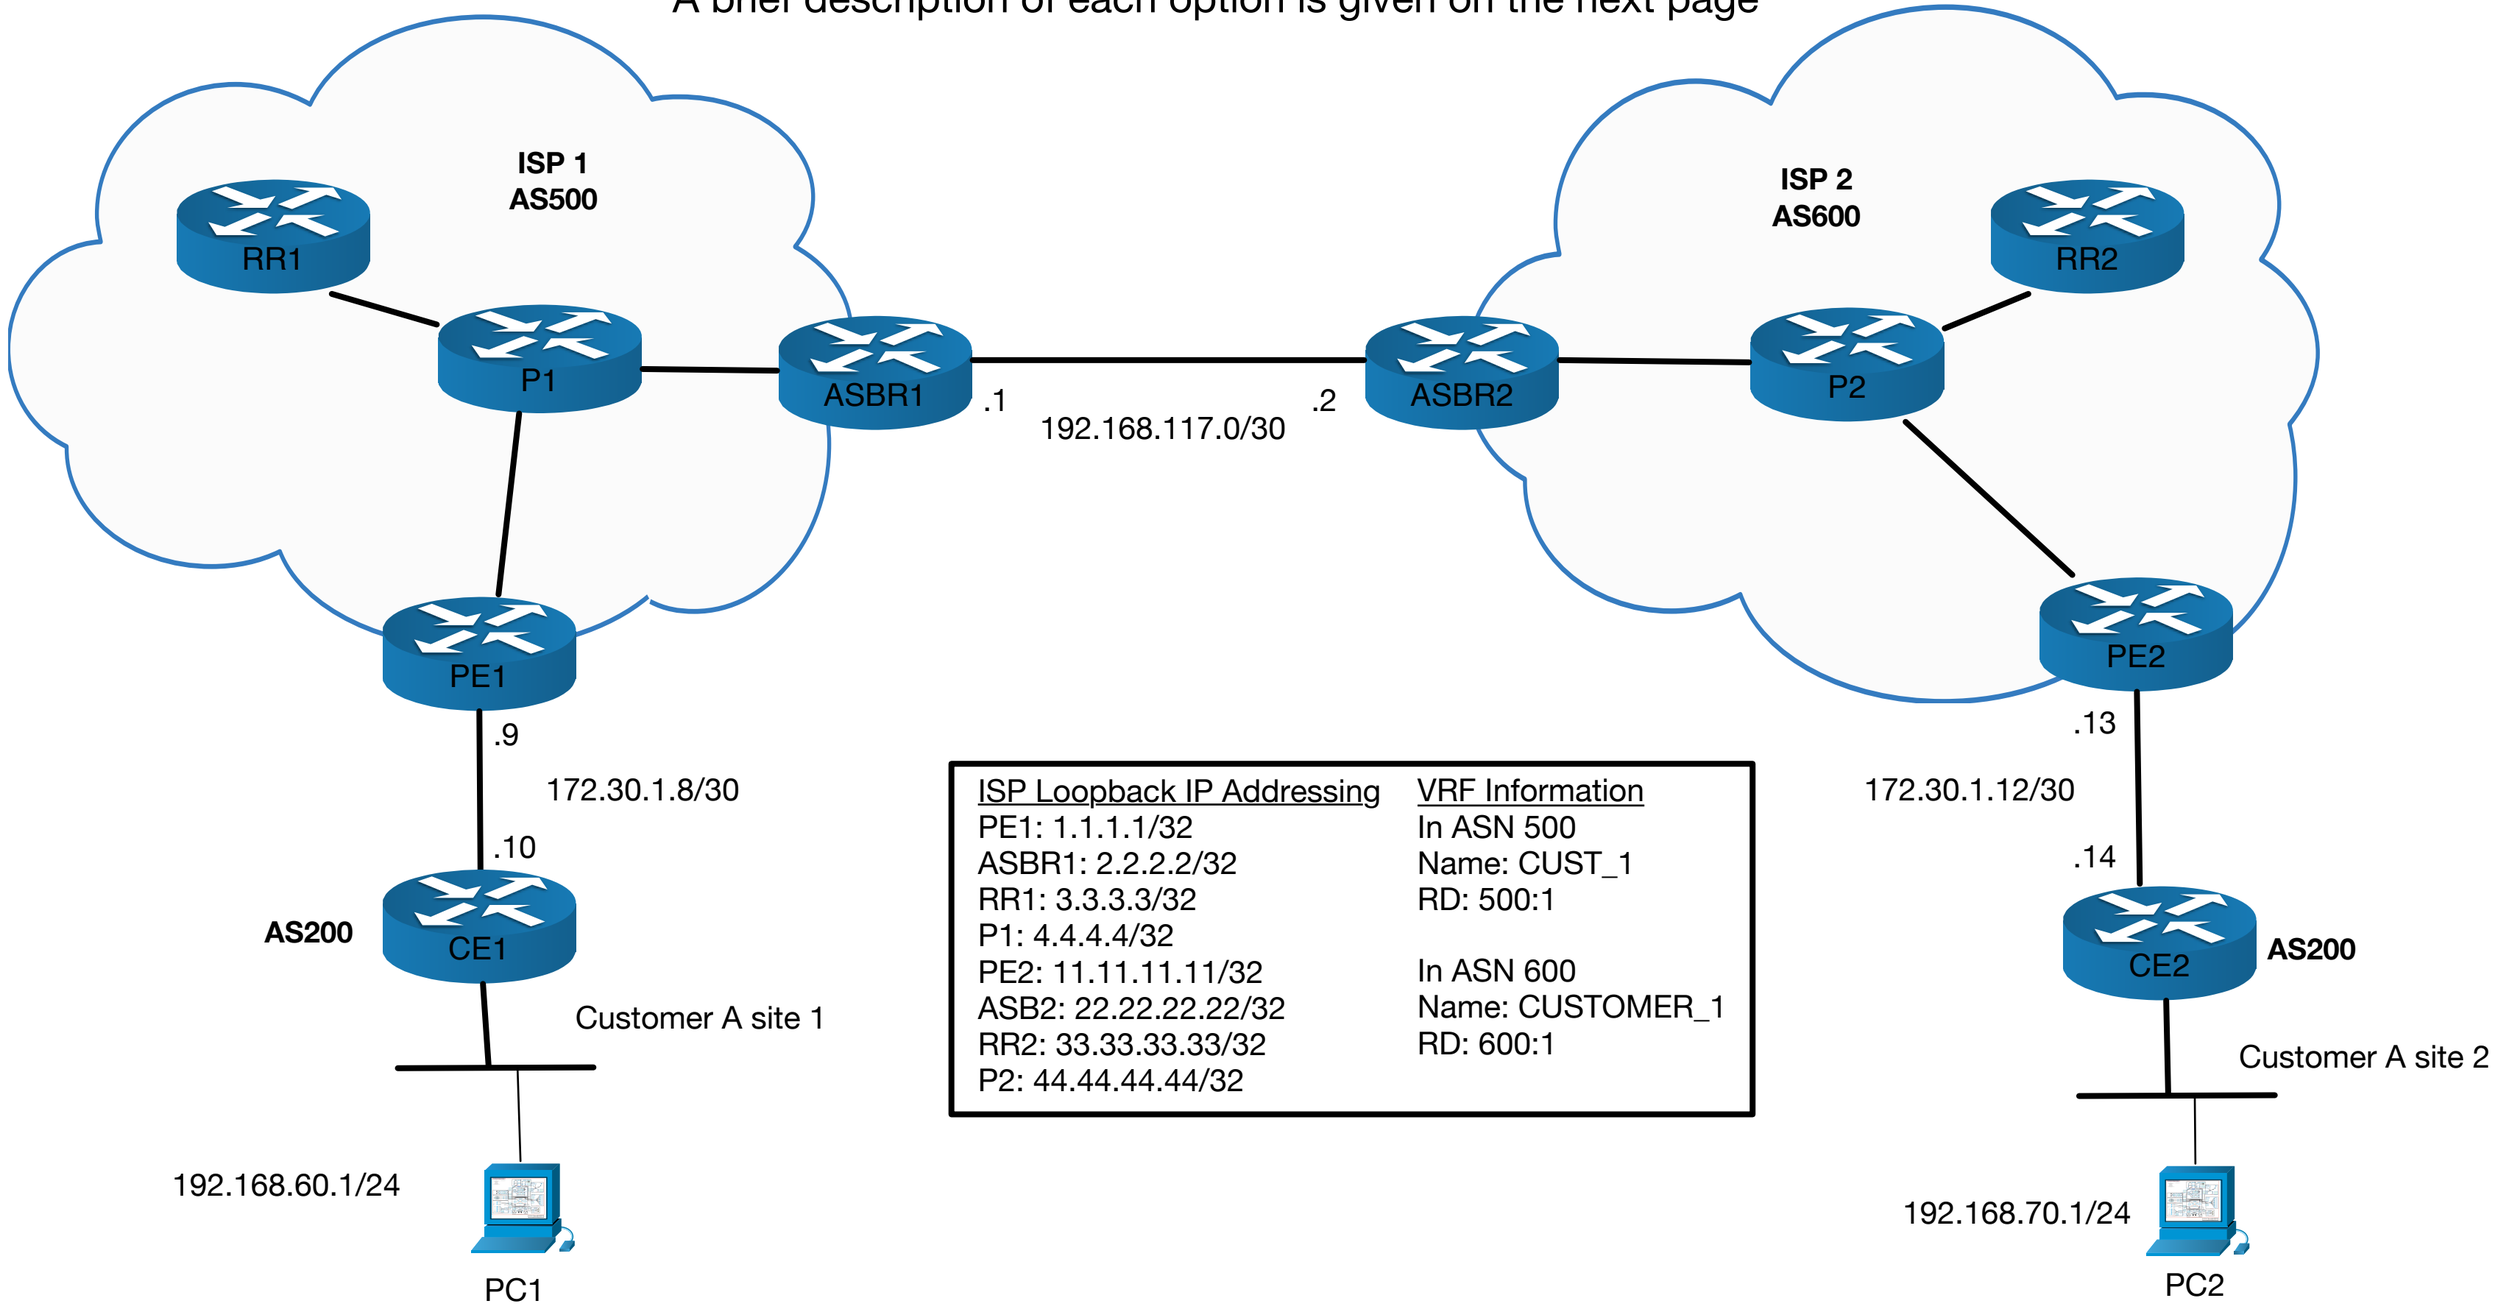

ISP Loopback IP Addressing	VRF Information
PE1: 1.1.1.1/32	In ASN 500
ASBR1: 2.2.2.2/32	Name: CUST_1
RR1: 3.3.3.3/32	RD: 500:1
P1: 4.4.4.4/32	
PE2: 11.11.11.11/32	In ASN 600
ASB2: 22.22.22.22/32	Name: CUSTOMER_1
RR2: 33.33.33.33/32	RD: 600:1
P2: 44.44.44.44/32	

Options Types

Option A:
In option A there is an IPv4 eBGP session between the ASBRs of each Autonomous System. Labels are not exchanged and there is one BGP session, and typically one sub-interface, per VRF. VRFs are locally configured on each ASBR meaning they have to store the routing tables of each VRF. This is the least scalable option and is analogous to what is usually done between PE and CE routers.

Option B:
VPNv4 routes, including associated labels are advertised between the ASBRs over a VPNv4 eBGP session. The ASBRs are configured to accept VPNv4 routes without needing the VRF locally configured (using ***no bgp default route-target filter***). The VPNv4 routes are then advertised into the local iBGP and assigned labels. There are two ways to accomplish this:

1) Using the iBGP ***next-hop-self*** command to make the local ASBR the next hop
2) Using the default behaviour of leaving the next-hop unchanged. In this scenario a host route, and corresponding label, to the neighbouring ASBR is local generated.

Both options are explored in this document.

Option C:
In this option a complete end-to-end LSP is created from one PE to another. An eBGP multi-hop session is established between the router reflectors of each Autonomous System. This session is used to exchange VPNv4 routes. For this to work, the next-hop addresses (typically loopbacks) are needed from the other Autonomous System. Both the route reflectors next-hop address (for the multi-hop eBGP session) and the egress PE next-hop addresses (to complete the end to end LSP) are needed. These next-hop addresses are advertised over an IPv4 eBGP session (with ***send-label*** enabled) between the ASBRs. From there, there are two methods the local AS can use to distribute the remote next-hop reachability information:

1) Advertise the remote next-hop addresses over eBGP and redistribute into the local IGP
2) Advertise the remote next-hop addresses over eBGP and then use iBGP and next-hop-self

Both are explored in this document

Option A

Concept

All next-hop/transport labels are assigned by LDP based on the IGPs route for the next-hop as seen in the BGP Update.

