JUNIPER NETWORKS | Engineering Simplicity

Configure DHCP Relay in EVPN-VXLAN Fabric Architecture

Published
2023-09-13

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Network Configuration Example

This document describes the configuration required for working with DHCP Relay in an EVPN-VXLAN fabric. The document also highlights key considerations when configuring DHCP Relay with EVPN-VXLAN.

RELATED DOCUMENTATION

No Link Title

CHAPTER 1

Configure DHCP Relay in EVPN-VXLAN Fabric

About This Network Configuration Example

This network configuration example (NCE) describes how to deploy DHCP relay in an EVPN-VXLAN fabric architecture in an enterprise networking environment. The document highlights key considerations required for preparing your network to support DHCP relay in EVPN-VXLAN fabric.

Use Case Overview

IN THIS SECTION

Enterprise networks are adopting IP Fabric at the core and distribution layers as underlay and use an EVPN-VXLAN as overlay network to provide Layer 3 or Layer 2 connectivity.

Data center environments require Dynamic Host Configuration Protocol (DHCP) relay to pass DHCP messages between DHCP Clients and a DHCP Server. The DHCP Relay is typically configured at the subnet default gateway.

Benefits

DHCP is an essential component of any data center deployment. DHCP Relay requirements in the data center fabric may vary based on the unique requirements in a given data center deployment. Juniper's data center fabric architectures based on EVPN-VXLAN provide the necessary flexibility to customize DHCP Relay configuration to address most of these requirements.

Technical Overview

IN THIS SECTION

DHCP Relay

Dynamic Host Configuration Protocol (DHCP) relay allows the network to forward DHCP messages between DHCP Clients and DHCP Servers. The configuration used in this topic should work whether the fabric is deployed in CRB or ERB model.

The DHCP Server and Clients communicate with each other over the existing network without further configuration when the DHCP Client and Server are in the same VLAN. When a DHCP Client and Server are in different VLANs, DHCP traffic between the client and server is forwarded between the VLANs through the IRB interfaces on spine/leaf devices. You must configure the IRB interfaces on the spine-and-leaf devices to support DHCP Relay.

When DHCP Relay is used with EVPN-VXLAN, you cannot configure any binding/snooping features that require the relay to maintain state for the DHCP packets. For latest updates on DHCP features, see DHCP User Guide.

Use the `forward-only` option for the DHCP Relay configuration. The forward–only option ensures that DHCP packets are forwarded on the switch without creating DHCP Server Client bindings.

Location of the DHCP Server

The DHCP Server is connected to the service leaf on the fabric or in a different DC and reachable over Layer 3 network. On the device which is performing the relay, the DHCP Server is reachable in a directly connected VLAN or over a Type 5 tunnel, see Figure 1 on page 4 and Figure 2 on page 5.

