



DHCP Relay in VXLAN BGP EVPN

This appendix contains the following sections:

- [DHCP Relay in VXLAN BGP EVPN Overview, on page 1](#)
- [Guidelines and Limitations for DHCP Relay, on page 2](#)
- [DHCP Relay in VXLAN BGP EVPN Example, on page 2](#)
- [Configuring VPC Peers Example, on page 19](#)
- [vPC VTEP DHCP Relay Configuration Example, on page 20](#)

DHCP Relay in VXLAN BGP EVPN Overview

DHCP relay is supported by VXLAN BGP EVPN and is useful in a multi-tenant VXLAN EVPN deployment to provision DHCP service to EVPN tenant clients.

In a multi-tenant EVPN environment, DHCP relay uses the following sub-options of Option 82:

- Sub-option 151(0x97) - Virtual Subnet Selection

(Defined in RFC#6607.)

Used to convey VRF related information to the DHCP server in an MPLS-VPN and VXLAN EVPN multi-tenant environment.

- Sub-option 11(0xb) - Server ID Override

(Defined in RFC#5107.)

The server identifier (server ID) override sub-option allows the DHCP relay agent to specify a new value for the server ID option, which is inserted by the DHCP server in the reply packet. This sub-option allows the DHCP relay agent to act as the actual DHCP server such that the renew requests will come to the relay agent rather than the DHCP server directly. The server ID override sub-option contains the incoming interface IP address, which is the IP address on the relay agent that is accessible from the client. Using this information, the DHCP client sends all renew and release request packets to the relay agent. The relay agent adds all of the appropriate sub-options and then forwards the renew and release request packets to the original DHCP server. For this function, Cisco's proprietary implementation is sub-option 152(0x98). You can use the **ip dhcp relay sub-option type cisco** command to manage the function.

- Sub-option 5(0x5) - Link Selection

(Defined in RFC#3527.)

The link selection sub-option provides a mechanism to separate the subnet/link on which the DHCP client resides from the gateway address (giaddr), which can be used to communicate with the relay agent by the DHCP server. The relay agent will set the sub-option to the correct subscriber subnet and the DHCP server will use that value to assign an IP address rather than the giaddr value. The relay agent will set the giaddr to its own IP address so that DHCP messages are able to be forwarded over the network. For this function, Cisco's proprietary implementation is sub-option 150(0x96). You can use the **ip dhcp relay sub-option type cisco** command to manage the function.


Guidelines and Limitations for DHCP Relay

The following are the guidelines and limitations for DHCP Relay in VXLAN BGP EVPN:

- Beginning in Cisco NX-OS Release 9.2(2), support is added for Cisco Nexus 3636C-R and 36180YC-R.
- IPv6 DHCP is not supported for Cisco Nexus 3636C-R and 36180YC-R switches.

DHCP Relay in VXLAN BGP EVPN Example

Figure 1: Example Topology



Topology characteristics:

- Switches 3636-1 and 3636-2 are VTEPs connected to VXLAN fabric.
- Client1 and client2 are DHCP clients in vlan1001. They belong to tenant VRF vxlan-900001.

- The DHCP server is ASR1K, a router that sits in vlan10.
- DHCP server configuration

```
ip vrf vxlan900001
ip dhcp excluded-address vrf vxlan900001 172.16.16.1 172.16.16.9
ip dhcp pool one
 vrf vxlan900001
 network 172.16.16.0 255.255.240.0
 defaultrouter 172.16.16.1
```

Basic VXLAN BGP EVPN Configuration

- 3636-1

```
version 7.0(3)I1(3)

hostname 3636C-R

nv overlay evpn
feature vn-segment-vlan-based
feature nv overlay

fabric forwarding anycast-gateway-mac 0000.1111.2222

vlan 101
 vn-segment 900001
vlan 1001
 vn-segment 2001001

vrf context vxlan-900001
 vni 900001
 rd auto
 address-family ipv4 unicast
  route-target both auto
  route-target both auto evpn

interface Vlan101
 no shutdown
 vrf member vxlan-900001
 ip forward

interface Vlan1001
 no shutdown
 vrf member vxlan-900001
 ip address 172.16.16.1/20
 fabric forwarding mode anycast-gateway

interface nve1
 no shutdown
 source-interface loopback1
 host-reachability protocol bgp
 member vni 10000 associate-vrf
 mcast-group 224.1.1.1
 member vni 10001 associate-vrf
 mcast-group 224.1.1.1
```

```

member vni20000
suppress-arp
mcast-group 225.1.1.1
member vni 20001
suppress-arp
mcast-group 225.1.1.1

interfaca Ethernet1/49
switchport mode trunk
switchport trunk allowed vlan 10,1001
spanning-tree port type edge trunk

interface Ethernet1/50
no switchport
ip address 192.1.33.2/24
ip router ospf 1 area 0.0.0.0
ip pirt sparse-mode
no shutdown

interface loopback0
ip address 1.1.1.1/32
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode

interface loopback1
vrf member vxlan-900001
ip address 11.11.11.11/32

router bgp 65535
router-id 1.1.1.1
log-neighbor-changes
neighbor 2.2.2.2 remote-as 65535
update-source loopback0
address-family l2vpn evpn
send-community both
vrf vxlan-900001
address-family ipv4 unicast
network 11.11.11.11/32
network 192.1.42.0/24
advertise l2vpn evpn
evpn
vni 2001001 12

rd auto
route-target import auto
route-target export auto