VRFs are configured on the ASBR routers and vpnv4 routes are imported into customer VRFs based on the IPv4 eBGPs session between the ASBRs.
On the link between the ASBRs, sub-interfaces or service instances are used for each VRF.
The BGP peering with the neighbouring ASBR is IPv4 only and no labels are exchanged (much like an eBGP exchange between a PE and a CE)

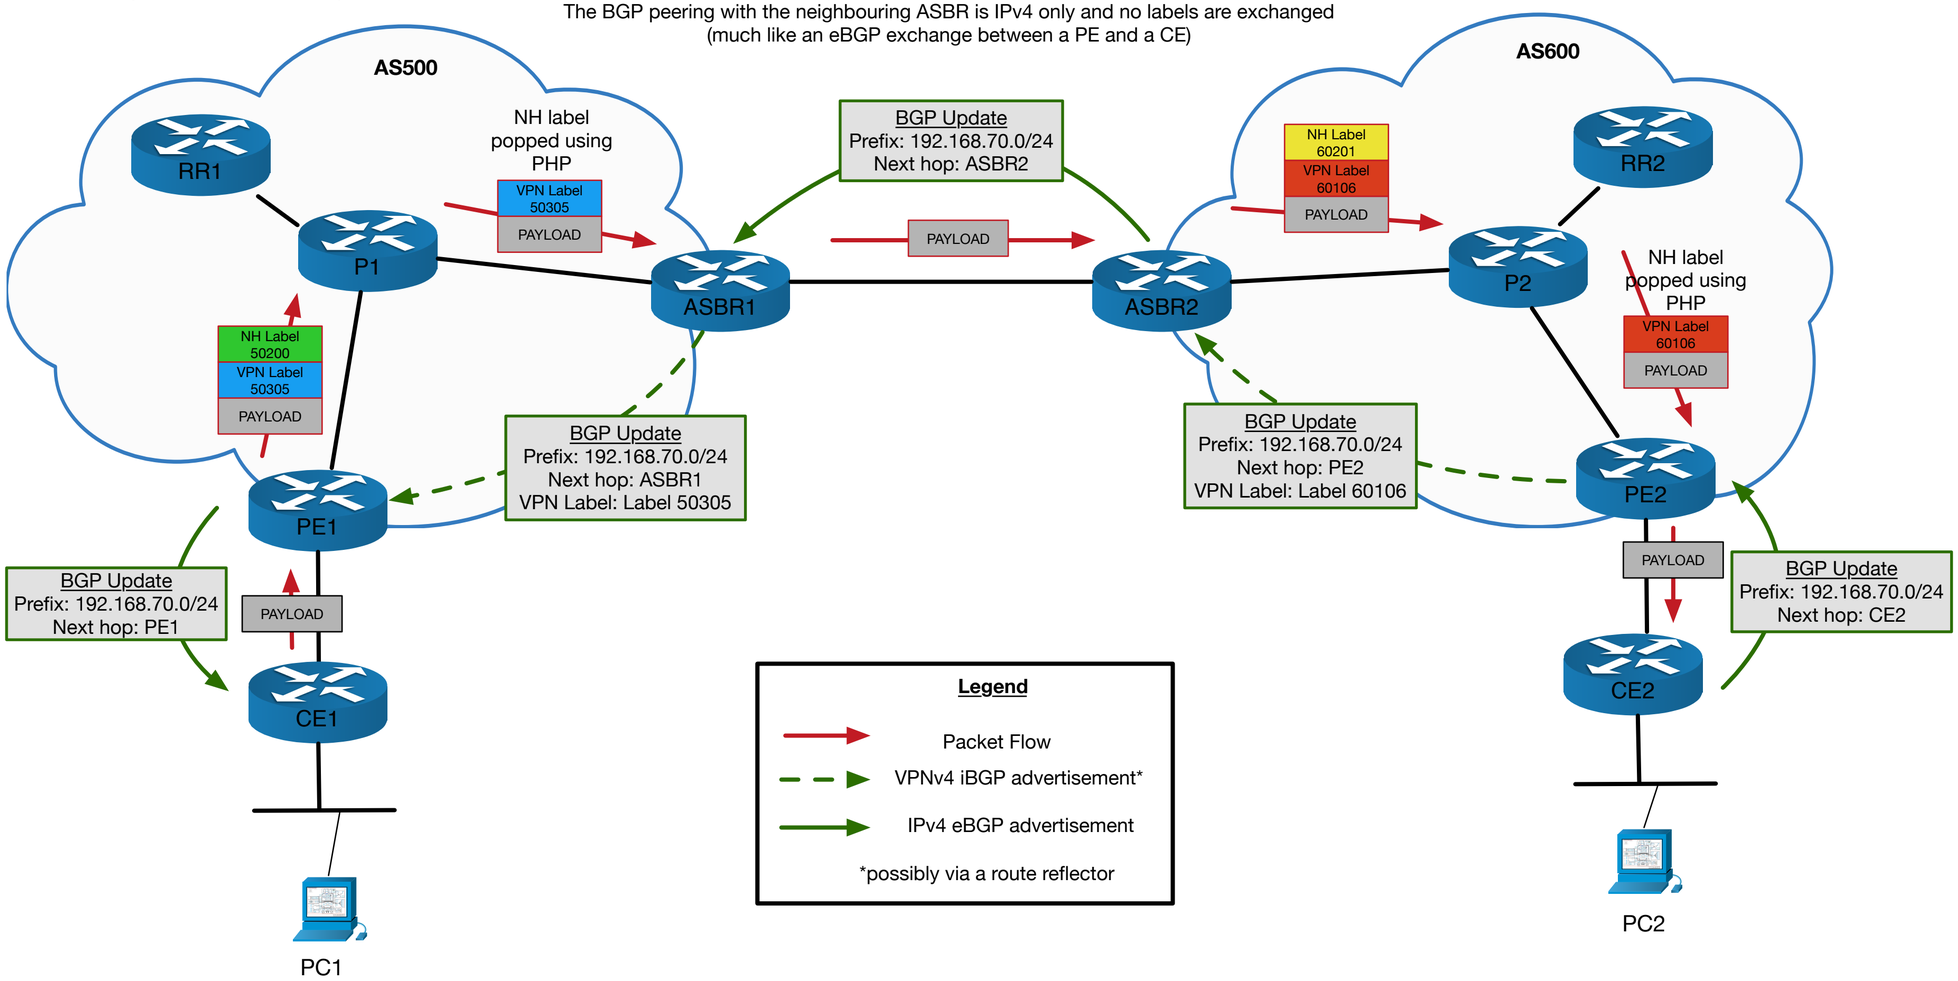

Option A

Output

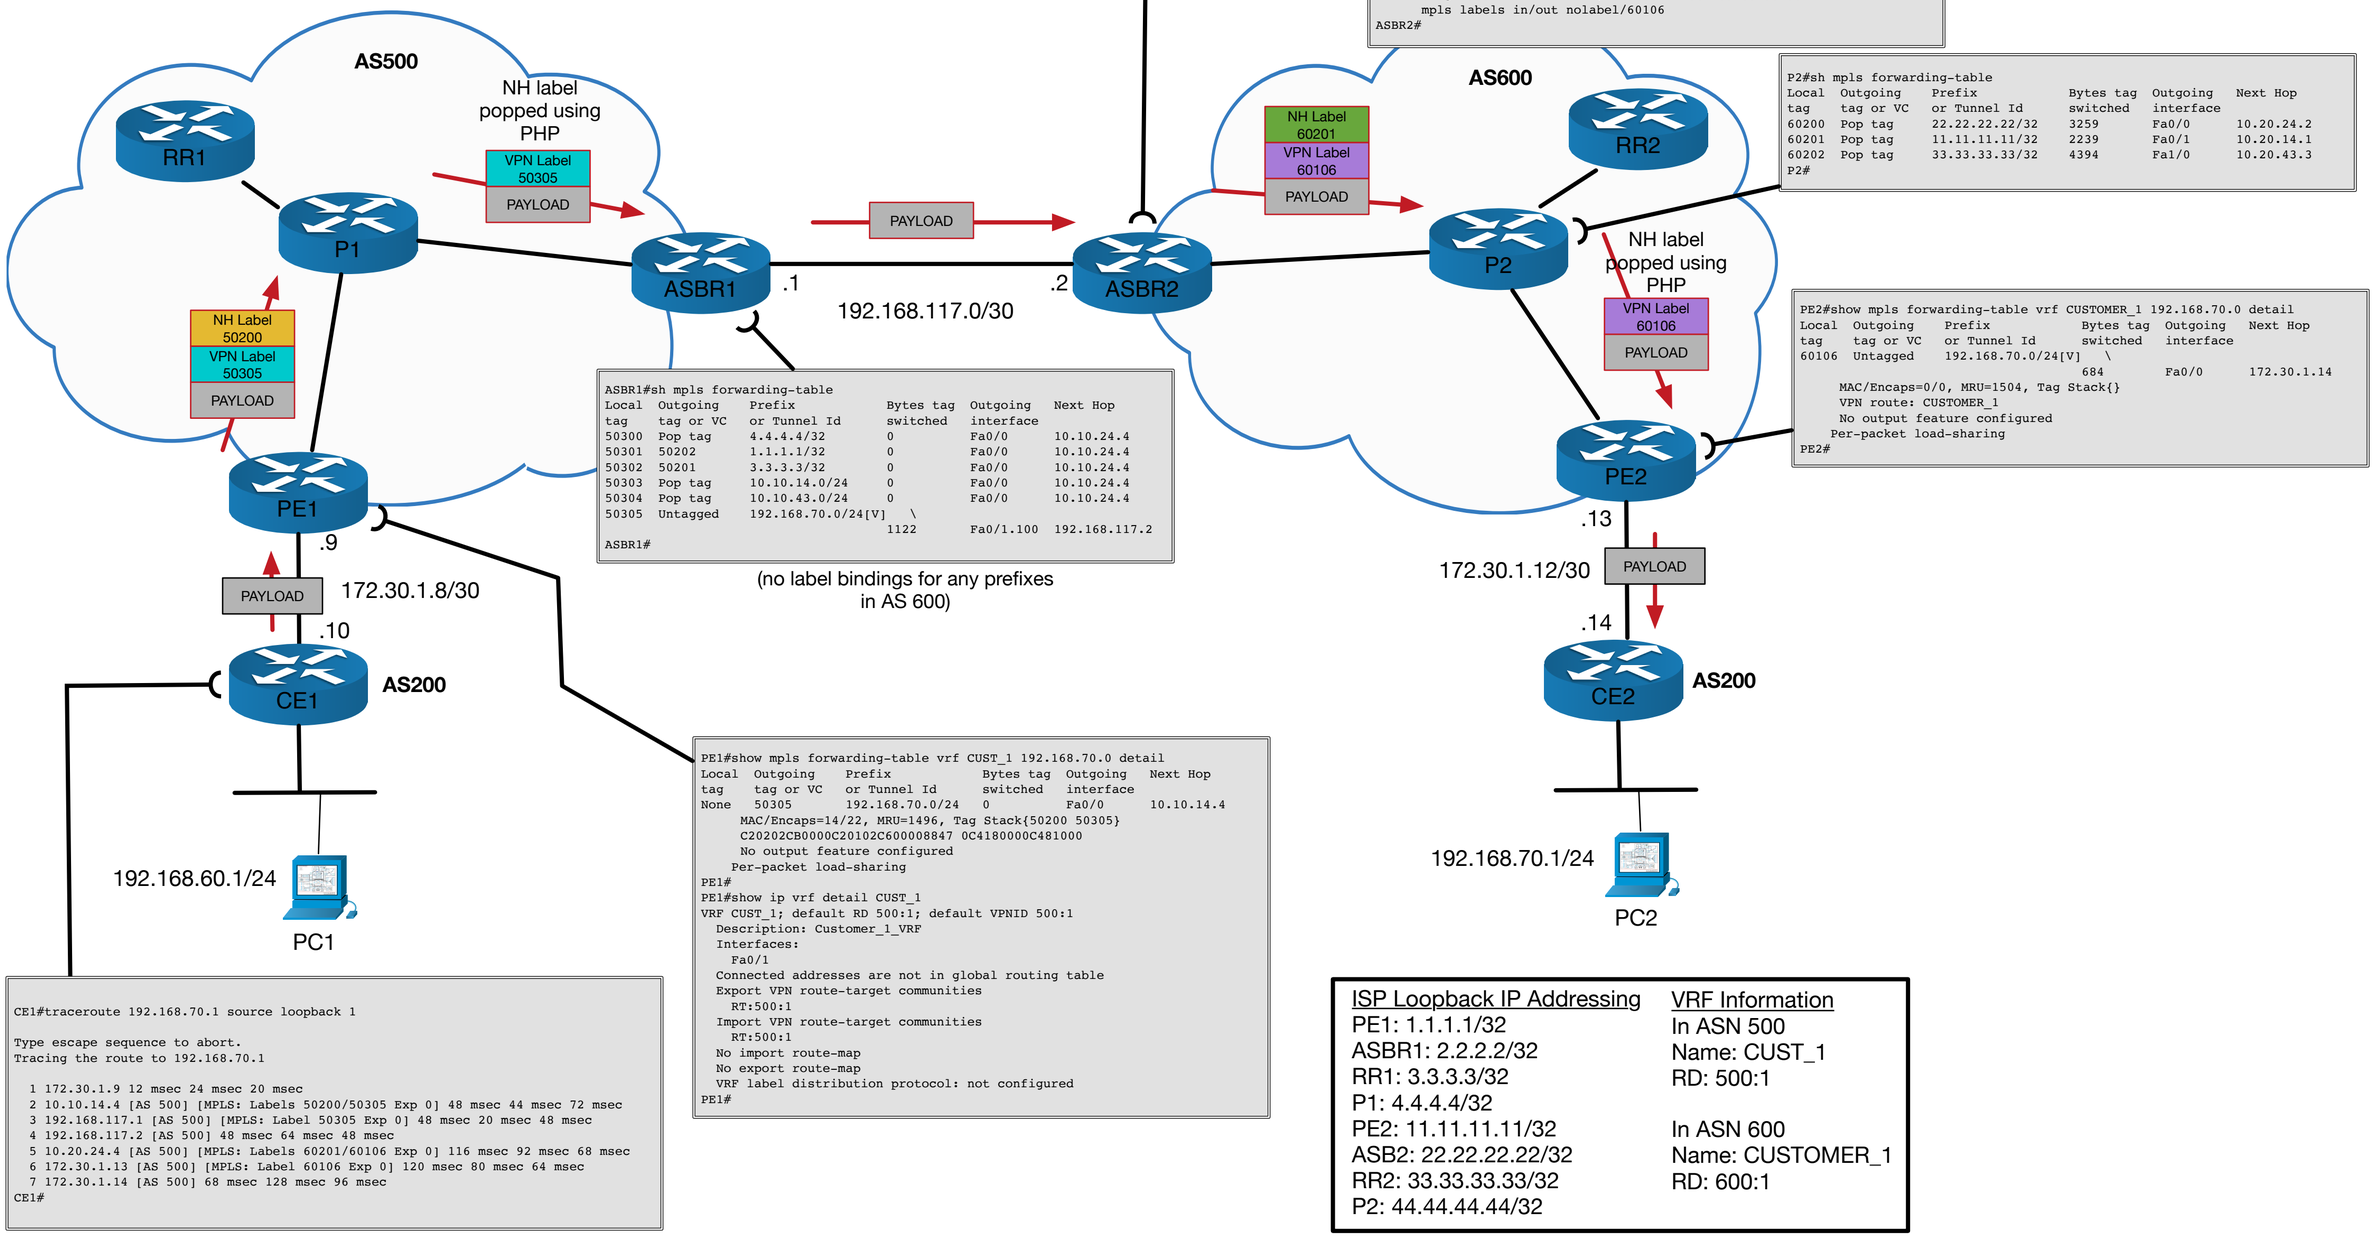


```
ASBR2#sh bgp vpnv4 unicast rd 600:1 192.168.70.1
BGP routing table entry for 600:1:192.168.70.0/24, version 5
Paths: (1 available, best #1, table CUSTOMER_1)
  Advertised to update-groups:
     1
  200
    11.11.11.11 (metric 20) from 33.33.33.33 (33.33.33.33)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:600:1
      Originator: 11.11.11.11, Cluster list: 33.33.33.33
      mpls labels in/out nolabel/60106
ASBR2#
```

```
P2#sh mpls forwarding-table
Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
60200  Pop tag     22.22.22.22/32   3259       Fa0/0       10.20.24.2
60201  Pop tag     11.11.11.11/32   2239       Fa0/1       10.20.14.1
60202  Pop tag     33.33.33.33/32   4394       Fa1/0       10.20.43.3
P2#
```

```
ASBR1#sh mpls forwarding-table
Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
50300  Pop tag     4.4.4.4/32       0          Fa0/0       10.10.24.4
50301  50202       1.1.1.1/32       0          Fa0/0       10.10.24.4
50302  50201       3.3.3.3/32       0          Fa0/0       10.10.24.4
50303  Pop tag     10.10.14.0/24    0          Fa0/0       10.10.24.4
50304  Pop tag     10.10.43.0/24    0          Fa0/0       10.10.24.4
50305  Untagged    192.168.70.0/24[V]   \
                                    1122       Fa0/1.100   192.168.117.2
ASBR1#
```

