MPLS L3 VPN and OSPF Domain-ID

« MPLS L3 VPN and OSPF Sham-Links (/workbook/view/service-provider-v4/task/mpls-l3-vpn-and-ospf-sham-links-Mjg2NA%3D%3D) | Multi-VRF CE (VRF Lite) (/workbook/view/service-provider-v4/task/multi-vrf-ce-vrf-lite-Mjg2Ng%3D%3D) »

Last updated: April 23, 2016

Note:

Initial Configuration & Diagrams: Load the initial configuration files for the section named **L3VPN Backdoor Links**, 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 **L3VPN Backdoor Links Diagram** in order to complete this task.

Task

- Create a new Loopback1 interface on R1 with the IP address 10.1.1.1/32.
- Create a new Loopback1 interface on XR2 with the IP address 10.20.20.20/32.
- Configure a VRF on R2 and XR1 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 links connecting to R1 and XR2 respectively.
- Configure OSPF routing for the VRF as follows:
  - Use Process-ID 100.
  - Enable OSPF Area 0 between R1 & R2.
  - Enable OSPF Area 0 between XR1 & XR2.
  - Enable OSPF Area 120 between R1 & XR2.
  - Advertise the Loopback0 network of R1 into OSPF Area 1.
  - Advertise the Loopback0 network of XR2 into OSPF Area 20.
  - Redistribute the Loopback1 networks of R1 and XR2 into OSPF.
  - Modify the link between R1 and XR2 to have an OSPF cost of 100.
  - Configure the OSPF Domain-ID on R2 and XR1 for the VRF aware process as type 0005 with value 0x000000640200.
- Configure BGP on R2 and XR1 as follows:
  - Use BGP AS 100.
  - R2 and XR1 should be iBGP peers for the VPNv4 Address Family.
  - Use their Loopback0 interfaces as the source of the BGP session.
  - Redistribute between BGP and the VRF aware OSPF process.
- Once complete, the following reachability should be achieved:
  - R1 and XR2 should have reachability to all of each other's networks.
  - Traffic between the Loopback0 networks of R1 and XR2 should prefer to use the MPLS network.

CONTENTS

- Traffic between the Loopback1 networks of R1 and XR2 should prefer to use the backdoor link.
- If either R1 or XR2 lose connectivity to the MPLS cloud or the backdoor link, traffic should be automatically be rerouted accordingly.

Configuration **Click to *collapse***

CONTENTS

```
R1:
interface Loopback1
 ip address 10.1.1.1 255.255.255.255
!
interface GigabitEthernet1.120
 ip ospf cost 100
!
ip prefix-list LOOPBACK1 seq 5 permit 10.1.1.1/32
!
route-map CONNECTED_TO_OSPF permit 10
 match ip address prefix-list LOOPBACK1
!
router ospf 100
 redistribute connected subnets route-map CONNECTED_TO_OSPF
 network 1.1.1.1 0.0.0.0 area 1
 network 10.1.2.1 0.0.0.0 area 0
 network 10.1.20.1 0.0.0.0 area 120


R2:
vrf definition VPN_A
 rd 100:1
 route-target export 100:1
 route-target import 100:1
 !
 address-family ipv4
 exit-address-family
!
interface GigabitEthernet1.12
 vrf forwarding VPN_A
 ip address 10.1.2.2 255.255.255.0
!
router ospf 100 vrf VPN_A
 redistribute bgp 100 subnets
 network 10.1.2.2 0.0.0.0 area 0
 domain-id type 0005 value 000000640200
!
router bgp 100
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 19.19.19.19 remote-as 100
 neighbor 19.19.19.19 update-source Loopback0
 !
 address-family vpnv4
  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 match internal external 1 external 2
 exit-address-family
!


XR1:
```

```
vrf VPN_A
 address-family ipv4 unicast
  import route-target 100:1
 !
  export route-target
   100:1
  !
 !
!
interface GigabitEthernet0/0/0/0.1920
 vrf VPN_A
 no ipv4 address
 ipv4 address 10.19.20.19 255.255.255.0
!
router bgp 100
 address-family vpnv4 unicast
 !
 neighbor 2.2.2.2
  remote-as 100
  update-source Loopback0
  address-family vpnv4 unicast
  !
!
vrf VPN_A
 rd 100:1
 address-family ipv4 unicast
   redistribute ospf 100
  !
 !
!
router ospf 100
 vrf VPN_A
  domain-id type 0005 value 000000640200
  redistribute bgp 100
  area 0
   interface GigabitEthernet0/0/0/0.1920
   !
  !
 !
!


XR2:
interface Loopback1
 ipv4 address 10.20.20.20 255.255.255.255
!
route-policy CONNECTED_TO_OSPF
 if destination in (10.20.20.20/32) then
    pass
 endif
end-policy
!
router ospf 100
 redistribute connected route-policy CONNECTED_TO_OSPF
  area 0
```

