

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

MPLS L3 VPN and OSPF Sham-Links

« MPLS L3 VPN VPNv4 Route Reflection w/ IOS XR (/workbook/view/service-provider-v4/task/mpls-l3-vpn-vpnv4-route-reflection-w-ios-xr-Mjg2Mw%3D%3D) | MPLS L3 VPN and OSPF Domain-ID (/workbook/view/service-provider-v4/task/mpls-l3-vpn-and-ospf-domain-id-Mjg2NQ%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 \(http://labs.ine.com/workbook/view/service-provider-v4/task/ccie-spv4-topology-diagrams-initial-configs\)](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

- 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 R1 & XR2.
 - Enable OSPF Area 0 between XR1 & XR2.
 - Advertise the Loopback0 networks of R1 & XR2 into OSPF Area 0.
 - Modify the link between R1 and XR2 to have an OSPF cost of 100.
- 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.
- Configure an OSPF Sham Link between R2 and XR1 as follows:
 - Create a new Loopback interface on R2 with the address 10.2.2.2/32.
 - Create a new Loopback interface on XR1 with the address 10.19.19.19/32.
 - Assign these new Loopbacks to VRF VPN_A.
 - Advertise them into VPNv4 BGP, but **not** into the VRF aware OSPF process.
 - Configure an OSPF Sham Link in Area 0 between R2 and XR1 using these new Loopback interfaces.
- Once complete, the following reachability should be achieved:

<https://t.me/learningnets>

- 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.
- If either R1 or XR2 lose connectivity to the MPLS cloud, traffic should be automatically be rerouted over the backdoor link.

Configuration [Click to collapse](#)

```

R1:
interface GigabitEthernet1.120
 ip ospf cost 100
 !
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
 !
interface Loopback100
 vrf forwarding VPN_A
 ip address 10.2.2.2 255.255.255.255
 !
ip prefix-list SHAM_LINK_ENDPOINTS seq 5 permit 10.2.2.2/32
ip prefix-list SHAM_LINK_ENDPOINTS seq 10 permit 10.19.19.19/32
 !
route-map BGP_TO_OSPF deny 10
 match ip address prefix-list SHAM_LINK_ENDPOINTS
 !
route-map BGP_TO_OSPF permit 20
 !
router ospf 100 vrf VPN_A
 area 0 sham-link 10.2.2.2 10.19.19.19
 redistribute bgp 100 subnets route-map BGP_TO_OSPF
 network 10.1.2.2 0.0.0.0 area 0
 !
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
 no synchronization
 redistribute ospf 100
 network 10.2.2.2 mask 255.255.255.255
 exit-address-family

```

```

!
XR1:
vrf VPN_A
  address-family ipv4 unicast
    import route-target
      100:1
  !
  export route-target
    100:1
  !
  !
  !
interface GigabitEthernet0/0/0.1920
  vrf VPN_A
  no ipv4 address
  ipv4 address 10.19.20.19 255.255.255.0
!
interface Loopback100
  vrf VPN_A
  ipv4 address 10.19.19.19 255.255.255.255
!
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
    network 10.19.19.19/32
  !
  !
  !
  prefix-set SHAM_LINK_ENDPOINTS
    10.2.2.2,
    10.19.19.19
  end-set
  !
  route-policy BGP_TO_OSPF
    if destination in SHAM_LINK_ENDPOINTS then
      drop
    else
      pass
    endif
  end-policy
  !
router ospf 100
  vrf VPN_A

```

```

redistribute bgp 100 route-policy BGP_TO_OSPF

area 0

interface GigabitEthernet0/0/0/0.1920

sham-link 10.19.19.19 10.2.2.2

!

!

XR2:

router ospf 100

area 0

interface Loopback0

!

interface GigabitEthernet0/0/0/0.1920

!

interface GigabitEthernet0/0/0/0.120

cost 100

!

!

!