(no label bindings for any prefixes in AS 600)

```
PE2#show mpls forwarding-table vrf CUSTOMER_1 192.168.70.0 detail
Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
60106  Untagged    192.168.70.0/24[V]   \
                                    684        Fa0/0       172.30.1.14
     MAC/Encaps=0/0, MRU=1504, Tag Stack{}
     VPN route: CUSTOMER_1
     No output feature configured
   Per-packet load-sharing
PE2#
```

```
PE1#show mpls forwarding-table vrf CUST_1 192.168.70.0 detail
Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
None   50305       192.168.70.0/24  0          Fa0/0       10.10.14.4
     MAC/Encaps=14/22, MRU=1496, Tag Stack{50200 50305}
     C20202CB0000C20102C600008847 0C4180000C481000
     No output feature configured
   Per-packet load-sharing
PE1#
PE1#show ip vrf detail CUST_1
VRF CUST_1; default RD 500:1; default VPNID 500:1
  Description: Customer_1_VRF
  Interfaces:
    Fa0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:500:1
  Import VPN route-target communities
    RT:500:1
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
PE1#
```

```
CE1#traceroute 192.168.70.1 source loopback 1

Type escape sequence to abort.
Tracing the route to 192.168.70.1

  1 172.30.1.9 12 msec 24 msec 20 msec
  2 10.10.14.4 [AS 500] [MPLS: Labels 50200/50305 Exp 0] 48 msec 44 msec 72 msec
  3 192.168.117.1 [AS 500] [MPLS: Label 50305 Exp 0] 48 msec 20 msec 48 msec
  4 192.168.117.2 [AS 500] 48 msec 64 msec 48 msec
  5 10.20.24.4 [AS 500] [MPLS: Labels 60201/60106 Exp 0] 116 msec 92 msec 68 msec
  6 172.30.1.13 [AS 500] [MPLS: Label 60106 Exp 0] 120 msec 80 msec 64 msec
  7 172.30.1.14 [AS 500] 68 msec 128 msec 96 msec
CE1#
```

ISP Loopback IP Addressing	VRF Information
PE1: 1.1.1.1/32	In ASN 500
ASBR1: 2.2.2.2/32	Name: CUST_1
RR1: 3.3.3.3/32	RD: 500:1
P1: 4.4.4.4/32	
PE2: 11.11.11.11/32	In ASN 600
ASB2: 22.22.22.22/32	Name: CUSTOMER_1
RR2: 33.33.33.33/32	RD: 600:1
P2: 44.44.44.44/32	

Option B - Next Hop Self

Concept

In option B a VPNv4 eBGP session is used to exchange VPNv4 routes with the opposing AS. Here ASBR1 is setting the next-hop to itself when advertising this route internally over iBGP.

All next-hop/transport labels are assigned by LDP based on the IGPs route for the next-hop as seen in the BGP Update.

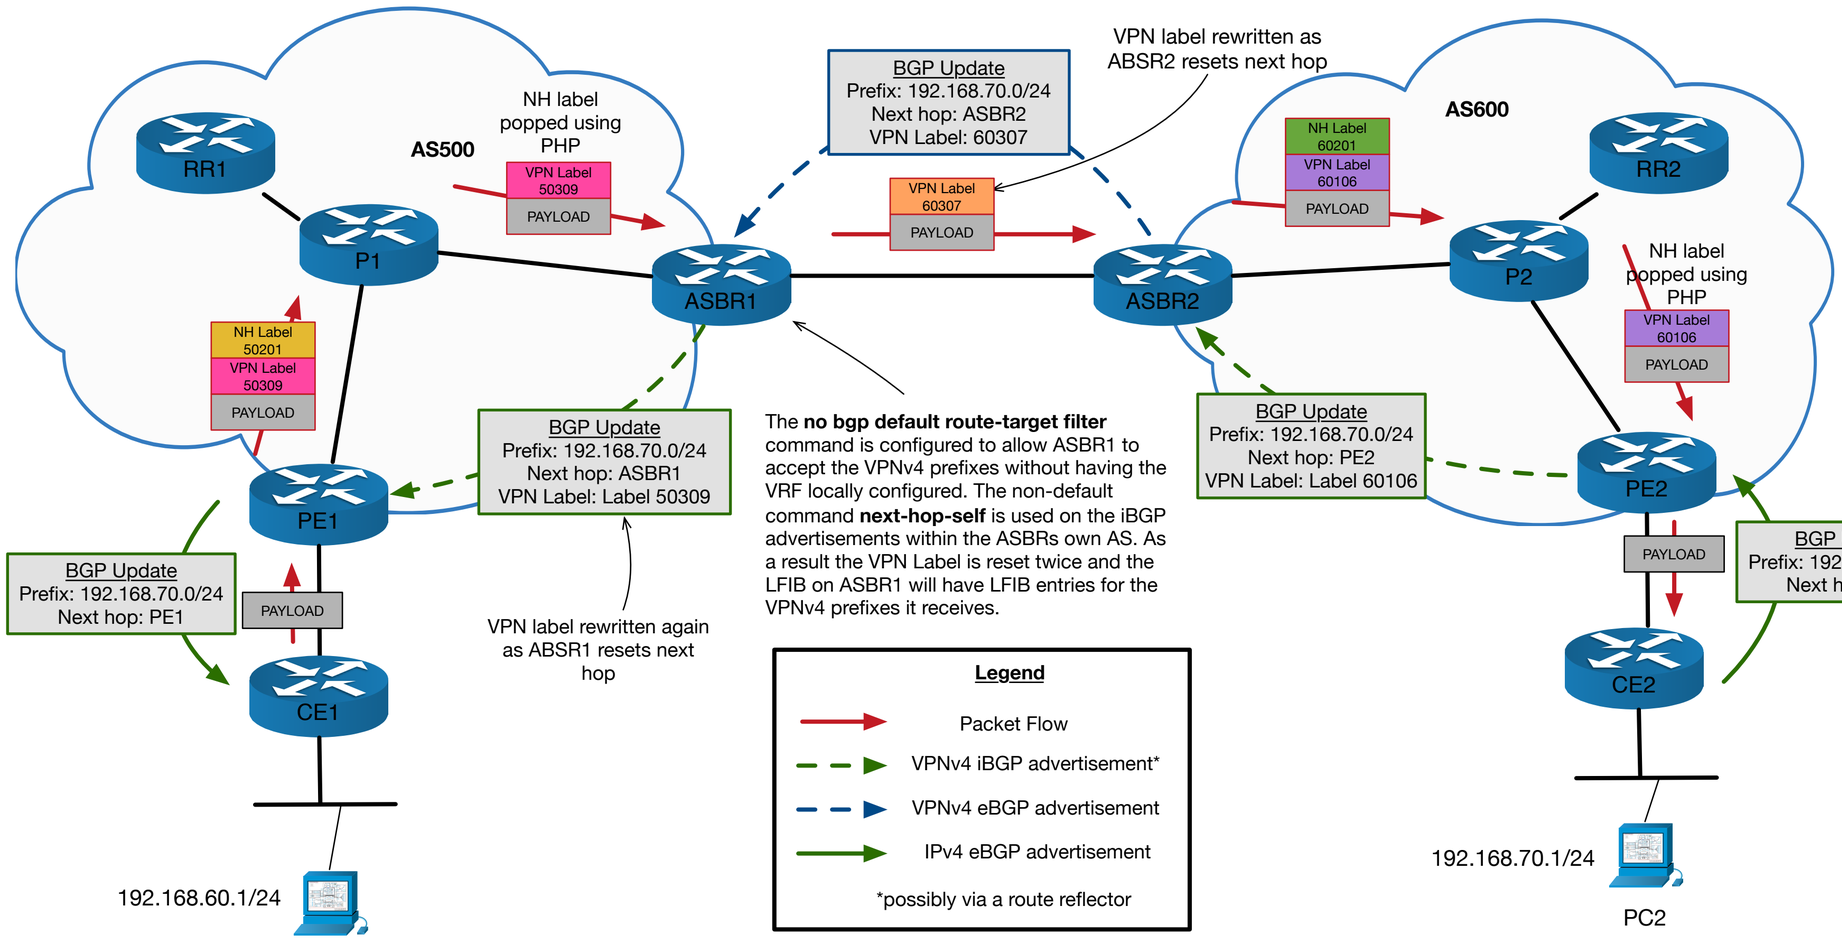

The **no bgp default route-target filter** command is configured to allow ASBR1 to accept the VPNv4 prefixes without having the VRF locally configured. The non-default command **next-hop-self** is used on the iBGP advertisements within the ASBRs own AS. As a result the VPN Label is reset twice and the LFIB on ASBR1 will have LFIB entries for the VPNv4 prefixes it receives.

Legend

- Packet Flow
- VPNv4 iBGP advertisement*
- VPNv4 eBGP advertisement
- IPv4 eBGP advertisement