```

- 3636-2

```

version 7.0(3)I1(3)
hostname 3636-1

nv overlay evpn
feature vn-segment-vlan-based
feature nv overlay

fabric forwarding anycast-gateway-mac 0000.1111.2222

vlan 101

```

```
vn-segment 900001
vlan 1001
  vn-segment 2001001

vrf context vxlan-900001
  vni 900001
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn

interface Vianl01
no shutdown
vrf member vxlan-900001
ip forward

interface Vlan1001
no shutdown
vrf member vxlan-900001
ip address 172.16.16.1/20
fabric forwarding mcde anycast-gateway

rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn

interface Vianl01
no shutdown
vrf member vxlan-900001
ip forward

interface Vlan1001
no shutdown
vrf member vxlan-900001
ip address 172.16.16.1/20
fabric forwarding mcde anycast-gateway

interface nve1
no shutdown
source-interface loopback1
host-reachability protocol bgp
member vni 10000 associate-vrf
mcast-group 224.1.1.1
member vni 10001 associate-vrf
mcast-group 224.1.1.1
member vni20000
suppress-arp
mcast-group 225.1.1.1
member vni 20001
suppress-arp
mcast-group 225.1.1.1

interface Ethernet1/49
switchport mode trunk
switchport trunk allowed vlan 10,1001
spanning-tree port type edge trunk

interface Ethernet1/50
no switchport
ip address 192.1.34.2/24
```

```

ip router ospf 1 area 0.0.0.0
ip pim sparse-mode
no shutdown

interface loopback0
ip address 2.2.2.2/32
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode

interface loopback1
vrf member vxlan-900001
ip address 22.22.22.22/32

router bgp 65535
router-id 2.2.2.2
log-neighbor-changes
neighbor 1.1.1.1 remote-as 65535
update-source loopback0
address-family l2vpn evpn
send-community both
vrf vxlan-900001
address-family ipv4 unicast
network 22.22.22.22/32

advertise l2vpn evpn
evpn
vni 2001001 12

rd auto
route-target import auto
route-target export auto

```

DHCP Relay on VTEPs

The following are common deployment scenarios:

- Client on tenant VRF and server on Layer 3 default VRF.
- Client on tenant VRF (SVI X) and server on the same tenant VRF (SVI Y).
- Client on tenant VRF (VRF X) and server on different tenant VRF (VRF Y).
- Client on tenant VRF and server on non-default non-VXLAN VRF.

The following sections below move vlan10 to different VRFs to depict different scenarios.

Client on Tenant VRF and Server on Layer 3 Default VRF

Put DHCP server (192.1.42.3) into the default VRF and make sure it is reachable from both 3636-1 and 3636-2 through the default VRF.

```

3636-1# sh run int vl 10

!Command: show running-config interface Vlan10
!Time: Mon Aug 7 07:51:16 2018

version 9.2(1)

```

```
interface Vlan10
  no shutdown
  ip address 192.1.42.1/24
  ip router ospf 1 area 0.0.0.0

3636-1# ping 192.1.42.3 cou 1

PING 192.1.42.3 (192.1.42.3): 56 data bytes
64 bytes from 192.1.42.3: icmp_seq=0 ttl=254 time=0.593 ms
- 192.1.42.3 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
roundtrip min/avg/max = 0.593/0.592/0.593 ms

3636-2# ping 192.1.42.3 cou 1
PING 192.1.42.3 (192.1.42.3): 56 data bytes
64 bytes from 192.1.42.3: icmp_seq=0 ttl=252 time=0.609 ms
- 192.1.42.3 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 0.609/0.608/0.609 ms
```

DHCP Relay Configuration

- 3636-1

```
3636-1# sh run dhcp

!Command: show running-config dhcp
!Time: Mon Aug 6 08:26:00 2018

version 9.2(1)
feature dhcp

service dhcp
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ipv6 dhcp relay

interface Vlan1001
  ip dhcp relay address 192.1.42.3 use-vrf default
```

- 3636-2

```
3636-2# sh run dhcp

!Command: show running-config dhcp
!Time: Mon Aug 6 08:26:16 2018

version 9.2(1)
feature dhcp

service dhcp
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ipv6 dhcp relay

interfaoe Vlan1001
```

```
ip dhcp relay address 192.1.42.3 use-vrf default
```

