

DMVPN Support for IWAN

DMVPN supports Cisco Intelligent WAN architecture to provide transport independence through overlay routing. The DMVPN Multiple Tunnel Termination feature enables support for secondary paths for the supported routing protocols in the Routing Information Base (RIB).

- Prerequisites for DMVPN Support for IWAN, on page 1
- Restrictions for DMVPN Support for IWAN, on page 1
- Information About DMVPN Support for IWAN, on page 1
- How to Configure DMVPN Support for IWAN, on page 4
- Configuration Examples for DMVPN Support for IWAN, on page 5
- Additional References for DMVPN Support for IWAN, on page 10
- Feature Information for DMVPN Support for IWAN, on page 11

Prerequisites for DMVPN Support for IWAN

For DMVPN Multiple Tunnel Termination feature to work, the following prerequisites must be considered

- DMVPN Multiple Tunnel Termination feature requires support from crypto maps and DMVPN.
- Only BGP and EIGRP routing protocols are supported on this feature. One of the two routing protocols, BGP and EIGRP, must be enabled for this feature to work.

Restrictions for DMVPN Support for IWAN

For DMVPN Multiple Tunnel Termination feature the overlay routing should be active-passive in nature.

Information About DMVPN Support for IWAN

Transport Independence

DMVPN supports Cisco IWAN by providing transport independence through overlay routing. Overlay routing simplifies the WAN transport (dial-up, leased circuits, MPLS, and IPsec VPNs), by deploying and supporting

consistent routing protocol across any transport, controlling traffic and load sharing. Overlay routing provides transport independence so that the user can select any WAN technology.

Transport independence eases change in transport options and service providers. Changing transports does not impact the overlay routing design. This technology supports use of multiple WAN transports, as the transport type is associated to the underlay network and is not relevant to the overlay network which is consistent to the DMVPN tunnel.

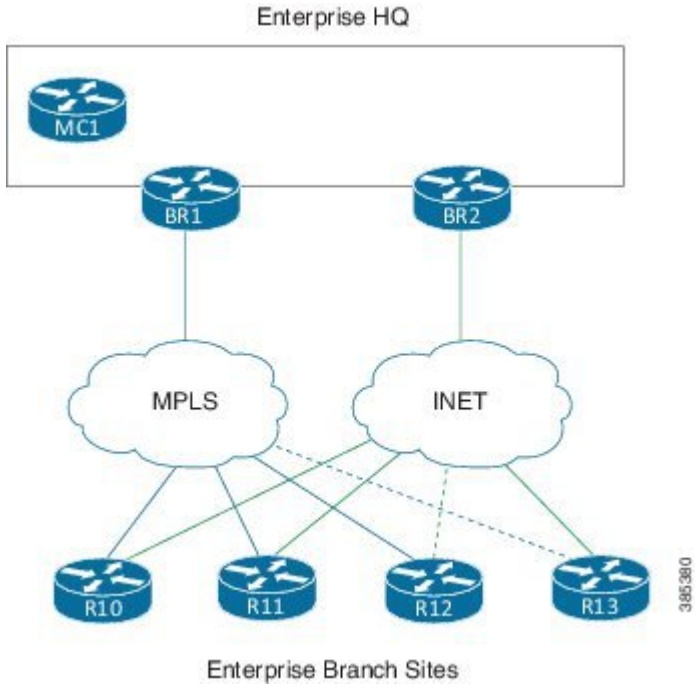
Transport independence provides single routing domain, consistent troubleshooting and topology for WAN transports. As long as the transport network delivers the DMVPN packets between the hub and the spoke, the transport device topology is not relevant to the traffic flowing across the DMVPN tunnel.

DMVPN for IWAN

DMVPN uses multipoint generic routing encapsulation (mGRE) tunnels to interconnect the hubs and all of the spokes. For IWAN deployments, DMVPN provides integration with PfR and simplifies route control across any transport. DMVPN supports full mesh connectivity over any carrier transport with a simple hub-and-spoke configuration. DMVPN also supports spoke that have dynamically assigned IP addresses.

The following figure shows IWAN deployments with multiple WAN transports. This design enables convergence across WAN transports when all channels in a given transport fail or reach their maximum bandwidth limits.

Figure 1: DMVPN for IWAN



Secondary Paths

For a single tunnel case, the routing method installs multiple paths in the RIB, one or more leaving each tunnel. Based on the configuration, this includes some or all of the available free paths. The paths can be classified into following classes:

- Regular next-hops/paths are the most common kind of paths. They are also referred to as primary paths; other alternate next-hops are sometimes referred to as secondary paths.
- Repair next-hops/paths forward traffic during a routing transition and are not used as long as one or more regular next hops are active.
- Secondary next-hops/paths are special loop free paths that is used as an alternate to regular and repair paths.