```

Verification

OSPF Sham Links, similar to Virtual Links, are multihop unicast adjacencies between OSPF neighbors that are used for the purpose of traffic engineering. Due to the inherent nature of OSPF Path Selection per the RFC specification, OSPF always prefers Intra Area routes over Inter Area routes over External Routes over NSSA External Routes, regardless of Administrative Distance or Metric.

OSPF Sham Links are needed in designs where backdoor connections exist between customer sites, such as legacy T1, Frame Relay, or any other point-to-point connection, in addition to the MPLS L3VPN connection, and the MPLS connections need to be preferred over the backdoor links. Due to the behavior of the OSPF “Superbackbone” of MPLS, routes on one customer site that are Intra-Area (O routes) will appear in other customer sites as Inter-Area (O IA routes) as long as the OSPF Domain-ID field matches in the BGP VPNv4 route. If the Domain-ID field does not match, the routes will appear as either External Type 1 (E1 routes) or External Type 2 (E2 routes). This means that if there is a backdoor link between the customer sites that allows routes to be exchanged as Intra-Area (O routes), these will always be preferred over the O IA, E1, or E2 routes coming from the MPLS L3VPN.

To fix this, a Sham Link, similar to a Virtual Link, extends the OSPF flooding domain of an area over a logical multi-hop adjacency. This essentially makes the PE routers no longer appear as ABRs or ASBRs in the OSPF database, but regular routers that are all in the same area. The result is that Intra-Area (O routes) can be learned from the PE routers, which can then be preferred over backdoor links simply based on changing the OSPF cost of interfaces in the topology per your desired traffic engineering goals.

Unlike virtual-links, sham-links can be assigned to any area, and a reachable IPv4 source and destination need to be manually configured in order for the adjacency to be established.

In this particular example R2 and XR1 form a Sham Link adjacency over the MPLS network. To do this, they first need a new /32 Loopback interface that is a member of the VRF, and is advertised into VPNv4 BGP. In this specific case these are the 10.2.2.2/32 and 10.19.19.19/32 prefixes.

```
R2#show ip route vrf VPN_A 10.2.2.2
```

```
Routing Table: VPN_A
```

```
Routing entry for 10.2.2.2/32
```

```
Known via "connected", distance 0, metric 0 (connected, via interface)
```

```
Redistributing via ospf 100
```

```
Advertised by bgp 100
```

```
Routing Descriptor Blocks:
```

```
* directly connected, via Loopback100
```

```
Route metric is 0, traffic share count is 1
```

```
R2#show ip route vrf VPN_A 10.19.19.19
```

```
Routing Table: VPN_A
```

```
Routing entry for 10.19.19.19/32
```

```
Known via "bgp 100", distance 200, metric 0, type internal
```

```
Redistributing via ospf 100
```

```
Last update from 19.19.19.19 00:22:44 ago
```

```
Routing Descriptor Blocks:
```

```
* 19.19.19.19 (default), from 19.19.19.19, 00:22:44 ago
```

```
Route metric is 0, traffic share count is 1
```

```
AS Hops 0
```

```
MPLS label: 16000
```

```
MPLS Flags: MPLS Required
```

```
RP/0/0/CPU0:XR1#show route vrf VPN_A 10.2.2.2
```

```
Fri May 8 01:10:30.351 UTC
```

```
Routing entry for 10.2.2.2/32
```

```
Known via "bgp 100", distance 200, metric 0, type internal
```

```
Installed May 8 00:47:24.236 for 00:23:06
```

```
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 0
```

```
No advertising protos.
```

```
RP/0/0/CPU0:XR1#show route vrf VPN_A 10.19.19.19
```

```
Fri May 8 01:10:46.730 UTC
```

```
Routing entry for 10.19.19.19/32
```

```
Known via "local", distance 0, metric 0 (connected)
```

```
Installed May 8 00:46:25.840 for 00:24:20
```

```
Routing Descriptor Blocks
```

```
directly connected, via Loopback100
```

```
Route metric is 0
```

```
No advertising protos.
```

Only once connectivity is established between these new interfaces can the Sham Link be formed.

```
R2#ping vrf VPN_A 10.19.19.19 source 10.2.2.2
```

