

CCIE Service Provider Lab Workbook v4.0

(<http://labs.ine.com/workbook/toc/service-provider-v4>) »

CCIE SP v4 Advanced Technology Labs - Layer 3 VPN

Multi-VRF CE (VRF Lite)

« [MPLS L3 VPN and OSPF Domain-ID \(/workbook/view/service-provider-v4/task/mpls-l3-vpn-and-ospf-domain-id-Mjg2NQ%3D%3D\)](#) | [MPLS L3VPN Inter-AS Option A - Back-to-Back VRF Exchange \(/workbook/view/service-provider-v4/task/mpls-l3vpn-inter-as-option-a-back-to-back-vrf-exchange-Mjg2Nw%3D%3D\)](#) »

Last updated: April 23, 2016

Note:

Initial Configuration & Diagrams: [Load the initial configuration files for the section named Multi VRF CE, which can be found in CCIE SPv4 Topology Diagrams & Initial Configurations \(<http://labs.ine.com/workbook/view/service-provider-v4/task/ccie-spv4-topology-diagrams-initial-configs>\).](#) [Refer to the Multi VRF CE Diagram in order to complete this task.](#)

Task

- Configure a VRF on R1, R2, R4, XR1, and XR2 as follows:
 - VRF Name: VPN_A
 - Route Distinguisher: 100:1
 - Route Target Import: 100:1
 - Route Target Export: 100:1
 - Assign the VRF to the following links:
 - On both R1 and R2 on the link with the address 10.1.2.0/24
 - On both XR1 and XR2 on the link with the address 10.19.20.0/24
 - On R1's link to R3
 - On R4's link to R7
 - On XR2's link to R6
- Configure a VRF on R1, R2, R5, XR1, and XR2 as follows:
 - VRF Name: VPN_B
 - Route Distinguisher: 100:2
 - Route Target Import: 100:2
 - Route Target Export: 100:2
 - Assign the VRF to the following links:
 - On both R1 and R2 on the link with the address 30.1.2.0/24
 - On both XR1 and XR2 on the link with the address 30.19.20.0/24
 - On R1's link to R9
 - On R5's link to R8
 - On XR2's link to R9
- Configure OSPF routing for VRF VPN_A as follows:
 - R3, R6, and R7 are preconfigured with OSPF.
 - Configure OSPF Process-ID 100 on R1, R2, R4, XR1, and XR2.
 - Enable this OSPF process on all links in VRF VPN_A.

- Configure the OSPF Domain-ID on XR1 as type 0005 with value 0x000000640200.
- Configure the VRF aware OSPF process of R1 and XR2 to ignore the OSPF Down Bit in Type 3 LSAs.
- Configure RIP routing for VRF VPN_B as follows:
 - R8 and R9 are preconfigured with RIP.
 - Configure RIPv2 on R1, R2, R5, XR1, and XR2.
 - Enable this RIP process on all links in VRF VPN_B.
- Configure BGP on R2, R4, R5, and XR1 as follows:
 - Use BGP AS 100.
 - Configure a full mesh of iBGP VPNv4 peerings between the PE routers.
 - Use their Loopback0 interfaces as the source of the BGP sessions.
 - Redistribute between the VRF aware OSPF process and VPNv4 BGP.
 - Redistribute between the VRF aware RIP process and VPNv4 BGP.
- Once complete, the following reachability should be achieved:
 - R3, R6, and R7 should have full reachability to all networks in VRF VPN_A.
 - R8 and R9 should have full reachability to all networks in VRF VPN_B.

Configuration [Click to collapse](#)

```
R1:
vrf definition VPN_A
  rd 100:1
  !
  address-family ipv4
  exit-address-family
  !
vrf definition VPN_B
  rd 100:2
  !
  address-family ipv4
  exit-address-family
  !
interface GigabitEthernet1.13
  vrf forwarding VPN_A
  ip address 10.1.3.1 255.255.255.0
  !
interface GigabitEthernet1.112
  vrf forwarding VPN_A
  ip address 10.1.2.1 255.255.255.0
  !
interface GigabitEthernet1.120
  vrf forwarding VPN_B
  ip address 30.1.20.1 255.255.255.0
  !
interface GigabitEthernet1.212
  vrf forwarding VPN_B
  ip address 30.1.2.1 255.255.255.0
  !
router ospf 100 vrf VPN_A
  capability vrf-lite
  network 0.0.0.0 255.255.255.255 area 0
  !
router rip
  !
address-family ipv4 vrf VPN_B
  network 30.0.0.0
  no auto-summary
  version 2
  exit-address-family
  !
R2:
vrf definition VPN_A
  rd 100:1
  !
  address-family ipv4
  route-target export 100:1
  route-target import 100:1
  exit-address-family
  !
vrf definition VPN_B
  rd 100:2
```

```

!
address-family ipv4
route-target export 100:2
route-target import 100:2
exit-address-family
!
interface GigabitEthernet1.112
vrf forwarding VPN_A
ip address 10.1.2.2 255.255.255.0
!
interface GigabitEthernet1.212
vrf forwarding VPN_B
ip address 30.1.2.2 255.255.255.0
!
router ospf 100 vrf VPN_A
redistribute bgp 100 subnets
network 0.0.0.0 255.255.255.255 area 0
!
router rip
!
address-family ipv4 vrf VPN_B
redistribute bgp 100 metric 5
network 30.0.0.0
no auto-summary
version 2
exit-address-family
!
router bgp 100
no bgp default ipv4-unicast
neighbor 4.4.4.4 remote-as 100
neighbor 4.4.4.4 update-source Loopback0
neighbor 5.5.5.5 remote-as 100
neighbor 5.5.5.5 update-source Loopback0
neighbor 19.19.19.19 remote-as 100
neighbor 19.19.19.19 update-source Loopback0
!
address-family vpv4
neighbor 4.4.4.4 activate
neighbor 4.4.4.4 send-community extended
neighbor 5.5.5.5 activate
neighbor 5.5.5.5 send-community extended
neighbor 19.19.19.19 activate
neighbor 19.19.19.19 send-community extended
exit-address-family
!
address-family ipv4 vrf VPN_A
redistribute ospf 100 vrf VPN_A
exit-address-family
!
address-family ipv4 vrf VPN_B
redistribute rip
exit-address-family
!