CONTENTS

```
 interface GigabitEthernet0/0/0/0.1920
 !
!
area 20
 interface Loopback0
 !
!
area 120
 interface GigabitEthernet0/0/0/0.120
   cost 100
  !
 !
!
```

Verification

As previously discussed, the OSPF Domain-ID is encoded as a BGP Extended Community when redistribution of OSPF and VPNv4 BGP occurs. Once redistribution of OSPF and VPNv4 BGP is complete, and VPNv4 routes are exchanged between PE routers, the OSPF Domain-ID of the received BGP routes is compared against the OSPF Domain-ID of the local OSPF process. If these values match, the MPLS network can be treated as the OSPF "Superbackbone", which is considered a hierarchy above Area 0. This allows the PE routers to be treated as ABRs instead of ASBRs, and encode routes that are being redistributed from VPNv4 BGP into OSPF as Network Summary LSAs (LSA Type 3) as opposed to External LSAs (LSA Type 5). Note that this behavior only takes place if the route in question was not already an External route to begin with.

The details about the specific values that are encoded in the OSPF Route Type, OSPF Domain ID, and OSPF Router ID fields of the VPNv4 BGP Extended Communities can be found in RFC 4577: OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) (http://www.ietf.org/rfc/rfc4577.txt).

In regular IOS, the Domain-ID is automatically inherited from the OSPF Process-ID number. In this example a Process-ID of decimal 100 equals a Domain-ID of hexadecimal 0x64. In IOS XR, the Domain-ID is not automatically set, which is why the command domain-id is needed. The actual encoding of the BGP extended community can be seen as follows.