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

```
Sending 5, 100-byte ICMP Echos to 10.19.19.19, timeout is 2 seconds:
```

```
Packet sent with a source address of 10.2.2.2
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/6/10 ms
```

```
RP/0/0/CPU0:XR1#ping vrf VPN_A 10.2.2.2 source 10.19.19.19
```

```
Fri May 8 01:11:24.677 UTC
```

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

```
Sending 5, 100-byte ICMP Echos to 10.2.2.2, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 9/21/29 ms
```

Similar to a Virtual Link, the Sham Link forms a unicast multi-hop adjacency between the neighbors.

R2#show ip ospf 100 neighbor

Neighbor ID	Pri	State	Dead Time	Address	Interface
1.1.1.1	1	FULL/DR	00:00:30	10.1.2.1	GigabitEthernet1.12
19.19.19.19	0	FULL/ -	-	10.19.19.19	OSPF_SL0

R2#show ip ospf 100 neighbor 19.19.19.19

Neighbor 19.19.19.19, interface address 10.19.19.19

In the area 0 via interface OSPF_SL0

Neighbor priority is 0, State is FULL, 6 state changes

DR is 0.0.0.0 BDR is 0.0.0.0

Options is 0x0 in Hello

Options is 0x72 in DBD (E-bit, L-bit, DC-bit, O-bit)

LLS Options is 0x1 (LR)

Neighbor is up for 00:07:42

Index 2/2, retransmission queue length 0, number of retransmission 1

First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)

Last retransmission scan length is 1, maximum is 1

Last retransmission scan time is 0 msec, maximum is 0 msec

R2#show ip ospf 100 sham-links

Sham Link OSPF_SL0 to address 10.19.19.19 is up

Area 0 source address 10.2.2.2

Run as demand circuit

DoNotAge LSA allowed. Cost of using 1 State POINT_TO_POINT,

Timer intervals configured, Hello 10, Dead 40, Wait 40,

Hello due in 00:00:04

Adjacency State FULL (Hello suppressed)

Index 2/2, retransmission queue length 0, number of retransmission 1

First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)

Last retransmission scan length is 1, maximum is 1

Last retransmission scan time is 0 msec, maximum is 0 msec

RP/0/0/CPU0:XR1#show ospf vrf VPN_A neighbor

Fri May 8 01:12:33.992 UTC

* Indicates MADJ interface

Neighbors for OSPF 100, VRF VPN_A

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.2.2.2	1	FULL/ -	-	10.2.2.2	OSPF_SL0
Neighbor is up for 00:08:02					
20.20.20.20	1	FULL/DR	00:00:31	10.19.20.20	GigabitEthernet0/0/0/0.1920
Neighbor is up for 00:06:36					

Total neighbor count: 2

RP/0/0/CPU0:XR1#show ospf vrf VPN_A neighbor 10.2.2.2

Fri May 8 01:13:01.381 UTC

* Indicates MADJ interface

```
Neighbors for OSPF 100, VRF VPN_A
```

```
Neighbor 10.2.2.2, interface address 10.2.2.2
```

```
  In the area 0 via interface OSPF_SL0
```

```
  Neighbor priority is 1, State is FULL, 6 state changes
```

```
  DR is 0.0.0.0 BDR is 0.0.0.0
```

```
  Options is 0x72
```

```
  LLS Options is 0x1 (LR)
```

```
  Neighbor is up for 00:08:29
```

```
  Number of DBD retrans during last exchange 0
```

```
  Index 1/1, retransmission queue length 0, number of retransmission 1
```

```
  First 0(0)/0(0) Next 0(0)/0(0)
```

```
  Last retransmission scan length is 1, maximum is 1
```

```
  Last retransmission scan time is 0 msec, maximum is 0 msec
```

```
  LS Ack list: NSR-sync pending 0, high water mark 0
```

```
Total neighbor count: 1
```

