Inter-AS MPLS VPNs

VPLS and L3VPN using Carrier Supporting Carrier (CSC)

Topology:

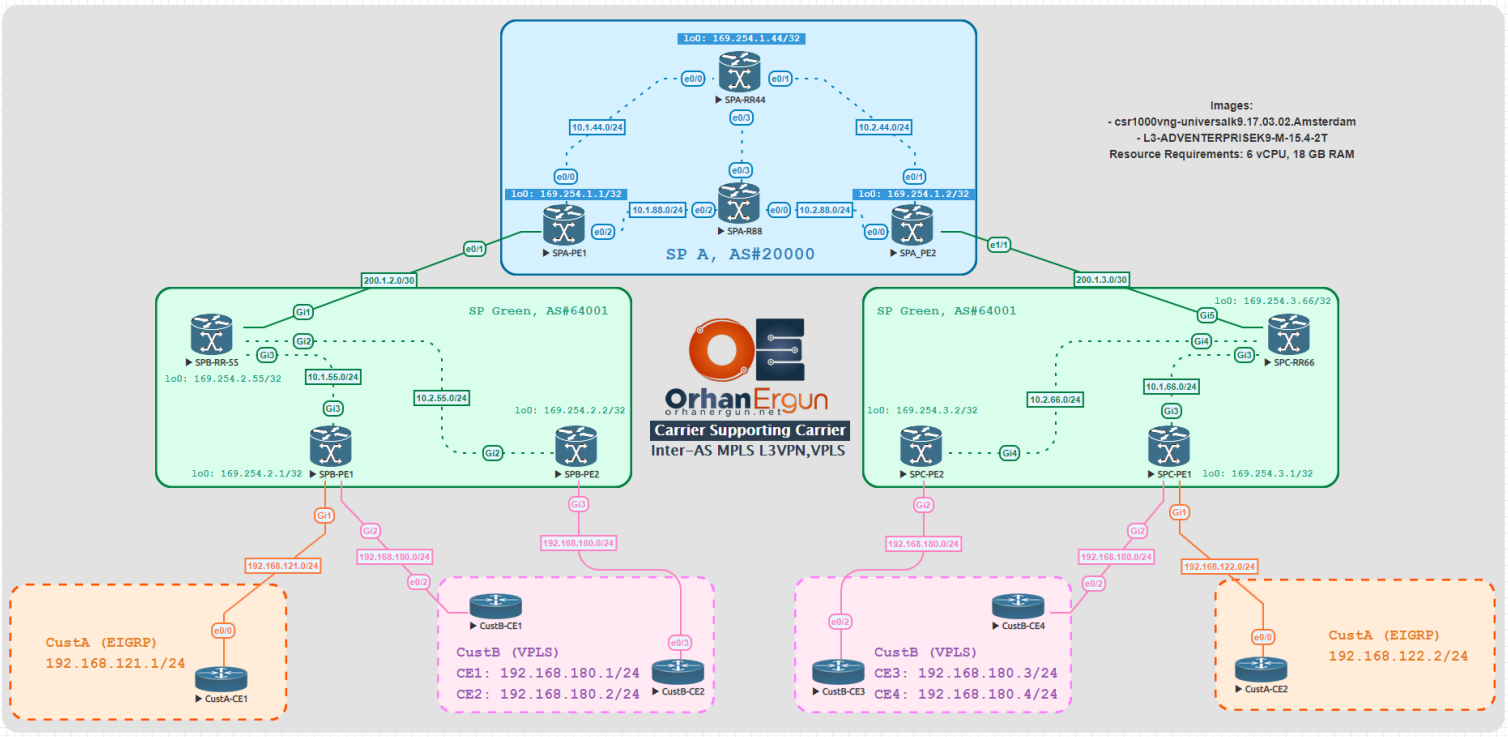

MP-BGP AF/SAF: VPNv4 Unicast, L2VPN VPLS

PE-CE Routing Protocols:

- VPLS (Martini with BGP Auto-Discovery)
- L3VPN PE-CE routing protocol: EIGRP

Task 01:

- Provide L3VPN service for the CustA (Use EIGRP and the PE-CE Routing Protocol)
- Provide VPLS service for CustB (in SP side use Auto Discovery and LDP as the signaling protocol between PE devices)

Solution:

Please take a look at the topology, we have a service provider connected to the Customer Edge devices (SP Green), this service provider has two routing domain which is in separate from each other.

