MPLS L3 VPN with OSPF

« MPLS L3 VPN with EIGRP (/workbook/view/service-provider-v4/task/mpls-l3-vpn-with-eigrp-Mjg1Nw%3D%3D) | MPLS L3 VPN with BGP (/workbook/view/service-provider-v4/task/mpls-l3-vpn-with-bgp-Mjg1OQ%3D%3D) »

Last updated: April 23, 2016

CONTENTS

Note:

Initial Configuration & Diagrams: Load the initial configuration files for the section named **OSPFv2 and LDP**, 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 **Base IPv4 Diagram** in order to complete this task.

Task

- 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.
  - Advertise the Loopback0 networks of R1 & XR2 into OSPF Area 0.
- 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 OSPF.
- Once complete R1 and XR2 should have reachability to each other’s Loopback0 interfaces and PE-CE links.

Configuration **Click to *collapse***

```
R1:
router ospf 100
 network 0.0.0.0 255.255.255.255 area 0

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
 network 0.0.0.0 255.255.255.255 area 0
 redistribute bgp 100 subnets
!
router bgp 100
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 19.19.19.19 remote-as 100
 neighbor 19.19.19.19 update-source Loopback0
 !
 address-family ipv4
 exit-address-family
 !
 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
 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
 ipv4 address 10.19.20.19 255.255.255.0
```

CONTENTS

```
!
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
  redistribute bgp 100
  area 0
   interface GigabitEthernet0/0/0/0.1920
  !
 !
!
end

XR2:
router ospf 100
 area 0
  interface Loopback0
  !
  interface GigabitEthernet0/0/0/0.1920
  !
 !
!
```


Verification

In regular IOS OSPF requires one unique process for the global table and each subsequent VRF table. However in IOS XR multiple VRF tables can share the same OSPF process, similar to how RIP or EIGRP works in IOS XR.

The first step in verifying that this design works is to ensure that the OSPF adjacencies are functional. Since we now have both a global OSPF process for the core of the MPLS network and an OSPF process for VRF VPN_A, R2 should reference the unique OSPF Process-ID during verification to clarify which attributes apply to the global table and which apply to the VRF table, as seen below.

OSPF Process-ID 1 is used for the global table, while 100 is used for VRF VPN_A. If we were to look at the **show ip ospf neighbor** command without referencing a PID, it is very difficult to determine which options apply to the global table vs. the VRF. Hence the verification should include the PID.

```
R2#show ip ospf neighbor


Neighbor ID     Pri   State           Dead Time   Address         Interface
4.4.4.4           1   FULL/DR         00:00:39    20.2.4.4        GigabitEthernet1.24
3.3.3.3           1   FULL/DR         00:00:34    20.2.3.3        GigabitEthernet1.23
 1.1.1.1           1   FULL/BDR        00:00:30    10.1.2.1        GigabitEthernet1.12


R2#show ip ospf 1 neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
4.4.4.4           1   FULL/DR         00:00:39    20.2.4.4        GigabitEthernet1.24
3.3.3.3           1   FULL/DR         00:00:34    20.2.3.3        GigabitEthernet1.23



R2#show ip ospf 100 neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
 1.1.1.1           1   FULL/BDR        00:00:34    10.1.2.1        GigabitEthernet1.12
```

In IOS XR this verification is clearer as the VRF's name is referenced in the show commands.

```
RP/0/0/CPU0:XR1#show ospf vrf VPN_A neighbor
Sun May  3 23:28:38.935 UTC


* Indicates MADJ interface


Neighbors for OSPF 100, VRF VPN_A


Neighbor ID     Pri   State           Dead Time   Address         Interface
 20.20.20.20     1     FULL/BDR        00:00:34    10.19.20.20     GigabitEthernet0/0/0/0.1920
    Neighbor is up for 00:06:26


Total neighbor count: 1
```

Since the PE routers participate in the same area as the CE router's advertisements, the PEs should be learning their attached CE's routes as OSPF Intra-Area prefixes.