```

```
R4:
vrf definition VPN_A
  rd 100:1
  !
  address-family ipv4
    route-target export 100:1
    route-target import 100:1
  exit-address-family
!
interface GigabitEthernet1.47
  vrf forwarding VPN_A
  ip address 10.4.7.4 255.255.255.0
!
router ospf 100 vrf VPN_A
  redistribute bgp 100 subnets
  network 0.0.0.0 255.255.255.255 area 0
!
router bgp 100
  no bgp default ipv4-unicast
  neighbor 2.2.2.2 remote-as 100
  neighbor 2.2.2.2 update-source Loopback0
  neighbor 5.5.5.5 remote-as 100
  neighbor 5.5.5.5 update-source Loopback0
  neighbor 19.19.19.19 remote-as 100
  neighbor 19.19.19.19 update-source Loopback0
!
  address-family vpnv4
    neighbor 2.2.2.2 activate
    neighbor 2.2.2.2 send-community extended
    neighbor 5.5.5.5 activate
    neighbor 5.5.5.5 send-community extended
    neighbor 19.19.19.19 activate
    neighbor 19.19.19.19 send-community extended
  exit-address-family
!
```

```
  address-family ipv4 vrf VPN_A
    no synchronization
    redistribute ospf 100 vrf VPN_A match internal external 1 external 2
  exit-address-family
```

```
R5:
vrf definition VPN_B
  rd 100:2
  route-target export 100:2
  route-target import 100:2
  !
  address-family ipv4
    route-target export 100:2
    route-target import 100:2
  exit-address-family
!
interface GigabitEthernet1.58
  vrf forwarding VPN_B
  ip address 30.5.8.5 255.255.255.0
```

```

!
router rip
!
address-family ipv4 vrf VPN_B
  redistribute bgp 100 metric 5
  network 30.0.0.0
  no auto-summary
  version 2
exit-address-family
!
router bgp 100
  no bgp default ipv4-unicast
  neighbor 2.2.2.2 remote-as 100
  neighbor 2.2.2.2 update-source Loopback0
  neighbor 4.4.4.4 remote-as 100
  neighbor 4.4.4.4 update-source Loopback0
  neighbor 19.19.19.19 remote-as 100
  neighbor 19.19.19.19 update-source Loopback0
!
address-family vpnv4
  neighbor 2.2.2.2 activate
  neighbor 2.2.2.2 send-community extended
  neighbor 4.4.4.4 activate
  neighbor 4.4.4.4 send-community extended
  neighbor 19.19.19.19 activate
  neighbor 19.19.19.19 send-community extended
exit-address-family
!
address-family ipv4 vrf VPN_B
  no synchronization
  redistribute rip
exit-address-family
!

```

```

vrf VPN_A
  address-family ipv4 unicast
    import route-target
      100:1
    !
    export route-target
      100:1
    !
  !
vrf VPN_B
  address-family ipv4 unicast
    import route-target
      100:2
    !
    export route-target
      100:2
    !
  !

```

```
!  
interface GigabitEthernet0/0/0/0.11920  
  vrf VPN_A  
  no ipv4 address  
  ipv4 address 10.19.20.19 255.255.255.0  
!  
interface GigabitEthernet0/0/0/0.21920  
  vrf VPN_B  
  no ipv4 address  
  ipv4 address 30.19.20.19 255.255.255.0  
!  
route-policy BGP_TO_RIP  
  set rip-metric 5  
end-policy  
!  
router ospf 100  
  vrf VPN_A  
  domain-id type 0005 value 000000640200  
  redistribute bgp 100  
  area 0  
    interface GigabitEthernet0/0/0/0.11920  
      !  
      !  
      !  
      !  
router bgp 100  
  address-family vpnv4 unicast  
  !  
  neighbor 2.2.2.2  
  remote-as 100  
  update-source Loopback0  
  address-family vpnv4 unicast  
  !  
  !  
  neighbor 4.4.4.4  
  remote-as 100  
  update-source Loopback0  
  address-family vpnv4 unicast  
  !  
  !  
  neighbor 5.5.5.5  
  remote-as 100  
  update-source Loopback0  
  address-family vpnv4 unicast  
  !  
  !  
  vrf VPN_A  
  rd 100:1  
  address-family ipv4 unicast  
    redistribute ospf 100  
  !  
  !  
  vrf VPN_B  
  rd 100:2
```

```
address-family ipv4 unicast
  redistribute rip
  !
  !
  !
router rip
vrf VPN_B
  interface GigabitEthernet0/0/0/0.21920
  !
  redistribute bgp 100 route-policy BGP_TO_RIP
  !
  !
```