They are providing L3VPN and L2VPN (VPLS) service to the Customers.

The basic configuration is easy ans straight forward.

Just like MPLS Lab 01 and MPLS Lab 02 we will follow the same steps in each SP domain to configure the L3VPN and VPLS configuration, if you need more explanation about how to configure them, please refer to those two labs.

```
SPB-PE1:
vrf definition CustA
 rd 64012:1
 route-target export 169.254.2.1:1
 route-target import 169.254.3.1:1
 !
 address-family ipv4
 exit-address-family
!
!
interface GigabitEthernet1
 vrf forwarding CustA
 ip address 192.168.121.254 255.255.255.0
 negotiation auto
 no mop enabled
 no mop sysid
!
router eigrp CustA
!
```

```
address-family ipv4 unicast vrf CustA autonomous-system 1
  !
  topology base
   redistribute bgp 64001
  exit-af-topology
  network 192.168.121.0
 exit-address-family
!
router bgp 64001
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 169.254.2.55 remote-as 64001
 !
!
 address-family vpnv4
  neighbor 169.254.2.55 activate
  neighbor 169.254.2.55 send-community extended
 exit-address-family
 !
 address-family l2vpn vpls
  neighbor 169.254.2.55 activate
  neighbor 169.254.2.55 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf CustA
  redistribute eigrp 1
 exit-address-family
interface GigabitEthernet2
 no shutdown
 service instance 2 ethernet
  encapsulation untagged
  l2protocol forward cdp
 !
```

```
l2vpn vfi context CustB-VPLS
 vpn id 2
 autodiscovery bgp signaling ldp
  vpls-id 64012:2
  rd 64012:2
  route-target export 169.254.2.1:2
  route-target import 169.254.3.1:2
  route-target import 169.254.3.2:2
  route-target export 169.254.2.2:2
!
bridge-domain 2
 member GigabitEthernet2 service-instance 2
 member vfi CustB-VPLS
!
```

This time, we are using Martini VPLS, LDP is going to be used as the signaling protocol in order to form pseudowires and distribute labels, LDP itself does not have any Auto-Discovery, so the original Martini does not have any Auto-Discovery mechanism in order to provide PE to PE discovery. Nowadays thanks to the MP-BGP we can have Auto-Discovery feature.

Take a look at the command: `autodiscovery bgp signaling ldp`

- The first part implies: Auto-Discovery is BGP
- The second part implies: Signaling is LDP

Let's configure the rest of the other PE devices as well:

SPB-PE2:

```
router bgp 64001
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 169.254.2.55 remote-as 64001
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 169.254.2.55 activate
```

```
  neighbor 169.254.2.55 send-community extended
 exit-address-family
 !
 address-family l2vpn vpls
  neighbor 169.254.2.55 activate
  neighbor 169.254.2.55 send-community extended
 exit-address-family
l2vpn vfi context CustB-VPLS
 vpn id 2
 autodiscovery bgp signaling ldp
  vpls-id 64012:2
  rd 64012:2
  route-target import 169.254.2.1:2
  route-target export 169.254.2.2:2
  route-target import 169.254.3.1:2
  route-target import 169.254.3.2:2
!
interface GigabitEthernet3
 no shutdown
 service instance 2 ethernet
  encapsulation untagged
  l2protocol forward cdp
 !
!
bridge-domain 2
 member GigabitEthernet3 service-instance 2
 member vfi CustB-VPLS
!
```

We could use the Same RT import and export values, that is the recommended way for simplicity and ease of management and troubleshooting. But this is just a lab invironment and a training workbook, so we are free to play with RT values to understand what they are actually doing!

If you don't like to play with them just skip that step and use the same RT import/export values.