CONTENTS ›

```
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        1.1.1.1 [110/2] via 10.1.2.1, 00:10:01, GigabitEthernet1.12

RP/0/0/CPU0:XR1#show route vrf VPN_A ospf
Sun May  3 23:30:08.239 UTC

O    20.20.20.20/32 [110/2] via 10.19.20.20, 00:07:25, GigabitEthernet0/0/0.1920
```

OSPF is then redistributed from the PE's VRF aware OSPF process into VPNv4 BGP.

```
R2#show bgp vpnv4 unicast all
BGP table version is 66, 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: 100:1 (default for vrf VPN_A)
 *>  1.1.1.1/32       10.1.2.1                 2         32768 ?
 *>  10.1.2.0/24      0.0.0.0                  0         32768 ?
 *>i 10.19.20.0/24    19.19.19.19              0    100      0 ?
 *>i 20.20.20.20/32   19.19.19.19              2    100      0 ?

RP/0/0/CPU0:XR1#show bgp vpnv4 unicast
```

Sun May 3 23:31:11.284 UTC BGP router identifier 19.19.19.19, local AS number 100 BGP generic scan interval 60 secs BGP table state: Active Table ID: 0x0 RD version: 0 BGP main routing table version 66 BGP scan interval 60 secs

CONTENTS

```
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network            Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf VPN_A)
*>i1.1.1.1/32         2.2.2.2                  2    100      0 ?
*>i10.1.2.0/24        2.2.2.2                  0    100      0 ?
 *> 10.19.20.0/24      0.0.0.0                  0         32768 ?
 *> 20.20.20.20/32     10.19.20.20              2         32768 ?

Processed 4 prefixes, 4 paths
```

Like EIGRP, certain attributes of the routes are maintained and encoded as BGP extended communities. This can be seen in the details of the VPNv4 tables of the PE routers below. Note that in addition to the next-hop and the VPN label values, which are ultimately used to build the label stack in the data plane, OSPF specific attributes such as the OSPF Domain-ID, Route Type, and Router-ID are encoded.

```
R2#show bgp vpnv4 unicast all 1.1.1.1/32
BGP routing table entry for 100:1:1.1.1.1/32, version 62
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:0.0.0.0:2:0 OSPF ROUTER ID:10.1.2.2:0
      mpls labels in/out 37/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 7
Paths: (1 available, best #1, table VPN_A)
    Not advertised to any peer
    Local
        19.19.19.19  (metric 4) from 19.19.19.19 (19.19.19.19)
            Origin incomplete, metric 2, localpref 100, valid, internal, best
            Extended Community: RT:100:1 OSPF RT:0.0.0.0:2:0
              OSPF ROUTER ID:19.19.19.19:0
            mpls labels in/out nolabel/16013

RP/0/0/CPU0:XR1#show bgp vpnv4 unicast vrf VPN_A 1.1.1.1/32
Sun May  3 23:33:27.475 UTC
BGP routing table entry for 1.1.1.1/32, Route Distinguisher: 100:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 62          62
Last Modified: May  3 23:19:23.451 for 00:14:04
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 37
      Origin incomplete, metric 2, localpref 100, valid, internal, best, group-best, import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 62
      Extended community: OSPF domain-id:0x5:0x000000640200 OSPF route-type:0:2:0x0 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
Sun May  3 23:34:04.582 UTC
BGP routing table entry for 20.20.20.20/32, Route Distinguisher: 100:1
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 66          66
    Local Label: 16001
Last Modified: May  3 23:22:43.451 for 00:11:21
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 66

    Extended community: OSPF route-type:0:1:0x0 OSPF router-id:19.19.19.19 RT:100:1
```

The most notable of these encoded attributes is the OSPF Domain-ID, which is used to determine if the MPLS VPNv4 BGP network should be considered as the OSPF “super backbone”, which is treated as one hierarchy above OSPF area 0. In the case that the OSPF Domain-ID of a received VPNv4 route matches the OSPF Domain-ID of the local VRF aware OSPF process, then the OSPF Route-Type is examined to determine how the OSPF LSA should be encoded in the database when BGP to OSPF redistribution occurs. In other words, the OSPF Domain-ID makes the determination if the PE router should be treated as an OSPF ABR, which originates Type-3 Network Summary LSAs (Inter Area Routes) or an OSPF ASBR, which would originate Type-5 External LSAs for the redistributed routes, regardless if the route-type attribute is set to 1 or 2 (for internal routes).

Note that in the case of regular IOS the OSPF Domain-ID is automatically encoded from the OSPF Process-ID, but in IOS XR it is not - following RFC 4577 more closely than IOS. RFC 4577 states that each OSPF instance MUST be associated with one or more Domain IDs - which MUST be configurable - and the default value (if none is configured) SHOULD be NULL. As can be observed, IOS has automatically configured a Domain-ID value which is not NULL. However, IOS-XR has defaulted to NULL. Additionally, the RFC states that when the Domain-ID value is set to NULL, it is not necesary to include the value in the BGP update. Once again, IOS-XR follows the RFC closely and does not attach the Domain-ID community to the BGP update as the value is NULL.

R2 is currently setting the Domain-ID as follows:

```
DOMAIN ID:0x0005:0x000000640200
```

The first 4 bytes, 0x00000064, can be decoded into decimal to derive the OSPF process ID (100).