*possibly via a route reflector

Option B - Next Hop Self

Output

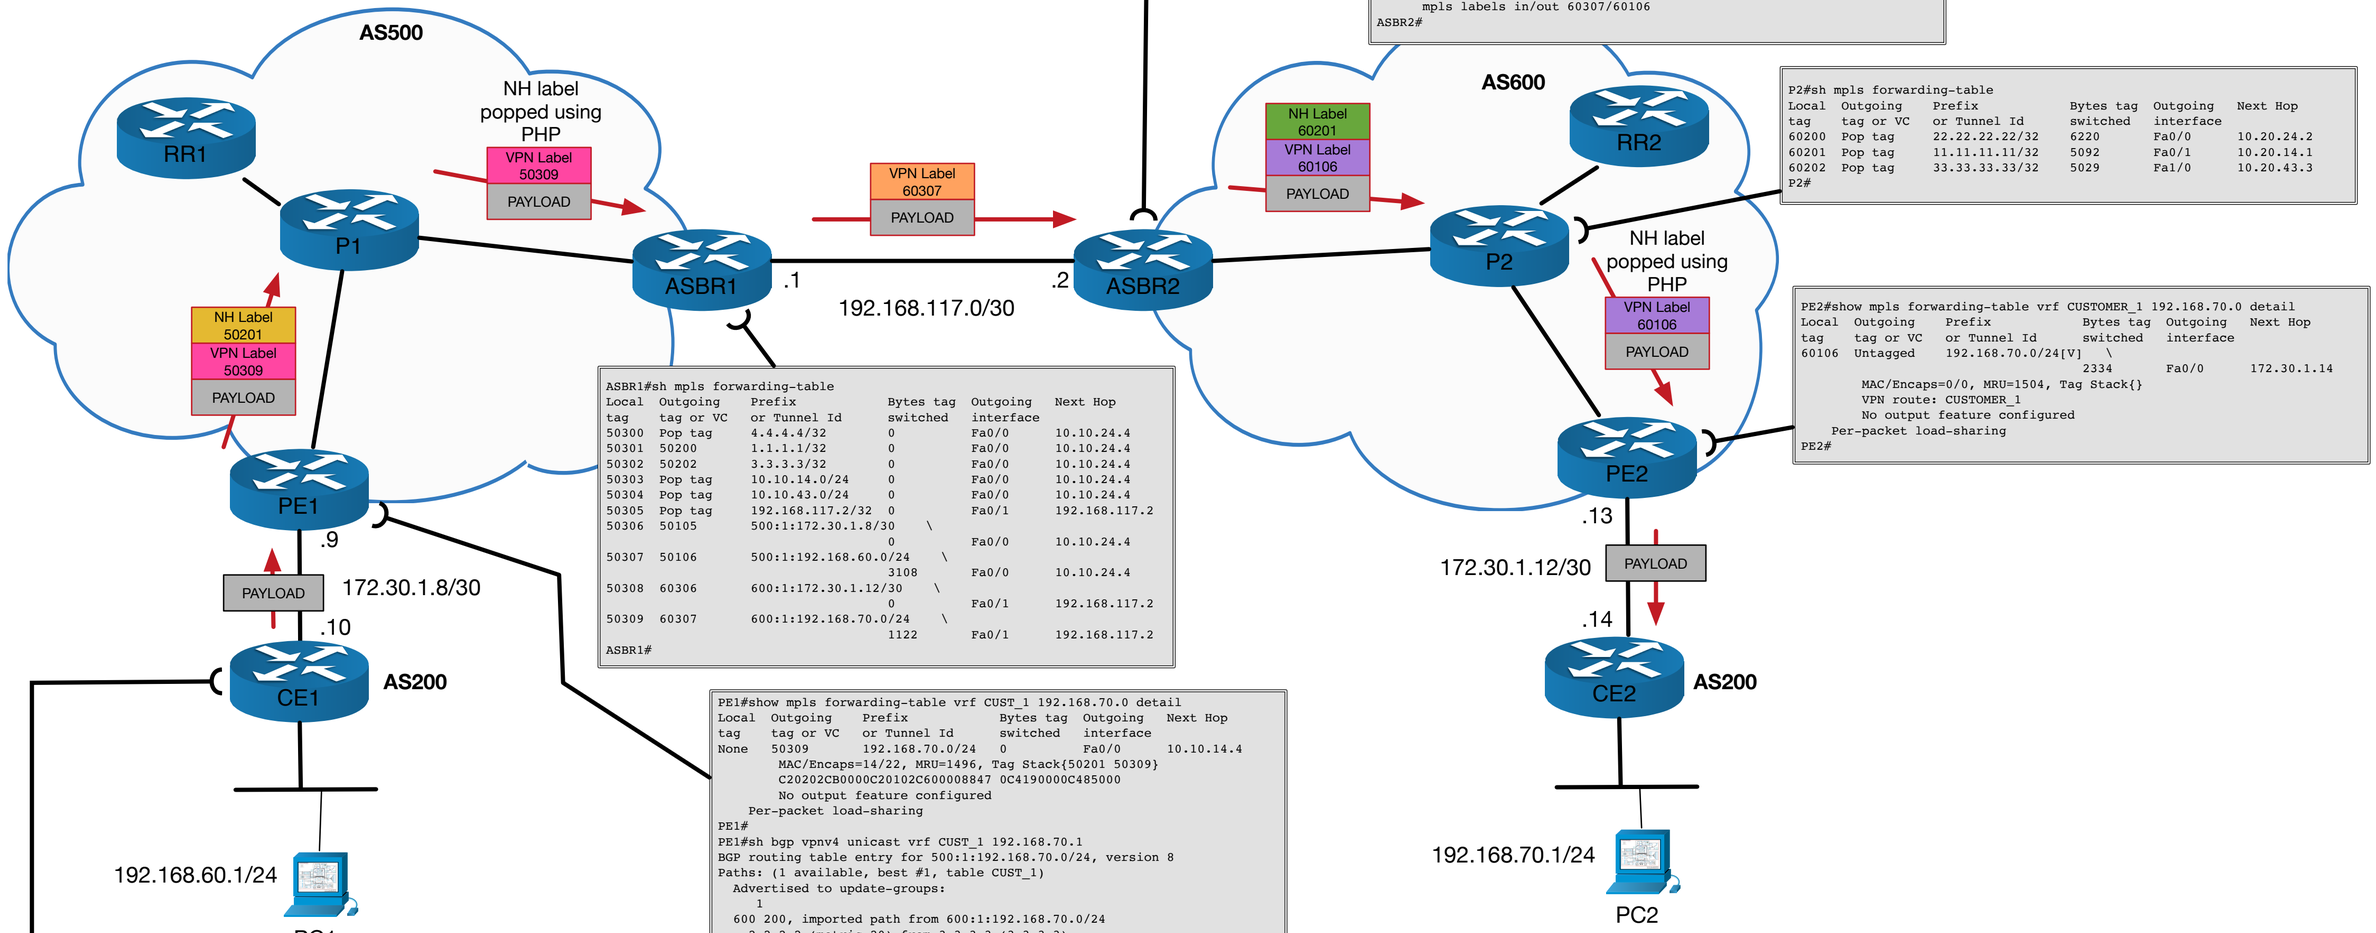


```
ASBR2#sh bgp vpnv4 unicast rd 600:1 192.168.70.1
BGP routing table entry for 600:1:192.168.70.0/24, version 3
Paths: (1 available, best #1, no table)
  Advertised to update-groups:
     1
  200
    11.11.11.11 (metric 20) from 33.33.33.33 (33.33.33.33)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:600:1
      Originator: 11.11.11.11, Cluster list: 33.33.33.33
      mpls labels in/out 60307/60106
ASBR2#
```

```
P2#sh mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
60200  Pop tag     22.22.22.22/32    6220       Fa0/0       10.20.24.2
60201  Pop tag     11.11.11.11/32    5092       Fa0/1       10.20.14.1
60202  Pop tag     33.33.33.33/32    5029       Fa1/0       10.20.43.3
P2#
```

```
PE2#show mpls forwarding-table vrf CUSTOMER_1 192.168.70.0 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
60106  Untagged    192.168.70.0/24[V]   \
                                      2334       Fa0/0       172.30.1.14
        MAC/Encaps=0/0, MRU=1504, Tag Stack{}
        VPN route: CUSTOMER_1
        No output feature configured
    Per-packet load-sharing
PE2#
```

```
ASBR1#sh mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
50300  Pop tag     4.4.4.4/32        0          Fa0/0       10.10.24.4
50301  50200       1.1.1.1/32        0          Fa0/0       10.10.24.4
50302  50202       3.3.3.3/32        0          Fa0/0       10.10.24.4
50303  Pop tag     10.10.14.0/24     0          Fa0/0       10.10.24.4
50304  Pop tag     10.10.43.0/24     0          Fa0/0       10.10.24.4
50305  Pop tag     192.168.117.2/32  0          Fa0/1       192.168.117.2
50306  50105       500:1:172.30.1.8/30   \
                                     0          Fa0/0       10.10.24.4
50307  50106       500:1:192.168.60.0/24    \
                                     3108       Fa0/0       10.10.24.4
50308  60306       600:1:172.30.1.12/30     \
                                     0          Fa0/1       192.168.117.2
50309  60307       600:1:192.168.70.0/24    \
                                     1122       Fa0/1       192.168.117.2
ASBR1#
```

```
PE1#show mpls forwarding-table vrf CUST_1 192.168.70.0 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
None   50309       192.168.70.0/24   0          Fa0/0       10.10.14.4
       MAC/Encaps=14/22, MRU=1496, Tag Stack{50201 50309}
       C20202CB0000C20102C600008847 0C4190000C485000
       No output feature configured
    Per-packet load-sharing
PE1#
PE1#sh bgp vpnv4 unicast vrf CUST_1 192.168.70.1
BGP routing table entry for 500:1:192.168.70.0/24, version 8
Paths: (1 available, best #1, table CUST_1)
  Advertised to update-groups:
     1
  600 200, imported path from 600:1:192.168.70.0/24
    2.2.2.2 (metric 20) from 3.3.3.3 (3.3.3.3)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:600:1
      Originator: 2.2.2.2, Cluster list: 3.3.3.3
      mpls labels in/out nolabel/50309
PE1#
PE1#show ip vrf detail CUST_1
VRF CUST_1; default RD 500:1; default VPNID 500:1
  Description: Customer_1_VRF
  Interfaces:
    Fa0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:500:1
  Import VPN route-target communities
    RT:500:1                  RT:600:1
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
PE1#
```

```
CE1#traceroute 192.168.70.1 source loopback 1

Type escape sequence to abort.
Tracing the route to 192.168.70.1

  1 172.30.1.9 16 msec 12 msec 16 msec
  2 10.10.14.4 [AS 500] [MPLS: Labels 50201/50309 Exp 0] 112 msec 112 msec 84 msec
  3 10.10.24.2 [AS 500] [MPLS: Label 50309 Exp 0] 92 msec 84 msec 60 msec
  4 192.168.117.2 [AS 500] [MPLS: Label 60307 Exp 0] 80 msec 96 msec 80 msec
  5 10.20.24.4 [AS 500] [MPLS: Labels 60201/60106 Exp 0] 124 msec 88 msec 60 msec
  6 172.30.1.13 [AS 600] [MPLS: Label 60106 Exp 0] 76 msec 96 msec 72 msec
  7 172.30.1.14 [AS 600] 104 msec 120 msec 96 msec
CE1#
```

ISP Loopback IP Addressing	VRF Information
PE1: 1.1.1.1/32	In ASN 500
ASBR1: 2.2.2.2/32	Name: CUST_1
RR1: 3.3.3.3/32	RD: 500:1
P1: 4.4.4.4/32	
PE2: 11.11.11.11/32	In ASN 600
ASB2: 22.22.22.22/32	Name: CUSTOMER_1
RR2: 33.33.33.33/32	RD: 600:1
P2: 44.44.44.44/32	

by Steven Crutchley

Option B - Next Hop Unchanged

Concept

All next-hop/transport labels are assigned by LDP based on the IGPs route for the next-hop as seen in the BGP Update.

In option B a VPNv4 eBGP session is used to exchange VPNv4 routes with the opposing AS. Here ASBR1 is leaving the next-hop unchanged when advertising this route internally over iBGP.

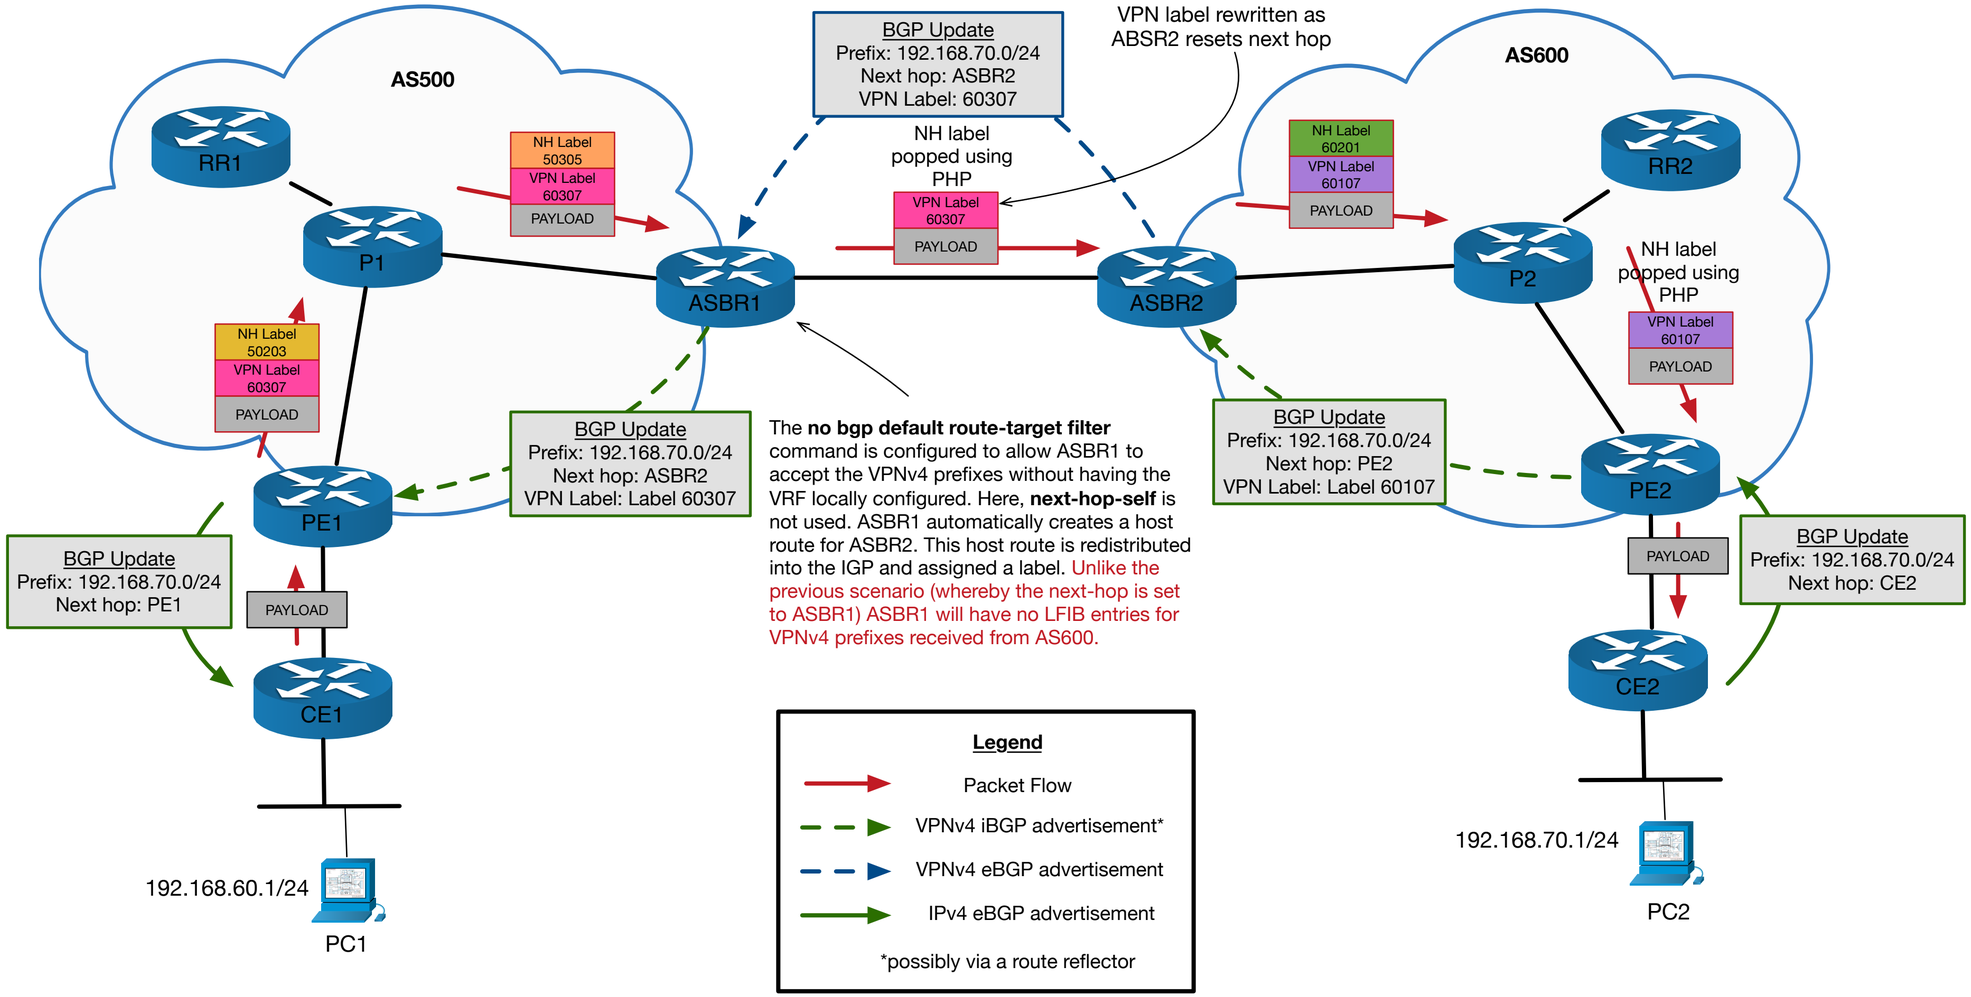

The **no bgp default route-target filter** command is configured to allow ASBR1 to accept the VPNv4 prefixes without having the VRF locally configured. Here, **next-hop-self** is not used. ASBR1 automatically creates a host route for ASBR2. This host route is redistributed into the IGP and assigned a label. Unlike the previous scenario (whereby the next-hop is set to ASBR1) ASBR1 will have no LFIB entries for VPNv4 prefixes received from AS600.