SPC-PE2:

```
router bgp 64001
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 169.254.3.66 remote-as 64001
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 169.254.3.66 activate
  neighbor 169.254.3.66 send-community extended
 exit-address-family
 !
 address-family l2vpn vpls
  neighbor 169.254.3.66 activate
  neighbor 169.254.3.66 send-community extended
 exit-address-family
!
interface GigabitEthernet2
 service instance 2 ethernet
  encapsulation untagged
  l2protocol forward cdp
 !
!
l2vpn vfi context CustB-VPLS
 vpn id 2
 autodiscovery bgp signaling ldp
  vpls-id 64012:2
  rd 64012:2
  route-target import 169.254.2.1:2
  route-target import 169.254.2.2:2
  route-target import 169.254.3.1:2
  route-target export 169.254.3.2:2
!
```

```
bridge-domain 2
 member GigabitEthernet2 service-instance 2
 member vfi CustB-VPLS
!
```

SPC-PE1:

```
router bgp 64001
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 169.254.3.66 remote-as 64001
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 169.254.3.66 activate
  neighbor 169.254.3.66 send-community extended
 exit-address-family
 !
 address-family l2vpn vpls
  neighbor 169.254.3.66 activate
  neighbor 169.254.3.66 send-community extended
 exit-address-family
 !
vrf definition CustA
 rd 64012:1
 route-target export 169.254.3.1:1
 route-target import 169.254.2.1:1
 !
 address-family ipv4
 exit-address-family
!
!
interface GigabitEthernet1
 vrf forwarding CustA
```

```
 ip address 192.168.122.254 255.255.255.0
 negotiation auto
 no mop enabled
 no mop sysid
!
router eigrp CustA
!
address-family ipv4 unicast vrf CustA autonomous-system 1
  !
  topology base
   redistribute bgp 64001
  exit-af-topology
  network 192.168.122.0
 exit-address-family
!
router bgp 64001
 !
 address-family ipv4 vrf CustA
  redistribute eigrp 1
 exit-address-family
!
interface GigabitEthernet2
 no shutdown
 service instance 2 ethernet
  encapsulation untagged
  l2protocol forward cdp
 !
!
l2vpn vfi context CustB-VPLS
 vpn id 2
 autodiscovery bgp signaling ldp
  vpls-id 64012:2
  rd 64012:2
  route-target import 169.254.2.1:2
```

```
  route-target import 169.254.2.2:2
  route-target export 169.254.3.1:2
  route-target import 169.254.3.2:2
!
bridge-domain 2
 member GigabitEthernet2 service-instance 2
 member vfi CustB-VPLS
!
```

All Green SP PE devices configuration are done.

Before we continue configuring Green SP Route-Reflector configuration let's start the task two and configure The SP A (AS 20000) instead.

Task 02:

- Customer sites connected to both Customer Carrier networks (Green SP) should have reachability (Use Carrier Supporting Carrier)

Solution:

If you are already familiar and comfortable with L3VPNs configuration, this step is going to be very easy to understand, if you are not good at L3VPN configuration please refer to the MPLS Lab 01 and complete all the tasks, then you can come back and configure this task!

The Blue SP (SP A, AS#20000) is going to be configured as a top level carrier for the SP Green domains (low level carrier that is connected to the customer, we call it customer carrier).

The Blue SP needs to provide L3VPN service to the SP Green. That simple!

We will configure a VRF for the Customer carrier (SP Green), we will enable VPNv4 unicast neghborship inside Blue SP between PEs and the Route-Reflector:

```
SPA-RR44:
router bgp 20000
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 169.254.1.1 remote-as 20000
 neighbor 169.254.1.1 update-source Loopback0
 neighbor 169.254.1.2 remote-as 20000
 neighbor 169.254.1.2 update-source Loopback0
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 169.254.1.1 activate
  neighbor 169.254.1.1 send-community extended
  neighbor 169.254.1.1 route-reflector-client
  neighbor 169.254.1.2 activate
  neighbor 169.254.1.2 send-community extended
  neighbor 169.254.1.2 route-reflector-client
 exit-address-family
```