XR2:

```
vrf VPN_A
  address-family ipv4 unicast
  !
  !
vrf VPN_B
  address-family ipv4 unicast
  !
  !
interface GigabitEthernet0/0/0/0.120
  vrf VPN_B
  no ipv4 address
  ipv4 address 30.1.20.20 255.255.255.0
  !
interface GigabitEthernet0/0/0/0.620
  vrf VPN_A
  no ipv4 address
  ipv4 address 10.6.20.20 255.255.255.0
  !
interface GigabitEthernet0/0/0/0.11920
  vrf VPN_A
  no ipv4 address
```

```
  ipv4 address 10.19.20.20 255.255.255.0
  !
interface GigabitEthernet0/0/0/0.21920
  vrf VPN_B
  no ipv4 address
  ipv4 address 30.19.20.20 255.255.255.0
  !
router ospf 100
  vrf VPN_A
  disable-dn-bit-check
  area 0
    interface GigabitEthernet0/0/0/0.620
    !
    interface GigabitEthernet0/0/0/0.11920
    !
    !
    !
  !
router rip
```

```
vrf VPN_B
  interface GigabitEthernet0/0/0/0.120
  !
  interface GigabitEthernet0/0/0/0.21920
  !
  redistribute bgp 100
  !
  !
```

Verification

The Multi VRF CE feature, or what is more commonly referred to as “VRF Lite”, simply means that the router has been configured with multiple Virtual Routing and Forwarding instances (VRFs), but does not have MPLS or VPNv4 BGP configured. Typically this configuration is used to separate a single router or switch into multiple logical devices, as interfaces assigned to different VRF tables are then independent of each other. A common practical application of this feature is for managed CE devices that connect to multiple customers. For example a service provider may service multiple customers in the same office building, with their connections from the PE to CE aggregated by the managed CE.

Configuration of this feature is identical to the other L3VPN configs that have been done up to this point, with the exception that the BGP and MPLS portions are left out. As we see below on R1 and XR2, there are two separate VRFs configured to service VPN_A and VPN_B. Note that no Route Target import or export policies are configured on R1 or XR2, nor is the Route Distinguisher value set on XR2. These values can be configured, but are not needed, since they are only relevant within the scope of VPNv4 BGP.

```

R1#show ip vrf detail

VRF VPN_A (VRF Id = 1); default RD 100:1; default VPNID <not set>

  New CLI format, supports multiple address-families

  Flags: 0x180C

  Interfaces:

    Gi1.13          Gi1.112
Address family ipv4 unicast (Table ID = 0x1):

  Flags: 0x0

  No Export VPN route-target communities

  No Import VPN route-target communities

  No import route-map

  No global export route-map

  No export route-map

  VRF label distribution protocol: not configured

  VRF label allocation mode: per-prefix

VRF VPN_B (VRF Id = 2); default RD 100:2; default VPNID <not set>

  New CLI format, supports multiple address-families

  Flags: 0x180C

  Interfaces:

    Gi1.120         Gi1.212
Address family ipv4 unicast (Table ID = 0x2):

  Flags: 0x0

  No Export VPN route-target communities

  No Import VPN route-target communities

  No import route-map

  No global export route-map

  No export route-map

  VRF label distribution protocol: not configured

  VRF label allocation mode: per-prefix

RP/0/0/CPU0:XR2#show vrf all detail

Tue May 12 00:04:13.658 UTC

```

```

VRF **nVSatellite; RD not set; VPN ID not set

VRF mode: Regular

Description not set

Interfaces:

  nV-Loopback0

Address family IPV4 Unicast

  No import VPN route-target communities

  No export VPN route-target communities

  No import route policy

  No export route policy

Address family IPV6 Unicast

  No import VPN route-target communities

  No export VPN route-target communities

  No import route policy

  No export route policy

VRF VPN_A; RD not set; VPN ID not set

VRF mode: Regular

Description not set

```

```

Interfaces:

  GigabitEthernet0/0/0/0.620

  GigabitEthernet0/0/0/0.11920
Address family IPV4 Unicast

  No import VPN route-target communities

  No export VPN route-target communities

  No import route policy

  No export route policy
Address family IPV6 Unicast

  No import VPN route-target communities

  No export VPN route-target communities

  No import route policy

  No export route policy

VRF VPN_B; RD not set; VPN ID not set

VRF mode: Regular

Description not set

Interfaces:

  GigabitEthernet0/0/0/0.120

  GigabitEthernet0/0/0/0.21920
Address family IPV4 Unicast

  No import VPN route-target communities

  No export VPN route-target communities

  No import route policy

  No export route policy
Address family IPV6 Unicast

  No import VPN route-target communities

  No export VPN route-target communities

  No import route policy

  No export route policy

```

One of the potential design problems with this configuration (VRFs without MPLS/BGP) specifically relates to how the OSPF loop prevention process changes when a VRF aware OSPF process is used as opposed to a process in the global routing table.

Per RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) (<http://www.ietf.org/rfc/rfc4577.txt>), loop prevention of VRF aware OSPF is enhanced by adding a new bit in the LSA Options field known as the “Down Bit” or the DN Bit, and by adding an automatically derived route tag. The Down Bit is set in OSPF when a PE router redistributes VPNv4 BGP into OSPF and originates a Type 3 Network Summary LSA (an Inter-Area route). For Type 5 External LSAs (an E1 or E2 route) a VPN Route Tag value is used in place of the Down Bit, but essentially serves the same purpose. The goal of the addition of the Down Bit is to use it determine if an OSPF route was originated from the local customer site or a remote customer site. Note that on current version of IOS and IOS-XR, such as the versions used during these labs, the DN bit is set in both Type-3 and Type-5 LSA - more closely following the RFC. The DN bit was only set on Type-3 LSAs in previous versions.

CONTENTS

Since the Down Bit is set when a Type-3 or Type-5 LSA is generated during VPNv4 BGP to OSPF redistribution, and only routes originated by a remote customer site would have been redistributed from VPNv4 BGP into OSPF, the Down Bit can then be used to figure out if the route is from the local site or a remote site. Feedback loops can then be prevented for routes being redistributed from VPNv4 BGP to OSPF, then from OSPF back to VPNv4 BGP, then from VPNv4 BGP back to OSPF, etc. by using the Down Bit as an indicator if the route should be candidate for redistribution or not. From an implementation point of view the Down Bit is specifically used to control whether a route in the OSPF database can be installed in the routing table.