Legend

- Packet Flow
- VPNv4 iBGP advertisement*
- VPNv4 eBGP advertisement
- IPv4 eBGP advertisement

*possibly via a route reflector

Option B - Next Hop Unchanged

Output

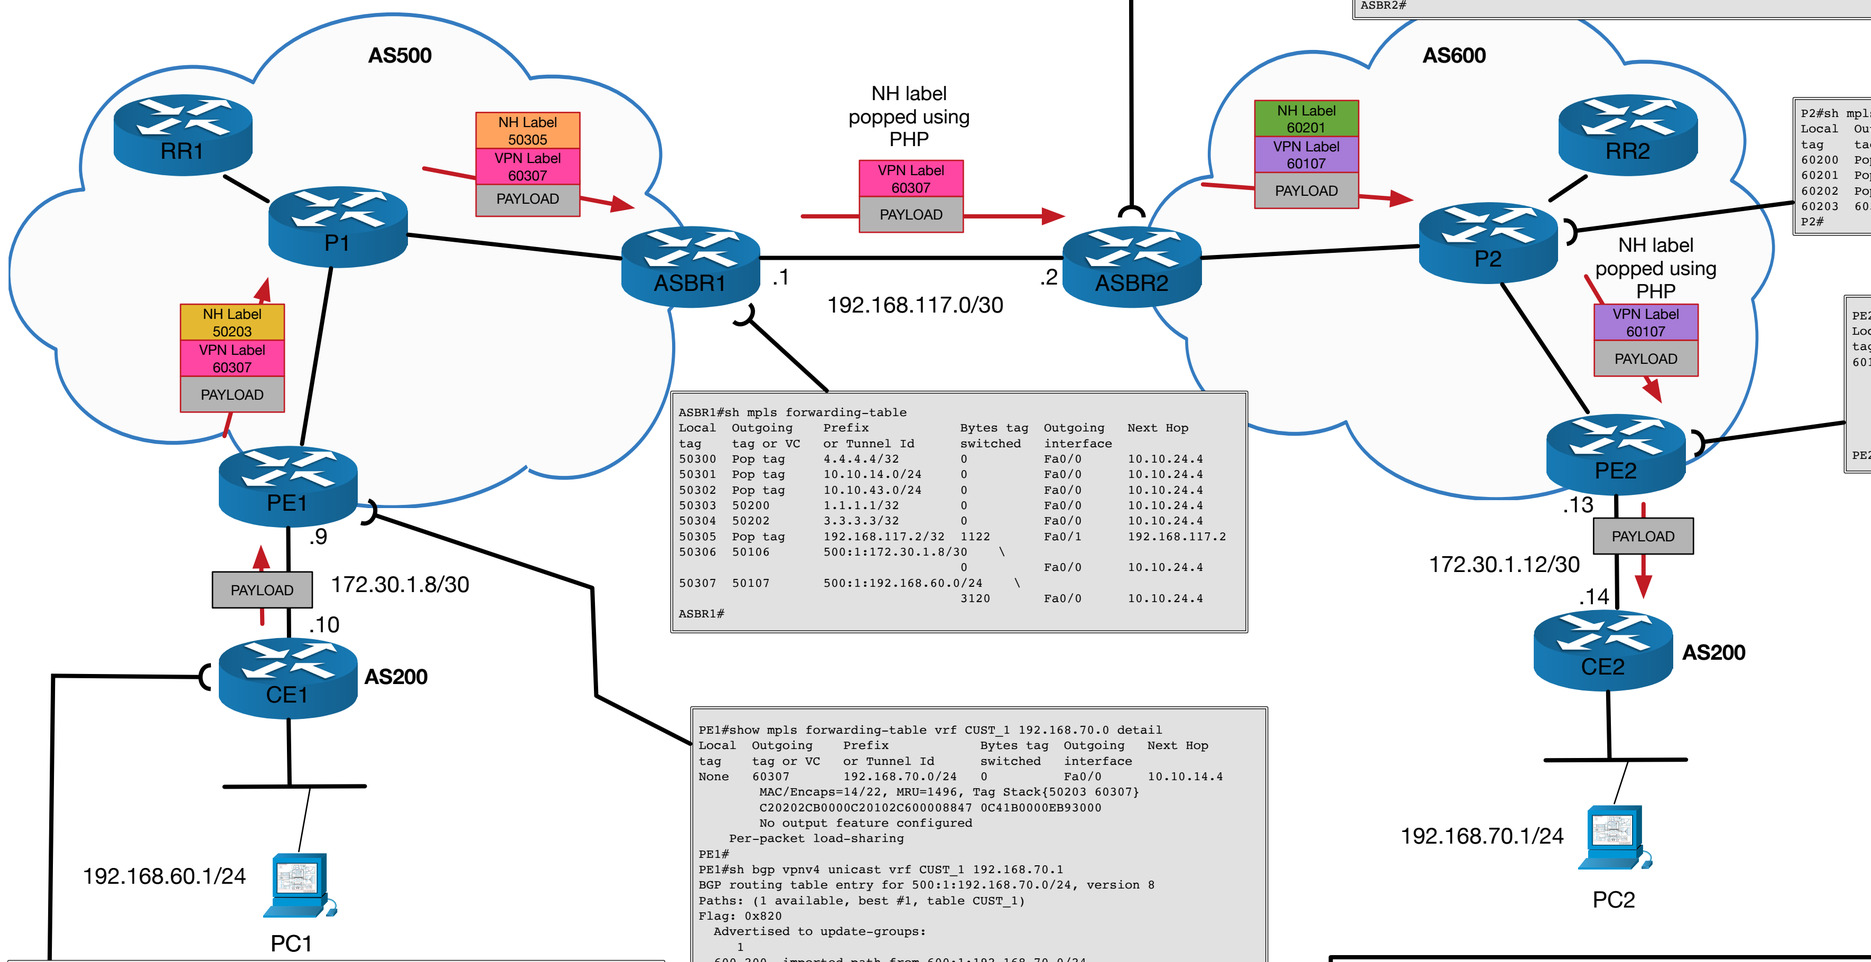


```
ASBR2#sh bgp vpnv4 unicast rd 600:1 192.168.70.1
BGP routing table entry for 600:1:192.168.70.0/24, version 3
Paths: (1 available, best #1, no table)
Flag: 0x820
  Advertised to update-groups:
     1
  200
    11.11.11.11 (metric 20) from 33.33.33.33 (33.33.33.33)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:600:1
      Originator: 11.11.11.11, Cluster list: 33.33.33.33
      mpls labels in/out 60307/60107
ASBR2#
```

```
P2#sh mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
60200  Pop tag     22.22.22.22/32    1344       Fa0/0       10.20.24.2
60201  Pop tag     11.11.11.11/32    3534       Fa0/1       10.20.14.1
60202  Pop tag     33.33.33.33/32    2959       Fa1/0       10.20.43.3
60203  60305       192.168.117.1/32  3120       Fa0/0       10.20.24.2
P2#
```

```
ASBR1#sh mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
50300  Pop tag     4.4.4.4/32        0          Fa0/0       10.10.24.4
50301  Pop tag     10.10.14.0/24     0          Fa0/0       10.10.24.4
50302  Pop tag     10.10.43.0/24     0          Fa0/0       10.10.24.4
50303  50200       1.1.1.1/32        0          Fa0/0       10.10.24.4
50304  50202       3.3.3.3/32        0          Fa0/0       10.10.24.4
50305  Pop tag     192.168.117.2/32  1122       Fa0/1       192.168.117.2
50306  50106       500:1:172.30.1.8/30    \
                                     0          Fa0/0       10.10.24.4
50307  50107       500:1:192.168.60.0/24  \
                                     3120       Fa0/0       10.10.24.4
ASBR1#
```

```
PE2#show mpls forwarding-table vrf CUSTOMER_1 192.168.70.0 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
60107  Untagged    192.168.70.0/24[V]    \
                                     2346       Fa0/0       172.30.1.14
        MAC/Encaps=0/0, MRU=1504, Tag Stack{}
        VPN route: CUSTOMER_1
        No output feature configured
    Per-packet load-sharing
PE2#
```

```
PE1#show mpls forwarding-table vrf CUST_1 192.168.70.0 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
None   60307       192.168.70.0/24   0          Fa0/0       10.10.14.4
        MAC/Encaps=14/22, MRU=1496, Tag Stack{50203 60307}
        C20202CB0000C20102C600008847 0C41B0000EB93000
        No output feature configured
    Per-packet load-sharing
PE1#
PE1#sh bgp vpnv4 unicast vrf CUST_1 192.168.70.1
BGP routing table entry for 500:1:192.168.70.0/24, version 8
Paths: (1 available, best #1, table CUST_1)
Flag: 0x820
  Advertised to update-groups:
     1
  600 200, imported path from 600:1:192.168.70.0/24
    192.168.117.2 (metric 20) from 3.3.3.3 (3.3.3.3)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:600:1
      Originator: 2.2.2.2, Cluster list: 3.3.3.3
      mpls labels in/out nolabel/60307
PE1#
```

```
CE1#traceroute 192.168.70.1 source loopback 1

Type escape sequence to abort.
Tracing the route to 192.168.70.1

  1 172.30.1.9 24 msec 4 msec 16 msec
  2 10.10.14.4 [AS 500] [MPLS: Labels 50203/60307 Exp 0] 104 msec 92 msec 64 msec
  3 10.10.24.2 [AS 500] [MPLS: Labels 50305/60307 Exp 0] 84 msec 88 msec 120 msec
  4 192.168.117.2 [AS 500] [MPLS: Label 60307 Exp 0] 104 msec 76 msec 76 msec
  5 10.20.24.4 [AS 500] [MPLS: Labels 60201/60107 Exp 0] 124 msec 120 msec 88 msec
  6 172.30.1.13 [AS 500] [MPLS: Label 60107 Exp 0] 60 msec 112 msec 72 msec
  7 172.30.1.14 [AS 500] 116 msec 100 msec 124 msec
CE1#
```

ISP Loopback IP Addressing	VRF Information
PE1: 1.1.1.1/32	In ASN 500
ASBR1: 2.2.2.2/32	Name: CUST_1
RR1: 3.3.3.3/32	RD: 500:1
P1: 4.4.4.4/32	
PE2: 11.11.11.11/32	In ASN 600
ASB2: 22.22.22.22/32	Name: CUSTOMER_1
RR2: 33.33.33.33/32	RD: 600:1
P2: 44.44.44.44/32	

Option C - Redistributing remote ASN next-hop prefixes into local IGP

Concept

All next-hop/transport labels are assigned by LDP based on the IGPs route for the next-hop as seen in the BGP Update.

Note: Only the PE and RR loopbacks need to be exchanged. The RRs need their next hops to establish the multihop eBGP session. The PE loopbacks are needed for the end-to-end LSP. P and ASBR loopbacks need not be exchanged.

The create an end-to-end LSP, Route Reflectors do not reset the next hop (***neighbour x next-hop-unchanged***).