Debug Output

- The following is a packet dump for DHCP interact sequences.

```
3636-1# ethanalyzer local interface inband display-filter
"udp.srcport==67 or udp.dstport==67" limit-captured frames 0

Capturing on inband
20150824 08:35:25.066530 0.0.0.0 -> 255.255.255.255 DHCP DHCP Discover - Transaction
ID 0x636a38fd
20150824 08:35:25.068141 192.1.42.1 -> 192.1.42.3 DHCP DHCP Discover - Transaction ID
0x636a38fd
20150824 08:35:27.069494 192.1.42.3 -> 192.1.42.1 DHCP DHCP Offer Transaction - ID
0x636a38fd
20150824 08:35:27.071029 172.16.16.1 -> 172.16.16.11 DHCP DHCP Offer Transaction - ID
0x636a38fd
20150824 08:35:27.071488 0.0.0.0 -> 255.255.255.255 DHCP DHCP Request Transaction - ID
0x636a38fd
20150824 08:35:27.072447 192.1.42.1 -> 192.1.42.3 DHCP DHCP Request Transaction - ID
0x636a38fd
20150824 08:35:27.073008 192.1.42.3 -> 192.1.42.1 DHCP DHCP ACK Transaction - ID
0x636a38fd
20150824 08:35:27.073692 172.16.16.1 -> 172.16.16.11 DHCP DHCP ACK Transaction - ID
0x636a38fd
```



Note Ethanalyzer might not capture all DHCP packets because of inband interpretation issues when you use the filter. You can avoid this by using SPAN.

- DHCP Discover packet 3636-1 sent to DHCP server.

giaddr is set to 192.1.42.1 (ip address of vlan10) and suboptions 5/11/151 are set accordingly.

```
Bootp flags: 0x0000 (unicast)
client IP address: 0.0.0.0 (0.0.0.0)
Your (client) IP address: 0.0.0.0 (0.0.0.0)
Next server IP address: 0.0.0.0 (0.0.0.0)
Relay agent IP address: 192.1.42.1 (192.1.42.1)
client MAC address Hughes_01:51:51 (00:00:10:01:51:51)
client hardware address padding: 00000000000000000000
Server host name not given
Boot file name not given
Magic cookie: DHCP
Option: (53) DHCP Message Type
  Length: 1
  DHCP: Discover (1)
Option: (55) Parameter Request List
  Length: 4
  Parameter Request List Item: (1) Subnet Mask
  Parameter Request List Item: (3) Router
  Parameter Request List Item: (58) Renewal Time Value
  Parameter Request List Item: (59) Rebinding Time Value
Option: (61) client identifier
  Length: 7
```



```

Hardware type: Ethernet (0x01)
Client MAC address: Hughes_01:51:51 (00:00:10:01:51:51)
Option: (82) Agent Information Option
Length: 47
Option 82 Suboption: (1) Agent Circuit ID
Length: 10
Agent Circuit ID: 01080006001e88690030
Option 82 Suboption: (2) Agent Remote ID
Length: 6
Agent Remote ID: f8c2882333a5
Option 82 Suboption: (151) VRF name/VPN ID
Option 82 Suboption: (11) Server ID Override
Length: 4
Server ID Override: 172.16.16.1 (172.16.16.1)
Option 82 Suboption: (5) Link selection
Length: 4
Link selection: 172.16.16.0 (172.16.16.0)

```

```

ASR1K-DHCP# sh ip dhcp bin
Bindings from all pools not associated with VRF:
IP address ClientID/ Lease expiration Type State Interface
      Hardware address/
      User name

```

```

Bindings from VRF pool vxlan900001:
IP address ClientID/ Lease expiration Type State Interface
      Hardware address/
      User name
172.16.16.10 0100.0010.0175.75 Aug 25 2015 09:21 AM Automatic Active GigabitEthernet2/1/0
172.16.16.11 0100.0010.0151.51 Aug 25 2015 08:54 AM Automatic Active GigabitEthernet2/1/0

```

```

3636-1# sh ip route vrf vxlan900001
IP Route Table for VRF "vxlan900001"
'*' denotes best ucast nexthop
'***' denotes best mcast nexthop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

11.11.11.11/32, ubest/mbest: 2/0, attached
  *via 11.11.11.11, Lo1, [0/0], 18:31:57, local
  *via 11.11.11.11, Lo1, [0/0], 18:31:57, direct
22.22.22.22/32, ubest/mbest: 1/0
  *via 2.2.2.2%default, [200/0], 18:31:57, bgp65535,internal, tag 65535 (evpn)segid:
900001 tunnelid: 0x2020202
encap: VXLAN

```

```

172.16.16.0/20, ubest/mbest: 1/0, attached
  *via 172.16.16.1, Vlan1001, [0/0], 18:31:57, direct
172.16.16.1/32, ubest/mbest: 1/0, attached
  *via 172.16.16.1, Vlan1001, [0/0], 18:31:57, local
172.16.16.10/32, ubest/mbest: 1/0
  *via 2.2.2.2%default, [200/0], 00:00:47, bgp65535,internal, tag 65535 (evpn)segid:
900001 tunnelid: 0x2020202
encap: VXLAN

```

```

172.16.16.11/32, ubest/mbest: 1/0, attached
  *via 172.16.16.11, Vlan1001, [190/0], 00:28:10, hmm

```

```

3636-1# ping 172.16.16.11 vrf vxlan900001 count 1
PING 172.16.16.11 (172.16.16.11): 56 data bytes
64 bytes from 172.16.16.11: icmp_seq=0 ttl=63 time=0.846 ms
- 172.16.16.11 ping statistics -