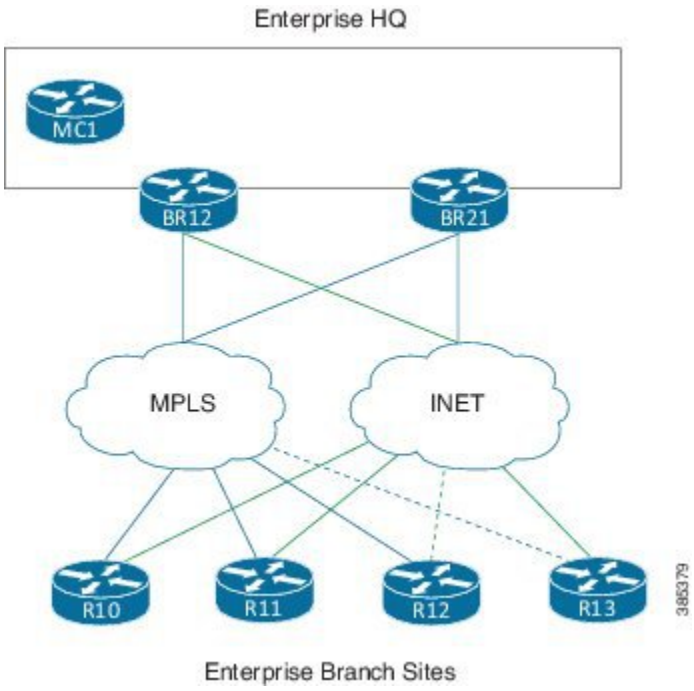
When at least one of the primary paths are in use, the secondary paths are not used for regular forwarding. The secondry paths should be distinguishable from other regular and alternate paths. The secondary paths can still be overridden using next hop overrides. The routing protocol computes "n" secondary paths with the following requirement from RIB:

- Allow the routing protocol to install the "n1" primary paths as a regular path
- Allow the routing protocol to install the "n2" secondary paths as alternate paths.

DMVPN Multiple Tunnel Termination

Network access resiliency at a single hub in Cisco IWAN without having to add any network devices, involves terminating multiple WAN links on the same device. The DMVPN Multiple Tunnel Termination feature provides support for multiple tunnel terminations (interfaces) in the same VRF on the same hub device.

Figure 2: DMVPN Multiple Tunnel Termination



The DMVPN Multiple Tunnel Termination feature also provides transport resilience to DMVPN. Using one tunnel per-transport provides better visibility to Performance Routing (PfR), about the conditions in the underlying transport and still being transport independent. IWAN as a whole is transport independent along with the services running on the overlay.

DMVPN Multiple Tunnel Termination feature brings in support for secondary paths for the supported routing protocols in the RIB. The routing protocols are configured in such a way that there is only one primary/regular path and one or more secondary paths for a network. When PfR is used in conjunction with this feature, PfR

is used as the primary as well the secondary path so that all paths can be used in an active-active manner. Use the **maximum-secondary-paths** [**eigrp** | **ibgp**] *path* command to configure this feature, where the path indicates the number of secondary paths a routing protocol is allowed to install. The range for path is from zero to 32.

How to Configure DMVPN Support for IWAN

Configuring DMVPN Support for IWAN

Perform this task to configure IPsec profile on the device.

```
crypto ikev2 keyring keyring1
peer peer1
  address 0.0.0.0 0.0.0.0
  pre-shared-key key1

crypto ikev2 proposal proposal1
 encryption aes-cbc-128
 prf sha256 sha512
 group 14

crypto ikev2 policy proposal1
 match fvrf vrf1
 proposal proposal1

crypto ikev2 profile profile1
 description This is an IKEv2 profile
 match fvrf vrf1
 match identity remote address 10.0.0.1
 identity local address 10.0.0.0
 authentication remote pre-share
 authentication local pre-share
 keyring local key1

crypto ipsec transform-set transform1 esp-gcm 256
 mode transport

crypto ipsec profile profile2
 set transform-set esp-gcm 256
 set ikev2-profile profile1

crypto ipsec security-association replay window-size 15
```

Perform this task to configure the tunnel.

```
interface Tunnel 10
 tunnel source GigabitEthernet0/0/0
 tunnel mode gre multipoint
 tunnel key 10000
 tunnel vrf vrf1
 tunnel protection ipsec profile profile2
```