If the local router is running the VRF aware OSPF process, it is assumed that this router is a PE that is doing OSPF to VPNv4 BGP redistribution, as this is normally the default design logic of VRFs since they are used for MPLS L3VPN applications. In this case, it is advantageous to prevent received routes with the Down Bit set to be installed in the routing table, mitigating the possibility of redistributing the route into another protocol. The problem with this default logic though is that in the application of VRF Lite (i.e. Multi VRF CE), the router is not performing OSPF to VPNv4 BGP redistribution, and hence this native loop prevention mechanism is not needed.

In this example, the Down Bit loop prevention must be disabled on the Multi VRF CEs of R1 and XR2, otherwise they will not install Type 3 or Type 5 LSAs that are advertised by PE routers R2 and XR1. All of the Type-3 and Type-5 LSAs currently advertised by R2 and XR1 have the DN bit set, as they are routes that originated from another OSPF domain. This is what the commands **capability vrf-lite** and **disable-dn-bit-check** do under the VRF aware OSPF processes of IOS and IOS XR respectively. In the below output we see the final result of this configuration, where the VRF VPN_A and VPN_B tables are fully populated everywhere.

```
R3#show ip route ospf
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
ia - IS-IS inter area, * - candidate default, U - per-user static route
```

```
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
```

```
a - application route
```

```
+ - replicated route, % - next hop override
```

```
Gateway of last resort is not set
```

```
6.0.0.0/32 is subnetted, 1 subnets
```

```
O IA 6.6.6.6 [110/5] via 10.1.3.1, 00:01:14, GigabitEthernet1.13
```

```
7.0.0.0/32 is subnetted, 1 subnets
```

```
O E2 7.7.7.7 [110/20] via 10.1.3.1, 00:03:35, GigabitEthernet1.13
```

```
10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
```

```
O 10.1.2.0/24 [110/2] via 10.1.3.1, 00:45:54, GigabitEthernet1.13
```

```
O IA 10.4.7.0/24 [110/3] via 10.1.3.1, 00:03:35, GigabitEthernet1.13
```

```
O IA 10.6.20.0/24 [110/4] via 10.1.3.1, 00:01:14, GigabitEthernet1.13
```

```
O IA 10.7.7.7/32 [110/4] via 10.1.3.1, 00:03:35, GigabitEthernet1.13
```

```
O IA 10.19.20.0/24 [110/3] via 10.1.3.1, 00:01:14, GigabitEthernet1.13
```

```
R6#show ip route ospf
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
ia - IS-IS inter area, * - candidate default, U - per-user static route
```

```
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
```

```
a - application route
```

```
+ - replicated route, % - next hop override
```

```
Gateway of last resort is not set
```

```
3.0.0.0/32 is subnetted, 1 subnets
```

```
O IA 3.3.3.3 [110/5] via 10.6.20.20, 00:01:34, GigabitEthernet1.620
```

```
7.0.0.0/32 is subnetted, 1 subnets
```

```
O E2 7.7.7.7 [110/20] via 10.6.20.20, 00:23:00, GigabitEthernet1.620
```

```
10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
```

```
O IA 10.1.2.0/24 [110/3] via 10.6.20.20, 00:01:34, GigabitEthernet1.620
```

```
O IA 10.1.3.0/24 [110/4] via 10.6.20.20, 00:01:34, GigabitEthernet1.620
```

```
O IA 10.4.7.0/24 [110/3] via 10.6.20.20, 00:01:34, GigabitEthernet1.620
```

```
O IA 10.7.7.7/32 [110/4] via 10.6.20.20, 00:01:34, GigabitEthernet1.620
```

```
O 10.19.20.0/24 [110/2] via 10.6.20.20, 00:41:44, GigabitEthernet1.620
```

```
R7#show ip route ospf
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
a - application route
+ - replicated route, % - next hop override

Gateway of last resort is not set

```
3.0.0.0/32 is subnetted, 1 subnets
O IA   3.3.3.3 [110/4] via 10.4.7.4, 00:04:15, GigabitEthernet1.47
6.0.0.0/32 is subnetted, 1 subnets
O IA   6.6.6.6 [110/4] via 10.4.7.4, 00:01:54, GigabitEthernet1.47
10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O IA   10.1.2.0/24 [110/2] via 10.4.7.4, 00:04:20, GigabitEthernet1.47
O IA   10.1.3.0/24 [110/3] via 10.4.7.4, 00:04:15, GigabitEthernet1.47
O IA   10.6.20.0/24 [110/3] via 10.4.7.4, 00:01:54, GigabitEthernet1.47
O IA   10.19.20.0/24 [110/2] via 10.4.7.4, 00:01:54, GigabitEthernet1.47
```

The key prefixes to check reachability are networks 7.7.7.7/32 and 10.7.7.7/32, both of which are being originated by R7.