CONTENTS ›

```
R2#show bgp vpnv4 unicast all 1.1.1.1/32
BGP routing table entry for 100:1:1.1.1.1/32, version 2
Paths: (1 available, best #1, table VPN_A)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    10.1.2.1 (via vrf VPN_A) from 0.0.0.0 (2.2.2.2)
      Origin incomplete, metric 2, localpref 100, weight 32768, valid, sourced, best
       Extended Community: RT:100:1 OSPF DOMAIN ID:0x0005:0x000000640200
         OSPF RT:0.0.0.0:3:0 OSPF ROUTER ID:10.1.2.2:0
      mpls labels in/out 32/nolabel
      rx pathid: 0, tx pathid: 0x0

R2#show bgp vpnv4 unicast all 20.20.20.20/32
BGP routing table entry for 100:1:20.20.20.20/32, version 14
Paths: (1 available, best #1, table VPN_A)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    19.19.19.19 (metric 4) (via default) from 19.19.19.19 (19.19.19.19)
      Origin incomplete, metric 2, localpref 100, valid, internal, best
       Extended Community: RT:100:1 OSPF DOMAIN ID:0x0005:0x000000640200
         OSPF RT:0.0.0.0:3:0 OSPF ROUTER ID:19.19.19.19:0
      mpls labels in/out nolabel/16007
      rx pathid: 0, tx pathid: 0x0

RP/0/0/CPU0:XR1#show bgp vpnv4 unicast vrf VPN_A 1.1.1.1/32
Mon May 11 01:11:36.580 UTC
BGP routing table entry for 1.1.1.1/32, Route Distinguisher: 100:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                  9           9
Last Modified: May 11 01:08:08.451 for 00:03:28
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local
    2.2.2.2 (metric 4) from 2.2.2.2 (2.2.2.2)
      Received Label 32
      Origin incomplete, metric 2, localpref 100, valid, internal, best, group-best, import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 9
       Extended community: OSPF domain-id:0x5:0x000000640200 OSPF route-type:0:3:0x0 OSPF router-id:10.1.2.2 RT:100:1
      Source VRF: VPN_A, Source Route Distinguisher: 100:1

RP/0/0/CPU0:XR1#show bgp vpnv4 unicast vrf VPN_A 20.20.20.20/32
Mon May 11 01:11:58.998 UTC
BGP routing table entry for 20.20.20.20/32, Route Distinguisher: 100:1
Versions:
  Process           bRIB/RIB  SendTblVer
```

```
  Speaker                14          14
    Local Label: 16007
Last Modified: May 11 01:08:36.451 for 00:03:22
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    2.2.2.2
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    2.2.2.2
  Local
    10.19.20.20 from 0.0.0.0 (19.19.19.19)
      Origin incomplete, metric 2, localpref 100, weight 32768, valid, redistributed, best, group-best, import-candidate
      Received Path ID 0, Local Path ID 1, version 14
       Extended community: OSPF domain-id:0x5:0x000000640200 OSPF route-type:0:3:0x0 OSPF router-id:19.19.19.19 RT:100:1
```

Since the OSPF Domain-ID matches between the local OSPF process and the VPNv4 BGP route, and the OSPF Route Type is 3 (meaning the routes were Inter-Area to begin with), when VPNv4 BGP to OSPF redistribution occurs these routes should be encoded in the database as Network Summary LSAs (LSA Type 3). In other words, the PE routers will be seen as ABRs that are advertising Inter-Area routes.

CONTENTS

```
R1#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

     10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
O IA     10.19.20.0/24 [110/2] via 10.1.2.2, 00:04:48, GigabitEthernet1.12
O E2     10.20.20.20/32
           [110/20] via 10.1.20.20, 00:04:21, GigabitEthernet1.120
     20.0.0.0/32 is subnetted, 1 subnets
 O IA     20.20.20.20 [110/3] via 10.1.2.2, 00:04:20, GigabitEthernet1.12
```

```
R1#show ip ospf database

            OSPF Router with ID (10.1.1.1) (Process ID 100)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        401         0x80000002 0x00F601 1
10.1.2.2        10.1.2.2        402         0x80000002 0x00E210 1

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.2.2        10.1.2.2        402         0x80000001 0x00748B

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
1.1.1.1         10.1.1.1        462         0x80000001 0x00F535
10.1.20.0       10.1.1.1        462         0x80000001 0x009A12
10.19.20.0      10.1.2.2        320         0x80000001 0x004B30
20.20.20.20     10.1.2.2        292         0x80000001 0x00FD5D

                Summary ASB Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
 20.20.20.20     10.1.1.1        292         0x80000001 0x005C1E
<snip>
```

```
RP/0/0/CPU0:XR2#show route ospf
Mon May 11 01:14:13.050 UTC
```

```
 O IA 1.1.1.1/32 [110/3] via 10.19.20.19, 00:05:36, GigabitEthernet0/0/0/0.1920
O E2 10.1.1.1/32 [110/20] via 10.1.20.1, 00:05:36, GigabitEthernet0/0/0/0.120
O IA 10.1.2.0/24 [110/2] via 10.19.20.19, 00:05:36, GigabitEthernet0/0/0/0.1920


RP/0/0/CPU0:XR2#show ospf database
Mon May 11 01:14:42.108 UTC



            OSPF Router with ID (20.20.20.20) (Process ID 100)


                Router Link States (Area 0)


Link ID         ADV Router      Age         Seq#       Checksum Link count
19.19.19.19     19.19.19.19     366         0x80000002 0x00fe0f 1
20.20.20.20     20.20.20.20     365         0x80000003 0x00be45 1


                Net Link States (Area 0)


Link ID         ADV Router      Age         Seq#       Checksum
10.19.20.19     19.19.19.19     366         0x80000001 0x00d438


                Summary Net Link States (Area 0)


Link ID         ADV Router      Age         Seq#       Checksum
 1.1.1.1         19.19.19.19     393         0x80000001 0x00abbe
10.1.2.0        19.19.19.19     393         0x80000001 0x002b37
10.1.20.0       20.20.20.20     369         0x80000003 0x00abbb
20.20.20.20     20.20.20.20     375         0x80000001 0x009dfd


                Summary ASB Link States (Area 0)


Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        20.20.20.20     369         0x80000001 0x006911


<snip>
```