SPA-PE1:

```
vrf definition SPB
 rd 20000:23
 route-target export 169.254.1.1:23
 route-target import 169.254.1.2:23
 !
 address-family ipv4
 exit-address-family
!
interface Ethernet0/1
 vrf forwarding SPB
 ip address 200.1.2.1 255.255.255.252
 mpls bgp forwarding
!
router bgp 20000
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 169.254.1.44 remote-as 20000
 neighbor 169.254.1.44 update-source Loopback0
 neighbor 200.1.2.2 remote-as 64001
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 169.254.1.44 activate
  neighbor 169.254.1.44 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf SPB
  neighbor 200.1.2.2 remote-as 64001
  neighbor 200.1.2.2 activate
  neighbor 200.1.2.2 as-override
  neighbor 200.1.2.2 send-label
 exit-address-family
```

SPA-PE1 is providing an L3VPN service to the Customer carrier (SP Green) and it is forming eBGP session with the SPB-RR-55.

But there is an additional step needed in this situation compared to the simple L3VPN service. That is enabling Labeled Unicast between SPA-PE1 and SPB-RR-55, because label-switching needed between those two routers.

SPA-PE1 is going to take the routes from SPB-RR-55 (RR and PE devices Loopback IP addresses) and advertise them to the SPA-PE2, so SPA-PE2 will advertise them t the SPC-RR-66.

Just like Inter-AS MPLS VPNs using Option C.

NOTE: If you are not familiar with Option A, B and C, please go back and complete those labsfirst! Otherwise you cannot understand what is happening in CSC.

The high level overview:

- SPB-RR-55 is an eBGP IPv4 Unicast (Labeled unicast) neighbor with SPA-PE1
- SPB-RR-55 is going to advertise it's own loopback IP address along with the othe two PE Loopback IP addresses to the SPA-PE1
- SPA-PE1 takes those routes and advertises them to the SPA-RR-44
- SPA-RR-44 advertises those routes to the SPA-PE2
- SPA-PE2 advertises the routes to the SPC-RR-66

Now, two RRs and even 4 SP Green PE devices have end to end (Loopback to Loopback) LSPs to each other.

RRs (Green SP RRs) can form VPNv4 Unicast with each other (Just like Option C).

Let's configure the remaining devices:

```
SPA-PE2:
vrf definition SPB
 rd 20000:23
 route-target export 169.254.1.2:23
 route-target import 169.254.1.1:23
 !
 address-family ipv4
 exit-address-family
!
!
interface Ethernet1/1
```

```
 vrf forwarding SPB
 ip address 200.1.3.1 255.255.255.252
 mpls bgp forwarding
!
router bgp 20000
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 169.254.1.44 remote-as 20000
 neighbor 169.254.1.44 update-source Loopback0
 neighbor 200.1.3.2 remote-as 64002
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  neighbor 169.254.1.44 activate
  neighbor 169.254.1.44 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf SPB
  neighbor 200.1.3.2 remote-as 64001
  neighbor 200.1.3.2 activate
  neighbor 200.1.3.2 as-override
  neighbor 200.1.3.2 send-label
 exit-address-family
```

Why we even need as-override command? It is because of the same AS number of SP Green.