Figure 1: Centrally-routed bridging

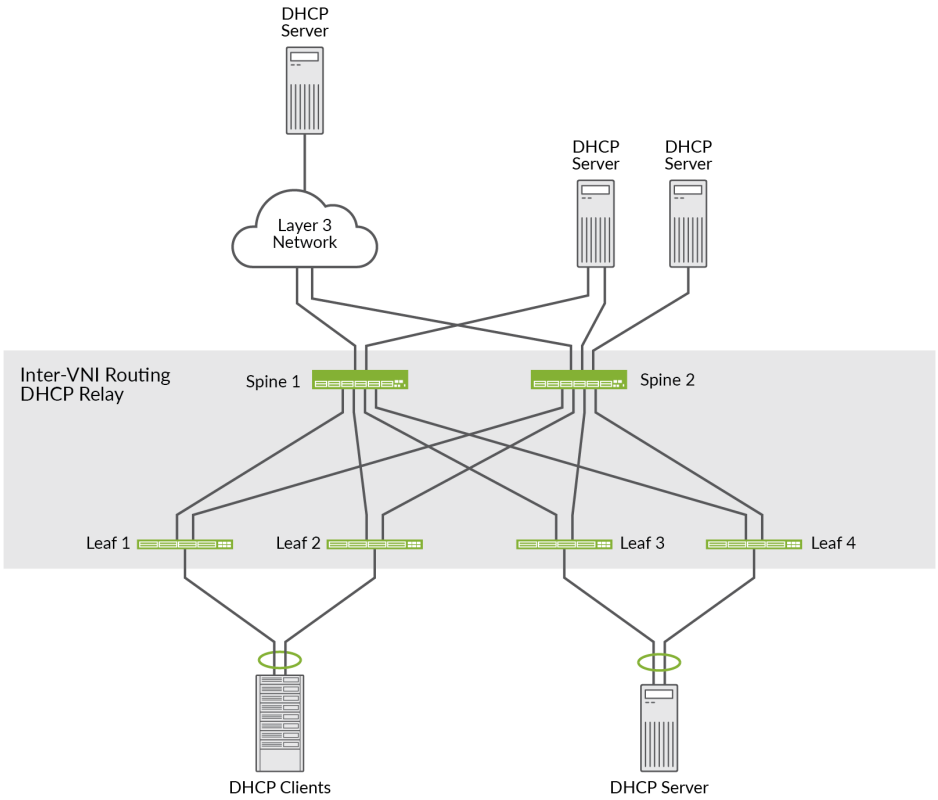

Centrally-routed Bridging

Figure 2: Edge-routed bridging

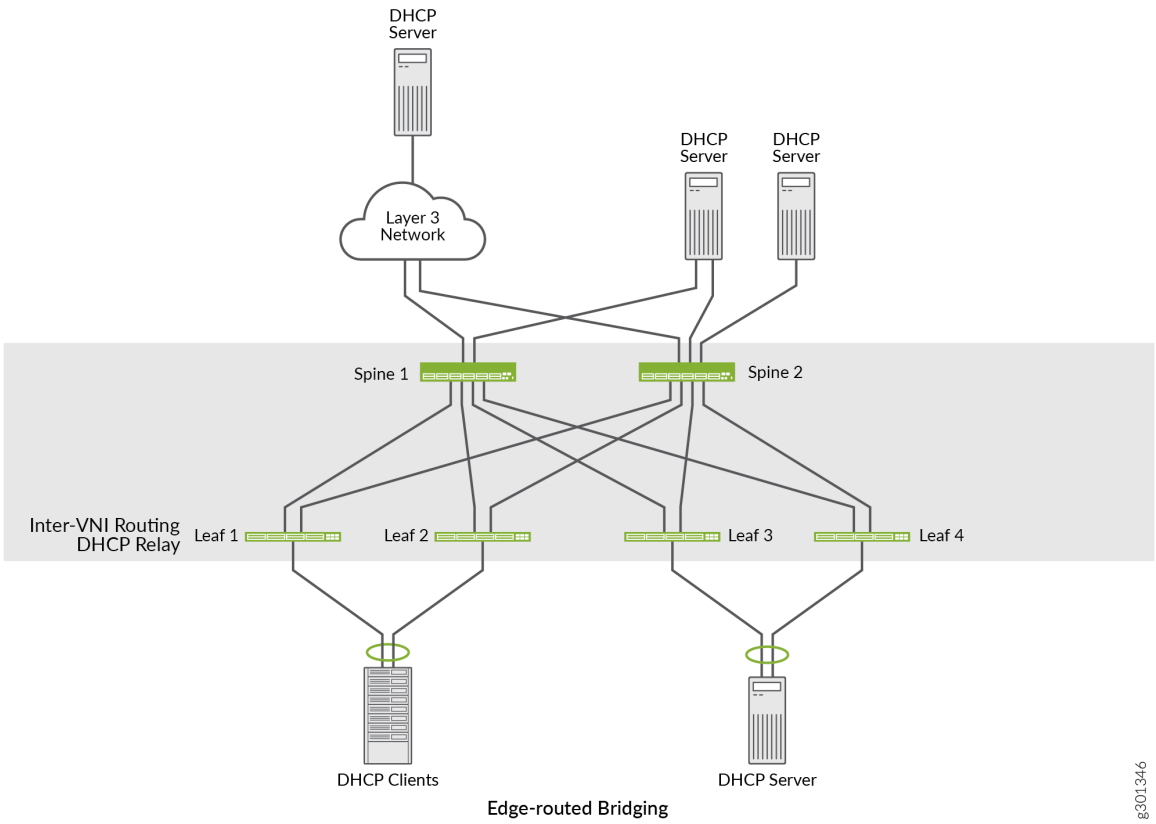

Default Gateway Function

You can use one of the following two models to enable the default gateway function:

IRB Interface with a VGA with Anycast IP Address

In this model, for a given VXLAN network identifier(VNI), configure an IRB interface on each provideredge(PE) device with a unique IP address and an anycast IP address. The Layer 3 VXLAN gateway automatically generates a MAC address.

Example:

```
set interfaces irb unit 202 virtual-gateway-accept-data
set interfaces irb unit 202 family inet address 192.168.202.2/24 primary
set interfaces irb unit 202 family inet address 192.168.202.2/24 preferred
```

```
set interfaces irb unit 202 family inet address 192.168.202.2/24 virtual-gateway-address
192.168.202.1
```

When you configure the DHCP Relay in this model, the source address of the DHCP Relay packets is the unique IP address of the IRB interface. The Relay agent IP address (giaddr) in the DHCP request message is same as the unique IP address of the IRB interface.

The DHCP Server uses the giaddr field to:

- Identify the pool for the DHCP request. DHCP Server looks for a pool that matches the subnet for the IP address available in the giaddr field.
- Identify the destination IP address of the DHCP reply message. DHCP Server sends the DHCP reply message to the IP address available in the giaddr field.

IRB Interface with Anycast IP Address

Many data center deployments use IRB Anycast model to enable default gateway function. In this model, for a given VNI, configure an IRB interface on each PE device with the same Anycast IP address.

```
set interfaces irb unit 202 family inet address 192.168.202.1/24
set interfaces irb unit 202 mac 00:0:02:02:00:01
```

When you enable DHCP Relay in the Anycast IRB model, the source address of the DHCP Relay packets is the Anycast IP address of the IRB interface. The Relay agent IP address (giaddr) in the DHCP request message is the Anycast IP address of the IRB interface. The DHCP Server will send the DHCP reply message with the destination IP address using the address from giaddr field in the DHCP request message. DHCP reply message might go back to a leaf that did not relay the DHCP Client request. This is because all PE devices with an IRB in that VNI have the same IP address.

To address this issue, consider the following when you use IRB interface with Anycast IP address:

- Relay agent IP address (giaddr) must be unique for each leaf or unique for each VRF inside a leaf in the case of multi-tenancy.
- When you use the loopback address as the Relay agent IP address (giaddr), the DHCP reply messages returns to the same leaf that initiated the DHCP relay.
- DHCP Server uses the giaddr field for pool selection. If the giaddr field is used as the loopback address of the leaf, it does not reflect actual subnet from which the IP address should be allocated. In this case, the DHCP Server is unable to select a pool to assign an IP address. To address this, include the `option-82` attribute link-selection (suboption 5) in the DHCP Relay request. The option 82 attribute includes the IP address of the IRB interface on which the client request was received. The DHCP Server can reuse the information to identify the pool for the DHCP request.

For more details on default gateway configuration, see Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network

Read VXLAN Constraints on QFX Series and EX Series Switches be aware of the constraints when you configure Virtual Extensible LANs (VXLANs) on QFX Series and EX Series switches.