From the point of view of the CE routers, the PE routers R2 and XR1 are now all part of the same OSPF area, and hence the same flooding domain. Note that there are no Network Summary LSAs (LSA Type 3) or External LSAs (LSA Type 5) because there are no ABRs or ASBRs. The entire topology is treated as one flat area 0 now.

```
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	600	0x80000003	0x00C87B	3
10.2.2.2	10.2.2.2	681	0x80000003	0x004041	2
19.19.19.19	19.19.19.19	2 (DNA)	0x80000005	0x007177	2
20.20.20.20	20.20.20.20	565	0x80000003	0x00B22B	3

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

Link ID	ADV Router	Age	Seq#	Checksum
10.1.2.1	1.1.1.1	78	0x80000002	0x00AF63
10.1.20.1	1.1.1.1	600	0x80000001	0x001AA7
10.19.20.20	20.20.20.20	565	0x80000001	0x00AC5B

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

```
Fri May 8 01:16:50.615 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
1.1.1.1	1.1.1.1	659	0x80000003	0x00c87b	3
10.2.2.2	10.2.2.2	739	0x80000003	0x004041	2
19.19.19.19	19.19.19.19	623	0x80000005	0x007177	2
20.20.20.20	20.20.20.20	622	0x80000003	0x00b22b	3

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

Link ID	ADV Router	Age	Seq#	Checksum
10.1.2.1	1.1.1.1	137	0x80000002	0x00af63
10.1.20.1	1.1.1.1	659	0x80000001	0x001aa7
10.19.20.20	20.20.20.20	622	0x80000001	0x00ac5b

If we look at this from a route calculation and path selection point of view, R1 now performs a full SPF run over both the backdoor link and over the MPLS L3VPN network in order to reach XR2's Loopback0, and vice versa. R1 simply chooses the lowest cost path to reach XR2, and vice versa, as there is no difference in route types anymore.

Based on the Router LSAs (LSA Type 1) R1 knows that it can reach XR2 with Router-ID 20.20.20.20 two ways, directly over the backdoor link or via the MPLS L3VPN. The SPF cost via the MPLS L3VPN is lower, hence that is the shortest path.

```
R1#show ip ospf database router 1.1.1.1
```

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

```
Router Link States (Area 0)
```

```
LS age: 801
```

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

```
LS Type: Router Links
```

```
Link State ID: 1.1.1.1
```

```
Advertising Router: 1.1.1.1
```

```
LS Seq Number: 80000003
```

```
Checksum: 0xC87B
```

```
Length: 60
```

```
Number of Links: 3
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 1.1.1.1
```

```
(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: 10.1.20.1
```

```
(Link Data) Router Interface address: 10.1.20.1
```

```
Number of MTID metrics: 0
```

```
TOS 0 Metrics: 100
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 10.1.2.1
```

```
(Link Data) Router Interface address: 10.1.2.1
```

```
Number of MTID metrics: 0
```

```
TOS 0 Metrics: 1
```

```
R1#show ip ospf database router 10.2.2.2
```

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

```
Router Link States (Area 0)
```

```
LS age: 935
```

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

```
LS Type: Router Links
```

```
Link State ID: 10.2.2.2
```

```
Advertising Router: 10.2.2.2
```

```
LS Seq Number: 80000003
```

```
Checksum: 0x4041
```

```
Length: 48
```

```
Area Border Router
```

```
AS Boundary Router
```

```
Number of Links: 2
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.1.2.1
(Link Data) Router Interface address: 10.1.2.2
Number of MTID metrics: 0
  TOS 0 Metrics: 1
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 19.19.19.19
(Link Data) Router Interface address: 0.0.0.20
Number of MTID metrics: 0
  TOS 0 Metrics: 1
```

```
R1#show ip ospf database router 19.19.19.19
```

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

```
      Router Link States (Area 0)
```

```
LS age: 2 (DoNotAge)
```

