Step 1:

IP Configuration:

On R1:

```
R1(config)#int s1/2
R1(config-if)#ip addr 12.1.1.1 255.255.255.0
R1(config-if)#no shu

R1(config-if)#int lo0
R1(config-if)#ip addr 1.1.1.1 255.0.0.0

R1(config-if)#int lo1
R1(config-if)#ip addr 10.1.1.1 255.255.255.0
```

On R2:

```
R2(config)#int s1/1
R2(config-if)#ip addr 12.1.1.2 255.255.255.0
R2(config-if)#no shu

R2(config)#int s1/3
R2(config-if)#ip addr 23.1.1.2 255.255.255.0
R2(config-if)#no shut

R2(config)#int lo0
R2(config-if)#ip addr 2.2.2.2 255.255.255.255
```

On R3:

```
R3(config)#int s1/2
R3(config-if)#ip addr 23.1.1.3 255.255.255.0
R3(config-if)#no shut

R3(config)#int lo0
R3(config-if)#ip addr 3.3.3.3 255.255.255.255

R3(config)#int s2/0
R3(config-if)#ip addr 34.1.1.3 255.255.255.0
R3(config-if)#no shu
```

On R4:

```
R4(config)#int s1/3
R4(config-if)#ip addr 34.1.1.4 255.255.255.0
R4(config-if)#no shut

R4(config)#int s2/1
R4(config-if)#ip addr 45.1.1.4 255.255.255.0
R4(config-if)#no shu

R4(config)#int lo0
R4(config-if)#ip addr 4.4.4.4 255.255.255.255
```

On R5:

```
R5(config)#int s2/0
R5(config-if)#ip addr 45.1.1.5 255.255.255.0
R5(config-if)#no shut

R5(config)#int s2/2
R5(config-if)#ip addr 56.1.1.5 255.255.255.0
R5(config-if)#no shu

R5(config)#int lo0
R5(config-if)#ip addr 5.5.5.5 255.255.255.255
```

On R6:

```
R6(config)#int lo0
R6(config-if)#ip addr 6.6.6.6 255.0.0.0

R6(config)#int s2/1
R6(config-if)#ip addr 56.1.1.6 255.255.255.0
R6(config-if)#no shu
```

After each step we MUST verify:

On R2:

```
R2#ping 12.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
```

```
R2#ping 23.1.1.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 23.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
```

On R4:

```
R4#ping 34.1.1.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 34.1.1.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms

R4#ping 45.1.1.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 45.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
```

On R6:

```
R6#ping 56.1.1.5
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 56.1.1.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/9 ms
```

Perfect let's go to the second step:

Step 2

Configure IGP in the core:

The core routers in this topology are R2 – R5, therefore, we will run OSPF on all these routers and their loopbacks:

On R2:

```
R2(config)#router ospf 1
R2(config-router)#netw 23.1.1.2 0.0.0.0 a 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 a 0
```

NOTE: We are NOT running OSPF on the interface facing R1, the customer router.

On R3:

```
R3(config)#router ospf 1
R3(config-router)#netw 23.1.1.3 0.0.0.0 a 0
R3(config-router)#netw 3.3.3.3 0.0.0.0 a 0
R3(config-router)#netw 34.1.1.3 0.0.0.0 a 0

*Sep 19 18:39:50.510: %OSPF-5-ADJCHG: Process 1, Nbr 2.2.2.2 on
Serial1/2 from LOADING to FULL, Loading Done
```

On R4:

```
R4(config)#router ospf 1
R4(config-router)#netw 34.1.1.4 0.0.0.0 a 0
R4(config-router)#netw 4.4.4.4 0.0.0.0 a 0
R4(config-router)#netw 45.1.1.4 0.0.0.0 a 0

*Sep 19 18:40:35.574: %OSPF-5-ADJCHG: Process 1, Nbr 34.1.1.3 on
Serial1/3 from LOADING to FULL, Loading Done
```

On R5:

```
R5(config)#router ospf 1
R5(config-router)#netw 45.1.1.5 0.0.0.0 a 0
R5(config-router)#netw 5.5.5.5 0.0.0.0 a 0

*Sep 19 18:41:23.433: %OSPF-5-ADJCHG: Process 1, Nbr 4.4.4.4 on
Serial2/0 from LOADING to FULL, Loading Done
```

Let's verify:

On R2:

```
R2#sh ip rou ospf | b Gate
Gateway of last resort is not set

      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/65] via 23.1.1.3, 00:00:05, Serial1/3
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/129] via 23.1.1.3, 00:01:53, Serial1/3
      5.0.0.0/32 is subnetted, 1 subnets
O        5.5.5.5 [110/193] via 23.1.1.3, 00:01:05, Serial1/3
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/128] via 23.1.1.3, 00:02:38, Serial1/3
      45.0.0.0/24 is subnetted, 1 subnets
O        45.1.1.0 [110/192] via 23.1.1.3, 00:01:53, Serial1/3
```