RELATED DOCUMENTATION

- Data Center EVPN-VXLAN Fabric Architecture Guide
- Junos DHCP Relay Agent Overview

Configure a DHCP Relay in EVPN-VXLAN Fabric Architecture

IN THIS SECTION

DHCP Relay is an essential feature in most data center deployments. This example shows how to configure DHCP Relay in an EVPN-VXLAN-based data center fabric. The document also covers other common deployment models for DHCP Relay depending on how the DHCP Server is connected to the network. See "Technical Overview" on page 3 for details.

Requirements

This example uses the following hardware and software components:

- QFX5120 switches or QFX10002 switches

- Junos OS Release 18.4R2-S5

We've tested the configuration example using Junos OS Release 18.4R2-S5.

For details on supported platforms and Junos or Junos Evolved release support for DHCP relay, see Feature Explorer.

Overview

The DHCP Relay agent operates as the interface between DHCP Clients and the Server. DHCP Relay agent forwards incoming requests from DHCP Clients to a specified DHCP Server. In this example, we are using the Edge-routed bridging (ERB) topology as shown in Figure 3 on page 8.

Figure 3: DHCP Relay Overview

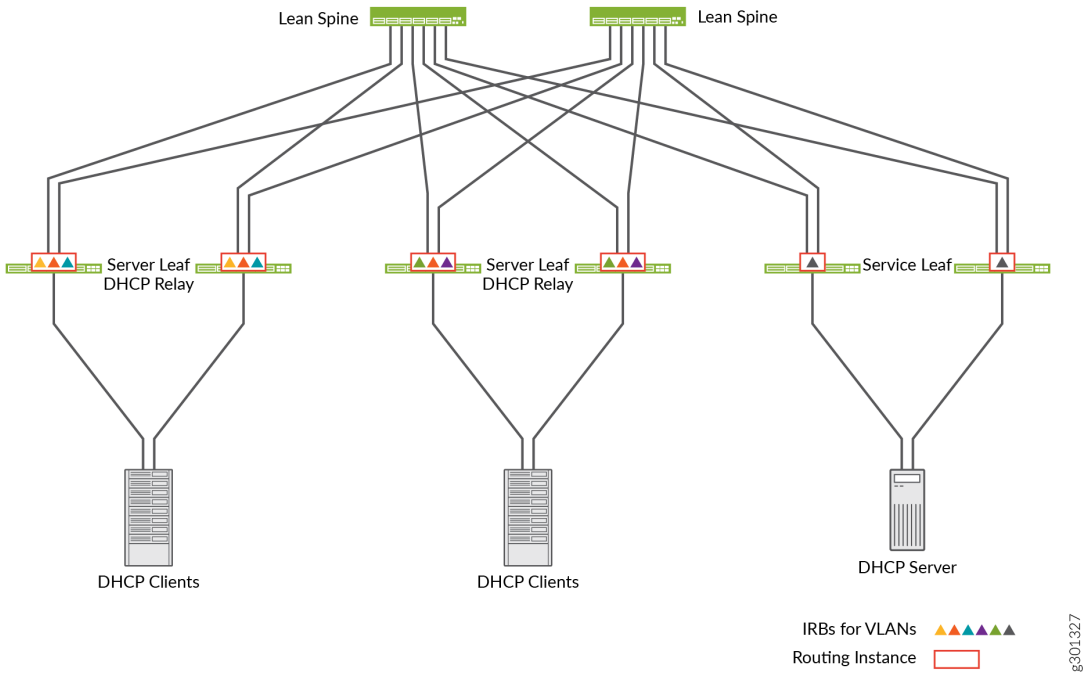

In the case of ERB, the inter-VLAN routing happens at the server leaf layer. The server leaf switches are configured to perform DHCP Relay function for the VLANs that have IRB interfaces configured on those switches.

You can apply the similar DHCP Relay configuration to a centrally-routed bridging (CRB) topology as well. In the case of CRB, Inter-VLAN routing happens at the spine switches level. So, the DHCP Relay must be configured on the spine switches.

Configuration

IN THIS SECTION

DHCP Relay Configuration with Virtual Gateway Address (VGA)

Step-by-Step Procedure

Do the steps that follow to configure DHCP relay for the IRB interface configured with VGA.

1. Enable DHCP Relay with `forward-only` option. The forward–only option ensures that DHCP packets are forwarded on the switch and that no DHCP Server Client bindings are created.

   ```
   set routing-instances TENANT_1_VRF forwarding-options dhcp-relay forward-only
   ```

   Do not use any other DHCP Relay overrides.

2. Create and activate the DHCP Relay server group.

   ```
   set routing-instances TENANT_1_VRF forwarding-options dhcp-relay server-group Server_Group1
   10.101.10.31
   set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 active-
   server-group Server_Group1
   ```

   The DHCP Relay server group include one or more DHCP Servers—individually identified by IP address—and a user-defined name for the servers. In this example, one DHCP server—10.101.10.31—is assigned into a DHCP server group named Server_Group1.

3. Associate the server group with the IRB interfaces on the leaf devices.

```
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 interface
irb.110
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 interface
irb.120
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 interface
irb.130
```

DHCP Relay Configuration with Anycast IRB

Step-by-Step Procedure

Do the steps below to configure DHCP Relay for IRB interface configured with Anycast IP address.

1. Configure the DHCP Relay with the loopback address that will be used as relay source.

```
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 overrides
relay-source lo0.101
```

2. Enable DHCP Relay option 82 with **server-id-override** option.

```
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 relay-
option-82 server-id-override
```

3. Create the DHCP Relay server group and associate the server group with the IRB interfaces on the leaf devices. Same as in procedure "DHCP Relay Configuration with Virtual Gateway Address (VGA)" on page 9.

```
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay forward-only
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay server-group Server_Group1
10.101.10.31
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 active-
server-group Server_Group1
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 interface
irb.110
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 interface
irb.120
```

```
set routing-instances TENANT_1_VRF forwarding-options dhcp-relay group Relay_Group1 interface
irb.130
```

4. Check packet capture on the DHCP Server to verify the contents of the relayed DHCP packet. Figure 4 on page 11 shows the sample packet capture file.

