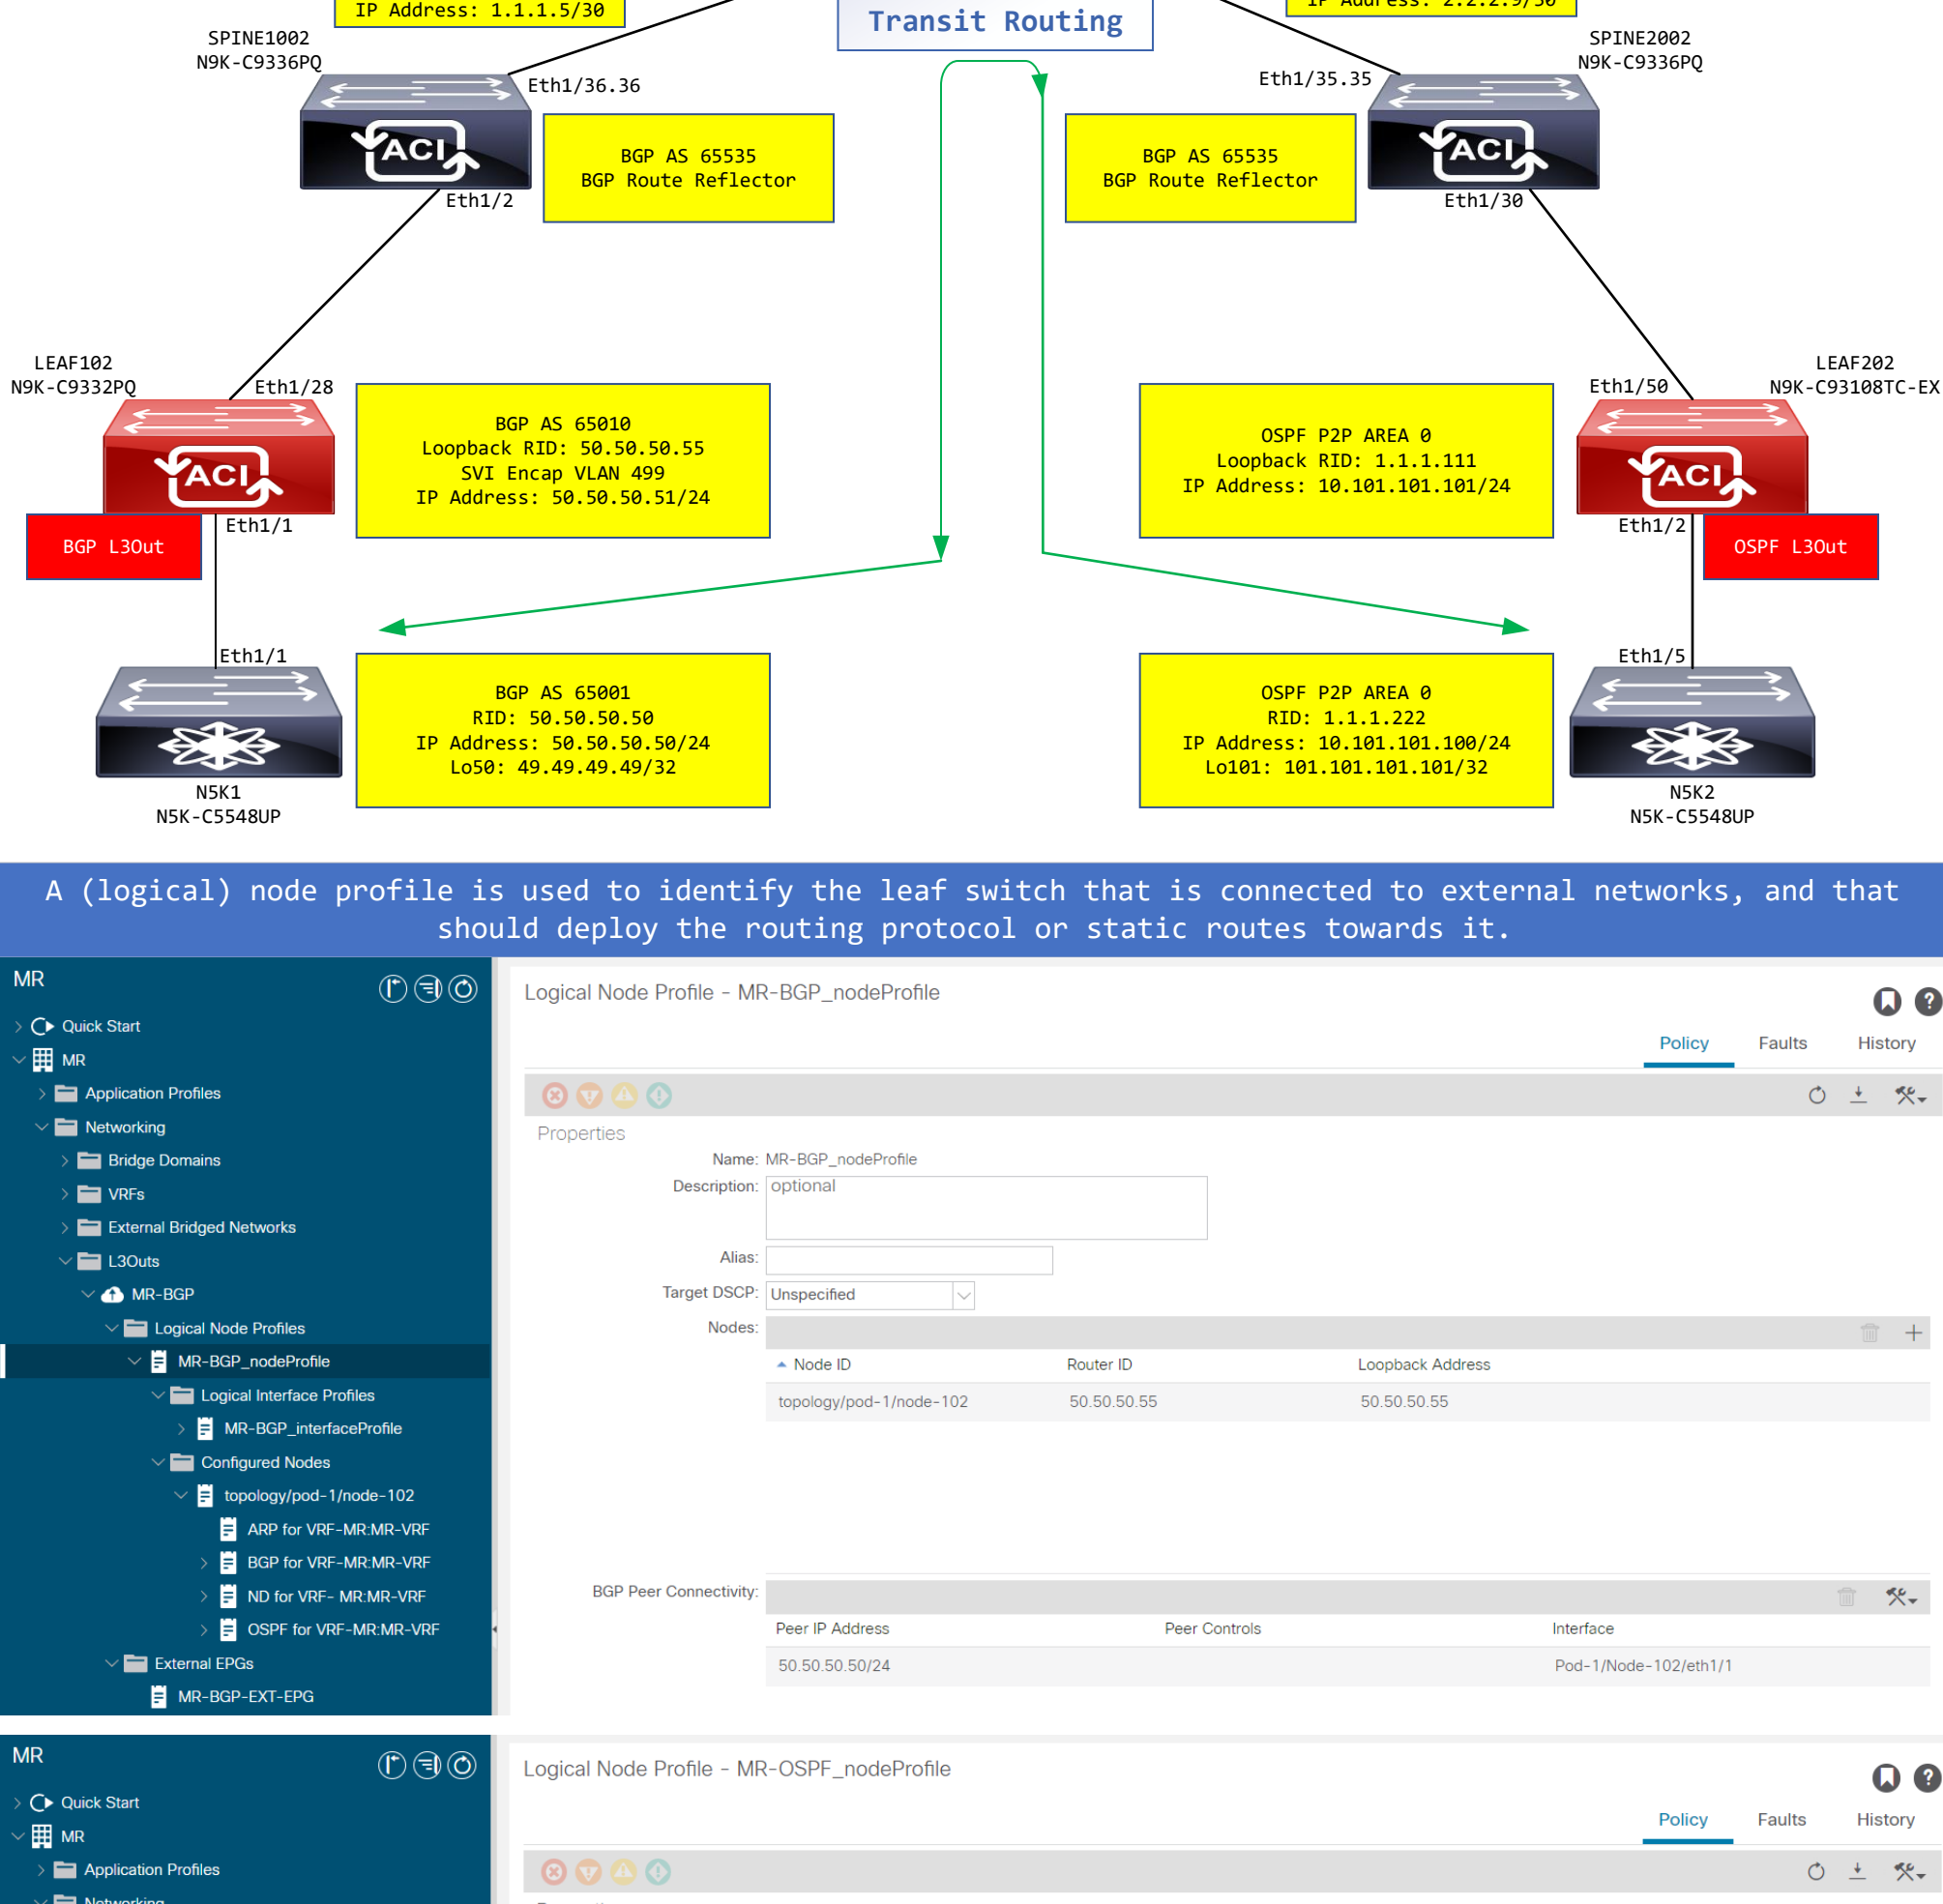


ACI Transit Routing

In transit routing, the Cisco Application Centric Infrastructure (Cisco ACI) fabric advertises the routes that are learned from one Layer 3 Out (L3Out) connection to another L3Out connection. The external Layer 3 domains peer with the fabric on the border leaf switches. The fabric is a transit Multiprotocol Border Gateway Protocol (MP-BGP) domain between the peers.

Topology:



A (logical) node profile is used to identify the leaf switch that is connected to external networks, and that should deploy the routing protocol on static routes towards it.

MR

Logical Node Profile - MR-BGP_nodeProfile

Name: MR-BGP_nodeProfile
Description: optional

Target OSPF: Unspecified

Nodes:

Node ID	Router ID	Loopback Address
Topology/leaf-2node-102	50.50.50.50	50.50.50.50

BGP Peer Connectivity:

Peer IP Address	Peer Configs	Interface
50.50.50.50/24		Port-100h-102h/eth1/1

MR

Logical Node Profile - MR-OSPF_nodeProfile

Name: MR-OSPF_nodeProfile
Description: optional

Target OSPF: Unspecified

Nodes:

Node ID	Router ID	Loopback Address
Topology/leaf-2node-202	1.1.1.1/32	1.1.1.1/32

Create BGP Protocol Profile:

A (logical) interface profile is used to identify the L3Out interface that connects to the external device. You will see several function elements that are defined for virtual routing and forwarding (VRF): Address Resolution Protocol (ARP), BGP, Neighbor Discovery, and OSPF as a consequence of both profiles.