```

```

1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 0.846/0.845/0.846 ms

3636-1# ping 172.16.16.10 vrf vxlan900001 count 1
PING 172.16.16.10 (172.16.16.10): 56 data bytes
64 bytes from 172.16.16.10: icmp_seq=0 ttl=62 time=0.874 ms
- 172.16.16.10 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 0.874/0.873/0.874 ms

```

Client on Tenant VRF (SVI X) and Server on the Same Tenant VRF (SVI Y)

Put DHCP server (192.1.42.3) into VRF of vxlan-900001 and make sure it is reachable from both 3636-1 and 3636-2 through VRF of vxlan-900001.

```

3636-1# sh run int vl 10

!Command: show running-config interface Vlan10
!Time: Mon Aug 6 09:10:26 2018

version 9.2(1)

interface Vlan10
 no shutdown
 vrf member vxlan-900001
 ip address 192.1.42.1/24

```

Because 172.16.16.1 is an anycast address for vlan1001 configured on all the VTEPs, we need to pick up a unique address as the DHCP relay packet's source address to make sure the DHCP server can deliver a response to the original DHCP Relay agent. In this scenario, we use loopback1 and we need to make sure loopback1 is reachable from everywhere of VRF vxlan-900001.

```

3636-1# sh run int lo1

!Command: show running-config interface loopback1
!Time: Mon Aug 6 09:18:53 2018

version 9.2(1)

interface loopback1
 vrf member vxlan-900001
 ip address 11.11.11.11/32

3636-1# ping 192.1.42.3 vrf vxlan900001 source 11.11.11.11 cou 1
PING 192.1.42.3 (192.1.42.3) from 11.11.11.11: 56 data bytes
64 bytes from 192.1.42.3: icmp_seq=0 ttl=254 time=0.575 ms
- 192.1.42.3 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 0.575/0.574/0.575 ms

3636-2# sh run int lo1

!Command: show running-config interface loopback1
!Time: Mon Aug 6 09:19:30 2018

version 9.2(1)

```

```

interface loopback1
  vrf member vxlan900001
  ip address 22.22.22.22/32

3636-2# ping 192.1.42.3 vrf vxlan-900001 source 22.22.22.22 cou 1
PING 192.1.42.3 (192.1.42.3) from 22.22.22.22: 56 data bytes
64 bytes from 192.1.42.3: icmp_seq=0 ttl=253 time=0.662 ms
- 192.1.42.3 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 0.662/0.662/0.662 ms

```

DHCP Relay Configuration

- 3636-1

```

3636-1# sh run dhcp

!Command: show running-config dhcp
!Time: Mon Aug 6 08:26:00 2018

version 9.2(1)
feature dhcp

service dhcp
ip dhcp relay
ip dhcp relay information option
I4ip dhcp relay information option vpn
ipv6 dhcp relay

interface Vlan1001
  ip dhcp relay address 192.1.42.3
  ip dhcp relay source-interface loopback1

```

- 3636-2

```

3636-2# sh run dhcp

!Command: show running-config dhcp
!Time: Mon Aug 6 08:26:16 2018

version 9.2(1)
feature dhcp

service dhcp
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ipv6 dhcp relay

interface Vlan1001
  ip dhcp relay address 192.1.42.3
  ip dhcp relay source-interface loopback1

```

Debug Output

- The following is a packet dump for DHCP interact sequences.

```

3636-1# ethanalyzer local interface inband display-filter
"udp.srcport==67 or udp.dstport==67" limit-captured frames 0

Capturing on inband
20150824 09:31:38.129393 0.0.0.0 -> 255.255.255.255 DHCP DHCP Discover - Transaction
ID 0x860cd13
20150824 09:31:38.129952 11.11.11.11 -> 192.1.42.3 DHCP DHCP Discover - Transaction ID
0x860cd13
20150824 09:31:40.130134 192.1.42.3 -> 11.11.11.11 DHCP DHCP Offer - Transaction ID
0x860cd13
20150824 09:31:40.130552 172.16.16.1 -> 172.16.16.11 DHCP DHCP Offer - Transaction ID
0x860cd13
20150824 09:31:40.130990 0.0.0.0 -> 255.255.255.255 DHCP DHCP Request - Transaction ID
0x860cd13
20150824 09:31:40.131457 11.11.11.11 -> 192.1.42.3 DHCP DHCP Request - Transaction ID
0x860cd13
20150824 09:31:40.132009 192.1.42.3 -> 11.11.11.11 DHCP DHCP ACK - Transaction ID
0x860cd13
20150824 09:31:40.132268 172.16.16.1 -> 172.16.16.11 DHCP DHCP ACK - TransactionID
0x860cd13

```



Note Ethanalyzer might not capture all DHCP packets because of inband interpretation issues when you use the filter. You can avoid this by using SPAN.