Figure 4: Packet Capture on the DHCP Server

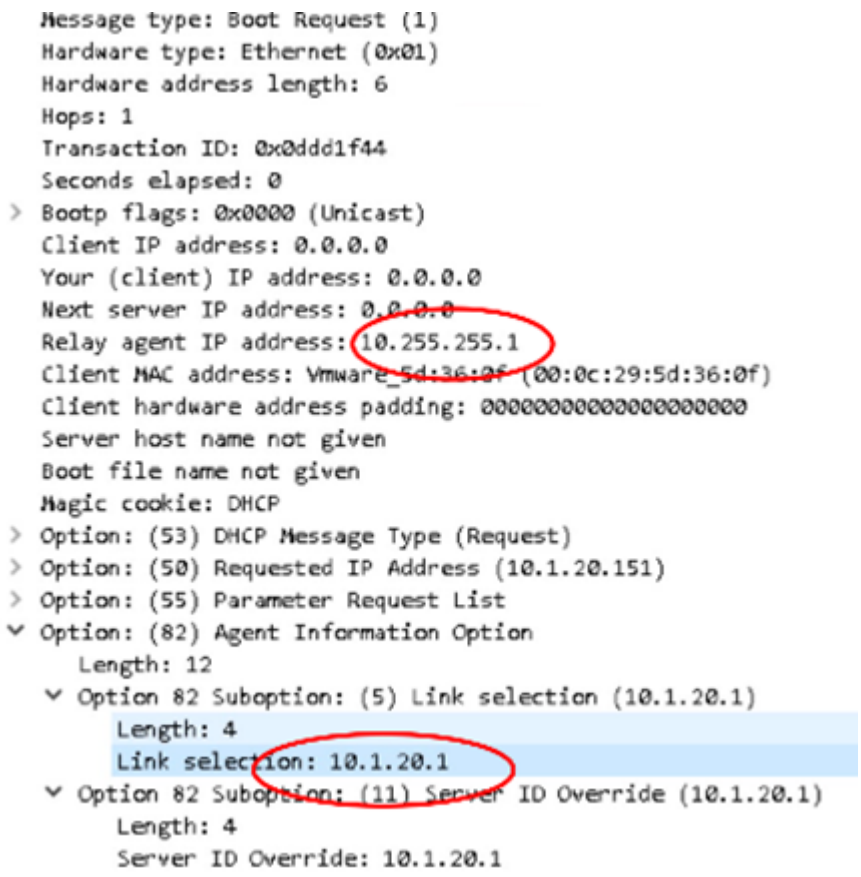

In the sample, you can notice that the relay agent IP address is the loopback IP address and the link selection attribute shows the IP address of the IRB interface.

DHCP Server Reachable only in a Service VRF

Step-by-Step Procedure

Do the steps below to configure the DHCP Server in a Service VRF.

1. Configure the loopback interfaces.

```
set interfaces lo0 unit 99 family inet address 1.1.6.2/32
set interfaces lo0 unit 100 family inet address 1.1.7.2/32
```

   Apart from the lo0.0 interface used for VTEP, you must use a separate loopback interfaces for every routing instance. In this case, the loopback interface lo0.110 is associated with the DHCP Server VRF. The loopback interface lo0.120 is associated with the DHCP Client VRF.

2. Configure the routing instance where the DHCP Server is located. The DHCP Server is located in VLAN 99 with IRB.99. The IRB.99 is placed in TENANT_SERVICE_VRF.

   Complete the following configurations in the Service VRF:

   - Configure the `dhcp-relay forward-only-replies` option to enable DHCP response packets forwarded to the DHCP Clients in the other VRF.
   - Configure the `auto-export` command along with vrf-target export and import policies that also import routes from the DHCP Client VRFs.

```
set routing-instances TENANT_SERVICE_VRF description VRF for DHCP server
set routing-instances TENANT_SERVICE_VRF instance-type vrf
set routing-instances TENANT_SERVICE_VRF interface irb.99
set routing-instances TENANT_SERVICE_VRF interface lo0.99
set routing-instances TENANT_SERVICE_VRF route-distinguisher 1.1.6.2:1099
set routing-instances TENANT_SERVICE_VRF vrf-import TENANT_SRV-IMPORT
set routing-instances TENANT_SERVICE_VRF vrf-export TENANT_SRV-EXPORT
set routing-instances TENANT_SERVICE_VRF vrf-target target:99:65001
set routing-instances TENANT_SERVICE_VRF vrf-table-label
set routing-instances TENANT_SERVICE_VRF routing-options auto-export
set routing-instances TENANT_SERVICE_VRF forwarding-options dhcp-relay forward-only-replies
set policy-options policy-statement TENANT_SRV-EXPORT term Direct-Routes from protocol direct
set policy-options policy-statement TENANT_SRV-EXPORT term Direct-Routes then community add
com-vrf-Tenant_SRV
set policy-options policy-statement TENANT_SRV-EXPORT term Direct-Routes then accept
set policy-options policy-statement TENANT_SRV-IMPORT term vs-Tenant_SRV from community com-
vrf-Tenant_SRV
set policy-options policy-statement TENANT_SRV-IMPORT term vs-Tenant_SRV then accept
set policy-options policy-statement TENANT_100-IMPORT term vs-Tenant_100 from community com-
vrf-Tenant_100
set policy-options policy-statement TENANT_100-IMPORT term vs-Tenant_100 then accept
```

```
set community com-vrf-Tenant_SRV member target:99:65001
set community com-vrf-Tenant_100 member target:100:65001
```

3. Configure the routing instances where the DHCP Clients are located.

   In this case, the DHCP Clients are located in VLAN 10 and VLAN 20 with corresponding IRB interfaces—IRB.10 and IRB.20. The IRB.10 and IRB.20 are part of the routing instance TENANT_CLIENT_VRF1.