MR

Logical Interface Profile - MR-OSPF_interfaceProfile

Name: MR-OSPF_interfaceProfile
Description: optional

Target OSPF: Unspecified

Nodes:

Node ID	Router ID	Loopback Address
Topology/leaf-2node-202	1.1.1.1/32	1.1.1.1/32

Create BGP Protocol Profile:

An external network instance profile (External EPG, L3Out EPG) represents a group of external subnets that have the same security behavior. Other subnets can also associate to other scopes, which define the routing behavior for that subnet.

MR

External EPG Instance Profile - MR-BGP_EXT-EPG

Name: MR-BGP_EXT-EPG
Description: optional

Contract Exception Tag: 49159

Resolved VRF Name: MR-VRF

Resolved VRF: MR-VRF

QoS Class: Unspecified

Target OSPF: Unspecified

Configuration Status: applied

Preferred Group Member:

Subnets	IP Address	Scope	Name	Aggregate	Route Control Profile	Route Summarization Policy
<input type="checkbox"/>	49.49.49.32	External Subnets for th...				

MR

External EPG Instance Profile - MR-OSPF_EXT-EPG

Name: MR-OSPF_EXT-EPG
Description: optional

Contract Exception Tag: 49159

Resolved VRF Name: MR-VRF

Resolved VRF: MR-VRF

QoS Class: Unspecified

Target OSPF: Unspecified

Configuration Status: applied

Preferred Group Member:

Subnets	IP Address	Scope	Name	Aggregate	Route Control Profile	Route Summarization Policy
<input type="checkbox"/>	101.101.101.32	External Subnets for th...				

The MR-PERMIT-ICMP contract is applied as both a provided and consumed contract in both external EPGs.

MR

External EPG Instance Profile - MR-BGP_EXT-EPG

Name	Tenant	Tenant Alias	Contract Type	Provided / Consumed	QoS Class	State	Label	Subject Label
MR-PERMIT-ICMP	MR		Contract	Provided	Unspecified	Format		
MR-PERMIT-ICMP	MR		Contract	Consumed	Unspecified	Format		

MR

External EPG Instance Profile - MR-OSPF_EXT-EPG

Name	Tenant	Tenant Alias	Contract Type	Provided / Consumed	QoS Class	State	Label	Subject Label
MR-PERMIT-ICMP	MR		Contract	Provided	Unspecified	Format		
MR-PERMIT-ICMP	MR		Contract	Consumed	Unspecified	Format		

On LEAF102, BGP is established with neighbor 50.50.50.50 and is receiving the external network 49.49.49.32.

MR

BGP Peer Entry - 50.50.50.50

Properties:

Neighbor ID: 50.50.50.50
Remote ID: 50.50.50.50
BGP State: Established
Neighbor Capability: 40000
Hold Time: 180
Keepalive Interval: 60
Neighbor: 50.50.50.50
Peer Link: 49SP
Peer Role: 1
ShutDown Reason: Unspecified
State Reason: none
Droptail: 100
Top Most Administrative: disabled
Connection Established: 1
Permitted Advertisements: 14
Connection Attempts: na

Message Statistics:

Opens	Sent	Recv
1	1	1

On LEAF202, OSPF is established with neighbor 1.1.1.222 and is receiving the external network 101.101.101.32.

MR

OSPF - MR-MR-VRF

Properties:

Name: MR-MR-VRF
Route ID: 1.1.1.1/32
Distance: 110
Max Hops: 8
Bandwidth Reference: 4000
Operational Status: Up

Neighbors:

Neighbor ID	State	Peer IP	Interface
1.1.1.222	Full	10.101.101.100	eth1/2

Interior Protocol Route Leak into OSPF:

Name	Distribution Protocol	Route Map	Scope	Cost
MR-MR-VRF	BGP	exp-ctrl-2555906	Inter-protocol leak	6501
MR-MR-VRF	OSPF	exp-ctrl-2555906	Inter-protocol leak	1
MR-MR-VRF	OSPF	exp-ctrl-2555906	Inter-protocol leak	1
MR-MR-VRF	OSPF	exp-ctrl-2555906	Inter-protocol leak	1

On both LEAF102 and LEAF202, the MR-BGP table on the VRF shows the external BGP network, 49.49.49.32, but it appears as external on LEAF102 and internal on LEAF202.

The OSPF external network 101.101.101.32 also appears in the BGP tables on both leaves, on LEAF202 as redistributed from OSPF and on LEAF102 as internal.