The Customer Carrier in both domains is using the same AS#64001, the as-override command is needed to replace the Carrier Carrier AS# with the Customer Carrier AS# when it is advertised to the Customer Carrier CE devices. We discussed about this stuff previously on MPLS Lab 01, you can refer to that lab for more information.

```
SPB-RR-55:
router bgp 64001
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 169.254.2.1 remote-as 64001
```

```
neighbor 169.254.2.1 update-source Loopback0
neighbor 169.254.2.2 remote-as 64001
neighbor 169.254.2.2 update-source Loopback0
neighbor 169.254.3.66 remote-as 64001
neighbor 169.254.3.66 ebgp-multihop 255
neighbor 169.254.3.66 update-source Loopback0
neighbor 200.1.2.1 remote-as 20000
!
address-family ipv4
 network 169.254.2.1 mask 255.255.255.255
 network 169.254.2.2 mask 255.255.255.255
 network 169.254.2.55 mask 255.255.255.255
 neighbor 200.1.2.1 activate
 neighbor 200.1.2.1 send-label
exit-address-family
!
address-family vpnv4
 neighbor 169.254.2.1 activate
 neighbor 169.254.2.1 send-community extended
 neighbor 169.254.2.1 route-reflector-client
 neighbor 169.254.2.2 activate
 neighbor 169.254.2.2 send-community extended
 neighbor 169.254.2.2 route-reflector-client
 neighbor 169.254.3.66 activate
 neighbor 169.254.3.66 send-community extended
exit-address-family
!
address-family l2vpn vpls
 neighbor 169.254.2.1 activate
 neighbor 169.254.2.1 send-community extended
 neighbor 169.254.2.1 route-reflector-client
 neighbor 169.254.2.2 activate
 neighbor 169.254.2.2 send-community extended
 neighbor 169.254.2.2 route-reflector-client
```

```
  neighbor 169.254.3.66 activate
  neighbor 169.254.3.66 send-community extended
 exit-address-family
!
interface GigabitEthernet1
 ip address 200.1.2.2 255.255.255.252
 negotiation auto
 mpls bgp forwarding
!
```

SPC-RR-66:

```
interface GigabitEthernet5
 ip address 200.1.3.2 255.255.255.252
 mpls bgp forwarding
!
router bgp 64001
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 169.254.2.55 remote-as 64001
 neighbor 169.254.2.55 ebgp-multihop 255
 neighbor 169.254.2.55 update-source Loopback0
 neighbor 169.254.3.1 remote-as 64001
 neighbor 169.254.3.1 update-source Loopback0
 neighbor 169.254.3.2 remote-as 64001
 neighbor 169.254.3.2 update-source Loopback0
 neighbor 200.1.2.1 remote-as 20000
 neighbor 200.1.3.1 remote-as 20000
 !
 address-family ipv4
  network 169.254.2.55 mask 255.255.255.255
  network 169.254.3.1 mask 255.255.255.255
  network 169.254.3.2 mask 255.255.255.255
  network 169.254.3.66 mask 255.255.255.255
  neighbor 200.1.2.1 activate
```

```
 neighbor 200.1.3.1 activate
 neighbor 200.1.3.1 send-label
exit-address-family
!
address-family vpnv4
 neighbor 169.254.2.55 activate
 neighbor 169.254.2.55 send-community extended
 neighbor 169.254.3.1 activate
 neighbor 169.254.3.1 send-community extended
 neighbor 169.254.3.1 route-reflector-client
 neighbor 169.254.3.2 activate
 neighbor 169.254.3.2 send-community extended
 neighbor 169.254.3.2 route-reflector-client
exit-address-family
!
address-family l2vpn vpls
 neighbor 169.254.2.55 activate
 neighbor 169.254.2.55 send-community extended
 neighbor 169.254.3.1 activate
 neighbor 169.254.3.1 send-community extended
 neighbor 169.254.3.1 route-reflector-client
 neighbor 169.254.3.2 activate
 neighbor 169.254.3.2 send-community extended
 neighbor 169.254.3.2 route-reflector-client
exit-address-family
```

The configuration is done.

All RRs and PE devices in Customer Carrier domains (SP Green) have loopback to loopback LSP (Label Switched Path) reachability. Also they are not aware that is something happening in the underlay! SP A takes care of them transparently.