   Configure the following in the DHCP Client VRF:

   - Configure the `dhcp-relay forward-only routing-instance <name>` option. This configuration specifies the routing instance where the DHCP Server is located. In this case, it is the "TENANT_SERVICE_VRF".
   - Configure the `auto-export` command to enable the routes from the DHCP Client VRF exported into the DHCP Server VRF.

```
set routing-instances TENANT_CLIENT_VRF1 description "VRF for DHCP Clients in VRF1”
set routing-instances TENANT_CLIENT_VRF1 instance-type vrf
set routing-instances TENANT_CLIENT_VRF1 interface irb.10
set routing-instances TENANT_CLIENT_VRF1 interface irb.20
set routing-instances TENANT_CLIENT_VRF1 interface lo0.100
set routing-instances TENANT_CLIENT_VRF1 route-distinguisher 1.1.7.2:1100
set routing-instances TENANT_CLIENT_VRF1 vrf-target target:100:65001
set routing-instances TENANT_CLIENT_VRF1 vrf-table-label
set routing-instances TENANT_CLIENT_VRF1 routing-options auto-export
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay forward-only routing-
instance TENANT_SERVICE_VRF
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay forward-only-replies
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay server-group
DHCP_SERVER_GROUP_1
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay group Relay_Group1
active-server-group DHCP_SERVER_GROUP_1
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay group Relay_Group1
overrides relay-source lo0.100
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay group Relay_Group1
relay-option-82 server-id-override
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay group Relay_Group1
interface irb.10
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay  group  Relay_Group1
interface irb.20
```

DHCP Relay with a Single Loopback IP Address for the Entire Chassis

Step-by-Step Procedure

Do the steps that follow to configure DHCP Relay with a single loopback IP address for the entire chassis.

1. Configure the loopback interfaces.

```
set interfaces lo0 unit 0 family inet address 1.1.5.2/32 preferred
set interfaces lo0 unit 99 family inet address 1.1.5.2/32
set interfaces lo0 unit 100 family inet address 1.1.5.2/32
set interfaces lo0 unit 101 family inet address 1.1.5.2/32
```

2. Configure the routing instance where the DHCP Server is located.

```
set routing-instances TENANT_SERVICE_VRF description "VRF for DHCP server“
set routing-instances TENANT_SERVICE_VRF instance-type vrf
set routing-instances TENANT_SERVICE_VRF interface irb.99
set routing-instances TENANT_SERVICE_VRF interface lo0.99
set routing-instances TENANT_SERVICE_VRF route-distinguisher 1.1.5.2:1099
set routing-instances TENANT_SERVICE_VRF vrf-import TENANT_SRV-IMPORT
set routing-instances TENANT_SERVICE_VRF vrf-export TENANT_SRV-EXPORT
set routing-instances TENANT_SERVICE_VRF vrf-target target:99:65001
set routing-instances TENANT_SERVICE_VRF vrf-table-label
set routing-instances TENANT_SERVICE_VRF routing-options auto-export
set routing-instances TENANT_SERVICE_VRF forwarding-options dhcp-relay forward-only-replies
set policy-options policy-statement TENANT_SRV-EXPORT term Direct-Routes from protocol direct
set policy-options policy-statement TENANT_SRV-EXPORT term Direct-Routes then community add
com-vrf-Tenant_SRV
set policy-options policy-statement TENANT_SRV-EXPORT term Direct-Routes then accept
set policy-options policy-statement TENANT_SRV-IMPORT term vs-Tenant_SRV from community com-
vrf-Tenant_SRV
set policy-options policy-statement TENANT_SRV-IMPORT term vs-Tenant_SRV then accept
set policy-options policy-statement TENANT_100-IMPORT term vs-Tenant_100 from community com-
vrf-Tenant_100
set policy-options policy-statement TENANT_100-IMPORT term vs-Tenant_100 then accept
set policy-options policy-statement TENANT_100-IMPORT term vs-Tenant_101 from community com-
vrf-Tenant_101
set policy-options policy-statement TENANT_100-IMPORT term vs-Tenant_101 then accept
set community com-vrf-Tenant_SRV member target:99:65001
```

```
set community com-vrf-Tenant_100 member target:100:65001
set community com-vrf-Tenant_100 member target:101:65001
```

3. Configure the routing instances where the DHCP Clients are located.

```
set routing-instances TENANT_CLIENT_VRF1 description "VRF for DHCP Clients in VRF1”
set routing-instances TENANT_CLIENT_VRF1 instance-type vrf
set routing-instances TENANT_CLIENT_VRF1 interface irb.10
set routing-instances TENANT_CLIENT_VRF1 interface irb.20
set routing-instances TENANT_CLIENT_VRF1 interface lo0.100
set routing-instances TENANT_CLIENT_VRF1 route-distinguisher 1.1.5.2:1100
set routing-instances TENANT_CLIENT_VRF1 vrf-target target:100:65001
set routing-instances TENANT_CLIENT_VRF1 vrf-table-label
set routing-instances TENANT_CLIENT_VRF1 routing-options auto-export
set routing-instances TENANT_CLIENT_VRF1 forwarding-options  dhcp-relay  forward-only
routing-instance TENANT_SERVICE_VRF
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay forward-only-replies
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay server-group
DHCP_SERVER_GROUP_1
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay group Relay_Group1
active-server-group DHCP_SERVER_GROUP_1
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay  group  Relay_Group1
overrides relay-source lo0.100
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay group Relay_Group1
relay-option-82 server-id-override
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay  group  Relay_Group1
interface irb.10
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay  group  Relay_Group1
interface irb.20
set routing-instances TENANT_CLIENT_VRF2 description "VRF for DHCP Clients in VRF2”
set routing-instances TENANT_CLIENT_VRF2 instance-type vrf
set routing-instances TENANT_CLIENT_VRF2 interface irb.20
set routing-instances TENANT_CLIENT_VRF2 interface lo0.101
set routing-instances TENANT_CLIENT_VRF2 route-distinguisher 1.1.5.2:1101
set routing-instances TENANT_CLIENT_VRF2 vrf-target target:101:65001
set routing-instances TENANT_CLIENT_VRF2 vrf-table-label
set routing-instances TENANT_CLIENT_VRF2 routing-options auto-export
set routing-instances TENANT_CLIENT_VRF2 f orwarding-options  dhcp-relay  forward-only
routing-instance TENANT_SERVICE_VRF
set routing-instances TENANT_CLIENT_VRF2 forwarding-options dhcp-relay forward-only-replies
set routing-instances TENANT_CLIENT_VRF2 forwarding-options dhcp-relay server-group
DHCP_SERVER_GROUP_1
```

```
set routing-instances TENANT_CLIENT_VRF2 forwarding-options dhcp-relay group Relay_Group1
active-server-group DHCP_SERVER_GROUP_1
set routing-instances TENANT_CLIENT_VRF2 forwarding-options dhcp-relay  group  Relay_Group1
overrides relay-source lo0.101
set routing-instances TENANT_CLIENT_VRF2 forwarding-options dhcp-relay group Relay_Group1
relay-option-82 server-id-override
set routing-instances TENANT_CLIENT_VRF2  forwarding-options dhcp-relay  group  Relay_Group1
interface irb.20
```