Perform the following task to configure BGP routing process.

```
router bgp 45000
 bgp router-id 172.17.1.99
 bgp log-neighbor-changes
 timers bgp 70 120
 neighbor 192.168.1.2 remote-as 40000
 neighbor 192.168.3.2 remote-as 50000
 neighbor 192.168.3.2 description finance
 !
 address-family ipv4
  neighbor 192.168.1.2 activate
  neighbor 192.168.3.2 activate
  no auto-summary
  no synchronization
  network 172.17.1.0 mask 255.255.255.0
  exit-address-family
 !
```

Configuring DMVPN Multiple Tunnel Termination

Perform the following task to configure DMVPN Multiple Tunnel Termination

```
router bgp 1
bgp log-neighbor-changes
bgp listen range 192.168.0.0/16
peer-group SPOKES2
bgp listen range 190.168.0.0/16
peer-group SPOKES network 192.168.0.0
aggregate-address 192.168.0.0 255.255.0.0 summary-only
timers bgp 10 30
neighbor SPOKES2 peer-group
neighbor SPOKES2 remote-as 1
neighbor SPOKES2 next-hop-self
maximum-secondary-paths eigrp 1
```

Configuration Examples for DMVPN Support for IWAN

Example: DMVPN Support for IWAN

The following is an example for configuring DMVPN on hub.

```
router eigrp DMVPN
!
address-family ipv4 unicast autonomous-system 100
!
af-interface Tunnel0   \
summary-address 192.168.0.0 255.255.0.0
no split-horizon
exit-af-interface
!
af-interface Tunnel1
summary-address 192.168.0.0 255.255.0.0
no split-horizon
exit-af-interface
```

```
!
topology base
maximum-secondary-paths 4
fast-reroute per-prefix all
fast-reroute tie-break interface-disjoint 1
exit-af-topology
network 10.0.0.0
network 20.0.0.0
network 192.168.149.0
exit-address-family
!
```

The following is an example for configuring DMVPN on spoke 1.

```
router bgp 1
bgp log-neighbor-changes
bgp listen range 20.0.0.0/8
peer-group SPOKES2
bgp listen range 10.0.0.0/8
peer-group SPOKES
network 192.168.149.0
aggregate-address 192.168.0.0 255.255.0.0 summary-only
timers bgp 10 30
neighbor SPOKES peer-group
neighbor SPOKES remote-as 1
neighbor SPOKES next-hop-self
neighbor SPOKES2 peer-group
neighbor SPOKES2 remote-as 1
neighbor SPOKES2 next-hop-self
maximum-secondary-paths eigrp 1
!
```

The following is an example for configuring DMVPN on spoke 2.

```
router bgp 1
bgp log-neighbor-changes
bgp listen range 20.0.0.0/8
peer-group SPOKES2
bgp listen range 10.0.0.0/8
peer-group SPOKES
bgp additional-paths install
network 192.168.149.0
aggregate-address 192.168.0.0 255.255.0.0 summary-only
timers bgp 10 30
neighbor SPOKES peer-group
neighbor SPOKES remote-as 1
neighbor SPOKES next-hop-self
neighbor SPOKES2 peer-group
neighbor SPOKES2 remote-as 1
neighbor SPOKES2 next-hop-self
maximum-secondary-paths eigrp 1
!
```

The following is the sample output for the **show ip bgp** command.

```
Device# show ip bgp

BGP table version is 10, local router ID is 192.168.149.1
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,
              t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
     Network          Next Hop            Metric LocPrf Weight Path
```

```
*>   192.168.0.0/16   0.0.0.0                              32768 I
s i t192.168.40.0     20.0.0.41               0    100      0 I
s>i                   10.0.0.41               0    100      0 I
s>i  192.168.50.0     10.0.0.51               0    100      0 I
s i t                 20.0.0.51               0    100      0 I
s>   192.168.149.0    0.0.0.0                 0         32768 I

Device#
```