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

```
LS Type: Router Links
```

```
Link State ID: 19.19.19.19
```

```
Advertising Router: 19.19.19.19
```

```
LS Seq Number: 80000005
```

```
Checksum: 0x7177
```

```
Length: 48
```

```
Area Border Router
```

```
AS Boundary Router
```

```
Number of Links: 2
```

```
Link connected to: another Router (point-to-point)
```

```
(Link ID) Neighboring Router ID: 10.2.2.2
```

```
(Link Data) Router Interface address: 0.0.0.2
```

```
Number of MTID metrics: 0
```

```
  TOS 0 Metrics: 1
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 10.19.20.20
```

```
(Link Data) Router Interface address: 10.19.20.19
```

```
Number of MTID metrics: 0
```

```
  TOS 0 Metrics: 1
```

```
R1#show ip ospf database router 20.20.20.20
```

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

```
      Router Link States (Area 0)
```

```
LS age: 989
```

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

```
LS Type: Router Links
```

```
Link State ID: 20.20.20.20
```

Advertising Router: 20.20.20.20

LS Seq Number: 80000003

Checksum: 0xB22B

Length: 60

Number of Links: 3

Link connected to: a Stub Network

(Link ID) Network/subnet number: 20.20.20.20

(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: 10.1.20.1

(Link Data) Router Interface address: 10.1.20.20

Number of MTID metrics: 0

TOS 0 Metrics: 100

Link connected to: a Transit Network

(Link ID) Designated Router address: 10.19.20.20

(Link Data) Router Interface address: 10.19.20.20

Number of MTID metrics: 0

TOS 0 Metrics: 1

Notice that the BGP next-hop of routes learned over the Sham-Link is the Sham-Link endpoint, when looking at an IOS-XR device's local OSPF RIB.

```
RP/0/0/CPU0:XR1#show ospf 100 vrf VPN_A routes
```

```
Fri May 8 01:58:19.694 UTC
```

```
Topology Table for ospf 100, VRF VPN_A with ID 19.19.19.19
```

```
Codes: 0 - Intra area, 0 IA - Inter area
```

```
0 E1 - External type 1, 0 E2 - External type 2
```

```
0 N1 - NSSA external type 1, 0 N2 - NSSA external type 2
```

```
0 1.1.1.1/32, metric 3 area 0.0.0.0
```

```
10.2.2.2, from 1.1.1.1,
```

```
0 10.1.2.0/24, metric 2 area 0.0.0.0
```

```
10.2.2.2, from 1.1.1.1,
```

```
0 10.1.20.0/24, metric 101 area 0.0.0.0
```

```
10.19.20.20, from 1.1.1.1, via GigabitEthernet0/0/0.1920
```

```
0 10.19.20.0/24, metric 1 area 0.0.0.0
```

```
10.19.20.19, directly connected, via GigabitEthernet0/0/0.1920
```

```
0 20.20.20.0/32, metric 2 area 0.0.0.0
```

```
10.19.20.20, from 20.20.20.20, via GigabitEthernet0/0/0.1920
```

On IOS, the local OSPF RIB displays the Sham-Link interface as the outgoing interface for routes learned over the Sham-Link.

```
R2#show ip ospf 100 rib
```

```
OSPF Router with ID (10.2.2.2) (Process ID 100)
```

```
Base Topology (MTID 0)
```

```
OSPF local RIB
```

```
Codes: * - Best, > - Installed in global RIB
```

```
*> 1.1.1.1/32, Intra, cost 2, area 0
    via 10.1.2.1, GigabitEthernet1.12
* 10.1.2.0/24, Intra, cost 1, area 0, Connected
    via 10.1.2.2, GigabitEthernet1.12
*> 10.1.20.0/24, Intra, cost 101, area 0
    via 10.1.2.1, GigabitEthernet1.12
*> 10.19.20.0/24, Intra, cost 2, area 0
    via 19.19.19.19, OSPF_SL0
*> 20.20.20.0/32, Intra, cost 3, area 0
    via 19.19.19.19, OSPF_SL0
```

Although these routes are installed in the RIB via OSPF, they do not get redistributed back into BGP due to them resolving the next-hop/outgoing interface to the Sham-Link.