```
R3#show ip route 7.7.7.7
Routing entry for 7.7.7.7/32
  Known via "ospf 100", distance 110, metric 20
  Tag Complete, Path Length == 1, AS 100, , type extern 2, forward metric 2
  Last update from 10.1.3.1 on GigabitEthernet1.13, 00:05:01 ago
  Routing Descriptor Blocks:
  * 10.1.3.1, from 10.1.2.2, 00:05:01 ago, via GigabitEthernet1.13
    Route metric is 20, traffic share count is 1
    Route tag 3489661028
```

```
R3#show ip route 10.7.7.7
Routing entry for 10.7.7.7/32
  Known via "ospf 100", distance 110, metric 4, type inter area
  Last update from 10.1.3.1 on GigabitEthernet1.13, 00:05:31 ago
  Routing Descriptor Blocks:
  * 10.1.3.1, from 10.1.2.2, 00:05:31 ago, via GigabitEthernet1.13
    Route metric is 4, traffic share count is 1
```

```
R3#ping 7.7.7.7 source 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.7.7.7, timeout is 2 seconds:
Packet sent with a source address of 3.3.3.3
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/8/25 ms
```

```
R3#ping 10.7.7.7 source 3.3.3.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.7.7.7, timeout is 2 seconds:
Packet sent with a source address of 3.3.3.3
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/6/19 ms
```

The first prefix, 7.7.7.7/32, is a Type 5 External LSA that is being originated by R7.

```

R7#show ip ospf database external 7.7.7.7

      OSPF Router with ID (10.7.7.7) (Process ID 100)

      Type-5 AS External Link States

LS age: 1636
Options: (No TOS-capability, DC, Upward)
LS Type: AS External Link
Link State ID: 7.7.7.7 (External Network Number )
Advertising Router: 10.7.7.7
LS Seq Number: 80000002
Checksum: 0xB4AF
Length: 36
Network Mask: /32

    Metric Type: 2 (Larger than any link state path)
    MTID: 0
    Metric: 20
    Forward Address: 0.0.0.0
    External Route Tag: 0

```

R7's PE (R4) takes this OSPF route and redistributes it into VPNv4 BGP. Since the route is external to start, it is encoded with the OSPF Route Type of 5 for external.

```

R4#show bgp vpnv4 unicast all 7.7.7.7/32

BGP routing table entry for 100:1:7.7.7.7/32, version 2
Paths: (1 available, best #1, table VPN_A)
Advertised to update-groups:
  1
Refresh Epoch 1
Local
  10.4.7.7 (via vrf VPN_A) from 0.0.0.0 (4.4.4.4)
  Origin incomplete, metric 20, localpref 100, weight 32768, valid, sourced, best
  Extended Community: RT:100:1 OSPF DOMAIN ID:0x0005:0x000000640200
  OSPF RT:0.0.0.0:5:1 OSPF ROUTER ID:10.4.7.4:0
  mpls labels in/out 36/nolabel
  rx pathid: 0, tx pathid: 0x0

```

When the remote PEs of R2 and XR1 redistribute VPNv4 BGP back into OSPF, they originate this route into the database as a Type 5 External LSA, and set both the DN Bit and VPN Route Tag. On previous versions of IOS, only the VPN Route Tag would have been set on this LSA.

```
R2#show ip ospf 100 database external 7.7.7.7
```

```
OSPF Router with ID (10.1.2.2) (Process ID 100)
```

```
Type-5 AS External Link States
```

```
LS age: 574
```

```
Options: (No TOS-capability, DC, Downward)
```

```
LS Type: AS External Link
```

```
Link State ID: 7.7.7.7 (External Network Number )
```

```
Advertising Router: 10.1.2.2
```

```
LS Seq Number: 80000004
```

```
Checksum: 0xE5D6
```

```
Length: 36
```

```
Network Mask: /32
```

```
Metric Type: 2 (Larger than any link state path)
```

```
MTID: 0
```

```
Metric: 20
```

```
Forward Address: 0.0.0.0
```

```
External Route Tag: 3489661028
```

```
RP/0/0/CPU0:XR1#show ospf vrf VPN_A database external 7.7.7.7
```

```
Tue May 12 00:43:39.206 UTC
```

```
OSPF Router with ID (19.19.19.19) (Process ID 100, VRF VPN_A)
```

```
Type-5 AS External Link States
```

```
LS age: 462
```

```
Options: (No TOS-capability, DC, DN)
```

```
LS Type: AS External Link
```

```
Link State ID: 7.7.7.7 (External Network Number)
```

```
Advertising Router: 19.19.19.19
```

```
LS Seq Number: 80000003
```

```
Checksum: 0x2858
```

```
Length: 36
```

```
Network Mask: /32
```

```
Metric Type: 2 (Larger than any link state path)
```

```
TOS: 0
```

```
Metric: 20
```

```
Forward Address: 0.0.0.0
```

```
External Route Tag: 3489661028
```

This value 3489661028 is where the BGP AS number of the PE router is encoded. Converted to hexadecimal this number is 0xD0000064. The last two bytes, 0x0064, is the BGP AS number of the PE router, which is specifically AS 100 in this case. This value is then used for loop prevention where the OSPF route will not be redistributed back into VPNv4 BGP if the destination BGP AS number is 100. However, since the Multi VRF CEs of R1 and XR2 are not doing OSPF to VPNv4 BGP redistribution, this value is not examined, and the prefix can be installed without any issues. This can be seen below as both routers have the route installed and can reach it.

```
R1#show ip route vrf VPN_A 7.7.7.7

Routing Table: VPN_A
Routing entry for 7.7.7.7/32
  Known via "ospf 100", distance 110, metric 20
  Tag Complete, Path Length == 1, AS 100, , type extern 2, forward metric 1
  Last update from 10.1.2.2 on GigabitEthernet1.112, 00:11:27 ago
  Routing Descriptor Blocks:
  * 10.1.2.2, from 10.1.2.2, 00:11:27 ago, via GigabitEthernet1.112
    Route metric is 20, traffic share count is 1
    Route tag 3489661028
```

```
R1#ping vrf VPN_A 7.7.7.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.7.7.7, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/6/19 ms
```

```
RP/0/0/CPU0:XR2#show route vrf VPN_A 7.7.7.7/32

Tue May 12 00:45:31.959 UTC

Routing entry for 7.7.7.7/32
  Known via "ospf 100", distance 110, metric 20
  Tag 3489661028, type extern 2
  Installed May 12 00:14:30.716 for 00:31:01
  Routing Descriptor Blocks
  10.19.20.19, from 19.19.19.19, via GigabitEthernet0/0/0/0.11920
    Route metric is 20
  No advertising protos.
```

```
RP/0/0/CPU0:XR2#ping vrf VPN_A 7.7.7.7

Fri Mar 16 21:13:23.650 UTC

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 7.7.7.7, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/11/29 ms
```

The second prefix that is generated by R7, 10.7.7.7/32, is advertised as a normal internal OSPF route in Area 0.