The following is the sample output for the **show ip bgp** command in two different interfaces.

```
Device# show ip bgp 192.168.40.0

BGP routing table entry for 192.168.40.0/24, version 6
Paths: (2 available, best #2, table default, Advertisements suppressed by an aggregate.)
  Not advertised to any peer
  Refresh Epoch 1
  Local    20.0.0.41 from *20.0.0.41 (192.168.40.2)
      Origin IGP, metric 0, localpref 100, valid, internal, secondary path
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  Local    10.0.0.41 from *10.0.0.41 (192.168.40.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx pathid: 0, tx pathid: 0x0

Device# show ip bgp 192.168.50.0
BGP routing table entry for 192.168.50.0/24, version 10
Paths: (2 available, best #1, table default, Advertisements suppressed by an aggregate.)
  Not advertised to any peer
  Refresh Epoch 1
  Local    10.0.0.51 from *10.0.0.51 (192.168.50.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 1
  Local    20.0.0.51 from *20.0.0.51 (192.168.50.2)
      Origin IGP, metric 0, localpref 100, valid, internal, secondary path
      rx pathid: 0, tx pathid:
```

The following is the sample output for the **show ip route** command.

```
Device# show ip route

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, p - overrides from PfR
Gateway of last resort is not set
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.0.0.0/8 is directly connected, Tunnel0
L        10.0.0.149/32 is directly connected, Tunnel0
     20.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        20.0.0.0/8 is directly connected, Tunnel1
L        20.0.0.149/32 is directly connected, Tunnel1
B     192.168.0.0/16 [200/0], 00:02:26, Null0
B     192.168.40.0/24 [200/0] via 10.0.0.41, 00:02:26
B     192.168.50.0/24 [200/0] via 10.0.0.51, 00:01:55
      192.168.149.0/24 is variably subnetted, 2 subnets, 2 masks
```

```
C        192.168.149.0/24 is directly connected, Ethernet1/0
L        192.168.149.1/32 is directly connected, Ethernet1/0
```

The following is the sample output for the **show ip route** command for the secondary path.

```
Device# show ip route

sec 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, p - overrides from PfR
Gateway of last resort is not set
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.0.0.0/8 is directly connected, Tunnel0
L        10.0.0.149/32 is directly connected, Tunnel0
      20.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        20.0.0.0/8 is directly connected, Tunnel1
L        20.0.0.149/32 is directly connected, Tunnel1
B     192.168.0.0/16 [200/0], 00:02:26, Null0
B     192.168.40.0/24 [200/0] via 10.0.0.41, 00:02:26
                      [SEC][200/0] via 20.0.0.41, 00:02:26
B     192.168.50.0/24 [200/0] via 10.0.0.51, 00:01:55
                      [SEC][200/0] via 20.0.0.51, 00:01:55
      192.168.149.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.149.0/24 is directly connected, Ethernet1/0
L        192.168.149.1/32 is directly connected, Ethernet1/0
```

The following is the sample output for the **show ip cef** command.

```
Device# show ip cef 192.168.40.0 detail

192.168.40.0/24, epoch 0, flags [rib only nolabel, rib defined all labels]
  recursive via 10.0.0.41
    attached to Tunnel0

Device#show ip cef 192.168.40.0 int

192.168.40.0/24, epoch 0, flags [rnolbl, rlbls], RIB[B], refcnt 5, per-destination sharing
  sources: RIB
  feature space:
    IPRM: 0x00018000
  ifnums:
   Tunnel0(19): 10.0.0.41
  path list F3BDA6DC, 3 locks, per-destination, flags 0x269 [shble, rif, rcrsv, hwcn, bgp]
    path F3BDABAC, share 1/1, type recursive, for IPv4
     recursive via 10.0.0.41[IPv4:Default], fib F693B80C, 1 terminal fib,
v4:Default:10.0.0.41/32
     path list F3BDA72C, 2 locks, per-destination, flags 0x49 [shble, rif, hwcn]
          path F3BDAC14, share 1/1, type adjacency prefix, for IPv4
            attached to Tunnel0, IP midchain out of Tunnel0, addr 10.0.0.41 F555A5E0
  output chain:   IP midchain out of Tunnel0, addr 10.0.0.41 F555A5E0
                  IP adj out of Ethernet0/0, addr 11.0.0.41 F3CCCC10
```

The following is the sample output for the **show ip route** command for the repair paths.

```
Device# show ip route repair-paths

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, p - overrides from PfR
Gateway of last resort is not set
       10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.0.0.0/8 is directly connected, Tunnel0
L        10.0.0.149/32 is directly connected, Tunnel0
      20.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        20.0.0.0/8 is directly connected, Tunnel1
L        20.0.0.149/32 is directly connected, Tunnel1
B     192.168.0.0/16 [200/0], 00:02:26, Null0
B     192.168.40.0/24 [200/0] via 10.0.0.41, 00:00:10
                      [RPR][200/0] via 20.0.0.41, 00:00:10
B     192.168.50.0/24 [200/0] via 10.0.0.51, 00:00:10
                      [RPR][200/0] via 20.0.0.51, 00:00:10
      192.168.149.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.149.0/24 is directly connected, Ethernet1/0
L        192.168.149.1/32 is directly connected, Ethernet1/0
```