DHCPv6 Relay

Step-by-Step Procedure

Do the steps below to configure the DHCP6 Relay.

1. Configure the DHCPv6 Relay in the routing instance.

```
set interfaces irb unit 110 virtual-gateway-accept-data
set interfaces irb unit 110 family inet6 address 2001:db8::10:1:110:1/112
set interfaces lo0 unit 110 family inet address 192.168.110.1/32
set interfaces lo0 unit 110 family inet6 address 2001:db8::192:168:110:1/128
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay dhcpv6 overrides relay-
source lo0.110
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay dhcpv6 forward-only
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay dhcpv6 forward-only-
replies
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay dhcpv6 group all
interface irb.110
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay dhcpv6 group all
interface irb.120
set routing-instances TENANT_CLIENT_VRF1 forwarding-options dhcp-relay dhcpv6 group all
interface irb.130
```

2. Configure the IRB interfaces to send router advertisement messages with a default gateway address and the IPv6 prefix length information.

   The configuration enables the DHCPv6 Clients to identify the prefix length and the default gateway because the DHCPv6 Server does not provide the information.

```
set protocols router-advertisement interface irb.110 preference high
set protocols router-advertisement interface irb.110 max-advertisement-interval 60
```

```
set protocols router-advertisement interface irb.110 managed-configuration
set protocols router-advertisement interface irb.110 other-stateful-configuration
set protocols router-advertisement interface irb.110 solicit-router-advertisement-unicast
set protocols router-advertisement interface irb.110 prefix 2001:db8:0:0:10:1:110::/112 no-
autonomous
set protocols router-advertisement interface irb.110 prefix ::/0 no-autonomous
set protocols router-advertisement interface irb.120 preference high
set protocols router-advertisement interface irb.120 max-advertisement-interval 60
set protocols router-advertisement interface irb.120 managed-configuration
set protocols router-advertisement interface irb.120 other-stateful-configuration
set protocols router-advertisement interface irb.120 solicit-router-advertisement-unicast
set protocols router-advertisement interface irb.120 prefix 2001:db8:0:0:10:1:120::/112 no-
autonomous
set protocols router-advertisement interface irb.120 prefix ::/0 no-autonomous
set protocols router-advertisement interface irb.130 preference high
set protocols router-advertisement interface irb.130 max-advertisement-interval 60
set protocols router-advertisement interface irb.130 managed-configuration
set protocols router-advertisement interface irb.130 other-stateful-configuration
set protocols router-advertisement interface irb.130 solicit-router-advertisement-unicast
set protocols router-advertisement interface irb.130 prefix 2001:db8:0:0:10:1:130::/112 no-
autonomous
set protocols router-advertisement interface irb.130 prefix ::/0 no-autonomous
```

Step-by-Step Procedure

Verification for DHCPv6 Relay

1. Verify DHCPv6 Solicit Message

   Use the packet capture details of the DHCPv6 solicit message, see Figure 5 on page 18.

Figure 5: Packet Capture of the DHCPv6 Solicit Message

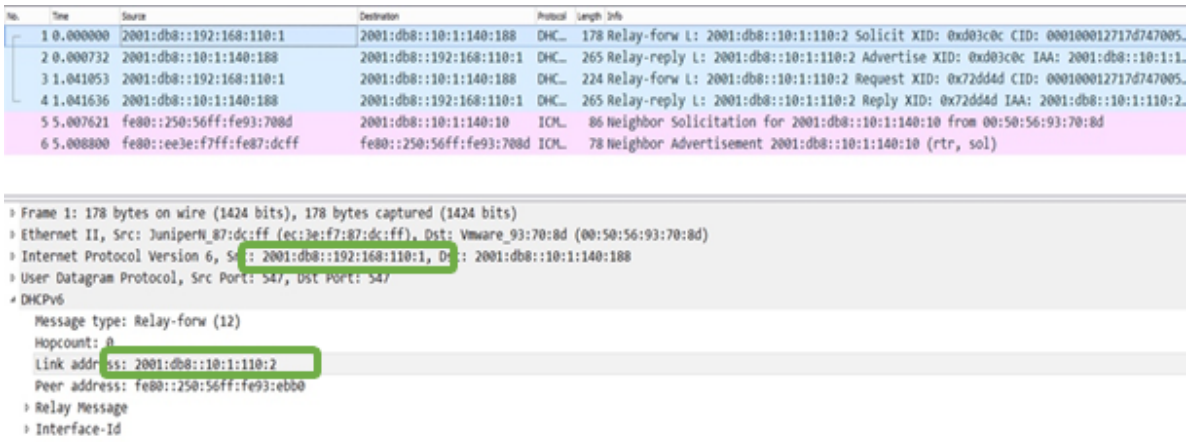

In the packet capture file, you can see information of the DHCPv6 solicit message on the DHCPv6 Server. The output indicates that the source address of the DHCPv6 relay packet is the loopback IPv6 address of the VRF on the leaf device. The link address field indicates the prefix pool that needs to be selected by the DHCPv6 Server for address assignment.