```

R7#show ip ospf database router 10.7.7.7

      OSPF Router with ID (10.7.7.7) (Process ID 100)

      Router Link States (Area 0)

LS age: 1264
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 10.7.7.7
Advertising Router: 10.7.7.7
LS Seq Number: 80000005
Checksum: 0x284
Length: 48
AS Boundary Router
Number of Links: 2

Link connected to: a Stub Network
(Link ID) Network/subnet number: 10.7.7.7
(Link Data) Network Mask: 255.255.255.255
Number of MTID metrics: 0
TOS 0 Metrics: 1

Link connected to: a Transit Network
(Link ID) Designated Router address: 10.4.7.7
(Link Data) Router Interface address: 10.4.7.7
Number of MTID metrics: 0
TOS 0 Metrics: 1

```

This means that when R7's PE router, R4, redistributes OSPF into BGP, the OSPF Route Type is set to 2, as this route is learned via the Type 2 LSA that the DR on the Ethernet link between R4 and R7 is generating.

```

R4#show bgp vpnv4 unicast all 10.7.7.7/32

BGP routing table entry for 100:1:10.7.7.7/32, version 4
Paths: (1 available, best #1, table VPN_A)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  Local
    10.4.7.7 (via vrf VPN_A) from 0.0.0.0 (4.4.4.4)
      Origin incomplete, metric 2, localpref 100, weight 32768, valid, sourced, best
      Extended Community: RT:100:1 OSPF DOMAIN ID:0x0005:0x000000640200
        OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:10.4.7.4:0
      mpls labels in/out 34/nolabel
      rx pathid: 0, tx pathid: 0x0

```

When the remote PEs R2 and XR1 receive this VPNv4 route, they have to decide if they are going to encode it as a Type 3 LSA (Inter-Area route) or a Type 5 LSA (External route). As previously discussed, the route can be encoded as Type 3 if it was an Intra-Area or Inter-Area route to begin with, which the OSPF Route Type field tells us, and if the OSPF Domain-ID of the VPNv4 route matches the local OSPF process. In the case of R2 and XR1 both of these checks pass, and the route is then redistributed as a Type 3 LSA (Inter Area route). Note that when this origination occurs, the OSPF Down Bit is set in the LSA Options field.

<https://t.me/learningnets>

```
R2#show bgp vpnv4 unicast all 10.7.7.7/32
BGP routing table entry for 100:1:10.7.7.7/32, version 13
Paths: (1 available, best #1, table VPN_A)
Not advertised to any peer
Refresh Epoch 1
Local
  4.4.4.4 (metric 2) (via default) from 4.4.4.4 (4.4.4.4)
  Origin incomplete, metric 2, localpref 100, valid, internal, best
  Extended Community: RT:100:1 OSPF DOMAIN ID:0x0005:0x00000640200
    OSPF RT:0.0.0.2:0 OSPF ROUTER ID:10.4.7.4:0
  mpls labels in/out nlabel/34
  rx pathid: 0, tx pathid: 0x0
```

```
R2#show ip ospf 100 database summary 10.7.7.7
```

```
OSPF Router with ID (10.1.2.2) (Process ID 100)
```

```
Summary Net Link States (Area 0)
```

```
LS age: 944
Options: (No TOS-capability, DC, Downward)
LS Type: Summary Links(Network)
Link State ID: 10.7.7.7 (summary Network Number)
Advertising Router: 10.1.2.2
LS Seq Number: 80000004
Checksum: 0x2960
Length: 28
Network Mask: /32
  MTID: 0      Metric: 2
```

```
RP/0/0/CPU0:XR1#show bgp vpnv4 unicast vrf VPN_A 10.7.7.7/32
```

```
Tue May 12 00:49:45.511 UTC
```

```
BGP routing table entry for 10.7.7.7/32, Route Distinguisher: 100:1
```

```
Versions:
```

Process	bRIB/RIB	SendTblVer
Speaker	37	37

```
Last Modified: May 12 00:14:30.451 for 00:35:15
```

```
Paths: (1 available, best #1)
```

```
Not advertised to any peer
```

```
Path #1: Received by speaker 0
```

```
Not advertised to any peer
```

```
Local
```

```
4.4.4.4 (metric 3) from 4.4.4.4 (4.4.4.4)
```

```
Received Label 34
```

```
Origin incomplete, metric 2, localpref 100, valid, internal, best, group-best, import-candidate, imported
```

```
Received Path ID 0, Local Path ID 1, version 37
```

```
Extended community: OSPF domain-id:0x5:0x00000640200 OSPF route-type:0:2:0x0 OSPF router-id:10.4.7.4 RT:100:1
```

```
Source VRF: VPN_A, Source Route Distinguisher: 100:1
```

```
RP/0/0/CPU0:XR1#show ospf 100 vrf VPN_A database summary 10.7.7.7
```

```
Tue May 12 00:50:41.917 UTC
```

```
OSPF Router with ID (19.19.19.19) (Process ID 100, VRF VPN_A)
```

```
Summary Net Link States (Area 0)
```

```
LS age: 885
```

```
Options: (No TOS-capability, DC, DN)
```

```
LS Type: Summary Links (Network)
```

```
Link State ID: 10.7.7.7 (Summary Network Number)
```

```
Advertising Router: 19.19.19.19
```

```
LS Seq Number: 80000001
```

```
Checksum: 0x6fdf
```

```
Length: 28
```

```
Network Mask: /32
```

```
TOS: 0 Metric: 2
```

The Multi VRF CE routers R1 and XR2 now learn these Type 3 LSAs. Since the Down Bit is manually being ignored, the routes *can* be installed in the routing table.