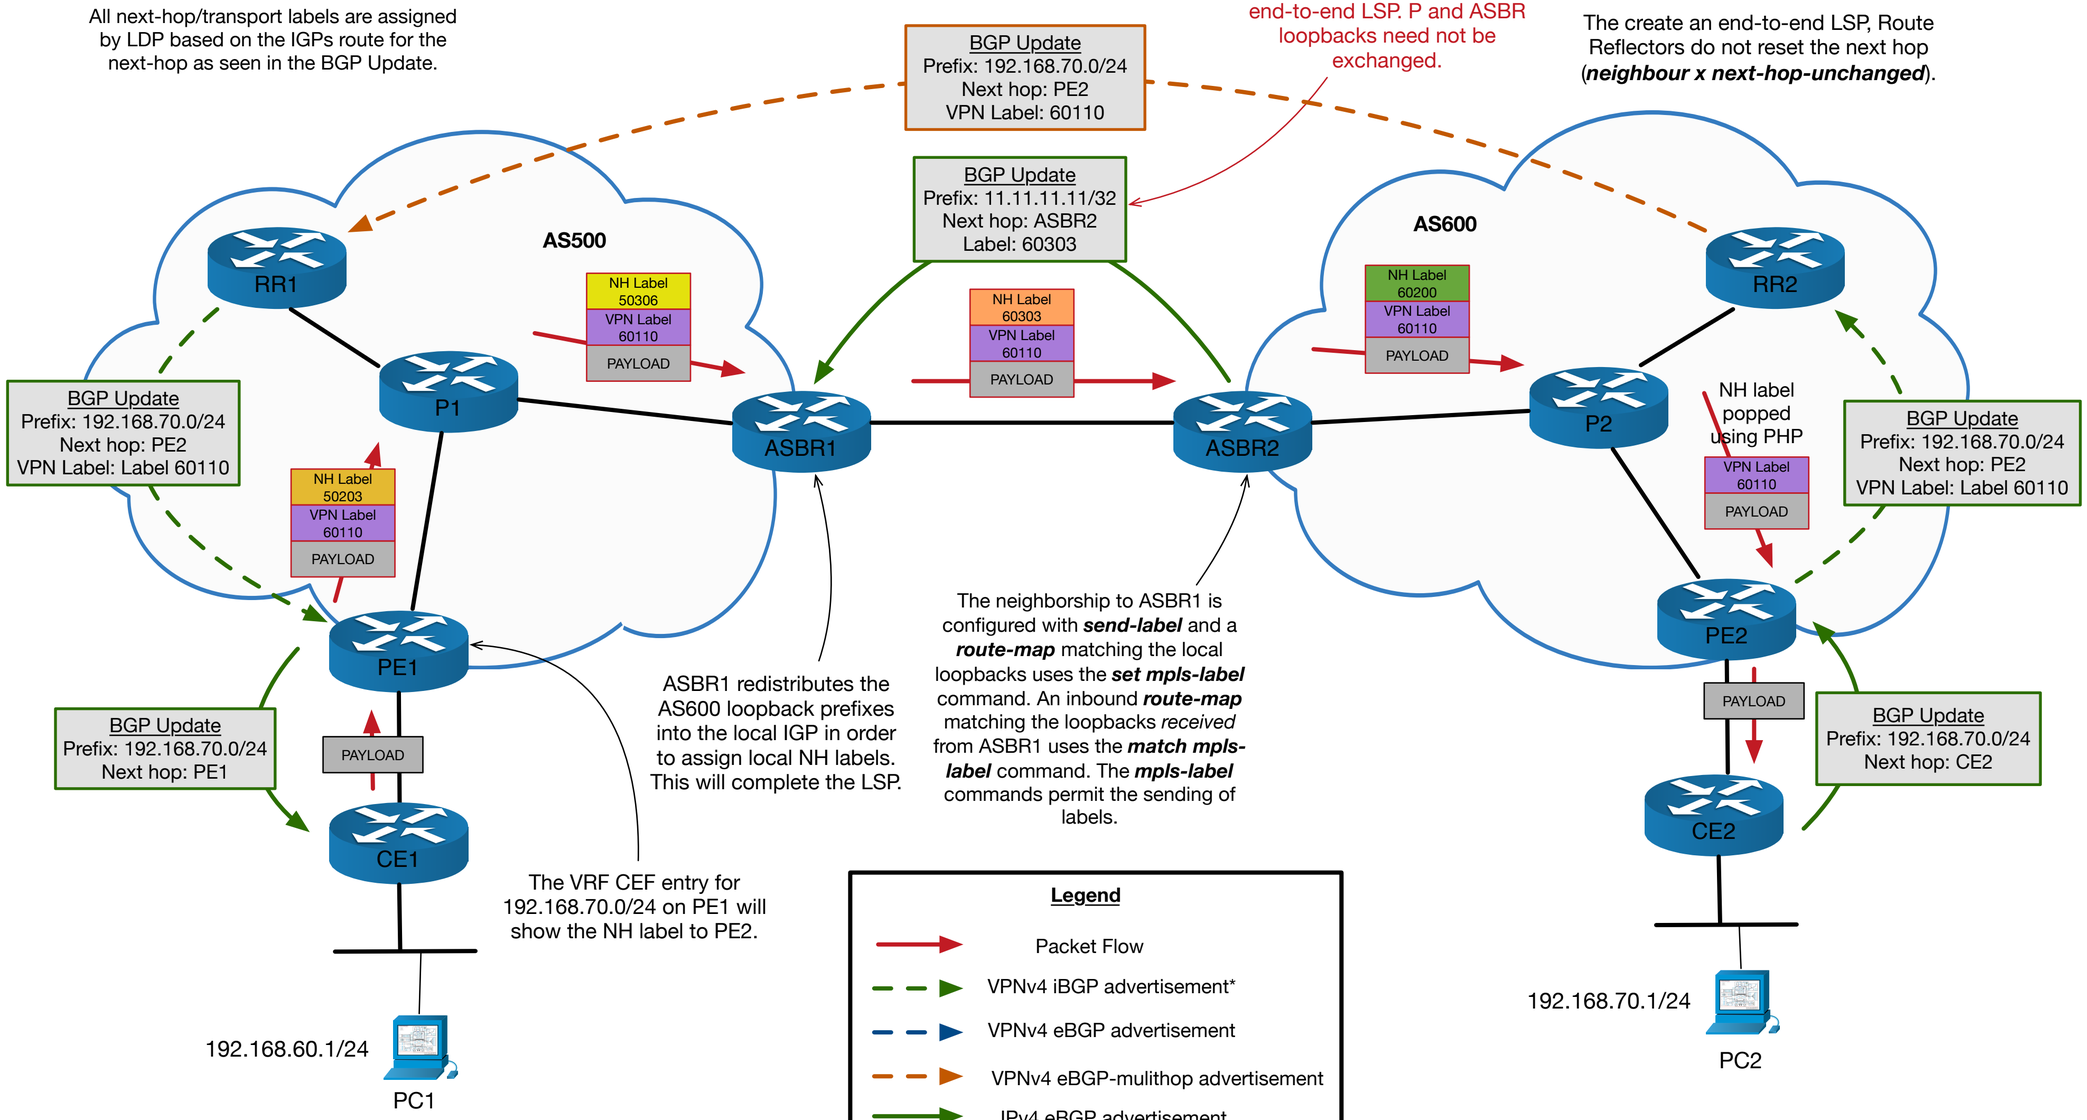

ASBR1 redistributes the AS600 loopback prefixes into the local IGP in order to assign local NH labels. This will complete the LSP.

The neighborship to ASBR1 is configured with ***send-label*** and a ***route-map*** matching the local loopbacks uses the ***set mpls-label*** command. An inbound ***route-map*** matching the loopbacks *received* from ASBR1 uses the ***match mpls-label*** command. The ***mpls-label*** commands permit the sending of labels.

The VRF CEF entry for 192.168.70.0/24 on PE1 will show the NH label to PE2.

Legend

- Packet Flow
- VPNv4 iBGP advertisement*
- VPNv4 eBGP advertisement
- VPNv4 eBGP-mulithop advertisement
- IPv4 eBGP advertisement

*possibly via a route reflector

Option C - Redistributing remote ASN next-hop prefixes into local IGP

All next-hop/transport labels are assigned by LDP based on the IGPs route for the next-hop as seen in the BGP Update.

Output

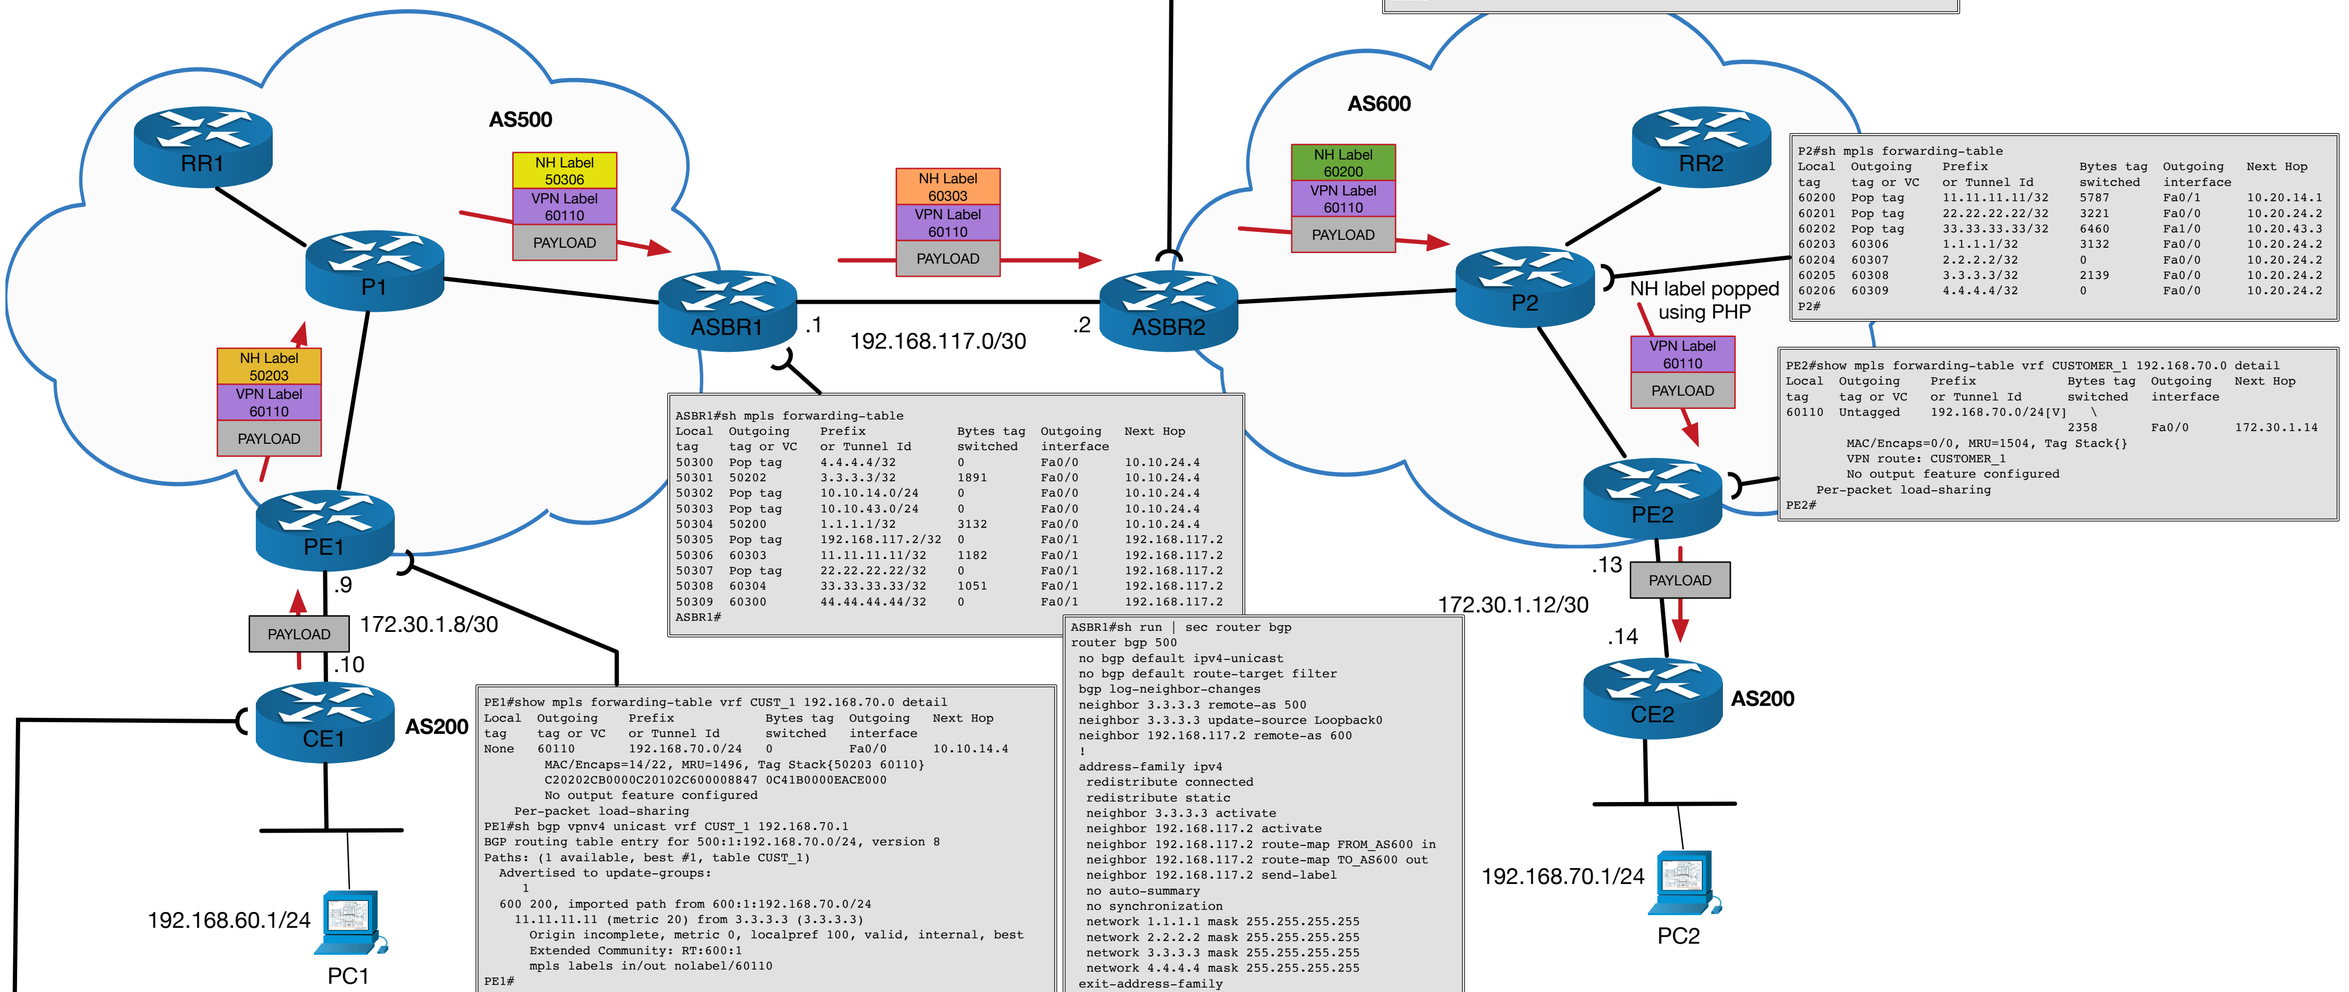


```
ASBR2#sh bgp vpnv4 unicast rd 600:1 192.168.70.1
BGP routing table entry for 600:1:192.168.70.0/24, version 3
Paths: (1 available, best #1, no table)
  Not advertised to any peer
  200
    11.11.11.11 (metric 20) from 33.33.33.33 (33.33.33.33)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:600:1
      Originator: 11.11.11.11, Cluster list: 33.33.33.33
      mpls labels in/out nolabel/60110
ASBR2#
```

```
P2#sh mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
60200  Pop tag     11.11.11.11/32    5787       Fa0/1       10.20.14.1
60201  Pop tag     22.22.22.22/32    3221       Fa0/0       10.20.24.2
60202  Pop tag     33.33.33.33/32    6460       Fa1/0       10.20.43.3
60203  60306       1.1.1.1/32        3132       Fa0/0       10.20.24.2
60204  60307       2.2.2.2/32        0          Fa0/0       10.20.24.2
60205  60308       3.3.3.3/32        2139       Fa0/0       10.20.24.2
60206  60309       4.4.4.4/32        0          Fa0/0       10.20.24.2
P2#
```

```
PE2#show mpls forwarding-table vrf CUSTOMER_1 192.168.70.0 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
60110  Untagged    192.168.70.0/24[V]    \
                                      2358       Fa0/0       172.30.1.14
        MAC/Encaps=0/0, MRU=1504, Tag Stack{}
        VPN route: CUSTOMER_1
        No output feature configured
    Per-packet load-sharing
PE2#
```

```
ASBR1#sh mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
50300  Pop tag     4.4.4.4/32        0          Fa0/0       10.10.24.4
50301  50202       3.3.3.3/32        1891       Fa0/0       10.10.24.4
50302  Pop tag     10.10.14.0/24     0          Fa0/0       10.10.24.4
50303  Pop tag     10.10.43.0/24     0          Fa0/0       10.10.24.4
50304  50200       1.1.1.1/32        3132       Fa0/0       10.10.24.4
50305  Pop tag     192.168.117.2/32  0          Fa0/1       192.168.117.2
50306  60303       11.11.11.11/32    1182       Fa0/1       192.168.117.2
50307  Pop tag     22.22.22.22/32    0          Fa0/1       192.168.117.2
50308  60304       33.33.33.33/32    1051       Fa0/1       192.168.117.2
50309  60300       44.44.44.44/32    0          Fa0/1       192.168.117.2
ASBR1#
```

```
PE1#show mpls forwarding-table vrf CUST_1 192.168.70.0 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
None   60110       192.168.70.0/24   0          Fa0/0       10.10.14.4
        MAC/Encaps=14/22, MRU=1496, Tag Stack{50203 60110}
        C20202CB0000C20102C600008847 0C41B0000EACE000
        No output feature configured
    Per-packet load-sharing
PE1#sh bgp vpnv4 unicast vrf CUST_1 192.168.70.1
BGP routing table entry for 500:1:192.168.70.0/24, version 8
Paths: (1 available, best #1, table CUST_1)
  Advertised to update-groups:
     1
  600 200, imported path from 600:1:192.168.70.0/24
    11.11.11.11 (metric 20) from 3.3.3.3 (3.3.3.3)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:600:1
      mpls labels in/out nolabel/60110
PE1#
```

```
ASBR1#sh run | sec router bgp
router bgp 500
 no bgp default ipv4-unicast
 no bgp default route-target filter
 bgp log-neighbor-changes
 neighbor 3.3.3.3 remote-as 500
 neighbor 3.3.3.3 update-source Loopback0
 neighbor 192.168.117.2 remote-as 600
 !
 address-family ipv4
  redistribute connected
  redistribute static
  neighbor 3.3.3.3 activate
  neighbor 192.168.117.2 activate
  neighbor 192.168.117.2 route-map FROM_AS600 in
  neighbor 192.168.117.2 route-map TO_AS600 out
  neighbor 192.168.117.2 send-label
  no auto-summary
  no synchronization
  network 1.1.1.1 mask 255.255.255.255
  network 2.2.2.2 mask 255.255.255.255
  network 3.3.3.3 mask 255.255.255.255
  network 4.4.4.4 mask 255.255.255.255
 exit-address-family
 !
 address-family vpnv4
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 send-community extended
  neighbor 3.3.3.3 next-hop-self
 exit-address-family
```

```
CE1#traceroute 192.168.70.1 source lo1

Type escape sequence to abort.
Tracing the route to 192.168.70.1

  1 172.30.1.9 16 msec 12 msec 12 msec
  2 10.10.14.4 [AS 500] [MPLS: Labels 50203/60110 Exp 0] 92 msec 104 msec 92 msec
  3 10.10.24.2 [AS 500] [MPLS: Labels 50306/60110 Exp 0] 120 msec 124 msec 56 msec
  4 192.168.117.2 [AS 500] [MPLS: Labels 60303/60110 Exp 0] 92 msec 116 msec 108 msec
  5 10.20.24.4 [AS 500] [MPLS: Labels 60200/60110 Exp 0] 124 msec 100 msec 72 msec
  6 172.30.1.13 [AS 600] [MPLS: Label 60110 Exp 0] 120 msec 56 msec 68 msec
  7 172.30.1.14 [AS 600] 72 msec 140 msec 100 msec
```