On R5:

```
R5#sh ip rou ospf | b Gate
Gateway of last resort is not set

      2.0.0.0/32 is subnetted, 1 subnets
O        2.2.2.2 [110/193] via 45.1.1.4, 00:01:36, Serial2/0
      3.0.0.0/32 is subnetted, 1 subnets
O        3.3.3.3 [110/129] via 45.1.1.4, 00:00:26, Serial2/0
      4.0.0.0/32 is subnetted, 1 subnets
O        4.4.4.4 [110/65] via 45.1.1.4, 00:01:36, Serial2/0
      23.0.0.0/24 is subnetted, 1 subnets
O        23.1.1.0 [110/192] via 45.1.1.4, 00:01:36, Serial2/0
      34.0.0.0/24 is subnetted, 1 subnets
O        34.1.1.0 [110/128] via 45.1.1.4, 00:01:36, Serial2/0
```

Step 3:

Configure LDP:

On R2:

```
R2(config)#mpls label range 200 299
R2(config)#mpls label protocol ldp

R2(config)#router ospf 1
R2(config-router)#mpls ldp autoconfig are 0
```

On R3:

```
R3(config)#mpls label range 300 399
R3(config)#mpls label proto ldp

R3(config)#router ospf 1
R3(config-router)#mpls ldp auto area 0

*Sep 19 18:45:39.229: %LDP-5-NBRCHG: LDP Neighbor 2.2.2.2:0 (1) is UP
```

We can see that the LDP session came up between R2 and R3.

On R4:

```
R4(config)#mpls label range 400 499
R4(config)#mpls label protocol ldp
```

```
R4(config)#router ospf 1
R4(config-router)#mpls ldp auto are 0

*Sep 19 18:46:44.902: %LDP-5-NBRCHG: LDP Neighbor 3.3.3.3:0 (1) is UP
```

On R5:

```
R5(config)#mpls label range 500 599
R5(config)#mpls label protocol ldp

R5(config)#router ospf 1
R5(config-router)#mpls ldp auto are 0

*Sep 19 18:47:36.213: %LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (1) is UP
```

To verify the configuration:

On R3:

```
R3#sh mpls inter

Interface              IP            Tunnel   BGP Static Operational
Serial1/2              Yes (ldp)     No       No  No     Yes
Serial2/0              Yes (ldp)     No       No  No     Yes

R3#sh mpls ldp nei

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.64035
        State: Oper; Msgs sent/rcvd: 13/14; Downstream
        Up time: 00:02:59
        LDP discovery sources:
          Serial1/2, Src IP addr: 23.1.1.2
        Addresses bound to peer LDP Ident:
          12.1.1.2        23.1.1.2        2.2.2.2
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 4.4.4.4.37128 - 3.3.3.3.646
        State: Oper; Msgs sent/rcvd: 12/12; Downstream
        Up time: 00:01:53
        LDP discovery sources:
          Serial2/0, Src IP addr: 34.1.1.4
        Addresses bound to peer LDP Ident:
          34.1.1.4        45.1.1.4        4.4.4.4
```

On R5:

```
R5#sh mpls inter

Interface              IP            Tunnel   BGP Static Operational
Serial2/0              Yes (ldp)     No       No  No     Yes

R5#sh mpls ldp nei

    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 5.5.5.5:0
        TCP connection: 4.4.4.4.646 - 5.5.5.5.37072
        State: Oper; Msgs sent/rcvd: 12/11; Downstream
        Up time: 00:01:23
        LDP discovery sources:
          Serial2/0, Src IP addr: 45.1.1.4
        Addresses bound to peer LDP Ident:
          34.1.1.4        45.1.1.4        4.4.4.4
```

Step 4:

Configure iBGP:

This iBGP peer session MUST be established between the PE routers:

On R2:

```
R2(config)#router bgp 100
R2(config-router)#no bgp default ipv4-unicast
R2(config-router)#neigh 5.5.5.5 remote 100
R2(config-router)#neigh 5.5.5.5 update lo0
```

On R5:

```
R5(config)#router bgp 100
R5(config-router)#no bgp default ipv4-unicast
R5(config-router)#neigh 2.2.2.2 remote 100
R5(config-router)#neigh 2.2.2.2 update lo0
```

NOTE: The peer session is NOT up. This is because we configured the two BGP routers NOT to form a peer session using IPv4. This means that we need to activate the peers under the address-family VPNv4:

On R2:

```
R2(config)#router bgp 100
```

```
R2(config-router)#address-family vpnv4
R2(config-router-af)#neigh 5.5.5.5 act
```

On R5:

```
R5(config)#router bgp 100
R5(config-router)#address-family vpnv4
R5(config-router-af)#neigh 2.2.2.2 act

*Sep 19 18:55:18.118: %BGP-5-ADJCHANGE: neighbor 2.2.2.2 Up
```

And sure enough the peer session between the PE routers is up.