```
R1#show ip route vrf VPN_A 10.7.7.7
```

```
Routing Table: VPN_A
```

```
Routing entry for 10.7.7.7/32
```

```
Known via "ospf 100", distance 110, metric 3, type inter area
```

```
Last update from 10.1.2.2 on GigabitEthernet1.112, 00:17:43 ago
```

```
Routing Descriptor Blocks:
```

```
* 10.1.2.2, from 10.1.2.2, 00:17:43 ago, via GigabitEthernet1.112
```

```
Route metric is 3, traffic share count is 1
```

```
R1#show ip ospf 100 database summary 10.7.7.7
```

```
OSPF Router with ID (10.1.3.1) (Process ID 100)
```

```
Summary Net Link States (Area 0)
```

```
LS age: 1079
```

```
Options: (No TOS-capability, DC, Downward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 10.7.7.7 (summary Network Number)
```

```
Advertising Router: 10.1.2.2
```

```
LS Seq Number: 80000004
```

```
Checksum: 0x2960
```

```
Length: 28
```

```
Network Mask: /32
```

```
MTID: 0 Metric: 2
```

```
R1#ping vrf VPN_A 10.7.7.7
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.7.7.7, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/19 ms
```

```
RP/0/0/CPU0:XR2#show route vrf VPN_A 10.7.7.7/32
```

```
Tue May 12 00:52:19.581 UTC
```

```
Routing entry for 10.7.7.7/32
```

```
Known via "ospf 100", distance 110, metric 3, type inter area
```

```
Installed May 12 00:35:56.178 for 00:16:23
```

```
Routing Descriptor Blocks
```

```
10.19.20.19, from 19.19.19.19, via GigabitEthernet0/0/0.11920
```

```
Route metric is 3
```

```
No advertising protos.
```

```
RP/0/0/CPU0:XR2#show ospf vrf VPN_A database summary 10.7.7.7
```

```
Tue May 12 00:52:41.529 UTC
```

```
OSPF Router with ID (20.20.20.20) (Process ID 100, VRF VPN_A)
```

Summary Net Link States (Area 0)

```
Routing Bit Set on this LSA
LS age: 1006
Options : (No TOS-capability, DC, DN)
LS Type: Summary Links (Network)
Link State ID: 10.7.7.7 (Summary Network Number)
Advertising Router: 19.19.19.19
LS Seq Number: 80000001
Checksum: 0x6fdf
Length: 28
Network Mask: /32
    TOS: 0 Metric: 2
```

```
RP/0/0/CPU0:XR2#ping vrf VPN_A 10.7.7.7
```

```
Tue May 12 00:54:13.963 UTC
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.7.7.7, timeout is 2 seconds:
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/5/9 ms
```

If we were to revert back to the default behavior of using the Down Bit for filtering, these prefixes would still be advertised on to the C routers R3 and R6, but they would not actually be able to reach them, as the Multi VRF CE routers R1 and XR2 would not install them in the routing table. This can be demonstrated as follows.

Note that currently both Inter-Area and External (Type-3 and Type-5) are installed in the RIB. Both of these routes have the DN bit set as previously shown.

```
R1#show ip route vrf VPN_A 10.7.7.7
```

```
Routing Table: VPN_A
```

```
Routing entry for 10.7.7.7/32
```

```
Known via "ospf 100", distance 110, metric 3, type inter area
```

```
Last update from 10.1.2.2 on GigabitEthernet1.112, 00:22:57 ago
```

```
Routing Descriptor Blocks:
```

```
* 10.1.2.2, from 10.1.2.2, 00:22:57 ago, via GigabitEthernet1.112
```

```
Route metric is 3, traffic share count is 1
```

```
R1#show ip route vrf VPN_A 7.7.7.7
```

```
Routing Table: VPN_A
```

```
Routing entry for 7.7.7.7/32
```

```
Known via "ospf 100", distance 110, metric 20
```

```
Tag Complete, Path Length == 1, AS 100, , type extern 2, forward metric 1
```

```
Last update from 10.1.2.2 on GigabitEthernet1.112, 00:23:04 ago
```

```
Routing Descriptor Blocks:
```

```
* 10.1.2.2, from 10.1.2.2, 00:23:04 ago, via GigabitEthernet1.112
```

```
Route metric is 20, traffic share count is 1
```

```
Route tag 3489661028
```

```
R1#config t
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
R1(config)#router ospf 100 vrf VPN_A
```

```
R1(config-router)#no capability vrf-lite
```

```
R1(config-router)#end
```

```
R1#
```

```
%OSPF-5-ADJCHG: Process 100, Nbr 10.1.2.2 on GigabitEthernet1.112 from FULL to DOWN, Neighbor Down: Interface down or detached
```

```
%OSPF-5-ADJCHG: Process 100, Nbr 3.3.3.3 on GigabitEthernet1.13 from FULL to DOWN, Neighbor Down: Interface down or detached
```

```
%OSPF-5-ADJCHG: Process 100, Nbr 10.1.2.2 on GigabitEthernet1.112 from LOADING to FULL, Loading Done
```

```
%OSPF-5-ADJCHG: Process 100, Nbr 3.3.3.3 on GigabitEthernet1.13 from LOADING to FULL, Loading Done
```