- DHCP Discover packet 3636-1 sent to DHCP server.
giaddr is set to 11.11.11.11(loopback1) and suboptions 5/11/151 are set accordingly.

```

Bootstrap Protocol
Message type: Boot Request (1)
Hardware type: Ethernet (0x01)
Hardware address length: 6
Hops: 1
Transaction ID: 0x0860cd13
Seconds elapsed: 0
Bootp flags: 0x0000 (unicast)
Client IP address: 0.0.0.0 (0.0.0.0)
Your (client) IP address: 0.0.0.0 (0.0.0.0)
Next server IP address: 0.0.0.0 (0.0.0.0)
Relay agent iP address: 11.11.11.11 (11.11.11.11)
Client MAC address: Hughes_01:51:51 (00:00:10:01:51:51)
Client hardware address padding: 00000000000000000000
Server host name not given
Boot file name not given
Magic cookie: DHCP
Option: (53) DHCP Message Type
Length: 1
DHCP: Discover (1)
Option: (55) Parameter Request List
Option: (61) Client Identifier
Option: (82) Agent Information Option
Length: 47
Option 82 suboption: (1) Agent Circuit ID
Option 82 suboption: (151) Agent Remote ID
Option 82 suboption: (11) Server ID Override
Length: 4

```

```

Server ID override: 172.16.16.1 (172.16.16.1)
Option 82 suboption: (5) Link selection
Length: 4
Link selection: 172.16.16.0 (172.16.16.0)

```

```

ASR1K-DHCP# sh ip dhcp bin
Bindings from all pools not associated with VRF:
IP address ClientID/Lease expiration Type State Interface
Hardware address/
User name

Bindings from VRF pool vxlan-900001:
IP address ClientID/Lease expiration Type State Interface
Hardware address/
User name

172.16.16.10 0100.0010.0175.75 Aug 25 2015 10:02 AM Automatic Active GigabitEthernet2/1/0
172.16.16.11 0100.0010.0151.51 Aug 25 2015 09:50 AM Automatic Active GigabitEthernet2/1/0

3636-1# sh ip route vrf vxlan-900001
IP Route Table for VRF "vxlan-900001"
'*' denotes best ucast nexthop
'**' denotes best mcast nexthop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

11.11.11.11/32, ubest/mbest: 2/0, attached
  *via 11.11.11.11, Lo1, [0/0], 19:13:56, local
  *via 11.11.11.11, Lo1, [0/0], 19:13:56, direct
22.22.22.22/32, ubest/mbest: 1/0
  *via 2.2.2.2%default, [200/0], 19:13:56, bgp65535,internal, tag 65535 (evpn)segid:
900001 tunnelid: 0x2020202
encap: VXLAN
172.16.16.0/20, ubest/mbest: 1/0, attached
  *via 172.16.16.1, Vlan1001, [0/0], 19:13:56, direct
172.16.16.1/32, ubest/mbest: 1/0, attached
  *via 172.16.16.1, Vlan1001, [0/0], 19:13:56, local
172.16.16.10/32, ubest/mbest: 1/0
  *via 2.2.2.2%default, [200/0], 00:01:27, bgp65535,
internal, tag 65535 (evpn)segid: 900001 tunnelid: 0x2020202
encap: VXLAN
172.16.16.11/32, ubest/mbest: 1/0, attached
  *via 172.16.16.11, Vlan1001, [190/0], 00:13:56, hmm
192.1.42.0/24, ubest/mbest: 1/0, attached
  *via 192.1.42.1, Vlan10, [0/0], 00:36:08, direct
192.1.42.1/32, ubest/mbest: 1/0, attached
  *via 192.1.42.1, Vlan10, [0/0], 00:36:08, local
9372-1# ping 172.16.16.10 vrf vxlan-900001 cou 1
PING 172.16.16.10 (172.16.16.10): 56 data bytes
64 bytes from 172.16.16.10: icmp_seq=0 ttl=62 time=0.808 ms
- 172.16.16.10 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 0.808/0.808/0.808 ms

3636-1# ping 172.16.16.11 vrf vxlan-900001 cou 1
PING 172.16.16.11 (172.16.16.11): 56 data bytes
64 bytes from 172.16.16.11: icmp_seq=0 ttl=63 time=0.872 ms
- 172.16.16.11 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 0.872/0.871/0.872 ms

```

Client on Tenant VRF (VRF X) and Server on Different Tenant VRF (VRF Y)

The DHCP server is placed into another tenant VRF vxlan-900002 so that DHCP response packets can access the original relay agent. We use loopback2 to avoid any anycast ip address that is used as the source address for the DHCP relay packets.

```
3636-1# sh run int vl 10
!Command: show runningconfig interface Vlan10
!Time: Tue Aug 6 08:48:22 2018

version 9.2(1)
interface Vlan10
  no shutdown
  vrf member vxlan900002
  ip address 192.1.42.1/24

3636-1# sh run int lo2
!Command: show runningconfig interface loopback2
!Time: Tue Aug 7 08:48:57 2018
version 9.2(1)
interface loopback2
  vrf member vxlan900002
  ip address 33.33.33.33/32

3636-2# sh run int lo2
!Command: show runningconfig interface loopback2
!Time: Tue Aug 7 08:48:44 2018
version 9.2(1)
interface loopback2
  vrf member vxlan900002
  ip address 44.44.44.44/32

9372-1# ping 192.1.42.3 vrf vxlan-900002 source 33.33.33.33 cou 1
PING 192.1.42.3 (192.1.42.3) from 33.33.33.33: 56 data bytes
64 bytes from 192.1.42.3: icmp_seq=0 ttl=254 time=0.544 ms
- 192.1.42.3 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 0.544/0.544/0.544 ms

3636-2# ping 192.1.42.3 vrf vxlan-900002 source 44.44.44.44 count 1
PING 192.1.42.3 (192.1.42.3) from 44.44.44.44: 56 data bytes
64 bytes from 192.1.42.3: icmp_seq=0 ttl=253 time=0.678 ms
- 192.1.42.3 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 0.678/0.678/0.678 ms
```