2. Verify DHCPv6 Reply Message Details

Use the packet capture details of the DHCPv6 reply message, see Figure 6 on page 18

Figure 6: Packet Capture of the DHCPv6 Reply Message

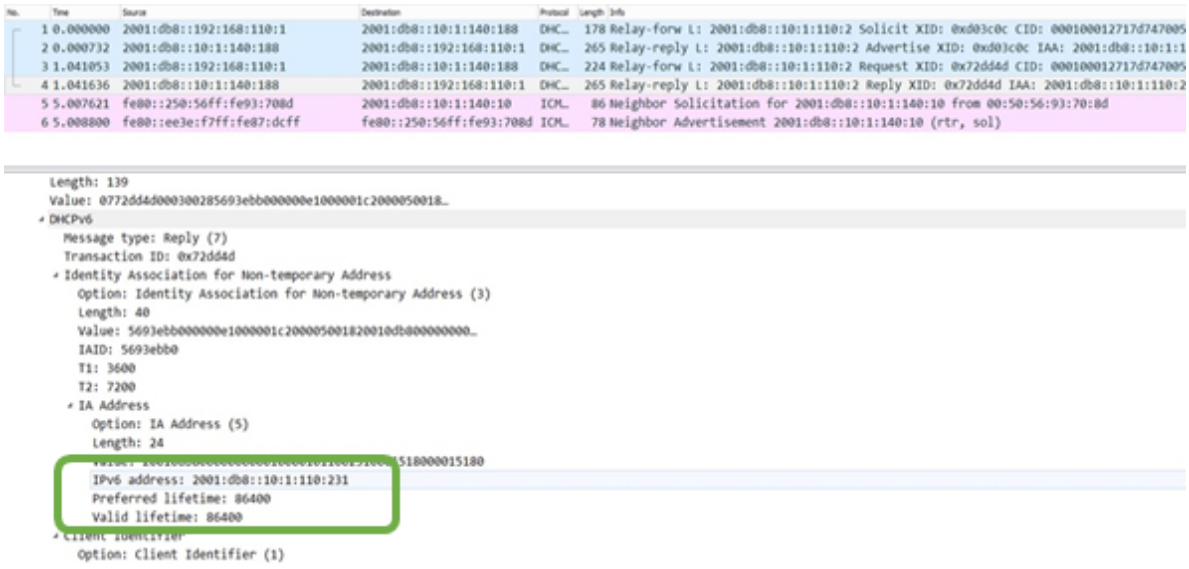

In the packet capture file, you can see information of the DHCPv6 reply message sent by the DHCPv6 Server. The DHCPv6 reply is sent to the loopback address in the VRF on the leaf device.

The fields indicate that the DHCPv6 Server is sending the assigned IP prefix for the endpoint and lifetime information. The reply message does not include prefix length and the default gateway address.

3. Verify DHCPv6 Client Details

   Verify router advertisement messages on the DHCPv6 Client.

   Use the packet capture details of the DHCPv6 solicit message, see Figure 7 on page 19

Figure 7: Packet Capture of the DHCPv6 Solicit Message

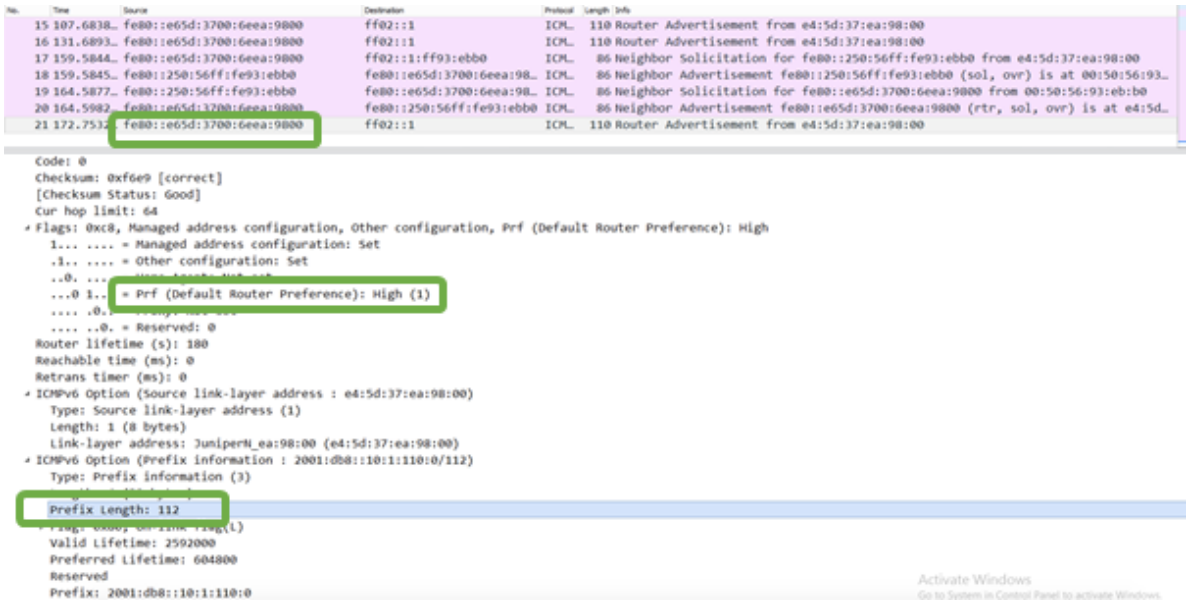

In the packet capture file, you can see information of the DHCPv6 Client. Here, you can see that the router advertisement message is providing the prefix length information along with the default gateway IP address as the link local IP address of the IRB interface.