Proper labeling still takes place, as these routes end up recursing to their corresponding BGP next-hop once they are entered into the FIB. Notice that in IOS's case, the 1.1.1.1/32 prefix is considered a RIB-Failure due to this route being installed via OSPF.

```
R2#show bgp vpnv4 unicast all
BGP table version is 23, 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 10.1.20.0/24	19.19.19.19	101	100	0	?
*>	10.1.2.1	101		32768	?
*> 10.2.2.2/32	0.0.0.0	0		32768	i
*>i 10.19.19.19/32	19.19.19.19	0	100	0	i
r>i 10.19.20.0/24	19.19.19.19	0	100	0	?
r>i 20.20.20.20/32	19.19.19.19	2	100	0	?

```
R2#show ip cef vrf VPN_A 20.20.20.20/32 detail
20.20.20.20/32, epoch 0, flags [rib defined all labels]
recursive via 19.19.19.19 label 16007
nextthop 20.2.3.3 GigabitEthernet1.23 label 37
nextthop 20.2.4.4 GigabitEthernet1.24 label 28
```

```
RP/0/0/CPU0:XR1#show bgp vpnv4 unicast
advertised attribute-key
RP/0/0/CPU0:XR1#show bgp vpnv4 unicast
Fri May 8 02:08:40.822 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 36
BGP scan interval 60 secs
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	?
*> 10.1.20.0/24	10.19.20.20	101		32768	?
* i	2.2.2.2	101	100	0	?
*>i10.2.2.2/32	2.2.2.2	0	100	0	i
*> 10.19.19.19/32	0.0.0.0	0		32768	i
*> 10.19.20.0/24	0.0.0.0	0		32768	?
*> 20.20.20.20/32	10.19.20.20	2		32768	?

Processed 7 prefixes, 8 paths

```
RP/0/0/CPU0:XR1#show cef vrf VPN_A 1.1.1.1/32 detail
```

```
Fri May 8 02:08:54.331 UTC
```

```
1.1.1.1/32, version 5, internal 0x14004001 0x0 (ptr 0xa0edc1f4) [1], 0x0 (0x0), 0x208 (0xa13f6280)
```

```
Updated May 8 00:47:24.256
```

```
Prefix Len 32, traffic index 0, precedence n/a, priority 3
```

```
gateway array (0xa0d30040) reference count 3, flags 0x4038, source rib (6), 0 backups
```

```
[1 type 1 flags 0x48089 (0xa1410320) ext 0x0 (0x0)]
```

```
LW-LDI[type=0, refc=0, ptr=0x0, sh-ldi=0x0]
```

```
via 2.2.2.2, 3 dependencies, recursive [flags 0x6000]
```

```
path-idx 0 NHID 0x0 [0xa145f5f4 0x0]
```

```
next hop VRF - 'default', table - 0xe0000000
```

```
next hop 2.2.2.2 via 16005/0/21
```

```
next hop 20.5.19.5/32 Gi0/0/0/0.519 labels imposed {28 21}
```

```
next hop 20.6.19.6/32 Gi0/0/0/0.619 labels imposed {24 21}
```

```
Load distribution: 0 (refcount 1)
```

Hash	OK	Interface	Address
0	Y	Unknown	16005/0

The final result is that all inter-site traffic goes over the MPLS L3VPN as opposed to the backdoor link

```
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, 5 subnets, 2 masks
O    10.19.20.0/24 [110/3] via 10.1.2.2, 00:20:01, GigabitEthernet1.12
    20.0.0.0/32 is subnetted, 1 subnets
O    20.20.20.20 [110/4] via 10.1.2.2, 00:20:01, GigabitEthernet1.12
```

```
R1#traceroute 20.20.20.20
```

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 28/16007 Exp 0] 14 msec 11 msec 12 msec
 3 20.4.5.5 [MPLS: Labels 36/16007 Exp 0] 12 msec 8 msec 8 msec
 4 20.5.19.19 [MPLS: Label 16007 Exp 0] 20 msec 20 msec 20 msec
 5 10.19.20.20 20 msec * 11 msec
 5 10.19.20.20 136 msec * 0 msec
```