When R1 and XR2 make their path selection decision, they now see two possible paths to each other's Loopback0 networks, one via the MPLS L3VPN PE and one via the backdoor link.

CONTENTS >

```
R1#show ip ospf database summary 20.20.20.20

        OSPF Router with ID (10.1.1.1) (Process ID 100)

            Summary Net Link States (Area 0)

LS age: 441
Options: (No TOS-capability, DC, Downward)
LS Type: Summary Links(Network)
Link State ID: 20.20.20.20 (summary Network Number)
 Advertising Router: 10.1.2.2
LS Seq Number: 80000001
Checksum: 0xFD5D
Length: 28
Network Mask: /32
    MTID: 0         Metric: 2



            Summary Net Link States (Area 1)

LS age: 440
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 20.20.20.20 (summary Network Number)
Advertising Router: 10.1.1.1
LS Seq Number: 80000001
Checksum: 0x9C40
Length: 28
Network Mask: /32
    MTID: 0         Metric: 3



            Summary Net Link States (Area 120)

LS age: 440
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 20.20.20.20 (summary Network Number)
Advertising Router: 10.1.1.1
LS Seq Number: 80000001
Checksum: 0x9C40
Length: 28
Network Mask: /32
    MTID: 0         Metric: 3


LS age: 450
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 20.20.20.20 (summary Network Number)
 Advertising Router: 20.20.20.20
LS Seq Number: 80000001
Checksum: 0x9DFD
Length: 28
```

```
Network Mask: /32
    MTID: 0         Metric: 1




RP/0/0/CPU0:XR2#show ospf database summary 1.1.1.1
Mon May 11 01:18:36.162 UTC


           OSPF Router with ID (20.20.20.20) (Process ID 100)


               Summary Net Link States (Area 0)


  Routing Bit Set on this LSA
  LS age: 627
  Options: (No TOS-capability, DC, DN)
  LS Type: Summary Links (Network)
  Link State ID: 1.1.1.1 (Summary Network Number)
   Advertising Router: 19.19.19.19
  LS Seq Number: 80000001
  Checksum: 0xabbe
  Length: 28
  Network Mask: /32
        TOS: 0  Metric: 2



               Summary Net Link States (Area 20)


  LS age: 599
  Options: (No TOS-capability, DC)
  LS Type: Summary Links (Network)
  Link State ID: 1.1.1.1 (Summary Network Number)
  Advertising Router: 20.20.20.20
  LS Seq Number: 80000001
  Checksum: 0x1fc6
  Length: 28
  Network Mask: /32
        TOS: 0  Metric: 3



               Summary Net Link States (Area 120)


  LS age: 770
  Options: (No TOS-capability, DC)
  LS Type: Summary Links (Network)
  Link State ID: 1.1.1.1 (Summary Network Number)
   Advertising Router: 10.1.1.1
  LS Seq Number: 80000001
  Checksum: 0xf535
  Length: 28
  Network Mask: /32
        TOS: 0  Metric: 1
```

```
LS age: 599

Options: (No TOS-capability, DC)

LS Type: Summary Links (Network)

Link State ID: 1.1.1.1 (Summary Network Number)

Advertising Router: 20.20.20.20

LS Seq Number: 80000001

Checksum: 0x1fc6

Length: 28

Network Mask: /32

      TOS: 0  Metric: 3
```