The OSPF Route Type Extended Communities Attribute is encoded as: 'Type-Field - Area Number - Route-Type - Options'. The table below describes the possible values found in the route-type portion of the extended community value.

Encoding	Route Type
1 or 2	Intra-Area routes
3	Inter-Area routes
5	External routes
7	NSSA routes

More details can be found in RFC-4577 (http://tools.ietf.org/html/rfc4577)

This means that in the current network setup the OSPF Domain-IDs do not match (R2 set it as DOMAIN ID:0x0005:0x000000640200, and XR1 as NULL), even though the OSPF Process-IDs are the same. The final result of this is that the CE routers see each other’s routes as Type-5 External LSAs, not Inter-Area routes.

```
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, 3 subnets, 2 masks
O E2     10.19.20.0/24 [110/1] via 10.1.2.2, 00:28:27, GigabitEthernet1.12
      20.0.0.0/32 is subnetted, 1 subnets
O E2     20.20.20.20 [110/2] via 10.1.2.2, 00:27:13, GigabitEthernet1.12


RP/0/0/CPU0:XR2#show route ipv4 ospf
Mon May  4 00:07:10.387 UTC

O E2 1.1.1.1/32 [110/2] via 10.19.20.19, 00:44:27, GigabitEthernet0/0/0/0.1920
O E2 10.1.2.0/24 [110/1] via 10.19.20.19, 00:44:27, GigabitEthernet0/0/0/0.1920
```

In the OSPF database the CE routers view the PE routers as ASBRs.

```
R1#show ip ospf database


          OSPF Router with ID (1.1.1.1) (Process ID 100)


              Router Link States (Area 0)


Link ID         ADV Router      Age         Seq#       Checksum Link count

1.1.1.1         1.1.1.1         930         0x80000003 0x00B343 2

10.1.2.2        10.1.2.2        906         0x80000003 0x00E011 1


              Net Link States (Area 0)


Link ID         ADV Router      Age         Seq#       Checksum

10.1.2.2        10.1.2.2        906         0x80000002 0x000FF8


              Type-5 AS External Link States


Link ID         ADV Router      Age         Seq#       Checksum Tag

 10.19.20.0      10.1.2.2        906         0x80000002 0x002A92 3489661028

 20.20.20.20     10.1.2.2        645         0x80000002 0x00DCBF 3489661028


RP/0/0/CPU0:XR2#show ospf database

Mon May  4 00:08:07.463 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     777         0x80000003 0x00fc10 1

20.20.20.20     20.20.20.20     714         0x80000003 0x005451 2


               Net Link States (Area 0)


Link ID         ADV Router      Age         Seq#       Checksum

10.19.20.19     19.19.19.19     777         0x80000002 0x00d239


               Type-5 AS External Link States


Link ID         ADV Router      Age         Seq#       Checksum Tag

 1.1.1.1         19.19.19.19     777         0x80000002 0x008a21 3489661028

 10.1.2.0        19.19.19.19     777         0x80000002 0x000a99 3489661028
```

Note that this does not affect connectivity, as the customer sites still have reachability to each other as seen below, it simply affects how the OSPF path selection occurs.

```
R1#ping 20.20.20.20 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 20.20.20.20, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/12 ms

R1#ping 10.19.20.20

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.19.20.20, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/6/12 ms
```

If the OSPF Domain-IDs would match, the PE routers would appear as ABRs instead of ASBRs, and the CEs would see each other’s routes as Type-3 Network Summary LSAs (Inter-Area routes). This can be accomplished in this case by manually defining the Domain-ID on IOS XR to match what the regular IOS process is encoding, as seen below. Note that the 0x00000064 portion of the Domain-ID is the OSPF Process-ID 100 represented in hexadecimal.

```
RP/0/0/CPU0:XR1#show bgp vpnv4 unicast vrf VPN_A 1.1.1.1/32
Mon May  4 00:34:18.435 UTC
BGP routing table entry for 1.1.1.1/32, Route Distinguisher: 100:1
Versions:
 Process           bRIB/RIB  SendTblVer
 Speaker                 62          62
Last Modified: May  3 23:19:23.451 for 01:14:55
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 37
     Origin incomplete, metric 2, localpref 100, valid, internal, best, group-best, import-candidate, imported
     Received Path ID 0, Local Path ID 1, version 62
     Extended community: OSPF domain-id:0x5:0x000000640200 OSPF route-type:0:2:0x0 OSPF router-id:10.1.2.2 RT:100:1
     Source VRF: VPN_A, Source Route Distinguisher: 100:1



RP/0/0/CPU0:XR1#
RP/0/0/CPU0:XR1#conf t
Mon May  4 00:35:41.039 UTC
RP/0/0/CPU0:XR1(config)#router ospf 100
RP/0/0/CPU0:XR1(config-ospf)#vrf VPN_A
RP/0/0/CPU0:XR1(config-ospf-vrf)#domain-id type 0005 value 000000640200
RP/0/0/CPU0:XR1(config-ospf-vrf)#commit
Mon May  4 00:36:00.778 UTC
RP/0/0/CPU0:May  4 00:36:01.088 : config[65813]: %MGBL-CONFIG-6-DB_COMMIT : Configuration committed by user 'admin'. Use 'show configuration
 commit changes 1000000159' to view the changes.
RP/0/0/CPU0:XR1(config-ospf-vrf)#end
RP/0/0/CPU0:May  4 00:36:06.138 : config[65813]: %MGBL-SYS-5-CONFIG_I : Configured from console by admin
RP/0/0/CPU0:XR1#
RP/0/0/CPU0:XR1#
```