```
R1#
```

```
R1#show ip route vrf VPN_A 10.7.7.7
```

```
Routing Table: VPN_A
```

```
% Subnet not in table
```

```
R1#show ip route vrf VPN_A 7.7.7.7
```

```
Routing Table: VPN_A
```

```
% Subnet not in table
```

```
R3#show ip route 10.7.7.7
```

```
Routing entry for 10.7.7.7/32
```

```
Known via "ospf 100", distance 110, metric 4, type inter area
```

```
Last update from 10.1.3.1 on GigabitEthernet1.13, 00:26:24 ago
```

```
Routing Descriptor Blocks:
```

```
* 10.1.3.1, from 10.1.2.2, 00:26:24 ago, via GigabitEthernet1.13
```

```
Route metric is 4, traffic share count is 1
```

```
R3#show ip route 7.7.7.7
Routing entry for 7.7.7.7/32
  Known via "ospf 100", distance 110, metric 20
  Tag Complete, Path Length == 1, AS 100, , type extern 2, forward metric 2
  Last update from 10.1.3.1 on GigabitEthernet1.13, 00:26:29 ago
  Routing Descriptor Blocks:
  * 10.1.3.1, from 10.1.2.2, 00:26:29 ago, via GigabitEthernet1.13
    Route metric is 20, traffic share count is 1
    Route tag 3489661028
```

R3 is still learning the routes, as the LSAs are flooded throughout the area, but R1 doesn't have them installed in the routing table. The result is that R3 forwards packets which are dropped by R1, causing a traffic black-hole.

```
R3#debug ip icmp
ICMP packet debugging is on
R3#ping 10.7.7.7

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.7.7.7, timeout is 2 seconds:
U
ICMP: dst (10.1.3.3) host unreachable rcv from 10.1.3.1.U
ICMP: dst (10.1.3.3) host unreachable rcv from 10.1.3.1.U
Success rate is 0 percent (0/5)
R3#
ICMP: dst (10.1.3.3) host unreachable rcv from 10.1.3.1
```

Likewise, the same would occur on R6 if XR2 reverted back to its default behavior of checking the Down Bit. This can be seen as follows.

```
RP/0/0/CPU0:XR2#show route vrf VPN_A 10.7.7.7
```

```
Tue May 12 01:02:55.237 UTC
```

```
Routing entry for 10.7.7.7/32
```

```
Known via "ospf 100", distance 110, metric 3, type inter area
```

```
Installed May 12 00:35:56.178 for 00:26:59
```

```
Routing Descriptor Blocks
```

```
10.19.20.19, from 19.19.19.19, via GigabitEthernet0/0/0/0.11920
```

```
Route metric is 3
```

```
No advertising protos.
```

```
RP/0/0/CPU0:XR2#show route vrf VPN_A 7.7.7.7
```

```
Tue May 12 01:02:59.867 UTC
```

```
Routing entry for 7.7.7.7/32
```

```
Known via "ospf 100", distance 110, metric 20
```

```
Tag 3489661028, type extern 2
```

```
Installed May 12 00:14:30.716 for 00:48:29
```

```
Routing Descriptor Blocks
```

```
10.19.20.19, from 19.19.19.19, via GigabitEthernet0/0/0/0.11920
```

```
Route metric is 20
```

```
No advertising protos.
```

```
RP/0/0/CPU0:XR2#config t
```

```
Tue May 12 01:04:19.321 UTC
```

```
RP/0/0/CPU0:XR2(config)#router ospf 100
```

```
RP/0/0/CPU0:XR2(config-ospf)#vrf VPN_A
```

```
RP/0/0/CPU0:XR2(config-ospf-vrf)#no disable-dn-bit-check
```

```
RP/0/0/CPU0:XR2(config-ospf-vrf)#commit
```

```
Tue May 12 01:04:22.321 UTC
```

```
RP/0/0/CPU0:May 12 01:04:22.681 : config[65710]: %MGBL-CONFIG-6-DB_COMMIT : Configuration committed by user 'admin'. Use 'show configuration commit changes 1000000110' to view the changes.
```

```
RP/0/0/CPU0:XR2#show route vrf VPN_A 10.7.7.7
```

```
Tue May 12 01:04:50.689 UTC
```

```
% Network not in table
```

```
RP/0/0/CPU0:XR2#show route vrf VPN_A 7.7.7.7
```

```
Tue May 12 01:04:55.049 UTC
```

```
% Network not in table
```

Now that XR2 no longer installs the routes in the routing table, traffic is blackholed as it transits this router.

```
R6#debug ip icmp
```

```
ICMP packet debugging is on
```

```
R6#ping 10.7.7.7
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.7.7.7, timeout is 2 seconds:
```

```
U
```

```
ICMP: dst (10.6.20.6) net unreachable rcv from 10.6.20.20 .U
```

```
ICMP: dst (10.6.20.6) net unreachable rcv from 10.6.20.20.U
```

```
Success rate is 0 percent (0/5)
```

« [MPLS L3 VPN and OSPF Domain-ID \(/workbook/view/service-provider-v4/task/mpls-l3-vpn-and-ospf-domain-id-Mjg2NQ%3D%3D\)](/workbook/view/service-provider-v4/task/mpls-l3-vpn-and-ospf-domain-id-Mjg2NQ%3D%3D) | [MPLS L3VPN Inter-AS Option A - Back-to-Back VRF Exchange \(/workbook/view/service-provider-v4/task/mpls-l3vpn-inter-as-option-a-back-to-back-vrf-exchange-Mjg2Nw%3D%3D\)](/workbook/view/service-provider-v4/task/mpls-l3vpn-inter-as-option-a-back-to-back-vrf-exchange-Mjg2Nw%3D%3D) »