Transit DHCP Relay

Validation

IN THIS SECTION

To confirm that the configuration is working properly, perform the following tasks:

Check DHCP Relay Statistics

Purpose

Verify that the DHCP relay is functioning.

Action

From operational mode:

```
user@host>set dhcp relay statistics
show dhcp relay statistics routing-instance Tenant1_VRF
Packets dropped:
    Total                     0

Messages received:
    BOOTREQUEST               1
    DHCPDECLINE               0
    DHCPDISCOVER              0
    DHCPINFORM                0
    DHCPRELEASE               0
    DHCPREQUEST               1
    DHCPLEASEACTIVE           0
    DHCPLEASEUNASSIGNED       0
    DHCPLEASEUNKNOWN          0
    DHCPLEASEQUERYDONE        0
```

```
Messages sent:
    BOOTREPLY                   1
    DHCPOFFER                   0
    DHCPACK                     1
    DHCPNAK                     0
    DHCPFORCERENEW              0
    DHCPLEASEQUERY              0
    DHCPBULKLEASEQUERY          0

Packets forwarded:
    Total                       2
    BOOTREQUEST                 1
    BOOTREPLY                   1
```

Enable Tracing Options for DHCP

Purpose

Enable tracing options for the DHCP Relay agent.

Action

Use the following commands from edit mode:

```
user@host> set system processes dhcp-service traceoptions file jdhcpd
```

```
user@host> set system processes dhcp-service traceoptions file size 1g
```

```
user@host> set system processes dhcp-service traceoptions level all
```

DHCP Log Files

Purpose

View DHCP log files to get DHCP services details.

Action

From operational mode:

```
user@host>show log jdhcpd

Nov 19 00:43:00.972038 [MSTR][DEBUG][default:Tenant1_VRF][RLY][INET][irb.110]
jdhcpd_io_process_ip_packet: LOCAL: recv pkt; sa 0.0.0.0; da 255.255.255.255; src_port 68;
dst_port 67; len 300
Nov 19 00:43:00.972074 [MSTR][DEBUG][default:Tenant1_VRF][RLY][INET][irb.110] --[ DHCP/BOOTP
from == 0.0.0.0, port == 68 ]--
Nnt1_VRF][RLY][INET][irb.110] --[ OPTION code  50, len   4, data 0a 01 6e c9 ]--



.........output truncated.............]




19 00:43:00.972320 [MSTR][DEBUG][default:Tenant1_VRF][RLY][INET][irb.110] --[ OPTION code  12,
len  12, data 6b 72 69 73 68 6e 61 6e 2d 76 6d 31 ]--
Nov 19 00:43:00.972345 [MSTR][DEBUG][default:Tenant1_VRF][RLY][INET][irb.110] --[ OPTION code
55, len  19, data 01 1c 02 79 0f 06 0c 28 29 2a 1a 77 03 79 f9 21 fc 2a 11 ]--
Nov 19 00:43:00.972361 [MSTR][INFO] [default:Tenant1_VRF][RLY][INET][irb.110] --[ OPTION code
255, len   0 ]--
Nov 19 00:43:00.972386 [MSTR][DEBUG] sus_name_get: Extracted ifd_name = lo0
Nov 19 00:43:00.972407 [MSTR][DEBUG] client_key_compose: Composing key (0xa74e2c0) for cid_l 0,
cid NULL, mac 00 50 56 93 eb b0, htype 1, subnet 192.168.110.1, ifindx 0, opt82_l 0, opt82 NULL
Nov 19 00:43:00.972424 [MSTR][DEBUG] client_key_compose: Successfully composed
CK_TYPE_HW_ADDR_ON_SUBNET (2) client key object.

Nov 19 00:43:00.991497 [MSTR][DEBUG] jdhcpd_hex_dump: 00 00 01 8b 00 00 02 2d 00 00 02 42 00 00
00 00
Nov 19 00:43:00.991509 [MSTR][DEBUG] jdhcpd_hex_dump: ab ad d0 0d 00 00 00 00 00 00 45 00 01 55
61 c6
Nov 19 00:43:00.991521 [MSTR][DEBUG] jdhcpd_hex_dump: 00 00 00 00 00 00 00 00 45 00 01 55 61 c6
40 00
Nov 19 00:43:00.991533 [MSTR][DEBUG] jdhcpd_hex_dump: 3f 11 2f 6b 0a 01 70 bc c0 a8 6e 01 00 43
00 43
Nov 19 00:43:00.991546 [MSTR][DEBUG] jdhcpd_hex_dump: 01 41 55 d4
```

```
Nov 19 00:43:00.991560 [MSTR][INFO] [irb.110] jdhcpd_io_get_ifs: The L3 interface is 557 and L2
interface is 578, using the L3 interface
Nov 19 00:43:00.992565 [MSTR][INFO] [default:Tenant1_VRF][RLY][INET][irb.110]
jdhcpd_io_send_packet_legacy: Set the outgoing if to 557
Nov 19 00:43:00.992640 [MSTR][INFO] [default:Tenant1_VRF][RLY][INET][irb.110]
jdhcpd_io_send_packet_legacy: DHCP PDU from 192.168.110.1 to 10.1.110.201 port 68 out interface
557 len 329



.........output truncated.............]
```

Meaning

The sample output shows the DHCP log messages in the messages file. The output command shown in the document is truncated for easy readability.

RELATED DOCUMENTATION

- Data Center EVPN-VXLAN Fabric Architecture Guide