DHCP Relay Configuration

- 3636-1

```
3636-1# sh run dhcp

!Command: show running-config dhcp
!Time: Mon Aug 6 08:26:00 2018

version 9.2(1)
feature dhcp

service dhcp
```

```

ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ipv6 dhcp relay

interface Vlan1001
 ip dhcp relay address 192.1.42.3 use-vrf vxlan-900002
 ip dhcp relay source-interface loopback2

```

- 3636-2

```

!Command: show running-config dhcp
!Time: Mon Aug 6 08:26:16 2018

version 9.2(1)
feature dhcp

service dhcp
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ipv6 dhcp relay

interface Vlan1001
 ip dhcp relay address 192.1.42.3 use-vrf vxlan-900002
 ip dhcp relay source-interface loopback2

```

Debug Output

- The following is a packet dump for DHCP interact sequences.

```

3636-1# ethanalyzer local interface inband display-filter "udp.srcport==67 or
udp.dstport==67" limit-captured-frames 0
Capturing on inband
20180806 08:59:35.758314 0.0.0.0 -> 255.255.255.255 DHCP DHCP Discover - Transaction
ID 0x3eebccae
20180806 08:59:35.758878 33.33.33.33 -> 192.1.42.3 DHCP DHCP Discover - Transaction ID
0x3eebccae
20180806 08:59:37.759560 192.1.42.3 -> 33.33.33.33 DHCP DHCP Offer - Transaction ID
0x3eebccae
20180806 08:59:37.759905 172.16.16.1 -> 172.16.16.11 DHCP DHCP Offer - Transaction ID
0x3eebccae
20180806 08:59:37.760313 0.0.0.0 -> 255.255.255.255 DHCP DHCP Request - Transaction ID
0x3eebccae
20180806 08:59:37.760733 33.33.33.33 -> 192.1.42.3 DHCP DHCP Request - Transaction ID
0x3eebccae
20180806 08:59:37.761297 192.1.42.3 -> 33.33.33.33 DHCP DHCP ACK - Transaction ID
0x3eebccae
20180806 08:59:37.761554 172.16.16.1 -> 172.16.16.11 DHCP DHCP ACK - Transaction ID
0x3eebccae

```

- DHCP Discover packet 3636-1 sent to DHCP server.

giaddr is set to 33.33.33.33 (loopback2) and suboptions 5/11/151 are set accordingly.

Bootstrap Protocol

```

Message type: Boot Request (1)
Hardware type: Ethernet (0x01)
Hardware address length: 6
Hops: 1
Transaction ID: Ox3eebccae
Seconds elapsed: 0
Bootp flags: 0x0000 (unicast)
Client IP address: 0.0.0.0 (0.0.0.0)
Your (client) IP address: 0.0.0.0 (0.0.0.0)
Next server IP address: 0.0.0.0 (0.0.0.0)
Relay agent IP address: 33.33.33.33 (33.33.33.33)
Client MAC address: i-iughes_01:51:51 (00:00:10:01:51:51)
Client hardware address padding: 00000000000000000000
Server host name not given
Boot file name not given
Magic cookie: DHCP
Option: (53) DHCP Message Type
  Length: 1
  DHCP: Discover (1)
Option: (55) Parameter Request List
Option: (61) client identifier
Option: (82) Agent Information option
  Length: 47
Option 82 Suboption: (1) Agent circuit W
Option 82 suboption: (2) Agent Remote 10
Option 82 suboption: (151) VRF name/VPN ID
Option 82 Suboption: (11) Server ID Override
  Length: 4
  Server ID Override: 172.16.16.1 (172.16.16.1)
Option 82 Suboption: (5) Link selection
  Length: 4
  Link selection: 172.16.16.0 (172.16.16.0)

```

Client on Tenant VRF and Server on Non-Default Non-VXLAN VRF

The DHCP server is placed into the management VRF and is reachable through the M0 interface. The IP address changes to 10.122.164.147 accordingly.

```

3636-1# sh run int m0
!Command: show running-config interface mgmt0
!Time: Tue Aug 7 09:17:04 2018
version 9.2(1)
interface mgmt0
  vrf member management
  ip address 10.122.165.134/25

3636-1# ping 10.122.164.147 vrf management cou 1
PING 10.122.164.147 (10.122.164.147): 56 data bytes
64 bytes from 10.122.164.147: icmp_seq=0 ttl=251 time=1.024 ms
- 10.122.164.147 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 1.024/1.024/1.024 ms

3636-2# sh run int m0
!Command: show running-config interface mgmt0
!Time: Tue Aug 25 09:17:47 2015
version 7.0(3)I1(3)
interface mgmt0
  vrf member management
  ip address 10.122.165.148/25