```
RP/0/0/CPU0:XR2#show route ospf
```

Fri May 8 01:28:51.456 UTC

```
O    1.1.1.1/32 [110/4] via 10.19.20.19, 00:22:22, GigabitEthernet0/0/0.1920
O    10.1.2.0/24 [110/3] via 10.19.20.19, 00:22:22, GigabitEthernet0/0/0.1920
```

```
RP/0/0/CPU0:XR2#traceroute 1.1.1.1
```

Fri May 8 01:29:09.305 UTC

Type escape sequence to abort.

Tracing the route to 1.1.1.1

```
 1 10.19.20.19 9 msec 0 msec 0 msec
 2 20.5.19.5 [MPLS: Labels 28/21 Exp 0] 9 msec 0 msec 0 msec
 3 20.4.5.4 [MPLS: Labels 27/21 Exp 0] 0 msec 0 msec 0 msec
 4 10.1.2.2 [MPLS: Label 21 Exp 0] 0 msec 0 msec 0 msec
 5 10.1.2.1 0 msec * 9 msec
```

If a connection to the MPLS network is lost, traffic will reroute over the backdoor link. Note that this path has a higher OSPF cost than the previous one.

```
RP/0/0/CPU0:XR2#configure
Fri May 8 01:31:38.134 UTC
RP/0/0/CPU0:XR2(config)#interface GigabitEthernet0/0/0.1920
RP/0/0/CPU0:XR2(config-subif)#shutdown
RP/0/0/CPU0:XR2(config-subif)#commit
Fri May 8 01:31:48.724 UTC
RP/0/0/CPU0:May 8 01:31:48.754 : ospf[1015]: %ROUTING-OSPF-5-ADJCHG : Process 100, Nbr 19.19.19.19 on GigabitEthernet0/0/0.1920 in area
0 from FULL to DOWN, Neighbor Down: interface down or detached, vrf default vrfid 0x60000000
RP/0/0/CPU0:May 8 01:31:48.974 : config[65710]: %MGBL-CONFIG-6-DB_COMMIT : Configuration committed by user 'admin'. Use 'show configuration
commit changes 1000000098' to view the changes.
RP/0/0/CPU0:XR2(config-subif)#end
RP/0/0/CPU0:May 8 01:31:55.573 : config[65710]: %MGBL-SYS-5-CONFIG_I : Configured from console by admin

RP/0/0/CPU0:XR2#show route ospf
Fri May 8 01:31:59.703 UTC

0 1.1.1.1/32 [110/101] via 10.1.20.1, 00:00:10, GigabitEthernet0/0/0.120
0 10.1.2.0/24 [110/101] via 10.1.20.1, 00:00:10, GigabitEthernet0/0/0.120
```

```
RP/0/0/CPU0:XR2#traceroute 1.1.1.1
Fri May 8 01:32:05.223 UTC

Type escape sequence to abort.
Tracing the route to 1.1.1.1

 1 10.1.20.1 0 msec * 0 msec
RP/0/0/CPU0:XR2#

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, 5 subnets, 2 masks
O      10.19.20.0/24 [110/3] via 10.1.2.2, 00:01:28, GigabitEthernet1.12
      20.0.0.0/32 is subnetted, 1 subnets
O      20.20.20.0 [110/101] via 10.1.20.20, 00:02:07, GigabitEthernet1.120

R1#traceroute 20.20.20.20

Type escape sequence to abort.
Tracing the route to 20.20.20.20
```

« MPLS L3 VPN VPNv4 Route Reflection w/ IOS XR (/workbook/view/service-provider-v4/task/mpls-l3-vpn-vpnv4-route-reflection-w-ios-xr-Mjg2Mw%3D%3D) | MPLS L3 VPN and OSPF Domain-ID (/workbook/view/service-provider-v4/task/mpls-l3-vpn-and-ospf-domain-id-Mjg2NQ%3D%3D) »