Now all routes that were redistributed from OSPF to VPNv4 BGP on either R2 or XR1 have the same OSPF Domain-ID, as seen below.

```
R2#show bgp vpnv4 unicast all 1.1.1.1/32 | include DOMAIN
    Extended Community: RT:100:1 OSPF DOMAIN ID:0x0005:0x000000640200


R2#show bgp vpnv4 unicast all 20.20.20.20/32 | include DOMAIN
    Extended Community: RT:100:1 OSPF DOMAIN ID:0x0005:0x000000640200


RP/0/0/CPU0:XR1#show bgp vpnv4 unicast vrf VPN_A 1.1.1.1/32 | include domain
```

Mon May 4 00:37:50.130 UTC Extended community: **OSPF domain-id:0x5:0x000000640200** OSPF route-type:0:2:0x0 OSPF router-id:10.1.2.2 RT:100:1

```
RP/0/0/CPU0:XR1#show bgp vpnv4 unicast vrf VPN_A 20.20.20.20/32 | include domain
Mon May  4 00:38:17.709 UTC
     Extended community: OSPF domain-id:0x5:0x000000640200 OSPF route-type:0:1:0x0 OSPF router-id:19.19.19.19 RT:100:1
```

CONTENTS

The final result of this is that the CE routers now see the PE routers as ABRs for these prefixes instead of ASBRs.

CONTENTS

```
R1#show ip ospf database

            OSPF Router with ID (1.1.1.1) (Process ID 100)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
1.1.1.1         1.1.1.1         1032        0x80000004 0x00B144 2
10.1.2.2        10.1.2.2        986         0x80000004 0x00DE12 1

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
10.1.2.2        10.1.2.2        986         0x80000003 0x000DF9

                Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
 10.19.20.0      10.1.2.2        399         0x80000001 0x004B30
 20.20.20.20     10.1.2.2        399         0x80000001 0x00FD5D

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, 3 subnets, 2 masks
 O IA     10.19.20.0/24 [110/2] via 10.1.2.2, 00:07:33, GigabitEthernet1.12
      20.0.0.0/32 is subnetted, 1 subnets
 O IA     20.20.20.20 [110/3] via 10.1.2.2, 00:07:33, GigabitEthernet1.12

RP/0/3/CPU0:XR2#show ospf database
Mon May  4 00:46:27.855 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     1073        0x80000004 0x00fa11 1
20.20.20.20     20.20.20.20     1001        0x80000004 0x005252 2

                Net Link States (Area 0)
```


CONTENTS

```
Link ID         ADV Router      Age         Seq#       Checksum
10.19.20.19     19.19.19.19     1073        0x80000003 0x00d03a

            Summary Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
 1.1.1.1         19.19.19.19     627         0x80000001 0x00abbe
 10.1.2.0        19.19.19.19     627         0x80000001 0x002b37


RP/0/3/CPU0:XR2#show route ipv4 ospf
Mon May  4 00:46:55.083 UTC

 O IA 1.1.1.1/32  [110/3] via 10.19.20.19, 00:10:53, GigabitEthernet0/0/0/0.1920
 O IA 10.1.2.0/24  [110/2] via 10.19.20.19, 00:10:53, GigabitEthernet0/0/0/0.1920
```

« MPLS L3 VPN with EIGRP (/workbook/view/service-provider-v4/task/mpls-l3-vpn-with-eigrp-Mjg1Nw%3D%3D) | MPLS L3 VPN with BGP (/workbook/view/service-provider-v4/task/mpls-l3-vpn-with-bgp-Mjg1OQ%3D%3D) »

CONTENTS