```



```

3636-2# ping 10.122.164.147 vrf management cou 1
PING 10.122.164.147 (10.122.164.147): 56 data bytes
64 bytes from 10.122.164.147: icmp_seq=0 ttl=251 time=1.03 ms
- 10.122.164.147 ping statistics -
1 packets transmitted, 1 packets received, 0.00% packet loss
round-trip min/avg/max = 1.03/1.03/1.03 ms

```

DHCP Relay Configuration

- 3636-1

```

3636-1# sh run dhcp 3636-2# sh run dhcp

!Command: show running-config dhcp
!Time: Mon Aug 6 08:26:00 2018

version 9.2(1)
feature dhcp

service dhcp
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ipv6 dhcp relay

interface Vlan1001
 ip dhcp relay address 10.122.164.147 use-vrf management

```

- 3636-2

```

3636-2# sh run dhcp
!Command: show running-config dhcp
!Time: Tue Aug 7 09:17:47 2018

version 9.2(1)
feature dhcp

service dhcp
ip dhcp relay
ip dhcp relay information option
ip dhcp relay information option vpn
ipv6 dhcp relay

interface Vlan1001
ip dhcp relay address 10.122.164.147 use-vrf management

```

Debug Output

- The following is a packet dump for DHCP interact sequences.

```

3636-1# ethanalyzer local interface inband display-filter "udp.srcport==67 or
udp.dstport==67" limit-captured-frames 0
Capturing on inband
20180806 09:30:54.214998 0.0.0.0 -> 255.255.255.255 DHCP DHCP Discover - Transaction
ID 0x28a8606d
20180806 09:30:56.216491 172.16.16.1 -> 172.16.16.11 DHCP DHCP Offer - Transaction ID

```

```

0x28a8606d
20180806 09:30:56.216931 0.0.0.0 -> 255.255.255.255 DHCP DHCP Request - Transaction ID
  0x28a8606d
20180806 09:30:56.218426 172.16.16.1 -> 172.16.16.11 DHCP DHCP ACK - Transaction ID
  0x28a8606d

```

```

3636-1# ethanalyzer local interface mgmt display-filter "ip.src==10.122.164.147 or
ip.dst==10.122.164.147" limit-captured-frames 0
Capturing on mgmt0
20180806 09:30:54.215499 10.122.165.134 -> 10.122.164.147 DHCP DHCP Discover - Transaction
  ID 0x28a8606d
20180806 09:30:56.216137 10.122.164.147 -> 10.122.165.134 DHCP DHCP Offer - Transaction
  ID 0x28a8606d
20180806 09:30:56.217444 10.122.165.134 -> 10.122.164.147 DHCP DHCP Request - Transaction
  ID 0x28a8606d
20180806 09:30:56.218207 10.122.164.147 -> 10.122.165.134 DHCP DHCP ACK - Transaction
  ID 0x28a8606d

```

- DHCP Discover packet 3636-1 sent to DHCP server.

giaddr is set to 10.122.165.134 (mgmt0) and suboptions 5/11/151 are set accordingly.

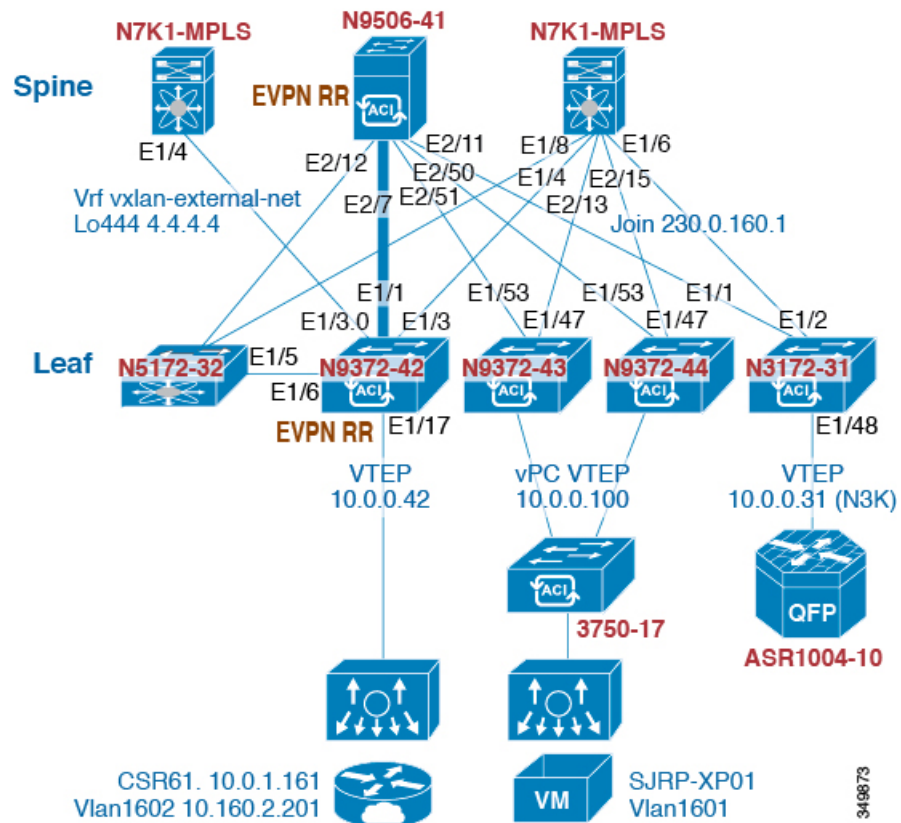
```

Bootstrap Protocol
  Message type: Boot Request (1)
  Hardware type: Ethernet (0x01)
  Hardware address length: 6
  Hops: 1
  Transaction ID: 0x28a8606d
  Seconds elapsed: 0
  Bootp flags: 0x0000 (Unicast)
  Client IP address: 0.0.0.0 (0.0.0.0)
  Your (client) IP address: 0.0.0.0 (0.0.0.0)
  Next server IP address: 0.0.0.0 (0.0.0.0)
  Relay agent IP address: 10.122.165.134 (10.122.165.134)
  Client MAC address: Hughes_01:51:51 (00:00:10:01:51:51)
  Client hardware address padding: 00000000000000000000
  Server host name not given
  Boot file name not given
  Magic cookie: DHCP
  Option: (53) DHCP Message Type
    Length: 1
    DHCP: Discover (1)
  Option: (55) Parameter Request List
  Option: (61) Client identifier
  Option: (82) Agent Information Option
    Length: 47
    Option 82 Suboption: (1) Agent Circuit ID
    Option 82 Suboption: (2) Agent Remote ID
    Option 82 Suboption: (151) VRF name/VPN ID
    Option 82 Suboption: (11) Server ID Override
      Length: 4
      Server ID Override: 172.16.16.1 (172.16.16.1)
    Option 82 Suboption: (5) Link selection
      Length: 4
      Link selection: 172.16.16.0 (172.16.16.0)