Since these routes are now both the same type (i.e. Type 3 LSA vs. Type 3 LSA) the OSPF cost is the tie breaker. On R1 the shortest path to the ABR 10.1.2.2 (R2) is closer than the shortest path to the ABR 20.20.20.20 (XR2) because the cost of the link to XR2 was increased to 100. Likewise XR2 prefers to route via the ABR 19.19.19.19 (XR1) due to the lower cost. The final result is that traffic between these links follows the MPLS L3VPN path.

CONTENTS ›

```
R1#show ip ospf border-routers

            OSPF Router with ID (10.1.1.1) (Process ID 100)



                Base Topology (MTID 0)


Internal Router Routing Table
Codes: i - Intra-area route, I - Inter-area route

i 20.20.20.20 [ 100 ] via 10.1.20.20, GigabitEthernet1.120, ABR/ASBR, Area 120, SPF 2
i 10.1.2.2 [ 1 ] via 10.1.2.2, GigabitEthernet1.12, ABR/ASBR, Area 0, SPF 3



RP/0/0/CPU0:XR2#show ospf border-routers
Mon May 11 01:19:44.477 UTC

OSPF 100 Internal Routing Table

Codes: i - Intra-area route, I - Inter-area route

i 10.1.1.1 [ 100 ] via 10.1.20.1, GigabitEthernet0/0/0/0.120, ABR/ASBR , Area 120, SPF 6
i 19.19.19.19 [ 1 ] via 10.19.20.19, GigabitEthernet0/0/0/0.1920, ABR/ASBR , Area 0, SPF 4



R1#traceroute 20.20.20.20 source 1.1.1.1
Type escape sequence to abort.
Tracing the route to 20.20.20.20
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.2.2 4 msec 1 msec 1 msec
  2 20.2.4.4 [MPLS: Labels 21/16007 Exp 0] 14 msec 11 msec 12 msec
  3 20.4.6.6 [MPLS: Labels 22/16007 Exp 0] 12 msec 12 msec 20 msec
  4 20.6.19.19 [MPLS: Label 16007 Exp 0] 20 msec 20 msec 23 msec
  5 10.19.20.20 17 msec *  14 msec



RP/0/0/CPU0:XR2#traceroute 1.1.1.1 source 20.20.20.20
Mon May 11 01:22:30.166 UTC

Type escape sequence to abort.
Tracing the route to 1.1.1.1

 1  10.19.20.19 0 msec  0 msec  0 msec
 2  20.5.19.5 [MPLS: Labels 26/32 Exp 0] 9 msec  9 msec  9 msec
 3  20.4.5.4 [MPLS: Labels 19/32 Exp 0] 9 msec  9 msec  9 msec
 4  10.1.2.2 [MPLS: Label 32 Exp 0] 9 msec  9 msec  9 msec
 5  10.1.2.1 9 msec  *  19 msec
```

Both R1 and XR2 install the path to each other's Loopback0 with a cost of 2. 1 to reach the ABR (as seen in the border-router output), plus 2 as advertised by each ABR's Type-3 LSA.

