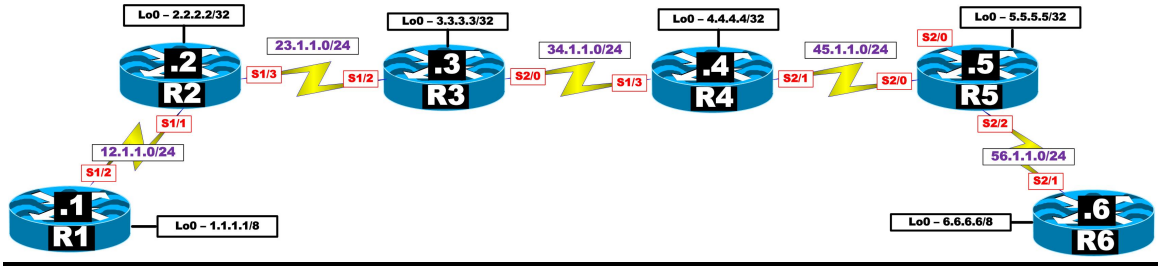

The Topology:

R1 – R2 – R3 – R4 – R5 – R6
R1 and R6 are the CE routers.
R2 and R5 are the PE routers.
R3, and R4 are the P routers.

Everything is configured correctly but the CE routers (R1 and R6) see the routes from the other CE router as "IA".

On R1:

```
R1#sh ip rou ospf | b Gate
Gateway of last resort is not set

      6.0.0.0/32 is subnetted, 1 subnets
O IA     6.6.6.6 [110/129] via 12.1.1.2, 00:00:06, Serial1/2
      56.0.0.0/24 is subnetted, 1 subnets
O IA     56.1.1.0 [110/65] via 12.1.1.2, 00:00:06, Serial1/2
```

We need to configure sham-link:
The loopback interfaces are configured on the PE routers (R2's lo100 is configured with an IP address of 100.1.1.1/32, and on R5, lo100 is configured with an IP address of 200.1.1.1/32) and they are running in the customer VRF.

In this case, we are advertising the two loopback interfaces in OSPF:

```
R2#sh ip int br lo100
Interface              IP-Address      OK? Method Status                Protocol
Loopback100            100.1.1.1       YES manual up                    up

R2#sh run | s router ospf 2 vrf 11
router ospf 2 vrf 11
 area 0 sham-link 100.1.1.1 200.1.1.1
 redistribute bgp 100 subnets
```

```
 network 12.1.1.2 0.0.0.0 area 0
 network 100.1.1.1 0.0.0.0 area 0
```

On R5:

```
R5#sh ip int br lo100
Interface      IP-Address      OK? Method Status                Protocol
Loopback100    200.1.1.1       YES manual up                    up

R5#sh run | s router ospf 2 vrf 66
router ospf 2 vrf 66
 area 0 sham-link 200.1.1.1 100.1.1.1
 redistribute bgp 100 subnets
 network 56.1.1.5 0.0.0.0 area 0
 network 200.1.1.1 0.0.0.0 area 0
```

Since we have mutually redistributed between OSPF and BGP, BGP should be aware of the loopback interfaces.

```
R5#sh ip route vrf 66 bgp | b Gate
Gateway of last resort is not set

      100.0.0.0/32 is subnetted, 1 subnets
B        100.1.1.1 [200/0] via 2.2.2.2, 00:19:54
```

On R2:

```
R2#sh ip route vrf 11 bgp | b Gate
Gateway of last resort is not set

      200.1.1.0/32 is subnetted, 1 subnets
B        200.1.1.1 [200/0] via 5.5.5.5, 00:19:47
```

But this will not work and the sham-link interface will keep on flapping.
So all I did was to use the "ip ospf prefix-suppression" on the lo100 interfaces so OSPF will not be the source of advertisement.

If you do it this way, you won't even have to filter the loopback interfaces on the PE routers because you have done it already with the "ip ospf prefix-suppression" command.

```
R1#sh ip route ospf | b Gate
Gateway of last resort is not set

      6.0.0.0/32 is subnetted, 1 subnets
O        6.6.6.6 [110/130] via 12.1.1.2, 00:29:01, Serial1/2
      56.0.0.0/24 is subnetted, 1 subnets
```

```
O        56.1.1.0 [110/129] via 12.1.1.2, 00:29:01, Serial1/2

R6#sh ip route ospf | b Gate
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
O        1.1.1.1 [110/130] via 56.1.1.5, 00:29:14, Serial2/1
      12.0.0.0/24 is subnetted, 1 subnets
O        12.1.1.0 [110/129] via 56.1.1.5, 00:29:14, Serial2/1

R6#ping 1.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 41/42/43 ms

R6#ping 1.1.1.1 source lo0
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 6.6.6.6
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 41/42/43 ms

R6#trace 1.1.1.1 nu
Type escape sequence to abort.
Tracing the route to 1.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 56.1.1.5 9 msec 8 msec 9 msec
  2 45.1.1.4 [MPLS: Labels 16/21 Exp 0] 40 msec 42 msec 42 msec
  3 34.1.1.3 [MPLS: Labels 16/21 Exp 0] 42 msec 42 msec 40 msec
  4 12.1.1.2 [MPLS: Label 21 Exp 0] 34 msec 34 msec 33 msec
  5 12.1.1.1 46 msec *  41 msec
```