```

Configuring VPC Peers Example

The following is an example of how to configure routing between VPC peers in the overlay VLAN for a DHCP relay configuration.



- Enable DHCP service.

```
service dhcp
```

- Configure DHCP relay.

```
ip dhcp relay
ip dhcp relay information option
ip dhcp relay sub-option type cisco
ip dhcp relay information option vpn
```

- Create loopback under VRF where you need DHCP relay service.

```
interface loopback601
 vrf member evpn-tenant-kk1
 ip address 160.1.0.43/32
 ip router ospf 1 area 0 /* Only required for VPC VTEP. */
```

- Advertise LoX into the Layer 3 VRF BGP.

```
Router bgp 2
vrf X
  network 10.1.1.42/32
```

- Configure DHCP relay on the SVI under the VRF.

```
interface Vlan1601
vrf member evpn-tenant-kk1
ip address 10.160.1.254/24
fabric forwarding mode anycast-gateway
ip dhcp relay address 10.160.2.201
ip dhcp relay source-interface loopback601
```

- Configure Layer 3 VNI SVI with **ip forward**.

```
interface Vlan1600
vrf member evpn-tenant-kk1
ip forward
```

- Create the routing VLAN/SVI for the VPC VRF.



Note Only required for VPC VTEP.

```
Vlan 1605
interface Vlan1605
vrf member evpn-tenant-kk1
ip address 10.160.5.43/24
ip router ospf 1 area 0.0.0.41
```

- Create the VRF routing.



Note Only required for VPC VTEP.

```
router ospf 1
vrf evpn-tenant-kk1
  router-id 10.160.5.43
```

vPC VTEP DHCP Relay Configuration Example

To address a need to configure a VLAN that is allowed across the MCT/peer-link, such as a vPC VLAN, an SVI can be associated to the VLAN and is created within the tenant VRF. This becomes an underlay peering, with the underlay protocol, such as OSPF, that needs the tenant VRF instantiated under the routing process.

Alternatively, instead of placing the SVI within the routing protocol and instantiate the Tenant-VRF under the routing process, you can use the static routes between the vPC peers across the MCT. This approach ensures that the reply from the server returns to the correct place and each VTEP uses a different loopback interface for the GiAddr.

The following are examples of these configurations:

- Configuration of SVI within underlay routing:

```
/* vPC Peer-1 */

router ospf UNDERLAY
vrf tenant-vrf

interface Vlan2000
  no shutdown
  mtu 9216
  vrf member tenant-vrf
  ip address 192.168.1.1/30
  ip router ospf UNDERLAY area 0.0.0.0

/* vPC Peer-2 */

router ospf UNDERLAY
vrf tenant-vrf

interface Vlan2000
  no shutdown
  mtu 9216
  vrf member tenant-vrf
  ip address 192.168.1.2/30
  ip router ospf UNDERLAY area 0.0.0.0
```

- Configuration of SVI using static routes between vPC peers across the MCT:

```
/* vPC Peer-1 */

interface Vlan2000
  no shutdown
  mtu 9216
  vrf member tenant-vrf
  ip address 192.168.1.1/30

vrf context tenant-vrf
ip route 192.168.1.2/30 192.168.1.1

/* vPC Peer-2 */

interface Vlan2000
  no shutdown
  mtu 9216
  vrf member tenant-vrf
  ip address 192.168.1.2/30

vrf context tenant-vrf
ip route 192.168.1.1/30 192.168.1.2
```