Note that this conversion process with the OSPF Domain-ID only occurs when the route was first redistributed as an Intra-Area OSPF or Inter-Area OSPF route to being with. In this specific example when R2 redistributes OSPF into BGP, it is learning the route 1.1.1.1/32 as an OSPF Inter-Area route,

but it’s learning the route 10.1.1.1/32 as an OSPF External route.

```
R2#show ip route vrf VPN_A ospf

Routing Table: VPN_A
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

      1.0.0.0/32 is subnetted, 1 subnets
O IA     1.1.1.1 [110/2] via 10.1.2.1, 00:19:38, GigabitEthernet1.12
      10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks
O E2     10.1.1.1/32 [110/20] via 10.1.2.1, 00:19:38, GigabitEthernet1.12
O IA     10.1.20.0/24 [110/101] via 10.1.2.1, 00:19:38, GigabitEthernet1.12
O E2     10.20.20.20/32 [110/20] via 10.1.2.1, 00:17:48, GigabitEthernet1.12
```

This means that when OSPF to BGP redistribution occurs, R2 will encode the Inter-Area route 1.1.1.1/32 with an OSPF Route Type 3 (meaning the route was Inter-Area to start), but will encode the External route 10.1.1.1/32 with an OSPF Route Type 5 (meaning the route was External to start).

CONTENTS

```
R2#show bgp vpnv4 unicast all 1.1.1.1/32
BGP routing table entry for 100:1:1.1.1.1/32, version 2
Paths: (1 available, best #1, table VPN_A)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    10.1.2.1 (via vrf VPN_A) from 0.0.0.0 (2.2.2.2)
      Origin incomplete, metric 2, localpref 100, weight 32768, valid, sourced, best
      Extended Community: RT:100:1 OSPF DOMAIN ID:0x0005:0x000000640200
        OSPF RT:0.0.0.0:3:0 OSPF ROUTER ID:10.1.2.2:0
      mpls labels in/out 32/nolabel
      rx pathid: 0, tx pathid: 0x0



R2#show bgp vpnv4 unicast all 10.1.1.1/32
BGP routing table entry for 100:1:10.1.1.1/32, version 3
Paths: (2 available, best #2, table VPN_A)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    19.19.19.19 (metric 4) (via default) from 19.19.19.19 (19.19.19.19)
      Origin incomplete, metric 20, localpref 100, valid, internal
      Extended Community: RT:100:1 OSPF DOMAIN ID:0x0005:0x000000640200
        OSPF RT:0.0.0.0:5:1 OSPF ROUTER ID:19.19.19.19:0
      mpls labels in/out 24/16016
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  Local
    10.1.2.1 (via vrf VPN_A) from 0.0.0.0 (2.2.2.2)
      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:10.1.2.2:0
      mpls labels in/out 24/nolabel
      rx pathid: 0, tx pathid: 0x0
```

Notice the difference in the Route-Type Extended Community Attribute of both of these routes. R2 sees 1.1.1.1/32 as an Inter-Area route, and 10.1.1.1/32 as an external route, thus the routes are encoded with Route-Type 0.0.0.0:**3** and 0.0.0.0:**5** respectively. When the routes are learned by the remote PE (XR1) and redistributed from VPNv4 BGP back into OSPF, only the route with both the matching Domain ID and the Route Type of 3 can be advertised as a Network Summary LSA. The External route will remain an External route regardless of any other settings.

In this example this only becomes apparent when the backdoor link between R1 and XR2 is disabled, due to the fact that the PE routers will not re-originate a Type 5 LSA that someone else in their area is already originating.