Step 5:

Configure VRFs:

NOTE : This is configured on the PE routers:

On R2:

```
R2(config)#ip vrf 11
R2(config-vrf)#rd 1:10
R2(config-vrf)#route-target both 1:100
```

Before applying the VRF to the interface:

```
R5#sh run int s1/1

Building configuration...

Current configuration : 86 bytes
!
interface Serial1/1
 ip address 12.1.1.2 255.255.255.0
 serial restart-delay 0
end

R2(config)#int s1/1
R2(config-if)#ip vrf forwarding 11
```

You should see the following console message:

```
% Interface Serial1/1 IPv4 disabled and address(es) removed due to
disabling VRF 11
```

```
R2(config-if)#ip address 12.1.1.2 255.255.255.0
```

To verify:

```
R2#ping vrf 11 12.1.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 12.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/9/13 ms
```

On R5:

```
R5(config)#ip vrf 66
R5(config-vrf)#rd 1:10
R5(config-vrf)#route-target both 1:100
```

You should see the following console message:

```
*Sep 19 18:59:18.584: %SYS-5-CONFIG_I: Configured from console by
console

R5#sh run int s2/2
Building configuration...

Current configuration : 86 bytes
!
interface Serial2/2
 ip address 56.1.1.5 255.255.255.0
 serial restart-delay 0
end

R5(config)#int s2/2
R5(config-if)#ip vrf forwarding 66


% Interface Serial2/2 IPv4 disabled and address(es) removed due to
disabling VRF 66

R5(config-if)#ip addr 56.1.1.5 255.255.255.0
```

To verify the configuration:

```
R5#ping vrf 66 56.1.1.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 56.1.1.6, timeout is 2 seconds:
```

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/9/10 ms
```

Step 6:

Configuring Routing Protocols between PE and the CE routers:

On R1:

```
R1(config-if)#route-map tst
R1(config-route-map)#match ip addr 1

R1(config-route-map)#router ospf 1
R1(config-router)#router-id 0.0.0.1
R1(config-router)#netw 1.1.1.1 0.0.0.0 a 0
R1(config-router)#netw 12.1.1.1 0.0.0.0 a 0
R1(config-router)#redistri conn route-map tst sub
```

On R2:

```
R2(config)#router ospf 2 vrf 11
R2(config-router)#netw 12.1.1.2 0.0.0.0 a 0

*Sep 19 19:03:24.304: %OSPF-5-ADJCHG: Process 2, Nbr 0.0.0.1 on
Serial1/1 from LOADING to FULL, Loading Done
```

To verify:

On R2:

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

      1.0.0.0/32 is subnetted, 1 subnets
O        1.1.1.1 [110/65] via 12.1.1.1, 00:00:42, Serial1/1
      10.0.0.0/24 is subnetted, 1 subnets
O E2     10.1.1.0 [110/20] via 12.1.1.1, 00:00:42, Serial1/1
```

On R6:

```
R6(config)#router ospf 1
R6(config-router)#netw 6.6.6.6 0.0.0.0 a 0
R6(config-router)#netw 56.1.1.6 0.0.0.0 a 0
```

On R5:

```
R5(config)#router ospf 2 vrf 66
R5(config-router)#netw 56.1.1.5 0.0.0.0 a 0

*Sep 19 19:05:22.801: %OSPF-5-ADJCHG: Process 2, Nbr 6.6.6.6 on
Serial2/2 from LOADING to FULL, Loading Done
```

To verify:

On R5:

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

      6.0.0.0/32 is subnetted, 1 subnets
O        6.6.6.6 [110/65] via 56.1.1.6, 00:00:26, Serial2/2
```

Step 7:

Redistribution:

NOTE : You should perform a one way redistribution and verify if everything is working before configuring redistribution in the other direction:

On R2:

```
R2(config)#router bgp 100
R2(config-router)#address-family ipv4 vrf 11
R2(config-router-af)#redistr ospf 2
```

Let's verify:

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

      1.0.0.0/32 is subnetted, 1 subnets
O        1.1.1.1 [110/65] via 12.1.1.1, 00:05:01, Serial1/1
      10.0.0.0/24 is subnetted, 1 subnets
O E2     10.1.1.0 [110/20] via 12.1.1.1, 00:05:01, Serial1/1


R2#sh ip bgp vpnv4 all
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
             r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
             x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf 11)
 *>  1.1.1.1/32      12.1.1.1                 65         32768 ?
 *>  12.1.1.0/24     0.0.0.0                   0         32768 ?
```