Verification:

```
SPB-RR-55#show bgp vpnv4 uni all summary | begin Neigh
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
169.254.2.1     4        64001     106     111        5    0    0 01:30:46        2
169.254.2.2     4        64001     105     111        5    0    0 01:30:44        0
169.254.3.66    4        64001     107     107        5    0    0 01:29:58        2

SPB-RR-55#show bgp l2vpn vpls all summary | begin Ne
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
169.254.2.1     4        64001     107     112        5    0    0 01:31:16        1
169.254.2.2     4        64001     105     112        5    0    0 01:31:15        1
169.254.3.66    4        64001     108     108        5    0    0 01:30:28        2

SPB-PE1#show bgp l2vpn vpls all summary | begin Neigh
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
169.254.2.55    4        64001     112     107        5    0    0 01:31:38        3



SPB-PE1#show bgp vpnv4 uni all summary | begin Neigh
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
169.254.2.55    4        64001     113     108        7    0    0 01:32:20        2



SPC-PE1#show bgp l2vpn vpls all summary | begin Neigh
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
169.254.3.66    4        64001     113     107        5    0    0 01:32:12        3



SPB-PE1#ping 169.254.3.1 source lo 0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 169.254.3.1, timeout is 2 seconds:
Packet sent with a source address of 169.254.2.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/3 ms
```

```
SPB-PE1#traceroute 169.254.3.1 source lo 0 probe 1
Type escape sequence to abort.
Tracing the route to 169.254.3.1
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.55.55 [MPLS: Label 645003 Exp 0] 12 msec
  2 200.1.2.1 [MPLS: Label 21006 Exp 0] 2 msec
  3 10.1.88.88 [MPLS: Labels 16/22006 Exp 0] 2 msec
  4 200.1.3.1 [MPLS: Label 22006 Exp 0] 1 msec
  5 200.1.3.2 [MPLS: Label 642601 Exp 0] 2 msec
  6 10.1.66.1 3 msec

CustA-CE1#show ip route eigrp | begin Gate
Gateway of last resort is not set

D     192.168.122.0/24 [90/1029120] via 192.168.121.254, 01:32:34, Ethernet0/0
      192.168.255.0/32 is subnetted, 2 subnets
D        192.168.255.2 [90/1029760] via 192.168.121.254, 01:32:34, Ethernet0/0

CustA-CE1#trace 192.168.122.2 source lo 0 probe 1
Type escape sequence to abort.
Tracing the route to 192.168.122.2
VRF info: (vrf in name/id, vrf out name/id)
  1 192.168.121.254 5 msec
  2 10.1.55.55 [MPLS: Labels 645003/642105 Exp 0] 7 msec
  3 200.1.2.1 [MPLS: Labels 21006/642105 Exp 0] 2 msec
  4 10.1.44.44 [MPLS: Labels 24001/22006/642105 Exp 0] 2 msec
  5 10.2.44.2 [MPLS: Labels 22006/642105 Exp 0] 2 msec
  6 200.1.3.2 [MPLS: Labels 642601/642105 Exp 0] 2 msec
  7 192.168.122.254 3 msec
  8 192.168.122.2 3 msec
```

Customer packets are label switched end to end (From SPB-PE1 to the SPC-PE1) and there are 3 labels in the label stack at the middle of the path (Step 4) (when packet is traversing SP A).

An interresting thing about this type of implementation is the VPN label (642105) stays the same along the path until it reaches to the egress PE.

```
CustB-CE1#show ip int br | ex unas
Interface                  IP-Address      OK? Method Status                Protocol
Ethernet0/2                192.168.180.1   YES NVRAM  up                    up

CustB-CE1#ping 192.168.180.255 re 1
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 192.168.180.255, timeout is 2 seconds:

Reply to request 0 from 192.168.180.2, 1 ms
Reply to request 0 from 192.168.180.4, 2 ms
Reply to request 0 from 192.168.180.3, 2 ms

SPB-PE1#show bridge-domain
Bridge-domain 2 (4 ports in all)
State: UP                    Mac learning: Enabled
Aging-Timer: 300 second(s)
Maximum address limit: 65536
    GigabitEthernet2 service instance 2
    vfi CustB-VPLS neighbor 169.254.2.2 2
    vfi CustB-VPLS neighbor 169.254.3.1 2
    vfi CustB-VPLS neighbor 169.254.3.2 2
   AED MAC address    Policy  Tag       Age  Pseudoport
   0   AABB.CC00.0C20 forward dynamic   293  GigabitEthernet2.EFP2
   0   AABB.CC00.1020 forward dynamic   274  CustB-VPLS.404012
   0   AABB.CC00.0B30 forward dynamic   294  CustB-VPLS.404011
   0   AABB.CC00.0F20 forward dynamic   274  CustB-VPLS.404013
```

VPLS service also works fine! We have the MAC addresses in the bridge-domain output and BUM traffic is being sent from a CE device to all of the other CE devices.