MR

LEAF102# show ip bgp vrf MR:MR-VRF

BGP routing table information for VRF MR:MR-VRF, address family IPv4 Unicast

BGP table version is 119, local router ID is 10.101.101.100

Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best

Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected

Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 102:2555906 (VRF MR:MR-VRF)					
*49.49.49.32	50.50.50.50	0	100	0	65010 65001 i
*101.101.101.32	20.0.248.0	41	100	0	?

MR

LEAF202# show ip bgp vrf MR:MR-VRF

BGP routing table information for VRF MR:MR-VRF, address family IPv4 Unicast

BGP table version is 95, local router ID is 20.0.248.0

Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best

Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected

Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 202:2555906 (VRF MR:MR-VRF)					
*149.49.49.32	10.0.232.68	100	0	0	65010 65001 i
*101.101.101.32	0.0.0.0	41	100	0	32768 ?

The BGP IPv4 table contains equivalent information.

MR

LEAF102# show ip bgp vrf MR:MR-VRF

BGP routing table information for VRF MR:MR-VRF, address family IPv4 Unicast

BGP table version is 37, local router ID is 50.50.50.55

Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best

Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected

Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 102:2555906 (VRF MR:MR-VRF)					
*49.49.49.32	50.50.50.50	0	100	0	65010 65001 i
*101.101.101.32	20.0.248.0	41	100	0	?

MR

LEAF202# show ip bgp vrf MR:MR-VRF

BGP routing table information for VRF MR:MR-VRF, address family IPv4 Unicast

BGP table version is 31, local router ID is 1.1.1.111

Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best

Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected

Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 102:2555906 (VRF MR:MR-VRF)					
*149.49.49.32	10.0.232.68	100	0	0	65010 65001 i
*101.101.101.32	0.0.0.0	41	100	0	32768 ?

However, the OSPF external network 101.101.101.32 is not in the routing table of NSK1.

Similarly, the BGP external network 49.49.49.32 is not in NSK2's RIB.

MR

NSK1# show ip route vrf MR-BGP

IP Route Table for VRF "MR-BGP"

* denotes best ucst next-hop

*x/y denotes [preference/metric]

*<string> in via output denotes VRF <string>

```
49.49.49.32, ubest/mbest: 2/0, attached
  *via 49.49.49.49, Lo50, [0/0], 1d08h, local
  *via 49.49.49.49, Lo50, [0/0], 1d08h, direct
50.50.50.0/24, ubest/mbest: 1/0, attached
  *via 50.50.50.50, Vlan499, [0/0], 1d07h, direct
50.50.50.32, ubest/mbest: 1/0, attached
  *via 50.50.50.50, Vlan499, [0/0], 1d08h, local
101.101.101.32, ubest/mbest: 1/0
  *via 50.50.50.50, Vlan499, [0/0], 1d07h, local
```

In the BGP L3Out, enter the IP address of the external subnet received from the OSPF L3Out, 101.101.101.32, choose Export Route Control Subnet in the Route Control section and clear the External Subnets for the External EPG classification. Click Submit.

The Export Route Control Subnet option allows a network to be exported (advertised) to the external peer.

MR

External EPG Instance Profile - MR-BGP_EXT-EPG

Preferred Group Member: Exclude Include

Subnets:

Subnets	IP Address	Scope	Name	Aggregate	Route Control Profile	Route Summarization Policy
<input type="checkbox"/>	49.49.49.32	External Subnets for th...				

Create Subnet:

IP Address: 101.101.101.32

Name: 100h/101

Route Control:

Export Route Control Subnet

Import Route Control Subnet

Shared Route Control Subnet

Aggregate:

Aggregate Exact

Aggregate Import

Aggregate Shared Routes

Route Summarization Policy: select an option

Route Summarization Policy:

Route Control Profile:

Name	Direction

Route control is used for filtering external routes advertised out of the fabric, allowed into the fabric, or leaked to other VRFs within the fabric.

External EPG classification:

External Subnets for External EPG

Shared Security Import Subnet

External EPG classification is used to identify the external networks associated with this external EPG for policy enforcement (Contracts).

Cancel Submit

On NSK1, the OSPF external network is now received over BGP.