ISP Loopback IP Addressing	VRF Information
PE1: 1.1.1.1/32	In ASN 500
ASBR1: 2.2.2.2/32	Name: CUST_1
RR1: 3.3.3.3/32	RD: 500:1
P1: 4.4.4.4/32	
PE2: 11.11.11.11/32	In ASN 600
ASB2: 22.22.22.22/32	Name: CUSTOMER_1
RR2: 33.33.33.33/32	RD: 600:1
P2: 44.44.44.44/32	

Option C - Advertise remote ASN next-hop prefixes using iBGP

Concept

All next-hop/transport labels are assigned by LDP based on the IGPs route for the next-hop as seen in the BGP Update.

The create an end-to-end LSP, Route Reflectors do not reset the next hop (***neighbour x next-hop-unchanged***).

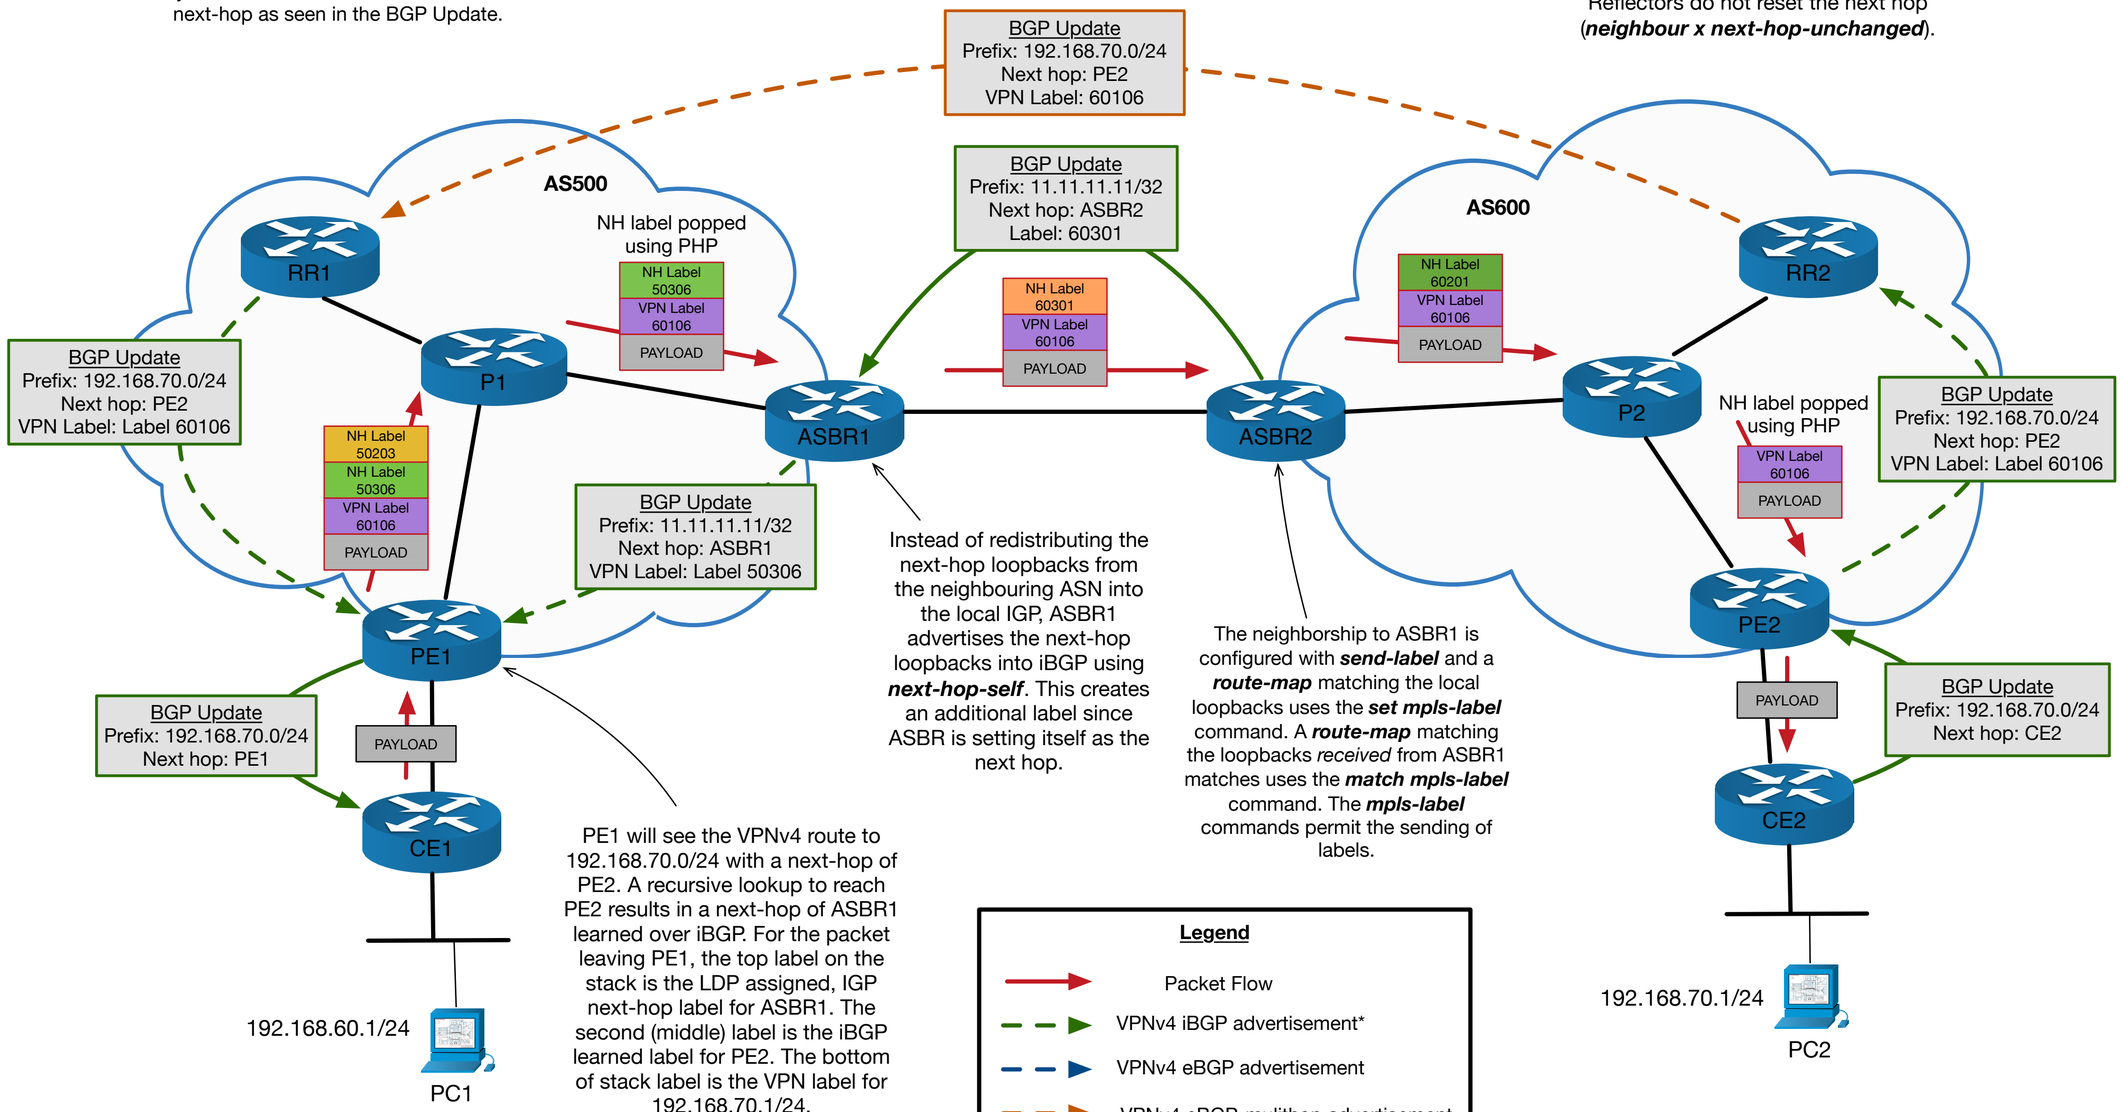

Instead of redistributing the next-hop loopbacks from the neighbouring ASN into the local IGP, ASBR1 advertises the next-hop loopbacks into iBGP using ***next-hop-self***. This creates an additional label since ASBR is setting itself as the next hop.