NOTE : the external OSPF route is not redistributed, in order to redistribute the external routes we have to do the following:

```
R2(config)#router bgp 100
R2(config-router)#address-family ipv4 vrf 11
R2(config-router-af)#redistr ospf 2 match internal external 1 external 2
```

To verify:

```
R2#sh ip bgp vpnv4 all

BGP table version is 4, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
             r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
             x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf 11)
 *>  1.1.1.1/32      12.1.1.1                 65         32768 ?
 *>  10.1.1.0/24     12.1.1.1                 20         32768 ?
 *>  12.1.1.0/24     0.0.0.0                   0         32768 ?
```

Perfect. These three routes should be in R5's BGP table, let's verify:

On R5:

```
R5#sh ip bgp vpnv4 all
BGP table version is 7, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
             r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
             x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf 66)
 *>i 1.1.1.1/32      2.2.2.2                  65    100      0 ?
```

```
 *>i 10.1.1.0/24      2.2.2.2                  20    100      0 ?
 *>i 12.1.1.0/24      2.2.2.2                   0    100      0 ?
```

This is good. Now, we need to redistribute BGP into OSPF so R6 can see the routes:

On R5:

```
R5(config)#router ospf 2 vrf 66
R5(config-router)#redistribute bgp 100 sub
```

To verify:

On R6:

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

      1.0.0.0/32 is subnetted, 1 subnets
O IA     1.1.1.1 [110/129] via 56.1.1.5, 00:00:19, Serial2/1
      10.0.0.0/24 is subnetted, 1 subnets
O E2     10.1.1.0 [110/20] via 56.1.1.5, 00:00:19, Serial2/1
      12.0.0.0/24 is subnetted, 1 subnets
O IA     12.1.1.0 [110/65] via 56.1.1.5, 00:00:19, Serial2/1
```

The routes are there, but do we have reachability?

On R1:

```
R1#deb ip icmp
ICMP packet debugging is on
```

On R6:

```
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 0 percent (0/5)
```

We knew that this ping was not going to be successful, but let's see if R1 can see the ICMP echo messages:

On R1:

```
*Sep 19 19:13:44.496: ICMP: echo reply sent, src 1.1.1.1, dst 56.1.1.6,
topology BASE, dscp 0 topoid 0
R1#
*Sep 19 19:13:46.495: ICMP: echo reply sent, src 1.1.1.1, dst 56.1.1.6,
topology BASE, dscp 0 topoid 0
R1#
*Sep 19 19:13:48.496: ICMP: echo reply sent, src 1.1.1.1, dst 56.1.1.6,
topology BASE, dscp 0 topoid 0
R1#
*Sep 19 19:13:50.497: ICMP: echo reply sent, src 1.1.1.1, dst 56.1.1.6,
topology BASE, dscp 0 topoid 0
R1#
*Sep 19 19:13:52.496: ICMP: echo reply sent, src 1.1.1.1, dst 56.1.1.6,
topology BASE, dscp 0 topoid 0
```

We can see that R1 did receive the ICMP messages, now we can redistribute going in the other direction:

On R5:

```
R5(config)#router bgp 100
R5(config-router)#address-family ipv4 vrf 66
R5(config-router-af)#redistr ospf 2
```

To verify:

```
R5#sh ip bgp vpnv4 all
BGP table version is 9, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf 66)
 *>i 1.1.1.1/32       2.2.2.2                  65    100      0 ?
 *>  6.6.6.6/32       56.1.1.6                 65         32768 ?
 *>i 10.1.1.0/24      2.2.2.2                  20    100      0 ?
 *>i 12.1.1.0/24      2.2.2.2                   0    100      0 ?
 *>  56.1.1.0/24      0.0.0.0                   0         32768 ?
```

The routes are redistributed, these routes should show up on R2. Let's verify:

On R2:

```
R2#sh ip bgp vpnv4 all

BGP table version is 8, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf 11)
 *>  1.1.1.1/32       12.1.1.1                 65         32768 ?
 *>i 6.6.6.6/32       5.5.5.5                  65    100      0 ?
 *>  10.1.1.0/24      12.1.1.1                 20         32768 ?
 *>  12.1.1.0/24      0.0.0.0                   0         32768 ?
 *>i 56.1.1.0/24      5.5.5.5                   0    100      0 ?
```

Perfect! Now we need to redistribute BGP into OSPF so the customer router (R1) can see the routes:

```
R2(config)#router ospf 2 vrf 11
R2(config-router)#redistr bgp 100 sub
```

To verify:

On R1:

```
R1#show ip route 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, 00:00:26, Serial1/2
      56.0.0.0/24 is subnetted, 1 subnets
O IA     56.1.1.0 [110/65] via 12.1.1.2, 00:00:26, Serial1/2

R1#ping 6.6.6.6
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 37/40/42 ms
```