Below we see that R1 and XR2 prefer to use the backdoor link to reach their Loopback1 networks, which are the External routes that came from **redistribute connected**.

```
R1#show ip route 10.20.20.20

Routing entry for 10.20.20.20/32

  Known via "ospf 100", distance 110, metric 20, type extern 2, forward metric 100

  Last update from 10.1.20.20 on GigabitEthernet1.120, 00:54:14 ago

  Routing Descriptor Blocks:

  * 10.1.20.20, from 20.20.20.20, 00:54:14 ago, via GigabitEthernet1.120

      Route metric is 20, traffic share count is 1


R1#traceroute 10.20.20.20 source 10.1.1.1


Type escape sequence to abort.

Tracing the route to 10.20.20.20


1 10.1.20.20 4 msec *   0 msec


RP/0/0/CPU0:XR2#show route 10.1.1.1

Mon May 11 02:03:13.408 UTC


Routing entry for 10.1.1.1/32

  Known via "ospf 100", distance 110, metric 20, type extern 2

  Installed May 11 01:08:36.983 for 00:54:36

  Routing Descriptor Blocks

    10.1.20.1, from 10.1.1.1, via GigabitEthernet0/0/0/0.120

      Route metric is 20

  No advertising protos.


RP/0/0/CPU0:XR2#traceroute 10.1.1.1 source 10.20.20.20

Fri Mar 16 16:09:35.648 UTC


Type escape sequence to abort.

Tracing the route to 10.1.1.1


1    10.1.20.1  5 msec    *   2 msec
```

Once the backdoor link is disabled, then the external route can transit over the MPLS L3VPN.

CONTENTS >

```
RP/0/0/CPU0:XR2#config
Mon May 11 02:03:59.535 UTC
RP/0/0/CPU0:XR2(config)#interface GigabitEthernet0/0/0/0.120
RP/0/0/CPU0:XR2(config-subif)#shutdown
RP/0/0/CPU0:XR2(config-subif)#exit
RP/0/0/CPU0:XR2(config)#commit
Mon May 11 02:04:17.614 UTC
RP/0/0/CPU0:May 11 02:04:17.654 : ospf[1015]: %ROUTING-OSPF-5-ADJCHG : Process 100, Nbr 10.1.1.1 on GigabitEthernet0/0/0/0.120 in area 120 f
rom FULL to DOWN, Neighbor Down: interface down or detached, vrf default vrfid 0x60000000
RP/0/0/CPU0:May 11 02:04:17.914 : config[65710]: %MGBL-CONFIG-6-DB_COMMIT : Configuration committed by user 'admin'. Use 'show configuration
 commit changes 1000000104' to view the changes.
RP/0/0/CPU0:XR2(config)#
RP/0/0/CPU0:XR2(config)#end
RP/0/0/CPU0:May 11 02:04:21.683 : config[65710]: %MGBL-SYS-5-CONFIG_I : Configured from console by admin
RP/0/0/CPU0:XR2#


RP/0/0/CPU0:XR2#show route 10.1.1.1
Mon May 11 02:06:06.256 UTC


Routing entry for 10.1.1.1/32
  Known via "ospf 100", distance 110, metric 20
   Tag 3489661028, type extern 2
  Installed May 11 02:04:18.114 for 00:01:48
  Routing Descriptor Blocks
     10.19.20.19, from 19.19.19.19, via GigabitEthernet0/0/0/0.1920
       Route metric is 20
  No advertising protos.


RP/0/0/CPU0:XR2#traceroute 10.1.1.1 source 10.20.20.20
Mon May 11 02:06:43.404 UTC


Type escape sequence to abort.
Tracing the route to 10.1.1.1


 1  10.19.20.19 0 msec  0 msec  0 msec
 2  20.6.19.6 [MPLS: Labels 25/24 Exp 0] 9 msec  9 msec  0 msec
 3  20.3.6.3 [MPLS: Labels 20/24 Exp 0] 9 msec  0 msec  9 msec
 4  10.1.2.2 [MPLS: Label 24 Exp 0] 9 msec  9 msec  0 msec
 5  10.1.2.1 9 msec  *  9 msec


R1#show ip route 10.20.20.20
Routing entry for 10.20.20.20/32
 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 GigabitEthernet1.12, 00:02:04 ago
 Routing Descriptor Blocks:
 * 10.1.2.2, from 10.1.2.2, 00:02:04 ago, via GigabitEthernet1.12
     Route metric is 20, traffic share count is 1
     Route tag 3489661028


R1#traceroute 10.20.20.20 source 10.1.1.1
Type escape sequence to abort.
Tracing the route to 10.20.20.20
```

```
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.2.2 5 msec 1 msec 1 msec
  2 20.2.3.3 [MPLS: Labels 24/16017 Exp 0] 11 msec 12 msec 12 msec
  3 20.3.6.6 [MPLS: Labels 22/16017 Exp 0] 12 msec 12 msec 12 msec
  4 20.6.19.19 [MPLS: Label 16017 Exp 0] 16 msec 20 msec 20 msec
  5 10.19.20.20 16 msec *  14 msec
```