MR

NSK1# show ip route vrf MR-BGP

IP Route Table for VRF "MR-BGP"

* denotes best ucst next-hop

*x/y denotes [preference/metric]

*<string> in via output denotes VRF <string>

```
49.49.49.32, ubest/mbest: 2/0, attached
  *via 49.49.49.49, Lo50, [0/0], 1d08h, local
  *via 49.49.49.49, Lo50, [0/0], 1d08h, direct
50.50.50.0/24, ubest/mbest: 1/0, attached
  *via 50.50.50.50, Vlan499, [0/0], 1d08h, direct
50.50.50.32, ubest/mbest: 1/0, attached
  *via 50.50.50.50, Vlan499, [0/0], 1d08h, local
101.101.101.32, ubest/mbest: 1/0
  *via 50.50.50.51, [20/0], 00:00:03, BGP-65001, external, tag 65010,
```

In the OSPF L3Out, enter the IP address of the external subnet received from the BGP L3Out, 49.49.49.32, choose Export Route Control Subnet in the Route Control section and clear External Subnets for the External EPG classification. Click Submit.

MR

External EPG Instance Profile - MR-OSPF_EXT-EPG

Preferred Group Member: Exclude Include

Subnets:

Subnets	IP Address	Scope	Name	Aggregate	Route Control Profile	Route Summarization Policy
<input type="checkbox"/>	101.101.101.32	External Subnets for th...				

Create Subnet:

IP Address: 49.49.49.32

Name: 100h/101

Route Control:

Export Route Control Subnet

Import Route Control Subnet

Shared Route Control Subnet

Aggregate:

Aggregate Exact

Aggregate Import

Aggregate Shared Routes

Route Summarization Policy: select an option

Route Summarization Policy:

Route Control Profile:

Name	Direction

Route control is used for filtering external routes advertised out of the fabric, allowed into the fabric, or leaked to other VRFs within the fabric.

External EPG classification:

External Subnets for External EPG

Shared Security Import Subnet

External EPG classification is used to identify the external networks associated with this external EPG for policy enforcement (Contracts).

Cancel Submit

Now on NSK2, the BGP external network is now received over OSPF.

MR

NSK2# show ip route vrf MR-OSPF

IP Route Table for VRF "MR-OSPF"

* denotes best ucst next-hop

*x/y denotes [preference/metric]

*<string> in via output denotes VRF <string>

```
1.1.1.111/32, ubest/mbest: 1/0
  *via 101.101.101.100, Eth1/5, [110/41], 2d05h, ospf-1, intra
10.101.101.0/24, ubest/mbest: 1/0, attached
  *via 10.101.101.100, Eth1/5, [0/0], 6d22h, direct
10.101.101.32, ubest/mbest: 1/0, attached
  *via 10.101.101.100, Eth1/5, [0/0], 6d22h, local
49.49.49.32, ubest/mbest: 1/0
  *via 10.101.101.101, Eth1/5, [110/1], 00:01:59, ospf-1, type-2, tag 4294967295,
101.101.101.101/32, ubest/mbest: 2/0, attached
  *via 101.101.101.101, Lo101, [0/0], 2d05h, local
  *via 101.101.101.101, Lo101, [0/0], 2d05h, direct
```

Ping works between the two networks because of the MR-PERMIT-ICMP contract that was applied to both external EPGs earlier.

MR

NSK1# ping 101.101.101.101 vrf MR-BGP source 49.49.49.49

```
PING 101.101.101.101: 56 data bytes
64 bytes from 101.101.101.101: icmp_seq=0 ttl=252 time=3.059 ms
64 bytes from 101.101.101.101: icmp_seq=1 ttl=252 time=2.963 ms
64 bytes from 101.101.101.101: icmp_seq=2 ttl=252 time=2.928 ms
64 bytes from 101.101.101.101: icmp_seq=3 ttl=252 time=2.984 ms
64 bytes from 101.101.101.101: icmp_seq=4 ttl=252 time=2.982 ms
-- 101.101.101.101 ping statistics --
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 2.954/3.077/3.928 ms
```

MR

NSK2# ping 49.49.49.49 vrf MR-OSPF source 101.101.101.101

```
PING 49.49.49.49: 56 data bytes
64 bytes from 49.49.49.49: icmp_seq=0 ttl=252 time=1.107 ms
64 bytes from 49.49.49.49: icmp_seq=1 ttl=252 time=2.99 ms
64 bytes from 49.49.49.49: icmp_seq=2 ttl=252 time=2.98 ms
64 bytes from 49.49.49.49: icmp_seq=3 ttl=252 time=2.98 ms
64 bytes from 49.49.49.49: icmp_seq=4 ttl=252 time=2.99 ms
-- 49.49.49.49 ping statistics --
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 2.983/3.013/3.107 ms
```