The neighborship to ASBR1 is configured with ***send-label*** and a ***route-map*** matching the local loopbacks uses the ***set mpls-label*** command. A ***route-map*** matching the loopbacks *received* from ASBR1 matches uses the ***match mpls-label*** command. The ***mpls-label*** commands permit the sending of labels.

PE1 will see the VPNv4 route to 192.168.70.0/24 with a next-hop of PE2. A recursive lookup to reach PE2 results in a next-hop of ASBR1 learned over iBGP. For the packet leaving PE1, the top label on the stack is the LDP assigned, IGP next-hop label for ASBR1. The second (middle) label is the iBGP learned label for PE2. The bottom of stack label is the VPN label for 192.168.70.1/24.

Legend

- Packet Flow
- VPNv4 iBGP advertisement*
- VPNv4 eBGP advertisement
- VPNv4 eBGP-mulithop advertisement
- IPv4 eBGP advertisement

*possibly via a route reflector

Option C - Advertise remote ASN next-hop prefixes using iBGP

Output

All next-hop/transport labels are assigned by LDP based on the IGPs route for the next-hop as seen in the BGP Update.

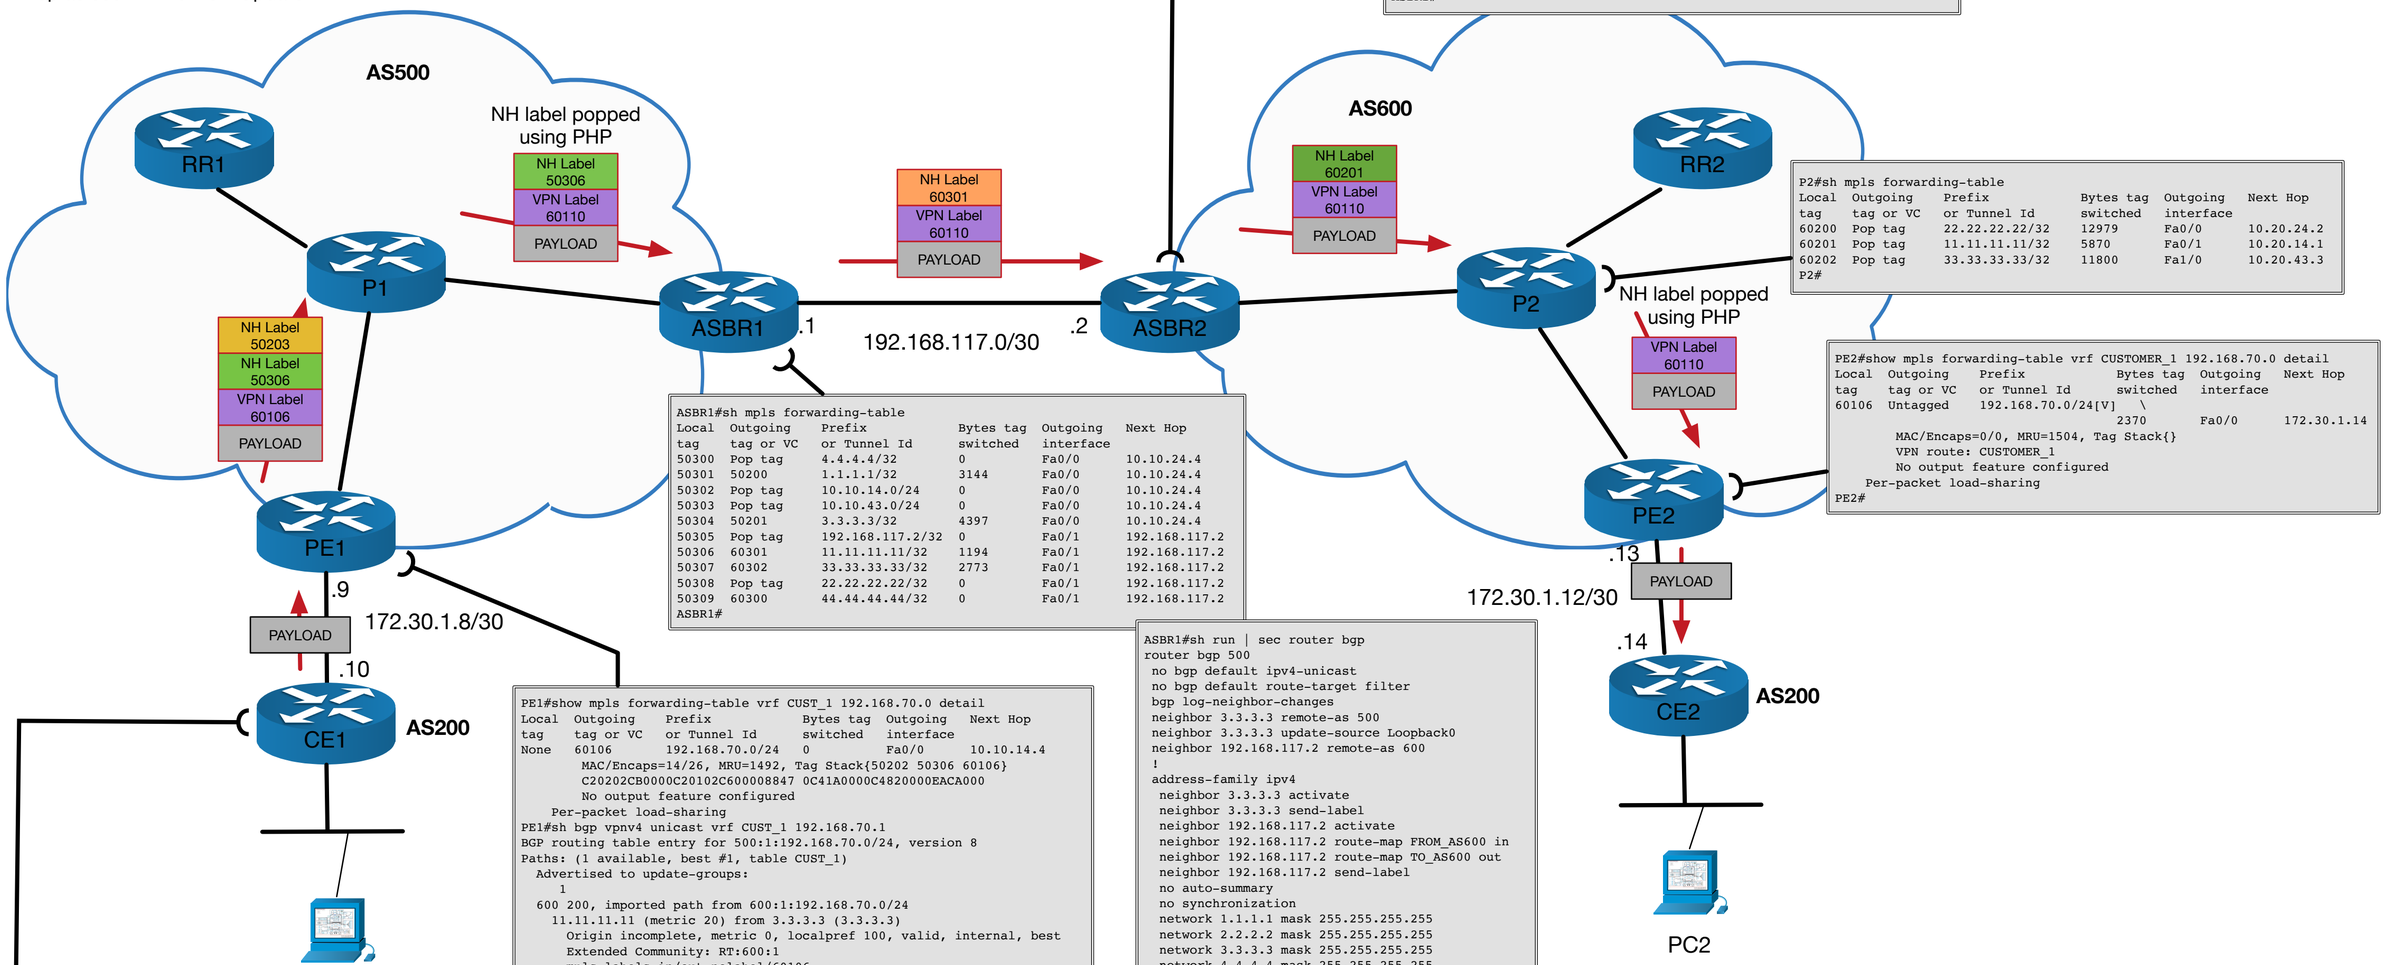


```
ASBR2#sh bgp vpnv4 unicast rd 600:1 192.168.70.1
BGP routing table entry for 600:1:192.168.70.0/24, version 5
Paths: (1 available, best #1, no table)
  Not advertised to any peer
  200
    11.11.11.11 (metric 20) from 33.33.33.33 (33.33.33.33)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:600:1
      Originator: 11.11.11.11, Cluster list: 33.33.33.33
      mpls labels in/out nolabel/60106
ASBR2#
```

```
P2#sh mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
60200  Pop tag     22.22.22.22/32    12979      Fa0/0       10.20.24.2
60201  Pop tag     11.11.11.11/32    5870       Fa0/1       10.20.14.1
60202  Pop tag     33.33.33.33/32    11800      Fa1/0       10.20.43.3
P2#
```

```
ASBR1#sh mpls forwarding-table
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
50300  Pop tag     4.4.4.4/32        0          Fa0/0       10.10.24.4
50301  50200       1.1.1.1/32        3144       Fa0/0       10.10.24.4
50302  Pop tag     10.10.14.0/24     0          Fa0/0       10.10.24.4
50303  Pop tag     10.10.43.0/24     0          Fa0/0       10.10.24.4
50304  50201       3.3.3.3/32        4397       Fa0/0       10.10.24.4
50305  Pop tag     192.168.117.2/32  0          Fa0/1       192.168.117.2
50306  60301       11.11.11.11/32    1194       Fa0/1       192.168.117.2
50307  60302       33.33.33.33/32    2773       Fa0/1       192.168.117.2
50308  Pop tag     22.22.22.22/32    0          Fa0/1       192.168.117.2
50309  60300       44.44.44.44/32    0          Fa0/1       192.168.117.2
ASBR1#
```

```
PE2#show mpls forwarding-table vrf CUSTOMER_1 192.168.70.0 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
60106  Untagged    192.168.70.0/24[V]   \
                                        2370       Fa0/0       172.30.1.14
       MAC/Encaps=0/0, MRU=1504, Tag Stack{}
       VPN route: CUSTOMER_1
       No output feature configured
   Per-packet load-sharing
PE2#
```

```
PE1#show mpls forwarding-table vrf CUST_1 192.168.70.0 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
None   60106       192.168.70.0/24   0          Fa0/0       10.10.14.4
       MAC/Encaps=14/26, MRU=1492, Tag Stack{50202 50306 60106}
       C20202CB0000C20102C600008847 0C41A0000C4820000EACA000
       No output feature configured
   Per-packet load-sharing
PE1#sh bgp vpnv4 unicast vrf CUST_1 192.168.70.1
BGP routing table entry for 500:1:192.168.70.0/24, version 8
Paths: (1 available, best #1, table CUST_1)
  Advertised to update-groups:
     1
  600 200, imported path from 600:1:192.168.70.0/24
    11.11.11.11 (metric 20) from 3.3.3.3 (3.3.3.3)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:600:1
      mpls labels in/out nolabel/60106
PE1#
```

```
ASBR1#sh run | sec router bgp
router bgp 500
 no bgp default ipv4-unicast
 no bgp default route-target filter
 bgp log-neighbor-changes
 neighbor 3.3.3.3 remote-as 500
 neighbor 3.3.3.3 update-source Loopback0
 neighbor 192.168.117.2 remote-as 600
 !
 address-family ipv4
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 send-label
  neighbor 192.168.117.2 activate
  neighbor 192.168.117.2 route-map FROM_AS600 in
  neighbor 192.168.117.2 route-map TO_AS600 out
  neighbor 192.168.117.2 send-label
  no auto-summary
  no synchronization
  network 1.1.1.1 mask 255.255.255.255
  network 2.2.2.2 mask 255.255.255.255
  network 3.3.3.3 mask 255.255.255.255
  network 4.4.4.4 mask 255.255.255.255
 exit-address-family
 !
 address-family vpnv4
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 send-community extended
  neighbor 3.3.3.3 next-hop-self
 exit-address-family
```

```
CE1#trace 192.168.70.1 source lo1

Type escape sequence to abort.
Tracing the route to 192.168.70.1

  1 172.30.1.9 8 msec 24 msec 12 msec
  2 10.10.14.4 [AS 500] [MPLS: Labels 50202/50306/60106 Exp 0] 124 msec 120 msec 96 msec
  3 10.10.24.2 [AS 500] [MPLS: Labels 50306/60106 Exp 0] 116 msec 64 msec 80 msec
  4 192.168.117.2 [AS 500] [MPLS: Labels 60301/60106 Exp 0] 120 msec 120 msec 96 msec
  5 10.20.24.4 [AS 500] [MPLS: Labels 60201/60106 Exp 0] 104 msec 100 msec 96 msec
  6 172.30.1.13 [AS 600] [MPLS: Label 60106 Exp 0] 72 msec 72 msec 64 msec
  7 172.30.1.14 [AS 600] 108 msec 76 msec 88 msec
```

ISP Loopback IP Addressing	VRF Information
PE1: 1.1.1.1/32	In ASN 500
ASBR1: 2.2.2.2/32	Name: CUST_1
RR1: 3.3.3.3/32	RD: 500:1
P1: 4.4.4.4/32	
PE2: 11.11.11.11/32	In ASN 600
ASB2: 22.22.22.22/32	Name: CUSTOMER_1
RR2: 33.33.33.33/32	RD: 600:1
P2: 44.44.44.44/32	