Notice that the routes now have a Route-Tag value that was automatically generated. This is a 32 bit tag which is used for loop prevention, similar to the function provided by the OSPF Down Bit, however this tag is only present in External routes injected into the OSPF domain by redistribution from MP-BGP into OSPF by a PE router. When using 16 bit BGP ASNs, as such is the case in our example, the last 16 bits of the 32 bit tag are used to encode the BGP ASN.

XR1 receives a VPNv4 update for 10.1.1.1/32. The route has been advertised by R2, and the Route-Type Extended Community value has been encoded as External (Type-5) by R2.

```
RP/0/0/CPU0:XR1#show bgp vpnv4 unicast rd 100:1 10.1.1.1/32
Mon May 11 22:51:35.787 UTC
BGP routing table entry for 10.1.1.1/32, Route Distinguisher: 100:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 28          28
Last Modified: May 11 02:27:54.451 for 20:23:41
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
  Local
    2.2.2.2 (metric 4) from 2.2.2.2 (2.2.2.2)
      Received Label 44
      Origin incomplete, metric 20, localpref 100, valid, internal, best, group-best, import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 28
      Extended community: OSPF domain-id:0x5:0x000000640200 OSPF route-type:0:5:0x1 OSPF router-id:10.1.2.2 RT:100:1
      Source VRF: VPN_A, Source Route Distinguisher: 100:1
```

This route is then installed in VRF_A by XR1 as a BGP route, causing it to be redistributed into OSPF.

```
RP/0/0/CPU0:XR1#show route vrf VPN_A 10.1.1.1
Mon May 11 22:51:44.876 UTC

Routing entry for 10.1.1.1/32
  Known via "bgp 100", distance 200, metric 20, type internal
  Installed May 11 02:27:54.526 for 20:23:50
  Routing Descriptor Blocks
    2.2.2.2, from 2.2.2.2
      Nexthop in Vrf: "default", Table: "default", IPv4 Unicast, Table Id: 0xe0000000
      Route metric is 20
  No advertising protos.
```

Since this was an external route that originated from an OSPF domain, the route is tagged.

CONTENTS >

```
RP/0/0/CPU0:XR1#show ospf 100 vrf VPN_A database external 10.1.1.1

Mon May 11 23:00:14.191 UTC



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


                Type-5 AS External Link States


  LS age: 1502

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

  LS Type: AS External Link

  Link State ID: 10.1.1.1 (External Network Number)

   Advertising Router: 19.19.19.19

  LS Seq Number: 80000027

  Checksum: 0x7feb

  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
```

If this OSPF route is then advertised over the backdoor link, the tag will prevent R2 from installing the route, much like the Down Bit.

```
> 3489661028 is the decimal representation of hex value 0xD0000064. The last 2 bytes (16 bits) of this hex value, 0x0064, are the hex repres

entation of decimal value 100 - the BGP ASN.
```

More information about the tagging mechanism can be found on Section 4.2.5.2 of RFC 4577 (http://tools.ietf.org/html/rfc4577#section-4.2.5.2).

« MPLS L3 VPN and OSPF Sham-Links (/workbook/view/service-provider-v4/task/mpls-l3-vpn-and-ospf-sham-links-Mjg2NA%3D%3D) | Multi-VRF CE (VRF Lite) (/workbook/view/service-provider-v4/task/multi-vrf-ce-vrf-lite-Mjg2Ng%3D%3D) »

CONTENTS