The following is the sample output for the **show ip route** command.

```
Device# show ip route
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, p - overrides from PfR
Gateway of last resort is not set
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.0.0.0/8 is directly connected, Tunnel0
L        10.0.0.149/32 is directly connected, Tunnel0
      20.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        20.0.0.0/8 is directly connected, Tunnel1
L        20.0.0.149/32 is directly connected, Tunnel1
D     192.168.0.0/16 is a summary, 00:08:53, Null0
D     192.168.40.0/24 [90/30378666] via 20.0.0.41, 00:08:45, Tunnel1
D     192.168.50.0/24 [90/30378666] via 20.0.0.51, 00:08:34, Tunnel1
      192.168.149.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.149.0/24 is directly connected, Ethernet1/0
L        192.168.149.1/32 is directly connected, Ethernet1/0
```

The following is the sample output for the **show ip route** command for the secondary path.

```
Device#  show ip route sec
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, p - overrides from PfR
Gateway of last resort is not set
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.0.0.0/8 is directly connected, Tunnel0
L        10.0.0.149/32 is directly connected, Tunnel0
      20.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
```

```
C        20.0.0.0/8 is directly connected, Tunnel1
L        20.0.0.149/32 is directly connected, Tunnel1
D     192.168.0.0/16 is a summary, 00:08:53, Null0
D     192.168.40.0/24 [90/30378666] via 20.0.0.41, 00:08:45, Tunnel1
                      [SEC][90/31232000] via 10.0.0.41, 00:08:45, Tunnel0
D     192.168.50.0/24 [90/30378666] via 20.0.0.51, 00:08:34, Tunnel1
                      [SEC][90/31232000] via 10.0.0.51, 00:08:34, Tunnel0
      192.168.149.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.149.0/24 is directly connected, Ethernet1/0
L        192.168.149.1/32 is directly connected, Ethernet1/0
```

Troubleshooting

```
NHRP-INT: Multipath nexthop lookup requested(if_in:, netid:1) for 192.168.50.1 in vrf
global(0x0)
NHRP-INT: Multipath recursive lookup for 192.168.50.1 (192.168.50.0/24)
NHRP-INT:  Path(1/1): [0x1]192.168.50.0/24 via 20.0.0.51, Tunnel1
NHRP-INT: Current first level nexthop: 20.0.0.51
NHRP-INT: Path(1) for 192.168.50.1 in vrf global(0x0) recursively resolved to 20.0.0.51,
Tunnel1, path metric: D/O/, l1_nhop 20.0.0.51
NHRP-INT: Found a better path(old: X/X//0, new: D/O//1); updating nhop: 20.0.0.51, Tunnel1
NHRP-INT: Updated best path to 20.0.0.51, Tunnel1(D/O/)
NHRP-INT:  Path(2/1): [0x100]192.168.50.0/24 via 10.0.0.51, Tunnel0
NHRP-INT: Current first level nexthop: 10.0.0.51
NHRP-INT: Path(2) for 192.168.50.1 in vrf global(0x0) recursively resolved to 10.0.0.51,
Tunnel0, path metric: S/C/, l1_nhop 10.0.0.51
NHRP-INT: Found a better path(old: D/O//1, new: S/C//1); updating nhop: 10.0.0.51, Tunnel0
NHRP-INT: Updated best path to 10.0.0.51, Tunnel0(S/C/)
NHRP-INT: Multipath recursive path walk(if_in:, netid:1) for 192.168.50.1(pfx:192.168.50.0/24)
 in global(0x0) yielded 10.0.0.51, Tunnel0
NHRP-DETAIL: Multipath recursive nexthop lookup(if_in:, netid:1) for 192.168.50.1 in
global(0x0) yielded 10.0.0.51, Tunnel0
..
```

Additional References for DMVPN Support for IWAN

Related Documents

Related Topic	Document Title
Cisco Intelligent WAN - An SD-WAN Solution	Cisco Intelligent WAN - An SD-WAN Solution

MIBs

MIB	MIBs Link
• CISCO-NHRP-MIB	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Feature Information for DMVPN Support for IWAN

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for DMVPN Support for IWAN

Feature Name	Releases	Feature Information
DMVPN Multiple Tunnel Termination	Cisco IOS XE Denali 16.3.2 Cisco IOS XE Everest 16.4.1	The following command was introduced by this feature: **maximum-secondary-paths**.